

Multiple objective project portfolio selection based on reference points

Theodor J. Stewart^{1,2}

Published online: 15 January 2016
© Springer-Verlag Berlin Heidelberg 2015

Abstract The problem of selecting a portfolio of research and development projects from a set of proposed projects subject to resource constraints is formulated as a multiobjective optimization problem. Three categories of objectives are identified, namely quantitative, qualitative and balance/distributional. The first two are linear in the binary decision variables, but the third involves nonlinear measures of discrepancy between desired and actual distributions of activity amongst defined categories. The result is a nonlinear combinatorial multiobjective optimization problem. A form of reference point approach is motivated and developed as a solution method, and a special purpose genetic algorithm is implemented for obtaining the solution numerically. Extensions to interactive approaches for exploration of the efficient frontier are presented and discussed.

Keywords Project portfolios · Multiobjective programming · Reference point method

JEL Classification 90B50 · 90C29

This research was partly supported by a grant from the South African National Research Foundation.

✉ Theodor J. Stewart
theodor.stewart@uct.ac.za

¹ Department of Statical Sciences, University of Cape Town, 7701 Rondebosch, Cape Town, South Africa

² Alliance Manchester Business School, University of Manchester, Manchester M15 6PB, UK

1 Introduction

The problem described here arose within the Innovation and Research Centre of ESKOM, the largest electricity generating company in Africa. Refer <http://www.eskom.co.za> for details on the company and its operations. The problem arises as part of the annual budgeting cycle, during which decisions have to be made regarding which of a set of competing projects are to be selected for funding in the next budget year. Budgetary and other resource constraints limit project selection, while a number of diverse company objectives need to be taken into consideration

An earlier approach to the same problem setting was described in Stewart (1991). At that time, both the availability of sophisticated algorithms and computational power limited solution methods to a very simple greedy heuristic, which still took a considerable time to solve (10 min for a problem of 200 projects, with computational times growing at least quadratically with number of projects). As a result, the original model fell into disuse, but the need for a project prioritization system (PPS) resurfaced more recently in the light of tightly constrained budgets. We thus revisited the problem, leading to some adaptations to the formulation, and perhaps more critically to algorithmic design based on a reference point approach and a specifically designed genetic algorithm. It is this redesign which forms the main body of the present paper. In Sect. 2 we describe the modelling structure and formulation of the PPS problem, followed in Sect. 3 by the reference point approach adopted and the design of a genetic algorithm for its solution. Section 4 discusses how the approach may be used interactively to explore the Pareto optimal set, and adaptations to the algorithm are proposed. Conclusions regarding future work are presented in Sect. 5.

2 Problem formulation

We suppose that N projects have been proposed, and we define binary variables x_i for $i = 1, 2, \dots, N$ to indicate whether project i is included in the portfolio for the planning period. Choice is constrained by the availability of scarce resources (typically budgets, equipment availability, personnel required in certain categories). Let a_{ij} be the amount of resource j required for execution of project i , and A_j the total availability of resource j ; this implies the constraints:

$$\sum_{i=1}^N a_{ij}x_i \leq A_j \quad \text{for } j = 1, \dots, J. \quad (1)$$

It is worth commenting here that in our context, cost was never an objective, but was rather a budgetary constraint. (In fact at one stage it was even suggested that we maximize costs in order to ensure maximum utilization of the available budget!).

The critical issue was that of identifying and capturing management objectives. Extensive discussions with the client led to an appreciation that these objectives could be grouped into three categories:

- *Directly quantifiable benefits* These may include economic benefits, improvements to capacity, etc., and are labelled by $k = 1, \dots, K$. An estimate of the benefit of type k accruing by execution of project i is given by b_{ik} , and we shall shortly discuss aggregation across projects.
- *Qualitative criteria* These may include impacts on safety, consumer satisfaction, etc. It is assumed that the assessment of the contributions of a particular project i is provided on a 5-pt Likert scale, but extended to include a 0 option if the project makes no contribution. Let $c_{i\ell}$ be the Likert scale assessment of the contribution of project i to qualitative objective ℓ , for $\ell = 1, \dots, L$.
- *Balance or distributional criteria* A corporate research and development department needs to ensure also that “justice is being seen to be done”, in the sense, for example, that levels of activity do not unfairly favour one client group over another, are balanced in terms of short, medium and long term objectives, and are aligned with strategic high level management objectives such as security of supply, environmental impact etc., which we shall expand in greater detail below.

