Available online at www.sciencedirect.com
SCIENCE DIRECT®
ELSEVIER
International Journal of Thermal Sciences 41 (2002) 1121–1129
International Journal of Thermal Sciences
www.elsevier.com/locate/ijts

Optimization of a building's cooling plant for operating cost and energy use

Darrell D. Massie

Department of Civil and Mechanical Engineering, United States Military Academy, West Point 10996-1792, USA

Received 20 February 2001; accepted 16 November 2001

Abstract

An optimal neural network-based controller for an ice thermal storage system has been developed and tested. The controller consists of four neural networks, three of which map equipment behavior and one that acts as a global controller. The controller self-learns equipment responses to the environment and then determines the control settings required to minimize operating cost. It has the advantage over other controllers in that it always remains calibrated. Since it does not rely upon rules or assumptions, it is able to provide optimal control under any utility pricing and operating condition. Although originally designed to minimize operating costs, simulation and optimization techniques often determine minimum energy use as well.

Keywords: Optimal control; Thermal energy storage; Chiller; Cooling; Least cost; Minimize energy; Neural network

1. Introduction

Cooling of buildings in the United States is a major contributor to the peak electrical load. By some estimates it contributes up to 35 percent of the total electrical demand in the United States, Henze [1]. As such, there is tremendous potential to reduce operating cost and increase energy efficiency with improved control. With the development of advanced computer control techniques, costs can be lowered without the need for trained technicians to continually monitor equipment. This paper will address how neural networks (NN) have been implemented in a laboratory cooling plant to reduce cost and improve energy efficiency.

The chiller is a major Vapor Compression Refrigeration Cycle (VCRC) component of a building's cooling plant that removes energy from chilled water, which is distributed to cooling coils within the building's mechanical system. Operating costs associated with a building's cooling plant are often the highest in comparison with other mechanical components (fans and pumps) within the building. VCRC equipment is also generally more expensive to operate during the day than at night due to time of use electricity pricing used in many countries and higher condenser temperatures due to higher ambient air temperatures. In general, rates are highest during the day when electrical consumption is the greatest.

Using ice storage to cool commercial buildings, termed thermal energy storage (TES), is a load management strategy that can reduce electrical power or energy costs. Although TES systems have been historically used so that smaller chillers could be installed, most of today's installed thermal storage systems are employed to shift the cost of electricity from on-peak to off-peak periods, thus reducing demand and energy charges. Unfortunately, many facility owners are often disappointed with system performance since these systems are not providing the expected load shifting. Poor control has been identified as the primary reason for their insufficient performance, Potter et al. [2]. Optimal control of thermal storage with most of today's price structures is difficult to establish because of the requirement to determine building loads and equipment operating characteristics over a planning horizon.

This paper describes a neural network controller that learns the complex behavior of VCRC equipment and building loads and then optimally controls the system for least cost. Neural networks, are a well-known tool among artificial intelligence techniques. They can reproduce the existing relationship between input and output variables of complex non-linear systems. Thus, they can be used to learn the behavior of complex cooling plants and then be used to control

E-mail address: id4747@usma.edu (D.D. Massie).

1290-0729/02/$ – see front matter
PII: S1290-0729(02)01398-4

Nomenclature

k	hour of each day	1–24
J	monthly power cost	$\$\cdot$month^{-1}
$P(k)$	total power demand due to the cooling and non-cooling load at hour k	kW
$P_{\max,\nu}$	maximum power demand during period	kW
$P_{\text{non-cooling}}$	non-cooling electrical load including plug and lighting loads	kW
P_{plant}	cooling plant power demand including compressors and the air-cooled condenser	kW
$r_{d,\nu}$	demand rate incurred during hour of the month	$\$\cdot$kW^{-1}
r_e	energy charge at hour k of month	$\$\cdot$kWh^{-1}
$SCAP$	ice tank storage capacity	kW–hr
x_k	fraction of storage tank state of charge at start of time k	0–1
x_{k+1}	fraction of storage tank state of charge of at the end of time k	0–1
$x_{\max}$	maximum allowable charge state, set to 1.0.	
$x_{\min}$	minimum allowable charge state, set to 0.	
Δt	unit time step, one hour for this study	hr
ν	demand period	hr
μ	number of days in the month	1–31

them. Because the controller "learns" equipment operating characteristics, equipment models change as equipment ages or undergoes retrofit and there is no need to adjust the controller once installed. It has advantages over other controllers in that it provides optimal solutions and always remains calibrated.

2. Laboratory

The laboratory, Kreider et al. [3], where all equipment testing was accomplished, is a full scale heating ventilating and air conditioning (HVAC) laboratory, with a chiller that is capable of satisfying 236 kW of cooling load and representative of a typical floor of up to 930 m^2 of a commercial building. The laboratory incorporates a central hydronic heating and cooling plant, ice storage tank, air-handling unit with variable frequency drives on the fans, outside air conditioning station and four load simulator zones, two of which are full scale. The laboratory has a computer control and data acquisition system that accommodates analog and digital inputs and outputs. The chiller is a packaged dual-circuit unit with semi-hermetic helical-rotary (screw) compressor. The ice storage tank is a 6560 L, 2.4 GJ nominal capacity ice-on-coil with internal melt storage system. Fig. 1 is a simplified schematic diagram of the cooling plant.

