

This item was submitted to Loughborough's Institutional Repository (<https://dspace.lboro.ac.uk/>) by the author and is made available under the following Creative Commons Licence conditions.



For the full text of this licence, please go to:
<http://creativecommons.org/licenses/by-nc-nd/2.5/>

SOLUTION ANALYSIS IN MULTI-OBJECTIVE OPTIMIZATION

Alexander Brownlee and Jonathan Wright

School of Civil and Building Engineering,
Loughborough University, Loughborough,
Leicestershire, LE11 3TU, UK

A.E.I.Brownlee@lboro.ac.uk

ABSTRACT

Recent years have seen a growth in the use of evolutionary algorithms to optimize multi-objective building design problems. The aim is to find the Pareto optimal trade-off between conflicting design objectives such as capital cost and operational energy use. Analysis of the resulting set of solutions can be difficult, particularly where there are a large number (possibly hundreds) of design variables to consider.

This paper reviews existing approaches to analysis of the Pareto front. It then introduces new approach to the analysis of the trade-off, based on a simple rank-ordering of the objectives, together with the correlation between objectives and problem variables. This allows analysis of the trade-off between the design objectives and variables. The approach is demonstrated for an example building, covering the different relationships that can exist between variables and the objectives.

INTRODUCTION

Building design is an inherently multi-objective process, there being a trade-off to be made between two or more conflicting design objectives (such as between minimising both operating and capital cost). This has led to research into the application of simulation-based multi-objective optimization methods that identify the Pareto optimum trade-off between conflicting design objectives (Caldas, 2008; Flager et al., 2008; Geyer, 2009; Perfumo et al., 2010, Hamdy et al., 2011a; Villa and Labayrade, 2011). In this approach, the trade-off is represented by a set of equally optimal solutions, from which a single design solution must be selected for construction. Therefore, the benefit of the optimization process can only be realised if the results of the optimization can be analysed in a way that aids the decision-making process and the selection of the final design solution.

The analysis of multi-objective optimization results is non-trivial, in that the problem is multi-dimensional with several interacting relationships being of interest, particularly:

1. the trade-off between the design objectives;
2. the extent to which the problem variables drive the trade-off;
3. and the extent to which elements of building performance change along a trade-off and are influenced by the problem variables.

This paper focuses on the first two of these points. The difficulty of such an analysis is apparent when compared to the complexity of analysing the simulation results for a single design solution alone. Not only are there multiple simulation results to analyse, but also multiple design solutions consisting of many design variables (perhaps as many as 100 or more variables). Given the scale of the task, in terms of the number of design objectives, variables, and solutions to be analysed, it is probable that any approach to the analysis will be based on both quantitative metrics, and qualitative graphical procedures. The decision-making workflow is also likely to be iterative.

This paper reviews existing approaches to the analysis of solutions from multi-objective optimization problems, the majority of which use a visual plot of the Pareto front or set in the objective space to choose solutions for analysis. There are few existing approaches that identify systematically the impact of variables on the trade-off, particularly for problems with more than a few variables.

The paper then introduces an approach that allows analysis of the relationship (trade-off) between the design objectives, and also the extent to which the trade-off is driven by certain design variables. The approach is based on a simple rank-ordering of the objectives, together with the correlation between the objectives and the problem variables. The correlation allows an additional check for variables that do not appear to drive the trade-off. In addition, the relative impact of driver variables can be determined at a glance. This provides useful knowledge of the problem to inform the decision making process when selecting a final design from the trade-off. This is demonstrated for a two-objective optimization problem formulated for a five zone building.

Conclusions are drawn about the practicability of the approach, as well as observation being drawn about the types of problem variables and the characteristic relationships that can exist between the variables and the objective functions in building optimization problems.

A REVIEW OF EXISTING APPROACHES TO MULTI-OBJECTIVE ANALYSIS

There are two broad approaches to analysis of results from a multi-objective parameter optimization; quantitative metrics for comparison of the Pareto fronts or sets, and qualitative analysis based on observation of trends among the solutions in the fronts. There exist a large number of quantitative metrics designed for comparing Pareto fronts in the objective space. These measure three different aspects of the sets: the distance of the set from the “true” Pareto-optimal set, how uniformly distributed the members of the set are, and the extent of the set (that is, how wide a range of values in each objective is covered by the set). Such metrics include the commonly used hypervolume, generational distance and spread metrics – and there are several comparative surveys of these in the evolutionary computing literature including (Knowles and Corne 2002; Zitzler et al., 2003). Hypervolume is used to choose one front for analysis from those found by multiple runs in (Perfumo et al., 2010). Beyond this, while such metrics are useful for comparing the relative performance of different algorithms (or configurations of the same algorithm) on a problem, they are of limited use for decision-making or analysis of the trends in a particular front. Here, it is important to relate trends in the variables to the trade-off in the objectives.

