



UNIVERSITÀ DEGLI STUDI DI GENOVA

POLYTECHNIC SCHOOL

**Department of Electrical, Electronics, Telecommunication
Engineering and Naval Architecture**

INTELLIGENT ELECTRIC ENERGY SYSTEMS LABORATORY

XXXVIII Cycle Ph.D in Sciences and Technologies for
Electrical Engineering and Complex Systems for Mobility –
Curriculum: Electrical Engineering

**Data-Driven Energy Management for
Renewable Energy Communities: Leveraging
Smart Monitoring and Machine Learning for
Grid Flexibility**

Giovanni Baglietto

Supervisors:

Prof. Ing. Stefano Massucco

Prof. Ing. Federico Silvestro

July, 2026

Abstract

The paradigm shift necessitated by the global imperative for decarbonization and for energy security is fundamentally reshaping the architecture of electric power systems, transitioning from a centralized, dispatchable model to a decentralized ecosystem largely influenced by stochastic Renewable Energy Sources (RES). This structural transformation has caused a critical "flexibility gap," characterized by the increasing divergence between variable generation profiles and inelastic demand, which threatens grid stability at both transmission and distribution levels. This thesis addresses this challenge by proposing an integrated, data-driven framework designed to unlock Demand-Side Flexibility (DSF) within the emerging context of Renewable Energy Communities (RECs), carrying out evaluations based on the acquisition and processing of real data from different types of case studies.

Against the backdrop of the evolving European and Italian regulatory landscapes, specifically the establishment of Configurations for Self-Consumption and Sharing of Renewable Energy and of Local Flexibility Markets (MLF), the technological enablement of Distributed Energy Resources (DERs) constitutes a fundamental necessity for effective market participation. While the regulatory architecture establishes the principle of technological neutrality, the operational reality requires scalable platforms capable of bridging the gap between physical assets and market mechanisms through robust observability and control.

To overcome these barriers, the thesis presents a multi-layered methodology focused on the enabling technologies for aggregation and flexibility. First, it investigates diverse advanced energy monitoring architectures suitable for heterogeneous environments, ranging from complex infrastructures to residential units. Within this monitoring framework, specific attention is dedicated to Non-Intrusive Load Monitoring (NILM) as a promising technique to enhance data granularity where sub-metering is economically inconvenient, particularly in the residential sector. This data acquisition layer serves as the foundation for a data-driven Energy Management System (EMS), which integrates Machine Learning-based forecasting models with optimization algorithms.

The research culminates in the design and development of an intelligent platform for the optimized management of energy resources, capable of real-

time control over the dispatch of flexible assets. By maximizing shared energy incentives while simultaneously offering ancillary services to the grid, this work provides a technical pathway for making RECs active nodes within the smart grid ecosystem.

Preface

This thesis is submitted in partial fulfilment of the requirements for the degree of Philosophiae Doctor (Ph.D.) at the University of Genoa. The research work has been carried out at the Department of Electrical, Electronics, Telecommunication Engineering, and Naval Architecture (DITEN) in the IEES laboratory. The work has been supervised by Prof. Stefano Massucco and Prof. Federico Silvestro and was financially supported by a PhD scholarship co-funded by the European Union – NextGenerationEU, under the Italian National Recovery and Resilience Plan Mission 4, Component 1 “Strengthening the provision of education services: from nurseries to University”.



Ministero
dell'Università
e della Ricerca



Italiadomani
PIANO NAZIONALE
DI RIPRESA E RESILIENZA



Borsa di dottorato cofinanziata con risorse dell'Unione europea-*NextGeneration EU*
Piano Nazionale di Ripresa e Resilienza Missione 4, componente 1 “*Potenziamento
dell’offerta dei servizi di istruzione: dagli asili nido all’Università*”

Acknowledgements

First and foremost, I wish to express my deepest appreciation to Prof. Stefano Massucco and Prof. Federico Silvestro for their constant guidance and invaluable insights during this entire research journey.

I am equally thankful to my colleagues in the IEES laboratory research group for the productive discussions, successful teamwork, and the time we have spent together.

I would like to express my gratitude to IESolutions s.r.l., and specifically to Dr. Andrea Bagnasco and Dr. Giovanni Roso, for giving me the opportunity to approach the industrial sector through “technology transfer” from the research world. In particular, I thank Dr. Andrea Vinci for his teachings, advice, and heartfelt support, even beyond the duration of my internship.

I also extend my thanks to WellD Sagl, particularly to Dr. Luca Camerini, for allowing me to discover a different and surprising reality, and to all the colleagues I met during the six-month experience at WellD in Lugano.

Last but not least, I want to deeply thank my entire family and the friends who have been accompanying me on this journey for several years now.

Contents

Abstract	iii
Preface	v
Acknowledgements	vi
Contents.....	vii
Acronyms.....	xiii
List of Figures	xvii
List of Tables.....	xxi
1 Introduction	1
1.1 The European Policy Response for the Energy Transition.....	2
1.2 The Italian Transposition.....	4
1.3 The Technical Challenges of the New Electric Power System .	6
1.3.1 Transmission System Challenges	6
1.3.2 Distribution System Challenges.....	7
1.4 The Thesis Proposition: Unlocking Demand-Side Flexibility...	8

1.4.1	Pillar 1: Energy Communities as Aggregators	9
1.4.2	Pillar 2: Demand Response and Flexibility Markets	10
1.4.3	Pillar 3: Advanced Monitoring and Optimization	10
1.5	Research Objectives and Thesis Structure	12
2	The Evolution of the Italian Electricity Dispatching System: Local Flexibility and the TIDE Framework	15
2.1	The Italian Electricity Market architecture	16
2.1.1	Day-Ahead and Intra-Day markets	17
2.1.2	The Ancillary Services Market	18
2.2	The path toward the Local Flexibility Markets	19
2.2.1	The UVAM experience	19
2.2.2	The active role of the DSO and Local Flexibility Markets	20
2.3	The key innovations of the TIDE framework	24
2.3.1	Technological neutrality and new unit taxonomy	24
2.3.2	New market roles	25
2.3.3	Harmonization of Local and Global Services	26
2.4	The technological enablement	27
3	Energy Communities: Aggregate Users to Enable the Energy Transition	29
3.1	The Regulatory Framework and REC State of the Art	31
3.1.1	From EU Directives to the Italian Regulations	31
3.1.2	The Italian Scenario: Current Status, Barriers, and Drivers	37
3.2	The REC Scenario and Feasibility Analysis Tool	38
3.2.1	The Production Block	40
3.2.2	The Consumption Block	42
3.2.3	The Energy Analysis Block	42

3.2.4	The Economic Analysis Block.....	45
3.3	The Role of Advanced Monitoring and Digitalization	47
3.3.1	Operational Optimization Strategies.....	48
3.3.2	Data Availability and Advanced Monitoring.....	50
4	Energy Monitoring: case studies.....	53
4.1	Hardware and Software Tools.....	54
4.1.1	Energy Meters and Gateways	55
4.1.2	Data Acquisition, Storage, and Visualization Tools	59
4.2	The ANTES Project: Energy Impact Verification in Greenhouse Cultivation.....	62
4.3	Energy Management and Advanced Control in Healthcare Facilities	67
4.3.1	Data-driven energy analysis: the “Don Orione” case study	67
4.3.2	Towards Optimized Control: the Humanitas “Casa Suore” case study.....	71
4.4	Energy Management and Optimization in a Zero Energy Building.....	75
4.4.1	Monitoring and Control Architecture	77
4.4.2	Optimization scenarios	78
4.4.3	Formalization of the optimization problem	81
5	Residential Energy Monitoring: Optimisation Model for Demand Flexibility Based on Non-Intrusive Load Disaggregation.....	85
5.1	Case-Study	87
5.2	The NILM Algorithm.....	90
5.2.1	Events Detection.....	92
5.2.2	Patterns Merging.....	94
5.2.3	Patterns Filtering.....	95
5.2.4	Patterns Clustering.....	96

5.3	The REC Optimization Model.....	96
5.4	Assessment of the NILM Algorithm	100
5.5	Assessment Of The Optimization Model.....	104
5.6	Conclusion.....	107
6	An Intelligent Platform for Monitoring and Optimization of RECs.....	109
6.1	The SPARKle Project: Platform Scope and Architecture.....	111
6.2	Forecasting Algorithms	114
6.2.1	PV generation forecasting.....	115
6.2.2	Load forecasting.....	122
6.3	Optimization strategies.....	127
6.3.1	REC scenario	128
6.3.2	Flexibility Scenario.....	134
6.4	Algorithmic Integration within the Platform Architecture	136
6.5	Conclusion.....	140
7	Broadening the Horizon: European Projects on Energy Communities	143
7.1	LIGHTNESS	144
7.1.1	Digitalization in the Engagement Process.....	145
7.1.2	Upskilling citizens: lessons learnt.....	147
7.1.3	Project findings	148
7.2	HESTIA.....	149
7.2.1	Baseline Formulation Framework and the IPMVP Methodology.....	150
7.2.2	Pilot Validation Outcomes	154
7.2.3	Project findings	157
7.3	NEON.....	158
7.3.1	Energy Performance Contracting Schemes.....	159

7.3.2	Financial Barriers and Opportunities for CECs	161
7.4	Conclusions	163
8	Conclusions	165
	Bibliography.....	169
	List of Publications	185
	List of Projects	187
	List of attended Courses and Conferences	191
	International and Industrial Experiences.....	193

Acronyms

AC	Air Conditioning
AI	Artificial Intelligence
AMI	Advanced Metering Infrastructure
ANN	Artificial Neural Networks
API	Application Programming Interface
BAS	Building Automation Systems
BESS	Battery Energy Storage System
BEM	Basic Ensemble Method
BMS	Building Management System
BRP	Balance Responsible Party
BSP	Balancing Service Provider
CART	Classification And Regression Tree
CC	Cloud Cover index
CCSM	Corrected Clear-Sky Model

CEC	Citizen Energy Community
CNN	Convolutional Neural Networks
CSC	collective self-consumption
CSM	Clear-Sky Model
CT	Current Transformer
DER	Distributed Energy Resources
DR	Demand Response
DSF	Demand-Side Flexibility
DSM	Demand Side Management
DSO	Distribution System Operators
EC	Energy Community
EES	Energy Efficiency Services
EMS	Energy Management System
EPC	Energy Performance Contracting
ESCo	Energy Services Companies
EV	Electric Vehicle
GHG	Greenhouse Gas
GME	Gestore dei Mercati Energetici (Energy Markets Operator)
GSE	Gestore dei Servizi Energetici (Energy Services Manager)
GUI	Graphical User Interface
HV	High Voltage
HVAC	Heating, Ventilation and Air Conditioning
IBR	Inverter-Based Resources
IEMD	Internal Electricity Market Directive
IoT	Internet of Things
IPMVP	International Performance Measurement and Verification Protocol

KPI Key Performance Indicator

LSTM Long Short-Term Memory

LV Low Voltage

MB Mercato di Bilanciamento (Balancing Market)

MGP Mercato del Giorno Prima (Day-Ahead Market)

MI Mercato Infragiornaliero (Intra-Day Market).

MILP Mixed-Integer Linear Programming

ML Machine Learning

MLF Mercati Locali della Flessibilità (Local Flexibility Markets)

MPC Model Predictive Control

MSD Mercato per il Servizio di Dispacciamento (Italian Ancillary Services Market)

MV Medium Voltage

M&V Measurement and Verification

NILM Non-Intrusive Load Monitoring

PLC Power-line Communication

POD Point Of Delivery

P2P Peer-to-Peer

P4P Pay-for-Performance

PUN Prezzo Unico Nazionale (Single National Price)

PV Photovoltaic

REC Renewable Energy Community

RED Renewable Energy Directive

RES Renewable Energy Sources

SME Small and Medium-sized Enterprise

SP Service Provider

TIDE Testo Integrato del Dispacciamento Elettrico (Integrated Text on Electricity Dispatching)

TSO Transmission System Operator

UAS Unità Abilitate Singolarmente (Individually Qualified Units)

UPR Unità di Produzione Rilevanti (Relevant Production Units)

UVAM Unità Virtuali Abilitate Miste (Virtually Aggregated Mixed Units)

UVAN Unità Virtuali Abilitate Nodali (Virtually Aggregated Nodal Units)

UVAZ Unità Virtuali Abilitate Zonali (Virtually Aggregated Zonal Units)

VPP Virtual Power Plant

ZEB Zero Energy Building

List of Figures

Figure 2.1 – Italy’s market zones and connections [adapted from Terna]	17
Figure 3.1 – Flow diagram of the REC Scenario and Feasibility Analysis Tool	39
Figure 3.2 – Typical PV produced power daily profiles for each month of the year	41
Figure 3.3 – Monthly average daily load profiles for the prosumer	43
Figure 3.4 – Monthly average daily load profiles for the consumer	43
Figure 3.5 – Monthly energy production KPIs	44
Figure 3.6 – Monthly energy consumption KPIs	45
Figure 3.7 – Discounted cash flow and NPV over 20 years horizon	46
Figure 4.1 – Shelly 3EM wiring diagram	56
Figure 4.2 – Socomec DIRIS A-40	56
Figure 4.3 – Socomec DIRIS Digiware multipoint power monitoring system	57
Figure 4.4 – InHand IoT Edge Gateway IG502	58

Figure 4.5 – Intesis 700 Series Air Gateway for Panasonic to Modbus systems.....	59
Figure 4.6 – Geothermal heat pump and geothermal probes in the greenhouse	63
Figure 4.7 – Heating mats on the cultivation benches	63
Figure 4.8 – LED lighting in the experimental greenhouse	64
Figure 4.9 – Details of automated louvers and roof windows of the experimental greenhouse	64
Figure 4.10 – Scheme of the ANTES monitoring system architecture ...	65
Figure 4.11 – Screenshot of the dashboard with general energy consumption data of the experimental greenhouse	65
Figure 4.12 – Screenshot of the homepage dashboard of the Paverano facility	68
Figure 4.13 – Day/night electrical consumption and daily average temperature	69
Figure 4.14 – Correlation between daily electricity consumption of the heating system and average daily outdoor temperature	69
Figure 4.15 – ABC curve of Paverano facility.....	70
Figure 4.16 – Heatmap of thermal power plant 15-minutes electricity consumption.....	71
Figure 4.17 – Scheme of the “Casa Suore” monitoring system architecture	72
Figure 4.18 – Rule-based policy flowchart for HVAC control in “Casa Suore”	74
Figure 4.19 – Plan of the laboratory and external view of the laboratory	76
Figure 4.20 – Monitoring and control system architecture of the ZEB...	77
Figure 4.21 – Schematic representation of the Model 1 ZEB architecture	79
Figure 4.22 – Schematic representation of the Model 4 ZEB architecture in the REC scenario	80

Figure 5.1 – System architecture	89
Figure 5.2 – REC configuration to test the optimization model	90
Figure 5.3 – Flow diagram of the NILM algorithm	91
Figure 5.4 – Flow diagram of the Events Detection function	93
Figure 5.5 – Clustering metrics trend with the number of clusters for House 1	100
Figure 5.6 – Output of K-means algorithm with the optimal number of clusters for House 1: the silhouette coefficient for each single load pattern (the red dashed line represents the average) on the left and the cluster centroids on the right.....	101
Figure 5.7 – Daily load profile of 20th November 2022 of House 1	102
Figure 5.8 – Daily load profile of 21th November 2022 of House 2	102
Figure 5.9 – Daily load profile of 5th November 2022 of House 3	103
Figure 5.10 – Comparison between Eloadh of House 1, both the actual one and the optimized one, and Einjh of 9th and 10th May 2024	106
Figure 5.11 – Optimized load profile of House 1 of 10th May 2024, with the profiles of the suggested appliances, identified by the cluster number.	106
Figure 6.1 – Flow diagram of the optimisation in the REC scenario.....	113
Figure 6.2 – Flow diagram of the optimisation in the flexibility scenario	113
Figure 6.3 – General scheme of the implemented PV generation forecasting technique [54]	117
Figure 6.4 – Time horizons for forecasting techniques [99].....	123
Figure 6.5 – Logical scheme of the autoregressive CNN-LSTM architecture	125
Figure 6.6 – Hourly energy and its daily cumulative value over 24 hours	138
Figure 6.7 – Energy data chart over time showing daily energy over a month.....	138

Figure 6.8 – Forecast chart of the aggregated production and consumption for the entire REC	138
Figure 6.9 – Simulation output of the activation of a flexibility service.	139
Figure 7.1 – Engagement screen with gamification widgets of the LIGHTNESS platform.....	146
Figure 7.2 – Test comparison between the winter period baseline and the actual measurement of hourly energy imported from the grid, for a Dutch pilot household.....	153
Figure 7.3 – Test comparison between the summer period baseline and the actual measurement of hourly energy imported from the grid, for a Italian pilot household	153
Figure 7.4 – Daily average load profiles considering all days (blue), only weekdays (green) and only weekends and holidays (orange) of a French pilot household.....	154

List of Tables

Table 1.1 – Key Italian 2030 Energy Transition Targets (PNIEC 2024).	4
Table 2.1 – Pilot projects’ dimensions [36]	23
Table 3.1 – REC and CEC definitions	32
Table 3.2 – Parameters for the TIP calculation in €/MWh	36
Table 3.3 – TIP geographical correction factor	36
Table 3.4 – Input parameter for the Clear-Sky Model	41
Table 3.5 – Economic parameters used in the economic analysis example	47
Table 4.1 – Layers and components of the proposed software architecture	62
Table 5.1 – List of Case Studies	88
Table 5.2 – Optimized Cluster Centroids	101
Table 5.3 – Algorithm Performance for the Case Study	104
Table 6.1 – Configuration of the layers of the proposed CCN-LSTM model	127

Table 7.1 – Correlation between user platform logins and energy performance in the Dutch Pilot 155

Table 7.2 – Correlation between user platform logins and energy performance in the Italian Pilot..... 156

CHAPTER 1

Introduction

The 21st-century global energy system is navigating a profound and irreversible structural transformation, propelled by the convergence of two co-equal imperatives.

The first is the long-standing and existential mandate for environmental sustainability, driven by the unequivocal scientific consensus on anthropogenic climate change. The Intergovernmental Panel on Climate Change (IPCC) in its Sixth Assessment Report (AR6) presents definitive evidence that the combustion of fossil fuels, the bedrock of the modern energy sector, is the primary driver of the atmospheric greenhouse gas (GHG) concentrations responsible for accelerating global warming and climate disruption [1]. This scientific consensus was translated into a global political mandate by the 2015 Paris Agreement, which commits signatories to limit global warming to well below 2 °C, and preferably to 1.5 °C, above pre-industrial levels, necessitating a rapid and deep decarbonization of the entire energy system [2].

The second, more recent and acute driver is the imperative for economic and geopolitical energy security. The long-term climate agenda was violently accelerated by the global energy crisis that began in 2021 and peaked in 2022. This period of extreme price volatility and supply disruption, exacerbated by geopolitical conflicts, exposed the profound economic and

strategic vulnerabilities of regions heavily dependent on imported fossil fuels [3]. For nations within the European Union (EU), including Italy, this crisis served as a structural break, demonstrating in stark economic terms that energy dependence is a critical threat to national security and industrial competitiveness.

Consequently, the energy transition has been reframed. It is no longer perceived merely as a long-term environmental objective but as an urgent, short-term strategy for ensuring energy sovereignty, economic stability, and geopolitical autonomy. This convergence of the climate and security imperatives has created the political will for an unprecedented acceleration of the transition toward clean energy sources.

1.1 The European Policy Response for the Energy Transition

In response to this dual imperative, the European Union has established one of the world's most ambitious and legally binding policy frameworks to orchestrate a systemic decarbonization. This framework has evolved through several key legislative phases, reflecting a progressive deepening of commitment and a reaction to geopolitical events.

The initial comprehensive legislative effort to build this framework was the Clean Energy for all Europeans Package (CEP), adopted between 2018 and 2019. This extensive legislative package established a complete legal foundation for the energy transition, focusing on critical areas such as energy efficiency, renewable energy integration, strategic electricity market design, and mechanisms to guarantee security of supply. Crucially, the CEP comprised key legal acts that introduced essential concepts for new actors and decentralized energy structures. Chief among these were the 2018 Renewable Energy Directive (RED II) [4] and the 2019 Internal Electricity Market Directive (IEMD) [5]. These two directives provided the formal, EU-level definitions for Renewable Energy Communities (RECs) and Citizen Energy Communities (CECs), a pivotal step that empowered citizens to evolve from passive consumers into active participants, or "prosumers", in the energy transition.

Building on this foundation, the European Union developed the current framework, which is constructed around three core legislative pillars:

1. The European Green Deal [6]: the overarching strategy committing the EU to becoming the first climate-neutral continent by 2050,

fundamentally reframing economic growth as a function of resource efficiency and sustainability.

2. The ‘Fit for 55’ Package [7]: the primary legislative engine designed to align EU law with the 2030 target of reducing net GHG emissions by at least 55% compared to 1990 levels. This package initiated the comprehensive revision of the EU's entire energy and climate rulebook.
3. The REPowerEU Plan [8]: launched in 2022 as a direct response to the energy security crisis, this plan aims to rapidly phase out dependence on Russian fossil fuels. Its primary effect was to accelerate the ‘Fit for 55’ timeline, mandating faster deployment of renewable energy and more aggressive energy efficiency measures as a matter of urgent security policy.

This policy acceleration has been codified into new, binding directives that fundamentally alter the technical and economic parameters for the energy sector. The most critical of these, resulting from the ‘Fit for 55’ and REPowerEU impetus, are the 2023 recasts of the Renewable Energy Directive (RED III) [9] and the Energy Efficiency Directive (EED) [10], and the 2024 reform of the Electricity Market Design (EMD) [11].

The revised Renewable Energy Directive, which entered into force in November 2023, raises the binding EU-wide renewable energy target for 2030 from 32% to a minimum of 42.5%, with a collective aspiration to reach 45% [12]. This directive legally mandates the massive integration of variable Renewable Energy Sources (RES) that forms the core of the technical problem this dissertation addresses.

Concurrently, the recast Energy Efficiency Directive which entered into force in October 2023, establishes a new, binding EU-level target to reduce final energy consumption by 11.7% by 2030 compared to 2020 projections [13]. Critically, this directive legally enshrines the "energy efficiency first" principle, elevating demand-side solutions from a peripheral option to a central, legally-mandated component of energy policy and system planning.

Furthermore, the EU's 2024 reform of the Electricity Market Design, a direct response to the 2022 price crisis, aims to decouple short-term electricity prices from the volatility of fossil fuels. Its provisions explicitly strengthen the market framework for non-fossil flexibility, promoting long-term contracts and creating new rules to encourage participation in Demand Response (DR) and the adoption of energy storage.

This comprehensive set of EU laws sets the goals for the energy transition. Member States must transpose these mandates into national law, primarily via the National Energy and Climate Plans (NECPs). Required by the Governance Regulation (EU)2018/1999 [14], these NECPs are the key strategic tool to deliver a climate-neutral Europe by outlining precise actions, achieving 2030 targets, and steering the necessary public and private investments. Originally drafted in 2020 and recently updated, the NECPs incorporate the ambitious targets from the European Green Deal, Fit for 55 Package, and REPowerEU plan.

1.2 The Italian Transposition

In Italy, the EU's top-level mandates are translated into national strategy through the *Piano Nazionale Integrato per l'Energia e il Clima* (PNIEC). The final updated PNIEC [15], submitted in 2024, internalizes the heightened ambitions of the "Fit for 55" and REPowerEU packages, setting quantitative targets for 2030 that represent a radical acceleration of the country's energy transition.

The scale of this challenge is captured in the plan's core targets, summarized in Table 1.1. The plan mandates an increase in solar and wind capacity from approximately 50 GW in 2024 to 107 GW by 2030, aiming for renewables to cover 63% of all electricity consumption. To manage the intermittency of this massive RES fleet, the PNIEC also sets a crucial new target for 72 GWh of energy storage capacity.

Table 1.1 – Key Italian 2030 Energy Transition Targets (PNIEC 2024)

Metric	2024 (Baseline)	2030 (Target)
Total Solar & Wind Capacity	~50 GW	107 GW
% RES in Final Electricity Consumption	~41%	63%
Energy Storage Capacity	~7 GWh	72 GWh
Projected Annual Overgeneration (Curtailment)	~0.3 TWh	~5 TWh
Total RES Production (Electrical)	129 TWh	227 TWh

Notably, scenarios from Terna, the Italian Transmission System Operator (TSO), supporting the PNIEC, forecast that, without a fundamental shift in system operation, this RES fleet will result in approximately 5 TWh of annual ‘overgeneration’ by 2030 [16], as shown in Table 1.1. This amount represents a quantification of the energy that should be curtailed, and wasted, due to the current system inability to absorb it, underscoring the magnitude of the emerging ‘flexibility gap’.

To bridge this gap and enable the necessary investments, the Italian system is deploying a new set of market and regulatory tools designed to de-risk the transition. As highlighted by Terna's strategic analysis, the high capital intensity (CAPEX) of renewables and storage, combined with the volatility of spot markets, makes purely merchant investments financially unsustainable. Therefore, Italy is implementing mechanisms to separate commercial risk (price/volume) from industrial risk:

1. MACSE (Mechanism for the Procurement of Electricity Storage Capacity): to address the flexibility gap, Italy targets 72 GWh of storage capacity by 2030, including 50 GWh of new utility-scale storage. The MACSE mechanism, as required by Article 18 of the Italian Legislative Decree 210/2021 [17], auctions long-term contracts where operators receive a fixed premium to build and maintain storage capacity, while the TSO manages the utilization of this capacity in the ancillary services market. This ensures the physical adequacy of the system and facilitates the integration of non-programmable renewables.
2. FER X (Decree for Renewable Incentives) [18]: to reach the 107 GW target, this mechanism supports new renewable capacity, stabilizing revenues for investors and lowering the cost of capital.
3. Capacity Market: essential for ensuring system adequacy during the phase-out of thermal generation, this mechanism remunerates the availability of programmable capacity to cover peak demand. It could serve as an alternative market tool for storage systems unable to participate in the MACSE, in order to meet the 72 GWh target.

This legal-technical cascade creates the central research challenge: these "enabling factors" provide the legal structure to build new distributed renewable capacity and storage for the transition, but it does not provide the technical tools to integrate these new assets into the grid and operate them optimally, both from an efficiency and economic point of view.

1.3 The Technical Challenges of the New Electric Power System

The rapid, policy-driven transition from a power system dominated by centralized, dispatchable, synchronous generators to one reliant on decentralized, non-dispatchable, and inverter-based resources (IBRs) creates a critical "flexibility gap" [19]. This gap, the system's inability to match variable supply with fluctuating demand across all timescales and geographical locations, manifests as two distinct but interconnected sets of problems at the transmission and distribution levels, discussed in the next paragraphs.

1.3.1 Transmission System Challenges

At the bulk system level, the displacement of large, rotating thermal generators by non-dispatchable resources, like solar photovoltaic (PV) and wind, fundamentally degrades the grid's physical stability and controllability [20], leading to the following main challenges:

- **Loss of Rotational Inertia:** system inertia, the stored kinetic energy in large rotating generators, is the grid's primary defence against sudden frequency deviations. IBRs are decoupled from the grid via power electronics and inherently provide zero physical inertia. As the penetration of IBRs reaches high levels, the system's inertia progressively declines. In a low-inertia system, any supply-demand imbalance (e.g., a generator trip or a sudden cloud cover event) causes a much faster Rate of Change of Frequency (ROCOF) [21]. This high ROCOF can trigger protective relays and lead to cascading failures before slower-acting primary frequency reserves can be deployed, as demonstrated in the 2016 South Australia blackout [22].
- **Balancing and Ramping Deficits:** the generation profile of solar PV creates the "duck curve" phenomenon. This profile introduces two critical challenges: (1) an oversupply of energy at midday, which pushes net demand to record lows (the "minimum demand" problem), creates negative market prices, and threatens to force essential, inertia-providing synchronous generators offline [23]; (2) an extremely steep "ramping" requirement in the evening as solar generation simultaneously collapses and system load peaks [24].

- Decreasing of System Dispatchability and Real-Time Controllability: the massive penetration of passive distributed resources fundamentally compromises the system operator's real-time visibility and active power control mechanisms [23]. Because this aggregated generation operates outside centralised control, it renders existing dispatch optimization engines increasingly inadequate.

1.3.2 Distribution System Challenges

The emergence of the "prosumer" fundamentally inverts the traditional operating model of the distribution network. Historically, this network was designed as a passive, "fit-and-forget" system for unidirectional power flow from the high-voltage grid down to end consumers. The proliferation of Distributed Energy Resources (DERs), primarily rooftop PV systems, destabilizes this structure because it was not engineered for local generation. When local DER generation exceeds local consumption, power flows "backwards" or "up" from the Low Voltage (LV) grid to the Medium Voltage (MV) grid. This reverse power flow is the primary cause of sustained voltage rise (overvoltage) at the grid edge. Such overvoltage violates statutory limits (e.g., EN 50160), potentially damaging customer equipment and forcing PV inverters to automatically disconnect. The mass-tripping of DERs during peak production presents as a sudden loss of generation to the TSO, thereby directly coupling a localized voltage problem in the distribution system to the TSO's system-wide frequency stability crisis. [23]

The distribution grid is thus transformed into an active arena of bidirectional flows and rapid fluctuations that traditional planning cannot accommodate without prohibitively expensive reinforcement. This complex, two-way power exchange introduces several interconnected challenges for Distribution System Operators (DSOs), including the need for extensive grid stability and protection adaptation, as new power flow patterns significantly affect local voltage and fault current levels, necessitating the costly redesign of protection systems. Furthermore, high penetration of intermittent DERs can lead to network congestion and losses, as well as increased costs of imbalances if not operated correctly. The lack of real-time monitoring in the LV network, unlike in transmission grids, also creates critical data and observability gaps, which have to be solved to effectively manage voltage-quality issues and aggregated DERs. Finally, the use of power electronic converters to interface some DER technologies can compromise power quality by generating harmonics. [25]

Consequently, the concept of *flexibility*, defined by International Energy Agency (IEA) as “the ability of a power system to reliably and cost-effectively manage the variability and uncertainty of supply and demand across all relevant timescales” [26], has become a critical operational necessity for DSOs. Indeed, DERs, when properly controlled, can enhance grid stability and reliability by providing network support and ancillary services, and intentional islanding can protect clusters of customers against power outages. Crucially, Demand-Side Flexibility (DSF), the ability of users to modify their load profile (e.g., via smart electric vehicle (EV) charging or managing battery energy storage systems (BESS)), is a critical resource. The coordinated operation of PVs, EVs, and BESS enhances grid stability, power quality, and network flexibility, mitigating the technical issues that arise when these devices are managed individually.

DER integration fundamentally changes the DSO's role into an active system manager, who, as a neutral market facilitator without proprietary energy assets, has to manage challenges like overvoltage and protection by leveraging the newfound opportunities for coordination and flexibility across PV, EV, and BESS assets owned and operated by prosumers and third-party aggregators, as further detailed in subsection 2.2.2 of Chapter 2.

1.4 The Thesis Proposition: Unlocking Demand-Side Flexibility

This dual TSO/DSO flexibility gap cannot be solved efficiently, economically, or rapidly enough by supply-side solutions alone. Traditional network reinforcement (i.e., replacing conductors and transformers) is economically prohibitive and faces intractable permitting delays. While utility-scale storage is a necessary component (as recognized by the PNIEC's 72 GWh target), it is capital-intensive and primarily addresses the TSO's temporal balancing problem, not the DSO's localized voltage and congestion issues.

The critical, largely untapped resource required to balance this new energy paradigm is the inherent flexibility of the end-users themselves: Demand-Side Flexibility. The 20th-century operational paradigm of "supply follows load" is no longer viable in a system dominated by non-dispatchable renewables. The 21st-century grid must operate on a new principle: "load follows supply". This requires transforming end-users from passive,

stochastic "consumers" into active, responsive "prosumers" whose consumption patterns can be modulated.

The DSF is a uniquely powerful resource for DSOs. By modulating consumption (e.g., intelligently scheduling EV charging, heat pump cycles, or water heaters), an aggregated block of prosumers can be dispatched to solve the DSO's local challenges. The DSF can absorb local overgeneration from rooftop PV, increasing self-consumption to mitigate reverse power flows and correct local overvoltage issues. In addition to this primary function, the aggregated block can also provide fast-acting reserves and balancing services, which contribute to system-wide stability by participating in ancillary service markets. [26]

Activating this demand-side resource is not a simple technical problem; it requires the creation of an integrated socio-technical-economic framework. A single household's flexibility is negligible. It only becomes a viable grid resource when aggregated and controlled collectively, and it is presumable that would only participate if economically y incentivized to do so.

The central proposition of this dissertation is that activating DSF requires the seamless integration of three interdependent pillars. This thesis will design, develop, and validate the technological core (Pillar 3) that makes the entire socio-economic structure (Pillars 1 and 2) operationally viable.

1.4.1 Pillar 1: Energy Communities as Aggregators

The legal frameworks for Renewable Energy Communities and Citizen Energy Communities, as defined in Italy by Legislative Decree 199/2021 [28] and technically regulated by the Integrated Text for widespread self-consumption (TIAD), introduced by the Italian Regulatory Authority for Energy, Networks and Environment (ARERA) Resolution 727/2022/R/eel [29], provide the novel socio-technical structure for the aggregation of single users. The REC acts as the legal and operational entity that bundles a set of individual prosumers, producers or consumers into a single, virtual entity. This aggregation achieves the critical mass necessary to interact meaningfully with the grid operator and participate in energy markets [30]. In Italy, under the current regulatory framework, REC participation is subject to specific constraints. Notably, individual renewable energy plants are limited to a maximum installed capacity of 1 MW, and the incentive measure for virtual energy sharing is geographically bounded by the perimeter of the underlying primary substation. A comprehensive analysis of these regulatory details and limitations is provided in Chapter 3.

1.4.2 Pillar 2: Demand Response and Flexibility Markets

Aggregation alone is insufficient; participation must be monetized. This is achieved through new market mechanisms that remunerate flexibility. This thesis will focus on two primary economic drivers in the Italian context:

1. **Shared Energy Incentives:** the primary incentive defined in [28] and TIAD, which provides a tariff for energy virtually ‘shared’ within the REC’s primary substation perimeter. The internal sharing mechanism is entirely virtual and relies on the DSO’s smart metering infrastructure, without the need for third-party measuring equipment.
2. **Local Flexibility Markets (MLF, Mercati Locali della Flessibilità):** historically, flexibility was the exclusive domain of the Ancillary Services Market (MSD, Mercato per il Servizio di Dispacciamento), managed by Terna, the national TSO. The MSD’s primary function is to balance supply and demand in real-time across the national transmission grid, while also managing inter-zonal congestion. Participation in this process was limited exclusively to large, programmable production units (UPP) capable of providing substantial volumes of regulating power. While reforms such as the UVAM (Virtually Aggregated Mixed Units) pilot project opened the MSD to distributed resources, the MSD remains fundamentally TSO-centric. Its signals are derived from national or zonal balancing needs, which do not necessarily correlate with, and may even contradict, the needs of the local distribution grid. So, recognizing the specific needs of DSOs, ARERA enacted Resolution 352/2021/R/eel [31], which established the framework for “pilot projects for the procurement of local ancillary services”, effectively birthing Local Flexibility Markets in Italy. Unlike the TSO’s frequency reserves, these local markets exclusively target DSO-specific products, primarily local congestion management and local voltage control.

1.4.3 Pillar 3: Advanced Monitoring and Optimization

This is the technical core of the thesis and the primary research gap. To optimize against the shared energy incentive and bid into the Local Flexibility Market, the REC’s aggregator requires a sophisticated Energy

Management System (EMS) [32]. A functional EMS is fundamentally reliant on two critical data streams:

- Monitoring: real-time, granular data on aggregate or, better, appliance-level consumption to identify the available, shiftable loads (e.g., water heaters, EV chargers, HVAC (heating, ventilation and air conditioning) systems) and on energy controllable assets (e.g. PV plants, storage systems) that constitute the REC's flexibility potential.
- Forecasting: accurate, machine-learning-based short-term forecasts of both energy generation and aggregate load consumption [33].

The primary and most critical barrier is the availability of data, and consequently the cost and complexity of monitoring. Deploying intrusive, per-appliance sub-meters can be a viable solution for companies or large consumers, but it is economically non-viable, logistically complex, and unacceptable to residential consumers.

