



Towards scalable and generic data-driven solutions for energy flexibility control in buildings

Ana David

► To cite this version:

Ana David. Towards scalable and generic data-driven solutions for energy flexibility control in buildings. Automatic. Université Grenoble Alpes [2020-..], 2023. English. NNT : 2023GRALT017 . tel-04121671

HAL Id: tel-04121671

<https://theses.hal.science/tel-04121671>

Submitted on 8 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

UNIVERSITÉ GRENOBLE ALPES

THÈSE

pour obtenir le grade de

DOCTEUR DE L'UNIVERSITÉ DE GRENOBLE ALPES

Spécialité : **Automatique-Productique**

Arrêté ministériel : 7 août 2006

Présentée par
Ana DAVID

Thèse dirigée par **Mazen ALAMIR**
et co-encadrée par **Claude LE PAPE-GARDEUX**

préparée au sein du laboratoire **GIPSA-Lab** et
Schneider-Electric dans l'école doctorale **EEATS**

Towards scalable and generic data-driven solutions for energy flexibility control in buildings

Vers des solutions extensibles et génériques,
basées sur les données, pour le contrôle de
la flexibilité énergétique dans les bâtiments

Thèse soutenue publiquement le **8 mars 2023**,
devant le jury composé de:

Olivier SENAME

Grenoble INP, Examineur, Président du jury

Geert DECONINCK

KU Leuven, Rapporteur

Hervé GUEGUEN

Centrale Supélec Rennes, Rapporteur

Anna ROBERT

TotalEnergies, Examineur



Remerciements

Malgré ma petite expérience j'ai le sentiment que dans une vie professionnelle il est rare d'être entouré d'autant de personnes remarquables que celles que j'ai eu l'occasion de côtoyer pendant mon doctorat. Je profite donc de cette occasion pour leur adresser ma gratitude.

Je tiens tout d'abord à remercier les membres de mon jury de thèse : Olivier Sename, Geert Deconinck, Hervé Gueguen et Anna Robert. Si je garde un si beau souvenir de l'aboutissement de ces trois années de thèse c'est à grâce à vous. Merci pour vos commentaires constructifs, votre implication et votre gentillesse.

Un très grand merci à Mazen Alamir, mon directeur de thèse, pour son énergie, son soutien invariable et tout le savoir transmis. Merci à Claude Le Pape-Gardeux et Peter Pflaum pour avoir eu confiance en moi et de m'avoir offert cette opportunité. Merci aussi pour votre encadrement par la suite, nos discussions ont toujours été un plaisir. Puisque j'ai la chance de pouvoir continuer mon expérience au sein de Schneider Electric, j'aimerais remercier de nouveau Claude mais également Didier Pellegrin, qui ont pris le soin de m'intégrer durablement dans l'équipe.

Merci à tous les experts et amis de Schneider Electric qui m'ont apporté une inestimable aide technique mais aussi un cadre stimulant et bienveillant où l'on a toujours envie d'y aller le matin. Plus particulièrement merci à Patrick Béguery, Alejandro Yousef Da Silva, Lap Van Ngo, Théo Lagarde, Tristan Rigaud, Alain Malot, Henri Obara pour tous leurs conseils. J'aimerais également remercier chaleureusement Anna Stavrianou pour son soutien et sa bienveillance lors de la transition entre le doctorat et mon nouveau poste.

Je ne pourrais pas finir sans mentionner également mes amis de Gipsa-lab. J'en connais au moins un qui sera déçu de ne pas voir son nom explicitement listé, mais j'espère que vous ne m'en tiendrez pas rigueur, je pense à vous tous. Malgré la musique de Cascada en soirée, je suis si heureuse de vous avoir rencontré !

Enfin, mes plus tendres pensées vont à ma famille (française et moldave) à qui je dois beaucoup, sinon tout.

Abstract

Energy flexibility of an electric utility customer is the capability of changing its power consumption from the normal consumption pattern. Energy flexibility on the demand side is a major resource for ensuring the grid stability, avoiding activation of carbon-intensive energy resources, avoiding costly investments for the reinforcement of the grid infrastructure and construction of new power plants. In this context, buildings have a great potential due to their thermal inertia (capacity of storing and discharging thermal energy).

The objective of this thesis is to study and propose efficient solutions for the estimation and optimal control of energy flexibility in buildings. The diversity of buildings in terms of thermal characteristics, Heating, Ventilation and Air-Conditioning (HVAC) systems, meters, sensors and actuators is considerable. It follows that the flexibility can be generated by different manners, and the exact impact of the different available options is often unknown. Given this complexity, available methods are rarely scalable nor easily implementable.

The solutions investigated in this thesis are designed to address real-life implementation issues: availability of data, scalability, genericity, integration to available building management systems, thermal comfort requirements. Consequently, they are based on commonly available data (sensors, meters and actuators) in buildings and do not impact the thermal comfort levels defined by the building managers.

A solution that addresses the above objective and implementation constraints is HVAC load control via indoor temperature setpoint adjustments. The first part of this thesis is therefore focused on exploration of methods for day-ahead forecasting of the total HVAC power consumption profile, based on commonly available features: global indoor temperature setpoints, indoor temperature, weather data, time related data. Since indoor temperature reflects the thermal state of the building, similar methods have been investigated for forecasting its dynamics. Cascaded predictions have been tested, which imply using indoor temperature forecasts for predicting the power consumption. Regarding this option, results show that the propagation of the indoor temperature forecasts errors is prohibitive.

Based on the first investigations, a predictive model for power consumption forecasting has been implemented, based on supervised learning. To capture the dynamical behavior of power consumption, without forecasting the evolution of indoor temperature, actual and past values of the selected features are used for predicting the actual HVAC load.

A second part of the thesis is focused on refining and integrating the predictive model in a demand response framework, consisting in load shedding incentives. Load shedding is achieved using inherent energy storage capability of the building (preheating/precooling). Based on a simplistic demand response program that compensates the shedded load, an optimization-based strategy has been proposed to perform optimized load shedding operations. An estimation of the possible income based on the given program has been carried out.

This thesis has been prepared as part of a partnership bewteen GIPSA-lab laboratory, Schneider-Electric and MIAI Grenoble Alpes Institute, supported by Association Nationale de la Recherche et de la Technologie (ANRT).

Résumé

La flexibilité des clients, vis-à-vis du système électrique, désigne la capacité de modifier temporairement leur profil de consommation. Il s'agit d'une ressource non-négligeable pour la gestion efficace du réseau électrique : réduction de la consommation d'énergie carbonée, maintien de la stabilité du réseau, atténuation des courbes de la demande électrique, limitation des investissements dans la construction des nouveaux moyens de production et dans l'extension ou la consolidation du réseau électrique. Dans ce cadre, les bâtiments ont un grand potentiel qui peut être exploité grâce à l'inertie thermique qui les caractérise (capacité de stocker et restituer l'énergie sous forme de chaleur).

L'objectif de cette thèse est la conception de solutions pour estimer et piloter de manière optimisée la flexibilité énergétique des bâtiments. La diversité des bâtiments en termes de caractéristiques thermiques, systèmes de Chauffage, Ventilation et Climatisation (CVC), actionneurs et capteurs accessibles est considérable. Il en résulte que la flexibilité peut être produite de différentes manières et l'impact exact des nombreuses options disponibles est souvent mal connu. Compte tenu de cette complexité, les méthodes qui peuvent être conçues et déployées facilement font défaut.

Les solutions étudiées cherchent à répondre à l'objectif fixé tout en considérant des problématiques d'implémentation réelles : accès aux données, généricité, intégration aux systèmes de gestion existants, barrières d'acceptation liées au confort thermique. Ainsi, elles sont basées sur des données (capteurs, compteurs et actionneurs) généralement disponibles et communs à la plupart des bâtiments, et ne dégradent pas le niveau de confort thermique établi par les gestionnaires de bâtiments.

Une solution qui répond aux objectifs et contraintes fixées consiste dans le pilotage de la consommation énergétique par l'intermédiaire du contrôle de la consigne de température intérieure. La première partie de cette thèse est ainsi focalisée sur l'exploration des méthodes visant à prédire, pour un horizon de 24 heures, le profil de puissance totale consommée dans un bâtiment, en fonction de variables explicatives disponibles : consignes globales de la température intérieure, température intérieure, données météo. Puisque la température intérieure reflète l'état thermique dans un bâtiment, des méthodes similaires pour prédire sa dynamique ont été étudiées. Des solutions de prédiction en cascade ont notamment été testées, impliquant l'utilisation des prédictions de la température intérieure pour la prédiction du profil de puissance consommée. Sur ce point, les résultats ont montré un impact rédhibitoire des erreurs de prédictions de la température.

Sur la base des premières investigations, un modèle prédictif de la puissance consommée, basé sur des outils d'apprentissage supervisé a été adopté. Afin de capturer la dynamique, sans prédire l'évolution de la température intérieure, les valeurs actuelles et passées des variables explicatives sont utilisées pour la prédiction de la puissance consommée à un instant donné.

La deuxième partie de la thèse est dédiée au raffinement et l'utilisation de l'outil de prédic-

tion développé, dans le cadre d'un programme de réponse à la demande, incitant à la réduction de la consommation énergétique. L'effacement de la consommation est accompli en utilisant les capacités intrinsèques de stockage thermique d'un bâtiment (préchauffage/pré-climatisation). Sur la base d'un programme rudimentaire, rémunérant l'effacement de la consommation, une stratégie d'optimisation a été mise en place pour effectuer des effacements de manière optimisée. Une estimation des bénéfices tirés de l'intégration de ce programme a été faite.

Cette thèse a été réalisée dans le cadre d'un partenariat entre Schneider Electric, GIPSA-Lab et l'Institut MIAI Grenoble Alpes, avec le soutien de l'Association Nationale de la Recherche et de la Technologie (ANRT).

Contents

1	General introduction	1
1.1	Electrical grid transformation and the active role of buildings	1
1.2	Aims of the thesis	2
1.3	Outline of the thesis	3
I	Background and problem setting	5
2	Control and optimization of HVAC systems - objective and challenges	7
2.1	Buildings in the new energy landscape	7
2.2	A short overview of HVAC systems	10
2.3	Background on control and optimization of HVAC systems	11
2.3.1	Model-based optimal control	12
2.3.2	Model-free optimal control	19
2.4	Conclusion	21
3	Problem setting and adopted approach synthesis	23
3.1	Investigated use-case and adopted approach	23
3.2	Case studies	30
3.2.1	An IDA ICE simulated building	30
3.2.2	A Synthetic Building Operation Dataset	31
3.2.3	A Digital Twin of a large office building	32
3.2.4	An Equivalent Thermal Parameter simulated building	32
3.3	Adopted approach synthesis	33
II	Exploratory work on HVAC load and indoor temperature data-driven	

modelling	37
4 Data-driven modelling for power consumption and indoor temperature forecasting	39
4.1 Errors metrics	41
4.2 Linear state-space modelling	42
4.2.1 Linear state-space modelling of the thermal dynamics	43
4.2.2 Linear state-space modelling using indoor temperature setpoints	47
4.3 Autoregressive Random Forest regression predictive modelling	50
4.4 Random Forest regression predictive modelling	57
4.4.1 Structured data-driven modelling of indoor temperature	60
4.4.2 RF regression predictive modelling of power consumption only	64
4.5 Conclusion	67
III Data-driven modelling for HVAC load forecasting in a DR context	69
5 Predictive modelling using boosted trees	71
5.1 Structure of the solution	71
5.2 Validation on available building case-studies	73
5.3 Construction of the Demand Response learning dataset	76
5.4 Zoom on preheating and shedded energy depending on the preheating parameters	78
5.5 Zoom on forecasting accuracy during Demand Response events	80
6 Economical evaluation of the energy flexibility potential	83
6.1 Demand Response mechanism and participation strategy	84
6.2 Optimization of Demand Response flexibility	86
6.2.1 Delve into optimization results and impact of forecasting errors	89
6.2.2 Building characteristics impact on flexibility capability	91

Contents	ix
6.3 Comparison with a classical load shedding approach	92
6.4 Conclusion	94
General conclusion and perspectives	95
Appendices	99
A Schematic view of the IDA ICE simulated building	101
B AR RF model for power consumption and indoor temperature estimation	103
B.1 ETP building	103
B.2 IDA ICE building with night setback temperature control	105
B.3 IDA ICE building with DR representative step setpoints	106
C Forecasting accuracy during Demand Response events	109
D Load shedding by lowering the temperature setpoints	111
E Résumé en français	113
E.1 Introduction	113
E.1.1 Contexte général	113
E.1.2 État de l'art	114
E.1.3 Objectif des travaux et approche investiguée	115
E.2 Étude exploratoire des méthodes de prédiction de la consommation énergétique et de la température intérieure	116
E.3 Modélisation prédictive pour l'exploitation de la flexibilité énergétique	118
E.4 Évaluation économique du potentiel de flexibilité énergétique	119
E.5 Conclusion générale et perspectives	120
Bibliography	127

List of Figures

2.1	Types of DR programs.	8
2.2	An example of power system balancing in a deregulated electricity market. . . .	9
3.1	An example of Demand Response, given a comfort range $[T_{sp}^{min}, T_{sp}^{max}]$. The global temperature setpoints are increased to 25°C before a DR event. This allows passive accumulation of heat in the air and the building structure, that is released at the beginning of the DR event, by resetting the setpoints back to their nominal value.	25
3.2	A simplified building example. The building is composed of several zones and a global HVAC system with production and distribution elements.	26
3.3	BMS coupling in the simplified building example. Data from sensors and meters are collected by the BMS. Through the BMS, the user defines the default temperature setpoints for each zone. Given a DR incentive, the integrated optimization brick computes and communicates a new global setpoint schedule that is applied to the individual zones.	27
3.4	Mean temperature in the building (15 zones) compared to three distinct zone temperatures. The power consumption variation regarding the relation between between global temperature setpoints and mean temperature is less obvious due to zone disparities.	28
3.5	Test building representation. The building represents a medium-size office building with 15 zones.	30
4.1	Explored modelling approaches and the modelling/forecasting targets. Six modelling approaches are discussed, designed for indoor temperature, total HVAC load, or both. Starting with state of the art linear models, then cascaded predictions of indoor temperature and HVAC load, followed by HVAC load prediction only.	40
4.2	R^2 and WAPE errors. Subfigure (a) illustrates the case where the predictions are a translation of an original signal. Subfigure (b) illustrates the case where predictions include a noise sine signal (same amplitude and frequency) added to three different original signals.	42
4.3	Linear state-space identification of the thermal dynamics model. Schematic representation of the approach, input variables and outputs.	43

4.4	State-space modelling given partially known heat gains affecting the zones, with $u = (T_{in}, T_{ext}, \Phi_{irr})^T$ and output $y = P$. ETP building.	45
4.5	State-space modelling given partially known heat gains affecting the zones, with $u = (P, T_{ext}, \Phi_{irr})^T$ and output $y = T_{in}$. ETP building	45
4.6	State-space modelling given partially known heat gains affecting the zones, with $u = (T_{in}, T_{ext}, \Phi_{irr})^T$ and output $y = P$. IDA ICE building.	46
4.7	State-space modelling given partially known heat gains affecting the zones, with $u = (P, T_{ext}, \Phi_{irr})^T$ and output $y = T_{in}$. IDA ICE building	46
4.8	Linear state-space model identification using indoor temperature setpoints. Schematic representation of the approach, input variables and outputs.	47
4.9	State-space modelling given partially known heat gains affecting the zones, with $u = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$ and output $y = P$. ETP building.	48
4.10	State-space modelling given partially known heat gains affecting the zones, with $u = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$ and output $y = T$. ETP building.	49
4.11	State-space modelling given unknown disturbances, with $u = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$ and output $y = P$. IDA ICE building.	49
4.12	State-space modelling given unknown disturbances, with $u = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$ and output $y = T$. IDA ICE building.	50
4.13	Auto-regressive Random Forest regression predictive modelling of power consumption and indoor temperature. Schematic representation of the approach, input variables and outputs.	50
4.14	Error propagation through the 24h prediction horizon for indoor temperature (left) and power consumption (right) models. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on a 9-zone ETP building with night setback temperature control.	51
4.15	Error propagation through the 24h prediction horizon for indoor temperature (left) and power consumption (right) models. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on IDA ICE building with night setback temperature control.	52
4.16	Error propagation through the 24h prediction horizon for indoor temperature (left) and power consumption (right) models. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on IDA ICE building with simulated preheating operations.	52

- 4.17 1-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on a 9-zone ETP building with night setback temperature control. 54
- 4.18 24-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on ETP building with night setback temperature control. 54
- 4.19 1-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on IDA ICE building with night setback temperature control. 55
- 4.20 24-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on IDA ICE building with night setback temperature control. 55
- 4.21 1-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on IDA ICE building with simulated preheating operations. 56
- 4.22 24-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on IDA ICE building with simulated preheating operations. 56
- 4.23 Random Forest regression predictive modelling of power consumption and indoor temperature. Schematic representation of the approach, input variables and outputs. 57
- 4.24 RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1, \dots, t-N})$, where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$ and $N=8$ (2 hours), and power consumption: $P_t \approx f_{RF}(X_t)$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$. Validation on ETP building with night setback temperature control. 59
- 4.25 RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1, \dots, t-N})$, where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$ and $N=8$ (2 hours), and power consumption: $P_t \approx f_{RF}(X_t)$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$. Validation on IDA ICE building with night setback temperature control. 59

4.26	RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N})$, where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$ and $N=8$ (2 hours), and power consumption: $P_t \approx f_{RF}(X_t)$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$. Validation on IDA ICE building simulated using step setpoints representative of DR control. . . .	60
4.27	Structured data-driven modelling of indoor temperature and RF regression predictive modelling of power consumption. Schematic representation of the approach, input variables and outputs.	60
4.28	Indoor temperature predictions based a nonlinear explicit model. RF estimator for power consumption: $P_t \approx f_{RF}(X_t)$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$. Validation on IDA ICE building simulated using ramp preheating setpoints. . .	63
4.29	RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N})$, where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$ and $N=8$ (2 hours), and power consumption: $P_t \approx f_{RF}(X_t)$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$. Validation on IDA ICE building simulated using ramp preheating setpoints.	63
4.30	RF regression predictive modelling of power consumption. Schematic representation of input variables and output.	64
4.31	RF power consumption forecasting results based on past and current feature values: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N})$, where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$ and $N=8$ (2 hours). Validation on ETP building with night setback temperature control.	65
4.32	RF power consumption forecasting results based on past and current feature values: $y_t \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N})$, where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$ and $N=8$ (2 hours). Validation on IDA ICE building with night setback temperature control.	65
4.33	RF power consumption forecasting results based on past and current feature values: $y_t \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N})$, where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$ and $N=8$ (2 hours). Validation on IDA ICE building with representative DR preheating operations.	66
5.1	Feature set for power consumption forecasting.	72
5.2	Augmenting the original dataset with rolling window averages of the data in order to smoothen the forecasts and reinforce the dependency between the target and the features.	72
5.3	Fitting results of the adopted solution. Validation on IDA ICE building simulated with two random preheating scenarios each working day. Train data: November 2001 - March 2002, test data: November 2002 - March 2003.	73

5.4	Power consumption forecasting results of the adopted solution. Validation on IDA ICE building simulated with two random preheating scenarios each working day. Train data: November 2001 - March 2002, test data: November 2002 - March 2003.	73
5.5	HVAC electrical load forecasting results of the adopted solution. Validation on Digital Twin building. Test data: 30% of the available June-October 2019 data.	74
5.6	HVAC electrical cooling load forecasting results of the adopted solution. Validation on a Synthetic building operation dataset. The dataset corresponds to a building with standard efficiency, San Francisco location and typical meteorological year. For testing purposes, 30% of the available data was used.	75
5.7	HVAC electrical heating load forecasting results of the adopted solution. Validation on a Synthetic building operation dataset. The dataset corresponds to a building with standard efficiency, Chicago location and typical meteorological year. Two different winter seasons are used for training and testing purposes.	76
5.8	Power consumption profile given preheating operations and baseline load profile (under normal temperature setpoint control) which reflect the lasting effects of frequent preheating operations.	77
5.9	Example of preheating scenarios for one day. Given the fixed start of shedding at 17:00, preheating operations with temperature setpoints [22, 23, 24, 25 °C] and duration of [0.5, 1, 1.5, 2, 2.5, 3 hours] were simulated. The baseline scenario is defined by a temperature setpoint of 21 °C during the day and 17 °C during the night and weekends.	78
5.10	Preheating and shedded energy variation depending on the preheating temperature setpoint and duration. The considered shedding duration is 1 hour	79
5.11	Preheating and shedded energy variation depending on the day. The considered shedding duration is 1 hour and the preheating temperature setpoint is 22 °C.	79
5.12	Power consumption forecasting accuracy on selected preheating timesteps.	80
5.13	Power consumption forecasting accuracy on selected shedding timesteps (1 hour time-interval after a preheating operation).	81
5.14	Baseline power consumption forecasting accuracy on selected preheating and shedding timesteps.	81
6.1	Example of a DR event. Participation to a given DR event is determined by the balance between the preheating energy $\Delta E^{preheat}$ and the remunerated shedded energy ΔE^{shed}	84

6.2	Day-ahead DR participation optimization which consists in determining the optimal preheating duration, preheating temperature setpoint and start of the shedding.	86
6.3	DR participation income relative to the baseline energy cost. Training on Nov. 2001 - March 2002. DR participation during Nov. 2002 - March 2003.	87
6.4	Number of triggered events (estimated positive benefit $\hat{B} > 0$).	87
6.5	DR participation income relative to the baseline energy cost. Initial training on Nov. 2002, and retraining every month. The continuous learning case income is compared to the default one, where a whole winter season dataset is used for learning before integrating the DR program. DR participation during Dec. 2002 - March 2003.	88
6.6	Optimality of the decisions with respect to the start of shedding t_{start}^{shed} . Subfigure (a) illustrates results for the Demand Response (DR) events in the morning and subfigure (b), for the DR events in the evening. According to the matrix, $t_{start}^{shed} = 19:00$ was indeed optimal in 38 evening DR cases. In 15 cases, $t_{start}^{shed} = 18:30$ would have been the optimal decision, instead of 19:00.	90
6.7	Optimality of the decisions with respect to ΔT_{min}^{sp} ($\Delta t^{preheat}$ turned out to be optimal in all cases). According to this matrix, preheating with $\Delta T_{min}^{sp} = 4^\circ\text{C}$ was indeed optimal in 59 DR events. In 7 cases $\Delta T_{min}^{sp} = 4^\circ\text{C}$ would have been optimal, instead of the decided $\Delta T_{min}^{sp} = 3^\circ\text{C}$	91
A.1	Schematic view of the building.	101
A.2	Schematic view of the primary system.	102
A.3	Schematic view of the AHU system.	102
B.1	Error propagation through the 24h prediction horizon for indoor temperature (left) and power consumption (right) models. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{T}_{in,t-1,\dots,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{P}_{t-1,\dots,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, $N=8$ (2 hours). Validation on ETP multizone building with night setback temperature control.	103
B.2	1-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{T}_{in,t-1,\dots,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{P}_{t-1,\dots,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, $N=8$ (2 hours). Validation on ETP multizone building with night setback temperature control.	104

- B.3 24-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}, \bar{T}_{in,t-1,...,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}, \bar{P}_{t-1,...,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, N=8 (2 hours). Validation on ETP multizone building with night setback temperature control. 104
- B.4 Error propagation through the 24h prediction horizon for indoor temperature (left) and power consumption (right) models. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}, \bar{T}_{in,t-1,...,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}, \bar{P}_{t-1,...,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, N=8 (2 hours). Validation on IDA ICE building with night setback temperature control. 105
- B.5 1-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}, \bar{T}_{in,t-1,...,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}, \bar{P}_{t-1,...,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, N=8 (2 hours). Validation on IDA ICE building with night setback temperature control. 105
- B.6 24-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}, \bar{T}_{in,t-1,...,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}, \bar{P}_{t-1,...,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, N=8 (2 hours). Validation on IDA ICE building with night setback temperature control. 106
- B.7 Error propagation through the 24h prediction horizon for indoor temperature (left) and power consumption (right) models. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}, \bar{T}_{in,t-1,...,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}, \bar{P}_{t-1,...,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, N=8 (2 hours). Validation on IDA ICE building with simulated preheating operations. 106
- B.8 1-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}, \bar{T}_{in,t-1,...,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}, \bar{P}_{t-1,...,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, N=8 (2 hours). Validation on IDA ICE building with simulated preheating operations. 107
- B.9 24-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}, \bar{T}_{in,t-1,...,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}, \bar{P}_{t-1,...,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, N=8 (2 hours). Validation on IDA ICE building with simulated preheating operations. 107
- C.1 Mean power consumption forecasting accuracy during preheating. 109

C.2	Mean power consumption forecasting accuracy during shedding (1 hour time-interval after a preheating operation).	109
C.3	Mean baseline power consumption forecasting accuracy during preheating. . . .	110
C.4	Mean baseline power consumption forecasting accuracy during shedding.	110
D.1	Power consumption profile under shedding operations by lowering the temperature setpoints, and baseline load profile (under normal temperature setpoint control).	111
E.1	Un exemple d'opération de DR. La consigne T_{sp}^{min} dans le bâtiment est augmentée à 25°C avant un évènement de DR. Cela permet une accumulation d'énergie dans l'air et les éléments structurels du bâtiment, qui peut être déchargée à la demande, en remettant la consigne à sa valeur de 21°C par défaut.	115
E.2	Méthodes de modélisation explorées et variables à prédire. Les flèches indiquent une prédiction en cascade, de la température intérieure qui est utilisée ensuite pour la prédiction de la puissance consommée.	117

List of Tables

4.1	Simulation with night setback temperature control. First three approaches are based on chain indoor temperature and HVAC power consumption forecasting, while the last two methods bypass usage of indoor temperature predictions. . . .	68
4.2	Simulation with step setpoints, equivalent to preheating operations for DR. First three approaches are based on chain indoor temperature and HVAC power consumption forecasting, while the last two methods bypass usage of indoor temperature predictions.	68
4.3	Simulation with ramp setpoints. Approaches are based on chain indoor temperature and HVAC power consumption forecasting. The first approach uses RF regression indoor temperature prediction, while the second one uses a structured dynamical model. Power consumption is forecasted using RF regression. .	68

Abbreviations

BMS Building Management System.

BRP Balance Responsible Party.

BSP Balancing Service Provider.

DOW Day Of Week.

DR Demand Response.

ETP Equivalent Thermal Parameter.

HVAC Heating, Ventilation, and Air-Conditioning system.

ML Machine Learning.

MPC Model Predictive Control.

NEBEF Notification d'Échanges de Blocs d'Effacement.

RC Resistance and Capacitance.

RF Random Forest.

RL Reinforcement learning.

TOD Time Of Day.

TOU Time Of Use.

TSO Transmission System Operator.

WAPE Weighted Absolute Percentage Error.

General introduction

1.1 Electrical grid transformation and the active role of buildings

The electrical grid is evolving, from centralized production assets and passive interactions between the main electrical system and end-users, towards distributed smaller power sources and complex energy transfers among grid participants. The sector is subject to increasing uncertainties - electricity market deregulation, intermittent nature of certain grid assets, multiplication of electric vehicles - that are all additional challenges for operating the grid. In this context, a strategic shift from supply that adapts to demand, to demand that also adapts to supply, has been engaged.

In a potentially highly electrified future, the importance of DR, as a mean of inducing changes in the electric usage of end-users, is enhanced. Applications are possible on the main electrical grid as well as in local microgrids. In the first case the benefits are multiple, for instance reducing carbon intensive power plants usage, consolidating grid stability reserves, avoiding investments in grid extensions and construction of new power plants. In the second case, DR allows optimising local energy usage and maximising the grid resilience and reliability.

Some DR programs have been around since a long time now, for instance Time Of Use (TOU) electricity rates, that discourage energy consumption when the cost of electricity generation is high. DR, and more generally Demand Side Management, is however evolving. New opportunities are confirmed, markets rules are established, new control strategies are adapted for optimizing grid management. On the end-user side, the capability to answer to DR incentives depends, among other things, on the DR program requirements and the customer operational constraints.

Buildings, in this context, have been identified as a major source of energy flexibility due to several elements: the significant part of the total energy consumption that the building sector accounts for, the inherent ability to store energy under a thermal form, the low investment needed for energy flexibility management. Potential value exists in all sorts of buildings: conventional or energy efficient, residential or tertiary buildings, etc. Energy efficiency at a local level has been the major focus of optimization strategies in conventional buildings. DR is inciting buildings to participate actively to the grid management, thus bringing new challenging problems from a broader perspective.

Accounting for the major energy consumption in a building, the contribution from Heating, Ventilation, and Air-Conditioning system (HVAC) systems is crucial. Adjusting the HVAC load might involve different subsystems and actuators: chilled/hot water temperature setpoints, mass flow rates, air flow rates, supply air temperatures, zone temperature setpoints and humidity levels, etc. The interactions between the subsystems, and the various nonlinear behaviours, makes the control very complex. Furthermore estimating the load impact of a particular control action is challenging. HVAC load control problems have been tackled in numerous ways using classical model-based or model-free optimal control frameworks. The methods found in literature and applied in practice differ greatly in modelling complexity, optimal policy training duration, the diversity of actuators and sensors, occupant's comfort impact, precision, etc.

1.2 Aims of the thesis

The main focus of the thesis is the design of a suitable data-driven model, based on the available sensors and meters, that allows accurate estimations of HVAC load with respect to the control actions. Because the indoor temperature contains important information about the past dynamics of the building, attempts for its accurate estimation in a day-ahead framework were also made. In a second stage, evaluation of DR benefits is carried out on a complex simulated building.

