

POLITECNICO
MILANO 1863

Department of Electronics, Information and Bioengineering
Doctoral Programme In Information Technology

Managing Uncertainty in Real-Time Optimisation and Control of Power Grids

Alessandro Del Duca
Doctoral Dissertation

HITACHI
Inspire the Next
Hitachi Energy

Advisor: Prof. Fredy Ruiz
Tutor: Prof. Fabio Dercole
Year 2025 - XXXVII Cycle



Finanziato dall'Unione europea
NextGenerationEU

To the one who is
my beacon in every storm



Abstract

The growing integration of Renewable Energy Sources (RESs) and Distributed Energy Resources (DERs) is reshaping the global power infrastructure landscape. This integration introduces a complex set of operational challenges: the inherent unpredictability and variability of renewable energy generation, combined with communication uncertainties and the presence of heterogeneous participants, make optimal power system management increasingly critical.

This thesis addresses these challenges by developing scalable and robust energy management methodologies tailored to power systems composed of a large number of medium-to-low-power resources, including Battery Energy Storage Systems (BESSs), small-scale solar generators, and prosumers participating in the energy market through Demand Response (DR).

At its core, this research establishes a theoretical framework for managing uncertainty in distributed multi-agent energy systems. To this end, it integrates scenario-based techniques, distributed and decentralised optimisation, Set Membership Estimation (SME), and adaptive Model Predictive Control (MPC) strategies, enabling robust system operation, mitigating performance degradation, and supporting risk-aware decision-making.

The proposed framework supports the coordination of Microgrids (MGs), Virtual Power Plants (VPPs), and aggregations of Flexible Loads (FLs), ensuring the optimal utilisation of available resources and compliance with operational constraints under uncertainty, while emphasising scalability, robustness, and safety.

In summary, this thesis provides advanced tools for designing resilient, adaptive, and efficient Energy Management System (EMS) capable of managing uncertainties in power generation, communication, and model mismatches. These tools support the integration of RESs and DERs, facilitating the transition towards low-carbon smart grids.

Keywords: Uncertain Systems; Distributed Optimisation; Adaptive Control; Active Learning; Microgrids; Virtual Power Plants; Demand Response.

Sommario

L'integrazione crescente di Fonti Energetiche Rinnovabili (RESs) e Distribuite (DERs) sta trasformando radicalmente il panorama delle infrastrutture elettriche a livello globale. Tale integrazione introduce complesse sfide operative: l'imprevedibilità e la variabilità intrinseche della generazione rinnovabile, le possibili incertezze di comunicazione e la presenza di partecipanti eterogenei rendono cruciale la corretta gestione della rete elettrica.

Questa tesi affronta tali sfide sviluppando metodologie di gestione energetica scalabili e robuste, pensate per reti elettriche composte da un numero elevato di risorse di media e bassa potenza, tra cui Sistemi di Accumulo a Batteria (BESSs), generatori eolici e fotovoltaici di piccola scala, e prosumer che partecipano al mercato energetico attraverso programmi di Demand Response (DR).

Il fulcro di questa ricerca è la definizione di un quadro teorico per la gestione dell'incertezza in ambienti di ottimizzazione multiagente e distribuita. A tale scopo, vengono integrate tecniche basate su scenari, ottimizzazione distribuita e decentralizzata, Set Membership Estimation (SME) e strategie di Controllo Predittivo (MPC). Questi approcci consentono al sistema di migliorare le prestazioni mantenendo un funzionamento risk-aware.

Le metodologie proposte trovano applicazione nel coordinamento di Microreti (MGs), delle Centrali Elettriche Virtuali (VPPs) e dell'aggregazione di carichi flessibili (FLs). Esse garantiscono l'uso ottimale delle risorse disponibili e il rispetto dei vincoli operativi anche in presenza di incertezze, enfatizzando al contempo scalabilità, robustezza e sicurezza.

In sintesi, questa tesi propone strumenti avanzati per la progettazione di Sistemi di Gestione dell'Energia (EMS) resilienti, adattivi ed efficienti, capaci di affrontare le incertezze nella produzione, nella comunicazione e nei modelli di sistema. Tali sistemi non solo favoriscono l'integrazione di RESs e DERs, ma contribuiscono anche alla transizione verso reti elettriche intelligenti e sostenibili.

Parole chiave: Sistemi Incerti; Ottimizzazione Distribuita; Controllo Adattativo; Apprendimento Attivo; Microreti; Centrali Elettriche Virtuali; Demand Response.

Contents

Abstract **i**

Sommario **iii**

Contents **v**

1 Introduction **1**

1.1 The Energy Sector in Climate-Change Mitigation and Resilience 1

1.2 The Smart Grid 3

1.2.1 The multi-agent nature of the Smart Grid 4

1.2.2 The Role of Uncertainty in Energy Management 6

Research question and thesis outline **11**

2 Scalable Scenario Approach in distributed multi-agent settings **15**

2.1 Scenario approach preliminaries 17

2.2 Scenario Approach in a multi-agent setting 20

2.2.1 Problem Statement 21

2.2.2 Structured uncertainty space 25

2.2.3 Structured decision space 26

3 Active Learning and Dual MPC **31**

3.1 Set-Membership Estimation Theory preliminaries 33

3.1.1 SME problem formulation for ARX models 33

3.2 Fast SME-MPC with Probabilistic Guarantees 38

3.2.1 Infinite-Horizon Optimal Control Problem 38

3.2.2 Scenario-MPC 39

3.2.3 Uncertainty model adaptation through bounding box 40

3.3 Set-Membership Dual MPC 42

3.3.1 Problem formulation 42

3.3.2 Exploratory-term in the cost function . . . 43
3.3.3 Multiobjective MPC . . . 45
3.4 SME Dual MPC theoretical guarantees . . . 45
3.4.1 FPS convergence to a singleton Set . . . 45
3.4.2 Volume shrinking and certainty equivalence . . . 48
3.5 SME Dual MPC Practical Implementation . . . 49
3.5.1 Forgetting Factor . . . 50
3.5.2 Volume minimisation and Warm-start . . . 51
3.5.3 Terminal FPS . . . 52
3.5.4 Arx Assumption and structural model mismatch . . . 54
3.5.5 MPC algorithm . . . 54
3.6 Numerical example . . . 55
3.6.1 Stability Region in Parameter Space . . . 56
3.6.2 Set-Based Representation: Strips and Feasible Parameter Set (FPS) 57

4 Smart Grid Energy Management 61
4.1 European Energy Market . . . 61
4.1.1 Day-Ahead Market (DAM) . . . 63
4.1.2 Intraday Market (IDM) . . . 63
4.1.3 Balancing Market (BM) . . . 64
4.1.4 Ancillary Services Market (ASM) . . . 64
4.2 Hierarchical Architecture for Frequency Control and Restoration . . . 64
4.3 Energy Management Systems . . . 66
4.3.1 Architectures for EMSs . . . 67
4.3.2 Microgrid Energy Management preliminaries . . . 69

5 Case Study: Hierarchical Energy Management System 77
5.1 Contributions . . . 77
5.2 Hierarchical Architecture . . . 78
5.2.1 Supervisory Model Predictive Control . . . 79
5.2.2 Low-level controller . . . 83
5.3 Simulation case study . . . 86
5.3.1 Microgrid Parameters and Simulation . . . 87
5.3.2 Economic optimisation . . . 87
5.3.3 Low-level balancing . . . 90

6 Case-study: Distributed VPP architecture through scenario approach 93
6.1 Contributions . . . 93

6.2 Virtual Power Plant modelling . . . 94
6.3 VPP architecture . . . 98
6.3.1 Centralised optimisation . . . 98
6.3.2 Distributed Algorthm . . . 100
6.4 Simulation results . . . 103
6.4.1 Algorithm for the randomised generation of VPPs . . . 103
6.4.2 Aggregated power production probabilistic guarantees . . . 104
6.4.3 Convergence of the Distributed Algorithm . . . 106

7 Dual SME MPC for Aggregated TER Demand Responsee 109
7.1 Contributions . . . 110
7.2 Problem Definition . . . 111
7.2.1 Thermo-Electric (TEC) module heat flows . . . 111
7.2.2 Thermal dynamic equations . . . 112
7.2.3 Aggregated dynamics and problem statement . . . 112
7.3 Simulations . . . 114
7.3.1 Initial Feasible Parameter Set . . . 115
7.3.2 Simulation setting . . . 116

8 Conclusions 123
8.1 Scientific contributions . . . 123
8.2 Future developments . . . 125
8.3 Final considerations . . . 126

List of Figures 127

List of Tables 129

Acronyms 131

Acknowledgements 135

1 | Introduction

1.1. The Energy Sector in Climate-Change Mitigation and Resilience

Over the past few decades, a substantial body of scientific evidence from both national [NAS23] and international agencies [Mas+21; Uni19; Foo22; Wor23] has demonstrated that anthropogenic activities have increased the concentration of GreenHouse Gases (GHGs) in the atmosphere with effects that directly correlate, with a high degree of confidence, with rises in global temperatures and related environmental impacts. It is widely discussed in the scientific literature and in climate change research that these environmental hazards pose a societal and global crisis that can be overcome only through international cooperation.

The primary international cooperation project that gives life to the International Panel for Climate Change (IPCC) periodically produces reports to inform policymakers on climate change and mitigation strategies. In the 6th IPCC assessment report on climate change: Mitigation of Climate Change (group 3) [Shu+22], the authors identify the main classes of human activity that are the primary contributors to GHGs production. Figure 1.1 shows the contribution of different sectors in the global GHGs production, specifically the combustion of fossil fuels in transportation (15%) and energy production (34%), the large-scale use of land for Agriculture, Forestry, and Other Land Use (AFOLU) and livestock farming (22%), and Industrial (24%) and Building (6%) Activities.

As shown in Fig. 1.1, the energy sector contributes to more than one-third of the total amount of anthropogenic GHGs emissions, primarily through the combustion of coal, oil, and natural gas for electricity production and heating. At the same time, it is also the sector with the highest potential for emission reductions, given the availability of both mature and emerging technologies that enable decarbonization. In fact, as shown in Fig. 1.2, Renewable Energy Sources (RESs) and Distributed Energy Resources (DERs), such as Photovoltaic (PV), wind energy, and Battery Energy Storage Systems (BESSs), have experienced both a decrease in costs due to technological advancements and an

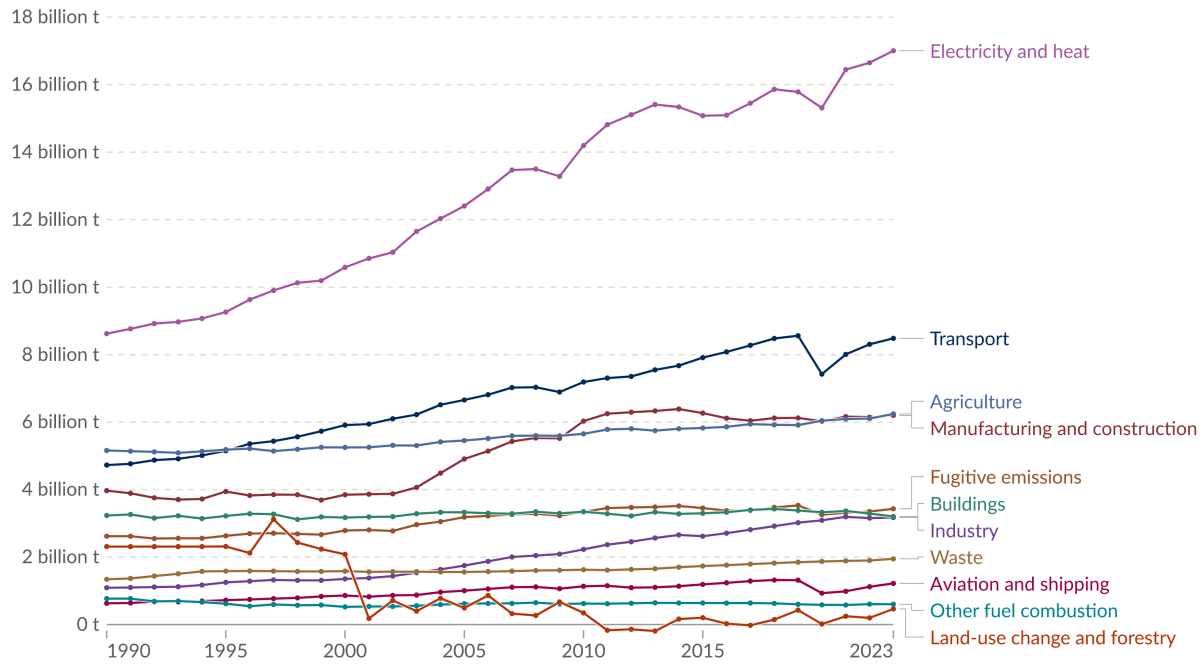


Figure 1.1: Total global GHGs emission by sector. Data source [Our26]

increase in usage thanks to higher profitability and supportive policy-making.

According to the IPCC, scenarios that aim to limit global warming to 1.5°C require renewable energy to account for 70% to 85% of electricity generation by 2050 [Shu+22]. Increasing the share of RESs for power generation is crucial for limiting the energy sector's carbon footprint. Wind, solar photovoltaics, and hydropower have already been cost-competitive with fossil fuels in many markets, and their deployment has accelerated significantly over the past years [REN16].

In addition, the large-scale integration of renewable energy has become a matter of national energy security, as it reduces reliance on fossil fuel markets that can be unstable due to conflicts and geopolitical tensions. This was evident during the European crisis in the winter of 2022, following Russia's embargo on natural gas after tensions on Ukrainian territory [Int23b], which caused significant disruptions to Europe's power generation and heating systems.

However, the shift toward RESs alone is not sufficient. While they help reduce GHGs emissions, they also pose challenges during deployment. These challenges include increased uncertainty due to their non-dispatchable nature, reduced system inertia, and a decline in overall flexibility [Shu+22; REN16].

Enhancing efficiency across power systems by improving grid design and modernising infrastructure not only provides a framework to address the challenges posed by RESs

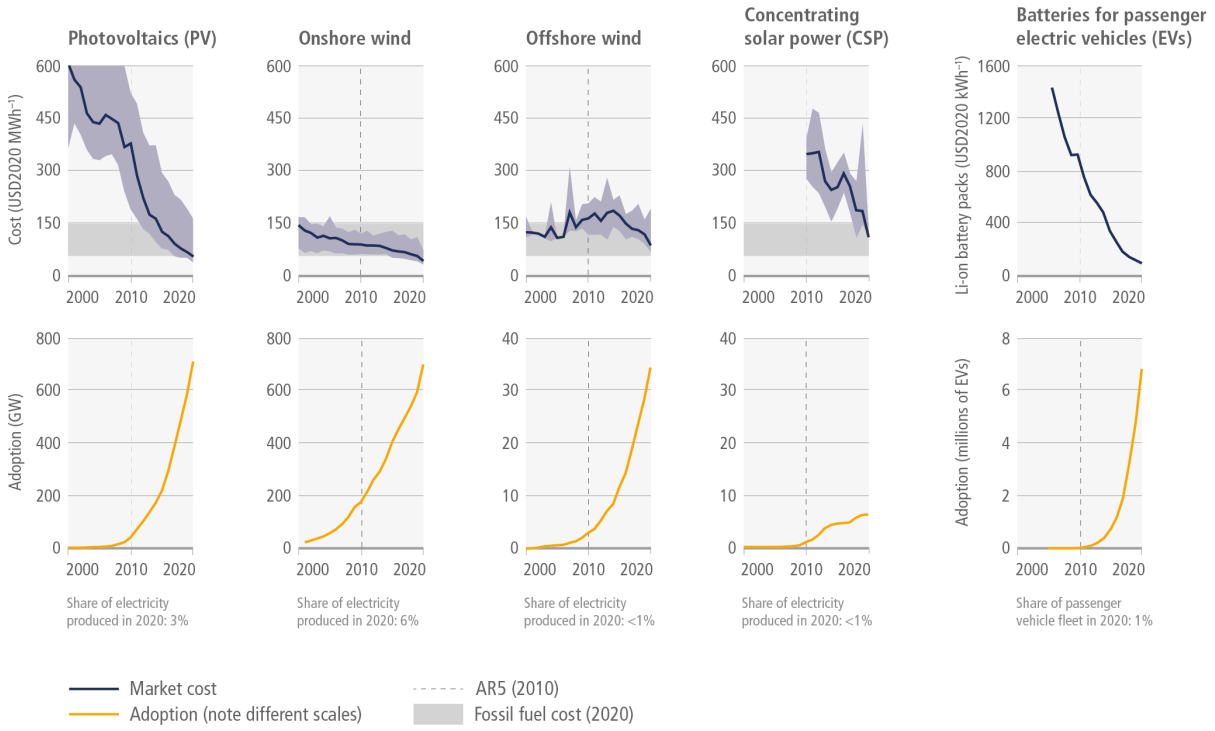


Figure 1.2: RES and DER cost decrease and increase adoption [Shu+22].

but also increases their penetration in the energy mix. Additionally, this approach is a fast and cost-effective way to mitigate GHGs emissions caused by electricity losses in transmission and distribution, which account for approximately 15% of total losses [Int23a]. On the demand side, implementing efficient appliance and device standards, along with improving demand-side flexibility, can enable demand-response programs and consumption pattern adjustments. These measures could further enhance the system's overall efficiency.

1.2. The Smart Grid

The urgency in climate change adaptation and mitigation has driven a fast-paced, transformative evolution of the global energy sector. The traditional power infrastructure, historically designed as a monolithic, unidirectional network with consumers unable to participate in the energy market, is proving ineffective in addressing the complexities and variations of modern energy demands; the traditional power infrastructure is giving way to a more efficient digital infrastructure, known as the Smart Grid (SG), revolutionising the way power is generated, distributed, and consumed [Fan+12; Mom12; IEE11].

The SG has emerged as a revolutionary framework that can integrate advanced optimisation, control, communication, and sensing technologies to enhance the reliability, effi-

ciency, and sustainability of power infrastructure. Unlike conventional power grids, which depend on centralised decision-making, SGs can integrate decentralised intelligence, allowing for two-way flows, both in electricity and information, between producers, consumers, and system operators [PS16; MM18].

In this new infrastructure paradigm, users are no longer passively relying on the power grid; with the integration of DERs and Demand Response (DR) programs, they can now adjust their energy consumption according to availability. This shift enables them to become "prosumers," a term that combines the roles of both consumer and producer, by participating in the energy market [PS16; MM18; Str08]. User market participation not only enhances consumer engagement but also contributes to system resilience by enabling consumers to provide energy back to the grid during periods of peak demand.

Key enabling technologies for SG encompass Advanced Metering Infrastructure (AMI), Supervisory Control and Data Acquisition (SCADA) systems, and Internet of Energy (IoE) devices. Advanced optimisation and control technologies, Artificial Intelligence (AI), and AI-powered analytics are fundamental tools for real-time monitoring and management of the grid, enabling predictive maintenance, regenerative capabilities, and advanced adaptation to respond dynamically to power fluctuations in supply and demand [Lun+15; Fan+12].

The advantages of SG extend beyond mere operational and technical enhancement. By facilitating market participation and encouraging demand-side engagement, SG paves the way for innovative, fairer, and more collaborative business models, such as Energy Communities (ECs) and Virtual Power Plants (VPPs). These new frameworks empower energy consumers to become energy citizens, leading to greater community engagement and participation in sustainable energy practices. Moreover, SGs are crucial in supporting the transition to a low-carbon economy, as they enable the reliable integration of high shares of RESs without compromising power grid stability and security [MM18; PS16].

In Fig. 1.3, the Smart Grid architecture is depicted.

1.2.1. The multi-agent nature of the Smart Grid

The transition toward Smart Infrastructure Systems (Smart Grids) is a dominant driver of societal transformation in the coming decades [AJP23]. This challenge requires optimising systems that share some crucial features:

Multi-agent: The system is the aggregation of different interacting "agents": decision-making entities that have partial information about the overall system and other agents.

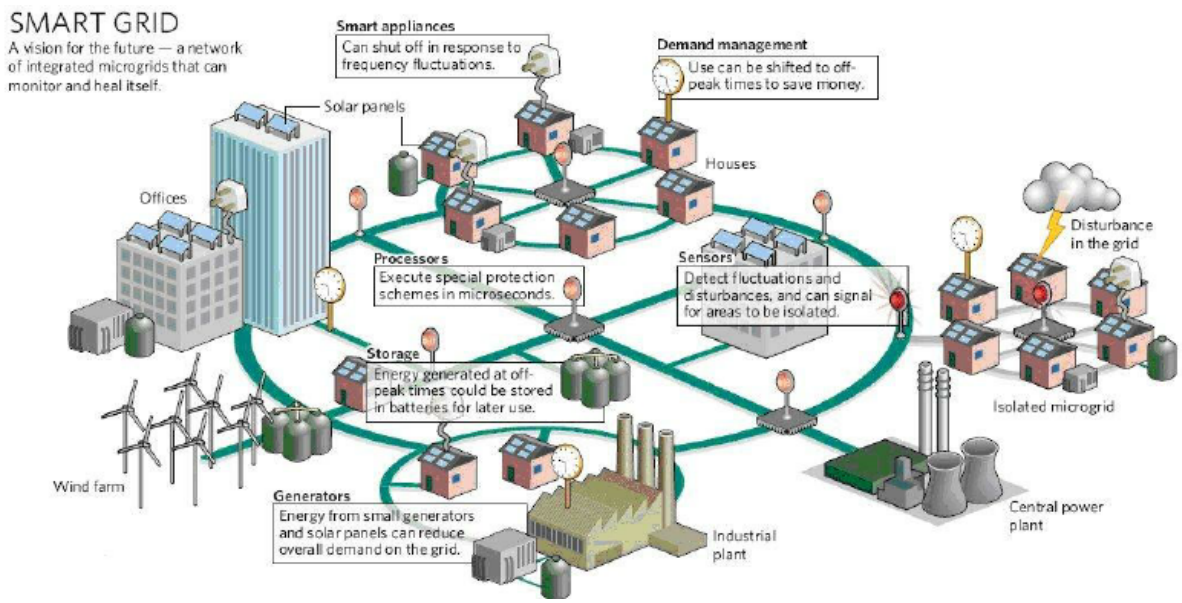


Figure 1.3: Smart Grid architecture [Wol12].

Each of the latter can only act on individual decisions and decision variables, but can account for the overall network through consensus mechanisms and communication with the neighbours.

Welfare maximisation: The overall system's goal, typically controlled by an aggregating operator, is to steer the agents' decisions toward an operating point that maximises participants' benefit (network welfare).

Large-scale: The high number of agents reflects in a high number of possible decisions which influence one another. This makes the optimisation of such systems remarkably complex.

Uncertainty handling: Finally, agents are affected by internal and network-wide uncertainty and have partial information on how their decisions affect the global system, making the decision-making process complex. Moreover, this complexity in handling uncertainty is coupled with the network size.

The renewed interest in distributed and multi-agent optimisation has become evident in recent years, as this framework effectively addresses these features. Significant contributions appear in the literature. The work by [Boy+11] is a landmark study in this area, while [NO09] introduce a distributed sub-gradient method. Additionally, [Mar+16] and [FP21] present algorithms based on the proximal dual decomposition of optimisation problems. These methods are particularly beneficial for solving non-smooth convex problems by distributing computational tasks among agents across potentially time-varying

communication networks.

1.2.2. The Role of Uncertainty in Energy Management

While the conceptual role of an SG is clear, its practical implementation is significantly complicated by pervasive uncertainties in the system [Kha24; Sha+25]. Additionally, factors such as communication requirements, scalability, and energy security become critical [Yu+25].

Modern electrical grids integrate an increasing and unprecedented share of RESs, which, although essential for decarbonisation of the energy sector, are inherently stochastic: the production heavily depends on weather conditions such as wind speed for wind turbines, radiation and cloud coverage for PV panels [Kha24; Sha+25]. Such variability and unpredictability introduce challenges in balancing supply and demand at both short- and long-term time scales [Gao+15].

Uncertainties also extend to load consumption, another primary source of unpredictability. The growth of heterogeneous loads, such as Electric Vehicles (EVs), Thermo-Electric Refrigerators (TERs),Heat Pumps (HPs), and innovative Heating, Ventilation, and Air Conditioning (HVAC) systems, adds significant variability to consumer profiles, making forecasting difficult due to intrinsic correlations with uncertainties in weather forecasts [Sha+25].

In addition, allowing prosumers to participate in DR programs introduces a new source of uncertainty in system dynamics [Gao+15]. When the system operator requests a DR service, the utility control system must modulate the appropriate amount of power while adhering to technical and operational constraints. One example of this uncertainty occurs with BESSs, where the amount of power that can be delivered and the control system's response time are unpredictable for the system operator [Kha24]. Another example involves TERs, whose aggregated response relies on various uncertain factors, such as external temperatures, control logic, and physical component variability [Gao+15].

Uncertainty is also present in the communication infrastructure underpinning the SG. Effective and timely information exchange is essential, but in practice, communication networks are prone to delays, packet dropouts, asynchronous updates, and cyberattacks such as Denial of Service (DoS) attacks [Yu+25; Fus+21]. Such non-idealities can significantly degrade performance in distributed decision-making, and in worst cases, destabilise control loops. For example, in consensus-based schemes for power grid coordination, even moderate communication failures can slow convergence and lead to suboptimal operation [Yu+25].

Real-world systems are also affected by measurement noise and sensor inaccuracies. Imperfect state estimation can propagate errors across the control hierarchy, leading to suboptimal or unsafe decisions [Fus+21]. Limited computational resources further constrain algorithmic complexity, forcing trade-offs between optimality, scalability, and robustness [Fus+21].

Given this spectrum of uncertainties, designing SGs requires not only optimising operations under ideal conditions but also ensuring robustness and resilience [Kha24; Sha+25]. In particular, three dimensions of robustness must be considered:

- **Operational robustness**: ability to guarantee secure operation despite stochastic variations in renewable production, load demand, and DERs dynamics.
- **Communicational robustness**: tolerance to packet loss, delays, and network failures, ensuring effective coordination among prosumers.
- **Computational robustness**: ensuring efficiency and reliability of optimisation and control algorithms under limited processing power and time constraints.

This thesis aims to explore all the latter dimensions of robustness, while also considering practical implementation aspects, real data, and theoretical guarantees under uncertainty. In the following section, the research questions are presented, and an overview of the thesis is provided.

Chapter 1 Bibliography

[AJP23] Anuradha M. Annaswamy, Karl H. Johansson, and George J. Pappas. "Control for Societal-scale Challenges: Road Map 2030". In: *IEEE Control Systems Society Publication* (2023). URL: https://ieeecss.org/control-societal-scale-challenges-roadmap-2030.

[Boy+11] Stephen Boyd et al. "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers". In: *Foundations and Trends® in Machine Learning* 3.1 (2011), pp. 1–122. ISSN: 1935-8237. DOI: 10.1561/2200000016.

[Fan+12] Xi Fang et al. "Smart Grid — The New and Improved Power Grid: A Survey". In: *IEEE Communications Surveys and Tutorials* 14.4 (2012), pp. 944–980. DOI: 10.1109/SURV.2011.101911.00087.

[Foo22] Food and Agriculture Organization of the United Nations. *The State of Food Security and Nutrition in the World 2022*. Tech. rep. Rome: Food and Agriculture Organization (FAO), 2022. DOI: 10.4060/cc0639en. URL: https://www.fao.org/publications/sofi/2022/en/.

[FP21] Alessandro Falsone and Maria Prandini. "A Distributed Dual Proximal Minimization Algorithm for Constraint-Coupled Optimization Problems". In: *IEEE Control Systems Letters* 5.1 (2021), pp. 259–264. DOI: 10.1109/LCSYS.2020.3001427.

[Fus+21] G Fusco et al. "On the use of robust techniques in smart grid control". In: *Energies* 15.1 (2021), p. 7. DOI: 10.3390/en15010007.

[Gao+15] D Gao et al. "A robust demand response control of commercial buildings for smart grid under load prediction uncertainty". In: *Energy* 82 (2015), pp. 1–12. DOI: 10.1016/j.energy.2015.01.076.

[IEE11] IEEE Smart Grid Initiative. *IEEE Smart Grid: Vision for the Future*. Tech. rep. Accessed: 2025-10-03. IEEE, 2011. URL: https://smartgrid.ieee.org/resources/reports/ieee-smart-grid-vision.

[Int23a] International Energy Agency. *World Energy Outlook 2023*. Tech. rep. Paris: International Energy Agency (IEA), 2023. URL: https://www.iea.org/reports/world-energy-outlook-2023.

[Int23b] International Energy Agency (IEA). *IEA warns of gas market tightness in 2023 as Russia further reduces supplies to Europe*. https://www.aa.com.tr/en/economy/iea-warns-of-gas-market-tightness-in-2023-as-russia-further-reduces-supplies-to-europe/2701179?utm_source=chatgpt.com. Accessed: 2025-10-03. 2023.

[Kha24] M Khalid. "Smart grids and renewable energy systems: Perspectives and challenges". In: *Renewable and Sustainable Energy Reviews* 154 (2024), p. 111819. DOI: 10.1016/j.rser.2021.111819.

[Lun+15] Henrik Lund et al. "Smart Energy and Smart Energy Systems". In: *Energy* 86 (2015), pp. 1–10. DOI: 10.1016/j.energy.2015.02.067.

[Mar+16] Kostas Margellos et al. "Proximal minimization based distributed convex optimization". In: *2016 American Control Conference (ACC)*. 2016, pp. 2466–2471. DOI: 10.1109/ACC.2016.7525287.

[Mas+21] V. Masson-Delmotte et al., eds. *Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, 2021. DOI: 10.1017/9781009157896. URL: https://www.ipcc.ch/report/ar6/wg1/.

[MM18] Thomas Morstyn and Malcolm D. McCulloch. "Distributed Energy Resources and Virtual Power Plants: A Review of Enabling Technologies and Market Participation". In: *Renewable and Sustainable Energy Reviews* 92 (2018), pp. 605–618. DOI: 10.1016/j.rser.2018.04.046.

[Mom12] James A. Momoh. *Smart Grid: Fundamentals of Design and Analysis*. Hoboken, NJ, USA: John Wiley and Sons, 2012.

[NAS23] NASA Global Climate Change. *Global Temperature*. https://climate.nasa.gov/vital-signs/global-temperature/. Accessed: 2025-10-03. 2023.

[NO09] Angelia Nedic and Asuman Ozdaglar. "Distributed Subgradient Methods for Multi-Agent Optimization". In: *IEEE Transactions on Automatic Control* 54.1 (2009), pp. 48–61. DOI: 10.1109/TAC.2008.2009515.

[Our26] Our World in Data. *Our World in Data*. 2026. URL: https://ourworldindata.org/.

[PS16] Yael Parag and Benjamin K. Sovacool. "Electricity Market Design for the Prosumer Era". In: *Nature Energy* 1 (2016), p. 16032. DOI: 10.1038/nenergy.2016.32.

[REN16] REN21. *Renewables 2016 Global Status Report*. Tech. rep. Accessed: 2025-10-03. Paris, France: REN21 Secretariat, 2016. URL: https://www.ren21.net/reports/global-status-report/.

[Sha+25] H Shayeghi et al. "Uncertainty aware energy management in microgrids with electric bicycle and electric vehicle charging stations". In: *Scientific Reports* 15.1 (2025), p. 12328. DOI: 10.1038/s41598-025-12328-5.

[Shu+22] P. R. Shukla et al., eds. *Climate Change 2022: Mitigation of Climate Change. Contribution of Working Group III to the Sixth Assessment Report of the*

Intergovernmental Panel on Climate Change. Cambridge, UK and New York, NY, USA: Cambridge University Press, 2022. DOI: 10.1017/9781009157926. URL: https://www.ipcc.ch/report/ar6/wg3/.

[Str08] Goran Strbac. "Demand Side Management: Benefits and Challenges". In: *Energy Policy* 36.12 (2008), pp. 4419–4426. DOI: 10.1016/j.enpol.2008.09.030.

[Uni19] United Nations Environment Programme. *Emissions Gap Report 2019*. Tech. rep. Nairobi: United Nations Environment Programme (UNEP), 2019. URL: https://www.unep.org/resources/emissions-gap-report-2019.

[Wol12] Maarten Wolsink. "The research agenda on social acceptance of distributed generation in smart grids: Renewable as common pool resources". In: *Renewable and Sustainable Energy Reviews - RENEW SUSTAIN ENERGY REV* 16 (Jan. 2012). DOI: 10.1016/j.rser.2011.09.006.

[Wor23] World Meteorological Organization. *State of the Global Climate 2022*. Tech. rep. Geneva, Switzerland: World Meteorological Organization (WMO), 2023. DOI: 10.35614/2218-6050. URL: https://public.wmo.int/en/our-mandate/climate/wmo-statement-state-of-global-climate.

[Yu+25] Y Yu et al. "Robust frequency estimation in smart grids with communication network uncertainties". In: *Scientific Reports* 15.1 (2025), p. 16533. DOI: 10.1038/s41598-025-16533-0.

Research question and thesis outline

The thesis aims to investigate the following research question:

How can distributed energy management methods be designed for systems with large populations of medium-to-low power resources, operating under communication delays and structural uncertainties, to ensure optimal utilisation of assets while maintaining scalability, robustness, and adaptability?

The developed methods ensure optimal utilisation of available assets, including BESSs, flexible loads (FL), and RESs, while handling uncertainty and structural variations within the system.

The thesis is organised into two main parts.

The first part introduces the mathematical foundations for robustifying the energy management problem in the presence of system uncertainties. It also presents the two main theoretical contributions:

Chapter 2: Scenario approach for distributed networks with global uncertain coupling constraints. This chapter presents the main result from [DRS24b]. By leveraging the scenario approach to approximate the coupling chance-constraint, it is proven that for distributed multi-agent optimisation with global uncertain coupling constraints, the number of scenarios needed to guarantee a violation probability is independent of the number of agents—thus ensuring scalability.

Chapter 3: Set Membership Estimation based Dual Model Predictive Control. This chapter leverages a polytopic parametrisation of model uncertainty using Set Membership Estimation (SME) to design an adaptive Model Predictive Control (MPC), where robustness is imposed through a scenario formulation of risk-constrained stochastic optimisation [Mez+25]. The algorithm is then modified to include an active learning criterion in the cost, quantified by the reduction in the volume of the Feasible Parameter Set (FPS) along the prediction horizon.

The second part applies the developed theory to distributed energy management problems.

Chapter 4: Energy Management Problem. This chapter introduces the general formulation of energy management; the resulting models, and the results from Chapter 2 are used to develop hierarchical, decentralised, and distributed decision-making algorithms that compensate for variability in agent characteristics by exploiting historical data and quantifying uncertainty.

Chapter 5: Robust Micro-grid Energy Management System through a Scenario Approach. This chapter presents a two-level hierarchical Energy Management System (EMS) architecture with layers operating at different time scales. The supervisory MPC optimises micro-grid costs, while the low-level controller handles uncertainties to reduce unplanned grid interventions. Computational complexity is kept low by adopting the scenario approach [DRS24a].

Chapter 6: Distributed Virtual Power Plant Optimisation via Scenario Approach. This chapter presents a scalable distributed stochastic optimisation framework for VPP management. The method is tested across various network sizes and real data. Results show that the probability of constraint violation at minimum service level remains constant as the number of agents grows—even with a fixed number of scenarios [DR25]. The proposed frameworks are scalable and robust to variations in network topology, communication delays, and asset availability.

Chapter 7: Thermo-Electric Loads aggregated control with exploration exploitation tradeoff. This chapter presents a case study on SME-based Dual MPC for aggregated control of TERs. The method estimates uncertain dynamics and adapts control actions while leveraging historical operational data for system identification.

List of Publications

[DR25] A. Del Duca and F. Ruiz. "Distributed Virtual Power Plant Optimization via Scenario Approach". In: *2025 IFAC Workshop on Smart Energy Systems for Efficient and Sustainable Smart Grids and Smart Cities (SENSYS)*. Bari, Italy, 2025.

[DRS24a] A. Del Duca, F. Ruiz, and R. Scattolini. "Robust Micro-Grid Energy Management System Through a Scenario Approach". In: *2024 American Control Conference (ACC)*. Toronto, ON, Canada, 2024, pp. 1801–1806. DOI: 10.23919/ACC60939.2024.10644278. URL: https://hdl.handle.net/11311/1273510.

[DRS24b] A. D. Del Duca, F. Ruiz, and R. Scattolini. "Distributed Stochastic Optimisation with Uncertain Coupling Constraints". In: *2024 IEEE 63rd Conference on Decision and Control (CDC)*. Milan, Italy, 2024, pp. 4431–4436. DOI: 10.1109/CDC56724.2024.10886359. URL: https://hdl.handle.net/11311/1286946.

[Mez+25] Erick Meza et al. "Adaptive Scenario-Based Predictive Control: A Set Membership Learning Approach". In: *IFAC-PapersOnLine* 59.26 (2025). 7th IFAC Conference on Intelligent Control and Automation Sciences ICONS 2025, pp. 323–328. ISSN: 2405-8963. DOI: https://doi.org/10.1016/j.ifacol.2025.12.055. URL: https://www.sciencedirect.com/science/article/pii/S2405896325027314.

2 | Scalable Scenario Approach in distributed multi-agent settings

In engineering and decision-making processes, it is often necessary for decisions to remain valid across a wide range of potential uncertainties. A key class of these problems is addressed using the well-known convex optimisation framework. The objective is to identify the values of decision variables x that minimise a convex cost function $f(x)$ while adhering to convex constraints, which are represented by the inclusion of the decision variables in a convex set X.

