ANNEX 81

EBC
Energy in Building and Communities Programme

International Energy Agency

Subtask B Report

IEA EBC Annex 81 - Data-Driven Smart Buildings

Energy in Buildings and Communities Technology Collaboration Programme

June 2025

Technology Collaboration Programme
by iea

ANNEX 81

EBC
Energy in Building and Communities Programme

International Energy Agency

Subtask B Report

IEA EBC Annex 81 - Data-Driven Smart Buildings

Energy in Buildings and Communities Technology Collaboration Programme

June 2025

Authors

José A. Candanedo, Université de Sherbrooke, Canada
Igor Sartori, SINTEF, Norway
David Blum, Lawrence Berkeley National Laboratory, USA
Henrik Madsen, Technical University of Denmark (DTU), Denmark
Christian Thilker, Technical University of Denmark (DTU), Denmark
Zhe Wang, The Hong Kong University of Science and Technology, Hong Kong SAR
Harald Taxt Walnum, SINTEF Community, Norway
Yuan Gao, Kyushu University, Japan
Zheng O'Neill, Texas A&M University, USA
Saman Mostafavi, SRI, USA

Published by CSIRO, PO Box 330, Newcastle, NSW 2300, Australia

ISBN (13-digit) … [*to be arranged by Annex XX Operating Agent organisation*]

Participating countries in the EBC TCP: Australia, Austria, Belgium, Brazil, Canada, P.R. China, Czech Republic, Denmark, Finland, France, Germany, Ireland, Italy, Japan, Republic of Korea, the Netherlands, New Zealand, Norway, Portugal, Singapore, Spain, Sweden, Switzerland, Turkey, United Kingdom and the United States of America.

Additional copies of this report may be obtained from: EBC Executive Committee Support Services Unit (ESSU), C/o AECOM Ltd, The Colmore Building, Colmore Circus Queensway, Birmingham B4 6AT, United Kingdom
www.iea-ebc.org
essu@iea-ebc.org

Preface

The International Energy Agency

The International Energy Agency (IEA) was established in 1974 within the framework of the Organisation for Economic Co-operation and Development (OECD) to implement an international energy programme. A basic aim of the IEA is to foster international co-operation among the 30 IEA participating countries and to increase energy security through energy research, development and demonstration in the fields of technologies for energy efficiency and renewable energy sources.

The IEA Energy in Buildings and Communities Programme

The IEA co-ordinates international energy research and development (R&D) activities through a comprehensive portfolio of Technology Collaboration Programmes (TCPs). The mission of the IEA Energy in Buildings and Communities (IEA EBC) TCP is to support the acceleration of the transformation of the built environment towards more energy efficient and sustainable buildings and communities, by the development and dissemination of knowledge, technologies and processes and other solutions through international collaborative research and open innovation. (Until 2013, the IEA EBC Programme was known as the IEA Energy Conservation in Buildings and Community Systems Programme, ECBCS.)

The high priority research themes in the EBC Strategic Plan 2019-2024 are based on research drivers, national programmes within the EBC participating countries, the Future Buildings Forum (FBF) Think Tank Workshop held in Singapore in October 2017 and a Strategy Planning Workshop held at the EBC Executive Committee Meeting in November 2017. The research themes represent a collective input of the Executive Committee members and Operating Agents to exploit technological and other opportunities to save energy in the buildings sector, and to remove technical obstacles to market penetration of new energy technologies, systems and processes. Future EBC collaborative research and innovation work should have its focus on these themes.

At the Strategy Planning Workshop in 2017, some 40 research themes were developed. From those 40 themes, 10 themes of special high priority have been extracted, taking into consideration a score that was given to each theme at the workshop. The 10 high-priority themes can be separated into two types: 'Objectives' and 'Means'. These two groups are distinguished to better understand the different themes.

Objectives - The strategic objectives of the EBC TCP are as follows:

– reinforcing the technical and economic basis for refurbishment of existing buildings, including financing, engagement of stakeholders and promotion of co-benefits;
– improvement of planning, construction and management processes to reduce the performance gap between design stage assessments and real-world operation;
– the creation of 'low tech', robust and affordable technologies;
– the further development of energy efficient cooling in hot and humid, or dry climates, avoiding mechanical cooling if possible;
– the creation of holistic solution sets for district level systems taking into account energy grids, overall performance, business models, engagement of stakeholders, and transport energy system implications.

Means - The strategic objectives of the EBC TCP will be achieved by the means listed below:

– the creation of tools for supporting design and construction through to operations and maintenance, including building energy standards and life cycle analysis (LCA);
– benefitting from 'living labs' to provide experience of and overcome barriers to adoption of energy efficiency measures;
– improving smart control of building services technical installations, including occupant and operator interfaces;
– addressing data issues in buildings, including non-intrusive and secure data collection;
– the development of building information modelling (BIM) as a game changer, from design and construction through to operations and maintenance.

The themes in both groups can be the subject for new Annexes, but what distinguishes them is that the 'objectives' themes are final goals or solutions (or part of) for an energy efficient built environment, while the 'means' themes are instruments or enablers to reach such a goal. These themes are explained in more detail in the EBC Strategic Plan 2019-2024.

The Executive Committee

Overall control of the IEA EBC Programme is maintained by an Executive Committee, which not only monitors existing projects, but also identifies new strategic areas in which collaborative efforts may be beneficial. As the Programme is based on a contract with the IEA, the projects are legally established as Annexes to the IEA EBC Implementing Agreement. At the present time, the following projects have been initiated by the IEA EBC Executive Committee, with completed projects identified by (*) and joint projects with the IEA Solar Heating and Cooling Technology Collaboration Programme by (☼):

Annex 1: Load Energy Determination of Buildings (*)
Annex 2: Ekistics and Advanced Community Energy Systems (*)
Annex 3: Energy Conservation in Residential Buildings (*)
Annex 4: Glasgow Commercial Building Monitoring (*)
Annex 5: Air Infiltration and Ventilation Centre
Annex 6: Energy Systems and Design of Communities (*)
Annex 7: Local Government Energy Planning (*)
Annex 8: Inhabitants Behaviour with Regard to Ventilation (*)
Annex 9: Minimum Ventilation Rates (*)
Annex 10: Building HVAC System Simulation (*)
Annex 11: Energy Auditing (*)
Annex 12: Windows and Fenestration (*)
Annex 13: Energy Management in Hospitals (*)
Annex 14: Condensation and Energy (*)
Annex 15: Energy Efficiency in Schools (*)
Annex 16: BEMS 1- User Interfaces and System Integration (*)
Annex 17: BEMS 2- Evaluation and Emulation Techniques (*)
Annex 18: Demand Controlled Ventilation Systems (*)
Annex 19: Low Slope Roof Systems (*)
Annex 20: Air Flow Patterns within Buildings (*)
Annex 21: Thermal Modelling (*)
Annex 22: Energy Efficient Communities (*)
Annex 23: Multi Zone Air Flow Modelling (COMIS) (*)
Annex 24: Heat, Air and Moisture Transfer in Envelopes (*)
Annex 25: Real time HVAC Simulation (*)
Annex 26: Energy Efficient Ventilation of Large Enclosures (*)
Annex 27: Evaluation and Demonstration of Domestic Ventilation Systems (*)
Annex 28: Low Energy Cooling Systems (*)
Annex 29: ☼ Daylight in Buildings (*)
Annex 30: Bringing Simulation to Application (*)
Annex 31: Energy-Related Environmental Impact of Buildings (*)
Annex 32: Integral Building Envelope Performance Assessment (*)
Annex 33: Advanced Local Energy Planning (*)
Annex 34: Computer-Aided Evaluation of HVAC System Performance (*)
Annex 35: Design of Energy Efficient Hybrid Ventilation (HYBVENT) (*)
Annex 36: Retrofitting of Educational Buildings (*)
Annex 37: Low Exergy Systems for Heating and Cooling of Buildings (LowEx) (*)
Annex 38: ☼ Solar Sustainable Housing (*)
Annex 39: High Performance Insulation Systems (*)
Annex 40: Building Commissioning to Improve Energy Performance (*)
Annex 41: Whole Building Heat, Air and Moisture Response (MOIST-ENG) (*)
Annex 42: The Simulation of Building-Integrated Fuel Cell and Other Cogeneration Systems (FC+COGEN-SIM) (*)
Annex 43: ☼ Testing and Validation of Building Energy Simulation Tools (*)
Annex 44: Integrating Environmentally Responsive Elements in Buildings (*)
Annex 45: Energy Efficient Electric Lighting for Buildings (*)
Annex 46: Holistic Assessment Tool-kit on Energy Efficient Retrofit Measures for Government Buildings (EnERGo) (*)
Annex 47: Cost-Effective Commissioning for Existing and Low Energy Buildings (*)
Annex 48: Heat Pumping and Reversible Air Conditioning (*)
Annex 49: Low Exergy Systems for High Performance Buildings and Communities (*)
Annex 50: Prefabricated Systems for Low Energy Renovation of Residential Buildings (*)
Annex 51: Energy Efficient Communities (*)
Annex 52: ☼ Towards Net Zero Energy Solar Buildings (*)
Annex 53: Total Energy Use in Buildings: Analysis and Evaluation Methods (*)
Annex 54: Integration of Micro-Generation and Related Energy Technologies in Buildings (*)
Annex 55: Reliability of Energy Efficient Building Retrofitting - Probability Assessment of Performance and Cost (RAP-RETRO) (*)
Annex 56: Cost Effective Energy and CO2 Emissions Optimisation in Building Renovation (*)
Annex 57: Evaluation of Embodied Energy and CO2 Equivalent Emissions for Building Construction (*)

Annex 58: Reliable Building Energy Performance Characterisation Based on Full Scale Dynamic Measurements (*)
Annex 59: High Temperature Cooling and Low Temperature Heating in Buildings (*)
Annex 60: New Generation Computational Tools for Building and Community Energy Systems (*)
Annex 61: Business and Technical Concepts for Deep Energy Retrofit of Public Buildings (*)
Annex 62: Ventilative Cooling (*)
Annex 63: Implementation of Energy Strategies in Communities (*)
Annex 64: LowEx Communities - Optimised Performance of Energy Supply Systems with Exergy Principles (*)
Annex 65: Long-Term Performance of Super-Insulating Materials in Building Components and Systems (*)
Annex 66: Definition and Simulation of Occupant Behavior in Buildings (*)
Annex 67: Energy Flexible Buildings (*)
Annex 68: Indoor Air Quality Design and Control in Low Energy Residential Buildings (*)
Annex 69: Strategy and Practice of Adaptive Thermal Comfort in Low Energy Buildings (*)
Annex 70: Energy Epidemiology: Analysis of Real Building Energy Use at Scale (*)
Annex 71: Building Energy Performance Assessment Based on In-situ Measurements (*)
Annex 72: Assessing Life Cycle Related Environmental Impacts Caused by Buildings (*)
Annex 73: Towards Net Zero Energy Resilient Public Communities (*)
Annex 74: Competition and Living Lab Platform (*)
Annex 75: Cost-effective Building Renovation at District Level Combining Energy Efficiency and Renewables (*)
Annex 76: ☼ Deep Renovation of Historic Buildings Towards Lowest Possible Energy Demand and CO_2 Emissions (*)
Annex 77: ☼ Integrated Solutions for Daylight and Electric Lighting (*)
Annex 78: Supplementing Ventilation with Gas-phase Air Cleaning, Implementation and Energy Implications (*)
Annex 79: Occupant-Centric Building Design and Operation (*)
Annex 80: Resilient Cooling (*)
Annex 81: Data-Driven Smart Buildings
Annex 82: Energy Flexible Buildings Towards Resilient Low Carbon Energy Systems
Annex 83: Positive Energy Districts
Annex 84: Demand Management of Buildings in Thermal Networks
Annex 85: Indirect Evaporative Cooling
Annex 86: Energy Efficient Indoor Air Quality Management in Residential Buildings
Annex 87: Energy and Indoor Environmental Quality Performance of Personalised Environmental Control Systems
Annex 88: Evaluation and Demonstration of Actual Energy Efficiency of Heat Pump Systems in Buildings
Annex 89: Ways to Implement Net-zero Whole Life Carbon Buildings
Annex 90: EBC Annex 90 / SHC Task 70 Low Carbon, High Comfort Integrated Lighting
Annex 91: Open BIM for Energy Efficient Buildings
Annex 92: Smart Materials for Energy-Efficient Heating, Cooling and IAQ Control in Residential Buildings
Annex 93: Energy Resilience of the Buildings in Remote Cold Regions
Annex 94: Validation and Verification of In-situ Building Energy Performance Measurement Techniques
Annex 95: Human-centric Building Design and Operation for a Changing Climate
Annex 96: Grid Integrated Control of Buildings

Working Group - Energy Efficiency in Educational Buildings (*)
Working Group - Indicators of Energy Efficiency in Cold Climate Buildings (*)
Working Group - Annex 36 Extension: The Energy Concept Adviser (*)
Working Group - HVAC Energy Calculation Methodologies for Non-residential Buildings (*)
Working Group - Cities and Communities
Working Group - Building Energy Codes

Executive Summary

Advanced HVAC control is one of the most promising applications of “big data” from building automation systems, energy meters, and other sources. These advanced control strategies leverage access to data from sensors and building mechanical systems, that would have been unimaginable a few decades ago.

Considering the significant potential of data-driven approaches to optimise building performance, Subtask B of Annex 81 was devoted specifically to this topic. Other components of Annex 81 addressed complementary themes: Subtask A focused on data governance, platforms, and dataset utilization; Subtask C on advanced applications such as benchmarking, fault detection, and building–grid integration; and Subtask D on case studies, innovation strategies, and dissemination.

This report summarises the contributions, activities, major findings, and conclusions of Subtask B, which focused on the application of emerging data for model-based predictive control (MPC) and other data-driven advanced control strategies, such as reinforcement learning (RL).

MPC uses a mathematical model of a building and its systems to optimise the building’s operation over a control horizon, typically ranging from a few hours to a few days. Unlike conventional building operation, which is often “reactive” (e.g. heating activates only when temperatures drop below a setpoint), MPC adopts a “proactive” approach. It anticipates future conditions, such as weather, occupancy or pricing signals, and prepares accordingly. For example, MPC can optimise start times for heating or cooling, close motorised blinds in anticipation of solar gains, or plan energy storage charging when favourable. Similarly, RL can leverage historical or synthetic data (from digital models) to learn how various control actions affect performance and make informed decisions based on expected scenarios. Data plays a crucial role in developing optimal control strategies in both cases.

Subtask B activities, included: the collection of data from relevant case studies (Activity B1); a discussion on suitable modelling techniques for control applications (Activity B2); the use of a detailed simulation engine as a platform for *in silico* testing of control strategies (Activity B3); and a reflection on the importance of incorporating uncertainty (e.g. forecasts, model parameters, unknown factors) within control-oriented models to accelerate MPC deployment (Activity B4).

Additional research and discussion with industry practitioners, focussed on developing recommendations for policymakers on how to support industry adoption. The output of this work is included in a previous Annex report (*Opportunities for Government Leadership on Data-Driven Smart Buildings.pdf*). This work helped launch a survey initiative within ASHRAE (the leading North American professional organisation for heating, ventilation and air conditioning professionals) aimed at identifying and addressing barriers to the adoption of MPC.

The collective experience of the Annex participants confirmed the strong potential of advanced control to improve both energy efficiency and occupant comfort.

Among Subtask B’s recommendations to support wider adoption of data-driven control strategies is the need for a dedicated effort in properly structuring and managing data from building automation systems, thus making it easily accessible and available to the relevant users, while taking measures for IT. Supporting this recommendation, data was collected from case study sites in Activity B1. This data has been published and made available to facilitate exploring modelling approaches and advanced control techniques.

A key conclusion from the collection of data-driven models in Activity B2 is the pressing need to accelerate (even automate) the development of models for control applications. These models—typically classified as “white-box,” “grey-box,” or “black-box” based on their complexity and dependence on data—capture the diversity of real-world scenarios, including differences in building types, control applications, data availability, system configurations, and energy market contexts.

Activities B3 and B4 explored two distinct approaches to accelerating the deployment of data-driven building operation: (1) developing, testing, and comparing control strategies using an offline platform; and (2) employing simplified models that trade complexity for the ability to account for uncertainty.

Offline high-fidelity testing environments, such as BOPTEST, which served as the primary platform in Activity B3, enable the creation of standardised control solutions using a detailed digital twin in a "closed-loop" as a testing environment. This approach has been applied in initiatives like the ASHRAE Guideline 36 to develop systematic control sequences for energy-efficient performance.

Subtask B participants reported outcomes from several control methods, including MPC and RL, tested under standardised conditions. These showed that both MPC and RL data-driven controllers typically out-performed the traditional baseline controller by around 20%. MPC generally performed better than RL controllers, although RL controllers have the advantage of not needing a control-oriented model.

These results were obtained from a limited number of comparable studies. More such studies are required, covering more diverse building configurations and use-case scenarios. The BOPTEST library of building emulator models should be expanded to cover these extra cases.

Activity B4 explored a complementary perspective, focusing on the use of low-order data-driven models (simple resistive-capacitive thermal models based on a bare minimum of parameters calibrated with real data). Simple models are flexible enough to account for the limitations of the underlying data implicitly, the uncertainty of forecast variables like temperature, solar radiation, and occupancy, and even the influence of unforeseen factors, while delivering meaningful information for decision-making within an optimal (or near-optimal) control strategy. Moreover, simple models facilitate the assessment of the energy flexibility potential of buildings, i.e. their capacity to respond to pricing signals, and the adoption of a hierarchical control approach, suitable for optimising energy exchanges in buildings, communities, cities and regions.

Finally, building on these insights, a main takeaway of Subtask B is that data-driven control will be an essential factor in integrating smart buildings within the smart grid. MPC, RL and other data-driven control strategies have significant potential to enhance energy efficiency and support decarbonization, particularly from the perspective of load management, energy flexibility, and resilience. As buildings become more integrated with the smart grid, they will help seamlessly coordinate centralised and distributed energy resources. Consequently, data-driven building operations will be increasingly important in this integration. This topic will be the focus of future R&D efforts.

Table of Contents

Figures

Tables

Abbreviations

Abbreviations	Meaning
AR	Autoregressive Model
ARIMA	Autoregressive Integrated Moving Average
ARX	Autoregressive Model with Exogenous Inputs
BOPTEST	Building Optimisation Testing Framework
C&I	Commercial and Institutional Buildings
DDPG	Deep Deterministic Policy Gradient
DER	Distributed Energy Resources
DKSQP	Dieter Kraft's Sequential Quadratic Programming
DQN	Double-deep Q-network (RL technique)
DRL	Deep Reinforcement Learning
DSF	Demand-side Flexibility
DSL	Domain-Specific Language
FCU	Fan-coil Unit
FF	Flexibility Function
FMI	Functional Mock-Up Interface
FMU	Functional Mock-Up Unit
GDM	Gradient Descent Methods
GRU	Gated Recurrent Unit
HVAC	Heating, Ventilating and Air-Conditioning
KPI	Key Performance Indicator
LSTM	Long-short Term Memory
MDP	Markov Decision Process
MLP	Multi-layer Perceptron (a type of neural network)
MSE	Mean Square Error
MPC	Model Predictive Control
PPO	Proximal Policy Optimisation (RL technique)
Q-network	Neural network used in RL
QRDQN	Quantile Regression Deep Q-network (RL technique)
RBC	Rule-Based Control
RC	Resistive-Capacitive model
RNN	Recurrent neural network
RES	Renewable Energy Sources
RL	Reinforcement Learning
SAC	Soft Actor-Critic algorithm (RL technique)
SDE	Stochastic Differential Equation
SEOS	Smart Energy Operating System
SQP	Sequential Quadratic Programming

1. Introduction: Data-Driven Control

New requirements for building HVAC systems have emerged as part of global efforts to decarbonise energy systems and improve building resilience in response to a changing climate. These requirements include flexible load management, enhanced coordination with electric grids and other district energy systems, and tighter integration of distributed energy resources (DER). Consequently, this development calls for innovative HVAC control strategies that are more responsive to current and future operating conditions while optimising for multiple objectives. These new control strategies must rely on an increased awareness of current conditions, an understanding of system behaviour based on past performance, and the ability to anticipate the system's future behaviour over a relevant time horizon. Such control strategies, which are made possible through the availability of large datasets, are hereafter referred to as Data-Driven Control. These strategies have therefore become the focus of numerous recent research, development, and demonstration efforts.

While Annex 81 focused primarily on Model-based Predictive Control, our work included other significant data-driven strategies (most notably Reinforcement Learning). This section defines some basic terminology and summarises the structure of Subtask B activities.

1.1 Model-based Predictive Control

Model-based predictive control (MPC) is one of the most promising applications of data to improve the operation of buildings (Figure 1.1), although its potential has not been fully tapped (Drgoňa *et al.*, 2020). In this type of control approach (also known as “model predictive control”), a suitable mathematical *model* of a building and its systems provides forecasts, which then allows for control strategies to be scheduled in advance, provided that reliable forecasts of future conditions (which usually include weather conditions, building occupancy rates, energy prices, among other foreseeable events) are available.

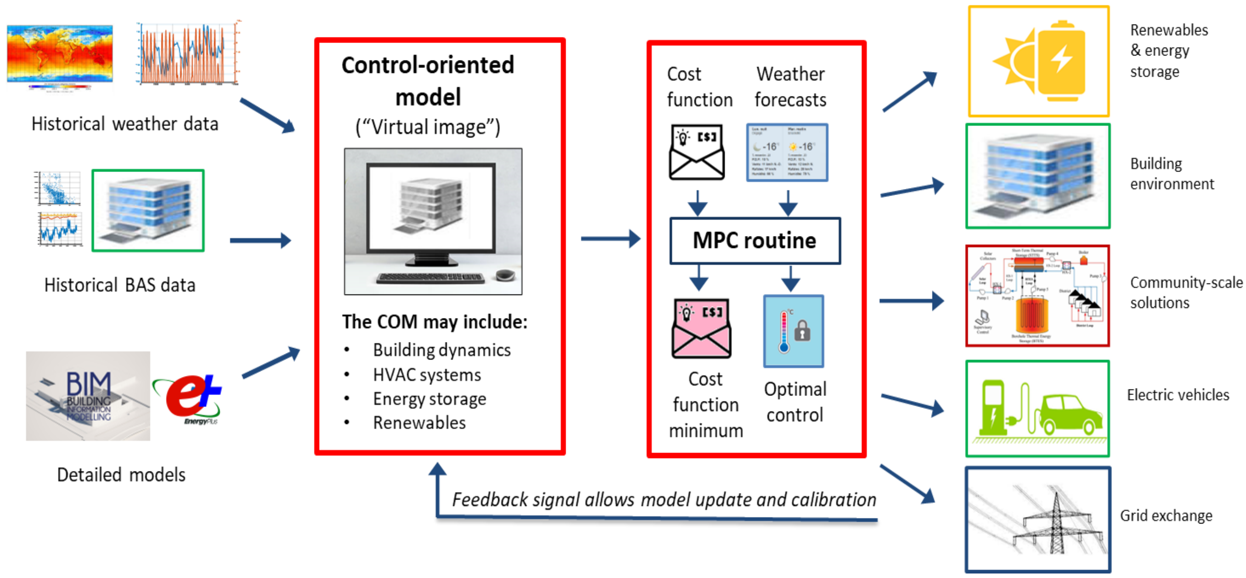

Figure 1.1: Schematic representation of model-based predictive control.