We now examine the means by which these different categories of objectives are formulated for purposes of optimization. The first category is straightforward. If there are no interdependencies between projects, then the aggregate level of achievement for quantitative objective k is given by:

$$V_k^{\text{QN}} = \sum_{i=1}^N b_{ik}x_i. \quad (2)$$

We shall comment shortly on the issue of possible interdependencies. The second category of objectives is similarly treated, except that the overall contribution of a Likert scale score for a project will need to be related to the magnitude of the project. The aggregate level of achievement for qualitative objective ℓ (again, for now assuming no interdependencies) is defined as follows:

$$V_\ell^{\text{QL}} = \frac{\sum_{i=1}^N c_{i\ell}w_{i\ell}x_i}{\sum_{i=1}^N c_{i\ell}w_{i\ell}} \quad (3)$$

being a weighted proportion of the available contributions which are realized. Here, the weight $w_{i\ell}$ indicates the magnitude of the contribution of project i to the total. In the system, we allowed the user for each objective ℓ to set all weights equal, or to equate the weight to the usage of a specified resource j (i.e., $w_{i\ell} = a_{ij}$).

Optimization of the V_k^{QN} and V_ℓ^{QL} subject to the resource constraints (1) is a standard multiobjective project portfolio problem (in fact, a multiobjective knapsack problem) as surveyed for example in Yu et al. (2012). In Yu et al. (2012), interdependencies are also explicitly modelled, but it seemed that in our context, management would have difficulties in providing the necessary parameter estimates. For this reason, and recognizing that substantial interdependencies only occurred in a small number of cases, we adopted a simpler modelling approach. Suppose there exists a small subset of of projects, say \mathcal{I} exhibiting substantial

interdependencies. We then create a set of artificially constructed project definitions, or “metaprojects”, each of which is a feasible combination of projects from \mathcal{I} . The resource requirements and benefits from each metaproject can then separately be assessed (where this is typically a judgmental issue). Call this set of metaaprojects \mathcal{I}^* . The original projects in \mathcal{I} are then replaced by the set of metaprojects \mathcal{I}^* . The optimization then proceeds as for the real projects, except that for each such set \mathcal{I}^* we would require the following mutual exclusivity constraint:

$$\sum_{i \in \mathcal{I}^*} x_i \leq 1. \quad (4)$$

We note in passing that the same structure of artificially constructed projects subject to mutual exclusivity constraints can also be applied to a single project but potentially implemented at different levels of activity.

At this point, we also make the comment that we saw little reason formally to include modelling of uncertainties either in parameter estimates or in the project execution. If needed, the work here could perhaps be extended along the lines presented in Hassanzadeh et al. (2014). For our purposes, sensitivity analysis coupled to a provision for the user to fix certain projects in or out appeared to suffice.

The seriously confounding factor in the formulation related to the third category of objectives (balance and distributional criteria). The concept of balance in portfolio construction is discussed for example in Karsu and Morton (2014). They create a bicriterion model to combine efficiency with a measure of balance. Balance in their model is defined by a categorization of projects, with a desired proportion of inputs to be allocated to each category. An aggregate measure of imbalance may then be defined by the sum or maximum of absolute deviations between actual and desired amounts of inputs across all categories.

Our situation extends that of Karsu and Morton (2014), as many different categorizations are simultaneously considered. As mentioned earlier, categorizations may relate to different client groups, different time horizons (short, medium and long term benefits), and to alignment with various high level management objectives. Different inputs (resources) may be relevant to different categorization sets. For purposes of the multicriteria comparisons we thus made use of a dimensionless statistical measure for representation of imbalance, namely a form of chi-squared statistic which we now describe.

