Evolutionary Many-Objective Optimisation: An Exploratory Analysis

Robin C. Purshouse and Peter J. Fleming
Department of Automatic Control and Systems Engineering
University of Sheffield
Mappin Street, Sheffield, S1 3JD, UK.
{r.purshouse, p.fleming@sheffield.ac.uk}

Abstract – This inquiry explores the effectiveness of a class of modern evolutionary algorithms, represented by NSGA-II, for solving optimisation tasks with many conflicting objectives. Optimiser behaviour is assessed for a grid of recombination operator configurations. Performance maps are obtained for the dual aims of proximity to, and distribution across, the optimal trade-off surface. Classical settings for recombination are shown to be suitable for small numbers of objectives but correspond to very poor performance as the number of objectives is increased, even when large population sizes are used. Explanations for this behaviour are offered via the concepts of dominance resistance and active diversity promotion.

1 Introduction

Much of the research into multi-objective evolutionary algorithms (MOEAs) concentrates on optimisation tasks with two conflicting objectives. However, the real-world challenges to which these algorithms are applied often feature many more objectives (Coello *et al*, 2002). Hence, there is a clear need to extend evolutionary multi-objective optimisation (EMO) research into the realm of *many-objectives* (Farina and Amato, 2002).

Interactions often arise between objectives. These can be classified as *conflicting* or *harmonious*, and both interactions may co-exist between two objectives in the context of a single optimisation problem (Purshouse and Fleming, 2003a). In the case of conflict, a solution modification that will improve performance in one objective is seen to cause deterioration in a second objective. In the case of harmony, the modification causes simultaneous improvement to both objectives. The conflict that exists in a many-objective optimisation task has been identified as a serious challenge for contemporary EMO researchers, and it is this relationship that is explored in this paper.

If the only assumption concerning decision-maker (DM) preferences is that a unidirectional *line of preference* (Edgeworth, 1932) exists for each objective then performance comparisons between solutions can be based on the notion of *Pareto dominance*. For further details refer to Coello *et al* (2002). In these conditions, the optimal solution to an *M*-objective task, in which all objectives conflict, is an (*M*-1)-dimensional hypersurface. The number of samples required to represent the surface at a fixed resolution is exponential in *M*. Even if such an *approximation set* (Zitzler *et al*, 2003) could be achieved, the quantity of information contained within the set may overwhelm the DM, who must ultimately select a single solution.

The inherent difficulties in solving many-objective problems have lead EMO researchers to incorporate preference-based schemes into their algorithms, as comprehensively reviewed by Coello *et al* (2002). The fundamental aim of these methods is to limit the search requirements of the optimiser to a sub-region of overall objective-space. However, as argued by Knowles (2002), the potential for an exclusively Pareto-based solution to the many-objective optimisation problem remains a matter of some interest. Indeed, if the resolution of the obtained approximation set is regarded as a function of some maximum limit imposed on the size of the set (such as the population size of an MOEA), then there is no *a priori* restriction that prevents the achieved set from being globally non-dominated and optimally distributed across the trade-off surface. But is it possible to design an evolutionary algorithm that is capable of generating such results, given finite resources, when faced with many conflicting objectives?

A family of tractable, real-parameter optimisation tasks that are scalable to any number of conflicting objectives was proposed by Deb *et al* (2002a) to stimulate research into many-objective optimisation. In the first known study of its kind, this test suite was used by Khare *et al* (2003) to investigate the scalability of some contemporary MOEAs. An implementation of Corne *et al*'s (2000) *PESA* was found to generate approximation sets with good proximity and poor distribution as *M* increases, whilst both Deb *et al*'s (2002b) *NSGA-II* and Zitzler *et al*'s (2001) *SPEA2* were found to produce the opposite result. However, since a single design-space instance of each algorithm was used, and each algorithm is itself a complicated structure of basic EMO components, it is not immediately clear from the study which components and processes are critical from the many-objective optimisation perspective.

This paper focuses on the behaviour of a single MOEA, NSGA-II, which can be considered generally representative of a larger class of EMO optimisers. Behaviour can then be explained through decompositions of the algorithm in terms of fundamental search components and processes. In addition, results are generated for a *map* of variation operator configuration settings. This permits analysis to be made in terms of the *exploration-exploitation* (EE) trade-offs in EMO (Bosman and Thierens, 2003) and for performance *sweet-spots* to be identified (Goldberg, 1998).

The remainder of the paper is organised as follows. The fundamental processes that comprise the NSGA-II optimiser are described in Section 2, with reference to similar components of other MOEAs. In Section 3, the design of the empirical inquiry is introduced. The optimisation task and performance indicators considered are described and a new framework for studies into many-objective optimisation, based on the methodologies of Laumanns *et al* (2001) and Purshouse and Fleming (2002) is also proposed. The results of the inquiry are presented in Section 4, and these are subsequently analysed in Section 5. The paper concludes in Section 6 with the key issues raised by the study and suggested directions for future work. Note that an extended research report into evolutionary many-objective optimisation is also available (Purshouse and Fleming, 2003b).

