Preprint ANL/MCS-P1728-0210

PROACTIVE ENERGY MANAGEMENT FOR NEXT-GENERATION BUILDING SYSTEMS

Victor M. Zavala[1], Jianhui Wang[2], Sven Leyffer[1]
Emil M. Constantinescu[1], Mihai Anitescu[1], and Guenter Conzelmann[2]
[1]Mathematics and Computer Science Division
[2]Decision and Information Sciences Division
Argonne National Laboratory, 9700 S Cass Ave, Argonne, IL 60439

ABSTRACT

We present a proactive energy management framework that integrates predictive dynamic building models and day-ahead forecasts of disturbances affecting efficiency and costs. This enables an efficient management of resources and an accurate prediction of the daily electricity demand profile. The strategy is based on the on-line solution of mixed-integer nonlinear programming problems. The framework is able to integrate forecasts of weather conditions, fuel prices, heat gains, and utility demands. In addition, it can capture net-metering interactions using agent-based market models. We claim that a large adoption level of this proactive technology can improve the predictability of the overall electricity demand at high-level power grid operations such as unit commitment and economic dispatch which can be used to minimize the overall reserves.

INTRODUCTION

In an attempt to minimize energy consumption and costs, next-generation building systems will need to become more active participants of the electricity market by reducing and shaping their electricity demand profile. In addition, they will need to manage resources in an efficient manner under highly volatile conditions. These requirements can be met through the adoption of net-metering schemes and the installation of cogeneration and storage units. While these actions can lead to significant reductions in the overall costs of the power grid, building operations will also become significantly more complex.

A key conceptual problem of current energy management (EM) technology is that it is reactive, in the sense that the operational decisions (controls) are updated based only on current information of disturbances such as the weather conditions and electricity prices. This type of strategy is followed, for instance, in widely used simulation packages such as EnergyPlus or TRNSYS (Crawley et al. 1999; Klein, Duffie, and Beckman 1976). This lack of proactiveness has important effects on the costs and responsiveness of the automation system, mainly because the dynamic responses of the building zones are slow (on the order of hours). This also limits the possibility of exploiting daily disturbance trends and of using storage components (batteries, thermal, and ice storage) to manipulate the electricity demand profile. In addition, the lack of proactiveness limits the robustness of feedback control loops and increases the frequency of startups and shutdowns of pumps, chillers, and fans, leading to decreased equipment lifetime.

In this work, we propose a proactive, optimization-based EM framework able to incorporate disturbance forecasts and predictive building models. This will enable the building automation system to (1) manage multiple generation and storage components systematically, (2) exploit future disturbance trends affecting costs, (3) predict the daily electricity demand profile, and (4) increase the robustness of the feedback control loops. The core of the proposed EM framework is a mixed-integer optimal control problem that can be cast as a mixed-integer nonlinear programming (MINLP) problem. The proactive EM approach is motivated from the fact that state-of-the-art optimization tools are capable of solving large-scale problems with hundreds of thousands of variables on standard personal computers. This makes the proposed technology suitable for massive deployment. The paper is structured as follows: We first motivate the development of the proposed EM framework through a brief analysis of the future operational environments of building systems. We then describe the components of the EM framework including a general mixed-integer optimal control formulation, forecasting capabilities, and agent-based market models.

NEXT-GENERATION BUILDING SYSTEMS

In this section, we describe the operational environments of building systems in order to motivate the components and characteristics of the proposed EM framework. We discuss issues related to the effect of cogeneration and storage technologies on costs, hierarchical operations, disturbances arising in buildings, the formulation of the EM problem, the effect of using physics and empirical models, strategies to forecast disturbances, and solution strategies for the EM problem.

Cogeneration and Storage Technologies

Building automation systems can reduce electricity costs by *shaping* the demand profile. For instance, since in hot regions the electricity demand is dominated mostly by air conditioning, it is possible to shift the peak demand by making use of the thermal mass of the building (Braun 1990; Braun, Montgomery, and Chaturvedi 2001). In this case, the use of electricity is maximized when it is cheaper, which is usually at night. With the introduction of the smart grid in unregulated markets and with a major deployment of solar and wind resources, large-scale storage, and plug-in hybrid electric vehicles, electricity pricing structures will become significantly more complex and volatile. Consequently, peak shifting strategies will not be sufficient to minimize costs. However, the conceptual idea of shaping the demand profile can still be exploited by installing multiple cogeneration and storage facilities in the building. Co-generation technologies include integrated heat and power systems, including gas and diesel turbines, solid-oxide and molten carbonate fuel-cells and microturbines; solar and wind power; solar heating and cooling; ice and thermal storage; and batteries. With this added flexibility, buildings will be able to sell and buy electricity from the grid, becoming much more active participants of the grid and forming local markets with other buildings.

