On using decision maker preferences with ParEGO

Hakanen, Jussi; Knowles, Joshua

DOI:

[10.1007/978-3-319-54157-0_20](https://doi.org/10.1007/978-3-319-54157-0_20)

License:

None: All rights reserved

Document Version

Peer reviewed version

Citation for published version (Harvard):

Hakanen, J & Knowles, J 2017, On using decision maker preferences with ParEGO. in *Evolutionary Multi-Criterion Optimization: 9th International Conference, EMO 2017, Munster, Germany, March 19 - March 22, 2017, Proceedings*. vol. 10173, Lecture Notes in Computer Science, Springer, pp. 282-297, 9th International Conference on Evolutionary Multi-Criterion Optimization, Munster, Germany, 19/03/17.
https://doi.org/10.1007/978-3-319-54157-0_20

[Link to publication on Research at Birmingham portal](#)

Publisher Rights Statement:

Checked 26/1/2017

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

On using decision maker preferences with ParEGO

Jussi Hakanen¹ and Joshua D. Knowles²

¹ University of Jyväskylä, Faculty of Information Technology, P.O. Box 35(Agora),
FI-40014 University of Jyväskylä, Finland

`jussi.hakanen@jyu.fi`

² School of Computer Science, University of Birmingham, UK, Edgbaston B15 2TT
`j.knowles.1@cs.bham.ac.uk`

Abstract. In this paper, an interactive version of the ParEGO algorithm is introduced for identifying most preferred solutions for computationally expensive multiobjective optimization problems. It enables a decision maker to guide the search with her preferences and change them in case new insight is gained about the feasibility of the preferences. At each interaction, the decision maker is shown a subset of non-dominated solutions and she is assumed to provide her preferences in the form of preferred ranges for each objective. Internally, the algorithm samples reference points within the hyperbox defined by the preferred ranges in the objective space and uses a DACE model to approximate an achievement (scalarizing) function as a single objective to scalarize the problem. The resulting solution is then evaluated with the real objective functions and used to improve the DACE model in further iterations. The potential of the proposed algorithm is illustrated via a four-objective optimization problem related to water management with promising results.

Keywords: Surrogate-based optimization, interactive multiobjective optimization, preference information, computational cost, visualization

1 Introduction

Multiobjective optimization problems have several conflicting objective functions to be optimized simultaneously. Due to the conflict, optimal solutions have inherent trade-offs between the objectives and are called Pareto optimal solutions where none of the objectives can be improved without impairing some other one [15]. Without any additional information, all the Pareto optimal solutions are equally good. On a general level, one can see two approaches for solving multi-objective optimization problems: 1) approximate the whole Pareto front, that is, the set of all Pareto optimal solutions or 2) identify the most preferred Pareto optimal solution(s) utilizing preference information from a human decision maker. In this paper, we concentrate on the latter approach.

Decision maker preferences are nowadays widely used to find most preferred solutions in both multiple criteria decision making and evolutionary multiobjective optimization communities (see e.g. [2,4,15,21]). Preferences can be used

before optimization, after optimization, or iteratively during the optimization process [15]. The methods belonging to the latter approach are called interactive multiobjective optimization methods where the decision maker is iteratively involved in the solution process. The benefits of interactive methods include learning about the interdependencies between the conflicting objectives and ones preferences especially for problems with a high number of objectives. This new knowledge gained during the interactive solution process may necessitate adjusting decision maker's preferences. Interactive methods indeed allow the decision maker to change her preferences which is not possible if the preferences are asked before optimization (see, e.g., [22]). Therefore, they have been found promising for solving real-world problems [17,18].

In computationally expensive multiobjective optimization problems, a single evaluation of the objective and/or constraint functions takes considerable amount of time. To be able to tackle such problems, surrogate-based approaches are commonly used where the idea is to approximate the problem with simpler functions that are faster to evaluate [10,13,24,26]. There are different ways of using surrogates in multiobjective optimization [24], for example, 1) to approximate each objective function separately and apply any suitable multiobjective optimization algorithm by using the approximated functions (the most widely used approach, see e.g. [3,10,24]), 2) to approximate a scalarizing function that converts multiple objectives together with decision maker preferences into a single objective optimization problem [7,12], or 3) to approximate directly the Pareto front [6,9,19]. However from decision making point of view, providing most preferred solutions to the decision maker for computationally expensive multiobjective optimization problems was identified as a future research challenge in a recent survey [24].

