

THE SIMULTANEOUS OPTIMIZATION OF BUILDING FABRIC CONSTRUCTION, HVAC SYSTEM SIZE, AND THE PLANT CONTROL STRATEGY

Jonathan Wright and Raziye Farmani
Department of Civil and Building Engineering
Loughborough University
Loughborough
Leicestershire, LE11 3TU, UK

ABSTRACT

The simultaneous optimization of building's fabric construction, the size of heating, ventilating and air conditioning (HVAC) system, and the HVAC system supervisory control strategy, would automatically account for the thermal coupling between these building design elements. This paper describes the formulation of such an optimization problem for a single air conditioned zone. The problem formulation is described in terms of the optimization problem variables, the design constraints, and the design objective function.

The optimization problem has been solved using a Genetic Algorithm (GA) search method. It is concluded that the GA is able to find a feasible solution and it exhibits an exponential convergence on a solution. The solutions obtained are near-optimal, the lack of final convergence being related to variables having a secondary effect on the energy cost objective function. Further research is required to investigate methods for improving the handling of equality constraints and to reduce the number of control variables (which will also improve the robustness of the algorithm).

INTRODUCTION

The conventional design of buildings is conducted in a sequential manner. The building construction is selected first to conform to local building regulations. The choice of building materials then dictates the required capacity of the heating, ventilating and air conditioning (HVAC) system. Finally, the HVAC system supervisory control strategy is implemented to ensure that the system is functional. However, this approach to the thermal design of buildings is not practicable where the thermal storage in the building fabric is used to offset the peak thermal loads on the HVAC system. For these buildings, the fabric thermal storage is not only dependent on the thermal properties of the building materials, but also the manner in which the HVAC system is controlled. The effectiveness of the thermal storage subsequently dictates the extent to which the peak thermal load, and therefore capacity of the HVAC system, is reduced. This paper describes the application of an optimization method in simultaneously optimiz-

ing the design of the building fabric construction, the HVAC system size, and the HVAC system supervisory control strategy. Since the approach simultaneously optimizes all three design elements (fabric construction, HVAC system size, control strategy) the coupling between them is automatically considered during the design process.

Among other research, optimization methods have been used to investigate the optimum thermal design of building fabric construction (Bouchlaghem, 1990), the size of HVAC systems (Wright, 1996), and the supervisory control of HVAC systems in buildings with fabric thermal storage (Keeney and Braun, 1996; Ren, 1997). Although an approach to optimizing both the plant capacity and control strategy has also been investigated (Arkin et al, 1997), to date, no research has been conducted to investigate the simultaneous optimization of building fabric design, HVAC system size and the supervisory control strategy (termed here, "whole building design"). This paper describes the formulation of such an optimization problem, and its solution by a genetic algorithm (GA) search method.

PROBLEM FORMULATION

The aim of this study is to investigate the simultaneous optimization of the HVAC control strategy, HVAC system size, and building fabric construction. This is achieved through the optimum design of a simple "all outside air" HVAC system. The system consists of a regenerative heat exchanger, cooling coil, heating coil, and supply fan (Figure 1). The control system is open loop to the zone, since the optimization seeks to determine the supply air temperature (θ_s), and flow rate (M_s), in each hour such that the zone conditions remain comfortable. This control configuration has been selected since it can be applied to buildings having ventilated floor slabs as well as conventional air conditioned buildings. The controller has two setpoint's, one for the supply air temperature (θ_{sp}) and a second for the supply air flow rate (M_{sp}). For a given supply air setpoint (θ_{sp}), the control supervisor determines if the heat recovery device should be operational, and whether active heating or cooling is required in order to meet the setpoint.

The supply air flow rate is controlled through a separate fan speed controller.