Usually, in multi-objective optimization, we wish to make comparisons using more than one objective, but it is possible to use quantitative comparisons on one objective at a time. For example, in (Hamdy et al., 2011b), the optimization algorithm is divided into two steps, between which the variable bounds and penalties for constraints are adjusted, and comparisons of solutions are performed one objective at a time. Within this framework, the authors noted that some design variables have little or no influence on the results in some parts of the solution space. This was also previously observed by Wright et al (2002).

Qualitative methods focus on either a visualisation of the Pareto set in either the objective or variable space, or analysis of the raw variable values among the Pareto set. A common approach for 2-objective problems is to plot the objective values of the Pareto set in 2D, such as in (Farmani et al., 2005; Nassif et al., 2005). From such a plot, it is possible to visually select one “trade-off” solution for analysis (Perfumo et al., 2010).

A 2D plot of the Pareto set can also be simply presented with solutions identified with their variable in a table (Shi, 2011) or as rendered images (Caldas, 2008).

Similar to the idea of sorting in this paper, in (Hamdy et al., 2011a), the Pareto set is sorted by one of the objectives, with bar charts given for the values of each variable among the set. Coloured bands can be used to map between plots of the Pareto set in the objective and variable spaces (Villa and Labayrade, 2011), or groups of solutions in the Pareto set plot can be identified to highlight common features (Geyer, 2009). Plots showing the relationship of one or two variables with an objective are used by (Flager et al. (2008), with additional colouring to identify trends. A Pareto set plot can also include indicators for uncertainty and robustness (Hoes et al., 2011).

The Phi-array (Mourshed et al., 2011) is used to incorporate information from suboptimal solutions in the decision-making process. Two variables (positions for primary and secondary luminaire) are used to plot solutions on a grid; size and colour reflect fitness. This means it is possible to show multiple solutions in the same position, and identify connections between optimality and variable values.

It is difficult to compare more than two objectives at a time; even a 3D surface plot is hard to interpret. In (Jin et al., 2011), a 3D plot of the Pareto set has projections of points on the three axes to show precise locations with respect to the objectives. Points are also coloured to show window-to-wall ratio (one of the variables) and ranges of objective values among the set are given. Further, tables give variable values for parts of the Pareto set representing trade-offs between pairs of objectives.

Pairs of objectives are often compared for many-objective problems, for example (Geyer, 2009). In (Kim and Park, 2009), groups of three objectives were compared and the authors gave a table of the variable values for the whole Pareto set, sorted by one of the objectives. In (Suga et al., 2010), four objectives were each plotted in histograms, and cluster analysis was used to aid trend-finding among the variables. Parallel coordinate plots (Flager et al., 2008) can also help identify broad trends for many-objective problems; an example is given in Figure 1.

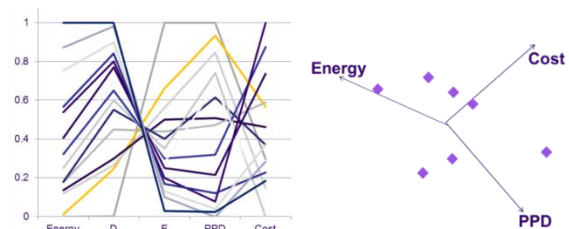


Figure 1. Parallel coordinate plots (left), and the Promethee / Gaia Method (right)

Outside of building optimization, there exist a wide number of techniques for visualisation and analysis

of the Pareto set such as the Gaia/Promethee (Brans and Mareschal, 1994) (Figure 1), and summary attainment surfaces for comparing multiple Pareto sets (Knowles, 2006). Automated approaches to learning the trends in variables (“design principles”) in the Pareto set are suggested as an area of future research in (Deb and Srinivasan 2008); these features are identified by visual and statistical comparisons.

Wang et al. (2009) recently reviewed multi-criteria decision analysis in sustainable energy decision-making. Recent reviews of visualisation techniques for the Pareto sets found by MOEA include (Korhonen and Wallenius, 2008; Lotov and Miettinen, 2008), and discussion of uncertainty and interactive decision making for MO problems is presented in (Bonissone et al. 2009).