Dynamic Disturbances Affecting Costs and Efficiency

The basic operational objective in buildings is to provide comfort to occupants at minimum energy costs. This implies balancing heat in the building zones to control temperatures and to balance air to control pressure and pollutant (e.g., carbon dioxide) concentrations at appropriate levels. This task is complicated because buildings are operated in the presence of persistent and volatile *disturbances* such as external temperature, radiation, wind and humidity conditions; heat gains due to equipment, lighting, and heating; and heat and air losses. In addition, complex occupant *behaviors* have a strong influence on this. Finally, changing prices of natural gas, diesel, and electricity can also be thought as disturbances affecting the economic performance of the building. Most of these disturbances exhibit periodic trends on time-scales from days to months. For instance, occupant and lighting heat gains exhibit daily and weekly periods, whereas solar radiation, wind speed, and temperature exhibit both daily and monthly periods. Weather conditions are particularly critical because, with the deployment of cogeneration facilities, they will affect not only the demand but also the supply of electricity to the building. In addition, electricity prices traditionally have been handled as an *exogenous* disturbance, in the sense that the EM system has no influence over it. However, in next-generation buildings, this disturbance will become *endogenous* because the EM system demand will have an effect on the electricity market, which in turn dictates the prices.

Hierarchical Decision-Making

The operational decisions in building automation systems can be *conceptually* decomposed, in analogy to power grid operations, into hierarchical levels such as unit commitment, economic dispatch, and feedback control. This decision hierarchy is illustrated in Figure 1. The *unit commitment* decisions consist of setting up the schedule of startup and shutdown tasks for the equipment units. Once this on/off schedule is set, an economic dispatch layer determines the operating set-points based on an economic metric. The set-points are sent to a feedback control layer. Commitment and dispatch tasks are

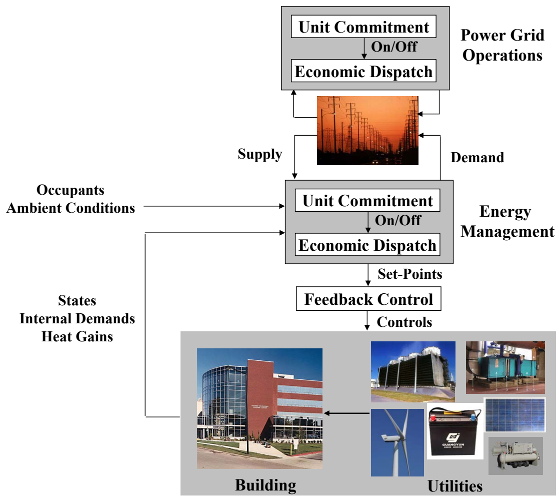

Figure 1: Hierarchical representation of building operations.

the most critical steps because they implicitly *manage* all of the building energy resources (such as passive and active storage devices; photovoltaic, wind, and grid power; and fuel) to satisfy occupancy HVAC and lighting demands while trying to minimize the overall energy costs. Consequently, the commitment/dispatch layer can also be seen as an *energy manager*. This manager can consist of a logic-based system or human operator that determines the operational decisions based on experience in response to hourly variations of disturbances such as external ambient conditions, electricity and fuel costs, and heat gains. Logic-based management systems are based on decision tree rules (e.g., if-then-else statements) and are widely used for EM in literature reports and in EnergyPlus (Ellis, Torcellini, and Crawley 2007; Vosen and Keller 1999; Ulleberg 2004). A problem with these strategies is that design and training (tuning) of the decision-making tree structure and thresholds can become intractable in tightly coupled systems. In addition, these strategies are not robust in the presence of situations not considered during the training phase. To avoid these limitations, the energy manager can use an optimization-based approach. This approach uses a physical model coupled to an optimization solver to compute the set-points and on/off status of the units *all at once*. This approach has been widely used in chemical process operations where tight energy integrations exist, and it has generated annual savings on the order of several million dollars a year per facility (White 1998). We believe that the use of optimization-based EM technology will become critical in next-generation building systems where generation and storage components will be heterogeneous and where the number of operational degrees of freedom and the degree of uncertainty will significantly increase. We illustrate this increasing level of operational complexity in Figure 2.