To accommodate potential uncertainties in the optimisation, both f and X may depend on an unknown or partially known variable that belongs to an uncertain set $\delta \in \Delta$. The formal representation of the optimisation problem can be expressed as

$$x^\star \in \arg \min_{x \in X(\delta)} f(x, \delta). \tag{2.1}$$

A potential solution to this problem is to consider the complete uncertain set Δ and solve the optimisation problem in a manner that ensures the solution remains valid across the entire set Δ. This strategy is known as Robust Optimisation (RO) [BN98; Ben+04; BS04]. When the uncertain set can be estimated, we can apply RO techniques to effectively solve the problem, regardless of the specific value of δ. Because these methods provide guarantees on optimality throughout the entire uncertain set, they are often referred to as "worst-case" methods.

While these methods are practical, they do come with some limitations:

- The RO framework does not fully utilise available data and statistical knowledge regarding δ. In fact, δ often has inherently stochastic behaviour that could be exploited.
- The framework relies on worst-case scenarios and overapproximations of the uncertain set Δ, leading to overly conservative solutions. Additionally, it requires

additional effort to include mechanisms for balancing robustness and performance.

- The RO framework may be demanding from a computational point of view

To address these issues while facilitating the use of statistical knowledge regarding δ, we can utilise the concept of chance-constrained optimisation. The optimisation problem can be expressed as

$$x^\star \in \text{argmin } f(x,\delta) \qquad \text{s.t.} \tag{2.2a}$$

$$\mathbb{P}[\delta \in \Delta : x^\star \in X(\delta)] < 1-\epsilon. \tag{2.2b}$$

This optimisation is known as chance-constrained optimisation program [Küç22; JG12; GL19]. In such problems, we aim to find the value of the decision variable x which minimises a convex cost function $f(x,\delta)$. However, we enhance the performance of the optimal solution by allowing for potential constraint violations under very unlikely realisations of δ, which occur with a low probability denoted by ϵ, also referred to as the risk parameter. As depicted in Fig. 2.1, by tolerating the violation of a percentage ϵ of constraints, induced by unlikely uncertainty realisation, we can improve the cost function $f(x^\star(\epsilon))$. By adjusting the value of ϵ, we can fine-tune the trade-off between performance

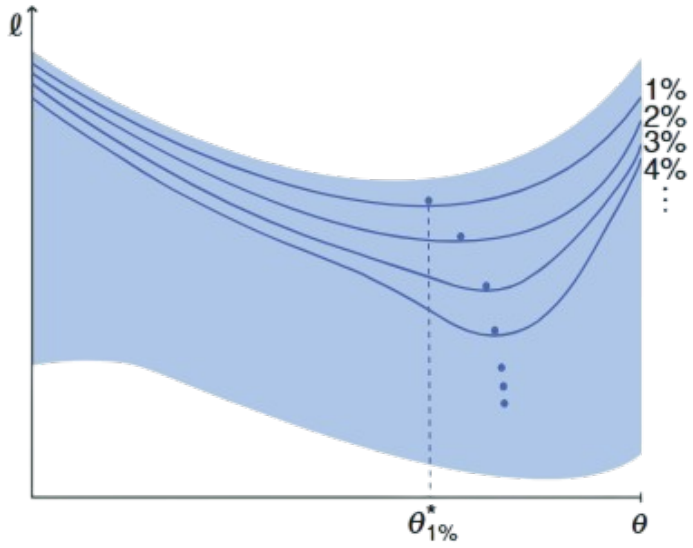


Figure 2.1: Optiman value $x^\star(\epsilon)$ for increasing level of ϵ [GC13].

and robustness. Specifically, as $\epsilon \to 1$, we revert to robust optimisation (RO), while decreasing ϵ can lead to improved optimal performance $f(x^\star)$, at the cost of constraint violations affecting a larger proportion of the uncertainty mass.

It is important to note that not all domains can tolerate constraint violations, as these can

have severe consequences. However, many problems can withstand temporary violations if these are associated with improbable events. Particularly in power systems, where the power balance can always be recovered through the use of power reserves from power plants and battery storage.

Chance-constrained optimisation is a valuable method for modelling decision-making under uncertainty. However, it can be challenging to implement in practice, especially when the uncertainty variable, denoted as δ, lacks identifiable structural properties, such as a precise distribution [Küç22].

One approach to solving instances of convex stochastic programming is the scenario approach proposed by various authors such as [CG08; ATC09; CGP19; CC06; TDC13]. In this method, the uncertain set Δ is approximated by sampling N instances of the uncertainty, referred to as scenarios, which can be expressed as $\delta_\kappa \in \Delta$ for $\kappa = 1, \ldots, N$. These scenarios can be drawn from a known probability distribution over Δ that mimics the actual, unknown distribution. Alternatively, scenarios can be selected from historical data. This second method is particularly advantageous because it circumvents the complex task of identifying a probability distribution for δ, making it purely data-driven and distribution-free, with the only sensible hypothesis being the "identically distributed and independent" assumption for the uncertainty realisations δ_κ. After sampling the scenarios, we solve an approximate optimisation problem in which the optimisation variables must lie within the intersection of the constraints derived from each scenario, denoted by X_κ. Thus, the approximation can be expressed as $X(\delta) \approx \cap_{\kappa=1}^{N} X_\kappa(\delta_\kappa)$. Provided that the optimisation program resulting in fixing an uncertainty realisation $\delta = \delta_\kappa$ is convex, the deterministic approximation built by considering the intersection of multiple scenarios is simpler to solve than chance-constraint problems, as it is also convex.

2.1. Scenario approach preliminaries

The main result of the scenario theory can be summarised as the following theorem:

Theorem 2.1 (Scenario Approach [CG08]). *Consider the convex uncertain optimization problem with the decision variable $x \in \mathbb{R}^d$:*

$$\min_x \; f(x) \quad s.t. \quad x \in X(\delta), \quad \forall \delta \in \Delta, \tag{2.3}$$

where the uncertainty δ is distributed according to an unknown probability measure $\mathbb{P}$ over the set Δ. Let N independent and identically distributed (i.i.d.) samples $S^N =$

$\{\delta_1, \ldots, \delta_N\}$ be drawn from Δ. Let $x_N^\star$ denote the solution of the corresponding scenario program, which is defined as:

$$\min_x \; f(x) \quad s.t. \quad \bigcap_{\kappa=1}^{N} X_\kappa(\delta_\kappa), \quad \forall \delta_\kappa \in S^N. \tag{2.4}$$

The violation probability *of $x_N^\star$, defined as*

$$V(x_N^\star) := \mathbb{P}\left\{\delta \in \Delta : x_N^\star \notin X(\delta)\right\}, \tag{2.5}$$

is so that, for any $\epsilon \in [0, 1]$, the following probabilistic guarantee holds:

$$\mathbb{P}^N\left\{V(x_N^\star) > \epsilon\right\} \leq \sum_{i=0}^{\zeta-1} \binom{N}{i} \epsilon^i (1-\epsilon)^{N-i}, \tag{2.6}$$

where $\mathbb{P}^N$ denotes the product probability measure that generates the N samples.

Some clarifications of key concepts:

- ϵ represents the *violation probability level*, which indicates the maximum tolerable probability mass of unseen constraints that the scenario solution may violate.
- N denotes the *number of sampled scenarios*, which consists of independent and identically distributed (i.i.d.) uncertainty realizations.
- ζ refers to the *Helly's dimension*, and it is linked with the dimension of the support set.
- The right-hand side of the inequality represents the tail of a *Beta distribution*, which provides a bound on the risk that the scenario solution violates more than an ϵ-fraction of the constraints.

From this theorem, we can determine the number of scenarios needed to ensure that the violation probability remains within a prescribed limit ϵ, regardless of the randomness of the sampled scenario. This is achieved with a high confidence level of $1 - \beta$, where $\beta \in (0, 1)$.

Theorem 2.2 (Scenario Approach)**.** *For any $\epsilon \in (0, 1)$ (risk parameter) and $\beta \in (0, 1)$*

(confidence parameter), if the number of scenarios N satisfies the condition

$$\bar{\epsilon} = 1 - \sqrt[N-\zeta]{\frac{\beta}{\binom{N}{\zeta}}}, \tag{2.7}$$

then, with a probability of at least $1-\beta$, the scenario solution $x_N^\star$ is guaranteed to meet the ϵ-risk criteria with high confidence $1-\beta$.

In other words

$$\mathrm{P_N}[S^N \in \Delta^N : P[\delta \in \Delta : x \notin X = \cap_{\kappa=1}^{N} X_\kappa] < \epsilon] \geq 1-\beta. \tag{2.8}$$

A more practical way to find the correct number of scenarios after fixing ϵ, β is to compute the over-approximation [CGP09], i.e.

$$N \geq \frac{2}{\epsilon} \cdot [\frac{1}{\beta} + \zeta - 1]. \tag{2.9}$$

As anticipated, the number of scenarios N highly depend on the violation probability ϵ, and the so-called problem complexity ζ. In convex problems, the support dimension corresponds to the maximum number of constraints that can "support" the optimal solution $x_N^\star$, and it is typically limited by the dimension of the decision variable $x \in \mathcal{R}^d$.

To better understand the concept of support dimension, we define the support constraints and the support set [SFM13].

Definition 2.1 (Support constraint)**.** *Given a general scenario optimisation program, i.e.*

$$\min_x f(x), \tag{2.10a}$$

$$s.t.,$$

$$h(x, \delta_\kappa) \leq 0, \qquad \kappa = 1..N, \tag{2.10b}$$

with solution $x^\star$. A **support constraint** *is a constraint $\bar{h_\kappa}(x^\star, \delta_\kappa) \leq 0$ such that if eliminated from the problem, the $x_{N-1}^\star$ solution corresponds to a decrease in the cost function $f(x^\star)$*

$$f(x_{N-1}^\star) < f(x_N^\star). \tag{2.11}$$

Definition 2.2 (Support set)**.** *The* **support set** *$SP(x_N^\star)$ is defined as the set containing*

all the support scenarios

$$SP(x_N^\star) = \bigcup_{\kappa=1}^{N} \bar{h_\kappa}((x^\star, \delta_\kappa)). \tag{2.12}$$

It is worth noting that the support dimension, defined as the cardinality of the support set $dim(SP(x_N^\star))$, is inherently a random variable. The Helly's dimension, instead, is the best upper bound for the support dimension.

Definition 2.3. *(Helly's dimension) Denoting $|SP(x^\star)|$ the random cardinality of the support set $SP(x^\star)$, the* **Helly's dimension** *is the smallest integer ζ such that*

$$ess\ sup_{S^N \in \Delta^N} |SP(x^\star)| \leq \zeta. \tag{2.13}$$

For the class of fully supported problems, interestingly, the Helly's dimension is equal to the dimension of the optimisation variable vector d, i.e.

$$|SP(x^\star)| <= \zeta = d. \tag{2.14}$$

In contrast, for the broader class of non-fully-supported problems, Helly's dimension is difficult to assess. The problem of bounding the support dimension is crucial for two reasons: First, using a non-tight bound requires more scenarios to ensure a-priori violation guarantees, and, most importantly, it leads to an overly conservative solution.

2.2. Scenario Approach in a multi-agent setting

The scenario approach has been utilised as a framework in distributed multi-agent settings to identify a shared feasible solution that is optimal for the entire uncertainty set, except for a subset with an arbitrarily small probability measure. In [YTX19] and [Cha+21], the authors leverage the scenario approach to address uncertainty in a distributed multi-agent framework. They employ a distributed random projection algorithm to achieve a solution with high confidence of satisfying the constraints. This algorithm relies on the communication of local "support sets" among neighbouring agents to ensure a solution that meets the global constraint, which represents the intersection of all local constraints.

To protect agents' privacy and avoid disclosing scenarios among them, the authors in [Mar+18] and [Fal+20] develop a new framework. This framework preserves local information and privacy at the cost of requiring more scenarios by sharing only dual variables among agents. In [Fal+20], they present two families of problems that can be resolved

using their method: **(1) decision-coupled uncertain programs:** where agents must reach a consensus on a shared vector of decision variables while satisfying local uncertain constraints; **(2) constraint-coupled uncertain problems:** where agents aim to find the optimal solution while adhering to uncertain local constraints and deterministic global ones, which can be viewed as a shared resource.

In the works mentioned above, uncertainty affects only the agents' local constraints, and none of the proposed methods address uncertainty within a global coupling constraint.

In this section, we present one of the primary contributions of this thesis, as outlined in [DRS24]. This contribution presents a scenario-theory framework explicitly designed for multi-agent uncertain optimisation problems, where uncertainty directly influences the global coupling constraints.

We utilise a scenario-based optimisation approach to transform the stochastic problem into a deterministic one, and then propose a distributed solution for the robustified problem. The presented algorithm ensures convergence to the exact solution of the centralised chance-constrained problem. We also demonstrate that, for specific structures of the chance constraint, it is possible to bound the problem's support dimension, thereby showing that the number of scenarios required to guarantee a prescribed violation probability does not depend on the number of agents. The latter property is then used in chapter 6 to develop a scalable distributed algorithm for optimising a VPP network.

2.2.1. Problem Statement

Consider a multi-agent cooperative setting with m agents. The system aims to minimise a global cost function built by summing the private agent's cost functions $f_i(x_i) : \mathcal{R}^{d_i} \to \mathcal{R}$, by acting on the local decision vector $x_i \in R^{d_i}$, and while satisfying local constraints, the optimisation program becomes

$$min_{\{x_i\}_{i=1}^m} \sum_{i=1}^{m} f_i(x_i), \tag{2.15a}$$

$$s.t.,$$

$$x_i \in X_i \ \forall i = 1..m. \tag{2.15b}$$

The agents are coupled by a set of r global and uncertain coupling constraints of the form

$$\sum_{i=1}^{m} g_i(x_i, \delta) \leq 0, \tag{2.16}$$

with $g_i(x_i, \delta) : \mathcal{R}^{d_i} \to \mathcal{R}^r$.

Assumption 2.1 (Convexity). *The stochastic optimisation problem is convex (for fixed δ):*

- *The cost functions $f_i(x_i)$ are convex.*
- *The local sets X_i are convex.*
- *The feasible set defined by* (2.16) *is convex, for fixed δ.*

Although it is a strict assumption, many engineering problems can be approximated in this form, and they are also the best class of problems for which we have tools to handle uncertainty.

Eq. (2.16) represents a global uncertain constraint which depends on the decision of all the agents x_i and the realisation of the random variable $\delta \in \Delta$ taken from the uncertainty set Δ with probability measure P_δ, which is unknown.

The system aims to find the optimal solution $x^\star = [x_1^{*T}..x_m^{*T}]^\top \in \mathcal{R}^d$ which is robust, in a probabilistic sense, towards the possible realisations of the uncertainty $\delta \in \Delta$ with a predefined risk level $\epsilon \in (0, 1)$. The following chance constraint characterises this robustness requirement, i.e.

$$\mathbf{P}\{\delta \in \Delta : \sum_{i=1}^{m} g_i(x_i^\star, \delta) > 0\} < \epsilon. \tag{2.17}$$

The chance-constraint centralised problem becomes

$$\mathcal{P} : min_{\{x_i\}_{i=1}^m} \sum_{i=1}^{m} f_i(x_i), \tag{2.18a}$$

$$s.t.,$$

$$x_i \in X_i, \ \forall i = 1..m, \tag{2.18b}$$

$$\mathbf{P}\{\delta \in \Delta : \sum_{i=1}^{m} g_i(x_i, \delta) > 0\} < \epsilon. \tag{2.18c}$$

The latter chance-constrained optimisation program is relevant for all systems in which agents cooperate to optimise the network's welfare in the presence of a common, limited, uncertain resource, or a global aggregated performance level that the agents must satisfy. Both problem classes arise, for example, in power grid management, electric vehicle fleet charging, distributed energy resources, and the integration of energy storage.

The resulting scenario program can be solved in a centralised fashion; however, as noted in the literature, the agents wish to preserve private information (f_i, x_i, $g_i(x_i, \delta)$, and X_i). Foremost, the number of scenarios, N, depends on the number of agents, m, making uncertainty management scale with the network size.

A possible solution is to exploit the problem's separable structure to develop a distributed scheme. In such a scheme, each agent optimises a local problem and communicates a signal to its neighbours, used to iteratively steer the agents' decision variables towards a consensus operative point, without disclosing any private information.

However, such schemes do not address the underlying issue with the scenario approach. In fact, in multi-agent settings, the number of variables d of the centralised problem directly depends on the number of agents m; in particular, considering homogeneous agents, the number of decision variables of the program becomes $d = d_i \cdot m$. Therefore, the number of scenarios N needed to ensure a risk level $\bar{\epsilon}$ grows linearly with m and thus is not scalable with the number of agents.

As an example, consider a network with $m = 5$ agents, each one with $d_i = 2$ decision variables, the centralised optimisation program has a support dimension of $d = m \cdot d_i = 10$, requiring $N = 254$ to guarantee a violation probability under $\bar{\epsilon} = 10\%$ with confidence $1 - 10^{-4}$. If the network has $m = 100$ agents, the number of scenarios increases to $N = 2540$. Increasing scenarios can affect both the numerical stability and real-time capability of the algorithm used to solve the problem, making it challenging to design scenario-based optimisation and control algorithms for distributed systems.

The novel idea presented in [DRS24] is to exploit the geometry of the global coupling constraint to provide tighter bounds on the Support dimension of the problem, thereby making the framework scalable. Two different bounds are presented now for two classes of coupling constraints $g(x, \delta)$.

We assume that the agents can extract N realisations of the uncertainty $\delta \in \Delta$ from historical data and use them to robustify the solution by providing off-sample certificates for the solution, as per standard scenario theory. Problem (2.15) is approximated as a deterministic one, where the chance constraint (2.17) is replaced by the envelope of N deterministic constraints defined by scenarios δ_κ, $\kappa = 1, \dots, N$, i.e.,

$$\mathcal{F}_\kappa = \{x = [x_1^\top .. x_m^\top]^\top : \sum_{i=1}^{m} g_i(x_i, \delta_k) \leq 0\}. \tag{2.19}$$

The approximated centralised scenario program becomes

$$\mathcal{P}_{\mathcal{N}} : min_{\{x_i\}_{i=1}^m} \sum_{i=1}^{m} f_i(x_i), \tag{2.20a}$$

$$s.t.,$$

$$x_i \in X_i, \ \forall i = 1..m, \tag{2.20b}$$

$$x \in \bigcap_{k=1}^{N} \mathcal{F}_\kappa. \tag{2.20c}$$

Provided that the agents pick N scenarios, $S^N = \{\delta_1, ..\delta_N\} \in \Delta^N$, problem $\mathcal{P}_{\mathcal{N}}$ becomes deterministic. Additionally, given a confidence parameter $\beta \in (0, 1)$ and a risk parameter $\bar{\epsilon} \in (0, 1)$, we can compute the required number of scenarios using Eq. 2.7 where d is the dimension of the aggregated decision vector of the problem ($d = \sum d_i$), it follows that with confidence $1 - \beta$, the probability that the optimal solution $x^\star = [x_1^{*T}..x_m^{*T}]^\top$ of $\mathcal{P}_{\mathcal{N}}$, violates the constraint (2.16) for a new uncertainty realization δ, is less than $\bar{\epsilon}$, i.e.

$$\mathbf{P_N}\{S^N = \{\delta_1..\delta_N\} \in \Delta^N : \mathbf{P}\{\delta \in \Delta \ \ s.t. \ \ \sum_{i=1}^{m} g_i(x_i, \delta) > 0\} \leq \epsilon\}\} \geq 1 - \beta. \tag{2.21a}$$

A naive solution can be derived by applying the reasoning in [Fal+20] to the global constraint (2.16) by substituting it with a set of convex local constraints

$$\bigcap_{i=1}^{m} \{x_i : (g_i(x_i, \delta) \leq 0)\}. \tag{2.22}$$

This solution is conservative since it implies that every agent contributes to the global constraint with a non-positive term $g_i(x_i, \delta) <= 0$, which is not true in general; the optimal solution should also contemplate cases in which only a part of the agents provide a non-positive contribution to the global constraint which is enough to compensate other "violating" agents.

Unfortunately, this is the best we can do for generic constraints as in Eq. 2.17; however, for two specific structured constraint families, we can improve this bound.

2.2.2. Structured uncertainty space

If the uncertainty space is structured, i.e. it respects the structure

$$g(x, \delta) = G(x)Q(\delta) + H(x) + S(\delta), \tag{2.23}$$

with $G : \mathcal{R}^d \to \mathcal{R}^{r\times q}$, $Q : \Delta \to \mathcal{R}^q$, $H : \mathcal{R}^d \to \mathcal{R}^r$, and $S(\delta) : \Delta \to \mathcal{R}^r$, we can use the results from [Zha+15] to derive a better bound for the Helly's dimension.

For constraints as in Eq. 2.23, theorem 2.3 holds [Zha+15].

Theorem 2.3. *Let $g(x, \delta)$ be as in (2.23). Then Helly's dimension is upper bounded by $r \cdot (q + 1)$ almost surely.*

The bound $r \cdot (q + 1)$ is independent of m in the proposed framework.

Now we want to prove that, if each $g_i(\ )$ belongs to the class of structured uncertainty constraints, then the global coupling constraint $g(x, \delta) = \sum_{i=1}^{m} g_i(x_i, \delta_k) \leq 0$ has a Support dimension upper bound independent of m.

Lemma 2.1. *Let the local constraint $g_i(x_i, \delta)$ be*

$$g_i(x_i, \delta) = G_i(x_i)Q(\delta) + H_i(x_i) + S_i(\delta), \tag{2.24}$$

with $G_i : \mathcal{R}^{d_i} \to \mathcal{R}^{r\times q_i}$, $Q : \Delta \to \mathcal{R}^{q_i}$, $H_i : \mathcal{R}^d_i \to \mathcal{R}^r$, and $S(\delta) : \Delta \to \mathcal{R}^r$. Then, the coupling constraint

$$g(x, \delta) = \sum_{i=1}^{m} g_i(x_i, \delta_k) \leq 0, \tag{2.25}$$

has Helly's dimension upper bounded by $r \cdot (q_i + 1)$ almost surely.

Proof. We can rewrite

$$g(x, \delta) = \sum_{i=1}^{m} g_i(x_i, \delta_k), \tag{2.26a}$$

$$= \sum_{i=1}^{m} (G_i(x_i)Q(\delta) + H_i(x_i) + S_i(\delta)), \tag{2.26b}$$

$$= \hat{G}(x)Q(\delta) + \hat{H}(x) + \hat{S}(\delta). \tag{2.26c}$$

By defining

$$\hat{G}(x)Q(\delta) = (\sum_{i=1}^{m} G_i(x_i)) \cdot Q(\delta), \tag{2.27a}$$

$$\hat{H}(x) = \sum_{i=1}^{m} H_i(x_i), \tag{2.27b}$$

$$\hat{S}(\delta) = \sum_{i=1}^{m} S_i(\delta), \tag{2.27c}$$

with $\hat{G}(x) : \mathcal{R}^d \to \mathcal{R}^{r \times q_i}$, $Q : \Delta \to \mathcal{R}^{q_i}$, $\hat{H}(x) : \mathcal{R}^d \to \mathcal{R}^r$, and $\hat{S}(\delta) : \Delta \to \mathcal{R}^r$, we can rewrite constraint (2.16) as in (2.23), thereby using theorem 2.3 to prove that

$$\zeta \leq (r+1) \cdot q_i, \tag{2.28}$$

concluding the proof. □

The dimensions r and q_i do not depend on m, making the bound scalable to an increasing number of agents. Moreover, considering the case where the constraint has the special structure

$$G_i(x_i) = 0 \Rightarrow g_i(x_i, \delta) = H_i(x_i) + S_i(\delta), \tag{2.29}$$

that is, the uncertainty enters as an additive term in the constraint; then the support dimension is bounded by r. This is an important class of constraints that includes the case of a structured decision space defined by convex functions $H_i(x_i)$, e.g., matrices A_i independent of δ, as shown in the next section.

2.2.3. Structured decision space

We consider scenario problems with affine constraints, where the uncertainty δ enters only additively, i.e.

$$g_i(x_i, \delta) = A_i \cdot x_i + b_i(\delta), \tag{2.30a}$$

$$g(x, \delta) = \sum_{i=1}^{m} (A_i \cdot x_i + b_i(\delta)), \tag{2.30b}$$

with $A_i : \mathcal{R}^{d_i} \to \mathcal{R}^r$ and $b_i(\delta) : \Delta \to \mathcal{R}^r$.

For this family of constraints, we can rely upon the concept of support rank ρ, as defined in [SFM13], by realising that, whichever the scenario realisation δ is, the constraint is not able to restrain directions normal to the row-space of matrix A due to its independence from δ.

For constraints with the structure of (2.30), lemma 2.2 holds.

Lemma 2.2. *If $r < d$, the support dimension ζ is upper bounded by r.*

Proof. We can rewrite the global coupling constraint as

$$g(x, \delta) = \sum_{i=1}^{m} (A_i \cdot x_i + b_i(\delta)) = \hat{A} \cdot x + \hat{b}(\delta), \tag{2.31}$$

where $\hat{A} = [A_1 \; .. \; A_m]$ and is such that $\hat{A} : \mathcal{R}^d \to \mathcal{R}^r$ and $\hat{b}(\delta) = \sum_{i=1}^{m} b_i(\delta)$ such that $\hat{b}(\delta) : \Delta \to \mathcal{R}^r$.

After fixing δ, the support rank is equal to the dimension of the row space of $\hat{A}$, which is bounded by $min(r, d)$.

$$\dim(row(\hat{A})) = rank(\hat{A}) \leq \min(r, d). \tag{2.32}$$

Therefore, if $r < d$ the following chain of inequalities holds $\zeta \leq \rho \leq r$. □

By establishing an upper bound on the support dimension for the centralised problem that is independent of the number of agents, we have demonstrated that managing uncertainty is independent of the network's size or complexity, drastically reducing the number of scenarios to be used compared to the case of non-global uncertainty presented in [Fal+20]. In the proposed formulation, the number of scenarios directly correlates with a reduction in the problem's constraints, making the mathematical optimisation easier to solve with standard solvers for convex programs.

A relevant concern is how the framework behaves when the true uncertainty differs from the structure for which we have modelled it, and the assumption about the support bound is invalid. In this case, the theoretical bound on the violation probability may not hold tightly, and we may observe a higher violation level than predicted by the theory. However, the scenario approach allows us to assess the a posteriori validity of such bounds by computing the number of support constraints of the actual realisation of the uncertainty, enabling us to correct and redefine assumptions and increase the number of scenarios to return under the desired violation level. Therefore, although the theoretical guarantees degrade if the assumptions on the constraint structure are not satisfied, the framework allows practical adaptation through a posteriori data-driven validation.

Finally, we stress that the analysis is valid for a narrow class of problems in which the uncertainty enters the global coupling constraint and is structured as previously described.

However, this class of problems remains important because it includes situations in which agents aggregate their services into a single service, or when a common, uncertain resource is shared across agents. An important engineering problem belonging to this class is the optimisation of interconnected microgrids, which provide an aggregate, uncertain service to the TSO. In chapter 6, we apply the developed theory to build a scalable architecture for a Virtual Power Plant (VPP) that aggregates microgrids for service provision. In that chapter, we strengthen the proof by showing the separation between uncertainty complexity N and the number of agents m numerically.

Chapter 2 Bibliography

[ATC09] Teodoro Alamo, Roberto Tempo, and Eduardo F. Camacho. "Randomized Strategies for Probabilistic Solutions of Uncertain Feasibility and Optimization Problems". In: *IEEE Transactions on Automatic Control* 54.11 (2009), pp. 2545–2559. DOI: 10.1109/TAC.2009.2031207.

[Ben+04] Aharon Ben-Tal et al. "Adjustable Robust Solutions of Uncertain Linear Programs". In: *Mathematical Programming* 99.2 (2004), pp. 351–376. DOI: 10.1007/s101070400300.

[BN98] Aharon Ben-Tal and Arkadi Nemirovski. "Robust Convex Optimization". In: *Mathematics of Operations Research* 23.4 (1998), pp. 769–805. DOI: 10.1287/moor.23.4.769.

[BS04] Dimitris Bertsimas and Melvyn Sim. "The Price of Robustness". In: *Operations Research* 52.1 (2004), pp. 35–53. DOI: 10.1287/opre.1030.0065.

[CC06] Giuseppe C. Calafiore and Marco C. Campi. "The Scenario Approach to Robust Control Design". In: *IEEE Transactions on Automatic Control* 51.5 (2006), pp. 742–753. DOI: 10.1109/TAC.2006.875041.

[CG08] Marco C. Campi and Simone Garatti. "The Exact Feasibility of Randomized Solutions of Uncertain Convex Programs". In: *SIAM Journal on Optimization* 19.3 (2008), pp. 1211–1230. DOI: 10.1137/080711019.

[CGP09] Marco C. Campi, Simone Garatti, and Maria Prandini. "The scenario approach for systems and control design". In: *Annual Reviews in Control* 33.2 (2009), pp. 149–157. ISSN: 1367-5788. DOI: https://doi.org/10.1016/j.arcontrol.2009.07.001. URL: https://www.sciencedirect.com/science/article/pii/S1367578809000479.

[CGP19] Marco C. Campi, Simone Garatti, and Maria Prandini. "Scenario Optimization for Model Predictive Control". In: *Handbook of Model Predictive Control* (2019), pp. 1–23. DOI: 10.1007/978-3-030-18358-3_1.

[Cha+21] Mohammadreza Chamanbaz et al. "Randomized Constraints Consensus for Distributed Robust Mixed-Integer Programming". In: *IEEE Transactions on Control of Network Systems* 8.1 (2021), pp. 295–306.

[DRS24] A. D. Del Duca, F. Ruiz, and R. Scattolini. "Distributed Stochastic Optimisation with Uncertain Coupling Constraints". In: *2024 IEEE 63rd Conference on Decision and Control (CDC)*. Milan, Italy, 2024, pp. 4431–4436. DOI: 10.1109/CDC56724.2024.10886359. URL: https://hdl.handle.net/11311/1286946.

[Fal+20] Alessandro Falsone et al. "A scenario-based approach to multi-agent optimization with distributed information∗∗Research was partially supported by EPSRC UK, grant EP/P03277X/1." In: *IFAC-PapersOnLine* 53.2 (2020). 21st IFAC World Congress, pp. 20–25. ISSN: 2405-8963. DOI: https://doi.org/10.1016/j.ifacol.2020.12.034.

[GC13] Simone Garatti and Marco C. Campi. "Modulating Robustness in Control Design: Principles and Algorithms". In: *IEEE Control Systems Magazine* 33.2 (2013), pp. 36–51. DOI: 10.1109/MCS.2012.2234964.

[GL19] X. Geng and J. Luedtke. "Theory and Applications of Chance-Constrained Optimization". In: *Mathematical Programming* 177.1 (2019), pp. 1–38. DOI: 10.1007/s10107-018-1224-4.

[JG12] R. Jiang and Y. Guan. "Data-driven Chance Constrained Stochastic Program". In: *Mathematical Programming* 158.1 (2012), pp. 291–327. DOI: 10.1007/s10107-016-1002-2.

[Küç22] S. Küçükyavuz. "Chance-Constrained Optimization under Limited Distributional Information: A Review of Reformulations Based on Sampling and Distributional Robustness". In: *Mathematical Programming* 190.1 (2022), pp. 1–38. DOI: 10.1007/s10107-021-01688-1.

[Mar+18] Kostas Margellos et al. "Distributed Constrained Optimization and Consensus in Uncertain Networks via Proximal Minimization". In: *IEEE Transactions on Automatic Control* 63.5 (2018), pp. 1372–1387. DOI: 10.1109/TAC.2017.2747505.

[SFM13] Georg Schildbach, Lorenzo Fagiano, and Manfred Morari. "Randomized Solutions to Convex Programs with Multiple Chance Constraints". In: *SIAM Journal on Optimization* 23.4 (2013), pp. 2479–2501. DOI: 10.1137/120878719.

[TDC13] Roberto Tempo, Fabrizio Dabbene, and Giuseppe C. Calafiore. *Randomized Algorithms for Analysis and Control of Uncertain Systems: With Applications*. Communications and Control Engineering. Springer London, 2013. ISBN: 978-1-4471-4609-4. DOI: 10.1007/978-1-4471-4610-0.

[YTX19] Keyou You, Roberto Tempo, and Pei Xie. "Distributed Algorithms for Robust Convex Optimization via the Scenario Approach". In: *IEEE Transactions on Automatic Control* 64.3 (2019), pp. 880–895. DOI: 10.1109/TAC.2018.2828093.

[Zha+15] Xiaojing Zhang et al. "On the sample size of random convex programs with structured dependence on the uncertainty". In: *Automatica* 60 (2015), pp. 182–188. ISSN: 0005-1098. DOI: https://doi.org/10.1016/j.automatica.2015.07.013.

3 | Active Learning and Dual MPC

When controlling a system, a key source of uncertainty arises from the discrepancies between reality and the model being used. This intrinsic uncertainty is consistently present and has led to significant efforts in developing control strategies that can adapt the model during operations. These strategies, which focus on learning from data and refining the model, belong to the broad family of Adaptive Control (AC) theory.

Moreover, the controller needs to address model uncertainty by balancing immediate control performance (exploitation) with the ability to identify a better model through probing the system (exploration), all while adhering to operational constraints. The problem of simultaneously controlling and identifying a dynamic system, commonly referred to as the dual control problem, was addressed in [Fel63]. The author noted that achieving optimal performance from a controller requires solving a stochastic optimal control problem in an augmented space that includes both the system's states and parameters. In this augmented control problem, the dual nature of the control action is highlighted, clarifying how the control policy can actively learn a better model by gaining information and use that information to gain a future advantage in exploiting control performance. This problem is linked with the Dynamic programming concept of the multi-armed bandit problem [Bel57; BR56].

In [Fel63], the author also points out the intractability of solving such problems exactly due to the explosion in complexity, known as the curse of dimensionality, which is inherited from the Bellman equations for solving the augmented stochastic optimal control problem. As the time horizon increases, the computation of the information gained depends on future outputs, which are uncertain.

Classic control theory bypasses this problem by utilising the Certainty-Equivalence control framework, in which the system model is identified and the controller is designed as if there were an equivalence between the system's actual underlying dynamics and its model. The resulting controller must be designed robust enough to compensate for differences between

the model and the system.

Instead, in the Dual Control theory, this balance between exploration and exploitation is directly accounted for by the controller. An overview of Dual Control methods is presented in the work [Wit95]; the author points out that different techniques can be used to address the problem's intractability. In the work [Mes18] the author divided those methods into two classes: explicit dual control, where the control action is directly augmented with a probing signal, which is used to perturb and learn the system, and implicit dual control, which instead relies on introducing approximations in the stochastic optimal control problem.

A Bayesian perspective provides a natural framework for dual control. In this setting, the system's uncertain parameters are modelled as random variables with a prior probability distribution, which is updated recursively using Bayes' rule as new data becomes available. The resulting control problem is formulated as a stochastic optimal control problem in the belief space, where the state now includes both the physical system and the posterior distribution of the uncertain parameters [Åst65]. The Bayesian formulation enables principled reasoning about uncertainty and facilitates the design of controllers that explicitly account for the value of information.

A significant advantage of Bayesian dual control lies in its theoretical optimality and its ability to achieve near-optimal performance with limited prior knowledge. It allows the controller to naturally balance risk and learning, avoiding overly conservative or overly aggressive behaviours. Moreover, this approach provides interpretable probabilistic confidence in decisions, which is especially relevant in safety-critical systems such as autonomous vehicles, robotics, or power networks under uncertainty [KH16].

However, the main drawback of Bayesian dual control is its computational intractability. The belief-state dynamics are infinite-dimensional, and exact solutions require propagating and optimising over probability distributions. As a result, exact Bayesian dual control is only tractable for very low-dimensional or linear-Gaussian systems, where analytical updates are possible. Practical implementations, therefore, rely on approximations such as linearisation, Gaussian processes, or Monte Carlo sampling. These approximations reduce computational burden but may introduce biases or limit robustness when the underlying uncertainty distribution deviates from the assumed model.

An alternative to Bayesian statistics is the Set Membership Estimation (SME) theory from the works [MV91], [MB82] and [BM15], which treats uncertainty as Unknown-but-Bounded rather than stochastic, allowing the application of Set Theory. The SME theory has advantages over the Bayesian framework, including greater computational tractability

and being agnostic regarding the underlying probability measures governing the uncertain set.

The SME framework has been formulated within Model Predictive Control (MPC) schemes, where it has demonstrated strong performance, providing robustness against disturbances, satisfaction of constraints, and recursive feasibility. There, the model uncertainty is quantified through the size of predicted state tubes [Par+23]. A different approach to reducing the computational requirements of joint SME and MPC strategies is to combine SM uncertainty sets with scenario-based MPC. In this work, we propose actively reducing uncertainty by minimising the volume of the feasible parameter set.

In this chapter, we introduce two novel Set-Membership MPC algorithms. The first algorithm, based on the work in [Mez+25], is a passive SME MPC that adapts an Autoregressive with Exogenous input (ARX) model and robustifies the control action using the scenario approach. The second algorithm, which extends the latter, includes minimising the Feasible Parameter Set (FPS) volume as an indicator of system uncertainty. This enables the controller to optimise the exploration/exploitation tradeoff and provides the control action with probing features. The effectiveness of the developed algorithm is presented in chapter 7, where it is used for aggregate control of a group of Thermo-Electric Refrigerators (TERs) for Demand Response (DR).

