A Non-Parametric Harmony-Based Objective Reduction Method for Many-Objective Optimization

Alan R. R. Freitas
Graduate Program in Electrical Engineering,
Federal University of Minas Gerais,
Av. Antônio Carlos 6627,
31270-901, Belo Horizonte, MG, Brazil
alandefreitas@gmail.com

Peter J. Fleming
Department of Automatic Control
and Systems Engineering,
University of Sheffield
Sheffield, S1 3JD, U.K.
p.fleming@sheffield.ac.uk

Frederico G. Guimarães
Department of Electrical Engineering,
Federal University of Minas Gerais,
Av. Antônio Carlos 6627,
31270-901, Belo Horizonte, MG, Brazil
fredericoguimaraes@ufmg.br

***Abstract*—Multiobjective optimization has been applied successfully to real-world optimization problems with few objectives. However, the performance of current algorithms for multiobjective optimization reduces exponentially as the number of objectives grows. In this paper we present a non-parametric harmony-based approach for objective reduction in order to deal with this issue in many-objective optimization problems. The proposed approach has many advantages such as the independence of the relationship between the objectives as long as they are harmonious and the possibility to visualize conflict and trade-off in the reductions performed.**

***Index Terms*—Objective Reduction, Multiobjective Optimization, Many-objective Optimization, Non-Parametric Objective Reduction**

I. Introduction

In this work we present an approach for objective reduction in multiobjective optimization. Multiobjective optimization treats objectives as non-comparable and a number of solutions can be found to represent the trade-off between those objectives. Multiobjective evolutionary algorithms have been successful for real-world problems but their scalability for many-objective problems is poor and their performance reduces significantly. In this context, we present a strategy for reducing the number of objectives in such problems according to a measure of their harmony.

We present forms of measuring conflict and harmony, the latter being used to group iteratively the most harmonious objectives into a new compound objective. In contrast with other simple measures of conflict, the harmony measure presented is non-parametric and therefore can be robust no matter the relationship between the objectives is.

We then present the result of the application of the algorithm that includes the main contributions of this paper:

- A method that enables the simple visualization of the relationships between the objectives with an iterative one-by-one reduction scheme and calculation of conflict;
- Independent of any linearity relationship between the objectives;
- Can work with any number of objectives, subsequently to be reduced defined by the decision maker;
- Easy to implement.

We then conclude our paper with discussions about the method and suggestions concerning possible future work.

II. Objective Reduction in Multiobjective Optimization

In engineering applications and systems design we often face multiple and conflicting objectives and quality criteria. Usually, these are combined into a single objective function that reflects the utility of each objective, while some objectives might be turned into constraints. In multiobjective optimization we treat the objectives separately as non-comparable objectives, which are assumed to be conflicting, and a number of solutions can be found to represent the trade-off between those objectives [1], [2], [3].

The solutions that represent the trade-off form the Pareto-optimal set, solutions that lead to an optimal set of non-comparable points in the objective space. That means that there is no solution in the feasible set that is better than (dominates) a solution in the Pareto set.

First-generation multiobjective evolutionary algorithms (MOEAs) for solving multiobjective optimization problems were developed in the 1990s, see [2], and incorporated Pareto-based ranking and comparison for computing fitness values. Since then, MOEAss have been widely recognized as suitable and successful methods in this context. One commonly reported advantage of using evolutionary algorithms for solving multiobjective optimization problems is the fact that they work with an evolving population of candidate solutions. This set-based evolution offers the possibility of searching for a set of estimates in a single run.

The second generation of MOEAs are more efficient from the computational point of view. Moreover, these algorithms have introduced an explicit elitism by means of an external archive population, for storing the best estimates of the Pareto set. NSGA-II (Non-Dominated Sorting Genetic Algorithm - II) [4], SPEA2 (Strength Pareto Evolutionary Algorithm 2) [5], and PAES (Pareto Archived Evolution Strategy) [6] are good examples of this generation.

