An Adaptive Divide-and-Conquer Methodology for Evolutionary Multi-Criterion Optimisation

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. Improved sample-based trade-off surface representations for large numbers of performance criteria can be achieved by dividing the global problem into groups of independent, parallel sub-problems, where possible. This paper describes a progressive criterion-space decomposition methodology for evolutionary optimisers, which uses concepts from parallel evolutionary algorithms and nonparametric statistics. The method is evaluated both quantitatively and qualitatively using a rigorous experimental framework. Proof-of-principle results confirm the potential of the adaptive divide-and-conquer strategy.

1 Introduction

Treatments for the evolutionary optimisation of large numbers of criteria can potentially be discovered through analysis of the relationships between the criteria [1]. This paper considers *divide-and-conquer* strategies based on the concept of *independence* between a pair of criteria, in which performance in each criterion is entirely unrelated to performance in the other.

An *independent set* is herein defined as a set whose members are linked by dependencies, and for which no dependencies exist with elements external to the set. Consider a problem with n independent sets of criteria $[\mathbf{z}_1,\ldots,\mathbf{z}_n]$ and associated independent sets of decision variables $[\mathbf{x}_1,\ldots,\mathbf{x}_n]$. If knowledge of these sets is available then the global problem, $\mathbf{p}$, can be decomposed into a group of parallel sub-problems $[p_1,\ldots,p_n]$ that can be optimised independently of each other to ultimately yield n independent trade-off surfaces.

This paper demonstrates the benefit of using such a divide-and-conquer strategy when the correct decompositions are known in advance. It also proposes a general methodology for identifying, and subsequently exploiting, the decomposition during the optimisation process. An empirical framework is described in Sect. 2, which is then used to establish the case for divide-and-conquer in Sect. 3. An on-line adaptive strategy is proposed in Sect. 4 that exploits the iterative, population-based nature of the evolutionary computing paradigm. Independent sets of criteria are identified using nonparametric statistical methods of independence testing. Sub-populations are assigned to the optimisation of each set, with migration between these occurring as the decomposition is revised over the course of the optimisation. Proof-of-principle results are presented in Sect. 5, together with a discussion of issues raised by the study.

2 Experimental Methodology

2.1 Baseline Algorithm

The baseline evolutionary multi-criterion optimiser chosen in this work is an elitist multi-objective genetic algorithm (MOGA) [2]. An overview is shown in Table 1. Parameter settings are derived from the literature and tuning has not been attempted.

Table 1. Baseline multi-objective evolutionary algorithm (MOEA) used in the study

EMO component	Strategy
General	Total population = 100*n*, Generations = 250
Elitism	Ceiling of 20%-of-population-size of non-dominated solutions preserved. Reduction using SPEA2 clustering [3].
Selection	Binary tournament selection using Pareto-based ranking [4].
Representation	Concatenation of real number decision variables. Accuracy bounded by machine precision.
Variation	Uniform SBX crossover with η_c = 15, exchange probability = 0.5 and crossover probability = 1 [5]. Element-wise polynomial mutation with η_m = 20 and mutation probability = (chromosome length)$^{-1}$ [6].

2.2 Test Functions

A simple way to create independent multi-criterion test functions is to concatenate existing test problems from the literature, within which dependencies exist between all criteria. In this proof-of-principle study, only the test function *ZDT-1* proposed in [7] is used. The recently proposed scalable problems [8] will be used in later studies.

Definition 1. Concatenated ZDT-1 test function.

$$\text{Min.}\quad \mathbf{z} = \left[z_1\left(x_1\right), z_2\left(x_1,\ldots,x_m\right),\ldots, z_{2n-1}\left(x_{(n-1)m+1}\right), z_{2n}\left(x_{(n-1)m+1},\ldots,x_{nm}\right)\right],$$

$$\text{with}\quad z_{2i-1} = x_{(i-1)m+1},$$

$$z_{2i} = 1 - \sqrt{z_{2i-1} \Big/ \left(1 + \left[9/(m-1)\right] \sum_{j=(i-1)m+2}^{im} x_j \right)},$$

where $i = [1,\ldots,n]$ is a particular sub-problem, m = 30 is the number of decision variables per sub-problem, and $x_j \in [0,1] \forall j$. An instance is denoted by *C-ZDT-1(i)*.

The global solution to this problem is a set of bi-criterion trade-off surfaces (z_1 versus z_2, z_3 versus z_4, and so forth). Each trade-off surface is a convex curve in the region $[0\ 1]^2$ (for which the summation in Definition 1 is zero). The *ideal vector* is [0 0]. The *anti-ideal vector* of worst possible performance in each criterion is [1 10].

2.3 Performance Metrics

Hypervolume. A quantitative measure of the quality of a trade-off surface is made using the hypervolume *S* unary performance metric [9]. The hypervolume metric measures the amount of criterion-space dominated by the obtained non-dominated front, and is one of the best unary measures currently available, although it has limitations [10][11]. The anti-ideal vector is taken as the reference point. The metric is normalised using the hypervolume of the ideal vector, as illustrated in Fig. 1a.

