

Validation of a Model Predictive Control Strategy on a High Fidelity Building Emulator

Original

Validation of a Model Predictive Control Strategy on a High Fidelity Building Emulator / Fop, D., Yaghoubi, A.R., Capozzoli, A.. - In: ENERGIES. - ISSN 1996-1073. - ELETTRONICO. - 17:20(2024), pp. 1-20. [10.3390/en17205117]

Availability:

This version is available at: 11583/2995395 since: 2024-12-15T11:43:07Z

Publisher:

Multidisciplinary Digital Publishing Institute (MDPI)

Published

DOI:10.3390/en17205117

Terms of use:



This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Article

Validation of a Model Predictive Control Strategy on a High Fidelity Building Emulator

Davide Fop , Ali Reza Yaghoubi and Alfonso Capozzoli * 

TEBE Research Group, BAEDA Lab, Department of Energy “Galileo Ferraris”, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Turin, Italy; davide.fop@polito.it (D.F.)

* Correspondence: alfonso.capozzoli@polito.it

Abstract: In recent years, advanced controllers, including Model Predictive Control (MPC), have emerged as promising solutions to improve the efficiency of building energy systems. This paper explores the capabilities of MPC in handling multiple control objectives and constraints. A first MPC controller focuses on the task of ensuring thermal comfort in a residential house served by a heat pump while minimizing the operating costs when subject to different pricing schedules. A second MPC controller working on the same system tests the ability of MPC to deal with demand response events by enforcing a time-varying maximum power usage limitation signal from the electric grid. Furthermore, multiple combinations of the control parameters are tested in order to assess their influence on the controller performance. The controllers are tested on the BOPTTEST framework, which offers standardized test cases in high-fidelity emulation models, and pre-defined baseline control strategies to allow fair comparisons also across different studies. Results show that MPC is able to handle multi-objective optimal control problems, reducing thermal comfort violations by between 66.9% and 82% and operational costs between 15.8% up to 20.1%, depending on the specific scenario analyzed. Moreover, MPC proves its capability to exploit the building thermal mass to shift heating power consumption, allowing the latter to adapt its time profile to time-varying constraints. The proposed methodology is based on technologically feasible steps that are intended to be easily transferred to large scale, in-field applications.

Keywords: energy flexibility; Model Predictive Control; building energy management; high fidelity building emulator



Citation: Fop, D.; Yaghoubi, A.R.; Capozzoli, A. Validation of a Model Predictive Control Strategy on a High Fidelity Building Emulator. *Energies* **2024**, *17*, 5117. <https://doi.org/10.3390/en17205117>

Academic Editor: Antonio Gagliano

Received: 2 September 2024

Revised: 4 October 2024

Accepted: 11 October 2024

Published: 15 October 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Research by the International Energy Agency (IEA), the United Nations Environment Program, and the Global Alliance for Buildings and Construction reveals that the building sector alone is responsible for 37% of global energy-related CO₂ emissions [1], amounting to 13.6 Gt CO₂-eq/year and 36% of global energy consumption. Notably, over half of the energy consumption within buildings is attributed to heating, ventilation, and air conditioning (HVAC) systems [2]. A complementary solution to the renovation of building envelopes is the improvement of the energy management systems, particularly heating, ventilation, and air conditioning (HVAC) systems, since they account for almost 40% of a house’s average energy consumption [3].

The thermal inertia inherent in buildings offers flexibility in operation, which is particularly valuable when the building is integrated in a power grid accommodating intermittent renewable energy sources like wind and solar power [4]. Considering the points mentioned earlier, the focus is on improving control of building HVAC systems. This heightened focus is driven by imperatives such as CO₂ emission reduction, the integration of renewable and distributed energy sources into grids, adaptation to emergencies, and management of complex system architectures [5]. The challenge lies in ensuring indoor comfort, safety, energy efficiency, and flexibility through suitable control mechanisms [6].

Numerous control techniques have been devised to manage a building's operation and its technical mechanisms, which adds complexity to choosing the most fitting control strategy for each unique situation. Nonetheless, these techniques can generally be grouped into two primary categories: traditional control strategies (TCS) and advanced control strategies (ACS). Traditional control strategies can be classified into sequency control strategies and process control strategies [7]. The former class, also referred to as *rule-based* control, is based on pre-defined *if-then* conditions which make the system switch between operative modes. The definition of such conditional rules relies on the designer's domain experience rather than on an optimization process. The latter class adjusts the control variables to achieve the process objectives against disturbances by reacting to the evolution of state and/or disturbance variables. Process objectives mostly represent the desired attainment of a setpoint for a given controlled variable. The TCS approach is simple and robust but cannot adapt to external factors like weather or electricity prices efficiently. It also depends on fine-tuning PID parameters, which are complex [8]. The main drawbacks of the TCS approach are:

1. Poor control accuracy, resulting in thermal discomfort and low energy efficiency [9].
2. Inability to predict the future states of the system and find an optimal control signal sequence [8].
3. Inability to leverage buildings' thermal capacity to shift loads [10].
4. Difficulty in dealing with external grid requirements [8].

To overcome the drawbacks of traditional controllers, attention has shifted to advanced control methods. These methods offer the potential to reduce environmental impact and enhance the integration of renewable energy by optimizing how buildings manage their energy usage [11]. ACS can handle external elements like weather and variable electricity rates, as well as internal demands imposed on consumers [12]. Monitoring data on energy consumption and user behavior is gathered in ACS and used as inputs and feedback to help the system adapt to the needs of the occupants [13]. Initiatives like ASHRAE's Guideline 36, [14] the development of grid-friendly strategies, and growing interest in data-driven control [15] underscore the drive to improve control algorithms. However, even in new constructions, traditional rule-based controls prevail, though the potential benefits of integrating prediction and optimization techniques for operational efficiency are gaining recognition [16]. Recent decades have witnessed a surge in research on advanced control techniques, including Model Predictive Control (MPC), for optimal HVAC system operation [16–18]. The MPC is a model-based controller that adjusts the input of a dynamic model to generate the optimal control strategy over a time frame that begins at the present and extends over a specified prediction horizon in the future [19,20].

The system's mathematical model, present state measurements, and weather forecast are utilized to forecast and optimize the building's future behavior [16]. The MPC has been used in building energy control for a variety of objectives, including indoor comfort [21] and energy cost reduction [22], thanks to advancements in computer technology. During a two-month test period, Blum et al. achieved roughly 40% energy savings on an HVAC system in a real office building in the United States by testing the influence of the MPC algorithm [23]. Freund et al. showed the Model Predictive Control (MPC) algorithm's potential for energy savings by successfully implementing it in a sizable building. The team reported a 75% reduction in heating energy use throughout a three-month test period [24]. Employing MPC in building applications might boost energy efficiency by 10% to 30% and enhance inhabitants' indoor comfort [25]. The ability to forecast a building's future behavior using a mathematical model of the structure is what makes MPC so effective. With the help of these projections, MPC can systematically and adaptably select the best course of action for control actions based on a specified aim, comfort and technology limitations, and weather forecasts [16].