2 NSGA-II Processes

The search mechanism, often described in terms of exploration and exploitation, employed by the class of evolutionary optimiser considered in this inquiry can be summarised by Equation 1.

$$P[t+1] = s_s\left(v\left[s_v\left(P[t]\right)\right], P[t]\right) \tag{1}$$

$P[t]$ is the population at iteration t, s_v is the *selection-for-variation* operator, v is the *variation* operator, and s_s is the *selection-for-survival* operator.

s_v, usually known simply as the *selection* operator, selects candidate solutions from the current population to form the *mating pool*. Variation operators are applied to the solutions in the mating pool to create a set of new candidate solutions. These new solutions then compete with the population of current solutions in the s_s stage to determine the composition of the subsequent population. The s_s operator is also known as the *population management* or *reinsertion* operator.

The different multi-objective evolutionary optimisers proposed in the literature are generally categorised by the manner in which selection is performed. One particular broad category of selection scheme can be represented by the NSGA-II. Whilst this algorithm has its own particularities, it can be considered broadly similar to other MOEAs that (i) use the concepts of dominance and density estimation as bases for fitness assignment, (ii) consistently assign equal selection probabilities to solutions with equal dominance and density measures, and (iii) respect a specific integer bound on population size.

2.1 Discriminators in Selection Processes

NSGA-II uses Pareto dominance comparisons between pairs of candidate solutions to form a partial ordering across a set of solutions, via Goldberg's (1989) *non-dominated sorting* method. This is a coarse-grained relative of *Pareto-based ranking* (Fonseca and Fleming, 1993) and the *strength* approach (Zitzler *et al*, 2001) used in other MOEAs.

Density estimation is also used as a basis for discrimination in NSGA-II. The density at a solution location is calculated using the first-nearest-neighbour (1NN) *crowding distance* estimator (Deb *et al*, 2002b), which is defined as the mean side length of the hypercube formed using 1NN values in each objective as vertices. In this inquiry, the boundary condition for an objective is set to the maximum non-boundary value calculated for that objective. This ensures that the estimator is unbiased for the equilibrium condition of a perfectly distributed approximation set. Crowding distance is a low complexity estimator with limited accuracy, and its effectiveness has been questioned (Laumanns *et al*, 2001; Deb *et al*, 2003). Nevertheless, the general effects of its inclusion are argued to be broadly representative of the effects of any estimator used in EMO.

2.2 Selection-for-Variation in NSGA-II

NSGA-II uses a type of binary tournament selection, defined by Deb *et al* (2002b) as the *crowded-comparison operator*, for s_v. From two solutions chosen at random (with replacement) from the population, selection is made via a primary comparison of non-dominated ranks. If the ranks are equal, a secondary comparison is made between density estimates. If the densities are equal, the tertiary stage selects one of the solutions at random for inclusion in the mating pool.

2.3 Selection-for-Survival in NSGA-II

In the s_s stage, both non-dominated sorting and density estimation are undertaken for the combined P[t] and post-variation solution sets. This allows a new partial ordering to be constructed, based primarily on dominance, and secondarily on density. This is a similar technique to both *SPEA2* fitness assignment (Zitzler *et al*, 2001) and *intra-ranking* (Purshouse and Fleming, 2002). A complete ordering can then be established by applying a random ordering within any shared equivalence class. P[t+1] is then deterministically designated as the best of this hierarchy. Note that Deb *et al* (2002b) describe an efficient implementation of the above *elitist* concept.

Motivated by the desire to control the EE trade-off via elitism, some modifications to the above strategy have been proposed (Deb and Goel, 2001; Laumanns *et al*, 2001). For simplicity, this inquiry assumes that the EE trade-off can be adequately controlled by the variation operator settings.

2.4 Variation

A two-parent recombination operator, *simulated binary crossover* (SBX), is used at the v stage (Deb and Agrawal, 1995). This is a popular operator that has been previously used in the context of NSGA-II for real-parameter function optimisation (Deb *et al*, 2002a; Khare *et al*, 2003). SBX generates child solutions according to a symmetric, self-adaptive, distribution about the parent values, with standard deviation based on (i) the distance between parents and (ii) a distribution parameter η_c (a larger value for η_c corresponds to a smaller expected relative closeness of children to their parents). Child values are then swapped with probability p_e. The operator is applied to a parent pair with probability p_c. In a uniform recombination scheme, given that SBX is to be applied to a decision vector, the probability of applying SBX to each element of the vector is p_{ic}. By setting p_c = 1, the probability of applying recombination to a vector is then equivalent to the probability of applying a standard mutation operator, thus simplifying comparisons with such schemes (Purshouse and Fleming, 2003b). This approach is adopted in the inquiry. Also, for simplicity, $p_e = 0$.

3 Experimental Design

3.1 Many-Objective Optimisation Task

This inquiry considers a real-parameter function optimisation task known as *DTLZ2*, which is defined in Equation 2. The task is taken from a highly tractable set of problems developed by Deb *et al* (2002a) specifically for studies into many-objective optimisation. The global Pareto front is continuous and non-convex. Distance from the front is determined by a single, unimodal cost function, *g*.

$$\begin{aligned} \min \quad & z_1 = (1+g)\cos(x_1\pi/2)\ldots\cos(x_{M-1}\pi/2), \\ \min \quad & z_2 = (1+g)\cos(x_1\pi/2)\ldots\sin(x_{M-1}\pi/2), \\ \vdots \quad & \vdots \\ \min \quad & z_M = (1+g)\sin(x_1\pi/2), \\ \text{where} \quad & g = \sum_{x_i \in \mathbf{x}_M} (x_i - 0.5)^2. \end{aligned} \tag{2}$$

M is the number of objectives, $n=M+k-1$ is the number of decision variables, *k* is a difficulty parameter set to 10 in this study, $\mathbf{x}_M = [x_M,\ldots,x_n]$, and $0 \le x_i \le 1$ for $i = [1,\ldots,n]$.