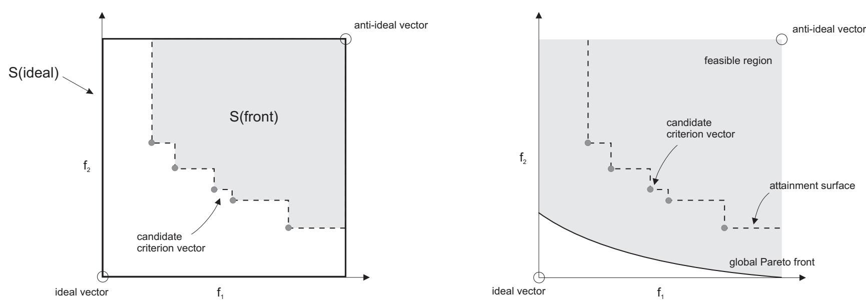

Fig. 1. (a) Hypervolume metric, **(b)** Attainment surface

Attainment surfaces. Given a set of non-dominated vectors produced by a single run of an MOEA, the attainment surface is the boundary in criterion-space that separates the region that is dominated by or equal to the set from the region that is non-dominated. An example attainment surface is shown in Fig. 1b. Performance across multiple runs can be described in terms of regions that are dominated in a given proportion of runs and can be interpreted probabilistically. For example, the 50%-attainment surface is similar to the median, whilst the 25%- and 75%-attainment surfaces are akin to the quartiles of a distribution. Univariate statistical tests can be performed on attainment surfaces using the concept of auxiliary lines [12][13]. However, in this paper, the technique is simply used to provide a qualitative comparison of the median performance of two algorithms, supported by quantitative hypervolume-based significance testing.

2.4 Analysis Methods

For the type of MOEA described in Table 1, the final population represents an appropriate data set upon which to measure performance. 35 runs of each algorithm configuration have been conducted in order to generate statistically reliable results. Quantitative performance is then expressed in the distribution of obtained hypervolumes. A comparison between configurations is made via the difference between the means of the distributions.

The significance of the observed result is assessed using the simple, yet effective, nonparametric method of *randomisation testing* [14]. The central premise of the method is that, if the observed result has arisen by chance, then this value will not appear unusual in a distribution of results obtained through many random relabellings of the samples. Let *S1* be the distribution of hypervolume metrics for *algorithm_1*, and let *S2* be the corresponding distribution for *algorithm_2*. The randomisation method for *algorithm_1 versus algorithm_2* proceeds as follows:

- Subtract the mean of *S2* from the mean of *S1*: this is the observed difference.
- Randomly reallocate half of all samples to one algorithm and half to the other. Compute the difference between the means as before.
- Repeat Step 2 until 5000 randomised differences have been generated, and construct a distribution of these values.
- If the observed value is within the central 99% of the distribution, then accept the null hypothesis that there is no performance difference between the algorithms. Otherwise consider the alternative hypotheses. Since optimal performance is achieved by maximising hypervolume, if the observed value falls to the left of the distribution then there is strong evidence to suggest that *algorithm_2* has outperformed *algorithm_1*. If the observed result falls to the right, then superior performance is indicated for *algorithm_1*. This is a two-tailed test at the 1%-level.

2.5 Presentation of results

Comparisons of *algorithm_1* versus *algorithm_2* for n = [1,…,4] are summarised within a single figure such as Fig. 2, for which *algorithm_1* is the baseline algorithm of Table 1 (with no decomposition) and *algorithm_2* is a parallel version with an a priori decomposition of both criterion-space and decision-space. Region (a) shows the validation case of one independent set, *C-ZDT-1(1)*, whilst regions (b), (c), and (d) show two, three, and four sets respectively. Within each region, each row indicates a bi-criterion comparison. The left-hand column shows the results of the randomisation test on hypervolume (if the observed value, indicated by the filled circle, lies to the right of the distribution then this favours *algorithm_1*), whilst the right-hand column shows the median attainment surfaces (the unbroken line is *algorithm_1*).

3 The Effect of Independence

The potential of a divide-and-conquer strategy can be examined by comparing a global solution to the concatenated ZDT-1 problem to a priori correct decompositions in terms of decision-space, criterion-space, or both. Consider a scheme in which a sub-population of 100 individuals is evolved in isolation for each independent set. Each EA uses only the relevant criteria and decision variables. This is compared to a global approach, with a single population of size 100n, using all criteria and decision variables in Fig. 2. Substantially improved performance is shown for the divide-and-conquer scheme.