This dissertation first proposes and validates different lightweight monitoring solutions for industrial settings, demonstrating how to analyse the readily available data to significantly improve energy efficiency. Building upon this foundation, the work developed for this thesis subsequently proposes Non-Intrusive Load Monitoring (NILM) as a potential enabling technology to bridge this gap in the residential sector. NILM, also known as energy disaggregation, is a family of data-analytic techniques that apply machine learning (ML) algorithms to the single, aggregated data stream from a building's main smart meter to infer the status and consumption of individual appliances. This approach provides a low-cost, scalable, and non-invasive method to acquire the granular, appliance-level data necessary for the EMS. These data allow the EMS to build a bottom-up model of its available flexibility and execute optimal Demand Side Management (DSM) strategies to maximize the REC's economic objectives.

A second critical point concerns how these data streams are effectively leveraged to offer operational solutions for energy efficiency and management. To overcome this limitation, this thesis proposes the use of advanced Machine Learning and Artificial Intelligence (AI) tools for short-term forecasting of both generation from renewable sources and aggregate consumption. These forecasts will serve as essential input for optimization algorithms, based on Mixed-Integer Linear Programming (MILP), that control the aggregator's flexible and programmable assets (e.g., local storage,

electric vehicle charging) in real-time to maximize the energy and economic benefits of the Energy Community (increase in shared energy and revenues from flexibility markets).

1.5 Research Objectives and Thesis Structure

The principal objective of this doctoral research is to design, develop, and validate an integrated, data-driven framework for the optimal energy management of Renewable Energy Communities within the new Italian regulatory context. This framework will leverage advanced energy monitoring and data analytics, specifically unsupervised NILM, machine learning-based forecasting and optimization algorithms, to unlock demand-side flexibility, enabling RECs to maximize their shared energy incentives and participate effectively in ancillary service markets.

Specifically, the thesis aims to:

1. Analyze the Electricity Market: critically review the Italian electricity market, with a specific focus on the new MLF and the related pilot projects, evaluating barriers and opportunities for distributed resources.
2. Analyze the REC regulatory framework: examine the transposition of EU directives into Italian law regarding Energy Communities.
3. Present design and management tools: introduce and evaluate software tools and methodologies available for the dimensioning, economic simulation, and management of energy communities.
4. Explore enabling technologies: investigate the role of energy monitoring and machine learning as key enablers. Explore NILM as a non-invasive technique to acquire granular consumption data and Machine Learning for forecasting generation and load.
5. Develop and validate an optimization framework: design and implement an integrated optimization framework for RECs. This framework will leverage data and ML-based forecasts to optimize the community's energy flows, maximizing shared energy incentives and providing flexibility services to enhance grid stability.

The thesis is organized into the following chapters to systematically address these aims:

- Chapter 2 analyses the structure of the Italian electricity market, detailing the MSD and the evolution of flexibility mechanisms, including the local ancillary services projects and the future opening of the market to demand response.
- Chapter 3 provides a comparative analysis of the regulatory frameworks for RECs and CECs in Europe and Italy. It details the Italian transposition and the technical rules for widespread self-consumption. It also presents the necessary tools to optimally design and manage RECs, from the feasibility study to the monitoring.
- Chapter 4 presents case studies where an architecture for energy monitoring was implemented with different scopes and objectives.
- Chapter 5 proposes an innovative monitoring system, based on the Chain2 feature of the 2G Italian DSO's smart meter and coupled with a new NILM methodology, to be applied within energy management systems of residential RECs.
- Chapter 6 details the forecasting and optimization algorithms developed for the EMS of a REC and dives into the integration of this optimization framework into a comprehensive platform, developed within a funded research project, for energy monitoring and the management of real-world Energy Communities.
- Chapter 7 briefly presents three EU funded projects about the development of Energy Communities in Europe, shifting the focus from the national level to the International and EU Community framework.
- Chapter 8 summarizes the main findings and highlights the original contributions of the thesis. It also critically discusses the limitations of the current work and suggests directions for future research.

CHAPTER 2

The Evolution of the Italian Electricity Dispatching System: Local Flexibility and the TIDE Framework

The European energy transition is not merely a technological substitution of fossil fuel generation with renewable sources; it represents a fundamental structural paradigm shift in the physics and economics of power systems. As established in the preceding chapter, the massive integration of Renewable Energy Sources, specifically photovoltaic and wind, into the distribution grid has caused a critical "flexibility gap." This gap is defined as the widening divergence between the stochastic, weather-dependent nature of modern generation and the rigid, inelastic profile of traditional demand. In physical terms, this manifests as grid congestion, voltage violations, and a significant decrease in system inertia, challenging the traditional "load-following" dispatch model. The aim of Chapter 2 is to address the parallel evolution required to bridge this gap: the transformation of the electricity market architecture from a centralized, supply-following model to a decentralized, flexibility-centric ecosystem.

The necessity of market reform is driven by the realization that physical assets alone cannot stabilize the grid; economic signals are required to unlock the latent flexibility of distributed resources. This evolution is mandated by the European Union's Clean Energy for All Europeans Package, specifically

Directive (EU) 2019/944 on the common rules for the internal market for electricity [5]. Article 32 of this Directive explicitly revolutionizes the role of DSOs, mandating a transition from passive network management ("Fit and Forget") to active system operation. It requires DSOs to procure flexibility services to manage local constraints efficiently, rather than relying solely on grid reinforcement.

In the Italian context, this directive interacts with the ambitious targets of the PNIEC, which envisions a renewable capacity exceeding 130 GW by 2030, with a substantial portion connected to medium and low voltage grids [15]. To put into operation this vision, the Italian regulator has orchestrated a phased regulatory reform. This began with a "pilot projects" phase, which opened the ancillary services market to non-traditional resources, [34] and is culminating in the organic structural reform known as Integrated Text on Electricity Dispatching (TIDE, Testo Integrato del Dispacciamento Elettrico), established by Resolution 345/2023/R/eel [35].

This chapter provides a review of this transition, starting from examining the traditional Italian market architecture to identify the structural rigidities that exacerbate the flexibility gap. It then explores the emerging landscape of Local Ancillary Services, analysing the results of pioneer projects such as UVAM, and more recent pilot projects for local flexibility markets like RomeFlex, EDGE and Mindflex. Finally, it provides a critical analysis of the TIDE framework, evaluating how the innovations that it has introduced create the necessary market signals to unlock demand-side flexibility.

2.1 The Italian Electricity Market architecture

To comprehend the magnitude of the required flexibility reform, the existing market structure must first be thoroughly analysed.

The Italian wholesale electricity market is structured as a sequence of trading venues managed by the Energy Markets Operator (GME, Gestore dei Mercati Energetici) and the TSO, Terna. This architecture was originally designed for an unidirectional system dominated by programmable thermal generation, where the primary objective was to dispatch large power plants to follow a relatively predictable load curve. But the efficiency of this "As-Is" model is being challenged by the increasing penetration of DERs, as described in Chapter 1.

2.1.1 Day-Ahead and Intra-Day markets

The commercial negotiation of energy occurs primarily in the Day-Ahead Market (MGP, Mercato del Giorno Prima) and the Intra-Day Market (MI, Mercato Infragiornaliero). These markets facilitate the scheduling of energy production and consumption but, historically, have not been designed to value flexibility as a distinct service.

The Day-Ahead Market

The MGP acts as the primary arena for energy trading, closing at 12:00 on the day prior to delivery (D-1). It operates as an implicit auction where supply and demand curves are matched to determine the system marginal price for each Market Time Unit (MTU). While the Italian MGP has historically operated on an hourly resolution, the ongoing market evolution, driven by European integration and implemented by the TIDE reform, mandates a transition towards a 15-minute MTU.

A defining characteristic of the Italian MGP is its zonal structure, designed to reflect the physical transmission constraints of the National Transmission Grid. Italy is segmented into geographical market zones (e.g., North, Centre-North, Centre-South, South, Calabria, Sicily and Sardinia), as shown in Figure 2.1, and virtual zones for foreign interconnections.




Figure 2.1 – Italy's market zones and connections [adapted from Terna]

Crucially, while producers are paid the zonal clearing price, reflecting local congestion and generation costs, consumers pay a unified national price

(PUN, Prezzo Unico Nazionale), which is the weighted average of zonal prices. This mechanism smooths geographical price disparities for consumers but masks the true locational value of flexibility. For example, a consumer in a constrained zone like Sicily pays the same PUN as a consumer in Lombardy, effectively removing the market incentive for demand response to alleviate local grid stress. This decoupling of price from physical location for loads is a significant distortion that the TIDE reform aims to address by exposing consumers to more granular price signals.

The Intra-Day Market

Following the MGP, the MI allows operators to adjust their positions closer to real-time, correcting forecast errors, a function that has become critical with the rise of RES. Historically, the MI operated through discrete auction sessions (MI1 to MI7). However, aligning with the European XBID (Cross-Border Intraday) project, Italy has transitioned to the Single Intraday Coupling (SIDC) model, which replaces discrete auctions with a hybrid model consisting of continuous trading interspersed with three auctions, allowing positions to be adjusted up to one hour before physical delivery (H-1). This continuous trading capability is a fundamental prerequisite for integrating RES, as it reduces the volume of imbalances that must be managed by the TSO in real-time.

2.1.2 The Ancillary Services Market

While MGP and MI are commercial markets for energy commodities, MSD is the technical venue where Terna procures the flexibility resources necessary to ensure grid security and physical balance. Unlike MGP and MI, where GME acts as the counterparty, in the MSD, Terna is the central counterparty. Crucially, the Ancillary Services Market operates on a "Pay-as-Bid" basis, meaning accepted resources are remunerated at the price they offered, rather than a marginal clearing price.

The MSD is structurally divided into two distinct phases:

1. MSD ex-ante: it is subdivided in six sessions and the offers presentation phase occurs only in MSD1, which closes at 17:00 on the day prior to delivery (D-1). On the basis of the last available results of the electricity market, Terna selects offers to resolve inter-zonal congestions and create reserve margins (secondary and tertiary reserve). This phase is predictive and, by imposing constraints to the following sessions of the electricity market, it

aims to position the grid in a secure state before real-time operations begin.

2. **Balancing Market (MB, Mercato di Bilanciamento):** operating in continuous trading up to one hour before physical delivery (H-1), this mechanism allows Terna to correct instantaneous deviations between supply and demand, ensuring the system balance and safety. Specifically, the MB manages the real-time activation of the secondary and tertiary reserve margins that were previously procured during the MSD ex-ante phase. Conversely, the primary frequency reserve is excluded from market bidding, being a mandatory service automatically activated by local frequency deviations.

Historically, the MSD was designed around "Relevant Production Units" (UPR, Unità di Produzione Rilevanti), typically large thermal or hydro plants (>10 MVA), which were the only ones enabled to participate in the MSD. Being the smaller distributed resources excluded, the flexibility potential of the new dominant resources, e.g. distributed PV, storage, and demand response, which remained locked out of the market due to high entry thresholds and rigid technical requirements, was not fully exploited.

In the current scenario, this market architecture would lead to a system with abundant installed capacity but scarce flexibility, resulting in curtailment of renewables and increased dispatching costs (uplift fees) borne by end-users. The TIDE reform aims to resolve this by transitioning from a "relevant unit" model to a "technology-neutral" model where any asset capable of modulating its profile can provide services.

2.2 The path toward the Local Flexibility Markets

Responding to the evolution of the national power system, the Italian regulatory framework began an extensive reform in 2017. The strategy was to open the MSD to non-traditional resources through pilot projects [34] and subsequently to empower DSOs to manage local constraints [31].

2.2.1 The UVAM experience

The most significant breakthrough of the pilot phase was the creation of Virtual Aggregated Mixed Units. Resolution 300/2017/R/eel [34] dismantled the barrier of "relevance," allowing production units <10 MVA, consumption units (with the demand response), and storage systems to

aggregate and participate in the MSD.

A UVAM functions as a single virtual dispatchable node managed by an aggregator, who aggregates diverse individual assets, such as industrial loads, storages, and backup generators, to meet the 1 MW minimum bid requirement. To incentivize the formation of these new entities, the regulator introduced a remuneration scheme composed by an annual premium (€/MW/year) for the commitment to be available during critical hours (typically late afternoon/evening), i.e. a capacity payment, and by a pay-as-bid remuneration for the energy actually delivered when activated by Terna, i.e. an activation payment.

The UVAM pilot project demonstrated the feasibility of aggregation but also the fragility of its economic model under current market conditions. Results indicate that UVAMs have not yet significantly contributed to overall system flexibility. Despite improved reliability, over the years activated energy volumes have stagnated while participation has contracted. As reported in [36], the number of qualified UVAMs dropped to 161 by the end of 2024, representing a loss of 22 units compared to the previous year. More importantly, this decline is reflected in the contracted capacity: out of a theoretically available contingent of 1,000 MW, the actually contracted capacity plummeted from a historical peak of nearly 100% (in 2019-2020) to roughly 17% of the available contingent by the end of 2024. Furthermore, data reveal a persistent dichotomy between “availability” and “activation”. In 2023, the average energy offered during contracted evening hours (18:00–22:00) was approximately 90.6 MWh/h, compared to just 21.9 MWh/h in free periods. This trend intensified in 2024, with offers averaging 61.6 MWh/h during contracted hours versus a mere 10.8 MWh/h in non-contracted periods. So the actual energy offered on the market dropped precipitously outside of the mandatory contracted hours and this suggests that the pure merchant value of flexibility on the Italian MSD is currently insufficient to sustain the aggregator business model independent of capacity payments.

2.2.2 The active role of the DSO and Local Flexibility Markets

The second phase of this market evolution is characterized by the pivotal role of DSOs. As congestion migrates from the transmission to the distribution network, the role of the DSO has undergone a profound evolution, formally established by the European Union's Directive 2019/944 [5] and formalized in Italy by Resolution 352/2021/R/eel [31], moving from a passive grid manager to a key player with a dual regulatory mandate.

DSO as a Neutral Facilitator

The first role, initially tested under the UVAM pilot projects, is to support the TSO, Terna, in procuring national global ancillary services (like reserves and balancing) from resources connected to the distribution network. In this capacity, the DSO acts as a neutral facilitator, ensuring that the TSO's activation of these resources does not violate local grid security constraints.

This is managed through the "Traffic Light" Model:

- **Static Traffic Light:** during the initial qualification of an aggregated unit (e.g., UVAM), the DSO verifies for structural grid risks ex-ante. A red light excludes the resource; a yellow light imposes technical operating limits (capability); and a green light permits full modulation. Due to the "silence-is-consent" rule, a green light is automatically assigned and the qualification request is considered accepted if the DSO does not respond within five days.
- **Dynamic Traffic Light:** a subsequent pilot project focused on Dynamic TSO-DSO Coordination introduced an additional operational layer, which does not bypass the initial static validation but acts directly prior to market activation. For each delivery period, in time for the submission of bids by the Service Provider (SP), the DSO performs a verification of the overall modulation capacity of the flexibility resources connected to specific DSO perimeters. Depending on mutable local grid conditions, the DSO can dynamically allow full modulation, restrict it to a partial capability limit, or completely prevent the activation.. Terna, during the activation of the requests, is responsible for respecting any dynamic limits imposed.

DSO as an Active Buyer

The second fundamental role involves the active procurement of national local ancillary services. This strategy is designed to manage non-structural constraints, such as local congestion and voltage fluctuations, while potentially deferring the need for expensive grid infrastructure expansions. Following the formal framework established by ARERA in 2021 via resolution 352/2021/R/eel [31], these competitive procedures led to the creation of Local Flexibility Markets. In this framework, the DSO acts strictly as a market participant, specifically, the buyer of flexibility, while, in this initial phase, the market interface platforms are operated by independent institutional or private entities, such as GME or Piclo¹. In recognition of the

¹ <https://www.piclo.com/>

experimental nature of these markets, the resolution mandated transparency and technological neutrality, effectively inviting DSOs to put in place the transition through dedicated pilot projects. Furthermore, to ensure sufficient cash flow to cover local flexibility needs at the national scale in this experimental phase, the procurement costs incurred by DSOs are financially backed by a dedicated national fund for exceptional events, resilience, and other special projects.

Currently, in the Italian landscape there are three major pilots that serve as a comparative laboratory for future regulation: RomeFlex (managed by Areti), EDGE (e-distribuzione), and MindFlex (Unareti). While these projects operate in diverse environments, they share several unifying principles. Most notably, all three respect the principle of technological neutrality, enabling participation from a wide spectrum of assets, including PV systems, BESS, electric vehicle charging stations, and industrial, commercial, and residential loads. Also, the economic architecture of these pilots relies on the same dual remuneration scheme designed to balance availability with actual performance. Participants typically receive a capacity payment (€/MW/year) for their commitment to remain available during specified windows, alongside an activation payment (€/MWh) based on the energy actually delivered.

Despite these shared goals, the specific technical requirements reveal different strategic priorities. For instance, the threshold for resource participation varies significantly: while RomeFlex facilitates high granularity with a minimum requirement of just 300 W, EDGE originally targeted larger-scale contributions with a 100 kW minimum for auctioned products. MindFlex, which initially set a 20 kW modulation capacity for larger connection points, has recently moved toward alignment with Areti's requirements as of 2025. Thus, based on the pilot project, even the smallest consumers can take part in flexibility services, effectively enabling the engagement of any type of user, including residential ones that, in meantime, might participate in an Energy Community.

The activation strategies further distinguish these pilots. While most began by establishing utilization prices through term contracts, Areti has introduced a more dynamic approach by incorporating a spot market. This evolution allows the utilization price to be determined for each individual session, moving the system closer to a real-time, market-responsive flexibility framework.

To systematically quantify the scale of these experimentations in 2024, the pilot projects can be compared across three primary dimensions: the

maximum requested modulation power, the annual hours of experimentation, and the allocated financial budget. In Table 2.1 the dimensions of the three pilot projects are reported.

Table 2.1 – Pilot projects' dimensions [36]

	<i>Power [MW]</i>	<i>Hours/year</i>	<i>Budget [€]</i>
Areti	20	2190	4,950,000
e-distribuzione	10.35	7392	1,287,000
Unareti	9	364	717,000

The first-year results provided crucial insights into the viability of the MLF model [36]:

- RomeFlex: demonstrated robust liquidity in the urban environment, successfully assigning more than the 18 MW of flexibility requested across its auctions. To provide context on the volume of controllable assets, RomeFlex engaged between 81 and over 150 registered resources across different auctions, managed by up to 13 Service Providers. The use of a dedicated device for data exchange between stakeholders and the evolution of the project, with the introduction of a spot market, confirmed the feasibility of reliable, highly-observable flexibility services at the low voltage level. The high success rate of activation (80-90%) indicated operational reliability, but the hardware requirement suggested a higher barrier to entry.
- EDGE: experienced low liquidity, contracting only 1.16 MW against a requirement of over 10 MW. However, a significant number of assets were initially accredited, with 495 enabled resources and 5 participating SPs. The low participation in the actual contract allocation highlighted the difficulty in aggregating resources in semi-rural areas, often because the enabled resources were not geographically located on the critical grid nodes where flexibility was actually needed. Furthermore, the risk-based procurement model, with activation contingent on non-structural, specific grid conditions, provided less predictable revenue certainty to the BSPs, making the business case more fragile than in Rome.
- MindFlex: successfully procured its target of 5 MW of "upward" flexibility (demand reduction) to manage anticipated summer peaks. However, minimal activation (only 4 hours) due to mild weather conditions revealed a key economic insight: the primary value of the service is often in the capacity payment. The commitment to be

available acts as a necessary insurance policy for the DSO against volatile and extreme weather events, even if the resource is rarely dispatched.

The diversity of these pilots demonstrated the technical feasibility of distributed flexibility and confirmed the market's interest in these services, despite numerous initial difficulties. A significant number of participants engaged in the experimentation but were not always able to actually provide the service, which is considered a natural part of the learning curve.

However, the pilots also highlighted that the economic model for aggregators remains highly sensitive to factors like location (urban vs. rural), procurement model (term vs. conditional), and the need for standardized, low-cost verification technology. Furthermore, operational issues such as the difficulty in impacting critical nodes and forming viable aggregates revealed a gap between theoretical market design and field reality. While the upcoming TIDE framework addresses the regulatory side of this complexity, it does not solve the technological deficit. This specific observation motivates the research focus of this thesis (Pillar 3), namely the development of the advanced monitoring and optimization tools necessary to transform these pilot experiences into a scalable, automated reality.

2.3 The key innovations of the TIDE framework

The TIDE, which entered into force, as a transitory phase, on January 1, 2025, codified the structural reorganization of the renovated Italian electricity market. It integrates previous regulations with innovative mechanisms to govern the modern grid, constituting a comprehensive revision of the dispatching rules. Its objective is to integrate local and global resources into a unified merit-order stack, allowing a number of new players to participate in the energy market and then to contribute to the flexibility and the overall stability of the grid. This initiative marks the shift from an experimental stage to a robust, industrial-scale market framework, though it is currently still in a transitional phase.

2.3.1 Technological neutrality and new unit taxonomy

The defining pillar of TIDE is technological neutrality. The historical distinction between "Relevant" (programmable, >10 MVA) and "Non-Relevant" units is abolished. TIDE introduces a new taxonomy based on capability rather than size or technology:

- Individually Qualified Units (UAS): single units capable of bidding individually (evolution of the old UPR). These are typically larger plants connected to HV (high voltage) or MV grids.
- Virtually Aggregated Nodal Units (UVAN): aggregates of resources located at the same transmission node. Because their grid position is precise, UVANs can provide local services (congestion relief) to both TSO and DSO. This category is crucial for industrial clusters or energy communities connected to the same primary substation.
- Virtually Aggregated Zonal Units (UVAZ): aggregates spread across a wider market zone. Since their location is wider, they are restricted to providing balancing services (frequency restoration) where specific location is less critical, but cannot resolve nodal congestions.

This classification allows diverse technologies to compete on equal footing, provided they meet the technical performance requirements (ramping, duration). It removes discriminatory barriers, formally opening the MSD to any resource capable of modulating its consumption or production profile, including small-scale production units, storage systems and consumption units. The concept of aggregation thus becomes pivotal, as it allows distributed energy resources to be transformed into a unified virtual asset of system flexibility.

2.3.2 New market roles

Another important innovation, incorporated in the TIDE framework, is the aforementioned DSO's evolving dual role and the birth of the local flexibility markets. While the regulation of MLF is still being defined through ongoing pilot projects, TIDE clarifies the counterparts to the network operator (TSO or DSO) for flexibility procurement, represented by Service Providers. A key feature of TIDE is thus the introduction of a critical unbundling of market roles, a separation that aims to encourage competition and specialization by distinguishing between the commercial responsibility for energy programming and the technical responsibility for providing system services:

1. the Balance Responsible Party (BRP) retains the commercial responsibility for ensuring that the energy injected and withdrawn by its associated portfolio of clients and assets matches its submitted schedule (or "program"). The BRP is financially accountable for any imbalances in the energy market. This is typically the trader or wholesaler.

2. the Balancing Service Provider (BSP) is the entity responsible for the physical provision of flexibility services to the TSO or DSO. The BSP aggregates individual resources and manages their technical pre-qualification and activation for ancillary services.

Critically, TIDE provides a mechanism to sterilize the BRP's financial position when a resource under its portfolio is moved by a BSP at the TSO/DSO's command, ensuring the BRP is not penalized for a grid service. In the traditional model, the user of the grid was responsible for both roles. Separating them allows a specialized aggregator (BSP) to optimize the flexibility of a portfolio of assets without taking on the massive commercial risk of wholesale energy trading (BRP). This is particularly important for Renewable Energy Communities, which may want to offer flexibility services without becoming full-fledged energy traders.

2.3.3 Harmonization of Local and Global Services

In order to integrate the new local flexibility needs and to align with European standards, the TIDE reform fundamentally restructures the ancillary services landscape as well. It introduces a clear distinction between Standard Products and Specific Products, harmonizing the procurement of resources for both global (TSO) and local (DSO) requirements.

TIDE reclassifies ancillary services into four main categories:

1. Frequency Ancillary Services: these incorporate the standard European balancing products, including Frequency Containment Reserve (FCR), automatic Frequency Restoration Reserve (aFRR, traded on PICASSO), manual Frequency Restoration Reserve (mFRR, traded on MARI), and Replacement Reserve (RR, traded on TERRE).
2. Non-Frequency Ancillary Services: such as voltage regulation and black-start capability, which remain largely location-specific.
3. Redispatching Services: dedicated to managing grid congestions and nodal constraints.
4. Extraordinary Modulation: a new service replacing the traditional interruptibility scheme to manage extreme adequacy risks.

In this case, the core innovation of TIDE, however, is the harmonization of these services to allow resources, particularly distributed ones, to stack value streams. The adoption of Standard Products is mandatory for global markets to enable cross-border exchange on European platforms, ensuring

maximum liquidity and competition. Conversely, when standard products cannot satisfy specific local constraints (e.g., voltage issues in a specific feeder), TSOs and DSOs can procure Specific Products.

To prevent conflicts—such as the TSO activating a resource that inadvertently causes a local overload—TIDE institutionalizes coordination mechanisms like the "Traffic Light" models piloted in earlier phases. This ensures that a resource's participation in global markets respects local grid limits, effectively reconciling the "global" need for balancing with the "local" need for security. Ultimately, this structure transitions the Italian system from a compliance-based model to a market-based model, where flexibility is a commodity traded across different timeframes and locations under a coherent regulatory umbrella.

2.4 The technological enablement

The evolution of the Italian electricity market, from the rigidity of the historical monopoly to the dynamic architecture of TIDE, represents a necessary adaptation to the physics of the energy transition.

As analysed in this chapter, the regulatory structures are now largely in place: the pilot projects, from UVAM to the Local Flexibility Markets, have successfully demonstrated that distributed resources can provide carrier-grade grid services, effectively validating the principle of "technological neutrality".

However, a critical misalignment remains. While the market design is becoming increasingly established, the challenge of technological enablement remains a subject of ongoing investigation. The TIDE framework constructs the "venue" for flexibility trading, but it does not provide the "ticket" to enter. Specifically, for the mass participation of residential prosumers envisioned by EU Directives, the current barriers, high investment costs for monitoring, complexity of aggregation, and lack of observability, are still prohibitive. This creates the central problem that this thesis addresses in the following chapters. A robust market design is ineffective if the transaction costs for small-scale flexibility providers outweigh the benefits. Consequently, regulatory innovation requires a matching technological foundation. Low-cost, scalable energy monitoring systems and automated analytics constitute the essential digital infrastructure required to implement the TIDE framework and ensure the practical viability of Energy Communities.

CHAPTER 3

Energy Communities: Aggregate Users to Enable the Energy Transition

In recent years in Europe, the regulatory efforts to enforce the technological neutrality principle and to open electricity markets to every type of resource have been huge. Among these, there was the introduction of Energy Communities (ECs), formally born at the EU level with the Clean Energy Package. ECs empower citizens making them active, no longer passive, users of the distribution grid, who have the possibility to become producers and owners of plants, capable of generating economic, social and environmental value in their territory. Indeed, the Energy Community represents a paradigm shift from individual optimization to collective intelligence. It serves as the primary layer of aggregation, absorbing local imbalances through collective self-consumption (CSC) before exposing the residual flexibility to the national grid.

Concurrently, aggregation through energy communities enables small-scale users to participate in local ancillary services, from which they would otherwise be excluded due to prohibitive transaction costs and technological fragmentation. An individual household with a 3 kW PV system and a residential battery is, from the perspective of the grid operator, electrically invisible and operationally uncontrollable, but an EC, collecting distributed

resources and serving as Service Provider, could offer valuable flexibility services to the DSO.

Therefore, the Energy Community represents an opportunity for grid operators by acting as a flexibility provider. It facilitates the local optimization of power flows and enhances the quality of service due to economies of scale, specifically, the more efficient management of aggregated assets compared to individual users [37][38]. Conversely, flexibility offers ECs an opportunity to enhance their business model by providing ancillary services to the grid, thereby generating a new revenue stream [39].

This dissertation posits that the solution to the flexibility gap is not solely technological, but it is also organizational. Thus, this chapter introduces the "First Pillar" of the proposed flexibility architecture: the Energy Community. Being RECs increasingly recognized as a key topic in the development of MLF, EC is not simply defined as a legal entity for sharing energy, but also as a techno-economic aggregation vehicle to foster end-users participation in energy markets [40]. By clustering disparate small users under a single virtual node, Energy Communities provide the necessary scale, legal personality, and economic incentive structure to transform passive consumers into active market participants.

This chapter is structured to provide a comprehensive examination of the Energy Community landscape, moving from the regulatory framework to the practical implementation. In Section 3.1 an exhaustive analysis of the regulatory genealogy of ECs, tracing the transposition of European Directives into the Italian legal framework, is provided. The mechanics of the "Virtual Self-Consumption" model and the incentive structures are examined, and the current deployment status in Italy, identifying the critical barriers hindering mass adoption, is assessed. Section 3.2, building on the regulatory constraints, introduces a simulation tool developed to perform scenario and feasibility analyses of Energy Communities within the Italian context. Finally, in Section 3.3, it is illustrated how advanced monitoring techniques can provide the granular visibility required to manage an aggregated portfolio of users and how data-driven technologies can allow maximizing the benefits of RECs, deriving from energy sharing and from the participation in flexibility markets, by automatically optimizing the energy management of the community itself.

3.1 The Regulatory Framework and REC State of the Art

The formal introduction of Energy Communities in the Italian energy landscape is not a spontaneous market phenomenon but the result of a deliberate, top-down legislative transposition process. To comprehend the operational potential of these entities, and their limitations, it is first necessary to acquire an in-depth understanding of the complex regulatory framework by which they are governed. This framework determines everything, from the maximum size of a PV plant to the geographic radius of membership, and ultimately, the economic viability of the aggregation itself.

3.1.1 From EU Directives to the Italian Regulations

The legal genesis of Energy Communities lies in the European Union's "Clean Energy for All Europeans" package, a comprehensive legislative suite designed to facilitate the energy transition. Two directives within this package are fundamental, establishing a dual definition of "community" that has profound implications for national implementation. First, Renewable Energy Communities were introduced in the Renewable Energy Directive II (RED II) in 2018 [4], and then, Citizen Energy Communities were defined in 2019 via the Internal Electricity Market Directive (IEMD) [5]. In Table 3.1 the two definitions are presented.

Key differences exist between RECs and CECs, stemming from their definitions and directives. Firstly, CECs have no technology limitations, while RECs are restricted to renewable energy technologies. Secondly, REC membership requires proximity to the energy projects, unlike CECs, which have no such restrictions. Thirdly, concerning ownership and governance, small and medium-sized enterprises (SMEs) can control a REC if their primary activities are non-energy related, but this is limited to small and micro enterprises for CECs. Nevertheless, large companies can participate in EC decision-making, provided social and citizen action strictly controls the regulatory basis. Finally, CECs activities are limited to the electricity sector, whereas RECs ones can operate across all energy sectors. [41]

With the two definitions, the EU has set the basic regulatory framework for ECs. However a lot of freedom is given to the member states to transpose the EU legislative framework into national law and, in particular, this dichotomy presents a strategic choice for national regulators. In Italy, the

Table 3.1 – REC and CEC definitions

	Renewable Energy Communities (RECs)	Citizen Energy Communities (CECs)
<i>Form</i>	“a legal entity” [RED II, Article 2, para 16]	“a legal entity” [IEMD, Article 2, para 11]
<i>Membership Eligibility</i>	“in accordance with the applicable national law, is based on open and voluntary participation, is autonomous, and is effectively controlled by shareholders or members that are located in the proximity of the renewable energy projects that are owned and developed by that legal entity; the shareholders or members of which are natural persons, SMEs or local authorities, including municipalities” [RED II, Article 2, para 16]	“that is based on voluntary and open participation and is effectively controlled by members or shareholders that are natural persons, local authorities, including municipalities, or small enterprises” [IEMD, Article 2, para 11]
<i>Primary Purpose</i>	“the primary purpose of which is to provide environmental, economic or social community benefits for its shareholders or members or for the local areas where it operates, rather than financial profits” [RED II, Article 2, para 16]	“has for its primary purpose to provide environmental, economic or social community benefits to its members or shareholders or to the local areas where it operates rather than to generate financial profits” [IEMD, Article 2, para 11]
<i>Activities</i>	The RED II also states that RECs shall be entitled to produce, consume, store and sell renewable energy, including through renewables power purchase agreements. [RED II, Article 22, para 2(a)]	“and may engage in generation, including from renewable sources, distribution, supply, consumption, aggregation, energy storage, energy efficiency services or charging services for electric vehicles or provide other energy services to its members or shareholders” [IEMD, Article 2, para 11]

focus has heavily favored the REC model, primarily because it is linked to renewable generation national targets, establishing incentive mechanisms

and thus making the REC an economically dominant vehicle compared to the CEC.

3.1.1.1 The Italian Transposition

After a transitory phase, which had been enacted by Law 8/2020 [42] and during which Italy had introduced the concept of "Collective Self-Consumption" and had activated the first RECs, the full transposition was achieved with Legislative Decree 199/2021 (implementing RED II) [28] and Legislative Decree 210/2021 (implementing IEMD) [17], entering into force in December 2021.

To make the decrees operational, several regulations were published in the following years. In 2022, ARERA Resolution 727/2022/R/eel [29], alongside the issuance of the Integrated Text on Collective Self-Consumption (TIAD), later integrated and modified by Resolution 15/2024/R/eel [43], established the functional mechanisms and economic remuneration schemes related to collective self-consumption configurations. Subsequently, the incentive mechanism was defined by the Ministry of Environment and Energy Security (MASE) via Decree No. 414 of December 7, 2023 [44], which was amended by MASE Decree No. 127 of May 16, 2025 [45]. Finally, the technical rules regarding incentive allocation, drafted by the GSE (Gestore dei Servizi Energetici), were approved by MASE Directorial Decree on February 23, 2024 [46], and updated by Decree 228 on July 17, 2025 [47].

Although both RECs and CECs have been regulated, the current Italian regulatory status provides an incentive mechanism only for collective self-consumption configurations related to decree 199/2021 and therefore, regarding ECs, for RECs. Specifically, the configurations admitted to the incentive mechanism are three:

1. RECs: are legal entities made up of groups of subjects (such as natural persons, local authorities, SMEs but not large enterprise, territorial bodies, research and training bodies, religious bodies) located in the same market area (the market zones shown in Figure 2.1) of renewable energy production plants that group on a voluntary basis to produce, share and consume clean electricity. Selling energy cannot be the main business activity of REC members; the primary purpose of the REC is to provide environmental, economic or social community benefits. Moreover, members or partners of the community retain the rights of end

customers, including the right to choose their own seller, and can exit the configuration at any time.

2. Groups of collective self-consumers: are groups of final users located in the same building or condominium, and ruled by a private law agreement, that group on a voluntary basis to produce, share and consume clean electricity. The renewable power plants must be under the full control of the members and selling energy cannot be the main business activity of the members. Moreover, members or partners of the configuration retain the rights of end customers, including the right to choose their own seller, and can exit the configuration at any time.
3. Remote individual self-consumers: is a single end customer who produces and consumes energy, produced by renewable power plants under his full control.

The renewable power plants admitted to the configurations must have entered into operation after the entry into force of the decree (hereinafter referred to as “new power plants”), pre-existing ones can be admitted too, but in the case of RECs only up to 30% of the total production capacity of the community.

To share the energy among the participants of the configurations, a virtual model for self-consumption is adopted, in order to separate the physical flow of electricity from the economic valuation of sharing, granting members independence but introducing complexity in settlement. In a physical self-consumption scheme, a private cable connects the generator directly to the load behind the meter and the grid only sees the net flow. In the virtual model adopted for RECs, members maintain their existing individual connection points (PODs) to the public grid, which is used to “virtually” exchange energy, thus eliminating the need for private bilateral contracts among community members. Consumers continue to buy electricity from their chosen retail suppliers and pay their full bills, while producers continue to inject electricity into the grid and sell energy in the electricity market. Shared energy is defined as the hourly minimum between the total energy injected and the total energy withdrawn by the community's members and is determined *ex-post* on a monthly basis by the central counterparty, the GSE. Furthermore, prosumers, end-users equipped with behind-the-meter RES generation assets, can actively contribute to this energy sharing mechanism by acting simultaneously as both consumers and producers. This model allows for unlimited scalability, since it leverages the DSO’s existing smart

meters for energy measurements and, then, connecting a new member is merely a database update, eliminating the need for new physical installations. However, it requires a robust data infrastructure to synchronize hourly meter readings from a multitude of users and to optimize energy consumptions.

Although the geographic scope defining a REC is very broad, covering an entire market zone, the area narrows for the purpose of energy sharing eligible for valorization and incentives: PODs associated with community members must be connected to the same primary substation. This subset of the shared energy is referred to as virtually self-consumed or collectively self-consumed. Moreover, incentive tariffs only reward energy produced by new power plants with a maximum nominal power of 1 MW; only a valorisation contribution is recognized for the energy produced and shared by the remaining plants that don't comply with these two last requirements.

