Preface

Optimization is the task of finding one or more solutions which correspond to minimizing (or maximizing) one or more specified objectives and which satisfy all constraints (if any). A single-objective optimization problem involves a single objective function and usually results in a single solution, called an optimal solution. On the other hand, a multiobjective optimization task considers several conflicting objectives simultaneously. In such a case, there is usually no single optimal solution, but a set of alternatives with different trade-offs, called Pareto optimal solutions, or non-dominated solutions. Despite the existence of multiple Pareto optimal solutions, in practice, usually only one of these solutions is to be chosen. Thus, compared to single-objective optimization problems, in multiobjective optimization, there are at least two equally important tasks: an optimization task for finding Pareto optimal solutions (involving a computer-based procedure) and a decision-making task for choosing a single most preferred solution. The latter typically necessitates preference information from a decision maker (DM).

1 Modelling an Optimization Problem

Before any optimization can be done, the problem must first be modelled. As a matter of fact, to build an appropriate mathematical or computational model for an optimization problem is as important or as critical as the optimization task itself. Typically, most books devoted to optimization methods tacitly assume that the problem has been correctly specified. However, in practice, this is not necessarily always the case. Quantifying and discussing the modelling aspects largely depend on the actual context of the underlying problem and, thus, we do not consider modelling aspects in this book. However, we wish to highlight the following points.

First, building a suitable model (that is, the formulation of the optimization problem with specifying decision variables, objectives, constraints, and variable bounds) is an important task. Second, an optimization algorithm (single or multiobjective, alike) finds the optima of the model of the optimization problem specified and not of the true optimization problem. Due to these reasons, the optimal solutions found by an optimization algorithm must always be analyzed (through a post-optimality analysis) for their 'appropriateness' in the context of the problem. This aspect makes the optimization task iterative in the sense that if some discrepancies in the optimal solutions obtained are found in the post-optimality analysis, the optimization model may have to be modified and the optimization task must be performed again. For example, if the DM during the solution process of a multiobjective optimization problem learns that the interdependencies between the objectives do not correspond to his/her experience and understanding, one must get back to the modelling phase.

2 Why Use Multiple Objectives?

It is a common misconception in practice that most design or problem solving activities must be geared toward optimizing a single objective, for example, bringing maximum profit or causing the smallest cost, even though there may exist different conflicting goals for the optimization task. As a result, the different goals are often redefined to provide an equivalent cost or a profit value, thereby artificially reducing the number of apparently conflicting goals into a single objective. However, the correlation between objectives is usually rather complex and dependent on the alternatives available. Moreover, the different objectives are typically non-commensurable, so it is difficult to aggregate them into one synthetic objective. Let us consider the simple example of choosing a hotel for a night. If the alternatives are a one-star hotel for 70 euros, or a zero-star hotel for 20 euros, the user might prefer the one-star hotel. On the other hand, if the choice is between a five-star hotel for 300 euros, and a four-star hotel for 250 euros, the four-star hotel may be sufficient. That is, stars cannot be simply weighted with money. How much an extra star is valued depends on the alternatives. As a consequence, it may be very difficult to combine different objectives into a single goal function a priori, that is, before alternatives are known. It may be comparatively easier to choose among a given set of alternatives if appropriate decision support is available for the DM. Similarly, one cannot simply specify constraints on the objectives before alternatives are known, as the resulting feasible region may become empty, making the optimization problem impossible to solve.

It should be clear that multiobjective optimization consists of three phases: model building, optimization, and decision making (preference articulation). Converting a multiobjective optimization problem into a simplistic single-objective problem puts decision making before optimization, that is, before alternatives are known. As explained above, articulating preferences without a good knowledge of alternatives is difficult, and thus the resulting optimum may not correspond to the solution the user would have selected from the set of Pareto optimal solutions. Treating the problem as a true multiobjective problem means putting the preference articulation stage after optimization, or interlacing optimization and preference articulation. This will help the user gain a much better understanding of the problem and the available alternatives, thus leading to a more conscious and better choice. Furthermore, the resulting multiple Pareto optimal solutions can be analyzed to learn about interdependencies among decision variables, objectives, and constraints. Such knowledge about the interactions can be used to redefine the model of the optimization problem to get solutions that, on the one hand, correspond better to reality, and, on the other hand, satisfy better the DM's preferences.