The contribution of this paper is to introduce an interactive multiobjective optimization method for computationally expensive multiobjective optimization problems to partly answer the challenge mentioned above. It is based on the ParEGO algorithm [12] developed for approximating the whole Pareto front for computationally expensive multiobjective optimization problems. ParEGO uses a weighted Tchebycheff function to convert multiple objectives into a single objective optimization problem and randomly generates weights at each iteration to be able to approximate the whole Pareto front. Since ParEGO uses a surrogate to approximate a scalarization function, it is convenient to add also decision maker preferences and, thus, try to approximate only preferred parts of the Pareto front. This is done by replacing the weighted Tchebycheff function with an achievement scalarizing function and using reference points instead of weights to include decision maker preferences. *The decision maker guides the search by specifying preferred ranges for each objective function. By specifying ranges, the decision maker deals directly with objective function values (as opposed to e.g. weights that have no clear meaning) and it allows the decision maker to explore more if she does not have exact values in mind.* When compared to approaches that use surrogates to approximate directly the whole Pareto front (e.g. [6,9,19]), the benefit of interactive ParEGO is that it does not need any pre-computed

Pareto optimal solutions as an input. On the other hand, when compared to the approaches that use surrogates for each objective separately, the number of surrogates used is always one independently of the number of objectives used.

The rest of the paper is organized as follows. Section 2 presents the problem formulation along with brief introduction of the basics of the ParEGO algorithm. In Section 3, ideas of incorporating decision maker preferences into ParEGO are discussed and the interactive ParEGO algorithm is introduced. Section 4 illustrates the potential and ability of interactive ParEGO to utilize changing preference information through a four-objective example related to water management. Finally, the paper ends with conclusions and future research ideas presented in Section 5.

2 Basics of ParEGO

We consider a multiobjective optimization problem

$$\begin{aligned} & \text{minimize} \quad \{f_1(\mathbf{x}), \dots, f_k(\mathbf{x})\} \\ & \text{subject to} \quad \mathbf{x} \in S, \end{aligned} \tag{1}$$

where all k objective functions $f_i : \mathbb{R}^n \rightarrow \mathbb{R}$, $i = 1, \dots, k$, are to be minimized. The feasible region $S \subset \mathbb{R}^n$ denotes the set of feasible decision variable values $\mathbf{x} = (x_1, \dots, x_n)^T$. Here we assume that the function evaluations are costly in the sense of taking a long time.

The ParEGO algorithm [12] was developed for computationally expensive multiobjective optimization problems and it is based on the efficient global optimization (EGO) algorithm [11]. In ParEGO, the multiobjective problem (1) is converted into a single objective optimization problem by using the augmented Tchebycheff scalarizing function

$$\begin{aligned} & \text{minimize} \quad \max_{i=1, \dots, k} [w_i f_i(\mathbf{x})] + \rho \sum_{i=1}^k w_i f_i(\mathbf{x}) \\ & \text{subject to} \quad \mathbf{x} \in S, \end{aligned} \tag{2}$$

where $\mathbf{w} = (w_1, \dots, w_k)^T$, $w_i \geq 0$ and $\sum_{i=1, \dots, k} w_i = 1$. The term containing $\rho > 0$ is called an augmentation term and it is used to guarantee that the solutions of problem (2) are Pareto optimal [15]. An important reason for choosing the Tchebycheff scalarization was that it can find Pareto optimal solutions also in the non-convex part of the Pareto front. The overall goal in ParEGO is to find an approximation for the whole Pareto front. Next, the basic idea of ParEGO is briefly described.

ParEGO starts by using the latin hypercube sampling to find $11n - 1$ points in the decision space which are then evaluated by the real functions. Those points are then used to train a Kriging-based design and analysis of computer experiments (DACE) model to approximate the objective function in problem

(2). To train the DACE model, Nelder and Mead downhill simplex algorithm is used to maximize the likelihood. Then at each iteration of ParEGO, a new point to be evaluated with the real functions is determined by using a single objective genetic algorithm to maximize the expected improvement. After finding the new point, it is evaluated with the real functions and, then, the DACE model is re-trained by considering also the new point. The updated DACE model is then used to find the next point and this iteration continues until the budget for function evaluations with the real functions is exhausted.

We can make some important observations of ParEGO. Note that at each iteration, one new point is evaluated with the real functions. It is possible to use all the points evaluated so far with the real functions to train the DACE model. However, when the number of the points increases, the more time training takes. It is also possible to use only some subset of the points in training in order to reduce the training time while compromising the accuracy if needed [5]. Further, for each iteration, the weight vector $\mathbf{w}$ used in problem (2) is randomly generated in order to finally cover the whole Pareto front. In practice, ParEGO does not set any limitations for the number of objective functions to be considered since it only affects on the generation of the weight vectors. However, the current implementation found in https://github.com/CristinaCristescu/ParEGO_Eigen supports only problems up to four objective functions [5]. More details of ParEGO can be found in [12].