AN APPROACH TO THE ANALYSIS OF SOLUTIONS TO BI-OBJECTIVE OPTIMIZATION PROBLEMS

Reviewing the literature, we see that often, once the Pareto set is found, the decision making process takes place entirely within the objective space. Analysis of the front tends to focus on individual solutions rather than trends within the front for each of the problem variables. We propose an approach which combines a table showing the values for each variable, with a visual bar showing the relative values, combined with the statistical correlation between each variable and one of the objectives. The correlation used is Spearman’s rank correlation (Lukey, 2002).

The technique makes use of conditional formatting in MS Excel, with the variable values normalised for easy comparison. There are two ways to approach the normalisation; to the lower and upper bounds of the problem variable, and to the minimum and maximum values present in the Pareto set. The former allows for the user to easily see how much coverage of each variable’s range there is, which would be useful if the optimization problem is to be revised. The latter can highlight trends that only cover a small part of a variable’s range but still nonetheless important.

Before demonstrating the approach, we will discuss the broad categories of trends that can be identified among the problem variables.

Variable Types

In defining a suite of test problems for evolutionary multi-objective optimization (EMO) algorithms, Huband et al. (2006) identify three categories of problem variable. These are defined by the influence they have on the position of solutions relative to the optimal trade-off (the Pareto front):

- The values given to *distance variables* determine how close to the Pareto front a solution lies (illustrated by “d” in Figure 2). We would expect that these would be constant along the front – although constant variables

could simply have no effect (Hamdy et al., 2011b; Wright et al., 2002).

- The values given to *position variables* determine where along the Pareto front a solution lies (illustrated by “p” in Figure 2).
- *Mixed variables* are a combination of both.

Huband et al. also identify *extremal* and *medial* variables, for which those in the Pareto front are all at the extreme or in the middle of the variable’s range respectively. In an iterative optimization process, the bounds of such variables could be adjusted to allow the search to focus on a wider range or more detail. It is simple to spot such variables when the values are normalised to the lower and upper bounds for the visualisation.

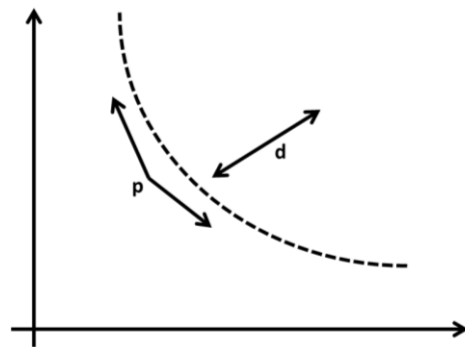


Figure 2. The impact of position (p) and distance (d) variables. The dotted line represents the Pareto front; the axes are the two problem objectives.

In this paper, we give examples of these for the Pareto front found for a building optimization problem. We also extend these definitions:

- Position variables are further categorised into primary and secondary position variables. This depends on whether they exhibit a single trend along the whole Pareto front, or a periodic trend, influenced by another variable.
- Floating variables, the values of which are unimportant (the objective functions are insensitive to these).
- Composite variables are a mixture of any of the above variable types.

These definitions are important because, in identifying the influence on the objectives that variables have relative to each other, we are able to better understand the problem, and make informed decisions about the final solution to be chosen.

The picture is further complicated by interactions between variables; such as a set-point having no effect if the system is not in operation (applies to example problem in terms of out-of-hours operation).

EXAMPLE OPTIMIZATION PROBLEM

The example optimization problem is based on a mid-floor of a commercial office building (Figure 3). Although the example has 5 zones, in this study, only the design variables associated with the perimeter zones are considered and optimized. The two end

zones are 24m x 8m, and the three middle zones 30m x 8m. The floor to ceiling height of all zones is 2.7m. The working hours are 9:00 to 17:00. Each zone has typical design conditions of 1 occupant per 10m² floor area and equipment loads of 11.5 W/m² floor area. Maximum lighting loads are set at 11.5 W/m² floor area, with the lighting output controlled to provide an illuminance of 500 lux at two reference points located in each of the perimeter zones. Infiltration is set at 0.1 air change per hour, and ventilation rates at 8 l/s per person. The heating and cooling is modelled using an idealized system that provides sufficient energy to offset the zone loads and meet the zone temperature set-point during hours of operation; free-cooling is available through natural and mechanical ventilation. Heating and cooling operation is restricted to separately identified seasons. The internal zone has been treated as a passive unconditioned space. The building performance has been simulated using EnergyPlus (V7). The building is nominally located in Birmingham UK, with the CIBSE test reference year used in simulating the annual performance (CIBSE, 2002).

