Building Simulation 2009
Eleventh International IBPSA Conference
Glasgow, Scotland
July 27-30, 2009

SIMULATION BASED PREDICTIVE CONTROL OF LOW-ENERGY BUILDING SYSTEMS USING TWO-STAGE OPTIMIZATION

Zhen Yu and Arthur Dexter
Department of Engineering Science, University of Oxford, UK

ABSTRACT

Simulation based control schemes for a low-energy building system are introduced and compared in this paper. The simulation of a low-energy system is firstly constructed and a fast two-stage optimisation method is proposed to find the optimal control policy in short time. A Model Predictive Control (MPC) scheme and a Hierarchical Fuzzy Rule based Control (HFRC) scheme that is tuned online by a reinforcement learning (RL) agent are introduced. The MPC scheme runs the simulation online to predict the future behaviour in order to make long-term optimal decisions. On the other hand, the HFRC+RL scheme run the simulation offline to generate prior knowledge for the RL agent.

The performances of the different schemes are evaluated by comparing energy consumption, thermal comfort and computing time.

INTRODUCTION

In the last few decades, computer simulation has enabled more detailed analysis of building energy system that could not be achieved with theoretical or experimental methods. Various information about the building, such as thermal response, energy consumption and even occupancy behaviour, can be acquired conveniently using simulation. This valuable information can be used for system design, equipment sizing, energy audit, fault identification, etc. Particularly, building simulation is very useful in terms of optimising the operation and optimal control of building energy systems. The use of the information about the controlled system is crucial for the controller design. Different ways of using this information exist. The building simulation can be used online as a model of the building energy system therefore a MPC scheme can be constructed based on it. Otherwise, it can be used offline, thus the optimisation results can be used to generate rules or performance maps to control the system online, or can provide good initial values or policies for online learning methods.

In this paper, different ways of using the information provided by a building simulation are discussed and a low-energy building system is used as an example to demonstrate the properties of the different approaches. Firstly, a low-energy building system and its simulation are described. Then a fast two-stage optimisation method is introduced which can find optimal control commands given an associated cost function. Two different ways of using the simulation are described and compared: the simulation and two-stage optimizer is used online as a MPC scheme; or the simulation is used offline to generate a HFRC controller and an accelerated RL method is introduced to tune the HFRC online. The discussion and conclusions are given at the end of this paper.

A TYPICAL LOW-ENERGY SYSTEM

The low energy building system used for simulation is based on a real building in central England (Zhang & Hanby, 2006). The original system is a heating system using solar energy, heat recovery and thermal storage. To provide proper cooling for the summer, a solar driven absorption chiller and an evaporative cooling system, which uses the cooling tower as cooling source, are added (Yu and Dexter, 2008).

The building has three different zones: an exhibition room ($217.9m^2$), a dinning area ($74.0m^2$) and a class room ($178.6m^2$).

The system composition is given in Figure 1 and 2. The building energy system is composed of a water circuit and an air circuit. The water circuit includes a heating source, a cooling source, a thermal storage system, a chilled water system, and a cooling tower. The air circuit is composed of a VPV unit, a heat recovery heat exchanger, an AHU, three heated zones, air supply fans, air dampers and an independent heat recovery unit. The heating and cooling coils of the air handling unit (AHU) act as the interfaces between the water circuit and the air circuit.

A boiler, a solar water collector and a VPV (Ventilated Photovoltaic) panel can be used as sources of energy to heat the system. The hot water stored in the tank can be heated by the boiler using purchased energy or it can be heated by the solar collector if the solar radiation and outdoor temperature permit. When the VPV is not used to heat the inlet air, it is also possible to use the VPV to heat the hot water through an air-to-water heat exchanger. The choice of heat source depends on the water temperature in the storage tank, solar collector and VPV panel, along with the solar radiation, heating and cooling load and outdoor air temperature.

A stratified hot-water storage tank is used for thermal storage.