DTLZ2 is comprised of decision variables of two distinct functional types: those that control convergence towards the globally optimal surface ($x_1,\ldots,x_{M-1}$) and those that control distribution in objective-space ($x_M,\ldots,x_n$). The convergence-variables define the distance of the solution vector from the true front via a *k*-dimensional quadratic bowl, *g*, with global minimum $x_{M,\ldots,n}$ = 0.5. The distribution-variables describe position on the positive quadrant of the unit hypersphere. An *M*-objective instance of DTLZ2 is denoted as DTLZ2(*M*).

3.2 Performance Indicators

In the context of this inquiry, *performance* relates to the quality of the trade-off surface discovered by an optimiser, given a finite number of candidate solution evaluations. Quality is generally expressed in terms of (i) the proximity of the obtained locally non-dominated vectors to the true Pareto surface and (ii) the distribution of those vectors across the surface (Bosman and Thierens, 2003). Ideally, the optimiser should obtain solutions that are Pareto optimal (are of distance zero from the global front), that extend across the full range of optimal objective values, and that are as near uniformly distributed as the true surface permits.

Various performance indicators have been proposed to measure the different aspects of quality (Deb, 2001). Many indicators are unary (they describe the absolute performance of one approximation set), although a few are binary (they describe the relative performance of two sets). Zitzler *et al* (2003) have shown that no finite combination of unary measures can indicate whether one approximation set is superior to another (from the perspective of the dominance relation). Thus, care must be taken when making statements about global performance.

This study adopts the *functional* approach described by Deb and Jain (2002). Specific unary indicators are used to evaluate specific aspects of performance. There is no attempt to describe global performance using a unary indicator or indeed a combination of such indicators.

3.2.1 Proximity indicator

The proximity indicator measures a median level of distance of the approximation set, $\boldsymbol{Z_A}$, from the global trade-off surface. In terms of attainment across the objectives, an objective vector for DTLZ2 will respect Equation 3. The equality condition will only hold for a globally optimal vector. Thus, a specialised proximity indicator, I_P, for DTLZ2 can naturally be described by Equation 4. This is essentially the same as Veldhuizen's (1999) *generational distance* metric, for the case of a continuous globally optimal reference set, $\boldsymbol{Z_T}$.

$$\left[\sum_{m=1}^{M}(z_m)^2\right]^{1/2} \ge 1 \tag{3}$$

$$I_P = \underset{\boldsymbol{z_A}\in \boldsymbol{Z_A}}{\text{median}}\left\{\left[\sum_{m=1}^{M}\left(\boldsymbol{z}_{\mathbf{A}_m}\right)^2\right]^{1/2} - 1\right\} \tag{4}$$

In the inquiry, in order to clearly show the direction of optimiser evolution, proximity is calculated by subtracting the I_P calculated for the first generation from the median I_P obtained for the final 100 generations of the optimiser.

3.2.2 Spread indicator

To achieve high quantisation of the non-dominated set, it would be advantageous to express both the extent and uniformity aspects of distribution within a single indicator. This approach has been implemented in Deb *et al*'s (2002b) *Δ* metric. Unfortunately, it can become unclear which aspect of the distribution is responsible for the observed indicator value. For example, errors on the spread of the distribution can potentially mask difficulties with uniformity. To manage the complexity of the inquiry, only the spread of solutions is considered further.

The study uses a variant of Zitzler's (1999) *maximum spread* indicator. This metric measures the length of the diagonal of the hypercube with vertices set to the extreme objective values observed in the achieved approximation set, as defined in Equation 5.

$$D = \left[\sum_{m=1}^{M} \left(\max_{\mathbf{z}_A \in \mathbf{Z}_A} \mathbf{z}_{A_m} - \min_{\mathbf{z}_A \in \mathbf{Z}_A} \mathbf{z}_{A_m} \right)^2 \right]^{1/2} \qquad (5)$$

It is possible to achieve too much or too little spread. In the former case, the vectors span regions that are not part of the global trade-off surface, (highlighting a relationship between spread and proximity). In the latter case, the optimiser has converged to a sub-region (that *may* be globally optimal). To highlight the requirement for an intermediate spread value, the indicator, I_S, is formed by normalising D with respect to the optimal spread, as indicated in Equation 6. I_S values decreasing from unity to zero now represent increasing levels of convergence to a sub-region. Thus, globally optimal regions of the surface are certain to be missing. Indicator values increasing from unity demonstrate widespread dispersal of vectors throughout non-optimal objective-space.

$$I_s = D \Bigg/ \left[\sum_{m=1}^{M} \left(\max_{\mathbf{z}_T \in \mathbf{Z}_T} \mathbf{z}_{T_m} - \min_{\mathbf{z}_T \in \mathbf{Z}_T} \mathbf{z}_{T_m} \right)^2 \right]^{1/2} \qquad (6)$$

In the inquiry, spread is calculated as the median value of I_S for the final 100 generations of the optimiser.