Table 1 give the optimization **variables**. The building is orientated with the longest façades facing north (and south) when the orientation is set at 0°. The dead band has been optimized instead of the cooling set-point to ensure that problem formulation does not result in an overlap of the heating and cooling set-points. The window-to-wall ratios are applied by dividing the total window area across 6 windows placed in three groups across each façade (Figure 3). The names given to the window-to-wall ratios in Table 1 reflect the general orientation of the façade for the base solution (approximately that illustrated in Figure 3). The start and stop times are hours of the day.

The value of the categorical construction variables corresponds to a particular type of construction. Three construction types are available for the external wall construction, a heavy weight, medium weight, and light weight. Similarly two floor and ceiling constructions (heavy and light weight), and three internal wall constructions (heavy, medium, and light weight) have been defined. The alternative constructions have been taken from the ASHRAE

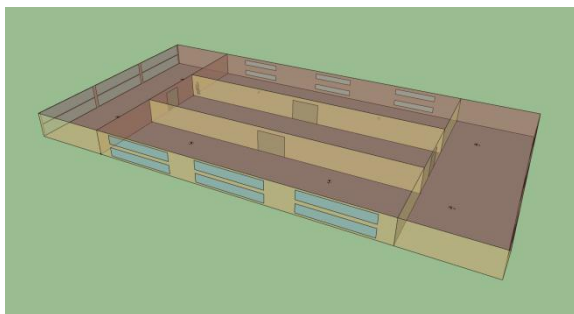


Figure 3. Example building.

handbook (ASHRAE, 2005). Two double glazed windows types are available, one having plain glass, and the second, low emissivity glass.

The **objective functions**, to be minimized by the optimization, are:

- the total annual energy use for heating, cooling, extractor fans, and artificial lighting;
- the capital cost of the building, using a model derived from cost estimating data.

The **design constraints** on thermal comfort and air quality during working hours in each of the perimeter zones are as follows. Air temperature must not exceed 25°C for more than 150 hours per annum, more than 28°C for more than 30 hours, or less than 18°C for more than 30 hours. CO₂ concentration should not exceed 1500ppm.

Table 1 (end of paper) details the 52 optimization problem variables, with the lower and upper bounds. These include orientation, heating and cooling set-points (via the dead band), window-to-wall ratios, start and stop times, and construction type.

OPTIMIZATION ALGORITHM

The optimization run was carried out using the Non-dominated Sorting Genetic Algorithm II (NSGA-II) (Deb et al., 2002). The specific implementation details were:

- Gray-coded bit-string encoding of the problem variables (163 bits)
- Population size 15
- Binary tournament selection
- Uniform crossover for every offspring
- Bit-flip mutation at a rate of 1/163
- Limit of 5000 unique simulations

The output set of non-dominated solutions was derived from the set of all solutions generated over the run, rather than the final population (the latter being limited by the population size). For the purposes of the analysis presented here, the results from a single run of the algorithm are adequate.

The run found 49 solutions in the trade-off, which are plotted in the objective space in Figure 4.

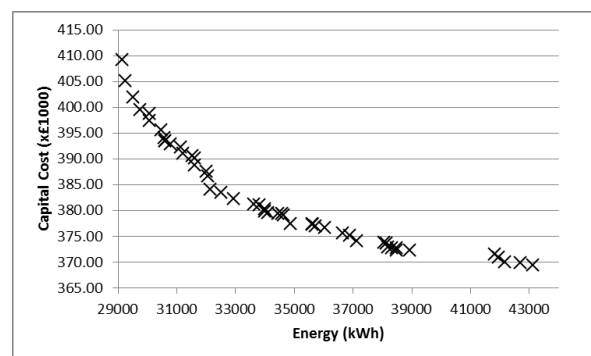


Figure 4. The Pareto optimal set found by the optimization, plotted in objective space.

