Proceedings of Building Simulation 2011:
12th Conference of International Building Performance Simulation Association, Sydney, 14-16 November.

MULTI-OBJECTIVE OPTIMISATION OF THE CONFIGURATION AND CONTROL OF A DOUBLE-SKIN FACADE

Ralph Evins[1,2], Philip Pointer[1], Ravi Vaidyanathan[2]
[1]Buro Happold, 17 Newman Street, London, W1T 1PD, UK.
[2]University of Bristol, Tyndall Avenue, Bristol, BS8 1TH, UK.

ABSTRACT

We present a new approach to the optimisation of Double-Skin Facades (DSFs). Parameters defined possible geometries, shading devices, openings and ventilation paths, as well as control schedules for their operation. A genetic algorithm was used to discover the best configuration and control strategies for a given scenario from scratch, rather than using a particular configuration type. The algorithm performed a thermal and air-flow simulation of each proposed solution using EnergyPlus. The optimisation process has been illustrated with a case study. In addition, the process has been applied to a range of use types and the results examined graphically to derive innovative design guidelines (a process known as "innovization").

INTRODUCTION

Double-Skin Facades

Double-Skin Facades (DSFs) may be beneficial in reducing the energy used in buildings for heating and / or cooling. A DSF consists of two glazing layers with an air space between. The air space may be sealed or may be ventilated in a range of configurations. For example, if air is drawn into the room through the facade, it will be pre-heated by solar gain to the air space; this should reduce the heating load. Alternatively hot air from the room may be exhausted via the air space, where it will receive the solar gain that would otherwise have been directly transmitted to the room; this should reduce the cooling load.

DSFs are an expensive addition to a building, and must be justified by improved performance. Because of the complex nature of their operation, careful simulation is required to assess their benefits. Performance is dependent upon the geometric configuration of the facade, the operational mode governing air flow, and the control system used to activate different modes. This presents a complex engineering challenge. This paper aims to use an optimisation algorithm to derive the best configuration and control for a DSF for a given scenario, to aid in the resolution of this design problem.

Multi-objective optimisation

Computational optimisation is a rapidly emerging discipline for aiding engineering design. Multi-objective optimisation is particularly useful as it involves the consideration of several objectives simultaneously, with no weightings or aggregations, allowing the robust resolution of complex trade-offs between conflicting objectives. This involves finding the *non-dominated-* or *Pareto-front*, a set of points in the objective space for which no point performs better in all objectives (see Figure 1).

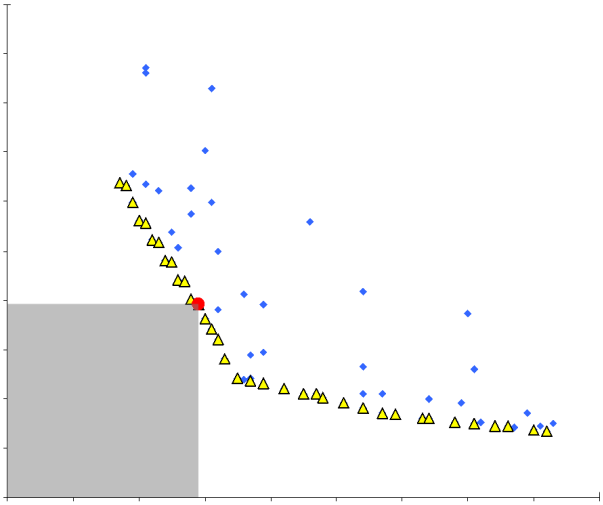

Figure 1: An example of a Pareto front for the minimisation of two objectives, one on each axis. Triangles are members of the Pareto front; dots are not. For each point in the front there is no other point which performs better, e.g. for the highlighted point there is no point within the shaded area.

Genetic algorithms are a common means of achieving multi-objective optimisation. The principle follows that of Darwinian evolution: a population of possible solutions is maintained, with the "fittest" allowed to progress to the subsequent generation. Fitness for multi-objective problems attempts to quantify distance from the non-dominated front. Solutions are altered by crossover (splicing variable values with another solution) and by mutation (changing variable values randomly). The algorithm used in this work is the Non-dominated Sorting Genetic Algorithm (NSGA-II) of Deb et al. (2002). This algorithm selects solutions for subsequent generations based firstly on non-domination rank[1], and secondly by crowd-

[1]Rank 1 solutions are the non-dominated front. These are removed and domination is recalculated to form a new front, which

ing distance[2]. The parameters of NSGA-II used were: populations size 20, number of generations 20, crossover probability 0.9, mutation probability 0.7.

