Loughborough University

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/

Multi-objective Optimization of Cellular Fenestration by an Evolutionary Algorithm

Jonathan Wright, Alexander Brownlee, Monjur Mourshed, and Mengchao Wang

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

Multi-objective Optimization of Cellular Fenestration by an Evolutionary Algorithm

Abstract

This paper describes the multi-objective optimized design of fenestration that is based on the façade of the building being divided into a number of small regularly spaced cells. The minimization of energy use and capital cost by a multi-objective genetic algorithm was investigated for; two alternative problem encodings (bit-string and integer); the application of constraint functions to control the aspect ratio of the windows; and the seeding of the search with feasible design solutions. It is concluded that the optimization approach is able to find near locally Pareto optimal solutions that have innovative architectural forms. Confidence in the optimality of the solutions was gained through repeated trail optimizations and a local search and sensitivity analysis. It was also concluded that seeding the optimization with feasible solutions was important in obtaining the optimum solutions when the window aspect ratio was constrained.

Keywords: façade optimization; evolutionary algorithms; multi-objective optimization; local and global sensitivity analysis.

1. Introduction

Fenestration and its design have a significant influence on the energy performance of buildings; it impacts on daylight penetration, artificial lighting energy use, and heating and cooling energy use. To date, optimization methods have been used to minimize building energy use by optimizing the dimensions of the windows and shading elements, but with the window shape typically being rectangular and only minor movement in the position of the windows on the façade being allowed (among others; (Caldas and Norford, 2002). Although this approach can result in reduced building energy use (and associated carbon emissions), the extent to which the energy use can be minimized is limited by the fact that the shape and position of the windows on the façade is fixed. Restricting the shape of the windows to be rectangular also inhibits the architectural form of the fenestration.

The objective of the research presented in this paper is to investigate an approach to fenestration optimization in which the shape, number, and position of windows can be optimized for the trade-off between minimum building energy use and minimum capital cost and in particular to investigate the performance of an Evolutionary Algorithm in solving this optimization problem. Flexibility in the shape, position, and number of windows is achieved by dividing the building façade into a number of small rectangular cells, each cell being defined by the optimization to be either a solid construction or glazed. The paper also investigates the use of a window aspect ratio constraint that allows the designer to control the general shape of the windows. The aspect ratio metric can also be used to indirectly constrain the number of windows.

Earlier work (Wright and Mourshed, 2009) found that a genetic algorithm was able to find near-optimal designs when applied to the minimisation of building energy use for both unconstrained and constrained shapes of windows. Design constraints for aspect ratio and number of windows were found to be effective in leading the optimization to produce different design solutions. The flexibility to define a range of constrained and unconstrained optimization problems, and to be able to solve them

using a single genetic algorithm is an important feature of this class of algorithm. Optimization for a single objective results in a single optimized design solution. In contrast, multi-objective evolutionary algorithms provide designers with a set of equally optimal design solutions that lie on the trade-off between two or more conflicting design goals. The alternative design solutions not only provide the designer with choice, but also information about the relationship between the conflicting design goals, and the importance of the design variables (Brownlee and Wright, 2012). This paper extends an earlier single-objective optimization study by introducing a shading variable associated with each glazed cell and extending the problem to two objectives, minimum energy use and minimum capital cost.

Previous research (Brownlee et al, 2011) established that among several evolutionary algorithms evaluated, the NSGA-II algorithm (Deb et al, 2002) had the best performance in solving a cellular window optimization problem. This paper extends this research by further examining the behaviour of NSGA-II in solving this class of problem, particularly the impact of the design constraints on the ability of the algorithm to find optimum solutions. Conclusions are drawn about the impact of the variable encoding on the ability of the algorithm to find optimized solutions, as well as the need to seed the search with feasible solutions when the aspect ratio of the windows is constrained.

2. Related Research

Genetic and evolutionary algorithms (EA's) have been proven to work well for a wide variety of building optimization problems. Their stochastic, population-driven operation allows them to efficiently overcome problems having local optimum solutions and large discontinuous search spaces, this having been shown to be important in the solution of potentially non-smooth building optimization problems (Wetter and Wright, 2004). EA's have been applied to both single and multi-objective problems including HVAC configuration and design (Hanby and Wright, 1989; Wright et al, 2008; Wright and Zhang, 2008; Sanaye and Hajabdollahi, 2010); optimal control strategies (Yang and Wang, 2011); building form (Caldas, 2008); structural design (Geyer, 2009); and other aspects of construction (Shi, 2011; Sambou et al, 2009).

EA's have also been applied to the optimization of glazing area and position for improved construction cost, operational energy use and occupant comfort (Evins, 2010; Oh et al, 2011; Jin et al, 2011; Caldas and Norford, 2002; Caldas, 2008; Suga et al, 2010), although typically the optimization focuses on glass type and glazed areas or the position of fixed shape windows.

Two approaches that have a near-freeform approach to glazing optimization are (Shea et al, 2006) and (Grierson, 2008). Both divide a façade into cells, with the former assigning opaque or clear glass to cells to balance lighting and cost, and the latter varying the transparency of the cells to balance light and heat gain and direct sunlight penetration. Both are similar in concept to the approach described here with the exception of the constraints on overall window shape and the application of shading. Unlike previous work, this paper also provides an analysis of the impact of the constraints on the ability of a multi-objective EA to find the Pareto optimum solutions.

3. Formulation of the Optimization Problem

In Wright and Mourshed (2009), the optimization of building fenestration was formulated by the building façade being divided into a number of rectangular cells. Each cell had one of two possible states, a window element, or an opaque construction

element. The number and location of the cells having a window construction was set by the optimization algorithm with the objective of the optimization being to minimize building energy use (and in some cases, while satisfying a constraint on the form of window geometry). The approach investigated in this paper extends this optimization problem to the minimisation of both energy use and capital cost, and the inclusion of 0.3m deep shading overhangs located above each of the glazed cells (each of which may be independently added or removed by the optimization process).
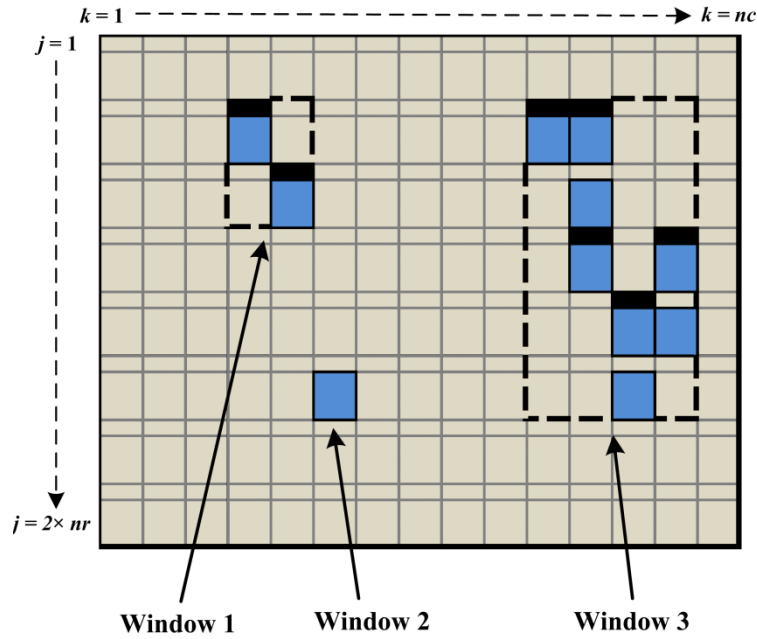

Figure 1. Cellular window façade

Figure 1, illustrates the cellular window approach. The individual window cells are shaded blue and the overhangs black; the remaining grey cells represent the opaque construction. A "window" is defined by a set of adjoining window cells there being 3 such "windows" in Figure 1. Note that the thickness of the overhang is exaggerated in the diagram and that two adjacent glazed cells are considered to be part of the same window whether adjoined by an overhang or not. Note also that as indicated for "Window 1", cells that are adjacent at their corner point are considered to be part of the same "window".

3.1 Problem variables and encoding

The optimization problem can be formulated in one of two ways. The first approach extends that described by Wright and Mourshed (2009), in which a vector of variables was mapped ($\rightarrow$), to a *nc* by *nr* matrix of binary decision variables $x_{k,j}$, one variable for each potentially glazed cell on the façade. This is extended here to provide separate vectors of binary decision variables for the glazed cells ($X^{glazed-opaque}$), and for the overhang cells ($X^{overhang}$). The extension to include the overhangs doubles the number of rows in the window matrix to $2 \times nr$ (Figure 1). The complete set of variables X is defined by Equation group (1):

$$X = (X^{glazed\_opaque}, X^{overhang});$$

where,

(1)

$$X^{glazed\_opaque} \rightarrow \begin{pmatrix} x_{1,2}, x_{2,2}, \cdots\cdots, x_{nc,2} \\ x_{1,4}, x_{2,4}, \cdots\cdots, x_{nc,4} \\ \vdots \\ x_{1,2\times nr}, x_{2,2\times nr}, \cdots\cdots, x_{nc,2\times nr} \end{pmatrix};$$

$$\forall j \in \{2, 4, \cdots, 2 \times nr\}$$

and,

$$X^{overhang} \rightarrow \begin{pmatrix} x_{1,1}, x_{2,1}, \cdots\cdots, x_{nc,1} \\ x_{1,3}, x_{2,3}, \cdots\cdots, x_{nc,3} \\ \vdots \\ x_{1,(2\times nr-1)}, x_{2,(2\times nr-1)}, \cdots\cdots, x_{nc,(2\times nr-1)} \end{pmatrix};$$

$$\forall j \in \{1, 3, \cdots, (2 \times nr - 1)\}$$

and

$$x_{k,j} \in \{0,1\};$$

$$\forall k \in \{1, \cdots, nc\};$$

Given that the variables $x_{k,j}$ are binary (a value of 0 resulting in an opaque construction or absence of an overhang and a value of 1 the presence of a glazed unit or overhang), this approach naturally lends itself to a binary chromosome encoding in the genetic algorithm.