3.1. Set-Membership Estimation Theory preliminaries

The SME framework provides an alternative to classical probabilistic estimation methods for system identification. The key underlying difference is that, instead of assuming a probability distribution for noise and disturbances, SME relies on the concept of Unknown-But-Bounded (UBB) uncertainty, meaning that the uncertainty $\delta \in \Delta$ is unknown, with unknown probability measure $\mathbb{P}$, but with the set Δ being bounded. The SME theory is a general framework; in this thesis, we focus solely on the problem of estimating an ARX Model from data.

3.1.1. SME problem formulation for ARX models

We consider the problem of estimation and control of a dynamical system whose behaviour is unknown but can be described by an ARX model of order (n_a, n_b),

$$y_{k+1} = a_1 \cdot y_k + ... \; a_{n_a} \cdot y_{k-na+1} + b_1 \cdot u_k + ... \; b_{n_b} \cdot u_{k-nb+1} + e_k, \tag{3.1}$$

which can be compactly rewritten as

$$y_{k+1} = \theta^{*\top} \varphi_k + e_k, \tag{3.2}$$

where:

- $\varphi_k = [y_k, ... \; y_{k-na+1}, u_k, ... \; u_{k-nb+1}]^\top$ is the regressor vector
- $\theta^* \in \mathbb{R}^n, \; \theta^* = [a_1, ... \; a_{n_a}, b_1, ... \; b_{n_b}]^\top$ is the unknown parameter vector, containing the true $n = n_a + n_b$ model parameters.

The disturbance process $e = \{e_k\}_{k=0}^{\infty}$ is assumed to be UBB, i.e. the signal e is norm-bounded, $||e||_p^p <= \bar{e}$, for some p-norm. We specifically consider point-wise bounded disturbance.

Definition 3.1 (Point-wise bounded disturbance)**.** *The disturbance process $e = \{e_k\}_{k=0}^{\infty}$ is* **point-wise bounded***,i.e. the signal remains inside a finite bound set $[-\bar{e}, \bar{e}]$ with $\bar{e}$ a finite value which can be estimated from data, i.e.*

$$||e||_\infty = \sup_{k \in \mathbb{N}^+} |e_k| <= \bar{e}. \tag{3.3}$$

Thanks to the bounded nature of the systems' disturbance, by knowing the bound $\bar{e}$, we can characterise the set of all the models which are coherent with a-priori knowledge of the system and data up to the present time-instant k. This set is called the Feasible Parameter Set (FPS) at time k.

Definition 3.2 (Feasible Parameter Set)**.** *The* **Feasible Parameter Set** *at time k, denoted by Θ_k, is the subset of $\mathbb{R}^n$ defined as*

$$\Theta_k := \mathcal{K}_0 \cap \bigcap_{i=1}^{k} \mathcal{S}_i, \tag{3.4}$$

where:

- *$\mathcal{K}_0 \subseteq \mathbb{R}^n$ is a closed and convex polytope encoding prior knowledge about the parameter, such as physical bounds or structural properties:*

$$\mathcal{K}_0 = \{\theta : A_0 \cdot \theta \leq b_0\}, \tag{3.5}$$

- *$\mathcal{S}_i$ is a strip in $\mathbb{R}^n$ induced by the i-th measurement y_i and corresponding regression vector φ_{i-1}, i.e.,*

$$\mathcal{S}_i := \left\{\theta \in \mathbb{R}^n \;\middle|\; |y_i - \varphi_{i-1}^\top \theta| \leq \bar{\varepsilon}\right\}, \tag{3.6}$$

with $\bar{\varepsilon} > \bar{e}$ being the best known bound on the noise process e_k.

If the true parameter $\theta^* \in \mathbb{R}^n$ belongs to the initial set $\mathcal{K}_0$, and the available data at time k, $\mathcal{H}_k = \{(y_1, u_1), ..(y_k, u_k)\}$, is generated in accordance to the unknown system (3.2), by construction the true parameter $\theta^* \in \mathbb{R}^n$ is always contained inside the FPS, i.e.

$$\theta^* \in \Theta_k \;\forall\; k. \tag{3.7}$$

The SME theory aims at finding the best interpolatory estimate of the true parameter $\theta^\star$ based on the available data $\mathcal{H}_k$, a prior information about the FPS $\mathcal{K}_0$, and the knowledge about the bound on the disturbance $\bar{e}$.

Definition 3.3 (Interpolatory Estimate of the FPS Θ_k)**.** *Let $\Theta_k \subset \mathbb{R}^n$ be the FPS at time k. An* **interpolatory estimate** *$\theta_k \in \Theta_k$ is a point estimate of the true parameter vector θ^* that lies within Θ_k.*

By choosing an interpolatory estimate as our best guess, we accept a maximum error that is at most equal to the diameter of the information. In other words, the information content of the FPS is linked to the problem's intrinsic uncertainty.

Definition 3.4 (Diameter of information of the FPS Θ_k)**.** *The* **diameter of information** *of the FPS is defined as:*

$$dia(\Theta_k) = \max_{\theta_1, \theta_2 \in \Theta_k} ||\theta_1 - \theta_2||_p. \tag{3.8}$$

To further improve the accuracy of the estimate, we can choose a central estimate for $\theta^\star$, defined as the interpolatory estimate that minimises the information radius.

Definition 3.5 (Information radius of the FPS Θ_k)**.** *The* **information radius** *of the FPS Θ_k is defined as:*

$$rad(\Theta_k) = \max_{\theta_2 \in \Theta_k} ||\theta_c^k - \theta_2||_p. \tag{3.9}$$

The central estimator, which is the point of the FPS for which the uncertainty, i.e. the information radius, is minimised, coincides with the Chebyshev Centre θ_c^k of the FPS.

Definition 3.6 (Chebyshev Centre of the FPS Θ_k))**.** *The* **Chebyshev centre** *θ_c^k of Θ_k*

is defined as the centre of the largest Euclidean p-sphere inscribed in Θ_k.

$$\theta_c^k = \arg(\min_{\theta_1 \in \Theta_k, \theta_2 \in \Theta_k} \max ||\theta_1 - \theta_2||_p). \tag{3.10}$$

The maximum error between the Chebyshev Centre θ_c^k and the true parameter $\theta^\star$ is always less than the radius of the FPS $rad(\Theta_k)$

$$||\theta - \theta_c^k|| \leq rad(\Theta_k) \quad \forall \quad \theta \in \Theta_k. \tag{3.11}$$

By choosing the Chebyshev Centre θ_c^k as an estimator, we are simultaneously selecting the best point estimate for $\theta^\star$ based on the available information, and characterising our current belief over the parameter vector with the radius $rad(\Theta_k)$, and thus the size, of the FPS.

When a new measurement, $\tilde{y}_k$, is informative enough, we can reduce the radius of the FPS, decreasing the maximum error between the estimate and the true parameter, leading to a more accurate central estimate. The process of updating our belief to reduce the uncertainty consists of intersecting the previous FPS with the strip induced by the latest data $\mathcal{S}_k(\tilde{y}_k)$.

$$\mathcal{S}_k(\tilde{y}_k) := \left\{\theta \in \mathbb{R}^n \;\middle|\; |\tilde{y}_k - \varphi_{k-1}^\top \theta| \leq \bar{\varepsilon}\right\}, \tag{3.12}$$

Geometrically, each strip $\mathcal{S}_k$ correspond to a pair of parallel hyperplanes enclosing a slab in $\mathbb{R}^n$. The equations can be expressed as a polyhedral constraint of the kind:

$$\mathcal{S}_k = \{\theta : A_k \cdot \theta <= b_k\}, \tag{3.13a}$$

$$\text{with} \quad A_k = [\varphi_{k-1}, -\varphi_{k-1}]^\top, \tag{3.13b}$$

$$\text{and} \quad b_k = [\tilde{y}_k + \bar{\varepsilon}, -\tilde{y}_k + \bar{\varepsilon}]^\top. \tag{3.13c}$$

Hence, the FPS is recursively updated via intersection:

$$\Theta_k = \Theta_{k-1} \cap \mathcal{S}_k = (\bigcap_{i=1}^{k-1} \mathcal{S}_i) \cap \mathcal{K}_0, \tag{3.14a}$$

$$\Theta_k = \{\theta : [A_0, ...A_k]^\top \theta \leq [b_0, ...b_k]^\top\} = \{\theta : \bar{A}_k\, \theta \leq \bar{b}_k\}, \tag{3.14b}$$

where $\bar{A}_k$ and $\bar{b}_k$ are obtained by stacking all A_i, b_i up to k columnwise.

In Fig. 3.1, we illustrate an example of the FPS update following Eq. 3.14, with dimension of the parameter space $n = 2$. The initial FPS is defined as $\mathcal{K}, = \mathcal{I}([-1, 1], [1, 2])$,

representing an interval set that is updated iteratively using strips induced by the newest data $\mathcal{S}_k(\tilde{y}(k))$.

The FPS continues to shrink until iteration 4, at which point it lies within the strip $\mathcal{S}_4(\tilde{y}(4))$. This indicates that the data $\tilde{y}(4)$ is not informative enough to enhance our belief in the estimate of $\theta^\star$. However, the FPS at iteration $k = 4$ is small, suggesting that the uncertainty between the central estimate θ_c^4 and $\theta^\star$, represented by the radius of the FPS at iteration 4, is also small.

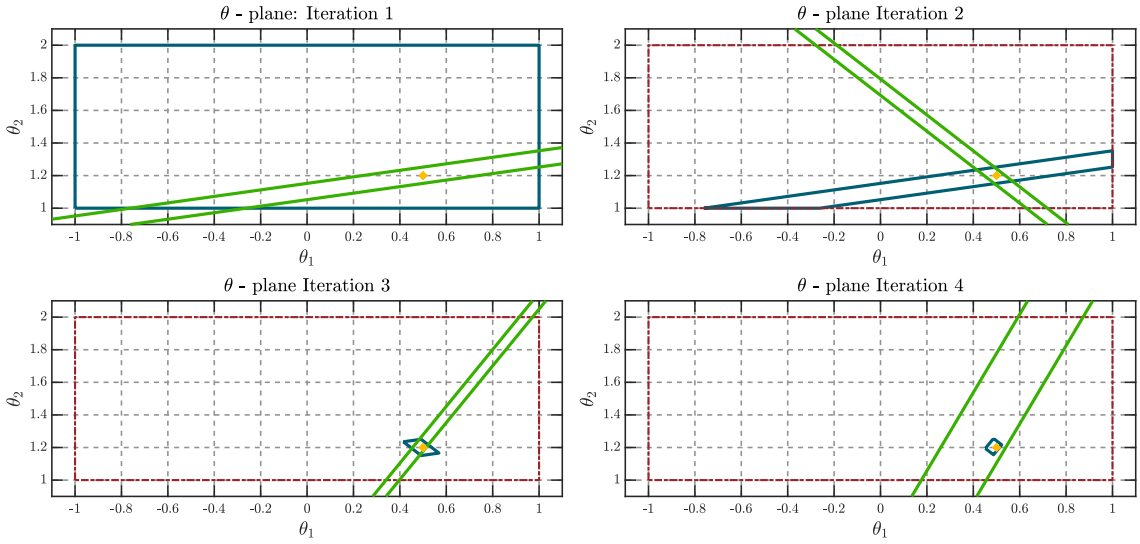


Figure 3.1: Feasible Parameter Set evolution. In dotted red the intial FPS $\mathcal{K}_0$; in blue the current FPS Θ_{k-1}; in green the strip induced by the $k-th$ measurement $\mathcal{S}_k$; in yellow, the true parameter $\theta^\star$.

The adaptation mechanism based on the SM framework ensures that the FPS is non-increasing and, under some mild hypotheses on the strong-tightness of the bound $\bar{e}$ and Persistent Excitation (PE) of the control signal u_k, the FPS becomes a singleton Set consisting only of the true parameter as k tends to infinity, i.e.

$$\mathbb{P}[\Theta_k = \{\theta^*\}] \to 1 \quad k \to \infty. \tag{3.15}$$

Consequently, the radius of information $rad(\Theta_k)$ tends to 0, meaning that our central estimate θ_c^k tends to the true value $\theta^\star$, without any bias, as the data increases:

$$\mathbb{P}[\theta_c^k = \theta^*] \to 1 \quad k \to \infty. \tag{3.16}$$

The proof of the previous two statements is postponed to section 3.4.

3.2. Fast SME-MPC with Probabilistic Guarantees

In Sec. 3.1, we described the general Set-Membership framework for parameter estimation of an ARX system. In this section, we present a receding-horizon algorithm that exploits this framework to adapt the model online.

The algorithm leverages a bounding-box over-approximation to update the FPS, incorporating a scenario-based approach to robustify the solution while keeping computational complexity minimal. This design not only enhances efficiency but also allows for effective real-time performance optimisation.

3.2.1. Infinite-Horizon Optimal Control Problem

The algorithm addresses the infinite-horizon optimal tracking of a reference signal r_{k+i}; the goal of the control system is to determine the optimal control action u_k^*, at each time instant k, to strike a balance between the closed-loop tracking error and the control action energy. The algorithm must account for model uncertainties and ensure the system operates within specified limits.

This is accomplished by solving a Robust Infinite-Horizon Optimal Control Problem

$$u_{k|k}^* = \arg\min_{u_k} \sum_{i=0}^{\infty} Q(y_{k+i|k} - r_{k+i})^2 + R u_{k+i|k}^2, \tag{3.17a}$$

$$\text{s.t. } y_{k+i+1|k} = \theta^{*\top} \varphi_{k+i|k}, \ \forall\, i, \tag{3.17b}$$

$$y_{k+i|k} \in \mathcal{Y} \ , \ u_{k+i|k} \in \mathcal{U}, \ \forall\, i \tag{3.17c}$$

$$y_{k|k} = \tilde{y}_k, \tag{3.17d}$$

where $\mathcal{U}$, $\mathcal{Y}$ are convex safe sets representing operational constraints of the system, r_{k+i} is the reference signal over the infinite horizon, and $\tilde{y}_k$ is the measured output at time k.

Considering an infinite-horizon problem is motivated by the theoretical optimality of the controller and the fact that, over an infinite horizon, the optimal policy must be stationary, thereby simplifying the analysis and providing an evaluation benchmark for finite-horizon approximations (MPC).

An optimal closed-loop policy $u_k = \Pi(y_k)$ for solving this IHOCP is through the so-called Receding Horizon. After solving the optimal control problem, we apply the first optimal control action $u_{k|k}^\star$ to the plan, then repeat the process at the next time instant $k+1$, updating the measure $\tilde{y}_k$ and the regressor vector φ_{k-1}.

Problem (3.17) is not easily solvable since, as previously stated, the true parameter vector θ^* is unknown. At time k, we only have access to the central estimator $\theta_c^k \in \Theta_k$.

3.2.2. Scenario-MPC

Not having probabilistic information on $\theta^\star$, we can rely on sampling N equiprobable values $D_N = \{\theta_1, \theta_2, ...\theta_N\}$ of possible parameters in the FPS, called scenarios. The $\kappa - th$ scenario θ_κ represents a pointwise estimate of the true parameter $\theta^\star$ that is coherent with historical data and previous knowledge about the system. For each scenario, the worst-case error is at most the information diameter: $|\theta_\kappa - \theta^\star| = d(\Theta_k) \leq 2 \cdot rad(\Theta_k)$.

The N random models $\theta_k \in \Theta_k$ drawn from the FPS are employed to generate N potential trajectories $\{y^{(\kappa)}_{k+0|k},\ y^{(\kappa)}_{k+1|k}, \ldots, y^{(\kappa)}_{k+N-1|k}\}$ for the system dynamics, with $\kappa = 1, 2, 3, \ldots, N$, which are then used to approximate the control problem, to a deterministic one.

The control action is chosen by minimising the cost function while satisfying safe constraints on input and state across all scenario realisations of potential models. This approach robustifies the problem with respect to model uncertainty and ensures probabilistic guarantees of constraint violation levels.

The scenario-based optimisation problem behind the MPC strategy, solved at each iteration, is stated as follows:

$$\min_{u_{k|k},\ldots,u_{M-1|k}} \sum_{\kappa=1}^{N} \sum_{i=0}^{M-1} Q \cdot (y^{(\kappa)}_{k+i|k} - r_{k+i})^2 + R \cdot u^2_{k+i|k}, \tag{3.18a}$$

s.t.

$$y^{(\kappa)}_{k+i+1|k} = \theta_\kappa^\top \varphi^{(\kappa)}_{k+i|k}, \forall i = 0, \ldots, M-1, \kappa = 1, \ldots, N, \tag{3.18b}$$

$$\varphi^{(\kappa)}_{0|k} = \varphi_k \quad \forall \kappa = 1, \ldots, N, \tag{3.18c}$$

$$\theta_\kappa \in \Theta_k \quad \forall \kappa = 1, \ldots, N, \tag{3.18d}$$

$$y^{(\kappa)}_{k+i+1|k} \in \mathcal{Y} \quad \forall i = 1, \ldots, M-1, \kappa = 1, \ldots, N, \tag{3.18e}$$

$$u_{k+i|t} \in \mathcal{U} \quad \forall i = 0, \ldots, M-1, \tag{3.18f}$$

where M denote the prediction horizon and the weighting matrices Q and R are positive definite and square of dimension $M \times M$, that is, $Q \succeq 0, R \succ 0$. Moreover, the reference signal r_{k+i} is known over the prediction horizon M, and the initial system state φ_k is constructed based on the ARX structure of the model.

By choosing a sufficiently large value for N, we can ensure a priori probabilistic guarantees regarding the violation level of the safe constraints. This way, we can limit the instances

in which our control action violates the safe constraints to a small probabilistic mass ε of the FPS. We define the average probability of constraint violation over time.

Definition 3.7. *The* **average probability of closed-loop constraint violation over time** *is the average-in-time violation of the robust constraints given the actual state of the system encoded in the regressor φ_k:*

$$E^k[V_k \mid \varphi_k] = E^k[\mathbb{P}\{\theta^\star \in \Theta_k : y_{k+1} \notin \mathcal{Y} \leq \varepsilon \mid \varphi_k\}. \tag{3.19}$$

By using the result presented in chapter 2:

$$E^k[V_k \mid \varphi_k] \leq \frac{\rho}{N+1} = \varepsilon, \tag{3.20}$$

we can compute the number of scenarios N for which the closed-loop constraint violation probability is smaller than a desired risk-level $\varepsilon \in (0, 1)$. Here we remark that ρ is the support rank of the chance constraint at time $k+1$ and for a Single Input Single Output (SISO) system in closed-loop operation, ρ equals 1 as proven in [SFM13].

The scenario-based approach was selected because it efficiently approximates the uncertainty on the value $\theta^\star$ by sampling possible model realisations from Θ_k. This feature reduces computational complexity and improves control performance compared to other robust methods. Furthermore, this approach allows for balancing the trade-off between performance and the probability of constraint violations over time, depending on the number of scenarios considered.

3.2.3. Uncertainty model adaptation through bounding box

Given that the system can be represented by an ARX model influenced by UBB noise, we can utilise the adaptation mechanism described in Section 3.1.1 to refine the FPS at each time step. This process involves eliminating any parameters that become inconsistent with the historical data $\mathcal{H}_k$ and our prior knowledge of the bounds $\bar{e}$.

One drawback of this adaptation mechanism is that, over time, the polytope becomes increasingly complex as additional strips are added. This complexity makes computing the central estimate θ_c^k more resource-intensive and requires greater memory storage. It may happen that many strips become inactive; determining which remain active may be computationally intensive.

To simplify the FPS, we over-approximate it with the bounding-box set $\hat{\Theta}_k$ constructed by

taking the maximum and minimum values for each parameter in Θ_k. This simplification yields an efficient algorithm that reduces the volume of the uncertainty set. However, this approach may be more conservative, since $\Theta_k \subseteq \hat{\Theta}_k$, potentially encompassing a broader range of parameter values than what is strictly dictated by the available information.

To find the intersection of each new strip with the bounding box, it is not necessary to compute the exact intersection set. Instead, we only need to determine the maximum and minimum values of the intersection along each coordinate direction, by solving two Linear Programs (LP), i.e.,

$$\begin{aligned}
&\theta_{i,\min}^{(k+1)} = \min \theta_i, \quad \forall i = 1, ..., na + nb, \\
&\text{s.t.} \quad \theta^\top \varphi_{k-1} \leq y_k + \bar{e}, \\
&\qquad \theta^\top \varphi_{k-1} \geq y_k - \bar{e}, \\
&\qquad \theta_i \in \Theta_k, \\
&\theta_{i,\max}^{(k+1)} = \max \theta_i, \quad \forall i = 1, ..., na + nb, \\
&\text{s.t.} \quad \theta^\top \varphi_{k-1} \leq y_k + \bar{e}, \\
&\qquad \theta^\top \varphi_{k-1} \geq y_k - \bar{e}, \\
&\qquad \theta_i \in \Theta_k.
\end{aligned}$$

The aforementioned optimisation problems aim to establish the new limit values for each parameter θ_i within the existing set $\hat{\Theta}_k$. If the new bounds are identical to those from the previous iteration, the set $\hat{\Theta}_k$ remains unchanged. However, if new extreme points are identified, the set is updated to

$$\hat{\Theta}_{k+1} = \{\theta_i : \theta_{i,\min}^{(k+1)} \leq \theta_i \leq \theta_{i,\max}^{(k+1)}, \qquad \forall i = 1, 2, \ldots, na + nb\}, \tag{3.22}$$

resulting in a reduction of the FSS volume.

Besides, the central estimate of the box set $\hat{\Theta}_k$ is easily computed by computing the component-wise central point, i.e.,

$$\theta_{i,k}^c = \frac{\theta_{i,\max}^{(k)} + \theta_{i,\min}^{(k)}}{2}, \qquad \theta_c^k = [\theta_{1,k}^c, \theta_{2,k}^c, ..\theta_{n,k}^c]^\top. \tag{3.23}$$

Once this estimation process is complete, the algorithm uses the latest FPS to sample the model scenarios. It then solves the scenario MPC in (3.18) to recompute the control signal.

The complete Fast SME-MPC algorithm is presented in 3.1.

Algorithm 3.1 Fast Adaptive SME MPC

Require: Model parameters: $\bar{e}$, n_a, n_b, N, M, Q, R

Require: Initial regressor vector φ_{k-1}, feasible parameter set $\hat{\Theta}_k$, central estimate θ_c^k, historical data $\mathcal{H}_k$, latest output $\tilde{y}_k$

1: **Model update:**
2: Solve the LP problem ▷ see Eq. (3.21)
3: Update the feasible parameter set $\hat{\Theta}_k$ ▷ see Eq. (3.22)
4: Compute the central estimate θ_c^k ▷ see Eq. (3.23)
5: **Scenario-based MPC:**
6: Sample a set of models $D_N = \{\theta_1, \theta_2, \ldots, \theta_N\}$ with $\theta_i \in \hat{\Theta}_k$
7: Solve the Scenario MPC problem ▷ see Eq. (3.18)
8: **Receding horizon:**
9: Apply the first control input $u_k = u_{k|k}^*$

3.3. Set-Membership Dual MPC

The previous SME Scenario MPC provides only an approximation of the IHOCP. This limitation arises not only from the incomplete representation of the entire set Θ_k, but also because it neglects the exploration-exploitation trade-off. In fact, current control actions can affect future information gathering, and as a result, the optimal solution must balance exploiting present performance with exploring the parameter space. This balance is essential for learning the model and achieving improved performance in the future.

In this section, we develop a tractable Set-Membership Dual MPC that minimises the FPS volume ($Vol(\Theta_k)$) as a proxy for model uncertainty, aiming to balance exploration of the FPS with control objectives. While the algorithm belongs to the indirect family, it retains certain advantages over other methods in the literature, particularly in terms of computation and complexity.

3.3.1. Problem formulation

To address this challenge, we propose a simplified approach that incorporates a model uncertainty penalty directly into the cost function. Specifically, we modify the objective

to include a function $V_{k+i}(y,u)$ that quantifies model uncertainty:

$$u^*_{k|k} = \arg\min_{u_k} \sum_{i=0}^{M} Q(y_{k+i|k} - r_{k+i})^2 + R u^2_{k+i|k} + \lambda \cdot V_{k+i}(y_{k+i|k}, u_{k+i|k}), \tag{3.24a}$$

$$\text{s.t. } y_{k+i+1|k} = \theta_c^{k\top} \varphi_{k+i|k},\ \theta_c^{k\top} \in \Theta_k, \forall\ i, \tag{3.24b}$$

$$y_{k+i|k} \in \mathcal{Y}\ ,\ u_{k+i|k} \in \mathcal{U}\ \forall\ i, \tag{3.24c}$$

$$y_{k|k} = \tilde{y_k} \tag{3.24d}$$

The probing capabilities of the control action are enhanced by minimising the function $V_{k+i}(y,u)$ that accounts for the uncertainty in the system's model. By adjusting the parameter λ, we can effectively balance the exploration-exploitation trade-off.

The model uncertainty is represented by the radius of information $rad(\Theta_k)$ (see sec. 3.1.1). In the following sec. 3.3.2 we are going to link the information radius to the volume of the FPS ($Vol(\Theta_k)$).

3.3.2. Exploratory-term in the cost function

Intuitively, minimising the size of the FPS is equivalent to minimising the uncertainty in the model parameter $\theta^\star$; by construction, the maximum error between our current best estimate θ_c^k and the true parameter θ^* is upper-bounded by the radius of information of Θ_k, i.e.,

$$rad(\Theta_k) = max||\theta - \theta_c^k|| \ \forall\ \theta \in \Theta_k. \tag{3.25}$$

At each time step, when a new data $\tilde{y}_k$ becomes available, the FPS is updated by adding two parallel hyper-semiplanes $\mathcal{S}_i = \{\theta : A_i\theta \leq b_i\}$. The direction of such semiplanes is orthogonal to the vector φ_k; thus, by changing φ_k through u_k, we can rotate the strip to choose a direction which shrinks the FPS the most, or in other words, the direction of maximum information gain on the model parameter θ^*.

In Fig. 3.2, an example of strips $\mathcal{S}_k$ in $n = 2$ dimension, for different values of u_k, shows the ability of u_k to rotate the semiplanes, also clear from the strip equation

$$\mathcal{S}_{k+1}(\tilde{y}_{k+1}) := \left\{\theta \in \mathbb{R}^n \ \middle|\ |\tilde{y}_{k+1} - \varphi_k^\top \theta| \leq \bar{\varepsilon}\right\} = \tag{3.26a}$$

$$\{\theta \in \mathbb{R}^n \mid |\tilde{y}_{k+1} - y_k\,\theta_1 + u_k\,\theta_2| \leq \bar{\varepsilon}\} \tag{3.26b}$$

Here, the only free variable is u_k, which can rotate the semi-planes around the centre θ_c^k.

By minimising the volume $vol(\Theta_k)$, we can shrink the FPS in all the directions, thus

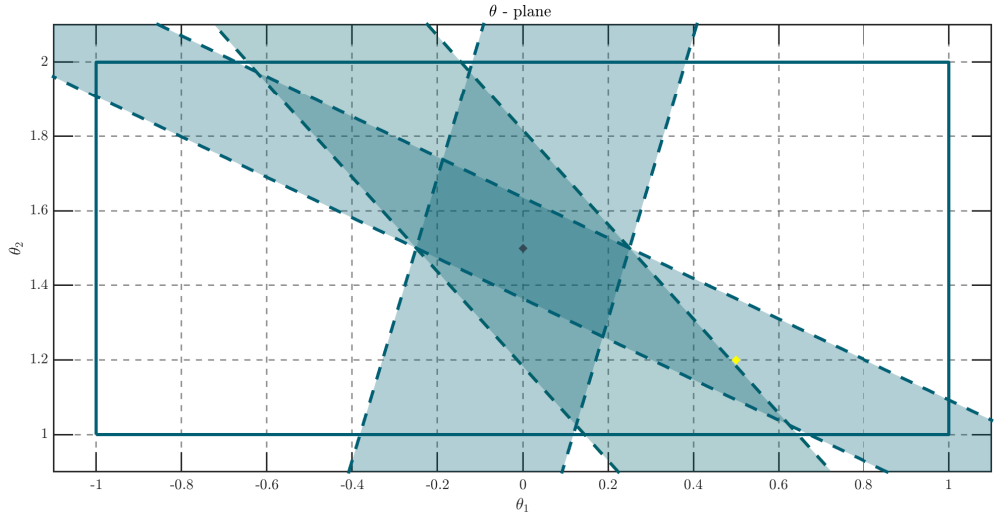


Figure 3.2: Strips $\mathcal{S}_k$ for different values of u_k.

reducing $rad(\Theta_k)$ and consequently the uncertainty. Therefore, by minimising the volume $vol(\Theta_k)$ over the prediction horizon M, we promote control actions that probe the system in directions that are the most informative for reducing model uncertainty while also striking performance optimality.

Unfortunately, the calculation of the volume $V_{k+i|k}(y_{k+i|k}, u_{k+i|k})$ depends on future measurements $y_{k+i|k}$. However, since we do not know the true parameter $\theta^\star$ that generates these data, a sensible approximation we can achieve is to minimise the volume under the assumption that the system behaves according to our central estimate model $\hat{V}_{k+i|k}(y^{(c)}_{k+i|k}, u_{k+i|k})$. While this approach may not be exact, it tends to yield an overestimation of the volume, meaning that the realised FPS has a lower volume than the expected one.

3.3.3. Multiobjective MPC

The final form of the MPC is:

$$\min_{u_{k|k},\dots,u_{M-1|k}} \sum_{i=0}^{M-1} Q(y^{(c)}_{k+i|k} - r_{k+i})^2 + Ru^2_{k+i|k} + \lambda \cdot \hat{V}_{k+i|k}(y^{(c)}_{k+i|k}, u_{k+i|k}), \tag{3.27a}$$

s.t.

$$y^{(\kappa)}_{k+i+1|k} = \theta_j^\top \varphi^{(\kappa)}_{k+i|k}, \forall i = 0,\dots,M-1, \kappa = 1,\dots,N, \tag{3.27b}$$

$$y^{(c)}_{k+i+1|k} = \theta_c^{k\top} \varphi^{(c)}_{k+i|k}, \forall i = 0,\dots,M-1, \tag{3.27c}$$

$$\varphi^{(\kappa)}_{0|k} = \varphi_k \quad \forall \kappa = 1,\dots,N, \tag{3.27d}$$

$$\theta_\kappa \in \Theta_k \quad \forall \kappa = 1,\dots,N, \tag{3.27e}$$

$$y^{(\kappa)}_{k+i+1|k} \in \mathcal{Y} \quad \forall i = 1,\dots,M-1, \kappa = 1,\dots,N, \tag{3.27f}$$

$$u_{k+i|k} \in \mathcal{U} \quad \forall i = 0,\dots,M-1, \tag{3.27g}$$

Where λ is a tunable parameter that weights the exploratory action of the controller, a wise choice could be to make λ depend on the value of the FPS Θ_k's volume at the current time $\lambda = \frac{\hat{\lambda}}{V_k}$. This way, as the volume of Θ_k shrinks, the exploratory action becomes more and more influential.

Here, we anticipate that the volume $\hat{V}_{k+i|k}(y^{(c)}_{k+i|k}, u_{k+i|k})$ is non-convex, but we postpone the discussion to section 3.5.2.

As before, we compute N scenarios to ensure probabilistic guarantees on violation levels.

3.4. SME Dual MPC theoretical guarantees

This section sec. 3.4 is dedicated to presenting proofs for the theoretical guarantees of SME Dual MPC.

3.4.1. FPS convergence to a singleton Set

In this subsection, we prove that, under some assumptions on the tightness of the bound and regressor persistent excitation, the FPS Θ_k converges to a singleton set solely consisting of the true parameter $\theta^\star$, i.e.,

$$lim_{k\to\infty}\Theta_k = \{\theta^*\}. \tag{3.28}$$

Definition 3.8 (Tightness of the bound $\bar{e}$)**.** *The noise bound $\bar{e}$ is* **tight**, *i.e., for any $\delta > 0$ there exists $\eta(\delta) > 0$ such that*

$$\mathbb{P}(|e_k| > \bar{e} - \delta) \geq \eta(\delta), \quad \forall k. \tag{3.29}$$

Definition 3.9. *The regressor sequence $\{\varphi_k\}$ is* **persistently exciting (PE)** *if there exist $\alpha > 0$, $\beta > 0$, and an integer $m > 0$ such that for all $i_0 \geq 0$,*

$$\alpha\, I \leq \frac{1}{m} \sum_{i=i_0}^{i_0+m-1} \varphi_i \varphi_i^\top \leq \beta\, I. \tag{3.30}$$

Theorem 3.1 (FPS convergence to a singleton set containing θ^*)**.**

Consider the discrete-time scalar system in (3.2)*, and the FPS Θ_k updated as in* (3.14)*. Assume the following conditions hold:*

(i) The noise sequence $\{e_k\}$ consists of independent random variables, each bounded in absolute value by $\bar{e}$.

(ii) The noise bound $\bar{e}$ is tight.

(iii) The regressor sequence $\{\varphi_k\}$ is PE.

Then, the sequence of feasible parameter sets $\{\Theta_k\}$ is non-increasing, i.e.,

$$\Theta_{k+1} \subseteq \Theta_k \quad \forall k \geq 0, \tag{3.31}$$

And converges almost surely to the singleton containing the true parameter vector θ^* with probability one,

$$\mathbb{P}[\lim_{k\to\infty} \Theta_k = \{\theta^*\} = 1], \tag{3.32}$$

or, in other words, the diameter of the FPS Θ_k approaches zero as time approaches infinity,

$$d(\Theta_k) = sup||\theta_1 - \theta_2||\forall \theta_1, \theta_2 \in \Theta_k, \tag{3.33}$$

$$d(\Theta_k) \to 0 \; as \quad k \to \infty. \tag{3.34}$$

In particular:

$$d(\Theta_k) = \sup ||\theta_1 - \theta_2|| \;\forall\; \theta_1, \theta_2 \in \Theta_k \Rightarrow ||\theta_c^k - \theta^*|| \to 0 \; as \quad k \to \infty. \tag{3.35}$$

Proof. First, by contradiction, we prove that a generic but fixed parameter $\hat{\theta} \in \mathbb{R}^n$ different from the true parameter $\hat{\theta} \neq \theta^*$ is excluded from Θ_k as $k \to \infty$. The parameter $\hat{\theta}$ is excluded from Θ_k if, with probability one, the prediction error exceeds $\bar{e}^2$.

$$z_k = max\{\frac{1}{m}\sum_{i=1}^{m}(y_i - \varphi_i^\top \hat{\theta})^2; \quad \frac{1}{m}\sum_{i=m+1}^{2m}(y_i - \varphi_i^\top \hat{\theta})^2; \quad \frac{1}{m}\sum_{i=(k-1)m+1}^{km}(y_i - \varphi_i^\top \hat{\theta})^2\} \tag{3.36}$$

$$\mathbb{P}[z_k > \bar{e}^2] \to 1 \quad , k \to \infty. \tag{3.37}$$

We define the bias between $\hat{\theta}$ and the true θ^*: $\tilde{\theta} = \theta^* - \hat{\theta} \neq 0$, we compute the prediction error as

$$y_i - \varphi_i^\top \hat{\theta} = y_i - (\varphi_i^\top \theta - \varphi_i^\top \tilde{\theta}) = e_i + \varphi_i^\top \tilde{\theta}, \tag{3.38}$$

and therefore:

$$\frac{1}{m}\sum_{i=i_0+1}^{i_0+m}(y_i - \varphi_i^\top \hat{\theta})^2 = \frac{1}{m}\sum_{i=i_0+1}^{i_0+m}((\varphi_i^\top \tilde{\theta})^2 + 2\varphi_i^\top \tilde{\theta} v_i + v_i^2). \tag{3.39}$$

Now we show that the probability of the prediction error being greater than $\bar{e}$ is greater than zero:

$$\mathbb{P}[\frac{1}{m}\sum_{i=i_0+1}^{i_0+m}(y_i - \varphi_i^\top \hat{\theta})^2 \geq \bar{e}^2] = \mathbb{P}[\frac{1}{m}\sum_{i=i_0+1}^{i_0+m}((\varphi_i^\top \tilde{\theta})^2 + 2\varphi_i^\top \tilde{\theta} v_i + v_i^2) \geq \bar{e}^2] \geq p > 0. \tag{3.40}$$

The first term $(\varphi_i^\top \tilde{\theta})^2$ is greater than zero thanks to the hypothesis on PE, i.e.,

$$\frac{1}{m}\sum_{i=i_0+1}^{i_0+m}(\varphi_i^\top \tilde{\theta})^2 \geq \alpha ||\tilde{\theta}||^2 > 0, \tag{3.41}$$

and the sum of the other terms $2\varphi_i^\top \tilde{\theta} v_i + v_i^2$ has a non-zero probability of being greater than$\bar{e}^2$,

$$\mathbb{P}[2\varphi_i^\top \tilde{\theta} v_i + v_i^2 > \bar{e}^2] = p > 0. \tag{3.42}$$

This is easily proven by the fact that $v_i^2 \leq \bar{e}^2$ and that the probability of $2\varphi_i^\top \tilde{\theta} v_i$, meaning the likelihood that $sign(\varphi_i^\top \tilde{\theta}) = sign(v_i)$ is greater than zero, thanks to the hypothesis on the tightness of the noise bound. So the probability of $\hat{\theta}$ being excluded, i.e.,

$$\mathbb{P}[z_k > \bar{e}] = 1 - \mathbb{P}[z_k \leq \bar{e}] = 1 - \prod_k (1 - \mathbb{P}[\frac{1}{m}\sum_{i=i_0+1}^{i_0+m}(y_i - \varphi_i^\top \hat{\theta})^2 \leq \bar{e}^2]) \tag{3.43}$$

$$= 1 - (1-p)^k \to 1 \quad k \to \infty. \tag{3.44}$$

In other words, any $\hat{\theta} \neq \theta^*$ is excluded from Θ_k in probability. This implies that $d(\Theta_k) \to 0$ in probability. By contradiction, consider that Θ_k contains other points $\hat{\theta}$ in probability; this contradicts the fact that any $\hat{\theta} \neq \theta^*$ is excluded in probability. This means that at least one subsequence of $\{d(\Theta_k)\}_k^\infty$ converges to zero. Moreover, $\Theta_{k+1} \subseteq \Theta_k$ implies that $\{d(\Theta_k)\}_k^\infty$ is a non-increasing sequence converging to zero, concluding the proof. □

The convergence results rely on PE, which could be enforced through a probing component on the control signal. In practice, the objective is not necessarily to converge to a singleton parameter set $\Theta = \{\theta^\star\}$. Instead, it is sufficient to reduce uncertainty to a level where further active learning does not significantly improve control performance, as explained in section 3.5.3. In fact, while the PE assumption provides theoretical convergence guarantees, enforcing excitation on the system may conflict with performance objectives without a beneficial effect. In practice, an external excitation is not needed; once the uncertainty is sufficiently small for control performance, i.e., when the FPS is small enough, there is no further incentive to explore the model space. Thus, PE should be only interpreted as a theoretical condition to ensure the correct identification of the true underlying model rather than an operational requirement.