Previous work

Park et al. (2003) applied nonlinear programming to the optimisation of daylighting criteria for a DSF. This was later incorporated into a real-time optimisation program using a lumped model and parameter estimation (Park et al. (2004)), and using a genetic algorithm (Yoon et al. (2011)). Saelens et al. (2005) looked at the introduction of control strategies to improve the performance of DSFs. Stec and Paassen (2005) looked at the symbiosis of the DSF with the HVAC system, including a discussion of predictive control. Charron and Athienitis (2006) optimised the performance of a DSF incorporating photo-voltaics by algebraic means. Gagne and Andersen (2011) optimised for illuminance and glare by varying glazing properties, geometries and shading devices using a genetic algorithm.

DOUBLE SKIN FACADE SIMULATION

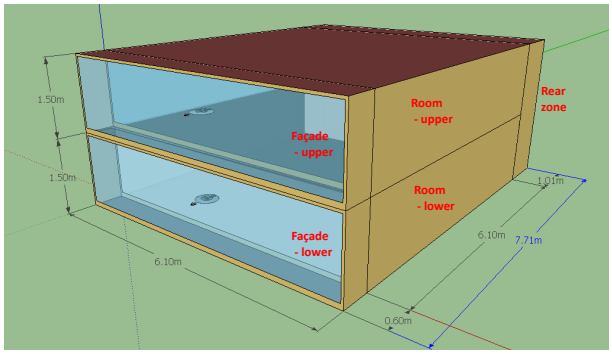

Figure 2: Model geometry. A single room and facade section were simulated, with adiabatic surfaces above, below and to the sides. The front facade was exposed to the outdoor environment; the rear facade had an opening exposed to outdoor air, but no sun or wind.

The simulation of DSFs has been conducted using Energy Plus (see Crawley et al. (2000)). Geometry of the model can be found in Figure 2; construction details are given in Table 1.

All zones had internal gains of $10W/m^2$ for lights and $10W/m^2$ for small power (see 5 for occupancy details). A daylighting control reference point was located in the centre of the lower room space. It was used to modulate the lights to achieve 500 lux during working hours.

An idealised HVAC system was used to control temperatures in the upper and lower room zones; the heating set point was 18^oC and the cooling set point was 23^oC, with night setbacks of 10^oC and 30^oC.

is given rank 2. This process continues until all solutions are ranked. This ensures that the algorithm progresses towards the true non-dominated front.

[2]A measure of the distance of a solution from its neighbors in the objective space. Solutions in less crowded regions are preferred, ensuring that the algorithm explores the whole front.

Table 1: Model construction details. Thermal mass was exposed to the zone.

External Wall U-value	0.25 W/m^2K
External Wall thermal mass	1600 kJ/m^3K
Floor thermal mass	1800 kJ/m^3K
Roof thermal mass	1800 kJ/m^3K
Partition thermal mass	1300 kJ/m^3K
Clear U-value	5.7 W/m^2K
Clear-Clear U-value	2.6 W/m^2K
Clear-LowE U-value	2.3 W/m^2K
LowE-Clear U-value	2.3 W/m^2K
LowE-LowE U-value	2.0 W/m^2K
Clear g-value	0.9
Clear-Clear g-value	0.8
Clear-LowE g-value	0.4
LowE-Clear g-value	0.1
LowE-LowE g-value	0.05

There is no outdoor air requirement specified; it is assumed that if fresh air is supplied via the HVAC system, use of heat recovery will make any load increase negligible.

The simulation was conducted for a 7 day period in early May using London climate data (London Gatwick IWEC weather file); Figure 3 shows the outdoor temperature over the year and for the run period. There were five timesteps per hour.

Air flow simulation

Air flow simulation used the AirFlowNetwork module in EnergyPlus to model forced and buoyancy flows; details of the possible openings and fans are given in Figure 4. The room and facade spaces have been divided into upper and lower parts in order to allow air to be drawn upwards (or downwards) through the facade. Exhaust fans were used to drive air flow through the facade and / or room. This provided a simple, highly controllable means of examining different air flow patterns; a naturally-ventilated solution would require a much greater level of simulation accuracy, and if necessary could be investigated at a later stage. The aim of this work is to investigate the effects of potential air flow regimes and controls, rather than a detailed examination of how these may be achieved.