A restriction on this particular problem formulation, which results from the limitations of the building performance modelling, is that an overhang can only exist above a glazed window unit. The impact of this is that the value of overhang variables associated with cells that have an opaque construction are redundant and are in effect "floating" in that they have no influence over either of the optimization objective functions. Although (as is evident from the results presented here), EA are able to cope with a degree of redundancy in the problem variables, this paper also investigates an approach that removes the variable redundancy. The alternative approach uses a vector of variables mapped to an nc by nr matrix of integer decision variables, each variable governing the choice of opaque construction, glazed construction without an overhang, or glazed construction with an overhang (the corresponding integer values being, 0, 1, and 2):

$$X \rightarrow \begin{pmatrix} x_{1,1}, x_{2,1}, \cdots\cdots, x_{nc,1} \\ x_{1,2}, x_{2,2}, \cdots\cdots, x_{nc,2} \\ \vdots \\ x_{1,nr}, x_{2,nr}, \cdots\cdots, x_{nc,nr} \end{pmatrix};$$

(2)

$$x_{k,j} \in \{0,1,2\};$$

$$\forall k \in \{1, \cdots, nc\};$$

$$\forall j \in \{1, \cdots, nr\}$$

In order to distinguish between the two approaches, the first approach in which the overhangs are defined by a separate set of variables is referred to here as the *bit-string encoding*, and the second approach in which the use of glazing and the overhang is specified through a single variable is referred to as the *integer encoding*. The reasons for pursuing both approaches in this study are that although the *integer encoding* halves the number of variables (and for the example problem investigated here has a $1\text{x}10^{15}$ smaller search space), the *bit-string encoding* is more generic in that it can be applied to any variable type, continuous or discrete, which will therefore allow the problem to be expanded in the future to include other design features in the optimization. The *integer encoding* also has a property that the choice of glazed area has twice the probability of occurring than an unglazed area (since an integer value of both 1 and 2 result in a window, a value of 0 being the only unglazed solution); the probability of selecting a window is unbiased in the *bit-string encoding*.

3.2 Objective functions

Two objective functions are simultaneously minimized by the optimization, the building energy use ($f^{energy}(X)$), and the capital cost of the façade ($f^{cost}(X)$). The energy use (kWh), is a function of the simulated annual heating (Q_l^{heat}), cooling (Q_l^{cool}) and electrical energy use (Q_l^{elect}):

$$\min f^{energy}(X) = \sum_{l=1}^{nl}\left(Q_l^{heat} + Q_l^{cool} + Q_l^{elect}\right) \tag{3}$$

where; l, is a building load condition, and nl is the number of load conditions; the electrical energy use includes the energy used in artificial lighting.

The capital cost of the façade is a function of the cost of the opaque construction, glazed construction, and cost of the shading overhang. Since the size of each cell is fixed, the cost can be calculated directly from the number of glazed units n_g, and the number of shading overhangs n_s. If n_t, is the total number of possible window cells ($n_c \times n_r$), then the capital cost of the façade is given by:

$$\min f^{cost}(X) = a\times\left(n_t - n_g\right) + b\times n_g + c\times n_s \tag{4}$$

where; a is the cost coefficient for the opaque construction; b the cost coefficient for the glazed construction; and c the cost coefficient for the shading overhangs.

3.3 Constraint functions

Wright and Mourshed (2009), investigated the use of constraint functions in restricting the optimized form of the windows. Two constraint functions where investigated, one to limit the total number of "windows", and a second to constrain the aspect ratio of the "window". It was concluded that constraining the aspect ratio also restricted the maximum number of windows, and so in this paper only aspect ratio is used as a constraint function. The research described here differs from previous research in that it investigates the impact of the constraints on the solution of the multi-objective problem rather than the single-objective optimization problem.

The aspect ratio of a window is defined here as the height divided by width. The height and width of a window are described in terms of the number of cells contained within the rectangular box that forms a boundary on all cells of the "window" (regardless of whether the cell is glazed or not). The bounding-box of the windows in Figure 1 is illustrated with a dashed line. The aspect ratio of a window i, is given by:

$$AR_i = \frac{N_i^{cellsHigh}}{N_i^{cellsWide}} \qquad (5)$$

where; AR_i is the aspect ratio of window i, $N_i^{cellsHigh}$ is the number of cells forming the height of the bounding-box, and $N_i^{cellsWide}$ is the number of cells forming the width of the bounding-box. As such, "Window 1" in Figure 1 has an aspect ratio of 1.0, "Window 2", 1.0, and "Window 3", 1.25.

In order to relax the optimization task, rather than constrain the solutions to have an absolute aspect ratio, a range of possible aspect ratios has been defined. The feasible range of aspect ratio is defined through two inequality constraints:

$$g(X)_{mx} - ub \le 0$$
$$lb - g(X)_{mn} \le 0 \qquad (6)$$

where, lb and ub are lower and upper bounds of the feasible aspect ratio; $g(X)_{mx} = \max(AR_i, \forall_i \in \{1, \ldots, n_w\})$ and $g(X)_{mn} = \min(AR_i, \forall_i \in \{1, \ldots, n_w\})$; and n_w is the number of windows found for a particular optimization solution.

4. The optimization algorithm

The multi-objective algorithm selected for use in this study is an implementation of the Non-Dominated Sorting Genetic Algorithm II (NSGA-II) first proposed by Deb et al. (2002). This is a well-known evolutionary algorithm, which has been successfully applied to a number of building optimization problems (Daum and Morel 2010; Sanaye and Hajabdollahi, 2010; Chantrelle et al, 2011; Jin et al, 2011; Evins et al, 2011).

In common with most evolutionary algorithms, NSGA-II starts with a randomly generated population of candidate solutions. The process of generating a new population of solutions is based on the selection of promising solutions which are then recombined and mutated (these operations being probabilistic). The process is repeated until either a predefined maximum number of generations is reached or a predefined maximum number of solutions evaluated.

The selection of solutions (individuals), for recombination is based on the solutions "fitness". The fitness of an individual is assigned in NSGA-II through the non-dominated sorting of the individuals in the population. The non-dominated sorting ensures that solutions which are better in both objectives, or violate fewer constraints, are given a higher rank (fitness) in the population. In order to maintain the diversity of

the Pareto front, NSGA-II uses a crowding distance penalty that is applied to individuals which have similar objective values to others in the population. The Pareto optimum front forming the solution to the optimization problem is simply the set of non-dominated individuals in the final population.

The appropriateness of this algorithm for use in solving this particular optimization problem was determined by comparing its performance with that of a number of other evolutionary multi-objective optimization algorithms (Brownlee et al, 2011). The comparison was based on the algorithms performance over 32 repeat optimization runs of each algorithm (each run starting with a different randomly generated population of solutions). Prior to the comparison of algorithm performance, factorial experiments were performed to optimally tune the parameters of each algorithm (including population size, mutation and crossover rate).,

The ideal algorithm resulting from the comparison and which is used in this study, can be summarised as having:
- random initialization of the problem variables, with the option to seed the initial population with viable solutions;
- a binary or integer encoding of the variables;
- binary tournament selection (De Jong, 2006), in which solutions of a lower rank are preferred, and in the case of equal rank, a higher crowding distance);
- 100% probability of chromosome crossover and a 50% probability of binary gene crossover (known as "uniform" crossover; De Jong, 2006);
- a probability of 1 gene mutation per chromosome;
- elitism: the top 50% of the combined set of parents and offspring are carried from one generation to the next;
- search completion after 5000 unique solutions have been evaluated;
- a population size of 30 solutions.

4.1 Seeding

Seeding the initial population with viable solutions is an established technique for improving the optimization efficiency of EA's (Julstrom, 1994), and which has also been applied to building optimization problems (Evins, 2010; Kayo, 2010). Seeding is used in this study to improve the feasibility of the optimized solutions. Experiments were performed in which a set of 20 feasible solutions that covered a range of glazing areas from the minimum to the maximum area where used as a seed in the initial population. The seed solutions were manually generated to produce an arbitrary set of solutions that contained a variety of forms and covered a range of glazed areas (Appendix A); the only criterion applied to all seeds was that they satisfied the aspect ratio constraints. The remaining 10 solutions of the 30 individual population were randomly generated.

5. Example building and performance simulation

The example optimization is based on an atrium of a three-storey commercial building (Wright and Mourshed, 2009). The atrium is 15m wide by 15m long by 8.2m high with only the southern façade being exposed to the external environment. The other three walls of the atrium are connected to interior spaces that are controlled to have the same thermal conditions as the atrium. The internal and external walls, roof, and floor are all of a medium to heavy weight construction; the window cells are double-glazed with plain glass.

The façade is divided into $120 \times 1m^2$ cells, 15 cells wide by 8 cells high (Figure 1). The higher the number of cells, the greater the flexibility of the search to optimize the number, shape, and position of the windows. This is similar to the impact of the number of pixels in a digital image, the higher the number of pixels the greater the clarity and resolution of the image. The number of cells has been capped at 120 by both the practicable size of the individual glazed units and a desire to limit the size of the search space (the search space increases in the order of 3^n for the integer encoding, where n is the number of cells); a 1 m^2 glazed unit has been taken as the practicable lower limit of unit size for the example façade.

5.1 Energy simulation

The building performance and energy use objective function (Equation 3), has been modelled and simulated using the EnergyPlus whole building performance simulation program (v 6.0.0.023; Crawley et al. 2001). The atrium has been modelled as an independent zone with the internal partition walls being treated as adiabatic heat transfer surfaces. The performance of the atrium heating, ventilating, and air-conditioning (HVAC) system has been modelled using a pseudo-system having an idealized 100% efficiency (the EnergyPlus "ideal loads air system"). The system is operated 24 hours/day with cooling and heating set-points for the occupied periods being set to 25.6 °C and 20 °C respectively; night setback set-point temperatures are 30.0 °C for cooling operation and 15.0 °C for heating operation.

Artificial lighting (and its energy use), is used to supplement daylight illuminance levels by an amount that maintains the illuminance set-point (500lux), calculated at two reference points (both reference points having equal weight in the control of the artificial lighting). The two reference points are located along the mid-point of the façade width, and at a distance of 25% and 75% of the depth of the atrium. The building is fully occupied from 09:00 to 17:00; the occupancy is reduced during the periods from 07:00 to 09:00, and 17:00 to 22:00.

The building is nominally located in Chicago, USA which experiences a continental type climate. This ensure that the optimization study relates to a building that has both an active heating and cooling demand.

5.2 Construction cost

The relative weights of the opaque and glazed constructions as described by Equation 4, have been derived from data presented by Spon (2011).The resulting coefficients for Equation 4 $a = £112$, $b = £350$, and $c = £128$. It is evident that the cost of the glazed construction is over three times higher than for the opaque construction.

5.3 Design constraints

The lower lb, and upper bounds ub, of the constraint functions (Equations 6), have been set to 1.5 and 1.75 respectively. This ensures that the windows are taller than they are wide and that they give the impression of being vertically orientated. The choice of aspect ratio has been abstracted from the Golden Ratio (Height/Width=1.62), this ratio often being used to define the geometric proportions of a building . Given the 15 x 8 cell structure of the façade, this results in 8 possible window geometries that satisfy the constraints (Table 1).

Height (cells)	Width for an approximate Golden Ratio (cells)	Aspect Ratio (-)
8	5	1.60
7	4	1.75
6	4	1.50
5	3	1.67
3	2	1.50

Table 1. Possible range of solutions with aspect ratio constrained

5.4 Inherent characteristics of the trade-off between energy and cost

The inherent characteristic of the trade-off between energy use and capital cost has been investigated here through an incremental traverse across the façade. The traverse incrementally adds window cells starting at the top right-hand corner (top-east corner in the example), and ends in the bottom-right corner (bottom-west corner). In relation to Figure 1, the traverse starts at position ($k=1$, $j=1$), and ends at ($k=nc$, $j=nr$), with the cells being filled first across the face of the façade before the traverse moves down to the next row of cells (in this way, the glazed cells are populated from the top of the façade to the bottom of the façade). All windows in the traverse are automatically shaded by overhangs. Since the traverse ignores the constraints on the window aspect ratio, the results only serve as an indication of the inherent relationship between the energy use and capital cost.