3 Multiple Criteria Decision Making

The research field of considering decision problems with multiple conflicting objectives (or goals or criteria) is known as multiple criteria decision making (MCDM) or multiple criteria decision aiding (MCDA). It covers both discrete problems (with a finite set of alternatives, also called actions or solutions) and continuous problems (multiobjective optimization). Traditionally, in multiobjective optimization (also known as multicriteria optimization), mathematical programming techniques and decision making have been used in an intertwined manner, and the ultimate aim of solving a multiobjective optimization problem has been characterized as supporting the DM in finding the solution that best fits the DM's preferences. The alternating stages of decision making and optimization create typically an interactive procedure for finding the most preferred solution. The DM participates actively in this procedure, particularly in the decision-making stage. Decision making on alternatives discovered by optimization requires a more or less explicit model of DM's preferences, so as to find the most preferred solution among the alternatives currently considered, or to give indications for finding better solutions in the next optimization stage. Many interactive methods have been proposed to date, differing mainly in the way the DM is involved in the process, and in the type of preference model built on preference information elicited from the DM.

The origin of nonlinear multiobjective optimization goes back almost 60 years, when Kuhn and Tucker formulated optimality conditions. However, for example, the concept of Pareto optimality has a much earlier origin. More information about the history of the field can be found in Chap. 1 of this book. It is worth mentioning that biannual conferences on MCDM have been regularly organized since 1975 (first by active researchers in the field, then by a Special Interest Group formed by them and later by the International Society on Multiple Criteria Decision Making). In addition, in Europe a Working Group on Multiple Criteria Decision Aiding was established in 1975 within EURO (European Association of Operational Research Societies) and holds two meetings per year (it is presently in its 67th meeting). Furthermore, International Summer Schools on Multicriteria Decision Aid have been arranged since 1983. A significant number of monographs, journal articles, conference proceedings, and collections have been published during the years and the field is still active.

4 Evolutionary Multiobjective Optimization

In the 1960s, several researchers independently suggested adopting the principles of natural evolution, in particular Darwin's theory of the survival of the fittest, for optimization. These pioneers were Lawrence Fogel, John H. Holland, Ingo Rechenberg, and Hans-Paul Schwefel. One distinguishing feature of these so-called evolutionary algorithms (EAs) is that they work with a population of solutions. This is of particular advantage in the case of multiobjective optimization, as they can search for several Pareto optimal solutions simultaneously in one run, providing the DM with a set of alternatives to choose from.

Despite some early suggestions and studies, major research and application activities of EAs in multiobjective optimization, spurred by a unique suggestion by David E. Goldberg of a combined EA involving domination and niching, started only in the beginning of 1990s. But in the last 15 years, the field of evolutionary multiobjective optimization (EMO) has developed rapidly, with a regular, dedicated, biannual conference, commercial software, and more than 10 books on the topic. Although earlier studies focused on finding a representative set of solutions on the entire Pareto optimal set, EMO methodologies are also good candidates for finding only a part of the Pareto optimal set.

5 Genesis of This Book

Soon after initiating EMO activities, the leading researchers recognized the existence of the MCDM field and commonality in interests between the two fields. They realized the importance of exchanging ideas and engaging in collaborative studies. Since their first international conference in 2001 in Zurich, EMO conference organizers have always invited leading MCDM researchers to deliver keynote and invited lectures. The need for cross-fertilization was also realized by the MCDM community and they reciprocated. However, as each field tried to understand the other, the need for real collaborations became clear.

In the 2003 visit of Kalyanmoy Deb to the University of Karlsruhe to work on EMO topics with Jürgen Branke and Hartmut Schmeck, they came up with the idea of arranging a Dagstuhl seminar on multiobjective optimization along with two MCDM leading researchers, Kaisa Miettinen and Ralph E. Steuer.

The Dagstuhl seminar organized in November 2004 provided an ideal platform for bringing in the best minds from the two fields and exchanging the philosophies of each other's methodologies in solving multiobjective optimization problems. It became obvious that the fields did not yet know each other's approaches well enough. For example, some EMO researchers had developed ideas that have existed in the MCDM field for long and, on the other hand, the MCDM field welcomed the applicability of EMO approaches to problems where mathematical programming has difficulties.

The success of a multiobjective optimization application relies on the way the DM is allowed to interact with the optimization procedure. At the end of the 2004 Dagstuhl seminar, a general consensus clearly emerged that there is plenty of potential in combining ideas and approaches of MCDM and EMO fields and preparing hybrids of them. Examples of ideas that emerged were that more attention in the EMO field should be devoted to incorporating preference information into the methods and that EMO procedures can be used to parallelize the repetitive tasks often performed in an MCDM task. By sensing the opportunity of a collaborative effort, a second Dagstuhl seminar was organized in December 2006 and Roman Słowiński, who strongly advocated for inclusion of preference modelling into EMO procedures, was invited to the organizing team. The seminar brought together about 50 researchers from EMO and MCDM fields interested in bringing EMO and MCDM approaches closer to each other. We, the organizers, had a clear idea in mind. The presence of experts from both fields should be exploited so that the outcome could be written up in a single book for the benefit of both novices and experts from both fields.