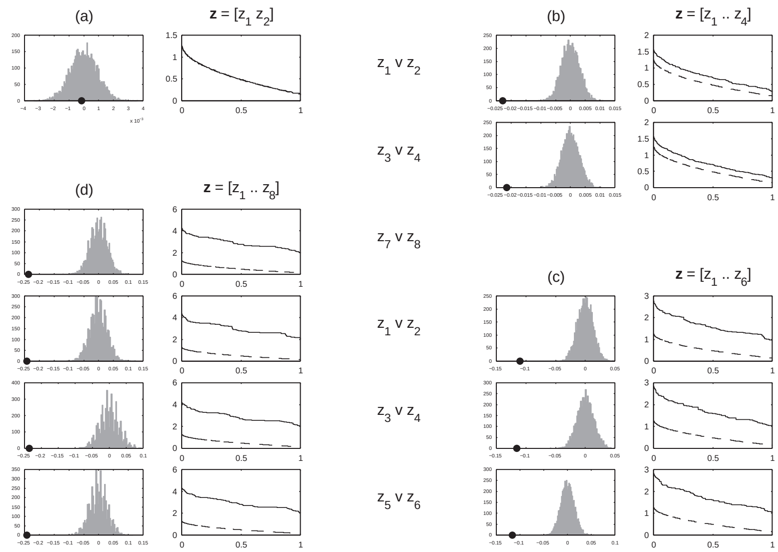

Fig. 2. Global model versus decomposition of both criterion-space and decision-space

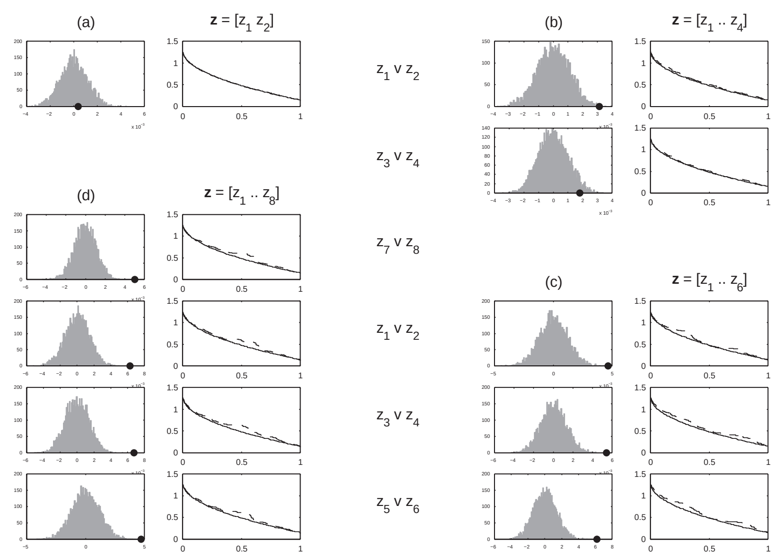

Fig. 3. Decomposition of both criterion-space and decision-space versus decision-space decomposition alone

To clarify which parts of the decomposition are important, sub-population schemes that decompose decision-space whilst treating criterion-space globally and vice versa are now considered. In the decision-space scheme, each 100 individual sub-population operates on the correct subset of decision variables. However, fitness values (and elitism) are globally determined. Elite solutions are reinserted into the most appropriate sub-population depending on their ranking on local criterion sets. Assignment is random in the case of a tie. Performance is compared to the ideal decomposition in Fig. 3. It is evident that decision-space decomposition alone is not responsible for the results in Fig. 2, and that the quality of the trade-off surfaces deteriorates with *n*. The attainment surfaces for cases (c) and (d) suggest that the global treatment of criteria may be affecting the shape of the identified surface.

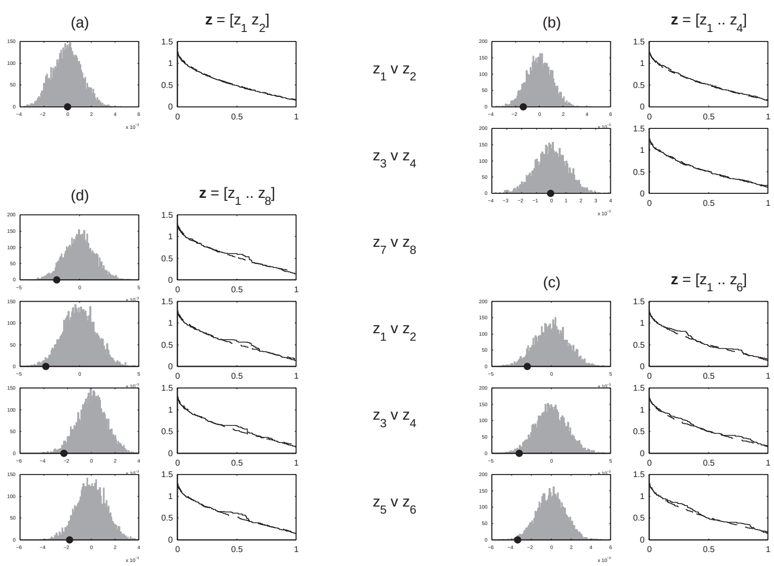

Fig. 4. Decision-space decomposition versus criterion-space decomposition

In the criterion-space scheme, each sub-population operates on the correct subset of criteria (fitness and elitism is isolated within the sub-population), but the EA operates on the global set of decision variables. A comparison with the decision-space method is shown in Fig. 4. No statistically significant performance difference is evident in any of the cases. Thus, criterion-space decomposition alone is also not responsible for the achievement in Fig. 2. Note that if single-point rather than uniform crossover had been used then results would have been much worse for the global treatment of decision-space since, for the former operator, the probability of affecting any single element of the chromosome (and thus the relevant section) decreases with chromosome length.