3.1.1.2 Incentive and valorisation mechanisms

In Italy, shared energy that is virtually self-consumed, as defined in the previous section at an hourly resolution, is rewarded with a feed-in premium and is valorised with the restitution of some tariff components on a monthly basis. This valorisation varies for different configuration types.

In particular in REC, since shared energy is consumed locally (under the same primary substation), it does not traverse the high-voltage transmission network and the system avoids transmission losses and certain grid costs. Therefore, a portion of the variable transmission charges paid by members, proportional to shared energy, is refunded to the community. The grant for the valorisation of the shared energy C_{VSE} is obtained from the product between the variable component of the transport fee for low voltage users (CU_{tr}), expressed in €/MWh, and the energy shared under the same HV/MV substation (E_{VSE}), as shown in (3.1).

$$C_{VSE} = CU_{tr} \times E_{VSE} \quad (3.1)$$

The value of the transport fee is determined annually by ARERA and for 2026 it was equal to 11.90 €/MWh [48].

The feed-in premium (C_{FIP}), conversely, is not attributable to avoided network costs; rather, its purpose is to promote the construction of new generation facilities that serve the community and it is the primary economic driver for RECs. This premium is granted exclusively for the virtually self-consumed energy E_{VSC} and is guaranteed for a period of 20 years. For each RES plant contributing to the collective self-consumption, C_{FIP} is calculated as the product of E_{VSC} and an incentive tariff (TIP), as expressed in (3.2):

$$C_{FIP} = TIP \times E_{vsc} \quad (3.2)$$

Regarding the allocation of shared energy among multiple RES plants, they are ranked by their commissioning dates, where the oldest facilities are the first to contribute to the energy sharing. The *TIP* depends on the nominal power (P_n) of the power plant from which the energy is produced and shared and it is determined according to the following equation, expressed in €/MWh:

$$TIP = x + \max(0; 180 - Pz) \quad (3.3)$$

where:

- x is the fixed part of *TIP* and depends on P_n
- Pz is the hourly zonal price of the electric energy determined in the MGP

The *TIP* also has an upper limit depending on P_n . In Table 3.2 the values in €/MWh of x and of the limits are shown.

Table 3.2 – Parameters for the *TIP* calculation in €/MWh

	$P_n \leq 200 \text{ kW}$	$200 \text{ kW} < P_n \leq 600 \text{ kW}$	$P_n > 600 \text{ kW}$
x	80	70	60
<i>TIP max value</i>	120	110	100

Furthermore, for PV plants there is a correction factor, which is added to the *TIP*, depending on the geographical area, as reported in Table 3.3.

Table 3.3 – *TIP* geographical correction factor

Geographical area	Correction factor
Regions of central Italy (Lazio, Marche, Toscana, Umbria, Abruzzo)	+4 €/MWh
Regions of northern Italy (Emilia-Romagna, Friuli-Venezia Giulia, Liguria, Lombardia, Piemonte, Trentino-Alto Adige, Valle d'Aosta, Veneto)	+10 €/MWh

This economic architecture creates a strong signal for synchronization, since the REC maximizes revenue only when generation and consumption happen simultaneously, favouring behaviour required to solve the "flexibility gap" introduced in Chapter 1.

3.1.2 The Italian Scenario: Current Status, Barriers, and Drivers

With the regulatory framework fully defined as of early 2024, which has significantly expanded the operational scope of RECs by raising the capacity limit for individual plants to 1 MW and extending the aggregation perimeter to the primary substation, Italy is moving from the experimental phase to mass deployment. However, data reveal a market that is growing but remains far from the multi-gigawatt scale required by national targets.

As analysed in the Electricity Market Report 2025 [49], the sector is currently at a crossroads between theoretical potential and practical execution. The latest census indicates 876 active configurations of collective self-consumption in Italy, of which 48% are RECs, and a total involved power generation capacity of around 85 MW. While the nineteen fold increase in configurations represents a significant year-on-year growth, the absolute numbers suggest that the market is experiencing a "ramp-up" phase, but it's still in its infancy compared to the 5 GW target set for 2027. The average size of these initiatives remains small, in terms of both members number and installed power capacity. The median installed power for RECs is approximately 17-20 kW, with only a handful exceeding 200 kW. This "micro-scale" prevalence is a legacy of the experimental phase (which capped plants at 200 kW) and reflects a caution among early adopters.

Indeed, the widespread adoption of RECs in Italy faces a complex landscape of technical and non-technical hurdles [50]. regulatory and administrative complexities remain a primary bottleneck; the authorization processes for renewable installations are often fragmented and slow, which discourages potential investors and aggregators. On the technical front, the digitalization of the distribution grid is still uneven. Specifically, the incomplete rollout of second-generation (2G) smart meters in certain regions restricts the granular data availability required to calculate shared energy precisely and to implement advanced demand-side management strategies. Beyond these structural issues, Khorrami et al. [51] highlight significant social and economic barriers, including high initial investment costs and the challenge of engaging citizens. Lacking specific technical expertise, many potential participants perceive the community model as financially risky or overly complicated. This issue of "social acceptance" is particularly critical because the REC model relies fundamentally on the active behavioural change of prosumers to balance local generation and consumption profiles.

Nevertheless, the opportunities arising from the effective deployment of RECs are substantial and align perfectly with the broader decarbonization

and efficiency goals of the energy transition. Economically, the current incentive schemes, combining feed-in premiums on virtually self-consumed energy with capital grants from the National Recovery and Resilience Plan (PNRR), which can cover 40% of initial investment in municipalities with less than 50,000 inhabitants, provide a solid return on investment, especially for smaller municipalities [50]. However, the value of RECs extends beyond simple financial returns; they constitute a fundamental architecture for unlocking demand-side flexibility. By aggregating fragmented loads and generation units, energy communities can technically operate as Virtual Power Plants (VPPs), providing essential local ancillary services such as congestion management to the DSO [49]. This capability effectively transforms passive consumers into active grid assets, thereby increasing the hosting capacity of the distribution network. Furthermore, RECs offer profound social benefits by addressing energy poverty through the redistribution of revenues to vulnerable members [51].

In this context, the integration of advanced monitoring systems and data analytics, the core focus of this thesis, emerges as a technological enabler required to convert these theoretical opportunities into measurable operational performance.

3.2 The REC Scenario and Feasibility Analysis Tool

As stated in the previous section, the established legislative architecture introduces an incentive mechanism where the REC economic value is strictly correlated with the synchronicity between generation and consumption within the perimeter of a primary substation. Consequently, traditional deterministic sizing methodologies, which often rely on annual energy averages or simplified "net metering" assumptions, are mathematically insufficient for predicting the financial viability of a REC.

The effectiveness of any optimization or management strategy is, therefore, primarily contingent on the community's configuration itself. Performance metrics such as collective self-consumption and individual bill savings are profoundly influenced by input parameters like the members' power profiles, the percentage of prosumers, and the inherent heterogeneity of the group. Indeed, the main factor influencing the performance of an EC is the people forming it, with the diversity of their load power profiles proving to be more beneficial than simply oversizing the generation or storage asset capacity. This diversity contributes to the natural decrease of

demand peaks, increasing the chances for local self-consumption. Specifically, studies suggest that configurations comprising a strategic proportion of prosumers, such as around 75%, and exhibiting heterogeneous load profiles yield the most favourable economic outcomes. [52]

The dependence on factors such as the type of user and the load profile makes it desirable to use support tools for the technical-economic evaluation of scenarios and for feasibility analyses across a wide spectrum of member configurations to predict the economic and technical performance of a REC and then to design the optimal one.

So a "REC Scenario and Feasibility Analysis Tool" was developed, integrating Python within Excel sheets with the goal to provide an user-friendly tool to simulate energy flows and financial outcomes for a community over a 20-year horizon, aligning with the duration of the Italian incentive tariff. The architectural philosophy of the tool is predicated on modularity and sequential data processing. As shown in Figure 3.1, it is structured into four distinct, interdependent logic blocks: *Production*, *Consumption*, *Energy Evaluation*, and *Economic Analysis*. The data flow moves from the physical determination of energy vectors (production and consumption) to their temporal integration (energy evaluation), culminating in financial valuation (economics).




Figure 3.1 – Flow diagram of the REC Scenario and Feasibility Analysis Tool

3.2.1 The Production Block

The Production block is responsible for estimating the renewable energy availability. While the tool is agnostic regarding the generation source, the current implementation focuses on PV technology, given its dominance in the Italian distributed energy landscape. To ensure accuracy, the tool implements two parallel computational methodologies: one to define the monthly average daily profile of hourly energy under clear-sky conditions, and the other to define the total energy production for each month of the year. Both of them were developed in python and integrated within Excel using the *xlwings* library, which makes it easy to call Python from Excel [53].

The first one is based on a physical model to simulate PV production. This model, referred to as the Clear-Sky Model (CSM), is a deterministic approach assuming no cloud cover and minimal uncertainty in the PV output profile, making it highly accurate for clear-sky days [54]. By taking the PV system's technical parameters and the site's monthly average hourly meteorological data as inputs, the core of the CSM is the estimation of the PV output, expressed by equation (3.4):

$$P_{system} = \frac{E_{g,pv} \cdot P_{peak} \cdot \eta_{panel} \cdot \eta_{inv}}{E_{STD}} \quad (3.4)$$

where:

- P_{peak} , the peak nominal power and η_{inv} , the inverter efficiency are fixed technical parameters of the PV plant,
- E_{STD} is the irradiance of standard test conditions set to 1000 W/m²
- $E_{g,pv}$ is the global irradiance reaching the plane of the array,
- η_{pan} is the relative efficiency factor of the panels.

$E_{g,pv}$ is determined as the sum of three components, beam $E_{b,pv}$, diffuse $E_{d,pv}$, and reflected $E_{r,pv}$ irradiance, each adjusted by site-specific shading factors. Each irradiance component derives from the global horizontal and diffuse irradiances adjusted by the ground albedo and the tilt angle β of the PV array. Specifically, $E_{g,pv}$ is modelled by equation (3.5) [55]:

$$E_{g,pv} = E_{b,pv} \cdot (1 - S_{dir}(\alpha, \gamma)) + E_{d,pv} \cdot (1 - S_{diff}) + E_{r,pv} \quad (3.5)$$

where S_{dir} and S_{diff} are the direct and diffuse shading factors that account for surrounding obstructions, and α and γ are the sun's azimuth and elevation, respectively.

The panel efficiency is primarily influenced by the module temperature T_m , calculated as shown in (3.6) [56]:

$$\eta_{pan} = 1 + \beta_c \cdot (T_m - 25^\circ\text{C}) \quad (3.6)$$

where β_c is the module temperature coefficient.

The temperature of the module is then estimated based on the ambient temperature T_a and wind speed W_s using the empirical relation in equation (3.7) [57]:

$$T_m = T_a + \frac{E_{g,pv}}{E_{STD}} \cdot (0.0712 \cdot W_s^2 - 2.411 \cdot W_s + 32.96) \quad (3.7)$$

In Figure 3.2 the output of the CSM is shown for the inputs in Table 3.4 and not considering any shadows.

Table 3.4 – Input parameter for the Clear-Sky Model

peak power	80 kW
latitude	44.36
longitude	8.57
azimuth	120°
tilt	10°
inverter efficiency	0.971




Figure 3.2 – Typical PV produced power daily profiles for each month of the year

Complementing the physical model, the tool includes a module that interfaces directly with the European Commission's Joint Research Centre

(JRC) PVGIS database via application programming interface (API) [58]. This function retrieves monthly energy data specific to the site's coordinates and to the PV system's nominal parameters, applying a fixed system loss factor (typically 14%) to account for cabling, inverter inefficiency, and soilage. The average monthly energy production, resulting from real-world solar radiation, is then used to recalibrate the hourly profile generated by the Clear-Sky Model, which inherently assumes perpetually clear-sky conditions.

In summary, the CSM is employed to derive the theoretical hourly generation profile, which is cross-referenced with consumption data to calculate the hourly energy sharing percentage relative to the CSM production. This percentage is then applied to the actual monthly energy yield retrieved from PVGIS, effectively recalibrating the hourly indices to align with the real-world monthly production totals. It is possible to have multiple PV plants with different nominal parameters in the same REC.

3.2.2 The Consumption Block

The tool ingests data regarding the composition of the community, categorized by user type (e.g., residential, office, commercial, industrial) and quantity. For each category, the tool assigns a specific standard load profile, derived from normalized datasets provided by GSE [59], that represent the average hourly consumption behaviour of a typical user class over a year, relative to their monthly consumption levels. Moreover, the tool allows the customization of load profiles, to ensure more realistic and heterogeneous representations, and distinguishes between members with and without a physical connection to the PV plant, to correctly compute the shared energy.

By taking the members user type and the monthly consumption as input, this block sums the individual profiles to create a community load curve, i.e. a matrix of hourly consumption values for the entire year, segmented by month, as shown in Figure 3.3 and Figure 3.4 for members respectively with and without the PV plant behind the meter, i.e. prosumers and consumers.

3.2.3 The Energy Analysis Block

The Energy Analysis block is the computational core where the "virtual" nature of the REC is resolved into physical and regulatory quantities.

This block takes the synchronized time-series data from the Production and Consumption blocks and applies the regulatory logic to partition energy flows. Here, shared energy and virtual self-consumption coincide, as all members of the REC are assumed to be connected under the same primary

substation. So virtually self-consumed energy E_{vsc} , which is the basis for the feed-in premium, is calculated for each hour according to the definition (3.8):

$$E_{vsc} = \min(E_i; E_w) \quad (3.8)$$

where E_i and E_w are respectively the energy injected to and the withdrawn from the grid for the purpose of the sharing.




Figure 3.3 – Monthly average daily load profiles for the prosumer




Figure 3.4 – Monthly average daily load profiles for the consumer

E_i is obtained by subtracting the load profile of physically connected users to the respective PV production profile and then aggregating all production profiles; instead, E_w is obtained by aggregating the load profiles of all members without a direct connection to the production plants, along with the non-physically self-consumed load portion of the remaining members.

Once the hourly E_{vsc} , with the others quantities of interest, is calculated for a typical day of each month, the monthly shared energy and the key performance indicators (KPIs) are computed. This monthly representative day approach is selected because the input consumption profiles are defined on a monthly basis. Furthermore, this method avoids excessive computational load within the spreadsheet environment, ensuring a responsive tool while still providing an effective visualization of seasonal trends for decision support. The tool then generates visualization graphs, as the ones in Figure 3.5 and Figure 3.6, that explicitly highlight the "wasted" energy (injected but not virtually self-consumed) and the "residual" demand (energy withdrawn from the grid that doesn't contribute to virtual self-consumption), guiding the user toward optimization strategies such as load shifting or storage integration.




Figure 3.5 – Monthly energy production KPIs

For example, in the case analysed in Figure 3.5, during the winter period, the generated and injected energy is fully absorbed by the community loads, whereas a generation surplus is observed during the summer months, particularly in May. Indeed, Figure 3.6 depicts a reduced summer load

profile; while the increased generation enables a greater volume of shared energy compared to winter, this lower demand simultaneously prevents the full exploitation of the total available injected energy.




Figure 3.6 – Monthly energy consumption KPIs

Ultimately, participation in the REC enables a partial valorisation of the energy injected by facilities that are oversized relative to the loads they are directly connected to (with the self-generation ratio always remaining below 10% in the example case).

3.2.4 The Economic Analysis Block

The final module, the Economic Analysis block, translates the energy flows into financial metrics over a 20-year horizon, corresponding to the lifespan of the Italian incentive mechanisms. The tool explicitly models the annual revenue streams available to Italian RECs, which are here assumed to be constant over the years and are composed of three primary components:

1. Savings from self-consumption: behind-the-meter self-consumed energy never enters the public grid, then its economic value is the full avoidance of variable grid costs, which are simply represented by a single voice named energy cost.
2. Revenue from energy sales: the energy injected into the grid can be sold and GSE offers a simple mechanism, called *Ritiro Dedicato*,

that values the energy at the hourly market zonal price (an average value is considered in the tool)

3. Incentive and valorisation: E_{vsc} is rewarded with a feed-in premium and is valorised with the restitution of some tariff components as described in subsection 3.1.1.2

Then, cash flow and discounted cash flow, given a discount rate, are calculated based on an initial capital expenditure (CAPEX), which is modelled using parametric benchmarks for specific PV pricing (€/kW_p) and installation-related outlays (engineering and commissioning), alongside annual operating expenses (OPEX), inclusive of maintenance costs and REC management fees. The tool also accounts for fiscal incentives, such as the PNRR capital grant available for plants installed in small municipalities, modelling the tax recovery over the years as a positive cash flow.

In addition to the graphical visualizations, like the one shown in Figure 3.7, the output of the tool includes standard financial KPIs like Net Present Value (NPV) and Payback Period (PBP).




Figure 3.7 – Discounted cash flow and NPV over 20 years horizon

For the case presented in Figure 3.7, where economic parameters reported in Table 3.5 were considered, revenues from energy sales and incentives for energy sharing contribute equally to the community's earnings, resulting in a PBP of approximately 10 years and in a 20 years NPV of about 75,000 € against an initial investment of about 120,000 €. The inflation rate was not taken into account.

Table 3.5 – Economic parameters used in the economic analysis example

energy cost	0.3 €/kWh
energy zonal price	0.1 €/kWh
PV plant capital cost	1550 €/kW
ordinary maintenance cost	25 €/kW/year
extraordinary maintenance cost	100 €/kW
management cost	1000 €/year
discount rate	4%

3.3 The Role of Advanced Monitoring and Digitalization

While simulation tools, such as the one presented in Section 3.2, are crucial for conducting scenario and feasibility analyses, evaluating initial capital expenditure, and forecasting the theoretical long-term economic viability of RECs, they remain intrinsically static planning instruments. Traditional deterministic sizing methodologies, which rely on historical averages and aggregated standard load profiles, are mathematically insufficient for predicting and managing the real-time operational dynamics of a community ecosystem. The operational efficiency of a REC, and consequently the maximization of its shared energy incentives and market revenues, relies entirely on the dynamic, real-time coordination of its underlying energy resources. This dynamic operational framework requires a sophisticated layer of technological enablement driven by digitalization.

The deployment of Internet of Things (IoT) architectures and Advanced Metering Infrastructure (AMI) provides the connectivity required to monitor power generation, energy storage, and consumption patterns with high temporal resolution. Central to this infrastructure are smart meters, which enable the bidirectional communication necessary to capture real-time data on consumption and renewable production. This data stream empowers local Energy Management Systems to implement sophisticated Demand-Side Management strategies, actively regulating demand to match local generation. DSM strategies within this context typically encompass Energy Efficiency Services (EES) and Demand Response programs. EES utilize technologies such as Building Automation Systems (BAS) and advanced data analytics to identify inefficiencies and reduce baseline consumption, for instance, through the optimization of HVAC systems. Conversely, DR programs focus on shifting flexible loads in response to external signals, utilizing either centralized or distributed control schemes to adapt to

renewable energy availability. Furthermore, the integration of IoT with Machine Learning algorithms enables the application of innovative techniques, among which Non-Intrusive Load Monitoring that allows for the disaggregation of electrical loads with minimal privacy intrusion, facilitating both the optimization of specific appliance usage and the equitable distribution of community revenues. [41]

Consequently, the integration of these technologies transforms the REC from a mere legal entity designed for administrative energy sharing into an active, intelligent aggregation vehicle. The acquisition of real-time data constitutes only the prerequisite capability. To fully realize the economic and technical benefits of the community, ranging from the internal maximization of virtually self-consumed energy to external participation in Local Flexibility Markets, these raw data must serve as the input for advanced operational optimization algorithms.

3.3.1 Operational Optimization Strategies

The transition from passive monitoring to active control requires an optimization framework capable of navigating the complex interplay between heterogeneous assets, stochastic generation, and regulatory constraints. As highlighted by international bodies such as the IEA, the digitalization of energy systems is not merely about data collection but about helping ensure that energy is consumed when and where it is needed through automated decision-making [60]. In the context of RECs, this involves determining the optimal dispatch of flexible resources to maximize collective self-consumption and economic revenues. Sangarè et al. [61] showed that by optimizing the collective self-consumption rate in energy communities by scheduling members' loads it is possible to reduce the total energy drawn from the primary grid by at least 15%. Fonseca et al. [62] proved that scheduling controllable loads in order to optimize the usage of locally generated energy in individual and collective terms allows the average cost of energy to be reduced by 6.5%.

But the challenges that advanced optimization strategies must address to unlock these benefits are diverse, from the design of the control architecture and the management of uncertainty to the integration of multi-vector assets.

A fundamental structural decision in designing a REC optimization strategy lies in the trade-off between computational efficiency and user autonomy, leading to a dichotomy between centralized and decentralized architectures. In a centralized framework, a community operator acts as a single decision-maker, managing the entire portfolio of flexible assets to

achieve a global optimum. While this approach theoretically maximizes total community welfare, it raises significant issues regarding data privacy and the fairness of benefit redistribution, necessitating robust ex-post allocation mechanisms to equitably redistribute the generated surplus [63]. Conversely, decentralized architectures preserve the autonomy of individual end-users, who manage their assets independently to minimize their own bills, interacting with the community through market-based mechanisms such as Peer-to-Peer trading. However, since these may converge to sub-optimal solutions compared to centralized coordination [64], recent advancements in distributed optimization, specifically the Alternating Direction Method of Multipliers (ADMM), are employed to bridge this gap, allowing communities to approximate global optimality while retaining local data privacy [65].

Regardless of the chosen architecture, the optimization logic must inevitably confront the inherent stochasticity of RES and the volatility of human behaviour. Deterministic models, which assume perfect foresight, fail when faced with real-world prediction errors. Consequently, the literature advocates for two-stage management strategies that decouple planning from execution. As proposed by Mustika et al. [66], this hierarchical approach begins with a day-ahead scheduling phase to commit to an optimal baseline, which is subsequently refined by a real-time adjustment phase where an online controller compensates for forecast deviations. This decoupling transforms the REC from a chaotic aggregate into a predictable entity capable of reliable participation in ancillary service markets.

To further enhance this reliability and fully exploit the community's flexibility potential, it is beneficial to extend the optimization perimeter beyond the electrical vector to encompass thermal systems. The integration of the heating and cooling sector represents a massive, often underutilized source of demand-side flexibility. Doroudchi et al. [67] demonstrate that by coupling electrical generation with common thermal energy storage, communities can significantly increase their self-sufficiency. By utilizing the thermal inertia of water tanks as virtual batteries, the EMS can divert surplus PV generation to pre-heat storage systems, thereby reducing grid exports during peak hours and avoiding expensive imports during the evening. This multi-vector approach (Power-to-Heat) reduces the reliance on capital-intensive electrochemical batteries, leveraging the lower marginal cost of thermal storage to buffer electrical variability.

However, the aggressive utilization of these distributed resources for economic minimization frequently creates a fundamental conflict with the

physical constraints of the distribution grid, such as voltage limits and transformer capacity. Uncoordinated cost minimization can induce synchronous peaks that stress the network. Coignard et al. [68] argue that sustainable REC operation therefore requires Multi-Objective Optimal Power Flow (MO-OPF) algorithms that explicitly model these constraints. Their analysis of Pareto fronts reveals that slight deviations from the economic optimum can yield significant improvements in grid stability, proving that by internalizing network constraints, RECs can evolve from passive grid users into active providers of flexibility services.

3.3.2 Data Availability and Advanced Monitoring

The effectiveness of any Demand-Side Management strategy is entirely contingent upon knowing what is consuming and what is producing power and when. So, the translation of sophisticated operational optimization strategies into functional, real-world deployments is severely bottlenecked by the chronic lack of granular empirical data regarding actual consumer behaviour and individual resource usage.

Due to this profound scarcity of real-world, disaggregated datasets, exacerbated by strict privacy regulations and the high cost of monitoring infrastructure, the current academic literature focusing on REC optimization is frequently forced to rely on public dataset or on synthetic approximations. In [69] the application of generative artificial intelligence, specifically Generative Adversarial Networks (GANs), to fabricate artificial load and generation profiles when exploring community clustering and sizing is demonstrated. In this computational framework, deep learning models are trained on the limited public datasets available to produce artificial daily load curves that mimic the statistical properties of actual residential consumers. These synthetic profiles are subsequently fed into evolutionary algorithms to perform multi-objective clustering operations, allowing the identification of optimal REC configurations by grouping virtual prosumers.

While the use of public dataset or on synthetic approximations is necessary for theoretical simulation and policy design, it constitutes a significant operational vulnerability when transitioning from the laboratory to the field. Synthetic data, no matter how statistically refined, inherently smooths over the chaotic reality of the grid edge. It cannot accurately capture the true, stochastic behavioural anomalies, the real-time physical constraints of specific hardware, or the sudden, unpredictable lifestyle shifts of actual human participants. An EMS trained, calibrated, and optimized purely on synthetic, aggregated data will inevitably struggle to execute reliable real-

time adjustments when confronted with the noisy, highly unpredictable realities of a physical microgrid.

Therefore, to bridge the gap between theoretical mathematical optimization and robust, real-life implementation, the energy management system architecture must incorporate an advanced monitoring system. This technological imperative paves the way for the subsequent chapters of this dissertation. Chapters 4 and 5 present applied research on advanced monitoring architectures and the development of NILM techniques suitable for the residential sector. This monitoring framework provides the essential data foundation that subsequently fuels the specific optimization algorithms and the intelligent EMS platform detailed in Chapter 6.

Energy Monitoring: case studies

Advanced monitoring systems are the key enablers for the implementation of innovative energy management systems, and, consequently, for the intelligent control of flexible assets available to the monitored users. Tracking the consumption and generation profiles of individual end-users over time enables a thorough analysis of their behaviour and a quantification of their interaction with the electrical grid. Therefore, monitoring systems, as a fundamental part of the third pillar (the technological one) mentioned in the first chapter of this dissertation, are the prerequisite for automated Demand-Side Management, providing the data foundation required to transform physical grid assets into dispatchable resources governed by advanced control algorithms.

These systems rely on a robust hardware measurement infrastructure for data acquisition, coupled with a scalable and reliable software layer for data historicization and processing. To meet these requirements, modern implementations widely employ IoT devices, thereby defining a distinct edge-to-cloud topology designed to support real-time decision-making processes [70]. IoT measurement devices and smart meters continuously acquire raw data, process, and transmit them using standard protocols, among which one of the most important is currently the Message Queuing Telemetry Transport (MQTT) protocol. MQTT provides lightweight

communication, reducing bandwidth overhead, and ensures reliable data delivery even under unstable network conditions [71]. Upon reaching the central server, the transmitted telemetry is historicized within different kinds of databases and then made available to logic processes, such as visualization and analytics layers, which transform aggregated datasets into configurable dashboards.

To validate the scalability and adaptability of such cloud-based monitoring architectures, the subsequent sections present three distinct empirical case studies developed in the context of this thesis. The proposed architectures, operating in different sectors and serving distinct purposes, may vary in their measurement acquisition methods depending on the specific needs of individual users, but they share the same tools for data historicization and visualization.

Thus, before presenting the case studies, Section 4.1 describes the primary hardware and software tools utilized. Then, in Section 4.2, the ANTES project is examined, focusing on the deployment of distributed monitoring within agricultural environments. The architecture correlates energy consumption directly with crop growth parameters to quantify efficiency deficits and trace accurate load profiles. Section 4.3 transitions to the complex infrastructure of healthcare facilities. Here, advanced Energy Management Systems serve for the systematic analysis of consumption profiles to identify structural efficiency opportunities and to lay the groundwork for the implementation of smart control strategies for automated HVAC regulation. Finally, in Section 4.4, the monitoring system implemented in an experimental laboratory, built as a Zero Energy Building (ZEB) is presented. In this case study, the architecture serves as the basis for Model Predictive Control (MPC) algorithms, with the final goal of executing optimal dispatch strategies for flexible assets, namely, BESS and heat pumps, to maximize the economic value of the building's energy footprint, also within the context of Energy Communities.

4.1 Hardware and Software Tools

A monitoring platform is generally developed as a potential multi-client, multi-site, and multi-user software architecture, wherein data are collected from the field and transmitted to a central server, where they are stored in a database and made available to logic processes. Platform's functionalities are accessible to users based on their specific access privileges via a web

service and the interface layout is designed to provide seamless access to data, graphs, tables, and frequently used analyses for each monitored unit, which the operator can select through a hierarchical diagram.

The monitoring platform is integrated with the field through interfacing instrumentation that enables bidirectional communication via various communication protocols, such as Modbus RTU or TCP, BACnet, TCP/IP, and others. The system allows the acquisition of field measurements from sensors and meters, as well as the storage of data (both automatically scheduled and manually on-demand), creating historical time series of the recorded measurements.

4.1.1 Energy Meters and Gateways

Monitoring systems can take into consideration one or more energy vectors, such as the thermal and electrical vectors, alongside climatic data (e.g., indoor and outdoor ambient temperatures) and other exogenous parameters. Within the case studies presented here, measurement acquisition primarily and preferably occurs by interfacing existing field devices with the proposed monitoring system; newly installed, dedicated instrumentation is employed only where strictly necessary and feasible. It was possible to adopt this approach as the proposed monitoring architecture is based on open protocols and is sufficiently flexible to allow this, significantly reducing commissioning and maintenance costs.

The hardware utilized specifically relates to the measurement of the electrical vector and to data concentrators (gateways) for transmitting measurements to the system's central server. The instruments used for electrical measurements in some of the projects presented in the subsequent sections are detailed below:

- 3-phase energy meter Shelly 3EM: this device calculates two-way consumption, tracking both produced and used energy for each of the three phases. It can be configured to independently measure three separate points of a single-phase electrical system. It measures active power, voltage, current, total energy, and power factor. The device provides a wireless access point for configuration and data visualization. Furthermore, it is enabled for Wi-Fi network connection and transmits measurements via the MQTT protocol. In Figure 4.1 the wiring diagram of the Shelly 3EM measuring three separate points is shown.




Figure 4.1 – Shelly 3EM wiring diagram¹

- Socomec DIRIS A-40 Multifunction Meter: this unit is a performance monitoring device designed for measuring, monitoring and managing electrical energy. The DIRIS A-40 offers a range of functions for measuring voltage, current, power, energy and quality (however, power quality functions are not used within the scope of this thesis). It enables analysis of a single-phase or three-phase load and allows the visualization of all electrical parameters via a backlit display. It provides Modbus RTU, via RS485, and Modbus TCP, via Ethernet, communication capabilities.



Figure 4.2 – Socomec DIRIS A-40²

- Socomec Multi-point power monitoring devices (DIRIS Digiware D50 + U30 and I45 modules): this configuration involves a display (D50) with a power supply module, a voltage measurement module (U30), and several cascaded current measurement modules (I45), each equipped with 3, 4, or 6 inputs, as shown in Figure 4.3. Communication with the monitoring and control system can be established via RS485 or Ethernet through Modbus protocol.

¹ <https://www.shelly.com/it/products/shelly-3em>

² <https://www.socomec.co.uk/en-gb/p/diris-a-40>

Voltage, current, power, and energy measurements are provided for each individual input.




Figure 4.3 – Socomec DIRIS Digiware multipoint power monitoring system³

These meters, both Socomec and Shelly, require Current Transformers (CTs) to operate, which enable the actual measurement of the monitored electrical quantities. CTs are used in measurement systems for high currents to step them down to more easily measurable values. They consist of a toroidal core through which the current-carrying cable passes (acting as the primary winding), while the secondary wire is wound around the core itself. CTs can be either split-core or solid-core. The former can be installed without disconnecting the current-carrying cable they will measure, whereas the latter requires a service interruption to disconnect and reconnect the cable into its housing. CTs feature various primary nominal current levels, generally ranging between 5 A and 6 kA, and must therefore be carefully selected to avoid measurement errors or excessively high secondary currents.

Regarding the gateway devices, the following equipment was utilized, with each device selected according to the specific project objectives and the existing field communication protocols or to the infrastructural and economic requirements of the end-user:

- InHand IoT Edge Gateway: this device is equipped with hardware for large-scale IoT edge computing, including a high-performance quad-core ARM Cortex-A8 processor, high-capacity up to 1 GB RAM, and 8 GB storage flash memory. These specifications grant

³ <https://www.socomec.co.uk/en-gb/c/multi-circuit-power-monitoring-system>

the gateway powerful edge computing capabilities, facilitating data optimization, real-time response, agile connectivity, and intelligent analysis at the edge nodes. Figure 4.4 resumes some of the functionalities of the InHand IoT Edge Gateway. To adapt to the highly diversified IoT industry, the IG502 is compatible with major real-time Ethernet and industrial fieldbus protocols, including Modbus TCP and Modbus RTU, as well as MQTT, making it widely applicable in the industrial IoT. Furthermore, the device features a configurable development platform supporting Python code, allowing for the customization of applications. Various versions of the device exist, differing in available RAM, disk space, the presence of digital and analogue inputs, and the option to insert a 4G SIM card for communication in the absence of a wired network.



Figure 4.4 – InHand IoT Edge Gateway IG502⁴

- **Intesis 700 Series Air Gateway:** it is specially designed to allow bidirectional control and monitoring of multi-brand HVAC systems from a BMS, SCADA, PLC, or any other device working as a client based on building automation protocols, such as Modbus TCP, Modbus RTU, KNX, BACnet/IP, BACnet MS/TP, or Home Automation. The gateway acts as a server device of the installation itself, accessing all signals from each air conditioner unit, both indoor and outdoor, and controlling the whole Air Conditioning (AC) network. Furthermore, it integrates a Power Estimation Algorithm (PEA), developed in collaboration with the leading AC brands, to calculate the energy consumption of each indoor unit in a VRF installation based on the total electrical energy consumption

⁴ <https://www.inhand.com/en/products/iot-edge-gateways/>

of the outdoor unit. To ensure proper functioning, it must be coupled with an energy meter, which can only be of two types: a Modbus TCP meter or a pulse meter. In Figure 4.5 a diagram of the Intesis Gateway for Panasonic AC to Modbus system is shown.




Figure 4.5 – Intesis 700 Series Air Gateway for Panasonic to Modbus systems⁵

4.1.2 Data Acquisition, Storage, and Visualization Tools

As anticipated in the chapter’s introduction, while the case studies presented herein differ in their specific data acquisition methods at the edge level, they adopt a unified software architecture for data historicization and visualization. This architecture, studied, defined and developed in the context of this thesis, relies on a stack of open-source applications, chosen for their reliability, scalability, and widespread adoption in industrial IoT environments. Each component serves a specific function within the data pipeline, bridging the gap from edge-level data transmission to higher-level visualization and analytics.

At the core of this architecture, acting universally across all implementations, InfluxDB⁶ was chosen as the primary database solution for the historicization of the acquired measurements. InfluxDB is a specialized timeseries database engineered specifically to handle the high write throughput and query loads associated with timestamped sensor data. It organizes data using a combination of indexed metadata, referred to as tags (such as meter ID or location), and the actual measured values, known as fields (like voltage or power). Its primary advantage is its good write and read performance for sequential, time-centric data, which outperforms standard relational databases for this specific use case. Furthermore, it natively features data retention policies and continuous queries. These are

⁵ HMS Industrial Networks, Leaflet 700 Series Air gateways, 2026, available at <https://hmsnetworks.blob.core.windows.net/nlw/docs/default-source/products/intesis/brochures-and-datasheets/n-a/leaflet-700series-air-gateways.pdf>

⁶ <https://www.influxdata.com/>

essential for automatic downsampling, a necessary process for aggregating high-frequency energy profiles for long-term historical storage without exhausting server disk space. However, due to its specialized nature, InfluxDB is not well-suited for handling highly relational data or performing complex joins. Additionally, poorly designed data schemas that result in high series cardinality can lead to excessive memory consumption and degraded query performance.

Coupled with the database is Grafana⁷, which serves as the top-level interface for data visualization and basic analytics. Operating as an observability platform, Grafana connects directly to InfluxDB as a data source and executes queries to populate configurable, real-time dashboards. It effectively translates raw time-series data into actionable insights through various graphical representations, including line charts, gauges, heatmaps, and tables. Among its strengths, Grafana offers deep, native integration with InfluxDB alongside an intuitive user interface for dashboard creation. It also supports robust alerting mechanisms, multi-tenancy, and role-based access control, making it highly versatile for sharing customized energy dashboards with different stakeholders, such as facility managers or members of an Energy Community.

