Model-based predictive control methods for distributed energy resources in smart grids

Thèse N° 9264

LUCA FABIETTI

2019

EPFL
ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

"If you look at what you have in life, you'll always have more. If you look at what you don't have in life, you'll never have enough."— Oprah Winfrey

... To my parents.

Acknowledgements

This thesis would not have been possible without the support of many people who, in one way or another, have helped me to get to the end of this fantastic four-year journey.

First of all, I would like to warmly thank my supervisor, Colin Jones, for mainly two reasons. For believing in me and giving me the possibility to carry out my doctoral studies. That choice you made has changed my life in so many positive ways and I am deeply grateful to you for that. Second, for being the best mentor one could hope for. Your guidance, positive attitude, and invaluable advises have helped me improving on both a professional and personal level.

I would also like to thank the members of my thesis committee Prof. Paolone, Prof. Karimi, Prof. Zeilinger, Dr. Alet for dedicating their valuable time to read my thesis giving great suggestions to improve the quality of the manuscript. I would also like to thank them for letting my defence be an enjoyable and interesting experience.

I always like to say that my Ph.d. has been an atypical one in that rarely I have been working on my own. On the contrary, I was lucky enough to extensively collaborate with many bright people. I am sure that this allowed me to accomplish far more than I would have alone. A special thanks to Tomasz, Faran, Altug and Georgios for your friendship and for being as second supervisors to me. You are really great people and I hope we will keep in touch in the future.

A big big thank you to ALL the LA (laboratoire d'automatique) family from the professors to the secretaries, and of course all my amazing colleagues. From the very first day I stepped into LA I felt at home thanks to warm and welcoming atmosphere one can breathe in the laboratory. This is a very precious thing that I am sure the lab will (and definitely should) preserve in the years to come thanks to the great people who populate it. I would like to reserve a special thanks to Christophe Salzmann, the pillar of LA. Thank you for your immense patience and for always finding the time to help us solving all sorts of problems.

I would also like to thank all the people I met in Lausanne who contributed to make these four

years so pleasant. I refrain from listing names for fear of forgetting some dear friends. A special mention to my volleyball buddy Amadou with whom I have won some many matches (and lost at least as many). I would also like to reserve a very warm thank-you to my extended family here in Lausanne. Thank you to Christiane, Roland, Jess, Marc and the little Ellie and Alessia for always making me feel welcome in your beautiful family.

My gratitude and love goes to my family, especially my parents. First, for teaching me, through your optimism and positive attitude, the difference between important and futile things in life. Then for always supporting me in achieving what mattered the most to me. Without you I would not be even close to where I am today.

Finally, my last thanks go to my significant other, Audrey. Thank you for having taken my already beautiful life and having made it even more special.

Lausanne, January 2019   *L.F.*

Abstract

Power grids are undergoing massive changes to reach ambitious targets in terms of reduced carbon dioxide emissions, higher energy efficiency, economic competitiveness and increased security of supply. However, the increasing share of intermittent renewable energy sources connected to the grid challenges the current power grid stabilization paradigms. The increasing need for reserve power, which is now mostly provided by carbon-based generation resources, has brought attention to the provision of regulation services by Distributed Energy Resources (DER)s such as commercial and residential loads, storage devices, electric vehicles, etc. In this respect, this thesis develops optimization-based techniques for the control of DERs to provide multiple services to the power network. It is divided into three parts.

The first part of this thesis focuses on the development of a framework for the efficient control of a single resource that is subject to the effect of periodic stochastic uncertainties. More specifically, resources that can be described by the general class of periodic constrained linear systems are considered and a method, based on Stochastic MPC, to control the over-time-average constraint violations is developed. Finally, the effectiveness of the control framework is tested, by means of a simulation analysis, for the case of the climate control of a building.

The second part of the thesis introduces the required background for the electric power grid, energy markets, and distributed energy resources providing grid support services. First, the control problem of scheduling the operation of a set of energy resources offering multiple services to the grid is formally stated as a multi-stage uncertain optimization problem. In particular, the problem is designed so as to maximize the provision of a shared tracking service while enforcing the satisfaction of the operational constraints on both the individual resources, as well as on the hosting distribution network. Two computationally tractable approximated solution methods are then presented, which are based on robust-optimization techniques and on a linearization of the power flow equations around a general linearization point. A simulation-based analysis demonstrates the capability of the proposed framework to adapt to different levels of uncertainty acting on the overall system. Finally, a quantitative and qualitative comparison between the two approximation schemes is presented and general guidelines are given.

The last part of the thesis demonstrates the practical relevance of the control framework developed in part II. In particular, the aggregation of an electrical battery system and of an office building is

considered, and two case studies are investigated. In both cases, we show how to adapt the general framework of Part II so as to accommodate the given application. Then, we design a hierarchical multi-timescale controller in order to reliably deliver the service by coordinating the resources during real-time operation.

The first case study deals with the provision of secondary frequency control in the Swiss market. An experimental platform comprised of a fully-occupied auditorium building on the EPFL campus and of an emulated electrical battery is developed. We run experiments, in full compliance with the current Swiss regulations, and for nine consecutive days. The second case study deals with the problem of dispatching the operation of an active distribution feeder characterized by the presence of stochastic prosumers. The experimental platform is represented by the EPFL 20 kV distribution feeder and it is comprised of five uncontrollable office buildings, PV generation, a utility-scale electrical battery, and a living multi-zone office building which has been completely retrofitted. Closed-loop experiments are conducted over an extended period of time (12 full days). The results of both experimental campaigns demonstrate the effectiveness and robustness of the control methodology against the wide range of uncertainty involved. In fact, in both cases, high-quality tracking performance could be achieved without jeopardizing the occupants' comfort in the building nor violating the operational constraints of the battery. Finally, the results also show the benefit of combining resources with complementary technical capabilities as the building and the battery.

Keywords: MPC, predictive control, robust and stochastic optimization, model identification, smart grid, frequency control, battery energy storage system, building control

Résumé

Les réseaux électriques subissent des changements massifs pour atteindre des objectifs ambitieux en termes de réduction des émissions de dioxyde de carbone, d'amélioration de l'efficacité énergétique, de compétitivité économique et de sécurité de l'approvisionnement. Toutefois, la part croissante des sources d'énergie renouvelable intermittentes raccordées au réseau remet en question les paradigmes actuels de stabilisation du réseau électrique. Le besoin croissant d'électricité de réserve, qui est maintenant principalement fourni par des ressources de production à base de carbone, a attiré l'attention sur la fourniture de services de régulation par des ressources énergétiques distribuées (DERs) comme les charges commerciales et résidentielles, les dispositifs de stockage, les véhicules électriques, etc. À cet égard, cette thèse développe des méthodes basées sur l'optimisation pour le contrôle des DERs afin de fournir de multiples services au réseau électrique. Elle est divisée en trois parties.

La première partie porte sur le développement d'une méthode de contrôle efficace d'une ressource unique soumise à l'effet d'incertitudes périodiques et stochastiques. Plus précisément, les ressources qui peuvent être décrites par la classe générale des systèmes linéaires périodiques avec contraintes sont prises en compte et une méthode, basée sur "Stohastic MPC", est développée pour contrôler la moyenne dans le temps de violations des contraintes. L'efficacité de la méthode de contrôle est testée, par une analyse numérique, pour le cas de la climatisation d'un bâtiment.

La deuxième partie de la thèse présente les concepts du réseau électrique, des marchés de l'énergie, et des ressources énergétiques distribuées fournissant des services de soutien au réseau. Première-ment, le problème de la planification de l'exploitation d'un ensemble de ressources énergétiques, offrant des services multiples au réseau, est formellement énoncé comme un problème d'optimisation incertain multi-périodes. En particulier, le problème est conçu de manière à maximiser la fourni-ture d'un service de suivi partagé tout en veillant à la satisfaction des contraintes opérationnelles tant sur les ressources individuelles que sur le réseau de distribution. Deux méthodes offrant des solutions approximatives calculables et traçables sont ensuite présentées. Elles sont basées sur des techniques d'optimisation de la robustesse et sur une linéarisation des équations de flux de puis-sance autour d'un point général de linéarisation. Une analyse fondée sur des simulations démontre la capacité du système de contrôle proposé à s'adapter à différents niveaux d'incertitude agissant sur l'ensemble du système. Enfin, une comparaison quantitative et qualitative entre ces deux méthodes d'approximation est présentée et des recommandations générales sont données.

La dernière partie de la thèse montre la pertinence pratique du cadre de contrôle développé dans la partie II. En particulier, l'agrégation d'un système de batteries électriques et d'un bâtiment commercial est considérée, et deux études de cas sont examinées. Dans les deux cas, nous montrons comment adapter le cadre général de la partie II afin de tenir compte de l'application donnée. Ensuite, nous concevons un contrôleur hiérarchique multi-échelles de temps afin de fournir le service de manière fiable en coordonnant les ressources en temps réel.

La première étude de cas porte sur la fourniture d'une régulation secondaire de fréquence sur le marché suisse. Une plateforme expérimentale composée d'un bâtiment d'audience entièrement occupé sur le campus de l'EPFL et d'une batterie électrique émulée est développée. Nous effectuons des expériences, dans le respect de la réglementation suisse en vigueur, pendant neuf jours consécutifs. La deuxième étude de cas porte sur le problème de la répartition de l'exploitation d'un réseau de distribution actif caractérisé par la présence des consommateurs et producteurs d'énergie stochastiques. La plateforme expérimentale est représentée par la ligne de distribution de 20 kV de l'EPFL et se compose de cinq bâtiments non-contrôlés, d'une production photovoltaïque, d'une batterie électrique à échelle utilitaire ainsi que d'un bâtiment multizone contrôlable. Les expériences en boucle fermée se déroulent sur une période prolongée de 12 jours complets. Les résultats des deux campagnes expérimentales montrent l'efficacité et la robustesse de la méthode contre le large éventail d'incertitudes. En fait, dans les deux cas, il était possible d'obtenir des performances de suivi de haute qualité sans compromettre le confort des occupants de l'immeuble, ni violer les contraintes d'exploitation de la batterie. Enfin, les résultats montrent également l'avantage de combiner des ressources avec des capacités techniques complémentaires comme le bâtiment et la batterie.

Mots-clés : MPC, contrôle prédictif, optimisation robuste et stochastique, identification de modèles, réseau intelligent, contrôle de fréquence, système de stockage d'énergie sur batterie, contrôle du bâtiment

Contents

List of Figures

List of Tables

List of abbreviations

AGC Area Generation Control

AHU Air Handling Unit

ALD ASHRAE Likelihood of Dissatisfied

AS Ancillary Services

ASP Ancillary Services Provider

BEMS Building Energy Management System

BESS Battery Energy Storage Systems

BMS Battery Management System

DER Distributed Energy Resource

EPFL École Polytechnique Fédérale de Lausanne

HLC High Level Controller

HP Heat Pump

HVAC Heating, Ventilation, Air conditioning

LLC Low Level Controller

LP Linear Programming

MPC Model Predictive Controller

PCC Point of Common Coupling

PRBS Pseudorandom Binary Signal

PV Photovoltaic

PF Power Flow

PFC Primary Frequency Control

PHEV Plugin Hybrid Electric Vehicles

SFC   Secondary Frequency Control

SMPC Stochastic Model Predictive Controller

SoC   State of Charge

TFC   Tertiary Frequency Control

TCL   Thermal Controllable Loads

TSO   Transmission System Operator

Notation

Throughout the thesis, scalars and vectors are denoted with lower case letters $(a, b, \ldots)$, matrices are denoted with upper case letters $(A, B, \ldots)$. Sets are denoted with upper case blackboard bold letters $(\mathbb{R}, \mathbb{N}, \ldots)$ for number sets, and with calligraphic upper case letters $(\mathcal{A}, \mathcal{B}, \ldots)$ for general sets. a_i represents the value of vector a at time i, whereas bold letters are used to denote sequences over time, $e.g.$, $\mathbf{a} = [a_0^T, a_1^T, \ldots, a_{N-1}^T]^T$.

General Operators and Relations

$\cdot$	general placeholder
j	the imaginary unit, $j := \sqrt{-1}$
$:=$	left-hand side is defined by the right-hand side
$\|$	such that
$\in$	is element of (belongs to)
$\forall$	for all
$\exists$	there exists
$\notin, \nexists$	/ denotes negation
$\wedge$	and
$\Rightarrow$	implies

Sets, Spaces, and Set Operator

$\{\cdot,\dots\}$	set or sequence
$\varnothing$	empty set
$\mathbb{R}$	real numbers
$\mathbb{N}$	natural numbers
$\mathbb{C}$	complex numbers
$\mathbb{N}_+$	set of non-negative natural numbers
$\mathbb{N}_j^k$	set of consecutive non-negative integers $\{j,\dots,k\}$
$\mathbb{R}^n$	space of n-dimensional (column) vectors with real entries
$\mathbb{R}^{n\times m}$	space of n by m matrices with real entries
$\mathbb{C}^n$	space of n-dimensional (column) vectors with complex entries
$\mathbb{C}^{n\times m}$	space of n by m matrices with complex entries
$(\subset)\subseteq$	(strict) subset
$(\supset)\supseteq$	(strict) superset
$\times$	Cartesian product, $\mathbb{X}\times\mathbb{Y}:=\{(x,y)\,|\,x\in\mathbb{X},\,y\in\mathbb{Y}\}$

Operators on Vectors and Matrices

$[\cdot,\dots]$	a matrix (or a vector)
$A\otimes B$	Kronecker's product between matrix A and B
v^T	row vector, transpose of vectors
$\|v\|$	(any) vector norm
$\|v\|_1$	l_1-norm or vector 1-norm (sum of absolute values)
$\|v\|_2$	l_2-norm or vector 2-norm (Euclidian norm)
$\|v\|_\infty$	l_∞-norm or vector ∞-norm (largest absolute element)
$\mathbf{I}_{n\times n}$	identity matrix of dimensions $n\times n$
$\mathbf{0}_{n\times n}$	zero matrix of dimensions $n\times n$
$\mathbf{0}_n$	n-dimensional vector of zeros
$\mathbf{1}_n$	n-dimensional vector of ones
M^T	transpose of a matrix
M^{-1}	inverse of a square matrix
$M>(\geq)0$	all elements of the matrix M are positive (non-negative)
$M\succ(\succeq)0$	the matrix M is positive definite (positive semidefinite)
$\bar{M}$	the complex conjugate of the matrix M.
$\|M\|$	(any) matrix norm
$\|M\|_1$	induced matrix 1-norm
$\|M\|_2$	induced matrix 2-norm
$\|M\|_\infty$	induced matrix ∞-norm

Systems and Control Theory

n_x	number of states, $n_x \in \mathbb{N}_+$
n_u	number of inputs, $n_u \in \mathbb{N}_+$
n_w	number of disturbances, $n_w \in \mathbb{N}_+$
n_y	number of outputs, $n_y \in \mathbb{N}_+$
x	state, $x \in \mathbb{R}^{n_x}$
u	control input, $u \in \mathbb{R}^{n_u}$
w	disturbance, $w \in \mathbb{R}^{n_w}$
y	output, $w \in \mathbb{R}^{n_y}$
$\mathbb{X}$	set of state vectors, $\mathbb{X} \subseteq \mathbb{R}^{n_x}$
$\mathbb{U}$	set of input vectors, $\mathbb{U} \subseteq \mathbb{R}^{n_u}$
$\mathbb{W}$	set of disturbance vectors, $\mathbb{W} \subseteq \mathbb{R}^{n_w}$
$\mathbb{Y}$	set of output vectors, $\mathbb{Y} \subseteq \mathbb{R}^{n_y}$

Probability Theory

$\mathbb{P}(\cdot)$	probability measure
$\mathbb{E}\{\cdot\}$	expectation with respect of the probability measure $\mathbb{P}$
$\mathbb{I}[x > 0]$	indicator function for the event $[x > 0]$

Chapter 1

Introduction

The need to modify the current structure of the electric grid into a more sustainable configuration motivates a rapid increase in renewable generation. However, renewable energy sources are intrinsically volatile and uncertain and this poses new engineering challenges in order to guarantee the proper and safe functioning of the grid [93]. The balance between demand and generation of active power, at any time, is typically guaranteed by a set of fossil fuel-based spinning reserves which are kept on standby and activated to compensate for any mismatch between generated and consumed power. The increasing share of uncontrollable power units, *e.g.*, photovoltaic panels or windmills, will jeopardize the current controlling mechanism resulting in an electric grid that strongly relies on the active control of physical as well as *virtual* storage systems.

Potential candidates could be represented by, *e.g.*, battery energy storage systems (BESS) or thermal storage systems by controlling the heating, ventilation and air conditioning (HVAC) of large commercial buildings. BESSs are very well suited for the provision of services to the grid since they are highly controllable devices that exhibit very fast ramp rates [105]. In a landscape where the overall rotational inertia of the grid is decreasing, having such fast-responding Ancillary Services Providers (ASPs) could help reducing frequency deviations and, thus, better stabilizing the operation of the electric grid. However, when providing grid services, BESSs can be required to absorb significant energy biases over time which makes the management of the State of Charge (SoC) a challenging task.

Recent research has also shown the potential of the demand-side as a provider of services to the grid. These services are commonly referred to as Demand-Response (DR) programs and they require the load-side to vary its energy consumption profiles to meet the current needs of the grid. In particular, buildings represent a tremendous source of untapped energy storage. This is due to the fact that they are responsible for 40% of total energy consumption worldwide [29]; furthermore, buildings are inherently characterized by a large thermal inertia and this allows one to think of them as *virtual* storage devices capable of decoupling production from consumption. In particular, commercial buildings are good candidates for providing services to the grid for the following reasons: 1) they are typically characterized by a large HVAC system with respect to residential buildings. This

corresponds to larger energy consumption per unit which in turn means a smaller cost of acquisition. 2) Most commercial buildings are already equipped with Building Energy Management Systems (BEMS) that can be readily used to monitor and control the operation of their HVAC systems [87]. All these factors make commercial buildings an excellent target for providing flexibility to the grid. At the same time, however, there are also a number of challenges related to it. First of all, the primary objective of any controller should be to maintain occupant comfort which is typically specified in terms of a temperature range. This should be guaranteed against a wide range of disturbances that typically display periodic cycles such as occupancy, outside temperature, and solar radiation. Second, typical HVAC systems are quite complex in nature with multiple cascade control loops, self-correcting behaviors, and physical limitations of the equipment. For this reason, their power consumption cannot, in general, be modulated reliably and at very high frequencies.

Motivated by the aforementioned reasons, in recent years, the research community has started to explore the technical and economic feasibility of providing grid services with distributed energy resources such as BESSs and commercial buildings. This thesis follows these directions by examining the following research questions:

1. How to efficiently control resources, such as buildings, that are subject to external perturbations which are characterized by periodic patterns?

2. How can one formally determine the amount of services that an aggregation of distributed energy resources can offer to the power grid?

3. Is it possible to demonstrate in practice the feasibility of providing such services, and which resources should be aggregated in order to maximize their full exploitation?

4. When providing ancillary services, is there a noticeable deterioration to the primal purpose that these resources should individually fulfill?

We report the research that was conducted, over the course of the Ph.D., in order to address all the questions above. The thesis is structured in three main parts, and the main contributions of all the following chapters are summarized below.

Part I - MPC for distributed energy resources

Chapter 2 - Preliminaries

This part provides the foundations for the remainder of the thesis. First, it provides some background and mathematical material that is relevant for the comprehension of the following chapters. Second, it reviews the current practices for the control of energy resources. Finally, it introduces the formulation and main components of Model Predictive Control (MPC) giving an example for the building climate control case.

Chapter 3 - Stochastic MPC for Controlling the Average Constraint Violation of Periodic Linear Systems with Additive Disturbances

This chapter deals with the problem of enforcing the satisfaction of state constraints for general linear systems subject to periodically time-varying uncertainty. In particular, we consider additive disturbances that are bounded and described by periodically time-dependent probabilistic distributions. The aim of the chapter is to develop a strategy that controls the over-time-average constraint violations. The effectiveness of the proposed algorithm is demonstrated through an extensive simulation study on a building climate control case.

The main novelties introduced in this part are:

- the generalization of the concept of periodic controlled invariant sets, available in the robust framework, to the stochastic case.

- the formulation of a receding horizon control scheme that enforces recursive feasibility for the closed-loop trajectories of a system subject to periodic disturbances.

The chapter is based on the following paper:

- L. Fabietti and C. N. Jones, "Stochastic MPC for controlling the average constraint violation of periodic linear systems with additive disturbances", 2016 American Control Conference (ACC), Boston, MA, 2016, pp. 5395-5400. doi: 10.1109/ACC.2016.7526515

Part II - Provision of grid services using distributed energy resources: Theory

Chapter 4 - Multi-service provision through coordination of energy-constrained resources

Our objective is to investigate the provision of multiple services to the grid using a set of heterogeneous energy resources. In particular, in this chapter, the problem of formally characterizing the amount of services that can be offered to the grid is tackled. The general planning problem is first introduced. It accounts for both local and network constraints while maximizing the amount of fast regulation that the resources can collectively offer. Then, two tractable reformulations, based on data-driven robust optimization methods, are proposed. The efficacy of the control framework is demonstrated in simulations.

The main novelties introduced in this part are:

- the formulation of the optimal day-ahead planning problem that determines the allocation of the available control power across multiple local and shared services. The problem is formulated as a multi-stage uncertain optimization problem and it differs from existing works in: i) it accounts for both the uncertainty across the network as well as the uncertainty of the tracking signal received from the TSO, ii) it explicitly considers power flow (PF) constraints, iii) to increase the offerable flexibility, it encodes directly the possibility to adjust the baseline power consumption in the intraday market.

- the formulation of two novel approximated solutions of the planning problem. The approximations are based on data-driven robust optimization techniques and rely on a multi-stage linearization of the PF constraints around a general point of operation.

- the formal investigation of the properties of the considered PF linearization which is used to explicitly express voltages and currents across the network as a function of the controllable active and reactive injections.

- the presentation of an extensive simulation study that demonstrates the capability of the proposed framework to adapt to different levels of uncertainties while guaranteeing the provision of both local and shared services.

The chapter is based on the following papers, and most of the text and content in Chapter 4 has appeared in these papers.

- L. Fabietti, E. Namor, P. Nahata, M. Paolone, C. N. Jones, "Multi-service provision through coordination of energy-constrained resources", Applied Energy (to be submitted), 2018.

- Tomasz T. Gorecki, Luca Fabietti, Faran A. Qureshi, Colin N. Jones, "Experimental demonstration of buildings providing frequency regulation services in the Swiss market", Energy and Buildings, Volume 144, 2017, Pages 229-240, ISSN 0378-7788, https://doi.org/10.1016/j.enbuild.2017.02.050.

Chapter 5 - Comparison of the two approximation schemes

This chapter complements chapter 4 by comparing the previously proposed approximation schemes both from a qualitative as well as from a quantitative perspective. In particular, advantages an disadvantages of each method are described and an extensive simulation and experimental analysis is conducted to build some intuition on which method is most appropriate for a given application.

The content of this chapter is mostly taken from the following paper:

- L. Fabietti, T. T. Gorecki, F. A. Qureshi, A. Bitlislioglu, I. Lymperopoulos and C. N. Jones, "Experimental Implementation of Frequency Regulation Services Using Commercial Buildings", in IEEE Transactions on Smart Grid, vol. 9, no. 3, pp. 1657-1666, May 2018. doi: 10.1109/TSG.2016.2597002

Part III - Provision of grid services using distributed energy resources: Experiments

This part demonstrates the technical feasibility of deploying the control methods developed in Part II by means of two extensive experimental campaigns. Specifically, we consider two campus-scale experimental setups that comprise both controllable resources, stochastic on-site renewable generation, and uncertain consumption.

Chapter 6 - Multi-Time Scale Coordination of Complementary Resources for the Provision of Ancillary Services

This chapter presents a predictive control scheme for coordinating a set of complementary resources for the provision of fast regulation services. In particular, a commercial building (slow resource) and a battery energy storage systems (fast resource) are combined to augment the flexibility that can be provided to the grid compared to the flexibility that any of these resources can provide individually. An extensive experimental campaign (9 consecutive days) is conducted using a large auditorium building and an electrical storage system on the EPFL campus. The control method of Chapter 4 is adapted to this specific case in order to formally compute the flexibility of the aggregated system. A multi-level control scheme based on data-driven robust optimization methods is also developed to coordinate the resources in real-time.

The key outcomes of the study are:

- By optimally splitting the regulation signal, each resource can track the components that best suit its technical capabilities. In particular, this enlarges the number of commercial buildings that could be considered for such an application irrespective of the HVAC system in place. Moreover, the energy requirements on the BESS can be reduced by tapping the abundant thermal storage of the building.

- Combining complementary resources (fast and slow) can significantly augment the flexibility that can be provided to the grid.

- Closed-loop performance can be greatly improved when exploiting the time-correlation of the regulation signal in order to predict its future realization.

- The occupants' comfort in the building is not jeopardized when providing ancillary services.

- Optimally combining such resources allows one to achieve excellent tracking performance. Thus, the experimental results, that are in full compliance with the current Swiss regulations, show the practical feasibility of the developed method.

The chapter is based on the following paper, and most of the text and content in Chapter 6 has already appeared in this paper.

- L. Fabietti, F. A. Qureshi, T. T. Gorecki, C. Salzmann, C. N. Jones, "Multi-time scale coordination of complementary resources for the provision of ancillary services", Applied Energy, Volume 229, 2018, Pages 1164-1180, ISSN 0306-2619, https://doi.org/10.1016/j.apenergy.2018.08.045.

Chapter 7 - Enhancing the dispatchability of distribution networks through utility-scale batteries and flexible demand

This chapter investigates the problem of dispatching the operation of an active distribution feeder which is comprised of PV generation, uncertain consumption, and two controllable resources. More precisely, the 20kV distribution feeder on the EPFL campus is considered. The experimental

platform, developed with other two Ph.D. students (Tomasz Gorecki, Emil Namor) and in collaboration with the DESL laboratory, comprises: five office uncontrollable buildings (350kWp), a roof-top photovoltaic installation (90kWp), a grid-connected electrical storage (720kVA-500kWh), and a fully-occupied multi-zone office building (45 kWp). The objective is to precisely track a power trajectory, called the *dispatch plan*, which is fixed the day before the beginning of operation. To this aim, the two-stage method described in Chapter 4 is adapted so as to compute a suitable dispatch plan, and closed-loop experiments are carried out for an extended period of time (12 full days of operation). A hierarchical real-time controller is designed to coordinate the BESS and the building so as to track the dispatch plan while maintaining comfort and respecting the operational constraints of the battery.

The main novelties of this chapter are:

- The adaptation of the general control method of Chapter 4 to a particular case of practical interest. In particular, the resulting day-ahead problem optimizes the operation of the feeder while allocating enough local reserves to absorb real-time deviations.

- The design of a hierarchical multi-timescale MPC controller to accurately track in real-time the committed dispatch plan by coordinating the BESS and the building.

- The experimental validation of the control method of Chapter 4 against a wide range of modeled and unmodeled uncertainties.

The content of this chapter is mostly taken from the following papers:

- L. Fabietti, T. T. Gorecki, E. Namor, F. Sossan, M. Paolone, C. N. Jones, "Enhancing the dispatchability of distribution networks through utility-scale batteries and flexible demand", Energy and Buildings, Volume 172, 2018, Pages 125-138, ISSN 0378-7788, https://doi.org/10.1016/j.enbuild.2018.04.056.

- L. Fabietti, T. T. Gorecki, E. Namor, F. Sossan, M. Paolone and C. N. Jones, "Dispatching active distribution networks through electrochemical storage systems and demand side management," 2017 IEEE Conference on Control Technology and Applications (CCTA), Mauna Lani, HI, 2017, pp. 1241-1247. doi: 10.1109/CCTA.2017.8062629

Additional Publications

The following paper was published during the Ph.D., but is not included in this thesis.

- R. Gupta, F. Sossan, E. Scolari, E. Namor, L. Fabietti, C. N. Jones, and M. Paolone (2018). An ADMM-Based Coordination and Control Strategy for PV and Storage to Dispatch Stochastic Prosumers: Theory and Experimental Validation. 2018 Power Systems Computation Conference (PSCC), 1-7.

This paper is inspired by the results provided in Chapter 7. In particular, it proposes a control framework to coordinate the operation of a curtailable PV facility and a battery system in order

to achieve the dispatchability of an active distribution feeder. A multi-time scale hierarchical controller, similar to the one of Chapter 7, is designed. The lower layer is represented by an MPC controller executed at 10 s resolution to achieve fine tuning of a given setpoint. The upper layer is a slower MPC controller coordinating the two resources. It runs at a 5 minutes resolution and it is solved in a distributed fashion. The control framework is validated experimentally using the same platform of Chapter 7 with the only difference being the replacement of the controllable building with a 13 kWp curtailable PV facility.

My contribution to the paper was the development of the hierarchical controller, the formulation of the high-level centralized MPC problem, and the reformulation of the latter in order to be solved using the alternating direction method of multipliers (ADMM). The tuning, deployment, and testing of the algorithm was instead conducted by the co-authors.

Part I

MPC for distributed energy resources

Chapter 2

Preliminaries

This chapter serves as a reference for the rest of the manuscript as the main building blocks used in the thesis are presented. It is divided in two parts, the first presents some mathematical preliminaries and the second describes MPC while keeping in mind the objective of applying it to control distributed energy resources.

2.1 Mathematical Preliminaries

2.1.1 Set Theory

Definition 1 (Convex set). A set $\mathcal{X} \subseteq \mathbb{R}^n$ is *convex* if by selecting any pair of points in $\mathcal{X}$ the segment that connects them lies entirely in $\mathcal{X}$, *i.e.*

$$x_1, \, x_2 \in \mathcal{X}, \quad \lambda x_1 + (1 - \lambda)x_2 \in \mathcal{X} \quad \text{for } \lambda \in [0, 1].$$

Definition 2 (ϵ-ball). The open n-dimensional ϵ-ball in $\mathbb{R}^n$ around a fixed point (center) $x_c \in \mathbb{R}^n$ is defined as the set

$$\mathcal{B}_\epsilon(x_c) := \{x \in \mathbb{R}^n \mid ||x - x_c|| < \epsilon\},$$

where the radius of the ball is $\epsilon > 0$ and $|| \cdot ||$ can be any norm (typically the Euclidean norm).

Definition 3 (Closed set). A set $\mathcal{X} \subseteq \mathbb{R}^n$ is *closed* if every point not belonging to $\mathcal{X}$ has a disjoint neighborhood from $\mathcal{X}$, *i.e.*

$$\forall x \notin \mathcal{X}, \, \exists \, \epsilon > 0 \quad \text{such that} \quad \mathcal{B}_\epsilon(x) \cap \mathcal{X} = \varnothing.$$

Definition 4 (Bounded set). A set $\mathcal{X} \subseteq \mathbb{R}^n$ is *bounded* if it is contained inside some ball $\mathcal{B}_r(\cdot)$ of finite radius r, *i.e.*

$$\exists \, r < \infty, \, x \in \mathbb{R}^n \quad \text{such that} \quad \mathcal{X} \subseteq \mathcal{B}_r(x)$$

Definition 5 (Compact set). A set $\mathcal{X} \subseteq \mathbb{R}^n$ is *compact* if it is closed and bounded.

2.1.2 Systems and Control Theory

Definition 6 (Discrete-time system).

$$
\begin{aligned}
x_{i+1} &= f(x_i, u_i, w_i) && \text{Non-linear time-invariant} \\
x_{i+1} &= Ax_i + Bu_i + Ew_i && \text{Linear time-invariant (LTI)}
\end{aligned}
\tag{2.1}
$$

with time step $i \in \mathbb{N}$, system state $x_i \in \mathbb{R}^{n_x}$, control input $u_i \in \mathbb{R}^{n_u}$, and disturbance $w_i \in \mathbb{R}^{n_w}$.

Definition 7 (Constrained system). A discrete-time system is a *constrained system* if the state and/or the control inputs are enforced to belong to the sets

$$
\begin{aligned}
u_i &\in \mathbb{U} \subset \mathbb{R}^{n_u} \\
x_i &\in \mathbb{X} \subset \mathbb{R}^{n_x}
\end{aligned}
\tag{2.2}
$$

where it is assumed that the constraint sets $\mathbb{U}$ and $\mathbb{X}$ are compact and both of them contain the origin.

Definition 8 (Disturbance set). For bounded disturbances, w, the *disturbance set* $\mathbb{W}$ is defined by

$$
w_i \in \mathbb{W} \subset \mathbb{R}^{n_w}
$$

where the set $\mathbb{W}$ is assumed to be compact and contain the origin.

Definition 9 (Positive invariant set). A set $\mathcal{X} \in \mathbb{R}^{n_x}$ is a *positive invariant* set for the autonomous system $x_{i+1} = f(x_i)$, if $f(x_i) \in \mathcal{X}$ for all $x_i \in \mathcal{X}$.

Definition 10 (Robust positive invariant set). A set $\mathcal{X} \in \mathbb{R}^{n_x}$ is a *robust positive invariant* set for the system $x_{i+1} = f(x_i, w_i)$, if $f(x_i, w_i) \in \mathcal{X}$ for all $x_i \in \mathcal{X}$ and for all $w_i \in \mathcal{W}$.

Definition 11 (Controlled invariant set). A set $\mathcal{X} \in \mathbb{R}^{n_x}$ is a *controlled invariant* set for the system $x_{i+1} = f(x_i, u_i)$, if $\exists u_i \in \mathbb{U}$ such that $f(x_i, u_i) \in \mathcal{X}$ for all $x_i \in \mathcal{X}$.

Definition 12 (Robust controlled invariant set). A set $\mathcal{X} \in \mathbb{R}^{n_x}$ is a *controlled invariant* set for the system $x_{i+1} = f(x_i, w_i, u_i)$, if $\exists u_i \in \mathbb{U}$ such that $f(x_i, w_i, u_i) \in \mathcal{X}$ for all $x_i \in \mathcal{X}$ and for all $w_i \in \mathcal{W}$.

Definition 13 (Maximum controlled invariant set). A set $\mathcal{X} \in \mathbb{R}^{n_x}$ is the *maximum controlled invariant* for the system $x_{i+1} = f(x_i, u_i)$, if $\mathcal{X}$ is a controlled invariant set as defined in Definition 11, and any other controlled invariant set, $\mathcal{S}$ is a subset of $\mathcal{X}$, *i.e.* $\mathcal{S} \subseteq \mathcal{X}$.

Definition 14 (Maximum robust controlled invariant set). A set $\mathcal{X} \in \mathbb{R}^{n_x}$ is the *maximum controlled invariant* for the system $x_{i+1} = f(x_i, u_i)$, if $\mathcal{X}$ is a robust controlled invariant set as defined in Definition 12, and any other robust controlled invariant set, $\mathcal{S}$ is a subset of $\mathcal{X}$, *i.e.* $\mathcal{S} \subseteq \mathcal{X}$.

2.2 Receding Horizon Control

This section provides a brief description of the main building blocks of an MPC controller. Sometimes also called receding horizon control, MPC is an optimization-based control strategy that has been extensively and successfully used in many areas to perform constrained control [91].

The main idea is the following, at every discrete sampling time i, a measurement of the current state of the system, x_i, is obtained. Then, the system evolution is computed, over a pre-defined prediction horizon, N, by propagating the initial state through a model such as (2.1). Finally, an optimization problem is solved that minimizes some optimality criterion given the system evolution and constraints (2.2). By solving such an optimization problem, a sequence of inputs and states over the prediction horizon is obtained. In particular, these sequences respect the system dynamics and constraints while minimizing the given optimality criterion. The optimal input sequence is given by

$$\mathbf{u}^*(x_0) := \{u_0^*, \ldots, u_{N-1}^*\}$$

According to the receding-horizon structure of the MPC controller, the first input u_0^* of the sequence is applied to the system whereas the rest is discarded. At the next time instant, the new state is measured and the whole process is repeated with the prediction horizon shifted by one time step. In particular, this receding-horizon structure of the controller is the key mechanism that introduces feedback into the control loop in order to compensate for modeling errors, external perturbations, etc.. A sketch of the resulting MPC algorithm is reported in the following

---
Algorithm 1 MPC controller algorithm

1: Obtain the current state of the system x_0
2: Obtain the optimal input sequence $\mathbf{u}^*(x_0)$ by solving an optimization problem over the prediction horizon, N.
3: Apply the first input, u_0^* of the sequence to the system
4: Repeat from 1 at the next time iteration.

Remark 1. *As described in this section, the basic structure of the MPC problem consists in optimizing directly over input sequences. Another possibility, that will be exploited later in this thesis, is to optimize over control laws, i.e., over functions of the system states or disturbances. To keep this introduction as simple as possible, we focus herein on input sequences. The reader is referred to Chapter 4 for the formulation using control laws.*

2.2.1 MPC for distributed energy resources

There is a vast literature that supports the efficacy of MPC for the optimization and control of complex energy systems. Possible applications range widely from industrial processes [42], wind turbines [77], fuel cells [151], residential and commercial building climate control [123], just to cite a few. Thus, in this part, rather than providing an in-depth literature review across all applications, we focus on describing the main motivations for using MPC for the control of flexible resources. While doing this, we also summarize the most important findings appearing throughout the literature. In particular, as it will be one of the most recurrent systems considered in this thesis, we will focus on the particular example of the development of MPC controllers for building control.

2.2.2   MPC for Building Control

Traditionally, the control of commercial building is obtained by means of PID controllers which are coordinated by a number of rule-based controllers. In most cases, the resulting overall controller can achieve good performance by means of an iterative tuning procedure which aims at occupants' comfort. However, this usually comes at the cost of high energy expenditure and commissioning. Moreover, these controllers are not suitable for the interaction with a smart grid in order to provide support services such as secondary frequency control, peak shaving or load shifting.

Recently, the common trend is to improve HVAC system performance using more advanced predictive controllers such as MPC that are inherently able to incorporate several pieces of information. Thus, many contributions have appeared that consider a variety of different objectives, such as total cost minimization [83, 100], peak power reduction [84, 103], optimal energy use [53], etc.. The conclusions of these studies is that significant performance improvements can be obtained if an appropriate model for the building is available, as well as reliable forecasts for weather conditions and occupancy.

In particular, in an important contribution [106] *Olderwurtel et al.* present an investigation into the energy saving potential of integrated room automation through a large-scale simulation study. In the paper, different control schemes are considered and compared for several simulation set-ups varying building type, HVAC system, and environmental perturbations. The main outcome of this study is that, in most cases, MPC outperforms industry-standard rule-based controllers. A particular class of MPC controllers, named Stochastic MPC (SMPC), has also been shown to achieve better performance with respect to its deterministic or robust counterpart. This is due to its ability to directly account, in the problem formulation, for uncertainty in environmental conditions and to formulate comfort constraints in a probabilistic fashion, allowing violations under certain situations.

In most SMPC schemes for HVAC control, in order to simplify the tractability of the problem formulation, the disturbances acting on the system are assumed to be extracted from a normal distribution. A different approach is taken in [110, 111], where a tractable-scenario-based approximation of a chance constraint problem is considered. Another scenario-based approach can be found in [146], where *Zhang et al.* propose an iterative bilinearization of the building model around nominal trajectories and sample occupancy pattern from a set of historical data.

The design and the real-world implementation of a predictive controller for a campus building connected to thermal storage are analyzed in [85, 86]. The resulting nonlinear MPC is shown to improve plant performance by 19% with respect to the previously installed baseline controller, and this percentage corresponds to a total of US$1,280 weekly saving. This is obtained by fully exploiting the active control of thermal storage operating the chilling system when economically convenient and, therefore, enhancing the efficiency of the chiller and cooling tower.

Summarizing the findings in the literature, the main advantages of predictive control, in the context of building control, are:

- The ability to exploit all the available information: MPC offers a very natural way to incorporate a wide range of factors also including future quantities such as weather conditions, electricity prices, occupancy.

- The capability of directly accounting for operational and comfort constraints.

- The possibility to specify and optimize over complex economic costs in an intuitive manner.

Optimization problem

Starting from the introduction given above, we provide herein an overview of the main ingredients used in MPC for buildings. We start from a standard MPC formulation

Problem 1 (General MPC Problem).

$$J(\mathbf{u}^*) := \underset{\mathbf{x},\mathbf{u}}{\text{minimize}} \quad J(\mathbf{u}) \tag{2.3}$$

$$x_{i+1} = f(x_i, u_i, w_i) \tag{2.4}$$

$$y_i = g(x_i) \tag{2.5}$$

$$(x_i, u_i) \in (\mathbb{X}, \mathbb{U}) \tag{2.6}$$

$$y_i \in \mathbb{Y} \tag{2.7}$$

$$x_0 = \hat{x} \tag{2.8}$$

$$i = 1, \ldots, N$$

The term (2.3) represents the cost function to be minimized. Equations (2.4) and (2.5) capture the dynamics of the system and the effect of external perturbations that may affect the system. Equations (2.6), (2.7) express the constraints on the states, inputs, and the outputs of the considered system. Equation (2.8) fixes the initial state of the system. Finally, N represents the numbers of steps in the future over which the system trajectories are optimized. To build some intuition, in the following sections, we provide possible examples of how each component can be chosen for the case of buildings.

Modeling

We focus on commercial buildings served by an HVAC system that is either fully or partially controllable through an Building Energy Management (BEMS) system. In such a scenario, it is quite common to divide the model into two components: the model for the thermodynamics of the building, and the model for the HVAC system.

The building thermodynamics can be typically characterized by a set of linear equations of the form:

$$x_{i+1} = Ax_i + Bu_i + Ew_i$$
$$y_i = Cx_i \tag{2.9}$$

where the output, y, is represented by the air temperature in different zones of the building, and the inputs are the control variable of the HVAC system such as massflow rates, supply temperatures, temperature setpoints, etc..

As buildings are affected by large disturbances that can have a significant impact on the internal dynamics of the temperature, it is of paramount importance to capture these dependencies in the model. This is reflected in the term w_i that can encompass a wide range of disturbances such as occupancy, outside temperature, sun irradiance, and internal gains.

The HVAC system consumes electrical energy, p^{hvac}, and produces the required thermal energy to maintain occupants' comfort in the building.

$$p_i^{\text{hvac}} = h_{\text{hvac}}(u_i, d_i) \tag{2.10}$$

where $h_{\text{hvac}}(\cdot)$ is, in general, a non-linear static map that captures the dependence of the electric consumption on the input commands and the external perturbations.

It is quite well-known that obtaining models such as (2.9) and (2.10) represents the most time-consuming and challenging task in the design of MPC controllers for buildings. Thus, this has been the topic of an intense research activity and many different approaches have been proposed [2, 74]. Broadly speaking, these approaches can be divided in three categories - first principle models (white-box), data-driven (black-box), and a combination of the previous two (grey-box). Each of them has its own advantages and disadvantages and the ultimate decision on which method to consider should depend on the particular case at hand. In this thesis, the model has either been obtained by means of available modeling software [50], or, as in the experimental part, has been identified using data-driven approaches on available input-output data.

Current state

As detailed in (2.8), the dynamics of the system are initialized with the current state which can be either directly measurable or it can be estimated by means of a global observer relying on a state estimation technique (*e.g.* weighted least squares, Discrete Kalman filter, etc.).

Constraints

As already mentioned, one of the main reasons to consider an advanced control technique such as MPC, is its inherent ability to handle constraints on the inputs and states of the system. In the case of a building, comfort constraints on the indoor air temperature of the building zones can be seamlessly included by specifying a temperature range $[T_{ref} - \beta_i, \ T_{ref} + \beta_i]$ within which the indoor

Minimum Energy	$J(\mathbf{u}) = \sum_{i=0}^{N} e_i$
Minimum Cost	$J(\mathbf{u}) = \sum_{i=0}^{N} c_i^{\mathrm{ener}} e_i$
Peak charge	$J(\mathbf{u}) = c^{\mathrm{peak}} \max_{i \in [T_{\mathrm{start}},\, T_{\mathrm{end}}]} p_i$

Table 2.1 – Cost function in typical MPC control problems for buildings.

temperature should be maintained. The term T_{ref} is the ideal temperature to be tracked and β_i is a user-defined parameter which determines the maximum allowed deviation from T_{ref}. Please notice how β is allowed to be time-varying (typically periodic) as temperature constraints are often times relaxed during non-working/unoccupied hours. In the same manner, operational constraints are considered to guarantee the safe operation of the HVAC system. Possible examples are represented by maximum limit on the electrical power drawn from available compressors, minimum running time to avoid significant wear of mechanical components, etc..

Objective function

The second advantage of MPC is represented by the possibility of directly considering various types of cost function which is chosen to fulfill two purposes:

- *Stability* which is the property of a system to respond to bounded inputs with bounded outputs. In general, to achieve this goal, the cost function is chosen so that the corresponding optimal cost function, $J(\mathbf{u}^*)$ is a Lyapunov function for the system (2.4). In particular, this guarantees the closed-loop stability of the system [91].

- *Performance* the cost captures a given optimality criterion that we wish to minimize over time. Multiple objectives are also possible and the cost should be designed to directly express the relative importance of these different objectives.

When dealing with buildings, being an inherently stable system, the stability requirement is typically relaxed, and, as a consequence, the focus is purely on performance. One example of such cost is, for instance, the minimization of energy use:

$$J(\mathbf{u}) = \sum_{i=0}^{N} e_i$$

with e_i representing the energy use of the resource, *e.g.* energy consumption of an HVAC system over a one-day period. Similarly, the minimum cost is expressed as:

$$J(\mathbf{u}) = \sum_{i=0}^{N} c_i^{\mathrm{ener}} e_i$$

with c^{ener} the cost of energy that is, in general, time-varying. In particular, depending on the type of resource considered the price of energy could vary according to a fixed schedule or continuously.

An example of fixed schedule is represented by on-peak/off-peak tariffs that are typically applied to commercial buildings. In this situation a higher price is fixed during hours with overall high consumption whereas a lower price is set during non-peak hours. Hence, while in the static case the cost function can be readily formed using the pre-determined tariffs, in the case of dynamic prices, one would need to forecast them before setting up the MPC problem.

Another common objective is to minimize the peak demand over a specified period of time:

$$J(\mathbf{u}) = c^{\text{peak}} \max_{i \in [T_{\text{start}}, \, T_{\text{end}}]} p_i \tag{2.11}$$

where c^{peak} is a fixed cost associated to the peak power consumption over the time period $[T_{\text{start}}, \, T_{\text{end}}]$.

Finally, it is quite common to consider multi-objective cost functions. For instance, if the HVAC system of a building is controlled to reduce the peak charges, besides considering the economic metric (2.11) one could also include a term penalizing large deviations from the reference temperature, T_{ref}. In such situations the cost function would be of the form:

$$J(\mathbf{u}) = J^{\text{econ}}(\mathbf{u}) + \rho(\mathbf{u})$$

where $J^{\text{econ}}(\mathbf{u})$ represents one of the economic costs previously introduced, and the term $\rho(\mathbf{u})$ is a penalization term.

2.2.3   Additional preliminaries for MPC

We conclude this chapter by providing some additional definitions that will be essential for the comprehension of the following chapters.

Definition 15 (Admissible). A control input, sequence, or control law is called *admissible* if it satisfies the input constraints.

Definition 16 (Feasible set). The feasibility set $\mathcal{X}_0$ is defined as the set of initial states $x_0 \in \mathbb{R}^n$ for which the MPC Problem (1) with horizon N is feasible, *i.e.*.

$$\mathcal{X}_0 := \{\hat{x} \in \mathcal{R}^{n_x} \mid \exists (u_0, \ldots, u_{N-1}) \text{ s.t. } (x_i, u_i) \in \mathbb{X} \times \mathbb{U}, \; y_i \in \mathbb{Y}, \; i = 1, \ldots, N\} \; .$$

Definition 17 (Recursive feasibility). The MPC Problem 1 is called *recursively feasible*, if for all initial states feasibility is guaranteed at every state along the closed-loop trajectory.

Chapter 3

Stochastic MPC for Controlling the Average Constraint Violation of Periodic Linear Systems with Additive Disturbances

3.1   Introduction

As underlined in the previous chapter, when dealing with the control of complex distributed energy resources such as buildings, SMPC has been shown to be the best option with respect to standard rule-based controllers or robust controllers. This is due to the fact that SMPC is designed to directly account for the stochastic nature of environmental disturbance and it is, therefore, capable of reducing the conservatism of robust control schemes. Moreover, in many situations, constraint specification is, by nature, probabilistic so that it is reasonable to allow violations to occur with a certain frequency or when the amount by which the constraint is violated is small.

Standard methods for SMPC consider point-wise in time constraint specifications such as expectation, probabilistic or integrated chance constraints. These constraints are then typically enforced by implementing a mixed stochastic-worst-case tightening procedure [23, 24, 69]. However, these approaches do not consider past trajectories of the state process and this generally results in a conservative formulation. A different approach was first proposed in [67] and then extended in [68]. The main idea is to reduce the conservatism of previous methods by looking at the whole history of the state trajectory. Hence, rather than controlling the amount of constraint violation at each time separately, the quantity of interest is the average over time of constraint violation.

All the aforementioned approaches focused on LTI systems subject to time-invariant disturbances which is a quite restrictive assumption in many applications. Periodic linear systems offer a useful generalization of time-invariant systems providing a natural framework for modeling various

phenomena [11]. A relevant example of this is represented by building climate control where the system is subject to time-varying environmental perturbations that typically present periodic, seasonal/daily patterns. MPC of linear/nonlinear periodic systems has been tackled in many contributions [15, 62, 73]. The case of linear periodic systems subject to additive uncertainties has been considered, yet in a robust setting, in [45]. Thus, the control of periodic systems has not been addressed in an SMPC framework which was one of the main motivations of this research.

The contribution, and novelty of this chapter is twofold. First we generalize the concept of periodic controlled invariant sets, available in the robust framework [45], to the stochastic case. Second we provide the extension of the least-conservative approach of [68] to the powerful and more general class of discrete-time periodic linear systems with periodically time-varying system dimensions and subject to additive time-varying disturbances. Third, we demonstrate in simulation the effectiveness of the control strategy on a building climate control case.

3.2 Problem Statement

First, we provide an extension of the general linear time-invariant state-space model defined in Chapter 2 to the general class of periodic systems:

Definition 18. A discrete time linear periodic system is defined by

$$
\begin{aligned}
x_{i+1} &= A_{\sigma_i} x_i + B_{\sigma_i} u_i + w_i \\
\sigma_i &:= \mathrm{mod}(i, p) \quad \sigma_i \; : \; \mathbb{N} \to \mathbb{N}_0^{p-1}
\end{aligned}
\tag{3.1}
$$

with system time step $i \in \mathbb{N}$, period length $p \in \mathbb{N}_+$, intra-period step index function $\sigma(\cdot)$, state $x_i \in \mathbb{R}^{n_{\sigma_i}}$, input $u_i \in \mathbb{R}^{m_{\sigma_i}}$, disturbance $w_i \in \mathbb{R}^{l_{\sigma_i}}$ and $A_{\sigma_i} \in \mathbb{R}^{n_{\sigma_{i+1}} \times n_{\sigma_i}}$, $B_{\sigma_i} \in \mathbb{R}^{n_{\sigma_{i+1}} \times m_{\sigma_i}}$.

Remark 2. *The dimensions of state, input and disturbance vectors are allowed to be periodically time-dependent and have to satisfy $n_j, m_j, l_j \in \mathbb{N}_+ \forall j \in \mathbb{N}_0^{p-1}$. The possibility to consider systems characterized by time-varying dimensions can be useful in many practical situations, e.g., to model systems with asynchronous control inputs [45]. Another evidence of this fact is reported in Section 3.4.2.*

We assume that the state of the system is perfectly known at time i. The inputs are subject to hard constraints of the form:

$$
u_i \in \mathbb{U}_{\sigma_i}, \quad i \in \mathbb{N}_+
\tag{3.2}
$$

At each time step k, we assume the support of w_i to belong to a compact polyhedron $\mathbb{W}_{\sigma_i}$. Finally the closed-loop state of the process is required to satisfy, in a probabilistic fashion, the following constraint:

$$
g_{\sigma_i}^T x_i \leq h_{\sigma_i}
\tag{3.3}
$$

where $g_{\sigma_i} \in \mathbb{R}^{n_{\sigma_i}}$ and $h_{\sigma_i} \in \mathbb{R}$.

Remark 3. *To simplify the development, a single linear constraint is considered in the following. We highlight, however, that more complex constraints can be easily considered using the same methodology [68].*

In order to quantify the amount of violation occurring at time k we introduce the concept of a *loss function*.

Definition 19. A function $l : \mathbb{R} \rightarrow \mathbb{R}$ is a loss function if it is non-decreasing and lower-semicontinuous and it is zero at the origin.

In loose terms, one wants the state to remain in the half space defined by (3.3) "most of the time" or alternately not to exceed it "very much". Note that the *state constraint* can be time-varying as well.

A possible way to formalize this requirement is to impose *point-wise in time* constraints such as

$$\mathbb{E}\{l(g_{\sigma_i}^T x_i - h_{\sigma_i})\} \leq \xi, \quad i \in \mathbb{N}_+ \tag{3.4}$$

where $\xi \geq 0$. In general, due to the fact that constraints of this type are hard to deal with, one wants to enforce (3.4) by means of the *one-step* conditional constraint.

$$\mathbb{E}\{l(g_{\sigma_{i+1}}^T x_{i+1} - h_{\sigma_{i+1}})|x_i)\} \leq \xi, \quad i \in \mathbb{N}_+ \tag{3.5}$$

However, constraint (3.5) is, in general, very conservative which reduces the potential benefit of the probabilistic constraint specification. Instead of focusing on point-wise in time probabilistic constraints such as (3.4), we enforce constraints directly on the closed loop state process as a whole.

To this aim, we first introduce the weighted cumulative loss up to time i

$$v_i := \sum_{k=0}^{i} \gamma^{i-k} l(g_{\sigma_k}^T x_k - h_{\sigma_k}) \quad i \in \mathbb{N}_+ \tag{3.6}$$

where $\gamma \in [0,1]$ determines the forgetting rate of past losses.

Secondly, we define the normalization factor s_i

$$s_i := \sum_{k=0}^{i} \gamma^{i-k} = \begin{cases} \frac{1-\gamma^{i+1}}{1-\gamma} & \gamma \in [0,1) \\ i+1 & \gamma = 1 \end{cases} \tag{3.7}$$

Hence, the ratio v_i/s_i represents the weighted average loss up to time i and it is the main quantity

of interest of this chapter. In particular we require that:

$$
\begin{cases}
\mathbb{E}_i\{\dfrac{v_{i+1}}{s_{i+1}}\} \leq \xi & \text{if} \quad \dfrac{v_i}{s_i} \leq \xi & (3.8\text{a})\\[3ex]
\lim_{k \mapsto \infty} \dfrac{v_{\min(i+k,\tau_i)}}{s_{\min(i+k,\tau_i)}} \leq \xi & \text{if} \quad \dfrac{v_i}{s_i} > \xi & (3.8\text{b})
\end{cases}
$$

where the integer number τ_i is the first time of return. More precisely, if at time i the amount of violation exceeds the maximum allowed value, *i.e.* $v_i > \xi s_i$, τ_i represents the first time after i when the condition $v_{\tau_i} \leq \xi s_{\tau_i}$ is satisfied again, that is:

$$
\tau_i := \inf\{k \geq i \mid v_k/s_k \leq \xi\} \in \{i, i+1, \ldots\} \tag{3.9}
$$

3.3   Main Result

In this section a recursively feasible receding horizon control policy that enforces the constraint (3.8) for the closed-loop state process is presented. The main idea is to apply feedback on the ratio v_i/s_i acting on the *one-step* conditional constraint (3.5). If the quantity v_i/s_i is "small" we will loosen it, on the other hand, if it is "large" we will enforce (3.5) as it is.

As a first step, it can be observed that, by exploiting the state dynamic (3.1), the conditional constraint (3.5) can be written as

$$
\mathbb{E}\{l(g_{\sigma_{i+1}}^T(A_{\sigma_i}x_i + B_{\sigma_i}u_i + w_i) - h_{\sigma_{i+1}})\} \leq \xi \tag{3.10}
$$

The evaluation of (3.10) is in general difficult, hence to obtain a sufficient condition for its satisfaction, we first observe that the constraint (3.10) can be arranged as

$$
\mathbb{E}\{l(\underbrace{g_{\sigma_{i+1}}^T(A_{\sigma_i}x_i + B_{\sigma_i}u_i) - h_{\sigma_{i+1}}}_{\mu} + g_{\sigma_{i+1}}^T w_i)\} \leq \xi
$$

now, considering the function

$$
f_{\sigma_i}(\mu) := \mathbb{E}\{l(\mu + g_{\sigma_{i+1}}^T w_i)\} \tag{3.11}
$$

one immediately recognizes that it can be written as

$$
f_{\sigma_i}(\mu) := \int_{-\infty}^{\infty} l(\mu + y)\text{pdf}_{g_{\sigma_{i+1}}^T w_i}(y)dy
$$

Finally, the inequality (3.5) is satisfied if and only if

$$
g_{\sigma_{i+1}}^T(A_{\sigma_i}x_i + B_{\sigma_i}u_i) \leq q_{\sigma_i}(\xi) + h_{\sigma_{i+1}} \tag{3.12}
$$

where $q_{\sigma_i} : \mathbb{R} \to \mathbb{R}$ is defined as

$$
q_{\sigma_i}(\xi) := \sup\{\mu \in \mathbb{R} \mid f_{\sigma_i}(\mu) \leq \xi\}
$$

and its existence is guaranteed by the assumptions on the loss function $l(\cdot)$.

In the following we will introduce key concepts that are useful for the understanding of the proposed algorithm.

