Energy and Buildings 125 (2016) 196–204

Contents lists available at ScienceDirect

Energy and Buildings

journal homepage: www.elsevier.com/locate/enbuild

Demand response potential of model predictive control of space heating based on price and carbon dioxide intensity signals

Michael Dahl Knudsen*, Steffen Petersen

Department of Engineering—Indoor Climate and Energy, Aarhus University, Inge Lehmanns Gade 10, 8000 Aarhus C, Denmark

ARTICLE INFO

Article history:
Received 3 February 2016
Received in revised form 19 April 2016
Accepted 22 April 2016
Available online 30 April 2016

Keywords:
Demand response
Model predictive control
Space heating
Carbon dioxide emissions

ABSTRACT

This paper reports on a simulation-based study that investigated the demand response potential of a model predictive controller (MPC) for space heating defined to minimize a weighted sum of electricity costs and CO_2 emissions. The performance of the MPC was compared to a traditional controller and the results showed that an MPC with no weight on CO_2 emissions reduced the total electricity costs, shifted consumption from high to low load periods and reduced consumption in the hour with the yearly maximum grid load; but it could also cause an increase in CO_2 emissions. Contrary, the MPC with no weight on electricity costs reduced CO_2 emissions; but it only reduced total costs marginally, it could cause a shift of consumption from low to high load periods and it increased consumption in the hour with the yearly maximum grid load. Finally, if the MPC used a weighted sum of electricity costs and CO_2 emissions a range of intermediate results were obtained. The weighting factor can thus be used either to balance the performance of the MPC with respect to all performance indicators or to maximise it with respect to one indicator of particular interest.

1. Introduction

A characterizing feature of the electricity grid is that there always has to be an instantaneous balance between demand and supply. This is currently ensured by adjusting electricity production and network facilities to meet demands. However, sheer supply-side management of the balance becomes an increasing challenge as more renewable energy sources (RES) are introduced into the system because the production from the RES is intermittent and uncontrollable by nature. There is, therefore, a growing interest in exploring the potential of Demand Responses (DR) [1–7]. DR refers to actions on the demand side in response to certain conditions in the electric grid [7] and can be used to serve a variety of purposes. DR can lower peak demands and hence reduce the need for expensive peak load capacity and network facilities or it can serve as a means to flatten the demand profile to obtain efficient operational conditions for the generators [3,4,6,7]. DR can also be used to shift demands from one period to another, i.e. to achieve cost savings or to increase the utilization of supply from RES.

Households account for 26% of the total gross energy consumption in the European Union [8] whereof 50% is used to operate heating, ventilation and air conditioning (HVAC) systems for space conditioning [9]. Energy consumption for household HVAC thus constitutes a large DR potential if it can be made flexible e.g. through the concept of model predictive control (MPC) [10–15]. Nevertheless, household electricity loads are seldom used for DR, mainly because the DR potential is distributed over a large number of small loads. This requires investment in a large number of meters, communication and control devices, and a management system that coordinates and aggregates the small individual loads into a DR resource of significant size to electric system planners and operators (PO). However, this is likely to change as equipment costs decrease and the PO introduce demand response programs that facilitates participation of small household loads [5].

1.1. Demand response programs

DR can be established through different types of DR programs and are often divided into direct and indirect control programs [1,4,5]. In direct control programs the consumer give the PO direct control of electrical loads while in indirect control programs the consumer retains full control of the electrical loads and the PO can only try to change the consumption pattern indirectly. An example of indirect control is price-based DR programs where the flexibility demand from the PO is reflected in a data signal containing future time-varying electricity prices. This is an attempt to motivate con-

* Corresponding author.
E-mail addresses: mdk@eng.au.dk (M. Dahl Knudsen), stp@eng.au.dk (S. Petersen).

http://dx.doi.org/10.1016/j.enbuild.2016.04.053
0378-7788/

sumers to perform DR by enabling them to save money by reducing consumption in high price periods or by shifting consumption from high to low price periods.

It is common to subdivide price-based DR programs according to their pricing structure which is either static, dynamic or a mixture of both. One example of a static pricing structure is Time-Of-Use (TOU) tariffs which divide a day or a week into periods with different average electricity prices [1,2,16]. The objective of such programs is to flatten out consumption and reduce peaks [1]. DR programs based on TOU tariffs are relatively simple as there is no need for an ongoing communication of a price signal to the consumers. However, TOU tariffs are not guaranteed to always have an appropriate effect on consumption because an average electricity price does not reflect the actual state and needs of the electric system in periods that diverge significantly from the average. Dynamic pricing, on the other hand, has no pre-fixed price for specific time blocks. Instead they vary dynamically over short time periods, e.g. minutes or hours, and can, therefore, better indicate the actual state of the grid. An example of a dynamic pricing structure is day-ahead Real-Time Prices (RTP) where the wholesale prices are announced one day ahead in time [1]. Price-based DR programs can also be a combination of static and dynamic price structures and one example of this is Critical Peak Pricing (CPP) which is based on static TOU tariffs but dynamically adds an additional peak price during critical events or very high wholesale prices [1,17].

A limitation of a dynamic price signal based solely on the spot market price is that it does not necessarily reflect the environmental impact of the power production in a given time block [18]. An additional data signal containing the future time-varying CO_2 intensity of the available electricity could be used as a supplement to the price signal to take environmental impact into account.

1.2. Related work and aim of this paper

Several studies make use of price-based DR in relation to space conditioning, e.g. [12–15], but there are, to the knowledge of the authors, no reported studies involving a combination of price and CO_2 signals. The few studies found that investigate the use and potential of applying a combination of price and CO_2 signals in indirect DR programs are related to electricity consumption in household appliances. Tsagarakis et al. [19] solve a bi-objective optimization problem where the objectives are to minimize daily electricity costs and CO_2 emissions associated with residential wet load appliances (dishwasher, washing machine and dryer). They use day-ahead RTP data from the United Kingdom and find that there is a trade-off between costs and CO_2-emissions: When a relatively high weight is assigned to one of the objectives, it will reduce the saving on the other. A similar study by Paridari et al. [20] applies mixed integer linear programming to optimize the daily scheduling of wet loads and charging/discharging of a battery. Their objectives are to minimize electricity costs and CO_2 emissions using Swedish data for CO_2 intensity and RTP from Nordpool [21]. They also find a trade-off effect between electricity costs and CO_2 emissions, but one that is more significant than in Tsagarakis et al. [19] as their optimized controller will perform worse on one of the objectives compared to a non-optimized baseline when a sufficiently large weight is put on the other. This means that a controller that seeks solely to minimize costs can lead to a performance that increase the total CO_2 emissions compared to a traditional control strategy.

