Applied Energy 298 (2021) 117227

Contents lists available at ScienceDirect

Applied Energy

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

Experimental test of a black-box economic model predictive control for residential space heating

Michael Dahl Knudsen [a,*], Laurent Georges [b], Kristian Stenerud Skeie [c], Steffen Petersen [a]

[a] Department of Civil and Architectural Engineering, Aarhus University, Inge Lehmanns Gade 10, 8000 Aarhus C, Denmark
[b] Department of Energy and Process Engineering, Faculty of Engineering, NTNU - Norwegian University of Science and Technology, Kolbjørn Hejes vei 1a, 7034 Trondheim, Norway
[c] Department of Architecture and Technology, Faculty of Architecture and Design, NTNU - Norwegian University of Science and Technology, Alfred Getz vei 3, 7491 Trondheim, Norway

HIGHLIGHTS

- An economic model predictive control was tested for residential space heating.
- Good predictive performance of the applied black-box state-space model.
- A single thermal zone model provided acceptable temperature control.
- The E-MPC successfully shifted consumption from high to low cost periods.

ARTICLE INFO

Keywords:
Economic model predictive control
Black-box model
State-space model
Residential space heating
Price-based demand response

ABSTRACT

Previous studies have identified significant demand response (DR) potentials in using economic model predictive control (E-MPC) of space heating to exploit the inherent thermal mass in residential buildings for short-term energy storage. However, the economically viable realisation of E-MPC in residential buildings requires an effort to minimise the need for additional equipment and labour-intensive modelling processes. This paper reports on an experiment where a novel E-MPC setup was used for thermostatically control of a hydronic radiator in a highly-insulated residential building located on the NTNU Campus in Trondheim, Norway. The E-MPC utilized data from a heating meter, two temperature sensors and an existing weather forecast web service to train a linear black-box model. The results showed that the precision of model trained on excitation data that was generated using setpoints of either 21 or 24 °C was sufficient to obtain good control of the indoor air temperature while shifting consumption from high to low price periods. The findings of the experiment indicate that a minimal E-MPC setup is able to realize the significant DR potential that lies in utilizing the inherent thermal mass in residential buildings.

1. Introduction

The penetration of weather-dependent renewable electricity production such as wind and solar power technologies complicates the task of continuously balancing supply and demand in the grid, for example, due to the inherent uncertainties of weather forecasting. Demand response (DR) is often mentioned as a solution for grid balancing in systems with a high penetration of volatile electricity production [1–3]. There are various terminologies and definitions of DR strategies in current literature but common for DR programs is that they seek to make temporary adjustments in consumption at times when this has value for the energy system as a whole. IEA EBC Annex 67 defines building energy flexibility as the ability of a building to manage its demand and generation according to local climate conditions, user needs and grid requirements [4].

The building fabric can be used as thermal short-term heat storage to adjust the space-heating load in time and has proved to have large energy flexibility potential [5,6]. In this respect, previous studies have identified significant DR potential in using economic model predictive control (E-MPC) of space heating to exploit the thermal mass of

* Corresponding author.
E-mail address: mdk@cae.au.dk (M.D. Knudsen).

https://doi.org/10.1016/j.apenergy.2021.117227
Received 15 February 2021; Received in revised form 29 May 2021; Accepted 30 May 2021
Available online 9 June 2021

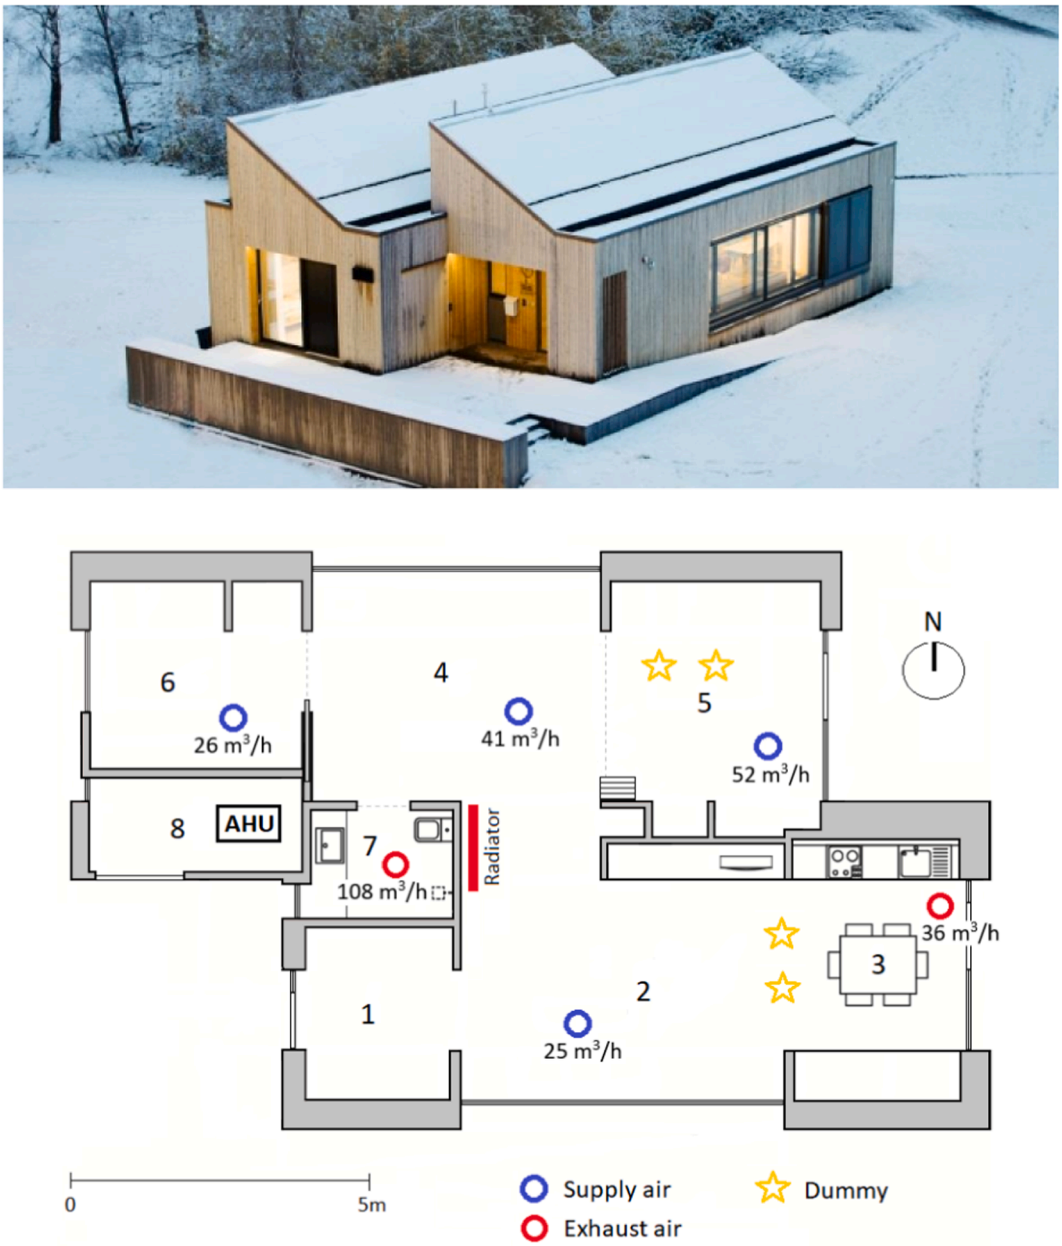

Fig. 1. (Top) ZEB Living Lab on the Gløshaugen campus, NTNU. (Bottom) Floor plan of Living Lab (1: Entrance; 2: Living room south; 3: Kitchen; 4: Living room north; 5: Bedroom east; 6: Bedroom west; 7: Bathroom; 8: Technical room) [39].

residential buildings. The vast majority of these studies were simulation-based, e.g. [7–10]. Some simulation-based studies have also combined MPC of space heating with control of other building energy systems [11,12], minimisation of CO_2 emissions [13], and investigations of energy tax structures [14]. Some studies report on field investigations of E-MPC in residential buildings [15–17], but physical implementation and tests of E-MPC have mainly been reported for applications in laboratory test cells [18,19], office buildings [20–22], or educational buildings [23–25]. This limited set of field investigations is also confirmed in the review paper by Kathirgamanathan et al. [26] on predictive control for building energy flexibility.