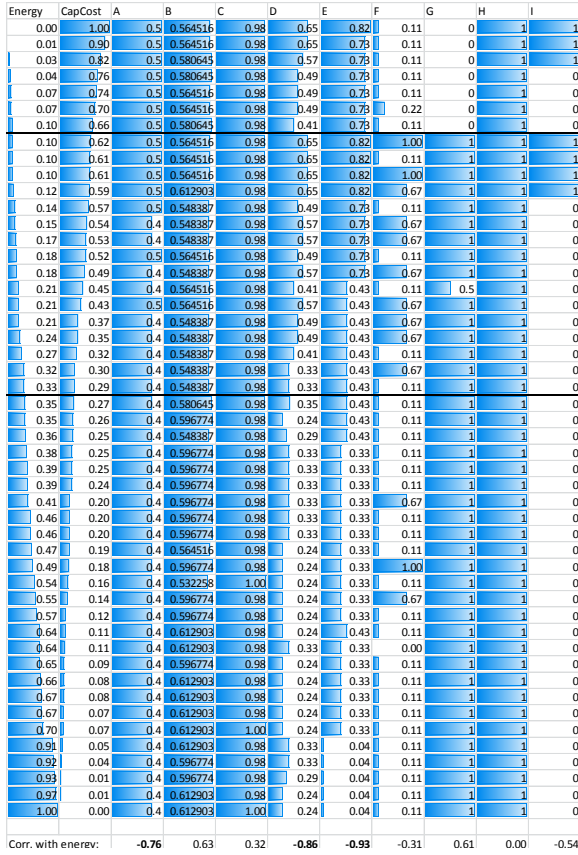


Figure 5. Selected variables for the Pareto front, normalised to the lower and upper bounds in the definition of the optimization problem.

EXAMPLE ANALYSIS

We now look at a sub-set of the problem variables which fall in to each of the categories identified earlier. For this, we refer to Figure 5 and Figure 6, in which we have visualized the values for each variable (columns) and each solution (rows) as a bar. The normalised numerical values for variables are also included, but are of little importance for our analysis. The bar lengths are normalised to the lower and upper bounds for the optimization problem in Figure 5, and to the range of variable values within the Pareto set in Figure 6. Normalisation of a specific value x_i of a variable x is conducted according to:

$$n(x_i) = \frac{x_i - \min(x)}{\max(x) - \min(x)}$$

The solutions are sorted in order of ascending energy use and hence descending capital cost. From left to right, the columns are:

Energy – the simulated annual energy usage

CapCost – the modelled capital cost for the design

A – HVAC heating set point for occupied hours

B – HVAC cooling set-point for unoccupied hours

C – min, outdoor temperature for natural ventilation

D – glazed area for the north upper window

E – glazed area for the south upper window

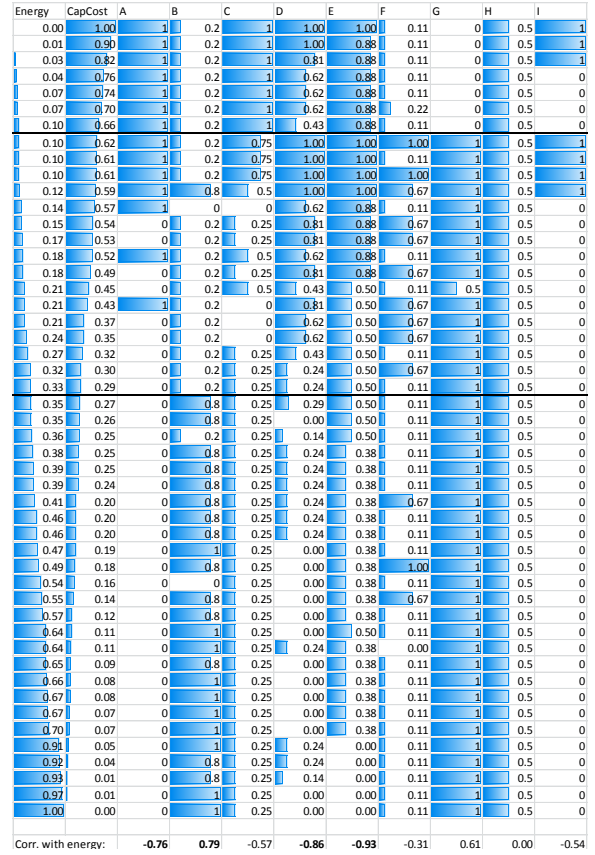


Figure 6. Selected variables for the Pareto front, normalised to the minimum and maximum values present in the Pareto front itself.

F – mechanical ventilation rate for the interior zone

G – external wall construction type

H – ceiling and floor construction type

I – shading overhang present on south lower window

The construction types are represented by numbers: 0 to 1 representing heavy to light weight constructions.

The final row in both figures is the statistical rank correlation between that variable and energy use. Figures over a magnitude of 0.7 (a strong correlation according to Moore (2010)) are in bold.

Primary position variables

The glazed area for the south upper window (E), HVAC heating set-point (A), and external wall construction (G) may all be regarded as primary position variables. In Figure 6, we can see that all three show broad trends in line with increasing energy usage. A lower window area, lower heating set-point, and a light-weight wall all lead to lower energy use for the building. In making an analysis of the trends, it is important to note that only the south upper window glazing area is easily identified as a position variable by both normalisation approaches. In Figure 5, the HVAC heating set-point (A) appears to vary very little because the Pareto set has converged to a small set of the allowed range of values for this variable. Only when normalised among the Pareto set (Figure 6) is it clear that this

variable still has an impact; this can also be seen in the strong statistical correlation (-0.76) between the variable and energy. In contrast, while the external wall construction appears to have an impact on energy use, it does not have a strong statistical correlation with energy (0.61), and may be filtered out by an approach that relies on this metric alone. In both cases, it could be that the comfort constraints have forced these variables to take on values within a narrow range, but the variable still has an impact on energy use, that is revealed by the trend. This emphasises the need for both visual and metric analysis of the set.