The considered approaches, in contrast with methods found in the state of the art, have the following technical features:

- They are based only on common available data such as indoor zone temperatures, indoor temperature setpoints, HVAC total power consumption, weather conditions
- They use indoor temperature setpoints as control actions

It follows that the designed model is not tailored for a specific building, HVAC system or DR mechanism, does not require extensive system data, is capable of capturing complex system nonlinearities, enables explicit comfort control/maintenance, allows simple implementation with no additional investments for installing sensors and actuators. The targeted use-case is a day-ahead DR program that offers financial compensation for load reduction during strategic time slots. In order to maintain indoor thermal comfort, load reduction is achieved by load shifting, i.e. preheating or precooling the building. The associated benefit is determined by the balance between the load reduction and the increase of the total power consumption occurring with the precooling/preheating. The data-driven model is therefore used to estimate the optimal precooling/preheating operation.

1.3 Outline of the thesis

The thesis manuscript is organized as follows:

A first introductory part is articulated around the new challenges for energy management in buildings and state of the art solutions.

Chapter 2 Control and optimization of HVAC systems - objective and challenges

This first chapter introduces the subject of Demand Response programs. Follows an overview of HVAC system processes and components. The subject of optimal control of HVAC systems is then introduced, with examples of classical methods. The key features/pain points of each category, with respect to the desired application, are emphasized, motivating the adopted approach.

Chapter 3 Problem setting and adopted approach synthesis

In this second chapter, the adopted approach for enabling the energy consumption flexibility in buildings is briefly described, along with the data used for the assessment, and a summary of the final solution and results.

In a second part, HVAC power consumption and indoor temperature forecasting is investigated in a quite experimental way, covering questions such as building behaviours under temperature setpoint changes, model structure and explanatory variables, sources of modelling errors.

Chapter 4 Data-driven modelling for power consumption and indoor temperature forecasting

In this chapter, experimental modelling approaches are tested for short term forecasting of HVAC load. Since indoor temperature reflects the current and past thermal states of the building, equivalent techniques have been studied for indoor temperature forecasting as well. More precisely, have been tested linear state-space models, autoregressive random forest predictive models, "classical" random forest predictive models, as well as structured explicit models (for indoor temperature only).

In a third part, predictive modelling is investigated mainly under the angle of Demand Response.

Chapter 5 Predictive modelling using boosted trees

Given the elements from the previous part, the adopted solution is described: features, learning experiments, validation on available buildings.

Chapter 6 Economical evaluation of the energy flexibility potential

In this chapter the assumed DR mechanism is described, along with the optimal participation strategy. Economical results of the energy flexibility valued in such a program is provided. Besides, the impact of some thermal characteristics of the building is discussed. An economical comparison with the usual way of ensuring load shedding, by lowering the

setpoints, is also provided.

General conclusion

The general conclusion discusses the main results and highlights the challenges and uncertainties that still have to be answered in order to provide a operational solution.

Part I

Background and problem setting

Control and optimization of HVAC systems - objective and challenges

2.1 Buildings in the new energy landscape

The evolution of communication technologies, the development of concepts such as smart buildings together with Building Management Systems (BMS) spreading and the subsequently large amount of data generated, fostered research studies which focus on modelling, forecasting, retrofitting and controlling issues in buildings. For a very high-level view of the new opportunities that the data "deluge" leveraged, readers may refer to (Alanne et al. 2022). Large amounts of data and highly connected devices allow functions such as: monitoring, predicting, controlling, interacting and managing (Alanne et al. 2022). Besides the possible enhancements of energy performance and comfort, efforts have been deployed to enable buildings to become truly active components of the electrical grid. The importance of DR and the need of flexibility on the consumer side has been highlighted since the increasing share of renewable energy and distributed storage systems in the electrical grid. Indeed, the inherent uncertainty and limited controllability of some grid assets motivated stakeholders to seek balance between demand and supply by intelligent management of the customer power consumption. DR programs were thus introduced to motivate users to transition to an active role in planning and operating the power systems, by adapting their consumption to some signals that depend on the electrical grid constraints or the market. Rigorously, the Federal Energy Regulatory Commission defines DR as "Changes in electric usage by demand-side resources from their normal consumption patterns in response to changes in the price of electricity over time, or to incentive payments designed to induce lower electricity use at times of high wholesale market prices or when system reliability is jeopardized".

DR programs are based on motivations offered for reducing the power consumption. Usually DR programs are categorized based on the type of motivation, as illustrated in fig. 2.1. Implicit motivations (price-based and non-dispatchable) are based on the end-user reaction to price signals, for instance TOU, Real-Time Pricing, Critical Peak Pricing, etc. Explicit motivations (committed and dispatchable) are based on rewards offered for the reduction of the power consumption at certain moments, for instance Direct Load Control, interruptible load programs, capacity reserves market, etc. By reflecting in this way the supply and demand balance, the marginal cost of electricity and more generally the global cost of the utility, DR programs encourage users to adapt their energy consumption and avoid activation of carbon-

intensive energy resources during peak hours. On a longer timescale, this helps avoiding costly investments for the reinforcement of the grid infrastructure and construction of new power plants. Extensive information about DR schemes, stakeholders and load management strategies in smart grids is given in (Vardakas et al. 2014) and (Alizadeh et al. 2012).

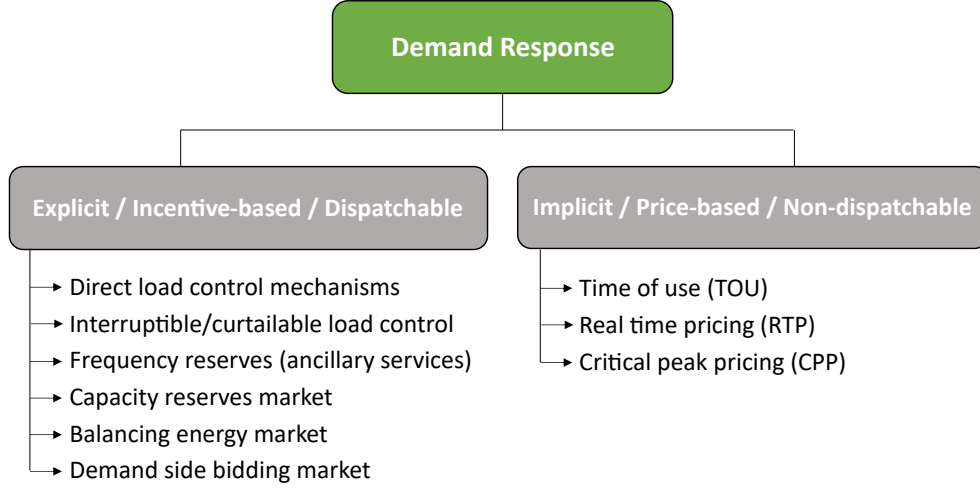


Figure 2.1: Types of DR programs.

Buildings in particular have been successfully integrated in Direct Load Control programs, given the significant electricity consumption and their flexible demand. In France, for example, Direct Load Control schemes for buildings are integrated in ancillary services programs such as frequency control. For this service, buildings are usually grouped by an intermediary market participant, called an aggregator, in order to form a power reserve. The national transmission system operator (RTE) contracts participation agreements with reserve providers and activates the reserves as needed in order to ensure grid balance in real time. Apart from this example, flexibility products can be remunerated for different services and through different market mechanisms. There is no unified common architecture of such mechanisms among the European countries, and even less along continents. Even though several DR programs have common features among many countries, their implementation in practice varies substantially. In France, the explicit flexibility products can take the form of capacity or energy products and can be valued on several, possibly interdependent markets. In the first category the participants are remunerated based on a fixed price (in €/MW) for modulating their consumption during a given time-slot. The energy is necessarily remunerated on an energy market. Several mechanisms are available: demand response tenders, capacity market, manual frequency restoration and replacement reserves tenders, frequency containment reserve tenders, interruptability tenders. In the second category, participants are remunerated (in €/MWh) for an effective load reduction or generation for a determined power and duration. The energy can be valued via frequency containment or restoration mechanisms, balancing mechanisms, Notification d'Échanges de Blocs d'Effacement (NEBEF) mechanism.

Fig. 2.2 below illustrates the interactions among participants for power system balancing in the French context, where the electricity system is a deregulated market. Two key roles

are played by the Balance Responsible Parties (BRPs) and the Balancing Service Providers (BSPs). The first balancing level is ensured by the Balance Responsible Party (BRP). Any activity on the French electricity market (generation, consumption, selling, purchasing) must be performed under the responsibility of a BRP, who is responsible for settling imbalances between injections and withdrawals in its balance perimeter. A second level is provided by the Balancing Service Provider (BSP) and the Transmission System Operator (TSO) who ensures the balancing via dedicated mechanisms. An end-user for instance can participate to the power system balancing through different paths. The participation by contracting directly with the BRP or the BSP is possible as well as through intermediate participants such as aggregators or retailers. Globally the involved market mechanisms, the interactions between the stakeholders, the contractual commitments (duration of events, notification periods, events, etc.) and the remuneration for the products are complex and under constant evolution.

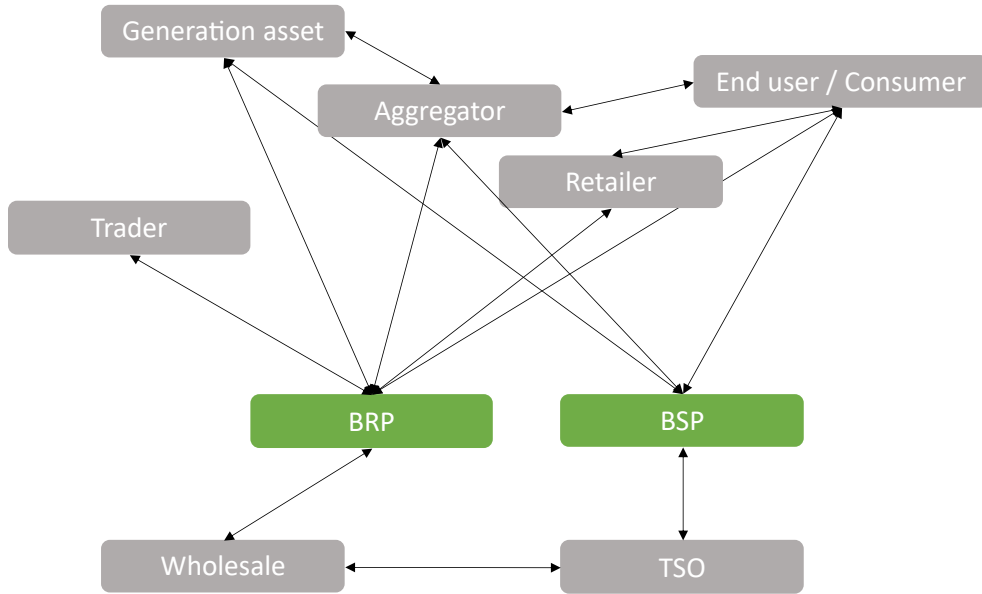


Figure 2.2: An example of power system balancing in a deregulated electricity market.

Despite the complexities, the flexibility potential of buildings thanks to their thermal inertia has been investigated and pointed out. One advantage of integrating systems like buildings in DR programs is the cost effectiveness (Olsthoorn et al. 2017), compared, for instance, to building new power plants to ensure the demand at all times. Furthermore, given the advances in building management systems and building standards, newly constructed buildings are very energy-efficient and sometimes energy self-sufficient. As a consequence, energy efficiency is not the primary goal anymore. DR schemes allow in this case new streams of revenue, ways of lowering the energy bill or efficient management of microgrids.

In this work, the focus is put on thermal mass activation by precooling or preheating (in summer and winter respectively). In the case of precooling, the energy storage is activated by lowering the indoor temperature setpoints, for instance when the energy price is low. When the prices become high, the stored thermal energy is released by resetting the temperature setpoints to their nominal values, resulting in lower energy demand. In this case the cooling or

the heating load is shifted in time according to the DR incentive. This operation can be seen as a form of short-term energy storage that can be charged and discharged when convenient in a control strategy that operates zone air temperatures. From a physical point of view, thermal mass activation is a consequence of conditioning the thermal zone air, which enables heat transfer to the building mass. The potential of shifting the power consumption, especially in office buildings due to their typically high thermal mass, can be quite important, but limited however to ensure the thermal comfort of the occupants (Olsthoorn et al. 2017).

Building an optimal strategy for managing the HVAC system is not trivial, especially when occupants comfort is at stake. As pointed out in (Vázquez-Canteli et al. 2019), given the actual low utility prices, DR programs that impact comfort and preferred electricity usage might not be widely accepted. This has been also mentioned in (Schubnel et al. 2021) and highlighted by many clients of Schneider Electric. Therefore, as soon as we leave the domain of rule-based or expert-based strategies for DR, the complexity of finding an optimal solution requires advanced techniques from the field of optimal control and modelling of the considered system. The next part will address this question by referencing studies that are focused on optimal control of HVAC systems, and investigating which techniques might be suited in the DR context.

2.2 A short overview of HVAC systems

The goal of the HVAC system is maintaining in an efficient manner a comfortable and productive environment in buildings, such as residential houses, schools, retails stores, offices factories, etc. This implies maintaining indoor temperature and humidity in a comfortable range, but also ensuring that the air is free from contaminants. Therefore HVAC processes include:

1. heating and cooling, by adding or removing thermal energy to a space in order to raise, lower or maintain temperature at comfortable levels;
2. humidifying and dehumidifying, by adding or removing water vapor to the air in order to raise, lower or maintain the moisture levels in a comfortable range;
3. cleaning, by removing dust, biological contaminants or other particles, to ensure a good air quality;
4. ventilating, by bringing fresh air from outdoor in order to dilute the gaseous contaminants, and more generally circulating and mixing air, to ensure good ventilation and thermal energy transfers.

Each process depends on external factors such as the environment, the climate or the time of the year.

HVAC components fall usually in four categories:

1. primary heating and cooling system, for example boilers, furnaces, expansion coils, chillers, cooling towers, etc.;
2. distribution system, which ensures the transport of heating and cooling to the indoor spaces, for example pumps, pipes, coils, valves, fans, dampers;
3. delivery components which introduce the heating and cooling into the spaces, such as fan-coil units, radiators, diffusers;
4. control components, which manage the operation of all the components for comfort, process and energy efficiency. Some controllers can be connected together in a building management system to optimize comfort energy and facility management.

HVAC systems can be further categorized depending on the used medium for distributing heating and cooling: all-air systems, all-water systems and air-and-water systems. In a particular building we can find central HVAC systems, meaning that primary heating and cooling components provide the service for the entire building, while the other components might be distributed through the building. In contrast, decentralized systems pack all the components into a complete unit which can serve a particular zone.

What is challenging when dealing with HVAC systems is the fact that there are a myriad of possible configurations and components. The task of designing optimal control strategies is therefore a complex one and the approaches can be as numerous as the possible configurations. In the next section we will review some classical methods and highlight their limitations when dealing with such complex systems.

2.3 Background on control and optimization of HVAC systems

The HVAC system of a building is usually the system that affects the most the energy use of the building. As a result, a lot of efforts are concentrated on controlling it efficiently. Optimization of the HVAC processes allows energy cost, CO₂ or energy consumption efficiency, while satisfying the desired indoor comfort. It's worth noting that the cost efficiency doesn't lead necessarily to energy efficiency and that indoor comfort is an antagonist objective with respect to energy consumption and cost.

A detailed overview of different supervisory control classes such as model-based or performance map-based can be found in (S. Wang and Ma 2008). As mentioned in this review, the classical approach to the problem of optimal control is considering overall system characteristics and interactions among HVAC subsystems. Hence the optimization problem might be fundamentally different depending on the HVAC system, the building characteristics and the targeted application, as the references that follow will show.

When dealing with optimization of large and complex systems, the conventional approach is model-based optimization. This implies identification of models that describe the HVAC system energy consumption and the thermal dynamics of the building. Reinforcement learning

(RL) is emerging as a model-free alternative, as it offers a paradigm of learning directly an optimal control policy, based on the quality of the results obtained with controls undertaken in a wide set of circumstances.

Note that expert systems (that can be considered in some cases as a model-free technique) or performance map control are also viable solutions, and are usually implemented as a first step toward efficient control of HVAC systems. Expert systems are based on specific human expert knowledge and consist in determining energy or cost-efficient control settings for HVAC system operation given the operating conditions. Even though such systems are usually automated, the expert rules are often static and depend on the richness of the expert knowledge database. Performance maps are used for controlling efficiently HVAC components and are obtained through simulation or real-world tests of the system over the expected range of operating conditions. Both methods reveal their limitations when dealing with large and complex HVAC systems, and therefore will not be further discussed.

2.3.1 Model-based optimal control

Model-based control is achieved by using models to estimate various variables of interest needed to formulate the underlying optimization problem (objective and constraints). Modelling approaches for buildings and their sub-systems dynamics are usually classified in three categories: white-box, black-box, and grey-box modelling. (X. Li et al. 2014).

White-box modelling The approach referred to as white-box modelling or physical modelling is based on whole-building simulating software that incorporates knowledge based modelling of the underlying phenomena. For example in (Zhao et al. 2015), the authors used a detailed physical model in a Model Predictive Control (MPC) framework, that targets minimization of the HVAC energy consumption (cost function) while maintaining indoor thermal comfort (operational constraints).

In practice, the drawbacks of such an approach are substantial. In the case where a digital twin is developed, the model is built based on accurate system information. Regarding the structural elements, this implies that the most accurate information should be collected, such as thermal characteristics of materials, windows size, zone layout, etc. All the elements of the HVAC system and the corresponding low level controllers should also be accurately modelled. The whole modelling process is time consuming and requires advanced domain expertise. In case the model is slightly simplified, the calibration phase to match simulation outputs to true data can be quite challenging. It goes without saying that such approaches are hardly scalable given the fact that the modelling phase starts from scratch for each investigated building. Moreover, the simulation can be computationally very demanding, and complex to integrate in standard efficient and robust optimization frameworks.

Grey-box modelling Grey-box modelling uses simplified physical descriptions of a building in the form of explicit equations. The modelling is based on physical insights of the thermal behaviour of the building and the HVAC system operation. Therefore the structure of the model is defined by the user, through dynamical differential equations that represent the conduction, convection and capacitance phenomena. Usually the building thermal dynamics can also be represented by using an Equivalent Thermal Parameter (ETP) concept involving a thermal Resistance and Capacitance (RC) network. The complexity of the model varies with the targeted precision in the description of the thermal coupling and layers. For example a building can be considered as a single zone and its external envelope modeled through only one resistance and capacitance, or it can be modelled as multi-zone structure where each external wall has a multi-layer model with a given number of resistors and capacitors.

For an extensive review of the method, including naming conventions, fundamental aspects, modelling approaches and applications, one may refer to (Y. Li et al. 2021) and (Maasoumy et al. 2016). Details about model structure, parameters and RC network order estimations can be found in (Braun et al. 2002).

As an example of such modelling, the authors in (Lee et al. 2008) have used a model-based method to limit peak-demand by thermal mass activation. The considered zones were modeled as a single zone with four structural elements: an external wall, an internal wall, a roof and a floor. Using RC equivalent networks, an energy balance is deduced for each of these four structural elements. The resulting thermal network can also be arranged in a state-space form where the ensemble of state variables is given by two nodal temperatures for each of the four structural elements. The input variables are the indoor, outdoor and ground temperatures, solar radiation absorbed on external walls and on the roof, solar radiation transmitted through windows and absorbed on the floor, internal radiative gains for the floor, external walls, and internal walls, and internal convective gains for the interior air. The output of the model is the rate of instantaneous heat gain to the building air. This model is used to estimate, for a specific zone temperature, the heat transferred to the air, associated to the building structure. A second energy balance is introduced, describing the dynamics of the zone temperature under sensible cooling load and heat gains associated to the structural elements. The derived models are used in an optimization problem to determine the temperature setpoint trajectory that minimizes the peak cooling load during a demand-limiting period. The cooling loads for zones operated with a demand-limiting control are compared to night setup control loads.

Another example of equivalent RC model is given in (S. Wang and X. Xu 2006). In the proposed model, the associated heat transfers for the main elements (walls, roof, windows, internal masses and internal air) are considered, as well as latent, infiltration, occupancy, lights and equipment's heat gains. The paper describes the numerous complexities linked to heat gains estimations, for instance estimating convective and radiative components of internal gain from occupancy, lighting and equipments; estimating solar heat gains on different external wall surfaces or transmitted through windows; etc. The building model was validated by using it to predict the cooling load given the indoor air temperature and, in reverse mode, to predict the uniform indoor air temperature given the measured cooling load and other estimated heat gains.

In (Kelman et al. 2011) authors studied a MPC approach in a typical commercial building

for widely known energy savings programs. They proposed a first-order thermal model to approximate the energy balance and zone temperature dynamics. The controlled variables are the mass flow rates and supply temperatures to each zone by the air handling unit, the fraction of supply flow recirculated from zones and the air handling unit cooling coil outlet air temperature setpoint. The indoor air temperature dynamics is approximated with a first-order model, meaning that the thermal capacitance of the indoor air and the structural elements of the zones are all lumped in a single thermal capacitance. Heat gains associated with occupancy, equipment, and heat transfer with the ambient/exterior air are assumed to be known in advance. The cooling and heating load is estimated using an energy balance of the air handling unit. Finally the system dynamics and other variable equalities (such as power used by the supply fan and the heating and cooling coils, cooling coil inlet and return air temperatures) are expressed in a bilinear form. Bilinear state-space models were also employed in (Lamoudi et al. 2011) and (Kolokotsa et al. 2009), where authors have used physical insights for building modelling intended for control of heating, cooling, lighting and blinds. The complexity of such models is variable, for instance in (Lamoudi et al. 2011) each zone of the building was modelled separately in order to describe the dynamics of the temperature, CO₂ and indoor illuminance. In order to use the model in an energy optimization framework, the power consumption was assumed to exhibit a linear behavior with respect to the control action.

In (Berthou et al. 2014) several RC structures were compared with respect to their ability to predict heating and cooling demands and indoor air temperature. Beyond zone cooling and heating loads, the only heat gains considered were solar heat gains on the walls and transmitted through the windows. The two solar heat gains were estimated using a specific model that uses geometrical information of the building, cloud cover data, solar radiation from the weather file, and time of the day data. The HVAC system is not modeled but is assumed to have a known limited heating/cooling power.

More detailed modelling allows explicit representation of the control influence on the power consumption. For instance in (Sun, Luh, Jia, and Yan 2013) and (Sun, Luh, Jia, Jiang, et al. 2013) a system specific control was investigated. Two HVAC components are controlled: the general fresh air unit and the fan coil unit of each zone. The corresponding actions are fresh air unit outlet air temperature, fan coil unit outlet air temperatures in each zone, fresh air unit air flow rates and fan coil unit air flow rates in each zone. To solve the optimization problem, the state of the system (indoor air temperature, wall temperature, indoor humidity, and indoor CO₂ concentration for each zone) is estimated based on equations of energy and mass conservation, which depend on the control actions. In addition, the energy consumption of chillers, pumps, and cooling towers and the fans in the fresh air unit and fan coil units, is estimated using specific dedicated models and physical insights.

Grey-box modelling still faces several challenges. Depending on the HVAC system, the building configuration and the optimization objective, one has to decide what physical phenomena should be modelled in detail. Depending on the complexity of the model, the number of parameters to estimate can be very high. Finding a good balance between accuracy and simplicity is a difficult task in this kind of applications.

The main difficulty comes however from the fact that the paradigm relies on modelling heat transfers in a zone. The final equations are describing evolution of indoor temperature in a zone given a set of heat gains (cooling/heating, occupancy, irradiance, etc.) or inversely - the cooling/heating gains given indoor temperature dynamics and heat gains from auxiliary sources. This requires the estimation of the heating or cooling power transferred to the zone, as well as the rest of the heat gains, which is far from being a straightforward task. For example, estimating the heating or cooling load per zone, when the production system is centralized, might be impossible without specific sensors. If the heating system is regulated using a heating curve depending on the outside temperature, the estimation of quantities such as maximum power is very difficult. Therefore, the operation complexity, the interaction between the different HVAC subsystems, the system constraints and the nonlinearities of the production and distribution system makes it very difficult to build a mathematical knowledge-based model to predict the energy consumption.

When integrated in a control framework, modelling the dynamics of power consumption with respect to the chosen set of actuators is necessary. Most frequently this is dealt with by assuming control of heat load directly, or assuming a simple law, for instance a linear behavior, between the actuator and the power consumption. Several approaches can be found in literature that deal with this issue: building a mathematical model of the HVAC system, that integrates a wide set of actuators and sensors (i.e. hot or chilled water heat flows, air flow rates, valve positions and temperatures in various locations of the HVAC system); assuming a simple law, for instance a linear behavior, between the actuator and the power consumption; assuming control of heat load directly; assuming a simple control strategy such as switching on or off the heating/cooling system. The following examples show that these methods are adopted as well in black-box or model-free approaches.

Black-box modelling Another approach is using purely statistical or data-driven techniques, with no knowledge needed regarding the mathematical structure of the system and the corresponding dynamics. For an in-depth view of data-driven techniques for estimation of building energy consumption, readers may refer to (Seyedzadeh et al. 2018). Machine Learning (ML) and statistical methods are used to predict energy demand of a building and other related variables with different objectives: long and short term forecasting (hourly, daily, monthly, etc.) without controllable inputs, optimisation of structural parameters (thermal insulation, glazing properties, etc.), model-based control, etc.

(a) Forecasting only models

It's worth mentioning that a considerable fraction of data-driven techniques focuses on forecasting, with no inherent control application possible. The objective in this case is to accurately predict short-term (hourly or daily) or long-term (monthly or yearly) energy consumption based on past data and sometimes weather or usage forecasts. Such models, trained with actual building energy consumption data, are valuable and have been adopted for measurement, verification or ongoing commissioning of building performance. For instance in (Dong et al. 2005), the monthly energy consumption is predicted using Support Vector Machine algorithms trained on three types of weather data: monthly mean outdoor dry-bulb temperature, relative humidity and global solar radiation.

Dealing with short-term predictions, in (Sha et al. 2019) authors investigated daily energy consumption predictions by using three features: degree-day (based on outdoor temperature and the energy profile pattern of the considered building), day type (1–8, where 8 represents holidays), and month type (1–12). For this task they tested the prediction accuracy of three popular ML algorithms were: multi-variable linear regression, support vector machine, and artificial neural network. The authors in (Fan et al. 2014) focused on the next day energy consumption and peak power demand predictions. They investigated a feature selection strategy and compared a set of popular algorithms, for instance multiple linear regression, autoregressive integrated moving average, support vector regression, random forests, multi-layer perceptron, etc.

A long short-term memory neural network model in (Sendra-Arranz et al. 2020) forecasts the next day multi-step power consumption using relevant data from previous time-steps. Neural network models were also used in (Karatasou et al. 2006), where authors discuss predictive performance of the models for single-step and multi-step forecasting of energy consumption when using a set of independent inputs (environmental or calendar) or when past values of the load are considered. In (Guo et al. 2014) a statistical method based on two separate time-indexed autoregressive models with exogenous inputs is investigated for hourly cooling demand predictions. A first model is used for predictions of 1 to 6 hours ahead, while the second one is used for estimating 7 to 24 hours ahead. The model uses two exogenous input variables - outdoor temperature and relative humidity. Time lags for the autoregressive inputs and exogenous variables were chosen based on literature and analysis of the cooling load.

Black-box modelling can also be applied for commissioning or in the planning phase. For example, optimizing design parameters based on ML energy consumption predictions (Solmaz 2020).

The aforementioned approaches are not explicitly control oriented. They can be efficiently used for predicting loads under normal conditions, where recent information or daily repeating patterns are sufficient for a good estimation. Given the structure of the models and the importance of previous steps measurements, the accuracy of the predictions in dynamical conditions should be carefully analysed. The propagation of the forecasting errors through a prediction horizon is not usually evaluated, although it could condition the use of such methods in a control framework.

(b) Control-oriented models

Proceeding to explicit control-oriented models, the examples hereafter give an idea of the variety of specific problems related to control that were tackled in the literature and the corresponding appropriate data-driven techniques.

Some classical system identification approaches to derive models for control in buildings have a long history. In (Mustafaraj et al. 2010) authors investigated linear parametric models such as box-jenkins, autoregressive with external inputs, autoregressive moving average with external inputs, and output error models, that could be used for optimizing the HVAC control. They tested the ability of the models to output up to two hours ahead predictions of room temperature and relative humidity by considering outside temperature and relative humidity, supply air flow-rate, temperature and relative humidity, chilled and hot water temperature. In (Jiménez et al. 2008) an overview of some nonlinear and linear models, appropriate for modelling the thermal dynamics in buildings, is given. The difference between the two families is demonstrated in a case study, where the example belonging to the nonlinear type falls into the category of grey-box modelling. Regarding the model choice, parametric models or linear regressions limitations have been pointed out in the case of high nonlinearities that characterises buildings and HVAC systems (Y. Chen et al. 2022), (Berthou et al. 2012).

Nonlinear black-box models have yielded promising results. Given the diversity of buildings, HVAC systems and targeted applications, the feature selection of such models differ accordingly. In (Kim 2020) authors adopted neural networks techniques to model the thermal dynamics of the building. More specifically neural networks were trained to predict the indoor temperature of each zone based on controllable (power consumption), feedback (indoor temperature), and environmental input variables (can include exterior temperature, solar irradiance, wind speed, occupancy or thermal load). Neural networks were used to schedule the optimal power inputs of the HVAC system during the next 24 hours under variable energy prices. The approach has two main limitations. As mentioned for the grey-box modelling, energy measurements for building sub-systems or zones are not commonly available. Moreover the control of the power consumption per zone is a non realistic assumption.