Figure 2, illustrates the relationship between energy use and capital cost resulting from the traverse. The solution with the minimum capital cost and the maximum energy use corresponds to a zero glazed area. As the number of glazed cells increases, the capital cost increases until a point of inflection is reached beyond which the energy use increases with increasing cost (and the number of window cells).
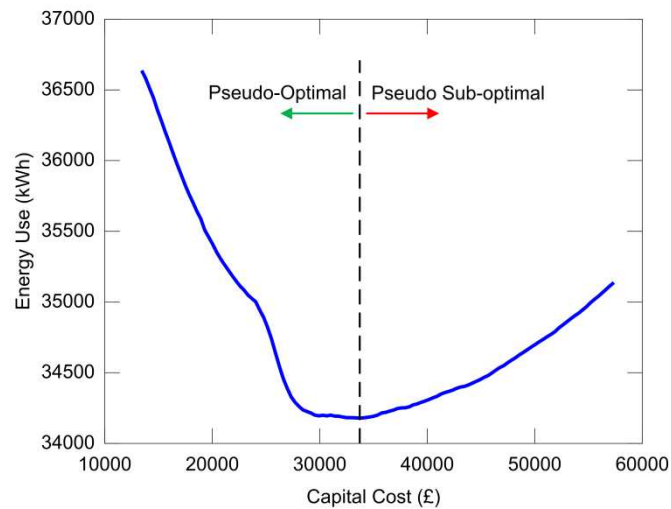
Figure 2. Relationship between energy use and capital cost

The minimum energy use resulting from the traverse corresponds to 55 glazed cells and a glazed area of 45.8%. Given that the energy use increases with cost beyond this point, it could be considered that the minimum energy use solution delineates between the pseudo optimal, and pseudo sub-optimal region of solutions. Note that the solutions illustrated in Figure 2 are not Pareto-optimal, and although it is expected that the window cells would be distributed towards the top of the façade (as this maximizes daylight penetration and reduces lighting energy use), it is likely that the optimized number and distribution of window cells will be different to that found from the simple incremental traverse of the solution space (Wright and Mourshed, 2009). This is particularly true when the aspect ratio constraint is applied as this forces the windows to be tall and thin rather than filling the full width of the façade, as is the case with the traverse.

The traverse also provides the minimum and maximum values for cost and the range of energy use. The cost ranges from £13,440 for no glazing and £57,360 for a fully glazed façade, and the energy use from 34,180kWh per annum to 36,637kWh per annum.

Figure 3 illustrates the separate energy use for heating, cooling and electrical lighting. The energy use is dominated by the heating of the atrium until more than 70% of the façade is glazed and beyond which cooling dominates. The pseudo optimum window area is dictated by the trade-off between the reduction in heating and lighting (electrical) energy and the increase in cooling energy with the number of window cells. The greatest rate of reduction in lighting (electrical) energy use occurs between 1 and 40 window cells, with no significant reduction for a further increase in number of cells. The region in which no further reduction in lighting energy use is achieved also marks a point at which the total energy use begins to increase with the glazed area (and the solutions become pseudo sub-optimal).
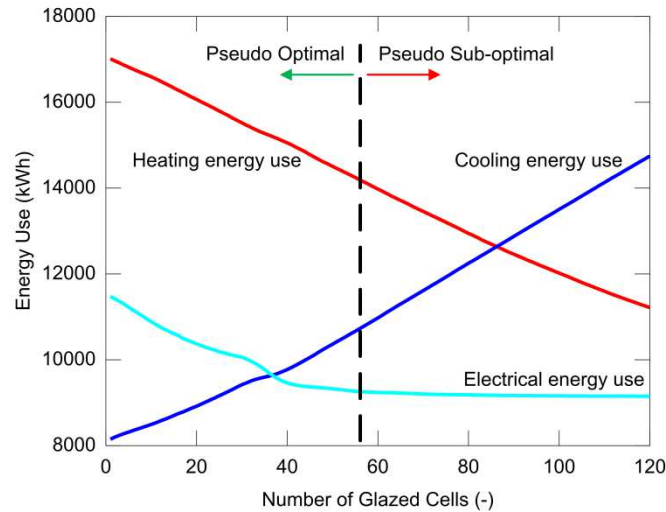
Figure 3. Components of energy use
(electrical energy includes artificial lighting)

6. Optimization experiments and methods of analysis

The optimum solutions to building optimization problems are unknown and are subject to uncertainty in the problem boundary conditions (Hopfe and Hensen, 2011). EA are also probabilistic optimizers that have uncertain convergence. The uncertainty in the problem boundary conditions (such as that due to changing climate and occupant behaviour) is not addressed in this paper, and as such the optimization presented here is for a single set of uncertain boundary conditions. The true optimum solutions for a single set of boundary conditions can only be determined using an evaluation of all solutions in the search space (a so called "exhaustive" or "brute-force" search). Given that for the integer encoding the search space contains $3^{120} = 1.797 \times 10^{57}$ possible solutions, it is impracticable to conduct an exhaustive search and therefore the true optima remain unknown. The optimality of the solutions must therefore be examined by other means. In this paper, we explore the optimality of the solutions through; the comparison of the solutions to those found from the traverse of the search space (section 5.4); a subjective assessment of whether the solutions appear to concur with building physics; and finally, and most important, through a statistically significant number of repeat runs of the algorithm for each optimization problem. An analysis based on a statistically significant number of repeat runs of the algorithm also addresses the uncertainty due to the probabilistic convergence of NSGA-II. For the unconstrained optimization problem, further confidence in the extent to which the solutions may be optimal has been gained through a local and global sensitivity analysis.

Eight different sets of experiments have been performed, the different experiments resulting from the alternative combinations of problem formulation (Table 2). Each of the 8 experiments and the associated analysis is based on 32 different runs of the algorithm. This results in 128 different optimization results for each of the unconstrained and constrained problems.

The analysis of the characteristics of the optimization problem and the behaviour of NSGA-II in solving the problem has been performed using both graphical (qualitative) and quantitative methods. The quantitative methods are used to indicate the relative optimality of the solutions obtained from a particular algorithm configuration (the "optimality" being judged in relation to other solutions found in the experiments), whereas the qualitative analysis has been used to gain an understanding of the underlying characteristics of solutions.

6.1 Quantitative methods of analysis

The quality of solutions to a multi-objective optimization problem can be judged in relation to: the convergence of the solutions onto the true Pareto front; the distribution of the solutions along the Pareto front; and the range of objective function values found between the extremes of the Pareto front. The convergence to the true Pareto front and the spread of the solutions is commonly measured by the hypervolume metric (Zitzler, 1999). Although this metric measures both the position of the front and the spread of the solutions, it is known to provide a good balance between measuring the different traits of the solutions (Knowles and Corne, 2002). Given that for this building optimization problem the true Pareto front is unknown, the hypervolume of the multi-dimensional space has been measured between the solutions of a Pareto front and a specific reference (*nadir*) point that has been taken to be equal to the maximum value of each objective function value found in all experiments. This ensures that the nadir solution equates to the least optimal solution found in any of the experiments and therefore that it is dominated by all solutions.

Figure 4 illustrates the hypervolume as the shaded area between the points of the Pareto set of solutions and the nadir reference point. Note that solutions having a higher hypervolume are desirable, as these solutions have Pareto fronts that are either further from the nadir point or have more spread and a greater range of objective function values, both of which are desirable. Note also that given our definition of the nadir point, it can lie beyond the range of objective function values found in the final Pareto front.
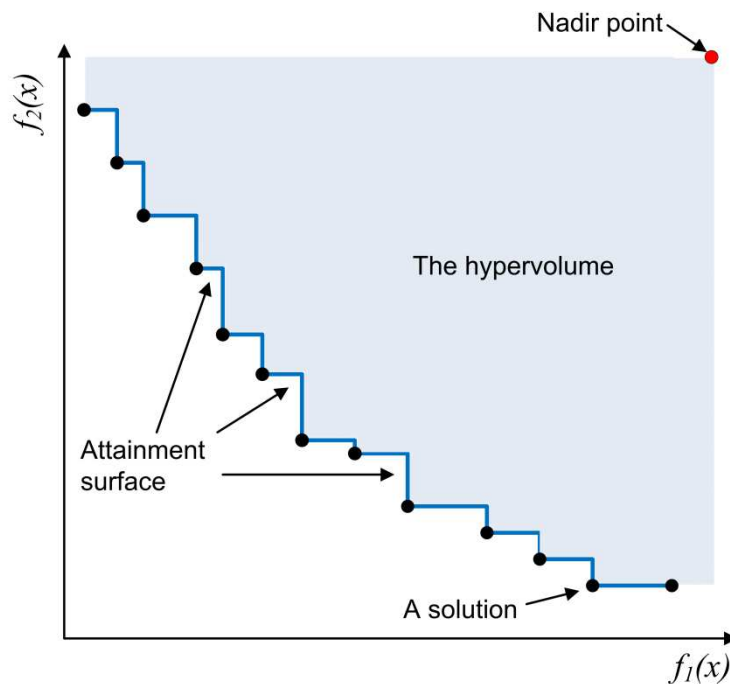

Figure 4. Attainment curve and hypervolume in a bi-objective space, with both objectives being minimized

Given that the hypervolume is sensitive to the spread of the solutions as well as the distance of the solutions from the nadir point, two further metrics have been used in this study to examine the spread of the solutions. First, the minimum values for both objectives are examined, these lying at opposite ends of the Pareto front (the minimum energy use corresponds to the maximum cost, and vice versa).

Second, the distribution of the solutions along the Pareto front is evaluated using the spread metric (Deb, 2001). The spread of solutions along the Pareto front is important in design practice as the higher the spread the greater the choice of design solutions. Spread is a measure of the variation in distances between neighbouring solutions in the front with values closer to zero being desirable (zero being the condition that all distances are equal). In (Deb, 2001), spread, Δ, is formally defined as:

$$\Delta = \frac{d_f + d_l + \sum_{i=1}^{N-1}|d_i - \bar{d}|}{d_f + d_l + (N-1)\bar{d}} \tag{9}$$

where, the parameters d_f and d_l are the Euclidean distances between the solutions at each end of the obtained non-dominated set, and the limits for the objective values. The parameter $\bar{d}$ is the average of all distances d_i, $i=1,2,3...,N-1$, assuming that there are N solutions on the best non-dominated front. The spread metric also requires the limits for the objective values to be supplied and again these were taken as the minimum and maximum found over all runs of all the algorithms.

The hypervolume, spread, and minimum objective values are reported in terms of the mean and standard deviations of the results for the 32 repeat optimizations. The ability of the algorithm to find a feasible solution is also examined and reported as a percentage success rate.