Despite the success of those algorithms for addressing real-world problems, researchers recently began to investigate

978-1-4799-0652-9/13/$31.00 ©2013 IEEE

the scalable behavior of these MOEAs and when problems have more than three objectives, it was observed that the performance of those algorithms reduces dramatically. This fact led to the definition of a new term to better refer to this context, namely, many-pbjective problems [7]. In many-objective problems, Pareto-based ranking ceases to be an effective discriminator and the majority of solutions are non-comparable, thus compromising the convergence of the algorithms. The main issues with many-objective optimization are:

- Loss of selective pressure because most of the population becomes non dominated.
- Dimensionality and computational cost - increase in the population size and the number of points required to represent the Pareto front.
- Visualization - difficult to visualize trade-offs in many dimensions.
- Issues on the scalability of decision-making techniques.

Schemes for solving tmany-objective problems include different dominance relations to improve discrimination at each generation [8], other fitness measures involving indicators [9], and simultaneously optimizing many scalarized objective functions [10]. Another technique for solving these problems is dimensionality reduction through identification of redundant objectives [11], [12], [13].

Brockhoff and Zitzler [11], [12] propose a method of objective reduction geared towards direct integration into the evolutionary search. Their algorithm depends on a change in the dominance structure. They also present the problem of finding a minimum objective subset, maintaining the given dominance structure with a given error. A greedy algorithm removes objectives if the dominance relations do not change with the removal of the objective because it is considered a non-conflicting one.

Saxena et al. [13] employ Principal Component Analysis (PCA) of the objective values in order to identify related objectives. The PCA is adapted to handle non-linear relations between the objectives and the technique is then embedded into an NSGA-II. Other papers have also presented new techniques based on maximum variance unfolding and PCA for nonlinear reduction of objectives [14]. Those techniques remove dependencies in the non-dominated solutions.

Jaimes et al. [15] employ unsupervised feature selection to reduce the number of objectives according to their correlation. Distant objectives are treated as conflicting ones. Similarly to [12], the algorithm also seeks the minimum objective subset with minimum error.

Singh et al. [16] present a Pareto Corner Search. This algorithm focuses only on corners among the non-dominated solutions. After that, the number of resulting non-dominated solutions from a reduction of the achieved solutions indicates if the objectives are important or if they can be reduced.

III. Harmony-Based Objective Reduction

A. *Conflict*

In simple terms, there is conflict [17], [18] between two objectives when good values for one of them imply bad values

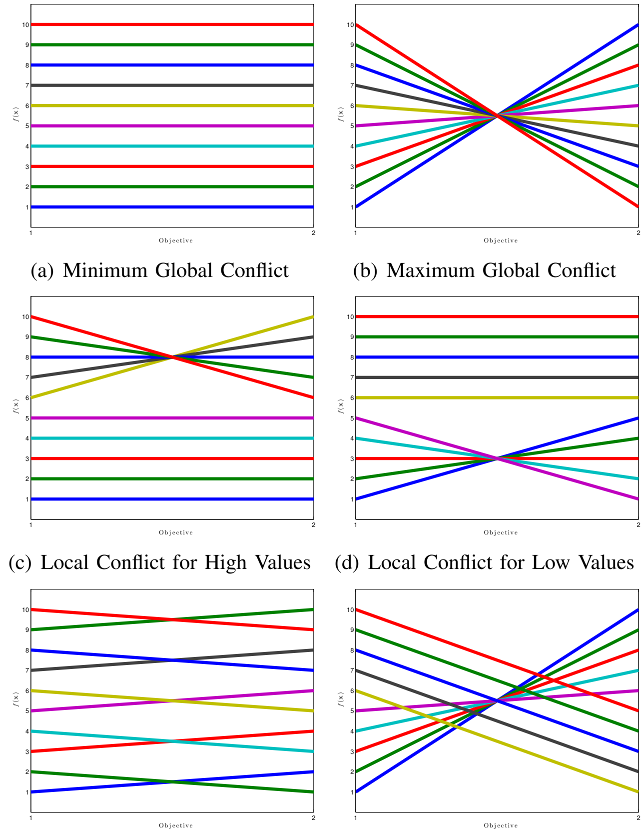

(a) Minimum Global Conflict (b) Maximum Global Conflict (c) Local Conflict for High Values (d) Local Conflict for Low Values (e) Global Medium Conflict (f) Local Conflict in Different Locations for Each Objective