In (Manjarres et al. 2017) authors investigate optimal control in a building by switching on/off the HVAC system and scheduling the operation of the mechanical ventilation based on the weather forecast and estimated indoor temperature. A Random Forest (RF) model is used to estimate internal temperature based on historical outdoor and indoor temperature and humidity, occupancy data and weather forecasts. The model is used in an iterative manner to predict if the comfort would be satisfied or not with respect to the computed optimal control. The method can be applied successfully for energy saving strategies such as optimal stop or optimal start. However for DR applications, additional opportunities provided by preheating or precooling cannot be exploited.

A common trend is identification of models based on a relatively large set of sensor data and assuming access to an also large or system specific set of actuators. In (G. Xu 2012) for instance, a multi layer perceptron was used to model energy consumption and temperature based on a set of controllable and uncontrollable inputs and their respective lagged or forecasted values. The model was used in a continuous optimization framework searching for the optimal supply air temperature and the supply air duct static pressure

setpoints. In parallel, in (Tang 2010), different ML algorithms were tested to model the indoor temperature, indoor humidity, indoor CO₂ concentration and the energy consumption of an air handling unit, based on 10 relevant inputs such as air handling unit supply air temperature setpoint, supply air duct static pressure setpoint, chilled water coil valve position, etc. Furthermore the modelling and optimization of the energy consumption of specific HVAC components (chillers, fans, pumps, variable air volume box) was studied using six different data-mining algorithms.

The major drawback in this case is that in real life applications, energy measurements for building sub-systems are not commonly available and installing sensors, actuators, and other expensive equipment might be prohibitive. This also implies that the model is inherently building-specific and that its structure in terms on inputs and outputs should be tailored for each considered building.

Integrating instantaneous feedback, such as in (C. Zhang et al. 2019), has also been frequently considered. In this study, authors seek optimal control of the zone supply fan air mass flow rate and temperature setpoints, to maintain temperature at desired level while reducing the power consumption. The proposed method relies on a long short-term memory neural network to estimate zone temperatures and power demand based on the previous observations and control actions. The model is used in a MPC framework to determine optimal controls. In a second step, a neural network is trained on observation and action pairs returned from the MPC, in order to imitate the MPC optimal output and speed up the control decision making. The approach is legitimate in continuous optimal control. However, if the application implies optimal scheduling for a longer horizon, the result will be conditioned by the model's predictive ability when used iteratively on a longer horizon.

To overcome some limitations mentioned above, regarding the modelling efforts needed for detailed mathematical models or realistic implementation of models based on large sensor and actuator sets, global setpoint change rules have been considered. For instance authors in (Nghiem et al. 2017) use an autoregressive gaussian process to model the building's power demand in response to DR control signal $u_t \in [-1, 1]$. Each building implements its own strategy to increase ($u_t > 0$) or decrease ($u_t < 0$) the power consumption proportionally to the received DR control signal u (for instance adjusting the global temperature setpoints by $-2u$ (°C) implies raising or lowering the setpoints by up to 2°C compared to the nominal values). The implemented strategy is supposed to satisfy the occupant comfort, while reacting to u_t . The inputs of the model are the current hour-of-day, the current and past control inputs (DR control signal u_t), the current and past outside air temperature and humidity as well as past power demands. The model was tested successfully for multistep-ahead predictions, by propagating the expected power consumption values. The data was obtained by simulating a month building operation with random DR signals $u_t \in [-1, 1]$ during weekdays. The model was trained on three weeks of data, and one week was used for validation. Finally the model was used in a MPC demand tracking control problem. The objective is to track, during several hours, a power consumption profile provided by the DR program, which is below the baseline.

A battery is also considered, to improve the tracking quality by limiting the model uncertainty effects.

Adjusting power demand to DR by setpoint control was also investigated in (Behl et al. 2016). Different data-driven models were investigated for predicting the power consumption of the building based on data such as: weather (outdoor temperature, humidity etc.), schedules (setpoints of chilled water supply temperatures, supply air temperature and zone air temperature, time of day, etc.), state data (chilled and hot water supply temperatures, indoor air temperature, supply air temperature, lighting levels). Several methods based on regression trees are proposed for different objectives: DR baseline prediction (in the absence of a demand response event), DR strategy evaluation and DR optimal control. The baseline power consumption model, trained on weather and time related variables, was evaluated on real hourly data. For DR evaluation, the model features included also schedules and zone temperatures. More specifically auto-regressive trees were used to predict power consumption for a one-hour horizon (duration of the DR event). Zone temperature were estimated using additional auto-regressive trees. The models were used to estimate the best DR strategy and the associated power consumption, from a set of rule based strategies. Each strategy has defined control schedules for chilled water set-point, zone temperature set-point, and lighting set-point during the DR event. In a similar simulation experiment, the models were used to determine, in a closed-loop fashion, the control actions that maximize the power curtailment while maintaining occupant comfort.

The last two methods are scalable and can be seamlessly implemented in a Building Management System (BMS). The propagated errors induced by the use of auto-regressive models are limited, although in the second reference only a short term horizon of 1 hour was studied. Both methods implement power shedding strategies which consist, for instance, in increasing the indoor cooling temperature setpoints during a DR event. Such strategies however might not be accepted by building managers. Indeed, for energy efficiency reasons, buildings are usually operated at the comfort limits, and little or no deviation beyond is accepted. In such cases, instead of load curtailment, a load shifting strategy by precooling the building can be implemented. This strategy has the advantage of respecting predefined comfort bounds.

2.3.2 Model-free optimal control

RL has emerged as a different approach that tries to overcome some of the model-based method's pain points (modelling and control complexity, scalability), by learning optimal control from interacting with the system directly. Although methods based on tuning control settings in order to improve performance or implementing expert rules can be also considered as model-free, their capabilities are quickly limited in complex situations and systems. RL paradigm implies supervisory control without building a model of the system and has been considered as a promising method for model-free optimal control design. The method relies on interacting with the system and learning the optimal control strategy to be adopted. For applications considered here, the learning is achieved by continuously evaluating the results of

control actions in terms of state (observations/measurements of the environment or the system, such as occupancy, indoor or outdoor temperature and humidity) and reward (incentive reflecting the optimization goal, such as a penalty for the energy consumption of the HVAC system or a penalty for the occupant thermal discomfort). A non negligible advantage of RL is being able to adapt seamlessly to environmental, structural or system changes in a building, where model-based or expert-based control needs retraining, remodelling or rebuilding the knowledge database.

For an overview of the subject and examples related to DR readers may refer to (Vázquez-Canteli et al. 2019). The examples hereafter are focused on optimal control for energy efficiency and some DR applications.

RL techniques have been often investigated in use-cases for energy efficiency optimization by on/off operations. In (Barrett et al. 2015) authors propose a method that learns to optimally switch on or off the heating/cooling. To ensure occupants comfort, Bayesian inference is used to predict occupancy. The method was tested on a first-order equivalent thermal building model with no additional stochastic heat gains other than those from the heating/cooling system. Compared to an expert rule control that defines an optimal start and stop of the system, the RL technique allowed an improvement of 10% in costs. A similar approach was investigated in (Ruelens, Iacovella, et al. 2015) for finding optimal start and stop schedules of heat pumps in residential houses that respect comfort ranges during occupancy.

The next examples have more complex or broader action spaces. In (Henze et al. 2011) authors investigated the performance of RL techniques for the operation of electrical thermal storage systems in buildings. The RL controller was trained to optimally adjust the charging rate of the thermal storage system under variable cost of electricity.

Also specific to DR applications, authors in (Ruelens, Claessens, et al. 2017) proposed a batch RL method to construct control policies for two types of thermostatically controlled loads - heat-pumps and electric water heaters. The method is applied first in a dynamic pricing DR model, where the agent learns an optimal close-loop control policy that minimizes the electricity cost. In a second DR case, the RL agent constructs and tracks an optimal day-ahead scheduling consumption program.

The authors in (Y. Wang et al. 2017) investigated advanced RL techniques for optimization of energy consumption by controlling room cooling setpoints. The set of state observations gathered at every control time-step are outdoor temperature, indoor temperature, previous time-step cooling power load, solar irradiance. The control action is a discrete room cooling temperature setpoint, the cost function is composed of a weighted sum of two objectives: thermal comfort and energy demand. A similar objective has been investigated in (Gao et al. 2019), where a deep RL technique is used to dynamically control the temperature and humidity setpoints while seeking energy efficiency and ensuring occupants comfort. The set of indoor and outdoor temperature and humidity is considered as state observations, while the actions are temperature and humidity setpoints .

A substantial number of methods are evaluated on simple building architectures (single or very few zones) or simplified physical models (for example a second order model of the indoor temperature dynamics) with little complex dynamics or constraints. However, the bottleneck of model-free RL approaches is the prohibitive training process duration, especially when facing complex situations, systems or non-stationary environments. For instance in (Wei et al. 2017) the training is performed on 100 months data (more than 8 years). (Mocanu et al. 2019) achieved a convergence of the optimal controller after 3500 episodes of on-line training, where every episode represents an average of 20 random days of the database (more than 9 years). The RL agent in (Azuatalam et al. 2020) needs 60000 hours of training for convergence (more than 6 years).

To reduce the training time, techniques based on a physical or data-driven model first for training offline the RL controller have been proposed. Although this is a conceivable solution, one may argue that this step cuts off RL methods of their main advantage. The examples hereafter fall into this category.

Authors in (Schubnel et al. 2021) proposed the use of a reduced model, such as a neural network architecture, which is first trained on data generated with a coarse model of the system and then shortly retrained on historical data via transfer learning techniques. The obtained reduced model is used in a RL framework to train an agent offline to control blind opening in each zone and heating temperature supply, cooling temperature supply, ventilation temperature supply of the central HVAC system. In (Nikovski et al. 2013) a markov decision process framework is investigated for indoor temperature setpoints optimization using a third-order thermal equivalent model and an estimated coefficient of performance of the HVAC system. Authors in (Urieli et al. 2013) have used a linear regression based on 14 relevant features in order to model the effects of the actions, in other words - the transition functions. Based on the learned model, the RL agent uses a lookahead policy to choose the optimal actions. The framework was applied for finding optimal running schedules of heat pumps and auxiliary heating systems in residential houses, that respect comfort ranges during occupancy. The training in (Ding et al. 2019) is executed on a realistic physical model using E+ simulation software. The whole-building modelling approach was adopted as well in (Wei et al. 2017), (Azuatalam et al. 2020) and (Z. Zhang et al. 2019).

2.4 Conclusion

To summarize, methods differ substantially depending on the fixed objective, the comfort constraints, the available data from sensors or technical documentation and the accessible actuators. Given all possible configurations of HVAC systems, a rigorous comparison between all the modelling and optimization approaches in the state of the art is very tricky if not impossible. This is especially true for solutions that are tailored for specific systems. Some key features in each method allow to make the choice in a particular context.

The vast majority of RL methods tackle the optimal control task in a closed-loop fashion, i.e the control policy depends on the current measurement of the state and direct feedback from users. Several DR programs however expect in advance an optimal control plan of the HVAC system and the corresponding estimation of the power consumption. If such estimations are possible for our problem, it's still unclear what is the expected precision. Together with the prohibitive training process duration, this seems to indicate that RL might not be adapted for all applications.

Black-box models have been widely investigated thanks to their ability to capture nonlinearities, that often characterize systems such as buildings. Their precision is conditioned by the set of inputs, that should be carefully selected. For example, it is generally known that power consumption is autocorrelated and that it depends, among other factors, on the indoor temperature. Therefore, single-step models that depend on past time-step measurements might be very precise but unusable in some frameworks where long term open-loop predictions are required, i.e. when continuous measurements of power and indoor temperatures are not available.

Grey-box modelling is based on physical insights of building's thermal dynamics and accurate model of the HVAC equipment. Therefore, a first model describes the evolution of the indoor temperature, depending on different heat transfers that occur in a certain zone. A second model is meant to describe the consumed power of the HVAC system depending on a set of external conditions and control variables. Usually, the metered consumed power of the HVAC systems is not equivalent to the input of energy in the zone, therefore the weakness of the approach lies at the welding spot of the two necessary models. In some cases, a valid simple relation can be found between the power consumption and the heat input in a building. In real life situations this is often not feasible due to the complexity of the HVAC system. This induces the need of modelling the functioning of the specific equipment and bridging the gaps between the modelled control, power consumption and heat transferred to the indoor space. In most cases, the required set of sensors and actuators needed for a precise global model is large and tailored for a specific system.

Physical modelling requires huge amounts of real system data to be collected: building layout, physical properties of the construction materials, HVAC components and functioning. The modelling process itself requires advanced expert knowledge. In almost all cases the exact system components cannot be faithfully modeled, which requires tailored solutions to represent the desired phenomena or components. This normally leads to a tricky calibration phase. Finally, depending on the building size and the sophistication of the model, the simulation time can be substantial. For instance the digital twin of the new energy efficient building of Schneider Electric - IntenCity - requires 15 hours for a whole year simulation. Using such a model to test control scenarios for one-day horizon, as in the proposed solution, would require more than 3 hours of computation time.

Problem setting and adopted approach synthesis

3.1 Investigated use-case and adopted approach

This research focuses on the use-case of office buildings participating in a day-ahead DR program that values energy flexibility. Flexibility can be generated by different manners, however, the exact impact of the available options is often unknown. The task is difficult given the large diversity of buildings in terms of thermal characteristics, HVAC configurations and components, and available or accessible actuators and sensors. Moreover, the global behaviour of the building and the interaction between different HVAC subsystems are complex and highly nonlinear functions, explaining why the task of optimal control of energy consumption in buildings is a difficult one.

We have seen in the previous chapter some classical approaches to deal with operational optimization of HVAC systems. The cornerstone of model-based techniques is the model itself, and one usually tries to balance the complexity and the precision. RL methods that try to overcome modelling difficulties are still subject to training duration capabilities. Accordingly, this motivates a black-box modelling approach, adapted to a set of constraints, challenges and specific use-case expectations.

The first source of complexity that shaped the proposed methodology is linked to the **comfort of the occupants in an office building**. Facility managers are often reluctant to adopt control strategies that might negatively impact the comfort. A convenient way of treating this constraint is by ensuring that eventual control operations will cause indoor temperature to vary in a certain comfort range, predefined by the facility operator. Incidentally, the comfort range defines also the setpoints for heating or cooling activation: $[T_{min}^{sp}, T_{max}^{sp}]$. T_{min}^{sp} being the setpoint below which the heating is activated and T_{max}^{sp} the setpoint above which the cooling is activated.

A second source of complexity lies in the need for a **sufficiently reliable estimation of the flexibility potential**. Indeed, some DR programs require the power consumption schedule or an estimation of the shedded energy. This is also of interest if the building is integrated in a microgrid or aggregated with other grid assets, for which a global optimized schedule is computed in advance. Consequently, if power consumption can be lowered for a certain duration, with respect to a baseline, then it is essential to quantify the result.

Other challenges were set based on Schneider Electric’s past experiences with operational optimization of HVAC systems: prioritize **scalable methods** that are easy to duplicate on different buildings, **that can be integrated in a BMS** in a simple manner with minor or no equipment investments and that can adapt easily to changes that might occur in buildings (HVAC system or structural modifications).

An intuitive approach that could satisfy all the above constraints is controlling indoor temperature setpoints¹. Firstly, setpoint modification inside comfort bounds guarantees comfort at all times. Secondly, the choice of the inputs and outputs for the data-driven model is such that the model is agnostic to the building specific set of actuators and sensors as well as to the degree of accessibility of the actuators internal tuning and characteristics, and therefore potentially applicable to different buildings without much effort besides data collection for training. In order to design such a scalable method, only data from commonly available sensors and meters are to be considered: temperature setpoints, indoor temperature, total power consumption, and exogenous data such as outdoor temperature, solar irradiance and time related information. Moreover, the integration in a BMS is non-intrusive and can be achieved by establishing a simple communication channel. The BMS delivers continuous information such as total power consumption, indoor and outdoor temperatures, and other relevant available measures. The designed model-based optimization module, sends daily, or more frequently if needed, the optimized indoor temperature setpoint schedule intended to overwrite the default one.

The method investigated in this work is novel in regard to the following features and minimalist assumptions:

- the modelling is purely data-based and system agnostic;
- data comes from common available sensors and meters (indoor temperature, indoor temperature setpoints, weather conditions, total HVAC power consumption);
- the indoor setpoint control can be implemented in a straightforward way;
- comfort is ensured at all times;
- the method is suitable for explicit DR programs, where an evaluation of control actions is needed for a given horizon (24 hours for instance).

Given the comfort constraints, flexibility activation by load shifting will be studied in this work. In practice this is achieved by modifying the baseline indoor temperature setpoints to preheat/precool the indoor space before a DR event. The effect is activation of the inherent thermal storage capacity of the building. When the setpoints are reset to their baseline values, the thermal inertia of the building, or in other words discharging of the stored thermal energy, will induce a reduction of power consumption.

¹In a general case, a control strategy for each individual zone in the building can be applied. In this work, however, the control concerns global setpoints (as an average of the individual zone setpoints for instance). The proposed adjustments of the setpoints will consist of lowering or raising the setpoints of all zones (or a predefined set of zones) by the same amount. Accordingly, a mean indoor temperature and a total power consumption of the HVAC system are considered.

Simulation results in Fig. 3.1 illustrate a preheating example in winter. The global lower comfort bound (heating temperature setpoint T_{min}^{sp}) is set to 21°C. In order to react efficiently to a DR event that starts at 7PM, without impacting occupants comfort, the temperature setpoints of the building are increased to 25°C during 1 hour before the DR event. This allows passive accumulation of heat in the air and the building structure, that is released during a DR event, by resetting the heating setpoint back to its nominal value.




Figure 3.1: An example of Demand Response, given a comfort range $[T_{sp}^{min}, T_{sp}^{max}]$. The global temperature setpoints are increased to 25°C before a DR event. This allows passive accumulation of heat in the air and the building structure, that is released at the beginning of the DR event, by resetting the setpoints back to their nominal value.

This type of operation can be employed in different DR programs: TOU tariffs, critical peak pricing, supply of ancillary services. However, the benefits of power consumption limitation must be balanced with the increase of the total power consumption occurring with the precooling/preheating, and any associated benefit can only be determined with respect to a baseline estimate and the DR remuneration associated to the reduction in the power use induced by the operation. In this work we assume that DR incentives materialize as remunerations for a reduced amount of energy at the specified interval of time. In practice the problem can be formulated in the exact same way for TOU electricity rates. Subsequently, an optimization framework has to be implemented in order to determine the optimal control based on the power reduction during the DR event, the increase of power during the preheating/precooling,

the baseline energy cost and the DR remuneration.

An abstract view of the framework is illustrated hereafter on a simplified building example in the heating season. The example building is composed of several zones and a global HVAC system with production and distribution elements as in Fig. 3.2. Sensors are available in each zone i for measuring the indoor temperature $T_{in,i}$, which is controlled by default temperature setpoints schedules $T_i^{sp\ baseline}$ sent by the BMS to the local controllers. The heating is produced by the central plant and distributed according to individual thermostat controllers. Energy meters are available for the global production and distribution systems, recording the associated power consumption $P_{production}$ and $P_{distribution}$.

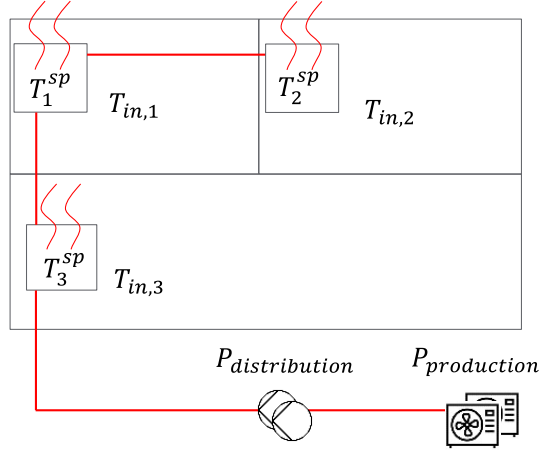


Figure 3.2: A simplified building example. The building is composed of several zones and a global HVAC system with production and distribution elements.

Fig. 3.3 illustrates the BMS-induced coupling. In this example a trained model estimates total HVAC power consumption based on features such as heating temperature setpoints and weather data (exterior temperature T_{ext} and solar irradiance Irr). To enable the learning of such a model with a reasonable amount of data, mean temperature setpoints over the zones are considered. Given a DR incentive, the optimization module uses power model predictions to perform a day-ahead optimization of the preheating, and outputs an indoor global temperature setpoint schedule T^{sp} . The BMS can adapt it depending on zone particular constraints, and communicate the resulting zone schedules T_i^{sp} to the local controller of each zone i .

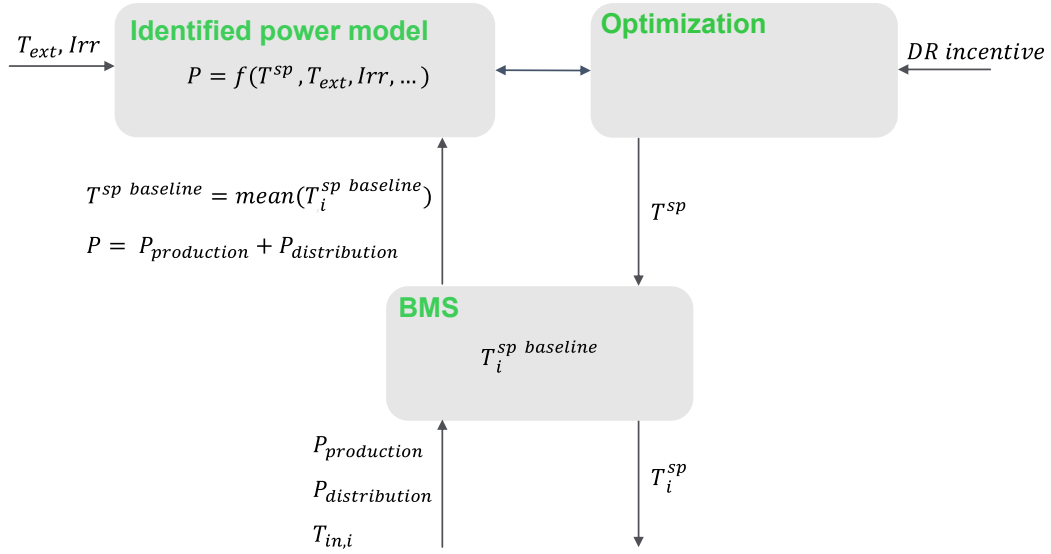


Figure 3.3: BMS coupling in the simplified building example. Data from sensors and meters are collected by the BMS. Through the BMS, the user defines the default temperature setpoints for each zone. Given a DR incentive, the integrated optimization brick computes and communicates a new global setpoint schedule that is applied to the individual zones.

For the sake of clarity, only the winter/heating season will be treated, although the method can be applied without loss of generality to the summer/cooling season. Learning different models for heating/cooling and summer/winter seasons is also justifiable by the fact that HVAC cooling and heating systems might be fundamentally different and display distinctive power consumption profiles.

For simplicity sake, only the mean average temperature and setpoints among zones (balanced by the surface area) are used for the results showed here. This is also motivated by the fact that handling only a mean temperature, especially if the number of zones is high, reduces the complexity of the model. Averaging the temperatures, however, leads inevitably to a loss of information. For instance at a particular time, mean indoor temperature being above the mean heating temperature setpoint does not imply that the heating system is off. Indeed, the temperature variation between zones can be substantial and some particular zones can still be heated while the temperature in others are varying inside the comfort temperature bounds, as illustrated in Fig. 3.4. Nonetheless, the investigated methods allow integration of zone specific data, without loss of generality.

As already mentioned, the only variable considered, that accounts for the power consumption of the HVAC system, is the total power consumption as an aggregation from all the components (ventilation, heating, cooling, etc.). This aggregation is particularly important to investigate given the variety of HVAC systems and the different ways of metering the power consumption in real buildings. Nonetheless, as for the indoor temperature, the method can

be adapted seamlessly if finer information is available.

The value of the solution is highlighted in this work with respect to the electrical grid. It is assumed therefore that such a solution is intended to be implemented in buildings equipped with HVAC systems that have a predominant electrical power consumption. Even though the diversity of HVAC systems can imply different energy carriers, the implementation of such a solution is not prejudiced. On the contrary, as long as a meaningful connection exists between global thermostat settings and the electrical power consumption, the approach is meant to deal with the complex and nonlinear relations between the energy input and the heat transfer to the conditioned space.




Figure 3.4: Mean temperature in the building (15 zones) compared to three distinct zone temperatures. The power consumption variation regarding the relation between global temperature setpoints and mean temperature is less obvious due to zone disparities.

It should be noted that weather forecast uncertainties have not been considered in this thesis. Since the assumed DR program is based on day-ahead notifications, experiments are still needed to test the sensibility of the resulting models with respect to weather forecast uncertainty.

Based on the concepts from above, the next chapters will detail the final approach used for power consumption estimations (features, regression model) and the optimization framework (DR mechanism, load shifting strategy). Results obtained with a set of other modelling approaches, that shaped the final one, are also discussed. To test the different methods, several datasets have been used throughout this study, corresponding to simulated buildings of different complexities. The next section describes these buildings/datasets.

3.2 Case studies

The method developed in this work was tested in several ways. The modelling precision was tested on simulation data from different publicly available sources and internally developed simulation buildings. Economical evaluation of the energy flexibility, in a DR framework, was estimated using a building simulated in IDA ICE simulation software. In all the case studies, a perfect weather forecast has been considered.

3.2.1 An IDA ICE simulated building

A physical model of a typical medium-size office building constitutes the reference for this research work. The simulated building was used in the complete procedure of modelling and DR optimization.

Note that IDA ICE is a high fidelity simulating environment - IDA Indoor Climate and Energy (IDA ICE) simulation software provided by EQUA Simulation AB. IDA ICE is a building design tool based on an extensible equation-based simulation framework, in which way it resembles to the well known Modelica-based simulation frameworks. IDA ICE offers detailed and dynamic multi-zone simulation capabilities, with true modelling of control loops, advanced HVAC modelling. It also offers capabilities in terms of model integration for example co-simulation and functional mock-up interface integration, which we have used in this study. It has been successfully tested and adopted by industry, as well as Schneider Electric, in a wide range of projects for calibration or digital-twin applications (for example the EU projects TRIBUTE Béguery, Petit-Pierre, et al. 2017 and HOMES Béguery, Lamoudi, et al. 2011). Readers may refer to Sahlin et al. 2004 for more details about the IDA ICE characteristics.

The designed simulation building model is 3-level building with 11 offices, 2 meeting rooms and 2 open spaces. The simulated climate is based on Chicago historical weather data.




Figure 3.5: Test building representation. The building represents a medium-size office building with 15 zones.

The HVAC system includes two central subsystems: a central air handling unit and the primary heating and cooling system, represented by a boiler and a chiller. The boiler supplies, with a constant performance, hot water to the zones and to the central air handling unit. The supply hot water temperature setpoint is a function of outside air temperature. Therefore, the total available heating load is varying with the outside temperature. The chiller supplies cold water to the zones and to the central air handling unit, depending on a variable coefficient of performance.

The central air handling unit is designed to supply fresh air to the zones. The supply air is conditioned using heating and cooling coils with water mass flow control. The heating and cooling water is supplied by the boiler and the chiller. An air-to-air heat exchanger allows to recover heat from exhaust air and to transfer it to the supply air.

Fan coils in each zone ensure thermal conditioning based on waterborne heating and cooling. The nominal power in each zone is specified for a set of design conditions. Therefore the effective heating or cooling power will vary depending on the indoor temperature and the difference between inlet and outlet water temperature.

The simulation tool allows faithful modelling of control loops and integrates complex phenomena such as solar radiation to different surfaces depending on orientation, internal and external shadings, thermal bridges, various heat exchanges between construction elements, internal gains and masses. The level of complexity of the building represents a real challenge, and hence a good validating framework for building agnostic control design as the one claimed in this contribution.

Note that the same temperature setpoints schedule has been applied for all zones. Similarly, occupant presence schedules (described by a working day $[0, 1]$ profile) are the same for all zones and working days, although the assigned number of occupants differ among the zones.

Schematic views of the building are presented in Appendix A.

3.2.2 A Synthetic Building Operation Dataset

In order to test the relevance of the power model's structure, a publicly available dataset based on the United States Department of Energy (U.S. DOE) detailed medium-sized reference office building was used. The simulated building represents an office with different types of zones spread on three floors. The total floor area is equal to 4890 m². Each floor is equipped with an air handling unit and each zone has a variable volume air terminal unit with electric reheat coils. Annual simulations are available, obtained using 30 years of historical weather data in three different regions including Miami, San Francisco, and Chicago. Data contains system operating conditions: HVAC, various electric loads, occupancy, lighting, weather data and total energy consumption at 10-minute intervals. The dataset is described in detail in (H. Li et al. 2021).

3.2.3 A Digital Twin of a large office building

Another tested dataset was available on the Schneider Electric Exchange platform. The building has a total floor area of 10250 m² and is located in Grenoble. The dataset includes time-series for over 2700 sensors, recording occupancy, energy, comfort and environment variables, system setpoints.

3.2.4 An Equivalent Thermal Parameter simulated building

In order to confirm some intuitions regarding the performance of data-driven models and to build up the complexity of the final solution, a simplified building, represented by an Equivalent Thermal Parameter (ETP) simulation model, was used. Since the thermal dynamics are fully understood, the performance of state-space models for instance, investigated in the next chapter, can be easily explained.