As illustrated in Figure 1.1, the Model Predictive Controller works in the following steps:

- **Step 1: Data collection**. In this step, diverse data streams are obtained, including historical weather conditions (ambient temperature, solar irradiance, humidity, etc.), recorded data from building automation systems and other sources, and, if necessary, complementary synthetic data from highly detailed models (if available). These datasets are used for the creation of a model that can be used for MPC. The datasets should be as *information rich* as possible, and reflect a panoply of weather conditions, modes of operation, setpoint transitions, etc.

- **Step 2: Development of control-oriented model**. Once enough data has been collected, an appropriate *control-oriented model* is developed. This model aims to provide sufficient information for a control algorithm to decide on the best control action within an upcoming period (denominated the control horizon). Ideally, this model should be accurate enough to provide reliable information for decision-making, yet simple enough that optimisation problems can be easily formulated and solved in real time.
- **Step 3: Optimisation routine (MPC).** Once an appropriate model has been developed, an *optimisation routine* is carried out to determine (via one of a vast range of optimisation algorithms) the optimal solution under a cost function or objective function. This cost function can include total costs, energy use, electric peak demand, GHG emissions, thermal comfort limits, wear and tear of equipment, etc.
- **Step 4: Sending signals to systems being controlled**. Finally, the control decisions made in the previous step are sent to the different systems being controlled, such as the HVAC system, EV chargers, energy storage devices, activation of renewables. Feedback from these systems can then be used to adjust the model in Step 2, thus having an impact on the optimisation routine.

MPC is useful in virtually any situation where knowledge of the future within a reasonable timeframe allows for better decision-making. For example, a solar radiation forecast enables the planning of electricity use based on the anticipated production of solar panels. Foreseeing a sudden drop in outdoor temperature allows for adjusting the building's indoor temperature setpoints to "pre-heat" it without exceeding comfort limits to avoid energy use during peak hours. Anticipating the human occupancy of an office makes it possible to activate an underfloor heating system. A thermal energy storage system (e.g. a water tank) allows for “stockpiling” energy collected during a particularly favourable period. The expectation of possible overheating in a space due to excessive solar gains can enable the gradual closing of motorised blinds. The examples are numerous and varied, and can be applied to all kinds of scales, from a small office, a medium-sized building (e.g. a school), a larger one (such as a hospital or a university campus) or an entire building cluster or neighbourhood.

Even today, most control strategies implemented in commercial and institutional buildings are still based on basic rules for supervisory level control (winter/summer mode, schedule-based setpoints) and on some basic "feedback control" loops, such as on/off thermostatic control, and "PID" loops (proportional, integral, derivative control). Although MPC has not yet taken off as a standard practice in the building mechanical systems control industry, there is widespread recognition of its potential, confirmed in numerous studies, for managing electricity demand (and consequently the relationship between buildings and the power grid), reducing greenhouse gas emissions, providing economic benefits to building operators, and significantly improving the well-being and comfort of building occupants.

The advent of "big data" (ever more abundant, with better resolution and containing additional variables) and the possibility of having better and easier access to this data are key pieces that allow us to glimpse the first steps towards the popularization of MPC. Subtask B focused on investigating how the availability of "big data" could be instrumental in driving the standardization of model creation, implementation procedures and ultimately the widespread adoption of MPC.

In a research context, MPC has been studied following two approaches, illustrated respectively in Figure 1.2 and Figure 1.3. Firstly, an implementation of MPC in an existing building (which is the final objective) implies the collection of data in a real building, based on the instrumentation of a BAS, energy meters, readings from IoT devices (Figure 1.2). Data acquisition could also include unconventional sources, such as Wi-Fi connections, readings from security systems, etc.

1.1.1 Dataflow in the Implementation of MPC in a real building

As illustrated in Figure 1.2, continuous communication between the real building and a control system will enable data to be transferred from the building automation system and used by the control algorithm. This control algorithm is written using a programming tool. Programming languages could be either commercial domain specific languages (DSL) within the building automation system or a general tool (e.g., Python, MATLAB, etc.). This control algorithm will use a *control-oriented model*, a mathematical representation of the building behaviour at appropriate time scales. The control-oriented model is usually (but not always) a simplified mathematical model, "grey box" or "black box" type, with relatively few parameters, which allows for short-term prediction of the building's behaviour. This type of model facilitates formulating and solving complex optimisation problems.

MPC also requires a forecast of relevant variables (e.g., weather and occupancy) over a horizon ranging from a few minutes up to a few days. Finally, the control strategy is selected according to an objective function. The objective function varies depending on the goal pursued by the operator: cost, comfort, GHG reduction, peak demand reduction, energy savings, etc. Information travels towards the real building, conveying the optimal control actions calculated within the optimisation algorithm using the control-oriented model. Conversely, measured data is collected from the real building and used to improve the accuracy of the control-oriented model.

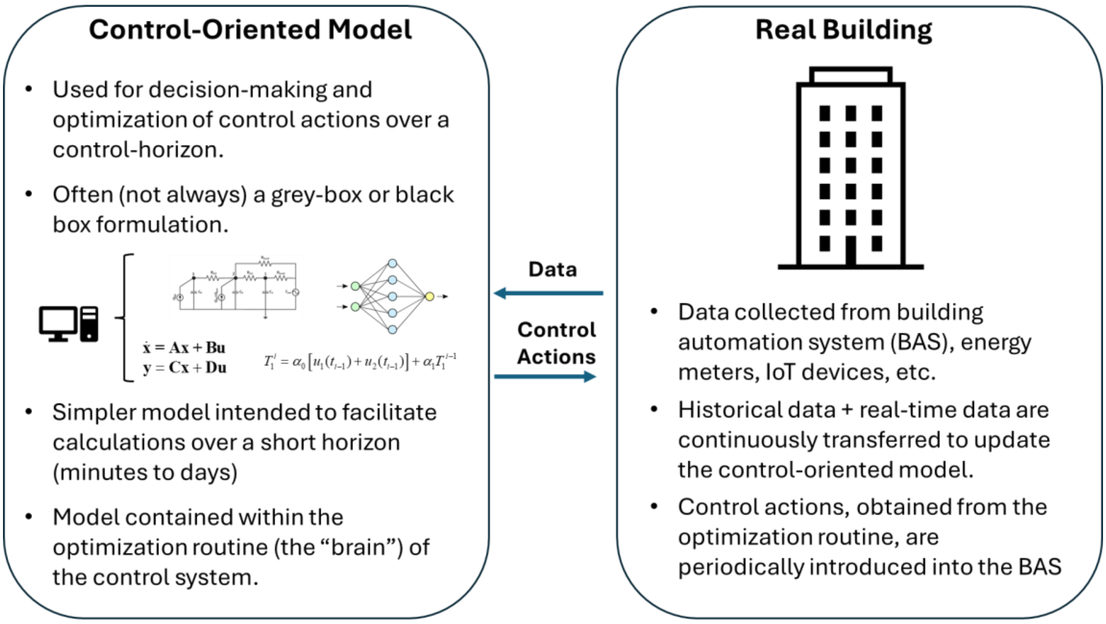

Figure 1.2: MPC implementation in a real building.

1.1.2 Testing MPC within an Emulator

It is sometimes convenient to study the application of an MPC algorithm in the context of a "closed loop" simulation, i.e. without interaction with a real system. For this type of study, carried out completely in a numerical simulation framework (Figure 1.3: Implementation of MPC with an emulator.

), one must have a mathematical model known as an "emulator," a simulation model that faithfully reproduces, with as much fidelity as possible, a real building. Colloquially, this type of model is called a *digital twin,* although this terminology is less often used in more formal settings.

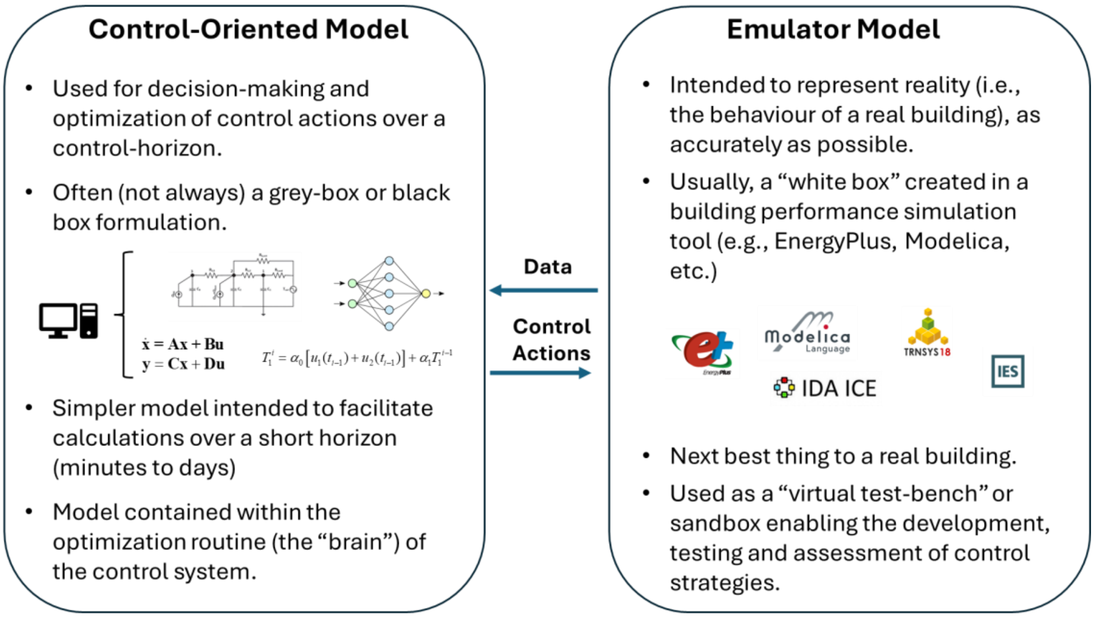

Figure 1.3: Implementation of MPC with an emulator.

Similarly to how the MPC algorithm interacts with a real building, control decisions are sent to the emulator, which then generates synthetic operational data to update and fine-tune the control-oriented model.

Using an emulator to study and compare control strategies has several advantages. For example, "virtual experiments" can be conducted under reproducible conditions and using fair comparison criteria. Implementing MPC in a real-world context faces the challenge of comparing the effect of a real control action with the hypothetical results of a "benchmark" set of control actions. In contrast, an emulator allows replicating the same conditions, as many times as necessary, under alternative control sequences and providing consistent, reliable results. An emulator also enables different research teams to tackle the same problem, to provide a way to compare fairly and rigorously the results of their proposed strategies. It is worth noting that control-oriented models are needed both in a real implementation and in a study using an emulator. These types of models (not quite "digital twins", but perhaps "digital cousins") are part of the controller itself and an essential element for control decision-making.

MPC consistently yields significant benefits; a recent review article of field implementations found average savings of 25% in electricity use and energy costs, and about 17% in GHG emissions (Saloux et al., 2025).

1.2 Reinforcement Learning

Reinforcement learning (RL) is another data-driven control approach that has received significant attention recently. In reinforcement learning, a branch of machine learning, a controller (or control agent) applies a set of control actions to a building system with the goal of maximising a *reward*, such as energy cost, energy flexibility, thermal comfort or GHG emissions (Vázquez-Canteli and Nagi, 2019; Wang and Hong, 2020). Reinforcement learning can be described through a Markov Decision Process (MDP), represented by a *tuple* (S, A, P, R):

- **S**: *State* is a mathematical description of the current situation of a system (the "environment), such as a building, by means of relevant parameters (e.g., room temperatures).
- **A**: *Action* is the execution of a decision on how to operate the system, such as valve positions, temperature setpoints, etc.
- **P**: *Transition probability* is the mathematical description of the likelihood of the next state, given the current state and the action.

- **R**: *Reward function:* describes the immediate reward for taking at a given time step.

The rules that govern which action to take under a given are described by a *policy function* (π), which maps states to actions:

$$\pi : S \rightarrow A$$

The main goal of the reinforcement learning controller or agent is to find the optimal π that will maximise the rewards. This policy is typically derived by collecting a large dataset of examples illustrating the outcomes of specific actions taken under given situations. This approach, similar to humans learn through "trial and error", has been successfully used in areas such as robotics.

While the main challenge of MPC is the laborious procedure involved in model creation, in principle, reinforcement learning can develop effective operational policies *without* a model. Nevertheless, the main challenge consists in obtaining sufficient data to derive an appropriate control policy. The collection of examples is often carried out within a "virtual environment", which may include a model or collection of models that provide an accurate representation of what will happen when an action takes place. This type of "training environment" is sometimes called a "gym", metaphorically expressing how the agent develops by intensively applying control actions to a large set of situations.

1.3 Some Terminology

Some key terms mentioned in this report are succinctly defined to help the reader navigate jargon that might sometimes become confusing.

- **Agent.** In the context of reinforcement learning, an agent is a controller or entity which decides which action to take as a function of the conditions in the environment and the optimal policy (π), developed through trial and error to optimise a certain reward.
- **Control-oriented model**. A simplified mathematical model of the building (which is baked into the controller itself) to forecast the future state and energy consumption of the building if different control actions were to be taken. The MPC controller can use the control-oriented model to optimise decision-making, considering both current and future conditions.
- **Controller**. A control algorithm (e.g., PID controller, MPC controller, or RL controller).
- **Emulator**. A high-fidelity simulation model of a specific building system which is used to replicate, as closely as possible, the behaviour of a real system at appropriate time scales for building operation. It enables testing and evaluation of advanced control strategies in a closed loop (i.e., "offline", within a confined computer simulation), as it is the "next best thing" to a real building.
- **Hyperparameter:** In reinforcement learning (RL), hyperparameters are configuration settings that are not learned during training but are manually defined before training begins. They influence how the learning process unfolds, such as learning rate, discount factor, or network architecture**.**
- **Optimiser**: Algorithm used to find the optimal solution according to an objective function or reward.
- **Solver**: This term refers to an algorithm that finds a solution to a mathematical problem. In the context of optimisation, it specifically refers to an algorithm whose purpose is to find the optimal solution (in this sense, it will be a synonym for "optimiser").
- **Surrogate modelling**: simpler model that reproduces the behaviour of a more complex system. It is akin to a control-oriented model, but it can have a more general purpose.

1.4 Subtask B Structure

Figure 1.4 shows how the activities within Subtask B were organised. These activities concern various aspects of predictive control, namely: the collection and organisation of data (B1), the creation of control-oriented models (B2), studies of MPC performance within an emulator environment (B3), and investigations relating to the challenges of implementing predictive control in real time. The most relevant results of these activities are presented in the following pages. Activities B1-B4 fed into the development of a policy roadmap (activity B5), which has been detailed in an earlier report.

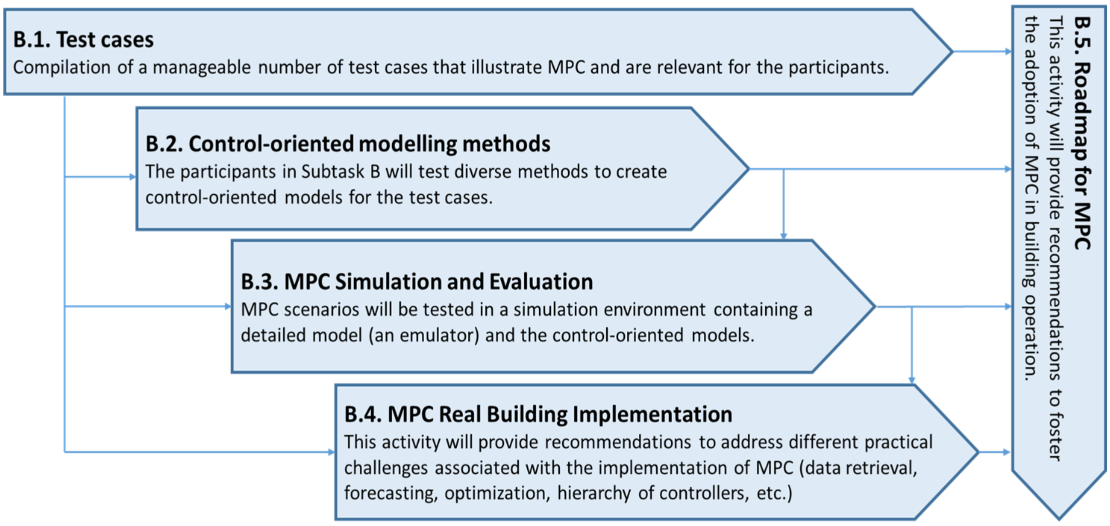

Figure 1.4: Subtask B Structure.

1.5 Organisation of this Report

The first chapter of this report presents background information on data-driven control, the basic concepts of model-based predictive control (MPC), its operation in a real building and within an emulator, a brief introduction to reinforcement learning, and an overview of the main structure of the work carried out in Subtask B. Chapter 2 describes the work conducted in Activity B1 on test case collection. Chapter 3 focuses on Activity B2, discussing diverse modelling approaches used in MPC case studies. Chapter 4 covers Activity B3 and the studies conducted within BOPTEST, an emulator framework, on different control approaches. Chapter 5 outlines Activity B4's approach to implementing an MPC algorithm in buildings in the context of a smart grid, incorporating uncertainty treatment and a hierarchical approach. Finally, Chapter 6 presents concluding remarks.

1.6 References

Abdelrahman, M., Macatulad, E., Lei, B., Quintana, M., Miller, C., & Biljecki, F. (2025). What is a Digital Twin anyway? Deriving the definition for the built environment from over 15,000 scientific publications. Building and Environment, 112748.

Drgoňa, J., Arroyo, J., Figueroa, I. C., Blum, D., Arendt, K., Kim, D., Ollé, E.P., Oravec, J., Wetter, M., Vrabie, D.L. & Helsen, L. (2020). All you need to know about model predictive control for buildings. *Annual Reviews in Control, 50*, 190-232.

Saloux E., Candanedo J., Vallianos C., Morovat N., Zhang K. (2025). From theory to practice: a critical review of model predictive control field implementations in the built environment. Accepted in *Applied Energy*.

Vázquez-Canteli, J. R., & Nagy, Z. (2019). Reinforcement learning for demand response: A review of algorithms and modeling techniques. *Applied Energy, 235*, 1072-1089.

Wang, Z., & Hong, T. (2020). Reinforcement learning for building controls: The opportunities and challenges. *Applied Energy, 269*, 1-19.

2. Activity B1: Test Cases

2.1 Introduction

This activity has been a joint effort by the Annex participants to compile a diverse range of datasets suitable for advanced control applications of indoor climate and energy use in buildings. Initially, more than twenty test cases were collected for consideration from all the participants in the Annex who provided the necessary information through a common template. These cases corresponded to real buildings, from experimental/lab setups or simulations. It was then decided to continue working with 6 test cases from real buildings, where both the available dataset itself and the *metadata* (relevant contextual information) were of high quality.

All test case datasets had either been collected for developing a model to be used in MPC applications or were judged to be suitable for the purpose. The datasets have been made publicly available for download at the Mendeley Data repository:

- https://data.mendeley.com/datasets/xztfbtsgys/3

The datasets are described in detail an open access data paper:

- Igor Sartori, Harald Taxt Walnum, Kristian S. Skeie, Laurent Georges, Michael D. Knudsen, Peder Bacher, José Candanedo, Anna-Maria Sigounis, Anand Krishnan Prakash, Marco Pritoni, Jessica Granderson, Shiyu Yang, Man Pun Wan (2023) Sub-hourly measurement datasets from 6 real buildings: Energy use and indoor climate, *Data in Brief*, 48-109149, ISSN 2352-3409. DOI: https://doi.org/10.1016/j.dib.2023.109149.

This report focuses on describing the test cases, i.e. the buildings and the most relevant technical system information. A brief description of how the data was collected is provided below, while the above-mentioned paper should be consulted for a description of the datasets.

The datasets were collected independently for six real buildings from different buildings in different countries and climates. The data was acquired from energy meters, both consumption and PV generation, and sensors of indoor climate variables, such as temperature, flow rate, relative humidity, CO_2 level, and illuminance. Weather variables, such as temperature, solar radiation and wind speed, were either acquired by local sensors or obtained from a close by meteorological station.

Most data is raw data, although some variables in some datasets have been filtered, such as the indoor temperature being the average of several sensors placed at different heights or different points. Two experimental designs were used: longer observation periods (2 weeks to 2 months) with normal operation, and shorter observation periods (ca. 1 week) with Pseudo-Random Binary Sequence (PRBS) experiments aimed at exciting the thermal response of the building, when there was no occupation. The data has a time resolution varying between 1 min and 15 min; in some cases, the highest resolution data are also averaged at larger intervals, up to 30 min.

2.2 Collected Test Cases

2.2.1 Dataset 1: Office building in Oslo, Norway

This office building, built in the 1960s, has a concrete construction with a brick façade; it has eight floors plus a basement, with a total floor area of about 3,800 m^2. The basement and 1st floor are laboratories and common areas. The 2nd to 7th floors are office areas, a mix of cell offices and open-plan configurations, while the 8th floor consists of meeting rooms. A schematic of the building and its HVAC system is given in Figure 2.1.

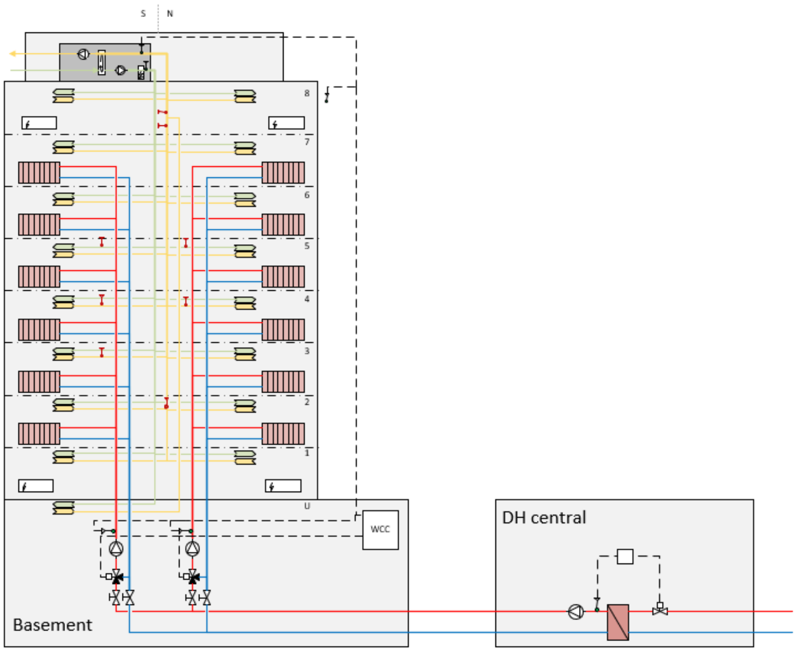

Figure 2.1: Schematic of the building's HVAC system.

The building's air handling unit (AHU) has a rotating wheel heat recovery unit and an electric reheating element. During normal operation, the air volume flow is controlled by a daily schedule to 0%, 67% or 100%. The supply air temperature is controlled by an extract air temperature compensation curve. The air is distributed and extracted to the different floors in one common shaft. In each shaft, the air is split into two branches, covering the northern and southern parts of the building, respectively.

Except for the 1st and 8th floor, the room heating in the building is covered by a radiator system supplied by district heating. The district heating substation is located in a nearby building, and the heat is transferred to the central heating unit in the basement. From there it is split into two radiator circuits supplying the eastern and western façade. An outdoor temperature compensation curve controls the supply temperatures of the radiator circuits, and the two circuits have individual outdoor temperature sensors. The radiators have no thermostatic control, only manual control valves. This means that the outdoor temperature compensation curve only controls the building's heat supply in a so-called weather-compensated control (WCC). Further details on the building are available in Walnum *et al.* (2024).