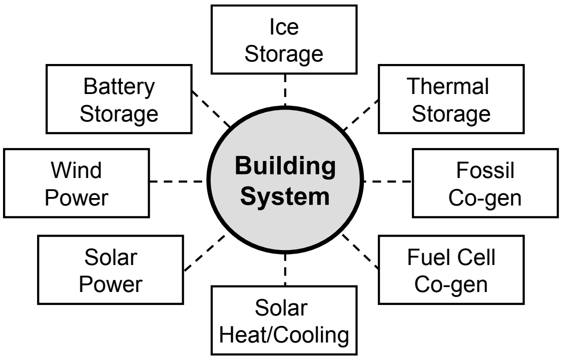

Figure 2: Generation and storage technologies portfolio for next-generation building system.

PROACTIVE ENERGY MANAGEMENT

We propose a proactive, optimization-based EM framework able to incorporate disturbance forecasts and predictive building models. The core of the proposed proactive EM framework is a *generalized* mixed-integer optimal control problem of the form:

$$\min_{w(\tau),u(\tau)} \frac{1}{N_s}\sum_{j=1}^{N_s}\left[\int_{t_k}^{t_k+T}\varphi(z_j(\tau),y_j(\tau),w(\tau),u(\tau),\chi_j(\tau))d\tau\right] \tag{1a}$$

$$\left.\begin{aligned} \text{s.t.}\ \frac{dz_j}{d\tau} &= \mathbf{f}(z_j(\tau),y_j(\tau),w(\tau),u(\tau),\chi_j(\tau))\\ 0 &= \mathbf{g}(z_j(\tau),y_j(\tau),w(\tau),u(\tau),\chi_j(\tau))\\ 0 &\geq \mathbf{h}(z_j(\tau),y_j(\tau),w(\tau),u(\tau),\chi_j(\tau)) \end{aligned}\right\}\ \tau\in[t_k,t_k+T] \tag{1b}$$

$$z_j(t_k)=x_k, \quad j=1,\ldots,N_s. \tag{1c}$$

The problem consists of finding optimal policies for a set of binary $w(\tau)$ and continuous control variables $u(\tau)$ that minimize the future expected cost of the system over the time horizon $\tau\in[t_k,t_k+T]$. The binary controls represent the on/off status for the equipment units at a given time. The continuous controls representing fuel, electric power, cooling water, air flow rates, and so on. The cost can be a composite objective function including fuel and electricity costs and storage utilization. The model is represented in general form as a set of differential and algebraic equations (DAEs) where $x(\tau)$ and $y(\tau)$ are the differential and algebraic state variables, respectively. The states represent zone and wall temperatures, concentrations, pressures, and humidities, among others. The current state of the building system at time t_k is given by x_k. The building system is affected by a set of uncertain disturbances $\chi_j(\tau)$, $\tau\in[t_k,t_k+T], j=1,\ldots,N_s$ that need to be forecasted. The proactive energy manager accounts for uncertainty in these forecasts by considering N_s realizations or *scenarios* which are obtained by sampling a given probability distribution $\mathcal{P}$. Incorporating uncertainty information can be critical to enhance the robustness of the operating policies. The EM optimization problem is solved in a closed-loop manner by updating the current building states and forecasts on a moving horizon window.

We have used simplified variants of the proposed EM framework in a building system equipped with electric and natural gas heating and in a photovoltaic hybrid system. These studies have focused on analyzing the impact of the forecast horizon length on economic performance (Zavala et al. 2009; Zavala, Anitescu, and Krause 2009). In the building system, we found that combining forecasts of electricity prices and of the ambient temperature leads to cost savings. The optimal timing at which to start the cooling at night directly depends on the ambient temperature expected the next day. In addition, special care needs to be taken to stay within the thermal comfort zone at all times. To analyze the effect of increasing the forecast horizon of the ambient temperature, we formulate an optimal control problem of the form in (1). The model considers the dynamic response of the building internal temperature and of the building wall and trying to find the optimal temperature set-point that minimizes the heating and cooling costs under a given electricity price structure. The ambient temperature enters the model through a boundary condition at the external face of the wall.

