



UNIVERSITÀ
DEGLI STUDI
FIRENZE

UNIVERSITÀ DEGLI STUDI DI FIRENZE
DIPARTIMENTO DI INGEGNERIA DELL'INFORMAZIONE (DINFO)
CORSO DI DOTTORATO IN INGEGNERIA DELL'INFORMAZIONE
CURRICULUM: EEE

SHAPING A MORE RESILIENT GRID: MULTI-LAYER DISTRIBUTED INTELLIGENCE FOR THE GREEN TRANSITION

Candidate

Lorenzo Becchi

Supervisors

Prof. Francesco Grasso

Prof. Gabriele Maria Lozito

PhD Coordinator

Prof. Stefano Ricci

Università degli Studi di Firenze, Dipartimento di Ingegneria
dell'Informazione (DINFO).

Thesis submitted in partial fulfillment of the requirements for the
educational and research PhD Program in Information Engineering.
Copyright © 2026 by Lorenzo Becchi.

*To all the people who care about me, to those who do not
and to those I will never meet.*

Acknowledgments

I wish to express my sincere gratitude to my supervisors, Prof. Francesco Grasso and Prof. Gabriele Maria Lozito, whose guidance have been instrumental throughout the course of this research. I am equally indebted to my colleagues at the LEM (Laboratorio di Elettrotecnica e Misure), whose professionalism and support consistently strengthened the work carried out. My particular thanks go to Dr. Matteo Intravaia, Dr. Marco Bindi and Prof. Antonio Luchetta, whose collaboration on core aspects of this study was essential and deeply appreciated.

I would also like to extend my gratitude to Prof. Igor Aizhemberg for his generous support and hospitality during my stay at Manhattan College in New York, as well as to Prof. Yajuan Guan, Dr. Wenfa Kang, Dr. Mahshid Javidsharifi and the CROM group at Aalborg University for the valuable assistance and opportunity they provided during my time there.

The author acknowledges the use of AI-assisted language tools (including ChatGPT) to support English language refinement and improve clarity and readability. All scientific content, interpretations, and conclusions remain the sole responsibility of the author.

1

¹This project was funded by the European Union, Next Generation EU, Mission 4, Component 1, CUP J53D23000760006. The authors also acknowledge the financial support of the National Research Center in High-Performance Computing, Big Data and Quantum Computing (<https://www.supercomputing-icsc.it/en/icsc-home/>, accessed on 1 October 2024) foreseen within Mission 4 (Education and Research) of the “National Recovery and Resilience Plan” (NRRP) that is part of the Next Generation EU (NGEU) program (<https://www.italiadomani.gov.it/en/>, accessed on 1 October 2024).

Contents

Contents	v
1 Introduction	1
1.1 The objective	3
1.2 Contributions	3
1.3 Thesis Organization	4
2 Literature review	5
2.1 Scope and Structure	6
2.2 The Italian Electricity Market	7
2.3 National Electrical Infrastructure	10
2.4 Renewable Energy Communities (RECs)	13
2.5 Energy Management Systems (EMS)	17
2.6 Resilience in Distribution Networks	25
2.7 Discussion and Research Gaps	31
3 Proposed Architecture	33
4 Forecasting of Consumption and Generation Profiles	41
4.1 Introduction to Electrical Load Forecasting	42
4.2 The Role of Forecasting in Energy Management Systems . . .	44
4.3 Characteristics of Electrical Load and PV Generation Time Series	45
4.4 Low-Complexity Forecasting Models for EMS Operation . . .	46
4.5 Complex Valued Neural Networks in Load Forecasting	52
4.6 Load Profiles Stochastic Generation	64
4.7 Conclusion and Final Remarks	73

5	Primary-Level Energy Management (HEMS)	75
5.1	Introduction	76
5.2	Notation and Preliminary Hypotheses	76
5.3	Linear Programming Benchmark Definition	82
5.4	Rule-Based Algorithm Definition	85
5.5	LP vs. Rule-based	91
5.6	Price Overlapping Relaxation	98
5.7	Controllable Loads	107
5.8	Handling Forecasting Uncertainty	109
5.9	Conclusion and Final Remarks	123
6	Secondary-Level Energy Management (CEMH)	125
6.1	Introduction	126
6.2	Choice of the Optimization Approach	126
6.3	Notation and Preliminary Consideration	128
6.4	Definition of the MILP Benchmark	131
6.5	Definition of the Proposed Approach	140
6.6	Load Shedding and Resilience	163
6.7	Conclusion and Final Remarks	171
7	Scaling Up Control: Transmission-Level EMS	173
7.1	Introduction	174
7.2	Power Flow Analysis	175
7.3	Ethiopian Transmission Network (E-TN)	183
7.4	Operational Management System (OMS) for E-TN	209
7.5	Preliminary Results	220
7.6	Conclusion & Future Development	224
8	Conclusion	227
9	Future Developments	231
	Bibliography	237
A	Datasets and Information Sources	247
B	Publications	249

Chapter 1

Introduction

The urgency of climate change mitigation has placed the energy sector at the center of global decarbonisation strategies [1]. International agreements and national policies converge on the need to drastically reduce greenhouse gas emissions, pushing power systems toward rapid and large-scale integration of renewable energy sources [2]. This transition is not merely a technological upgrade, but a structural transformation of the electricity system, affecting generation, transmission, distribution and consumption paradigms simultaneously. The traditional model, based on centralised and dispatchable generation that supplies passive end users, is progressively being replaced by a highly distributed ecosystem characterised by intermittent production, active consumers, and bidirectional power flows.

This evolution has accelerated the development of Smart Grids, conceived as electricity networks enhanced by advanced sensing, communication and control capabilities [3]. Smart grids aim to ensure reliability, efficiency and sustainability in a context where uncertainty and variability are intrinsic features of system operation. However, the deployment of intelligence across the grid introduces new layers of complexity. Real-time observability, fast decision-making and coordination among heterogeneous actors become essential requirements, particularly at the distribution level where renewable penetration is most pronounced [4].

In Europe climate and energy policies have led to a sharp increase in non-programmable renewable generation over the past decade. In Italy in particular photovoltaic, wind power plants and renewables now represent a significant share of the electricity mix, covering more than 40% of the de-

mand [5, 6], contributing substantially to emission reduction targets while simultaneously challenging conventional network operation. Issues such as voltage regulation in medium and low voltage feeders, congestion management, protection coordination and reverse power flows have become increasingly frequent.

In response to these challenges, recent European legislation, along with a complete package of regulations [7, 8], introduced Renewable Energy Communities (RECs) as a regulatory and organisational framework to promote local energy sharing and active participation of consumers and prosumers. By fostering local balance of production and consumption, RECs shift part of operational flexibility from the transmission system to the distribution network, where control actions can be faster and more targeted.

This transition, however, cannot be relying solely on large-scale infrastructure upgrades or centralised control solutions [9]. Achieving meaningful decarbonisation requires the active participation of a wide set of different typologies of users, including individual citizens and small prosumers [10]. As power system operation becomes increasingly decentralised and algorithm-driven, the interaction between humans and complex energy management tools emerges as a non-negligible component of system effectiveness. For such reason the proposed technologies must not only be technically effective, but also accessible, transparent and usable by non-expert users, ensuring that participation in energy communities does not remain limited to a restricted group of highly skilled operators.

Within this framework, Energy Management Systems (EMSs) emerge as a key enabling technology [11]. An EMS collects measurements, processes forecasts and computes control actions for distributed resources such as battery energy storage systems and flexible loads. When applied to energy communities and distribution networks, traditional EMS architectures reveal important limitations. The heterogeneity of user behaviour, limited observability, forecast uncertainty and the coexistence of individual and collective objectives hinder the effectiveness of centralised and rigid control strategies. These constraints motivate the development of EMS architectures that are scalable, adaptive, and capable of coordinating a large number of distributed actors while preserving autonomy at the local level.

Beyond efficiency and economic optimisation, resilience has become a central concern in modern power systems. Climate change itself increases the frequency and severity of extreme weather events [12], while ageing infras-

structure and the growing interdependence between energy and digital systems expose distribution networks to new types of disturbances.

In this context, resilience requires fast detection of abnormal conditions, adaptive reconfiguration capabilities and the ability to exploit distributed energy resources to maintain supply continuity under degraded operating conditions [13]. Power flow analysis in its classical and real-time formulations plays a fundamental role in this process. It provides a consistent framework to assess system states, validate control strategies and ensure that decisions taken at different control layers remain compatible with network constraints.

1.1 The objective

Against this background, the objective of this dissertation is to develop an intelligent, distributed, reliable and cost-effective monitoring and control system capable of supporting the green transition and reinforcing resilience in distribution networks. The proposed system integrates advanced forecasting techniques, low-complexity neural models, distributed state estimation and multi-level EMS strategies. The architecture has been designed to operate with realistic measurement infrastructures and computational constraints, making it suitable for practical deployment in microgrids, RECs and medium-voltage distribution networks.

The following literature review situates the proposed approach within the current state of the art, highlighting existing solutions and open challenges that motivate the contributions of this work.

1.2 Contributions

The main contributions achieved throughout this research are:

- the design of a distributed monitoring architecture combining edge intelligence and central coordination;
- the development of forecasting models crafted for small-scale users and energy communities, with low computational cost and high robustness;
- a multi-level EMS framework that coordinates heterogeneous prosumers while respecting network constraints;
- algorithms for adaptive network reconfiguration to enhance resilience;

- an integrated simulation and decision-support platform, including a conversational AI tool, to evaluate REC operation under realistic regulatory conditions.

1.3 Thesis Organization

The remainder of this dissertation is organised as follows.

Chapter 2 reviews the background on modern distribution networks, distributed generation, renewable energy communities, Energy Management System (EMS) architectures and Reliability.

Chapter 3 presents the system architecture and the methodological framework adopted throughout this work.

Chapter 4 discusses the forecasting models developed during this research.

Chapter 5 introduces the energy management methodology designed for the single-user level.

Chapter 6 provides an in-depth analysis of the second-layer control system, the one that stands at the community level

Chapter 7 extends the energy management strategy to the transmission interface by implementing an Optimal Power Flow (OPF) formulation.

Chapter 8 discusses the broader implications, limitations, and potential future developments.

Finally, Chapter 9 concludes the dissertation.

Chapter 2

Literature review

This literature review analyzes the evolution of distribution networks in the presence of high penetration of renewable energy sources, distributed generation, and active consumers. It examines the interaction between electricity market mechanisms, network infrastructure, regulatory frameworks, and control architectures, with a particular focus on Renewable Energy Communities. The review highlights the central role of Energy Management Systems as coordination layers enabling forecasting, optimization, and multi-level control. Current approaches, limitations, and open challenges related to scalability, uncertainty, and resilience are discussed, providing the conceptual foundation for the methodological framework developed in the subsequent chapters.

2.1 Scope and Structure

This literature review focuses on the operational and control implications of the rapid changes and evolution that the electric power systems are experiencing at the distribution level.

In particular, the chapter surveys the body of work addressing how market mechanisms, network infrastructure, regulatory frameworks, and control architectures interact in modern distribution systems characterized by distributed energy resources (DERs) and active consumers. Rather than treating these aspects in isolation, the literature is reviewed with the explicit goal of highlighting their interdependencies and the resulting need for advanced coordination and energy management solutions.

The review begins by examining the Italian electricity market as a representative case study, emphasizing how wholesale and balancing market price signals reflect network conditions and influence the behavior of producers and consumers. It then analyzes the physical structure of the national electrical infrastructure and the technical challenges introduced by widespread distributed generation. Building on this foundation, the chapter reviews the regulatory and operational framework of Renewable Energy Communities (RECs), with particular attention to incentives, constraints, and coordination requirements.

Subsequently, the role of Energy Management Systems (EMS) is discussed as a unifying control layer capable of coordinating distributed resources across multiple scales, from individual users to communities and system-level operators. Finally, the chapter addresses the concept of resilience in distribution networks highlighting the role of DER coordination, microgrids and EMS-based control in enhancing the system's ability to withstand and recover from extreme events.

Overall, this literature review aims to establish the state of the art and identify open challenges that motivate the development of intelligent, distributed and scalable EMS architectures for modern distribution networks and renewable energy communities.

2.2 The Italian Electricity Market

2.2.1 Structure

Italy's electricity sector has undergone liberalization over the past two decades, evolving into a competitive market with multiple actors and trading sessions. The wholesale market, managed by the Italian Power Exchange (IPEX), is organized into a Day-Ahead Market (MGP – Mercato del Giorno Prima) as the primary session where the majority of energy transactions are scheduled, along with Intra-Day Market (MI) and forward and balancing markets [14].

In the MGP, generators and consumers (or their retailers) submit bids for each hour of the next day. A market-clearing algorithm matches supply and demand for each hour and determines the dispatched volumes and prices. Italy is divided into several bidding zones (geographical regions such as North, Center-North, Center-South, South, Sardinia, Sicily, etc.) to account for transmission constraints. When transmission limits between zones bind, the market splits and produces different zonal prices in the constrained areas, reflecting localized supply-demand conditions.

A national reference price, the Prezzo Unico Nazionale (PUN), is then computed as the average of the zonal prices weighted by volumes. In practice, generators are paid the price in their zone, while many consumers pay the PUN (for example, in regulated contracts or certain wholesale purchasing arrangements) [15]. The PUN serves as a single benchmark price that many retail contracts reference, although underlying regional price differences can be significant.

In addition to the MGP and MI sessions, Italy also operates a real-time Balancing Market (Mercato per il Servizio di Dispacciamento, MSD) in which the TSO (Terna) procures last-minute adjustments and ancillary services to maintain grid stability. Participation in balancing is mandatory for certain large units (e.g. generators >10 MW) and open to others on a voluntary basis. This multi-tier market structure provides critical price signals at different timescales for market participants to adjust their production or consumption.

2.2.2 Price Signals

The hourly clearing prices emerging from the day-ahead market represent the marginal cost of electricity for each hour and act as signals of scarcity or surplus guiding generation dispatch. High prices signal scarcity (encouraging

more generation or reduced consumption), whereas low or negative prices signal surplus (encouraging flexible loads or storage uptake) [16]. Italy's market has historically exhibited a pronounced spread between day and night prices, and between different zones, reflecting factors such as solar generation peaks, wind availability, and transmission limits. These zonal price differences are a feature, not a flaw, of the market design, as they signal where the grid is constrained and electricity is more costly to deliver [17].

The increasing penetration of renewables has led to greater price volatility: for instance, strong midday solar output even led to occasional zero or negative prices in those hours, whereas evening hours, when solar is offline but demand is still high, see price spikes. This day-night price spread has become more pronounced with more solar in the system. Such encreasing of the intra-day price volatility present operational implications for the grid [18].

On the positive side, price volatility incentivizes flexibility in the system, like schedule of energy-intensive processes of the industrial users or the adoption of energy storage and demand response resources that can arbitrage between low and high price periods. By doing so, these resources not only earn money but also help the system by absorbing excess generation when it's abundant and providing extra supply or reduced load when the system is tight.

On the other hand, rapid price fluctuations can pose challenges for generators and retailers in terms of risk management and can create uncertainty for distributed energy resources that lack direct market access. This loss in predictability has driven interest in improved risk management and forecasting tools in Italy's power sector [19].

2.2.3 Operational Implications

A key challenge for DER like rooftop solar, small wind or electric vehicles is that many of these resources do not have direct access to the wholesale market price signals. These prosumers, those who both consume and produce electricity, typically do not participate directly in the wholesale market so they result partially blinded regarding the real-time wholesale price signal. Traditional retail tariffs in Italy have included a mix of fixed charges and time-of-use energy prices, the latter giving some incentive for consumers to shift consumption to off-peak periods.

Another operational concern is that uncoordinated demand response to price signals can cause new problems, such as rebound peaks, phenomenon

in which a synchronized price-responsive response shiwinzer2024tariffsft or concentrate shift peaks of excessive or scarcity of productions to others hours instead of truly alleviate the stress on the whole system. This underscores that price signals need to be paired with intelligent coordination (through smart appliances, staggered response, or grid-aware control algorithms) to avoid simply moving a peak from one hour to another [20].

2.2.4 Towards the Flexibility

As the Italian system is becoming more and more dynamic, flexibility is both encouraged and required. Unfortunately, the majority of Italy's smaller consumers (households and small businesses) do not have access to the energy markets. The purchase of electricity pass through retail contracts, which historically were often regulated or fixed-price offering.

This means that in most end-user electricity bills the real-time price signal is only partially reflected, or not at all. To overcome this issue there has been movement in Italy toward introducing more dynamic pricing at the retail level. In compliance with EU Electricity Market Design regulations [14], retail suppliers in Italy are now obliged to offer at least one dynamic tariff option, where the price is linked to wholesale market outcomes, to consumers with suitable smart meters. These meters record hourly consumption, which is a prerequisite for dynamic pricing and participation in flexibility programs.

Italian regulators are working to better align prosumer incentives with market signals. One approach has been to enable aggregators and new market mechanisms so that even small resources can respond to system conditions. Rather than each prosumer acting in isolation under fixed incentives, aggregators can pool many small resources and respond to price signals on their behalf, effectively bringing them into the market.

Several pilot programs and regulatory innovations has been implemented in recent years. A cornerstone initiative is Terna's introduction of Virtual Aggregated Units in the balancing market. Starting around 2017–2018, Terna opened up the Ancillary Services Market (MSD) to aggregated resources under schemes called UVAC, UVAP, and later UVAM. UVAM (Unità Virtuali Aggregated Miste) stands for "mixed" virtual aggregated units, meaning a combination of demand reduction, generation increase and storage can be aggregated to offer a single block of flexibility [21].

The success of these pilots has been important. By opening the door to independent aggregators, Italy signaled that resources like factory demand

response, supermarket refrigeration loads or clusters of home batteries can support the grid just as traditional power plants do. Italian UVAM pilot represent a very promising start and show that the flexibility market is mature and ready to expand to more participants, including possibly residential-scale aggregations. Still challenging however is the economic access barrier for small consumers and prosumers. In practical terms, that means while the technology such as smart meters and home energy management may be ready, the economic payoff for a residential customer to adjust usage is still low in many cases [22].

In addition to nationwide market integration, Italy is experimenting with Local Flexibility Markets (LFM) at the distribution level to bridge this gap. These are mechanisms where the local distribution system operator (DSO) can procure flexibility from resources in a specific area to solve local network issues created an online marketplace for local ancillary services [23] [24].

2.3 National Electrical Infrastructure

2.3.1 Layers of the Network: HV, MV and LV

The Italian electrical infrastructure is traditionally organized in a hierarchical structure with three main voltage levels: High Voltage (HV) for transmission, Medium Voltage (MV) for sub-transmission and primary distribution, and Low Voltage (LV) for final distribution to end-users.

The HV network, managed by Terna as TSO, consists of meshed 380 kV, 220 kV and 150 kV lines interconnecting large power plants, cross-border interconnections, and major load centers across the country. Substations then step down voltage to MV (typically 10–30 kV), which is used by regional distribution networks operated by Distribution System Operators (DSOs) to supply cities, towns, and industrial areas. Finally, local transformers further step down MV to LV (400 V three-phase or 230 V single-phase) for delivery to households, small businesses, and other low-voltage consumers [25].

Historically, power flow within this topology was predominantly one-directional. Electricity generation was concentrated in a limited number of centralized power stations connected to the HV grid, transmitted over long distances, and progressively distributed radially outward to consumers at MV and LV levels. Under this paradigm, distribution networks were designed as passive systems, characterized by radial feeders and limited real-time monitoring

or control capabilities, whose primary function was the reliable delivery of energy from the upstream network to end users [26].

2.3.2 Challenges of Distributed Generation

The rapid diffusion of distributed generation (DG) in recent years has fundamentally challenged this traditional architecture. Italy has experienced a significant proliferation of DG units connected at MV and LV levels, particularly photovoltaic (PV) systems. As a consequence, approximately one quarter of the total installed generation capacity is now connected to distribution networks, ranging from residential rooftop PV installations of a few kilowatts to medium-scale wind farms, solar plants, and cogeneration units of several megawatts. This decentralization implies that electricity can be injected into the grid from thousands of nodes within the MV/LV network, at times exceeding local demand. As a result, power flows are no longer strictly top-down and reverse power flows toward higher voltage levels may occur. The traditional distribution network, once a passive receiver of energy, must evolve into an active part of the system.

Indeed, the entire electricity infrastructure is transforming from a paradigm of few large generation nodes to one of many distributed generation sources spread across the grid. This evolution introduces several technical challenges for distribution system operation. One of the most critical issues concerns voltage regulation along MV and LV feeders. Traditional radial feeders were designed under the assumption of monotonically decreasing voltage from the substation to the end user. However, high DG penetration near feeder endpoints can lead to voltage rise beyond regulatory limits when local generation exceeds consumption. Maintaining voltage within admissible bounds therefore requires increasingly coordinated control of on-load tap changers, voltage regulators and inverter-based DG resources [27].

Protection and fault management are also significantly affected. Conventional protection schemes in passive radial networks rely on assumptions regarding fault current direction and magnitude that are no longer valid in the presence of DG. Bidirectional fault currents can reduce protection selectivity or lead to miscoordination, thus requiring adaptive protection strategies and revised coordination schemes. In parallel, reverse power flow at HV/MV transformer interfaces, once a rare operating condition, is now observed with increasing frequency in several areas of Italy. Sustained reverse flows indicate that distributed generation within a primary substation area can exceed to-

tal local consumption, potentially leading to thermal and voltage constraints and necessitating network reinforcements or active operational measures.

Along side these issues there is the limited real-time observability and controllability historically present in distribution networks. Unlike the HV transmission system, which is extensively monitored through Supervisory Control and Data Acquisition (SCADA) systems, many MV and LV networks still lack comprehensive real-time measurement infrastructure. In Italy, a substantial number of medium and small scale DG units are connected without providing real-time telemetry to the TSO, limiting system-wide situational awareness. Although DSOs are progressively deploying advanced metering infrastructure (AMI), visibility of the instantaneous state of the distribution network remains partial in many areas, complicating both local operation and coordination between transmission and distribution system operators [28].

Collectively, these factors necessitate a transition from the traditional “fit-and-forget” approach to DG integration toward a more actively managed distribution network. Solutions currently being implemented or investigated include:

- dynamic voltage control through inverter-based resources;
- active power curtailment or redispatch of DG during congestion;
- enhanced monitoring and state estimation techniques, and the use of economic;
- market-based mechanisms to incentivize behaviors that support network operation.

In this context, distribution grids are evolving toward the paradigm of Active Distribution Networks (ADNs), capable of managing bidirectional power flows and coordinating distributed resources.

This evolution may further involve the deployment of new operational architectures such as microgrids and virtual power plants, in which clusters of DERs are coordinated to provide local and system-level services. Planning and developing such architectures is widely recognized as essential to preserving reliability and power quality as DG penetration increases. Italy’s DSOs, in line with regulatory directives, are moving toward this active network paradigm, investing in smart grid technologies (automation, sensors, advanced control systems) to meet the challenges posed by widespread distributed generation [29].

2.4 Renewable Energy Communities (RECs)

2.4.1 Concept and Regulatory Framework

One innovative concept that has emerged to enhance the integration of distributed renewables and engage consumers in the energy transition is that of Renewable Energy Communities (RECs). Introduced at the European policy level through the Clean Energy Package with the Internal Electricity Market Directive (EU) 2019/944 [30] and then formally defined with the Renewable Energy Directive (EU) 2018/2001 (RED II) [31], RECs are envisioned as legal entities that allow citizens, businesses and local authorities to cooperatively engage in producing, sharing and consuming renewable energy.

The EU directives define broad principles for energy communities, such as open and voluntary participation, autonomy and primacy of environmental and social objectives over profit, and member states have been tasked with incorporating these into national law [32]. In essence, a REC is an autonomous legal entity that bring together private and public actors (citizens, enterprises, local authorities, etc.) to collaboratively generate, share and consume green energy to deliver environmental, economic and social benefits at the community level.

In such vision these communities are seen as a key vector of the clean energy transition, aimed at increasing public acceptance of renewable projects and attracting private investment into the energy sector [33]. By enabling collective generation and local consumption of energy, RECs allow citizens to directly receive benefits through improved energy efficiency, lower electricity bills, and reduced energy poverty.

RECs also empower consumers to become prosumers, fostering a more democratic and decentralized energy model. By enabling distributed generation and local energy sharing, RECs will help improve the resilience of the grid at the community level and can enhance demand-side flexibility to maximize self-consumption [34]. RECs have the potential to contribute to a more sustainable and flexible energy system that is aligned with EU climate goals.

This initiative can further promote sustainable local development and a circular economy. They encourage the use of underutilized spaces, such as disused industrial land or public building rooftops, for installing renewable generation, turning dormant assets into community value. Prosumers with large roofs can install renewable capacity, self-consume the portion of energy

needed and share the surplus with community members. Every member will then benefit from lower energy bills and greater energy independence as a result of this collective action [35].

RECs were formally introduced into the Italian regulation through Legislative Decree No. 199/2021 [36]. A distinctive feature of the Italian implementation is that energy sharing is virtual rather than through private wires. Each member remains physically connected to the public distribution grid and maintains their own supplier contract, but energy sharing is accounted from the grid data. Specifically, the energy community's contribution is evaluated by measuring, for each hour, the minimum between the total renewable energy injected into the grid by the community's generators and the total energy withdrawn by the community's consumers.

This minimum value is termed collective self-consumption or "virtually shared energy". This concept of building such entities on the already present distribution infrastructure simplifies the process of aggregation of the members and the Italian system operator uses these measurements to determine the community's performance and to allocate incentives.

2.4.2 Constraints and Incentives

Such incentives are a crucial component of the REC framework, designed to make community projects financially attractive. In Italy, the implementation of Renewable Energy Communities is further regulated by the CACER and TIAD operational rules issued by the Gestore dei Mercati Energetici [37]. While the whole regulation saw numerous changes during the years, the current version allows any end-user type (residential, small and medium commercial, public) to become members.

Administrative requirements include establishing a legal entity, often a cooperative or association, and registering the REC with authorities while the aggregation may be done on national level, even if the shared energy is evaluated only at the level of the primary distribution substation. The advantage is in sharing the economic cost among different subconfigurations. Another limitation set by Italian law is a size limit on generation capacity, as each renewable installation participating in the REC must be no larger than $1MW_p$ (1 megawatt peak). In practice, this means RECs are oriented toward small to medium renewable projects while larger plants fall under different schemes.

Italy provides a dedicated economic reward for the shared renewable en-

ergy within a community. Such incentive is paid by the GSE (Gestore Servizi Energetici, the national energy service manager) and can be allocated among the community members according to rules they agree upon in their statute. Importantly, this incentive is time-limited as they are currently guaranteed for 20 years for each new REC under the scheme and funded by national energy policy programs aimed at promoting renewables and efficiency.

In the first form of the regulation, the community receives approximately 110 € per MWh (0.11 € per kWh) of electricity that is generated and consumed within the community on an hourly basis. The current and updated regulation introduced a revised incentive structure with a variable component that adjusts according to market electricity prices, among other factors [38]. This change aim to ensure the economic sustainability of RECs in the long run and aligning the incentives with real-time grid needs, by providing higher support when market prices are low, and tapering support when electricity prices are high and self-consumption is already economically attractive.

2.4.3 Operation and Opportunities

From an operational perspective, the goal of a REC is to maximize the matching of local renewable generation with local consumption. In other words, to maximize the shared energy that yields both economic and environmental benefits. Achieving a high self-consumption ratio requires careful coordination of the community's load and generation profiles. Flexible resources play a pivotal role in this. For example, Battery Energy Storage Systems (BESS) and controllable loads. By optimizing such behaviors, a community can substantially increase the fraction of renewable energy that is consumed on-site versus spilled to the grid [39].

The Italian framework implicitly encourages this by calculating the incentive only on the simultaneous portion of generation and consumption, effectively rewarding real-time coordination between prosumers and consumers.

In particular, BESS play a pivotal role in enhancing a community's energy flexibility. By optimally scheduling charging and discharging cycles, a BESS can store surplus renewable energy when generation exceeds immediate consumption and then supply that stored energy during later periods of high demand or low generation.

This not only helps REC members cover a larger share of their needs with locally produced green energy, but also improves overall efficiency and eases pressure on the main grid. Deploying storage at the community level creates

new opportunities both for the grid and the REC members: a sufficiently sized community battery can act as a flexible resource providing ancillary support to the network (voltage regulation, peak shaving, etc.) while also empowering citizens by increasing energy autonomy, enabling them to participate in ancillary service and balancing markets [40].

Such participation could be rewarded at higher rates than the base self-consumption incentive, potentially unlocking new revenue streams for communities while assisting the power system. Furthermore, advances in peer-to-peer energy trading are being explored, where community members could directly trade energy or negotiate the use of each other's flexible resources based on price signals. These concepts remain at pilot stage in Italy indicate the future direction where communities become active cells of the grid able to interact with the broader market and grid for mutual benefit.

2.4.4 Challenges

Beyond their technical and economic aspects, the success of RECs in Europe critically depends on human and social factors. Studies highlight that social acceptance, trust, and member engagement are decisive for long-term sustainability. Participation in community energy projects is strongly correlated with perceived fairness, transparency and trust in both peers and community management [41,42]. Equally important is the fair and transparent distribution of economic rewards among members. Research shows that perceived inequity in benefit-sharing can undermine participation, while transparent and participatory governance increases trust and acceptance and various allocation mechanisms have been proposed, all seeking to balance efficiency and equity [43,44].

The presence of an energy community can however introduce the typical DER challenges to be managed. Community members might be encouraged to alter their normal consumption patterns to increase self-consumption, but care must be taken to ensure this does not create new peaks, rebound peaks or issues in the local grid. It is even possible that a REC configuration led to reverse power flow during peak PV production hours despite improving local self-consumption. That mean that there is an interplay between maximizing economic incentives and maintaining network limits, which requires active energy management within the community.

2.5 Energy Management Systems (EMS)

2.5.1 Role and Working Paradigm

The integration of renewables, distributed generation, and flexible loads into power networks has elevated the importance of Energy Management Systems (EMS) as essential tools for optimizing operations [45]. Broadly defined, an EMS is a control and optimization system that monitors energy resources, essentially generation, storage and loads, and orchestrates their operation to achieve certain objectives, such as cost minimization, efficiency improvement, emission reduction or reliability enhancement. the EMSs are the “brains” that enable intelligent decision-making, using data and algorithms to balance supply and demand inside the monitored system in an optimal way. As these systems operate on energy flows, they typically function on time granularities ranging from minutes to hours.

A typical EMS operates with a rolling horizon control paradigm or Model Predictive Control (MPC). This structure reflects the need to continuously adapt decisions to updated system information and forecast revisions and can be split in a cycle of four key functions:

1. **Data Acquisition:** The EMS continuously collects data from the field, including measurements of power, voltage and current from smart meters and sensors, as well as external information like weather forecasts, crucial for forecasting the production due RES and electricity market prices. Modern EMS implementations rely on an AMI to gather granular consumption and generation data, and often IoT devices to monitor the state of distributed resources.
2. **Forecasting:** Using the acquired data and the known history of the system, the EMS forecasts future system behavior over a given planning horizon. This usually involves predicting the load demand of consumers and the generation expected from renewable sources for the next few hours or days. Forecasting accuracy is a critical factor, as it directly affects the quality and robustness of subsequent control decisions.
3. **Optimization and Scheduling:** Given the forecasts and the current state of the system, the EMS computes an optimal schedule for controllable devices like batteries, electric vehicles charging, flexible industrial processes and heating/cooling systems. This step determines a control

trajectory over the prediction horizon while respecting operational constraints and aims to achieve specified objectives like the maximizing economic benefit and maximizing self-consumption of renewables.

4. Dispatch/Control: Finally, the EMS implements the first control action of the schedule by dispatching control commands to devices in real time. This could involve sending setpoints to battery inverters, signals to smart appliances via a home automation system, or even requests to distributed generators to adjust output. The process is then repeated at the next control interval, incorporating new measurements and updated forecasts. This rolling-horizon mechanism allows the EMS to compensate for forecasting errors and unmodeled disturbances while maintaining computational tractability.

Collectively, these steps allow an EMS to act as a tool for optimal energy flow management at the local level, maintaining a balance between generation and demand and steering the system towards desired performance metrics.

2.5.2 Optimization Techniques

The optimization stage is the core of an EMS and has been extensively studied in the literature. A wide range of optimization techniques has been proposed reflecting trade-offs between optimality, computational complexity, robustness and scalability [46–48]. Following the outline the main categories and their characteristics.

Heuristic Approach (Rule-Based Control): Many early EMS, especially at the home or microgrid level, used rule-based logic due to its simplicity and transparency. A rule-based EMS is typically a set of if-then-else rules or a priority order for resources. Such deterministic rule heuristics are easy to implement and robust, but they usually do not guarantee optimal use of resources. They might lead to suboptimal economic results because they cannot perfectly trade off different objectives.

Nonetheless, rule-based EMS remain popular in commercial building management systems and some microgrids where reliability and simplicity are prioritized over optimality. In hybrid power systems like microgrids or vehicles, deterministic rule-based strategies, often derived from human expertise, have been widely adopted and sometimes enhanced by tuning via optimization algorithms.

Also rule-based EMS in the “priority” form manifest greater reliability under uncertainties, whereas more complex optimal strategies require accurate data and computation.

Deterministic Optimization (Mathematical Programming): A large portion of EMS research formulates the energy scheduling problem as a mathematical optimization problem. Techniques include linear programming (LP), mixed-integer programming for on/off decisions of generators or discrete battery states, quadratic programming for AC power flow optimizations.

MILP (Mixed-Integer Linear Programming) has been a workhorse for microgrid EMS optimization, used to find minimum-cost dispatch schedules over a forecasted horizon [49]. Deterministic optimization methods excel at explicitly handling constraints (like generator limits, ramp rates, network flow limits) and optimizing a clear objective (cost, emissions, etc.).

Even if the near-optimality of the solutions given the inputs is guaranteed, it strongly depends on forecast accuracy and can be computationally intensive as the size of systems or the length of the horizon scale up. Deterministic optimization is often combined with decomposition techniques to handle large scales or with heuristics to obtain faster solutions when exact methods are too slow.

Metaheuristic Algorithms: Metaheuristics such as Genetic Algorithms (GA) and Particle Swarm Optimization (PSO) have been used to solve complex EMS scheduling problems that are non-linear, non-convex or combinatorial [50].

These algorithms do not guarantee a global optimum but can find good solutions in reasonable time and are easier to adapt to different objective formulations. For instance, if an EMS must optimize a multi-objective function (cost, emissions, and battery life simultaneously), formulating a linear MILP may be challenging, so researchers have applied multi-objective GAs to evolve a set of Pareto-optimal operating schedules.

Pure metaheuristics might need long computation or fine tuning to converge, so integrating domain knowledge (rules) provides a shortcut to good solutions. The main downside is the lack of guarantee on optimality. However, for real-time control or very large systems where a good solution now is preferable than the optimal solution too late, metaheuristics algorithms

are more than valuable tools.

Artificial Intelligence and Machine Learning Based Control: In recent years, AI techniques have been increasingly integrated into EMS algorithms. Briefly, AI-based approaches include machine learning models for forecasting and optimization and reinforcement learning (RL) for control.

For example, in many modern EMS designs, neural networks or other regression models forecast the short-term load and renewable generation with high accuracy, feeding those predictions into an optimization engine. More ambitiously, reinforcement learning algorithms are used to let an EMS learn optimal control policies by interacting with the environment, simulated or real, over time. Multi-agent reinforcement learning has recently emerged as a promising approach enabling distributed agents to autonomously learn coordination strategies under dynamic and uncertain conditions [51].

AI-based control offers strong adaptability and can handle non-linearities and high-dimensional state spaces that are challenging for classical methods. The downsides are the need for extensive training data or exploration, which can be risky on real systems, and the “black-box” nature of some AI models that makes them less interpretable and possibly less trustworthy in mission-critical grid operations. A hybrid trend is emerging where AI is combined with model-based methods, like using neural networks to approximate parts of the system model or to warm-start an optimization, or conversely using optimization outputs to safely guide the RL training.

Stochastic and Robust Optimization: To account for uncertainties in renewable output or loads, stochastic EMS approaches incorporate multiple scenarios or probabilistic constraints optimizing the expected cost over many possible forecasts, or ensure constraints hold with a certain probability. Robust optimization takes a worst-case approach, finding a solution that minimizes cost while guaranteeing feasibility for all realizations within a defined uncertainty set. These methods improve reliability, at the cost of a more conservative or higher-cost operation and an increase of the weight of the computational effort. The trend is to combine these with decomposition or heuristics to remain tractable.

In summary, typical EMS control algorithms range from simple heuristics to advanced AI, and there is a clear evolution towards more intelligent,

adaptive strategies. The most effective EMSs are often integration of multiple strategies to combine the strenghts of each of them. The state-of-the-art is characterized by multi-layered control algorithms with forecasting modules, optimization solvers and learning-based controllers working in tandem.

2.5.3 EMS and electrical network

EMSs exist at different control levels of the electrical network, corresponding to the hierarchy of the grid.

At the low-voltage (LV) level, we have Home or building EMS, which focus on individual users or facilities [?, 52]. Their objectives typically include cost minimization, self-consumption maximization, and comfort preservation. Control decisions are largely local, involving appliances, batteries, and EV charging systems.

At the medium-voltage level, EMS operate at the scale of microgrids, campuses, or energy communities. These systems coordinate multiple users and distributed resources, often with additional constraints related to network operation, such as voltage limits and feeder congestion [53]. Community-level EMS act as intermediaries between individual users and the distribution network, aggregating flexibility and aligning local objectives with grid requirements.

At the high-voltage level, EMS functions are integrated into transmission system operation, focusing on system-wide balancing, frequency regulation and congestion management [54, 55]. While traditionally centralized, these EMS increasingly rely on flexibility originating from lower voltage levels through aggregators and coordinated interfaces with distribution systems.

The interaction between EMS operating at different levels has become a central research topic, particularly as flexibility shifts from centralized generators to distributed assets.

2.5.4 Architecture and Coordination

Architectural design about who makes decisions and how information is shared is a fundamental aspect of EMS. The literature identifies three broad EMS architectures, with hybrid variations: centralized, decentralized, and distributed [56, 57].

Centralized EMS: In a centralized architecture, a single control entity gathers data from all assets in the system and computes the optimal control actions for every controllable device. This top-down approach ensures global optimization of the objective since all information is available in one place. Centralized EMS were historically common at utilities and in micro-grid controllers. The advantage is optimal coordination, the downside is that centralized EMS suffer from scalability limits and single-point-of-failure risk.

As the number of devices grows, computational requirements become bottlenecks. In present-day contexts with thousands of distributed assets, fully centralized EMS are often impractical beyond a certain scale. They also raise data privacy concerns, since individual consumers may be unwilling to stream fine-grained energy usage data to a central authority.

Decentralized EMS: A decentralized architecture pushes decision-making to individual devices or local controllers, with minimal central coordination. Each DER or consumer may have its own EMS that operates autonomously based on local objectives and conditions, without explicit communication with neighbors or a central optimizer. Decentralized EMS enhance autonomy and resilience as if one controller fails, others are unaffected and the system can be inherently robust to localized issues.

They also naturally scale with the number of devices, since control is distributed. However, a pure decentralized approach can sacrifice optimality and system-wide coordination. Without oversight, individual controllers might act greedily or in conflict, causing suboptimal or even unstable outcomes for the grid. Thus, decentralized EMS often rely on indirect coordination mechanisms to achieve some level of harmonization and simple decentralized schemes may not fully exploit global optimization opportunities.

Distributed (Multi-Agent) EMS: Between these extremes lies the distributed or multi-agent EMS architecture, often implemented in a hierarchical manner. Here, control is spread across multiple controllers (agents) that communicate and coordinate to achieve a near-global optimum. A popular form is hierarchical EMS, where control is organized in layers. Each layer makes decisions based on its scope, and passes information or setpoints up and down. This approach attempts to combine the benefits of both centralized and decentralized schemes.

Another paradigm is fully distributed multi-agent systems (MAS), where

multiple peer agents, representing DER units or microgrids, negotiate or iteratively coordinate to reach a cooperative solution. In this architecture, no single agent has global authority.

This type of framework often relies on the Alternating Direction Method of Multipliers (ADMM), a distributed optimization algorithm that decomposes large-scale convex problems into smaller subproblems coordinated through dual variable updates.

With this methodology, independent nodes attempt to converge toward an optimal solution through a consensus mechanism. Consensus is enforced by penalizing deviations from shared variables and updating dual multipliers, thereby driving local solutions toward global feasibility and optimality.

The advantages of distributed and hierarchical EMS are improved scalability and fault tolerance as no single failure kills the whole system. On the other side, distributed EMS depend heavily on communication infrastructure and protocols. Coordination requires frequent exchange of data or signals between agents or layers, which introduces latency and potential communication failures.

In practice, many EMS implementations are hybrid combining the different approaches, dividing the problems on multiple layers. One interesting emerging architecture is blockchain-enabled transactive energy platforms, where prosumers and consumers trade energy or flexibility through a distributed ledger system. Such blockchain-based EMS architectures allow peer-to-peer energy transactions and distributed decision-making with inherent security and transparency [58].

2.5.5 EMS for Microgrids and RECs

Microgrids constitute a well-established reference framework for the integration and coordinated operation of DERs [59]. They are commonly defined as localized electrical systems capable of operating either grid-connected or in islanded mode, providing enhanced flexibility and resilience at the distribution level. Due to their limited scale and high penetration of power-electronics-interfaced resources, microgrids have driven the development of dedicated control architectures that differ from those adopted in traditional centralized power systems.

Such structures typically rely on a hierarchical control paradigm, which decomposes control tasks across different time scales and functional objec-

tives [60,61]. This hierarchy is generally organized into three levels:

- The primary control level ensures local stability and power sharing through fast, decentralized mechanisms, such as droop control, based solely on local measurements.
- The secondary control level operates at a slower time scale and focuses on restoring voltage and frequency deviations introduced by primary control actions, often relying on limited communication among controllers.
- The tertiary control level addresses system-wide coordination and optimization objectives. Operating on time horizons ranging from minutes to hours, this level manages power exchanges with the main grid, schedules generation and storage units, and coordinates flexible loads in order to achieve economic, operational, or regulatory goals.

Within this hierarchical framework, the Energy Management System (EMS) is naturally located at the tertiary control level, as it does not directly interact with fast electrical dynamics but instead computes optimal operating strategies based on forecasts, system constraints, and external signals.

The control concepts developed for microgrids provide a direct foundation for the operation of RECs. While RECs differ from microgrids in terms of regulatory structure and physical coupling, they share the same need for coordinated management of distributed resources and flexible demand. In this context, the EMS-centric tertiary control layer can be extended from the microgrid domain to RECs, enabling coordinated energy scheduling and interaction with the distribution network while preserving the autonomy of individual participants.

REC-oriented EMS must simultaneously address multiple objectives such as maximizing collective self-consumption, ensuring fair benefit allocation among participants and respecting distribution network constraints [62–64]. This multi-objective nature coupled with a strong need of scalability, makes centralized control impractical and highlights the need for distributed or hierarchical coordination architectures.

To summarize, existing EMS coordination strategies for Renewable Energy Communities can be broadly categorized into three main paradigms: centralized MILP-based formulations, fully distributed multi-agent systems often relying on ADMM, and decentralized heuristic approaches.

Aspect	Centralized	ADMM-based MAS	Heuristic
Intelligence	Central	Distributed	Local
Edge Control	Central setpoints	Local optimization	Rule-based
Optimization	MILP / convex	Convex + consensus	Metaheuristic
Forecasting	Central	Local / shared	Local
Scalability	Limited	High*	High
Communication	Moderate	Iterative*	Low
Computation	High	Distributed	Low

*Convergence dependent on coupling strength and tuning.

Table 2.1: Structural comparison of EMS coordination paradigms in REC literature

Centralized methods ensure global optimality but suffer from scalability and single-point-of-failure issues. Fully distributed ADMM-based architectures enhance scalability and privacy but may require significant iterative coordination and careful parameter tuning to guarantee convergence. Decentralized heuristic approaches reduce computational burden but typically lack global coordination guarantees.

2.6 Resilience in Distribution Networks

2.6.1 Definition

Resilience in electrical distribution systems is commonly defined as the ability of the system to withstand, adapt to and rapidly recover from high-impact-low-frequency (HILF) disruptive events, such as extreme natural disasters or large-scale blackouts [65]. A fundamental aspect is the temporal dimension. The evaluation is not limited to the system's capability to resist an initial disturbance but also includes the speed and effectiveness of recovery. This depends on the availability of reserve resources, operational flexibility, automation and well-defined emergency procedures.

In practical terms, a resilient power network is designed not only to reduce the likelihood of outages but, more importantly, to absorb severe disturbances limiting their consequences and duration, while minimizing service interruptions, preserving the supply to critical loads whenever possible and restoring normal operating conditions within a short time frame.

This concept goes beyond traditional reliability and explicitly targets rare but catastrophic events capable of simultaneously affecting multiple components of the grid. Examples include natural disasters, coordinated physical attacks or technical failures. As in recent years climate change and system complexity have raised both the probability and severity of such events the relevance of resilience has increased significantly.

2.6.2 Metrics

To effectively design and enhance the resilience of electrical distribution networks, it is crucial to quantify resilience through suitable quantitative indicators. Resilience metrics provide a systematic means to evaluate system performance under adverse conditions and are used both during the planning phase and during operation, for example to monitor the real-time state of the network in emergency situations.