6.2 Qualitative methods of analysis

Graphical (qualitative) analysis is used to describe the shape and optimality of the Pareto front, and to investigate the difference in distribution of the solutions along the Pareto front resulting from the different experiments.

The trade-off between two objectives and the shape of the Pareto front is commonly represented as a continuous line between solution points. However, this can be misleading as there is no guarantee that the any solutions exist between the solutions found by the search, or that the trade-off is smooth. An attainment surface (Fonseca and Fleming, 1996), addresses these difficulties by presenting the trade-off as a surface that divides the objective space into the region dominated by a Pareto set, and the region not dominated by it. This is easier to interpret than simply plotting the individual points for multiple fronts. For a two-objective space, the attainment surface becomes an attainment curve. Figure 4 illustrates an attainment curve as the blue line connecting the set of Pareto-optimal points represented by black dots.

The hypervolume, spread metric, and the attainment surface are all representations of the objective function space. The trend in the distribution of solutions along the Pareto front in the variables space has also been analysed by identifying the probability that each cell on the façade will be glazed, and presenting these results graphically. As well as enabling an analysis of differences in the solution trends between the different experiments, these results indicate the general characteristics of the solutions in that they illustrate the region of the façade that has highest probability of being glazed (regardless of the particular solution adopted along the trade-off between energy use and capital cost).

6.3 Sensitivity analysis

To enhance the understanding of the problem characteristics and the extent to which the solutions may be optimal a local and global sensitivity analysis has been performed for the unconstrained optimization problem. The local (differential) sensitivity analysis has been performed about three solutions taken from the best of the unconstrained optimization results, the highest energy use solution, the median energy use solution,

and the lowest energy use solution. Using the given solution as a base-case, the state of each cell is changed in turn (from either opaque to glazed, or glazed to opaque), and the difference in energy use evaluated; the state of the cell is then reset before the next cell on the façade is changed. Since the lowest and median energy use solutions have glazed cells that are shaded, a change in the glazed state of the cell also changes the state of the overhang (the overhang being present if the cell is glazed). No shading is applied in the case of the highest energy use solution as the base-case is un-shaded. Note that since the cells have only two states (glazed or opaque), the local sensitivity analysis is also in effect a local search, the results from which can be used to confirm the local optimality of the solution.

The local sensitivity analysis has been supplemented with a global sensitivity analysis. 3,000 randomly initialised solutions were evaluated with each cell having a 50% probability of being glazed in each solution. Each glazed cell was automatically shaded. The global sensitivity analysis is based on stepwise linear regression analysis of rank transformed variable values. The model was identified using a bi-directional selection and the Bayesian information criterion (this criterion reducing the risk of over-fitting).

It is not necessary to evaluate the sensitivity of the capital cost objective function in the same way as for the energy use since the sensitivity of the capital cost is known to be uniform across the façade and linear across the search space (Equation 4). That is, it is known that all cells are of equal importance in terms of the capital cost.

7. Results and analysis

Table 2 gives the solution metrics for the different optimization experiments. The values for the hypervolume, spread, minimum energy use and cost are given as the means of the solutions with the standard deviation of the solutions in brackets. For the experiments in which the aspect ratio constraints where applied, the means and standard deviations are calculated for the feasible solutions only, infeasible solutions being excluded. For reference, the maximum values of the objective functions found in all experiments which are used to represent the nadir solution are: 36,637kWh and £31,374 for the unconstrained problem; and 35,787kWh and £29,050 for the constrained problem.

The extent to which the results from the 32 repeat optimization runs for each experiment are normally distributed (and therefore the extent to which some confidence can be placed in the analysis), was confirmed by applying a Shapiro-Wilk test to each data set. In each case, the test resulted in a p-value of > 0.05, confirming that there was no significant departure from normality in the results. Similarly, in a two-tailed unpaired t-test on the results of the seeded versus the unseeded results all had a p-value of < 0.0001 and so the differences in the results can be considered to be statistically very significant.

Experiment			Feasibility Success Rate (%)	Hyper-volume (-)	Spread (-)	Minimum Energy Use (kWh)	Minimum Cost (£)
Constrained?	Encoding	Seeded?					
no	bit-string	no	-	0.754 (0.022)	0.700 (0.047)	34,262 (16)	15,503 (744)
		yes	-	0.798 (0.005)	0.663 (0.005)	34,286 (25)	**13,596** (155)
	integer	no	-	0.281 (0.063)	1.032 (0.026)	34,265 (26)	25,811 (1162)
		yes	-	**0.819** (0.002)	**0.580** (0.038)	**34,154** (8)	13,939 (130)
yes	bit-string	no	43.8	0.329 (0.072)	0.850 (0.048)	35,419 (94)	22,178 (1066)
		yes	100	0.681 (0.021)	0.668 (0.040)	34,450 (55)	14,158 (23)
	integer	no	0	-	-	-	-
		yes	100	**0.705** (0.029)	**0.653** (0.079)	**34,342** (109)	**14,154** (0)

Table 2. Experiments and solution metrics
(the values are given as mean with standard deviations in brackets;
the best solutions in each category are shown in bold)

7.1 Unconstrained optimization

It can be observed from Table 2 that for the unconstrained and unseeded optimization, the integer encoding of the variables results in a mean hypervolume, spread, and minimum capital cost that all have worse values than for the binary encoding. However, the use of solution seeding results in the integer coding having the best solutions as measured by all metrics except the minimum capital cost. Seeding also resulted in an improvement to the results for the bit-string encoding, but to a lesser extent than the integer encoding. In comparison to the solutions found from the traverse of window cells (section 5.4), the lowest energy use solution from the seeded integer encoding has a marginally (<0.1%) lower energy use, but with this energy use being achieved with a 14.2% lower capital cost, this suggesting that the optimization has found a significantly more optimal solution than that given by the traverse of window cells. Since the highest energy use solution from the traverse is unglazed, and the solution obtained from the optimization has only one glazed cell (Figure 6), the difference in the optimized objective function values and the values from the traverse are small (the optimized energy use being 1.1% lower than that from the traverse and the capital cost 1.8% higher).

The improvement in solutions from both the bit-string and integer encodings with seeding can be attributed to the seeded initial population having a greater variety of solutions than the randomly initialised population. The seeded initial population contains solutions ranging from 2.5% glazed to 83% glazed, whereas the randomly initialized population has a bias towards solutions that are 50% glazed for the bit-string encoding, and 66.7% glazed for the integer encoding.

The impact of the seeding is also evident from the attainment curves. Figure 5, illustrates the attainment curve for the most successful and unsuccessful variable encodings for the unconstrained optimization (these being the seeded and unseeded

integer encodings). As confirmed by the standard deviation of the hypervolume and spread (Table 2), there is very little difference across the solutions found for the seeded initial population. However, this is not the case for the unseeded integer encoding, there being a range of convergence and hypervolume across the unseeded solutions.
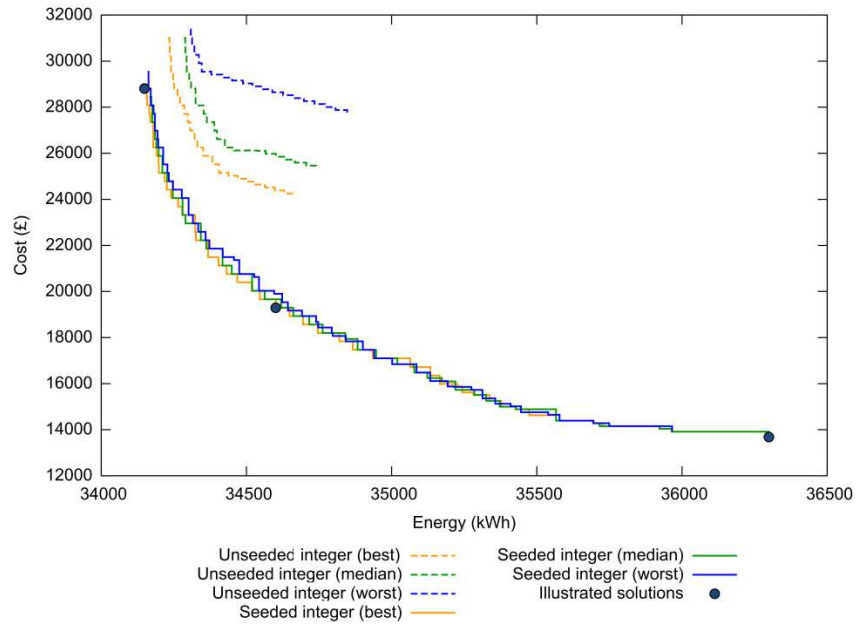

Figure 5. Attainment curve for the unconstrained optimization
(the most and least successful variable encodings)

Figure 6, illustrates three solutions taken from the best of the seeded integer solutions, the solutions being for the lowest capital cost, lowest energy use, and median energy use and cost (Figure 5). It is perhaps unsurprising that the median and lowest energy solutions have glazed cells positioned towards the top of the façade as this increases daylight penetration onto the artificial lighting control points and therefore tends to limit electrical energy use in artificial lighting. The window(s) have a wide but short form that tends to cover the width of the façade, with all glazed cells for the median and lowest energy solution having overhangs. The solution for the lowest capital cost is marginally sub-optimal as it is known that the minimum cost solution corresponds to a fully opaque wall construction.

(a)
Lowest capital cost;
highest energy use

(b)
Median capital cost
and energy use

(c)
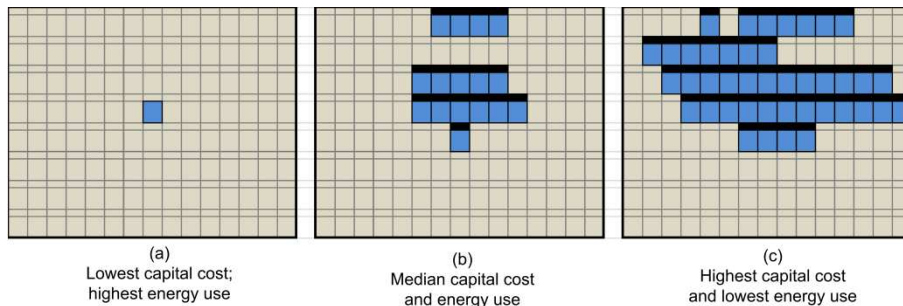Highest capital cost
and lowest energy use

Figure 6. Selected solutions from the unconstrained optimization
(overhangs are illustrated as solid black blocks above the blue window cells)

The variation of the solutions across the Pareto set is further examined in Figure 7, which illustrates the probability that a particular cell is glazed for any solution in the

Pareto set. Solutions for both of the bit-string encoding and the seeded integer encoding have a wide variation in the glazed cells among the Pareto set (ranging from cells that are always glazed, through to cells that are never glazed). In all three cases, the most frequent placement is towards the centre of the façade but with the some probability that the cells towards the top of the façade are also glazed; this is explored further in the local sensitivity analysis. The probabilities for the unseeded integer encoding show far less variety among the Pareto set. The population has converged on a highly similar set of solutions, there being a high probability of one large tall window and a number of single-celled or other small windows being present in all solutions along the trade-off. The reason for this is unclear, but is most likely related to the fact that the integer encoding results in a 66.7% chance of a cell being glazed, particularly in the initial population; the seeded integer encoding results in a greater variety of solutions due to the wider range of solutions in the initial population.

