

# Model Predictive Control for Market-Based Demand Response Participation <sup>\*</sup>

Faran A. Qureshi <sup>\*</sup> Tomasz T. Gorecki <sup>\*</sup> Colin N. Jones <sup>\*</sup>

<sup>\*</sup> *Laboratoire d'automatique, EPFL, Lausanne, CH (e-mail: {faran.qureshi, tomasz.gorecki, colin.jones}@epfl.ch)*

---

**Abstract:** In this study, we investigate the maximum possible profit for a commercial office building participating in New York's Day-Ahead Demand Response (DADR) program. We formulate an optimal control problem, assuming perfect knowledge of future weather, occupancy, and day-ahead electricity price predictions to examine this potential benefit. Then, a practical control strategy based upon the framework of Model Predictive Control (MPC) is proposed, which enables a building to participate in the DADR program. The controller decides once every day, whether or not to participate in the Demand Response (DR) event, and then optimizes the electric consumption to increase savings. A simulation study is carried out using a building model extracted from an EnergyPlus model, real measured weather data, and real day-ahead spot market price data for New York. Savings in the range of 23% to 33% are reported.

Keywords: Demand Response; Building Control; Model Predictive Control.

---

## 1. INTRODUCTION

The growing addition of renewable energy sources in the electricity grid increases the variability of generation. This increased variability and the expected growth of the electricity consumption, require higher levels of regulating power in order to meet the peak load demands and to balance production and consumption on the electric grid. Traditionally, the required operating reserves needed to react to contingencies on the network were provided by energy producers. The independent system operators (ISOs) are looking more and more to the demand side to perform this regulation. A variety of demand response (DR) programs have been created and can be classified in price-based programs where a variable tariff is the driver, and incentive-based programs where temporary financial incentives are given for participation [Albadi and El-Saadany, 2008]. The main purpose of these programs is to incentivize participation of demand-side resources in the regulation of the grid. Early programs include direct control of flexible loads such as electric boilers, *e.g.*, the one proposed by the company Voltalis [Hull, 2010]. More recently, control of more sophisticated loads have been proposed to provide DR services. [Oldewurtel et al., 2013] highlights that buildings are fit for providing services on slow timescales in the range of hours. We consider in this work the participation of commercial buildings in DR programs.

Model Predictive Control (MPC) is a popular advanced control technique in which control decisions are taken via an optimization problem over future trajectories in receding horizon. Numerous authors have considered MPC for control of building Heating, Ventilation, and Air Conditioning (HVAC) systems. It is particularly appreciated for its natural ability to deal with constraints and incorporate

future prediction of weather [Oldewurtel et al., 2012]. [Ma et al., 2012] reports significant savings in the operation of the cooling system of the Merced University campus and overall coefficient of Performance (COP) increase of 19% by using MPC controllers. Savings reported are sometimes not deemed sufficient to offset the costs involved in the implementation of these complex controllers. Offering services to the electric grid is a way of leveraging additional revenues that further justify their use. Participation of real commercial buildings in DR programs has been reported in [Page et al., 2011]. It mainly involves using simple strategies to dim lights and change temperature setpoints. Modest peak demand reductions are reported and the need for tailored DR strategies is underlined.

The contribution of this paper is two-fold: firstly, the maximum benefit for a commercial office building participating in the day-ahead demand response (DADR) market of the New-York ISO (NYISO) [NYISO, 2003] is evaluated. This market-driven program allows consumers to bid demand reductions on the energy market. Secondly, a practical control architecture based on MPC is proposed for participation. The control architecture comprises two decision layers: one, executed every day, decides on the participation of the building in the DR market. A second layer operates the building in real-time, taking into account the decisions of the upper layer. This control scheme helps evaluate the saving potential under realistic settings.

The rest of the paper is structured as follows. Section 2 summarizes the main characteristics of the NYISO DR program. The building model used in this study is described in Section 3. The control setup is presented in Section 4. Section 5 gives an overview of the results of the simulation study. Conclusions are drawn in Section 6.