In Figure 3 we present the effect of the forecast horizon of the proactive EM problem (1) on the relative energy costs for two insulation levels. As can be seen, for a purely reactive strategy, the relative costs (compared to the optimal policy with an infinite horizon) can go as high as 24%. In addition, we observe that a horizon of 1 day is sufficient to achieve the minimum potential costs. The reason is that the thermal mass of the building cannot be used for a long time because there exist losses through the wall. In fact, we found that if the building insulation is enhanced, the costs for the purely reactive strategy increase significantly. On the other hand, when the building is poorly insulated, increasing the forecast horizon does not reduce the costs. In other words, *the economic potential of adding forecast information is tightly related to the ability to store energy in the system and to use it during off-peak times.* Another potential benefit of using forecast information is to minimize the number of startups and shutdowns of equipment such as light bulbs, water chillers, and gas furnaces, thereby *enhancing the responsiveness of the automation system, avoiding saturation of actuators, and reducing equipment wearing.* To illustrate the effect of productiveness on storage management and equipment wearing, we also present simulation results on a photovoltaic cogeneration system. The photovoltaic system is coupled to two storage options described in (Ulleberg 2004; Zavala, Anitescu, and Krause 2009) and is sketched in Figure 4. The first storage option has a large capacity but low round-trip efficiency (hydrogen with 70% efficiency), while the second has a small capacity but high efficiency (battery with 90% efficiency). The operating principle of this system is similar to that of other multistorage systems, such as photovoltaic-compressed air-battery or wind-hydrogen-hydrothermal systems. In these systems, it is necessary to decide the best strategy to store the intermittent power input in order to minimize power losses and satisfy a given load. To compute the optimal policies, we formulate problem (1) with forecast horizons ranging from 1 hour to 14 days. In the middle graph of Figure 4, we present the effect of increasing the horizon on the relative operating costs (using a one-year forecast policy as reference). In the bottom graph of this same figure we show the effect of the forecast horizon on the power profiles of the fuel cell system. Note that: (1) the relative operating costs decay quickly to zero as the horizon is increased; (2) for a purely reactive strategy (1 hr), the relative costs can go as high as 300%; and (3) the close-to-optimal costs can be obtained with a relatively short forecasts (1-14 days). The economic penalty of using a forecast of 1 day is just an increase of 10% in relative costs, whereas the penalty for a forecast of 12 hr goes up to 31%. For this system, as the horizon is increased, it is possible to exploit the more efficient battery storage system to reduce the power losses of the hydrogen storage loop. Note also that, as the horizon is increased, the power profiles of the fuel cell become more smooth, reducing equipment wearing.

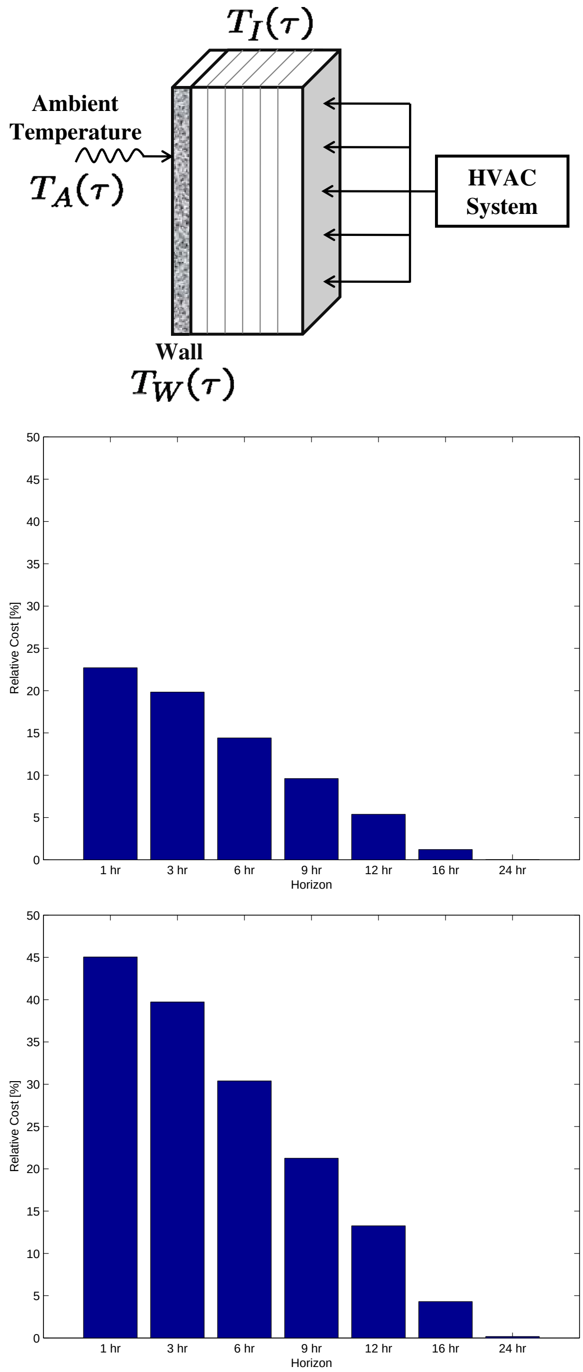

Figure 3: Schematic representation of building system (top). Impact of forecast horizon on energy costs of building HVAC system with base insulation (middle) and enhanced insulation (bottom).