Fig. 1. Example of many possible kinds of conflict

for another. This may sound simple but it is a concept that depends on the set of possible solutions, the shape of the Pareto front, and, indeed, the expectations of the decision maker. Figure 1 represents many kinds of conflict in parallel coordinate graphs of two objectives. Each line represents a solution while the x-axis represents each of the objectives and the y-axis represents the objective function values. We can see from the graphs how not only the amount of conflict varies but the conflict can also be concentrated at certain values.

B. *Harmony*

Harmony [17], [18], [12], [19], on the other hand, is a concept that applies when an improvement in one objective would lead to an improvement in another objective. Harmonious objectives are indicated on parallel coordinates by non-crossing lines. Thus, it is related to the possibility of joining the objectives through summation without loss of quality in the Pareto front. If we are to group two objectives in a new compound objective, it is best to group those objectives with greater harmony even if there is some level of conflict between them. Figure 1(a) shows an example of complete harmony.

It is commonly understood that if there is harmony between objectives there is no conflict. However, harmony and conflict are not perfect antonyms and one example would be when conflict happens because the solutions are in the range from

9 to 10 for objective f_1 and the range 1 to 2 for f_2. Even though there is conflict between the objectives if the expected range of values is from 0 to 10, it may make sense to group those objectives if there is harmony between them. Harmony between those objectives means that improvement in one objective would still lead to improvement in the other.

In this paper, we propose a simple non-parametric objective reduction technique that considers the fact that the more harmonious two objectives are, the more they can be reduced or grouped into a single new scalarized objective value.

IV. Conflict and Harmony Measures

A. Direct Conflict

The first measure of conflict is the difference of gain or loss in absolute terms for absolute objectives. This measure may be particularly interesting if both objectives use the same units. This conflict measure c is the absolute difference between the values for the objectives. If the range of values is very different for each objective, the objective values of $\mathbf{X}$ may be normalized to zero in a new $\mathbf{X}'$ formed by the subtraction of the minimum values in $\mathbf{X}$ for each objective.

Thus, if $\mathbf{X}_{ij}$ is the value for objective j in the solution i, the direct conflict C_{ab} between objectives a and b is:

$$\begin{aligned} C_{ab} &= \sum_i |\mathbf{X'}_{ia} - \mathbf{X'}_{ib}| \\ \mathbf{X'}_{ij} &= \mathbf{X}_{ij} - min(\mathbf{X}_{.j}) \end{aligned} \tag{1}$$

With this normalization, the direct conflict measure is:

- Insensitive to summation or subtractions in the original objective values
- Useful when the objectives use the same units
- The summation of the objectives could equally solve the problem because all objectives are equally important
- Decision maker wants to understand the trade-offs involved

The range of possible values of conflict goes from $c_{\min} = 0$ to a problem dependent $c_{\max}$ because absolute values are directly being compared. The value $c_{\max}$ is then related to the original range of values in the objectives.

The direct conflict measure answers the following question:

- Given that the gains in all objectives in absolute terms (and not relative in terms of known ranges) are equally important, what is the conflict existent between those objectives?

As multiobjective optimization is meant to treat incomparable objectives in the general case, it seems that is does not make sense to directly compare the objective values. However, this can be useful when the objectives have the same units.

For example, there may be decisions in a company that can be monetized. The objectives probably could be grouped by summation as the company is willing to make money in both objectives. However, the decision maker may want to analyze the results from a multiobjective perspective so that he or she can understand the trade-off involved in the decision.