During the Easter holiday, a pseudorandom binary signal (PRBS) experiment was performed on the building to generate a rich dataset for model identification.

2.2.2 Dataset 2: ZEB Living Lab, Norway

The experiments performed in this study were carried out in the ZEB Living Lab, a single-family, zero-emission house with a heated floor area of about 100 m^2 on the Norwegian University of Science and Technology (NTNU) campus in Trondheim. The building envelope has a wooden frame with mineral wool insulation measuring 35–40 cm and a window-to-wall ratio of 0.23, or a glazing ratio of 0.14 (when the opaque window panels and door frames are subtracted). More information on the envelope characteristics is provided in Skeie (2022). Space heating can be floor heating, a central radiator, or air ventilation. The balanced mechanical ventilation system is equipped with a heat recovery unit, and the additional ventilation heating battery was disabled. By operating the doors in the building, four zones can be created (bedroom west, bedroom east, bathroom, and living areas). The building and the internal layout of the Living Lab are shown in Figure 2.2.

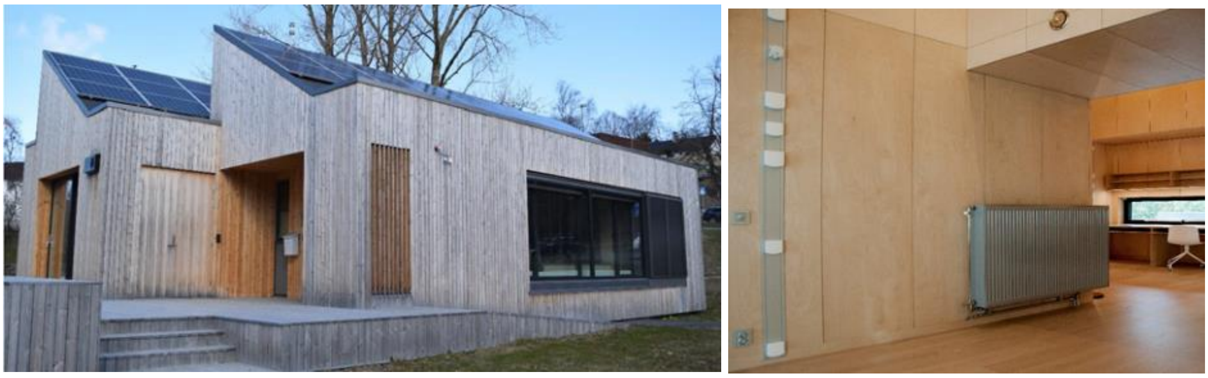

Figure 2.2: (left) Exterior view of the ZEB Living Lab (right) Interior view with hydronic radiator and example of wall-mounted sensor.

2.2.3 Dataset 3: FlexHouse, Denmark

FlexHouse is a part of the experimental distributed energy system (Syslab) at DTU campus. The building is controlled by one central server, where, among other things, it is possible to record the temperature in each room and implement control of the installed electrical heaters. Furthermore, a climate station is located right next to the building; see Figure 2.3 Bacher and Madsen (2010) provide more details on the building and the data.

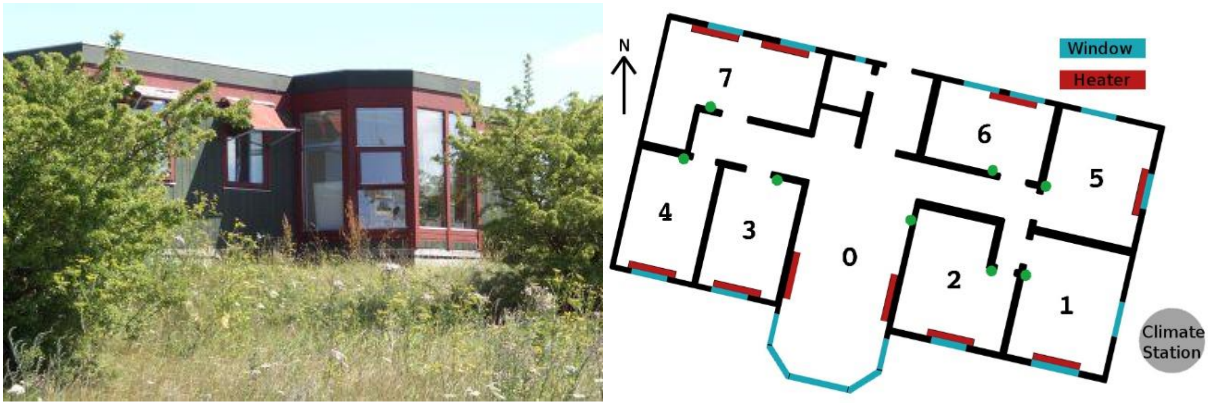

Figure 2.3: (left) View from the south side of FlexHouse and (right) Floorplan: The green dots indicate the position of the wall-mounted temperature measurements.

The data is from experiments carried out with the objective of estimating the building's thermal dynamics. The designs of the experiments vary from conditions optimised for modelling to more common living conditions, i.e., high variation of the indoor temperature, thermostatic temperature control and human activities in the building.

2.2.4 Dataset 4: Varennes Library, Canada

The Varennes Net-Zero Energy library, completed in 2014, is the first net-zero institutional building in Canada (Figure 2.4). This building is located in the town of Varennes, 30 km northeast of Montreal (ASHRAE Climate Zone 6). It is a two-storey building with a floor area of about 2,100 m^2. It comprises a large basement containing the mechanical room and a mezzanine between the two floors. Both floors consist of mainly open plans used primarily as library spaces, with some closed offices on the first floor.

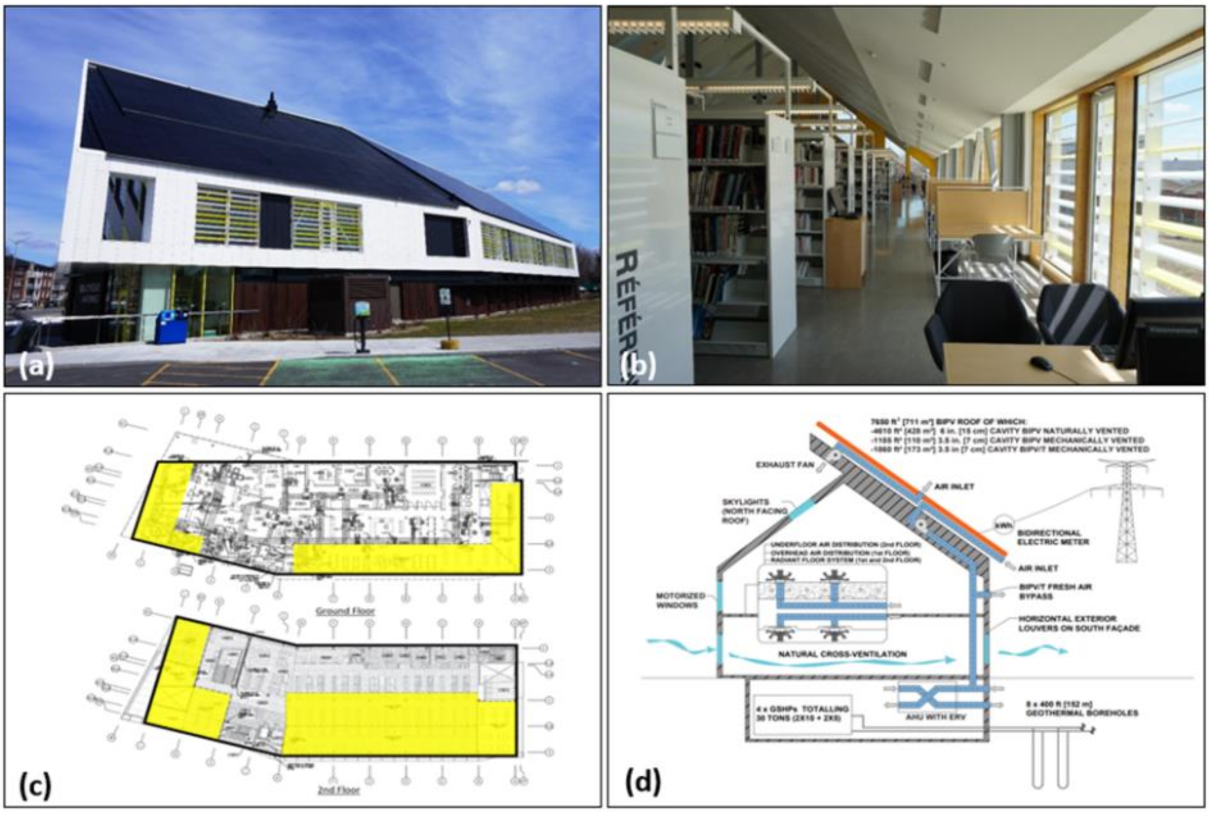

Figure 2.4: (a) Varennes Net-zero Energy Library exterior, (b) Varennes Net-zero Energy Library interior, (c) Plan highlighting areas supplementary served by a hydronic radiant concrete circuit, (d) Schematic of library and integrated energy systems (Dermardiros *et al.*, 2019).

The library has a 110.5 kWp, roof-mounted, building-integrated photovoltaic (BIPV) array on the south-facing section of the roof. From the total BIPV area (706m^2) about 16% also incorporates thermal recovery from where air is collected and used as pre-heated ventilation air during the heating season. The building's heating and cooling needs are covered by a geothermal system of 8 boreholes feeding four heat pumps with a capacity of 105 kW. Space heating and cooling are distributed by hydronic radiant slabs in the south, east and west perimeter zones. The building is well instrumented, with hundreds of variables recorded in its BAS located in its mechanical room. The data was collected under normal operation of the library for two consecutive months during the heating period.

2.2.5 Dataset 5: FLEXLAB, USA

FLEXLAB is a well-instrumented experimental test facility at Lawrence Berkeley National Laboratory in Berkeley, California, United States (Figure 2.5). It has a modular and flexible design with options to modify the building envelope, HVAC and lighting systems, indoor room arrangement etc. It consists of four testbeds, each containing two identical rooms and allowing for a side-by-side comparison of a novel technology against a baseline scenario.

The testbeds also have solar panels installed on the roof-top along with multiple Tesla Powerwall batteries. With these technologies, FLEXLAB allows for accurate measurement and evaluation of the behaviour of integrated systems (e.g., HVAC and batteries) in providing grid services (e.g. shift/shed load when requested) and building services (e.g., visual and thermal comfort).

The data in this test was collected from testbed "XR," which contains two identical and adjacent rooms (Room CellA and Room CellB). Each represents a small commercial office space with a floor area of 57 m^2 and a large south-facing window.

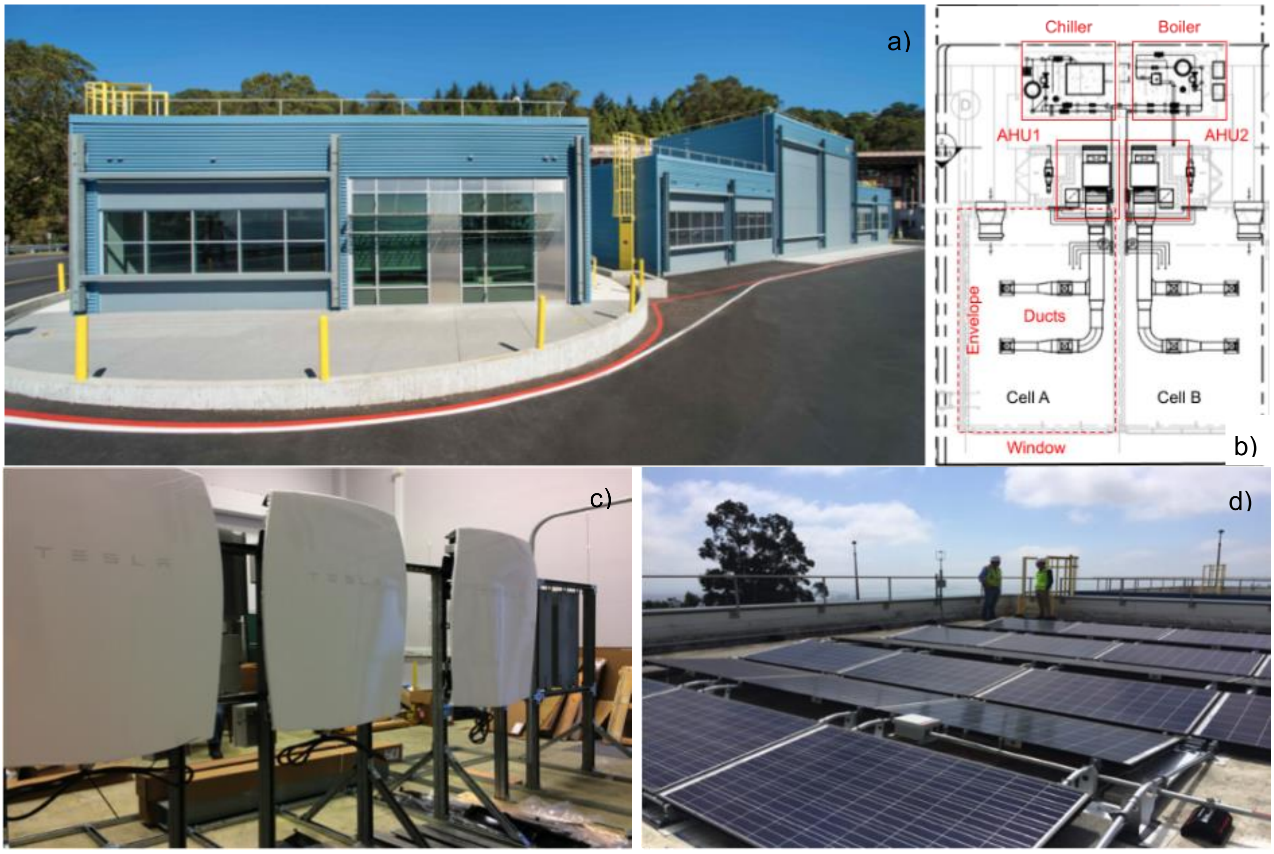

Figure 2.5: (a) Outside view of FLEXLAB, (b) Mechanical drawing of FLEXLAB envelope and HVAC, top view, (c) 7.2 kWh local energy storage capacity with a peak output of 3.3 kW: only two of the three shown have been used in the test (d) 3.64 kW local solar power generation (Touzani *et al.*, 2021).

As part of this experiment, the HVAC system in FLEXLAB was reconfigured to be a single-zone variable-capacity AHU that conditioned each room in the testbed. It had a single variable-speed fan and two dampers for outdoor air and return air mixing. The water side of the two AHUs was served by a small chiller and boiler shared by the two cells. Each room had six ceiling-mounted light fixtures, six plug-load stations (i.e., desks with desktop computers), and six heat-generating mannequins that emulate occupants. The distributed energy resources on the testbed comprise a 3.64 kilowatt (kW) photovoltaic system and a Tesla Powerwall battery storage, with a capacity of 7.2 kWh and a peak power output of 3.3 kW.

2.2.6 Dataset 6: Office space at Research Techno Plaza, Singapore

The testbed office is located at level 3 of Research Techno Plaza building on the campus of Nanyang Technological University, Singapore, as shown in Figure 2.6. The dataset contains time-series data of variables in the heating, ventilation, and air conditioning (HVAC) system and room space.

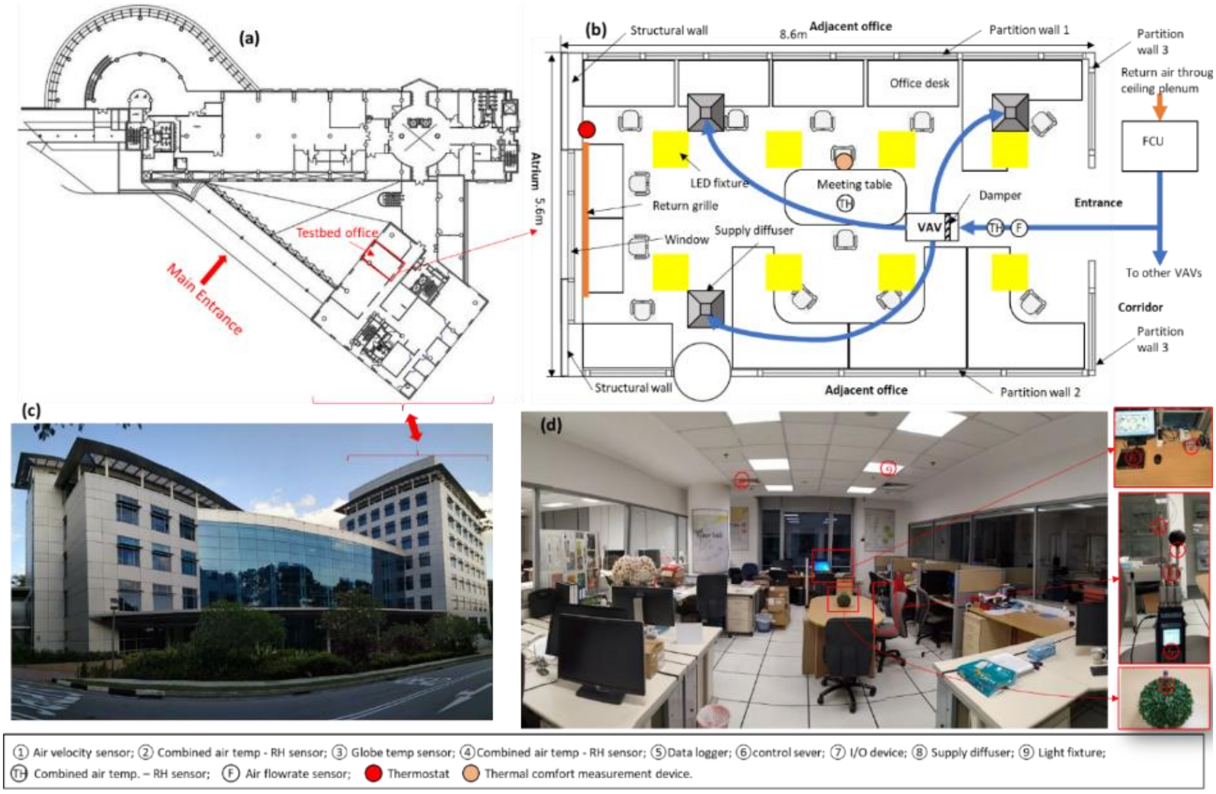

Figure 2.6: Overview of the testbed (a) Location of the testbed office in the test building, (b) internal setup of the testbed office, (c) external view of the test building, and (d) internal view of the testbed office The testbed office has a floor area of 46 m2, with a floor-to-ceiling height of 3 m, and is air-conditioned by a variable air volume (VAV) box. Further description and details of the testbed office setup are given in (Yang *et al.*, 2020).

2.3 References

Bacher, P., & Madsen, H. (2010). Experiments and data for building energy performance analysis: Financed by The Danish Electricity Saving Trust. Technical University of Denmark, DTU Informatics, Building 321. *IMM-Technical Report-2010-03*. https://orbit.dtu.dk/en/publications/experiments-and-data-for-building-energy-performance-analysis-fin

Dermadiros, V., Athienitis, A., & Bucking, S. (2019). Energy performance, comfort, and lessons learned from an institutional building designed for net zero energy. *ASHRAE Transactions, 125*(1). Available here.

Skeie, K. S. (2022). Building energy performance evaluation of a Norwegian single-family house applying ISO-52016. *BuildSim Nordic 2022, Copenhagen*. Available here

Touzani, S., Prakash, A. K., Wang, Z., Agarwal, S., Pritoni, M., Kiran, M., Brown, R., & Granderson, J. (2021). Controlling distributed energy resources via deep reinforcement learning for load flexibility and energy efficiency. *Applied Energy, 304*, 117733. https://doi.org/10.1016/j.apenergy.2021.117733

Walnum, H. T., Sartori, I., Ward, P., & Gros, S. (2024). Demonstration of a low-cost solution for implementing MPC in commercial buildings with legacy equipment. *Applied Energy*. In publication.

Yang, S., Wan, M. P., Chen, W., Ng, B. F., & Dubey, S. (2020). Model predictive control with adaptive machine-learning-based model for building energy efficiency and comfort optimisation. *Applied Energy, 271*, 115147. https://doi.org/10.1016/j.apenergy.2020.115147

3. B2: Control-Oriented Modelling

3.1 Introduction

The mathematician George E.P. Box famously said: "*we make tentative assumptions about the real world which we know are false but which we believe may be useful nonetheless*" (Box, 1976). This phrase, often rendered as “all models are false, but some are useful”, contains a profound truth: A model is meant to be a simplified representation of reality, intended to target a specific purpose, such as gaining insight about a phenomenon, informing project management, selecting a specific design or planning control decisions.

A critical challenge -arguably the main hurdle- for implementing data-driven MPC is the development of control-oriented models that are sufficiently accurate and reliable for the purpose at hand, but are also manageable within an optimisation routine (or at least allow a comparison of different control scenarios). Activity B2 focused on collecting and discussing different approaches to model creation contributed by the Annex participants. These examples include diverse building types and climates, using both data from real buildings and synthetic data from detailed simulations. Diverse control-oriented modelling approaches were used.

Chapter 3 is organized as follows: First, it discusses the most used approaches for the creation of MPC models in Sections 3.2 through 3.5. Secondly, Section 3.6 briefly discusses the examples provided by the Annex participants. Lastly, future steps and research opportunities for control-oriented models are discussed.

3.2 Control-Oriented Modelling Methods

Modelling approaches are traditionally classified using the categories of "white-box," "grey-box," and "black-box" models (Li *et al.*, 2021). This classification is rather coarse, and it can only hint at the vast variety of ways a mathematical model can be created starting from a dataset. Paraphrasing G.E.P. Box, we could say that "all categorisations are false, but some are useful." This white-grey-black classification, limited as it is, provides a practical primer on how a model can be deployed for control applications.

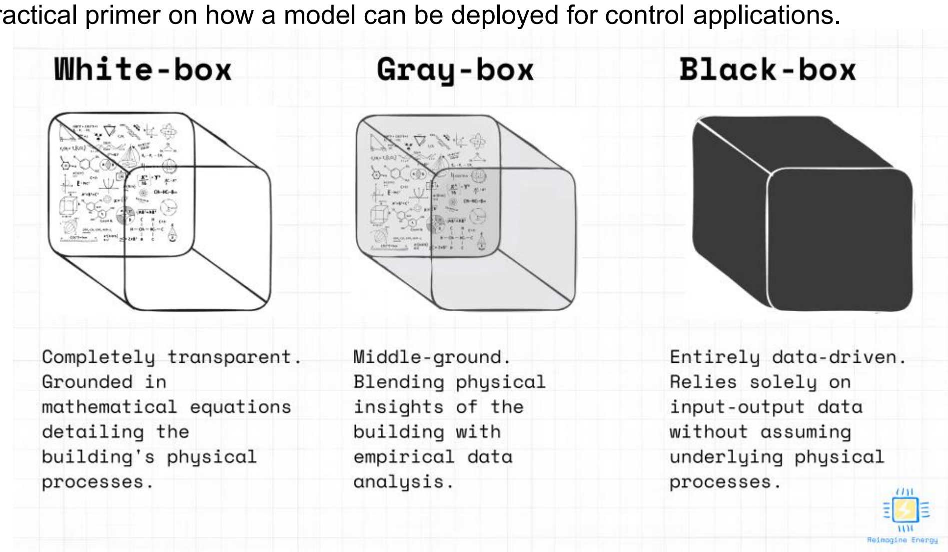

Figure 3.1: White-box, gray-box and black-box models (image credit: Benedetto Grillone, https://benedettogrillone.substack.com)

3.2.1 White-Box Models

As illustrated in Figure 3.1, a white-box model is based on parameters with clear physical meanings. The equations rely on an aggregation of well-understood physical phenomena, such as energy balance equations at specific locations, heat transfer coefficient correlations, and radiative heat transfer models. Both the equations and their associated parameters are clearly and explicitly accessible. In buildings, white-box models refer to models created using building performance simulation tools: EnergyPlus, TRNSYS, Modelica, among others. White-box models have several advantages: interpretability, extrapolation under unusual or unforeseen situations and incorporation of novel technologies. Therefore, a white-box model can play the role of a "digital twin" in the sense of closely reproducing a building's behaviour in real time for decision-making.