In 2018, residential buildings represented 26.1% of the total energy consumption in the EU and most of this was used for space heating (63.6%) [27]. The aggregated DR potential in residential buildings is thus very large and simulation-based studies have demonstrated that realising the potential can be a valuable asset to operational challenges in e.g. urban district heating systems [28–31]. However, the DR potential of individual residential buildings is small in absolute terms and it is therefore crucial that the cost of implementing E-MPC is low to be economically justified. The current challenges in this relation are:

- According to Killian and Kozek [32], existing home automation systems and advances in wireless technology have provided practical and flexible means for collecting data needed for MPC but they are a significant extra investment cost if not already present for other purposes. Realising E-MPC in residential buildings, therefore, requires an effort to minimise the additional equipment needed, e.g. as seen in [33] where the need for investing in equipment for local weather data measurements was eliminated.
- The development of a control model is known to be the most time-consuming and thus expensive part of implementing an MPC [20,34,35]. Efforts to make automated modelling procedures should therefore be a priority to keep the investment costs at a minimum. Modelling for MPC can be divided into three main categories [36]:

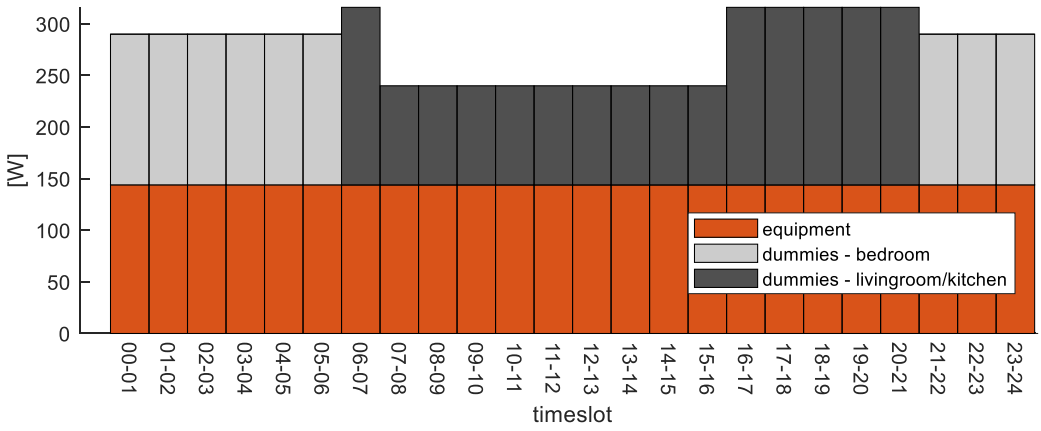

Fig. 2. The daily internal heat gain (in Watts) repeatedly distributed over the 24 h of the day. The heat gain consisted of gains from electrical equipment and cylindrical dummies. The dummies mimicked two adults that were either in the bedroom, the living-room/kitchen or away.

Table 1
User choices for the *n4sid* and *pem* algorithms.

Algorithm	Parameter	Choice
n4sid:	Focus	Prediction
	N4Weight	CVA
	N4Horizon	[20 20 20]
pem:	Focus	Prediction
	SearchMethod	Auto
	Tolerance	0.01
	MaxIterations	1000

white-, grey- and black-box models. White box models are physics-based models (i.e. first principle models) and require a large amount of input and expert knowledge to setup the model. Also, the building can change in time which may require frequent recalibration. Black-box models are pure data-driven methods based on measured input–output data. The quality of black-box models relies strongly on the quantity and quality of the data. The extrapolation properties of black-box models are lower than for white and grey-box models. Finally, grey-box models are intermediate between white- and black-box models: the model structure is defined by physical laws while model parameters are identified using measurement data. The model structure of grey-box models is often expressed as a network of thermal resistances and capacitances (i.e., RC models). A review of grey-box models applied to MPC of buildings can be found in e.g. Viot et al [37]. Both white- and grey-box models require physical knowledge of the building and may therefore be a challenge for realisation of low-cost models for MPC of residential buildings.

This paper reports on an experiment where a hydronic radiator in a residential building was controlled by a novel low-cost E-MPC setup. The controller was designed to rely on already existing equipment in modern residential buildings; the exhaust and outdoor air temperature sensors of the ventilation system, a smart heat meter, and a weather web service to minimize equipment investment costs. The control model was a linear time-invariant black-box model suited for low-cost automated system identification where the original contribution is that an appropriate model for E-MPC was trained on data from a period of approx. two weeks with an excitation signal where the indoor air temperature was kept within typical thermal comfort limits. The model can therefore be trained while occupants are present and comfortable. The experimental approach of the study addresses the current knowledge gap that is holding back the practical realization of the potential identified in the many simulation-based studies and adds to the very limited number of studies on the performance of black-box MPC for residential space heating [26].

2. Method

2.1. Case building

The case building was the Living Lab (Fig. 1) which is an

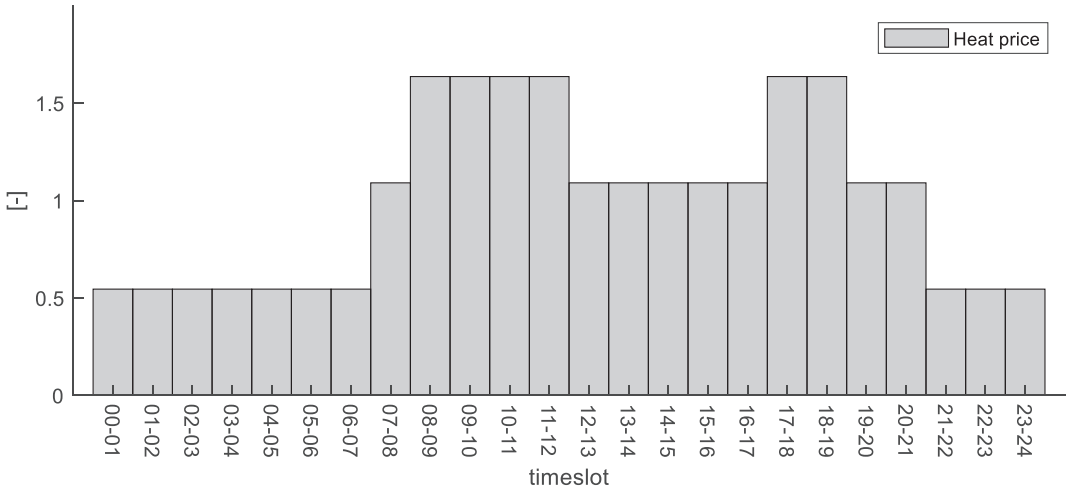

Fig. 3. Daily schedule of time-varying heating prices. The price is unitless and normalized to get a mean of one.

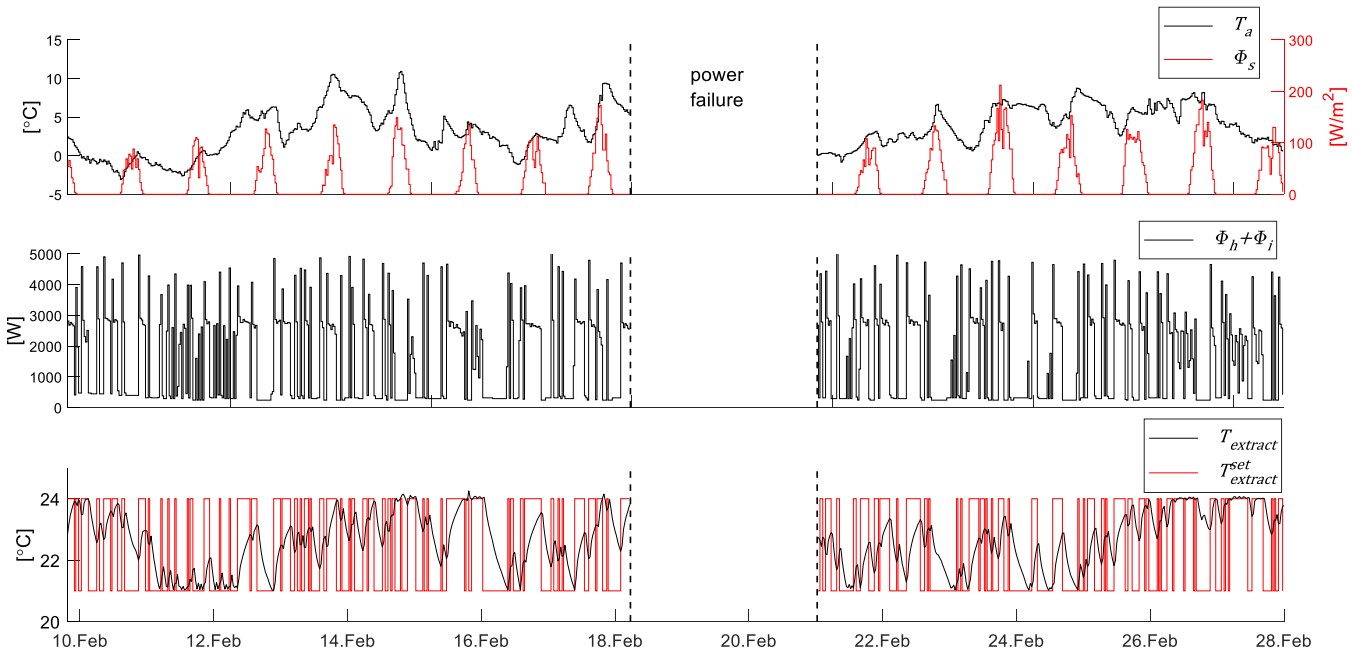

Fig. 4. Training and validation data used for system identification. The top plot shows the ambient temperature and global solar irradiation. The middle plot shows the lumped heat input (space heating and internal heat gain). The bottom plot shows the ventilation extract air temperature as well as the random binary sequence of temperature setpoints.