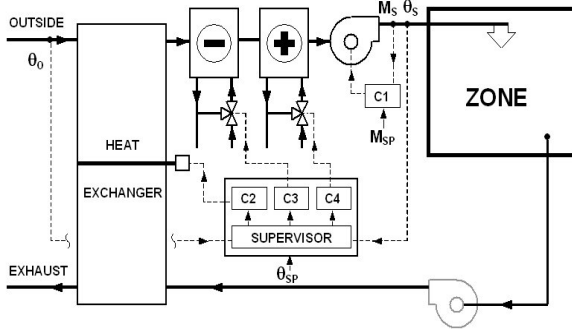


Figure 1: The Example HVAC System

The formulation of an optimization problem for this system can be described through the choice of problem variables, the design constraint, and the objective function to be optimized.

Problem Variables

The problem variables are formed from the control system setpoints, the size of the HVAC components, and the construction of the building fabric. It is necessary to evaluate the performance of the HVAC system for at least two design days if the problem formulation is to reflect the coupling between the control strategy and HVAC system size. For instance, the size of the cooling coil will be dictated by the peak summer operation, and the heating coil by the peak winter operation. In this study, three days of operation have been used to evaluate the performance of a given system design and control strategy, a winter design day, a day in the “swing” season, and a summer design day.

The influence of the system control strategy on the fabric thermal storage is included in the problem formulation by having separate problem variables for the supply air temperature and flow rate setpoints in each hour of the day (giving a total of 48 setpoints for each day of operation). Since it may not be necessary for the system to operate continuously outside of the occupied period, separate system status (ON/OFF) variables have been included for each hour of the unoccupied period (giving a further 15 control variables for each day of operation).

Given three design days, and 63 control variables for each day, there are 189 control variables in the problem formulation. Previous research indicates that in some instances, the number of control variables could be significantly reduced to 15 for each day of operation (Ren, 1997; Ren and Wright, 1997). This is particularly true for the slow react-

ing “heavy-weight” buildings and buildings having ventilated floor slabs, but reducing the number of setpoints may prove less effective for fast reacting “light-weight” buildings. Further, the number of control variables defined here represents the maximum number likely to be required and therefore poses a suitable problem against which to evaluate the performance of the optimization algorithm.

The size of the HVAC system is represented by the width, height, number of rows, and number of water circuits of each coil and the supply fan diameter. The maximum water flow rate to each coil is also a problem variable. This adds a further 11 problem variables. The size of the heat recovery device has been fixed as has the return fan diameter.

Three building constructions have been defined as problem variables, a “light-weight” building, a “medium-weight” building, and a “heavy-weight” building. Each building construction, may also have one of two possible glazing types, clear glass or low emissivity glass. Finally, three possible areas of glazing have been defined, 10%, 20% or 30% of the facade. This gives a further 3 problem variables (building weight, glazing type, glazed area), and a total of 203 problem variables.

The range of each type of problem variable is given in Table 1. Fan sizes 0, 1 and 2 correspond to fan diameters of 0.38 m, 0.445 m, and 0.508 m (this non-linear series being represented by the evenly incremented integer values). The integer values for building weight correspond to 0 for a “light-weight” construction, 1 for a “medium-weight” construction and 2 for a “heavy weight” construction. The clear and low emissivity glass is represented by the integer values 0 and 1 respectively, and the 10%, 20% and 30% glazed areas by integer values of 0, 1, and 2 respectively. In this problem formulation, the integer values for building weight, glazing type, and glazed area together identify the building configuration input file used by the building performance simulation.

Variable Type	Lower Bound	Upper Bound	Discrete Increment
ON/OFF Status (-)	0	1	1
M_{sp} (kg/s)	0.18	1.0	0.01
θ_{sp} ($^{\circ}$ C)	10.0	40.0	0.24
Coil Width (m)	0.50	2.50	0.05
Coil Height (m)	0.50	2.50	0.05
Rows (-)	1	10	1
Water Circuits (-)	2	129	1
Water Flow Rate (kg/s)	0.0	10.23	0.01
Fan Size (-)	0	2	1
Building Weight (-)	0	2	1
Glazing Type (-)	0	1	1
Glazed Area (-)	0	2	1

Table 1: Problem Variables

Design Constraints

The constraints on the optimization problem are derived from the need to maintain a comfortable environment and to ensure that the HVAC system is correctly designed. The zone comfort conditions are represented here by the “predicted percentage of dissatisfied” (PPD). It is not necessary to include a comfort constraint for every hour of the day, it is only necessary to ensure that the highest PPD occurring during the occupied period remains below a specified maximum. Hence only one comfort constraint is required for each day of operation (giving 3 constraints with 3 design days).