In the context of our discussion on "data-driven" approaches, white-box models can also be calibrated with real measurements: temperatures, flow rates, electric power measurements, etc. After fine-tuning it with real data, a reliable and accurate white-box model is a powerful tool for assessing control strategies under different scenarios.

A white-box model specifically created to mimic a particular building configuration as closely as possible to reality is called an "emulator." This kind of model, arguably the next best thing to an actual building, enables testing and evaluating control strategies under a "level playing field." This report discusses such an approach in Section 4, in which emulators developed for the BOPTEST environment enable the comparison of control strategies.

White-box models may also be used as control-oriented models within an MPC controller. However, control-oriented models tend to be simpler grey-box or black-box models, with fewer input variables and a reduced number of parameters.

3.2.2 Black Box Models

In contrast with white-box models, black-box models describe mathematical relationships between a set of inputs and outputs. They are purely data-reliant, and their parameters do not have an explicit physical meaning. There is a wide variety of black-box modelling approaches. Here are just a few examples:

- **Linear regressions** (Bünning et al., 2021; Ciulla & D'Amico, 2019) and **non-linear regressions.** The oldest black-box modelling approach, the least squares method, was created in the 18th century. It involves finding the coefficients of a function (which can be either a linear or non-linear function of the inputs) that will minimise the difference (typically measured using the mean square error as a parameter) between predicted values and measured output variables. For a linear regression, the basic equation is:

$$y = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \ldots + \beta_n x_n + \varepsilon \tag{3.1}$$

  Where the output is a linear function of inputs (x_1, $x_2 \ldots x_n$). In the case of a non-linear regression, the general equation is:

$$y = f(\mathbf{x}, \boldsymbol{\theta}) + \varepsilon \tag{3.2}$$

  In this case, a set of parameters $\boldsymbol{\theta}$ is used to predict the output y based on a vector of inputs $\mathbf{x}$.

- **Time-series methods** (Deb et al., 2017; Zhuang et al., 2023)**.** Time-series models are a mathematical framework in which the output variable depends on *past* values of the output variable collected at regular intervals (*autoregressive model*) and on present and past values of input variables (exogenous variables). For example, an autoregressive exogenous (ARX) model has the form:

$$y_t = a_1 y_{t-1} + a_2 y_{t-2} + a_3 y_{t-3} \ldots + \beta_1 x_{t-1} + \beta_2 x_{t-2} + \beta_3 x_{t-3} \tag{3.3}$$

  Many other schemes exist, such as autoregressive models (AR), autoregressive integrated moving average models (ARIMA), etc. Time-series models share some similarities with linear regressions, but they implicitly embed information about time delays. They can be derived from physical models (e.g., grey-box models), or they can be obtained by applying system identification techniques to calibrate the values of the coefficients in the models.

- **Machine learning methods** include decision trees (Yu et al., 2010), neural network models (Afram et al., 2017; Macarulla et al., 2017), support vector machines (Dong et al., 2005), Gaussian process régressions (Maddalena et al., 2022), gradient boosting models (Miller et al., 2020), etc.
- **Deep learning methods** (LeCun et al., 2015) : recurrent neural networks (Fan et al., 2019), long short-term memory (LSTM) networks (Mtibaa et al., 2020).

Black-box models are a popular option for implementing predictive control, especially within the building operation industry. They can be rapidly deployed, require less expertise about building systems, and they can achieve a high degree of accuracy. However, black-box models also have shortcomings: they require a significant amount of high-quality data for their creation, some expertise and skill is required in selecting appropriate input variables, and they lack extrapolation capability.

3.2.3 Grey Box Models

A grey-box model represents a sort of compromise between a white-box and a black-box model. While it does not include as many details as a white-box model, its parameters still retain physical meaning, typically as "effective" or "equivalent" values. Moreover, unlike white-box models, where parameters are derived from physical properties or the building's geometry, a grey-box model's parameters are calibrated from collected data, akin to a black-box model.

The classic grey-box model (albeit not the only one) is a thermal RC network, or resistive-capacitive network. Figure 3.2 shows an example of an RC network for a small building (e.g., a small commercial building or a house). The nodes 1, 2 and 3 represent respectively the temperatures of the indoor air, the inner walls and floors and the building envelope. The thermal capacitances represent the effective energy storage capacity of these components, and the resistances represent the effective thermal resistance between them, including the resistance between specific nodes and the outdoor air. For instance, R_{1ext} represents the equivalent thermal resistance between the indoor air and the ambient air (this kind of direct link can be the result of infiltration and heat loss through the fenestration). Q_{SG}, Q_{IG} and $Q_{heating}$ respectively represent the thermal contributions of solar gains, internal gains and the heating system.

In this case, the RC network in Figure 3.2 is a "third order" network, as each capacitance defines a differential equation. The actual value of these parameters (e.g., the resistances in K/W or the capacitances in MJ/K) may be calibrated by collecting data from a real system.

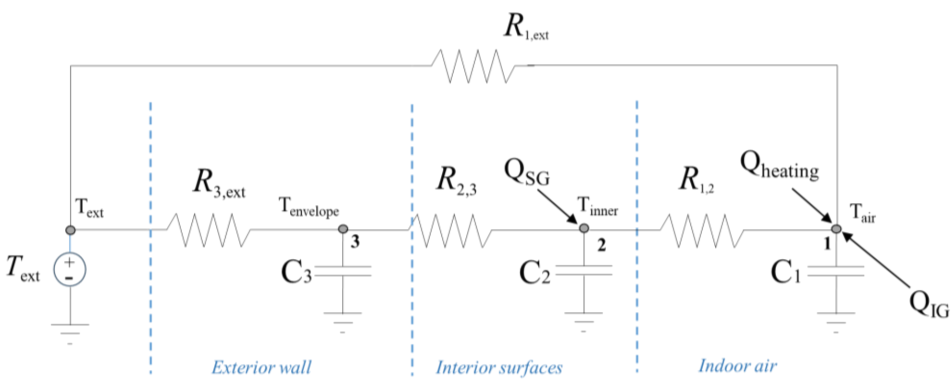

Figure 3.2. Example of grey-box model: RC representation of a simple building.

A keen observer might point out that the model structure is somewhat arbitrary, and that other parameters could have been included. This is a valid point—there is indeed an element of subjectivity in defining this RC network. However, it is important to remember that the primary objective of this type of model is to offer a simplified representation that supports decision-making.

Grey-box models must be *structural-identifiable (*Madsen *et al.*, 2006). This implies that depending on the number of input and output variables, the structural complexity must be balanced, and the number of differential equations must be limited to the most important. Some work has been carried out to create generic archetypes capable of representing typical configurations (Candanedo et al., 2022).

The heat balance in each of the nodes with capacitances may be described by ordinary differential equations. For instance, the heat balance in node 1 in Figure 3.2 can be written as follows:

$$C_1 \frac{dT_1}{dt} = \frac{T_{\text{ext}} - T_1}{R_{1,\text{ext}}} + \frac{T_2 - T_1}{R_{1,2}} + Q_{\text{IG}} + Q_{\text{heating}} \tag{3.4}$$

In this equation, the parameters C_1, R_{1ext}, and R_{12} (which loosely stand for the thermal energy stored in the indoor space or thermal resistance between specific areas) can be calibrated using collected data. In general, the grey-box framework bridges the gap between models based on first principles (white-box models) and models based solely on data (black-box models), as shown schematically in Figure 3.3. Boundaries between modelling approaches can be blurry, depending on the model's parameterisation level, and how these parameters can be interpreted as representations of physical phenomena. As discussed in a paper on the identification of control-oriented models, there can be numerous "shades of grey" (Leprince *et al*., 2022).

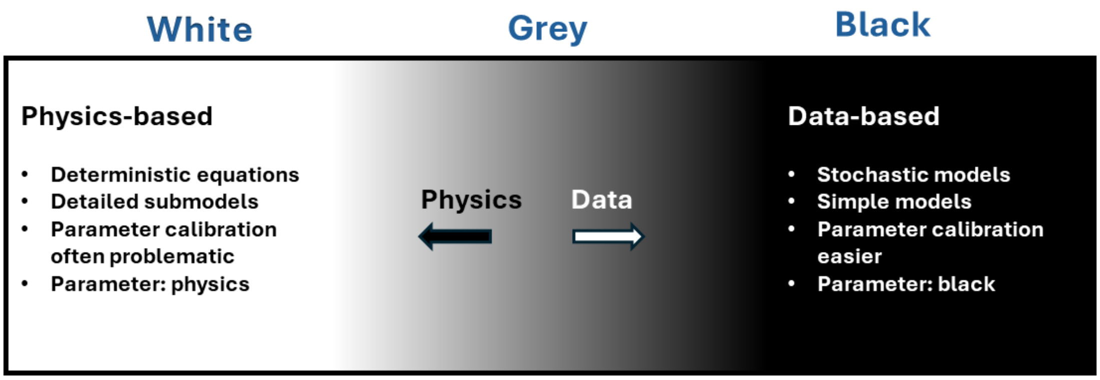

Figure 3.3: Grey-box modelling bridges the gap between white- and black-box modelling.

3.2.4 Incorporating uncertainty in grey-box models

Characterising uncertainty stemming from sensors, model approximations, and unrecognised disturbances is essential for implementing sensor-driven and real-time controllers. Grey-box models provide a useful framework for handling such uncertainties, as they can be formulated as SDEs (*stochastic differential equations*, i.e., differential equations that incorporate randomness). When the models are also formulated as SDEs, information from sensors can be assimilated into the model parameters in almost real-time. Consequently, the information obtained from sensors can be used to update the predictions and, consequently, for making better control decisions.

A grey-box model, as illustrated in Figure 3.3, provides a link between white-box models and black-box models. Grey-box models, which may be represented with resistive-capacitive networks, are typically formulated as state-space models, where the dynamics of the states are described in *continuous time* by a set of stochastic differential equations (SDEs) (system equations).

A grey-box model may be formulated as a continuous-discrete time stochastic state space model in the form:

$$dx(t) = \underbrace{f(x(t),u(t),d(t),t)dt}_{\text{Drift (deterministic component)}} + \underbrace{g(x(t),u(t),d(t),t)dw(t)}_{\text{Diffusion (random component)}} \tag{3.5}$$

$$y_k = h(x(t_k)) + v_k \qquad v_k \sim N(0, R_v) \tag{3.6}$$

In Equation 3.5, which describes the evolution of the states:

- x is the system vector (describing the states, e.g., temperatures)
- u and d are, respectively, the inputs and disturbances
- ω is a standard Wiener process (also called a Brownian motion)
- f and g are the *drift* and *diffusion* functions, respectively

The drift function is the *deterministic* part of the SDE. The diffusion function describes all the system-related *uncertainties* that the drift function does not adequately describe.

Equation 3.6 describes how discrete-time observations are related to the states, by a set of observation functions, which are static equations (i.e., without time-dependence), where:

- y_k are the observed outputs,
- h is the observation function
- v_k is the observation noise

3.2.5 Sources of uncertainty

Modelling physical systems using SDEs provides a natural method to represent the phenomenon as it evolves in continuous time[1]. The methodologies for readily handling uncertainty from sensors and other sources are one of the benefits of grey-box modelling when using stochastic differential equations. There are many reasons for introducing the system noise (the diffusion term):

- *Modelling approximations*. The dynamics, as described by the drift term, are necessarily an approximation of reality.
- *Unrecognised and un-modelled inputs*. These include variables which are not considered, such as wind speed, may affect the system.
- *Noise in measurements of input variables*: the measured input signals are regarded as the actual input to the system, and the deviation from the true input is described by the noise term.

In the observation equation, a noise term is also introduced. The reason for this noise term is:

- *Noise in measurements of output variables*. The sensors that measure the output signals are affected by noise and drift.

It can be reasonably assumed that the system and measurement noise are independent.

3.2.6 Examples of grey-box modelling with uncertainty

In Madsen and Holst (1995), a simple second-order grey-box model for the thermal dynamics of a building using the so-called RC-formulation is used. The thermal capacity is lumped into two states, and each of these states has an associated thermal mass. The states of the model are given by the temperature Tm of the dominant heat accumulating medium, and by the temperature Ti of the room air and possibly the inner part of the walls. This model has been identified in (Madsen & Holst, 1995). It is concluded that for the considered building, this second-order model provides a good description of all the variations in the data since the residuals are white noise; this paper describes how the parameters are estimated using a Maximum Likelihood method.

More recently, Thilker *et al.* (2021a) estimated an SDE-based model for a building using a water-based heating system with radiators to distribute the heat. The nonlinearities arise from the varying water flow and the radiator thermostats. Hou et al. (2022a) presents a nonlinear grey-box model and MPC of a university building in Norway. They also model the varying mass flow of the air and water.

[1] In a continuous-time formulation, variables evolve continuously over time and are described by differential equations. Discrete-time formulations describe how the system changes at successive time steps and are described by difference equations.

A linear, fourth-order RC model is identified and estimated in Zhang *et al.* (2022), where the system has an on-site battery with a distributed PV system. The sensitivity of grey-box model parameter estimates have been studied in Brastein *et al.* (2018). Here, they identify and estimate grey-box models for a small wooden building and find that the parameter estimates are robust towards initial conditions. An automated and scalable method for stochastic grey-box model identification of building heat dynamics is suggested by Leprince *et al.* (2022a). The method is implemented on 247 Dutch residential buildings Leprince *et al.* (2022b). This method was used by Vallianos et al. (2024) for the estimation of the model parameters for 60,000 homes in the USA and Canada.

Continuous-time models have many strengths concerning control:

- They describe the state-evolution of the system continuously. This is compared to discrete-time methods where the system state between samples is unknown.
- For control, continuous-time models are very flexible and can be used in many setups. First off, they allow operators to use controllers that minimises integral-cost functions. Secondly, continuous-time models can be discretised with arbitrary time samples to fit into any discretisation.
- It is important to emphasise that the physical interpretability of the grey-box models is essential in a control perspective, as it enhances their ability to extrapolate beyond the training data (which is paramount for control).

3.3 MPC Test Cases

To gain insight into the various approaches used to create predictive models for control applications, Subtask B participants contributed their different modelling approaches to different case studies. Table 3.1 summarises the different approaches, which respond to factors such as data availability, the need or not to generalise results, the incorporation of uncertainty, and familiarity with different modelling tools. The MPC projects compiled in Table 3.1 contain a wide array of control-oriented modelling approaches.

Table 3.1: Control-oriented modelling approaches by Annex 81 participants.

PROJECT	FEATURES	CONTROL-ORIENTED MODEL	INSTITUTION	COUNTRY
Varennes Library	NZEB with PV panels, geothermal heat pumps, radiant floor	Detailed grey-box models	Concordia University /CanmetENERGY-Varennes	Canada
Low energy building	MPC of PCM-based ventilation cooling	Grey-box	University of Southern Denmark	Denmark
Office building in a tropical climate	Office at Level # 3, research building in a tropical country	Black-Box (ANN)	Nanyang Technological University	Singapore
Small office building	Energy forecasting model for whole building	Black-box model with Energy+ as virtual testbed	Drexel University	USA
Large hospital central cooling plant	HVAC plant with thermal and electrical chillers	Hybrid: white-box for equipment, grey-box for building and storage, black-box forecasting model	CSIRO	Australia
Sol4City, EXCESS	Local control of room thermostats and valve positions	Grey-box with white-box for synthetic data	AEE-INTEC	Austria
Office-like test facility	Rule-based control algorithm for adaptive building envelope	White-box model using co-simulation	University College London (UCL), other institution	UK

Varennes Library, Canada. The Varennes Net-Zero Energy Library is a 2,100 m^2 building located in Varennes, Québec, Canada, approximately 30 km northeast of Montreal (Jalilov, 2021). Opened in 2014, it was designed the first net-zero institutional building in Canada, featuring a building-integrated photovoltaic thermal (BIPV/T) roof with a nominal electrical capacity exceeding 110 kW. Other features include 10 geothermal heat pumps, displacement ventilation and a carefully planned passive solar design. A detailed (10th order) *grey-box* RC thermal network was used to model an MPC strategy.

Small office building, USA. This is a 511 m^2 virtual building developed by NREL (National Renewable Energy Laboratory) in the USA. The virtual testbed was developed in EnergyPlus. The primary goal of the overall MPC is to reduce the building's total electricity costs over a forecasting horizon while ensuring thermal comfort by keeping the zone temperature within an acceptable range. The MPC strategy for this building uses a *black-box* model, which focuses on generating short-term electricity forecasts for the entire building using Multivariate Adaptive Regression Splines (MARS) (Zhang and Wen, 2019).

Echuca Hospital, Australia. This case study, which focuses on the largest building among the case studies considered, focuses on the cooling plant of a hospital (Heidari et al., 2024). This cooling plant includes a 4 MW chiller which provides the cold water for the air handling units distributed in the hospital. This case study includes an active thermal storage system includes cold water storage tanks that can be operated in charge, discharge, or standby modes; cooling towers that remove heat from the building's return water stream; and two solar thermal fields supplying thermal energy to drive two absorption chillers in the chiller plant. In this case, the MPC strategy applied a combination of modelling approaches: a *white-box* model for the chiller, complemented with operational data provided by the manufacturer, *grey-box* models for most of the equipment, and *black-box* models to forecast cooling load and solar thermal generation.

Low-Energy Building, Denmark. This building, denominated OU44, is an 8,500 m^2 teaching building located at Odense, in the University of Southern Denmark (Clausen et al., 2021). Some features of this building include a ventilation system with embedded phase change material. In the context of an MPC study, a virtual testbed was developed in the Modelica language, within the Dymola platform. To implement the MPC formulation, a grey-box model (RC network) was used. The calibration of the model was carried out using ModestPy. Training data included ambient air temperature, PCM temperature, damper position, valve positions, etc.

Dynamic building envelope in testing facility (MATELab), UK. This case study focused on a dynamic façade installed in a 30-m^2 office-like facility located in Cambridge, UK (Borkowski, 2021; Luna-Navarro et al., 2022). The south, east and west facades of this facility consist of modular components, which enable testing different glazing configurations. In this case study, the south-facing envelope consisted of two high-performance double-glazed units, and automated venetian blinds installed internally. In this case, a *white-box* model was used to model the impact of the façade on the test facility. For this purpose, this study used EnergyPlus for the building itself, a Modelica (Dymola) control algorithm, and the functional mock-up interface standard to exchange information between Dymola and EnergyPlus.

Office building in a tropical climate, Singapore. The test setup is an office located on the third floor of the Research Techno Plaza building at Nanyang Technological University in Singapore, covering 46 m^2 with a height of 3 m, equipped with eleven desks, a meeting table, and eight manually operated LED light fixtures (Yang et al., 2020). Climate control is handled by a ceiling-mounted VAV box with a 1320 m^3/h capacity, connected to an external fan coil unit that also serves nearby spaces, with a thermostat set to maintain 23°C during weekday working hours. Sensors in the supply duct and office area monitor air temperature and humidity in real-time, transmitting data to the control system, while a thermal comfort station near the meeting table collects data for offline analysis. In this case, a *black box* model (a recurrent neural network) is used to predict the indoor temperature of the office within an MPC approach targeting the real-time control of the cooling power supplied into the office. The RNN is continuously updated during real-time control.

Sol4City/EXCEES Projects, Austria. These projects focused on the application of MPC concepts to several net-zero and net-positive residential buildings incorporating solar energy (Putz et al., 2023). The buildings were simulated within an emulator created in the IDA-ICE software tool. The objective of the MPC model was to maximise the use of local solar generators (photovoltaic, photovoltaic thermal and solar thermal generators). Considering the diversity of buildings, both *grey-box* models (2nd-order RC thermal networks as the basic models, with additional nodes as per each case) and *black-box* models (ARX time series were used). The models and the controllers were created with the GEKKO tool, a Python optimisation package (Beal et al., 2018).

3.4 Conclusions

This chapter has briefly presented and discussed some methods used in building modelling and the modelling of building sections aimed at predictive control applications. The traditional classification of models into "white-box," "black-box," and "grey-box" categories was introduced. Each of these methods has its advantages and disadvantages. Generally, "grey-box" and "black-box" models, by their very nature, allow for a more immediate calibration of their parameters through data collection. That said, "white-box" models more easily facilitate the handling of new technologies or specific cases for which sufficient data has not yet been gathered.

A particular advantage of "grey-box" models is their ability to easily incorporate the treatment of uncertainty—be it in predictions, parameters, or the inevitable limitations that come with developing a mathematical model. Since a "perfect" model for control applications is simply unattainable, the ability not only to tolerate but even to embrace uncertainty allows for the rapid implementation of MPC in real buildings, as will be discussed in Chapter 4.

The examples presented by participants in Annex 81 are very diverse in nature, including real buildings or high-fidelity emulators, and range from small experimental buildings to large institutional or commercial buildings, and dealt with the control of different technologies (chillers, dynamic façades, heat pumps, etc.). "Grey-box" or "hybrid" models have been predominantly used by the researchers participating in Annex 81. It is worth noting that most Annex participants are researchers from government agencies or academic institutions. Information about the modelling approaches used in the commercial sector is limited due to confidentiality.

The case studies presented by the participants indicate that the most common practice for creating a "control-oriented" model remains addressing each case individually. Given the vast diversity of existing buildings and the near limitless variety of mechanical configurations, technologies, and building types to address, the large-scale deployment of data-driven MPC remains an outstanding challenge. However, the standardization of data structures and collection, combined with well-defined control practices and the selection of models suited to the specific task (specifically, determining the optimal control action over a defined horizon) will progressively simplify the development of control-oriented models grounded in real-world system data and facilitate the large-scale implementation of MPC solutions.

3.5 References

Afram, A., Janabi-Sharifi, F., Fung, A. S., & Raahemifar, K. (2017). Artificial neural network (ANN) based model predictive control (MPC) and optimisation of HVAC systems: A state-of-the-art review and case study of a residential HVAC system. *Energy and Buildings*, *141*, 96–113. https://doi.org/10.1016/j.enbuild.2017.02.012

Beal, L. D., Hill, D. C., Martin, R. A., & Hedengren, J. D. (2018). Gekko optimisation suite. *Processes*, *6*(8), 106.

Box, G. E. (1976). Science and statistics. *Journal of the American Statistical Association*, *71*(356), 791-799.

Borkowski, E. (2021). *Integrated modelling of control and adaptive building envelope: development of a modelling solution using a co-simulation approach* (Doctoral dissertation, University College London).