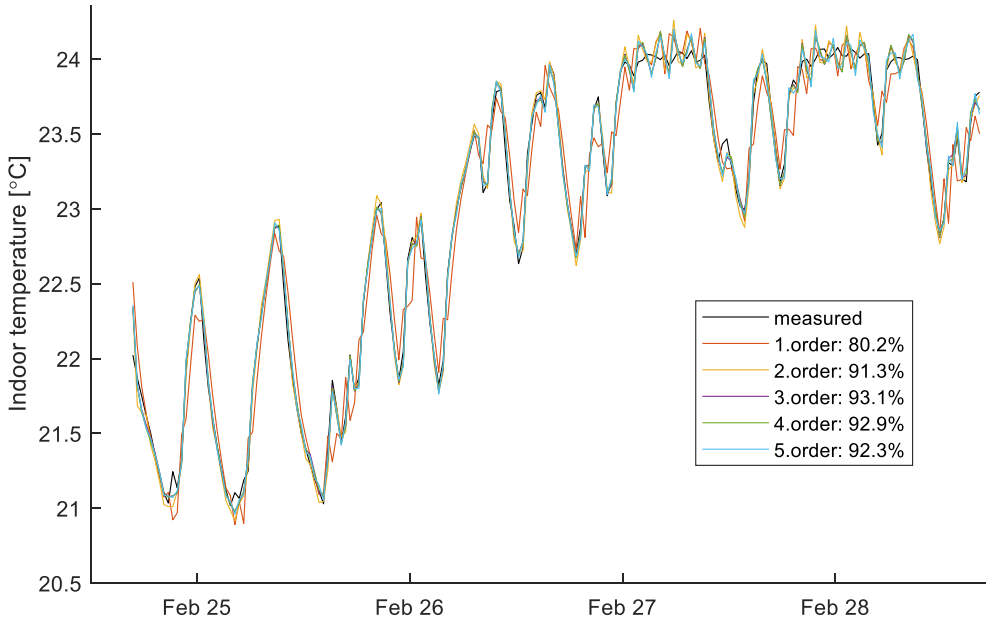

Fig. 5. Comparison of one step ahead predictions on validation data for the refined models of order one to five.

experimental test facility located at the Norwegian University of Science and Technology (NTNU) in Trondheim (temperate oceanic climate) built as a part of the Research Centre on Zero Emission Buildings (ZEB) [38]. The laboratory is constructed as a light-weight single-family detached house with a floor area of roughly 102 m^2. The envelope consists of highly insulated wooden-frame structures (U-values $\approx$ 0.11 W/m^{2}K) with a glass-to-envelope ratio of 23%, and low-energy windows (U-values between 0.65–1.00 W/m^{2}K). A total of 90 m^2 phase-changing materials boards are mounted on the sloped ceilings just behind the wooden cladding. For this experiment, all internal door openings (dotted lines in Fig. 1) were kept open, all windows were kept closed, and external shadings were retracted at all times.

An internal heat gain (Φ_i) consisting of heat gains from electric equipment and four human body sized metal cylinders (dummies) equipped with incandescent lightbulbs to mimic occupants was established. The four dummies were placed as shown in Fig. 1 and were operated to imitate the daily heat gain profile from two adults (Fig. 2).

An air-handling unit (AHU) located in the technical room provided balanced ventilation of approx. 144 m^3/h and the location of inlets and outlets and airflow rates are shown in Fig. 1. A heat recovery effectiveness of approximately 87% was the only pre-heating of the supply air. Space heating was handled by the central hydronic radiator in the living room (see location in Fig. 1). It has previously been shown that a simple heating system with only one radiator can provide acceptable

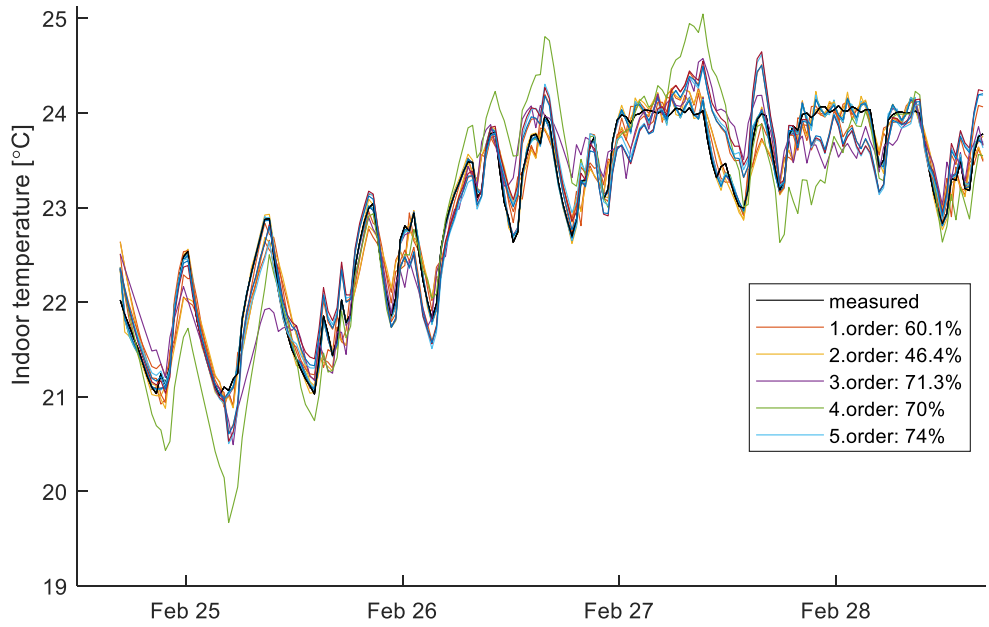

Fig. 6. Comparison of 24 step ahead predictions for the refined models on validation data.

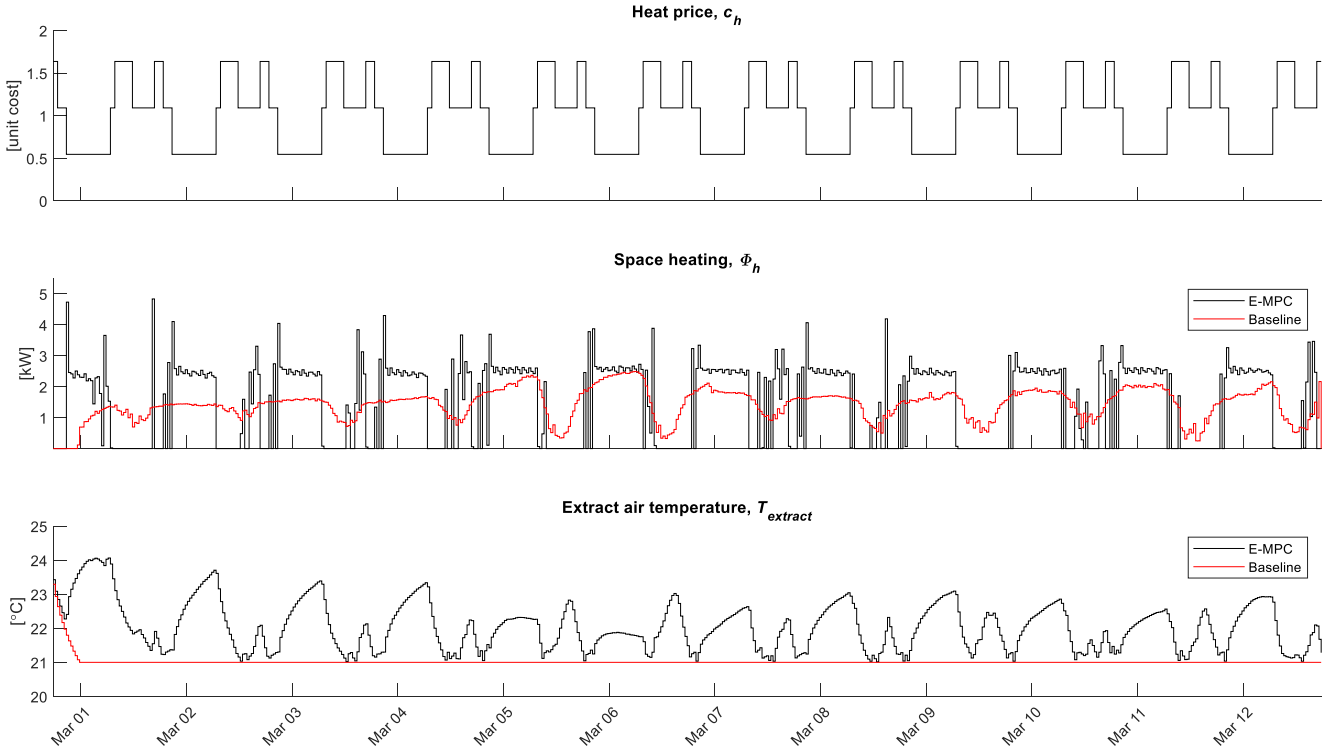

Fig. 7. Comparison between E-MPC and baseline controller with heat price (top), heat consumption for the E-MPC and baseline (middle) and zone air temperature (bottom) for the E-MPC and baseline, respectively.