The development of innovative control algorithms requires them to be tested in order to assess their effectiveness and convenience. The direct application of newly conceived controllers to real-world buildings might result in improper operation, occupant discomfort,

or equipment damage; moreover, experimental campaigns are time-consuming [5]. Simulation offers a safe environment in which control strategies can be tested, avoiding costs and risks associated with real-world implementation, while allowing diverse scenarios to be tested at will. On the other hand, the performance assessment of new controllers must be benchmarked against either state-of-the-art controllers or other innovative strategies. To address these challenges, the Building Optimization Performance Testing (BOPTTEST) framework [5] was created to provide a publicly available software that allows “*simulation-based benchmarking of building HVAC control algorithms using high-fidelity building emulators*”. Along with the software infrastructure, BOPTTEST provides a set of test cases based on high-fidelity models equipped with reactive, state-of-the-art control logic, so that any new strategies can be fairly compared to the same baseline controller. Finally, the tool computes a set of Key Performance Indicators (KPI) that further ease comparison between the tested solutions. The BOPTTEST platform has been increasingly used by researchers to test advanced energy management strategies for buildings.

Table 1 showcases a variety of advanced control strategies applied to different case studies using the BOPTTEST framework. Model Predictive Control (MPC) consistently demonstrates strong performance, especially when handling uncertainties and optimizing system operations against baseline controllers. Reinforcement Learning (RL) shows promise, especially when combined with MPC (RL-MPC), indicating a potential for adaptive and intelligent control systems. The studies emphasize the shift from traditional rule-based to advanced control strategies, highlighting the importance of realistic simulation environments for controller development and validation. Overall, advanced controllers, particularly those integrating MPC and RL, offer promising avenues for improving building energy management systems.

In the present work, a set of Model Predictive Control based controllers are conceptualized and tested on a high fidelity building emulator. The emulation model is provided by a publicly available platform that allows researchers to test innovative control strategies on a limited set of shared testcases. The main contributions of this work can be summarized as follows:

- The work, albeit simulated, aims at conceptualizing a methodology that could realistically be transferred to a real-life applications. Indeed, the choice of a high-fidelity emulator allows the consideration of control actions and measurement signals that are commonly found in ordinary HVAC systems. Furthermore, the non linear and dynamic behaviour of the energy system components are accounted for. The emulator accurately reproduces the actuators’ characteristics, the thermal inertia of the emission system, and the heat pump variable efficiency. As pointed out in [26], “*many load-based simulators compute the energy to be provided to meet a room temperature setpoint, then try to dispatch HVAC equipment at some fictitious part load operation that provides the required energy*”. In particular, the adopted emulator reproduces the coefficient of performance of the heat pump according to manufacturer’s data, along with the characteristics of the circulation pump and the evaporator’s fan. On the other hand, a real control system reads a feasibly measurable quantity, such as room temperature or air flow rate, and then computes a control signal for an actuator such as a compressor, a damper, or a fan variable speed drive. In this context, this work presents a methodology that, considering a wide range of problems that can be encountered in real life applications, addresses them despite being conducted in simulative fashion. For this reason, the choice of a standardized, well established emulation platform that is precisely intended to benchmark different control strategies with each other gives the obtained results both the reliability that comes from a realistic simulation and the possibility of being compared against future works that would employ the same test case. On top of that, the selected test case is representative of a widely adopted configuration, so that conclusions drawn can be extended to a large share of the existing building stock. In the authors’ opinion, this further renders the outcome of the present work more relevant in bridging the gap between research and the industrial-scale adoption of advanced control strategies.

The resulting MPC-based controllers do not represent a notable novelty in terms of their formulation or the methodology adopted for identifying their control-oriented models per se. However, this work contributes by demonstrating the capabilities of the BOPTTEST framework and, more generally, the use of detailed building emulators to test advanced controllers.

- The flexibility potential of a building energy system controlled by a predictive, model based optimal controller is assessed using a high fidelity emulation model. In buildings that do not have energy storage systems, the energy flexibility that they offer relies on the dynamic behaviour of the building. Therefore, assessing the ability of a predictive controller to exploit the flexibility of a building requires a simulation model that is detailed enough to provide a realistic dynamic response to control actions. The BOPTTEST platform provides such emulation capabilities. Flexibility is evaluated in a scenario of demand response, proving that the predictive capabilities of a properly formulated MPC can leverage the thermal mass of the building to shift energy usage in time. To the best of the authors' knowledge, this is the first instance of usage of this BOPTTEST test case to evaluate the performance of an advanced controller under a demand response event, with the purpose of evaluating the energy flexibility of the building. Additionally, the work provides an example of how that the tool can be adapted to more specific needs by considering an additional KPI and a boundary condition profile.

Finally, results are benchmarked against a standardized, reactive baseline controller, so that the performance of the proposed advanced controllers can be fairly compared to that of control strategies brought forth by other researchers in the field.

Table 1. Summary of case studies on building control.

Author (s)	Case Study Name	Controller Type	Controller Summary	Final Results
Arroyo, Manna, et al., 2022	Single-zone residential hydronic	RL vs. PI	RL algorithms outperform a baseline PI controller in operational cost and thermal discomfort.	RL achieves better trade-off between cost reduction and comfort levels.
Wang et al., 2023	BOPTTEST Hydronic Heat Pump	RL (DDPG, DDQN, SAC) vs. MPC	Among RL, only DDPG outperforms the baseline. MPC excels over all in typical and peak scenarios.	MPC is superior in optimizing actions and managing system state.
Arroyo, Spiessens, et al., 2022	BESTEST Hydronic Heat Pump	MPC, RL, RL-MPC	RL underperforms in complex environments. Hybrid RL-MPC shows improved performance.	Hybrid RL-MPC bridges the gap between model-free and model-based strategies.
Bünning et al., 2021	bestest-hydronic	ARX, Random Forest, Input Convex Neural Network	ARX model outperforms others in sample efficiency, predictive accuracy, and computational efficiency.	ARX model is better suited for standard residential building dynamics.
Blum et al., 2021	BESTEST Hydronic Heat Pump	MPC vs. PI	MPC significantly outperforms the PI controller in all pricing scenarios, optimizing comfort and costs.	MPC achieves up to 90.8% reduction in discomfort and 27.2% in costs.
Zanetti et al., 2023	Two-room apartment in Milan	Various MPC formulations (QP to MINLP)	All MPC strategies reduce discomfort, QP, and NLP are more efficient than MINLP.	Simpler linear constraints with nonlinear objectives provide a balanced approach.
Arroyo, Manna, et al., 2022	Two-room apartment in Milan	RL-MPC vs. MPC	RL-MPC integrates MPC predictability with RL adaptability. MINLP MPC is comparable to successful NLP MPCs.	RL-MPC outperforms traditional RL and delivers results akin to MPC in deterministic settings.
Maier et al., 2023	Two-Zone Apartment	AMPC vs. MPC vs. RBC	AMPC reduces operational costs by up to 33% and discomfort by 70%, less computational time than MPC.	While increasing thermal discomfort slightly, AMPC nearly matches MPC performance.
Marzullo et al., 2022	DOE's Reference Small Office Building	Rule-based controls, MPC, RL	ACTB integrates high-fidelity models with advanced controllers, moving from rule-based to advanced strategies.	ACTB offers realistic testing and potential cost reduction for advanced control strategies.
Gao et al., 2023	BESTEST Case 900	Tube-based MPC	Tube-based MPC effectively reduces operational costs by at least 24% while better managing indoor temperatures under uncertainties.	Superior performance in reducing operational costs and improving temperature control.

