

# Conflict, Harmony, and Independence: Relationships in 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.** This paper contributes a platform for the treatment of large numbers of criteria in evolutionary multi-criterion optimisation theory through consideration of the relationships between pairs of criteria. In a conflicting relationship, as performance in one criterion is improved, performance in the other is seen to deteriorate. If the relationship is harmonious, improvement in one criterion is rewarded with simultaneous improvement in the other. The criteria may be independent of each other, where adjustment to one never affects adjustment to the other. Increasing numbers of conflicting criteria pose a great challenge to obtaining a good representation of the global trade-off hypersurface, which can be countered using decision-maker preferences. Increasing numbers of harmonious criteria have no effect on convergence to the surface but difficulties may arise in achieving a good distribution. The identification of independence presents the opportunity for a divide-and-conquer strategy that can improve the quality of trade-off surface representations.

## 1 Introduction

Theoretical evolutionary multi-criterion optimisation (EMO) studies generally consider a small number of objectives or criteria. The bi-criterion case is by far the most heavily studied. EMO applications, by contrast, are frequently more ambitious, with the number of treated criteria reaching double figures in some cases [1, pp207-290]. Hence, there is a very clear need to develop an understanding of the effects of increasing numbers of criteria on EMO. The recently proposed set of benchmark problems, which are scalable to any number of conflicting criteria, represent an important early step towards this aim [2].

This paper establishes a complementary platform for research into increasing numbers of criteria via consideration of the different types of pair-wise relationships between the criteria. A classification of possible relationships is offered in Sect. 2, and the notation used in the paper is introduced. Conflict between criteria is discussed in Sect. 3, whilst Sect. 4 considers harmonious objectives. The aim of a multi-objective evolutionary algorithm (MOEA) is generally regarded as to generate a sample-based representation of the Pareto optimal front, where the samples lie close to the true front and are well distributed across the front. The effects of increasing numbers of each type of criteria on both aspects of the quality of the trade-off surfaces produced are described, together with a review of methods for dealing with the difficulties that

arise. The case where criteria can be optimised independently of each other is introduced in Sect. 5. Qualitative studies of pair-wise relationships between criteria are not uncommon in the EMO community, especially in the case of real-world applications. These are discussed in Sect. 6, alongside similar quantitative methodologies from the multi-criterion decision-making (MCDM) discipline. Conclusions are drawn in Sect. 7.

Some of the concepts described in this paper are illustrated using an example result from a recently proposed multi-objective genetic algorithm (MOGA) [3] solving the 3-criterion *DTLZ2* benchmark problem [2]. The equations for this test function are provided in Definition 1. Note that all criteria are to be minimised.

*Definition 1.* 3-criterion DTLZ2 test function [2].

$$\text{Min. } z_1(\mathbf{x}) = \left[1 + g(x_3, \dots, x_{12})\right] \cos(x_1 \pi/2) \cos(x_2 \pi/2),$$

$$\text{Min. } z_2(\mathbf{x}) = \left[1 + g(x_3, \dots, x_{12})\right] \cos(x_1 \pi/2) \sin(x_2 \pi/2),$$

$$\text{Min. } z_3(\mathbf{x}) = \left[1 + g(x_3, \dots, x_{12})\right] \sin(x_1 \pi/2),$$

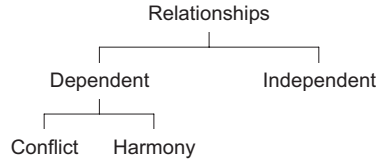
$$\text{where } 0 \leq x_i \leq 1, \text{ for } i = 1, 2, \dots, 12,$$