Table 2
Cost and energy for space heating during the experiment.

	Cost [-]		Energy [kWh]	
Baseline	387		412	
E-MPC	300	(–22.5%)	417	(+1.3%)

thermal comfort in super-insulated buildings [39]. The water supply temperature was fixed to approx. 55 °C during the experiment and the water flow was adjusted by a thermostatic radiator valve to track the setpoint determined by the E-MPC. The ventilation extract air temperature ($T_{extract}$) was used to represent the zone air temperature of the entire building and the radiator thermostat was therefore configured to control this. $T_{extract}$ was measured in the return duct near the AHU and was thus a flow-weighted average of the extract air from the kitchen and the bathroom, respectively. The extract air temperature was a reliable

Table 3
Shows how much energy for space heating is used in the peak, medium and low price periods, respectively.

	Peak [kWh]		Medium [kWh]		Low [kWh]	
Baseline	93		112		207	
E-MPC	17	(−82.0%)	100	(−11.4%)	301	(+45.4%)

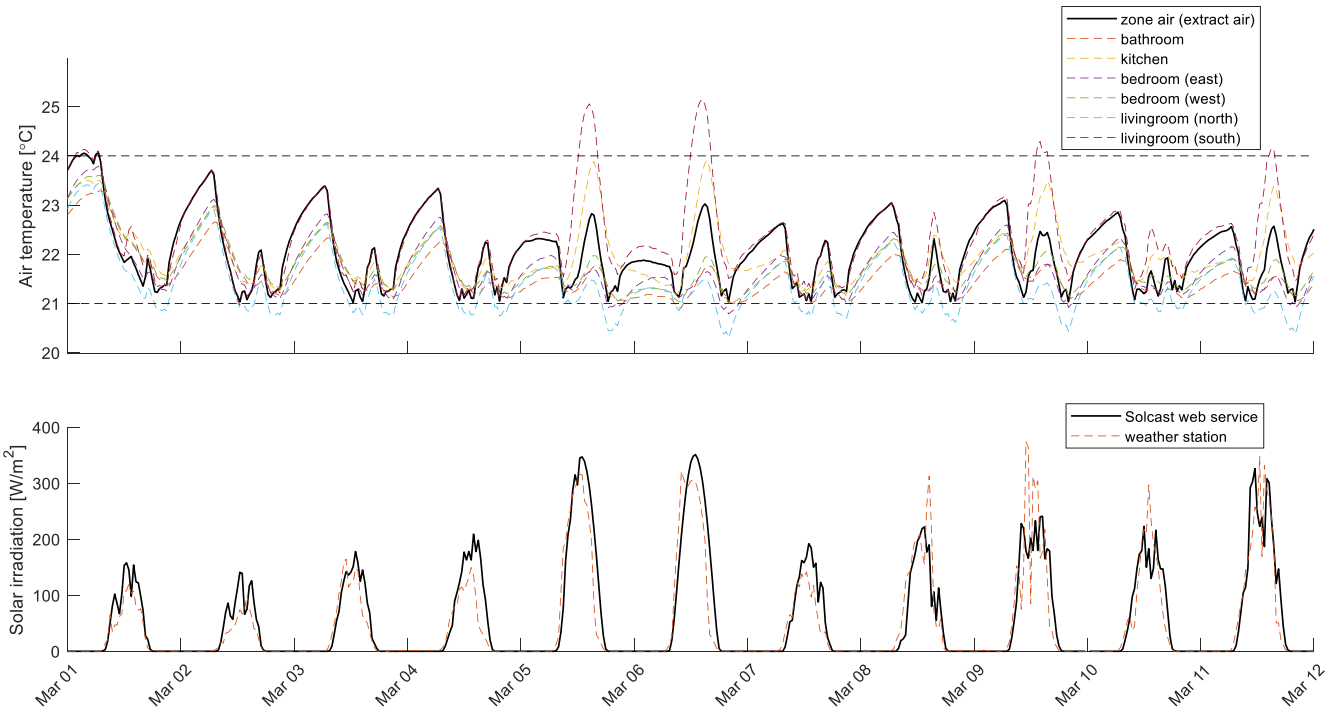

Fig. 8. Air temperatures (top) at different locations in the building and solar irradiation (bottom) from a web service (Solcast) and the local weather station on top of the roof (bottom). All temperature sensors were placed 1.6 m above the floor except the zone air sensor, which was placed in the extraction duct just before the AHU.

representation of the overall temperature throughout the building, as is shown in the results section. One benefit of this implementation was that only a single air temperature sensor was required.

A Kamstrup heat meter (Multical 602) determined the radiator heat power (Φ_h) based on measured flow and temperature difference across the radiator. The outdoor air temperature was measured by a weather station on the roof of the building and forecasts were retrieved from the Norwegian Meteorological Institute weather API [40]. The outdoor temperature forecasts were adjusted to reduce the deviation between the forecast and the measurement (see Fig. 10). Solar irradiation data were retrieved from the Solcast web service [41] that provided both real-time data and forecasts based on satellite images. All data were averaged over one time-step (30 min) before it was used by the E-MPC. The data acquisition was handled by cRIO real-time controllers from National Instrument (model cRIO-9068 and NI-9148) and the LabVIEW software [42].

2.2. Economic model predictive control

The E-MPC was programmed in MATLAB [43] while LabVIEW logged sensor data and updated the thermostat setpoint. LabVIEW was scheduled to call MATLAB every half hour to execute the E-MPC function with updated sensor data and weather forecasts. The following sections describe the model structure, system identification process, and the optimisation problem solved by the E-MPC.

2.2.1. Model structure

The task of generating an appropriate model is often highlighted as the biggest challenge and the main reason that MPC has not been used for building control in practice [20,21,26,44]. Previous research projects have successfully implemented MPC in buildings, but the approaches often relied on expert modelers to build and tune complex models [44]. This makes the implementation unprofitable in practice, especially in single-family houses. In this context, the use of black-box models could have an advantage as they do not require manual collection of prior knowledge of the building to establish a model but instead are calibrated solely from operational data in a process that can be executed using automated procedures. A downside of black-box models is that they are generally known to require exciting training data for longer periods to be reliable and have lower extrapolation properties than e.g. grey-box models. On the other hand, the identification of linear black-box state-space models can be based on subspace methods that are non-iterative and hence there are no convergence problems [45]. For a further description of the advantages and disadvantages of different modeling techniques, see e.g. the review by Kathirgamanathan [26]. Finally, a stochastic subspace method is used so that the identified Kalman gains are directly used in the Kalman filler of the MPC. In other words, no additional work is required to tune the Kalman filter as the Kalman gains are a result of the model identification process.