2. Methodology and Methods

2.1. Methodology

The purpose of this work is to prove the efficacy of a model-based optimal controller in handling multiple objectives. In particular, the work is divided into two stages:

- A first Model Predictive Control problem (MPC1) is formulated, with its objectives being the minimization of power drawn from the electric grid by the building energy system, while ensuring internal temperature constraints are respected. This MPC controller is tested with different combinations of control horizon lengths and control timesteps, which are thus treated as hyperparameters for the control problem.
- After the selection of the most fitting combination of the forementioned hyperparameters, a second MPC controller (MPC2) is formulated, with the additional objective of adhering to a maximum electric power usage restriction, to simulate a demand response grid requirement.

The methodology adopted consists of the following steps. Firstly, a simulated case study is selected. Simulation allows multiple scenarios, in terms of control problem, parameters, and objectives, to be tested under the same boundary conditions. The choice of test case was made from the BOPTTEST repository, since this emulation platform allows the performance of experiments that can be easily reproduced by other researchers. The case study is described in depth in Section 2.2. Secondly, a control oriented model is defined. Such a model is intended to inform the controller of the controlled system behaviour. In particular, two models are considered: a differential model for the building dynamics, and an algebraic model for the heat pump. The envelope model must account for the dynamical behaviour of the building elements and the air mass, which offer a source of energy flexibility due to their thermal energy storage potential. The modelling paradigm chosen to capture the building dynamics is the grey box modelling approach (Section 2.3). As to the heat pump model, its behaviour was considered static, that is, the relationship between the system input and output variables was considered to be purely algebraic. The same section details motivations behind the choices regarding the model order and the selection of the variables to use as predictors. Once the control oriented models are obtained, a first MPC problem (MPC1) is defined with the aim of minimizing the power drawn from the electric grid, while ensuring thermal comfort for the occupants. The objective of this control is twofold. The first objective is to minimize the comfort violations, that is, to keep the internal temperature inside a comfort band. The second objective is to minimize the amount of energy drawn from the grid. These objectives are, in general, contrasting, and therefore the problem can be considered as a multi-objective optimization problem. The adopted paradigm for the formulation is that of the economic MPC, since it does not focus on the attainment of a precise setpoint level. Indeed, inside the comfort band, all temperature values are equally admissible, and the ability to choose any temperature inside the comfort band is what gives the system the required flexibility. MPC1, identified by a control-oriented model and a cost function, is deployed on the emulator varying the control horizon (assumed as equal to the prediction horizon) and the control timestep. These quantities are therefore treated as hyperparameters for the control problem and their influence on the controller performance is analysed in depth. Moreover, three different pricing scenarios are tested for each combination of control timestep and horizon. These pricing scenarios, namely the *constant*, *dynamic*, and *highly dynamic* scenarios, are pre-determined pricing schedules provided by the BOPTTEST framework. MPC1 also helps in finding the best combination of hyperparameters to then be adopted in MPC2. Subsequently, the capability of the MPC approach to exploit the energy flexibility of buildings is tested by formulating a second optimal control problem (MPC2). In this case, a further restriction is imposed on the optimal problem in order to limit the power consumption of the HVAC system below a time varying signal coming from the grid. The formulation of this MPC keeps the economic-MPC structure of MPC1, but adds an additional soft constraint on the maximum power that can be drawn by the grid, in order to attain the goal set by the demand response grid signal. All simulations are run in the same period, in order to carry out a fair comparison

between all control solutions. Results are compared through a set of Key Performance Indicators (KPIs) and benchmarked against a baseline scenario, where control is performed by a traditional PI-based controller.

2.2. Case Study

The considered case study is a single-zone residential building located in Brussels, Belgium. The model represents a dwelling for a family of five, featuring a single zone with a rectangular floor measuring 12 by 16 m and a height of 2.7 m. Details on the envelope materials are reported in Table 2.

Table 2. Envelope materials and properties.

Envelope Element	Layer	Thickness [m]	Specific Thermal Capacity [J/(kgK)]	Thermal Conductivity [W/(mK)]	Density [kg/m ³]
External Walls	wood	0.009	900	0.14	530
	insulation	0.0615	1400	0.04	10
	concrete	0.1	1000	0.51	1400
Floor	concrete	0.15	840	1.4	2100
	insulation	0.2	1470	0.02	30
	screed	0.05	840	0.6	1100
	tile	0.01	840	1.4	2100
Roof	roof deck	0.019	900	0.14	530
	fiberglass	0.1118	840	0.04	12
	plasterboard	0.01	840	0.16	950

A deterministic, pre-defined occupancy schedule considers the zone as occupied by five people before 7 a.m. and after 8 p.m. during weekdays, and all day during weekends. No other internal heating loads are considered. The model utilizes a year's worth of weather data for Brussels, Belgium, to simulate realistic environmental conditions for testing and analysis.

The HVAC system consists of an air-to-water heat pump feeding a heating floor. The heat pump nominal capacity is 15 kW and is equipped with an evaporator fan. An ON/OFF controlled pump of 0.5 kg/s mass flow rate ensures the circulation of the heat carrying fluid (water) across the heating floor piping system. Both the evaporator fan and the circulation pump are ON when the heat pump is operating. The system layout is depicted in Figure 1.

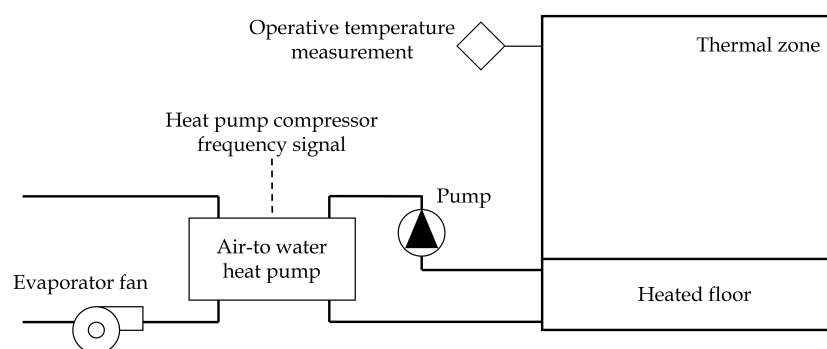


Figure 1. Schematic of the system, modified from [5].