Over the past few years, a wide range of resilience indices and assessment frameworks has been proposed in the technical and scientific literature [66–68]. However, a universally accepted metric has yet to be established, as different approaches tend to focus on specific aspects of resilience, including structural robustness, recovery speed, and social impact. In the following, the main categories of resilience metrics and evaluation methods adopted are reviewed.

Qualitative and Quantitative Metrics

A first fundamental distinction in resilience assessment is between qualitative and quantitative evaluations.

Qualitative metrics typically rely on descriptive analyses, checklists and expert judgment to estimate the level of preparedness and response capability of a system, without reducing resilience to a single numerical value. A network operator may qualitatively assess resilience by verifying the existence of formalized emergency plans, the availability of backup power supplies for critical components or the level of staff training for crisis management.

Studies often adopt rating scales, like from 1 to 5, across multiple factors or use qualitative descriptors such as low, medium, or high to characterize infrastructural, organizational and community resilience. Although this approach is less mathematically rigorous, it is particularly useful when quantitative data are limited.

Quantitative metrics express resilience through numerical values derived from system data and models. A widely adopted approach consists in representing the temporal evolution of electrical system performance during an adverse event, commonly referred to as the resilience curve or resilience trapezoid.

From this curve meaningful indicators are derivable. Before a disruptive event, the system operates at full functionality meeting the entire demand then, when the disaster level event occurs, system performance drops to a lower level and then gradually recovers as resources are restored.

Three key parameters emerge from this dynamic behavior:

1. the magnitude of performance degradation immediately after the event, such as the percentage of unserved load or the number of customers without electricity;
2. the rate at which performance declines at the onset of the event, which reflects the system's ability to absorb the initial shock;
3. the recovery time, or resilience time, defined as the interval required to restore service to near-normal levels, like recovering 90% of the demand.

The evaluation of the area defined by the resilience trapezoid curve and the one corresponding to uninterrupted operation represent a common example of a composite resilience indices. This indicator provides a compact representation of the total service loss experienced by the system as a smaller area correspond directly to higher resilience.

Static and Dinamic Metrics

Another way to classify resilience metrics is to distinguish between static and dynamic metrics.

Static metrics assess resilience a priori, based on the intrinsic properties of the network before a disruptive event occurs. This category typically includes indicators of structural robustness, such as the number of alternative paths available between generation sources and loads in the event of line outages and the amount of available spare capacity.

These indices often rely on probabilistic analysis methods—such as N-1 or N-k contingency analyses and probabilistic power flow models or on graph-theoretic approaches, providing a measure of the network's vulnerability.

Dynamic metrics, by contrast, refer to the temporal response of the system during and after a disruptive event. The metrics discussed previously, such as the extent of service loss and recovery times, fall within this category. Examples of a dynamic metric is the restoration duration curve or the recovery speed metrics quantify the rate at which service is restored per unit of time. Such dynamic indicators enable a direct comparison of different restoration strategies.

Due to the multidimensional nature of resilience however, no single metric can fully capture all its aspects. Consequently, resilience assessment often relies on a set of complementary indicators to address different perspectives. For example, a more comprehensive and informative resilience assessment can be achieved by combining these elements:

- percentage of critical customers demand guaranteed during a blackout;
- average restoration time under disaster scenarios;
- amount of energy not supplied during extreme events;
- a qualitative score of organizational resilience.

Resilience metrics are increasingly integrated into distribution network planning and regulations. Scenario-based analyses are used to identify critical assets, prioritize investments and quantify the benefits of resilience-oriented solutions such as microgrids or automation [69]. In this context, resilience metrics support cost–benefit analyses by explicitly accounting for avoided social and economic losses during extreme events. Their adoption is opening the way toward regulatory mechanisms that reward utilities for resilience improvements, complementing traditional reliability-based incentives.

2.6.3 Challenges

Challenges to resilience in distribution systems are multifaceted. Unlike transmission networks which are often meshed and can tolerate single contingencies by redirecting flows, many distribution grids are traditionally designed with radial topologies and minimal redundancy.

This means that a single component fault may compromise all downstream customers supply until it is repaired. Also, distribution infrastructure are often exposed long stretches of overhead lines run through forests or along roads, making them vulnerable to high winds, ice accretion or falling trees.

Climate change is exacerbating these risks by increasing the frequency and severity of extreme weather events.

Meanwhile, the constant increasing electrification of society means that the socio-economic costs of such outages is increasing, raising the importance of enhancing resilience.

The growing penetration of DERs is creating both opportunities and new challenges [70, 71]. As already discussed about, electrical distribution networks are undergoing a structural transformation from centralized architectures toward highly decentralized systems characterized by widespread distributed generation. This paradigm shift is widely recognized as an opportunity to enhance resilience.

DERs, when properly coordinated, can maintain local supply even when the upstream grid is compromised, reducing dependency on transmission networks and centralized assets. In an advanced scenario, the main power system can be conceptualized as a “network of microgrids,” where local sub-systems are capable of autonomous operation and fault isolation. Such architectures can preserve critical services during major disruptions and prevent system-wide collapse by confining failures to limited areas.

On the other side of the coin, high penetration of non-dispatchable renewable sources like photovoltaic and wind generation depend on weather conditions and may not align with local demand profiles, leading to undesired stress on the network. The key point is that distributed resources improve resilience only if they are intelligently managed. Otherwise, their stochastic nature can exacerbate system vulnerability during critical conditions and, paradoxically, lower its own resilience.

2.6.4 Improving the Resilience

To bolster distribution network resilience, a variety of mitigation strategies are employed or under development:

Network Reinforcement: The first way is the hardening of the infrastructure starting by strengthening physical components to better withstand stress. Examples include using stronger poles or burying lines underground in areas prone to wind damage or elevating substations above flood levels.

Another way is to increase the network reconfiguration capability and redundancy. Modern distribution networks are increasingly designed with tie switches between feeders and automated switchgear that can re-route power

and use sensors and intelligent switches to detect faults and automatically isolate the faulted section to restore the supply in the unaffected parts in seconds or minutes.

While such improvements can significantly improve the overhaul Resilience of the system, it is often capital-intensive and thus is applied selectively to critical segments or where cost-benefit analyses justify it.

Enhanced Monitoring and Control: As noted, situational awareness is key in crisis scenarios. Distribution automation systems and EMS are being deployed to get real-time data on fault locations and system status. Smart meters can report loss of power at customer premises, helping pinpoint outage extents. Feeder telemetry can detect variations that signal a fault or imminent problem. With better data, operators or automated schemes can make faster decisions to isolate problems and guide repair crews efficiently.

Also, as stated already, the presence of a monitoring system able to coordinate the DER is getting more and more relevant to their proper integration in the grid and to counter their eventual problematics. In essence, a smart grid with advanced EMS can significantly enhance resilience by reducing both the detection time and restoration time of outages, while improving the integration of DERs.

Microgrids and Islanding Capability: Perhaps one of the most promising strategies for distribution resilience is the use of microgrids as partitions of the overall grid. By doing so, such subsets are able to disconnect from the main grid and operate autonomously using local generation and storage. During normal operation, each microgrid is grid-connected and may exchange power.

However, during an outage in the main grid, the microgrid controller can isolate the system and continue supplying its internal loads using local resources. This islanding capability limits outage propagation and preserves electricity supply for critical infrastructures such as hospitals, emergency shelters, and communication centers.

Island-capable microgrids must support black start functionality, enabling them to recover autonomously after a blackout without external power. Once operational, microgrids can also assist in the restoration of surrounding network sections, acting as anchors for system recovery. In this way, they provide not only resilience to the participants within the microgrid but can also

support the wider power network.

Real-world evidence confirms these benefits. During recent hurricanes in the United States, community microgrids powered by photovoltaic systems and battery storage remained operational while large portions of the grid were offline, serving as resilience hubs for affected populations [72].

2.6.5 Resilience and RECs

RECs represent a complementary organizational model that enhances resilience through collective and cooperative action. In the Italian regulatory framework, RECs operate through virtual energy sharing and, while this mechanism promotes load shifting, increases self-consumption and reduces peak demand on the distribution network, they are not inherently designed for islanded operation.

However their widespread deployment of distributed generation and storage creates latent resilience capacity, working both on the social acceptancy and on the penetration and coordination of the DERs. With appropriate control layers and emergency protocols, RECs could evolve toward hybrid configurations combining virtual sharing with physical microgrid functionalities.

2.7 Discussion and Research Gaps

The literature reviewed in this chapter clearly shows that modern distribution networks are undergoing a profound transformation driven by the combined effects of market liberalization, high penetration of distributed renewable generation, and increasing demand-side flexibility. Electricity markets increasingly rely on granular price signals to reflect system conditions, while distribution infrastructures, originally designed for passive operation, are now required to accommodate bidirectional power flows, local congestion and voltage regulation challenges.

RECs emerge in this context as a promising organizational and regulatory model to foster local energy sharing, social acceptance and citizen participation in the energy transition. However, the literature also highlights that RECs introduce nontrivial operational challenges as maximize collective self-consumption, ensuring fair allocation of economic benefits and maintaining compliance with distribution network constraints require active coordina-

tion of the resources. Incentive-based mechanisms alone are insufficient to guarantee system-optimal behavior, particularly under high penetration of flexible loads and storage.

EMSs are consistently identified as a key enabling technology to address these challenges. As DERs and microgrids are widely recognized as enablers of resilience, their effective contribution critically depends on the presence of intelligent control layers capable of coordinating resources under both normal and emergency conditions.

These gaps motivate the need for practical and implementable EMS frameworks that explicitly account for scalability, regulatory constraints and real-world deployment considerations. A multi-level energy management architecture designed for renewable energy communities and distribution networks proposed in the following chapters is build on the insights derived from all of this consideration, with particular emphasis on simplicity of implementation, robustness and resilience-oriented operation.

Chapter 3

Proposed Architecture

This work proposes a multi-layer Energy Management System for Renewable Energy Communities, designed to be scalable, resilient, and deployable in real electrical networks. Intelligence is distributed across domestic, secondary, and tertiary control levels, combining lightweight rule-based strategies with coordinated optimization. Local controllers operate on low-complexity forecasting and decision logic, while higher layers exploit digital twins, price-based coordination, and linearized optimal power flow to ensure efficient and reliable system-wide operation. A human-machine interface based on a mobile application and LLM-assisted interaction is explored to improve accessibility, enabling non-expert users to interact with complex control systems through natural language while preserving deterministic decision-making.

The proposed architecture, showed in Figure 3.1, aim to clear strategic need: building a control system that is modular, scalable and realistically implementable in a real electrical network. The philosophy followed was simple: in real time control, a good solution delivered quickly always beats a perfect one delivered too late. Considering that a realistic deployment may involve hundreds of distributed nodes, scalability had to be addressed from the very beginning.




Figure 3.1: Proposed Architecture

To achieve these objectives and to enhance the resilience and adaptability of the overall system the choice fell on a compromise among the typical solution presented in literature. The proposed approach is to divide the coordination problem in a hierarchical form and then distribute intelligence across multiple levels, mirroring the hierarchical structure of actual power grids.

The architecture is therefore organized into three layers, named and structured as follow:

- **Home Energy Management System (HEMS)**, a domestic layer;
- **Community Energy Management Hub (CEMH)**, a secondary-substation layer that work as coordinator;
- **Overlord Energy Management System (OEMS)**, a higher-level supervisory layer associated with primary substations or large-scale energy infrastructures.




Figure 3.2: First Control Layer Architecture

Combining these layers allows the system to start from lightweight, local decision blocks and progressively integrate higher-level coordination mechanisms.

Primary Control Layer

At the lowest level, the domestic controller implemented on the prosumer must remain computationally light. Every node in the network includes forecasting and planning/decision blocks, both designed to operate with minimal computational resources. Figure 3.2 show the architecture of the implemented HEMS.

For this reason, classical heavy optimization methods, such as linear programming or MILP, were deliberately avoided. Instead, the design relies on rule-based techniques and simple decision logic that can be executed on low-power devices such as microcontrollers or smart meters.

Similar effort has been placed on the choice of the forecasting methodology. Over the years, several lightweight forecasting modules were developed specifically for this purpose, aiming to find a solution compatible with hardware constraints.

At the domestic layer, the controller receives local measurements of consumption, production, and any available contextual information and combines them with external signals such as buying/selling energy prices. Based on these inputs, the unit performs its own optimization to protect the user's




Figure 3.3: Second Control Layer Architecture

interest. This optimization also integrates a scheduling function for controllable loads, implemented through a heuristic algorithm.

Secondary Control Layer

The second layer of control introduces a crucial architectural choice that significantly improves scalability, resilience and communication efficiency: the integration of digital twins within each supervisory controller. Instead of relying on continuous data exchange among all domestic EMSs, which would scale poorly as the number of nodes increases, the second-level controller locally emulates their behavior through lightweight digital twins. This enables the optimization algorithm, based on a variant of Particle Swarm Optimization (PSO), to evaluate multiple scenarios rapidly and adjust particle values without saturating the communication infrastructure. Figure 3.3 show the architecture of such CEMH.

At this intermediate layer, the controller assumes responsibility for coordinating the prosumers assigned to its domain. It manages both energy production and storage, such as BESS scheduling, and demand flexibility including controllable loads. When necessary, the second-level controller can also execute selective load shedding. This function is essential for system resilience as it allows the preservation of critical users, such as essential services or medical facilities, while preventing broader blackouts in situations

of energy scarcity.

Coordination at this layer follows an important design principle: the system does not force individual EMSs into predetermined behaviors. Instead, it influences their decisions by modifying the perceived energy price used inside their local optimizations. This strategy ensures that users remain free to act according to their own economic best interest, thereby maintaining acceptability in real residential deployments. No homeowner would accept a system that increases their electricity bill or forces decisions that contradict their preferences.

To preserve fairness within an Energy Community, the second-level controller can redistribute part of the collective incentives earned at the community level. This compensates and reward users for adopting the “virtuous” behaviors suggested by the monitoring system, such as sharing surplus production or adjusting load profiles, that may be individually disadvantageous in the short term but beneficial for the community as a whole. Through this mechanism, the system ensures both cooperative behavior and user satisfaction.

The second-level controller thus serves as a bridge between local autonomy and community wide optimization. It receives forecasts and performance indicators from the domestic layer, incorporates higher-level directives when necessary, and shapes the collective response through incentive signals rather than coercive commands. Its objectives may include maximizing virtual self-consumption, performing peak shaving, reducing network stress or implementing community-level policies defined at higher hierarchical layers.

Tertiary Control Layer

A third level of control has also been developed and explored over the years, completing the hierarchical architecture. Unlike the lower layers, which must operate under strict computational constraints, this upper-level controller can rely on significantly higher computational resources. This allows the execution of more sophisticated algorithms such as linearized optimal power flow (OPF) formulations while maintaining real-time applicability. In Figure 3.4 such architecture is visually represented.

At this point scalability is no longer a limiting factor: the controller supervises only the second-level coordinators, rather than individual prosumers, and therefore handles a problem of much smaller dimensionality. Its purpose is not to impose direct decisions on users, but to guide the overall system




Figure 3.4: Third Control Layer Architecture

toward globally optimal operational points.

The third level extends this philosophy to a higher voltage domain. Its natural positioning is above primary substations, interfacing between multiple communities and the transmission network operating at a scale where multiple primaries and their associated second-level controllers are coordinated simultaneously. The use of OPF-based methods becomes not only feasible but advantageous, as it enables accurate modelling of network losses, transfer limitations, and the physical cost of moving energy across the grid.

This level controller receives aggregated forecasts of both consumption and production from the lower layers and combines them with the predicted availability of higher-cost or fuel-based generation resources. In the prototype implementation, this framework was validated using the Ethiopian power system, but the approach is general and can be applied to any transmission network.

From these inputs, the controller computes an optimal global strategy. Its objectives include:

- minimizing electrical losses across the network
- reducing the use of expensive or carbon intensive generation

- satisfying the largest possible portion of the demand
- ensuring that the coordinated system remains within operational limits

Its outputs to the second-level controllers are of two types:

- Profile shaping commands, expressing the desired aggregate behavior of each community in order to optimize global system performance
- Price-signal adjustments, when operating in price-based coordination mode, to influence the behavior of local EMSs without overriding their autonomy

Additionally, it is possible to implement a resilience function such as issuing load-shedding commands when energy scarcity makes it impossible to serve all loads. Because such system has a full overview of network constraints and limited resources, it can determine which areas or communities must reduce demand to maintain system stability.

This layer acts as the strategic orchestrator of the entire system. It integrates forecasts, evaluates network constraints, coordinates multiple primary substations and shapes the behavior of subordinate controllers. All of it to ensure efficient, resilient and sustainable operation of the power system as a whole.

Human-Machine Interface and LLM Integration

While Energy Management Systems traditionally focus on optimization and control, their practical deployment in real communities highlights the need for effective human-system interfaces. As EMS architectures become more layered and distributed, the ability to explain system behavior, simulate alternative scenarios, and support user-level decision making becomes as important as achieving optimal control performance.

To complete the proposed architecture, the envisioned solution is a mobile application that users can download on their smartphones. This app is intended to expose selected functionalities of both the first and second level controllers, providing visibility and control without overwhelming the user with the underlying technical complexity.

An aspect explored during this research is the integration of Large Language Models (LLMs), such as ChatGPT, into the interface. The goal is not merely to provide a chatbot, but to create an intelligent AI-agent capable of acting as an “expert intermediary” between inexperienced users and the

monitoring and control system. Many individuals interacting with energy management tools are not trained professionals, and conventional interfaces often fail to bridge that gap. In Figure 3.5 the concept is displayed.




Figure 3.5: LLM Integration Architecture

The challenge lies in combining two fundamentally different worlds. On one hand, LLMs offer exceptional flexibility in interpreting natural language inputs, handling ambiguity and incomplete information with ease. On the other hand, traditional control software requires deterministic, rigid, and precisely structured inputs; even minor deviations can cause failure. The research addressed this mismatch by adopting JSON as a unifying communication layer. In this configuration, the LLM does not replace the control logic; instead, it acts as an adaptive, user-friendly interface that expands accessibility while preserving correctness and reliability.

Although this integration has not yet been extended to the full multi-level controller, it has been successfully implemented and tested within a dedicated simulator of Renewable Energy Communities, yielding promising results.

Chapters 4, 5, and 6 provide detailed descriptions of the first, second and third-level controllers, respectively. The final chapter includes an overview of the work carried out on LLM integration, outlining its current state and future research directions.

Chapter 4

Forecasting of Consumption and Generation Profiles

*This chapter investigates forecasting techniques for electrical load and generation profiles within Energy Management Systems, with particular focus on robustness, computational efficiency, and real-time applicability. Low-complexity feed-forward neural networks are proposed and compared with seasonal persistence and LSTM benchmarks, demonstrating comparable accuracy with significantly reduced computational cost. The chapter further explores complex-valued neural networks, highlighting their ability to encode phase and magnitude information and their affinity with periodic load patterns. Finally, a stochastic load profile generation method is introduced to support robust optimisation, enabling the creation of realistic scenario flocks under uncertainty. The results confirm the effectiveness of the proposed approaches for resource-constrained EMS applications.*¹

¹Part of this chapter has been published as *Low-Complexity Neural Networks for Electrical Load Forecasting in Renewable Energy Communities* [73] and as *Load Profile Generation for Robust Optimization: A Stochastic Approach Based on Conditional Probability Approximation* [74]. Another part of this work was conducted while the author was a visiting Ph.D. student at Manhattan College, New York (USA) from March to May 2024 (working with Prof. Igor Aizhemberg)

4.1 Introduction to Electrical Load Forecasting

The forecasting of time series has been extensively studied, with methodologies ranging from classical statistical models to modern machine learning and deep learning architectures. Forecasting is the activity of attempt to predict the future trend of a profile based on its own historical behaviour and the most recent past.

A common baseline for evaluating forecast performance is the persistence model, which assumes that the next value will be equal to the most recent measurement. In multi-step forecasting, the analogue is the seasonal persistence model, where each future point is predicted by repeating the value at the same hour of the previous day or week. These baselines are simple, but surprisingly competitive, and any forecasting method must demonstrate consistent improvements over them to be considered effective.




Figure 4.1: Concept of Forecasting

The forecasting problem for energy applications is a long-standing and well-studied problem in power systems, yet it remains inherently challenging, especially at fine temporal and spatial resolutions. Although forecasting energy production, at least in the case of photovoltaic systems, largely reduces to the availability and accuracy of meteorological predictions, electrical load forecasting is fundamentally different.

Load time series are characterised by a strong stochastic component, driven by human behaviour, lifestyle variability and external factors that are difficult to model explicitly. As a result, load forecasting is widely regarded a more complex and uncertain task, especially when the objective is to forecast the residential consumption of single households [75].

Statistical models such as ARMA, ARIMA, and ARIMAX have historically shown strong performance for time series with stable autocorrelation and identifiable seasonal patterns [76,77]. These models are computationally lightweight and interpretable, which makes them suitable for environments with limited computational resources. However, their linear structure limits their ability to capture nonlinear dependencies and abrupt changes typical of prosumer electrical loads.

With the advent of modern computing power, Artificial Intelligence (AI) and Deep Learning (DL) approaches have become central to short-term load forecasting. Techniques investigated in recent years include:

- Feed-Forward Neural Networks (FFNNs) [78,79]
- Jordan–Elman Recurrent Neural Networks (RNNs) [80]
- Long Short-Term Memory (LSTM) networks [81]
- Transformer architectures [82,83]

Among these, LSTM networks are widely regarded as the state of the art for short-term electrical load forecasting due to their ability to capture temporal dependencies over long horizons and to handle complex nonlinear relationships. Figure 4.2 show the internal architecture of a cell of such networks.




Figure 4.2: LSTM recurrent cell. The σ and $\tanh$ blocks indicate respectively the application of the sigmoid function and the hyperbolic tangent function.

However, increasing complexity does not necessarily translate into better performance. Due to the high stochasticity of electrical load time series,

large neural architectures often experience overfitting when trained on limited data, and their computational cost can hinder real-time EMS operation. This motivates the investigation of low-complexity neural predictors that balance prediction accuracy with training speed, inference time, and robustness.

4.2 The Role of Forecasting in Energy Management Systems

Forecasting is a fundamental component of any Energy Management System (EMS). Based on historical measurements of a time series, the objective of forecasting is to predict its future evolution over a selected time horizon.

In the context of electrical systems, forecasting is typically performed with a sampling granularity between 15 minutes and one hour, depending on the speed and resolution required by the control logic. Accurate short-term forecasting becomes essential whenever an EMS is deployed, because the efficiency of the control strategy is heavily dependent on the system's ability to anticipate future operating conditions. Figure 4.3 show the planned place of the forecasting in the developed architecture.




Figure 4.3: Forecasting blocks in an HEMS architecture

The primary goal of an EMS is to optimise the operation of available flexible resources, such as controllable loads (CLs) and Battery Energy Storage Systems (BESSs), with respect to a defined criterion, often economic or energy-related. These optimisation problems are based on predictive information about consumption and generation.

For example, scheduling the charging or discharging of a BESS requires projections of future net load to avoid costly grid imports or to maximise the use of renewable generation. Similarly, shifting or curtailing flexible loads is meaningful only if the system has a reasonable expectation of what the consumption or production profile will look like in the near future.

Because EMS architectures are commonly implemented using a rolling-horizon approach, forecasting and optimisation are executed repeatedly over time: at each step the forecasted trajectories are updated with the latest measurements, and the optimisation is re-evaluated. As a result, the entire control quality is deeply tied to forecasting performance: inaccurate forecasts lead to suboptimal decisions, unnecessary battery cycling, and reduced economic benefit.

In contrast, more precise predictions extend the effective foresight of the EMS and improve its ability to organise energy flows efficiently. Forecasting is not an optional add-on, but a structural pillar of any EMS designed to operate in real-world conditions.

4.3 Characteristics of Electrical Load and PV Generation Time Series

Electrical load forecasting is notoriously challenging due to the complex statistical behaviour of consumption signals. By examining typical prosumer load profiles, it is possible to distinguish at least three characteristic daily patterns: one for weekdays, one for Saturdays, and one for Sundays or holidays.

This reflects the presence of strong daily and weekly seasonal patterns. However, these periodic components coexist with a high level of irregularity. Electrical load data are characterised by:

- high short-term variance
- sudden unpredictable spikes
- weak correlation among consecutive hours
- dependence on human habits and exogenous factors

These aspects significantly reduce the predictive power of purely deterministic or low-order stochastic models and motivate the adoption of more flexible data-driven methods.

In contrast, PV generation exhibits more regular behaviour. Its trend is dictated by solar irradiance, which induces a strong deterministic daily seasonality and, under clear-sky conditions, a smooth and physically constrained profile. PV time series typically show:

- strong correlation among adjacent samples
- very low variance under stable weather
- predictable sunrise and sunset patterns

When meteorological forecasts are available, predicting PV production reduces to combining irradiance predictions with a plant model. However, in scenarios where weather input is not available, the EMS must rely on offline forecasting methods. Even under these circumstances, PV generation is usually easier to predict than electrical consumption due to its more structured and less stochastic nature.

Understanding these intrinsic characteristics is crucial as they determine which forecasting techniques can be effectively applied and how to adjust the model's complexity for operational use within an EMS.

4.4 Low-Complexity Forecasting Models for EMS Operation

In many Energy Management Systems, the forecasting module must operate under strict timing and computational constraints. This is particularly true in real-time or near-real-time control architectures, where the EMS must repeatedly acquire measurements, compute updated forecasts, run an optimisation routine, and dispatch control commands.

In such environments, and when the monitored system increases in dimension, highly complex forecasting architectures may become impractical: their training time, inference latency, and computational footprint can quickly erode the time budget available for the rest of the control pipeline.

For this reason, low-complexity neural models have been studied and presented as an attractive compromise. The model is designed to capture the essential temporal regularities of electrical load and PV production while avoiding excessive architectural depth or recurrent structures. Their main advantages can be summarised as follows:

- Fast inference: Small models require only a few matrix multiplications, enabling predictions to be computed in milliseconds. This ensures that forecasting never becomes the bottleneck of the EMS cycle.
- Rapid and frequent retraining: Because the models contain a limited number of parameters, they can be retrained online using sliding windows of recent data. This allows the system to quickly adapt to changes in user behaviour, seasonal shifts, or unexpected consumption patterns without interrupting control operations.
- Reduced risk of overfitting: Electrical load signals are inherently noisy and only partially predictable. High-capacity models tend to overfit this randomness, learning spurious correlations that degrade forecast's accuracy. Low-complexity models, on the other hand, naturally regularise the prediction task and often generalise better under realistic data conditions.
- Simplified deployment and scalability: In multi-user or community-level EMS scenarios, a lightweight predictor can be installed on each household EMS or edge device. This decentralisation enables scalable architectures with hundreds or thousands of users without requiring dedicated high-performance servers.

The low-complexity architecture adopted in this work is a Feed-Forward Neural Network (FFNN) with one or two hidden layers of limited size. The input vector is constructed using a fixed-width temporal window of past measurements, in our case the past three days, and the output is a multi-step forecast generated in a sequence-to-sequence fashion, allowing all future samples within the prediction horizon to be estimated simultaneously.

4.4.1 Explored Approach

To evaluate the effectiveness of the proposed forecasting method, seasonal persistence was adopted as a baseline, while an LSTM neural network served as a state-of-the-art reference model. The methodology was tested using both MISO (Multiple-Input Single-Output) and MIMO (Multiple-Input Multiple-Output) approaches to allow a fair comparison across forecasting strategies.

In the MISO architecture, the forecasting technique is based on a set of 24 Feed-Forward Neural Networks (FFNNs), one for each hour of the day. This choice exploits the strong daily seasonality of electrical load signals. Each

FFNN is specifically trained to predict the value associated with a particular hour. As a result, both this method and the one described in the next section employ 24 neural networks to obtain the 24-step forecast.

Defining E_h as the measurement of the energy consumed in the last temporal slot before performing the forecast, Figure 4.4 represents a visual example of the use of the described set of neural networks.




Figure 4.4: On the left, the concept of the multiple FFNN to predict different section of the profile, while on the right the forecasting process employing the set of MISO FFNN is visualized

Each MISO network consists of a single hidden layer with 20 neurons and logsig activation, receives an input vector of size $n_p = 72$, and outputs a single predicted value. Every network is assigned a network label $h_n \in 1, \dots, 24$, corresponding to the predicted hour. Similarly, each input sequence is assigned a sequence label $h_s \in 1, \dots, 24$, based on the hour of its last sample.

For instance, $h_s = 1$ is assigned to all sequences ending at 01:00 a.m. The h_n -th network is trained to predict the next sample of the series labelled $h_s = h_n - 1$. The network with $h_n = 1$ is trained on the sequences with the label $h_s = 24$, with wrap-around for $h_n = 1$, where the sequences with $h_s = 24$ are used. Given the initial data set containing 17,520 sequences, each network is trained on 730 sequences.

During inference, the MISO predictors are executed sequentially: starting from the network associated with the current sequence label, the 24 networks are applied one after another to generate the full 24-sample forecast.

In the MIMO architecture shown in Figure 4.5, the model again employs 24 FFNNs, one per hour of the day, but this time each network has a MIMO




Figure 4.5: Forecasting process employing the set of MIMO FFNN

structure. Each MIMO network contains 48 hidden neurons with logsign activation, takes as input the same n_p samples, but outputs the full 24-step forecast in a single execution. Sequence and network labels follow the same rules as in the MISO approach.

The primary advantage of the MIMO architecture is the significantly shorter inference time, as only one network execution is required to obtain the entire forecast horizon.

4.4.2 Results and Discussion

Figures 4.6, 4.7, 4.8, 4.9 present representative examples of predictions obtained with each forecast strategy. In every plot:

- the blue curve shows the input sequence (past 72 samples)
- the red curve shows the ground-truth future sequence (target)
- the green curve shows the model prediction




Figure 4.6: Example of prediction on the test-set with persistence.




Figure 4.7: Example of prediction on the test-set with LSTM




Figure 4.8: Example of prediction on the test-set with MISO FFNNs.




Figure 4.9: Example of prediction on the test-set with MIMO FFNNs.

A recurring behaviour can be observed: while the LSTM effectively captures the average daily profile, it often fails to reproduce sharp load peaks—particularly the early-morning peak around the 80th–85th sample. In contrast, both FFNN-based approaches more successfully identify secondary peaks, likely because each individual FFNN specialises in a specific hour-of-day pattern.

Table 4.1 reports aggregated results throughout the test dataset, averaged among the 127 users. The metrics include the MAPE for:

- the first predicted sample
- the last predicted sample
- the entire 24-step horizon

Table 4.1: Prediction results of the proposed forecasting methods obtained on the testset.

Forecasting technique	Seasonal persistence	FFNNs (MISO)	FFNNs (MIMO)	LSTM
Inference time	$< 10^{-3}s$	$\approx 0.2s$	$\approx 0.01s$	$\approx 0.5s$
Training time	—	$\approx 10s$	$\approx 10s$	$> 1 \text{ min}$
MAPE 1 st point	0.394	0.326	0.298	0.329
MAPE 24 th point	0.394	0.329	0.308	0.340
MAPE sequence	0.393	0.328	0.305	0.332

Figure 4.10 shows the MAPE disaggregated by hour of day. From these results, several observations emerge:

- All neural models outperform seasonal persistence, improving accuracy by approximately 10

- Despite their differences, MISO FFNNs, MIMO FFNNs, and LSTMs perform similarly, with the MIMO FFNN showing slightly superior accuracy. When no exogenous variables, such as temperature, are available, even simple architectures such as FFNNs can achieve accuracy comparable to that of deeper recurrent models.
- The persistence model exhibits uniform error across the prediction horizon, while neural approaches perform better on the first few forecasted points. This indicates that neural networks can adjust predictions in response to recent deviations in the load profile.
- Training time strongly favours FFNNs: Both MISO and MIMO FFNNs require approximately 10 seconds to train, whereas LSTMs require minutes to tens of minutes, making them less suitable for scenarios requiring frequent retraining.
- Inference time differs between models: LSTM and MISO FFNNs have similar inference times, since MISO FFNNs require 24 sequential executions.
- The MIMO FFNN is the fastest, producing the entire 24-point forecast in a single execution, with inference times of only a few tens of milliseconds.




Figure 4.10: Prediction accuracy (MAPE) obtained on each hour of the day for the considered predictors.

This study demonstrates the effectiveness of FFNN-based methods for

short-term electrical load forecasting using only load time series as input. Both architectures—MISO and MIMO—leverage the intrinsic daily seasonality of consumption profiles by training 24 specialised networks. Compared to an LSTM, FFNNs achieve similar forecasting accuracy (MAPE $\approx 30\%$), while being significantly faster to train and deploy.

Among the explored methods, the MIMO FFNN emerges as a more than suitable choice for EMS applications, primarily due to its optimal balance between forecasting performance, computational efficiency, and extremely low inference time. These characteristics make it particularly attractive for real-time and rolling-horizon energy management systems.

4.5 Complex Valued Neural Networks in Load Forecasting

During the research period abroad at Manhattan College in New York, a dedicated investigation was carried out on the application of Multi-Level Multi-Valued Neurons (MLMVNs) to the electrical load forecasting problem. Unlike traditional real-valued neural networks, MLMVNs operate entirely in the complex domain: both weights and activation functions are complex-valued.

This formulation introduces an additional degree of expressiveness and non-linearity, enabling the network to model intricate relationships within the data that real-valued architectures may struggle to capture.

Each neuron in a MLMVN is effectively characterized by two parameters (real and imaginary components), which expands the representational capacity of the model while preserving manageable architectural depth. This richer parametrization allows the network to “bend” the feature space more flexibly, potentially capturing subtle correlations or periodic components that are typical in electrical load signals.

A fundamental motivation behind studying MLMVNs lies in the natural affinity between complex arithmetic and many structures commonly encountered in energy systems. Load profiles, especially at high resolution, present strong periodic components that resonate with Fourier-domain representations. In this sense, reasoning in the complex domain provides a modelling framework that aligns more closely with the intrinsic mathematical nature of the signals, much like complex phasors in AC circuit analysis capture amplitude and phase simultaneously.

Furthermore, the complex domain offers inherent advantages for modelling multi-feature interactions. Phase relationships, correlations across different time scales, and frequency-domain dynamics can be represented in a way that is both compact and expressive. By extending classical architectures such as feed-forward networks, recurrent networks, or sequence-to-sequence models into their complex-valued variants, it becomes possible to explore forecasting performance beyond the limitations of real-valued models.

The New York research activities were therefore structured along two main lines:

- Evaluating different MLMVN architectural configurations and assessing their suitability for short-term load forecasting
- Conducting a series of experiments using complex-valued deep architectures, aimed at quantifying whether the richer representational power of the complex domain offers practical benefits in accuracy, stability, or generalization.

These investigations provided valuable insights into how complex-valued models behave when confronted with highly stochastic residential load data, and how they compare to conventional predictors within the broader forecasting framework.

4.5.1 CVNN Proposed architecture

The architecture developed during the research activity is a complex-valued feed-forward neural network (CVNN) composed of two hidden layers written in matlab code. The first hidden layer contains 75 complex parameters, while the second includes 100 complex parameters, resulting in a compact yet expressive structure suitable for experimentation. Such architecture is displayed in Figure 4.11.

The objective of the model is to forecast the active power consumption for the next three days, using as input the consumption measured during the previous three days, with an hourly sampling time. So, the number of Inputs N_I and the Number of Outputs N_O are taken equally at 72.




Figure 4.11: CVNN used architecture

The architectural choice follows an intentional design strategy: the first hidden layer performs a compression of the input data, projecting the high-dimensional time series into a compact complex subspace; the second hidden layer expands this representation to reconstruct the temporal structure required for multi-day forecasting. This compression–expansion mechanism proved to be effective in extracting relevant patterns while keeping the overall complexity manageable.

Parameter tuning, such as the number of neurons, the choice of activation function, and the learning rate, was carried out empirically through iterative experimentation, given the lack of consolidated guidelines for complex-valued architectures applied to load forecasting.

Each layer employs a complex logarithmic-normalisation activation function, applied to the modulus of the complex-valued signals flowing through the network. This choice stabilises the magnitude of activations, mitigates gradient explosion, and preserves the phase information that is essential when operating in the complex domain. Such activation functions are commonly used in CVNNs because they balance non-linearity with the need to maintain meaningful geometric transformations in the complex plane.

Training is performed using supervised learning with early stopping based on a validation set. A fixed 5% portion of the training dataset is always kept

and used as validation to determine convergence and prevent overfitting. Training is halted when the validation loss fails to improve for a predefined number of epochs, ensuring that the model does not over-adapt to noise.

This architecture, although relatively simple in terms of depth and parameter count, enables the exploration of the unique representational capabilities of complex-valued neural networks and sets the foundation for the more advanced experiments described in the following sections.

4.5.2 Evaluated Data Processing

During the experimentation phase, several preprocessing strategies were assessed to determine their influence on the performance and stability of the CVNN. Domestic load data are inherently noisy, highly irregular and strongly influenced by daily and weekly periodicities. For this reason, preprocessing plays a critical role in shaping the input representation before it is passed to the CVNN.

The following subsections describe the different strategies explored, each designed to enhance the compatibility of the data with complex-domain operations and to facilitate the learning dynamics of the network.

A notable advantage of a custom train code is the ability to visualise how information propagates through the layers of the CVNN. Because the network operates in the complex domain, the transformation applied to the samples can be observed as rotations, distortions, and twists within the complex plane.

This offers valuable insight into the internal behaviour of the model. For every normalisation approach considered, the evolution of the training sequences across the network layers will be shown. In the resulting graphs, each “plane” corresponds to one sequence of the training set, enabling a direct visualisation of how CVNN manipulates the data in the complex domain.

Unit Circle Normalization

The first normalisation strategy consists of encoding the active power information into the phase of complex numbers. In Figure 4.12 a graphical representation of this approach is shown. The preprocessing pipeline involves:

1. applying a logarithmic compression to the load values, in order to reduce the dynamic range of the signal;

$$S_{log} = \log(S_{Input} + 1) \quad (4.1)$$

2. mapping the resulting values to a sector of the unit circle, specifically within the interval $[0.35\pi, 0.6\pi]$ radians. This angular interval was selected empirically to ensure a good balance between resolution and phase separability.

$$S_{UC} = S_{log} * \frac{(\vartheta^{Max} - \vartheta^{min})}{S_{log}^{MaxTrain} - S_{log}^{minTrain}} + \vartheta^{min} \quad (4.2)$$




Figure 4.12: Representation of Unit Circle Normalization

Figure 4.13 and Figure 4.14 show how the data are processed by the network. In this case, the transformation experienced by the data as they propagate through the network layers is relatively simple. At a macroscopic level, the samples preserve their angular-sector shape after each layer and continue rotating along the unit circle.

Since in this normalisation, all information is encoded exclusively in the phase, the geometric evolution remains constrained to that sector. The output of the network eventually settles in the same angular region as the input, confirming that the CVNN preserves and reshapes the phase structure.




Figure 4.13: Unit circle preprocessing train data transformation: 3D representation




Figure 4.14: Unit circle preprocessing train data transformation: planar representation

Normalization via Fourier Transform

A second preprocessing approach consists of transforming the input time series using the Fast Fourier Transform (FFT), showed in Figure 4.15. Working in the frequency domain naturally aligns with the mathematical structure of CVNNs, given their strong affinity for complex arithmetic and harmonic components. The methodology is simple:

1. apply the logarithmic compression
2. compute the FFT of the load profile;
3. normalise the spectrum by its maximum magnitude;
4. pass the resulting complex vector directly to the CVNN.

As the information is already expressed in amplitude and phase, the format is inherently compatible with complex-valued operations.




Figure 4.15: Example of the Fourier preprocessing

Figure 4.16 and Figure 4.17 show how, in this configuration, the geometric transformation applied by the network becomes more significant: the FFT spectrum initially appears as a “pillar” in the complex plane, then the activations collapse into a compact angular sector in the unit circle after both hidden layers, before expanding again toward the final output. This behaviour reflects how the network contracts and reorganises frequency components during processing.




Figure 4.16: FFT preprocessing train data transformation: 3D representation




Figure 4.17: FFT preprocessing train data transformation: planar representation

Simplified FFT Normalization

A variation of the previous method exploits two well-known properties of the Fourier Transform:

- the DC component (first coefficient) can be removed and later reintroduced after prediction
- for real-valued time-domain signals, the Fourier spectrum exhibits conjugate symmetry, meaning that only half of the coefficients carry unique information.

, showed in Figure 4.15. Following these principles, this method uses:

1. logarithmic compression
2. FFT computation
3. removal of the DC component from the input
4. retention of only half of the spectrum
5. re-mirroring and reinsertion of the DC term in post-processing to reconstruct the full signal

The effect on the data is showed in Figure 4.18.




Figure 4.18: Example of the Simplified Fourier preprocessing

This reduces the input dimensionality and computational burden, while preserving all relevant frequency information. Geometrically, Figure 4.19 and Figure 4.20 allow to see how the transformations across the CVNN layers

resemble those obtained in the full FFT method, though the points are less tightly concentrated around the unit circle due to the reduced spectral content.




Figure 4.19: Simplified FFT preprocessing train data transformation: 3D representation




Figure 4.20: Simplified FFT preprocessing train data transformation: planar representation

3.5.2.4 Encoding the Day of the Week

To enrich the input and strengthen the representation of the seasonality of the week, the day of the week was also integrated into the complex-valued input, as showed in Figure 4.21. Complex numbers are naturally suited to encode two-dimensional information, one variable in magnitude and the other in the phase.

The procedure is as follows:

1. the active power is mapped to the modulus of the complex number, using logarithmic compression to linearise its distribution
2. the day of the week, encoded as an integer from 1 to 7, is mapped to the phase

This angular encoding is particularly effective due the phase is a circular variable: the mapping covers the full 2π range, ensuring smooth transitions between days (from Sunday to Monday).




Figure 4.21: Example of the processing with the day feature

Figure 4.22 and Figure 4.23 show how, in this configuration, the input dataset presents a seven-point star shape when represented in three dimensions. After passing through the first and second layers, this structure is transformed into a spiral, reflecting how the CVNN manipulates the phase–magnitude coordinates. At the output, the seven-point star re-emerges, demonstrating the network’s ability to reconstruct the encoded weekly periodicity.




Figure 4.22: Day feature preprocessing train data transformation: 3D representation




Figure 4.23: Day feature preprocessing train data transformation: planar representation

4.5.3 Comparison with persistency and LSTM

To properly evaluate the performance of the CVNN architectures developed, two benchmark models were identified. The first is a seasonal persistence predictor, selected as the baseline to surpass. The second is a Long Short-Term Memory (LSTM) network, used as the gold standard to match or approach.

The seasonal persistence baseline was defined on a weekly horizon: the predictor assumes that the future load value at each time step will be equal to the value observed exactly one week earlier. This choice reflects the strong weekly periodicity often present in residential consumption patterns and provides a simple, yet surprisingly competitive, reference point.

The LSTM network was designed to take advantage of both active power measurements and the day-of-week index, allowing the model to internalise weekly seasonality. The day-of-week variable was transformed using its sine and cosine components, producing a smooth circular encoding and avoiding discontinuities between consecutive days, like from Sunday to Monday. Active power was logarithmically normalised, as done in the CVNN configurations, to compress the dynamic range.

Figure 4.24 show the four layer LSTM architecture, implemented in MATLAB, used in this study.




Figure 4.24: Implemented LSTM Architecture

The number of hidden units in the LSTM layer was set at 40, following empirical tuning aimed at balancing representational capacity and training stability.

Defining as $\hat{y}_h$ the response of the network in an input sequence x_h and y_h as the expected ground truth, the performances of all predictors, CVNN variants, the LSTM model and seasonal persistence were evaluated on the test set of N_{Test} samples using two complementary metrics:

- RMSE, which quantifies the magnitude of the prediction error across the forecast horizon. The metric has been evaluated for each step of the forecasting horizon, showing how the error vary as the predicted points get closer to the present. It has been evaluated for each step of the sequence by applying the following formulation.

$$RMSE_h = \sqrt{\frac{1}{N_{Test}} \sum_{i=1}^{N_{Test}} (y_{h,i} - \hat{y}_{h,i})^2} \quad (4.3)$$

- ERR(Error Reduction Ratio), which evaluates the ability of the predictor to adapt its forecast trajectory as new data become available. denote the number of steps between the moment the forecast was generated and the present time. To quantify how effectively the predictor reduces the absolute prediction error as the target samples in the sequence y_h become closer to the present, the following formulation is adopted:

$$y_h^{ERR} = |y_h - \hat{y}_{h,T}| \quad (4.4)$$

$$ERR_T = \frac{1}{N_{Samples}} \sum_{k=1}^{N_{Samples}} (y_{h,k}^{ERR_1} - y_{h,k}^{ERR_T}) \quad (4.5)$$

The following preliminary results shown in Figure 4.25 compare the performance of all proposed CVNN configurations against the two benchmark models. The evaluation was carried out on a load profile selected from the Ausgrid dataset. Each network was trained on three months of data and tested on the following twenty days.

The performance obtained is comparable to that of the LSTM, both in terms of RMSE and ERR, while relying on significantly simpler architectures. A smaller RMSE indicates a smaller magnitude of the prediction

error, a higher ERR reflects a stronger ability of the network to correct its forecast trajectory as new data become available. This indicates that complex-valued representations may offer expressive modelling capabilities without the computational overhead typically associated with recurrent networks.

Although these results are promising, further investigation is required to determine whether this behaviour is consistent across different users, datasets, and time horizons. Future research will focus on validating the robustness and generality of these findings to confirm whether this is a reliable assumption.