Previous simulation-based studies have shown that linear and time-invariant (LTI) models approximate the thermal dynamics of many buildings with sufficient accuracy for E-MPC purposes [13,11,8,14,34,46,47]. We therefore implemented a black-box LTI state-space model in innovations form:

$$\boldsymbol{x}[\boldsymbol{k}+1] = \boldsymbol{A}\boldsymbol{x}[\boldsymbol{k}] + \boldsymbol{B}\boldsymbol{u}[\boldsymbol{k}] + \boldsymbol{K}\boldsymbol{e}[\boldsymbol{k}] \quad (1)$$

$$T_{extract}[k] = \boldsymbol{C}x[k] + e[k] \quad (2)$$

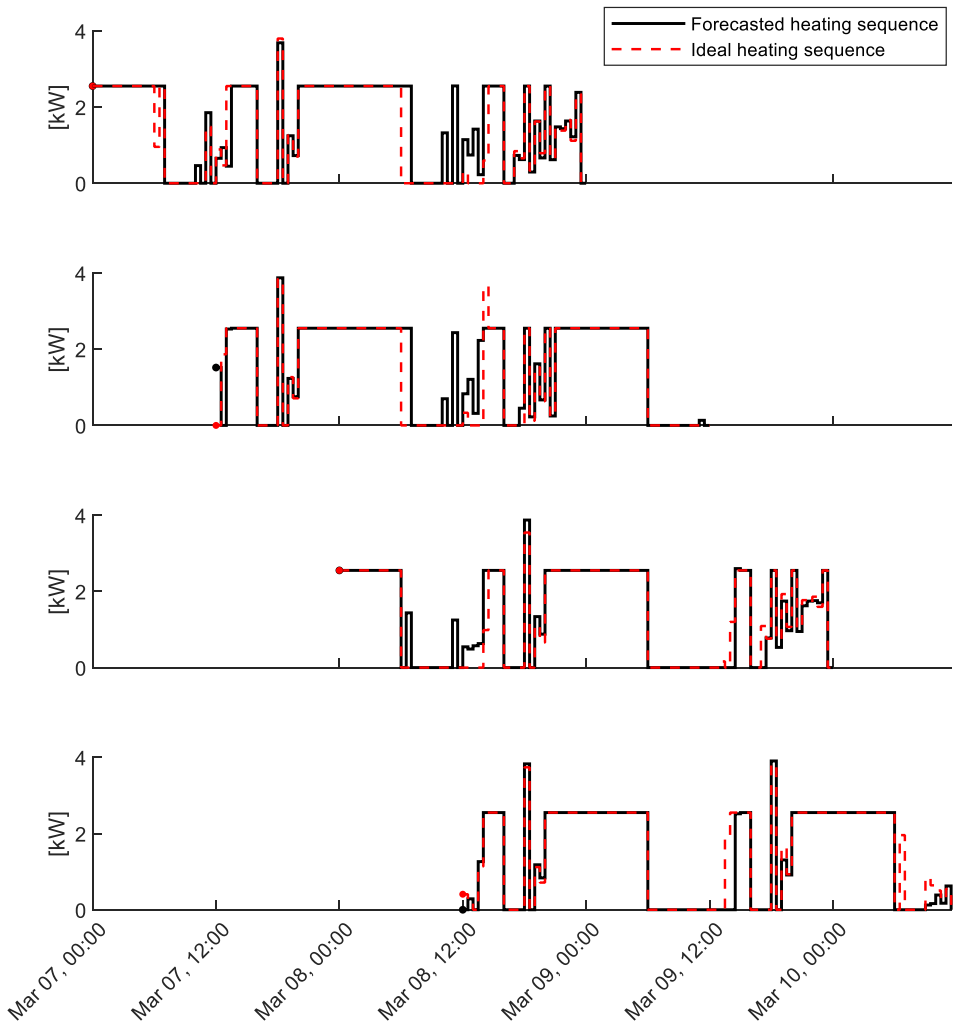

Fig. 9. Optimal heat sequence for prediction horizons determined at different points in time (12 h interval).

where k is the time index, $x[k] \in \mathbb{R}^n$ are the system states, $T_{extract}[k]$ is the measured ventilation extract air temperature [°C], $u[k] \in \mathbb{R}^m$ are inputs and $e[k] \in \mathbb{R}^p$ is the output prediction error (or innovation) [°C]. $\boldsymbol{A} \in \mathbb{R}^{n\times n}$ is the state matrix, $\boldsymbol{B} \in \mathbb{R}^{n\times m}$ is the input matrix,$\boldsymbol{K} \in \mathbb{R}^{n\times p}$ is the Kalman gain matrix and $\boldsymbol{C} \in \mathbb{R}^{p\times n}$ is the output matrix. The input vector contained the following values:

$$u[k] = \begin{bmatrix} \Phi_h[k] + \Phi_i[k] \\ T_a[k] \\ \Phi_s[k] \end{bmatrix} \tag{3}$$

The first input was the total internal heat gains and was a sum of space heating Φ_h [W] and heat gains from people and equipment Φ_i [W], T_a was ambient temperature [°C] and $\Phi_s$was global solar irradiation on a horizontal plane [W/m^2]. The input matrix $\boldsymbol{B}$ was thus comprised of three columns corresponding to each input:

$$\boldsymbol{B} = [\boldsymbol{B_{hi}} \quad \boldsymbol{B_a} \quad \boldsymbol{B_s}] \tag{4}$$

The state-space model in (1) can thus be rewritten as follows:

$$\boldsymbol{x}[\boldsymbol{k}+1] = \boldsymbol{A}\boldsymbol{x}[\boldsymbol{k}] + \boldsymbol{B_{hi}}\boldsymbol{\Phi_h}[\boldsymbol{k}] + \boldsymbol{B_{hi}}\boldsymbol{\Phi_i}[\boldsymbol{k}] + \boldsymbol{B_a}\boldsymbol{T_a}[\boldsymbol{k}] + \boldsymbol{B_s}\boldsymbol{\Phi_s}[\boldsymbol{k}] + \boldsymbol{K}\boldsymbol{e}[\boldsymbol{k}] \tag{5}$$

$$\boldsymbol{T_{extract}}[\boldsymbol{k}] = \boldsymbol{C}\boldsymbol{x}[\boldsymbol{k}] + \boldsymbol{e}[\boldsymbol{k}] \tag{6}$$

Notice that space heating and internal gains were assumed to enter the system through the same input matrix $\boldsymbol{B_{hi}}$ which reduced the number of unknown model parameters and hence avoid overfitting [48].

2.2.2. System identification

Several simulation studies have followed similar procedures to identify building models for E-MPC as this experimental study albeit with some differences, see e.g. [13,11,8,14]. The applied system identification procedure consisted of the following five steps:

Step 1: Generate data.

It is of practical interest to investigate whether a reliable black-box model can be identified using excitation of the building thermal dynamics within the comfortable indoor temperature limits. If possible, training data can be generated while the residential building is occupied. Therefore, a random sequence of half-hourly binary temperature setpoints of either 21 or 24 °C was sent to the radiator thermostat. This induced an exciting sequence of radiator heat inputs. This data set was then split into a training and a validation dataset so that cross-validation could be used to choose an appropriate model order.

Step 2: Identify models using subspace system identification (n4sid).

The training data was used to identify five models with model order ranging from 1 to 5 states using subspace system identification. This method can be applied to identify unstructured LTI state-space models as (1–2), i.e. where no internal structure or constraints are imposed on the parameters in the system matrices. The theory behind this approach is described in detail in [45] and implemented in the *n4sid* function that is a part of the MATLAB System Identification Toolbox [49]. This function is associated with many manual user settings that can have a substantial influence on the model quality [50]. Table 1 lists the settings used in this study.