6 Topics Covered

Before we discuss the topics covered in this book, we mention a few aspects of the MCDM field which we do not discuss here. Because of the large amount of research and publications produced in the MCDM field during the years, we have limited our review. We have mostly restricted our discussion to problems involving continuous problems, although some chapters include some extensions to discrete problems, as well. However, one has to mention that because the multiattribute or multiple criteria decision analysis methods have been developed for problems involving a discrete set of solution alternatives, they can directly be used for analyzing the final population of an EMO algorithm. In this way, there is a clear link between the two fields. Another topic not covered here is group decision making. This refers to situations where we have several DMs with different preferences. Instead, we assume that we have a single DM or a unanimous group of DMs involved.

We have divided the contents of this book into five parts. The first part is devoted to the basics of multiobjective optimization and introduces in three chapters the main methods and ideas developed in the field of nonlinear multiobjective optimization on the MCDM side (including both noninteractive and interactive approaches) and on the EMO side. This part lays a foundation for the rest of the book and should also allow newcomers to the field to get familiar with the topic. The second part introduces in four chapters recent developments in considering preference information or creating interactive methods. Approaches with both MCDM and EMO origin as well as their hybrids are included. The third part concentrates with Chap. 8 and 9 on visualization, both for individual solution candidates and the whole sets of Pareto optimal solutions. In Chap. 10-13 (Part Four), implementation issues including meta-modelling, parallel approaches, and software are of interest. In addition, various real-world applications are described in order to give some idea of the wide spectrum of disciplines and problems that can benefit from multiobjective optimization. Finally, in the last three chapters forming Part Five, some relevant topics including approximation quality in the EMO approaches and learning perspectives in decision making are studied. The last chapter points to some future challenges and encourages further research in the field. All 16 chapters matured during the 2006 Dagstuhl seminar. In particular, the last six chapters are outcomes of active working groups formed during the seminar.

7 Main Terminology and Notations Used

In order to avoid repeating basic concepts and problem formulations in each chapter, we present them here. We handle *multiobjective optimization problems* of the form

$$\begin{array}{ll} \text{minimize} & \{f_1(\mathbf{x}), f_2(\mathbf{x}), \ldots, f_k(\mathbf{x})\} \\ \text{subject to} & \mathbf{x} \in S \end{array} \tag{1}$$

involving k (≥ 2) conflicting *objective functions* $f_i : \mathbf{R}^n \rightarrow \mathbf{R}$ that we want to minimize simultaneously. The *decision (variable) vectors* $\mathbf{x} = (x_1, x_2, \ldots, x_n)^T$ belong to the nonempty *feasible region* $S \subset \mathbf{R}^n$. In this general problem formulation we do not fix the types of constraints forming the feasible region. *Objective vectors* are images of decision vectors and consist of *objective (function) values* $\mathbf{z} = \mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}),\ f_2(\mathbf{x}), \ldots, f_k(\mathbf{x}))^T$. Furthermore, the image of the feasible region in the objective space is called a *feasible objective region* $Z = \mathbf{f}(S)$.

In multiobjective optimization, objective vectors are regarded as optimal if none of their components can be improved without deterioration to at least one of the other components. More precisely, a decision vector $\mathbf{x}' \in S$ is called *Pareto optimal* if there does not exist another $\mathbf{x} \in S$ such that $f_i(\mathbf{x}) \leq f_i(\mathbf{x}')$ for all $i = 1, \ldots, k$ and $f_j(\mathbf{x}) < f_j(\mathbf{x}')$ for at least one index j. The set of Pareto optimal decision vectors can be denoted by $P(S)$. Correspondingly, an objective vector is Pareto optimal if the corresponding decision vector is Pareto optimal and the set of Pareto optimal objective vectors can be denoted by $P(Z)$. The set of Pareto optimal solutions is a subset of the set of weakly Pareto optimal solutions. A decision vector $\mathbf{x}' \in S$ is *weakly Pareto optimal* if there does not exist another $\mathbf{x} \in S$ such that $f_i(\mathbf{x}) < f_i(\mathbf{x}')$ for all $i = 1, \ldots, k$. As above, here we can also denote two sets corresponding to decision and objective spaces by $WP(S)$ and $WP(Z)$, respectively.

