Argonne NATIONAL LABORATORY

ANL/MCS-TM-313

Techno-Economic Evaluation of a Next-Generation Building Energy Management System

Mathematics and Computer Science Division
Technical Memorandum ANL/MCS-TM-313

About Argonne National Laboratory
Argonne is a U.S. Department of Energy laboratory managed by UChicago Argonne, LLC under contract DE-AC02-06CH11357. The Laboratory's main facility is outside Chicago, at 9700 South Cass Avenue, Argonne, Illinois 60439. For information about Argonne and its pioneering science and technology programs, see www.anl.gov.

Availability of This Report
This report is available, at no cost, at http://www.osti.gov/bridge. It is also available on paper to the U.S. Department of Energy and its contractors, for a processing fee, from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
phone (865) 576-8401
fax (865) 576-5728
reports@adonis.osti.gov

ANL/MCS-TM-313

Techno-Economic Evaluation of a Next-Generation Building Energy Management System

by
V. M. Zavala,[1] D. Skow,[2] T. Celinski,[3] and P. Dickinson[3]
[1] Laboratory for Advanced Numerical Simulations,
Mathematics and Computer Science Division, Argonne National Laboratory, USA
[2] Computing and Information Systems Division, Argonne National Laboratory, USA
[3] BuildingIQ, Suite 103/19a Boundary Street, Rushcutters Bay, NSW 2011, Australia

May 2011

Contents

List of Figures

Techno-Economic Evaluation of a Next-Generation Building Energy Management System

Victor M. Zavala*
Laboratory for Advanced Numerical Simulations
Mathematics and Computer Science Division
Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL 60439, USA
Dane Skow†
Computing and Information Systems Division
Argonne National Laboratory, 9700 South Cass Avenue, Argonne, IL 60439, USA
Tom Celinski and Peter Dickinson
BuildingIQ, Suite 103/19a Boundary Street, Rushcutters Bay, NSW 2011, Australia

Abstract

We perform a technological evaluation of BuildingIQ's next-generation energy management (EM) system and present a preliminary energy savings analysis for a commercial-sized building located at Argonne National Laboratory. The EM system uses a model predictive control framework that incorporates black-box machine learning dynamic models.

The system enables the optimization of operational set-points for the HVAC system and for the building zones in a proactive manner by foreseeing weather conditions and electricity prices. In addition, it can foresee dynamic limitations arising from the thermal mass of the building, can perform proactive demand-response, and can trade off between electricity cost, comfort satisfaction, and CO_2 emissions resulting from energy consumption. The black-box models are learned in real-time by using available sensor data and do not require building and HVAC topology information. This approach enables short deployment times, low installation and maintenance costs, and widespread deployment.

Our preliminary studies at Argonne indicate that significant amounts of energy can be saved by exploiting the ambient air to delay the start-up of the HVAC system and by exploiting unoccupied periods to relax the zone set-points. We analyze how to use the system to achieve energy savings of up to 45% with minimal impact on comfort conditions. We also identify several technological limitations of the EM system that restrict its performance and we propose directions of future research.

Keywords: building automation, energy management, proactive, comfort, temperature, set-points, energy costs, CO_2 emissions.

*E-mail: `vzavala@mcs.anl.gov`, Phone: +1 (630) 252-3343.
†E-mail: `skow@anl.gov`, Phone: +1 (630) 252-8724.

1 Introduction

Energy Management (EM) is a high-level supervisory task performed as part of real-time building operations. EM can be performed by a human operator or by an automated EM system interfaced to the building management system (BMS). It is important to distinguish an EM system from the BMS, the latter is an interface providing communication access to the different building components such as the thermostats, humidity and CO_2 sensors, and the units of the heating ventilation and air conditioning (HVAC) system such as the variable air volume (VAV) boxes, air-handling units (AHUs), chillers, cooling towers, and boilers.

The objectives of an EM system are to maximize the operational efficiency, comfort, and economy of the building. To do so, the EM system adjusts the operational conditions (set-points) of the building components as the external and internal conditions change in time, in order to maintain efficient operations. External conditions include the electricity prices, ambient temperature, wind speed, solar radiation, and sun position. Internal conditions include variable heat and CO_2 gains due to occupancy and heat gains due to computing and lighting equipment.

There exists a wide range of EM system solutions that vary in their degree of automation. In most buildings, EM is performed manually by operators who use their experience to adapt the set-points. Some BMS are equipped with logic rules or expert systems [15, 8, 12] aiming at saving energy by adapting the set-points according to the time of the day, occupancy schedules, and ambient conditions. Some limitations of logic-based systems are that they quickly exploit in complexity as the number of decisions increases, they are difficult to tune, and they cannot handle constraints and multiple performance objectives efficiently.

A more efficient and systematic EM strategy, which has been widely studied in academia, is simulation-based (or model-based) optimization [3, 4, 19, 24, 10, 7, 20]. This strategy makes use of a simulation model for the building and HVAC system to compute optimal set-points and is effective in dealing with many decision variables, constraints, and complex objectives. Model-based systems have been shown to achieve significant energy savings in many academic studies but have experienced very low adoption levels in industry. In our opinion, the main reason has been the complexity of existing building simulation tools [13]. In addition, the large number of operational decisions and tight interconnection between units make model-based EM systems computationally intensive and difficult to implement, deploy, and maintain [9]. We believe that the lack of more practical technologies has significantly blocked the widespread deployment of these EM systems.

In this report, we perform a techno-economic analysis of the next-generation EM system of BuildingIQ (BIQ). This system has been commercialized using technology developed at the Commonwealth Scientific and Industrial Research Organisation (CSIRO) in Australia [19]. The system uses a model predictive control (MPC) framework, a technique widely applied in the petrochemical and aerospace industries [16, 22, 2]. Recently, studies have reported significant energy savings in buildings using MPC [11, 1, 24]. A key novelty of BIQ's system is that it couples the MPC framework to black-box machine learning models [18]. In addition, it focuses on the interface between the HVAC system and the VAV boxes. This makes the technology practical, highly flexible, and easy to deploy in a wide range of buildings. For instance, the system does not require knowledge of the HVAC layout and building topology. We analyze these and some other features that make the system attractive for large-scale deployment and describe an on-going implementation at Argonne National Laboratory. We use the gained operational experience to analyze the technical limitations and energy savings bottlenecks of the system and to identify directions of future research. We believe that this operational experience can be valuable to other building automation companies in identifying technological bottlenecks of their EM systems.

The report is structured as follows. In Section 2 we provide background information on EM and building automation. In Section 3 we present our position on the ideal features of a next-generation EM system. In Section 4 we provide a high-level description of the technology behind BIQ's system. In Section 5 we describe the implementation at Argonne and present an energy savings analysis. In Section 6 we perform a technological evaluation of the system. In Section 7 we discuss our next steps and research needs.

2 Energy Management Background

In this section we provide background information on the basic configuration of the interface between the building envelope and the HVAC system. This description will highlight some of the interactions between the building components and the information flow through the automation system.

2.1 Building-HVAC Layout

The layout of modern buildings can be generally partitioned as follows. The building envelope is composed of several zones, each separated into subzones. In Figure 1 we present a typical layout for a given zone in cooling mode. Each subzone is conditioned through a VAV box. The subzone VAV boxes request supply air from the zone AHU in order to keep the temperatures at their predetermined set-points as heat gains from occupants, equipment, solar radiation, and convection change in time. The supply air is low in CO_2 and colder than the subzone temperatures in cooling mode or warmer than the subzone temperature in heating mode. The total cooling or heating load of each subzone can be satisfied by changing the supply temperature or the AHU or the volume of air flowing into the subzone. A zone fan extracts the return air from the subzones and sends it back to the AHU. The return air is higher in CO_2 as a result of occupant emissions and warmer than the subzone temperatures in cooling mode or colder in heating mode.