The above results indicate that a sub-population-based decomposition of either criterion-space or decision-space can significantly benefit performance. Best results are obtained when both domains are decomposed. Given that, in general, the correct decomposition for either domain is not known in advance, the choice of domain will depend on which is less demanding to analyse. Note that if criterion-space is decomposed then decision-space decomposition is also required at some point in order to synthesise a global solution. However, the converse is not the case. Decomposition may be a priori, progressive, or a posteriori with respect to the optimisation. A correct early decomposition in both spaces would be ideal but this may not be achievable.

4 Exploiting Independence via Criterion-Space Decomposition

4.1 Overview of the Methodology

A progressive decomposition of criteria, together with a retrospective decomposition of decision variables, is proposed in this paper. This is appropriate for problem domains where the number of criteria is significantly fewer than the number of decision variables. A sub-population approach is taken, in which the selection probability of an individual in a sub-population is determined using only the subset of criteria assigned to that sub-population. The topology of this parallel model can vary over the course of the optimisation.

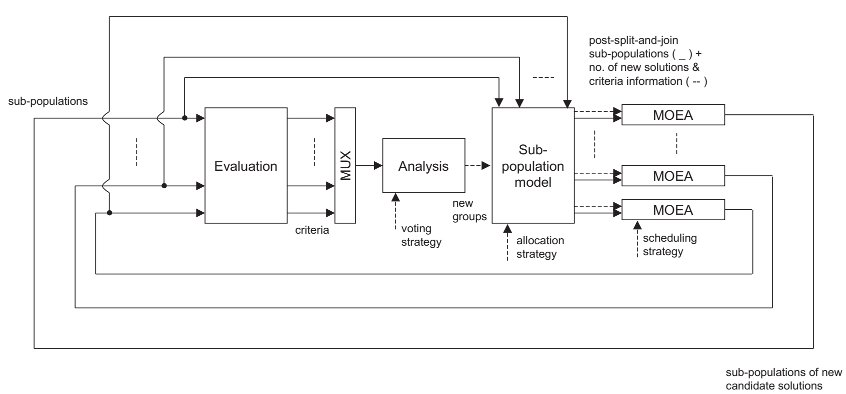

Fig. 5. Schematic overview of the progressive decomposition methodology

An overall schematic of the technique is given in Fig. 5. The process begins with a global population model. The multi-criterion performance of each candidate solution is then obtained. From the perspective of a single criterion, the population provides a set of observations for that criterion. Pair-wise statistical tests for independence are then performed for all possible pairs of criteria to determine between which criteria dependencies exist. Linkages are created for each dependent relationship. A sub-problem is then identified as a linked set of criteria. This concept is illustrated for an ideal decomposition of *C-ZDT-1(2)* in Fig. 6. Of all pair-wise dependency tests, significant dependencies have been identified between z_1 and z_2, and between z_3 and z_4.

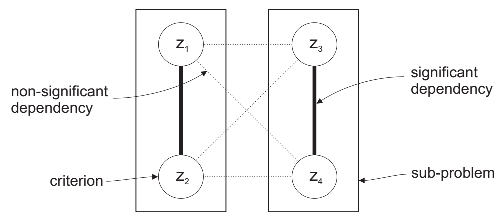

Fig. 6. Identification of sub-problems [z_1 z_2] and [z_3 z_4] via linkage

The new topology of the population model follows from the decomposition. *Split* and *join* operations are implemented to allow criteria (and associated candidate solutions) to migrate between sub-problems as appropriate.

When each new sub-population has been formed, selection probabilities and the identification (and management) of elites are determined using the current subset of criteria. Performance across all other criteria is ignored. Selection and variation operators are then applied within the boundaries of the sub-population. The size of the resulting new sub-population is pre-determined by the population management process.

All new solutions are evaluated across the complete set of criteria. This new data is then used to determine an updated decomposition. In this study, the update is performed at every generation, although in general it could be performed according to any other schedule. The process then continues in the fashion described above.

4.2 Population Management

The population topology is dependent on the identified decomposition. This can change during the course of the optimisation, thus requiring some sub-problem resources to be reallocated elsewhere. Operations are required to split some portion of the candidate solutions from a sub-population and subsequently join this portion on to another sub-population.