While InfluxDB and Grafana form the invariable core of the software stack, specific data acquisition layers are deployed conditionally, depending on the communication protocols utilized in the field. In scenarios where measurements are transmitted from the edge via the MQTT protocol, the central server requires an intermediate brokering layer to handle the incoming telemetry. For this purpose, the open-source Eclipse Mosquitto⁸ is employed. Mosquitto acts as the central hub of a publish/subscribe messaging model, receiving messages from edge measurement devices (acting as publishers) and routing them to downstream applications (subscribers) based on predefined hierarchical topics. Written in C, Mosquitto is highly lightweight and efficient, demanding a minimal footprint in terms of CPU and RAM. This makes it suitable for the low-bandwidth, high-latency environments typical of sensor networks. It also fully supports MQTT Quality of Service (QoS) levels, guaranteeing reliable message delivery under varying network conditions. The main limitation of its standard open-source version is the lack of out-of-the-box clustering capabilities. While perfectly adequate for the scale of the case studies

⁷ <https://grafana.com/>

⁸ <https://mosquitto.org/>

presented in this work, achieving high availability and massive horizontal scalability in larger enterprise deployments typically requires additional configuration or the adoption of enterprise-grade alternatives.

Finally, to systematically extract the telemetry from the Mosquitto broker and route it into the database, Telegraf⁹ is utilized. Telegraf is a server-based agent specifically designed for metric collection and processing. Within the proposed MQTT architecture, it is configured with a consumer input plugin to subscribe to the broker, parse the incoming payloads, converting them into the line protocol format required by InfluxDB, and utilize an output plugin to systematically write the structured data into the database. Its primary strength lies in its modularity; boasting an ecosystem of over 400 community-tested plugins, it can seamlessly interface with nearly any data source or destination. It features a very low memory footprint and can perform fast, on-the-fly data formatting and filtering. Nevertheless, managing its heavily file-based configuration across a highly distributed architecture can become cumbersome without dedicated orchestration tools. Furthermore, while it can perform basic data processing, Telegraf is not intended for complex, stateful data transformations, leaving such operations to more suitable data processing frameworks.

To synthesize the proposed software architecture, Table 4.1 outlines the three primary layers of the system, (Data Acquisition, Storage and Visualization) alongside their respective components and key technical characteristics.

Overall, the adoption of this open-source software stack yields significant practical benefits that extend beyond the technical capabilities of the individual components. Economically, it eliminates software licensing fees while supporting standard communication protocols, which enables the integration of existing field meters and avoids the capital expenditure of deploying new hardware infrastructure. Technically, the native interoperability among the selected tools minimizes the need for custom middleware, accelerating deployment and shifting the focus toward core data analysis and energy optimization. Finally, the architecture ensures high accessibility: Grafana's role-based web interface facilitates secure data sharing, delivering intuitive and tailored information directly to both technical operators and non-expert end-users.

⁹ <https://www.influxdata.com/time-series-platform/telegraf/>

Table 4.1 – Layers and components of the proposed software architecture

<i>Architectural Layer</i>	<i>Component</i>	<i>Primary Function</i>	<i>Key Characteristics</i>
Data Acquisition & Processing	Eclipse Mosquitto	MQTT Broker	Lightweight, low resource footprint; supports MQTT QoS for reliable delivery in high-latency sensor networks.
	Telegraf	Metric Collection Agent	Highly modular; handles on-the-fly payload parsing and format conversion.
Data Storage (Database)	InfluxDB	Time-Series Database	High write-throughput optimization; native data retention policies and continuous queries for efficient downsampling.
Visualization	Grafana	Observability Platform	Real-time dashboarding; native integration with InfluxDB; robust alerting mechanisms.

4.2 The ANTES Project: Energy Impact Verification in Greenhouse Cultivation

The ANTES project¹⁰, successfully funded under the cross-border cooperation program INTERREG-ALCOTRA 2014-2020, aimed to capitalize on the outcomes and activities developed in two predecessor projects. By bringing together previous partners and exploiting structural synergies, ANTES seeks to expand knowledge and development opportunities within two specific agricultural supply chains: aromatic plants and edible flowers. These chains share several critical features, including sustainable cultivation practices, specialized post-harvest treatments, innovative packaging techniques, and high nutritional value linked to human health. Specifically, the core objectives of the ANTES project include the creation of a scalable model for emerging agri-food supply chains, the promotion of circular economy paradigms (e.g., replacing conventional packaging with compostable alternatives), and the enhancement of overall production sustainability. Within this framework, a crucial aspect is the

¹⁰ <https://www.interregantea.eu/about-us/antes-project>

precise assessment of the energy impact associated with protected cultivation. In response to the recent international energy crisis, which saw electricity prices surge from approximately 60 €/MWh in late 2020 to peaks exceeding 500 €/MWh in the summer of 2022, the project dedicated targeted efforts to monitoring and optimizing energy consumption in greenhouses.

To empirically assess the energy footprint of greenhouse cultivation, experimental trials on selected species of edible flowers and aromatic plants were conducted at the CERSAA¹¹ facilities. The primary goal was to correlate plant growth (biomass production) with electrical energy consumption, providing a quantitative basis for energy efficiency improvements. The experimental site consisted of a greenhouse, equipped with four main controllable and energy-intensive assets:

- A geothermal heat pump connected to seven geothermal probes, as shown in Figure 4.6, responsible for the climate control (heating and cooling) of the cultivation environment.



Figure 4.6 – Geothermal heat pump and geothermal probes in the greenhouse

- Basal heating systems (the heating mats depicted in Figure 4.7) placed directly on the cultivation benches to maintain optimal root temperatures.



Figure 4.7 – Heating mats on the cultivation benches

¹¹ <https://www.cersaa.it/>

- Artificial LED lighting (the three lines of 200 W Valoya lamps with AP673L spectrum in Figure 4.8).



Figure 4.8 – LED lighting in the experimental greenhouse

- Automated louvers and roof windows for ventilation and air exchange, as the details in Figure 4.9 show.




Figure 4.9 – Details of automated louvers and roof windows of the experimental greenhouse

To accurately track the energy absorption of these systems, a cloud-based real-time monitoring architecture, whose scheme is illustrated in Figure 4.10, was designed and deployed.

The architecture is structured into three main layers:

- Hardware layer (edge): IoT-enabled energy measurement devices (Shelly 3EM configured for single-phase measurements) were installed directly on the electrical cables powering the different greenhouse assets. These devices continuously measure current, voltage, frequency, active power, and total energy. Connected to the internet via Wi-Fi, they securely transmit the telemetry data to a central server using the MQTT protocol.



Figure 4.10 – Scheme of the ANTES monitoring system architecture

- Cloud/Data layer: a virtual machine hosts the core software components, including an MQTT broker (Eclipse Mosquitto) to receive field messages, and a time-series database (InfluxDB) where data undergo processing and are historicized.
- Visualization layer: a web-based graphical interface (Grafana) was implemented to query the database and present the data through a series of customized dashboards. This layer provides stakeholders with real-time insights, as shown in Figure 4.11, allowing them to identify inefficiencies and perform economic and energetic evaluations.



Figure 4.11 – Screenshot of the dashboard with general energy consumption data of the experimental greenhouse

The cultivation trials took place between December 2022 and February 2023, focusing on three representative species: viola (*Viola x wittrockiana*), various cultivars of basil (*Ocimum basilicum*), and salvia (*Salvia discolor*). Growth metrics such as the number and diameter of flowers, pruning weight, and plant height were meticulously recorded to quantify biomass production. These agronomic and growth parameters were monitored and provided by a partner research team within the project framework, serving as external validation data for the energy performance analyses presented in this work.

Energy consumption was continuously tracked over the winter monitoring period. The total energy absorbed over 66 operative days amounted to 1,677.5 kWh. A detailed load analysis revealed that the artificial lighting was rarely utilized during this specific period. Consequently, the total consumption was essentially divided between the basal heating mats (accounting for 51% of the total load) and the geothermal heat pump (49%).

Data analysis highlighted specific operational patterns: while at least one line of the heating mats operated continuously (averaging a daily consumption of 14 kWh), the geothermal heat pump exhibited a variable power draw ranging from 2 kW to 7 kW. Furthermore, a correlation analysis between the heat pump's energy consumption and external environmental conditions revealed a negative correlation coefficient (-0.56), quantitatively demonstrating the expected decrease in electrical absorption as external temperatures rose.

The most significant outcome of this study was the cross-referencing of crop development data with energy telemetry to establish specific indicators. These KPIs successfully mapped the energy cost required to produce a single unit of biological product:

- Salvia: cultivation required approximately 500 Wh per gram of useful biomass, or 7.5 kWh per plant.
- Viola: production necessitated roughly 250 Wh per gram, translating to 17.5 kWh per plant.
- Basil: depending on the specific cultivar and the developmental stage of the seedling, energy requirements varied significantly, ranging from 6 kWh to 40 kWh per centimetre of growth.

By establishing these direct correlations between energy input and biomass output, the implemented monitoring architecture proved to be an invaluable tool to clearly assess the energy impact of greenhouse cultivation and then the economic viability under high energy prices.

4.3 Energy Management and Advanced Control in Healthcare Facilities

Healthcare facilities and associated hospital infrastructures are characterized by high structural complexity and continuous, energy-intensive operations. The implementation of advanced EMS in such environments is critical to enable energy-savings opportunities, optimize the operation of technological plants, and reduce the overall environmental footprint without compromising occupant comfort and patient care. This section presents two distinct applications of cloud-based monitoring and control architectures deployed in healthcare contexts: a data-driven analytical framework for some “Don Orione”¹² facilities in Genoa, and an advanced control system for the “Casa Suore” building at the Humanitas Gradenigo¹³ hospital in Turin.

4.3.1 Data-driven energy analysis: the “Don Orione” case study

The first healthcare case study involves the “Piccolo Cottolengo Genovese” managed by the Don Orione institute, focusing on three facilities in Genoa: Paverano, Castagna, and Camaldoli. The primary objective of this project was to provide a technological support service for the definition of energy management policies, aimed at characterizing the consumption baselines of energy-intensive assets and identifying structural efficiency opportunities.

From an architectural perspective, the deployed solution relies on the extensive utilization of the existing Building Management System (BMS), specifically the Siemens Desigo platform, to gather field data. Rather than installing new edge measurement hardware, the architecture extracts data directly from the BMS. A customized software pipeline was developed to periodically handle the periodic export of telemetry related to electrical energy, natural gas, and water consumption (comprising over 70 distinct measurement points with a 15-minute sampling granularity). Specifically, the data collected over the previous day are automatically transferred once a day in csv format to an FTP server hosted on a dedicated cloud-based analytical platform, provided as a Software as a Service (SaaS). Subsequently, a Python routine executes the parsing and processing of the acquired files, systematically loading the structured data into the time-series database. Consistent with the technological stack presented in the previous sections, this platform leverages InfluxDB for time-series data

¹² <https://www.donorione-genova.it/>

¹³ <https://www.gradenigo.it/>

historicization and Grafana for the Man-Machine Interface (MMI), enabling both off-line systematic analysis and near-real-time monitoring via customizable dashboards, of which an example is shown in Figure 4.12.




Figure 4.12 – Screenshot of the homepage dashboard of the Paverano facility

The continuous ingestion of structured data into the cloud platform lays the foundation for comprehensive energy audits and the systematic computation of specific KPIs. Specifically, the implemented monitoring architecture enables several advanced analytical methodologies:

- **Load profiling and baseline definition:** the system allows for the dynamic establishment of hourly and daily load profiles for each facility, emphasizing the variance between daytime/nighttime operations and weekday/weekend patterns. It is also possible to characterize significant daytime peaks or extract day/night electrical consumption. As an example, Figure 4.13 highlights a regular consumption pattern characterized by a steady nighttime baseline and a daytime energy usage that varies between weekdays and weekends, showing no correlation with the external temperature.
- **Weather normalization and correlation:** to properly evaluate HVAC efficiency, the architecture enables the correlation of energy demand with external temperatures. Linear regression models formalize the negative correlation for heating (e.g., thermal power plants) and the positive correlation for cooling, serving as a baseline for high-level energy auditing. Figure 4.14 illustrates the correlation between energy consumption of one of the buildings' thermal power plant and the mean external temperature from September to December 2023. While the

regression line confirms the decreasing trend as temperature rises, the real data distribution highlights a step-controlled operational logic, characterized by four distinct capacity levels. Although this discrete behaviour results in a lower coefficient of determination for a continuous linear fit, such a representation allows for the immediate identification of operational anomalies or baseload inefficiencies.




Figure 4.13 – Day/night electrical consumption and daily average temperature




Figure 4.14 – Correlation between daily electricity consumption of the heating system and average daily outdoor temperature

- ABC Analysis: the granular data collection supports the execution of ABC analyses on the monitored assets to strategically prioritize

energy-saving interventions. This methodology can highlight whether a minor fraction of the equipment drives the vast majority of the consumption. Figure 4.15 provides an empirical example of this: the analysis reveals that, during the heating season, the thermal power plant alone accounts for the largest share of the facility’s consumption, perfectly isolating the most critical asset for targeted interventions




Figure 4.15 – ABC curve of Paverano facility

- **Heatmap Visualization:** the continuous data stream allows for the generation of temporal heatmaps, which prove highly effective in visually pinpointing anomalous behaviours, such as unexpected night-time operations or the continuous running of air conditioning units in specific zones. Ultimately, these visualization tools provide the technical staff with direct, actionable insights to optimize the manual scheduling of the technological plants. As an example, Figure 4.16 displays the quarter-hourly consumption of a thermal power plant. This carpet plot perfectly outlines the daily programmed schedule, active from approximately 06:00 to 22:00, and the exact onset of the heating season. Furthermore, it immediately highlights a distinct operational deviation, showing a significant daytime load reduction during the first week of November, demonstrating the tool’s effectiveness for schedule verification.




Figure 4.16 – Heatmap of thermal power plant 15-minutes electricity consumption

4.3.2 Towards Optimized Control: the Humanitas “Casa Suore” case study

While the Don Orione project focused on data analysis to suggest potential efficiency intervention to the management, the project developed for the "Casa Suore" office building, at the Humanitas Gradenigo hospital in Turin, represents a step towards fully automated, intelligent regulation. The building's energy footprint is heavily dominated by its HVAC system, a complex Panasonic heat pump system operating on a three-way Variable Refrigerant Flow (VRF) logic, which accounts for an estimated 60% of the facility's total electrical consumption. The system comprises 30 indoor units distributed across 27 rooms over 5 floors. These indoor units are served by a single outdoor unit featuring a nominal cooling capacity of 85 kW and a heating capacity of 95 kW, at an electrical power input of approximately 23 kW.

Given the limitations of standard manual scheduling via local controllers, the objective of this project is to overlay an advanced control architecture capable of dynamically calculating and enforcing the optimal temperature setpoints for each indoor unit (or thermal zone) throughout the day. The goal is to maximize the operational efficiency of the HVAC system in response to varying environmental conditions, strictly within the boundaries of user

thermal comfort, targeting a significant reduction in the baseline HVAC energy consumption.

To interface the proprietary HVAC network with external optimization algorithms, a dedicated edge-to-cloud hardware and software architecture was designed, of which a scheme is shown in Figure 4.17.




Figure 4.17 – Scheme of the “Casa Suore” monitoring system architecture

At the edge layer, an Intesis gateway (specifically designed for HVAC integration) was installed to interface directly with the proprietary Panasonic P-Link/S-Link communication bus, making the signals of real-time operational parameters available in Modbus TCP. The controllable parameters collected from each indoor unit are:

- the on/off status
- the operating mode,
- the internal ambient temperature T_{amb}
- the temperature setpoint SP
- the vane position
- the fan speed

Furthermore, to accurately track the power consumption, a Socomec A-40 smart energy meter was installed on the outdoor VRF unit, whose active/reactive power and energy, voltage, current, and frequency are measured, and was connected to the gateway, which then proportionally distributes this total energy consumption to the respective indoor units based on their operational state.

At the cloud layer, on a server under the same local area network of the gateway, a specific service was deployed to collect the telemetry, acting as a Modbus client, and store it in a time-series database (InfluxDB). Then, using Grafana, the data thus collected are made easily accessible via browser. Crucially, the system also ingests actual and forecast external weather data (temperature, humidity, and cloudiness) via API from external providers.

Since, at the time of drafting this thesis, the project has been recently commissioned, the control logic is still in a transitional learning phase and solid operational results cannot yet be discussed. Initially, the system operates using an empirical, rule-based logic, which treats each room as an independent thermal zone by neglecting any heat transfer with adjacent spaces, to gather a sufficient dataset of thermal transient behaviours. As illustrated in Figure 4.18, this logic dynamically modulates the HVAC state and thermal deadband (T_{min} , T_{max}) by evaluating real-time temporal and environmental data and occupancy, which follows a predefined schedule. To optimize energy conservation, the algorithm relaxes the comfort constraints during unoccupied periods and tightens them for predictive pre-conditioning as occupancy approaches. Every $t = 15$ min, the operating mode is determined by baseline T_{amb} , while the comfort boundaries are linearly adjusted based on the average outdoor temperature T_{ext} . Before actuation, these adjusted thresholds are cross-validated against the room's calculated thermal inertia ($\Delta T_{amb}/\Delta t$) to mitigate temperature. Ultimately, binary On/Off commands are dispatched based on the real-time T_{amb} .

Although this preliminary rule-based control logic could theoretically be executed directly at the edge layer (on the gateway), the deployment of a cloud-based architecture remains necessary to enable long-term measurement historization, ensure real-time data visualization through advanced dashboards, and provide the high-performance computing environment required to support the deployment of computationally demanding optimization algorithms.

Indeed, the historical dataset generated during this transitional phase will subsequently be used to train a data-driven control algorithm, based on MPC or Reinforcement Learning (RL) logics.

In the MPC scenario, the future evolution of T_{amb} would be predicted based on the current state and external disturbances:

$$T_{amb,i} = f(SP_i, T_{amb,i-1}, T_{ext,i}, CC_i, d_i, h_i) \quad (4.1)$$

where CC_i is the cloudiness, d_i represents the day type (weekday/holiday) and h_i the hour of the day.




Figure 4.18 – Rule-based policy flowchart for HVAC control in “Casa Suore”

By simulating the thermal response of the building over a 24-hour horizon, the optimizer would dynamically evaluate the optimal setpoint trajectory (SP_{opt}) that minimizes the total energy consumption while strictly adhering to the comfort constraints required during working hours. Thermal comfort would be evaluated employing the PMV-PPD model (the predicted mean vote and the predicted percentage dissatisfied index, respectively) recommended by ISO 7730 and the constraint would then be enforced by requiring the PPD index, i.e. the average percentage of occupants dissatisfied with the thermal environment, to remain below 20%. Also, the direct correlation between the applied setpoints and the resulting energy consumption of the HVAC assets would be modelled with a data-driven approach.

Alternatively, in the Reinforcement Learning scenario, an RL agent would learn optimal control policies through continuous interaction with the building's thermal environment rather than solving a predefined optimization problem. In this framework, the agent would regulate the indoor temperature setpoints (SP_i) by aiming to maximize a cumulative reward function. The reward logic would be specifically designed to assign strict penalties for any deviation from the predefined thermal comfort boundaries during working hours, while actively assigning positive rewards for control actions that minimize the overall HVAC energy consumption.

Regardless of the algorithmic path ultimately prioritizing execution, the optimizer will dynamically resolve the delicate balance between maximum thermal comfort and optimum energy efficiency. The calculated setpoints will then be transmitted back to the edge layer and written directly to the HVAC controllers via the Intesis gateway, thereby effectively closing the automated control loop.

4.4 Energy Management and Optimization in a Zero Energy Building

The transition towards decentralized and sustainable energy paradigms requires the development of infrastructures capable of actively participating in the energy market. In this context, the third case study focuses on an experimental Zero Energy Building developed within the framework of the Italian PNRR, specifically, the RAISE¹⁴ (Robotics and AI for Socio-

¹⁴ <https://www.raiseliguria.it/en/spoke-3-raise/>

economic Empowerment) ecosystem, Spoke 3, dedicated to sustainable environmental caring and protection technologies.

The ZEB laboratory, whose plan is shown in Figure 4.19, is located near Milan and consists of an existing single-storey structure that simulates an office building. The external dimensions are about 7 x 8 x 4 m (length x width x height), with 2 windows on the South-East facade and 1 window on the North-West facade. The laboratory is internally divided into three spaces: two are intended for office use (A1 and A2) and the small one (A3) is dedicated for the installation of plants. The envelope and plant systems have been designed in order to achieve the ZEB target, with insulated roof, floor and façades.



Figure 4.19 – Plan of the laboratory and external view of the laboratory

The HVAC system consists of a full inverter monobloc air-to-water (A2W) heat pump for heating, cooling and the production of domestic hot water (DHW), operating with outdoor temperatures down to $-20\text{ }^{\circ}\text{C}$ and producing hot water up to $65\text{ }^{\circ}\text{C}$. It is connected to different emission systems for the two rooms: to a fan-coil unit and to a radiant floor system for A1 and to a radiant floor system only for A2. The hydraulic circuit includes two water tanks: one serving the heating and cooling system and another dedicated to DHW production. Furthermore, the ZEB laboratory is equipped with a solar panel for DHW production and a 3 kW peak power PV panels for electricity production, connected to a 5 kWh lithium battery placed in the technical room of the ZEB laboratory through a dedicated inverter. The panels are installed on the rooftop.

This experimental ZEB is designed as an active prosumer that can interact dynamically with the electrical grid. The synergistic operation of the integrated local renewable energy generation, coupled with highly flexible

assets, which specifically include the battery the heat pump, allows the building to achieve a nearly zero net energy footprint. Furthermore, it empowers the facility to function as a flexible node within broader energy sharing schemes, such as Renewable Energy Communities, where the temporal synchronization of generation and consumption acquires significant economic and environmental value.

4.4.1 Monitoring and Control Architecture

To fully exploit the flexibility of the ZEB and support the implementation of advanced energy management policies, a comprehensive edge-to-cloud monitoring architecture was deployed. Consistent with the technological approaches utilized in the previous case studies, this infrastructure is designed to acquire, aggregate, and historicize heterogeneous data streams efficiently, creating a reliable foundation for decision-making algorithms.

Figure 4.20 shows a graphic representation of the deployed architecture.




Figure 4.20 – Monitoring and control system architecture of the ZEB

At the field layer, the central element of the architecture is an InHand IoT Edge Gateway, which is responsible for interfacing a distributed network of smart meters and environmental sensors, among which the following instruments:

- Socomec Digiware multi-point energy meters to measure energy consumption of the heat pump, the recuperator of the HVAC system, of lights and pumps and, then, to send the measurements to the gateway via Modbus TCP;
- the Huawei inverter¹⁵, through a specific device for Modbus TCP communication (Huawei Smart Dongle WLAN-FE¹⁶), to read the PV energy production values and battery operating parameters and to write the appropriate registers to operate the inverter and battery control.
- the Eneren heat pump¹⁷, through its native controller, and the A1 and A2 room controllers, through the Newtohm digital regulators, model DDC-mPID3¹⁸, connected to the gateway with a serial bus, to read heat pump operating values and current ambient conditions of the rooms and to write the appropriate registers to operate room temperature control, via Modbus RTU protocol.

The gateway securely transmits the data collected from the field via MQTT protocol to the central system, which has been installed on a local server and where a MQTT broker has been configured.

In the same local server, which represents the cloud layer, the continuous stream of data is ingested by a time-series database (InfluxDB). Concurrently, a visualization layer built on Grafana provides operators with real-time, highly customizable dashboards for performance monitoring and structural auditing. This robust data pipeline is essential for supplying accurate, real-time boundary conditions to the downstream predictive models, needed for the implementation of advanced energy management policies.

4.4.2 Optimization scenarios

Indeed, the ultimate purpose of the deployed ZEB monitoring infrastructure is to enable the execution of MPC strategies for optimal energy dispatch. Unlike heuristic or rule-based controllers, the implemented MPC approach relies on data-driven models, such as Artificial Neural Networks (ANN), to forecast the PV generation and the building's electrical load demand and on a physical model to predict the thermal response of the building over a 24 hours horizon. At each control step, leveraging these

¹⁵<https://solar.huawei.com/en/products/SUN2000-3-4-5-6KTL-L1/>

¹⁶<https://solar.huawei.com/en/products/SmartDongle-WLAN-FE/>

¹⁷<https://eneren.it/en/products/naw/>

¹⁸<https://www.newtohm.it/mpid3/>

predictive inputs and the real-time telemetry acquired by the monitoring system, the MPC formulates a constrained mathematical optimization problem, based on MILP formulation and characterized by the following objectives in two different scenarios:

1. Optimization considering the ZEB individually: the objective is to minimize purchasing costs of the single building and, if possible, to maximize profits from sales, through optimized control of the battery and the HVAC system and considering constraints related to building thermal comfort.
2. Optimization based on incentives related to RECs: in this case, the ZEB is considered to be a member of a REC, alongside three other houses whose electric loads are simulated; the objective is to maximize profits and then the incentives for collective self-consumption.

For each scenario, two different models were defined, for a total of four models, as specified below:

Model 1: ZEB in the first scenario with supervisory control of the battery and no supervisory control of HVAC, which is regarded as an exogenous input, as illustrated in Figure 4.21;




Figure 4.21 – Schematic representation of the Model 1 ZEB architecture

Model 2: ZEB in the second scenario with supervisory control of the battery and no supervisory control of HVAC, which is regarded as an exogenous input;

Model 3: ZEB in the first scenario with supervisory control of both battery and HVAC;

Model 4: ZEB in the second scenario with supervisory control of both battery and HVAC, as illustrated in Figure 4.22.




Figure 4.22 – Schematic representation of the Model 4 ZEB architecture in the REC scenario

In the respective model, by solving the optimization problem, the algorithm determines the optimal operational trajectory for the dispatchable assets. Specifically, the optimizer computes the most advantageous charging and discharging cycles for the BESS and, under models 3 and 4, modulates the power absorption of the heat pumps, ensuring that the self-consumption of locally generated renewable energy (first scenario) or the collective self-consumption (second scenario) is maximized. Finally, the calculated optimal setpoints are transmitted back to the edge layer for physical execution, guaranteeing that the ZEB operates at maximum economic efficiency while satisfying all structural limits and internal comfort constraints.

However, it should be noted that the experimental testing and validation of this optimization architecture are currently ongoing at the time of drafting this dissertation. Consequently, the empirical results are not reported in detail in this section, and they will be the subject of a forthcoming dedicated scientific publication.

4.4.3 Formalization of the optimization problem

For the implementation of the supervisory control, the objective function is formulated to minimize the net operating costs over the optimization horizon. This formulation is strategically designed to incentivize internal energy sharing within the REC while ensuring the economic valorisation of the renewable energy injected into the grid. To enable the evaluation of the different models, a generalized minimization problem is defined as follows by eq. (4.2):

$$Y = \sum_{t=1}^T (c_{net,in_t} \cdot P_{net,in_t} - c_{net,out_t} \cdot P_{net,out_t} - c_{inc_t} \cdot P_{sh_t}) \cdot \Delta t \quad (4.2)$$

where:

- t is the time step and T is a number of time intervals constituting the optimization horizon.
- $c_{net,in}$ is the cost of energy purchased [€/kWh]
- $c_{net,out}$ is the remuneration of energy sold [€/kWh]
- $P_{net,in}$ is the imported power from the grid [kW]
- $P_{net,out}$ exported power to the grid [kW]
- P_{sh} is the shared power [kW]
- c_{inc} is the value of incentive for shared energy [€/kWh]

For $t=1, \dots, T$, the optimization is subject to the energy balance equation (4.3):

$$P_{pv_t} + P_{st,out_t} + P_{net,in_t} = P_{load_t} + P_{st,in_t} + P_{net,out_t} \quad (4.3)$$

where:

- P_{pv} is the photovoltaic power [kW]
- $P_{st,in}$ is the battery charging power [kW]
- $P_{st,out}$ is the battery discharging power [kW]
- P_{load} is the total electrical demand [kW], as defined in (4.4):

$$P_{load_t} = P_{hvac_t} + P_{unc,load_t} + P_{v,load_t} \quad (4.4)$$

where:

- P_{hvac} is the HVAC load [kW]
- $P_{unc,load}$ is the uncontrolled ZEB load [kW]
- $P_{v,load}$ is the virtual load [kW] of the three simulated houses

Thus, the shared power P_{sh} is defined by eq. (4.5):

$$P_{sh_t} = \min(P_{pv_t} + P_{st,out_t}, P_{load_t} + P_{st,in_t}) \quad (4.5)$$

The decision variables that mathematically determine (4.2), subject to (4.3), are P_{net,in_t} , P_{net,out_t} , P_{sh_t} , P_{st,in_t} and P_{st,out_t} and the absence of virtual EC demand in models 1 and 3 is imposed by setting $P_{v,load_t} = 0$ and $P_{sh_t} = 0$. In models 3 and 4 also P_{hvac_t} becomes a decision variable.

Furthermore, two sets of binary variables, denoted by z_{net} and w_{st} , that respectively prevent simultaneous import and export at the grid connection, and simultaneous charging and discharging of the battery are added.

The following constraints must be met:

- Mutual exclusivity of the exchange with the grid:

$$P_{net,in_t} \leq z_{net_t} \cdot P_{net,in_t}^{max} \quad (4.6)$$

$$P_{net,out_t} \leq (1 - z_{net_t}) \cdot P_{net,out_t}^{max} \quad (4.7)$$

where $z_{net_t} \in [0,1]$ and $z_{net_t} = 0$ indicates no energy import from the grid and P_{net,in_t}^{max} and P_{net,out_t}^{max} are respectively the maximal imported power and the maximal exported power.

- Mutual exclusivity of the exchange with the battery:

$$P_{st,in_t} \leq w_{st_t} \cdot P_{st,in_t}^{max} \quad (4.8)$$

$$P_{st,out_t} \leq (1 - w_{st_t}) \cdot P_{st,out_t}^{max} \quad (4.9)$$

where $w_{st_t} \in [0,1]$ and $w_{st_t} = 0$ indicates that the battery is not being charged and P_{st,in_t}^{max} and P_{st,out_t}^{max} are respectively the maximal battery charging power and the maximal battery discharging power.

- The shared energy is either set to zero (models 1 and 3) or it's determined imposing constraints (4.10) and (4.11) (in models 2 and 4).

$$P_{sh_t} \leq P_{pv_t} + P_{st,out_t} \quad (4.10)$$

$$P_{sh_t} \leq P_{load_t} + P_{st,in_t} \quad (4.11)$$

The specification of the objective function (4.2) is such that P_{sh_t} is pushed towards its upper bounds, so that its optimized value will meet (4.5).

For every sequence of optimization variables, the battery state of charge SoC is modelled dynamically by eq. (4.12).

$$SoC_t = SoC_{t-1} + \frac{100 \cdot \eta_{in} \cdot \Delta t}{E_{st}^{nom}} P_{st,int} - \frac{100 \cdot \eta_{out} \cdot \Delta t}{E_{st}^{nom}} P_{st,out} \quad (4.12)$$

where:

- η_{in} and η_{out} are respectively the battery charge and discharge efficiency
- E_{st}^{nom} is the nominal battery capacity [kWh]

The state of charge is bounded, $SoC^{min} \leq SoC_t \leq SoC^{max}$, and the terminal condition of eq. (4.13) is imposed to keep the final and initial states close to each other.

$$-e \leq SoC_T - SoC_0 \leq e \quad (4.13)$$

with e representing a predefined tolerance margin.

This constraint prevents the optimizer from obtaining artificially favourable results by ending the current horizon with a depleted or overcharged battery.

Models 3 and 4 use an extended formulation in which the HVAC operation is jointly optimized with battery dispatch.

In this case, the controller determines not only the electrical flows but also the HVAC set-point trajectory T_{sp} , constrained by eq. (4.14) within both admissible (T_{sp}^{min} and T_{sp}^{max}) and thermal comfort bounds (T_{comf}^{min} and T_{comf}^{max}), and the corresponding thermal power delivered by the system.

$$\max(T_{sp}^{min}, T_{comf}^{min}) \leq T_{sp_t} \leq \min(T_{sp}^{max}, T_{comf}^{max}) \quad (4.14)$$

To include the set-point trajectory in the MILP problem, we obtained a linear state-space representation of an RC model as in eq. (4.15).

$$\mathbf{x}_{t+1} = \mathbf{A} \times \mathbf{x}_t + \mathbf{B} \times \mathbf{u}_t \quad (4.15)$$

where the state vector $\mathbf{x}$ collects the nodes temperature, i.e. the room air temperature and the internal and external surface temperatures of each wall, and the vector $\mathbf{u}$ collects independent inputs, such as those shown in eq. (4.16).

$$\mathbf{u} = [T_{ext}, q_{s,1}, \dots, q_{s,h}, q_{hvac}] \quad (4.16)$$

where $q_{s,1}, \dots, q_{s,h}$ represent the solar gains transposed on each combination of surface orientation and tilt angle.

The state matrix A and input matrix B represent the discrete-time dynamics derived from an RC model, whose thermal properties were identified using empirical data.

The thermal power produced by the HVAC system q_{hvac} at time t is linked to the difference between the set-point and the indoor temperature through the linear relation in eq. (4.17).

$$q_{hvac_t} = \frac{\gamma \cdot (T_{sp_t} - x_{1_t}) + offset_{hvac}}{1000} \quad (4.17)$$

where x_{1_t} is the internal air temperature in the discretized version of the state space model of (4.15), γ is a conversion coefficient and $offset_{hvac}$ is a constant offset.

Minimum and maximum nominal bounds q_{hvac}^{min} and q_{hvac}^{max} are also set to the HVAC thermal power through constraint (4.18).

$$q_{hvac}^{min} \leq q_{hvac_t} \leq q_{hvac}^{max} \quad (4.18)$$

The corresponding electrical power contributing to (4.4) is obtained through a coefficient of performance (COP) by eq. (4.19), so that the HVAC demand becomes a decision-dependent load rather than an input profile.

$$P_{hvac_t} = q_{hvac_t} / COP \quad (4.19)$$

Finally, auxiliary slack variables and a discomfort penalty coefficient C_{disc} are also included in the objective function, whose extended form is in eq. (4.20).

$$Y = \sum_{t=1}^T (c_{net,in_t} \cdot P_{net,in_t} - c_{net,out_t} \cdot P_{net,out_t} - c_{inc_t} \cdot P_{sh_t} + c_{disc} \cdot (s_t^+ + s_t^-)) \cdot \Delta t \quad (4.20)$$

where s_t^+ and s_t^- are nonnegative auxiliary variables associated with thermal-comfort management. In particular, they are the positive slack variables of the constraint (4.21).

$$T_{comf}^{min} - s_t^- \leq x_{1_t} \leq T_{comf}^{max} + s_t^+ \quad (4.21)$$

If s_t^+ and s_t^- are 0, it means the internal air temperature is inside the limits of thermal comfort; they are positive when the limits are violated. This ensures that the system can work also when violating the thermal comfort limits, but at the cost of a penalty in the objective function.

Residential Energy Monitoring: Optimisation Model for Demand Flexibility Based on Non-Intrusive Load Disaggregation

Residential users have acquired a central role in the energy transition following the formal establishment of RECs and, to maximise the resulting economic and technical benefits, an optimal management of decentralised energy assets is required.

The preceding chapter demonstrated the efficacy of advanced, edge-to-cloud monitoring architectures across diverse operational environments. Those empirical case studies relied upon the extensive deployment of dedicated sub-metering hardware and the direct integration of existing BMS to achieve necessary data granularity. While this hardware-centric approach is highly effective in commercial, agricultural, and industrial settings, its extrapolation to the residential sector encounters intractable barriers. Installing distributed sensor networks within individual dwellings is economically prohibitive and it is also logistically complex and intrusive for the end-user. Consequently, unlocking the latent demand-side flexibility of residential prosumers requires a shift in the monitoring paradigm.

An optimized EMS fundamentally relies on energy data to orchestrate

DSM strategies, but, given the physical constraints of sub-metering, these granular data must be derived analytically rather than measured directly. The study of energy disaggregation, to detect the specific consumption of the appliances subtended to the main load, is called Non-Intrusive Load Monitoring and NILM techniques are becoming more and more widespread in recent years. Typically, NILM techniques are divided into four stages: data collection, load detection, features extraction and load identification [72]. The quality of the performance achieved by these techniques strongly depends on the sampling frequency of the data collected: higher sampling frequencies allow for better performance, but increase data volume linearly, which can create significant communication burdens due to continuous real-time data streams [73], and hardware costs exponentially [74].

Performance also depends on the approaches used for the load identification, such as supervised or unsupervised machine learning algorithms. Unsupervised algorithms, compared to supervised one, generally show worse performances, but they have the advantage of not requiring labelled data of consumption for the training of the model. This makes them preferable for real-time application [75]. As reported in [72] and [76], many supervised learning techniques have been applied in NILM, such as Artificial Neural Networks, Support Vector Machines (SVM), K-Nearest Neighbours (k-NN). Regarding unsupervised methods, clustering techniques and Hidden Markov Methods (HMM) have been proposed in literature.