Definition 20. The stochastic feasibility set $\mathcal{X}_j^s$ corresponding to the intra-period index $j = 0, \ldots, p-1$ is defined as

$$
\begin{aligned}
\mathcal{X}_j^s = \{x \in \mathbb{R}^{n_j} \; : \; \exists u \in \mathbb{U}_j \; | \\
g_{\sigma_{j+1}}^T (A_j x + B_j u) \le q_{\sigma_j}(\xi) + h_{\sigma_{j+1}}\}
\end{aligned}
\tag{3.13}
$$

Essentially, X_j^s represents the set of all states from which there exists an admissible input for the intra-period index j such that the state process will satisfy the constraint (3.5) for all the possible disturbance realizations contained in the current disturbance support set $\mathbb{W}_j$.

In many practical situations, besides the stochastic constraint (3.8), it is desirable that the amount of violation is constrained by a maximum admissible loss at each time iteration.

$$
l(g_{\sigma_i}^T x - h_{\sigma_i}) \le \bar{\xi}_{\sigma_i}
\tag{3.14}
$$

Typically the parameter $\bar{\xi}_{\sigma_i}$ is derived from the problem specification and can be, in general, time-varying as well. This requirement defines the *feasible set* $\bar{\mathcal{X}}_j$ for each $j = 0, \ldots p-1$

$$
\bar{\mathcal{X}}_j := \{x \in \mathbb{R}^{n_j} \,|\, g_{\sigma_j}^T x \le h_{\sigma_j} + l^{-1}(\bar{\xi}_{\sigma_j})\}
\tag{3.15}
$$

where

$$
l^{-1}(a) := \sup\{y \in \mathbb{R} | l(y) \le \xi\} \in [-\infty, \infty]
$$

The second important concept that needs to be introduced is that of *Stochastic Periodic Controlled Invariance sequence* (SPCI).

Definition 21. A collection of sets $(S_0, S_1, \ldots, S_{p-1})$, is an SPCI sequence if it satisfies for each $j = 0, 1 \ldots, p-1$ $S_j \subseteq \mathcal{X}_j^s \cap \bar{\mathcal{X}}_j$ and

$$
\begin{aligned}
\forall x \in S_j \quad \exists u \in \mathbb{U}_j : \; A_j x + B_j u + w \in S_{\sigma_{j+1}} \quad \forall w \in \mathbb{W}_j \\
\wedge \quad g_{\sigma_{j+1}}^T (A_j x + B_j u) \le q_j(\xi) + h_{\sigma_{j+1}}
\end{aligned}
\tag{3.16}
$$

In the following, the parameter that will adjust the loosening of the *one-step* conditional constraint (3.5) is presented. Starting from the first line of (3.6) one can observe that the expected value $\mathbb{E}_i\{v_{i+1}\}$ can be written as

$$
\mathbb{E}_i\{v_{i+1}\} = \gamma v_i + \mathbb{E}_i\{l(g_{\sigma_{i+1}}^T x_{i+1} - h_{\sigma_{i+1}})\}
\tag{3.17}
$$

The right side of the inequality (3.7), in turn

$$\xi s_{i+1} = \xi(\gamma s_i + 1)$$

Therefore (3.8a) will be satisfied at time $i+1$ if the following condition holds

$$\mathbb{E}_i\{l(g_{\sigma_{i+1}}^T x_{i+1} - h_{\sigma_{i+1}})\} \leq \gamma(\xi s_i - v_i) + \xi \tag{3.18}$$

Definition 22. The probability leeway $\beta_i \in [\xi, \infty)$ at time $i \in \mathbb{N}$ is

$$\beta_i := \max\{\gamma(\xi s_i - v_i) + \xi, \xi\}^1 \tag{3.19}$$

Essentially, we will exploit β_i to define a control policy which enforces the satisfaction of (3.8) as a whole. To this effect, at each time iteration, depending on the amount of previous constraint violation, we will enforce the constraint

$$\mathbb{E}\{l(g_{\sigma_{i+1}}^T x_{i+1} - h_{\sigma_{i+1}})|x_i\} \leq \beta_i \tag{3.20}$$

Remark 4. *Please note that when $\beta_i > \xi$, enforcing (3.20), guarantees the satisfaction of the first line of (3.8). On the contrary, this is not guaranteed whenever $\beta_i = \xi$. Despite this, it is still possible to define a control policy which achieves to enforce satisfaction of (3.8).*

To this aim, we introduce an auxiliary state that controls the definition of the parameter β

$$\chi_i = \xi s_i - v_i \tag{3.21}$$

For each $i \in \mathbb{N}_+$, we further define the set

$$\tilde{\mathcal{U}}_{\sigma_i}(x_i, \chi_i) := \{u \in \mathbb{U}_{\sigma_i} :$$
$$A_{\sigma_i} x_i + B_{\sigma_i} u + w \in S_{\sigma_{i+1}} \; \forall w \in \mathbb{W}_{\sigma_i} \tag{3.22a}$$
$$\mathbb{E}\{l(g_{\sigma_{i+1}}^T (A_{\sigma_i} x_i + B_{\sigma_i} u + w_{\sigma_i}) - h_{\sigma_{i+1}})\} \leq \beta_i\} \tag{3.22b}$$

A basic single-layer set-valued control policy that will lead to the satisfaction of (3.8) as a whole is then defined as

$$\tilde{\kappa}_{\sigma_i}(x_i, \chi_i) \in \tilde{\mathcal{U}}_{\sigma_i}(x_i, \chi_i), \quad i \in \mathbb{N}_+ \tag{3.23}$$

---

[1]Please note that the probabilistic leeway parameter β_i could be equivalently defined as

$$\beta_i := \begin{cases} \xi & if \quad \xi s_i < v_i \\ \gamma(\xi s_i - v_i) + \xi & if \quad \xi s_i \geq v_i \end{cases}$$

The proof of this statement is given in the following section where a multi-layer version of the control policy is presented.

3.3.1 Multi-layer version

As already underlined, enforcing (3.8) is, in general, less conservative with respect to standard point-wise probabilistic constraints such as (3.5). Still the invariance constraint (3.22a) is independent of the amount of past violations, therefore it reduces the potential benefits of loosening the constraint (3.22b). Hence, in the following, we will relax the condition (3.22a) allowing the state process to move within a sequence of nested sets built around the SPCI sequence whenever the number of past violations is "small" enough. More precisely let's first observe that, at time i and when $x_{i+1} \in \bar{\mathcal{X}}_{\sigma_{i+1}}$

$$
\begin{aligned}
s_{i+1}\xi - v_{i+1} &= s_{i+1}\xi - \gamma v_i - l(g_{\sigma_{i+1}}^T x_{i+1} - h_{\sigma_{i+1}}) \\
&= \gamma(s_i\xi - v_i) + \xi - l(g_{\sigma_{i+1}}^T x_{i+1} - h_{\sigma_{i+1}}) \\
&\geq \gamma(s_i\xi - v_i) + \xi - \bar{\xi}_{\sigma_{i+1}}
\end{aligned}
$$

Continuing, we obtain

$$
\begin{aligned}
s_{i+2}\xi - v_{i+2} &= s_{i+2}\xi - \gamma v_{i+1} - l(g_{\sigma_{i+2}}^T x_{i+2} - h_{\sigma_{i+2}}) \\
&= \gamma(s_{i+1}\xi - v_{i+1}) + \xi - l(g_{\sigma_{i+2}}^T x_{i+2} - h_{\sigma_{i+2}}) \\
&\geq \gamma^2(s_i\xi - v_i) + \gamma(\xi - \bar{\xi}_{\sigma_{i+1}}) + \xi - \bar{\xi}_{\sigma_{i+2}}
\end{aligned}
$$

the same argument can be used to obtain a condition k steps ahead

$$
s_{i+k}\xi - v_{i+k} \geq \gamma^k(s_i\xi - v_i) - \sum_{t=1}^{k} \gamma^{k-t}(\bar{\xi}_{\sigma_{i+t}} - \xi)
$$

Consequently, if $x_{i+t} \in \bar{\mathcal{X}}_{\sigma_{i+t}}$ for each $t \in \{1, \ldots, k\}$, the requirement $s_{i+k}\xi - v_{i+k} \geq 0$ is satisfied if

$$
(s_i\xi - v_i) \geq \sum_{t=1}^{k} \gamma^{-t}(\bar{\xi}_{\sigma_{i+t}} - \xi) \tag{3.24}
$$

In particular, if at time i, condition (3.24) is met, we are guaranteed that $v_{i+k} \leq s_{i+k}\xi$ is satisfied without the imposition of any constraints besides $x_{i+t} \in \bar{\mathcal{X}}_{\sigma_{i+t}}$ for $t \in \{1, \ldots, k\}$. This opens the possibility of letting the system state temporarily leave the SPCI sequence but still making sure that it will return at time $i + k$.

To this end we define the concept of pre-set

$$
\begin{aligned}
\text{Pre}(\mathcal{M}_{j+1}) =: \{x \in \mathbb{R}^{n_j} : \ &\exists u \in \mathbb{U}_{\sigma_j} \ | \\
&A_j x + B_j u + w \in \mathcal{M}_j \quad \forall w \in \mathbb{W}_{j+1}\}
\end{aligned} \tag{3.25}
$$

The sequence of nested family of length n_s is then obtained through

$$S_j^1 := S_j, \qquad\qquad \forall j = 0, 1, \ldots, p-1 \qquad\qquad (3.26a)$$

$$S_j^{k+1} := \text{Pre}(S_{j-1}^k), \qquad\qquad \forall j = 0, 1, \ldots, p-1 \qquad\qquad (3.26b)$$

$$\forall k = 2, \ldots, n_s - 1$$

Now let's assume the current time to be i and the state of the system to belong to the set S_{σ_i}. If condition (3.24) holds then the state is free to move up to the set $S_{\sigma_{i+1}}^{k+1}$ from which we are guaranteed to get back to $S_{\sigma_{i+k+1}}^1$ at time $i+k+1$ where it might be necessary to enforce (3.22b). According to this argument, it is possible to introduce an index $r_i \in \mathbb{N}_1^{n_s}$ that determines to which layer the state is allowed to move. Define $\tilde{r}_i$ as
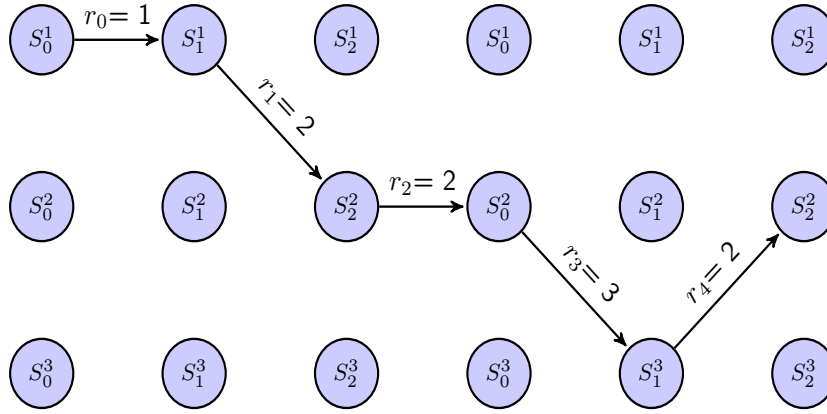
$$v_0 = 0, \qquad v_1 = 0, \qquad v_2 = 0, \qquad v_3 = 0, \qquad v_4 = 1,$$
$$r_0 = 1 \qquad r_1 = 2 \qquad r_2 = 2 \qquad r_3 = 3 \qquad r_4 = 2$$

Figure 3.1 – Possible evolution of the state over several time steps when starting from set S_0^1 at time $k = 0$, with p $= 3$, $n_s = 3$, $\gamma = 1$ and $l(\cdot) = \mathbb{I}[x > 0]$ which corresponds to the classic chance constraint formulation. Note how the state process is forced to move through the periodic invariant chain at each time iteration (from left to right) while allowed to climb up the family of nested sets when a low rate of past violations occurs (from top to bottom).

$$\tilde{r}_i := \max\{k \geq 1 \mid s_i \xi - v_i \geq \sum_{t=1}^{k} \gamma^{-t}(\bar{\xi}_{\sigma_{i+t}} - \xi)\} \qquad\qquad (3.27)$$

The index r_t is then defined by

$$r_i := \max\{\min\{\tilde{r}_i, n_s\}, 1\} \qquad\qquad (3.28)$$

Finally, we can introduce the multi-layer control policy.

To this aim, we define for all $i \in \mathbb{N}$ the sets

$$
\mathcal{U}_{\sigma_i}(x_i, \chi_i) := \{ u \in \mathbb{U}_j :
$$
$$
A_{\sigma_i} x_i + B_{\sigma_i} u + w \in S^{r_i}_{\sigma_{i+1}} \ \forall w \in \mathbb{W}_{\sigma_i}, \tag{3.29a}
$$
$$
\mathbb{E}\{ l(g^T_{\sigma_{i+1}}(x_{i+1}) - h_{\sigma_{i+1}}) \,|\, x_i \} \leq \beta_i \}. \tag{3.29b}
$$

$$
\Pi_{\sigma_i} := \{ (x_i, \chi_i) \,|\, \mathcal{U}_{\sigma_i}(x_i, \chi_i) \neq \varnothing \} \tag{3.30}
$$

The multi-layer control law is then defined as

$$
\kappa_{\sigma_i}(x_i, \chi_i) \in \mathcal{U}_{\sigma_i}(x_i, \chi_i), \quad k \in \mathbb{N} \tag{3.31}
$$

We are finally ready to state the following theorem:

Theorem 1. *Under the control law $u_i = \kappa_i(x_i.\chi_i)$ the following holds:*

(i) *If $x_0 \in S_0$ then $(x_0, \chi_0) \in \Pi_0$ (initial feasibility)*

(ii) *If $(x_i, \chi_i) \in \Pi_i$ then $(x_{i+1}, \chi_{i+1}) \in \Pi_{i+1}$ (recursive feasibility)*

(iii) *If $(x_0, \chi_0) \in \Pi_0$ then x_i satisfies the constraint (3.8) (closed-loop satisfaction)*

The proof of theorem 1 can be found in the Appendix A.

3.4   Implementation

In this section, we describe the implementation in an MPC framework of the theory previously presented. The general problem formulation assumes the form

$$
\min\{ J_{\sigma_i} \,|\, u_i \in \mathcal{U}_{\sigma_i}(x_i, \chi_i) \}
$$

where the cost function J_{σ_i} is completely arbitrary as well as the policy parametrization and the prediction horizon.

3.4.1   SPCI Parametrization

Explicit computation of the SPCI sequence

One possibility to parametrize the SPCI sequence, is to explicitly construct the sets by adapting the original algorithm [12] to the periodic case. To this end, we first present an adaptation of the pre-set operator that will be useful for later derivations.

Definition 23. Given a set $\mathcal{M}_{j+1} \subseteq \mathbb{R}^{n_{j+1}}$ at intra-period $j+1$, the set $\mathrm{Pre}^s(\mathcal{M}_{j+1})$ is defined as

$$
\mathrm{Pre}^s(\mathcal{M}_{j+1}) =: \{ x \in \mathbb{R}^{n_j} : \ \exists u \in \mathbb{U}_j \,|
$$
$$
A_j x + B_j u + w \in \mathcal{M}_{j+1} \quad \forall w \in \mathbb{W}_j
$$
$$
g^T_{\sigma_{j+1}}(A_j x + B_j u) \leq q_j(\xi) + h_{\sigma_{j+1}} \}
$$

Algorithm 2 SPCI Computation

1: Initialize the sets $S_j := X_j^s \cap \bar{\mathcal{X}}_j$, $j = 0, 1, \ldots, p - 1$ and set $i := 0$
2: Let $h = \mathrm{mod}(i, p)$. Compute $Q(S_h) =: \mathrm{Pre}^s(S_h) \cap S_{\sigma_{h-1}}$.
3: **if** $i \leq -p$ and $S_{\sigma_{h-1}} = Q(S_h)$ **then** stop. The maximal SPCI sequence has been found.
4: **if** $Q(S_h)$ is empty **then** stop. The maximal SPCI sequence does not exist.
5: Update $S_{h-1} = Q(S_h)$
6: Set $i =: i - 1$, and **goto** Step 2.

Assuming that a polyhedral SPCI sequence has been successfully computed together with the nested family sets at each period j, the constraint $u_i \in \mathcal{U}_{\sigma_i}(x_i, \chi_i)$ can be enforced as follows.

Being a polyhedron, each set, $S_{\sigma_{i+1}}^{r_i}$ can be described as $S_{\sigma_{i+1}}^{r_i} := \{x \mid v_{\sigma_{i+1}, r_i}^T x \leq b_{\sigma_{i+1}, r_i}\}$. Now the two conditions, (3.29a) and (3.29b), can be enforced as:

$$v_{\sigma_{i+1}, r_i}^T (A_{\sigma_i} x_i + B_{\sigma_i} u_i) + \max_{w \in \mathbb{W}_{\sigma_i}} v_{\sigma_i, r_i}^T w \leq b_{\sigma_{i+1}, r_i}$$

$$g_{\sigma_{i+1}}^T (A_{\sigma_i} x_i + B_{\sigma_i} u_i) \leq h_{\sigma_{i+1}} + q_{\sigma_i}(\beta_i)$$

The maximum on the left-hand side of the first constraint can be computed offline as it solely depends on the value of a single random variable, w_i.

Implicit Parametrization of the SPCI sequence

The definition of the multi-layer control policy (3.29) requires the parametrization of the SPCI sequence and the family of nested sets $\mathcal{S}_{\sigma_{i+1}}^{r_i}$. As seen in the previous section, one possible approach to this end is the explicit parametrization of the maximal sequence. However, the explicit computation of invariant sets suffers the so-called "curse of dimensionality" problem so that it might be difficult to compute the SPCI explicitly in large dimensions or when a long system period is considered. In such cases, it is possible to obtain an implicit parametrization as explained in the following.

Consider an MPC problem formulation with prediction horizon N and current time i. The predictions for the control input u_{i+k} are provided by an explicit policy parametrization for all $k \in \mathbb{N}_0^{N-1}$ whereas, for $k \geq N$ a fixed controller is assumed. To give a possible instance, let

$$u_{i+k} = \pi_{\sigma_{i+k}}(x_i^{i+k}), \quad k \in \mathbb{N}_0^{N-1} \tag{3.32}$$

be the explicit control policy (*e.g.* affine feedback policy). Further assume the terminal regulator to be a state-feedback controller

$$u_{i+k} = \kappa_{\sigma_{i+k}}^f(x_{i+k}), \quad k \geq N.$$

The constraint satisfaction is enforced explicitly for (3.32) through constraints on the policy $\pi_{\sigma_{i+1}}$ and implicitly for $i \geq N$ by requiring that the state x_{i+N} lands in a predetermined sequence of

invariant sets associated to $\kappa^f_{\sigma_{i+k}}$. More precisely:

Definition 24. A collection of sets $(\mathcal{X}^f_0, \ldots, \mathcal{X}^f_{p-1})$ is a sequence of terminal robust periodic invariant sets if it satisfies for each $j = 0, \ldots, p-1$, $\mathcal{X}^f_j \subseteq \mathcal{X}^s_j \cap \bar{\mathcal{X}}_j$ and

$$\forall x \in \mathcal{X}^f_j \quad A_j x + B_j \kappa^f_j(x) + w \in \mathcal{X}^f_{\sigma_{j+1}} \quad \forall w \in \mathbb{W}_j$$
$$g^T_{\sigma_{j+1}}(A_j x + B_j \kappa^f_j(x)) \leq q_j(\xi) + h_{\sigma_{j+1}}$$
$$\kappa^f_j(x) \in \mathbb{U}_j.$$

Therefore, at time i, the constraint on the terminal state assumes the form $x_{i+N} \in \mathcal{X}^f_{\sigma_{i+N}}$ which implicitly ensures satisfaction of both (3.29a) and (3.29b) for all w_k where $k \geq N$. Regarding the constraint on the policy parametrization for $k \in \mathbb{N}^{N-1}_0$, the constraint (3.29b) is imposed as it is whereas the invariance constraint (3.29a) is enforced implicitly as follows. We observe that $x_{i+1} \in \mathcal{S}^{r_i}_{\sigma_{i+1}}$ is guaranteed if $x_{i+k} \in \bar{\mathcal{X}}_{\sigma_{i+k}}$ for $k \in \mathbb{N}_+$ and

$$\mathbb{E}\{l(g^T_{\sigma_{i+r_i+k}} x_{\sigma_{i+r_i+k}} - h_{\sigma_{i+r_i+k}}) \mid x_{i+r_i+k-1}, x_i, v_i\} \leq \xi$$

The previous constraint is enforced explicitly along the prediction horizon for a given (x_i, v_i), all $k \in \mathbb{N}_+$ and all the possible trajectories generated by all possible $w^{i+r_i+k-2}_i$ under the control policy π_{σ_i} and the terminal controller, κ_{σ_i}. For $k > N - r_i$ it is enforced implicitly by enforcing that $x_{i+N} \in \mathcal{X}^f_{\sigma_{i+N}}$. The same approach is considered for the standard robust constraint, $x_{i+k} \in \bar{\mathcal{X}}_{\sigma_{i+k}}$.

3.4.2 Deactivation of the counter

In many practical applications, especially when time-varying state constraints are considered, it is desirable to account for constraint violations only during a specific sub-period, $\bar{p}$, of the system period p. This is the case, for example, of building climate control where the comfort constraint on the air temperature is typically relaxed during non business hours. If the counter for the process v_i/s_i was kept active during night this would result in a large accumulation of non-violations which would lead to a large amount of violations at the beginning of the next working day. This aspect, which represents a limitation of the approach proposed in [68], perfectly suits the proposed extension to periodic systems with periodically time-dependent dimensions.

For the sake of understanding, we assume the system to be described by a LTI system of the form

$$x_{i+1} = A x_i + B u_i + D w_i$$

with $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$ and $w \in \mathbb{R}^l$. We consider a time-varying state constraint with period p and we want to control the trajectory of the process v_i/s_i on the sub-period $\bar{p} \leq p$. In the following, we describe how to modify the construction of the SPCI sequence in order to accomplish this task. For the other quantities of interest, (3.13) and (3.26), the procedure follows the same principles and it is not reported.

For each $j \in \mathbb{N}_0^{\bar{p}-2}$ the definition (3.16) does not change whereas the set $S_{\bar{p}-1}$ is defined as

$$\forall x \in S_{\bar{p}-1} \quad \exists u \in \hat{\mathbb{U}}_{\bar{p}-1} :$$
$$\hat{A}_{\bar{p}-1}x + \hat{B}_{\bar{p}-1}u + \hat{D}_{\bar{p}-1}w \in S_0 \;\; \forall w \in \hat{\mathbb{W}}_{\bar{p}-1}$$
$$\wedge \quad g_0^T(\hat{A}_{\bar{p}-1}x + \hat{B}_{\bar{p}-1}u) \le \hat{q}_{\bar{p}-1}(\xi) + h_0$$

where $u \in \mathbb{R}^{m \cdot (p-\bar{p}+1)}$, $w \in \mathbb{R}^{l \cdot (p-\bar{p}+1)}$, $\hat{A}_{\bar{p}-1} = A^{p-\bar{p}+1}$, $\hat{B}_{\bar{p}-1} = [A^{p-\bar{p}}B, \dots, B]$, $\hat{D}_{\bar{p}-1} = [A^{p-\bar{p}}D, \dots, D]$, $\hat{\mathbb{U}}_{\bar{p}-1} = \mathbb{U}_{\bar{p}-1} \times \dots \times \mathbb{U}_{N-1}$, and $\hat{\mathbb{W}}_{\bar{p}-1} = \mathbb{W}_{\bar{p}-1} \times \dots \times \mathbb{W}_{N-1}$.

Hence, the resulting SPCI sequence $(S_0, \dots, S_{\bar{p}-1})$ is such that the controller does not update the cumulative loss v_i and the normalization factor s_i for each time k such that $\sigma_i \in \mathbb{N}_{\bar{p}}^{p-1}$.

3.5 Numerical Example

The system under analysis is a single zone for a commercial building modeled by a three state LTI model of the form

$$x_{i+1} = Ax_i + Bu_i + Dw_i$$

The model is an adaptation of the one described in [52] and discretized with a sampling period of 15 min which provides a nice compromise between temporal resolution of the control and computational complexity of the problem formulation. The control input, expressed in MW, is constrained to $\mathbb{U} = [0, 0.2]$. The states represent room air temperature, interior wall temperature and exterior wall temperature. The single comfort constraint for the first dimension of the state, $x_i^{(1)}$, takes the form $x_i^{(1)} \ge T_{\min}$ and is assumed to be time-varying. This reflects the fact that typically the building controller is asked to provide a comfortable work environment just during business hours. Therefore, $T_{\min}$ varies as follows

$$T_{\min} = \begin{cases} 21.5^o C & \text{from 8 am to 6 pm} \\ 18^o C & \text{otherwise} \end{cases}$$

The disturbance vector w_i (Figure 3.2) models environmental perturbations such as outside air temperature, solar radiation and internal heat sources. It consists of two terms $w_i = \delta_i + \epsilon_i$, where δ is deterministic and periodically time-dependent with a period of 24h whereas ϵ is stochastic, bounded and subject to periodically time-dependent bounds with the same period. The deterministic component represents the known fluctuation of the disturbances and it could be provided by, *e.g*, weather forecast and historical occupancy patterns. The random term ϵ models the uncertainty related to the prediction of the perturbation and is assumed to be distributed as a truncated normal random vector with zero mean and time varying variance. The truncation interval is periodically time-variant as well with period 24h. The selected cost function is the sum of control inputs over the simulation time which corresponds to the minimization of energy consumption. The number of nested sets to which the state is allowed to climb in case of low amount of past constraint violations is equal to $n_s = 4$ which has been found to be a good compromise between computational burden and performance. The SPCI sequence was determined explicitly for the two case studies in 34s on a 3.4 GHz Intel Core i7 processor. In the simulation, we study the performance of the controllers

Table 3.1 – Parameters for the control policies under analysis.

Policy	Probab0.95	Integ0.95	Integ1.0
loss function $l(x)$	$\mathbb{I}[x > 0]$	max $\{x, 0\}$	max $\{x, 0\}$
allowed violation	0.2	0.1	0.1
maximum violation	1	0.5^{o}C	0.5^{o}C
forgetting factor	0.95	0.95	1

resulting from different constraint specifications as summarized in Table (3.1). We compared them by means of 100 Monte Carlo simulations each of 5 full days long (424 steps). Table 3.2 and Figures 3.3 - 3.6 show how all the three stochastic specifications fully exploit the available flexibility in order to bring some relevant cost improvement with respect to the robust approach.
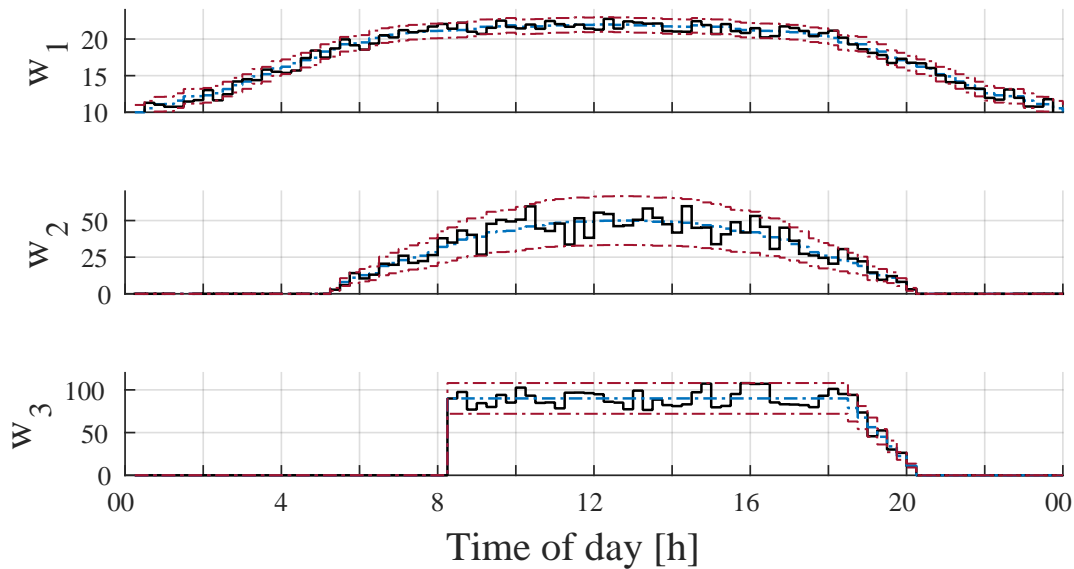
Figure 3.2 – Deterministic component δ (dashed blue), time-varying bounds on the stochastic term ϵ (dash-dotted red) and a possible realization (solid black).

Remark 5. *Note that, since we are interested in violations/non-violations occurring just during working hours, (3.13), (3.16) and (3.26) have been defined as described in Section 3.4.2 which shows the practical capabilities of the proposed method.*

Table 3.2 – Average cost improvement over the 100 runs of the control policies with respect to the robust controller obtained setting $\gamma = 0$ and converting the stochastic constraint (3.5) to its robust counterpart.

Policy	Probab0.95	Integ0.95	Integ1.0
Cost improvement	5.8%	15.3 %	17.2 %
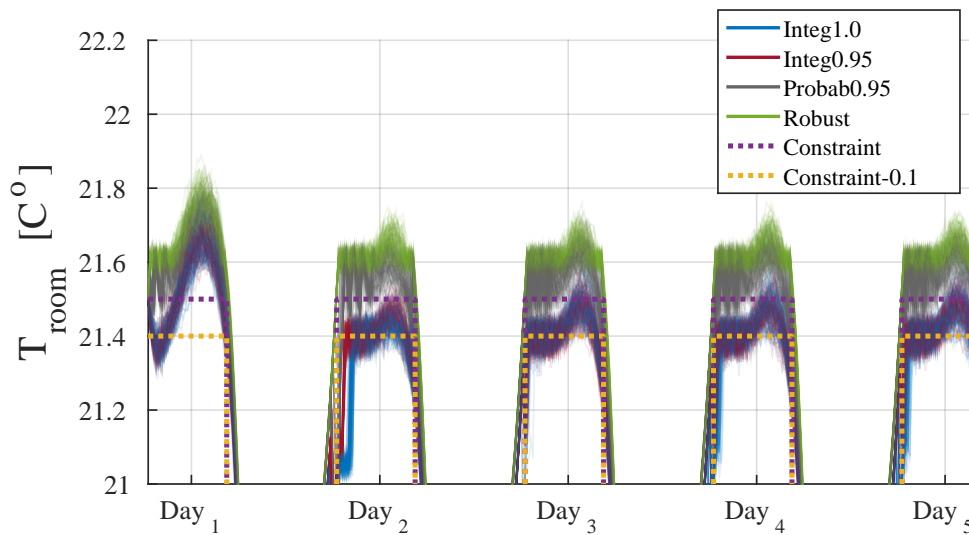
Figure 3.3 – One hundred Monte Carlo simulations for the controllers under analysis. Air temperature variation for each simulation. Note that, for all constraint specifications, most of the violations occur during the early business hours of the day. This phenomenon is due to the mean value of the environmental perturbation which tends to heat the room air temperature from 13pm to 17pm so that the controller accumulates large number of non-violations.
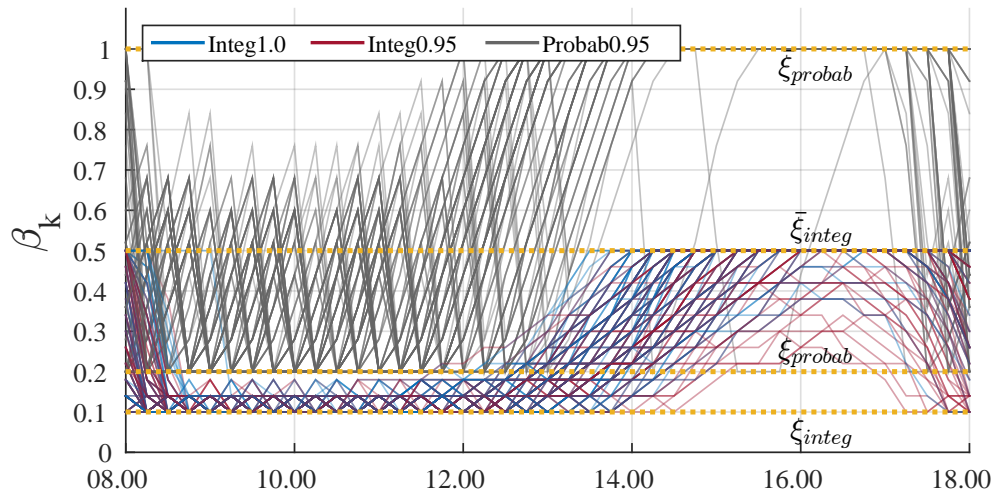
Figure 3.4 – One hundred Monte Carlo simulations for the controllers under analysis. The right side of the one-step conditional constraint β_i for one particular day of the simulation (Day 3). As observed in the state trajectories, most of violations occur in the early business hours of the day due to the mean value of the uncertainty.
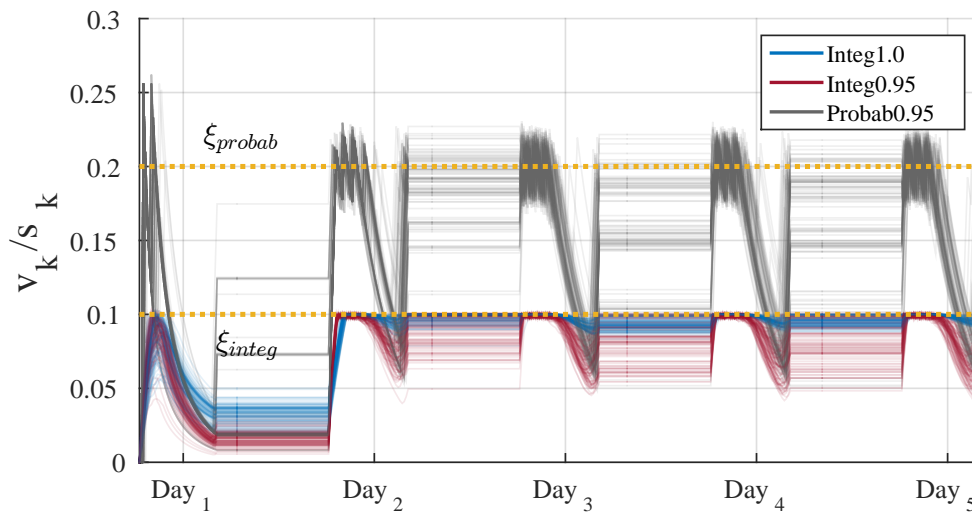
Figure 3.5 – One hundred Monte Carlo simulations for the controllers under analysis. The process v_i/s_i representing the average value over time of the loss function $l(\cdot)$. Please notice the deactivation of the loss counter during non-working hours.
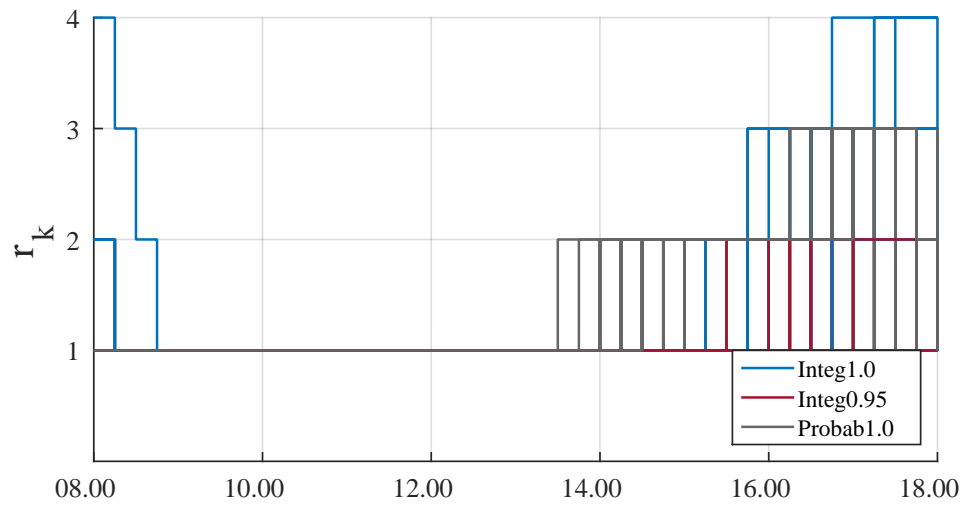
Figure 3.6 – One hundred Monte Carlo simulations for the controllers under analysis. The layer index r_i for one particular day of the simulation (Day 3).

3.6 Conclusions

In this chapter, we have presented a receding horizon control scheme that enforces recursive feasibility for the closed loop process of a periodic linear system when subject to stochastic constraints. The class of considered systems is wide and represents a powerful modeling tool for many real-life applications. This is true, in particular, since it allows to consider periodic inputs and states constraints as well as periodic disturbances that are characterized by time-varying probability distributions. The developed approach has been applied, in simulation, on a building temperature control case showing its flexibility and effectiveness with respect to robust MPC control schemes.

Appendices

Appendix A

Proof of Theorem 1

proof. (i) At time zero the layer index $r_0 = 1$ and $\beta_0 = \xi$. We need to show that $\mathcal{U}_0(x_0, \chi_0) \neq \varnothing$. But this is guaranteed by the condition $x_0 \in S_0$ and the definition of S_0.

(ii) At time i we assume $r_i = 1$. Hence, feasibility at time i implies that the state will land in $S^1_{\sigma_{i+1}}$ at time $i + 1$. To prove feasibility at time $i + 1$ one needs to show that the two constraints (3.29a), (3.29b) are satisfied for, at least, an admissible input $u_{i+1} \in \mathbb{U}_{j+1}$. Once again, this is ensured by the definition of $S^1_{\sigma_{i+1}}$. Note further that $r_i = 1$ is the only case when the second constraint (3.29b) is not redundant since $\beta_{i+1} < \bar{\xi}_{\sigma_{i+1}} \Rightarrow r_i = 1$.

For the case $r_i > 1$ we note that feasibility at time i implies that the state process is in $S^{r_i}_{\sigma_{i+1}} = \mathrm{Pre}(S^{r_i-1}_{\sigma_{i+2}}) \bigcap \bar{\mathcal{X}}_{\sigma_{i+2}}$ at the next time iteration. As we have assumed that the maximum violation is $\bar{\xi}_{\sigma_i}$, we have that r_{i+1} can decrease by at most one unit with respect to r_i, i.e., $r_{i+1} \geq r_i - 1$. Thus, we know that constraint (3.29a) will be satisfied at time $i + 1$ whereas constraint (3.29b) is redundant, as already underlined.

(iii) Let i be the current time instant and σ_i the correspondent intra-period time. First consider the case $v_i/s_i \leq \xi$ which can also be written as $\xi s_i - v_i \geq 0$. From the definition of β_i we have $\beta_i = \gamma(\xi s_i - v_i) + \xi$. Hence

$$\mathbb{E}_i(v_{i+1}) = \gamma v_i + \mathbb{E}\{l(g^T_{\sigma_{i+1}} x_{i+1} - h_{\sigma_{i+1}} | x_i)\}$$
$$\leq \gamma v_i + \gamma(\xi s_i - v_i) + \xi = \xi s_{i+1}$$

as required.

Consider now a time instant when $v_i/s_i > \xi$ and let τ_i be the first time of return under the threshold ξ. Obviously, whenever $\tau_i < \infty$ the second line of (3.8) is satisfied. In the case when $\tau_i = \infty$ we can define a new process

$$\eta_i := v_{i+k} - s_{i+k}\xi, \quad i \in \mathbb{N}$$

As a first thing, we show that η_i is a supermartingale

$$
\begin{aligned}
& \mathbb{E}_{i+k}\{\eta_{i+1} - \eta_i\} \\
&= \mathbb{E}_{i+k}\{(\eta_{i+1} - \eta_i)\} \\
&= \mathbb{E}_{i+k}\{(\gamma - 1)\eta_i + l(g^T_{\sigma_{i+k+1}} x_{i+k+1} - h_{\sigma_{i+k+1}}) - \xi)\} \\
&\leq \underbrace{(\gamma - 1)}_{\leq 0} \underbrace{\eta_i}_{\geq 0} \leq 0
\end{aligned}
$$

where the penultimate inequality derives from the fact that, when $\tau_i = \infty$, we have $\mathbb{E}\{l(g^T_{\sigma_{i+k+1}} x_{i+k+1} - h_{\sigma_{i+k+1}})\} \leq \xi$. Therefore η_i is a supermartingale process and for Doob's martingale convergence theorem it converges with probability one to some finite random variable η_∞.

To conclude the proof we need to demonstrate that, in the case $\tau_i = \infty$, we have $v_{i+i}/s_{i+i} \to \xi$. To this aim two cases need to be considered.

When $\gamma = 1$, $s_{i+k} = i + k \to \infty$ as $i \to \infty$ and

$$
\begin{aligned}
\lim_{k \to \infty} \frac{v_{i+k}}{i+k} - \xi &= \lim_{k \to \infty} \frac{v_{i+k} - (i+k)\xi}{i+k} \\
&= \lim_{k \to \infty} \frac{\eta_i}{i+k} = \lim_{k \to \infty} \frac{\eta_\infty}{i+k} = 0
\end{aligned}
$$

For the case $\gamma \in [0, 1)$ it is sufficient to prove that $\eta_i \to 0$. To this aim, given the convergence of η we just need to prove

$$
\mathbb{P}\left(\bigcup_{t=1}^{\infty} \left\{ \inf_{k \geq 0}(\eta_k) > 1/t \right\} \right) = 0
$$

which essentially requires that the process η_i crosses from above any positive level with probability 1. In other words, $\eta_i \to 0$.

Exploiting Boole's inequality, a sufficient condition for this to hold is

$$
\mathbb{P}(\tau(t) = \infty) = 0 \quad \forall t \in \mathbb{N}_+
$$

with $\tau(t) := \inf\{k \geq 0 \,|\, \eta_k \leq 1/t\}$. To show this we over-bound the trajectories of η_k by a random walk with a drift. Thanks to the assumptions on the loss function $l(\cdot)$, it is possible to write

$$
\begin{aligned}
l(g^T_{\sigma_{i+k+1}} x_{i+k+1} &- h_{\sigma_{i+k+1}}) \\
&\leq l(q_{\sigma_{i+k}}(\xi) + g^T_{\sigma_{i+k+1}} w_{\sigma_{i+k}})
\end{aligned}
$$

Therefore

$$
\begin{aligned}
\eta_{k+1} &= \gamma \eta_k + l(g^T_{\sigma_{i+k+1}} x_{i+i+1} - h_{\sigma_{i+i+1}}) - \xi \\
&\leq \gamma \eta_k + l(q_{\sigma_{i+k}}(\xi) + g^T_{\sigma_{i+k+1}} w_{\sigma_{i+k}}) - \xi
\end{aligned}
$$

which means that the trajectories of η_i are bounded by the AR(1) process of the form

$$X_{k+1} = \gamma X_k + z_k \quad X_0 = \eta_0 > 0$$

with $z_k := l(q_{\sigma_{i+k}}(\xi) + g^T_{\sigma_{i+k+1}} w_{\sigma_{i+k}}) - \xi$ that is an i.i.d innovation with non-positive mean. Moreover, if $\tau(t) = \infty$ we have $\eta_k > 1/t$ which means that the trajectories of X_k (and hence η_k) are over bounded by the random walk with a drift

$$Y_{k+1} = Y_k + z_k - (1 - \gamma)/t, \quad Y_0 = X_0 = \eta_0 > 0$$

The drift of this random walk, $\mathbb{E}\{z_k\} - (1 - \gamma)/t$, is strictly negative and bounded away from zero since $\mathbb{E}\{z_k\} \le 0$ and $(1 - \gamma)/t > 0$; this implies that the expected return time below $1/t$ is finite and, as a result, $\mathbb{E}\{\tau(t)\} < \infty$ which implies $\mathbb{P}(\tau(t) = \infty) = 0$. This finishes the proof.

Part II

Provision of grid services using distributed energy resources: Theory

Chapter 4

Provision of multiple grid services with distributed energy resources: Theory

4.1 Introduction

This section briefly summarizes the main actors and mechanisms underlying the operation of today's power systems. Despite the fact that the majority of the concepts introduced herein are common across different countries, we will tailor the description to the Swiss power grid as it will be the main focus of the third part of this thesis. The references [58, 64, 70] have been used for this section.

4.1.1 Traditional power system - overview

Electricity is the most versatile and easily controlled form of energy and, for this reason, it is considered an absolute necessity in modern society. However, with respect to other commodities, electricity cannot be efficiently and economically stored in large quantities and, as a consequence, it must be consumed at the same time is it generated. The transportation of electric power from power generators to consumers is achieved via the so-called *electric power systems*.

A power system, depicted in Figure 4.1, is composed of three main sub-parts: the generation, the transmission, and the distribution system. The *generation system* comprises generation units that supply power to the system. Some examples of traditional generators are thermal (oil, coal, natural gas, etc.), hydro, or nuclear power plants. Even though significant differences exist between them both in terms of controllability and responsiveness, the most important feature of these units is that they can be considered as being deterministic, i.e., they can be operated so as to output predefined power levels. The *transmission system* is composed of electrical lines that transport the power from the generation side to the distribution system. As high voltage AC transmissions provide a very efficient way to carry power with reduced losses, typical voltage levels range between 130 kV and 400kV. To guarantee the proper functioning of the network, a number of mechanisms and electrical devices are deployed that maintain voltage levels and phases (*e.g.* voltage-regulators, step-up/-down transformers, etc.). Beside voltage and phases, also the power balance and the frequency
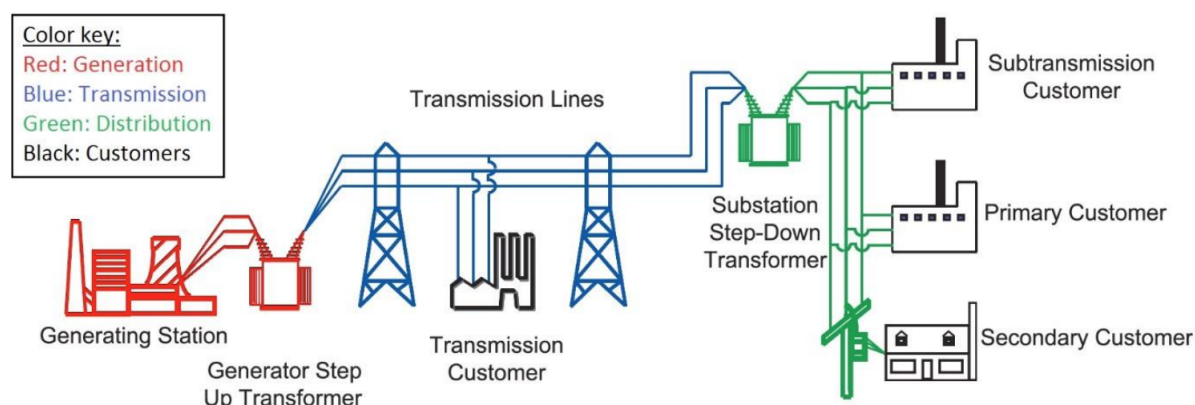
Figure 4.1 – Overview of a traditional electrical power system. Source: [58]

must be controlled at all times. In particular, this is achieved by specialized generation units that exhibit fast reaction times. Finally, the *distribution system* is in charge of feeding power to the loads such as households, hospitals, commercial buildings, and small to medium sized industries, among others. It has a similar structure to the transmission network but it works at lower voltage levels and it covers a significantly smaller geographical area.

In order to understand the operation of an electrical power grid, beside its physical structure, one should also consider another important component: the electricity market. The function of electricity markets is to facilitate the trading of electric power products between consumers and producers. Most electricity markets worldwide are deregulated meaning that competition between different actors is encouraged in order to increase the overall efficiency of the system. In Switzerland, electric energy is traded in the European integrated wholesale market, EPEX SPOT, that groups France, Germany, Britain, Switzerland, Austria, Belgium, the Netherlands, and Luxembourg. Every day, D, energy bids are collected until 12:00 of the current day for each hourly spot of the next day, D+1. Then, at market closure, following a clearing mechanism, the best bids are accepted and the final index of prices together with production and consumption schedules are communicated by the EPEX to each participant. At the beginning of the next day, D+1, these schedules come into action and should be respected by the participants. However, possible adjustments can be made, either at one hour or 15 minutes resolution and up to 45 minutes before delivery. In the central European area, such modifications take place in the form of energy trades, called *intraday transactions*, in the so-called *intraday market*.

All the previously described markets are energy markets. However, even after intraday adjustments, discrepancies between generation and consumption during real-time operation can arise and should be appropriately managed. Balancing imbalances at this faster time scale is the responsibility of Transmission System Operators (TSOs) that are typically state-owned monopolies in charge of guaranteeing the proper functioning of transmission systems. The Swiss TSO, SwissGrid, achieves this goal by procuring reserve generation capabilities, called Ancillary Services (AS), that function
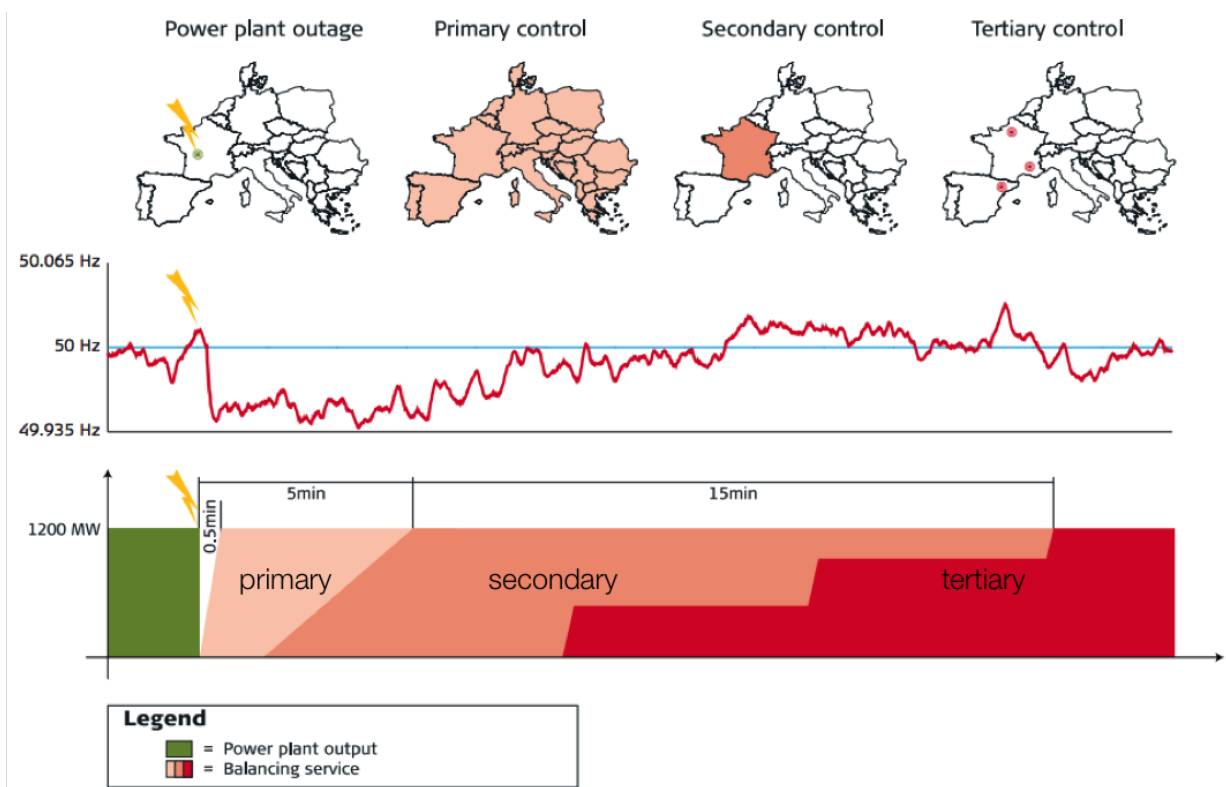
Figure 4.2 – Sequential activation of frequency control reserves for a power plant outage in France.
Source: www.swissgrid.ch

as a backup to cover the real-time mismatch between generation and consumption. Although
different TSOs might procure different types of AS, the following categories are generally present:

- Frequency control (active power reserves)

- Voltage control (reactive power reserves)

- Black start

- Compensation of active power losses

The technical specifications of such services are described in great detail in [119] while the economic
aspects are investigated in [120].

In this thesis the main focus is on frequency control reserves, which is the service specifically
designed to maintain the system frequency at its rated value (50 Hz in Europe). The frequency
is, in fact, a direct and instantaneous measure of the imbalance present in the network so that a
value above 50 Hz indicates a surplus of energy and a value below 50 Hz indicates a shortage of
energy. To make sure that grid contingencies are managed at various time scales, TSOs control the

frequency in three steps which are activated in a sequential fashion (Fig. 4.2): primary, secondary, and tertiary control.

Primary Frequency Control (PFC): It represents the fastest regulating layer after a contingency. For instance, in Switzerland an activation requirement of 30 seconds is imposed. Also, it is completely decentralized with each provider responding to local frequency measurements in a proportional fashion according to its droop characteristic [70]. Due to the proportional control and the lack of communication, PFC is employed to solely stabilize the frequency without removing steady-state frequency errors. In theory, PFC should be released after 5 minutes when the next layer is activated. For this reason, the energy requirements of PFC are quite small and, the providers of this service are compensated only for the provided capacity (and not for the energy).

Secondary Frequency Control (SFC): After the frequency is stabilized, and at most after 5 minutes, SFC is activated to restore the nominal frequency of 50 Hz. As opposed to the decentralized structure of PFC, SFC is centrally regulated by the TSO responsible for the control area where the contingency took place. For the Swiss control area, SwissGrid computes the so-called *area generation control* signal (AGC) which is the output of a Proportional-Integral (PI) controller with the area control error as input [64]. The signal is then broadcast every 2-4 seconds to the SFC providers which react accordingly by modulating their active power injection. If the AGC assumes a positive value, the provider should increase the power production and vice versa. In Switzerland, SFC reserves are contracted in a market setting where generators bid their capacity in weekly auctions. SwissGrid requires symmetric capacity meaning that the providers should be able to both increase and decrease their active power injection with respect to a pre-defined baseline injection.

Tertiary Frequency Control (TFC): It is the last layer to be activated in case of a major and persistent contingency. Its main purpose is to relieve SFC reserves after a period of 15 minutes. As of today, TFC is activated in a manual or semi-automatic fashion by special electronically transmitted messages sent by the TSO to large generating units that are required to adapt their power production levels.

4.1.2   Motivation

Due to the increasing connection of renewable energy resources, the physical stress on the power grid is increasing [93]. In fact, as opposed to more traditional power plants, renewable energy resources are inherently uncertain so that not only the consumption side but also the generation side is becoming more and more volatile and difficult to predict. For this reason, real-time imbalances and, thus, the need for AS, are expected to increase drastically in the near future [35]. Moreover, beside frequency deviations at the transmission level, renewable energy resources may also create issues, such as overvoltages and/or overloading of cables and transformers, at the distribution level where they are typically connected.

All the aforementioned reasons have compelled grid operators to look beyond traditional generators

for the provision of ASs. In this direction, Distributed Energy Resources (DER)s, such as controllable loads, often referred to as Demand Response (DR), and electrical storage, have been identified as credible alternatives [22]. In particular, BESSs represent perfect candidates thanks to their high responsiveness and ability to provide a wide range of grid services [112]. This has induced, in the recent years, a rapid increase in the number of connected batteries that are typically deployed for services characterized by a low energy content (*e.g.* PFC) [65, 66]. However, BESSs are expensive devices which still prevents their usage for numerous applications [105]. On the other hand, loads are ubiquitous and, from a conceptual standpoint, they can also provide ASs by adjusting their consumption patters, i.e., decrease consumption whenever the frequency level is low, and increase consumption when the frequency is over 50 Hz. Moreover, thanks to their inherent distributed nature, controllable loads can also be exploited to reduce the stress at the distribution level by absorbing locally extreme fluctuations of the uncertain generation.

4.1.3 Grid services with DERs

When providing grid services, the main characteristic and challenge that differentiates DERs from conventional generators is represented by the presence of two conflicting objectives - on one side the resource should be able to deliver a high quality service to the grid operator while, on the other side, fulfilling its primary purpose to the level expected by the customers. Thus, only flexible resources have the potential to offer services to the grid. Throughout the years, a number of different resources, such as BESSs, Thermostatically Controlled Loads (TCLs), Plugin Hybrid Electric Vehicles (PHEVs), and commercial buildings, have been identified as being suitable for the provision of grid services both at the distribution as well as at the transmission level. In [105] a comprehensive assessment of the capabilities of these devices is conducted. The reference also provides guidelines, based on technical and economic criteria, to determine the most appropriate type of resource for a given service.

In the following, we start by providing a broad overview of the recent advancements in the control of DERs for the provision of single services either at the transmission or at the distribution level. In particular, few selected contributions are described and their main conclusions are summarized. After that, the focus will shift on a recently introduced paradigm, typically referred to as *multitasking*, where, in order to maximize their exploitation, the controllable resources are employed for the simultaneous provision of multiple services. Finally, a novel model-based control framework that formally characterizes the amount of services that a set of DERs can offer to the grid is presented, and similarities and differences with respect to existing works are highlighted.

Distribution level services

Thanks to their high controllability and the fact that their location can be arbitrarily chosen, BESSs are starting to become a very competitive solution to relieve congestion at the distribution level. Possible applications are: upgrade deferral as in [102], where a rule-based controller for a 1.2 MW battery is used to reduce the peak load on a distribution transformer; peak shaving using optimization-based control methods [72, 124]; absorption of local deviation as in [127] where an MPC controller is designed to dispatch the operation of a distribution feeder characterized by the

presence of stochastic prosumers; and voltage regulation using both centralized [99] and distributed approaches [147].

The potential of the demand side to provide services at the distribution level has been the subject of many studies. Despite DR programs have different peculiarities depending on the considered country, they can, in general, be categorized in two groups - price-driven programs, and direct load control programs.