However, the indicated existence of a trade-off effect between electricity costs and CO_2 emissions in DR programs for electricity consumption in household appliances cannot immediately be transferred to control of space heating. The reason is that scheduling of appliances is optimized using a prediction horizon of a single day whereas MPC of space heating probably needs a longer prediction horizon due to the time constant of the thermal capacity of the building constructions. The aim of this paper is, therefore, to report on a simulation-based study that investigated the performance of an MPC for space heating when applying a combination of RTP and CO_2 signals. Various performance indicators were used. It was investigated whether there was any trade-off effect between electricity costs and CO_2 emissions using the MPC with different prediction horizons compared to more traditional controllers. The RTP and CO_2 signals used in the study were from the Danish electricity grid which has a large share of wind power. It was therefore also investigated whether an MPC using a combination of RTP and CO_2 signals would shift consumption towards periods with a high share of wind power production. Finally, the ability of the MPC using RTP and CO_2 signals to flatten out load profiles, cutting off peaks, and induce a lower usage of non-RES was compared to more traditional space heating control systems.

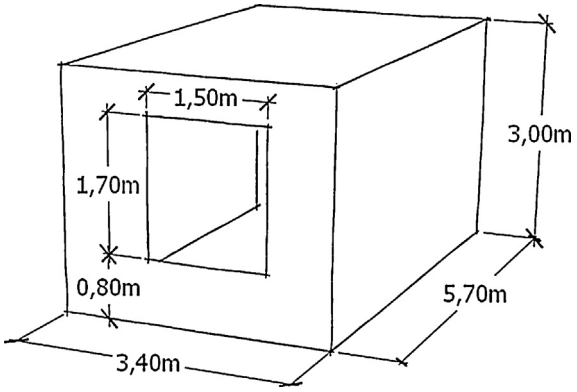

Fig. 1. Geometry of the EnergyPlus model (internal dimensions).

2. Method

The study reported in this paper is based on results from co-simulations of a test case modeled in EnergyPlus (EP) [22] and a heating system controlled by an MPC defined in MATLAB [23]. The EP model and the MPC are connected and exchange data in run-time via the Building Controls Virtual Test Bed [24,25]. The simulations are carried out for the period January 1 to February 14 in 2013 and 2014, respectively, which is within what is considered to be the general heating season in Denmark. The following sections explain the applied models in more details.

2.1. EnergyPlus model

The test case is featuring a dormitory apartment in the Grundfos Dormitory located in Aarhus, Denmark and modeled in EP. Fig. 1 shows the geometry of the apartment.

The model consists of a single thermal zone. All surfaces are internal and assumed to be adiabatic except for the south-facing wall which also has a low-energy window with no recess. The mechanical ventilation is always active with a constant airflow rate of 1.1 h^{-1} and a heat recovery of 75%. The infiltration rate is set to 0.05 h^{-1}. The heat source is assumed to be an electric baseboard with no thermal capacity. There are included no other internal heat loads to make the results easier to interpret. The Conduction Finite Difference algorithm is used to calculate the construction heat balances which induce the need for a small time step; we therefore apply 60 s time steps. The applied weather data is the EP weather file for Copenhagen [26]. Table 1 provides the details regarding the building constructions.

Table 1
Data for constructions used in the EnergyPlus model.

Construction type	Material	Properties
Window	Low-e coated clear double glazing w. argon in cavity4–14(Ar)-4(LowE)	U = 1.1 W/(m^2 K), g = 0.63
	Frame: Vinyl (width = 0.06 m)	U = 1.7 W/(m^2 K)
External wall	0.100 m concrete (exterior)	λ = 1.1 W/(m K), c = 736 kJ/(m^3 K)
	0.250 m insulation	λ = 0.037 W/(m K), c = 52 kJ/(m^3 K)
	0,200 m concrete (interior)	λ = 1.1 W/(m K), c = 736 kJ/(m^3 K)
Internal walls	0.180 m concrete	λ = 1.1 W/(m K), c = 736 kJ/(m^3 K)
Ceiling/Floor	0.025 m wood floor	λ = 0.15 W/(m K), c = 991 kJ/(m^3 K)
	0.050 m air space	R = 0.095 (m^2 K)/W
	0.220 m concrete	λ = 1.1 W/(m K), c = 736 kJ/(m^3 K)

2.2. Control model

The MPC uses a state space model (SSM) of the EP model to predict future output and to find the optimal sequence of control inputs. It is assumed that a linear time-invariant SSM can approximate the EP model since the EP room model has a constant air change. The advantage of this assumption is that the identification of the SSM and the MCP optimization becomes relatively simple to perform. This study applies the following discrete time linear SSM in the controller:

$$\boldsymbol{x_{k+1}} = \boldsymbol{Ax_k} + \boldsymbol{B}u_k + \boldsymbol{Ed_k} \quad \text{(1) State equation}$$

$$y_k = \boldsymbol{Cx_k} \quad \text{(2) Output equation}$$

where $\boldsymbol{x_k}$ is a 5-by-1 vector with the state variables in time step k, and $\boldsymbol{x_{k+1}}$ contains the state variables in the following time step $(k+1)$. u_k is a scalar with the controllable heat input from the electric heater, $\boldsymbol{d_k}$ is a 2-by-1 vector with the uncontrollable disturbances (outdoor air temperature and solar irradiation) while y_k is a scalar with the room temperature in time step k. A is a 5-by-5 state transition matrix, B is a 5-by-1 input matrix, E is a 5-by-2 disturbance matrix, and C is a 1-by-5 output matrix. The time step T_s between k and $k+1$ is one hour.

A pseudo random binary sequence [28] excited the EP model and the resulting room temperatures were registered and used for identification of the dimensions and parameters of A,B,E and C using the N4SID method in the MATLAB System Identification Toolbox [29]. In this specific case, a fifth order model turned out to be suitable to capture the dynamics of the EP model. This is a so-called black box identification process and it is therefore not possible to give a direct physical interpretation of the states $\boldsymbol{x_k}$ or the matrix parameters.

