Experimental Analysis of Model-Based Predictive Optimal Control for Active and Passive Building Thermal Storage Inventory

Gregor P. Henze, PhD, PE
Member ASHRAE

Doreen E. Kalz

Simeng Liu
Student Member ASHRAE

Clemens Felsmann, PhD

Received December 9, 2003; accepted August 3, 2004

This paper demonstrates model-based predictive optimal control of active and passive building thermal storage inventory in a test facility in real time using time-of-use differentiated electricity prices without demand charges. A novel supervisory controller successfully executed a three-step procedure consisting of (1) short-term weather prediction, (2) optimization of control strategy over the next planning horizon using a calibrated building model, and (3) post-processing of the optimal strategy to yield a control command for the current time step that can be executed in the test facility. All primary and secondary building mechanical systems were effectively orchestrated by the model-based predictive optimal controller in real time while observing comfort and operational constraints. It was determined that even when the optimal controller is given imperfect weather forecasts and when the building model used for planning control strategies does not match the actual building perfectly, measured utility cost savings relative to conventional building operation can be substantial. Central requirements are a facility that lends itself to passive storage utilization and a building model that includes a realistic plant representation. Savings associated with passive building thermal storage inventory proved to be small in this case because the test facility is not an ideal candidate for the investigated control technology. Moreover, the facility's central plant revealed the idiosyncratic behavior that the chiller operation in the ice-making mode was more energy efficient than in the chilled-water mode. Field experimentation is now required in a suitable commercial building with sufficient thermal mass, an active TES system, and a climate conducive to passive storage utilization over a longer testing period to support the laboratory findings presented in this study.

INTRODUCTION

Background

Commercial buildings contribute a substantial 18% or 17.4 out of 97.4 quadrillion Btu ("quads") to the total US primary energy consumption (EIA/DOE 2002). Aggravated by a surge in the use of office equipment combined with the associated demand for cooling energy, electricity is responsible for 75% of the primary energy consumption in commercial buildings, about 800 million metric tons of carbon emissions per year, and over $65 billion of utility cost. Harnessing the efficiency potential in current and future buildings will be instrumental in attenuating the growth of energy consumption and electrical demand as well as the nation's dependency on an uninterrupted supply of fossil fuels. Unlocking this potential constitutes the motivation for this work.

Gregor Henze is an associate professor and **Doreen Kalz** and **Simeng Liu** are graduate students in architectural engineering at the University of Nebraska-Lincoln. **Clemens Felsmann** is a senior research associate at the Institute of Thermodynamics and Building Systems Engineering, Technical University of Dresden, Germany.

Equipment and systems providing thermal comfort and indoor air quality consume 42% of the total energy used in buildings (A.D. Little Inc. 1999). Energy use and utility cost can be reduced by increasing the efficiency of building systems, by distributing thermal energy more efficiently, and by more closely meeting the needs of building occupants. The energy efficiency of system components for heating, ventilating, and air conditioning (HVAC) has improved considerably over the past 20 years (ARI 1999; American Standard Inc. 1999).

In contrast to energy conversion equipment, less improvement has been achieved in thermal energy distribution, storage, and control systems in terms of energy efficiency and peak load reduction potential. Advancements are also needed to improve thermal storage systems, improve control systems, and improve systems integration from a whole building perspective while meeting occupant comfort and performance requirements (NETL/DOE 2003).

Purpose of Research

This paper discusses the implementation and evaluation of a novel supervisory real-time control strategy in a test facility in order to tackle several of these needed advancements.

Our approach employs the simultaneous utilization of active and passive building thermal storage inventory under model-based predictive optimal supervisory control, i.e., a model representation of the building is provided forecasts of relevant future parameters and used to determine optimal building operational setpoints in real time. In the definition of this article, "active" denotes that thermal storage systems, such as chilled-water or ice storage, call for an additional fluid loop to charge and discharge the storage tank and to deliver cooling to the existing chilled water loop. The use of building thermal capacitance through nighttime precooling is "passive," since it requires no additional heat exchange fluid other than the conditioned airstream. Moreover, a control strategy is considered optimal if it minimizes an objective function of choice. The objective function selected for the model-based predictive controller discussed in this investigation is the minimization of building operating costs while observing constraints on occupant thermal comfort, indoor environmental quality, and HVAC equipment operation.

Literature Review

Cooling of commercial buildings contributes significantly to the peak demand placed on an electrical utility grid. Time-of-use electricity rates encourage shifting of electrical loads to off-peak periods at night and weekends. Buildings can respond to these pricing signals by shifting cooling-related electrical loads either by precooling the building's massive structure, by the use of an active energy storage system only, or by a combination of both thermal reservoirs. Henze et al. (1997) developed a predictive optimal controller for active thermal energy storage (TES) systems and investigated the potential benefits of optimal control for ice storage systems under real-time pricing in order to minimize the cost of operating a central cooling plant. It was found that in the presence of complex rate structures, i.e., real-time pricing rates that change on an hourly basis, the proposed optimal controller has a significant performance advantage over conventional control strategies while requiring only simple predictors. Daryanian et al. (1992) and Daryanian and Norford (1994) describe experimental minimum-cost control of active thermal storage systems under real-time pricing (RTP) rates using an RTP-based energy management system.

Braun (2003) surveyed research on passive building thermal storage utilization, i.e., the precooling of a building's thermal mass during nighttime in order to shift and reduce peak cooling loads in commercial buildings. He identified considerable savings potential for operational costs, even though the total zone loads may increase. Opportunities for reducing operating expenses are due to four effects: reduction in demand costs, use of low-cost off-peak electrical energy, reduced mechanical cooling resulting from the use of cool nighttime air for ventilation

precooling, and improved mechanical cooling efficiency due to increased operation at more favorable part-load and ambient conditions. However, these benefits must be balanced with the increase in the total cooling requirement that occurs with the precooling of the thermal mass. Therefore, the savings associated with load shifting and demand reduction are very sensitive to utility rates, building and plant characteristics, weather conditions, occupancy schedules, operating condition, the method of control, and the specific application. In general, better opportunities for effective precooling exist for higher ratios of on-peak to off-peak rates, longer on-peak periods, heavy-mass building construction with a small ratio of the external area to the thermal mass, and for cooling plants that have good part-load characteristics for which the best performance occurs at about 30% of the design load. Braun et al. (2002) demonstrated control of passive thermal storage in the same facility as used for this work, the Energy Resource Station (ERS) of the Iowa Energy Center.

The combined usage of both active and passive building thermal storage inventory under optimal control has recently been investigated by Henze et al. (2004a) for the reduction of electrical utility cost in the context of common time-of-use rate differentials. The objective function used in the optimization is the total utility bill including the cost of heating and a time-of-use electricity rate without demand charges. The analysis showed that the utility cost savings are significantly greater than either storage but less than the sum of the individual savings, and the cooling-on-peak electrical demand can be drastically reduced. This statement is true for the case when the optimal controller is given perfect weather forecasts and when the building model used in the model-based predictive control perfectly matches the actual building.

While Henze et al. (2004a) established the theoretical maximum performance of this novel control strategy, further research by Henze et al. (2004b) explored how strongly prediction uncertainty in the required short-term weather forecasts affects the controller's cost-saving performance. The best prediction accuracy was found for a bin model that develops a characteristic daily profile from observations collected over the past 30 or 60 days.

It was determined that the predictive optimal control of active and passive building thermal storage inventory involving weather predictions leads to utility cost savings that are only marginally inferior compared to a hypothetical perfect predictor that exactly anticipates the weather during the next planning horizon. This finding rests on the assumption that the building thermal response is perfectly represented by the building model, i.e., there is no mismatch between the modeled and actual building behavior. The primary finding is that it takes only very simple short-term prediction models to realize almost all of the theoretical potential of this storage control technology.

Liu and Henze (2004) investigated the impact of five categories of building modeling mismatch on the performance of model-based predictive optimal control of combined thermal storage using perfect prediction. It was found that for an internal heat gain dominated commercial building, the deviation of building geometry and zoning from the reference building only marginally affects the optimal control strategy; reasonable simplifications are acceptable without loss of cost-saving potential. In fact, zoning simplification may be an efficient way to improve the optimizer performance and save computation time. The mass of the internal structure did not strongly affect the optimal results; however, it did change the building cooling load profile, which, in turn, will affect the operation of the active storage (TES) system. Exterior building construction characteristics were found to impact building passive thermal storage capacity. Thus, it is recommended to ensure the construction material is well modeled. Furthermore, zone temperature setpoint profiles and TES performance are strongly affected by mismatches in internal heat gains; thus, internal gain mismatches should be avoided. Efficiency of the building energy system has no direct impact on the building cooling load, but it affects both zone temper-

ature setpoints (passive storage) and active storage operation because of the coupling to the cooling equipment. Mismatch in this category may be significant.

On the background of these findings, a predictive optimal controller for the combined usage of active and passive thermal storage that accounts for uncertainty in predictive variables and model mismatch was developed and verified in the context of the presented work. Once the supervisory controller was implemented in the laboratory setting, the test facility was controlled by the optimizer in real time, which to the authors' knowledge has not been done before. This paper describes the implementation of the real-time control strategy and evaluates its benefits with respect to HVAC energy consumption and cost reduction. In addition, model accuracy and constraint compliance will be examined. The paper concludes with a re-creation of the experiment in a simulation environment during which previously experienced problems, such as the interruption of the communication to the building automation system, were avoided.

DESCRIPTION OF TEST FACILITY

General Background on ERS

To investigate the potential of the optimal controller, the Energy Resource Station (ERS) operated by the Iowa Energy Center (IEC) is used in this study as the experimental test bed. The ERS is a unique demonstration and test facility wherein laboratory-testing capabilities are combined with real building characteristics. It is capable of simultaneously testing two full-scale commercial building systems side by side with identical thermal loading. The ERS building, a single-story structure with a concrete slab-on-grade, has a height of 4.6 m and a total floor area of 855 m^2, divided into a general area (office space, service rooms, media center, two classrooms, etc.) and two sets of identical test rooms, labeled "A" and "B," adjacent to the general area. Eight test rooms are organized in pairs with three sets of zones having one exterior wall (east, south, and west) and one set that is internal (Price and Smith 2000). Figure 1 presents a layout or the ERS including the four sets of identical test rooms used for the experiment.