The design of the coils are restricted by 3 constraints (giving 6 constraints for both coils). First, the air face velocity is restricted to limit noise, and moisture carry-over in the cooling coil. Second, the water velocity per water circuits is limited to prevent excessive pipe erosion. Finally, the coil water circuits configuration is restricted to ensure that the flow and return headers are on the same side of the coil. This requires that the number of tubes per circuit is an even number (Equation 1).

$$\text{mod}(\text{tubes_per_circuit}, 2) = 0 \quad (1)$$

Since the fan model is derived from manufacturers product data, it is unwise to extrapolate the fan performance beyond the limits of this data (Wright, 1991). Therefore, 4 constraints have been implemented to ensure that the fan remains within its performance envelope. The first two constraints restrict the fan speed while the second two restrict the volume flow rate through the fan.

Finally, there is a need to ensure that the HVAC system has sufficient capacity to meet the supply air temperature and flow rate setpoints. This is ensured by 2 constraints, one on the error in the supply air temperature, and a second on the error in the supply air flow rate. There are therefore a total of 15 constraints on the optimization, 3 comfort constraints, 6 coil design constraints, 4 fan envelope constraints, and 2 setpoint constraints.

Table 2 indicates that the majority of the constraints are inequality constraints ($g(\mathbf{X}) \leq \text{limit}$ or $g(\mathbf{X}) \geq \text{limit}$), although there are also three equality constraints ($h(\mathbf{X}) = \text{limit}$). Note that the fan performance envelope constraints have been normalized to be in the range 0.0 to 1.0. Also, satisfaction of the temperature and flow rate setpoints is indicated by an equality constraint with a value of 0 (failure being indicated by a value of 1). This form of constraint has been adopted since it is also used to represent instances where the HVAC system simulation has failed to con-

verge on a solution.

Constraint Function	Constraint Form	Constraint Limit
PPD (%)	\leq	10.0
Face Velocity (m/s)	\leq	2.5
Water Velocity (m/s)	\leq	1.8
Water Circuits	$=$	0.0
Fan Speed (-)	\leq	1.0
Fan Speed (-)	\geq	0.0
Volume Flow (-)	\leq	1.0
Volume Flow (-)	\geq	0.0
Temperature Setpoint (-)	$=$	0
Flow Setpoint (-)	$=$	0

Table 2: Constraint Functions

Objective Function

In general, building designers seek to minimize either the building capital or operating cost. In some instances, the life-cycle cost of the building is minimized. In this paper, the energy (operating) cost of the system is minimized, since this is the objective function for which there is greatest coupling between the design elements (supervisory control, building construction, and HVAC system size). The minimization of capital and life cycle cost objective functions will be studied in future research.

The operating cost of the HVAC system is defined here simply as the sum of the system energy cost over the three design days; an approach to estimating the annual energy cost from a day of operation in each season requires further research. Hot water has been considered to have been derived from a gas fired boiler and the chilled water from an electric powered chiller. In this respect, a two part electricity tariff has been applied, the high tariff occurring from 8:00am to 12:00pm. The gas is charged at a flat rate over the day.

Optimization Procedure

Figure 2 illustrates the general procedure for solving the optimization problem. Since there are a large number of control variables, it can be difficult to manually define a control strategy that will satisfy the zone PPD constraints, and in this respect, the initial solution to the problem has been randomly generated. It is highly unlikely that a randomly generated solution would be feasible either, which leads to a two part optimization procedure. The initial stage of optimization is to determine a feasible solution; once a feasible solution is found, the optimization can continue to minimize the objective function subject to maintaining the feasibility of the solution(s). Note that the performance of each candidate solution is evaluated through a thermal simulation of the building and its HVAC system.