3.4.2. Volume shrinking and certainty equivalence

Here, we prove that, if the Volume shrinks to 0, our control scheme converges to the Certainty Equivalence MPC.

Theorem 3.2 (Volume Shrinking and Certainty Equivalence)**.** *Let the true system be described by*

$$y_k = \theta^{*\top} \varphi_k + e_k, \quad k \in \mathbb{Z}_{k \geq 0}, \tag{3.45}$$

where $\theta \in \mathcal{K}_0 \subseteq \mathbb{R}^n$ is the true (unknown) parameter vector. Let the FPS Θ_k at time k be updated as in (3.14)*. The scenario* (3.27) *becomes asymptotically equivalent to the certainty-equivalent MPC that assumes $\theta = \theta^*$ is known:*

$$\min_{u_{k|k},\dots,u_{M-1|k}} \sum_{i=0}^{M-1} Q(y_{k+i|k} - r_{k+i})^2 + R u_{k+i|k}^2, \tag{3.46}$$

s.t.

$$\varphi_{0|k} = \varphi_k, \quad y_{k+i+1|k} = \theta_j^\top \varphi_{k+i|k} \forall i = 0, \dots, M-1, \tag{3.47}$$

$$y_{k+i+1|k} \in \mathbb{Y}, u_{k+i|k} \in \mathbb{U} \quad \forall i = 1, \dots, M-1. \tag{3.48}$$

Proof. This theorem is a direct consequence of Theorem 1 and the fact that

$$lim_{k\to\infty}\Theta_k = \{\theta^*\}. \tag{3.49}$$

To prove the theorem, we can see that both the cost function and the constraints of the robust scenario problem collapse to the CE cost function and constraints.

As time approach infinity, choosing K random scenarios $\theta_j \quad j = 1..K$, the chosen scenario dataset contains only the true parameter $D = \{\theta^*, ...\theta^*\}$, meaning that all the trajectories $y^{(j)}_{k+i|k}$ collapse to a singleton trajectory of the true system $y_{k+i|k}$. Additionally, we prove that the cost function is equal to the CE cost function by demonstrating that the Volume of the FPS Θ_k tends to zero.

By definition of the diameter of the FPS, we always construct a Hypercube with length $d(\Theta_k)/\sqrt{n}$ centred in θ^k_c such that C^n_k majorates Θ_k, i.e.,

$$C^n_k \subset \mathbb{R}^N = \{\theta : |\theta - \theta^k_c| \leq d(\Theta_k)/\sqrt{n}\} \tag{3.50}$$

Thanks to Theorem 1 we know that as $k \to \infty$, $d(\Theta_k) \to 0$, meaning that the hyper-cube volume $V(C^n_k) = (d(\Theta_k)/\sqrt{n})^n$ tent to zero:

$$lim_{k\to\infty}V(C^n_k) = lim_{k\to\infty}(d(\Theta_k)/\sqrt{n})^n = 0. \tag{3.51}$$

Since the cost function is continuous to the parameter θ, as θ_k approach θ^* it follows that:

$$lim_{k\to\infty}J_k(\{u_i\}^N_{i=1}; \theta_k) = J_{CE}(\{u_i\}^N_{i=1}; \theta)s, \quad \text{as } k \to \infty, \tag{3.52}$$

with

$$J_{CE}(\{u_i\}^N_{i=1}; \theta^*) = \sum_{i=0}^{N-1} Q(y_{k+i|k} - r_{k+i})^2 + Ru^2_{k+i|k}. \tag{3.53}$$

Hence, the robust MPC controller converges to the certainty-equivalent MPC controller as $k \to \infty$, completing the proof. □

3.5. SME Dual MPC Practical Implementation

This section explores the practical challenges associated with implementing the SME Dual MPC and offers potential solutions. By addressing these challenges, we aim to provide a comprehensive understanding of the algorithm, including its implementation details.

3.5.1. Forgetting Factor

As we update the FPS, we continue to intersect it with the most recent strips, which increasingly complicates the polytope over time. In principle, we could simplify the polytope by including only the more stringent strips. However, applying this method can be challenging in practice, as finding the minimal hyperplane representation of a polytope is an NP-hard problem.

A more straightforward approach is to use a forgetting factor. This involves initially enlarging the strips and subsequently removing the oldest ones. This strategy not only reduces the dependence of the feasible polytope set Θ_k on outdated data but also makes the method more suitable for systems with slowly varying parameters.

We define the importance metric for each $k-i$ strip as:

$$I(i) = \begin{cases} 1, & \text{if } i \leq i_0, \\ e^{\alpha(i_0 - i)}, & \text{if } i_0 < i \leq i_1, \\ 0, & \text{if } i > i_1. \end{cases} \tag{3.54}$$

We can then update the FPS Θ_k using the following update rule:

$$\mathcal{S}_{k-i} := \left\{ \theta \in \mathbb{R}^n \;\middle|\; \left| \tilde{y}_{k-i+1} - \varphi_{k-1}^\top \theta \right| \leq \frac{\bar{\varepsilon}}{I(i)} \right\}, \tag{3.55}$$

In this update mechanism, the most recent i_0 strips are kept unchanged, since $I(i) = 1$ for $i \leq i_0$. For intermediate data ($i_0 < i \leq i_1$), the strips become progressively less important according to an exponential decay law. Finally, strips older than i_1 are completely forgotten by setting $I(i) = 0$. Note that we can avoid 0-divisions by directly eliminating the oldest i_1 strip at each time instant.

This mechanism ensures that recent information remains tight, intermediate information is gradually relaxed, and obsolete information is eventually discarded.

In Fig. 3.3, we can see the same strip $\mathcal{S}_k$ for increasing values of $I(k)$; the figure show that by raising $I(k)$ we can control the size of $\mathcal{S}_k$, accounting for the fact that we trust newest strips more than older ones.

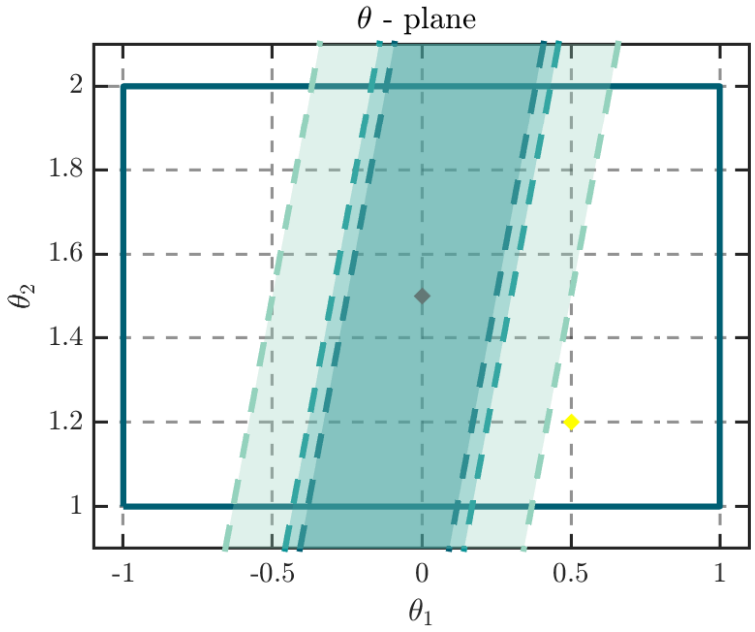


Figure 3.3: Strips $\mathcal{S}_k$ for increasing values of $I(k)$.

3.5.2. Volume minimisation and Warm-start

The volume of the FPS $V_{k+i|k}(y^{(c)}_{k+i|k}, u_{k+i|k})$ is a non-convex function which depend on the control action, and the next predicted state $y^{(c)}_{k+i|k}$, which is computed assuming the best known estimate of the true parameter $\theta^\star$, i.e the centroid at time k, θ^k_c.

The non-convexity of the problem necessitates the use of a specialised gradient-based algorithm, such as the interior-point method, to find a solution. Additionally, the initialization of the variables $\{u_i\}_{i=1}^N$ as $u_{k|k}, \dots, u_{N-1|k}$ is crucial for identifying a good solution.

On one hand, we aim for our control actions to be as close as possible to the optimal solution of the Certainty Equivalence (CE) problem, denoted as u^*, to optimise performance. On the other hand, we want to encourage exploration when the system is at equilibrium.

To achieve this, we define the initial condition of the multi-objective problem as

$$u^0 = [u^*_0 + \alpha \cdot u_{vol}, u^*_1, \dots, u^*_N]^\top. \tag{3.56}$$

Here, u_{vol} represents the control action that minimises the volume of Θ_{k+1} at the next time step, i.e.,

$$u_{vol} \in \arg\min_{\Pi_{\mathcal{U}}(u)} V_{k+1}(y^{(c)}_{k+i|k}, u). \tag{3.57}$$

In this equation, $y^{(c)}_{k+i|k}$ is computed using the current best estimate θ^k_c (central estimate) of the true parameter $\theta^\star$. This optimisation problem is single-variable and can be solved using either a bisection method or a grid-based algorithm, followed by projection into the set $\mathcal{U}$.

For polytopes in small-dimensional spaces ($\theta \in \mathbb{R}^n$ with $n \leq 5$, exact methods based on the vertex simplex decomposition can be used to compute the volume of a polytope efficiently. For polytopes in higher-dimensional space, we need to rely on approximate numerical methods based on Monte Carlo Sampling, such as log-concave sampling methods.

3.5.3. Terminal FPS

An interesting question is whether we can always gather additional information about the system through control actions, or whether there is a point at which we can no longer actively increase our knowledge.

At any given time, we can gain information actively if and only if there exists a control action u_k such that the set $\mathcal{S}_i(u_k, y^c_{k+1})$ intersects with Θ_k, i.e.,

$$\Theta_{k+1} = \mathcal{S}_{k+1}(u_k, y^{(c)}_{k+1}) \cap \Theta_k \subset \Theta_k. \tag{3.58}$$

We can also compute the convex envelope of all possible observations $\mathcal{S}_{k+1}(u)$ and define it

$$\mathcal{F}_k := \bigcap_{u \in \mathcal{U}} \mathcal{S}_k(u) = \left\{ \theta \in \mathbb{R}^n : \forall u \in \mathcal{U},\ \left| y^{(c)}_k(u) - \varphi_{k-1}(u)^\top \theta \right| \leq \bar{e} \right\}. \tag{3.59}$$

If $\Theta_k \subseteq \mathcal{F}_k$, it indicates that, for our best knowledge of the system θ^k_c it does not exist no control action $u \in \mathcal{U}$ capable of eliminating any non-empty section from Θ_k; thus, the active learning has reached its limit.

At this point, the only available gain in information comes from the error e approaching its bound $\bar{e}$ in all directions of Θ_k. At that point, we can safely turn off the volume minimisation process and focus solely on optimising the nominal cost function, which is significantly faster and can be addressed as a Quadratic Program (QP).

To ensure that we can still learn the true underlying parameter vector θ^*, we must confirm that the regressors φ_k are Persistently Exciting. One way to achieve this is by introducing a small amount of white noise into the control action:

$$u_k = u_k + \nu_k \quad \nu_k \sim \mathcal{N}(0, \sigma^2). \tag{3.60}$$

This adjustment guarantees that the feasible parameter set converges in probability to the true parameter θ^* as $k \to \infty$, leading to $\Theta_k \to \{\theta^*\}$, but we stress again that, even without the persistent excitation, the algorithm remain a robust MPC with model adaptation and guaranteed performances and risk levels.

Theorem 3.3 (Excitation via input noise)**.** *Consider a discrete-time linear time-invariant system of ARX form:*

$$y_{k+1} = \theta^{*\top} \varphi_k + e_k, \tag{3.61}$$

Assume the following:

1. *The system is BIBO stable.*
2. *The control input is of the form $u_k = \bar{u_k} + \nu_k$, where $\nu_k \sim \mathcal{N}(0, \sigma^2)$ is a zero-mean white Gaussian noise.*
3. *The noise process ν_k is independent of the measurement noise e_k, i.e. $\nu_k \perp e_k$.*
4. *without loss of generalisation, assume that the system is controlled at the $(u, y) = (0, 0)$ equilibrium point.*

Then the regressor sequence $\{\varphi_k\}$ is persistently exciting **in expectation**, *i.e., there exists $\alpha > 0$ and a window length $m \in \mathbb{N}$ such that:*

$$\mathbb{E}\left[\frac{1}{m} \sum_{i=k}^{k+m-1} \varphi_i \varphi_i^\top\right] \succeq \alpha I \quad \forall k. \tag{3.62}$$

Proof. Under the asymptotic stability assumption, the output y_k is a stable linear combination of past inputs and outputs. Since $u_k = \nu_k$ is white noise with positive variance, the excitation propagates through the system dynamics, and each y_k becomes a filtered (stable) combination of the noise terms ν_{k-i}. Hence, the regressor φ_k can be expressed as a linear combination of past ν_i, i.e.,

$$\varphi_k = \mathcal{H} \cdot \nu_k^{\text{hist}}, \quad \text{with } \mathcal{H} \text{ linear operator}, \tag{3.63}$$

where $\nu_k^{\text{hist}} := [\nu_k, \nu_{k-1}, \dots]^\top$ input noise history, and $\mathcal{H}$ contains the impulse response coefficients that generate φ_k from ν_k.

Because ν_k is zero-mean and independent, the induced process φ_k has a full-rank covari-

ance matrix, i.e.,

$$R := \mathbb{E}[\varphi_k \varphi_k^\top] = \mathcal{H}\,\mathbb{E}[\nu_k^{\text{hist}}(\nu_k^{\text{hist}})^\top]\,\mathcal{H}^\top = \sigma^2 \mathcal{H}\mathcal{H}^\top \succ 0, \tag{3.64}$$

since $\mathbb{E}[\nu_k \nu_k^\top] = \sigma^2 I$ and $\mathcal{H}$ is full-rank due to the system structure and the input entering all modes.

Finally, since φ_k is stationary (the system is LTI and the excitation is white noise), the time-averaged Gramian satisfies

$$\mathbb{E}\left[\frac{1}{m}\sum_{i=k}^{k+m-1} \varphi_i \varphi_i^\top\right] = R \succeq \sigma^2 I, \tag{3.65}$$

for some $\alpha > \sigma^2 > 0$ and $m > 0$, concluding the proof. □

3.5.4. Arx Assumption and structural model mismatch

The previously proven convergence guarantees on the FPS are valid if the true system belongs to the ARX class, which is rarely the case for real systems with non-linearities and non-idealities. However, if the mismatch between the underlying non-linear system generating the data and the ARX model approximation can be interpreted as a UBB disturbance, the theoretical framework still applies. In such cases, the knowledge of the error bound ϵ must account for this uncertainty, which could make the strings too wide, making it difficult for the FPS to shrink and leading to a more conservative controller. This fact does not invalidate the violation probability guarantees, which enable the safe use of the controller in real-world scenarios. At the same time, the algorithm improves the model knowledge, but it can result in conservative behaviour and reduced closed-loop performance.

3.5.5. MPC algorithm

The full algorithm is:

Algorithm 3.2 SME Dual MPC with Active Exploration

Require: Initial regressor vector φ_0, initial FPS $\mathcal{K}_0$, initial central estimate θ_c^0, exploration flag `volFlag` $\leftarrow$ `true`

for $k = 0, 1, 2, \ldots$ **do**
  if `volFlag` == `true` **then**
    Solve the *volume-minimizing* MPC problem ▷ see Eq.(3.27)
    Compute control input: $u_k = u^*_{k|k}$
  else
    Solve the *nominal* MPC problem ▷ see Eq. (3.18)
    Add excitation noise: $u_k = u^*_{k|k} + \nu_k$
  end if
  Apply the control input u_k
  Measure the system output $\tilde{y}_{k+1}$
  Compute the corresponding strip:

$$\mathcal{S}_{k+1} = \left\{\theta \in \mathbb{R}^n : \left|\tilde{y}_{k+1} - \varphi_k^\top \theta\right| \leq \bar{e}\right\} \tag{3.66}$$

  Update the FPS: $\Theta_{k+1} \leftarrow \Theta_k \cap \mathcal{S}_{k+1}$
  if $\Theta_{k+1} \subseteq \mathcal{F}_k$ **then**
    Set `volFlag` $\leftarrow$ `false` ▷ Switch to nominal MPC
  end if
end for

3.6. Numerical example

In this section, we present a simple control task to grasp the method's geometrical intuition better. Consider the first-order linear time-invariant (LTI) system, i.e.,

$$G(s) = \frac{\mu}{\tau s + 1}, \tag{3.67}$$

where $\mu > 0$ denotes the steady-state gain and $\tau > 0$ is the time constant; both parameters are assumed to be unknown but constant.

Applying a Zero-Order Hold (ZOH) discretisation with sampling period T_s, we obtain

$$x_{k+1} = ax_k + bu_k, \qquad y_k = x_k, \tag{3.68}$$

where

$$a = e^{-T_s/\tau}, \qquad b = \mu(1 - e^{-T_s/\tau}). \tag{3.69}$$

Hence, the discrete-time transfer function reads

$$G(z^{-1}) = \frac{bz^{-1}}{1 - az^{-1}}. \tag{3.70}$$

Equation (3.68) can be rewritten in the ARX form:

$$y(k+1) = \theta^\top \phi(k) + \epsilon(k), \tag{3.71}$$

where

$$\theta = \begin{bmatrix} -a \\ b \end{bmatrix}, \qquad \phi(k) = \begin{bmatrix} y(k) \\ u(k) \end{bmatrix}, \tag{3.72}$$

and $\epsilon(k)$ represents the modelling or measurement uncertainty, assumed UBB. This compact form provides a convenient basis as all dynamics are parameterised by the vector θ. Also, the presence of only two parameters $n = na + nb = 2$ is beneficial for visualising the FPS and strips, which lie on the (θ_1, θ_2) plane.

3.6.1. Stability Region in Parameter Space

For the discrete-time system (3.70), asymptotic stability requires that

$$|a| < 1. \tag{3.73}$$

Expressing this condition in terms of the ARX parameter vector $\theta = [-a,\, b]^\top$ yields to

$$\mathcal{K}_\prime = \left\{ \theta \in \mathbb{R}^2 \;\middle|\; -1 < -\theta_1 < 1 \right\}, \tag{3.74}$$

that is, $\theta_1 \in (-1, 1)$. The gain parameter b does not affect the pole location and therefore does not influence stability. However, for simplicity, we made the hypothesis of knowing the sign of the gain and an upper value, i.e.,

$$\mathcal{K}_\prime = \left\{ \theta \in \mathbb{R}^2 \;\middle|\; -1 < -\theta_1 < 1, \underline{\mu} < \theta_2 \leq \bar{\mu} \right\}. \tag{3.75}$$

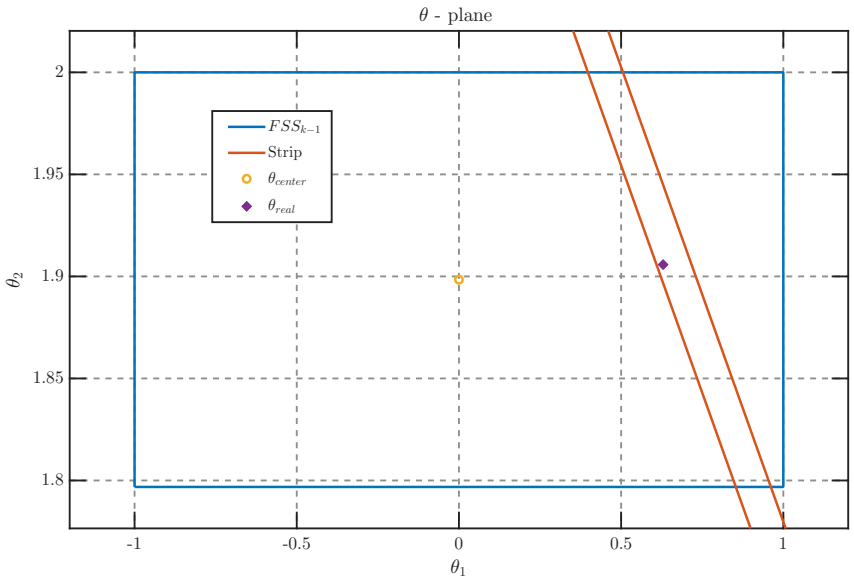


Figure 3.4: Initial FPS $\mathcal{K}$, region in the (θ_1, θ_2) parameter space, corresponding to the inequality (3.75). Real $\theta^\star$(purple cross) and central estimate θ_c^k (yellow dot). First realised strip $\mathcal{S}_1$

To simulate the discrete dynamics, we consider the recursive equation

$$\begin{cases} y(0) = 0, \\ y(k+1) = -a\, y(k) + b\, u(k), \end{cases} \tag{3.76}$$

where $u(k)$ is a known input sequence, e.g. a unit-step $u(k) = 1$.

3.6.2. Set-Based Representation: Strips and Feasible Parameter Set (FPS)

In the presence of measurement noise, model (3.1) can be expressed as

$$y(k+1) = \theta_1\, y(k) + \theta_2\, u(k) + \epsilon(k), \qquad |\varepsilon(k)| \leq \bar{\epsilon}. \tag{3.77}$$

At each time instant k, after we measure the newest data $y(k)$, the pair $(y(k), \varphi(k-1) = [y(k-1), u(k-1)]^\top)$ defines a new strip in the FPS, i.e.,

$$\mathcal{S}_k = \left\{ \theta \in \mathbb{R}^2 \;\middle|\; |y(k) - \theta_1\, y(k-1) - \theta_2\, u(k-1))| \leq \bar{\epsilon} \right\}. \tag{3.78}$$

The FPS at time k is the intersection of all strips up to that instant

$$\Theta_k = \bigcap_{i=1}^{k} \mathcal{S}_i. \tag{3.79}$$

This intersection progressively narrows the admissible region for the unknown parameters (τ, μ), yielding guaranteed bounds consistent with the available data.

As a remark, from Eq. (3.78) it is clear that, if $y(k-1)$ is fixed, the only available degree of freedom is the control action $u(k-1)$, rotating the strip (counterclockwise) as it increases.

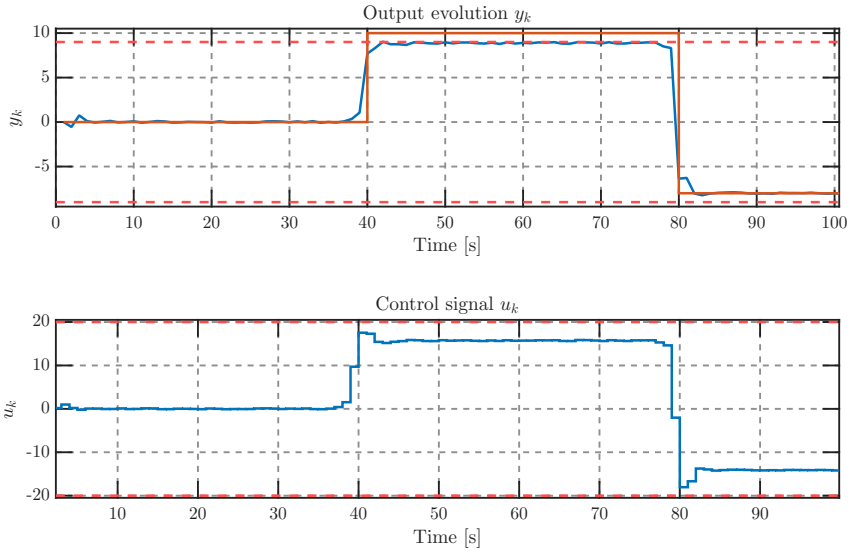


Figure 3.5: Output y_k and input u_k closed-loop evolution (blue line), reference value $y_{ref,k}$(orange line), output constraints (red-dotted line).

In Fig. 3.5, the trajectory of the control system is depicted. When the system is at equilibrium, the MPC actively perturbs the system to improve the plant's knowledge.

Finally, Fig. 3.6 shows the parameter trajectory. Here, we can see that in the first 5s, the system reaches values close to the true underlying parameters $\theta^\star = [\theta_1^\star, \theta_1^\star]^\top$

Chapter 7 presents the simulation results and a case study in which the Dual SME-MPC is implemented to control Thermo-Electric Refrigerators for providing aggregated Demand Response. In the latter, the active and passive learning algorithms are tested against each other, demonstrating the need for active probing when controlling an uncertain system and thus the superiority of the Dual SME-MPC.

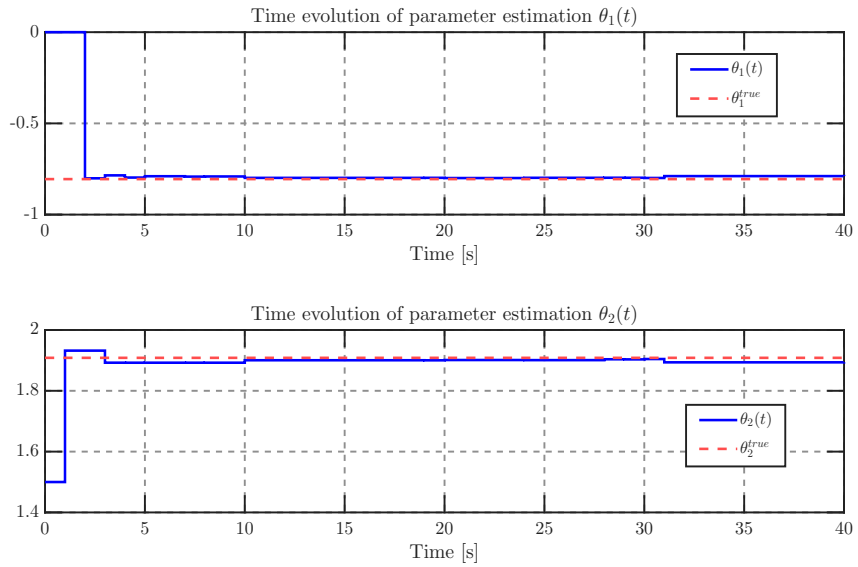


Figure 3.6: Parameter θ_1^k, θ_2^k trajectory (blue line). Real value $\theta_1^\star, \theta_2^s tar$ (red-dotted line).

Chapter 3 Bibliography

[Åst65] Karl J. Åström. "Optimal Control of Markov Processes with Incomplete State Information". In: *Journal of Mathematical Analysis and Applications* 10.1 (1965), pp. 174–205. DOI: 10.1016/0022-247X(65)90154-X.

[Bel57] Richard Bellman. *Dynamic Programming*. Princeton, NJ: Princeton University Press, 1957. ISBN: 978-0-691-07951-6. DOI: 10.1515/9781400874668.

[BM15] Franco Blanchini and Stefano Miani. "Set-theoretic methods in control". In: *Systems and Control: Foundations and Applications* 0 (2015), pp. 1–630.

[BR56] Richard Bellman and M. Roth. "A Problem in the Sequential Design of Experiments". In: *Econometrica* 24.1 (1956), pp. 39–57. DOI: 10.2307/1907743.

[Fel63] Á.Á. Feldbâum. "Dual control theory problems". In: *IFAC Proceedings Volumes* 1.2 (1963). 2nd International IFAC Congress on Automatic and Remote Control: Theory, Basle, Switzerland, 1963, pp. 541–550. ISSN: 1474-6670. DOI: https://doi.org/10.1016/S1474-6670(17)69687-3.

[KH16] Eliot D. Klenske and Philipp Hennig. "Dual Control for Approximate Bayesian Reinforcement Learning". In: *Journal of Machine Learning Research* 17.6 (2016), pp. 1–30. URL: https://jmlr.org/papers/v17/14-552.html.

[MB82] Mario Milanese and Gustavo Belforte. "Estimation Theory and Uncertainty Intervals Evaluation in Presence of Unknown But Bounded Errors: Linear Families of Models and Estimators". In: *IEEE Transactions on Automatic Control* 27.2 (1982), pp. 408–414. DOI: 10.1109/TAC.1982.1102926.

[Mes18] Ali Mesbah. "Stochastic model predictive control with active uncertainty learning: A Survey on dual control". In: *Annual Reviews in Control* 45 (2018), pp. 107–117. ISSN: 1367-5788. DOI: https://doi.org/10.1016/j.arcontrol.2017.11.001.

[Mez+25] Erick Meza et al. "Adaptive Scenario-Based Predictive Control: A Set Membership Learning Approach". In: *IFAC-PapersOnLine* 59.26 (2025). 7th IFAC Conference on Intelligent Control and Automation Sciences ICONS 2025, pp. 323–328. ISSN: 2405-8963. DOI: https://doi.org/10.1016/j.ifacol.2025.12.055. URL: https://www.sciencedirect.com/science/article/pii/S2405896325027314.

[MV91] M. Milanese and A. Vicino. "Optimal estimation theory for dynamic systems with set membership uncertainty: An overview". In: *Automatica* 27.6 (1991), pp. 997–1009. ISSN: 0005-1098. DOI: https://doi.org/10.1016/0005-1098(91)90134-N.

[Par+23] Anilkumar Parsi et al. "Dual adaptive MPC using an exact set-membership reformulation". In: *IFAC-PapersOnLine* 56.2 (2023). 22nd IFAC World Congress, pp. 8457–8463. ISSN: 2405-8963. DOI: https://doi.org/10.1016/j.ifacol.2023.10.1132.

[SFM13] Georg Schildbach, Lorenzo Fagiano, and Manfred Morari. "Randomized Solutions to Convex Programs with Multiple Chance Constraints". In: *SIAM Journal on Optimization* 23.4 (2013), pp. 2479–2501. DOI: 10.1137/120878719.

[Wit95] Björn Wittenmark. "Adaptive Dual Control Methods: An Overview". In: *IFAC Proceedings Volumes* 28.13 (1995). 5th IFAC Symposium on Adaptive Systems in Control and Signal Processing 1995, Budapest, Hungary, 14-16 June, 1995, pp. 67–72. ISSN: 1474-6670. DOI: https://doi.org/10.1016/S1474-6670(17)45327-4.

4 | Smart Grid Energy Management

The previous chapters established the theoretical foundations for handling uncertainty in distributed optimisation and optimal control.

This chapter presents the general Energy Management Problem, which refers to the coordinated scheduling and control of Distributed Energy Resources (DERs), such as storage units, flexible loads, and renewable generation, with the objective of optimising system performance while ensuring reliability, efficiency, and sustainability. In the Smart Grid context, as presented in chapter 1, the Energy Management System (EMS) serves as the supervisory layer responsible for balancing supply and demand at local scale, allocating resources, and mitigating the impacts of uncertainty in non-dispatchable generation and consumption, communication delays, and dynamic network structures.

This chapter formalises the mathematical framework of the Energy Management Problem. It highlights how its structure naturally aligns with the distributed and robust optimisation methods developed in the previous chapters 2, 3. The formulation presented here serves as the basis for the hierarchical and distributed control architectures discussed in chapter 5 and 7.

4.1. European Energy Market

The operation of any EMS is tightly coupled with the structure and rules of the electricity market in which it operates. In Europe, the energy market plays a central role in coordinating the generation, consumption, and exchange of electricity among a large number of heterogeneous participants, from traditional centralised plants to distributed renewable generators and prosumers. The European electricity market is not only a medium for trading energy but also a complex system that must ensure stability, resilience, and real-time operational efficiency. In fact, unlike traditional grids, comprised of large centralised plants and unidirectional flows, the European energy system has evolved to accommodate

prosumers and distributed resources actively participating in energy exchanges. Understanding the European market framework is therefore fundamental for designing effective control and optimisation mechanisms that align with market incentives, regulatory constraints, and real-time operational requirements, to enhance smart-grid operations and exploit flexibility reliably.

The energy market comprises several layers, each operating on a different time scale: the Day-Ahead and Intraday markets, which manage energy trading; the Balancing Market, which maintains real-time system balance; and the Ancillary Services Markets, which procure additional stability services. These layered markets provide the operational context in which distributed energy resources and flexibility services can be monetised and managed.

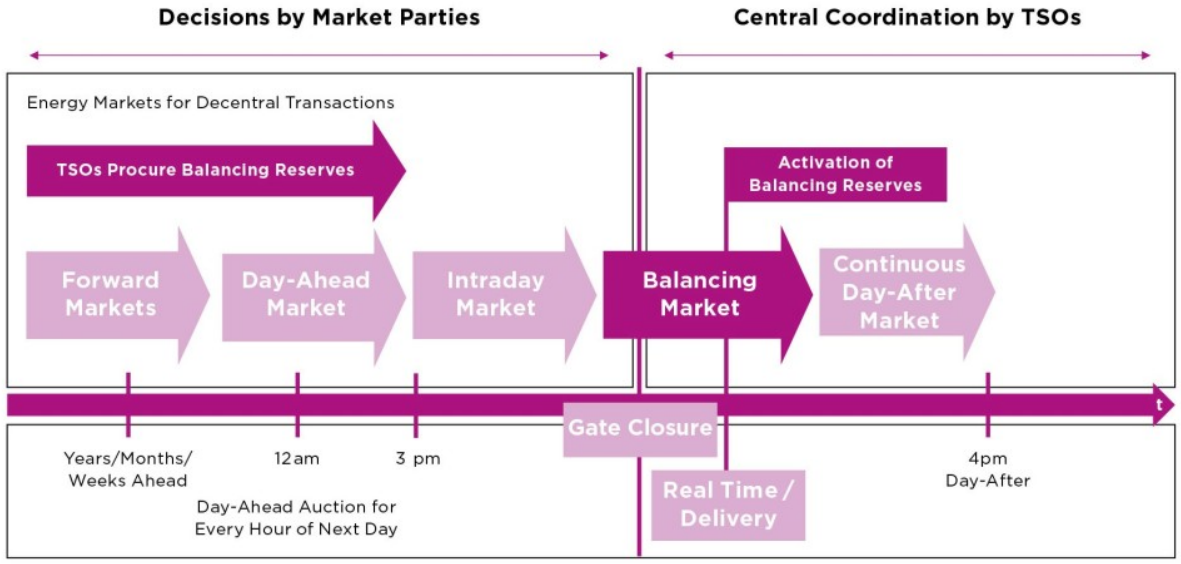


Figure 4.1: Overview of the main electricity markets in Europe and their temporal sequence. [Gmb24].

The diagram in Fig. 4.1 provides a concise representation of the temporal organisation of the European electricity markets. The structure is based on a sequence of partially overlapping trading layers, each operating at a different time scale and with a specific operational objective.

The process begins with long-term contracts and forward markets, where participants hedge price risks and secure capacity months (even years) ahead. The Day-Ahead Market (DAM) follows, determining the baseline energy schedules for the next operating day through marginal pricing mechanisms. Subsequently, the Intraday Market (IDM) allows participants to continuously adjust their positions as updated forecasts for demand and

renewable generation become available. In real-time, the Balancing Market (BM), operated by Distribution System Operators (DSO), ensures system stability by activating upward or downward reserves when deviations occur between scheduled and actual power flows. Finally, Ancillary Services Market (ASM) provide additional system support functions such as frequency containment, automatic and manual frequency restoration, and voltage control.

This layered market architecture enables progressive refinement of the system's energy balance as delivery time approaches, balancing market efficiency and operational reliability. Understanding the interactions among these markets is fundamental for the design of advanced EMS: the hierarchical and distributed control frameworks proposed in this thesis are conceived to operate coherently with these market structures and timescales.

4.1.1. Day-Ahead Market (DAM)

The DAM is the cornerstone of the European electricity trading mechanism. Every day, a few hours before delivery (typically noon of the day before), all participants submit bids and offers for each hourly slot (and in some countries for 15-minute slots) of the following day. The market clears by matching supply and demand through marginal pricing: the highest accepted offer determines the market-clearing price, which is then applied to all accepted offers. The DAM thus establishes the baseline power production and consumption schedule [Sha+20; Klæ+22].