Opaque exterior surfaces of the ERS are composed of several layers of construction materials with a thermal mass outside of the insulation. The percentage of the window area to exterior wall area is 15% on the east side, 16% on the west side, 32% on the south, and no windows on the north.

The TRNSYS building model was experimentally validated in the context of a research project conducted by the International Energy Agency Solar Heating and Cooling Program (IEA-SHC) Building Energy Analysis Tools Experts Group, Task 22, Subtask A. Reference IEA-SHC (2001) documents empirical validation testing for models related to the thermal behavior of buildings and commercial HVAC equipment installed in typical commercial buildings. In particular, the motivation for the Iowa ERS validation exercise work was to create a suite of test cases for evaluating the capability of building energy analysis tools to model HVAC system and realistic buildings of commercial construction. One of the building energy analysis tools was TRNSYS, and the ERS TRNSYS model was made available for this research. We ran additional tests to confirm values such as internal gains, COPs of the two chillers, and charge/discharge behavior of the TES system and used these data to enhance the building model.

Primary and Secondary HVAC Systems

The primary pieces of mechanical equipment in a central heating plant are a natural gas-fired boiler and a cooling plant with three nominal 35 kW air-cooled chillers that operate in both chilled-water and ice-making modes. The chilled-water loop is filled with 22% propylene glycol water solution. In addition, the building includes a 440 kWh internal melt ice-on-tube thermal energy storage tank as well as pumps and auxiliary equipment needed to provide cooling. Dis-

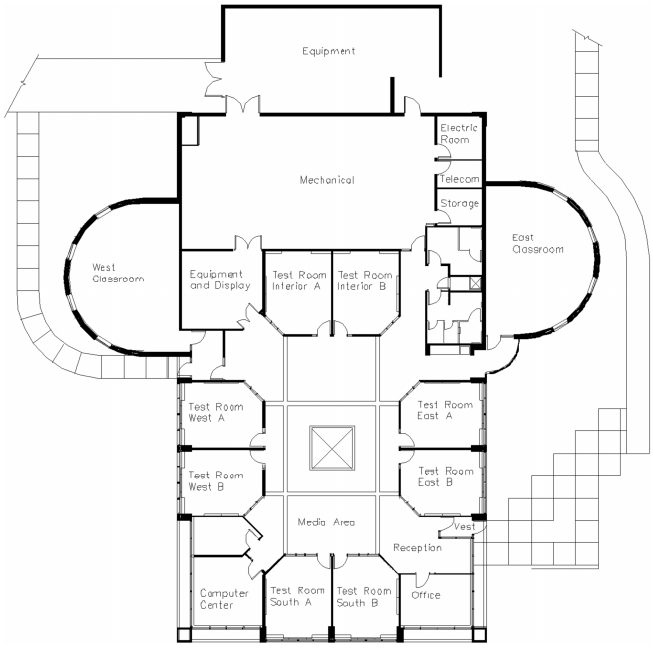

Figure 1. Layout of the test facility, the Energy Resource Station (ERS), Ankeny, Iowa.

trict cooling can be provided by the campus chilled water plant, which was not used in this experiment. Hence, several modes of operation between these sources of cooling are possible in order to supply chilled water to the air-handling units (AHU). A primary-secondary flow arrangement with dedicated constant-volume chiller pumps and secondary variable-flow distribution pumps in the AHU loop using variable-frequency drive (VFD) control is installed in the ERS.

Secondary HVAC systems include three air-handling units (AHU) that condition the building: test rooms A and B are served by two similar single-duct variable air volume (VAV) with reheat AHU systems A and B, and the general area is served by a similar but larger AHU-1. An overhead air distribution system utilizing pressure-independent VAV boxes supplies air to each test room using hydronic or three-stage electrical resistance reheat.

Finally, there is an on-site weather station with measurements of outdoor air dry-bulb temperature, relative humidity, wind speed and direction, atmospheric pressure, total normal incident solar flux, and global horizontal solar flux.

Investigated Test Rooms

Test rooms A and B, each with a net floor area of 24.8 m^2 and carpeted floor, were used for the experiment. The ceiling height is 2.6 m and there is a plenum above the suspended ceiling with a height of 1.7 m. Having the same geometry and construction specifications, the identical pairs of zones A and B experience the same heating and cooling load. Glazing area of the exterior zones consists of double-pane 6.4 mm clear insulating glass and measures 6.9 m^2. During the test, these windows were covered with fully open external blinds. Furthermore, to reduce thermal coupling to the general area, the interior windows between the test rooms and the gen-

eral area were covered with 12.7 mm of dry wall paper. Following the suggestion of Braun et al. (2002) at the ERS, additional mass was added to the interior test rooms A and B in the form of two rows of standard concrete cinder block, 3.05 m long and each stacked three layers high. The walls were located near the middle of each interior room.

Test rooms were unoccupied; however, false internal heat gains can be introduced using baseboard heaters and lights to simulate the occupancy schedule of a typical building. Test rooms A are equipped with two-stage lighting, whereas test rooms B are fitted with dimming electronic ballasts, both with a maximum wattage of 585 W. The baseboard heater at each zone can operate in two stages with a maximum output of 1.8 kW (900 W per stage).

One comfort sensor measuring the air temperature, humidity, and wind speed was placed in the middle of each of the rooms. Conditioned air at a temperature of 13°C was supplied to the test rooms by two ceiling-mounted diffusers in order to maintain the room temperature within a range of 20°C and 24°C during time of occupancy. The interior flow rate throughout the occupied period was characterized by a minimum flow of 94 L/s and a maximum flow of 189 L/s. Finally, all test rooms were kept locked throughout the period of the experiment in order to avoid disturbance and interruptions. These conditions were applied to all eight test rooms.

The ERS is not a particularly good candidate for the use of building thermal mass, as documented by Braun (2003), for two reasons: (1) it is a lightweight single-story structure with a high exterior surface area to volume ratio and (2) significant thermal coupling with the ground, the ambient environment, and the zones adjacent to the test rooms is present. Furthermore, the test zones are not equipped with a representative amount of furniture, and the floor is carpeted, which reduces thermal coupling to the massive structure.

Assumptions for Predictive Optimal Control

The simulated occupied period extends from 8 a.m. to 5 p.m. each day including weekends. During this time, baseboard heaters are applied at one stage (0.9 W) and they are turned off during the remaining hours. Furthermore, one stage of lighting (360 W) is employed from 7 a.m. to 6 p.m. The applied utility rate structure assumes an on-peak electricity rate of $0.20/kWh from 9 a.m. to 7 p.m., and an off-peak electricity rate of $0.05/kWh the remaining hours. Demand charges are not levied. Typically, both energy and demand charges are billed to commercial customers.

Of the available equipment, the HVAC system during the test consists of two chillers, namely, a main and a dedicated precooling chiller, and the ice-based TES system. Charging of the TES tank and meeting on-peak cooling loads was accomplished by the main chiller, which operates in the chilled-water mode with a coefficient of performance (COP) of 2.1 and in the ice-making mode (charging the TES) with a COP of 2.4. The COP relates the amount of heat extracted to the mechanical work input of a refrigeration system. COPs were validated through repeated tests at the ERS. Consequently, meeting cooling loads through the use of ice storage is more attractive from an energy consumption perspective than standard chilled-water operation. The second law of thermodynamics suggests that efficiency drops with a reduction in evaporation temperature; thus, the measured COP values are peculiar to the ERS and not common.

Initial tests investigating conventional control strategies revealed that charging of the TES system takes noticeably more time than estimated by the controller, which was traced back to a significantly reduced chiller capacity to only 50% in the ice-making mode. It is common for refrigeration capacity to decline with a reduction in evaporation temperature. Further, it was determined that the ice storage system behaves very nonlinearly below 20% and above 90% state of charge (SOC) because the measurement of ice level in the tank is highly inaccurate at both very high and very low states of charge. Since the model employed in the predictive optimal control assumes a linear change in SOC with the charging and discharging rates, the SOC

was limited to an available range of 25% to 75% because the actual TES system was sufficiently linear in this range. Effectively, these limits cut the available storage capacity in half to 220 kWh.

A dedicated precooling chiller with a measured COP of 3.4 is assigned to flush the building with cool air during nighttime and, consequently, to precool the building's massive structure and furniture. Both chillers cannot simultaneously supply chilled water to the AHU.

Outdoor air ventilation is governed by a return air temperature economizer that allows for free cooling when the ambient air conditions are favorable. The minimum outdoor intake damper is restricted to a position of 45% open for AHU A and 37.5% open for AHU B to ensure 20% of ventilation air at design airflow conditions.

Simulations and experiments in the same facility conducted by Braun et al. (2002) revealed that there exists significant thermal coupling between the test rooms and the adjacent general area. As a result, there would be significant energy transfer between zones when utilizing different zone temperature strategies. Therefore, the decision was made to condition both—test rooms A with AHU-A and test rooms B with AHU-B—using predictive optimal control, while the general area was conditioned by AHU-1 and its own dedicated chiller under standard nighttime setback control without active storage utilization. The general area and related energy consumption was not considered in this work. As a result, the evaluation of optimal and conventional control strategies is accomplished by comparing measured results for the eight ERS test rooms under optimal control and simulated results for the eight ERS test rooms using the calibrated building model and conventional control. The alternative approach of comparing measured results from test rooms A under optimal control and simultaneously measured results from test rooms B under conventional control was not adopted.

The general area was conditioned from 7 a.m. to 5 p.m. with a zone temperature setpoint of 22°C. During unoccupied periods, temperatures were allowed to float between 15°C and 30°C. Outside air intake for the third AHU serving the general area was controlled by an economizer, restricting the minimum damper position to 10%.

ERS personnel verified that all sensors were sufficiently calibrated and over 750 monitoring points at minute-by-minute intervals were recorded during the experiments.