At each AHU, a fraction of the returned air is dumped by the AHU control system into the environment to purge impurities such as CO_2 and is mixed with ambient air. The resulting mix is cooled down or warmed up to a predetermined supply temperature set-point that matches the current VAV boxes cooling or heating demands. If extra cooling capacity is necessary to reach the supply temperature, this can be obtained from electricity or chilled water in cooling mode, depending on the configuration. In heating mode, any necessary heating capacity can be obtained from electricity or steam, depending on the configuration. If at any time during the day the ambient conditions are such that no additional cooling or heating capacity is needed, the AHU uses energy only to run the fans to mix return and ambient air to purge CO_2 and to keep constant the internal building envelope pressure. Because of this, it is desirable to exploit the ambient air as much as possible.

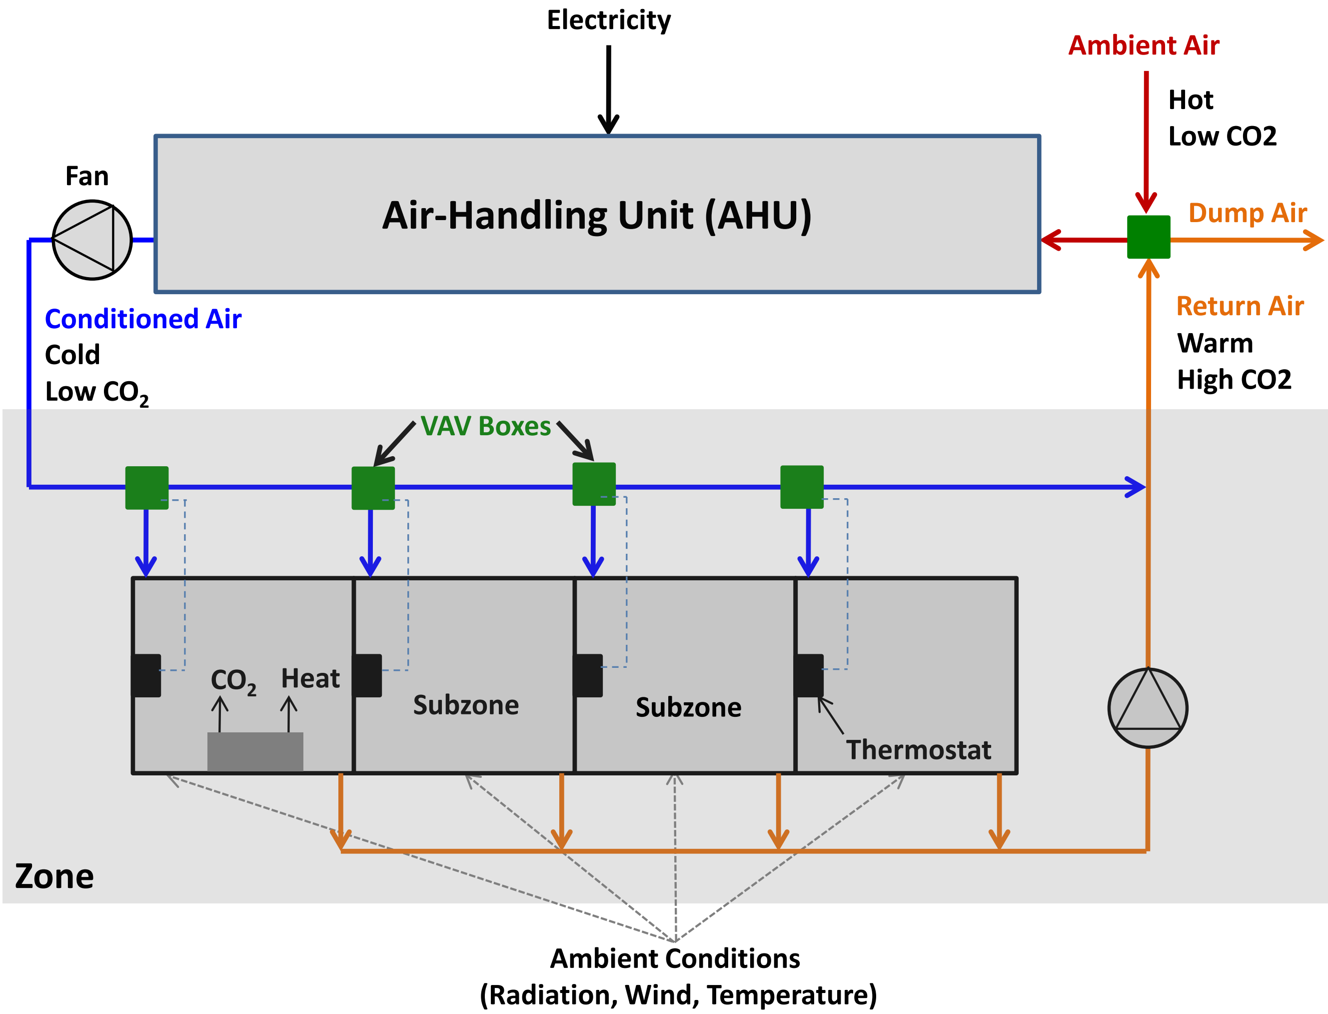

Figure 1: Schematic representation of the interface between the building envelope and the HVAC system (cooling mode).

2.2 Building Automation System Layout

A building automation system comprises two major components: the basic control layer and the supervisory control layer. The basic layer is composed of fast single input-single output controllers such as proportional-integral-derivative controllers. These controllers are used for set-point tracking at the different building zones and subzones and at the equipments units making up the entire HVAC system. The controllers reject high-frequency disturbances (on the range of milliseconds to seconds) in trying to bring a given sensed variable to its corresponding set-point. An example is the thermostat-VAV box loop at each subzone. The thermostat senses the room temperature and sends a signal to the controller in the VAV boxes, which positions the damper to adjust the volume of supply air flowing into the subzone in trying to bring the room temperature to the desired set-point. Similar control loops are encountered in chillers, where the output set-point is the delivery temperature of chilled water and AHUs, where the set-point is the supply air temperature. The interface between the building envelope and the HVAC system sketched in Figure 2.

The supervisory control layer comprises either a human operator or an automated EM system that adjusts the set-points in order to save energy and costs and to ensure appropriate comfort levels as building conditions change. These conditions include weather, occupancy rates, room schedules, time of the day, and energy prices. The set-points are updated at lower frequencies (minutes to hours) and sent to the control loops of the basic layer. The supervisory layer takes care of satisfying *central, building-wide* objectives such as minimizing total energy costs and consumption and maximizing comfort while the basic layer takes care of *locally* tracking the set-points set by the supervisory layer.