Figure 4.25: RMSE and ERR comparison between the proposed architectures and the LSTM

4.6 Load Profiles Stochastic Generation

Short-term load forecasting, especially on the scale of individual households, is a challenging task. While numerous forecasting techniques have been proposed, achieving high accuracy remains difficult due to the strong stochasticity of user behaviour, hardware limitations, and the scarcity of high-resolution data. In the context of Energy Management Systems (EMS), this uncertainty is not a secondary issue: it directly impacts the quality of control actions.

When the goal is to implement a robust optimisation strategy, relying on a single forecast is fundamentally insufficient. Robust planning requires the ability to anticipate not only the most likely evolution of the load but also a set of plausible alternatives. For this reason, every EMS iteration must operate on a flock of future load scenarios, representing different but credible

realisations of the user’s behaviour.

These synthetic trajectories must satisfy two core requirements: they must be consistent with the user’s historical consumption patterns and cover a wide spectrum of possible future variations to effectively support robust decision-making. This need motivates the formulation of the load profile generation problem.

4.6.1 Proposed Methodology

This section addresses these challenges by introducing a stochastic approach to load profile generation specifically designed to enable robust optimisation. The method is conceived with three primary objectives: explainability, ensuring transparency of the generative mechanism; computational efficiency, allowing real-time integration into EMS control cycles; and scenario quality, producing a diverse and realistic flock of future trajectories capable of supporting reliable robust scheduling.

The proposed methodology involves creating a set of probable future profiles by analysing recent historical data for effective optimisation. Specifically, the goal is to identify and leverage the profile distribution to improve forecasting and consequently enhance the system’s performance.

It encompasses two distinct phases: first, estimating the conditional probability distributions that characterise the time series in question, and second, producing profiles from the previously gathered data. The next paragraph expands on these procedures, while Figure 4.30 shows flock profile generation diagram.

The method was validated using the Ausgrid datasets to confirm its effectiveness. It includes 127 profiles of production and consumption from prosumers, each covering a year. Table I describes the dataset’s contents.

Step 1: Conditional probabilities approximation

The first step is to approximate the conditional probabilities to feed the profile generator. The underlying idea is that stochastic patterns govern a load consumption time series.

As example The “mining” of the probabilistic properties of a profile starts from its historical data. The history consists of N available days, each counting 24 hours. Thus, it is possible to define Y , the history length, as follows:

$$Y = N * 24[hours] \quad (4.6)$$

Then, we define knowledge K as the number of past hours used to generate future profiles and the generation horizon G as the number of hours used to generate those profiles. This allows us to calculate how many sequences we need and their length for segmenting the historical data.

$$l_{seq} = K + G[hours] \quad (4.7)$$

$$N_{seq} = Y - l_{seq} + 1 \quad (4.8)$$

The N_{seq} sequences can be grouped on the basis of similarity, using the property of periodicity, leading to 24 distinct groups. Each group will count N_{seq}^H sequences and correspond to a specific hour $H \in [1, 24]$. This procedure is shown in Figure 4.26.




Figure 4.26: Sequences definition in the Dataset and split in H groups.

Histograms can be created to approximate the distribution of values for each group H at each hour. Each histogram consists of a defined number of bins, N_{bins} , with the first and last centers corresponding to the minimum and maximum values observed during that hour, as shown in Figure 4.27. The bin edges are set to ensure equal widths for all bins. For the final G hours, the expected values for each bin are determined by means of the two edges.

The next step is to investigate the relationship between each of K 's past time steps and the target hour. Hence, the correlation and influence within the past and the future bins are analysed. Also, to make the understanding easier, let us consider $G = 1$ from this point on.





Figure 4.27: Histogram bins definition.

Figure 4.28: Correspondence example between histograms at a past step k and a g step, used as visual representation of conditioning of the values in the step g respect to the value assumed in the step k .

The correlations $Corr_{(g|k)}^H$, with $k \in [1, K]$ and $g \in [1, G]$ are computed and stored in an array. The correspondence between the bins of each of the k past steps and the bins of the landing g is analysed. A visual representation of the conditional pattern is represented in Figure 4.28. The colors of the bars of the landing histogram, the one on the left, are used to represent the dependence to the bars of the departing histogram, the one on the right.

At this point, we have the empirical absolute distribution. To obtain an approximation of the conditional probability $P_{(g|k)}^H$, the probability of the value assumed at step g , knowing the value assumed at step k , we must normalise these results by counts. An example is shown in Figure 4.29.

Step 2: Flock profile generation

After determining the last K known past values of a load consumption time series, labelled X_{past} , we can integrate these into the generation of a *flock*




Figure 4.29: Passage from an absolute relationship to a conditional landing probability. The values in the example are taken from Figure 4.28.

of F possible and plausible future trajectories. Each of these profiles extends into the future for a horizon L steps long. The complete procedure is described in the flow chart in Figure 4.30

The elements of each possible future profile are generated step by step. For each $l \in [1, L]$, starting from the first, the corresponding group H is identified. Then, the 10% of the past steps more correlated to this position are selected and defined $K_{10\%}$. For each of the $Nk_{10\%}$ steps, a candidate $X_{land}^{(l,k)}$ is randomly generated using the corresponding $P_{g|k}^H$ by sampling a bin and assigning its associated expected value.

$$X_{land}^{l,k} \sim Discrete(P_{g|k}^H(\cdot | X_{past}(k))) \quad \forall k \in K_{10\%} \quad (4.9)$$

Finally, the value X_{land}^l is obtained as a weighted average of $X_{land}^{l,k}$, using correlation values $Corr_{g|k}^H$ as weights.

$$X_{land}^l = \frac{1}{Nk_{10\%}} * \sum_{k \in K_{10\%}} (X_{land}^{l,k} * Corr_{g|k}^H) \quad (4.10)$$

To evaluate the probability of occurrence of X_{land}^l , the conditional probability is reversed using $P_{g|k}^H$.

The probability P_{land}^l is assigned based on the bin to which X_{land}^l belongs within $P_{g|k}^H$. The generated bin is considered “impossible” if the associated probability is zero. In that case, X_{land}^l is replaced with the value corresponding to the bin with the highest conditioned probability of the most recent step k and is replaced with this value.

The procedure is iteratively repeated by removing the oldest sample and appending the newly generated one as the most recent observation. At the




Figure 4.30: Step 2 detailed flowchart

end of the process, what is left is a newly generated profile X_{land}^f of L steps, produced based on the original known history.

Once the profile has been generated, an overall probability of occurrence can be associated by concatenating all the conditional probabilities of each generated step.

$$P_{land}^f = P_{l=1|K}^H * \prod_{l \in [1, L-1]} P_{l+1|l}^H \quad (4.11)$$

The procedure is repeated for each of the $f \in [1, F]$ desired profiles, resulting in a flock of F distinct and plausible future scenarios, each coupled with its own probability of occurrence.

Our methodology introduces two exceptions to the procedure. The first exception, explained in Figure 4.31, addresses the arrival of a new input value that cannot be assigned to any admissible bin, making it impossible to define the conditional probability $P_{g|K}^H$ for the subsequent step.




Figure 4.31: Exception 1 flowchart explanation

As a bin can be seen as a cluster, it is reasonable to consider the value part of the closest bin_{j^*} in terms of edge and utilising the associated probability of that bin.

$$Let bin_{j^*} = \arg \min_j |edges(bin_j) - X_{past}(K)| \quad (4.12)$$

$$X_{land}^{l,K} \sim Discrete(P(\cdot|X_{past}(K))) = P_{g|Kj^*}^H \quad (4.13)$$

Limitations of Step 2

The second issue arises when a generated value X_{land}^l falls into a bin that is not accessible from the bin of the previous step $X_{past}(K)$. In such instances, we solve the problem by selecting the bin with the highest probability of occurrence relative to the preceding bin. the solution is illustrated in Figure 4.32.

$$P_{g|K}^H(X_{land}^l | X_{past}(K)) = 0 \quad (4.14)$$

$$X_{land}^l = E[bin(j^*)] with j^* = \arg \max_j P_{g|K}^H(X_{land}^l | X_{past}(K)) \quad (4.15)$$

4.6.2 Results and Discussion

The procedure was condensed into a block that, based on a known past, allows for generating a flock of possible future profiles. This tool can be used for the implementation of advanced optimisation techniques to handle




Figure 4.32: Exception 2 flowchart explanation

uncertainty in forecasts. The following results were obtained by testing the procedure on MATLAB R2024a, running on a machine with the following characteristics:

- Processor, Intel Core i9-139000K
- RAM, 32 GB.

One of the advantages of this approach is the simplicity in extracting information from the time series. The approximation of the conditional probability is a fast procedure, and its computational cost is negligible compared to training a neural network. This aspect becomes critical in online time-series applications, where brand new data continues to flow and must be continuously incorporated by updating the information available.

In the case of household load-consumption especially, it is essential to consider a high variability in user behaviour as the time passes. That means that such instruments must be updated frequently to address these changes.

The time required for the information “mining” depends on the size of the dataset representing the user’s history to date, from which the stochastic information will be extracted. That size will now be addressed as history depth. In all tests so far, this time remained in the 0.5-0.6 s window: 0.5 s for a history depth of 30 days and about 0.6 s for a history depth of 180 days.

The generation time cost shows an almost linear dependency on the number of full profiles generated at every iteration. In Figure 4.33 an example of the test performed is shown. 200 profiles of 72 steps ahead have been generated, using as input the most recent 72 past, taking 8.1 ms.

The memory cost to store this information manifests a quadratic dependence from N_{bins} and a linear dependence on K , H and N_{Byte} . The latter




Figure 4.33: Example of generation of the flock of profiles, whit the intensity of the blue that represents the density of the generated profiles.

represents the size of the chosen variable type. For this work it was considered a $K = 72$ hours, $N_{bins} = 10$ and the chosen variable type (or single in Matlab), with a weight $N_{Byte} = 4$ bytes. The resulting memory occupancy will be ~ 200 KB, potentially allowing implementation of this mechanism on a microcontroller.

$$Memory_{cost} \cong 2 * (N_{bins})^2 * Nk_{10\%} * H * N_{Byte} \quad (4.16)$$

$$with Nk_{10\%} \approx K/10$$

To Conclude, the stochastic approach to profile generation proves to be a valuable alternative, especially when computational resources are limited. The methodology is, in fact, designed for ease of implementation and does not necessitate specialised hardware, such as a GPU

By avoiding any explicit training phase, the method achieves significant computational cost savings while delivering performance comparable to more complex techniques.

As outlined, this method was designed for integration within a robust optimisation framework and will be extended in future work to enable real-time, robust planning for domestic prosumers.

4.7 Conclusion and Final Remarks

This chapter has addressed the problem of forecasting consumption and generation profiles from a system-oriented perspective, with the explicit objective of supporting Energy Management Systems operating under uncertainty and computational constraints.

Different forecasting paradigms were analysed, ranging from low-complexity feed-forward neural networks to more expressive complex-valued neural architectures, highlighting how model simplicity and architectural choices strongly affect real-time applicability without necessarily degrading performance.

A comparative analysis shows that lightweight predictors can achieve accuracy comparable to state-of-the-art recurrent models while offering substantial advantages in training speed and scalability.

In parallel, the investigation of complex-valued neural networks demonstrated how phase-magnitude representations naturally align with the periodic structure of electrical load signals, opening interesting directions for compact and expressive modelling. However, forecasting accuracy alone is insufficient when the ultimate goal is operational decision-making.

For this reason, the chapter introduced a stochastic load profile generation framework, capable of producing a flock of plausible future scenarios together with associated probabilities. This approach enables the explicit handling of uncertainty and directly supports robust optimisation strategies, which are essential in real-world residential energy systems.

These forecasting and generation tools form the foundation upon which higher-level control strategies can be built. In the next chapter, they are integrated into a Home Energy Management System (HEMS), where forecasts and stochastic scenarios are exploited to coordinate flexible resources, manage uncertainty, and enable intelligent, user-orientated energy control in domestic environments.

Chapter 5

Primary-Level Energy Management (HEMS)

*This work presents a comprehensive Home Energy Management System (HEMS) for residential prosumers equipped with photovoltaic generation and battery storage. A rule-based approach is proposed for battery scheduling, enabling high computational efficiency while preserving near-optimal economic performance with experimental validation on large real-world datasets. The proposed architecture also includes strategies for controllable load scheduling and overlapping price scenarios. To address forecasting uncertainty, the methodology integrates both knowledge-based countermeasures and robust optimization. Preliminary results show that the hybrid approach significantly preserve economic outcomes while remaining computationally lightweight, making it suitable for large-scale deployment in energy communities and smart grid environments.*¹

¹Part of this chapter has been published as *A Computationally Efficient Rule-Based Scheduling Algorithm for Battery Energy Storage Systems* [73], as *A Computationally Efficient Rule-Based BMS Countermeasures for Price Overlapping Issues* [84] and as *Application of Control Algorithms for Battery Scheduling in Grid-Connected Energy Prosumers* [85].

5.1 Introduction

This chapter presents a rule-based control strategy for the Home Energy Management System (HEMS) of a prosumer connected to a low-voltage distribution network. The main objective is to propose a computationally efficient algorithm capable of managing energy flows between the distribution network and a prosumer equipped with a photovoltaic (PV) energy production system. The goal of the EMS is to maximize the prosumer's economic revenue by optimizing the use, storage, sale and purchase of PV energy based on electricity market information and daily production/consumption curves.

To achieve this goal, the method proposed in this paper consists of developing a rule-based algorithm that manages the prosumer's BESS and controllable loads. The rule-based approach in this type of problem allows for the reduction of computational costs, which is of fundamental importance in contexts where many users will be coordinated simultaneously. This means that the EMS presented in this work could play a vital role in the context of RECs as well.

From a general point of view, the method requires an algorithm to process the load and generation profiles of the prosumer for the following three days, together with the hourly price curve. The output is a plan for battery and load scheduling for the forecasted horizon, which is updated every hour.

The algorithm is validated in terms of economic performance achieved and computational times on two experimental datasets with different scenarios characterized by real productions and loads of prosumers for over a year. The annual economic results of the proposed rule-based approach has then been compared with a linear programming optimization algorithm used as benchmark. The comparison highlights similar performance in terms of economic revenue, but the rule-based approach guarantees 30 times lower processing time.

5.2 Notation and Preliminary Hypotheses

5.2.1 Energy Profile

Considering an electricity user equipped with a PV system, the power exchange with the electrical grid can be represented, without losing generality, as an alternation of injection intervals and withdrawal intervals, as shown in the example in Figure 5.2).




Figure 5.1: HEMS Architecture

In the following, we refer to the injection periods (green areas in Figure 5.2) as positive intervals and the withdrawal periods (red areas in Figure 5.2) as negative intervals. Let T_S be the control time step in hours, i.e., the time granularity of the actions performed by the EMS, and N_H be the control future horizon, i.e., the number of future actions scheduled by the EMS at each algorithm execution.

Thus, the system's future time horizon is $T_H = N_H T_S$. In this study, $T_S = 1$ h and $N_H = 72$, yielding $T_H = 72$ h; however, the proposed energy management approach is generalizable to any T_S and N_H values. Since T_S is equal to one hour, we consider the hourly energy exchange with the power grid.

As an example, Figure 5.3 shows the hourly energy exchange corresponding to the power curve in Figure 5.2. The hourly energy exchange is, therefore, represented as an array $\bar{E}$ of length N_H , with values E_t being the energy exchanged in the t -th hour, $t = 1, \dots, N_H$.

The energy exchange sequence $\bar{E}$ is also composed of alternating positive and negative intervals. We denote with N_I the total number of positive and negative intervals, and we indicate the intervals as I_j , with $j = 1, \dots, N_I$. Moreover, we denote with I_H the union of all the intervals:

$$I_H = \bigcup_{j=1}^{N_I} I_j \quad (5.1)$$

For each interval, we define:

- the start index $t_{start}^{(j)}$, with $j = 1, \dots, N_I$, as the position of the first element of the j -th interval in the array $\bar{E}$




Figure 5.2: Example of 24 h profile representing the power exchange with the electrical grid for a domestic prosumer. The green areas represent power injections from the user into the grid (positive intervals), while in the red periods, the user takes power from the grid (negative intervals).

- the stop index $t_{stop}^{(j)}$, with $j = 1, \dots, N_I$, as the position of the last element of the j -th interval in the array $\bar{E}$
- the interval energy $E_{tot,j}$, with $j = 1, \dots, N_I$, as the total energy exchanged in the j -th interval. $E_{tot,j}$ is computed as

$$E_{tot,j} = \sum_{t=t_{start}^{(j)}}^{t_{stop}^{(j)}} E_t \quad (5.2)$$

Conventionally, energy is considered positive when injected into the grid and negative when withdrawn from the grid. Note that $t_{start}^{(1)}$ and $t_{stop}^{(N_I)}$ correspond respectively to the array indices 1 and N_H . Using the introduced notation, each generic sequence $\bar{E}$ of N_H samples can be mapped in a matrix containing the $t_{start}^{(j)}$ and $t_{stop}^{(j)}$ indices, and the interval energies $E_{tot,j}$. Table. 5.1 shows the mapping matrix obtained for the example sequence shown in Figure 5.3)). In this case, $N_I = 6$, and therefore, the mapping matrix has 18 elements.

5.2.2 Economic Price Profiles

As the EMS objective is to maximize the economic benefits based on the trend of the electricity market, we introduce the following important hy-




Figure 5.3: Hourly energy exchange with the electrical grid for a domestic prosumer (corresponding to the example in Figure 1). Each bar represents the energy injected into (or taken from) the grid during that hour.

Table 5.1: Mapping matrix resulting from the sequence in Figure 5.3)).

	$j = 1$	$j = 2$	$j = 3$	$j = 4$	$j = 5$	$j = 6$
$t_{\text{start}}^{(j)}$	1	5	8	13	17	20
$t_{\text{stop}}^{(j)}$	4	7	12	16	19	24
$E_{\text{tot},j}$ (kWh)	2.544	-1.336	2.547	-1.273	2.289	-1.909

pothesis: in each considered control horizon of N_H actions, the minimum purchase price is higher than the maximum selling price. Such condition can be written as

$$\min_{t \in I_H} C_t^{(p)} > \max_{t \in I_H} C_t^{(s)} \quad (5.3)$$

where $C_t^{(p)}$ and $C_t^{(s)}$ are respectively the purchase price and the selling price, $t = 1, \dots, N_H$. This condition implies that it is always more convenient to store the excess production in the battery if needed for future usage in negative intervals, instead of selling the overproduced energy and buying it back later.

Such an assumption usually holds, even though in recent years, mainly due to the consequences of COVID-19, there have been cases where the high instability of the energy market has led to a violation of this hypothesis.

In this work, we do not take into account such exceptions. To help visualize the condition expressed by Ineq. 5.3, Figure 5.4) and Figure 5.5) show, respectively, a situation where the hypothesis is respected and a situation in which it is violated.

In this study, we refer to the Italian framework and therefore consider the following energy prices, according to the Italian regulation of the energy market and to what was established by the GSE (Gestore dei Servizi Energetici, i.e., the Italian energy services manager):

- The selling price $C_t^{(s)}$ is set equal to the PUN (Prezzo Unico Nazionale).
- The purchase price is calculated as $C_t^{(p)} = 1.21C_t^{(s)} + 0.088/kWh$, where the second constant term accounts for taxes and system charges.

5.2.3 Final Consideration

Before introducing the LP cost function and constraints, some additional definitions are needed: $\bar{E}_0$ is the original sequence representing the energy exchange with the power grid $\bar{E}$ of length N_H , without the contributions from the BESS. The elements of $\bar{E}_0$ are $E_{0,t}$ with $t = 1, \dots, N_H$.

$\bar{e}_B$ is the sequence of actions on the energy storage scheduled for the next N_H time steps, i.e., is an array of length N_H containing the energy exchanged with the storage. The elements of $\bar{e}_B$ are $e_{B,t}$ with $t = 1, \dots, N_H$. Conventionally, $\bar{e}_B$ is positive when the energy is stored in the battery and negative when the energy is drawn from the battery.




Figure 5.4: Example of selling price and purchase price trends respecting condition 5.3.




Figure 5.5: Example of selling price and purchase price trends violating condition 5.3. In this case, the minimum purchase price (red line) is below the maximum selling price (blue line).

$\bar{E}_B$ is the sequence $\bar{E}_0$ considering the contributions of $\bar{e}_B$, as established by the scheduling algorithm. In other terms, the actions $\bar{e}_B$ modify the original energy exchange with the grid vector $\bar{E}_0$, yielding a new energy exchange vector $\bar{E}_B$. The elements of $\bar{E}_B$ are $E_{B,t}$, $t = 1, \dots, N_H$. As such, it holds:

$$E_{B,t} = E_{0,t} - e_{B,t} \quad t = 1, \dots, N_H \quad (5.4)$$

or, in vector notation:

$$\bar{E}_B = \bar{E}_0 - \bar{e}_B \quad (5.5)$$

Note that, to ease readability, quantities indicated with capital “ E ” refer to energy exchanges with the electrical grid, while quantities indicated with lowercase “ e ” refer to energy exchanges with the BESS.

We denote, respectively, with I_{neg} and I_{pos} the set of negative intervals and the set of positive intervals. I_H is defined as in Equation Eq. 5.1. Therefore, it also holds: $I_H = I_{neg} \cup I_{pos}$. The energy price array $\bar{C}$, with elements C_t , $t = 1, \dots, N_H$, is defined as follows:

$$C_t = \begin{cases} C_t^{(p)} & \text{se } t \in I_{neg} \\ C_t^{(s)} & \text{se } t \in I_{pos} \end{cases} \quad (5.6)$$

Figure 5.6 shows how the $\bar{C}$ array is derived from purchase and selling prices, considering, as an example, the intervals in Figure 5.3.

SOC_j , $j = 1, \dots, N_I$, represents the BESS charge level at the end of the j -th interval; SOC_M is the maximum battery charge level; SOC_m is the minimum battery charge level. As such, the maximum possible energy exchange with the storage at each control time step is $e_{B,full} = SOC_M - SOC_m$.

5.3 Linear Programming Benchmark Definition

Given the premises of the previous section, the LP cost function is formulated as

$$\begin{aligned} \min_{E_{B,t}} & \left(- \sum_{t=1}^{N_H} E_{B,t} C_t \right) \\ \text{subject to} & \quad \gamma_1, \gamma_2 \end{aligned} \quad (5.7)$$



Figure 5.6: Example of derivation of the energy price array $\bar{C}$, given the energy profile in Figure 5.3

γ_1 and γ_2 are the problem constraints, as defined in the following. The first constraint γ_1 establishes the lower and upper boundaries for $\bar{E}_B$ as

$$\min(E_{0,t}, 0) \leq E_{B,t} \leq \max(E_{0,t}, 0) \quad \forall t \in I_H \quad (5.8)$$

Such a constraint implies that the battery can never exchange energy directly with the grid. This entails, at the same time, the impossibility of changing a positive interval into a negative one (and vice versa) by selling or buying additional energy. Figure 5.7 shows the limits defined in 5.8 for the energy profile shown in Figure 5.3.

Secondly, we define, for each interval, a γ_2 constraint as a linear inequality to limit the battery usage so that it always operates under the safe operating conditions (in terms of charge level) recommended by the manufacturer.

As such, γ_2 is a set of linear inequality constraints, one for each interval. This is achieved by ensuring that the battery can never overcharge (over the SOC_M limit) in a positive interval or be depleted completely (under the SOC_m limit) in a negative interval. Therefore, we can express the γ_2 a constraint for the j -th interval as

$$\begin{cases} SOC_j \leq SOC_M & \text{if } I_j \subseteq I_{\text{pos}} \\ SOC_j \geq SOC_m & \text{if } I_j \subseteq I_{\text{neg}} \end{cases} \quad (5.9)$$




Figure 5.7: Visualization of the limits imposed by the constraint 5.8

$$\text{SOC}_j = \text{SOC}_{j-1} + \sum_{t=t_{\text{start}}^{(j)}}^{t_{\text{stop}}^{(j)}} e_{B,t}, \quad j = 1, \dots, N_I \quad (5.10)$$

Note that SOC_j for $j = 0$ represents the battery state of charge at the beginning of the scheduling sequence, as determined by previously performed actions. It is possible to express the generic SOC_j in terms of SOC_0 by recursively expanding Eq. 5.10, yielding:

$$\text{SOC}_j = \text{SOC}_0 + \sum_{t=t_{\text{start}}^{(1)}}^{t_{\text{stop}}^{(j)}} e_{B,t}, \quad j = 1, \dots, N_I \quad (5.11)$$

Rewriting 5.9 using 5.11 and also introducing the relationship given by 5.4, we obtain

$$\begin{cases} \text{SOC}_0 + \sum_{t=t_{\text{start}}^{(1)}}^{t_{\text{stop}}^{(j)}} E_{0,t} - \sum_{t=t_{\text{start}}^{(1)}}^{t_{\text{stop}}^{(j)}} E_{B,t} \leq \text{SOC}_M & \text{if } I_j \subseteq I_{\text{pos}} \\ \text{SOC}_0 + \sum_{t=t_{\text{start}}^{(1)}}^{t_{\text{stop}}^{(j)}} E_{0,t} - \sum_{t=t_{\text{start}}^{(1)}}^{t_{\text{stop}}^{(j)}} E_{B,t} \geq \text{SOC}_m & \text{if } I_j \subseteq I_{\text{neg}} \end{cases} \quad j = 1, \dots, N_I \quad (5.12)$$

Rearranging the terms in 5.12 yields the following expression for the γ_2 constraints:

$$\begin{cases} \sum_{t=t_{\text{start}}^{(1)}}^{t_{\text{stop}}^{(j)}} E_{B,t} \geq \text{SOC}_0 - \text{SOC}_M + \sum_{t=t_{\text{start}}^{(1)}}^{t_{\text{stop}}^{(j)}} E_{0,t} & \text{if } I_j \subseteq I_{\text{pos}} \\ \sum_{t=t_{\text{start}}^{(1)}}^{t_{\text{stop}}^{(j)}} E_{B,t} \leq \text{SOC}_0 - \text{SOC}_m + \sum_{t=t_{\text{start}}^{(1)}}^{t_{\text{stop}}^{(j)}} E_{0,t} & \text{if } I_j \subseteq I_{\text{neg}} \end{cases} \quad j = 1, \dots, N_I \quad (5.13)$$

5.4 Rule-Based Algorithm Definition

In this section, the proposed rule-based approach is introduced. The algorithm is formulated to achieve the same objective expressed by the LP cost function 5.7, i.e., the maximization of the prosumer revenue derived from the management of the battery. The implemented rule-based algorithm consists of the following three main steps.

5.4.1 Step 1

Determination of the energy possibly requested from the battery and the energy possibly available for storage in the battery in each interval I_j . These quantities, denoted respectively as $\bar{e}_R$ and $\bar{e}_A$, are vectors of size N_I , with elements $e_{R,j}$ and $e_{A,j}$, computed as follows:

$$e_{R,j} = \begin{cases} 0 & \text{if } I_j \subseteq I_{\text{pos}} \\ \max(E_{\text{tot},j}, -e_{B,\text{full}}) & \text{if } I_j \subseteq I_{\text{neg}} \end{cases} \quad j = 1, \dots, N_I \quad (5.14)$$

$$e_{A,j} = \begin{cases} \min(E_{\text{tot},j}, e_{B,\text{full}}) & \text{if } I_j \subseteq I_{\text{pos}} \\ 0 & \text{if } I_j \subseteq I_{\text{neg}} \end{cases} \quad j = 1, \dots, N_I \quad (5.15)$$

In each interval the energy $E_{\text{tot},j}$ represent the total energy exchanged between the user and the grid.

$$E_{\text{tot},j} = \sum_{t=t_{\text{start}}^{(j)}}^{t_{\text{stop}}^{(j)}} E_t \quad (5.16)$$

In a positive interval, the energy request from the battery must be zero, whereas the available energy surplus potentially storable in the battery is equal to the quantity $E_{\text{tot},j}$. The upper limit of $E_{\text{tot},j}$ is the full battery capacity $e_{B,\text{full}}$.

Similarly, in a negative interval, the energy request from the battery equals $E_{\text{tot},j}$, limited below by $-e_{B,\text{full}}$, while the available energy is zero.

Note that, consistently with the convention adopted for $e_{B,t}$, the quantity $e_{R,j}$ is always nonpositive, since it represents energy that may be drawn from the battery, whereas $e_{A,j}$ is nonnegative, indicating energy that may be stored in the battery.



Figure 5.8: Flowchart of the first step of the rule-based algorithm. The left and right branches illustrate, respectively, the procedure to determine $\bar{e}_R$ and $\bar{e}_A$.

$$e_{R/A,j} = e_{R,j} + e_{A,j} = \begin{cases} \min(E_{\text{tot},j}, e_{B,\text{full}}) & \text{if } I_j \subseteq I_{\text{pos}} \\ \max(E_{\text{tot},j}, -e_{B,\text{full}}) & \text{if } I_j \subseteq I_{\text{neg}} \end{cases} \quad j = 1, \dots, N_I \quad (5.17)$$

5.4.2 Step 2

Determination of the target energy amount that the system needs to move from or to the battery in each interval. We denote this quantity as $\bar{e}_{\text{target}}$, with elements $e_{\text{target},j}$ for $j = 1, \dots, N_I$.

We also introduce the vector $\bar{e}_R^{(\text{next})}$, with elements $e_{R,j}^{(\text{next})}$, containing, for each interval, the cumulative unsatisfied energy request remaining from all future intervals. The calculation proceeds backwards, starting from the last interval I_{N_I} , as illustrated in the flowchart in Figure 5.9. In this way, the algorithm correctly computes $\bar{e}_R^{(\text{next})}$ for all previous intervals.

For the last interval, we set

$$e_{R,N_I}^{(\text{next})} = 0$$

$$e_{\text{target},N_I} = \begin{cases} \max(e_{R/A,N_I}, E_{\text{pos}}^{(\text{avg})}) & \text{if } I_{N_I} \subseteq I_{\text{pos}} \\ \min(e_{R/A,N_I}, E_{\text{neg}}^{(\text{avg})}) & \text{if } I_{N_I} \subseteq I_{\text{neg}} \end{cases} \quad (5.18)$$

- $e_{R,N_I}^{(\text{next})}$, i.e., the unsatisfied energy request from time instants beyond the system's future horizon T_H , is set to zero.
- If $I_{N_I} \subseteq I_{\text{pos}}$, then e_{target,N_I} is forced to be at least $E_{\text{pos}}^{(\text{avg})}$, which represents the average energy surplus observed in positive intervals from historical data. This accounts for the possibility that future positive intervals may extend further, potentially providing additional energy.
- If $I_{N_I} \subseteq I_{\text{neg}}$, then e_{target,N_I} is forced to be at most $E_{\text{neg}}^{(\text{avg})}$, representing the average energy deficiency observed in negative intervals. This assumption similarly accounts for the possibility that the last negative interval may extend further and require more energy.




Figure 5.9: Flowchart of the second step of the rule-based algorithm.

For the other intervals, $e_{\text{target},j}$ and $e_{R,j}^{(\text{next})}$ are computed as follows:

$$e_{R,j}^{(\text{next})} = \begin{cases} \max\left(e_{R,j+1}^{(\text{next})} + e_{\text{target},j+1}, -e_{B,\text{full}}\right) & \text{if } I_j \subseteq I_{\text{pos}} \\ \min\left(e_{R,j+1}^{(\text{next})} + e_{\text{target},j+1}, 0\right) & \text{if } I_j \subseteq I_{\text{neg}} \end{cases} \quad j = 2, \dots, N_I - 1 \quad (5.19)$$

$$e_{\text{target},j} = \begin{cases} \min(e_{R/A,j}, -e_{R,j}^{(\text{next})}) & \text{if } I_j \subseteq I_{\text{pos}} \\ e_{R/A,j} & \text{if } I_j \subseteq I_{\text{neg}} \end{cases} \quad j = 2, \dots, N_I - 1 \quad (5.20)$$

Finally, for the first interval I_1 , the initial state of charge of the battery SOC_0 must be considered:

$$e_{R,1}^{(\text{next})} = \begin{cases} \max(e_{R,2}^{(\text{next})} + e_{\text{target},2}, -e_{B,\text{full}}) & \text{if } I_1 \subseteq I_{\text{pos}} \\ \min(e_{R,2}^{(\text{next})} + e_{\text{target},2}, 0) & \text{if } I_1 \subseteq I_{\text{neg}} \end{cases} \quad (5.21)$$

$$e_{\text{target},1} = \begin{cases} \min(e_{R/A,1}, \max(-e_{R,1} - SOC_0, 0)) & \text{if } I_1 \subseteq I_{\text{pos}} \\ \min(e_{R/A,1}, -SOC_0) & \text{if } I_1 \subseteq I_{\text{neg}} \end{cases} \quad (5.22)$$

5.4.3 Step 3

Determination of the final battery scheduling $\bar{e}_B$ that the system will try to follow. This is achieved by distributing the quantities $e_{\text{target},j}$, computed in Step 2, within each interval. Unlike Step 2, this operation is performed *forward*, starting from the first interval.

In this way, the computation of $\bar{e}_B$ correctly accounts for the fact that, in some intervals, it might be impossible to move the entire energy amount $e_{\text{target},j}$ due to the battery capacity limits.

The criteria adopted for determining the scheduling are as follows.

Positive intervals. In positive intervals, the algorithm schedules the battery recharge by distributing the energy $e_{\text{target},j}$ among the time slots of the interval, starting from those associated with the lowest energy selling price. This process proceeds iteratively until the entire quantity $e_{\text{target},j}$ is transferred to the battery or until the battery becomes fully charged.

Figure 5.10 shows an example of scheduling in a positive interval, assuming a battery capacity of 3 kWh and an initially discharged battery. In this case, the algorithm identifies the third slot as the least convenient for selling energy and therefore gives priority to this slot for charging the battery. The process continues until the desired energy amount is stored in the battery ($e_{\text{target},j} = 2.5$ kWh in this example). Any remaining energy is then sold to the GSE at the price C_t .




Figure 5.10: Example of scheduling in a positive interval.

Negative intervals. In negative intervals, the algorithm schedules the battery discharge by distributing the energy $e_{\text{target},j}$ among the interval time slots, starting from those with the highest energy purchase price. This continues iteratively until the entire amount $e_{\text{target},j}$ is provided by the battery or until the battery becomes completely depleted.

Figure 5.11 illustrates an example of scheduling in a negative interval. In this case, priority is given to the slot with the highest purchase price. At the fourth iteration, the scheduling process ends because the battery becomes fully discharged. Any additional energy requirement must therefore be purchased from the grid at the cost C_t .




Figure 5.11: Example of scheduling in a negative interval.

Flowchart representation. The operations described in this third step are visually represented in the flowcharts reported in Figure 5.12 and 5.13. These flowcharts detail the procedure used to determine the scheduling for positive and negative intervals starting from a given $\bar{e}_{\text{target}}$ vector.

5.5 LP vs. Rule-based

5.5.1 Dataset

The described LP and rule-based approaches are tested on two datasets containing historical production and consumption data from real prosumers. None of the users is equipped with a BESS; therefore, the presence of a storage system is assumed in the simulation. Additionally, the historical trends of the energy market, made available by the GSE in the form of PUN, are used as the reference for the energy price. All data are provided in hourly format.

The datasets used in this study are summarized below:

- **Ausgrid:** this dataset contains one year (8760 samples per user) of production and consumption data from 126 prosumers. In the following, this dataset is indicated as the *AUS* dataset.
- **Planet Smart City:** this dataset contains five months of data (about 3550 samples per user) related to 286 users in the Milan area (Italy). These users are pure consumers; therefore, production data are artificially added by assuming specific PV and BESS sizes (listed below) and by using solar irradiation data from PVGIS. In the following, this dataset is referred to as the *PSC* dataset.

Historical energy price data are available from the GSE website. This study uses PUN values extracted from 2020, stable around 0.05 *Euro/kWh*. Over this period, the hypothesis expressed in Eq. 5.3 always holds.

Furthermore, the datasets are augmented by rescaling the production and consumption profiles of each user to achieve annual consumption levels of 3 MWh or 6 MWh, PV sizes of 3 kW or 6 kW, and BESS capacities of 6 kWh or 12 kWh. These options are combined in all possible ways, yielding eight test scenarios.

In the following, scenarios are indexed by $s = 1, \dots, 8$. Each scenario includes 126 rescaled users from the AUS dataset and 286 rescaled users




Figure 5.12: Flowchart of the third step of the rule-based algorithm. The “NEG Scheduling” and “POS Scheduling” blocks are illustrated in Figure 5.13.



Figure 5.13: Operations performed in the “NEG Scheduling” block (on the left) and “POS Scheduling” block (on the right) of Figure 5.12.

from the PSC dataset. The dataset is denoted by d .

5.5.2 Results

The results of the two scheduling algorithms on the AUS and PSC datasets are reported in this section, comparing computational performance and economic benefits achieved per user. All algorithms are implemented in MATLAB R2024a and executed on a machine equipped with an Intel Core i9-13900K processor and 32 GB RAM.

Both the LP and rule-based scheduling algorithms, implemented within a BMS based on a receding-horizon framework, are tested across the eight scenarios described. The battery scheduling is computed once per hour (i.e., per dataset sample) to determine the actions for the next N_H hours (equal to 72 in this study). Only the first scheduled action is applied, after which the horizon is shifted forward by one hour and a new scheduling solution is computed.

Tables 5.2 and 5.3 summarize the results obtained on the PSC and AUS datasets. In the tables, T_{RB} and T_{LP} represent, respectively, the time required by the rule-based algorithm and the LP algorithm to compute a full 72-hour battery schedule.

Considering the LP schedule as the optimal reference, the quantity $\Delta_{\text{rev},s}$ denotes the average percentage error made by the rule-based algorithm in terms of final economic revenue, for each scenario s . More precisely:

$$\Delta_{\text{rev},[i,s]} = \frac{\text{rev}_{\text{RB},[i,s]}^{(\text{final})} - \text{rev}_{\text{LP},[i,s]}^{(\text{final})}}{\text{rev}_{\text{LP},[i,s]}^{(\text{final})}} \cdot 100 \quad (5.23)$$

$$\Delta_{\text{rev},s} = \frac{1}{N_{\text{user}}} \sum_{i=1}^{N_{\text{user}}} \Delta_{\text{rev},[i,s]} \quad (5.24)$$

where i is the user index, s is the scenario index, and N_{user} is the total number of users. The values $\text{rev}_{\text{RB},[i,s]}^{(\text{final})}$ and $\text{rev}_{\text{LP},[i,s]}^{(\text{final})}$ represent the final economic revenue obtained for user i in scenario s after the complete simulation (8760 h for AUS scenarios, 3550 h for PSC scenarios).

By definition, $\Delta_{\text{rev},[i,s]}$ is negative when the rule-based approach performs worse than the LP benchmark.

Tables 5.2 and 5.3 report the average results for each scenario. Figure 5.14 shows such results in terms of distribution of $\Delta_{\text{rev},[i,s]}$ and the average revenue

trends over time for all scenarios across both datasets, comparing the rule-based and LP algorithms. Note that the percentage difference between the final points of each pair of curves corresponds to $\Delta_{\text{rev},s}$.

Table 5.2: Results on PSC dataset. The values reported are averaged over all users in each scenario.

Load (kW)	PV Size (kW)	BESS Size (kWh)	T_{RB} (ms)	T_{LP} (ms)	$\Delta_{\text{rev},s}$ (%)
3	3	6	~ 20	~ 720	-1.55
3	3	12	~ 20	~ 720	-4.31
3	6	6	~ 20	~ 720	-0.71
3	6	12	~ 20	~ 720	-1.80
6	3	6	~ 20	~ 720	-0.19
6	3	12	~ 20	~ 720	-0.53
6	6	6	~ 20	~ 720	-0.39

Table 5.3: Results on AUS dataset. The values reported are averaged over all users in each scenario.

Load (kW)	PV Size (kW)	BESS Size (kWh)	T_{RB} (ms)	T_{LP} (ms)	$\Delta_{\text{rev},s}$ (%)
3	3	6	~ 20	~ 720	-0.02
3	3	12	~ 20	~ 720	-4.13
3	6	6	~ 20	~ 720	-0.18
3	6	12	~ 20	~ 720	-0.54
6	3	6	~ 20	~ 720	-0.12
6	3	12	~ 20	~ 720	-0.56
6	6	6	~ 20	~ 720	-0.004
6	6	12	~ 20	~ 720	-1.02

These results demonstrate that, although slightly suboptimal, the rule-based algorithm closely approximates the performance of the LP benchmark across all scenarios. However, the rule-based approach requires only a fraction of the computational time needed by LP, making it preferable in applications where computational efficiency is critical.




Figure 5.14: On the left: probability distribution of $\Delta_{\text{rev},[i,s]}$ including both datasets; on the right: comparison of the average revenue achieved by the users in each scenario when using either the LP or the rule-based approach. The two curves are visually indistinguishable because the error introduced by the rule-based algorithm is too small to be appreciated at the scale of the plot.

Information regarding the execution time of the algorithms is particularly relevant from the perspective of the scalability of the management system. As shown in Figure 5.15, when increasing the number of scheduled actions N_H to achieve a finer time granularity in the scheduling the rule-based approach becomes progressively and significantly preferable to LP in terms of computational efficiency.

This is due to the fact that the computational time of the LP optimization does not scale linearly with the length of the scheduled sequence.

Figure 5.15 illustrates this advantage graphically and highlights how it becomes substantial for large values of N_H . For example, for a one-day schedule with a resolution of one minute (1440 time steps), the rule-based algorithm is already more than 100 times faster than the LP-based approach.




Figure 5.15: Comparison between the computational time of the rule-based and the LP algorithms for increasing values of N_H .

5.5.3 Discussion

The proposed EMS aims to maximize the prosumer's economic gains by optimizing the use, storage, and trading of photovoltaic energy according to market prices and daily production and consumption profiles.

A rule-based approach was adopted to achieve low computational costs in determining the battery scheduling strategy, while maintaining economic revenues comparable to those obtained with more classical optimization-based methods.

For this reason, the paper presents a comparison between two different approaches for developing a EMS for a single prosumer: the rule-based strategy designed and implemented and a classical approach based on linear programming.

The performance of the two methods is evaluated under the simplifying assumptions of perfect predictions of the electric load and PV generation over the scheduling window and absence of a battery degradation model. These simplifications were introduced to isolate the evaluation of the scheduling algorithms from external factors such as forecasting uncertainty.

The results can be summarized as follows:

- Both techniques provide extremely similar outcomes in terms of economic revenue, with the LP approach typically yielding slightly better results.
- In terms of computational time, the rule-based algorithm is approximately 30 times faster than its LP counterpart when scheduling 72 future battery actions. This difference becomes even more significant as the number of scheduled actions N_H increases.

These findings indicate that the proposed rule-based algorithm is a viable alternative to more common LP-based approaches for developing a MES for a single prosumer equipped with renewable generation and storage. While economic performance is comparable between the two methods, the rule-based approach clearly outperforms LP optimization in computational efficiency.

This characteristic becomes fundamental when system scalability is required, such as in energy communities where hundreds or thousands of users must be controlled and coordinated simultaneously.

Existing studies have shown that, in such scenarios, commonly used approaches such as MILP rapidly become too complex and ultimately impractical. Rule-based strategies, such as the one proposed in this work, are inherently more suitable for large-scale implementation.

5.6 Price Overlapping Relaxation

The technique presented in the previous section emphasizes cost efficiency and the explainability of the resulting scheduling strategy. However, it does not account for the scenario of *price overlapping*, in which the selling price may temporarily exceed the buying price. By studying the behavior of the

problem, the proposed approach meant to find a solution close to the optimum in a fraction of the time and without recurring to computational expensive techniques.

On the other edge, relying on a list of rules to take decisions is less flexible than use an optimization technique. As the proposed approach was developed under the hypothesis that the selling price is never higher than the buying price.




Figure 5.16: Comparison between 2020 and 2022 Selling prices and variable part of the Buying prices

Such events arise from a combination of the uncontrollability of renewable energy sources and the declining stability of fossil fuel supply chains, both of which increase uncertainty within the energy market.

Historically, fossil fuel-based generation was the main driver of energy prices, but the 2022 energy crisis, triggered by geopolitical instability, disrupted supply chains and significantly raised thermal generation costs. In Figure 5.16 a comparison between the two different type of year is showed. As a consequence, situations occurred in which selling excess energy became more profitable than storing it.




Figure 5.17: Example of prices overlapping. The overlap is highlighted in yellow

This section addresses this limitation by proposing and implementing simple countermeasures to preserve the quality of the cashflow generated by an EMS that relies on the previously described approach.

5.6.1 Implemented Countermeasures

The countermeasure introduced in this section is based on the following concept: if, during a specific hour within a forecasted injection interval, the selling price of energy is higher than the lowest buying price in the subsequent forecasted absorption interval, then that hour is *marked*. Marking an hour implies that, in the presence of any energy surplus, the entire surplus is sold to the grid rather than stored, as it is showed in Figure 5.17.

This procedure has been integrated into the first stages of the approach. Maintaining the structure of the full rule-based methodology previously described, the proposed extensions can be summarized as follows:

1. The presumption $\bar{E}_0$, defined as the difference between the user's forecasted production and consumption profiles, is processed to extract information on injection and absorption intervals, exactly as in the original methodology.
2. For each positive interval, the selling price of every hour within the interval is compared with the lowest buying price of the subsequent negative interval. All hours satisfying this condition are marked.
3. A dummy presumption profile is constructed. It matches the original one except for the marked hours, where the energy value is replaced

with an arbitrarily small positive number.