Building and HVAC System Models

The performance of each candidate design solution is evaluated using a thermal performance simu-

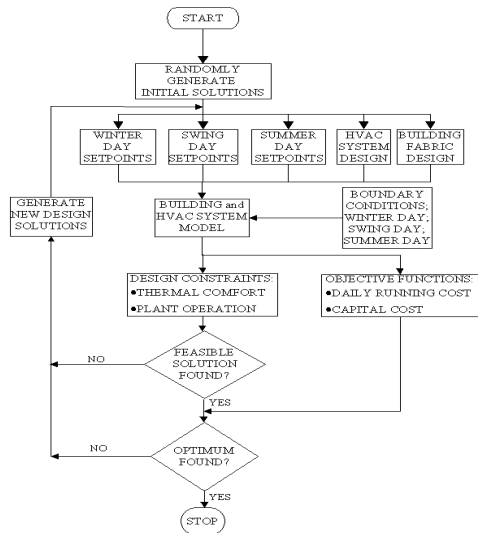


Figure 2: The Optimization Procedure

lation of the building and associated HVAC system. A single zone lumped capacitance model has been used here to represent the thermal response of the zone (Ren and Wright, 1998). The model addresses the thermal storage in a ventilated slab system as well as the building fabric. The model may also be used to simulate the performance of a conventional building constructed without ventilated floor and ceiling slabs. The zone has been modelled as a south facing mid-level zone in a multi-storey building. The zone is subject to a daily profile of adjacent zone temperatures and ambient conditions (temperature and solar irradiance). Three design days of ambient conditions have been selected from UK weather data, a “summer” day, a “swing” season day, and a “winter day”. A profile of internal gains during occupancy has also been defined. The model initialization is ensured by running the model five times with each day of boundary data; however, only the results from the last day in the sequence are used in the optimization. Five consecutive runs of the model has been found to result in acceptable initialization of the model, even for the “heavy-weight” building construction.

The zone air temperature and radiant star temperatures are used in the calculation of the room thermal comfort (PPD). All other thermal comfort parameters (metabolic rate, air velocity, and vapour pressure), have been assigned constant values.

The building construction for each weight of building is given in Table 3. The corresponding mate-

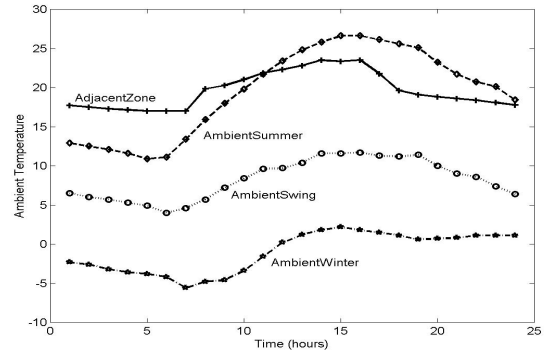


Figure 3: Boundary Conditions

rial properties for each construction element are given in Spitler and Rees (1998).

Fabric Element	Ligh Weight	Medium Weight	Heavy Weight
External Wall	Facing brick, air gap, gypsum sheathing, insulation, wall board	Facing brick, air gap, insulation, concrete block, plaster	Facing brick, air gap, insulation, solid concrete, block plaster
Internal Wall	Wall board, insulation, wall board	Plaster, concrete block, plaster	Plaster, concrete block plaster
Floor and Ceiling	Floor tile, concrete	Floor tile, concrete	Floor tile, concrete

Table 3: Building Construction

The HVAC system performance has been simulated using steady state component models. Both the heating and cooling coil models have been based on a transfer unit, capacity ratio, effectiveness relationship for a counter flow heat exchanger. The overall conductance for the coils has been taken from Holmes (1982). The latent heat transfer in the cooling coil has not been modelled. The fan model is a simple polynomial curvefit of the manufacture’s normalized performance data (Wright, 1991). The pressure drop due to the ductwork has been modelled as a simple quadratic resistance; the quadratic resistance of the coils is also a function of the number of coil rows. The heat recovery device has been modelled as having a fixed temperature effectiveness, regardless of air flow rate. Although in some instances, simplified models have been implemented (such as ignoring latent heat transfer), the models retains sufficient detail and characteristics to allow evaluation of the optimization procedure.