3 Preferences with ParEGO

Next we discuss how to include decision maker preferences into ParEGO. A natural way for this would be to apply the ideas of the interactive weighted Tchebycheff procedure introduced in [23] where the idea is to reduce the space of feasible weight vectors based on preferences from the decision maker. At each iteration, the decision maker sees some Pareto optimal solutions and is asked to select the most preferred one. Based on the selection, the weight vector space is reduced and the new Pareto optimal solutions are generated by using the reduced space of weight vectors. The limitations of the interactive weighted Tchebycheff approach include that one can not go back to the part of the space that has been eliminated (i.e. change preferences). It means that an important property of interactive approaches, a possibility to change one's mind, is not available. For that reason, we do not go towards that path but take another approach utilizing achievement (scalarizing) function (ACH) [25] and preferred ranges for the objective function values as a form of preference information that will enable the decision maker to change her mind during the interactive decision making process. Next, our approach is discussed in more details.

3.1 A priori preferences

We start by discussing how to include a priori preferences to ParEGO, that is, how to handle fixed preferences expressed before optimization. That will also

Fig. 1. Examples of different types of hyperboxes. The solutions in the Pareto Front (black curve) that can be reached from each hyperbox by using ACH are marked with Xs. The dotted lines denote the projection direction determined by the weights in ACH.

serve as a building block for interactive ParEGO described later. As mentioned above, we chose to use preferred ranges for objectives as decision maker preferences since dealing directly with objective function values is found cognitively easy way of expressing preferences for the decision maker [14]. Preferred ranges will result in a k -dimensional hyperbox $H = \{z \in \mathbb{R}^k | a_i \leq z_i \leq b_i\}$ in the objective space. In practice, this hyperbox can be considered as either 1) *optimistic* which means that it does not contain any feasible solutions, 2) *pessimistic* where all the solutions are dominated or 3) *neither optimistic nor pessimistic* which means that it intersects with the Pareto front [7]. Different types of hyperboxes are illustrated in Figure 1.

Further, multiple objectives are scalarized by using ACH instead of the weighted Tchebycheff scalarizing function used in original ParEGO. The mathematical formulation of the augmented ACH used here is

$$\begin{aligned} & \text{minimize} \quad \max_{i=1, \dots, k} w_i (f_i(\mathbf{x}) - \bar{z}_i) + \rho \sum_{i=1}^k w_i f_i(\mathbf{x}) \\ & \text{subject to} \quad \mathbf{x} \in S. \end{aligned} \tag{3}$$

The weights typically used with ACH are $w = (z^{nad} - z^{ideal})^{-1}$ where z^{ideal} and z^{nad} denote vectors containing the best and the worst values for the objectives in the Pareto front, respectively. Note that the weights used in the weighted Tchebycheff scalarizing function (2) are used as a preference information to get different Pareto optimal solutions as opposed to the weights in ACH (3) that

are used to normalize the scales of the objective function values. In the case of ParEGO, the objective functions are normalized based on the set of evaluated solutions and, therefore, weights $w_i = 1, i = 1, \dots, k$ are used. For the augmentation term, we here set $\rho = 0.05$ as has been done in original ParEGO. The solution of problem (3) is Pareto optimal [15] and different Pareto optimal solutions can be obtained by varying the reference point $\bar{z} \in \mathbb{R}^k$. Thus, the reference point $\bar{z}$ is the way to include preference information to ACH and, in our approach, to ParEGO. These reference points will be sampled from the hyperbox H determined based on the decision maker preferences in analogy of sampling the weights for the augmented weighted Tchebycheff function in the original ParEGO.

The basic idea here is that decision maker fixes the preferred ranges for the objectives resulting into a hyperbox H in the objective space. Then, ParEGO with ACH is run by using reference points sampled within H . Note that it is not needed to sample reference points from the whole H since it is known that all the reference points along the direction specified by the weight vector w give the same Pareto optimal solution. Therefore, it is enough to sample reference points in H such that they lie in the plane orthogonal to w .