Because the chiller plant is the primary focus of this paper, it will be described in more detail. The 236 kW chiller, uses R-22 refrigerant and has an overall full load EER of 10.1 MBH$\cdot$kW^{-1}. It is a packaged dual-circuit unit with two semi-hermetic helical-rotary (screw) compressors, a shell-and-tube evaporator and electronic expansion valves. The unit comes with its own packaged controls that allow low temperature operation for icemaking purposes, and provides a low ambient lock-out and low water flow cut-out. The typical operating temperature for the primary loop is 3 °C. Heat is rejected from the chiller through an air-cooled condenser located outside the laboratory. The air-cooled condenser was originally designed to serve a single refrigerant circuit but has been retrofitted to serve two independent and isolated circuits. It has a rated capacity, at sea-level, of 176 kW. Each circuit is designed to turn on and off with one of the compressors in the chiller. Each circuit has three fans and three stages. The 176 kW rated capacity is below the 236 kW capacity of the chiller at design

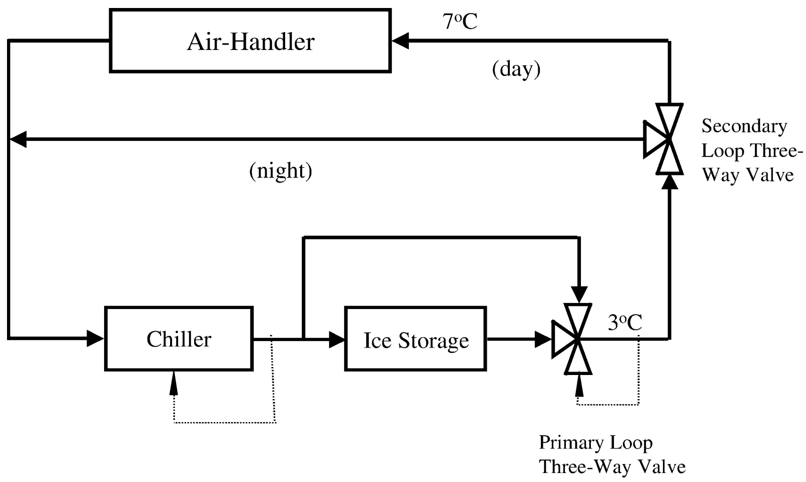

Fig. 1. Cooling plant configuration of Larson Laboratory.

conditions, so unless the ambient temperature is below the design temperature, the overall cooling plant capacity is limited to 176 kW under design conditions.

The chiller plant is located upstream of a 665 kWh of nominal capacity ice-on-coil storage packaged unit tank. A three-way valve located in the primary loop controls the amount of chilled water that flows through the ice storage tank. A secondary loop allows chilled water to flow through the air-handler that places the thermal load on the system. The secondary loop is typically maintained at 7 °C. There are two constant speed pumps located in the primary chilled water loop and one constant speed pump in the secondary chilled water loop.

3. Reducing energy cost using thermal storage

Building owners are motivated to incorporate TES in a cooling system to reduce operating costs. They are rarely knowledgeable about the energy consumption rates of using TES, which, usually consume more energy than direct cooling systems due to the inefficiency of making ice at subfreezing evaporator temperatures, Henze et al. [4], Kintner-Meyer and Emery [5]. Subsequently, when there is no economic incentive to freeze ice for later use, the building's cooling load can be met with direct cooling, which generally consumes less energy.

A significant problem with controlling TES systems is that ice must be formed in advance during periods of low energy cost so that the ice is available for cooling when energy cost is high. This implies that accurate cooling plant equipment models must be available to predict performance and energy consumption over a wide variety of conditions. Energy cost and equipment performance can then be estimated over a planning horizon to determine if it is advantageous to freeze ice in the near term to be used for later cooling.

Control strategies implemented in the field today do not consider variations in building use and equipment operation that can change from year to year, season to season or even day to day. As a result, much of the potential cost savings of using thermal storage systems is lost as cooling plants operate to meet conditions that are assumed and do not occur. For example, many algorithms assume that a full charge of ice might be needed during shoulder months, whereas this is often not the case. Sohn [6] also showed that equipment models that are developed by using manufacturer's data or from laboratory testing can vary significantly from field installed equipment. Even though these difficulties have been recognized, optimal control has not been implemented because of complications accommodating the complex interactions between equipment and the requirement for sensors. Equipment behavior is highly non-linear and varies by location, requiring experts to fine-tune and control. Even for experts with vast experience in installing cool storage equipment, models are complex and require significant effort to calibrate. Furthermore, as equipment ages or undergoes retrofit, models that describe equipment behavior must be changed, requiring further expert assistance.

Several studies have addressed the need for improved TES control. Akbari and Sezgen [7] point out that few TES systems take advantage of daily variations in climate and operating conditions so that charging and discharging are optimized. This work also stresses the need for continued research in TES systems. Different approaches have been used to find optimal TES solutions. Braun [8] used an index of performance over a one-day period to minimize energy cost and Drees and Braun [9] developed a rule-based controller to minimize cost without consideration of energy consumption. Simmonds [10] investigated energy consumption but excluded the effect of price structure, which is the primary concern of building owners. Kintner-Meyer and Emery [11] investigated the sizing of thermal storage components and their impact on the overall system cost and in another study [12], investigated the use of an ice storage facility in conjunction with the building thermal capacitance. Henze et al. [13] developed a simulation environment that used a realistic plant model to investigate the theoretical limit of operating cost savings achieved by cool storage. In another study, Henze and Krarti [14] showed that it was possible to implement TES in such a way that operating costs could be minimized while reducing energy consumption. This work improves on previous work in that it incorporates price structure, equipment forecast modeling and calibration, and an optimizer that determines what combination of the above should be used to minimize cost.