Secondary position variables

The glazed area for the north upper window (D), and south lower window overhang (I) can be regarded as secondary position variables. These have less impact on the energy use than the primary position variables, so appear to vary periodically, in line with changes in those. For the glazed area variable, there are three trends within the Pareto set: two in which the energy use increases with increasing glazed area, and a third in which the glazed area floats around a low value, while energy use continues to increase (affected by changes in other variables). Horizontal lines on both tables indicate the divisions between these regions. The first line clearly falls on the point where solutions change from having heavy weight external construction to light weight. The second falls at the point where the unoccupied cooling set-point changes from low to high.

The window overhang variable has two regions, with an overhang present for lower energy use buildings. The division between these regions also lies at the point where the external construction switches from heavy to light weight. As this point represents a large change in the makeup of the buildings, for robustness during final decision making it may be better to avoid solutions around this point in favour of those further from the transition.

As both variables exhibit a periodic trend along the Pareto front, rather than a single linear trend, the correlation coefficients show only a weak correlation with energy, when clearly they do have an impact which should be considered when decision making.

Fixed variables

Along the Pareto set, values for the floor and ceiling construction have become fixed on light-weight construction. As the variable has only one value for the whole set, correlation with energy is zero.

This implies that to reach the region of the Pareto front, solutions must have light-weight construction. If this applies only because of energy and cost considerations, then this is what Huband et al refer to as a distance variable; however it may also be that a lightweight construction means that the building is more easily able to meet the comfort constraints than with a heavyweight construction. An alternative

explanation can be a phenomenon known as hitch-hiking (Schraudolph and Belew, 1992), where a genetic algorithm assigns a value to an unimportant variable simply because it shared a solution with another variable value which was important. To be certain of the explanation, further exploration such as a formal sensitivity analysis around the points represented by the Pareto set would be required.

Floating variables

Our model deliberately includes one variable which has no impact on energy use or cost. The ventilation fan to the internal zone is switched off, so changing the flow rate for this fan has no effect on energy use. This can be seen in the values for the variable among the Pareto set – it “floats” around the whole range of possible values, and has a weak correlation with energy (-0.31). Note that this is greater than the correlation for the floor and ceiling construction (which was zero), despite these having an influence on the energy and cost of the building. Variables determined to be floating can be set to any value for the final design solution.

Composite variables

In practice, many variables will be a mixture of different types. A clear example is the minimum outdoor temperature for natural ventilation (C). Divided into the same three sections described earlier under “secondary position variables”, there are two constant regions and region where the variable floats.

DISCUSSION AND CONCLUSIONS

There are a large number of variables to consider, and it is difficult to analyse their characteristics over all 49 solutions in the Pareto front together. A typical approach may be to use a quantitative metric such as the correlation coefficient to filter out the solutions that show a lower correlation with the objectives for more detailed analysis.

A problem with this is that it fails to distinguish between distance variables (which tend to have fixed values in the front) and floating variables. The former are important as they have an impact on the objectives and constraints (energy, cost and comfort in this case). In contrast, floating variables do not and can have one of many values assigned to them at the final stage of the decision making process. A formal sensitivity analysis may provide a solution to this problem, filtering out variables found to be insensitive.

The correlation coefficient between objectives and variables should also be used with care as it fails to detect periodic trends such as that exhibited by composite and secondary position variables, as well as possibly failing to detect some of the primary position variables.

In this context, the graphical approach is useful. While for space reasons we cannot show all of the optimization variables here, it was possible at a

glance to see the fixed (distance) and position variables, and select a subset for closer analysis. Normalising to the values present in the Pareto set rather than the variable bounds made this process simpler, although the latter approach is still useful to see where variables cover the whole or a small part of their range. This would allow for improvement of the problem definition in subsequent optimization runs if the ranges initially chosen were inadequate.