---

<sup>\*</sup> This work has received support from the Swiss National Science Foundation under the GEMS project (Green Energy Management of Structures, grant number 200021\_137985)

Min. duration of participation	1h
Max. duration of participation	8h
Payment	Day-ahead Clearing price
Minimum bid	1 MW
Bid increment	1 MW
Bid price floor	75 \$/kWh

Table 1. Characteristics of DADR Program

## 2. NEW YORK'S DAY-AHEAD DR PROGRAM

A variety of DR programs have recently appeared in energy markets. The New York ISO is one of the pioneering system operators proposing DR programs, with four programs for medium and large consumers. In this work, we focus on their Day-Ahead DR program (DADR). It allows customers to bid consumption reductions on the day-ahead market similarly to energy producers. The main reason to focus on this program is that it allows the bidders to choose the participation times freely. Therefore it is one of the most flexible programs from the consumer's point of view. Characteristics of the DADR program are summarized in Table 1 and all regulations can be found in [NYISO, 2003]. Features of this program are outlined below:

- Consumers post bids on the day-ahead market comprising a reduction level and a price above which the reduction can be performed.
- If the bid is accepted, the consumer has to reduce its consumption by the bid amount during the participation time.
- The reduction is computed with respect to a baseline consumption, which is known when the bid is posted. This baseline basically depends on consumption during days preceding the DR participation. Hence, the lower the consumption during the past days, the more difficult it is to reduce consumption. Additionally, a correcting factor can be taken into account to adjust for weather on the day of the event. Namely, the baseline is adjusted up or down by comparing the consumption during the hours preceding the DR event with the same hours during previous days. Details are given in Section 4.
- The payment for the reduction is the locational day-ahead clearing price (and not the bid price). In case of failure to decrease consumption, a penalty has to be paid.

A challenging feature of the DADR market is that even when the building does not participate on the DADR market, the current consumption still impacts the potential benefits of future participation through the baseline. This occurs on a timescale that exceeds usual prediction capabilities for the weather or price. Hence the objective of this work is to first evaluate the maximum potential benefits of DADR participation and secondly to propose a more practical participation strategy using receding horizon control on a short horizon, but which still tries to capture the longer term effects.

## 3. BUILDING MODEL

This section presents the building model used in this study. EnergyPlus[Crawley et al., 2001] is a widely used detailed building energy simulation tool used for the thermodynamic modeling of buildings and HVAC systems. It is

considered to give a good approximation of the thermal behavior of buildings, but is not suitable for control. We use the toolbox OpenBuild [Gorecki et al., 2014] to extract a linear state space model from an EnergyPlus building model. The toolbox uses a physical modeling approach to model the thermodynamics of the building. The effect of the HVAC system is modeled directly as independent heat fluxes input to the zones. This is justified by the fact that the building is equipped with independent terminal cooling units. The central heating is not explicitly modeled and we consider it has a fixed COP. For more details, we refer the reader to [Gorecki et al., 2014].

The model obtained with OpenBuild is reduced using classical balanced truncation method to model with a few states per zone. This gives a model of the form:

$$\begin{aligned} x_{k+1} &= Ax_k + B_u u_k + B_w w_k \\ y_k &= Cx_k \end{aligned} \quad (1)$$

where  $x_k \in \mathbb{R}^n$  is the state,  $u_k \in \mathbb{R}^m$  is the thermal energy input to the rooms,  $w_k \in \mathbb{R}^p$  is the disturbance input (including weather, long-wave radiations, solar gains, and internal gains), and  $y_k \in \mathbb{R}^q$  is the output containing the room air temperature of each zone at time step  $k$ . The conversion from electrical energy to thermal energy is modeled using a constant COP for the heating ( $\eta_h$ ) and cooling ( $\eta_c$ ) systems, and is given as:

$$e_k = \eta_h \sum_{u_k^i \geq 0} u_k^i - \eta_c \sum_{u_k^i \leq 0} u_k^i \quad (2)$$

where  $e_k$  is the consumption of electrical energy (kW) at time step  $k$ .