Brastein, O., Perera, D., Pfeifer, C., & Skeie, N. O. (2018). Parameter estimation for grey-box models of building thermal behaviour. *Energy and Buildings, 169,* 58–68. https://doi.org/10.1016/j.enbuild.2018.03.057

Candanedo, J. A., Vallianos, C., Delcroix, B., Date, J., Saberi-Derakhtenjani, A., Morovat, N., John, C. & Athienitis, A. K. (2022). Control-oriented archetypes: a pathway for the systematic application of advanced controls in buildings. *Journal of Building Performance Simulation*, *15*(4), 433-444.

Clausen, A., Arendt, K., Johansen, A., Sangogboye, F. C., Kjærgaard, M. B., Veje, C. T., & Jørgensen, B. N. (2021). A digital twin framework for improving energy efficiency and occupant comfort in public and commercial buildings. *Energy Informatics*, *4*, 1-19.

Ciulla, G., & D’Amico, A. (2019). Building energy performance forecasting: A multiple linear regression approach. *Applied Energy*, *253*, 113500. https://doi.org/10.1016/j.apenergy.2019.113500

Deb, C., Zhang, F., Yang, J., Lee, S. E., & Shah, K. W. (2017). A review on time series forecasting techniques for building energy consumption. *Renewable and Sustainable Energy Reviews*, *74*, 902–924. https://doi.org/10.1016/j.rser.2017.02.085

Dong, B., Cao, C., & Lee, S. E. (2005). Applying support vector machines to predict building energy consumption in tropical region. *Energy and Buildings*, *37*(5), 545–553. https://doi.org/10.1016/j.enbuild.2004.09.009

Fan, C., Wang, J., Gang, W., & Li, S. (2019). Assessment of deep recurrent neural network-based strategies for short-term building energy predictions. *Applied Energy*, *236*, 700–710. https://doi.org/10.1016/j.apenergy.2018.12.004

Heidari, R., Dioguardi, E., Sethuvenkatraman, S., & Braslavsky, J. H. (2024). Evaluating advanced HVAC control benefits in operational buildings using historic data—A case study. *Applied Thermal Engineering*, 123611.

Hou, J., Li, H., & Nord, N. (2022a). Nonlinear model predictive control for the space heating system of a university building in Norway. *Energy, 253,* 124157. https://doi.org/10.1016/j.energy.2022.124157

Jalilov, E. (2021). *Development of heuristic model-based predictive control strategies for an institutional net-zero energy building* (Master’s Thesis, Concordia University).

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436–444. https://doi.org/10.1038/nature14539

Leprince, J., Madsen, H., Miller, C., Real, J. P., van der Vlist, R., Basu, K., & Zeiler, W. (2022a). Fifty shades of grey: Automated stochastic model identification of building heat dynamics. *Energy and Buildings, 266,* 112095. https://doi.org/10.1016/j.enbuild.2022.112095

Leprince, J., Miller, C., Madsen, H., Basu, K., Van Der Vlist, R., & Zeiler, W. (2022b). Grey-brick buildings, an open data set of calibrated RC models of Dutch residential building heat dynamics. *Proceedings of the 20th ACM Conference on Embedded Networked Sensor Systems (Sensys ’22),* 1067–1071. https://doi.org/10.1145/3560905.3567760

Li, Y., O'Neill, Z., Zhang, L., Chen, J., Im, P., & DeGraw, J. (2021). Grey-box modeling and application for building energy simulations-A critical review. *Renewable and Sustainable Energy Reviews*, *146*, 111174.

Macarulla, M., Casals, M., Forcada, N., & Gangolells, M. (2017). Implementation of predictive control in a commercial building energy management system using neural networks. *Energy and Buildings*, *151*, 511–519. https://doi.org/10.1016/j.enbuild.2017.06.027

Maddalena, E. T., Müller, S. A., Dos Santos, R. M., Salzmann, C., & Jones, C. N. (2022). Experimental data-driven model predictive control of a hospital HVAC system during regular use. *Energy and Buildings*, *271*, 112316. https://doi.org/10.1016/j.enbuild.2022.112316

Madsen, H., Holst, J., & Lindström, E. (1995). Estimation of continuous-time models for the heat dynamics of a building. *Energy and Buildings, 22,* 67–79.

Miller, C., Arjunan, P., Kathirgamanathan, A., Fu, C., Roth, J., Park, J. Y., Balbach, C., Gowri, K., Nagy, Z., Fontanini, A. D., & Haberl, J. (2020). The ASHRAE Great Energy Predictor III competition: Overview and results. *Science and Technology for the Built Environment*, *26*(10), 1427–1447. https://doi.org/10.1080/23744731.2020.1795514

Mtibaa, F., Nguyen, K.-K., Azam, M., Papachristou, A., Venne, J.-S., & Cheriet, M. (2020). LSTM-based indoor air temperature prediction framework for HVAC systems in smart buildings. *Neural Computing and Applications*, *32*(23), 17569–17585. https://doi.org/10.1007/s00521-020-04926-3

Putz, D., Gumhalter, M., & Auer, H. (2023). The true value of a forecast: Assessing the impact of accuracy on local energy communities. *Sustainable Energy, Grids and Networks*, 33, 100983.

Thilker, C. A., Bacher, P., Bergsteinsson, H. G., Junker, R. G., Cali, D., & Madsen, H. (2021a). Non-linear grey-box modelling for heat dynamics of buildings. *Energy and Buildings, 252,* 111457. https://doi.org/10.1016/j.enbuild.2021.111457

Vallianos, C., Candanedo, J., & Athienitis, A. (2024). Thermal modeling for control applications of 60,000 homes in North America using smart thermostat data. *Energy and Buildings*, *303*, 113811. https://www.sciencedirect.com/science/article/pii/S0378778823010411

Zhang, L., & Wen, J. (2019). A systematic feature selection procedure for short-term data-driven building energy forecasting model development. *Energy and Buildings*, *183*, 428-442.

Zhang, K., Prakash, A., Paul, L., Blum, D., Alstone, P., Zoellick, J., Brown, R., Pritoni, M. (2022). Model predictive control for demand flexibility: Real-world operation of a commercial building with photovoltaic and battery systems. *Advances in Applied Energy, 7,* 100099. https://doi.org/10.1016/j.adapen.2022.100099

Yang, S., Wan, M. P., Chen, W., Ng, B. F., & Dubey, S. (2020). Model predictive control with adaptive machine-learning-based model for building energy efficiency and comfort optimisation. *Applied Energy*, *271*, 115147.

Yu, Z., Haghighat, F., Fung, B. C. M., & Yoshino, H. (2010). A decision tree method for building energy demand modeling. *Energy and Buildings*, *42*(10), 1637–1646. https://doi.org/10.1016/j.enbuild.2010.04.006

4. Activity B3: Evaluation of Data-Driven Control Strategies in a Closed-Loop Environment

4.1 Background

IEA Annex 81 presented an opportunity to gather an international community of data-driven control experts and developers to compare closed-loop performance of various data-driven control approaches proposed by each participant. Activity B3 focused on simulation-based performance evaluation, where data-driven controllers are tested against high-fidelity simulation models of building HVAC systems meant to stand in for real buildings. Such models are also known as building *emulators*. This approach is common in the literature and has been chosen over tests in actual buildings since simulations can be run faster than real time, are more flexible in the definition of the building, system, and data access, and have completely controllable and replicable operating conditions. At the same time, recent developments in simulation technology have allowed for significantly improved opportunities for simulation models to represent real buildings realistically.

Two prominent categories of data-driven control strategies explored within Activity B3 are Model Predictive Control (MPC) and Reinforcement Learning (RL). Though recently brought into more focus, MPC has been studied for many decades in the context of building control, dating back to initial work in the early 1990s within mechanical and process engineering communities. RL is a relatively newer approach, especially in building control, focusing on machine learning techniques developed mainly within data and computer science communities. There exist differences between MPC and RL approaches, and there also exist further differences between specific implementations of each regarding modelling approaches, algorithm choices, and related additional tuning of each of the components of the approaches. Meanwhile, researchers have contributed numerous studies demonstrating the value and advantages of various approaches over chosen baselines to compare against, with comparative baselines often being conventional rule-based control. Though valuable in their own right for identifying promising approaches, individualised case studies are difficult to compare and contrast in broader contexts that would identify overarching conclusions of advantages and disadvantages of approaches and variations for given applications or conditions. Such conclusions are key to reducing the risk for industry adoption and owner investment and guiding continued development

4.2 Objectives

This Section presents the activities of B3 participants in their evaluation of data-driven control strategies in a closed-loop environment and efforts to draw conclusions or observe trends, where possible, about the performance of the various control strategies employed. The rest of the report is structured as follows. Section 4.3 overviews methods used within the B3 studies. Next, Section 4.4 summarises five studies undertaken and their key results. Section 4.5 presents results that are comparable across studies. Finally, Section 4.6 summarises the conclusions of B3 and describes the needs for future work.

4.3 Simulation Methods

To make results generated by one Annex participant comparable to another, and to enable such results to be continually used as baselines for future studies, a test environment is required that separates the virtual building simulation (emulator) and KPI generation from the controller under test, to isolate any change in building and KPI results such that they can be attributed only to the controller under test. A visualisation of such a test environment is shown in Figure 4.1 below.

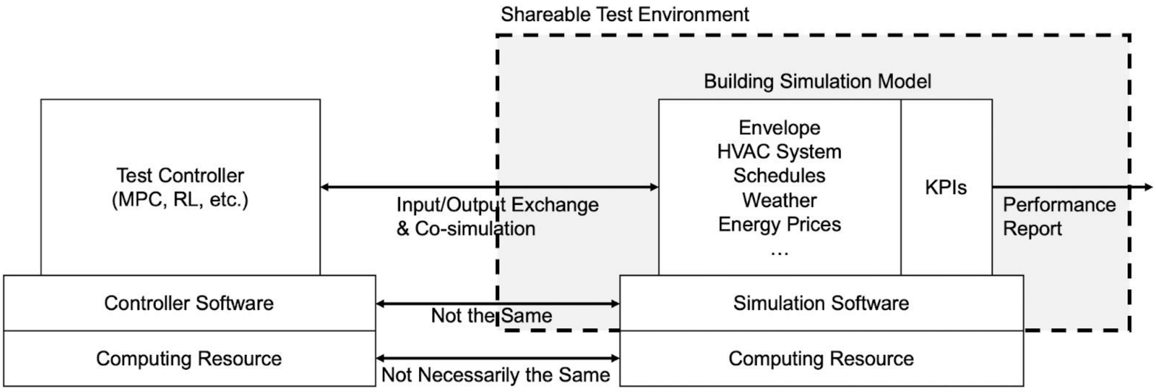

Figure 4.1: Elements of a test environment for benchmarking control performance.

Note that the test environment encapsulates the building simulation model with any boundary condition data and parameters needed to run the model, along with the simulation software and solvers needed to integrate the simulation model through time. The test environment also uses data generated by the simulation model and the necessary software to calculate KPIs unambiguously concerning the user or test controller. Finally, the environment has an API which the user can use to manage tests and input/output data between the test controller and building simulation model.

The environment depicted in Figure 4.1 has been developed as the Building Optimisation Testing Framework (BOPTEST), described in detail in Blum *et al.* (2021). See also the homepage at https://boptest.net. BOPTEST consists of a rapidly deployable and accessible run-time environment (RTE), based in Python and packaged and deployed using Docker, that supports simulating building simulation models in the form of Functional Mockup Units (FMU), calculating KPIs using stored simulation data, and exposing a RESTful HTTP API for users to manage simulations, report KPIs, and interact with the simulated building using a test controller using input/output data. The API also enables predictive controllers to retrieve forecasts of boundary conditions (e.g. weather and electricity prices) and users to retrieve a history of measurement and control data trajectories.

BOPTEST also makes building simulation models (*emulators*) publicly available. These emulators (denominated test cases) are high-fidelity models written in Modelica and using open-source Modelica libraries extending from the Modelica IBPSA Library (https://github.com/ibpsa/modelica-ibpsa), such as Modelica Buildings Library (Wetter *et al.*, 2014) and IDEAS (Jorissen *et al.*, 2018). Test cases may also use Spawn (Wetter *et al.*, 2024) to integrate envelope models written for EnergyPlus with HVAC and control models written in Modelica. Note that BOPTEST uses Modelica-based HVAC and control simulation so that the HVAC and control dynamics are represented explicitly, making the virtual buildings, and associated input/output behaviour, perform as close to the real world as possible.

Each test case contains embedded baseline control, so that test controllers can overwrite any subset of control signals at the supervisor or actuator levels. BOPTEST currently offers a selection of seven publicly available test case emulators. Additional development of framework capabilities and new test cases are ongoing within IBPSA Project 2. For more information, see https://ibpsa.github.io/project1-boptest/ibpsa/index.html.

4.4 Studies Conducted in BOPTEST

4.4.1 Study 1 – Using an RC network as a control-oriented model for MPC

This study implemented a three-state linear RC model as a control-oriented model in an MPC framework. The principal layout of the control model is shown in Figure 4.2 below. A more detailed description of the MPC implementation is provided by Walnum *et al.* (2020).

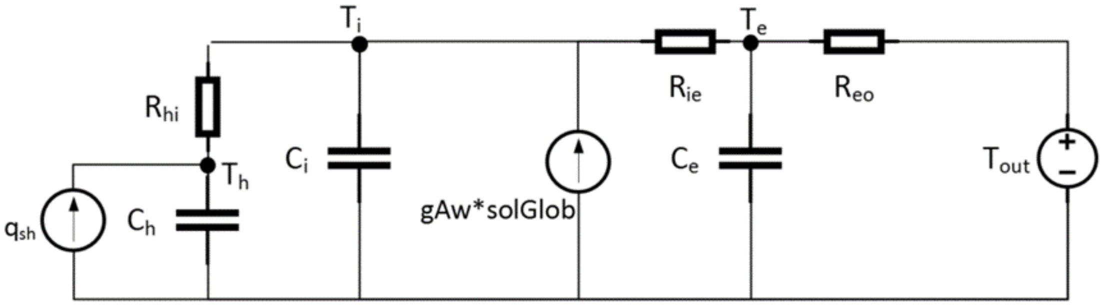

Figure 4.2: The three-state linear RC model used in the MPC framework, as described in Walnum et al. (2020).

To assess the MPC framework's performance and scalability across various building types, it was tested on two BOPTEST cases:

- *bestest_hydronic_heat_pump*
- *singlezone_commercial_hydronic*

The *bestest_hydronic_heat_pump* test case is a single-zone residential radiant floor heating system served by an air-to-water heat pump. In contrast, the singlezone_commercial_hydronic test case is a single-zone approximation of a commercial office building with radiator heating and a ventilation system delivering outside air. More details about these test cases can be found on the BOPTEST home page.

In both cases, room temperature setpoint is the control variable, as it is commonly accessible in commercial heating systems. The implementation approach is nearly identical: models are trained on one month of baseline operation data (excluding test periods), using only forecasts of outdoor air temperature and global horizontal solar radiation, as these inputs are typically available in real buildings.

For the *singlezone_commercial_hydronic* case, ventilation is excluded from the MPC formulation. In both setups, the MPC algorithm determines the optimal heat flow into the building, which is then "translated" into a temperature setpoint. This conversion is tuned separately for each case due to the slower thermal response of the radiant floor system compared to the radiator system. This conversion must be tuned slightly differently for the two applications, as the floor heating system in the *bestest_heat_pump_hydronic* test case reacts much slower than the radiator system in the *singlezone_commercial_hydronic* test case.

For each test case, the controllers were evaluated using BOPTEST's predefined two-week heating scenarios (*typical_heat_day* and *peak_heat_day*) along with all three predefined pricing schemes: constant, dynamic, and highly dynamic. The MPC setup was tested using both 15-minute and 1-hour control intervals.

Overall, the results indicate that MPC outperforms the baseline BOPTEST controllers, with the 15-minute control step yielding better KPI performance than the 1-hour step. However, a closer look (especially for the *bestest_hydronic_heat_pump* test case) reveals that the 15-minute control step leads to significant oscillations in the setpoints, which may increase wear on the equipment. Directly controlling the compressor speed could probably reduce the oscillation effect. However, this is typically unavailable in commercial heat pumps, or one could risk breaching the manufacturer's warranty. A key functionality of predictive controllers is their ability to adapt to variable prices by shifting loads in time. However, price variability in the test cases is too limited to assess this effect.

4.4.2 Study 2 – Comparing RL and MPC

Deep Reinforcement Learning (RL) and Model Predictive Control (MPC) are widely employed optimisation-based control strategies for enhancing building energy performance. Numerous studies have investigated the control performance of RL and MPC individually, but comparing their performance directly is challenging due to the absence of studies conducted in the same environment. This study, detailed in Wang *et al.*, 2023, implemented RL and MPC controls and compared their performance with traditional rule-based controls using closed-loop simulation with the BOPTEST *bestest_hydronic_heat_pump* test case under the *typical_heat_day* and *peak_heat_day* time-period scenarios, and *highly_dynamic* pricing scenario. This approach allowed for a fair comparison between RL and MPC.

Three widely used RL algorithms, Deep Deterministic Policy Gradient (DDPG), Dueling Deep Q-Networks (DDQN), and Soft Actor-Critic (SAC), were used to develop the RL controllers. Given DRL's sensitivity to hyperparameter choices, we performed thorough tuning using Optuna.

The MPC controller was developed using a reduced-order (1R1C) thermal RC network model for the building envelope and a regression-based heat pump performance model. The MPC controller was implemented using a 12-hour prediction horizon and a 60-minute control step. The building envelope and heat pump models are trained using a two-week dataset of normal operational conditions under the control of the baseline controller and outside the control test period. Given the data-intensive nature of training a DRL agent, we employed an extended training period where the DRL agent interacted with BOPTEST throughout the heating period, excluding the two control testing periods.

The performance of DRL and MPC controllers is evaluated in terms of control effectiveness, data efficiency, implementation effort, and computational requirements. All RL agents outperformed the baseline controller in the *peak_heat_day* scenario, but only DDPG surpassed the baseline in the *typical_heat_day* scenario. Specifically, the DDPG agent exhibited a 92% reduction in thermal discomfort[2] and an 11% operational cost saving during the *peak_heat_day* scenario, and an 18% reduction in thermal discomfort and a 14% cost saving during the *typical_heat_day* scenario. The MPC controller outperformed the baseline controller with a 99.8% reduction in discomfort and a 21.9% cost saving in the *peak_heat_day* scenario, and a 12% reduction in discomfort and a 26.6% cost saving in the *typical_heat_day* scenario.

Comparing the MPC with the best-performing RL controller (DDPG), the MPC demonstrated superior performance. During the *peak_heat_day* scenario, the MPC nearly eliminated thermal discomfort and achieved a 12% electricity cost saving compared to the RL controller. During the *typical_heat_day* scenario, the MPC achieved a 14% electricity cost saving but resulted in an 8% higher thermal discomfort. These findings indicate that the model-based approach of MPC is more effective in identifying an optimal solution than the model-free approach of RL (DDPG). Regarding computational demands, the RL and MPC controllers required higher online computational resources than the baseline controller. The RL controller necessitated more computational resources during offline training, while the MPC controller required higher online computational demand for real-time optimisation at each control time step.

Future research endeavours should focus on comparing advanced controllers in real buildings, considering the challenges posed by insufficient operational data for training RL controllers, as well as the uncertainties associated with measurement and disturbance forecasts for both RL and MPC controllers.

4.4.3 Study 3 – Deriving surrogate models from BOPTEST emulators

Implementing MPC at scale is difficult due to the high cost of modelling. To address this, Mostafavi et al. (2023) propose using machine learning and automatic differentiation (an efficient method for computing derivatives) to formulate and solve nonlinear HVAC MPC problems in buildings, enhancing scalability and robustness. The study tests this approach on two BOPTEST cases: the complex *multizone_office_simple_air* scenario (a 5-zone VAV system with terminal box reheat) and the simpler *bestest_air* case (a single-zone fan coil unit system).

For surrogate modelling (i.e. simplified models used in place of more detailed models, typically used to facilitate the implementation of model-based control strategies)[3], PyTorch was used to implement and evaluate three architectures: Linear, Multilayer Perceptron (MLP), and Long Short-Term Memory (LSTM) models (Table 4.1). To ensure robustness and accuracy, a train-validation-test split was adopted. The validation set was used to tune the models during training, while the test set assessed final performance on unseen data.

[2] Thermal discomfort quantifies the total deviation of the zone's operative temperature beyond the defined comfort limits (upper and lower) over a specified period, expressed in degree-hours (K·h).
[3] In this context, surrogate models play the role of control-oriented models, facilitating the formulation of optimisation problems.

Training data was generated by sampling random control inputs at each time step and passing them to the BOPTEST simulation, which returned the resulting observations and disturbances. In other words, data from a detailed model (BOPTEST emulators) was used to train the surrogate (i.e., control-oriented) models. This process was repeated over multiple trajectories with varying initial conditions.

- In the *bestest_air* case, 120 trajectories of 500 time steps (15-minute intervals) were generated, with 100 used for training, 10 for validation, and 10 for testing.
- In the *multizone_office_simple_air* case, 600 trajectories of 1000 time steps were generated, with 500 for training, 50 for validation, and 50 for testing.

Table 4.1: Mean Squared Error (MSE) for different model choices in the zone temperature bestest_air test case (Mostafavi et al., 2023).

Model	Train MSE (×10^-5)	Val MSE (×10^-5)	Test MSE (×10^-5)
Linear	699.5	566.8	780.3
MLP	8.846	12.70	17.56
LSTM	1.418	1.726	2.145

Nonlinear MPC problems were formulated as receding-horizon controllers with a 10-step (15-minute) look-ahead, computing setpoints for low-level HVAC control loops. Optimisation was performed using Sequential Quadratic Programming (SQP) and Gradient Descent Methods (GDM). SQP followed Dieter Kraft's implementation in SciPy, using a quasi-Newton approach with automatic differentiation for function and Jacobian evaluations. The GDM involved updating decision variables iteratively, using projected gradient descent for input constraints. This framework ensured efficient, accurate optimisation, enabling effective HVAC control while meeting all constraints.

Performance was evaluated over two one-week long periods (January 4-11 and August 7-14). In the single-zone case, LSTM models yielded the best results; in the multi-zone case, MLP models performed better, illustrating the trade-off between model complexity and accuracy.

This approach (i.e. obtaining surrogate models from the emulators) successfully reduced cooling and heating loads while maintaining occupant comfort. It is adaptable to various models and control algorithms, providing a robust solution for real-world applications. For the two weeks, the KPIs focused on total energy use, thermal discomfort, and MPC computation time (MacBook Pro 2017, 16GB RAM) across different model and solver combinations. Table 4.2 below summarizes the results. Discomfort is given in terms of Kelvin-hours (Kh).

Table 4.2: KPIs for bestest_air (single-zone) and multizone_office_simple_air (multi-zone) test cases, with testing periods from January 4 to January 11 and August 7 to August 14 (Mostafavi et al., 2023).

Scenario	Model	Solver	Energy (kWh/m2)	Computation Time (sec)	Discomfort (Kh/zone)
Single-zone	Linear	GDM	0.0189	1.607	1556
Single-zone	Linear	SLSQP	0.2551	0.933	1528
Single-zone	MLP	GDM	4.804	1.694	2.935
Single-zone	MLP	SLSQP	5.059	1.684	5.207
Single-zone	**LSTM**	**GDM**	**4.818**	**0.620**	**2.081**
Single-zone	LSTM	SLSQP	4.943	0.661	4.415
Multi-zone	Linear	GDM	2.807	1.504	10.44
Multi-zone	Linear	SLSQP	2.487	1.600	11.40
Multi-zone	MLP	GDM	3.458	1.782	4.054
Multi-zone	**MLP**	**SLSQP**	**2.778**	**2.144**	**3.154**
Multi-zone	LSTM	GDM	2.222	0.570	124.7
Multi-zone	LSTM	SLSQP	2.880	0.818	35.48

This study demonstrated that accurate ML models can be derived from building emulators, though data generation remains challenging, particularly for complex systems. Additionally, automatic differentiation can significantly streamline the formulation process, allowing for more efficient and generalised solutions without custom designs. Results indicate that model complexity can affect control performance, and predictive accuracy should not be the only metric.