The simplified building consists in two heat flow equations that describe the thermal dynamics. The first one describes a given zone indoor temperature evolution, submitted to different heat transfers: direct heat exchanges with the outdoor air, heat transfers through external walls, heat exchanges with adjacent zones, solar irradiance, equipment and occupants and HVAC system heat gains. The indoor temperature dynamics are described mathematically through a first differential equation that accounts for the above mentioned heat transfers. Accordingly, the mathematical equation can be expressed for a given zone j as follows:

$$\frac{1}{C_{in,j}}\dot{T}_{in,j} = \frac{1}{R_{in\ w,j}}(T_{w,j} - T_{in,j}) + \frac{1}{R_{in\ ext,j}}(T_{ext} - T_{in,j}) + \sum_{i \neq j} \frac{1}{R_{adj\ i,j}}(T_{in,i} - T_{in,j}) + \Phi_{irr,j} + \Phi_{occ,j} + \Phi_{h,j} \quad (3.1a)$$

Where $T_{in,j}$ is the indoor temperature, $T_{w,j}$ is the temperature associated with the external walls, T_{ext} is the outdoor temperature, $T_{in,i}$ is the indoor temperature of the adjacent zone i ; $\Phi_{irr,j}$, $\Phi_{occ,j}$, $\Phi_{h,j}$ are heat gains associated to solar irradiance, occupancy and HVAC heating; $C_{in,j}$ represents the thermal mass parameter of the air; $R_{in\ w,j}$, $R_{in\ ext,j}$ and $R_{adj\ i,j}$ are lumped thermal resistances between the interior and the exterior walls, between the interior and the exterior and between the interior and the adjacent zones respectively.

A second differential equation describes exterior wall temperature dynamics, submitted to heat exchanges with the outdoor atmosphere and the indoor air. This allows to simulate a more realistic indoor temperature behavior linked to the thermal inertia of the envelope.

$$\frac{1}{C_{w,j}}\dot{T}_{w,j} = \frac{1}{R_{in\ w,j}}(T_{in,j} - T_{w,j}) + \frac{1}{R_{w\ ext,j}}(T_{ext} - T_{w,j}) \quad (3.1b)$$

Where $C_{w,j}$ represents the lumped thermal mass parameter of the external wall, $R_{in\ w,j}$ and $R_{w\ ext,j}$ are lumped thermal resistances between the external wall and the interior air and

between the external wall and the exterior air respectively. All the thermal parameters have been tuned in order to represent realistic thermal dynamics.

The HVAC heat gain in a given zone is controlled by a Proportional Integral (PI) controller:

$$\Phi_{h,j} = \max(0, \min(\overline{\Phi_{h,j}}, \int K_I(T_j^{sp} - T_{in,j})dt + K_P(T_j^{sp} - T_{in,j}))) \quad (3.1c)$$

Where $\overline{\Phi_{h,j}}$ is the HVAC maximum heating power, K_I and K_P are PI controller parameters and T_j^{sp} is the temperature heating setpoint. Since in this particular case there is no additional equipment, the total power consumption P of the HVAC system is equivalent to the total heating gains multiplied by the Coefficient of Performance (COP): $COP * \sum_j \Phi_{h,j}$.

Additional features were implemented in order to enhance the complexity of the building. A multi-zone configuration was enabled, with varying parameters among zones. The solar heat gains depend therefore on the configuration (orientation, glazing area, inter-zone connections). The simulated solar heat gains are also impacted by a varying cloud cover. Additionally, a varying occupancy profile was implemented.

3.3 Adopted approach synthesis

The adopted approach for enabling the energy flexibility in a building is based on regression predictive modelling of the HVAC power consumption. This section offers a summary of the modelling, optimization framework and main results. The case study is a medium office building modeled in IDA ICE simulation software.

Learning and testing dataset Two distinct meteorological years have been used for training and testing purposes, corresponding to Chicago location. Experiments showed the interest of adopting a seasonal approach, therefore this work treats the winter/heating season only: November 2001 - March 2002 for training and November 2002 - March 2003 for testing. The data sampling time is 15 minutes.

The shedded HVAC load can only be determined with respect to a baseline estimate. Simulations showed that the preheating operations can have a lasting effect on power consumption profile, therefore the experiment for model learning should be carefully designed. To enable an unbiased baseline modelling and DR outcome estimation, the learning dataset has been constructed by simulating two preheating operations per day, every other week. Random preheating operations have been simulated in terms of setpoint increase with respect to the baseline values, duration and time of the day. For a DR purpose, two distinct models are learned. The main one, based on weeks with simulated preheating operations, is used for predicting the HVAC load during preheating and shedding. A secondary one, based on weeks where the default operating temperature setpoints are applied, is used for predicting the baseline energy consumption.

Model features and algorithm To capture the dynamical behaviour of the building the following features have been considered (cf. Section 5.1):

- Weather data: outdoor temperature and solar irradiance;
- Time related data: time of the day, day of the week;
- Past indoor temperature setpoints and weather conditions.

The set of features is meant to contain the necessary information for recovering the current thermal state of the building. It should be noted here that depending on the thermal characteristics and behaviour of a given building, lag values for the past inputs should be optimized accordingly.

For this particular problem, the extreme gradient boosting algorithm XGBoost (T. Chen et al. 2016) turned out very convenient, in terms of tuning and accuracy.

Model Validation The testing dataset consists in two random preheating operations during working days. The accuracy of the main model are $R^2=0.95$ and Weighted Absolute Percentage Error (WAPE)=0.10. Fig. 5.3 and 5.4 in Section 5.1 illustrate the prediction errors.

DR market and optimization framework HVAC load flexibility is assumed to be activated in the context of a DR program. The DR program consists in day-ahead incentives for reducing the HVAC load during specific intervals of time (events) next day. Two fixed DR notifications each day are assumed: from 7 to 10AM and from 5 to 8PM. For details reader can refer to Section 6.1.

Participation in a given DR event is decided by estimating the balance between the preheating cost and the remuneration for the shedded energy. For a given DR event, the load profile (and consequently the preheating/shedded energy) depends on the preheating settings: preheating duration $\Delta t^{preheat}$, start of shedding t_{start}^{shed} , temperature setpoint increase with respect to the default/baseline profile ΔT_{min}^{sp} .

Let $x = (\Delta T_{min}^{sp}, \Delta t^{preheat}, t_{start}^{shed})$ be the set of decision variables for a given DR event. The shedding duration $\theta = \Delta t^{shed}$ is fixed to one hour, so x and θ define the preheating settings for that event.

Consequently the maximum predicted net benefit \hat{B} ² for a given DR event is defined by the difference between the DR participation reward and the cost of the preheating operation:

$$\hat{B} = \max_x (r |\Delta \hat{E}^{shed}(x, \theta)| - \Delta \hat{E}^{preheat}(x, \theta)) \quad (3.2)$$

Where r is the ratio between the reward for DR participation and the energy cost during the preheating time-interval. In the DR framework, r is supposed to represent the intrinsic

²The estimate of a variable x is noted \hat{x}

"value" of the shedded energy, in a period of grid stress, with respect to the cost of the energy in a normal period. This is based on a simplistic assumption according to which the baseline energy price and the reward level are known and fixed by the grid operator/DR program. It follows also that the price of the energy is constant during preheating and similarly, during DR, the compensation rate for the shedded energy is also constant. Although the defined reward factor does not rely on existing DR programs, one may interpret the following results as compensation levels for which buildings may be motivated to adopt a proactive position and participate to grid services.

The real and estimated energy consumption difference, $\Delta\hat{E}^{(\sigma)} = \hat{E}^{(\sigma)} - \hat{E}^{(\sigma)baseline}$ and $\Delta E^{(\sigma)} = E^{(\sigma)} - \hat{E}^{(\sigma)baseline}$, where $\sigma \in \{\text{preheat}, \text{shed}\}$, for a preheating setting (x, θ) , are defined with respect to the baseline load profile, i.e. the load profile in absence of any DR participation. $\Delta\hat{E}^{(\sigma)}$ is predicted using the two regression predictive models detailed in the paragraphs above. The observed energy consumption difference $\Delta E^{(\sigma)}$ still needs estimation of the baseline energy consumption. The corresponding baseline model is used for this intent.

Participation to a DR event is triggered if the predicted maximum income is positive. Subsequently, the total income for the triggered set of DR events is computed a posteriori based on the effective power consumption profile.

Given the two daily DR notifications, the optimal preheating settings x are computed, for a fixed shedding duration of one hour. The income optimization is carried out by discretizing the decision variables and carrying an exhaustive search of the optimal set of preheating parameters x .

Estimation of DR income Three economical results have been deducted, which are detailed in Section 6.2. In real life situations, the DR program integration would be conditioned by the availability of precise enough forecasting models, trained on representative instances of DR. The results hereafter are related to different assumptions regarding the available amount of learning data.

In the first case, the forecasting models are trained on a whole winter season before integrating the DR mechanism. The total benefit for the next winter season, relative to the total baseline energy cost, is 0.2% and 7.9%, for reward factors $r = 2$ and $r = 5$ respectively (see Fig. 6.3).

A second investigation shows similar DR income if the main model is retrained every month during the DR program, by integrating new available data with preheating instances in the training dataset.

A third experiment considers the case where the DR mechanism is integrated faster. Only one month of data is used for training the main model, before integrating the DR program. The model is then retrained periodically by incorporating new data from DR actions. The experiment supposes however a baseline prediction tool which, in this case, is the same as in the first two cases above. Results show that the income loss is limited to 22% for $r = 2$, and 2.5% for $r = 5$ (see Fig. 6.5).

Additional results of DR income with respect to building thermal characteristics and DR program are given in Section 6.2 and Section 6.3.

Part II

Exploratory work on HVAC load and indoor temperature data-driven modelling

Data-driven modelling for power consumption and indoor temperature forecasting

Contents

4.1	Errors metrics	41
4.2	Linear state-space modelling	42
4.2.1	Linear state-space modelling of the thermal dynamics	43
4.2.2	Linear state-space modelling using indoor temperature setpoints	47
4.3	Autoregressive Random Forest regression predictive modelling	50
4.4	Random Forest regression predictive modelling	57
4.4.1	Structured data-driven modelling of indoor temperature	60
4.4.2	RF regression predictive modelling of power consumption only	64
4.5	Conclusion	67

In this chapter, several methods for indoor temperature and power consumption prediction will be investigated, starting with simple linear models and moving towards more complex regression models. These investigations have been very experimental in nature, starting from initial understanding and intuitions about the building behaviour, and then in reaction to the identified sources of model errors. With the objective of developing good-enough models using as less data as possible, the results presented in this chapter show not only the best results obtained but also the drawbacks of simpler, intermediate models that have been tested.

Given the considered DR mechanism, the optimal indoor temperature setpoint schedule should be chosen for the next day. The corresponding power consumption modelling accuracy is investigated from this angle. The proposed control of indoor temperature setpoints, inside defined comfort bounds, implies that comfort is guaranteed at all times. Nevertheless, predicting indoor temperature can be useful for estimating the power consumption, and will also be investigated in this chapter. Indeed, indoor temperature contains information about present and past thermal dynamic states of the building, which is confirmed by the results in the following sections.

To better identify potential modelling limitations, both simulation data issued with ETP and IDA ICE buildings (cf. section 3.2.4 and 3.2.1) will be used. First, results for state-space model estimation are presented. State of the art references suggest that, despite the nonlinearities characterizing a building, such a linear model can capture the main thermal dynamics. Linear models are especially tempting given their simplicity. Moreover, since the indoor temperature is controlled by a temperature setpoint schedule, our intuition was that linear models could perform well enough on average in closed loop. Based on ETP building data, the exact impact of different heat gains (HVAC, solar, ambient, occupancy) on indoor temperature has been studied. The limitation of such models, given partial disturbance/inputs knowledge, is pointed out, additionally to limitations regarding nonlinearities.

In a second step, more complex approaches are tested (see Fig. 4.1 below), based on supervised learning approaches.

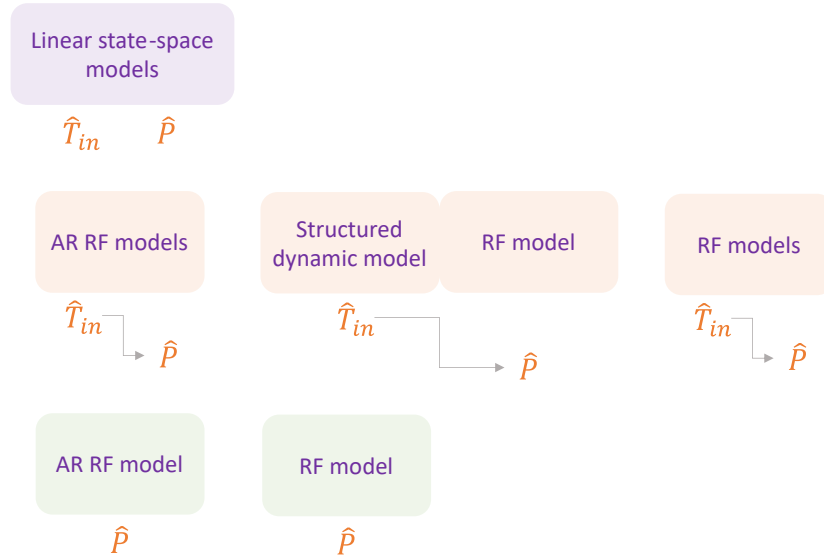


Figure 4.1: Explored modelling approaches and the modelling/forecasting targets. Six modelling approaches are discussed, designed for indoor temperature, total HVAC load, or both. Starting with state of the art linear models, then cascaded predictions of indoor temperature and HVAC load, followed by HVAC load prediction only.

More specifically, an autoregressive random forest approach (current output depends on past inputs and outputs) has been studied, where indoor temperature predictions are fed in to the power consumption model. Results indicate that the accuracy is low due to indoor temperature forecasting errors and error propagation through the prediction horizon, and suggest that a larger lag should be used for the autoregressive features. On this basis, a "classical" RF modelling approach has been tested (in contrast to the autoregressive approach, the current output does not depend on past inputs). Results show that the accuracy is improved compared to the autoregressive approach where previous step input and outputs are considered in the feature vector. In parallel, a structured nonlinear model has been designed in order to improve the indoor temperature predictions. The accuracy of the model is comparable

to a RF approach, however the method relies on finely tuned ramp temperature setpoint profiles for the preheating operations. Additional investigations are needed to determine how the indoor temperature evolves depending on the fixed ramp and what is the impact on the passive stored energy, compared to step preheating schedules such as presented in Part I. Finally, it seems that, due to the sensitivity of the power consumption forecasts to the errors of the indoor temperature predictions, the best strategy is to bypass the indoor temperature modelling. This is achieved by considering current and past inputs for the power consumption modelling, which most probably allows reconstitution of past states of the building, normally reflected by the indoor temperature.

4.1 Errors metrics

$$R^2 = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2}$$

The coefficient of determination R^2 provides a measure of goodness of fit. A higher value indicates a better accuracy. It compares the model errors to the variance of the target y and represents therefore the proportion of variance that can be explained by the features. Since the variance is dataset dependent, R^2 score may not be comparable across different datasets.

$$WAPE = \frac{\sum_{i=1}^n |y_i - \hat{y}_i|}{\sum_{i=1}^n |y_i|}$$

WAPE measures the overall deviation of forecasted values from observed values. A lower value indicates a better accuracy. It has the advantage of being scale-independent, and can be used to compare forecast performance across different data sets. As it is sensitive to relative errors, WAPE is great for measuring a model's performance when the dataset has low or intermittent values. However, WAPE needs a meaningful zero value for the target variable. So it may not be suited for predicting variables such as exterior temperatures, for instance.

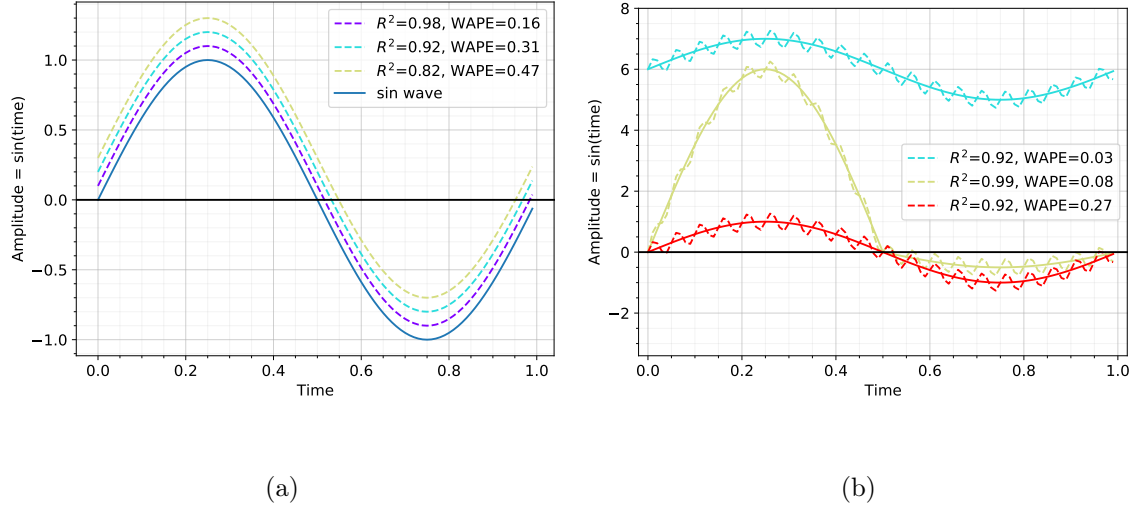


Figure 4.2: R^2 and WAPE errors. Subfigure (a) illustrates the case where the predictions are a translation of an original signal. Subfigure (b) illustrates the case where predictions include a noise sine signal (same amplitude and frequency) added to three different original signals.

4.2 Linear state-space modelling

As mentioned in Chapter 2, grey-box models have been widely used to represent the dynamics in buildings. The models usually describe the thermal dynamics, i.e. the dynamics of indoor temperature given the set of heat gains (cooling/heating, occupancy, irradiance, etc.) or inversely, the cooling/heating gains given indoor temperature dynamics and heat gains from auxiliary sources (namely occupancy, irradiance, outdoor). Since heating/cooling power is not usually directly controllable, additional steps are needed in order to describe the function between the applied control and the power consumption. As an alternative that enables the use of a model in a straightforward way in control strategies, introducing the temperature setpoints in the set of inputs can be considered, as it has been studied in (Berthou et al. 2012) and (Berthou et al. 2014). To test this idea and evaluate how temperature setpoints can be included in the model, an evaluation of state-space models was performed. By tackling different system complexities, the limitations of such an approach were highlighted.

In its discrete-time process form, a stochastic state-space model can be written as follows:

$$x_{t+1} = Ax_t + Bu_t + w_t \quad (4.1a)$$

$$y_t = Cx_t + Du_t + v_t \quad (4.1b)$$

Where x represents the state of the system (in contrast to grey-box modelling, x has not necessarily a physical meaning here); u is the system input; y is the vector of outputs (will be considered power and temperature); w_t and v_t are state and output measurement noises respectively; A, B, C and D are appropriate size matrices to be identified, t denotes the time-step

(15 minutes sampling time). Note that depending on the modelling approach, the terminology might differ. Therefore, in the following sections, the AI/ML terms are also employed, that refer to X as the set of inputs/features/predictive variables and to y as the output/target.

For the identification procedure, an open-source system identification package for Python (Armenise et al. 2018), freely available at GITHUB: <https://github.com/CPCLABUNIFI/SIPPY>, was used. More specifically the subspace identification algorithm N4SID was applied.

The identification was performed on two different buildings: an ETP 9-zone building and an IDA ICE building (cf. section 3.2.4 and 3.2.1). As mentioned in the problem statement, in the case of a multi-zone building, all zone specific variables are averaged and the total power used by the building is considered. To facilitate the analysis, the identification and the testing were performed on the same 3 months winter period. Note however that ETP and IDA ICE buildings have been simulated with different weather profiles.

4.2.1 Linear state-space modelling of the thermal dynamics




Figure 4.3: Linear state-space identification of the thermal dynamics model. Schematic representation of the approach, input variables and outputs.

In a first step, the ability of linear models to capture the thermal dynamics in a building was tested. This follows the reasoning behind parametric modelling found in literature (see for instance (Mustafaraj et al. 2010), (Braun et al. 2002), (Kolokotsa et al. 2009)) although in this case no particular structure for the state-space model matrices, i.e. system parameters, is assumed. The approach falls therefore in the category of black box modelling.

ETP building case. One advantage of considering a parametric simulation building is allowing an easy analysis of how unknown disturbances affect the thermal dynamics model. The results hereafter show the modelling precision in two different cases: when all inputs affecting the system are known, i.e. exact heat gains from all sources, and when some unknown/partially known disturbances are considered. The latter is motivated by the fact that, in practice, the exact heat gains in each zone are complex to estimate; this is especially true for solar and occupancy heat gains. Therefore, in the second tested case, the occupancy heat gains are not included in the input set, and a global solar irradiance profile, as expected from a weather provider, is used instead of precise value of the solar heat gains affecting the zones.

a) Fully known disturbances

When all heat gains affecting the zones were considered, i.e. P , Φ_{irr} , heat transfers with ambient air and Φ_{occ} , the state-space modelling was able to capture precisely the system dynamics in the case of an ETP building. The identification of a heating power model in this case is performed with inputs $u = (T_{in}, T_{ext}, \Phi_{irr}, \Phi_{occ})^T$ and output $y = P$.

The identification of a temperature model is based on inputs $u = (P, T_{ext}, \Phi_{irr}, \Phi_{occ})^T$ and output $y = T_{in}$.

This is an expected result given that the identified model is equivalent to the linear system formed by eq. (3.1a) and (3.1b) in Section 3.2.4, describing the building thermal dynamics. It can be noted here that dealing with an average indoor temperature and a total heating power has negligible effects.

b) Unmeasured disturbances

Assuming unknown occupancy heat gains and solar irradiance from a weather provider, the identified model precision is reduced to $R^2=0.94$ and $WAPE=0.12$, compared to a perfect identification in the previous case where precise solar and occupancy heat gains were used. Identification results are illustrated in Fig. 4.4 below.

Identification of an indoor temperature model, with inputs $u = (P, T_{ext}, \Phi_{irr})^T$ and output $y = T_{in}$, results in a quite bad modelling precision: $R^2=0.12$ and $WAPE=0.04$. The fit between the estimated and observed values are illustrated in Fig. 4.5 below.

IDA ICE building case. For the IDA ICE building, only the case with partially known heat gains was considered. The error statistics are $R^2=0.83$ and $WAPE=0.25$ for the power model.

The temperature model identification results in $R^2=0.69$ and $WAPE=0.04$. Fitting results can be seen in Fig. 4.6 and Fig. 4.7.

In these first tests, the impact of unknown disturbances affecting the zones was evaluated. The results show that the system can be perfectly identified if all the heat gains affecting the thermal system are considered. However the impact of unknown disturbances can be considerable. This is especially critical for the indoor temperature model identification, as indicated by the corresponding mediocre accuracy. Note here that the identification results depend highly on the considered time interval. For example in the ETP building case, an identification with inputs $u = (P, T_{ext}, \Phi_{irr})^T$ and output $y = T_{in}$, performed on January-February instead of February-March leads to $R^2=0.64$ and $WAPE=0.04$ (instead of $R^2=0.12$ and $WAPE=0.04$).

Since HVAC power is not usually directly controllable, we move forward to the fixed objective of designing a convenient model that describes power consumption as a function of indoor temperature setpoints, among other available data. In the next subsection, the effects on adding the control actions to the input set is investigated.




Figure 4.4: State-space modelling given partially known heat gains affecting the zones, with $u = (T_{in}, T_{ext}, \Phi_{irr})^T$ and output $y = P$. ETP building.




Figure 4.5: State-space modelling given partially known heat gains affecting the zones, with $u = (P, T_{ext}, \Phi_{irr})^T$ and output $y = T_{in}$. ETP building




Figure 4.6: State-space modelling given partially known heat gains affecting the zones, with $u = (T_{in}, T_{ext}, \Phi_{irr})^T$ and output $y = P$. IDA ICE building.




Figure 4.7: State-space modelling given partially known heat gains affecting the zones, with $u = (P, T_{ext}, \Phi_{irr})^T$ and output $y = T_{in}$. IDA ICE building

4.2.2 Linear state-space modelling using indoor temperature setpoints




Figure 4.8: Linear state-space model identification using indoor temperature setpoints. Schematic representation of the approach, input variables and outputs.

Similarly as in (Berthou et al. 2012) and (Berthou et al. 2014), a state-space modelling was performed using temperature setpoints. The resulting model, in this case, could be integrated directly in an optimization framework.

ETP building case.

a) Fully known disturbances

When all heat gains were precisely considered, the state-space model with $u = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, \Phi_{occ})^T$ and output $y = P$ describes no longer a linear system, given the nonlinearity induced by the heating power limits. Identifications lead to $R^2=0.82$ and WAPE=0.23 for the heating power model identification. If indoor temperature is included in the input set, i.e. $u = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr}, \Phi_{occ})^T$, the identification procedure recovers the linear model and leads to an almost perfect accuracy. This is an ideal case however, as in practice indoor temperature is unknown in advance.

The indoor temperature modelling with $u = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, \Phi_{occ})^T$ and output $y = T_{in}$ leads to $R^2=0.86$ and WAPE=0.01.

b) Unmeasured disturbances

Additionally to the characteristic nonlinearity, if unknown inputs are assumed, the accuracy drops to $R^2=0.76$ and WAPE=0.29 for the heating power model, as illustrated in Fig. 4.9.

It's worth mentioning here that if indoor temperature is included in the set of inputs, the model precision is improved: $R^2=0.86$ and WAPE= 0.20.

An equivalent modelling procedure applied to the indoor temperature (inputs $u = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$ and output $y = T_{in}$), results in $R^2=0.85$ and WAPE=0.01. The Fig. 4.10 illustrates the results.

IDA ICE building case. The identification leads to $R^2=0.62$ and MAPE=0.38 for the power consumption model with inputs $u = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$ (see Fig. 4.11) and $R^2=0.84$ and WAPE=0.24 if the indoor temperature is included in the set of inputs.

The temperature model precision is $R^2=0.84$ and $WAPE=0.03$. The goodness of fit between estimated and observed values is illustrated in Fig. 4.12 below.

Overall, this section confirms the linear model limitations when dealing with unmeasured disturbances and complex systems that exhibit nonlinearities, especially when dealing with the HVAC load. Moreover, models turned out highly sensible to the identification time period, requiring a further investigation on how much data should be used for identification, how the model parameters are varying and in which cases the models underperform. For instance, despite the limitations of a linear model, in some cases results suggest that since the indoor temperature is controlled, a temperature model based on temperature setpoints can lead to a satisfying model. If this is possible in some limited cases, linear models are clearly not accurate enough for our application. This motivates exploration of more complex power consumption modelling approaches and precise temperature models, estimations of which could be used for in the power model. The next section will describe the steps taken in this direction.




Figure 4.9: State-space modelling given partially known heat gains affecting the zones, with $u = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$ and output $y = P$. ETP building.




Figure 4.10: State-space modelling given partially known heat gains affecting the zones, with $u = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$ and output $y = T$. ETP building.

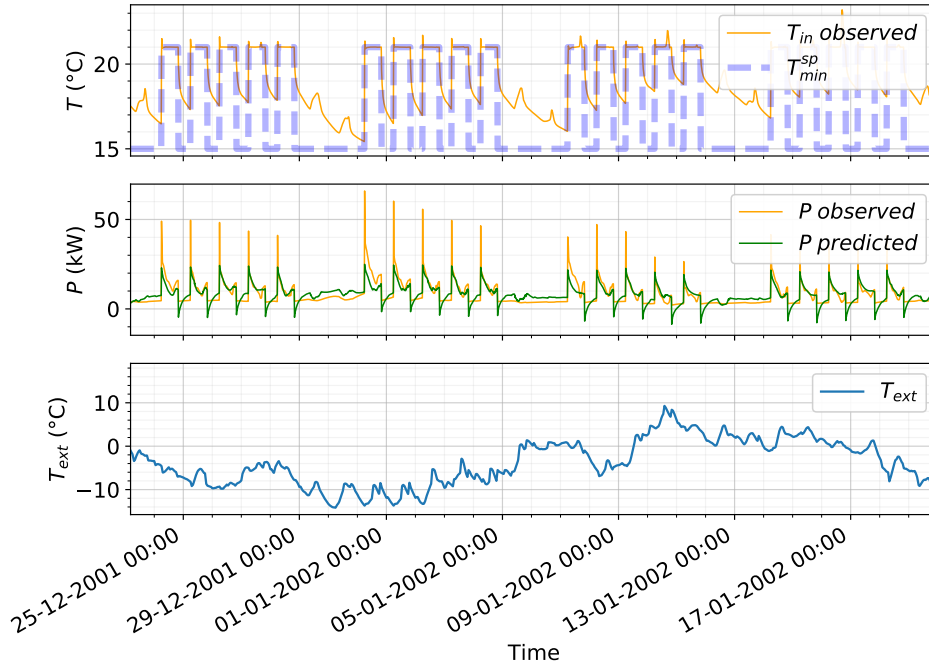


Figure 4.11: State-space modelling given unknown disturbances, with $u = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$ and output $y = P$. IDA ICE building.




Figure 4.12: State-space modelling given unknown disturbances, with $u = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$ and output $y = T$. IDA ICE building.

4.3 Autoregressive Random Forest regression predictive modelling




Figure 4.13: Auto-regressive Random Forest regression predictive modelling of power consumption and indoor temperature. Schematic representation of the approach, input variables and outputs.