The BOPTTEST test case provides a reactive set of controllers which ensure a robust operation of the system, while providing a baseline control logic against which innovative logics will be benchmarked. To control the indoor temperature profile and ensure comfort, a PI (Proportional–Integral) controller modulates the heat pump's output based on the zone's operative temperature. The controller adjusts the compressor frequency to maintain the desired temperature setpoint, which varies based on occupancy schedules. A second controller governs the operation of the fan and floor heating system pump, activating

them when the heat pump is modulating and deactivating them when it is not. As far as the model output, it must be noted that the emulator provides the measurement of the operative temperature: it will be implied henceforth that zone temperature refers to the operative temperature. The choice to provide the operative temperature rather than the dry bulb air temperature for the zone is motivated by the developers of the emulator because the emission system is a radiant system [5]. The building envelope, the HVAC systems along with all their components, and the low level controllers are modelled in the Modelica language [27], which is a well established modelling language that allows the modeling of physical components across different domains, as well as control logic units, in a causal fashion. In particular, the building model is based on the BESTEST case 900 from the IDEAS Library [28]. Details on the model and the computation of the operative temperature can be found in the documentation of the IDEAS Library. The model for the heat pump is the IDEAS.Fluid.HeatPumps.ScrollWaterToWater, which has been adapted to model an air-to-water heat pump instead. The heat pump model uses manufacturer data (Carrier air-to-water heat pump model 30AW015) to obtain a performance model, as per the calibration procedure explained in [29]. The circulation pump and the evaporator fan are modelled with components IDEAS.Fluid.Movers.FlowControlleddp of the same library. A two-week period representing peak heating demands from 17 January to 31 January are the subject of this study's analysis and will be henceforth referred to as the peak heating period. As far as energy tariffs, three different pricing models—constant, dynamic, and highly dynamic—are assessed at these intervals. While the dynamic model uses a variable rate of 0.2666 euros per kWh during peak hours (7:00 to 22:00) and 0.2383 euros per kWh during off-peak hours, the constant model uses a fixed rate of 0.2535 euros per kWh. The highly dynamic model uses 2019 Belgian day-ahead energy prices from the BELPEX wholesale electricity market.

The simulation software infrastructure is sketched in Figure 2. The main purpose of BOPTTEST is to enable the evaluation of the performance of control strategies, to facilitate benchmarking between different controllers. The main elements making up the platform are:

- A set of models, which provide the emulation capabilities of the software. BOPTTEST provides a publicly available repository of test cases for application in high fidelity models of buildings. Models are written in the Modelica language and are compiled into Functional Mockup Units to allow co-simulation from an external interface.
- A run-time environment, which manages the interaction between the emulator and an external code. This functionality allows controllers to be written in an external code, in the present instance a Python program, and to interact with the emulator at each timestep, to allow a co-simulation that recreates the actual deployment on a real building.
- A set of Key Performance Indicators, that due to being standardized, allow fair comparisons between different control solutions.

The model, written in Modelica, is compiled into a Functional Mockup Unit (FMU) and its interaction is managed by the BOPTTEST backend. Simulations are run from a master code based in python, interacting with the BOPTTEST interface with a local computing resource. BOPTTEST runs the simulation one timestep at a time, allowing the sending of control signals to the emulator. At the same timestep, the master code reads measurement signals (representing feedback) and obtains the disturbance variables prediction for the desired future horizon (which is necessary for the correct prediction). Finally, BOPTTEST provides the KPIs for the simulated time period. The MPC problem is formulated by means of the Python library `do_mpc` [30].




Figure 2. Simulation software infrastructure.

2.3. Control Oriented Model

As a necessary requirement for a model-based controller [16], a control-oriented model of the system was defined and calibrated, striking a balance between model complexity and prediction accuracy. The present section details the modeling assumptions adopted and presents the results of the calibration carried out.

The building thermal dynamics model is based on a energy balance equation (Equation (1)):

$$C \frac{dT_z}{dt} = -\frac{(T_z - T_{out})}{R} + A_w \dot{Q}_s + \dot{Q}_g + \dot{Q}_{HVAC} \quad (1)$$

where T_z is the zone temperature, T_{out} is the external air temperature, A_w is the window area, \dot{Q}_s is the global horizontal solar radiation, \dot{Q}_g is the power due to internal gains (namely occupants), and \dot{Q}_{HVAC} is the thermal power injected into the zone by the HVAC system. As far as internal gains, the BOPTTEST testcase documentation specifies that people occupancy is the only source of heat gains [31]. Temperatures are expressed in units of kelvin and thermal powers are expressed in watts, while solar radiation is expressed in $[W/m^2]$. R and C represent the envelope thermal resistance and the building thermal capacity, respectively. The resistance is expressed in $[K/W]$ as it accounts for the envelope surface, while the thermal capacity is expressed in $[J/K]$. The circuit equivalent circuit is represented in Figure 3, according to the electrical-thermal analogy.




Figure 3. Grey box model equivalent circuit.

A training and testing dataset were generated in order to tune the grey box model parameters and validate its accuracy on unseen data.

The specific focus of the data generation is the period leading up to the peak heat day scenario, spanning from 2 January to 16 January. This timeframe was selected to analyze

the system's performance prior to the highest demand conditions: the peak heating period from 16 January to 30 January. Choosing a training and testing dataset from a time window preceding the deployment period is necessary to ensure the fairness of the results, since training a model on the deployment period would reward the overfitting of the model itself.

The conditions under which the training and testing data are gathered are those of ordinary operation with temperature setpoints compatible with occupants' comfort. While some authors suggest training the model on data which explore the system under great variability of the monitored variables, often as a result of pseudo random control signals, others employ data generated under regular operation. Indeed, while simulation opens up the possibility to apply the former approach with ease, ensuring occupants' comfort makes the latter a much more realistic choice. Data for the training and validation of the model are plotted in Figure 4.




Figure 4. Training and validation period for grey box model parameters with baseline controller.

Upon completion of the optimization process, the model's predictions are compared against measured data, and various error metrics are computed to quantify its performance. The model reaches a Root Mean Square Error (RMSE) of 0.305 K and a Mean Absolute Percentage Error (MAPE) of 0.012. These indicators suggest a good prediction accuracy with respect to the output variable (zone air temperature).

Since the aim of the proposed controller is to minimize the final use of energy drawn from the power grid, or, in the case of flexibility event, optimize a function of it, a mathematical relationship between the thermal power provided by the heat pump and the corresponding electrical power must be established. This way, the final use of power at all instants can be expressed as a function of the control problem states and actions. While a heat pump's electrical power consumption is in general a function of source temperature, sink temperature, and partial load ratio [32], only two regressors were chosen to predict the thermal and electrical consumption. Namely, these are the heat pump control signal u_{hp} (controlling the compressor speed with a signal between 0 and 1) and the outdoor temperature T_{out} . These choices were the result of forward selection of candidate variables and polynomial orders. Ultimately, a first-order polynomial of these two variables was selected as it yielded satisfactory prediction results while ensuring the following advantages:

- The variables are easily available. Indeed, T_{out} is a disturbance whose prediction is provided by BOPTTEST and, in real-world application, is provided by forecast services. The heat pump control signal u_{hp} is a control action, therefore it is directly yielded by the controller itself.
- A linear formulation of the MPC problem is enabled. Non-linearity is avoided; in particular, the heat pump water production temperature, while being a good predictor,

forces a bilinear formulation of the optimal problem that slows down the optimization process.