In our example, we identified the categories of variable and how these appear as trends in the visualisation. Knowing which variables drive the trade-off between the objectives (both primary and secondary) is useful both in deciding on the values for the final chosen design, and understanding the characteristics of the model. Floating variables can be confidently set to any convenient value, whereas primary and secondary position variables may be set to mutually compatible values. The range of solutions found may increase confidence in the optimality, or at least the improvement gained through optimization. Further, if a variable which should be of little importance appears to be a primary position or a distance variable, there may be an issue with the model worth further investigation.

Further work is needed to extend this approach to three or more objectives, and to obtain better understanding of the influence of the variables on the optimization constraints. The optimization algorithm used here retained only “feasible” solutions (those meeting the constraints) for the final Pareto set. Analysis of “infeasible” solutions could be useful in providing information on areas of the solution space to avoid – this requires an approach in which infeasible solutions are allowed to be generated (such as a three objective approach), or via post processing.

ACKNOWLEDGEMENT

The research described in this paper was funded under UK EPSRC grant TS/H002782/1.

REFERENCES

- ASHRAE, 2005, “ASHRAE Handbook of Fundamentals”, Chp. 30, Table 19 and Table 22.
- BONISSONE, P.P., SUBBU, R. and LIZZI, J., 2009. Multicriteria decision making: a framework for research and applications. *Computational Intelligence Magazine, IEEE*, 4(3), pp. 48.
- BRANS, J. and MARESCAL, B., 1994. The PROMCALC & GAIA decision support system for multicriteria decision aid. *Decision Support Systems*, 12(4-5), pp. 297-310.
- CALDAS, L., 2008. Generation of energy-efficient architecture solutions applying GENE_ARCH: An evolution-based generative design system. *Adv. Engineering Informatics*, 22(1), pp. 59-70.
- CIBSE, 2002. *Guide J: Weather, solar, and illuminance data* CIBSE London.
- DEB, K., PRATAP, A., AGARWAL, S. and MEYARIVAN, T., 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. on Evol. Comp.*, 6, pp. 182-197.
- DEB, K. and SRINIVASAN, A., 2008. Innovization: Discovery of Innovative Design Principles Through Multiobjective Evolutionary Optimization. In: J. KNOWLES, D. CORNE, K. DEB and D.R. CHAIR, eds, *Multiobjective Problem Solving from Nature*. Springer Berlin Heidelberg, pp. 243-262.
- FARMANI, R., SAVIC, D.A. and WALTERS, G.A., 2005. Evolutionary multi-objective optimization in water distribution network design. *Engineering Optimization*, 37(2), pp. 167-183.
- FLAGER, F., WELLE, B., BANSAL, P., SOREMEKUN, G. and HAYMAKER, J., 2008. Multidisciplinary Process Integration and Design Optimization of a Classroom Building. *ITcon*, 14(August), pp. 595-612.
- GEYER, P., 2009. Component-oriented decomposition for multidisciplinary design optimization in building design. *Advanced Engineering Informatics*, 23(1), pp. 12-31.
- HAMDY, M., HASAN, A. and SIREN, K., 2011a. Applying a multi-objective optimization approach for design of low-emission cost-effective dwellings. *Building and Environment*, 46(1), pp. 109-123.
- HAMDY, M., HASAN, A. and SIREN, K., 2011b. Impact of adaptive thermal comfort criteria on building energy use and cooling equipment size using a multi-objective optimization scheme. *Energy and Buildings*, 43(9), pp. 2055-2067.
- HOES, P., TRCKA, M., HENSEN, J.L.M. and BONNEMA, B.H., 2011. Multi-objective optimisation of the configuration and control of a double-skin facade, *Proc. of the Building Sim. Conf.*, Sydney, Australia, pp. 1710-1717.
- HUBAND, S., HINGSTON, P., BARONE, L. and WHILE, L., 2006. A review of multiobjective test problems and a scalable test problem toolkit. *IEEE Trans. on Evol. Comp.*, 10(5), pp. 477.
- JIN, Q., OVEREND, M. and THOMPSON, P., 2011. A whole-life value based assessment and optimisation model for high-performance glazed facades, *Proc. of the Building Simulation 2011 Conf.*, Sydney, Australia 2011, pp. 1017-1024.
- KIM, Y. and PARK, C., 2009. Multi-criteria optimal design of residential ventilation systems, *Proc. of the Building Sim. 2009 Conf.*, pp. 1806-1813.
- KNOWLES, J., 2006. ParEGO: a hybrid algorithm with on-line landscape approximation for expensive multiobjective optimization problems. *IEEE Trans. on Evol. Comp.*, 10(1), pp. 50.