Active cooling is provided by a solar absorption chiller, which generates chilled water from the hot water in the storage tank. The cooling tower of the absorption chiller can also be used as a direct cooling source when the building needs cooling and the outdoor wet bulb temperature is low enough for evaporative cooling to be used. In evaporative cooling mode, a water-to-water heat exchanger is used to transfer heat directly between the chilled water and the water from the cooling tower, instead of using the absorption chiller for cooling.

The system can operate in different modes according to the positions of the dampers and the use of the fans. For example, the main ventilation system can be run by turning on Fan 1 and Fan 2; or outside air alone can be supplied to the building through the fresh air heat recovery unit by turning on Fan 4. If the main ventilation mode is chosen, the inlet air can be preheated by the VPV panel, warmed by the heat recovery unit, heated by the AHU, or a combination of them, by selecting different positions of the dampers. The VPV unit can work in four modes: discharge mode, preheat mode, storage mode and bypass mode (Cartmell et al., 2004).

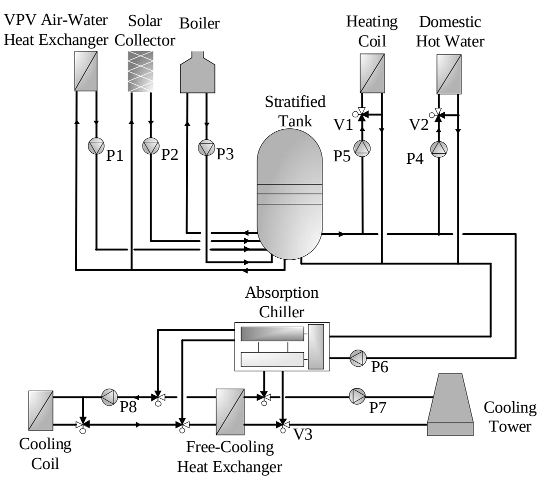

Figure 1 Water sub-system

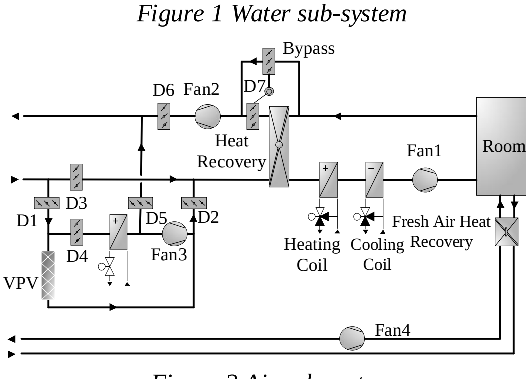

Figure 2 Air sub-system

An equation based simulation of the low-energy building is implemented in a Matlab/Simulink® environment. The simulation uses a simplified first-order lumped-parameter room model , which is identified from measured data taken from the real building (Zhang & Hanby, 2006). Most of the equipment is modelled using component models from the library of an HVAC simulation toolbox (CSTB, 1998). The absorption chiller is modelled by fitting curves (Muneer & Uppal, 1985) to the manufacturer's data (EAW, 2006). The expert rules are an extension of those used to control the heating only system studied previously (Zhang & Hanby, 2006). A manual check was made to ensure that the rules are complete and consistent.

TWO-STAGE OPTIMISATION

Even after the control problem has been redefined and the output space has been greatly reduced, the optimisation search space is $O(2^5)^{24}$ (the value when all of the control signals are assumed to be binary; 5 is the number of outputs after redefinition; 24 is the number of the time steps), which is not feasible to be searched directly for optimal control commands. Besides, the simulation problem is stiff because some equipment has very small time constants compared to the building. The size of the search space, the nonlinearity of the system and the stiffness of the simulation result in an optimisation problem that exceeds the current capability of most computers. A method of simplifying the problem is needed.