Achieving practical MPC necessitates further tests and improved control formulations. Future research will explore fast solvers and alternative control formulations, collaborating with BOPTEST developers to control scaled-up models, including multiple buildings and incorporating robust sampling procedures for data generation. The ultimate goal is to develop scalable, robust predictive control methods.

4.4.4 Study 4 – Incorporating uncertainty in MPC and forecast information into RL training

This study (a two-part study) investigates the impact of time series prediction results on MPC and RL control algorithms. For MPC, the effectiveness of optimising control strategies is highly dependent on the accuracy of the prediction model. However, any prediction model inherently contains uncertainties.

Uncertainty propagation. The first research objective of this study is to explore how the MPC algorithm can effectively mitigate the impact of prediction uncertainties to produce more robust operational strategies. Towards this end, using a first-order thermal RC network grey-box model (1R1C), a “tube-based” [4] MPC algorithm was used to compare the optimisation results under prediction uncertainty for a single peak heating day in the BOPTEST *bestest_air* test case (this test case represents a single-zone FCU system where hot water is supplied by a boiler and chilled water by a chiller). Since HVAC prediction accuracy declines over time, we modified the tube-based MPC algorithm to let the uncertainty space gradually expand along the prediction horizon. This better captures the accumulation of errors and ensures more robust and realistic control. Compared to traditional MPC approaches, the tube-based MPC can reduce operational costs by up to 42.6% and 48.8% under outdoor temperature prediction errors of 10% to 20%, while achieving better indoor temperature control performance. However, each MPC demonstrated issues with maintaining the indoor temperature set point, which should be considered further. More details about this study can be found in (Gao *et al.*, 2023).

Forecast information into RL training. The second research objective focuses on how forecast information can be integrated into reinforcement learning (RL) to improve control performance. In this context, an RL agent learns to make decisions, such as when to heat or cool, based on the system's current state. This study examines whether adding predictive information about future conditions to the agent’s state vector can support more effective decision-making. The Soft Actor-Critic (SAC) algorithm, which applies the maximum entropy principle to promote exploration and avoid local optima, served as the foundation. Modifications included replacing more common neural network techniques with recurrent neural networks (RNN), specifically, Gated Recurrent Units (GRUS), and incorporating forecast data into the state representation. The enhanced RL configuration was evaluated against the traditional SAC setup across four BOPTEST *bestest_air* scenarios: *peak_heat_day*, *typical_heat_day*, *peak_cool_day*, and *typical_cool_day*. Results show that integrating RNN with forecast inputs in a DRL formulation can improve thermal comfort by up to 46% and reduce operational costs by as much as 20%. Gao et al. (2024) provide additional details.

4.4.5 Study 5 - How close to optimal are RL and MPC control strategies?

Both model predictive control (MPC) and deep reinforcement learning control (DRL) have been presented as pathways to approximate the true optimality of a dynamic programming formulation[5]. These two methods have shown significant operational cost-saving potential for building energy systems. However, thorough quantitative evaluations of their closeness to true optimality are still limited, especially in the building sector.

[4] Tube-based MPC is a robust control approach that accounts for uncertainty in future system inputs or disturbances by maintaining the predicted state trajectories within a bounded “tube” (allowable deviations) around a nominal trajectory over the prediction horizon
[5] Dynamic programming, based on Bellman’s equation, guarantees global optimality under certain conditions, such as full knowledge of the system and a discrete set of possible actions.

To address this research gap, this study, detailed in (Fu *et al.*, 2023), developed a *containerised framework* (separate, but similar to BOPTEST) that enables evaluation of the optimality levels of different controllers for building energy systems, mainly focusing on optimal control and learning-based control for standardised building environments. This framework, shown in Figure 4.4: KPI results from BOPTEST for the *peak_heat_day* (top) and *typical_heat_day* (bottom) time-period scenarios and *highly_dynamic* electricity price scenario for the *bestest_hydronic_heat_pump* test case of BOPTEST for different data-driven control strategies.

, leverages the development of a Functional Mock-up Unit (FMU) for generic dynamic modelling and simulation, which can be flexibly extended to integrate with user-specific simulation environments that support FMU exporting. The framework is open-source and available from: https://github.com/BE-HVACR/FMU-DRL-DOCKER.

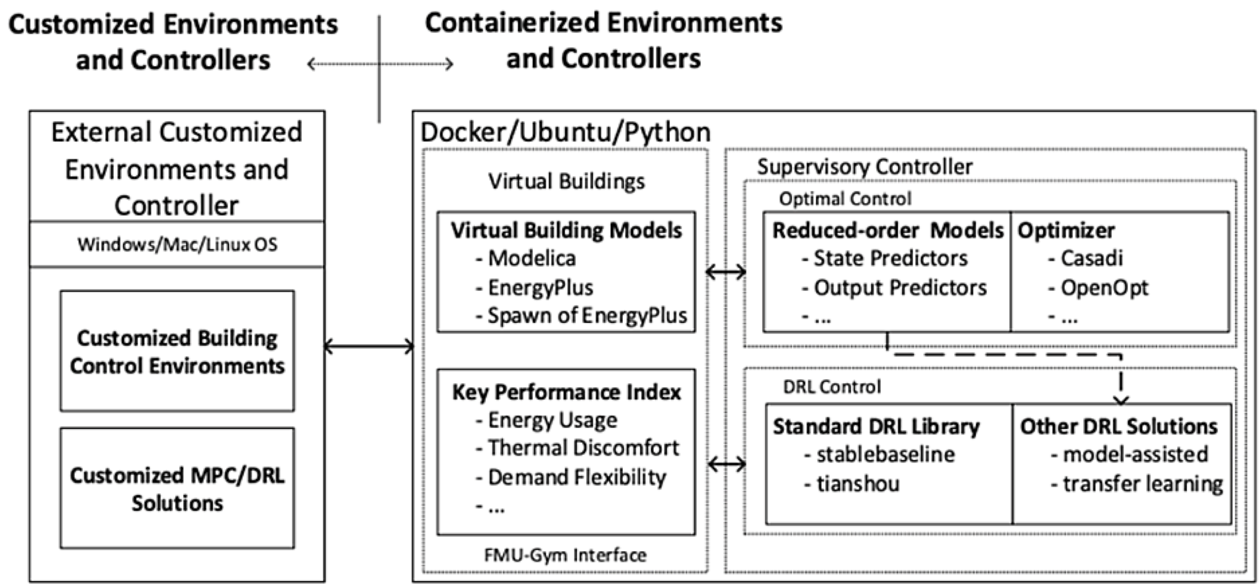

Figure 4.3: The proposed flexible containerised framework for evaluations of building controllers (Fu et al., 2023).

This framework was used to compare the optimal control performance of both MPC and DRL controllers with given computation budgets for a single zone fan coil unit system. The performance of different controllers with DRL agents during 1-week testing in July, using Chicago, IL (USA) weather data, is summarised in Table 4.3. Agents were trained in a different 1-week period from testing. Optimality is calculated according to an objective function based on a user-specific selection of trade-off weights among *energy costs, thermal comfort, and control slew rates* (i.e., the rate at which the control signal changes over time).

Compared to the best achievable performance (OPT), obtained through costly dynamic programming (DP), the top-performing deep reinforcement learning (DRL) agent can reach up to 96.54% of the optimality. This surpasses the best model predictive control (MPC) solution, which achieves 90.11%. However, due to the stochastic nature of the training process, DRL agents are expected to achieve, on average, 90.42% of the optimality, similar to the best MPC performance. Except for the Proximal Policy Optimisation (PPO) agent, all DRL agents demonstrate better potential for approximating the optimal solution than the best MPC. They also outperform MPC configured with a 32-step (H=32) prediction horizon (15 minutes per step), although the MPC results improve with longer prediction horizons.

When evaluating practical performance metrics like energy cost and thermal discomfort, MPC outperforms rule-based control (RBC) by 18.5% to 25.4%, while DRL is expected to offer gains of 18.9% to 25.7%. The best DRL policy achieves improvements from 20.3% to 29.7% over RBC.

Although these comparisons assume idealised conditions (e.g., perfect models for MPC and flawless offline training and deployment for DRL), they provide valuable insight into each method’s ability to approximate the true optimal solution defined by the DP benchmark.

Table 4.3: Control performance of different controllers (Fu et al., 2023). H corresponds to the number of time steps in the prediction horizon of MPC strategies.

Agent	MPC or RL Scenario	Energy [kWh]	Energy Cost [$]	Total Thermal Discomfort [K·h]	Maximum Temperature Violations [K]	Action Changes [-]	Episodic Rewards [-]
RBC	-	57.37	7.16	2.0	2.32	1.84	-746.94
MPC	H=4	47.08	5.9	14.8	0.5	1.4	-609.0
	H=8	47.14	5.9	13.14	0.38	1.4	-606.33
	H=16	47.1	5.8	11.31	0.5	1.43	-595.52
	H=32	50.92	5.58	9.48	0.47	3.24	-578.38
	H=48	56.38	5.31	8.0	0.56	5.38	-559.35
	H=96	63.9	5.2	11.68	1.21	3.24	-556.95
DRL	DDQN	58.22	5.27	12.0	0.96	3.45	-563.97
	QRDQN	64.0	5.11	9.68	0.81	4.34	-549.47
	PPO	46.92	5.78	18.24	0.92	2.32	-613.68
	SAC	56.16	5.72	3.43	0.59	2.04	-580.71
OPT	-	66.08	4.92	7.51	0.68	3.35	-506.82

Notes
• Rule-based control (RBC): Fan speed is controlled by a PI controller to track the zone temperature setpoint • "True" Optimal Control (OPT): Optimise over the whole simulation period using the emulator model; • Model Predictive Control (MPC): Optimise recedingly over the prediction horizon using the emulator model; • Deep Reinforcement Learning (DRL): Optimise future accumulative returns. • Acronyms: DQN: Double deep Q-network; QRDQN: Quantile regression deep Q-network; PPO: Proximal policy optimisation; SAC: Soft actor-critic

This study assumed perfect training and deployment processes for both MPC and DRL, which can shed insight into their possible maximum performance. While the control-oriented models used in MPC are assumed to be free of modelling errors, the optimisation process itself is not. The resulting optimisation problem is inherently nonlinear (and often non-convex), thus making it difficult to solve optimally. Since no "global" optimiser can guarantee an optimal solution within practical computational limits, the MPC's performance may still fall short of its theoretical best. Evaluating upper bounds for nonlinear MPC performance remains challenging without stepwise optimality, making it an open research topic.

Future work will compare both controllers in a realistic setting, with MPC using error-prone reduced-order models and DRL adopting non-episodic training for imperfect environments

4.5 Discussion: Direct Comparisons Across Studies

Based on the studies performed, we can compare the MPC and RL approaches from Study 1 and Study 2 directly with each other and with a benchmark from a previous study (Arroyo et al., 2022). These studies tested control approaches on the *bestest_heat_pump_hydronic* test case during the *peak_heat_day* and *typical_heat_day* time-period scenarios and the *highly_dynamic* electricity price scenario.

In Table 4.4 and Figure 4.4 below, the controllers are categorised as MPC or RL, along with variations in the control step size taken by each MPC. The Baseline approach is the baseline controller embedded in the BOPTEST test case.

Table 4.4: Comparison of data-driven approaches for the bestest_heat_pump_hydronic test case in BOPTEST.

bestest_heat_pump_hydronic		**Peak Heat Day Period**		**Typical Heat Day Period**	
Highly Dynamic Electricity Price		Operational Costs (EUR/m2)	Thermal Discomfort (K·h/zone)	Operational Costs (EUR/m2)	Thermal Discomfort (K·h/zone)
Baseline		0.91	8.38	0.41	9.44
Benchmark (Arroyo et al., 2022)	RL (DDQN)	0.82	2.80	0.51	180.84
	MPC 15 min Step	0.66	1.15	0.28	7.24
	MPC 60 min Step	0.76	2.67	0.30	7.06
Study 1 [Walnum et al., 2020]	MPC 15 min Step	0.80	0.02	0.37	6.39
	MPC 60 min Step	0.82	1.15	0.34	7.99
Study 2 [Wang et al., 2023]	RL (DDPG)	0.81	0.87	0.35	7.73
	MPC 60 min Step	0.71	0.00	0.31	8.29

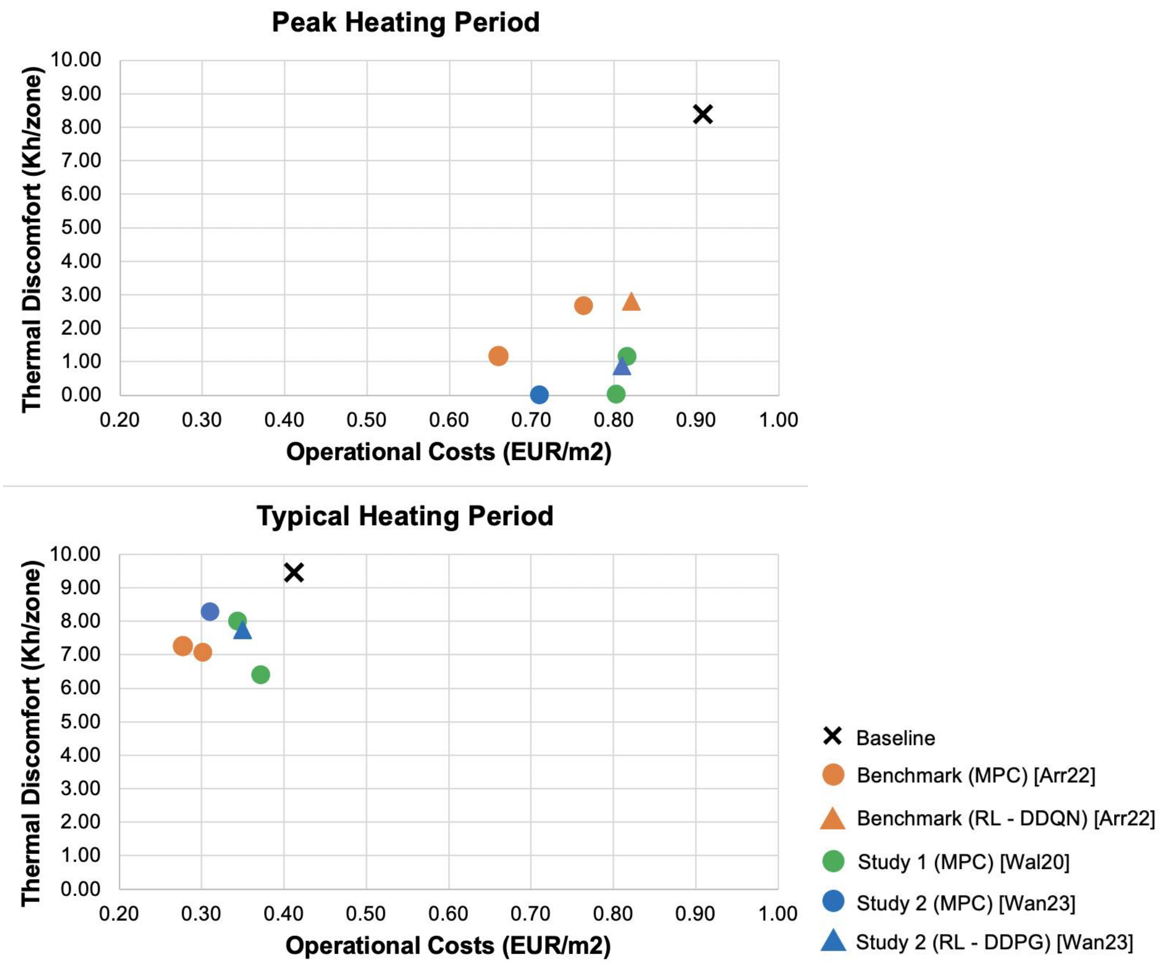

Figure 4.4: KPI results from BOPTEST for the *peak_heat_day* (top) and *typical_heat_day* (bottom) time-period scenarios and *highly_dynamic* electricity price scenario for the *bestest_hydronic_heat_pump* test case of BOPTEST for different data-driven control strategies.

A direct comparison of results from Study 1, Study 2, and (Arroyo et al., 2022) on the *bestest_hydronic_heat_pump* test case in BOPTEST shows that the most effective MPC strategies (those with short control steps and long prediction horizons) outperformed the best RL methods (DDPG and DDQN) used in those studies. This aligns with findings from both (Arroyo et al., 2022) and (Wang et al., 2023).

It is worth noting that the MPC models varied: Study 1 used a three-state RC network, Study 2 used a one-state RC network, and (Arroyo et al., 2022) used a five-state RC model. Meanwhile, the MPC controllers in Study 1 used the MPC solution to control zone temperature set points, while the MPC controllers in Study 2 and (Arroyo et al., 2022) directly controlled the heat pump compressor speed.

Study 5 used a different environment and scenarios than BOPTEST, though it used a similar model as the BOPTEST *bestest_air* test case. It found that the best MPC approach applied in the study performed approximately as well as the best RL approach on average. However, RL training involves a randomised process, and considering the best RL training scenario, it outperformed the best MPC approach.

4.6 Conclusions and Future Work

Activity B3 aimed to test data-driven control strategies in closed-loop simulation with virtual buildings. The open-source BOPTEST framework and an open-source FMU-DRL framework made such virtual buildings available. Five studies were carried out; they tested variations of MPC and RL, including grey-box and black-box models for MPC, different prediction horizons and control steps, and RL learning algorithms. Two of the studies were directly comparable since they were tested on the same BOPTEST test case (*bestest_hydronic_heat_pump*) and testing scenarios; they are also comparable to other tests in the literature using the same test case and scenarios. Study 4 utilised a different BOPTEST test case (*bestest_air*) and standard test scenarios, enabling the comparison of results with future studies.

The studies' findings consistently show that both MPC and RL significantly outperform rule-based control strategies in maintaining thermal comfort while minimising energy costs under dynamic electricity pricing. In the *bestest_hydronic_heat_pump* test case (Studies 1 and 2) and a related study from the literature (Arroyo et al., 2022), MPC consistently outperformed RL. However, in another study focusing on exploring the limits of control performance, in a different building type (Study 5), the top-performing RL method matched MPC performance and even surpassed it in one stochastic learning scenario.

Among the RL methods evaluated in Study 2, Study 5, and (Arroyo et al., 2022), DDPG, DDQN, and QRDQN emerged as the most effective. Meanwhile, Study 3 explored the integration of nonlinear differentiable machine learning models within an MPC framework, where MLP models performed well across both single-zone (*bestest_air*) and multi-zone (*multi-zone_office_simple_air*) test cases. If MLP performance is insufficient, LSTM is a viable alternative. Overall, the results show significant variation based on how each data-driven control strategy is designed.

Future work should continue benchmarking different approaches for data-driven control on the same test cases and scenarios to identify and refine trends leading to effective performance. Case study setups must be consistent to allow meaningful comparisons across studies and advance shared knowledge in the building operation community, as inconsistencies can hinder the identification of broader trends. Additional future work should evaluate control approaches on more complex HVAC systems, diverse building types, and real-world settings. Efforts like BOPTEST and IBPSA Project 2 aim to address these needs by creating new virtual models, adding uncertainty features to BOPTEST, and coordinating further benchmarking of data-driven control strategies.

4.7 References

Arroyo, J., Spiessens, F., & Helsen, L. (2022). Comparison of optimal control techniques for building energy management. *Frontiers in Built Environment, 8*.

Blum, D., Arroyo, J., Huang, S., Drgona, J., Jorissen, F., Walnum, H. T., Chen, T., Benne, K., Vrabie, D., Wetter, M., & Helsen, L. (2021). Building Optimisation Testing Framework (BOPTEST) for simulation-based benchmarking of control strategies in buildings. *Journal of Building Performance Simulation, 14*(5), 586–610. https://doi.org/10.1080/19401493.2021.1986574

Fu, Y., Xu, S., Zhu, Q., O'Neill, Z., & Adetola, V. (2023). How good are learning-based control vs model-based control for load shifting? Investigations on a single zone building energy system. *Energy, 273*, 127073. https://doi.org/10.1016/j.energy.2023.127073

Gao, Y., Miyata, S., & Akashi, Y. (2023). Energy saving and indoor temperature control for an office building using tube-based robust model predictive control. *Applied Energy, 341*, 121106. https://doi.org/10.1016/j.apenergy.2023.121106

Gao, Y., Shi, S., Miyata, S., & Akashi, Y. (2024). Successful application of predictive information in deep reinforcement learning control: A case study based on an office building HVAC system. *Energy, 291*, 130344. https://doi.org/10.1016/j.energy.2024.130344

Jorissen, F., Reynders, G., Baetens, R., Picard, D., Saelens, D., & Helsen, L. (2018). Implementation and verification of the IDEAS building energy simulation library. *Journal of Building Performance Simulation, 11*(6), 669–688. https://doi.org/10.1080/19401493.2018.1428361

Mostafavi, S., Song, C., Sharma, A., Goyal, R., & Brito, A. E. (2023). Benchmarking model predictive control algorithms in Building Optimisation Testing Framework (BOPTEST). In *Proceedings of the International Building Simulation Conference 2023* (pp. 1371–1380). Shanghai, China. https://publications.ibpsa.org/proceedings/bs/2023/papers/bs2023_1371.pdf

Walnum, H. T., Sartori, I., & Bagle, M. (2020). Model predictive control of district heating substations for flexible heating of buildings. In *Proceedings of the International Conference Organized by IBPSA-Nordic* (pp. 45–54). Oslo, Norway. https://sintef.brage.unit.no/sintef-xmlui/handle/11250/2683181

Wang, D., Zheng, W., Wang, Z., Wang, Y., Pang, X., & Wang, W. (2023). Comparison of reinforcement learning and model predictive control for building energy system optimisation. *Applied Thermal Engineering, 228*, 120430. https://doi.org/10.1016/j.applthermaleng.2023.120430

Wetter, M., Zuo, W., Nouidui, T. S., & Pang, X. (2014). Modelica Buildings Library. *Journal of Building Performance Simulation, 7*(4), 253–270. https://doi.org/10.1080/19401493.2013.765506

Wetter, M., Benne, K., Tummescheit, H., & Winther, C. (2024). Spawn: Coupling Modelica Buildings Library and EnergyPlus to enable new energy system and control applications. *Journal of Building Performance Simulation, 17*(2), 274–292. https://doi.org/10.1080/19401493.2023.2266414

5. Activity B4: Uncertainty-aware hierarchical control of energy systems

5.1 Introduction

Data from sensors and actuators at all levels (appliances, buildings, districts, regions, etc.) will be important for controlling buildings to unlock demand-side flexibility (Madsen *et al.*, 2015). Controlling individual buildings in isolation is not enough to effectively address electricity grid issues. A *hierarchy* of control algorithms is needed. This need calls for novel demand response methodologies, which must be based on control technologies operating at a wide range of spatial-temporal scales and aggregation levels. This chapter, based on the work of Activity B4, briefly outlines methods for real-time implementation of forecasts and controllers for grid-integrated smart buildings, focusing on how to use sensors in real-time controllers for buildings.