*Remark 1.* The time step index  $\{k\}$  can equivalently be written as the index  $\{d, h\}$  referring to the day  $d$ , and hour  $h$ . In the sequel we use both these notations interchangeably.

## 4. CONTROL SETUP

This section presents the optimal control problem, formulated to investigate the maximum potential benefit of a building participating in the New York's DADR program, and an MPC controller enabling a building to participate in this program.

### 4.1 Optimal Control problem

Our objective is to maximize the total profit of a building participating in the DADR program while maintaining acceptable comfort constraints in the building. Our cost include both the cost of the energy consumed, and the payments obtained during the DR participations. The cost of energy consumption is given by:

$$V_k^e = c_k^t e_k T_s, \quad (3)$$

where  $c_k^t$  is the electricity tariff (\$/kWh) at time step  $k$ , and  $T_s$  the sampling time in hours. DR payments are given by:

$$V_k^{dr} = c_k^d T_s (B_{d,h} - e_{d,h}), \quad (4)$$

where  $c_k^d$  is the day-ahead clearing price of electricity (\$/kWh) and,  $B_{d,h}$  is the baseline consumption at time step  $k$ . The baseline consumption for an hour  $h$  is the average energy consumption during hour  $h$  over a set of previous days  $\mathcal{S}_{d,h}$ , and is given by:

$$B_{d,h} = \beta_{d,h} \frac{1}{|\mathcal{S}_{d,h}|} \sum_{j \in \mathcal{S}_{d,h}} e_{d-j,h}, \quad (5)$$

where  $\mathcal{S}_{d,h}$  is the set of days used to compute the baseline, and  $\beta_{d,h}$  is a weather correction factor. In case  $d$  is a week day (a similar procedure applies for weekend days), the steps to select the days  $\mathcal{S}_{d,h}$  are:

- Starting two days before  $d$ , select the ten previous weekdays during which the building did not participate in a DR event during hour  $h$ .
- Among these, select the five days with highest consumption during hour  $h$ .

The weather correction factor  $\beta_{d,h}$  adjusts the baseline for the weather conditions on the day of the event. It is computed as the ratio of the energy consumption during a time period preceding the participation period (hours  $h_1$ , and  $h_2$ ) of the day  $d$ , to the average energy consumption during the same period for the days used in the baseline computation:

$$\beta_{d,h} = \frac{e_{d,(h_1,h_2)}}{\frac{1}{|\mathcal{S}_{d,h}|} \sum_{j \in \mathcal{S}_{d,h}} (e_{d-j,(h_1,h_2)})} \quad (6)$$

Furthermore, it is saturated so that  $0.8 \leq \beta_{d,h} \leq 1.2$ . Comfort is maintained by enforcing temperature constraints on the zone temperatures:

$$T_{min,k} \leq y_k \leq T_{max,k}, \quad (7)$$

where  $T_{min,k}$  and  $T_{max,k}$  are the minimum and the maximum allowable temperature in each zone at time step  $k$ , respectively. Constraint on the input actuators are given by

$$u_{min} \leq u_k \leq u_{max}, \quad (8)$$

where  $u_{min}$  and,  $u_{max}$  are the minimum and the maximum allowable actuator limits. The optimal control problem is formulated as:

$$\begin{aligned} V_{OC}(x_0, \{w_k\}, \{c_k^d\}) = \\ = \underset{\mathbf{x}, \mathbf{u}, \mathbf{s}, \delta}{\text{minimize}} \quad & \sum_{k=0}^{N_{OC}-1} V_k^e - \delta_k V_k^{dr} + V_k^s \\ \text{subject to} \quad & (1), (7), (8), \\ & \delta_k \in \{0, 1\}, \quad k = 0, \dots, N_{OC} - 1 \end{aligned} \quad (9)$$