The analysis in Figure 7 is useful in design as it indicates the general characteristics of the optimised solutions, in that a designer seeking more freedom of design choice will understand that the central cells should always be glazed, and that in order to maintain a degree of optimality, the designer should place the majority of glazed cells in the upper half of the façade.
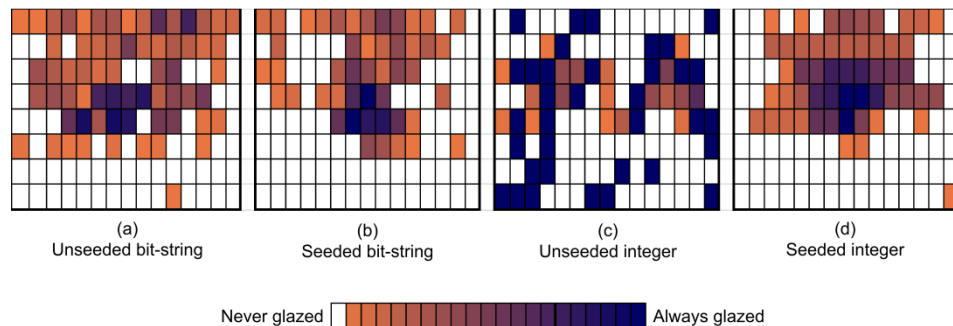

| (a) | (b) | (c) | (d) |
| Unseeded bit-string | Seeded bit-string | Unseeded integer | Seeded integer |

Never glazed ▭▭▭▭▭▭▭▭▭▭▭▭▭ Always glazed

Figure 7. Probability of a cell being glazed for any solution in the Pareto set (unconstrained optimization)

7.1.1 Sensitivity Analysis

Figure's 8 (a) to (c) illustrate the results of the local sensitivity analysis and Figure 8 (d) the global sensitivity analysis. The local sensitivity figures are shaded according to the absolute value of the differential sensitivity from the base case. The arrows indicate the direction of the change in energy use resulting from the change in state of the cell; a downward pointing arrow indicates that the change of state resulted in reduction in energy use and an upward pointing arrow an increase in energy use. The base case window placements are indicated with a dashed-line, these being the same, low energy, median energy, and high energy solutions as illustrated in Figures 6 (a) to (c)
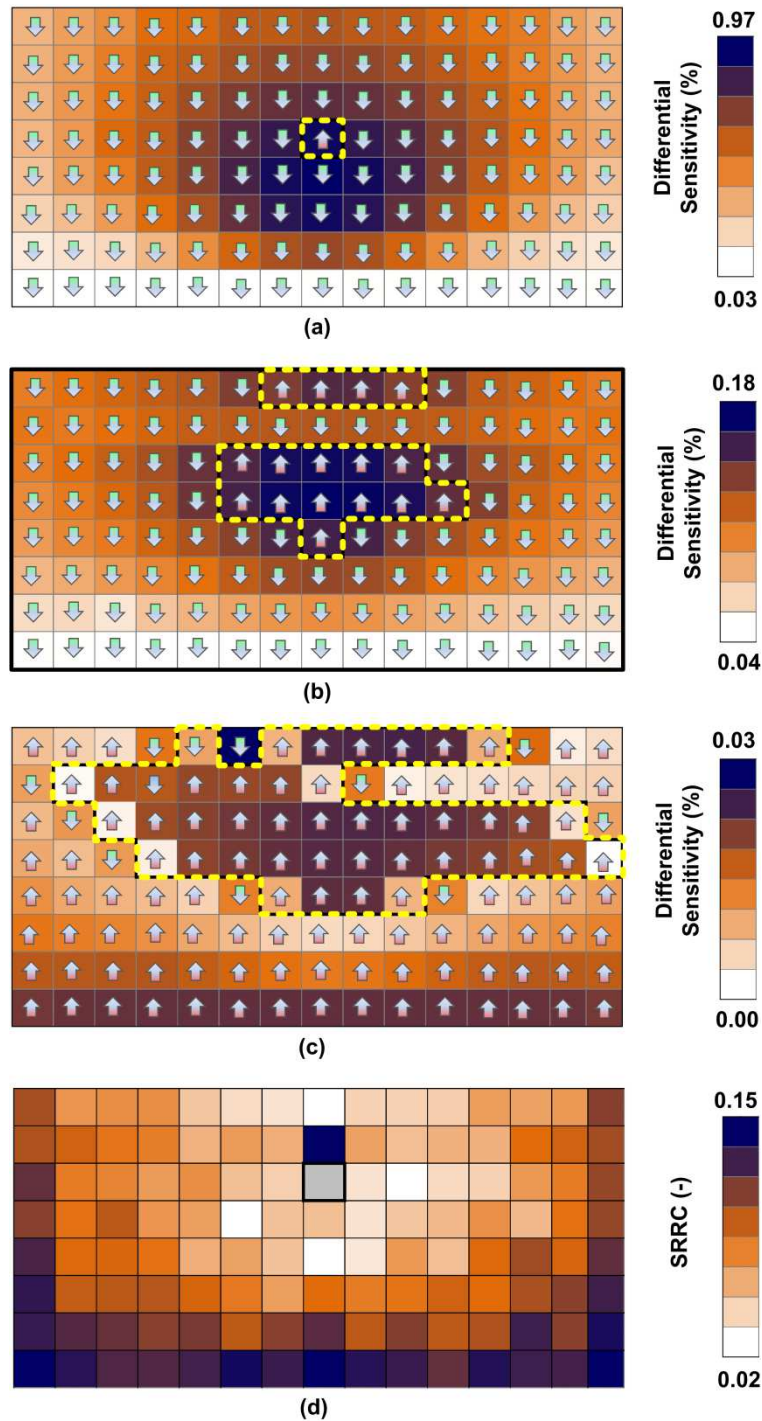
Figure 8. Energy use local and global sensitivity analysis (unconstrained optimization; (a) maximum energy use solution local sensitivity; (b) median energy use solution local sensitivity; (c) minimum energy use solution local sensitivity; (d) global sensitivity analysis)

Figure 8 (a) illustrates the differential sensitivity for the highest energy use solution. Given that it is known that the lowest energy use requires a higher number of windows, adding a window to every opaque cell results in a reduction in energy use, while removing the base-case glazed cell increases energy use. This characteristic is repeated for the median energy use solution illustrated in Figure 8 (b), as this solution also has a glazed area that is lower than that required for minimum energy use.

However, since the solution for the minimum energy use (Figure 8 (c) has a glazed area that is optimized for minimum energy use, changing the state of an opaque cell to glazed in general results in an increase, rather than decrease, in energy use. This is not the case for all opaque cells as changing the state of those on the border of the window results in a reduction in energy use. For all three solutions, the sensitivity of the capital cost is known to be uniform across the façade, with a change of an opaque cell to glazed resulting in an increase in cost (and vice versa).

A comparison of Figures 8 (a) to (c) indicates that as the glazed area increase (from Figure 8 (a) to (c)), the sensitivity of the solution to a change in the state of a cell is reduced from a maximum of 0.97% to 0.03%. This occurs as the differential sensitivity represents the local gradient of the search space, the gradient being reduced as the minimum energy use is approached (the change in gradient is evident in Figure 2). It is also evident that the centre of the sensitivity shifts upwards as the glazed area increases (from Figure 8 (a) to (c)). The reason for this requires further investigation, but is likely to relate to the fact that single glazed cell positioned at the top of the façade is unlikely to provide sufficient daylight penetration onto the artificial lighting control sensor that is positioned closest to the façade.

The absolute sensitivities for the capital cost are 1.74% for the maximum energy use solution, 1.90% for the median energy use solution, and 1.27% for the minimum energy use solution. In all cases the sensitivity is uniform across the façade with a change of a cell from opaque to glazed resulting in an increase in cost; correspondingly, replacement of a glazed cell with an opaque construction results in a reduction in cost. As well as being uniform across façade for a given solution, the change in cost is known to be linear across the entire search space (Equation 4). The change in percentage sensitivities between the different solutions given here is a result of the cost of the base case solution increasing with the size of the window (the sensitivity being given as a percentage change against the base case cost). The magnitude of the energy use for the base case solutions also affects the energy use sensitivity, but by a small amount (these sensitivities being driven largely by the local gradient for the change in energy use). The cost sensitivities also differ between the solutions since the median and minimum energy use sensitivities include the effect of the overhangs, whereas the maximum energy use solution does not. Finally, note that the magnitude of the sensitivies is small due to the fact that each cell occupies only 0.83% of the façade, so it might be expected that a change in the state of the cell would have a small effect.

The procedure for performing the local sensitivity analysis is equivalent to that of a local search of the solutions around the base case designs. Each sensitivity analysis results in 120 new solutions (the state of all 120 cells being changed in the procedure). These results can be used to examine the optimality of the base case solutions by the extent to which they dominate, or are dominated by, the 120 new solutions (the domination being examined for both the energy use and capital cost objective functions). The maximum energy (minimum cost) solution is sub-optimal and dominated by just one of the 120 solutions (and therefore by a single glazed cell); the new non-dominated solution has a fully opaque façade. The median energy use solution dominates all 120 new solutions and so it can therefore be concluded that this solution is locally optimal. Two new non-dominated solutions were found for the minimum energy use base case, one for an additional glazed cell, and a second for the removal of a single glazed cell from the base-case window; the base case solution was dominated by both of the new solutions. It can therefore be concluded that the optimization results in locally near Pareto optimal solutions, with the median energy use solution in this case being locally fully optimal.

Figure 8 (d), illustrates the global sensitivity analysis, with the shading indicating the magnitude of the standard rank regression coefficients (SRRC); no SRRC is available for the boxed cell located in the middle of the façade as this cell was automatically removed from the regression model. All coefficients have a positive value which indicates that the use of a glazed cell rather than an opaque cell results in an increase in energy use. The sensitivity as represented by the SRRC's increases towards the bottom of the façade. This is a result of each solution used in the analysis having an approximately 50% glazed area (since each cell has a 50% chance of being glazed or unglazed). Given that the minimum energy use solution (Figure 8 (c)), has a glazed area of 35%, a 50% glazed façade is bound to result in a higher energy use. Although unconfirmed, the effect is likely to be greatest when glazed cells are positioned at the bottom of the façade, as these are likely to have the lowest impact on energy use due to artificial lighting. This effect is also apparent from the local sensitivity analysis for the minimum energy use, as in comparison to the median and maximum energy use solutions, the bottom row of cells has an increased sensitivity. Given the premise that placing cells towards the bottom of the facade results in the highest increase in energy use, the global sensitivity analysis confirms that the optimization is placing the windows in a generally optimal location of towards the middle to upper region of the façade.