DESCRIPTION OF IMPLEMENTED PREDICTIVE CONTROL STRATEGY

Overview

In this study a sequential approach to model-based control is employed: (1) short-term forecasting, (2) optimization, and (3) post-processing and control command implementation as shown in Figure 2. A real-time weather station provides the current weather data to the short-term weather predictor. If an on-site weather station is not available, the nearest available weather station operated by the National Weather Service can be used. This predictor provides an improved forecast for the next planning horizon to the optimal controller, which adjusts the control variables in the model according to Figure 3 until convergence is reached. The optimal solution is passed to a post-processor that interprets the optimal results and turns them into commands understood by the building automation system of the facility under control. The building is modeled in the transient systems simulation program TRNSYS (2003), while the general purpose technical computing environment Matlab (2004), including the optimization toolbox, was used to interface with the building simulation program. In the current investigation, the building

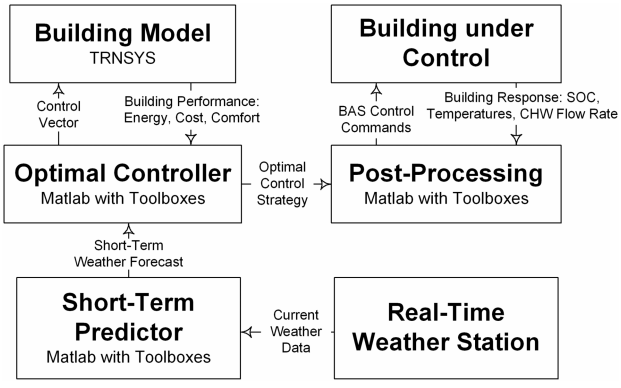

Figure 2. Real-time predictive optimal control schematic.

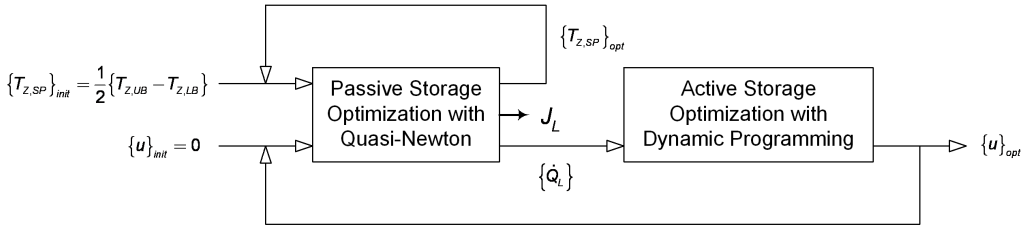

Figure 3. Iterative sequential optimization of utility cost C_L.

model is static and no system identification is employed to estimate the model parameters during operation.

Prediction

A 30-day bin predictor model was found to provide the most accurate weather forecasts for a range of models tested in previous research by Henze et al. (2004b). Predicted variables include ambient air dry-bulb temperature, relative humidity, global solar radiation, and direct normal solar radiation. The assumption underlying the prediction procedure is that the actual time series will exhibit a behavior similar to a reference pattern, developed by rendering bin estimates. For a planning horizon of L = 24 hours, the bin model develops the characteristic profile on the basis of observations collected over the past 30 days. The forecast is made by shifting the L-hour profile such that the predicted value for the current hour k^* coincides with the actual measured value by the weather station at the ERS. Hence, the bin values are computed from

$$\{\hat{X}_t\} = \frac{1}{d}\sum_{n=1}^{d} X_{t-24n}, \qquad t \in [k^*, k^* + L], \tag{1}$$

where d is the number of past days used to compute the bins, and X_t is the observed variable, i.e., either dry-bulb outdoor air temperature, relative humidity, global horizontal insolation, or direct

normal insolation. The 24-hour forecast is handed over to the optimal controller that uses these values among others to estimate the building cooling load profile for the next L hours.

Real-Time Model-Based Predictive Optimal Control

The optimal controller governing the two sources of thermal energy storage can minimize any objective function of choice including total energy consumption, energy cost, occupant discomfort, or a combination of these. In this study, the real-time controller was charged to minimize operating cost for time-of-use differentiated electricity and fixed-cost natural gas by adjusting global zone temperature setpoints $T_{Z,SP}$ for the passive storage and a dimensionless charge/discharge rate u for the active storage. Other objective functions such as occupant discomfort will lead to different optimal strategies.

Optimal control is defined as that control trajectory that minimizes the total monthly utility bill C_m for electricity and heating:

$$J_m = \min C_m = \min\{C_{elec,m} + C_{heat,m}\}, \text{ where}$$
$$C_{elec,m} = \sum_{k=1}^{K_m} r_{e,k} P_k \Delta t_h; \qquad C_{heat,m} = \sum_{k=1}^{K_m} r_h \dot{Q}_{heat,k} \Delta t_h \tag{2}$$

where $r_{e,k}$ is the energy rate for electricity according to the utility tariff in effect for time k, K_m is the number of hours in the current month, Δt_h is a time increment of one hour, r_h is the unit cost of heat delivered, and $\dot{Q}_{heat,k}$ is the heating demand from zone reheat in hour k.

To apply fixed-horizon optimal control to an infinite-horizon problem such as the given real-time control (it could go on indefinitely), closed-loop optimization (CLO) is employed, i.e., the predictive optimal controller carries out an optimization over a predefined planning horizon L (here L = 24 hours), and of the generated optimal strategy, only the first action is executed. At the next time step the process is repeated. The final control strategy of this near-optimal controller over a total period of K steps is thus composed of K initial control actions of K optimal strategies of horizon L, where $L < K$. As an example, assume the proposed controller is used to govern the operation of a building for a typical month of 8760/12 = 730 hours. At every one of these K = 730 hours, the controller plans ahead over the next L = 24 hours but uses only the first of these L-hour control strategies. At the end of this month, the controller implemented a strategy of K = 730 actions, each of which is the first action of all K optimal strategies over L = 24 hours.

By moving the time window of L time steps forward and updating the control strategy after each time step, a new forecast is introduced at each time step and yields a control strategy that is different from the strategy found without taking new forecasts into account. Since we optimize over a planning horizon of L hours, we can only minimize an approximate cost function C_L, which allows for the determination of a near-optimal strategy, whose cumulative utility cost approaches the desired J_m at the end of the billing period.

Initial attempts to simultaneously optimize both active and passive thermal storage in 48 dimensions (24 temperature setpoints and 24 TES charge/discharge rates) led to unacceptably long optimization times and poor solutions resulting from local minima. It was decided to avoid the "curse of dimensionality," i.e., the exponential growth of optimization time with the number of optimization parameters, by separating the passive from the active portion of the real-time control problem. The cost over the next planning horizon C_L is affected by both the active and passive building thermal storage strategy. The choice of zone temperature setpoints will affect the cooling load, which has to be known for the active storage to be controlled properly. Therefore, there is a causal relationship from the passive to the active storage, which allows us to

solve the passive storage first, followed by the optimization of the active thermal storage inventory on the basis of the previously determined optimal building cooling load profile.

Due to the presence of simple upper and lower zone temperature bounds, the passive thermal storage (building mass) component of the control problem proved to be solved effectively with the help of a common implementation of the quasi-Newton method. The use of a direct search method (Nelder-Mead Simplex) led to an excessive number of function evaluations (TRNSYS calls) because of cost penalties arising from bound violations. The active storage (TES) optimization problem is characterized by complex and nonlinear constraints yet simple state transitions (change of TES inventory). This class of problem is most readily solved using dynamic programming (DP). The optimal AHU cooling loads calculated by the complete TRNSYS building model are passed to a central plant model used by the DP algorithm. This small submodel mirrors the relevant parameters of the complete building model and is suitable for the repeated calls from the DP routine.

Figure 3 illustrates how the minimal utility cost J_L over planning horizon L, i.e., the minimum of C_L, is determined. At time zero and starting with initial zone temperature setpoints $\{T_{Z,SP}\}_{init}$ halfway between the upper and lower bounds and no active storage utilization $\{u\}_{init} = 0$, the passive storage inventory is optimized to minimize C_L. As a result, the optimal building cooling load profile is computed and handed over to the active storage optimization, which calculates an optimal TES charge/discharge strategy. In a second pass, the optimal active storage utilization strategy and the previously found optimal zone temperature setpoint profile are employed to determine the new optimal zone temperature setpoint profile and optimal utility cost J_L. This cycle is repeated until the optimal cost J_L converges. Typically, convergence is attained after two to three iterations. Previously optimal solutions are stored as starting values for subsequent optimizations to reduce execution time. We refer to Henze et al. (2004a) for a detailed description of the model-based predictive optimal controller for building thermal storage inventory.

At each time step k^*, the model-based controller derives the following four operational parameters for the active TES system from the optimal charge/discharge rate u_{k*}: charging load for the main chiller (Q_{charge}), discharging load for the active TES system ($Q_{discharge}$), remaining cooling load for the main chiller (Q_{main}), and cooling load met by the precooling chiller ($Q_{precool}$). Rules incorporated in the building model ensure that (a) charging and discharging cannot occur simultaneously; (b) when the main chiller charges the active TES tank, any building cooling load has to be met by the precooling chiller; and (c) when the TES system is discharged, any remaining building cooling load has to be met by the main chiller.

Post-Processing

A post-processing computer program was developed for the ERS test facility to translate the optimal results produced by the model-based controller into commands that can be understood by the building automation system and executed by the ERS HVAC system. The post-processing program sketched in Figure 4 sequentially executes the following operations every hour: (a) set up of a communication channel between the optimal controller environment and the BAS using a proprietary general-purpose communication software, (b) reading the optimal results from the optimal control and the required values from the BAS, (c) conversion of optimal results into control commands, and (d) sending the new control commands to the BAS.

The following post-processing procedure is executed: First, the room air temperature setpoint is sent directly to the BAS. Next, the cooling discharge rate of the TES is accomplished by sending the TES leaving water temperature as a setpoint for the TES mixing valve local loop control.