where  $\delta_k$  is the binary variable indicating the status of DR participation at time step  $k$ . It is one if the building is participating at time step  $k$ , zero otherwise. The comfort constraint (7) is implemented as a soft constraint, and  $V_k^s$  is the quadratic penalty on the constraint violation. Solving (9) over a long period gives us the best participation strategy in DR. If the problem is solved on a long enough horizon, then it allows to both optimize our consumption during DR participation, and prepare the baseline when not participating. This means that a long horizon (more than a week to optimize the baseline) is desirable. In the case studied here, the dynamics of the system are relatively fast because we consider an air based system, but decisions have long lasting impact because of the effect of the baseline.

#### 4.2 Model Predictive Control Problem

In MPC, a measurement is obtained at time step  $k$ , a prediction model is used to predict the future trajectories of the system, and an optimization problem is solved

to compute a sequence of optimal control inputs over a predefined prediction horizon. The first optimal control input is applied and at the next time step, the procedure is repeated. An MPC controller is proposed to solve the problem presented in the previous section.

*Demand Response MPC* We propose an MPC strategy allowing the building to participate in the DADR program. Using a shorter receding horizon reflects the fact that, in practice, the predicting ability is limited by the lack of confidence in the long term (typically longer than two days) future weather predictions. The computational complexity of the resulting optimization problem also increases with the horizon. The MPC optimization problem solved at each time step  $k$  is exactly the same as in (9), except that the baseline  $B_k$  is a known constant and is an input to the optimization problem, due to the fact that the prediction horizon is shorter than two days. Since the prediction horizon is not long enough, the MPC controller is not able to optimize the future baseline. At each time step  $k$  the baseline  $B_k$  is computed using past energy consumption (5).

*Suboptimal Demand Response MPC* In practice, the participation in DR for the next day is decided once a day when the market is cleared. There is then no possibility to reoptimize for them as in a classical MPC scheme. Therefore, we propose a sub-optimal solution where the demand response optimization problem with binary variables is solved only once each day and the decision variables  $\delta_k$  are fixed over the prediction horizon. At each time step, a simplified version of the optimization problem (9) is solved, where the decision variables  $\delta_k$  are fixed as inputs to the optimization problem. In reality, if the horizon is more than a day, there still is the possibility to reoptimize for participation decisions beyond the first day of prediction, however, we keep this structure for the MPC controller because it reduces the computational load significantly, and as we will detail in the next sections, the controller decides most of the time to participate systematically, which implies that its decisions in closed loop are very consistent.

*Baseline increase* With a short horizon, the MPC controller cannot optimize the future baseline and DR participation status, thus limiting our ability to participate in DR. When not participating in DR, our consumption influences future baseline computations. Given the expression of the baseline computation (5), we make the approximation that our current consumption will influence the future profits from DR participation in a linear fashion, yielding:

$$V_k^b = \alpha_k e_k T_s \quad (10)$$

where  $\alpha_k > 0$  captures how the future payments for DR participation will change and indicates our confidence in future DR participation. A larger value of  $\alpha_k$  implies a higher confidence in future DR participation. For simplicity, we focus here on constant values of  $\alpha$ . The optimization problem solved at each time step is given by:

Building Type	Office Building
Location	Representing eight climate zones, including New York
Floor Area ( $m^2$ )	4979
No. of Zones	15
No. of Floors	3
Window Fraction	33%
Peak Occupancy (people/100 $m^2$ )	5.4

Table 2. Building Characteristics

$$\begin{aligned}
 &V_{DR}(x_0, w_k, c_k^d, \delta_k, B_k) = \\
 &= \underset{\mathbf{x}, \mathbf{u}, \mathbf{s}}{\text{minimize}} \quad \sum_{k=0}^{N_{DR}-1} V_k^e - (1 - \delta_k)V_k^b - \delta_k V_k^{dr} + V_k^s \\
 &\text{subject to} \quad (1), (7), (8), \quad k = 0, \dots, N_{DR} - 1
 \end{aligned} \tag{11}$$