3.1. *Cooling plant control*

This portion of the study describes the development of an optimal controller that minimizes operating cost and also minimizes energy usage. Equipment performance is estimated using self-calibrating NN models developed by Massie et al. [15]. Since these models self-calibrate to installed operating performance, equipment modeling and calibration problems are eliminated.

The chiller load is controlled by adjusting the cooling plant temperature setpoint, which is the chilled water supply temperature at the evaporator exit. The primary loop three-way valve position determines the ice tank flow. The primary loop three-way mixing valve located at the thermal storage tank outlet determines how much of the water coming from the chiller is circulated through the tank. When the valve is set to 100% (termed 100% open), all fluid leaving the chiller circulates through the ice tank. To charge the tank, the chiller setpoint temperature must be below 0 °C and the valve must be opened. The lower the setpoint temperature and more fully opened the valve, the faster the charge. When discharging the tank, if the majority of cooling load is to be met by the chiller, a combination of lowering the chiller setpoint temperature (but still above freezing) and closing the valve position will shift the load onto the chiller. Likewise if more of the cooling load is to be met by the ice tank, a combination of raising the chiller setpoint temperature and opening the primary loop three-way valve must be accomplished.

3.2. Planning horizon and cost function

In optimizing TES systems, the period of time (planning horizon) that energy estimates are to be made into the future must be selected. Henze et al. [16] found that a planning horizon of 21 hours yielded cost differences of less than two percent regardless of the predictor method used. That study used a variety of estimators to include perfect prediction, bin model, random walk harmonic model and auto-regressive neural network. For the current study, a conservative planning horizon of 24 hours was chosen due to the cyclic nature of TES operation.

The objective function in a traditional cost function consists of two parts, the cost of electrical energy [$\$\cdot\mathrm{kWh}^{-1}$] consumed over the billing period and a cost for peak electrical demand [$\$\cdot\mathrm{kW}^{-1}$]. For most US locations, electricity is billed on a monthly basis using two distinct rate periods. The cost J (expressed in monthly units) of operating the cooling plant for one day can be simulated from

$$J = \sum_{1}^{\mu} \sum_{k=1}^{24} P(k) r_e(k) \Delta t + \sum_{\nu=1}^{2} P_{\max,\nu} r_{d,\nu} \quad [\$\cdot\mathrm{month}^{-1}] \tag{1}$$

where μ is the number of days in the month and k is the hour of each day. $P(k)$ is the total power demand due to the cooling and non-cooling load at hour k, r_e is the energy charge at hour k of month and Δt is the unit time step, which has been set to one hour in this study, although it could be any period of time. Demand charges are computed by taking the product of the maximum power consumption $P_{\max,\nu}$ of the demand period ν and the demand rate $r_{d,\nu}$, that is incurred during that hour of the month. For a utility tariff that has two demand periods, ν would take on a value of 1 or 2.

An alternate price structure, real-time pricing (RTP), is designed to charge the consumer more for electricity during periods when electricity is more costly to produce. As a result, there may or may not be a demand charge [$\$\cdot\mathrm{kW}^{-1}$] as found in traditional rate structures. Indications are that most true real-time pricing rate structures will consist of only an energy charge [$\$\cdot\mathrm{kWh}^{-1}$] that will vary for each hour of the day. Typically, the electricity supplier informs users of next day rates in advance so that decisions on how to manage electrical costs (to include ice storage) can be implemented. The cost function for true RTP can then be written as

$$J = \sum_{1}^{\mu} \sum_{k=1}^{24} P(k) r_e(k) \Delta t \quad \mathrm{month}^{-1} \tag{2}$$

Power demand $P(k)$ includes the total of the cooling plant and the non-cooling load and can be computed using

$$P(k) = P_{\text{non-cooling}} + P_{\text{plant}} \tag{3}$$

where $P_{\text{non-cooling}}$ is the non-cooling electrical load including plug and lighting loads, P_{plant} is the power demand of the cooling plant including compressors, the air-cooled condenser, all pumps in the primary and secondary loops and to drive the variable speed fans of the air-handler units.

3.3. Simulation environment

The key feature of thermal storage is to minimize power cost by bridging the temporary difference between cooling load supply and demand. The power consumption, however, is not a control variable, but is instead a result of operation of the cooling system and non-cooling loads. In a system without storage, the building load must be met immediately by the chiller. With TES, the ice storage can be used to meet the building load and the cost of replenishing the storage moved to a period when electricity is less expensive. Therefore, in a system with thermal storage, there is a choice as to which source of cooling will be used at any particular time. Cooling can be taken from the storage, the chiller, or some combination of the two. This decision is based on a comparison of operating costs.