The air flow rate through the fans was selected from a large range: the minimum of 0.05m^3/s corresponded to a fresh air rate of 12l/p/s for 4 people; the maximum of 0.8m^3/s limited the velocity at the facade openings (area: 1.8m^2) to below 0.5m/s. In order to ensure a viable air flow path, the following rules were enforced:

- If no fan was specified, all windows were kept closed.
- Wherever the fan was specified, the adjoining window was kept closed.
- If the fan was specified in the rear position, an opening was required in both the facade and the

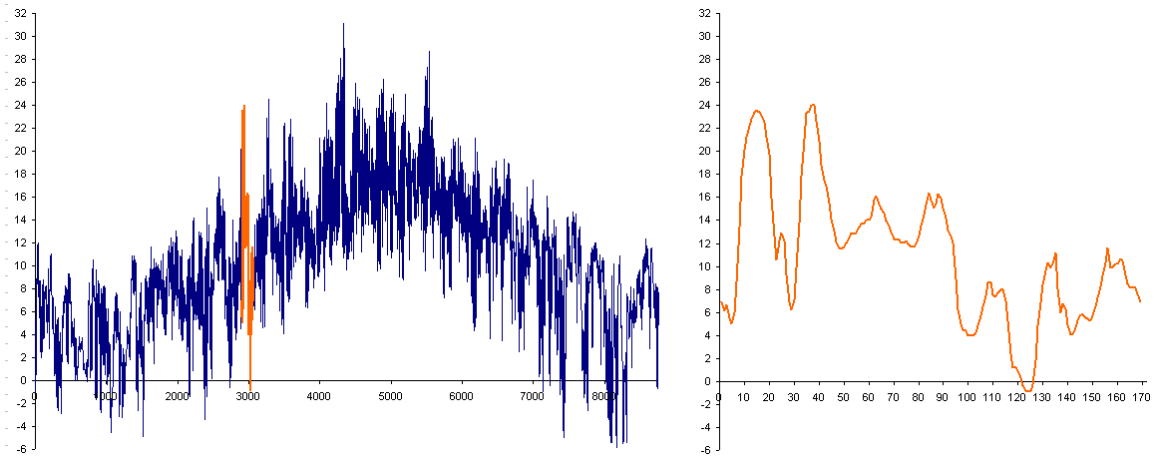

Figure 3: Outdoor temperatures. The left plot shows outdoor temperatures for the whole year, with the simulation period highlighted in orange; the right plot shows the simulation period in detail.

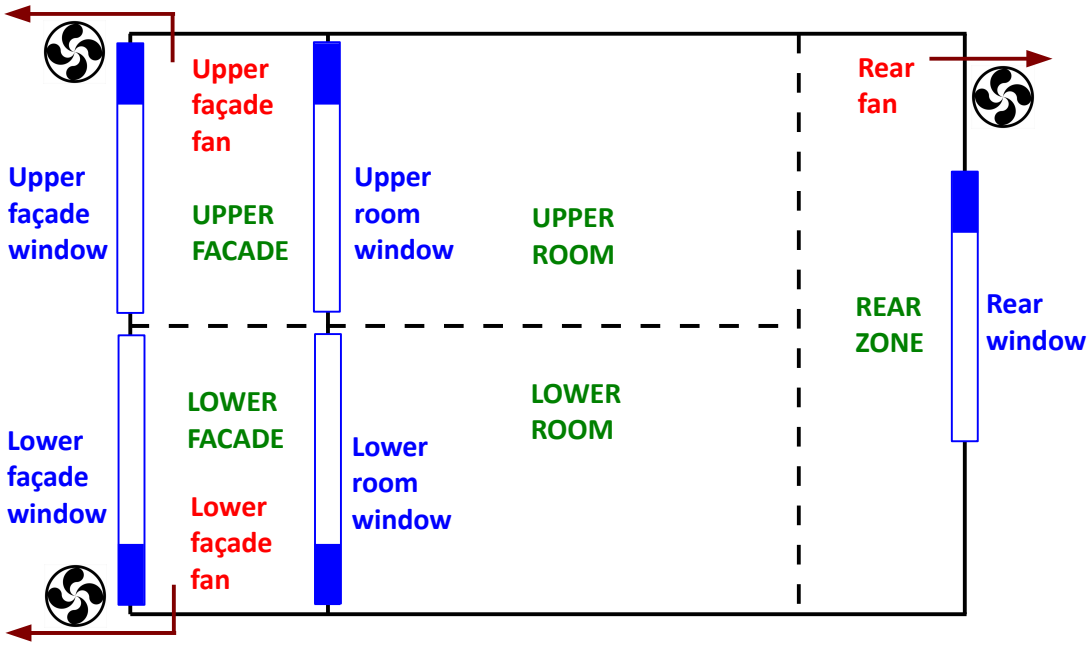

Figure 4: Schematic of openings and fans. Dashed lines indicate divisions between model zones that do not exist physically, and are modelled by openings that are always open. Glazed areas are shown as blue boxes; the shaded region denotes the opening portion (20% of the height, full width). Fans operate in the direction indicated only.

room partitions. If not, the fan was removed.