OPTIMIZATION ALGORITHM

The choice of optimization algorithm is influenced by the characteristics of the optimization problem. The “whole building optimization” problem has several notable characteristics. The problem variables are a mixture of integer and continuous variables, the problem has non-linear inequality and equality constraints, and the objective functions can be discontinuous (particularly for the capital

cost objective function). Further, some problems have variables that are not explicitly included in the objective function and only serve to satisfy one or more of the constraint functions. For instance, the maximum water flow rate to a coil would not be a direct parameter of the coils capital cost, but would simply serve to ensure that the coil has sufficient capacity. Finally, the number of problem variables is relatively high.

Traditional optimization methods can be broadly classified as being either gradient based methods or direct search methods. Since the derivatives of the objective and constraint functions are not available, they would have to be obtained numerically by a gradient based method. The effectiveness to which this can be achieved is limited by the discrete variables (some of which are highly discrete, such as the series of fan diameters). Further, where a variable is not a parameter of the objective function, the derivative of the function in relation to the variable has a zero gradient. For these reasons, it can be concluded that gradient based methods are not suitable for solving “whole building optimization” problems.

Direct search methods use heuristic rules to search through the solution space. They have been used successively in solving HVAC supervisory control problems with 24 control variables (Keeney and Braun, 1996; Wheatley, 1999), but have been found to be less effective for problems having 48 control variables (Ren, 1997). The discrete nature of the problem variables has also been found to restrict the extent to which direct search methods can traverse the constraint boundaries if the boundaries are encountered early in the search (Lambert and Wright, 1991). Direct search methods are therefore of limited use in solving “whole building optimization” problems.

Unlike gradient based and direct search methods, “genetic algorithms” (GA’s) are less sensitive to the characteristics of an optimization problem and are reported to be effective in solving a range of optimization problems (Goldberg, 1989). Of particular importance to “whole building optimization” problems is that their relative effectiveness increases with size of the solution space (number of problem variables). This is an important factor in the choice of algorithm for solving “whole building optimization” problems since the number of problem variables can be large. GA’s also process or move all variables in parallel which enables movement in variables not directly linked to the objective function.

The Genetic Algorithm

GA’s operate on a *population* of trial solutions that are initially generated at random. The GA

seeks to maximize the *fitness* of the population by selecting the “fittest” *individuals* from the population and using their “genetic” information in “mating” and “mutation” operations to create a new population of solutions. Many variations of the fundamental GA have been developed, but the algorithm implemented here is derived from the simple GA described by Goldberg (1989). The form of GA implemented can be summarized as being a simple binary encoded GA with “roulette wheel” selection, uniform cross-over, an elitist replacement strategy, and a non-overlapping population.

GA’s are unconstrained search methods, so that enhancements to the basic algorithm must be made to deal with infeasible solutions. This is most commonly done by assigning an individual’s fitness based on both the objective function value and the extent to which the individual satisfies the design constraints. In this respect, one major advantage of GA’s is that the value of the objective and constraint functions is not of direct importance to the search, the GA simply needs to know how “fit” an individual is in relation to the fitness of the other solutions in the population. This allows the use of simple rules for constraint handling. For instance, the simplest means of dealing with infeasible constraints is to assign them a zero fitness (the so called “death penalty”). Since the solutions having zero fitness should not be selected for “mating” and inclusion in the next population, the GA can only continue to make progress towards a solution provided that some of the solutions are feasible and therefore have a fitness greater than zero. This approach to constraint handling is therefore ineffective for highly constrained problems with populations containing a high number of infeasible solutions. This has led to much recent research being directed towards more advanced constraint handling procedures (summarized in Michalewicz and Fogel, 2000). None however specifically address the need to automatically find an initial feasible solution for highly constrained optimization of the type represented by “whole building optimization” problems. This has been addressed by an algorithm developed in this research.