7.2 Constrained optimization

Table 2, indicates that the seeding of the initial population with a range of feasible solutions had a significant impact on the optimization results for the constrained optimization (in which the aspect ratio is constrained to be approximately equal to the Golden Ratio). In particular, no feasible solutions where found for the integer encoding for the randomly initialised population (the unseeded initial population), and without seeding, only 43.8% of the solutions for the bit-string encoding where feasible, feasibility taking between 20 and 100 generations to be achieved and with no improvement in feasibility being found beyond 100 generations. In contrast, all solutions for both the bit-string and integer encodings where feasible when the initial population was seeded. The success of the seeding is in-part due to the seeded solutions of the initial population being feasible (feasible meaning that all seeds satisfy the constraints). This eliminates the need for the search to find feasible solutions and allows more of the search effort to be focused on minimising the objective function values. The seeded integer encoding results in the best performance as measured by all metrics (hypervolume, spread, minimum energy use, and minimum cost). However, no solutions had an energy use or capital cost lower than that found by the incremental traverse. This is unsurprising in that the traverse results in window geometries that violate the aspect ratio constraints, these constraints resulting in both a higher minimum cost (in this case, a small window rather than no window), and higher energy use (due to constraints producing tall-thin windows rather than short-wide windows).

Figure 9, illustrates the attainment curves for the most and least successful algorithm configurations that resulted in any feasible solutions (these being the seeded integer encoding and the unseeded bit-string). The trade-offs found for the seeded integer-encoding are more consistent with the median and best fronts these covering a similar range in both of the objectives. The worst solution is convergent on the best solution, but is restricted to the low cost designs. In contrast, the unseeded and bit-string-encoded solutions show a greater variation, both in extent and in closeness to the origin (this is reflected in Table 2 by a higher standard deviation for the hypervolume).

This illustrates the usefulness of the hypervolume metric in selecting a good Pareto front from a number alternative optimization runs.
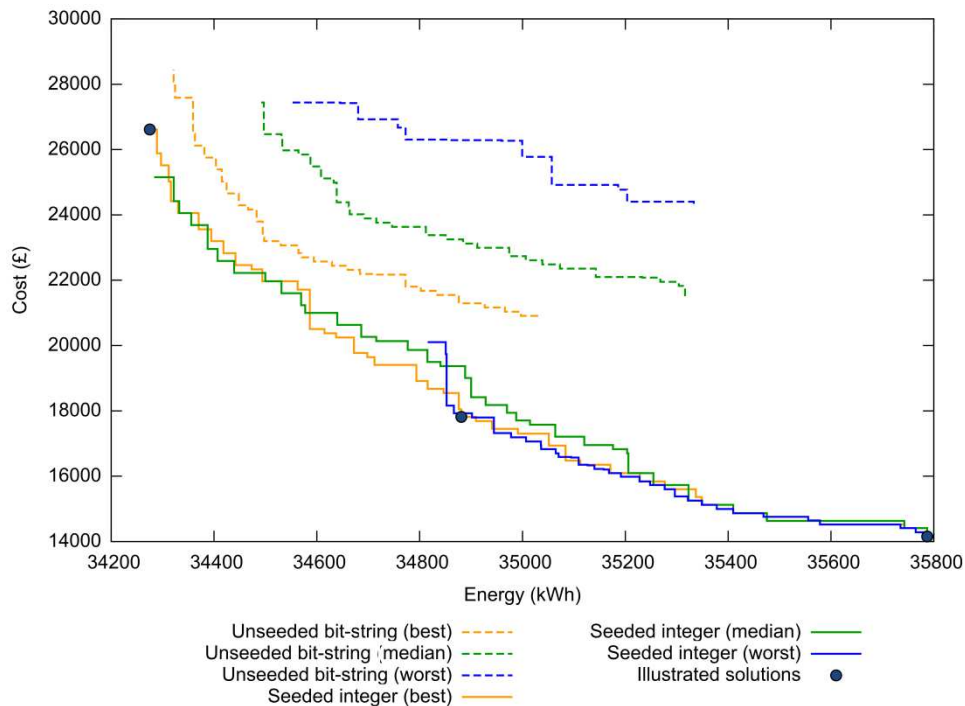

Figure 9. Attainment curves for the constrained optimization
(for the most and least successful variable encodings)

Figure 10, illustrates three solutions taken from the best Pareto front, the minimum cost solution, the minimum energy use solution, and the median cost and energy use solution (as indicated in Figure 9). The figure illustrates that as the glazed area increases, the aspect ratio constraint results in multiple windows and that the window cells also have a greater tendency to be shaded by overhangs (all window cells being shaded for the highest cost solution). This is driven by the trade-off in energy use, in that as the cost increases with glazed area, the energy use must reduce (the shading contributing to the reduction in energy use).

As is the case for the unconstrained solutions, for the higher glazed solutions, the window cells tend to be located towards the top of the façade in order to capitalize on daylight penetration. Figure 10, also indicates that for the high cost solution there is a slight bias in the cell positions towards the top-right side of the façade (this being the top-west side as viewed from inside the building). This is similar to the solutions found in a single objective optimization study in which the bias was thought to be a result of the distribution of the solar beam radiation onto the internal surfaces of the room (Wright and Mourshed, 2009). This has not been investigated further here, as conversely, it was also noted that across the different repeat optimization runs, although the optimum energy use and cost found were often similar, the distribution of individual glazed cells varied across the façade. This implies that position of the cells is of secondary importance to the number of glazed cells and the general areas of the façade that is glazed.
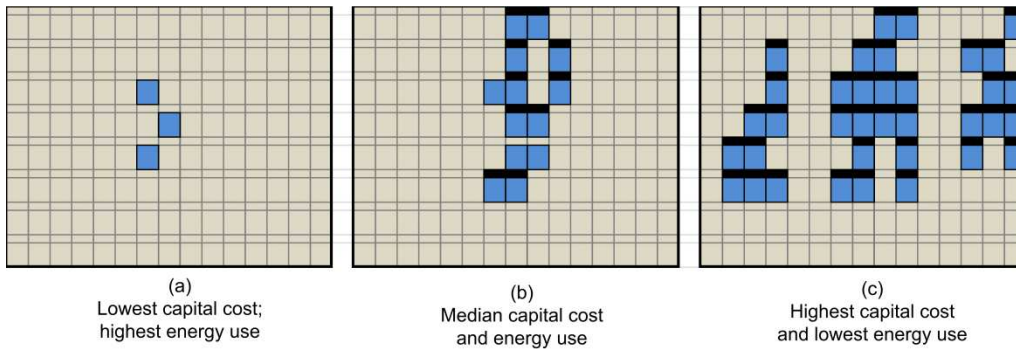
Figure 10. Selected solutions from the constrained optimization
(overhangs are illustrated as solid black blocks above the blue window cells)

Figure 11, illustrates the probability that a cell is glazed for any one solution on the best Pareto front found for the alternative variable encodings and use of seeding. The solutions for the unseeded integer encoding have been omitted since no feasible solutions were found for this configuration.

The figure for the unseeded bit-string encoding illustrates that there is a poor distribution of probabilities, this indicating that regardless of the position of the solution on the Pareto front, the same cells tend to be glazed. Further analysis of the individual optimization runs indicated that once the algorithm has found one solution that met the aspect ratio constraint, that solution dominated the population, and therefore limited variation in solutions amongst the Pareto set. In contrast, there is a greater variation in probabilities and therefore diversity of solutions for both of the seeded optimization results. As for the unconstrained solutions, the probabilities are highest towards the centre of the façade.
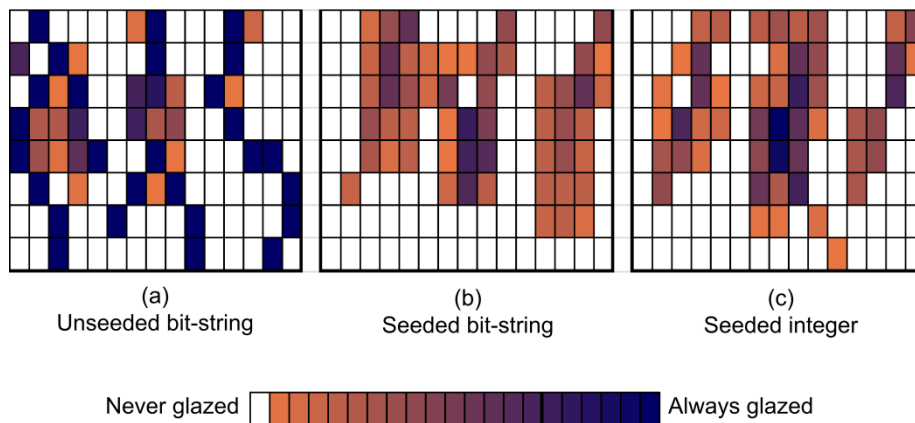

Figure 11. Probability of a cell being glazed for any solution in the Pareto set
(constrained optimization)

7.3 The impact of the aspect ratio constraints

Figure 12, compares the attainment curves for the best solutions found from the unconstrained, and constrained optimization experiments. The difference in the results is most noticeable as the glazed area (and capital cost) increases, the solutions being convergent for the low glazed areas (and low cost). The reason for the divergence is that the constraints have increasing impact on the shape of the windows as the glazed area increases. The high cost unconstrained windows have a short-wide aspect ratio (Figure 6(c)), whereas the high cost constrained windows have a tall-thin aspect ratio (Figure

10(c)). The convergence of the solution as the capital cost reduces is due to a reduction in the sensitivity of the energy use to the number of glazed cells. An increase from 1 cell (Figure 6(a)) in the unconstrained case, to three cells in the constrained optimization (Figure 10(a)) having a limited impact on capital cost.
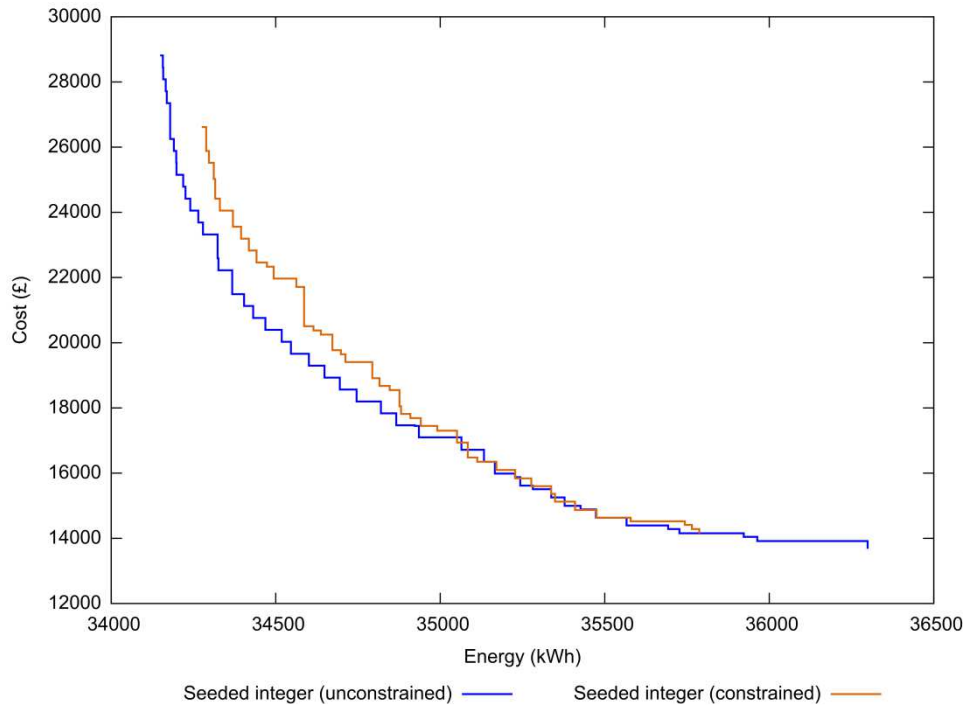

Figure 12. Best attainment surfaces from the unconstrained and constrained optimization

7.4 Seeded Solutions in the Final Population

Given that for this problem, seeding has such a strong impact on the success of the optimization, the extent to which the seeds survived through the search and dominate the final population has been investigated by the distance (in the variable space), between any solution and the seed solutions. The analysis is based on the bit-string encoded solutions, as the distance between solutions can be easily measured using the Hamming distance (this being the number of bits that are different in comparing a solution from the search, and any one of the seed solutions). The Hamming distance was only computed between the bits that represent the glazing, as the shading overhangs were all omitted in the seeds. The smallest Hamming distance (that is, the distance from a given solution to the most similar seed), was recorded and aggregated over all trial runs of the algorithm.