4.1.2. Intraday Market (IDM)

Owing to the inherent uncertainties in generation, especially from renewable sources, and in consumption patterns, the Day-Ahead schedule must often be adjusted in real time. The IDM provides the mechanism to perform such adjustments, operating from the closure of the DAM up to real-time delivery. In many European market areas, the IDM trades products with a temporal resolution of 15 or 30 minutes, allowing participants to fine-tune their positions to correct forecast errors and achieve a more flexible re-balancing of supply and demand. Depending on the regional market design, the IDM operates either through continuous trading, where bids and offers are matched as they arrive, or via discrete intraday auctions held at regular intervals [OCo+25; Šro+24].

4.1.3. Balancing Market (BM)

The DSO manages the BM to maintain the real-time balance and stability of the power system. When deviations occur between scheduled and actual generation or demand, the DSO activates balancing services provided by market participants offering upward or downward regulation capacity. These real-time operations are essential to correct unexpected imbalances between supply and demand and preserve system stability. The pricing of such services may follow a pay-as-bid mechanism or a marginal pricing scheme, depending on the national market design [Klæ+22]. Settlement in the balancing market is typically performed by charging or crediting participants based on imbalance prices and the activation of balancing services, ensuring that market participants are incentivised to maintain accurate forecasts and reliable operations [Kau+19].

4.1.4. Ancillary Services Market (ASM)

Beyond the balancing of active power, the stability and security of the electrical grid rely on a set of additional services collectively known as Ancillary Services. These are traded through the ASM, which ensures that TSO can access reserves and capabilities required for reliable system operation. Among the most important ancillary services are those dedicated to frequency control—namely, the Frequency Containment Reserve (FCR), the Automatic Frequency Restoration Reserve (aFRR), and the Manual Frequency Restoration Reserve (mFRR), which collectively ensure frequency stability under varying load and generation conditions [Pol20; Kau+19]. Voltage support services are also procured to manage reactive power flows and maintain acceptable voltage levels across the network [Kau+19]. Additionally, black-start capability is a critical service that enables power system restoration after a total or partial blackout, ensuring resilience against large-scale outages [Kau+19].

4.2. Hierarchical Architecture for Frequency Control and Restoration

While electricity markets coordinate the economic dispatch of resources across Europe, the physical power system must continuously ensure that supply and demand remain balanced in real time. Market-based mechanisms govern electricity trading, but they rely on a robust technical control framework to ensure operational stability and security.

In practice, this coupling between economic and physical layers is achieved through a hierarchical control architecture, coordinated by TSO and DSO. Such architecture regulates

system frequency, voltage, and reserve deployment across multiple timescales—from milliseconds to hours—ensuring both reliability and cost-efficiency in grid operations [ENT23; Reb+07].

The following section describes the hierarchical structure for frequency control and restoration in the European power system, highlighting the role of each control layer and its relationship with market mechanisms.

The European interconnected power system operates at a nominal frequency of 50 Hz. To maintain this frequency within acceptable limits, a hierarchical control scheme is applied. The structure is based on three distinct but coordinated layers: **primary**, **secondary**, and **tertiary control**. These layers act across different temporal and spatial scales, from the sub-second response of local controllers to the multi-hour optimisation of dispatch and market coordination [ENT23; Reb+07].

Primary Control, also known as FCR, consists of active power reserves that are immediately available in the event of a frequency deviation. These reserves act as the system's first line of defence against frequency drift following an imbalance, providing automatic stabilisation. Operators require delivery of 50% of the reserve within 15 s and full activation within 30 s, with the ability to sustain the service for up to 15 minutes.

Secondary Control, or aFRR, comprises active power reserves used to restore the system frequency to its nominal value and bring power exchanges back to scheduled levels. This control layer is automatically activated, typically within 5 to 15 minutes, and can remain active for up to 15 minutes, allowing the system to recover from longer disturbances.

Finally, **Tertiary Control**, also known as Replacement Reserves (RR), is activated to replenish the secondary control capacity and prepare the system for future imbalances. The activation time for tertiary reserves is approximately 30 minutes. These reserves support the gradual re-optimisation of operating setpoints and allow the grid to return to an economically optimal operating state despite persistent uncertainties [ENT23].

In Fig. 4.2, the hierarchical control architecture based on the European Network of Transmission System Operators for Electricity (ENDSO-E) coordination across member states is depicted. With the emergence of smart grids and digital platforms, distributed resources such as residential batteries, electric vehicles, and flexible loads are progressively participating in all three control layers, transforming consumers into active market participants [ENT23; Sch+80].

Conceptually, there is an overlap in functions between the Tertiary control level, which provides the high-level coordination of generation, storage, and demand-side resources,

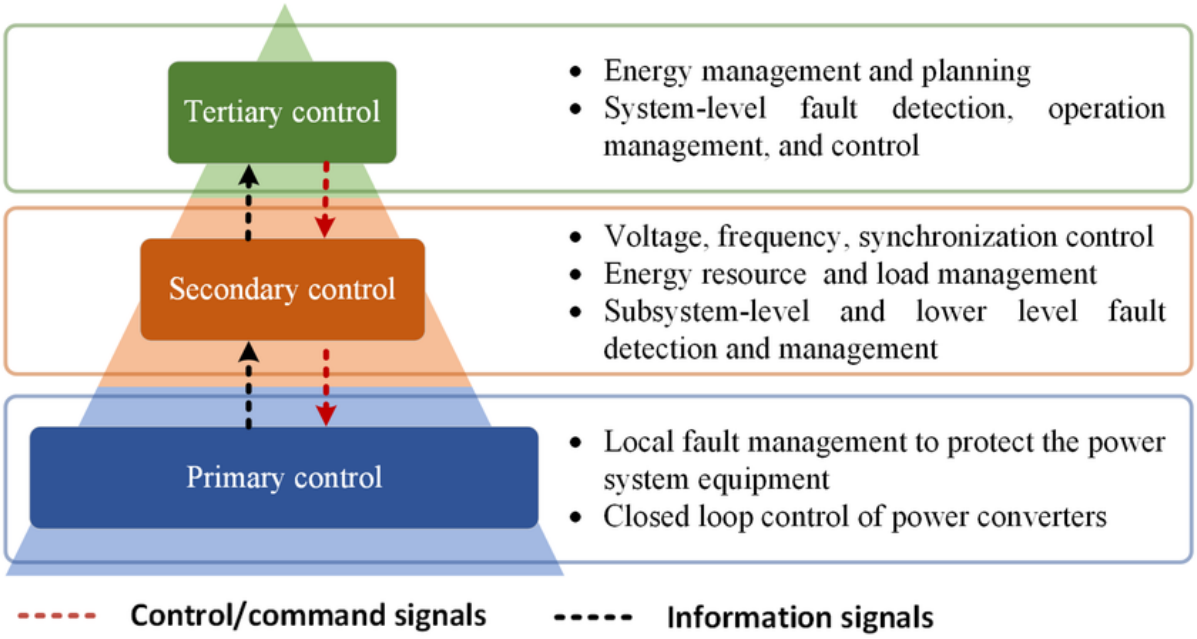


Figure 4.2: Hierarchical control architecture. [Sah+23].

and the EMS, particularly those designed for MGs or distributed VPP.

The EMS is the operational implementation of tertiary control, dispatching power across microgrids and VPP by translating tertiary control objectives and market signals into asset-level control actions, ensuring that both system-level constraints and local operational limits are satisfied.

4.3. Energy Management Systems

Historically, power grid coordination has relied on a centralised architecture. This was necessary because power was generated in dislocated power plants, geographically distant from cities where most power was consumed, meaning that power flows from large power plants in rural areas to numerous loads in cities.

While this centralised architecture optimises operations by providing complete information on production and consumption and enabling global coordination, it also has significant drawbacks: it becomes increasingly complex as the number of participants grows and their characteristics become heterogeneous, and it suffers from single points of failure.

These challenges are central to the debate on which EMS architecture is most suitable for future smart grids. The evaluation of EMS architectures must be carried out not only in terms of nominal efficiency, but also by assessing how each framework copes with the

multiple uncertainties inherent in modern power systems. Designing EMS solutions that are both optimal and robust is a prerequisite for the reliable integration of renewable sources, the stability of microgrids, and the efficiency of power exchanges in increasingly complex and dynamic networks.

4.3.1. Architectures for EMSs

The increasing complexity of the modern Smart grid has led to the development of various architectures for *Energy Management System (EMS)*. Broadly, three main architectures are commonly adopted: **hierarchical**, **decentralised**, and **distributed**. Each approach has specific advantages and drawbacks, depending on the scale, technological maturity, and system objectives.

A *hierarchical* EMS architecture decomposes the overall control problem into multiple levels, typically including a supervisory layer, a coordination layer, and a local control layer. The supervisory level provides long-term planning and global objectives, such as minimising total energy costs or carbon emissions. The coordination layer ensures consistency among interconnected subsystems by exchanging aggregate information. Finally, local controllers operate individual DERs or microgrids based on locally available data and implement real-time control actions. This layered structure enhances scalability and facilitates modular system design while maintaining global optimality through top-down coordination.

In contrast, *distributed* and *decentralised* EMS frameworks aim to reduce dependence on a central authority by enabling agents—each representing a microgrid, a building, or a DERs—to make decisions through local computation and limited information exchange. In distributed schemes, agents cooperatively solve an optimisation problem using consensus-based or decomposition techniques, ensuring convergence to a global or near-global optimum. Decentralised architectures further relax communication requirements by relying solely on local measurements and control rules, without explicit coordination. These approaches are desirable for large-scale systems characterised by communication delays, privacy concerns, and heterogeneous ownership structures.

The integration of hierarchical and distributed paradigms has recently emerged as a promising direction for Smart Grid (SG) control. Hierarchical distributed architectures combine the benefits of hierarchical information flow and multi-time-scale operation, with distributed optimisation at each layer, enabling coordination among agents while preserving autonomy. Such hybrid frameworks provide an effective trade-off between optimality, robustness, and scalability, making them suitable for next-generation smart grids where

flexibility, resilience, and cybersecurity are paramount.

Hierarchical Architecture

In a hierarchical EMS, control is organised in a layered structure. Higher levels perform long-term optimisation and coordination, while lower levels handle local dynamics and fast responses. The traditional power system control paradigm inspires this architecture.

Advantages and disadvantages: The hierarchical control paradigm in energy management systems features distinct layers: primary, secondary, and tertiary control. The primary layer focuses on local stability and regulation, the secondary layer restores nominal conditions, and the tertiary layer optimises coordination across the network. This structure enables high-level subsystem coordination and aligns well with existing regulatory frameworks, making integration into current grids manageable.

However, hierarchical control has limitations. Centralised decision-making at higher levels can create scalability issues as systems grow. Effective operation relies on a reliable communication infrastructure, which can be costly and challenging, especially in diverse networks. Additionally, this structure can create single points of failure, risking system resilience during communication breakdowns or cyber disruptions.

Decentralized Architecture

In a decentralised EMS, each agent (e.g., microgrid, generator, or prosumer) optimises its operation independently, based on local objectives and constraints, with little or no communication with other agents.

Advantages and disadvantages: Decentralised architecture exhibits several advantages that enhance its performance and reliability. First, it is characterised by high robustness and resilience, ensuring that there is no single point of failure. This design feature significantly reduces dependency on the communication infrastructure. Additionally, it provides the flexibility to operate autonomously in the event of grid disconnection, such as in islanded microgrids.

On the other hand, there are notable disadvantages to consider. The lack of global coordination may result in suboptimal solutions at the system level. This can lead to potential conflicts between local and global objectives, such as balancing price and stability. Furthermore, it can be challenging to ensure system-wide constraints, including maintaining voltage and frequency stability across the network.

Distributed Architecture

A distributed EMS lies between hierarchical and decentralised approaches. In this case, multiple agents optimise their local operation while exchanging limited information with neighbours or a coordinating entity. This enables **consensus-based optimisation** and cooperative control strategies.

Advantages and disadvantages: Distributed architecture is scalable and flexible, therefore it is well-suited for large-scale smart grids. It allows local autonomy while still ensuring global coordination, helping reduce the computational burden compared to a fully centralised system. Additionally, this architecture naturally aligns with multi-agent and peer-to-peer trading paradigms.

However, there are also some disadvantages to consider. It requires reliable communication, and convergence of distributed optimisation algorithms can be slow. The design is notably complex, as it necessitates careful definition of protocols and consensus mechanisms to maintain the system's safety in the presence of adversarial agents.

Table 4.1 summarises the main characteristics of the three EMS architectures, and Fig. 4.3 depicts the three possible EMS architectures.

Architecture	Advantages	Disadvantages
Hierarchical	Global optimality, clear structure, compatible with existing systems	Scalability issues, high communication needs, single point of failure
Decentralised	Robust, no central coordinator, autonomous operation	Suboptimal solutions, local vs. global conflicts, hard to enforce system-wide constraints
Distributed	Balance between autonomy and coordination, scalable, supports peer-to-peer	Communication and algorithmic complexity, convergence speed issues

Table 4.1: Comparison of EMS architectures.

4.3.2. Microgrid Energy Management preliminaries

A microgrid is a localised energy system that integrates DERs within a clearly defined electrical boundary. It can operate either connected to the main grid or in islanded mode, maintaining power supply to local users even during grid outages. By combining energy generation, storage, and demand-side management, a microgrid enhances energy

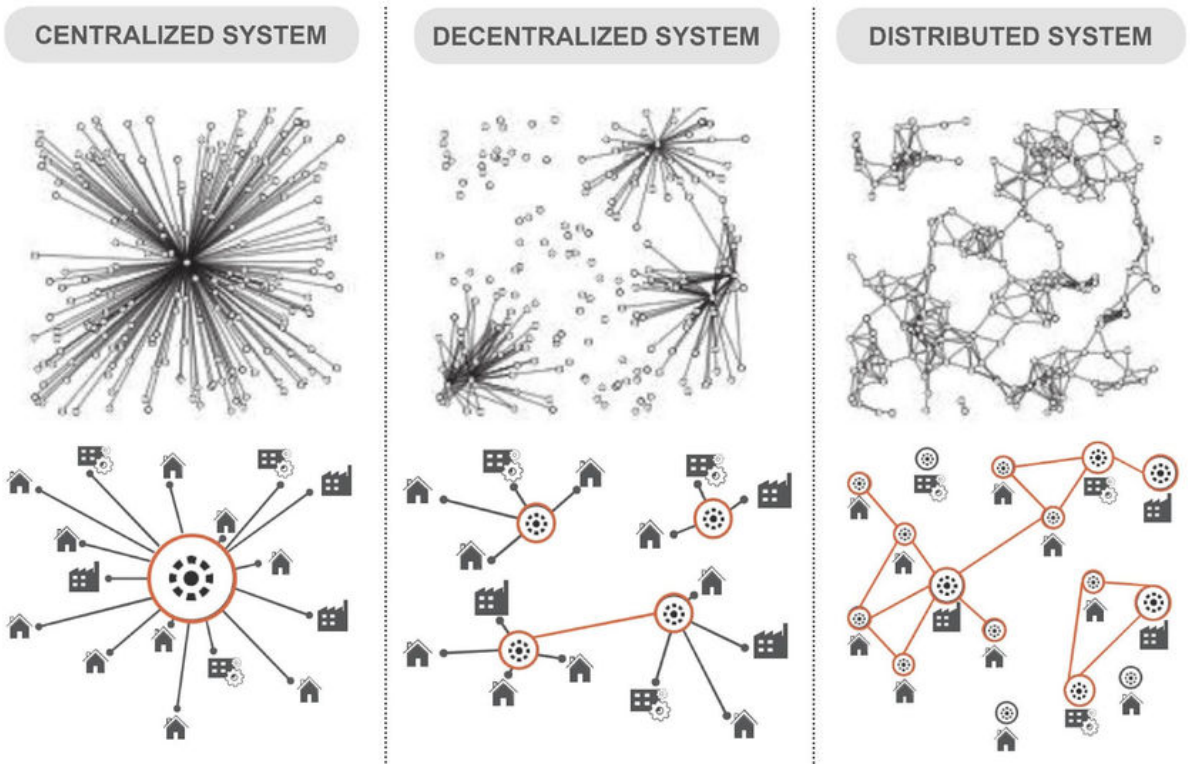


Figure 4.3: Hierarchical, Decentralised, and Distributed EMS control architecture. [PS17].

resilience, efficiency, and sustainability. It also enables the integration of renewable resources at a community scale.

From a control and optimisation perspective, a microgrid functions as a cyber-physical system, where energy flows and information exchange are closely linked. The EMS of a microgrid coordinates the operation of its components to minimise costs, reduce emissions, and maintain real-time power balance. Depending on its architecture, decision-making within the microgrid can be centralised, distributed, or decentralised, offering varying degrees of autonomy for individual units.

Microgrids are fundamental components of the smart grid, providing flexibility, supporting ancillary services, and facilitating the transition to low-carbon energy systems. Their modular design enables scalable deployment and the formation of interconnected microgrid networks, paving the way for future self-organising, resilient energy ecosystems.

In a grid-connected microgrid, the EMS coordinates the dispatch of active power across subsystems by computing optimal set points for controllable units, such as distributed generators and storage devices, to meet power demand at minimum cost. The microgrid is assumed to operate in parallel with the utility grid, which serves as a balancing resource to mitigate forecast errors and maintain power quality. Under this assumption, which is

standard for non-islanded microgrids (see e.g. [Rai+18; Han+20; Han+15; Mel+21]), the network is considered to have sufficient reactive support for local voltage regulation.

The system under study includes the following components:

- **Thermal distributed generators**, which supply the baseline active power P_{gen}.
- **Photovoltaic panels**, providing renewable generation P_{pv} during daylight hours.
- **Battery Energy Storage System (BESS)**, capable of charging or discharging with power P_{bess}.
- **Building load**, representing the local demand P_{load}.
- **Grid interface**, which exchanges power P_{grid} with the main distribution network.

The power balance constraint ensures that, at each discrete time instant $t = k \cdot \tau \quad k = 0, 1, 2..$, with sampling time τ, generation and imports match the net demand, i.e.,

$$P_{gen}(t) + P_{bess}(t) + P_{grid}(t) = P_{load}(t) - P_{pv}(t). \tag{4.1}$$

Thermal Generation Unit

Standard operational constraints characterise the thermal generator:

1. **Capacity limits:**

$$\begin{cases} 0 < P_{gen}(t) \leq P_{gen}^{max}, & \text{if } isON(t) = 1, \quad (4.2a) \\ P_{gen}(t) = 0, & \text{if } isON(t) = 0, \quad (4.2b) \end{cases}$$

   where $isON(t)$ is a binary status variable.

2. **Ramping constraint:**

$$|P_{gen}(t) - P_{gen}(t-1)| \leq P_{ramp}^{max}. \tag{4.3}$$

The operating cost of the generator is represented by a piecewise affine function c_{gen} of the output power:

$$C_{gen} = \sum_{t=1}^{T} \tau \cdot c_{gen}(P_{gen}(t), t). \tag{4.4}$$

Battery Energy Storage System

The BESS State Of Charge (SOC), denoted $s(t)$, is a dimensionless state variable describing the amount of energy stored in the BESS as a percentage of the maximum capacity Q_c; it evolves according to the charging/discharging power commands $u_{CH}(t)$ and $u_{DCH}(t)$ (in [KW]), with gains $k_{CH} = k_{ch}^{eff} \cdot \frac{\tau}{Q_c}$ and $k_{DCH} = \frac{\tau}{k_{dch}^{eff} \cdot Q_c}$ depending on the sampling time τ, charging and discharging efficiencies $k_{ch}^{eff}, k_{dch}^{eff}$, and maximum BESS capacity Q_c. The discrete-time dynamical system is

$$\begin{cases} s(t+1) = s(t) - k_{DCH}\, u_{DCH}(t) + k_{CH}\, u_{CH}(t), & \text{(4.5a)} \\ s(0) = s_0. & \text{(4.5b)} \end{cases}$$

Binary variable $b(t)$ prevents simultaneous charging and discharging,

$$\begin{cases} 0 \leq u_{DCH}(t) \leq u_{DCH}^{max}\, b(t), & \text{(4.6a)} \\ 0 \leq u_{CH}(t) \leq u_{CH}^{max}\, (1 - b(t)). & \text{(4.6b)} \end{cases}$$

The net exchanged power is

$$P_{bess}(t) = u_{DCH}(t) - u_{CH}(t), \tag{4.7}$$

and it is bounded by

$$-u_{CH}^{max} \leq P_{bess}(t) \leq u_{DCH}^{max}. \tag{4.8}$$

Finally, the state of charge of the battery must remain within safe operational limits:

$$s(t) \in [s_{min}, s_{max}]. \tag{4.9}$$

Net Demand

The net demand is defined as the difference between the inelastic load and the uncontrollable PV production:

$$P_{net}(t) = P_{load}(t) - P_{pv}(t), \tag{4.10}$$

for which historical data are assumed to be available.

Interaction with the Utility Grid

After local set-points for $P_{gen}(t)$ and $P_{bess}(t)$ are established, the remaining imbalance is compensated via the utility grid. This means that the microgrid may either purchase or sell electricity at the market price, depending on whether $P_{grid}(t)$ is positive (import) or negative (export). The resulting cost of exchanged energy is given by

$$C_{MG} = \sum_{t=1}^{T} \tau \, c_{buy}(t) \, \max(0, P_{grid}(t)) + \sum_{t=1}^{T} \tau \, c_{sell}(t) \, \min(0, P_{grid}(t)), \tag{4.11}$$

where $c_{buy}(t)$ and $c_{sell}(t)$ denote the time-varying purchase and selling prices, assumed to be exogenous and known, and τ is the sampling time of the EMS (generally $\tau = 15min$).

The DSO enforces technical limits on the power exchange, i.e.,

$$P_{grid}^{min} \leq P_{grid}(t) \leq P_{grid}^{max}, \tag{4.12}$$

as well as ramping limits on successive transactions,

$$\left| P_{grid}(t) - P_{grid}^{Sl}(t) \right| \leq \delta P_{grid}^{max}, \tag{4.13}$$

where $P_{grid}^{Sl}(t)$ is the scheduled exchange defined by the supervisory EMS. Violations of this constraint incur penalties.

The proposed EMS model serves as a foundational framework for developing advanced energy management methodologies. Formalising the system dynamics, decision variables, cost functions, and operational constraints enables the systematic design and validation of optimisation and control strategies under uncertainty. This model provides a versatile reference structure that supports extensions in future chapters. The model serves as the basis for the microgrid model in chapter 5 and for the microgrids comprising the VPP in chapter 6.

Chapter 4 Bibliography

[ENT23] ENTSO-E. *Frequency Containment, Restoration and Replacement Reserves in the European Power System*. Accessed: October 2025. European Network of Transmission System Operators for Electricity, 2023. URL: https://www.entsoe.eu/network_codes/.

[Gmb24] Amprion GmbH. *Overview of the European Electricity Market Structure*. Accessed: October 22, 2025. Amprion GmbH. 2024. URL: https://www.amprion.net/Bilder/Strommarkt/Marktbericht/Overview/figure-a_720x0.jpg.

[Han+15] Christian A. Hans et al. "Scenario-based model predictive operation control of islanded microgrids". In: *2015 54th IEEE Conference on Decision and Control (CDC)*. 2015, pp. 3272–3277. DOI: 10.1109/CDC.2015.7402711.

[Han+20] Christian A. Hans et al. "Risk-Averse Model Predictive Operation Control of Islanded Microgrids". In: *IEEE Transactions on Control Systems Technology* 28.6 (2020), pp. 2136–2151. DOI: 10.1109/TCST.2019.2929492.

[Kau+19] A. Kaushal et al. "An Overview of Ancillary Services and HVDC Systems in Europe". In: *Energies* 12.18 (2019), p. 3481. DOI: 10.3390/en12183481.

[Klæ+22] Gro Klæboe et al. "Day-ahead market bidding taking the balancing power market into account". In: *Optimization and Engineering* 30 (2022), pp. 683–703. DOI: 10.1007/s11750-022-00645-1.

[Mel+21] Meryem Meliani et al. "Energy management in the smart grid: State-of-the-art and future trends". In: *International Journal of Engineering Business Management* 13 (July 2021), p. 184797902110329. DOI: 10.1177/18479790211032920.

[OCo+25] Ciaran O'Connor et al. "A Review of Electricity Price Forecasting Models in the Day-Ahead, Intra-Day, and Balancing Markets". In: *Energies* 18.12 (2025), p. 3097. DOI: 10.3390/en18123097.

[Pol20] Michael G. Pollitt. "Competition in Markets for Ancillary Services? The Case of Frequency and Voltage Support". In: *Energy Policy* 147 (2020), pp. 111–127. DOI: ?.

[PS17] Iana Perez and Aguinaldo dos Santos. "Distributed economies through open design and digital manufacturing". In: Aug. 2017, pp. 261–267. DOI: 10.5151/sbds-issd-2017-042.

[Rai+18] Stefano Raimondi Cominesi et al. "A Two-Layer Stochastic Model Predictive Control Scheme for Microgrids". In: *IEEE Transactions on Control Systems Technology* 26.1 (2018), pp. 1–13. DOI: 10.1109/TCST.2017.2657606.

[Reb+07] Yann G. Rebours et al. “A Survey of Frequency and Voltage Control Ancillary Services – Part I: Technical Features”. In: *IEEE Transactions on Power Systems* 22.1 (2007), pp. 350–357.

[Sah+23] Diptish Saha et al. “Multiple Microgrids: A Review of Architectures and Operation and Control Strategies”. In: *Energies* 16 (Jan. 2023), p. 600. DOI: 10.3390/en16020600.

[Sch+80] Fred C. Schweppe et al. *Spot Pricing of Electricity*. Kluwer Academic Publishers, 1980.

[Sha+20] D. Shah et al. “A comprehensive review on day-ahead electricity market”. In: *International Transactions on Electrical Energy Systems* 30 (2020). DOI: 10.1002/2050-7038.12360.

[Šro+24] Jakub Šrom et al. “Study of the Imbalance and Trading Settings on the Impact of the Intraday Electricity Trading on the Czech Transmission System Operation”. In: *Preprints* (2024). DOI: ?.

5 | Case Study: Hierarchical Energy Management System

In this chapter, we present the work in [DRS24], which explores the benefits of the Scenario approach in Energy Management System (EMS) with a Hierarchical architecture. The EMS balances and controls a single, grid-connected microgrid with battery storage and uncertain power production and consumption. It is designed as a two-layer control system, with each layer based on the Model Predictive Control (MPC) framework, but it operates at different time scales and pursues distinct goals.

5.1. Contributions

The hierarchical architecture separates the problem into two distinct objectives. The first objective, handled by the supervisory level, focuses on the economic optimisation of operations based on forecasts of net demand and market information. The second objective, managed by the low-level controller, addresses uncertainties within the system. This is done by adjusting participants' power set points and ensuring that operational constraint violations remain below a customizable risk level.

By utilising the two-layer architecture, we can separate the optimisation of economic objectives from the management of uncertainty. This approach allows us to exploit the convexity of the resulting stochastic problem. As a result, we can employ scenario theory (see Chapter 2) to establish a priori probabilistic guarantees against constraint violations arising from unplanned grid interventions.

The second contribution focuses on utilising the support rank of chance constraints, as presented in [SFM13] and chapter 2, to significantly reduce the number of scenarios needed to provide a priori probabilistic guarantees about closed-loop constraint violations. By minimising the number of scenarios, we can compute control actions more quickly and efficiently, enabling a higher sampling frequency while maintaining real-time capabilities.

5.2. Hierarchical Architecture

The EMS has two primary objectives: minimising the costs associated with the micro-grid and ensuring that all subsystems operate within their constraints despite uncertainties. This EMS is structured hierarchically, which helps separate economic goals from the robust operation of the micro-grid. A supervisory controller is responsible for optimising financial operations by utilising nominal forecasts of Photovoltaic (PV) production and load consumption.

At the supervisory level, the controller employs an Economic MPC approach, with a sampling time of $\tau_{Sl} = 15$ minutes and a prediction horizon of $T_{Sl} = 4$ hours. The sampling time aligns with the intraday market bidding time frame, while the prediction horizon balances the need for accuracy with the problem's complexity.

To address high-frequency disturbances within the supervisory level's sampling time, a low-level controller is implemented. This low-level controller operates as a Scenario MPC, using a sampling time of $\tau_{ll} = 1$ minute and a prediction horizon that matches the supervisory level's sampling time, $T_{ll} = \tau_{Sl} = 15$ minutes. This design reflects the controller's hierarchical structure and relies on historical data to generate scenarios for uncertainty.

Given the hierarchical and multi-rate nature of this EMS, it is helpful to define two distinct time indices:

- The index t refers to the discrete-time variable measured in minutes.
- The time index k indicates the k-th 15-minute interval of the day.

When both indices appear in the same equation, the variable with index k is treated as constant during the k-th 15-minute period:

$$P(t) = P(k), \quad \forall t \in [\tau_{Sl} \cdot k, \tau_{Sl} \cdot (k+1) - 1]. \tag{5.1}$$

The overall structure of the hierarchical Energy Management System is presented in Fig. 5.1.

At the upper layer, the Supervisory MPC receives forecasted information, denoted as $P_{\text{net}}^{\text{for}} = \hat{F}(\Phi(k))$, which includes predicted load demand and renewable generation based on previous historical data $\Phi(k)$. Based on these forecasts, the market prices, and technical and operational constraints, the controller computes optimal reference values for the nominal generation P_{gen} and the grid intervention P_{grid}. These setpoints define the desired operating conditions of the microgrid over a prediction horizon, ensuring economic

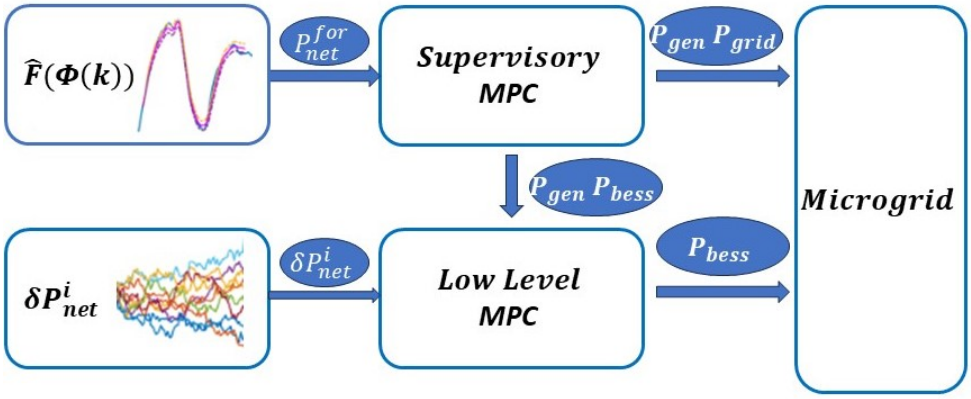


Figure 5.1: Hierarchical EMS scheme

efficiency while complying with operational constraints.

The Low-Level MPC operates on a faster time scale and refines the decisions made by the supervisory layer. It accounts for short-term fluctuations in net demand, represented by δP^i_{net}, arising from uncertainties or forecast errors and leading to rapid variations in net power. The low-level controller adjusts the Battery Energy Storage System (BESS) power P_{bess} to maintain system balance and stability in real time, while minimising the unplanned grid intervention δP_{grid}.

The introduction of a two-layer (hierarchical) architecture introduces some degree of sub-optimality compared to a centralised formulation. This trade-off is intentional and necessary: the hierarchical structure improves scalability, modularity, and, often, the computational tractability of the problem itself, with reduced performance loss compared to the global optimal solution.

5.2.1. Supervisory Model Predictive Control

The supervisory level is an Economic MPC with a prediction horizon of $n_{Sl} = 16$, corresponding to 4 hours, and a sampling time $\tau_{Sl} = 15$ min. This introduces feedback into the system by measuring the BESS's State Of Charge (SOC) and updating the forecast with more recent power demand information.

Nominal forecast

The net demand is non-controllable and uncertain due to its correlation with various environmental and temporal factors, such as solar radiation, temperature, time of day, and season; its low-frequency trend can be accurately forecast. Predicting a The EMS uses a Neural Autoregressive with eXogenous inputs (NARX) Model $\hat{\mathbf{F}}$ to predict the

future net demand based on available past data with a sampling time of 15 minutes,

$$P_{net}^{for}(j) = \hat{\mathbf{F}}(\Phi(k)), \forall j \in [k : k + T_{for}], \tag{5.2}$$

where T_{for} is the forecasting horizon and the regressor $\Phi(k)$ is formed as

$$\Phi(k) = [P_{net}(k-1) \dots P_{net}(k - D_\Phi)]^\top,$$

where D_Φ is the number of input delay of the model.
The model $\hat{\mathbf{F}}(\Phi(k))$ is trained on historical data to minimise prediction error and is used by the supervisory level to make informed decisions about future power. The forecasting model (5.2) is a NARX model with 24 feedback delays and a hidden layer containing 10 neurons. This model is trained on 1,000 days of data and validated on a separate dataset comprising 400 days, achieving a Mean Square Error (MSE) of 12.79 kW. The data is taken from a real-world microgrid [VAK19]. An example of net-demand prediction, with error levels and true realisation, is depicted in Fig. 5.2

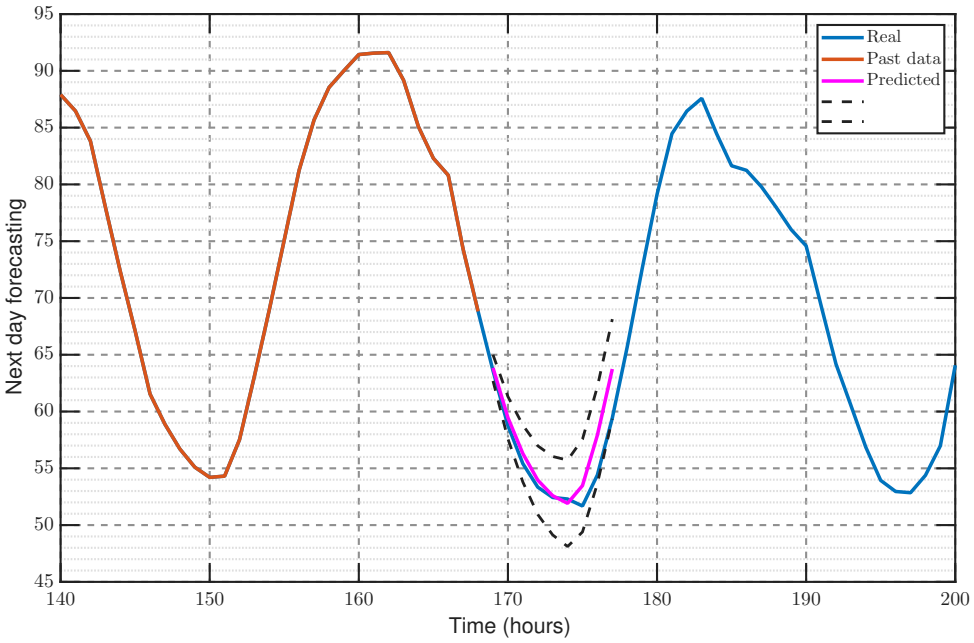


Figure 5.2: Past net-demand data (orange), NARX net-demand prediction (pink-line) with 3σ interval (black-dotted), true net-demand (blue-line)

Supervisory level optimization

The optimisation aims to minimise the operational costs of the micro-grid during the time horizon by acting on the schedule of the nominal power set-points u_{Sl}

$$\begin{aligned} &u_{Sl} = [P^{Sl}_{gen}(k), P^{Sl}_{bess}(k), P^{Sl}_{grid}(k)], \\ &\forall k \in [1, n_{Sl}]. \end{aligned} \tag{5.3}$$

Cost function: the cost function is a multi-objective economic function, i.e.,

$$f(u_{Sl}) = C_{gen} + C_{MG} + \mathcal{F}_{gen} + \mathcal{F}_{bess} + \mathcal{F}_{grid} \tag{5.4}$$

The cost function comprises the local generation costs C_{gen}, the costs for energy trading C_{MG}, as defined in (4.4) and (4.11); additional filtering terms penalises excessive variation in power commands:

$$\mathcal{F}_h = q_h \cdot \sum_{k=2}^{n_{Sl}} (P_h(k) - P_h(k-1))^2, \tag{5.5}$$

with $h \in \{gen,\ bess,\ grid\}$.