From physical insights and from the results presented in the previous section, it can be seen that the indoor temperature carries important information about past states of the building that describes the power consumption profile at a given instant. Therefore, both power and temperature modelling are dealt with in this section. To model the dynamical system, an auto-regressive feature is implemented, that consists in integrating past input and output

values in the regressor set. For this modelling task we use a RF regression and a one step lag value $N = 1$ (15 minutes) for the auto-regressive variables. The function to approximate can be expressed therefore as follows:

$$P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1}) \quad (4.2)$$

where $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$ and the output to estimate is the power consumption. For the power consumption data-driven model, the set of features X is composed of manipulated and non manipulated variables (weather data, indoor temperature estimations, temperature setpoints). As mentioned above, in addition to learning a power consumption model, a second similar model is introduced to predict the indoor temperature in the building based on weather data and temperature setpoints.

For finite horizon predictions the estimated values provided by the models are used recursively until the end of the prediction horizon, i.e. at every time-step of the prediction horizon, first the indoor temperature is predicted. Then the power consumption model uses the past and current temperature forecast, along with past power forecast in order to compute the next time-step forecast.

ETP building case. The models were trained on November-December simulation data, obtained with the 9-zone ETP building. For testing, January-February data was used. Fig. 4.14 below illustrates how the estimation errors are propagated through the prediction horizon.

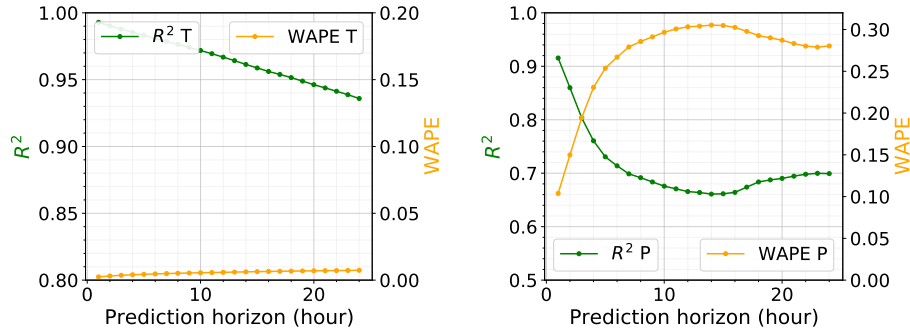


Figure 4.14: Error propagation through the 24h prediction horizon for indoor temperature (left) and power consumption (right) models. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on a 9-zone ETP building with night setback temperature control.

Firstly, an error compensation can be observed in Fig. 4.14, towards the end of the horizon. Overall, given that indoor temperature is tightly following the temperature setpoint, the error metrics are good. However, the power consumption is varying greatly even with small indoor temperature errors and depending on if the indoor temperature is correctly predicted below or above the heating temperature setpoint. Adds up the disturbances during the day (occupancy, solar heat gains), causing the power consumption forecast errors to be concentrated mostly

on daytime periods. The goodness of fit between 1-hour ahead forecasts and observed values is illustrated in Fig. 4.17. The equivalent results for 24-hour ahead forecasts are illustrated in Fig. 4.18.

IDA ICE building case. In this second case the models were trained on one winter season (November - March) and the next winter season was used for testing.

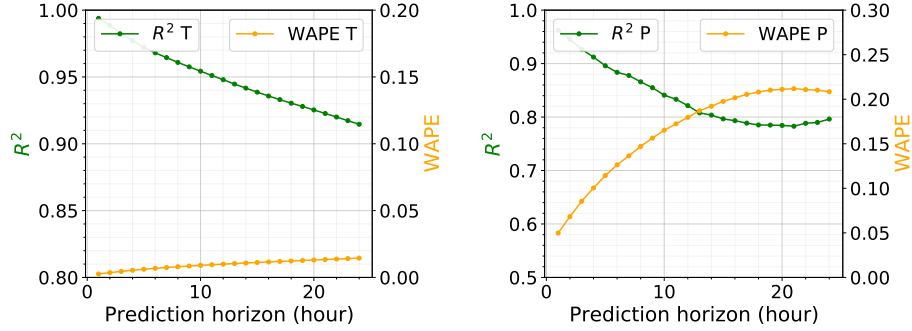


Figure 4.15: Error propagation through the 24h prediction horizon for indoor temperature (left) and power consumption (right) models. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on IDA ICE building with night setback temperature control.

For this second building, the error levels are comparable to the previous building case (Fig. 4.15). It can be seen though that the temperature is not controlled at all times, and the transients observed during night setback and weekends seem more complex to model, resulting in slightly worse error metrics. The goodness of fit between 1-hour ahead forecasts and observed values is illustrated in Fig. 4.19. The equivalent results for 24-hour ahead forecasts are illustrated in Fig. 4.20.

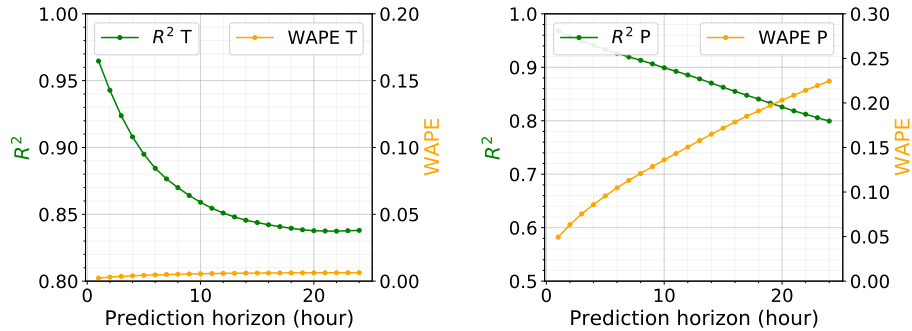


Figure 4.16: Error propagation through the 24h prediction horizon for indoor temperature (left) and power consumption (right) models. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on IDA ICE building with simulated preheating operations.

To validate further the modelling approach, a second simulation dataset was used, with representative instances of DR events, i.e. random preheating operations represented by step temperature setpoints. The results in Fig. 4.16 show the error propagation through the prediction horizon and suggest that the method struggles to reflect correctly the HVAC load variations through the prediction horizon. Fitting results for 1-hour ahead and 24-hour ahead forecasts are illustrated in Fig. 4.21 and Fig. 4.16 respectively.

Given the model structure here, including autoregressive terms in the feature set might lead to a naive learning. A possible solution is modelling, for instance, the variation of power consumption based on previous observations:

$$\Delta P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1}) \quad (4.3)$$

where $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$ and the output to estimate is the power consumption variation with respect to the previous observation. A similar procedure is applied for indoor temperature modelling. Tested on the dataset with simulated preheating operations, the results indicate a similar modelling performance.

Overall the observations suggest that the approach is fragile. The indoor temperature predictions impact highly the power consumption predictions. Also the error propagation through the prediction horizon, due to the autoregressive feature, leads to non-negligible errors. In order to mitigate the effects of the indoor temperature prediction errors, one possibility is considering a larger lag value. The power consumption model could therefore extract information about the thermal state of the building directly from the weather data instead of relying on indoor temperature predictions. Here a lag value $N=8$ of two hours was considered. Since the dataset sampling frequency is 15 minutes, the resulting high number of features could lead to an overfitting problem. The applied solution to alleviate this phenomena consists in considering the average of the autoregressive terms - $\bar{X}_{t-1,...,t-N}$, $\bar{T}_{t-1,...,t-N}$, $\bar{P}_{t-1,...,t-N}$ - over the given lag interval $t-1, ..., t-N$. This can also have the effect of reducing error propagation through the prediction horizon, by breaking the relation between current and previous step observations. The results are presented in Appendix B, and confirm indeed the benefit of considering a larger lag value.

The results above motivate the following bifurcations. One of the further investigations (see subsection 4.4.2) is intended to alleviate the effects of the indoor temperature forecasting errors on the HVAC load predictions, by eliminating the indoor temperature from the set of inputs. The second one is an attempt to alleviate the error propagation through the prediction horizon, by considering a more classical approach, in which no autoregressive terms are considered for the HVAC load forecasting (see section 4.4 below).




Figure 4.17: 1-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on a 9-zone ETP building with night setback temperature control.




Figure 4.18: 24-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on ETP building with night setback temperature control.




Figure 4.19: 1-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on IDA ICE building with night setback temperature control.




Figure 4.20: 24-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on IDA ICE building with night setback temperature control.




Figure 4.21: 1-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on IDA ICE building with simulated preheating operations.



Figure 4.22: 24-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, X_{t-1}, T_{in,t-1})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, X_{t-1}, P_{t-1})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$. Validation on IDA ICE building with simulated preheating operations.

4.4 Random Forest regression predictive modelling




Figure 4.23: Random Forest regression predictive modelling of power consumption and indoor temperature. Schematic representation of the approach, input variables and outputs.

In this section a static power map is identified based on a set of features encompassing weather conditions, indoor temperature, indoor temperature setpoints, time-based information such as Time Of Day (TOD) and Day Of Week (DOW). Adding time-based information to the set of inputs seemed important, since autoregressive terms are no longer considered. Note also that the effect of improved accuracy is not of magnitude of invalidating further comparisons and conclusions. Indoor temperature are supposed to accurately describe the thermal state of the building while the time-based variables, such as the hour of the day or the day of the week, are meant to capture time-related patterns, for instance occupancy or weekday/weekend transitions.

Again, a RF regression algorithm was used. The power consumption function to approximate can be expressed as follows:

$$P_t \approx f_{RF}(X_t) \quad (4.4)$$

where $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$ and the output to estimate is the power consumption P .

The indoor temperature is estimated based on a similar model trained on features such as weather conditions, indoor temperature setpoints, TOD and DOW. In the case of the indoor temperature however, the past inputs are crucial for describing the current output. As suggested at the end of the last section, a lag value N of two hours is applied. The function to approximate can be represented as follows:

$$T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}) \quad (4.5)$$

where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$, the output to estimate is the indoor temperature T_{in} , and $\bar{X}_{t-1,...,t-N}$ is the average of the autoregressive terms over the considered lag interval $t-1, ..., t-N$.

The same training and testing datasets as in the last section were used.

ETP building case. In the ideal case where indoor temperature is known the power consumption model errors metrics are $R^2=0.98$ and $WAPE=0.06$.

The error metrics for the temperature model are $R^2=0.94$ and $WAPE=0.01$. As a result, if forecasted indoor temperature is used to predict the power consumption, the error metrics for the power model become $R^2=0.96$ and $WAPE=0.08$. Fig. 4.24 illustrates the corresponding results.

IDA ICE building case. In the case of IDA ICE building with night setback setpoints, the errors metrics when using observed indoor temperature are $R^2=0.97$ and $WAPE=0.08$.

The error metrics for the temperature model are $R^2=0.91$ and $WAPE=0.01$. If forecasted indoor temperature is used to predict the power consumption, the error metrics for the power model become $R^2=0.86$ and $WAPE=0.17$. Fig. 4.25 illustrates the corresponding results.

To validate further the modelling approach, a second simulation dataset was used, with representative instances of DR events. More specifically, random preheating operations were simulated, represented by step temperature setpoints. When using the observed temperatures, the power model error metrics are $R^2=0.96$ and $WAPE=0.08$. The temperature model however results in $R^2=0.83$ and $WAPE=0.01$. Subsequently, when using forecasted temperatures, the power model inherits the temperature modelling errors, leading to an accuracy $R^2=0.84$ and $WAPE=0.19$.

In this last case, it seems that the diverse set of control actions is leading to a less accurate temperature model. Indeed, the thermal dynamics to be captured depend on the temperature setpoint profile, and in this specific case they are complex. For instance, when night setback is applied, the stored energy in the building is released during the night. As a consequence, the impact of solar or occupancy heat gains during the day will not necessarily lead to a visible variation with respect to the heating temperature setpoints, in contrast to simulating preheating operations case, with no night setback. Additionally in the latter case, the temperature is often in transient phases and a set of behaviours have to be captured by the model. For example, depending on the ambient conditions the preheating temperature setpoints are not always reached, or not reached with the same speed. Also the observed inertia after a preheating episode will depend on the preheating duration and the preheating temperature setpoints.

Given the mediocre accuracy in this last case, where preheating scenarios are simulated, and since this is exactly the type of dynamics we target when dealing with flexibility for DR mechanisms, one may wonder if indoor temperature could be modelled using domain knowledge. For instance, by assuming ramp preheating setpoints, in such a way that the temperature increases by following tightly the setpoint profile, the temperature profile during the preheating operation can be considered as known. Based on this type of control and knowledge of the thermal dynamics, an explicit structured model can be developed. This idea is developed further in the next section.




Figure 4.24: RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N})$, where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$ and $N=8$ (2 hours), and power consumption: $P_t \approx f_{RF}(X_t)$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$. Validation on ETP building with night setback temperature control.




Figure 4.25: RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N})$, where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$ and $N=8$ (2 hours), and power consumption: $P_t \approx f_{RF}(X_t)$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$. Validation on IDA ICE building with night setback temperature control.




Figure 4.26: RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N})$, where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$ and $N=8$ (2 hours), and power consumption: $P_t \approx f_{RF}(X_t)$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$. Validation on IDA ICE building simulated using step setpoints representative of DR control.

4.4.1 Structured data-driven modelling of indoor temperature




Figure 4.27: Structured data-driven modelling of indoor temperature and RF regression predictive modelling of power consumption. Schematic representation of the approach, input variables and outputs.

In the previous sections, data-driven modelling approaches have been investigated and the importance of the indoor temperature when describing the power consumption variation has been illustrated. For the targeted application, the control actions have to be decided in advance, hence, the indoor temperature needs to be estimated. The data-driven models how-

ever struggle to capture correctly the temperature dynamics, especially when a diverse set of temperature setpoint modifications is considered. The temperature dynamics are however more or less difficult to model, depending on how the temperature setpoints are modified. The approach investigated in this section is based on this observation. Given that the indoor temperature is a controlled variable, it is reasonable to assume that the preheating can be performed in such a way that the indoor temperature variation is closely controlled. For instance, slow ramp setpoints allow to rise the indoor temperature tightly at the desired temperature setpoint, therefore the temperature profile during the preheating operation can be considered as known. Given this system knowledge and based on physical insights about the drivers of the indoor temperature dynamics, the following nonlinear explicit model was developed:

$$T^+ = \text{Sat}_{T_{min}^{sp}, T_{max}^{sp}} \left[p_1 T + p_2 \bar{T}_{ext} + p_3 T_{ext} + p_4 \Phi_{irr} + p_5 + e \right] \quad (4.6)$$

$$\bar{T}_{ext}^+ = p_6 \bar{T}_{ext} + (1 - p_6) T_{ext} \quad (4.7)$$

$$e^+ = e + \beta (T_{meas} - T) \quad (4.8)$$

- $p_2 \bar{T}_{ext}$ Inertia
- $p_3 T_{ext}$ Exterior heat gains
- $p_4 \Phi_{irr}$ Solar heat gains
- p_5 Basic permanent heat gains
- e Slow correction term

The saturation expresses the fact that the comfort bounds $[T_{min}^{sp}, T_{max}^{sp}]$ are closely respected. Given the major heat exchanges affecting a zone, the evolution of the indoor temperature is influenced by its current value, by an inertia term \bar{T}_{ext} (associated to slow dynamics of the building's thermal mass that depend on the ambient temperature), by heat exchanges with the exterior T_{ext} , by solar heat gains Φ_{irr} , and eventual unknown auxiliary heat gains. A correction term e allows slow rectification of the predicted temperature based on on-line observations. The corresponding parameter β is chosen such that the correction term accounts for persistent errors. For instance, given the 15 minutes sampling time, $\beta = 0.01$ allows a correction of 0.04°C per degree of error that lasts 1 hour.

It's worth mentioning here that the modelling approach described in this section results from several iterations. A general modelling approach namely, describing all the temperature dynamics, was tested. Instead of the saturation operator a set of parameters are identified for each of the following contexts :

- $T_{min}^{sp} < T_{in} < T_{max}^{sp}$
- $T_{in} > T_{max}^{sp}$ (cooling is on)
- $T_{in} < T_{min}^{sp}$ (heating is on)

The advantage is that the preheating dynamics is modelled precisely for all temperature setpoint profiles. The difficulties arise from different sources. Firstly, the model accuracy is sensible to seasonality effects. To improve the accuracy, identification on short time-intervals was tested as an option, but turned out complicated in practice. Indeed, depending on the considered time intervals, the different contexts are not encountered with the same frequency, which leads often to a poor accuracy and unrealistic parameters. The option of identifying parameters based on ambient temperature, in order to reflect the seasonality effects, doesn't seem to be sufficient. A second major difficulty was the transition between contexts, which the model is not able to perform smoothly.

As before, the model parameters are identified using on a first winter season data, and the identification results are investigated based on a second season data. For the parameter identification procedure, a nonlinear optimization algorithm was used, publicly available at <https://github.com/mazenalamir/torczon>.

The model is tested first in a closed-loop fashion, i.e. considering on-line observations of the indoor temperature and correcting the next-step temperature estimations. The resulting error metrics for the temperature model are $R^2=0.95$ WAPE=0.01. Subsequently, for the power model using estimated temperatures, the errors are $R^2=0.86$ and WAPE=0.20 (in the ideal case, where observed indoor temperatures are used, results are $R^2=0.90$ and WAPE=0.15). However, given that the models are supposed to be used for forecasting a horizon of 24 hours, the correction term cannot be updated continuously. By assuming a correction every 24 hours, the error metrics for the temperature are $R^2=0.93$ and WAPE=0.01. The errors for the power are unchanged with respect to the on-line corrections case: $R^2=0.86$ and WAPE=0.20. The forecasting results are illustrated in Fig. 4.28.

The interest of such an explicit temperature model should be compared with the ML approach seen at the beginning of this section. In the case where a RF temperature model is used ($R^2=0.91$ and WAPE=0.01), the errors of the power model are $R^2=0.88$ and WAPE=0.18, see Fig. 4.29 below.

The interpretability and the accuracy of such a structured explicit model are non-negligible advantages. However, since the model performance is conditioned by the hypothesis of a fine temperature control, additional investigations are necessary to determine how exactly the preheating should be performed. More precisely, should be studied the maximum steepness of the preheating ramps (in order to guarantee that the indoor temperature can follow the setpoints) depending on the exogenous variables, as well as the impact in terms of inertia/passive energy storage. In what follows however, it has been decided to simplify this task by assuming step preheating setpoints. This has been also motivated by the results from the next subsection, showing comparable precision of the HVAC load modelling without predicting the indoor temperature (more precisely, when eliminating the indoor temperature from the regressor and using current and past values of the features), the errors of the power model being $R^2=0.88$ and WAPE=0.18 for this dataset.




Figure 4.28: Indoor temperature predictions based a nonlinear explicit model. RF estimator for power consumption: $P_t \approx f_{RF}(X_t)$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$. Validation on IDA ICE building simulated using ramp preheating setpoints.




Figure 4.29: RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N})$, where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$ and $N=8$ (2 hours), and power consumption: $P_t \approx f_{RF}(X_t)$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$. Validation on IDA ICE building simulated using ramp preheating setpoints.

4.4.2 RF regression predictive modelling of power consumption only




Figure 4.30: RF regression predictive modelling of power consumption. Schematic representation of input variables and output.

To mitigate the effects of the indoor temperature prediction errors, and the cumulative effects of the autoregressive feature errors, the solution investigated in this section consists in eliminating the autoregressive features and circumventing the use of indoor temperature predictions by considering a large enough feature lag value N for power consumption forecasting. The information about the past states of the building could therefore be directly extracted from the weather data instead of relying on indoor temperature predictions. The function to approximate can be expressed therefore as follows:

$$P_t \approx f_{RF}(X_t, X_{t-1}, \dots, X_{t-N}) \quad (4.9)$$

where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$ and the output to estimate is the power consumption P .

Here a lag value N of two hours was considered. To avoid dealing with a high number of features that could lead to an overfitting problem, the average of the autoregressive terms are considered - $\bar{X}_{t-1,...,t-N}$ - over the lag interval $t-1, \dots, t-N$.

The results for the ETP multizone building are $R^2=0.97$, $WAPE=0.07$, for the IDA ICE building with night setback temperature setpoint schedules: $R^2=0.94$, $WAPE=0.11$, in the case of IDA ICE building with preheating events: $R^2=0.94$, $WAPE=0.11$. Fig. 4.31, Fig. 4.32 and Fig. 4.33 below illustrate these results.

In the case where autoregressive terms, with a lag N of 2 hours, are considered, as expressed by the equation hereafter, the accuracy is improved for the beginning of the prediction horizon. The improvement varies however among the datasets. In the case of IDA ICE building with DR preheating operations, the accuracy ranged from $R^2=0.96$ and $WAPE=0.06$ to $R^2=0.93$ and $WAPE=0.13$ at the 24 hours horizon, with an equivalent error as the non-autoregressive variant above, at the middle of the prediction horizon.

$$P_t \approx f_{RF}(X_t, \bar{X}_{t-1,...,t-N}, \bar{P}_{t-1,...,t-N}) \quad (4.10)$$




Figure 4.31: RF power consumption forecasting results based on past and current feature values: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1, \dots, t-N})$, where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$ and $N=8$ (2 hours). Validation on ETP building with night setback temperature control.




Figure 4.32: RF power consumption forecasting results based on past and current feature values: $y_t \approx f_{RF}(X_t, \bar{X}_{t-1, \dots, t-N})$, where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$ and $N=8$ (2 hours). Validation on IDA ICE building with night setback temperature control.




Figure 4.33: RF power consumption forecasting results based on past and current feature values: $y_t \approx f_{RF}(X_t, \bar{X}_{t-1, \dots, t-N})$, where $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr}, TOD, DOW)^T$ and $N=8$ (2 hours). Validation on IDA ICE building with representative DR preheating operations.

4.5 Conclusion

Preliminary studies were performed using linear state-space modelling approaches. Based on two different simulated buildings, the results show the impact of unknown disturbances on one hand, and on the other hand, the limitations of such models when dealing with nonlinearities specific to buildings. These tests also confirmed that modelling the power consumption as a function of indoor temperature (besides other exogenous variables) is much easier than modelling the power consumption as a function of indoor temperature setpoints. On the contrary, since indoor temperature is a controlled variable, it can be accurately modelled based on indoor temperature setpoints. Nevertheless, the robustness and the global precision of this approach is not sufficient for the targeted application.

Moving to more advanced techniques, two different approaches were then tested for temperature and power consumption forecasting: a RF regression with autoregressive features and a "classical" RF regression. In the first case, the propagation of errors through the prediction horizon as well as the impact of indoor temperature prediction errors causes a mediocre accuracy. Considering large lag values attenuates these effects and improves considerably the results. Additionally, an attempt was made to improve the indoor temperature forecasts based on a nonlinear structured dynamic model. The resulting precision was comparable to a RF model. However, the model performance is conditioned by the hypothesis of a fine temperature control with ramp setpoints and additional investigations are necessary to determine how exactly the preheating should be performed. Finally, results indicate that a convenient approach consists in bypassing the indoor temperature forecasting and learning a power consumption model based on current and past feature values (T_{sp}^{min} , T_{ext} , Φ_{irr} , TOD , DOW). Thus, past thermal state information is extracted directly from temperature setpoint and exogenous data instead of, possibly erroneous, predicted indoor temperature. This idea is adopted and refined in the next chapter.

The modelling accuracy for these approaches, at the end of the 24 hours prediction horizon, is summarized in Table 4.1, 4.3 and 4.2.

Experiments suggest also that efforts should be concentrated on precise modelling of the targeted dynamics instead of modelling the whole range of observed phenomena in a building. On this matter, results presented above are already one step in this direction, as the models are learned using heating/winter data only. It should be noted that the mentioned approaches were tested on full year datasets and that the resulting global models are sensibly less precise. This suggests that the best approach is to learn distinct models for winter and summer seasons, and possibly an additional one for dealing with inter-seasonal periods, when both cooling and heating might occur. This is also supported by the fact that HVAC cooling and heating systems might be fundamentally different in terms of technology, efficiency, energy vectors, and consequently display distinctive power consumption profiles.

	ETP building				IDA ICE building			
	Temperature		Power		Temperature		Power	
	R^2	WAPE	R^2	WAPE	R^2	WAPE	R^2	WAPE
AR RF (15 min lag)	0.94	0.01	0.70	0.27	0.91	0.01	0.80	0.20
AR RF (2 h lag)	0.97	0.01	0.92	0.13	0.92	0.01	0.93	0.11
RF	0.94	0.01	0.96	0.07	0.91	0.01	0.85	0.17
RF			0.97	0.06			0.95	0.10
AR RF (2 h lag)			0.97	0.07			0.95	0.10

Table 4.1: Simulation with night setback temperature control. First three approaches are based on chain indoor temperature and HVAC power consumption forecasting, while the last two methods bypass usage of indoor temperature predictions.

	IDA ICE building			
	Temperature		Power	
	R^2	WAPE	R^2	WAPE
AR RF (15 min lag)	0.84	0.01	0.80	0.23
AR RF (2 h lag)	0.81	0.01	0.87	0.17
RF	0.84	0.01	0.87	0.26
RF			0.94	0.11
AR RF (2 h lag)			0.93	0.10

Table 4.2: Simulation with step setpoints, equivalent to preheating operations for DR. First three approaches are based on chain indoor temperature and HVAC power consumption forecasting, while the last two methods bypass usage of indoor temperature predictions.

	IDA ICE building			
	Temperature		Power	
	R^2	WAPE	R^2	WAPE
RF	0.91	0.01	0.88	0.18
Explicit temperature model	0.93	0.01	0.86	0.20

Table 4.3: Simulation with ramp setpoints. Approaches are based on chain indoor temperature and HVAC power consumption forecasting. The first approach uses RF regression indoor temperature prediction, while the second one uses a structured dynamical model. Power consumption is forecasted using RF regression.

Part III

Data-driven modelling for HVAC load forecasting in a DR context

Predictive modelling using boosted trees

Contents

5.1	Structure of the solution	71
5.2	Validation on available building case-studies	73
5.3	Construction of the Demand Response learning dataset	76
5.4	Zoom on preheating and shedded energy depending on the preheating parameters	78
5.5	Zoom on forecasting accuracy during Demand Response events . . .	80

In the previous chapter, several approaches were tested, which aim for short-term forecasting (upcoming day in this case) of the power consumption profile, based on indoor temperature setpoints and other commonly available data such as weather forecasts. The adopted approach stems from the last chapter's conclusions. Here, the approach will be described in detail: features, learning data, validation results on additional buildings, before moving to the next chapter describing the application in the DR context.

5.1 Structure of the solution

Model features As suggested by the results from the last section, instead of estimating the indoor temperature profile, the thermal state of the building is recovered by incorporating information about the past conditions into the regressor, i.e. past mean values of the ambient air temperature and solar irradiance, as well as the past temperature setpoint profile. Fig. 5.1 hereafter illustrates this for the IDA ICE building case.

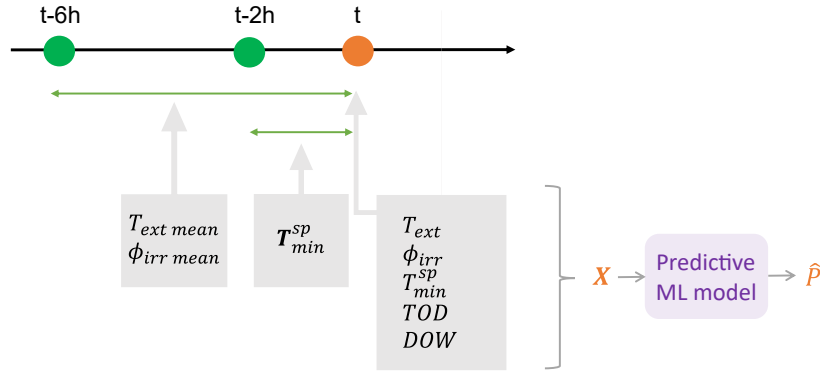


Figure 5.1: Feature set for power consumption forecasting.

Data enhancement In some cases, adding rolling window average values of the originally 15-minutes-sampled data (see Fig. 5.2 below), to the learning dataset, improves the model by smoothing the forecasts, reinforcing the dependency between the target and the features and inferring additional meaning about the features. Are applied here rolling averages up to two hours.




Figure 5.2: Augmenting the original dataset with rolling window averages of the data in order to smoothen the forecasts and reinforce the dependency between the target and the features.

Learning algorithm We use a ML-based algorithm, XGboost (extreme gradient boosting) (T. Chen et al. 2016), for the prediction of the HVAC load. The XGboost algorithm has advantageous characteristics: it does not require fine-tuning of the hyperparameters or intensive pre-processing of the data, it can incorporate interactions between features and capture complex nonlinear relationships between the target and features, it is not prone to over-fitting, it allows model interpretability, and compared to previously used RF algorithms - it is highly efficient in terms of computation time. Note that several other regression ML algorithms have been tested: support vector machines, k nearest neighbors, RF, multi-layer perceptron, long short-term memory networks. Among them, RF and long short-term memory networks models were showing comparable precision.

5.2 Validation on available building case-studies

IDA ICE building (c.f. Section 3.2.1) As in the previous chapter, the model is validated using simulation data consisting in two random preheating events during working hours, one in the first and another in the second half of the day. The accuracy is $R^2=0.95$ and $WAPE=0.10$. Fig. 5.3 and 5.4 below illustrate the prediction errors.




Figure 5.3: Fitting results of the adopted solution. Validation on IDA ICE building simulated with two random preheating scenarios each working day. Train data: November 2001 - March 2002, test data: November 2002 - March 2003.




Figure 5.4: Power consumption forecasting results of the adopted solution. Validation on IDA ICE building simulated with two random preheating scenarios each working day. Train data: November 2001 - March 2002, test data: November 2002 - March 2003.