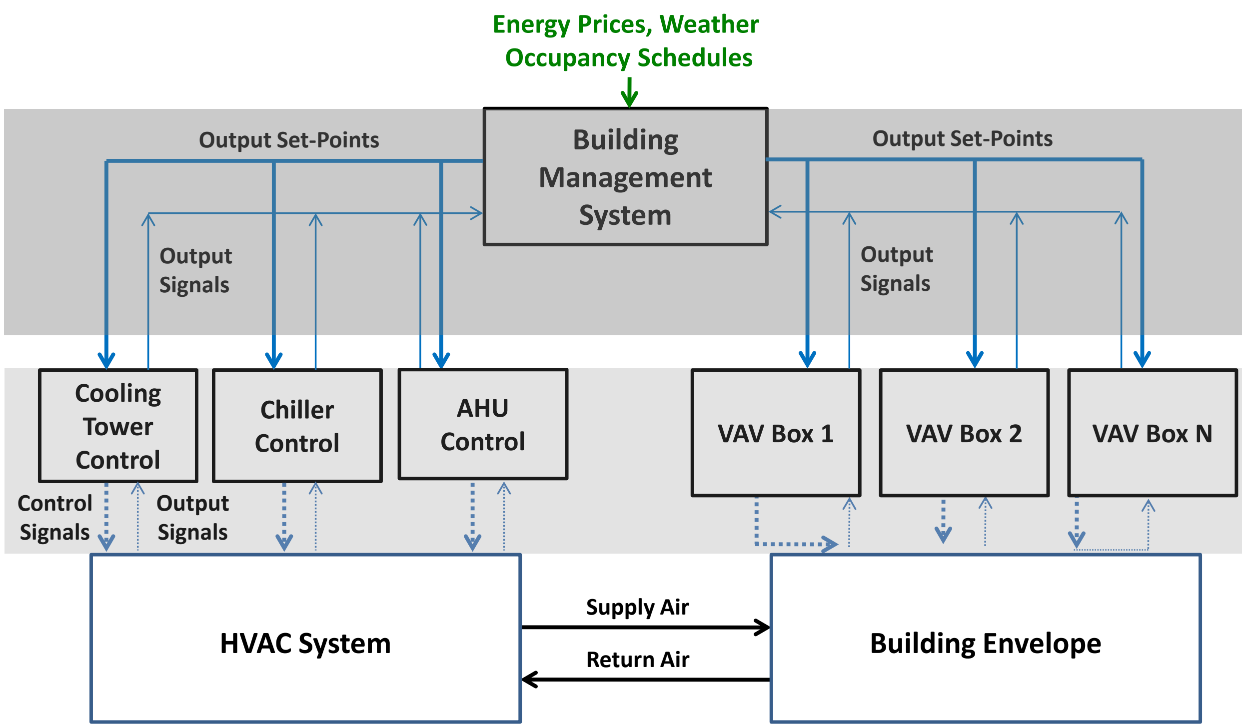

Figure 2: Schematic representation of BMS system. Dark zone is supervisory level and light zone is basic level.

3 Next-Generation Energy Management

One of the key sources of inefficiency in modern buildings is the limited degree automation of the EM system. EM is a complicated task because of the highly dynamic ambient conditions and prices, the many variables and sensors involved, and the close interconnections between the HVAC units and the building envelope. In addition, the complexity of the HVAC layouts is increasing as these systems become more tightly integrated to recover energy and expanded with storage and co-generation units

[1, 9, 17]. Operational complexity is also increasing as buildings are operated under demand-response programs and more volatile energy prices. In dealing with these challenges, this section discusses some of the ideal features of an effective next-generation EM system.

3.1 Models and Integration

Set-points should be computed in a way that captures as many interconnections as possible between the HVAC units and the building envelope and trade offs between energy and comfort. This strategy enables consistency and synchronization of the multiple components. Model-based strategies are thus essential. A model implicitly captures interconnections between components. Most deployed EM systems use logic rules that are not capable of capturing interconnections and complex objectives efficiently. Models can range from first-principles to data-driven black-box models (e.g., regression, neural networks, machine learning) or combinations of both. Usually, the more complex the model, the higher the predictive precision but also the higher the development cost and computational complexity. Higher complexity translates in higher installation and maintenance costs, deployment times, and technology sophistication.

In certain buildings the energy savings can justify the development of high-fidelity energy models such as EnergyPlus, eQuest, TRNSYS [6, 5] or Modelica [21]. A model development task can easily take three months to one year of effort of a highly trained engineer. The costs are related to data collection, model building, and tuning. These tasks currently translate into costs as high as $50,000-$100,000, but these costs might decrease as model development tasks become more automated. If the energy savings can justify this investment, first-principles models are preferred. If the model is already available as part of the design process, as in newly deployed buildings, then the development cost can be saved. The savings in small and medium-size buildings, however, are unlikely to justify the development of sophisticated models.

An advantage of first-principles models is that they are not as sensitive to the availability of sensor information and can provide predictions in non sensed zones. This capability can significantly expand their robustness to sensor faults and their optimization scope, respectively. The computational complexity of first-principles models, however, is still an important barrier to avoid in real-time optimization applications such as EM. One of the reasons, as discussed in [1], is that existing simulation tools do not provide enough capabilities to apply efficient optimization solvers.

Because of the tightness of comfort and operational limits, the building operating conditions remain within a relatively well-defined window. Consequently, building models do not necessarily require wide prediction domains such as those captured in first-principles models. Hence,, black-box data-based models present an interesting alternative for EM. These models learn adaptively the building and HVAC responses using sensor information. Consequently, they are highly adaptive and inexpensive and can be deployed over a wide range of buildings of different sizes. In addition, these models are computationally more tractable. A limitation of black-box models is that they rely entirely on available sensor information. Thus, their performance can be affected more strongly by sensor availability and faults.

A combination of both first-principles and black-box modeling paradigms can provide solutions that trade off costs and precision. Currently, however, building modeling environments that enable such constructs are not commonplace.

3.2 Constraints and Objectives

The EM set-points should satisfy operational constraints such as maximum equipment capacities and dynamic limitations due to equipment thermal capacities, low-level controllers, and the building thermal mass. This capability enables feasibility and reachability of the set-points by the basic control layer.

This feature requires of optimization strategies coupled to *dynamic* models with appropriate precision. Optimization enables the system to enforce constraints and to trade off building-wide objectives systematically. Model precision enhances the chances that the set-points sent to the basic level controllers are feasible. The model should also capture major dynamic limitations of the HVAC units and of the building envelope limiting performance. This requirement is important since the operating conditions of most units and of the building zones cannot be ramped up and down instantaneously.

3.3 Foresight and Set-Point Update Frequency

Set-points should be updated at a frequency that is consistent with major disturbances affecting building-wide objectives. In particular, strong variations of weather and occupancy require higher-frequency updates. Real-time pricing will also require of higher-frequency updates. The computational complexity of the EM system constrains this updating frequency. If the EM system is equipped with appropriate forecasting capabilities and dynamic models, however, the updating frequency becomes less relevant since the variations can be anticipated. On the other hand, if a steady-state model is used in the EM system, the strong variations can lead to unreachable set-points (e.g., due to slow building responses not captured by the model), suboptimal performance, and drifts from comfort conditions. Because of these limitations, recent studies have shown that using proactive EM systems with dynamic models is preferred [24, 23]. In particular, these systems enable a more efficient exploitation of external ambient air, not only to shift demands to save energy costs [3], but also to reduce HVAC energy consumption.

3.4 Sensor Information

In any automated system, it is critical to have informative sensors that can be used by the EM system to map basic variables such as temperature and humidity to building-wide objectives such as power consumption in the AHUs and overall comfort. The information content in sensors is related to number, position, and type. The number of sensors and their required reliability is directly related to the type of model to be used in the EM system. It is also critical to have a reliable information infrastructure to capture external factors affecting performance such as pricing and weather.

4 BIQ's Energy Management System

In this section, we describe the features of BIQ's EM system. This system presents several of the desirable features discussed in the previous section.

4.1 Black-Box Modeling Approach

The physical optimization domain of BIQ's system includes the AHUs and the building zones. The system follows a black-box modeling approach to capture the effect of the AHU energy consumption (power), ambient conditions, and other factors on the average zone temperature (average of all subzone temperatures).

We consider a set of zones $j \in \mathcal{Z}$. We denote the average temperature at zone j at time τ as $T_j(\tau)$. A black-box dynamic model for each average zone temperature can be represented as:

$$\frac{dT_j(\tau)}{dt} = f_j^z(T_j(\tau), P_j(\tau), T^{amb}(\tau), p_j^z) + w_j^z(\tau),\ \tau \in [t, t+T], \tag{4.1}$$

where $P_j(\tau)$ is the consumed AHU power at time τ, $T^{amb}(\tau)$ is the ambient temperature at time τ, p_j^z are the model parameters for the zone, and $w_j^z(\tau)$ are the model errors at time τ. The model errors $w_j^z(\tau)$ account for unmodeled effects such as heat gains from occupants and solar radiation (thermal loads). The model structure $f_j^z(\cdot)$ can be constructed by using machine learning techniques such as support vector machines or by using system identification techniques [19]. Using system identification with second-order polynomial transfer functions is often preferred computationally.

At time t_k the system can collect values for the measured variables $T_j(t_i), T^{amb}(t_i), P_j(t_i)$ over the previous times $t_i,\ i = k, k-1, ..., k-N$ where N is the number of measurements, to *train* the model by computing estimates for the parameters p_j^z, model errors $w_j^z(t_i), i = 0, ..N$ and initial conditions $T_j(t_k)$.