The MPC solves a minimization problem in each time step k to find the optimal sequence of control inputs for a finite prediction horizon of N time steps but it only executes the first input in the EP model. In the next time step $k+1$, the optimization problem is solved again but with the prediction horizon shifted one time step forward. This is called a receding horizon procedure [27,30] and it enables the controller to implement feedback from measurements of e.g. room temperature and apply continuously updated forecasts on weather, price and CO_2 signals. Feedback regarding room temperature in the EP model to the MPC goes into a Kalman filter state observer to correct the states from prediction errors in each time step. The optimization problem solved can be described as a deterministic linear program:

$$\min_u J = \sum_{k=1}^{N} s_k u_k \quad \text{(3) Cost function}$$

$$\text{Subject to} \begin{matrix} \bar{x}_{k+1} = A\bar{x}_k + Bu_k + E\bar{d}_k \\ y_k = C\bar{x}_k \end{matrix} \quad \text{(3.a) System dynamics}$$

$$\bar{x}_0 = \hat{x}_{ini} \quad \text{(3.b) Initial condition}$$

$$0 \le u_k \le 500 \quad \text{(3.c) Control input}$$

$$y_k \ge y_{k,min} \quad \text{(3.d) Output}$$

The cost in time step k is the product of the heat input,u_k, and the combined price and CO_2 signal, s_k (see Section 2.3). The total cost J Eq. (3) is thus the sum of the cost in each time step for the entire prediction horizon of N time steps and the objective is to find the control sequence, u_{opt}, that minimize J. Any feasible solution is subject to four types of constraints Eq. (3-a-d). The system dynamic constraints Eq. (3-a) force the solution to comply with the system dynamics described by the SSM model in Eqs. (1)–(2). The initial condition constraint Eq. (3-b) sets the initial condition equal to the initial states estimated by the Kalman filter. The control constraint Eq. (3-c) ensures that the solution only applies physical sensible heat inputs which, in this study, lie in the interval [0 W, 500 W]. Finally, the output constraint Eq. (3-d) forces any feasible solution to produce an indoor temperature,y_k, that lies above a minimum set point temperature which, in this study, is 21 °C during daytime and 18 °C from 12:00 am to 7 am. We also apply slack variables as described by Halvgaard et al. [13] to ensure that a feasible solution is always available, but the above formulation does not include these for the sake of simplicity. In this study, the weather forecasts are assumed to be perfect but forecast uncertainties could be incorporated e.g. as a stochastic MPC formulation [31,32].

2.3. RTP and CO_2 intensity signals

The combined signal s used in this study is composed of two separate signals, c and e, where $c = [c_1, c_2, \ldots, c_N]^T$ is the prediction of RTP and $e = [e_1, e_2, \ldots, e_N]^T$ is the prediction of the CO_2 intensity associated with the electricity production. The problem is really a bi-objective optimization problem, but a simple and often used approach is to transform this into a single-objective problem by taking a weighted sum of each objective [19,20]. This study applies the following weighted sum:

$$s_k = (1-\lambda)c_k + \lambda e_k \quad (4)$$

The combined signal, s, is thus a weighted sum of the two original signals, c and e, and λ is a weighting constant. One must typically assign a λ-value between [0,1]where λ = 0 means that all weight is put on the RTP signal while λ = 1 means that all weight is put on the CO_2 intensity.

The Danish electricity market is divided into two price areas and the RTP data applied in this study are hourly day-ahead RTP from Western Denmark [21] and are made available by the Danish TSO, Energinet.dk [33]. The CO_2 intensity data are calculated by Energinet.dk based on the actual mix of electricity production in Denmark and corrected for net import or export to neighbouring countries [34]. Fig. 4 (bottom) depicts an example of RTP and CO_2 data for a week in the simulation period in 2013.

Table 2 shows the mean values and standard deviations for the two simulation periods. Notice that the mean values are of the same order of magnitude as the standard deviations, which makes it reasonable to apply c and e directly in Eq. (4) without any prior normalization.

Table 2
The mean value and standard deviation for RTP and CO_2 intensity.

	2013		2014	
	RTP [DKK/MWh]	CO_2 [kg/MWh]	RTP [DKK/MWh]	CO_2 [kg/MWh]
Mean	299	416	223	291
Standard deviation	88	86	74	68

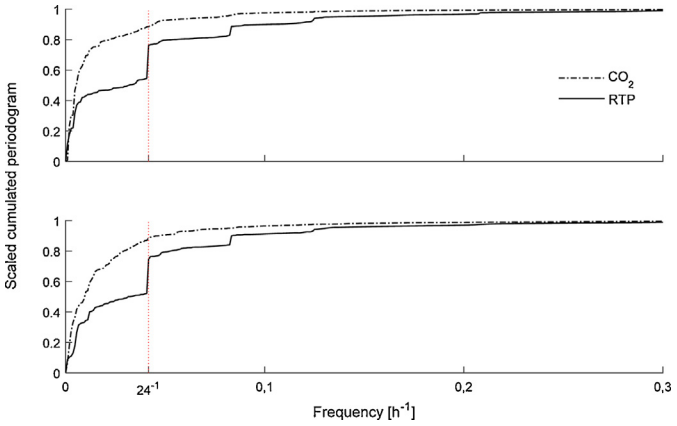

Fig. 2. The scaled cumulated periodograms of fluctuations in the RTP and CO_2 intensity time series with respect to mean values. Top: 2013, bottom: 2014.

However, there are differences in the frequency content of the RTP and CO_2 intensity time series, which is seen in Fig. 2 displaying the cumulated periodograms as a function of frequency measured in oscillations per hour. There is jump in the RTP plot at a frequency of $1/24\,h^{-1}$ revealing that the RTP has a relatively high content of diurnal oscillations. This is not the case for the CO_2 intensity. Fig. 2 also shows that the CO_2 intensity time series have a relatively high content of frequencies below $1/24\,h^{-1}$ compared to the RTP and that can help us understand some of the phenomena observed in the simulation results (see Section 3).

The bars in Fig. 3 indicate the standard winter load periods as defined by Energinet.dk [34,36] and include low, high and peak load periods. It will, generally, be advantageous to shift consumption from peak to low load periods and a reasonable performance criterion of a controller is its ability to perform these shifts. The line plots in Fig. 3 depict the hourly mean values in the simulation period for the total electricity consumption in Western Denmark (the grid load), RTP and CO_2 intensity, respectively. The values are normalized with respect to the yearly maximum values. Notice that the mean of RTP and the total consumption in the grid have very similar daily profiles and are in accordance with the standard load periods. In contrast, the mean CO_2 intensity has no diurnal pattern. The hour at which the maximum RTP, CO_2 intensity and total consumption occur is also indicated in Fig. 3. The maximum consumption in the electricity grid (circle) is located at 6 p.m. in January for both years and the maximum RTP (cross) is located at 6 p.m. and 7 p.m. in 2013 and 2014, respectively, and are thus in good agreement. In contrast, the maximum CO_2 intensity (square) is located at 4 a.m. and at midnight which are low load periods. These observations suggest that RTP is a good indicator of the load in the electricity grid while CO_2 intensity is not.