$$\text{and } g(x_3, \dots, x_{12}) = \sum_{i=3}^{12} (x_i - 0.5)^2.$$

## 2 Relationships Between Criteria

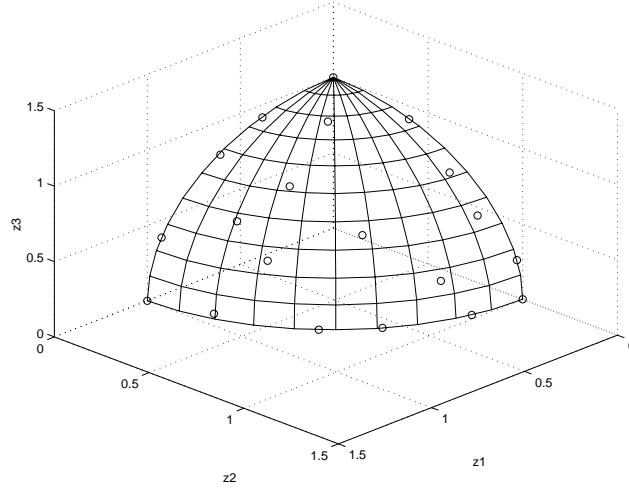
### 2.1 Classification

In theoretical EMO studies, the criteria are generally considered to be in some form of conflict with each other. Thus, in the bi-criterion case, the optimal solution is a one-dimensional (parametrically speaking) trade-off surface upon which conflict is always observed between the two criteria. However, other relationships can exist between criteria and these may vary within the search environment. A basic classification of possible relationships is offered in Fig. 1. These relationships are explained in the remainder of the paper.



**Fig. 1.** Classification of relationships between criteria

The dependency classifications are not necessarily mutually exclusive. For example, in the case of three conflicting criteria, there may be regions where two criteria can be improved simultaneously at the expense of the third. This is illustrated in Fig. 2 for the final on-line archive of a MOGA solving the 3-criterion DTLZ2 problem (see Definition 1). For example, ideal performance in  $z_2$  and  $z_3$  (evidence of harmony) can be achieved at the expense of nadir performance in  $z_1$  (evidence of conflict), as indicated by the left-most criterion vector in the figure. However, on the far right of the figure,  $z_1$  and  $z_3$  are now in harmony and are both in conflict with  $z_2$ . Thus, the nature of the relationships change across the Pareto front.



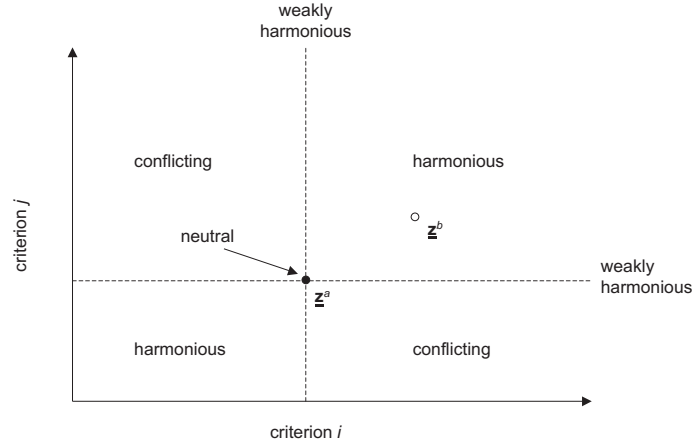
**Fig. 2.** Final on-line archive of MOGA (depicted as circles) solving DTLZ2, superimposed on the global trade-off surface

## 2.2 Notation

The following notation is used in the remainder of the paper:  $M$  is the number of criteria to be considered in the optimisation procedure,  $Z$  is the set of all realisable criterion vectors  $\underline{z} \in \mathbb{R}^M$ , and  $Z_R$  is a particular region of interest in criterion-space,  $Z_R \subseteq Z$ . If  $Z_R = Z$  then the relationship is said to be *global*, otherwise it is described as *local*. The case  $Z_R = Z^*$ , where  $Z^*$  is the Pareto optimal set, may be of particular interest since these are typically the relationships that will be presented to the decision-maker (DM). The DM is the entity that expresses preferences within the optimisation and selects an acceptable solution from the set returned by the optimisation process. The DM is usually one or several humans.

Let  $i$  and  $j$  be indices to particular criteria:  $i, j \in [1, \dots, M]$ . Let  $a$  and  $b$  be indices to individual criterion vector instances:  $a, b \in [1 \dots |Z_R|] : \underline{z}^a, \underline{z}^b \in Z_R$ . Also let  $(a, b)$  denote a pair of instances for which  $a \neq b$ . Minimisation is assumed throughout the paper without loss of generality.

The dependency relationships that can be identified via pair-wise analysis are summarised in Fig. 3. They are based on the position of criterion vector  $\underline{z}^b$  relative to the position of  $\underline{z}^a$ . These relationships are explored in more detail in Sect. 3 and Sect. 4 to follow.



**Fig. 3.** Dependency relationship regions between a pair of criteria,  $i$  and  $j$ , identified using the location of sample vector  $\underline{z}^b$  relative to that of  $\underline{z}^a$

### 3 Conflicting Criteria

#### 3.1 Definitions of Conflict

A relationship in which performance in one criterion is seen to deteriorate as performance in another is improved is described as *conflicting*. This is summarised by Definition 2 below and can be related to the  $\underline{z}^b$ -relative-to- $\underline{z}^a$  regions marked as such in Fig. 3.