The size of the split is decided using an allocation strategy. Since the number of candidate solutions required to represent a trade-off surface grows exponentially with the number of criteria, *k*, it would seem a reasonable heuristic to use an exponential allocation strategy such as 2^k. As an example, consider a four-criterion problem, $k = 4$, with an initial decomposition of {[z_1 z_2],[z_3 z_4]}. The suggested new decomposition is {[z_1],[z_2 z_3 z_4]}. The situation is depicted in Fig. 7.

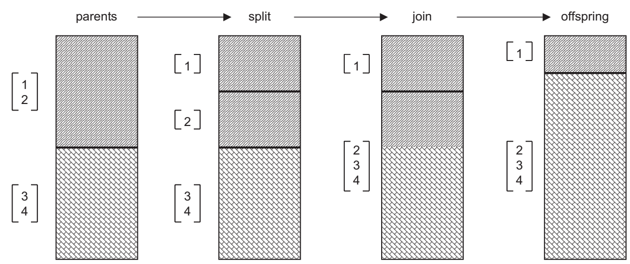

Fig. 7. Example of the split and join operations

Prior to reallocation, each sub-population should have a proportion $2^2/(2^2+2^2) = 1/2$ of the available resources. z_2 is now to be linked with z_3 and z_4, and must thus be split from its grouping with z_1. Both z_1 and z_2 will receive $2^1/(2^1+2^1) = 1/2$ of the resources in this sub-population. The actual candidate solutions to be assigned to each part of the split are determined randomly. The resources allocated to z_2 are then added to the [z_3 z_4] sub-population. Now z_1 has 1/4 of the resources, whilst [z_2 z_3 z_4] has 3/4 of the resources. The selection and variation operations are then used to return to the required proportions of $2^1/(2^1+2^3) = 1/5$ for z_1 and $2^3/(2^1+2^3) = 4/5$ for [z_2 z_3 z_4].

4.3 Tests for Independence

A sub-problem group is generated by collecting all criteria that are linked by observed pair-wise dependencies. In this sub-section, the tests used to determine if a connection should be made are introduced. Several tests for variable independence based on sample data exist in the statistics literature [15]. Two nonparametric procedures, the first based on the *Kendall K* statistic and the second on the *Blum-Kiefer-Rosenblatt D* statistic, are used in this work.

Both methods require special care for the handling of tied data. This is of concern in an evolutionary algorithm implementation since a particular solution may have more than a single copy in the current population. Large-sample approximations to each method have been implemented. This is possible because reasonably large population sizes have been used (100 individuals per independent bi-criterion set). All significance tests are two-tailed at the 1%-level, the null hypothesis being that the criteria are independent.

Kendall *K*. A distribution-free bivariate test for independence can be made using the Kendall sample correlation statistic, *K*. This statistic measures the level of *concordance* (as one variable increases/decreases, the other increases/decreases) against the level of *discordance* (as one variable increases/decreases, the other decreases/increases). This is somewhat analogous to the concepts of harmony and conflict in multi-criterion optimisation [1]. The standardised statistic can then be tested for significance using the normal distribution $N(0,1)$. Ties are handled using a modified paired sign statistic. A modified null variance is also used in the standardisation procedure. For further details, refer to [15, pp363-394].

The main concern with this method is that if $K = 0$ this does not necessarily imply that the two criteria are independent (although the converse is true). This restricts the applicability of the method beyond bi-criterion dependencies, where relationships may not be monotonous.

Blum-Kiefer-Rosenblatt *D*. As an alternative to the above test based on the sample correlation coefficient, Blum, Kiefer, and Rosenblatt's large-sample approximation to the *Hoeffding D* statistic has also been considered in this study. This test is able to detect a much broader class of alternatives to independence. For full details, refer to [15, pp408-413].

4.4 Decision-Space Decomposition: An Aside

Discussion. In the above methodology, and throughout the forthcoming empirical analysis of this method in Sect. 5, different sub-populations evolve solutions to different sets of criteria. Decision-space decomposition is not attempted. Thus, at the end of the optimisation process, complete candidate solutions exist for each criterion set. It is now unclear which decision variables relate to which criterion set. In order to finalise the global solution, solutions from each trade-off surface must be synthesised via partitioning of the decision variables.

An a posteriori decomposition, as described below, is simple to implement but has two clear disadvantages: (1) some reduction in EA efficiency will be incurred because the operators search over inactive areas of the chromosome (operators that are independent of chromosome length should be used), and (2) further analysis is required to obtain the global solution.

Method. A candidate solution should be selected at random from the overall final population. Each variable is then perturbed in turn and the effect on the criteria should be observed. The variable should be associated with whichever criteria are affected. Then, when the decision-maker selects a solution from the trade-off featuring a particular set of criteria, the corresponding decision variables are selected from the sub-population corresponding to this set.