To demonstrate how the modified ParEGO can handle a priori preferences and how our basic building block for the interactive ParEGO works, Figures 2 and 3 show two example runs for the three-objective DTLZ2 problem used in [12], which has the following formulation:

$$\begin{aligned}
 &\text{minimize } f_1 = (1 + g) \cos(x_1^\alpha \pi/2) \cos(x_2^\alpha \pi/2), \\
 &\text{minimize } f_2 = (1 + g) \cos(x_1^\alpha \pi/2) \sin(x_2^\alpha \pi/2), \\
 &\text{minimize } f_3 = (1 + g) \sin(x_1^\alpha \pi/2), \\
 &\text{subject to } g = \sum_{i \in \{3, \dots, 8\}} (x_i - 0.5)^2, \\
 &\quad x_i \in [0, 1], i \in \{1, \dots, 8\},
 \end{aligned} \tag{4}$$

where $\alpha = 1$. Figure 2 illustrates a case where ParEGO was run for 100 real function evaluations by using the preferred ranges $0.4 \leq f_1 \leq 0.5$, $0.5 \leq f_2 \leq 0.6$, and $0.4 \leq f_3 \leq 0.5$. Parallel coordinate plot of the non-dominated points obtained after the initial sampling together with the preferred ranges are shown in Figure 2 by using blue and red color (together with symbol X), respectively. Similarly, Figure 3 shows a case where the preferred ranges $0.1 \leq f_1 \leq 0.3$, $0.7 \leq f_2 \leq 0.9$, and $0.0 \leq f_3 \leq 0.2$ were used. In both the figures, one can observe that the obtained solutions follow the given preferences quite well, although not perfectly since the given ranges are not feasible.

3.2 Interactive ParEGO

In this paper, we are not satisfied with fixed preferences since our overall goal is to develop an interactive version of ParEGO where the decision maker is able to change her mind if needed to better steer the solution process towards preferred solutions. Therefore, we will use progressive preference articulation that is used

Fig. 2. Non-dominated solutions (in blue) obtained for the DTLZ2 problem by using 100 real function evaluations and the preferred ranges (in red and with symbol X) $0.4 \leq f_1 \leq 0.5$, $0.5 \leq f_2 \leq 0.6$, and $0.4 \leq f_3 \leq 0.5$.

Fig. 3. Non-dominated solutions (in blue) obtained for the DTLZ2 problem by using 100 real function evaluations and the preferred ranges (in red and with symbol X) $0.1 \leq f_1 \leq 0.3$, $0.7 \leq f_2 \leq 0.9$, and $0.0 \leq f_3 \leq 0.2$.

in interactive multiobjective optimization approaches [15,17,18]. Algorithm 1 shows the steps of the interactive ParEGO algorithm which has three input parameters. The first is the budget of function evaluations FE_{max} available for the decision making process. The second parameter fixes the frequency of interaction, that is, how many iterations to run ParEGO with given preferences before the next interaction. The last parameter determines how many solutions the decision maker wants to see at each interaction. The output of the algorithm will be the solution most preferred by the DM.

After initialization step 1, $11n - 1$ points are sampled in the decision space by using latin hypercube sampling in the same way than in the original ParEGO algorithm in step 2. The resulting points are evaluated by using the real functions and added to an archive A that will contain all the points evaluated with the real functions. In step 3, non-dominated solutions of A are determined and shown to the decision maker. If there are more than N_S non-dominated solutions in A , the number will be reduced to N_S . This can be done e.g. by using clustering. Otherwise, all the non-dominated solutions are shown to decision maker. The solutions can be visualized to the decision maker with parallel coordinate plots or some other suitable visualization technique depending on the number of objectives [16]. Step 4 is the termination step of the algorithm and it involves two different criteria: stopping if decision maker so desires or if the budget of function evaluations is about to be exceeded. Thus in order to have any interactions, the

Algorithm 1 Interactive ParEGO

- Input:** FE_{max} = budget for function evaluations; F_{IA} = frequency of interactions; N_S = a number of solutions to be shown to decision maker at each interaction
- Output:** The solution most preferred by decision maker
- 1: Initialize the archive of solutions evaluated with real functions $A = \emptyset$, interaction counter $it = 0$, and function evaluation counter $fe = 0$.
 - 2: Create initial population A^0 by latin hypercube sampling of $11n - 1$ points in the decision space, evaluate them by using the original functions, and update $A = A \cup A^0$ $fe = 11n - 1$.
 - 3: Find non-dominated solutions of A and show N_S of them to decision maker.
 - 4: If decision maker wants to stop or $fe + F_{IA} > FE_{max}$, go to step 8.
 - 5: Ask decision maker to indicate preferred ranges for the objectives.
 - 6: Run F_{IA} iterations of ParEGO with given preferences to obtain a set A^{it} of F_{IA} solutions evaluated with real functions.
 - 7: Update $A = A \cup A^{it}$, $fe = fe + F_{IA}$, $it = it + 1$ and go to step 3
 - 8: Ask decision maker to indicate the most preferred solution as the final solution
-