A two-stage optimisation scheme is proposed here to reduce the computational demands. Because the equipment can reach steady state much faster than the building, the control time step can be chosen so that the devices will have reached steady state at every step and the dynamics of building are represented adequately. On this time scale, the only linkage between the present and future is the energy stored in the building and the hot-water tank. One hour is often used as the time step by researchers for long term energy analysis of buildings (Braun, 1990, 2003; Henze, 2005; Henze et al., 2004; Kintner-Meyer & Emery, 1995; Nagai, 1999; Zhang & Hanby, 2006). For short term device simulation, a 20 seconds time step is short enough. Thus the problem can be partitioned into maximizing the long term performance in terms of suitable sub goals and achieving the sub goals by optimizing the equipment control variables. The energy stored in the building and tank can be used as sub goals because they are the only linkages between the present and the future in the hourly analysis. Using this approach, the searching space of the problem can be reduced to $O(2^{24})$.

MPC SCHEME

The two-step optimisation is fast enough to be implemented online when the techniques described above are used. A (Dynamic Programming) DP algorithm uses the long-term information to calculate the optimum set points and a short-term optimisation program (and the local controllers) uses these set points together with the short-term information to determine the optimal mode of the operation of the equipment. The performance of this online optimisation based control scheme is compared to the performance of an expert rule based supervisory control scheme in the following part of this paper.

Figure 3 is a block diagram of the online MPC control scheme. The feedback form the system is used to dealt with the uncertainty associated with the predictions. At every time step, only the first step in the optimal command trajectory is used. The optimisation runs once every hour with the updated long-term and short-term inputs and system state. This MPC based control scheme has been widely used by other researchers to study the optimal supervisory control problem of conventional building systems (Braun, 1990; Henze, 2003; Huang et al., 2006; Zhang & Hanby, 2006).

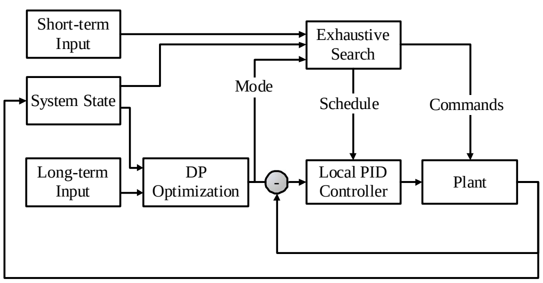

Figure 3 A MPC scheme

HFRC SCHEMES

Non-adaptive HFRC scheme

In the previous part, the online implementation of a dynamic optimisation based MPC control scheme becomes feasible by employing the two-stage optimisation approach. However, the computational demand of MPC is still high. It takes 2 minutes for a long-term optimisation and 10 minutes for a short-term optimisation on a Pentium® 3.4GHz PC in Matlab/Simulink environment. Besides, the optimisation requires detailed models that generally take a long time to develop, which is often unacceptable in practical. On the other hand, fuzzy rule based controllers are widely used for systems with high uncertainties and can be interpreted linguistically. It should also be easier to adapt the rule base online rather than to identify an accurate system model online.

However, since this task is a complex problem with a high dimensional input space, the number of rules is so large that the rules will be difficult to understand and generate. The inputs to the controller include the current state of the system, the current operating environment, as well as the future state and the future operating environment. Those inputs include time, date, occupancy load, building temperatures, tank water temperatures, outdoor dry bulb temperatures, outdoor wet bulb temperatures, and solar radiation. The results of simulation and optimisation are used to generate the fuzzy rules. Ideally, rules should be generated from optimisation results calculated for all combinations of values of the inputs at the centres of the input fuzzy sets, so that the accuracy of the consequence can be guaranteed for every possible input combination. In that case, even if it is assumed that each of the input variables is described by only three fuzzy sets, the total number of rules would still be extremely large (> 3^{73}) and the processing power required to generate them would far exceed that currently available. A hierarchical presentation of the fuzzy rules has been proposed to reduce the number of fuzzy rules and fuzzy decision tree has been used to study the sensitivity of rules further. In this way, the number of rules has been successfully reduced to less than 100 (Yu & Dexter, 2008).