The model is therefore in the following form:

$$\dot{Q}_{hp} = a_1 \cdot u_{hp} + b_1 \cdot T_{out} + c_1 \quad (2)$$

$$P_{el} = a_2 \cdot u_{hp} + b_2 \cdot T_{out} + c_2 \quad (3)$$

where:

- \dot{Q}_{hp} : heat pump power.
- P_{el} : electrical power consumption of the heat pump.
- u_{hp} : heat pump signal between 0 and 1.
- T_{out} : outdoor temperature.

The curve fitting was carried out by means of a polynomial regression on simulated data points. Data for the regression were generated by running the model with a random signal for the heat pump control in order to get a wide range of combinations of the regressors and the predicted variables. The final result of the process for the heat pump power is described in Table 3.

Table 3. Heat pump model regression coefficients.

Coefficient	Value
a_1	0.7489
b_1	1.1320
c_1	0.7489
a_2	0.7489
b_2	1.1320
c_2	0.7489

A scatter plot of the model results against the simulated data can be seen in Figures 5 and 6 for the thermal power yielded and the electrical consumption, respectively. The $\pm 20\%$ error lines in the scatter plots give a visual indication of the prediction accuracy.




Figure 5. Predicted vs. measured thermal power based on the heat pump signal and outdoor dry bulb temperature.




Figure 6. Predicted vs. measured electric power based on the heat pump signal and outdoor dry bulb temperature.

2.4. Model Predictive Control Formulation

The mathematical model describing the building's dynamics is provided by Equation (1). Here, the heating power \dot{Q}_{HVAC} is expressed as a function of the control signal u_{hp} through Equations (2) and (3).

The constraint on zone temperature is expressed according to the economic MPC paradigm, so that the controller can leverage the entire temperature range considered acceptable for occupants' comfort to fully exploit the flexibility offered by the envelope thermal mass. For this purpose, slacking variable s_{T_z} is defined so that the constraint on zone temperature becomes a soft constraint:

$$T_{z,min} - s_{T_z} \leq T_z \leq T_{z,max} + s_{T_z} \quad (4)$$

where $T_{z,min}$ is the lower temperature boundary, being 21 °C during occupied hours and 15 °C during unoccupied hours, and $T_{z,max}$ is the upper temperature boundary, being 24 °C during occupied hours and 30 °C during unoccupied hours. These values are provided by BOPTTEST and are not meant to be modified, so that all future and past works on the same test case can be fairly compared. The slacking variable itself thus becomes an additional control action for the optimization problem.

The heat pump is controlled by a continuous signal between 0 and 1, and the maximum electrical power is 4.5 kW:

$$\begin{aligned} 0 &\leq u_{hp} \leq 1 \\ 0 &\leq P_{el} \leq 4.5 \text{ kW} \end{aligned}$$

Overall, the time-varying variables involved in the control problem can be classified as:

- States: states describe the dynamic system conditions at all time instants. The one state in this problem is zone temperature $\underline{x} = [T_i]^T$, with $\underline{x} \in \mathbb{R}^1$.
- Actions: control actions are the heat pump control signal and the slacking variable, $\underline{u} = [u_{hp}, s_{T_z}]^T$, with $\underline{u} \in \mathbb{R}^2$
- Disturbances: the uncontrolled input variables are $\underline{d} = [T_{out}, \dot{Q}_s, \dot{Q}_g]^T$ (external air temperature, global horizontal solar radiation, internal heat gains), with $\underline{d} \in \mathbb{R}^3$

The objective of any MPC controller is encapsulated in its cost function. In this case, the aim is to optimize the heat pump's operation over a prediction horizon from instant t_i

to instant t_h . In the present work, two different control problems are formulated, namely MPC1 and MPC2. Both problems are formulated as economic MPC optimal problems.

MPC1 focuses on minimizing electrical power consumption P_{el} while maximizing thermal comfort. This is translated into the cost function J_1 :

$$J_1 = \int_{t=t_i}^{t_h} (w_{el}P_{el} + w_s s_{T_z}) dt \quad (5)$$

where:

- w_{el} represents the cost coefficient related to electrical power consumption.
- w_s stands for the weight associated with maintaining the indoor temperature T_z within a comfortable range. In particular, it weights slacking variable s_{T_z} .

Controller MPC2 was conceived in order to assess the efficacy of the MPC controller in leveraging energy flexibility. A time-varying constraint on the heat pump maximum power consumption was introduced to limit the power drawn from electric grid during peak hours. This constraint was formulated as a soft constraint in order to allow the system to violate the boundary if the thermal load required it. For this purpose, a new slacking variable s_{hp} was introduced:

$$P_{el} \leq P_{hp,max} + s_{hp}$$

The slacking variable was added to cost function J_1 , multiplied by a properly selected weighting factor:

$$J_2 = \int_{t=t_i}^{t_h} (w_{el}P_{el} + w_s s_{T_z} + w_{hp} s_{hp}) dt \quad (6)$$

where:

- w_{el} represents the cost coefficients related to electrical energy.
- w_s is the weight associated with maintaining the indoor temperature T_z within a comfortable range.
- w_{hp} is the weight associated with limiting energy consumption during on-peak hours.

The initial step involved understanding the power grid's typical requirements specific to the building's location, situated in Belgium. Subsequently, the net energy demand and associated pricing profiles were derived, drawing insights from the literature [33], as depicted in Figure 7.




Figure 7. Net demand energy with consideration of the different scenarios [33].

Drawing from data related to the Belgian grid and theoretical conjectures concerning the building, specific energy boundaries were delineated to prescribe energy constraints

for the case study. As illustrated in Figure 8, accounting for the maximum potential energy consumption of the heat pump at 15 kW, these recommended energy boundaries were established using logical coefficients derived from the relevant literature [33].




Figure 8. Net energy demand of the case study modified from [33].

2.5. Key Performance Indicators

The BOPTTEST platform yields a set of KPIs at each run. These KPIs are intended to enable all researchers to compare results using the same metrics. Among the KPIs offered by the platform, two were selected as significant for the present study. Further details on the available KPIs can be found in [5].

- Thermal discomfort is expressed in units of [Kh] and it is computed as the cumulative deviation of the zone temperature from the lower and the upper boundary comfort limits. These limits are pre-defined for each testcase, so that future results on the same testcase will be fairly comparable. The thermal discomfort is computed as:

$$K_{tdis} = \frac{\sum_{z=1}^N \int_{t_s}^{t_f} [(\max(T_z(t) - T_{z,coo}(t)), 0) + (\max(T_{z,hea}(t) - T_z(t)), 0)] dt}{N} \quad (7)$$

where $T(z)$ is the temperature of zone z , $T_{z,hea}$ and $T_{z,coo}$ are the lower and upper setpoints respectively (all temperatures in kelvin), t_s and t_f are the start and end time of the simulation, and N is the number of zones. The purpose of this KPI is to quantify how well the control system is able to keep the temperature between the set boundaries. In particular, it accounts for both the magnitude of the violation and the amount of time the temperature is outside the boundaries.