In a price-driven scheme, the general idea is to affect the consumption pattern by appropriately setting energy or power prices. Common strategies being explored are: standard day-night tariffs; critical peak pricing for small business customers [134], large commercial buildings [43] or TCLs [56], where the prices can experience a manifold increment during peak hours; dynamic pricing where the price tag is chosen to reflect the marginal cost of energy provision; peak-power reduction where, beside the consumed energy, the load also pays for the maximal power drawn over a pre-determined period of time (*e.g.* one month). Regardless of the particular scheme being considered, it is clear that the key to optimize the operation cost of a resource participating in such programs is the capability of adapting to rapidly changing external conditions. Because of this, MPC has emerged in the literature as one of the most promising control techniques thanks to its ability of directly incorporating time-varying prices in the problem formulation [104, 105].

In a direct load control, the loads are incentivized to reduce/modify their consumption at specific times of the day. The remunerations are then based on the amount of reduction the loads can provide. Reference [116] studies the potential using MPC of an office building participating in New York's DR program. A similar study is represented by [137] where a model-based method is used to provide DR services with a building equipped with an additional thermal storage. The work [138] investigates the usage of TCLs for the provision of voltage regulation.

Transmission level services: frequency control

BESSs clearly represent a potential candidate for services at the transmission level thanks to the high ramp rates they can achieve, which typically exceed the requirements for AS provision. However, the main challenge is represented by the management of the SoC since regulation signals can display significant biases over time. Thus, in most approaches, it is assumed that BESSs are allowed to only track the fast and zero-mean components of the regulation signal, while low-frequency and biased components are passed to slower units [19, 108]. In particular, most applications have been focusing on PFC due to the smaller energy requirements involved.

When dealing with loads, the provision of AS presents specific challenges and, as a consequence, requires more sophisticated controllers. First of all, the participating resource should declare, ahead of time and over the whole regulation period, its baseline consumption as well as the power capacity around the baseline. The first determines the energy the resource commits to consume in absence

of any regulation requests. The second characterizes the maximum deviations around the baseline the resource is willing to sustain. Finally, the last challenge is represented by the strict tracking requirements that are typically imposed during online operation. In the following, we summarize the main contributions appearing in the literature to address these problems. We focus mainly on theoretical/simulation-based works leaving the review of experimental contributions to Chapter 6.

A min-max approach to compute the electrical flexibility and the baseline of a commercial building for the provision of SFC was investigated in [87]. However, the applicability of the method is restricted to a single zone served by a single electrical fan. A robust-based method for the control of multi-zone systems is proposed in [47] where the load's flexibility is characterized as a virtual battery with power and energy limitation. Again for commercial buildings, a two-stage approximation scheme was designed in [118] and extended in [117] to account also for the non linear models of HVAC systems. Regarding TCLs, most of recent research in the area has focused on the development of aggregated system models to be used in the controller design [54, 90, 138]

Regardless of the type of resource or service considered, one of the shared conclusions of these studies is that a potential obstacle for the exploitation of DERs as providers of grid services is their limited storage capacity, i.e., their inability to store or release energy over extended periods of time. For instance, for the case of a BESS, the energy capacity is not only one of the defining specifications, but it also represents the main driver for its cost. Thus, it would be desirable to keep the energy/power ratio as small as possible. Similarly, for the case of an aggregation of loads (PHEVs, TLCs, or commercial buildings), and for a fixed amount of grid service to be provided, one would like to keep the aggregation as small as possible in order to minimize the dispersion of the economic return among the participants. Nevertheless, due to the worst-case energy requirements and/or the conservative prequalification rules implemented by many TSOs, the sizing of a BESS (respectively loads aggregation) for a particular service is often dictated by few extreme scenarios seldomly encountered in practice. In particular, this could cause an under exploitation of the resources which consequently reduces the return on the invested capital (equivalently an increment of the operating cost). A possible way to improve the economics is to stack multiple services and optimize their provision. Thus, when the available control power of a resource is not used to fulfill the main service (*e.g.* alleviation of local grid congestions, peak shaving), it could be assigned to provide a secondary service (*e.g.* frequency control). Such a way of coupling different services and dynamically allocating the available power across them is typically referred to as *multi-tasking* [38] and it has been the subject of recent intense research activity.

4.1.4   Review on the provision of multiple grid services by DERs

Several recent works have proposed the simultaneous provision of multiple services in order to optimize the exploitation of distributed battery storage systems, as well as other distributed controllable resources, such as thermostatically controlled loads [26, 33, 34, 37, 60, 61, 92, 94–96, 113, 126, 135, 142–144].

For instance, [92] proposes the usage of distributed storage to provide simultaneously primary frequency control and minimize PV curtailment. In [144], the authors formulate a stochastic dynamic programming model to schedule the operation of distributed energy storage devices simultaneously participating energy and regulation markets. Similarly, [143] proposes a rolling horizon optimization scheme and considers an energy storage device generating revenue from energy arbitrage, balancing service, distribution system deferral and outage mitigation. Table 4.1 summarizes the existing literature, highlighting the service synergies analyzed by the works cited above.

Independently of the specific control scheme or objective, the shared conclusion of these works is that the provision of multiple services simultaneously results in an operation that is more effective than the single service provision either in terms of higher revenues or satisfaction of technical objectives.

Most of the aforementioned works aim at optimizing the revenue of the considered resource. For instance, [61] propose the joint participation to the wholesale electricity market (i.e. doing energy arbitrage) and the ancillary services market (*e.g.* providing primary reserve), while [126] aims at maximizing the revenue of a battery providing primary reserve and jointly performing peak-shaving in a tariff scheme including a charge for the peak demand. Such schemes are subject to market rules that may change as the penetration of storage devices increase (notably, invalidating the assumption that such resources can be considered as price-takers). On the other side, BESSs and other distributed resources are most often installed in distribution networks to provide *local* services that satisfy *technical* objectives. Among these, there are energy management (*e.g.* peak shaving [109] or load levelling [97]), voltage control for active distribution networks [28, 140, 141] or congestion management [55].

In this chapter, we aim at designing a scalable and general model-based method to formally characterize the amount of local (*e.g.* congestion management of a MV network, peak shaving for a residential building) and shared services (*e.g.* SFC to the upper grid layer) that a set of distributed resources, otherwise employed individually and each for a single service, can provide when controlled in a coordinated fashion.

Moreover, we observe that in most of the aforementioned literature the proposed control schemes focus on determining only the active power schedules and/or real-time injection of the controlled devices, although BESSs, as well as other distributed controllable resources, are interfaced to the grid through power electronics capable of injecting reactive power as well. The works [95] and [34] do schedule the reactive power injections; however, the objective in these works is limited to constraining the apparent power injection and voltage levels at the connection point of the storage system, rather than managing the voltage levels across the considered distribution network, which is in fact a more ambitious goal attainable through the control of distributed controllable resources (see *e.g.* [30, 129]). With respect to such literature, we propose a scheme to control both active and reactive power injections of such devices, so as to take advantage of their reactive power capabilities.

Author, year	Considered services	Methods, contributions	Main outcomes
Wu 2015, [143]	Energy arbitrage, balancing services	Receding horizon optimization, includes outage mitigation in the revenue streams	Multiple revenue streams can be captured simultaneously
Kazemi 2017, [60, 61]	Simultaneous offering in day-ahead energy, spinning reserve, and regulation markets	Robust optimization considering uncertainties related to market prices and energy deployment and, in [60], the battery lifetime	Participation in multiple markets increases profits
Drury 2011, [33]	Operating reserves in addition to energy arbitrage	Based on a heuristic method, specific to compressed air energy storage (CAES)	Providing operating reserves simultaneously to arbitrage increases annual net CAES revenues
Cheng 2016, [26]	Energy arbitrage and frequency regulation	Multi-scale dynamic programming	Focuses on assessing the algorithm performance rather than its economic benefits
Megel 2015, [92]	Frequency regulation and load smoothing or minimization of PV curtailment	MPC scheduling, considers transformer overheating	Multitasking can almost double a storage system's profits as compared with a single-service approaches
Shi 2017, [126]	Peak shaving and frequency regulation	Joint optimization considering battery degradation, operational constraints, and uncertainties in customer load and regulation signals	The saving from joint optimization is often larger than the sum of the optimal savings for the two individual applications, with the gain being superlinear
Xi 2014, [144]	Energy arbitrage and frequency regulation, backup service vs outages	Stochastic dynamic programming, market and system uncertainties	Batteries can achieve much larger economic benefits than previously thought if they jointly provide multiple services
Moreno 2015, [95]; Perez 2016, [113]	Distribution network congestion management, energy price arbitrage and various reserve and frequency regulation services	Mixed integer linear programming, constraints the power flow at the electrical substation; [113] considers as well battery degradation	Distributed storage revenues associated with frequency control services are significantly more profitable
Namor 2018, [96]	Dispatch of a MV feeder and primary frequency regulation	Day-ahead robust optimization accounting for uncertainties in the load and regulating signal forecasts; experimental validation of the control framework	The simultaneous provision of the two services allows to fully exploit the BESS energy capacity
Engels 2017, [37]	Frequency control and maximization of self-consumption for a PV-storage installation	Chance-constrained robust optimization	optimally combining the two services increases value from batteries significantly
Trovato 2016, [135]	Frequency services and energy arbitrage	Controls aggregates of thermostatically controlled loads, linear optimization model	Clustering of appliances with similar capabilities can significantly enhance the flexibility available to the system
Dutrieux 2013, [34]	Energy arbitrage, removal of grid constraints on a reactive power management	Two stages approach defining a priori a BESS operating domain respecting grid constraints	Demonstrates the feasibility of providing simultaneously the mentioned services
Moreira 2016, [94]	Energy arbitrage, peak demand shaving and various balancing services	Rather than focusing on control, it assesses the synergies and conflicts among possibly concurrent services	Services interact differently depending on markets and system operating conditions

Table 4.1 – Review on the provision of multiple grid services by DERs.

In particular, this is exploited to guarantee the satisfaction of network constraints across the hosting distribution grid, as well as performing other local (peak shaving) and global (frequency regulation) services. We do so by integrating, in the proposed scheduling problem, a set of constraints based on the power flow (PF) equations of the grid which are linearized around the current operating condition (as detailed in Section 4.6) to preserve the tractability of the planning problem.

4.1.5 Contribution of the chapter

We consider an aggregation of controllable energy resources that are distributed in an active distribution network. The problem we look at is how to schedule the operation of these resources so as to maximize the provision of a shared tracking service while ensuring the satisfaction of local constraints of the network as well as the provision of the primary service of each resource.

The contributions of this chapter with respect to existing works in the literature are:

- the formulation of the optimal day-ahead planning problem that determines the allocation of the available control power across multiple local and shared services. The problem is formulated as a multi-stage uncertain optimization problem and it differs from existing works in: *i*) it accounts for both the uncertainty across the network as well as the uncertainty of the tracking signal received from the TSO, *ii*) it explicitly considers PF constraints, *iii*) to increase the offerable flexibility, it encodes directly the possibility to adjust the baseline power consumption in the intraday market.

- the formulation of two approximated solutions which are based on robust optimization techniques and on the multi-stage linearization of the PF constraints around a general point.

- the formal investigation of the properties of the considered PF linearization which is used to explicitly express voltages and currents across the network as a function of the controllable active and reactive power injections.

- the presentation of an extensive simulation study that demonstrates the capability of the proposed framework to adapt to different levels of uncertainties acting on the network while guaranteeing the provision of both local and shared services.

4.1.6 Preliminaries and notation

As commonly done in this thesis, we use uppercase letters for matrices and lower case letters for vectors. r_i represents the value of vector r at time i. Bold letters are used to denote sequence over time, i.e., $\mathbf{r} = [r_0^T, \ldots, r_{N-1}^T]^T$. To denote fixed quantities, we use the hat notation, *e.g.*, $\hat{r}$ is a vector with fixed values. Solutions to an optimization problem are denoted with r^*. The cumulative sum of a signal over time is denoted as $\tilde{r}_{[l \to i]} = \sum_{t=l}^{i} r_t$. The bracket notation, $\mathbf{r}^{(k)}$ indicates the realization of the vector $\mathbf{r}$ corresponding to the k-th scenario.

Given $x \in \mathbb{R}^n$, let $\text{diag}(x) \in \mathbb{R}^{n \times n}$ be the associated diagonal matrix. For a given matrix $A \in \mathbb{R}^{n \times n}$, A is a Metzler or M-matrix if $A_{ij} \leq 0 \; \forall \; i \neq j$ and all eigenvalues of A have positive real parts. In this case $A^{-1} \geq 0$, with strict inequality if A is irreducible [8, Chapter 6]. For a real symmetric matrix $A \succeq 0$, $x^T A x = 0$ if and only if $x \in kernel(A)$ [20, Appendix C].

4.2 Problem Statement

We consider a distribution network hosting uncontrollable demand and renewable generation as well as a set of controllable energy resources. The latter are employed jointly to provide a set of local services and, at the same time, to track a signal r provided by an exogenous entity (*e.g.* a regulation signal provided by the grid operator). The problem consists of determining a *baseline* active power injection, **D**, for the aggregation of controllable resources over the future horizon, N^1, as well as a *power flexibility*, γ, determining the maximum deviation that the aggregation can sustain around the baseline. More specifically, if a bid of 1 MW is accepted, the aggregation will receive a tracking request from the TSO with a maximum magnitude of ± 1 MW. Thus, the values of **D** and γ should be appropriately determined so as to respect the operational limits of the controllable resources as well as the grid's constraints. This should be guaranteed against any occurrence of the uncertainty acting on the system.

Formally, the controllable resources are modeled through a set of time-invariant linear models, as described in Section 4.3. The distribution grid is assumed to be balanced, composed of n buses, and modeled via its direct sequence circuit. To characterize the relation among the active power p and reactive power q and the voltage magnitude v and phase θ at each node, we consider the power flow equations described in Section 4.4. To ensure the proper and reliable operation of the grid, limits on the magnitudes of node voltages and currents are also considered. The forecasts of the power injections at all grid nodes as well as the tracking signal, r, are uncertain at the moment of the scheduling. As a consequence, the values of **D** and γ, over the next regulation period, are determined by solving an uncertain optimization problem, detailed in Section 4.5

4.3 Resource Models and Constraints

In this section, we introduce the model describing the dynamical behavior of a single resource. Rather than providing the details related to a particular type of resource, we only describe the class of models that are considered in the chapter. In the simulations Section 4.7, a particular instance for the case of an electrical battery, is provided.
We consider time-invariant linear systems of the form:

$$
\begin{aligned}
x_{i+1} &= Ax_i + Bu_i + Ed_i \\
y_i &= Cx_i \\
s_i &= [\eta \quad \xi]^T u_i
\end{aligned}
\tag{4.1}
$$

where $x_i \in \mathbb{X} \subseteq \mathbb{R}^{n_x}$ represents the state, $u_i \in \mathbb{U} \subseteq \mathbb{R}^{n_u}$ represents the input, $d_i \in \mathbb{R}^{n_d}$ the external disturbance acting on the resource, and $s_i \in \mathbb{S} \subseteq \mathbb{R}^2$ the admissible apparent power injection to the resource expressed in rectangular coordinates, i.e., $s_i = [p_i, q_i]^T$ with p_i and q_i representing active and reactive power injection, respectively. The set $\mathbb{X}$ is used to model state constraints, the set $\mathbb{Y}$ models output constraints (*e.g.* state of charge (SoC) constraints for a BESS). The convex set $\mathbb{U}$

---

[1]The considered temporal horizon is discretized in N intervals of equal length Δ_T. As a possible example, we could consider a 24-hours period with discretization time of 15 minutes which gives $N = 96$.

is used to model actuators operational constraints such as limits on the maximum apparent power for the inverter of a BESS. The term Ed_i models the perturbation of external quantities (outside temperature, solar radiation, etc.) on the state dynamics. Finally, as it is quite common [107, 121], the quantities $\eta, \xi \in \mathbb{R}^m$ are conversion factors that translate the input to the active and reactive power injection at the resource (expressed in MW, Mvar).

With these definitions, we can the introduce the set of admissible input trajectories, which comply with the system dynamics and constraints, over the considered prediction horizon, N:

$$
\mathcal{U}(\hat{x}, \hat{\mathbf{d}}) = \left\{ \mathbf{u} \left|
\begin{array}{l}
x_{i+1} = Ax_i + Bu_i + E\hat{d}_i \\
y_i = Cx_i \\
s_i = [\eta \quad \xi]^T u_i \\
u_i \in \mathbb{U}, \quad y_i \in \mathbb{Y}, \quad s_i \in \mathbb{S} \\
x_0 = \hat{x} \\
\forall i = 1, \ldots, N
\end{array}
\right. \right\}
\tag{4.2}
$$

In the definition of the set $\mathcal{U}$, both the initial state, $\hat{x}$, as well as the vector of external perturbations, $\mathbf{d}$, along the prediction horizon, are considered as parameters and are, therefore, fixed. Regarding the initial state we assume it to be either directly measurable or that it can be estimated by means of a global observer relying on state estimation techniques (*e.g.* Discrete Kalman filter, least square estimation, etc.). For $\mathbf{d}$ we assume the availability of reliable predictions, denoted as $\hat{\mathbf{d}}$, that can be obtained, *e.g.*, from external entities (weather stations).

Given that there might be more than one controllable resource in the aggregator portfolio, we further introduce the notation, $\mathbf{u}^j$, to denote the input trajectories for the j-th resource together with the corresponding admissible set, $\mathcal{U}^j(\hat{x}^j, \hat{\mathbf{d}}^{\mathbf{j}})$.

We refer the reader to Section 4.7 for an example of the model for the case of a BESS.

4.4   Power Network Models and Constraints

We consider a balanced distribution network with n buses with only one slack bus, with index 0, which is assumed to be the Point of Common Coupling (PCC) with the upstream network. Thus, all other buses are indexed sequentially starting from the PCC. The network can be modelled as a linear circuit represented by a connected weighted graph $O(\mathcal{V}, \mathcal{E}, \mathcal{W})$ where $\mathcal{V} = \{0, \cdots, n-1\}$ is the set of vertices (buses) and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ is the set of edges (branches).

Let $z_{ht} = r_{ht} + jx_{ht} \in \mathbb{C}$ be the impedance between bus h and t, where $r_{ht} \in \mathbb{R}_{>0}$ is the resistance and $x_{ht} \in \mathbb{R}_{>0}$ is the inductive reactance. The edge weights of the associated graph are the associated admittances $y_{ht} = g_{ht} + jb_{ht} \in \mathbb{C}$, where $g_{ht} = r_{ht}/(r_{ht}^2 + x_{ht}^2) \in \mathbb{R}_{>0}$ is the associated conductance

and $b_{ht} = -x_{ht}/(r_{ht}^2 + x_{ht}^2) \in \mathbb{R}_{<0}$ the susceptance. The network is represented by symmetric admittance matrix $Y \in \mathbb{C}^{n \times n}$, where the off-diagonal elements are given by $Y_{ht} = -y_{ht}$ for each branch $\{h, t\} \in \mathcal{E}$ (0 if $h, t \notin \mathcal{E}$), and the diagonal elements are given by $y_h^{sh} + \sum_{t=0, t \neq h}^{n-1} y_{ht}$, where y_h^{sh} is the shunt admittance at bus h. We represent $Y = G + jB$, where G and B respectively are the conductance and susceptance matrices. It should be noted that $G_{hh} = g_h^{sh} + \sum_{t=0, t \neq h}^{n-1} g_{ht} > 0$ and $G_{ht} = -g_{ht} < 0$.

Remark 6. *For the sake of simplicity, we assume in the following that no voltage regulating transformers are present in the branches. Under this assumption, the matrices G and B are symmetric. Moreover, G is a Laplacian matrix (loopy Laplacian if the shunt conductances are not equal to zero) [32]. We highlight however that such an assumption is not strictly necessary for the control framework proposed in the following.*

As we are interested in the scheduling problem, we consider the network to be in steady state and operating in perfect sinusoidal conditions. Thus, to each node $h \in \mathcal{V}$, we can associate a phasor voltage $\vartheta^h := v^h e^{\theta^h} \in \mathbb{C}$, with v^h, $\theta^h \in \mathbb{R}$ and complex apparent power $s^h := p^h + jq^h \in \mathbb{C}$, with p^h, $q^h \in \mathbb{R}$. Considering the aforementioned notation, it is possible to concisely write the AC power flow equations relating voltages and power injections. The power flow equations descend directly from Kirchhoff's and Ohm's laws [70] as

$$s = \text{diag}(\vartheta) \circ (\overline{Y\vartheta}). \tag{4.3}$$

where $\circ$ denotes element-wise vector multiplication, and the overline notation denotes complex conjugation. By defining the grid state

$$z := [v, \theta, p, q]^T \in \mathbb{R}^{4n}$$

one can rewrite the PF equation (4.3) implicitly as $\text{F}(z) = 0$ where the map is defined as $\text{F}: \mathbb{R}^{4n} \to \mathbb{R}^{2n}$ and it is obtained by rearranging the power flow equation after expressing it in rectangular or polar coordinates.

4.4.1 Partition of the grid state

In the following section, we consolidate the notation that will be used in the remainder of the manuscript. In particular, we provide herein a partition of the grid state, z, that will be helpful for later derivations.

Let us define by $z^{cont} \in \mathbb{R}^{n_c}$ the components of z that are *directly controllable*[2] by the aggregator, and by $z^w \in \mathbb{R}^{n_w}$ the *uncontrollable and uncertain* components of the grid state. Together, these two sets of components form the vector of exogenous variables. We define by

---

[2]Without loss of generality, we are implicitly assuming that the controllable resources do not have errors in the deployed control setpoints. In fact, if this is not true, the uncertainty related to the discrepancy between the setpoints and the implemented power injections could be treated as all other sources of uncertainty.

$$z^{ex} := \begin{bmatrix} z^{cont} \\ z^w \end{bmatrix} \in \mathbb{R}^{2n}$$

the vector formed by stacking all the exogenous variables at each time instant.

The remaining $2n$ components are, in turn, dependent variables which can be derived, through the power flow map, $\mathrm{F}(z)$, from the exogenous components. For this reason, we define by $z^{end} \in \mathbb{R}^{2n}$ the vector containing all dependent variables to which we refer to as endogenous variables.

Finally, by utilizing the newly introduced notation, we can rewrite the set of power system states that satisfy equation (4.3) as:

$$\mathcal{M} := \{(z^{ex}, z^{end}) \mid \mathrm{F}(z^{ex}, z^{end}) = 0\} \tag{4.4}$$

which, as it was shown in [16], describes a 2n-dimensional smooth manifold in $\mathbb{R}^{4n}$ to which the grid state, z, belongs at each time instant, i.

Remark 7. *For the sake of simplicity and clarity of explanation, in the remainder of the chapter, we make the assumption that every bus is a PQ bus. Under this assumption, the state partition is such that for every node, h, the exogenous variables are represented by active, p^h and reactive, q^h power injections that can be either controllable or uncertain depending on the particular configuration. In turn, the endogenous variables are represented by the voltage magnitude, v^h and angle, θ^h. We emphasize that all the results presented hereafter can be easily extended to accommodate other configurations.*

Remark 8. *In what follows, we will assume the voltage at the slack bus to be fixed to a known constant level. Therefore, we are neglecting the possibility, often used in distribution networks to perform voltage control, to modulate the slack voltage by means of On Load Tap Changers (OLTC). This is done principally for two reasons 1) we focus on a scenario in which the network experiences overvoltages at very specific locations which motivated the deployment (or retrofitting) of DERs, 2) optimizing over the transformer tap's positions would result in a mixed integer problem which is difficult to handle in an uncertain setting like the one considered in this chapter. We highlight, however, that the method detailed in the following sections could be extended to account for OLTC as well by considering, for instance, the tap's positions as pseudo-continuous variables [27].*

4.4.2 Grid constraints

Equation (4.4) enforces the grid state, z to lay on the PF manifold, $\mathcal{M}$, which describes the physical relation between voltages and apparent power. In a multi-OPF setting, this map needs to be satisfied at each time instant, i. To compactly write this, we denote the PF constraints along the prediction horizon, N, as:

$$\mathbf{F}(\mathbf{z}^{ex}, \mathbf{z}^{end}) := \begin{bmatrix} F_1(z_0^{ex}, z_0^{end}) \\ \vdots \\ F_N(z_N^{ex}, z_N^{end}) \end{bmatrix} = \mathbf{0} \tag{4.5}$$

which is obtained by simply stacking all the constraints of the form (4.4) along the horizon. A similar definition is assigned to the variables, $\mathbf{z}^{ex}$ and $\mathbf{z}^{end}$.

We assume that the network is subject to additional constraints which define the safe region of operation for the system. Thus, at each bus, h, and for each time instant, i, we impose:

$$\underline{v} \leq v_i^h \leq \bar{v}$$
$$i_i^{ht} \leq \bar{I}^{ht}$$

where $\bar{v}$, $\underline{v}$, denote upper and lower limits on the voltage magnitudes, i_i^{ht} the magnitude of the current flowing in the branch from h to t, and $\bar{I}^{ht}$ the current flow limit for the branch.

Finally, by condensing all the constraints along the horizon, we can define the feasible region for the endogenous variables:

$$\mathcal{Z} := \left\{ \mathbf{z}^{end} \middle| \begin{array}{ll} \underline{v} \leq v_i^h \leq \bar{v} & h = \{1, \ldots, n\} \\ i_i^{ht} \leq \bar{I}^{ht} & \forall (h,t) \in \mathcal{E} \\ \forall i = \{1, \ldots, N\} \end{array} \right\} \tag{4.6}$$

4.5 Problem Formulation

At the time of planning, the most important objective for the aggregator is to allocate enough controllable power so as to guarantee the satisfaction of all the physical constraints of the network. This should be ensured against all load's and generation's forecasts, and even in presence of disturbances forcing the system to deviate from its nominal (or predicted) mode of operation. According to the previously defined notation, we capture this source of uncertainty with the term $\mathbf{z}^w$, which, represents the *uncontrollable and uncertain* components of $\mathbf{z}^{ex}$. In particular, we assume $\mathbf{z}^w$ to belong to a compact uncertainty set $\mathcal{Z}^w \subseteq \mathbb{R}^{N \cdot n_w}$.

Concurrently, the aggregator aims at maximizing its economic return by offering additional ancillary services to the TSO in order to fully exploit all the available controllable power. To this aim, the aggregator is required to compute an appropriate schedule for the controllable resources, in order to ensure a satisfactory level of tracking of the signal, r, received by the TSO. In particular, the tracking should be ensured for *all values of the uncertainty* in a set $\mathcal{R} \subseteq \mathbb{R}^N$.

Finally, we denote with $\Xi := \mathcal{Z}^w \times \mathcal{R}$ the composed uncertainty set which is the result of the Cartesian product between $\mathcal{Z}^w$ and $\mathcal{R}$.

The Planning Problem for the controllable resources is formulated as follows:

Problem 2 (Planning Problem).

$$\underset{\gamma,\bar{\mathbf{D}},\pi_{\mathbf{u}j},\pi_{\mathbf{D}}}{\text{minimize}} \quad c_{\text{energy}}^T \bar{\mathbf{D}} - c_{\text{reward}}\gamma \tag{4.7}$$

$$\text{s.t.}$$

$$\text{(Resource constraints)} \quad \mathbf{u}^j \in \mathcal{U}^j(\hat{x}^j, \hat{\mathbf{d}}^{\mathbf{j}}) \quad \forall j = 1, \ldots, N_{cont} \tag{4.8}$$

$$\text{(Grid constraints)} \quad \mathbf{F}(\mathbf{z}^{ex}, \mathbf{z}^{end}) = \mathbf{0}, \quad \mathbf{z}^{end} \in \mathcal{Z} \tag{4.9}$$

$$\text{(Apparent power controllable res.)} \quad \mathbf{z}^{cont} = \Lambda\mathbf{U} \tag{4.10}$$

$$\text{(Power tracking)} \quad \Gamma\mathbf{z}^{cont} = \mathbf{D} + \gamma\mathbf{r} \tag{4.11}$$

$$\text{(Power flexibility)} \quad \gamma \geq 0 \tag{4.12}$$

$$\text{(Control Policies)} \quad \mathbf{u}^j = \pi_{u^j}(\mathbf{r}, \mathbf{w}), \quad \mathbf{D} = \bar{\mathbf{D}} + \pi_{\mathbf{D}}(\mathbf{z}^w, \mathbf{r}) \tag{4.13}$$

$$\text{(Uncertainty)} \quad \forall(\mathbf{z}^w, \mathbf{r}) \in \Xi. \tag{4.14}$$

where the cost function is composed of a term related to the cost of energy for the purchased baseline, $\bar{\mathbf{D}}$ and a reward for the provided flexibility, γ. Notice how the cost function can accommodate both fixed as well as time-varying flexibility. In particular, the latter can be encoded by simply augmenting the decision space and defining $\boldsymbol{\gamma} = [\gamma_1, \ldots, \gamma_N]^T$.

The set of constraints (4.8) guarantees the satisfaction of all local constraints for the controllable resources. As described in Section 4.3, these might include, depending on the type of resource, comfort constraints and HVAC limits (for a commercial building), SoC limitations and limits on the inverter (BESS), or even local requirements such as peak shaving. Examples of this constraint will be provided in the simulation Section 4.7.

The first equation in (4.9) enforces the grid state to belong to the PF manifold for all time instants along the prediction horizon. In turn, the second constraint guarantees the satisfaction of all constraints on the endogenous variables, $\mathbf{z}^{end}$. Equation (4.10) simply states that active and reactive power injections at the controllable resources can be obtained as a linear combination of their input variables. The term $\mathbf{U}$ is obtained by simply stacking the individual resource components in one vector, i.e., $\mathbf{U} := [(\mathbf{u}^1)^T, \ldots, (\mathbf{u}^{N_{cont}})^T]^T$. In a similar manner, the matrix Λ, obtained directly from (4.2), maps the input components at each resource into their active and reactive power. The expression for Λ can be found in the appendix.

The constraint (4.11) imposes the tracking requirement. It states that the total active power consumption of the controllable resources, given by $\Gamma\mathbf{z}^{cont}$, should be equal to the total power

request which represented by the term $\mathbf{D} + \gamma \mathbf{r}$. More precisely, at each time instant, i, it enforces the following equality:

$$p_i^1 + p_i^2 + \ldots + p_i^{N_{cont}} = D_i + \gamma r_i$$

Thus, the matrix Γ (derivation in the appendix) is constructed so as to select and sum the active power components of the vector $\mathbf{z}^{cont}$.

Finally, equations (4.13) encode the fact that, as the uncertainty $(\mathbf{z}^w, \mathbf{r})$ is revealed in real-time, both the input trajectories, as well as the baseline (according to market rules) profile can be causally adjusted accordingly. Thus, the only decision variables which are fixed in the planning phase are the baseline $\bar{\mathbf{D}}$, and the offered power flexibility γ. The remaining decision variables are represented by the control policies $\pi_{\mathbf{u}^j}, \pi_{\mathbf{D}}$, that are enforced to be causal, i.e., decisions at stage i only depend on observations of the uncertainty up to time i.

With the current formulation, Problem 2 is a non-convex infinite-dimensional optimization problem that cannot be readily solved by existing solvers. This is due to the presence of the non-convex multi-period PF constraints (4.9), the infinite-dimensionality of the decision space (4.13) (optimizing over functions rather than over control sequences), and to the infinite number of constraints (4.14).

To address these issues and obtain a tractable approximation of the planning problem, the following elements need to be specified:

- an approximation of the non-convex PF constraint (Section 4.6.1)

- the choice of the uncertainty set, $\Xi := \mathcal{Z}^w \times \mathcal{R}$ (Section 4.6.2)

- the parametrization of the control policies, $\pi_{u^j}, \pi_{\mathbf{D}}$. (Section 4.6.3)

Remark 9. *In the Planning Problem, we assumed that the aggregator incurs a cost only for the day-ahead baseline $\bar{\mathbf{D}}$. However, in reality, a cost/reward is also associated to the actual baseline, $\mathbf{D}$, after the intraday transactions. The latter depends on the realization of the uncertainty and it was not included in the problem formulation for the sake of simplicity. However, different methods exist in the literature to account for these types of costs. We refer the reader to chapter 7 where an example of how to include costs at subsequent stages is presented.*

4.6 Problem approximation

4.6.1 PF approximation

Several methodologies to deal with the non-linearity of the exact PF equations have been recently proposed in the literature. Common approaches typically rely either on convex relaxation techniques (resulting in semidefinite programming [82]), or on the linearization of the PF constraints.

Regarding the latter, first attempts focused on balanced distribution networks [17, 31, 130]. A more general approach is represented by [27] which proposes an efficient method to compute sensitivity coefficients of the node voltages and line currents in the context of unbalanced networks with radial topology. General grid topology and bus models are considered in [16] where a derivation of the best linear approximant around a general voltage condition is obtained. Finally, the extension to generic load models with wye and delta connections is introduced in [9].

In order to obtain a tractable reformulation of the original Problem 2, we consider, in this chapter, the implicit linearization introduced in [16] which gives a nice trade-off between computational efficiency, ease of notation, and general applicability. We highlight once more that we could have equivalently considered the approaches [9, 27].

Due to the fact that $\mathcal{M}$ is a smooth manifold, it is always possible to define the tangent space at a given grid state, $\hat{z}$. More precisely, one can define the linear manifold:

$$\mathcal{A}_{\hat{z}}(z - \hat{z}) = \mathbf{0}_{2n}, \tag{4.15}$$

where the matrix $\mathcal{A}_{\hat{z}}$ only depends on the chosen general voltage profile, $\hat{\vartheta}$, and it can be expressed as follows

$$\mathcal{A}_{\hat{z}} = \left[\left(\langle \mathrm{diag}\, \overline{Y\hat{\vartheta}} \rangle + \langle \mathrm{diag}\, \hat{\vartheta} \rangle N \langle Y \rangle \right) R(\hat{\vartheta}) \quad -\mathbf{I} \right], \tag{4.16}$$

with

$$\hat{\vartheta} := \hat{v} e^{j\hat{\theta}}$$

$$N := \begin{bmatrix} \mathbf{I}_{n \times n} & 0_{n \times n} \\ 0_{n \times n} & -\mathbf{I}_{n \times n} \end{bmatrix}$$

$$\langle A \rangle := \begin{bmatrix} \mathrm{Re}\, A & -\mathrm{Im}\, A \\ \mathrm{Im}\, A & \mathrm{Re}\, A \end{bmatrix} \qquad .$$

$$R(u) := \begin{bmatrix} \mathrm{diag}(\cos\theta) & -\mathrm{diag}(v)\mathrm{diag}(\sin\theta) \\ \mathrm{diag}(\sin\theta) & \mathrm{diag}(v)\mathrm{diag}(\cos\theta) \end{bmatrix}$$

In the following sections, we will analyze the properties of such a linearization. In particular, we are interested in deriving conditions under which one can explicitly express the relation between endogenous (voltage magnitudes and angles) and exogenous components of the grid state (active and reactive power injections). As it will be detailed, this corresponds to the invertibility of a particular submatrix of (4.16) for which we provide an explicit expression.

For the sake of clarity, we start by focusing on the resulting submatrix for the particular case of the flat voltage profile. We show that, for this case, such conditions can be easily derived. Finally, we investigate how these conditions are modified for the case of practical interest of a generic

linearization point.

Flat voltage profile

By fixing the linearization point to the *flat voltage profile* (*i.e.* $\hat{\vartheta} = \mathbf{1}$), the general linear approximant of equation (4.16) can be further simplified as

$$A_{\text{flat}} = [\langle \text{diag } \overline{y^{sh}} \rangle + N \langle Y \rangle \quad -\mathbf{I}] \tag{4.17}$$

By assuming zero shunt admittances, one obtains the following relation:

$$\begin{bmatrix} \text{Re}Y & -\text{Im}Y \\ -\text{Im}Y & -\text{Re}Y \end{bmatrix} \begin{bmatrix} v \\ \theta \end{bmatrix} = \begin{bmatrix} p \\ q \end{bmatrix} \tag{4.18}$$

Moreover, as previously described, we set the voltage magnitude and angle of slack bus (bus 0) to fixed values, i.e., $u^0 = 1$ [3]. On partitioning the buses into $\{0\}$ and $\mathcal{F} := \{1, \ldots n-1\}$, and utilizing $\text{Re}Y = G$, $\text{Im}Y = B$, the set of linear equations (4.18) can be rewritten as

$$\begin{bmatrix} G_{00} & G_{0\mathcal{F}} & -B_{00} & -B_{0\mathcal{F}} \\ G_{0\mathcal{F}} & G_{\mathcal{F}\mathcal{F}} & -B_{0\mathcal{F}} & -B_{\mathcal{F}\mathcal{F}} \\ -B_{00} & -B_{0\mathcal{F}} & -G_{00} & -G_{0\mathcal{F}} \\ -B_{0\mathcal{F}} & -B_{\mathcal{F}\mathcal{F}} & -G_{0\mathcal{F}} & -G_{\mathcal{F}\mathcal{F}} \end{bmatrix} \begin{bmatrix} 1 \\ v^{\mathcal{F}} \\ 0 \\ \theta^{\mathcal{F}} \end{bmatrix} = \begin{bmatrix} p^0 \\ p^{\mathcal{F}} \\ q^0 \\ q^{\mathcal{F}} \end{bmatrix}$$

On eliminating the equations corresponding to the slack bus (which does not contribute to the solution of the remaining buses), one obtains

$$\underbrace{\begin{bmatrix} G_{\mathcal{F}\mathcal{F}} & -B_{\mathcal{F}\mathcal{F}} \\ -B_{\mathcal{F}\mathcal{F}} & -G_{\mathcal{F}\mathcal{F}} \end{bmatrix}}_{\mathcal{M}_{\mathcal{F}\mathcal{F}}^{\text{flat}}} \begin{bmatrix} v^{\mathcal{F}} \\ \theta^{\mathcal{F}} \end{bmatrix} = \begin{bmatrix} p^{\mathcal{F}} \\ q^{\mathcal{F}} \end{bmatrix} + \begin{bmatrix} \overline{v}^{\mathcal{F}} \\ \overline{\theta}^{\mathcal{F}} \end{bmatrix}.$$

The constant quantities on the right-hand side are obtained by simply rearranging the terms (in the second and fourth columns) corresponding the the slack bus.

Lemma 1. $\mathcal{G}_{\mathcal{F}\mathcal{F}}$ is a positive definite M-matrix. $\mathcal{G}_{\mathcal{F}\mathcal{F}}^{-1}$ is well defined, positive and positive definite.

proof. Matrix $G_{\mathcal{F}\mathcal{F}}$ is symmetric with positive diagonal and non-negative off-diagonal entries. Since the graph O is connected, $G_{\mathcal{F}\mathcal{F}}$ has at least one row with strictly positive row sum. $G_{\mathcal{F}\mathcal{F}}$ is nothing but a loopy Laplacian (sum of a Laplacian and a diagaonal matrix) and can be written as:

$$G_{\mathcal{F}\mathcal{F}} = G'_{\mathcal{F}\mathcal{F}} + \text{diag}(G'_{\mathcal{F}0})$$

where $G'_{\mathcal{F}0} = -G_{\mathcal{F}0} \in \mathbb{R}^{m+n \times 1}$ is a non-negative vector. $G'_{\mathcal{F}\mathcal{F}} \succeq 0$ is a Lapalcian matrix and $\text{diag}(G'_{\mathcal{F}0}) \succeq 0$. Since $kernel(G'_{\mathcal{F}\mathcal{F}}) \cap kernel(\text{diag}(G'_{\mathcal{F}0})) = \phi$, $G_{\mathcal{F}\mathcal{F}} \succ 0$. All the eigen values of $G_{\mathcal{F}\mathcal{F}}$ are positive. Hence, $G_{\mathcal{F}\mathcal{F}}$ is a positive definite M-matrix. Positive definiteness of $G_{\mathcal{F}\mathcal{F}}$ implies

---

[3]Considering the voltage magnitude as fixed to 1 can be done without any loss of generality as a numerical re-normalization of the voltage levels can always be performed.

that $G_{\mathcal{FF}}^{-1}$ is well defined and is also positive definite. Additionally, $G_{\mathcal{FF}}$ is an irreducible M-matrix implying $G_{\mathcal{FF}}^{-1}$ is positive.

Lemma 2. The matrix $\mathcal{M}_{\mathcal{FF}}^{\mathrm{flat}}$ is invertible.

proof. The matrix $\mathcal{M}_{\mathcal{FF}}^{\mathrm{flat}}$ is a block matrix and is invertible if and only if both $-\mathcal{G}_{\mathcal{FF}}$ and the Schur complement $S = \mathcal{G}_{\mathcal{FF}} + \mathcal{B}_{\mathcal{FF}}\mathcal{G}_{\mathcal{FF}}^{-1}\mathcal{B}_{\mathcal{FF}}$ are invertible [145]. Lemma 1 guarantees the nonsingularity of $-\mathcal{G}_{\mathcal{FF}}$. Thus, we only need to show that the matrix S is positive definite. But this is guaranteed by the fact that both $G_{\mathcal{FF}}$ and $G_{\mathcal{FF}}^{-1}$ are positive definite matrices.

Approximation around a general voltage profile

The results above imply that it is always possible, around the flat voltage condition, to explicitly relate the endogenous variables, $[v^{\mathcal{F}} \ \theta^{\mathcal{F}}]$, as a function of the exogenous ones, $[p^{\mathcal{F}} \ q^{\mathcal{F}}]$. In the following, we show that this property is not lost when one considers, as it is often the case in distribution networks, small perturbations to the flat voltage condition.

Let us define the function

$$\vartheta^{\mathcal{F}} \in \mathbb{C}^{n-1} \to \mathcal{M}_{\mathcal{FF}}(\vartheta^{\mathcal{F}}) \in \mathbb{C}^{n-1 \times n-1}$$

which characterizes the family of matrices that can be obtained by varying the linearization point, $\vartheta^{\mathcal{F}}$. According to this definition, the matrix at flat voltage can be computed as $\mathcal{M}_{\mathcal{FF}}^{\mathrm{flat}} = \mathcal{M}_{\mathcal{FF}}(\mathbf{1}_{n-1})$. Moreover, by expanding (4.16), it is also possible to explicitly express $\mathcal{M}_{\mathcal{FF}}(\vartheta^{\mathcal{F}})$ as

$$\mathcal{M}_{\mathcal{FF}}(\vartheta^{\mathcal{F}}) = \left(\langle \mathrm{diag}(\overline{\varsigma^{\mathcal{F}}}) \rangle + \langle \mathrm{diag}(\vartheta^{\mathcal{F}}) \rangle \mathcal{M}_{\mathcal{FF}}^{\mathrm{flat}} \right) R(\vartheta^{\mathcal{F}}) \tag{4.19}$$

where $\varsigma^{\mathcal{F}} := i^{\mathcal{F}} e^{j\alpha^{\mathcal{F}}}$ is the vector of the nodal currents for buses in $\mathcal{F}$. With this notation we can state the following

Proposition 1. Let $\mathcal{M}_{\mathcal{FF}}(\vartheta^{\mathcal{F}})$ be the family of matrices defined in (4.19), $\vartheta^{\mathrm{flat}} = \mathbf{1}_{n-1}$ the vector of flat voltages, and $\mathcal{M}_{\mathcal{FF}}^{\mathrm{flat}} = \mathcal{M}_{\mathcal{FF}}(\vartheta^{\mathrm{flat}})$. Then

  (i) There exists a neighborhood U of $\vartheta^{\mathrm{flat}}$, such that $\forall \vartheta^{\mathcal{F}} \in U$, the matrix $\mathcal{M}_{\mathcal{FF}}(\vartheta^{\mathcal{F}})$ is invertible.

 (ii) Consider the ∞-norm $|| \cdot ||_{\infty}$ in $\mathbb{C}^{n-1}$, and $\sigma_{\min}(\cdot)$ the minimum singular value of a given matrix. If

$$||i^{\mathcal{F}}||_{\infty} < \sigma_{\min}\left(\langle \mathrm{diag}(\vartheta^{\mathcal{F}}) \rangle \mathcal{M}_{\mathcal{FF}}^{\mathrm{flat}} \right) \tag{4.20}$$

then the matrix $\mathcal{M}_{\mathcal{FF}}(\vartheta^{\mathcal{F}})$ is invertible.

proof.   (i) The function $\vartheta^{\mathcal{F}} \to \mathcal{M}_{\mathcal{FF}}(\vartheta^{\mathcal{F}})$ can be easily shown to be continuous as it is obtained as the composition of continuous functions of the term $\vartheta^{\mathcal{F}}$. This implies [59] that also the spectrum of $\mathcal{M}_{\mathcal{FF}}$ is a continuous function of the linearization point, $\vartheta^{\mathcal{F}}$.

This and the eigenvalues of $\mathcal{M}_{\mathcal{FF}}^{\mathrm{flat}}$ being non-zero, implies that it is possible to find a neighborhood, U, of $\vartheta^{\mathrm{flat}}$ where $\mathcal{M}_{\mathcal{FF}}(\vartheta^{\mathcal{F}})$ is invertible.

(ii) The matrix $R(\vartheta^{\mathcal{F}})$ is invertible. Thus, to determine the invertibility of $\mathcal{M}_{\mathcal{F}\mathcal{F}}$ it suffices to investigate the properties of the matrix

$$\langle \text{diag}(\overline{\varsigma^{\mathcal{F}}}) \rangle + \underbrace{\langle \text{diag}(\vartheta^{\mathcal{F}}) \rangle \mathcal{M}_{\mathcal{F}\mathcal{F}}^{\text{flat}}}_{H} \tag{4.21}$$

Lemma 2 guarantees the invertibility of $\mathcal{M}_{\mathcal{F}\mathcal{F}}^{\text{flat}}$. Similarly, it can be easily shown that $\langle \text{diag}(\vartheta^{\mathcal{F}}) \rangle$ is invertible, which implies that the product of these two matrices is of full rank $n-1$. From standard arguments of matrix perturbation theory [59], we can characterize the perturbation, Δ, of minimum 2-norm that can cause the resulting matrix to lose its rank

$$\min_{\Delta \in \mathbb{C}^{n-1 \times n-1}} \{ ||\Delta||_2 \mid \Delta + H \quad \text{loses rank} \} = \sigma_{\min}(H)$$

Moreover, we notice the matrix $||\langle \text{diag}(\overline{\varsigma^{\mathcal{F}}}) \rangle||_2$ is by construction a skew symmetric matrix obtained, through the operator $\langle \cdot \rangle$, from the diagonal matrix, $\text{diag}(\varsigma^{\mathcal{F}})$. Thus, by simply using the definition of the induced 2-norm, one can easily prove the following identity $||i_{\mathcal{F}}||_\infty = ||\langle \text{diag}(\overline{\varsigma^{\mathcal{F}}}) \rangle||_2$. This and the aforementioned condition proves (4.20).

The first part of Proposition 1 guarantees that small perturbations of the flat voltage condition, often encountered in distribution networks, will not make the matrix $\mathcal{M}_{\mathcal{F}\mathcal{F}}$ lose its rank. The second part provides a sufficient condition that can be checked, for the problem at hand, upon linearization.

Under the conditions of Proposition 1, it is possible to invert the matrix $\mathcal{M}_{\mathcal{F}\mathcal{F}}$, therefore explicitly expressing the relation between the endogenous and exogenous components of the grid state through a linear static map. In particular, at each stage, this relation reads:

$$z^{end} = A_{\hat{z}}^c z^{cont} + A_{\hat{z}}^w z^w + \bar{z}^{end}$$

The linearization is performed around the voltage profile, u_i, corresponding to the current prediction stage, i, and for nominal operating conditions, i.e., active and reactive power injections without any fluctuation ((r_i, w_i) is set to zero). Finally, by assembling the linear static map along the horizon we obtain:

$$\mathbf{z}^{end} = \mathbf{A}_{\hat{z}}^c \mathbf{z}^{cont} + \mathbf{A}_{\hat{z}}^w \mathbf{z}^w + \bar{\mathbf{z}}^{end} \tag{4.22}$$

Having such an explicit relation is particularly useful as it allows one to apply standard robust optimization techniques in order to obtain a tractable reformulation of Problem 2.

4.6.2   Parametrization of the uncertainty sets

In this section, two methods for the construction of the uncertainty set Ξ are presented.

In the first method, the uncertainty set is obtained by directly considering past realizations of the uncertainty, $(\mathbf{z}^{w,(k)}, \mathbf{r}^{(k)})$. The analytic description of the resulting set, Ξ_{ts}, is then simply:

$$\Xi_{\text{ts}} = \left\{ (\mathbf{z}^{w,(k)}, \mathbf{r}^{(k)}), \ k = 1, \dots, N_{\text{scen}} \right\} \tag{4.23}$$

As it has been highlighted in the literature [115, 125], constructing the uncertainty set using scenarios allows one to capture the potential correlation between the uncertain components.

For the second method, a data-driven approach is also used to construct Ξ. First, we focus on the construction of the set, $\mathcal{R}$. It is known *a priori* that the normalized tracking request has a magnitude of at most one. Therefore, we have $\mathcal{R} \subseteq \mathcal{B}_\infty(1) = \{\mathbf{r} \,|\, \|\mathbf{r}\|_\infty \leq 1\}$. However, using the uncertainty set $\mathcal{R} = \mathcal{B}_\infty(1)$ disregards statistical information about the uncertainty available under the form of scenarios. An approach inspired by [89] can be used for uncertainty set design. The key idea is to use the scenarios to fit the uncertainty set in an optimization problem. In particular, we focus on the integral of the tracking signal which represents the energy content the resources will need to absorb in order to provide the service. The intuition behind this choice is that, this is a key factor in determining if the resources can actually support the tracking signal.

We propose to build the set, $\mathcal{R}$, as follows:

$$\mathcal{R}_{\text{ms}} = \{\mathbf{r} \,|\, \|\mathbf{r}\|_\infty \leq 1, \ \|\tilde{r}_{[1 \to N]}\|_\infty \leq s_{\max}\} \tag{4.24}$$

where $\tilde{r}_{[1 \to i]}$ denotes the cumulative sum of the signal $\mathbf{r}$ up to time i ($\tilde{r}_{[1 \to i]} = \sum_{t=1}^{i} r_t$). $s_{\max}$ is chosen as the worst-case cumulative sum on a number N_{scen} of previously observed realizations.

Regarding the modeling of $\mathcal{Z}^w$, we first notice that, in most practical cases, the components of z_i^w are correlated so that its evolution along the prediction horizon can be described in a lower-dimensional space, $\mathbb{R}^{n_r}$:

$$\mathbf{z}^w = \mathbf{T}\mathbf{w}$$

where the matrix $\mathbf{T}$ effectuates the lifting into the higher dimensional space, $\mathbb{R}^{n_w}$, and the components of w_i are now uncorrelated and assumed to belong to the composed uncertainty set:

$$\mathcal{W}_{\text{ms}} := \{\mathcal{W}^1 \times \dots \times \mathcal{W}^{n_r}\}$$

where $\mathcal{W}^t$ is a polytopic set for each t.

Clearly, the choice of $\mathbf{T}$ and the sets $\mathcal{W}^t$ depends on the particular application at hand. Generally speaking, they should be designed in order to contain the actual realization of the uncertainty with high confidence and different methods are present in the literature [10, 89]. We refer the reader to the simulation Section 4.7 for a possible instance of $\mathcal{W}_{\text{ms}}$.

Remark 10. *The modeling of the uncertain grid state, z^w, as driven by a lower dimensional disturbance w is motivated by many practical situations. To provide a possible example, let us consider a network with two neighboring PV installations. Given their geographical proximity, the uncertainty related to the PV production of the two installations will likely be correlated and it will depend on some shared external conditions such as, e.g., the cloud coverage.*
We want to emphasize that the proposed framework has no restrictions and, in the most general scenario, the components of z_w can be modeled as uncorrelated so that $w \in \mathbb{R}^{n_w}$ and $\mathbf{z}^w = \mathbf{I}\mathbf{w}$.

Finally, the composed uncertainty set, Ξ_{ms}, is described by:

$$\Xi_{\mathrm{ms}} = \{(\mathbf{T}\mathbf{w}, \mathbf{r}) \mid \mathbf{w} \in \mathcal{W}_{\mathrm{ms}}, \ \mathbf{r} \in \mathcal{R}_{\mathrm{ms}}\} \tag{4.25}$$

4.6.3 Control Policies and intraday

In this section, we propose two different methods, to parametrize the control policies (4.13) and, thus, obtain a tractable reformulation of the planning problem (2).

Two-stage approximation

This solution method approximates the planning problem 2 by a two-stage robust optimization problem [125] that uses Ξ_{ts} as the uncertainty set. Thus, instead of having multi-stage control policies parametrized by $\mathbf{u}^j = \pi_{u^j}(\mathbf{z}^w, \mathbf{r})$ and $\mathbf{D} = \bar{\mathbf{D}} + \pi_{\mathbf{D}}(\mathbf{z}^w, \mathbf{r})$, a two-stage control policy is considered. This implies that the causality requirements are relaxed and it is assumed that the uncertainty is revealed for the whole prediction horizon at once. More precisely, the first stage variables are represented by the power flexibility, γ, and the baseline profile, $\bar{\mathbf{D}}$, which are decided before the realization of the random parameters $(\mathbf{z}^w, \mathbf{r})$. On the other hand, the second-stage decisions are represented by $\mathbf{u}^j$ and $\mathbf{D}$ and they can be chosen after the uncertainty has been fully revealed.

As an example, if we focus on the input realization for a single resource, $\mathbf{u}^j$, considering a two-stage structure allows the optimizer to select different input trajectories, denoted as $\mathbf{u}^{j,(k)}$, for each possible realization of the uncertainty contained in Ξ_{ts}. A similar approach could also be considered to parametrize the intraday control policy, $\pi_{\mathbf{D}}$. However, as underlined in [118], considering this additional degree of flexibility in the problem formulation actually comes at a cost. In fact, due to the relaxation of the causality requirement, the solution of the approximated problem can potentially overfit the uncertainty realizations giving rise to unrealistic behaviours. As a consequence, we decided herein to separate the intraday control policy from the planning problem and to optimize it independently. In particular, we considered the offline causal policy firstly introduced in [118] which is reported in the following.

The tracking request, $\mathbf{r}$, received by the TSO might exhibit considerable bias in either direction over short periods of time. When dealing with energy-constrained resources, such as BESSs or buildings, these energy biases could have a detrimental effect on the resources causing a reduction

in the offerable flexibility, γ. Thus, the focus of the intraday policy is to reduce such biases by modifying the future baseline. As a consequence, rather than depending on both uncertain parameters, $(\mathbf{z}^w, \mathbf{r})$, the policy is just a function of $\mathbf{r}$, i.e., $\pi_{\mathbf{D}}(\mathbf{z}^w, \mathbf{r}) = \pi_{\mathbf{D}}(\mathbf{r})$.

We first define the *residual tracking signal*, $\mathbf{a}$, as the sum of the received tracking signal, $\mathbf{r}$ and the intraday transaction, $\boldsymbol{\tau}$:

$$\mathbf{a} := \mathbf{r} + \boldsymbol{\tau}$$

where $\boldsymbol{\tau}$ is the normalized intraday transaction corresponding to the normalized tracking request, $\mathbf{r}$. The total intraday transaction is then given by $\gamma\boldsymbol{\tau}$.

The intraday policy is then given by:

$$\pi_{\mathbf{D}}^{\mathrm{ts}}(\mathbf{r}) := \left\{ \boldsymbol{\tau} \left|
\begin{array}{ll}
\tau_{i+\Delta_{\mathrm{intra}}} = \\
\underset{\tau_{i+\Delta_{\mathrm{intra}}}}{\operatorname{argmin}} & | \ \tilde{a}_{[0\to i]} + \tau_{i+\Delta_{\mathrm{intra}}} + \mathbb{E}_r \left\{ \tilde{a}_{[i\to i+\Delta_{\mathrm{intra}}-1]} + r_{i+\Delta_{\mathrm{intra}}} \right\} \ | \\
\mathrm{s.t.} & \tilde{a}_{[0\to i]} = \tilde{a}_{[0\to i-1]} + \tau_i + r_i \\
& \tau_0 = 0, \tilde{a}_{[0\to -1]} = 0 \\
& \forall i = 0, ..., N-1.
\end{array}
\right. \right\} \tag{4.26}$$

where Δ_{intra} denotes the number of steps, in the chosen sampling time, for which the baseline is fixed according to the specific market rules (*e.g.* 45 minutes in current Swiss regulations).

Essentially, the objective of the intraday control policy is to reduce the energy content of the residual tracking signal, $\mathbf{a}$. Thus, the policy is obtained by minimizing for each time instant, i, the expected absolute value, in Δ_{intra} sampling times, of the cumulative sum (energy content) of $\mathbf{a}$. This is achieved through the following steps: 1) measure the current energy content, $\tilde{a}_{[0\to i]}$, 2) choose a control action, $\tau_{i+\Delta_{\mathrm{intra}}}$, such that the expected energy content of the signal in the future is minimized. At time i the uncertainty, and for the next Δ_{intra}, the tracking signal has not yet been realized. As a consequence, the policy minimizes the energy content in expectation. In particular, the expected value can be estimated at each step using scenarios of the tracking signal, $\mathbf{r}$.