B. Maxmin Conflict

This second measure of conflict also assumes comparability of the objectives. It is useful when we have any reference to imply that all the objectives are equally important. The comparability and the importance of each objective is proportional to their known range of values. That implies that the objectives are not completely incomparable but equally important in regard to their range of values. The mathematical formulation of this maxmin conflict is:

$$\begin{aligned} C_{ab} &= \sum_i |\mathbf{X'}_{ia} - \mathbf{X'}_{ib}| \\ \mathbf{X'}_{ij} &= \frac{\mathbf{X}_{ij} - min(\mathbf{X}_{.j})}{max(\mathbf{X}_{.j}) - min(\mathbf{X}_{.j})} \end{aligned} \tag{2}$$

This normalization leaves us with values from 0 to 1 for all objective values and the maxmin

- Insensitive to any previous linear normalization
- Useful when all the objectives are equally important
- The importance of each objective is linearly proportional to its known range of values in the preference area

The maxmin conflict answers the following question:

- Given that all the objectives are equally important in proportion to their range of values, what is the conflict existent between those objectives?

With this measure of conflict, we have more loss of information because of the lack of direct comparability in the same units. On the other hand, it is a more robust measure because there are fewer assumptions, such as the very possibility of having different units of measurement.

With a small adaptation, the maxmin normalization can also include the preferences of the decision maker. Given that $\mathbf{P}_j$ is the decision maker's goal for each objective j that defines the preference area, a simple adaptation of this conflict measure for minimization problems would be:

$$\begin{aligned} C_{ab} &= \sum_i |\mathbf{X'}_{ia} - \mathbf{X'}_{ib}| \\ \mathbf{X'}_{ij} &= \frac{\mathbf{X}_{ij} - min(\mathbf{X}_{.j})}{\mathbf{P}_j - min(\mathbf{X}_{.j})} \end{aligned} \tag{3}$$

In this case, the importance of the objectives defined by the range of known values also considers the expectations of the decision maker.

C. Logarithmic or Exponential Conflict

There may be cases in which it does not make sense to look at the non-normalized results and the direct interpretation of results may cognitively deceive us. For instance, on the Pareto optimal sensitivity bounds for a strongly interacting 2-loop brazing temperature control system [20], we have 3 sets of solutions for 3 different systems in Figure 2(a). From those results we can see that there is almost no conflict at all for the first set and there is more conflict for the third set.

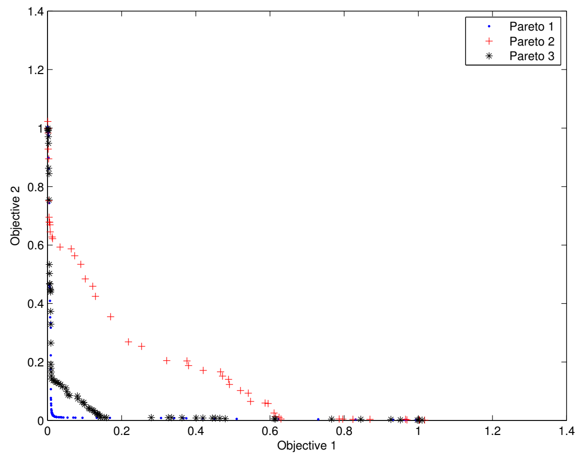

(a) Pareto Fronts

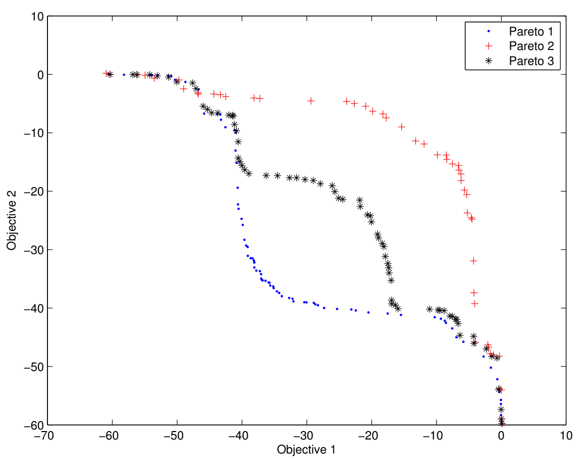

(b) Pareto Fronts in Log Scale

Fig. 2. Pareto Optimal Sensitivity Bounds for a Strongly Interacting 2-Loop Brazing Temperature Control System