3.3 Inquiry Framework

The inherent high dimensionality of many-objective optimisation presents both conceptual and computational challenges to the analysis of algorithm behaviour. Thus, the inquiry framework is aimed towards exploratory data analysis rather than statistically significant performance comparison. Following the methodology of Laumanns *et al* (2001), single replication results are generated for a wide variety of configuration instances (each representing a particular EE trade-off setting) to yield a *response map* in optimiser design-space. The use of multiple replications is still regarded as preferable, but this is computationally impractical for this inquiry. Note that spatial similarity between optimiser responses arguably provides some support for statistical confidence (or otherwise) in the observed behaviour.

The configuration of the recombination operator, via p_{ic} and η_c, provides suitable control over the EE trade-off. Optimiser responses have been obtained for all pair-wise permutations from sample sets of p_{ic} and η_c, with elements chosen according to a heuristic, pseudo-logarithmic scale that helps to show relativity within and between different response maps. The maps themselves portray scalar summary statistics for each overall response, such as proximity and spread indicator values.

In Section 4 optimiser responses, measured over 1000 generations, are generated for varying M with the population size fixed at 100. Deb (2001) has suggested that a main method for coping with large M is to increase the population size, since this will tend to reduce the proportion of the population that is non-dominated and thus provide improved dominance-based discrimination. Whilst this approach is unlikely to be practical in many real-world applications, where the computational cost of evaluating a candidate solution may be very high, the effect of population size is considered at the analysis stage of the inquiry in Section 5.

3.4 Presentation of Results

An example response map for I_P is shown in Figure 1. Performance for each $\{p_{ic}, \eta_c\}$ setting is indicated by a grey-scale square at the appropriate location. Lighter shades of grey indicate better proximity, as shown by the colour-bars of indicator values to the right of the map. A region of good proximity is evident for η_c values in the range [0 50] together with variation probabilities in the order of [0.01 1]. Conversely, a relatively poor value of proximity is evident for $\{p_{ic} = 5\times10^{-05}, \eta_c = 5\}$. The grid squares highlighted by a solid boundary correspond to configurations that exceed a performance threshold for proximity of -0.5: $\{p_{ic} = 0.05, \eta_c = 100\}$ is one such example in Figure 1. For the spread response maps, such as that provided in Figure 4, performance is highlighted for I_S in the range [0.75 1.25].

4 Results

A broad region of good proximity for intermediate to high p_{ic} and $1/\eta_c$ is evident for the map obtained for DTLZ2(3) shown in Figure 1. As the number of objectives, M, is increased to six, as shown in Figure 2, this sweet-spot contracts to areas of intermediate p_{ic}. For high p_{ic}, proximity values are worse than those that would be obtained from a random sample. This divergence behaviour becomes even more extensive as M is increased still further. Proximity results for DTLZ2(12) are shown in Figure 3.

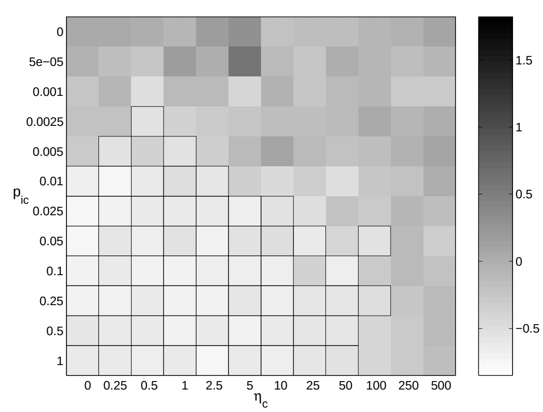

Figure 1: NSGA-II proximity map for DTLZ2(3)

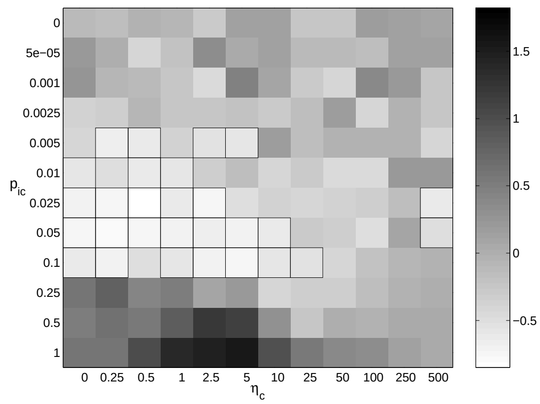

Figure 2: NSGA-II proximity map for DTLZ2(6)

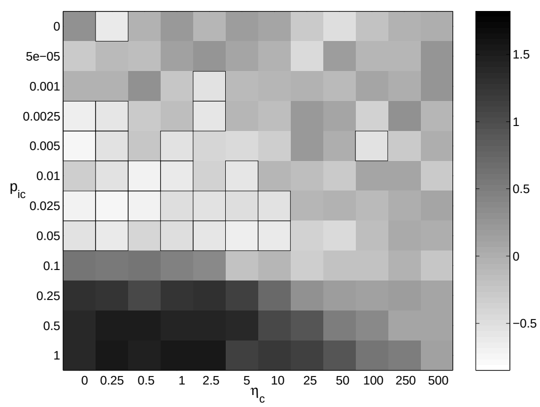

Figure 3: NSGA-II proximity map for DTLZ2(12)

The NSGA-II spread map for DTLZ2(3) is shown in Figure 4. A large sweet-spot is evident, largely corresponding to the regions of good proximity identified in Figure 1. Very small spread values are obtained for low p_{ic}, indicating that the approximation set represents a highly limited section of objective-space.