- Cost is expressed in units of [€/m²] and it is computed as:

$$K_{cost} = \frac{\sum_{j \in J} \sum_{i \in I} \int_{t_s}^{t_f} p_j^\tau(t) P_{ij}(t) dt}{A} \quad (8)$$

where $P_{ij}(t)$ is the power consumed by device i with energy source j (in units of watts), J is the set of the energy sources available, I_j is the set of the pieces of equipment using source j , $p_j^\tau(t)$ is the price of energy source j with tariff τ , and A is the building total floor area (in units of m²).

Additionally, two further KPIs are computed for MPC2 results in order to better assess the performance with respect to its specific task, namely that of exploiting the building thermal mass to increase its energy flexibility. These KPIs are automatically computed by the BOPTTEST platform, but are additions made by the authors. The two KPIs are defined as follows.

- In order to assess the ability of MPC2 to adhere to the power usage constraint, the violations of the constraint itself are computed as:

$$P_{viol} = \int_{t_s}^{t_f} [\max(P_{el}(t) - P_{hp,max}(t), 0)] dt \quad (9)$$

where $P_{el}(z)$ is the actual power consumption of the heat pump, and $P_{hp,max}(z)$ is the upper bounding constraint for its usage. The KPI is expressed in unit of [kWh].

- The Flexibility Factor indicator, expressed as dimensionless, is a measure of the ability of the system to shift energy usage from periods where usage is intended to be lowered to period of higher convenience. The definition here adopted is a slight modification to the indicator commonly employed in the literature (see [34]). Indeed, the non-peak period ($P_{el,nonpeak}$ in Equation (10)) power consumption refers to the power used when the power consumption limit is at the highest value; conversely, $P_{el,nonpeak}$ refers to the power consumed when power reduction is required by the grid.

$$FF = \frac{\int_{t_s}^{t_f} P_{el,nonpeak}(t)dt - \int_{t_s}^{t_f} P_{el,peak}(t)dt}{\int_{t_s}^{t_f} P_{el,nonpeak}(t)dt + \int_{t_s}^{t_f} P_{el,peak}(t)dt} \quad (10)$$

3. Results

This section presents the outcomes of the Model Predictive Control (MPC) implementation, along with a thorough analysis of various scenarios relative to MPC1, evaluated against the Key Performance Indicators (KPIs) supplied by the Building Optimization Performance Testing Framework (BOPTTEST). The effectiveness of MPC1 is validated through these results, identifying the most successful combinations of control horizon and timestep for the particular case at hand. Finally, validation results related to the effectiveness of the MPC2 problem for the energy flexibility are provided.

Based on Table 4, which details various performance metrics across different MPC1 hyperparameter combinations, a comparison can be drawn between the advanced controller and the baseline controller performances. The lowest thermal discomfort is achieved with an 8 h horizon and 300 s time steps in the highly dynamic electricity price scenario. When evaluating both cost and thermal discomfort together, the 8 h horizon and 300 s time steps in the highly dynamic electricity price scenario offer the best balance. The percentage improvement over the baseline for combined cost and comfort efficiency would require a multi-objective analysis, considering both factors simultaneously. Results are reported in Table 4 as percentages of improvement with respect to the baseline conditions. All MPC1 scenarios reduce thermal discomfort, with the most significant improvement in the 8 h horizon and 300 s time steps scenario (82.2% reduction). Cost is lower in all MPC1 scenarios, with the 8 h horizon and 600 s time steps showing the greatest reduction (20.3%). Energy use is reduced in all scenarios, with the 8 h horizon and 600 s time steps using the least energy (20.3% less). In summary, the best scenario varies for each KPI. For thermal discomfort and energy use, the 8 h horizon with 600 s time steps is most effective.

Figure 9 showcases the MPC controller's outcomes during a peak heating period under a constant pricing scenario, with a time horizon (T_h) of 8 h and a simulation time step of 5 min. Results show that MPC1 tends to keep the internal temperature close to the lower boundary, as that is the solution guaranteeing the lowest energy demand. The system reduces energy consumption during periods of inactivity and leverages predictive data to precondition the building for occupancy, ensuring thermal comfort upon arrival. An abrupt temperature spike on 19 January is observed, which can be attributed to significant fluctuations in outdoor temperature, solar radiation, and internal heat gains. When compared to the baseline PI controller, the MPC shows a better ability to anticipate setpoint changes that correspond to the beginning of occupation time windows. Similarly, the MPC allows internal temperature to decrease during unoccupied hours more so than the baseline controller does, which results in a lower energy demand over the entire period. Indeed, the prediction capabilities of the MPC controller allow it to anticipate setpoint changes. This also justifies the choice of an economic formulation, since a tracking MPC would not allow the leveraging of the permitted temperature range to find an optimal temperature profile.



Figure 9. Simulation results for the peak heating period with constant pricing scenario. $T_h = 8$ h and time step = 5 min.

Table 4. Cost and thermal discomfort improvements under different scenarios compared to the baseline controller. Results are expressed as percentage reduction with respect to the baseline.

Horizon (H)	Time Step (S)	Scenario Price	Cost [EUR/m ²] (%)	Thermal Discomfort [Kh/Zone] (%)
8	300	constant	15.8%	82.2%
12	300	constant	16.5%	81.2%
24	300	constant	17.3%	80.9%
8	600	constant	19.8%	66.9%
12	600	constant	18.6%	77.6%
24	600	constant	18.8%	79.7%
8	300	dynamic	17.5%	80.9%
12	300	dynamic	16.3%	81.2%
24	300	dynamic	17.1%	80.9%
8	600	dynamic	20.0%	66.9%
12	600	dynamic	18.6%	77.6%
24	600	dynamic	18.9%	79.7%
8	300	highly dynamic	17.9%	80.9%
12	300	highly dynamic	16.7%	81.2%
24	300	highly dynamic	17.5%	80.9%
8	600	highly dynamic	20.1%	66.9%
12	600	highly dynamic	18.9%	77.6%
24	600	highly dynamic	19.1%	79.7%

Analysis of the Effectiveness of the Designed MPC for Demand Response in the Belgian Grid

MPC2 was implemented with a control horizon of 8 h and a control timestep of 5 min. Simulations were conducted for the peak heat day period, that is, from 16 January to 30 January. The results are illustrated in Figure 10. The green dashed line represents the upper boundary on power consumption for the HVAC system, considering the on-peak and off-peak hours of the grid. The results demonstrate the effectiveness of the MPC in enhancing energy flexibility, while also maintaining thermal comfort. The Key Performance Indicators (KPIs) in Table 5 quantify the improvements.

Table 5. KPI metrics comparison between baseline controller and MPC2.

KPI	Cost (EUR/m ²)	Energy (kWh/m ²)	Thermal Discomfort (Kh/Zone)
Baseline	0.87	3.33	7.35
MPC2	0.79	3.13	3.79




Figure 10. Demand response validation result of the MPC.