However, once those results are obtained, it does not usually make sense to interpret the values directly. For better visualization, a logarithmic scale is used. By normalizing all results $\mathbf{X}_ij$ to $\mathbf{X}_ij = 20\log_{10}\mathbf{X}_ij$, we obtain the representation in Figure 2(b). In this second representation of the solutions, we could interpret that there is conflict in all the fronts. However, the perception of conflict only comes from the cognitive representation of the same solutions. This illustrates how the perception of conflict may depend on preferences of the decision maker.

For this type of problem in which more advanced normalization techniques are required, only the interference of the decision maker can lead to an appropriate measure of conflict because when a maxmin or direct conflict measure is applied afterwards, that would involve comparability of the objectives.

Also, it would not be appropriate to further extend towards more general normalization conflict measures in the direction of logarithmic or exponential measures in the context of optimization, as we have done with summation and multiplication in the last subsections. The reason for this is that, in general, those sorts of transformations could make initially increasing values decrease after the normalization.

Thus, further normalizations could change the original harmony relationship of the objectives. We are going to call these *disruptive normalizations*. However, they illustrate the necessity for a more robust measure that would recognize conflict in all fronts in Figure 2, independently of deceptive cognitive interpretations of the results.

D. Non-parametric Rank Conflict

In the example of Figure 2, different normalization techniques would also have different results but a non-parametric measure of conflict could be insensitive to the perceptions of the decision maker. This robustness, however, comes at a cost: that of losing information in a ranking process, something that might need to be compensated by increasing the granularity of the solution set.

A non-parametric measure of conflict works without the assumption of comparability between the objectives. It is useful when we do not have any reference to imply a relation of importance between the objectives and we have the goal of understanding the relationship between non-comparable objectives without wrongly implying or stating that they are equally important. Rank differences in the objectives are used to compare results without considering the distance between those values. Thus, the mathematical formulation of the non-parametric rank conflict is:

$$
\begin{aligned}
C_{ab} &= \sum_i |\mathbf{X'}_{ia} - \mathbf{X'}_{ib}| \\
\mathbf{X'}_{ij} &= \mathbf{R}_{ij} \\
\mathbf{R}_{ij} &= \text{rank of } \mathbf{X}_{ij} \text{ within } \mathbf{X}_{.j}
\end{aligned} \tag{4}
$$

If n is the number of solutions being analyzed, this normalization leaves us with values from 0 to n for each objective value and the non-parametric rank conflict measure is:

- Insensitive to any previous non-disruptive normalization
- Useful when the objectives use different units and are not comparable
- Useful when a value of importance of each objective cannot be inferred but we want to understand the relationship between them

The range of possible values of conflict between two objectives goes from $c_{\min} = 0$ to:

$$
\begin{aligned}
c_{\max} &= \sum_{i=1}^{n} |2i - n - 1| \\
c_{\max} &= 2\left(\lceil n/2\rceil(n+1) - \frac{2(1+\lceil n/2\rceil)\lceil n/2\rceil}{2}\right)
\end{aligned} \tag{5}
$$

The non-parametric rank conflict measure answers the following question:

- Independently of any previous non-disruptive normalization, would an increase in one objective imply in a loss in another objective inside the preference area?

According to the definition given in Section IV and the question related to this conflict measure, we can see that the non-parametric rank conflict measure can also be used as a measure of harmony, which is the corresponding relation of increase or decrease altogether for the objectives, independently of the shape of the Pareto front.