This method requires as many extra candidate solution evaluations as there are decision variables in the problem. For the 4-set concatenated ZDT-1 test function, this is 120 evaluations or 30% of a single generation of the baseline algorithm.

Two special cases must be addressed: (1) If the perturbation of a decision variable affects criteria in more than one criterion subset, this indicates an invalid decomposition of criterion-space. Information of this kind could be used progressively to increase the robustness of the decomposition. (2) It is possible that no disturbance of criteria is seen when the decision variable is perturbed. Here, the alternatives are to consider another candidate solution or to consider more complicated variable interactions. This may also be an indication that the variable is globally redundant.

5 Preliminary Results

Proof-of-principle results for the adaptive divide-and-conquer strategy devised in Sect. 4 are presented herein for the concatenated ZDT-1 test function (Definition 1) with $n = [1,\ldots,4]$. A summary of the chosen strategy is given in Table 2. Both the Kendall K method and the Blum-Kiefer-Rosenblatt D method have been considered.

Table 2. Divide-and-conquer settings

EMO component	Strategy	
Independence test	(either)	Blum-Kiefer-Rosenblatt D
	(or)	Kendall K
Resource allocation	2^k	
Schedule	Revise the decomposition every generation	

5.1 Blum-Kiefer-Rosenblatt *D* Results

The performance of this strategy when compared to the baseline case of no decomposition is shown in Fig. 8. Both the hypervolume metric results and the attainment surfaces indicate that the divide-and-conquer strategy produces trade-off surfaces of higher quality in cases where independence exists. However, the attainment surfaces also show that the absolute performance of the method degrades as more independent sets are included.

The degradation can be partially explained by considering the percentage of correct decompositions made by the algorithm at each generation (measured over the 35 runs) shown in Fig. 9a. As the number of independent sets increases, the proportion of correctly identified decompositions decreases rapidly. Note that this does not necessarily mean that the algorithm is making invalid decompositions or no decomposition: other valid decompositions exist, for example $\{[z_1\ z_2], [z_3\ z_4\ z_5\ z_6]\}$ for $n = 3$, but these are not globally optimal (also the number of possible decompositions increases exponentially with n.). Indeed, on no occasion did the test produce an invalid decomposition (identified independence when dependency exists). This is evident from plots of the decomposition history over the course of the optimisation. A typical history is depicted in Fig. 9b. Each criterion is labelled on the vertical axis, whilst the horizontal axis depicts the current generation of the evolution. At a particular generation, criteria that have been identified as an independent set are associated with a unique colour. Thus, as shown in Fig. 9b, at the initial generation z_1 and z_2 have been identified as a cluster (white), as have $[z_3\ z_4]$ (black), $[z_5\ z_6]$ (light grey), and $[z_7\ z_8]$ (dark grey). At generation 200 all the criteria have been grouped together, as indicated by the complete whiteness at this point in the graph. Note that there is no association between the colours across the generations.