## 5. SIMULATION CASE STUDY

### 5.1 Simulation Setup

We use the standard ASHRAE EnergyPlus model of a medium-size modern commercial office building for this study, taken from the reference building data base of the U.S. department of energy<sup>1</sup>. The characteristics of this building are summarized in Table 2. The building is described with typical usage schedules for equipment, lights, and occupancy. It is equipped with one centralized cooling system for each floor with electrically powered cooling coils. The zones are each equipped with individual terminal VAV units controlling the airflows. Therefore, it is realistic in our model to consider that we control the energy input to the rooms independently. Specific details of the simulation are outlined below:

- Simulations are performed on a two month period from June, 5. to August, 18. 2012.
- We use real measured weather data for New York, provided by the company WeatherAnalytics.
- Real day-ahead prices for Long Island, the most constrained part of the New York network are used in this study. This price data is obtained from the New York's ISO.
- The electricity tariffs published by the Long Island power authority are used, *i.e.*,  $c_k^t = 0.15\$/kWh$ .
- A maximum of eight hours of DR participation is allowed each day (*i.e.*, between 1pm and 9pm). The hours  $h_1$  and  $h_2$  used in (6) are hour 10, and 11.
- Participation is not allowed if the day-ahead price is below a minimum floor price level, which is  $75\$/MWh$  in this case.
- The comfort constraint  $T_{min,k}$ , and  $T_{max,k}$  is  $20^\circ C$ , and  $23^\circ C$  during office hours, and  $15^\circ C$ , and  $27^\circ C$  otherwise.

### 5.2 Validation of the model

We aim to compare the thermal model of the building since the cooling system is not modeled. We simulate the EnergyPlus model with the building's built-in controller

<sup>1</sup> [http://www1.eere.energy.gov/buildings/commercial/ref\\_new\\_construction.html](http://www1.eere.energy.gov/buildings/commercial/ref_new_construction.html)

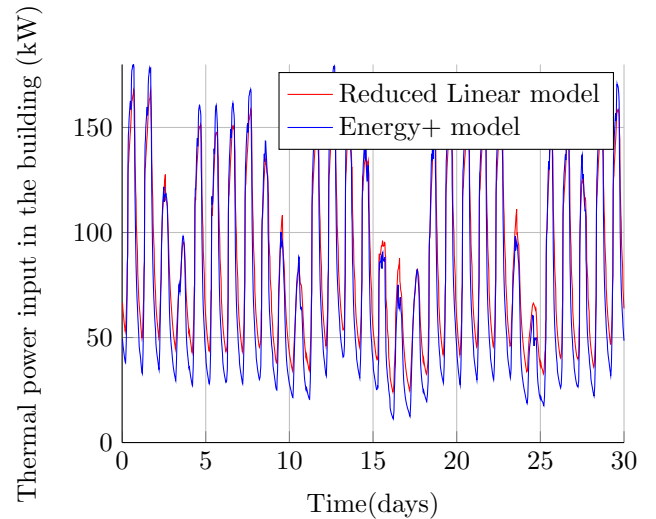


Fig. 1. Validation of the extracted model

to maintain a constant indoor temperature of  $23^\circ C$  in the control zones. On the other hand, a simulation is performed with our linear model and an MPC controller maintaining the same temperature. The simulation covers the full month of July 2012. Figure 1 shows the aggregated cooling power in the two cases. We see that the power requirements of the building are successfully captured by our linear model. A tendency to overestimate the lowest consumption peak during the day is noticeable. The model reduction itself impacts the power consumption by less than one percent in the worst case. This justifies the use of the reduced linear model as our prediction model.

### 5.3 Simulations

The study aims at investigating the following questions:

**Q1: Bound on maximum savings:** What is the maximum potential benefit for a building participating in New York's DADR program, utilizing only the available thermal storage of the building, and what is the impact of the participation strategy on this maximum saving potential? This will help us decide if this DR program is beneficial for a typical commercial office building.