The non-parametric rank conflict measure is more robust because, in general, it is more insensitive to any previous normalization. As in any non-parametric measure, this involves loss of information at the cost of fewer assumptions. The loss of information comes from the lack of direct comparability between the objectives, which seems appropriate for the general case of multiobjective optimization. On the other hand, it is a more robust measure because there are fewer assumptions, such as the very possibility of having different units of measure without any conversion between them.

Together with the direct and maxmin measures of conflict, this is one of the three relevant and general kinds of normalization considered in this work. However, only the rank normalization is robust enough to also measure harmony and to work with any sort of data, in general.

V. Objective Reduction Algorithm

As mentioned in section III, the more harmonious two objectives are, the more they can be reduced or grouped into a single new scalarized objective value without deforming the Pareto-shape. For this reason, for each group of objectives we find the two most harmonious objectives and group them as a new scalarized objective. Thus, the decision maker or the user can make a decision on the number of objectives they want to consider with the following algorithm:

1) For each objective from m objectives
    a) Rank the values in that objective with cost $O(n \log n)$ or $O(n)$ in relation to the n points
2) While the user still wants to the reduce one more objective for k reduced objectives
    a) Calculate C_{ij} for each pair of objectives ij with cost $O(nm^2)$
    b) Find the smallest value of C_{ij} with cost $O(m^2)$
    c) Reduce the objectives with minimum non parametric conflict (or maximum harmony) with cost $O(n)$

Ranking the values of one objective has the same cost as sorting the elements, that is, $O(n \log n)$ for comparison-based algorithms and $O(n)$ for other algorithms such as Radix sort. Thus, the first step of the algorithm costs $O(mn \log n)$ or $O(mn)$. Each iteration in the second step costs $O(\max(nm^2, m^2, n)) = O(nm^2)$. As we have k reductions, the cost of the second step is $O(knm^2)$. Thus, the complexity of the algorithm is either $O(max(mn \log n, knm^2))$ for comparison-based sorting algorithms or $O(max(mn, knm^2)) = O(knm^2)$ for other algorithms. For practical purposes, if we consider that k and n are constants, the cost of the algorithm becomes $O(m^2)$ for any case. If we want to reduce all the objectives to a single objective, then $k = m$ and the algorithm is $O(nm^3)$ or $O(\max(mn \log n, nm^3))$ for comparison-based algorithms. For $k = 1$, with one reduction at a time, the tool can be easily and efficiently embedded into a MOEA with cost $O(nm^2)$.

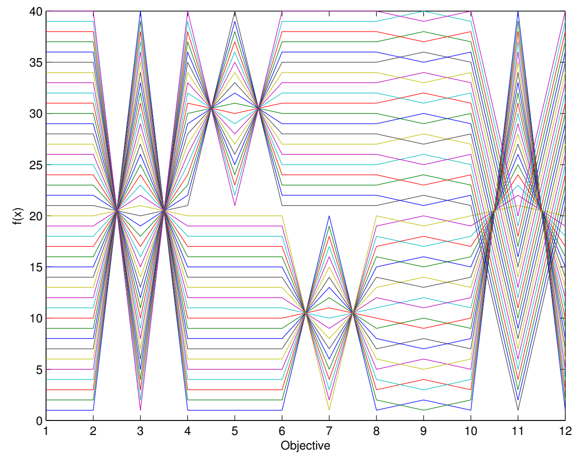

Fig. 3. A Set of Solutions

As an example, we consider the set of solutions represented in the Parallel Coordinates of Figure 3. Every pair of objectives in this set of solutions contains all the 6 kinds of conflict already presented in Figure 1 with a granularity of 40 points. Objectives 1 and 2 are in complete harmony. Objectives 2, 4, 6, 8, 10, and 12 are also in complete harmony because for each pair of objectives representing a sort of conflict, the values return the ones in objective 1 before going to the next adjacent type of conflict.