4. The updated prosumption is processed to recompute the interval information matrix. The original rule-based procedure is then applied up to the scheduling stage of the BESS actions.
5. The resulting prosumption after the BESS actions is obtained as the pointwise difference between the original prosumption and the scheduled sequence of battery actions.

Note that in step 3, in correspondence of the marked hours, it is important to not place the value of the dummy prosumption equal to zero. Not following this guideline will cause a fraction of the positive interval each fraction separated by a new negative interval.

This cause an increase in the number of intervals to consider in the successive steps of the method, causing a severe loss of performance.

An example of the consequences of the proposed counter measures is showed in Figures 5.18 and 5.19, while a detailed flow chart is displayed in Figure 5.20.




Figure 5.18: BESS scheduling comparison without overlapping countermeasures




Figure 5.19: BESS scheduling comparison with overlapping countermeasures




Figure 5.20: Overlapping countermeasures flowchart

To verify the functionality of the proposed method, it was tested on the two datasets previously introduced: *Ausgrid* and *Planet Smart City*.

Regarding energy prices, two years of PUN (Prezzo Unico Nazionale) were used:

- 2020, a year in which no price overlapping is present;
- 2022, characterized by high volatility and episodes where selling electricity becomes more profitable than buying it.

The selling price was assumed equal to the PUN, while the buying price was computed from a realistic retail contract.

The purchasing cost have been modelled following a realistic contract and consists of several components, of which only the variable ones can be influenced by electric load management and BESS scheduling. For this reason, the hourly buying price used in this study considers only the variable components. The cost definitions are as follows:

$$E_{\text{cost}} = \text{PUN}_{\text{avg}}^{\text{month}} + F_{\text{cost}} \quad (5.25)$$

$$\text{Buy}_{\text{cost}} = (E_{\text{cost}} + T_{\text{cost}} + S_{\text{cost}}) \cdot \text{VAT} \quad (5.26)$$

with:

$$\text{PUN}_{\text{avg}}^{\text{month}} = \text{monthly average of the PUN} \quad [\text{Euro/kWh}]$$

$$F_{\text{cost}} = \begin{cases} 0.011 \text{ Euro/kWh}, & \text{Sunday and from 23:00 to 07:00,} \\ 0.0286 \text{ Euro/kWh}, & \text{otherwise} \end{cases}$$

$$T_{\text{cost}} = 0.012 \text{ Euro/kWh}$$

$$S_{\text{cost}} = \begin{cases} 0.038 \text{ Euro/kWh} & \text{if YEAR} = 2020 \\ 0 \text{ Euro/kWh} & \text{if YEAR} = 2022 \end{cases}$$

$$\text{VAT} = 1.1.$$

Here, $\text{PUN}_{\text{avg}}^{\text{month}}$ represents the monthly average PUN, F_{cost} is an additional hourly-dependent tariff component, T_{cost} is the transport cost, S_{cost} is the

system cost (zeroed in 2022 due to policy interventions), and VAT accounts for billing.

5.6.2 Results and Discussion

To validate the effectiveness of the improvements introduced in the proposed approach, both the original and the upgraded versions of the method were tested in a MATLAB-based simulated environment.

All profiles in the two available datasets were rescaled to simulate a contractual load of 3 kW and a photovoltaic system of 3 kW. This scaling was achieved under the assumption that these values correspond to an annual consumption and production of approximately 3 MWh.

The two methods, the original and the upgraded version, were evaluated on all profiles across their respective time spans, first using 2020 prices and then using 2022 prices as references. In parallel, using the same datasets, the performance of two additional benchmarks was collected:

- *Brainless* strategy, which charges and discharges the battery as soon as possible without considering price dynamics, and
- LP-based formulation.

Even in these simulations, the forecasting block of the EMS was modeled as an *Oracle*, while the forecasts assumed to be perfect. A forecasting horizon of 72 steps ahead was adopted, and the optimization was performed once every hour. Figures 5.21 and 5.22 show the resulting trend and distributions on the two dataset.

For completeness Table 5.4 repeats the contents and structure of the datasets used.

Table 5.4: Summary of datasets used for simulation.

Dataset	Type of Profiles	N. of Profiles	Profile Length
Planet Smart	Production	127	8760 hours
City	consumption	127	8760 hours
Ausgrid	Consumption only	287	3528 hours
PVGIS	Irradiation and PV output	1	3528 hours
GME (PUN)	Energy prices	2020 & 2022	8760 hours




Figure 5.21: Ausgrid Dataset average cashflow comparison and final result distribution with PUN 2020(left) and PUN 2022(right)




Figure 5.22: Planet SC Dataset average cashflow comparison and final result distribution with PUN 2020(left) and PUN 2022(right)

Table 5.5: Performance comparison of the original and upgraded methods across datasets.

Dataset & Price	Method	Percent Avg. Result Comp. Cost (%) (ms)
PSC – PUN 2020	No Counters	91.46 0.135
	With Counters	91.46 0.143
PSC – PUN 2022	No Counters	63.51 0.134
	With Counters	81.75 0.143
AUS – PUN 2020	No Counters	92.11 0.148
	With Counters	92.11 0.155
AUS – PUN 2022	No Counters	88.46 0.148
	With Counters	95.73 0.155

Table 5.5 summarizes the performance of the original and upgraded methods across the two datasets. Simulations on the Ausgrid dataset span 8160 hours (340 days), while those on the Planet Smart City dataset cover 3360 hours (140 days). The final cashflow values for each user are expressed as percentages. For every user u in each dataset, the *Brainless* strategy is used as the baseline reference, whereas the LP solution acts as the gold standard, as defined in (3):

$$\text{Result}_{\%}^u = \frac{\text{Result}^u - \text{Result}_{\text{Brainless}}^u}{\text{Result}_{\text{LP}}^u - \text{Result}_{\text{Brainless}}^u}. \quad (5.27)$$

The tests were executed again in MATLAB R2024a on a machine equipped with an Intel Core i9-13900K processor and 32 GB of RAM.

The obtained results show that, in scenarios without price overlapping (e.g., 2020), the performance of the two versions of the method is identical. However, under conditions of extreme volatility, such as in 2022, the upgraded method preserves a significantly higher level of performance with only a negligible increase in computational cost.

To conclude, the instability of the energy market and the growing penetration of renewable energy sources introduce new challenges in energy management and this section extends the rule-based method previously proposed to correctly handle situations of price overlapping, without increasing computational cost or compromising the explainability of the methodology.

The proposed countermeasure is simple to implement and enables the user to optimize the cashflow even when the selling price temporarily exceeds the

buying price with negligible additional computational effort, preserving the speed and efficiency that characterize the rule-based approach.

This makes the upgraded methodology a valuable candidate for computationally constrained or large-scale applications where responsiveness and scalability are essential.

5.7 Controllable Loads

Another important degree of freedom at the single-user level is represented by controllable loads. The ability to shift their activation in time offers a valuable countermeasure to the inherent unpredictability of renewable energy production.

By reallocating consumption to more advantageous periods, these flexible appliances can help compensate for fluctuations in local generation and improve the overall balance between demand and supply.

For this reason, incorporating the control and scheduling of controllable loads becomes a highly relevant enhancement to the HEMS, enabling smarter operation and a more effective exploitation of available renewable resources.

In our modelling of the controllable loads each of them is characterized by:

- a load consumption CL_v , an energy expressed in kWh;
- a window of time in which it can be used during the day CL_w ;
- the frequency of usage a day CL_r .

The management of controllable loads is handled by a dedicated add-on module designed to schedule each device within its allowed activation window. Figures 5.23 show the flowchart of the implemented method. The objective is fulfilling the daily activation requirements of every load and ensuring that the resulting profile respects the maximum hourly grid absorption limit.

The optimisation criterion remains economic, aiming to minimise total daily expenses by exploiting favorable market conditions.

For pure consumers, scheduling decisions depend solely on the Buy_{price} curve, naturally preferring low-price hours for load activation. For prosumers equipped with local generation, the logic changes during injection intervals: the $Sell_{price}$ is used instead, representing the opportunity cost of missing potential revenues.




Figure 5.23: Flow chart of the Controllable Loads Scheduling

Activating a controllable load when energy could otherwise be sold effectively “costs” the user that selling price. This mechanism is integrated into the objective function so that the system selects time slots that reflect the real economic trade-off. Figure 5.24 illustrates an example for a controllable load requiring 1 kWh, with an activation window from 9:00 to 22:00 and a daily requirement of one activation, under a grid-absorption limit of 3 kWh.




Figure 5.24: Example of scheduling of a controllable load characterized by $CLv = 1kWh, CLw = 9 - 22, CLr = 1$ for a user with $Load_{Max} = 3kWh$.

When a prosumer has both controllable loads and a BESS, the scheduling process follows a strict priority: controllable loads are always scheduled before the battery. This ensures that self-consumption is maximized, preventing the storage system from absorbing energy that would be more profitably used to operate local loads. Once the load positions are fixed, the HEMS proceeds to compute the optimal battery schedule.

5.8 Handling Forecasting Uncertainty

The methodology presented so far relies heavily on the assumption of perfect forecasting. Unfortunately, reality hits hard on that front: once forecasting errors inevitably appear these inaccuracies propagate through the control loop and degrade the quality of the resulting optimization.

The rolling-horizon strategy is introduced precisely to mitigate forecasting uncertainty; however, for this approach to be truly effective, it must rely on predictors that not only limit the magnitude of the forecasting errors but are

also capable of promptly self-correcting their estimates as new measurements become available.

In the context of load forecasting, and particularly at the household level, the potential for improving predictive accuracy is inherently limited. Even when augmenting active power measurements with additional contextual parameters, the accuracy cannot be improved beyond a certain threshold due to the highly stochastic and user-driven nature of domestic consumption patterns.

To increase the overall robustness of the architecture, two complementary strategies have been explored. The first is the integration of a set of knowledge-based rules derived from an in-depth analysis of the system's dynamics.

These rules act as safeguards, allowing the controller to react sensibly in situations where the forecasts deviate significantly from reality. The second strategy is the adoption of robust optimization.

As discussed in Chapter 4.6, implementing this approach requires the generation of multiple plausible future trajectories, capturing the uncertainty affecting both load and generation forecasts. The development and evaluation of such stochastic trajectories play a key role in enabling the HEMS to remain effective even under imperfect predictive conditions.

In the following section, the studies conducted so far on this topic are presented. These preliminary investigations form the foundation for a more extensive research effort that will continue in the near future, with the aim of consolidating the approach and developing results suitable for publication.

The long-term objective is to design a forecasting-based energy management framework that remains effective even in the presence of significant uncertainty, an essential requirement for real-world deployment.

To correctly implement these corrective strategies, it is essential to clarify how the EMS commands are interpreted by the Battery Management System (BMS), the control layer responsible for translating the high-level scheduling decisions into the actual charging or discharging actions of the battery. In the scenarios discussed so far, the BMS receives two reference values from the EMS: an upper and a lower threshold for the allowable State of Charge (SOC) during the current time step.

At the beginning of each time step, if the planned action is to charge the BESS, the current SOC acts as the lower bound, while the target SOC provided by the EMS defines the upper bound. Conversely, if the planned

action is to discharge the battery, the roles of the two levels are swapped: the EMS target becomes the lower limit, and the current SOC becomes the upper limit.

Within this bounded interval, the BMS regulates the charging or discharging power according to its internal safety rules, ensuring that the physical constraints of the battery are always respected.

Forecasting uncertainty is addressed by adjusting these SOC thresholds. By expanding or contracting the allowable SOC interval, the EMS provides the BMS with the flexibility needed to react to unexpected deviations in load or generation. A wider interval increases the system's robustness, allowing more freedom to respond to unforeseen conditions, while still preserving the overall goal set by the optimization.

5.8.1 Knowledge Based Countermeasures

One of the main challenges of EMS-level planning is the temporal resolution, typically ranging from several tens of minutes to one hour. Within such intervals, the controller has no opportunity to revise its decision if unexpected events occur. For this reason, it is crucial to complement the day-ahead and intra-day scheduling with emergency mechanisms capable of reacting to unforeseen deviations.

By analysing the dynamics and structural properties of the problem, it is possible to identify a series of “safe” behaviours, decisions that are always preferable regardless of the specific scenario. These behaviours can then be implemented by dynamically adjusting the SOC thresholds sent to the BESS, enabling the battery to operate in a more flexible and resilient manner whenever uncertainty arises. An example of the architecture is showed in Figure 5.25




Figure 5.25: EMS Architecture with knowledge based counter measures

In our study, the following good practices have been identified as essential to improving robustness:

- Corrective Offset on Charging Target During Positive Intervals:** A corrective term is added to the target charging level during positive slots, this offset is derived from the mean forecasting error on the aggregated expected surplus within the interval, yielding a more conservative charging request when predictions are unreliable. During a positive interval $t \in \mathcal{T}^+$, a corrective offset δ_{corr} is added to the scheduled target SOC to account for average forecasting errors in previous negative intervals:

$$SOC_t^{\text{target}} = SOC_t^{\text{sched}} + \delta_{\text{corr}}, \quad \delta_{\text{corr}} = \mu(\varepsilon_t \mid t \in \mathcal{T}^-) \quad (5.28)$$

where $\mu(\cdot)$ denotes the mean prediction error of surplus energy in negative slots.

- Emergency Load Supply During Expensive Negative Intervals:** If the first interval is negative, the EMS identifies the two most expensive buying hours, If the current hour corresponds to one of them, the EMS overrides the planned schedule and allocates all the required power to serve local consumption. This avoids significant economic losses when the real load exceeds the prediction.

If the first time slot t_0 is negative ($t_0 \in \mathcal{T}^-$) and $p_{t_0}^{\text{buy}}$ is among the two highest prices within that interval:

$$p_{t_0}^{\text{buy}} \in \text{Top2}(\{p_t^{\text{buy}} : t \in \mathcal{T}^-\}) \Rightarrow P_{t_0}^{\text{grid}} = \min(P_{t_0}^{\text{load}}, P^{\text{grid}, \text{max}}) \quad (5.29)$$

This allows bypassing the plan to fully cover unexpected loads.

- **Charge Release During Cheapest Positive Hours:** If the first interval is positive, the two cheapest selling hours are identified, Should the current hour coincide with one of these, the charging constraints are relaxed to allow full absorption of any unexpected overproduction, preventing profitable energy from being curtailed.

If $t_0 \in \mathcal{T}^+$ and the selling price is among the two lowest in the interval:

$$p_{t_0}^{\text{sell}} \in \text{Bottom2}(\{p_t^{\text{sell}} : t \in \mathcal{T}^+\}) \Rightarrow SOC_{t_0+1} \leq SOC^{\text{max}} \quad (5.30)$$

This relaxes charging constraints to capture unexpected surplus energy.

- **Critical False Surplus Detection:** For positive intervals, the EMS must also detect the critical case where a surplus is predicted but does not occur. As these intervals typically correspond to high market prices, such errors can be particularly costly. If the current selling price exceeds the average buying price of the upcoming negative interval, the EMS is pre-alerted to switch the battery from charging to discharging as soon as the first inconsistency between forecast and measurement is observed.

If a surplus is forecasted in a positive interval but actual production is zero or negligible, and the current selling price is higher than the average future buying price:

$$p_{t_0}^{\text{sell}} > \mu(p_t^{\text{buy}} \mid t \in \mathcal{T}^-) \text{ and } E_{t_0}^{\text{prod}} \approx 0 \Rightarrow u_{t_0}^{\text{batt}} = \text{discharge} \quad (5.31)$$

The battery must switch mode to discharge to avoid purchasing at high prices later.

- **Surplus Energy Check in Negative Intervals:** In negative intervals, the EMS continually monitors the battery energy, if the stored energy exceeds the amount forecasted to be required for the remainder of the interval, the discharge constraints are relaxed. This allows the system to compensate for unexpected load increases instead of adhering to an unnecessarily conservative discharge plan.

If $t_0 \in \mathcal{T}^-$ and the stored energy exceeds the remaining predicted demand in that interval:

$$SOC_{t_0} - SOC^{\text{min}} > \sum_{t \in \mathcal{T}_{t_0}^-} P_t^{\text{load}} \cdot \Delta t \Rightarrow P_{t_0}^{\text{batt}, \text{dis}} \leq P^{\text{batt}, \text{dis}, \text{max}} \quad (5.32)$$

This allows relaxing the discharge constraint to meet unforeseen load increases.

Together, these strategies form a defensive layer that limits the impact of forecasting inaccuracies. By adjusting the SOC bounds in real time, the HEMS remains capable of responding to unforeseen disturbances, maintaining operational efficiency even when predictions deviate significantly from reality.

Figure 5.26, 5.27 and 5.28 provide a visual comparison of how forecasting uncertainty affects the yearly cashflow of a representative prosumer. Five different scheduling strategies are evaluated:

1. Receding-horizon optimization over a three-day window, updated hourly
2. Receding-horizon optimization over a one-day window, updated hourly
3. Daily optimization with a three-day horizon
4. Daily optimization with a one-day horizon
5. A naïve First-Charge First-Discharge (FCFD) policy, used as a lower-bound benchmark and unaffected by forecast quality

Figure 5.26 presents the results under perfect foresight. In this ideal case, only strategy 4 fails to outperform the brain-less benchmark, while strategies 1 and 3 yield identical performances. This equivalence is expected: when the forecast is flawless, the initial scheduling is already optimal, and no corrective updates are needed—making the rolling-horizon mechanism redundant.




Figure 5.26: Scheduling Methodology in presence of perfect forecasting

In Figure 5.27 introduces forecasting errors into the loop. Here the impact is severe: all optimization methods drop below the naïve benchmark. The comparison, however, clearly highlights the benefit of rolling-horizon formulations in mitigating the consequences of uncertainty. Even though the performance remains sub-optimal, frequent re-planning prevents the EMS from drifting too far from the economically sensible trajectory.

Figures 5.28 shows the effect of enabling the corrective countermeasures described earlier. Their contribution is evident: when these rules are added on top of the rolling-horizon framework, the yearly cashflow is lifted back above the lower benchmark. The improvement is especially visible for strategies that already rely on frequent updates, confirming that rule-based corrections and receding-horizon operation are a good a combination to restore robustness in the presence of imperfect forecasts.




Figure 5.27: Scheduling Methodology in presence of forecasting uncertainty




Figure 5.28: Scheduling Methodology in presence of forecasting uncertainty and countermeasures

5.8.2 Robust Optimization

Robust optimization addresses forecasting uncertainty from a fundamentally different perspective. Instead of relying on a single predicted trajectory, this approach evaluates the scheduling problem across multiple plausible future scenarios and identifies the sequence of actions that remains effective in most of them. The underlying idea is to explore a sufficiently wide ensemble of possible evolutions so that at least part of this “scenario cloud” will resemble the actual future conditions.

While robust formulations can traditionally be implemented through Linear Programming, in our study we adopt a simplified yet practical strategy. Leveraging the high computational speed of the developed rule-based EMS, the method introduced in chapter reference is used to generate a large “storm” of potential future profiles. The scheduling algorithm is then executed independently on each scenario, resulting in a collection of candidate SOC upper and lower thresholds.

The final SOC bounds applied in real time are selected statistically through a Confidence Interval (CI), expressed as a percentage. The enforced lower SOC limit is chosen as the value that is lower than CI% of all generated lower-threshold commands.

Symmetrically, the high SOC limit corresponds to the value that is higher than CI% of the upper-threshold commands. In this way, the dispatched

thresholds represent a robust compromise. Tight enough to be economically meaningful, yet sufficiently conservative to remain valid across most of the plausible future trajectories. An example of architecture is showed in Figure 5.29.




Figure 5.29: EMS Architecture with Robust Optimization

$$SOC_H^{CI} \geq IC \cdot N_{seq} \cdot SOC_H \quad (5.33)$$

$$SOC_L^{CI} \leq IC \cdot N_{seq} \cdot SOC_L \quad (5.34)$$

To evaluate the performances, the methodology has been tested on the full Ausgrid dataset for a total of $N_{profiles} = 127$. The scheduling and planning were performed over a total of nine months of the full year, as the first three months were used as a training set for the predictors.

Given a set of yearly cash flows for each user $CF_{Method}^{u,abs}$ and identifying $CF_{Method}^{u,final}$ as the final cash flow reached by that specific user managed with that method, and defining the case relying on perfect forecasting as “Oracle”, the resulting cash flows of each method have been normalized at each point of the sequence using the following operation:

$$CF_{Method}^{u,rel} = \frac{CF_{Method}^{u,abs} - CF_{Brain-less}^{u,abs}}{CF_{Oracle}^{u,final} - CF_{Brain-less}^{u,final}} \quad (5.35)$$

Then the average result across the whole set as been calculated as:

$$CF_{Method}^{set,rel} = \frac{1}{N_{profiles}} \sum_{i=1}^{N_{profiles}} CF_{Method}^{u,rel} \quad (5.36)$$

The operation has been repeated for the 2020 PUN and the 2024 PUN, and the results are shown in Figure 5.30 and Figure 5.31. Even if not as impactful as the knowledge-based rule set previously introduced, the implementation of this approach still has a sensible impact on the quality of the economic results of the prosumers, allowing them to considerably increase their expected economic return, marking it as a good practice.




Figure 5.30: Average Relative cashflow comparison on PUN 2020




Figure 5.31: Average Relative cashflow comparison on PUN 2024

5.8.3 Hybrid Methodology

An idea that comes out naturally is to combine the two methodologies exposed above. Since both approaches address uncertainty and improve planning from complementary perspectives, merging them offers the opportunity to exploit their strengths simultaneously, as displayed in An example of architecture is showed in Figure 5.32.

The concept remains simple: each profile generated in the flock of the robust optimization is also processed through the knowledge-based block, allowing the rule set to refine or correct the scheduling decisions before the final evaluation. Then, the real thresholds to send to the storage are determined through Eq. 5.33 and 5.34.




Figure 5.32: EMS Hybrid Architecture: Robust Optimization Knowledge-based Countermeasures

The idea has been tested in the same way exposed for the robust optimisation, using the same energy profiles and the same two price cases. The results have been processed through Eq. 5.35 and 5.36, and the outcomes are shown in Figure 5.33 and Figure 5.34.

What is possible to appreciate from these results is that the hybrid methodology leads to results similar to, or even better than, those reached by the two methods applied individually, implying that the two methodologies are actually complementary and that it makes sense to combine them.




Figure 5.33: Average Relative cashflow comparison on PUN 2020




Figure 5.34: Average Relative cashflow comparison on PUN 2024

5.8.4 Conclusion and Future Work

This section has presented some preliminary results on the impact of forecasting uncertainty on HEMS and the strategies developed to mitigate their effects. As discussed throughout the chapter, the presence of forecasting errors in household-level consumption and renewable generation cannot be overlooked when aiming to design an energy management system suitable for real-world deployment. Forecasts, by their nature, are subject to variability and imprecision, and thus it becomes essential to incorporate mechanisms capable of reacting to and compensating for such deviations.

The proposed methodology addresses this issue through the dual approach of integrating knowledge-based corrective strategies and applying robust optimisation principles. The former provides fast and computationally lightweight responses to common error patterns, while the latter offers a structured way to plan actions that remain effective across multiple possible future scenarios.

The initial results are promising, demonstrating that the quality of the optimisation output can be substantially restored even under non-ideal prediction conditions, while keeping the computational burden low enough for practical application in embedded or distributed systems.

Given the relevance of uncertainty handling for any realistic deployment of HEMS, particularly in contexts involving distributed energy resources, battery storage and prosumer dynamics, this line of research will remain a key focus in the future work. Ongoing investigations will aim to rigorously validate the performance of the proposed methods under a wider range of conditions, including different levels of forecast accuracy, weather variability and load profiles.

In addition, attention will be paid to formalising the proposed strategies within scalable optimisation frameworks and evaluating their integration into multi-agent or community-level architectures. The final objective is to develop a reliable, adaptive and cost-effective HEMS framework that is effectively deployable in real-world energy communities or in Smart Grids in general.

5.9 Conclusion and Final Remarks

In conclusion, this chapter has presented a lightweight yet effective HEMS architecture designed to operate at the residential level under realistic conditions of uncertainty and limited computational resources. By combining rule-based scheduling with a modular design, the system supports the integration of controllable loads and battery storage without relying on computationally intensive or complex optimisation methods.

This simplicity is not a drawback, but rather a strategic asset: in distributed energy systems, where coordination must scale across dozens or even hundreds of households, a fast and elementary local controller is a necessary building block.

Preliminary studies on the management of forecasting uncertainty demonstrate how the inclusion of knowledge-based correction mechanisms and robust planning strategies enables the system to handle prediction errors effectively, maintaining both operational feasibility and economic efficiency.

The true strength of the proposed HEMS lies in its autonomy and ability to produce reliable and interpretable scheduling decisions. These local decisions can be aggregated and coordinated at higher levels to achieve community-objectives.

The next chapter builds directly on this foundation, exploring how multiple distributed HEMSs, each running the proposed logic, can be orchestrated to achieve collective goals such as peak shaving, cost balancing, and energy self-sufficiency. Emphasis will be still placed on scalability and a non-coercive coordination method through price signal modulation will be presented.

Chapter 6

Secondary-Level Energy Management (CEMH)

*This work presents a coordinated energy management framework for Renewable Energy Communities based on a hierarchical architecture combining distributed HEMS units with a central Community Energy Management Hub. Coordination is achieved through price-based signaling and metaheuristic optimization, enabling the alignment of individual user objectives with community-level goals while preserving scalability and real-time feasibility. Digital twins are employed to emulate user behavior and support centralized decision-making without direct control actions. As a preliminary extension, a priority-based load-shedding mechanism is introduced at the secondary control level to enhance system resilience under grid constraints. Simulation results show that coordinated operation significantly improves overall performance, while controlled load shedding mitigates blackout conditions and increases the fraction of satisfied demand.*¹

¹Part of this chapter has been published as *Distributed EMS Coordination via Price-Signal Control for Renewable Energy Communities* [86]

6.1 Introduction

The Community Energy Management Hub (CEMH) represents the supervisory layer responsible for coordinating the actions of individual Home Energy Management Systems (HEMSs) within a Renewable Energy Community (REC). Its role is to provide a unified framework for collective decision-making, enabling the community to behave as an integrated entity rather than as a set of independent users.

Although an HEMS focuses exclusively on optimising the energy profile of a single household, the CEMH is tasked with enhancing the overall performance of the community. This includes aligning aggregated consumption and generation patterns and increasing the level of virtual self-consumption, which is a key source of economic benefit for RECs.

A major challenge in the design of a CEMH lies in the balance of individual and collective benefits. In many cases, the optimal community-level strategy does not coincide with the interests of specific members; therefore, ensuring fairness, transparency, and long-term participation becomes essential.

This chapter examines the structure, operating principles, and strategic relevance of the CEMH in contemporary energy communities. Particular attention is devoted to its coordination methods, its management of distributed Battery Energy Storage Systems (BESSs), and its capacity to guide the REC toward efficient, stable, and future-oriented operation.

The proposed approach addresses this issue by avoiding direct central control and instead employing incentive-driven coordination mechanisms, where the CEMH influences without overriding the decisions executed at the household level. This paradigm supports scalability, preserves user autonomy, and enhances social acceptance. The proposed architecture is displayed in Figure 5.29

6.2 Choice of the Optimization Approach

The work presented in [87] outlines an architecture similar in spirit to the one proposed in this study, featuring a centralised manager coordinating distributed local controllers. These local agents operate independently at the user level, with the goal of minimising household energy costs, while the central controller oversees the collective optimisation potentially enabling peer-to-peer energy exchanges among participants. Although such energy




Figure 6.1: Proposed architecture for two-layer EMS for the RECs showing the main information exchanges.

sharing is not permitted under current Italian regulation and is therefore excluded from this work, the reference model provides a valuable baseline.

The referenced study solves both levels of control using Mixed-Integer Linear Programming (MILP), achieving promising results under different configurations of load, renewable generation and storage. However, their analysis is limited to a community of just four users. In contrast, this study aims to simulate and coordinate communities of the order of one hundred users, more representative of actual energy communities with potential grid-level impact. This scale introduces challenges in both the problem's dimensionality and the model's complexity.

Although MILP provides mathematical rigour and certified optimality and is supported by mature solvers, it requires a strict linear problem formulation. Nonlinear dynamics must be approximated, typically through piecewise-linear models using auxiliary variables and big-M constraints. As the problem's size increases, so does the computational burden and the complexity of model tuning, often limiting applicability of MILP in large-scale or real-time scenarios.

For this reason, metaheuristic approaches, such as Particle Swarm Optimisation (PSO), have gained popularity as a more scalable alternative. PSO is a population-based, derivative-free optimisation method that excels in handling nonlinear, nonconvex problems with limited modelling effort. Each candidate solution is encoded by a particle whose position evolves based on

individual experience and collective swarm behaviour.

It supports flexible constraint handling, allows encoding of discrete decisions, and is inherently parallelizable. These features make it well-suited for problems involving a large number of agents and complex interactions, especially when near-real-time responsiveness is needed.

Studies in the domain of polygeneration highlight the recurring limitations of MILP, including the need for piecewise linearization, full-horizon coupling and clustering techniques to reduce dimensionality, factors that inevitably increase solver burden and modelling effort [88]. Further evidence from intraday hydropower scheduling with fine-grained time discretization shows that MILP struggles to remain tractable as physical detail and system size grow, especially when compared to heuristic alternatives [89].

PSO trades optimality guarantees for model agnosticism, ease of implementation and computational tunability. Such traits make it highly attractive for large-scale EMS coordination where digital twins, uncertainty and dynamic feedback are present in the loop. Hardware-level studies have demonstrated that PSO can operate effectively even on low-cost microcontrollers, with predictable timing and small memory footprints, thanks to bounded population sizes and iteration caps [90].

Therefore, due to its scalability, modelling flexibility, tunability and compatibility with potential hardware constraints PSO has been adopted in this study as the primary coordination method. MILP remains used as a benchmark reference to represent the theoretical optimal achievable under idealised conditions, providing a valuable baseline for performance comparison in the studied scenarios.

6.3 Notation and Preliminary Consideration

In this chapter only domestic users are considered, with different degrees of flexibility depending on their typology: prosumers are equipped with BESS units, while simple users are modelled through controllable loads. The data set from which the load and production profiles are drawn to simulate the user group is described in detail in the corresponding section.

The focus of this study is solely on the optimisation technique; therefore, the forecasting process is not described in detail. The optimisation strategy relies on a rolling-horizon architecture with a one-hour time step, meaning that user behaviour is forecast and then optimised at the beginning of each

step. For the remainder of the chapter, predictions are assumed to be perfect, a choice made once again to focus the analysis on the performance of optimisation techniques without the noise introduced by forecasting inaccuracy.

The impact of these studies is left for future research. Throughout the chapter, the term *prosumption* is used to represent the energy exchange between a user and the grid. The sign convention is defined such that the value is positive when energy is injected into the grid and negative when it is drawn from it.

The proposed architecture relies heavily on data exchange between the nodes of the distributed system. The use of digital twins in the CEMH helps minimise the number of such exchanges. Moreover, since the Energy Management problem is a “slow” one, with time steps of one hour or, at minimum, fifteen minutes, latency and bandwidth are not expected to pose significant issues. It can therefore be assumed that, especially in the case of many users, the most critical time requirement and the main contributor to computational load will be the time needed to perform the optimisation.

As will be shown in the following section, important considerations must be made regarding the prices at which energy is bought and sold. The selling price *Sell_price* at which the prosumers sell their excess energy is assumed to be equal to the PUN (*Prezzo Unico Nazionale*). The PUN is an array that contains the result of the day-ahead market listing, hour by hour, for the following day.

The actual energy price is only one component of the total purchase cost. Some components are fixed and independent of consumption, while the variable components are proportional to it. Since only the variable part can affect load planning and BESS scheduling, the hourly buying price considers only the variable components:

$$E_{\text{price}} = \text{PUN}_{\text{avg}}^{\text{month}} + F_{\text{price}}, \quad (6.1)$$

$$\text{Buy_price} = (E_{\text{price}} + T_{\text{price}} + S_{\text{price}}) \cdot \text{VAT}. \quad (6.2)$$

where:

- $\text{PUN}_{\text{avg}}^{\text{month}}$ = monthly average PUN [Euro/kWh];
- $F_{\text{price}} = \begin{cases} 0.011 \text{ [Euro/kWh]}, & \text{Sunday and from 23:00 to 07:00,} \\ 0.0286 \text{ [Euro/kWh]}, & \text{otherwise;} \end{cases}$

- $T_{\text{price}} = 0.012$ [Euro/kWh] (transport cost);
- $S_{\text{price}} = 0.038$ [Euro/kWh] (system cost);
- VAT = 1.1.

Another important aspect to consider is virtual self-consumption within the community which is briefly recalled here for completeness. The Italian legislation defines it as the minimum, calculated hour by hour, between the total energy injected into the grid from renewable production by the community and the total energy requested by the community from the grid.

This parameter is relevant because the community receives a reward for each kWh self-consumed in this way. Although the exact value may vary depending on the hour and geographical area, it is typically close to 0.10 Euro/kWh. Therefore, in this study a constant value of 0.10 Euro/kWh is assumed.

Regarding the REC, an internal compensation mechanism is employed. In the simulated communities, the cumulative incentive obtained through virtual self-consumption is partially allocated to reward and compensate individual users who incur losses while contributing to collective objectives. The remaining share is retained at the community level.

To apply both the benchmark and the proposed methods, the following quantities must be known:

- H : prediction horizon (number of steps ahead of the forecasting),
- N_u : number of users.

Then, for each user $u \in [1, N_u]$, the following data must be collected:

- p_t : forecasted prosumption;
- $\text{SOC}_{\text{max}}, \text{SOC}_{\text{min}}, \text{SOC}_0$: BESS SOC limits and initial value;
- N_C : number of controllable loads;
- Buy_price, Sell_price: hourly buying and selling prices;
- Load_{max}, Prod_{max}: consumption and production at peak hourly rates.

For each load $n \in \{1, \dots, N_C\}$, the following parameters must be specified:

- CLv_n : energy absorption [kWh];
- CLw_n : set of hours during which the load can be activated;
- CLr_n : number of hours that the load must operate.

The maximum request from controllable loads is therefore:

$$CLv_{\max} = \sum_{n=1}^{N_C} CLv_n. \quad (6.3)$$

For each controllable load n , the following additional parameters must be defined:

- $CLreq_n$: minimum number of activations required;
- $CLTW_n$: set of allowable scheduling hours;
- $CLoffset_n$: size of the first scheduling slot;
- $CLblocked_n$: parameter indicating whether the load request has already been satisfied in the initial slot.

6.4 Definition of the MILP Benchmark

Mixed-Integer Linear Programming (MILP) is one of the most common and reliable approaches in the literature for addressing the problem of energy management. The first steps are the definition of a set of variables representing the state of the problem, followed by a cost function and a set of constraints describing the relationships among the defined variables. In this work, a mixed-integer linear formulation of the problem was developed, starting from the linearisation of the single user–prosumer case and then extended to cover the entire community.

6.4.1 Definition of the MILP for Single Users

The optimisation problem for a single user aims at scheduling the operation of the BESS and/or the available controllable loads to maximise the user's cash flow. To formulate the problem, the following variables are introduced, ordered as shown in Figure 6.2:

	Horizon	Horizon	Horizon	N_Ctrl_Loads * Horizon	N_Ctrl_Loads * Horizon	Horizon	Horizon	Horizon
Variable's Names	User Presumption Post Optimization	BESS Actions	Ctrl Loads Tot	Ctrl Loads Bin	Ctrl Loads Value	User Presumption Injection	User Presumption Absorbtion	User Presumption Sign
Variable's Labels	x_t	b_t	CLT_t	$CLB_{n,t}$	$CLV_{n,t}$	Inj_t	Abs_t	xs_t
Upper Bound	$\max(p_t, 0)$	$\max(p_t, 0)$	CLV_{Max}	1	CLV_n	$\max(p_t, 0)$	$\max(-p_t, 0)$ $+ CLV_{Max}$	1
Lower Bound	$\min(p_t, 0)$ $- CLV_{Max}$	$\min(p_t, 0)$	0	0	0	0	0	0

Figure 6.2: Visual representation of the order of the variables for the single-user problem.

- x_t : presumption forecast after optimisation;
- b_t : BESS action after optimisation;
- CLT_t : total consumption of the controllable loads;
- $CLB_{n,t}$: hourly scheduling of controllable load n (binary);
- $CLV_{n,t}$: hourly consumption of controllable load n ;
- Inj_t : grid injection after optimisation;
- Abs_t : grid absorption after optimisation;
- x_t^s : supporting binary variable for sign handling.

Counting all variables, the total number of decision variables in the optimisation model is:

$$N_{\text{var}} = H (5 + 2N_c),$$

where H is the optimisation horizon and N_c is the number of controllable loads.

1) Definition of Inequality constraints (total 12*H)

BESS Constraint (2H constraints) The cumulative BESS action must keep the State of Charge within its admissible limits for all $t^* \in \{1, \dots, H\}$:

$$\begin{cases} SOC_0 + \sum_{t=1}^{t^*} b_t \leq SOC_{\max}, \\ SOC_0 + \sum_{t=1}^{t^*} b_t \geq SOC_{\min}. \end{cases} \quad (6.4)$$

Equivalently,

$$\begin{cases} \sum_{t=1}^{t^*} b_t \leq SOC_{\max} - SOC_0, \\ -\sum_{t=1}^{t^*} b_t \leq -SOC_{\min} + SOC_0, \end{cases} \quad (6.5)$$

which introduces two linear constraints for each t^* .

Prosumption Sign Constraint (2H constraints) The binary variable $x_{t^*}^s$ identifies the sign of the optimised prosumption x_{t^*} :

$$x_{t^*}^s = \begin{cases} 0, & x_{t^*} < 0, \\ 1, & x_{t^*} > 0, \end{cases} \quad \forall t^* \in \{1, \dots, H\}. \quad (6.6)$$

Using a big- M linearisation (with $M = \max(\text{Load}_{\max}, \text{Prod}_{\max})$):

$$\begin{cases} \frac{x_{t^*}}{M} \leq x_{t^*}^s, \\ \frac{x_{t^*}}{M} + 1 \geq x_{t^*}^s, \end{cases} \quad (6.7)$$

which yields the linear form:

$$\begin{cases} \frac{x_{t^*}}{M} - x_{t^*}^s \leq 0, \\ -\frac{x_{t^*}}{M} + x_{t^*}^s \leq 1. \end{cases} \quad (6.8)$$

Grid Injection Constraint (4H constraints) Linearization of

$$\text{Inj}_{t^*} = \max(0, x_{t^*}). \quad (6.9)$$

The standard big- M formulation is:

$$\begin{cases} x_{t^*} \leq \text{Inj}_{t^*} \leq x_{t^*} + M(1 - x_{t^*}^s), \\ 0 \leq \text{Inj}_{t^*} \leq Mx_{t^*}^s. \end{cases} \quad (6.10)$$

This expands to:

$$\begin{cases} x_{t^*} - \text{Inj}_{t^*} \leq 0, \\ -x_{t^*} + \text{Inj}_{t^*} + Mx_{t^*}^s \leq M, \\ -\text{Inj}_{t^*} \leq 0, \\ \text{Inj}_{t^*} + Mx_{t^*}^s \leq M. \end{cases} \quad (6.11)$$

Grid Absorption Constraint (4H constraints) Linearization of

$$\text{Abs}_{t^*} = \max(0, -x_{t^*}). \quad (6.12)$$

Using big- M :

$$\begin{cases} -x_{t^*} \leq \text{Abs}_{t^*} \leq -x_{t^*} + Mx_{t^*}^s, \\ 0 \leq \text{Abs}_{t^*} \leq M(1 - x_{t^*}^s). \end{cases} \quad (6.13)$$

Equivalent linear constraints:

$$\begin{cases} -x_{t^*} - \text{Abs}_{t^*} \leq 0, \\ x_{t^*} + \text{Abs}_{t^*} - Mx_{t^*}^s \leq 0, \\ -\text{Abs}_{t^*} \leq 0, \\ \text{Abs}_{t^*} + Mx_{t^*}^s \leq M. \end{cases} \quad (6.14)$$

2) Equality Constraints ((1 + 3N_c)H)

Actions Constraint (H constraints) For every time step $t^* \in \{1, \dots, H\}$, the optimised prosumption must satisfy:

$$x_{t^*} + CLT_{t^*} + b_{t^*} = p_{t^*}. \quad (6.15)$$

Controllable Loads Total Constraint (N_cH constraints) The total controllable load consumption must be equal to the sum of the consumption of each controllable load:

$$CLT_{t^*} = \sum_{n=1}^{N_c} CLV_{n,t^*} \quad \implies \quad CLT_{t^*} - \sum_{n=1}^{N_c} CLV_{n,t^*} = 0. \quad (6.16)$$

Controllable Loads Value Constraint (N_cH constraints) Each controllable load n consumes a fixed amount of energy CLv_n at every hour in which it is scheduled:

$$CLV_{n,t^*} = CLv_n CLB_{n,t^*} \quad \implies \quad CLV_{n,t^*} - CLv_n CLB_{n,t^*} = 0. \quad (6.17)$$




Figure 6.3: Visual representation of the variables of controllable load n with $CLreq_n = 24$, $CLTW_n$ from 2 to 15, and $CLoffset_n = 19$.

Controllable Loads Schedule and Slots ($N_c H$ constraints) The binary variable CLB_{n,t^*} determines whether the load n is scheduled at hour t^* . Figure 6.3 illustrates the interaction between the parameters $CLreq_n$, $CLTW_n$, and $CLoffset_n$ for an example load.

From these parameters, a matrix A_{CLn} is constructed with N_{SLOT} rows and H columns (Figure 6.4). Each row identifies the time window of a slot, containing ones in the admissible hours and zeros elsewhere.

A trigger vector b_{CLn} is defined as:

$$b_{CLn} = \begin{bmatrix} CLtrigg_{1,n} \\ \vdots \\ CLtrigg_{N_{\text{SLOT}},n} \end{bmatrix}, \quad (6.18)$$

where

$$CLtrigg_{s,n} = \begin{cases} 1, & \text{if the load has not yet been activated in slot } s, \\ 0, & \text{otherwise.} \end{cases} \quad (6.19)$$

For each slot s , the scheduling constraint becomes:

$$\sum_{t^*=1}^H A_{CLn}(s, t^*) CLB_{n,t^*} = CLtrigg_{s,n}, \quad \forall s \in \{1, \dots, N_{\text{SLOT}}\}. \quad (6.20)$$

3) Definition of the cost function

Since the objective of the single user is to maximise the resulting cash flow, the associated cost function f_{cost} assigns a value only to the grid exchanges. Specifically, for each slot of the hour t :

$$f_{\text{cost}}(\text{Inj}_t) = \text{Sell_price}_t \quad f_{\text{cost}}(\text{Abs}_t) = \text{Buy_price}_t \quad (6.21)$$



Figure 6.4: Example of the architecture of the matrix $A_C L n$

and all remaining entries of the cost vector are equal to zero.

Thus, the economic objective becomes:

$$\max \sum_{t=1}^H (\text{Sell_price}_t \cdot \text{Inj}_t - \text{Buy_price}_t \cdot \text{Abs}_t) \quad (6.22)$$

which corresponds to maximising the user's cash flow over the prediction horizon.

6.4.2 Definition of the MILP for the REC

The purpose of community-level optimisation is to coordinate the actions of N_u users to maximise community self-consumption and perform peak shaving. The formulation of the REC problem extends directly from the single-user formulation: for each user, the complete set of local variables is defined, and an additional set of community-level variables is introduced.

REC-Level Variables For every time step $t \in \{1, \dots, H\}$, the following community variables are added:

- x_t^{REC} : total community prosumption after optimisation,
- $\text{Inj}_t^{\text{REC}}$: total community grid injection,
- $\text{Abs}_t^{\text{REC}}$: total community grid absorption,
- $\text{Self}_t^{\text{REC}}$: virtual self-consumption of the community,
- $\text{Gap}_t^{\text{REC}}$: gap between two consecutive points of the prosumption profile,
- $(\text{Gap}_t^{\text{REC}})^{\text{mod}}$: absolute value of the prosumption gap.

Figure 6.5 illustrates the ordering of these variables within the global decision vector.

Variable's Names	VAR USER 1	...	VAR USER N	REC Prosumption	REC Prosumption Injection	REC Prosumption Absorption	REC Selfconsumption	REC Prosumption Gap	REC Prosumption Gap module
Variable's Labels				x_{RECt}	Inj_{RECt}	Abs_{RECt}	$Self_{RECt}$	Gap_{RECt}	Gap_{RECt}^{mod}
Upper Bound	← <i>User's Variables</i> →			1	$\sum_{u=1}^{N_u} Inj_{u,t}^{UB}$	$\sum_{u=1}^{N_u} Abs_{u,t}^{UB}$	Inj_{RECt}^{UB}	$Prod_{Max}^u$	$\max \left(Prod_{Max}^u, Load_{Max}^u \right)$
Lower Bound				0	0	0	Abs_{RECt}^{UB}	$-Load_{Max}^u$	0

← *REC's Variables* →

Figure 6.5: Visual representation of the order of the variables for the whole REC problem.

Community Maximum Limits The theoretical peak consumption and peak production of the REC are defined as the sum of the corresponding limits of each user:

$$Load_Max^{REC} = \sum_{u=1}^{N_u} Load_Max^u, \quad Prod_Max^{REC} = \sum_{u=1}^{N_u} Prod_Max^u. \quad (6.23)$$

Decision Variable Count Let N_{user}^{vars} denote the number of variables required for a single-user problem. The total number of variables in the REC-level optimisation becomes:

$$N_{REC}^{vars} = N_u \cdot N_{user}^{vars} + 6H. \quad (6.24)$$

REC Constraint Structure Community-level constraints must include the following:

- the full set of constraints defined for each user of the REC,
- the additional constraints defining the REC variables and their relationships (e.g., aggregation of prosumptions, injection, absorption, virtual self-consumption, and peak-shaving constraints).