It is beyond the scope of this paper to describe the constraint handling algorithm in detail. However, the algorithm uses a normalized sum of the constraint violations as a measure of a solutions infeasibility (the normalization being described in, Wright and Farmani, 2001). The solutions infeasibility is used to rank the solutions in the population, the ranking forming the fitness of the solution. For instance, if all of the solutions are infeasible then the ranking could be obtained for

each individual by counting the number of solutions with a higher (worse) infeasibility. Since the GA seeks to maximise the fitness of the population, bias would be given to the solutions with the higher rank (lower infeasibility) as the search progresses. A two part ranking (objective and infeasibility) can be applied once a feasible solution is found.

RESULTS AND ANALYSIS

The ability of the GA to solve the “whole building optimization” problem is investigated here by first establishing the pattern of “optimal” energy cost for each building construction. These have been obtained by simultaneously optimizing the control strategy and HVAC system size for each building construction in turn. Table 4 indicates that the results give the expected pattern of energy costs, the energy cost increasing with window area and the “lightness” of the construction. As would be expected, the low emissivity glazing also results in a reduction in energy costs.

No feasible solutions were found for the light weight 30% clear glazed construction since thermal gains during summer operation were higher than could be offset by the largest available system operating with its maximum output; this resulted in failure to satisfy the PPD comfort constraint. This indicates that the variable ranges set in Table 1 are too narrow (an increase in the upper bound of the air mass flow rates would resolve the problem).

The infeasible solution for the lightweight building with 10% clear glazing is due to the search being unable to satisfy the coil water circuit configuration constraints (Equation 1). Such equality constraints have proved to be particularly difficult to solve, with methods for their solution needing to be addressed in future research.

Glass Type	Building Weight	Glazed Area (%)		
		10	20	30
Clear	Heavy	0.9	1.77	2.08
	Medium	0.97	2.10	2.75
	Light	2.10	5.15	-
Low Emissivity	Heavy	0.28	0.63	0.98
	Medium	0.29	0.64	0.99
	Light	-	1.48	2.51

Table 4: Energy Cost Solutions (£) for Control Strategy and HVAC System Size Optimization

Having identified the pattern of energy costs, the optimization problem was extended to simultaneously optimize the form of building construction, the HVAC system size, and the HVAC system control strategy. The optimality of the solutions can be considered in relation to the three design elements, the building construction, the control strategy, and the HVAC system size. The GA correctly identified the construction with the minimum energy cost (heavy-weight, low emissivity

glazing, 10% glazed). Inspection of the control strategies for winter and swing season operation suggested that they are near optimal. No off-peak heating occurred for winter operation since with a flat energy tariff there is no incentive for thermal storage during winter. Inspection of the system operation for the swing season also suggested that it was near-optimal, with much of the thermal load offset by free cooling.

Figure 4, illustrates the optimum supply air temperature and flow rate setpoints for the summer design day. Limited use has been made of the fabric thermal storage, although there is a low tariff period between mid-night and 8:00am. Only one period of free cooling occurred between 4:00am and 5:00am. System operation did not begin again until the start of occupancy at 8:00am, and finished at 5:00pm (the end of occupancy). Inspection of the supply temperature setpoints indicates that during most of the morning, the setpoint matches the ambient temperature so that the system is operating in a free cooling mode. However, Figure 5 illustrates that active cooling was required in the late afternoon in order to maintain thermal comfort (Figure 6). The lack of fabric thermal storage and active cooling in the late afternoon is repeated in the solutions for the other building constructions studied here. This suggests that it is a characteristic of the optimization problem as formulated here, and may in fact represent optimal operation of the system (this will be confirmed in further research). The period of heating between 9:00am and 10:00am is clearly sub-optimal, although for a low heating output. Finally, the supply air flow rate setpoints are close to their lower bound throughout the occupied period, which suggests that these setpoints are optimal.