Once the PF linearization and the intraday policy are fixed, the planning problem 2 is approximated by the two-stage optimization problem resulting in the following:

Problem 3 (Two-stage approximation).

$$\underset{\gamma,\bar{\mathbf{D}},\mathbf{u}^{j,(k)}}{\text{minimize}} \quad c_{\text{energy}}^T \bar{\mathbf{D}} - c_{\text{reward}}^T \gamma$$

s.t.

$$\text{(Resource constraints)} \quad \mathbf{u}^{j,(k)} \in \mathcal{U}^j(\hat{x}^j, \hat{\mathbf{d}}^{\mathbf{j}}) \quad \forall j = 1, \dots, N_{cont} \tag{4.27}$$

$$\text{(Linearized PF)} \quad \mathbf{z}^{end,(k)} = \mathbf{A}_{\hat{z}}^c \mathbf{z}^{cont,(k)} + \mathbf{A}_{\hat{z}}^{w,(k)} \mathbf{z}^w + \bar{\mathbf{z}}^{end} \tag{4.28}$$

$$\text{(Grid constraints)} \quad \mathbf{z}^{end,(k)} \in \mathcal{Z} \tag{4.29}$$

$$\text{(Apparent power controllable res.)} \quad \mathbf{z}^{cont,(k)} = \Lambda \mathbf{U}^{(k)} \tag{4.30}$$

$$\text{(Power tracking)} \quad \Gamma \mathbf{z}^{cont,(k)} = \bar{\mathbf{D}} + \gamma \pi_{\mathbf{D}}^{\text{ts}}(\mathbf{r}^{(k)}) + \gamma \mathbf{r}^{(k)} \tag{4.31}$$

$$\text{(Power flexibility)} \quad \gamma \geq 0 \tag{4.32}$$

$$\text{(Scenarios)} \quad \forall (\mathbf{z}^{w,(k)}, \mathbf{r}^{(k)}) \in \Xi_{\text{ts}}. \tag{4.33}$$

where, as already underlined, an implicit parametrization of the second-stage variables, $\mathbf{u}^j$ is defined by having separate trajectories, $\mathbf{u}^{j,(k)}$, for each considered scenario. Please also notice how the pre-determined intraday policy in (4.31), which was determined for the normalized tracking signal, $\mathbf{r}$, is scaled by the capacity γ in order to obtain the total intraday transaction.

The resulting optimization problem is a linear program and can, therefore, be solved efficiently using available optimization software.

Multi-stage approximation

A second approximation scheme is proposed in this section. Its main advantage with respect to the two-stage approximation lies in its capability to retain the multi-stage structure of the original problem. In particular, to reduce the infite-dimensional decision space into a finite dimensional one, we consider *affine feedback policies* which offer a nice trade-off between performance and computational properties [6].

Thus, we define $\mathbf{u}^j := \mathbf{M}_r^j \mathbf{r} + \mathbf{M}_w^j \mathbf{w} + \mathbf{m}^j$ and $\mathbf{D} := \bar{\mathbf{D}} + \mathbf{K}_r \mathbf{r} + \mathbf{K}_w \mathbf{w}$ where, in order to ensure causality, we impose the following constraints on the policies (for both r and w):

$$\begin{aligned} M_{[k,l]}^j &= 0 \text{ for } l > k \qquad\qquad j = \{1, \dots, N_{\text{cont}}\} \\ K_{[k,l]} &= 0 \text{ for } l > k - \Delta_{\text{intra}} \end{aligned} \tag{4.34}$$

where, once again, Δ_{intra} is imposed by the current market regulations.

The uncertainty set considered in this solution method is Ξ_{ms}, and the reformulated problem takes the form:

Problem 4 (Multi-stage approximation)**.**

$$\underset{\gamma, \bar{\mathbf{D}}}{\text{minimize}} \quad c_{\text{energy}}^T \bar{\mathbf{D}} - c_{\text{reward}}^T \gamma$$

$$\text{s.t.}$$

(Resource constraints)	$\mathbf{u}^j \in \mathcal{U}^j(\hat{x}^j, \hat{\mathbf{d}}^{\mathbf{j}}) \quad \forall j = 1, \ldots, N_{cont}$	(4.35)
(Linearized PF)	$\mathbf{z}^{end} = \mathbf{A}_{\hat{z}}^c \mathbf{z}^{cont} + \mathbf{A}_{\hat{z}}^w \mathbf{z}^w + \bar{\mathbf{z}}^{end}$	(4.36)
(Grid constraints)	$\mathbf{z}^{end} \in \mathcal{Z}$	(4.37)
(Apparent power controllable res.)	$\mathbf{z}^{cont} = \Lambda \mathbf{U}$	(4.38)
(Power tracking)	$\Gamma \mathbf{z}^{cont} = \mathbf{D} + \gamma \mathbf{r}$	(4.39)
(Input Control Policy)	$\mathbf{u}^j = \mathbf{M}_r^j \mathbf{r} + \mathbf{M}_w^j \mathbf{w} + \mathbf{m}^j$	(4.40)
(Intraday Control Policy)	$\mathbf{D} = \bar{\mathbf{D}} + \mathbf{K}_r \mathbf{r} + \mathbf{K}_w \mathbf{w}$	(4.41)
(Power flexibility)	$\gamma \geq 0$	(4.42)
(Uncertainty)	$\forall (\mathbf{z}^w, \mathbf{r}) \in \Xi_{\text{ms}}.$	(4.43)

With the considered restrictions on the control policies (4.40), (4.41), the reformulation of the PF constraint through the linear approximation (4.36), and the considered polytopic uncertainty set (4.43), the robust optimization Problem 4 can be cast as a linear program using standard dualization arguments [6]. The derivation of the dualized problem can be found in the Appendix to this chapter.

The solution of the resulting problem is represented by the baseline profile, $\bar{\mathbf{D}}^*$, the power flexibility offered to the TSO γ^*, and the affine feedback policies $(\mathbf{M}_r^j)^*$, $(\mathbf{M}_w^j)^*$, $\mathbf{K}_\mathbf{r}^*$, $\mathbf{K}_\mathbf{w}^*$, $(\mathbf{m}^j)^*$.

4.7 Simulations and Results

This section presents the simulation results to demonstrate the effectiveness of the proposed framework. In particular, the solution method (4) was considered for the following reasons: 1) it retains the original multi-stage structure and the causality requirement with a limited sacrifice of performance, 2) it directly optimizes over the intraday policy as opposed to the two-stage approximation where $\pi_{\mathbf{D}}^{\text{ts}}$ has to be pre-determined, 3) the policy parametrization (4.41) allows one to include the effect of $\mathbf{z}^w$ otherwise ignored in (4.26).

The reader is referred to Chapter 5 for an in-depth comparison between the two approximation schemes on a specific case study.
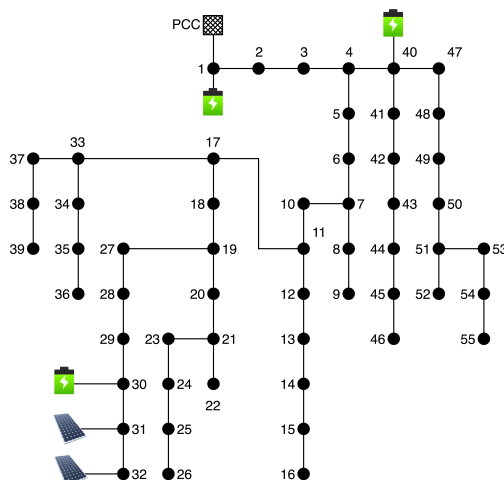
Figure 4.3 – Scheme of the 56 buses test feeder.

4.7.1 Simulation Setup

We consider a symmetric and balanced distribution network inspired by the IEEE 123 test feeder consisting of 56 buses. The testbed has been augmented in order to obtain an interesting set of simulations and highlight the capabilities of the framework. Figure 4.3 shows the topology of the network and the location of the PV plants and BESSs.

The power injection profile at each bus is randomly selected among four heterogeneous building profiles which are taken from the reference database of the U.S. Department of Energy [36]. To equip the network with stochastic renewable generation, we consider two relatively large and identical PV installations which are placed at nodes 31 and 32. The active power injection of the PV is assumed to be uncertain. In particular, given their geographical proximity, the uncertainty associated with the PVs production is constructed as being linearly dependent on the solar radiation which is assumed to belong to a time-varying interval around the mean prediction. The resulting uncertainty set for the active power production of each PV is shown in Figure 4.4. The AGC requests are modeled through the uncertainty set $\mathcal{R}$ that is constructed, as detailed in (4.24), starting from the 365 historical daily realizations of the AGC signal for the year 2013. To have an out-of-sample validation of the reliability of the obtained solution, all simulations are performed on a different set of scenarios corresponding to the year 2014. All the statistics regarding the load and generation profiles can be found in Table 4.2. Finally, three BESSs of different sizes are connected to the network at nodes 1, 30, and 40, respectively. As it will be apparent in the following, the BESSs are strategically placed in the network so as to accomplish both local as well as global objectives. The technical specifications of the controllable resources can be found in Table 4.3.

4.7.2 Service Provision

The set of services that are considered in the simulation are summarized in the following.

Type	Max injection [kW]	# Units connected
Warehouse	-60	11
Medium office	-118	12
Small office	-11	27
PV	1000	2

Table 4.2 – Technical specifications of the connected loads and distributed generation.
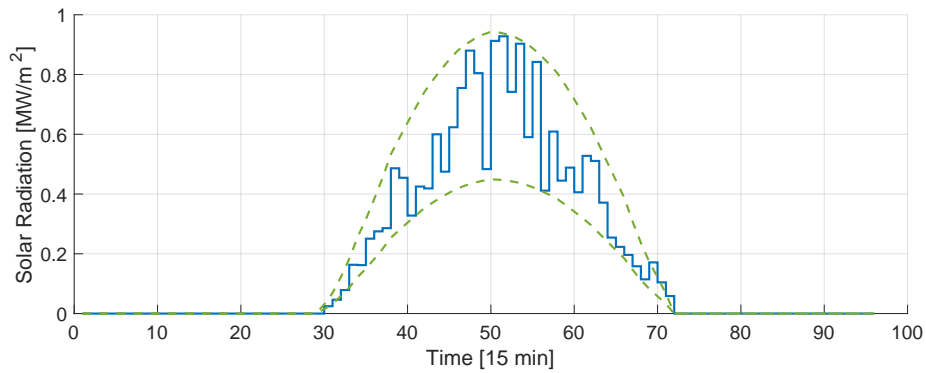
Figure 4.4 – Uncertainty set for the active power production of each PV for the next 24 hours of operation. The time-varying uncertainty bounds are displayed with dashed green lines whereas a possible realization of the uncertainty is displayed with a solid blue line.

Resource	Capacity [kWh]	Inverter limit [kVar]	R_{BESS} [Ω]
BESS node 1	1000	500	0.012
BESS node 30	800	600	0.014
BESS node 40	400	200	0.005

Table 4.3 – Technical specifications of the controllable BESSs.
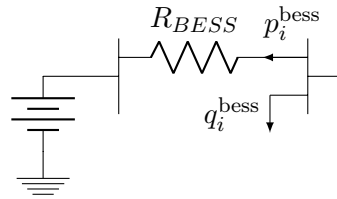
Figure 4.5 – BESS resistance-based model.

- *Voltage regulation*: The voltage magnitude, v_i^h, at each bus and time instant is constrained to satisfy $0.9\,p.u. \leq v_i^h \leq 1.05\,p.u.$. Even though the voltage constraint is enforced across the whole network, this service is of a local nature as the most critical nodes are expected to be in the neighborhood of the two PV installations (node 31 and 32).

- *Peak shaving*: A constraint on the maximum active power injection is considered for a specific selected bus (node 40). The constraint reads $|p_i^{40}| \leq \bar{p}^{40} \quad \forall i$, where $\bar{p}^{40}$ is a predetermined threshold for the local injection.

- *Secondary frequency regulation*: The controllable resources are aggregated into a virtual entity that acts as an ancillary service provider. As described in Section 4.1, the virtual entity receives, during real-time operation, a TSO request (called AGC signal) which determines the deviation (scaled by the offered flexibility) of the injected active power from the contracted baseline. To facilitate the provision of the service for energy-constrained resources such as batteries, it was assumed that a time-varying flexibility can be offered. More precisely, the optimizer is allowed to vary the offered capacity every 6 hours.

4.7.3   Resource Modeling

The BESSs are modeled by means of the resistance-based model proposed in [129], where all losses are condensed into a single resistance, R_{BESS}. Thus, it is possible to represent the battery with the equivalent model shown in Figure 4.5.

The model is comprised of an ideal battery (unitary round-trip efficiency) connected to a virtual node. The virtual node is then connected to the physical system by means of a purely resistive line which models the losses of the BESS. As it can be noted in Figure 4.5, the reactive power injection

is connected directly to the physical node of the network rather than the virtual one. This is due to the fact that the reactive power depends only on the property of the power conversion. The considered modeling choice is particularly appealing in the proposed framework as it allows one to get rid of the non-linearities associated to the BESS models by integrating the equivalent model directly into the admittance matrix of the network (which is augmented with as many virtual nodes as the number of BESSs). The resulting augmented PF equation is then linearized following the same procedure of Section 4.6.1.

After adopting the aforementioned simplification, each BESS, which is connected to its respective virtual node, can be described by means of a simple linear time-invariant system of the form:

$$\text{SoC}_{i+1} = \text{SoC}_i + p_i^{\text{bess}}$$
$$s_i^{\text{bess}} = \sqrt{(p_i^{\text{bess}})^2 + (q_i^{\text{bess}})^2} \tag{4.44}$$

Both the SoC and the apparent power injection of the battery are constrained to lay within their feasible operation range at each time instant, k:

$$\text{SoC}_{\text{min}} \leq \text{SoC}_i \leq \text{SoC}_{\text{max}} \tag{4.45}$$

$$||s_i^{\text{bess}}||_2 \leq s_{\text{max}}^{\text{bess}} \tag{4.46}$$

where, referring to the general model of Section 4.3, we have now defined $u_i^{\text{bess}} := [p_i^{\text{bess}} \ q_i^{\text{bess}}]$, $B = [1 \ 0] \ C = 1$

Remark 11. *Please note that the power inverter constraint* (4.46) *is expressed as a non-linear map of the decision variables q_i^{bess} and q_i^{bess}. This, however, does not pose any problem as it can be easily approximated (to arbitrary accuracy) by means of a polytopic inner approximation [5], therefore, allowing one to retain the linearity of the approximated optimization problem.*

4.7.4 Computations

All simulations are performed in MATLAB. The original Planning Problem 2 is approximated as an LP as described in Section 4.6 with condition (4.20) empirically found to be respected at all stages. The resulting optimization problem is then formulated and solved in MATLAB using the parsing tool YALMIP [80]. After examining the optimal solutions for smaller instances of the problem, it was found that limiting the recourse for the feedback policies to 6h does not affect significantly the optimality of the solution while considerably reducing the computational demand. The resulting optimization problem is solved with GUROBI on a 3.4 Ghz intel Core i7 with an average solving time of 280s.

4.7.5 Results

To demonstrate the capability of the proposed framework to adapt to different levels of renewable generation, we present herein three possible scenarios.
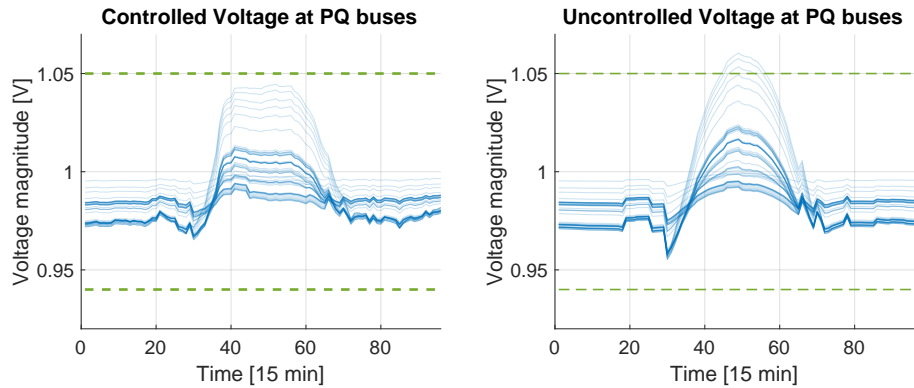
Figure 4.6 – Left: Voltage magnitude in the uncontrolled setting. Right: Voltage magnitude in the controlled setting for a randomly selected (out-of-sample) AGC signal and the worst-case scenario for the PV injection (maximum injection). In both subplots, each thin line corresponds to the voltage magnitude profile at a specific node which is obtained by solving the exact PF equations (with MATPOWER).

The first set of simulations (Scenario A) corresponds to a *summer day* and it is characterized by a significant power injection from the PV installations. On the contrary, the second set (Scenario B) corresponds to a *winter day* with a relatively low solar injection. The third set of simulations (Baseline scenario) provides the results for the case in which all three BESSs are operated separately.

Scenario A

Figure 4.6 shows the predicted voltage profiles for all nodes, obtained by solving the exact PF equation with the open-source Matlab simulation package MATPOWER [150]. It can be observed that the voltage profiles for the nodes in the proximity of the PVs exceed the allowed limits if no control is implemented.

The results in the controlled setting are depicted in the left-hand side plot of Figure 4.6 which corresponds to a randomly selected AGC signal and the worst-case scenario for the PV injection (maximum injection). Thus, by solving Problem 2, the aggregator can allocate enough control power so as to guarantee the respect of all network constraints. In particular, in the uncontrolled case, the predicted PV and load injection profiles are considered with the power injection of the BESSs at zero. On the contrary, in the controlled case, the BESSs power injections are set by solving Problem 2 and by computing their evolution for a particular realization of the uncertainty.

Figure 4.7 depicts the active power injection at node 40. As it can be observed, the original load profile at the node (solid brown line) exceeds the imposed peak limit. On the contrary, the combined profiles (BESS + load) for all considered scenarios are well within the limits. Moreover, by looking at the distribution of the active power profiles, it can be noticed that, whenever feasible, the remaining available controllable power of the BESS is exploited to complement the other BESSs
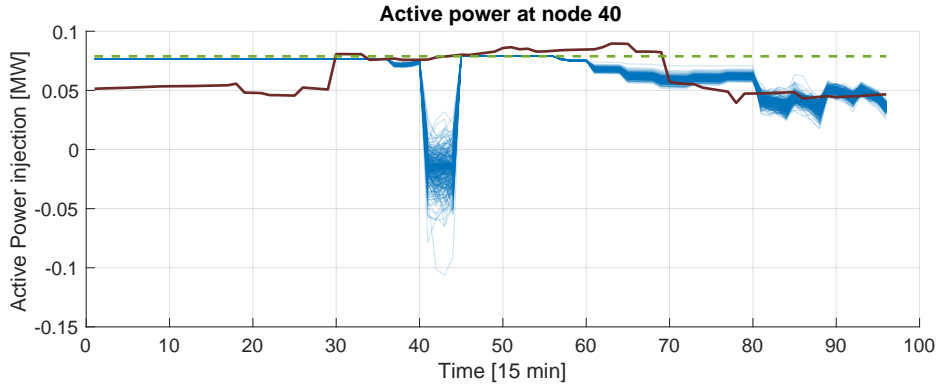
Figure 4.7 – Active power profile at a selected node. The brown solid line represents the active power injection of the load in the absence of the BESS. Each thin blue line corresponds to the profile obtained in the controlled setting and corresponding to 365 randomly selected uncertainty scenarios (AGC + PV).

to offer the frequency control service. This is particularly visible during the central hours of the day, and towards the end of the day, when each uncertainty realization results in a different aggregated power profile. In fact, during these periods, there is enough slack between the aggregated power and the peak limit. Consequently, the active power injection of the BESS, that otherwise would not be fully exploited, can be modulated to contribute to the provision of SFC.

The offered electrical flexibility, γ^*, the baseline consumption, $\bar{\mathbf{D}}^*$, the out-of-sample AGC scenarios, and the energy trades placed in the intraday market, are depicted in Figure 4.8. Moreover, in order to account for typical delays imposed by market rules, the baseline profile is kept fixed for the first eight steps (2h) as it can be observed in the second subplot of Figure 4.8.

A single realization of the AGC signal together with the corresponding intraday trades is depicted in Figure 4.9 which displays also how the controller coordinates the three BEESs to precisely tracking the received AGC signal. Also, notice how the intraday trades are used as a mechanism to reduce the large energy bias of the AGC that would otherwise result in a net discharge of the BESSs. Finally, the contracted capacity is reported in the bottom-most plot. As it can be observed, the Planning Problem reduces the offered flexibility in the central hours of the day so to satisfy the local constraints. The capacity is then restored towards the end of the day when local constraints are no longer active.

Finally, the SoC trajectories together with active and reactive power injection of all controllable resources are depicted in Figure 4.10 which shows how the solution of the Planning Problem optimally allocates the available power to manage both local and global constraints while simultaneously managing all operational constraints of the controlled resources. Examining the profiles of the three BESSs a few observations are in order. 1) Even though the voltage regulation is managed prin-
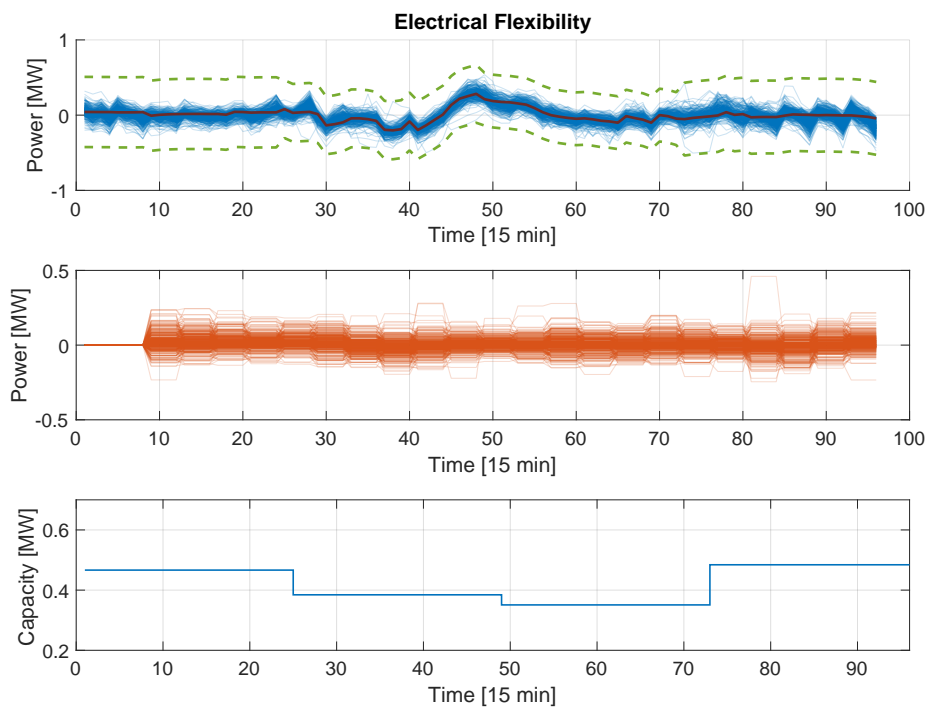
Figure 4.8 – Top: Baseline power profile (solid brown line), offered symmetric electrical flexibility (two dashed green line around the baseline), out-of-sample AGC scenarios scaled by the contracted flexibility (blue thin lines). Middle: Intraday trades corresponding to the worst-case PV injection scenario and to the 365 out-of-sample AGC scenarios. Bottom: Contracted electrical flexibility.
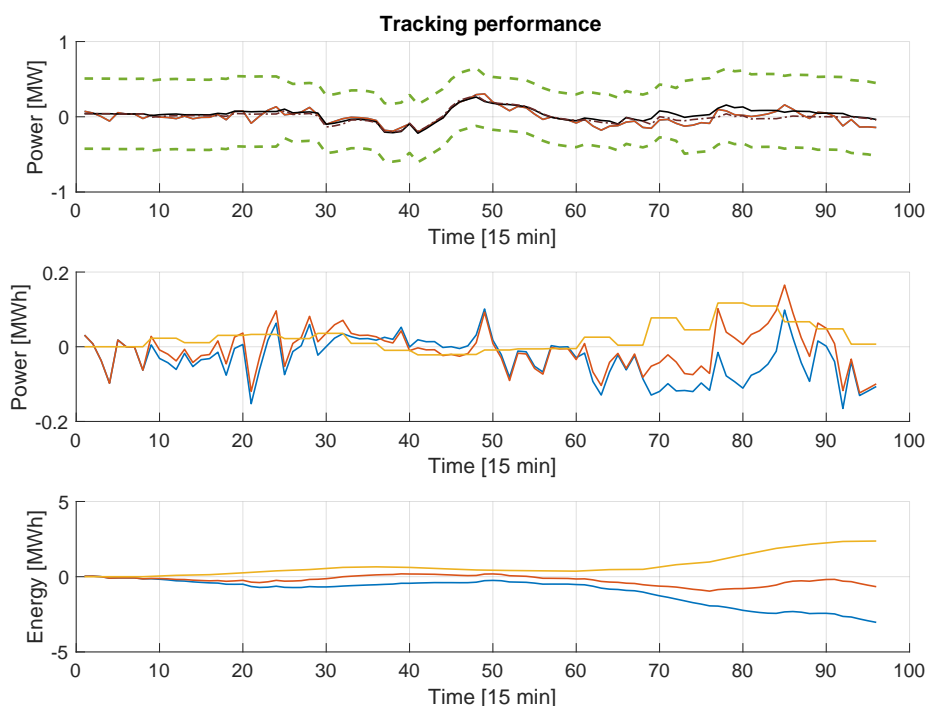
Figure 4.9 – Tracking for one randomly selected AGC scenario. Top: The original baseline power profile (dashed brown line), the final baseline profile after the intraday trades (solid black line), the offered symmetrical electrical flexibility around the final baseline (dashed green lines), the resulting tracking signal (computed as $\bar{\mathbf{D}}^* + \mathbf{K}_r^* \mathbf{r}_{\text{sel}} + \mathbf{K}_w^* \mathbf{w}_{\text{sel}} + \gamma^* \mathbf{r}_{\text{sel}}$ where $\mathbf{r}_{\text{sel}}$ and $\mathbf{w}_{\text{sel}}$ represent the selected scenarios for the AGC and the PV injection, respectively), and the aggregated power consumption of the controllable resources (solid blue). Middle: The AGC signal (solid blue), the intraday trades (solid yellow), and the filtered AGC signal, *i.e.*, AGC + intraday trades (solid red). Bottom: The cumulative energy of the original AGC signal (solid blue), the filtered AGC signal after intraday trades (solid red), and the intraday transactions (solid yellow).
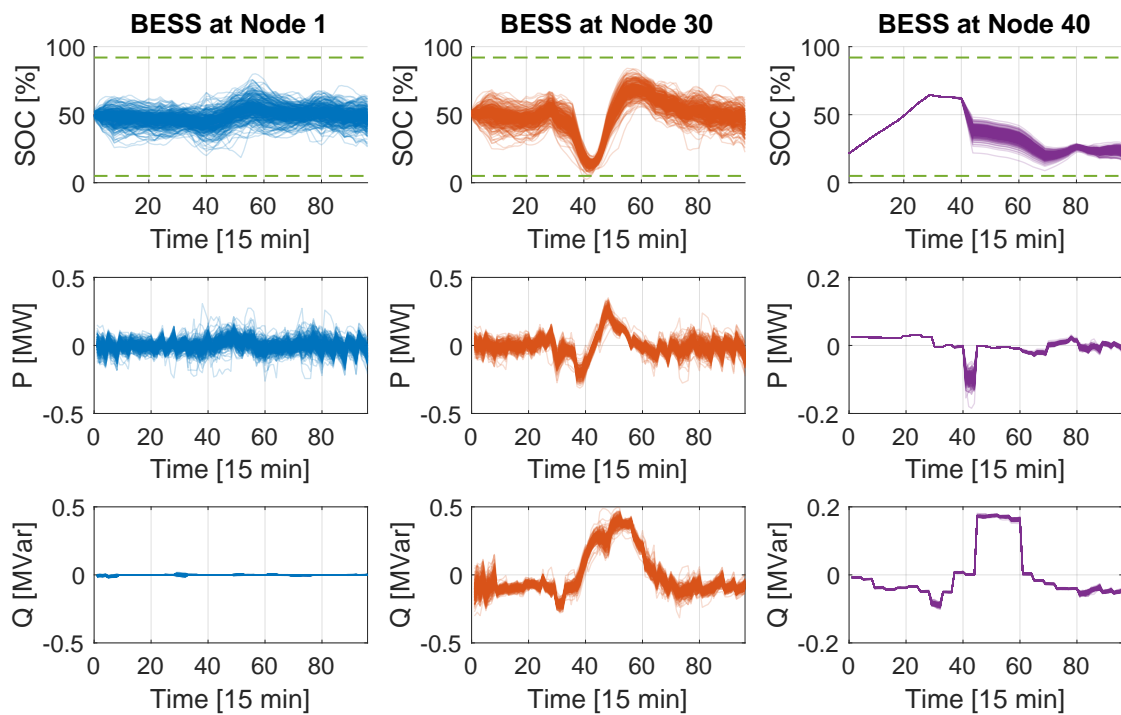
Figure 4.10 – **Scenario A**: Time evolution of the three controllable resources. First row: SoC trajectories. The dashed green lines represent SoC constraints. Second row: Active power injections. Third row: Reactive power injections. In all plots, each line corresponds to a different realization of the uncertainty (AGC + PV).
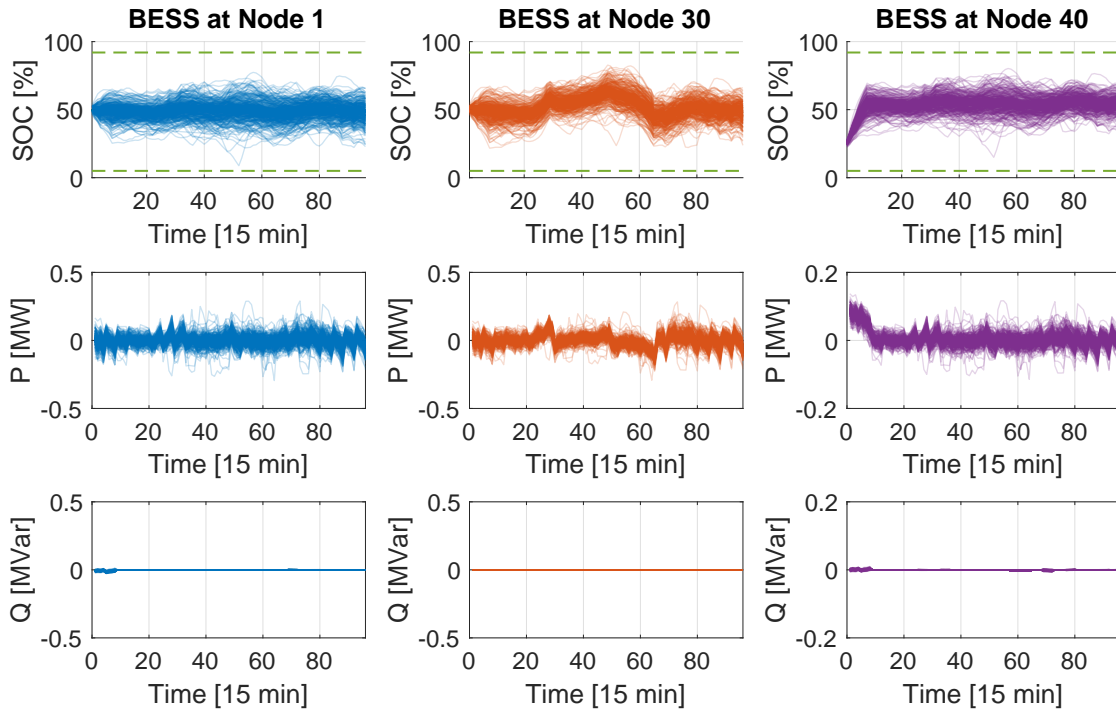
Figure 4.11 – **Scenario B**: Time evolution of the three controllable resources. First row: SoC trajectories. The dashed green lines represent SoC constraints. Second row: Active power injections. Third row: Reactive power injections. In all plots, each line corresponds to a different realization of the uncertainty (AGC).

cipally by the BESS placed at node 30, the BESS at 40 contributes during the central hours of the day by modulating its reactive power injection that would be otherwise not exploited. 2) As expected, the BESS placed at node 1 is the one contributing the most to the SFC since it is not subject to stringent local constraints as is the case for the BESS at node 30 (voltage regulation) and the one at node 40 (peak shaving). Nonetheless, both BESS 30, as well as BESS 40, are used, whenever possible, to increase the aggregated flexibility. 3) As the initial SoC of the BESS at node 40 is too low to absorb the peak consumption, the Planning Problem appropriately computes a baseline consumption, $\bar{\mathbf{D}}$, so as to bring the SoC to a sufficient charging level.

Scenario B

The second scenario is characterized by a relatively low PV injection so that, even in the uncontrolled case, the network does not experience overvoltages. Moreover, we also consider that no peak power limits are enforced at node 40. Thus, as opposed to scenario A, the BESSs at node 30 and 40 can now be fully used to provide SFC. This is displayed in Figure 4.11 where both the SoC and the active power injections of the three BESSs reveal how the Planning Problem successfully adapts
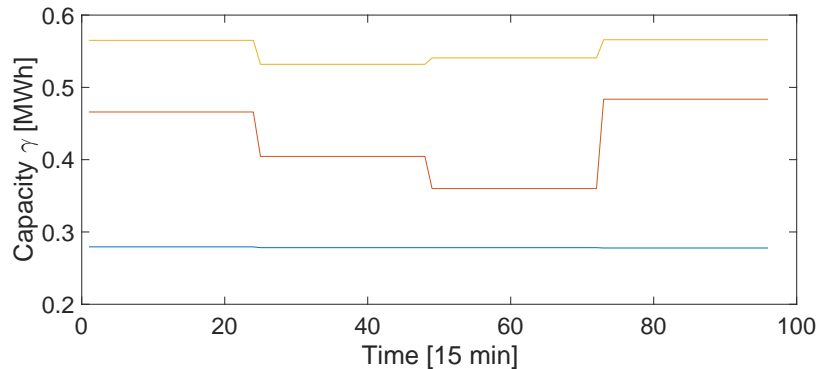
Figure 4.12 – Offered flexibility, γ for the three considered scenarios. For scenario A (red solid line), for scenario B (yellow solid line), and for the case of BESS at node 1 providing the service alone (blue solid line).

the solution depending on the level of local services required. As a comparison metric between the different scenarios, we considered the flexibility that was offered on the ancillary service market.

Baseline scenario

To better compare the results, we also computed γ for a **baseline scenario** in which the three BESSs are operated separately. More precisely, BESS 1 is the only resource providing SFC whereas BESS 30 and BESS 40 are used solely to provide their respective local services. The results are shown in Figure 4.12. Thanks to the resource aggregation, in both scenario A and B, the offered capacity is significantly increased with respect to the baseline scenario. Also, by adapting to the level of renewable generation and local services, the full exploitation of the resources can be achieved in both scenarios. In fact, the average value of γ for scenario B is equal to 0.58 MW which represents a 26% increment with respect to the average capacity of scenario A (0.43 MW).

4.8 Conclusions

In this chapter, we investigated the problem of scheduling the provision of multiple services to the grid using a set of controllable resources integrated into an active distribution network.

A general formulation of the planning problem as a multi-stage uncertain optimization problem was proposed. The problem directly accounts for the uncertainty acting on the system in order to coordinate the resources and satisfy both local as well as network constraints. Additionally, the problem also maximizes the amount of fast regulation services that can collectively be offered by the resources.

To obtain a tractable reformulation of the original non-convex infinite-dimensional problem, a

linearization of the PF constraints was considered. The original linearization provides only an implicit relation between controlled (injections of the controllable resources) and derived variables (currents and voltages). A thorough analysis of the properties of such a linearization provided formal conditions under which the relation can be explicitly stated. This enabled us to apply standard robust-optimization results and reformulate the original planning problem as a tractable linear programming.

An extensive simulation study was performed on the IEEE 123 feeder equipped with stochastic renewable generation, heterogeneous demand, and three controllable electrical batteries. The study demonstrated the ability of the proposed control framework to adapt to different levels of uncertainty affecting the network in order to maximize the exploitation of the available power. The method was tested against a large number of out-of-sample uncertainty scenarios (tracking signal + renewable production), and showed satisfactory performance in guaranteeing the provision of both local and shared services, while satisfying the operational constraints on the resources.

Appendices

Appendix B

Derivation of the transition matrices

In this section, we provide the explicit expression of the transition matrices, Γ and Λ appearing in the planning problem 2.

Derivation of Γ

We start by restating the definition of $\mathbf{z}^{cont}$:

$$\mathbf{z}^{cont} := \begin{bmatrix} z_1^{cont} \\ \vdots \\ z_N^{cont} \end{bmatrix}$$

where each component along the prediction horizon is defined as $z_i^{cont} = [p_i^1, q_i^1, \ldots, p_i^{N_{cont}}, q_i^{N_{cont}}]^T$.

Thus, the total active power injection at time, i is given by:

$$\sum_{j=1}^{N_{cont}} p_i^j = \Gamma_i z_i^{cont}$$

with $\tilde{\Gamma} := \mathbf{1}_{N_{cont}} \otimes [0, 1]$. Finally, the matrix Γ is obtained as:

$$\Gamma = \mathbf{I}_{N \times N} \otimes \tilde{\Gamma}_i$$

so that $\Gamma \mathbf{z}^{cont}$ provides the total active power injection of the controllable resources at each time instant.

Derivation of Λ

For a given resource, the relation between the input vector, u_i, and the apparent power injection, s_i is given by:

$$s_i = \begin{bmatrix} p_i \\ q_i \end{bmatrix} = \begin{bmatrix} \eta^T \\ \xi^T \end{bmatrix} u_i$$

Thus, considering all resources, the term z_i^{cont} can be obtained as

$$z_i^{cont} = \begin{bmatrix} s_i^1 \\ \vdots \\ s_i^{N_{cont}} \end{bmatrix} = \underbrace{\begin{bmatrix} (\eta^1)^T \\ (\xi^1)^T \\ & \ddots \\ & & (\eta^{N_{cont}})^T \\ & & (\xi^{N_{cont}})^T \end{bmatrix}}_{\tilde{\Lambda}} \begin{bmatrix} u_i^1 \\ \vdots \\ u_i^{N_{cont}} \end{bmatrix} \tag{B.1}$$

Stacking the terms Z_i^{cont} along the prediction horizon gives:

$$\mathbf{z}^{cont} = (\mathbf{I}_{N \times N} \otimes \tilde{\Lambda})\Upsilon\mathbf{U}$$

where $\mathbf{U} = [(\mathbf{u}^1)^T, \ldots, (\mathbf{u}^{N_{cont}})^T]^T$, and the matrix Υ performs a permutation of the coordinates to preserve the structure of (B.1).

Finally, we highlight that all the expressions can be modified to accommodate the case of time-varying conversion factors, i.e., $s_i = [\eta_i, \xi_i]^T u_i$.

Appendix C

Robust Solution

C.1 Robust counterpart

In this section we present a set of results, available in the robust optimization literature [6], that are used in this thesis to reformulate robust constraints. In particular, let's consider the following

$$F\mathbf{M}\mathbf{r} \leq b \quad \forall \mathbf{r} \in \mathcal{R} \tag{C.1}$$

where the matrix F and the vector b are problem parameters and, therefore, fixed. The term $\mathbf{M}$ is the decision variable, and $\mathbf{r}$ is the uncertain parameter which is assumed to belong to a compact polytopic set $\mathcal{R}$.

$$\mathcal{R} = \left\{ \mathbf{r} \in \mathbb{R}^N \middle| S_r \mathbf{r} \leq h_r \right\}$$

The first step is to notice that the constraint (C.1) is equivalent to the following worst-case reformulation

$$\max_{\mathbf{r} \in \mathcal{R}} F\mathbf{M}\mathbf{r} \leq b \tag{C.2}$$

where the maximization is taken row-wise.

By dualizing each maximization problem, and introducing the composed dual variable, Ψ, that is obtained by stacking the dual variables corresponding to each maximization problem, one obtains the following equivalent problem

$$\min_{\Psi} \{ \Psi^T h_r \mid \Psi^T S_r = F\mathbf{M}, \quad \Psi \geq 0 \} \leq b$$

It is important to observe that the minimization can be omitted as it is sufficient that the constraint is satisfied for at least one Z. Hence, the final formulation of the dualized robust constraint reads

$$\exists \Psi \mid \Psi^T h_r \leq b, \quad \Psi^T S_r = F\mathbf{M}, \quad \Psi \geq 0$$

We highlight that, despite the fact that, in the steps above, we have exploited the polytopic assump-

tion on the uncertainty set, similar approaches can be considered for different and more general classes of uncertainty sets. We refer the reader to [6] for a complete overview.

In the following, we show how the dualization method reported above can be used to replace the universal quantifier, $\forall (\mathbf{z}^w, \mathbf{r}) \in \Xi_{\mathrm{ms}}$, in the multi-stage approximation (4).

C.2 Reformulation of the multi-stage planning problem

As a first step, we simplify the notation by stacking the variables corresponding to each controllable resource into a single variable. More precisely, we define $\mathbf{U} := [(\mathbf{u}^1)^T, \ldots, (\mathbf{u}^{N_{\mathrm{contr}}})^T]^T$, $\mathbf{M}_r := [(\mathbf{M}_r^1)^T, \ldots, (\mathbf{M}_r^{N_{\mathrm{contr}}})^T]^T$, $\mathbf{M}_w := [(\mathbf{M}_w^1)^T, \ldots, (\mathbf{M}_w^{N_{\mathrm{contr}}})^T]^T$, $\mathbf{m} := [(\mathbf{m}^1)^T, \ldots, (\mathbf{m}^{N_{\mathrm{contr}}})^T]^T$. Moreover, we explicitly characterize the constraints sets for the input variables, $\mathcal{U}$:

$$
\begin{aligned}
\mathcal{U}^{\mathrm{all}} &:= \{ \mathcal{U}^1 \times \ldots \times \mathcal{U}^{N_{\mathrm{contr}}} \} \\
&= \{ \mathbf{U} \mid \mathbf{L}\mathbf{U} \leq \mathbf{b} \}
\end{aligned}
$$

and the exogenous variables, $\mathcal{Z}$:

$$
\mathcal{Z} := \left\{ \mathbf{z}^{end} \;\middle|\; \mathbf{V}\mathbf{z}^{end} \leq \mathbf{v} \right\}
$$

With the newly introduced notation, the planning problem for the multi-stage reformulation reads:

$$
\begin{aligned}
\underset{\gamma, \bar{\mathbf{D}}}{\mathrm{minimize}} \quad & c_{\mathrm{energy}}^T \bar{\mathbf{D}} - c_{\mathrm{reward}}^T \gamma \\
\mathrm{s.t.} \quad &
\end{aligned}
$$

(Resource constraints)	$\mathbf{L}\mathbf{U} \leq \mathbf{b}$	(C.3)
(Linearized PF)	$\mathbf{z}^{end} = \mathbf{A}_{\hat{z}}^c \mathbf{z}^{cont} + \mathbf{A}_{\hat{z}}^w \mathbf{z}^w + \bar{\mathbf{z}}^{end}$	(C.4)
(Grid constraints)	$\mathbf{V}\mathbf{z}^{end} \leq \mathbf{v}$	(C.5)
(Apparent power controllable res.)	$\mathbf{z}^{cont} = \Lambda \mathbf{U}$	(C.6)
(Power tracking)	$\Gamma \mathbf{z}^{cont} = \mathbf{D} + \gamma \mathbf{r}$	(C.7)
(Input Control Policy)	$\mathbf{U} = \mathbf{M}_r \mathbf{r} + \mathbf{M}_w \mathbf{w} + \mathbf{m}$	(C.8)
(Intraday Control Policy)	$\mathbf{D} = \bar{\mathbf{D}} + \mathbf{K}_r \mathbf{r} + \mathbf{K}_w \mathbf{w}$	(C.9)
(Power flexibility)	$\gamma \geq 0$	(C.10)
(Uncertainty)	$\forall (\mathbf{z}^w, \mathbf{r}) \in \Xi_{\mathrm{ms}}.$	(C.11)

Also, for the sake of clarity, we state the polytopic parametrization of the uncertainty sets, $\Xi_{\mathrm{ms}} = \mathcal{Z}^w \times \mathcal{R}$:

$$
\mathcal{Z}^w := \{ \mathbf{T}\mathbf{w} \mid \mathbf{S}_w \mathbf{w} \leq \mathbf{h}_w \}, \quad \mathcal{R} := \{ \mathbf{r} \mid \mathbf{S}_r \mathbf{r} \leq \mathbf{h}_r \}
$$

where $\mathbf{S}_w \in \mathbb{R}^{n_{s1} \times N}$, $\mathbf{h}_w \in \mathbb{R}^{n_{s1} \times 1}$, $\mathbf{S}_r \in \mathbb{R}^{n_{s2} \times N}$, $\mathbf{h}_r \in \mathbb{R}^{n_{s2} \times 1}$.

In the following, we will demonstrate how to dualize one of the constraint appearing in the problem formulation above. In particular, we will focus on the resource constrain (C.3) but a similar approach can also be followed for the grid constraint (C.5).

On substituting (C.8) in (C.3), one obtains:

$$\mathbf{L}\mathbf{M}_r \mathbf{r} + \mathbf{L}\mathbf{M}_w \mathbf{w} + \mathbf{L}\mathbf{m} \leq \mathbf{b} \quad \forall(\mathbf{z}^w, \mathbf{r}) \in \Xi_{\mathrm{ms}}$$

Considering the dualizing procedure introduced in the previous Section C.1, one obtains the equivalent set of constraints:

$$
\begin{aligned}
\mathbf{L}\mathbf{M}_r &= \Psi_1^T \mathbf{S}_r \\
\mathbf{L}\mathbf{M}_w &= \Psi_2^T \mathbf{S}_w \\
\mathbf{L}\mathbf{h}_r &+ \mathbf{L}\mathbf{h}_w + \mathbf{L}\mathbf{m} \leq \mathbf{b} \\
\Psi_1 \geq 0, &\quad \Psi_2 \geq 0
\end{aligned}
$$

Finally, the linear equalities of tracking constraint (C.7) imposes the following additional constraints on the policies:

$$
\begin{aligned}
\mathbf{\Gamma}\Lambda\mathbf{M}_r &= \mathbf{K}_r + \gamma\mathbf{I} \\
\mathbf{\Gamma}\Lambda\mathbf{M}_w &= \mathbf{K}_w \\
\mathbf{\Gamma}\Lambda\mathbf{m} &= \bar{\mathbf{D}}
\end{aligned}
$$

The reformulated problem takes the form:

$$
\begin{aligned}
\underset{\gamma, \bar{\mathbf{D}}}{\text{minimize}} \quad & c_{\mathrm{energy}}^T \bar{\mathbf{D}} - c_{\mathrm{reward}}^T \gamma \\
\text{s.t.} \quad & \mathbf{L}\mathbf{M}_r = \Psi_1^T \mathbf{S}_r \\
& \mathbf{L}\mathbf{M}_w = \Psi_2^T \mathbf{S}_w \\
& \Psi_1^T \mathbf{h}_r + \Psi_2^T \mathbf{h}_w + \mathbf{L}\mathbf{m} \leq \mathbf{b} \\
& \Theta\mathbf{M}_r = \Psi_3^T \mathbf{S}_r \\
& \Theta\mathbf{M}_w + F = \Psi_4^T \mathbf{S}_w \\
& \Psi_3^T \mathbf{h}_r + \Psi_4^T \mathbf{h}_w + \Theta\mathbf{m} + \mathbf{V}\bar{\mathbf{z}}^{end} \leq \mathbf{v} \\
& \mathbf{\Gamma}\Lambda\mathbf{M}_r = \mathbf{K}_r + \gamma\mathbf{I} \\
& \mathbf{\Gamma}\Lambda\mathbf{M}_w = \mathbf{K}_w \\
& \mathbf{\Gamma}\Lambda\mathbf{m} = \bar{\mathbf{D}} \\
& \Psi_1 \geq 0, \ \Psi_2 \geq 0, \ \Psi_3 \geq 0, \ \Psi_4 \geq 0 \\
& \gamma \geq 0
\end{aligned}
$$

where $\Theta := \mathbf{V}\mathbf{A}_{\hat{z}}^c \Lambda$, $F := \mathbf{V}\mathbf{A}_{\hat{z}}^w \mathbf{T}$, and the dual variables, Ψ_1, Ψ_2, Ψ_3, Ψ_4 are also decision variables

with appropriate dimensions.

As already mentioned, the solution of the dualized problem provides the optimal baseline, $\bar{\mathbf{D}}^*$, the power flexibility, γ^*, the input control law defined by $\mathbf{M}_r^*$, $\mathbf{M}_w^*$, $\mathbf{m}^*$, and the intraday control policy defined by $\mathbf{K}_r^*$, $\mathbf{K}_w^*$.

Chapter 5

Comparison between the two-stage and the multi-stage approximation schemes

5.1 Introduction

In the previous chapter, we introduced a novel control framework to schedule the operation of a set of heterogeneous energy resources connected at the distribution level and offering multiple services to the power grid. The main features of the framework are: its generality as it can be easily adapted to accommodate several practical situations characterized by different resources and/or type of services, and its computational tractability as the resulting optimization problem to be solved is an LP. In particular, the latter property is a direct consequence of the proposed approximation schemes, *i.e.* multi-stage and two-stage approximations, that allow one to simplify the infinite-dimensional structure of the original planning problem while preserving the robustness of the solution against the uncertainty acting on the system.

The objective of this part is to complement the previous one by presenting a set of guidelines on which approximation scheme is the most appropriate depending on the particular situation at hand. More specifically, we do this by providing both a quantitative as well as a qualitative comparison between the two approximations.

5.1.1 Structure of the chapter

The content presented in this chapter is the result of a collaborative effort with other three Ph.D. students (Tomasz T. Gorecki, Faran A. Qureshi, and Altug Bitlislioglu) for the development of the experimental platform LADR. Thus, some of the content reported herein has already appeared in their thesis.

The rest of this chapter is organized as follows. In Section 5.2 we present the LADR platform that is used as a case study for this chapter. In Section 5.3 we introduce a simplified instance of the planning problem 2 that will allow us to easily compare the level of conservatism related to

each approximation scheme. In Section 5.4, the two tractable reformulations for the considered problem are derived. Then, a quantitative comparison is conducted by means of both simulations, Section 5.5, and experiments, Section 5.6. Finally, in Section 5.7, we draw upon the obtained numerical and experimental results, in order to summarize the main advantages and disadvantages of the two approximation schemes.

5.2 LADR experimental platform

This section introduces the platform LADR (Laboratoire d'Automatique Demand Response). The purpose of the platform is to validate the control methodologies developed in the laboratory, and various experiments were performed throughout the years [39, 41, 48]. We only provide herein a brief description of the platform, while the reader is referred to chapter 7 for more details.

5.2.1 Testbed description

We have equipped offices in the lab with wireless temperature sensors and customized electric radiators (Figure 5.1). For the sake of experiments, the original heating system consisting of hot water based radiators was switched off. In order to perform closed-loop experiments, a communication platform was developed that allows the flow of data from the sensors and the heaters to the main controlling unit, and vice-versa.

The heaters are rated at a power, $P_{\max}$, of 1950 Watts at 230 Volts, summing up to a total maximum power capacity of 7800 Watts. They are normally equipped with a thermostat and a switch to adjust the level of heating between three distinct levels. In order to be able to modulate their power consumption continuously, the heaters were customized with additional hardware that allows pulse width modulation (PWM) at 4 Hz.

The controlling algorithm (offline and online) runs in MATLAB. It retrieves the current temperatures in the rooms and it determines an appropriate activation of the heaters so as to respect the system constraints.

5.2.2 Building model and constraints

As previously mentioned, thanks to the presence of electric heaters, the model of the HVAC system in our experimental setup is relatively simple. The control input, u_i, is represented by the pulse-width modulation ratio to each heater. Therefore, $u_i \in \mathbb{U} = [0,1]^4$. This directly results in an electric power consumption which is a linear map of the control input, $i.e.$, $p_i = P_{\max} u_i$. Finally, the heaters being resistive elements, the power consumption directly translates into a thermal input, Q, to the building, $i.e.$, $Q = p = uP_{\max}$.

The thermodynamic model of the offices was identified using standard black-box linear system identification [78]. Each office was identified separately since the thermal coupling between them is
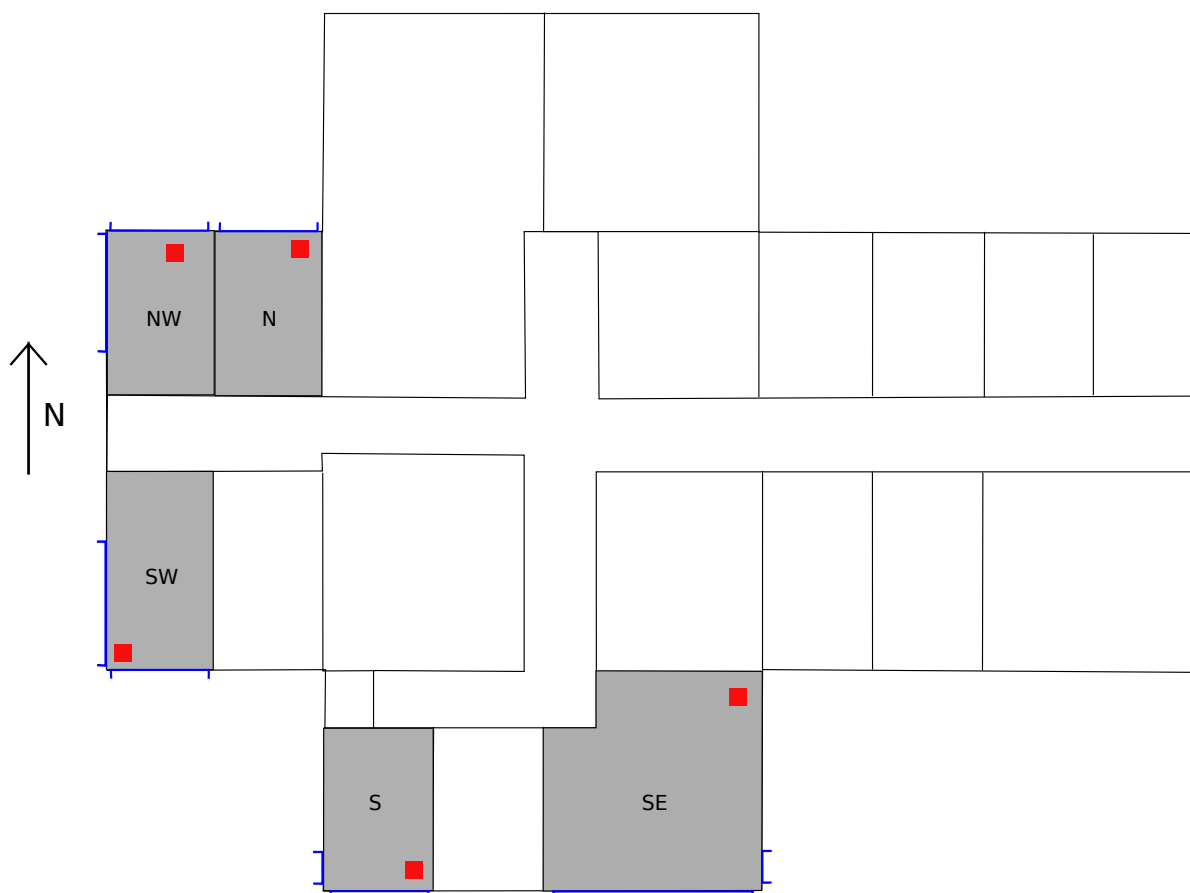
Figure 5.1 – Floor map of the offices used in LADR. The offices' shaded area have been equipped with sensors and electric radiators in order to use them for the experiments. Red squares show the position of the heating units.

weak. The model of each room (2^{nd}-order ARX models) has one control input, the thermal input, and one output, the indoor temperature. No disturbances are considered in this first experimental campaign as only night experiments are performed. In a second experimental campaign, the model was extended to account for the impact of external perturbations and, therefore, was used for daytime experiments as well. Please refer to Chapter 7.

The full model of the building is obtained by combing the individual models of the rooms and it has four control inputs (thermal input in each room) and four outputs (temperatures). It is transformed into the following state-space form as required in (4.1)

$$
\begin{aligned}
x_{i+1} &= Ax_i + B_u u_i \\
y_i &= Cx_i
\end{aligned}
\tag{5.1}
$$

where we recall that $u = \frac{Q_i}{P_{\max}}$, with Q_i the thermal input to the rooms.

The model can be used to define the set of all feasible input trajectories that the building can follow while respecting the constraints. It takes the following form:

$$
\mathcal{U}(\hat{x}) = \left\{ \mathbf{u} \left|
\begin{array}{l}
x_{i+1} = Ax_i + B_u u_i \\
y_i = Cx_i \\
|y_i - T_{\text{ref}}| \leq \beta \\
u_i \in \mathbb{U} = [0,1]^{n_u} \\
x_0 = \hat{x}, \\
\forall i = 0, \ldots, N-1,
\end{array}
\right. \right\}
\tag{5.2}
$$

where N is the prediction horizon, T_{ref} is the optimal temperature and β a parameter controlling the allowed comfort level deviation from the optimum. The constraint $u_i \in \mathbb{U}$ captures the constraint on the maximum and minimum PWM ratio to control the heaters.

5.3   Problem statement

In the following, we present a particular instance of the original planning problem 2, *i.e.*, the case of a single commercial building providing secondary frequency control is considered. In particular, we assume the service provision to only take place during night-time when external disturbances (occupancy, solar radiation, etc.) acting on the building are at the minimum level. Moreover, we do not consider the possibility for the provider to participate in the intraday market. The main idea behind these simplifications is to reduce the complexity of the general problem to a minimal example so as to easily compare the two approximations in terms of their main characteristics and performance.

As already introduced in the previous chapter, the provision of SFC involves two distinct phases,

bidding and tracking, that are illustrated in Figure 5.2 and briefly described in the following.