Both equality and inequality constraint matrices follow the block-structured organisation shown in Figure 6.6, where user-level constraint blocks are combined with REC-level constraints to form the complete MILP system.



Figure 6.6: Visual representation of the content on the constraint matrices of the REC problem.

1) Definition of Inequality constraints (total $(12 \cdot N_u + 4) \cdot H$)

Self-Consumption Constraint ($2H$ constraints) The REC's virtual self-consumption at each time step is defined as the minimum between the total grid injection and the total grid absorption:

$$\text{Self}_{t^*}^{\text{REC}} = \min(\text{Inj}_{t^*}^{\text{REC}}, \text{Abs}_{t^*}^{\text{REC}}), \quad \forall t^* \in \{1, \dots, H\}. \quad (6.25)$$

The minimum operator is linearised through the following inequalities:

$$\begin{cases} \text{Self}_{t^*}^{\text{REC}} \leq \text{Inj}_{t^*}^{\text{REC}}, \\ \text{Self}_{t^*}^{\text{REC}} \leq \text{Abs}_{t^*}^{\text{REC}}. \end{cases} \implies \begin{cases} \text{Self}_{t^*}^{\text{REC}} - \text{Inj}_{t^*}^{\text{REC}} \leq 0, \\ \text{Self}_{t^*}^{\text{REC}} - \text{Abs}_{t^*}^{\text{REC}} \leq 0. \end{cases} \quad (6.26)$$

To enforce a correct solution, the objective function must be formulated so that $\text{Self}_{t^*}^{\text{REC}}$ is maximised on whole prediction horizon.

Absolute Value of Prosumption Gap ($2H$ constraints) For each time step, the REC prosumption gap is defined as:

$$\text{Gap}_{t^*}^{\text{REC}} = x_{t^*+1}^{\text{REC}} - x_{t^*}^{\text{REC}}, \quad \forall t^* \in \{1, \dots, H\}. \quad (6.27)$$

Its absolute value is modelled through the auxiliary variable $(\text{Gap}_{t^*}^{\text{REC}})^{\text{mod}}$, which satisfies the following:

$$(\text{Gap}_{t^*}^{\text{REC}})^{\text{mod}} = |\text{Gap}_{t^*}^{\text{REC}}|. \quad (6.28)$$

This is linearised through:

$$\begin{cases} \text{Gap}_{t^*}^{\text{REC}} \leq (\text{Gap}_{t^*}^{\text{REC}})^{\text{mod}}, \\ -\text{Gap}_{t^*}^{\text{REC}} \leq (\text{Gap}_{t^*}^{\text{REC}})^{\text{mod}}. \end{cases} \implies \begin{cases} \text{Gap}_{t^*}^{\text{REC}} - (\text{Gap}_{t^*}^{\text{REC}})^{\text{mod}} \leq 0, \\ -\text{Gap}_{t^*}^{\text{REC}} - (\text{Gap}_{t^*}^{\text{REC}})^{\text{mod}} \leq 0. \end{cases} \quad (6.29)$$

To produce the intended peak-shaving effect, the objective function must be designed to minimise $(\text{Gap}_{t^*}^{\text{REC}})^{\text{mod}}$.

REC Grid Injection (H constraints) The total REC grid injection is the sum of the grid injections of all users:

$$\text{Inj}_{t^*}^{\text{REC}} = \sum_{u=1}^{N_u} \text{Inj}_{u,t^*} \implies \text{Inj}_{t^*}^{\text{REC}} - \sum_{u=1}^{N_u} \text{Inj}_{u,t^*} = 0. \quad (6.30)$$

REC Grid Absorption (H constraints) Similarly, the total REC absorption from the grid is the sum of the absorptions of the users:

$$\text{Abs}_{t^*}^{\text{REC}} = \sum_{u=1}^{N_u} \text{Abs}_{u,t^*} \implies \text{Abs}_{t^*}^{\text{REC}} - \sum_{u=1}^{N_u} \text{Abs}_{u,t^*} = 0. \quad (6.31)$$

REC Prosumption (H constraints) The REC prosumption is defined as the difference between the total grid injection and the total grid absorption:

$$x_{t^*}^{\text{REC}} = \text{Inj}_{t^*}^{\text{REC}} - \text{Abs}_{t^*}^{\text{REC}} \implies x_{t^*}^{\text{REC}} - \text{Inj}_{t^*}^{\text{REC}} + \text{Abs}_{t^*}^{\text{REC}} = 0. \quad (6.32)$$

REC Prosumption Gap (H constraints) For peak-shaving purposes, the REC prosumption gap is defined as the difference between two consecutive REC prosumption values, for all $t^* \in \{1, \dots, H-1\}$:

$$\text{Gap}_{t^*}^{\text{REC}} = x_{t^*}^{\text{REC}} - x_{t^*+1}^{\text{REC}} \implies \text{Gap}_{t^*}^{\text{REC}} - x_{t^*}^{\text{REC}} + x_{t^*+1}^{\text{REC}} = 0. \quad (6.33)$$

3) REC-Level Cost Function

The objective of REC combines both the individual goals of each user and the collective goals of the community. For the components of the cost vector associated with user-level variables, the REC cost function simply replicates

the corresponding single-user's costs. For community-level variables, the objective incorporates the reward for virtual self-consumption and the penalty associated with prosumption variations.

The resulting REC cost vector $f_{\text{cost}}^{\text{REC}}$ is therefore defined as follows:

$$f_{\text{cost}}^{\text{REC}}(\text{Self}_t^{\text{REC}}) = 0.10, \quad (6.34)$$

corresponding to the incentive (€/kWh) granted for each unit of virtual self-consumption within the community.

To promote peak shaving and smooth the aggregated profile of the REC, the absolute value of the prosumption gap is assigned a small penalisation:

$$f_{\text{cost}}^{\text{REC}}((\text{Gap}_t^{\text{REC}})^{\text{mod}}) = 0.01, \quad (6.35)$$

a value selected through preliminary calibration to ensure that maximising the self-consumption reward remains the primary objective.

For user-related variables, the REC cost vector inherits the corresponding single-user cost weights:

$$f_{\text{cost}}^{\text{REC}}(\text{Var}_{\text{user},N}) = f_{\text{cost}}^{(N)}. \quad (6.36)$$

All remaining entries in the REC cost vector are set to zero.

$$f_{\text{cost}}^{\text{REC}}(\cdot) = 0 \quad \text{for all other positions.} \quad (6.37)$$

6.5 Definition of the Proposed Approach

The proposed architecture combines a rule-based approach with metaheuristic optimisation techniques, implemented in a distributed system. Each prosumer is equipped with a local EMS that interacts with a central coordination node. The distributed HEMS units send their forecasted prosumption profiles to this central CEMH.

In CEMH, a digital twin of each prosumer's HEMS is maintained. These digital twins allow the system to emulate the behaviour of the entire community locally, minimising communication delays. Once all forecasts are collected, a metaheuristic optimisation method is applied to determine the optimal strategy to achieve the REC objectives. In this work, the PSO technique was used.

Importantly, this optimisation does not directly alter the actions of individual users. Instead, it indirectly influences HEMS behaviour by adjusting

the perceived buy and sell prices, effectively steering the distributed systems toward the desired global outcome. These updated commands, expressed as price modifications, are then sent back to the distributed HEMS network.

Note that the label *CMD* assigned to some of the following variables indicates that they are evaluated after accounting for the effect of storage management within the system.

6.5.1 Energy Management Hub

Figure 6.7 shows the information flow between the blocks of the system, while Figure 6.8 illustrates the CEMH logic. After collecting the distributed load and production forecasts, the flow splits into two branches: load coordination (left, blue) and BESS Coordination (right, green).

Both rely on digital twins to emulate the user's HEMSs behaviour and employ PSO as a metaheuristic optimisation technique. The load coordination is achieved by adjusting the $\text{Buy}_{\text{price},t}$, while the BESS coordination is driven by altering the $\text{Sell}_{\text{price},t}$. Defining $\text{Sell}_{\text{price},t}^u$ and $\text{Buy}_{\text{price},t}^u$ as the price profiles received as input by the HEMS of the user u , the following relationships hold:

$$\text{Sell}_{\text{price},t}^u = \text{Sell}_{x,t}^u + \text{Sell}_{\text{price},t} \quad (6.38)$$

$$\text{Buy}_{\text{price},t}^u = \text{Buy}_{x,t}^u + \text{Buy}_{\text{price},t} \quad (6.39)$$

where $\text{Sell}_{x,t}^u$ and $\text{Buy}_{x,t}^u$ are the decision variables optimised by CEMH to achieve coordination. Note that the HEMS digital twins are exact clones of the distributed ones. To complete a full coordination cycle, both paths must be executed once. The priority between them is determined by evaluating the variability of the controlled variables, defined as follows:

$$\text{Buy}_{\text{price}}^{\text{GAP}} = \max(\text{Buy}_{\text{price},t}) - \min(\text{Buy}_{\text{price},t}) \quad (6.40)$$

$$\text{Sell}_{\text{price}}^{\text{GAP}} = \max(\text{Sell}_{\text{price},t}) - \min(\text{Sell}_{\text{price},t}) \quad (6.41)$$

As noted, the outputs of this process are not direct control commands to individual users but price alterations perceived by their EMSs. Since these alterations translate into costs in real-world scenarios, it is desirable to minimise them. Therefore, the coordination process starts with the branch




Figure 6.7: Architecture of the CEMH and information flow to and from the distributed EMS network.

that require smaller price adjustments to influence user behaviour toward the REC objectives, assumed to be the one associated with the smaller GAP.

Both paths start with the same two steps: a preliminary digital-twin planning, followed by a forecast of REC behaviour. First, the optimisation and scheduling of all controlled users are executed without any alterations to the price profile, which will only be optimised later within the branch. This operation simulates the REC as it would behave without coordination, providing a baseline forecast of the community's behaviour.

The result of this step provides, for each user u , the following forecasted profiles:

- $\text{Inj}_{\text{flow},t}^u$: grid injection profile scheduled in absence of BESS
- $\text{Abs}_{\text{flow},t}^u$: grid absorption profile scheduled in absence of BESS
- $\text{Inj}_{\text{CMDflow},t}^u$: grid injection profile scheduled in presence of BESS
- $\text{Abs}_{\text{CMDflow},t}^u$: grid absorption profile scheduled in presence of BESS

From these, the cumulative injection and absorption profiles of all users can be evaluated as follows:

$$\text{Inj}_{\text{flow},t}^{\text{tot}} = \sum_{u=1}^{N_u} \text{Inj}_{\text{flow},t}^u \quad (6.42)$$

$$\text{Inj}_{\text{CMDflow},t}^{\text{tot}} = \sum_{u=1}^{N_u} \text{Inj}_{\text{CMDflow},t}^u \quad (6.43)$$




Figure 6.8: Flowchart of the decision three of the optimization process performed inside of the CEMH.

$$\text{Abs}_{\text{flow},t}^{\text{tot}} = \sum_{u=1}^{N_u} \text{Abs}_{\text{flow},t}^u \quad (6.44)$$

$$\text{Abs}_{\text{CMDflow},t}^{\text{tot}} = \sum_{u=1}^{N_u} \text{Abs}_{\text{CMDflow},t}^u \quad (6.45)$$

Note that these profiles differ from those of the REC itself. A further step is required to compute the REC profiles, both in the case without BESS and after scheduling:

$$\text{REC}_{\text{flow},t} = \text{Inj}_{\text{flow},t}^{\text{tot}} - \text{Abs}_{\text{flow},t}^{\text{tot}} \quad (6.46)$$

$$\text{REC}_{\text{CMDflow},t} = \text{Inj}_{\text{CMDflow},t}^{\text{tot}} - \text{Abs}_{\text{CMDflow},t}^{\text{tot}} \quad (6.47)$$

$$\text{Inj}_{\text{flow},t}^{\text{REC}} = \max(\text{REC}_{\text{flow},t}, 0) \quad (6.48)$$

$$\text{Inj}_{\text{CMDflow},t}^{\text{REC}} = \max(\text{REC}_{\text{CMDflow},t}, 0) \quad (6.49)$$

$$\text{Abs}_{\text{flow},t}^{\text{REC}} = \max(-\text{REC}_{\text{flow},t}, 0) \quad (6.50)$$

$$\text{Abs}_{\text{CMDflow},t}^{\text{REC}} = \max(-\text{REC}_{\text{CMDflow},t}, 0) \quad (6.51)$$

From this point on, the two branches diverge; their steps are described in detail below. From the BESS coordination, the steps are the following.

REC Ideal Injection Profile Definition

Assuming that BESS units are used only for the needs of their respective owners, virtual self-consumption can occur only during time intervals where a surplus of production and injection is forecasted. Under this assumption, the maximum possible injection at each hour cannot exceed the cumulative injection forecast without storage:

$$\text{Inj}_{\text{flow},t^*}^{\text{tot}} \geq \text{Inj}_{\text{CMDflow},t^*}^{\text{tot}} \quad \forall t^* \in \{1, \dots, H\} \quad (6.52)$$

During the forecasting horizon, the energy profile can be segmented into intervals characterised by energy exchanges of the same sign, alternating

between injection and absorption periods. Defining TW_{Inj}^{tot} as a time window of adjacent hours with $Inj_{flow,t}^{tot} > 0$, the injected and absorbed energy within this interval can be evaluated.

$$E_{TW}^{Inj} = \sum_{t^* \in TW_{Inj}^{tot}} Inj_{CMDflow,t^*}^{tot} \quad (6.53)$$

$$E_{TW}^{Abs} = \sum_{t^* \in TW_{Inj}^{tot}} Abs_{CMDflow,t^*}^{tot} \quad (6.54)$$

The target injection profile for each such window is defined as:

$$REC_{CMDflow,t^*}^{Target} = Inj_{CMDflow,t^*}^{tot,Target} - Abs_{CMDflow,t^*}^{tot} \quad \forall t^* \in TW_{Inj}^{tot} \quad (6.55)$$

The energy available in each window is:

$$E_{TW}^{disp} = E_{TW}^{Inj} - E_{TW}^{Abs} \quad (6.56)$$

A first approximation of the target profile is:

$$Inj_{CMDflow,t^*}^{tot,Target} = \frac{E_{TW}^{disp}}{|TW_{Inj}^{tot}|} + Abs_{CMDflow,t^*}^{tot} \quad \forall t^* \in TW_{Inj}^{tot} \quad (6.57)$$

Since the feasibility condition in 6.52 must always hold, the profile is iteratively corrected by redistributing excess energy from unfeasible hours between the others. This operation is repeated for each injection interval in the REC, obtaining a target profile that serves as a reference for subsequent optimisation steps.

Figure 6.9 shows a visual representation of this procedure, while Figure 6.10 illustrates the corresponding logic flowchart.




Figure 6.9: Examples of the definition of the Target Injection Profile.




Figure 6.10: Flowchart of the evaluation process of the Target Injection Profile.

6.5.2 BESS Coordination Through PSO Optimization

In this branch, the search for the optimal solution is carried out using Particle Swarm Optimisation (PSO) in two distinct phases. In both phases, the digital twins of the community HEMSs are included in the cost functional, meaning that the entire REC is fully simulated for each particle of the swarm and at every iteration of the algorithm.

The optimised variables are the alterations $Sell_x^u$ applied to the selling price $Sell_{price}^u$ perceived by each user. For consumers, $Sell_{x,t}^u$ is used at every time step, while for prosumers an upper boundary is imposed. Since $Sell_x^u$ essentially represents an incentive for paying the prosumer to act against their individual interest to achieve community objectives, it is important to guarantee sustainability. In this work, such boundary is set to be equal to the reward of self-consumption, 0.10 €/kWh.

The choice of a rule-based technique is essential here as it reduces computational cost while still allowing multiple full REC simulations within an acceptable time frame. For each particle p of the swarm, there is an $Inj_{CMDflow,t^*}^{(tot,p)}$ and a corresponding $CER_{CMDflow,t^*}^p$.

The first phase is the *Aiming* step, in which user selling prices are altered to drive the particle injections $Inj_{CMDflow,t^*}^{(tot,p)}$ as close as possible to the target profile $Inj_{CMDflow,t^*}^{(tot,Target)}$ for all $t^* \in \{1, \dots, H\}$. Figure 6.11 offers a visual example of this step. The proximity is evaluated through RMSE, so the functional to minimise becomes the following:

$$Aim_{cost}^p = \sqrt{\frac{\sum_{t^*=1}^H \left(Inj_{CMDflow,t^*}^{(tot,p)} - Inj_{CMDflow,t^*}^{(tot,Target)} \right)^2}{H}} \quad (6.58)$$

At the end of this run of the PSO is expected a set of $Sell_{x,t}^{(u,aim)}$ for all $u \in \{1, \dots, N_u\}$ that corresponds to an $Inj_{CMDflow,t^*}^{(tot,aim)}$ very close to the Target Injection Profile, with a corresponding Aim_{cost} .

It is possible to define $Sell_x^{aim}$ as the sum of all contributions $Sell_{x,t}^{(u,aim)}$ in both t and u , which represents the total set of promises made to achieve a given solution. A challenging property of the energy management problem at the user level is its discrete nature: to alter the scheduling of a storage unit in a given hour, the corresponding price adjustment must exceed a minimum threshold.




Figure 6.11: Example of the effect of the alteration on the selling prices and result after the aiming run of the PSO

An adjustment smaller than this value does not trigger any scheduling change, while one that is larger but still below a second threshold simply increases the price without effect. Both cases result in wasted resources, and thus the solution obtained at the end of the Aiming step does not take into account this behaviour.

For this reason, a second *Refining* run of PSO is introduced. In this phase, the sets of $\text{Sell}_{x,t}^{(u,\text{aim})}$ obtained in the first step are used both as initial conditions and as upper bounds for the corresponding variables. An example of this phase is shown in Figure 6.12. The optimisation functional is also modified. While it continues to include and rely on the digital twins of the community EMSs, the cost function is defined as follows:

$$\text{Ref}_{\text{cost}}^{(p,\text{app})} = \sqrt{\frac{\sum_{t^*=1}^H \left(\text{Inj}_{\text{CMDflow},t^*}^{(\text{tot},p)} - \text{Inj}_{\text{CMDflow},t^*}^{(\text{tot Target})} \right)^2}{H}} \quad (6.59)$$

$$\text{Sell}_x^{(p,\text{ref})} = \sum_{u=1}^{N_u} \sum_{t^*=1}^H \text{Sell}_{x,t^*}^{(p,u)} \quad (6.60)$$




Figure 6.12: Example of the effect of the alteration on the selling prices and result after the refining run of the PSO.

$$\text{Ref}_{\text{cost}}^{\text{P}} = \alpha \sqrt{\left(\text{Ref}_{\text{cost}}^{(\text{p},\text{app})} - \text{Aim}_{\text{cost}}\right)^2} + \beta \text{Sell}_x^{(\text{p},\text{ref})} \quad (6.61)$$

By choosing $\alpha \gg \beta$, the swarm is driven to identify a solution equivalent to the one obtained at the end of the Aiming step in terms of proximity to the target injection profile, but with a significantly lower global cost in terms of the promised alterations. This effectively eliminates unnecessary overpricing.

The resulting set of N_u alteration profiles $\text{Sell}_{x,t}^{(\text{ref},u)}$, obtained at the end of the Refining step, is considered the optimal solution for the set $\text{Sell}_{x,t}^u$ to be used as commands to be sent to the distributed HEMSs.

The decision to split the optimization into two phases resulted from trial-and-error experimentation. Tests showed that this two-step process outperforms a single-step approach, both in terms of solution quality and computational efficiency.




Figure 6.13: Visual representation of the two PSO runs employed in the BESS coordination process, illustrating how they interact with each other.

Two additional strategies applied to sensibly reduce computation time are described as follows:

Interval independence. Given the assumption that prosumers always prioritise self-consumption, storing sufficient energy in their batteries to cover future demand regardless of market conditions, the total energy injected into the grid in each injection interval remains fixed, independent of actions taken in other intervals.

Consequently, the optimisation of each interval can be treated as an independent problem, decomposing a large, complex problem into a sequence of smaller subproblems, each focused only on the relevant variables for that interval. This approach allows the forecasting horizon to be divided into a series of intermittent injection and absorption intervals, each characterised by its own duration I_{length} . The counting of these intervals starts from the one closest to the present moment.

Asymmetric resource allocation. Due to the receding-horizon structure of the problem, the same points and intervals are optimised multiple times before actual execution. To improve efficiency, more computational effort, expressed as the maximum number of PSO iterations, is allocated to intervals closer to real-time execution, whereas fewer resources are devoted to points further in the future. This reflects both practical considerations of rolling optimisation and the real-world effect (though not addressed in this work) of increasing forecast uncertainty at longer horizons.

To preserve progress across successive temporal steps, at the start of a new optimisation cycle, 10% of the swarm particles are initialised from the best solution obtained at the end of the previous step, shifted by one temporal step. This anchoring effect stabilises the optimisation process while enabling adaptation to updated forecasts.

For both phases, the PSO hyperparameters were tuned experimentally and heuristically, and the chosen ones are shown in Tables 6.5.2 and 6.5.2, while Figure 6.13 represents the concept of the two optimisation runs and how they are connected. Some of them were kept constant, while others were adjusted to implement the two strategies described above, aimed at improving computation time.

This is particularly true for the PSO used in the *Aiming* phase, where computational resources are focused on the search for solutions within the injection intervals closer to the present moment. In both phases, the swarm size depends on the length of the interval; however, the upper bound used in the Refining phase is smaller, as this problem is easier to solve.

$$\text{Swarm}_{\text{Size}}^{\text{Aim}} = \min(\max(N_u, I_{\text{length}}), 200) \quad (6.62)$$

$$\text{Swarm}_{\text{Size}}^{\text{Ref}} = \min(\max(N_u, I_{\text{length}}), 100) \quad (6.63)$$

Table 6.1: PSO hyperparameters for the Aiming Run.

N_{Interval}	Max_{iter}	$Max_{\text{iter}}^{\text{stall}}$	Tol.	In. <i>range</i>	W^{Self}	W^{Social}
$I = 1$	30	10	0.001	[0.1, 1.1]	1.49	1.49
$I = 2$	20	8				
$I = 3$	10	8				
$I > 3$	5	5				

Table 6.2: PSO hyperparameters for the Refining Run.

Max_{iter}	Max_{iter}^{stall}	Tol.	In.range	W^{Self}	W^{Social}
5	5	0.001	[0.1, 1.1]	1.49	1.49

6.5.3 Load Coordination Through PSO Optimization

This branch is essentially constituted by the PSO for Load Coordination, carried out in a single phase. Figure 6.14 helps to recapture the details of this passage. In this case, the digital twins of the community HEMSs are included in the optimisation functional, but only for the controllable load component. The optimised variables are again the alterations Buy_x^u applied to the buying price Buy_{price}^u perceived by each user.

The objective of the CEMH in this branch is to distribute controllable loads across the time horizon to limit peaks in grid absorption and maximise self-consumption. REC simulation is first performed without considering controllable loads, obtaining $CER_{CMDflow,t}$, which allows us to identify in which hours self-consumption can be achieved and in which grid absorption will be critical. For each swarm particle, the load scheduling of all users is then performed, producing sets of grid absorption profiles due to controllable loads, $Abs_{CL,t}^{(u,p)}$.

The total grid absorption due to controllable loads for each particle is given by:

$$Abs_{CL,t}^{(tot,p)} = \sum_{u=1}^{N_u} Abs_{CL,t}^{(u,p)} \quad (6.64)$$

By combining this with the forecasted REC profile, an approximation of the community behaviour that take in account the presence of controllable loads is obtained:

$$REC_{CL,t}^p = REC_{CMDflow,t} - Abs_{CL,t}^{(tot,p)} \quad (6.65)$$

$$Abs_{CL,t}^{(REC,p)} = \min\left(REC_{CL,t}^p, 0\right) \quad (6.66)$$




Figure 6.14: Representation of the PSO run employed in the load coordination process.

The functional cost is defined as the RMSE of the REC grid absorption:

$$CL_{cost}^p = \sqrt{\frac{\sum_{t^*=1}^H \left(Abs_{CL,t^*}^{(REC,p)} \right)^2}{H}} \quad (6.67)$$

This cost function naturally reduces total grid absorption by pushing controllable loads to be scheduled during hours of forecasted REC injection. At the same time, it penalises high peaks, leading to a smoother load distribution across the horizon. The rolling-horizon framework, already described in previous branches, is also applied here.

At the end of the optimisation cycle, the result is a set of N_u alteration profiles $Buy_{x,t}^{(CL,u)}$, considered the optimal solution for the set $Buy_{x,t}^u$, and subsequently used as commands to be sent to the distributed HEMSs.

Table 6.5.3 shows the hyperparameters used for PSO in the load coordination phase:

$$Swarm_Size^{Ref} = \min(\max(N_u, I_{length}), 100) \quad (6.68)$$

Table 6.3: PSO hyperparameters for Load Coordination.

Max_{iter}	Max_{iter}^{stall}	Tol.	In.range	W^{Self}	W^{Social}
5	10	0.001	[0.1, 1.1]	1.49	1.49

6.5.4 Results

To evaluate performance, both the MILP benchmark and the proposed methodology were applied to simulated RECs over a one-year test period, performing the optimisation problem three days ahead, under the constraint of prioritising storage for self-consumption. To emulate the behaviour of users and prosumers, historical data of loads and productions were taken from the Ausgrid dataset. Meanwhile, to take into account the day-ahead energy market, the 2024 PUN history was downloaded directly from the GSE portal.

- **Ausgrid:** dataset from which the load and production profiles of prosumers from Australia are collected, with half-hour resolution. From the set, the data for a whole year were extracted and, among them, 126 profiles with fewer missing values were selected and converted to hourly-based resolution. Then these profiles were used to represent the first set of prosumers.
- **PUN:** downloaded from the GSE portal, representing the historical result of the MGP. This data represents the cost reached by an MWh of energy, hour by hour, in every zonal market in Italy. The PUN is the hourly mean of all zonal prices. For this work, the year 2024 of PUN was used.

To combine these datasets, the Ausgrid dataset must be shifted by six months to align the seasonal cycles. This adjustment ensures that both load and production profiles correspond to the appropriate price curves, resulting in a more realistic case study.

To normalise the behaviour of prosumers, all profiles were rescaled in the same way. All load profiles were adjusted to reach an annual consumption of 3 MWh. The same process was applied to the production profiles to simulate 3 kWp plants, scaling them to simulate an annual production of 3 MWh. For each prosumer, a 6 kWh BESS was considered.

For pure users, the presence of two controllable loads was also considered. The first has a consumption of 0.5 kWh, schedulable once per day within the time window from 8 a.m. to 6 p.m. The second has a request of 1 kWh, schedulable once every two days within the time window from 3 p.m. to 11 p.m.

The algorithms were implemented in MATLAB R2024a and executed on a machine equipped with an Intel Core i9-13900K processor (Santa Clara, CA, USA) and 32 GB of RAM. The simulated RECs consisted of 100 domestic users, including 60 pure consumers and 40 prosumers. User profiles were randomly drawn from the Ausgrid dataset until the required community size was reached. Three optimisation approaches were tested under the different REC scenarios:

- The MILP benchmark, used as the golden standard to provide the best achievable results in terms of cash flow;
- Distributed HEMSs without CEMH, representing the baseline behaviour of the REC in the absence of coordination;
- Distributed HEMSs with active CEMH, i.e., the proposed methodology.

In all tests, the results showed that the performance of the defined benchmark are consistent with what is presented in the literature, taking an average time of 20 s to solve the problem in each single iteration, leading to an improvement in the self-consumption of the community of 35% on average with respect to the case without any coordination. In Figure 6.15, an example of a full year of energy exchanges between the REC and the grid is shown.

It can be observed that, in the case of active coordination, most of the peaks are shaved, both for the EMH ON and MILP cases. Figure 6.16 shows the same example; however, this time, the curves are not overlapped, allowing for a clearer comparison of the effects of the two different techniques.

Therefore, the tests show that the introduction of the CEMH, in place of the MILP, improved the REC performance to a similar extent, both in terms of the regularity of the community profile and the incentives collected. The latter parameter is easier to quantify and provides an indirect indication of the quality of coordination. The proposed approach achieved results within the range of $8_{-1.6}^{+6.2}$ % of the optimum obtained by the MILP benchmark, with a confidence of 95%.

In Figure 6.17, an example of the trend of self-consumption incentive is presented, illustrating the resulting REC behaviour under different coordination approaches. Table 4 instead reports the results obtained for all the RECs tested, using the following definitions:




Figure 6.15: Visual comparison of the energy exchange between REC and Grid in case of different methodology. The profiles are shown overlapped.




Figure 6.16: Visual comparison of the energy exchange between REC and Grid in case of different methodology. The profiles are shown on different plots.

$$\text{Inc}_{\text{MILP}}^{\text{abs}} = \text{Inc}_{\text{MILP}}^{\text{final}} - \text{Inc}_{\text{CEMH.off}}^{\text{final}} \quad (6.69)$$

$$\text{Inc}_{\text{CEMH.on}}^{\text{abs}} = \text{Inc}_{\text{CEMH.on}}^{\text{final}} - \text{Inc}_{\text{CEMH.off}}^{\text{final}} \quad (6.70)$$

$$\text{Inc}_{\text{MILP}}^{\text{rel}} = \frac{\text{Inc}_{\text{MILP}}^{\text{abs}}}{\text{Inc}_{\text{CEMH.off}}^{\text{final}}} \quad (6.71)$$

$$\text{Inc}_{\text{CEMH.on}}^{\text{rel}} = \frac{\text{Inc}_{\text{CEMH.on}}^{\text{abs}}}{\text{Inc}_{\text{CEMH.off}}^{\text{final}}} \quad (6.72)$$

The results are expressed as distributions and, due to the limited number of tests, the standard deviation s was evaluated using the Bessel correction. The chi-square distribution was then used to estimate the boundaries of the confidence interval. Here, x_i represents the observed value, $\bar{x}$ is the mean, n is the number of observations, and $\text{df} = n - 1$ denotes the degrees of freedom of the chi-square distribution.

$$s = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2} \quad (6.73)$$

$$\sigma_{1-\alpha} = \left[\sqrt{\frac{\text{df } s^2}{\chi_{1-\alpha/2, \text{df}}^2}}, \sqrt{\frac{\text{df } s^2}{\chi_{\alpha/2, \text{df}}^2}} \right] \quad (6.74)$$

Table 6.4: Reward from self-consumption of the tested RECs and averages.

REC	1	2	3	4	5	6	Avg. $\pm \sigma_{95\%}$
$\text{Inc}_{\text{MILP}}^{\text{final}} (\text{€})$	3052	3102	3042	2961	3128	3084	$3061 \pm_{37}^{144}$
$\text{Inc}_{\text{CEMH.on}}^{\text{final}} (\text{€})$	2936	2928	2779	2789	2926	2951	$2914 \pm_{49}^{193}$
$\text{Inc}_{\text{CEMH.off}}^{\text{final}} (\text{€})$	2301	2287	2163	2120	2254	2281	$2234 \pm_{47}^{183}$
$\text{Inc}_{\text{MILP}}^{\text{abs}} (\text{€})$	751	816	880	841	873	803	$827 \pm_{30}^{117}$
$\text{Inc}_{\text{CEMH.on}}^{\text{abs}} (\text{€})$	635	641	616	669	671	670	$650 \pm_{14}^{56}$
$\text{Inc}_{\text{MILP}}^{\text{rel}} (\%)$	32.6	35.7	40.7	39.7	38.7	35.2	$37.1 \pm_{1.9}^{7.8}$
$\text{Inc}_{\text{CEMH.on}}^{\text{rel}} (\%)$	27.6	28.0	28.5	31.6	29.8	29.4	$29.1 \pm_{0.9}^{3.5}$
$\text{Gap}_{\text{CEMH.on}}^{\text{MILP}} (\%)$	5.0	7.6	12.2	8.1	9.0	5.9	$8.0 \pm_{1.6}^{6.2}$




Figure 6.17: Visual comparison of the cash flow obtained after one year of usage of the different optimization techniques on the same CER.

All results were represented with a confidence level of 95% ($\alpha = 0.05$), evaluating the chi-square distribution with $df = 5$. The following values are used:

$$\chi_{0.975,5}^2 \approx 12.833, \quad \chi_{0.025,5}^2 \approx 0.831. \quad (6.75)$$

In Figure 6.18, the results in terms of $\text{Gap}_{\text{CEMH}_{\text{on}}}^{\text{MILP}}$ are presented throughout the year in a box plot, allowing an evaluation of how the difference in quality between the two methods varies throughout the year.

However, such improvement in performance is not cost-free. Coordination imposes a burden on users, particularly prosumers, who must reschedule the operation of their BESS to sell energy at different, and sometimes less favourable, hours. Since these lost opportunities are the result of deliberate choices made for collective benefit, it is essential that the members who take them are adequately rewarded for their collaboration and compensated for their losses.

In this regard, tracking the price alterations generated by the proposed method provides a fair and real-time evaluation of each user's contribution, facilitating redistribution of incentives and rewarding those who prioritise community objectives over individual gains.

These values ultimately make it possible to estimate the lower bound of




Figure 6.18: Boxplot of $\text{Gap}_{CEMH_{on}^{MILP}}$ during the year.

the incentive percentage that should be redistributed among worthy users, while the remaining share can be retained by the REC and used at its own discretion.

Defining $\text{Users}_{\text{loss}}^{\text{final}}$ as the difference between the cumulative cash flow of all REC members with and without CEMH, $\text{Users}_{\text{Sell rew}}^{\text{final}}$ as the total incentive allocated to reward users who shift their injection scheduling, and $\text{Users}_{\text{Buy rew}}^{\text{final}}$ as the cumulative incentive used to reward users who modify their controllable load scheduling, the results of the rewards evaluation are shown in Table 6.5.

Each reward is evaluated by valuing the energy exchanged by the users using the historical price alterations computed through the CEMH coordination, instead of the original market prices. Let $x_t^{(u, \text{hist})}$ denote the historical prosumption of user u at time t , and let $\text{Sell}_{x,t}^{(u, \text{hist})}$ and $\text{Buy}_{x,t}^{(u, \text{hist})}$ denote, respectively, the historical sell- and buy-price alteration signals assigned to user u . We can then define the rewards, $\forall t \in \text{hist}$, as follows:

$$\text{Sell}_{\text{Rew},t}^u = x_t^{(u, \text{hist})} \cdot \text{Sell}_{x,t}^{(u, \text{hist})} \quad (6.76)$$

$$\text{Buy}_{\text{Rew},t}^u = x_t^{(u, \text{hist})} \cdot \text{Buy}_{x,t}^{(u, \text{hist})} \quad (6.77)$$

$$\text{Sell}_{\text{Rew}}^{(u, \text{final})} = \sum_{t \in \text{hist}} \text{Sell}_{\text{Rew},t}^u \quad (6.78)$$

$$\text{Buy}_{\text{Rew}}^{(\text{u},\text{final})} = \sum_{t \in \text{hist}} \text{Buy}_{\text{Rew},t}^{\text{u}} \quad (6.79)$$

$$\text{Users}_{\text{Sell rew}}^{\text{final}} = \sum_{u \in N_u} \text{Sell}_{\text{Rew}}^{(\text{u},\text{final})} \quad (6.80)$$

$$\text{Users}_{\text{Buy rew}}^{\text{final}} = \sum_{u \in N_u} \text{Buy}_{\text{Rew}}^{(\text{u},\text{final})} \quad (6.81)$$

It should be noted that such evaluations are easy to perform either at the end of each optimisation cycle in real time or over longer aggregation periods, such as weekly or monthly intervals.

Table 6.5: Users' losses due to coordination and evaluated rewards.

REC	1	2	3	4	5	6	Avg. $\pm \sigma_{95\%}$
$\text{Users}_{\text{loss}}^{\text{final}} (\text{€})$	-93.6	-104.7	-94.5	-103.1	-102.3	-69.3	-94.6 $\pm 32.5_{8.3}$
$\text{Users}_{\text{Sell rew}}^{\text{final}} (\text{€})$	235.2	214.2	221.4	219.1	237.1	240.5	228.2 $\pm 27.0_{6.9}$
$\text{Users}_{\text{Buy rew}}^{\text{final}} (\text{€})$	89.6	89.6	89.0	88.8	90.0	89.0	89.3 $\pm 1.1_{0.3}$

Another key aspect is the computational cost. As mentioned in the introduction, MILP formulations tend to grow rapidly in size and require significant computational resources. The MATLAB implementation required sparse matrix formulations and, more importantly, exhibited longer execution times. The main tests showed that MILP required, on average in the case studies, roughly three times longer to complete an optimisation cycle compared to the proposed method, for communities structured as described above.

The computation time depends on the number of users N_u and the proportion of prosumers within the community. This trend was further investigated and confirmed by additional tests in which the number of community members N_u varied from 50 to 1000 users, and the prosumer-to-consumer ratio PC_{ratio} was adjusted from 20% to 80% of the total.

The composition of the community was randomly generated and then tested for a full day of operation, which is a total of 24 iterations. The rolling horizon of the optimisation was kept to three days ahead, i.e., $H = 72$ h.

Tables 6.5.4 and 6.5.4 present the results of varying N_u and PC_{ratio} with respect to the average optimisation time, allowing a direct comparison between the two methodologies. The number of constraints increases linearly with the number of users.

Talking about the number of inequality constraints, which mainly define the controllable loads possessed only by pure consumers, it decrease as the percentage of prosumers increases. It should also be noted that a higher share of prosumers implies that more variable bounds are relaxed, such as those related to battery operations, allowing MILP to explore a larger feasible region and thereby increasing the overall complexity of the problem.

Table 6.6: Average time cost of the proposed method to complete one optimization iteration.

	PC_ratio = 20%	40%	60%	80%
$N_u = 50$	5.46 s	5.85 s	4.85 s	4.55 s
$N_u = 100$	5.53 s	6.44 s	8.91 s	8.21 s
$N_u = 500$	77.68 s	100.03 s	113.14 s	124.14 s
$N_u = 1000$	175.18 s	213.24 s	234.30 s	254.34 s

Table 6.7: Average time cost of MILP benchmark per iteration and number of inequality and equality constraints.

	PC_ratio = 20%	40%	60%	80%
$N_u = 50$	3.40 s 604×72 h	7.11 s 294×72 h	5.80 s 604×72 h	48.88 s 234×72 h
$N_u = 100$	4.96 s 1204×72 h	10.79 s 584×72 h	31.02 s 1204×72 h	704.23 s 464×72 h
$N_u = 500$	88.52 s 6004×72 h	294.08 s 2904×72 h	407.92 s 6004×72 h	3897.4 s 2304×72 h
$N_u = 1000$	483.66 s $12\,004 \times 72$ h	1048.1 s 5804×72 h	1077.6 s $12\,004 \times 72$ h	7242.5 s 4604×72 h

6.5.5 Discussion

As expected, the MILP benchmark outperformed the proposed methodology in terms of solution optimality. Nevertheless, the proposed approach achieved results sufficiently close to the theoretical optimum to justify a discussion of the trade-offs involved. As recognised in the literature, real-time energy management must operate under strict time constraints, where obtaining a near-optimal solution within an acceptable time is preferable to computing the true optimum too late for practical implementation.

The proposed hybrid framework, which combines a PSO metaheuristic with a rule-based control layer, was specifically designed for real-time deployment with particular attention to both computational efficiency and hardware constraints. Its light implementation and limited computational demand make it suitable for edge-computing environments and large-scale community deployment.

Beyond performance efficiency, the proposed coordination mechanism introduces a transparent price-based reward system that allows the fair compensation of individual users contributing to community objectives at a personal cost. This transparency and fairness in value sharing are essential design choices, recognising that Renewable Energy Communities depend on the active cooperation of real human participants.

By ensuring that collective optimisation does not penalise individual members, the proposed methodology fosters trust and increases user acceptance in delegating part of their control to the EMS. It is important to note, however, that the results shown were obtained by applying the proposed method to a specific case study.

Although the methodology is expected to be able to approach the optimum under different conditions, the value of $\text{Gap}_{\text{CEMH}_{\text{on}}}^{\text{MILP}}$ may vary in terms of percentage points. Further analysis, applying the proposed method to different configurations, will be carried out in future works.

A further step toward real-world implementation of the proposed architecture will involve addressing the communication layer and validating the system in operational environments and the reliance on perfect forecasting. Unfortunately, in real-world scenarios, this assumption is unrealistic and the impact of uncertainty on performance will be explicitly evaluated. The implementation of techniques such as robust optimisation and stochastic model predictive control has already been evaluated for integration both at the HEMS and within the digital twin environment of the CEMH, to effectively address this uncertainty.

Overall, the proposed framework represents a pragmatic and scalable step toward the realisation of resilient, fair and human-centric Energy Management Systems for next-generation Renewable Energy Communities.

6.6 Load Shedding and Resilience

Load shedding is one of the key functionalities that can be implemented at the secondary control level of an energy management architecture. The underlying idea is to modulate the load requested by users according to the amount of energy available within the system and the maximum power that can be safely drawn from the grid. In practical terms, load shedding acts as a security layer that selectively disconnects the minimum amount of electrical loads needed to prevent the system from exceeding its operational limits.

This mechanism plays a crucial role in enhancing the overall resilience of the grid. When the capacity of the grid is exceeded, the typical consequence is a cascading failure leading to a widespread blackout and, therefore, the complete loss of all loads served. By contrast, a controlled load-shedding strategy mitigates this risk by ensuring that only a limited subset of loads is temporarily curtailed, allowing the system to remain operational.

In this sense, resilience can be quantified through metrics such as the Load Served Ratio (LSR), defined as the ratio between the effectively supplied load and the load originally requested by users, and is also one of the core components of the broader System Resilience Index (SRI), widely used in the analysis of smart grids and microgrids. In this work, we focus specifically on improving the LSR, as it directly captures the effectiveness of the shedding mechanism in maintaining service continuity.

$$\text{LSR} = \frac{\text{Load Served}}{\text{Load Requested}} \quad (6.82)$$

Furthermore, a hierarchical or priority-based load-shedding scheme can be implemented to determine the order in which the loads are disconnected. This introduces a structured and socially meaningful approach to resilience. For example, in a residential smart grid, critical loads such as hospitals or hydraulic pumping stations would be classified as high-priority assets and disconnected only under extreme conditions. Conversely, non-critical residential loads would be the first candidates for curtailment, ensuring that essential services remain operational for as long as possible.

In our work, this mechanism has been incorporated as part of the CEMH functionality and is executed prior to the scheduling process. This ensures that the optimisation phase begins from a feasible operating point, where the risk of violating grid constraints has already been mitigated through controlled and prioritised load reduction.

When CEMH schedules the shedding, it does so over the entire duration of the forecasted horizon, and the process occurs in two steps. In the first step, all internal production is omitted, and the cumulative load profile is compared with the grid's availability to evaluate a preliminary shedding plan. The second step involves the use of a PSO, which starts from both the solution obtained in the previous CEMH iteration and the one found at the end of step one.

Load shedding has been implemented in a discrete form, and each user has been characterised by:

- a shedding priority level $Shedd_p^u$,
- the minimum load level that must remain available before complete disconnection $Shedd_{BS_{lvl}}^u$,
- the total number of shedding levels $Shedd_{N_{lvl}s}^u$.

6.6.1 Research of a Preliminary Solution

To perform the planning, users have been divided according to their $Shedd_p^u$, with the idea of handling users with the same level of priority in the same way. The planning is performed over the whole forecasting horizon, and the total number of variables of the problem is equivalent to the horizon length multiplied by the number of different priority levels in the system:

$$N_{Shedd}^{var} = H \cdot \max(Shedd_p^u) \quad (6.83)$$

Given the requested load profiles of each user of the community, $Load_{req,t}^u$, for the entire forecasting horizon, each variable LS_t^p represents the percentage of load that will be supplied to all users belonging to the priority class p , above their minimum guaranteed level $Shedd_{BS_{lvl}}^u$. If at any time LS_t^p is set equal to zero, this means that the user will be completely shut down.

The remaining load above the base service level is computed as:

$$Load_{rem,t}^u = Load_{req,t}^u - Shedd_{BS_{lvl}}^u \quad (6.84)$$

The resulting supplied load after shedding is then defined as:

$$Load_{shed,t}^u = \begin{cases} LS_t^p \cdot Load_{rem,t}^u + Shedd_{BS_{lvl}}^u & \text{if } LS_t^p > 0 \\ 0, & \text{if } LS_t^p = 0. \end{cases} \quad (6.85)$$

Given a set that covers the entire forecast horizon, the Grid Availability profile $Grid_t^{Ava}$ represents the maximum amount of energy that can be supplied by the grid at each time step. The methodology implemented to obtain a first candidate solution is straightforward: acting step by step, and starting from the lowest priority class, the value of LS_t^p is iteratively decreased by one unit until the cumulative load demand of the community is reduced sufficiently to not exceed the grid's availability at time t .

$$Load_{shedd,t}^{CER} = \sum_{u=1}^{N_{users}} Load_{shedd,t}^u \quad (6.86)$$

$$Load_{shedd,t}^{CER} < Grid_t^{Ava} \quad (6.87)$$

This procedure is repeated for every time step of the forecast horizon, generating the complete set of shedding levels for all user categories over the entire sequence of time steps.