**Q2: Bound on maximum savings under limited forecasts:** What is the maximum potential benefit under a more realistic setting, *i.e.*, without assuming long-term future knowledge of weather, occupancy, and price predictions, and what is the impact of  $\alpha$  (confidence in future participation) on this maximum saving potential? This will indicate how close to the optimum we can get with limited knowledge of the future.

To answer these questions, we consider three simulation cases:

**Minimum Cost Trajectory without DR [no DR]:** This is the minimum cost trajectory without DR over a simulation period of two months. It is found by solving the optimal control problem over the complete two months period (9) with  $\delta_k = 0$  being enforced, *i.e.*, no DR participation. We assume perfect knowledge of the weather, and occupancy predictions on this period. This is the base case to compare the cost of energy consumption and potential DR income. The results in this case are very close to using the built-in EnergyPlus controller, which comforts

its choice as our base case.

**Minimum Cost Trajectory with DR [DR]:** This is the minimum cost trajectory among all possible trajectories with DR over a simulation period of two months. We use this simulation to compute the bound on the maximum savings achievable. It is obtained by solving (9), which is a large scale mixed-integer bi-linear problem. There, the sets  $S_{d,h}$  depend on the participation decisions  $\delta_k$ . To circumvent this difficulty, we use fixed participation hours in advance, which leaves a bi-linear problem to solve. For the optimization, we consider the weather rule inactive, *i.e.*,  $\beta_k = 1$ . DR payments with the weather rule are then post-computed. The problem then boils down to a large quadratic program. It is solved over the two months, assuming perfect knowledge of future weather, occupancy, and day-ahead price.

**Minimum Cost Operation with DR and Limited Forecast [DR & MPC]:** This is the minimum cost trajectory under closed-loop MPC control with limited forecast horizon. This simulation is useful to evaluate the maximum savings under limited forecasts. We solve (9) over an horizon of two days, once per day. Notice that with such horizon, the baseline (5) becomes a linear function of  $e_{d,(h_1,h_2)}$ . In turn, the payment (4) is a linear function of the optimization variables and (9) becomes a MIQP with a limited number of integer variables, which is therefore tractable. The participation decisions for the first day are then fixed and fed to a second MPC controller (11) that optimizes the consumption over an horizon of one day in a receding horizon fashion. Perfect knowledge of future weather, occupancy, and day-ahead price predictions is again assumed, but only on an horizon of two days.

#### 5.4 Analysis of Results

**Q1: Bound on maximum savings:** The maximum savings achievable for the commercial office building participating in the DADR program is 33%. This result is obtained by comparing the results of the minimum cost trajectory without DR [no DR], with the minimum cost trajectory with DR [DR], as depicted in Fig. 2, which shows the cumulative expenses over the simulation period. The results are normalized with respect to the cost of the minimum cost trajectory without DR [no DR]. The total DR payment (top red line) amounts to 40% of the total minimum energy cost resulting in a net saving of 33%, as illustrated in Fig. 3. In Fig. 2, the cumulative profit generated by DR participation progresses in steps, *i.e.*, only when the building is participating in a DR event. The cumulative DR profit increase sharply between days 30 to 34, and days 41 to 45. The reason is that these days have particularly high day-ahead prices. On the contrary, towards the end of the simulation period, the cumulative DR profit increases very slowly because of low day-ahead prices.