The state of charge x of the storage tank can be represented with a single variable that defines the fraction of maximum ice formation. At any point in time, a decision is needed to either charge, discharge or leave the ice inventory unchanged. For ice storage systems the state transition equation is

$$x_{k+1} = x_k + u_k \frac{\Delta t}{SCAP} \tag{4}$$

subject to the constraints

$$x_{\min} \leqslant x_{k+1} \leqslant x_{\max} \tag{5}$$

where x_{k+1} is the state of charge of storage at the end of time k, Δt is the time step (one-hour), *SCAP* is the storage capacity of the ice tank (e.g., kWh, BTU or ton-hours) and u_k is the charging (+) or discharging (−) rate of storage for time step k. The minimum state of charge, $x_{\min}$ can be set to zero if only the latent heat of fusion is to be considered or to a negative value if sensible heat is to be used. The maximum state of charge, $x_{\max}$, is 1.0.

The rate of charge u_k is automatically set by the NN equipment models if the control settings for plant operation are within the same range as those used to develop the models. For example, if the range used for training a neural network was between $-4\,^\circ\mathrm{C}$ and $7\,^\circ\mathrm{C}$, then the models should not be applied outside of this range. Adjusting the setpoint temperature at the chiller evaporator outlet controls the chiller. A combination of chiller setpoint temperature and the primary three-way valve determine the ice tank rate of charge u_k.

3.4. Neural network controller architecture

Neural networks, are a well-known tool among artificial intelligence techniques. They can reproduce the existing relationship between input and output variables of complex non-linear systems. Thus, they can be used to learn the behavior of complex cooling plants and then to control them. Neural networks are particularly well suited for these types of problems since they are easily configured to map several input variables to multiple output variables. Neural networks are trained to map an input vector to an output vector such that error is minimized (e.g., Bishop [17] or Massie and Curtiss [18]).

The setup of a neural network requires the choice of the number of layers, the number of neurons in each layer, the activation (transfer) function of each layer and the training algorithm, Wasserman [19]. Two phases are then required to make the neural network operative. The first phase is the training (or learning) phase, in which the neural network is taught to match a known set of corresponding input and output values, in order to "learn" the relationship existing between them. Training is achieved through the modification of the weights associated with each neural connection. This is done by the training algorithm, which aims to minimize the error between predicted and actual values in the training set. Training is the most time-consuming phase and it is critical for the success of the neural network as a predictive model. The second phase is called generalization (or testing). The neural network is tested using another known set of corresponding input and output values (none of which belong to the training set) and its performance is evaluated.

Neural networks offer the potential for control of processes through predictive techniques. Jordan and Rumelhart [20] describe the construction of a composite learning system of state-dependent mapping from inputs to predicted sensations in a forward looking network. Anderson [21] described a control algorithm where an inverted pendulum was balanced using a neural network that controlled the movement of the base. Curtiss [22] developed an algorithm in which neural networks were used to control a heating coil and Curtiss et al. [23] implemented a neural network-based energy management program that successfully performed on-line setpoint resets in an actual HVAC system without TES and Jeannette et al. [24] developed a NN to control a heating coil and boiler.

For this study, the neural network-based supervisory controller used to determine hourly setpoints is a recurrent network that computes output sequentially in time. Controlled and uncontrolled input values are fed into the controller network and the network modifies the weights associated with the controlled input to minimize cost. For this network the activation function is a combination of the NN equipment models coupled with the hourly cost function.

The supervisory controller consists of two networks, a training and predictor network, working in parallel (see Fig. 2). The training network is used to learn the relationship between the controlled and uncontrolled variables and the plant characteristics, such as chiller power consumption and tank charge/discharge rate. For example, the controlled variables would be chiller setpoint temperature and ice tank valve position over a 24 hour planning horizon. Examples of uncontrolled variables would be outside air temperature, utility tariffs and building cooling load. The training network weights are then passed to the predictor network where they are used in the activation function for the predictor network. The predictor network subsequently finds values for the control variables that minimize operating cost. Each network operates independently depending on the need to find values