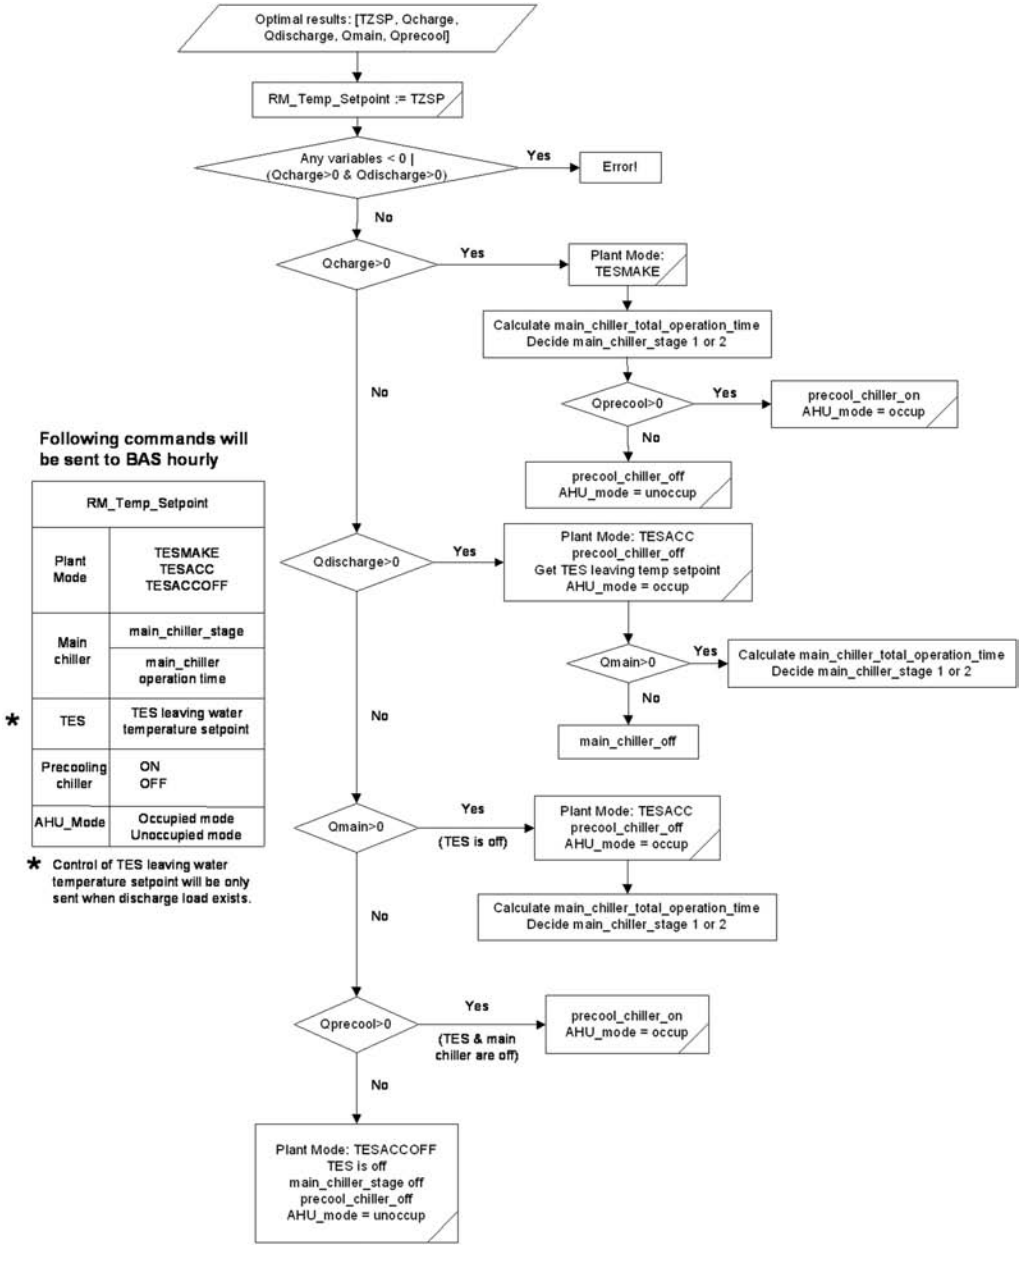

Figure 4. Post-processing program flow chart.

The leaving water temperature $T_{LW,TES}$ is calculated from $\dot{Q}_{discharge} = \dot{m}c_p(T_{EW,TES} - T_{LW,TES})$, where flow rate ($\dot{m}$) and entering water temperature ($T_{EW,TES}$) are read from the BAS.

Main chiller operation is possible in one of two stages (50% and 100% cooling capacity). During occupancy, the cooling output of the main chiller is accomplished by pulse width modulation (PWM). The PWM algorithm translates the optimal control result Q_{main} into a chiller

stage and minutes of operating time during the next hour. The PWM time period was 20 minutes. Thus, the total chiller operating time is distributed over three PWM periods per hour. Operational constraints have also been taken into account. For example, there are at least five minutes between two periods of chiller operation in order to avoid the chiller cycling too frequently. Moreover, if the calculated main chiller load results in an operating time of less than five minutes, than the chiller will operate five minutes. The precooling chiller is operated by the existing on-off control algorithm without PWM to maintain the global zone temperature setpoint $T_{Z,SP}$ in the building.

Control Command Execution

Proper ERS test facility operation is facilitated by the building automation system selecting one of five plant modes that had been previously defined and had been modified to accommodate the predictive optimal control. A plant mode has to be selected before a command from the post-processing program can be sent to the HVAC system. Three of the available plant modes used in the context of the experiment were: TESMAKE plant mode, which represents charging of the active TES system; TESACC plant mode, in which the test rooms are conditioned by the main chiller and/or the active TES system; and TESACCOFF, where the entire HVAC system is turned off. In addition, the BAS enables the air-handling unit fans based on the existence of a cooling load on the AHU cooling coils. In summary, the post-processing program obtains the optimal results from the controller, converts the values into comprehensible commands, selects a plant mode, and forwards the commands using the communication channel to the BAS.

DESCRIPTION OF CONVENTIONAL CONTROL STRATEGIES

Before the real-time control experiments were conducted, two additional tests—a reference case and a base case—were carried out. Both tests were required to calibrate the simulation model with respect to the building thermal response, the mechanical systems, and the operational schedules. Figure 5 shows the system configuration for the reference case and the base case. The active TES system is bypassed in the reference case.

Reference Case

The reference case represents the standard case of a cooling system with one sufficiently sized chiller (35 kW) that serves the air-handling units (AHU) A and B during occupancy, with nighttime setback during unoccupied periods and with neither active nor passive building thermal storage utilization. The test was performed under the same schedule of occupancy and temperature setpoints as in the experiment: during occupied hours the zone temperature cooling setpoint was 24°C and the heating setpoint was 20°C, while the space temperature was allowed to float within the range of 15°C to 30°C during unoccupied periods.

Base Case

In the base case test, the zone temperature setpoints were identical to the reference case test; however, the main chiller was downsized to only 25% of its nominal capacity (8.8 kW) in order to achieve a sufficiently high number of on-peak hours during which the TES system is used. During those periods when the cooling load exceeds the reduced chiller capacity, the remaining cooling is taken from the active TES system. At night the active TES system is recharged with full capacity of 35 kW to the upper inventory level of 75% state-of-charge. As in the reference case, the passive building thermal storage inventory is not utilized.

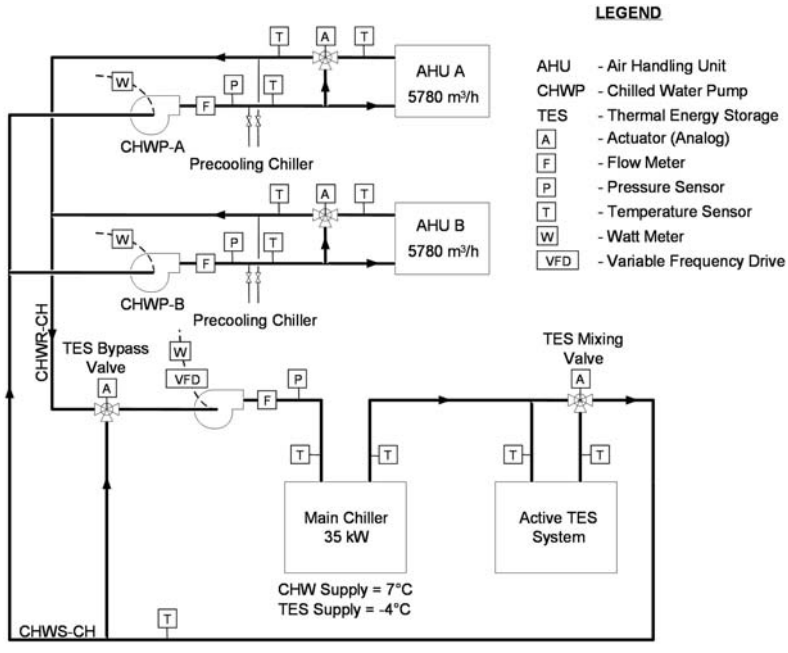

Figure 5. System schematic of ERS HVAC equipment.

RESULTS

Modeling Accuracy

Model-based predictive control relies on a mathematical model to emulate the system under control. It is therefore essential that the model comes as close as possible to reality. The real-time optimal control experiment was executed over a period of five days (128 hours), from midnight on September 13 until 8 a.m. on September 18, 2003. In order to verify the experimental results, the accuracy of the building model must first be evaluated. Minute-by-minute measurements taken at the test facility from September 14 to September 17, 2003 (subscript ERS), are compared to the optimal results determined by the controller during the real-time control experiment (subscript ActPredOpt, actual predictive optimal results). Preliminary tests conducted during June and July of 2003 facilitated the calibration of the simulation model. As supported by evidence provided below, the accuracy of the building model was adequately high. Moreover, the post-processing program correctly translated and transferred the optimal control results, and the HVAC systems and components were successfully orchestrated.

The experiment experienced two interruptions due to server crashes that made the communication channel to the building automation system unavailable on September 16 at 18:00 and on September 17 at 12:00 noon. During these interruptions, the BAS returned zero values for all properties, and deviations between model and measured data necessarily occurred.