Each form of categorization thus defines a criterion for evaluating portfolios. Let M be the number of such categorizations, and for categorization m ($= 1, 2, \dots, M$) let n_m be the number of categories defined.

For each project i , let ρ_{imv} be the degree to which project i is associated with category v from categorization m . We require that $\sum_{v=1}^{n_m} \rho_{imv} = 1$, but two distinct cases can be distinguished as follows:

- *Simple binary classification* such that $\rho_{imv} = 0$ or 1 only; in this case we referred to the criterion as a *balance* criterion.

- *Varying degrees of contribution* such that ρ_{imv} is a real number in the interval $[0, 1]$; in this case we referred to the criterion and a *distributional* criterion.

In either case we can define the proportion of project activity associated with category v (for categorization m) by:

$$p_{mv} = \frac{\sum_{i=1}^N w_{im} \rho_{imv} x_i}{\sum_{i=1}^N w_{im} x_i} \quad (5)$$

where again we require some weighting w_{im} as indicator of project magnitude, which as in the case of qualitative criteria we set to be either equal or to equate the weight to the usage of a specified resource j (i.e., $w_{im} = a_{ij}$).

Let the desired proportion in this case be specified by management to be π_{mv} . For each categorization, a measure of discrepancy between the actual and desired proportions is provided by the relevant chi-squared statistic for deviation between the two distributions, namely:

$$D_m = \sum_{v=1}^{n_m} \frac{(p_{mv} - \pi_{mv})^2}{\pi_{mv}}. \quad (6)$$

The D_m represents the performance measure to be minimized for each m .

We thus have a non-linear combinatorial multiobjective problem involving $P = K + L + M$ objectives (where in practice this total number of objectives may add to 10 or more), solution methods for which are discussed in the next section.

3 Reference point genetic algorithm approach

A wide variety of approaches have been suggested for the solution of multiobjective optimization problems. A useful survey is provided by the pair of papers (Miettinen 2008; Miettinen et al. 2008). Approaches are sometimes classified by the timing of elicitation of preference information. At two extremes are (a) approaches that first characterize all Pareto optimal solutions from which the decision maker selects a final choice, and (b) those in which a complete preference model is first elicited and then applied to define a scalarized optimization problem maximizing preferences. In between the two extremes are interactive methods in which efficient solutions are explored progressively guided by local preference information.

In the present context, direct display of the Pareto frontier for evaluation appeared to be impracticable because of the large number of objectives and the availability of management time. At the outset, it was also not evident that decision makers would always be available for lengthy interactive processes, although it was possible that at a later stage some interaction may be included. For this reason, we chose to implement a *reference point approach* [see Miettinen et al. (2008) for a fuller description], as the reference point (a set of aspiration levels for each objective) is a simple representation of preferences, but can be extended to multiple reference points applied sequentially in an interactive method. In this section we discuss the

implementation of a reference point approach with a well-defined set of aspiration levels, while in the next section we extend the discussion to interactive options.

In order to present our reference point approach, we define $f_p(\mathbf{x})$ as the value for objective function i , for $p = 1, 2, \dots, P$ given the vector \mathbf{x} of binary decision variables. As seen previously, the first $K + L$ functions need to be maximized, while the rest are to be minimized, but we shall largely be able to express the structure without differentiating between minimization and maximization. For each objective function, define I_p as the ideal (best value) achievable amongst feasible solutions. For purposes of the algorithm, an exact value of I_p is not really needed, as it serves primarily in scaling the objectives and a good approximation serves our needs. In this project selection problem:

- For the quantitative and qualitative objectives, the ideal is easily estimated by the simple heuristic of rank ordering projects by the corresponding values, and taking from the top, observing mutual exclusivity constraints, until resources are exceeded.
- For the balance and distribution objectives, a value of 0 can never be improved upon, and even though not quite achievable, serves as a good estimate for the algorithm.