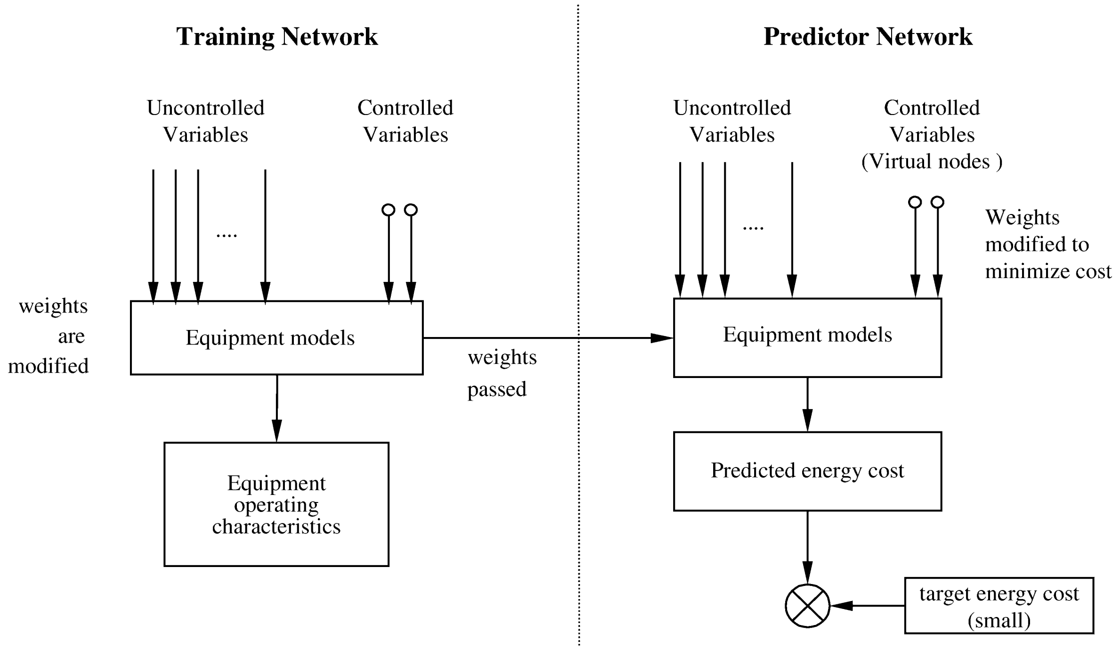

Fig. 2. Architecture of global controller.

for control variables or the need to improve equipment and forecasting models.

The predictor network's goal is to operate the cooling plant for minimum cost. By fixing the equipment model weights, the system is required to find an optimal solution by varying only the control settings. Input for the controlled settings come from "virtual nodes" whose activation is always unity. Weights from the virtual nodes are trained (i.e., modified) to attempt to reach an operating cost of zero. A training goal of zero cannot be attained (unless the plant is turned off), but a value of zero simplifies the mathematics by removing it from the network and the back-propagated error. If a linear activation function is used for the virtual nodes, the weights become analogous to the setpoints. The weights (setpoints) can also be constrained to meet any real-world restrictions or desired limits.

3.4.1. *Training network*

The training network "learns" how a vector of controlled and uncontrolled inputs will affect chiller and ice tank operation. In short, this network trains the equipment models. Training is accomplished by collecting actual input and output data over discrete time periods. The NN equipment models are then compared to actual results and if the predicted output varies from actual equipment performance by more than an allowable error, the data pair is included into the training data set and the network retrained. The allowable error is set to 10% of equipment capacity. A 10% error tolerance was chosen based on findings of Drees and Braun [25].

3.4.2. *Predictor network*

The structure of the predictor network is based on the work of Nguyen and Widrow [26] and Curtiss et al. [27], where each demonstrated how a recurrent network was used to minimize a future error. The predictor network receives weights from the training network and determines the sequence of control actions that minimizes total cost over a 24 hour horizon.

Calculations proceed by determining power consumption and ice charge starting with the first hour and throughout the planning window. Current conditions, at hour k, include information such as monthly peak demand and ice tank state of charge. For each hour into the planning horizon, uncontrolled variables such as hourly ambient temperature and anticipated building load are estimated. Hourly outside air temperature can be estimated sufficiently well using National Weather Service high and low temperature predictions and the ASHRAE model discussed in Chapter 28, Table 2, ASHRAE [28]. Building loads may be estimated by a variety of methods, such as those listed in Kreider and Haberl [29]. The expected combined error of these estimates will vary by building type and location, and will lead to a loss in optimization accuracy of approximately 10%. Each hour's calculation estimates a cost and ice charge at the end of the hour. If a control action attempts to use ice inventory that is not available or fails to meet the buildings cooling load, a cost penalty is applied and back propagated through the network. The current hour's final ice charge becomes the next hour's initial ice charge and so on. At the end of the planning window, hourly costs are summed and compared to the desired cost of zero.

4. Results

The NN controller described here has been thoroughly tested by computer simulation and on a full scale HVAC system. Complete results that demonstrate the robustness of this controller can be found in Massie [30] and Massie and Bailey [31]. Provided here are examples to demonstrate how both operating cost and energy consumption can be reduced.

During the spring and fall seasons, cooling loads are typically smaller than during summer months. If the electrical rate structure encourages load shifting, but the building's cooling load is smaller than the storage capacity, then only enough ice should be made during the off-peak period to meet the building load for the next day. Freezing more ice than required for the cooling load would incur additional cost as well as energy penalties associated with reduction in efficiency while making ice at subfreezing evaporator temperatures. In this example, the on-peak demand [$\$\cdot\mathrm{kW}^{-1}$] and energy [$\$\cdot\mathrm{kWh}^{-1}$] rates were set to five times that of off-peak rates—a rate structure that strongly encourages load shifting to off-peak periods. The on-peak period runs from 0800 until 1700 daily and the building's cooling load is nearly flat and occurs during these same hours. There is no cooling load during the off-peak period.

In this example, see Fig. 3, the building's daily cooling load was set to 60 percent of maximum ice storage capacity. The controller recognized that with an on-peak to off-peak cost ratio of five to one, that there was sufficient economic incentive to move the cooling to the off-peak period and use ice exclusively during the on-peak period. Additionally since there is an off-peak demand charge, it was determined that partial-loading the chiller at a nearly constant six to seven kW was more cost effective than running the chiller at a full load for a shorter period of time. This is the optimal economic solution for the constraints given in this example and also reduces energy consumption by only charging the ice tank to 60 percent. There were some energy inefficiencies associated with the part loading of the chiller, but they were small in comparison to the economic gain. In general screw type compressors have better part load efficiencies than reciprocating compressors.