Physics-Based Vs. Empirical Building Models

The DAE model constraints can contain detailed heat and mass transfer in the building zones, HVAC balances, and thermodynamic relations (e.g., steam tables). Alternatively, the model might incorporate only aggregated en-

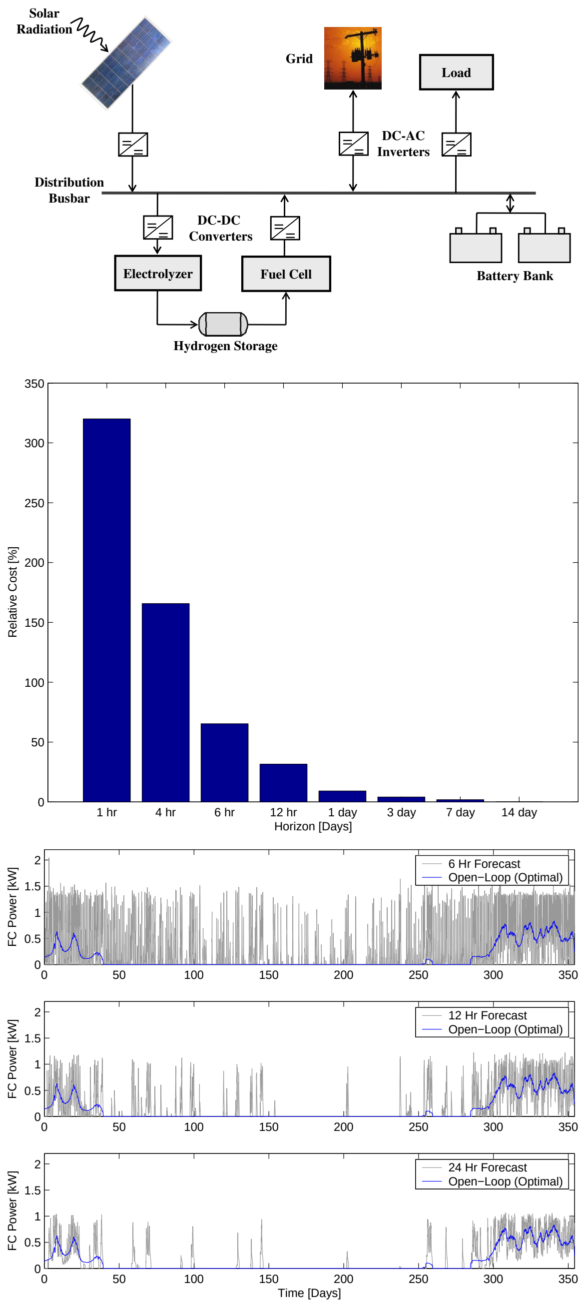

Figure 4: Schematic representation of photovoltaic cogeneration system (top). Effect of forecast horizon on operational costs of multi storage hybrid system (middle). Impact of forecast horizon on fuel cell power output (bottom).

ergy and mass balances. The degree of model sophistication is system dependent and limited by implementation and development costs. Next-generation building systems will be designed by using advanced models such as EnergyPlus or TRNSYS. These models can then be reused and exploited by the EM system. On the other hand, the development of a sophisticated model for an already installed building system might not be justifiable. In this case, it is possible to resort to aggregated or *lumped* energy balances that are relatively easy to build and implement using online temperature measurements (Braun 1990). However, since some physical information is lost in the aggregation procedure, translating the aggregated state variables to implementable set-points can be complicated.

Trade-Offs between Investment Costs and Operational Savings

In previous studies, we have found that the potential operational rewards of advanced EM technology are directly related to the ability of building to *isolate itself* from the power grid. In other words, if electricity can be produced on-site at a lower price compared to that of the power grid, then the investment in EM technology may pay off. Clearly, it is clear that the EM should be provided with extra equipment to beat the power grid prices. Therefore, large investments in additional cogeneration, heat integration, and storage units are necessary. As a consequence, there exists a trade-off between investment costs and operational savings. In order to justify the installation of additional units, the annualized investment cost must be lower that the annual operational savings achieved by the EM. These savings in turn depend on the price margins. In addition, natural gas and diesel prices come into play if integrated heat and power cogeneration systems are considered. All these imply that design and retrofitting tasks of the building and of the utility system must be performed in conjunction with detailed closed-loop EM studies considering historical disturbance data and climate change forecasts. Capturing climate changes is particularly critical because the seasons are likely to displace in both space and time. This can render installed technologies useless throughout their lifetime.