- If the fan was specified in either facade position, an opening was required in either the facade (to allow circulation through the facade space) or in the room partition and the rear (to allow air to be drawn through the room and out through the facade). If not, the fan was removed.

Control of facade ventilation

The opening or activation of windows and exhaust fans was controlled using logic functions evaluated every time step by the Building Controls Virtual Test Bed developed by Wetter (2008). This allowed a very high degree of control, combined with the flexibility to easily change control parameters. Three functions were generated: ϕ_C was applied to an opening or fan if that element was specified for cooling operation; similarly ϕ_H was used for heating operation; ϕ_B, activated if either ϕ_C or ϕ_H was on, was applied if an element was specified for both heating and cooling. The functions depended on the internal and external temperatures T_{int} and T_{ext} (which were taken from the simulation at each time step), and offsets $\Delta_{H,C}$ and set points $\eta_{H,C}$ (which were variables of the genetic algorithm, and held constant for each simulation). The definitions are given in Equations 1-3. The set point was used to determine whether the schedule becomes active before, at the same time as, or after the HVAC system; the offset was used to specify the difference between external and internal temperatures needed for activation.

$$\phi_C = 1 \; if \; \Delta_C + T_{ext} < T_{int} \; and \; T_{int} > \eta_C \quad (1)$$

$$\phi_H = 1 \; if \; \Delta_H + T_{ext} < T_{int} \; and \; T_{int} > \eta_H \quad (2)$$

$$\phi_B = 1 \; if \; \phi_H = 1 \; and \; \phi_C = 1 \quad (3)$$

OPTIMISATION PROCESS

Twenty variables were included in the optimisation. These are given in Table 2, along with the permissible values used. The decision to keep the problem as general as possible, and hence the high number of variables and permutations, may have increased the time taken for optimal solutions to be found. However, this enabled a more comprehensive exploration of the design space, rather than simply selecting between preconceived solution forms.

There were two objectives: cooling load and heating load over the simulation period. Although the two objectives could have been combined as total heating and cooling load, this would require assumptions about system efficiencies. More information is available to the designer by keeping the two separate; the trade-off between cooling and heating performance is immediately apparent.

It is not possible to completely dissociate heating and cooling performance, for example by running winter and summer design periods. This is because the two affect each other during mid-season operation due to control issues (whether the facade is cooled or heated at different times or to different levels from the HVAC set points) as well as dynamic effects (thermal mass becoming hotter or cooler). It is therefore desirable to simulate the performance of the facade for a mid-season design period, in order to assess its performance at heating and cooling concurrently.

As shown in Figure 3, a mid-season week was selected that combined low minimum temperatures (0^oC) and high maximum temperatures (24^oC). Whilst ideally the simulation would have been conducted for the entire year, this is not practical as part of the optimisation process. The run time for a single week it was $\sim$10 seconds, whereas for a whole year was $\sim$5 minutes; up to 400 evaluations were required for each optimisation. Therefore the use of a single week kept the total runtime to around 1 hour, whereas a full yearly simulation would have taken more than a day.

The initial optimisation addressed the scenario of a conventional office: set points, occupancies and internal gains followed a schedule of typical office hours. Two subsequent optimisations were performed for comparison: the first looked at a meeting room with high but intermittent occupancy; the second considered the use of the space as a luxury hotel. Figure 5 gives the set points and occupancy schedules for each case; other internal gains were kept fixed, and followed the occupancy schedule.

RESULTS

Ventilation patterns for the solutions have been classified according to the types given in Figure 6. Figures 7-9 show the non-dominated (i.e. optimal) solutions found for the three cases along with the solutions for the single-skin designs. Crosses indicate the performance of single facade solutions for the five glazing types; the single facade cases were simulated without air flow, to represent the situation when acoustic concerns prevent the use of opening vents with a single facade design. Figure 10 gives details of the parameters of the optimal double-skin solutions for the three cases.

Office case

For the initial case of a standard office, there were 19 non-dominated solutions after the removal of duplicates. Overall the solutions did not outperform the single facade case, but allowed a much greater range of variability between heating and cooling load. However, this would be achievable using a single facade by fine-tuning the g-value of the glazing. Four DSF solutions had completely closed facades and differed only by glazing type, indicating that the air flow permutations available for the DSF were not greatly helpful in this case. This was particularly true for cooling, for which there were only two solutions that did not use a closed facade. For

Table 2: Optimisation variables.