budget for function evaluations FE_{max} should be more than $11n - 1 + F_{IA}$. If the decision maker wants to continue and there is budget for function evaluations remaining, then the decision maker is asked to provide preferred ranges for all the objective functions in step 5. The given ranges are then treated as preference information in a way described in Section 3.1. Next in step 6, F_{IA} iterations of ParEGO are run with the given preference information resulting into the same amount of new points evaluated with the real functions. Then in step 7, those points are added to the archive A , the function evaluations used and the number of interactions are updated, and the algorithm continues from step 3.

4 Numerical example

To illustrate the performance and potential of the interactive ParEGO algorithm, we here describe an interactive solution process for the four-objective optimization problem related to water management. We chose this problem because the objective functions in that have some real meaning unlike the objectives in the synthetic test problems often used to evaluate the performance of multiobjective optimization methods. In this way, the interactive solution process becomes more meaningful and easier to follow. In addition, although the example problem is not computationally expensive, it enables faster testing and the limited budget of available function evaluations reflects the challenge with computationally expensive problems. Before describing the water management problem along with the interactive solution process, we illustrate how the interactive ParEGO reacts to changing preferences through the three-objective DTLZ4 problem.

4.1 DTLZ4

Here, we consider the three-objective DTLZ4 problem which is a modification of the DTLZ2 problem described in Equation (4) by using $\alpha = 100$. The Pareto

Table 1. Preferred ranges used with DTLZ4.

	f_1	f_2	f_3
1	[0.0, 0.2]	[0.8, 1.0]	[0.4, 0.6]
2	[0.4, 0.6]	[0.0, 0.2]	[0.8, 1.0]
3	[0.8, 1.0]	[0.4, 0.6]	[0.0, 0.2]
4	[0.4, 0.6]	[0.4, 0.6]	[0.4, 0.6]

Fig. 4. Illustration of changing preferences for the DTLZ4 problem.

front for both DTLZ2 and DTLZ4 is one eighth of a sphere of radius 1, centred on $(0, 0, 0)$ but using $\alpha = 100$ severely biases the density distribution of solutions. To demonstrate the effect of changing preferences, four different preferred ranges (shown in Table 1) were consequently used after initialization.

Parameter values used were $FE_{max} = 200$ and $F_{IA} = 30$. The resulting non-dominated solutions of a single run are shown as a scatter plot in Figure 4 where different colors/symbols denote solutions obtained after different interactions, i.e., with different ranges. As can be seen, the solutions obtained follow the ranges given despite the limited number of function evaluations used.

4.2 Water management problem

The problem deals with water management in the (hypothetical) Fast Water Valley on a stretch of 100 river miles presented in [20]. There is a Fresh Fishery situated near the head of the valley causing pollution to the river. The city of Fortuna having population of 300 000 produces municipal waste pollution and is located 50 miles downstream from the fishery. The measure for water quality is expressed in terms of dissolved oxygen (DO) concentration and both the pollutants are described in pounds of biochemical oxygen demanding material (BOD). There exist primary treatment facilities that reduce the BOD in the gross discharge of the fishery and the city by 30 percent. Additional treatment facilities would increase the tax rate in Fortuna and decrease the return on investment (ROI) from the Fresh Fishery. The decision maker is interested in four objective functions: f_1 = water quality in the fishery, f_2 = water quality in the city, f_3 = ROI from the fishery, and f_4 = addition to the city tax rate resulting in to the following optimization problem

$$\begin{aligned}
& \text{maximize } f_1 = 4.07 + 2.27x_1, \\
& \text{maximize } f_2 = 2.6 + 0.03x_1 + 0.02x_2 + \frac{0.01}{1.39-x_1^2} + \frac{0.3}{1.39-x_2^2}, \\
& \text{maximize } f_3 = 8.21 - \frac{0.71}{1.09-x_1^2}, \\
& \text{minimize } f_4 = -0.96 + \frac{0.96}{1.09-x_2^2},
\end{aligned} \tag{5}$$

subject to $x_1 :=$ Proportionate amount of removed BOD at fishery,
 $x_2 :=$ Proportionate amount of removed BOD at city,
 $0.3 \leq x_1, x_2 \leq 1$.