Having a given amount of power $P_j(\tau)$ for the AHU system, the system back calculates set-points for the AHUs. BIQ's system focuses on the supply air temperature $T_j^{in}(\tau)$ and volumetric flow $V_j(\tau)$ of the AHU. This can be done using a black-box AHU model of the form:

$$P_j(\tau) = g_j^{AHU}(T_j^{in}(\tau), V_j(\tau), T^{amb}(\tau), p_j^{AHU}) + w_j^{AHU}(\tau),\ \tau \in [t, t+T]. \tag{4.2}$$

where p_j^{AHU} and $w_j^{AHU}(\tau)$ are the parameters and errors of the AHU model for zone j. The model structure $g_j^{AHU}(\cdot)$ is also constructed by using black-box modeling techniques.

To predict the effect of the zone temperature $T_j(\tau)$ on human comfort, BIQ's system uses ASHRAE's percentage of people dissatisfied (PPD) metric [14]. To obtain real-time PPD measurements, BIQ's system contains a feature called ComfortIQ, which reports the comfort status of the occupants at current conditions. This makes the PPD metric adaptable to the current building configuration and nature of the occupants. This is important since comfort is a subjective metric. The resulting PPD model can be posed as:

$$PPD_j(\tau) = g_j^{PPD}(T_j(\tau), p_j^{PPD}),\ \tau \in [t, t+T], \tag{4.3}$$

where p_j^{PPD} are the model parameters.

The entire dynamic model for zone j can be posed as a system of differential and algebraic equations (DAEs) of the form:

$$\frac{dx_j(\tau)}{d\tau} = f_j(x_j(\tau), y_j(\tau), z_j(\tau), u_j(\tau), p_j) + w_j^x(\tau) \quad (4.4a)$$

$$0 = g_j(x_j(\tau), y_j(\tau), z_j(\tau), u_j(\tau), p_j) + w_j^y(\tau), \quad (4.4b)$$

where $x_j(\cdot)$ is the dynamic state given by the average zone temperature, $y_j(\cdot)$ are the algebraic states given by the AHU power and PPD variables, $u_j(\cdot)$ is the control given by the supply air temperature, p_j are the parameters, and $w_j^x(\cdot), w_j^y(\cdot)$ are the model errors. Variables $z_j(\cdot)$ are exogenous disturbances given by the ambient temperature and the electricity prices. BIQ's system models *separately* each of the building zones j. The interactions between zones are handled implicitly through the model errors $w_j^x(\cdot), w_j^y(\cdot)$.

4.2 Optimization Structure

Once a dynamic model is learned, BIQ's system computes optimal set-point policies for the AHU over a future horizon $\tau \in [t_k, t_{k+T}]$ to minimize a future performance objective. Here, t_k is the current time and T is the length of the prediction horizon. This problem can be cast as an optimal control problem of the following form:

$$\min_{u_j(\tau)} \int_{t_k}^{t_{k+T}} \varphi_j(x_j(\tau), y_j(\tau), z_j(\tau), u_j(\tau)) d\tau \quad (4.5a)$$

$$\text{s.t.} \quad (4.5b)$$

$$\frac{dx_j(\tau)}{d\tau} = f_j(x_j(\tau), y_j(\tau), z_j(\tau), u_j(\tau), p_j) + w_j^x(\tau) \quad (4.5c)$$

$$0 = g_j(x_j(\tau), y_j(\tau), z_j(\tau), u_j(\tau), p_j) + w_j^y(\tau) \quad (4.5d)$$

$$0 \leq h_j(x_j(\tau), y_j(\tau), z_j(\tau), u_j(\tau)) \quad (4.5e)$$

$$x_j(t_k) = \text{given}, \ \tau \in [t_k, t_{k+T}]. \quad (4.5f)$$

Here, the initial conditions $x_j(t_k)$ are given by the current state of the average zone temperature of zone $j \in \mathcal{Z}$. This optimization formulation is a special case of the general conceptual EM framework presented in [1].

4.2.1 Performance Objective

The objective of BIQ's optimization problem (4.5a) is a weighted sum of three competing metrics. The cost has the form:

$$\varphi_j(x_j(\tau), y_j(\tau), z_j(\tau), u_j(\tau)) := \alpha \cdot PPD_j(\tau) + \beta \cdot P_j(\tau) \cdot \lambda(\tau) + \gamma \cdot CO_2 f \cdot P_j(\tau). \quad (4.6)$$

The terms correspond to comfort, energy cost, and CO_2 emissions (indirect measure of energy consumption). The time-varying energy prices are given by $\lambda(\cdot)$. The symbol $CO_2 f$ is a conversion factor, and α, β, and γ are weighting factors. We note that minimizing energy cost and minimizing consumption are not equivalent goals because it is possible to minimize energy cost by shifting the demand according to the price structure and still use the same amount of energy (or greater). The weighting factors can be tuned based on the operation priorities of the site.

As more priority is assigned to PPD, the system will try to keep the average zone temperatures close to the specified set-points. When PPD is relaxed and more priority is assigned to energy consumption (e.g., CO_2 emissions), the system will allow for drifts from the zone temperature set-points. As more preference is given to energy costs, the system will allow for drifts in the zone temperatures that shift the demand profile according to the price structure along the day.

4.2.2 Degrees of Freedom

The explicit degrees of freedom of the EM system are the supply air temperature and air volume of the AHUs. The system also allows one to implicitly manipulate the drifting of the average zone temperatures. In other words, the system does not explicitly manipulate the zone temperature set-points.

4.2.3 Operational Constraints

The system can explicitly handle bounds on the degrees of freedom, PPD, and on the zone temperatures. These operational limits are contained in the constraints (4.5e) or the optimization problem (4.5). In addition, the system implicitly handles dynamic limitations due to the thermal mass of the building. These are part of the dynamic model (4.5c).

4.2.4 Sensors Needs

The system requires basic temperature sensors in the zones, supply air temperatures in the AHUs, air demands of VAV boxes, and power meters for the AHUs.

4.3 Real-Time Implementation

The system follows a receding-horizon MPC implementation. At time t_k it solves the optimization problem (4.5) over the horizon $[t_k, t_{k+T}]$ using the current state of the building as initial conditions and the forecasts for prices and ambient conditions. The computed set-points for the AHU are implemented over a period of five minutes. At the next time step t_{k+1}, the state of the building is measured, and the optimization problem is solved over the horizon $[t_{k+1}, t_{k+1+T}]$ by using the updated forecast information. The implementation is illustrated in Figure 3. The model is currently retrained daily, but the training frequency can be adapted as necessary.

The receding horizon implementation enables the introduction of feedback in order to account for modeling and forecast errors. The prediction horizon length used by BIQ's system ranges from 3 to 24 hours. This horizon length is recommended to be as long as possible to capture the periodicity of the ambient temperature and occupancy. This recommendation is supported by research experience reported in [24].

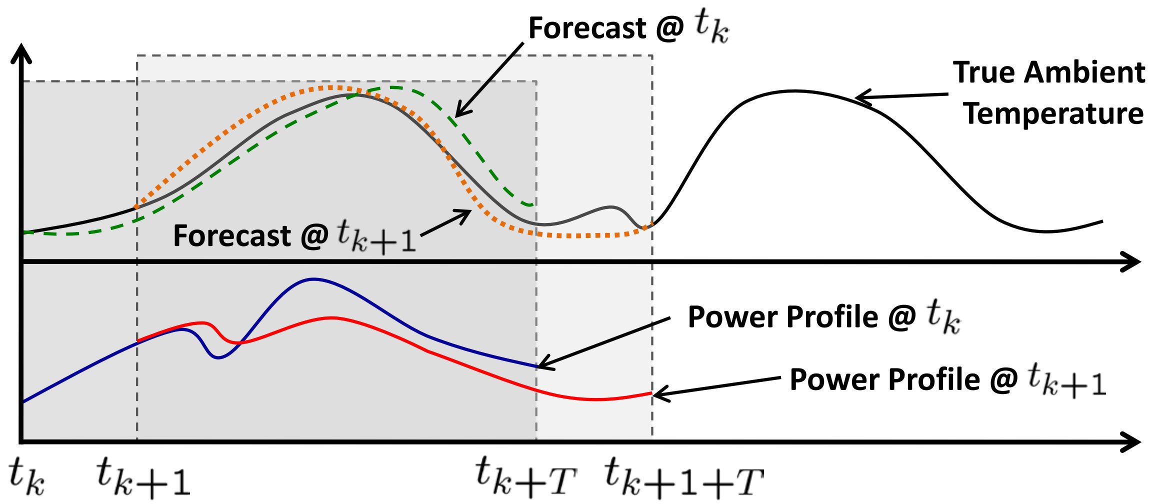

Figure 3: Schematic representation of receding-horizon MPC implementation.