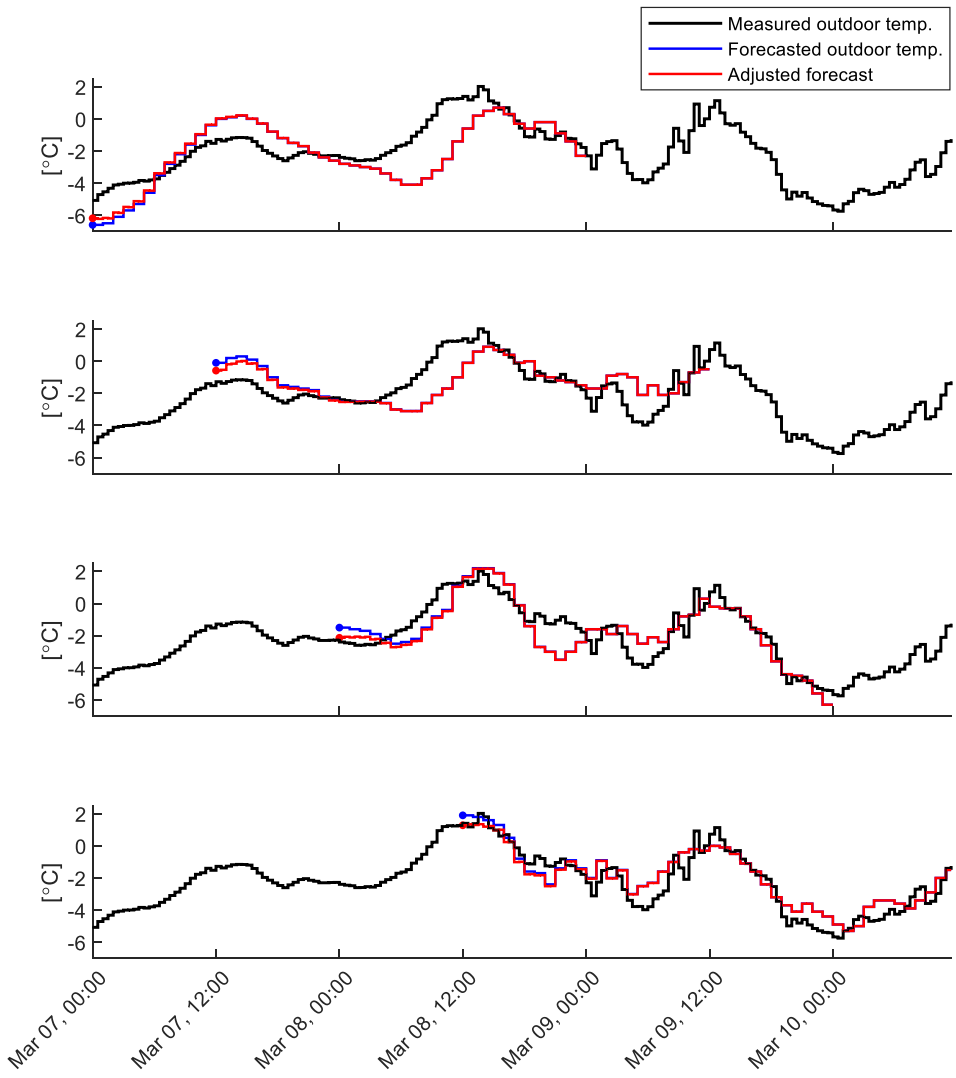

Fig. 10. Comparison of measured outdoor air temperature and two days forecasts forecasted at four different four time points.

Step 3: Refine models using the prediction error method (pem).

The training data was then used to refine the five subspace models using the prediction error method [51]. This method is implemented in the *pem* MATLAB function that is also part of the MATLAB System Identification Toolbox [49]. The prediction error method can often identify models with better predictive performance than subspace methods. However, one of the disadvantages is that they rely on iterative search schemes, which means that if initiated with a poor initial model it can take a long time to converge and risk being stuck in local minima. This can be avoided by generating initial models using subspace identification (step 2), which is based on linear projections and thus very fast execute and is not stuck in local minima.

Step 4: Choose a model.

The order of the state-space model was identified by comparing the one-step-ahead prediction errors of the five refined models to the validation data.

Step 5: Final refinement using all data.

All data (i.e. training and validation) was used to refine the model picked in step 4 using *pem*. This step ensured that all information in the data was incorporated into the final model. The results of these five steps are presented in the results section.

2.2.3. Optimization problem

A linear program Eq. (7a)-(7f) was solved to determine the sequence of heat inputs $\Phi_h[k]$ minimizing total heating costs over all time steps $k \in 0, 1, \cdots, N-1$ in the prediction horizon. The *linprog* function in MATLAB was used to solve this problem and it took less than a second using the dual-simplex algorithm [43]. In this experiment, the time step was 30 min and the prediction horizon was two days, i.e. $N = 95$, based on the experience from previous simulation studies [13].

$$\min_{\Phi_h} \sum_{k=0}^{N-1} c_h[k]\Phi_h[k] \tag{7a}$$

Subject to

$$x[k+1] = \boldsymbol{A}x[k] + \boldsymbol{B_h}\Phi_h[k] + \boldsymbol{B_h}\Phi_i[k] + \boldsymbol{B_a}T_a[k] + \boldsymbol{B_s}\Phi_s[k], k \in \mathbb{N}_0^{N-1} \tag{7b}$$

$$T_{extract}[k] = \boldsymbol{C}x[k], k \in \mathbb{N}_0^{N-1} \tag{7c}$$

$$21^{\circ}C \leq T_{extract}[k] \leq 24^{\circ}C, k \in \mathbb{N}_0^{N-1} \tag{7d}$$

$$0 \leq \Phi_h[k] \leq \Phi_{h,max}[k], k \in \mathbb{N}_0^{N-1} \tag{7e}$$

$$x[0] = \hat{x}[0] \tag{7f}$$

A receding horizon procedure was implemented, which means the linear program was solved every half hour for a receded prediction horizon and with updated initial state estimates ($\hat{x}[0]$) and updated weather forecasts. The resulting output temperature in the first time step

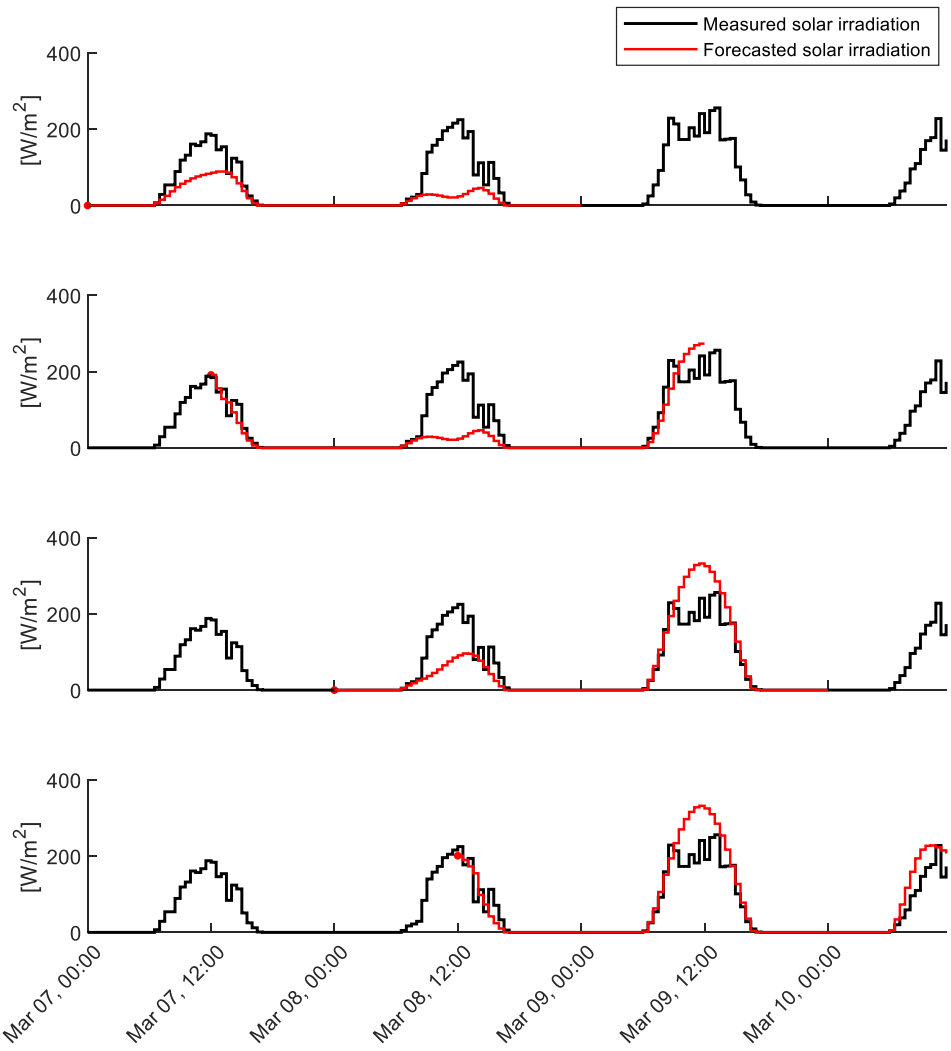

Fig. 11. Comparison of measured solar irradiation and two days forecasts forecasted at four different four time points.

($T_{extract}[1]$) was then used as a setpoint for the radiator thermostat. Slack variables were used to obtain soft constraints in the implemented linear program [7,11], but omitted in Eq. (7a)-(7f) for simplicity.

Fig. 3 shows the time-varying heat prices (c_h) used in the experiment. The prices were repeated daily and consisted of three price levels corresponding to typical levels of demand in district heating systems. Note that the optimal solution is independent of the particular unit (or currency) of cost and is thus normalized in Fig. 3 to get a mean heat price of one.

3. Results

The experiment lasted 32 days of which 18 days were used to generate data for system identification and 14 days with the E-MPC in control. Fig. 4 shows the 18 days used for system identification and includes: ambient temperature T_a and solar irradiation Φ_s (top), space heating Φ_h and internal heat gains Φ_i lumped together (middle), ventilation extract air temperature $T_{extract}$ and the random sequence of setpoints $T^{set}_{extract}$ (bottom). The NTNU Campus experienced a power failure in the electricity network, which stopped the controller and hence the data collection for approximately three days.