Forecasting Dynamic Disturbances

A proactive EM system relies on accurate predictions of disturbances. These disturbances can be predicted by constructing empirical (e.g.; autoregressive, neural networks, agents) and mechanistic (e.g., physics-based) models. Ambient conditions such as temperature, humidity, cloud cover, and wind speed can be predicted reliably by numerical weather prediction (NWP) models. These predictions are subject to uncertainties stemming both from initial atmospheric conditions and from imperfections in the numerical models (Kalnay 2002). For ac-

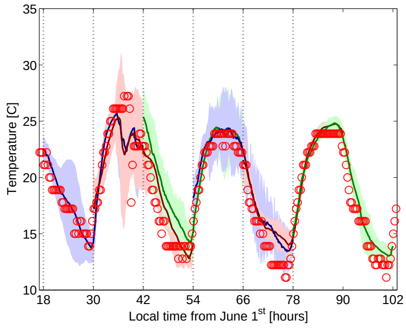

Figure 5: Temperature predictions using WRF.

curate predictions, the simulation is restarted at least every day from a new state corrected for errors by incorporating observations. This task is called data assimilation. Forecasting with NWP models is challenging because it is computationally expensive and requires large computational resources. In addition, national weather centers do not provide high-resolution forecasts for uncertainty information (Chen et al. 2006). In a recent study, we have performed an experiment for temperature forecasting using the Weather Research and Forecasting (WRF) model (Zavala et al. 2009). These forecasts were used in proactive EM strategies for building systems. In Figure 5 we show the temperature forecasts for four days in June 2006 together with their uncertainty for a location in Chicago, IL together with real observations (denoted by circles). We remark that the uncertainty interval (colored shade) closely encapsulates the observations and therefore the uncertainty model describes the forecast errors. This is important for achieving robustness of a proactive EM system.

The construction of empirical models to forecast ambient conditions relies on the availability of *on-site* measurement data. In addition, our experience is that empirical models cannot compute accurate day-ahead forecasts of volatile trends such as the ambient temperature. Empirical forecasts remain accurate only for a few hours, which might nevertheless be sufficient depending on the application, for instance, in Figure 3 it is clear that a forecast horizon of a few hours can reduce costs. Empirical models are also the only alternative to forecast disturbances when no mechanistic model exists. For instance, for forecasting natural gas prices, occupancy rates, and electricity and hot water demands, and occupant behaviors. To forecast these disturbances, one can resort to neural networks (Miller, Sutton, and Werbos 1990), autoregressive time-series models (Box, Jenkins, and Reinsel 1994), or kernel-based models (Rasmussen and Williams 2006). Complex occupant behaviors can be captured using agent-based models (ABMS).

Capturing Electricity Market Effects

The EM system must account for two-way market interactions with the power grid. Buildings vary widely in terms of physical, functional, occupational, ownership, management, decision-making, and market structure aspects. All of these factors will influence the total electricity demand, its dynamic pattern, the electricity contract with the distribution company, and eventually, wholesale power prices. The *real-time* price signals emerging from these supply and demand interactions will influence the optimal operating strategies of the EM system.

Modeling techniques are required to analyze and forecast market interactions and to obtain deeper insights into the impacts and various feedbacks arising between the building and the power grid. This information will be particularly critical with the advent of the smart grid. Market models can be used to analyze the energy consumption pattern of EM systems and the interaction with the power grid under various conservation scenarios and demand-side schemes, including reliability-based (load response) and market-based (price response) programs. These tools can also be used to capture smart grid interactions. Market behavior can be predicted by using multilevel optimization (game theoretical) problems that capture economic interactions among multiple players. However, these type of models assume that the market is in equilibrium, which might not be a reasonable assumption in building operations. An alternative are agent-based models such as the Electricity Market Complex Adaptive System (EMCAS) model (Conzelmann et al. 2005).

The market model lying at the interface between the grid and the EM system (see Figure 1) can provide price-demand curve information that can be exploited by the EM system. The market model can *switch* to a new contract with the utilities depending on the season or changing building activities thus updating the price-demand curve. At the same time, the EM system can provide candidate demand profiles so that the market model can determine the new contract. This two-way interaction will be adapted based on the most current market conditions that rely mostly on the season and on fuel prices at the unit commitment and energy dispatch level.