Cooling mode	Fan position	None, Upper Facade, Lower Facade, Rear
	Fan flow rate, m^3/s	0.05, 0.1, 0.2, 0.4, 0.8
	Lower facade window	Closed, Open
	Upper facade window	Closed, Open
	Lower room window	Closed, Open
	Upper room window	Closed, Open
	Rear window	Closed, Open
Heating mode	Fan position	None, Upper Facade, Lower Facade, Rear
	Fan flow rate, m^3/s	0.05, 0.1, 0.2, 0.4, 0.8
	Lower facade window	Closed, Open
	Upper facade window	Closed, Open
	Lower room window	Closed, Open
	Upper room window	Closed, Open
	Rear window	Closed, Open
Control	η_H, ^{o}C	16, 18, 20
	η_C, ^{o}C	21, 23, 25
	Δ_H, ^{o}C	0, 10, 20
	Δ_C, ^{o}C	0, -10, -20
Glazing	Outer	Clear, Clear-Clear, Clear-LowE, LowE-Clear, LowE-LowE
	Inner	Clear, Clear-Clear, Clear-LowE, LowE-Clear, LowE-LowE

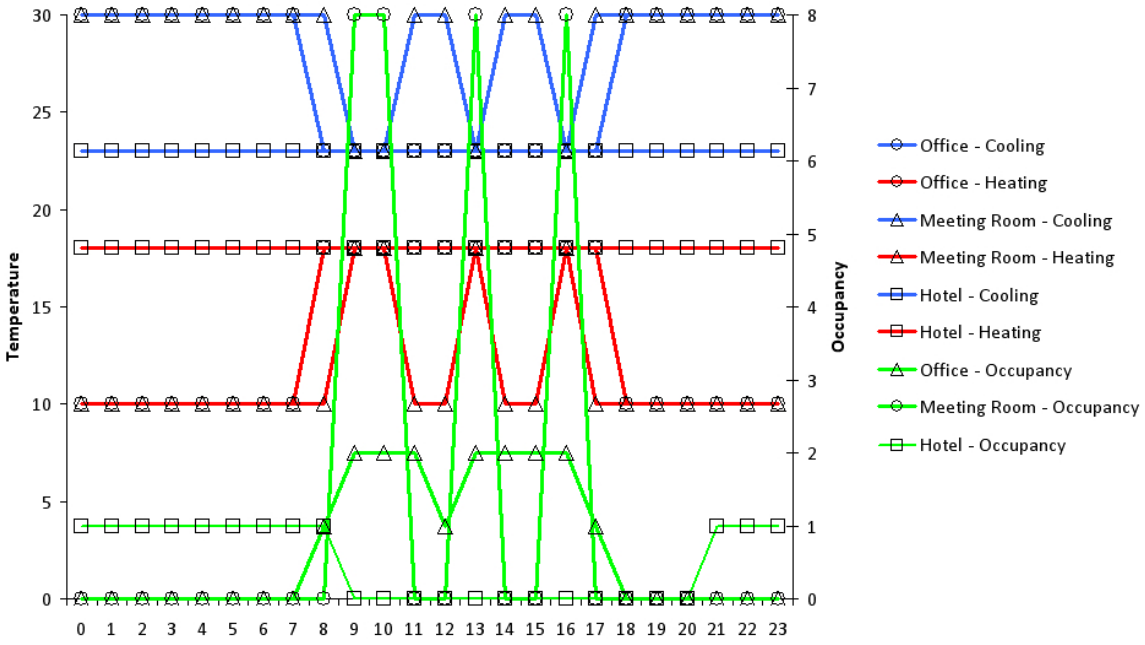

Figure 5: Heating and cooling set points and occupancies for the three scenarios: standard office, meeting room and hotel.

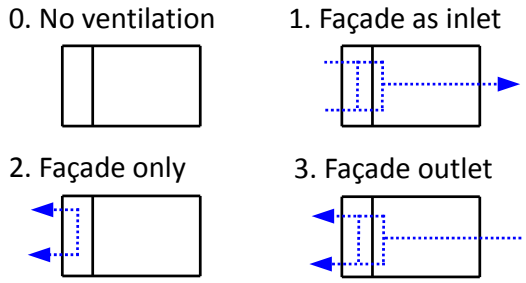

Figure 6: Classification of facade ventilation types.