The units for f_1 and f_2 are $[mg/L \text{ of DO}]$ while the unit for f_3 is $[\%]$. The unit for the addition to the city tax rate f_4 is $[1/1000 \text{ of } \$]$, i.e., $f_4 = \delta$ results in the tax rate of about $\$ \delta$ per $\$1000$ assessed value at Fortuna. The decision variables x_1 and x_2 describe the proportionate amount of BOD removed from water discharge at the fishery and at the city, respectively.

4.3 Interactive solution process

Parameter values used in this study were $FE_{max} = 100$, $F_{IA} = 20$, and $N_S = 5$. These values were chosen just to illustrate the performance and a study on their effect in the performance of the method is left as a topic for future study. In addition, k-means clustering with squared Euclidean distance was used to reduce the number of produced solutions if it was more than N_S and, from each N_S cluster identified, the solution closest to cluster centroid was chosen. An alternative for k-means would be to use multiobjective clustering techniques where the number of clusters is not fixed but determined by the clustering algorithm [8]. Exploring this option is left for future research.

An important aspect related to interactive multiobjective optimization methods is the graphical user interface that implements the interaction between the

Fig. 5. Non-dominated solutions shown to the decision maker after the initial sampling.

Fig. 6. Non-dominated solutions (in blue) shown to the decision maker after the first interaction. The preferred ranges from the first interaction are shown in red with symbol X.

decision maker and the method. In this example, we have used the parallel coordinate plot tool developed in Prof. Patrick Reed's group in Cornell [1] to visualize the non-dominated solutions to the decision maker and for her/him to analyze them in order to adjust the preferred ranges if necessary. It is important to note that the figures of parallel coordinate plots used to illustrate the solutions in this paper do not really provide the same possibilities to analyze them as the tool itself. In other words, by using the tool the decision maker can interact with the visualizations and get a better understanding by e.g. filtering the solutions or changing the positions of different objectives to have better understanding of the trade-offs. Developing a more enhanced user interface is one of the future research topics.

To start with, latin hypercube sampling for this problem corresponded to 21 points in the two dimensional decision space and a set of five representative non-dominated solutions among those is shown in Figure 5. From the figure, the decision maker could also observe the initial ranges for the objectives that help in determining preferred ranges.

Our aim here was to show how changes in preferences during the solution process can be taken into account with the interactive ParEGO and we illustrate this by considering preferences from different stakeholders' points of view. We start from a citizen's perspective who is living in the city of Fortuna. The first goal is to try to maximize the water quality in the city and, thus, the preferred ranges $5.5 \leq f_1 \leq 6.2$, $3.3 \leq f_2 \leq 4.0$, $4.0 \leq f_3 \leq 6.0$, and $3.0 \leq f_4 \leq 5.0$ were used.

The resulting non-dominated solutions after clustering are shown in Figure 6 which also shows the preferred ranges from the first interaction (in red and with symbol X). One can observe that it was not possible to improve water quality in the city for more than around 3.45 which resulted in the city tax rate increase of over 9. This was not found satisfactory and, therefore, the aim was

Fig. 7. Non-dominated solutions (in blue) shown to the decision maker after the second interaction. The preferred ranges from the second interaction are shown in red with symbol X.

Fig. 8. Non-dominated solutions (in blue) shown to the decision maker after the third interaction. The preferred ranges from the third interaction are shown in red with symbol X.

to find satisfactory water quality levels in the city without too large tax rate increase. To continue with, the decision maker gave new ranges $5.5 \leq f_1 \leq 6.2$, $3.0 \leq f_2 \leq 3.5$, $2.0 \leq f_3 \leq 6.0$, and $5.0 \leq f_4 \leq 7.0$.

The set of five non-dominated solutions together with the preferred ranges are shown in Figure 7. Of the solutions shown to the decision maker, there were two within the preferred ranges for the city tax increase: $(5.00, 3.32, 7.44, 5.70)$ and $(6.09, 3.34, 5.82, 5.70)$. Both the solutions have similar values for f_2 and f_4 while the values differed for the other objectives. Thus, with these levels for the water quality in the city and the city tax increase, there exists a clear trade-off between the water quality in the fishery and the fishery ROI. For the next interaction, we looked at the problem from the fishery shareholder perspective and the aim was to maximize the ROI from the fishery and see what happens to the other objectives. Thus, the new ranges given were $4.0 \leq f_1 \leq 6.5$, $2.5 \leq f_2 \leq 3.5$, $6.0 \leq f_3 \leq 8.0$, and $0.0 \leq f_4 \leq 10.0$.