*Definition 2.* Criteria  $i$  and  $j$  exhibit evidence of conflict according to the condition  $(\underline{z}_i^a < \underline{z}_i^b) \wedge (\underline{z}_j^a > \underline{z}_j^b)$ . If  $\exists (a, b)$  for which the condition holds then there is *no*

*conflict*, if  $\exists(a,b)$  then there is *conflict*, whilst if the condition holds  $\forall(a,b)$  then there is *total conflict*.

Note that no attempt has been made to define intermediate levels of conflict (or harmony, as discussed in Sect. 4) since this requires further DM preference information.

### 3.2 Effect on EMO

For  $M$  conflicting criteria, an  $(M - 1)$ -dimensional trade-off hypersurface exists in criterion-space. The number of samples required to achieve an adequate representation of the surface is exponential in  $M$ . Given a finite population, an evolutionary optimiser will encounter intractable difficulties in representing the surface when large numbers of conflicting criteria are considered. Even if such a representation were possible, the value to the DM of such a large number of candidate solutions is questionable.

Deb [4, pp400-405] has shown that the proportion of locally non-dominated criterion vectors in a finite randomly-generated sample becomes very large as the number of criteria increases (interestingly, it would appear to be possible to optimise all of the criteria independently in this particular study). Similar results were reported in [5b] for the final non-dominated set of a real-world MOEA application. Since dominance is used to drive the search towards the true Pareto front, there may be insufficient selective pressure to make such progress. The use of a large population can help reduce the proportion, but this is impractical for many real-world problems in which evaluation of a single candidate solution is very time-consuming. Also, the benefit would appear to become progressively weaker as the number of criteria increases.

Many MOEAs use some method of population density estimation, either to modify the selection probability of an individual or as part of the on-line archive acceptance procedure, to achieve a good distribution of solutions. Density estimation also becomes increasingly difficult as the dimensionality of the problem increases (the number of criteria, for density estimation in criterion-space). Due to the ‘curse of dimensionality’ (the sparseness of data in high dimensions), the ability to fully explore surfaces in greater than five dimensions is regarded as highly limited [6]. Statisticians generally use dimensionality reduction techniques prior to application of the estimator. This assumes that the ‘true’ structure of the surface is of lower dimension, but the potential for reduction may be limited for a trade-off surface in which all criteria are in conflict with each other.

### 3.3 Remedial Measures

**Preferences.** The exploitation of DM preferences is arguably the current best technique for handling large numbers of conflicting criteria. In this case, the aim of EMO is to achieve a good representation of trade-off regions of interest to the DM (essentially limiting the ambition of the optimiser by requiring it to represent only a subspace of the trade-off hypersurface). This sub-section provides a brief overview of

preference articulation within MCDM, before examining two popular techniques that use DM preferences to handle large numbers of conflicting criteria.

*Preference articulation overview.* Preference information can be classified into four distinct types:

- *definitions* of the criteria and problem domain,
- *requirements* of a solution,
- *abstract judgements* on the relative importance of criteria,
- *specific judgements* on a set of candidate solutions.

Definitions of the problem domain and the criteria to be optimised are essential to the optimisation process. However, this information is not necessarily readily available a priori and may contain a degree of uncertainty. Problem domain definitions include limits on the range of decision variables. Development of a good set of criteria is critical to the performance of the optimiser.

Requirements of a solution tend to be specified at criterion level, resulting in the formulation of an ideal or acceptable global solution. Definitions of unacceptability may also be given. In this approach, the DM could specify goals for particular objectives. These goals may be a requirement that must be met (a hard constraint) or may represent softer aspirations.

The DM may wish to provide abstract information about the relative importance of criteria. This may include precise, or perhaps rather vague, information about the level of trade-off that the DM is willing to accept between two criteria (for example, an acceptance of ‘an improvement of  $\Delta_1$  in criterion  $i$  in exchange for a detriment of  $\Delta_2$  in criterion  $j$ ’). Priority information may also be provided, in which the DM expresses a preference for some criteria being more important than others. This information may be imprecise. It may also be qualitative (‘much more important’) or quantitative (‘twice as important’). Priority information can be used to build a partial ordering of criteria, perhaps to be optimised lexicographically, or to determine weights in an aggregation of criteria.

Given a set of candidate solutions, the DM may be able to express preferences for some candidate solutions over others (perhaps allowing a partial ordering of potential solutions to be generated). Again, this information may be qualitative or quantitative, and is likely to be somewhat imprecise.