The aspiration levels, say g_p , needed to define the reference point is required. In the system provided to the users:

- For quantitative and qualitative objectives, the user specified a proportion of the ideal to act as a target, say γ_p ($0 < \gamma_p < 1$), so that the aspiration level becomes $g_p = \gamma_p I_p$. A default of $\gamma_p = 0.8$ was suggested, but users were free to experiment with other levels (described to them as “importance levels”).
- For the non-parametric chi-squared measures, the aspiration level was to be specified directly. A default of $g_p = 0.1$ was suggested, but users were again free to experiment with other levels.

We chose a smoother scalarizing function than the conventional augmented Chebychev measure, as this seemed to simplify computations. We thus sought to minimize the function $S(\mathbf{x})$ defined by:

$$S(\mathbf{x}) = \sum_{p=1}^P \left[\frac{\mathbf{I}_p - \mathbf{f}_p(\mathbf{x})}{\mathbf{I}_p - \mathbf{g}_p} \right]^4. \quad (7)$$

Note that the terms in brackets are expressed in the same form for all objectives, whether minimization or maximization, as $f_p(\mathbf{x}) < \mathbf{I}_p$ and $g_p < I_p$ for maximizing objectives and vice versa for minimizing objectives. Notice also that the bracketed terms are dimensionless, so that no rescaling of objectives is required. Users sometimes expect to see a weighting term included in (7), but the introduction of weights is redundant; the function is in fact already a weighted sum, with weights proportional to $[I_p - g_p]^{-4}$, i.e. determined by choice of reference level. In fact the

scalarizing function can be viewed as weighted distance from the ideal, with distance measured by an L_4 metric (which more strongly penalizes large deviations). Such weighted distances from ideal are in fact used in other methods to generate Pareto optimal solutions.

Minimization of $S(\mathbf{x})$ is now a uni-dimensional non-linear combinatorial optimization problem. The structure of the problem lends itself to a very simple special purpose form of genetic algorithm, which can rapidly be solved. General principles of genetic algorithms may for example be found in Chapter 4 of Deb (2001), and are taken here as familiar to the reader. The special purpose algorithm can then be described as follows:

Initial population generation For each solution to be randomly generated, projects are placed in random order, and then taken from the top, respecting mutual exclusivities, and excluding projects whose addition to the portfolio would violate a resource constraint.

Fitness Fitness is defined simply by the scalarizing function (7), as all constructed solutions are feasible so that no penalty terms are required.

Parent selection and crossover Each of the two parents is selected as best in a tournament of four randomly selected solutions in the population. To construct the child solution, any common assignments are retained, and the remaining assignments selected in the same way as for the initial population generation, subject to the constraints implied by the common assignments.

Mutation Projects are selected randomly for mutation. Assignments for the selected projects in the child solution are undone, and the initial population generation is repeated, with the constraint that the assignments for the unselected projects are retained.

Retention of population members Elitist selection.

After some experimentation, the parameters defining the algorithm were chosen as population size: 200; number of children per generation: 200; and probability of selecting a project for “mutation”: 0.05. Results were, however, quite robust to choice of these parameters. As an indication of computation properties, we display in Fig. 1 the rate of convergence of the algorithm for a problem involving 250 projects, three quantitative objectives, three qualitative objectives and three balance and distribution objectives.

The algorithm was coded in Pascal, and running time [on a Toshiba Portege computer with an Intel(R) Core(TM) i7-3540M CPU @ 3.00 GHz processor] for the above problem was approximately 60 ms per generation.

4 Effective exploration of the Pareto front

As has been mentioned previously, the user is free to modify the reference point at will, and thus to explore alternative Pareto optimal solutions. This may be rather ‘hit and miss’, and it may be desirable to provide more systematic guidance to the decision maker. One approach may be to employ a standard interactive procedure.