The resulting solutions are shown in Figure 8. One can observe, that in the solution with the best ROI, $(5.34, 2.86, 7.30, 0.00)$, the tax rate in the city does not increase since no additional water treatment is used. On the other hand, the water quality in both the city and the fishery is degraded. Thus, in order to keep the ROI from the fishery at a relatively high level and, simultaneously, not to compromise the water quality too much, the city tax rate should be increased.

Finally, we looked at the problem from perspective of the mayor of Fortuna. From her point of view, all the objectives are important since the quality of the water in the city directly affects to the living conditions, the fishery brings jobs and food to people in the city, and the tax rate should be kept at a relatively low level in order to keep the citizens happy. Therefore, she aims at a balanced

Fig. 9. Non-dominated solutions (in blue) shown to the decision maker after the fourth interaction. The preferred ranges from the fourth interaction are shown in red with symbol X .

Fig. 10. All the non-dominated solutions found during the interactive solution process.

solution between all the objectives and the preferred ranges $4.0 \leq f_1 \leq 6.5$, $2.5 \leq f_2 \leq 3.5$, $6.0 \leq f_3 \leq 8.0$, and $0.0 \leq f_4 \leq 10.0$ were taken around the middle of the ranges of each objective obtained among all solutions evaluated so far.

The resulting set of five representative non-dominated solutions along with the preferred ranges are shown in Figure 9. The first observation is that there is not a single solution that satisfies all the preferred ranges and it might be the case that such a solution does not exist. The closest one has the values (5.95, 3.27, 6.45, 4.11) which seems to balance all the objectives quite nicely although it has quite high ROI from the fishery. However, if one looks at all the non-dominated solutions found during the solution process shown in Figure 10, there are not many solutions having ROI from the fishery below 5.0. At this point, the interactive solution process ends.

5 Conclusions

In this paper, an interactive ParEGO algorithm utilizing decision maker preferences in an iterative manner is proposed. As opposed to the original ParEGO algorithm, the aim is not to approximate the whole Pareto front but to help the decision maker to find her most preferred solution. The key modifications to the existing ParEGO algorithm are 1) incorporating decision maker preferences to focus the search, 2) interacting with the decision maker by showing her/him intermediate non-dominated solutions and asking for preferred ranges for the objective functions, and 3) replacing the Tchebycheff scalarization with the achievement scalarizing function enabling the usage of reference points inside the algorithm. The potential of the proposed algorithm was demonstrated

through the four objective optimization problem in water management and its ability to consider preference information that changes during the solution process seems promising. The topics for future research include a study on the influence of the three parameters of the algorithm, that is, the budget for real function evaluations, the frequency of interaction, and the number of solutions shown to the decision maker, building a more enhanced graphical user interface for implementing the interaction, using multiobjective clustering to reduce the number of solutions shown to the decision maker, and finally, applying the method to computationally expensive real-world problems.

Acknowledgment

This work was supported by the FiDiPro project DeCoMo funded by TEKES, The Finnish Funding Agency for Innovation. The authors want to thank mr. Tinkle Chugh, prof. Yaochu Jin, prof. Kaisa Miettinen, and dr. Karthik Sindhya for discussions related to the ideas presented in the paper. In addition, the authors would like to acknowledge Prof. Patrick Reed's group Decision Analytics for Complex Systems (Cornell University) related to their parallel coordinate plotting tool.

References

1. Web-based parallel coordinate plot, <https://reed.cee.cornell.edu/parallel-axis-categories/parallel/>, visited on 14.10.2016
2. Branke, J., Deb, K., Miettinen, K., Slowinski, R. (eds.): Multiobjective Optimization: Interactive and Evolutionary Approaches. Springer-Verlag, Berlin (2008)
3. Chugh, T., Jin, Y., Miettinen, K., Hakanen, J., Sindhya, K.: A surrogate-assisted reference vector guided evolutionary algorithm for computationally expensive many-objective optimization. *IEEE Transactions on Evolutionary Computation* (2017), to appear
4. Coello, C.: Handling preferences in evolutionary multiobjective optimization: A survey. In: 2000 IEEE Congress on Evolutionary Computation (CEC), pp. 30–37. IEEE (2000)
5. Cristescu, C., Knowles, J.: Surrogate-based multiobjective optimization: ParEGO update and test. Workshop on Computational Intelligence (UKCI) (2015)
6. Eskelinen, P., Miettinen, K., Klamroth, K., Hakanen, J.: Pareto navigator for interactive nonlinear multiobjective optimization. *OR Spectrum* 32, 211–227 (2010)
7. Haanpää, T.: Approximation method for computationally expensive nonconvex multiobjective optimization problems. Jyväskylä Studies in Computing 157, PhD thesis, University of Jyväskylä, Jyväskylä (2012), available online, <http://urn.fi/URN:ISBN:978-951-39-4968-6>
8. Handl, J., Knowles, J.: An evolutionary approach to multiobjective clustering. *IEEE Transactions on Evolutionary Computation* 11(1), 56–76 (2007)
9. Hartikainen, M., Miettinen, K., Wiecek, M.M.: PAINT: Pareto front interpolation for nonlinear multiobjective optimization. *Computational Optimization and Applications* 52(3), 845–867 (2012)