Preference articulation schemes are generally classified according to when the preference data is elicited from the DM:

- *a priori*, in which preference data is incorporated prior to execution of the optimiser,
- *progressive*, in which the information is requested and exploited during the optimisation process,
- *a posteriori*, where a solution is chosen from a group of results returned by the completed optimisation process.

Refer to [1, pp321-344] for a review of preference articulation schemes in the EMO literature.

*Aggregation.* One method for reducing the number of conflicting performance criteria is to combine several of them into a single optimisation criterion. Aggregation may be achieved by means of a weighted-sum, or a more complicated function. In this approach, the DM pre-specifies the trade-offs between the combined subset of criteria. This eliminates the requirement for the optimiser to represent this portion of the global trade-off surface. The inherent disadvantage of the approach is that the DM must be able to specify the required trade-off a priori. Also, any adjustments to the preferences will require a complete re-run of the optimisation. Nevertheless, this may be an appropriate technique, especially when faced with very large numbers of criteria.

*Goals and Priorities.* Greater flexibility can be achieved in terms of criterion reduction by exploiting goal values and priorities for various criteria, if these can be elicited from the DM. The *preferability* relation developed in [5a] unifies various classical operations research (OR) schemes based on goals and priorities and applies them within the context of EMO. In essence, the method adaptively switches on or off different criteria, from the perspective of the dominance relation, for each pair of vectors considered. The iterative nature of the EA paradigm can be exploited to update the preferences as information becomes progressively available to the DM.

**Dimension Reduction.** Existing dimensionality reduction techniques could be used to transform criterion-space into a lower dimension. This could be done prior to the optimisation, based on some preliminary analysis, or could be updated on-line as the MOEA evolves. The key benefit of the latter approach is that, as the MOEA progressively identifies the trade-off surface, the reduction is performed on a space more relevant to both the EA and the DM. If the reduction is to be performed iteratively then the balance between capability and complexity of the applied technique must be considered. For example, curvilinear component analysis [7] has good general applicability but a significant computational overhead, whilst principal components analysis [8] has the opposite features.

Dimension reduction methods can be applied directly to the density estimation process to preserve trade-off diversity in information-rich spaces. However, since the methods do not respect the dominance relation, they cannot be used directly in the Pareto ranking process without modification.

**Visualisation.** Note that the ability to visualise the developing trade-off surface becomes increasingly difficult as the number of criteria increases. The method of parallel coordinates is a popular countermeasure for large numbers of criteria. Scatter-plots with brushing and glyph approaches, such as Chernoff faces [9], are amongst the possible alternatives [6][10]. Parallel coordinates and scatter-plots are both closely linked to the concepts of conflict and harmony described in this paper, and are discussed further in Sect. 6.

## 4 Harmonious Criteria

### 4.1 Definitions of Harmony

A relationship in which enhancement of performance in a criterion is witnessed as another criterion is improved can be described as *harmonious*. If performance in the criterion is unaffected, the relationship is described as *weakly harmonious*. Complete definitions are provided below and can be related to the relevant  $\underline{z}^b$ -relative-to- $\underline{z}^a$  regions and lines in Fig. 3.

*Definition 3.* Levels of harmony are determined by the condition  $(\underline{z}_i^a < \underline{z}_i^b) \wedge (\underline{z}_j^a < \underline{z}_j^b)$ . If  $\nexists(a, b)$  for which the condition holds then there is *no harmony*, if  $\exists(a, b)$  then there is *harmony*, and if the condition holds  $\forall(a, b)$  then there is *total harmony*.

*Definition 4.* Levels of weak harmony are determined by the condition  $\left[ (\underline{z}_i^a < \underline{z}_i^b) \wedge (\underline{z}_j^a = \underline{z}_j^b) \right] \vee \left[ (\underline{z}_i^a = \underline{z}_i^b) \wedge (\underline{z}_j^a < \underline{z}_j^b) \right]$ . If  $\nexists(a, b)$  for which the condition holds then there is *no weak harmony*, if  $\exists(a, b)$  then there is *weak harmony*, and if the condition holds  $\forall(a, b)$  then there is *total weak harmony*.

*Definition 5.* *Neutrality* is determined by the condition  $(\underline{z}_i^a = \underline{z}_i^b) \wedge (\underline{z}_j^a = \underline{z}_j^b)$ . If  $\nexists(a, b)$  for which the condition holds then there is *no neutrality*, if  $\exists(a, b)$  then there is *neutrality*, and if the condition holds  $\forall(a, b)$  then there is *total neutrality*.