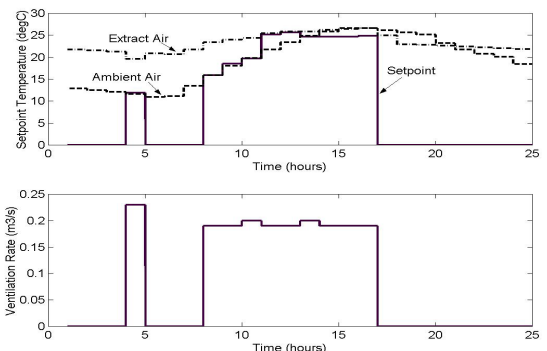


Figure 4: Optimum Summer Setpoints

Figure 6, illustrates the profile of thermal comfort for the summer design day. It is evident that the control strategy is maintaining thermal comfort,

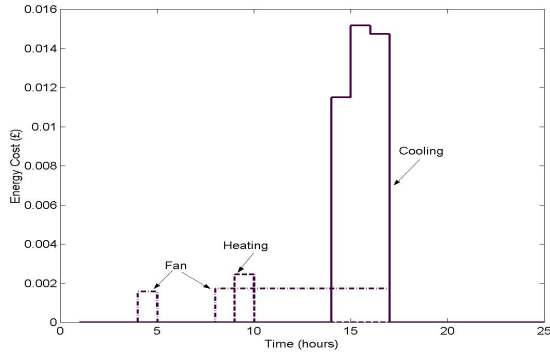


Figure 5: Optimum Energy Costs

with the peak PPD occurring in the late afternoon when the thermal gains to the zone are highest.




Figure 6: Optimum Summer PPD

In terms of the optimality of the size of the HVAC system components, for minimum energy cost, it would be expected that the coil face area would tend to its maximum and that the number of coil rows would tend to a minimum, both of which would result in a reduction in the coil air pressure drop and fan duty. The effect of the total system resistance and supply air flow rate on the choice of fan is less clear (but will be investigated in the future for the size of fans specified). The solutions for the coil dimensions are that the face areas of the coils have been increased to reduce the face velocities to very low levels (0.04 m/s for the heating coil, and 0.09 m/s for the cooling coil). However, the number of coil rows is still higher than expected for each coil (10 rows for the heating coil and 9 rows for the cooling coil). This is sub-optimal but is not surprising considering that the number of rows has a “secondary” effect on the fan energy cost, and that movement in the number of rows is restricted by the coil water circuit configuration constraint (Equation 1).

Due to the probabilistic nature of GA's it is difficult to define criteria for the convergence of the search so that the GA was simply allowed to run here for a 1000 generations (each generation having a population of 80 solutions). Figure 7, illustrates the mean value of the objective function for the population in each generation. The search took 40 generations to find an initial feasible solution, followed by an exponential decrease in the mean objective function value. A measure of the effectiveness of the GA is that it was able to find a feasible solutions in such a short period with 203 problem variables and such a highly constrained variables space, (no feasible solutions were found in 100,000 randomly generated solutions).

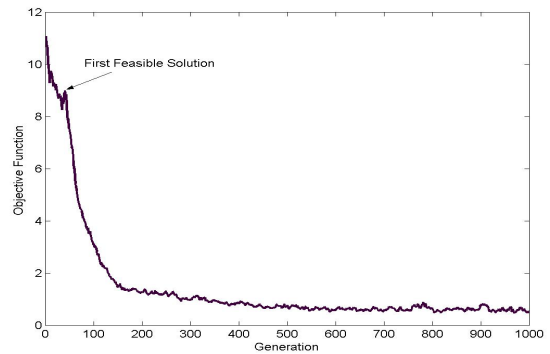


Figure 7: Convergence of the GA

CONCLUSIONS

In order to design a building with fabric thermal storage systems, it is necessary to account for the effect of the thermal storage and its control strategy on the HVAC system capacity. This can be achieved by simultaneously optimizing the choice of building construction, HVAC system size, and HVAC system supervisory control strategy (the simultaneous optimization automatically accounting for the thermal coupling between the design elements). This paper describes an example formulation for a single zone HVAC system. The problem formulation is described in terms of the optimization variables, the design constraints, and an energy cost objective function. The problem variables are formed from the supply air temperature and flow rate setpoints, the heating and cooling coil dimensions, the supply fan diameter, and the form of building construction. The design constraints arise from fluid velocity limits, the coil water circuit configurations, the zone thermal comfort, the envelope of fan performance, and the need for the control system to meet the supply air setpoints.