MIQP programming: The resulting problem is a Mixed Integer Quadratic Program (MIQP) due to the presence of binary variables for the BESS, network and generator logic and the

cost function being quadratic in u_{Sl}.

$$\min_{u_{Sl}} \quad \sum_{t=1}^{T} \tau_{Sl} \cdot [\ c_{gen}(P_{gen}(t), t) + c_{buy}(t)\ \max(0, P_{grid}(t)) + c_{sell}(t)\ \min(0, P_{grid}(t))] \tag{5.6a}$$

$$+ \mathcal{F}_{gen} + \mathcal{F}_{bess} + \mathcal{F}_{grid}, \tag{5.6b}$$

s.t. **Power balance**

$$P_{\text{gen}}(t) + P_{\text{bess}}(t) + P_{\text{grid}}(t) = P_{\text{load}}(t) - P_{\text{pv}}(t), \tag{5.6c}$$

Generator constraints

$$\begin{cases} 0 < P_{\text{gen}}(t) \leq P_{\text{gen}}^{\max}, & \text{if isON}(t) = 1, \\ P_{\text{gen}}(t) = 0, & \text{if isON}(t) = 0, \end{cases} \tag{5.6d}$$

$$|P_{\text{gen}}(t) - P_{\text{gen}}(t-1)| \leq P_{\text{ramp}}^{\max}, \tag{5.6e}$$

BESS BESS dynamics and limits

$$\begin{cases} s(t+1) = s(t) - k_{\text{DCH}}\, u_{\text{DCH}}(t) + k_{\text{CH}}\, u_{\text{CH}}(t), \\ s(0) = s_0, \end{cases} \tag{5.6f}$$

$$s_{\min} \leq s(t) \leq s_{\max}, \tag{5.6g}$$

BESS power limits

$$\begin{cases} 0 \leq u_{\text{DCH}}(t) \leq u_{\text{DCH}}^{\max}\, b(t), \\ 0 \leq u_{\text{CH}}(t) \leq u_{\text{CH}}^{\max}\, (1 - b(t)), \end{cases} \tag{5.6h}$$

$$P_{\text{bess}}(t) = u_{\text{DCH}}(t) - u_{\text{CH}}(t), \tag{5.6i}$$

$$-u_{\text{CH}}^{\max} \leq P_{\text{bess}}(t) \leq u_{\text{DCH}}^{\max}, \tag{5.6j}$$

Net-Demand forecast

$$P_{\text{net}}^{\text{for}}(j) = \hat{\mathbf{F}}(\Phi(k)), \quad \forall j \in [k : k + T_{\text{for}}], \tag{5.6k}$$

Grid transmission power constraints

$$P_{\text{grid}}^{\min} \leq P_{\text{grid}}(t) \leq P_{\text{grid}}^{\max}, \tag{5.6l}$$

$$\left|P_{\text{grid}}(t) - P_{\text{grid}}^{Sl}(t)\right| \leq \delta P_{\text{grid}}^{\max}. \tag{5.6m}$$

Binary and integer variables are fundamental for capturing the behaviour of complex logic, switching dynamics (such as battery charge/discharge cycles), and non-convex functions approximated by Piece-Wise Affine functions. Moreover, commercial solvers can efficiently solve MILP-based EMSs.

Once the MIQP is solved, the first control action $u_{Sl}^{*}(k) = [P_{gen}^{Sl*}(k), P_{bess}^{Sl*}(k), P_{Grid}^{Sl*}(k)]^{\top}$ is applied to the system in a receding horizon way, and the MPC time window is moved forward.

5.2.2. Low-level controller

The low-level controller is responsible for regulating and maintaining the power balance within the microgrid by addressing uncertainties in power production and consumption. This low-level controller operates using a scenario-based MPC approach, with a sampling time of $\tau_{ll} = 1$ minute and a prediction horizon of $n_{ll} = 15$, corresponding to $T_{ll} = \tau_{Sl}$. It can only adjust the amount of power exchanged by the BESS but does not control its operational state—whether it is charging, discharging, or idle. The supervisory level determines this operational state.

Short-term uncertainties

While the long-term trend of net demand can be predicted effectively using $\hat{\mathbf{F}}$, the actual net demand over a short timescale of $\tau_{ll} = 1$ minute is influenced by various factors such as noise and high-frequency components, which the model cannot accurately forecast. Rapid changes in net demand can lead to a significant increase in unplanned interventions in the distribution system, highlighting the importance of effective EMS uncertainty management for proper system operation. We can decompose $P_{PV}(t)$ and $P_{load}(t)$ into a predictable component plus an unknown short-term uncertainty.

$$P_{PV}(t) = P_{PV}^{For}(k) + \delta P_{PV}^{ll}(t), \tag{5.7a}$$
$$P_{load}(t) = P_{load}^{For}(k) + \delta P_{load}^{ll}(t). \tag{5.7b}$$

The unknown realisation of uncertainty $\delta = [\delta P_{pv}(t), \delta P_{load}(t)]$ is drawn from the uncertainty set $\mathbf{\Delta}$ defined over an unknown probability measure $\mathbf{P}$. The EMS must operate the system robustly to compensate for possible realisations of uncertainty $\delta \in \Delta$.

Stochastic Finite Horizon Optimal Control Problem

After setting optimal power set points every 15 minutes at the supervisory level and assuming the generator set point is fixed, the EMS must rely on either the BESS or grid intervention to balance the remaining stochastic components of the net demand. By defining the variations

$$\delta P_{bess}^{ll}(t) = P_{bess}(t) - P_{bess}^{Sl*}(k), \tag{5.8a}$$
$$\delta P_{grid}^{ll}(t) = P_{grid}(t) - P_{grid}^{Sl*}(k), \tag{5.8b}$$

$\forall t \in [15 \cdot (k-1) + 1, 15 \cdot k]$, the power balance (4.1) can be rewrite as:

$$P_{gen}^{Sl*}(k) + P_{bess}^{Sl*}(k) + P_{grid}^{Sl*}(k) + \delta P_{bess}^{ll}(t) + \delta P_{grid}^{ll}(t) = P_{net}^{for}(k) + \delta P_{net}(t). \tag{5.9}$$

Recalling that the supervisory level updates the power set-points at a 15-minute rate, the power

balance for the low level simplifies as

$$\delta P^{ll}_{bess}(t) + \delta P^{ll}_{grid}(t) = \delta P_{net}(t) \tag{5.10}$$

Using a stochastic formulation, the EMS can permit temporary violations of the maximum power-variation constraint on grid inflows by imposing a risk constraint with a risk parameter ϵ.

$$\mathrm{P}[\delta \in \Delta : |\delta P^{ll}_{grid}(t)| \geq \delta P^{max}_{grid}] \leq \epsilon \tag{5.11}$$

Additionally, the EMS aims to minimise the worst-case variance of the grid intervention to minimise the unplanned grid interventions, i.e.,

$$\begin{gathered} \min_{\delta P^{ll}_{bess}(q)} \; max_{\delta}[\textstyle\sum_{q=t}^{t+n_{ll}} ||\delta P^{ll}_{grid}(q)||_2^2] \\ q \in [t \ldots, t + n_{ll} - 1] \end{gathered} \tag{5.12}$$

The latter problem can be stated as the following Stochastic Finite Horizon Optimal Control Problem (FHOCP):

$$\delta P^{ll*}_{\text{bess}}(q) = \arg \min_{\delta P^{ll}_{\text{bess}}(q)} \max_{\delta} \left[\sum_{q=t}^{t+n_{ll}} \|\delta P^{ll}_{\text{grid}}(q)\|_2^2 \right], \tag{5.13a}$$

s.t. **Unplanned intervention chance constraint**

$$\mathbb{P}\big[\delta \in \Delta : |\delta P^{ll}_{\text{grid}}(q)| \geq \delta P^{\max}_{\text{grid}}\big] \leq \epsilon, \tag{5.13b}$$

BESS SoC dynamics and limits

$$\begin{cases} s(q+1) = s(q) - k_{\text{DCH}}\, u_{\text{DCH}}(q) + k_{\text{CH}}\, u_{\text{CH}}(q), \\ s(0) = s_0, \end{cases} \tag{5.13c}$$

$$s_{\min} \leq s(q) \leq s_{\max}, \tag{5.13d}$$

BESS power limits

$$\begin{cases} 0 \leq u_{\text{DCH}}(q) \leq u^{\max}_{\text{DCH}}\, b(q), \\ 0 \leq u_{\text{CH}}(q) \leq u^{\max}_{\text{CH}}\, (1 - b(q)), \end{cases} \tag{5.13e}$$

$$-u^{\max}_{\text{CH}} \leq P_{\text{bess}}(q) \leq u^{\max}_{\text{DCH}}, \tag{5.13f}$$

Uncertainty model

$$P_{\text{PV}}(q) = P^{\text{for}}_{\text{PV}}(q) + \delta P^{ll}_{\text{PV}}(q), \qquad P_{\text{load}}(q) = P^{\text{for}}_{\text{load}}(q) + \delta P^{ll}_{\text{load}}(q), \tag{5.13g}$$

$$\delta P^{ll}_{\text{bess}}(q) + \delta P^{ll}_{\text{grid}}(q) = \delta P_{\text{net}}(q), \qquad q \in [t,\, t + n_{ll} - 1]. \tag{5.13h}$$

The hierarchical structure eliminates the need for binary decision variables, which were previ-

ously introduced to manage the BESS logic in the low-level controller. The supervisory level defines the BESS logic and other binary decision variables through nominal MIQP optimisation. This includes determining whether the BESS is in charging or discharging mode. Once established, these decisions are communicated to the low-level controller. Consequently, the low-level program becomes convex, enabling the use of generalisation properties of the scenario approach based on the probability of constraint violation for unseen scenarios.

Convex stochastic programming instances can be addressed using the scenario approach described in chapter 2, where the uncertain set, denoted by Δ, is approximated by sampling N scenarios, δ_i, for $i = 1 \dots N$.

After sampling the scenarios, we solve an approximate program where the optimisation variables must lie within the intersection of the constraints derived from each scenario, represented as $X = \cap_{i=1}^{N} X_i$. This approximation is deterministic and can be solved with ease. In the context of scenario theory, to ensure the satisfaction of a chance constraint with a risk level of ϵ and a confidence level of $1 - \beta$,

$$\mathcal{P}[\delta \in \Delta : x \notin X = \cap_{i=1}^{N} X_i] < \epsilon. \tag{5.14}$$

Given the power set points generated by the supervisory controller at the k^{th} 15-minute time interval, the low-level controller addresses an approximated version of the stochastic programming problem (5.13) with a fixed configuration of the BESS, which can either be in a charging (CH) or discharging (DCH) state. Let the superscript i denote the i^{th} scenario, and let the operational state of the BESS be represented by $j \in \{CH, DCH\}$. At each time instant t, the low-level controller solves the following deterministic optimisation

$$\delta P^{*}_{bess}(q) = \arg \min_{\delta P_{bess}(q)} \left\{ \max_{i} \left[\sum_{q=t}^{t+n_{ll}} \|\delta P^{i}_{grid}(q)\|_2^2 + r \cdot \|\gamma\|^2 \right] \right\}, \tag{5.15a}$$

s.t. **Unplanned intervention chance constraint**

$$\delta P^{i}_{grid}(q) = \delta P^{i}_{net}(q) - \delta P_{bess}(q), \tag{5.15b}$$

Grid power limits

$$|\delta P^{i}_{grid}(q)| \leq \delta P^{\max}_{grid} + \gamma^2, \tag{5.15c}$$

BESS SoC dynamics

$$\begin{cases} s^i(q+1) = s^i(q) - K^i_j \cdot \left(P^{Sl}_{bess} + \delta P_{bess}(q)\right), \\ s^i(0) = s^i_0, \end{cases} \tag{5.15d}$$

BESS power limits

$$\delta P^{\min,j}_{bess} \leq \delta P_{bess}(q) \leq \delta P^{\max,j}_{bess}, \tag{5.15e}$$

BESS SoC constraints

$$s^i(q) \in S, \qquad i \in \{1, \dots, N\}, \quad q = t, \dots, t + n_{ll} - 1. \tag{5.15f}$$

In the scenario program eq. (5.15) the original cost function in eq. (5.13) is relaxed by minimising the maximum cost among all the scenarios, in other words, minimising the worst-case cost. Eq. (5.15) can be written as an equivalent epigraph form, preserving convexity; an additional variable λ is added to soften the constraint on δP_{grid}^{max}, and it is weighted in the cost function by a positive coefficient r.

Once the MPC solves the scenario program eq. (5.15) at each time step, the first control action $\delta P_{bess}(t) = \delta P_{bess}^{ll*}(t)$ is used to update the power command of the BESS in a receding horizon way. The scenario MPC time window is then moved forward.

Statistical guarantees and choice of N

By exploiting the closed-loop architecture, one can derive tighter probabilistic guarantees on the expected time-average probability of constraint violation, as proposed in [SFM14]. By denoting M_t as the violation condition of the chance-constraint for an unseen scenario, only at time t, that is, at the first step of the prediction horizon of the MPC:

$$M_t = 1 \; if \; |\delta P_{grid}(t)| > \delta P_{grid}^{max}, \tag{5.16}$$

the expected time-average probability of violation occurrence remains below an arbitrary risk level ϵ is

$$\mathrm{E}_\delta \; [\frac{1}{T} \sum_{t=1}^{T} M_t] \leq \epsilon, \tag{5.17}$$

provided that the number of scenarios used to construct the MPC program, N_{mpc}, is chosen to be

$$N_{mpc} = \frac{\rho}{\epsilon} - 1, \tag{5.18}$$

where ρ is the support rank of the chance-constraint, [SFM14]. The support rank for the chance-constraint is upper-bounded by the dimension of the input channel of the system, that is, $\rho = 1$. In other words, by properly choosing N_{mpc}, the percentage of violations over the operating time is bounded. For example, if the system accepts a percentage of violations $\epsilon = 5\%$, the number of scenarios to be picked is $N_{mpc} = 19$. This number is way less than the $N = 1020$ scenarios needed to ensure the bound on the violation over the entire horizon.

5.3. Simulation case study

This section presents a case study on the implementation of the Hierarchical EMS using historical data from an actual microgrid research building located in Tsukuba, Japan, as described in [VAK19]. The microgrid is simulated for 20 hours (denoted as $T_{sim} = 20$ hours) to demonstrate the effectiveness of the proposed method.

5.3.1. Microgrid Parameters and Simulation

The microgrid consists of several components:

- A Battery Energy Storage System (BESS) with a capacity of 326 kWh and a maximum charging and discharging power of $u_{ch}^{max} = u_{dch}^{max} = 90\,\text{kW}$, starting with an initial SOC of $s(0) = 0.5$.
- A thermal generator with a nominal power of $P_{gen}^{nom} = 80\,\text{kW}$, a maximum power of $P_{gen}^{max} = 120\,\text{kW}$, and a maximum ramp rate of $P_{ramp}^{max} = 5\,\text{kW}$.
- Solar panels with a combined peak power of $P_{pv}^{max} = 90\,\text{kW}$.
- A building that can consume up to $P_{load}^{max} = 140\,\text{kW}$.
- A connection to the primary grid, which has a maximum exchange limit of $P_{grid}^{max} = 100\,\text{kW}$ and allows for a maximum unplanned intervention of $\delta P_{grid}^{max} = 5\,\text{kW}$.

The power profiles and cost information are sourced from the database referenced in [VAK19]. The generator cost $c_{gen}(P_{gen})$ and the buying prices are depicted in Fig. 5.3. The function representing the generator cost is non-linear and non-convex, reflecting the economics of real power plants; the MILP optimisation is solved by approximating the non-convex function through a piecewise affine one. During periods of high demand, prices increase in a three-tiered range.

Examples of load consumption and PV production are depicted in Fig. 5.4. By computing the error between forecasted and actual net demand, we find that the variance increases during periods of solar power production, indicating that the significant uncertainty stems from solar availability. This is shown in Fig. 5.5, where we can see that the $3 \cdot \sigma$ interval in our power prediction increases between 6 a.m. and 6 p.m..

5.3.2. Economic optimisation

The MIQP for the supervisory level, as represented in eq. (5.6), along with the scenario MPC stated in eq. (5.15), is modelled using the Yalmip language [Löf04] and solved with the commercial multipurpose solver, Gurobi [Gur22].

The scenarios are generated by applying the predictive model to each day in the dataset and calculating the residuals, which serve as the realisation of uncertainty, as described in equation (5.7).

Figure 5.6 illustrates the trajectory of the power schedule over an entire day. It is evident that the EMS effectively utilises market information by discharging the BESS during periods of high prices. Furthermore, the EMS ensures the system remains within operational constraints through MPC feedback.

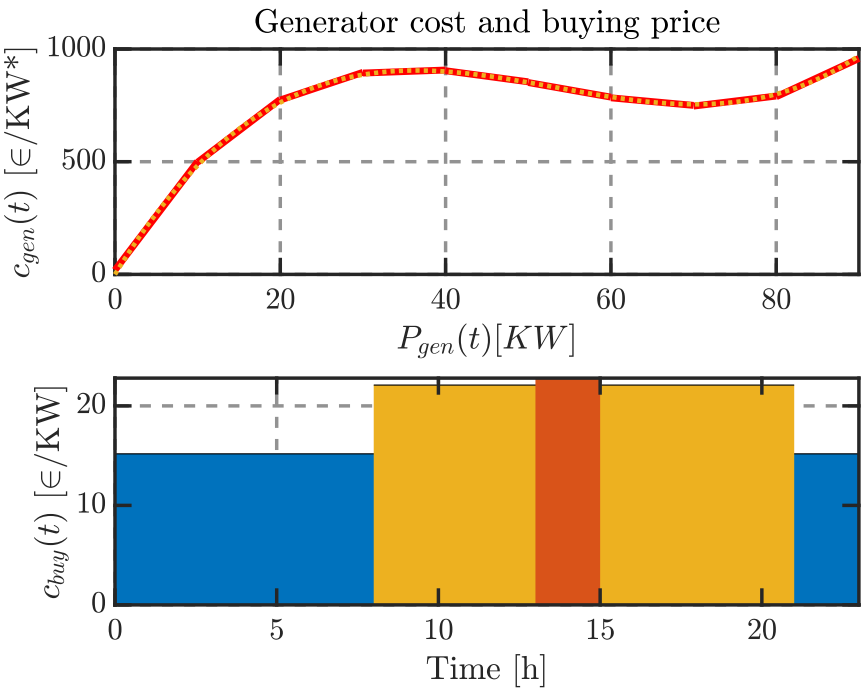


Figure 5.3: Upper plot, cost of generation c_{gen} as a function of generated power P_{gen}. Lower plot, Pricing Ranges during the day: low cost (blue), medium cost (orange), high cost (red).

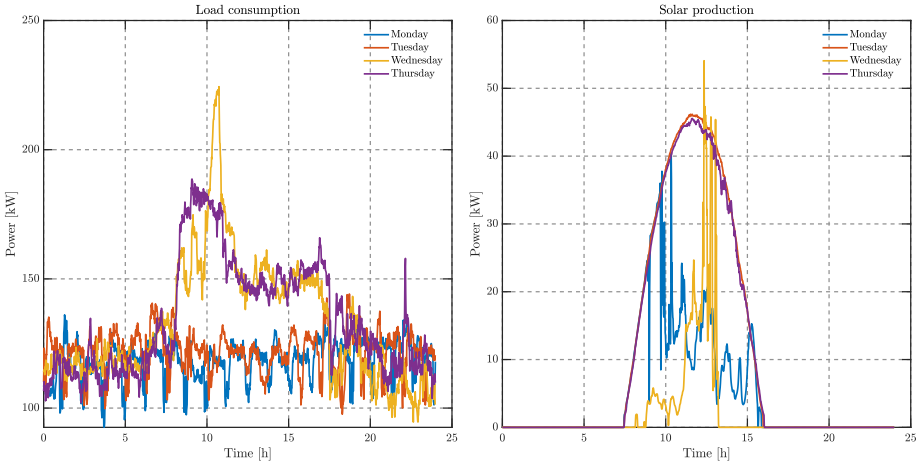


Figure 5.4: Left plot: load -, right plot PV production

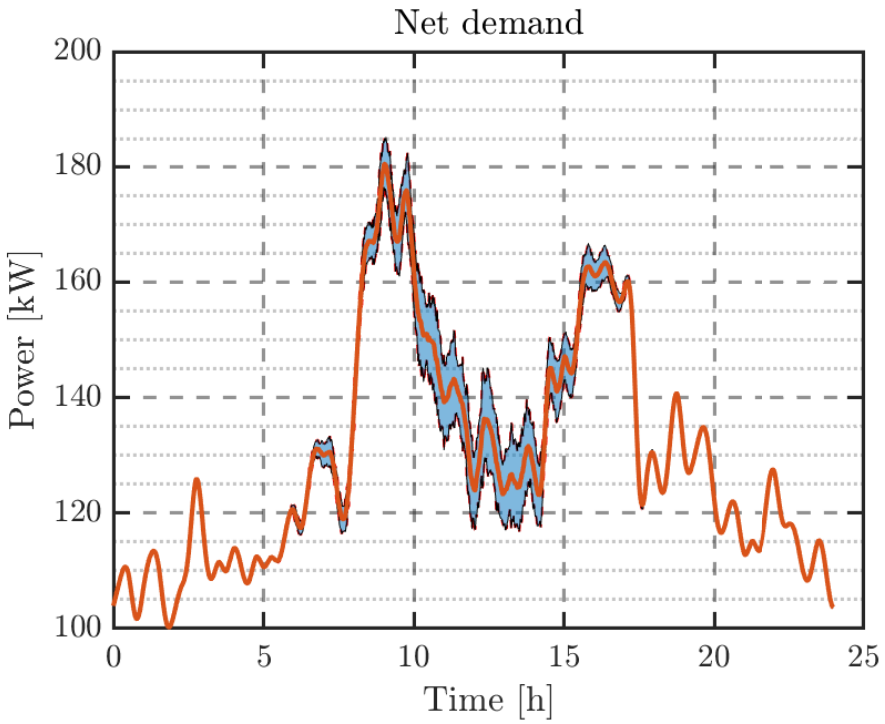


Figure 5.5: Net-demand with $3 \cdot \sigma$ intervals.

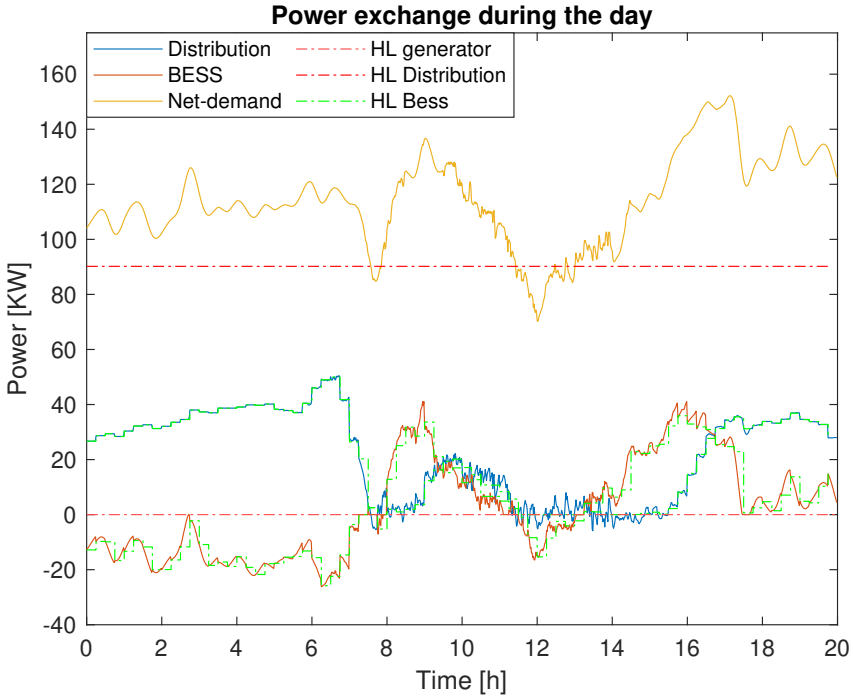


Figure 5.6: Daily Power profiles for the micro-grid.

5.3.3. Low-level balancing

The low-level controller is responsible for managing micro-grid operations, keeping them close to the nominal values set by the supervisory level. This helps reduce unexpected energy exchanges with the primary grid. When low-level control is absent, all uncertainties in net demand directly lead to increased unplanned interventions in the grid's power flow. However, the hierarchical EMS effectively minimises grid interventions, ensuring a lower level of variability in P_{grid}.

Figure 5.7 displays a histogram showing the distribution of $||\delta P_{grid}(t)||_2^2$ for scenarios involving only the supervisory level EMS (in blue) and the full hierarchical EMS (in orange). The results indicate that both the mean and variance of $||\delta P_{grid}||_2^2$ decrease when the low-level controller is utilised, with the majority of the distribution density falling below 1 kW. This reduction results in a significant decrease in unexpected energy exchanges with the primary grid, from 70.83 kWh to 20.89 kWh, a 71.5% reduction.

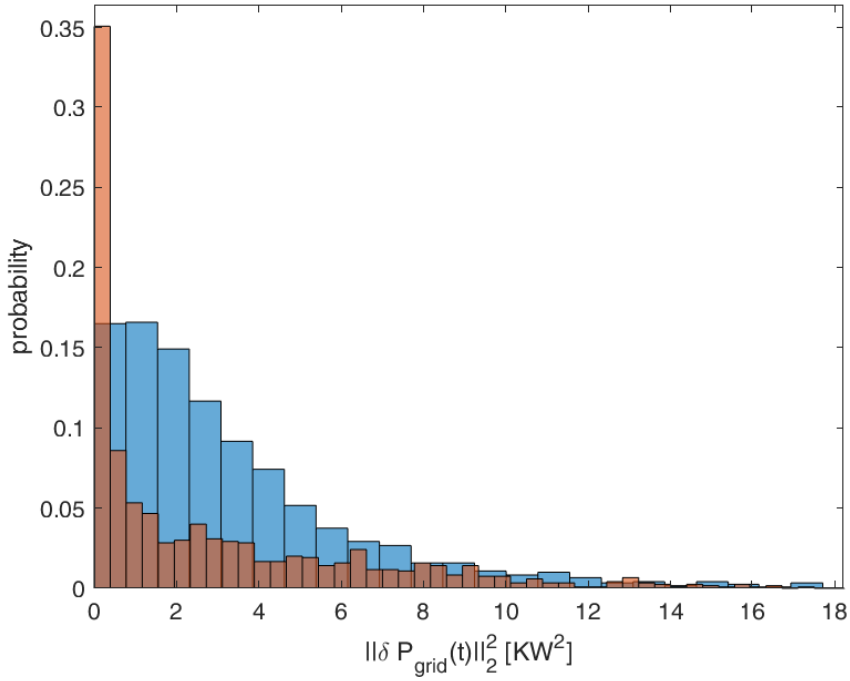


Figure 5.7: Histograms of power flows exchanged with the grid, $||\delta P_{grid}(t)||_2^2$, for EMSs with (orange) and without (blue) the low-level risk-constrained MPC.

The sampled violation occurrence of the chance-constraint along the simulation is computed as

$$M = \frac{1}{N_{Sim}} \sum_{t=1}^{N_{sim}} M_t, \tag{5.19}$$

where N_{Sim} is the number of simulation steps. The EMS simulation has been performed for

different numbers of scenarios, and the resulting violation occurrences are summarised in Table 5.1, together with the upper bound provided by eq. (5.18).

Table 5.1: Violation occurrence

N_{mpc}	40	60	100
M	2.33%	2.12%	1.36%
ϵ	2.50%	1.69%	1.00%

As the theory predicts, increasing the number of scenarios reduces the average in-time occurrence of violations M_t. Since this is an asymptotic result, discrepancies between ϵ and M_t are observed. These discrepancies would disappear with longer simulation intervals.

Chapter 5 Bibliography

[DRS24] A. Del Duca, F. Ruiz, and R. Scattolini. "Robust Micro-Grid Energy Management System Through a Scenario Approach". In: *2024 American Control Conference (ACC)*. Toronto, ON, Canada, 2024, pp. 1801–1806. DOI: 10.23919/ACC60939.2024.10644278. URL: https://hdl.handle.net/11311/1273510.

[Gur22] Gurobi Optimization, LLC. *Gurobi Optimizer Reference Manual*. 2022. URL: https://www.gurobi.com.

[Löf04] J. Löfberg. "YALMIP : A Toolbox for Modeling and Optimization in MATLAB". In: *In Proceedings of the CACSD Conference*. Taipei, Taiwan, 2004.

[SFM13] Georg Schildbach, Lorenzo Fagiano, and Manfred Morari. "Randomized Solutions to Convex Programs with Multiple Chance Constraints". In: *SIAM Journal on Optimization* 23.4 (2013), pp. 2479–2501. DOI: 10.1137/120878719.

[SFM14] Georg Schildbach, Lorenzo Fagiano, and Manfred Morari. "Energy Management for Building Climate Comfort in Uncertain Smart Thermal Grids with Aquifer Thermal Energy Storage". In: *Energy* 68 (2014), pp. 1–11. DOI: 10.1016/j.energy.2014.02.069.

[VAK19] Karina Vink, Eriko Ankyu, and Michihisa Koyama. "Multiyear microgrid data from a research building in Tsukuba, Japan". In: *Scientific Data* 6.1 (Feb. 2019), p. 190020. ISSN: 2052-4463. DOI: 10.1038/sdata.2019.20. URL: https://doi.org/10.1038/sdata.2019.20.

6 | Case-study: Distributed VPP architecture through scenario approach

In this chapter, we present a Virtual Power Plant (VPP) architecture based on the distributed coordination of multiple microgrids. The VPP uses results proposed in [DRS24][DR25] and presented in chapter 2 to handle uncertainty in the global uncertain constraint, which represents the minimum amount of aggregated power reserves of the VPP. The VPP enables microgrids to operate as a single, cohesive generation unit, allowing them to participate in the primary market; otherwise, they may struggle to do so individually.

6.1. Contributions

The presented VPP architecture is evaluated through an extensive simulation campaign 6.4, verifying that the framework in 2 works for networks of increasing size, confirming the possibility of decouple uncertainty handling from network complexity, supporting the claims made in [DR25].

The VPP operations are optimised by means of distributed Augmented Lagrangian algorithms, as presented in [FP23]. The algorithm ensures local balance across the networks, meets the minimum aggregated capacity required by the Transmission System Operator (TSO), and accounts for the operational constraints and economics of the single microgrid.

The proposed architecture provides several advantages: microgrids retain their independence as decision-making entities and do not need to disclose any sensitive information about their assets, including operational constraints and costs. As explained in 6.3.2, they only need to communicate dual variables, which are linked with the global coupling constraints, and therefore, it is not private information.

The absence of a central unit reduces communication and burden by reducing the need for communication between microgrids and the VPP. In this sense, the proposed VPP is entirely virtual, as it emerges solely from the aggregation of microgrids. The resulting VPP is inherently robust against microgrid disconnections, communication delays, and topology reconfigurations,

making it suitable for aggregating microgrids, even during islanded operations.

6.2. Virtual Power Plant modelling

We consider a VPP comprising of m microgrids (agents), interconnected through a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$. Here, $\mathcal{V} = \{1, 2, \ldots, m\}$ represents the set of vertices (microgrids), while $\mathcal{E} = \{(i, j) : i \in \mathcal{N}_j \text{ and } j \in \mathcal{N}_i\}$ denotes the set of edges.

The network's edges $(i, j) \in \mathcal{E}$ serve two distinct roles: A communication service between microgrids i, j, and the possibility of exchanging power $P_{(i,j)}$ through a physical bus. Microgrids interconnected in such a way, meaning that for $(i, j) : (i, j) \in \mathcal{E}$, are neighbors; we can define the set of neighbors of microgrid i as $\mathcal{N}_i = \{j \in V : (i, j) \in \mathcal{E}\}$.

Microgrid modeling

Each microgrid is a decision-making entity that operates based on local and private information, such as local decision variables $x_i(t)$, representing the power dispatch of the microgrid over the optimisation horizon T $t \in [0, T-1]$, costs $f_i(t)$, and operational constraints X_i.

Each microgrid is a discrete-time linear time-invariant (LTI) system with sampling time $\Delta t = 1h$. As presented in chapter 4, microgrids generally consist of a dispatchable generator, some Photovoltaic generators, and some loads that consume power. Most microgrids also include a Battery Energy Storage System (BESS) to enhance flexibility and provide energy storage.

In this section, we present a microgrid model for the network's agents, based on the model developed in 4.

Dispatchable generation units: A dispatchable generator, such as diesel generators and backup generators, produces a controllable amount of power $P^i_{gen} \in [0, \overline{P}^i_{gen}]$ by burning fuel, where $\overline{P}^i_{gen}$ denotes the maximum amount of power that can be produced. The cost of production is a quadratic function of the amount of power generated this way: $c^i_{gen}(t) = \alpha S P^i_{gen}(t)^2 + \beta P^i_{gen}(t) + \gamma$. The total cost over the optimisation horizon T is

$$c^i_{gen} = \sum_{t=1}^{\top} c^i_{gen}(t) = \sum_{t=1}^{\top} (\alpha_i \cdot (P^i_{gen}(t))^2 + \beta_i \cdot P^i_{gen}(t) + \gamma_i). \tag{6.1}$$

Dispatchable generations consist of rotary machines and are therefore subject to maximum ramp-rate constraints to avoid excessive mechanical stresses. Such constraints can be modelled as

$$|P^i_{gen}(t) - P^i_{gen}(t-1)| \leq \overline{P}^i_{ramp}. \tag{6.2}$$

Battery Energy Storage Systems: Battery Energy Storage Systems (BESSs) store electrical energy in the form of chemical energy. The State Of Charge (SOC) of a BESS, $s^i(t)$, indicates the amount of energy stored. The dynamics of the BESS are modelled as a Linear Time-Invariant (LTI) system, with the power exchange with the microgrid denoted as $P^i_{bess}(t)$, i.e.,

$$\begin{cases} s^i(t+1) = s^i(t) - \frac{\Delta t}{Q_i} \cdot P^i_{bess}(t), \\ s^i(0) = s^i_0, \end{cases} \tag{6.3}$$

where s^i_0 is the initial energy level of the battery, and the term $P^i_{bess}(t)$ is considered positive when the battery is discharging and negative when charging. This equation differs from the one presented in (4.5) due to the absence of different efficiencies. This simplification is needed to avoid binary variables, which turn the optimisation problem into a Mixed Integer Problem and make it difficult for distributed schemes to converge to the global solution. However, this approximation is not crucial, since the VPP's scope is to provide reference values for microgrids with economic rather than operational insight.

Furthermore, $P^i_{bess}(t)$ must remain within safe operational constraints: $P^i_{bess}(t) \in [\underline{P}^i_{bess}, \overline{P}^i_{bess}]$. The SOC cannot exceed the operational limits: $s^i(t) \in [\underline{s}, \bar{s}]$. To ensure enough residual capacity at the end of the day, we impose an additional constraint on the final SOC. The final SOC must be equal to the initial SOC: $s^i(0) = s^i(T-1)$. Finally, each BESS is associated with a cost to prevent excessive degradation resulting from rapid charge and discharge cycles,

$$c^i_{bess} = \sum_{t=0}^{T-1} (\delta_i \cdot (P^i_{bess}(t) - P^i_{bess}(t-1))^2). \tag{6.4}$$

Power exchange with neighbours microgrids: The i-th microgrid can exchange power with all its neigbours $j \in \mathcal{N}_i$ through the connection $(i,j) \in \mathcal{E}$. This power exchange must adhere to operational constraints, remaining within a range $P_{(i,j)}(t) \in [\underline{P}_{(i,j)}, \overline{P}_{(i,j)}]$, and the cost is associated to the line power losses, consequently scaling quadratically with the exchanged power, i.e.,

$$c_{(i,j)} = \sum_{t=0}^{T} (R_{(i,j)} \cdot P^2_{(i,j)}(t)). \tag{6.5}$$

By incorporating this penalisation term in the VPP cost function, we can encode the linearised power flow approximation into the optimisation, as outlined in [Nau+21].

Non-dispatchable production and consumption: The main source of uncertainty in each microgrid is the presence of Photovoltaic (PV) power production $\hat{P}^i_{pv}(t)$ and local consumption $\hat{P}^i_l(t)$, for which forecasts are available.

Microgrid internal balance: To be balanced, each microgrid has to ensure that, at each time instant, generated, consumed, and received/transmitted power are balanced. In other words, each microgrid has to satisfy this local power balance

$$P^i_{gen}(t) + P^i_{bess}(t) + \hat{P}^i_{pv}(t) - \hat{P}^i_l(t) + \sum_{j \in \mathcal{N}_i} P^i_j(t) \geq 0. \tag{6.6}$$

In Eq. (6.6), a positive sign indicates either that the power is generated (from dispatchable and non-dispatchable sources), discharged from the BESS, or imported from neighbouring microgrids. Conversely, a negative sign indicates that power is consumed by the load, used to charge the BESS, or exported to other neighbours.

This constraint must always be non-negative and is zero if the VPP does not furnish virtual power to the TSO. When, instead, the TSO makes an aggregated Virtual ramp-up request to the VPP, some microgrids have a (strictly) positive contribution. In this way, the VPP can export power to the main grid, fulfilling the obligations.

Based on the model previously depicted, each microgrid has the following decision variables:

$$x_i(t) = [P^i_{gen}(t), P^i_{bess}(t), s^i(t), \{P^i_j(t)\}_{j \in \mathcal{N}_i}] \ \forall t, \tag{6.7}$$

that must lie within the microgrid's operational constraints can be compactly rewritten as a convex set: $x_i \in X_i$. The total operational cost each microgrid must minimise is

$$f_i(x_i) = c^i_{gen} + c^i_{bess} + 0.5 \cdot \sum_{j \in \mathcal{N}_i} c_{(i,j)}. \tag{6.8}$$

If the microgrids were not interconnected, each agent would have determined its optimal power schedule x^*_i as follows:

$$\mathcal{P}_i : min_{x_i} f_i(x_i), \tag{6.9a}$$

$$s.t.$$

$$x_i \in X_i. \tag{6.9b}$$

Coupling constraints

To encode the networked nature of the system, enabling coordination of multiple microgrids as a single VPP entity, it is necessary to explicitly state the constraints that couple each microgrid with its neighbours and with the TSO. In the presented VPP architecture, we identified two different coupling constraints. The first one encodes the power flow within the network, and the second one encodes the aggregated service that the VPP must provide to the TSO, which is uncertain due to the presence of uncertain production and consumption.