5 Implementation at Argonne

BIQ's system has been implemented at a commercial-sized building at Argonne National Laboratory. The building occupies a total floor space 200,000 sq.ft. with 7 AHUs, and holds 700 occupants. The building also hosts significant amounts of computing equipment ranging from laptops and personal computers to high-performance computers.

The building under study was opened in 2009, and it is equipped with a standard configuration of state-of-the-art sensors and uses Johnson's Controls Metasys system as the BMS. The implementation of BIQ's system required the following tasks.

- *Data gathering and storage*. Before this study, the Metasys system was not configured to store operational data for more than one day. This configuration severely limited data analysis and collection required to perform preliminary learning tasks of BIQ. The BIQ team installed their system to perform data collection from Metasys and to perform storage.

- *Power meter installation*. The building was equipped only with a central power meter. It was necessary to install power meters at the AHUs in order to monitor and learn the AHU power consumption under different operating conditions.

- *Model learning and validation*. Collected data over a coupled of weeks was used to perform preliminary model training of the BIQ's system. The model was validated against real measurements.

- *Deployment in static mode*. Prior to full deployment, the system was deployed in static mode at a single AHU in order for the operator to gain experience. This deployment has also enabled us to perform a preliminary energy savings analysis, which we report here.

- *Deployment in live mode*. Full deployment is pending on approval by the building owners.

The validation and energy savings analysis was performed by using data collected during the first two weeks of October 2010. In Figure 4, we present the model fit of an average zone temperature over 14 days of operation. Note that the range of operation is 65-80^{o}F. These variations are typical during seasonal transitions. The model is able to track the signal closely. The errors are in the 1-2^{o}F range. The sudden spikes in the measured temperatures reflect a sensor malfunction or bad reading (it is not physical possible for the zone temperatures to move that quickly), the model fitting exercise filters out these errors.

Figure 4: Validation of an average zone temperature for the Argonne building. Green line is real measurements and red line is the model fit.

As a first case study (Scenario 1), we consider a day in which there is a wide ambient temperature difference during the day (cool day), between 45^{o}F and 85^{o}F. The EM system was run in simulation mode using the learned model of the building and considering a case for 3% PPD. In all cases we assume that the electricity price remains fixed, which is the case at Argonne. Hence, we can focus on energy savings and comfort. The results are presented in Figure 5. The light blue line is the ambient temperature, the pink line is the current power consumption of the AHU (in kW), and the black line is the optimized power from BIQ's system. The dark blue line is the current average zone temperature, and the red line is BIQ's signal. We summarize our findings as follows:

- The BIQ's system is able to drop the power consumption during the entire day by 31% compared with current operation.
- We have discovered that the AHU is currently run at night at nearly 30% capacity. This represents a significant amount of energy lost. This discovery has been made possible by our ability to track the power consumption of the AHU with the newly installed power meters.

- In current operation, cooling is started at around 5 a.m. BIQ's system delays this start to nearly 9:45 a.m., since occupancy typically starts at 8:30 a.m. The system is able to predict the ambient temperature trend, which remains below 60^{o}F until 9:45 a.m. The system exploits this cold ambient air to cool down the building instead of recurring to electrical cooling in the AHU.

- The BIQ's system decreases power consumption by exploiting the thermal mass of the building as it foresees closing hours starting at 6 p.m. This can be seen from the shape of the peaks. The current peak is high and flat with a peak of 320 kW whereas that of the BIQ system is much lower and exhibits a consumption of 250 kW two hours before the closing time.

- BIQ's system keeps the average zone temperature always less than 1^{o}F away from current operation during the *unoccupied* times. This enables the system to further save energy. During occupied times, the temperature is kept equal to the baseline, enabling large savings without sacrificing comfort. The tight drift of 1^{o}F results from the high priority on comfort (3% PPD). The savings can further be increased by allowing larger drifts during the unoccupied periods.

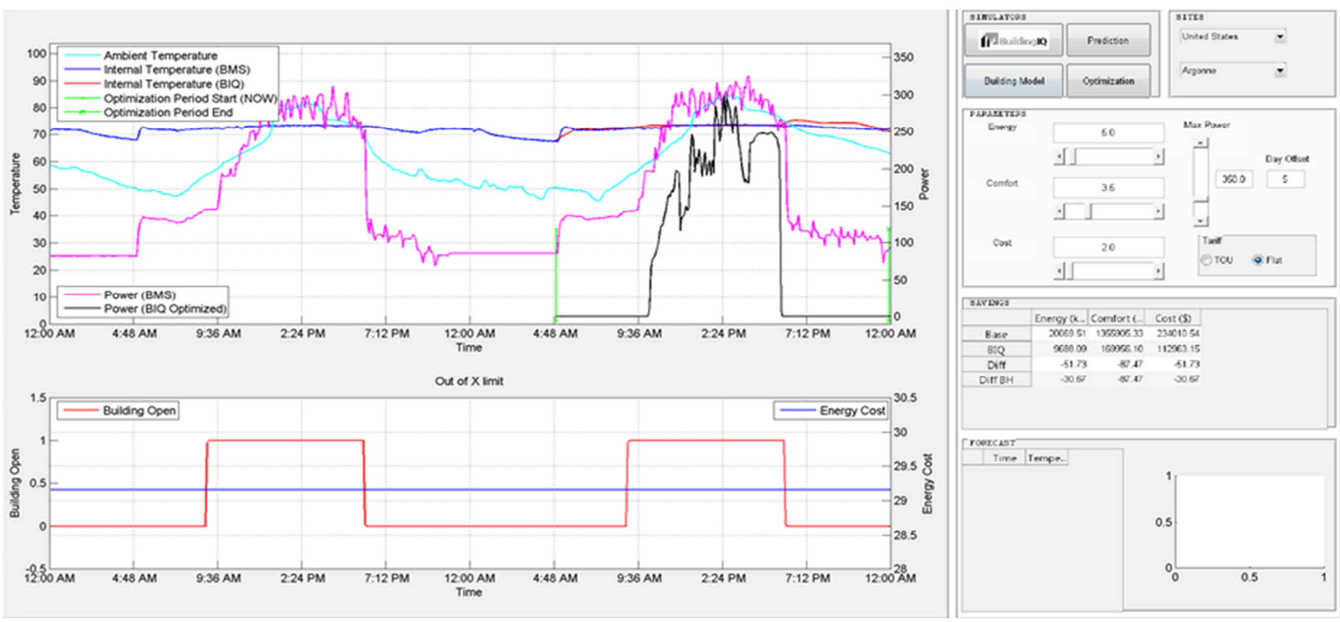

Figure 5: Scenario 1: Energy savings analysis for cool day, 3% PPD.

In the second case study (Scenario 2), we consider the same day but with very high priority on PPD (0.1%). Comparing the profiles to those of Scenario 2 we can draw the following conclusions:

- The energy savings in this case drop to 29%. This can be explained from the high PPD priority, which does not permit temperature drifts during the unoccupied times.

- Despite the high PPD priority, the power profiles for the AHU remain nearly identical to those of Scenario 1. This result reflects the large effect that the ambient temperature profile has on the energy savings.