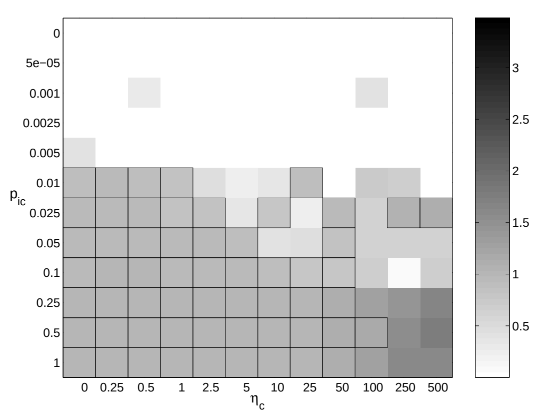

Figure 4: NSGA-II spread map for DTLZ2(3)

As the number of objectives is increased, the region of good spread becomes limited to intermediate p_{ic}. The sweet-spot is broadest in the area of high expected recombination perturbation (corresponding to low η_c) for DTLZ2(6) as shown in Figure 5. As *M* is increased still further the band of good spread becomes even thinner. The spread map for DTLZ2(12) is shown in Figure 6. For high p_{ic}, the approximation set is spread widely throughout non-optimal regions of objective-space. The relationship to poor proximity is clear through comparison with Figure 3.

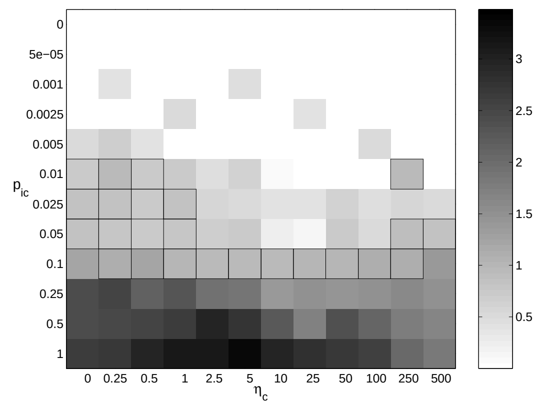

Figure 5: NSGA-II spread map for DTLZ2(6)

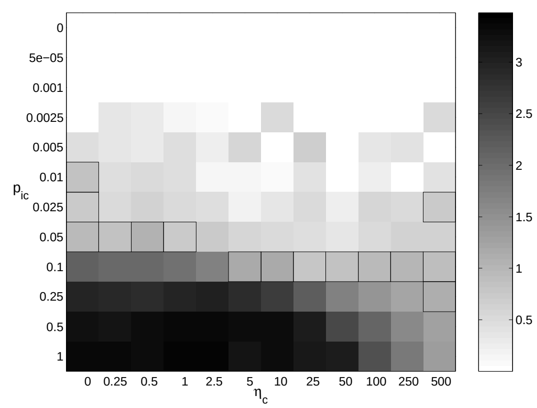

Figure 6: NSGA-II spread map for DTLZ2(12)

The standard SBX configuration for real-parameter optimisation tasks is $\{p_{ic} = 0.5, \eta_c = 15\}$ with $p_c = 1.0$ (Deb *et al*, 2002a). As evident from Figures 1 and 4, these choices are appropriate for small *M*. However, as *M* is increased the settings would appear to become unacceptable. Figures 3 and 6 suggest that the forthcoming approximation set would be spread widely through regions of objective-space with very poor proximity to the true Pareto front.

5 Analysis

Explanations for the observed results in Section 4 can be found through consideration of the selection and variation processes of NSGA-II. The effect of the diversity-promoting mechanisms is considered in Section 5.1. The effect of population size on the quality of achieved approximation sets is studied in Section 5.2.

5.1 Active Diversity Promotion

Diversity promotion mechanisms are present in both the selection-for-variation and selection-for-survival aspects of NSGA-II, as detailed in Section 2. The DTLZ2(6) proximity map obtained for the optimiser when the diversity mechanisms are removed is shown in Figure 7. Comparing this to the equivalent results for the full version of NSGA-II in Figure 2, it can be seen that the proximity for high p_{ic} is much improved. Elsewhere, performance remains largely equivalent.