Harmonious relationships have been observed in several EMO application papers, where they are indicated by non-crossing lines between pairs of criteria on a parallel coordinates plot (see Sect. 6), including the following:

- passenger cabin acceleration versus control voltage in electromagnetic suspension controller design for a maglev vehicle [11],
- gain margin versus phase margin, and 70% rise time versus 10% settling time, in the design of a Pegasus low-pressure spool speed governor [5b].

### 4.2 Effect on EMO

In either form of total harmony, one of the criteria can be removed without affecting the partial ordering imposed by the Pareto dominance relation on the set  $Z_R$  of candi-

date solutions. This type of relationship has received some consideration in the classical OR community, usually for  $Z_R = Z^*$ , where one member of the criterion pair is known variously as *redundant*, *supportive*, or *nonessential* [12][13][14]. It remains an open question whether or not to include redundant criteria in the optimisation process. Reasons to keep such criteria include:

- knowledge of the relationship may be of interest to the DM, especially if the rate of harmonious behaviour changes over the course of the search space,
- the relationship may not be apparent from a random finite sample of the search space,
- inclusion does not, necessarily, harm the search,
- the DM may be more comfortable with the inclusion of the criterion.

Reasons to remove redundant criteria include:

- to eliminate the extra burden on the DM, who must inspect and make decisions on matters that do not affect the search and may be misleading,
- to reduce the computational load, in terms of both performance evaluations and comparisons.

The inclusion of a redundant criterion does not affect the partial ordering of candidate solutions imposed by the Pareto dominance operator. Thus, progress towards the global Pareto front is unaffected. It is, however, possible that such an inclusion could affect the diversity in the representation of the trade-off hypersurface. This depends on the definition of distance between criterion vectors used by the density estimator. For example, any procedure using Euclidean distances or the *NSGA-II* crowding algorithm [15] could suffer from potential bias. Consider the case of three criteria: where  $z_1$  and  $z_2$  totally conflict,  $z_1$  and  $z_3$  totally conflict, and  $z_2$  and  $z_3$  are in total harmony. The resulting trade-off surface is one-dimensional, and can be represented by the conflict between  $z_1$  and  $z_2$ . A uniform distribution in the Euclidean sense may not be arrived at across the normalised trade-off surface, even if such a distribution is achievable, because the Euclidean distance calculation is biased in favour of  $z_2$  since  $\{z_2, z_3\}$  has greater influence on the Euclidean distance measure than  $z_1$ . Thus, a diversity preservation technique would bias in favour of diversity in  $z_2$  on the trade-off surface. The overall effect of this depends on the trade-off surface in question: sometimes, good diversity in  $z_2$  will naturally lead to good diversity in  $z_1$  but this is not guaranteed to be the case.

### 4.3 Remedial Measures

Redundant criteria may be identified by using the sample set contained within the EA population for each criterion and looking for large positive correlations between the data sets for each pair of criteria. Redundant criteria may be removed if this is felt appropriate for the problem in-hand. Alternatively, the criteria may be selectively ignored in the density estimation process (and also the ranking process in order to reduce the number of unnecessary comparisons) and yet still be presented to the DM.

## 5 Independent Criteria

### 5.1 Independence in the context of EMO

In this paper, *independence* refers to the ability to decompose the global optimisation problem into a group of sub-problems that can be solved separately from each other. Thus, different criteria and decision variables will be allocated to different sub-problems.

In the context of the relationship between a pair of criteria, independence means that the criteria can, in theory, be optimised completely separately from each other. As with a harmonious relationship, it is possible to make improvements to both criteria simultaneously (from the perspective of the complete solution). The difference between independence and harmony is that appropriate adjustments must be made to two distinct parts of the complete solution in the former case, whilst in the latter case a single good decision modification for one of the criteria will *naturally* produce improvement in the second criterion.

If two criteria are independent then they do not form part of the same trade-off surface. Thus multiple, distinct, trade-off surfaces exist, each of which should be represented separately for inspection by the DM.

### 5.2 Effects of Independence on EMO

Consider a global problem,  $\mathbf{p}$ , comprised of  $n$  independent sub-problems  $[p_1, \dots, p_n]$  with associated independent sets of criteria  $[z_1, \dots, z_n]$  and independent sets of decision variables  $[x_1, \dots, x_n]$ . If advance knowledge of these sets is available then the global problem can be decomposed into the groups of sub-problems prior to optimisation. Then a proportion of the total available resources (candidate solution evaluations) could be exclusively allocated to the optimisation of each sub-problem. Both a global approach and the aforementioned *divide-and-conquer* method should yield the same solution of  $n$  independent trade-off surfaces. From an EMO perspective, it then becomes a matter of interest as to which technique produces superior results in terms of trade-off surface quality. Is the effort expended identifying and exploiting the correct decompositions rewarded with improved results?