Since 2010 residential buildings are becoming an important sector of NILM applications and the number of works involving NILM techniques in this area is growing. Despite this growing interest, real-world applications of NILM methods lack [73] and NILM techniques have primarily been evaluated using publicly available datasets [74], whose sampling rate varies from 1 Hz to 100 kHz [77]. Only few NILM algorithms have been tested on a real data acquisition architecture. In [78] and [79] measurement systems with a high sampling frequency (≥ 1 Hz) are proposed and the identification of the electrical loads is based on a set of measurements available in a database, that is used to classify the detected loads. In [80] energy data are collected at a sampling rate of 1 Hz and analysed by a microcontroller with embedded software and a pretrained model is used to identify the appliance loads. Also [81] and [82] presented results of a real embedded system, where a supervised NILM method using high sampling frequency (≥ 1 Hz) signals is applied. In [83] an unsupervised technique for NILM is developed and applied to public datasets with a measurement sampling time of one minute.

It appears clear that a high sampling frequency (generally ≥ 1 Hz) is fundamental to realize load disaggregation. However, this asks the installation of an appropriate measuring instrument by qualified personnel, thus introducing significant costs and preventing users to adopt this kind of technological solution. One key idea is to use the measurements provided by the already installed DSO's meters, which are experiencing an upgrade in Europe, becoming smart tools able to provide the user with real-time data, only requiring inexpensive plug & play gateways. However, in the specific case of Italy, the available measurements are provided with a regular sampling time of 15 minutes (thus, a frequency of $1/900$ Hz $\ll 1$ Hz) mixed to spot measurements delivered when the total power consumption moves between 300 W intervals over a 1 Hz monitoring.

In this chapter the work done in [84] is presented and the main contribution of it, an extension of [85], is therefore the development and the validation on real data of a novel NILM method based on unsupervised techniques, suitably designed to use the specific measurement provided by the DSO's smart meters in Italy. The final practical result is a lightweight and cost-effective technological architecture that is minimally invasive and simple as possible for a larger real-world application.

Moreover, the proposed architecture is assumed to be applied within the energy management system of a REC to show the potential benefit of using such a system. Indeed, this methodology is ultimately used to define the inputs of an optimization algorithm that aims to determine the optimal load curve of the monitored case study in order to maximize the shared energy within a REC, identifying the best times to use the utilities detected by applying the NILM algorithm.

In Section 5.1 the case study and the data acquisition and monitoring system architecture are presented. The NILM methodology is presented in Section 5.2. Section 5.3 proposes the house power demand optimization method. Results of the NILM algorithm are shown in Section 5.4 and those of the optimization in Section 5.5. Finally, conclusions and future development are summarized in Section 5.6.

5.1 Case-Study

The proposed NILM methodology was applied on three houses of different types, located in Italy, in the Liguria region. The main features are listed in Table 5.1.

Table 5.1 – List of Case Studies

	House 1	House 2	House 3
Location	Genova (GE)	Cogorno (GE)	Genova (GE)
Number of inhabitants	2	6	4
Type	Apartment	Detached house	Apartment
Surrounding context	urban	rural	urban
Number of floors	1	3	1
Available power [W]	3000	3000	4500
Living area [m²]	80	200	140
Main appliances	Dishwasher, oven, washing machine, iron	Dishwasher, oven, washing machine, iron, water pump	Dishwasher, oven, washing machine, iron, electric vehicle

Figure 5.1 shows the system architecture, which is divided into a hardware part and in a software part. The measurements are acquired directly from the II generation (2G) smart meter, already installed in each house, of the Italian DSO using plug & play gateways to be connected directly to a power outlet, thus eliminating the cost of installation by qualified personnel. The 2G smart meter measures the total electrical load of a household. In particular, it measures: the quarter-hourly electric energy in Wh, the average power of every quarter of an hour in W, and the instant power in W at a frequency of one sample per second. However, instant power measurements are not accessible. Indeed, the 2G smart meter can provide data in real-time through the Chain2 feature, that allows sending every 15 minutes electric energy and average and instant power measurements to a gateway, a simple home automation device, via Power-line Communication (PLC). Furthermore, it sends the measurement of the instant power every time its value moves from one power band to another. The power bands are defined every 300 W starting from 0 W, as defined by CEI (Comitato Elettrotecnico Italiano) technical standard [86][87]. Thus, the sampling frequency of the instant power acquired by the gateway is not constant: it is a combination of a 1/900 Hz regular sampling and “event triggered” spot measurements. It is worth remarking that the format of these data has been established by the Italian

energy authority (ARERA) and are therefore the unique that can be acquired in real-time from the 2G smart meters.




Figure 5.1 – System architecture

The gateways are specifically designed for residential and small business users, they can easily be configured via smartphone app, communicate directly with the 2G meter via the Chain2 protocol and can read the energy values both produced and consumed, in order to have efficient monitoring of the entire energy flow. The gateways, after converting the analog PLC signal to digital, connected to the Wi-Fi network, send data to a central server over MQTT protocol. The data sent to the broker are stored in a non-relational timeseries database and are made available to a web-based visualization tool and to the data analysis layer.

The system architecture follows the one presented in the previous chapter:

- Eclipse Mosquitto as MQTT broker;
- Telegraf as data exchange layer between the broker and the database;
- InfluxDB as database;
- Grafana as visualization software;
- Python as code for algorithms and data analysis.

Besides the three houses on which the NILM methodology was applied, a fourth house (House 4), with a rooftop 3.2 kWp PV plant, was added to the monitoring system. This additional house thus provides real measurements of energy fed into the grid necessary for simulating REC scenarios among the monitored users and therefore, in a REC configuration with at least one of the three houses of Table 5.1, allows the sharing of the energy itself. Indeed, House 4 is configured as a prosumer, *i.e.* it not only consumes energy but also produces energy, and excess energy, *i.e.* not self-consumed, is fed into the grid. Since 2G smart meters are bidirectional (they measure both the energy drawn and the energy injected), measurements of energy fed into the grid are acquired with the same architecture described in Figure 5.1.

To evaluate the optimization model cited in the introduction of this chapter, which aims to maximize the shared energy, a REC configuration consisting of House 4 and one of the three other households was considered, as shown in Figure 5.2.




Figure 5.2 – REC configuration to test the optimization model

Since House 4 was added to simulate a REC using real PV energy data, in this work there was no interest in applying the NILM algorithm also to it.

5.2 The NILM Algorithm

The main two objectives of the data analysis layer are to identify the electricity consumption of household appliances from the data gathered by

the proposed acquisition system and to create a database of devices classified according to their electrical load properties. For this purpose, an algorithm in Python has been developed from scratch that automatically recognizes and classifies electrical loads without a priori knowledge about the household devices and avoiding involving users in providing information on the use of appliances, simplifying the process and reducing time and costs due to the intrusive use of specific measuring instruments.

The developed load disaggregation algorithm consists in four main functions, as shown in Figure 5.3 (where the function blocks have a rectangular shape and the outputs have a rhomboid shape), that are subsequently discussed in this section:

- *Events Detection*: it takes raw data stored in the database as input and aims to identify whenever an appliance has run, providing as output the power profile of each identified load (load pattern).
- *Patterns Merging*: it receives the identified load patterns as input, aiming to recognize whether two adjacent events belong to the same appliance operation cycle and, if so, merging them into a single load pattern.
- *Patterns Filtering*: it receives the power profiles of the identified loads, both merged and not, and filters them to take into account only the significant ones.




Figure 5.3 – Flow diagram of the NILM algorithm

- *Patterns Clustering*: it receives the filtered load patterns and clusters them, with the aim of assigning the load profiles of the same appliance to the same cluster and creating a database of appliance power profiles.

5.2.1 Events Detection

The first function, called *Events Detection*, aims to identify the events triggered by all the significant controllable and non-periodic electrical loads and their relative load pattern. Appliances such as lights or the fridge are indeed not taken into account as they have low consumption or are present periodically during the whole day and then they are not useful for DSM purposes. Therefore, these types of electrical loads are those that make up the so-called baseline, *i.e.* the load profile without considering the appliance loads.

The first part of the algorithm takes the instant (without a fixed sampling frequency) and the quarter-hourly average power (P_i , $P_{ave15,i}$) and the quarter-hourly energy ($EE_{ave15,i}$) measurements as input and evaluates for each time sample i whether: a significant load is running; a possible electrical load has started; or a possible electrical load is starting, as indicated in the diagram in Figure 5.4.

The primary outputs of this function consist of the time indexes marking the start and end of each detected load, alongside the power baseline $P_{bl,i}$ for each time sample i , to obtain the disaggregated power profile isolated for each significant load.

In order to identify the presence of an appliance load in the aggregate load profile, a combination of dynamic and adaptive thresholds is applied on:

- the difference between P_i or $P_{ave15,i}$ and $P_{bl,i}$, as defined in (5.1);
- the absolute instant power P_i ;
- the number of measurements acquired by the gateway over time (the sampling frequency), which varies between 1 and 1/900 Hz, as described in Section 5.1.

So the start and end events of each identified load are tracked and the baseline is defined. $P_{bl,i}$ is the same for all samples within the same quarter of an hour and is calculated only at the first sample every 15 minutes. For each time step, in the absence of any significant load, $P_{bl,i}$ is equal to $P_{ave15,i}$ of the closest previous quarter of an hour. Instead, when an appliance load is present, it is equal to the minimum value in the quarter of an hour of the instant power, if lower than a pre-set threshold, or, otherwise, to an estimated

value based on the baseline values of the previous instants and on the time of day:

$$P_{bl,i} = \begin{cases} P_{ave15,i} & \text{no load} \\ \min_{15 \text{ mins}} (P_i) & \text{with load if } \min_{15 \text{ mins}} (P_i) < P_{bl,i-1} + 200 \\ P_{bl,i-1} + \Delta P_t & \text{with load if } \min_{15 \text{ mins}} (P_i) \geq P_{bl,i-1} + 200 \end{cases} \quad (5.1)$$

where ΔP_t is the baseline variation that depends on the time of day and is obtained from the average daily profile of power demand.




Figure 5.4 – Flow diagram of the Events Detection function

Once the baseline has been computed, the disaggregated profiles of the detected appliance loads are obtained. Specifically, for each time window delimited by a tracked pair of start and end events, the load's power profile is extracted by subtracting the baseline outlined above from the instant power measurements. The set L_D of these extracted load patterns serves as the input for the following function.

5.2.2 Patterns Merging

The *Patterns Merging* function has the objective of aggregating those load patterns in L_D , initially detected as separated, that actually constitute a single load pattern. Indeed, some appliances, that have two or more operating cycles, separated by a significant time interval (such as the dishwasher), might be identified by the *Events Detection* function with two individual different load patterns. For this purpose, this function aims to merge any pair of adjacent load patterns consisting of two load patterns with repetitive load profiles and whose mutual time distance is always comparable. Therefore, a clustering algorithm is applied to all the detected load patterns in L_D , in order to assign similar load profiles to the same source of consumption. Moreover, for each pattern, the time distance between its end and the start of the next one is computed. Finally, adjacent load patterns with the same time distance belonging to the same pair of clusters, and that are repeated at least a predefined number of times (*i.e.* 5) in the analysed consumptions, are merged. Therefore, a new set L_M of load patterns, with merged ones, is obtained as output of this function.

For clustering, the K-means algorithm was adopted, which is one of the simplest and most popular clustering algorithms used in the analysis of data acquired by smart meters [88][89]. The K-means algorithm divides a set of N patterns X into K disjoint clusters C , each described by the mean μ_j of the patterns in the cluster, commonly called the cluster centroids. The K-means algorithm aims to choose centroids that minimise the inertia I , a measure of how internally coherent clusters are:

$$I = \sum_j^K \sum_{i \in C_j} \|x_{i,j} - \mu_j\|^2 \quad (5.2)$$

Normally, it requires as input the number K of clusters to divide the load patterns into. Since, in this case, this information is not *a priori* known, an optimization to select the optimal number K is performed. In particular, the clustering algorithm is applied with K varying between 2 and 10 (upper limit

assumed for the number of significant household appliances present in a home) and the clustering performance is evaluated using metrics such as:

- the Silhouette coefficient S [90]: it is a measure of how similar an object is to its own cluster (cohesion) compared to other clusters (separation) and it is bounded between -1 for incorrect clustering and +1 for highly dense clustering, so the higher the better. For a set of samples it is given as the mean of the Silhouette coefficient for each sample.
- the Calinski-Harabasz index CH [91]: a higher score relates to a model with better-defined clusters and it is defined as the ratio of the sum of between-clusters dispersion and of within-cluster dispersion for all clusters.
- the Davies-Bouldin index DB [92]: it compares the distance between clusters with the size of the clusters themselves; a lower value relates to a model with better separation between the clusters.

The optimal number of clusters K is the one involving clustering with the maximum decision score, defined as a combination of the three cited metrics:

$$decision\ score = (S \times CH)/DB \quad (5.3)$$

Furthermore, in order to apply the clustering algorithm, each load pattern is identified by its load profile with a vector of equal size. Therefore, each power profile over time is normalized on its duration and as many dummy samplings as necessary are added to make the size equal to that of the load pattern with the highest number of samplings; the holes between two real measurements are filled by propagating the first value forward to the next. Consequently, this procedure yields load patterns defined by fictitious sampling frequencies that, while varying from one profile to another, result in a time resolution significantly higher than the minimum of 15 minutes.

The K-means algorithm has been implemented using the scikit-learn library, an open-source machine-learning library that supports supervised and unsupervised learning [93].

5.2.3 Patterns Filtering

After obtaining the merged load patterns, the third function, called *Patterns Filtering*, is performed. It filters the load patterns in L_M , both merged and not, simply by imposing two limits: a minimum duration and a minimum energy consumption, in order to take into account only significant loads compared to the total electric consumption for the DSM. Then each pattern

is characterized by its own load profile and by a signature vector consisting of the following properties:

- Average power
- Duration
- Total electric energy
- Average sampling frequency
- Standard deviation of the instant power

As the previous functions, the output of *Patterns Filtering* is a new set L_F of filtered load patterns.

5.2.4 Patterns Clustering

The fourth and last function, called *Patterns Clustering*, aims to group the detected load patterns in L_F into different clusters, that should represent the major home appliances used. Therefore, the K-means clustering algorithm, the same as the *Patterns Merging* function, is applied to the filtered load patterns. For clustering, each pattern is identified with its normalized signature vector, setting the maximum and minimum of each property to one and zero respectively.

The set of K_2 cluster centroids (where K_1 represents the number of clusters obtained in the *Patterns Merging* phase) defines the average signature vector of each cluster. These centroids, along with their respective average power profiles P_{cl}^j , computed by averaging the individual load profiles determined as described in subsection 5.2.2 and belonging to the respective cluster j , compose the list of significant home appliances that represents the final output of the NILM algorithm.

5.3 The REC Optimization Model

Dispose of a set of K_2 main home appliances of the houses, *i.e.* the output of the NILM algorithm described in Section 5.2, is of significant support in DSM activities for users of an energy community, with particular reference to residential ones. Indeed, the main goal of a REC is to produce its own energy through renewable power plants, typically rooftop solar PV plants, and maximize self-consumption and shared energy within the community, in order to increase the revenue due to incentives. To carry out this, since the energy production is not controllable, it is necessary to optimize power demand by scheduling consumers' loads in order to match the excess power

produced.

In this section, the adopted optimization model, developed based on the one presented in [94], is described. Electric energy demand of a single house that is part of a REC, where excess renewable energy E_{inj} injected into the grid is available, is modelled. The objective of the optimization is to maximize the shared energy E_{sh} , while simultaneously balancing the withdrawn and the injected energy, by scheduling the use of the home appliances on a daily basis. This second goal is achieved by minimizing the absolute daily difference E_{diff} between E_{inj} and the energy consumed and drawn from the grid E_{load} by the house. The appliances are identified by the clusters centroids and their relative average power profiles, as defined in subsection 5.2.4.

Once defined the positive variables E_{sh_h} and E_{diff_h} on hourly basis with $h = 0, 1, \dots, 23$, reflecting the hourly definition of shared energy established by the Italian regulation cited in subsection 3.1.1.1, the objective function to be minimized is

$$Y = \sum_h (E_{diff_h} - E_{sh_h}). \quad (5.4)$$

The definitions of E_{sh_h} and E_{diff_h} are included in the model by the following inequalities:

$$E_{sh_h} \leq E_{inj_h} \quad (5.5)$$

$$E_{sh_h} \leq E_{load_h} \quad (5.6)$$

$$E_{diff_h} \geq (E_{inj_h} - E_{load_h}) \quad (5.7)$$

$$E_{diff_h} \geq (E_{load_h} - E_{inj_h}) \quad (5.8)$$

where also E_{inj_h} and E_{load_h} are defined on hourly basis.

Indeed, in virtue of objective function (5.4), E_{sh_h} and E_{diff_h} will be respectively maximized and minimized and therefore brought to be equal respectively to the minimum between E_{load_h} and E_{inj_h} and to the absolute difference between E_{load_h} and E_{inj_h} .

E_{load_h} is obtained by (5.9) from the daily load profile of the house P_{load_t} , which is defined at discrete time steps $t \cdot \Delta ts$ with $t = 0, 1, 2, \dots, 719$. Δts is the sampling time interval and, since in Italy it is possible to exceed the contractual available power P_{av} by more than a third for a maximum of two minutes, $\Delta ts = 2$ minutes is considered.

$$E_{load_h} = \sum_{t=30h}^{30(h+1)-1} P_{load_t} / 30 \quad (5.9)$$

Since in Italy power drawn from the grid cannot exceed more than 10% of P_{av} for more than three consecutive hours, to avoid power outages also the following constraint must be satisfied:

$$\sum_h^{h+2} E_{load_h} / 3 \leq 1.1 \cdot P_{av} \quad \forall h \leq 21. \quad (5.10)$$

P_{load_t} is characterized by basic consumption, consisting of lower power loads, and by the use of home appliances, so it is obtained by

$$P_{load_t} = \left(P_{bl_t} + \sum_j P_{cl_t}^j \right) < \frac{4}{3} P_{av} \quad (5.11)$$

where:

- P_{bl_t} is known and is the baseline as defined in subsection 5.2.1, but upsampled into 2 minutes bins, replacing missing values that appeared in the resampled data with the next value in the original sequence;
- $P_{cl_t}^j$ is the daily load profile for each cluster centroid j , which corresponds to a different appliance, with $j = 0, 1, \dots, K_2$ and with K_2 that depends on the number of clusters identified by *Patterns Clustering* function.

For each cluster j , the duration in time steps dur_j and the power profile $P_{cl_k}^j$, defined at discrete time steps $k \cdot \Delta t_s$ with $k = 0, 1, \dots, dur_j$, are known, since they are outputs of the NILM algorithm. $P_{cl_k}^j$ is downsampled into 2 minutes bins, averaging the values for each bin, to obtain Δt_s as time step.

However, the cluster power profiles over time during the day are unknown. So, to model the operation of the home appliances during the day and obtain $P_{cl_t}^j$, the following variables are defined:

- $\delta_{t,k}^j$, a binary variable equal to 1 if the k -th step is taking place at time t ;
- n_j an integer variable that defines the number of execution of each cluster.

Then, power $P_{cl_t}^j$ is given by

$$P_{cl_t}^j = \sum_{k=0}^{dur_j} \left(\delta_{t,k}^j \cdot P_{cl_k}^j \right) \quad (5.12)$$

and the following constraints must be satisfied

$$\sum_{k=0}^{dur_j} \delta_{t,k}^j \leq U_t, \quad (5.13)$$

where $U_t = 1$ when an appliance can work and $U_t = 0$ when it cannot,

$$\sum_t \sum_{k=0}^{dur_j} \delta_{t,k}^j = dur_j \cdot n_j, \quad (5.14)$$

$$\delta_{t+1,k+1}^j = \delta_{t,k}^j \quad \forall j < dur_j, \quad (5.15)$$

$$\sum_t^{t+dur_j+30} \delta_{t,k}^j \leq 1 \quad \forall j = 0, t < 690 - dur_j, \quad (5.16)$$

$$\sum_j^K n_j \leq N_{max}. \quad (5.17)$$

Constraint (5.13) imposes that each appliance can only work during the granted time window and that, at any time step t , just one step k can take place, *i.e.* an appliance cannot work twice at the same time; constraint (5.14) imposes that each appliance can work n_j times during the day and must be completed at the end of the day; constraint (5.15) imposes the order of the power profile steps; constraint (5.16) avoids the appliance re-start earlier than one hour after its termination; constraint (5.17) imposes an upper limit N_{max} on the total number of executions of the home appliances, *i.e.* no more than N_{max} appliances can work in one day.

The optimization is run once per day and the expected output is the optimal daily power profile of each cluster, *i.e.* for each day whether and when it is best for each appliance associated with a cluster to work. This information, on which type of appliance is best to use and what is the best time to use it during the day, can then be sent to the household as a suggestion to maximize shared energy.

It is worth mentioning that the optimization model described was inspired by the one presented in [94], but it differs from it since it was adapted to the proposed system and was applied to a real test case. Main differences are: the sampling time, here defined as 2 minutes instead of 1 hour to take into account power draw constraints; the appliances considered, here modelled only by the average load profile of each cluster to adapt the optimization to the outputs of the NILM methodology proposed; the objective function, which, in this case, does not consider the costs in order to evaluate the possible benefits in energy terms; the type of application, in this paper it is applied within a manual DSM, according to which suggestions are sent to users once a day, rather than a MPC, which runs every hour.

5.4 Assessment of the NILM Algorithm

In this section, the results of the proposed NILM algorithm applied on the data collected from three houses located in Italy (as described in Section 5.1) are presented. The case study considers a testing period of six months, from July to December 2022. To establish a reliable ground truth, the usage of the main household appliances was manually logged throughout the testing period; specifically, the date, start time, and end time of every operational event for each house were systematically recorded in an Excel spreadsheet.

First of all, events were detected and load patterns were identified, then the optimal number of clusters was estimated using the decision score defined in Section 5.2. Figure 5.5 shows the trend of the clustering metrics as the number of clusters varies for House 1, for which the maximum decision score is obtained with 6 clusters.




Figure 5.5 – Clustering metrics trend with the number of clusters for House 1

In Figure 5.6 the results of the *Patterns Clustering* function applied to House 1 are shown, on the left side it is shown how many patterns are grouped in each cluster (the more numerous they are, the wider the coloured shape) and the silhouette coefficient for each single load pattern is depicted, on the right side the normalized cluster centroids are plotted.

The centroid of each cluster, which represents the signature vector of the load patterns grouped into that cluster, and its relative average load profile were computed and stored. The set of average load profiles of the clusters composes the database of the home appliances, which will be used for the DSM, as described in the following sections. The cluster centroids obtained for each household in the case study are presented in Table 5.2.




Figure 5.6 – Output of K-means algorithm with the optimal number of clusters for House 1: the silhouette coefficient for each single load pattern (the red dashed line represents the average) on the left and the cluster centroids on the right

Table 5.2 – Optimized Cluster Centroids

	cluster	Average power [W]	Duration [min]	Total electric energy [Wh]	Average sampling frequency [1/min]	Standard deviation [W]
House 1	0	932	46	655	0.58	1009
	1	446	156	1127	3.86	700
	2	938	101	1571	0.50	1135
	3	221	71	252	4.03	381
	4	397	117	775	0.20	945
	5	149	158	396	3.89	339
House 2	0	814	84	1087	1.61	1008
	1	481	200	1565	6.12	696
	2	258	98	409	2.00	481
House 3	0	864	84	1130	2.68	864
	1	479	101	703	0.39	714
	2	777	208	2601	2.04	889

Once defined the number of clusters and the cluster centroids, each detected load pattern was assigned to the proper cluster. Figure 5.7, Figure 5.8 and Figure 5.9 show three representative daily load profiles for each of the three houses, where the dark blue line is the aggregate instant power

profile, the light blue line is the aggregate quarter-hourly average power profile, the black line is the baseline and the green and the red dots represent respectively the detected start and end events. The portion of the aggregate load profile where a load pattern was detected is coloured with the same colour as the cluster to which the pattern was assigned.

The difference between the daily load profiles of the three houses is very high, implying different performances of the algorithm. The aggregate instant power profile of House 2, as shown in Figure 5.8, has more noise than the one of House 1. Instead, House 3 has higher consumption, having a greater available power, and a high contemporaneity factor, so two or more appliances often work at the same time, making it difficult to isolate a single load pattern.




Figure 5.7 – Daily load profile of 20th November 2022 of House 1




Figure 5.8 – Daily load profile of 21th November 2022 of House 2




Figure 5.9 – Daily load profile of 5th November 2022 of House 3

When the aggregate consumption is low, significant electrical loads are well recognized and isolated, as shown in Figure 5.9 between 10:00 and 5:00. On the other hand, when consumption is high and two or more appliances are used at the same time, as shown in Figure 5.9 from 00:00 to 10:00 and from 5:00 to 24:00, the events are detected, since the baseline is less than the measured power (respectively the black line and the light blue line in the figure), but they are not recorded because the load disaggregation does not take place correctly (the loads are detected as a single one).

In order to rate the performance of the algorithm, its outputs were compared with the actual data:

- the optimal number of clusters against the optimal number of clusters determined empirically, *i.e.* the number of significant home appliances actually used, also subdivided by working mode (e.g. the different operating programs of a washing machine).
- the number of detected load patterns that correspond to actual usage of a device (correctly detected patterns) against the number of all electrical loads that actually occurred.
- the number of detected patterns grouped into the correct cluster (correctly clustered patterns) against the number of all detected load patterns.

The results of algorithm performance are summarized in Table 5.3. For House 1, the algorithm shows high performance, reaching 95% accuracy both in events detection and patterns clustering. Instead for House 2 and House 3, it has acceptable results in events detection (with an accuracy of about 80% and 70 % respectively), but bad results in patterns clustering (with

an accuracy of about 20-25%). In the case of House 2, because of noise and greater use of appliances (+66.5% of loads that actually occurred than House 1 in the same period), many load patterns are detected that do not correspond to significant loads or two consecutive load patterns are detected as a single one. On the other hand, for House 3, bad results in detecting events are mainly due to the high contemporaneity factor: appliances that are used at the same time are recognized as a single pattern. The high variability of the profiles of the detected patterns, due to the incorrect detection, implies an incorrect clustering too. In fact, both for House 2 and for House 3, the optimal number of clusters, that was selected, is lower than that is actually necessary, because the clustering algorithm tends to divide the load patterns into a few large clusters, rather than dividing them into many small groups. Thus, different appliances are grouped into the same cluster. By empirically identifying the optimal number of clusters, knowing the number and type of loads in advance, it is possible to increase the performance of the algorithm. For House 2, by manually setting the right number of clusters, the accuracy rate in clustering patterns increases from 24.5% to 48.3%. So, with a light user involvement, the results of the proposed methodology might improve.

Ultimately in common cases such as House 1, the performance of the algorithm is excellent, however in more complex cases it is worse.

Table 5.3 – Algorithm Performance for the Case Study

	House 1	House 2	House 3
Optimal number of clusters	6	3	3
Optimal number of clusters (empirical)	6	7	6
Detected patterns	173	327	168
Correctly detected patterns	173	241	119
Loads actually occurred	182	303	168
Correctly detected patterns / Loads actually occurred	95.1%	79.5%	70.8%
Correctly clustered patterns / Detected patterns	95.4%	24.5%	21.4%

5.5 Assessment Of The Optimization Model

To test the effectiveness of the proposed optimization model, the configuration described in the last paragraph of Section 5.1 was assumed (a

simple REC consisting of House 4 and House 1) and the model was applied to House 1, on which the NILM algorithm shows the highest performance, as reported in Section 5.4, and thus provides a reliable appliance database.

First, the home appliances database of House 1 was updated by applying the NILM algorithm to consumption data from January 1 to May 20, 2024. Then, the optimization method was tested using data of the last two weeks of that period. The following inputs were provided to the model:

- the duration dur_j and the load profile P_{clt}^j of each cluster identified for House 1 by applying the NILM algorithm previously presented;
- the power baseline P_{blt} of House 1, also obtained from the outputs of the NILM algorithm;
- the hourly energy fed into the grid E_{inj_h} by House 4;
- the available power P_{av} of House 1 (see Table 5.1);
- the maximum number of appliances N_{max} that can work in one day, set to 3;
- the time range in which appliances can work, set from 7 a.m. to the end of the day.

The model has been implemented in Python, integrated with GAMS¹, to write the optimization problem, which is solved by IBM CPLEX² solver.

To validate the results, the optimized hourly load profile and then the shared energy are compared with the actual ones. Figure 5.10 shows how the optimized hourly energy drawn from the grid E_{load_opt} , the green line, is closer to the injected energy curve E_{inj} , the orange line, than the actual energy drawn from the grid E_{load} , the blue line, thus maximizing shared energy, which corresponds hourly to the minimum of E_{load_opt} and E_{inj} . Indeed, considering the two-week period from 7th to 20th May 2024, optimization increased the shared energy between House 1 and House 4 by +263%, shifting from 8.734 kWh to 31.698 kWh.

It is important to underline that the purpose of optimization is not to automatically control the loads, but to provide suggestion to users. Furthermore, the user's habits and needs are not taken into account, so the actual increase in shared energy in a real-time application may be lower. Indeed, these unconstrained suggestions represent a theoretical upper bound of the REC efficiency, but, in a real-world scenario, strict adherence to this schedule might clash with rigid daily routines.

¹ <https://www.gams.com/>

² <https://www.ibm.com/products/ilog-cplex-optimization-studio/cplex-optimizer>




Figure 5.10 – Comparison between E_{load_h} of House 1, both the actual one and the optimized one, and E_{inj_h} of 9th and 10th May 2024

The hourly energy drawn from the grid is obtained from the power load profile, which is obtained by (5.11) by adding the baseline and the load profile over time of each appliance. The definition of the daily load profile of each appliance is of particular interest since it provides information to the user on when and which appliance is best to use during the day to optimize the aggregate load profile. Figure 5.11 shows the optimized aggregate load profile for one day, in which it is suggested to use the home appliances identified by clusters 1 and 5, starting them at 3 and 3.40 pm respectively.




Figure 5.11 – Optimized load profile of House 1 of 10th May 2024, with the profiles of the suggested appliances, identified by the cluster number.

5.6 Conclusion

The activity carried out in this part of the thesis proposed a novel energy monitoring system for residential buildings and investigated over real data the application of a novel unsupervised NILM methodology, aiming to be simple, economical and as non-invasive as possible rather than having high performance.

The use of plug & play gateways, that gather low-frequency data of the DSO's smart meters, combined with a cloud-based freeware architecture, enabled the development of a lightweight and cost-effective solution that has proven reliable despite its simplicity. The performance of the algorithm with a completely unsupervised approach, evaluated on a real test-case, is very good for low energy consumption users, showing load identification rates higher than 95%, and is worse (with an identification rate of about 20-25%) when there is more noise in the load profile caused by low energy appliances or aggregated energy consumption is high and has a high contemporaneity factor. A semi-supervised approach might be considered and tested to improve the results of the proposed methodology, indeed with such an approach it is possible to double the identification rate.

Furthermore, the proposed system architecture can be used for the development of optimized DSM strategies, for residential RECs to maximize self-consumption and shared energy. In the context of this work, an optimization model is presented and a demonstration of potential benefits of such strategies applied to a REC member is shown. By applying the proposed optimization method to the tested case study, the shared energy showed an increase of +263% compared to the actual case without optimization. However, at this stage, users' habits and needs, which could reduce this increase, were not considered for the optimization.

CHAPTER 6

An Intelligent Platform for Monitoring and Optimization of RECs

The preceding chapters outlined the regulatory and structural transformations driving the energy transition. Specifically, Chapter 2 detailed the legislative framework that opened ancillary services and local flexibility markets to Distributed Energy Resources and Chapter 3 introduced Renewable Energy Communities within both the European and Italian contexts. This established the structural paradigm shift from a centralised power grid to a decentralised, participatory architecture.

To maximise the economic and operational benefits of this shift, through market participation and energy sharing, the distributed assets require advanced management systems, which must optimize the operational setpoints while maintaining user comfort and guaranteeing the provision of grid services. But, while the regulatory environment has been defined, the implementation tools required for full exploitation are not yet mature. Consequently, Chapters 4 and 5 presented applied case studies of hardware and software architectures designed to capture the necessary data to enable smart EMS across diverse environments, spanning the industrial, tertiary, and residential sectors.

With the data acquisition and processing pipeline established, the physical layer must be integrated into a centralised infrastructure capable of

translating raw telemetry into actionable flexibility. Therefore, this chapter presents the platform, denominated SPARKIE¹ (Smart Platform for Aggregated Renewable Energy communities), that was developed as part of a targeted research initiative responding to the specific objectives of Spoke 8 within the NEST² (Network 4 Energy Sustainable Transition) project, funded by the European Union through the NextGenerationEU programme, under the auspices of the Italian National Recovery and Resilience Plan. The research and development activities underpinning SPARKIE were executed during a six-month research internship at Welld Sagl³, a software engineering firm based in Lugano, Switzerland. This collaboration provided the opportunity to deepen practical understanding of industry-standard development practices essential for integrating advanced algorithms in the platform's infrastructure. Indeed, the primary objective of the SPARKIE platform is to optimise the management of DERs, thereby facilitating the active participation of end-consumers in national energy markets.

The conceptualisation of SPARKIE directly addresses four distinct technical requirements outlined in the NEST Spoke 8 framework. First, it requires the development of a comprehensive platform for the monitoring and active optimisation of a REC. Second, it mandates the creation of systems to ensure the broader management and sustainability of the electrical grid. Third, the project requires the implementation of data-driven control strategies for the decentralised coordination of local energy systems. Finally, it necessitates the development of innovative control logics applicable to specific loads, such as batteries, to provide services to the DSO. By addressing these four pillars, the platform bridges the gap between local energy production and high-level grid requirements.

The chapter proceeds as follows. Section 6.1 defines the architectural scope and functional requirements of the SPARKIE platform. Section 6.2 outlines the forecasting algorithms implemented to predict both PV generation and aggregated load profiles. Section 6.3 details the optimisation strategies deployed within the system, distinguishing between standard REC self-consumption maximisation and active participation in local flexibility markets. Finally, Section 6.4 explains the software engineering principles used to integrate these smart algorithms into the platform's microservice architecture and Section 6.5 summarizes the conclusions

¹ <https://sparkle.energy/>

² <https://fondazionenest.it/>

³ <https://www.welld.ch/>

6.1 The SPARKIE Project: Platform Scope and Architecture

The fundamental operation of a modern REC relies on incentivising local energy exchanges and maximising self-consumption. Furthermore, it could promote the aggregation of resources to offer services to the DSO in the emerging local flexibility market, thus generating an additional form of revenue. SPARKIE was created to facilitate the practical execution of the principles of energy sharing and flexibility, and it aims to achieve this by automatically balancing generation and load. It employs advanced Demand Response and Demand-Side Management techniques to orchestrate this balance. This dynamic management allows for a significantly higher penetration of distributed generation from renewable sources without violating grid constraints.

The SPARKIE platform functions as an advanced, cloud-oriented Energy Management System designed to monitor, coordinate, and actuate distributed energy assets in real time. The primary scope of the platform is to translate raw electrical measurements into automated, data-driven control strategies. To achieve this, the software architecture is designed around a highly decoupled, microservice-based framework. This structural paradigm divides the system into independent, functionally specific layers, ensuring horizontal scalability and robust fault tolerance. By exposing standard APIs, the architecture guarantees interoperability not only with heterogeneous field devices but also with external digital ecosystems, such as smart city platforms or DSO interfaces.

At the field level, the data ingestion layer interfaces directly with the hardware infrastructure. Edge gateways and smart meters acquire electrical measurements across residential and industrial nodes. This layer is responsible for normalising diverse communication protocols and transmitting telemetry to the cloud backend. To manage the massive volume of incoming time-series data without latency bottlenecks, the ingestion pipeline is strictly isolated from the platform's computational core. Specifically, the platform is designed to concurrently ingest and process multiple time series per node, classified into three main streams: electrical telemetry (active/reactive power and energy, voltage, current, and frequency), thermal measurements (energy and temperature) and environmental metrics (external temperature, humidity, wind speed, cloudiness). The system architecture natively handles mixed time

resolutions, potentially processing high-frequency data (ranging from 1 Hz down to 10-second sampling) alongside standard quarter-hourly channels. These 15-minute intervals are primarily used within the platform for standard monitoring purposes, unless specific operational needs dictate otherwise. These data are stored, with an indefinite, multi-year historical time horizon to support seasonal optimization and machine learning training, within dedicated time-series databases, structured specifically to support the rapid read-write operations required by the analytical engines.