A comparison of these results will be made with the simulation environment developed by Henze et al. [11] and the rule based controller developed by Drees and Braun [7]. These studies were chosen because they represent the tremendous improvements in TES control that have been made over the past decade. The optimal control environment

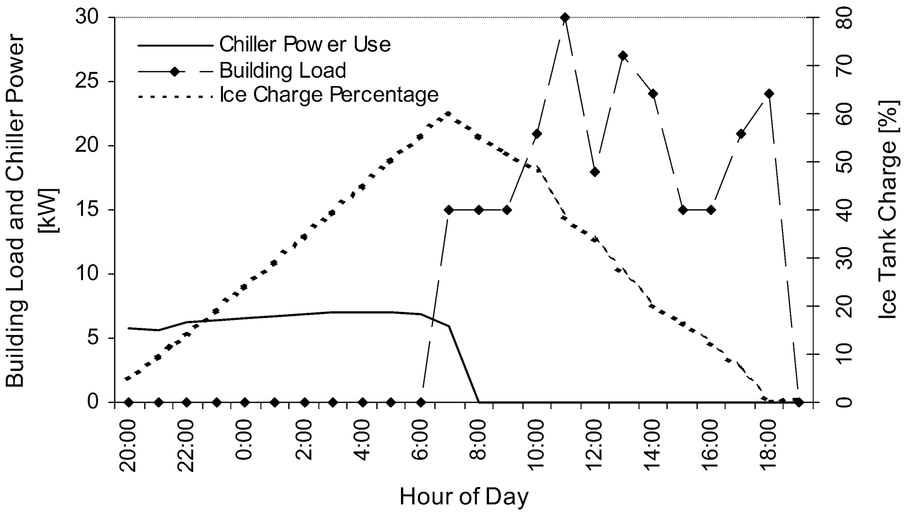

Fig. 3. Shoulder month cooling load when a full charge of ice is not required.

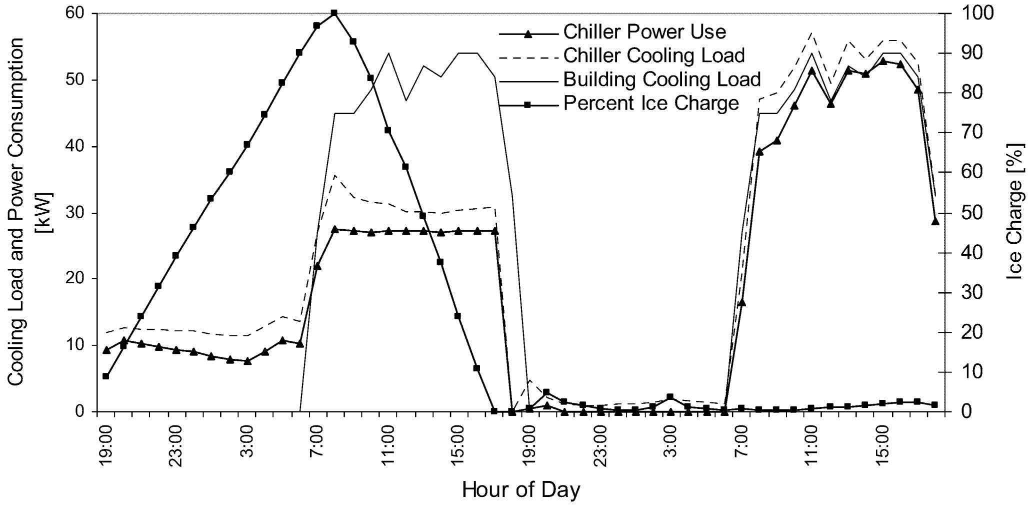

Fig. 4. Consecutive days with load-shifting price incentive followed by no incentive.

of Henze, would have provided an optimal solution, if equipment models had been accurately calibrated. For that study, un-calibrated models were developed for the sole purpose of comparing control strategies. However, since between this study and Henze's, the ice tank and chiller operated in a similar fashion, results are comparable. Results of this study are improved over the Drees and Braun solution. Their method assumes that a full charge is required for the following day and ice is made unconditionally. Their algorithm then discharges the tank so as to minimize cost. Had their assumption of the requirement for a full charge been correct, then their results would also be comparable. Unfortunately, the penalty for making ice can be substantial. Kintner-Meyer and Emery [10], for example, assumed a 39% performance penalty for making ice. Results of this study show that the increased cost varies according to environmental conditions and part-load ratios, however, a penalty in excess of 30% was observed.

In the next example, Fig. 4, a two-day period is considered. For both days, as might be found during the height of the cooling season, the building's cooling load is much greater than the storage capacity. The building's cooling load profile is unchanged for both days. A strong price incentive (as in the previous example) is used during the first day whereas there is no load shifting price incentive for the second day. This could occur if the first day were a Friday and the second a Saturday (assuming weekend rates are off-peak).

The NN controller determined that ice should be made for the first day and discharged during the period of high costs. The electrical demand during the on-peak period also remains flat to minimize cost. On the second day, the controller uses direct cooling since there is no price incentive to shift the cooling load and there is a thermodynamic (and economic) penalty if ice were frozen to meet the second day's load. This demonstrates the controller's ability to respond to price signals. As a side note, since the planning horizon for the NN controller is 24 hours, the controller considers each planning window as a separate event. However, because of the cyclic nature of TES use, the solution is optimal.