6.6.2 Refining of the Solution through PSO

To properly account for the presence of local production units and batteries, the PSO architecture relying on the Digital Twins framework within the functional has been reintroduced. Since PSO is an optimisation approach in which the variables evolve in a continuous space, preserving the discrete nature of the shedding levels requires discretising the variables inside the functional at each iteration.

During swarm initialisation, 10% of the particles are set equal to the simplified solution, which represents a good approximation of the minimal feasible shedding level. Another 10% is initialised using the solution obtained in the previous CEMH iteration, shifted by one hour, and an additional 10% is set to the maximum shedding level for each priority group and each time step. The remaining particles are initialised randomly across the feasible space. This 30% of structured initialisation provides a robust starting set for the search process and significantly speeds up convergence.

In this case, the entire forecast horizon is optimised in a single run, and each particle contains N_{Shedd}^{var} decision variables. At every PSO iteration, particle values are rounded to the nearest integer to preserve the discrete

structure of the shedding levels, and the load profiles of the users are modulated accordingly using Eq. 6.85. Then, the production profiles of the prosumers $Prod_t^u$ are incorporated and the complete energy profiles of all users $Energy_t^u$ are computed.

$$Energy_t^u = Prod_t^u - Load_{shed,t}^u \quad (6.88)$$

Finally, the digital twins of the community are used to assess how these profiles evolve in the presence of the BESS units or after the implementation of controllable commands in general $Energy_{CMD,t}^u$.

Summing all the generated user profiles yields a candidate REC-level profile.

$$Energy_{CMD,t}^{CER} = \sum_{u=1}^{N_{users}} Energy_{CMD,t}^u \quad (6.89)$$

The cost function is designed to reflect two main objectives:

1. Feasibility with respect to grid availability:

No time step of the sequence may exceed the corresponding grid availability. If this condition is violated, the entire sequence is considered unfeasible, and a heavy penalty is added to the cost function for each violating time step.

$$Obj_{points}^{Unfeasible} = \sum_{t=1}^H (Energy_{CMD,t}^u > Grid_t^{Ava}) \quad (6.90)$$

2. Maximisation of supplied load based on priority:

A sequence is considered *better* when more demand is actually supplied from higher-priority users. The PSO is therefore guided toward solutions that respect feasibility while maximising the supply of loads with higher priority.

$$Obj_{supply}^{Max} = \sum_{t=1}^H \sum_{u=1}^{N_{users}} Shed_d^u \cdot LS_t^p \quad (6.91)$$

Then, the two objectives are linearly combined to obtain a single formulation that represents the cost function. Note that the first objective, which evaluates the feasibility of the points, must always be satisfied. This means

that it should be significantly more weighted than the second one. The resulting formulation of the cost function is as follows:

$$\text{Shedd}_{\text{cost}} = W_1 \cdot \text{Obj}_{\text{points}}^{\text{Unfeasable}} + W_2 \cdot \text{Obj}_{\text{supply}}^{\text{Max}} \text{ with } W_1 \gg W_2 \quad (6.92)$$

6.6.3 Preliminary Result

The presence of this decision block shows an important impact on the resilience of the whole system. A simulation has been run on a community of 100 domestic users, of which 25 are defined as prosumers and, therefore, equipped with local production plants. The 75 pure consumers have been divided into three priority groups. The prosumers have been placed in the second priority level. This first test simply shows how essential it is to selectively shed the load demand in the community, or even in a smart grid, to avoid unnecessary full blackouts.

The simulation has been carried out over a 10 day period, corresponding to 240 hourly steps $N_{\text{step}}^{\text{test}}$, considering both a scenario in which load shedding is present and one in which the community is not reactive to an evident shortage of available energy. In the simulation, if at any time step the total request from the grid exceeds its capacity, it is assumed that the grid goes into protection mode and does not supply the community for the entire time step.

It is assumed, to simplify the study, that the grid is able to return to full functionality as soon as the constraint on availability is respected again, without considering the normal operational delays due to restoration of nominal operating conditions or black start procedures. Note that this hypothesis actually underestimates the importance of avoiding grid overload, as it reduces the practical impact of blackout occurrences.

The Load Satisfaction Ratio (LSR) achieved by the community has been evaluated hourly over the full experiment window:

$$LSR_t = \frac{\sum_{u=1}^{N_{\text{users}}} \text{Load}_{\text{shedd},t}^u}{\sum_{u=1}^{N_{\text{users}}} \text{Load}_{\text{req},t}^u} \quad (6.93)$$

We defined:

- **Hourly Average LSR (ALSR)**: the mean of the LSR evaluated at

each step over the entire time horizon of the test;

$$ALSR = \frac{1}{N_{step}^{test}} \sum_{t=1}^{N_{step}^{test}} LSR_t \quad (6.94)$$

- **Total LSR (TLSR)**: the LSR evaluated on the total energy requested and supplied over the test window;

$$TLSR = \frac{\sum_{t=1}^{N_{step}^{test}} \sum_{u=1}^{N_{users}} Load_{shedd,t}^u}{\sum_{t=1}^{N_{step}^{test}} \sum_{u=1}^{N_{users}} Load_{req,t}^u} \quad (6.95)$$

It is also possible to include the priority of the loads in the evaluation of the metrics. In this case, the metric becomes the Priority-Weighted Load Served Ratio (PW-LSR), defined as:

$$PWLSR_t = \frac{\sum_{u=1}^{N_{users}} Shedd_p^u Load_{shedd,t}^u}{\sum_{u=1}^{N_{users}} Shedd_p^u Load_{req,t}^u} \quad (6.96)$$

Then it is possible to define:

- **Hourly Average PWLSR (APWLSR)**: the mean of PWLSR evaluated at each step of the whole test horizon;

$$APWLSR = \frac{1}{N_{step}^{test}} \sum_{t=1}^{N_{step}^{test}} PWLSR_t \quad (6.97)$$

- **Total PWLSR (TPWLSR)**: the PWLSR evaluated on the total requested and supplied energy;

$$TPWLSR = \frac{\sum_{t=1}^{N_{step}^{test}} \sum_{u=1}^{N_{users}} Shedd_p^u Load_{shedd,t}^u}{\sum_{t=1}^{N_{step}^{test}} \sum_{u=1}^{N_{users}} Shedd_p^u Load_{req,t}^u} \quad (6.98)$$

Even without resorting to the priority-weighted formulation, it is already sufficiently clear how strongly the coefficient related to the satisfied load demand increases when load shedding is enabled. In the test window, the community's load demand reaches daily peaks of about 50 kWh, while, in our scenario, it is assumed that the grid can supply a maximum of 20 kWh. In such cases, even with the simplified *soft* blackout model adopted, the case

without load shedding fails badly to supply its members, and the LSR drops significantly.

Even if preliminary, the results achieved clearly show how community-level coordination can be impactful in regulating the interaction between a Renewable Energy Community and the grid. Through coordinated load shedding, a community can preserve the operability of key assets and critical services while avoiding a complete blackout. This coordinated behaviour, although potentially requiring temporary sacrifices from some users, ultimately benefits all participants. Even users whose loads are fully curtailed gain an indirect advantage, as a complete blackout would affect the entire community indiscriminately.

Although recovery time has not been explicitly considered in this study, it is reasonable to assume that, in the event of a grid collapse, the resulting disruption would be both longer and more severe than any partial and temporary shedding strategy. From this perspective, controlled and prioritised load shedding represents a form of collective, altruistic behaviour that enhances the overall resilience of the system. By preventing grid protection events and maintaining partial operability, the proposed approach shows how coordinated demand management can significantly reduce the severity of extreme events, laying the groundwork for more advanced resilience-oriented control strategies.




Figure 6.19: LSR metrics on the Example without active Load Shedding




Figure 6.20: LSR metrics on the Example with active Load Shedding

6.7 Conclusion and Final Remarks

This chapter presented a coordinated energy management framework for Renewable Energy Communities based on a hierarchical and distributed architecture. The proposed approach combines local autonomy at the HEMS level with centralised coordination performed by a Community Energy Management Hub, enabling the alignment of individual user objectives with community-level goals.

Coordination is achieved through indirect, price-based signals and meta-heuristic optimisation supported by digital twins, allowing scalable operation under realistic computational and communication constraints. The results showed that the proposed methodology can achieve performance close to that of a centralised MILP benchmark while maintaining compatibility with real-time operation.

In addition to coordination, a priority-based load-shedding mechanism was introduced at the secondary control level to enhance system resilience under grid availability constraints. Although this mechanism represents a preliminary contribution, the simulations highlighted its effectiveness in avoiding unnecessary full blackouts and in increasing the fraction of satisfied demand, particularly for higher-priority loads. Metrics such as the Load Satisfaction Ratio and its priority-weighted extension provided a clear and quantitative evaluation of the resilience improvements introduced by selective load curtailment.

While this chapter focused on coordination at the community level, the proposed architecture naturally fits within a broader, multi-layer control framework. The next chapter extends this perspective by addressing the tertiary control layer, where the interaction between multiple communities and the transmission network is considered.

In particular, a controller based on Optimal Power Flow is introduced to manage large-scale network constraints, enabling the coordination of Renewable Energy Communities with transmission-level objectives and ensuring system-wide feasibility and security.

Chapter 7

Scaling Up Control: Transmission-Level EMS

*This chapter presents a comprehensive framework for the modelling, simulation and operational management of a transmission network through the integration of a Digital Twin and a Linearised Optimal Power Flow based Operational Management System (OMS) of the Ethiopian Transmission Network (E-TN). The proposed methodology enables the detailed analysis of system behaviour under normal and stressed operating conditions, with a particular focus on resilience, flexibility, and the coordinated exploitation of renewable energy resources. The Digital Twin rely on an Optimal Power Flow (OPF) to provide a high-fidelity representation of E-TN. On top of this representation, the OMS operates as the highest decision-making layer. To ensure computational tractability while preserving sensitivity to network effects, the OMS formulates the optimization problem using a linearized OPF. This framework allows to investigate several critical case studies impossible to replicate on the physical system, including transmission line failures, generation outages and long-term demand growth scenarios and to evaluate the system response.*¹

¹The study presented in this chapter was conducted while the author was a visiting Ph.D. student at Aalborg University, Aalborg (Denmark) from April to July 2025 (working with. Prof. Yajuan Guan, Dr. Wenfa Kang and Dr. Mahshid Javidsharifi)

7.1 Introduction

In the previous chapters, the architecture of the lower-level management systems has been presented. The HEMS was designed as a lightweight controller capable of managing the needs of individual users with minimal computational effort. At the second layer, the CEMH was introduced to address issues arising at the community and distribution-network level. The core idea behind this architecture was to enable coordination among decentralized controllers without centralizing all the “intelligence” of the system. Here as well, the priority was to propose a methodology that remains simple, scalable, and easy to implement.

At the elementary level, the HEMS receives electrical measurements from the monitoring system and energy prices from the outside world. The secondary layer, instead, operates using forecasts provided by the HEMSs and higher-level objectives coming from external actors. To allow interaction between these layers, the CEMH influences the behavior of individual HEMSs by modifying the prices they perceive, effectively steering their decisions toward community-level objectives.

At this point, it becomes natural to ask how such community-wide objectives should be defined, and how a hypothetical third-level management system may coordinate multiple CEMHs at the scale of a distribution network, as shown in Figure 7.1. Essentially, two inputs can be used to influence the logic of the second layer: the desired shape of the community’s aggregated power profile and the availability of grid energy, which may trigger load shedding when required.

To properly integrate the transmission-level grid constraints into this hierarchical framework, the natural methodology to adopt is the Optimal Power Flow (OPF). OPF enables the evaluation of how different scheduling or generation patterns impact voltage levels, line loadings, and overall system feasibility, making it the appropriate tool for defining higher-level objectives.

The methodology presented in the following sections has been developed in collaboration with the CROM research group at Aalborg University and is based on the topology of the Ethiopian transmission network. This work was carried out within the “Last Wind” project [91], in which the Ethiopian TSO requested support from the university to identify optimal locations for future wind farms in order to increase the resilience of the national power system.




Figure 7.1: Third Level EMS Concept

7.2 Power Flow Analysis

Power flow analysis is one of the fundamental tools for understanding and operating modern electrical power systems. Although traditionally associated with transmission system studies, its relevance extends increasingly toward distribution networks and community-level management, especially in the presence of high renewable penetration and distributed control architectures.

A power flow model provides a steady-state snapshot of the electrical grid, describing how real and reactive power move across the network, how voltage magnitudes and angles are distributed among buses, and how heavily each transmission line is loaded.

In the context of this work, power flow analysis plays a bridging role. The previous chapters focused on decentralized decision layers, HEMS and CEMH, designed to shape local consumption and community-level profiles. However, when such local actions are scaled up or coordinated across multiple communities, their aggregated effects inevitably interact with the physical constraints of the transmission grid.

This makes power flow analysis not only a computational tool, but a necessary foundation for defining feasible system-wide objectives and for understanding how high-level EMS strategies propagate through the network.

As the following sections will show, this framework becomes particularly

relevant when evaluating or designing optimal strategies on real-world transmission systems such as the Ethiopian grid, where rapid expansion of renewable capacity and geographically heterogeneous load centers require precise and reliable network assessments.

The objective of this subsection is not to provide a comprehensive derivation of power flow theory, but to summarise the core concepts taken from [] required to understand the modelling choices presented later.

Problem Definition

A power system, like the one in Figure 7.2 can be represented as a set of buses connected through transmission lines and transformers. Each bus i is characterised by four quantities: real power P_i , reactive power Q_i , voltage magnitude V_i , and voltage angle ϑ_i . The power quantities may represent either a demand or an injection, depending on their sign.

Each line ij , representing the branch that connects buses i and j , is characterised by its impedance z_{ij} , composed of the resistance r_{ij} and the reactance x_{ij} , and by its admittance y_{ij} , composed of the conductance g_{ij} and the susceptance b_{ij} . The power flow components, P_{ij} and S_{ij} , represent the active and complex power flows on the branch ij .

Solving the power flow problem means determining V_i , ϑ_i , P_i , and Q_i for each bus i . As in general $P_{ij} \neq P_{ji}$ and $Q_{ij} \neq Q_{ji}$ due to the losses on the lines, the power flows injected in each line ij are defined as

$$P_{ij}^{inj} = \max(|P_{ij}|, |P_{ji}|) \quad (7.1)$$

$$Q_{ij}^{inj} = \max(|Q_{ij}|, |Q_{ji}|) \quad (7.2)$$

Addressing these tasks requires solving a nonlinear system of equations derived from Kirchhoff's laws and the network admittance matrix.

Buses Definition

Depending on which quantities are known and which must be solved for, buses are classified into three categories:

- Slack (or reference) bus: voltage magnitude and angle are specified; this bus balances system losses.
- PV buses (generator buses): real power and voltage magnitude are specified, while Q and ϑ are unknown.




Figure 7.2: Representation of a power system

- *PQ* buses (load buses): real and reactive power consumptions are specified, while V and ϑ are unknown.

When the network model is defined, at least one bus must be designated as the Slack (or reference) bus. This node sets the system voltage angle reference and absorbs any mismatch between scheduled injections and the actual solution of the power flow equations. If required by the topology or by numerical robustness, multiple slack buses can be assigned, allowing the error compensation to be distributed more evenly across the system.

Admittance Matrix

The Admittance matrix contain the full topology and electrical characteristic of the grid. For an electrical network with N_b buses, the relationship between bus currents and voltages is:

$$I = YV \quad \Longleftrightarrow \quad I_i = \sum_{j=1}^{N_b} Y_{ij} V_j \quad (7.3)$$

where $I = [I_1, I_2, \dots, I_{N_b}]^T$ is vector of injected currents, $V = [V_1, V_2, \dots, V_{N_b}]^T$ is vector of bus voltages and Y is bus admittance matrix (Y-bus), of $N_b \times N_b$ dimension

Defining y_{ij} as the admittance of the line between the buses i and j and




Figure 7.3: Bus Definition

y_i^{sh} the shunt admittance toward earth of each bus i , is possible to build the admittance matrix as:

- Diagonal elements Y_{ii}
 $Y_{ii} = \sum_{k \neq i} y_{ik} + y_i^{sh}$
 the diagonal elements represents how strongly bus i is connected to the network.
- Off-diagonal elements Y_{ij} (for $i \neq j$)
 $Y_{ij} = -y_{ij}$
 If there is no direct line between i and j , then $Y_{ij} = 0$ The negative sign indicates that a voltage increase at bus j causes a current flowing out of bus i .

Power Flow Equations

Starting from Kirchhoff's Current Law (KCL), the complex power injected at bus i is:

$$S_i = P_i + jQ_i = V_i \sum_{j=1}^N V_j^* Y_{ij}^*, \quad (7.4)$$

where V_j^* and Y_{ij}^* denote the complex conjugates of the bus voltage and bus admittance, respectively. Keep in mind that Y_{ij} is the i, j -th element of the admittance matrix.

$$Y = \begin{bmatrix} Y_{11} & -y_{12} & -y_{13} \\ -y_{12} & Y_{22} & -y_{32} \\ -y_{13} & -y_{32} & Y_{33} \end{bmatrix}$$

Figure 7.4: Example of Admittance matrix of a three-bus system. The diagonal terms are defined as:

$$Y_{11} = y_{12} + y_{13} + y_1^{sh}, \quad Y_{22} = y_{12} + y_{23} + y_2^{sh}, \quad Y_{33} = y_{13} + y_{23} + y_3^{sh}.$$




Figure 7.5: Admittance in the power system

Expanding into real and imaginary parts yields the standard power flow equations for characterization of the bus i :

$$P_i = V_i \sum_{j=1}^N V_j (G_{ij} \cos \theta_{ij} + B_{ij} \sin \theta_{ij}), \quad (7.5)$$

$$Q_i = V_i \sum_{j=1}^N V_j (G_{ij} \sin \theta_{ij} - B_{ij} \cos \theta_{ij}), \quad (7.6)$$

with

$$\theta_{ij} = \theta_i - \theta_j.$$

Once the buses are got defined is possible to evaluate the component of the power that flow on the branches ij with the following equations:

$$P_{ij} = V_i^2 g_{ij} - V_i V_j (g_{ij} \cos \theta_{ij} + b_{ij} \sin \theta_{ij}) \quad (7.7)$$

$$Q_{ij} = -V_i^2 b_{ij} - V_i V_j (g_{ij} \sin \theta_{ij} - b_{ij} \cos \theta_{ij}) \quad (7.8)$$

It must kept in mind that in 7.7 and 7.8 g_{ij} and b_{ij} are the once that characterize the branch and not the element of the admittance matrix.

Per-Unit System

The per-unit (p.u.) system is a normalized representation widely used in power flow and power system analysis. By expressing electrical quantities with respect to predefined base values, the per-unit formulation offers several advantages:

- **Simplifies numerical computation:** avoids extremely large or small values, improving numerical stability and algorithm convergence.
- **Voltage-level independence:** makes quantities comparable across different voltage levels.
- **Uniform component representation:** lines, transformers, and machines are expressed on the same scale.
- **Well-scaled Y-bus matrix:** all admittances become dimensionless and numerically consistent.

The general definition of a per-unit quantity is:

$$\text{p.u. value} = \frac{\text{actual value}}{\text{base value}}.$$

Typical base quantities include:

$$S_{\text{base}}, \quad V_{\text{base}}, \quad Z_{\text{base}} = \frac{V_{\text{base}}^2}{S_{\text{base}}}, \quad I_{\text{base}} = \frac{S_{\text{base}}}{\sqrt{3} V_{\text{base}}}.$$

One of the most remarkable properties of the per-unit representation is its effect on transformer modelling. When all quantities are converted to a common per-unit base transformer turns ratios effectively disappear from the model.



Figure 7.6: Representation of a transformer in p.u.

The normalization implicitly accounts for the voltage scaling introduced by transformers, eliminating the need to include explicit turns-ratio factors in the power flow equations. This greatly simplifies the modelling of multi-level networks and contributes to the robustness of the per-unit system.

Iterative Methods for Power Flow

Solving the power flow problem requires computing bus voltage magnitudes and angles that satisfy the nonlinear algebraic equations describing real and reactive power injections. Due to the nonlinear nature of these equations, iterative numerical methods are employed. Among the classical approaches, the Gauss-Seidel and Newton-Raphson algorithms represent the two fundamental solution strategies.

Gauss-Seidel is simple and computationally light, but it converges slowly and is rarely used in modern large-scale systems. Newton-Raphson, instead, achieves fast and robust convergence through a Jacobian-based linearization of the power flow equations and remains the industry-standard method for transmission networks.

In the context of this work, these methods are recalled only for completeness, as the subsequent sections focus on the Optimal Power Flow (OPF) formulation. OPF extends the power flow problem by embedding it within an optimization framework, where operating costs or system objectives are minimized subject to the physical constraints of the network.

Introduction to Optimal Power Flow

The classical power flow problem determines the steady-state operating point of an electrical network by computing bus voltages and power flows that satisfy Kirchhoff's laws. However, real power systems must operate not only in a feasible state, but in an *optimal* one according to economic, technical, or security criteria. This leads to the *Optimal Power Flow* (OPF), which embeds




Figure 7.7: Newton-Raphson flowchart.




Figure 7.8: Gauss-Seidel flowchart.

the AC power flow equations within a constrained optimization framework.

The general mathematical structure of an OPF problem can be written as:

$$\begin{aligned}
 \min_{x, u} \quad & f(x, u) \\
 \text{s.t.} \quad & g(x, u) = 0 \quad (\text{Power flow equations}) \\
 & h(x, u) \leq 0 \quad (\text{Operational limits}).
 \end{aligned} \tag{7.9}$$

Here, the variable sets are:

- **Controllable variables** u : generator active/reactive power setpoints (P_g, Q_g) , voltage setpoints, transformer tap positions, shunt injections, and in some formulations even controllable loads or storage.
- **Dependent variables** x : bus voltage magnitudes and angles, branch currents, line power flows, and all quantities determined by the AC power flow model.

The equality constraints $g(x, u) = 0$ enforce the nonlinear AC power flow equations, ensuring physical feasibility, while the inequality constraints $h(x, u) \leq 0$ include generator limits, voltage bounds, thermal line ratings, and other operational restrictions.

The objective function $f(x, u)$ encodes the operational goal. Common examples include:

$$f = \sum_{i=1}^{N_g} (a_i P_{G,i}^2 + b_i P_{G,i} + c_i) \quad (\text{Minimization of generation cost}), \quad (7.10)$$

$$f = \sum_{(i,j) \in \text{lines}} R_{ij} I_{ij}^2 \quad (\text{Minimization of network losses}). \quad (7.11)$$

In the context of this work, OPF provides the foundation for defining feasible and optimal system-level behaviors on large-scale networks, such as the Ethiopian transmission grid.

7.3 Ethiopian Transmission Network (E-TN)

7.3.1 Panoramic and Topology

Ethiopia's national power system is evolving rapidly in response to growing electricity demand, industrial development, and nationwide electrification initiatives. The transmission network extends over a wide and topographically diverse territory, organized around several voltage levels: 66 kV and 132 kV for regional transmission, 230 kV and 400 kV for corridors that plays a critical role in long-distance power transfers. In addition, Ethiopia has a 500 kV high-voltage direct current (HVDC) interconnection with Kenya, enabling large-scale cross-border energy exchange and supporting regional power market integration.

Hydropower remains the dominant source of electricity in Ethiopia, supported by a series of large-scale projects developed over the past two decades [92]. The Gibe III plant, located relatively close to the main load center of Addis Ababa, contributes significantly to supply security in the central region.

More recently, the Grand Ethiopian Renaissance Dam (GERD), situated near the Sudanese border, has become the largest hydropower installation in the country and one of the most significant in Africa. With a multi-gigawatt capacity, GERD dramatically reshapes the northern transmission landscape and introduces new operational challenges and opportunities con-

cerning long-distance bulk power transfers, voltage regulation, and system stability.

Despite these major infrastructure investments, the Ethiopian grid still presents structural vulnerabilities. Long transmission distances, uneven meshing across regions and geographical clustering of generation contribute to congestion risks and reactive power management challenges. Seasonal variability in river flows further complicates operational planning, making the system sensitive to both hydrological conditions and fluctuating renewable production.

As a result, maintaining system reliability requires not only operational coordination but also forward-looking infrastructure planning. Unfortunately, the Ethiopian power system is also weak from a regulatory standpoint, as grid codes are still incomplete and undergoing continuous evolution, which further complicates the integration of large-scale renewable generation [93,94].

As demand continues to rise and electrification progresses in rural areas, comprehensive transmission planning becomes essential. The integration of new renewable sources, particularly wind power, plays a crucial role in this strategy. In this context, the ongoing *Last Wind* project, developed in collaboration with Aalborg University, aims to identify optimal locations and capacities for future wind farms.

By diversifying the generation mix and distributing production more evenly across the territory, the project seeks to enhance system resilience, reduce reliance on hydropower-dominated corridors and mitigate the effects of seasonal variability. Assessing these impacts requires analytical tools capable of capturing network constraints and operational limits, making OPF-based studies a central component of the planning process.

For the purpose of this study, the modeled transmission network focuses on the 230 kV and 400 kV layers, which represent the primary high-voltage backbone responsible for long-distance bulk transfers. The topology of the network has been taken from Open Infrastructure map [95] and was simplified in order to have a 230 kV layer with 17 Buses and 17 lines and a 400 kV of 10 nodes and 10 lines.

Note that the 500 kV has been integrated within this layer in order to use the Kenya bus as Slack. In Figure 7.9 and Figure 7.10, the topologies of the two transmission layers are shown. The 230 kV and 400 kV networks are interconnected through transformers located at six substations, identifiable in the diagrams by the presence of a black dot.




Figure 7.9: 230 kV Network topology




Figure 7.10: 400 kV Network topology

Table 7.1: Network topology used parameters

	N_{nodes}	$N_{branches}$	$R[\Omega/km]$	$X[\Omega/km]$
230 kV	17	17	0.0678	0.36
400 kV	10	10	0.0276	0.32

7.3.2 Load Centers

The electrical demand in Ethiopia exhibits a strongly uneven spatial distribution. Due to the concentration of population, commercial activity, and industrial development, only a few large demand nodes dominate the national profile, the largest of which is Addis Ababa. This metropolitan area is surrounded by vast rural regions, many of which are still undergoing progressive electrification and therefore display lower and less stable consumption patterns [96].

Administratively, the country is divided into eleven regions, each characterized by different levels of industrialization, demographic density, and penetration of the national grid. The western and north-western regions, where most of the major hydropower plants are located, are supplied through the 400 kV transmission backbone redunded with the 230 kV layer.

In contrast, the eastern part of the country relies on the 230 kV network only, which extends across long lines. While the center of the nation surrounding Addis Ababa benefits from a relatively well-interconnected grid, the eastern regions face structural limitations due to their distance from major generation sites and the limited redundancy of local transmission infrastructure. As it will be discussed in the next subsection, the geographical imbalance between generation (mostly in the west) and demand (heavily concentrated in the center and increasingly in the east) creates operational challenges for power transfers and voltage stability.

In [97], the authors analyze the projected evolution of national electricity demand over multiple decades. Overall, Ethiopia's load demand is increasing rapidly. According to data provided by the Ethiopian Transmission System Operator (TSO), the total national electricity consumption reached approximately ~ 13.7 TWh in 2019, already exceeding earlier forecasts.

Current projections indicate that by 2030 total annual consumption may surpass 27 TWh, reflecting both demographic growth and accelerated electrification efforts. Table 7.2 summarizes the expected evolution of sectoral demand.

Table 7.2: Electricity Demand by Sector (TWh)

Sector	2012	2015	2020	2025	2030
Urban	1.17	1.50	2.35	3.68	5.84
Rural	0.42	0.90	1.86	3.22	5.22
Agriculture	0.04	0.05	0.08	0.13	0.18
Service	1.03	1.36	1.65	2.58	3.02
Industry	1.54	3.09	4.58	8.26	12.00
Transport	0	0.41	0.60	0.82	1.07
Total (TWh)	4.19	7.30	11.11	18.69	27.32

To realistically emulate the national load behaviour, the Ethiopian TSO provided the CROM research group with the complete regional load time series for the year 2019. For modeling purposes, a representative load center was identified within each of the eleven regions and assigned the corresponding regional load profile. This approach preserves the geographical asymmetry of demand and will allow the OPF framework to reproduce the pattern of power flows and the operational stress points that characterize the Ethiopian transmission network more realistically.

Table 7.3: Load centers by region and 2019 Total Demand

Region	Load Center	Total Demand (TWh)
Oromia	Addis Ababa	4.55
Sidama	Awasa	0.37
Amhara	Bahir Dar	2.04
Dire Dawa	Dire Dawa	1.28
Gambela	Gambela	0.03
Benishangul-Gumuz	Asosa	0.04
Tigray	Mekelle	1.32
Afar	Semera	0.17
Somali	Jijiga	0.30
Harari	Harar	0.11
SNNPR	SNNPR	3.55




Figure 7.11: Load Centers map

7.3.3 Production Centers

Ethiopia is one of the countries that relies most heavily on renewable sources to meet its national electricity demand. This highly unbalanced generation mix has historically been, and still is, dominated by large hydropower plants supplied by extensive river basins [98]. Over time, the structure of the national transmission grid has effectively developed around this hydro-based supply, with the primary function of transporting power from the generation-rich western regions toward the central and eastern load centers.

Although the ability to meet domestic demand almost entirely through renewable resources represents a significant strategic asset, depending for roughly 90% of total generation on a single energy source exposes the system to several vulnerabilities. Hydropower, while more controllable and predictable than other renewable sources such as solar or wind, remains strongly influenced by seasonal variations, hydrological uncertainty, and periods of drought. These factors can substantially affect generation availability and, consequently, the reliability of the entire system.

For these reasons, Ethiopia has initiated a progressive transition toward a more diversified generation portfolio. Project by Last Wind show the em-

phasis that has been placed on wind power, which represents a promising alternative for supplying regions where the transmission grid is comparatively weak or structurally constrained. Expanding the national renewable mix also aligns with long-term policy goals aimed at strengthening energy security, enhancing regional market participation, and supporting future electricity exports to neighboring countries.

The following subsections provide an overview of the historical evolution of Ethiopia's hydropower and wind power sectors, with particular attention to their geographical distribution and their implications for the operation and planning of the national transmission system.

Hydropower Generation

Hydropower has historically constituted the backbone of Ethiopia's electricity generation. The country is crossed by major rivers like the Omo, Abay (Blue Nile), and Tekeze alongside numerous smaller tributaries that form an extensive reservoir of potential low-cost renewable energy. Over the decades, national development plans have expanded this potential through the construction of increasingly ambitious hydroelectric projects.

Early installations such as Koka and Tis Abay were later followed by larger facilities equipped with substantial storage dams [99]. However, this rapid infrastructural growth was not always met with unanimous acceptance, as the creation of large reservoirs altered traditional agricultural practices and reshaped local ecosystems.

At present, Ethiopia operates fourteen major hydropower plants. Among these, the Gilgel Gibe cascade, especially Gilgel Gibe II and Gilgel Gibe III, plays a central strategic role. Together, these plants reinforce the southern-central generation corridor and serve as a vital hub supplying Addis Ababa, Oromia, and the surrounding regions. Collectively, the Gilgel Gibe plants can reach a peak nominal output, $Prod_p^{\max}$, of approximately ~ 2.5 GWe, supported by a reservoir capacity of nearly ~ 15 km³.

Another cornerstone in Ethiopia's hydropower list of plants is the Grand Ethiopian Renaissance Dam (GERD), located in the north-west near the Sudanese border. At the time of this study, only part of the installation was operational, providing a $Prod_p^{\max}$ of roughly ~ 0.75 GWe. Once fully commissioned, GERD is expected to reach a peak capacity of ~ 5.5 GWe, supported by an exceptionally large reservoir capable of storing up to ~ 74 km³ of water. This makes GERD a future cornerstone of Ethiopia's




Figure 7.12: Maps of the Ethiopian Hydro Power Plants

generation and a potential enabler for regional electricity exports. Each Hydro power plants of the list has been assigned to the close and better connected bus of the modelled network, as shown in Tab. 7.4.

Modelling the river flows

Despite their economic and environmental advantages, hydropower plants remain vulnerable to hydrological fluctuations. Seasonal variability, changes in rainfall patterns, and prolonged droughts can significantly reduce available generation, posing operational challenges to the national grid and highlighting the need for complementary energy sources.

Ethiopia experiences three distinct climatic seasons: *Bega*, the dry season, which accounts for less than 5% of annual rainfall; *Belg*, the short-rains season, contributing roughly 15%; and *Kiremt*, the main rainy season, during which more than 80% of annual precipitation occurs.

This strong seasonal pattern has a profound impact on river discharge across the country. During *Bega*, major river flows may drop to as little as 10–15% of their *Kiremt* peak values, while smaller rivers may dry up entirely. This underlines the importance of reservoir storage, both for ensuring con-

Table 7.4: HPP buses asigation: Plants and Buses

Plant	Bus
Beles	Bahir Dar
Tekeze	Merelle
Gilgel Gibe I	Jimma
Gilgel Gibe II	Jimma
Gibe III	Jimma
Melka Wakena	Awasa
Finchaa	Finchaa
Genale Dawa	Genale Dawa
GERD	GERD
Amerti Neshe	Finchaa
Tis Abay II	Bahir Dar
Koka	Addis Ebeba
Awash II	Addis Ebeba
Awash III	Addis Ebeba

Table 7.5: Main Ethiopian Hydropower Plants: Capacity, Reservoir Size, and River Basin

Plant	Peak Production	Reservoir Size	River
	[MWe]	[km ³]	
Beles	460	9.10	Beles
Tekeze	300	9.30	Tekeze
Gilgel Gibe I	184	0.92	Gilgel Gibe
Gilgel Gibe II	420	0.00	Gilgel Gibe
Gibe III	1870	14.70	Omo
Melka Wakena	153	0.75	Shebelle
Finchaa	134	0.63	Finchaa
Genale Dawa	254	2.60	Genale Dawa
GERD	5050	74.00	Blue Nile
Amerti Neshe	95	0.13	Amerti Neshe
Tis Abay II	84.4	0.00	Blue Nile
Koka	43	1.90	Awash
Awash II	32	0.00	Awash
Awash III	32	0.00	Awash

Table 7.6: Ethiopian Climatic Seasons

Season	Gregorian Months	Ethiopian Months
Bega (dry season)	October – January	Tqmt(i), Hdar, Tahsas, Trr(i)
Belg (short rains)	February – May	Yäkatit, Mägabıt, Miyazya, Gänbot
Kiremt (main rains)	June – September	Säne, Hamle, Nähasä, Mäskäräm*

*Mäskäräm marks the end of Kiremt and the transition into Bega.




Figure 7.13: Example of the modelled rivers flow

tinuous hydropower generation and for meeting agricultural and domestic water needs.

Since official continuous flow measurements were not available for all rivers considered in this study, a simplified hydrological model was adopted. Using the maximum, average, and minimum reported flows for each river, an annual discharge profile was constructed such that the minimum flow occurs during Bega, while the remaining volume is distributed across Belg and Kiremt according to their known seasonal proportions (approximately 85% during Kiremt and 15% during Belg).

Although approximate, these flow profiles provide a reasonable representation of seasonal water availability and therefore enable a simplified modelling of hydropower output variability throughout the year.

Table 7.7: Main Ethiopian Rivers Feeding Hydropower Plants

River	Max Flow	Avg Flow	Min Flow
	[m³/s]	[m³/s]	[m³/s]
Beles	440	318	17
Tekeze	743	171	0
Gilgel Gibe	134.4	52.8	13.12
Omo	1680	660	164
Shebelle	158.7	71.2	15.7
Finchaa	95.59	31.93	0.2
Genale Dawa	531.6	241.5	56.2
Blue Nile	6096	2106	424
Amerti Neshe	24.72	10.3	1.54
Awash	803	43.5	0.035

Modelling the power generation

Another important step required to properly emulate hydropower production is the definition of the relationship between electrical power output and water consumption. This can be derived beginning from the standard hydropower equation:

$$P = \eta \rho g H Q = \beta \eta Q \quad (7.12)$$

where:

- P is the electrical power output (W),
- η is the turbine-generator efficiency,
- ρ is the density of water (1000 kg/m³),
- g is the gravitational acceleration (9.81 m/s²),
- H is the hydraulic head (m),
- Q is the volumetric water flow rate (m³/s).
- $\beta = \rho g H$,

In this study the hydraulic head H is preliminarily assumed to be 50 m for each plant, which represents a reasonable average value for typical large-scale hydropower installations. Also the parameter β is assumed equal for all of the plants.

To better illustrate the relationship between power production and water consumption, it is useful to estimate the amount of water required to produce 1 kWh of electrical energy. A quantity of 1 kWh corresponds to maintaining an average electrical output of 1 kW for one hour. Assuming an efficiency of $\eta = 0.9$, the required power level can be fixed as:

$$P = 1000 \text{ W}.$$

Substituting this value into the hydropower equation yields:

$$1000 = 0.9 \cdot 1000 \cdot 9.81 \cdot 50 \cdot Q.$$

Solving for Q :

$$Q = \frac{1000}{0.9 \cdot 1000 \cdot 9.81 \cdot 50}.$$

The resulting flow rate needed to generate 1000 W under the assumed conditions is:

$$Q \approx 2.27 \times 10^{-3} \text{ m}^3/\text{s}.$$

With the required steady-state flow rate known, the total water consumption can be computed. Since $1 \text{ m}^3 = 1000 \text{ liters}$:

$$2.27 \times 10^{-3} \text{ m}^3/\text{s} \times 1000 = 2.27 \text{ liters/s}.$$

Thus, the total water required over a full hour is:

$$\text{Water per kWh} = 2.27 \text{ liters/s} \times 3600 \text{ s} = 8160 \text{ liters}.$$

Under the assumed hydraulic head and efficiency, producing 1 kWh of hydroelectric energy requires approximately 8160 liters of water to pass through the turbine during one hour of operation. Although this estimate is simplified, it provides an intuitive metric for evaluating the potential energy stored in Ethiopia's reservoirs and for understanding the relationship between hydropower production and seasonal water availability. In this estimation it is important to consider also that only the 80% of the volume is actually usable to store water for production.

Table 7.8: Main Ethiopian Hydropower Plants: Estimated Energy Capacity

Plant	Nom. Reserve [km ³]	Usable Reserve [km ³]	Est. Energy [TWh]
Beles	9.10	7.28	0.89
Tekeze	9.30	7.44	0.91
Gilgel Gibe I	0.92	0.74	0.09
Gibe III	14.70	11.76	1.44
Melka Wakena	0.75	0.60	0.07
Finchaa	0.63	0.50	0.06
Genale Dawa	2.60	2.08	0.25
GERD	74.00	59.20	7.25
Amerti Neshe	0.13	0.10	0.01
Koka	1.90	1.52	0.19
Total	114.0	91.22	11.18

Modelling the efficiency

Another important aspect required to complete the hydropower production model is the representation of turbine efficiency. The efficiency of a hydropower unit is a nonlinear function that link the volumetric flow rate Q at the inlet to the electrical power output P at the generator.

This behaviour depends strongly on the operating point of the turbine, and assuming a constant efficiency may lead to inaccurate energy estimations, especially during off-design conditions.

defining Q_{nom} as the flow required to bring the turbine to its own maximum power output generation it is possible to define Q_r , the relative flow, as:

$$Q_r = \frac{Q}{Q_{nom}} \quad (7.13)$$

that can be used to express the efficiency η as a function of such parameters.

In this work, all hydropower plants are modelled as equipped with Francis turbines. This assumption is consistent with the actual configuration of the largest Ethiopian plants, such as GERD and Gibe III, which indeed rely on large Francis units. To approximate the efficiency profile of $\eta(Q_r)$, a piecewise efficiency curve was adopted and uniformly applied to all turbines in the model.

Table 7.9: Points for the piecewise approximation of the efficiency

N_{point}	1	1	3	4	5
$\eta(Q_r)$	0.4	0.8	0.9	0.9	0.85
Q_r	0	0.5	0.8	0.9	1.2

- **Nominal operating range ($Q_r \in 80\text{--}90\%$):** Efficiency typically reaches **93–94%** at or near the design point, with stable hydraulic conditions and minimal energy losses. This is the optimal operating zone for Francis turbines.
- **Partial load ($Q_r \in 50\text{--}80\%$):** Efficiency progressively declines to approximately **80–90%** due to sub-optimal conditions. The turbine remains fully operable but less efficient.
- **Low-load region ($Q_r \in 40\text{--}50\%$):** The efficiency drops further to around **60–80%**, with the appearance of flow recirculation and draft-tube vortices. Operation becomes less stable and generally less desirable.
- **Very low load ($Q_r < 40\%$):** Efficiency falls below **60–70%**, and significant hydraulic instabilities, vibrations, and cavitation may occur. Continuous operation in this region is normally avoided.
- **Overload region ($Q_r \in 90\text{--}120\%$):** Operation above nominal flow is technically possible, but efficiency decreases by **2–5 percentage points**, with an increased risk of cavitation and mechanical stress. Beyond this range, performance deteriorates sharply.

For this study, a simplified piecewise approximation of the typical Francis efficiency profile was implemented, ensuring that efficiency decreases appropriately at low flows and near overload while remaining close to 90% within the nominal operating range.

In our model $\eta_{nom} \approx 0.88$ is the efficiency corresponding to Q_{nom} , or $\eta(Q_r = 1)$. Starting from $Prod_p^{max}$ it is possible to compute the corresponding nominal flow of each plant p as:

$$Q_{nom}^p = \frac{Prod_p^{max}}{\eta_{nom} \cdot \beta} \quad (7.14)$$




Figure 7.14: Simplified Francis Turbine efficiency curve

Wind Power Generation

As introduced before, wind energy represents the second major renewable resource in Ethiopia and is the center of particular interest due to its potential to scale production and diversify the mix of generation.

In recent decades, wind farms have seen a progressive expansion on the utility scale to reduce national dependence on hydrological conditions and strengthen the resilience of the power system and studies have been conducted to evaluate possible new sites [100].

At the moment, national efforts have been concentrated in three main locations where average wind speeds are adapt for large-scale power production. These locations are Ashegoda in the north, Adama in the center, and Aysha at the eastern edge of the country. In all three sites a wind farm is already present, although of moderate size.

A previous phase of the Last Wind project aimed specifically at evaluating how much these wind farms could be expanded to maximise their impact on the grid. Future investments in this field promise to provide an important non-hydrological asset capable of stabilise generation during the dry season.

From a system-operation perspective, wind generation represents a steady source of energy, available throughout the year and throughout the entire day, offering strong seasonal complementarity with hydropower.

Wind farms also improve geographic diversification by placing production nodes in regions where hydropower is absent, such as the eastern part of the




Figure 7.15: Maps of the Ethiopian wind farms

Table 7.10: Main Ethiopian Windpower Plants: Actual and Planned Size

Plant	Actual Size [MW]	Planned Size [MW]
Adama	204	261.5
Ashegoda	120	120
Aysha I	60	125
Aysha II	60	125

country, which is typically less well connected to the main grid.

On the other hand, integration of wind production introduces a new set of challenges. Wind is, by its nature, far less predictable than hydropower, and to ensure adequate usage the system must be equipped with sufficient reserve margins, robust forecasting tools, and transmission flexibility to manage rapid fluctuations in generation.

The previous phase of the Last Wind project [101] aimed to identify the locations where it would be most convenient to invest in the expansion of existing wind farms. The results of that study concentrated the recommended expansion mainly in the Adama area, but especially in Aysha in the far east of the country. This outcome aligns with both the national geography and

the structure of the transmission grid, as the eastern region lies far from the major hydropower generation sites and is characterised by a particularly weak network infrastructure.

Modelling the wind

The first key step in properly emulating the behaviour of wind farm production is to analyse the historical wind speed measurements available for the three selected regions. These data were acquired through the NASA wind database, which provides hourly wind speed records at a reference height of 50 m above ground for the chosen locations.

On a daily time scale, wind exhibits a high degree of variability, and any real-world operational framework requires appropriate forecasting tools to manage this uncertainty. However, in this work, the forecast is assumed to be perfect, as the focus is placed on the development of the monitoring architecture and the decision system. The integration of forecasting algorithms and uncertainty management is left to future extensions of the project.

The NASA portal provided multi-year wind profiles from 2020 to 2024, allowing us to characterize the seasonal behaviour of the wind in the different areas under study.