Table 3 shows the correlation coefficient r in the two simulation periods between the RTP and CO_2 intensity as well as the time series for the generation of wind power and the total electricity demand in Western Denmark [33]. There is a strong positive correlation between RTP and the total electricity demand in the grid suggesting that it is a good indicator of the grid load. On the other hand, there is only a moderate negative correlation between RTP and supply from wind turbines suggesting that it is a less good indicator of the amount of wind power in the grid. In contrast, the CO_2 intensity is not correlated with the total electricity demand, but there is a very strong negative correlation between CO_2 intensity and the generation of wind power. This suggests that the CO_2 intensity is a poor indicator of the grid load but a good indicator of the amount of wind power.

Table 3
Correlation coefficients between RTP, CO_2 intensity, wind power production and total electricity demand in the grid.

	2013				2014			
	RTP	CO_2	Wind	Demand	RTP	CO_2	Wind	Demand
RTP	1	0.41	−0.48	0.67	1	0.57	−0.34	0.62
CO_2 intensity		1	−0.79	−0.05		1	−0.81	0.13
Wind			1	0.04			1	0.09
Demand				1				1

Based on the above analysis we can expect that an MPC based on RTP alone will do a good job at shifting consumption away from periods with high loads on the grid, but it will not necessarily shift towards periods with a high production of wind power or reduce CO_2 emissions. In contrast, an MPC based solely on CO_2 intensity is expected to shift consumption towards periods with a high share of wind power, but it will not necessarily shift away from peak load periods or reduce costs.

2.4. Prediction horizon

The MPC determines the optimal sequence of control inputs for N future time steps, see Section 2.2 for details, and N is therefore the number of time steps in the prediction horizon. A proper choice of prediction horizon length depends on a number of conditions such as thermal characteristics of the building (i.e. the time constant) and characteristics of the RTP and CO_2 signals. In this study, we apply a prediction horizon of six days to ensure a full utilization of the thermal capacity of the building. However, the reliability of weather, RTP, and CO_2 forecasts for a prediction horizon of six days is currently questionable and often not even available. Outputs from simulations with a prediction horizon of six days can, therefore, be regarded as a theoretical performance bound, i.e. a benchmark for evaluating the performance of more realistic prediction horizon scenarios. We shall refer to this as the *ideal MPC*. A currently more realistic prediction horizon length is, according to Nordpool [21], approx. one day and simulations with a one-day prediction horizon are therefore also made. Furthermore, simulations are also executed with a prediction horizon of six days but where only the first day is assumed available as a perfect forecast and the remaining five days are obtained through the following coarse extrapolation:

$$s[k] = s[k-24] \tag{5}$$

where the unknown value $s[k]$ is set equal to the value 24 h earlier.

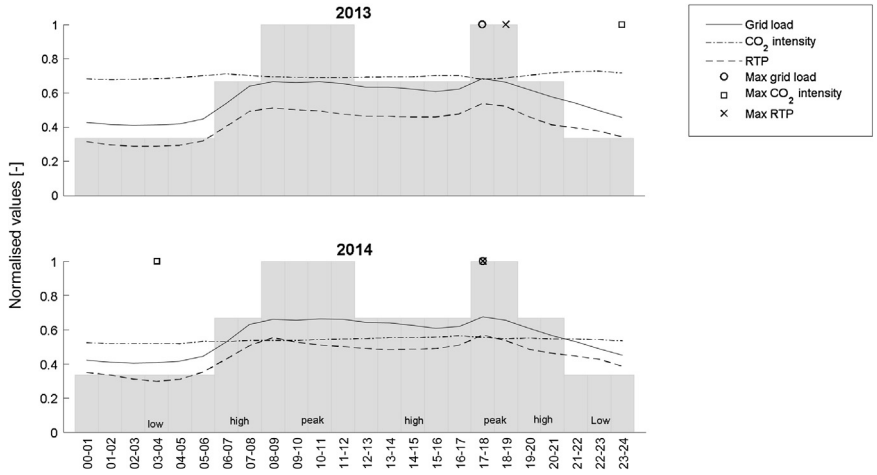

Fig. 3. Normalized hourly gross consumption, CO_2 intensity and RTP. Data is normalized with respect to maximum value in the periods. The light grey bars visualize the standard tariff periods; low, high and peak periods.