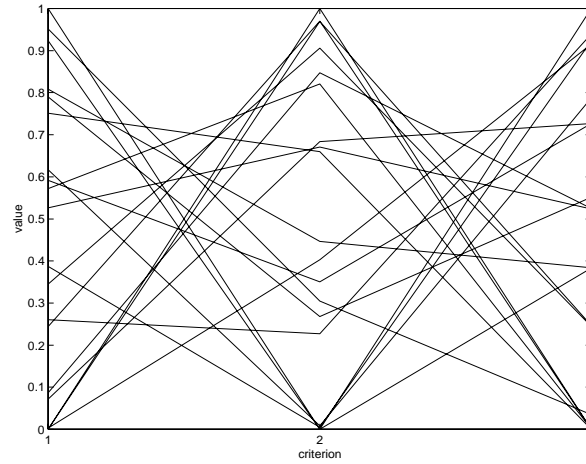
In the first study of its kind, an attempt has been made to answer this question in [16]. The study demonstrated that, for a simple test problem, a divide-and-conquer strategy could substantially improve MOEA performance. A priori decompositions were evaluated in criterion-space, decision-space, and both spaces simultaneously. Parallel EA models were applied to each sub-problem. All three methods led to significantly higher-quality trade-off surfaces than the global approach, with both-space decomposition proving the most attractive. Given that it may not be possible to accurately identify the sub-problems in advance of the optimisation, an on-line adaptive divide-and-conquer strategy for MOEAs was also proposed and evaluated in [16]. Bivariate statistical tests for independence were applied to the population sample data in order to identify the independence relationship.

## 6 Existing Methods for Identifying Pair-Wise Relationships

This paper considers the relationships that exist between pairs of criteria, by comparing pairs of criterion vectors. In this approach, composite relationships must be inferred from these simpler relations. However, the pair-wise methodology is very popular in multivariate studies and forms a good foundation for analysis, with many qualitative and quantitative techniques based on this approach. Methods that are closely linked to the definitions of conflict and harmony described earlier are discussed in the remainder of this section.

### 6.1 Qualitative Methods

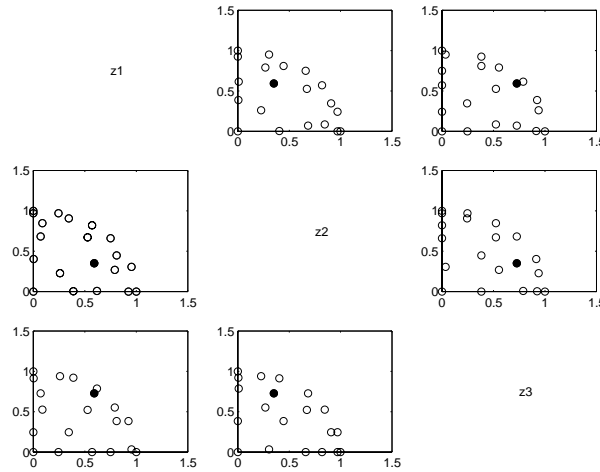
The method of parallel coordinates, first described in [17] and subsequently applied to EMO in [18], reduces an arbitrary high-dimensional space to two-dimensions. The parallel coordinates representation of the on-line archive of Fig. 2 is shown in Fig. 4. Criterion labels are located at discrete intervals along the horizontal axis (and these should be interchangeable). Normalised performance in each criterion is indicated on the vertical axis. A particular criterion vector is displayed by joining the performance levels in all adjacent criteria by straight lines. Then, considering two criterion vector instances for a pair of criteria, the lines representing the two instances will cross if conflict is exhibited according to Definition 2 or will fail to cross if harmony is observed according to Definitions 3 or 4 (in the case of Definition 5, the lines will be superimposed). Thus, the magnitude of conflict is heuristically visualised as ‘many’ crossing lines.



**Fig. 4.** Parallel coordinates representation of the data shown in Fig. 2

Wegman [19] presents some valuable insights and extensions towards using the parallel coordinates representation as a high-dimensional data analysis tool. Statistical interpretations of the plots are possible, with features such as marginal densities, correlations, clusters, and modes proving readily identifiable. Parallel coordinates plots can suffer from over-plotting for large data sets and thus a density plot variant is also presented in the paper to overcome this.