To evaluate the statistical behaviour of the wind flow during the years, the data set was divided by months, and the confidence interval was calculated using Eq. 7.15 and Eq. 7.16, which are repeated below for completeness.

$$s = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (x_i - \bar{x})^2} \quad (7.15)$$

$$\sigma_{1-\alpha} \left[\sqrt{\frac{\text{df}, s^2}{\chi^2_{1-\alpha/2, \text{df}}}}; \sqrt{\frac{\text{df}, s^2}{\chi^2_{\alpha/2, \text{df}}}} \right] \quad (7.16)$$

Table 7.11 shows the 95% confidence interval for the wind speed in each location during a period of a year while Figure 7.16 represents a box plot of the same wind speeds

Since turbines installed in Ethiopian wind farms operate at approximately 100 m of height, the wind speed data must be scaled accordingly. The vertical variation of wind speed can be estimated using empirical laws that approximate the wind profile as a function of surface roughness and atmospheric stability. One of the most widely used formulations is the *power law*:

Table 7.11: Monthly wind speed confidence intervals at 50 m for Aysha, Ashegoda, and Adama (95% CI)

Month	Aysha [m/s]	Ashegoda [m/s]	Adama [m/s]
January	7.0 ^(+1.9) _(-1.8)	4.5 ^(+1.8) _(-1.7)	5.2 ^(+1.7) _(-1.7)
February	7.0 ^(+2.1) _(-2.0)	4.7 ^(+2.0) _(-1.9)	5.5 ^(+1.9) _(-1.8)
March	6.6 ^(+2.5) _(-2.4)	5.0 ^(+2.2) _(-2.1)	5.2 ^(+2.1) _(-2.0)
April	5.8 ^(+2.3) _(-2.2)	4.5 ^(+2.0) _(-1.9)	4.5 ^(+2.2) _(-2.1)
May	5.2 ^(+2.3) _(-2.2)	3.9 ^(+2.0) _(-1.9)	4.1 ^(+2.2) _(-2.1)
June	6.4 ^(+2.9) _(-2.8)	4.5 ^(+1.9) _(-1.8)	4.2 ^(+2.0) _(-1.9)
July	7.9 ^(+2.9) _(-2.7)	4.6 ^(+1.3) _(-1.3)	5.6 ^(+2.2) _(-2.1)
August	7.5 ^(+2.7) _(-2.6)	3.8 ^(+1.6) _(-1.5)	4.8 ^(+2.2) _(-2.1)
September	5.7 ^(+2.7) _(-2.6)	3.5 ^(+1.7) _(-1.7)	3.3 ^(+1.8) _(-1.7)
October	6.0 ^(+2.0) _(-1.9)	4.2 ^(+1.7) _(-1.6)	4.7 ^(+2.2) _(-2.1)
November	7.1 ^(+1.6) _(-1.5)	5.1 ^(+1.6) _(-1.5)	5.3 ^(+1.8) _(-1.7)
December	7.2 ^(+1.6) _(-1.6)	4.9 ^(+1.8) _(-1.7)	5.2 ^(+1.8) _(-1.7)

$$v(h) = v(h_0) \left(\frac{h}{h_0} \right)^\alpha \quad (7.17)$$

where:

- $v(h)$ is the wind speed at height h
- $v(h_0)$ is the wind speed measured at the reference height h_0
- h and h_0 are the heights in meters
- α is the power-law exponent, dependent on surface roughness. In the Ethiopian case, it is possible to assume it as $\alpha = 0.15$

The new scaled values are represented in Table 7.12, which shows the updated 95% confidence interval of the wind speed in each site, while Figure 7.17 represents the boxplot of such distribution for the year.

Among the three analysed sites, Aysha consistently presents higher wind speeds throughout the year. An additional element of interest is the recurring




Figure 7.16: Boxplot of the wind speed at 50m in the three locations




Figure 7.17: Boxplot of the wind speed at 100m in the three locations

seasonal pattern: historical data show a noticeable decrease in wind speed during the transitions from Belg to Kiremt and from Kiremt to Bega.

This behaviour has been observed over multiple years and it must be considered both in long-term expansion planning and in the design of monitoring and control strategies, as the system must anticipate a significant reduction in wind generation during these periods.

Wind turbine modelling

To correctly address the behaviour of wind turbines, several aspects must be addressed. First, wind turbines operate within a well-defined range of wind speeds, outside of which their behaviour is governed by safety constraints.

In particular, when the wind speed exceeds a critical threshold, known as the *cut-out wind speed*, the turbine is automatically shut down to prevent mechanical and electrical damage. For modern utility-scale turbines, this limit typically lies between 20 and 25 m/s (approximately 72–90 km/h). Beyond this value, the aerodynamic loads on the blades and the structural components become potentially dangerous.

When cut-out conditions are reached, the turbine activates a dedicated safety mode: the blades are pitched into a feathered position to minimise

Table 7.12: Monthly wind speed confidence intervals at 50 m for Aysha, Ashegoda, and Adama (95% CI)

Month	Aysha [m/s]	Ashegoda [m/s]	Adama [m/s]
January	7.8 ^(+2.1) _(−2.0)	5.0 ^(+2.0) _(−1.9)	5.8 ^(+1.9) _(−1.9)
February	7.7 ^(+2.4) _(−2.3)	5.3 ^(+2.2) _(−2.1)	6.1 ^(+2.1) _(−2.0)
March	7.3 ^(+2.7) _(−2.6)	5.5 ^(+2.5) _(−2.4)	5.8 ^(+2.3) _(−2.2)
April	6.4 ^(+2.5) _(−2.4)	5.0 ^(+2.2) _(−2.1)	5.0 ^(+2.5) _(−2.3)
May	5.8 ^(+2.5) _(−2.4)	4.3 ^(+2.2) _(−2.1)	4.6 ^(+2.5) _(−2.4)
June	7.1 ^(+3.3) _(−3.1)	5.0 ^(+2.1) _(−2.0)	4.6 ^(+2.2) _(−2.1)
July	8.8 ^(+3.2) _(−3.0)	5.1 ^(+1.5) _(−1.4)	6.2 ^(+2.4) _(−2.3)
August	8.3 ^(+3.0) _(−2.9)	4.2 ^(+1.7) _(−1.7)	5.3 ^(+2.4) _(−2.3)
September	6.4 ^(+3.0) _(−2.9)	3.9 ^(+1.9) _(−1.9)	3.7 ^(+1.9) _(−1.9)
October	6.6 ^(+2.2) _(−2.1)	4.7 ^(+1.8) _(−1.8)	5.2 ^(+2.4) _(−2.3)
November	7.8 ^(+1.8) _(−1.7)	5.7 ^(+1.7) _(−1.6)	5.9 ^(+2.0) _(−1.9)
December	8.0 ^(+1.8) _(−1.7)	5.4 ^(+2.0) _(−1.9)	5.7 ^(+2.0) _(−1.9)

aerodynamic forces, the rotor is progressively braked, and the generator is disconnected from the grid by the conversion system.

This results in a complete stop in power production until wind conditions return within the operational range.

Table 7.13: Operating regimes of a modern wind turbine

Wind Speed Range [m/s]	Operating Regime	Power Output
0–3	Below cut-in	Turbine off (no production)
3–12	Cut-in to rated speed	Increasing, variable power
12–25	Rated operating region	Nominal (rated) power
> 25	Cut-out	Turbine shut down for safety

However, in the Ethiopian scenario, such extreme wind conditions are historically rare and occur significantly only in Aysha.

The second important step is to evaluate the power produced by the turbines. Wind turbines installed in Ethiopian wind farms, and in most wind

farms in general, use the Horizontal-Axis Wind Turbine (HAWT) design, and the following formulation has been used to evaluate the production.

Starting from the kinetic energy possessed by the moving air, the flux available to a wind turbine is expressed as:

$$E_c = \frac{1}{2} \rho A v^3 \quad (7.18)$$

where:

- ρ is the air density (approximately 1.225 kg/m^3),
- A is the swept area of the rotor ($A = \pi R^2$),
- v is the wind speed (m/s).

There exists a theoretical limit on the maximum kinetic energy that can be extracted by a turbine, known as the Betz limit, which fixes the maximum extractable fraction at 59.3%. Taking into account all practical considerations, modern utility-scale wind turbines achieve an overall efficiency of approximately 35–45%.

In this study, the onshore turbine model SG 5.0-145(2.0) has been used as a reference to ensure realistic parameters:

- $P_{nom} = 5 \text{ MW}$ nominal power output per turbine,
- $r = 72.5 \text{ m}$ blade radius, corresponding to a rotor area of $A \approx 1.6 \times 10^4 \text{ m}^2$,
- $\rho = 40\%$ expected conversion efficiency.

To emulate each wind farm, the single-turbine model was replicated until the total installed capacity of the plant was reached. Evaluating expected power production places even more emphasis on the criticality of seasonal wind speed drops: as shown in Eq. 7.18, the cubic relationship between wind speed and available energy makes wind generation extremely sensitive to variations occurring at seasonal transitions.

Each Wind power plants in the list has been assigned to the closest and better connected bus of the modelled network

Table 7.14: Monthly power generation of a single SG 5.0-145(2.0) turbine [MW] with 95% CI

Month	Aysha	Ashegoda	Adama
Month	[MW]	[MW]	[MW]
January	1.93 ^(+2.04) _(-1.15)	0.51 ^(+0.91) _(-0.39)	0.80 ^(+1.09) _(-0.55)
February	1.86 ^(+2.30) _(-1.20)	0.59 ^(+1.12) _(-0.47)	0.94 ^(+1.29) _(-0.64)
March	1.60 ^(+2.54) _(-1.17)	0.67 ^(+1.39) _(-0.55)	0.79 ^(+1.40) _(-0.61)
April	1.07 ^(+1.85) _(-0.82)	0.50 ^(+1.03) _(-0.41)	0.50 ^(+1.16) _(-0.42)
May	0.80 ^(+1.57) _(-0.64)	0.32 ^(+0.77) _(-0.28)	0.39 ^(+1.04) _(-0.35)
June	1.43 ^(+3.04) _(-1.18)	0.50 ^(+0.96) _(-0.40)	0.40 ^(+0.89) _(-0.34)
July	2.77 ^(+2.23) _(-1.99)	0.54 ^(+0.62) _(-0.34)	0.97 ^(+1.64) _(-0.73)
August	2.34 ^(+2.66) _(-1.69)	0.30 ^(+0.55) _(-0.24)	0.61 ^(+1.25) _(-0.50)
September	1.04 ^(+2.30) _(-0.87)	0.23 ^(+0.56) _(-0.20)	0.20 ^(+0.51) _(-0.17)
October	1.18 ^(+1.65) _(-0.81)	0.42 ^(+0.72) _(-0.32)	0.58 ^(+1.21) _(-0.47)
November	1.94 ^(+1.67) _(-1.02)	0.73 ^(+0.89) _(-0.47)	0.84 ^(+1.16) _(-0.58)
December	2.04 ^(+1.74) _(-1.07)	0.65 ^(+0.98) _(-0.47)	0.76 ^(+1.09) _(-0.53)

7.3.4 E-TN Digital Twin

With all of the information above, it is now possible to assemble the digital twin of the E-TN. The developed grid model includes 27 buses across both voltage layers, some of which are used only as transit nodes to properly emulate the topology of the network and the power flow.

Eleven of these nodes act as load centers for the various regions, assigned as in Figure 7.11, while production plants have been allocated as shown in Figure 7.12 and Figure 7.15, and according to Tables 7.3, 7.4 and 7.15.

Table 7.15: WPP buses asigation: Plants and Buses

Plant	Bus
Adama	Addis Ebeba
Ashegoda	Merelle
Aysha I	Aysha
Aysha II	Aysha

Each bus hosting at least one power plant has been defined as a PV bus, the Kenya bus represents the Slack bus, and the remaining nodes have been defined as PQ buses as shown in Table 7.16.

The network includes a total of 33 connections among the nodes, 6 of which are transformers, 10 belong to the 400 kV transmission layer, and the remaining lines form the 230 kV layer. Detailed list is shown in Table 7.17.

To properly emulate real-grid behavior, a ramp-up limitation has been introduced on the power deliverable by each hydroelectric plant. Since our simulations reproduce grid operation hour by hour, every plant is allowed to increase its power output by at most 10% of its nominal capacity relative to the set point of the previous time step.

In contrast, reductions in output are always permitted without restriction. This mechanism generates dynamic production bounds that depend on the operating state at the beginning of each step.

These dynamic boundaries are then used by the EMS to convey its strategy to the digital twin. Regardless of the cost function, we simply require the minimisation of total production. Why minimise production? Because when a given load must be met, reducing generation inherently reduces losses, guiding the system toward the most efficient operating point in terms of overall network losses.

A rudimentary form of load shedding has also been implemented to account for network contingencies. At each iteration, the topology of the grid is analyzed to verify whether the system remains fully connected or if line outages have split it into multiple electrically isolated sub-networks.

For each resulting sub-network, the available generation capacity is evaluated by considering the production limits and ramp-up constraints of the generators located within that island. If the locally producible energy is not sufficient to meet the total demand of the sub-network, a load shedding procedure is activated, reducing the load until the internal generation-demand balance is restored. This mechanism allows the model to emulate islanding operation and assess the impact of limited generation availability on supply continuity under degraded network conditions.

For each island $k \in \mathcal{I}$, load shedding is modelled through a non-negative decision variable s_i associated with each demand bus $i \in \mathcal{N}_k$, representing the fraction of load that is supplied. The served demand is therefore defined as:

$$Load_i^{\text{served}} = s_i Load_i, \quad 0 \leq s_i \leq 1, \quad \forall i \in \mathcal{D}_k. \quad (7.19)$$

Table 7.16: Ethiopian Transmission Grid Buses (PF Model)

ID	Name	Voltage	Bus Type	Plant List
1	Addis Ebeba	230 kV	PV	Koka, Awash II, Awash III AdamaWind I, AdamaWind II
2	Awasa	230 kV	PV	MelkaWakena
3	Jimma	230 kV	PV	GilgelG_I, GilgelG_II, GIBEIII
4	Finchaa	230 kV	PV	Finchaa, AmertiNeshe
5	Bahir Dar	230 kV	PV	Beles, TisAbayII
6	Dessie	230 kV	PQ	–
7	Dire Dawa	230 kV	PQ	–
8	Gode	230 kV	PQ	–
9	Gambela	230 kV	PV	Gambela
10	Asosa	230 kV	PQ	–
11	Merelle	230 kV	PV	Tekeze, AshegodaWind
12	Semera	230 kV	PQ	–
13	Aysha	230 kV	PV	AyishaWind
14	Jijiga	230 kV	PQ	–
15	Harar	230 kV	PQ	–
16	Deressa	230 kV	PQ	–
17	Debra Markos	230 kV	PQ	–
18	Addis Ebeba 400	400 kV	PQ	–
19	Awasa 400	400 kV	PQ	–
20	Jimma 400	400 kV	PQ	–
21	Bahir Dar 400	400 kV	PQ	–
22	Genale Dawa	400 kV	PV	GenaleDawa
23	SNNPR	400 kV	PQ	–
24	GERD	400 kV	PV	GERD
25	Deressa 400	400 kV	PQ	–
26	Debra Markos 400	400 kV	PQ	–
27	Kenya Slack	400 kV	Slack	–

Table 7.17: Transmission lines and transformers with resistance and reactance

From	To	Distance [km]	Type	R [Ω]	X [Ω]
1	2	220	230kV	14.92	79.20
1	3	260	230kV	17.63	93.60
1	4	170	230kV	11.53	61.20
1	6	250	230kV	16.95	90.00
1	7	345	230kV	23.39	124.20
2	8	580	230kV	39.32	208.80
3	4	210	230kV	14.24	75.60
3	9	260	230kV	17.63	93.60
4	17	90	230kV	6.10	32.40
5	6	250	230kV	16.95	90.00
5	17	140	230kV	9.49	50.40
6	11	260	230kV	17.63	93.60
6	12	170	230kV	11.53	61.20
7	13	150	230kV	10.17	54.00
7	15	40	230kV	2.71	14.40
10	16	240	230kV	16.27	86.40
14	15	75	230kV	5.09	27.00
18	20	260	400kV	7.18	83.20
18	25	245	400kV	6.76	78.40
18	26	185	400kV	5.11	59.20
19	20	190	400kV	5.24	60.80
19	22	210	400kV	5.80	67.20
19	27	600	400kV	16.56	192.00
20	23	190	400kV	5.24	60.80
21	24	255	400kV	7.04	81.60
21	26	140	400kV	3.86	44.80
24	25	285	400kV	7.87	91.20
1	18	0	Transf.	0.01	0.04
2	19	0	Transf.	0.01	0.04
3	20	0	Transf.	0.01	0.04
5	21	0	Transf.	0.01	0.04
16	25	0	Transf.	0.01	0.04
17	26	0	Transf.	0.01	0.04




Figure 7.18: Example of the E-TN Digital Twin

A feasibility check can be expressed by comparing the maximum producible power inside the island, also limited by ramp-up constraints, with the total served demand:

$$\Delta P_k = \sum_{i \in \mathcal{D}_k} Load_i^{\text{served}} - \sum_{g \in \mathcal{G}_k} Prod_g^{\text{max}}(t), \quad (7.20)$$

where $Prod_g^{\text{max}}(t)$ denotes the time-dependent maximum available generation within the sub-network, including ramp-up limitations. If $\Delta P_k > 0$, load shedding is activated and the coefficients s_i are gradually reduced until the feasibility condition is satisfied.

Once load shedding has been applied, the OPF is tasked with supplying the remaining demand within each sub-network by solving independently on each island k . The optimisation problem is formulated to minimise hydropower production, subject to power balance and operational constraints.

Since energy storage systems are not included in the current model, the system always exploits the full availability of wind generation first, treating it as a non-dispatchable and priority source. Hydropower plants are then

used only to cover the residual demand that cannot be satisfied by wind generation alone, while respecting ramp-up limits.

Denoting by $\mathcal{G}_k^H$ the set of hydro generators and by $\mathcal{G}_k^W$ the set of wind generators in island k , the optimisation problem reads:

$$\min_{x_k, u_k, s} \sum_{g \in \mathcal{G}_k^H} Prod_g \quad (7.21)$$

$$\text{s.t. } g_k(x_k, u_k) = 0 \quad (\text{power flow equations}) \quad (7.22)$$

$$h_k(x_k, u_k) \leq 0 \quad (\text{operational limits}) \quad (7.23)$$

$$\sum_{g \in \mathcal{G}_k} Prod_g = \sum_{i \in \mathcal{D}_k} Load_i^{\text{served}} + P_k^{\text{loss}} \quad (7.24)$$

$$0 \leq Prod_g \leq Prod_g^{\max}(t), \quad \forall g \in \mathcal{G}_k \quad (7.25)$$

$$Prod_g = Prod_g^{\max}(t), \quad \forall g \in \mathcal{G}_k^W \quad (\text{wind treated as must-take}) \quad (7.26)$$

where x_k are the dependent variables (voltages, angles, and line flows), and u_k are the control variables (generator set-points, voltage references, transformer taps, etc.). The constraint (7.26) enforces the full use of available wind power, while the objective (7.21) minimises hydropower usage to preserve water resources.

Minimising total generation while meeting the remaining load naturally leads to a reduction in network losses, as lower power injections result in reduced line currents.

As a consequence, the adopted objective function simultaneously promotes efficient network operation and conservative use of water resources, which is consistent with the long-term planning perspective of the study.

7.4 Operational Management System (OMS) for E-TN

The digital twin of E-TN is an important asset in studying the behaviour of the grid and then developing a monitoring system to coupled it with. In this way it will be possible to increase efficiency and robustness by coordinating the usage of the available power sources.

In the context of the Last Wind project, such a system has been defined as the Operational Management System (OMS), representing the highest

decision-making layer. Its primary role is to ensure secure and efficient system operation under network constraints, generation limitations and contingency conditions translating into medium and short term operational objectives attainable through coordinated control actions.

The OMS explicitly accounts for the following:

- Transmission constraints
- Inter-area power exchanges
- Ramp-up limitations of hydropower plants
- Availability of non-dispatchable renewable sources

In this context, its functions are to coordinate the dispatch of generation, managing load shedding when necessary and preserving critical resources such as water reserves.

The OMS can be interpreted as an Energy Management System which, based on forecasts of load center demand and wind power production, determines the set-points to be requested from hydropower plants. The OMS formulates its decision problem as a Linearised Optimal Power Flow (L-OPF) and this choice has been taken as a good compromise between computational cost and realistic consideration of the dynamics of the network.

Network losses are approximated using precomputed loss coefficients (the so-called α -coefficients), allowing the impact of power injections on losses to be represented in linear form. This approach enables fast optimisation over large networks and makes the OMS suitable for repeated execution within an operational time horizon.

By embedding the OPF within a rolling-horizon framework, the OMS can schedule control actions over time, anticipate ramp-up constraints, and mitigate future shortages. The resulting optimisation balances three competing objectives:

- Satisfying demand
- Minimising Network losses
- Limiting hydropower usage to support long-term sustainability and resilience of the transmission system

One of the strengths of a correct linearised formulation is that the search for the solution is carried out through the properties of the solution space itself, rather than relying solely on a physical interpretation of the problem.

Within a receding-horizon architecture, if all time steps are properly linked to each other, this allows the optimise over the entire horizon at once, rather than solving each step independently. As a result, the global optimum for the whole sequence of steps can always be identified in a single iteration.

In the following, the definition of the linearised problem is presented. Note that the ramp-up variables are used as coupling elements between consecutive time steps, linking all the steps in a single optimisation space. Initially, the variables and the problem formulation for a single-step horizon are presented in detail. Subsequently, the full multi-step horizon formulation is introduced.

7.4.1 Variables and the parameters for the OMS

The linearised OPF problem formulated within the Operational Management System relies on the following sets, parameters, and decision variables.

Sets

- $\mathcal{N}$: set of network nodes, indexed by i .
- $\mathcal{B}$: set of transmission branches, indexed by b .
- $\mathcal{G}$: set of all generation units, indexed by g .
- $\mathcal{G}^H \subset \mathcal{G}$: set of hydropower plants (HPPs), indexed by p .
- $\mathcal{G}^W \subset \mathcal{G}$: set of wind power plants (WPPs), indexed by w .
- $\mathcal{T}$: set of time steps within the optimisation horizon, indexed by t .

Parameters

- $Load_i^{req}(t)$: electrical demand requested at node i at time t .
- $Prod_g^{\max}(t)$: maximum power output available from generator g at time t .
- $Wind_w^{ava}(t)$: wind power available for WPP w at time t .
- $Prod_p^{\max}$: nominal power production of the hydropower plant p .
- Δ_p^{ru} : relative ramp-up limit of the hydropower plant p .
- $Prod_p^{prev}$: power output of the hydropower plant p in the previous time step.

- $Reserve_p(t)$: water stored in the reservoir associated with the plant p .
- $Flow_p(t)$: river flow associated with plant p .
- Q_{1MWh}^{nom} : nominal water volume required to produce 1 MWh.
- α_b : vector of linear loss coefficients used to approximate network losses.
- LS_{levels} : Integer variable, number of possible load shedding steps.
- M : large positive constant used in big- M linearisation constraints.
- N_{piece} : number of pieces used to approximate HPP efficiency.

Decision Variables

- $Prod_i^N(t)$: net power production at node i .
- $F_b(t)$: active power flow on branch b .
- $|F_b(t)|$: auxiliary variable representing the absolute value of branch flow.
- $Prod_g(t)$: power output of generator g .
- $s_i(t)$: load supply factor at node i , integer value $s_i(t) \in \{0, \dots, LS_{levels}\}$.
- $Load_i^{serv}(t)$: load served at node i , equal to $\frac{s_i(t)}{LS_{levels}} Load_i^{req}(t)$.
- $Q_p(t)$: water flow used by hydropower plant p .
- $Q_p^{ava}(t)$: water flow available for plant p .
- $RU_p(t)$: ramp-up bound variable for hydropower plant p .
- $B_{p,r}(t)$: binary variable selecting the r -th segment of the piecewise-linear efficiency curve of plant p .

7.4.2 Network modelling and loss approximation constraints

The transmission network is modelled through nodal power balance constraints, including branch power flows and a linear approximation of network losses. Let $\mathcal{N}$ be the set of buses and $\mathcal{B}$ the set of active transmission branches. For each branch $b \in \mathcal{B}$, let $i = \text{from}(b)$ and $j = \text{to}(b)$ denote its terminal buses, F_b the signed active power flow in the branch, and $|F_b|$ an auxiliary variable representing its magnitude.

Network losses are approximated using precomputed linear loss coefficients α_b , following a standard α -coefficient formulation. Due to the lack of detailed electrical data for the transmission lines, a simplifying assumption has been adopted by assigning a uniform loss coefficient $\alpha_b = 0.1$ (i.e., 10%) to all branches.

Although approximate, this choice allows the model to capture the qualitative impact of network losses while preserving linearity and computational efficiency.

Definition of nodal generations

Let $\mathcal{N}$ be the set of buses and $\mathcal{G}$ the set of generation units. A binary assignment parameter Γ_{ig} is defined such that:

$$\Gamma_{ig} = \begin{cases} 1 & \text{if generator } g \text{ is connected to bus } i \\ 0 & \text{otherwise} \end{cases}$$

For each bus $i \in \mathcal{N}$, the active nodal power exchange is defined as the sum of the power outputs of all generators connected to that bus:

$$Prod_i^N(t) = \sum_{g \in \mathcal{G}} \Gamma_{ig} Prod_g(t), \quad \forall i \in \mathcal{N}. \quad (7.27)$$

Absolute value modelling for branch flows

The auxiliary variables $|F_b|$ are linked to the signed branch flows through the linear constraints:

$$F_b(t) \leq |F_b(t)| \rightarrow F_b(t) - |F_b(t)| \leq 0 \quad (7.28)$$

$$-F_b(t) \leq |F_b(t)| \rightarrow -F_b(t) - |F_b(t)| \leq 0 \quad (7.29)$$

which, in combination with a cost function that minimises $|F_b(t)|$, ensures that $|F_b(t)|$ fall right on $|F_b(t)|_{\text{true}}$, allowing to linearise the loss term to depend on the magnitude of the branch flow.

Nodal power balance with linearized losses

For each bus $i \in \mathcal{N}$, the active power balance is enforced as:

$$P_i(t) + \sum_{j \in \mathcal{N}} F_{ij}(t) = 0 \quad (7.30)$$

where $P_i(t)$ is the net exchange of nodes (positive for injection), F_{ij} represents the flow incoming to the bus i from each of the connected buses $j \in \mathcal{N}$.

Nodal exchange $P_i(t)$ is defined as the difference between internal production and bus i load demand.

$$P_i(t) = Prod_i^N(t) - Load_i^{serv}(t) \quad (7.31)$$

For each branch $b = (i, j)$, F_b represents the flow that passes through it in case of no losses. It is considered positive if it moves from node i to node j and negative in the opposite direction. The power flow contribution, from the node end point of view, is then defined as:

$$F_{ij} = \begin{cases} F_b \\ (1 - \alpha_b) F_b \end{cases} \quad F_{ji} = \begin{cases} (1 - \alpha_b) F_b & \text{if } F_b \text{ } i \Rightarrow j \\ F_b & \text{if } F_b \text{ } i \Leftarrow j \end{cases} \quad (7.32)$$

Using the branch flow module $|F_b(t)|$ it is then possible to model such losses in a linear form.

$$\begin{aligned} F_{ij}(t) &= \frac{\alpha_b - 2}{2} F_b(t) - \frac{\alpha_b}{2} |F_b(t)| \\ F_{ji}(t) &= \frac{2 - \alpha_b}{2} F_b(t) - \frac{\alpha_b}{2} |F_b(t)| \end{aligned} \quad (7.33)$$

For each branch b of the network, it is possible to define the corresponding elements of the square matrices $\mathcal{N} \times \mathcal{N}$, A and C , to properly account for the distributed branch flows. Specifically, for a branch $b = (i, j)$:

$$A_{ij} = \frac{\alpha_b - 2}{2} \quad A_{ji} = \frac{2 - \alpha_b}{2} \quad (7.34)$$

$$C_{ij} = -\frac{\alpha_b}{2} \quad C_{ji} = -\frac{\alpha_b}{2} \quad (7.35)$$

while $A_{ij} = C_{ij} = 0$ for each couple of buses ij with no interconnection and no branch associated

For each bus $i \in \mathcal{N}$, the active power balance is enforced as:

$$Prod_i^N(t) - Load_i^{serv}(t) + \sum_{b \in \mathcal{B}} A_{ib} F_b(t) + \sum_{b \in \mathcal{B}} C_{ib} |F_b(t)| = 0, \quad (7.36)$$

where the coefficients A_{ib} and C_{ib} correspond to the i -th row of the matrices A and C

The linear loss term $\sum_b C_{ib} |F_b|$ provides a first-order approximation of active power losses by allocating the estimated branch losses equally to the terminal buses. This simplified approach preserves convexity and computational efficiency, making it suitable for large-scale and multi-step optimisation within the OMS.

Discrete load shedding formulation

Let $\mathcal{D} \subseteq \mathcal{N}$ be the set of load buses. For each load bus $i \in \mathcal{D}$, an integer decision variable π_i is introduced to represent the load shedding level.

The load shedded at node i is defined as:

$$Load_i^{served}(t) = \frac{s_i(t)}{LS_{levels}} Load_i^{req}(t), \quad (7.37)$$

As the variable $s_i(t)$ is constrained as:

$$s_i(t) \in \{0, 1, \dots, LS_{levels}\}, \quad (7.38)$$

the fraction $\frac{s_i(t)}{LS_{levels}}$ represents the percentage of load at node i that will be supplied. $\frac{1}{LS_{levels}}$ represent the smallest discrete steps applicable to the shedding.

In this study $LS_{max} = 5$ have been chosen, corresponding to shedding steps of 20% of the nominal load. This discrete formulation prevents arbitrarily small load reductions and reflects the coarse granularity typically available in real-world load shedding actions at the transmission level.

Hydropower ramp-up, water flow constraints, and piecewise efficiency approximation

The OMS models hydropower operational limitations through ramp-up coupling constraints and a piecewise-linear approximation of the relationship between generated power and water usage.

Ramp-up modelling

For each hydropower plant $p \in \mathcal{G}^H$, the increase in power output between consecutive time steps is limited. A dedicated variable $RU_p(t)$ is introduced to represent the time-dependent upper bound induced by ramp-up constraints. In this way it is also possible to connect the different steps of the horizon under optimisation.

- First time step: when optimisation is initialised at time $t = 1$, the ramp-up bound is computed from the known previous utilisation ratio $Prod_p^{prev}$:

$$RU_p(1) = Prod_p^{prev} + (\Delta_p^{ru}) P_p^{\max} \quad \forall p \in \mathcal{G}^H, \quad (7.39)$$

where Δ_p^{ru} is the relative ramp-up limit and $P_p^{\max}$ is the nominal power.

- Subsequent time steps: for $t > 1$, the bound is linked to the previous-step production set-point:

$$RU_p(t) - Prod_p(t-1) = \Delta_p^{ru} P_p^{\max}, \quad \forall p \in \mathcal{G}^H, t > 1. \quad (7.40)$$

Finally, the ramp-up limitation is enforced through:

$$Prod_p(t) \leq RU_p(t), \quad \forall p \in \mathcal{G}^H, \forall t \in \mathcal{T}. \quad (7.41)$$

Available water flow

The maximum usable water flow is defined as the sum of the stored reservoir volume and the river inflow:

$$Q_p^{ava}(t) = Reserve_p(t) + Flow_p(t), \quad \forall p \in \mathcal{G}^H, \forall t \in \mathcal{T}. \quad (7.42)$$

The flow used is constrained by availability:

$$Q_p(t) \leq Q_p^{ava}(t), \quad \forall p \in \mathcal{G}^H, \forall t \in \mathcal{T}. \quad (7.43)$$

Piecewise-linear approximation of hydropower efficiency

The relationship between generated power and required water flow is approximated through a piecewise-linear lower envelope. What is needed is a way to express the power output of the plant as a function of the flow entering the turbine. Starting by expressing η as a function of Q_r in Eq. 7.12, Eq. 7.12 leads to:

$$Prod_p(t) = \beta \cdot \eta(Q_{r,p}(t)) \cdot Q_p(t) = \beta \cdot \eta(Q_{r,p}(t)) \cdot Q_{r,p}(t) \cdot Q_{nom,p} \quad (7.44)$$

This leads to:

$$Prod_p(t) = \beta \cdot \eta'(Q_{r,p}(t)) \cdot Q_{nom,p} \quad (7.45)$$

Table 7.18: Points for the piecewise approximation of η'

N_{point}	1	2	3	4	5
η'	0	0.4	0.72	0.81	1
Q_r	0	0.5	0.8	0.9	1.2

where $\eta'(Q_{r,p}(t))$ collects all the nonlinearities and represents the relationship that must be approximated with a piecewise-linear function of the form $\eta'_{\text{piece}} \cdot Q_{r,p}(t)$.

To identify the breakpoints of the piecewise approximation, Table 7.9 can be used. The points of η' are computed as the product of the values of $\eta(Q_r)$ and the corresponding Q_r , as reported in Table 7.18.

Let $z \in \{1, \dots, N_{\text{piece}}\}$ index the linear segments, where $N_{\text{piece}} = N_{\text{points}} - 1$. The coefficients a_z and b_z are defined as:

$$a_z = \frac{\eta'(z+1) - \eta'(z)}{Q_r(z+1) - Q_r(z)} \quad (7.46)$$

$$b_z = \eta'(z) - a_z Q_r(z) \quad (7.47)$$

The values $Q_r(z)$ act as thresholds x_z to regulate the switch from one piece to another. A nominal conversion factor $Q_{1\text{MW}h}^{\text{nom}}$ is used to normalise water usage. Using the parameters just defined, Eq. 7.45 can be rewritten for each piece as:

$$\frac{\text{Prod}_p(t)}{\beta} = (a_z \cdot Q_{r,p}(t) + b_z) \cdot Q_{\text{nom},p} = a_z \cdot Q_p(t) + b_z \cdot Q_{\text{nom},p} \quad (7.48)$$

A binary variable $B_p^z(t) \in \{0, 1\}$ is introduced to keep track of which segment is active through a big- M logic. One binary variable is required for each transition between pieces. The coupling between $B_p^z(t)$ and the flow variable $Q_p(t)$ is defined as follows.

For $z \in \{2, \dots, N_{\text{piece}}\}$:

$$B_{z,p}(t) \leq 1 + \frac{Q_{r,p}(t) - x_z}{M} \Rightarrow B_{z,p}(t) - \frac{Q_p(t)}{M Q_{\text{nom},p}} \leq 1 - \frac{x_z}{M} \quad (7.49)$$

$$B_{z,p}(t) \geq \frac{Q_{r,p}(t) - x_z}{M} \Rightarrow -B_{z,p}(t) + \frac{Q_p(t)}{M Q_{\text{nom},p}} \leq \frac{x_z}{M} \quad (7.50)$$

with $M \gg 0$.

These constraints, combined with the bound $B_{z,p}(t) \in \{0, 1\}$, generate the required set of binary variables. This set, together with Eq. 7.48 and defining $z' = z + 1$, allows to write for each piece the following pair of constraints:

$$\frac{Prod_p(t)}{\beta} \leq a_z Q_p(t) + b_z Q_{nom,p} + MB_{z',p}(t) + M(1 - B_{z,p}(t)) \quad (7.51)$$

$$\frac{Prod_p(t)}{\beta} \geq a_z Q_p(t) + b_z Q_{nom,p} - MB_{z',p}(t) - M(1 - B_{z,p}(t)) \quad (7.52)$$

Rearranging all terms on the left-hand side yields:

$$\frac{Prod_p(t)}{\beta} - a_z Q_p(t) - MB_{z',p}(t) + MB_{z,p}(t) \leq b_z Q_{nom,p} + M \quad (7.53)$$

$$-\frac{Prod_p(t)}{\beta} + a_z Q_p(t) - MB_{z',p}(t) + MB_{z,p}(t) \leq -b_z Q_{nom,p} + M \quad (7.54)$$

Equations 7.53 and 7.54 must be modified for the first and last piece:

- If $z = 1$, the contributions involving $B_{z,p}(t)$ must be ignored, as this variable does not exist.
- If $z = N_{\text{piece}}$, the contributions involving $B_{z',p}(t)$ must be ignored for the same reason.

The defined set of constraints and support variables finally permit to fully implement the piecewise approximation of the efficiency.

7.4.3 Variable bounds and cost function

To ensure physical feasibility and numerical stability, the optimisation variables are constrained through explicit lower and upper bounds. These bounds encode operational limits of generators, admissible ranges for branch flows, and the discrete nature of load shedding.

Bounds

Nodal injections are bounded by the maximum generation that can be injected into each bus, obtained from the aggregation of the plants connected to the node. In particular, wind generation is treated as non-dispatchable (must-take), hence its lower and upper bounds coincide with the available wind power.

Hydropower outputs are bounded between zero and available capacity, which depends on the status of the plant and the availability of resources. Branch flows are bounded symmetrically to allow both directions of power transfer, while auxiliary absolute-value variables are constrained to be non-negative. The load shedding variables are bounded between zero and the nominal demand in each bus, and the corresponding discrete shedding level is bounded by the maximum number of shedding steps.

Cost function

The OMS objective is composed of three sub-objectives.

- It prioritises non-dispatchable renewable generation while preserving water resources and maintaining secure system operation. This is translated into minimising the water flow entering the hydropower plants and, consequently, hydropower production:

$$Obj_1 = \min \sum_{t \in \mathcal{T}} \sum_{p \in \mathcal{G}^H} Q_p(t) \quad (7.55)$$

- It aims to avoid curtailment of the served load by supplying the full demand whenever possible, allowing load shedding only when strictly necessary to preserve feasibility:

$$Obj_2 = \max \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{D}} Load_i^{serv}(t) = \min \sum_{t \in \mathcal{T}} \sum_{i \in \mathcal{D}} (-Load_i^{serv}(t)) \quad (7.56)$$

- It minimises the auxiliary variables introduced to represent absolute values, which, in combination with the associated constraints, ensure that these variables correctly match the true absolute values:

$$Obj_3 = \min \sum_{t \in \mathcal{T}} \sum_{b \in \mathcal{B}} |F_b(t)| \quad (7.57)$$

The complete objective is a linear weighted combination of these three parts:

$$Obj_{tot} = \alpha_1 Obj_1 + \alpha_2 Obj_2 + \alpha_3 Obj_3. \quad (7.58)$$

Since the primary purpose of the OMS is to serve the load demand as much as possible, to ensure that load shedding is selected only when serving the load is infeasible, it is important that $\alpha_2 \gg \alpha_1$.

Note that even if it is good practice to explicitly account Obj_3 it may be partially redundant with Obj_1 . Minimising the flow of used water while meeting the load demand implies minimising hydropower production, which tends to reduce network losses. Since such losses are strictly related to the magnitude of branch flows, this also promotes smaller values of $|F_b(t)|$, providing an effect similar to Obj_3 .

7.5 Preliminary Results

The previous sections have shown in detail how the Digital Twin of the E-TN and its OMS are defined. To obtain a complete simulated environment to testing the behaviour of the E-TN, these two components must be properly interconnected.

At the beginning of the simulation step k , the OMS receives measurements of the requests and served loads, as well as the production supplied by each plant at each node. Based on this information, the OMS forecasts the future behaviour of the grid and optimises the scheduling. It should be noted that, at this stage, the forecasting process is simplified and assumed to be perfect.

The output of the OMS optimisation is a detailed schedule of the actions to be implemented on the E-TN to achieve the efficiency objectives. However, this plan is overly detailed for direct implementation, as the first set of planned setpoints for each hydropower plant represents an exact production target. As these results are obtained through a linearised model, an additional processing step is required before their application.

The setpoints generated by the OMS are therefore used to define the boundaries of the feasible space within which the OPF of the Digital Twin can search for a solution, this time fully accounting for network nonlinearities. Given the OMS solution $Prod_g^{OMS}(1)$ for all $g \in \mathcal{G}$, the bounds for the Digital Twin generation are defined as:

$$Prod_{g,UBound}^{E-TN} = Prod_g^{OMS}(1) + 0.1 \cdot Prod_g^{OMS}(1), \quad (7.59)$$

$$Prod_{g,LBound}^{E-TN} = Prod_g^{OMS}(1) - 0.1 \cdot Prod_g^{OMS}(1). \quad (7.60)$$

The Load shedding parameters are also passed as inputs to the Digital Twin. Instead of relying on the simplified one-step shedding mechanism previously implemented, demand curtailment is now directly regulated by




Figure 7.19: Coupling of the Digital Twin and the OMS of the E-TN

the OMS output. For all $i \in \mathcal{N}$, this is expressed as:

$$s_i^{E-TN} = \frac{s_i^{OMS}(1)}{LS_{\text{levels}}}. \quad (7.61)$$

By constraining both the generation search space and the shed demand of the Digital Twin according to the OMS output, the control loop is effectively closed.

The resulting architecture enables a detailed simulation of the E-TN behaviour in a wide range of future scenarios and case studies. These scenarios typically involve major disturbances to the system that would be unrealistic, unsafe, or impractical to reproduce on the physical grid.

The proposed framework therefore allows the response of the entire transmission system to be evaluated even under highly critical conditions, providing valuable information on its robustness, flexibility, and operational limits.

Case Study 1: Failure of Transmission Lines

By simulating transmission line failures, it is possible to assess how changes in network topology affect the ability of the grid to supply the demanded loads. Although this represents a preliminary test, the results already highlight a strong spatial asymmetry in the system's resilience.

In particular, the eastern part of the country appears to be extremely exposed to this type of contingency, as it is essentially connected to the main grid through a limited number of corridors. In contrast, the presence of a double-layer transmission structure in the central and western regions




Figure 7.20: Case Study: Lines Failure

provides a higher degree of redundancy, improving their ability to withstand line outages. An example of such cases is presented in Figure 7.20

Case Study 2: Failure of Power Generation Plants

Another relevant scenario concerns the failure of major power production plants. In this case, it is of particular interest to evaluate the capability of the grid to absorb such severe disturbances through the activation of load shedding strategies and the re-dispatch of the remaining generation.

The non-ideal behaviour of production units, whose power output increases are constrained by ramp-up limits, plays a crucial role in determining how quickly the system can recover and compensate for the missing generation.

This case study therefore provides insight about the dynamic limitations of the system and the effectiveness of the OMS in coordinating a constrained response. Figure 7.21 show an example of such a scenario.

Case Study 3: Shift of Nominal Operating Conditions

A third important scenario is related to the anticipated long-term increase in load demand. In the coming years, national electricity consumption is




Figure 7.21: Case Study: Plants Failure

expected to grow significantly, potentially doubling with respect to current levels. It is therefore essential to study the behaviour of the entire system as this shift in nominal operating conditions occurs.

This analysis highlights the areas of the grid that are most stressed by the increasing demand and allows an assessment of whether the current and planned generation and transmission capabilities are sufficient to guarantee acceptable performance in terms of security of supply and quality of service.

In the example shown in Figure 7.22 it's possible to see how, in such a case, the cutoff of the loads may be needed even if no major failures are present into the network




Figure 7.22: Case Study: Shift of the nominal conditions

7.6 Conclusion & Future Development

In conclusion, this chapter presented a comprehensive methodology for building a simulated environment for the development and testing of a monitoring and operational management system aimed at improving the performance of a transmission network in terms of resilience and flexibility.

By coupling a high-fidelity Digital Twin OPF-based of the E-TN with a Linearised Optimal Power Flow-based Operational Management System, it is possible to analyse the behaviour of the grid under a wide range of operating conditions and severe contingencies.

One of the main strengths of the proposed approach lies in its generality and scalability. The same modelling framework can be applied to different power systems with limited adaptations, enabling the investigation of scenarios that are difficult, risky or outright impossible to test on a physical grid. In this sense, the Digital Twin becomes a powerful tool not only for operational studies, but also for the validation of control strategies and resilience-orientated assessments.

Future developments of this work may focus on a more systematic evalu-

ation of system resilience by introducing quantitative indicators such as the LSR, the total energy not served and the recovery time following major disturbances. By linking these metrics with the different simulated scenarios, it will be possible to objectively assess the impact of the OMS on the overall system performance and to compare alternative operational strategies.

In addition, to ensure full applicability in real-world conditions, the OMS should be extended by integrating dedicated forecasting blocks for both load demand and renewable generation. The inclusion of uncertainty-aware forecasting and probabilistic optimisation would further enhance the robustness of the decision-making process, allowing the system to proactively anticipate critical conditions rather than merely reacting to them.

These extensions represent a natural continuation of the methodology presented in this chapter and open the way to more advanced and realistic operational studies.

This chapter naturally builds upon the concepts and tools introduced in the previous chapters. While earlier sections focused on the modelling of local management systems and distributed control strategies to address what can happen inside of a bus of the network, here these elements are brought together within a unified framework that operates at transmission level.

In this sense, the Operational Management System acts as the logical extension of the lower-level control architectures, translating local flexibility and renewable variability into coordinated, system-wide decisions.

The integration of monitoring, optimisation and simulation within a closed-loop architecture highlights how the methodologies discussed throughout the thesis can be combined to address resilience, security of supply and efficient resource usage in a coherent manner.

Chapter 8

Conclusion

This chapter summarises the contribution and presents the conclusion of the work.

This work is framed within the context of the energy transition, at a historical moment in which the traditional model of energy generation and distribution is no longer sustainable. The paradigm based on large centralised power plants is progressively being replaced by a system dominated by renewable energy sources, particularly solar and wind. Although this transition is necessary to achieve long-term sustainability, it also introduces a set of significant operational challenges.

One of the main obstacles lies in the spatial distribution of the generation. Energy production is no longer concentrated in a few high-voltage nodes but is instead widely distributed throughout the distribution network through a large number of small-scale generation units.

This is further complicated by the intrinsic non-controllability of renewable sources, which generate energy only when environmental conditions allow it. As a result, there is a clear need for coordination mechanisms operating from the lowest levels of the system, enabling local consumption of locally produced energy and improving the alignment between generation and demand.

The primary contribution of this work does not consist in a single algorithmic solution but rather in the coherent multi-level architecture adopted to address the problem. The proposed framework is a compromise between the most used solutions available in literature and tackles coordination across different control layers, starting from the individual user level, extending to community-level and then reaching the network level, where similar issues emerge on a larger scale. This conceptual continuity represents a key element of the proposed approach.

Within this framework, particular emphasis has been placed on the development of lightweight and interpretable solutions. The focus on interpretability has been a deliberate design choice. By gaining a deeper understanding of the problem structure, it was possible to avoid the use of more powerful but computationally expensive optimisation techniques, which would have negatively affected scalability and practical deployability. In this sense, simplicity becomes an enabling factor rather than a limitation.

User coordination emerges as the central element of the proposed system. The architecture explicitly avoids rigid and fully centralised solutions that impose coordinated behaviour in a coercive manner. Instead, coordination is achieved through mechanisms that preserve individual interests whenever possible, while providing guarantees of fairness and transparency.