Fig. 1 Change in scalarizing function with numbers of generations

The NIMBUS method described in Miettinen et al. (2008) should easily be adapted to our problem setting. In this approach, a solution for a given reference point is obtained as described in the previous section. The decision maker then classifies performances on the criteria into five classes in terms of level of satisfaction, on the basis of which the reference point and the scalarizing function are modified, after which a new solution is obtained.

There was some concern that decision makers may not be available for frequent interactions of this nature, however. As an alternative, we were looking for a means of presenting a “snapshot” representation of the Pareto front, from which decision makers could obtain a more global perspective and subsequently to narrow the search to more desirable regions of the front. Evolutionary multiple objective (EMO) methods have been suggested for the “many objectives” problem, e.g. Wang et al. (2013), in which some preference information can restrict regions of the Pareto front, but the demands on management still to examine multidimensional regions was deemed in our case too excessive.

As an alternative approach we propose the following approach:

1. Select a small set of reference points, say R , bearing in mind the old maxim that subjects should not be exposed to more than “ 7 ± 2 ” simultaneous stimuli (Miller 1956), for each of which the solution minimizing (7) is obtained. We acknowledge in passing that the concept of multiple reference points has appeared in the EMO literature (e.g. Figueira et al. 2010), but this has been in the context of selecting sub-regions of the Pareto set, and not that of the creating a small number of discrete representative solutions.
2. Present the decision maker with the resulting set of solutions, requesting that they be classified as good, moderate or poor. This classification could perhaps be supported by use of a discrete choice multiple criteria decision analysis (MCDA). The reason for a three way classification is to aim at having a few (perhaps 3 or 4) in a better category, which can be assembled, depending on the user’s nature from either the good or the good + moderate categories.

3. Use this information to eliminate portions of the reference point space, after which a new set of reference points may be generated within the remaining space. The process is then repeated until the user is satisfied that a “good enough” solution has been found.

In many ways, this process is allied to that introduced in Steuer and Choo (1983). Their approach is based directly on randomly generating weights for the objectives. We have noted in the previous section that there is a close relationship between the reference points and equivalent weights, but the impacts of weighting vectors can be surprisingly non-intuitive. For this reason we prefer interaction with decision makers to be expressed in terms of the more intuitive concept of reference (aspiration) levels.

Our initial approach to generating reference points randomly (but this will be subject to further research studies) is based on the ideals I_p and some assessment of worst case performance which we shall term “nadirs”, say N_p . We simply derive these from the payoff table for quantitative and qualitative objectives, and by a more or less arbitrary worst case (set for the moment at 0.2) for the chi-squared statistics indicating performance on balance and distributional criteria. With the aim of obtaining a balance spread of reference points across the small number R (“ $\approx 7 \pm 2$ ”) chosen, our proposed method of randomly generating reference points may be described as follows. Randomly sort the objectives, and divide them into three groups, which we shall designate as H, M and L. Then assign reference levels to each objective in the form $g_p = (1 - \alpha)N_p + \alpha I_p$, where $\alpha = \alpha_H$ for objectives in group H, α_M for objectivess in group M and α_L for objectivess in group L. Here $\alpha_H, \alpha_M, \alpha_L$ are model parameters chosen such that $1 > \alpha_H > \alpha_M > \alpha_L > 0$.

Let the solutions obtained be indexed by $r = 1, \dots, R$, and the corresponding objective function values denoted by f_{rp} . The decision maker then classifies each solution r as good moderate or poor. In the light of this classification, the ideal and nadir points are replaced by \hat{I}_p and \hat{N}_p defined as follows:

- $\hat{I}_p = \beta_I \max_{r:\text{good}} f_{rp} + (1 - \beta_I)I_p$,
- $\hat{N}_p = \beta_N \min_{r:\text{moderate}} f_{rp} + (1 - \beta_N)N_p$,

where the β_I and β_N are also model parameters.

The process is repeated with the modified “ideal” and “nadir” until the decision maker is satisfied.