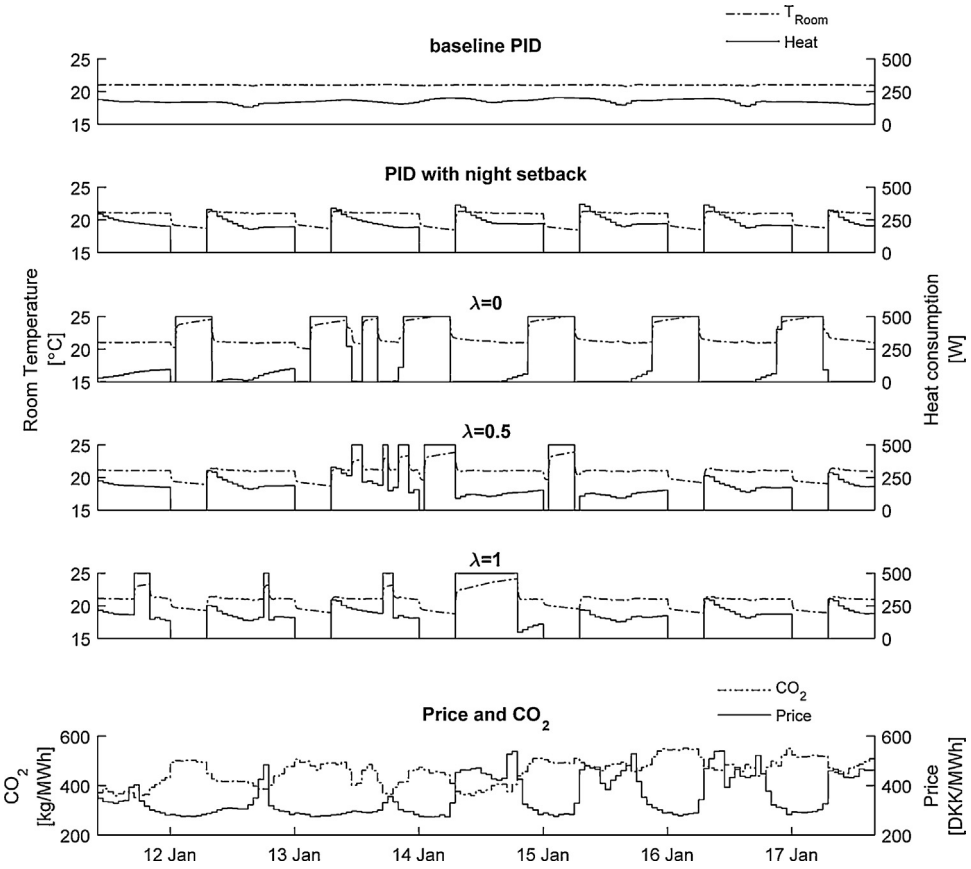

Fig. 4. One week of simulation results for the PID, the PID with night setback and three different ideal MPC control strategies. The bottom graph plots the RTP and CO_2 intensity for the given period (the tick marks on the x-axis are placed at 12.00 a.m.).

3. Results

The MPC simulation results were compared with two proportional-integral-derivative (PID) controllers [35]; one with a constant set point temperature of 21 °C and one with a 3 °C night setback from 12 am to 7 am. Notice that the minimum temperature constraints of the MPC Eq. (3-d) are the same as the temperature set points of the PID with night setback thus allowing the MPC to apply night setback if optimal.

Fig. 4 depicts one week of simulation results for the PID with constant set point, the PID with night setback and the ideal MPC for three different weighting constants (λ), respectively. The PID with a constant set point had a relatively steady heat consumption while the PID with night setback boosted heating in the morning to raise the night temperature to the daytime set point. The MPC with no weight on the CO_2 intensity ($\lambda = 0$) minimized cost by boosting heat consumption in periods with a relatively low RTP to use less in later periods with high RTP. In contrast, the MPC that solely minimized CO_2 emissions ($\lambda = 1$) boosted the heat consumption in

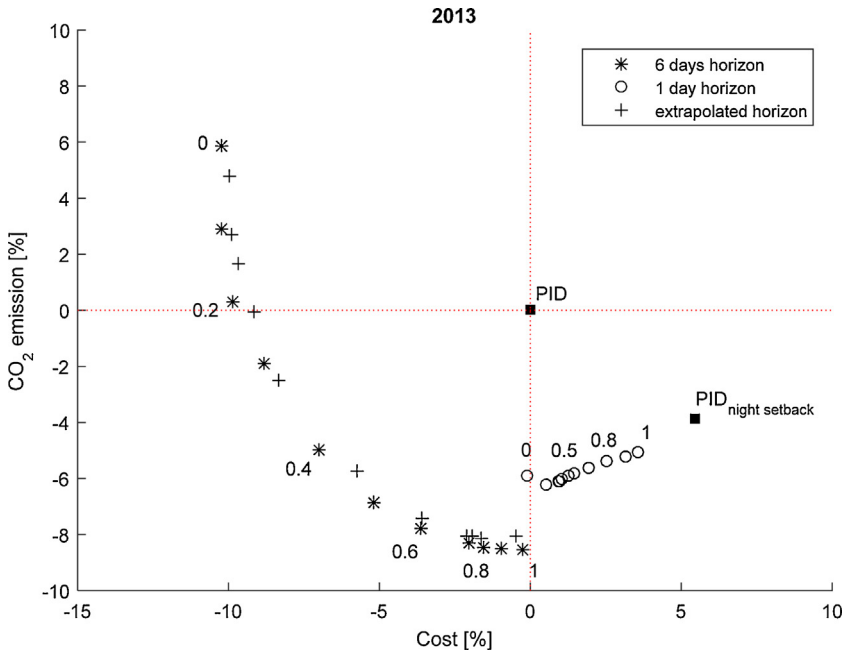

Fig. 5. Summary of the cost-emission trade-off for the different controllers in 2013. The number labels in the graph indicate the λ-value for the associated result.

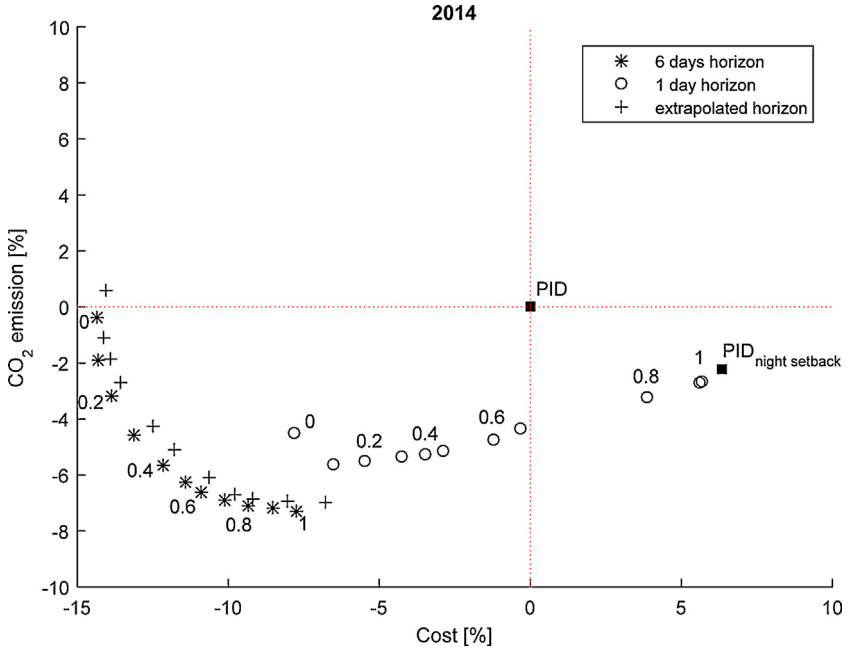

Fig. 6. Summary of the cost-emission trade-off for the different controllers in 2014. The number labels in the graph indicate the λ-value for the associated result.

periods with a low CO_2 intensity to use less in periods with a high CO_2 intensity. The MPC that put equal weight on costs and CO_2 emissions ($\lambda = 0.5$) boosted heating when both RTP and CO_2 intensity were relatively low. Fig. 4 also shows that the economic MPC ($\lambda = 0$) often did not reduce the night temperatures even though this would reduce the total heat consumption. The reason is that the nighttime RTP is often lower than the daytime RTP and the MPC therefore often found it more cost-effective to boost temperature during night time to save expensive electricity during daytime. On the other hand, when CO_2 emissions were given a high weight ($\lambda = 1$), night setback was often applied because CO_2 emissions did not have the same strong diurnal pattern as the RTP signal (see Figs. 2 and 3) and are, in general, not low at night and high during daytime.