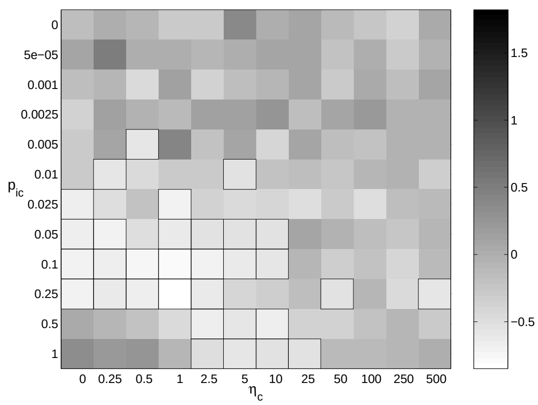

Figure 7: Non-crowding proximity map for DTLZ2(6)

For a fixed population size, the proportion of the population that is locally non-dominated is known to increase rapidly with *M*. Thus, for tasks with more than a small number of objectives, the secondary density-based selection mechanisms will be active in the original NSGA-II. As illustrated in Figure 8, the crowding distance density estimator will bias in favour of boundary solutions, other remote solutions, and the immediate neighbours of remote solutions. The volume of feasible objective-space increases with *M*, thus providing more opportunity for a solution to be remote, distant from the global surface, and still be locally non-dominated. In these circumstances, active diversity promotion will bias the search towards solutions with poor proximity. This produces difficulties in the overall optimiser if the variation operators are not capable of producing children that dominate their parents. The probability of this occuring is larger in regions of high p_{ic} for SBX on DTLZ2 and also increases with *M*. Hence, approximation sets with a high proximity cost are produced for NSGA-II in regions of high p_{ic} for anything other than low *M*. In regions of lower p_{ic}, SBX success rates are higher and so the search progresses towards the global surface. Removal of diversity promotion also restricts poor-proximity behaviour. Note that the difficulties in producing children that will dominate existing parents, a problem known as *dominance resistance*, was first identified by Ikeda *et al* (2001) and, in the context of many-objective optimisation, by Deb *et al* (2002a).

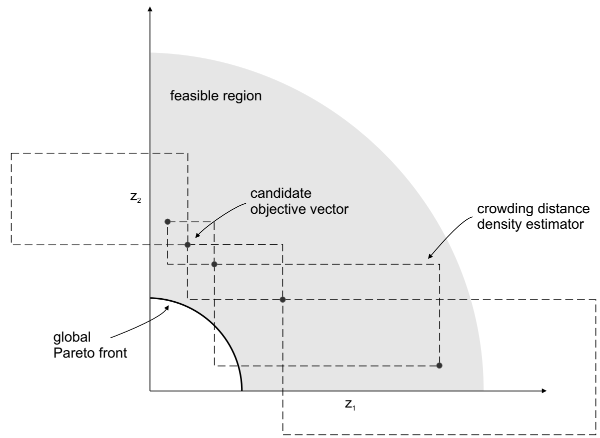

Figure 8: Crowding distance estimator

In regions of low p_{ic}, many of the children produced by recombination are copies of their parents. Since the density estimator used in NSGA-II is a type of first nearest-neighbour, as detailed in Section 2.1, these solutions will have the maximum possible density estimate (corresponding to a crowding distance of zero). This neutralises any density-dependent selection mechanisms because the densities of most solutions are identical. Thus, little difference is evident between algorithms that incorporate such a discriminator and those that do not for low p_{ic} configurations (compare the upper regions of Figure 7 to those of Figure 2).

In the absence of diversity-based selection, *genetic drift* causes population convergence unless the perturbations induced by variation are large. The resulting approximation set is only representative of a small area of objective-space, as shown by the spread map in Figure 9. In the lower-left of the map, p_{ic} and $1/\eta_c$ are large and thus the search retains a high degree of exploration.

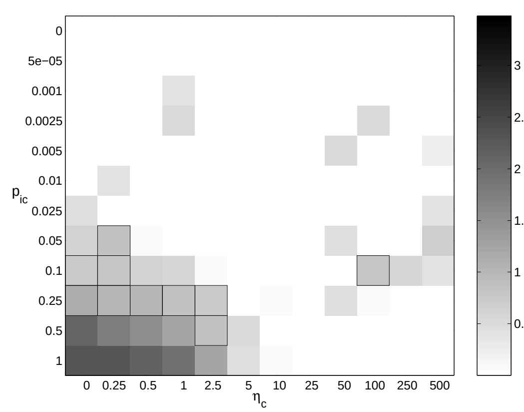

Figure 9: Non-crowding spread map for DTLZ2(6)

5.2 Population Size

Deb (2001) proposes the use of large population sizes as a potential method for achieving good many-objective optimisation results, since this will reduce the proportion of non-dominated solutions in the population and thus provide improved Pareto-based discrimation.

In practice, the use of large population sizes is often prohibitive in real-world applications because of the computational resources required to evaluate and process potential solutions. Nevertheless the benefits, in terms of approximation set quality, that can be obtained for larger population sizes remain a matter of interest.

NSGA-II proximity results, for a population size of 1000, solving DTLZ2(6) are shown in Figure 10. The sweet-spot extends further into regions of low p_{ic} than for the population of 100 shown in Figure 2. In areas of high p_{ic}, the proximity values are also slightly improved, since diversity-based selection will be less active for higher population sizes (discrimination is more likely to be based on dominance). Note that the classical SBX settings are still seen to correspond to poor proximity values.