System security and user profiling are managed by an integrated Identity and Access Management (IAM) module. This component enforces strict authentication protocols, ensuring that all data flow and the reverse actuation signals sent to physical devices are securely verified.

Observability constitutes a primary functional requirement of the architecture. The platform aggregates incoming telemetry to quantify instantaneous power flows, record historical consumption profiles, and evaluate the performance metrics of individual distributed nodes. To render these data accessible and actionable, SPARKIE features an intuitive Man-Machine Interface (MMI). This Graphical User Interface (GUI) provides a comprehensive overview of the REC's operational status, visualising real-time energy production, consumption, storage, and peer-to-peer exchanges. Dedicated dashboards allow system operators to track energy flows and accurately quantify the economic value of the shared energy. Furthermore, individual participants can access the platform to monitor their specific energy performance, receive efficiency recommendations, and track the remuneration accrued through incentivised flexibility mechanisms.

However, the technological value of SPARKIE lies beyond passive monitoring capabilities, as it functions as the computational core for DSM. For this purpose, the system hosts a suite of intelligent algorithms. The operational workflow within the platform follows a strict logical sequence, moving from data acquisition to prediction, and finally to optimisation.

First, pre-processed historical datasets and real-time telemetry are routed into forecasting modules. These algorithms predict the future states of the EC, computing forecasts for both non-dispatchable PV generation and user demand. These predictions establish the stochastic boundaries required for the subsequent control logic. Following the generation of these predictive profiles, data are processed by the optimisation engine.

Figure 6.1 and Figure 6.2 illustrate the interconnections between the algorithms and highlight the specific data inputs driving the optimisation process under two distinct operational scenarios.




Figure 6.1 – Flow diagram of the optimisation in the REC scenario




Figure 6.2 – Flow diagram of the optimisation in the flexibility scenario

In both cases, the system evaluates the forecasted profiles against the available flexibility of controllable assets, primarily batteries and PV plants. An algorithmic solver subsequently computes optimal dispatch schedules. Based on the designated objective function, the platform derives exact setpoints to either maximise REC self-consumption, as illustrated in Figure 6.1, or modulate power profiles for local flexibility market participation, as represented in Figure 6.2. Finally, the control loop is closed by transmitting these actuation signals back to the physical field devices.

This modular pipeline, monitoring, forecasting, and optimisation, enables the aggregation of passive locally distributed nodes into a dispatchable grid resource. While the overarching SPARKIE platform is the product of a collaborative consortium involving multiple industrial partners for hardware integration and cloud infrastructure deployment, the development of the intelligent algorithmic core constitutes the specific and direct contribution of this doctoral research. Consequently, the mathematical formulations and machine learning architectures governing the forecasting and optimisation engines are detailed sequentially in Section 6.2 and Section 6.3.

6.2 Forecasting Algorithms

The accurate forecasting of electrical consumption and renewable energy generation profiles represents an enabling factor for the intelligent and proactive management of RECs supported by the SPARKIE platform. In scenarios characterised by a high penetration of DERs, dynamic and non-stationary consumption behaviours, and increasing interaction with the distribution grid, the ability to predict the evolution of energy flows gets a strategic role [95].

A reliable forecast is the key for several smart-grid applications, such as optimal dispatch, active demand response, grid regulation, and intelligent energy management. Within the SPARKIE platform, electrical load and renewable production forecasting constitutes an advanced functionality supporting REC energy management, allowing for the improvement of collective self-consumption, the optimisation of energy sharing, and the maximisation of economic and environmental benefits for community members.

Machine-learning techniques are predominantly used to evaluate past observations in order to predict possible future values [96]. To build a predictive model, it is necessary to collect a sample set of real field data over

a sufficiently significant period to cover the most recurring cases (seasonality, typical production processes, work shifts, etc.). Subsequently, an algorithm is developed and trained to predict future data based on past values.

To manage the variety and complexity of problems to be addressed, various forecasting techniques have been developed over time. Each technique has its peculiar use, so it is necessary to carefully select the most suitable approach for each particular application. The selection of a methodology depends on several factors, including:

- context;
- availability and relevance of historical data;
- required level of precision;
- period for which historical data are available;
- forecasting period;
- computation time.

In recent years, forecasting techniques have found various applications in the energy sector, both in terms of consumption forecasting and in predicting energy production from renewable sources [97]. The increase in renewable generation, with the consequent impact on the operational needs of energy systems, has highlighted the requirement for accurate forecasts of energy demand and production capacity. For these reasons, the collection of historical data and the accuracy of forecasts play a central role in network management.

In SPARKIE, the forecasting modules operate as the primary analytical layer within the software architecture, converting historical telemetry and meteorological data into actionable predictive baselines. These algorithms generate load and generation profiles with a time horizon of 24-48 hours, which are necessary for day-ahead scheduling and real-time control. This section details the distinct predictive methodologies implemented within the platform, dividing the problem into PV generation forecasting and electrical load forecasting.

6.2.1 PV generation forecasting

Renewable sources, being highly dependent on environmental factors such as wind and irradiation, present an intrinsic randomness that challenges energy system management. Therefore, predicting energy production as accurately as possible is fundamental to optimally manage the grid through advanced strategies and innovative solutions.

Recently, neural network-based approaches have been widely used, facilitated by the spread of specific applications within numerical computing software (such as Python). Generally, for photovoltaic generation forecasting, standard datasets for neural network training consist of historical generation data and numerical meteorological parameters related to the plant's location (temperature, irradiation, humidity, wind speed, and sky coverage). These approaches are highly effective in short-term forecasting, typically outperforming physical models, and are comparatively easier to implement [98].

A further class of models widely exploited, especially to leverage the production of a photovoltaic plant under clear sky conditions, is represented by physical or deterministic models. These methods are essentially based on two physical variables: the irradiation on the solar panels and the temperature of the modules constituting the plant. The main advantage of this methodology is that it does not require historical data. However, physical models exhibit significant limitations when weather conditions are uncertain (e.g., cloudy skies, rain).

In this section, a hybrid forecasting technique for photovoltaic plant production, based on existing literature methodologies [54], is presented. It represents a mixture of physical/deterministic models and statistical/artificial intelligence-based approaches. Given the availability of day-ahead weather forecasts but the total absence of solar irradiance data, this combination aims to leverage the advantages of the different techniques while limiting their respective weaknesses, maintaining a computationally efficient approach.

The implemented algorithm is based on three distinct techniques, applied depending on the specific meteorological scenario. The methodology relies on the following assumptions:

- in clear or slightly cloudy sky conditions, the system utilises a deterministic model (Clear Sky Model - CSM), which is based on plant data and known irradiance equations and is applied when the sky is completely clear, and a correction of the CSM (Corrected Clear Sky Model - CCSM), when the sky is slightly cloudy;
- in uncertain weather conditions, the system employs a strategy that combines multiple artificial neural networks through a Basic Ensemble Method (BEM).

As illustrated in Figure 6.3, the approach consists of two decision rules. First, based on day-ahead weather forecasts, a decision is made regarding whether to use the BEM or a deterministic model. If a deterministic model

is selected, a second decision distinguishes between the CSM and the CCSM. These decision rules, *i.e.* classifying when the sky is slightly cloudy and when the weather conditions are uncertain, are not empirically determined a priori, but are systematically implemented through decision trees, as detailed in the following sections. This structure exploits the characteristics of the individual models: a deterministic model is more performant under clear skies, while artificial intelligence techniques are preferable under variable conditions. The forecasting method operates with a 24-hour time horizon and a 15-minute granularity, aligning with the resolution of the monitored data.




Figure 6.3 – General scheme of the implemented PV generation forecasting technique [54]

6.2.1.1 Deterministic models

When the sky is clear, the PV system is not shaded by any clouds. In this scenario, there is minimal uncertainty in the PV output profile and then a deterministic model can cover these weather conditions. The predicted PV output is modelled as described in subsection 3.2.1 of Chapter 3 of this thesis.

In situations with partly cloudy skies, a model derived from the CSM, called the Corrected Clear Sky Model, is preferred. This functions as a

regression model trained on PV historical data of partly cloudy days, *i.e.* when the sky coverage index, referred to as Cloud Cover (CC), is less than 50, where the regressors consist of the outputs P_{CSM} of the standard CSM and the CC index.

The CC index represents the percentage of the sky covered by clouds at a given instant and assumes values between 0 and 100, where:

- CC = 0 indicates a completely clear sky;
- CC = 100 indicates a completely overcast sky.

A linear regression model was implemented using a stepwise technique, which allows for the addition or removal of terms from a multi-linear model based on their statistical significance. Specifically, at each step, a potential term to be added or removed to the inputs is identified, and its p-value is calculated to test the model both by including and omitting this term. If the p-value is statistically significant, the term is retained in the model; otherwise, it is discarded. To account for non-linear behaviours between cloud cover and photovoltaic production, the input variables are considered up to the fifth power.

The method used for defining the variables can be summarised as follows:

1. the initial model is defined;
2. if any term not currently in the model has a p-value lower than the tolerance threshold (set to 0.05), the term with the lowest p-value is added, and this step is repeated; otherwise, the process proceeds to step 3;
3. if any term currently within the model has a p-value higher than the removal tolerance (set to 0.10), the term with the highest p-value is removed, and step 2 is repeated; otherwise, the algorithm terminates.

A general formula of the PV output P_{PV} for the CCSM is given by the following equation:

$$P_{PV} = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \dots + \beta_{10} x_{10} \quad (6.1)$$

where $x_1, x_2, \dots, x_n$ are the regressors (representing up to the fifth power of P_{CSM} and of CC index), β_0 is the intercept, and $\beta_1, \dots, \beta_{10}$ are the model coefficients.

6.2.1.2 Basic Ensemble Method

In the case of overcast skies, where Clear Sky models should be avoided, a model based on Multi-Layer Perceptrons (MLP) within the BEM is used.

The ensemble method in machine learning involves combining the predictions of multiple models to improve overall performance and robustness, mitigating instabilities by averaging out individual errors:

- the combined effect of different ANNs compensates for different random initialisations;
- each MLP employs a slightly different number of hidden units.

There are several typologies of ensemble methods. In the algorithm adopted in SPARKIE, the BEM output is defined in (6.2) as the average of the predictions generated by the individual neural networks:

$$f_{BEM}(t) = \frac{1}{n} \sum_{i=1}^n f_i(t) \quad (6.2)$$

where n is the total number of ANNs and $f_i(t)$ represents the output of a single network defined as a function of time t .

The proposed neural network architecture is composed of an MLP structure consisting of three levels: an input layer, a hidden layer, and an output layer. The number of neurons in the single hidden layer is automatically determined during the model training phase through a cross-validated grid search exploring an optimization range between 5 and 20 hidden units, ensuring a compact structure that avoids overfitting. Rather than predicting deviations from CSM profiles, which can introduce mathematical bias during heavily overcast conditions when the deterministic clear-sky baseline loses physical meaning, the MLP directly predicts absolute solar generation from atmospheric variables.

To design the day-ahead forecasting with quarter-hourly granularity, the following inputs are used:

- an integer between 1 and 96 defining the 15-minute interval of the day;
- an integer between 1 and 366 defining the day of the year on which the forecast is made;
- the predicted temperature [$^{\circ}\text{C}$] for the intervals in which generation is to be forecasted;
- the predicted relative humidity [%];
- the predicted wind speed [m/s];
- the predicted CC [%].

The numerical representation of time (day of the year and quarter-hour of the day) enables the model to effectively learn the temporal autocorrelations

of photovoltaic production, capturing both daily and seasonal patterns. Naturally, the accuracy in recognizing seasonal trends scales with the volume of available historical data. However, to guarantee the platform's adaptability to newly integrated power plants lacking historical measurements, the training pipeline is designed to operate incrementally. The model undergoes an initial training phase as soon as a minimum baseline of data is collected (e.g., after the first two weeks of operation) and is subsequently re-trained (every two weeks) on a rolling basis as the historical database grows, progressively refining its ability to capture long-term seasonal dynamics.

The meteorological variables are selected based on their physical influence on the yield of the photovoltaic systems. Specifically:

- temperature and wind speed affect module efficiency;
- relative humidity can indirectly impact irradiance and temperature;
- the cloud cover index is used as a synthetic proxy for the effect of clouds on solar irradiance, especially in the absence of reliable day-ahead irradiance data.

All meteorological inputs are obtained from weather forecasts and are normalised during the preprocessing phase to ensure proper scaling during neural network training.

6.2.1.3 Model Selection

As previously described, the proposed hybrid procedure involves two decision:

- Decision 1: choosing between a deterministic procedure (CSM/CCSM) and a neural network-based method (BEM technique);
- Decision 2: if Decision 1 selects the deterministic models, this choice determines whether to use CMS or CCSM.

This approach is implemented through a decision tree and using the daily average values of the available weather variables. Decision trees are an effective tool for classification and regression problems and consist of a series of "If/Else" rules applied to the regressors that determine the model's output. To understand which action to take, the algorithm traverses the decision tree from the root to the leaf node that contains the final decision.

Specifically, a Classification And Regression Tree (CART) algorithm is used. The CART algorithm builds decision trees starting from the maximum

size tree (i.e., considering all possible combinations of the variables that define the choice conditions) and subsequently prunes the branches that contribute less to overall performance. The CART algorithm evaluates different trees formed by various levels of branches to determine the optimal structure, based on predictive performance on a common test set.

Branching rules always pose a condition, the occurrence or non-occurrence of which selects a direction on the tree from the parent node to its child nodes. These conditions can take two possible forms:

- $X_i \leq C$ for continuous variables;
- $X_i \in \{\dots\}$ for categorical variables.

Decision trees can grow into a long and complex series of nodes, which is why the concept of pruning is introduced. Pruning reduces the risk of overfitting by removing branches that do not improve predictive accuracy on new data. The methodology adopted for choosing the best pruning level uses the tree corresponding to the smallest pruning level that improves the performance over the trivial tree on a validation set.

The overall process can be summarised in the following steps:

- Meteorological preprocessing: for each day, the average values of the considered meteorological variables (CC, temperature, humidity, wind speed, and pressure) are calculated, regardless of the day type.
- Selection among deterministic models (Decision Rule 2 – DR2): this phase is applied exclusively to clear or almost clear days:
 1. the performances of the deterministic models are evaluated in terms of the normalized Root Mean Squared Error (nRMSE);
 2. the difference between models performances is calculated: the sign of the difference indicates which methodology performs better;
 3. the numerical performance measure is transformed into a categorical label;
 4. the meteorological variables and the resulting labels are used to train a decision tree dedicated to DR2.
- Selection between deterministic approach and BEM (Decision Rule 1 – DR1): considering all day types:
 1. the performances of the deterministic model selected by DR2 and the BEM approach are compared;

2. the performance difference in terms of nRMSE is again transformed into a categorical label;
3. the meteorological variables and the resulting labels feed a second decision tree, dedicated to DR1.

In this way, SPARKIE utilises two distinct decision trees for the two decision rules: one to select the most appropriate deterministic model on clear days, and one to choose between the deterministic approach and the neural network-based model.

6.2.2 Load forecasting

In parallel to renewable energy generation forecasting, the prediction of electrical consumption constitutes the complementary element to complete the framework of energy flows within an energy community, forming the basis of the SPARKIE platform's optimisation strategies.

Within the context of the transition towards RECs, and more broadly the smart grid, energy demand is no longer considered as a passive variable. It transforms into a dynamic parameter whose anticipation is strictly required to effectively coordinate distributed energy resources. This challenge is further complicated by the necessity, in the examined scenario, to operate at the level of individual users or small residential and tertiary clusters. Unlike large industrial consumers or significant aggregations, the statistical smoothing effect diminishes in these contexts. Consequently, strong stochastic variability prevails, driven by user behaviours and the intermittent use of appliances. Therefore, the ability to accurately estimate future short-term requirements becomes a complex yet fundamental prerequisite. It is required not only to maximise collective self-consumption and adequately control storage systems but also to establish reliable baselines, which are indispensable for accessing flexibility markets and triggering Demand Response mechanisms.

Since a single forecasting technique cannot satisfy all operational requirements, various methodologies have been developed over the years for different purposes. A standard classification of these techniques relies on the prediction time horizon [99]:

- Very Short Term Load Forecasting (VSTLF – horizon less than a day);
- Short Term Load Forecasting (STLF – horizon from one day to two weeks);

- Medium Term Load Forecasting (MTLF – horizon from two weeks to three years);
- Long Term Load Forecasting (LTLF – horizon over three years).

A schematic representation of this conceptual division is illustrated in Figure 6.4.




Figure 6.4 – Time horizons for forecasting techniques [99]

Within the SPARKIE framework, specific attention is directed towards short-term forecasting, which is fundamental for providing operational support to RECs and enabling flexibility services.

To improve the accuracy of short-term predictions in such volatile environments, traditional analytical models often prove insufficient. However, the evolution of machine learning and artificial intelligence techniques has driven significant advancements in predicting electrical consumption and renewable generation. These approaches, applied for over thirty years in the energy domain, include ANNs, nearest-neighbour methods, kernel methods, regression trees, random forests, and, more recently, deep learning models [100]. Specifically, deep learning architectures have demonstrated superior performance in capturing the highly non-linear, volatile, and stochastic dynamics inherent in modern power systems. Unlike traditional statistical models or shallow machine learning algorithms, which often require extensive manual feature

engineering, deep learning networks can autonomously extract complex spatio-temporal representations from raw, high-dimensional time-series data [101]. This capability is particularly relevant for residential load forecasting, where aggregated consumption patterns exhibit severe and unpredictable fluctuations. By leveraging advanced network topologies, deep learning methodologies consistently yield lower forecasting errors and higher generalisation capabilities compared to traditional methods [102]. To further enhance these predictive capabilities, hybrid methodologies, which combine multiple approaches to strategically exploit the complementary strengths of distinct techniques, are also adopted.

Regardless of the specific architecture employed, the predictive accuracy of these advanced data-driven models depends on the integration of exogenous parameters alongside historical energy consumption [103]. These additional inputs act as explanatory variables, providing context to the stochastic nature of the load. Typically, these exogenous inputs encompass environmental conditions, calendar-dependent features, socio-economic factors, or the physical characteristics of the monitored building.

6.2.2.1 *The implemented hybrid CNN-LSTM model*

To address the complexities of forecasting residential or tertiary electrical loads, characterized by high volatility and irregular consumption patterns, a hybrid Deep Learning architecture was implemented within the SPARKIE project. The methodology integrates Convolutional Neural Networks (CNN) and Long Short-Term Memory (LSTM) networks, adapting approaches validated in the literature for short-term residential consumption forecasting [104][105].

The algorithm provides an hourly forecast of energy consumption for the subsequent 24 hours, using the data collected in the preceding 12 hours as the historical input observation window. This 12-hour lag was selected as a trade-off between computational efficiency and the retention of significant autoregressive information. It effectively captures immediate intra-day load dynamics without over-extending the input tensor or introducing unnecessary computational overhead during training. Furthermore, a formal sensitivity analysis on the look-back window was considered unnecessary, as it would be specific to the chosen test case, thereby compromising the model's generalizability. Fine-tuning the lag parameter on a single test case could lead to over-fitting, yielding peak performance for that specific dataset but failing to ensure predictive robustness across different, potentially heterogeneous, case studies. So, within this window, the hybrid CNN-LSTM

model effectively extracts local, short-term features via the convolutional layers, while learning medium-to-long-term temporal dependencies through the recurrent LSTM blocks.

The implemented model, along with the autoregressive flow of input and output data, is schematically illustrated in Figure 6.5. The algorithm implements a multi-step autoregressive forecasting strategy. Unlike a one-shot approach that estimates the entire horizon in a single vectorial step, the model iterates the forecast step by step: the predicted load value at time t is fed back as an input for the estimation at time $t+1$, alongside the exogenous variables known for that future instant (meteorological forecasts and temporal features).




Figure 6.5 – Logical scheme of the autoregressive CNN-LSTM architecture

In detail, the data processing pipeline and the neural network architecture are structured as follows:

- **Preprocessing and Feature Engineering:** the raw data undergo a process of normalization and hourly resampling. Despite the availability of 15-minute resolution data, an hourly temporal resolution was deliberately chosen as individual household load profiles exhibit extreme stochastic volatility due to the discrete activation of high-power appliances. Hourly aggregation effectively serves as a low-pass filter, smoothing out this high-frequency noise to ensure greater predictive robustness. Furthermore, a higher resolution is operationally redundant, as these load forecasts serve as inputs for a downstream optimization model aimed at maximizing the economic profits of RECs or

flexibility providers. Since REC energy-sharing balances and market settlements at this level are standardly regulated on an hourly basis, aligning the forecast to this timescale avoids unnecessary computational complexity in the optimization pipeline while maintaining strict relevance to the final control logic. The input vector to the model is not limited solely to the historical series of energy consumption but is enriched with exogenous variables that are critical for the load profile:

- meteorological data: temperature and cloud cover index, processed to align with the hourly resolution.
- cyclic temporal features: encoding of information related to the hour of the day, day of the week, and day of the year through sinusoidal transformations, in order to preserve the cyclic continuity of temporal variables.
- holiday indicator: a binary variable indicating the presence of national holidays specific to the reference country.
- Feature Extraction Block (CNN): the input, divided into batches of 32 elements, initially passes through three sequential convolutional stages. To steadily reduce dimensionality and extract relevant features, the number of filters decreases progressively (48, 32, and 16 filters respectively), maintaining a kernel size of 3. Each convolutional layer is followed by a Max Pooling operation (pool size = 2), to reduce the spatial dimensions of the feature maps and decrease the overall computational load, and activation via a Rectified Linear Unit (ReLU) function, with the interposition of a Dropout layer (rate 0.25) to mitigate the risk of overfitting.
- Sequential Learning Block (LSTM): the features extracted by the CNN component are processed by three cascaded LSTM layers, each consisting of 20 hidden units with the hyperbolic tangent (tanh) as the activation function. The first two layers return the entire temporal sequence, allowing the propagation of deep temporal dependencies up to the last LSTM layer, which compresses the information into a single context vector.
- Output and Optimization: the output signal from the recurrent component passes through an additional Dropout layer and two fully connected layers, the first of which consists of 20 hidden units while the last one returns the scalar forecast for the single time step. The training of the model is performed by minimizing the Mean

Squared Error (MSE) through the Adam optimizer, employing Early Stopping techniques to halt training if the validation metric shows no improvement for a predefined number of epochs.

Table 6.1 provides a summary of the configuration of the implemented hybrid CNN-LSTM model.

In the operational inference phase, the system combines recent historical data with future meteorological forecasts to generate the expected load profile, and then denormalizes the output to return the values in their original scale (kWh).

Table 6.1 – Configuration of the layers of the proposed CCN-LSTM model

Type	Filters	Kernel size	Activation function	Pool size	Units
Convolutional	48	3	ReLu	-	-
Max Pooling	-	-	-	2	-
Convolutional	32	3	ReLu	-	-
Max Pooling	-	-	-	2	-
Convolutional	16	3	ReLu	-	-
Max Pooling	-	-	-	2	-
Dropout	-	-	-	-	-
LSTM	-	-	tanh	-	20
LSTM	-	-	tanh	-	20
LSTM	-	-	tanh	-	20
Dropout	-	-	-	-	-
Dense	-	-	-	-	20
Dense	-	-	-	-	1

6.3 Optimization strategies

Within the framework of the SPARKLE project, the operational management of DERs, specifically BESS and PV plants, is governed by a deterministic optimisation algorithm based on MILP. This mathematical approach was selected for its capability to guarantee the identification of the global optimum in the presence of complex constraints and discrete variables, which are intrinsic characteristics when modelling energy systems and market logic.

The optimisation engine operates over a 24-hour day-ahead time horizon with an hourly resolution ($\Delta t = 1$ h). The software architecture is structured

to receive as inputs the generation and load forecasting data (described in the previous sections), market economic parameters, and the technical specifications of the assets. It subsequently returns as output the optimal power profile of the controllable assets and the optimal energy exchanged with the grid at the POD.

The mathematical model has been adapted into two distinct operational scenarios. These scenarios share the same core parameters and physical constraints but diverge in the formulation of their objective functions, reflecting different economic and regulatory purposes: the Renewable Energy Community scenario and the scenario involving an aggregated entity (whether a REC or not) participating in the local flexibility market.

6.3.1 REC scenario

In the REC management optimisation scenario, the MILP model is structured to maximise the aggregate economic benefit derived from energy sharing, in accordance with the Italian regulatory framework, which is detailed in subsection 3.1.1 of Chapter 3. The mathematical model considers an entire community, treating each POD as an energy balance node.

The algorithm receives a set of deterministic and forecasted parameters defined over the time horizon $t = 1 \div 24$ h:

- Economic data: hourly zonal electricity prices $c_{zon,t}$ (taken directly from MGP results on GME website), the unit valorisation tariff for the restitution of transmission charges CU_{tr} , as defined in subsection 3.1.1.2 of Chapter 3, and the energy withdrawal costs $c_{with,t,i}$. The latter are differentiated by site i , which identifies a community member's POD, according to the specific contract type (single-tier, two-tier, or multi-tier) across the three typical cost time bands (F1, F2, F3).
- Asset technical data: for each site i , the maximum contractual withdrawal power $P_{with_max,i}$ is known. For each inverter j within the REC, the nominal peak power $P_{nom,j}$, the plant commissioning date (crucial for incentive prioritisation), and the potential presence of capital grants $k_{grant,j}$ (which reduce the incentive tariff for withdrawal points not owned by territorial authorities, religious entities, third-sector entities, environmental protection groups, or natural persons) are defined. For each storage system l within the REC, the parameters include nominal capacity $E_{bess,l}$, maximum charge/discharge power $P_{ch_max,l}$ and $P_{disch_max,l}$, battery cycle efficiency $\eta_{bess,l}$, State of Charge operating limits SoC_{min} and SoC_{max} ,

and the initial state of charge $SoC_{0,i}$. A parameter γ_i for each site is also defined. This parameter is necessary to determine the renewable share of the energy injected into the grid by a single site when both a RES plant and a battery are present. Indeed, since only energy injected from renewable sources is eligible for incentives, any energy absorbed from the grid and re-injected by the storage system must be excluded from the shared energy accounting.

γ_i is defined by equation (6.3) using data from the month prior to the optimization time. If no battery is present at the site, it is equal to 0, and all the injected energy is accounted for sharing. If no RES generation plant is present, it is equal to 1, and no injected energy is accounted for sharing. Finally, if neither batteries nor RES plants are present, its value is undetermined, but there is no injected energy.

$$\gamma_i = \frac{E_{ch_grid,i}}{E_{site_inj,i}} \quad (6.3)$$

where:

- $E_{ch_grid,i}$ is the total energy absorbed from the grid by the storage systems within the site i , obtained every hour h as the portion of charging energy not covered by local PV generation, in the previous month:

$$E_{ch_grid,i} = \sum_h \max(0; E_{ch,h,i} - E_{PV,h,i}) \quad (6.4)$$

where $E_{ch,h,i}$ and $E_{PV,h,i}$ are respectively the hourly charging energy of batteries within the site and the hourly energy produced by the PV plants of site i in the previous month.

- $E_{site_inj,i}$ is the total energy provided to site i or injected into the grid by both the batteries and the RES plants in the previous month:

$$E_{site_inj,i} = \sum_h (E_{disch,h,i} + \max(0; E_{PV,h,i} - E_{ch,h,i})) \quad (6.5)$$

where $E_{disch,h,i}$ is the hourly discharging energy of batteries within the site i in the previous month.

- Forecast data: the hourly load profiles $E_{load,t,i}$ for each site i and the PV generation profiles $E_{pv,t,j}$ for each plant j over the subsequent 24 hours.

- Premium Tariff (TIP): the model pre-calculates the specific incentive tariff $TIP_{t,j}$ for each plant. This value is determined dynamically based on the plant size (<200 kW, 200-600 kW, >600 kW), adjusted by subtracting the zonal price (down to a minimum of zero) and applying any reductions linked to capital grants or regional geographic corrections, as described in subsection 3.1.1.2 of Chapter 3.

The problem variables are defined starting from the presence of a REC-managed battery at a given site. For each of these sites, charging power $P_{ch,t,l}$, discharging power $P_{disch,t,l}$ and State of Charge $SoC_{t,l}$ for each REC-managed battery belonging to the site are defined for each time step t . To ensure mutual exclusivity between charging and discharging operations, two binary variables, $a_{ch,t,l}$ and $a_{disch,t,l}$, are introduced and constraints in equations from (6.6) to (6.8) are set:

$$P_{ch,t,l} \leq a_{ch,t,l} \cdot P_{ch_max,l} \quad (6.6)$$

$$P_{disch,t,l} \leq a_{disch,t,l} \cdot P_{disch_max,l} \quad (6.7)$$

$$a_{ch,t,l} + a_{disch,t,l} \leq 1 \quad (6.8)$$

The battery dynamic is governed by the state-of-charge update equation (6.9):

$$(SoC_{t,l} - SoC_{t-1,l}) \times E_{bess,l} = \left(P_{ch,t,l} \cdot \eta_{bess,l} - \frac{P_{disch,t,l}}{\eta_{bess,l}} \right) \cdot \Delta t \quad (6.9)$$

and is subject to the capacity limits $SoC_{min} \leq SoC_{t,l} \leq SoC_{max}$.

For each site, the grid exchange $E_{grid,t,i}$ is defined by the energy balance among generation, load, and battery flows (where present), as shown in eq. (6.10):

$$E_{grid,t,i} = \sum_{j \in i} E_{pv,t,j} - E_{load,t,i} + \sum_{l \in i} (P_{disch,t,l} - P_{ch,t,l}) \cdot \Delta t \quad (6.10)$$

The withdrawn energy $E_{with,t,i}$ and the injected energy $E_{inj,t,i}$ of each site i are respectively defined as the negative part and the positive one of the grid exchange energy and are limited by the maximum contractual withdrawal power and by the nominal production power.

For batteries not managed by the REC, the energy absorbed and released by the battery itself is computed using a logic that sequentially maximises the physical self-consumption of the individual site at each instant. The

energy exchanged with the grid is therefore the direct result of a calculation, meaning no specific decision variables are required in this case.

Instead, for sites with REC-managed batteries, $E_{with,t,i}$ and $E_{inj,t,i}$ are variables and a binary variable $a_{grid,t,i}$ is also introduced to manage the sign of the net exchange, preventing the mathematical simultaneity of injection and withdrawal in the same interval. To enforce this, constraints from (6.11) to (6.15) are set:

$$a_{grid,t,i} \leq 1 + \frac{E_{grid,t,i}}{E_{grid_max,i}} \quad (6.11)$$

$$a_{grid,t,i} \geq \frac{E_{grid,t,i}}{E_{grid_max,i}} \quad (6.12)$$

$$E_{with,t,i} \leq E_{grid_max,i} \cdot (1 - a_{grid,t,i}) \quad (6.13)$$

$$E_{inj,t,i} \leq E_{grid_max,i} \cdot a_{grid,t,i} \quad (6.14)$$

$$E_{grid,t,i} = E_{inj,t,i} - E_{with,t,i} \quad (6.15)$$

where $E_{grid_max,i}$ is a generic energy value that has to be always higher than the absolute value of $E_{grid,t,i}$ (e.g. the maximum value between the maximum contractual withdrawal power and the maximum injection power multiplied by the interval of time Δt).

Then, an initial allocation of the energy injected by each plant for sharing purposes $E_{inj_pv0,t,j}$ is derived from the total injected energy of the site for sharing purposes $E_{inj_sh,t,i}$, which is defined as follows in eq. (6.16):

$$E_{inj_sh,t,i} = E_{inj,t,i} \cdot (1 - \gamma_i) \quad (6.16)$$

Constraint in eq. (6.17) enforces that, at each timestep t , the sum of the energy injected by all PV plants within the site i can be at most equal to the total injected energy of the site for sharing purposes:

$$\sum_{j \in i} E_{inj_pv0,t,j} \leq E_{inj_sh,t,i} \quad (6.17)$$

$E_{inj_sh,t,i}$ is subdivided according to a merit order: the contribution of the most recently installed plant is saturated first, since older plants are considered to contribute physically to self-consumption first. To enable this subdivision, for sites with at least one REC-managed battery, where $E_{inj_pv0,t,j}$ and $E_{inj_sh,t,i}$ are variables, an additional binary variable $a_{inj_pv,t,j}$ is introduced and, with it, constraints from (6.18) to (6.20), where j identifies each PV system of site i , sorted by commissioning date in descending order:

$$a_{inj_pv,t,j} - a_{inj_pv,t,j-1} \leq 0 \quad \forall j > 0 \quad (6.18)$$

$$a_{inj_pv,t,j} \cdot E_{pv,t,j} \leq E_{inj_pv0,t,j} \quad (6.19)$$

$$E_{inj_pv0,t,j} - a_{inj_pv,t,j-1} \cdot E_{pv,t,j} \leq 0 \quad \forall j > 0 \quad (6.20)$$

Beyond direct PV generation, REC-managed batteries can also inject energy into the grid, potentially contributing to the shared energy pool in accordance with the parameter γ_i . Specifically, if the total energy generated by the PV plants does not fully saturate the site's total injected energy eligible for sharing $E_{inj_sh,t,i}$, the residual amount represents the contribution of the storage systems. However, since the current Italian regulatory framework and REC incentive tariffs are related to the RES generation facilities rather than the storage assets, this battery-injected share must be formally reallocated among the existing PV units. Consequently, the contribution of the batteries to the final energy injected into the grid for sharing purposes $E_{inj_pv,t,j}$ is distributed equally across all generation units of the site, as defined by eq. (6.21):

$$E_{inj_pv,t,j} = E_{inj_pv0,t,j} + \frac{(E_{inj_sh,t,i} - \sum_{j \in I} E_{inj_pv0,t,j})}{\sum_{j \in I} j} \quad (6.21)$$

Once the variables for each site are established, the community variables are defined: total shared energy $E_{sh,t}$ and the shared energy attributed to a specific plant $E_{sh_pv,t,j}$.

The hourly shared energy is defined as the minimum between the total energy injected for sharing purposes and the total energy withdrawn by the community. Thus, it is constrained within the model as follows:

$$E_{sh,t} \leq \sum_i E_{with,t,i} \quad (6.22)$$

$$E_{sh,t} \leq \sum_i E_{inj_sh,t,i} \quad (6.23)$$

Since $E_{sh,t}$ is maximised during optimisation, these inequalities naturally yield the correct minimum value by definition.

To correctly allocate the shared energy to each plant, the model incorporates an inverter sorting mechanism based on the commissioning date, according to which the oldest facilities are the first to contribute to the energy sharing, as required by the regulation. The shared energy assigned to a single plant must be less than or equal to the energy injected by that plant (eq. (6.24)) and the sum of the shared energy from all plants must equal the global shared energy (eq. (6.25)):

$$E_{sh_pv,t,j} \leq E_{inj_pv,t,j} \quad (6.24)$$

$$\sum_j E_{sh_pv,t,j} = E_{sh,t} \quad (6.25)$$

To implement the shared energy allocation mechanism among all plants within the community, the binary variable $a_{sh_pv,t,j}$ and its associated constraints from (6.26) to (6.28) are introduced. In this case, unlike the definition of $E_{inj_pv,t,j}$, the variables in j refer to the PV plants sorted by commissioning date in ascending order:

$$a_{sh_pv,t,j} - a_{sh_pv,t,j-1} \leq 0 \quad \forall j > 0 \quad (6.26)$$

$$E_{sh_pv,t,j} \geq E_{inj_pv,t,j} \cdot a_{sh_pv,t,j} \quad (6.27)$$

$$E_{sh_pv,t,j} \leq E_{inj_pv,t,j} \cdot a_{sh_pv,t,j-1} \quad \forall j > 0 \quad (6.28)$$

Since the two multiplications in constraints (6.27) and (6.28) are between a continuous and a binary variable, two auxiliary variables, $E_{lin_sh_pv1,t,j}$ and $E_{lin_sh_pv2,t,j}$, are added to linearize them.

The objective of the optimisation is the maximisation of REC economic benefits, which translates into the minimisation of net operating costs, expressed as:

$$obj(t) = \min(\sum_i (E_{with,t,i} \cdot c_{with,t,i} - E_{inj,t,i} \cdot c_{zon,t}) - \sum_j (E_{sh_pv,t,j} \cdot TIP_{t,j}) - E_{sh,t} \cdot CU_{tr}) \quad (6.29)$$

where:

- the first term accounts for the net balance of standard grid exchanges for the individual sites, calculating the costs for the energy withdrawn from the grid and the revenues from the energy injected into the grid.
- the second term represents the economic incentive provided for the shared energy (*i.e.* collective self-consumption) for each PV plant.
- the third term accounts for the restitution of transmission charges. It is important to note that this specific parameter applies exclusively to the shared energy ($E_{sh,t}$), acting as a regulatory refund for avoiding the use of the transmission grid.