The bidding phase is done offline and consists in advertising to the grid operator two quantities, namely a *baseline* energy consumption $\bar{\mathbf{D}}$, and a *capacity bid* γ. The latter represents the highest deviation (in absolute value) in power consumption with respect to the purchased baseline the ASP is willing to track over the activation period.
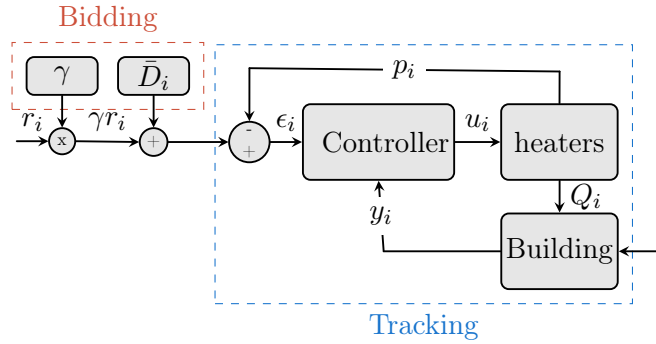
Figure 5.2 – Architecture of the control system for tracking service procurement with LADR.

The second phase is tracking. This phase is performed online as follows: considering the current state of the network, the power system operator will compute, at time i, a normalized AGC signal r_i, which will be sent to the service providers according to their respective accepted bids, so that the magnitude of the AGC signal they receive is proportional to their bid. The providers then have to modify their power consumption p_i so that

$$|\epsilon_i| := |p_i - \bar{D}_i - \gamma r_i| \leq m_e \gamma \tag{5.3}$$

where $\bar{D}_i$ is the baseline purchased for timestep i, p_i the actual power consumption at timestep i, r_i the normalized AGC signal and γ is the capacity bid. m_e represents the maximum tracking error allowed, as a percentage of the total submitted bid. This essentially states that the provider must consume the sum of the baseline power it has purchased and the scaled AGC signal it receives and that its tracking error ϵ_i must be less than a percentage of the capacity bid γ.

The goal of the controller is to choose the inputs to the HVAC system u in order to maintain the tracking error small enough as specified by equation (5.3), and, at the same time, satisfy constraints on the system.

At the time of bidding, the provider needs to solve the following uncertain optimization problem:
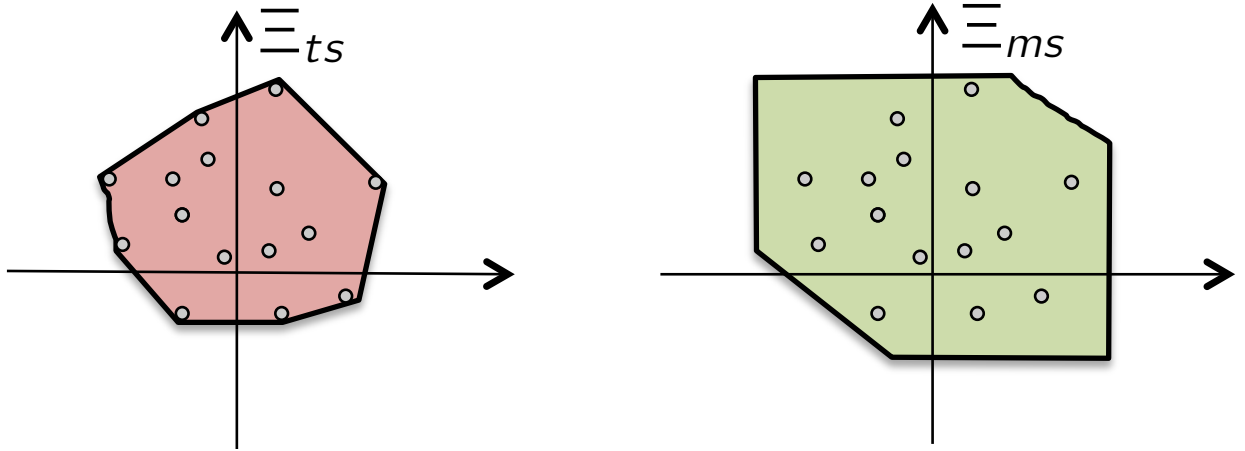
Figure 5.3 – Conceptual sketch of the construction of sets Ξ_{ts} and Ξ_{ms}. The grey dots, that represent possible realization of the the unknown AGC signal, are used to fit the two uncertainty sets.

Problem 5 (Planning problem).

$$\underset{\gamma,\bar{\mathbf{D}},\pi_{\mathbf{u}}}{\text{minimize}} \quad c_{\text{energy}}^T \bar{\mathbf{D}} - c_{\text{reward}}^T \gamma$$

$$\text{s.t.}$$

$$\text{(Building constraints)} \quad \mathbf{u} \in \mathcal{U}(\hat{x}) \tag{5.4}$$

$$\text{(Building active power)} \quad \mathbf{p} = \Lambda \mathbf{u} \tag{5.5}$$

$$\text{(Power tracking)} \quad ||\boldsymbol{\epsilon}||_\infty \leq m_e \gamma \tag{5.6}$$

$$\text{(Power flexibility)} \quad \gamma \geq 0 \tag{5.7}$$

$$\text{(Control Policies)} \quad \mathbf{u} = \pi_u(\mathbf{r}) \tag{5.8}$$

$$\text{(Uncertainty)} \quad \forall \mathbf{r} \in \Xi. \tag{5.9}$$

where, with respect to the general planning problem 2, many simplifications have taken place: only one resource is available, *i.e.* $N_{cont} = 1$, the provider does not participate in the intraday market, *i.e.* $\mathbf{D} = \bar{\mathbf{D}}$, and the hosting distribution network is assumed to be stiff meaning that all its constraints will be satisfied regardless of what the building will do. Thus, the grid constraints (4.9), can be discarded. Finally, the matrix $\Lambda = \mathbf{I}_{N \times N} \otimes \mathbf{1}_{n_u}^T$.

5.4 Solution to the planning problem

In this section we provide the reformulation of the planning problem using the two-stage and the multi-stage methods described in 4.6.3.

5.4.1 Two-stage approximation

For the two-stage approximation, the uncertainty parameter, r, is assumed to belong to the uncertainty set, Ξ_{ts} which is constructed, starting from historic AGC signals, as detailed in (4.23).

$$\Xi_{\text{ts}} = \left\{ \sum_{k=1}^{N_{\text{scen}}} \lambda^{(k)} \mathbf{r}^{(k)} \mid \sum_k \lambda^{(k)} = 1, \lambda^{(k)} \geq 0 \right\} \tag{5.10}$$

The first stage variables are the capacity, γ, and the baseline consumption, $\bar{\mathbf{D}}$ which are decided before the realization of the random parameter, $\mathbf{r}$, while the second-stage decisions are the realizations of $\mathbf{u}$ which depend on the uncertain parameter.

In particular, for each scenario of the uncertain parameter, there is a separate trajectory, $\mathbf{u}^{(k)}$, of the second-stage decision variables resulting in the following optimization problem:

$$\begin{aligned}
\underset{\gamma, \bar{\mathbf{D}}, \mathbf{u}^{(k)}}{\text{minimize}} \quad & c_{\text{energy}}^T \bar{\mathbf{D}} - c_{\text{reward}}^T \gamma \\
\text{s.t.} \quad & \\
& \mathbf{u}^{(k)} \in \mathcal{U}(\hat{x}), \\
& \mathbf{p}^{(k)} = \Lambda \mathbf{u}^{(k)} \\
& \|\boldsymbol{\epsilon}^{(k)}\|_\infty \leq m_e \gamma \\
& \forall \mathbf{r}^{(k)} \in \Xi_{\text{ts}}
\end{aligned}$$

5.4.2 Multi-stage approximation

In this case, the construction of the uncertainty set, Ξ_{ms} follows the algorithm described in (4.24).

$$\Xi_{\text{ms}} = \{ \mathbf{r} \mid \|\mathbf{r}\|_\infty \leq 1, \|\tilde{r}_{[1 \to N]}\|_\infty \leq s_{\text{max}} \} \tag{5.11}$$

where s_{max} denotes the worst-case cumulative sum on a number of historical realization of the AGC signal, and $\tilde{r}_{[1 \to i]}$ is the cumulative sum of the signal up to time i.

The multi-stage reformulation is then given by:

$$\begin{aligned}
\underset{\gamma, \bar{\mathbf{D}}, \mathbf{M}, \mathbf{v}}{\text{minimize}} \quad & c_{\text{energy}}^T \bar{\mathbf{D}} - c_{\text{reward}}^T \gamma \\
\text{s.t.} \quad & \mathbf{u} \in \mathcal{U}(\hat{x}) \\
& \mathbf{p} = \Lambda \mathbf{u} \\
& \mathbf{u} = \mathbf{M}\mathbf{a} + \mathbf{m} \\
& \|\boldsymbol{\epsilon}\|_\infty \leq m_e \gamma \\
& \forall \mathbf{r} \in \Xi_{\text{ms}}
\end{aligned} \tag{5.12}$$

with $\mathbf{M}$ respecting the causality requirement (4.34).

5.5   Simulation results

This section compares in simulation the two-stage and the multi-stage approximations discussed in Section 5.4.1, and 5.4.2, respectively.

5.5.1   Simulation setup

The model of the experimental setup, described in Section 5.2.2, is used for simulations. The sampling period is chosen equal to 15 minutes which provides a nice compromise between the temporal resolution of the control and computational complexity of the problem formulation. The comfort range for temperature is chosen as $21\,°\mathrm{C}$ to $25\,°\mathrm{C}$ ($\beta = 2^o\mathrm{C}$). Perfect tracking of the unknown AGC signal is also required ($m_e = 0$).

To solve the planning problem 2 and compute the maximum power capacity that the building can support over the activation period, we use both the two-stage and the multi-stage approach. The two uncertainty sets, Ξ_{ms} and Ξ_{ts}, are constructed offline. The scenarios used to perform this task are obtained by breaking the yearly normalized AGC signal of 2013 into 876 ten hour samples. The model for the building is obtained around an operating point of $5^o\mathrm{C}$ for the outside temperature. Solving the two planning problems results in different values for the optimal bid $\gamma^\star$ and baseline $\bar{\mathbf{D}}^\star$.

To test the robustness, and quality of the solution, historical realizations of the AGC signal of 2014 are considered for validation. The AGC for 2014 is also broken into 876 ten hours test instances. Each ten-hours test AGC sample is multiplied by the optimal power capacity γ and added to the baseline $\bar{\mathbf{D}}$ to obtain the total power signal to be tracked by the system

For the two-stage approximation, an open loop optimization problem is solved for each ten hour test sample to optimally distribute the power across the four zones while respecting the comfort constraints. Similarly, for the multi-stage approximation, the optimal affine control law parametrized by $\mathbf{M}^\star$ and $\mathbf{m}^\star$ is used to compute the open-loop trajectories of the zones temperature. The result is depicted in Figure 5.4.

5.5.2   Analysis of results

As seen in Figure 5.4 there are a few differences between the two approaches both in terms of bid capacity and of thermal response of the system while providing AGC tracking. The multi-stage approach is more conservative and results in a capacity bid of $\pm 1.85 kW$ while the two-stage approach results in a capacity bid of $\pm 3.2 kW$. This is visible in the bottom plots of Figure 5.4 where the AGC signals and their maximum amplitude are shown. The computed capacity represents 24% and 41% of the maximum available power, respectively. The resulting temperature trajectories in the four zones of the building for all the considered AGC test samples are shown in the top plots of Figure 5.4. For the multi-stage approach, the zone temperatures stay more closely around $23\,°\mathrm{C}$

Multi-Stage
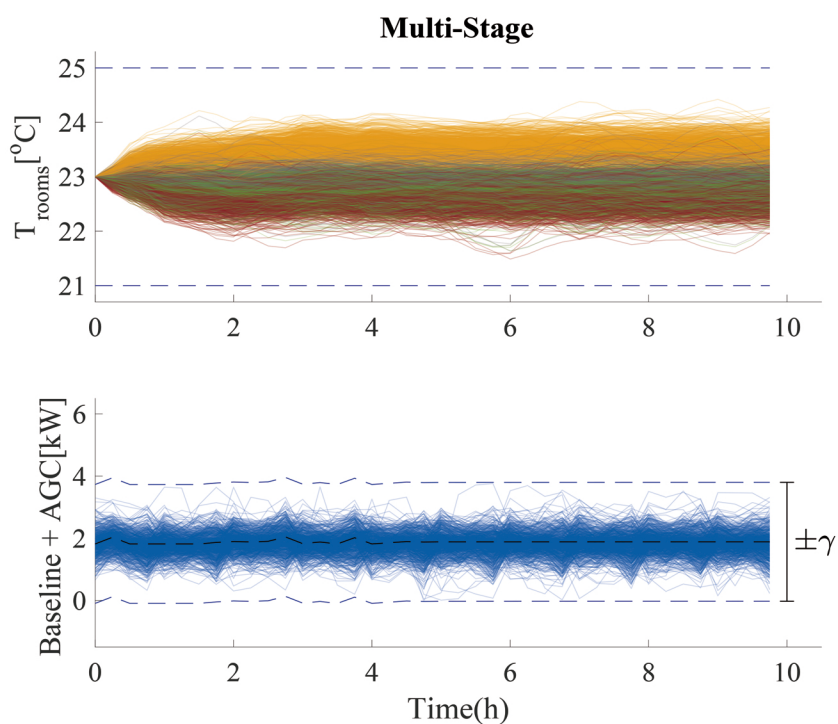
(a) Multi-Stage
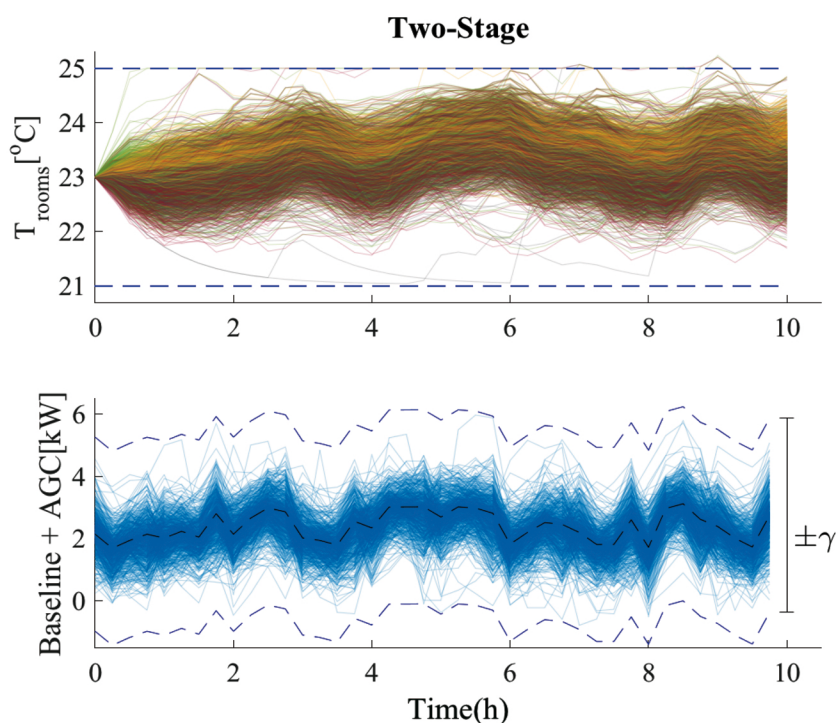
Two-Stage

(b) Two-Stage

Figure 5.4 – Open-loop predictions for the AGC multi-stage and two-stage controller for real AGC signals from 2014. The obtained capacity and baseline obtained off-line on data of 2013 are used to compare the two approaches applying 876 different signals extracted from the data set of 2014. Upper: Temperature in each of the 4 zones (different colors) during activation period. Lower: AGC signals superimposed on the computed baseline (black dashed line) and capacity bid bounds.

which represents the most robust state to be in to absorb both positive and negative realizations of the AGC. For the two stage approach, temperatures are closer to the constraints and violate the constraint slightly for a few AGC test samples.

5.6   Experimental results

5.6.1   On-line operation

In the preceding sections, the planning problem was detailed. This problem is solved at the beginning of the activation period to determine the capacity bid γ and the baseline $\bar{\mathbf{D}}$. These quantities are then fixed for the duration of the activation. In closed-loop, slightly modified versions of problems (5.11) and (5.12) are respectively solved at each time step with a shrinking horizon, and with the baseline power consumption $\bar{\mathbf{D}}$ and the power capacity γ being fixed. In the two-stage case, the problem is further modified such that the first time step of the input and the state trajectories are first stage variables while variables for the rest of the horizon are still second stage decision variables. In both cases, the aim of the closed-loop controller is to determine a feasible command input to the HVAC system, u_0, to each zone so that a high-quality tracking service is provided while making sure that thermal constraints are satisfied. The reader is referred to Chapters 6 and 7 for an explicit formulation of the real-time controller.

5.6.2   Experiments

Experiments have been performed to test the algorithms described in Sections 5.4.1 and 5.4.2. The tests have been conducted overnight to reduce the effect of disturbances acting on the building. More specifically, tests are performed for a period of 10 hours from 8 pm to 6 am on different days in February and March 2015. During the experiments, the outside conditions were relatively consistent with outdoor temperature ranging from 4 to 10^oC. In particular, since the outside temperatures were close to the ones experienced during the identification procedure, the identified steady-state model for outside temperature could be used. For the computation of the bid, it was assumed that the temperature at the beginning of the experiment is 23^oC to allow a meaningful comparison between different days, and with the simulation results. Therefore, the temperature was regulated to this value before each experiment. Since the same model and initial condition was used in simulation and in the experiment, the resuls of the planning problem were the same, as detailed in Section 5.5.1, with optimal bids that correspond to 24 and 41 % of the installed capacity, respectively.

For the experiments, different realizations of the AGC signal were used for testing. After the commitment of the bid and baseline, the computation of the control inputs, which determines how energy is split across the rooms, is performed with a time step of 15 minutes, therefore allowing to compensate partly for forecast errors. In practice, the frequency of update of the AGC signal is faster than 15 minutes, but the controller computed can be used to apply control actions at a faster rate. A rate of one minute was used in the experiments. A Kalman filter is used to estimate the state of the system.

After computing optimal bids solving respectively Problems (5.11) and (5.12), 4 closed-loop experiments were run, applying two different values of the AGC signal. Results are reported in Figures 5.5 and 5.6. For each experiment, four subplots are shown. The first one shows the evolution of the temperature in the four rooms, the second depicts the total power consumption in the four rooms and how it is split between the rooms. It can also be observed there how the energy dispatch in the four rooms is re-adjusted in closed-loop every 15 minutes. The third plot shows the scaled AGC signal that needs to be tracked and the fourth plot shows the integral of the AGC signal over time, which represents the energy stored in the system as a result of the tracking. In the case of the two-stage method, the computed bid is higher and, therefore, results in larger tracking requirements which drive the temperature closer to the comfort limits. This confirms the results obtained in simulations. Small constraint violations are observed in the case of the two-stage method. This is expected since already in the case of perfect predictions and no model mismatch in simulations, the two-stage method displays an "aggressive" behavior and runs very close to the constraints. The magnitude of those violations is however below $0.5^{o}C$.

5.7 Concluding remarks

In this chapter, we have looked at a specific instance of the general framework proposed in chapter 4. In particular, in order to easily compare the two proposed approaches, multi-stage and two-stage, we have focused on a relatively simple case study: a multi-zone office building, served by controllable electrical heaters, providing SFC in the Swiss market. Based on the results, a few observations are in order.

The first remark, confirmed by both the simulation analysis as well as the experiments, concerns the different level of conservatism of the two approaches. First, as it appears also from Figure 5.3, if the same set of samples is used in both approximations, $\Xi_{\mathrm{ts}} \subset \Xi_{\mathrm{ms}}$, and using the latter will lead to more conservative solutions. Second, the causality requirement is relaxed in Problem (5.11), which is an optimistic assumption in the sense that it improves the optimal value of the problem. It can be actually formally proven that the optimal objective value function is lower in Problem (5.11) than in Problem (5.12). On the other hand, the two-stage approximation has displayed, in simulations and experiments, a more aggressive behavior running very close to the temperature bound. In some cases, small temperature violations have also been observed.

The second thing to be noticed relates to the causal structure of the solution. In fact, the multi-stage approach has the advantage of retaining the original multi-stage structure and the causality requirement which could be a characteristic of paramount importance in certain applications. This is the case, for instance, of the simulation setup considered in Section 5.5.1 where the intraday control policy could be directly incorporated in the optimization problem for the multi-stage approach, while it had to be pre-determined for the two-stage approach. Moreover, for the latter, the pre-determined intraday policy did not accommodate the case of multiple disturbances acting on the system and a single parameter had to be considered (the AGC signal).
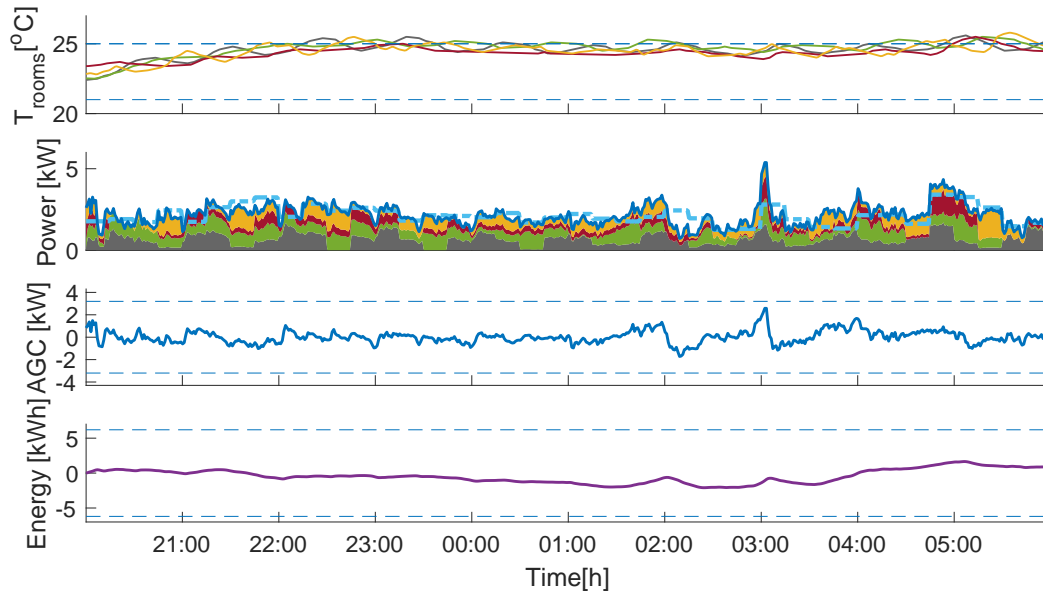
(a) 1$^{\text{st}}$ AGC signal
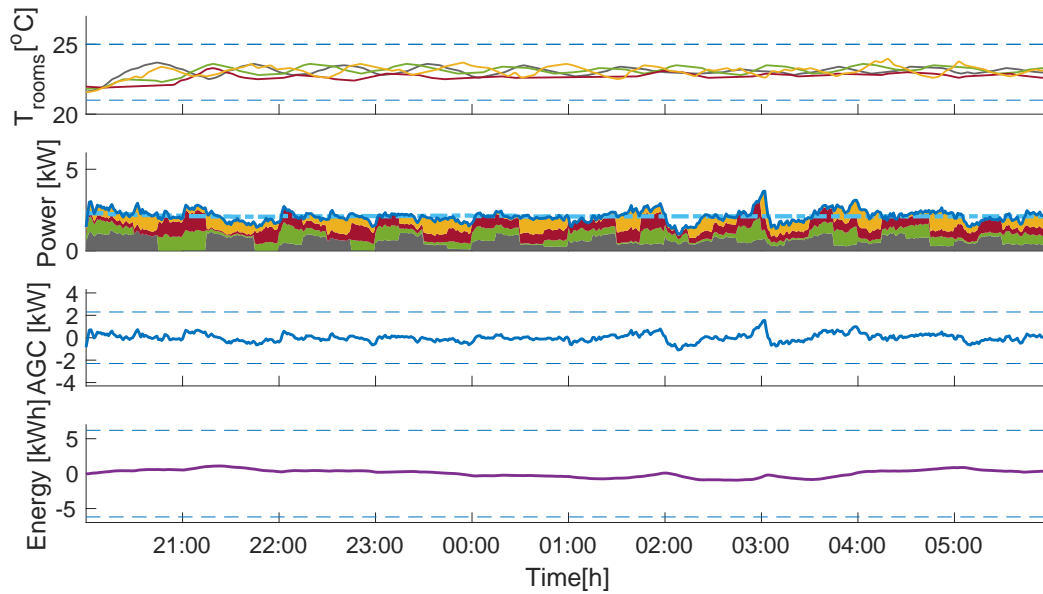
(b) 1$^{\text{st}}$ AGC signal

Figure 5.5 – Two experiments of AGC tracking. 1$^{\text{st}}$ selected AGC signal extracted from real data of 2013, and is used to test and compare the two controllers. Upper: Temperature variation for different zones. Each color corresponds to the measured temperature in each zone. Middle Up: Power distribution among zones. Middle Down: AGC signal variation and capacity bid. Lower: Integral of the AGC
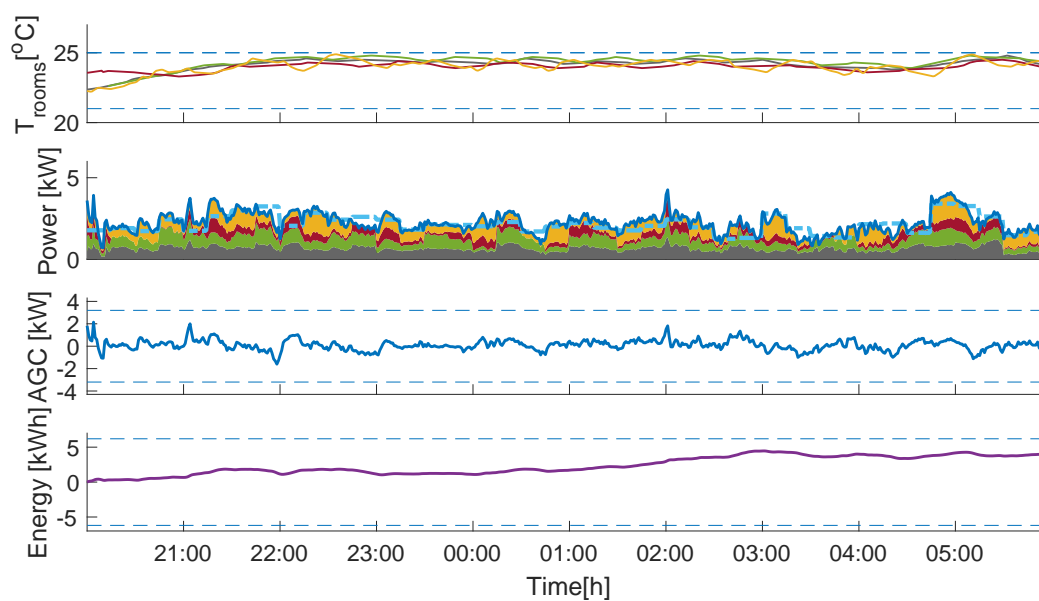
(a) 2$^{\text{nd}}$ AGC signal
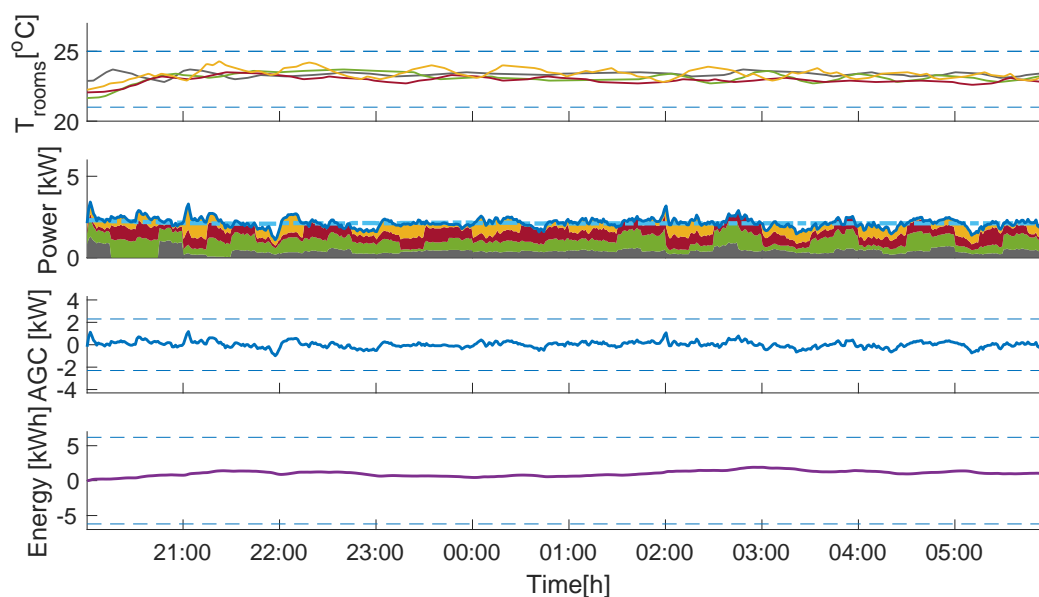
(b) 2$^{\text{nd}}$ AGC signal

Figure 5.6 – Two experiments of AGC tracking. 2$^{\text{nd}}$ selected AGC signal extracted from real data of 2013, and is used to test and compare the two controllers. Upper: Temperature variation for different zones. Each color corresponds to the measured temperature in each zone. Middle Up: Power distribution among zones. Middle Down: AGC signal variation and capacity bid. Lower: Integral of the AGC

Finally, the last observation relates to the assumption we made in Chapter 4 on the linearity of the optimization problem (after the linearization of the PF constraints), *i.e.*, both linear in the decision variables and the uncertain parameters. For the multi-stage approach, this is a strict requirement as it is needed in order to obtain, by means of duality arguments, a tractable robust counterpart of the uncertain optimization problem [51]. On the contrary, this assumption can be relaxed for the two-stage approach. Relying solely on the addition of constraints of the same type as the ones present in the original problem, the method could also be considered in a non-linear setting.

In a nutshell, the choice of the approximation scheme should depend on the main requirements associated with a particular application. Thus, the two-stage structure should be considered for all applications where performance is the main driver and/or when the relaxation of the multi-stage structure does not have a detrimental impact on the solution. On the other hand, the multi-stage method should be used in situations where it is desirable to have additional robustness margins (*e.g.*, to guarantee constraint satisfaction against unmodelled dynamics or disturbances), and/or when the temporal structure of the problem plays an important role.

Part III

Provision of grid services using distributed energy resources: Experiments

Chapter 6

Multi-Time Scale Coordination of Complementary Resources for the Provision of Ancillary Services

6.1 Introduction

In the previous chapter, we provided a quantitative comparison between the two control frameworks of Problem 3 and Problem 4. In particular, we considered a controlled environment composed of a single office building served by electric heaters. To reduce the complexity of the problem all the experiments were conducted solely overnight and in an unoccupied setting. Thus, as already highlighted, the focus was not to provide a full experimental validation of the proposed control methodologies but rather, gain an insight about their advantages and disadvantages.

On the contrary, the main objective of this chapter is to actually provide such an experimental validation. In particular, the study is designed so as to answer the following research questions:

Technical feasibility: Is it possible to coordinate the operation of a set of DERs to offer ancillary services to the grid over a realistic time range, and in full compliance with the current regulations?

Benefit of combining complementary resources: Is there a benefit in aggregating resources that display complementary technical specifications, and can this benefit be quantified?

Impact on the primary purpose of the resources: Is the primary purpose of the resources, *e.g.* occupants' comfort for a commercial building, affected when they are used to provide ancillary services to the grid?

To answer these questions, we conducted an experimental campaign emulating the provision of SFC service using a campus-scale platform that is composed of a fully-occupied building served by

an electric Heat Pump (HP), and an emulated grid-connected battery. The experiments were performed, for nine consecutive days, closely following all the current Swiss regulations for the provision of SFC. To the best of our knowledge, this work represented the first experimental demonstration of heterogeneous DR resources providing SFC over a realistic time range and in such realistic conditions.

6.2 Review of related experimental works

In this section, we provide an overview of the related experimental work that investigate the potential of DERs providing frequency reserves. In particular, due the considered experimental platform, we focus on work involving the control of BESSs, commercial buildings, or a combination of the two.

AS with BESS

The main challenge when proving ASs with BESSs is represented by the management of the SoC level. In fact, the control signal to be tracked can exhibit significant biases over prolonged periods of time which can rapidly lead to the complete charge or discharge of the BESSs. For this reason, in recent years different recharging strategies have appeared to optimize the provision of fast regulating services.

In [71] an off-peak-hours recharging strategy was considered that manages the SoC tapping the available capacity of conventional generators. *Oudalov et al.* [108] utilize a deadband around nominal system frequency to adjust the SoC. [19] propose a moving-average strategy to compensate for imbalances in the regulation signal, and efficiency losses in the storage system. This continuous adjustment is summed to the received tracking request in order to obtain the final power setpoint to the battery. The strategy was successfully tested on the grid-connected Zurich 1 MW electric battery [65, 66], that was used for different grid applications. All these studies typically focused on PFC due to the smaller energy throughput that is typically required with respect to, *e.g.*, SFC. Nevertheless, due to the worst-case energy requirements and/or the conservative prequalification rules recently implemented by many Transmission System Operators (TSOs), even in this case, BESSs do not represent, in general, an economically viable solution due to very large capital costs [65].

AS with commercial buildings

Many works have focused on the practical feasibility of actually deploying both simple and advanced strategies to reliably provide SFC using the HVAC systems of commercial buildings. A shared conclusion is that, due to their inherent complex nature with multiple cascade control loops, self-correcting behaviors, and physical limitations of the equipment, the power consumption of a standard HVAC systems cannot, in general, be modulated reliably and at a very high-frequencies.

To mitigate this problem, several approaches have been proposed that typically focus on identifying

specific HVAC components that can sustain such fast power changes. In [132], *Su et. al* propose a practical control framework to track a filtered version of the Automatic Generation Control (AGC) signal for secondary frequency regulation. This is achieved by acting on the chilled water supply setpoint of a chiller which, in turn, has a quantifiable effect on the electric power consumption of the HVAC system. A similar approach was also considered in [75, 88] where the power tracking was provided by adjusting the fan power consumption of the main air handling unit through either direct fan speed offset or by adjusting the mass flow setpoint. Also in this case, the building receives a filtered version of the AGC that is tracked in a [1/30s to 1/1min] or [1/1min 1/10min] frequency band, depending on the considered configuration. In the same direction, a more extensive experimental study is represented by [139] where *Vrettos et al.* analyze the potential of offering frequency regulating by modulating the fan power via speed control. The study includes also a formal computation of the regulation capacity that can be offered by the building. Experiments were made in single-zone unoccupied test cells equipped with a standard cooling system. The conclusion of the aforementioned papers, also confirmed in the extensive simulation studies [13, 14], is that controlling fast-responding Air Handling Unit (AHU) fans by means of either speed offset, or supply pressure/mass flow setpoint, could guarantee tracking performance in line with the typical TSO's requirements. However, despite their fast response rates, these methods have also some drawbacks. First, direct control of fan speed is not readily possible in many BEMS [4]. Thus, this would require some level of retrofitting adding cost and complexity. Moreover, due to the complex control architecture of commercial buildings, slower control loops will likely compensate for net changes to supply pressure of mass flow which limits the ability of these strategies to track reference signals with slow time-scales. A different approach has been proposed in [4, 44, 148], where authors propose to track the reference signal by adjusting the thermostat setpoint offset which has an indirect effect on the fan consumption through the corresponding change in the room damper and, therefore, the mass flow pressure. The advantages and disadvantages of this method are opposite of the previous ones [14]. In general only software modifications would be required since thermostat changes can be done through many BEMS. However, since the electric power of fans is only controlled indirectly, communication and mechanical latency can significantly impact the tracking performance [4].

6.2.1 Motivation

From the previous discussion, it is clear how BESSs and buildings are to some extent complementary resources. On one side BESSs are power-intensive devices with restrictive energy limitations; on the other side, buildings are energy-intensive devices with restrictive power limitations. Each type of resource could not provide ancillary services due to respectively economical and technical limitations but they have the potential to be operated together to provide reserves economically and reliably, and, in turn, improve the overall efficiency of the network. Thus, the idea is to overcome the limitations of these resources by combining them into a single virtual resource. We would like the BESS to only take care of high-frequency components of the AGC while the building takes care of slower and more energy-intensive components.

Similar ideas have been recently explored. In [3, 18] the authors propose control frameworks to

split the control signal at the TSO level so that the resulting components better fit the technical capabilities of different resources such as BESSs, supercapacitors, etc. However, the implementation of these frameworks would require a drastic modification to the way the power system is currently operated. On the contrary, we propose to combine complementary resources to provide fast regulating services while complying with current regulations. On the technical side, few contributions have also appeared. A rule-based controller to dispatch regulation services between a power generator and a flywheel storage device has been proposed in [57], where *Jin et al.* show, by means of simulation results, how this hybrid regulation resource outperforms standard generators in terms of tracking performance. Also in [63], authors consider the combination of an energy storage system and a variable-speed hear pump to provide frequency regulation. Feedback controllers to coordinate the two resources are designed and successfully tested in a power hardware-in-the-loop simulation. Despite providing clear evidence of the feasibility and potential of aggregating complementary slow and fast resources for the provision of frequency services, the aforementioned papers do not consider the planning phase to characterize the available flexibility which was instead derived by trial and error procedures. Moreover, they did not propose a dynamic coordination scheme of the two resources which updates the tracking setpoints based on their current states.

Notation: Beside the usual notation, we denote with $\mathbf{p}^{\text{res}}$ indicates the real power flow of the particular resource, res, whereas the bracket superscript notation, $\mathbf{p}^{\text{res},(k)}$ stands for the power trajectory corresponding the k-th scenario. For the sake of clarity, all the relevant notation is reported in Table 6.1.

6.3   Core Idea - Intuition

Fast regulating services to the electric grid, such as SFC, have been historically provided by traditional power plants such as hydropower plants, coal or gas stations, etc. The reason for this can be traced back to their relatively fast power responsiveness and to their inherent capabilities to absorb any energy bias of the regulation signal. The focus of this chapter is to provide a control framework in order to reliably provide these kinds of services by combining BESS and commercial buildings. In the following section, the core idea underpinning the whole proposition is provided.

6.3.1   Fast and Slow Resources

In this section, we provide a qualitative description of the physical capabilities of both the BESSs (fast resources) and commercial buildings (slow resources) with respect to the typical requirements for the provision of SFC to the grid.

In particular, we focus on a few key aspects that can be broadly categorized as *power requirements* and *energy requirements*.

Following the framework proposed in [105], we compare the main characteristics of each resource with respect to the typical requirements for SFC (in the Swiss case). In particular, we focus on five

Variable	Definition
x_i	Current state of the system for the building
d_i	Vector of disturbances acting on the building
y_i	Current zone indoor temperatures
SoC_i	Current state of Charge at the BESS
p_i^h	Active power consumption of the building
p_i^s	Active power injection of the BESS
p_i	Total electric power injection
r_i	AGC signal at time i
$a_i^{(k)}$	Residual AGC signal for the k-th scenario (bidding)
$\hat{r}_i^{(k)}$	AGC signal for the k-th scenario (real-time)
γ	Power flexibility for SFC
$\bar{D}_i$	Day-ahead baseline for SFC
ϵ_i	Tracking error for SFC
τ_i	Intraday transaction
$\pi_p^h, \quad \pi_p^s \quad \pi_{\mathbf{D}}$	Control policies
Parameter	Definition
N	Prediction horizon
N_{scen}	Number of scenarios in the two-stage stochastic approximation
J	Economic cost for the planning problem
T_{ref}	Reference temperature for the building zones
β_i	Allowed temperature variation for the building zones
m_e	Normalized maximum allowable tracking error
$\mathrm{SoC}_i^{\max}, \mathrm{SoC}_i^{\min}$	Upper and lower bound for the SoC of the BESS
$\mathrm{p}_{\max}^h, \mathrm{p}_{\min}^h$	Upper and lower bound for the electrical consumption of the HVAC system
$\mathrm{p}_{\max}^s, \mathrm{p}_{\min}^s$	Upper and lower bound for power injection of the BESS
Ξ	Uncertainty set for the AGC signal
T	Moving-block matrix for the building consumption

Table 6.1 – Nomenclature.
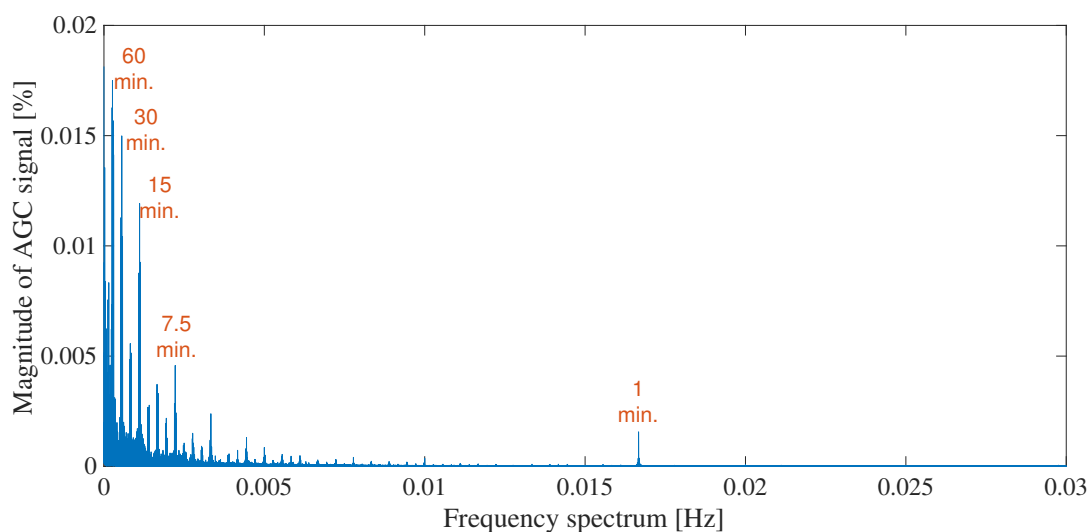
Figure 6.1 – Spectrum of the AGC signal

key characteristics that are described in the following.

- Energy capacity: the maximum energy that can be shifted or stored by the resource.

- Ramp rate: the maximum rate at which the resource can modify its power production/consumption from its maximum to its minimum value or vice-versa.

- Response granularity: the capability of the resource to implement any power setpoint between its minimum and maximum operating points.

- Response frequency: how often the power injection of the resource can be modified without wear to its physical equipment

- Response time: the maximum time elapsed between the power tracking request and the consequent modification of the power injection of the resource

To analyze the main characteristics of SFC with respect to the considered framework, we considered its frequency spectrum, shown in Figure 6.1. The spectrum is obtained from one year of historical realization of the AGC signal for 2014. A few observations are in order. First of all, high-frequency components are quite damped due to the effect of both the system inertia as well as primary frequency control. Many distinct peaks are then visible in the medium range of the spectrum which corresponds to particular instants of the day and are mainly due to the way the market is operated. The highest peaks are at a frequency corresponding to 60, 30, and 15 minutes. More generally, an overall significant presence of low-frequency peaks can be observed, that are due to the integral action of SFC. Thus, a resource providing this service will need to track a signal spanning a wide range of frequency components (from 1/4 seconds to 1/60 min.) and with significant energy requirements.

Property	SFC requirements	BESS	Commercial buildings
Energy capacity [kWh]	Significant	Limited (very expensive)	Significant (cheap)
Ramp rate [kW]	High	Very High	Moderate
Response granularity	Continuous	Continuous	Limited
Response frequency [Hz]	High	Arbitrary	Slow to Fast
Response time [s]	High	High	Slow to Fast

Table 6.2 – Qualitative description of the technical capabilities of BESSs and Commercial Buildings in relation to the requirements for the provision of SFC.

Referring to Table 6.2, BESSs are highly responsive devices that exhibit reaction times and ramp rates that are only limited by the capabilities of their power inverters. Thus, BESSs not only meet but even exceed, most of the requirements for the provision of SFC. However, due to the fact that the AGC signal can display significant biases over prolonged periods of time, the worst-case based dimensioning of BESSs typically represents a severe impediment to their widespread deployment. In fact, since the price of BESSs is typically determined by the required energy capacity and since the provision of SFC is rewarded in proportion to the power capacity offered, it is desirable to keep the energy/power ration as small as possible. Regarding commercial buildings, due to the large variety of different HVAC systems, it is difficult to provide a unique identifier for all the considered key characteristics. Nevertheless, general qualitative statements can be given. First of all, commercial buildings are inherently characterized by a large thermal inertia that can be exploited to cheaply store energy without perceptibly affecting the occupants' comfort. On the contrary, HVAC systems are in general not suitable for fast regulating services [105]. In particular, their response time can be relatively slow due to the mechanical, control, and communication latencies that are introduced by complex BEMS [4]. Also, strict constraints on ramp-rates and response frequency are typically imposed to prevent an excessive stress and wear of the equipment.

Comparing the characteristics of these two resources, their complementarity is apparent. On one side, the BESSs are power-intensive devices with restrictive energy limitations; on the other side, buildings are energy-intensive devices with restrictive power limitations. Each of these resources cannot readily provide ancillary services due to respectively economical and technical limitations. However, they have the potential to be operated together to provide reserves economically and reliably, and in turn, improve the overall efficiency of the network.

6.4   Experimental setup

To experimentally confirm the intuition of the previous section, a campus-scale platform was developed that comprises a fully-occupied building on campus and an emulated BESS. The experimental setup is depicted in Figure 6.2 and it is described in the following.
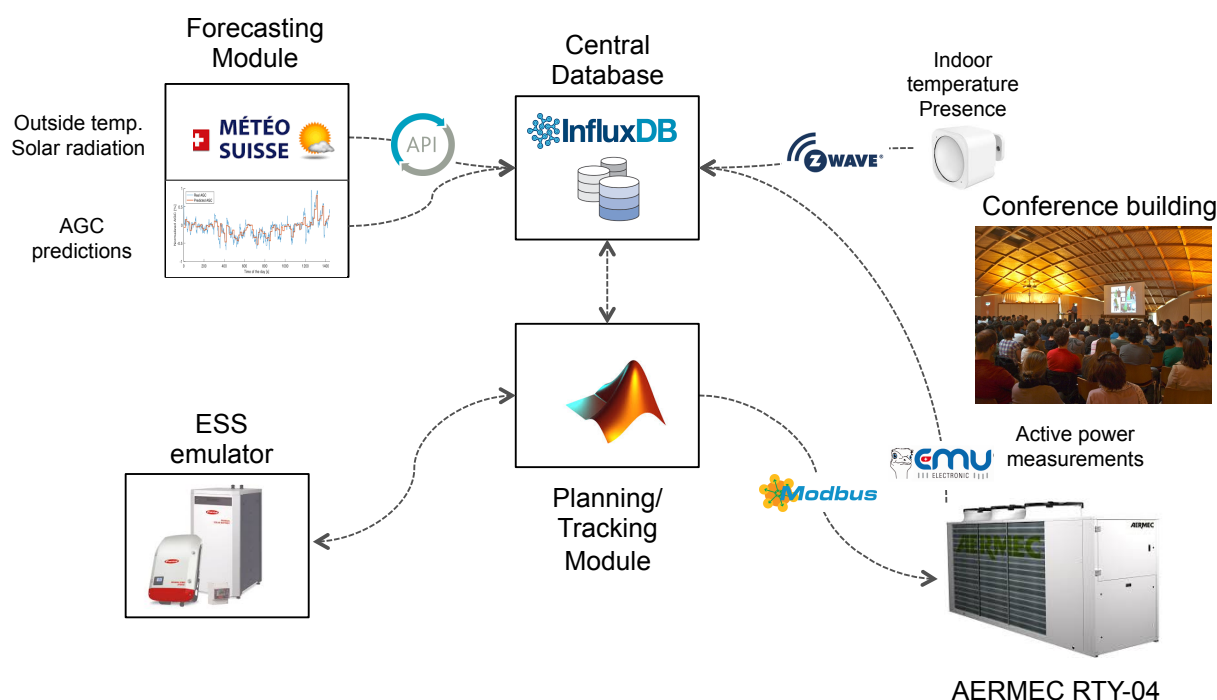
Figure 6.2 – Overview of the experimental setup.

6.4.1   Commercial Building

We run our experiments in a relatively newly constructed building on the EPFL campus. The building of roughly 600 m^2 is used as a large audience/lecture room and is occupied on a regular basis with a maximum of 200 occupants. The building is served by a forced air-system that works for both cooling and heating depending on the season. More precisely, a single compressor roof-top Heat Pump (HP), AERMEC RTY-04, is installed which accounts for 6kW active power at peak. A proprietary controller inside the HP continuously monitors the return air temperature coming from the building, it compares it with the reference temperature and determines the operating point of all its active components (compressor, fan, etc.). We decided not to overwrite the logic of the controller since this has been specifically designed by the manufacturer to reduce the stress and wear of mechanical components of the HP. Moreover, in order to re-design this internal controller, a certain level of retrofitting would have been needed as not all relevant internal parameters are readily available.

Thus, the control input to the system is represented by the reference indoor temperature. This reference change will be tracked by the internal controller of the HP with a consequent effect on the electric power consumption of the HVAC system. A rule-based control routine interfacing the HP controller to the proposed controller was developed. The routine receives as input a power setpoint and returns a sequence (1-minute resolution) of indoor temperature setpoints that are sent

to the internal controller of the HP so as to track the given power setpoint. This has the significant advantage of requiring, in general, only minor software modifications to the existing BEMS as most of them allow the remote control of thermostat setpoints.

For the measurements, the building has been retrofitted with three wireless Aeotec Multisensors that continuously monitor indoor temperature, and presence. Weather data were collected from a nearby weather station every 5 minutes as explained in Section 6.6, including measurements for outside temperature, horizontal global solar irradiance and weather observation. The monitoring and control of the HP are done through the serial communication protocol, Modbus. A central processing unit continuously receives all measurements and it uploads them into the database, InfluxDB, which is specifically designed to handle time-series data. The database is stored on the local network and it is connected to the open-source visualization platform, Grafana, which allows one to continuously supervise the overall functioning of the system. The optimization-based controllers (offline, online) are implemented in MATLAB, using the YALMIP parser [81], running on a 3.4 GHz Intel Core i7 iMac with 32 GB 1600 MHz DDR3 memory capabilities. The computer, connected to the local network, can access the latest measurements from the database and can send the temperature setpoints to the building and the power injection setpoints to the battery emulator.

Building model identification

The identification of the building model was performed using standard black-box linear identification techniques. In particular, we performed three weeks of open-loop experiments where, in order to excite the system dynamics, the temperature setpoint to the HP controller was modified using a mix of step and Pseudorandom Binary Signal (PRBS) [78] signals within a safe range of temperatures. Regarding the weather, a reasonably varying pattern was observed over the total duration of the experiments. In particular, the outside temperature varied in the range [13, 31]°C with a mean of 21°C and a standard deviation of 3.72°C. Concerning the solar radiation, the maximum attained value was 1.06 kW/m^2 with a mean of 0.23 kW/m^2 and standard deviation of 0.31 kW/m^2.

The experimental dataset was divided into two chunks, one for identification (two weeks), and one for validation (one week). A linear sub-space identification approach was used [78]. A state space dimension of order $n_x = 3$ was found to be sufficient to appropriately describe the system dynamics. The model comprises three inputs, *i.e.*, p^h, the electric power consumption of the HP, T_{out}, the outside temperature, and Q_{sun} the global horizontal irradiance. To simplify the notation, we introduce the external perturbation vector, $d := [T_{out}, Q_{sun}]$. The output, y, is the indoor temperature inside the building which is obtained as the average of all installed sensors. The identified model is of the following form

$$
\begin{aligned}
x_{i+1} &= Ax_i + B_u p_i^h + B_d d_i + Ke_i \\
y_i &= Cx_i + e_i
\end{aligned}
$$

(6.1)

where the term e_i represents the noise component and K the disturbance matrix.
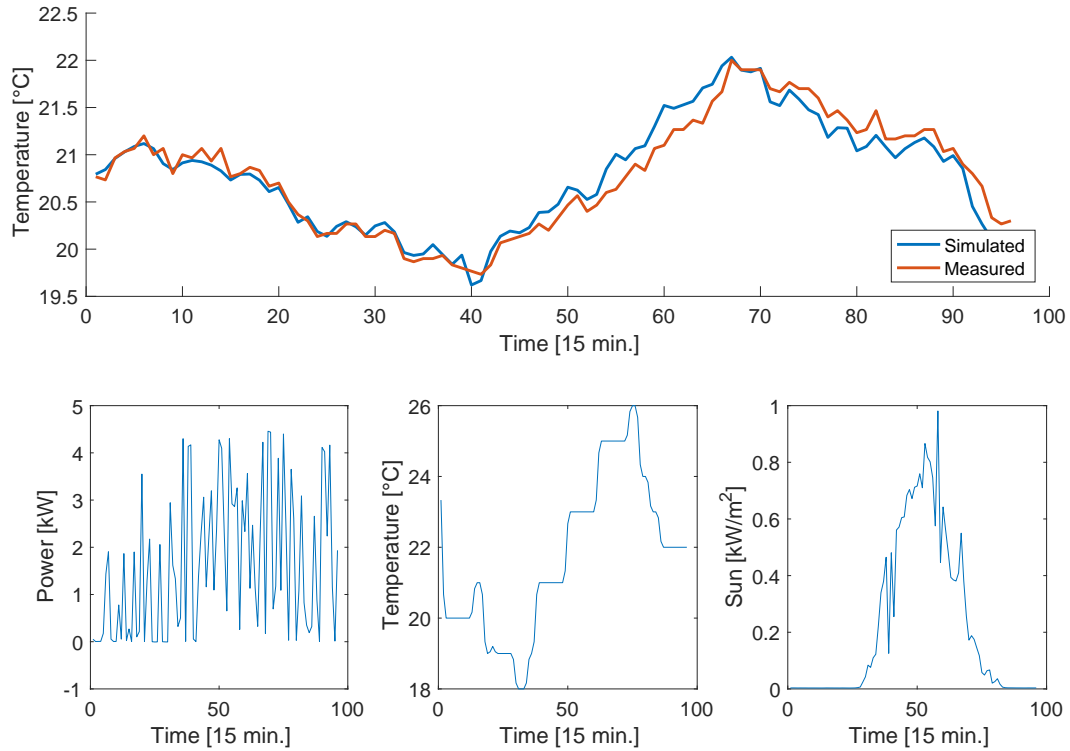
Figure 6.3 – Validation of the model. The prediction error over one day is typically below 1 C. Upper: The actual measured indoor temperature (red) vs. the predicted temperature of the identified model. Lower left: The electrical power consumption of the HP. Lower central: The outside temperature. Lower right: Solar irradiance.

As can be observed in Figure 6.3, the model shows relatively good performance. For a quantitative evaluation, we consider the fit of the model which is computed as follows:

$$\text{FIT} := \quad 100 \cdot \frac{1 - ||y - \hat{y}||}{||y - \bar{y}||}$$

where $\hat{y}$ is the predicted output, y the measured output, and $\bar{y}$ the mean of the measured output.

The FIT ranges always between 70 and 83 % on a one day prediction horizon for all days in the validation dataset.

An inspection of the model reveals a few key characteristics of the considered building. The impact of the solar irradiance is significant as the static gain from the solar input to the indoor temperature is $6.67^o\text{C}/(\text{kW}/m^2)$. The static gain from the outside temperature is less significant: $0.2^o\text{C}/^o\text{C}$. As

a three dimensional state-space model was selected, three different time constants are present: one relatively fast at 15 minutes and two slower ones at 1 and 1.5 hours respectively. A simple physical interpretation can be given for this result: the fast time constant corresponds to the fast indoor air dynamics, whereas the slower dynamics can be associated to the walls and floor temperatures which constitutes most of the thermal inertia of the building.

Building constraints

Based on the thermodynamic model presented in the previous section, it is possible to define, over a specified horizon, N, a set of electrical trajectories that the building can support without violating its dynamics and constraints. More precisely, we define:

$$\mathcal{P}^h(\hat{x}, \hat{\mathbf{d}}) = \left\{ \mathbf{p}^h \left| \begin{array}{l} x_{i+1} = Ax_i + B_u p_i^h + B_d \hat{d}_i \\ y_i = Cx_i \\ |y_i - T_{\text{ref}}| \le \beta_i \\ p_{\text{max}}^h \le p_i^h \le p_{\text{max}}^h \\ x_0 = \hat{x} \\ \forall i = 0, \ldots, N-1. \\ T\mathbf{p}^h = \mathbf{0} \end{array} \right. \right\} \tag{6.2}$$

where $\hat{x}$ is the initial state of the building, $\hat{\mathbf{d}}$ the current weather prediction, T_{ref} is a user-defined parameter which defines the ideal room zone temperature, and β_i is the maximum allowed deviation, that can be time-varying, of the zone temperature from the ideal value.

In the definition of $\mathcal{P}^h$, please notice the presence of a move blocking constraint (see *e.g.* [21]), $T\mathbf{p}^h = \mathbf{0}$. This constraint enforces that the electric power of the HVAC system is fixed over a certain number of time steps and is considered to encode the fact that, in general, it is not possible to adjust the power consumption too often either due to physical limitations of the equipment or to communication and mechanical latency that would inject significant fluctuations. Thus, the number of steps for which the power is blocked should be decided, depending on the particular application, in order to precisely track (on average) any given power setpoint. To provide an example, consider a building sampled at a 15 minute resolution for which the power consumption can be modulated only every 30 minutes. Assuming a horizon, N, equal to 6 steps, the matrix T would have the following form:

$$T = \begin{bmatrix} 1 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & -1 \end{bmatrix}$$

6.4.2 Electrical battery

In the experiments, a battery emulator was considered to propagate the battery SoC depending on the computed power injection.

Quantity	Value
Energy capacity	5 kWh
Power limits	-5/5 kVA
η	0.97
a	1

Table 6.3 – Technical specifications for the BESS considered in the experimental campaign.