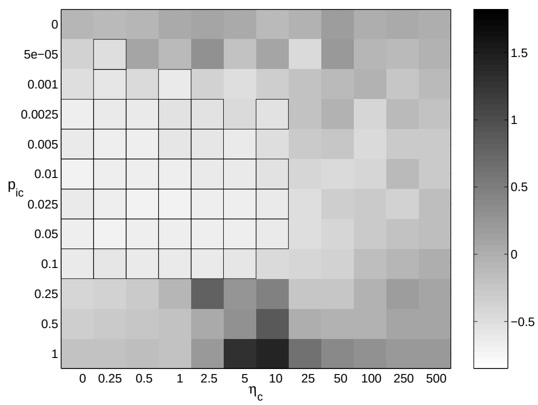

Figure 10: Large population proximity map for DTLZ2(6)

The DTLZ2(6) spread map obtained for a population size of 1000 is shown in Figure 11. The sweet-spot extends through all intermediate and low p_{ic}, being much larger than the corresponding region for a population size of 100 shown in Figure 5. The improved spread may reflect the increased diversity inherent in the use of larger sample sizes. Note also that the two-parent SBX operator requires sufficient population diversity in order to have exploratory properties.

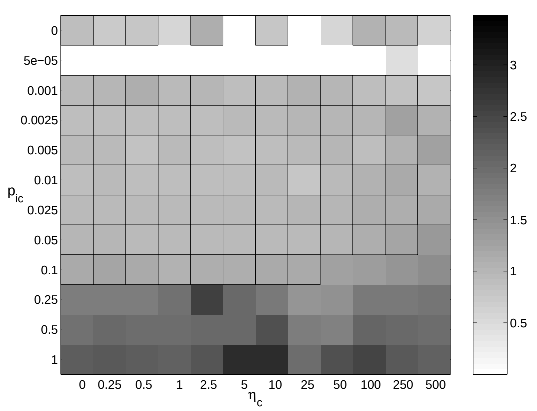

Figure 11: Large population spread map for DTLZ2(6)

6 Conclusions and Future Work

Insight into the many-objective optimisation process is required in order to assess the suitability of existing EMO techniques and to develop new methods for the simultaneous optimisation of many conflicting objectives.

This inquiry has considered the behaviour of NSGA-II on a many-objective task through analysis of the selection and variation processes that underpin the overall methodology. NSGA-II can be regarded as part of a family of algorithms that evaluate fitness via dominance comparisons and density estimation. Thus, generalised observations and insight into NSGA-II should be transferable to related algorithms, depending on the extent to which they share selection-for-variation and selection-for-survival mechanisms, and also how the EE trade-off within the variation operators relates to the EE trade-off in selection.

Simulations have indicated that the sweet-spot corresponding to good quality approximation sets contracts as the number of conflicting objectives is increased. The classical settings for recombination are shown to work well for small number of objectives, but become increasingly inappropriate as M increases.

EMO diversity promotion mechanisms can prove dangerous for many-objective optimisation. MOEAs are required to produce an approximation set with good diversity *in areas of good proximity*. For this reason, diversity promotion is generally regarded as a secondary selection operator (Bosman and Thierens, 2003). However, since the primary convergence-based operator uses the relative concept of Pareto dominance, if the proportion of non-dominated solutions is large then selection is based solely on diversity. As mentioned by Bosman and Thierens (2003), obtaining a good diversity is not a difficult task in itself, especially in many-objective space. Furthermore, the best diversity is often associated with very poor proximity values. Thus, if the current solutions are dominance resistant, then the many-objective search may evolve away from the true trade-off surface, with widespread dispersal of solutions in non-optimal objective-space.

Deb and Goldberg (1989) discovered that the recombination of spatially-dissimilar parent solutions often produced low-performance children, known as *lethals*. The authors were able to improve EA performance by only allowing recombination to occur between parents located within the same local neighbourhood. This process is known as *mating restriction*. In the context of many-objective optimisation, lethals could be considered to be non-dominated, remote solutions with a highly substandard component in one or more objectives. An exploration of the effect of mating restriction schemes in the context of diversity-promoting mechanisms for many-objective optimisation may prove rewarding. Also, a clear need exists to extend the analysis of evolutionary many-objective optimisation to include other classes of MOEA, other variation operators, and other test problems of a different underlying theme.

Acknowledgments

The authors would like to thank the anonymous reviewers for their helpful comments and suggestions. Robin Purshouse wishes to acknowledge the UK Engineering and Physical Sciences Research Council (EPSRC) for research studentship support.

Bibliography

Bosman, P.A.N. and Thierens, D., 2003. The balance between proximity and diversity in multiobjective evolutionary algorithms. *IEEE transactions on evolutionary computation*, 7 (2), 174-188.

Coello, C.A.C. *et al*, 2002. *Evolutionary algorithms for solving multi-objective problems*. New York: Kluwer Academic Publishers.

Corne, D.W. *et al*, 2000. The Pareto envelope-based selection algorithm for multiobjective optimization. *In:* M. Schoenauer *et al*, ed. *Proceedings of the parallel problem solving from nature VI conference,* Paris 16-20 September 2000. *Springer-Verlag, 839-848*.

Deb, K., 2001. *Multi-objective optimization using evolutionary algorithms*. Chichester: Wiley.

Deb, K. and Agrawal, R.B., 1995. Simulated binary crossover for continuous search space. *Complex systems,* 9 (2), 115-148.