Figure 6 compares the total simulated and measured AHU cooling loads of system B. It can be seen that the measured and modeled values are in good agreement, with the exception of a three-hour time lag during September 14. However, there are some peak cooling loads that are not represented well by the model-based controller. The AHU cooling loads are due to internal

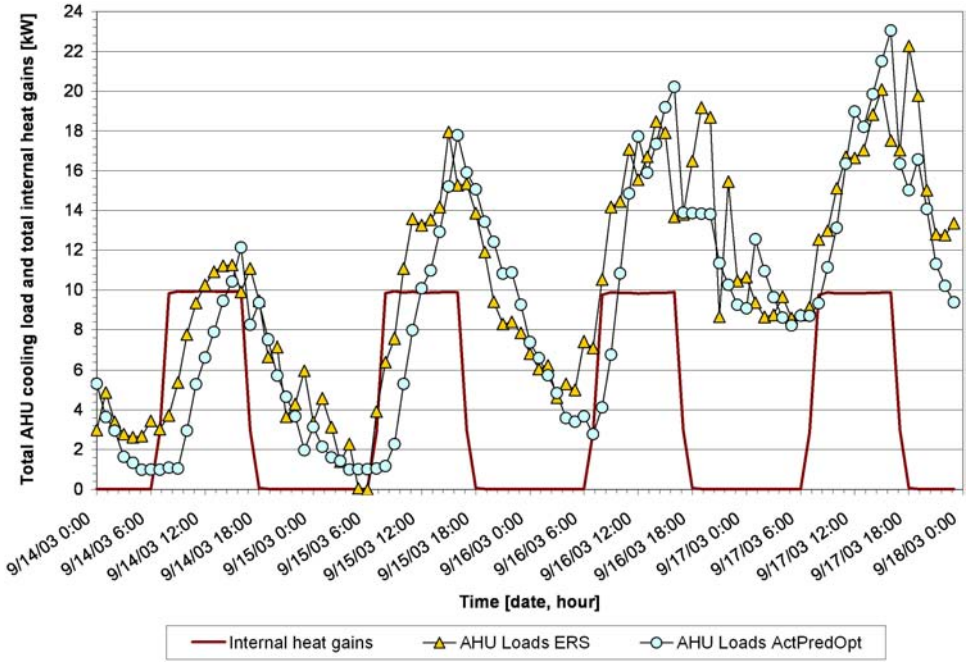

Figure 6. Measured and modeled cooling loads of AHU B (kW) as well as measured internal heat gains.

heat gains from baseboard heaters and lighting (which were chosen to be constant throughout the test days), solar gains, and the required intake of ventilation air.

Figure 7 illustrates the measured and modeled charging and discharging performance of the active TES system for September 15, 2003. It can be observed that the charging performance is accurately modeled; however, the discharge performance is modeled less precisely. Still, the discharge trend is captured well by the model used in the predictive optimal controller. While the charge/discharge performance appears to be adequately modeled, the profiles of the state-of-charge do not match well. Differences of up to 12% of active inventory can be observed and are attributed (a) to the compounded differences in the charge/discharge performance and (b) to the poor accuracy of the inventory sensor (claimed to be ±5%). To eliminate the discrepancy between the simulated and measured values of state-of-charge, the SOC was measured throughout the testing period and updated seven times in the simulation environment. This procedure implies that the SOC sensor reads inventory levels accurately, which is not the case. Thus, the compounded effect of modeling mismatch in the charge/discharge process is eliminated, yet at the same time the low SOC sensor accuracy is introduced to the optimization environment.

Figure 8 illustrates the measured and simulated performance of the precooling chiller. Compared to the main chiller cooling profiles (not shown), the modeling accuracy for the precooling chiller is inferior. Unlike the active TES system and the main chiller, the precooling chiller was not controlled to maintain a particular value of $Q_{precool}$ but to maintain a global zone temperature setpoint $T_{Z,SP}$.

Figure 9 shows the ambient air temperature, the average room air temperature, and the upper and lower temperature bounds selected for the operation of the ERS, and it confirms that the

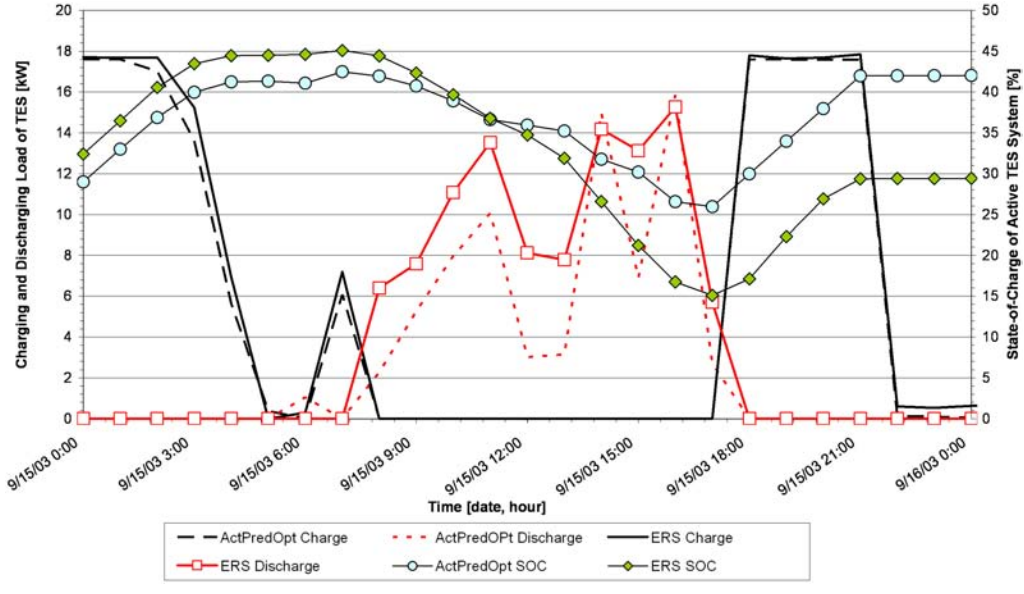

Figure 7. Measured and simulated charging and discharging load (kW) and state-of-charge (%).

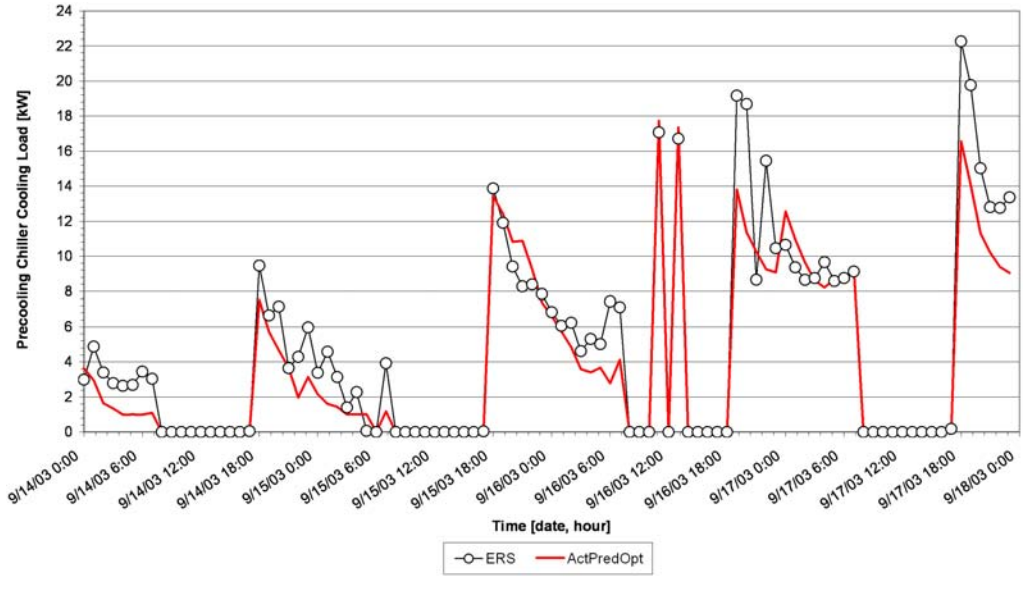

Figure 8. Measured and simulated precooling chiller cooling load (kW).

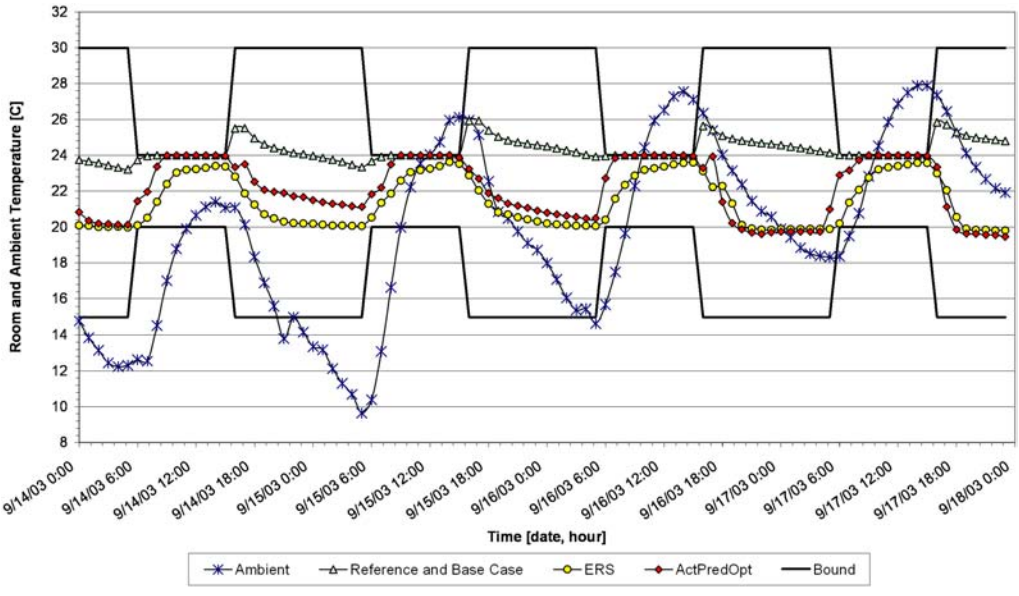

Figure 9. Average test room air temperature and ambient air temperature (°C).

model-based predictive optimal control complies with the operational constraints imposed in the model. During the real-time control experiment, the optimal controller decided on substantial nighttime precooling down to 20°C averaged over all test rooms. Had the controller decided to precool the building even lower, a need for heating would have occurred at the onset of the occupied period. When the temperatures were allowed to float as in the reference and base cases, the average test room temperature rises above 26°C during unoccupied periods. During occupied periods, the room temperature stayed within the required comfort range for all three control strategies investigated in this study.