In this example, our algorithm suggests the reduction of all objectives f_1, f_2, f_4, f_6, f_8, f_{10}, f_{12}. All those objectives are in complete harmony because they have the same objective values. By normalizing all these objectives into ranks again and summing their values into a new objective f_a we now have 6 objectives left. By applying the algorithm again, f_a is grouped with objective f_9 with 5% of non-parametric conflict.

By establishing that $f_b = f_9 + f_a$, from the 5 objectives left, our reduction algorithm will then group $f_c = f_{11} + f_3$ with 25% of conflict, $f_d = f_b + f_c$ with 0% of conflict, resulting in a three-objective problem. If we reduce one more objective, we would have $f_e = f_d + f_5$ with 26% of conflict and finally $f_f = f_e + f_7$, which represents a scalarization of the whole problem into a single objective problem.

VI. Discussion and Future Work

One advantage of this method is that it is easy to visualize the relationship between the objectives because at each step we can identify which pair of objectives is most harmonious and the amount of conflict between them. The strategy of grouping objectives one by one also makes it easy to interpret the results as to which would be the best choices of objective reduction given any desired number of final objectives. The measure of conflict also informs the decision maker of the size of the penalty involved in the reduction of those objectives.

Perhaps the most important feature of this method is that it is not restricted to linear relations between the objectives and it is not even dependent on any specific non-linear relation between the objectives because only their harmony is used as a criterion for deciding which objectives to reduce. This is a convenient feature because conflict may lead the decision maker to think there is no conflict where there might be. The characteristic that one objective can be reduced at a time makes it convenient for the decision maker to choose when to stop grouping objectives. The measures of conflict can make this decision easier. The proposed approach is also very simple to implement and has very low computational cost. The computational cost is low enough to allow it to be embedded into a MOEA with an interface for the decision maker to decide how many objectives should be reduced interactively.

When the algorithm is not to be embedded into a MOEA, an important question is whether we are to use (i) only non-dominated solutions to analyze the relationship between the objectives, (ii) any set of solutions, or (iii) even solutions limited to a preference area. There is no best answer to this question but we know that using non-dominated solutions would only show that the objectives can be optimized in a way that we improve both solutions while using any set of solutions could show that there is inherent harmony between the objectives. Either way, as the number of objectives grows, this difference tends to disappear because most solutions will be non-dominated.

Thus, as most solutions become non-dominated in many-objective optimization, explaining the relationships between those objectives becomes more important than optimizing them at this stage. It might be interesting also to use extreme solutions in each of the objectives to have a significant representation of the potential for each objective. When the objectives are paired one by one, these might be "greedy" choices that do not lead to an optimal reduction when we consider the objectives that have already been paired. In this sense, it might be that the order of reduction to this objective reduction problem is only a local optimum in relation to the number of objectives to be reduced. One last issue to be considered is the visualization of those reductions and how the final results will be presented to the decision maker as there is still a need to show absolute values and the reduced objectives.

Acknowledgements

This work has been supported by the Brazilian agencies CAPES, CNPq, and FAPEMIG; and the Marie Curie International Research Staff Exchange Scheme Fellowship within the 7th European Community Framework Programme.

References

[1] K. Deb, Multi-Objective Optimization Using Evolutionary Algorithms, Wiley-interscience series in systems and optimization, Wiley, 2009. URL http://books.google.com.br/books?id=U0dnPwAACAAJ

[2] C. Coello, G. Lamont, D. Veldhuizen, Evolutionary Algorithms for Solving Multi-Objective Problems, Genetic and evolutionary computation series, Springer, 2007. URL http://books.google.com.br/books?id=rXIuAMw3lGAC

[3] C. Fonseca, P. Fleming, et al., Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization, in: Proceedings of the fifth international conference on genetic algorithms, Vol. 1, San Mateo, California, 1993, p. 416.

[4] K. Deb, S. Agrawal, A. Pratap, T. Meyarivan, A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: Nsga-ii, Lecture notes in computer science 1917 (2000) 849–858.