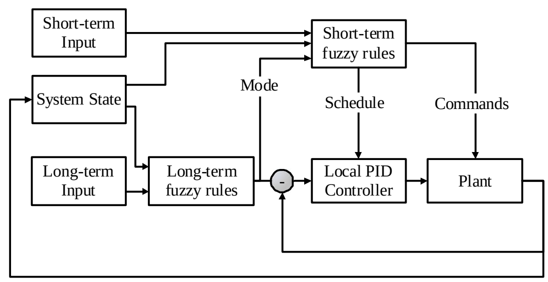

Figure 4 A HFRC scheme

As can be seen in Figure 4, the dynamic programming block in the MPC control scheme is replaced by the long-term fuzzy rule base, which is generated offline, and a short-term fuzzy rule base replaces the exhaustive search block in the MPC control scheme. By generating the fuzzy rules offline, the main computational demands are transferred from the online calculation to the offline optimisation. The employing of the hierarchical fuzzy rule representation make the offline generation of fuzzy rule base for a complex system feasible.

Performance comparison

The performance of the system is evaluated by a cost function R_i which takes account of both the energy consumption and the thermal comfort:

$$R_i = EnergyCost_i + ComfortCost_i \quad (2)$$

where, $EnergyCost_i$ is the energy in kWh purchased from the supplier during the ith hour and

$$\begin{cases} ComfortCost_i = 0, \quad 20 < \mathrm{T_{room}} < 24 \text{ or } \mathrm{i} \in \text{unoccupied period} \\ ComfortCost_i = \alpha \cdot \min\{(20 - \mathrm{T_{room}})^2, (24 - \mathrm{T_{room}})^2\}, \quad \text{otherwise} \end{cases} \quad (3)$$

α is an application dependent comfort-to-energy conversion factor.

Three very cold days are chosen to show the differences in the behaviour of the three control systems in winter. Table 1 compares the comfort cost and energy cost of the three controllers over the three winter days. The ERC refers to the expert rule based controller. The HFRC outperforms the ERC in terms of both comfort and energy use although not as good as the MPC.

As can be seen in Figure 5, the HFRC turns on the boiler to heat the hot water on most mornings no matter whether the water temperature in the hot water tank is low or high. This way of using boiler provides higher tank water temperature, which is essential to raise the room temperature in the cold mornings as quickly as possible. Similar decisions can also be found when the system is under control of the MPC. The benefit of this strategy can be seen is the lower comfort cost of these two control schemes in comparison with the ERC.

The HFRC turns off the main ventilation system and uses only the dedicated ventilation system when the weather is warm and room temperature is higher, e.g. day 1 and 2. To make such a decision, the controller must consider the risk of discomfort if it turns off the main temperature control system. The expert rules are incapable of making such a decision. On the other hand, this kind of decisions is also frequently used by the MPC.

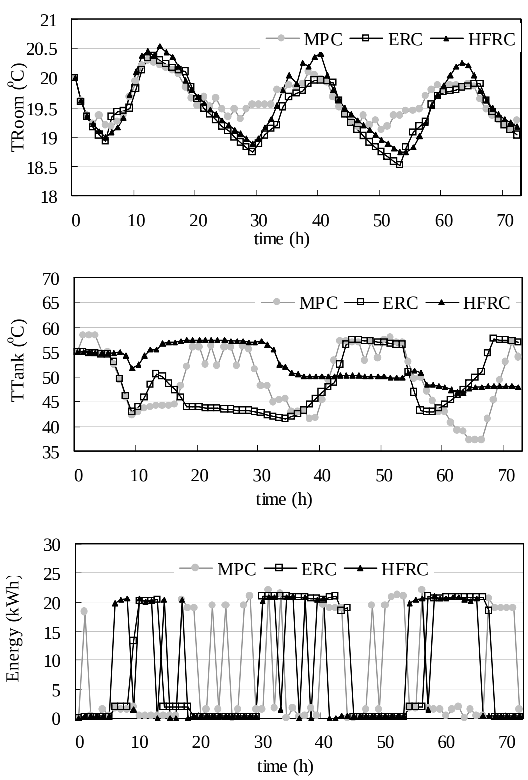

Figure 5 Comparison of room temperature, tank water temperature and energy consumption

The simulation results show that the HFRC can handle the trade-off between the comfort and energy relatively well. The MPC generated the best results among the three but the decision of the MPC is very sensitive to the current building condition and weather information. The commands may vary a lot from day to day even if there are only small changes in the values of the inputs. In contrast, the behaviour of the HFRC is more regular and easier to understand and predict.