The characteristics of the optimization problem can be summarized as being one with mixed-

integer problem variables, non-linear inequality and equality constraints, and a non-linear objective function. The high number of problem variables lead to a large solution space which is highly constrained. These characteristics suggest that solution of the problem by gradient based or direct search methods would be difficult, but that the problem lends itself to solution by a Genetic Algorithm (GA).

The results indicate that the GA is effective in finding a feasible solution from an initial randomly generated population of solutions; the search then exhibits an exponential convergence on a solution. Inspection of the solution revealed that the search correctly identified the optimum building construction. The supervisory control strategies are also near-optimal (the lack of off-peak thermal storage appears to be a characteristic of the particular example problem). The size of the HVAC system was found to be sub-optimal, although the non-optimal variables have a secondary and therefore minor effect on the objective function (energy cost).

It is concluded that the simultaneous optimization of the building fabric construction, the HVAC system size, and the HVAC system control strategy can be achieved through the use of the Genetic Algorithm search method. The algorithm is able to find a feasible solution and exhibits rapid convergence on a solution. The solutions obtained are near-optimal, the lack of convergence being due to variables having a secondary effect on the objective function. Further research is required to investigate methods for improving the handling of equality constraints, and in terms of the problem formulation, for the reduction in the number of control variables.

ACKNOWLEDGEMENTS

The authors acknowledge the UK Engineering and Physical Science Research Council for funding this work. The authors would also like to acknowledge the Research and Development Department of Ove Arup and Partners for their collaboration in this work.

REFERENCES

Arkin H., Navon R. and Burg I., "HVAC with thermal energy storage: Optimal design and optimal scheduling", *Proceeding of the CIBSE A: BSER&T*, 18(1), 31-38, 1997.

Bouchlaghem N.M., "Numerical Optimization Applied to the Thermal Design of Buildings" *Building and Environment*, 25(2), 117-124, 1990.

Goldberg, D.E., "Genetic Algorithms in Search, Optimization and Machine Learning", Addison-Wesley, 1989.

Holmes, M.J., "The Simulation of Heating and Cooling Coils for Performance Analysis", *Proceedings of the 1st International Conference on System Simulation in Buildings*, Liege, Belgium, 1982.

Keeney K. and Braun J.E., "Simplified method for determining optimal cooling control strategies for thermal storage in building mass", *HVAC&R Journal*, 2(1), 59-78, 1996.

Lambert G C and Wright J A, "Direct Search Optimisation of HVAC System Design", *Building Environmental Performance Conference*, Canterbury, UK, 196-206, 1991.

Michalewicz Z. and Fogel D.B., "How to Solve It: Modern Heuristics", Springer-Verlag, 2000.

Ren M.J., "Optimal Predictive Control of Thermal Storage in Hollow Core Ventilated Slab Systems", PhD Thesis, Loughborough University, UK, 1997.

Ren M.J., and Wright J.A., "A Ventilated Slab Thermal Storage System Model", *Building and Environment*, 33(1), 1998.

Ren M.J. and Wright J.A., "Predictive Optimal Control of Fabric Thermal Storage Systems", *Proceedings of the 5th International IBPSA Conference*, Prague, Czech Republic, 71-78, (2), 1997.

Spitler J.D and Rees S.J., "Quantitative Comparison of North American and U.K. Cooling Load Calculation Procedures - Methodology", *ASHRAE Transactions*, 36-46, 104(2), 1998.

Wheatley A.C., "Modelling and Performance Analysis of a Sub-Dew Point Chilled Beam in Mixed Mode Buildings", PhD Thesis, Loughborough University, UK, 1999).

Wright J A, and Farmani R, "Genetic Algorithms: A Fitness Formulation for Constrained Minimization", to be presented at the Genetic Algorithm and Evolutionary Computation Conference (GECCO-01), San Francisco, July, 2001.

Wright J.A., "HVAC optimisation studies: Sizing by genetic algorithm", *Proceeding of the CIBSE A: BSER&T*, 17(1), 7-14, 1996.

Wright J A, 1991, "HVAC Optimisation Studies: A Steady State Fan Model", *Building Services Engineering Research and Technology*, 129-135, 12(4), 1991.