On-Line Solution of Energy Management Problem

The EM problem (1) is a mixed-integer optimal control problem (MIOCP). This is one of the most complex classes of optimization problems. The complexity stems from the presence of integer variables, the large dimensionality of building models, and the large number of decision variables or degrees of freedom. Solution methods for mixed-integer optimal control problems such as (1) are not well developed, and existing approaches (Sager 2009) do not readily take advantage of the special structure of the EM problem. To solve these problems, it is necessary to extend existing techniques. Time discretization of problem (1) results in a mixed-integer nonlinear program (MINLP) with a large number of integer variables. Traditional methods for MINLPs, such as branch-and-bound (Goux and Leyffer 2002), outer approximation (Duran and Grossmann 1986), or LP/NLP-based branch-and-bound (Bonami et al. 2008; Abhishek, Leyffer, and Linderoth 2006), can in principle be applied to the resulting MINLP. These methods have in common that they perform a tree search to resolve the combinatorial nature of the binary decision variables. Because of this, these methods are unlikely to be successful for problems where either n_w or n_t is large, and new techniques must be developed. The EM problem, however, has special structure that can make it feasible to solve these problems in a reasonable amount of time. For example, the following techniques can be used to address this problem:

- **Warm-Starting Techniques for MINLP.** Traditionally, MINLPs are solved in isolation. However, as we move from subsequent solution times for the EM problem t_k to t_{k+1}, we can exploit the fact that the solution to these two problems is closely related. We will explore a windowing technique similar to techniques used in nonlinear model predictive control (Zavala and Biegler 2009). We will solve the resulting MINLP every hour over a 24 hour horizon, implement the resulting control over the next hour, and then solve the next MINLP, starting at t_{k+1} over a 24-hour horizon. This knowledge can be used in several ways. First, it is possible to update the current integer controls by solving a *continuous* control problem, and obtain a new upper bound. Second, one can reuse branching information such as pseudocosts (Gauthier and Ribière 1977) that improve the tree search. One can also use feasibility heuristics such as local branching (Fischetti and Lodi 2002).

- **Approximation and Hierarchical Approaches.** In the previous strategy we have assumed that the entire MIOCP has to be solved at each time. An alternative approximation that we will explore is obtained by hierarchical decomposition and linearization of the control problem (1). Discretizing the linearized problem leads to a mixed-integer linear program (MILP), which can be solved with solvers such as CPLEX or CBC. These solvers can handle problems several orders of magnitude larger than MINLPs. The drawback of this approach is that the nonlinear dynamics of the control problem are only approximately taken into account so the integer solution might not be optimal. It is possible to mitigate the effect of this approximation by fixing the binary controls and resolving (1) for the nonlinear controls by solving a nonlinear programming problem (NLP) with solvers such as KNITRO (Byrd, Gilbert, and Nocedal 2000) or IPOPT (Wächter and Biegler 2006). In other words, the MILP updates the binary controls while the NLP updates the continuous controls.

These strategies can enable the solution of large-scale EM problems with thousands of states and continuous controls and hundreds of binary controls *on standard personal computers*. Because of the large number of degrees of freedom encountered in buildings, however, the optimization solvers need access to derivative information to capture the effect of controls on the building states. Current building simulation packages such as EnergyPlus (Crawley et al. 1999; Ellis, Torcellini, and Crawley 2007) or TRNSYS (Klein, Duffie, and Beckman 1976) do not provide these capabilities. With this, the building models can be treated only as *black boxes* by optimization solvers, forcing the user to rely on logic-based strategies, genetic algorithms, or derivative-free optimization. These approaches use repetitive and time-consuming simulations of the building model and are highly inefficient because the number of simulations is related to the number of decision variables. *This seriously limits the scope of on-line energy management in building systems.* We believe that building simulation packages should incorporate capabilities to interface them with automatic differentiation capabilities (Griewank 2000). This will enable the implementation of highly sophisticated EM applications with detailed building models.

CONCLUSIONS

In this paper, we present a proactive energy management framework that integrates predictive dynamic building models and day-ahead forecasts of disturbances affecting efficiency and costs. We argue that physics-based building models will become crucial for the operation of next-generation buildings because they will enable an efficient management of resources and an accurate prediction of the daily electricity demand profile. One of the main advantages of the proactive framework is that it enables integration of forecast models of weather conditions, prices, and utilities demands. In addition, it can capture occupant behaviors and market interactions by using agent-based models. The framework is based on the *on-line* solution of mixed-integer nonlinear programming problems. We present strategies that can enable the solution of these challenging problems on standard personal computers. We claim that a large adoption level of this technology can improve the predictability of the electricity demand at high-level grid operations such as unit commitment and economic dispatch which can be used to minimize reserves. As part of future work, we are interested in testing the presented concepts in a comprehensive case study.

ACKNOWLEDGMENT

This work was supported by the Department of Energy, under Contract No. DE-AC02-06CH11357.