With real-time pricing (RTP), conventional control strategies become even less effective since they are based on assumptions that will not apply as rates vary throughout the day. Solutions under RTP are more difficult to verify because economics and equipment performance vary hourly, if not more frequently. To demonstrate the robustness of the NN controller, the RTP rate shown in Fig. 5 was used. Assuming utility tariffs are known 24 hours in advance, there is a three-hour period (from 1100 to 1300) where the utility rates are highest and use of ice storage should be maximized. We also note that the next highest tariff is found at 1400 and so an optimal solution would maximize use of the TES next. The building load profile is the same as that used for the last example.

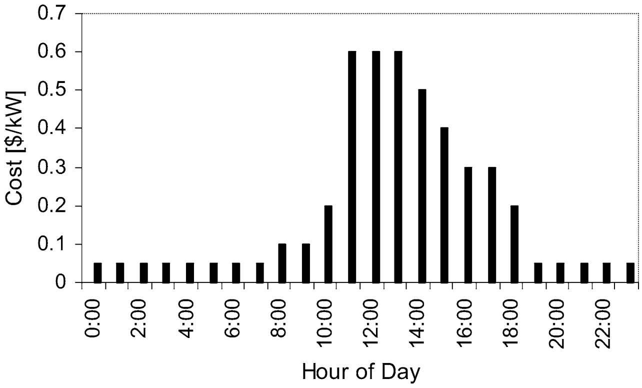

Fig. 5. RTP price structure.

The optimal solution is shown in Fig. 6. The solution shows that the optimal trajectory is to build ice inventory during the night when rates are low. During the three hours when rates are highest, the chiller is turned off completely and all cooling comes from the ice tank. Between low cost and peak cost hours the chiller is operated such that the building load is met and storage is for the most part neither increased nor decreased. The exception to this is at 1400 (second most costly per kilowatt-hour), when the remaining ice inventory, not used during the peak period, is depleted. These results agree with expected results for this pricing.

5. Conclusions

A neural network-based optimal controller has been developed to control a commercial ice storage system for least cost. It is predictive and considers building load forecasting, equipment behavior and utility rate structures. The controller uses dual networks, one as a training network that is used to create and train processes and a second that cascades the processes developed by the training network to determine setpoints over a planning window, here, 24 hours. Although designed to operate for least cost, it will often operate using minimal energy as well. Since the controller does not rely on assumptions, it is robust in finding solutions given any price structure, building cooling load and equipment operating conditions. Because of its ability to learn patterns, it self calibrates to equipment operating characteristics and does not require an expert to fine tune. This feature insures that the controller will operate optimally as a building or equipment undergoes retrofit.

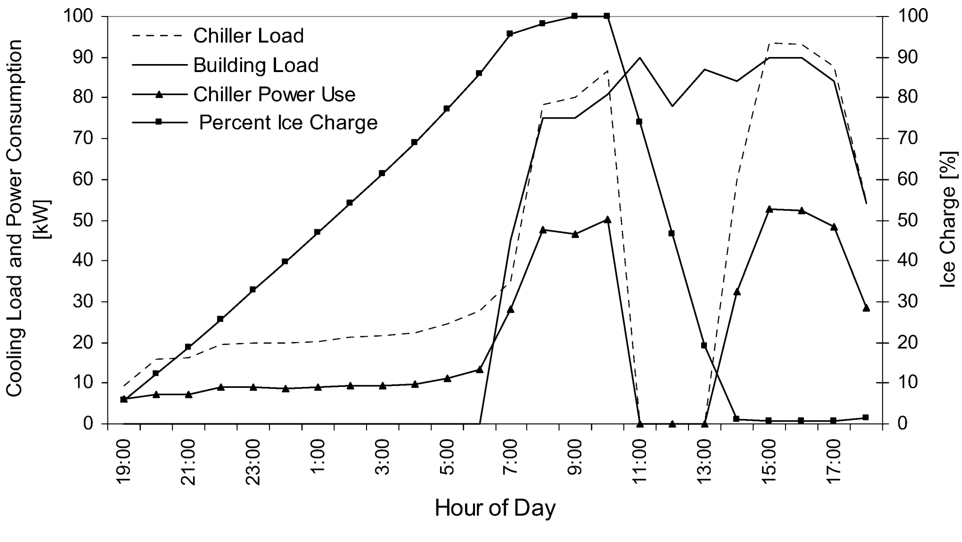

Fig. 6. Consecutive days with load-shifting price incentive followed by no incentive.

References

[1] G.J. Henze, Evaluation of optimal control for ice storage systems, Ph.D. Dissertation, University of Colorado, Boulder, CO: Department of Civil, Environmental and Architectural Engineering, 1995.

[2] R.A. Potter, D.P. Weitzel, D.J. King, D.D. Boettner, ASHRAE RP-766: Study of operational experience with thermal storage systems, ASHRAE Trans. 101 (2) (1995) 549–557.

[3] J.F. Kreider, P.S. Curtiss, D.D. Massie, E. Jeannette, A commercial scale University HVAC laboratory, ASHRAE Trans. 105 (1) (1999).

[4] Development of a predictive optimal controller for thermal energy storage systems, Internat. J. Heating, Ventilating, Air-Conditioning and Refrigeration Research 3 (3) (1997) 233–264.