Digital twin building In the case of the digital twin building (c.f. Section 3.2.3), 137 days in the summer period June-October 2019 happened to contain all necessary features and enough data for training and testing. Synthesizing indoor temperature setpoint data turned out difficult. Indeed, data for more than 180 zones was available, with often inconsistent values across the selected time period. Since a detailed investigation of all the respective time-series seemed time-consuming, indoor temperature setpoints of a single zone was selected in order to test the solution. Although this assumption doesn't allow a rigorous conclusion, given that the proposed method requires averaged data of all zones, it is realistic to assume for this building a certain similarity of temperature control across different zones. When testing on 30% of the available data (a suitable splitting procedure for time-series must be selected), the resulting accuracy is $R^2=0.89$ and $WAPE=0.17$. As illustrated in Fig. 5.5, the HVAC load is correctly captured on average, although it is clear that information is lacking to accurately describe the observed power consumption oscillations.




Figure 5.5: HVAC electrical load forecasting results of the adopted solution. Validation on Digital Twin building. Test data: 30% of the available June-October 2019 data.

A synthetic dataset The synthetic dataset (c.f. Section 3.2.2) offers time-series for different building efficiencies (low, standard or high, depending on factors such as envelope thermal resistance, COP of AHU, etc.), locations (San Francisco - moderate/mild weather, Miami - hot and humid weather, Chicago - cold winter and hot summer) and historical weather conditions. Although indoor temperature setpoints are available, zone areas are difficult to retrieve. Thus, it was decided here to consider a simple average of the setpoints over the zones, without balancing them by the respective zone surfaces.

In the case of the simulation dataset with standard efficiency, San Francisco location and typical meteorological year, the building was mostly in cooling mode. For the selected year, the testing data ratio was fixed to 30%, the resulting accuracy is $R^2=0.97$, WAPE=0.09. As illustrated in Fig. 5.6, the HVAC cooling load function is correctly captured on average. HVAC cooling load was also targeted in the case of Miami location simulation data. Using similar training/testing ratio, the accuracy is $R^2=0.98$, WAPE=0.07. In the Chicago case, given the marked seasonality, only the winter/heating mode was selected. Training and testing was realized using two different winter seasons (January - March), and without including rolling average windows in the final dataset. The resulting accuracy is $R^2=0.92$, WAPE=0.22. Fig. 5.7 below illustrates the results for Chicago location.




Figure 5.6: HVAC electrical cooling load forecasting results of the adopted solution. Validation on a Synthetic building operation dataset. The dataset corresponds to a building with standard efficiency, San Francisco location and typical meteorological year. For testing purposes, 30% of the available data was used.




Figure 5.7: HVAC electrical heating load forecasting results of the adopted solution. Validation on a Synthetic building operation dataset. The dataset corresponds to a building with standard efficiency, Chicago location and typical meteorological year. Two different winter seasons are used for training and testing purposes.

5.3 Construction of the Demand Response learning dataset

In order to design a realistic experimentation for data collection, and enable efficient learning of the HVAC load function, preheating scenarios had to be simulated in a particular fashion, as a result of thermal inertia effects. Indeed, frequent preheating operations induce, over time, an accumulation of passive stored energy, the effect of which can last up to several days. Fig. 5.8 bellow illustrates the difference between the power consumption profile, when preheating operations are performed, and the baseline power consumption profile. One can see that once the setpoints are reset to their baseline value, the power consumption attains its baseline level roughly 12h after the last preheating event.

As a result, past preheating events might influence the energy balance of the future ones. In order to simplify the analysis, this phenomena is avoided by allowing a sufficient delay between two preheating events. In practice, two preheating/shedding events per day are simulated, one in the first half of the day, from 7AM to 12AM, and another from 5PM to 8PM. In these time intervals, random preheating events in terms of start, duration, and temperature setpoint are simulated.

Additionally, the outcome of a given DR operation (DR revenue/modified power consumption) is computed based on the estimation of the baseline energy consumption, i.e. the HVAC load under normal operation. Therefore, in order to accurately estimate the baseline power

consumption profile, preheating operations are performed once every week. The weeks without any preheating operation constitute the baseline learning dataset.

As usually in the case of ML algorithms, an efficient model supposed a rich enough learning dataset. The whole methodology is therefore conditioned by the facility manager permission to modify indoor temperature setpoints, in order to constitute the learning dataset.

The corresponding baseline model accuracy is $R^2=0.90$ and $WAPE=0.09$. Training data corresponds to half of a winter season, and testing data corresponding to the next winter season simulated baseline HVAC load.




Figure 5.8: Power consumption profile given preheating operations and baseline load profile (under normal temperature setpoint control) which reflect the lasting effects of frequent preheating operations.

Several other approaches have been tested for estimating the baseline power consumption. For instance, evaluation of the baseline power consumption based on historical values, consisting in searching the last occurrence of similar operating conditions (indoor temperature setpoint, weather conditions, TOD, DOW), or computing the average of all such historical occurrences. In a more advanced manner, a regression based on k -nearest neighbors was tested. The accuracy of such methods did not surpass the baseline model presented above.

5.4 Zoom on preheating and shedded energy depending on the preheating parameters

A further insight of how the building responds to preheating operations and what inertia phenomena are expected is given first in this section. A specific set of simulations were performed, consisting in preheating scenarios for each working day in a month, with a fixed start of shedding. Fig. 5.9 below illustrates an example of such scenarios for one day.




Figure 5.9: Example of preheating scenarios for one day. Given the fixed start of shedding at 17:00, preheating operations with temperature setpoints [22, 23, 24, 25 °C] and duration of [0.5, 1, 1.5, 2, 2.5, 3 hours] were simulated. The baseline scenario is defined by a temperature setpoint of 21 °C during the day and 17 °C during the night and weekends.

Fig. 5.10 below illustrates, for a given day, how preheating and shedded energy varies depending on the preheating temperature setpoint and duration (the considered shedding duration is 1 hour). A linear relation appears between the total preheating energy and the preheating duration. The shedded energy displays a decreasing rate, due to the limited capacity of storing thermal energy.

The relation becomes more complex when looking at the preheating temperature setpoints. For instance, in some cases the available heating power is not enough to reach the fixed temperature setpoints, since the available heating power is directed by a weather compensation control and depends also on the indoor temperature.




Figure 5.10: Preheating and shedded energy variation depending on the preheating temperature setpoint and duration. The considered shedding duration is 1 hour

The preheating and shedded energy is additionally influenced by weather conditions, occupancy, as well as the past conditions. Fig. 5.11 below illustrates how the preheating and shedded energy varies among days, given a fixed temperature preheating setpoint of 22 °C.




Figure 5.11: Preheating and shedded energy variation depending on the day. The considered shedding duration is 1 hour and the preheating temperature setpoint is 22 °C.

5.5 Zoom on forecasting accuracy during Demand Response events

Preheating operations for passive energy storage, followed by the opposite phenomena when indoor temperature setpoints are reset to their default values, imply fast dynamics, especially since the proposed framework is based on step setpoint modifications. This subsection intends to illustrate the modelling accuracy precisely during preheating and shedding.

As mentioned in the previous section, the learning dataset is built by simulating two preheating/shedding events per day, one in the first half of the day, from 7AM to 12AM, and another from 5PM to 8PM. In this time intervals, random preheating events (in terms of start, duration, and temperature setpoint) are simulated. Preheating/shedding events are simulated every other week. A secondary model, intended for baseline HVAC load predictions, is trained on weeks with no preheating/shedding events. Fig. 5.12, 5.13, 5.14, below illustrates the accuracy on selected preheating timesteps, shedding timesteps (1 hour time-interval after a preheating operation), and the baseline power consumption during the preheating and shedding time-intervals.

In the economical evaluation that follows, the total energy consumption over the time-intervals of interest (preheating/shedding) will be considered. Appendix C illustrates the corresponding accuracy of the mean power consumption over the time-intervals of interest. The accuracy of finer-grained forecasts may not comply with actual requirements of some DR mechanisms. For instance, DR mechanisms organised by the French transmission system operator (RTE), are based on power consumption metering down to 10 minutes time-steps. Nevertheless, requirements in terms of accuracy are still to be adapted to office/commercial buildings end-users, which do not have the same constraints and capabilities as historical flexibility participants. Meanwhile other techniques can be used, in order to cope with accuracy requirements. For instance integrating energy storage devices, or managing a pool of buildings which can be mutually operated to compensate, behind the meter, the eventual commitment errors in a DR framework.




Figure 5.12: Power consumption forecasting accuracy on selected preheating timesteps.




Figure 5.13: Power consumption forecasting accuracy on selected shedding timesteps (1 hour time-interval after a preheating operation).




Figure 5.14: Baseline power consumption forecasting accuracy on selected preheating and shedding timesteps.

Economical evaluation of the energy flexibility potential

Contents

6.1	Demand Response mechanism and participation strategy	84
6.2	Optimization of Demand Response flexibility	86
6.2.1	Delve into optimization results and impact of forecasting errors	89
6.2.2	Building characteristics impact on flexibility capability	91
6.3	Comparison with a classical load shedding approach	92
6.4	Conclusion	94

In this chapter the flexibility value is estimated based on a DR program model. A simplified DR program is assumed, consisting in day-ahead notifications of load shedding. The participating end-user is compensated for the amount of shedded energy.

The compensation is expressed here as a reward factor r , a ratio between the reward level for DR participation and the cost of energy during the preheating time-interval. For instance, if applied in a TOU cost of electricity pricing framework, it represents the ratio between the on-peak and the off-peak energy cost. In this DR framework, r is supposed to represent the intrinsic "value" of the shedded energy, in a period of grid stress, with respect to the cost of the energy in a normal period. This is based on a simplistic assumption according to which the baseline energy price and the reward level are known and fixed by the grid operator/DR program. It follows also that the price of the energy is constant during preheating and similarly, during DR, the compensation rate for the shedded energy is also constant. Although the defined reward factor does not rely on existing DR programs, one may interpret the following results as compensation levels for which buildings may be motivated to adopt a proactive position and participate to grid services.

The shedded energy depends on how the preheating is performed. Hereafter are described the corresponding participation policy and the optimization strategy of the preheating parameters.

An economical estimation of the flexibility value is performed, showing the possible DR income depending on the remuneration for the flexibility service. Further insights are provided about the optimality of the decisions. A comparison is also performed with the classical shedding strategy by lowering the indoor temperature setpoints.

6.1 Demand Response mechanism and participation strategy

HVAC load flexibility is assumed to be activated in the context of a DR program. As already mentioned, the DR program that values energy flexibility "products", is supposed to be based on a day-ahead notification. Practically, DR incentives for specific intervals of time (events) during the next day are supposed to be sent to the end-user. In what follows, it is assumed that two DR notifications are sent each day, which correspond to usual peak hours for the electrical grid: from 7:00 to 10:00 and from 17:00 to 20:00.

Participation in a given DR event is decided by the end-user, by predicting the balance between the preheating cost and the remuneration for the shedded energy.

Fig. 6.1 below illustrates a DR scenario and the energy quantities that have to be estimated.




Figure 6.1: Example of a DR event. Participation to a given DR event is determined by the balance between the preheating energy $\Delta E^{preheat}$ and the remunerated shedded energy ΔE^{shed} .

For a given DR event, the load profile will depend on the preheating settings, more precisely: preheating duration $\Delta t^{preheat} = t_{start}^{shed} - t_{start}^{preheat}$, start of the shedding t_{start}^{shed} , and setpoint increase with respect to the baseline profile $\Delta T_{min}^{sp} = T_{min}^{sp} - T_{min}^{sp, baseline}$.

Let x be the set of decision variables for a given DR event, and θ represent a fixed shedding

duration $\Delta t^{shed} = t_{end}^{shed} - t_{start}^{shed}$ of one hour:

$$\begin{aligned} x &= (\Delta t^{preheat}, t_{start}^{shed}, \Delta T_{min}^{sp}) \\ \theta &= \Delta t^{shed} \end{aligned}$$

Therefore x and θ define the preheating and shedding settings for a given event.

Consequently, the maximum estimated net benefit for a given DR event is defined by the difference between the DR participation reward and the cost of the preheating operation:

$$\hat{B} = \max_x (r |\Delta \hat{E}^{shed}(x, \theta)| - \Delta \hat{E}^{preheat}(x, \theta)) \quad (6.1)$$

where r is the ratio between the reward level for DR participation and the cost of energy during the preheating time-interval.

The estimated and real energy consumption difference, for a preheating setting (x, θ) , are defined with respect to the baseline HVAC load profile, i.e. the load profile in absence of any DR participation, as follows:

$$\Delta \hat{E}^{(\sigma)} = \hat{E}^{(\sigma)} - \hat{E}^{(\sigma)baseline} \quad (6.2a)$$

$$\Delta E^{(\sigma)} = E^{(\sigma)} - \hat{E}^{(\sigma)baseline} \quad (6.2b)$$

where $\sigma \in \{preheat, shed\}$.

The baseline load can only be estimated, and the method used for this intent can be certified by the DR market operators beforehand. Since DR strategies usually consist in curtailing the HVAC system consumption, classical methods for baseline load calculation are mostly suited for this specific use-case. For instance, one of the methods applied by the French TSO consists in setting the baseline value based on metered power consumption before and after a curtailment event. In the case here, such a method is obviously not suited. Among other possibilities, certifying a prediction tool for default load estimation suits the best our case. For this intent, the prediction model is trained on the weeks of the training data without setpoint modifications, as detailed in section 5.3.

Participation to a DR event is triggered if the predicted maximum benefit is positive. The real event benefit is evaluated once the setpoint control has been applied:

$$B = (r |\Delta E^{shed}(x, \theta)| - \Delta E^{preheat}(x, \theta)) \quad (6.3)$$

Subsequently, the total income for a given set of DR events is defined as follows :

$$B_{total} = \sum_{\hat{B} > 0} B \quad (6.4)$$

In this chapter is evaluated, for a limited set of DR events, the ratio between the total DR benefit and the baseline energy cost during the whole winter season, where the DR benefit is the result of the application of computed optimal preheating parameters for each triggered DR event.

6.2 Optimization of Demand Response flexibility

Given the two daily DR notifications: from 7 to 10AM, and from 5 to 8PM, the end-user decides the optimal preheating settings x . The shedding duration is fixed to one hour.

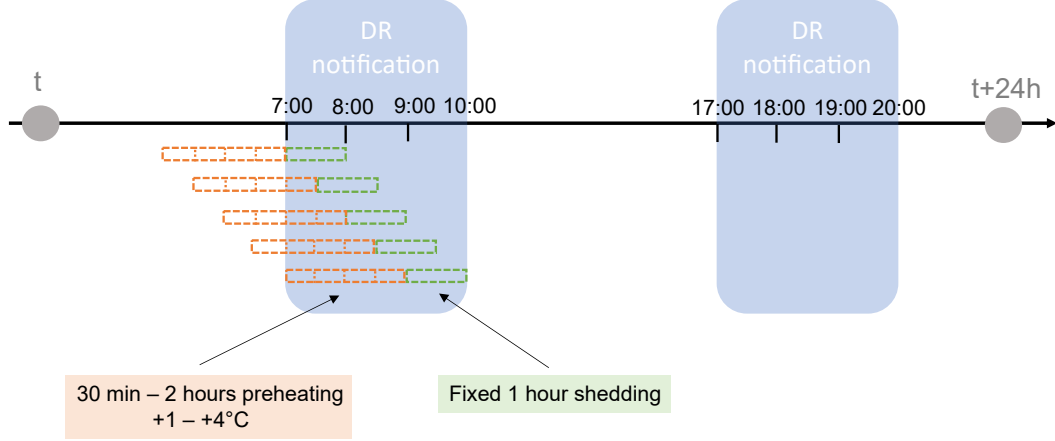


Figure 6.2: Day-ahead DR participation optimization which consists in determining the optimal preheating duration, preheating temperature setpoint and start of the shedding.

The benefit optimization is carried out by exhaustive search of the optimal set of preheating parameters x . More specifically, the decision variables are discretized as follows:

- ⇒ the beginning of the shedding can occur during the DR notifications, i.e. [7:00, 7:30, 8:00, 8:30, 9:00]/[17:00, 17:30, 18:00, 18:30, 19:00],
- ⇒ the preheating duration ranges from 0.5 hour to 2 hours, i.e. [0.5; 1; 1.5; 2 hours],
- ⇒ the setpoint increase for preheating, with respect to the fixed 21 °C baseline, ranges from 1 to 4 °C, i.e. [1, 2, 3, 4 °C].

A total of 80 scenarios are therefore to be evaluated for each DR event in order to determine the optimal setting. Obviously, discretization may not allow to find the global optimum of the expected benefit, but it has been mentioned as a good option to keep the integration in a BMS relatively simple.

Using the designed predictive models, the whole process can be synthesized as follows:

1. For each day-ahead DR notification, and scenario x , predict the corresponding energy consumption $\hat{E}^{(\sigma)}$, as well as the baseline energy consumption $\hat{E}^{(\sigma)baseline}$, where $\sigma \in \{preheat, shed\}$.
2. Estimate the shedded and preheating energy $\Delta\hat{E}^{(\sigma)}$, where $\sigma \in \{preheat, shed\}$ (Eq. 6.2).

3. Based on the reward factor r , evaluate the corresponding economic balance and determine the optimal preheating setting x (Eq. 6.1).
4. If a positive benefit is expected, apply the estimated optimal temperature setpoint modification and measure the shedded and preheating energy $\Delta E^{(\sigma)}$ (Eq. 6.2).
5. Compute a posteriori the actual DR event participation benefit (Eq. 6.1).

Three economical results are presented hereafter. All concern the IDA ICE simulated building.

In the first case, the predictive models (main and baseline) are trained on a whole winter season, November 2001 - March 2002, before integrating the DR mechanism. The total benefit represents the DR outcome for the next winter season, Nov 2002-March 2003. Fig. 6.3 illustrates, for different reward factors r , the ratio between the total DR income and baseline energy cost for the considered winter season.




Figure 6.3: DR participation income relative to the baseline energy cost. Training on Nov. 2001 - March 2002. DR participation during Nov. 2002 - March 2003.




Figure 6.4: Number of triggered events (estimated positive benefit $\hat{B} > 0$).

In a second case, the main predictive model is retrained every month once the DR program is launched. The evaluation, performed only for a reward factor $r = 5$, shows similar total relative income (7.8%), despite retraining with new data. This reveals that an eventual precision improvement in this case does not lead to different optimal decisions. A different conclusion might be drawn if a DR program with commitment error penalties is considered.

In the third case, the participation to DR starts right away and the main forecasting tool is retrained periodically, by incorporating new gathered DR observations in the initial learning dataset. More specifically, the initial learning is based on one month data, November 2002, then every month the predictive model is retrained. Note that, since November 2002 is used for learning only, the total income is computed for December 2002 - March 2003. Note also that the considered baseline predictive model is trained using the same data as in the first case. This assumption, of having prior baseline operation data, might be restrictive. Fig. 6.5 illustrates, for different reward factors r , the ratio between the total DR income and baseline energy cost for December 2002 - March 2003.




Figure 6.5: DR participation income relative to the baseline energy cost. Initial training on Nov. 2002, and retraining every month. The continuous learning case income is compared to the default one, where a whole winter season dataset is used for learning before integrating the DR program. DR participation during Dec. 2002 - March 2003.

By using the proposed method, the flexibility leverages non-negligible income based on realistic DR market assumptions. For a reward factor $r = 5$, the total income from DR participation represents 7.9 % of the total baseline energy cost for the considered winter season.

Seemingly due to the thermal inertia, no increase was observed in the overall energy consumption for the considered winter season November 2002 - March 2003 (less than 0.01%). The reason for this is that the difference between the energy increase during preheating is compensated by the energy decrease during shedding, observed for a longer period than the considered DR time interval.

To allow a degree of freedom regarding the shedding time-frame, it was optimistically assumed here that preheating can be performed during the DR notifications and all the shedding proposals are accepted. In practice, depending on the type of the DR contract, this might not be acceptable. Assuming such a case with random notifications inside previously defined DR time intervals that impose t_{start}^{shed} , preliminary results indicate a 20% benefit loss.

If the DR program is integrated straight away, based on one month learning data, the benefit loss for $r = 2$ is limited to 22 %, while for $r = 5$ the benefit loss is less than 2.5 %. This indicates that, provided with a baseline load prediction tool, the method can be functional in a short amount of time.

In this particular study, a constant indoor temperature setpoint was considered for the baseline calculations. Generally, a night setback temperature control is applied in buildings in order to reduce the total energy consumption. Technically, this could have the effect of yielding a higher relative DR income, although in the particular building studied here, the resulting baseline energy consumption was roughly the same when a reduced temperature setpoint of 15 °C between 22PM and 5AM was tested.

6.2.1 Delve into optimization results and impact of forecasting errors

The optimization process is based on imperfect predictions. Therefore, the optimality of the taken decisions can only be confirmed by simulating the whole set of 80 preheating scenarios. This task being tedious and time consuming, a partial check was carried out in two steps.

In a first step, the optimal start of shedding t_{start}^{shed} was investigated. To this end, all possible scenarios t_{start}^{shed} were simulated, whereas the preheating duration $\Delta t^{preheat}$ and the preheating temperature setpoint increase ΔT_{min}^{sp} , for each DR event, were fixed to the values found via the optimization process described above. The reward factor was set to $r = 5$.

It follows that the decision with respect to t_{start}^{shed} was indeed optimal in 114 cases, out of the total 201 DR events. The benefit loss corresponding to the suboptimal t_{start}^{shed} is 12.5%, corresponding to 1% of loss relative to the baseline energy cost. The matrix hereafter illustrates, for the 201 DR events, the true optimal t_{start}^{shed} versus the estimated optimal t_{start}^{shed} , and the corresponding count values.




Figure 6.6: Optimality of the decisions with respect to the start of shedding t_{start}^{shed} . Subfigure (a) illustrates results for the DR events in the morning and subfigure (b), for the DR events in the evening. According to the matrix, $t_{start}^{shed} = 19:00$ was indeed optimal in 38 evening DR cases. In 15 cases, $t_{start}^{shed} = 18:30$ would have been the optimal decision, instead of 19:00.

In a second step, decisions with respect to the determined preheating duration and temperature setpoint have been investigated. To this end, all the scenarios (a total of 16) of $\Delta t^{preheat}$ and ΔT_{min}^{sp} were simulated, by applying the optimal t_{start}^{shed} determined in the first step.

Except for one event, the estimated optimal preheating duration $\Delta t^{preheat}$ was indeed optimal, and equal to 30 minutes. Indeed, a longer preheating duration doesn't lead to a noticeable benefit given the sensible increase of the preheating energy. This characteristic is expected to vary depending on the thermal properties of the building. In the IDA ICE reference building case, it has been observed that the temperature transients are fast. In the training dataset, once the shedding starts, indoor temperature drops to the baseline setpoint value after roughly 30 minutes in most cases. This suggests that the building has a low flexibility capability and that there is no additional benefit of longer preheating duration.

The resulting benefit loss is 14.5 %, or 1.1 % relative to the baseline energy cost, and is therefore mainly due to suboptimal ΔT_{min}^{sp} and t_{start}^{shed} . This value accounts for prediction errors in all three preheating parameters, although it is not a upper indicator, as not all possible scenarios have been simulated.

An additional result is related to the preheating temperature setpoint. In the 114 cases out of the total 201 DR events, for which t_{start}^{shed} and $\Delta t^{preheat}$ were indeed optimal, the income loss induced by incorrect estimation of ΔT_{min}^{sp} is of 2.4 %. The matrix in Fig. 6.7 hereafter illustrates the true optimal versus the estimated optimal decision counts, for the 114 DR events.




Figure 6.7: Optimality of the decisions with respect to ΔT_{min}^{sp} ($\Delta t^{preheat}$ turned out to be optimal in all cases). According to this matrix, preheating with $\Delta T_{min}^{sp}=4$ °C was indeed optimal in 59 DR events. In 7 cases $\Delta T_{min}^{sp}=4$ °C would have been optimal, instead of the decided $\Delta T_{min}^{sp}=3$ °C.

Estimations of the benefit loss due to prediction errors show that the income loss is 14.5 %, or 1.1 % relative to the baseline energy cost. The study was performed for a reward $r = 5$. It follows that the errors are mostly due to suboptimal t_{start}^{shed} . The optimal preheating duration is correctly predicted in almost all cases.

6.2.2 Building characteristics impact on flexibility capability

In order to get an idea of how the building characteristics could influence the flexibility capacity of a building, several tests have been realized with modified characteristics of the reference IDA ICE building. More precisely, have been tested a version with decreased glazing insulation and a 6-floor version building (cf. Mohamad EL HALLAB, *Development and benchmarking of building energy optimization solutions for demand-response*. Master Thesis, Grenoble INP, 2022).

Decreased glazing insulation Besides a decreased insulation, a higher value of solar heat gain coefficient was set (fraction of the solar radiation that heats the indoor space, which includes both the radiation transmitted directly through the window and the part of solar radiation that is first absorbed by the glazing and passed further through convection/long-wave radiation phenomena).

It has been noticed that such thermal modification had the effect of higher HVAC heating power consumption at night/when solar irradiance is low, and lower power consumption during time-periods when solar irradiance is high. Indeed, during low irradiance periods, the internal

heat dissipates more through the low insulated glazing, which induces a higher HVAC heating power consumption compared to the default building version. When irradiance values are high however, the transmitted radiation through the window surpasses the heat dissipated through conduction, inducing a less power consumption compared to the default building.

An evaluation of the flexibility capability was performed, as for the default building, by optimizing the DR participation with a reward factor $r = 5$. On average, the shedded energy ΔE^{shed} was decreased by 8%, and the preheating energy $\Delta E^{preheat}$ was decreased by 3.6%. The total baseline power consumption for the considered winter season increased however by 1%. Compared to the default building, several instances were encountered where the cooling was activated, obviously inducing high prediction errors.

The total achieved income B_{total} was only 2% lower than in the default building case, being equal to 7.8% relative to the total baseline energy cost.

Increased number of floors By increasing the number of floors, from 3 to 6, the total envelope area relative to the total indoor volume is reduced. Therefore, compared to the default building, the relative thermal heat dissipation through the envelope is lower, which has a potential effect of increasing the thermal inertia.

The evaluation of the flexibility value, as above, allowed to notice an increase of 52% of total benefit B_{total} . Nevertheless, when put in perspective with the increase of 80% of the baseline power consumption, the total benefit relative to the baseline energy cost is 6.8%. Indeed, the preheating energy increased by 87% and the amount of shedded energy increased only by 57%. This behaviour is due to the fact that the building has a very low power consumption during the day. Limiting the heat dissipation by increasing the number of floors has a limited effect on shedding capabilities, as power consumption is already very low.

6.3 Comparison with a classical load shedding approach

In Chapter 2, several examples of data-driven demand response were cited. Among them, approaches proposed in (Nghiem et al. 2017) and (Behl et al. 2016) are also based on setpoint control. Power curtailment however is achieved there by lowering the setpoint with respect to the default schedule, although they do not handle comfort requirements in the same way. This section intends provides some elements of comparison with such a DR strategy. The full modelling and optimization framework is adapted to this type of DR mechanism hereafter.

Construction of the learning dataset The learning dataset was created by simulating, every other week, a sine temperature setpoint profile with a period of 3 hours and 1°C amplitude, such that the maximum value is 21 °C (the baseline value) while the minimum is 19 °C.

As before, two models were trained on alternate weeks data of a full winter season: one

for baseline predictions (based on constant setpoints data), and another on sine temperature setpoint profile. Accuracy of the model on unseen simulation data with sine setpoint profile is $R^2=0.93$, $WAPE=0.12$.

DR mechanism assumptions and participation strategy To ensure a base of comparison with the main results, the following DR market assumptions are taken. As before, it is assumed that two DR notifications are sent each day: from 7 to 10AM and from 5 to 8PM. In contrast to the previous case, the outcome of a shedding operation by lowering the setpoint can only be positive. For a given DR event, the load profile will depend on the start and the duration of the shedding, as well as the setpoint decrease. As before, a fixed shedding duration of one hour is considered. Lowering the setpoint implies comfort degradation, on which a cost must be imputed in order to enable a comparison with the main results. Given the subjectivity of such an estimation, they are not translated/quantified economically here. The setpoints are allowed to be lowered by only 1°C. As in previous approach, the comfort constraints and the margin of action are supposed to be decided beforehand by the facility manager.

Optimization of Demand Response flexibility Since ΔT_{min}^{sp} has been fixed to -1°C, and the shedding duration is fixed to 1 hour, the decision variable is the beginning of the shedding, which can take place during the DR notifications, i.e. [7:00, 7:30, 8:00, 8:30, 9:00]/[17:00, 17:30, 18:00, 18:30, 19:00].

Let $x = t_{start}^{shed}$ and the fixed parameters $\theta = (t_{end}^{shed} - t_{start}^{shed}, \Delta T_{min}^{sp})$, such that x and θ define the shedding for the given DR event. Consequently, the maximum estimated net benefit for a given DR event is defined by the participation reward:

$$\hat{B} = \max_x (-r |\Delta \hat{E}^{shed}(x, \theta)|) \quad (6.5)$$

where r is the ratio between the reward level for DR participation and the default energy cost.

The real and estimated energy consumption difference, for a preheating setting (x, θ) , are defined with respect to the baseline load profile, i.e. the load profile in absence of any DR participation.

Participation to a DR event is triggered if the estimated maximum benefit is positive (negative values may occur due to prediction errors). The real event benefit is evaluated once the setpoint control has been applied:

$$\hat{B} = |r \Delta E^{shed}| \quad (6.6)$$

As already mentioned, the forecasting models are trained on a whole winter season (Nov. 2001 - March 2002) dataset, the DR participation is tested on Nov. 2002 - March 2003. By fixing the reward factor r to 5, the relative income allowed by DR participation is 10.1 %. It's worth mentioning that, in this test, a strategy preventing the rebound effects was

not implemented. The rebound effect neutralized the shedding effect and resulted in 0.2 % increase in the total power consumption over the considered winter season.