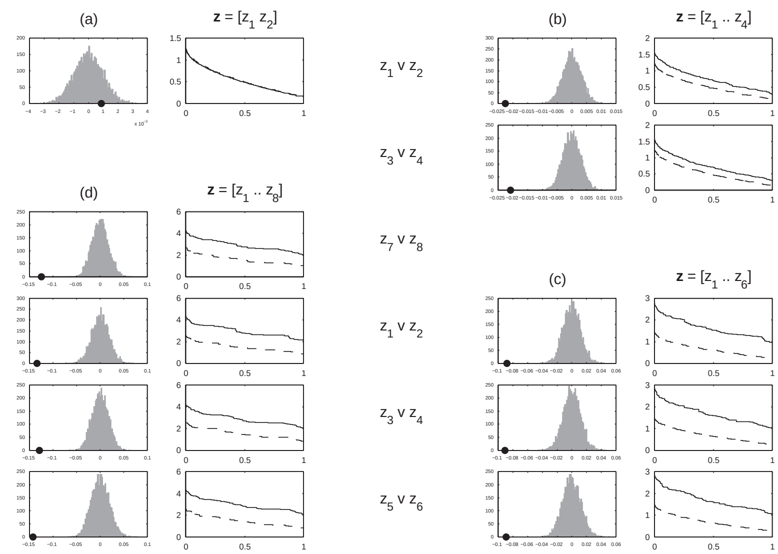

Fig. 8. No decomposition versus Blum-Kiefer-Rosenblatt *D* divide-and-conquer

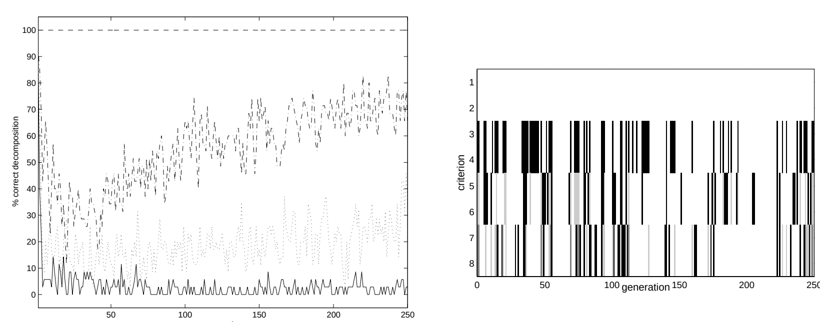

Fig. 9. Blum-Kiefer-Rosenblatt *D*: **(a)** Correct decompositions as a percentage of total runs over the course of the optimisation. ---- $[z_1\ z_2]$; -.-. $[z_1,\ldots,z_4]$; …. $[z_1,\ldots,z_6]$; ____ $[z_1,\ldots,z_8]$; **(b)** Typical decomposition history for a single replication. Each identified criterion cluster is represented by a colour (white, light grey, dark grey, black)

5.2 Kendall *K* Results

The performance of the divide-and-conquer algorithm with the Kendall *K* test for independence is compared to Blum-Kieffer-Rosenblatt *D* in Fig. 10. The former test appears to offer superior performance as the number of independent sets increases. No significant performance difference can be found for n = 2, but such a difference can be seen for two of the surfaces for n = 3, and every surface for n = 4.

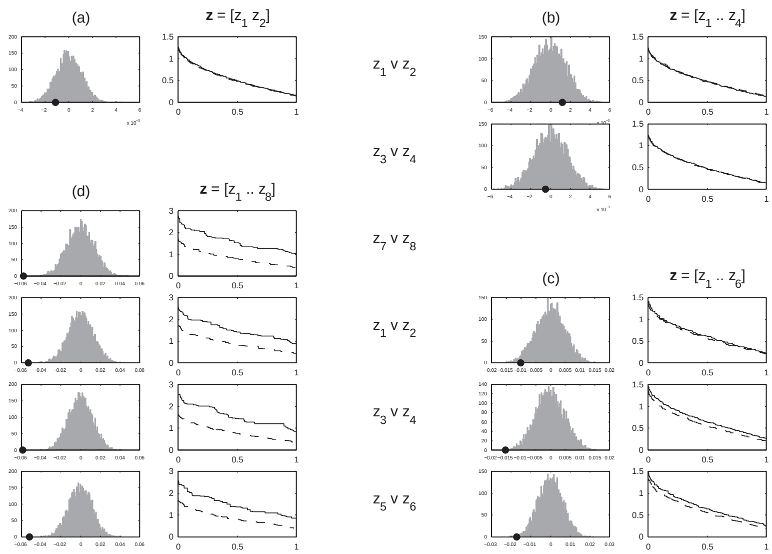

Fig. 10. Blum-Kiefer-Rosenblatt *D* divide-and-conquer versus Kendall *K* divide-and-conquer

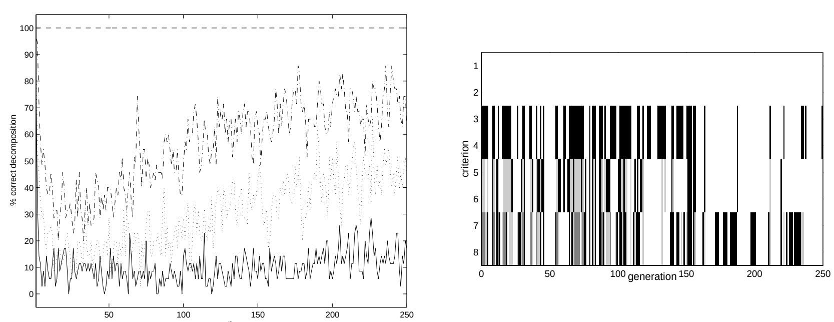

Fig. 11. Kendall *K*: **(a)** Correct decompositions as a percentage of total runs over the course of the optimisation. ---- [z_1 z_2]; -.-. [$z_1,\ldots,z_4$]; …. [$z_1,\ldots,z_6$]; ____ [$z_1,\ldots,z_8$]; **(b)** Typical decomposition history. Each identified criterion cluster is represented by a colour (white, light grey, dark grey, black)

This difference in performance can be explained by the plot of correct decompositions shown in Fig. 11a. Whilst the proportion of correct decompositions degrades as n increases, this degradation is not as severe as for Blum-Kiefer-Rosenblatt D (Fig. 9a). Also, from the typical decomposition history depicted in Fig. 11b, the valid decompositions tend to be of higher resolution than those developed by the alternative method (Fig. 9b).

5.3 Discussion

The obtained results show that the adaptive divide-and-conquer strategy offers substantially better performance than the global approach in terms of the quality of trade-off surfaces generated.