Power-flow coupling constraints: Each microgrid has a set of decision variables which dictate the power exchange with other neighbouring microgrids: $P_j^i(t)\ \forall j \in \mathcal{N}_i$.To ensure that the power balance across the network is satisfied, for each edge $(i,j) \in \mathcal{E}$, the following constraint must hold:

$$P_{(i,j)} = P_j^i(t) = -P_i^j(t)\ \forall (i,j) \in \mathcal{E}\ \forall t \tag{6.10}$$

This constraint enforces the correct power exchange between i and j; in fact, if j imports power from i, $P_i^j(t)$, the opposite is true for i, which instead imports the same amount of power $P_j^i(t) = -P_i^j(t)$. Here, the positive sign indicates that the power is flowing from j to i. With algebraic manipulations, it is possible to express such a constraint as the linear combination of all the microgrid variables $x_i(t)$, i.e.,

$$\sum_{i=1}^{m} \hat{A}_i \cdot x_i(t) = \underline{0}, \tag{6.11}$$

where the matrixA_i has size $|\mathcal{E}| \times |\mathcal{N}_i|$ and assigns the decision variable $P_j^i(t)$ to the appropriate constraint $(i,j) \in \mathcal{E}$. Consequently, each row of A_i corresponds to the balance across a specific link, and involves only two variables $P_i^j + P_j^i = 0$. The balancing constraint is then extended over the optimisation horizon as

$$\sum_{i=1}^{m}((\hat{A}_i \otimes I_T) \cdot [x_i^\top(0)..x_i^\top(T)]^\top = \sum_{i=1}^{m} A_i \cdot x_i = \underline{0}, \tag{6.12}$$

where I_T is the identity matrix with dimension $T \times T$, $\otimes$ is the Kronecker product and $(\hat{A}_i \otimes I_T) = A_i$.

Uncertain constraint on aggregated power generation: Finally, we present the constraints that encode the VPP's ability to consistently generate sufficient power to maintain network balance and to provide aggregated power to the TSO when necessary.

Due to the unpredictable nature of the aggregated PV production $\hat{P}_{pv}(t)$ and load consumption $\hat{P}_l(t)$,

$$\hat{P}_{pv}(t) = \sum_{i=1}^{m} P_{pv}^i(t), \quad \hat{P}_l(t) = \sum_{i=1}^{m} P_l^i(t), \tag{6.13}$$

the amount of power that is effectively produced is not known. It depends on the uncertain realisation of the random variable: $\delta(t) = \hat{P}_{pv}(t) - \hat{P}_l(t) \in \Delta$; therefore, the aggregated power the VPP can deliver to the TSO is limited to

$$P_n(t,\delta) = \sum_{i=1}^{m}(P_{gen}^i(t) + P_{bess}^i(t)) + \delta(t). \tag{6.14}$$

The latter constraint concerns the active power balance at a network scale. In fact, according to

Eq. (6.10), the sum of the power exchanged among the agents is always equal to zero, meaning that we do not need to consider internal exchanges when calculating the delivered power.

The VPP treat the amount of net production to be fulfilled as an exogenous signal $P_n^{ref}(t)$, sent by the TSO (which is known). The following chance constraint encodes the uncertain global constraint

$$\mathbf{P}_\delta\{\delta \in \Delta : P_n(t,\delta) - P_n^{ref}(t) \geq 0\} \geq 1 - \epsilon. \tag{6.15}$$

Eq. (6.15) encode the probability of fulfilling the TSO request, within a risk level ϵ. In a manner similar to equation (6.10), we can concisely rewrite this constraint for $t \in [0, T-1]$ as

$$\mathbf{P}\{\delta \in \Delta : \sum_{i=1}^{m} H_i \cdot x_i + H_\delta \delta \geq 0\} \geq 1 - \epsilon, \tag{6.16}$$

for proper matrices H_i and H_δ.

6.3. VPP architecture

In this section, we present the VPP distributed architecture.

6.3.1. Centralised optimisation

Before explaining the distributed architecture for optimising the VPP, it is didactic to state the centralised optimisation problem explicitly and use the scenario approach to approximate the stochastic solution.

The optimisation is conducted over a full day, as per the day-ahead market ($T = 24$; see chapter 4), within a receding-horizon framework. However, the proposed approach applies to different time horizons and sampling periods as well.

At each time instant, the stochastic centralised optimal control problem being solved is

$$\mathcal{P} : min_{\{x_i\}_{i=1}^m} \sum_{i=1}^{m} f_i(x_i) \quad s.t., \tag{6.17a}$$

$$x_i \in X_i \ \forall i = 1..m, \tag{6.17b}$$

$$\sum_{i=1}^{m} A_i \cdot x_i = 0, \tag{6.17c}$$

$$\mathbf{P}\{\delta \in \Delta : \sum_{i=1}^{m} H_i \cdot x_i + H_\delta \cdot \delta \geq 0\} \geq 1 - \epsilon. \tag{6.17d}$$

The problem outlined in Eq. (6.17) is challenging to solve even in a centralised way, due to the presence of the chance-constraint Eq. (6.15)

It is possible to employ the scenario theory (chapter 2) to derive an approximate solution by sampling a finite set of uncertainties δ from historical data to build the deterministic global coupling constraints for each scenario.

The now deterministic optimal control problem is built by approximating the stochastic constraint (6.16) through a set of N linear constraints, where N is the number of scenarios δ_κ $\kappa = 1...N$

$$\sum_{i=1}^{m} H_i \cdot x_i + H_\delta \cdot \delta_k \geq 0 \quad \forall \kappa = 1, \ldots, N, \tag{6.18}$$

or, equivalently,

$$\bigcap_{\kappa=1}^{N} (\mathcal{F}_\kappa) \; with \; \mathcal{F}_\kappa := \{x : \sum_{i=1}^{m} H_i \cdot x_i + H_\delta \cdot \delta_\kappa \geq 0\}. \tag{6.19}$$

By selecting a sufficiently large N, the solution of the scenario program $\{x_i^*\}_{i=1}^m$ satisfies a priori probabilistic guarantees under unseen realisations of the uncertainty δ.

Given the structure of constraint (6.16), based on [DRS24] and the reasoning in chapter 2, the following theorem is established:

Theorem 6.1. *Given a risk parameter $\epsilon \in (0,1)$ and a confidence parameter $\beta \in (0,1)$, the solution $x^* = [x_1^\top, \ldots x_m^\top]^\top$ to the scenario optimization problem* (6.17)*, replacing constraint* (6.16) *by* (6.18)*, is so that*

$$\mathbf{P}[\delta \in \Delta : \sum_{i=1}^{m} H_i \cdot x_i^* + H_\delta \cdot \delta < 0] \leq \epsilon, \tag{6.20}$$

with high confidence $1 - \beta$, if N is such that the relation

$$\epsilon = 1 - \sqrt[N-\rho]{\frac{\beta}{\binom{N}{\rho}}} \tag{6.21}$$

is satisfied, and $\rho = T$ is the Helly's dimension of the problem.

Given the scenario dataset $S^N = \{\delta_1, \ldots, \delta_N\} \in \Delta^N$, with scenarios independently chosen from the uncertain set $\delta_\kappa \in \Delta$, Theorem 6.1 states that the probability of violation remains within the prescribed risk level ϵ with confidence $1 - \beta$ provided that the number of scenarios satisfies Eq.(6.21).

However, in the VPP context, the support dimension ρ is strictly related to the number of variables of the problem, which scales with the network size m, making the computation of a solution challenging as the VPP's size increases.

However, by leveraging the chance-constraint special structure (Eq. (6.18), and employing the lemma 2 from [DRS24], we can find a much lower bound on ρ, which is only depends on the

row-span of the matrix H_i, that is the optimisation horizon: $\rho \leq r = T$.

This bound separates the complexity of handling uncertainty (N) from the network scale (m), making the method suitable for VPP regardless of size.

The resulting scenario program, coming from the approximation of the chance-constraint Eq. (6.17) with Eq. (6.18), becomes a Quadratic Program (QP), which is not only suitable to be solved using standard centralised techniques, but can also be decomposed into smaller problems, one for each agent, to be solved in a distributed way, avoiding all the complications of centralised schemes highlighted in chapter 4.

6.3.2. Distributed Algorthm

The problem Eq. (6.17) can be decomposed into smaller sub-problems, one for each of the m microgrids, and solved with a distributed algorithm, as in [FP23].

This algorithm is tailored to track both equality and inequality coupling constraints, with the same structure as the power flow balance (Eq. (6.10)) and the scenario global coupling constraint (Eq. (6.18)).

The procedure is summarised in Algorithm 0.

Algorithm initialisation: The algorithm begins with a preliminary initialisation phase. First, each microgrid extracts $N_i \geq \lceil N/m \rceil$ scenarios of PV generation and load profiles for the entire prediction horizon T from recorded data (*local scenario extraction*, step 2). Then, all microgrids broadcast their local scenarios to form the collective dataset $\tilde{D}$ (*scenario broadcast*, step 3). Subsequently, each microgrid initialises its dual variables λ_i and μ_i, the primal variable $x_i \in X_i$, and the associated constraint residuals $g_i(x_i)$, σ_i, and d_i (*initialisation*, steps 4–6).

Iterative convergence: After the initialisation, each microgrid iteratively updates the primal and dual variables in parallel for each iteration $\tau = 1, \ldots, \tau_{\max}$ until convergence requirements are met.

At each iteration, a *consensus step* (steps 8–10) is performed, during which each microgrid communicates its current dual variables λ_i, μ_i and residual estimates $g_i(x_i)$, σ_i, d_i to each neighbour. Then, each microgrid calculates the average of these quantities over the neighbourhood set $\mathcal{N}_i$, weighting the received information by the coefficients a_i^j which reflect the network topology.

Next, in the *primal solution step* (step 11), each microgrid solves a local optimisation problem augmented with proximal terms to handle the coupling constraints, thus obtaining a new candidate primal solution. Finally, in the *sub-gradient ascent step* (steps 12–16), each microgrid updates its residual estimates based on the latest primal solution and performs a gradient step on its local dual variables ℓ_i and m_i.

Algorithm 6.1 Augmented Lagrangian Tracking (ALT)

Initialization:

$\tilde{D}_i := \{\delta_1..\delta_{\lceil \frac{N}{m} \rceil + 1}\} \forall i = 1..m$ ▷ *Generate scenarios*

$\tilde{D} = \bigcap_{i=1}^{m} \tilde{D}_i$ ▷ *Broadcast scenarios*

$x_i^0 \in \mathbb{R}^{n_i}$, $\sigma_i^0 \in \mathbb{R}_+^q$ ▷ *Initialize variables*

$\lambda_i^0 \in \mathbb{R}^p$, $d_i^0 = -(A_i x_i^0 - b_i)$

$\mu_i^0 \in \mathbb{R}^q$, $g_i^0 = -(h_i(x_i^0) + \sigma_i^0)$

while not converged **do**

Consensus step

$\ell_i^\tau = \sum_{j \in \mathcal{N}_i} w_{ij} \lambda_j^\tau, \quad m_i^\tau = \sum_{j \in \mathcal{N}_i} w_{ij} \mu_j^\tau$

$\delta_i^\tau = \sum_{j \in \mathcal{N}_i} w_{ij} d_j^\tau, \quad \gamma_i^\tau = \sum_{j \in \mathcal{N}_i} w_{ij} g_j^\tau$

Primal update

$$x_i^{\tau+1} \in \arg\min_{x_i \in X_i} \Bigg\{ f_i(x_i) + (\ell_i^\tau)^\top A_i x_i + \frac{c}{2} \| A_i x_i - A_i x_i^\tau - \delta_i^\tau \|^2 + \frac{1}{2c} \| \max\{m_i^\tau + c(h_i(x_i) - h_i(x_i^\tau) - \sigma_i^\tau - \gamma_i^\tau), 0\} \|^2 \Bigg\}$$

Dual update

$\sigma_i^{\tau+1} = \gamma_i^\tau - h_i(x_i^{\tau+1}) + h_i(x_i^\tau) + \sigma_i^\tau - \frac{1}{c} m_i^\tau$

$\sigma_i^{\tau+1} = \max\left\{\sigma_i^{\tau+1}, 0\right\}$

$d_i^{\tau+1} = \delta_i^\tau - A_i x_i^{\tau+1} + A_i x_i^\tau$

$g_i^{\tau+1} = \gamma_i^\tau - (h_i(x_i^{\tau+1}) + \sigma_i^{\tau+1}) + (h_i(x_i^\tau) + \sigma_i^\tau)$

$\lambda_i^{\tau+1} = \ell_i^\tau - c d_i^{\tau+1}, \quad \mu_i^{\tau+1} = m_i^\tau - c g_i^{\tau+1}$

$\tau \leftarrow \tau + 1\text{s}$

end while

Assumptions: The following assumptions on the graph are needed to guarantee the convergence of the distributed optimisation scheme [FP23].

Hypothesis 6.1 (connectivity). *Graph $(\mathcal{V}, \mathcal{E}_\infty)$ is strongly connected, i.e., for any pair of nodes, there exists a path of directed edges that connects them. Moreover, there exists a time-period $\delta t > 1$ such that for every $(i,j) \in \mathcal{E}_\infty$ agent i receives information from a neighbour $j \in \mathcal{N}_i$ at least once every δt iterations.*

This assumption assumes that all other nodes influence every node and avoids situations where parts of the network are unreachable.

Hypothesis 6.2 (Information mixing). *It exists $\eta \in (0,1)$ such that for all $i,j \in \{1,\dots,m\}$ and all $t \geq 0$,*

$$a_{ij}^{(t)} \in [0,1), \quad a_{ii}^{(t)} \geq \eta, \quad \wedge \quad a_{ij}^{(t)} > 0 \Longrightarrow a_{ij}^{(t)} \geq \eta.$$

Moreover, for all $t \geq 0$,

1. $\sum_{j=1}^{m} a_{ij}^{(t)} = 1$ *for all* $i = 1,\dots,m$,
2. $\sum_{i=1}^{m} a_{ij}^{(t)} = 1$ *for all* $j = 1,\dots,m$.

This second hypothesis ensures information mixing at a non-decreasing rate, meaning that all agents are equally influenced by all their neighbours, avoiding situations in which a single neighbour can influence a node.

Both these hypotheses exclude degenerate cases that are not of interest, such as disconnected communication networks or pathological information flows. In practice, it is the communication network design that must ensure the connectivity of all components and sufficient information exchange. Therefore, such assumptions must be interpreted as design requirements to avoid degenerate configurations in which agents cannot reach consensus in a multi-agent setting.

Under these assumptions, the algorithm is proven to converge even in the presence of temporary miscommunication in the networks [FP21], making the VPP robust to temporary microgrid disconnections. These assumptions are paramount to the algorithm's convergence to the centralised solution, which in turn is necessary for the reasoning in 6.3.1 to be valid.

Finally, we stress the role of the new bound on the Support dimension ρ; if the bound were dependent on the number of agents, the number of scenarios required to ensure a-priori violation levels would increase with m, making all the algorithms non-scalable. Additionally, each agent would have to communicate a vector of dual variables, λ, that is increasingly large, limiting the network's communication capabilities.

6.4. Simulation results

To verify the proposed VPP, we tested it for a maximum violation parameter $\epsilon = 10\%$ and increasing values of m: $(m = 8, |\mathcal{E}| = 16)$, $(m = 20, |\mathcal{E}| = 70)$, $(m = 50, |\mathcal{E}| = 125)$, $(m = 100, |\mathcal{E}| = 150)$. These parameters imply a number of scenarios $N = 362$ (Eq. (6.21))

We simulate each VPP daily and repeat the simulation 1000 times, selecting different scenarios S^N, then we test the solution x_N^*, obtained by solving the N-scenario program in eq. (6.17) with the distributed algorithm 0, against $Mc = 1000$ unseen scenarios to estimate the violation probability. All the data is taken from the open-source database in [VAK19].

6.4.1. Algorithm for the randomised generation of VPPs

To simulate heterogeneous VPPs, we generate randomised network instances that capture both the topological structure and the operational characteristics of interconnected microgrids.

The network topology is constructed as an initially fully connected graph with m agents; then, $|m - \mathcal{E}|$ edges are randomly removed, ensuring that each node retains a minimum number of connections, i.e., the network's adjacency matrix W remains strongly connected. Finally, the network's adjacency matrix W is normalised and, optionally, processed using the Sinkhorn-Knopp algorithm to obtain a doubly stochastic consensus matrix, facilitating distributed coordination schemes.

Definition 6.1 (Doubly stochastic matrix)**.** *A matrix $W \in \mathbb{R}^{n\times n}$, with elements $a_{i,j}$, is said to be* doubly stochastic *if all its entries are non-negative and the sum of each row and each column equals one, i.e.,*

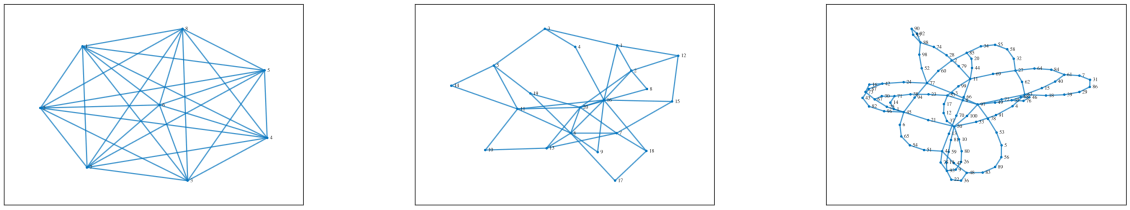
$$a_{i,j} \geq 0, \qquad \sum_{j=1}^{m} a_{i,j} = 1, \qquad \sum_{i=1}^{m} a_{i,j} = 1, \quad \forall j = 1, \dots, m. \tag{6.22}$$

This procedure yields a sparse yet connected graph suitable for decentralised coordination, which respects both hypotheses 6.1 and 6.2.

Each microgrid agent is then assigned generation and storage capabilities according to predefined aggregated levels (High, Medium, Low). Power ranges, ramp rates, and quadratic cost functions characterise Generation units. In contrast, BESSs are defined by capacity limits, operational ranges, costs, and initial state-of-charge. The assignment of these levels follows a probabilistic distribution consistent with the chosen aggregated level, allowing for variability across agents.

Power exchange capabilities between microgrids are represented by edge weights, which are randomly sampled from discrete cost and power ranges according to a prescribed probability distribution, reflecting different levels of exchange intensity.

This methodology enables the creation of realistic, heterogeneous VPP instances for testing decentralised optimisation and control algorithms, as depicted in Fig. 6.1.



(a) Fully connected VPP with $m = 8, \mathcal{E} = 28$.

(b) Partially connected VPP with $m = 20, \mathcal{E} = 30$.

(c) Sparsely connected VPP with $m = 100, \mathcal{E} = 90$.

Figure 6.1: VPP topologies for different values of $m, \mathcal{E}$.

6.4.2. Aggregated power production probabilistic guarantees

In simulation, the TSO requires an aggregated ramp-up service of $\Delta P_{ramp} = 40KWh$ between 4 a.m. and 2 p.m.: $t \in (4, 14)$. During the first part of the day, the VPP only needs to be self-sufficient, whereas after 2 p.m., it must transmit a constant power of $400kW$ for the remaining hours. The request is the dashed red line in Fig. 6.2.

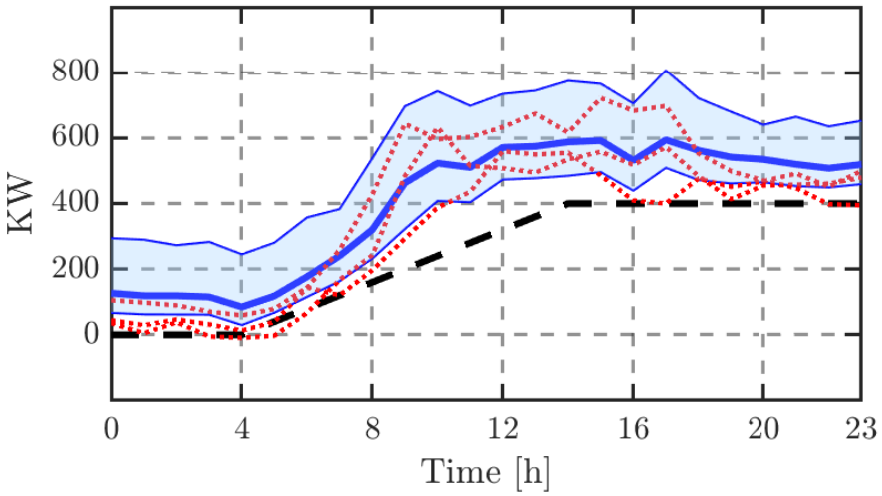


Figure 6.2: Aggregated power production $P_n(t, \delta)$: mean value (blue line), 1-ϵ level (light blue area), violating unseen scenarios (dotted red line). Minimum aggregated power required $P_n^{ref}(t)$ (dashed black line) [DR25]

Figure 6.2 shows an instance of the aggregated power production $P_n(t, \delta)$ for $m = 4$, the average transmitted power (blue line) is accompanied by the $1 - \epsilon$ robust levels (light blue area). The $1 - \epsilon$ levels represent the area containing all the possible trajectories of $P_n(t, \delta)$, excluding the

"worst" ϵ trajectories which violate the constraints. Examples of such trajectories are depicted in red dotted lines.

We observe that the VPP correctly schedules the aggregated power in advance, irrespective of uncertainty in production and consumption, while minimising production and transmission costs across the entire network and thereby optimising the costs.

To verify that the a-priori violation levels are satisfied, we can define the violation probability and approximate it through sampled violation levels for each daily simulation as

$$V_N(x^*) = \mathbf{P}[\delta \in \Delta : \sum_{i=1}^{m} H_i \cdot x_i^* + H_\delta \cdot \delta < 0]. \tag{6.23}$$

The resulting sampled violation $\hat{V}_N(x^*)$ is

$$\hat{V}_N(x^*) = \frac{1}{Mc} \sum_{j=1}^{Mc} [\mathbf{1}_{V_N}(x^*, \delta_j)], \tag{6.24}$$

where $1_{V_n}(x)$ is the indicator function, indicating if the scenario δ_i violates the constraint.

$$\mathbf{1}_{V_n}(x^*, \delta_j) := \begin{cases} 1 & \text{if } (\sum_{i=1}^{m} H_i \cdot x_i^* + H_\delta \cdot \delta_j < 0) \ , \\ 0 & \text{if } (\sum_{i=1}^{m} H_i \cdot x_i^* + H_\delta \cdot \delta_j \geq 0) \ . \end{cases}$$

Figure 6.3 illustrates the simulation results based on $N = 362$ scenarios (red boxes), and $N = 100$ scenarios (blue boxes).

Each box displays the probability distribution of the violation $V_N(x^*)$, sampled through the Monte Carlo simulations and in accordance with Eq.(6.15). The box plot shows the maximum and minimum violation levels, the mean, the median, and the σ confidence level for each Monte Carlo batch.

Notably, when using $N = 362$ scenarios (computed through Eq.(6.21)) we consistently remain way below the maximum prescribed violation level of $\epsilon = 10\%$ (red-dashed line), indicating an excess in conservativeness.

To reduce this conservativeness, we can lower N, pushing the beta distribution towards the limit, but renouncing the a-priori guarantees.

Our analysis also strengthens the theoretical result in chapter2, as the violation-level distribution remains stable as the network size increases, even if the number of scenarios is kept constant.

This finding suggests that the VPP can separate the complexity of uncertainty handling

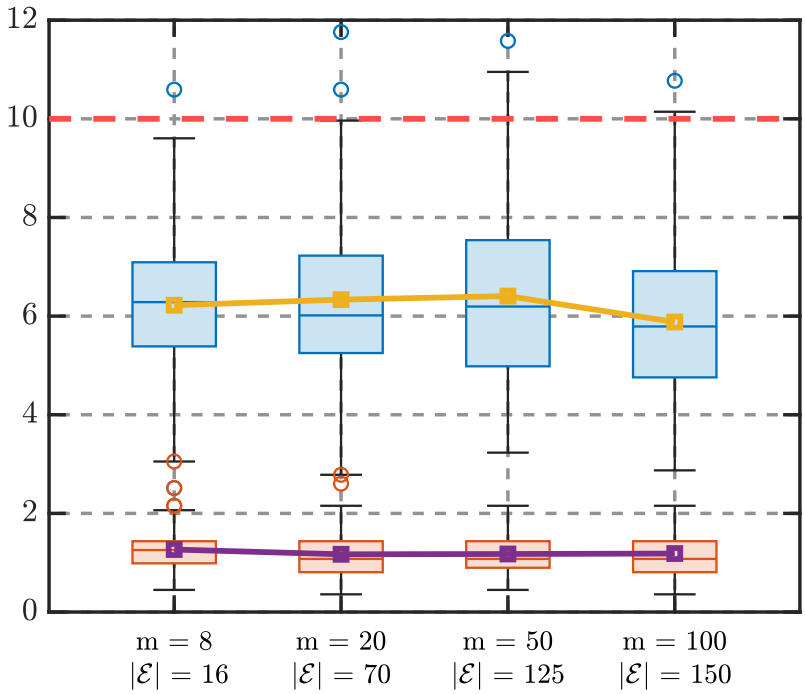


Figure 6.3: Violation probability $V_N(x^*, \delta)$, in red $N = 362$, in blue $N = 100$, yellow and purple line: average value of the violation [DRS24]

from the network scale, enabling it to maintain operational efficiency without constraining the VPP network size and scalability.

A physical interpretation of the probabilistic violation level can be given from the perspective of a TSO. When a violation occurs, i.e., when the aggregated power request is not satisfied, the TSO must activate reserve resources to maintain the system's power balance. The violation probability is interpreted as the likelihood that the TSO unexpectedly draws on balancing reserves, potentially resulting in higher costs and fees for the VPP and its microgrids.

6.4.3. Convergence of the Distributed Algorithm

The distributed scheme, thanks to assumptions 6.1 and 6.2, converges to the centralised solution as the number of iterations increases.

In Fig. 6.4, the upper plot shows the evolution of the normalised absolute error between the distributed solution $f_\tau(x_\tau)$ and the optimal centralised solution $f(x^*)$. The convergence ensures that the analysis on the scenario approach can be applied to the solution of the distributed algorithm as well.

The lower plot displays the residuals of the coupling equality constraints $e = P_i^j - P_j^i$ for all $(i, j) \in \mathcal{E}$, which represent the convergence of the power flow to a feasible one.

Both the residuals converge to a re small value ($\leq 10^{-4}$) within the first 100 iterations.

Both plots demonstrate that all residuals converge to a reasonably small value (less than or equal to 10^{-4}) within a maximum of $\tau = 100$ iterations.

It is worth noting that the convergence depends on both the step size of the algorithm c in 0, and the connectivity of the graph $\mathcal{G}$, which affect the consensus dynamics on the dual variables.

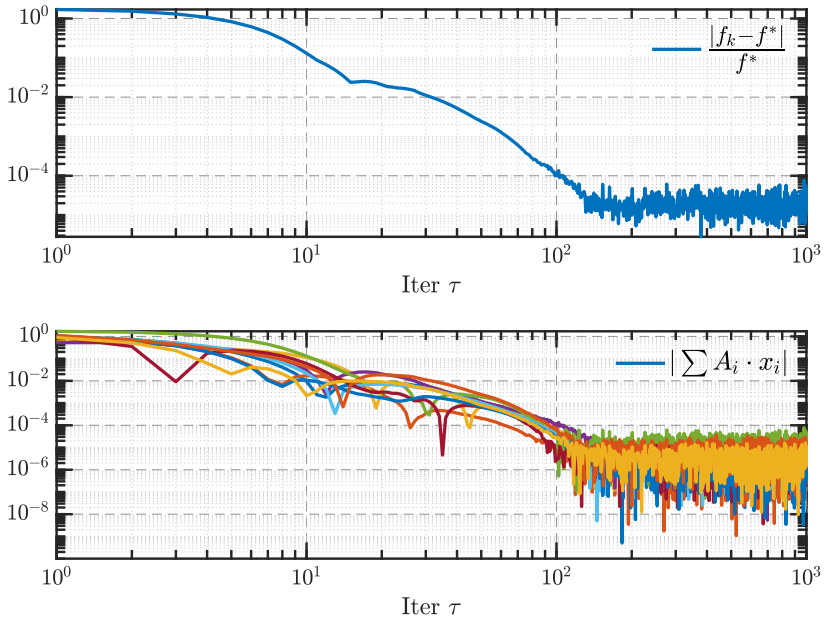


Figure 6.4: Convergence of the normalised absolute error

To conclude, we present a distributed VPP architecture based on the coordination of multiple interconnected microgrids. The algorithm leverages theoretical results on the scenario approach over distributed systems with global uncertain coupling constraints [DRS24] to make the system scalable with the number of agents m.

Chapter 6 Bibliography

[DR25] A. Del Duca and F. Ruiz. "Distributed Virtual Power Plant Optimization via Scenario Approach". In: *2025 IFAC Workshop on Smart Energy Systems for Efficient and Sustainable Smart Grids and Smart Cities (SENSYS)*. Bari, Italy, 2025.

[DRS24] A. D. Del Duca, F. Ruiz, and R. Scattolini. "Distributed Stochastic Optimisation with Uncertain Coupling Constraints". In: *2024 IEEE 63rd Conference on Decision and Control (CDC)*. Milan, Italy, 2024, pp. 4431–4436. DOI: 10.1109/CDC56724.2024.10886359. URL: https://hdl.handle.net/11311/1286946.

[FP21] Alessandro Falsone and Maria Prandini. "A Distributed Dual Proximal Minimization Algorithm for Constraint-Coupled Optimization Problems". In: *IEEE Control Systems Letters* 5.1 (2021), pp. 259–264. DOI: 10.1109/LCSYS.2020.3001427.

[FP23] Alessandro Falsone and Maria Prandini. "Augmented Lagrangian Tracking for distributed optimization with equality and inequality coupling constraints". In: *Automatica* 157 (2023), p. 111269. ISSN: 0005-1098. DOI: https://doi.org/10.1016/j.automatica.2023.111269.

[Nau+21] James Naughton et al. "Co-Optimizing Virtual Power Plant Services Under Uncertainty: A Robust Scheduling and Receding Horizon Dispatch Approach". In: *IEEE Transactions on Power Systems* 36.5 (2021), pp. 3960–3972. DOI: 10.1109/TPWRS.2021.3062582.

[VAK19] Karina Vink, Eriko Ankyu, and Michihisa Koyama. "Multiyear microgrid data from a research building in Tsukuba, Japan". In: *Scientific Data* 6.1 (Feb. 2019), p. 190020. ISSN: 2052-4463. DOI: 10.1038/sdata.2019.20. URL: https://doi.org/10.1038/sdata.2019.20.

7 | Dual SME MPC for Aggregated TER Demand Responsee

The technological shift toward renewable energy has transformed the power grid paradigm, moving from a generator-driven monolithic system to an innovative framework in which loads adapt their consumption to match available generation. In this context, consumers become active participants in energy dispatch through Demand Response (DR) programs, thereby enhancing grid flexibility [AE07; PM19]. To enable this kind of DR service, flexible resources are required to balance the grid while maintaining operational reliability. In such a context, Flexible Loads (FLs) enhance energy management by adapting their consumption to grid conditions [Alz+23].

As the electrification of Heating, Ventilation, and Air Conditioning (HVAC) systems rises, devices such as refrigerators, boilers, space heaters, and heat pumps are becoming increasingly relevant in the energy landscape [MCC24]. Additionally, with the rapid rise of artificial intelligence and the expansion of digital infrastructure, the transformation towards an AI-centric society is increasing the already high relevance of data centres, where cooling systems can account for 30-50% of total energy consumption [Oh+25]. The energy required for cooling and heating is increasing, positioning such thermo-electric systems as valuable assets and strong candidates for DR initiatives.

The broad class of Thermostatically Controlled Loads (TCLs) has gained particular attention across various types of FL. These devices consume electrical power to maintain a set temperature. Due to the thermal inertia of these systems, the temperature setpoint can be temporarily adjusted. This controlled deviation enables reduced power consumption without compromising performance or user comfort. Due to the small size of these devices, it is advantageous to aggregate them to provide a single service to the Distribution System Operator (DSO). This process is usually handled by an aggregator, a central unit that measures the total flexible response power and sends commands to individual TCLs.

These commands instruct the TCLs to adjust their reference temperature, enabling them to provide a unified service to the DSO [Dia+20; CDR22]. However, achieving this objective requires an adequate understanding of the dynamic behaviour of individual loads. Deriving analytical models that accurately capture these dynamics is often challenging, as it involves complex, nonlinear, and time-varying processes that are both difficult and time-consuming to model [BK22]. Learning these dynamics is essential for the aggregator to coordinate multiple devices effectively.

Given the inherent complexity of modelling individual systems, a more efficient approach for the aggregator is to learn the aggregate dynamics in real time. Moreover, it must handle model uncertainties by balancing immediate control performance with the active learning of the aggregate dynamics, all while respecting operational constraints such as power ramp-up and ramp-down limits. Balancing power control with the need to learn the system dynamics is a well-known challenge in control theory, as discussed in Ch. 3.

7.1. Contributions

In this chapter, we demonstrate the effectiveness of the adaptive Model Predictive Control (MPC) algorithm (Alg. 3.2) for identifying and controlling the aggregated dynamics of a group of TCLs in DR applications. The algorithm identifies two distinct models for the services the aggregator must provide to the Distribution System Operator (DSO): upward and downward reserves. The algorithm minimises the Feasible Parameter Set (FPS) volume, used as a metric of system uncertainty, while tracking a desired power request and satisfying stochastic constraints, extending the results developed in [Mez+25; SCR25].

The controller addresses two critical challenges: minimising tracking errors between the consumed power and the requested service, and accounting for uncertainties in the aggregated model. To tackle these issues, it actively learn a model for the upward/downward reserves exploiting the dual capabilities previously discussed, which emerge from the minimisation of the volume of the FPS. This approach balances achieving control objectives, such as optimising power delivery, with exploring the range of model parameters. By doing so, the controller enhances its predictive capabilities, leading to more reliable and efficient management of energy resources within the grid.

7.2. Problem Definition

Large populations of small TCLs can provide fast, reliable balancing services when coordinated by an aggregator. Thermo-Electric Refrigerators (TERs) are particularly appealing for DR due to their precise temperature control and the ability to shift steady-state power by adjusting set points. European balancing programs, such as Frequency Containment Reserve (FCR), Frequency Restoration Reserve (FRR), and Replacement Reserves (RR), impose specific timing and ramping requirements. For example, FCR requires a settling time of less than or equal to 30 seconds and 50% of the power to be provided within 15 seconds. These requirements must be adhered to by the aggregator and its low-level controllers. Our approach to TER modelling and aggregation strategy is based on the methods outlined in [Dia+20].

7.2.1. Thermo-Electric (TEC) module heat flows

This section summarises the TER thermal and electrical model used for control design and identification (models and parameters are those validated experimentally in [Dia+20]). Following standard Peltier module modelling, the heat flows Q_c, Q_h at the cold and hot side temperatures T_c, T_h are definied as

$$Q_c = \alpha\, T_c\, I_p - \tfrac{1}{2} R_p I_p^2 - K_p(T_h - T_c), \tag{7.1a}$$

$$Q_h = \alpha\, T_h\, I_p + \tfrac{1}{2} R_p I_p^2 - K_p(T_h - T_c), \tag{7.1b}$$

where α is the Seebeck coefficient, R_p the electrical resistance of the TEC, K_p the thermal conductance of the module and I_p the electrical current through the module. These expressions match the standard energy balances for a Peltier device.

The Seebeck voltage and electrical relations are

$$V = \alpha(T_h - T_c), \qquad P_{\text{in}} = V_{\text{in}} I_p, \qquad I_p = \frac{V_{\text{in}} - V}{R_p}. \tag{7.2}$$

where V_{in} is the DC voltage supplied to the TEC.

7.2.2. Thermal dynamic equations

Energy balances for the cold and hot thermal masses lead to

$$C_{TR}\,\dot{T}_c(t) = \frac{T_{\text{amb}}(t) - T_c(t)}{R_{TR}} - Q_c(t), \tag{7.3a}$$

$$C_{out}\,\dot{T}_h(t) = Q_h(t) - \frac{T_h(t) - T_{\text{amb}}(t)}{R_{out}}, \tag{7.3b}$$

where T_{amb} is the ambient temperature, R_{TR}, C_{TR} are respectively, the thermal resistance and the thermal capacily of the "Thermal Reservoir", which is the internal of the TER, and finally, R_{out}, C_{out} are respectively, the thermal resistance and the thermal capacily associated with the external ambient and the radiator. Experimental parameter identification and validation procedures are described in [Dia+20].

The dynamics of each TER can be represented as a nonlinear state-space model of the form

$$\begin{cases} \dot{x}_i = F_i(x_i(t), \beta_i(t)), \\ P_i(t) = G_i(x_i(t), \beta_i(t)), \end{cases} \tag{7.4}$$

where the functions $F()$ and $G()$ describe the dynamics of each thermoelectric system, including thermal loads, head flows, and Peltier module transduction, among others.

Locally, each TER can be regulated by means of an internal control loop, to track a set of three reference signals, encoded by the vector β_i:

$$\beta_i = \begin{cases} 0, & \text{Nominal set-point } T_{\text{set}} = T_{\text{nom}} \in \{3, ..7\}\,^\circ\text{C}, \\ 1, & \text{High set-point } T_{\text{set}} = 8\,^\circ\text{C}, \\ -1, & \text{Low set-point } T_{\text{set}} = 2\,^\circ\text{C}. \end{cases}$$

By changing the temperature set-point T_{set}, also the power consumed by the i^{th} TER, $P_i(t)$, changes. By coordinating multiple TER units, the aggregator can provide DR services to the DSO.