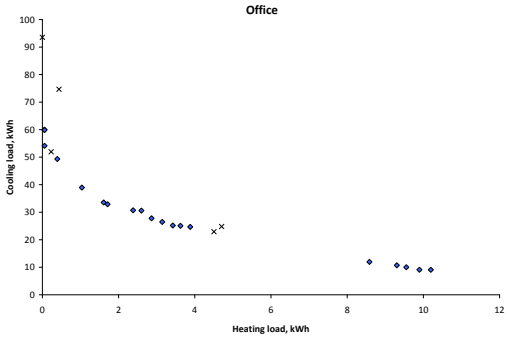

Figure 7: Office solutions.

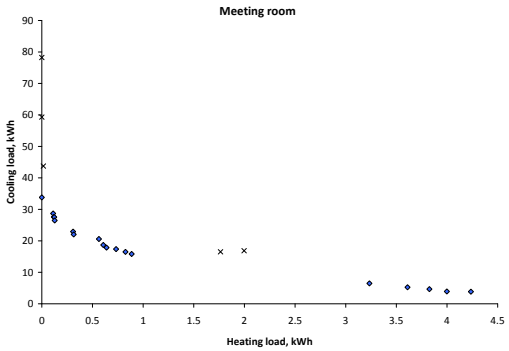

Figure 8: Meeting room solutions.

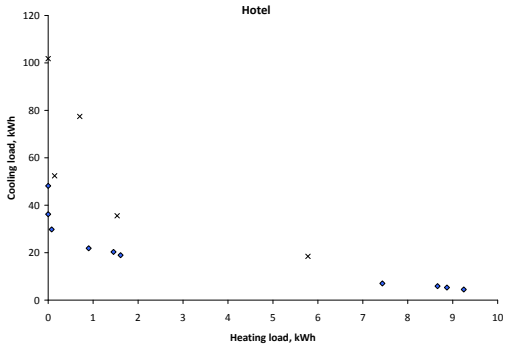

Figure 9: Hotel solutions.

heating ventilation, most of the non-closed cases used the "facade outlet" air flow configuration, with a few using the "facade only" configuration and two using "facade inlet" ventilation. For cooling ventilation, two solutions used "facade outlet" and one used "facade only". In general the solutions used fairly high flow rates; there were no solutions that used the minimum air flow rate. The solutions used a range of set points and offsets, with no specific pattern.

Meeting room case

For the meeting room case, the performance of the DSF solutions (of which there were 18) was better compared to the single facade designs, although not dramatically so. Again there were several entirely closed solutions, with increasing g-value of glazing as cooling load decreased at the expense of increased heating load. All of the non-closed solutions used the "facade outlet" air flow configuration. This is surprising, as it is more commonly used for cooling scenarios, whereas here it was successful at reducing heating loads (8 non-closed heating solutions compared to 2 non-closed cooling solutions). In this case the air flow rates used were lower, with many using the minimum rate and only one using the maximum flow rate. Most cases used the values for set points and offsets that corresponded to operation for the greatest amount of time.

Hotel case

For the scenario concerning use of the zone as a hotel room, the DSF solutions outperformed the single skin designs by a considerable amount. This is surprising, as this is the least common scenario for DSFs to be considered by designers. There were only 10 optimal solutions, only one of which was entirely closed (the case with minimum cooling and maximum heating). There were six non-closed heating cases, all of which used the "facade extract" configuration, and four non-closed cooling cases, of which one used "facade extract" and the rest used "facade only". Air flow rates were generally lower, with no designs using the maximum rate and several using the minimum. All cooling cases used an offset of 10^{o}C and all but one used a set point of 23^{o}C; most heating cases used an offset of 10^{o}C, and all used a set point of 16^{o}C.

CONCLUSIONS

This work has attempted to optimise the performance of a Double-Skin Facade in relation to two objectives: cooling load and heating load. It has shown that there is a large range of performance of a DSF, and that configuration of air flow and the parameters upon which it is controlled have a large impact. In some cases a DSF will only match the performance of a single-skin solution; in other scenarios it may go considerably beyond.

By setting up the optimisation in a very general way, the algorithm has discovered from scratch some of

the air flow configurations commonly employed with DSFs. However, they were not always employed in the usual manner, for example the extraction of air through the facade to reduce heating loads as well as cooling loads. It is important to note that almost all the solutions found beyond the basic closed system have different operation modes for cooling and heating: this highlights the need for precise control over the facade to allow hybrid operating modes. Some air flow configurations which occurred in optimal solutions were distinct from the standard archetypes: for example, in several solutions air from the room was mixed with external air in the facade space before being expelled.