The investigation of the active TES system state-of-charge data revealed that the simulated values remained consistently within the lower and upper bounds of 25% and 75%, respectively, while the measured TES inventory falls below the 25% mark due to a nonlinear discharging performance.

Average COPs of the main and precooling chillers were recalculated based on the data collected during the real-time control test in September of 2003. It was confirmed that the measured COPs deviated from the values in the building model by no more than ±0.1. Schedules for occupancy and the HVAC system as implemented in the building model and the building automation system proved to match identically.

Energy and Cost Savings Performance

As mentioned above, changes in energy consumption and utility cost will be expressed relative to a simulated reference case or base case using the same building model and the same weather data as occurred during the real-time control tests. A central performance metric for all cases is the utility cost for operating the entire HVAC system over a selected time horizon of four days, September 14-17, 2003. Results for September 13 were not considered to account for the transition from the uncontrolled to the controlled condition.

Raw Energy and Cost Savings Performance. The electrical utility rate structure features time-of-use differentiated energy charges, yet no demand charge is levied. On-peak periods extend daily from 9 a.m. to 7 p.m., with off-peak encompassing all remaining hours. Building occupancy is from 8 a.m. to 5 p.m. While the optimizer accounted for reheat energy at the VAV terminal boxes, no reheat was required in any of the test periods. Consequently, the discussion focuses on electrical energy consumption and costs.

Results in this section show measurements at the ERS during the real-time-control test and simulation results for the reference and base cases for four days. In this time period, the outdoor air temperature ranges from 10°C early in the morning to 27°C at 6 p.m. Figure 9 shows the increasing trend of average ambient air temperature over the course of the real-time control experiment.

In order to evaluate the results of the optimal control strategy with respect to cost and energy changes, the following tables and figures provide measured and calculated data for the following cases: (a) reference case under nighttime setback, (b) base case under chiller-priority control (labeled Base Case 87%), (c) the data measured at the facility (labeled ERS), and (d) the simulated data calculated during real-time control (labeled ActPredOpt). Both cases utilizing the active TES system, i.e., base and real-time control cases, started with an initial TES state-of-charge of 30%.

The original building model assumed a perfect, i.e., loss-free active TES system. This implies that 100% of the charging cooling load, Q_{charge}, is deposited in the storage tank. From measurements it was concluded that only about 87% of the cooling produced by the main chiller during the charging process contributes to changes of the active inventory storage; 13% is lost due to heat gains in the chilled-water distribution system and thermal transmission through the tank skin. For each charging period (five to seven hours per night) during the experiment, the change of the ice storage inventory was divided by the cumulative charging load that occurred over the same time period. The efficiency value of $\eta = 87\%$ was arrived at by averaging these five ratios.

$$\eta = \frac{1}{5} \sum_{p=1}^{5} \frac{\dfrac{(SOC_{final,p} - SOC_{initial,p})}{H_p}}{\sum_{h=1} Q_{charge,h,p}} \tag{3}$$

where SOC is the active TES system state-of-charge, Q_{charge} is the rate of heat extraction from the TES tank, p denotes the charging period, and H_p is the number of hours in charging period p. While the building model used for the ActPredOpt case assumed a perfect charging efficiency, the effect of heat gains and transmission losses on the state-of-charge was accounted for by repeatedly updating the SOC used in the simulation with measured SOC values at the test facility.

For an electrical utility rate structure without demand charges, daily profiles of HVAC utility costs can be plotted for the main chiller, the precooling chiller, and chilled water pumps. Figure 10 illustrates the HVAC electrical energy consumption for September 17. The area under each curve represents the total daily operating consumption. It can be seen that the reference case incurs the highest on-peak demand but, as a result of nighttime setback, does not consume any energy during the unoccupied period. Results for the base case manifested the second highest energy demand during the on-peak period. Although the on-peak energy consumption for the base case is significantly less than that for the reference case, the on-peak consumption is greater than that under optimal control. During the day shown in Figure 10, the building cooling load was moderate and only a small contribution from the active TES system was required in the base case. Consequently, only four hours of recharging were needed. Model-based predictive optimal

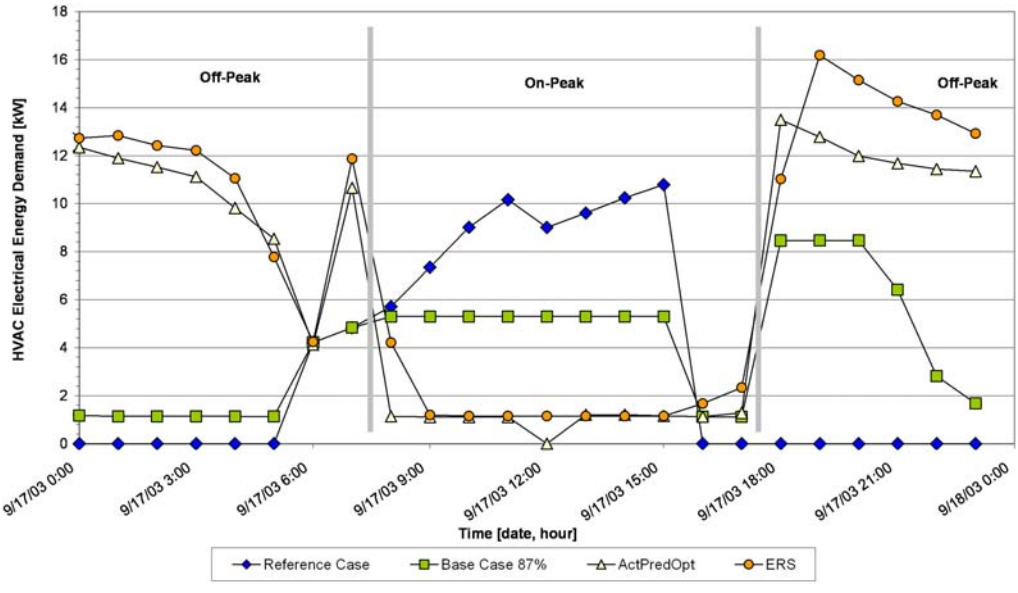

Figure 10. HVAC electrical energy consumption (kWh).

control successfully shifted building cooling loads to off-peak periods, and an excellent match between calculated (ActPredOpt) and measured (ERS) cooling load data can be observed.

The simultaneous utilization of active and passive building thermal storage inventory led to near-zero cooling-related electrical energy consumption during the on-peak period. Energy consumption during the on-peak period is caused by the chilled water pump operating continuously. During the off-peak period high values of energy consumption can be observed, which are due to precooling of the building structure (passive) and charging of the active TES system. It is obvious that on-peak energy consumption is reduced at the expense of increased off-peak energy consumption in the presence of an energy rate ratio of 4:1. Reducing on-peak electrical demand is a side effect of shifting expensive on-peak cooling loads to off-peak periods for an optimal controller, minimizing electrical energy cost without a demand charge.

Although the building model was extensively calibrated, the differences between measured and modeled hourly HVAC electrical energy consumption compounded to significant differences on a daily basis.

Table 1 provides the daily utility cost savings achieved during the experimental period. Relative to the reference case, measured savings of about 5% in total HVAC utility costs were achieved in the ERS, and about 10% of modeled savings (ActPredOpt). Compared to the base case with 87% charging efficiency, cost increases of about 7% and 1.4% were achieved for the ERS and the simulation, respectively. As shown, there are significant variations in the cost changes from one day to another relative to the reference and the base cases. This inconsistent pattern was caused by a number of reasons discussed in the next section. Only one week of test time was available in the ERS; longer test periods would lead to more representative results.

Corrected Energy and Cost Saving Performance. *Motivation for Correcting the Measured Results*—Previous research (Henze et al. 2004, 2004b) revealed that given strong load-shifting incentives, the benefits of the investigated predictive optimal control may be substantial. Therefore, we expected moderate daily savings, less fluctuation from day to day, and substantial

Table 1. Changes of Daily HVAC Electrical Utility Cost of the Optimal Control Strategy Compared to the Reference Case and the Base Case 87% (%)

	14-Sep	15-Sep	16-Sep	17-Sep	Cumulative
Cost changes relative to Reference Case (%)					
ERS	–19.7	+11.5	+9.1	–21.2	+5.0
ActPredOpt	–32.2	+20.4	+7.1	–34.4	–9.9
Cost changes relative to Base Case 87% (%)					
ERS	–30.3	+17.3	+32.1	+1.3	+7.0
ActPredOpt	–41.1	+26.6	+29.7	–15.6	+1.4

cumulative savings. However, control systems often do not work as desired in practice, even for mature systems, which has motivated the current interest in building commissioning. The optimal controller did not work as intended for two reasons, one due to data transfer and the other to a problem with the optimizer's algorithms (if getting stuck in a local minimum is an algorithm problem). Correcting these problems yielded larger savings, more representative of those achievable by a well-tuned and trouble-free control system. Corrected data discussed will be denoted by "ERS*cor*" and "ActPredOpt*cor*."

Description of Experimental Problems—There were two experimental problems encountered during the tests: First, invalid data were produced by the building model during two hours of the experiment caused by the interruptions of the communication channel on September 13 and 17. During one interruption, the building model received NaN data and crashed; during the second, the optimizer computed results that could not be transferred to the BAS. These erroneous data were eliminated by interpolating between the valid adjacent data points.