An issue still to be examined is whether the speed of computation for multiple reference points can be enhanced by simultaneous computations inspired by EMO approaches. The tentative procedure would be along the following lines:

- For each member of the population compute the scalarizing function for each of the R reference points. Rank order population members for each of these R cases, to create R separate lists.
- Cycling sequentially through the R lists, specify the fitness of the next solution in the list by corresponding scalarizing function value for that reference point,

provided that that the solution has not already had fitness allocated from another list, and *provided that* the number of fitness allocations from this list does not exceed $1/R$ of the total population. The second proviso ensures that solutions do not “crowd” round a small number of reference point cases.

The behaviour of such simultaneous generation will be investigated numerically in follow-up research.

5 Conclusions and future work

The software implementing the project prioritization system has been handed over to the client, and finally accepted, even though most of the features described in Sect. 4 were not included. Nevertheless, especially these last mentioned features do raise a number of interesting research questions which we will still want to address. These are the following:

- Exploration of alternative means of generating random sets of reference points that may lead to better characterization of the Pareto front.
- Exploration of alternative means of pruning the space of reference points for future iterations in the light of the classification of solutions provided by the decision maker.
- A simulation study of the quality of the solution obtained by the interactive process in comparison with hypothetically assumed “true” utility functions, rather along the lines of the research presented in Stewart (1999).
- Extension and implementation of the methods for simultaneous generation of solutions corresponding to each reference point.

Acknowledgments The assistance provided by Vince Micali, Mantombi Bashe and Awelani Phaswana of the ESKOM Research and Information Centre in evaluating and testing the model structures is gratefully acknowledged.

References

- Deb K (2001) Multi-objective optimization using evolutionary algorithms. Wiley, Chichester
- Figureira JR, Liefoghe A, Talibi E, Wierzbicki AP (2010) A parallel multiple reference point approach for multi-objective optimization. *Eur J Oper Res* 205:390–400
- Hassanzadeh F, Nemati H, Sun M (2014) Robust optimization for interactive multiobjective programming with imprecise information applied to R & D project portfolio selection. *Eur J Oper Res* 238:41–53. doi:[10.1016/j.ejor.2014.03.023](https://doi.org/10.1016/j.ejor.2014.03.023)
- Karsu Ö, Morton A (2014) Incorporating balance concerns in resource allocation decisions: a bi-criteria modelling approach. *Omega Intl J Manag Sci* 44:70–82. doi:[10.1016/j.omega.2013.10.006](https://doi.org/10.1016/j.omega.2013.10.006)
- Miettinen K (2008) Introduction to multiobjective optimization: noninteractive approaches. In: Branke J, Deb K, Miettinen K, Słowiński R (eds) *Multiobjective optimization: interactive and evolutionary approaches*, chap 1. Springer, New York, pp 1–26
- Miettinen K, Ruiz F, Wierzbicki AP (2008) Introduction to multiobjective optimization: interactive approaches. In: Branke J, Deb K, Miettinen K, Słowiński R (eds) *Multiobjective optimization: interactive and evolutionary approaches*, chap 2. Springer, New York, pp 27–57

- Miller GA (1956) The magical number seven plus or minus two: some limits on our capacity for processing information. *Psychol Rev* 63:81–97
- Steuer R, Choo E (1983) An interactive weighted Tchebycheff procedure for multiple objective programming. *Math Program* 26:326–344
- Stewart TJ (1991) A multicriteria decision support system for R&D project selection. *J Oper Res Soci* 42:17–26
- Stewart TJ (1999) Evaluation and refinement of aspiration-based methods in MCDM. *Eur J Oper Res* 113:643–652
- Wang R, Purshouse RC, Fleming PJ (2013) Preference-inspired coevolutionary algorithms for many-objective optimization. *IEEE Trans Evol Comput* 17:474–494
- Yu L, Wang S, Wen F, Lai KK (2012) Genetic algorithm-based multi-criteria project portfolio selection. *Ann Oper Res* 197:71–86. doi:[10.1007/s10479-010-0819-6](https://doi.org/10.1007/s10479-010-0819-6)