The MPC showed a significant improvement in thermal discomfort, with a reduction of approximately 48%. Additionally, there was a 6% improvement in both energy consumption and cost compared to the baseline controller. The final figure confirms that energy consumption remains within the boundaries set by the grid's energy limitations during both on-peak and off-peak hours. Such improvements were reached despite the stringent constraints on the maximum power consumption allowed. In order to satisfy both the consumption and the thermal comfort constraints, MPC2 exploited the dynamic behaviour of the building, that is, its ability to store internal energy. This was achieved by leveraging the admissible range for the internal temperature. Indeed, as Figure 10 shows, the controller drives the internal temperature between the lower and the upper boundary values, thus exploiting the building thermal capacity to store energy when a higher power could be drawn from the grid, and to release it when the power constraint forced the controller to reduce the HVAC instant consumption. MPC2 managed to keep the cumulated violation of the power consumption constraint to 12.6 kWh, computed as per Equation (9). As to the Flexibility Factor defined in Equation (10), the controller reached a value of 0.49, demonstrating a high capability of MPC to shift energy consumption to specific time windows of operation, while significantly reducing the power usage when grid constraint become more stringent.

Overall, these results validate the effectiveness of the designed MPC in meeting demand response objectives while maintaining occupant comfort and reducing energy costs and consumption.

4. Discussion and Conclusions

In this work, the potential of Model Predictive Control for the management of building energy systems was explored. Controller MPC1 was conceived with the purpose of minimizing the energy cost over a testing period, while ensuring thermal comfort conditions for the occupants. For MPC1, energy price was the sole source of interaction with the power grid, and it was treated as a boundary condition influencing the value of the optimal problem cost function. The implementation can be considered successful as MPC1 outperformed the PI-based baseline controller in terms of both cost savings and thermal comfort violations. As a variation of controller MPC1, MPC2 included in its formulation a constraint on the maximum power to draw from the grid by the HVAC system. Despite the addition of another constraint, the controller was still able to outperform the baseline controller in terms of cost and thermal comfort. The latter controller proved MPC as a viable solution to include grid signals into the design of a controller, as it successfully leveraged the thermal mass of the building as a means to reshape the power consumption profile to adapt it to grid requirements. Even though the present research was carried out in a simulated environment, the proposed methodology was carefully conceptualized so that

it could be easily implemented in the field, bearing in mind the associated technological and practical limitations. In particular, the following aspects can be underlined:

- The emulator on which the scenarios were tested is made up of high fidelity models, so the response of the system to the control input signals is representative of the dynamic response of an actual, real-world building.
- The level of detail of the emulator allows the consideration of feasible and realistic control actions and feedback signals in the design and testing of the controller, since the modeled actuators and sensors mirror those that are commonly found on real energy systems.
- Model training has been identified by the literature as one of the main challenges in designing a model-based controller. The methodology here adopted allowed us to obtain a plant model reliable enough for the MPC to perform satisfactorily while employing a small amount of easily obtainable data.
- The low order of the control oriented model, along with the strictly linear formulation of the control problem, guarantee a low computational burden in the solution of the optimal problem.

The aforementioned points make the proposed methodology suitable for in-field application; moreover, the results of the present work can be considered realistic in terms of the numerical assessment of the controller performance and its improvement with respect to the baseline. However, a few simplifying assumptions could be removed in future works, and some further aspects could be investigated.

- Uncontrolled disturbances and boundary conditions were here considered as known for a future window in a deterministic way; future work might investigate the effect of stochastic predictions on such variables and their influence on the MPC performance.
- Sensors in the emulator yield measurements that are unaffected by any uncertainty; the effect of non-ideal sensors might be subject to further analysis.

Author Contributions: Conceptualization, D.F. and A.C.; methodology, D.F.; software, A.R.Y. and D.F.; validation, A.C.; formal analysis, A.C.; investigation, D.F. and A.R.Y.; resources, A.C.; data curation, A.R.Y. and D.F.; writing—original draft preparation, D.F. and A.R.Y.; writing—review and editing, D.F., A.R.Y. and A.C.; visualization, A.R.Y. and D.F.; supervision, A.C.; project administration, A.C. All authors have read and agreed to the published version of the manuscript.

Funding: The present work was carried out within the PhD program of Davide Fop, financially supported by Comitato ICT.

Data Availability Statement: The original contributions presented in the study are included in the article, further inquiries can be directed to the corresponding author.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

MPC	Model Predictive Control
BOPTTEST	Building Optimization Testing framework
HVAC	Heating, Ventilation, and Air Conditioning
TCS	Traditional Control Strategies
ACS	Advanced Control Strategies
PID	Proportional Integral Derivative control
KPI	Key Performance Indicator
RL	Reinforcement Learning
DDPG	Deep Deterministic Policy Gradient
DDQN	Double Deep Q Network
SAC	Soft Actor Critic

ARX	Auto Regressive Exogenous model
QP	Quadratic Programming
MINLP	Mixed Integer Non Linear Program
AMPC	Approximating Model Predictive Control
NRMSE	Normalized Root Mean Square Error
MAPE	Mean Absolute Percentage Error