There are large differences between the computing times required by the different approaches. The MPC has high computational demands associated with running the online optimisation. In contrast, the computational demand of HFRC is very low and the time required by the controller is negligible compared to the time required to run the building simulation. The high computational demands of MPC not only make the design and validation of the controller near impossible, but also increase the field cost of the control system and make the maintenance and tuning of the MPC based control system difficult.

Table 1
Performance comparison over a 72-hour period during winter

Method	ERC	HFRC	MPC
Energy Cost (kWh)	656	574	540
Comfort Cost (kWh)	416	317	63
Total Cost (kWh)	1072	891	602
Computing Time	<120s	<120s	5 hours

HFRC scheme with online learning

The redefinition of the input and output space and the adoption of a hierarchical fuzzy rule-base are very effective in reducing the number of fuzzy rules (Yu et al., 2007). The control task is divided into supervisory control and the local control. The supervisory control focus on define the longer term energy profile by deciding the room temperature set point (*STroom*), the tank water temperature set points (*STtank*) and the working mode of the plant (*Mode*). After simplification, the number of inputs used by the supervisor is reduced to five: time (*Time*), indoor air temperature (*Troom*), tank water temperature (*Ttank*), an estimate of the current temperature increase (DT_{now}), and an estimate of future temperature increase (*DTfuture*). Therefore, the problem can be reduced to a 5 inputs / 3 outputs problem (Yu & Dexter, 2008).

Although the above simplification is successful in making the generation and implementation of a fuzzy rule based controller feasible, it is still impossible to adapt those fuzzy rules in the limited time that is available for the online adaptation. Even if every input variable is described by only two fuzzy sets, there are still 2^5=32 rules and the correct values of *Mode*, *STroom* and *STtank* must be found for every single one of these 32 different input conditions. The problem is a mixed integer and nonlinear problem. If an exhaustive search is used, the time need to try all possible values of the outputs will be at least 2^5 (Number of Rules)*8 (Number of *Mode*)*5 (Number of *STroom*)*5 (Number of *STtank*) days = 6400 days > 17 years.

Hence, the direct use of *Mode*, *STroom* and *STtank* as actions will lead to a large learning space and long learning time that will make the problem infeasible. To accelerate the learning process, a new two-stage approach based on the use of an indirect learning variable is proposed.

1. Use off-line optimisation to generate sets of supervisory rules for each of five different values of α (the Energy-Comfort trade-off parameter). Here the five values are chosen: 10, 50, 100, 150 and 200.

2. Use the RL scheme to find the optimal value of α given the current state of the system and select the rules that produce the best performance from the five pre-generated rule bases.

By changing the target of the learning from the rule consequences to a parameter representing the different rule bases, the learning time can be reduced significantly to 25*5=125 days. The structure of the proposed tuning process is shown in figure 6.

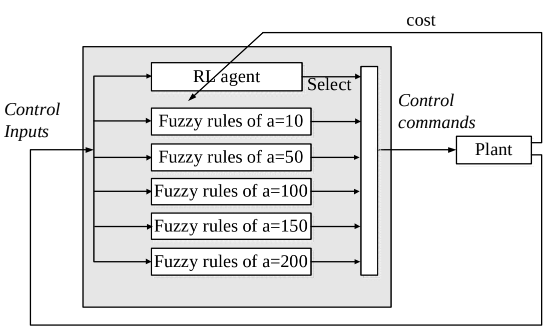

Figure 6 RL online tuning process with indirect input and output variables

Tests are performed to see whether the algorithm can achieve better performance by tuning the energy-comfort trade-off parameter α. A value of 200 is used for α when the "real" system is simulated. It represents the case when the occupancy is more concerned about the comfort than had been assumed when the controller was designed.