Another way of comparing the result hereby with the previous approach is by considering, in both cases, the total DR compensation (the compensation for one event is $r|\Delta\hat{E}^{shed}(x, \theta)|$). It appears therefore that the DR effective income is 25% lower compared to the previous approach.

6.4 Conclusion

In this section, an estimation of the flexibility capability, valued on in the context of a DR program, has been performed using the designed predictive models in an optimization framework. First, the possible DR income has been evaluated in the case where one year of training data is available. For shedding reward factors from 5 to 3, the benefit ranges from 7.9 % to 2.0%, relative to the baseline energy cost (energy cost when no DR actions are performed). Experiments also show that, provided a baseline load prediction tool, the method can be implemented quickly, by continuous model relearning. In this particular case, the loss due to less precise predictions is relatively limited.

Flexibility capacity has also been studied under the angle of building characteristics. Results show that less insulated glazing has negligible effects. The effect of doubling the building floors has the effect of increasing by only 52% the total DR income.

Finally, the proposed method is used to compare load shedding by passive energy storage with the case of a more classical method of load shedding, consisting in decreasing the temperature setpoints by 1 °C. The total DR compensation is in this case 25% lower compared to the preheating framework, and this is achieved at the cost of degraded occupant comfort.

General conclusion and perspectives

This work focuses on efficient activation of the energy flexibility in buildings. The research was carried on data-driven modelling and control of buildings for DR.

Part I provided the motivation for such a research. From the application side, are highlighted different cases where flexibility is valuable. A broad picture of key actors and Demand Response mechanisms is drawn. From an implementation point of view, a recall of the classical approaches for modelling and optimal control is given.

Two main pain-points are highlighted. The first one is the need of a precise model describing the dynamical behaviour of a building. This is often dealt with by designing a complex model, which needs precise knowledge of the building thermal characteristics and HVAC systems. Such models rely on a significant, and thus hardly implementable, amount of sensors and actuators. From a practical point of view, such an approach is time-consuming and implies scalability and integration issues. Simpler models are often based on unrealistic control actions or very simplistic assumptions regarding the HVAC power consumption dynamics. For example, control of the total HVAC power consumption is often assumed and used in energy optimization problems. In real life however this is possible only for a limited type of systems and buildings.

A second pain point is related to the indoor comfort. Classical load curtailment schemes imply, for example, switching off the heating/cooling, or lowering the heating temperature setpoints/increasing the cooling temperature setpoints for a limited amount of time. Building managers seem however reluctant to flexibility proposals that might impact the occupants comfort.

The proposed and investigated approach is put forth from these two perspectives. Firstly, the flexibility is supposed to be activated by indoor temperature setpoint control, inside predefined comfort bounds. This is achieved by shifting the energy consumption, i.e. preheating/precooling the building, therefore enabling passive energy storage, and releasing the stored energy during a DR event. Secondly the flexibility quantification and optimization is done by modelling the dynamical behaviour of the HVAC power consumption, without integrating any precise information about the building characteristics. The modelling is based on ML regression tools and uses commonly available data: indoor temperature setpoints, weather data such as exterior temperature and solar irradiance, total HVAC power consumption.

Part II dealt with exploration of modelling approaches of the HVAC power consumption and indoor temperature dynamics. The estimated HVAC power consumption function is meant to be used for evaluating the result of a setpoint control in a DR strategy. The interest of having an indoor temperature model is twofold. It incorporates information of past thermal states of the building, therefore it can be used to precisely estimate the power consumption at given moment. It can also be used in more classical optimization strategy for energy savings, such as optimal start or optimal stop. This exploration phase showed that a cascaded prediction approach of indoor temperature and power consumption leads to non-negligible errors. A subsequent model has been adopted that bypasses indoor temperature usage.

Part III focused on DR and the refinement of the model to fit this objective. Have been investigated here the experiment design for collecting the learning data, feature selection, baseline estimation for evaluation of the flexibility capability, model accuracy given the observed dynamics. For a simplistic DR program, an optimization framework has been designed for evaluating the DR potential. The study was based on a medium-office use-case, simulated in IDA ICE.

Based on a reward factor (remuneration for the shedded energy compared to the baseline energy cost, which can also be comprehended as on peak/off peak electricity prices ratio) the potential income of activating the flexibility in a DR program has been evaluated. For the reference building, the income for a winter season can represent 7.9 % with respect to the total baseline energy cost, for a reward factor of 5. The study has an interest both for the end-user and for a potential DR program manager. For the user it helps evaluating the interest of a DR program and how the income can differ depending on the building thermal characteristics. For the DR program managers, it provides an indication of what signal/incentive has to be sent to the end-user to effectively enable the wanted flexibility service.

Several challenges are yet to be treated in order to follow with an effective implementation of the framework:

1. **Flexibility optimization:** In this work, the optimization was performed based on an exhaustive search. A finer control could be tested, based on continuous optimization techniques. Since the problem is supposed to be solved offline, computing efficiency is not of special concern. Finer control however, if proven worthwhile, raises the question of model learning capabilities. It has been observed, for instance, that the accuracy of a model trained with preheating operations with temperature setpoint increase of only 1 °C above the baseline, is sensibly higher (compared to the investigated case with random preheating setpoints of 1,2,3 or 4 °C). It has also been observed that the model precision for unseen setpoint values can be quite bad. In the case of further enlargement of the explored setpoint space, one can wonder if an under-fitting problem will arise. In such a case, one may try circumventing the problem by performing an interpolation of two estimates, in order to get a value for an intermediate variable, such as indoor temperature setpoint, not present in the training dataset. Additionally, in the use-case treated here, in a majority of cases the optimization resulted in minimal preheating duration. It would be interesting to investigate how to efficiently explore the research space and refine the results.
2. **DR program:** The flexibility value has been studied here by assuming a simplistic DR program, based on shedded energy remuneration. Different DR program assumptions (variable compensation for shedded energy, variable energy tariffs, commitment penalties, communication delays, etc.) redesign the objective function. Although the income related to the program is expected to change, the proposed framework can be applied directly. The limitation of the proposed method however stems from the variable flexibility potential of the building. A too rigorous framework, for instance default commitment on an amount of shedded energy, with little notification time, seems hardly implementable in buildings. On this subject, some questions deserve a more in-depth investigation.

For instance, the hypothesis of commitment deviation penalties can be studied, together with corresponding optimization strategies. On this matter, slow on-line correction of the benefit estimations has been briefly investigated. First results suggest that the temporal link between the errors is not strong enough for such a method.

Moreover, the question of power consumption metering frequency and the according required modelling precision can be further looked upon. This opens perspectives in terms of optimisation of multiple resources (buildings, storage systems) to compensate the potential errors.

3. **Temperature setpoint control:** For the reference IDA ICE building, used in the complete learning and optimization framework, an identical temperature setpoint profile among the zones was considered, and the control decisions have been applied in an identical manner to the whole set of zones. Two questions can be further investigated. One is optimal control given heterogeneous setpoint profiles among zones. Here it might be relevant to include more specific model features, for instance setpoints for each zone group/type. On this matter, the modelling accuracy was validated on other available building datasets (c.f. Section 5.2), which have heterogeneous setpoint profiles. The second question is linked to controlling only a given set of zones, which can arise in practice. Although the method can be directly applied here, the control impact on the total HVAC power consumption should be especially looked upon.
4. **HVAC subsystem and energy carrier:** In the considered building use-cases, the energy carrier for the HVAC system is electricity. HVAC subsystem or energy carrier heterogeneity constitutes a serious problem to investigate from the modelling point of view. An example that constitutes a challenge is storage systems (hot water buffer tank for instance), which might decouple the energy consumption profile from the control actions. Differences among zones, in terms of installed systems, are also of interest (for instance zones having different or multiple heating/cooling systems installed). The applicable solutions will depend on the available data. If HVAC electrical power consumption data is available, the proposed method can be applied, even though the benefit level is expected to be impacted. Regarding underlying control mechanisms such as energy storage systems or alternating energy carriers, the interest of ML algorithms is specifically in their capability of capturing such complex behaviours. Although the proposed framework doesn't claim a general success, it can be easily tested even in these situations.
5. **Weather and occupancy uncertainties:** Investigations in this work are based on perfect knowledge of weather conditions. A natural extension of the work is looking for how uncertainties impact both modelling and the control strategy, and what solutions can be applied to mitigate the observed effects. A natural approach here is integrating statistical knowledge about weather predictions. The DR commitments could then be decided based on computed income expectation and respective lower and upper bounds. Occupancy uncertainties impact should first be evaluated. In case it is identified as a non-negligible source of errors, additional relevant features can be added to the prediction model, such as seasonal or holiday data. Dedicated forecasting techniques can also be

considered for estimating occupancy levels in buildings. This possibility will depend on the available relevant data.

6. **Building's flexibility potential:** the results presented in this work, regarding the flexibility value, were obtained using a simulated medium-size office building. This particular building seems however quite limited in terms of flexibility capabilities, due to two characteristics. Firstly, the thermal transients observed after a preheating operation are very fast, even for high temperature setpoints and long preheating duration. Indoor temperature reaches the baseline value generally in less than 30 min. Secondly, the building has a very high global thermal efficiency. Therefore during optimization, a significant number of DR events occurred during very low energy consumption, therefore when energy shedding is very low. This suggests that this building is not the one for which the flexibility value is the highest. Broadly speaking, for actors such as Schneider Electric, it would be interesting to explore a panel of different representative buildings (for instance small/medium/large size, high/low efficiency, high/low thermal mass) in order to establish which end-users can provide a flexibility service and at what extent. A study regarding the flexibility potential depending on the weather season or climate could also be of interest. A simulation framework seems an appropriate solution for such an evaluation.

Appendices

Schematic view of the IDA ICE simulated building

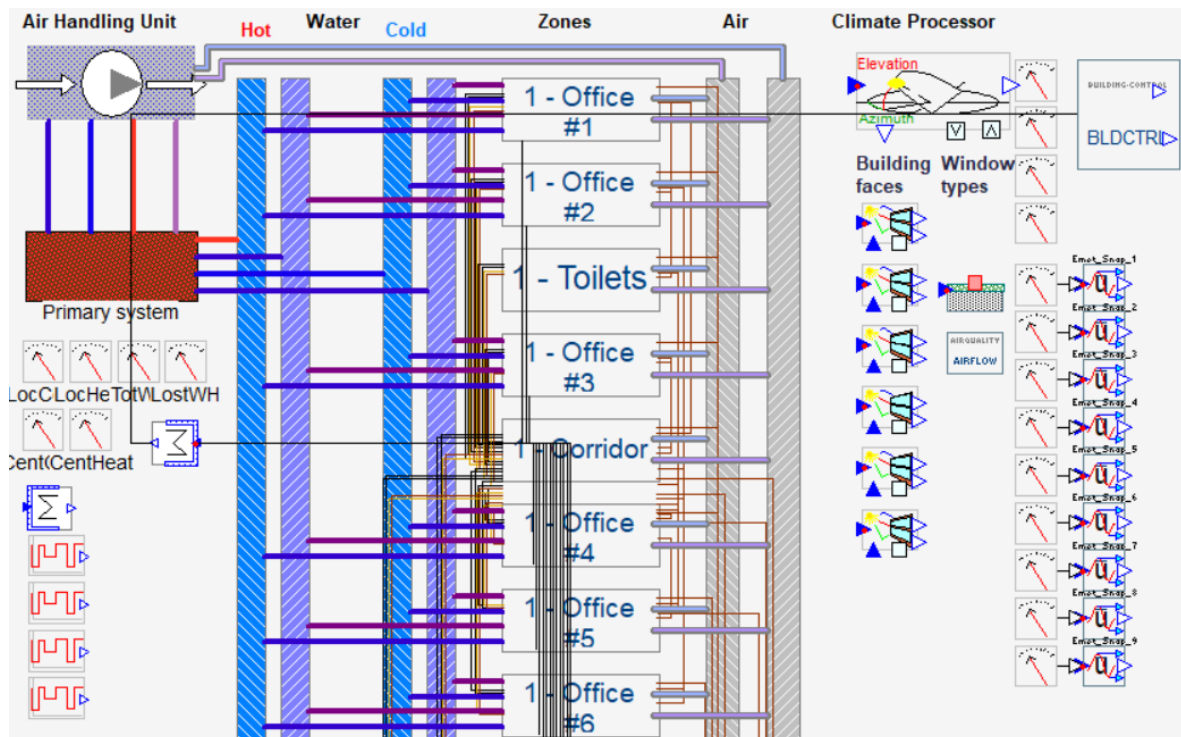


Figure A.1: Schematic view of the building.




Figure A.2: Schematic view of the primary system.




Figure A.3: Schematic view of the AHU system.

AR RF model for power consumption and indoor temperature estimation

In this appendix, results for Autoregressive Random Forest modelling with a lag value N of two hours are presented. Practically, are considered the average of the autoregressive terms - $\bar{X}_{t-1,\dots,t-N}$, $\bar{T}_{t-1,\dots,t-N}$, $\bar{P}_{t-1,\dots,t-N}$ - over the lag interval $t-1, \dots, t-N$.

B.1 ETP building

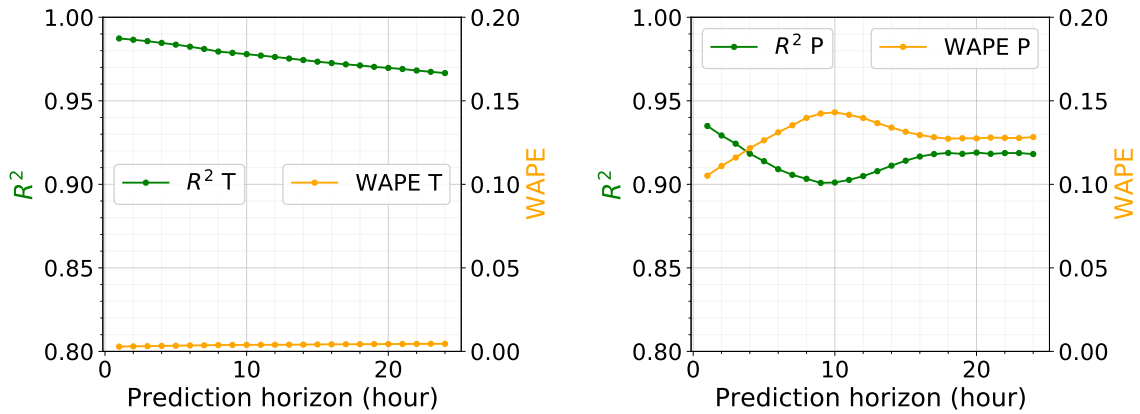


Figure B.1: Error propagation through the 24h prediction horizon for indoor temperature (left) and power consumption (right) models. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{T}_{in,t-1,\dots,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{P}_{t-1,\dots,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, $N=8$ (2 hours). Validation on ETP multizone building with night setback temperature control.




Figure B.2: 1-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{T}_{in,t-1,\dots,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{P}_{t-1,\dots,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, $N=8$ (2 hours). Validation on ETP multizone building with night setback temperature control.




Figure B.3: 24-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{T}_{in,t-1,\dots,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{P}_{t-1,\dots,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, $N=8$ (2 hours). Validation on ETP multizone building with night setback temperature control.

B.2 IDA ICE building with night setback temperature control

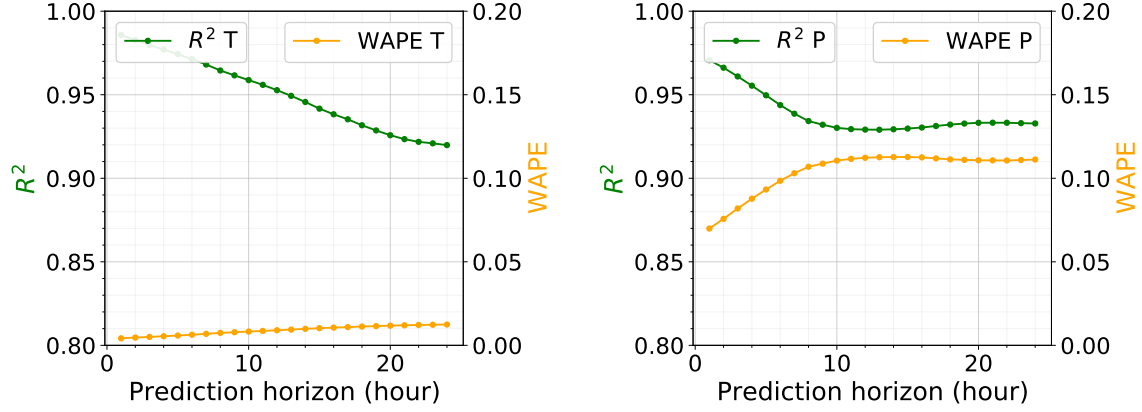


Figure B.4: Error propagation through the 24h prediction horizon for indoor temperature (left) and power consumption (right) models. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{T}_{in,t-1,\dots,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{P}_{t-1,\dots,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, $N=8$ (2 hours). Validation on IDA ICE building with night setback temperature control.




Figure B.5: 1-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{T}_{in,t-1,\dots,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{P}_{t-1,\dots,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, $N=8$ (2 hours). Validation on IDA ICE building with night setback temperature control.




Figure B.6: 24-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{T}_{in,t-1,\dots,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{P}_{t-1,\dots,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, $N=8$ (2 hours). Validation on IDA ICE building with night setback temperature control.

B.3 IDA ICE building with DR representative step setpoints




Figure B.7: Error propagation through the 24h prediction horizon for indoor temperature (left) and power consumption (right) models. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{T}_{in,t-1,\dots,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{P}_{t-1,\dots,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, $N=8$ (2 hours). Validation on IDA ICE building with simulated preheating operations.




Figure B.8: 1-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{T}_{in,t-1,\dots,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{P}_{t-1,\dots,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, $N=8$ (2 hours). Validation on IDA ICE building with simulated preheating operations.




Figure B.9: 24-hour ahead forecasting results. Auto-regressive RF estimators for indoor temperature: $T_{in,t} \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{T}_{in,t-1,\dots,t-N})$, $X = (T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, and power consumption: $P_t \approx f_{RF}(X_t, \bar{X}_{t-1,\dots,t-N}, \bar{P}_{t-1,\dots,t-N})$, $X = (T_{in}, T_{sp}^{min}, T_{ext}, \Phi_{irr})^T$, $N=8$ (2 hours). Validation on IDA ICE building with simulated preheating operations.

Forecasting accuracy during Demand Response events

Hereafter, accuracy is illustrated when considering mean power consumption over the time-intervals of interest (preheating/shedding).




Figure C.1: Mean power consumption forecasting accuracy during preheating.




Figure C.2: Mean power consumption forecasting accuracy during shedding (1 hour time-interval after a preheating operation).

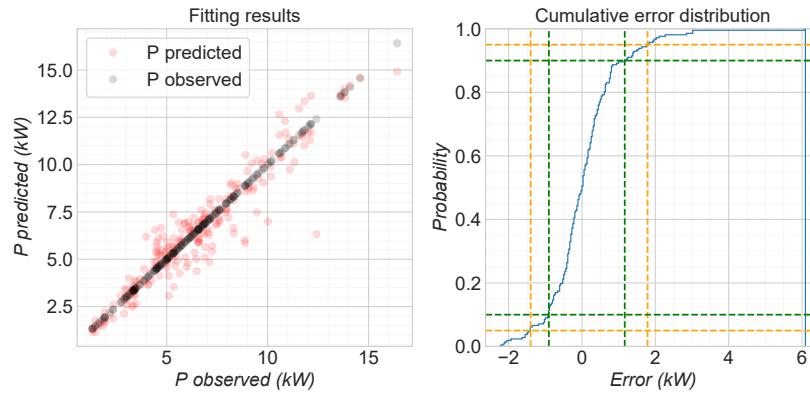


Figure C.3: Mean baseline power consumption forecasting accuracy during preheating.

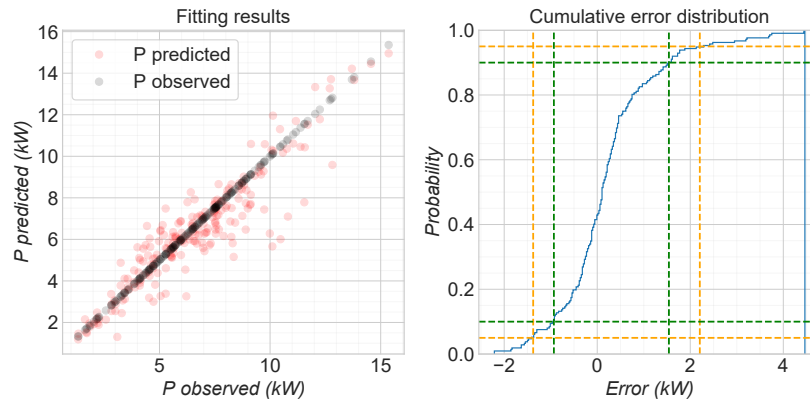


Figure C.4: Mean baseline power consumption forecasting accuracy during shedding.

Load shedding by lowering the temperature setpoints




Figure D.1: Power consumption profile under shedding operations by lowering the temperature setpoints, and baseline load profile (under normal temperature setpoint control).

Résumé en français

E.1 Introduction

E.1.1 Contexte général

Le monde énergétique aujourd'hui subit un changement total de paradigme. On observe une transition d'une production énergétique centralisée et des interactions linéaires entre les producteurs, les distributeurs les consommateurs, vers une interaction plus complexe où le consommateur joue un rôle actif. Cette transformation est due à plusieurs facteurs: la dérégulation du marché de l'énergie, la nature intermittente des nouveaux moyens de production de l'électricité, la multiplication de voitures électriques et des moyens de stockage de l'énergie. Dans ce contexte incertain, la flexibilité énergétique des consommateurs est devenu un moyen crucial pour la gestion efficace du réseau électrique et un vecteur de la transition énergétique.

La flexibilité de la consommation énergétique est appelée aussi réponse à la demande ou DR (Demand Response). Elle comprend les moyens mis en place pour modifier le profil de consommation énergétique des utilisateurs, à la demande de parties tierces, tels que le gestionnaire du réseau de transport ou de distribution de l'électricité, des agrégateurs, des opérateurs des réseaux électriques indépendants. L'intérêt de disposer d'une flexibilité au niveau des utilisateurs est multiple: réduire la consommation d'énergie carbonée, atténuer les courbes de la demande électrique, assurer la stabilité du réseau, éviter les investissements dans la construction des nouveaux moyens de production et l'extension ou la consolidation du réseau électrique.

Certains programmes de DR tels que les tarifs d'électricité heures creuses/heures pleines ont été adoptés à grande échelle. Néanmoins, le cadre lié à la flexibilité énergétique est en pleine évolution. Des nouvelles opportunités sont étudiées et confirmées, des nouveaux acteurs intègrent le marché, des cadres réglementaires sont établis et des stratégies d'exploitation de la flexibilité énergétique sont mises en place.

Les bâtiments, dans ce contexte, ont été identifiés comme des systèmes disposant d'une flexibilité non négligeable, compte tenu de l'inertie thermique les caractérisant. Cette ressource a l'avantage de ne pas nécessiter des coûts d'investissements ou d'exploitation importants. Le système HVAC (en français CVC pour Chauffage, Ventilation et Climatisation) compte généralement pour une grande partie de la consommation énergétique totale d'un bâtiment. L'opération efficace d'un tel système, en termes de coût, de consommation énergétique ou bien émissions de CO_2 , nécessite une adaptation du contrôle de la "flexibilité" au contexte

(météo, usages, prix d'énergie, etc.) dans lequel le système est opéré. La flexibilité peut être produite de différentes manières et l'impact exact des nombreuses options disponibles est souvent mal connu. En effet, la diversité des bâtiments en termes de caractéristiques thermiques, de systèmes HVAC, d'actionneurs et de capteurs accessibles est considérable. Le comportement global d'un bâtiment et l'interaction entre les sous-systèmes HVAC présentent des caractéristiques complexes et non linéaires, rendant difficile l'optimisation du pilotage de la consommation énergétique. Les méthodes traitées dans la littérature et appliquées en pratique diffèrent de plusieurs points de vue: objectif d'optimisation; méthode de contrôle; complexité de modélisation; diversité des informations, capteurs, ou actionneurs; impact sur le confort des occupants; précision de l'évaluation des actions de pilotage.

E.1.2 État de l'art

Les travaux ayant comme objectif le contrôle opérationnel optimal des systèmes HVAC sont, pour la plupart, basés sur une approche de modélisation du bâtiment et son système HVAC.

La modélisation diffère en termes de complexité, et donc de précision. Typiquement une modélisation relativement complexe est réalisée en utilisant un large panel de variables: ouverture des vannes, débits d'eau chaude/froide, vitesse des ventilateurs, consommation électrique des sous-systèmes HVAC, consommation électrique par zone, gains de chaleur secondaires (liées à l'irradiation solaire ou à l'occupation), etc. Dans la pratique, l'accès à une telle variété d'actionneurs/capteurs/compteurs est très souvent irréaliste ou demande des investissements conséquents. De plus, une telle approche n'est pas facilement reproductible, demande une expertise poussée et peut nécessiter un grand effort d'intégration dans un système informatique existant de gestion technique de bâtiment.

Une autre tendance consiste à approximer le comportement thermique du bâtiment via une modélisation moins complexe. L'approche repose généralement sur une modélisation simple des équations différentielles dynamiques représentant les principaux phénomènes thermiques (conduction, convection, phénomènes capacitifs). En fonction de la structure choisie pour modéliser ces phénomènes, la fonction mathématique peut avoir un sens physique. Les paramètres de cette fonction sont identifiés sur la base de données d'entrée et de sortie (les entrées étant les variables explicative et les sorties, les variables à expliquer). La modélisation du système HVAC n'est généralement pas proposée, la variable de contrôle étant typiquement la consommation énergétique, ou une fonction associée très simplifiée. Or en réalité la puissance consommée n'est souvent pas pilotable directement, et le lien entre la commande et la consommation est complexe.

Enfin, d'autres approches existent, ne nécessitant pas une modélisation préalable, parmi lesquelles l'apprentissage par renforcement. Plusieurs études ont montré la capacité de ces algorithmes à exploiter efficacement la flexibilité énergétique des bâtiments. Des efforts considérables sont cependant nécessaires afin de réduire la quantité de données nécessaire pour aboutir à des résultats satisfaisants.

E.1.3 Objectif des travaux et approche investiguée

Ces travaux ont eu pour but l'investigation de méthodes de modélisation et de contrôle optimal, visant à exploiter la flexibilité énergétique des bâtiments. Les méthodes considérées devaient répondre à plusieurs objectifs: permettre l'exploitation optimale de la flexibilité énergétique des bâtiments et une quantification, en termes de consommation énergétique, des options possibles; être facilement déployables sans connaissance des sous-systèmes spécifiques ou des caractéristiques thermiques à chaque bâtiment; utiliser uniquement des données généralement disponibles; ne pas impacter le confort thermique des occupants.

Compte tenu des difficultés d'implémentation réelle des méthodes mentionnées plus haut, une approche pratique a été envisagée, qui consiste à modifier les consignes de températures intérieure dans le but de piloter la consommation HVAC. L'enjeu conséquent est de déterminer comment, et dans quelle mesure, le profil de consommation énergétique peut être correctement approximé sur la base des consignes de température intérieure et d'autres données disponibles telles que les conditions météo. Pour cet objectif, les méthodes d'apprentissage Machine Learning (ML) sont alléchantes grâce à leur capacité d'approximer des dynamiques non-linéaires complexes.




Figure E.1: Un exemple d'opération de DR. La consigne T_{sp}^{min} dans le bâtiment est augmentée à 25°C avant un événement de DR. Cela permet une accumulation d'énergie dans l'air et les éléments structurels du bâtiment, qui peut être déchargée à la demande, en remettant la consigne à sa valeur de 21°C par défaut.

Par souci de clarté, uniquement la saison d'hiver/chauffage a été traitée, même si la saison d'été/climatisation peut être traitée de la même façon, sans perte de généralité. Il est supposé ici que le confort dans un bâtiment est défini par une consigne de température de chauffage T_{min}^{sp} et une température de consigne de climatisation T_{max}^{sp} . Ainsi, afin de répondre à l'objectif de maintien du confort thermique, tout en procédant à une baisse/effacement de consommation énergétique à la demande, le décalage de la consommation énergétique s'impose comme solution. Concrètement cela s'opère en préchauffant les espaces, préliminairement à une opération de DR visant une réduction de la consommation.

Le préchauffage peut être opéré en augmentant temporairement la consigne T_{min}^{sp} globale, dans la bande de confort prédéfinie par le gestionnaire du bâtiment. L'opération mène ainsi à un stockage passif d'énergie, grâce aux propriétés thermiques du bâtiment. Cette énergie peut être ensuite libérée lors d'un événement de DR visant la réduction de la consommation, en remettant la consigne T_{min}^{sp} à sa valeur par défaut. La Figure E.1 ci-dessus illustre une telle opération en hiver. La consigne globale de température de chauffage T_{min}^{sp} est fixée à 21°C. Pour répondre à une demande d'effacement de consommation d'énergie à partir du 7h, sans impacter le confort thermique, la consigne est augmentée à 25 °C pour une durée d'une heure. A 7h la consigne est remise à sa valeur par défaut, permettant ainsi une réduction de la consommation énergétique.

Ce type d'opération peut être envisagé dans le cadre de plusieurs programmes de DR: tarifs variables de l'électricité, participation aux services système, tarification des pointes. Néanmoins, le bénéfice potentiel pour le bâtiment participant doit être calculé par rapport au coût induit par l'opération de préchauffage. Par ailleurs, la compensation pour un service de flexibilité doit être calculée par rapport à l'estimation d'un profil de consommation de base.