DSF performance was highly dependent upon the use to which the zone was put: for example, configurations that worked well for a meeting room performed poorly for a standard office space. It is particularly interesting that the greatest gains over a single-skin solution were achieved for a hotel space; due to strict acoustic requirements, this could be a profitable area for the use of DSFs. There was also strong evidence of coupling between the heating and cooling operation of the design, for example when ventilation configurations activated during heating periods impacted upon cooling loads. This would not have been apparent if the two modes had been simulated separately.

Further work in this area could investigate in greater detail the links between other aspects of the design and the performance of a DSF. Possible areas include thermal mass, shading devices (including controls), HVAC systems, and the effect of adaptive comfort and other means of relaxing cooling requirements. The other major area in which DSFs have an impact is daylighting: it would be interesting to run a three-objective optimisation in order to incorporate this aspect of performance.

ACKNOWLEDGEMENTS

Funding has been provided by EPSRC and Buro Happold Ltd. The lead author is a Research Engineer with the Industrial Doctorate Centre in Systems, Universities of Bristol and Bath, UK.

REFERENCES

Charron, R. and Athienitis, A. K. 2006. Optimization of the performance of double-facades with integrated photovoltaic panels and motorized blinds. *Solar Energy*, 80(5):482–491.

Crawley, D. B., Lawrie, L. K., Pedersen, C. O., and Winkelmann, F. C. 2000. EnergyPlus: energy simulation program. *ASHRAE Journal*, 42(4).

Deb, K., Pratap, A., Agarwal, S., and Meyarivan, T. 2002. A fast and elitist multiobjective genetic algorithm: NSGA-II. *Evolutionary Computation, IEEE Transactions on*, 6(2):182–197.

Gagne, J. and Andersen, M. 2011. A generative facade design method based on daylighting performance goals. *Journal of Building Performance Simulation*.

Park, C., Augenbroe, G., and Messadi, T. 2003. Daylighting optimization in smart facade systems. *Proceedings: Building Simulation 2003*.

Park, C., Augenbroe, G., Sadegh, N., Thitisawat, M., and Messadi, T. 2004. Real-time optimization of a double-skin facade based on lumped modeling and occupant preference. *Building and Environment*, 39(8):939–948.

Saelens, D., Blocken, B., Roels, S., and Hens, H. 2005. Optimization of the energy performance of Multiple-Skin facades. *Proceedings: Building Simulation 2005*.

Stec, W. and Paassen, A. v. 2005. Symbiosis of the double skin facade with the HVAC system. *Energy and Buildings*, 37(5):461–469.

Wetter, M. 2008. A modular building controls virtual test bed for the integrations of heterogeneous systems.

Yoon, S., Park, C., and Augenbroe, G. 2011. On-line parameter estimation and optimal control strategy of a double-skin system. *Building and Environment*, 46(5):1141–1150.