This chapter also provides insight into how to consider uncertainties from sensor measurements, unmeasured or unrecognised disturbances, and model approximations. Enabling and unlocking the energy flexibility of buildings requires good predictions about occupancy and weather. Nonetheless, forecasts regarding future conditions come with inherent uncertainty; therefore, it is essential to employ methods that account for this uncertainty in controllers, necessitating stochastic controllers.

The weather plays a key role in the energy consumption of buildings. Weather components like solar radiation, ambient air temperature and wind speed are among the most important sources of uncertainty. Hence, short-term forecasts of the influence of weather variables that include uncertainties must be carefully modelled to achieve top-tier controllers (Thilker *et al.*, 2021b; Thavlov and Madsen, 2015). To achieve state-of-the-art control algorithms, occupancy-related disturbances must also be accounted for (Cali *et al.*, 2019; Zhang *et al.*, 2024).

This chapter overviews data- and sensor-driven controllers for real-time demand response in buildings. While machine learning models are widely used for quick load prediction and estimation of heating and cooling needs, grey-box models, which retain physical meaning in their parameters, offer a key advantage by better handling uncertainties from input variables, model approximations, and other influences.

Incorporating uncertainty information (from sensors, systems, disturbances, etc.) is not yet a standard practice. Most controllers in use today are *mean-value-based*. Since the uncertainty linked, e.g. to sensors for solar radiation and occupancy, might be considerable, and since some control-related problems are linked to constraints (such as comfort limits), it is important to take the uncertainties into account in sensor-based control of buildings. This chapter focuses on forecasting and the uncertainties associated with key disturbances, particularly those related to weather and occupancy. Additionally, it extends the discussion beyond smart building control to include real-time controllers for grid and balancing services. To achieve this, we introduce a hierarchical framework of controllers, each designed to address a specific grid or balancing challenge. This structured approach—integrating models, aggregators, and controllers — has been referred to as a Smart-Energy OS (Operating System) (Banaei *et al.*, 2023; Dognini *et al.*, 2022; Madina *et al.*, 2019).

5.2 Measured data: feedback from sensors and human interaction

The real-time collection of feedback data is at the core of the physical implementation of building controls. There are two main types of feedback: sensor data and human interaction.

Physical variables. The energy management system monitors a building's indoor environment using various sensors, such as temperature, humidity, illuminance, motion, and CO_2 sensors. These sensor networks provide necessary information for the intelligent control and management of building energy systems. The optimal placement of actuators and sensors in a building system is a challenging problem due to the complex fluid flow physics governed by coupled nonlinear partial differential equations subjected to disturbances, various sources of uncertainties, and complicated geometry (Vaidya *et al.*, 2012).

Human feedback. The concept known as “human-in-the-loop” control focuses on using the occupants as comfort sensors in the control loop. Kane (2018) presents the idea of human-perceived thermal comfort as a control objective, as opposed to the conventional approach of controlling the indoor air temperature. An important question to be answered is how to implement such a control scheme in practice. An attempt at this is made by Cali *et al*. (2022), presenting a human-in-the-loop-based controller, where a mobile app was developed that enables occupants to deliver feedback based on their perceived comfort. This way, *the human becomes the sensor* and gives feedback to the system rather than a sensor on a wall or in a thermostat. Information from occupants can be combined with information from traditional sensors.

5.3 Forecasting and incorporating disturbances

As Chapter 3 examines the suitability of grey-box models for data-driven MPC, further elaboration is unnecessary here. However, forecasting disturbances is a critical component of grey-box models. The term *disturbance* refers to an influential input beyond the operator’s control; weather variables are a good example. Disturbances (e.g., solar radiation and occupancy) can be significant and highly time-varying for buildings. Consequently, for real-time implementations, it is important to integrate short-term probabilistic perspectives. This section briefly describes methodologies for integrating disturbances into real-time control setups.

Future energy systems will be weather-driven; consequently, accurate knowledge and description of future weather events are crucial. Solar radiation plays a significant role in buildings. These fluctuations are critically important to a building’s HVAC system and the needed heating. High-quality forecasts are required to control energy systems on small scales (buildings, residential houses, batteries, etc.) and large scales (districts, electricity grids, etc.). The following sections will provide an overview of the most common ways of forecasting solar radiation and outdoor air temperature. Lastly, examples of embedded disturbance prediction models in controllers will be given.

5.3.1 Solar radiation forecasting

Solar radiation is one of the most important short-term disturbances. The literature proposes various grey-box-related models and linear stochastic models. AR and ARX models are popular stochastic modelling schemes for solar radiation estimation (Amaro e Silva and C. Brito, 2018; Boland, 2015; Bacher *et al.*, 2009). These studies also investigate the importance of meteorological forecasts related to short-term forecasting. They have found that the available solar radiation is the most important input to the model among the list of typically provided meteorological forecasts.

5.3.2 Outdoor air temperature forecasting

In well-insulated buildings, the short-term dynamic changes of the outdoor air temperature do not significantly influence the indoor air temperature. A building envelope can be thought of as *low-pass filter,* which filters out the fluctuating dynamics (Nielsen and Madsen, 2006). For this reason, in well-insulated buildings, the heating demand is primarily determined by the outdoor temperature signal after it has been low-pass filtered. Thus, in absence of a significant solar radiation input, the heat required to keep a comfortable indoor temperature during several hours is mostly constant. Dynamical models and auto-regressive (AR) and moving average (MA) models are natural choices for modelling outdoor air temperature. Florita and Henze (2009) estimate and compare different moving average models with neural network models. The comparison yielded no conclusive evidence of superior performance for either type

5.3.3 Occupancy forecasting

The goal of optimising the building's indoor air climate is to satisfy the building's users. Consequently, forecasting when and to what extent the building is occupied is crucial to scheduling its operation. Current standards for dealing with variable occupancy levels in buildings are, among others, fixed-schedule operation (e.g. nightly temperature setbacks to save energy during unoccupied periods) and occupation-detection-triggered activation (Nagy *et al.*, 2023) (using, e.g. radio- or video cameras to detect occupants (Adamopoulou *et al.*, 2016)). These control methods are straightforward to implement but have drawbacks (Dobbs and Hencey, 2014). Fixed-schedule operation cannot adjust for occupant numbers, and occupancy-triggered control fails to maintain a satisfactory indoor climate promptly due to thermal time constants.

The literature suggests that numerous methods can improve upon the baseline strategies mentioned above (i.e., fixed schedules and occupancy-triggered control). These methods include supervised and unsupervised machine learning (clustering, ANNs, LSTM networks, etc.), classical time series models such as ARIMA-types and discrete-state Markov models, and stochastic differential equations. Vazquez and Kastner (2011) use clustering techniques (such as C-means and self-organising maps) to forecast occupation levels. Neural network-based models such as ANNs (Schiele *et al.*, 2011) and (Hitimana *et al.,* 2021) have also been implemented and tested for a hospital room and a big office room, respectively, and suggests good performance, and can even supply satisfactory forecasts one week ahead (Pesic *et al.*, 2019). Discrete-time, discrete state-space Markov chains have been investigated and tested (Li and Dong, 2018; Andersen *et al.*, 2013) with promising results. Diurnal variations in occupant behaviour have been modelled using continuous-time Markov chains. (Zhang *et al.*, 2024). Wolf *et al.* (2019) use a stochastic differential equation approach to forecast the CO_2 levels in rooms, which are good indicators for occupancy.

5.3.4 Embedded disturbance models

Embedded disturbance models are mathematical representations within the control system that account for disturbances. Alternatively, mean value forecasts from third-party sources or persistent forecasts (assuming that conditions will remain unchanged) can be used. Embedded disturbance models in control are presented by (Hou et al., 2022b; Pippia et al., 2021; and Thilker et al., 2021b). Hou et al. (2022b) propose an error model that corrects and improves the supplied weather forecasts using easily accessible local weather observations. Pippia et al. (2021) introduce a scenario-based MPC where multivariate, probabilistic models generate coherent scenarios for control. Thilker et al. (2021b) briefly outlined an approach that included a description of the amount of cloud cover, solar radiation, net radiation, and ambient air temperature. Grey-box models, using stochastic differential equations (see Section 3.5.1), offer a unified framework for forecasting disturbances and describing both the building dynamics and disturbances.

5.4 Control Methods for Real-Time Implementation

This section reviews current standard stochastic controllers (i.e., control approaches) that can account for stochasticity in either the model or disturbances. Such controllers include, e.g., Gaussian controllers (where model noise is assumed to be Gaussian), chance-constrained controllers, and multi-stage controllers.

5.4.1 Model-based stochastic control: linear and Gaussian controllers

Due to their simplicity and broad applicability, a widely used class of controllers (i.e., control approaches) relies on a linear system model with additive Gaussian noise and a quadratic (or simpler) objective function. Examples can be found in West *et al.* (2014), Drgona *et al.* (2020), Cigler *et al.* (2013), and Siroky *et al.* (2011), where the quadratic penalty is used. Their popularity stems from several factors. For instance, building thermal dynamics is typically linear, which justifies using additive Gaussian noise (Bacher & Madsen, 2011). This assumption holds when disturbances are well-modelled. Quadratic or piecewise quadratic objective functions are suitable due to their smoothness and flexibility. Such a controller includes uncertainty in its formulation, and the solution is given in closed form by taking the *expected value* in the optimal control problem formulation. The objective function is given by:

$$\varphi_k = \mathrm{E}\left[\int_{t_k}^{t_{k+N}} l(x(\tau),u(\tau),d(\tau)) + l_b(x(t_{k+N})) \right] \tag{5.1}$$

where E is the expectation operator, l and l_b are respectively the stage cost and *cost-to-go* terms. Here l is either linear or quadratic (or a combination of both):

$$l(x,u,d) = x'Q_x x + q_x' x + u'Q_u u + q_u' u \tag{5.2}$$

An objective function balancing system deviation penalties and economic costs enables flexible trade-offs (Thilker et al., 2022). When stochastic grey-box models are used, uncertainty can be directly integrated into model predictive control (Wytock et al., 2017; Zhang et al., 2013). In linear-quadratic setups with Gaussian noise, forecast scenarios support multi-stage control; otherwise, mean-based controllers may suffice.

5.4.2 Linear vs non-linear models

The advantages mentioned above should be seen in the light of compromises (Eser *et al.*, 2021; Schirrer *et al.*, 2016). If the actual system is non-linear, the modeller faces a trade-off: (a) he/she can choose to include the non-linearity in the model but is left with a nonlinear optimal control problem, or (b) he/she can approximate the non-linearity with a linear model but is rewarded with a more robust and faster framework from the linear quadratic controller. One must determine case-by-case whether the advantages of the linear framework compensate for the performance loss caused by linearisation (Pčolka *et al.*, 2016). Zanetti et al. (2023) found that simpler solutions can often match the performance of non-linear systems.

Alternative ways to obtain the benefits of both linear and non-linear models have also been proposed. It is an active research field, and results point toward promising improvements. Pčolka *et al.* (2016); Mork *et al.* (2022); Ostadijafari and Dubey (2019) address the challenges posed by non-linear MPC, particularly the additional complexity it introduces into the MPC framework. They propose a way to get around this by formulating a linear *time-varying* model (which keeps the MPC problem convex and simple) and comparing it to MPC using the linear time-invariant and non-linear models. Results suggest significant improvement over the linear model and performance close to that of the non-linear model.

In the case of heating and cooling units, these are often operated on an *on/off* basis (Kim *et al.*, 2015; Kim and Braun, 2018; Kim *et al.*, 2022; LeBreux *et al.*, 2009). These kinds of problems are often solved as integer problems, since this is a natural approach to represent the on/off binary operation (Kim *et al.*, 2022). However, *continuous* approaches to the on/off operation exist in the literature (Brok *et al.*, 2022), where the optimisation variables are the periods between switches (Axelsson *et al.*, 2005). This approach has the advantage of being computationally efficient compared to integer programs since the continuous nature of the optimisation is preserved, and analytical gradients can be computed efficiently.

5.4.3 Implementing reinforcement learning controllers

MPC controllers have been widely used in different applications due to their simplicity and robustness in creating optimised controllers. However, one drawback of their current form is that formulating these controllers is a time-consuming task that must be performed by experts. Integrating reinforcement learning (RL) algorithms offers new possibilities for developing intelligent decision-making systems, as these methods depend solely on data collected from the environment. A popular introduction to the concept of RL can be found in Sutton and Barto (2018). Numerous examples of RL used for building climate control exist in the literature. Neural networks are popular as an underlying model for an RL controller (Elnour *et al.*, 2022; Yang *et al.*, 2021; Jin *et al.*, 2021). However, most RL-based optimal control studies for indoor climate control of buildings are simulation-based and hence not tested in real life. An experimental example is given in (Brandi *et al.*, 2020), where they achieve 5-15% performance improvement. In general, the literature on RL implementation is very scarce, which is partly due to the time and data-consuming learning phase. Wang and Hong (2020) report that only 11% of reinforcement learning controllers were implemented and tested in actual buildings, leaving a research gap between theory and practice for RL. Comparisons between MPC and RL have shown that each method offers specific benefits (Ernst *et al.*, 2009; Lin *et al.*, 2020). In their flexibility study, Mbuwir et al. (2019) found that both approaches yielded comparable results.

5.5 Hierarchical Control for Smart Grid Integration

Future smart energy systems will require linking local controllers (e.g. in buildings) with high-level electricity markets operating at larger scales. Rather than managing voltage, frequency, and congestion through distinct markets, energy flexibility research proposes addressing these services via integrated, hierarchical control structures (Madina et al., 2019; Madsen et al., 2016a; Brok et al., 2019). In this section, we first show how the controllers introduced earlier can function as low-level controllers within a hierarchical setup aimed at solving grid and ancillary service issues in future smart energy systems. We then briefly discuss how these principles can be generalised to broader multi-level control problems, and how such problems can be linked to existing electricity market structures.

5.5.1 Flexibility Function: a tool supporting hierarchical control

The previous sections have shown how to develop controllers for smart buildings based on forecasts of prices, weather conditions, and indoor comfort requirements. This approach can be further applied to generate *price signals* that subtly steer the demand in smart buildings (Halvgaard *et al.*, 2012). The basic concept is illustrated by Figure 5.1a, where a smart building takes a pricing signal (input) and, as a result, consumes a certain energy demand (output). This analysis enables the creation of a model, called the *Flexibility Function (FF)*, that *predicts* demand as a dynamic response to price.

This Flexibility Function could be any type of dynamic model. In (Junker *et al.*, 2018), a linear model (finite impulse response model) is suggested, but in (Junker *et al.*, 2020) it is shown that a grey-box model using SDEs is more appropriate. After defining an FF, a *second* supervisory controller can be implemented to manage energy demand according to certain objectives, using energy price as the primary input for decision-making. Figure 5.2: Hierarchical control and markets.

b shows that the FF can generate prices based on predefined reference signals or targets. Note how the demand returns information into the “price generator/controller”, thereby closing the control loop. In this context, the FF plays a role similar to that of a control-oriented model used for local MPC in buildings.

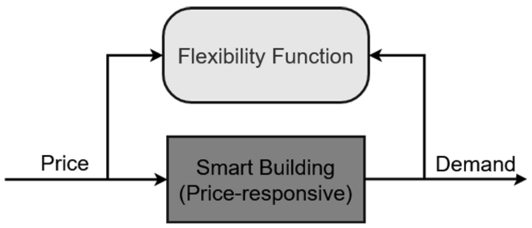

(a)

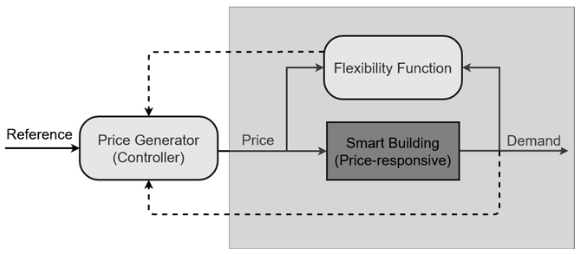

(b)

Figure 5.1: The demand of a smart building can be predicted as a function of price.

Let FF be the Flexibility Function, which takes a price signal as input (c_u) and outputs the expected demand, while r_l is a reference load. A simple *upper-level* optimisation can be formulated to minimise the difference between the reference load and the demand predicted by the FF function in response to a *price signal*. Combining this upper-level optimisation problem with an optimal control problem for a single building (*lower level*) reveals how the price signal (c_u) couples the two in an elegant fashion:

$$
\begin{array}{lll}
\min\limits_{c_u} & \left(FF(c_u)-r_l\right)^2 & \text{Upper level} \\
\hline
\min\limits_{u_k} & \sum\limits_k c_{u,k}^{\mathrm{T}} u_k & \text{Lower level} \\
s.t. & x_{k+1}=\phi(x_{k+1},u_k,d_k,\omega_k) & \\
 & c_e(x_{k+1},u_k)=a & \\
 & c_i(x_{k+1},u_k)\le b &
\end{array}
\tag{5.3}
$$

In the lower level, the inputs (u_k) are adjusted in response to the optimised price signal. Notice that the two optimisation problems are solved independently, thus preserving the building owners' autonomy and privacy while allowing an aggregator to utilise the energy flexibility. In practice, each aggregator will manage numerous smart buildings, each with its independent control challenges. This method scales well to this case since the computational burden for the upper-level remains constant, with the Flexibility Function simply representing the aggregated response from the smart buildings. The concept is one of the core elements of the **Smart Energy OS** framework, which is a hierarchy of controllers for providing grid and balancing services (Madina *et al.*, 2019).

5.5.2 Hierarchical control: opportunities and alternative approaches

The above-described hierarchical control setup couples the lowest level (buildings) and the upper level (e.g., a group of buildings or a district). However, hierarchical control can be utilised on different levels and for different tasks. Ancillary services address supply and demand imbalances to help operators maintain a reliable power grid. Building heating, ventilation, and air-conditioning (HVAC) systems are well-suited for such services due to their significant energy use and flexible demand. Qureshi and Jones (2018) proposed a two-level hierarchical control approach: a lower-level thermal flexibility controller adjusts temperature setpoints to maximise thermal flexibility, while a higher-level electrical flexibility controller manages HVAC operation to optimise grid support. Both layers use robust optimisation methods. Blum et al. (2017) introduced a method to define and recursively quantify the opportunity costs of ancillary services provided by building HVAC systems, allowing for predictive accounting of such services. Huang et al. (2019) proposed a two-level hierarchical control strategy for managing demand flexibility in large building groups. By modelling the group as a unified “virtual” building and applying a genetic algorithm to determine optimal performance, then coordinating each building’s operation through non-linear programming, their approach drastically cut computation time, requiring only one-eighth of the time needed by conventional decentralised methods.

Lankeshwara et al. (2022) proposed a hierarchical control framework to coordinate residential air conditioners across multiple households under uncertainty. The objective is to meet a specific combined power demand for the total consumption while meeting each household’s thermal constraints. Recognising the privacy and computational limitations of a fully centralised controller, which would require access to detailed thermal models, occupant preferences and real-time data, the authors adopted an *Alternating Direction Method of Multipliers* (ADMM) algorithm (Boyd et al., 2011). In this setup, the aggregator (the upper-level controller) knows only the total power consumption and computes a set of auxiliary and dual variables used in the decentralised (local) controllers. These variables are updated periodically between control intervals.

Building microgrids with distributed RESs and storage reduces transmission losses and grid dependency. Yamashita et al. (2021) defined hierarchical control (HC) in three layers: *primary control* manages power converters, *secondary control* corrects voltage and frequency deviations, and *tertiary control* optimises energy trading to balance supply and demand. RESs cause variability in the AC grid, which requires equilibrium between production and consumption. DC microgrids, powered by solar panels or fuel cells, offer a solution for stability. Jin et al. (2014) propose a three-level control strategy for DC microgrids: Level I (primary) maintains bus voltage, Level II (secondary) regulates voltage and manages battery charge/discharge, and Level III (tertiary) addresses extreme conditions like load or resource changes through load shedding or ballast.

5.5.3 Enabling energy markets at different scales through hierarchical control

Future smart energy systems will require a connection between controllers operating at local scales and high-level markets operating at large scales. A range of all relevant spatial aggregation levels (building, district, city, region, country, etc.) must be considered. At the same time, control or market solutions must ensure that the power system is balanced at all future temporal scales. As a result, data-driven solutions for managing flexible electrical energy systems should be implemented across all spatial and temporal scales. Several solutions have recently been proposed in the literature to address this challenge, including *Transactive Energy*, Peer-to-Peer, and *Control-Based* methods, as described in De Zotti *et al.* (2018).

Traditionally, power systems are operated through market-based bidding, on multiple markets corresponding to different horizons (e.g. day-ahead, intra-day, balancing and frequency regulation). These bids are usually static and include two main components: volume (the amount of power provided) and duration. Nevertheless, a fundamental shift is needed in how low-level flexibility is activated. Currently, supply and demand curves are established for each market horizon based on all submitted bids. In mathematical terms, these curves are typically *static* and *deterministic*. A merit-order dispatch is then applied to minimise generation costs. However, when generation comes from wind or solar sources, the supply curve becomes *stochastic*, and demand-side flexibility must instead be described dynamically using the *Flexibility Function*. Consequently, there is a need to introduce new digitised markets that are both dynamic and stochastic. Rather than relying on multiple markets for different purposes (e.g., frequency, voltage, congestion, etc.), across varying time horizons, we propose using concepts based on the Flexibility Function and stochastic control theory, as outlined in the previous section for the two-level case, in the framework called the Smart Energy OS (Madsen *et al.*, 2015, 2016b; Madina *et al.*, 2019; Tohidi *et al.*, 2024).

When we “zoom out” in space and time, such as by considering the load across a large area over several days or just the next day, the dynamics and randomness become less significant, allowing the use of conventional market principles, as shown in Figure 5.2. Conversely, when we “zoom in” on higher temporal and spatial resolutions (like, for instance, a house), the dynamics and stochasticity become important, and consequently, we will suggest using methods for increased flexibility as discussed in this report. The total setup consists of a combination of all these options, and the best option depends on the “zoom level.” In conclusion, future digitised market principles should incorporate a hierarchical structure, with traditional market-based bidding used at higher levels and a control-based approach at lower levels.

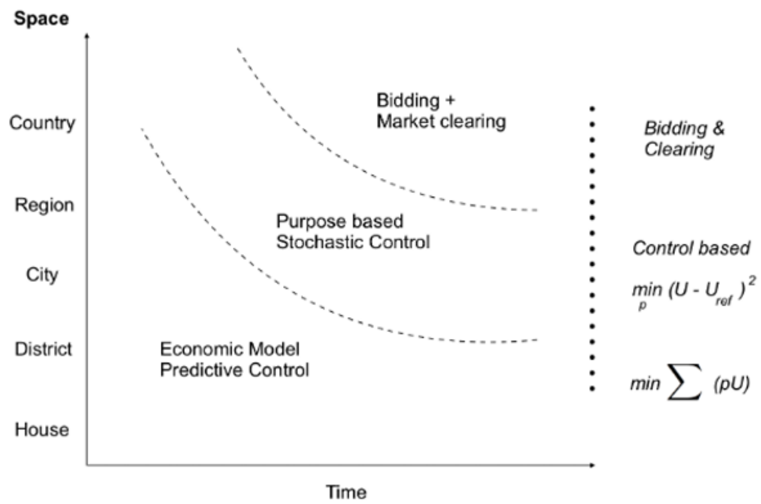

Figure 5.2: Hierarchical control and markets.