In order to investigate the impact of the participation strategy on the savings potential, we performed a set of simulations [DR] with different floor prices for participation ranging from 75\$/MWh to 200\$/MWh. Savings compared to the minimum cost trajectory without DR [no DR] are depicted in Figure 4. The maximum benefit is obtained by participating systematically above the minimum allowed floor price of 75\$/MWh. This suggests

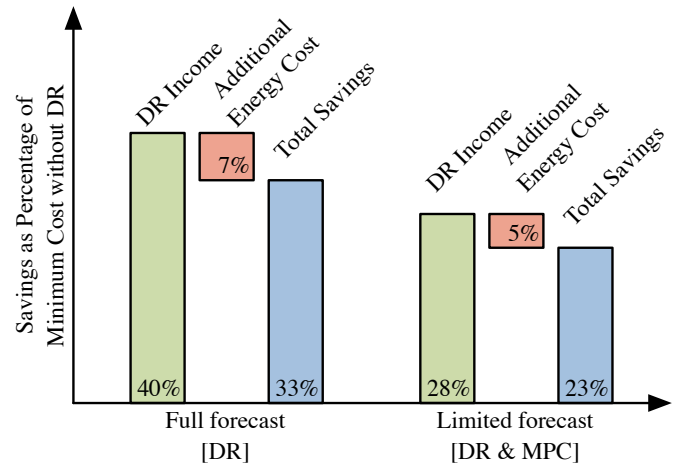


Fig. 3. Savings as percentage of minimum cost without DR

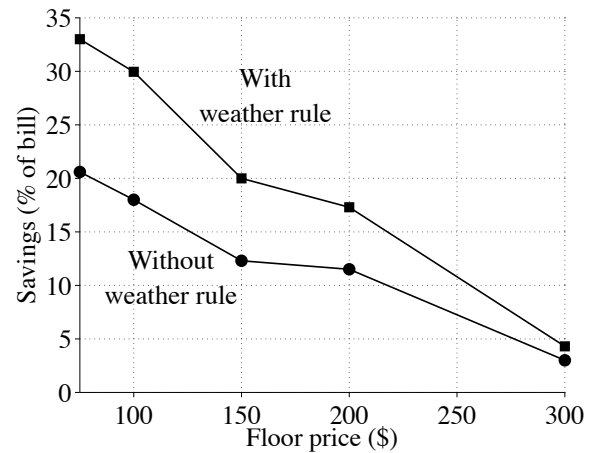


Fig. 4. Savings compared to the minimum energy trajectory [no DR] against minimum bidding level

that a systematic bidding strategy at the minimum price is reasonable, and this simple strategy is the one adopted in [DR] in the rest of the simulations.

**Q2: Bound on maximum savings under limited forecasts:** The maximum savings achievable for the commercial building, under limited forecast are 23%. This is obtained by comparing the minimum cost trajectory without DR [no DR] with the minimum cost operation with DR and limited forecast [DR & MPC], as depicted in Fig. 2. The cumulative DR income for [DR & MPC] (top blue line) shows a similar trend to the cumulative DR income for the minimum cost trajectory [DR] (top red line). This shows that the MPC controller almost always decides to participate in the DR event. However, the total DR profit for [DR & MPC] amounts to about 28% of the total minimum energy cost resulting in a net saving of 23%, as illustrated in Fig. 3.

Finally, the impact of the tuning parameter  $\alpha$  (our confidence in future DR participation) was studied. A set of simulations [DR & MPC] with uniformly increasing values of  $\alpha$  was performed. Results are depicted in Fig. 5. It can be seen that as expected, by increasing  $\alpha$  the total energy cost increases, whereas the profit increases to its maximum, and then starts to decrease again. Indeed, if the future baseline is increased too much, the overall energy consumption of the building will increase, thus reducing