The ranges of the Pareto optimal solutions in the feasible objective region provide valuable information about the problem considered if the objective functions are bounded over the feasible region. Lower bounds of the Pareto optimal set are available in the *ideal objective vector* $\mathbf{z}^\star \in \mathbf{R}^k$. Its components $z_i^\star$ are obtained by minimizing each of the objective functions individually subject to the feasible region. A vector strictly better than $\mathbf{z}^\star$ can be called a *utopian objective vector* $\mathbf{z}^{\star\star}$. In practice, we set $z_i^{\star\star} = z_i^\star - \varepsilon$ for $i = 1, \ldots, k$, where ε is some small positive scalar.

The upper bounds of the Pareto optimal set, that is, the components of a *nadir objective vector* $\mathbf{z}^{\mathrm{nad}}$, are usually difficult to obtain. Unfortunately, there exists no constructive way to obtain the exact nadir objective vector for nonlinear problems. It can be estimated using a payoff table but the estimate may be unreliable.

Because vectors cannot be ordered completely, all the Pareto optimal solutions can be regarded as equally desirable in the mathematical sense and we need a *decision maker* (DM) to identify the most preferred one among them. The DM is a person who can express preference information related to the conflicting objectives and we assume that less is preferred to more in each objective for her/him.

Besides a DM, we usually also need a so-called *analyst* to take part in the solution process. By an analyst we mean a person or a computer program responsible for the mathematical side of the solution process. The analyst may be, for example, responsible for selecting the appropriate method for optimization.

Acknowledgments

We would like to take this opportunity to thank all the participants of the 2004 and 2006 Dagstuhl seminars for their dedication and effort, without which this book would not have been possible. Andrzej Wierzbicki's suggestion to include a discussion on the importance of modelling issues in optimization in the preface is appreciated. We thank Springer for supporting our idea of this book. K. Deb and K. Miettinen acknowledge the support from the Academy of Finland (grant no. 118319) and the Foundation of the Helsinki School of Economics for completing this task.

The topics covered in this book are wide ranging; from presenting the basics of multiobjective optimization to advanced topics of incorporating diverse interactive features in multiobjective optimization and from practical real-world applications to software and visualization issues as well as various perspectives highlighting relevant research issues. With these contents, hopefully, the book remains useful to both beginners and current researchers including experts. Besides the coverage of the topics, this book will also remain a milestone achievement in the field of multiobjective optimization for another reason. This book is the first concrete approach in bringing two parallel fields of multiobjective optimization together. The 16 chapters of this book are contributed by 19 EMO and 22 MCDM researchers. Of the 16 chapters, six are written by a mix of EMO and MCDM researchers and all 16 chapters have been reviewed by at least one EMO and one MCDM researcher. We shall consider our efforts worthwhile if more such collaborative tasks are pursued in the coming years to develop hybrid ideas by sharing the strengths of different approaches.

June 2008

Jürgen Branke,
Kalyanmoy Deb,
Kaisa Miettinen,
Roman Słowiński

Most participants of the 2006 Dagstuhl seminar on "Practical Approaches to Multi-objective Optimization": 1 Eckart Zitzler, 2 Kalyanmoy Deb, 3 Kaisa Miettinen, 4 Joshua Knowles, 5 Carlos Fonseca, 6 Salvatore Greco, 7 Oliver Bandte, 8 Christian Igel, 9 Nirupam Chakraborti, 10 Silvia Poles, 11 Valerie Belton, 12 Jyrki Wallenius, 13 Roman Słowiński, 14 Serpil Sayin, 15 Pekka Korhonen, 16 Lothar Thiele, 17 Włodzimierz Ogryczak, 18 Andrzej Osyczka, 19 Koji Shimoyama, 20 Daisuke Sasaki, 21 Johannes Jahn, 22 Günter Rudolph, 23 Jörg Fliege, 24 Matthias Ehrgott, 25 Petri Eskelinen, 26 Jerzy Błaszczyński, 27 Sanaz Mostaghim, 28 Pablo Funes, 29 Carlos Coello Coello, 30 Theodor Stewart, 31 José Figueira, 32 El-Ghazali Talbi, 33 Julian Molina, 34 Andrzej Wierzbicki, 35 Yaochu Jin, 36 Andrzej Jaszkiewicz, 37 Jürgen Branke, 38 Fransisco Ruiz, 39 Hirotaka Nakayama, 40 Tatsuya Okabe, 41 Alexander Lotov, 42 Hisao Ishibuchi