Figs. 5 and 6 depict the performance of the different controllers with respect to the accumulated electricity costs and CO_2 emissions for the simulated periods in 2013 and 2014, respectively. The PID controller with a constant set point temperature was used as baseline and is therefore placed at the origin while the performance of the other controllers is plotted relative to this baseline. The controllers within the third quadrant dominated the baseline PID with respect to both objectives. The PID controller with night setback is in the fourth quadrant because it had lower accumulated CO_2 emissions than the baseline PID but higher accumulated costs. None of the two PID controllers can therefore be said to dominate the other.

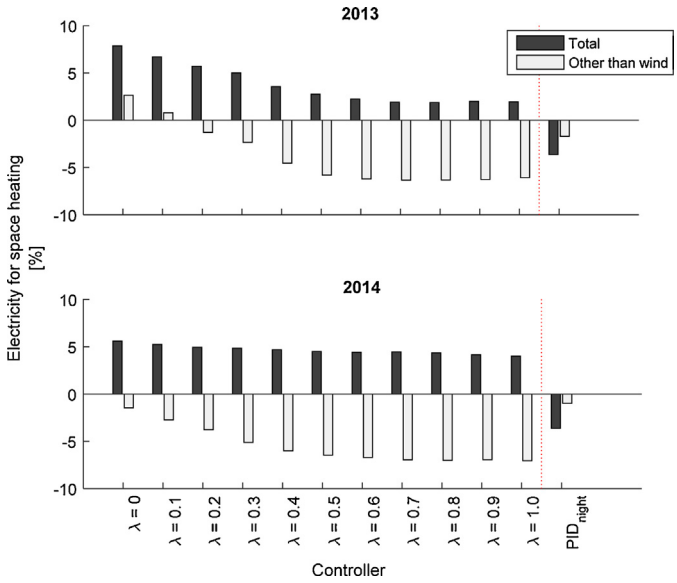

Fig. 7. Electricity consumption for different controllers. The dark bars are the total electricity consumption for space heating compared to the baseline PID and the light gray bars are the amount of the total electricity consumption covered by other generators than wind turbines.

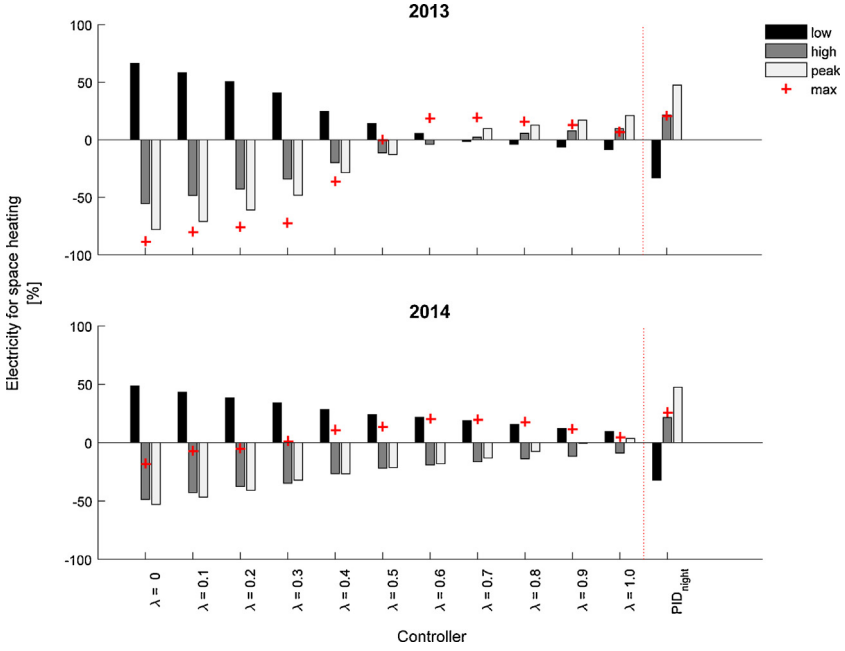

Fig. 8. Change in electricity consumption for space heating for different controllers compared to the baseline PID with constant set point. The black bars are changes in electricity consumption in the low load periods compared to the baseline, the grey bars are changes in the high load periods and the light grey bars are the changes in the peak periods. The plus signs show the changes in the maximum peak hour.

In 2013 there is a more profound trade-off effect compared to 2014 for the ideal MPC (six days of perfect forecasts). The reason is that there was a higher correlation between RTP and CO_2 intensity in 2014 than in 2013 (Table 3). In 2014 the ideal MPC dominated the baseline regardless of the chosen λ, while the ideal MPC in 2013 had higher emissions than the baseline when only a small weight was put on CO_2 ($\lambda \leq 0.2$).

For the MPC with a prediction horizon of just one day there was almost no trade-off to consider because, generally, the performance was improved for both objectives with decreasing weight on CO_2 emissions (except for $\lambda \approx 0$). The reason is that the CO_2 intensity varied at a slower rate compared to the RTP signal (Fig. 2) and there was, consequently, often only a small difference between the minimum and maximum values during a prediction horizon of one day. A high weight on CO_2 emissions would therefore not induce large shifts in heat consumption. In contrast, there was often a significant RTP difference within a day, and a high weight on RTP would, consequently, make the MPC shift significant amounts of heat consumption from high to low RTP periods. These shifts often also resulted in reduced CO_2 emissions because of the moderate positive correlation between RTP and CO_2 intensity (Table 3).

Reducing the horizon length from six to one day significantly decreased the obtainable savings and moved the trade-off front away from the third quadrant. However, the MPC that only applied forecasts for a single day but extrapolates the remaining five days

was able to approximate the theoretical performance bound of the six days prediction horizon of perfect forecasts.