5.6 Summary

Several good review papers already exist on simulation-based methods for model predictive control in smart grid-enabled buildings. However, these models are mostly *deterministic* and optimised for *long-term simulations* (e.g., for design purposes). On the other hand, there is a gap in the literature concerning real-time, sensor-based control. There is a critical need to focus on *short-term prediction horizons and* techniques to address sensor and disturbance-related uncertainties. This chapter has discussed strategies to address this need by integrating approaches suitable for real-time control, including stochastic modelling, forecasting methods and adequate control schemes. Furthermore, this section outlines methodologies for characterising and leveraging buildings' energy flexibility to provide services for energy grids. This section has discussed using grey-box models based on stochastic differential equations to represent partially observed states. These models bridge *physics-based and data-driven approaches* and are often seen as *data-driven digital twins*, optimised for real-time use and capable of integrating sensor uncertainty for improved forecasting, control, and optimisation. The MPC approaches discussed in this section incorporate this uncertainty. While current controllers rely on physical sensors, future approaches must also integrate occupant feedback. Likewise, energy grids require improved digital infrastructure to accommodate "smart grid-ready" buildings.

5.7 References

Adamopoulou, A. A., Tryferidis, A. M., & Tzovaras, D. K. (2016). A context-aware method for building occupancy prediction. *Energy and Buildings, 110,* 229–244. https://doi.org/10.1016/j.enbuild.2015.10.003

Amaro e Silva, R., & Brito, C. M. (2018). Impact of network layout and time resolution on spatio-temporal solar forecasting. *Solar Energy, 163,* 329–337. https://doi.org/10.1016/j.solener.2018.01.095

Andersen, P. D., Iversen, A., Madsen, H., & Rode, C. (2013). Dynamic modeling of presence of occupants using inhomogeneous Markov chains. *Energy and Buildings, 69,* 213–223. https://doi.org/10.1016/j.enbuild.2013.10.001

Axelsson, H., Egerstedt, M., Wardi, Y., & Vachtsevanos, G. (2005). Algorithm for switching-time optimisation in hybrid dynamical systems. *Proceedings of the IEEE International Symposium on Intelligent Control,* 256–261. https://doi.org/10.1109/.2005.1467024

Bacher, P., & Madsen, H. (2011). Identifying suitable models for the heat dynamics of buildings. *Energy and Buildings, 43,* 1511–1522. https://doi.org/10.1016/j.enbuild.2011.02.005

Bacher, P., Madsen, H., & Nielsen, H. A. (2009). Online short-term solar power forecasting. *Solar Energy, 83,* 1772–1783. https://doi.org/10.1016/j.solener.2009.05.016

Banaei, M., D'Ettorre, F., Ebrahimy, R., Pourmousavi, S., Blomgren, E., & Madsen, H. (2023). A stochastic methodology to exploit maximum flexibility of swimming pool heating systems. *International Journal of Electrical Power & Energy Systems, 145.* https://doi.org/10.1016/j.ijepes.2022.108643

Blum, D. H., Zakula, T., & Norford, L. K. (2017). Opportunity cost quantification for ancillary services provided by heating, ventilating, and air-conditioning systems. *IEEE Transactions on Smart Grid, 8,* 1264–1273. https://doi.org/10.1109/TSG.2016.2582207

Boland, J. (2015). Spatial-temporal forecasting of solar radiation. *Renewable Energy, 75,* 607–616. https://doi.org/10.1016/j.renene.2014.10.035

Boyd, S., Parikh, N., Chu, E., Peleato, B., & Eckstein, J. (2011). Distributed optimisation and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning, 3,* 1–122. https://doi.org/10.1561/2200000016

Brandi, S., Piscitelli, M. S., Martellacci, M., & Capozzoli, A. (2020). Deep reinforcement learning to optimise indoor temperature control and heating energy consumption in buildings. *Energy and Buildings, 224,* 110225. https://doi.org/10.1016/j.enbuild.2020.110225

Brok, N., Green, T., Heerup, C., Oren, S. S., & Madsen, H. (2022). Optimal operation of an ice-tank for a supermarket refrigeration system. *Control Engineering Practice, 119,* 104973. https://doi.org/10.1016/j.conengprac.2021.104973

Brok, N., Stentoft, P., Munk-Nielsen, T., & Madsen, H. (2019). Flexible control of wastewater aeration for cost-efficient, sustainable treatment. *Control of Smart Grid and Renewable Energy Systems (CSGRES 2019),* 494–499.

Cali, D., Kindler, E., Ebrahimy, R., Bacher, P., Hu, K. S., Østrup, M. L., & Bachalarz, M. (2019). climify.org: An online solution for easy control and monitoring of the indoor environment. *E3S Web of Conferences, 111,* 05006. https://doi.org/10.1051/e3sconf/201911105006

Cali, D., Thilker, C., Specht, S., Real, J., Madsen, H., & Olesen, B. (2022). Human in the loop: Perceived-based control as the key to enhance buildings' performance. *Building Simulation 2021 Conference (BS2021),* 1–3.

Cigler, J., Gyalistras, D., Siroky, J., Tiet, V., & Ferkl, L. (2013). Beyond theory: The challenge of implementing model predictive control in buildings. *11th REHVA World Congress (CLIMA),* 250, 1008–1018.

De Zotti, G., Pourmousavi, S. A., Madsen, H., & Poulsen, N. K. (2018). Ancillary services 4.0: A top-to-bottom control-based approach for solving ancillary services problems in smart grids. *IEEE Access, 6,* 11694–11706. https://doi.org/10.1109/ACCESS.2018.2805330

DNV. (2022). Demand-side flexibility in the EU: Quantification of benefits in 2030. *Technical Report.* Accessed: 2023-07-12.

Dobbs, J. R., & Hencey, B. M. (2014). Model predictive HVAC control with online occupancy model. *Energy and Buildings, 82,* 675–684. https://doi.org/10.1016/j.enbuild.2014.07.051

Dognini, A., Challagonda, C., Moro, E., Helmholt, K., Madsen, H., Daniele, L., Schmitt, L., Temal, L., Genest, O., & Calvez, P. (2022). Data spaces for energy, home, and mobility. *RWTH Aachen University.* https://doi.org/10.5281/zenodo.7193318

Drgoňa, J., Picard, D., & Helsen, L. (2020). Cloud-based implementation of white-box model predictive control for a geotabs office building: A field test demonstration. *Journal of Process Control, 88,* 63–77. https://doi.org/10.1016/j.jprocont.2020.02.007

Elnour, M., Himeur, Y., Fadli, F., Mohammedsherif, H., Meskin, N., Ahmad, A. M., Petri, I., & Rezgui, Y. (2022). Neural network-based model predictive control system for optimising building automation and management systems of sports facilities. *Applied Energy, 318,* 119153. https://doi.org/10.1016/j.apenergy.2022.119153

Ernst, D., Glavic, M., Capitanescu, F., & Wehenkel, L. (2009). Reinforcement learning versus model predictive control: A comparison on a power system problem. *IEEE Transactions on Systems, Man, and Cybernetics, Part B (Cybernetics), 39,* 517–529. https://doi.org/10.1109/TSMCB.2008.2007630

Eser, S., Stoffel, P., Kumpel, A., & Müller, D. (2021). Evaluation of linear and nonlinear system models in hierarchical model predictive control of HVAC systems. *Journal of Physics: Conference Series, 2042,* 012032. https://doi.org/10.1088/1742-6596/2042/1/012032

Florita, A. R., & Henze, G. P. (2009). Comparison of short-term weather forecasting models for model predictive control. *HVAC&R Research, 15,* 835–853.

Halvgaard, R., Poulsen, N. K., Madsen, H., & Jørgensen, J. B. (2012). Economic model predictive control for building climate control in a smart grid. *2012 IEEE PES Innovative Smart Grid Technologies (ISGT),* 1–6. https://doi.org/10.1109/ISGT.2012.6175631

Halvgaard, R. F., Vandenberghe, L., Poulsen, N. K., Madsen, H., & Jørgensen, J. B. (2016). Distributed model predictive control for smart energy systems. *IEEE Transactions on Smart Grid, 7,* 1675–1682. https://doi.org/10.1109/TSG.2016.2526077

Hitimana, E., Bajpai, G., Musabe, R., Sibomana, L., & Kayalvizhi, J. (2021). Implementation of IoT framework with data analysis using deep learning methods for occupancy prediction in a building. *Future Internet, 13,* 67. https://doi.org/10.3390/fi13030067

Hou, J., Li, H., Nord, N., & Huang, G. (2022b). Model predictive control under weather forecast uncertainty for HVAC systems in university buildings. *Energy and Buildings, 257,* 111793. https://doi.org/10.1016/j.enbuild.2021.111793

Huang, P., Fan, C., Zhang, X., & Wang, J. (2019). A hierarchical coordinated demand response control for buildings with improved performances at building group. *Applied Energy, 242,* 684–694. https://doi.org/10.1016/j.apenergy.2019.03.148

Jin, C., Wang, P., Xiao, J., Tang, Y., & Choo, F. H. (2014). Implementation of hierarchical control in DC microgrids. *IEEE Transactions on Industrial Electronics, 61,* 4032–4042. https://doi.org/10.1109/TIE.2013.2286563

Jin, Y., Yan, D., Zhang, X., An, J., & Han, M. (2021). A data-driven model predictive control for lighting system based on historical occupancy in an office building: Methodology development. *Building Simulation,* 219–235.

Junker, R. G. (2019). Characterisation and integration of energy flexibility through stochastic modelling and control.

Junker, R. G., Azar, A. G., Lopes, R. A., Lindberg, K. B., Reynders, G., Relan, R., & Madsen, H. (2018). Characterizing the energy flexibility of buildings and districts. *Applied Energy, 225,* 175–182.

Junker, R. G., Kallesøe, C. S., Real, J. P., Howard, B., Lopes, R. A., & Madsen, H. (2020). Stochastic non-linear modelling and application of price-based energy flexibility. *Accepted in Applied Energy.*

Junker, R. G., Relan, R., & Madsen, H. (2019). Designing individual penalty signals for improved energy flexibility utilisation. *IFAC-PapersOnLine, 52,* 123–128. https://doi.org/10.1016/j.ifacol.2019.08.166

Kane, M. B. (2018). Modeling human-in-the-loop behavior and interactions with HVAC systems. *2018 Annual American Control Conference (ACC),* 4628–4633. https://doi.org/10.23919/ACC.2018.8431913

Kim, D., Braun, J., Cai, J., & Fugate, D. (2015). Development and experimental demonstration of a plug-and-play multiple RTU coordination control algorithm for small/medium commercial buildings. *Energy and Buildings, 107,* 279–293. https://doi.org/10.1016/j.enbuild.2015.08.025

Kim, D., & Braun, J. E. (2018). Development, implementation, and performance of a model predictive controller for packaged air conditioners in small and medium-sized commercial building applications. *Energy and Buildings, 178,* 49–60. https://doi.org/10.1016/j.enbuild.2018.08.019

Kim, D., Wang, Z., Brugger, J., Blum, D., Wetter, M., Hong, T., Piette, M. A. (2022). Site demonstration and performance evaluation of MPC for a large chiller plant with TES for renewable energy integration and grid decarbonization. *Applied Energy, 321,* 119343. https://doi.org/10.1016/j.apenergy.2022.119343

Lankeshwara, G., Sharma, R., Yan, R., & Saha, T. K. (2022). A hierarchical control scheme for residential air-conditioning loads to provide real-time market services under uncertainties. *Energy, 250,* 123796. https://doi.org/10.1016/j.energy.2022.123796

LeBreux, M., Lacroix, M., & Lachiver, G. (2009). Control of a hybrid solar/electric thermal energy storage system. *International Journal of Thermal Sciences, 48,* 645–654. https://doi.org/10.1016/j.ijthermalsci.2008.05.006

Li, Z., & Dong, B. (2018). Short term predictions of occupancy in commercial buildings: Performance analysis for stochastic models and machine learning approaches. *Energy and Buildings, 158,* 268–281. https://doi.org/10.1016/j.enbuild.2017.09.052

Lin, Y., McPhee, J., & Azad, N. (2020). Comparison of deep reinforcement learning and model predictive control for adaptive cruise control. *IEEE Transactions on Intelligent Vehicles, PP,* 1–1. https://doi.org/10.1109/TIV.2020.3012947

Madina, C., Jimeno, J., Ortolano, L., Palleschi, M., Ebrahimy, R., Madsen, H., Pardo, M., Corchero, C., & Migliavacca, G. (2019). Technologies and protocols: The experience of the three smartnet pilots. *TSO-DSO Interactions and Ancillary Services in Electricity Transmission and Distribution Networks,* 141–183. https://doi.org/10.1007/978-3-030-29203-4_6

Madsen, H., Holst, J., Lindström, E. (2006). Modelling non-linear and non-stationary time series.

Madsen, H., Parvizi, J., & Bacher, P. (2016a). Smart-energy operating-system: A framework for implementing flexible electric energy systems in smart cities. *SUSTAIN-2016 Abstract L-1.*

Madsen, H., Parvizi, J., & Bacher, P. (2016b). Smart-energy operating-system: A framework for implementing flexible electric energy systems in smart cities. *SUSTAIN-2016.*

Madsen, H., Parvizi, J., Halvgaard, R. F., Sokoler, L. E., Jørgensen, J. B., Hansen, L. H., & Hilger, K. B. (2015). Control of electricity loads in future electric energy systems. *Handbook of Clean Energy Systems.*

Mbuwir, B. V., Geysen, D., Spiessens, F., & Deconinck, G. (2019). Reinforcement learning for control of flexibility providers in a residential microgrid. *IET Smart Grid, 3.* https://doi.org/10.1049/iet-stg.2019.0196

Mork, M., Materzok, N., Xhonneux, A., & Müller, D. (2022). Nonlinear hybrid model predictive control for building energy systems. *Energy and Buildings, 270,* 112298. https://doi.org/10.1016/j.enbuild.2022.112298

Nagy, Z., Gunay, B., Miller, C., Hahn, J., Ouf, M. M., Lee, S., Hobson, B. W., Abuimara, T., Bandurski, K., André, M., Lorenz, C. L., Crosby, S., Dong, B., Jiang, Z., Peng, Y., Favero, M., Park, J. Y., Nweye, K., Nojedehi, P., Stopps, H., Sarran, L., Brackley, C., Bassett, K., Govertsen, K., Koczorek, N., Abele, O., Casavant, E., Kane, M., O'Neill, Z., Yang, T., Day, J., Huchuk, B., Hellwig, R. T., & Vellei, M. (2023). Ten questions concerning occupant-centric control and operations. *Building and Environment, 242,* 110518. https://doi.org/10.1016/j.buildenv.2023.110518

Nielsen, H. A., & Madsen, H. (2006). Modelling the heat consumption in district heating systems using a grey-box approach. *Energy and Buildings, 38,* 63–71. https://doi.org/10.1016/j.enbuild.2005.05.002

Ostadijafari, M., & Dubey, A. (2019). Linear model-predictive controller (LMPC) for building's heating ventilation and air conditioning (HVAC) system. *2019 IEEE Conference on Control Technology and Applications (CCTA),* 617–623. https://doi.org/10.1109/CCTA.2019.8920657

Pešić, S., Tošić, M., Iković, O., Radovanović, M., Ivanović, M., & Bošković, D. (2019). Blemat: Data analytics and machine learning for smart building occupancy detection and prediction. *International Journal on Artificial Intelligence Tools, 28,* 1960005. https://doi.org/10.1142/S0218213019600054

Pippia, T., Lago, J., De Coninck, R., & De Schutter, B. (2021). Scenario-based nonlinear model predictive control for building heating systems. *Energy and Buildings, 247,* 111108. https://doi.org/10.1016/j.enbuild.2021.111108

Pčolka, M., Žáčeková, E., Robinett, R., Čelikovský, S., & Šebek, M. (2016). Bridging the gap between the linear and nonlinear predictive control: Adaptations for efficient building climate control. *Control Engineering Practice, 53,* 124–138. https://doi.org/10.1016/j.conengprac.2016.01.007

Qureshi, F. A., & Jones, C. N. (2018). Hierarchical control of building HVAC system for ancillary services provision. *Energy and Buildings, 169,* 216–227. https://doi.org/10.1016/j.enbuild.2018.03.004

Schiele, J., Koperna, T., & Brunner, J. O. (2011). Predicting intensive care unit bed occupancy for integrated operating room scheduling via neural networks. *Naval Research Logistics, 68,* 65–88. https://doi.org/10.1002/nav.21929

Schirrer, A., Brandstetter, M., Leobner, I., Hauer, S., & Kozek, M. (2016). Nonlinear model predictive control for a heating and cooling system of a low-energy office building. *Energy and Buildings, 125,* 86–98. https://doi.org/10.1016/j.enbuild.2016.04.029

Siroký, J., Oldewurtel, F., Cigler, J., & Prívara, S. (2011). Experimental analysis of model predictive control for an energy-efficient building heating system. *Applied Energy, 88,* 3079–3087. https://doi.org/10.1016/j.apenergy.2011.03.009

Sutton, R. S., & Barto, A. G. (2018). *Reinforcement learning: An introduction.* MIT Press.

Thavlov, A., & Madsen, H. (2015). A non-linear stochastic model for an office building with air infiltration. *International Journal of Sustainable Energy Planning and Management, 7,* 59–70. https://doi.org/10.5278/ijsepm.2015.7.5

Thilker, C. A., Jørgensen, J. B., & Madsen, H. (2022). Linear quadratic Gaussian control with advanced continuous-time disturbance models for building thermal regulation. *Applied Energy, 327,* 120086. https://doi.org/10.1016/j.apenergy.2022.120086

Thilker, C. A., Madsen, H., & Jørgensen, J. B. (2021b). Advanced forecasting and disturbance modelling for model predictive control of smart energy systems. *Applied Energy, 292,* 116889. https://doi.org/10.1016/j.apenergy.2021.116889

Tohidi, S. S., Madsen, H., & Tsaousoglou, G. (2024). Adaptive flexibility function in smart energy systems: A linearized price-demand mapping approach. *2024 European Control Conference (ECC),* 1315–1320.

Vaidya, U., Rajaram, R., & Dasgupta, S. (2012). Actuator and sensor placement in linear advection PDE with building system application. *Journal of Mathematical Analysis and Applications, 394,* 213–224. https://doi.org/10.1016/j.jmaa.2012.03.046

Vázquez, F. I., & Kastner, W. (2011). Clustering methods for occupancy prediction in smart home control. *2011 IEEE International Symposium on Industrial Electronics,* 1321–1328. https://doi.org/10.1109/ISIE.2011.5984350

Wang, Z., & Hong, T. (2020). Reinforcement learning for building controls: The opportunities and challenges. *Applied Energy, 269,* 115036. https://doi.org/10.1016/j.apenergy.2020.115036

West, S. R., Ward, J. K., & Wall, J. (2014). Trial results from a model predictive control and optimisation system for commercial building HVAC. *Energy and Buildings, 72,* 271–279. https://doi.org/10.1016/j.enbuild.2013.12.037

Wolf, S., Cali, D., Krogstie, J., & Madsen, H. (2019). Carbon dioxide-based occupancy estimation using stochastic differential equations. *Applied Energy, 236,* 32–41. https://doi.org/10.1016/j.apenergy.2018.11.078

Wytock, M., Moehle, N., & Boyd, S. (2017). Dynamic energy management with scenario-based robust MPC. *2017 American Control Conference (ACC),* 2042–2047. https://doi.org/10.23919/ACC.2017.7963253

Yang, S., Wan, M. P., Chen, W., Ng, B. F., & Dubey, S. (2021). Experiment study of machine-learning-based approximate model predictive control for energy-efficient building control. *Applied Energy, 288,* 116648. https://doi.org/10.1016/j.apenergy.2021.116648

Yassuda Yamashita, D., Vechiu, I., & Gaubert, J. P. (2021). Two-level hierarchical model predictive control with an optimised cost function for energy management in building microgrids. *Applied Energy, 285,* 116420. https://doi.org/10.1016/j.apenergy.2020.116420

Zanetti, E., Kim, D., Blum, D., Scoccia, R., & Aprile, M. (2023). Performance comparison of quadratic, nonlinear, and mixed integer nonlinear MPC formulations and solvers on an air source heat pump hydronic floor heating system. *Journal of Building Performance Simulation, 16,* 144–162. https://doi.org/10.1080/19401493.2022.2120631

Zhang, H., Thilker, C. A., Madsen, H., Li, R., Xiao, F., Ma, T., Xu, K. (2024). Stochastic occupancy modelling for spaces with irregular occupancy patterns using adaptive B-spline-based inhomogeneous Markov chains. *Building and Environment,* 111721. https://doi.org/10.1016/j.buildenv.2024.111721

Zhang, K., Prakash, A., Paul, L., Blum, D., Alstone, P., Zoellick, J., Brown, R., Pritoni, M. (2022). Model predictive control for demand flexibility: Real-world operation of a commercial building with photovoltaic and battery systems. *Advances in Applied Energy, 7,* 100099. https://doi.org/10.1016/j.adapen.2022.100099

Zhang, X., Schildbach, G., Sturzenegger, D., & Morari, M. (2013). Scenario-based MPC for energy-efficient building climate control under weather and occupancy uncertainty. *2013 European Control Conference (ECC),* 1029–1034. https://doi.org/10.23919/ECC.2013.6669664

6. Final Remarks

This report outlines the findings from Annex 81, focusing on how data can enhance building performance, specifically through Model Predictive Control (MPC), a control approach based on using a model to optimise control sequences. Some conclusions drawn from the activities within Subtask B are as follows:

- **Structuring data properly is necessary.** Control applications are very diverse, but in all of them, data collection is necessary to create and calibrate models and periodically update the model's state in an implementation carried out in real-time. With a well-organized database, it is possible to study different modelling methodologies.
- **Strategies for creating control-oriented models are as diverse as the types of buildings that exist**, mechanical configurations, as well as the different climatic and energy market conditions that exist around the world. However, "grey box" and "black box" models tend to be preferred because of the ease of parameter calibration, and in the case of grey-box models, how they can incorporate the treatment of uncertainty. Significant effort is required to standardise model development or at least establish uniform procedures for creating models.
- **A platform like BOPTEST is useful as it provides a standardised environment for comparing the performance of different control strategies.** Ongoing research is required to build a "library of emulators" with a sufficiently general application so that its results are informative and relevant for a vast number of cases. It is interesting to see how this strategy allows easy comparison of the results of control strategies. Moreover, since emulators only need to be created once, a very significant amount of effort can be saved in the long run when testing control strategies.
- **The implementation of MPC in real buildings involves dealing with uncertainty**. Instead of trying to obtain a "perfect" control-oriented model (a misleading and unnecessary goal), it is important to work with approximate models, as long as it is possible to incorporate uncertainty as a necessary reality. Advanced modelling techniques are provided for this. More case studies are required to demonstrate the efficacy of these techniques in real world applications
- Control is also required to better integrate building energy consumption with energy supply systems and improve energy reliability and energy market efficiency. This can be achieved by using **multi-level hierarchical control strategies**. A flexibility function is proposed as a suitable control-oriented model that bridges the building level with the energy system level (or other aggregation levels), taking into account uncertainty.

Much remains to be done to accelerate the deployment of MPC as a widespread solution in the control industry. Annex 81 has inspired the launch of an ASHRAE-sponsored project (ASHRAE 1934-TRP) focused on surveying the points of view of numerous industry practitioners. This study aims to identify barriers and opportunities for overcoming challenges in deploying advanced data-driven MPC and other control strategies. At the time of writing, this research project is being conducted under contract, with results expected by the end of 2025.

ANNEX 81

EBC

Energy in Building and Communities Programme

www.iea-ebc.org

Technology Collaboration Programme

by iea