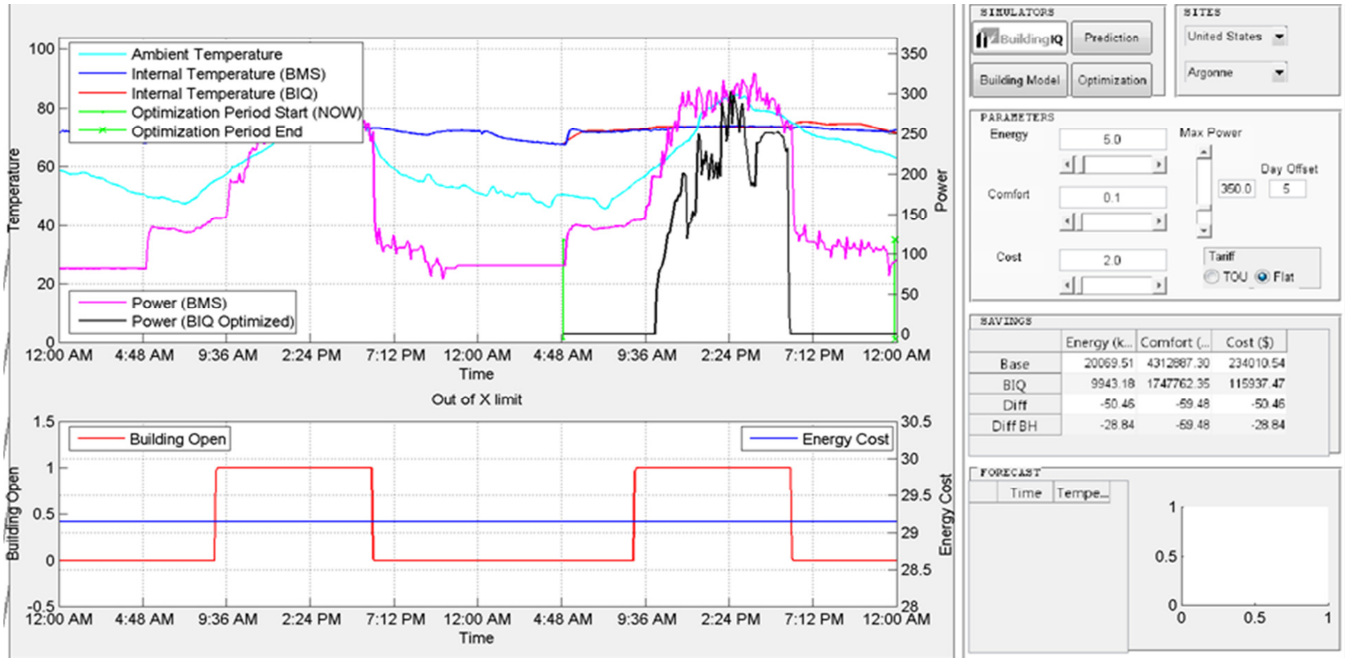

Figure 6: Scenario 2: Energy savings analysis for cool day, 0.1% PPD.

To further demonstrate the impact of the ambient temperature on energy savings, we consider a third case study (Scenario 3). This is a day with a narrower ambient temperature difference (warmer day between 55^{o}F and 85^{o}F). We have the following findings:

- The energy savings in this case drop significantly, to 6%. This strong effect is due to the higher night temperatures in this scenario compared with those of Scenario 1 (45^{o}F against 55^{o}F). Note, however, that the peak ambient temperatures of Scenarios 1 and 3 remain at 85^{o}F. This indicates that the magnitude of the savings is directly related to the ability to precool the building by using night ambient air and delay the start of the AHU.

- Comparing the shape of the power profiles of Scenarios 1 and 3, we observe that the power profiles remain nearly the same during the unoccupied periods. We can conclude that the savings during this time are around 5-6%. Consequently, for Scenario 1, the savings achieved during occupied times are nearly 25%, which represents almost 80% of the total savings.

- From the power profile of Scenario 3, we can see that BIQ's system peak consumption is above current operation levels. This indicates that the system is foreseeing the decay in the ambient temperature that can be exploited to shut down the AHU earlier than current operation.

In our last scenario, we consider the same ambient conditions of Scenario 3 (warmer day), but we relax the PPD to 5%.

- The energy savings in this case increase to 14% compared with savings of 6% enforcing a 3% PPD.

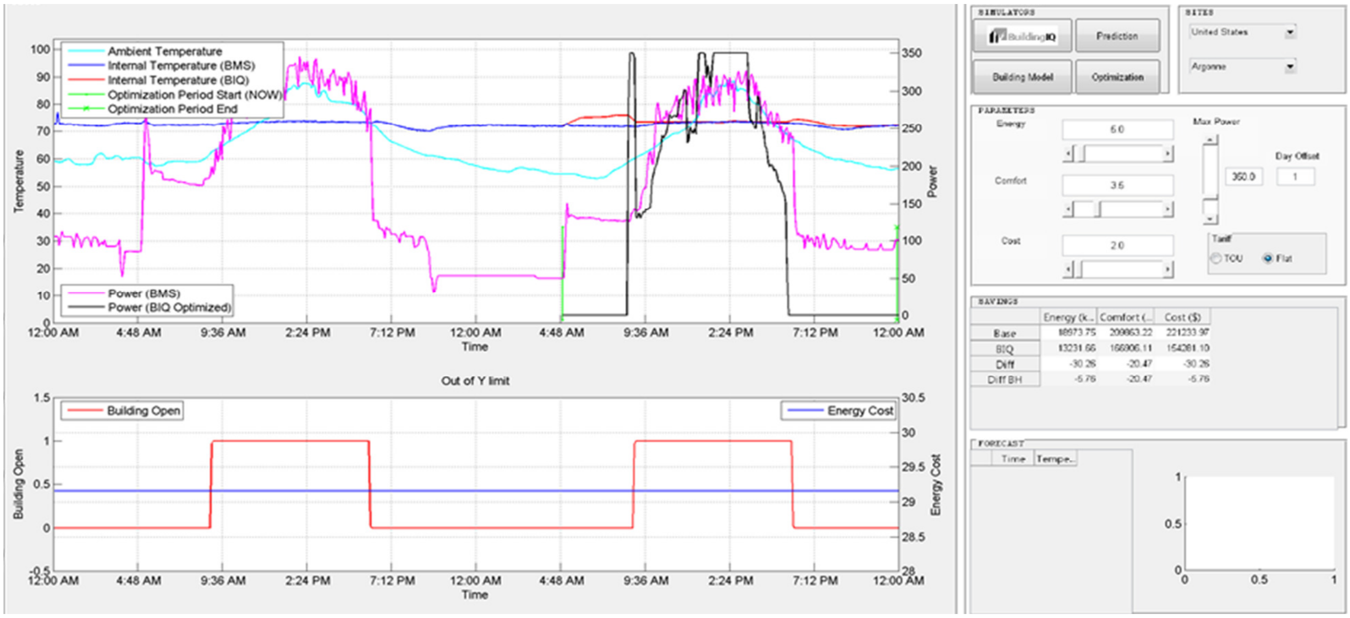

Figure 7: Scenario 3: Energy savings analysis for warmer day, 3% PPD.

- Note the larger drifts in the average zone temperature. Remarkably, however, the drifts occur only during the unoccupied times. This indicates that relaxing the average zone temperatures has dramatic effects on savings. Since the drifts occur during unocuppied periods, these savings *do not affect comfort conditions*.

We have observed that, by relaxing the PPD to levels of 10% during unoccupied times, it is possible to achieve savings of up to 45% for days with wide ambient temperature differences and to 25% in days with narrow differences. Again, since the zone temperature drifts do not occur during occupied times, these savings come at no expense in comfort. The savings, however, are limited by the thermal mass of the building. In other words, by allowing very large drifts at night, the system would need to ramp up the entire building temperature before occupancy starts, a situation that might not be feasible.