References

1. United Nations Environment Programme. *2021 Global Status Report for Buildings and Construction: Towards a Zero-Emission, Efficient and Resilient Buildings and Construction Sector*; United Nations Environment Programme: Nairobi, Kenya, 2021.
2. United Nations Environment Programme. *2022 Global Status Report for Buildings and Construction: Towards a Zero-Emission, Efficient and Resilient Buildings and Construction Sector*; United Nations Environment Programme: Nairobi, Kenya, 2022.
3. Wagner, A.; Ruellan, M.; Heinrich, M.; Bourdais, R. Towards a classification of the building stock in terms of potential energy savings through advanced control. *Energy Build.* **2024**, *310*, 114050. [\[CrossRef\]](#)
4. Arroyo, J.; Manna, C.; Spiessens, F.; Helsen, L. An Open-AI gym environment for the building optimization testing (BOPTTEST) framework. In *Proceedings of the Building Simulation 2021, 17th Conference of IBPSA, Bruges, Belgium, 1–3 September 2021; Volume 17*, pp. 175–182.
5. Blum, D.; Arroyo, J.; Huang, S.; Drgoňa, J.; Jorissen, F.; Walnum, H.T.; Chen, Y.; Benne, K.; Vrabie, D.; Wetter, M.; et al. Building optimization testing framework (BOPTTEST) for simulation-based benchmarking of control strategies in buildings. *J. Build. Perform. Simul.* **2021**, *14*, 586–610. [\[CrossRef\]](#)
6. Arroyo, J.; Spiessens, F.; Helsen, L. Comparison of optimal control techniques for building energy management. *Front. Built Environ.* **2022**, *8*, 849754. [\[CrossRef\]](#)
7. Gholamzadehmir, M.; Del Pero, C.; Buffa, S.; Fedrizzi, R.; Aste, N. Adaptive-predictive control strategy for HVAC systems in smart buildings—A review. *Sustain. Cities Soc.* **2020**, *63*, 102480. [\[CrossRef\]](#)
8. Wang, D.; Zheng, W.; Wang, Z.; Wang, Y.; Pang, X.; Wang, W. Comparison of reinforcement learning and model predictive control for building energy system optimization. *Appl. Therm. Eng.* **2023**, *228*, 120430. [\[CrossRef\]](#)
9. Lee, Y.M.; Horesh, R.; Liberti, L. Optimal HVAC control as demand response with on-site energy storage and generation system. *Energy Procedia* **2015**, *78*, 2106–2111. [\[CrossRef\]](#)
10. Afram, A.; Janabi-Sharifi, F.; Fung, A.S.; Raahemifar, K. Artificial neural network (ANN) based model predictive control (MPC) and optimization of HVAC systems: A state of the art review and case study of a residential HVAC system. *Energy Build.* **2017**, *141*, 96–113. [\[CrossRef\]](#)
11. del Mar Castilla, M.; Álvarez, J.D.; Rodríguez, F.; Berenguel, M. *Comfort Control in Buildings*; Springer: Berlin/Heidelberg, Germany, 2014.
12. Aswani, A.; Master, N.; Taneja, J.; Krioukov, A.; Culler, D.; Tomlin, C. Energy-efficient building HVAC control using hybrid system LBMPC. *IFAC Proc. Vol.* **2012**, *45*, 496–501. [\[CrossRef\]](#)
13. van de Bree, A.; von Manteuffel, B.; Ramaekers, L.; Offermann, M. Role of Building Automation Related to Renewable Energy in nZEB's. Project Report. 2014. Available online: <https://www.slideshare.net/slideshow/buin13782-report-final/38815644> (accessed on 1 September 2024).
14. American Society of Heating, Refrigerating and Air-Conditioning Engineers (ASHRAE). *Composite Index to the 2015 HVAC Applications, 2016 HVAC SYSTEMS and Equipment, 2017 Fundamentals, and 2018 Refrigeration Volumes*; ASHRAE: Peachtree Corners, GA, USA, 2018.
15. Vázquez-Canteli, J.R.; Nagy, Z. Reinforcement learning for demand response: A review of algorithms and modeling techniques. *Appl. Energy* **2019**, *235*, 1072–1089. [\[CrossRef\]](#)
16. Drgoňa, J.; Arroyo, J.; Figueroa, I.C.; Blum, D.; Arendt, K.; Kim, D.; Ollé, E.P.; Oravec, J.; Wetter, M.; Vrabie, D.L.; et al. All you need to know about model predictive control for buildings. *Annu. Rev. Control* **2020**, *50*, 190–232. [\[CrossRef\]](#)
17. Kathirgamanathan, A.; De Rosa, M.; Mangina, E.; Finn, D.P. Data-driven predictive control for unlocking building energy flexibility: A review. *Renew. Sustain. Energy Rev.* **2021**, *135*, 110120. [\[CrossRef\]](#)
18. Rockett, P.; Hathway, E.A. Model-predictive control for non-domestic buildings: A critical review and prospects. *Build. Res. Inf.* **2017**, *45*, 556–571. [\[CrossRef\]](#)
19. Arroyo, J.; Manna, C.; Spiessens, F.; Helsen, L. Reinforced model predictive control (RL-MPC) for building energy management. *Appl. Energy* **2022**, *309*, 118346. [\[CrossRef\]](#)
20. Serale, G.; Fiorentini, M.; Capozzoli, A.; Bernardini, D.; Bemporad, A. Model predictive control (MPC) for enhancing building and HVAC system energy efficiency: Problem formulation, applications and opportunities. *Energies* **2018**, *11*, 631. [\[CrossRef\]](#)
21. Pantazaras, A.; Lee, S.E.; Santamouris, M.; Yang, J. Predicting the CO₂ levels in buildings using deterministic and identified models. *Energy Build.* **2016**, *127*, 774–785. [\[CrossRef\]](#)
22. 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*, 372–385. [\[CrossRef\]](#)
23. Blum, D.; Wang, Z.; Weyandt, C.; Kim, D.; Wetter, M.; Hong, T.; Piette, M.A. Field demonstration and implementation analysis of model predictive control in an office HVAC system. *Appl. Energy* **2022**, *318*, 119104. [\[CrossRef\]](#)

24. Freund, S.; Schmitz, G. Implementation of model predictive control in a large-sized, low-energy office building. *Build. Environ.* **2021**, *197*, 107830. [[CrossRef](#)]
25. Thieblemont, H.; Haghighat, F.; Ooka, R.; Moreau, A. Predictive control strategies based on weather forecast in buildings with energy storage system: A review of the state-of-the art. *Energy Build.* **2017**, *153*, 485–500. [[CrossRef](#)]
26. Wetter, M.; Sulzer, M. A call to action for building energy system modelling in the age of decarbonization. *J. Build. Perform. Simul.* **2024**, *17*, 383–393. [[CrossRef](#)]
27. Mattsson, S.E.; Elmqvist, H. Modelica—An International Effort to Design the Next Generation Modeling Language. *IFAC Proc. Vol.* **1997**, *30*, 151–155. [[CrossRef](#)]
28. Jorissen, F.; Reynders, G.; Baetens, R.; Picard, D.; Saelens, D.; Helsen, L. Implementation and Verification of the IDEAS Building Energy Simulation Library. *J. Build. Perform. Simul.* **2018**, *11*, 669–688. [[CrossRef](#)]
29. Cimmino, M.; Wetter, M. Modelling of Heat Pumps with Calibrated Parameters Based on Manufacturer Data. In Proceedings of the 12th International Modelica Conference, Prague, Czech Republic, 15–17 May 2017; pp. 219–226.
30. Fiedler, F.; Karg, B.; Lüken, L.; Brandner, D.; Heinlein, M.; Brabender, F.; Lucia, S. do-mpc: Towards FAIR nonlinear and robust model predictive control. *Control Eng. Pract.* **2023**, *140*, 105676. [[CrossRef](#)]
31. Available online: https://ibpsa.github.io/project1-boptest/docs-testcases/bestest_hydrionic_heat_pump/index.html (accessed on 3 October 2024).
32. Mugnini, A.; Evens, M.; Arteconi, A. Model predictive controls for residential buildings with heat pumps: Experimentally validated archetypes to simplify the large-scale application. *Energy Build.* **2024**, *320*, 114632. [[CrossRef](#)]
33. Sheha, M.; Powell, K. Using real-time electricity prices to leverage electrical energy storage and flexible loads in a smart grid environment utilizing machine learning techniques. *Processes* **2019**, *7*, 870. [[CrossRef](#)]
34. Li, H.; Johra, H.; de Andrade Pereira, F.; Hong, T.; Le Dréau, J.; Maturo, A.; Wei, M.; Liu, Y.; Saberi-Derakhtenjani, A.; Nagy, Z.; et al. Data-driven key performance indicators and datasets for building energy flexibility: A review and perspectives. *Appl. Energy* **2023**, *343*, 121217. [[CrossRef](#)]

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.