Table 3, gives: the mean Hamming distance between each member of the Pareto set and the seed most similar to it; the number of solutions in the Pareto set in which the glazing pattern exactly matched a seed solution (ignoring shading); and the number of unique seeds present in each Pareto set.

Seeded solutions persisted most in the constrained optimization, with a 15% mean number of seeds remaining in the Pareto front. However, this equates to a mean of only 1.5 seeds. The solutions along the Pareto front only being different to the seeded solutions by an average of 8.6 bits (or 7% of the length of the bit-string).

Survival of the seeded solutions into the final Pareto front is lower for the unconstrained optimization with a mean of less than one seed surviving, and less than 1% of the final solutions being seeds. The distance between the solutions and the seed solution is also higher at 10.3 bit (9%) for the unconstrained solutions. The reason for the reduction in the number of surviving seeds is that the seeds have geometries that satisfy the constraints on the aspect ratio, whereas the solutions for the unconstrained aspect ratio are very different in the position and geometry of the windows (for instance, compare the solutions illustrated in Figure 6 and Figure 10). It might therefore be concluded that although the final Pareto set of solutions for the constrained optimization contained a high number of seeds, the results for the unconstrained problem suggest that the search is inherently able to benefit from seeding without being dominated by the seeding.

Problem Form	Hamming distance to closest seed (-)	Number of seeds in Pareto set		Number of unique seeds in Pareto set (-)
		(-)	(%)	
Constrained	8.6 (1.2)	7.6 (2.7)	15.0	1.5 (0.5)
Unconstrained	10.3 (1.1)	0.28 (0.73)	0.56	0.16 (0.37)

Table 3. Influence of seeding on the final Pareto set
(the values are given as the mean with standard deviations in brackets)

8.0 Conclusions

This paper investigates the multi-objective optimized design of fenestration that is based on the façade of the building being divided into a number of small regularly spaced cells, this formulation providing the search with the flexibility to optimize the glazed area, number of windows, and window shape (a window on the façade being defined by a set of adjoining glazed cells). Dividing the façade into a number of cells is similar in concept to that of a digital image in which the image is constructed from a number of different pixels. In the same manner as the resolution of a digital image being dependent on the number of pixels, the clarity of the design solutions, particularly the window shape, is dependent on the number of cells on the façade. The façade in the example optimization presented in the paper is divided into 120 cells, with each cell occupying less than 1% of the façade area. It was concluded that this should provide sufficient clarity as to the optimized position, number, and shape of the windows, without the search space being excessively large (the size of the search space increases in the order of 3^n, where n is the number of cells).

The objective functions minimized by the optimization in the paper are; the energy use associated with heating, cooling and artificial lighting; and the capital cost of the façade construction. The paper also investigates the use of constraints to control the aspect ratio of the windows (the use of this constraint also has an indirect impact on the number of windows).The optimization algorithm used to solve the optimization problem is an implementation of the non-dominated sorting genetic algorithm (NSGA-II). The energy performance of a candidate design found during the optimization was evaluated using the EnergyPlus performance simulation with the example building being located in in Chicago USA.

The paper investigates the impact of two alternative forms of variable encoding (bit-string and integer) and the use of seeding the initial population of solutions. Given the probabilistic behaviour of the genetic algorithm, the impact of a particular variable encoding and seeding configuration was evaluated through 32 repeat trial optimizations, this enabling the results to be reported using statistical means and standard deviations for hypervolume, solution spread, and the minimum energy use and capital cost.

Both the bit-string and integer encodings were able to find solutions for the unconstrained optimization problem, although the results for the integer encoding were noticeably worse when the initial population consisted entirely of randomly generated solutions. Seeding of the initial population with a range of possible solutions improved the optimality of the unconstrained results obtained for both the bit-string and integer encodings, this improvement being most significant in the case of the integer encoding. The use of constraints to control the aspect ratio of the windows resulted in failure to find any feasible solutions for the randomly initialised search and integer encoding. However, 43% of the of the randomly initialised optimization trials for the bit-string encoding resulted in feasible solutions. Seeding the initial population with feasible solutions resulted in all trail constrained optimizations having feasible solutions for both the bit-string and integer encodings. The extent to which the final solutions were dominated by the seeds was also investigated. The final solutions to the constrained optimization contained a small number of seed solutions, even fewer seeds being present in the unconstrained solutions. It was therefore inferred that the search was inherently able to find alternatives to the seed solutions and that the seeding enabled the search to spend more time in improving the optimality of the energy use and cost objective functions, and less time on finding feasible solutions.

An analysis of solutions obtained from the best of the unconstrained optimization confirmed that the search provided a set of alternative design solutions that ranged from an almost unglazed façade to a 35% glazed façade. The median and higher cost (most glazed), and lower energy use solution had windows positioned towards the top of the façade, and all had overhangs for shading. The lower cost (least glazed), solutions had no overhangs and were positioned centrally on the façade.

The optimality of the unconstrained solutions was tested by comparing them to the results from a binary traverse of the façade, a local search and sensitivity analysis, a global sensitivity analysis, and a subjective assessment of the probable window placement informed through building physics. In particular, through the local search and sensitivity analysis, it was possible to confirm that a solution taken from the middle of the trade-off curve dominated the solutions from the local search (in terms of both energy use and capital cost), but that the solutions at either extreme of the trade-off where dominated by one or two solutions found from the local search. It was therefore concluded that the optimization results were locally near Pareto optimal.

The range of solutions found for the constrained optimization problem and the general position of the windows and the use of shading, were, in general, similar to the solution characteristics of the unconstrained optimization problem. However, in comparison to the unconstrained solutions, constraining the aspect ratio resulted in an increase energy use for the higher glazed solutions, the difference in both cost and energy use diminishing for the lower glazed areas. Although the extent to which the solutions to the constrained optimization are locally optimal has not been explored using a local search and sensitivity analysis, the very low standard deviations for all quantitative metrics applied to the statistically significant repeat optimization trails suggests that the search consistently converged to at least a local Pareto optimal front.

It can be concluded, that a cellular façade, with the state of each cell (glazed or solid) being optimized using NSGA-II, is not only effective in the minimization of building energy use and capital cost, but that it can also result in design solutions that are atypical to common design practice and as such have innovative architectural forms. Using an integer encoding of the variables provided the best solutions, but only when the initial population was seeded with feasible solutions. Although in comparison to the solutions for the integer encoding the results for the bit-string encoding were sub-optimal, the bit-string encoding is robust in that it was able to find solutions to the constrained problem without seeding.

The multi-objective approach resulted in at least 30 unique solutions for both of the unconstrained and constrained optimizations, with all solutions lying on the trade-off between energy use and capital cost. Selection of a particular solution for construction can be achieved through inspection of the solutions and for example, the manner in which the energy use changes with capital cost. Given that in this example, each cell occupied less than 1% of the façade, the sensitivities of the energy use and capital cost to minor changes in the solutions are low. This allows the architect to modify the optimized design without there being a significant reduction in optimality of the solutions. For instance, say by exaggerating the rectangular form of the window or alternatively by giving the window edges a curvilinear rather than rectangular form.

Further research is required to investigate the application of the approach to architectural design, particularly the use of different constraints to control the window position and geometry (previous work by Wright and Mourshed (2009) described window density and window count metrics that could be used to produce alternative window forms). Given the success of the seeding approach demonstrated in this paper, it is reasonable to suggest that it would be equally successful in solving others constrained cellularwindow optimization problems. Further work is also necessary to extend the approach to integrated "whole-building" building design in which other elements of the building are optimized simultaneously with the fenestration.

References

BROWNLEE, A.E.I., WRIGHT, J.A. and MOURSHED, M.M., 2011. A multi-objective window optimisation problem, In: *Proceedings of the 13th annual conference companion on Genetic and evolutionary computation* 2011, Dublin. New York: ACM Press, pp. 89-90.

BROWNLEE, A.E.I., and WRIGHT, J.A., 2012. Solution Analysis in Multi-Objective Optimization, In: *"Proceedings of the 2012 Building Simulation and Optimization Conference";* Jonathan Wright and Malcolm Cook, (Editors); 10-11 September 2012, Loughborough, Leicestershire, UK; pp 317-324; ISBN 978-1-897911-42-6.

CALDAS, L., 2008. Generation of energy-efficient architecture solutions applying GENE_ARCH: An evolution-based generative design system. *Advanced Engineering Informatics,* **22**(1), pp. 59-70.

CALDAS, L.G. and NORFORD, L.K., 2002. A design optimization tool based on a genetic algorithm. *Automation in Construction,* **11**(2), pp. 173-184.

CHANTRELLE, F.P., LAHMIDI, H., KEILHOLZ, W., MANKIBI, M.E. and MICHEL, P., 2011. Development of a multicriteria tool for optimizing the renovation of buildings. *Applied Energy,* **88**(4), pp. 1386-1394.

CRAWLEY, D.B., LAWRIE, L.K., WINKELMANN, F.C., BUHL, W.F.A.H.,Y.J., PEDERSEN, C.O., STRAND, R.K., LIESEN, R.J.A.F.,D.E., WITTE, M.J. and GLAZER, J., 2001. EnergyPlus: creating a new generation building energy simulation program. *Energy and Buildings,* **33**(4), pp. 319-331.

DAUM, D. and MOREL, N., 2010. Identifying important state variables for a blind controller. *Building and Environment,* **45**(4), pp. 887-900.