In particular, at each fast sampling time (4 sec.), t, the battery was simulated by means of the following non-linear model:

$$\text{SoC}_{t+1} = a\text{SoC}_t + \eta(\text{p}_t^s)_+ + \frac{1}{\eta}(\text{p}_t^s)_-$$

where $(\text{p}_t^s)_+$ represents the power injected in the BESS and $(\text{p}_t^s)_-$ power extracted from the BESS. Both the cycle efficiency coefficient, η, as well as the temporal losses coefficient, a, have been identified based on experimental data for the Leclanché large-scale BESS on the EPFL campus. More information regarding the modeling of the battery can be found in [128] whereas all relevant parameters implemented in the BESS emulator are reported in Table 6.3. Please note that the technical specifications of the emulated BESS have been chosen by means of a simple simulation analysis, so as to have comparable capabilities for the the building and the battery. Thus, an in-depth optimal sizing of the BESS under the proposed control framework, has not been performed and could be an interesting line of future research.

BESS constraints

We assume that the power injection of the BESS can be modulated through a Battery Management System (BMS). Due to the high round trip efficiency of the battery, both in the Planning and Tracking Module the efficiency losses are neglected and the following model is considered:

$$\text{SoC}_{i+1} = \text{SoC}_i + p_i^s \tag{6.3}$$

Both SoC and the power injection are constrained to lay within a feasible operational range at each time instance, i:

$$\text{SoC}_{\min} \leq \text{SoC}_i \leq \text{SoC}_{\max} \tag{6.4}$$

$$p_{\min}^s \leq p_i^s \leq p_{\max}^s \tag{6.5}$$

The set of feasible power injections of the BESS is defined as

$$\mathcal{P}^s(\hat{\text{SoC}}) = \left\{ \mathbf{p}^s \left| \begin{array}{l} \text{SoC}_{i+1} = \text{SoC}_i + \text{p}_i^s \\ \text{SoC}_{\min} \leq \text{SoC}_i \leq \text{SoC}_{\max} \\ p_{\min}^s \leq p_i^s \leq p_{\max}^s \\ \text{SoC}_0 = \hat{\text{SoC}} \\ \forall i = 0, \dots, N-1. \end{array} \right. \right\} \tag{6.6}$$

where SoC_0 is the current SoC level as estimated by the BMS of the battery.

6.4.3   Total Power Consumption

As we are interested in combining these two resources into one single virtual entity, it is useful to introduce the total electrical power injection at time i, denoted as p_i, which is defined as the sum of the power injections at the two resources:

$$p_i = p_i^h + p_i^s \quad \forall i$$

Remark: Please notice that we consider a passive sign notation, *i.e.*, positive power values denote consumption.

6.5   Control Architecture

Referring to Figure 6.4, we provide herein an overview of the overall control architecture and how each component of it interacts with both the controlled resources as well as all external entities. The outputs and sampling times of the different modules in Figure 6.4 are given in Table 6.4.

A hierarchical control structure is considered which is comprised of the following modules.

Local Control & Measurements: We consider that both the building as well as the emulated battery are already equipped with local sensors and a control system which monitors and controls the low-level functioning of the resource. Regarding the building, as already mentioned, a proprietary controller monitors the building indoor temperature and decides on the activation of the HP based on the user-defined temperature setpoint. In the case of the emulated BESS, we assume the presence of a BMS which provides SoC estimation, management of the individual cells, and control of the local voltage and current constraints.

Forecasting Module: It is responsible for the prediction of all uncertain quantities that affect the performance of the system. Thus, this module is in charge of obtaining both weather as well as AGC power request predictions. These two quantities can be either locally computed, as in the case of the TSO future power requests, or simply retrieved from external weather stations, as in the case of weather forecasts. The outputs of this module are sampled with a sampling time of 15 minutes.

Module	Output	Sampling Time
Planning module	baseline power ($\bar{\mathbf{D}}$), and power capacity (γ)	1 day
Tracking Module High level controller Low level controller	power setpoint for building (p_i^h), and intraday transaction (τ_i) power setpoint for battery (p_t^s)	15 minutes 4 seconds
Forecasting module	AGC prediction scenarios ($\mathbf{r}$), and weather prediction ($\mathbf{d}$)	15 minutes

Table 6.4 – Outputs and sampling times of the modules in Figure 6.4.

More details can be found in Section 6.6

Planning Module: Activated once each day at a pre-specified time instant. Based on the most recent information, it solves an economic optimization problem to determine the aggregated power profile (day-ahead baseline) and the power capacity for the following day of operation. These quantities are then transmitted by means of an external secured channel to the Energy Market (baseline profile), and to the AS Market (power capacity) [1].

More information can be found in Section 6.5.1.

Tracking Module: Activated during the delivery of the service. It is in charge of performing a multi-timescale coordination of all the resources so as to track the power request coming from the TSO while ensuring local constraints at all resources. The tracking module consists of two controllers - high level and low level. The high level controller operates with a sampling time of 15 minutes and computes the setpoint for the slow resource local controller (BEMS) together with the intraday trades to modify the day-ahead baseline, while the low-level controller operates with a sampling time of 4 seconds and computes the setpoint for the fast resource local controller (BMS).

Please refer to Section 6.5.2 for more information.

6.5.1   Planning Module

This section presents the Planning Module which is activated at a pre-specified time instant each day to plan the next day of operation. More precisely, the objective of this module is to determine the day-ahead baseline which is the power profile that the set of resources should track in the absence of regulation service (normal operation). The baseline power profile is denoted by the sequence $\bar{D}_i$,

---

[1]We assume that the power flexibility is computed and transmitted daily. Clearly, other possibilities can also be encompassed, *e.g.* weekly bidding
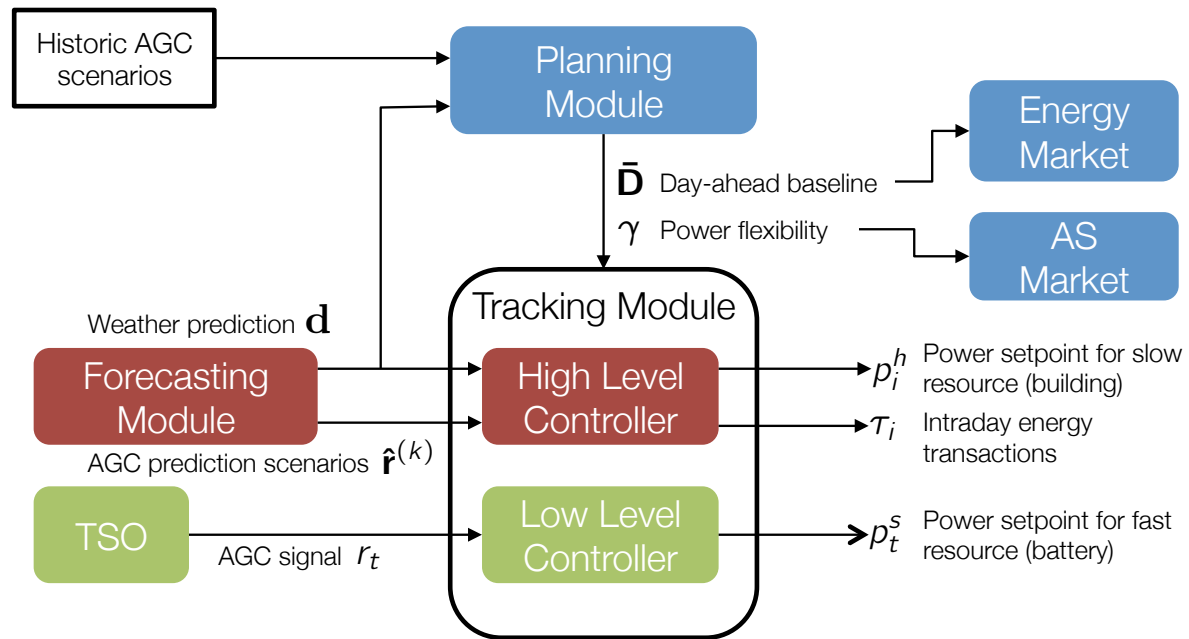
Figure 6.4 – Overview of the control architecture. The modules in blue, red, and green colors operate with a sampling time of one day, 15 minutes, and 4 seconds, respectively.

where the index i denotes a 15-minutes interval [2] for the next day of operation. The module also determines a power flexibility, γ, which represents the maximum allowable deviation around the baseline that the aggregation of resources is willing to sustain for the next day of operation.

This is done by solving the following adaptation of the planning problem 2 presented in chapter 4:

---

[2]The choice of the duration of this interval depends on the specific regulations for the country of interest. In the Swiss AS market, the baseline must be specified for each 15-minute slot in the day.

Problem 6 (Planning problem).

$$\underset{\gamma,\bar{\mathbf{D}},\pi_{p^h},\pi_{\mathbf{D}},\pi_{p^s}}{\text{minimize}} \quad \mathbb{E}_{\mathbf{r}}\left\{J(\gamma,\mathbf{D},\mathbf{r})\right\}$$

s.t.

(Building constraints)	$\mathbf{p}^h \in \mathcal{P}^h(\hat{x},\hat{\mathbf{d}})$	$\forall \mathbf{r} \in \Xi$	(6.7)
(BESS constraints)	$\mathbf{p}^s \in \mathcal{P}^s(\hat{\text{SoC}})$	$\forall \mathbf{r} \in \Xi$	(6.8)
(Total power)	$\mathbf{p} = \mathbf{p}^h + \mathbf{p}^s$	$\forall \mathbf{r} \in \Xi$	(6.9)
(Power tracking)	$\|\mathbf{p} - \bar{\mathbf{D}} - \pi_{\mathbf{D}}(\mathbf{r}) - \gamma\mathbf{r}\|_\infty \leq m_e\gamma$	$\forall \mathbf{r} \in \Xi$	(6.10)
(Power flexibility)	$\gamma \geq 0$		(6.11)
(Control Policies)	$\mathbf{p}^h = \pi_{p^h}(\mathbf{r}), \quad \mathbf{p^s} = \pi_{p^s}(\mathbf{r})$		(6.12)
(Intraday Policy)	$\mathbf{D} = \bar{\mathbf{D}} + \pi_{\mathbf{D}}(\mathbf{r}).$		(6.13)

where J represents the economic cost for the next day of operation which comprises different components such as the cost for the purchased baseline, the reward for the provided flexibility, a penalty cost for tracking violations, etc. As it depends on the uncertain parameter, $\mathbf{r}$, it is evaluated in expectation. A full overview of this cost function for the particular case of the Swiss market can be found in [118].

As usual, the equations (6.7) and (6.8) represent the constraints on the controllable resources. The total power consumption for the set of resources is defined in equation (6.9) whereas equation (6.10) imposes a minimum quality of tracking service. More precisely, it states that the aggregated power consumption during real-time operation, $\mathbf{p}$, should be adapted so as to guarantee a bounded tracking error:

$$|\epsilon_i| := |p_i - \bar{D}_i - \gamma\tau_i - \gamma r_i| \leq m_e\gamma \quad i = 1,\dots,N$$

where $\bar{D}_i$ is the committed baseline, and the term $\gamma\tau_i$ represents the possible modifications of the committed baseline through intraday energy transactions. Finally, the term $m_e\gamma$, determines the maximum allowable magnitude of tracking error in proportion to the offered capacity. In the Swiss market, SwissGrid fixes this term to 5% of the offered capacity, *i.e.*, $m_e = 0.05$.

The planning problem is approximated using the two-stage method of Section 4.6.3. This choice is motivated by the following reasons: 1) the main economic driver for this application is represented by the power flexibility that can be offered to the TSO [118]. As a consequence, we decided to consider a less conservative approximation scheme so as to maximize the offerable service; 2) only one uncertain parameter is considered, *i.e.* $\mathbf{r}$, meaning that the offline intraday policy (4.26) is sufficient in this case.

Thus, the multi-stage structure of the problem is relaxed to only two stages and a scenario-based

approach is considered to obtain a tractable reformulation. The intraday trades, $\boldsymbol{\tau} = \pi_{\mathbf{D}}(\mathbf{r})$, are optimized using the policy provided by (4.26). As the trades are fixed offline, they can be used to generate scenarios of residual tracking signal, $\mathbf{a}$, corresponding to given scenarios of the AGC signal, *i.e.*, $\mathbf{a}^{(k)} = \mathbf{r}^{(k)} + \boldsymbol{\tau}^{(k)}$.

Remark 12. *The intraday control policy defined in (4.26), is used only in the scheduling phase, not the tracking phase. During closed-loop operation, the system does not necessarily compensate for the energy biases of the AGC: the HLC of Problem 8 computes intraday actions that maintain the feasibility of operation by mitigating the wide range of uncertainty always present (occupancy, forecast errors, extreme AGC requests, etc..). For the bidding phase, we use the intraday policy detailed in (4.26) for tractability reasons as a proxy for an optimal strategy. Its primary purpose is to reduce the conservatism of the solution of the planning problem 1 against extreme realizations of the AGC signal, and allows to increase the available flexibility that otherwise would be limited by these few extreme realizations. In addition, we do not consider the opportunity of generating extra revenues trading energy on the intraday market, which is an entirely different topic that would require a much more detailed analysis of intraday price dynamics, volatility and liquidity. In the considered setup, adjusting the baseline is therefore only a way of ensuring the reliability of the tracking service.*

Scheduling algorithm

Once each day, the following steps are performed by the Planning Module to obtain the day-ahead baseline and power flexibility:

1. Retrieve the most recent weather forecast over the next day of operation.

2. Form the predicted vector of disturbances, $\hat{\mathbf{d}}$, over the considered prediction horizon.

3. Retrieve a set of filtered historic AGC signals, $\mathbf{a}^{(k)}$, as in (4.26).

4. Estimate the current state of the building, $\hat{x}$ by means of a standard Kalman filter.

5. Retrieve the current state of the battery, $\hat{\text{SoC}}$.

6. Solve the following tractable economic optimization problem:

Problem 7 (Tractable Planning Problem Formulation)**.**

$$\underset{\gamma,\bar{\mathbf{D}}}{\text{minimize}} \quad -\gamma$$

$$\text{s.t.}$$

$$\text{(Building constraints)} \quad \mathbf{p}^{h,(k)} \in \mathcal{P}^h(\hat{x}, \hat{\mathbf{d}}) \qquad \forall k = 0, \ldots, N_{\text{scen}} \qquad (6.14)$$

$$\text{(BESS constraints)} \quad \mathbf{p}^{s,(k)} \in \mathcal{P}^s(\hat{\text{SoC}}) \qquad \forall k = 1, \ldots, N_{\text{scen}} \qquad (6.15)$$

$$\text{(Total power)} \quad \mathbf{p}^{(k)} = \mathbf{p}^{h,(k)} + \mathbf{p}^{s,(k)} \qquad \forall k = 0, \ldots, N_{\text{scen}} \qquad (6.16)$$

$$\text{(Power tracking)} \quad \|\mathbf{p}^{(k)} - \bar{\mathbf{D}} - \gamma\mathbf{a}^{(k)}\|_\infty \leq m_e \qquad \forall k = 0, \ldots, N_{\text{scen}} \qquad (6.17)$$

$$\text{(Power flexibility)} \quad \gamma \geq 0 \qquad (6.18)$$

$$(6.19)$$

7. Transmit the optimal solution to the Energy Market (baseline) and the AS Market (power flexibility).

8. Wait for the next iteration.

Please notice how, in Problem 7, the intraday control policy has been directly incorporated in the residual signal, $\mathbf{a}^{(k)}$ as detailed in equation (4.26).

Finally, the general form of the cost function J introduced in Problem 2 has been simplified so that the problem only tries to maximize the offered power capacity. This was done mainly for the sake of simplicity. However, as detailed in [118], this is also desirable from an economic perspective as most of the economic benefits come from the reward for the offered capacity. We highlight that more complex cost functions can be easily incorporated.

The solution of the planning problem results in the optimal value of the baseline power $\bar{\mathbf{D}}^*$, and flexibility γ^*.

6.5.2 Tracking Module

Every day at 00:00 the committed baseline and capacity, computed by the Planning Module, is activated and the set of resources are required to deliver the SFC service every 4 seconds. This is done by means of a Tracking Module that optimally coordinates the two resources so that their total power consumption is equal, within the allowed error bounds, to the sum of the total baseline and the AGC signal scaled by the committed power capacity.

The Tracking Module is composed of two controllers continuously working at different timescales, one at 15 minutes and the other one at 4 seconds[3]. On one side, the High-Level Controller (HLC)

---

[3]These sampling times are specific to the Swiss market and will, in general, depend on the specific regulations for the SFC provision, and for the intraday market

decides the setpoints for the slow resource (building) as well as the future energy transactions to place in the intraday market. On the other side, the Low-Level Controller (LLC) computes the power injection of the BESS so that the set of resources precisely track the received AGC signal. In the following sections, the formulation of the two controllers is detailed.

High-level

The HLC operates at a 15 minute time step. It is responsible for computing an adequate power setpoint for the building for the following 15 minutes so as to: 1) respect comfort and operational constraints of the building 2) guarantee that the BESS will be operated within its operational constraints. Moreover, depending on the current status of the resources and the current AGC predictions, the HLC can decide to sell or buy energy in the intraday market so as to either guarantee a high tracking quality and to improve the overall efficiency of the system. This is done by continuously running the following algorithm at every slow iteration i:

1. Retrieve the most recent weather forecast over the considered prediction horizon, N.

2. Form the predicted vector of disturbances, $\hat{\mathbf{d}}$, over the prediction horizon.

3. Retrieve a set of possible scenarios, $\hat{\mathbf{r}}^{(k)}$ for the AGC signal from the Forecasting Module.

4. Estimate the current state of the building, $\hat{x}_i$ by means of a standard Kalman filter.

5. Retrieve the current state of the battery, $\hat{\mathrm{SoC}}_i$.

6. Solve the following Model Predictive Control (MPC) problem:

   Problem 8 (Tracking MPC Problem)**.**

   $$\underset{p_0^h, \boldsymbol{\tau}}{\text{maximize}} \quad \mathbb{E}_{\hat{\mathbf{a}}}\left\{J_{comfort}\right\}$$

   $$\text{s.t.}$$

   $$\text{(Building constraints)} \quad \mathbf{p}^{h,(k)} \in \mathcal{P}^h(\hat{x}_i, \hat{\mathbf{d}}) \qquad \forall k = 0, \ldots, N_{\text{scen}} \quad (6.20)$$

   $$\text{(BESS constraints)} \quad \mathbf{p}^{s,(k)} \in \mathcal{P}^s(\hat{\mathrm{SoC}}_i) \qquad \forall k = 1, \ldots, N_{\text{scen}} \quad (6.21)$$

   $$\text{(Total power)} \quad \mathbf{p}^{(k)} = \mathbf{p}^{h,(k)} + \mathbf{p}^{s,(k)} \qquad \forall k = 0, \ldots, N_{\text{scen}} \quad (6.22)$$

   $$\text{(Power tracking)} \quad \|\mathbf{p}^{(k)} - \bar{\mathbf{D}}^* - \boldsymbol{\tau} - \gamma^* \hat{\mathbf{r}}^{(k)}\|_\infty \leq m_e, \qquad \forall k = 1, \ldots, N_{\text{scen}} \quad (6.23)$$

   $$(6.24)$$

7. Send the computed setpoint, $(p_0^h)^*$, to the local controller of the building

8. Wait for the next iteration

In Problem 8, the quantity $\hat{\mathbf{r}}^{(k)}$ represents the scenarios of the future AGC signal (over the prediction horizon) as generated by the AGC predictor (Section 6.6.2), the cost function $J_{comfort}$ is user-defined and it can comprise different metrics as, *e.g.*, the comfort quality for the building occupants.

The solution of the MPC problem is the power setpoint, p_0^h, for the HVAC system which is sent to the local controller in charge of tracking this for the following 15 minutes. Moreover, the problem also determines the adjustment of the committed baseline profile, τ, which are placed as energy trades in the intraday market. Note that the first few time steps of τ are fixed by the previous iteration of the MPC controller and not an optimization variable to make sure that the appropriate delay of the intraday market is respected (45 minutes in the Swiss market).

Low-level

The low-level tracking controller computes the control input for the fast resource (BESS). Every 4 seconds, it measures the actual power consumption of the slow resource which, in general, might have small deviations with respect to the setpoint as computed by the HLC. Based on this information, the LLC computes the power setpoint at each fast iteration, t, for the BESS as follows:

$$p_t^s = \bar{D}_t^* + \tau_t^* + \gamma^* r_t - \tilde{p}_t^h \tag{6.25}$$

where τ_t^* is the intraday trade fixed by the HLC, r_t is the received AGC signal, and $\tilde{p}_t^h$ is the measured power consumption of the slow resource (building) at fast time step t.

The computed power setpoint for the BESS is then transmitted to the BMS controller which will be responsible to implement this given that all physical constraints of the BESS are respected. If this is not the case, the power injection is rejected and a tracking error will occur.

6.6 Forecasting Module

6.6.1 Weather forecast

Weather forecasts can be typically obtained through available web-services by simply specifying the geographic location of interest. This requires one to select the closest weather station in the same bioclimatic zone as the installation. In the experimental demonstrator of Section 6.4, the weather station was selected from four available stations in a 3 Km radius, based on the historical quality of the data it provides. (source MeteoSwiss, meteostation *Lausanne freiland* , GPS coordinates $6°38.56'$ $46°33.33'$). The forecasting module retrieves the most updated forecasts every 15 minutes for the next 24h.

6.6.2 AGC Predictor

The ability to anticipate future values of the power requests dispatched by the system operator helps greatly in improving the average economic performance of the controller. We start by highlighting some characteristics of the AGC:

- The AGC signal follows a periodic pattern with recurring daily and hourly patterns, as already discussed on the basis of Figure 6.1.

- The empirical probability distribution of the AGC is non Gaussian, and time-varying

- The AGC signal is strongly correlated over time.

In the view of using forecasts of the AGC to solve the multi-stage real-time MPC Problem 8, the availability of a probabilistic forecast of the AGC is desirable, as opposed to a simpler point forecast. These probabilistic forecasts should capture the time-correlated nature of the AGC appropriately. In view of these elements, a method based on a variation of the non-parametric probabilistic forecasting method presented in [115] has been adopted and is presented in this section.

We denote with r_t the AGC power request at time t as the realization of a random variable R_t. Following the observation that the mean of the AGC presents a strong daily pattern, its distribution around that mean is consistent over time, we use the following assumption:

$$r_t = \bar{r}_t + \tilde{r}_t, \quad \forall t \tag{6.26}$$

where $\bar{r}_t \in \mathbb{R}$ captures the daily mean, so that $\bar{r}_t = \bar{r}_{t+k*24h}$ for all k and $\tilde{r}_t$ originates from a single generating random variable denoted $\mathcal{R}$. The AGC is therefore generated by the sum of a daily mean and a single generating random variable $\mathcal{R}$.

We denote with f the probability density function of $\mathcal{R}$ and F the corresponding cumulative distribution function. Assuming F to be strictly increasing, we define the quantile of $\mathcal{R}$ at level α as $q^\alpha = F^{-1}(\alpha)$. A forecast of this quantile is denoted $\hat{q}^\alpha$. A non-parametric forecast of the density function is formed by collecting quantile forecasts as:

$$\hat{f} = \{\hat{q}^{\alpha_i} | 0 \leq \alpha_1 < \ldots < \alpha_m \leq 1\} \tag{6.27}$$

Based on m observed realizations $(\tilde{r}^{(k)})_{k=1,\ldots,n}$ of $\mathcal{R}$, unbiased estimates of the quantiles are formed as the empirical quantiles of the observed realizations, ie:

$$\hat{q}^{\alpha_i} = \min_x \frac{\#\{\tilde{r}^{(k)} < x\}}{n} \geq \alpha_i \tag{6.28}$$

In turn, $\hat{F}$ is obtained as the linear interpolation of the empirical cumulative distribution function between levels $(\alpha_i)_{i=1,\ldots,m}$.

It is desirable that the estimated and observed quantiles are as close as possible, and should asymptotically match exactly. This property is referred to as reliability. When the estimates $\hat{f}$ are reliable, the random variable Y whose realization at time t is defined as:

$$Y_t = \hat{F}(\tilde{r}_t) \tag{6.29}$$

is uniform, that is $Y \sim U([0, 1])$.

We will exploit the fact that predictions can be made Gaussian using a suitable transformation.

We can obtain a normally distributed variable function X from Y as follows:

$$X_t = \Phi^{-1}(Y_t), \ \forall t \tag{6.30}$$

where Φ^{-1} is the probit function defined as $\Phi^{-1} : p \to \sqrt{2}\,\mathrm{erf}^{-1}(2p-1)$ and erf^{-1} the inverse error function. This yields $X_k \sim \mathcal{N}(0,1)$, Gaussian-distributed with 0 mean and variance 1.

To capture the time correlation of the AGC, we assume for each time t that the random vector $\mathbf{X}_{t-K:t+N-1} = (X_{t-K}, \dots, X_{t+N-1})$ follows a multivariate Gaussian distribution $\mathcal{N}(\mathbf{0}, \Sigma_t)$, with K and N a horizon in the past and future that are long enough to capture the time-correlation of the AGC. It is assumed that N and K are fixed and we simply refer to $\mathbf{X}_{t-K:t+N-1}$ as $\mathbf{X}_t$. Once more, it is assumed based on observations that all $\mathbf{X}_t$'s originate from a single generating multi-variate Gaussian random variable $\mathcal{X}$. We can then estimate its covariance Σ by using past observations. An unbiased estimate is given by:

$$\hat{\Sigma} = \frac{1}{n-1} \sum_{t=1}^{n} \mathbf{X}_t \mathbf{X}_t^\top \tag{6.31}$$

The following procedure is then used to generate scenarios at time t.

1. The observations $\mathbf{r}_{t-K:t-1} = (r_{t-K}, \dots, r_{t-1})$ for the last K time steps are collected.

2. The corresponding realization of $X_\mathcal{I}$ for $\mathcal{I} = t-K, \dots, t-1$ are computed by removing the mean and transforming the result to normal as:

$$X_\mathcal{I} = \Phi^{-1}(\hat{F}(r_\mathcal{I} - \bar{r}_\mathcal{I})) \tag{6.32}$$

3. The parameters of the multi-variate Gaussian characterizing the prediction $(X_t, \dots, X_{t+N-1})$: $\hat{\mu}_{t:t+N-1}$ and $\hat{\Sigma}$ are calculated by marginalizing the estimated distribution $\hat{\mathbf{X}}_t$ with respect to the observation vector $(X_{t-k}, \dots, X_{t-1})$, which gives:

$$\begin{aligned} \hat{\mu}_N &= \hat{\Sigma}_{NK}\hat{\Sigma}_{KK}^{-1}\mathbf{X}_{t-K:t-1} \\ \hat{\Sigma}_N &= \hat{\Sigma}_{NN} - \hat{\Sigma}_{NK}\hat{\Sigma}_{KK}^{-1}\hat{\Sigma}_{NK}^\top \end{aligned} \tag{6.33}$$

where $\hat{\Sigma} = \begin{pmatrix} \hat{\Sigma}_{KK} & \hat{\Sigma}_{NK}^\top \\ \hat{\Sigma}_{NK} & \hat{\Sigma}_{NN} \end{pmatrix}$

4. N_{scen} scenarios $(\mathbf{X}_{t:t+N-1}^{(i)} = (X_t^{(i)}, \dots, X_{t+N-1}^{(i)}))_{i=1,..,N_{\text{scen}}}$ of this marginal are sampled in an iid fashion.

5. The inverse probit function is applied to each component of the scenarios to obtain $Y_\mathcal{I}^{(i)} = \phi(X_\mathcal{I}^{(i)})$, $\forall \mathcal{I} = t, \dots, t+N-1$, $i = 1, \dots, N_{\text{scen}}$.

6. The inverse of the estimated cumulative distribution function is finally used to compute the final forecast as $\hat{r}_\mathcal{I}^{(i)} = \hat{F}^{-1}(Y_\mathcal{I}^{(i)}) + \bar{r}_\mathcal{I}$
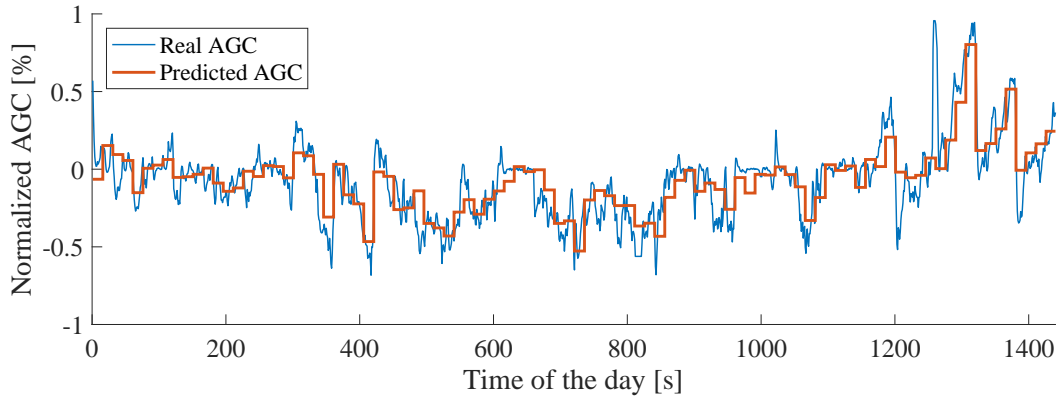
Figure 6.5 – Performance of the AGC predictor for one randomly selected AGC signal. The one-step-ahead AGC prediction as used in the HLC (red) on a 15-minutes timescale vs the actual realization of the AGC (blue).

This together provides forecasts for the AGC for the N next steps.

We have two years of AGC data available from the period 2013-2014. We use the data from the year 2013 to estimate the cumulative distribution function of the AGC, as well as its daily mean. The data from the year 2014 is kept for validation. Figure 6.5 illustrates the performance of the predictor to predict the first 15 minutes of the AGC signal. It shows a one-day realization together with the 15 minutes-averaged prediction generated by the predictor every 15 minutes for the next 15minutes. This is the prediction effectively taken into account in the MPC applied control actions. We see that it can capture the trend quite successfully.

Figure 6.6 shows the root mean square of the normalized prediction error comparing the predictor described in this section with two basic predictors. The Mean Predictor predicts future values of the AGC to be the corresponding daily mean $\bar{\mathbf{a}}$. The Persistent Predictor predicts that the future value of the AGC is the current observed value, representing the fact that the AGC is correlated over time. We see that our predictor achieves better predictions over the whole horizon, especially for horizons below 30 minutes. To further study the impact of having improved predictions in real-time, we conducted a 30 day-long simulation considering the provision of SFC with a 5 kWh battery. By focusing on a single device, it is possible to better isolate the effect of the predictor on the closed-loop performance. The considered evaluation metric is the tracking error for a fixed provision of 3.4 kW. As can be observed in Table 6.5, simpler predictors, such as the Mean Predictor or the Persistent Predictor, directly result in higher tracking errors and, consequently, higher economic penalties.

Remark 13. *Note that the prediction performance in Figure 6.6 is reported for the data of year 2014 while the data of the year 2013 was used for learning. While there are small statistical difference form one year to the next, the results show that the predictor generalizes quite well from one year*
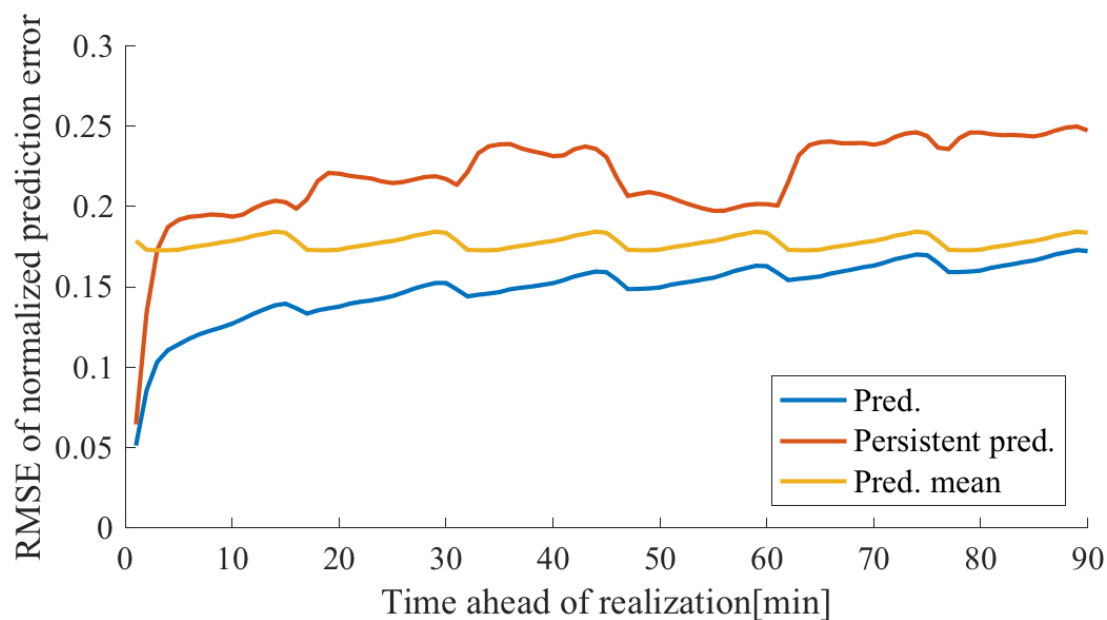
Figure 6.6 – RMSE of the normalized prediction error as a function of lead time of prediction

Predictor	Pred.	Mean Predictor	Persistent pred.
Tracking error [kWh / day]	0	0.057	0.057

Table 6.5 – Closed loop tracking performance using different predictors. The tracking error is measured as the integral of the tracking errors in kWh normalized by the number of days

to another one.

6.7 Experiments

In this section, we describe the experimental campaign that was conducted in August 2017. To fully test the robustness and reliability of the controller, we run a set of nine days of continuous experiments emulating the participation of the combined system (BESS + building) into the Swiss SFC according to the current regulations. Thus, we consider historic AGC for the year 2013/2014 that was obtained from the Swiss TSO, Swissgrid. This large set of data is split into two subsets. The first subset is used to construct the uncertainty set (4.23) used in the Planning Module as described in Section 4.6.3. The second subset is used to randomly pick a realization of the AGC signal that is used during the closed-loop operation.

Each day at 23:00, the Planning Module collects the most current weather predictions from the nearest MeteoSuisse weather station. The weather forecasts are comprised of outside temperature, and solar radiation over a 24h period. The Planning Module further retrieves the current SoC level of the BESS and the most updated state estimate of the building thermodynamics model from a continuously running Kalman filter. It then solves the planning problem (7), and determines the power baseline at a 15-minutes resolution, and power flexibility for the next 24 hours starting at 00:00 of the following day.

At 00:00 each day, the delivery begins and the Tracking Module is activated. It receives from the Planning Module the committed power profile, $\bar{\mathbf{D}}^*$ and capacity γ^*. As described in Section 6.5.2, at each slow iteration (15 minutes), it collects the current weather predictions over the prediction horizon together with the current state for each resource. The Tracking Module also obtains a set of 350 possible AGC scenarios generated by the AGC Predictor. It then solves the HLC problem with a horizon of 6 hours to determine both the power setpoint p^h for the controllable building as well as the amount of energy to trade on the intraday market τ. Finally, the LLC (6.25) computes the power injection setpoint for the BESS so as to track the received AGC signal.

Figure 6.7 displays one full day of operation. In particular, the topmost plot shows the pre-computed baseline power profile as computed by the Planning Module, the AGC request scaled by the bid capacity, and the power realization of the set of resources, *i.e.*, the HP and the BESS, during real-time operation. As can be observed, the Tracking Module is able to optimally coordinate the two resources so as to perfectly track the received AGC signal. This is done while, at the same time, respecting the physical limitation of each resource. In particular, it can be observed in the third and fourth subplot of Figure 6.7 how both the SoC and the indoor temperature of the building are within their operational constraints at all times. In the second subplot, the allocation of the power, as performed by the HLC of the Tracking Module, is displayed. In particular, it can be observed how the HLC determines, at the 15 minutes time-scale, the power setpoint to the local controller of the HP so as to simultaneously preserve comfort inside the building and reset the SoC

Quantity	BESS	BESS + HP
Avg. capacity offered	3.4 [kW]	5.96 [kW]
Max. capacity offered	3.4 [kW]	6.72 [kW]
Min. capacity offered	3.4 [kW]	4.38 [kW]

Table 6.6 – Statistics of the offered capacity over a period of 9 continuous days of experiments in August of 2017.

of the battery. This is particularly evident, for instance, between minutes 250 and 400 where the SoC is quite close to its upper limitation. At this time, the HLC decides to increase the power consumption of the HP with respect to the baseline profile in order to have a net discharging effect on the BESS and, therefore, reset its SoC to a safer value.

6.7.1 Multiplier effect

The statistics of the offered capacity, γ, overall experimental days are reported in Table 6.6. The comparison of the offered capacity between the case of a single BESS and the combination of the BESS and commercial building is performed. Please note that the offered capacity for the combined system (BESS + HP) depends on external varying factors affecting the system such as weather, occupancy pattern, etc. For this reason, the offered capacity displayed in Table 6.6 has some variability from one day to the next. On the contrary, the offered capacity with the BESS alone is computed as in Problem 2, where the total power (6.10) is obviously modified as $\mathbf{p} = \mathbf{p}^s$. Thus, the offerable capacity is only limited by the operational constraints of the battery and does not depend on other external conditions. This explains why the offered capacity is always the same each day. The results show that an overall substantial improvement is obtained through the proposed control method when combining a slow energy-intensive and a fast power-intensive resource in terms of the offered capacity to the grid. The experimental results confirm that exploiting the synergy between the slow and fast resources can increase the overall flexibility that can be provided to the grid.

Remark 14. *Please note that the case of the building and the BESS acting as independent providers on the AS market has not been considered in the comparison. This is due to the fact that, due to operational limits of the compressor, the power consumption of the HP cannot be modulated at a frequency higher than 1/15 minutes which clearly does not allow the building to track tracking requests on a 4-second timescale. As already mentioned, this represents one of the main strengths of the proposed method which allows virtually any type of building to participate in such applications irrespective of the particular HVAC system installed.*

6.7.2 Cost of Operation

We present here a comparison between two sets of experiments. In the first set of experiments (labeled as **Case A**), the building and the battery are operated in coordination to provide ancillary
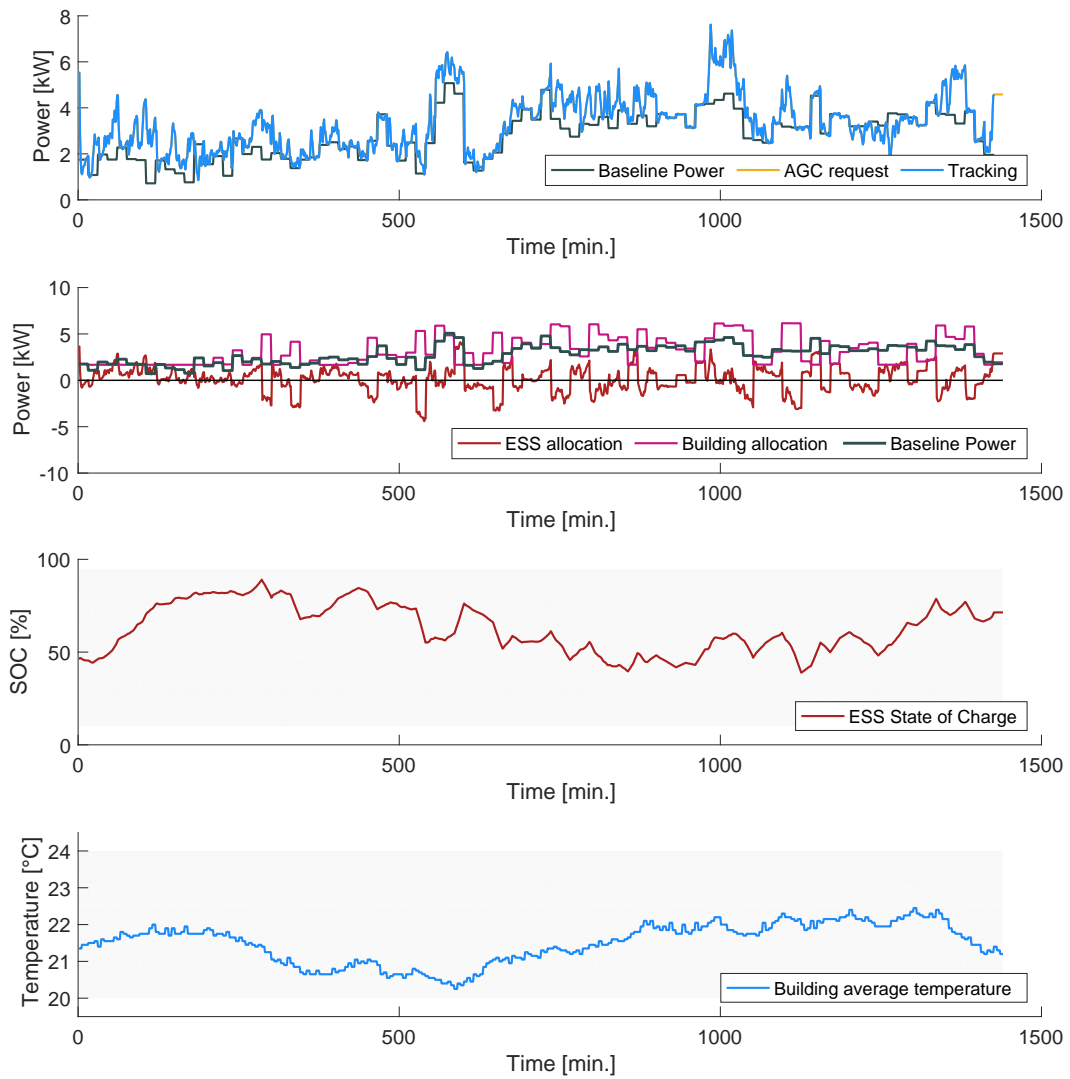
Figure 6.7 – Experimental results on the EPFL campus: One full day of operation emulating the participation to the SFC in the Swiss Market. Legend: The grey areas represent the operational limits of the resources.

services, as was described in the previous sections. In the second set of experiments (**Case B**), the building and the battery are operated independently. Since providing ancillary services requires fine control of the system, the building is not able to provide ancillary services alone. Therefore, we assume in the second set of experiments that the battery is operated to provide ancillary services, while the building is operated "normally", with a default controller which is designed to keep the temperature within a small range around the ideal temperature (22 ºC). To have a fair comparison, we compare days for which weather conditions are very similar. In particular, we considered for experiment set 'B' days for which the average outside temperature and solar radiation were in the same range as the minimum and maximum average values seen during the experiments form set 'A'.

The building needs to interact with different energy markets while providing SFC. A detailed description of the energy markets and a comprehensive economic analysis for the provision of SFC was performed in our previous publication [118]. Therefore, the focus here is not to perform an in-depth economic analysis, but rather to get an estimate of the operating costs in both cases.

There are four components of the total operational cost when providing SFC service. (1) Baseline energy cost: the cost of the energy purchased in the day-ahead and intraday markets; (2): Capacity Reward is the reward for providing the SFC capacity in the ancillary services market. (3): Tracking Penalty is the penalty for the errors in tracking the AGC signal. (4): Cost of tracking accounts for the fact that tracking the AGC results in a net energy consumption increase or decrease, that still need to be paid for, though at a discounted rate (note that this cost of tracking can be positive or negative depending on the trend of the AGC with respect to 0).

All the cost components for each experiment day in set 'A' are computed using the average Swiss price data for August 2014 to get an estimate of the total operating cost. The results are shown in Figure 6.8. Similarly, the total operating cost in set 'B' is shown in Figure 6.9, where the capacity reward and tracking cost are tied to the battery alone. The comparison shows that on average the operational cost is 5.94 CHF per day while providing SFC, and 8.15 CHF per day during normal operation. On average providing SFC using both the building and the battery results in about 26% reduction in the operating costs, thanks to a higher capacity reward, as explained in Section 6.7.1. Moreover, as a fixed average energy price is considered in the analysis, the baseline energy cost also provides an insight in terms of kWh used in the two configurations. Interestingly, the average baseline energy consumption does not significantly vary in the two cases with 9.56 kWh in scenario A and 10.23 kWh in scenario B.

6.7.3 Comfort

We studied the impact of providing SFC service to the grid on occupant comfort. To evaluate the occupant comfort, we used a standard measure, "Ashrae Likelihood of Dissatisfaction" (ALD) used in literature [25] (reported in the Appendix). ALD is a function of the absolute difference between the room temperatures and an ideal (most comfortable) temperature defined by the occupants. It
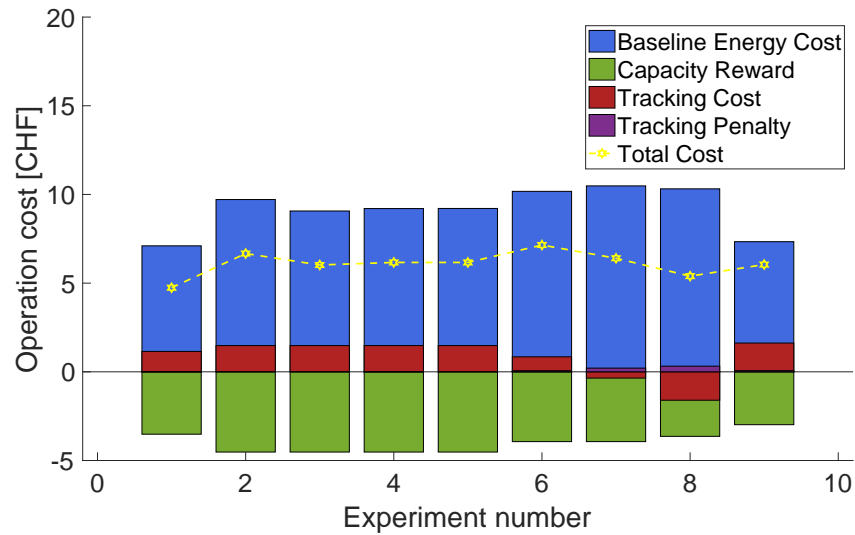
Figure 6.8 – Break down of all economic costs for the 9 days of experiments for **Case A** with the battery and the building operated together to provide SFC.
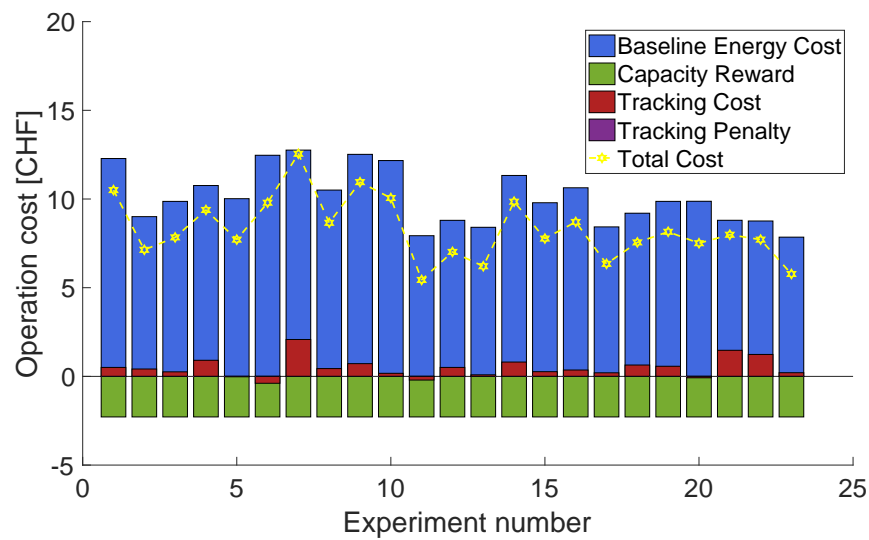

Figure 6.9 – Break down of all economic costs for the 23 days of experiments for **Case B** with the battery providing SFC and the building in normal operation.

represents the average percentage of people dissatisfied in an environment. The lower the ALD value, the higher the occupant comfort.

The temperature trajectories of each day in the experimental sets (both Case A and Case B) are used to compute the ALD comfort measure for each day. Moreover, the cost of operation is also computed, as explained in the previous section, taking into account all the cost components in accordance with the Swiss energy markets.

The results are depicted in Figure 6.10. Each day is denoted by a marker in the "cost of operation" against "comfort" plane. The blue markers represent the Case A days, while red markers represent the Case B days. Please notice that results can vary depending on the particular set of uncertainties acting on the system (received AGC, occupancy, outside temperature, etc.). Nevertheless, it appears that the blue markers are on average lower than the red ones for a comparable level of comfort. Given that each datapoint represents an entire day of operation, the total number of datapoints is limited and there is significant scatter in the data. Nonetheless, a clear pattern seems present and it can be concluded from the result that a similar level of occupant comfort can be provided at a reduced operating cost while providing SFC. In other words, the aggregated provision of SFC for complementary resources not only has the potential to improve the operational cost of the building and battery but also it does not affect the comfort level for the occupants.

6.8 Discussion on Scalability

This chapter proposes a predictive multi-timescale control framework to schedule and control the operation of a set of complementary resources such as commercial buildings and BESSs. Relying solely on the availability of dynamical models, such as (6.2) and (6.6), that describe the interactions between manipulated (zones temperature setpoints, massflow rates, etc.) and dependent (electrical consumption, zones temperature, etc.) variables, the method could, in principle, be applied to any real-life setting. However, obtaining reliable models, especially in the case of large commercial buildings with multiple zones and complex HVAC systems, can prove to be a very challenging task. As a matter of fact, this typically represents one of the main limitations in the application of predictive control to building climate control. Many research efforts have been made to address this problem using both first principle and data-driven methods such as in [122, 131, 149], only to cite a few. Depending on the particular application, one of these modeling techniques could be considered in our framework in order to obtain the required resource models.

Likewise, some of the recently introduced strategies to provide fast-regulating services using commercial buildings [14, 76, 133, 139], could be readily incorporated within our framework. These techniques typically focus on specific HVAC components that can be controlled in order to comply to the strict requirements for frequency regulations. However, a considerable advantage of the proposed method with respect to existing ones is that, by optimally splitting the TSO tracking requests into different components, the method can best fit the technical capabilities of each resource
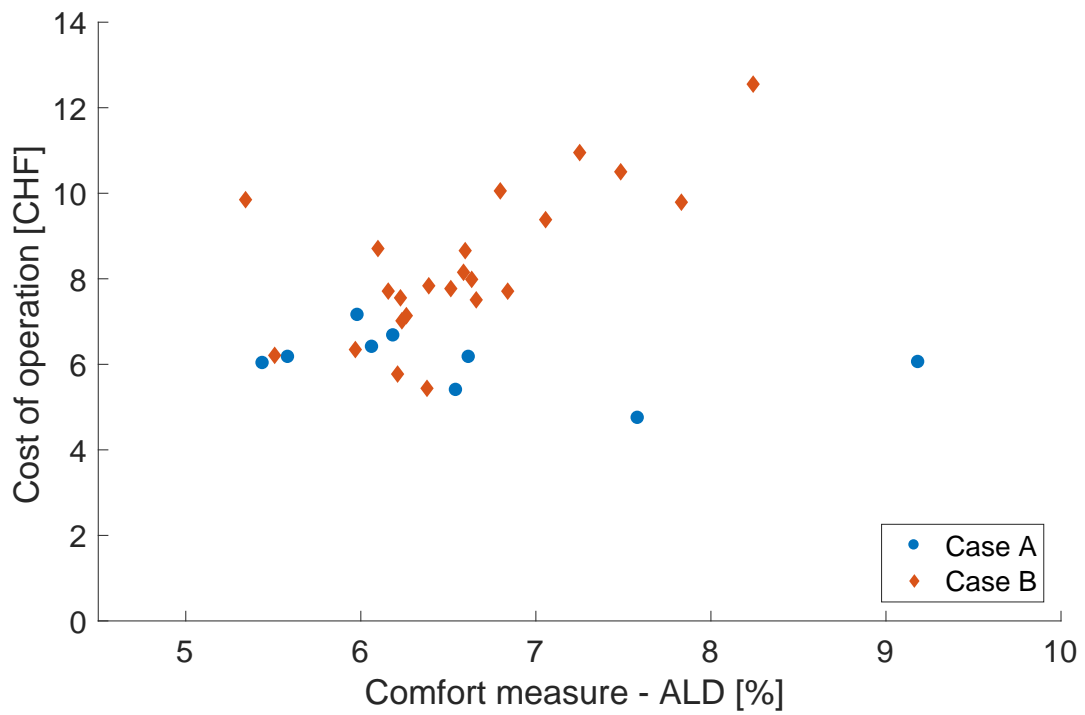
Figure 6.10 – Experiments: Cost of operation vs Comfort

according to its model and operational constraints. Thus, even very simple DR strategies could be considered as long as they are able to produce a net charging/discharging effect on the BESS in order to restore its SoC. Clearly, the degree of control and type of HVAC have a direct consequence on the required technical specifications of the BESS. This analysis has not been considered in this chapter and it represents an interesting direction for the future research.

6.9 Conclusion

In this chapter, we described a multi-level control framework to schedule and coordinate the operation of a set of heterogeneous resources, such as a building (slow) and an electrical battery (fast), in order to provide fast frequency regulation services.

The proposed control strategy was deployed on an experimental setup consisting of a fully-occupied controllable building on the EPFL campus, and an emulated grid-connected energy storage system. A set of nine days of continuous operation were conducted, emulating the provision, in full compliance with the Swiss regulations, of secondary frequency control.

The obtained results illustrated that: (a) the framework appropriately computed the aggregated flexibility that the set of resource can reliably offer to the grid; (b) the multi-timescale real-time controller yielded excellent tracking performance while respecting operational constraints of both resources; (c) combining complementary resources and exploiting their synergy can augment significantly the flexibility that can be provided to the grid with respect to the case of operating them separately; (d) in current Swiss market conditions, such an aggregation of resources can lead to a significant 26% reduction in operational costs.

Appendices

Appendix D

Comfort Modeling

One of the most important objectives of building control is to maintain or improve occupant comfort. Comfort is a human's perception of his environment, and therefore is difficult to measure. This perception of comfort is different for different people and might also vary for the same person at different times. Various measures of comfort have been reported in the literature, *e.g.*, the Predicted Mean Vote (PMV), the Predicted Percentage of Dissatisfied (PPD), etc. PMV is based on the model developed by Fanger [17] and is the predicted mean point rated by a large group of people. It is based on heat balance equations and empirical data that rates how a person would feel about a thermal condition. PPD is a function of PMV and analytical equations have been developed for this relationship [86]. The analytical equations defining PMV and PPD are complicated and are a function of many parameters, *e.g.*, operative temperature, relative humidity, air velocity, metabolic activity, and clothing resistance, etc. Therefore, it makes them difficult to use for control design.

Another similar, but slightly simpler measure developed by ASHRAE via a logistic regression analysis performed on the data collected in the ASHRAE RP-884 database is called ASHRAE Likelihood of Dissatisfied (ALD) [86] and is defined in literature as

$$\text{ALD}(T) = \frac{e^{0.008T^2 + 0.406T - 3.05}}{1 + e^{0.008T^2 + 0.406T - 3.05}} \in [0.05, 1.00) \tag{D.1}$$

where $T = |T_{\text{zone}} - T_{\text{comfort}}|$, T_{zone} is the zone temperature, and T_{comfort} is the optimal comfort temperature. Unlike, PMV and PPD, ALD is only a function of the absolute difference between the zone temperature and the optimal comfort temperature.

All these measures are for a specific building zone and for a specific point in time. A measure called Long-term Percentage of Dissatisfied (LPD) has been proposed for an average value of comfort throughout the building [86]. It accounts for the hourly-predicted ALD calculated for each zone and is weighted by the number of people inside the zone, and over time and is given as

$$\text{LPD}(\text{ALD}) = \frac{\sum_{t=1}^{T} \sum_{z=1}^{Z} (p_{t,z} \text{ALD}_{t,z})}{\sum_{t=1}^{T} \sum_{z=1}^{Z} (p_{t,z})} \tag{D.2}$$

where $\text{ALD}_{t,z}$ and $p_{t,z}$ are the ALD and normalized occupancy of the zone z at time t.

Although ALD is only a function of zone temperatures and is simpler than PMV and PPD, it is still difficult to use for control design because it is non-linear. In most of the MPC based control design found in the literature, the comfort is usually defined by a bound of temperatures around the optimal comfort temperature resulting in convex constraints for the MPC optimization problem. However, ALD together with LPD can easily be used in the post-processing to evaluate the occupant comfort.

Chapter 7

Enhancing the dispatchability of distribution networks through utility-scale batteries and flexible demand

7.1 Introduction

Motivated by the increasing need of control reserves, in the previous chapter, we provided an experimental validation of the technical feasibility of offering frequency support services to the grid using "unconventional" energy resources such as commercial buildings and BESSs.

In this chapter, we look at the same problem but from a different perspective. More specifically, rather than focusing on enlarging the set of potential providers of centralized services (e.g. SFC), the aim is to reduce the need for such control reserves altogether by absorbing the uncertainty associated to a smaller portion of the grid.

Thus, we focus on a paradigm, the so-called *dispatchability of distribution feeders* [127], that aims to achieve virtually perfect dispatchability of a distribution feeder comprising both uncontrollable loads as well as distributed generation (*prosumers*). More specifically, the proposed framework requires computing one day in advance a forecast power profile for the aggregated prosumers. During real-time operation, in order to track the committed profile, a BESS is operated by controlling its power injection in order to absorb any errors in the forecasts. Clearly, the success of such a control scheme relies on two factors: 1) an accurate forecasting tool for predicting the power profile of the prosumers one day in advance; 2) an appropriate energy/power sizing of the BESS.

For a given maximum prosumer power, it is desirable to achieve dispatchability with the smallest possible battery, as, despite the recent drop in prices, BESSs remain expensive devices. In the

current framework, the required specifications for the battery are dictated by the quality of the forecasting tool. However, achieving high-forecasting accuracy can prove quite challenging at a local scale where the effect of environmental factors and human behavior is not smoothed out as happens at higher levels of aggregation.