REFERENCES

Abhishek, K., S. Leyffer, and J. T. Linderoth. 2006. "FilMINT: An outer-approximation-based solver for nonlinear mixed integer programs." Preprint ANL/MCS-P1374-0906, Mathematics and Computer Science Division, Argonne National Laboratory.

Bonami, P., L. T. Biegler, A. R. Conn, G. Cornuéjols, I. E. Grossmann, C. D. Laird, J. Lee, A. Lodi, F. Margot, N. Sawaya, and A. Wächter. 2008. "An algorithmic framework for convex mixed integer nonlinear programs." *Discrete Optimization*. To appear.

Box, G. E. P., G. Jenkins, and G. Reinsel. 1994. *Time Series Analysis: Forecasting and Control*. New Jersey: Prentice-Hall.

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

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

Byrd, R. H., J. Ch. Gilbert, and J. Nocedal. 2000. "A trust-region method based on interior-point techniques for nonlinear programming." *Math. Programm.* 89:149–185.

Chen, F., M. Tewari, H. Kusaka, and T.T. Warner. 2006. "Current status of urban modeling in the community Weather Research and Forecast (WRF) model." *Joint with Sixth Symposium on the Urban Environment and AMS Forum*, Volume 1.

Conzelmann, G., G. Boyd, V. Koritarov, and T. Veselka. 2005. "Multi-agent power market simulation using EMCAS." *IEEE PES General Meeting*.

Crawley, D. B., L. K. Lawrie, C. O. Pedersen, R. J. Liesen, D. E. Fisher, R. K. Strand, R. D. Taylor, R. C. Winkelmann, W. F. Buhl, Y. J. Huang, and A. E. Erdem. 1999. "EnergyPlus: A New-Generation Building Energy Simulation Program." *Proceedings of Building Simulation '99*, Volume 1. 81–88.

Duran, M. A., and I. Grossmann. 1986. "An Outer-Approximation Algorithm for a Class of Mixed-Integer Nonlinear Programs." *Mathematical Programming* 36:307–339.

Ellis, P. G., P. A. Torcellini, and D. B. Crawley. 2007. "Simulation of energy management systems in EnergyPlus." *Building Simulation*.

Fischetti, M., and A. Lodi. 2002. "Local Branching." *Mathematical Programming* 98:23–47.

Gauthier, J. M., and G. Ribière. 1977. "Experiments in mixed-integer linear programming using pseudo-costs." *Mathematical Programming* 12:26–47.

Goux, J.-P., and S. Leyffer. 2002. "Solving large MINLPs on computational grids." *Optimization and Engineering* 3:327–346.

Griewank, A. 2000. *Principles and Techniques of Algorithmic Differentiation*. Philadelpia, PA: SIAM.

Kalnay, E. 2002. *Atmospheric Modeling, Data Assimilation and Predictability*. Cambridge University Press.

Klein, S. A., J. A. Duffie, and W. A. Beckman. 1976. "TRNSYS: A transient simulation program." *ASHRAE Transactions* 82:623–633.

Miller, W. T., R. S. Sutton, and P. J. Werbos. 1990. *Neural Networks for Control*. Cambridge: MIT Press.

Rasmussen, C. E., and C. K. Williams. 2006. *Gaussian Processes for Machine Learning*. Cambridge: MIT Press.

Sager, S. 2009. "Reformulations and algorithms for the optimization of switching decisions in nonlinear optimal control." *Journal of Process Control* 19 (8): 1238–1247.

Ulleberg, O. 2004. "The importance of control strategies in PV-hydrogen systems." *Solar Energy* 76:323–329.

Vosen, S. R., and J. O. Keller. 1999. "Hybrid energy storage systems for stand-alone electric power systems: optimization of system performance and cost through control strategies." *International Journal of Hydrogen Storage* 24 (12): 1139–1156.

Wächter, A., and L. T. Biegler. 2006. "On the Implementation of a Primal-Dual Interior Point Filter Line Search Algorithm for Large-Scale Nonlinear Programming." *Math. Program.* 106:25–57.

White, D. C. 1998. "Online Optimization: What have we learned?" *Hydrocarbon Processing* 77:55–69.

Zavala, V. M., M. Anitescu, and T. Krause. 2009. "On the Optimal On-Line Management of Photovoltaic-Hydrogen Hybrid Energy Systems." *In Proceedings of 10th International Symposium on Process Systems Engineering*.

Zavala, V. M., and L. T. Biegler. 2009. "Nonlinear programming strategies for state estimation and model predictive control." *In Nonlinear Model Predictive Control*, pp. 419–432.

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