Fig. 7 depicts the total electricity for space heating (dark bars) consumed by the different controllers and the total amount covered by other generators than wind turbines (light gray bars). The values are accumulated over the simulation periods in 2013 and 2014, respectively, and are relative to the baseline PID. The MPC generally used more energy than the baseline PID even though it was allowed to reduce the temperature at night. However, the MPC controllers also consumed less electricity from non-wind generators in all cases except for $\lambda \leq 0.1$ in 2013. This means that the MPC effectively invested wind power to shift consumption away from periods with a low share of wind power production. Both the extra total consumption and the usage of other sources than wind power decreased significantly when more weight was put on CO_2 emissions (λ increase). Notice that the PID controller with night setback was the only strategy that used less energy than the baseline PID but it achieved only a small reduction in the usage of non-wind produced electricity.

Fig. 8 shows how the controllers shifted consumption with respect to the three standard load periods (Fig. 3) compared to the baseline PID. The red plus signs indicate how the controllers changed consumption in the hour with the yearly maximum grid load (critical peak). When high weight was put on RTP ($\lambda \approx 0$), the MPC managed to shift consumption from the high/peak load periods to the low load periods and significantly reduce consumption in the critical peak hour. As more weight was put on CO_2 emissions, the shift from high/peak load periods to low load periods was reduced. In 2013, high weight on CO_2 intensity ($\lambda > 0.5$) resulted in a total shift from low to high/peak load periods and an increased consumption in the critical peak hour. The reason is that the MPC with high weight on CO_2 intensity often applied night setback and hence shifted consumption from night to daytime. This was not so much the case in 2014 because of the relatively high correlation between RTP and CO_2 intensity. Notice that the PID with night setback in general shifted consumption from low to high/peak load periods and increased the consumption in the maximum hour.

4. Conclusion

This paper presents results from a simulation-based study that investigated the performance of an MPC that applied different combinations of RTP and CO_2 intensity signals for space heating control compared to conventional PID control. The following performance indicators were used: Costs, CO_2 emissions, ability to shift consumption from high to low load periods, ability to reduce consumption in the critical peak hour and, finally, the ability to reduce the consumption of non-wind generated electricity. Simulations were performed using electricity system data for a specific period (1 January–14 February) in the Danish heating season in 2013 and 2014, respectively.

The results indicate that an economic MPC ($\lambda = 0$) will effectively reduce the total costs, shift consumption from high to low load periods, and reduce the hourly maximum consumption. However, it may cause an increase in the usage of non-wind generated electricity and a total increase in CO_2 emissions because price and CO_2 intensity were not very strongly correlated which is otherwise often assumed. In contrast, a CO_2-minimizing MPC ($\lambda = 1$) will effectively reduce the usage of non-wind generated electricity and hence reduce CO_2 emissions. However, it will only reduce the total costs marginally, potentially cause a shift in consumption from low to high load periods, and increase consumption in the critical peak hour. If the MPC uses a weighted sum of RTP and CO_2 intensity ($0 < \lambda < 1$), a range of intermediate results can be obtained with diverging performance with respect to the different performance indicators. The performance of the MPC is thus sensitive to the choice of λ which can be carefully chosen to either dominate PID in terms of all performance indicators or to make the MPC maximise performance with respect to certain performance indicators. The results indicate that it may be a good idea to put at least a small weight on CO_2 intensity ($\lambda \geq 0.2$) because this will give a relatively large reduction in CO_2 emissions compared to a pure economic MPC and only lead to a very small increase in cost (see Figs. 5 and 6). Likewise, a small weight on RTP ($\lambda \leq 0.8$) can give large cost reductions compared to a pure CO_2 minimizing MPC and leads to only a very small increase in CO_2 emissions.

The simulation results also showed that the potential of the MPC is improved by applying a prediction horizon longer than a single day. To do this the MPC needs forecasts beyond what is typically available, but the simulation results also showed that a simple extrapolation of daily forecasts can be used without a noticeable reduction of performance potential.

Future work is to investigate if other signals could supplement or replace the RTP and CO_2 intensity signals to further improve the MPC performance. For instance, a signal with information about critical peak periods may help an MPC with high λ-values to reduce consumption in the maximum hour. Reproduction of the analysis for periods later than 2014–and maybe even for future scenarios—may also be of interest because the electricity system keeps evolving, e.g. more wind-power is constantly introduced. The analysis could also be repeated for energy systems with a different composition of RES as conclusions are likely to be different for systems with e.g. a large share of solar power. Furthermore, the study reported in this paper should be considered as a potential study featuring a very simple case building using perfect weather forecasts and no upper limit for thermal comfort. Future work could be to investigate the MPC performance in more realistic case buildings, including the use of uncertain weather forecasts and occupancy data as well as an upper limit for thermal comfort. Finally, it would be interesting to investigate how different features of the building envelope, e.g. insulation level, would impact the performance of the MPC.

Acknowledgements

This paper is based upon work supported by The Danish Energy Agency project (no. 12019)—Virtual Power Plant for Smart Grid Ready Buildings (VPP4SGR) and Aarhus University, Denmark.

References

[1] M.H. Albadi, E.F. El-Saadany, A summary of demand response in electricity markets, Electron. Power Syst. Res. 78 (November (11)) (2008) 1989–1996, http://dx.doi.org/10.1016/j.epsr.2008.04.002.
[2] P. Khajavi, H. Abniki, A.B. Arani, The role of incentive based demand response programs in smart grid, Environ. Electr. Eng. 10 (2011) 1–4, http://dx.doi.org/10.1109/EEEIC.2011.5874702.
[3] M. Muratori, B. Schuelke-Leech, G. Rizzoni, Role of residential demand response in modern electricity markets, Renew. Sustain. Energy Rev. 33 (2014) 546–553, http://dx.doi.org/10.1016/jrser.2014.02.027.
[4] N. O'Connell, P. Pinson, H. Madsen, M. O'Malley, Benefits and challenges of electrical demand response: a critical review, Renew. Sustain. Energy Rev. 39 (November) (2014) 686–699, http://dx.doi.org/10.1016/j.rser.2014.07.098.
[5] P. Siano, Demand response and smart grids—a survey, Renew. Sustain. Energy Rev. 30 (2014) 461–478, http://dx.doi.org/10.1016/j.rser.2013.10.022.
[6] G. Strbac, Demand side management: benefits and challenges, Energy Policy 36 (December (12)) (2008) 4419–4426, http://dx.doi.org/10.1016/j.enpol.2008.09.030.
[7] J. Torriti, M.G. Hassan, M. Leach, Demand response experience in Europe: policies, programmes and implementation, Energy 35 (4) (2010) 1575–1583, http://dx.doi.org/10.1016/j.energy.2009.05.021.
[8] European Commission, EU energy in figures—statistical pocketbook (2014) 10.2833/24150.
[9] L. Pérez-Lombard, J. Ortiz, C. Pout, A review on buildings energy consumption information, Energy Build. 40 (3) (2008) 394–398, http://dx.doi.org/10.1016/j.enbuild.2007.03.007.