The objective of this chapter is to investigate the benefit of having, within the prosumer portfolio, an additional controllable resource that displays complementary technical capabilities with respect to the BESS. More specifically, as done in Chapter 6, we consider the presence of a commercial building whose electrical power consumption can be partially or fully controlled. Thus, the main contributions of these chapter are: 1) provide the adaptation of the control framework proposed in Part II for a particular case of practical interest; 2) design a hierarchical real-time controller to accurately track the committed dispatch plan while respecting the constraints of the building and the BESS; 3) quantify the benefit of aggregating complementary resources for a specific application; 4) provide an experimental demonstration of the control method of Chapter 4 on a real-scale 20 kV medium voltage distribution network which is comprised of a utility-scale battery, a fully occupied office building (LADR), conventional uncontrollable loads, and a roof-top PV generator.

7.1.1 Structure of the chapter

The rest of the chapter is structured as follows. In Section 7.2 we present the formulation of the dispatchability problem. In Section 7.4 we show how the general framework of Part II can be adapted to this application. Section 7.5 describes the real-time controller. In Section 7.6 a simulation analysis of the proposed algorithm is presented. Finally, in Section 7.7 the control architecture is experimentally validated on a 20kV medium voltage active distribution network.

Notation: The notation $\mathbf{p}^{\text{res}}$ indicates the real power flow of the particular resource, res, whereas the bracket subscript notation, $\mathbf{p}^{(k)}$ stands for the power trajectory corresponding the k-th scenario. Please refer to the Table 7.1 for the list of all mathematical symbols used in the chapter.

7.2 Problem Statement

The main objective of this chapter is to ensure that the aggregated power consumption of a cluster of resources follows a pre-established active power profile. The considered scenario is well exemplified by the experimental configuration depicted in Figure 7.1, which comprises three key elements:

- A set of heterogeneous, uncontrollable, and possibly unknown, resources. This typically comprises uncontrollable loads, as well as distributed generation, *e.g.*, photovoltaic generators.

- A grid-connected controllable electric energy storage system.

- A large commercial building already equipped with a Building Energy Management System (BEMS).
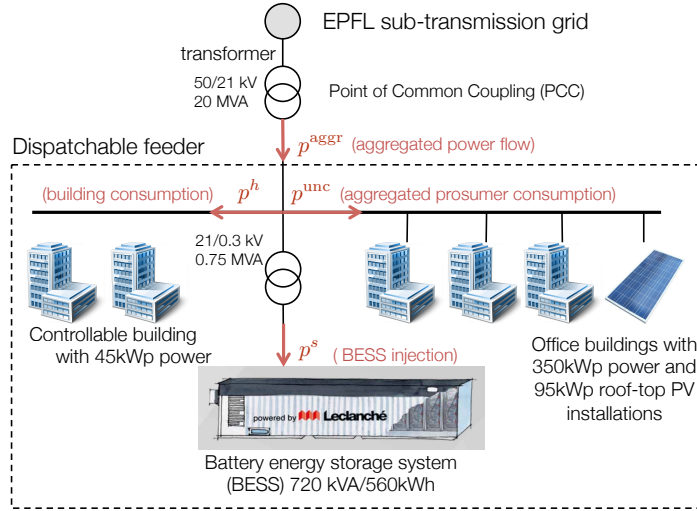
Figure 7.1 – The experimental setup of Section 7.7 that is used to validate the control architecture. It comprises a medium voltage feeder with a set of heterogeneous uncontrollable buildings, roof-top PV installations, an electrical storage, and a controllable office building. The control framework requires minimal invasive monitoring infrastructure, namely the aggregated power flow at the PCC, the battery power injection, and the buildings consumption (both controllable and uncontrollable).

Referring to Figure 7.1, in the following we use the following notation: p^{aggr} is the composite power consumption at the Point of Common Coupling (PCC), p^s is the bidirectional real power flow of the BESS, p^h is the power consumption of the controllable building and, p^{unc} is the aggregated power consumption of the uncontrollable resources, also referred as *prosumers*. For all resources, we assume a passive sign notation, *i.e.*, positive power values correspond to consumption, whereas negative values correspond to power injection towards the PCC.

Assumption 1. we assume the local distribution feeder to have adequate ample capacity and characteristics so as to operate within its voltage limits and technical constraints.

Similarly to the conventional way of planning a power system's operation [70], the dispatchability framework is comprised of two phases referred to as *day-ahead operation* and *real-time operation*.

- **Day-ahead operation**: Each day at a pre-specified time, an aggregated power profile for all elements connected at the PCC is computed for the following day, starting at 00.00. The power profile, called the *dispatch plan*, is a consumption profile with a T_{disp}-resolution that the feeder commits to track during real-time operation.

  The day-ahead phase is illustrated in Section 7.4.

- **Real-time operation**: Starting at 00:00 the next day of operation begins and the *dispatch plan* comes into effect. To correct for any forecasting error the dispatch feeder operator can

modulate both the power injection of the BESS, as well as the consumption of the flexible loads.

The real-time phase is described in Section 7.5.

7.3 Problem setup

This section presents the models and the constraints characterizing each resource.

7.3.1 Building model and constraints

As done in the previous chapters, we consider a building that has an HVAC system that is fully or partially controllable. With the usual notation, the set of feasible input trajectories for the building is defined as:

$$
\mathcal{U}^h(\hat{x}, \hat{\mathbf{d}}) = \left\{ \mathbf{u}^h \left|
\begin{array}{l}
x_{i+1} = Ax_i + Bu_i + Ed_i \\
y_{i+1} = Cx_i \\
p_i^h = h_{\text{hvac}}(u_i, d_i) \\
|y_i - T_{\text{ref}}| \le \beta_i \\
p_{\text{min}}^h \le p_i^h \le p_{\text{max}}^h \\
u_i^h \in \mathbb{U} \\
x_0 = \hat{x} \\
\forall i = 0, \dots, N-1,
\end{array}
\right. \right\} \tag{7.1}
$$

where the map $h_{\text{hvac}}(\cdot)$ captures the relation between the HVAC command and the external disturbance to the overall power consumption of the HVAC system. Thus, the set $\mathcal{U}^h$ characterizes the set of all power consumption trajectories for the building over the horizon N.

7.3.2 Battery model and constraints

We consider a grid-connected BESS, whose power injection is controlled by means of its Battery Management System (BMS). The BESS is capable of injecting or absorbing power values equal to the set-points demanded by the BMS, subject to operational constraints on the power values themselves and on the BESS SoC. The dynamics of the BESS SoC and DC voltage can be described by the discrete-time, generally non-linear, maps $k(\cdot)$ and $l(\cdot)$.

Variable	Definition
$\mathbf{D}^{\text{unc}}$	Dispatch plan for the uncontrollable prosumer
$\mathbf{D}^h$	Dispatch plan for the controllable building
$\mathbf{D}^s$	Dispatch plan for the BESS
$\mathbf{D}^{\text{con}}$	Dispatch plan for the controllable elements, namely the BESS and the building
$\mathbf{D}^{\text{aggr}}$	Aggregated dispatch plan at the Point of Common Coupling (PCC)
$\mathbf{p}^{\text{unc}}$	Real-time power realization of the uncontrollable elements
$\mathbf{p}^{\text{h}}$	Real-time power consumption of the controllable building
$\mathbf{p}^s$	Real-time power injection for the BESS
$\mathbf{p}^{\text{aggr}}$	Real-time aggregated power at the PCC
$\mathbf{w}$	Aggregated deviation from aggregated dispatch plan
$\boldsymbol{\epsilon}^{\text{unc}}$	Deviation from local dispatch for the uncontrollable elements
u_i	Actuation signal to the heating, ventilation, and air conditioning (HVAC) system
x_i	Current state of the building
d_i	Vector of disturbances acting on the building
y_i	Current zone indoor temperatures
$h_{\text{hvac}}(\cdot)$	Static map describing the power consumption of the HVAC system
SoC_i	Current state of charge at the BESS
v_i	Current DC voltage at the BESS
$k(\cdot)$	SoC transition map
$f(\cdot)$	DC voltage transition map

Parameter	Definition
T_{disp}	Resolution for the aggregated dispatch plan
N	Prediction horizon
N_{scen}	Number of scenarios in the two-stage stochastic approximation
c_{ener}	Vector of either user-defined or real energy prices
T_{ref}	Reference temperature for the building zones
β	Allowed temperature variation for the building zones
$v_{\text{max}},\ v_{\text{min}}$	Upper and lower bound for the voltage on the BESS DC
$p_{\text{max}}^h,\ p_{\text{min}}^h$	Upper and lower bound for the HVAC electric consumption
$p_{\text{max}}^s,\ p_{\text{min}}^s$	Upper and lower bound for the BESS electric injection
$\text{SoC}_{\text{max}},\ \text{SoC}_{\text{min}}$	Upper and lower bound for the SoC of the BESS
$\mathbb{U}_k$	Set describing the limits for the actuation commands to the HVAC system
l_{h}	Local cost function for the building
l_{BESS}	Local cost function for the BESS

Table 7.1 – Nomenclature.

The set of all feasible BESS power trajectories over the dispatching horizon N is:

$$
\mathcal{P}^s(\hat{\mathrm{SoC}}) = \left\{ \mathbf{p}^s \left|
\begin{array}{l}
\mathrm{SoC}_{i+1} = k(\mathrm{SoC}_i, p_i^s) \\
\mathrm{SoC}_{\min} \leq \mathrm{SoC}_i \leq \mathrm{SoC}_{\max} \\
v_{i+1} = f(v_i, p_i^s) \\
v_{\min} \leq v_i \leq v_{\max} \\
p_{\min}^s \leq p_i^s \leq p_{\max}^s \\
\mathrm{SoC}_0 = \hat{\mathrm{SoC}} \\
\forall i = 0, \ldots, N-1,
\end{array}
\right. \right\}
\tag{7.2}
$$

where $v_{\min}$ and $v_{\max}$ impose limitations on the voltage on the BESS DC bus. $p_{\min}^s$ and $p_{\max}^s$ limit the absolute value of the BESS active injection. The set $\mathcal{P}^s$ depends for the BESS only on the initial value of its state of charge, SoC_0.

The maps $k(\cdot)$ and $f(\cdot)$, modeling respectively the SoC and DC voltage as a function of the injected power can be of increasing complexity. Notably, in the day-ahead problem (Section 7.4) and in the real-time high-level controller (Section 7.5.1) the BESS faster dynamics are neglected, since the problem is solved with low time resolution. For such problems the voltage model and constraints can be neglected and the SoC model can be approximated by a linear model, neglecting the BESS charging efficiency and thermal effects. The low level controller of the BESS (Section 7.5.2) operates instead at faster sampling time. In this case the BESS faster voltage dynamics need to be taken into account and the $l(\cdot)$ function is therefore a non-linear model based on a battery equivalent electric circuit and identified through grey-box modelling techniques [127].

7.3.3 Uncontrollable prosumers

To obtain a power profile forecast we consider a forecasting that determines point-wise in time prosumption. It is worth noting that the proposed control methodology is independent of the adopted forecasting method, in other words the designer can use predictions from a forecasting tool of his own choice. The adopted forecasting method is not a contribution of this thesis. However, for the sake of completeness, we describe it in the following.

In general, the prosumption mix in low voltage distribution systems for residential and commercial areas normally consists of household and office demand (*e.g.* lightning, appliances, elevators, pumps) and, possibly, PV generation from roof-top installations. Assuming we observe the prosumption value in an aggregated fashion (*e.g.* from power flow measurements at the secondary substation level) without the knowledge of the power injections of the single units, it is convenient to decouple demand and generation and forecast them independently since they depend on different physical phenomena. The forecasting procedure we adopt for the experiments proposed in this chapter is:

1. historical prosumption measurements are disaggregated into the demand and PV generation

profiles according to the method described in [127]. The outputs of this process are the disaggregated profiles, as well as the values of the PV nominal capacity together with the tilt and azimuth configurations of the plants.

2. historical demand profiles are used in a vector auto-regression model to produce a sequence of point predictions of the demand for the next day, as described in [127];

3. Global horizontal irradiance predictions from a numerical weather predictions provider are coupled with a transposition model (according to the tilt/azimuth orientation identified in step 1) and a PV model to determine the forecast of the PV generation;

4. the forecast profiles of demand and PV generation are summed to obtain a sequence of point predictions of the prosumption at the PCC.

To model the related uncertainty, the adopted forecasting tool produces a number N_{scen} of scenarios. Let $\mathbf{L}^{(k)}$ be the sequence of prosumption point predictions for scenario k and $\mathbf{D}^{\text{unc}}$ the expected prosumption calculated as the component-wise average for all scenarios. A set of N_{scen} possible realizations for the forecast error, ϵ^{unc}, is thus created as:

$$\boldsymbol{\epsilon}^{\text{unc},(k)} = \mathbf{L}^{(k)} - \mathbf{D}^{\text{unc}} \quad k = 1, \dots N_{\text{scen}} \tag{7.3}$$

7.4 Day-Ahead Problem

As already explained in Section 7.2, the dispatchability framework is comprised of two different phases. The first one, referred to as *day-ahead operation*, requires determining at the PCC the so-called *dispatch plan*. The dispatch plan is an aggregated power profile, with a time resolution of T_{disp} (typical values could range from 5min to 15min), that the feeder commits to track during the operation of the next day.

In the following, the composite dispatch plan at the PCC, for time-interval k, is denoted by D_i^{aggr} whereas the sequence over the entire day of operation is denoted by $\mathbf{D}^{\text{aggr}}$. In the considered scenario (Figure 7.1), the dispatch plan is the sum of three distinct contributions:

$$\mathbf{D}^{\text{aggr}} := \mathbf{D}^h + \mathbf{D}^s + \mathbf{D}^{\text{unc}}$$

where $\mathbf{D}^h$ is the contribution of the flexible demand, $\mathbf{D}^s$ the contribution of the BESS element and, $\mathbf{D}^{\text{unc}}$ the portion due to the uncontrollable prosumer.

For later convenience, we also define the aggregated power profile for the controllable elements in the feeder portfolio, *i.e.* $\mathbf{D}^{\text{con}} := \mathbf{D}^h + \mathbf{D}^s$. Therefore, $\mathbf{D}^{\text{aggr}}$ can be simplified as:

$$\mathbf{D}^{\text{aggr}} = \mathbf{D}^{\text{con}} + \mathbf{D}^{\text{unc}} \tag{7.4}$$

In a similar way, it is possible to define the aggregated power consumption during real-time operation as the sum of all individual contributions:

$$\mathbf{p}^{\text{aggr}} := \mathbf{p}^h + \mathbf{p}^s + \mathbf{p}^{\text{unc}} \tag{7.5}$$

Finally, we introduce the *tracking error* as the difference between the committed composite dispatch plan (7.4) and the real-time realization (7.5):

$$\begin{aligned} \mathbf{w} :=& \mathbf{p}^{\text{aggr}} - \mathbf{D}^{\text{aggr}} \\ =& \mathbf{p}^h + \mathbf{p}^s - \mathbf{D}^{\text{con}} + \boldsymbol{\epsilon}^{\text{unc}} \end{aligned} \tag{7.6}$$

where the term $\boldsymbol{\epsilon}^{\text{unc}}$ is the dispatch forecast error for the uncontrollable load.

7.4.1 Day-ahead problem

At the beginning of the next day of operation, the dispatch plan comes into effect and it should be tracked accurately. To this aim, during real-time operation, the feeder can decide to operate the two controllable resources, *i.e.* the building and the BESS, in order to compensate for the fluctuations of the uncontrollable prosumer.

It is then clear that an adequate dispatch plan should be computed so as to: 1) achieve the best possible performance with respect to the considered application (minimization of the cost of purchased energy, reduction of tracking penalties, peak reduction, etc.) 2) guarantee enough reserves to absorb real-time deviations.

These requirements can be formally expressed by the following adaptation of Problem 2 given in Section 4.5:

Problem 9 (Day-ahead commitment problem)**.**

$$\underset{\substack{\mathbf{D}^{\text{con}} \\ \pi_{\text{h}}, \, \pi_s}}{\text{minimize}} \quad C(\mathbf{D}^{\text{aggr}}, \mathbf{p}^h, \mathbf{p}^s, \hat{\mathbf{d}}, \boldsymbol{\epsilon}^{\text{unc}}) \tag{7.7}$$

$$\text{s.t.} \tag{7.8}$$

$$\text{(Building Constraints)} \quad \mathbf{u} \in \mathcal{U}^h(\hat{x}, \hat{\mathbf{d}}) \tag{7.9}$$

$$\text{(Building Power)} \quad \mathbf{p}^h = h_{\text{hvac}}(\mathbf{u}, \hat{\mathbf{d}}) \tag{7.10}$$

$$\text{(BESS Constraints)} \quad \mathbf{p}^s \in \mathcal{P}^s(\hat{\text{SoC}}) \tag{7.11}$$

$$\text{(Control policies)} \quad \mathbf{p}^s = \pi_s(\boldsymbol{\epsilon}^{\text{unc}}), \quad \mathbf{p}^h = \pi_h(\boldsymbol{\epsilon}^{\text{unc}}) \tag{7.12}$$

$$\text{(Tracking Error)} \quad \mathbf{w} = \mathbf{p}^h + \mathbf{p}^s - \mathbf{D}^{\text{con}} + \boldsymbol{\epsilon}^{\text{unc}} \tag{7.13}$$

$$\text{(Uncertainty)} \quad \boldsymbol{\epsilon}^{\text{unc}} \sim \mathcal{F} \tag{7.14}$$

The cost function of equation (7.7), in its most general form, can be stated as:

$$C(\mathbf{D}^{\mathrm{aggr}}, \mathbf{p}^h, \mathbf{p}^s, \mathbf{d}, \boldsymbol{\epsilon}^{\mathrm{unc}}) := \mathbb{E}_{\boldsymbol{\epsilon}^{\mathrm{unc}}} \left(\overbrace{l_{\mathrm{h}}(\mathbf{p}^h, \mathbf{d})}^{\text{local cost}} + \overbrace{l_s(\mathbf{p}^s, \mathbf{d})}^{\text{local cost}} + \overbrace{c_{\mathrm{ener}}^T \mathbf{D}^{\mathrm{aggr}} + \alpha \|\mathbf{w}\|_2^2}^{\text{shared objective}} \right) \qquad (7.15)$$

where the terms $l_h(\cdot)$ and $l_s(\cdot)$ represent the local cost functions for the building and storage element, respectively. They capture the individual objectives that these two resources are pursuing. To provide possible examples (as considered later in the simulation and experimental analysis) of these individual costs, the building owner could aim at maximizing the comfort for occupants, *i.e.*, l_h can be set to $l_h := \|\mathbf{y} - T_{\mathrm{ref}}\|^2$. Regarding the BESS, one can aim to minimize the battery usage with respect to a pre-scheduled reference SoC, *i.e.*, $l_s := \|\mathbf{SoC} - \mathbf{SoC}_{\mathrm{ref}}\|^2$.

The last term of the cost is instead represented by the shared financial cost which comprises (in order of appearance) the energy cost for the aggregated power consumption at the PCC, and a tunable penalization term (through the weight α) for the tracking error. Due to the dependence on the uncertain parameter, $\boldsymbol{\epsilon}^{\mathrm{unc}}$, the cost function is evaluated in expectation.

Equations (7.9-7.11) enforce the satisfaction of all operational and comfort constraints for the building and the BESS. Finally, we assume the term, $\boldsymbol{\epsilon}^{\mathrm{unc}}$, to follow the probability distribution, $\mathcal{F}$ which is, in general, not normal.

Thus, with respect to Problem 2, we have two controllable resources, *i.e.* $N_{\mathrm{cont}} = 2$, and the tracking requirement is imposed with respect to the term, $\mathbf{w}$. Also, according to the structure proposed in [127], we do not consider the possibility of adjusting the dispatch plan during real-time operation using an intraday policy. Finally, the grid constraints are not considered thanks to Assumption 1 which is actually verified in the experimental setup.

7.4.2 Approximate solution

To obtain a tractable reformulation, we resort to the two-stage stochastic approximation scheme of Section 4.6.3. As done in the previous chapter, this is mainly done to obtain less conservative solutions. However, it should be noticed that considering the two-stage method allows one to accommodate also the general non-linear resource models of Sections 7.3.1 and 7.3.2.

Thus, for the considered problem, the first stage variable is the dispatch plan of the controllable resources, $\mathbf{D}^{\mathrm{con}}$, whereas the second stage decisions are the building, $\mathbf{u}^{(k)}$, and battery, $\mathbf{p}^{s,(k)}$, trajectories corresponding to each scenario, $\boldsymbol{\epsilon}^{\mathrm{unc},(k)}$. Also, the expectation appearing in the cost function (7.15) is replaced by its empirical expectation according to the well-known sample average approximation [125]. As a reminder, in our case each scenario is generated according to the procedure in 7.3.3. We highlight once more that any forecasting tool can be considered at this

stage.

The reformulated problem reads:

Problem 10 (Two-stage approximation)**.**

$$
\underset{\substack{\mathbf{D}^{\mathrm{con}} \\ \mathbf{p}^{s,(k)},\, \mathbf{u}^{(k)}}}{\text{minimize}} \quad c_{\mathrm{energ}}^T \mathbf{D}^{\mathrm{aggr}} + \frac{1}{N_{\mathrm{scen}}} \sum_{k=1}^{N_{\mathrm{scen}}} \{l_h(\mathbf{p}^{h,(k)}, \hat{\mathbf{d}}) + l_s(\mathbf{p}^{s,(k)}, \hat{\mathbf{d}}) + \alpha \|\mathbf{w}^{(k)}\|^2\}
$$

$$\text{s.t.} \tag{7.16}$$

$$\text{(Building Constraints)} \quad \mathbf{u}^{(k)} \in \mathcal{U}^h(\hat{x}, \hat{\mathbf{d}}) \tag{7.17}$$

$$\text{(Building Power)} \quad \mathbf{p}^{h,(k)} = h_{\mathrm{hvac}}(\mathbf{u}^{(k)}, \hat{\mathbf{d}}) \tag{7.18}$$

$$\text{(BESS Constraints)} \quad \mathbf{p}^{s,(k)} \in \mathcal{P}^s(\hat{\mathrm{SoC}}) \tag{7.19}$$

$$\text{(Tracking Error)} \quad \mathbf{w}^{(k)} = \mathbf{p}^{h,(k)} + \mathbf{p}^{s,(k)} - \mathbf{D}^{\mathrm{con}} + \boldsymbol{\epsilon}^{\mathrm{unc},(k)} \tag{7.20}$$

$$\text{(Scenarios)} \quad k = 1, \ldots N_{\mathrm{scen}} \tag{7.21}$$

$$\tag{7.22}$$

7.5   Real-time operation

At the beginning of a new day's operation, the aggregated dispatch plan comes into effect. The task of the feeder operator is thus to compensate for any real-time mismatch in order to track the committed profile precisely.

Intuitively, the real-time controller should take care of both slow and fast dynamics of the system. On one side, the tracking requirement for each sampling time, T_{disp}, calls for a fast modulation ($\sim$ 10-30 seconds) of the power consumption of two resources. On the other side, to maintain the comfort inside the building, long-term prediction and weather forecast should be taken into account. Clearly, managing these two separated time-scales simultaneously is, in general, computationally intractable. Moreover, due to the physical limitations of the equipment, for most commercial buildings, it is not possible to control the power consumption of the HVAC system at a very fast pace.

Similarly to the previous chapter, to deal with both slow and high frequency requirements, we consider a *hierarchical multi-scale* controller, sketched in Figure 7.2, that aims at exploiting the synergy between the two controllable elements. On one side, the BESS represents the master element which ultimately delivers the dispatchable service by correcting forecasting errors at a sub-minute time-scale; on the other side, the smart building operates at a slower time scale and its main goal is to try to restore the SoC of the BESS while respecting its operational constraints. The overall tracking problem is solved at two well-separated timescales: a Low-Level MPC controller operating at a fast sampling time continuously computes the power injection of the BESS so as to
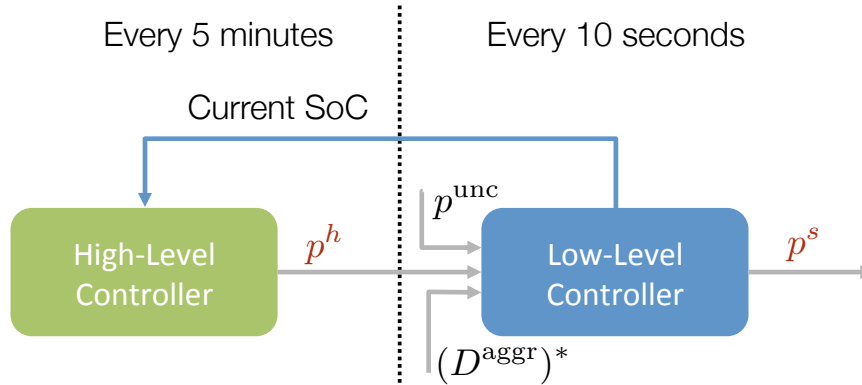
Figure 7.2 – Time separation for the overall real-time controller. Every slow time interval, the current SoC is transmitted by the BESS operator to the high-level controller which determines its next action in order to maneuver the SoC. On a faster time resolution (*i.e.* 10 seconds), the low-level controller measures the power realizations of both the controllable element, p^h, and the uncontrollable prosumer, p^{unc} and computes the BESS power injection for the following fast time interval, to track the committed dispatch plan, $(\mathbf{D}^{\mathrm{aggr}})^*$. Legend: Grey lines correspond to measured quantities, blue lines to transmitted information and, red terms represent decision variables at each controller level.

track the committed dispatch plan; at a slower resolution, a High-Level MPC problem is solved to determine the power modulation for the building so as to restore the SoC of the battery.

7.5.1 High-Level Controller

In this section, we describe the high-level control problem solved with a T_{disp}-minute time sampling over the prediction horizon, N (*e.g.* 12 hours). The objectives of this controller are, as in the day-ahead problem: 1) a high-level of comfort for the building occupants; 2) the satisfaction of operational constraints of the HVAC; 3) modulate the power consumption of the building so as to maneuver the SoC of the BESS.

To accomplish this, an MPC problem is solved at each iteration. The main steps for the MPC algorithm are the following:

1. Retrieve the most recent forecast for all perturbations over the considered prediction horizon, N.

2. Form the vector of disturbances, $\hat{\mathbf{d}}$, over the prediction horizon.

3. Reconstruct the current state of the building, $\hat{x}_i$ by means of a standard Kalman filter.

4. Retrieve the current state of the battery, $\hat{\mathrm{SoC}}_i$.

5. Solve the following optimization problem:

Problem 11 (Real-time operation: high-level controller).

$$\underset{\mathbf{u}^{(k)}}{\text{minimize}} \quad \frac{1}{N_{\text{scen}}} \sum_{k=1}^{N_{\text{scen}}} \{l_h(\mathbf{p}^{h,(k)}, \mathbf{d}) + l_s(\mathbf{p}^{s,(k)}, \mathbf{d}) + \alpha||\mathbf{w}^{(k)}||^2\} \tag{7.23}$$

$$\text{s.t.} \quad \mathbf{u}^{(k)} \in \mathcal{U}^h(\hat{x}_i, \hat{\mathbf{d}}) \tag{7.24}$$

$$\mathbf{p}^{h,(k)} = h_{\text{hvac}}(\mathbf{u}^{(k)}, \hat{\mathbf{d}}) \tag{7.25}$$

$$\mathbf{p}^{s,(k)} \in \mathcal{P}^s(\hat{\text{SoC}}_i) \tag{7.26}$$

$$\mathbf{w}^{(k)} = \mathbf{p}^{h,(k)} + \mathbf{p}^{s,(k)} - (\mathbf{D}^{\text{con}})^* + \boldsymbol{\epsilon}^{\text{unc},(k)} \tag{7.27}$$

$$p_0^{h,(1)} = \ldots = p_0^{h,(N_{\text{scen}})} \tag{7.28}$$

$$k = 1, \ldots, N_{\text{scen}} \tag{7.29}$$

6. Retrieve the optimal solution and apply the first power input for the building, *i.e.*, p_0^h.

Remark 15. *Weather forecast can be typically obtained through available web-services by simply specifying the geographic location of interest. In the experimental demonstrator of Section 7.7, the weather station was selected from 4 available stations in a 3 Km radius, based on the historical quality of the data it provides. (source MeteoSwiss, meteostation Lausanne freiland , GPS coordinates 6°38.56′ 46°33.33′).*

In equation (7.27), the constraint which contains the essence of the whole hierarchical controller is formulated. Essentially, this equality constraint states that the tracking error at the PCC can be controlled, by acting on the power consumption of the building. Since, the Low-Level controller will compensate for any residual tracking error at a higher-frequency, the capability of maneuvering the term **w** can be exploited to have a net effect on the power injection of the BESS. Finally, in order to apply the control input to the actual system, equation 7.28 guarantees the uniqueness of the first power input of the building for all considered scenarios.

In summary, depending on the chosen cost function, and the relative weights between each term, the High-Level controller guarantees a quality of comfort inside the building while, determining the average power injection of the BESS for the next T_{disp}-minutes.

Please note that in Problem 11 the minimization is not taken over $\mathbf{D}^{\text{con}}$ anymore since this has already been fixed during day-ahead operation. For this reason, the aggregated dispatch plan enters the optimization problem as a parameter.

Remark 16. *In the High-Level controller formulation, the simple linear reservoir model of Section 7.3.2 is considered for the BESS. Clearly the model does not capture complex dynamics of the battery such as charging/discharging losses, thermal effects, etc. that can be incorporated using a non-linear first-principle-based model [98]. Nevertheless, due to the hierarchical structure of the controller, the*

higher layer does not necessitate an exhaustive description of the BESS internal states which are instead considered at the lower layer. On the contrary, the simple linear model is exploited by the MPC to obtain a coarse prediction of the future SoC as a function of the power of the injected power according to (7.26) and (7.27).

7.5.2 Low-level Controller

After the building re-dispatch action performed by the high-level controller (which acts at T_{disp}-minute resolution), it is still necessary to compensate for the mismatch between the real-time realization and the aggregated dispatch plan. This is accomplished by the Low-Level controller, which determines the active power set-point of the BESS four quadrant power converter. Unlike the power consumption of building space heating systems which cannot be modified at a fast pace, the power injection of BESS can be typically controlled at a high-frequency and it is, therefore, suitable to perform fine power/energy adjustments. The Low-Level controller is the MPC algorithm described in [127] and is not a contribution of this work. In summary, it consists of solving an optimization problem at a 10 s resolution on a T_{disp} shrinking horizon, from the current time period until the end of the dispatch interval. The cost function is given by minimizing the energy mismatch between real-time realization (progressively known from real-time measurements) respect to the dispatch plan while being subject to BESS operation constraints on DC voltage, DC current and SoC limits. At this stage, the BESS is operated at unitary power factor, in other words there is no reactive power injection.

7.6 Simulations

In this section, we validate the proposed algorithm by means of an extensive simulation study.

7.6.1 Simulation Setup

The simulations are performed for the winter season for which the external disturbance sequence (internal gains, solar radiation, outside temperature) is assumed to be perfectly known. For the building model, we consider an ASHRAE standard five zone office building model taken from the reference database of the U.S. department of Energy [136]. We use the MATLAB toolbox OpenBuild [46] (freely available at [49]) to automatically extract a linear thermal dynamic model suitable for control purposes. A forced-air HVAC system served by an electric heat-pump is separately designed as in [101]. As for the BESS, we consider a simple linear model as in [40], which has been identified and validated on real measurements for the battery storage system described in Section 7.7.

The uncertain disturbance for the uncontrollable resources, ϵ^{unc}, is unknown at the time of day-ahead operation. To solve the day-ahead problem, we obtain a sequence of scenarios from historical realizations (2015-2016) of the real setup considered in Section 7.7. More precisely, each scenario is computed as $\epsilon^{\text{unc},(k)} = \mathbf{p}^{\text{unc},(k)} - \mathbf{D}^{\text{unc}}$ where each $\mathbf{D}^{\text{unc}}$ and $\mathbf{p}^{\text{unc},(k)}$ are realizations of the predicted power and its actual real-time counterpart for the uncontrollable resources of Section 7.7 for a particular day in 2015-2016. The distribution of the scenarios can be observed in Figure 7.3.
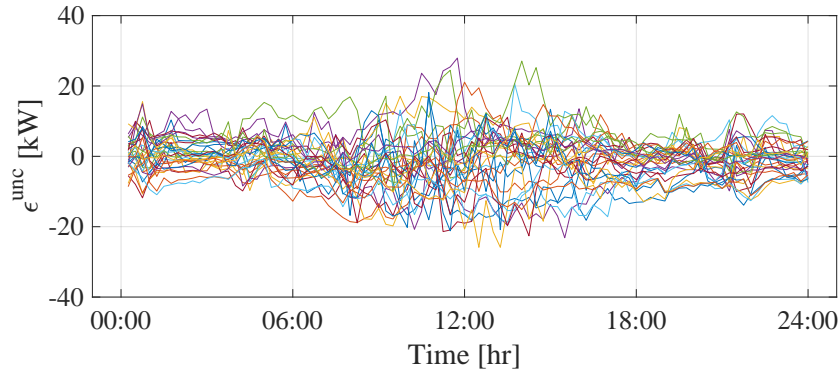
Figure 7.3 – 30 realizations of prosumer tracking errors, computed as $\boldsymbol{\epsilon}^{\text{unc},(k)} = \mathbf{p}^{\text{unc},(k)} - \mathbf{D}^{\text{unc}}$, for the experimental setup of Section 7.7. Peak prosumer power: 350kW. Peak deviation from forecast: 28 kW.

The full validation of the overall closed-loop controller, *i.e.*, dispatch computation and real-time control is provided in experimental Section 7.7. Hence, in this section, we focus on the planning phase, *i.e.*, the open-loop solution of the dispatch problem in order to highlight the main characteristics of the proposed algorithm.

The very general form of the cost function provided in (7.15) accounts for a wide range of different configurations. For the sake of space, we focus on the case of *minimization of the BESS usage*. The reasons behind this choice are: 1) it is easy to precisely quantify the performance of the controller by simply comparing the utilized SoC, 2) in general, for a given peak prosumer power, it is desirable to achieve dispatchability with the smallest possible battery, as BESS are expensive devices, 3) other configurations such as, *e.g.*, cost-effectiveness, depend heavily on local energy prices and government subsidies making the results very case-specific, and 4) this case already displays how the algorithm successfully exploits the synergy between the BESS and the building to absorb the forecast errors.

To assess the performance of the algorithm, two different simulation configurations are considered.

Configuration I [BESS]: In the first configuration, no flexible demand is present, *i.e.*, the BESS represents the only degree of freedom during real-time operation. Therefore, as in [127], the BESS compensates for all tracking errors so that the success of the scheme heavily depends on the accuracy of the prosumer forecast. In Figure 7.4, the requirements on the BESS in terms of both SoC and power are displayed for each considered scenario.

Configuration II [BESS + building]: The controller of Problem 10 is considered. Therefore, in the first stage, a composite dispatch plan, $\mathbf{D}^{\text{con}} = \mathbf{D}^h + \mathbf{D}^s$, is computed for the flexible demand and the BESS. In the second stage, the power realizations for the two resources are also determined
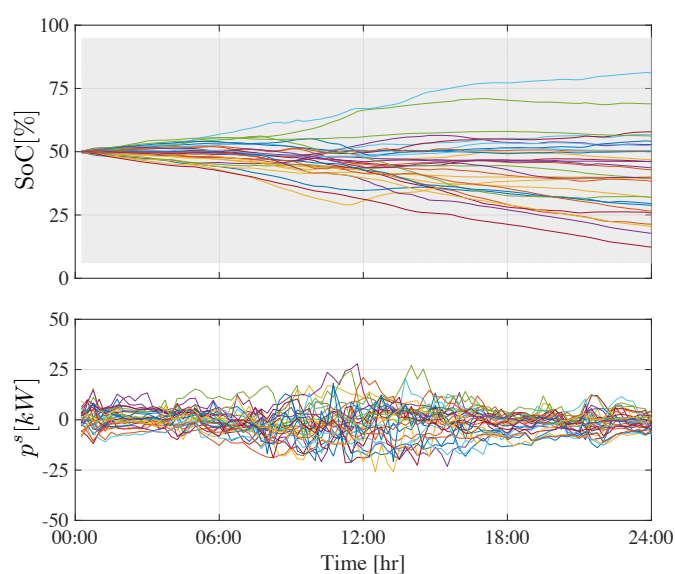
Figure 7.4 – **Configuration I.** Battery requirements for Configuration I for all considered scenarios. **Results**: Average BESS usage: 81.64 [kWh], Peak BESS usage = 213.74 [kWh]. *Upper*: SoC trajectories computed by directly injecting the prosumer power deviation into the battery model. Each color corresponds to a different scenario. The gray area represents the physical limit of the BESS. *Lower*: Power injection trajectories. As in this case the BESS is the only controllable element, the power injections simply correspond to the original prosumer deviations.

in order to absorb the uncertainty of the uncontrollable prosumer. Results are shown in Figure 7.5. In particular, Figure 7.5 displays how the algorithm successfully exploits the thermal inertia of the building to absorb the power mismatch introduced by the uncontrollable resources. This is achieved while still ensuring the satisfaction of the comfort constraint as no violations are experienced. In the upper plot of Figure 7.5a the battery requirements for the controller are also reported. Clearly, the presence of the building is crucial in reducing the required size of the battery as a 60% reduction with respect to Configuration I is obtained.
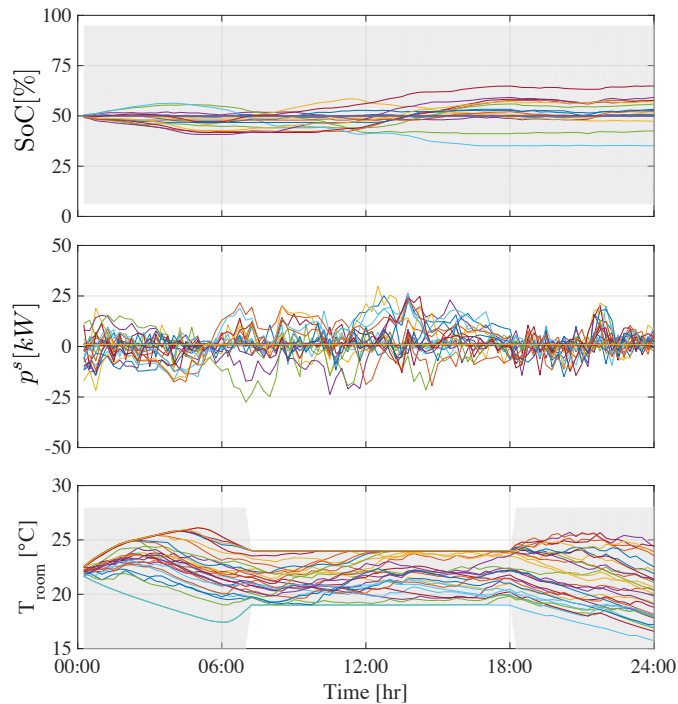
Figure 7.5 – Open-loop simulations for the BESS + building case. *Upper*: The battery SoC trajectories are reported. Each color corresponds to a different realization. The gray area displays the physical constraints of the battery. *Middle*: The power injection trajectories for each scenario. *Lower*: The indoor temperature realizations for the building. For the sake of visualization, instead of the individual zone temperature, the average zone temperature is reported for each scenario. As in the previous subplots, each color corresponds to a different sample of the uncertain deviation, ϵ^{unc}. Finally, the gray area depicts the time-varying comfort constraints.

7.7 Experiments

An experimental validation of the proposed algorithm on the EPFL medium-voltage (MV) distribution grid is presented in this section. It consists in the same infrastructure considered in [127] with an additional controllable unit represented by the LADR platform.

Configuration	BESS	BESS + building
Maximum SoC used [kWh]	213.74	84.28
Average SoC used [kWh]	81.64	18.02

Table 7.2 – Statistics of the battery requirements obtained with the two considered configurations.

More precisely and as depicted in Figure 7.1, the considered distribution grid includes a battery electric energy system, a set of uncontrollable loads, renewable generation units, and a controllable office building. First, we provide more details regarding each element.

7.7.1 Battery electric storage system (BESS)

We consider a Leclanché grid-connected Lithium Titanate BESS characterized by a 720kVA/560 kWh power/capacity ratio. The battery consists of 9 parallel racks (each composed of 15 modules in series, where each module is composed of a 20s3p cell pack), a four quadrant fully-controllable DC/AC converter, and a 0.3/20kV step-up transformer. The whole system is placed in a temperature controlled environment.

7.7.2 Non-controllable units

The second key element, which also represents the main source of uncertainty, is composed of: 1) an aggregation of non-controllable buildings with a 350 kW peak consumption; 2) a 95 kWp roof-top PV installation. The composite power for both the uncontrollable units and the BESS is measured at the feeder by means of high-frequency high-precision phasor measurement unit [114].

7.7.3 Controllable Building

The controllable demand is represented by the LADR platform that was introduced in Chapter 5. Thus, as a reminder, the required thermal heat is provided by customized electric heaters that are placed in each room. As a consequence, for this case, the map h_{hvac} is given by the simple linear relation $p_i^h = u_i P_{\max}$, where the term u_i is the pulse-width ration of the heaters, and $P_{\max} = 1950$ Watts. However, with respect to Chapter 5, experiments are conducted for several consecutive days and, therefore, under changing weather conditions and in occupied conditions. To capture these inputs, a second identification campaign was performed in the winter of 2015 resulting in a model that takes three inputs: the electric power input to the heater, the outside temperature, and one input capturing the effect of solar radiation on the room. Validation tests have also been conducted showing an average fit, depending on the room, between 70 % and 90%. Steady state calculations on the identified models suggest that the equivalent average U-value of the envelope of the building is between 1 and 1.5W/m2/K. This value is not particularly high, meaning that the thermal performance is acceptable. It is also not very low, meaning that this is not a high efficiency building. Therefore, this building can be considered as representative of an average office building in Switzerland. We refer the reader to [48] for more details on the modeling of LADR.

Scaling As already mentioned, the controllable building accounts for a $\sim 8kW$ peak consumption which represents only 2% of the 350 kW peak of uncontrollable units. In the experimental results, the power consumption of the office building was virtually scaled up by a factor of six, leading to a maximum peak of 45 kW. The reasons for this are twofold: 1) to have a rated power comparable with the one of the uncontrollable resources and, therefore, being able to draw meaningful conclusions on the impact of a controllable load in the dispatchable feeder framework and, 2) to have the minimum size to prevent previously experienced failures. In particular, among the set of experiments with no controllable building [127], we selected the days where the dispatchable feeder operator failed to track the committed dispatch plan due to a complete charge/discharge of the BESS. A simulation-based analysis was then performed in which the size of the controllable building was slowly increased until the aggregated system manages to successfully track the dispatch plan.

We highlight that the virtual scaling is in this case strictly equivalent to consider additional rooms characterized by the same thermal characteristics and served by individual electric heaters as in our setup. As highlighted in our previous contributions [39, 48], the linear electric response of the heaters and weak thermal coupling between the rooms allow us to model each room separately using standard auto regressive model with exogenous inputs [79]. An examination of the identified models reveals very similar characteristics for the rooms both in terms of the time constants involved as well as the static gains for all external perturbations. Thus, the assumption of a linear relation between the scaling and the number of rooms is reasonable in this situation.

7.7.4 Results and discussion

A set of experiments has been conducted during the period from December 2016 to March 2017. To assess the performance of the proposed control architecture, we compare the obtained experimental results to the case where the BESS is the only controllable element during real-time operation (as in [127]). Referring to Figure 7.6, for each day of operation, the following two quantities are compared:

- The actual SoC as it is *measured* during the experiments (orange solid line),

- The *simulated* SoC for the battery in absence of the building (blue solid line). More precisely, as in [127], we re-perform in simulation the experiment in the case where the only degree of freedom during real-time operation is represented by the BESS. Hence, starting from its actual initial condition, the SoC is propagated through its dynamical model with the input being the dispatch error for the uncontrollable elements.

In the following, we present two days of operation which are representative of the results obtained during the experimental campaign. We refer the reader to Table 7.3 for an overview of the characteristics of the recorded forecasting error. **Day 1** is characterized by a smaller energy content of the forecasting error (86.73 kWh) together with a higher initial SoC which is, in this case, beneficial to counteract the positive bias of the error. On the contrary, **Day 2** represents a more critic situation

(a) Experimental results of **Day 1**
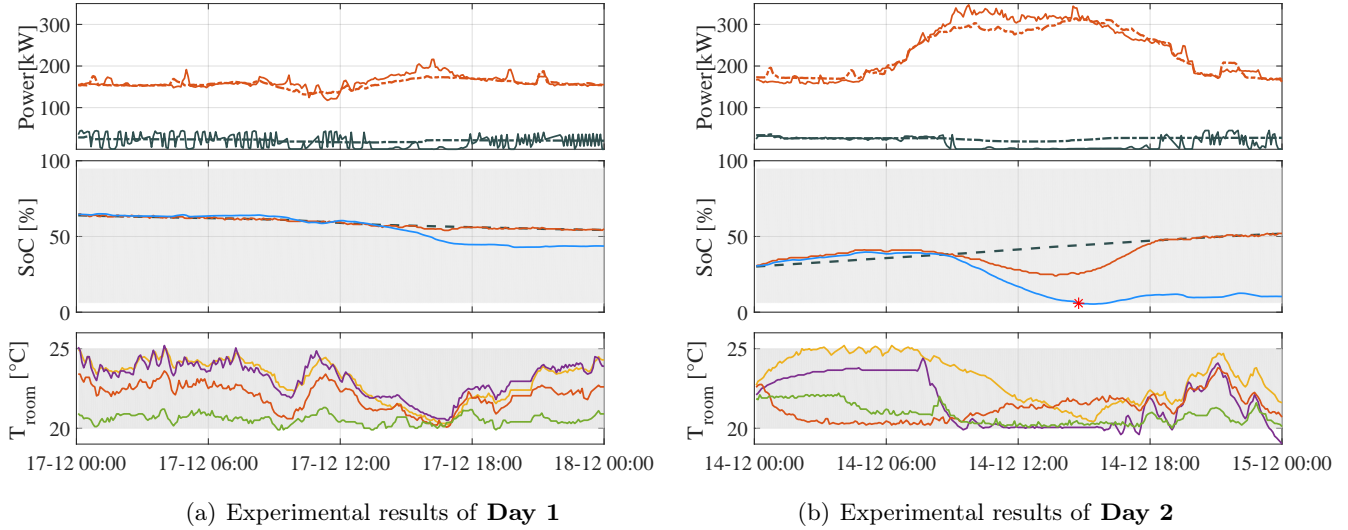
(b) Experimental results of **Day 2**

Figure 7.6 – Real-time operation for the dispatch tracking. *Upper*: The black dashed line represents the dispatch plan for the building, $\mathbf{D}^h$, whereas its actual realization, $\mathbf{p}^h$, is shown using a black solid line. Similarly, for the uncontrollable resources, the dashed orange line represents the day-ahead predicted power profile, $\mathbf{D}^{\text{unc}}$, and the solid line its measured value, $\mathbf{p}^{\text{unc}}$. *Middle*: The black dashed line represents the SoC reference. The experimental realization of the SoC is displayed in orange. The blue line is the simulated SoC in absence of the controllable building as previously explained. *Lower*: Temperature variation for the difference zones of the controllable building. Each color corresponds to the measured temperature in a zone. In both the middle and lower plots, the gray area represents the allowed ranges for the plotted quantities

due to a larger energy content of the forecasting error (262.31 kWh) and the initial SoC of the BESS which is already very close to the lower bound. As discussed in the following Section, in this case, the presence of the controllable load is strictly necessary to achieve dispatchability.

Figure 7.6(a) depicts the real-time operation for Day 1. Thanks to the more accurate prediction plan for the heterogeneous resources and the initial SoC, the BESS alone can easily compensate for the prediction error, *i.e.*, the difference between the prosumer dispatch plan, $\mathbf{D}^{\text{unc}}$, and its actual realization, $\mathbf{p}^{\text{unc}}$. Nevertheless, the battery requirements are drastically reduced when the controllable building is considered. This is particularly evident considering the central hours of the day, between 12:00 and 17:00. In fact, due to excessive prosumer consumption, the BESS would have experienced a partial discharge. On the contrary, the proposed algorithm prevents the discharge by exploiting the added flexibility represented by the thermal inertia of the building. Specifically, as shown in the top plot, the power consumption of the building is lowered for a limited period of time to counterbalance the negative forecasting error. This is done while still preserving a high-level of comfort as no temperature violations are observed (lower plot).

Experiment	Maximum [kW]	Minimum [kW]	Energy content [kWh]
Day 1	33.63	-16.72	86.73
Day 2	51.27	-33.88	262.31
Average over all experiments	44.96	-23.21	153.41

Table 7.3 – Statistics of the deviations for the selected experiments.

Quantity	BESS	BESS + building
Maximum SoC used [kWh]	262.31	149.71
Average SoC used [kWh]	153.41	74.52
Maximum BESS Power [kW]	82.51	65.23
Average Comfort Violation [°C/h]	n.a.	0.08
Maximum Comfort Violation [°C/h]	n.a.	0.97

Table 7.4 – Statistics over all 12 experiments conducted in the period between December 2016 and February 2017.

The second day of operation, displayed in Figure 7.6(b), represents a situation in which the presence of the deferrable load is crucial to achieve dispatchability. Once again, due to an unexpected higher prosumer consumption between roughly 10:00 and 14:00, the BESS would have experienced a rapid discharge until violating its lower SoC constraint. This event is displayed in the central plot of Figure 7.6(b) by a red cross around 15:00. Conversely, as in the previous case, the thermal inertia of the building is successfully exploited to decelerate the rate of discharge of the BESS and, therefore, prevent the failure. As it can be observed in the lowermost plot, in order to absorb the negative forecasting error, the controller almost entirely utilizes the flexibility of the building so that few modest comfort violations are experienced in this case.

As a summary, in both experiments of Figure 7.6, the presence of the controllable building has two positive effects: on one hand, it helps to reduce the required capacity for the BESS, and on the other hand it helps to track a predefined SoC reference. These results are achieved while still providing a high-level of comfort for occupants of the building. Finally, we report in Table 7.4 the statistics of all conducted experiments. As already observed in the simulation Section 7.6, by optimally exploiting the thermal inertia of the building, it is possible to drastically reduce the energy requirements on the BESS for a given size of uncontrollable prosumers.

Finally, to demonstrate the ability of the proposed control architecture to handle continuous operation, we report in Figure 7.7 the results obtained for three contiguous days in December 2016. Please note that according to the dispatchability framework, the dispatch plan for each resource (and thus the aggregated one) is recomputed every day one hour before delivery. The middle plot
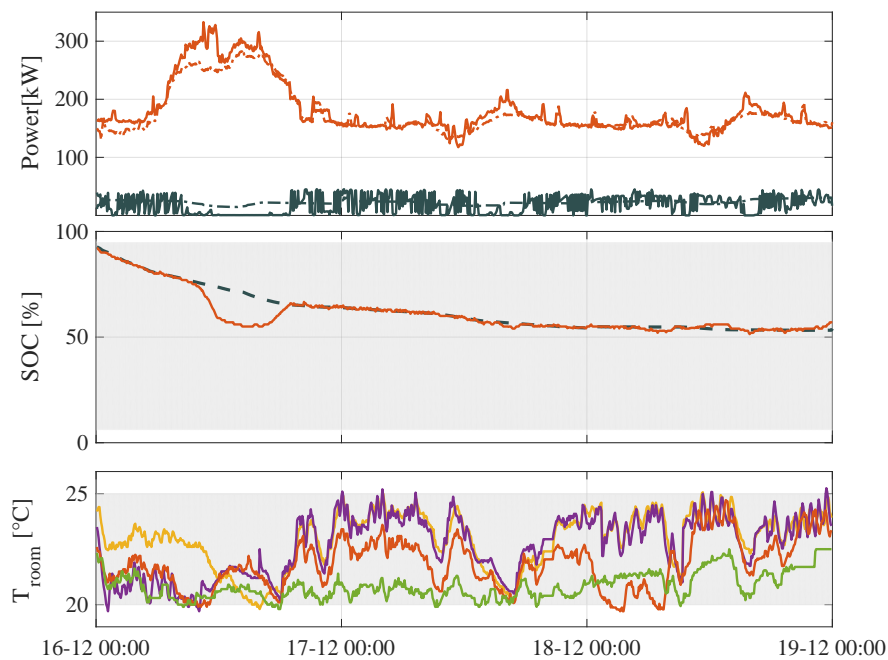
Figure 7.7 – **Continuous operation**: Three consecutive days of operation from 16th to 18th of December 2016. *Upper*: The black dashed line represents the dispatch plan for the building whereas its actual realization is shown using a black solid line. Similarly, for the uncontrollable resources, the dashed orange line represents the day-ahead predicted power profile and the solid line its measured value. *Middle*: The black dashed line represents the SoC reference. The experimental realization of the SoC is displayed in orange. The blue line is the simulated SoC in absence of the controllable building as previously explained. *Lower*: Temperature variation for the difference zones of the controllable building. Each color corresponds to the measured temperature in a zone. In both the middle and lower plots, the grey area represents the allowed ranges for the plotted quantities

of Figure 7.7 demonstrates how the real-time high-level controller successfully regulates the power consumption of the building so that the SoC accurately tracks the pre-scheduled reference SoC of the battery. This is particularly visible for the first of the three days when, due to an extreme forecasting error, the BESS is partially discharged and deviates from the reference trajectory (middle plot). During the same time of the day, the building lowers its power consumption so as to absorb part of the uncertainty and, therefore, restore the SoC. Once again, this is achieved without impacting the comfort in the building as temperature constraints are practically always respected.

7.8 Conclusion

This chapter considered the problem of ensuring that the aggregated power consumption of a cluster composed of a controllable building, a grid-connected electric storage, and a set of non-controllable heterogeneous prosumers follows a pre-defined profile, called a dispatch plan, which is established the day before operation.

To attain this aim, the control framework of Part II has been adapted in order to compute an aggregated dispatch plan that optimally splits the prosumer deviations among the two controllable resources. During real-time operation, a multi-scale hierarchical controller has also been designed that manages the delivery of the dispatchable service (by controlling the power injection of the electrical storage), while optimally exploiting the additional flexibility offered by the controllable load to facilitate continuous operation.

The effectiveness and generality of the proposed control architecture have been tested by means of extensive simulation results on a realistic setup. The practical relevance was also demonstrated by means of 12-full day experiments on a real-life real-scale medium voltage system of the EPFL campus, which comprises a number of uncontrollable office buildings (350kWp) equipped with 95 kWp roof-top generation, a utility-scale grid-connected electrical battery (720 kVA- 567 kWh) and, a controllable office building (45 kWp).

Both the simulations and experimental results showed that: 1) the proposed control scheme successfully achieves to allocate enough power reserves so as to absorb the real-time deviations of the prosumers while complying to local constraints at each flexible resource 2) the possibility to exploit the thermal inertia of a controllable building allows, when properly coordinated, to drastically reduce the energy requirements (57% in the experimental results) of the electric battery for the same size of uncontrollable prosumers.

Chapter 8

Conclusion

This chapter summarizes the main contributions of the thesis with respect to the research questions raised in Chapter 1.

Part I - MPC for distributed energy resources

This part answered the question of efficient control of resources that are subject to periodically time-varying stochastic disturbances.

A receding horizon control scheme that enforces recursive feasibility for the closed loop process of a periodic linear system subject to stochastic constraints was proposed in Chapter 3. The class of considered systems is wide and represents a powerful modeling tool for many real-life applications. This is true, in particular, since it allows the consideration of periodic inputs and state constraints as well as periodic disturbances that are characterized by time-varying probability distributions. The developed approach was applied, in simulation, on a building temperature control case showing its flexibility and effectiveness with respect to robust MPC control schemes.

Part II - Provision of grid services using distributed energy resources: Theory

This part answered the question of formally characterizing the amount of services that an aggregation of distributed energy resources can offer to the electric grid.

The general control problem, referred to as the *Planning Problem*, of distributed energy resources providing multiple local and shared services to the grid was first introduced. The problem determines the optimal allocation of the available controllable power so as to guarantee the fulfillment of the local services (peak shaving, congestion management, etc.), while maximizing the amount of fast regulating services (such as frequency control) that the resources can collectively provide. The novelty of the control framework lies in its capability of accounting for both the uncertainty related to the power injections at the network buses, as well as the uncertainty of the signal to be tracked during real-time operation. Two approximate solution methods for the solution of the planning

problem were proposed both relying on the convexification of the non-convex power flow equations around the current operating point. The first method relies on a two-stage approximation of the original problem which was solved using a scenario-based approach. The second method retains the original multi-stage structure of the problem while restricting the decision space to only affine control policies. In both cases, the reformulated problem is computationally tractable as a simple linear programming is to be solved.

An extensive simulation-based validation of the control methodology was undertaken on the IEEE 123 feeder, which was equipped with uncertain distributed generation, uncontrollable consumption, and three controllable electrical batteries. The simulation results showed that is indeed possible to formally characterize the amount of service that a set of resources can offer while respecting their operational constraints as well as the constraints of the hosting distribution network. The simulations also displayed how the method can successfully adapt to different levels of uncertainty acting on the network (*e.g.* uncertain generation).

To provide an insight into the main advantages and disadvantages of the two proposed approximation schemes, a quantitative comparison was performed in Chapter 5. Based on the simulations and the experimental results, the following conclusions could be drawn:

- Relying on the optimistic relaxation of the causality requirements, the two-stage reformulation resulted in a higher amount of offered service. At the same time, it also showed more aggressive closed-loop behavior, with the state trajectories running very close to the bounds and few constraint violations. On the contrary, the multi-stage reformulation resulted in a more conservative solution but it provided an extra robustness margin with the state trajectories well within the constraints.