Since the objective function minimizes the overall net costs, all revenue streams and economic benefits, namely energy sales, REC incentives, and transmission charge restitutions, are accounted for with a negative sign.

The formulation in (6.29) drives the algorithm to charge the batteries during periods of low prices or photovoltaic surplus, and to discharge them to cover local loads (reducing withdrawals) or to boost shared energy during hours when the remuneration (the sum of energy sale, incentive, and valorisation) peaks.

6.3.2 Flexibility Scenario

The second operational scenario adapts the mathematical framework to enable the participation of aggregated users (whether configured as a REC or not) in local flexibility markets. In this configuration, the objective function is formulated to maximize economic revenues while guaranteeing the reliable provision of local ancillary services dispatched by the DSO.

Compared to the REC scenario, in this scenario the MILP model is integrated with profile-tracking constraints. Controllable assets no longer dispatch power solely for price arbitrage or collective self-consumption maximization; instead, they are constrained to minimize the tracking error between the actual power exchange at the PODs and the flexibility target profile, $E_{flex,t}$, issued by the DSO. Accordingly, the objective function $obj(t)$ integrates a remuneration term for the activated flexibility, alongside a penalty term for tracking deviations. For a generic aggregate that is not a REC, the objective function is defined as in (6.30):

$$obj(t) = \min(\sum_i (E_{with,t,i} \cdot c_{with,t,i} - E_{inj,t,i} \cdot c_{zon,t}) - \Delta E_{base,t} \cdot c_{flex} + \Delta E_{flex,t} \cdot c_{pen}) \quad (6.30)$$

where:

- $\Delta E_{base,t} \cdot c_{flex}$ represents the revenue generated from flexibility provision. It is calculated based on the price c_{flex} paid by the DSO for the service (determined via auction) and the remunerable shifted energy $\Delta E_{base,t}$. This shift is defined by (6.31) as a strictly positive variable bounded by the flexibility request and is obtained from the difference between the aggregate's net grid exchange and the baseline $E_{base,t}$ (which is defined by the DSO and thus provided as a parameter):

$$\Delta E_{base,t} = \min\left(\pm(E_{flex,t} - E_{base,t}); \max\left(0; \pm(\sum_i (E_{with,t,i} - E_{inj,t,i}) - E_{base,t})\right)\right) \quad (6.31)$$

To linearise this relationship within the MILP framework, it is necessary to introduce a binary variable $a_{base,t}$, which identifies the sign of the

difference between the aggregate's net grid exchange profile and the baseline $E_{base,t}$, along with the relative constraints in equations from (6.32) to (6.35):

$$\Delta E_{base,t} \geq 0 \quad (6.32)$$

$$\pm(\sum_i(E_{with,t,i} - E_{inj,t,i}) - E_{base,t}) \leq \Delta E_{base_max} \cdot a_{base,t} \quad (6.33)$$

$$\Delta E_{base,t} \leq \pm(E_{flex,t} - E_{base,t}) \cdot a_{base,t} \quad (6.34)$$

$$\Delta E_{base,t} \leq \pm(\sum_i(E_{with,t,i} - E_{inj,t,i}) - E_{base,t}) - \Delta E_{base_min} \cdot (1 - a_{base,t}) \quad (6.35)$$

The operator $\pm$ in equations (6.31) and (6.33) is predetermined by the direction of the DSO's dispatch request: $+$ for $E_{flex,t} > E_{base,t}$ and $-$ for $E_{flex,t} < E_{base,t}$.

- $\Delta E_{flex,t} \cdot c_{pen}$ represents the penalty term for the non-compliance with the requested flexibility profile. It relies on an arbitrary penalty factor c_{pen} applied to the absolute tracking error $\Delta E_{flex,t}$. This error is modelled as a variable and is given by the difference between $E_{flex,t}$ and the aggregate's net grid exchange:

$$\Delta E_{flex,t} = abs(E_{flex,t} - \sum_i(E_{with,t,i} - E_{inj,t,i})) \quad (6.36)$$

Since $\Delta E_{flex,t}$ is subjected to minimization within the objective function, the absolute value in (6.36) is strictly linearized through a standard formulation, based on constraints (6.37) and (6.38):

$$\Delta E_{flex,t} \geq (E_{flex,t} - \sum_i(E_{with,t,i} - E_{inj,t,i})) \quad (6.37)$$

$$\Delta E_{flex,t} \geq (\sum_i(E_{with,t,i} - E_{inj,t,i}) - E_{flex,t}) \quad (6.38)$$

Furthermore, unlike the REC scenario, the hourly generation $E_{pv,t,j}$ of controllable photovoltaic units is decoupled from the pure deterministic forecast $E_{pv_forec,t,j}$ and is treated as an active continuous decision variable subjected to constraint (6.39). This allows the optimiser to apply active power curtailment to fulfil specific flexibility requests:

$$E_{pv,t,j} \leq E_{pv_forec,t,j} \quad (6.39)$$

Consequently, the net grid exchange of a site, as shown in eq. (6.10), becomes a variable dependent on both BESS dispatch and PV curtailment actions. If such curtailment is enabled, constraints (6.19) and (6.20)

introduce non-linear products between continuous and binary variables, which are exactly linearized by introducing auxiliary continuous variables to maintain the mixed-integer linear nature of the problem.

In all other respects, the optimisation model follows the architecture of the REC scenario, omitting the shared energy accounting mechanisms if not required by the specific aggregate configuration.

6.4 Algorithmic Integration within the Platform Architecture

The operational implementation of the forecasting methodologies (PV generation and electrical load) and the optimisation logic (REC management and flexibility markets), detailed in the preceding sections, was achieved by adopting a microservice architectural paradigm. This design choice addresses the fundamental requirements of modularity, independent scalability of computational resources, and rigorous process isolation, seamlessly aligning with the broader SPARKle platform ecosystem.

Each algorithm has been encapsulated within a dedicated microservice, developed in Python to leverage its extensive ecosystem of free and open-source scientific computing and machine learning libraries. Specifically, the neural network architectures for load and generation forecasting were implemented using TensorFlow [106] and Keras [107], while the MILP optimisation models were formulated via the linopy framework [106] and resolved using the free HiGHS solver [109].

The internal structure of each service has been standardised according to the principles of Clean Architecture [110], organising the codebase into distinct logical layers to enforce a strict separation of concerns and facilitate software maintainability. Consequently, each microservice is articulated into the following fundamental layers:

- **API and Entry Point Layer:** the external interface is managed via the FastAPI⁴ framework, which exposes RESTful API endpoints for synchronous interaction and health monitoring. Although specific routes are available for the manual triggering of jobs within each service, their execution in the production environment is fully orchestrated and automated. Specifically, service activation is

⁴ <https://fastapi.tiangolo.com/>

governed by CronJobs deployed on a Kubernetes⁵ cluster. These schedule the execution of requests that trigger the algorithmic logic according to the periodicities dictated by the specific use cases (e.g., hourly updates for forecasting and day-ahead optimisation, and bi-weekly updates for model retraining).

- **Service Layer (Core Logic):** this layer hosts the pure algorithmic implementation detailed in Section 6.2 and Section 6.3. It encapsulates the core business logic responsible for orchestrating data retrieval, processing the mathematical models, and generating the final predictive or optimal outputs.
- **Persistence and Data Model Layer:** data management is delegated to SQLAlchemy⁶, an Object-Relational Mapper (ORM) that simplifies interactions between the Python codebase and SQL databases through the use of native Python objects. Given the intrinsically temporal nature of energy data (e.g., historical production, consumption series, and meteorological variables), the storage backend relies on PostgreSQL databases extended with the TimescaleDB⁷ plugin. Database schema migrations are fully automated via Alembic⁸, ensuring consistent data versioning across both development and production environments.
- **Repository and External Integration Layer:** the microservices do not operate in isolation; rather, they interact continuously with other platform components to collect field telemetry, weather forecasts, technical and economic parameters. Access to these external resources is mediated by specific repositories encapsulating API communication logic. This includes the token-based authorisation mechanisms required to securely authenticate queries to other platform endpoints, centrally managed through a Keycloak⁹ server.

Ultimately, the efficacy of the aforementioned forecasting and optimization microservices relies on their accessibility and interpretability. The logical-mathematical architecture and the underlying communication infrastructure converge into the SPARKle platform's advanced GUI. This digital environment provides end-users and system operators with robust

⁵ <https://kubernetes.io/>

⁶ <https://sqlmodel.tiangolo.com/>

⁷ <https://www.tigerdata.com/timescaledb>

⁸ <https://alembic.sqlalchemy.org/en/latest/>

⁹ <https://www.keycloak.org/>

visual analytics tools for technical supervision, economic administration, and active participation in flexibility markets.

The assessment of temporal dynamics is facilitated by interactive charts that display the instantaneous active power alongside hourly or daily energy flows, as illustrated by Figure 6.6 and Figure 6.7. These visual tools enable operators to intuitively correlate load curves with generation profiles, easily identifying injection and withdrawal peaks.




Figure 6.6 – Hourly energy and its daily cumulative value over 24 hours




Figure 6.7 – Energy data chart over time showing daily energy over a month




Figure 6.8 – Forecast chart of the aggregated production and consumption for the entire REC

The SPARKIE frontend also integrates the AI-driven forecasting models detailed in Section 6.2. By overlaying predicted load and generation profiles for the upcoming 24/48-hour horizon, as shown in Figure 6.8, directly onto the main operational dashboard, the system provides the user with an essential predictive tool to anticipate potential criticalities and strategically plan load-shifting interventions.

Furthermore, the platform incorporates a dedicated module for the simulation and monitoring of local ancillary services provided to the DSO.




Figure 6.9 – Simulation output of the activation of a flexibility service

When the MILP optimisation algorithm (tailored for the flexibility use case described in subsection 6.3.2) is triggered, the GUI acts as a comprehensive diagnostic tool, comparing the optimised dispatch scenario against the explicit flexibility targets set by the DSO. A primary KPI within this module is the compliance ratio, i.e. the percentage of scheduled shifted energy relative to the requested volume, providing an immediate validation of the aggregate’s capability to fulfil the ancillary service. Power and energy charts, like the ones in Figure 6.9, also offer granular diagnostic insights.

The power chart overlays the baseline curve (the expected grid exchange profile defined by DSO rules in the absence of interventions) with the optimised exchange curve at the POD. A distinctly highlighted area within the graph isolates the intervention time-window, showcasing the precise orchestration of the distributed resources. It visually demonstrates how the system optimally dispatches assets, either through the modulation of BESS charge/discharge cycles (optimised battery output) or through the active curtailment of renewable generation (optimised produced power), to forcefully shift the actual exchange curve toward the flexibility objective. Concurrently, the energy bar chart quantifies these hourly differentials, strictly discriminating between the expected baseline energy and the optimised energy profile, thus certifying the scheduled shifted energy and the overall reliability of the aggregate.

6.5 Conclusion

This chapter presented the architectural and algorithmic foundation of the SPARKIE platform, proposing a data-driven Energy Management System designed to address the operational complexities of the modern energy transition. A distinguishing feature of this framework is the exclusive reliance on free, open-source computational tools and libraries, ensuring high accessibility, transparency, and reproducibility of the developed models. By integrating a data ingestion pipeline with machine learning techniques, specifically tailored hybrid and ensemble neural network architectures, the system establishes a short-term forecasting framework for both load profiles and renewable generation. This predictive baseline serves as the necessary input for the subsequent deterministic optimisation layer. The core formulation, based on a MILP solver, is mathematically structured to identify global optima for the active dispatch of DERs. Notably, the optimisation framework is designed to address two distinct operational

scenarios, which can be executed either individually or concurrently depending on the aggregate's specific configuration. It can maximise the economic viability of energy sharing within Renewable Energy Communities and, independently or jointly, enforce the profile-tracking constraints mandated for participation in emerging local flexibility markets. The encapsulation of these algorithms within a scalable microservice architecture, paired with an interactive User Interface, provides a comprehensive functional proof-of-concept that strongly validates the platform's practical applicability.

While the theoretical formulation and the software architecture have been successfully established, future research efforts must pivot towards the extensive testing and validation of the platform across diverse real-world pilot sites. This field validation is strictly necessary to reliably assess the robustness of the forecasting and optimisation algorithms when subjected to the stochastic unpredictability of real-world weather conditions, physical hardware latencies, and dynamic user behaviours. Once the algorithmic reliability is empirically proven, subsequent iterations will target the deep integration of the SPARKIE platform with broader digital ecosystems. Through the exposure of open, standardised APIs, the architecture will be designed to interoperate seamlessly with advanced grid management systems, third-party aggregation platforms, and smart city infrastructures. This standardisation effort will facilitate the scalability of the proposed framework, allowing the developed control logic to be extended to larger and more heterogeneous aggregations of users. Ultimately, continuous field validation, coupled with these integration capabilities, will be essential to consolidate the proposed architecture as a reliable tool for the coordinated management of distributed energy resources.

Broadening the Horizon: European Projects on Energy Communities

In response to the global decarbonization imperative, the European Union established the Clean Energy for All Europeans Package, a legislative suite designed to facilitate the energy transition. Within this package, the 2018 Renewable Energy Directive (RED II) and the 2019 Internal Electricity Market Directive (IEMD) formalized the legal definitions for Renewable Energy Communities and Citizen Energy Communities. These directives establish a continental mandate to empower citizens as active market participants, or “prosumers”. Nevertheless, the operational deployment of these communities remains highly fragmented across Member States, constrained by varying levels of grid digitalization, socio-technical readiness, and the structural complexity of aggregating residential flexibility.

Following the preceding chapters, which conducted specific algorithmic and regulatory evaluations of RECs within the Italian context, this chapter broadens the analytical scope by presenting the methodologies and empirical findings derived from three major Horizon 2020 European research projects: LIGHTNESS¹ (market uptake of citizen energy communities enabLing a

¹ <https://www.lightness-project.eu/>

high peneTratiON of renewable Energy SourceS), HESTIA² (Holistic demand response Services for European residenTIAL communities), and NEON³ (Next-Generation Integrated Energy Services for Citizen Energy Communities). These initiatives serve as advanced empirical laboratories, providing a comprehensive overview of the tri-dimensional, social, technological, and economic, enablement of Energy Communities. By synthesizing the outcomes of these projects, this chapter illustrates how theoretical demand-side flexibility models and energy monitoring systems are operationalized and scaled across diverse European regulatory and infrastructural environments.

This chapter systematically evaluates the multidimensional barriers to community deployment and the innovative solutions proposed in the projects to overcome them. Section 7.1 analyses the LIGHTNESS project, focusing on the socio-technical dimensions of citizen engagement, the deployment of user-centric digital dashboards, and the crucial role of reflexive learning in upskilling end-users to participate in flexibility services. Section 7.2 transitions to the HESTIA project, deeply examining the techno-analytical layer; it details the rigorous formulation of data-driven Measurement and Verification (M&V) frameworks, alongside the empirical validation of holistic Demand Response outcomes across heterogeneous pilot sites. Section 7.3 explores the NEON project, dissecting the economic and contractual architectures necessary for long-term viability. This includes the adaptation of Energy Performance Contracting (EPC) and Pay-for-Performance (P4P) schemes to the residential sector, and the implementation of transactive energy frameworks leveraging Distributed Ledger Technology (DLT) and smart contracts to ensure equitable Peer-to-Peer (P2P) remuneration. Finally, Section 7.4 synthesizes the overarching lessons learned and outlines future trajectories for scaling these community architectures.

7.1 LIGHTNESS

The analysis presented in this section is based on the conference paper entitled *LIGHTNESS – Engaging citizens community in the future of energy* [111], developed as a direct scientific contribution to the LIGHTNESS project.

² <https://hestia-eu.com/>

³ <https://neonproject.eu/>

LIGHTNESS is an EU-funded project under the Horizon 2020 Framework Programme. Its primary purpose is to support various stakeholders (including residential consumers, managers/occupants in tertiary buildings, professionals in the smart grid sector, Energy Services Companies (ESCOs), Aggregators, Utilities and DSOs), in their efforts to become engaged in the creation and management of CECs.

The project is based on six major objectives:

- Engage, motivate and educate consumers on CEC benefits;
- Create a roadmap for the uptake of CECs in use-case countries;
- Raise stakeholder acceptance via CEC impact assessment tools;
- Support the market uptake with a turnkey, cutting edge, low-cost technological package for CECs;
- Demonstrate CEC impact in real Energy Communities;
- Demonstrate CEC sustainability and business models.

To achieve these objectives, an extensive inquiry into the regulatory and policy framework was performed during the project in the partner countries: Poland, Spain, France, Italy, Bulgaria and the Netherlands [112]. The methodology adopted to investigate the existing regulations for energy communities was a Regulatory Impact Assessment. A final report was completed based on both desk research activity and a survey involving respondents from twelve EU countries, focusing on topics related to the policy and regulatory aspects of ECs, flexibility services and the use of blockchain technology. Special attention was paid to the regulation regarding the flexibility market and the accessibility for different actors in the energy value chain [113]. Pilot sites were deployed, engaging end users by implementing low-cost, customizable, user-friendly and secure hardware and software solutions to operate the Energy Community. For this reason, a platform was developed, resulting in a fast-penetration solution with data analytics capabilities, including blockchain, AI forecasting, optimization, model-predictive control and features designed to increase end-user engagement.

7.1.1 Digitalization in the Engagement Process

Digital technology can help citizens make aware decisions about their energy usage: the installation of smart meters, demand-side flexibility, etc. are some of the measures that can be adopted. In particular, tracking energy consumption in homes, i.e. energy monitoring, has enormous potential for saving energy because it provides information to final domestic consumers

about their actual energy consumption and the necessary changes in their habits and strategies [114]. To make this possible, citizen engagement is needed to allow users to decide on how their energy data are used and whether they want to actively participate in the energy community [115].

Since digital tools, such as phone apps or platforms, can facilitate citizen engagement as well, in the LIGHTNESS project the digitalization of the community was an important part of the engagement process and a user-centric community dashboard was deployed. The GUI application of the LIGHTNESS platform is provided in the form of a progressive web app, which is also packaged as a mobile app available in the Apple App Store and Google Play Store, in order to further contribute to the usability, versatility and ease of access for end users. It shows EC participants relevant insights on energy production and consumption, energy price savings and energy sharing with neighbours. Providing citizens with access to data, which are collected in real time using smart meters and gathered on a central server via the internet to make them accessible to end users, on their energy production and consumption is crucial to the success of energy communities, raising awareness about the impacts of their behaviour and highlighting the economic benefits of being part of the community [116].

In addition to energy monitoring, the platform includes a section to increase customer involvement and citizens co-creation in which gamification strategies are implemented, so the more the user behaves in a sustainable way the more points he will earn and thus the higher his rank will be in the community ranking to earn rewards. The dashboard is also equipped with an interaction system via contact form. Figure 7.1 illustrates an example of the dashboard's engagement screen.




Figure 7.1 – Engagement screen with gamification widgets of the LIGHTNESS platform

Beyond basic monitoring, digitalization enables the deployment of advanced energy services to enhance flexibility. Citizens adopting the role of prosumers can supply energy to the entire EC, reducing purchase costs and employing demand-side flexibility solutions. Digitalization facilitates this through technologies such as blockchain and smart contracts, which enable secure, decentralized P2P trading and automated transactions without reliance on traditional central authorities. This technological foundation also permits the integration of innovative financing schemes, such as equity crowdfunding and lending crowdfunding, which are a form of collective bottom-up financing and participatory tools, turning participants into active financial actors within the community.

7.1.2 Upskilling citizens: lessons learnt

The establishment and development of Energy Communities in the LIGHTNESS pilot sites were facilitated by stakeholders who encouraged residents showing willingness to join the energy project to participate, offering advice and information to upgrade their energy skills. This training process followed a reflexive learning approach, continuously improving engagement strategy based on resident feedback. The case studies covered varied contexts, including a village energy cooperative in Spain, a multi-building housing community in Poland, a smart condominium in Italy, and zero-energy building and renovation projects in the Netherlands. Empirical observations across these sites consistently showed that the effectiveness of digital tools is highly dependent on the pre-existing social fabric and baseline energy literacy of the participants.

For example, in the Woerden, Netherlands pilot, involving 109 households in social housing with a senior demographic, residents showed high energy awareness despite limited digital and formal energy literacy. Technical explanations proved confusing, necessitating a shift in communication to use direct, simple, and tangible language. The key recommendation for this context is to deploy appropriate leaders to implement reflexive learning strategies and facilitate shared learning through accessible, physical interactions. Conversely, the pilot in Quatre Bras, Amsterdam, included 17 households with a high level of energy technology knowledge, but their main barrier was the limited time availability due to work and young children. For this group, a proposed solution was to utilize an ambassador to minimize the necessary activities for all members and carefully select the preferred communication channels.

In Cagliari, Italy, a smart condominium of eight households engaged in

collective self-consumption benefited from strong community ties and daily interactions. A resident possessing in-depth technical knowledge successfully served as a trusted mediator for the whole group, accelerating the decision-making process and energy learning. The recommendation here is to actively support a learning process among residents to increase their understanding of energy matters within their household and the entire building.

Finally, the Spanish pilot in Alginet, an existing energy cooperative of 13 single-family homes, exhibited mixed motivations for participation, including financial reasons and climate concerns. A significant finding was the clear need for information on the current regulatory framework, indicating that technical and legal know-how is a critical requirement for implementation and organizational structure.

A definitive lesson learned is that the origin of the community is a fundamental factor; establishing an EC in environments lacking prior social bonds requires exponentially greater engagement effort. While digital applications scale efficiently to broad target groups, they are insufficient in isolation. Physical, tangible, and interactive tools, such as in-person workshops and informative gatherings, remain an absolute necessity for engaging residents in energy communities. Although energy literacy is essential, it is not a strict precondition for EC formation, as shared learning among residents is the most effective and pleasant way for participants to build their energy skills. This learning process should be facilitated by community leaders with strong social skills who are capable of gathering feedback from end-users and consequently adapting communication strategies and tools to suit diverse groups.

7.1.3 Project findings

The analysis of the citizen engagement process and the innovative tools developed in the LIGHTNESS project provides key conclusions for fostering the development of CECs.

The project's activities confirmed that the origin of the community is a paramount factor when designing engagement and upskilling strategies, as the creation of a community where no prior social bonds exist requires substantial extra effort. Furthermore, a persistent knowledge gap was identified regarding the benefits, regulation, and technology of ECs, suggesting that this gap could be filled by establishing advisory centres to provide necessary assistance in the early stages of a project.

In terms of methodology, physical, tangible, and interactive tools, such as

workshops, were found to be the most effective way to engage residents. The involvement of party leaders is also crucial, as their figure can facilitate the collection of feedback from end users. The overall outcome of the project confirms that citizen empowerment is the essential factor for unleashing the full potential of ECs. With appropriate engagement strategies and adequate energy literacy, citizens transform into active participants who can choose, produce, distribute, and share their own energy. Digitalization is a key enabler of this transition, offering powerful tools that facilitate the development of energy services, such as demand-side flexibility and P2P energy trading, which ultimately make end-users active participants in the electricity market. A just engagement process, supported by the right tools, ensures the acceptance of ECs and validates them as a critical instrument for making citizens active players in the energy transition.

7.2 HESTIA

The analysis presented in this section is based on the deliverable *D7.2 – Validation Analysis of Demonstration Scenarios* [117], developed as a direct technical contribution to the HESTIA project.

HESTIA is an EU-funded project under the Horizon 2020 Framework Programme, whose fundamental objective is to deliver a cost-effective, scalable architecture for next-generation DR services by treating energy supply and demand side operations in a holistic, multi-carrier manner. HESTIA specifically targets the residential sector, a massive yet largely untapped reservoir of flexibility, enabling consumers to play an active role in flexibility sharing and grid balancing through the systematic exploitation of demand flexibility and the valorisation of energy efficiency in multi-carrier energy dispatching. As a driving force, HESTIA leverages consumer engagement as part of a cooperative DR strategy at the community level, through participation in RECs and CECs.

To operationalize this strategy, the HESTIA architecture implements a suite of advanced analytical and operational services that converge on a centralized platform. These core services include intelligent algorithms for energy production and demand forecasting, integrative DR optimization, MPC for automated asset dispatch, smart DR settlement and remuneration modules, and a blockchain-based platform designed for secure flexibility exchange and community trading.

The platform and its associated services were deployed and rigorously

tested across three heterogeneous European pilot sites. The Dutch pilot in Voorhout consisted of highly efficient households equipped with PV arrays, hybrid inverters, heat pumps, and localized battery storage. The French pilot in Paris-Saclay involved households relying entirely on behavioural implicit DR without local generation or storage, utilizing standard Linky smart meters. The Italian pilot in Berchidda comprised households equipped with rooftop PV systems and lithium-ion batteries, focusing heavily on self-consumption dynamics. Sensors and smart meters were deployed in the pilot sites to collect energy data.

To rigorously assess the impact of the deployed services, monitor performance against defined KPIs, and quantify the exact volume of unlocked residential flexibility, comprehensive validation activities were carried out and a strict M&V methodology was applied.

7.2.1 Baseline Formulation Framework and the IPMVP Methodology

To validate the efficacy of the applied DR strategies and accurately compute the project's KPIs, a robust M&V framework was structured. The validation methodology relies extensively on the International Performance Measurement and Verification Protocol (IPMVP) [118]. The fundamental principle governing this approach is that energy savings, or unlocked flexibility, cannot be measured directly; instead, they must be mathematically computed as the difference between the actual measured consumption during a reporting period and a modelled baseline consumption, which represents the energy that would have been consumed under identical operational conditions had the DR interventions not occurred. Specifically, reference was made to Option C of the IPMVP methodology, which utilizes utility bills or whole-building metering data to analyse energy savings, requiring a longer period of data collection to account for variables like weather, occupancy, and operational changes.

To execute this before-after comparison accurately, the baseline must account for independent variables that cause normal, stochastic fluctuations in energy use. The standard methodology involves a six-step process:

1. Nominate a representative time period for the baseline that captures all normal variations of immeasurable independent variables.
2. Acquire empirical data for energy consumption and all accessible independent variables during this baseline period.
3. Execute a mathematical regression analysis to establish the predictive coefficients for each independent variable.

4. Nominate the formal reporting (validation) period.
5. Gather data for energy consumption and the independent variables during the reporting period.
6. Apply the established baseline coefficients to the reporting period variables to yield the estimated consumption, subtracting the measured consumption to quantify the absolute energy saving.

In HESTIA, for each household, two types of baseline were defined, one for the hourly energy consumed and one for the hourly energy withdrawn from the grid, in order to respectively compare the actual energy consumed and the actual energy withdrawn from the grid during the validation period with the one that would have been consumed and the one that would have been withdrawn under the same conditions but without special interventions and without using the HESTIA services. Clearly, for households where the energy consumed and the energy withdrawn coincide, a single baseline was defined.

To evaluate the effects of the HESTIA solution on each pilot as a community, baselines were also defined considering aggregated consumption and energy withdrawn, obtained by adding those of the single households of the pilot. A baseline can be modelled using several methods, typically classified into three categories: physics-based (white-box), hybrid (grey-box), and data-driven (black-box) approaches. Advancements in metering infrastructure enhance the appeal of the data-driven approach for conducting M&V analyses in building energy retrofit projects [119].

For pilots equipped with high-resolution historical telemetry and complex interacting assets (e.g., the Netherlands and Italy), data-driven black-box approaches were utilized. A feed-forward ANN model for both types of baseline was developed according to which the energy consumed or the one imported from the grid at a given instant of time is estimated, starting from: current weather data of the pilot (temperature, relative humidity and cloud cover index), current energy produced by the PV power plants and the type of day (holidays or weekdays). The model was trained with historical data stored in the databases. In this case the model was developed in Python using Keras, integrated into the TensorFlow library, and the architecture adopted is the Multi-Layer Perceptron, which consists of three parts: input layer, at least one hidden layer and output layer. Each layer receives the inputs from the preceding layer and, by means of weighting, translation, and a nonlinear transformation, passes them to the next layer. The input layer processes the original input vector, while the output layer passes the processed values to

the user.

The set of inputs chosen for the baseline definition is the following:

- temperature,
- relative humidity,
- cloud cover index,
- hour of the day,
- produced energy (if the household has a PV plant),
- type of day (0 for weekdays and 1 for holidays)

A single hidden layer with a variable number of neurons between 10 and 25 was considered, in order to have a model that showed good results and at the same time was simple enough to avoid long training time. The neuron number is optimized minimizing the mean squared error (MSE), which is also the objective function of the ANN during the training. The stochastic gradient descent (SGD) was used as the optimization algorithm and the Rectified Linear Unit (ReLU) as the activation function between layers.

Since originally the validation should have considered a period of a whole year to evaluate the energy performance of the HESTIA solution, depending on the availability of the data, for each type of baseline of each household, a baseline was modelled for the summer period and one for the winter period, which differ from each other only due to the data used to train the models. Figure 7.2 and Figure 7.3 compare the generated baseline against the actual energy imported from the grid during a test week of the winter period for a Dutch pilot household and of the summer period for an Italian pilot household, respectively. In both cases, while the baseline successfully tracks the general trend of the actual measurements, it fails to capture the sporadic peaks.

Instead, in scenarios where the variance driven by exogenous meteorological factors was statistically minimal regarding the targeted loads, and where energy data are present sporadically for each household involved (e.g., the French pilot), standard averaged quarter-hourly load profiles were constructed. As illustrated in Figure 7.4, these profiles were separated into weekday and weekend/holiday typologies to account for standard occupancy patterns, without the application of any regression.




Figure 7.2 – Test comparison between the winter period baseline and the actual measurement of hourly energy imported from the grid, for a Dutch pilot household




Figure 7.3 – Test comparison between the summer period baseline and the actual measurement of hourly energy imported from the grid, for a Italian pilot household




Figure 7.4 – Daily average load profiles considering all days (blue), only weekdays (green) and only weekends and holidays (orange) of a French pilot household

7.2.2 Pilot Validation Outcomes

Each pilot site underwent a validation analysis during the summer of 2024 to assess the effectiveness of the HESTIA solution in meeting the objectives set for each community. Furthermore, the analysis incorporated a comprehensive assessment of each pilot site, covering main achievements, challenges, and barriers encountered during the implementation.

7.2.2.1 The Dutch Pilot: Voorhout

The primary objective of the Dutch pilot was the maximization of domestic self-consumption and the minimization of grid import.

The validation analysis demonstrated a significant reduction in total energy consumption, exceeding 20% compared to the modelled baseline, alongside a 6% absolute reduction in energy drawn from the grid, directly translating to a reduction in the community's carbon emissions. The ratio of self-consumption to total consumption consistently remained above 30%, frequently peaking above 50% during high irradiance days, but it was lower than the self-consumption of the baseline, due to the lower energy demand.

Crucially, an analysis of user interaction logs revealed a direct, quantifiable correlation between platform engagement and energy performance. As demonstrated in Table 7.1, households that accessed the platform interface frequently (defined as ≥ 20 logins per month, equating to roughly once per working day) reduced their grid import by 15% compared to less active users and their self-consumption ratio was 7% higher.

Table 7.1 – Correlation between user platform logins and energy performance in the Dutch Pilot

User Engagement Level (Dutch Pilot)	Average Daily Grid Import (kWh)	Self-Consumption / Total Consumption Ratio
High Engagement (≥ 20 Logins/Month)	6.91	43%
Low Engagement (< 20 Logins/Month)	8.14	36%
Percentage Difference	-15%	+7%

However, the pilot also highlighted severe technical limitations inherent in early-stage residential aggregation. The complexity of integrating multi-vendor smart home energy management systems resulted in numerous initial hardware faults, including inverter communication errors, overheating batteries, and malfunctioning heating systems. While these issues were resolved toward the end of the project, they induced anxiety among residents, demonstrating that technology deployed at the residential level must prioritize resilience, standard interoperability, and ease of maintenance.

7.2.2.2 The French Pilot: Paris-Saclay

In stark contrast to the highly equipped Dutch site, the French pilot comprised households without any local generation (PV) or localized storage (batteries). Consequently, the DR strategy was entirely behavioural (implicit DR), focusing on load shifting and energy sobriety rather than the automated dispatch of local assets.

Following the energy crisis caused by the war in Ukraine, between October 2022 and February 2023, a public campaign titled "Every Gesture Counts," conducted as part of the French Governmental Sobriety Plan, was launched to raise awareness among the French population about energy conservation and achieve a consumption reduction of 10% at the national scale. This campaign was supported by the widespread deployment of the EcoWatt service, a demand response management program organized by the TSO that encourages users to adjust their electricity consumption based on supply and demand fluctuations. Consequently, the citizens' energy behaviour had already been changed before the validation, during the baseline period, preventing further improvements with the use of the HESTIA platform. Indeed, the value of daily energy consumption, as the sum

of the daily average energy consumption of each household, of the validation and of the baseline were almost identical.

The French pilot underscored the challenges of data acquisition in purely behavioural deployments. Collecting load curve data at a fine granularity via smart meters was necessary for algorithmic processing but posed severe challenges regarding edge data storage during power cuts, high-bandwidth data transmission over residential internet connections, and strict compliance with General Data Protection Regulation (GDPR) constraints.

7.2.2.3 The Italian Pilot: Berchidda

The Italian pilot was focused on maximizing the utilization of locally generated energy by providing recommendations to run high-load electrical appliances during the afternoon, thereby utilizing surplus PV generation and the electricity charged in the home batteries.

Daily energy consumption of the validation period was approximately the same as the baseline, but self-consumption was higher by 20% for the actual measurements than the baseline and therefore both energy imported from the grid, which decreased by 7%, and that injected into the grid were lower. Furthermore, similar to the Dutch pilot, higher platform engagement correlated with superior energy metrics, with highly active users importing 16% less energy from the grid than less active users, as reported in Table 7.2.

Table 7.2 – Correlation between user platform logins and energy performance in the Italian Pilot

User Engagement Level (Italian Pilot)	Average Daily Grid Import (kWh)	Self-Consumption / Total Consumption Ratio
High Engagement (≥ 20 Logins/Month)	3.67	66%
Low Engagement (< 20 Logins/Month)	4.38	64%
Percentage Difference	-16%	+2%

A major technical challenge encountered in the Italian pilot was the communication integration between the IoT monitoring instruments and the batteries. Persistent communication failures effectively prevented the integration of automated MPC into the system for a significant duration of the project. Consequently, the management of the batteries defaulted to manual user intervention based on platform notifications. This failure mode

underscores the critical need for standardized, robust communication protocols. Furthermore, the processing of individual load curves required the formulation of an extensive “Experiment Convention” to ensure strict GDPR compliance, requiring significant legal coordination to define equipment ownership and data handling responsibilities.

7.2.3 Project findings

The empirical validation of the HESTIA architecture provided critical insights into the physical deployment of residential DR frameworks. While advanced control algorithms serve as a powerful instrument for energy management for optimized energy dispatching, their residential application is heavily constrained by system resilience, maintenance costs and hardware usability. From an infrastructure perspective, minimizing upfront installation costs and the frequency of on-site commissioning is fundamental for delivering an economically viable solution. Consequently, edge computing devices and IoT sensor networks must be designed for high fault tolerance and operational continuity, incorporating plug-and-play configurability to reduce central maintenance interventions and ensure long-term system stability.

Furthermore, robust data acquisition is critical for deploying the energy services. The project highlighted that ensuring high-quality telemetry necessitates structural redundancy within the communication topology. Implementing multiple pathways for data collection effectively mitigates the risks of data loss or corruption, thereby substantially enhancing overall system reliability. Concurrently, the temporal granularity of data measurement must be aligned with the specific operational requirements of the underlying algorithms, ensuring that the acquired telemetry optimizes the effectiveness of the implemented analyses without inducing excessive computational overhead or bandwidth saturation.

From a socio-technical perspective, preliminary assessments of authentic user needs and anticipated benefits are essential for simplifying solution development while maintaining high adoption rates. The project demonstrated that digital interfaces, such as web and mobile platforms, must be highly customized to the distinct characteristics of each pilot site. This tailored approach prevents participant disorientation, caters to unique demographic requirements, and increases user engagement by providing an intuitive experience. Moreover, the gradual implementation of these advanced energy services emerged as a critical strategy to prevent user frustration, facilitating incremental adaptation and minimizing resistance.

Finally, engagement barriers can be effectively dismantled by streamlining bureaucratic procedures, specifically by minimizing the documentation required from end-users; a reduction in operational friction actively curtails scepticism and fosters enduring trust among participants.

7.3 NEON

The analysis presented in this section is based on collaborative research conducted within the framework of the NEON project, drawing specifically upon the co-authored journal article entitled *Key Aspects and Challenges in the Implementation of Energy Communities* [41] and the technical deliverable *D2.2 – Multi-actor energy performance contracting* [120].