DEB, K., 2001. Multi-Objective Optimization using Evolutionary Algorithms. Chichester : John Wiley & Sons.

DEB, K., PRATAP, A., AGARWAL, S. and MEYARIVAN, T., 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation,* **6(2)**, pp. 182-197.

DE JONG, K.A. 2006. Evolutionary Computation: A Unified Approach. Cambridge, MA: MIT Press, ISBN 0-262-04194-4.

EVINS, R., POINTER, P. and VAIDYANATHAN, R., 2011. Optimisation for CHP and CCHP decision-making, In: *Proceedings of the Building Simulation 2011 Conference, Sydney, Australia* 2011. IBPSA, pp. 1335-1342.

EVINS, R., 2010. Configuration of a genetic algorithm for multi-objective optimisation of solar gain to buildings, *In: Proceedings of the Genetic and Evolutionary Computation Conference* 2010, Portland, OR. New York: ACM Press, pp. 2003-2006.

EVINS, R., POINTER, P. and VAIDYANATHAN, R., 2011. Multi-objective optimisation of the configuration and control of a double-skin facade, *Proceedings of the Building Simulation 2011 Conference, Sydney, Australia* 2011. IBPSA, pp. 1343-1350.

FONSECA, C.M. and FLEMING, P.J., 1996. On the Performance Assessment and Comparison of Stochastic Multiobjective Optimizers, In: *Proceedings of the 4th International Conference on Parallel Problem Solving from Nature* 1996. Berlin: Springer-Verlag, pp. 584-593.

GEYER, P., 2009. Component-oriented decomposition for multidisciplinary design optimization in building design. *Advanced Engineering Informatics,* **23**(1), pp. 12-31.

GRIERSON, D.E., 2008. Pareto multi-criteria decision making. *Advanced Engineering Informatics,* **22**(3), pp. 371-384.

HANBY, V.I. and WRIGHT, J.A., 1989. HVAC optimisation studies: Component modelling methodology. *Building Services Engineering Research & Technology,* **10**(1), pp. 35-39.

HOPFE, C. J. and HENSEN, J. L., 2011. Uncertainty analysis in building performance simulation for design support. Energy and Buildings, 43 (10), pp. 2798 – 2805.

JIN, Q., OVEREND, M. and THOMPSON, P., 2011. A whole-life value based assessment and optimisation model for high-performance glazed facades, In: *Proceedings of the Building Simulation 2011 Conference, Sydney, Australia. IBPSA*, pp. 1017-1024.

JULSTROM, B.A., 1994. Seeding the population: improved performance in a genetic algorithm for the rectilinear Steiner problem, In: *Proceedings of the 1994 ACM symposium on Applied computing*, Phoenix, AZ. New York: ACM Press, pp. 222-226.

KAYO, G. and OOKA, R, 2010. Building energy system optimizations with utilization of waste heat from cogenerations by means of genetic algorithm, *Energy and Buildings*, **42**(7), pp. 985-991.

KNOWLES, J. and CORNE, D., 2002. On metrics for comparing nondominated sets, In Proceedings of the 2002 Congress on Evolutionary Computation, Honolulu, HI. Piscataway, NJ: IEEE, pp. 711-716.

MEZURA-MONTES, E. and COELLO COELLO, C.A., 2011. Constraint-handling in nature-inspired numerical optimization: Past, present and future. *Swarm and Evolutionary Computation,* **1**(4), pp. 173-194.

OH, S., KIM, Y., PARK, C. and KIM, I., 2011. Process-driven BIM-based optimal design using integration of energyplus, genetic algorithm, and pareto optimality, In: *Proceedings of the Building Simulation 2011 Conference, Sydney, Australia* 2011. IBPSA, pp. 894-901.

SAMBOU, V., LARTIGUE, B., MONCHOUX, F. and ADJ, M., 2009. Thermal optimization of multilayered walls using genetic algorithms. *Energy and Buildings,* **41**(10), pp. 1031-1036.

SANAYE, S. and HAJABDOLLAHI, H., 2010. Thermal-economic multi-objective optimization of plate fin heat exchanger using genetic algorithm. *Applied Energy,* **87**(6), pp. 1893-1902.

SHEA, K., SEDGWICK,A., and ANTONUNTTO, G., 2006. Multicriteria optimization of paneled building envelopes using ant colony optimization. *Lecture Notes In Artificial Intelligence,* **4200**. Berlin: Springer-Verlag, pp. 627-636.

SHI, X., 2011. Design optimization of insulation usage and space conditioning load using energy simulation and genetic algorithm. *Energy,* **36**(3), pp. 1659-1667.

SPON, 2011. SPON'S Architects and Builders Price Book. Spon Press, ISBN: 13: 978-0-415-58845-4.

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 Environment,* **45**(5), pp. 1144-1152.

WETTER M, and WRIGHT J A, 2004, "A Comparison of Deterministic and Probabilistic Optimization Algorithms for Non-Smooth Simulation-based Optimization", Building and Environment, 39(8), 989-999 ISSN: 0360-1323.

WRIGHT, J.A. and ZHANG, Y., 2008. Evolutionary synthesis of HVAC system configurations: experimental results. *HVAC&R Research,* **14**(1), pp. 57-72.

WRIGHT, J., ZHANG, Y., ANGELOV, P., HANBY, V. and BUSWELL, R., 2008. Evolutionary Synthesis of HVAC System Configurations: Algorithm Development (RP-1049). *HVAC&R Research,* **14**(1), pp. 33-55.

WRIGHT, J. and MOURSHED, M., 2009. Geometric Optimization of Fenestration, In: *Building Simulation 2009 - the Eleventh International IBPSA Conference*, July 2009, Glasgow, UK. IBPSA, pp. 920-927.

YANG, R. and WANG, L., 2011. Multi-objective optimization for decision-making of energy and comfort management in building automation and control. *Sustainable Cities and Society,* **2**(1), pp. 1-7.

ZITZLER, E., 1999. *Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications*. Thesis (Phd). Swiss Federal Institute of Technology (ETH).

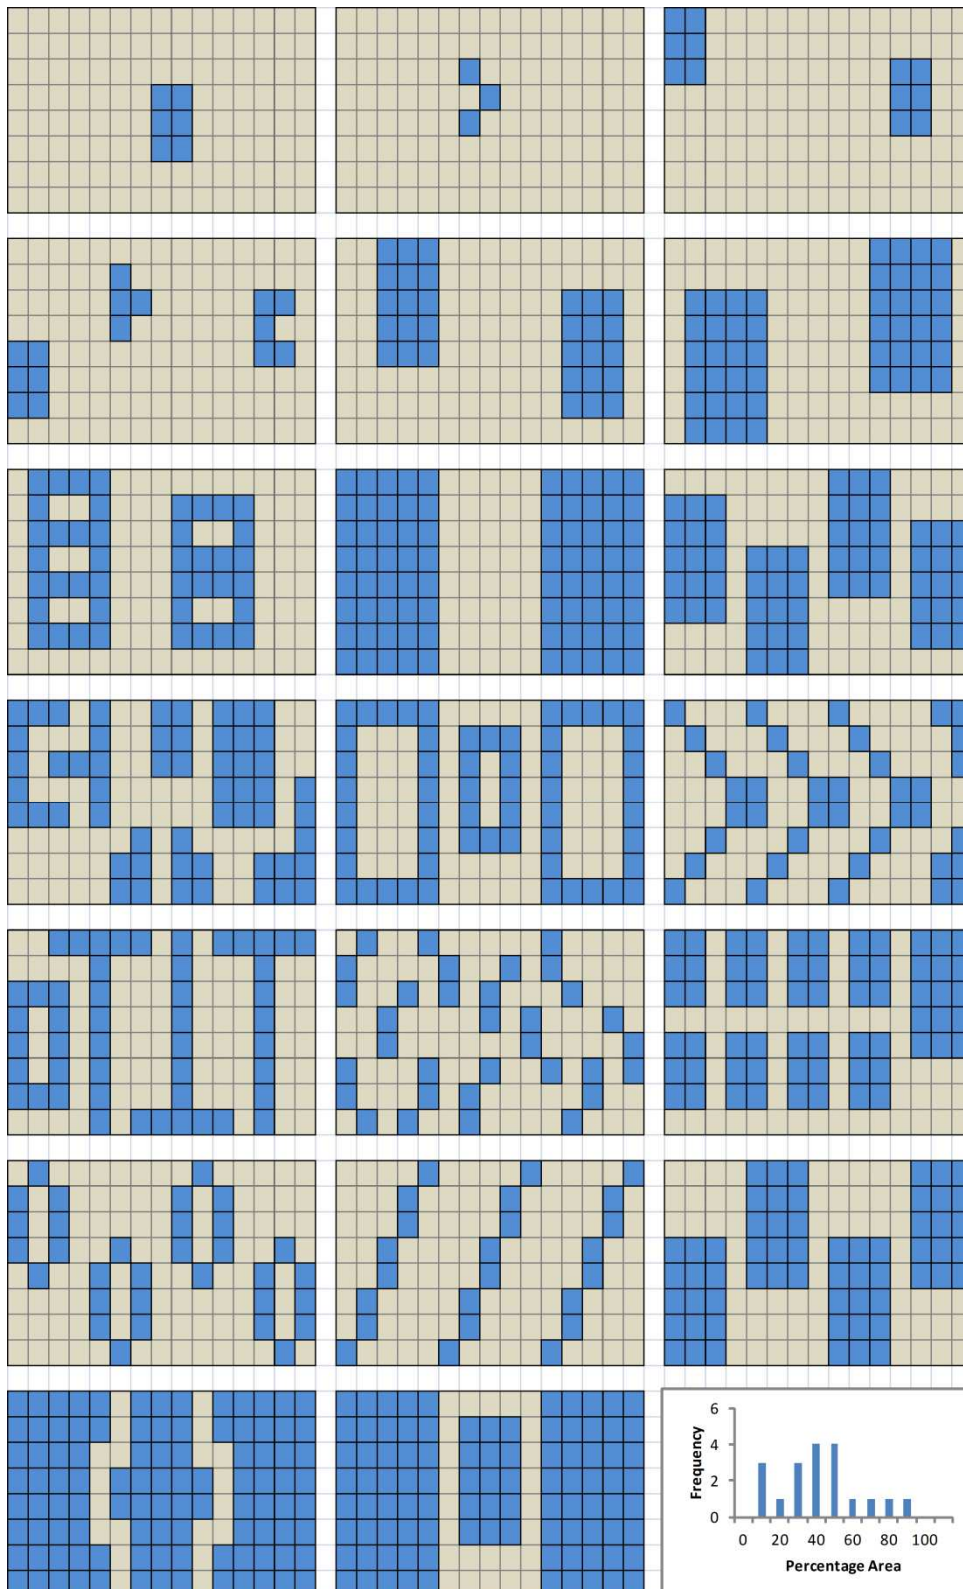**Appendix A: Feasible Seed Solutions**

Figure A1. Seed Solutions and the Frequency Distribution of the Glazed Area among the Seed Solutions