- The two-stage approach is not well suited for situations where the causality requirements have a prevalent influence on the structure of the solution. This is the case, *e.g.*, of the intraday policy which had to be separately determined before solving the planning problem. On the contrary, the causality requirements are respected for the multi-stage scheme and this allowed us to seamlessly include the intraday policy directly in the problem formulation.

Based on these observations, general guidelines on the method to be chosen for a given application were provided.

Part II - Provision of grid services using distributed energy resources: Experiments
This part answered the questions of the practical (technical) feasibility of deploying such advanced control techniques in a real-life setting, of what type of resources should be aggregated, and of the impact of the provision of such services on the primary purpose of the resources.

The control framework developed in Part II was adapted in order to accommodate two experimental case studies of practical interest. In Chapter 6, we considered an aggregation composed of a single-zone office building on the EPFL campus served by a commercial HP, and an electric battery. The

two resources acted as a virtual entity providing secondary frequency regulation in full compliance with the current Swiss regulations. The experimental results demonstrated that it is indeed possible to provide regulation services to the grid using distributed energy resources. In fact, excellent tracking performance could be achieved without compromising the comfort inside the building nor violating the operational constraints of the battery. Moreover, the significant benefit of combining complementary resources such as a building (slow) and a battery (fast) was shown:

- Aggregating the two resources in a single entity allowed a slow resource such as the building to participate in a high-frequency service.

- A 75% average increment in the offered flexibility was obtained with respect to the case of the battery alone providing the service.

- With respect to the current Swiss market conditions, a 26% reduction in the operational costs could be achieved with respect to the case of the battery providing the service and the building simply preserving comfort.

In Chapter 7, the problem of dispatching the operation of an active distribution feeder was investigated. An extensive experimental campaign was conducted on the EPFL MV feeder which is characterized by uncontrollable consumption, local renewable generation, a utility-scale battery, and a multi-zone office building. As for the previous application, the results showed once more the great benefit of considering complementary resources as the battery requirements, for the same size of uncontrollable prosumers, could be reduced by 57%. Moreover, also in this case, occupants' comfort was maintained at all times.

8.1   Future directions

There are many research directions worth exploring on the topics addressed in this thesis. In the following, we briefly present the most promising ones.

- Throughout the thesis, one of the underlying assumption has been the availability of models capturing the dynamical behavior of the resources to be controlled. In the experimental part, we have provided a few examples of how to identify or construct such models for specific cases. However, especially for complex systems, this process can prove to be time consuming and expensive, as it requires significant human expertise and a large amount of experimental data. In particular, in certain cases, this could even represent one of the main barriers to the utilization of an asset. It is therefore important to develop data-driven predictive control techniques to expedite the modeling and control design effort. A promising direction is represented by the recently introduced concept of *safe-learning* [1, 7] which has typically been considered for robotic applications. The main idea is to completely automatize the modeling step by through data-driven identification procedures that are conducted during regular operation of the resource. More specifically, at each iteration and given the available data, a model is constructed together with an associated confidence level. Based on this

information, appropriately computed excitation signals are applied to the system in order to maximize the learning rate, while also ensuring to not violate the constraints of the system. Developing similar techniques for DERs, with a particular focus in providing services to the power grid, is a subject surely worth of further investigation.

- Scale is typically required either by market regulations or in order to make a significant impact on the grid. We have shown in Chapter 4 how to efficiently schedule the operation of a pool of resources providing multiple services. However, there are a number of limitations of the framework that should be addressed. First, despite the linearity of the resulting optimization problem, computational limits might be reached for very large-scale aggregation of units (in the order of hundreds or thousands). Second, privacy could represent an issue. In fact, the structure of the control framework requires a single entity (the aggregator) to have access to the full details of any given resource, which, in certain situations, may not be realistic. Finally, especially for applications characterized by the presence of heterogeneous resources, there are a number of challenges for practical implementation, *e.g.*, asynchronous or unreliable communication, sudden disconnection of resources, etc. Thus, robust and scalable methods to decentralize or distribute the decision-making process across the network of resources, while addressing all the aforementioned issues, should be devised.

- The literature on DERs, particularly using the demand-side, for the provision of grid services lacks experimental demonstration and field tests which represent the key in the widespread development of these technologies. With the results shown in Chapter 6 and 7, we humbly believe to have made few important steps in this direction. However, more investigation is certainly required. For instance, the experimental campaigns presented in both chapters should be performed using larger commercial buildings with more complex HVAC systems. Year-long experiments should also be conducted to evaluate the impact of external factors, e.g. weather conditions, on the offerable amount of services. More importantly, the aggregation of a significant number of heterogeneous resources should be considered in order to test the scalability of the control methods. Finally, the optimal sizing of the resources for a given application should be examined by means of, *e.g.* sensitivity analysis.

Bibliography

[1] A. K. Akametalu, J. F. Fisac, J. H. Gillula, S. Kaynama, M. N. Zeilinger, and C. J. Tomlin. "Reachability-based safe learning with Gaussian processes". In: *53rd IEEE Conference on Decision and Control.* Dec. 2014, pp. 1424–1431. DOI: 10.1109/CDC.2014.7039601.

[2] E. Atam and L. Helsen. "Control-Oriented Thermal Modeling of Multizone Buildings: Methods and Issues: Intelligent Control of a Building System". In: *IEEE Control Systems Magazine* 36.3 (June 2016), pp. 86–111. ISSN: 1066-033X. DOI: 10.1109/MCS.2016.2535913.

[3] P. Barooah, A. Buic, and S. Meyn. "Spectral Decomposition of Demand-Side Flexibility for Reliable Ancillary Services in a Smart Grid". In: *2015 48th Hawaii International Conference on System Sciences.* Jan. 2015, pp. 2700–2709. DOI: 10.1109/HICSS.2015.325.

[4] I. Beil, I. Hiskens, and S. Backhaus. "Frequency Regulation From Commercial Building HVAC Demand Response". In: *Proceedings of the IEEE* 104.4 (Apr. 2016), pp. 745–757. ISSN: 0018-9219. DOI: 10.1109/JPROC.2016.2520640.

[5] Alberto Bemporad, Carlo Filippi, and Fabio D. Torrisi. "Inner and outer approximations of polytopes using boxes". In: *Computational Geometry* 27.2 (2004), pp. 151 –178. ISSN: 0925-7721. DOI: https://doi.org/10.1016/S0925-7721(03)00048-8. URL: http://www.sciencedirect.com/science/article/pii/S0925772103000488.

[6] A. Ben-Tal, L. El Ghaoui, and A.S. Nemirovski. *Robust Optimization.* Princeton Series in Applied Mathematics. Princeton University Press, Oct. 2009.

[7] F. Berkenkamp and A. P. Schoellig. "Safe and robust learning control with Gaussian processes". In: *2015 European Control Conference (ECC).* July 2015, pp. 2496–2501. DOI: 10.1109/ECC.2015.7330913.

[8] A. Berman and R. Plemmons. *Nonnegative Matrices in the Mathematical Sciences.* Society for Industrial and Applied Mathematics, 1994.

[9] A. Bernstein, C. Wang, E. Dall'Anese, J. Le Boudec, and C. Zhao. *Load-Flow in Multiphase Distribution Networks: Existence, Uniqueness, Non-Singularity and Linear Models.* 2017. eprint: arXiv:1702.03310.

[10] Dimitris Bertsimas, Vishal Gupta, and Nathan Kallus. "Data-driven robust optimization". In: *Mathematical Programming* 167.2 (Feb. 2018), pp. 235–292. ISSN: 1436-4646. DOI: 10.1007/s10107-017-1125-8. URL: https://doi.org/10.1007/s10107-017-1125-8.

[11] S. Bittanti and P. Colaneri. *Periodic Systems: Filtering and Control.* Springer Science & Business Media, 2008, p. 424. ISBN: 1848009119. URL: `https://books.google.com/books?id=AVdJAAAAQBAJ\&pgis=1`.

[12] F. Blanchini. "Set invariance in control". In: *Automatica* 35.11 (Nov. 1999), pp. 1747–1767. ISSN: 00051098. DOI: `10.1016/S0005-1098(99)00113-2`. URL: `http://www.sciencedirect.com/science/article/pii/S0005109899001132`.

[13] David H. Blum and Leslie K. Norford. "Dynamic simulation and analysis of ancillary service demand response strategies for variable air volume HVAC systems". In: *HVAC&R Research* 20.8 (2014), pp. 908–921. DOI: `10.1080/10789669.2014.958975`. eprint: `https://doi.org/10.1080/10789669.2014.958975`. URL: `https://doi.org/10.1080/10789669.2014.958975`.

[14] David H. Blum and Leslie K. Norford. "Dynamic simulation of regulation demand response by VAV HVAC systems". In: *ASHRAE/IBPSA-USA Buildings Simulation Conference, Atlanta, GA* (Jan. 2014), pp. 402–409.

[15] C. Böhm, S. Yu, and F. Allgöwer. "Predictive control for constrained discrete-time periodic systems using a time-varying terminal region". In: *Methods and Models in Automation and Robotics.* Vol. 14. Aug. 2009, pp. 537–542. ISBN: 978-3-902661-55-5. URL: `http://www.ifac-papersonline.net/Detailed/41135.html`.

[16] S. Bolognani and F. Dörfler. "Fast power system analysis via implicit linearization of the power flow manifold". In: *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton).* Sept. 2015, pp. 402–409. DOI: `10.1109/ALLERTON.2015.7447032`.

[17] S. Bolognani and S. Zampieri. "On the Existence and Linear Approximation of the Power Flow Solution in Power Distribution Networks". In: *IEEE Transactions on Power Systems* 31.1 (Jan. 2016), pp. 163–172. ISSN: 0885-8950. DOI: `10.1109/TPWRS.2015.2395452`.

[18] T. Borsche, A. Ulbig, and G. Andersson. "A new frequency control reserve framework based on energy-constrained units". In: *2014 Power Systems Computation Conference.* Aug. 2014, pp. 1–7. DOI: `10.1109/PSCC.2014.7038111`.

[19] T. Borsche, A. Ulbig, M. Koller, and G. Andersson. "Power and energy capacity requirements of storages providing frequency control reserves". In: *2013 IEEE Power Energy Society General Meeting.* July 2013, pp. 1–5. DOI: `10.1109/PESMG.2013.6672843`.

[20] A. van den Bos. *Parameter Estimation for Scientists and Engineers.* John Wiley and Sons, 2007, pp. 259–263. ISBN: 9780470173862.

[21] R. Cagienard, P. Grieder, E. C. Kerrigan, and M. Morari. "Move blocking strategies in receding horizon control". In: *2004 43rd IEEE Conference on Decision and Control (CDC) (IEEE Cat. No.04CH37601).* Vol. 2. Dec. 2004, 2023–2028 Vol.2. DOI: `10.1109/CDC.2004.1430345`.

[22] D. S. Callaway and I. A. Hiskens. "Achieving Controllability of Electric Loads". In: *Proceedings of the IEEE* 99.1 (Jan. 2011), pp. 184–199. ISSN: 0018-9219. DOI: `10.1109/JPROC.2010.2081652`.

[23] M. Cannon, B. Kouvaritakis, and D. Ng. "Probabilistic tubes in linear stochastic model predictive control". In: *Systems & Control Letters* 58.10-11 (Oct. 2009), pp. 747–753. ISSN: 01676911. DOI: `10.1016/j.sysconle.2009.08.004`. URL: `http://www.sciencedirect.com/science/article/pii/S0167691109001078`.

[24] M. Cannon, B. Kouvaritakis, S. V. Rakovic, and Q. Cheng. "Stochastic Tubes in Model Predictive Control With Probabilistic Constraints". In: *IEEE Transactions on Automatic Control* 56.1 (Jan. 2011), pp. 194–200. ISSN: 0018-9286. DOI: `10.1109/TAC.2010.2086553`. URL: `http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5599849`.

[25] Salvatore Carlucci. *Thermal comfort assessment of buildings.* Springer, 2013.

[26] Bolong Cheng and Warren Powell. "Co-optimizing battery storage for the frequency regulation and energy arbitrage using multi-scale dynamic programming". In: *IEEE Transactions on Smart Grid* (2016).

[27] K. Christakou, J. LeBoudec, M. Paolone, and D. Tomozei. "Efficient Computation of Sensitivity Coefficients of Node Voltages and Line Currents in Unbalanced Radial Electrical Distribution Networks". In: *IEEE Transactions on Smart Grid* 4.2 (June 2013), pp. 741–750. ISSN: 1949-3053. DOI: `10.1109/TSG.2012.2221751`.

[28] Konstantina Christakou, Dan-Cristian Tomozei, Maryam Bahramipanah, Jean-Yves Le Boudec, and Mario Paolone. "Primary voltage control in active distribution networks via broadcast signals: The case of distributed storage". In: *IEEE Transactions on Smart Grid* 5.5 (2014), pp. 2314–2325.

[29] Andrea Costa, Marcus M. Keane, J. Ignacio Torrens, and Edward Corry. "Building operation and energy performance: Monitoring, analysis and optimisation toolkit". In: *Applied Energy* 101 (Jan. 2013), pp. 310–316. ISSN: 03062619. DOI: `10.1016/j.apenergy.2011.10.037`. URL: `http://www.sciencedirect.com/science/article/pii/S030626191100691X`.

[30] Emiliano Dall'Anese, Kyri Baker, and Tyler Summers. "Chance-constrained AC optimal power flow for distribution systems with renewables". In: *IEEE Transactions on Power Systems* 32.5 (2017), pp. 3427–3438.

[31] S. V. Dhople, S. S. Guggilam, and Y. C. Chen. "Linear approximations to AC power flow in rectangular coordinates". In: *2015 53rd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. Sept. 2015, pp. 211–217. DOI: `10.1109/ALLERTON.2015.7447006`.

[32] Florian Dörfler and Francesco Bullo. "Kron Reduction of Graphs With Applications to Electrical Networks". In: *IEEE Transactions on Circuits and Systems I: Regular Papers* 60 (2013), pp. 150–163.

[33] Easan Drury, Paul Denholm, and Ramteen Sioshansi. "The value of compressed air energy storage in energy and reserve markets". In: *Energy* 36.8 (2011), pp. 4959–4973.

[34] Heloise Dutrieux, Gauthier Delille, Gilles Malarange, and Bruno Francois. "An energy supervision for distributed storage systems to optimize the provision of multiple services". In: *PowerTech (POWERTECH), 2013 IEEE Grenoble*. IEEE. 2013, pp. 1–6.

[35]   Erik Ela, Michael Milligan, and Brendan Kirby. *Operating reserves and variable generation.* Tech. rep. 2011, pp. 275–3000. URL: http://www2.econ.iastate.edu/tesfatsi/OperatingReservesVariableGenerationSurvey.NRELAug2011.pdf (visited on 11/22/2016).

[36]   US Department of Energy: Office of Energy Efficiency and Renewable Energy. *Commercial Reference Buildings.* 2016. URL: http://energy.gov/eere/buildings (visited on 08/14/2018).

[37]   Jonas Engels, Bert Claessens, and Geert Deconinck. "Combined stochastic optimization of frequency control and self-consumption with a battery". In: *IEEE Transactions on Smart Grid* (2017).

[38]   J Eyer and G Corey. "Energy storage for the electricity grid: Benefits and market potential assessment guide". In: (Jan. 2011), pp. 1–232.

[39]   L. Fabietti, T. Gorecki, F. Qureshi, A. Bitlislioglu, I. Lymperopoulos, and C. Jones. "Experimental Implementation of Frequency Regulation Services Using Commercial Buildings". In: *IEEE Transactions on Smart Grid* PP.99 (2016), pp. 1–1. ISSN: 1949-3053. DOI: 10.1109/TSG.2016.2597002.

[40]   Luca Fabietti, Tomasz Tadeusz Gorecki, Emil Namor, Fabrizio Sossan, Mario Paolone, and Colin Neil Jones. "Dispatching active distribution networks through electrochemical storage systems and demand side management". In: *1st IEEE Conference on Control Technology and Applications.* Big Island, Hawaii, USA, 2017.

[41]   Luca Fabietti, Tomasz T. Gorecki, Emil Namor, Fabrizio Sossan, Mario Paolone, and Colin N. Jones. "Enhancing the dispatchability of distribution networks through utility-scale batteries and flexible demand". In: *Energy and Buildings* 172 (2018), pp. 125 –138. ISSN: 0378-7788. DOI: https://doi.org/10.1016/j.enbuild.2018.04.056. URL: http://www.sciencedirect.com/science/article/pii/S0378778817325197.

[42]   Michael G. Forbes, Rohit S. Patwardhan, Hamza Hamadah, and R. Bhushan Gopaluni. "Model Predictive Control in Industry: Challenges and Opportunities". In: *IFAC-PapersOnLine* 48.8 (2015). 9th IFAC Symposium on Advanced Control of Chemical Processes ADCHEM 2015, pp. 531 –538. ISSN: 2405-8963. DOI: https://doi.org/10.1016/j.ifacol.2015.09.022. URL: http://www.sciencedirect.com/science/article/pii/S2405896315011039.

[43]   Pacific Gas and Electric Company. *Peak day pricing.* 2012. URL: http://www.pge.com/pdp.

[44]   G. Goddard, J. Klose, and S. Backhaus. "Model Development and Identification for Fast Demand Response in Commercial HVAC Systems". In: *IEEE Transactions on Smart Grid* 5.4 (July 2014), pp. 2084–2092. ISSN: 1949-3053. DOI: 10.1109/TSG.2014.2312430.

[45]   R. Gondhalekar, F. Oldewurtel, and C.N. Jones. "Least-restrictive robust periodic model predictive control applied to room temperature regulation". In: *Automatica* 49.9 (Sept. 2013), pp. 2760–2766. ISSN: 00051098. DOI: 10.1016/j.automatica.2013.05.009. URL: http://www.sciencedirect.com/science/article/pii/S0005109813002793.

[46] T. T. Gorecki, F. A. Qureshi, and C. N. Jones. "OpenBuild : An integrated simulation environment for building control". In: *2015 IEEE Conference on Control Applications (CCA)*. Sept. 2015, pp. 1522–1527. DOI: 10.1109/CCA.2015.7320826.

[47] T. T. Gorecki, A. Bitlislioğlu, G. Stathopoulos, and C. N. Jones. "Guaranteeing input tracking for constrained systems: Theory and application to demand response". In: *2015 American Control Conference (ACC)*. July 2015, pp. 232–237. DOI: 10.1109/ACC.2015.7170741.

[48] Tomasz T. Gorecki, Luca Fabietti, Faran A. Qureshi, and Colin N. Jones. "Experimental demonstration of buildings providing frequency regulation services in the Swiss market". In: *Energy and Buildings* 144 (2017), pp. 229 –240. ISSN: 0378-7788. DOI: http://doi.org/10.1016/j.enbuild.2017.02.050. URL: http://www.sciencedirect.com/science/article/pii/S0378778816311616.

[49] T.T Gorecki, F.A Qureshi, and C.N. Jones. *OpenBuild*. http://la.epfl.ch/openBuild. 2015.

[50] T.T. Gorecki, F.A. Qureshi, and C.N. Jones. "OpenBuild : An integrated simulation environment for building control". In: *Control Applications (CCA), 2015 IEEE Conference on*. Sept. 2015, pp. 1522–1527. DOI: 10.1109/CCA.2015.7320826.

[51] Bram L. Gorissen, İhsan Yanıkoğlu, and Dick den Hertog. "A practical guide to robust optimization". In: *Omega* 53 (2015), pp. 124 –137. ISSN: 0305-0483. DOI: https://doi.org/10.1016/j.omega.2014.12.006. URL: http://www.sciencedirect.com/science/article/pii/S0305048314001698.

[52] M. Gwerder and J. Todli. "Predictive Control for Integrated Room Automation". In: *8th REHVA World Congress for Building Technologies-CLIMA*. 2005, pp. 1–6. URL: http://www.sysecol2.ethz.ch/OptiControl/Lit/Gwer_05_Proc-Clima2005_244s.pdf.

[53] R. Halvgaard, N. K. Poulsen, H. Madsen, and J. B. Jørgensen. "Economic Model Predictive Control for building climate control in a Smart Grid". In: *2012 IEEE PES Innovative Smart Grid Technologies (ISGT)*. Jan. 2012, pp. 1–6. DOI: 10.1109/ISGT.2012.6175631.

[54] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent. "Aggregate Flexibility of Thermostatically Controlled Loads". In: *IEEE Transactions on Power Systems* 30.1 (Jan. 2015), pp. 189–198. ISSN: 0885-8950. DOI: 10.1109/TPWRS.2014.2328865.

[55] Reza Hemmati, Hedayat Saboori, and Mehdi Ahmadi Jirdehi. "Stochastic planning and scheduling of energy storage systems for congestion management in electric power systems including renewable energy resources". In: *Energy* 133 (2017), pp. 380–387.

[56] Karen Herter, Patrick McAuliffe, and Arthur Rosenfeld. "An exploratory analysis of California residential customer response to critical peak pricing of electricity". In: *Energy* 32.1 (2007), pp. 25 –34. ISSN: 0360-5442. DOI: https://doi.org/10.1016/j.energy.2006.01.014. URL: http://www.sciencedirect.com/science/article/pii/S0360544206000314.

[57] C. Jin, N. Lu, S. Lu, Y. V. Makarov, and R. A. Dougal. "A Coordinating Algorithm for Dispatching Regulation Services Between Slow and Fast Power Regulating Resources". In: *IEEE Transactions on Smart Grid* 5.2 (Mar. 2014), pp. 1043–1050. ISSN: 1949-3053. DOI: 10.1109/TSG.2013.2277974.

[58] Georgios Marios Karagiannis, Stamatios Chondrogiannis, Elisabeth Krausmann, and Zehra Irem Turksezer. *Power grid recovery after natural hazard impact.* Tech. rep. Dec. 2017.

[59] Tosio KATO. "On the perturbation theory of closed linear operators." In: *J. Math. Soc. Japan* 4.3-4 (Dec. 1952), pp. 323–337. DOI: 10.2969/jmsj/00430323. URL: https://doi.org/10.2969/jmsj/00430323.

[60] Mostafa Kazemi and Hamidreza Zareipour. "Long-term Scheduling of Battery Storage Systems in Energy and Regulation Markets Considering Battery's lifespan". In: *IEEE Transactions on Smart Grid* (2017).

[61] Mostafa Kazemi, Hamidreza Zareipour, Nima Amjady, William D Rosehart, and Mehdi Ehsan. "Operation Scheduling of Battery Storage Systems in Joint Energy and Ancillary Services Markets". In: *IEEE Transactions on Sustainable Energy* (2017).

[62] B. Kern, C. Böhm, R. Findeisen, and F. Allgöwer. *Nonlinear Model Predictive Control.* Vol. 384. Lecture Notes in Control and Information Sciences. Berlin, Heidelberg: Springer Berlin Heidelberg, May 2009, pp. 109–117. ISBN: 978-3-642-01093-4. DOI: 10.1007/978-3-642-01094-1. URL: http://www.researchgate.net/publication/225412965_Receding_Horizon_Control_for_Linear_Periodic_Time-Varying_Systems_Subject_to_Input_Constraints.

[63] Y. J. Kim and J. Wang. "Power Hardware-in-the-Loop Simulation Study on Frequency Regulation Through Direct Load Control of Thermal and Electrical Energy Storage Resources". In: *IEEE Transactions on Smart Grid* 9.4 (July 2018), pp. 2786–2796. ISSN: 1949-3053. DOI: 10.1109/TSG.2016.2620176.

[64] Brendan J Kirby. *Frequency Regulation Basics and Trends.* 12. 2004, pp. 1–32. ISBN: 1800553684.

[65] Michael Koller, Marina González Vayá, Aby Chacko, Theodor Borsche, and Andreas Ulbig. "Primary control reserves provision with battery energy storage systems in the largest European ancillary services cooperation". en. In: *Set of papers, CIGRE session 46 : 21-26 August 2016, Paris.* CIGRE Session 2016; Conference Location: Paris, France; Conference Date: August 21-26, 2016; . CIGRE, 2016, 361–NCA–C2–PS1.

[66] Michael Koller, Theodor Borsche, Andreas Ulbig, and Göran Andersson. "Review of grid applications with the Zurich 1MW battery energy storage system". In: *Electric Power Systems Research* 120 (2015). Smart Grids: World's Actual Implementations, pp. 128 –135. ISSN: 0378-7796. DOI: https://doi.org/10.1016/j.epsr.2014.06.023. URL: http://www.sciencedirect.com/science/article/pii/S0378779614002326.

[67] M. Korda, R. Gondhalekar, F. Oldewurtel, and C.N. Jones. "Stochastic model predictive control: Controlling the average number of constraint violations". In: *2012 IEEE 51st IEEE Conference on Decision and Control (CDC).* IEEE, Dec. 2012, pp. 4529–4536. ISBN: 978-1-4673-2066-5. DOI: 10.1109/CDC.2012.6426873. URL: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6426873.

[68] M. Korda, R. Gondhalekar, F. Oldewurtel, and C.N. Jones. "Stochastic MPC Framework for Controlling the Average Constraint Violation". In: *IEEE Transactions on Automatic Control* 59.7 (July 2014), pp. 1706–1721. ISSN: 0018-9286. DOI: 10.1109/TAC.2014.2310066. URL: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6756951.

[69] M. Korda, R. Gondhalekar, J. Cigler, and F. Oldewurtel. "Strongly feasible stochastic model predictive control". In: *IEEE Conference on Decision and Control and European Control Conference*. IEEE, Dec. 2011, pp. 1245–1251. ISBN: 978-1-61284-801-3. DOI: 10.1109/CDC.2011.6161250. URL: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6161250.

[70] P. Kundur. *Power System Stability and Control*. McGraw-Hill, 1994.

[71] H. J. Kunisch, K. G. Kramer, and H. Dominik. "Battery Energy Storage Another Option for Load-Frequency-Control and Instantaneous Reserve". In: *IEEE Transactions on Energy Conversion* EC-1.3 (Sept. 1986), pp. 41–46. ISSN: 0885-8969. DOI: 10.1109/TEC.1986.4765732.

[72] Jason Leadbetter and Lukas Swan. "Battery storage system for residential electricity peak demand shaving". In: *Energy and Buildings* 55 (2012). Cool Roofs, Cool Pavements, Cool Cities, and Cool World, pp. 685 –692. ISSN: 0378-7788. DOI: https://doi.org/10.1016/j.enbuild.2012.09.035. URL: http://www.sciencedirect.com/science/article/pii/S0378778812004896.

[73] J.H. Lee, S. Natarajan, and K.S. Lee. "A model-based predictive control approach to repetitive control of continuous processes with periodic operations". In: *Journal of Process Control* 11.2 (Apr. 2001), pp. 195–207. ISSN: 09591524. DOI: 10.1016/S0959-1524(00)00047-0. URL: http://www.sciencedirect.com/science/article/pii/S0959152400000470.

[74] Xiwang Li and Jin Wen. "Review of building energy modeling for control and operation". In: *Renewable and Sustainable Energy Reviews* 37 (2014), pp. 517 –537. ISSN: 1364-0321. DOI: https://doi.org/10.1016/j.rser.2014.05.056. URL: http://www.sciencedirect.com/science/article/pii/S1364032114003815.

[75] Y. Lin, P. Barooah, S. Meyn, and T. Middelkoop. "Experimental Evaluation of Frequency Regulation From Commercial Building HVAC Systems". In: *IEEE Transactions on Smart Grid* 6.2 (Mar. 2015), pp. 776–783. ISSN: 1949-3053. DOI: 10.1109/TSG.2014.2381596.

[76] Yashen Lin, Prabir Barooah, Sean Meyn, and Timothy Middelkoop. "Experimental Evaluation of Frequency Regulation from Commercial Building HVAC system". In: *IEEE Transaction on Smart Grid* 6.no. 2 (2015), pp. 776–783.

[77] Wai Hou Lio, J. A. Rossiter, and B. L. Jones. "A review on applications of model predictive control to wind turbines". In: *2014 UKACC International Conference on Control (CONTROL)*. July 2014, pp. 673–678. DOI: 10.1109/CONTROL.2014.6915220.

[78] L Ljung. *Ljung L System Identification Theory for User*. Vol. 25. PTR Prentice Hall Upper Saddle River NJ, 1987, pp. 475–476. ISBN: 0138816409. DOI: 10.1016/0005-1098(89)90019-8.

[79] Lennart Ljung. "System identification". In: *Signal Analysis and Prediction*. Springer, 1998, pp. 163–173.

[80] J. Löfberg. "YALMIP : A Toolbox for Modeling and Optimization in MATLAB". In: *In Proceedings of the CACSD Conference*. Taipei, Taiwan, 2004.

[81] Johan Lofberg. "YALMIP: A toolbox for modeling and optimization in MATLAB". In: *Computer Aided Control Systems Design, 2004 IEEE International Symposium on*. IEEE. 2004, pp. 284–289.

[82] S. H. Low. "Convex Relaxation of Optimal Power Flow—Part I: Formulations and Equivalence". In: *IEEE Transactions on Control of Network Systems* 1.1 (Mar. 2014), pp. 15–27. ISSN: 2325-5870. DOI: 10.1109/TCNS.2014.2309732.

[83] J. Ma, S. J. Qin, B. Li, and T. Salsbury. "Economic model predictive control for building energy systems". In: *ISGT 2011*. Jan. 2011, pp. 1–6. DOI: 10.1109/ISGT.2011.5759140.

[84] Jingran Ma, Joe Qin, Timothy Salsbury, and Peng Xu. "Demand reduction in building energy systems based on economic model predictive control". In: *Chemical Engineering Science* 67.1 (2012). Dynamics, Control and Optimization of Energy Systems, pp. 92 –100. ISSN: 0009-2509. DOI: https://doi.org/10.1016/j.ces.2011.07.052. URL: http://www.sciencedirect.com/science/article/pii/S0009250911005240.

[85] Jingran Ma, Joe Qin, Timothy Salsbury, and Peng Xu. "Demand reduction in building energy systems based on economic model predictive control". In: *Chemical Engineering Science* 67.1 (Jan. 2012), pp. 92–100. ISSN: 00092509. DOI: 10.1016/j.ces.2011.07.052. URL: http://www.sciencedirect.com/science/article/pii/S0009250911005240.

[86] Yudong Ma, Anthony Kelman, Allan Daly, and Francesco Borrelli. "Predictive Control for Energy Efficient Buildings with Thermal Storage: Modeling, Stimulation, and Experiments". In: *IEEE Control Systems* 32.1 (Feb. 2012), pp. 44–64. ISSN: 1066-033X. DOI: 10.1109/MCS.2011.2172532. URL: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6153586.

[87] M. Maasoumy, C. Rosenberg, A. Sangiovanni-Vincentelli, and D.S. Callaway. "Model predictive control approach to online computation of demand-side flexibility of commercial buildings HVAC systems for Supply Following". In: *American Control Conference (ACC), 2014*. June 2014, pp. 1082–1089.

[88] J. MacDonald and S. Kiliccote. "Commercial buildings loads providing ancillary services in PJM". In: *ACEEE Summer Study Energy Efficiency Buildings*. Aug. 2014, pp. 1–5. DOI: 10.1109/PESMG.2013.6672843.

[89] K. Margellos, P. Goulart, and J. Lygeros. "On the Road Between Robust Optimization and the Scenario Approach for Chance Constrained Optimization Problems". In: *IEEE Transactions on Automatic Control* 59.8 (Aug. 2014), pp. 2258–2263. ISSN: 0018-9286. DOI: 10.1109/TAC.2014.2303232.

[90] J. L. Mathieu, S. Koch, and D. S. Callaway. "State Estimation and Control of Electric Loads to Manage Real-Time Energy Imbalance". In: *IEEE Transactions on Power Systems* 28.1 (Feb. 2013), pp. 430–440. ISSN: 0885-8950. DOI: 10.1109/TPWRS.2012.2204074.

[91] D.Q. Mayne, J.B. Rawlings, C.V. Rao, and P.O.M. Scokaert. "Constrained model predictive control: Stability and optimality". In: *Automatica* 36.6 (June 2000), pp. 789–814. ISSN: 00051098. DOI: 10.1016/S0005-1098(99)00214-9. URL: http://www.sciencedirect.com/science/article/pii/S0005109899002149.

[92] Olivier Mégel, Johanna L Mathieu, and Göran Andersson. "Scheduling distributed energy storage units to provide multiple services under forecast error". In: *International Journal of Electrical Power & Energy Systems* 72 (2015), pp. 48–57.

[93] S. Meyn, M. Negrete-Pincetic, G. Wang, A. Kowli, and E. Shafieepoorfard. "The value of volatile resources in electricity markets". In: *49th IEEE Conference on Decision and Control (CDC)*. Dec. 2010, pp. 1029–1036. DOI: 10.1109/CDC.2010.5717327.

[94] Roberto Moreira, Rodrigo Moreno, and Goran Strbac. "Synergies and conflicts among energy storage services". In: *Energy Conference (ENERGYCON), 2016 IEEE International*. IEEE. 2016, pp. 1–6.

[95] Rodrigo Moreno, Roberto Moreira, and Goran Strbac. "A MILP model for optimising multi-service portfolios of distributed energy storage". In: *Applied Energy* 137 (2015), pp. 554–566.

[96] Emil Namor, Fabrizio Sossan, Rachid Cherkaoui, and Mario Paolone. "Control of Battery Storage Systems for the Simultaneous Provision of Multiple Services". In: *IEEE Transactions on Smart Grid* (2018).

[97] Emil Namor, Fabrizio Sossan, Rachid Cherkaoui, and Mario Paolone. "Load leveling and dispatchability of a medium voltage active feeder through battery energy storage systems: Formulation of the control problem and experimental validation". In: *PES Innovative Smart Grid Technologies Conference Europe (ISGT-Europe), 2016 IEEE*. IEEE. 2016, pp. 1–6.

[98] Emil Namor, Dimitri Torregrossa, Rachid Cherkaoui, and Mario Paolone. "Parameter identification of a lithium-ion cell single-particle model through non-invasive testing". In: *Journal of Energy Storage* 12 (2017), pp. 138 –148. ISSN: 2352-152X. DOI: https://doi.org/10.1016/j.est.2017.04.008. URL: http://www.sciencedirect.com/science/article/pii/S2352152X1630278X.

[99] H. Nazaripouya, Y. Wang, P. Chu, H. R. Pota, and R. Gadh. "Optimal sizing and placement of battery energy storage in distribution system based on solar size for voltage regulation". In: *2015 IEEE Power Energy Society General Meeting*. July 2015, pp. 1–5. DOI: 10.1109/PESGM.2015.7286059.

[100] T. X. Nghiem and G. J. Pappas. "Receding-horizon supervisory control of green buildings". In: *Proceedings of the 2011 American Control Conference*. June 2011, pp. 4416–4421. DOI: 10.1109/ACC.2011.5990995.

[101] T. X. Nghiem, A. Bitlislioğlu, T. Gorecki, F. A. Qureshi, and C. N. Jones. "OpenBuildNet framework for distributed co-simulation of smart energy systems". In: *2016 14th International Conference on Control, Automation, Robotics and Vision (ICARCV)*. Nov. 2016, pp. 1–6. DOI: 10.1109/ICARCV.2016.7838701.

[102] Ali Nourai. "Installation of the First Distributed Energy Storage System (DESS) at American Electric Power (AEP)". In: (Jan. 2007).

[103]   F. Oldewurtel, A. Ulbig, A. Parisio, G. Andersson, and M. Morari. "Reducing peak electricity demand in building climate control using real-time pricing and model predictive control". In: *Decision and Control (CDC), 2010 49th IEEE Conference on*. Dec. 2010, pp. 1927–1932. DOI: 10.1109/CDC.2010.5717458.

[104]   F. Oldewurtel, A. Ulbig, A. Parisio, G. Andersson, and M. Morari. "Reducing peak electricity demand in building climate control using real-time pricing and model predictive control". In: *49th IEEE Conference on Decision and Control (CDC)*. Dec. 2010, pp. 1927–1932. DOI: 10.1109/CDC.2010.5717458.

[105]   Frauke Oldewurtel, Theodor Borsche, Matthias Bucher, Philipp Fortenbacher, Marina Gonzalez Vaya Tobias Haring, Tobias Haring, Johanna L Mathieu, Olivier Mégel, Evangelos Vrettos, and Göran Andersson. "A framework for and assessment of demand response and energy storage in power systems". In: *Bulk Power System Dynamics and Control-IX Optimization, Security and Control of the Emerging Power Grid (IREP), 2013 IREP Symposium*. IEEE. 2013, pp. 1–24.

[106]   Frauke Oldewurtel, Alessandra Parisio, Colin N Jones, Dimitrios Gyalistras, Markus Gwerder, Vanessa Stauch, Beat Lehmann, and Manfred Morari. "Use of model predictive control and weather forecasts for energy efficient building climate control". In: *Energy and Buildings* 45 (2012), pp. 15–27.

[107]   Frauke Oldewurtel, Alessandra Parisio, Colin N. Jones, Dimitrios Gyalistras, Markus Gwerder, Vanessa Stauch, Beat Lehmann, and Manfred Morari. "Use of model predictive control and weather forecasts for energy efficient building climate control". In: *Energy and Buildings* 45 (2012), pp. 15 –27. ISSN: 0378-7788. DOI: https://doi.org/10.1016/j.enbuild.2011.09.022. URL: http://www.sciencedirect.com/science/article/pii/S0378778811004105.

[108]   A. Oudalov, D. Chartouni, and C. Ohler. "Optimizing a Battery Energy Storage System for Primary Frequency Control". In: *IEEE Transactions on Power Systems* 22.3 (Aug. 2007), pp. 1259–1266. ISSN: 0885-8950. DOI: 10.1109/TPWRS.2007.901459.

[109]   Alexandre Oudalov, Rachid Cherkaoui, and Antoine Beguin. "Sizing and optimal operation of battery energy storage system for peak shaving application". In: *Power Tech, 2007 IEEE Lausanne*. IEEE. 2007, pp. 621–625.

[110]   Alessandra Parisio, Marco Molinari, Damiano Varagnolo, and Karl Henrik Johansson. "A scenario-based predictive control approach to building HVAC management systems". In: *2013 IEEE International Conference on Automation Science and Engineering (CASE)* (Aug. 2013), pp. 428–435. DOI: 10.1109/CoASE.2013.6654024. URL: http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6654024.

[111]   Alessandra Parisio, Luca Fabietti, Marco Molinari, Damiano Varagnolo, and Karl H Johansson. "Control of HVAC Systems via Scenario-based Explicit MPC". In: *Conference on Decision and Control (CDC)* (2014), pp. 5201–5207.

[112]   A. Perez, R. Moreno, R. Moreira, M. Orchard, and G. Strbac. "Effect of Battery Degradation on Multi-Service Portfolios of Energy Storage". In: *IEEE Transactions on Sustainable Energy* 7.4 (Oct. 2016), pp. 1718–1729. ISSN: 1949-3029. DOI: 10.1109/TSTE.2016.2589943.

[113] Aramis Perez, Rodrigo Moreno, Roberto Moreira, Marcos Orchard, and Goran Strbac. "Effect of Battery Degradation on Multi-Service Portfolios of Energy Storage". In: *IEEE Transactions on Sustainable Energy* 7.4 (2016), pp. 1718–1729.

[114] Marco Pignati, Miroslav Popovic, Sergio Barreto, Rachid Cherkaoui, German Dario Flores, Jean-Yves Le Boudec, Maaz Mohiuddin, Mario Paolone, Paolo Romano, Styliani Sarri, et al. "Real-time state estimation of the EPFL-campus medium-voltage grid by using PMUs". In: *Innovative Smart Grid Technologies Conference (ISGT), 2015 IEEE Power & Energy Society*. IEEE. 2015, pp. 1–5.

[115] Pierre Pinson, Henrik Madsen, Henrik Aa Nielsen, George Papaefthymiou, and Bernd Klöckl. "From probabilistic forecasts to statistical scenarios of short-term wind power production". In: *Wind energy* 12.1 (2009), pp. 51–62.

[116] Faran A. Qureshi, Tomasz T. Gorecki, and Colin N. Jones. "Model Predictive Control for Market-Based Demand Response Participation". In: *IFAC Proceedings Volumes* 47.3 (2014). 19th IFAC World Congress, pp. 11153 –11158. ISSN: 1474-6670. DOI: `https://doi.org/10.3182/20140824-6-ZA-1003.02395`. URL: `http://www.sciencedirect.com/science/article/pii/S1474667016433884`.

[117] Faran A. Qureshi and Colin N. Jones. "Hierarchical control of building HVAC system for ancillary services provision". In: *Energy and Buildings* 169 (2018), pp. 216 –227. ISSN: 0378-7788. DOI: `https://doi.org/10.1016/j.enbuild.2018.03.004`. URL: `http://www.sciencedirect.com/science/article/pii/S0378778817315190`.

[118] Faran Ahmed Qureshi, Ioannis Lymperopoulos, Ali Ahmadi Khatir, and Colin Jones. *Economic Advantages of Office Buildings Providing Ancillary Services with Intraday Participation*. Tech. rep. 2016.

[119] Y. G. Rebours, D. S. Kirschen, M. Trotignon, and S. Rossignol. "A Survey of Frequency and Voltage Control Ancillary Services amp;mdash;Part I: Technical Features". In: *IEEE Transactions on Power Systems* 22.1 (Feb. 2007), pp. 350–357. ISSN: 0885-8950. DOI: `10.1109/TPWRS.2006.888963`.

[120] Y. G. Rebours, D. S. Kirschen, M. Trotignon, and S. Rossignol. "A Survey of Frequency and Voltage Control Ancillary Services amp;mdash;Part II: Economic Features". In: *IEEE Transactions on Power Systems* 22.1 (Feb. 2007), pp. 358–366. ISSN: 0885-8950. DOI: `10.1109/TPWRS.2006.888965`.

[121] Felix Rey, Xiaojing Zhang, Sandro Merkli, Valentina Agliati, Maryam Kamgarpour, and John Lygeros. "Strengthening the Group - Aggregated Frequency Reserve Bidding with ADMM". In: PP (Nov. 2017).

[122] Caleb Robinson, Bistra Dilkina, Jeffrey Hubbs, Wenwen Zhang, Subhrajit Guhathakurta, Marilyn A. Brown, and Ram M. Pendyala. "Machine learning approaches for estimating commercial building energy consumption". In: *Applied Energy* 208 (2017), pp. 889 –904. ISSN: 0306-2619. DOI: `https://doi.org/10.1016/j.apenergy.2017.09.060`. URL: `http://www.sciencedirect.com/science/article/pii/S0306261917313429`.

[123] Peter Rockett and Elizabeth Abigail Hathway. "Model-predictive control for non-domestic buildings: a critical review and prospects". In: *Building Research & Information* 45.5 (2017), pp. 556–571. DOI: 10.1080/09613218.2016.1139885. eprint: https://doi.org/10.1080/09613218.2016.1139885. URL: https://doi.org/10.1080/09613218.2016.1139885.

[124] R. T. de Salis, A. Clarke, Z. Wang, J. Moyne, and D. M. Tilbury. "Energy storage control for peak shaving in a single building". In: *2014 IEEE PES General Meeting | Conference Exposition.* July 2014, pp. 1–5. DOI: 10.1109/PESGM.2014.6938948.

[125] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on stochastic programming: modeling and theory.* 2009, p. 447. ISBN: 089871687X. DOI: http://dx.doi.org/10.1137/1.9780898718751.

[126] Yuanyuan Shi, Bolun Xu, Di Wang, and Baosen Zhang. "Using battery storage for peak shaving and frequency regulation: Joint optimization for superlinear gains". In: *arXiv preprint arXiv:1702.08065* (2017).

[127] F. Sossan, E. Namor, R. Cherkaoui, and M. Paolone. "Achieving the Dispatchability of Distribution Feeders Through Prosumers Data Driven Forecasting and Model Predictive Control of Electrochemical Storage". In: *IEEE Transactions on Sustainable Energy* 7.4 (Oct. 2016), pp. 1762–1777. ISSN: 1949-3029. DOI: 10.1109/TSTE.2016.2600103.

[128] F. Sossan, E. Namor, R. Cherkaoui, and M. Paolone. "Achieving the Dispatchability of Distribution Feeders Through Prosumers Data Driven Forecasting and Model Predictive Control of Electrochemical Storage". In: *IEEE Transactions on Sustainable Energy* 7.4 (Oct. 2016), pp. 1762–1777. ISSN: 1949-3029. DOI: 10.1109/TSTE.2016.2600103.

[129] Eleni Stai, Lorenzo Reyes-Chamorro, Fabrizio Sossan, Jean-Yves Le Boudec, and Mario Paolone. "Dispatching stochastic heterogeneous resources accounting for grid and battery losses". In: *IEEE Transactions on Smart Grid* (2017).

[130] B. Stott, J. Jardim, and O. Alsac. "DC Power Flow Revisited". In: *IEEE Transactions on Power Systems* 24.3 (Aug. 2009), pp. 1290–1300. ISSN: 0885-8950. DOI: 10.1109/TPWRS.2009.2021235.

[131] David Sturzenegger, Dimitrios Gyalistras, Manfred Morari, and Roy S Smith. "Model Predictive Climate Control of a Swiss Office Building : Implementation , Results , and Cost – Benefit Analysis". In: *IEEE Transactions on Control Systems Technology* PP.99 (2015), pp. 1–12. DOI: 10.1109/TCST.2015.2415411.

[132] Leo Su and Leslie K. Norford. "Demonstration of HVAC chiller control for power grid frequency regulation—Part 1: Controller development and experimental results". In: *Science and Technology for the Built Environment* 21.8 (2015), pp. 1134–1142. DOI: 10.1080/23744731.2015.1072449. eprint: https://doi.org/10.1080/23744731.2015.1072449. URL: https://doi.org/10.1080/23744731.2015.1072449.

[133] Leo Su and Leslie K Norford. "Demonstration of HVAC chiller control for power grid frequency regulation—Part 1: Controller development and experimental results". In: *Science and Technology for the Built Environment* 21.8 (2015), pp. 1134–1142.

[134] Jacopo Torriti, Mohamed G. Hassan, and Matthew Leach. "Demand response experience in Europe: Policies, programmes and implementation". In: *Energy* 35.4 (2010). Demand Response Resources: the US and International Experience, pp. 1575 –1583. ISSN: 0360-5442. DOI: `https://doi.org/10.1016/j.energy.2009.05.021`. URL: `http://www.sciencedirect.com/science/article/pii/S0360544209002060`.

[135] Vincenzo Trovato, Simon H Tindemans, and Goran Strbac. "Leaky storage model for optimal multi-service allocation of thermostatic loads". In: *IET Generation, Transmission & Distribution* 10.3 (2016), pp. 585–593.

[136] *US Department of Energy: Office of Energy Efficiency and Renewable Energy, Commercial Reference Buildings.* URL: `https://energy.gov/eere/buildings/commercial-reference-buildings`.

[137] E. Vrettos, K. Lai, F. Oldewurtel, and G. Andersson. "Predictive Control of buildings for Demand Response with dynamic day-ahead and real-time prices". In: *2013 European Control Conference (ECC)*. July 2013, pp. 2527–2534. DOI: `10.23919/ECC.2013.6669762`.

[138] Evangelos Vrettos and Goran Andersson. "Combined load frequency control and active distribution network management with thermostatically controlled loads". In: *Smart Grid Communications (SmartGridComm), 2013 IEEE International Conference on*. IEEE. 2013, pp. 247–252.

[139] Evangelos Vrettos, Emre C Kara, Jason MacDonald, Göran Andersson, and Duncan S Callaway. "Experimental Demonstration of Frequency Regulation by Commercial Buildings-Part I: Modeling and Hierarchical Control Design". In: *arXiv preprint arXiv:1605.05835* (2016).

[140] Pengfei Wang, Daniel H Liang, Jialiang Yi, Pádraig F Lyons, Peter J Davison, and Philip C Taylor. "Integrating electrical energy storage into coordinated voltage control schemes for distribution networks". In: *IEEE Transactions on Smart Grid* 5.2 (2014), pp. 1018–1032.

[141] Yu Wang, KT Tan, Xiao Yang Peng, and Ping Lam So. "Coordinated control of distributed energy-storage systems for voltage regulation in distribution networks". In: *IEEE Transactions on Power Delivery* 31.3 (2016), pp. 1132–1141.

[142] Bartholomäus Wasowicz, Simon Koopmann, Thomas Dederichs, Armin Schnettler, and Ulrich Spaetling. "Evaluating regulatory and market frameworks for energy storage deployment in electricity grids with high renewable energy penetration". In: *European Energy Market (EEM), 2012 9th International Conference on the*. IEEE. 2012, pp. 1–8.

[143] Di Wu, Chunlian Jin, Patrick Balducci, and Michael Kintner-Meyer. "An energy storage assessment: Using optimal control strategies to capture multiple services". In: *Power & Energy Society General Meeting, 2015 IEEE*. IEEE. 2015, pp. 1–5.

[144] Xiaomin Xi, Ramteen Sioshansi, and Vincenzo Marano. "A stochastic dynamic programming model for co-optimization of distributed energy storage". In: *Energy Systems* 5.3 (2014), pp. 475–505.

[145] F. Zhang. *The Schur Compliment and its Applications*. Springer, 2005.

[146] Xiaojing Zhang, G. Schildbach, D. Sturzenegger, and M. Morari. "Scenario-based MPC for energy-efficient building climate control under weather and occupancy uncertainty". In: *Control Conference (ECC), 2013 European.* Jan. 2013, pp. 1029–1034. URL: http://www. researchgate.net/publication/261025210_Scenario-based_MPC_for_energy-efficient_building_climate_control_under_weather_and_occupancy_uncertainty.

[147] Y. Zhang, J. Li, K. Meng, Z. Y. Dong, Z. Yu, and K. P. Wong. "Voltage regulation in distribution network using battery storage units via distributed optimization". In: *2016 IEEE International Conference on Power System Technology (POWERCON).* Sept. 2016, pp. 1–6. DOI: 10.1109/POWERCON.2016.7754056.

[148] Peng Zhao, Gregor P. Henze, Sandro Plamp, and Vincent J. Cushing. "Evaluation of commercial building HVAC systems as frequency regulation providers". In: *Energy and Buildings* 67 (2013), pp. 225 –235. ISSN: 0378-7788. DOI: https://doi.org/10.1016/j.enbuild.2013.08.031. URL: http://www.sciencedirect.com/science/article/pii/S0378778813005288.

[149] D. P. Zhou, Q. Hu, and C. J. Tomlin. "Quantitative comparison of data-driven and physics-based models for commercial building HVAC systems". In: *2017 American Control Conference (ACC).* May 2017, pp. 2900–2906. DOI: 10.23919/ACC.2017.7963391.

[150] R. D. Zimmerman, C. E. Murillo-Sanchez, and R. J. Thomas. "MATPOWER's extensible optimal power flow architecture". In: *2009 IEEE Power Energy Society General Meeting.* July 2009, pp. 1–7. DOI: 10.1109/PES.2009.5275967.

[151] Chrysovalantou Ziogou, Simira Papadopoulou, Michael C. Georgiadis, and Spyros Voutetakis. "On-line nonlinear model predictive control of a PEM fuel cell system". In: *Journal of Process Control* 23.4 (2013), pp. 483 –492. ISSN: 0959-1524. DOI: https://doi.org/10.1016/j.jprocont.2013.01.011. URL: http://www.sciencedirect.com/science/article/pii/S0959152413000218.

Av. des Collèges, 48b
1009, Pully, Switzerland
℘ +41 78 842 20 54
✉ fabietti.luca@gmail.com
↰ www.linkedin.com/in/luca-fabietti
Date of birth: 15.11.1989
Italian (Swiss B permit)

Luca Fabietti

Key skills

Control systems, optimization, model predictive control, data analysis and visualization, distributed energy resources, energy markets.

Education

2014–2018 **Ph.D. in Automatic Control**.
École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland
Thesis Title: Model-based predictive control methods for the provision of grid services with distributed energy resources

2011–2014 **Master in Automation Engineering**, *110/110 cum laude*.
Università degli Studi di Padova, Padova, Italy
Thesis title: Control of HVAC System via Implicit and Explicit MPC

2008–2011 **Bachelor in Engineering of Information Technology**, *110/110*.
Università degli Studi di Padova, Padova, Italy
Thesis title: Modeling, analysis, and control of an active suspension system

Experience

Jan. 2017 – **Co-founder**, *EnerSeer*, Switzerland.
April 2018 EnerSeer has created an innovative concept to boost the integration of storage systems to the power markets in order to maximize their revenues and optimize their usage.
   ○ Secured a proof-of concept seed investment (140 kCHF) and participation to MassChallenge accelerator program;
   ○ Wrote a patent protecting the core technology;
   ○ Gave more than 30+ presentations for potential utilities clients, investors, and expert panels;
   ○ Contributed to the market analysis and customer segmentation;
   ○ Worked on the creation of the business model.

Jan. 2018 – **Technical Consulting**, *CSEM (INDIGO Project)*, Switzerland.
March 2018 Worked in a team of four people to provide technical expertise on the development of the overall architecture and design of an advanced controller for a large-scale district cooling system.
   ○ Contributed to the design of a predictive controller at component scale to improve the district cooling efficiency;
   ○ Achieved a reduction of more than 50% for the computational burden of the control algorithm allowing its full deployment on the real system.

2014–2018 (expected)	**Doctoral Assistant**, *Automatic Control Laboratory, EPFL*, Switzerland.
	○ Developed data-driven based methods to automatically model, forecast, and control the dynamical behaviour of distributed energy resources such as commercial buildings, electro-chemical energy storage systems, and photovoltaic panels.
	○ Designed and developed a large scale experimental testbed to coordinate the operation of a smart building and of a Battery Electric Energy Systems. The platform features data retrieval, management (millions of datapoint), and visualization, real-time communication and control.
	○ Created ensemble model (XGBoost + NN) with state-of-the-art performance to forecast the aggregated load consumption in North of Italy.
2013–2014 (6 months)	**Research Engineer**, *KTH Royal Institute Of Technology*, Stockholm, Sweden.
	○ Contributed to the implementation and experimental validation of an advanced model-based predictive controller for an office building in KTH;

Languages

Italian	Native or bilingual proficiency	**French**	Full professional proficiency
English	Full professional proficiency	**German**	Limited working proficiency

Technical Skills

Mathematics
- ○ System ID, generalized linear models, gradient boosting, deep learning, random forests, Gaussian processes.
- ○ Model predictive control, optimal control, control theory, loop-shaping (PID).
- ○ Statistical signal processing, estimation, Kalman filtering, time series.
- ○ Optimization - convex, non-convex, uncertain (stochastic and robust optimization).
- ○ Optimization software (CPLEX, GUROBI, IPOPT, YALMIP, CasADi).

Programming
- ○ Languages: **Python (Tensor Flow, Keras, scikit-learn, Pandas), R, C++**.
- ○ Dynamical systems modeling and control: **Matlab, Simulink**.
- ○ Industrial protocols control, and monitoring: **Modbus, Labview**.
- ○ Data management and visualization tools: **SQL, Influxdb, Grafana, Indigo**.

Extra Curricular Activities

Sports: Beach volley, volleyball (currently playing in a local team), football (more than 10 years).
Music: Guitar (self-taught, played in a band for 2 years).
Seminar organization: Was the principal responsible for the seminar organization of the laboratory. Coordinated and managed the visit of 20+ visitors.
Others: Travelling, Cuban salsa.

Patents

L.Fabietti, T.T. Gorecki, Faran A. Qureshi, Ioannis Lymperopoulos, and C.N. Jones, "System and Method for Operating a Combination of Heterogeneous Energy Resources for fast grid Services". *European Patent, Submission Number: 6053730, 14 February 2018.*

Publications

L.Fabietti, E. Namor, P. Nahata, M. Paolone, and C.N. Jones, "Multi-service provision to the grid through coordination of energy-constrained resources". *In Applied Energy (submitted 2018).*

L.Fabietti, T.T. Gorecki, Faran A. Qureshi, C. Salzmann, and C.N. Jones, "Multi-time scale coordination of complementary resources for the provision of ancillary services". *Applied Energy, vol. 229, pp. 1164-1180, September 2018.*

L.Fabietti, T.T. Gorecki, E. Namor, F. Sossan, M. Paolone and C.N. Jones, "Enhancing the dispatchability of distribution networks through electric energy storage systems and flexible demand". *In Energy and Buildings, vol. 172, May 2018, pp. 125-138.*

L.Fabietti, T.T. Gorecki, E. Namor, F. Sossan, M. Paolone and C.N. Jones, "Dispatching active distribution networks through electrochemical storage systems and demand side management". *In Proceedings of 1st IEEE Conference on Control Technology and Applications, Big Island, Hawaii, USA, August 27-30, 2017.*

L.Fabietti, T.T. Gorecki, F.A. Qureshi, A. Bitlislioglu, I. Lymperopoulos and C.N. Jones, "Experimental Implementation of Frequency Regulation Services Using Commercial Buildings". *in IEEE Transactions on Smart Grid, vol. 9, no. 3, pp. 1657-1666, May 2018.*

L.Fabietti, and C.N. Jones, "Stochastic MPC for controlling the average constraint violation of periodic linear systems with additive disturbances". *In Proceedings 2016 American Control Conference (ACC), Boston, MA, 2016, pp. 5395-5400.*

T.T. Gorecki, **L. Fabietti**, F.A. Qureshi, C.N. Jones, "Experimental demonstration of buildings providing frequency regulation services in the Swiss market". *In Energy and Buildings, vol. 144, 1 June 2017, pp. 229-240.*

A. Parisio, **L. Fabietti**, M. Molinari, D. Varagnolo, K.H. Johansson, "Control of HVAC Systems via Scenario-based Explicit MPC". *53rd IEEE Conference on Decision and Control December 15-17, 2014, Los Angeles, California, USA.*

A. Parisio, D. Varagnolo, M. Molinari, G.Pattarello, **L. Fabietti**, K.H. Johansson, "Implementation of a scenario-based MPC for HVAC systems: an experimental case study". *In Energy and Buildings, vol. 144, 1 June 2017, pp. 229-240.*