The first 14 days of this data were used as training data and the last four days were reserved as validation data to determine the model order. Fig. 5 shows the one step ahead predictions for the five refined models (step 3) and the corresponding fit percentages (Normalized root mean square error) on the validation data. The third-order model was selected for the E-MPC since it had the highest fit (step 4). Fig. 6 compares the 24 step ahead predictions (half a day) to show how the models perform for a longer prediction horizon, although this was not used in the system identification procedure. Finally, the third-order model was refined using both training and validation data to utilize all available information to estimate the parameters (step 5).

Fig. 7 compares the performance of the E-MPC with a constant setpoint baseline controller. The baseline was never applied to the actual building but was a simulation in which an ideal (no offset) setpoint tracking controller maintained a constant setpoint of 21 °C. The baseline was simulated using a state-space model trained using all available data - including data from the MPC phase. The comparative results between the E-MPC and the baseline should therefore be taken with some reservations as it assumes that the baseline state-space model is the true building. The baseline can thus only be used to indicate how an actual setpoint tracking PID would perform. The figure shows the time-varying heat price (top), heat consumption (middle) and extract air temperature (bottom) for the E-MPC and the baseline controller, respectively. While the baseline simply took the necessary control actions to meet the setpoint without regard to the heat prices, the E-MPC shifted heating from peak to low price periods. As a result, the air temperature increased in periods with low prices but it never exceeded the upper constraints of 24 °C.

Table 2 summarizes the total heating costs and heat energy

consumed during the experiment. The E-MPC consumed 1.3% more energy than the baseline, but in turn, it reduced costs by 22.5%. This tendency of consuming more energy to reduce costs is also shown in previous simulation studies [13,11,8,14]. It should be noted that the cost savings depends largely on the building characteristics (insulation level and thermal mass) as well as the characteristic of the price signal as pointed out in [14].

Table 3 summarizes how much heat was consumed in peak, medium and low price periods, respectively. The E-MPC reduced the amount of heat consumed in peak price periods by 82% and increased the consumption in low price periods by 45% thus flattening out the system load profile.

Fig. 8 compares the extract air temperature (black solid) with the air temperature in different sections of the living lab. The room air temperature sensors were mounted in white casings on walls at a height of 1.6 m above the floor. It is seen, that the extract air temperature was a reasonable representation of the overall air temperature. There were, however, some significant temperature deviations, especially between the north and south parts of the living room. The differences were mainly in periods with large solar irradiation. The living room south sensor was mounted on an internal wall close to the radiator and it may have been exposed to direct solar radiation at periods. Contrary, the living room north sensor was mounted on an external wall close to a window and these circumstances could explain the relatively large difference. Fig. 8 also compares solar irradiation obtained from the web service and the solar irradiation measured by the weather station on top of the case building. Overall, there was a good agreement in terms of daily level of solar irradiation, but there was some short-term deviations due to local cloud movements especially on partially cloudy days.

Fig. 9 presents another way to gain confidence in the E-MPC and the applied model. It shows four optimal heat sequences determined by the E-MPC four different times with 12 h intervals. The sequences change, which was to be expected, but the overall pattern for the first 24–36 h of the sequences were stable, which is an important condition for a well-functioning MPC. Minor adjustments were expected due to the Kalman filter correcting the initial states every half hour and the regularly updating of weather forecasts. Finally, the prediction horizon was continuously receded in time, which causes the relatively large changes in the last part of the sequences. The stability of the first part of the sequence indicates that the model produced consistent predictions which in turn indicates that the initial states were only moderately corrected in each time step and the weather forecasts updates were not dramatic. The red dashed curve is the ideal heat sequence, i.e. the sequence obtained when optimized using the actual weather data instead of the forecasts. The two solutions are very similar which indicates that the weather prediction errors only had a relatively small effect on the optimal sequence at least for this time period.

Figs. 10 and 11 show the outdoor temperature and solar irradiation forecasts, respectively, at the same four times as Fig. 9. The blue curve in Fig. 10 shows the corrected forecast of the outdoor temperature.

4. Discussion

The findings of the experiment indicate that it is practically possible to realise the significant DR potential of utilising E-MPC to exploit the inherent thermal mass in residential buildings outlined in the introduction. Furthermore, the experiment suggests that it is practically possible to establish a black-box model from a random sequence of half-hourly binary temperature set points of either 21 or 24 °C, i.e. within a typical occupant comfort interval. The implemented E-MPC setup used a minimum of sensors; solar irradiation data was retrieved from a web service, the outdoor temperature was measured using the sensor already installed to control the AHU (could alternatively be retrieved from the web service together with the solar irradiation), and the temperature sensor in the extract air duct AHUs was used to represent indoor air temperature. This minimal setup can likely be replicated with a low investment cost in other residential buildings. However, several issues remain to be investigated and handled in future works:

1. *Longer test period* - The E-MPC was only tested for approx. two weeks. It therefore remains to be validated whether the black-box model performs satisfactorily over longer periods with changing weather conditions. It can also be expected that the surrounding building and the building itself will change over time. For this reason, it is therefore expected that the model needs to be recalibrated; appropriate strategies for this should be investigated.
2. *Side*-by-Side experiment – This experiment applied a simulated baseline for comparison. However, such a baseline assumes that the simulated building is an exact representation of the true building. A better baseline, but more difficult to realize, would be a completely identical building subjected to the baseline control strategy.
3. *Occupancy* - The robustness of the setup should be tested during occupancy.
4. *Multiple thermal zones* - When a building has more than a single radiator, it might be necessary to separate the building into several thermal zones. The E-MPC should be able to handle these cases.
5. *Supply temperature* – This experiment assumed a constant district heating supply temperature; however, it can be expected to vary to some extent during the year, which needs to be taken into account.
6. *Return temperature* – The return temperature was not considered in this experiment, but future works should include it in the optimization problem to avoid high return temperatures. One approach would be to include constraints regarding the return temperature in the E-MPC cost function. Another approach could be to include the return temperature in the cost function making it a multi-objective optimization problem; however, this requires a more advanced control model as it besides zone air temperature also must output the return temperature.

5. Conclusion

This study aimed to conduct an experiment on the performance of a novel low-cost E-MPC that utilise the thermal mass of a building as short-term heat storage to perform DR. The experiment was conducted in a highly insulated detached single-family house in Norway heated by a single hydronic radiator. The E-MPC used standard sensors usually available in buildings equipped with balanced mechanical ventilation, a linear black-box state-space thermal model suited for low-cost automated model identification and weather forecasts from an existing web-based application.

Data from the experiment showed that the predictive performance of the linear black-box model was sufficient for its purpose even though the model was trained on excitation data that was generated using setpoints of either 21 or 24 °C during a relatively short period of two weeks. Consequently, the thermal comfort remains acceptable during the excitation period which allows the identification to be performed during normal occupancy. Although the building was treated as a single thermal zone the E-MPC still managed to obtain good control of the air temperature in the entire building while shifting consumption from high to low price periods. The findings of the experiment indicate that a minimal E-MPC setup is able to realize the significant DR potential that lies in utilizing the inherent thermal mass in residential buildings.

CRediT authorship contribution statement

Michael Dahl Knudsen: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Data curation, Writing - original draft, Writing - review & editing, Visualization, Project administration. **Laurent Georges:** Resources, Writing - review & editing, Supervision, Project administration, Funding acquisition. **Kristian Stenerud Skeie:** Software, Investigation, Writing - review & editing. **Steffen Petersen:** Resources, Writing - review & editing, Supervision,

Funding acquisition.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

This work is supported by the REVALUE project financed by Innovation Fund Denmark, the PreHeat project funded by the Danish Energy Technology Development and Demonstration Program (EUDP), and the Norwegian Research Centre on Zero Emission Neighbourhoods in Smart Cities (FME ZEN). The authors would also like to acknowledge Odne Oksavik for providing us with technical assistance in setting up the living lab and Solcast [41] for providing solar irradiation forecasts.