Firstly, three of the five template fuzzy rule bases are used to control the simulated plant. The performances of the fuzzy rule bases are compared to that of an expert rule based controller in Table 2. It can be seen that the rules generated with the correct value of α produce the best performance. However, the default rules still perform better than the expert rule based controller even when the design parameter of α is very different to the actual one. The rules of α =10 perform worst in the table. It is because they give a low priority to the comfort, which is not the case in the real situation.

Table 3 gives the performances of the control policy after it has beeen updated using RL learning over different training times. Before the online learning process begins, the initial policy uses the rule base which is generated with an α value of 100. It can be seen that, at first, the performance improves as the agent continues to learn. The best performance is achieved after only three years. However, it is interesting to note that the policy after four years of training is not as good as the policy which only has three years of training.

Table 2
Performance of different rule set when α =200 is used to calculate the cost

	Comfort Cost	**Energy Cost**	**Total Cost**	**Cost/Cost (α =100)**
Expert Rules	14836	5288	20123	158%
Rules of α =10	68184	3395	71579	561%
Rules of α =100	8230	4531	12761	100%
Rules of α =200	2678	5186	7864	62%

Table 3
Performance of the updated control policy after different training time

	Comfort Cost	**Energy Cost**	**Total Cost**	**Cost/Cost (α =100)**
After 1 year	3187	4786	11160	87%
After 2 years	2057	4954	9069	71%
After 3 years	1481	5200	8163	64%
After 4 years	2750	5011	9511	75%

Figure 7 shows the costs produced by the control policy for different years of training. The costs are normalized to the cost produced by the initial policy. After 3 years of training, the policy had already become 64% of the initial cost. From the third year to the twelth year, the cost oscillates between 60%-70%. It shows that the further training is unnecessary and the optimal learning time is two to three years.

Figure 8 gives the average α value during different training periods. Because the comfort in the real building is more important than the designer estimated, online learning increases the frequency of using the rules that are generated with higher α values. This trend slows after three years and the average value of α begins to oscillate around a relatively constant value.

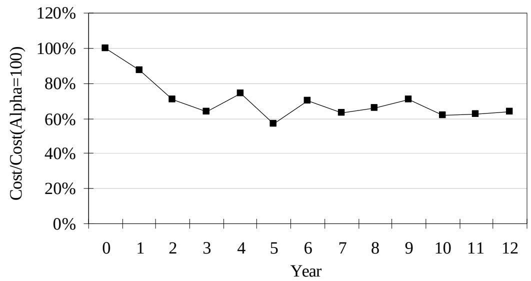

Figure 7 Relative cost of the control policy after on-line training

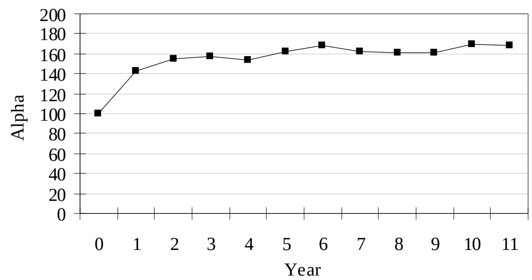

Figure 8 Changing of the average α value

DISCUSSION

It is difficult to guarantee that a simulation of a building energy system is of good quality. An attempt was made to validate the simulation used here by comparing it to results presented in the literature. Usually, only limited data can be found and the operation records are often incomplete. It is therefore nearly impossible to decide how reliable the simulation is under different operational and weather conditions. In addition, to describe the degree of the uncertainty associated with the simulation results is another near impossible task.

The MPC approach produced very good performance results compared to the ERC approach because of its efficient use of the information provided by the building simulation (see table 2). When the models it used are of acceptable quality, the MPC approach also outperformed the HFRC approach although the HFRC approach is better than the ERC. However, it is extremely difficult to know whether the models are good enough. The uncertainty associated with the results of the simulation must be taken into account.