[5] E. Zitzler, M. Laumanns, L. Thiele, E. Zitzler, E. Zitzler, L. Thiele, L. Thiele, Spea2: Improving the strength pareto evolutionary algorithm (2001).

[6] J. Knowles, D. Corne, Approximating the nondominated front using the pareto archived evolution strategy, Evolutionary computation 8 (2) (2000) 149–172.

[7] H. Ishibuchi, N. Tsukamoto, Y. Nojima, Evolutionary many-objective optimization: A short review, in: Evolutionary Computation, 2008. CEC 2008.(IEEE World Congress on Computational Intelligence). IEEE Congress on, IEEE, 2008, pp. 2419–2426.

[8] L. Batista, F. Campelo, F. Guimarães, J. Ramírez, Pareto cone ε-dominance: improving convergence and diversity in multiobjective evolutionary algorithms, in: Evolutionary Multi-Criterion Optimization, Springer, 2011, pp. 76–90.

[9] N. Beume, B. Naujoks, T. Wagner, Pareto-, aggregation-, and indicator-based methods in many-objective optimization.

[10] H. Ishibuchi, T. Doi, Y. Nojima, Incorporation of scalarizing fitness functions into evolutionary multiobjective optimization algorithms, Parallel Problem Solving from Nature-PPSN IX (2006) 493–502.

[11] D. Brockhoff, E. Zitzler, Are all objectives necessary? on dimensionality reduction in evolutionary multiobjective optimization, Parallel Problem Solving from Nature-PPSN IX (2006) 533–542.

[12] D. Brockhoff, E. Zitzler, Dimensionality reduction in multiobjective optimization: The minimum objective subset problem, Operations Research Proceedings 2006 (2007) 423–429.

[13] D. Saxena, K. Deb, Non-linear dimensionality reduction procedures for certain large-dimensional multi-objective optimization problems: Employing correntropy and a novel maximum variance unfolding, in: Evolutionary Multi-Criterion Optimization, Springer, 2007, pp. 772–787.

[14] D. Saxena, J. Duro, A. Tiwari, K. Deb, Q. Zhang, Objective reduction in many-objective optimization: Linear and nonlinear algorithms, Evolutionary Computation, IEEE Transactions on 17 (1) (2013) 77–99. doi:10.1109/TEVC.2012.2185847.

[15] A. López Jaimes, C. A. Coello Coello, D. Chakraborty, Objective reduction using a feature selection technique, in: Proceedings of the 10th annual conference on Genetic and evolutionary computation, GECCO '08, ACM, New York, NY, USA, 2008, pp. 673–680. doi:10.1145/1389095.1389228. URL http://doi.acm.org/10.1145/1389095.1389228

[16] H. Singh, A. Isaacs, T. Ray, A pareto corner search evolutionary algorithm and dimensionality reduction in many-objective optimization problems, Evolutionary Computation, IEEE Transactions on 15 (4) (2011) 539–556. doi:10.1109/TEVC.2010.2093579.

[17] R. Purshouse, P. Fleming, Conflict, harmony, and independence: Relationships in evolutionary multi-criterion optimisation, in: C. Fonseca, P. Fleming, E. Zitzler, L. Thiele, K. Deb (Eds.), Evolutionary Multi-Criterion Optimization, Vol. 2632 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2003, pp. 16–30. doi:10.1007/3-540-36970-8_2. URL http://dx.doi.org/10.1007/3-540-36970-8_2

[18] R. Purshouse, P. Fleming, On the evolutionary optimization of many conflicting objectives, Evolutionary Computation, IEEE Transactions on 11 (6) (2007) 770–784.

[19] I. Giagkiozis, R. Purshouse, P. Fleming, Generalized decomposition and cross entropy methods for many-objective optimization.

[20] E. Eitelberg, On multiloop interaction and relative and bristol gains, Journal of Dynamic Systems, Measurement, and Control 128 (4) (2006) 929–937. doi:10.1115/1.2363412. URL http://link.aip.org/link/?JDS/128/929/1