Second, after three hours, suboptimal solutions were found by the optimizer, which are believed to be due to a relaxed convergence criterion. As a result, main and precooling chiller activity occurred for three hours during the on-peak period of September 16 and drastically increased the electrical energy costs for that day. Precooling chiller activity can be observed for September 16 at 12 noon and 2 p.m. in Figure 8. The controller requested the main chiller to charge the TES with a very small charging load (not shown). As a result, the precooling chiller had to meet the daytime AHU cooling loads. What was the reason for peculiar behavior? At any point in time, meeting a cooling load is least expensive by discharging the active TES system (only pump energy is incurred), next by using the precooling chiller (COP = 3.4), and finally by using the main chiller in chilled-water mode (COP = 2.1). Since both chillers cannot operate at the same time, the optimizer decided to charge an insignificant amount in order to be able to use the precooling chiller to meet the on-peak cooling loads. During these hours, the optimal controller was likely caught in a local minimum; thus, it selected a suboptimal control strategy. In a typical building, the COP relationships are expected to be reversed and the spurious on-peak precooling chiller operation is not expected. These experimental defects affected both the measured and the simulated raw cost data, as shown in Table 1.

Elimination of Experimental Defects—In order to assess the potential of the model-based predictive optimal controller, measured and simulated raw data can be manually modified to account for the interruptions and spurious precooling chiller activity. In addition, the experiment in a simulation environment using the same building model, weather data, and initial state-of-charge (labeled RecPredOpt) was repeated and compared with the manually modified simulated data (labeled ActPredOpt*cor*). The expectation was that after removing the experimental defects from the raw data and repeating the simulation without the problem of local min-

ima, the results should match closely. Indeed, a repeated simulation run did not produce the same idiosyncrasies with respect to the precooling chiller operation, as can be seen in Figure 11.

The re-created experiment determines the cost savings one may have obtained without interruptions and local minima. Interestingly, there are minor differences between the results collected for the real-time optimization ActPredOpt*cor* and the recreated optimization RecPredOpt. Obviously, the controller does not find exactly the same optimal solutions, which can be attributed to the convergence criterion of the optimizer.

Cost Comparison Using the Modified Measured Results and the Re-created Simulation—Table 2 compares the daily savings of the corrected measurements at the ERS and of the corrected real-time simulation with the reference and base cases. It can be seen that the values for September 16 and 17 differ greatly from the cost calculation involving the raw data, as shown in Table 1. After the removal of the erroneous data, cost savings of 13% for the ERS data are obtained compared to the reference case and savings of 2% relative to the base case. The corrected real-time optimal results reveal higher cost savings as well. Cumulative savings of 18% are obtained when compared to the reference case and 7% are obtained relative to the base case. When comparing the re-created optimal results without local minima complications against the reference case, the same 18% savings as for corrected simulation results were obtained. The actual control achieved lower savings by about 5% than the simulation because the building model assumes chiller COP values to be constant. The COP values were identified during the model calibration process. In reality, however, part-load and ambient effects significantly affect the actual chiller COP values and lead to higher chiller power consumption than the model predicts. Future models will have to account for part-load performance and ambient conditions in the description of chiller performance.

The re-created simulation did not require any updates of the SOC values since it did not occur in real time and actual SOC data were not available. In order for the comparison of the re-created optimal and simulated base case to be valid, both have to use the same active TES model. We decided to assume a perfectly efficient charging process in the TES system for this comparison. The comparison yielded cost savings of about 7%.

Consideration of AHU Fan Power Consumption. *Motivation for Neglecting AHU Fan Operation*—Preliminary tests before the September 13-18 period had revealed that the global optimization of both active and passive building thermal storage inventory led to prohibitively long calculation times and inferior, i.e., often suboptimal, solutions. In response, we adopted the iterative sequential optimization approach depicted in Figure 3. This decision required the plant models of the passive and active optimization steps to be identical. To allow for easy plant model calibration, it was decided to include a simplified HVAC plant model characterized by constant COPs in each mode of operation excluding the operation and energy consumption of the fans.

Analysis of the measured data revealed that the fan energy consumption cannot be neglected and that fan operation has a significant impact on the decisions of the model-based predictive optimal controller. Therefore, the discussion of energy consumption and cost performance is now extended to take into account the fan power consumption and to highlight the differences in the optimal control decisions with and without fans.

Results with AHU Fan Power Consumption—On the basis of measured fan power consumption data, the supply and return fan electrical power consumption for AHUs A and B was approximated with second-order polynomials and integrated those in the building model. The simulated results, shown in Figure 12 below, present the hourly HVAC electrical demand on September 17 for the reference case, the base case, the measured data ERS*cor*, and the repeated optimal results RecPredOpt. The energy required by the reference case and the base case increased by the energy consumption of the fans during the occupied period.

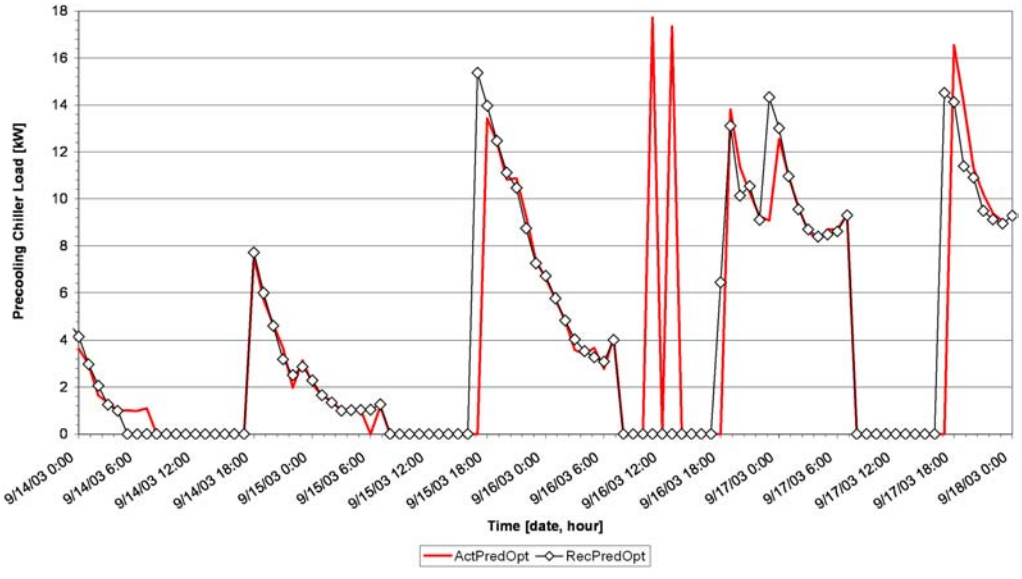

Figure 11. Simulated precooling chiller load as determined during the real-time experiment (ActPredOpt) and during recreated experiment (RecPredOpt) (kW).

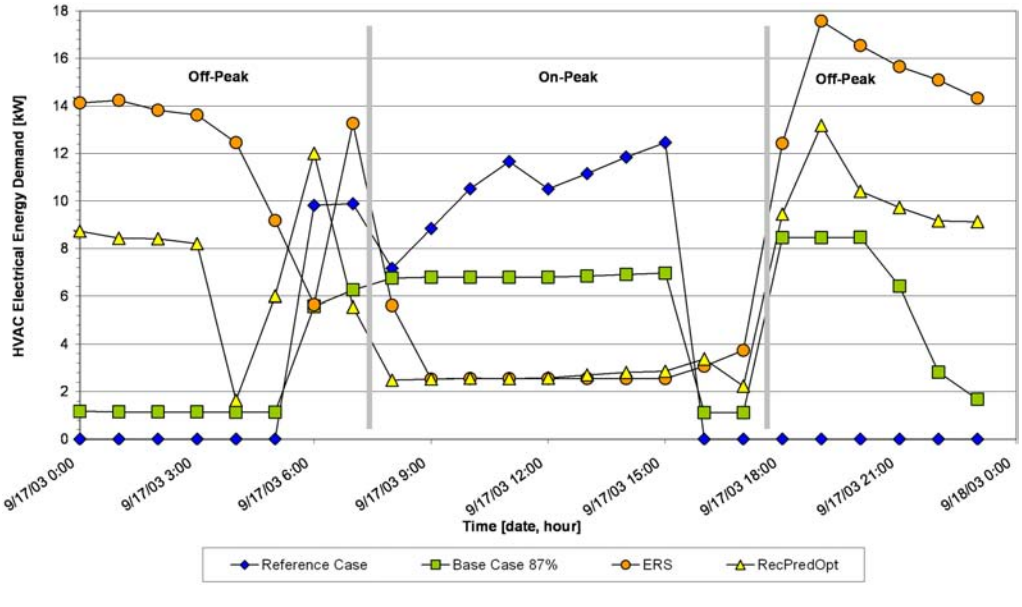

Figure 12. HVAC electrical demand including fan operation (kW).

Table 2. Changes of Daily Corrected HVAC Electrical Utility Cost of the Optimal Control Strategy Compared to the Reference Case and the Base Case (%)

	14-Sep	15-Sep	16-Sep	17-Sep	Cumulative
Cost changes relative to Reference Case (%)					
ERS*cor*	–19.7	+11.6	–19.0	–22.4	–13.6
ActPredOpt*cor*	–32.8	+19.8	–19.6	–34.3	–18.0
RecPredOpt	–35.3	+13.2	–18.2	–30.5	–18.2
Cost changes relative to Base Case 87% (%)					
ERS*cor*	–30.2	+17.4	–1.9	+1.2	–2.2
ActPredOpt*cor*	–41.7	+26.0	–2.6	–14.2	–7.2
Cost changes relative to Base Case 100% (%)					
RecPredOpt	–42.9	+19.0	0	–8.3	–6.6

Table 3. Changes of Daily Corrected HVAC Electrical Utility Cost of the Optimal Control Strategy Compared to the Reference Case and the Base Case Including the Fan Power Consumption (%)

	14-Sep	15-Sep	16-Sep	17-Sep	Cumulative
Cost changes compared to the Reference Case (%)					
ERS*cor*	–7.9	+13.5	–11.4	–13.2	–5.6
RecPredOpt	–30.0	–13.9	–26.4	–36.4	–27.3
Cost changes compared to Base Case 87% (%)					
ERS*cor*	–11.6	+25.4	+7.5	+9.6	+8.3
Cost changes compared to Base Case 100% (%)					
RecPredOpt	–32.9	–4.9	–10.7	–19.8	–16.7

As shown in Table 3, using the new plant model, the optimizer in a re-created experiment decides to make less use of the passive building thermal storage inventory, i.e., less precooling, during the night and, as a result, saves a notable 27% and 17% of electrical utility costs relative to the reference and base cases, respectively.