[5] M. Kintner-Meyer, A.F. Emery, Optimal control of an HVAC system using cold storage and building thermal capacitance, Energy and Buildings 23 (1995) 19–31.

[6] C.W. Sohn, Field performance of an ice harvester storage cooling system, ASHRAE Trans. 97 (2) (1991) 1187–1193.

[7] H. Akbari, O. Sezgen, Case studies of thermal energy storage (TES) systems, evaluation and verification of system performance, LBL-30852, Lawrence Berkeley Laboratory, University of California: Energy and Environment Division, 1992.

[8] J.E. Braun, A comparison of chiller-priority, storage-priority, and optimal control of an ice storage system, ASHRAE Trans. 98 (1) (1992) 893–902.

[9] K.H. Drees, J.E. Braun, Development and evaluation of a rule-based control strategy for ice storage systems, Internat. J. Heating, Ventilating, Air-Conditioning and Refrigeration Research 2 (4) (1996) 312–336.

[10] P. Simmonds, A comparison of energy consumption for storage priority and chiller priority for ice-based thermal storage systems, ASHRAE Trans. 100 (1) (1994) 1746–1753.

[11] M. Kintner-Meyer, A.F. Emery, Cost optimal analysis and load shifting potentials of cold storage equipment, ASHRAE Trans. 101 (2) (1995) 539–548.

[12] M. Kintner-Meyer, A.F. Emery, Optimal control of an HVAC system using cold storage and building thermal capacitance, Energy and Buildings 23 (1995) 19–31.

[13] G.J. Henze, M. Krarti, M.J. Brandemuehl, A simulation environment for the analysis of ice storage controls, Internat. J. Heating, Ventilating, Air-Conditioning and Refrigeration Research 3 (2) (1997) 128–148.

[14] G.J. Henze, M. Krarti, Ice storage system controls for the reduction of operating cost and energy, in: Proceedings of the ASME International Solar Energy Conference, American Society of Mechanical Engineers, 1996.

[15] D.D. Massie, P.S. Curtiss, J.F. Kreider, Predicting central plant HVAC equipment performance using neural networks—Laboratory System Test Results, ASHRAE Trans. 104 (2) (1998) 221–228.

[16] G.J. Henze, R.H. Dodier, M. Krarti, Development of a predictive optimal controller for thermal energy storage systems, Internat. J. Heating, Ventilating, Air-Conditioning and Refrigeration Research 3 (3) (1997) 233–264.

[17] C.M. Bishop, Neural Networks for Pattern Recognition, Clarendon Press, Oxford, 1995.

[18] D.D. Massie, P.S. Curtiss, Neural network basics for use in building mechanical systems, in: Proceedings from 2001 International Conference of Chartered Institution of Building Services Engineers and American Society of Heating Refrigeration and Air-Conditioning Engineers, Inc., 2001.

[19] P.D. Wasserman, Neural Computing: Theory and Practice, Van Nostrand Reinhold, New York, 1989.

[20] M.I. Jordan, D.E. Rumelhart, Forward models: Supervised learning with a distal teacher, Cognitive Sci. 16 (1992) 307–354.

[21] C.W. Anderson, Learning to control an inverted pendulum using neural networks, IEEE Control Systems Magazine (April 1989) 31–36.

[22] P.S. Curtiss, Experimental results from a network-assisted PID controller, ASHRAE Trans. 102 (1) (1996) 1157–1168.

[23] P.S. Curtiss, M.J. Brandemuehl, J.F. Kreider, Energy management in central HVAC plants using neural networks, ASHRAE Trans. 100 (1) (1994) 476–493.

[24] E. Jeannette, K. Assawamartbunlue, P.S. Curtiss, J.F. Kreider, Experimental results of a predictive neural network HVAC controller, ASHRAE Trans. 104 (2) (1998).

[25] K.H. Drees, Braun, Modeling of area-constrained ice storage tanks, Internat. J. Heating, Ventilating, Air-Conditioning and Refrigeration Research 1 (2) (1995) 143–159.

[26] D.H. Nguyen, B. Widrow, The truck backer-upper: An example of self-learning in neural networks, in: Proceedings of the International Joint Conference on Neural Networks, 1989, pp. 357–363.

[27] P.S. Curtiss, J.F. Kreider, M.J. Brandemuehl, Adaptive control of HVAC processes using predictive neural networks, ASHRAE Trans. 99 (1) (1993) 496–504.

[28] ASHRAE, ASHRAE Fundamentals Handbook, American Society of Heating, Refrigerating and Air-Conditioning Engineers, Atlanta, GA, 1997.

[29] J.F. Kreider, J.S. Haberl, Predicting hourly building energy use: The great energy predictor shootout-overview and discussion of results, ASHRAE Trans. 100 (2) (1994) 1104–1118.

[30] D.D. Massie, Neural network controller for commercial ice thermal storage systems, Ph.D. Dissertation, University of Colorado, Boulder, CO: Department of Civil, Environmental and Architectural Engineering, 1998.

[31] D.D. Massie, M.B. Bailey, Optimization and fault evaluation of a building's cooling plant using neural networks, in: ECOS 2000, International Conference on Efficiency, Cost, Optimisation, Simulation and Environmental Aspects of Energy and Process Systems, University of Twente, Enschede, The Netherlands, July 5–7, ASME, 2000, pp. 941–954.