10. Jin, Y.: Surrogate-assisted evolutionary computation: Recent advances and future challenges. *Swarm and Evolutionary Computation* 1(2), 61–70 (2011)
11. Jones, D.R., Schonlau, M., Welch, W.J.: Efficient global optimization of expensive black-box functions. *Journal of Global Optimization* 13(4), 455–492 (1998)
12. Knowles, J.: ParEGO: A hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Transactions on Evolutionary Computation* 10(1), 50–66 (2005)
13. Knowles, J., Nakayama, H.: Meta-modeling in multiobjective optimization. In: Branke, J., Deb, K., Miettinen, K., Slowinski, R. (eds.) *Multiobjective Optimization: Interactive and Evolutionary Approaches*, pp. 245–284. Springer-Verlag, Berlin (2008)
14. Larichev, O.: Cognitive validity in design of decision aiding techniques. *Journal of Multi-Criteria Decision Analysis* 1(3), 127–138 (1992)
15. Miettinen, K.: *Nonlinear Multiobjective Optimization*. Kluwer Academic Publishers, Boston (1999)
16. Miettinen, K.: Survey of methods to visualize alternatives in multiple criteria decision making problems. *OR Spectrum* 36(1), 3–37 (2014)
17. Miettinen, K., Hakanen, J., Podkopaev, D.: Interactive nonlinear multiobjective optimization methods. In: Greco, S., Ehrgott, M., Figueira, J. (eds.) *Multiple Criteria Decision Analysis: State of the Art Surveys*, pp. 931–980. Springer, New York, 2nd edn. (2016)
18. Miettinen, K., Ruiz, F., Wierzbicki, A.P.: Introduction to multiobjective optimization: Interactive approaches. In: Branke, J., Deb, K., Miettinen, K., Slowinski, R. (eds.) *Multiobjective Optimization: Interactive and Evolutionary Approaches*, pp. 27–57. Springer-Verlag, Berlin (2008)
19. Monz, M., Kufer, K.H., Bortfeld, T.R., Thieke, C.: Pareto navigation - algorithmic foundation of interactive multi-criteria imrt planning. *Physics in Medicine and Biology* 53(4), 985–998 (2008)
20. Narula, S.C., Weistroffer, H.R.: A flexible method for nonlinear multicriteria decision-making problems. *IEEE Transactions on Systems Man and Cybernetics* 19(4), 883–887 (1989)
21. Purshouse, R., Deb, K., Mansor, M., Mostaghim, S., Wang, R.: A review of hybrid evolutionary multiple criteria decision making methods. In: 2014 IEEE Congress on Evolutionary Computation (CEC), pp. 1147–1154. IEEE (2014)
22. Sindhya, K., Ojalehto, V., Savolainen, J., Niemistö, H., Hakanen, J., Miettinen, K.: Coupling dynamic simulation and interactive multiobjective optimization for complex problems: An APROS-NIMBUS case study. *Expert Systems with Applications* 41, 2546–2558 (2014)
23. Steuer, R.E., Choo, E.U.: An interactive weighted Tchebycheff procedure for multiple objective programming. *Mathematical Programming* 26(3), 326–344 (1983)
24. Tabatabaei, M., Hakanen, J., Hartikainen, M., Miettinen, K., Sindhya, K.: A survey on handling computationally expensive multiobjective optimization problems using surrogates: non-nature inspired methods. *Structural and Multidisciplinary Optimization* 52(1), 1–25 (2015)
25. Wierzbicki, A.P.: Reference point approaches. In: Gal, T., Stewart, T.J., Hanne, T. (eds.) *Multicriteria Decision Making: Advances in MCDM Models, Algorithms, Theory, and Applications*, pp. 9–1–9–39. Kluwer Academic Publishers, Boston (1999)
26. Zhou, A., Qu, B.Y., Li, H., Zhao, S.Z., Suganthan, P.N., Zhang, Q.: Multiobjective evolutionary algorithms: A survey of the state of the art. *Swarm and Evolutionary Computation* 1(1), 32–49 (2011)