- KNOWLES, J. and CORNE, D., 2002. On metrics for comparing nondominated sets, Proc. of the IEEE Congress on Evol. Comp., 2002., pp. 711.
- KORHONEN, P. and WALLENIS, J., 2008. Visualization in the Multiple Objective Decision-Making Framework. In: J. BRANKE, K. DEB, K. MIETTINEN and R. SLOWINSKI, eds, Multiobjective Optimization. Springer, pp. 195-212.
- LOTOV, A. and MIETTINEN, K., 2008. Visualizing the Pareto Frontier. In: J. BRANKE, K. DEB, K. MIETTINEN and R. SLOWINSKI, eds, Multiobjective Optimization. Springer, pp. 213-243.
- LUCEY, T., 2002. Quantitative Techniques. 6th edn. Thomson Learning.
- MOORE, D.S., 2010. The basic practice of statistics. W H Freeman and Company.
- MOURSHED, M., SHIKDER, S. and PRICE, A.D.F., 2011. Phi-array: A novel method for fitness visualization and decision making in evolutionary design optimization. Adv. Engineering Informatics, 25(4), pp 676-687.
- NASSIF, N., KAJL, S. and SABOURIN, R., 2005. Optimization of HVAC Control System Strategy Using Two-Objective Genetic Algorithm. HVAC&R Research, 11(3), pp. 459.
- PERFUMO, C., WARD, J.K. and BRASLAVSKY, J.H., 2010. Reducing energy use and operational cost of air conditioning systems with multi-objective evolutionary algorithms, Proc. of IEEE Congr. on Evol. Comp. 2010, pp. 3937-3944.
- SCHRAUDOLPH, N.N. and BELEW, R.K., 1992. Dynamic Parameter Encoding for Genetic Algorithms. Machine Learning, 9, pp. 9-21.
- SHI, X., 2011. Design optimization of insulation usage and space conditioning load using energy simulation and genetic algorithm. Energy, 36(3), pp. 1659-1667.
- SUGA, K., KATO, S. and HIYAMA, K., 2010. Structural analysis of Pareto-optimal solution sets for multi-objective optimization: An application to outer window design problems using Multiple Objective Genetic Algorithms. Building and Env., 45(5), pp. 1144-1152.
- VILLA, C. and LABAYRADE, R., 2011. Energy efficiency vs subjective comfort: a multiobjective optimisation method under uncertainty, Proc. of the Building Simulation 2011 Conf., Sydney, Australia, pp. 1905-1912.
- WANG, J., JING, Y., ZHANG, C. and ZHAO, J., 2009. Review on multi-criteria decision analysis aid in sustainable energy decision-making. Renewable and Sustainable Energy Reviews, 13(9), pp. 2263-2278.
- WRIGHT, J.A., LOOSEMORE, H.A. and FARMANI, R., 2002. Optimization of building thermal design and control by multi-criterion genetic algorithm. Energy and Buildings, 34(9), pp. 959-972.
- ZITZLER, E., THIELE, L., LAUMANN, M., FONSECA, C.M. and FONSECA, V.G.D., 2003. Performance Assessment of Multiobjective Optimizers: An Analysis and Review. IEEE Trans. on Evol. Comp., 7(2), pp. 117-132.

Table 1. Example problem optimization variables. Bounds are min / max for numerical variables, and allowed values for categorical variables.

NAME	TYPE	BOUNDS	NAME	TYPE	BOUNDS
Heating setpoint	Float	18-23	N/S/E/W Vent height diffs	Float	0.3375-0.8775
Heating unoccupied setpoint	Float	10-23	N/S/ core mech-vent rate	Float	0.06-0.24
Cooling setpoint	Float	19-31	E/W mech-bent rate	Float	0.04-0.2
Cooling unoccupied setpoint	Float	12-43	Winter / Summer Start Times	Integer	1-8
Mech-vent setpoint	Float	18-31	Winter / Summer Stop Times	Integer	17-23
Mech-vent unoccupied setpoint	Float	10-43	Heat-to-cool change month	Integer	4-6
Nat-vent indoor outdoor TempDiff	Float	0-15	Cool-to-heat change month	Integer	9-11
Nat-vent min outdoor Temp	Float	-5-15	External wall construction	Category	Heavy, Med., Light
Orientation	Float	270-0-90	Internal wall construction	Category	Heavy, Light
N/S/E/W Upper window areas	Float	2.43-43.74	Ceiling & floor construction	Category	Heavy, Med., Light
N/S/E/W lower window areas	Float	1.62-51.03	N/S/E/W upper/lower	Category	Y/N
N/S/E/W upper window heights	Float	1.08-1.89	Window overhangs	Category	Standard, LowE
N/S/E/W lower window heights	Float	0.54-1.89	Glazing type	Category	