Another popular method of pair-wise visualisation, which in its full form presents more simultaneous comparisons than the standard parallel coordinates plot, is the scatterplot matrix [10]. Such a plot for the MOGA on-line archive of Fig. 2 is shown in Fig. 5. Each element of the matrix of plots shows a particular bi-criterion section of the trade-off surface. For example, the upper central plot shows  $z_2$  on the horizontal axis and  $z_1$  on the vertical axis. It can sometimes be difficult to extract information from these plots, especially as the number of criteria increases. Highlighting of a particular criterion vector instance or group of instances – a technique known as *brushing* – can often aid higher-order understanding. The filled circle in Fig. 5 indicates one particular vector.



**Fig. 5.** Scatterplot matrix representation of the data shown in Fig. 2

## 6.2 Quantitative Methods

Several pair-wise methods exist for quantifying conflict between criteria that use similar concepts to the parallel coordinates notion of crossing lines. The *Kendall sample correlation statistic* measures the difference between the number of *concordant* samples (as one variable increases/decreases, the other follows suit) and the number

of *discordant* samples (as one variable increases/decreases, the other does the opposite) [20]. Thus, discordance produces crossing lines whilst concordance does not. Fuzzy measures of conflict also use this type of approach: see, for example, [21].

Schroder [22] developed a technique based directly on the method of parallel coordinates. In this approach, each criterion range is partitioned into a number of equally sized regions. The level of conflict is then defined as a weighted-sum of the crossings between pairs of regions (rather than actual solutions), where the weights are based on the separation between the regions. Crossings between distant regions are argued to be indicative of strong conflict. The method also normalises conflict levels with respect to population density. This may be appropriate if one particular region is thought to be over-sampled by the search, but may not in general be correct because if vectors are similar in two criteria this does not necessarily mean they are from the same region of the global trade-off surface. The method also requires additional preference information, unlike the previous techniques, since it is not based purely on ordinal data. However, more information can, potentially, be extracted using this method.

## 7 Conclusions

EMO applications have long considered the simultaneous optimisation of large numbers of criteria. However, EMO algorithm developers have tended to concentrate almost entirely on the bi-criterion domain. Thus, there is a present lack of understanding of how MOEAs cope with larger numbers of criteria. This paper has sought to lay foundations for future work in this direction by considering how increasing numbers of criteria affect MOEA search performance.

Three relationships – conflict, harmony, and independence – have been identified. It has been demonstrated how the relationship between two criteria can contain elements of both conflict and harmony, resulting from interaction with other criteria.

It has been argued that increasing numbers of conflicting criteria will severely hamper the ability of an MOEA to represent the global trade-off surface. Thus, in general, the oft-stated EMO aims of closeness to and diversity across the entire Pareto front could be little more than a pipedream. It is difficult to see how, when confronted with large numbers of criteria, increased requirements for preference information can be avoided. Even if an MOEA was capable of adequately representing the entire, complicated, high-dimensional trade-off surface, this amount of information is surely of little benefit to the human DM, who is faced with the task of selecting a single solution. The dimension of the problem must be kept reasonably low for human DM analysis to remain tractable: twelve criteria has been suggested as an upper limit [23]. Aggregation of criteria, whilst anathematic to many EMO researchers, may thus be necessary.

Even if the number of criteria is limited by DM considerations, it may still be prudent to consider further dimensionality reduction in the context of the MOEA search. This can be achieved by removing some criteria from certain comparisons (and thus preserving the dominance relation) or by applying some form of transformation to a new set of coordinates (the standard dimension reduction approach). The utility of methods based on the latter approach is limited because they do not respect the domi-

nance relation. They may, however, be used to help achieve good diversity in information-rich spaces (where ‘information’ is defined according to the chosen method).

Interactive preference articulation schemes, such as [5a], are particularly valuable in problems with large numbers of conflicting criteria and fit very nicely within the iterative EA framework. In such schemes, the attention of the optimiser is focused on various sub-regions of the trade-off surface as the search progresses. This is beneficial to the DM who may only be interested in learning about certain trade-offs within the global problem. The progressive nature of the scheme suits the often changing aspirations of the DM as more knowledge is uncovered. The main drawback of this approach is that it can be rather DM-intensive.

To summarise, increasing numbers of conflicting criteria in a problem transforms the aim of EMO from identification of a globally optimal solution set towards assisting the DM in learning about the trade-offs between criteria and finding an *acceptable* solution.