On the other hand, a RL method can be used to tune the HFRC. Few candidate fuzzy rule bases of different possible building models are generated using building simulation. The RL agent learns online to find proper choice of the rule bases according to operational experience. The learning process is transparent and quick in comparison to the online identification of a complex system, which has large number of parameters to learn.

CONCLUSIONS

Different control schemes are designed tested and compared using computer simulation for a low-energy building system. It has been shown that the building simulation can be used in different ways to construct different controllers for energy systems. The simulation results show that the information provided by the computer simulation is essential for good performance of the controller.

A MPC approach is designed by using the computer simulation online as the model of the controlled system to predict future response. The approach produces excellent performance if the building simulation is of reasonable quality. On the other hand, the computational demands, the black-box nature of the controller, and the reliance on an accurate simulation are inherent problems of the MPC approach.

The proposed HFRC+RL approach uses the building simulation offline to provide initial knowledge about the operation of the low-energy system. The RL based online tuning of the controller enable the system to perform well although an inaccurate building simulation was used. This approach also requires very small computational demand and can be understood by domain experts because of its rule-based presentation.

REFERENCES

Braun, J.E. 1990. Reducing Energy Costs and Peak Electrical Demand through Optimal Control of Building Thermal Mass. ASHRAE Transaction, 96(2), 876–888.

Braun, J.E. 2003. Load Control Using Building Thermal Mass. ASME Journal of Solar Energy Engineering, 125, 292-301.

Cartmell, B.P., Shankland, N.J., Fiala, D., Hanby, V. 2004. A Multi-Operational Ventilated Photovoltaic and Solar Air Collector: Application, Simulation and Initial Monitoring Feedback. Solar Energy, 76, 45–53.

CSTB. 1998. Simbad Building and Hvac Toolbox Version 3.0 User Manual: Centre Scientifique et Technique du Batiment.

Henze, G.P. 2003. An Overview of Optimal Control for Central Cooling Plants with Ice Thermal Energy Storage. ASME Journal of Solar Energy Engineering, 125, 302-309.

Henze, G.P. 2005. Energy and Cost Minimal Control of Active and Passive Building Thermal Storage Inventory. ASME Journal of Solar Energy Engineering, 127(3), 343-351.

Henze, G.P., Felsmann, C., Knabe, G. 2004. Evaluation of Optimal Control for Active and Passive Building Thermal Storage. International Journal of Thermal Sciences, 43, 173–183.

Huang, W.Z., Zaheeruddin, M., Cho, S.H. 2006. Dynamic Simulation of Energy Management Control Functions for Hvac Systems in Buildings. Energy Conversion and Management, 47, 926–943.

Kintner-Meyer, M., Emery, A.F. 1995. Optimal Control of an Hvac System Using Cold Storage Building Thermal Capacitance. Energy and Buildings, 23, 19-31.

Muneer, T., Uppal, A.H. 1985. Modelling and Simulation of a Solar Absorption Cooling System. Applied Energy, 19, 209-229.

Nagai, T. 1999. Dynamic Optimization Technique for Control of Hvac System Utilizing Building Thermal Storage. Building Simulation, Kyoto, Japan.CD-c18,

Yu, Z., Dexter, A.L. 2008. Hierarchical Fuzzy Control of Low-Energy Building Systems Using Fuzzy Decision Trees. Paper presented at the Windsor 2008, Air-conditioning & The Low Carbon Cooling Challenge, Windsor, UK

Yu, Z., Dexter, A.L., Zhou, Y. 2007. Hierarchical Fuzzy Rule-Based Control of Renewable Energy Building Systems. Paper presented at the CISBAT 2007, International Scientific Conference Lausanne, Switzerland

Zhang, Y., Hanby, V.I. 2006. Model-Based Control of Renewable Energy Systems in Buildings. HVAC&R Research, 12(3a), 739-760.