7.2.3. Aggregated dynamics and problem statement

We consider an aggregation of m heterogeneous TERs capable of providing upward/downward reserve services to the DSO. To provide the DR service, the aggregator issues a high-level command $u_k \in [-m, m], \quad u_k = \sum_{i=1}^{m} \beta_i$, representing the number of units that must adjust their temperature set point T_{set} to the scheduler, or TER selector. The scheduler

then determines which of the m TER units to query for the service, ensuring compliance with each TER's turn-on/turn-off and idle constraints, as detailed in [Dia+20]. It sends to each TER the corresponding vector $\beta_i(t) \in \{0, 1, -1\}^m$, indicating which TER units must modify their reference signal during the time interval $t \in [kT_s, (k+1)T_s]$.

By doing so, the aggregator can adjust the aggregate power consumption y_k to increase or decrease power demand as needed, while remaining within operational constraints.

$$y_k = \sum_{i=1}^{m} \big(P_{i,\text{nom}}(k\,T_s) - P_i(k\,T_s)\big). \tag{7.5}$$

Here, $P_i(k\,T_s)$ is the instantaneous electrical power of unit i (measured or inferred from voltage/current), $P_{i,\text{nom}}(k\,T_s)$ is the nominal baseline (corresponding to a temperature set-point of $T_{set} = T_{nom}$).

The TER control architecture is shown in Fig. 7.1. The scheme consists of three main components: the Dual MPC aggregator, which computes the control action u_k according to the required flexibility; the TER aggregator, which selects the TER units whose temperature reference β_i must be adjusted; and the individual TER units, which track the updated reference β_i and dynamically adapt their power consumption P_i accordingly.

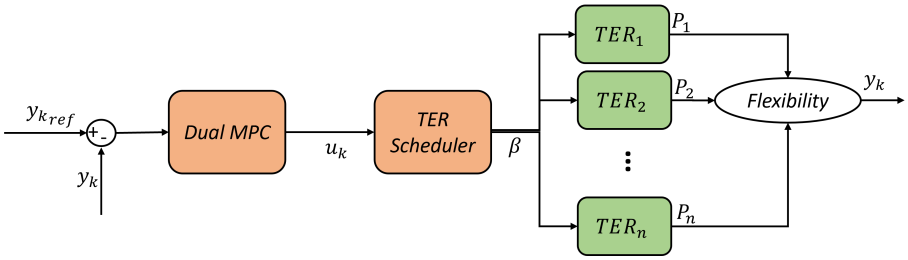


Figure 7.1: TER control architecture

The challenge the aggregator must address is that the aggregated dynamics are unknown a priori and cannot be identified offline without disrupting operations. This motivates the use of the Dual MPC framework described in chapter 3, since the presented algorithm can actively explore the aggregate dynamics, starting from a priori information, to learn a more accurate model of the system.

The Dual MPC optimization for the TER aggregation problem is:

$$\min_{u_{k|k},\dots,u_{N-1|k}} \sum_{i=0}^{N-1} Q(y^{(c)}_{k+i|k} - y^{ref}_{k+i|k})^2 + Ru^2_{k+i|k} + \tag{7.6a}$$

$$\lambda \cdot \hat{V}_{k+i|k}(y^{(c)}_{k+i|k}, u_{k+i|k}), \tag{7.6b}$$

s.t.

$$y^{(j)}_{k+i+1|k} = \theta_j^\top \varphi^{(j)}_{k+i|k}, \forall i = 0,\dots,N-1, j = 1,\dots,K, \tag{7.6c}$$

$$y^{(c)}_{k+i+1|k} = \theta_c^{k\top} \varphi^{(c)}_{k+i|k}, \forall i = 0,\dots,N-1, \tag{7.6d}$$

$$\varphi^{(j)}_{0|k} = \varphi_k \quad j = 1,\dots,K+1, \tag{7.6e}$$

$$\theta_j \in \Theta_k \quad \forall j = 1,\dots,K+1, \tag{7.6f}$$

$$y^{(j)}_{k+i+1|k} \in [y_{min}, y_{max}], \forall i = 1,\dots,N-1, j = 1,\dots,K, \tag{7.6g}$$

$$u_{k+i|k} \in [-m, m] \quad \forall i = 0,\dots,N-1. \tag{7.6h}$$

7.3. Simulations

The system comprises $m = 100$ independent TERs, each one with its own dynamics. The aggregated dynamics are discretised with sampling time $T_s = 1$s, which is the sampling time of the controller. The input/output aggregate system, with input u_k and output y_k behaves as a hybrid continuous time linear model of order 2 (two poles and two zeros), is identified in [Dia+20], and is used here as a baseline for the true system parameters $\theta^\star$ as follows:

$$G_{up}(s) = \frac{25.95s^2 + 0.017s + 9.3 \times 10^{-6}}{s^2 + 1.1 \times 10^{-3}s + 6.8 \times 10^{-7}}, \tag{7.7a}$$

$$G_{down}(s) = \frac{-17.35s^2 - 0.018s - 3.6 \times 10^{-6}}{s^2 + 2.3 \times 10^{-3}s + 3.7 \times 10^{-7}}. \tag{7.7b}$$

The manipulated variable corresponds to the aggregated control action u_k, e.g., the number of TERs used to increase or decrease the temperature reference, while the output represents the aggregated variation in the power response y_k.

The ARX model does not model each individual TER; instead, it relies on the smoothing effect of the aggregation, which compensates individual non-linearities through averaging, to model the resulting low-order linear dynamic, valid in the relevant operating region, correctly capturing the collective dynamic of the system.

The aggregated dynamics is sampled through a Sampler-and-Hold with the same sampling time $T_s = 1$s; the upward/downward dynamics can therefore be modelled as a linear

Autoregressive with Exogenous input (ARX) model

$$A(z^{-1})Y(z^{-1}) \approx B(z^{-1}) \cdot U(z^{-1}) + e(z^{-1}), \tag{7.8}$$

with $A(z^{-1})$ and $B(z^{-1})$ containing the parameters $\theta \in \mathbb{R}^n$ being the unknown parameter vector ($n = n_a + n_b = 5$), i.e.,

$$\begin{aligned} A(z^{-1}) &= 1 + a_1\, z^{-1} + a_2\, z^{-2}, \\ B(z^{-1}) &= b_0\, z^{-1} + b_1\, z^{-2} + b_2\, z^{-3}, \end{aligned}$$

and the signal e_k is unknown but (pointwise) bounded with bound $e_k \leq 5$ kW.

7.3.1. Initial Feasible Parameter Set

Considering the second-order ARX model in the causal form as in (7.8), the characteristic polynomial of the dynamical system is

$$p(z) = z^2 - \theta_1 z - \theta_2. \tag{7.9}$$

The discrete-time system is *asymptotically stable* (Schur stable) if and only if all roots λ_i of $p(z)$ lie strictly inside the unit disk, i.e.,

$$|\lambda_i| < 1, \qquad i = 1, 2. \tag{7.10}$$

For a quadratic polynomial $z^2 - \theta_1 z - \theta_2$, the classical Schur (Jury) conditions are

$$1 - \theta_1 - \theta_2 > 0, \tag{7.11}$$

$$1 + \theta_1 - \theta_2 > 0, \tag{7.12}$$

$$|\theta_2| < 1. \tag{7.13}$$

Thus the stability region in the (θ_1, θ_2)-plane is the intersection of the three half-spaces defined by (7.11)–(7.13). Equivalently, the region can be written in the matrix inequality form

$$A_\theta\, \theta \;\leq\; b_\theta, \tag{7.14}$$

with

$$A_\theta = \begin{bmatrix} -1 & -1 \\ 1 & -1 \\ 0 & 1 \\ 0 & -1 \end{bmatrix}, \qquad b_\theta = \begin{bmatrix} -1 \\ 1 \\ 1 \\ 1 \end{bmatrix}.$$

The three inequalities define a convex polytope (triangle) in the parameter plane. For any θ inside this polytope, the characteristic polynomial (7.9) has both roots inside the unit disk, and the corresponding autonomous second-order system is asymptotically stable. Instead, for parameters θ_i $\forall i = 3, 4, 5$ we have bound boxes, i.e.,

$$\theta_i \in [\theta_i^{min}, \theta_i^{max}], \qquad \forall i = 3, 4, 5. \tag{7.15}$$

We can rewrite all the prior information as the initial FPS $\mathcal{K}_0 = \{\theta : A_0\theta \leq b_0\}$.

The initial FPS $\mathcal{K}_0$ is shown in Fig. 7.4 as the blue area.

7.3.2. Simulation setting

During real-time operations, the DSO may request either upward or downward reserve from the TER aggregations, with a reference signal $y_k^{\text{ref}} \in [-1000, 1000]\,\text{kW}$. To accommodate this, the aggregator is designed to learn and employ two distinct models, G_{down} and G_{up}, depending on the required service type.

For operational safety, the aggregator must ensure that its actual output remains within the admissible range $y_k \in [-500, 500]\,\text{kW}$. As a result, there may be instances in which the requested reserve exceeds the maximum deliverable capacity. In such situations, the aggregator must robustly regulate the system while operating at the boundary of its safe constraints, due to the error e_k.

The active MPC developed in this work is tested against its passive counterpart, presented in [Mez+25], which cannot actively probe the system dynamics; instead, it relies on passively acquiring new data while controlling the system in a closed loop.

When $y_{k,ref} = 0\text{kW}$, meaning the DSO is not requiring any service, the passive aggregator controls the output power to remain at 0kW; instead, the active aggregator probes the system to reduce uncertainty in the model.

When a service is required, the active aggregator has a better model than the passive aggregator and, consequently, can better follow the step-reference.

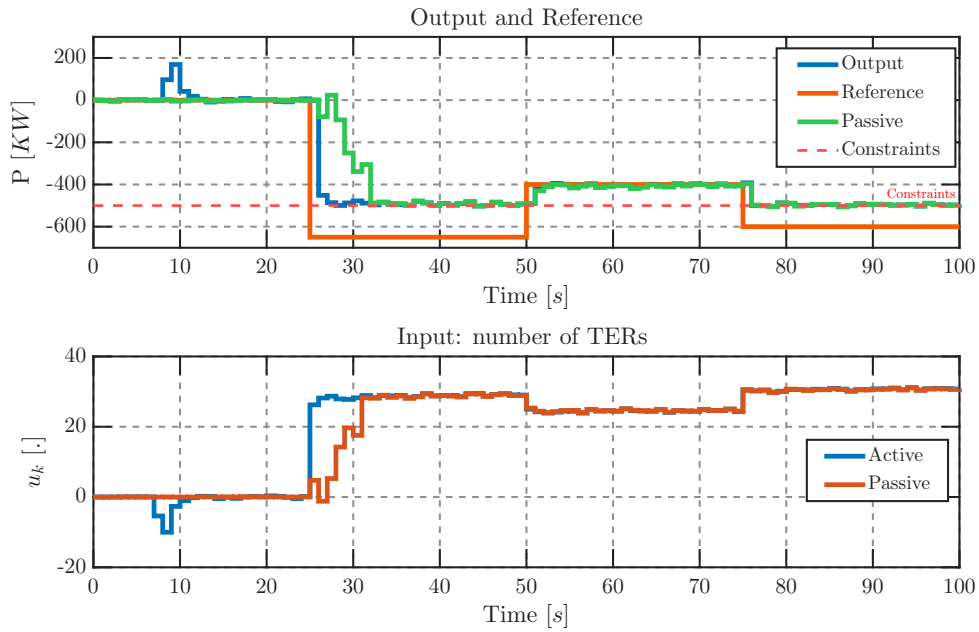


Figure 7.2: Top: Output trajectories of the passive (green) and active (blue) closed-loop systems, shown alongside the reference signal (orange) and the constraint (red, dotted). Bottom: Input trajectories of the passive (orange) and active (blue) closed-loop systems.

In Fig. 7.2, the trajectories of both active and passive closed-loop systems are depicted. At $k = 8$s, the active aggregator starts probing the dynamics to learn them; then, when the DSO asks for the reserve service at $k = 25$s, the active learning aggregator has better information about the system and follows the step signal closely.

Comparing the first step response for the active and passive controllers, we can compute important Key Performance Indicators (KPIs) such as the Integral Squared Error (ISE) and the Integral Absolute Error (IAE), respectively, defined as

$$ISE(k, k_0) = \sum_{i=k_0}^{k} (y_i - y_{i,ref})^2, \tag{7.16a}$$

$$IAE(k, k_0) = \sum_{i=k_0}^{k} |y_i - y_{i,ref}|. \tag{7.16b}$$

By computing these values, we can evaluate whether active learning is beneficial for better tracking of the reference, even during the initial learning phase between 8s and 15s. We compute such errors for $k_0 = 0$ and $k = 50$s, and we obtain the results presented in Table 7.1.

Despite an initial learning phase, the Active MPC significantly improves the closed-loop

Table 7.1: Performance metrics for the Active and Passive MPC comparison

KPI's	**ISE**	**IAE**
Active MPC	293,961.8	1,107.2
Passive MPC	996,898.9	2,632.3
Reduction [%]	**70.5**%	**57.9**%

step response compared to the Passive MPC. Even when the learning transient is included in the KPIs, the overall performance indices of the Active MPC remain substantially lower, demonstrating its superior control efficiency and adaptability. The first step response is depicted in Fig. 7.3.

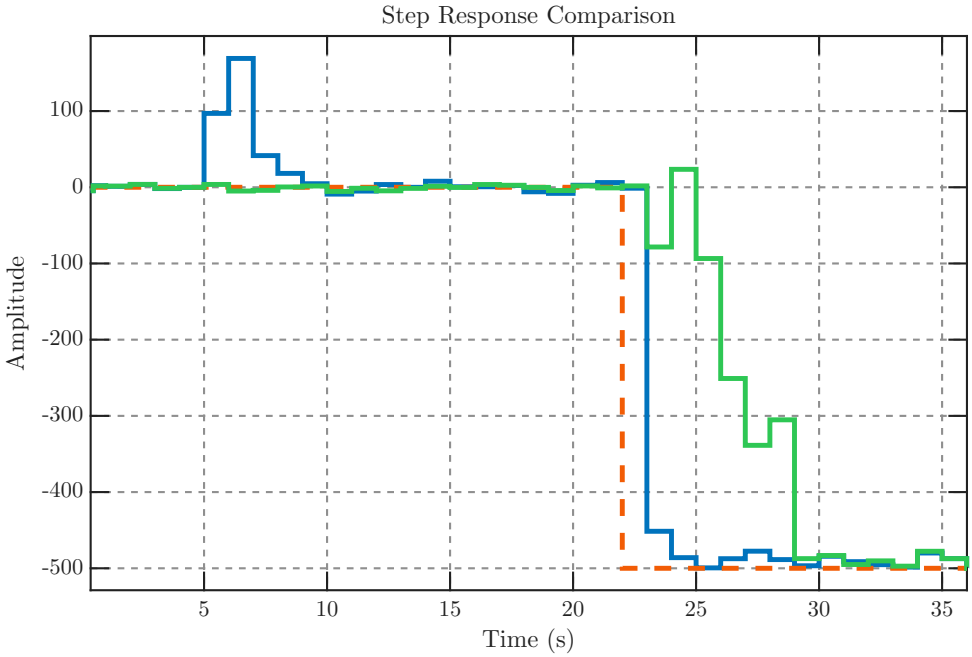


Figure 7.3: Input and output trajectory of the passive (green) and active (blue) closed-loop systems compared to the reference signal (dotted-red).

Figure 7.4 presents the FPS before and after the learning phase. The system can correctly reduce the distance between the real data $\theta^\star$ and the central estimate θ_c^k. With a lower FPS volume, the aggregator has better knowledge of the underlying system and, therefore, better performance in subsequent references. As a remark, since the passive aggregator does not probe the dynamic, the passive FPS remains equal to the initial FPS Θ_0, meaning that at the time instant in which the step start, i.e. $k = 22$s, the passive aggregator propagates the central estimate model θ_c^k, equal to the initial central estimate model θ_c^0, which does not represent the correct system dynamics. However, it is worth noting that, even when the reference signal exceeds the constraints, both active and passive MPC

remain within specified limits due to robustification of the scenario.

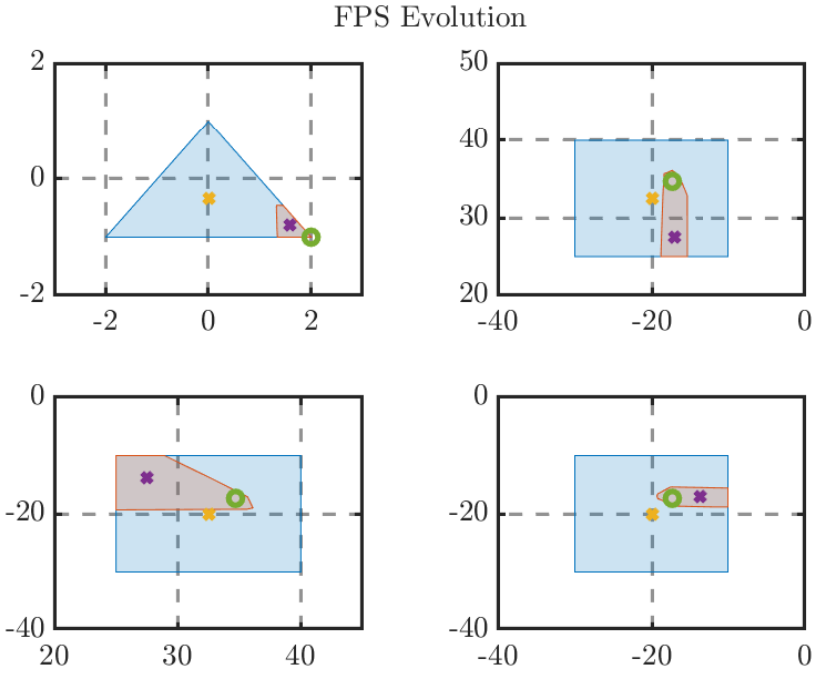


Figure 7.4: Initial FPS at $k = 1$ (blue area), FPS at $k = 25$ (orange area), central estimate θ_c^1 (yellow x) and θ_c^{24} (purple x); real parameter $\theta^\star$ (green circle).

Here, we present a computationally tractable approximation of the dual control problem based on the Set Membership Estimation (SME) framework. The proposed MPC dual algorithm minimises the volume of the feasible parameter set as a proxy for model-mismatch-induced uncertainty, thereby balancing exploration and performance within the control loop.

The algorithm is applied to an aggregated set of thermoelectric loads participating in demand response programs. In this setting, the DSO issues upward or downward reserve requests, which the aggregator meets by adjusting the temperature reference setpoints of individual devices, thereby modulating their power consumption. During periods with no reserve requests, the controller actively probes the system's switching dynamics to improve model accuracy, learning distinct models for each reserve direction. Simulation results show that the proposed dual MPC achieves better adaptability and performance than a passive MPC counterpart, particularly during reserve activations and model transitions.

Chapter 7 Bibliography

[AE07] M. H. Albadi and E. F. El-Saadany. "Demand Response in Electricity Markets: An Overview". In: *2007 IEEE Power Engineering Society General Meeting*. 2007, pp. 1–5. DOI: 10.1109/PES.2007.385728.

[Alz+23] Ahmad Alzahrani et al. "Demand Response for Optimal Power Usage Scheduling Considering Time and Power Flexibility of Load in Smart Grid". In: *IEEE Access* 11 (2023), pp. 33640–33651. DOI: 10.1109/ACCESS.2023.3263849.

[BK22] Steven L Brunton and J Nathan Kutz. *Data-driven science and engineering: Machine learning, dynamical systems, and control*. Cambridge University Press, 2022.

[CDR22] Andres Cordoba-Pacheco, Cesar Diaz-Londono, and Fredy Ruiz. "Data-Driven Aggregation Control for Thermoelectric Loads in Demand Response". In: *IFAC-PapersOnLine* 55.40 (2022). 1st IFAC Workshop on Control of Complex Systems COSY 2022, pp. 205–210. ISSN: 2405-8963. DOI: https://doi.org/10.1016/j.ifacol.2023.01.073.

[Dia+20] Cesar Diaz-Londono et al. "Experimental modeling and aggregation strategy for thermoelectric refrigeration units as flexible loads". In: *Applied Energy* 272 (2020), p. 115065. ISSN: 0306-2619. DOI: https://doi.org/10.1016/j.apenergy.2020.115065.

[MCC24] Zhihao Ma, Shuang Cui, and Jianli Chen. "Demand response through ventilation and latent load adjustment for commercial buildings in humid climate zones". In: *Applied Energy* 373 (2024), p. 123940. ISSN: 0306-2619. DOI: https://doi.org/10.1016/j.apenergy.2024.123940.

[Mez+25] E. Meza et al. "Adaptive Scenario-Based Predictive Control: A Set Membership Learning Approach". In: *2025 IFAC 7th Conference on Intelligent Control and Automation Sciences (ICONS)*. Padova, Italy, 2025.

[Oh+25] Hyunjun Oh et al. "Techno-economic performance of reservoir thermal energy storage for data center cooling system". In: *Applied Energy* 391 (2025), p. 125858. ISSN: 0306-2619. DOI: https://doi.org/10.1016/j.apenergy.2025.125858.

[PM19] J. Ponoćko and J.V. Milanović. "The Effect of Load-follow-generation Motivated DSM Programme on Losses and Loadability of a Distribution Network with Renewable Generation". In: *2019 IEEE PES GTD Grand International Conference and Exposition Asia (GTD Asia)*. 2019, pp. 165–170. DOI: 10.1109/GTDAsia.2019.8715988.

[SCR25] Qian Shi, Andres Cordoba-Pacheco, and Fredy Ruiz. "Efficient Learning Control through Zonotopic Set Membership Estimation and Scenario MPC". In: *2025 European Control Conference (ECC)*. 2025, pp. 1–6. DOI: 10.23919/ECC65951.2025.11186893.

8 | Conclusions

8.1. Scientific contributions

This dissertation has addressed several fundamental challenges in the modelling, control, and coordination of modern energy systems, with a particular focus on Microgrids (MGs), Virtual Power Plants (VPPs), and aggregations of Flexible Loads (FLs).

Over three years of research, the work has advanced the state of the art in uncertainty-aware control and optimisation for multi-agent energy systems by combining rigorous theoretical developments with algorithmic design, extensive simulation studies, and real–world oriented implementations. The results contribute to different layers of the energy management hierarchy, from local microgrid operation to distributed cooperative services among interconnected microgrids to adaptive control of aggregated flexible devices, adopting a principled, data-driven approach to uncertainty modelling.

The first significant contribution of this thesis concerns the economic dispatch and energy management of single microgrids in the presence of uncertain renewable production and demand. Traditional approaches to Energy Management System (EMS), such as scenario-tree methods and approximate dynamic programming, can become computationally intensive or overly conservative when attempting to capture the high-frequency fluctuations typical of real-world renewable sources. In contrast, this thesis proposes a hierarchical EMS architecture grounded in the Scenario Approach that treats uncertainties in an entirely data-driven, distribution-free manner. The two-level structure comprises a supervisory economic EMS responsible for long-horizon cost optimisation and a scenario-based EMS designed to mitigate high-frequency deviations. This architecture successfully reduces unexpected grid interventions while retaining computational efficiency. Extensive Monte Carlo simulations with real-world datasets demonstrate the robustness and scalability of the framework. These results, published in a leading international conference, advance the field by illustrating that high-frequency uncertainties can be integrated efficiently without resorting to overly conservative or computationally intractable methodologies.

Building upon this foundation, the second central pillar of the thesis advances the understanding of uncertainty handling in multi-agent distributed optimisation. Many emerging energy systems, particularly VPPs , microgrid clusters, and energy communities, require autonomous agents to coordinate to satisfy global coupling constraints, such as aggregated power limits or joint service obligations. Most existing distributed optimisation approaches scale poorly when uncertainty is incorporated into the global constraints, typically requiring a number of scenarios proportional to the network size. The theoretical work developed in this dissertation shows that, by exploiting the geometric structure of the global uncertain constraint, it is possible to derive a bound on the support dimension that is independent of the number of agents. This provides the first scalable, scenario-based method for robust multi-agent coordination, which has been formalised and published in a central control theory venue.

Leveraging this theoretical insight, the thesis develops novel distributed and decentralised VPP architectures capable of coordinating heterogeneous microgrids to provide aggregated services, such as ramp-up capability, to the transmission system operator. A dual-Lagrangian tracking algorithm is proposed to enforce the global coupling constraint while allowing each microgrid to optimise its local schedule. Crucially, the algorithm incorporates the developed scenario-theory extension for multi-agent systems with uncertain global coupling constraints, enabling robust coordination without compromising scalability. A comprehensive simulation framework was constructed to generate randomised VPPs with heterogeneous assets and uncertainty profiles, enabling extensive validation of the algorithmic properties. This combination of theoretical analysis, algorithm design, simulation, and implementation represents one of the thesis's significant achievements.

The final major research direction explored in this dissertation concerns adaptive control under parametric uncertainty for systems represented by Autoregressive with Exogenous input (ARX) models. Classical Model Predictive Control (MPC) methods either rely on nominal models or adapt the system's model passively, often lacking explicit guarantees for the exploration/exploitation tradeoff, robustness, and learning performance. To address these limitations, the thesis presents a Dual Adaptive MPC framework grounded in set-membership estimation theory. The passive version of the algorithm uses the Feasible Parameter Set (FPS) to represent all models consistent with the observed data and prior knowledge. Robust control is achieved by sampling models from the FPS using the Scenario Approach, thereby avoiding set propagation while preserving computational tractability. The active version extends this approach by incorporating the FPS volume into the MPC cost function, introducing exploratory behaviour that improves model learning while maintaining safety. Formal guarantees are provided for probabilistic constraint

satisfaction, robustness, and convergence to certainty equivalence as the FPS contracts. The algorithm is validated for the aggregated control of Thermo-Electric Refrigerators, a hybrid, partially unknown system of high relevance for demand-response applications. Here, the controller learns the unknown aggregated dynamics online while reliably delivering upward and downward reserve services, even with minimal prior knowledge.

8.2. Future developments

Several research directions emerge from the present work.

More work is needed on the study of microgrid architectures. In this context, decentralised, distributed, and uncertainty-aware control structures remain of central interest to the energy research community. Future developments should not only address the theoretical aspects of these architectures but also focus on their practical deployment at an industrial scale. In particular, communication delays, heterogeneous computational capabilities across nodes, and cybersecurity threats introduce additional layers of complexity and uncertainty that must be rigorously accounted for in both design and implementation.

A first extension to the work is the study of more advanced architectures, such as non-centralised VPP. In such architectures, each microgrid can operate in island mode or form coalitions with neighbouring microgrids to improve the formation's resilience. Future development is to propose online reconfiguration strategies for microgrids under uncertain production and consumption, based on robustifying existing heuristics and algorithms, such as simulated annealing, to address the problem's complexity.

Machine-learning techniques and big-data analytics need substantial improvements to enhance the predictive accuracy, generalisation capabilities, and robustness of forecasting models used in energy management systems. These advances should be complemented by scalable optimisation and control algorithms that can leverage high-volume, high-dimensional datasets without compromising real-time performance. Among machine-learning techniques, future work includes extending the Dual MPC algorithm to handle nonlinear and time-varying systems, where the system can be expressed as a linear system plus an unknown but bounded nonlinear component. Furthermore, the Dual MPC algorithm must be evaluated on real-world benchmarks to assess its capabilities.

The integration of game theory into distributed energy systems can enable the development of data-driven policies that steer the network toward optimal operating points, regardless of whether individual agents pursue local or conflicting objectives.

Finally, all the proposed approaches can be implemented and validated in multi-energy

networks, which serve as a natural testbed due to their inherent coupling among heterogeneous energy carriers. Such systems introduce additional uncertainties, modelling inaccuracies, and operational constraints that go beyond those encountered in single-energy microgrids. Validating distributed and uncertainty-aware control frameworks in these settings would provide strong evidence of their scalability, robustness, and relevance for future large-scale energy infrastructures.

8.3. Final considerations

Overall, this dissertation presents a unified and rigorous framework for uncertainty-aware modelling, control, and optimisation in next-generation energy systems. Through contributions spanning single-microgrid operation, distributed VPP coordination, adaptive control of flexible loads, and microgrid reconfiguration, the work consistently addresses two central challenges: the need for strong theoretical guarantees under uncertainty and the need for scalable, implementable algorithms suitable for deployment in real-world energy infrastructures. The research has been validated through peer-reviewed publications, simulation studies, and industry-embedded development, offering both scientific depth and practical relevance. Collectively, these results push forward the boundary of resilient, data-driven, and autonomous energy management and open promising directions for the design of future intelligent power systems.

List of Figures

1.1 Total global GHGs emission by sector. Data source [Our26] 2
1.2 RES and DER cost decrease and increase adoption [Shu+22]. 3
1.3 Smart Grid architecture [Wol12]. 5

2.1 Optiman value $x^\star(\epsilon)$ for increasing level of ϵ [GC13]. 16

3.1 Feasible Parameter Set evolution. In dotted red the intial FPS $\mathcal{K}_0$; in blue the current FPS Θ_{k-1}; in green the strip induced by the $k-th$ measurement $\mathcal{S}_k$; in yellow, the true parameter $\theta^\star$. 37
3.2 Strips $\mathcal{S}_k$ for different values of u_k. 44
3.3 Strips $\mathcal{S}_k$ for increasing values of $I(k)$. 51
3.4 Initial FPS $\mathcal{K}$, region in the (θ_1, θ_2) parameter space, corresponding to the inequality (3.75). Real $\theta^\star$(purple cross) and central estimate θ_c^k (yellow dot). First realised strip $\mathcal{S}_1$ 57
3.5 Output y_k and input u_k closed-loop evolution (blue line), reference value $y_{ref,k}$(orange line), output constraints (red-dotted line). 58
3.6 Parameter θ_1^k, θ_2^k trajectory (blue line). Real value $\theta_1^\star, \theta_2^s tar$ (red-dotted line). 59

4.1 Overview of the main electricity markets in Europe and their temporal sequence. [Gmb24]. 62
4.2 Hierarchical control architecture. [Sah+23]. 66
4.3 Hierarchical, Decentralised, and Distributed EMS control architecture. [PS17]. 70

5.1 Hierarchical EMS scheme 79
5.2 Past net-demand data (orange), NARX net-demand prediction (pink-line) with 3σ interval (black-dotted), true net-demand (blue-line) 80
5.3 Upper plot, cost of generation c_{gen} as a function of generated power P_{gen}. Lower plot, Pricing Ranges during the day: low cost (blue), medium cost (orange), high cost (red). 88
5.4 Left plot: load -, right plot PV production 88
5.5 Net-demand with $3 \cdot \sigma$ intervals. 89

5.6 Daily Power profiles for the micro-grid. 89
5.7 Histograms of power flows exchanged with the grid, $||\delta P_{grid}(t)||_2^2$, for EMSs with (orange) and without (blue) the low-level risk-constrained MPC. . . . 90

6.1 VPP topologies for different values of $m, \mathcal{E}$. 104
6.2 Aggregated power production $P_n(t, \delta)$: mean value (blue line), 1-ϵ level (light blue area), violating unseen scenarios (dotted red line). Minimum aggregated power required $P_n^{ref}(t)$ (dashed black line) [DR25] 104
6.3 Violation probability $V_N(x^*, \delta)$, in red $N = 362$, in blue $N = 100$, yellow and purple line: average value of the violation [DRS24] 106
6.4 Convergence of the normalised absolute error 107

7.1 TER control architecture 113
7.2 Top: Output trajectories of the passive (green) and active (blue) closed-loop systems, shown alongside the reference signal (orange) and the constraint (red, dotted). Bottom: Input trajectories of the passive (orange) and active (blue) closed-loop systems. 117
7.3 Input and output trajectory of the passive (green) and active (blue) closed-loop systems compared to the reference signal (dotted-red). 118
7.4 Initial FPS at $k = 1$ (blue area), FPS at $k = 25$ (orange area), central estimate θ_c^1 (yellow x) and θ_c^{24} (purple x); real parameter $\theta^\star$ (green circle). . 119

List of Tables

4.1 Comparison of EMS architectures. 69

5.1 Violation occurrence 91

7.1 Performance metrics for the Active and Passive MPC comparison 118

Acronyms

AC Adaptive Control. 31

AFOLU Agriculture, Forestry, and Other Land Use. 1

aFRR Automatic Frequency Restoration Reserve. 64, 65

AI Artificial Intelligence. 4

AMI Advanced Metering Infrastructure. 4

ARX Autoregressive with Exogenous input. 33, 38–40, 53, 54, 56, 114, 115, 124

ASM Ancillary Services Market. 63, 64

BESS Battery Energy Storage System. 72, 79, 81–83, 85–87, 95, 96

BESSs Battery Energy Storage Systems. i, iii, 1, 6, 11, 95, 103

BM Balancing Market. 63, 64

CE Certainty Equivalence. 49, 51

DAM Day-Ahead Market. 62, 63

DERs Distributed Energy Resources. i, iii, 1, 4, 7, 61, 67, 69

DoS Denial of Service. 6

DR Demand Response. i, iii, 4, 6, 33, 109–112

DSO Distribution System Operator. 63, 64, 73, 109, 110, 112, 116, 117, 119

ECs Energy Communities. 4

EMS Energy Management System. i, iii, 12, 61, 63, 66–70, 73, 77–79, 82–84, 86, 87, 90, 123

EVs Electric Vehicles. 6

FCR Frequency Containment Reserve. 64, 65, 111

FHOCP Finite Horizon Optimal Control Problem. 84

FLs Flexible Loads. i, iii, 109, 123

FPS Feasible Parameter Set. 11, 33–46, 48–51, 54–58, 110, 116, 118, 124, 125, 127

FRR Frequency Restoration Reserve. 111

GHGs GreenHouse Gases. 1–3, 127

HPs Heat Pumps. 6

HVAC Heating, Ventilation, and Air Conditioning. 6, 109

IAE Integral Absolute Error. 117

IDM Intraday Market. 62, 63

IoE Internet of Energy. 4

IPCC International Panel for Climate Change. 1, 2

ISE Integral Squared Error. 117

KPIs Key Performance Indicators. 117, 118

mFRR Manual Frequency Restoration Reserve. 64

MGs Microgrids. i, iii, 66, 123

MIQP Mixed Integer Quadratic Program. 81, 82, 85, 87

MPC Model Predictive Control. i, iii, 11, 12, 33, 39, 41, 42, 45, 48, 49, 58, 77–79, 82, 83, 86, 87, 90, 110, 113, 114, 116–119, 124, 125, 128, 129

MSE Mean Square Error. 80

NARX Neural Autoregressive with eXogenous inputs. 79, 80

PE Persistent Excitation. 48

PV Photovoltaic. 1, 6, 78, 87, 88, 95, 97, 100, 127

RESs Renewable Energy Sources. i, iii, 1, 2, 4, 6, 11

RR Replacement Reserves. 65, 111

SCADA Supervisory Control and Data Acquisition. 4

SG Smart Grid. 3, 4, 6, 7, 67

SME Set Membership Estimation. i, iii, 11, 32, 33, 35, 42, 45, 49, 58, 119

SOC State Of Charge. 72, 79, 87, 95

TCLs Thermostatically Controlled Loads. 109–111

TER Thermo-Electric Refrigerator. 111–114, 116

TERs Thermo-Electric Refrigerators. 6, 12, 33, 111, 112, 114

TSO Transmission System Operator. 64, 93, 96–98, 104, 106

UBB Unknown-But-Bounded. 33, 34, 40, 54, 56

VPP Virtual Power Plant. 12, 28, 66, 73, 93–100, 102–107, 124–126

VPPs Virtual Power Plants. i, iii, 4, 103, 123, 124

Acknowledgements

I would like to express my deepest gratitude to Professor Fredy Ruiz for his guidance and support throughout this work. I could never have reached this point without his understanding, encouragement, and constant presence, always offering help without ever pushing. I am also sincerely grateful to Professor Riccardo Scattolini, Professor Carlos Ocampo and Professor Alessio La Bella for their invaluable support and for the inspiring example they set.

I am profoundly thankful to my parents, Edoardo and Amelia, my sister Chiara, nonna Umbertina and nonno Mario, and my second family — Emanuela, Luca, Mattia, Raimondo, and Toby — whose love, support, and encouragement have been the foundation of everything I have achieved. Their patience and unconditional belief in me gave me the strength to pursue this path and grow as a person.

To the Chocolaterie, who stood by me during this PhD, offering encouragement, laughter, and perspective in both challenging and joyful moments. Special thanks — particularly to Lorenzo, Andres, Dave, and Laura — with whom I shared countless memorable moments that enriched this experience.

To my donnette — Chiara, Ema, Ame, Andre, Fabio, Marco, and Cami — thank you for being such an important part of my life. To Mario, Bubiz, Mitch, Rasto, Letizia, and all of my certezze, your support means the world to me. Thanks to Andre, Sara, and Marty, with whom I have shared unforgettable nights. To Francesca, thank you for helping me stay sane over the past year.

If you are reading this, I want to thank you as well — because in one way or another, you have contributed to this achievement, and I am deeply grateful for your friendship.

Finally, none of this would have been possible without my love, Danila, who has been my beacon through every storm and difficult moment, offering unconditional support, patience, and affection even at my lowest. Her love, encouragement, and belief in me have been a constant source of strength and inspiration, and I can never fully express my gratitude for all that she means to me.