Deb, K. and Goel, T., 2001. Controlled elitist non-dominated sorting genetic algorithms for better convergence. *In:* E. Zitzler *et al*, ed. *Proceedings of the first international conference on evolutionary multi-criterion optimization (EMO 2001),* Zurich 7-9 March 2001. *Springer-Verlag, 67-81*.

Deb, K. and Goldberg, D.E., 1989. An investigation of niche and species formation in genetic function optimization. *In:* J.D. Schaffer, ed. *Proceedings of the third international conference on genetic algorithms,* 1989. *Morgan Kaufmann, 42-50*.

Deb, K. and Jain, S., 2002. *Running performance metrics for evolutionary multi-objective optimization*. Kanpur: Indian Institute of Technology Kanpur, (KanGAL Report No. 2002004).

Deb, K. *et al*, 2002a. Scalable multi-objective optimization test problems. *In:* D.B. Fogel *et al*, ed. *Proceedings of the 2002 congress on evolutionary computation (CEC 2002),* Honolulu 12-17 May 2002. *IEEE Press, 825-830*.

Deb, K *et al*, 2002b. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE transactions on evolutionary computation,* 6 (2), 182-197.

Deb, K. *et al*, 2003. Towards a quick computation of well-spread Pareto-optimal solutions. *In:* C.M. Fonseca *et al*, ed. *Proceedings of the second international conference on evolutionary multi-criterion optimization (EMO 2003),* Faro 8-11 April 2003. *Springer-Verlag, 222-236*.

Edgeworth, F.Y., 1932. *Mathematical psychics. An essay on the application of mathematics to the moral sciences*. Reprint. London: The London School of Economics and Political Science.

Farina, M. and Amato, P., 2002. On the optimal solution definition for many-criteria optimization problems. *In:* J. Keller and O. Nasraoui, ed. *Proceedings of the NAFIPS-FLINT international conference 2002,* New Orleans 27-29 June 2002. *IEEE Press, 233-238*.

Fonseca, C.M. and Fleming, P.J., 1993. Genetic algorithms for multiobjective optimization: formulation, discussion and generalization. *In:* S. Forrest, ed. *Proceedings of the fifth international conference on genetic algorithms,* Urbana-Champaign June 1993. *Morgan Kaufmann, 416-423*.

Goldberg, D.E., 1989. *Genetic algorithms in search, optimization and machine learning*. Reading: Addison-Wesley.

Goldberg, D.E., 1998. *The race, the hurdle, and the sweet spot: lessons from genetic algorithms for the automation of design innovation and creativity*. Urbana: University of Illinois at Urbana-Champaign, (IlliGAL Report No. 98007).

Ikeda, K. *et al*, 2001. Failure of Pareto-based MOEAs: does non-dominated really mean near to optimal? *In:* J-H. Kim, ed. *Proceedings of the 2001 congress on evolutionary computation (CEC 2001),* Seoul 27-30 May 2001. *IEEE Press, 957-962*.

Khare, V. *et al*, 2003. Performance scaling of multi-objective evolutionary algorithms. *In:* C.M. Fonseca *et al,* ed. *Proceedings of the second international conference on evolutionary multi-criterion optimization (EMO 2003),* Faro 8-11 April 2003. *Springer-Verlag, 376-390*.

Knowles, J.D., 2002. *Local-search and hybrid evolutionary algorithms for Pareto optimization*. Thesis (PhD). University of Reading.

Laumanns, M. *et al*, 2001. On the effects of archiving, elitism, and density based selection in evolutionary multi-objective optimization. *In:* E. Zitzler *et al*, ed. *Proceedings of the first international conference on evolutionary multi-criterion optimization (EMO 2001),* Zurich 7-9 March 2001. *Springer-Verlag, 181-198*.

Purshouse, R.C. and Fleming, P.J., 2002. Why use elitism and sharing in a multi-objective genetic algorithm? *In:* W.B. Langdon *et al*, ed. *Proceedings of the genetic and evolutionary computation conference (GECCO 2002),* New York City 9-13 July 2002. *Morgan Kaufmann, 520-527*.

Purshouse, R.C. and Fleming, P.J., 2003a. Conflict, harmony, and independence: relationships in evolutionary multi-criterion optimisation. *In:* C.M. Fonseca *et al*, ed. *Proceedings of the second international conference on evolutionary multi-criterion optimization (EMO 2003),* Faro 8-11 April 2003. *Springer-Verlag, 16-30*.

Purshouse, R.C. and Fleming, P.J., 2003b. *An exploration of the many-objective optimisation process for a class of evolutionary algorithms*. Sheffield: University of Sheffield, (ACSE Research Report No. 837).

Veldhuizen, D.A.V., 1999. *Multiobjective evolutionary algorithms: classifications, analyses, and new innovations*. Thesis (PhD). US Air Force Institute of Technology.

Zitzler, E., 1999. *Evolutionary algorithms for multiobjective optimization: methods and applications*. Thesis (PhD). Swiss Federal Institute of Technology.

Zitzler, E. *et al*, 2001. *SPEA2: Improving the strength Pareto evolutionary algorithm.* Zurich: Swiss Federal Institute of Technology, (TIK-Report 103).

Zitzler, E. *et al*, 2003. Performance assessment of multiobjective optimizers: an analysis and review. *IEEE transactions on evolutionary computation,* 7 (2), 117-132.