Harmonious criteria, and the special case of redundancy, do not have the same severe impact on EMO as does conflict. Convergence to the Pareto front is unaffected by increasing numbers of totally harmonious criteria. Issues surrounding distribution of solutions across the surface do require some care however. The decision on whether to eliminate any identified redundant criteria from the search is perhaps best left to the discretion of the analyst and the DM.

The existence of independence within the global problem leads to multiple, separate, trade-off surfaces. If independence can be identified then the deployment of a divide-and-conquer strategy could potentially improve EMO performance [16].

The recently proposed suite of scalable test problems provides the opportunity to explore in full the behaviour of the MOEAs in a controlled and tractable manner [2]. Advancements in techniques for performance analysis and visualisation are required to aid understanding of the results from such problems. Assessment of the utility of contemporary dimension reduction methods for EMO is a rich area for future research. Alternative methods to the standard sample-based approach of describing the trade-offs may prove useful both as a decision aid and within the context of the search itself.

In conclusion, the simultaneous consideration of many criteria is arguably the greatest challenge facing the EMO community at the present time.

**Acknowledgments.** The authors would like to thank the anonymous reviewers for their valuable comments and advice given on the original submissions of this paper and [16]. Robin Purshouse also wishes to acknowledge the UK Engineering and Physical Sciences Research Council (EPSRC) for research studentship support.

## References

1. Coello, C.A.C., Van Veldhuizen, D.A., Lamont, G.B.: Evolutionary Algorithms for Solving Multi-Objective Problems. Kluwer Academic Publishers, New York Boston Dordrecht London Moscow (2002)

2. 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
3. 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
4. Deb, K.: Multi-Objective Optimization using Evolutionary Algorithms. Wiley, Chichester New York Weinheim Brisbane Singapore Toronto (2001)
5. Fonseca, C.M., Fleming, P.J.: Multiobjective optimization and multiple constraint handling with evolutionary algorithms (a) Part I: A unified formulation (b) Part II: Application example. IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans **28** (1998) 26-37 and 38-47
6. Scott, D.W.: Multivariate Density Estimation: Theory, Practice, and Visualization. Wiley, New York Chichester Brisbane Toronto Singapore (1992)
7. Demartines, P., Hérault, J.: Curvilinear Component Analysis: a Self-Organising Neural Network for Non-linear Mapping of Data Sets. IEEE Transactions on Neural Networks **8** (1997) 148-154
8. Jolliffe, I.T.: Principal Component Analysis. Springer, New York (1986)
9. Chernoff, H.: The Use of Faces to Represent Points in  $k$ -Dimensional Space Graphically. Journal of the American Statistical Association **68** (1973) 361-368
10. Cleveland, W.S.: Visualizing Data. Hobart Press, Summit (1993)
11. Dakev, N.V., Whidborne, J.F., Chipperfield, A.J., Fleming P.J.: Evolutionary H-infinity design of an electromagnetic suspension control system for a maglev vehicle. Proceedings of the Institution of Mechanical Engineers Part I **211** (1997) 345-355
12. Agrell, P.J.: On redundancy in multi criteria decision-making. European Journal of Operational Research **98** (1997) 571-586
13. Carlsson, C., Fullér, R.: Multiple Criteria Decision Making: The Case for Interdependence. Computers & Operations Research **22** (1995) 251-260
14. Gal, T., Hanne, T.: Consequences of dropping nonessential objectives for the application of MCDM methods. European Journal of Operational Research **119** (1999) 373-378
15. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II. IEEE Transactions on Evolutionary Computation **6** (2002) 182-197
16. Purshouse, R.C., Fleming, P.J.: An Adaptive Divide-and-Conquer Methodology for Evolutionary Multi-Criterion Optimisation. This volume. (2003)
17. Inselberg, A.: The Plane with Parallel Coordinates. The Visual Computer **1** (1985) 69-91
18. 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
19. Wegman, E.J.: Hyperdimensional Data Analysis Using Parallel Coordinates. Journal of the American Statistical Association **85** (1990) 664-675
20. Kendall, M.G.: A new measure of rank correlation. Biometrika **30** (1938) 81-93
21. Lee, J., Kuo, J.-Y.: Fuzzy decision making through trade-off analysis between criteria. Journal of Information Sciences **107** (1998) 107-126
22. Schroder, P.: Multivariable control of active magnetic bearings. Doctoral dissertation, University of Sheffield. (1998)
23. Bouyssou, D.: Building criteria: a prerequisite for MCDA. In: Bana e Costa, C.A. (ed.): Readings in Multiple Criteria Decision Aid. Springer-Verlag, Berlin Heidelberg New York (1990) 58-80