References

[1] O'Connell N, Pinson P, Madsen H, O'Malley M. Benefits and challenges of electrical demand response: A critical review. Renew Sustain Energy Rev 2014;39:686–99.
[2] Nolan S, O'Malley M. Challenges and barriers to demand response deployment and evaluation. Appl Energy 2015;152:1–10.
[3] Esther BP, Kumar KS. A survey on residential Demand Side Management architecture, approaches, optimization models and methods. Renew Sustain Energy Rev 2016;59:342–51.
[4] Jensen SØ, Marszal-Pomianowska A, Lollini R, Pasut W, Knotzer A, Engelmann P, et al. IEA EBC Annex 67 Energy Flexible Buildings. Energy Build 2017;155:25–34.
[5] Le Dréau J, Heiselberg P. Energy flexibility of residential buildings using short term heat storage in the thermal mass. Energy 2016;111:991–1002.
[6] Reynders G, Diriken J, Saelens D. Generic characterization method for energy flexibility: Applied to structural thermal storage in residential buildings. Appl Energy 2017;198:192–202.
[7] Halvgaard R, Poulsen NK, Madsen H, Jørgensen JB. Economic Model Predictive Control for Building Climate Control in a Smart Grid. In 2012 IEEE PES Innovative Smart Grid Technologies (ISGT), Washington DC; 2012.
[8] Knudsen MD, Hedegaard RE, Pedersen TH, Petersen S. System identification of thermal building models for demand response – A practical approach. Energy Procedia 2017;122:937–42.
[9] Awadelrahman MAA, Zong Yi, Li H, Agert C. Economic Model Predictive Control for Hot Water Based Heating Systems in Smart Buildings. Energy Power Eng 2017; 09(04):112–9.
[10] Ma J, Qin SJ, Li B, Salsbury T. Economic model predictive control for building energy systems. In ISGT 2011, Anaheim; 2011.
[11] Knudsen MD, Petersen S. Economic model predictive control of space heating and dynamic solar shading. Energy Build 2020;209:109661. https://doi.org/10.1016/j.enbuild.2019.109661.
[12] Kuboth S, Heberle F, König-Haagen A, Brüggemann D. Economic model predictive control of combined thermal and electric residential building energy systems. Appl Energy 2019;240:372–85.
[13] Knudsen MD, Petersen S. Demand response potential of model predictive control of space heating based on price and carbon dioxide intensity signal. Energy Build 2016;125:196–204.
[14] Knudsen MD, Hedegaard RE, Heidmann T, Petersen S. Model predictive control of space heating and the impact of taxes on demand response: A simulation study. CLIMA 2016 - Proceedings of the 12th REHVA World Congress. 2016.
[15] Finck C, Li R, Zeiler W. Optimal control of demand flexibility under real-time pricing for heating systems in buildings: A real-life demonstration. Appl Energy 2020;263:114671. https://doi.org/10.1016/j.apenergy.2020.114671.
[16] Finck C, Li R, Zeiler W. Economic model predictive control for demand flexibility of a residential building. Energy 2019;176:365–79.
[17] Zong Y, Böning GM, Santos RM, You S, Hu J, Han X. Challenges of implementing economic model predictive control strategy for buildings interacting with smart energy systems. Appl Therm Eng 2017;114:1476–86.
[18] Khanmirza E, Esmaeilzadeh A, Markazi AHD. Design and experimental evaluation of model predictive control vs. intelligent methods for domestic heating systems. Energy Build 2017;150:52–70.
[19] Yu Na, Salakij S, Chavez R, Paolucci S, Sen M, Antsaklis P. Model-based predictive control for building energy management: Part II – Experimental validations. Energy Build 2017;146:19–26.
[20] Sturzenegger D, Gyalistras D, Morari M, Smith RS. Model Predictive Climate Control of a Swiss Office Building: Implementation, Results, and Cost-Benefit Analysis. IEEE Trans Control Syst Technol 2016;24(1):1–12.
[21] Coninck RD, Helsen L. Practical implementation and evaluation of model predictive control for an office building in Brussels. Energy Build 2016;111(1): 290–8.
[22] Liao Z, Dexter AL. An Inferential Model-Based Predictive Control Scheme for Optimizing the Operation of Boilers in Building Space-Heating Systems. IEEE Trans Control Syst Technol 2010;18(5):1092–102.
[23] Široký J, Oldewurtel F, Cigler J, Prívara S. Experimental analysis of model predictive control for an energy efficient building heating system. Appl Energy 2011;88(9):3079–87.
[24] Prívara S, Široký J, Ferkl L, Cigler J. Model predictive control of a building heating system: The first experience. Energy Build 2011;43(2-3):564–72.
[25] Ferreira PM, Ruano AE, Silva S, Conceição EZE. Neural networks based predictive control for thermal comfort and energy savings in public buildings. Energy Build 2012;55:238–51.
[26] Kathirgamanathan A, De Rosa M, Mangina E, Finn DP. Data-driven predictive control for unlocking building energy flexibility: A review. Renew Sustain Energy Rev 2021;135:110120. https://doi.org/10.1016/j.rser.2020.110120.
[27] European Commission. eurostat: Your key to European statistics. Eurostat; 6 June 2020. [Online]. Available: https://ec.europa.eu/eurostat/data/database [Accessed 21 July 2020].
[28] Hedegaard Rasmus Elbæk, Kristensen Martin Heine, Pedersen Theis Heidmann, Brun Adam, Petersen Steffen. Bottom-up modelling methodology for urban-scale analysis of residential space heating demand response. Appl Energy 2019;242: 181–204.
[29] Dominković DF, Gianniou P, Münster M, Heller A, Rode C. Utilizing thermal building mass for storage in district heating systems: Combined building level simulations and system level optimization. Energy 2018;153:949–66.
[30] Cai Hanmin, Ziras Charalampos, You Shi, Li Rongling, Honoré Kristian, Bindner Henrik W. Demand side management in urban district heating networks. Appl Energy 2018;230:506–18.
[31] Hedegaard RE, Friedrichsen L, Tougaard J, Mølbak T, Petersen S. Building energy flexibility as an asset in system-wide district heating optimization models. uSIM2020 - Building to Buildings: Urban and Community Energy Modelling. 2020.
[32] Killian M, Kozek M. Ten questions concerning model predictive control for energy efficient. Build Environ 2016;105:403–12.
[33] Hedegaard Rasmus Elbæk, Pedersen Theis Heidmann, Knudsen Michael Dahl, Petersen Steffen. Towards practical model predictive control of residential space heating: Eliminating the need for weather measurements. Energy Build 2018;170: 206–16.
[34] Prívara Samuel, Cigler Jiří, Váňa Zdeněk, Oldewurtel Frauke, Sagerschnig Carina, Žáčeková Eva. Building modeling as a crucial part for building predictive control. Energy Build 2013;56:8–22.
[35] Atam E, Helsen L. Control-Oriented Thermal Modeling of Multizone Buildings: Methods and Issues: Intelligent Control of a Building System. IEEE Control Syst Mag 2016;36(3):86–111.
[36] Drgoňa Ján, Arroyo Javier, Cupeiro Figueroa Iago, Blum David, Arendt Krzysztof, Kim Donghun, et al. All you need to know about model predictive control for buildings. Ann Rev Control 2020;50:190–232.
[37] Viot H, Sempey A, Mora L, Batsale JC, Malvestio J. Model predictive control of a thermally activated building system to improve energy management of an experimental building: Part I—Modeling and measurements. Energy Build 2018; 172:94–103.
[38] ZEB. The Research Centre on Zero Emission Buildings. Xmedia; 2020. [Online]. Available: https://www.zeb.no/index.php/en/pilot-projects/158-living-lab-trondheim.
[39] Georges L, Alonso MJ, Woods R, Wen K, Håheim F, Liu P, et al. ZEB Project report 39: Evaluation of Simplified Space-Heating Hydronic Distribution for Norwegian Passive Houses. The Research Centre on Zero Emission Buildings (ZEB); 2017.
[40] Norwegian Meteorological Institute (MET Norway). Location Forecast; 2018. [Online]. Available: https://api.met.no/weatherapi/locationforecast/1.9/documentation.
[41] Solcast. Solar forecasting and real-time irradiance data; 2018. [Online].
[42] National Instruments. LabVIEW 2016 - Laboratory Virtual Instrument Engineering Workbench.
[43] The MathWorks Inc. MATLAB R2018a.
[44] Gwerder M, Gyalistras D, Sagerschnig C, Smith R. Final report: Use of weather and occupancy forecasts for optimal building climate control – Part II: Demonstration (OptiControl-II). ETH Zurich, Zurich: Automatic Control Laboratory; 2013.
[45] van Overschee P, Moor BD. Subspace Identification for Linear Systems: Theory, Implementation, Applications. New York City: Kluwer Academic Publishers; 1996.
[46] Bacher P, Madsen H. Identifying suitable models for the heat dynamics of buildings. Energy Build 2011;43(7):1511–22.
[47] Vogler-Finck P, Clauß J, Georges L, Sartori I, Wisniewski R. Inverse model identification of the thermal dynamics of a Norwegian zero emission house. Cold Climate HVAC 2018. 2018..
[48] Yu X, Georges L, Knudsen MD, Sartori I, Imsland L. Investigation of the Model Structure for Low-Order Grey-Box Modeling of Residential Buildings. In Building Simulation 2019, Rome; 2020.
[49] The MathWorks, Inc. MATLAB R2018a, System Identification Toolbox 9.8.
[50] Ljung L. Aspects and Experiences of User Choices in Subspace Identification Methods. IFAC Proceedings Volumes. 2013.
[51] Ljung L. System Identification: Theory for the User. 2nd ed. New Jersey: Prentice Hall; 1999.