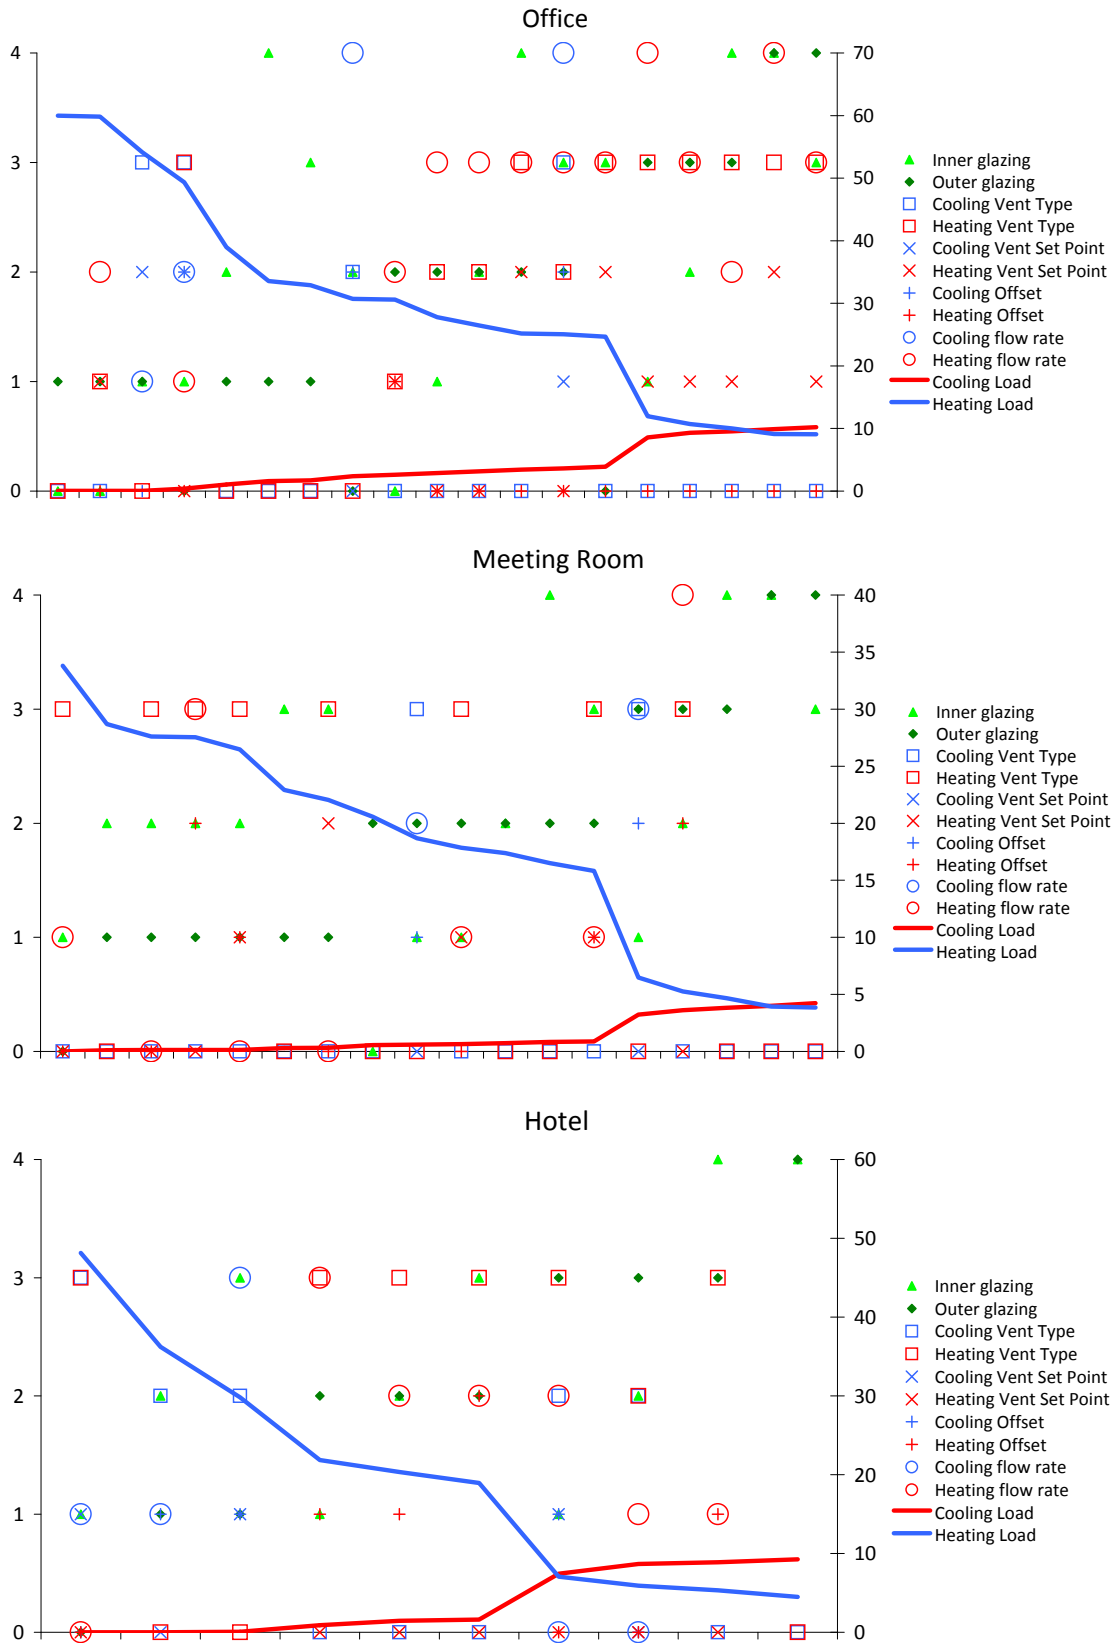

Figure 10: Results for the three cases. The red and blue lines show the heating and cooling loads (right axis, kWh for 7 day run period); red and blue symbols give details of the specification for each case (left axis). Ventilation type corresponds to Figure 6. For set points, offsets, flow rates and glazing types, the value shown (between 0 and 4) indicates the corresponding variable value given in Table 2.