In terms of energy costs, we consider Scenario 1 with savings of 30% of cooling power. The total power consumed by a single AHU in current operation is 3.94 MWh a day; BIQ's system brings it down to 2.90 MWh. If we consider 7 AHUs and an electricity rate of 80$/MWh, we have daily cost savings of around $550. If we assume that these savings can be achieved 200 days of the year, we have annual savings of $110,000. A service time of 200 days with similar savings is reasonable because the building under study runs on electric heating. The cost of the license of BIQ's system is around $36,000 a year. The license of the system is charged at a rate of 0.18$/sq.ft. per year for buildings around 200,000 sq.ft. The installation costs of the power meters was around $3,000. Thus, installating BIQ's system is profitable with savings above 10% with respect to current operation.

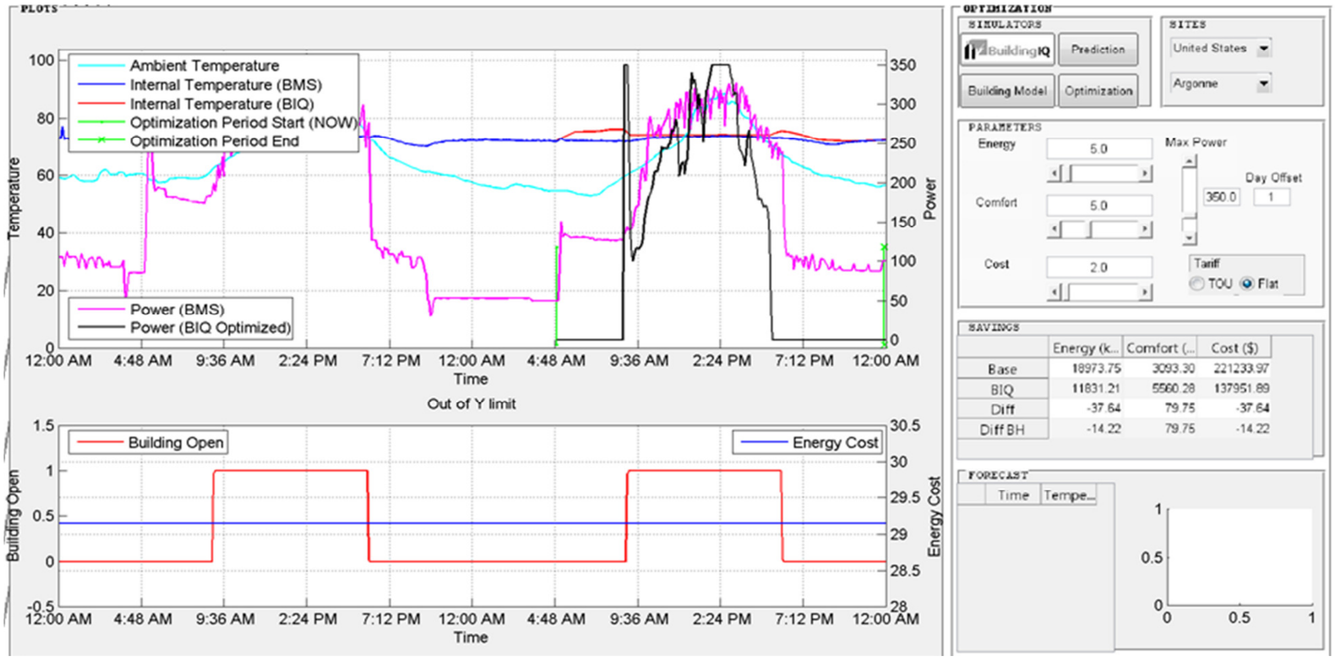

Figure 8: Scenario 4: Energy savings analysis for warmer day, 5% PPD.

6 Technological Evaluation

In this section we discuss the innovations of BIQ's system that we believe are critical in achieving important energy savings. In addition, we discuss different limitations of the technology.

6.1 Innovations

One of the key innovations of BIQ's system is the ability to learn the building dynamics adaptively. This enables the system to quantify and mitigate effects that are difficult to model from mechanistic principles, for instance, heat gains due to radiation and occupants. The use of black-box models enhances the adaptivity of the system to different HVAC configurations. This also enables fast deployment which can be as short as a month, and low cost for the technology. Consequently, this system can potentially be deployed over a wide range of commercial facilities and have significant impact on overall utility demands. For instance, if proactive EM systems are deployed over an entire city, the utility companies will be able to predict the overall city demand for the next day more accurately. This capability would result in much more efficient management of distribution resources and look-ahead planning of demand-response events.

Another key advantage of the system is the use of model predictive control concepts. This enables the system to account for dynamic limitations of the building; to take proactive actions in the presence of variations of weather, prices, and occupancy, and enables the system to handle constraints and complex performance objectives. The proactive feature makes the system particularly suitable for demand response tasks and also enables smoother (e.g., less agressive) control actions [1, 23].

The use of the PPD metric and measurements based on real-time polls from the occupants enables adaptation to changing comfort needs of the occupants, which might change in time due to clothing (difficult to model mechanistically) and from building to building.

6.2 Limitations

One of the key limitations of BIQ's system is that it captures the performance of the zones in an aggregated manner. In other words, it does not capture the granularity at the subzone configuration level. This strategy results in limited control capabilities that can ultimately lead to a wide spread of subzone temperatures and limited energy savings.

Another limitation is the fact that interactions between zones are not taken into account systematically. The system implicitly learns the effect of one zone on another through the model errors but it cannot predict them. If the interactions between zones are strong, this can lead to decreased overall performance as zones compete against each other.

The system currently does not account for occupancy at a zone by zone level. Occupancy, and more generally the building's variable thermal load associated with occupancy and related activities, are accounted for in an aggregated manner. This limitation is reflective of the fact that occupancy information is rarely available on a fine granular level (if at all). However, with the advent of advanced and cost-effective occupancy sensors, there is potential for significant improvements in energy savings to be derived from occupancy-driven EM strategies.

The system handles the multiple objectives in the optimization problem using weighting functions. This setting requires a different control mindset to traditional BMS control tuning. A poorly chosen set of weights will lead to sub-optimal performance and it is therefore required that operators of the system receive appropriate training. Opportunities exist to explore adaptive self-tuning approaches, which would lessen then operator skill level requirements.

The system does not account for the dynamic evolution of CO_2 and humidity but it is possible to extend the black-box modeling framework to account for these variables. It is expected that these extensions will lead to further savings and increased comfort.

7 Research Needs and Next Steps

The operational experience gained through the implementation of the BIQ's system at Argonne and the subsequent the energy savings analysis had pointed to several research directions.

As we have seen, allowing for temperature drifts during unoccupied periods leads to significant savings. We have also observed that 80% of the savings are obtained during peak times in regular office hours. Supported by these observations, our hypothesis is that significant savings can be obtained by monitoring occupancy during regular office hours (empty conference rooms and office spaces). In Figure 9 we present a snapshot of the location of the occupants at the Argonne building (using cell phone signals for tracking). As can be seen, there exist several zones in the building with nearly no occupants.

Real-time occupant locations can be tracked by using radio signals. We are currently installing this type of occupancy sensors. These will be acquired from Johnson Controls. With this information,

Figure 9: Snapshot of occupants locations at the Argonne building.

it will be possible to send signals to the BMS system to indicate which subzones are occupied. The BMS system will automatically position the set-points in set-back mode. This will generate lower demands in the VAV boxes that will be detected by the BIQ's system.

The deployment of the system at Argonne will be valuable in this study since it will allow us to have a baseline to quantify additional energy savings obtained from additional technologies such as occupancy sensors. In addition, since the EM system can be run in simulation mode, these savings can be quantified prior to purchase and deployment.

Questions that we plan to address in this experimental study are:

- How much does occupancy information increases energy savings?
- What is the effect of occupancy sensor on the optimal operating policies of BIQ's system?
- What are the optimal zone set-back temperatures that maximize savings without affecting comfort responsiveness?
- What is the effect of the number and position of occupancy sensors on energy savings? Is the installation justifiable from an energy savings point of view?

From a more technical stand point, several improvements can make EM technology significantly more efficient. First, it is necessary to capture zone and subzone interactions in the EM system model.