The real-time control experiment was conducted governed by a model-based predictive optimal controller that did not account for AHU fan power consumption. If one compares the measured total HVAC electrical energy consumption ERS*cor* with the fan consumption added to the reference and base cases using the modified plant model, the savings are reduced from 13.6% to 5.6% for the reference case and from 2.2% savings to cost increases of 8.3%. An important lesson learned was that the consideration of fan power consumption is mandatory for a successful implementation of passive thermal storage utilization. The optimal active TES system control strategy was not materially affected by the inclusion of the AHU fans in the plant model.

Energy consumption of the HVAC systems including fans relative to the reference case, increased by 6% for the base case and by 36% for the predictive optimal case.

SUMMARY, CONCLUSIONS, AND FUTURE WORK

This work investigates the demonstration of model-based predictive optimal control for active and passive building thermal storage inventory in a test facility in real time using time-of-use differentiated electricity prices with a ratio of 4:1 without demand charges. The novel supervi-

sory controller successfully executed a three-step procedure consisting of (1) short-term weather prediction, (2) optimization of control strategy over the next planning horizon using a calibrated building model, and (3) post-processing of the optimal strategy to yield a control command for the current time step that can be executed in the test facility. The primary and secondary building mechanical systems, consisting of two air-cooled chillers, an ice-based thermal energy storage system, two identical air-handling units, and auxiliary equipment, were orchestrated by the model-based predictive optimal controller in real time while observing comfort and operational constraints. The author believes that this has not been accomplished before.

Provided the building model includes a realistic plant model including AHU fan power consumption, the simulation results show that when the optimal controller is given imperfect weather forecasts and when the building model used for planning control strategies does not match the actual building perfectly, utility costs savings in the ERS can be on the order of 17% relative to the base case and 27% relative to the reference case.

However, the actual savings achieved by the optimal controller and measured at the ERS during the period of September 14-17, 2003, were disappointingly low. During the four-day testing period in 2003, the communication channel was interrupted twice and the optimizer was stuck in local minimum for three hours. Moreover, the assumption of constant COP values for the two chillers without accounting for part-load performance and ambient conditions, as well as the lack of modeling the VAV fan power consumption, proved to be inadequate.

The savings associated with passive building thermal storage inventory are also small because the ERS test facility is not an ideal candidate for the investigated control technology. The building structure is of lightweight construction, the test rooms are unfurnished, and significant thermal coupling exists between controlled test rooms and an uncontrolled adjacent area. Finally, the facility's central plant revealed the idiosyncratic behavior that the chiller operation in the ice-making mode was more energy efficient (COP = 2.4) than in the chilled-water mode (COP = 2.1).

The central conclusions of this article are:

1. Substantial utility costs savings relative to conventional building operation may be achieved even when the optimal controller is given imperfect weather forecasts and when the building model used for planning control strategies does not match the actual building perfectly.
2. A commercial building under model-based predictive optimal control must lend itself to passive storage utilization, i.e., have a significant amount of thermal capacitance available, which implies a heavyweight construction.
3. The building model used for model-based predictive optimal control must include a realistic model for the central chilled water plant and the secondary HVAC systems, including part-load performance, correction for off-design and ambient conditions, as well as fan power consumption.

Field experimentation is now underway in a suitable commercial building with sufficient thermal mass, an active TES system, and a climate conducive to passive storage utilization over a longer testing period to support the laboratory findings presented in this study. More work is necessary to address the practical problems of incorporating the proposed system, such as avoiding local minima in the optimization algorithm and developing fail-safe routines in case the communication between the optimizer and the building automation system breaks down. As an alternative approach, currently ongoing is research that attempts to create an optimal controller for the same control application that does not rely on a model description but learns to carry out the best control decisions based on reinforcement it received in response to past actions.

ACKNOWLEDGMENT

The authors gratefully acknowledge the financial support of this work through US Department of Energy Cooperative Agreement No. DE-FC26-01NT41255 and the outstanding on-site technical support at the Energy Resource Station of the Iowa Energy Center provided by Curtis J. Klaassen, Xiaohui Zhou, and David Perry.

NOMENCLATURE

AHU	=	air-handling unit
ActPredOpt	=	raw actual predictive optimal results during real-time simulation
ActPredOpt*cor*	=	corrected predictive optimal results during real-time simulation
BAS	=	building automation system
Base Case 87%	=	base case under chiller priority with 87% charging efficiency
Base Case 100%	=	base case under chiller priority with 100% charging efficiency
CHWP	=	chilled water pump
CLO	=	closed-loop optimization
COP	=	coefficient of performance for the chiller in ice-making and in chilled water mode
ERS	=	Energy Resource Station, test facility
ERS	=	raw measured data at ERS
ERS*cor*	=	corrected measured data at ERS
IEC	=	Iowa Energy Center
HVAC	=	heating, ventilation, and air conditioning
PWM	=	pulse width modulation
RecPredOpt	=	repeated predictive optimal results
SOC	=	state-of-charge for the inventory in the active thermal storage system
TES	=	active thermal storage system
TESACC	=	plant mode for conditioning the test room
TESACCOFF	=	plant mode for turning off HVAC system
TESMAKE	=	plant mode representing charging of the active TES system
VFD	=	variable-frequency drive
C_L	=	cost function
C_m	=	total monthly utility bill for electricity and heating
H_p	=	number of hours in charging period
J_m	=	utility cost in billing period
J_L	=	optimal utility cost
K_m	=	number of hours in current month
L	=	planning horizon for prediction
Q_{charge}	=	charging load for the main chiller
$Q_{discharge}$	=	discharging load for the active TES system
$Q_{heat,k}$	=	heating demand from zone reheat in hour k
Q_{main}	=	remaining cooling load for the main chiller
$Q_{precool}$	=	cooling load for the precooling chiller
$T_{EW,TES}$	=	entering water temperature
$T_{LW,TES}$	=	leaving water temperature
$T_{z,SP}$	=	global zone temperature setpoint
$\{\hat{X}_t\}$	=	forecasted time series
$\{X_t\}$	=	observed time series
c_p	=	specific heat capacitance
d	=	number of days
k^*	=	current hour
$\dot{m}$	=	flow rate
n	=	amount of values
$r_{e,k}$	=	energy rate for electricity
r_h	=	unit cost of heat delivered
t	=	time
u	=	charge/discharge rate
Δt_h	=	time increment for one hour
η	=	efficiency value for charging the TES

REFERENCES

ARI. 1999. *Statistical Profile of the Air-Conditioning, Refrigeration, and Heating Industry*. Arlington, Va.: American Refrigeration Institute, p. 28.

American Standard, Inc. 1999. *EarthWise Today*, Vol. 24, p. 3. LaCrosse, Wisc.

Arthur D. Little, Inc. 1999. *Guide for Evaluation of Energy Savings Potential*. Prepared for the Office of Building Technology, State and Community Programs (BTS), U.S. Department of Energy.

Braun J.E., T.M. Lawrence, C.J. Klaassen, and J.M. House. 2002. Demonstration of load shifting and peak load reduction with control of building thermal mass. *Proceedings of the 2002 ACEEE Conference on Energy Efficiency in Buildings, Monterey, CA.*

Braun, J.E. 2003. Load control using building thermal mass. *Journal of Solar Energy Engineering* 125(3):292-301. New York: American Society of Mechanical Engineers.

Daryanian, B., L.K. Norford, and R.D. Tabors. 1992. RTP-based energy management systems: Monitoring, communication, and control requirements for buildings under real time pricing. *ASHRAE Transactions* 98(1):1160-1170.

Daryanian, B., and L.K. Norford. 1994. Minimum-cost control of HVAC systems under real-time prices. *Proceedings of Third IEEE Conference on Control Applications*, Glasgow, U.K. (1855-1860).

EIA/DOE. 2002. *Annual Energy Review 2002*. U.S. Department of Energy. URL: http://www.eia.doe.gov/emeu/aer/enduse.html. October 2003. Energy Information Administration.

Henze, G.P., R.H. Dodier, and M. Krarti. 1997. Development of a predictive optimal controller for thermal energy storage systems. *International Journal of HVAC&R Research* 3(3):233-264.

Henze, G.P., C. Felsmann, and G. Knabe. 2004a. Evaluation of optimal control for active and passive building thermal storage. *International Journal of Thermal Sciences* 43(2):173-183.

Henze, G.P., D. Kalz, C. Felsmann, and G. Knabe. 2004b. Impact of forecasting accuracy on predictive optimal control of active and passive building thermal storage inventory. April 2004. *International Journal of HVAC&R Research* 10(2):153-178.

IEA-SHC. 2001. Empirical validation of Iowa energy resource station building energy analysis simulation models: A report of Task 22, Subtask A, Building Energy Analysis Tools, Project A.1, Empirical Validation. http://www.iea-shc.org/task22/reports/Iowa_Energy_Report.pdf.

Liu, S., and G.P. Henze. 2004. Impact of modeling accuracy on predictive optimal control of active and passive building thermal storage inventory. *ASHRAE Transactions* 110(1):151-163.

Matlab. 2004. Using Matlab v7. and Optimization Toolbox User's Guide. The MathWorks, Inc.

NETL/DOE. 2003. *Federal Assistance Solicitation for Energy Efficient Building Equipment and Envelope Technologies Round IV*. PS No. DE-PS26-03NT41635. p. 4. U.S. Department of Energy. National Energy Technology Laboratory.

Price, A.B., and T.F. Smith. 2000. Description of the Iowa Energy Center Resource Station: Facility Update III. Technical Report ME-TFS-00-001, Department of Mechanical Engineering, The University of Iowa, Iowa City.

TRNSYS. 2003. TRNSYS—A transient simulation program. SEL University of Wisconsin—Madison, http://sel.me.wisc.edu/trnsys/Default.htm.