Fig. 2. Normalized Cumulative Cost


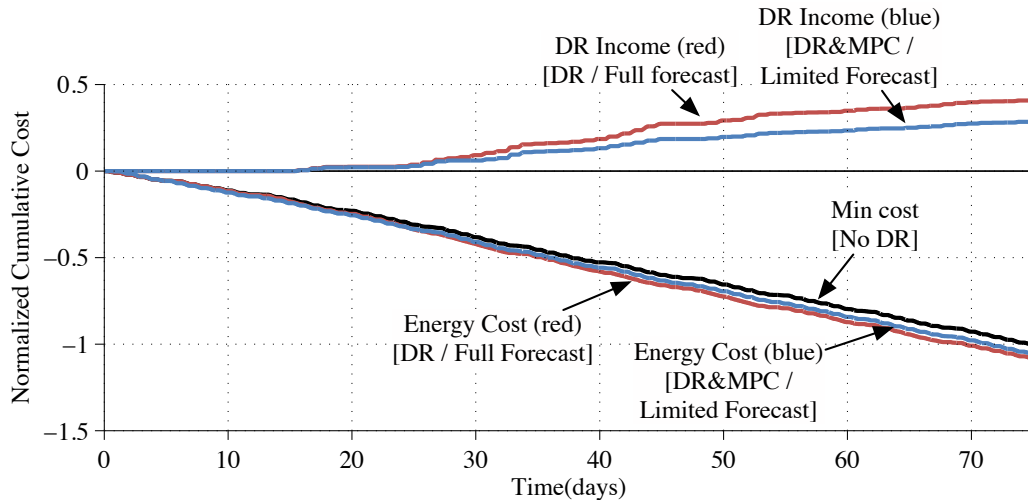


Fig. 5. Percentage increase in savings vs percentage increase in total energy cost (for values of  $\alpha$  between 0 and  $0.95c_k^t$ )

the total increase in profits. We used a fixed optimum value of  $\alpha = 0.8c_k^t$  in our simulation [DR & MPC] result depicted in Fig. 2. However, it is expected that a time-varying  $\alpha$  may improve the performance, especially if the building is not able to participate often.

## 6. CONCLUSION

We used optimal control to investigate the maximum possible savings for a commercial office building participating in New York's DADR program. The simulation results showed that significant savings are attainable. Furthermore, a practical MPC controller was proposed to enable a building to participate in the DADR program under realistic settings. It was shown that even with a limited availability of forecasts, the building was able to make significant savings. All these savings come by using only the thermal storage capacity of the building. The available storage in the electrical equipment of the building and the HVAC system is expected to increase this saving potential.

## REFERENCES

M.H. Albadi and E.F. El-Saadany. A summary of demand

response in electricity markets. *Electric Power Systems Research*, 78(11):1989–1996, November 2008.

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.

Tomasz T. Gorecki, Faran A. Qureshi, and Colin N. Jones. OpenBuild : An integrated simulation environment for building control. Technical report, 2014. URL <https://infoscience.epfl.ch/record/197863>.

L. Hull. Evaluating the business case for micro demand response and energy saving. Technical report, International Energy Agency Demand-Side Management Programme, May 2010.

Y. Ma, A. Kelman, A. Daly, and F. Borrelli. Predictive control for energy efficient buildings with thermal storage: Modeling, stimulation, and experiments. *IEEE Control Systems*, 32(1):44 –64, February 2012.

NYISO. Day-ahead demand response program manual. Technical report, NYISO Auxiliary market operations, 2003.

F. Oldewurtel, A. Parisio, C.N. Jones, D. Gyalistras, M. Gwerder, V. Stauch, B. Lehmann, and M. Morari. Use of model predictive control and weather forecasts for energy efficient building climate control. *Energy and Buildings*, 45:1527, 2012.

F. Oldewurtel, T. Borsche, M. Bucher, P. Fortenbacher, M.G.V.T. Haring, J.L. Mathieu, O. Megel, E. Vrettos, and G. Andersson. A framework for and assessment of demand response and energy storage in power systems. In *Bulk Power System Dynamics and Control - IX Optimization, Security and Control of the Emerging Power Grid (IREP), 2013 IREP Symposium*, pages 1–24, 2013.

J. Page, S. Kiliccote, J.H. Dudley, M.A. Piette, A.K. Chiu, B. Kellow, E. Koch, and P. Lipkin. Automated demand response technology demonstration project for small and medium commercial buildings. Technical report, Lawrence Berkeley National Laboratory, July 2011.