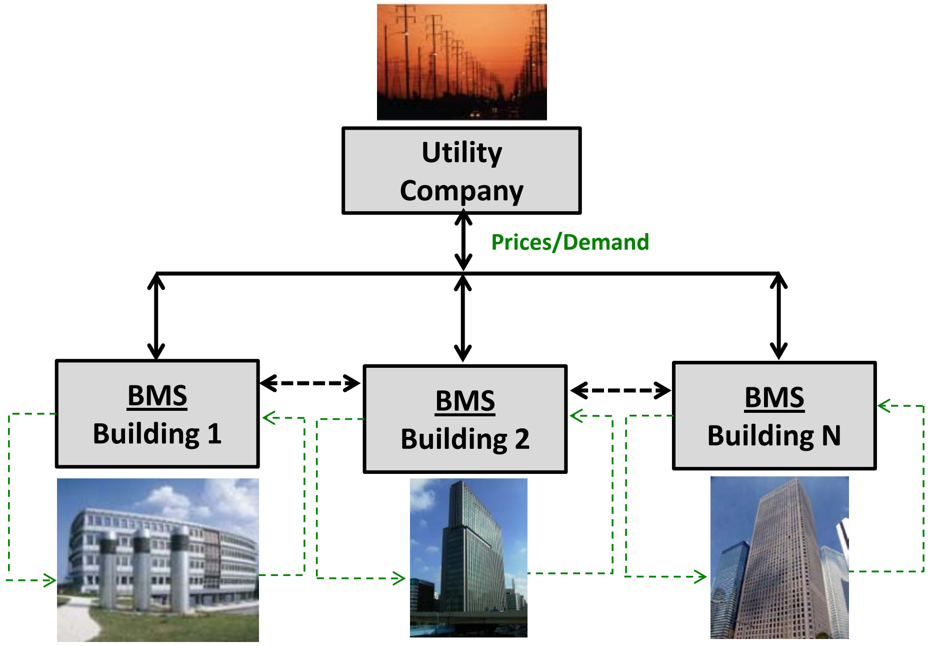

Figure 10: Schematic representation of multi-building energy management.

In addition, the prediction horizon should be extended as much as possible in order to anticipate strong temperature variations. These two enhancements represent a challenge from the computational point of view since the optimization problems solved by the EM system become much larger and more difficult to solve. This is particularly critical in real-time environments since the system needs to be run every few minutes ro reject strong dynamic variations. We will explore the use of more advanced optimization solvers to avoid these bottlenecks.

The use of more advanced optimization solvers will also enable us to expand the domain of the EM system to manage multiple buildings simultaneously. Since buildings have different demand patterns and physical limitations, this capability will enable the EM system to coordinate demand response events in order to maximize the total aggregated savings. This is sketched in Figure 10.

We will also explore the use of advanced multiobjective optimization algorithms in order to avoid the use of weighting factors. This will enable a more adaptive EM system with minimal intervention from the operator.

Acknowledgments

This work was supported by the U.S. Department of Energy, under Contract No. DE-AC02-06CH11357. We thank Dave Bruce from Johnson Controls for his assistance in deploying BIQ's system at Argonne.

References

[1] V. M. Zavala amd J. Wang, S. Leyffer, E. M. Constantinescu, M. Anitescu, and G. Conzelmann. Proactive energy management for next-generation building systems. *In Proceedings of 4th International Conference of IBPSA-USA*, 2010.

[2] J. T. Betts and W. P. Huffman. Application of sparse nonlinear programming to trajectory optimization. *J. Guidance Dyn. Cont.*, 1992.

[3] J. E. Braun. Reducing energy costs and peak electricty demand through optimal control of building thermal storage. *ASHRAE Transactions*, 96(2):876–888, 1990.

[4] J. E. Braun, K. W. Montgomery, and N. Chaturvedi. Evaluating the performance of building thermal mass control strategies. *HVAC&Research*, 7(4):403–428, 2001.

[5] D. B. Crawley, J. W. Hand, M. Kummert, and B. T. Griffith. Contrasting the capabilities of building energy performance simulation programs. *Building and Environment*, 43(4):661 – 673, 2008. Part Special: Building Performance Simulation.

[6] D. B. Crawley, L. K. Lawrie, F. C. Winkelmann, W. F. Buhl, Y. J. Huang, C. O. Pedersen, R. K. Strand, R. J. Liesen, D. E. Fisher, M. J. Witte, and J. Glazer. EnergyPlus: Creating a new-generation building energy simulation program. *Energy and Buildings*, 33(4):319 – 331, 2001.

[7] P. S. Curtiss, M. J. Brandemuehl, and J. F. Kreider. Energy management in central HVAC plants using neural networks. *ASHRAE Transactions*, 100:476–493, 1994.

[8] H. Doukas, K. D. Patlitzianasa, K. Iatropoulosa, and J. Psarrasa. Intelligent building energy management system using rule sets. *Building and Environment*, 42:3562–3569, 2006.

[9] M. W. Ellis and E. H. Mathews. Needs and trends in building and hvac system design tools. *Building and Environment*, 37(5):461 – 470, 2002.

[10] K.F. Fong, V.I. Hanby, and T.T. Chow. HVAC system optimization for energy management by evolutionary programming. *Energy and Buildings*, 38(3):220 – 231, 2006.

[11] D. Kolokotsa, A. Pouliezos, G. Stavrakakis, and C. Lazos. Predictive control techniques for energy and indoor environmental quality management in buildings. *Building and Environment*, 44(9):1850 – 1863, 2009.

[12] D. L. Loveday and G. S. Virk. Artificial intelligence for buildings. *Applied Energy*, 41(3):201 – 221, 1992.

[13] L. B. McElroy, J. W. Hand, and P. A. Strachan. Experience from a design advice service using simulation. *In Proceedings of the Fifth International IBPSA Conference*, 1:19–26, 1997.

[14] B. W. Olesen and G. S. Brager. A better way to predict comfort: The new ASHRAE standard 55-2004. *ASHRAE Journal*, pages 20–26, 2004.

[15] B. Paris, J. Eynard, S. Grieu, T. Talbert, and M. Polit. Heating control schemes for energy management in buildings. *Energy and Buildings*, 42(10):1908 – 1917, 2010.

[16] S. J. Qin and T. A. Badgwell. A survey of industrial model predictive control technology. *Control Engineering Practice*, 11:733–764, 2003.

[17] M. Trcka and J. L. M. Hensen. Overview of HVAC system simulation. *Automation in Construction*, 19(2):93 – 99, 2010.

[18] V. Vapnik. *The Nature of Statistical Learning Theory*. Springer-Verlag, 1995.

[19] J. K. Ward, J. Wall, S. West, and R. de Dear. Beyond comfort managing the impact of HVAC control on the outside world. *In Proceedings of Conference: Air Conditioning and the Low Carbon Cooling Challenge*, 2008.

[20] M. Wetter. *Simulation-Based Building Energy Optimization, Ph. D. thesis*. Department of Mechanical Engineering, University of California, Berkeley, Berkeley, CA, 2004.

[21] M. Wetter and P. Haves. A modular building controls virtual test bed for the integration of heterogeneous systems. *Technical Report, Lawrence Berkeley National Laboratory*, 2008.

[22] R. E. Young, R. D. Bartusiak, and R. W. Fontaine. Evolution of an industrial nonlinear model predictive controller. *In Proceedings of CPC-6, Tucson*, 2001.

[23] V. M. Zavala, M. Anitescu, and T. Krause. On the Optimal On-Line Management of Photovoltaic-Hydrogen Hybrid Energy Systems. In R. M. de Brito Alves, C. A. Oller do Nascimento, and E. C. Biscaia Jr., editors, *Proceedings of 10th International Symposium on Process Systems Engineering*, 2009.

[24] V. M. Zavala, E. M. Constantinescu, M. Anitescu, and T. Krause. On-line economic optimization of energy systems using weather forecast information. *Journal of Process Control*, 19:1725–1736, 2009.

Mathematics and Computer Science Division
Argonne National Laboratory
9700 South Cass Avenue, Bldg. 240
Argonne, IL 60439-4847

www.anl.gov