Of the two independence tests considered, Kendall K would appear more capable of finding good decompositions on the benchmark problem considered, especially as the number of independent sets increases. However, Kendall K may experience difficulties when the dimension of the trade-off surface increases, since it may incorrectly identify independence due to the variation in the nature of bi-criterion relationships over the surface (the relationship is not always conflicting, as it is for a bi-criterion problem). By contrast, Blum-Kiefer-Rosenblatt D offers a more robust search in these conditions, but is more conservative.

There is a clear need for the procedure to be robust (invalid decompositions should be avoided, although the progressive nature of the process may somewhat mitigate the damage from these), but conservativism should be minimised in order to increase the effectiveness of the methodology. Under these circumstances, it may be prudent to adopt a voting strategy, in which a decision is made based on the results from several tests for independence.

The adaptive divide-and-conquer strategy carries some overhead in terms of the test for independence and the sub-population management activity, which may be controlled using a scheduling strategy. This must be balanced against the improvements in the quality of the trade-off surfaces identified and the reduction in the complexity of the MOEA ranking and density estimation procedures.

6 Conclusion

This study has shown that, if feasible, a divide-and-conquer strategy can substantially improve MOEA performance. The decomposition may be made in either criterion-space or decision-space, with a joint decomposition proving the most effective. Criterion-space decomposition is particularly appealing because it reduces the complexity of the trade-off surfaces to be presented to the decision-maker. Furthermore, no loss of trade-off surface shape was observed for the ideal criterion-space decomposition as it was for the sole decomposition of decision variables.

An adaptive criterion-space decomposition methodology has been presented and proof-of-principle results on the concatenated ZDT-1 problem have been shown to be very encouraging. It should be noted that the approach is not confined to criterion-space: independence tests and identified linkages could equally well have been applied to decision variable data. In this case, the sub-populations would evolve different decision variables, whilst the evaluation would be global. It should also be possible to use the technique on both spaces simultaneously.

The main limitation of the methodology is the number of pair-wise comparisons that have to be conducted for high-dimensional spaces. In the concatenated ZDT-1 problem, analysis of the decision variables would be very compute-intensive. Further techniques for the progressive decomposition of high-dimensional spaces are required to complete the framework. Future work will also consider further concatenated problems, especially those of high dimension, with emphasis on real-world applications.

References

1. Purshouse, R.C., Fleming, P.J.: Conflict, Harmony, and Independence: Relationships in Evolutionary Multi-Criterion Optimisation. This volume. (2003)
2. Purshouse, R.C., Fleming, P.J.: Why use Elitism and Sharing in a Multi-Objective Genetic Algorithm? Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2002). (2002) 520-527
3. Zitzler, E., Laumanns, M., Thiele, L.: SPEA2: Improving the Strength Pareto Evolutionary Algorithm. TIK-Report 103, ETH Zürich. (2001)
4. Fonseca, C.M., Fleming, P.J.: Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization. Proceedings of the Fifth International Conference on Genetic Algorithms. (1993) 416-423
5. Deb, K., Agrawal, R.B.: Simulated binary crossover for continuous search space. Compex Systems **9** (1995) 115-148
6. Deb, K., Goyal, M.: A combined genetic adaptive search (GeneAS) for engineering design. Computer Science and Informatics **26** (1996) 30-45
7. Zitzler, E., Deb, K., Thiele, L.: Comparison of Multiobjective Evolutionary Algorithms: Empirical Results. Evolutionary Computation **8** (2000) 173-195
8. Deb, K., Thiele, L., Laumanns, M., Zitzler, E.: Scalable Multi-Objective Optimization Test Problems. Proceedings of the 2002 IEEE Congress on Evolutionary Computation (CEC 2002), Vol. 1. (2002) 825-830
9. Zitzler, E.: Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications. Doctoral dissertation, ETH Zürich. (1999)
10. Zitzler, E., Laumanns, M., Thiele, L., Fonseca, C.M., Grunert da Fonseca, V.: Why Quality Assessment Of Multiobjective Optimizers Is Difficult. Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2002). (2002) 666-674
11. Knowles, J.D., Corne, D.W.: On Metrics for Comparing Nondominated Sets. Proceedings of the 2002 IEEE Congress on Evolutionary Computation (CEC 2002), Vol. 1. (2002) 711-716
12. Fonseca, C.M., Fleming, P.J.: On the Performance Assessment and Comparison of Stochastic Multiobjective Optimizers. Parallel Problem Solving from Nature – PPSN IV. Lecture Notes in Computer Science, Vol. 1141 (1996) 584-593
13. Knowles, J.D., Corne, D.W.: Approximating the non-dominated front using the Pareto archived evolution strategy. Evolutionary Computation **8** (2000) 149-172
14. Manly, B.F.J.: Randomization and Monte Carlo Methods in Biology. Chapman and Hall, London New York Tokyo Melbourne Madras (1991)
15. Hollander, M., Wolfe, D.A.: Nonparametric Statistical Methods. 2nd edn. Wiley, New York Chichester Weinheim Brisbane Singapore Toronto (1999)