[10] A. Afram, F. Janabi-Sharifi, Theory and applications of HVAC control systems—a review of model predictive control (MPC), Build. Environ. 72 (2014) 343–355, http://dx.doi.org/10.1016/j.buildenv.2013.11.016.
[11] G.S. Pavlak, G.P. Henze, V.J. Cushing, Optimizing commercial building participation in energy and ancillary service markets, Energy Build. 81 (2014) 115–126, http://dx.doi.org/10.1016/j.enbuild.2014.05.048.
[12] M. Avci, M. Erkoc, A. Rahmani, S. Asfour, Model predictive HVAC load control in buildings using real-time electricity pricing, Energy Build. 60 (2013) 199–209, http://dx.doi.org/10.1016/j.enbuild.2013.01.008.
[13] R. Halvgaard, N.K. Poulsen, H. Madsen, J.B. Jørgensen, Economic model predictive control for building climate control in a smart grid, Innovative Smart Grid Technol. (2012) 1–6, http://dx.doi.org/10.1109/ISGT.2012.6175631.
[14] J. Ma, J. Qin, T. Salsbury, P. Xu, Demand reduction in building energy systems based on economic model predictive control, Chem. Eng. Sci. 67 (1) (2011) 92–100, http://dx.doi.org/10.1016/j.ces.2011.07052.
[15] E. Vrettos, K. Lai, F. Oldewurtel, G. Andersson, Predictive control of buildings for demand response with dynamic day-ahead and real-time prices, Eur. Control Conf. (July) (2013) 2527–2534.
[16] K. Train, G. Mehrez, The impacts of optional time-of-use prices: a case study, Energy Build. 22 (1995) 267–278, http://dx.doi.org/10.1016/0378-7788(95)00928-Q.
[17] K. Herter, S. Wayland, Residential response to critical-peak pricing of electricity: california evidence, Energy 35 (2010) 1561–1567, http://dx.doi.org/10.1016/j.energy.2009.07.022.
[18] P. Stoll, N. Brandt, L. Nordstrom, Including dynamic CO_2 intensity with demand response, Energy Policy 65 (2014) 490–500, http://dx.doi.org/10.1016/j.enpol.2013.10.044.
[19] G. Tsagarakis, R.C. Thomson, A.J. Collin, G.P. Harrison, A.E. Kiprakis, Assessment of the cost and environmental impact of demand side management on residential sector, Ecol. Veh. Renew. Energies 9 (2014) 1–7 (10.11097EVER.2014.6844154).
[20] K. Paridari, A. Parisio, H. Sandberg, K.H. Johansson, Energy and CO_2 efficient scheduling of smart appliances in active houses equipped with batteries, IEEE In. Conf. Autom. Sci. Eng. (August) (2014) 632–639, http://dx.doi.org/10.1109/CoASE.2014.6899394.
[21] Nord Pool, Day-Ahead Market Elspot, http://www.nordpoolspot.com/TAS/Day-ahead-market-Elspot (accessed 18.01.16.).
[22] US Department of Energy, EnergyPlus 8.1.0 Build 009 (December 2013).
[23] The MathWorks, Inc., MATLAB R2015a (8.5.0.197613) (February 2015).
[24] M. Wetter, Co-simulation of building energy and control systems with the building controls virtual test bed, J. Build. Perform. Simul. 4 (September) (2011) 185–203, http://dx.doi.org/10.1080/19401493.2010.518631.
[25] Lawrence Berkeley National Laboratory, Building Controls Virtual Test Bed (BCVTB) 1.5.0 (January 2015).
[26] U.S. Department of Energy, Energy Efficiency & Renewable Energy, EnergyPlus Weather Data -Copenhagen 061800 (IWEC) 2016 (accessed 17.01.16.). https://energyplus.net/weather.
[27] M. Gwerder, D. Gyalistras, C. Sagerschnig, R.S. Smith, D. Sturzenegger, Final Report: Use of Weather And Occupancy Forecasts For Optimal Building Climate Control -Part II: Demonstration (OptiControl-II) (2013).
[28] P. Bacher, H. Madsen, Identifying suitable models for the heat dynamics of buildings, Energy Build. 7 (2011) 1511–1522, http://dx.doi.org/10.1016/j.enbuil.2011.02.005.
[29] The Mathworks, Inc., MATLAB R2015b, System Identification Toolbox 9.2.
[30] F. Oldewurtel, Stochastic Model Predictive Control for Energy Efficient Building Climate Control (2011), Dissertation at ETH Zürich.
[31] F. Oldewurtel, C.N. Jones, A. Parisio, M. Morari, Stochastic model predictive control for building climate control, IEEE Trans. Control Syst. Technol. 22 (2014) 1198–1205, http://dx.doi.org/10.1109/TCST.2013.2272178.
[32] X. Zhang, G. Schildbach, D. Sturzenegger, M. Morari, Scenario-based MPC for energy-efficient building climate control under weather and occupancy uncertainty, Eur. Control Conf. (2013) 1029–1034.
[33] Energinet.dk Introduction to download of market data http://www.energinet.dk/EN/El/Engrosmarked/Udtraek-af-markedsdata/Sider/Om-markedsdata.aspx(accessed 18.01.16.).
[34] Energinet.dk CO_2 prognoser http://www.energinet.dk/DA/El/Engrosmarked/Udtraek-af-markedsdata/Sider/CO2-prognoser.aspx (accesed 18.01.16.).
[35] Association of Danish Energy Companies, Electric Utility tariffs and electricity prices—2014 (Danish title: Elforsyningens tariffer & elpriser—2014) (2015).
[36] K.H. Ang, G. Chong, L. Yun, PID control system analysis, design, and technology, IEEE Trans. Control Syst. Technol. 13 (4) (2005) 559–576, http://dx.doi.org/10.1109/tcst.2005.847331.