This aspect is particularly relevant, as the success of the energy transition

ultimately depends on the active and voluntary participation of end users, who must be able to contribute without being unfairly penalised.

The obtained results show that the desired system-level behaviour can be achieved through scalable, stable and low-complexity solutions without sacrificing coordination effectiveness. The relevance of these results lies not in the absolute level of optimisation achieved but in demonstrating that such an approach is practically feasible.

All these considerations are summarized in Table 8.1, where the proposed methodology is compared with the approaches available in the literature across a set of key aspects. It is important to note that the proposed methodology is not necessarily the best option for every individual aspect. Rather, it represents a balanced compromise that accounts for real-world constraints and is designed to enable practical system deployment.

Aspect	Centralized	ADMM MAS	Heuristic	Proposed
Intelligence	Central	Distributed	Local	Hierarchical
Edge Control	Central	Local	Rule-based	Rule-based
Optimization	setpoints MILP/ convex	optimization Convex + consensus	Metaheuristic	Mixed Approaches
Forecasting	Central	Local/ shared	Local	Local (lightweight)
Scalability	Limited	High*	High	High
Communication	Moderate	Iterative*	Low	Low
Computation	High	Distributed	Low	Low/ moderate

*Convergence dependent on coupling strength and tuning.

Table 8.1: Compact comparison of EMS coordination paradigms and proposed architecture

Future developments will begin by relaxing the simplifying assumptions introduced throughout this work, such as the assumption of perfect forecasting. Although these assumptions are not realistic, they allow the decoupling of forecasting and optimisation problems, enabling a clearer understanding and analysis of the mechanisms and dynamics of the two aspects independently while avoiding concealed structural interactions.

A key aspect of future research will also be the explicit treatment of fore-

casting uncertainty across all control levels, from local load and generation estimation to community and network coordination.

Then the focus will shift on the full integration of the proposed control layers, the deployment of the system on embedded platforms and experimental validation in real-world settings, such as operational energy communities.

In addition, communication constraints, latency, reliability and user acceptance will need to be systematically addressed, as they become critical when transitioning from simulation to real deployment.

In addition, the integration of conversational AI with Energy Management Systems will represent a complementary layer instead of an alternative to traditional control. By lowering the barrier to interaction and understanding, such tools enable broader participation in energy communities, facilitate informed decision making, and ultimately improve the social sustainability of technically optimal solutions.

In conclusion, this work is a starting point rather than a final outcome. The ultimate goal is to get the proposed architecture out of the simulation environment and confront it with real-world conditions, where delays, imperfections and constraints are not exceptions, but the rule.

Chapter 9

Future Developments

*This final chapter outlines the main future developments of the proposed multi-level Energy Management System, focusing on robustness, integration and real-world applicability. Key research directions include the mitigation of forecasting uncertainty across all control layers, the systematic quantification of resilience through dedicated metrics, and the full integration of local, community and transmission-level controllers. Particular attention is given to the deployment of lightweight control algorithms on embedded platforms and their experimental validation in real Renewable Energy Communities. Finally, the integration of Large Language Models as an expert-user layer is discussed as a promising approach to combine flexible human interaction with the transparency and reliability of deterministic control architectures.*¹

¹A preliminary version of the work presented in this chapter has been published as *GPT-Powered Chatbot for Dissemination and Simulation of Renewable Energy Communities* [102].

The work presented in this thesis defines a multi-level Energy Management System capable of coordinating local, community and grid-level objectives under realistic computational and operational constraints. While the proposed architecture already demonstrates solid performance in simulation, several research directions naturally emerge to further improve robustness, integration and real-world applicability.

Mitigation of Forecast Uncertainty

A first and fundamental line of future research concerns the explicit treatment of forecasting uncertainty across the different layers of the EMS. The effective deployment of the proposed architecture in real-world scenarios depends on the ability to manage both measurement noise and prediction errors in load and generation forecasts.

This research direction has already been partially addressed, particularly at the HEMS level, where preliminary corrective mechanisms have been introduced. The results obtained are encouraging, they remain preliminary and require validation through large-scale testing on more heterogeneous datasets. Such an evaluation is necessary to assess robustness and generalisation under diverse user behaviours and operating conditions.

At the community level, future work will focus on analysing how forecasting errors propagate to the Community Management Hub (CMH), affecting coordination quality, collective objectives, and resilience indicators. Understanding this propagation is essential to quantify the impact of uncertainty on coordinated control strategies.

At the third control level, corresponding to the Operation and Monitoring System (OMS), forecasting becomes a prerequisite for real-world applicability. In its current form, the OMS operates under idealised assumptions and is therefore limited to simulation studies. Future developments will require the explicit integration of forecasting blocks for nodal loads and renewable generation, such as wind power, enabling operational decision-making under uncertainty.

Overall, future work will focus on robust EMS formulations that explicitly account for uncertainty, including ensemble forecasting techniques, online learning strategies and uncertainty-aware scheduling policies, with the goal of transitioning from a simulation-orientated framework to an implementable control architecture.

Resilience Evaluation

Another relevant direction for future research concerns the systematic improvement and quantification of resilience metrics within the proposed energy management framework. Even if the concept of resilience has already been partially addressed in this work and preliminary mechanisms have been defined, no extensive studies have yet been conducted to evaluate the large-scale impact of the Energy Management System on the resilience of the monitored system.

At the CEMH level, future developments will focus on a more structured analysis of resilience indicators, such as the LSR, under a wide range of operating conditions. In particular, it will be necessary to move beyond isolated case studies and perform large-scale evaluations over longer time horizons and heterogeneous datasets.

This includes the analysis of different community configurations, load profiles and disturbance scenarios, with the objective of identifying representative and meaningful cases that deserve deeper investigation. Rather than relying exclusively on full datasets, intermediate or “typical” scenarios may be identified to balance statistical significance with computational feasibility.

Such an analysis would enable a clearer understanding of how the presence of a coordinated HEMS influences the system’s ability to absorb disturbances, prioritise critical loads and recover from stress conditions. In this context, resilience metrics would not only describe the system behaviour, but also serve as quantitative indicators to compare different coordination strategies and control policies.

At the third control level, corresponding to the OMS, the problem becomes substantially more complex. At this scale, resilience cannot be evaluated solely through community-level indicators, and new metrics and methodologies must be introduced to capture the interaction between coordinated energy management actions and transmission-level dynamics. Future work will therefore focus on defining appropriate resilience metrics at this level and integrating them into the algorithm.

This extension would enable the study of significantly more complex scenarios than those accessible at the CEMH level, including large-scale disturbances, cascading effects and interactions across multiple voltage levels. Ultimately, such a multi-level resilience assessment would make it possible to

quantitatively evaluate how coordinated energy management strategies influence system behaviour under extreme and realistic operating conditions, providing a solid basis for comparing architectures and control philosophies.

Full Integration of the Multi-Level EMS

A direction for future research concerns the full integration of the three control layers into a unified and deployable EMS architecture. Although the roles and interactions of the individual layers are clearly defined and the first two levels have already been correctly and fully integrated, further work is required to achieve complete system coherence across all layers.

In the current implementation, the interaction between the HEMSs and the CEMH is already operational and validated in simulation. However, the third control level, corresponding to the OMS, is still applied in isolation to the transmission network case study. A fundamental next step is therefore the full integration of the OMS as a component coherently connected with the CEMHs and capable of operating on the same coordinated system.

From an implementation perspective, an important development concerns the migration of the HEMS and CEMH control codes from MATLAB to lower-level programming languages, such as C. This transition would enable deployment on embedded platforms and low-performance hardware, including simple microcontrollers. Since the proposed algorithms have been explicitly designed to be lightweight and computationally efficient, such a migration represents a realistic and natural step toward edge-level intelligence, reducing reliance on centralised computation and allowing the evaluation of real-time constraints, communication delays and scalability limits.

Ultimately, the most significant milestone toward real-world applicability is experimental validation. Future work will therefore aim at testing the proposed EMS, in particular the first two control layers, on real Renewable Energy Communities or physical electrical systems. Such an experimental deployment would provide concrete evidence of the feasibility, robustness, and effectiveness of the proposed architecture under real operating conditions, representing a decisive step from simulation-based validation to practical implementation.

Integration of LLMs as an Expert User Layer in the HMI

A further and particularly promising research direction concerns the integration of Large Language Models as an expert user layer within the Human–Machine Interface of the EMS.

The proposed EMS architecture is inherently powerful but complex, as it involves multiple control layers, forecasting models and optimisation mechanisms. However, real-world users are rarely experts in power systems or control theory. Bridging this gap requires an intermediate layer capable of translating high-level human intentions into structured actions that can be safely processed by the underlying control system.

In this context, LLMs are not introduced as an auxiliary feature but as a natural extension of the EMS paradigm. Their role is to act as an intelligent interface between human users and complex energy management systems, mediating the interaction without interfering with deterministic control logic.

The key advantage of this architecture lies in the deliberate combination of two complementary properties: the high flexibility of LLMs in handling ambiguous and unstructured user inputs, and the transparency and reliability of classical simulation and control algorithms. This separation avoids the risks associated with black-box numerical reasoning, while significantly improving usability and accessibility 9.1.




Figure 9.1: Enter Caption

The preliminary results obtained with the chatbot-based REC simulator developed in this work demonstrate the feasibility of this approach. The LLM-driven interface is capable of interpreting highly flexible natural-language inputs, guiding users through complex configuration procedures and translating unstructured requests into structured inputs for a deterministic simulator. At the same time, all numerical computations and decision logic remain entirely delegated to conventional and transparent code 9.2.





Figure 9.2: Implemented chatbot for the REC simulation

At its current stage, this integration is intentionally limited to a supervised and non-autonomous role. The LLM role is not to perform control actions, nor does it replace optimisation or forecasting modules.

Although not yet fully deployed in operational environments, the extension of this concept toward a generalised expert user layer represents a serious and promising research direction. Addresses a critical barrier to the real-world adoption of advanced EMS architectures: the mismatch between system complexity and user expertise. By embedding intelligence at the interface level rather than in the control core, this approach preserves the rigour of the system while enabling broader applicability and acceptance.



Bibliography

- [1] M. McNutt, “Climate Change Impacts,” *Science*, vol. 341, no. 6145, p. 435, 2013.
- [2] J. Skea, P. Shukla, A. Reisinger, R. Slade, M. Pathak, A. Al Khourdajie, R. van Diemen, A. Abdulla, K. Akimoto, M. Babiker, Q. Bai, I. Bashmakov, C. Bataille, G. Berndes, G. Blanco, K. Blok, M. Bustamante, E. Byers, L. Cabeza, and H. Winkler, “Mitigating Climate Change,” 2022.
- [3] Y. R. Li, F. Nejabatkah, and H. Tian, *Smart Hybrid AC/DC Microgrids*. John Wiley & Sons, 2018.
- [4] A. Keyhani, M. N. Marwali, and M. Dai, *Integration of Green and Renewable Energy in Electric Power Systems*. John Wiley & Sons, 2010. [Online]. Available: <https://doi.org/10.1002/9780470556771>
- [5] Terna S.p.A., “Electricity demand,” 2025, press release, 16 January 2025. [Online]. Available: https://download.terna.it/terna/Terna_electricity_demand_2024_8dd361f3416a325.pdf
- [6] International Renewable Energy Agency (IRENA), “Italy – Renewable Energy Statistics 2024: Country Profile,” 2024. [Online]. Available: https://www.irena.org/IRENADocuments/Statistical_Profiles/Europe/Italy_Europe_RE_SP.pdf
- [7] A. M. Fagetan, “Transposition and Implementation of European Union Renewable Energy Legislation in France, Italy, and Germany: A Regulatory Perspective and a Comprehensive Analysis of Opportunities and Challenges,” *Laws*, vol. 15, no. 1, p. 3, 2026.
- [8] European Commission, “Renewable Energy Targets – Revised Directive (EU/2023/2413),” https://energy.ec.europa.eu/topics/renewable-energy/renewable-energy-directive-targets-and-rules/renewable-energy-targets_en,2023,eU_policy_webpage (accessed 2025).
- [9] P. Basilico, A. Biancardi, I. D’Adamo, and M. Gastaldi, “Energy Communities Toward Sustainable Development: The Role of Economic Factors in

- a Social Analysis,” *Sustainable Development*, vol. 33, no. 4, pp. 5587–5603, 2025.
- [10] M. Ertz, X. Cao, and J. M. Barragán-Maravilla, “The Prosumer,” *Encyclopedia*, vol. 4, no. 3, pp. 1263–1278, 2024.
 - [11] N. Castañeda-Arias, N. L. Díaz-Aldana, A. Luna Hernandez, and A. L. Jutínico, “Energy Management in Microgrid Systems: A Comprehensive Review Toward Bio-Inspired Approaches for Enhancing Resilience and Sustainability,” *Electricity*, vol. 6, no. 4, p. 73, 2025.
 - [12] GridWise Alliance, “Improving Electric Grid Reliability and Resilience: Lessons Learned from Superstorm Sandy and Other Extreme Events,” Grid-Wise Alliance, Technical Report, 2013.
 - [13] International Energy Agency (IEA), “Power Systems in Transition: Challenges and opportunities ahead for electricity security,” 2020, fuel report, IEA (27 Oct 2020), CC BY 4.0. [Online]. Available: <https://www.iea.org/reports/power-systems-in-transition>
 - [14] European Parliament and Council of the European Union, “Regulation (EU) 2019/943 on the internal market for electricity,” 2019, official Journal of the European Union, L158. [Online]. Available: <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32019R0943>
 - [15] International Energy Agency (IEA), “Italy 2023: Energy Policy Review,” 2023. [Online]. Available: <https://www.iea.org/reports/italy-2023>
 - [16] S. Stoft, *Power System Economics: Designing Markets for Electricity*. IEEE Press, 2002.
 - [17] M. H. Imani, “Empirical Analysis of Inter-Zonal Congestion in the Italian Electricity Market,” *Energies*, vol. 17, no. 23, p. 5901, 2024.
 - [18] D. Alessi, C. Bonaldo, and F. Fontini, “The Impact of Utility-Scale Renewable Energy Production on Italian Electricity Prices,” *Italian Economic Journal*, 2025.
 - [19] A. J. Conejo, M. Carrión, and J. M. Morales, *Decision Making Under Uncertainty in Electricity Markets*. Springer, 2010.
 - [20] C. Winzer and P. Hensler-Ludwig, “Design and Impact of Grid Tariffs,” *Energies*, vol. 17, no. 6, p. 1364, 2024.
 - [21] Terna S.p.A., “Distributed Resources in the Italian Ancillary Services Market: taking stock after two years,” 2020.
 - [22] I. Saviuc, A. P. Chema Lopez, K. Rollert, and P. Bertoldi, “Explicit Demand Response for Small End-Users and Independent Aggregators: Status, Context, Enablers and Barriers,” 2022. [Online]. Available: <https://publications.jrc.ec.europa.eu/repository/handle/JRC129745>

- [23] A. Scrocca, E. Daccò, D. Andreotti, G. Rancilio, F. Bovera, D. Falabretti, and M. Delfanti, “Local flexibility markets in Europe: A critical review of market designs, operational maturity and stakeholder perspectives,” *Renewable and Sustainable Energy Reviews*, vol. 226, p. 116434, 2026. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1364032125011074>
- [24] G. mercato Elettrico (GME), “LOCAL FLEXIBILITY MARKET REGULATIONS,” 2023. [Online]. Available: https://www.mercatoelettrico.org/Portals/0/Documents/en-us/20230807_Regolamento_MLF_EN.pdf#:~:text=,of%203%20August%202023%20no
- [25] A. I. R. A. for Energy, “Relazione Annuale 2023 (annual report on the state of services).”
- [26] I. E. Agency, “Italy Electricity Security Policy – Analysis,” 2022, summarizes Italy’s transmission distribution structure, cross-border links, and the need for grid investment to meet renewables targets. [Online]. Available: <https://www.iea.org/articles/italy-electricity-security-policy>
- [27] A. Dimovski, C. Caminiti, G. Rancilio, M. Ricci, B. Di Pietra, and M. Merlo, “Assessment of Italian Distribution Grids and Implications for Energy Communities’ Integration: A Focus on Reverse Power Flow and Energy Balance,” *Energies*, vol. 18, no. 5, p. 1255, 2023.
- [28] G. De Luca, L. Piscitelli, E. Rossi Scarpa Gregorj, and P. Viganò, “Italy, Electricity Grid Under Pressure: The Challenge of Saturation Between Input and Output,” *Advant Nctm Energy Law Insight*, June 24, 2025.
- [29] G. Kaldor, “Electricity Grids at the Heart of the Energy Transition,” *Renewable Matter*, no. 59, 2024.
- [30] European Parliament and Council of the European Union, “Internal Electricity Market Directive (EU) 2019/944,” 2019. [Online]. Available: <https://eur-lex.europa.eu/eli/dir/2019/944/oj/eng>
- [31] —, “Renewable Energy Directive (EU) 2018/2001 (RED II),” 2018. [Online]. Available: <https://eur-lex.europa.eu/eli/dir/2018/2001/oj/eng>
- [32] J. Lowitzsch, C. Hoicka, and F. van Tulder, “Renewable energy communities under the 2019 European Clean Energy Package – Governance model for the energy clusters of the future?” *Renewable and Sustainable Energy Reviews*, vol. 122, p. 109489, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1364032119306975>
- [33] European Commission, “Energy communities: citizen-driven energy actions that contribute to the clean energy

- transition,” EU Commission web page, 2025. [Online]. Available: https://energy.ec.europa.eu/topics/markets-and-consumers/energy-consumers-and-prosumers/energy-communities_en
- [34] D. Proto, “Renewable Energy Communities as an Enabling Framework to Boost Flexibility and Promote the Energy Transition,” *Energies*, vol. 15, no. 23, p. 8874, 2022.
- [35] Clean Air Task Force, “Clean Energy from the Ground Up: Energy Communities in the European Union,” Report, Clean Air Task Force, 2024. [Online]. Available: <https://www.catf.us/resource/clean-energy-ground-up-energy-communities-european-union/>
- [36] Repubblica Italiana, “Decreto Legislativo 8 novembre 2021, n. 199 - Attuazione della direttiva (UE) 2018/2001 sulla promozione dell’uso dell’energia da fonti rinnovabili,” 2021. [Online]. Available: <https://www.normattiva.it/uri-res/N2Ls?urn:nir:stato:decreto.legislativo:2021-11-08;199>
- [37] Gestore dei Mercati Energetici (GME), “Decreto CACER e TIAD – regole operative per l’accesso al servizio per l’autoconsumo diffuso e al contributo PNRR,” 2023. [Online]. Available: <https://www.mercatoelettrico.org/it/>
- [38] G. U. Magni, D. Bricca, and S. Familiari, “Economic incentives for Renewable Energy Communities: a scenario analysis in the transition process between the experimental and definitive Italian policy framework,” *J. of Sustainable Development of Energy, Water and Environment Systems*, vol. 13, no. 2, 2025.
- [39] BRIGHT Project Consortium, “Energy communities aggregation for optimal flexibility management (Italian Pilot, Terni),” H2020 BRIGHT Project website, 2021. [Online]. Available: <https://www.brightproject.eu/italy/>
- [40] T. Ferrucci, D. Fioriti, D. Poli, S. Barberis, A. Vannoni, F. Roncallo, C. Taccanelli, and V. Gambino, “Battery energy storage systems for ancillary services in Renewable energy communities,” in *E3S Web of Conferences (Proceedings of SUPEHR 2023)*, vol. 414. EDP Sciences, 2023, p. 03011.
- [41] M. Menegatto, A. Bobbio, G. Freschi, and A. Zamperini, “The Social Acceptance of Renewable Energy Communities: The Role of Socio-Political Control and Impure Altruism,” *Climate*, vol. 13, no. 55, 2025.
- [42] R. Caferra, A. Colasante, I. D’Adamo, A. Morone, and P. Morone, “Interacting locally, acting globally: Trust and proximity in social networks for the development of energy communities,” *Scientific Reports*, vol. 13, p. 16636, 2023.
- [43] G. Volpato, G. Carraro, E. D. Cin, and S. Rech, “On the Different Fair Allocations of Economic Benefits for Energy Communities,” *Energies*, vol. 17, p. 4788, 2024.

- [44] V. Casalicchio, G. Manzolini, M. G. Prina, and D. Moser, "Optimal Allocation Method for a Fair Distribution of Benefits in an Energy Community," *Solar RRL*, vol. 6, p. 2100473, 2021.
- [45] International Energy Agency (IEA), "Executive summary - Unlocking the Potential of Distributed Energy Sources," 2022. [Online]. Available: <https://www.iea.org/reports/unlocking-the-potential-of-distributed-energy-resources/executive-summary>
- [46] M. S. Nkambule, A. N. Hasan, and T. Shongwe, "A review of intelligent control strategies for energy management systems in microgrids," *Energy Conversion and Management: X*, vol. 28, p. 101323, 2025. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2590174525004556>
- [47] A. Elmouatamid, R. Ouladsine, M. Bakhouya, N. El Kamoun, M. Khaidar, and K. Zine-Dine, "Review of Control and Energy Management Approaches in Micro-Grid Systems," *Energies*, vol. 14, no. 1, 2021. [Online]. Available: <https://www.mdpi.com/1996-1073/14/1/168>
- [48] A. Esparza, M. Blondin, and J. P. F. Trovão, "A Review of Optimization Strategies for Energy Management in Microgrids," *Energies*, vol. 18, no. 13, 2025. [Online]. Available: <https://www.mdpi.com/1996-1073/18/13/3245>
- [49] H. T. Dinh, D. Kim, and D. Kim, "MILP-based optimal day-ahead scheduling for a system-centric community energy management system supporting different types of homes and energy trading," *Scientific Reports*, vol. 12, p. 18305, 2022.
- [50] A. Akter, E. I. Zafir, N. H. Dana, R. Joysoyal, S. K. Sarker, L. Li, S. M. Mueen, S. K. Das, and I. Kamwa, "A review on microgrid optimization with meta-heuristic techniques: Scopes, trends and recommendation," *Energy Strategy Reviews*, vol. 51, p. 101298, 2024.
- [51] P. Wilk, N. Wang, and J. Li, "Multi-Agent Reinforcement Learning for Smart Community Energy Management," *Energies*, vol. 17, no. 20, p. 5211, 2024.
- [52] Y. Parag and B. K. Sovacool, "Electricity market design for the prosumer era," *Nature Energy*, vol. 1, p. 16032, 2016.
- [53] H. T. Dinh, D. Kim, and D. Kim, "MILP-based optimal day-ahead scheduling for a system-centric community energy management system supporting different types of homes and energy trading," *Scientific Reports*, vol. 12, p. 18305, 2022.
- [54] P. Kundur, *Power System Stability and Control*. McGraw-Hill, 1994.
- [55] A. J. Wood, B. F. Wollenberg, and G. B. Sheblé, *Power Generation, Operation, and Control*, 3rd ed. Wiley, 2013.

- [56] A. Kudzin, S. Takayama, and A. Ishigame, "Energy Management Systems (EMS) for a Decentralized Grid: A Review and Analysis of the Generation and Control Methods Impact on EMS Type and Topology," *IET Renewable Power Generation*, vol. 19, no. 1, jan 2025. [Online]. Available: <https://ietresearch.onlinelibrary.wiley.com/doi/10.1049/rpg2.70008>
- [57] U. Kumar, F. Ding, C. Zhao, H. Padullaparti, and X. Jin, "A Model-Predictive Hierarchical-Control Framework for Aggregating Residential DERs to Provide Grid Regulation Services," in *Proceedings of the IEEE Conference on Innovative Smart Grid Technologies (ISGT)*. Washington, DC, USA: IEEE, Feb. 2020, preprint, National Renewable Energy Laboratory (NREL).
- [58] M. Elkazaz, M. Sumner, and D. Thomas, "A hierarchical and decentralized energy management system for peer-to-peer energy trading," *Applied Energy*, vol. 291, p. 116766, 2021. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0306261921002737>
- [59] R. Lasseter and P. Paigi, "Microgrid: a conceptual solution," pp. 4285–4290 Vol.6, 2004.
- [60] J. M. Guerrero, J. C. Vasquez, J. Matas, L. G. de Vicuna, and M. Castilla, "Hierarchical Control of Droop-Controlled AC and DC Microgrids—A General Approach Toward Standardization," *IEEE Transactions on Industrial Electronics*, vol. 58, no. 1, pp. 158–172, 2011.
- [61] J. C. Vasquez, J. Guerrero, J. Miret, M. Castilla, and L. Vicuna, "Hierarchical Control of Intelligent Microgrids," *Industrial Electronics Magazine, IEEE*, vol. 4, pp. 23 – 29, 01 2011.
- [62] A. Ahmadifar, M. Ginocchi, M. S. Golla, F. Ponci, and A. Monti, "Development of an Energy Management System for a Renewable Energy Community and Performance Analysis via Global Sensitivity Analysis," *IEEE Access*, vol. 11, pp. 4131–4154, 2023.
- [63] M. Pasqui, A. Felice, M. Messagie, T. Coosemans, T. T. Bastianello, D. Baldi, P. Lubello, and C. Carcasci, "A new smart batteries management for Renewable Energy Communities," *Sustainable Energy, Grids and Networks*, vol. 34, p. 101043, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S2352467723000516>
- [64] E. Cutore, R. Volpe, R. Sgroi, and A. Fichera, "Energy management and sustainability assessment of renewable energy communities: The Italian context," *Energy Conversion and Management*, vol. 278, p. 116713, 2023. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0196890423000596>

- [65] R. Arghandeh, L. Mehrmanesh, and L. Mili, "On the definition of cyber-physical resilience in power systems," *Renewable and Sustainable Energy Reviews*, vol. 58, p. 1060–1069, 01 2016.
- [66] A. Dehghani, M. Sedighizadeh, and F. Haghjoo, "An overview of the assessment metrics of the concept of resilience in electrical grids," *International Transactions on Electrical Energy Systems*, 2021. [Online]. Available: <https://doi.org/10.1002/2050-7038.13159>
- [67] S. Mishra, T. Kwasnik, and K. Anderson, "Microgrid Resilience: A Holistic and Context-Aware Resilience Metric," 06 2021.
- [68] A. Momen, M. A. Ranjbar, M. S. Sepasian, and M. T. Ameli, "Chapter 6 - Resilience metrics: Assessment of distribution systems performance," in *Future Modern Distribution Networks Resilience*, M. T. Ameli, K. Jalilpoor, S. Azad, M. Daneshvar, M. S. Sepasian, B. Mohammadi-Ivatloo, and M. Shafie-Khah, Eds. Elsevier, 2024, pp. 141–169. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/B9780443160868000075>
- [69] ARERA, "Orientamenti per la regolazione della resilienza del sistema elettrico di distribuzione," Autorità di Regolazione per Energia Reti e Ambiente, Documento di consultazione 322/2017/R/eel, 2017.
- [70] H. Raoufi, V. Vahidinasab, and K. Mehran, "Power Systems Resilience Metrics: A Comprehensive Review of Challenges and Outlook," *Sustainability*, vol. 12, no. 22, 2020. [Online]. Available: <https://www.mdpi.com/2071-1050/12/22/9698>
- [71] L. Das, S. Munikoti, B. Natarajan, and B. Srinivasan, "Measuring smart grid resilience: Methods, challenges and opportunities," *Renewable and Sustainable Energy Reviews*, vol. 130, p. 109918, 2020. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S1364032120302094>
- [72] Rocky Mountain Institute, "During a Historic Hurricane Season, Microgrids Kept Communities Running," RMI – Rocky Mountain Institute, Insight Article, Feb. 2025, published February 24, 2025.
- [73] L. Becchi, M. Bindi, M. Intravaia, L. Sabino, C. C. Garzon Alfonso, F. Grasso, F. R. Fulginei, and F. Crescimbeni, "Low-Complexity Neural Networks for Electrical Load Forecasting in Renewable Energy Communities," in *IEEE International Conference on Consumer Electronics*, 2024, pp. 61–66.
- [74] L. Becchi, C. Camerota, M. Intravaia, M. Bindi, A. Luchetta, and T. Pecorella, "Load Profile Generation for Robust Optimization: A Stochastic Approach Based on Conditional Probability Approximation," in *Proc. IEEE EEEIC / I&CPS Europe*, 2025, pp. 1–6.

- [75] A. Gasparin, S. Lukovic, and C. Alippi, "Deep learning for time series forecasting: the electric load case," *CAAI Transactions on Intelligent Technology*, vol. 7, no. 1, pp. 1–25, 2022.
- [76] J.-F. Chen, W.-M. Wang, and C.-M. Huang, "Analysis of an adaptive time-series autoregressive moving-average (ARMA) model for short-term load forecasting," *Electric Power Systems Research*, vol. 34, no. 3, pp. 187–196, 1995.
- [77] C.-M. Huang, C.-J. Huang, and M.-L. Wang, "A particle swarm optimization to identifying the ARMAX model for short-term load forecasting," *IEEE Transactions on Power Systems*, vol. 20, no. 2, pp. 1126–1133, 2005.
- [78] K. Y. Lee, Y. T. Cha, and J. H. Park, "Short-term load forecasting using an artificial neural network," *IEEE Transactions on Power Systems*, vol. 7, no. 1, pp. 124–132, 1992.
- [79] D. C. Park, M. A. El-Sharkawi, R. J. Marks, L. E. Atlas, and M. J. Damborg, "Electric load forecasting using an artificial neural network," *IEEE Transactions on Power Systems*, vol. 6, no. 2, pp. 442–449, 1991.
- [80] S. N. Yelamali and K. Byahatti, "Electricity short term load forecasting using Elman recurrent neural network, booktitle = Proceedings of the International Conference on Advances in Recent Technologies in Communication and Computing," 2010, pp. 351–354.
- [81] W. Kong, Z. Y. Dong, Y. Jia, D. J. Hill, Y. Xu, and Y. Zhang, "Short-term residential load forecasting based on LSTM recurrent neural network," *IEEE Transactions on Smart Grid*, vol. 10, no. 1, pp. 841–851, 2019.
- [82] H. Zhao, Y. Wu, L. Ma, and S. Pan, "Spatial and temporal attention-enabled transformer network for multivariate short-term residential load forecasting," *IEEE Transactions on Instrumentation and Measurement*, vol. 72, pp. 1–11, 2023.
- [83] A. Vaswani *et al.*, "Attention is all you need," in *Advances in Neural Information Processing Systems*, vol. 30, 2017.
- [84] L. Becchi, M. Bindi, M. Intravaia, F. R. Fulginei, and L. Sabino, "A Computationally Efficient Rule-Based BMS Countermeasures for Price Overlapping Issues," in *Proc. of ICECCME 2025*, 2025.
- [85] L. Becchi, M. Bindi, M. Intravaia, F. Grasso, M. Pasqui, and C. Carcasci, "Application of Control Algorithms for Battery Scheduling in Grid-Connected Energy Prosumers," in *Proc. IEEE MELECON*, 2024, pp. 932–937.
- [86] L. Becchi, M. Bindi, F. Grasso, M. Intravaia, G. M. Lozito, and A. Luchetta, "Distributed EMS Coordination via Price-Signal Control for Renewable Energy Communities," *Energies*, vol. 18, 2025.

- [87] M. Elkazaz, M. Sumner, and D. Thomas, "A Hierarchical and Decentralized Energy Management System for Peer-to-Peer Energy Trading," *Applied Energy*, vol. 291, p. 116766, 2021.
- [88] L. Urbanucci, "Limits and Potentials of Mixed Integer Linear Programming Methods for Optimization of Polygeneration Energy Systems," *Energy Procedia*, vol. 148, pp. 1199–1205, 2018.
- [89] R. C. Freibott, C. Gerbolés, Álvaro García-Sánchez, and M. Ortega-Mier, "MILP and PSO Approaches for Solving a Hydropower Reservoirs Intraday Economic Optimization Problem," *Central European Journal of Operations Research*, vol. 33, pp. 1343–1366, 2024.
- [90] M. I. M. Zakki, M. N. M. Hussain, and N. Seroji, "Implementation of Particle Swarm Optimization for Tuning of PID Controller in Arduino Nano for Solar MPPT System," *Journal of Electrical and Electronic Systems Research*, vol. 13, pp. 78–85, 2018.
- [91] Aalborg University. (2023) Large-Scale Integration of Wind Power Generation in Ethiopia (LAST-WIND). Department of Energy Technology, Aalborg University. [Online]. Available: <https://www.energy.aau.dk/research/research-groups/crom/research-and-projects/large-scale-integration-of-wind-power-generation-in-ethiopia-lastwind>
- [92] P. Juturu, T. Conlon, Y. Zhang, C. Avraam, S. Siddiqui, B. Simane, and B. Zaitchik, "Optimal grid expansion under future electricity demand for groundwater irrigation in Ethiopia," *Energy for Sustainable Development*, vol. 72, pp. 351–377, 02 2023.
- [93] B. Khan, J. M. Guerrero, S. K. Chaudhary, J. C. Vasquez, K. H. B. Frederiksen, and Y. Wu, "A Review of Grid Code Requirements for the Integration of Renewable Energy Sources in Ethiopia," *Energies*, vol. 15, no. 14, p. 5197, 2022.
- [94] Y. Yu, Y. Guan, B. Khan, J. C. Vasquez, and J. M. Guerrero, "A Review of International Grid Codes for Wind Power Integration," *Energy Conversion and Management: X*, vol. 28, p. 101278, 2025.
- [95] [Online]. Available: <https://www.lastwind.energy>
- [96] Y. T. Wassie and E. O. Ahlgren, "Understanding the load profiles and electricity consumption patterns of PV mini-grid customers in rural off-grid east africa: A data-driven study," *Energy Policy*, vol. 185, p. 113969, 2024. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0301421523005542>
- [97] M. A. H. Mondal, E. Bryan, C. Ringler, D. Mekonnen, and M. Rosegrant, "Ethiopian energy status and demand scenarios: Prospects to improve energy efficiency and mitigate GHG emissions," *Energy*, vol. 149, pp. 161–172, 2018.

- [98] H. Lakew, “Ethiopian energy sector review,” Addis Ababa, Ethiopia, 2010.
- [99] “Energy Sector Profile Projections: Volume 3 - Ethiopia,” Addis Ababa, Ethiopia, Jun. 2007.
- [100] M. Javidsharifi, H. P. Arabani, S. F. Aklilu, D. Anteneh, J. C. Vasquez, J. M. Guerrero, and Y. Guan, “Assessment of Wind Energy Potential in Ethiopia: A Case Study of the Sela Dingay Wind Farm,” in *Proc. IEEE Global Power, Energy and Communication Conference (GPECOM)*, 2025, pp. 657–662.
- [101] W. Kang, Y. Guan, Y. Yu, and S. Tan, “A GIS Integrated Framework for Optimal Planning of Ethiopian Onshore Wind Power Plants,” in *Proc. IEEE Power Systems and Energy Technologies Conference*, 2025.
- [102] L. Becchi, M. Bindi, M. Intravaia, G. M. Lozito, M. C. Piccirilli, R. Marconi, and V. Cacciavillani, “GPT-Powered Chatbot for Dissemination and Simulation of Renewable Energy Communities,” in *Proc. IEEE IEEEIC / I&CPS Europe*, 2025, pp. 1–6.

Appendix A

Datasets and Information Sources

This section summarizes the datasets and external information sources used throughout this study.

Ausgrid Dataset: This publicly available Australian dataset contains electricity consumption and photovoltaic production profiles of residential prosumers. In this work, one full year of data (8760 hourly samples per user) from 126 prosumers has been used. Additional complete yearly datasets are available online through the Ausgrid website¹.

Planet Smart City Dataset: This dataset is not publicly available, as it contains privacy-sensitive information provided by PLANET IDEA S.r.l. It includes approximately five months of data (about 3550 hourly samples per user) collected from 286 residential users located in the Milan metropolitan area (Italy). All users in this dataset are pure consumers, with no local generation assets. Information regarding data access can be requested via the official Planet Smart City website².

Ethiopian Regional Load Profiles: Hourly load profiles of the Ethiopian national electricity demand for the year 2019, disaggregated by region, were provided by the Ethiopian TSO to the CROM research group. These data

¹<https://www.ausgrid.com.au>

²<https://planetSMARTcity.it>

are not publicly available.

PVGIS: The Photovoltaic Geographical Information System (PVGIS) is an online platform developed by the European Commission Joint Research Centre. It provides solar radiation data and photovoltaic system performance estimates for almost any location worldwide, excluding polar regions. In this study, PVGIS has been extensively used to emulate the power production of photovoltaic plants at different geographical locations and for multiple historical years³.

GME: The Italian electricity market operator provides public access to historical day-ahead market prices, including the national single price (PUN) and zonal prices. In this work, price data from 2020 onward have been used⁴.

NASA POWER: NASA POWER is a global web-based platform providing meteorological and renewable energy-related data. In this study, it has been used to retrieve historical wind speed data for several locations in Ethiopia, enabling the emulation of wind power generation profiles⁵.

Open Infrastructure Map:

Open Infrastructure Map is an open-access, community-driven platform providing georeferenced information on critical infrastructure worldwide, including electricity transmission and distribution networks. In this study, the platform has been used to retrieve topographical and structural information on the Ethiopian power grid. The dataset covers multiple countries and is continuously updated through open data sources and user contributions, making it a valuable reference for large-scale comparative and planning studies.⁶

³https://joint-research-centre.ec.europa.eu/photovoltaic-geographical-information-system-pvgis_en

⁴<https://www.gme.it>

⁵<https://power.larc.nasa.gov/data-access-viewer/>

⁶<https://openinframap.org>

Appendix B

Publications

This research activity has led to several publications in international journals and conferences. These are summarized below.¹

International Journals

1. Becchi, L.; Bindi, M.; Grasso, F.; Intravaia, M.; Lozito, G.M.; Luchetta, A. “Distributed EMS Coordination via Price-Signal Control for Renewable Energy Communities”, *Energies*, vol. 18, 2025. [DOI: 10.3390/en18226072]
2. Bindi, M.; Intravaia, M.; Luchetta, A.; Lozito, G.M.; Becchi, L.; Grasso, F.; Manno, M.; Annigliato, C.M.; Ferri, F.; Bartoccini, E. “Predictive Analysis of Criticality for Underground Medium Voltage Lines”, *Electric Power Systems Research*, vol. 248, pp. 1–11, 2025. [DOI: 10.1016/j.epsr.2025.111947]
3. Becchi, L.; Belloni, E.; Bindi, M.; Intravaia, M.; Grasso, F.; Lozito, G.M.; Piccirilli, M.C. “A Computationally Efficient Rule-Based Scheduling Algorithm for Battery Energy Storage Systems”, *Sustainability*, vol. 16, 2024. [DOI: 10.3390/su162310313]
4. Pasqui, M.; Becchi, L.; Bindi, M.; Intravaia, M.; Grasso, F.; Fioriti, G.; Carcasci, C. “Community Battery for Collective Self-Consumption and Energy Arbitrage: Independence Growth vs. Investment Cost-Effectiveness”, *Sustainability*, vol. 16, 2024. [DOI: 10.3390/su16083111]

¹The author’s bibliometric indices are the following: H -index = 5, total number of citations = 111 (source: Google Scholar, Month 12, 2025).

International Conferences and Workshops

1. Becchi, L.; Bindi, M.; Intravaia, M.; Lozito, G. M.; Piccirilli, M. C.; Marconi, R.; Cacciavillani, V. “GPT-Powered Chatbot for Dissemination and Simulation of Renewable Energy Communities”, in *Proc. of the 2025 IEEE International Conference on Environment and Electrical Engineering and 2025 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I&CPS Europe 2025)*, pp. 1–6, 2025. [DOI: 10.1109/eeeic/icpseurope64998.2025.11169160]
2. Becchi, L.; Camerota, C.; Intravaia, M.; Bindi, M.; Luchetta, A.; Pecorella, T. “Load Profile Generation for Robust Optimization: A Stochastic Approach Based on Conditional Probability Approximation”, in *Proc. of the 2025 IEEE International Conference on Environment and Electrical Engineering and 2025 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I&CPS Europe 2025)*, pp. 1–6, 2025. [DOI: 10.1109/eeeic/icpseurope64998.2025.11169098]
3. Bindi, M.; Meshram, V. S.; Luchetta, A.; Becchi, L.; Intravaia, M.; Trivino, A.; Villagrasa, E. “Foreign Object Detection for Wireless Power Transfer Systems Using MLMVN”, in *Proc. of the 8th IEEE International Forum on Research and Technologies for Society and Industry Innovation (RTSI 2024)*, pp. 453–458, 2024. [DOI: 10.1109/rtsi61910.2024.10761762]
4. Corti, F.; Intravaia, M.; Bindi, M.; Lozito, G. M.; Becchi, L.; Reatti, A. “State of Charge Estimation Using a Neural Network for a Battery-Powered DC–DC Converter System”, in *Proc. of the 2025 International Conference on Clean Electrical Power (ICCEP 2025)*, pp. 561–567, 2025. [DOI: 10.1109/iccep65222.2025.11143649]
5. Becchi, L.; Belloni, E.; Bindi, M.; Intravaia, M.; Lozito, G. M.; Laudani, A. “Optical and Electrical Model for Vertical-Mounted Bifacial Solar Panels”, in *Proc. of the 10th IEEE International Symposium on Systems Engineering (ISSE 2024)*, pp. 1–6, 2024. [DOI: 10.1109/isse63315.2024.10741094]
6. Intravaia, M.; Becchi, L.; Bindi, M.; Corti, F.; Lozito, G. M.; Garzon Alfonso, C. C.; Luchetta, A.; Reatti, A. “LCC–S Compensated Wireless Power Transfer: System Optimization Using Genetic Algorithms”, in *Proc. of the 2024 International Conference on Electrical Machines (ICEM)*, pp. 1–7, 2024. [DOI: 10.1109/icem60801.2024.10700581]
7. Becchi, L.; Bindi, M.; Garzon Alfonso, C. C.; Intravaia, M.; Lozito, G. M.; Grasso, F. “A Neural Network-Based Control Strategy for Constant Current Battery Chargers with PV Source”, in *Proc. of the 8th IEEE International Forum on Research and Technologies for Society and Industry Innovation (RTSI 2024)*, pp. 459–464, 2024. [DOI: 10.1109/rtsi61910.2024.10761440]

8. Becchi, L.; Bindi, M.; Intravaia, M.; Grasso, F.; Pasqui, M.; Carcasci, C. "Application of Control Algorithms for Battery Scheduling in Grid-Connected Energy Prosumers", in *Proc. of the 2024 IEEE 22nd Mediterranean Electrotechnical Conference (MELECON)*, pp. 932–937, 2024. [DOI: 10.1109/melecon56669.2024.10608734]
9. Becchi, L.; Bindi, M.; Intravaia, M.; Sabino, L.; Garzon Alfonso, C.C.; Grasso, F.; Fulginei, F.R.; Crescimbeni, F. "Low-Complexity Neural Networks for Electrical Load Forecasting in Renewable Energy Communities", in *Digest of Technical Papers – IEEE International Conference on Consumer Electronics*, pp. 61–66, 2024. [DOI: 10.1109/isct62336.2024.10791183]
10. Intravaia, M.; Lozito, G.M.; Becchi, L.; Corti, F.; Luchetta, A.; Reatti, A. "Failure Prevention Based on Principal Component Analysis and Machine Learning for Wireless Power Transfer Systems", in *Proc. of the 2024 IEEE 22nd Mediterranean Electrotechnical Conference (MELECON)*, pp. 213–218, 2024. [DOI: 10.1109/melecon56669.2024.10608513]
11. Intravaia, M.; Becchi, L.; Bindi, M.; Costanzo, L.; Garzon Alfonso, C.C.; Meshram, V.S.; Reatti, A.; Vitelli, M. "Sensitivity Analysis of PV Produced Power in Presence of Measurement Uncertainty", in *Proc. of IEEE MetroInd4.0 & IoT*, pp. 245–250, 2024. [DOI: 10.1109/metroind4.0iot61288.2024.10584219]
12. Bindi, M.; Luchetta, A.; Grasso, F.; Piccirilli, M.C.; Becchi, L.; Zoni, M.; Bartemucci, L.; Borghesi, T. "Symbolic Analysis Role in the Reliability and Availability of Power Systems", in *Proc. of the 4th International Conference on Advanced Electrical and Energy Systems (AEES 2023)*, pp. 377–384, 2023. [DOI: 10.1109/aees59800.2023.10469260]
13. Intravaia, M.; Becchi, L.; Bindi, M.; Paolucci, L.; Grasso, F. "Autoencoders for Hourly Load Profile Reconstruction in Renewable Energy Communities", in *Proc. of EUROCON 2023 – 20th International Conference on Smart Technologies*, pp. 280–285, 2023. [DOI: 10.1109/eurocon56442.2023.10199058] [ISBN: 978-1-6654-6397-3]
14. Intravaia, M.; Bindi, M.; Becchi, L.; Luchetta, A.; Lozito, G.M.; Paolucci, L.; Grasso, F.; Iturrino-García, C. "Prognostic Analysis of Switching Devices in DC–DC Converters", in *Proc. of the 2023 IEEE International Conference on Metrology for Extended Reality, Artificial Intelligence and Neural Engineering (MetroXRaine 2023)*, pp. 224–229, 2023. [DOI: 10.1109/metroxraine58569.2023.10405672]



La borsa di dottorato è stata finanziata con risorse dell'Unione europea-NextGeneration EU Piano Nazionale di Ripresa e Resilienza Missione 4 – Componente 1 – “Dalla ricerca all’impresa” – Investimento “Centri Nazionali” – CUP J53D23000760006