Cette étude est basée sur un modèle simplifié de DR, qui se matérialise par des incitations financières transmises en amont, pour la réduction de la consommation énergétique pour des durées limitées de temps. Le travail est ainsi articulé pour répondre à cet objectif. Une première étape est la mise en place un outil qui permet d'estimer la dynamique de la puissance totale consommée par le système HVAC, en fonction de la consigne globale de température appliquée et d'autres données disponibles telles que les conditions météo. Une deuxième étape vise à intégrer l'outil prédictif dans un cadre d'optimisation, afin de déterminer l'opération optimale à effectuer, en fonction du niveau de rémunération pour le niveau de flexibilité atteint, et compte tenu des conditions exogènes dans lequel le bâtiment est opéré.

E.2 Étude exploratoire des méthodes de prédiction de la consommation énergétique et de la température intérieure

Plusieurs approches ont été testées pour la prévision de la consommation énergétique des bâtiments. Le cadre fixé dans la section précédente implique la garantie du maintien du confort à tout moment. Cependant, l'estimation de la dynamique de la température intérieure, en fonction du contrôle, permettrait d'étendre ce cadre à d'autres opérations d'optimisation

énergétique, tel que la commande optimale de marche/arrêt du système HVAC. De plus, la température intérieure reflète l'état thermique du bâtiment, et permet d'estimer précisément la consommation énergétique à un moment donné. Par conséquent, la prédiction de la température intérieure également été étudiée.

Ces études ont eu également pour objectif de forger une compréhension du comportement dynamique des bâtiments et de la puissance consommée associée, ou encore de l'impact des variables inconnues sur la précision des modèles. Deux bâtiments ont été utilisés pour les expérimentations: un bâtiment paramétrique basée sur des équations différentielles décrivant la dynamique thermique, et un bâtiment de bureaux simulé avec le logiciel IDA ICE.

La Figure E.2 ci-dessous regroupe l'ensemble des méthodes testées.

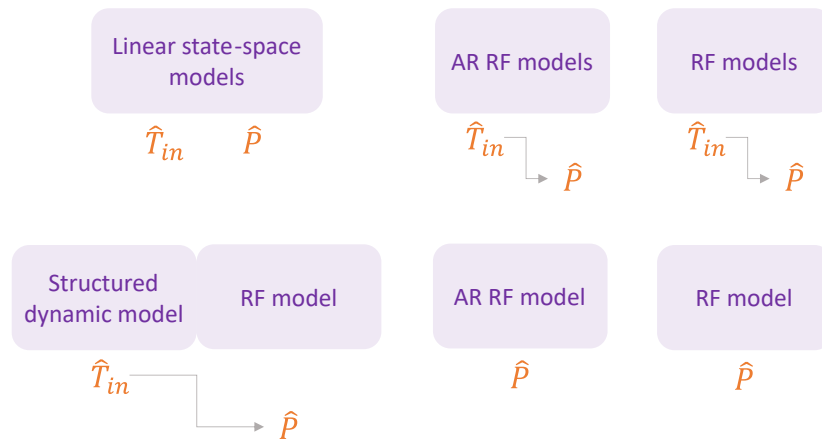


Figure E.2: Méthodes de modélisation explorées et variables à prédire. Les flèches indiquent une prédiction en cascade, de la température intérieure qui est utilisée ensuite pour la prédiction de la puissance consommée.

Dans un premier temps, des modèles linéaires espace-état type "boîte noire", ont été investigués. La capacité de ce type de modèles de représenter les principaux phénomènes thermiques a été testée, ainsi que la précision de modélisation lorsque des apports thermiques non-mesurables impactent le système étudié. Leur limitation par rapport à ce dernier point à été soulignée. La modélisation a également été testée en prenant en compte les consignes de température comme variable d'entrée. Au moins en ce qui concerne la température intérieure, l'intuition ici était basée sur le fait que celle-ci est contrôlée par les consignes de température. Un modèle, même linéaire, pourrait capter cette dynamique en boucle fermée. Globalement, les résultant de cette approche ont motivé la transition vers de modèles plus puissants.

Dans un second temps, des méthodes d'apprentissage supervisé ont été utilisées, basées notamment sur des forêts d'arbres décisionnels. Des versions autorégressives ont été testées également (la sortie actuelle du modèle dépend également des entrées et des sorties passées). Les modèles autorégressifs, dans une perspective de prévision sur un horizon de temps de 24 heures, entraîne une accumulation des erreurs non négligeable.

A été explorée également la prédiction en cascade, de la température intérieure puis de la puissance consommée. Cette exploration a indiqué une sensibilité importante des estimations de puissance aux erreurs de prédiction de la température intérieure. Enfin, les résultats montrent l'intérêt de prédire la puissance consommée actuelle à partir des entrées actuelles et passées. Cette méthode a été retenue pour passer à l'étape suivante où les prédictions de consommation sont utilisées dans le cadre d'un programme de DR. Les sections suivantes discutent donc de l'adaptation du modèle prédictif à objectif, ainsi que de l'estimation des revenus possibles grâce à l'intégration d'un tel programme.

E.3 Modélisation prédictive pour l'exploitation de la flexibilité énergétique

Compte tenu des observations précédentes, la configuration suivante a été adoptée en termes de modélisation:

Variables explicative

- Données météo: température intérieure et irradiation solaire
- Données temporelles: jour de la semaine, heure
- Variables de contrôle: consigne de température intérieure de chauffage
- Moyenne passée des variables météo et profil passé de la consigne de température

Cas d'étude Le cas d'étude est un bâtiment moyen de bureaux simulé avec le logiciel IDA ICE.

Données d'entraînement Afin de répondre à l'objectif de DR, l'expérimentation servant à collecter les données d'entraînement du modèle prédictif devait tenir compte plusieurs observations et contraintes. Premièrement, le résultat d'une opération d'effacement peut être calculée uniquement par rapport à l'estimation d'un profil de consommation de base. Or les résultats de simulation montrent une inertie thermique importante lorsque des opérations de préchauffage répétées sont effectuées. Un laps de temps relativement important doit être donc considéré, après des opérations de préchauffage, avant d'observer une consommation énergétique revenue à la normale.

De plus, le bilan énergétique d'une opération peut être influencé par les opérations précédentes, ce qui peut compliquer voire même fausser les analyses. En conséquence, la base de donnée d'entraînement a été constituée en simulant des opérations de préchauffage, uniquement deux fois par jour, une semaine sur deux. Ainsi, deux modèles ont été constitués: un servant à

l'objectif principal de prédire la puissance consommée résultant d'un changement de consigne de température, et un deuxième servant à la prédiction du profil de consommation de base.

Les données d'entraînement ont été collectées sur la base de simulation d'une saison d'hiver (Novembre 2001 - Mars 2002), avec les données climatiques de Chicago.

Algorithme utilisé Dans cette étude, XGBoost est utilisé comme modèle d'apprentissage.

Validation du modèle prédictif La précision du modèle a été testée sur plusieurs sets de données simulées publiquement accessibles, ainsi que sur le bâtiment de référence simulé avec IDA ICE. Pour ce dernier, la validation a été effectuée sur la saison d'hiver Novembre 2002 - Mars 2003. Le modèle affiche un coefficient de détermination $R^2 = 0.95$, et une erreur de pourcentage absolu pondéré WAPE=0.10.

E.4 Évaluation économique du potentiel de flexibilité énergétique

Hypothèses du programme de DR et stratégies de participation L'activation de la flexibilité énergétique est supposé avoir lieu dans le cadre d'un programme de DR. Dans ce cadre, les participants reçoivent des signaux journaliers pour le lendemain, incitant à l'effacement de la consommation énergétique durant des périodes de temps définies. Deux notifications par jour sont considérés dans cette étude: de 7h00 à 10h00 et de 17h00 à 20h00. La participation à un événement donné est décidé par le consommateur en estimant le bilan entre le coût l'énergie consommée pour effectuer un préchauffage, et la rémunération pour l'énergie effacée.

Pour chaque évènement, le profil de consommation va dépendre de la durée $\Delta t^{prechauffe}$ et de la température de préchauffage par rapport à la consigne ΔT_{min}^{sp} , ainsi que du moment de l'effacement t_{start}^{efface} . Les modèles prédictifs sont ainsi utilisés afin d'évaluer le bilan énergétique en fonction de ces paramètres de préchauffage. Les paramètres optimaux de préchauffage sont déterminés pour chaque évènement, compte tenu d'un taux de rémunération donné (rapport entre le niveau de rémunération pour l'énergie effacée et le prix de l'énergie consommée pendant le préchauffage).

En pratique, par rapport à une consigne de température intérieure de 21 °C, des scénarios de préchauffage de 30 minutes à 2 heures, avec de températures de 22°C à 25°C, ont été évalués. La durée d'effacement a été fixée à 1 heure.

Évaluation économique de l'intégration d'un programme de DR Un premier résultat montre les bénéfices réalisables dans le cas où une saison d'hiver est utilisée pour l'apprentissage des modèles. Le programme de DR est intégré la saison suivante et les ré-

sultats sont estimés par rapport à la facture énergétique de base totale (facture de l'énergie pour la toute la saison). Il en résulte un bénéfice allant de 0.2% à 7.9%, pour des taux de rémunération de $r = 2$ à $r = 5$ respectivement.

Dans un deuxième temps, les données initiales sont enrichies chaque mois avec les données résultant de la participation au mécanisme de DR, et le modèle de prédiction est réappris. Les résultats montrent, pour un facteur de rémunération $r = 5$, un niveau de bénéfice équivalent au premier cas.

Dans un troisième cas, il est supposé que le mécanisme de DR est intégré après seulement un mois de collecte de données pour le modèle principal. Celui-ci est ré-entraîné mensuellement en intégrant les données résultant de la participation au mécanisme de DR. Il est supposé néanmoins qu'un outil servant à estimer la consommation de base est disponible. Les résultats montrent une réduction du bénéfice total, liée aux erreurs de prédiction, limitée à 22% pour $r = 2$, et 2.5% pour $r = 5$.

Des études additionnelles ont été menées, concernant l'impact des caractéristiques thermiques du bâtiment (isolation et taille) sur la capacité de flexibilité. Des éléments de comparaison ont été également fournis par rapport à un effacement classique, impliquant une diminution de la consigne de température intérieure.

E.5 Conclusion générale et perspectives

Les travaux de cette thèse ont eu pour objectif la conception de solutions permettant d'estimer avec assez de précision et de piloter de manière optimale, la flexibilité énergétique des bâtiments. L'approche proposée cherche à répondre à l'objectif fixé, tout en considérant des problématiques d'implémentation réelles : accès aux données, répliquabilité, intégration aux systèmes de gestion existants, barrières d'acceptation liées au confort thermique. Ainsi, elle est basée sur des données (capteurs, compteurs et actionneurs) généralement disponibles et communes à la plupart des bâtiments, et ne dégrade pas le niveau de confort thermique établi par les gestionnaires de bâtiments.

Afin de répondre aux objectifs et contraintes fixées, la consommation énergétique est pilotée par l'intermédiaire du contrôle de la consigne de température intérieure. Une partie de cette thèse est ainsi focalisée sur l'exploration des méthodes visant à prédire, pour un horizon de 24 heures, le profil de puissance totale consommée dans un bâtiment. Uniquement des variables explicatives généralement disponibles sont considérées: consignes globales de la température intérieure, données météo, données temporelles (heure de la journée, jour de la semaine). Sur la base des premières investigations, un modèle prédictif de la puissance consommée, basé sur des outils d'apprentissage supervisé a été adopté.

Dans un second temps, la méthode de prédiction a été raffinée afin d'être utilisée dans un programme de réponse à la demande, incitant à la réduction de la consommation énergétique. L'effacement de la consommation est accompli en utilisant les capacités intrinsèques de stock-

age thermique d'un bâtiment (préchauffage/pré-climatisation).

Sur la base d'un programme rudimentaire, rémunérant l'effacement de la consommation, une stratégie d'optimisation a été mise en place pour effectuer les effacements de manière optimisée. Une estimation des bénéfices tirés de l'intégration de ce programme a été faite. Ce bénéfice atteint 7,9 % par rapport à la facture énergétique totale de base pour la saison d'hiver considérée, lorsque le taux de rémunération est égal à 5 (rapport entre le niveau de rémunération pour l'énergie effacée et le prix de l'énergie de base). Les recherches devraient néanmoins être élargies pour dresser une image complète du potentiel d'une telle approche. Il serait notamment souhaitable d'appliquer la méthode sur d'autres types de bâtiments, différents en termes de caractéristiques thermiques mais également en termes de systèmes HVAC. Des programmes plus complexes de DR devraient également être traités, afin d'obtenir des indications quant aux limites de la méthode proposée.

Bibliography

- Alanne, Kari and Seppo Sierla (Jan. 1, 2022). “An Overview of Machine Learning Applications for Smart Buildings.” In: *Sustainable Cities and Society* 76, p. 103445 (cit. on p. 7).
- Alizadeh, Mahnoosh et al. (Sept. 2012). “Demand-Side Management in the Smart Grid: Information Processing for the Power Switch.” In: *IEEE Signal Processing Magazine* 29.5, pp. 55–67 (cit. on p. 8).
- Armenise, Giuseppe et al. (Sept. 2018). “An Open-Source System Identification Package for Multivariable Processes.” In: *2018 UKACC 12th International Conference on Control (CONTROL)*. 2018 UKACC 12th International Conference on Control (CONTROL), pp. 152–157 (cit. on p. 43).
- Azuatalam, Donald et al. (Nov. 1, 2020). “Reinforcement Learning for Whole-Building HVAC Control and Demand Response.” In: *Energy and AI* 2, p. 100020 (cit. on p. 21).
- Barrett, Enda and Stephen Linder (Sept. 7, 2015). “Autonomous HVAC Control, A Reinforcement Learning Approach.” In: *Machine Learning and Knowledge Discovery in Databases. Joint European Conference on Machine Learning and Knowledge Discovery in Databases. Lecture Notes in Computer Science*. Cham: Springer International Publishing, pp. 3–19 (cit. on p. 20).
- Béguery, Patrick, Yacine Lamoudi, et al. (Jan. 1, 2011). “Simulation of Smart Building - HOMES Pilot Site.” In: *Proceedings of Building Simulation 2011: 12th Conference of International Building Performance Simulation Association* (cit. on p. 30).
- Béguery, Patrick, Melec Petit-Pierre, et al. (July 8, 2017). “Building Energy Simulation Coupled With Real Data For Enhanced Monitoring Analysis.” In: *Building Simulation 2017. Vol. 15. Building Simulation. IBPSA*, pp. 337–344 (cit. on p. 30).
- Behl, Madhur, Francesco Smarra, and Rahul Mangharam (May 15, 2016). “DR-Advisor: A Data-Driven Demand Response Recommender System.” In: *Applied Energy* 170, pp. 30–46 (cit. on pp. 19, 92).
- Berthou, Thomas et al. (June 2012). “Comparaison de Modèles Linéaires Inverses Pour La Mise En Place de Stratégies d’effacement.” In: *XXXe Rencontres AUGC-IBPSA*. Chambéry, France (cit. on pp. 17, 42, 47).
- (May 1, 2014). “Development and Validation of a Gray Box Model to Predict Thermal Behavior of Occupied Office Buildings.” In: *Energy and Buildings* 74, pp. 91–100 (cit. on pp. 14, 42, 47).
- Braun, James E. and Nitin Chaturvedi (Jan. 1, 2002). “An Inverse Gray-Box Model for Transient Building Load Prediction.” In: *HVAC&R Research* 8.1, pp. 73–99 (cit. on pp. 13, 43).
- Chen, Tianqi and Carlos Guestrin (Aug. 13, 2016). “XGBoost: A Scalable Tree Boosting System.” In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 785–794. arXiv: 1603.02754 [cs] (cit. on pp. 34, 72).
- Chen, Yongbao et al. (Nov. 1, 2022). “Physical Energy and Data-Driven Models in Building Energy Prediction: A Review.” In: *Energy Reports* 8, pp. 2656–2671 (cit. on p. 17).

- Ding, Xianzhong, Wan Du, and Alberto Cerpa (Nov. 13, 2019). “OCTOPUS: Deep Reinforcement Learning for Holistic Smart Building Control.” In: *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*. BuildSys ’19. New York, NY, USA: Association for Computing Machinery, pp. 326–335 (cit. on p. 21).
- Dong, Bing, Cheng Cao, and Siew Eang Lee (May 1, 2005). “Applying Support Vector Machines to Predict Building Energy Consumption in Tropical Region.” In: *Energy and Buildings* 37.5, pp. 545–553 (cit. on p. 16).
- Fan, Cheng, Fu Xiao, and Shengwei Wang (Aug. 15, 2014). “Development of Prediction Models for Next-Day Building Energy Consumption and Peak Power Demand Using Data Mining Techniques.” In: *Applied Energy* 127, pp. 1–10 (cit. on p. 16).
- Gao, Guanyu, Jie Li, and Yonggang Wen (Jan. 2019). “Energy-Efficient Thermal Comfort Control in Smart Buildings via Deep Reinforcement Learning.” In: *arXiv e-prints*, arXiv:1901.04693 (cit. on p. 20).
- Guo, Yin et al. (Apr. 1, 2014). “Hourly Cooling Load Forecasting Using Time-Indexed ARX Models with Two-Stage Weighted Least Squares Regression.” In: *Energy Conversion and Management* 80, pp. 46–53 (cit. on p. 16).
- Henze, Gregor P. and Jobst Schoenmann (Mar. 2, 2011). “Evaluation of Reinforcement Learning Control for Thermal Energy Storage Systems.” In: *HVAC&R Research* (cit. on p. 20).
- Jiménez, M. J. and H. Madsen (Feb. 1, 2008). “Models for Describing the Thermal Characteristics of Building Components.” In: *Building and Environment*. Outdoor Testing, Analysis and Modelling of Building Components 43.2, pp. 152–162 (cit. on p. 17).
- Karatasou, S., M. Santamouris, and V. Geros (Aug. 1, 2006). “Modeling and Predicting Building’s Energy Use with Artificial Neural Networks: Methods and Results.” In: *Energy and Buildings* 38.8, pp. 949–958 (cit. on p. 16).
- Kelman, Anthony and Francesco Borrelli (Jan. 1, 2011). “Bilinear Model Predictive Control of a HVAC System Using Sequential Quadratic Programming.” In: *IFAC Proceedings Volumes*. 18th IFAC World Congress 44.1, pp. 9869–9874 (cit. on p. 13).
- Kim, Young-Jin (Sept. 2020). “A Supervised-Learning-Based Strategy for Optimal Demand Response of an HVAC System in a Multi-Zone Office Building.” In: *IEEE Transactions on Smart Grid* 11.5, pp. 4212–4226 (cit. on p. 17).
- Kolokotsa, D. et al. (Sept. 1, 2009). “Predictive Control Techniques for Energy and Indoor Environmental Quality Management in Buildings.” In: *Building and Environment* 44.9, pp. 1850–1863 (cit. on pp. 14, 43).
- Lamoudi, Yacine, Mazen Alamir, and Patrick Béguey (Jan. 1, 2011). “Unified NMPC for Multi-Variable Control in Smart Buildings.” In: *IFAC Proceedings Volumes*. 18th IFAC World Congress 44.1, pp. 11024–11029 (cit. on p. 14).
- Lee, Kyoung-ho and James E. Braun (Oct. 1, 2008). “Model-Based Demand-Limiting Control of Building Thermal Mass.” In: *Building and Environment* 43.10, pp. 1633–1646 (cit. on p. 13).
- Li, Han, Zhe Wang, and Tianzhen Hong (Aug. 10, 2021). “A Synthetic Building Operation Dataset.” In: *Scientific Data* 8.1 (1), p. 213 (cit. on p. 31).

- Li, Xiwang and Jin Wen (Sept. 1, 2014). “Review of Building Energy Modeling for Control and Operation.” In: *Renewable and Sustainable Energy Reviews* 37, pp. 517–537 (cit. on p. 12).
- Li, Yanfei et al. (Aug. 1, 2021). “Grey-Box Modeling and Application for Building Energy Simulations - A Critical Review.” In: *Renewable and Sustainable Energy Reviews* 146, p. 111174 (cit. on p. 13).
- Maasoumy, Mehdi and Alberto Sangiovanni-Vincentelli (Mar. 9, 2016). “Smart Connected Buildings Design Automation: Foundations and Trends.” In: *Foundations and Trends® in Electronic Design Automation* 10, pp. 1–143 (cit. on p. 13).
- Manjarres, Diana et al. (Oct. 1, 2017). “An Energy-Efficient Predictive Control for HVAC Systems Applied to Tertiary Buildings Based on Regression Techniques.” In: *Energy and Buildings* 152, pp. 409–417 (cit. on p. 17).
- Mocanu, Elena et al. (July 2019). “On-Line Building Energy Optimization Using Deep Reinforcement Learning.” In: *IEEE Transactions on Smart Grid* 10.4, pp. 3698–3708 (cit. on p. 21).
- Mustafaraj, G., J. Chen, and G. Lowry (Mar. 1, 2010). “Development of Room Temperature and Relative Humidity Linear Parametric Models for an Open Office Using BMS Data.” In: *Energy and Buildings* 42.3, pp. 348–356 (cit. on pp. 17, 43).
- Nghiem, Truong X. and Colin N. Jones (May 2017). “Data-Driven Demand Response Modeling and Control of Buildings with Gaussian Processes.” In: *2017 American Control Conference (ACC)*. 2017 American Control Conference (ACC), pp. 2919–2924 (cit. on pp. 18, 92).
- Nikovski, D., J. Xu, and M. Nonaka (June 2013). “A Method for Computing Optimal Set-Point Schedule for HVAC Systems.” In: *REHVA World Congress (CLIMA)* (cit. on p. 21).
- Olsthoorn, Dave et al. (June 2017). “Abilities and Limitations of Thermal Mass Activation for Thermal Comfort, Peak Shifting and Shaving: A Review.” In: *Building And Environment* 118, pp. 113–127 (cit. on pp. 9, 10).
- Ruelens, Frederik, Bert J. Claessens, et al. (Sept. 2017). “Residential Demand Response of Thermostatically Controlled Loads Using Batch Reinforcement Learning.” In: *IEEE Transactions on Smart Grid* 8.5, pp. 2149–2159 (cit. on p. 20).
- Ruelens, Frederik, Sandro Iacovella, et al. (Aug. 2015). “Learning Agent for a Heat-Pump Thermostat with a Set-Back Strategy Using Model-Free Reinforcement Learning.” In: *Energies* 8.8 (8), pp. 8300–8318 (cit. on p. 20).
- Sahlin, Per et al. (Aug. 1, 2004). “Whole-Building Simulation with Symbolic DAE Equations and General Purpose Solvers.” In: *Building and Environment*. Building Simulation for Better Building Design 39.8, pp. 949–958 (cit. on p. 30).
- Schubnel, Baptiste et al. (Sept. 2021). “A Hybrid Learning Method for System Identification and Optimal Control.” In: *IEEE Transactions on Neural Networks and Learning Systems* 32.9, pp. 4096–4110 (cit. on pp. 10, 21).
- Sendra-Arranz, R. and A. Gutiérrez (June 1, 2020). “A Long Short-Term Memory Artificial Neural Network to Predict Daily HVAC Consumption in Buildings.” In: *Energy and Buildings* 216, p. 109952 (cit. on p. 16).
- Seyedzadeh, Saleh et al. (Oct. 2, 2018). “Machine Learning for Estimation of Building Energy Consumption and Performance: A Review.” In: *Visualization in Engineering* 6.1, p. 5 (cit. on p. 15).

- Sha, Huajing et al. (Nov. 1, 2019). “A Simplified HVAC Energy Prediction Method Based on Degree-Day.” In: *Sustainable Cities and Society* 51, p. 101698 (cit. on p. 16).
- Solmaz, Aslihan Senel (2020). “Machine Learning Based Optimization Approach for Building Energy Performance.” In: *ASHRAE Topical Conference Proceedings*. American Society of Heating, Refrigeration and Air Conditioning Engineers, Inc., pp. 69–76 (cit. on p. 16).
- Sun, Biao, Peter B. Luh, Qing-Shan Jia, Ziyang Jiang, et al. (July 2013). “Building Energy Management: Integrated Control of Active and Passive Heating, Cooling, Lighting, Shading, and Ventilation Systems.” In: *IEEE Transactions on Automation Science and Engineering* 10.3, pp. 588–602 (cit. on p. 14).
- Sun, Biao, Peter B. Luh, Qing-Shan Jia, and Bing Yan (Aug. 2013). “Event-Based Optimization with Non-Stationary Uncertainties to Save Energy Costs of HVAC Systems in Buildings.” In: 2013 IEEE International Conference on Automation Science and Engineering (CASE), pp. 436–441 (cit. on p. 14).
- Tang, Fan (Dec. 1, 2010). “HVAC System Modeling and Optimization: A Data-Mining Approach.” Master of Science. University of Iowa (cit. on p. 18).
- Urieli, Daniel and Peter Stone (May 6, 2013). “A Learning Agent for Heat-Pump Thermostat Control.” In: *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-Agent Systems*. AAMAS '13. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, pp. 1093–1100 (cit. on p. 21).
- Vardakas, John S., Nizar Zorba, and Christos V. Verikoukis (July 22, 2014). “A Survey on Demand Response Programs in Smart Grids: Pricing Methods and Optimization Algorithms.” In: *IEEE Communications Surveys & Tutorials* 17.1, pp. 152–178 (cit. on p. 8).
- Vázquez-Canteli, José R. and Zoltán Nagy (Feb. 1, 2019). “Reinforcement Learning for Demand Response: A Review of Algorithms and Modeling Techniques.” In: *Applied Energy* 235, pp. 1072–1089 (cit. on pp. 10, 20).
- Wang, Shengwei and Zhenjun Ma (Jan. 1, 2008). “Supervisory and Optimal Control of Building HVAC Systems: A Review.” In: *HVAC&R Research* 14.1, pp. 3–32 (cit. on p. 11).
- Wang, Shengwei and Xinhua Xu (Apr. 1, 2006). “Simplified Building Model for Transient Thermal Performance Estimation Using GA-based Parameter Identification.” In: *International Journal of Thermal Sciences* 45.4, pp. 419–432 (cit. on p. 13).
- Wang, Yuan, Kirubakaran Velswamy, and Biao Huang (Sept. 2017). “A Long-Short Term Memory Recurrent Neural Network Based Reinforcement Learning Controller for Office Heating Ventilation and Air Conditioning Systems.” In: *Processes* 5.3 (3), p. 46 (cit. on p. 20).
- Wei, Tianshu, Yanzhi Wang, and Qi Zhu (June 18, 2017). “Deep Reinforcement Learning for Building HVAC Control.” In: *Proceedings of the 54th Annual Design Automation Conference 2017*. DAC '17: The 54th Annual Design Automation Conference 2017. Austin TX USA: ACM, pp. 1–6 (cit. on p. 21).
- Xu, Guanglin (June 1, 2012). “Modeling and Optimization of HVAC Systems Using a Dynamic Neural Network.” In: *Energy* 42, pp. 241–250 (cit. on p. 17).
- Zhang, Chi et al. (Nov. 13, 2019). “Building HVAC Scheduling Using Reinforcement Learning via Neural Network Based Model Approximation.” In: *Proceedings of the 6th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation*, pp. 287–296. arXiv: 1910.05313 (cit. on p. 18).

- Zhang, Zhiang et al. (Sept. 15, 2019). “Whole Building Energy Model for HVAC Optimal Control: A Practical Framework Based on Deep Reinforcement Learning.” In: *Energy and Buildings* 199, pp. 472–490 (cit. on p. 21).
- Zhao, Jie et al. (May 4, 2015). “EnergyPlus Model-Based Predictive Control within Design–Build–Operate Energy Information Modelling Infrastructure.” In: *Journal of Building Performance Simulation* 8.3, pp. 121–134 (cit. on p. 12).

Résumé — La flexibilité énergétique des bâtiments est une ressource non-négligeable pour la gestion efficace du réseau électrique. La flexibilité peut être produite de différentes manières et l’impact exact des options disponibles est souvent mal connu. En effet, la diversité des bâtiments en termes de caractéristiques thermiques, systèmes CVC, actionneurs et capteurs, est considérable. Les travaux de cette thèse ont eu pour objectif la conception de solutions génériques et extensibles, dans un but d’estimer et d’optimiser la flexibilité énergétique des bâtiments. Les travaux ont été focalisés sur une stratégie de décalage de la consommation, accomplie par le contrôle de la consigne de température intérieure. L’effet en termes de consommation énergétique est estimé par l’intermédiaire d’outils d’apprentissage supervisé. Sur la base d’un programme rudimentaire valorisant l’effacement de la consommation énergétique, une stratégie a été proposée pour réaliser des effacements de manière optimisée.

Mots clés : Chauffage, Ventilation et Climatisation (CVC). Bâtiments. Flexibilité énergétique. Optimisation. Modélisation. Apprentissage supervisé.

Abstract — Energy flexibility of buildings is a major resource for efficient management of the electrical grid. Flexibility can be generated by different manners and the exact impact of various available options is often unknown. Indeed, the diversity of buildings in terms of thermal characteristics, HVAC systems, meters, sensors and actuators is considerable. The objective of this thesis is designing scalable and generic solutions, for precise estimation and optimization of the energy flexibility of buildings. The work is focused on a HVAC load shifting strategy, by indoor temperature setpoint control. The effect of the control in terms of energy consumption is estimated using supervised learning methods. Given a simplistic demand response program, that remunerates load shedding, a framework is proposed to optimize load shedding management.

Keywords: Heating, Ventilation and Air-Conditioning (HVAC). Buildings. Energy flexibility. Optimisation. Predictive modelling. Machine Learning.