Building on the results of other EU projects, such as LIGHTNESS and HESTIA, the NEON project aims to deliver advanced, cross-sectoral energy services to enhance the quality of life for European citizens while improving the overall performance of the power system. The objective is to actively exploit building energy efficiency, renewable energy generation and storage, and demand-side flexibility to increase energy savings, reduce CO₂ emissions, and provide certified cost savings across multiple sectors. The concept of CECs is leveraged to establish the necessary legal and business foundations to accelerate the adoption of the proposed services and support energy-efficient community development.

From a technological standpoint, the project architecture is centred on a Smart Grid Architecture Model (SGAM) compliant collaborative ICT platform. This infrastructure integrates smart metering, collecting high-resolution consumption data, with energy predictive analytics and digital control systems to optimize building performance. Furthermore, the platform utilizes distributed ledger technologies (blockchain) and automated smart contracts to securely manage data and facilitate P2P trading. This technological integration facilitates the execution of Demand Response programs and provides the data necessary to implement innovative performance contracting schemes in the residential sector.

To empirically validate these service concepts and business models, the NEON project established four demonstration pilot sites across three European countries, namely: CEC 1 Berchidda (Italy), CEC 2 Domaine de la Source (France), CEC 3 Polígono industrial las cabezas (Spain) and CEC 4 Stains city (France).

While the preceding sections, describing LIGHTNESS and HESTIA

projects, focused on socio-technical engagement models and on the validation of the outcomes of the implemented energy efficiency and DR services, this section primarily targets the analysis carried out within the NEON project concerning the economic and contractual architectures required to ensure the long-term viability of CECs. Despite the legally mandated non-profit nature of CECs, establishing a robust financial framework is essential for their operational sustainability and existence. Economic benefits within these communities typically derive from self-consumption efficiency, the valorisation of excess energy, and structural public incentives. To effectively distribute these benefits and mitigate initial capital investments, the project focuses on cross-fertilizing community-centric and grid-centric business models through innovative financing mechanisms and performance-based remuneration.

7.3.1 Energy Performance Contracting Schemes

A central element of the NEON framework is the adaptation of advanced energy services and contractual models, specifically EPC and P4P schemes, targeted at the community level. According to the European Directive 2012/27/EU [121], an EPC is a formalized arrangement between a beneficiary and a service provider, typically an ESCo, where investments in efficiency measures are compensated in direct relation to a contractually agreed level of energy performance improvement or financial saving. The involvement of a third party, like an ESCo, encourages private capital investment in energy efficiency measures since it minimises financial barriers by allowing individual users to partially own the renewable generation system and avoid large upfront capital investments.

There is a diverse spectrum of programs that fall into the EPC or P4P category, but essentially, these programs track and reward energy savings and emission reductions, or demand response actions as they occur. More specifically, these mechanisms remunerate end-users, aggregators and ECs for providing or achieving better energy performances, based on metered and verified data, against a baseline scenario and focus not only on the implementation of new equipment or systems, but also on their performance [122].

The EPC and P4P market is still under development, and therefore, there is not a strict and predefined structure of contracts, but there are different contractual models that can suit the individual client. These models differ from each other in how roles are assigned, risks are shared and cost savings are allocated among the participants [123]. The EPC primary structural

schemes include:

- **Guaranteed Savings Model:** the ESCo provides a turnkey contract and guarantees a specific volume of energy savings. The contractor commits to reimbursing the client for any savings not achieved (carrying the performance risk), while the client assumes the primary financing risk and retains ownership of the installed assets. This model is advantageous for the client as it provides greater access to loans and better terms, avoiding any additional interest rate added by the ESCo.
- **Shared Saving Model:** the ESCo is responsible for designing, financing, and implementing the project, taking on both performance and credit risks. The ESCo is paid a fixed portion of the energy cost savings over a determined period. A variation is "Fast-out EPCs," which lack a fixed term, with payments ceasing once the project cost has been paid in full to the ESCo.
- **Chauffage Model:** the ESCo retains permanent ownership of the energy conversion systems (e.g., chillers or boilers) installed on the customer's premises and subsequently sells the converted energy to the client at a reduced, predefined cost while assuming operational and maintenance responsibilities. Similarly, the Build-Operate-Own (BOO) model involves the permanent involvement of the ESCo from the initial investment phase through to the continuous operation and possession of the facility.

Prominently EPC schemes are contracts between two parties, a service provider and a client, used within the public sector or the private one for industry or large commercial buildings. To transition these established commercial models into the residential CEC environment, NEON integrated P4P mechanisms. P4P schemes' organizational framework is comparable to that of EPC, but unlike EPC, which has restricted use in residential structures due to high transaction costs, P4P utilizes metering tools and an aggregator approach, making it suitable for smaller customers. The P4P model is particularly suited for evaluating savings and motivating persistence in complex, multi-measure efficiency projects that involve continuous behavioural or operational changes, such as those found in an EC where savings are mathematically difficult to deem in advance.

In P4P schemes, the financial return is directly proportional to the metered energy savings, quantified by comparing metered energy consumption with modelled baseline consumption, constructed using historical data and

normalized for exogenous variables like weather and occupancy. This approach eliminates the risk of paying for potential energy savings that never materialize. Existing P4P programs generally fall into two categories [124]:

- Standard Offer Programs: the utility determines a fixed price for a measurable unit of savings and enters into long-term contracts (typically 5–15 years) with implementers (ESCO or aggregators) to ensure the savings are realized.
- Bidding Programs: implementers, aggregators, and clients compete for contracts through auctions or competitive solicitations, where selection is based on price and the cost reduction level that must be attained over several years.

Furthermore, to automate remuneration, energy exchanges, and the distribution of benefits among community members, the NEON project proposed a financial infrastructure that leverages smart contracts and blockchain technology. Such an architecture enables P2P energy trading, serving as an innovative energy management strategy that allows participants to trade electricity directly without the oversight of a centralized entity. Moreover, a P2P network can maintain its continuous operation even if a peer is added or removed [125]. This makes it ideal for ECs since it is in line with the regulations that state that a member of an EC is free to leave the community at any time. In a P2P energy application, a smart contract is used to specify the set of rules for the exchange of energy or flexibility and the price setting for the trading. Smart contracts function as autonomous, self-deployed protocols capable of monitoring and modifying transaction records according to pre-defined rules. Usually, they are stored inside a blockchain [126]. Blockchain is a Distributed Ledger Technology that provides a secure and immutable ledger, a data book of transactions and processes that is used to keep track of records of digital transactions. By ensuring transparency and equitable benefit sharing, blockchain technology fosters trust among the community's members [127].

7.3.2 Financial Barriers and Opportunities for CECs

Implementing innovative financing schemes within the collaborative architecture of CECs presents distinct economic opportunities alongside significant systemic barriers.

A primary obstacle to the mass deployment of CECs is the inherently high transaction cost associated with residential energy efficiency and distributed generation projects. The residential sector is characterized by a high degree

of fragmentation, meaning that traditional M&V methodologies, typically designed for large-scale industrial applications, present an unbalanced cost-benefit ratio when scaled down to individual households. Consequently, the lack of aggregation leads to extended payback periods and highly uncertain short-term returns, deterring risk-averse private investors and commercial credit institutions.

This is compounded by the stochastic nature of residential energy behaviour, which makes the formulation of accurate, long-term consumption baselines algorithmically complex. Since baselines are highly susceptible to unpredictable variations in climate, occupancy, and macroeconomic factors, the deterministic certification of guaranteed savings over a standard 10-to-20-year contract duration constitutes a critical technical risk.

From a socio-technical perspective, a persistent trust deficit acts as a significant barrier. Prosumers exhibit resistance toward potential vendor lock-in, specifically concerning over-dependence on a centralized ESCo and the probability of encountering tariff structures that exceed standard utility pricing models.

To mitigate these financial barriers, the integration of performance contracting schemes provides a viable mechanism, conditional on the aggregation of community demand. By clustering several buildings or households into a singular virtual or physical portfolio, an aggregator can achieve the critical mass necessary to offset high fixed transaction costs and improve the statistical reliability of baselines.

P4P models specifically mitigate the risk of underperformance for the end-user by linking financial remuneration strictly to metered and verified data. Because the implementer or ESCo assumes the performance risk, the end-user is shielded from the financial liability of unrealized savings. In addition, the competitive nature of these schemes encourages ESCos to deploy advanced, data-driven optimization tools (e.g., non-intrusive load monitoring and predictive HVAC control) to maximize their own returns, implicitly benefiting the overall energy efficiency of the community.

To address the social barrier of ESCo dependency and foster internal community trust, P2P trading frameworks executing on blockchain offer a robust solution. By utilizing smart contracts, P2P networks automate the fair remuneration and distribution of economic benefits among peers without requiring the oversight of a traditional centralized authority. This decentralized market design guarantees transparency and equitable profit-sharing, empowering citizens to actively trade flexibility and excess generation. Ultimately, P2P architectures decentralize the risk and

democratize the financial gains, providing a resilient operational model even as individual members freely join or exit the community.

Furthermore, from the customer point of view, a critical opportunity for unlocking sustainable finance lies in the capitalization of “non-energy benefits”, such as enhanced indoor environmental quality, improved thermal comfort, and localized job creation. From a strictly financial standpoint, these building upgrades directly increase real estate asset valuations. The resulting enhancement in collateral value improves the Loan-to-Value (LTV) ratio for community investments. A lower LTV ratio statistically correlates with a significantly reduced Probability of Default (PD) and Loss Given Default (LGD) on associated financing. Recognizing these metrics allows commercial banks to offer Energy Efficient Mortgages and unlocks alternative funding mechanisms such as green bonds and equity crowdfunding.

Finally, to trigger initial participation and secure early-stage financing, phasing project deployment to prioritize high-yield, short-term interventions is recommended. Strategic business planning within CECs demonstrates that services such as HVAC system optimization and smart EV charging from renewable sources yield short-to-medium-term payback periods. Prioritizing these cost-effective measures over highly capital-intensive, long-term projects builds the foundational trust and economic momentum necessary to transition communities from passive observers to active, financially self-sustaining participants in the energy transition.

7.4 Conclusions

This chapter expanded the analytical scope from the specific Italian regulatory framework of the preceding chapters to the broader European landscape, drawing upon empirical findings from the EU-funded LIGHTNESS, HESTIA, and NEON projects. The cross-project evaluation confirms that the practical deployment of Energy Communities requires a systemic approach integrating technical, social, and economic dimensions. From a technological standpoint, the outcomes demonstrate that the aggregation of Demand-Side Flexibility and the effectiveness of energy efficiency measures rely on interoperable ICT platforms and standardized data exchange architectures to manage distributed energy resources across varying national grids and technological readiness levels.

The operationalization of these communities highlights the relevance of

advanced data analytics, such as load forecasting and optimization algorithms (consistent with the methodologies developed in the preceding chapters of this thesis), which prove effective in residential demand response programs when coupled with user-centric interfaces. In this context, validation activities through M&V methods provide transparent feedback to users regarding the benefits obtained, in support of active citizen participation in the community. Concurrently, rigorous M&V protocols validate the flexibility services effectively provided to the grid compared to established baselines.

Furthermore, sustaining citizen participation requires structural engagement strategies to enable consumers to transition into active participants who can respond effectively to flexibility events. This objective can be achieved through the implementation of customized digital dashboards and reflexive learning mechanisms, which are methods that specifically aim to meet the needs of end users.

In parallel with technical enablement, the long-term sustainability of these initiatives depends on viable business models and access to early-stage financing. The aggregator nature of energy communities, coupled with the use of M&V methodologies, enables the application of EPC and P4P models to the residential sector. By combining these schemes with digital tracking tools such as blockchain and smart contracts, it is possible to mitigate financial risk and enhance the trust of community members. Additionally, the inclusion of "non-energy benefits," such as increased asset valuation, into financial risk assessments can improve Loan-to-Value ratios and facilitate access to institutional financing.

In summary, the synthesis of these European use cases indicates that algorithmic and data-driven optimization strategies must be embedded within scalable socio-economic frameworks. The integration of technical interoperability, structured user engagement, and validated financial models constitutes a comprehensive pathway for unlocking flexibility at the community level, supporting the broader objectives of the decentralized energy transition.

CHAPTER 8

Conclusions

In response to the global imperatives of environmental sustainability and energy security, the electric power system is transitioning toward a decentralized architecture characterized by a high penetration of non-dispatchable Renewable Energy Sources. This structural shift creates a critical flexibility gap, challenging both transmission and distribution grid stability. To address these challenges, the EU enacted a succession of aggressive policy frameworks that legally mandate the massive integration of RES and enshrine the active participation of citizens through the formalization of RECs and CECs. In Italy, the transposition of these European directives has led to the definition of incentive tools for RECs and the opening of ancillary services to DERs, introducing also new flexibility markets (such as the MLF).

This thesis explored the integration of demand-side flexibility, Renewable Energy Communities, and advanced data analytics as foundational solutions for the European and Italian energy transition. The central aim was to design and validate a data-driven Energy Management System capable of orchestrating distributed assets to maximize collective self-consumption and dispatch local flexibility.

The deployment of functional demand-side management requires robust empirical telemetry. Therefore, the research initially focused on deploying

diverse lightweight edge-to-cloud monitoring architectures across agricultural, healthcare, and tertiary facilities, demonstrating how continuous data ingestion enables both detailed energy consumption analysis, to identify efficiency measures, and the deployment of advanced control systems.

To extend this data acquisition to the residential sector without invasive sub-metering, Non-Intrusive Load Monitoring was investigated and a novel methodology was developed. Calibrated specifically for the low-frequency data streams generated by DSO's second-generation smart meters via the Chain2 feature, the algorithm executes sequential event detection. By applying unsupervised clustering to the detected load patterns, it extracts disaggregated appliance signatures which feed an optimization model to execute Demand-Side Management and generate precise day-ahead scheduling recommendations, thereby increasing theoretical shared energy within a simulated REC.

Building upon this data foundation, the thesis detailed the development of a microservice-based intelligent platform, the Smart Platform for Aggregated Renewable Energy communities (SPARKIE). This platform integrates hybrid forecasting models with deterministic MILP algorithms to manage the dispatch of flexible assets. The optimization logic orchestrates battery storage and curtailment operations in two different scenarios: one approach focused on profit maximization, taking into account incentives for collective self-consumption within a REC, and another balances profit maximization with the strict profile-tracking constraints required for ancillary service provision.

Finally, since the solutions proposed previously were developed within the Italian context, three European projects focused on the implementation of innovative solutions for ECs were presented. The analysis of empirical deployments within European frameworks provides a comprehensive overview of the tri-dimensional (social, technological, and economic) enablement of Energy Communities. The synthesis of these European use cases indicates that algorithmic and data-driven optimization strategies must be embedded within scalable socio-economic frameworks. The integration of technical interoperability, structured user engagement, and validated financial models constitutes a comprehensive pathway for unlocking flexibility at the community level, supporting the broader objectives of the decentralized energy transition.

The analytical and algorithmic architectures developed in this study directly support the evolving Italian regulatory landscape, specifically RECs and the emerging Local Flexibility Markets. By providing the technical

enablement for virtually aggregated nodal units, the proposed EMS proves that decentralized prosumers can be computationally transformed into active energy users and reliable service providers.

The contribution presented in this thesis has established a reliable and robust theoretical formulation and software architecture for data-driven REC management. However, future research must prioritize the extensive testing and validation of the EMS platform across diverse, real-world pilot sites. Such field validation is strictly required to assess the results in terms of the effectiveness of the solution and to further assess the resilience of the forecasting and optimization algorithms against stochastic weather conditions, physical hardware latencies, and dynamic user behaviours. Once empirical reliability is confirmed, subsequent development should target the integration of the platform with broader digital ecosystems. Standardizing the communication interfaces will enable seamless interoperability with advanced grid management systems and third-party aggregators, facilitating the scalability of the proposed framework to manage larger, more heterogeneous user aggregations.

Bibliography

- [1] IPCC Core Writing Team, H. Lee and J. Romero. *Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. IPCC, Geneva, Switzerland, pp. 35-115, 2023. doi: 10.59327/IPCC/AR6-9789291691647
- [2] IEA, *World Energy Outlook 2025*, IEA, Paris, 2025, <https://www.iea.org/reports/world-energy-outlook-2025>, Licence: CC BY 4.0 (report); CC BY NC SA 4.0 (Annex A)
- [3] European Commission website, accessed on November 17th, 2025, https://energy.ec.europa.eu/topics/energy-efficiency/energy-efficiency-targets-directive-and-rules/energy-efficiency-directive_en
- [4] European Union, *Directive (EU) 2018/2001 of the European Parliament and of the Council of 11 December 2018 on the promotion of the use of energy from renewable sources*. Official Journal of the European Union, L 328, 2018
- [5] European Union, *Directive (EU) 2019/944 of the European Parliament and of the Council of 5 June 2019 on common rules for the internal market for electricity and amending Directive 2012/27/EU*. Official Journal of the European Union, L 158, 2019

- [6] European Commission. *The European Green Deal*. Brussels, Belgium, COM(2019) 640 final, Dec. 11, 2019.
- [7] European Commission, *'Fit for 55': delivering the EU's 2030 Climate Target on the way to climate neutrality*. Brussels, Belgium, COM(2021) 550 final, Jul. 14, 2021.
- [8] European Commission, *REPowerEU Plan*. Brussels, Belgium, COM(2022) 230 final, May 18, 2022.
- [9] European Union, *Directive (EU) 2023/2413 of the European Parliament and of the Council of 18 October 2023 amending Directive (EU) 2018/2001, Regulation (EU) 2018/1999 and Directive 98/70/EC as regards the promotion of energy from renewable sources, and repealing Council Directive (EU) 2015/652*. Official Journal of the European Union, 2023.
- [10] European Union, *Directive (EU) 2023/1791 of the European Parliament and of the Council of 13 September 2023 on energy efficiency and amending Regulation (EU) 2023/955*. Official Journal of the European Union, L 231, 2023.
- [11] European Union, *Directive (EU) 2024/1711 of the European Parliament and of the Council of 13 June 2024 amending Directives (EU) 2018/2001 and (EU) 2019/944 as regards improving the Union's electricity market design*. Official Journal of the European Union, 2024.
- [12] European Commission website, accessed on November 17th, 2025, https://energy.ec.europa.eu/topics/renewable-energy/renewable-energy-directive-targets-and-rules/renewable-energy-targets_en
- [13] European Commission website, accessed on November 17th, 2025, https://energy.ec.europa.eu/topics/energy-efficiency/energy-efficiency-targets-directive-and-rules/energy-efficiency-targets_en
- [14] European Union, *Regulation (EU) 2018/1999 of the European Parliament and of the Council of 11 December 2018 on the Governance of the Energy Union and Climate Action*. Official Journal of the European Union, L 328, 2018.
- [15] Ministry of the Environment and Energy Security (MASE), *Piano Nazionale Integrato per l'Energia e il Clima*. Rome, Italy, Jul. 2024.

-
- [16] Terna and Snam, *Documento di Descrizione degli Scenari 2024*. Rome, Italy, Oct. 2024.
- [17] Italian Government, *Legislative Decree November 8, 2021, n. 210, implementation of Directive (EU) 2019/944 on common rules for the internal market for electricity*. Official Gazette of the Italian Republic, no. 294, Dec. 11, 2021.
- [18] Ministry of the Environment and Energy Security (MASE), *Decree of December 30, 2024: Transitional support mechanism for renewable energy plants with generation costs close to market competitiveness*. Rome, Italy, Dec. 30, 2024.
- [19] IEA, *Renewables 2024*, IEA, Paris, 2024, <https://www.iea.org/reports/renewables-2024>, Licence: CC BY 4.0
- [20] Eurelectric. *Ensuring a resilient and secure energy system*. Position paper, 2025. <https://www.eurelectric.org/publications/ensuring-a-resilient-and-secure-energy-system/>
- [21] S. Püschel-Løvengreen, M. Ghazavi Dozein, S. Low, and P. Mancarella, *Separation event-constrained optimal power flow to enhance resilience in low-inertia power systems*, Electric Power Systems Research, vol. 189, p. 106678, Dec. 2020, doi: 10.1016/j.epsr.2020.106678.
- [22] Australian Energy Market Operator, *Black System South Australia 28 September 2016 – Final Report*, 2017. https://www.aemo.com.au/-/media/files/electricity/nem/market_notices_and_events/power_system_incident_reports/2017/integrated-final-report-sa-black-system-28-september-2016.pdf
- [23] Australian Energy Market Operator, *Appendix A to the Renewable Integration Study Stage 1 report*, 2020, available at <https://www.aemo.com.au/energy-systems/Major-publications/Renewable-Integration-Study-RIS>
- [24] Australian Energy Market Operator, *Appendix C to the Renewable Integration Study Stage 1 report*, 2020, available at <https://www.aemo.com.au/energy-systems/Major-publications/Renewable-Integration-Study-RIS>

- [25] A. L'Abbate, G. Fulli, F. Starr and E. Peteves. *Distributed Power Generation in Europe: Technical Issues for Further Integration*. EUR 23234 EN. Luxembourg (Luxembourg): OPOCE; 2008. JRC43063
- [26] IEA, *Status of Power System Transformation 2018*, IEA, Paris, 2018 <https://www.iea.org/reports/status-of-power-system-transformation-2018>, Licence: CC BY 4.0
- [27] L. Del Greco, A. Losi and M. Mauro. *Demand Response in the Italian regulation and first results*, Energia Elettrica Supplement Journal, 2022, doi: 10.36156/ENERGIA07
- [28] Italian Government, *Legislative Decree November 8, 2021, n. 199, implementation of Directive (EU) 2018/2001 on the promotion of the use of energy from renewable sources*. Official Gazette of the Italian Republic, no. 285, Nov. 30, 2021.
- [29] Italian Regulatory Authority for Energy, Networks and Environment (ARERA), *Resolution 727/2022/R/eel: Approval of the Integrated Text on Widespread Self-consumption (TIAD)*. Milan, Italy, Dec. 27, 2022. [Online]. Available: www.arera.it
- [30] I. Reis, I. Gonçalves, M. A.R. Lopes, C. Henggeler Antunes, *Business models for energy communities: A review of key issues and trends*, Renew. Sustain. Energy Rev. 144, 2021, doi: 10.1016/j.rser.2021.111013
- [31] Italian Regulatory Authority for Energy, Networks and Environment (ARERA), *Resolution 352/2021/R/eel: pilot projects for the procurement of local ancillary services*. Milan, Italy, Aug. 04, 2021. [Online]. Available: www.arera.it
- [32] A. Esparza, M. Blondin, J.P.F. Trovão. *A Review of Optimization Strategies for Energy Management in Microgrids*. Energies, 2025, 18, 3245, doi: 10.3390/en18133245
- [33] L. Moosbrugger, V. Seiler, P. Wohlgenannt, et al., *Load Forecasting for Households and Energy Communities: Are Deep Learning Models Worth the Effort?*, arXiv preprint arXiv.2501.05000, Jan. 2025, doi: 10.48550/arXiv.2501.05000
- [34] Italian Regulatory Authority for Energy, Networks and

- Environment (ARERA), *Resolution 300/2017/R/eel: First opening of the Ancillary Services Market (MSD) to power demand, non-eligible production units (including renewable energy sources), and energy storage systems. Establishment of pilot projects in view of the Integrated Text on Electricity Dispatching (TIDE), consistent with the European Balancing Code*. Milan, Italy, May 05, 2017. [Online]. Available: www.arera.it
- [35] Italian Regulatory Authority for Energy, Networks and Environment (ARERA), *Resolution 345/2023/R/eel: Approval of the Integrated Text on Electricity Dispatching (TIDE)*. Milan, Italy, Jul. 28, 2023. [Online]. Available: www.arera.it
- [36] Italian Regulatory Authority for Energy, Networks and Environment (ARERA), *Report 343/2025/I/eel: Outcomes of pilot projects for the qualification of new resources for the provision of global and local ancillary services. Years 2023 and 2024*. Milan, Italy, Jul. 24, 2025. [Online]. Available: www.arera.it
- [37] M. Kubli and S. Puranik. *A typology of business models for energy communities: Current and emerging design options*. Renewable and Sustainable Energy Reviews, vol. 176, p. 113165, Apr. 2023, doi: 10.1016/j.rser.2023.113165.
- [38] M. Gržanić, T. Capuder, N. Zhang, and W. Huang. *Prosumers as active market participants: A systematic review of evolution of opportunities, models and challenges*. Renewable and Sustainable Energy Reviews, vol. 154, p. 111859, Feb. 2022, doi: 10.1016/j.rser.2021.111859.
- [39] smartEn, *Energy communities to increase local system efficiency*. Feb. 2022. Available: https://smarten.eu/wp-content/uploads/2022/02/ENERGY_COMMUNITIES_to_increase_local_system_efficiency_02-2022_ONLINE-final.pdf
- [40] A. Scrocca et al. *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, Oct. 2025, doi: 10.1016/j.rser.2025.116434.
- [41] G. Yiasoumas et al. *Key Aspects and Challenges in the Implementation of Energy Communities*. Energies, vol. 16, no. 12,

- pp. 4703–4703, Jun. 2023, doi: 10.3390/en16124703.
- [42] Italian Government. *Law No. 8 of Feb. 28, 2020, Conversion into law, with modifications, of Decree-Law No. 162 of Dec. 30, 2019, regarding urgent provisions on the extension of legislative deadlines, organization of public administrations, and technological innovation*. Official Gazette of the Italian Republic, no. 51, Feb. 29, 2020.
- [43] Italian Regulatory Authority for Energy, Networks and Environment (ARERA), *Resolution 15/2024/R/eel: Regulation of economic items relating to shared electricity and updating of the Integrated Text on Widespread Self-consumption (TIAD)*. Milan, Italy, Jan. 30, 2024. [Online]. Available: www.arera.it
- [44] Ministry of the Environment and Energy Security (MASE), *Decree no. 414 of December 7, 2023: Identification of an incentive tariff for renewable energy plants included in renewable energy communities and in remote and collective self-consumption configurations*. Official Gazette of the Italian Republic, no. 31, Feb. 7, 2024.
- [45] Ministry of the Environment and Energy Security (MASE), *Decree no. 127 of May 16, 2025: Amendments to the Decree of the Minister of the Environment and Energy Security of December 7, 2023, No. 414*. Jun. 25, 2025 [Online]. Available: www.mase.gov.it.
- [46] Ministry of the Environment and Energy Security (MASE), *Directorial Decree no. 22 of February 23, 2024: Approval of the Operational Rules regarding configurations for widespread self-consumption and NRRP grants*. Feb. 23, 2024. [Online]. Available: www.mase.gov.it.
- [47] Ministry of the Environment and Energy Security (MASE), *Directorial Decree no. 228 of July 17, 2025: Update of the operational rules pursuant to Article 11 of the Decree of the Minister of the Environment and Energy Security of December 7, 2023, No. 414*. July 17, 2025. [Online]. Available: www.mase.gov.it.
- [48] Italian Regulatory Authority for Energy, Networks and Environment (ARERA), *Resolution 573/2025/R/eel: Determination of the reference revenues for the transmission and dispatching*

- service and electricity transmission tariffs for the year 2026. Approval of the 2024 allowed revenues for the transmission and dispatching service. Amendments to RTTE and ROTE.* Milan, Italy, Dec. 23, 2025. [Online]. Available: www.arera.it
- [49] Energy & Strategy Group. *Electricity Market Report 2025. Renewable Energy Communities: myth or reality?*. Politecnico di Milano, Milan, Italy, Oct. 2025. [Online]. Available: <https://www.energystrategy.it>
- [50] Y. Zhu, G. Salvalai, and P. Zangheri. *Italian renewable energy communities: status and prospect development analysis*. Energy and Buildings, vol. 348, p. 116404, Dec. 2025, doi: 10.1016/j.enbuild.2025.116404.
- [51] S. Khorrami, M. C. Falvo, and M. Pompili. *Energy Communities Development: A Review of Challenges and Opportunities*. 2024 IEEE International Conference on Environment and Electrical Engineering and 2024 IEEE Industrial and Commercial Power Systems Europe (EEEIC / I&CPS Europe), Rome, Italy, 2024, pp. 1–6, doi: 10.1109/eeeic/icpseurope61470.2024.10751123.
- [52] M. V. Gasca, Remy Rigo-Mariani, V. Debusschere, and Yousra Sidqi. *Energy communities typologies and performances: Impact of members configurations, system size and management*. Energy Reports, vol. 14, pp. 173–184, Jun. 2025, doi: 10.1016/j.egyr.2025.05.049.
- [53] F. Zumstein et al., *xlwings*, version 0.33.0, Zoomer Analytics LLC, 2024. [Online]. Available: <https://www.xlwings.org>
- [54] S. Massucco, G. Mosaico, M. Saviozzi, and F. Silvestro. *A Hybrid Technique for Day-Ahead PV Generation Forecasting Using Clear-Sky Models or Ensemble of Artificial Neural Networks According to a Decision Tree Approach*. Energies, vol. 12, no. 7, p. 1298, Apr. 2019, doi: 10.3390/en12071298.
- [55] V. Quaschnig and R. Hanitsch. *Irradiance calculation on shaded surfaces*. Solar Energy, vol. 62, no. 5, pp. 369–375, May 1998, doi: 10.1016/s0038-092x(98)00018-8.
- [56] H.T. Yang, C.M. Huang, Y.C. Huang and Y.S. Pai. *A Weather-Based Hybrid Method for 1-Day Ahead Hourly Forecasting of PV*

- Power Output*. IEEE Transactions on Sustainable Energy, vol. 5, no. 3, pp. 917–926, July 2014, doi: 10.1109/TSTE.2014.2313600.
- [57] D.L. King. *Photovoltaic module and array performance characterization methods for all system operating conditions*. AIP Conference Proceedings, 15 Feb. 1997; 394 (1): 347–368. doi: 10.1063/1.52852
- [58] T. Huld, R. Müller, and A. Gambardella. *A new solar radiation database for estimating PV performance in Europe and Africa*. Solar Energy, vol. 86, no. 6, pp. 1803–1815, Jun. 2012, doi: 10.1016/j.solener.2012.03.006.
- [59] Energy Services Manager (GSE). *Standard load and injection profiles for Renewable Energy Communities and Self-Consumption Groups 2025*. 3 Feb. 2025. [Online]. Available: <https://www.gse.it/servizi-per-te/news/comunita-di-energia-rinnovabile-e-gruppi-di-autoconsumatori-pubblicati-i-profilo-standard-per-prelievo-e-immissione-2025>.
- [60] IEA, *Digitalisation and Energy*, IEA, Paris, 2017 <https://www.iea.org/reports/digitalisation-and-energy>, Licence: CC BY 4.0
- [61] M. Sangaré, E. Bourreau, B. Fortz, A. Pachurka, M. Poss. *Loads scheduling for demand response in energy communities*. Computers & Operations Research, vol. 160, 2023, doi: 10.1016/j.cor.2023.106358
- [62] T. Fonseca, L.L. Ferreira, J. Landeck, L. Klein, P. Sousa, F. Ahmed. *Flexible Loads Scheduling Algorithms for Renewable Energy Communities*. Energies, vol. 15, no. 23, 2022, doi: 10.3390/en15238875
- [63] M.V. Gasca, R. Rigo-Mariani, V. Debusschere, and Y. Sidqi. *Fairness in energy communities: Centralized and decentralized frameworks*. Renewable and Sustainable Energy Reviews, vol. 208, p. 115054, Nov. 2024, doi: 10.1016/j.rser.2024.115054.
- [64] R. Faia, J. Soares, T. Pinto, F. Lezama, Z. Vale, and J. M. Corchado. *Optimal Model for Local Energy Community Scheduling Considering Peer to Peer Electricity Transactions*. IEEE Access, vol. 9, pp. 12420–12430, Jan. 2021, doi:

- 10.1109/access.2021.3051004.
- [65] R. Rigo-Mariani and A. Ahmed. *Coordinated Management Under Uncertainty and Fair Revenue Distribution in a Multi Energy Community*. Energy, pp. 138802–138802, Oct. 2025, doi: 10.1016/j.energy.2025.138802.
 - [66] A. D. Mustika, R. Rigo-Mariani, V. Debusschere, and A. Pachurka. *A two-stage management strategy for the optimal operation and billing in an energy community with collective self-consumption*. Applied Energy, vol. 310, p. 118484, Mar. 2022, doi: 10.1016/j.apenergy.2021.118484.
 - [67] E. Doroudchi, H. Khajeh, and H. Laaksonen. *Increasing Self-Sufficiency of Energy Community by Common Thermal Energy Storage*. IEEE Access, vol. 10, pp. 85106–85113, 2022, doi: 10.1109/access.2022.3195242.
 - [68] J. Coignard, R. Rigo-Mariani, and V. Debusschere. *Individual and collective objectives in an energy community with network constraints*. Sustainable Cities and Society, vol. 101, p. 105083, Feb. 2024, doi: 10.1016/j.scs.2023.105083.
 - [69] F. Grasso, C. I. Garcia, G. M. Lozito and G. Talluri. *Artificial Load Profiles and PV Generation in Renewable Energy Communities Using Generative Adversarial Networks*. 2022 IEEE 21st Mediterranean Electrotechnical Conference (MELECON), Palermo, Italy, 2022, pp. 709-714, doi: 10.1109/MELECON53508.2022.9843062.
 - [70] I. Ficili, M. Giacobbe, G. Tricomi, and A. Puliafito, *From Sensors to Data Intelligence: Leveraging IoT, Cloud, and Edge Computing with AI*, Sensors, vol. 25, no. 6, pp. 1763–1763, Mar. 2025, doi: 10.3390/s25061763.
 - [71] A. Manowska, A. Wycisk, A. Nowrot, and J. Pielot, *The Use of the MQTT Protocol in Measurement, Monitoring and Control Systems as Part of the Implementation of Energy Management Systems*, Electronics, vol. 12, no. 1, p. 17, Dec. 2022, doi: 10.3390/electronics12010017.
 - [72] A. Ruano, A. Hernandez, J. Ureña, M. Ruano and J. Garcia, *NILM Techniques for Intelligent Home Energy Management and Ambient*

- Assisted Living: A Review*. Energies. 2019; 12(11):2203. doi: 10.3390/en12112203
- [73] L. Yan, M. Sheikholeslami, W. Gong, W. Tian and Z. Li, *Challenges for real-world applications of nonintrusive load monitoring and opportunities for machine learning approaches*, The Electricity Journal, Vol. 35, Issue 5, 2022, 107136, ISSN 1040-6190, doi: 10.1016/j.tej.2022.107136.
- [74] P. A. Schirmer and I. Mporas, *Non-Intrusive Load Monitoring: A Review*, in IEEE Transactions on Smart Grid, vol. 14, no. 1, pp. 769-784, Jan. 2023, doi: 10.1109/TSG.2022.3189598
- [75] H. Rafiq, P. Manandhar, E. Rodriguez-Ubinas, O.A. Qureshi and T. Palpanas, *A review of current methods and challenges of advanced deep learning-based non-intrusive load monitoring (NILM) in residential context*, Energy and Buildings, Vol. 305, 2024, 113890, ISSN 0378-7788, doi: 10.1016/j.enbuild.2024.113890.
- [76] F. Paradiso, F. Paganelli, D. Giuli, and S. Capobianco, *Context-Based Energy Disaggregation in Smart Homes*, 2016, Future Internet 8, no. 1: 4. doi: 10.3390/fi8010004
- [77] R. Gopinath, M. Kumar, C. P. C. Joshua and K. Srinivas, *Energy management using non-intrusive load monitoring techniques – State-of-the-art and future research directions*, Sustainable Cities and Society, Vol. 62, 2020, 102411, ISSN 2210-6707, doi: 10.1016/j.scs.2020.102411.
- [78] G. Bucci, F. Ciancetta, E. Fiorucci and S. Mari, *Load Identification System for Residential Applications Based on the NILM Technique*, 2020 IEEE I2MTC, Dubrovnik, Croatia, 2020, pp. 1-6, doi: 10.1109/I2MTC43012.2020.9128599.
- [79] M. Weiss, A. Helfenstein, F. Mattern and T. Staake, *Leveraging smart meter data to recognize home appliances*, IEEE PerCom 2012, Lugano, Switzerland, 2012, pp. 190-197, doi: 10.1109/PerCom.2012.6199866.
- [80] S. Biansoongnern and B. Plangklang, *An Alternative Low-Cost Embedded NILM System for Household Energy Conservation with a Low Sampling Rate*. Symmetry, 2022, 14, 279. doi: 10.3390/sym14020279

-
- [81] L. d. S. Nolasco, A. E. Lazzaletti and B. M. Mulinari, *DeepDFML-NILM: A New CNN-Based Architecture for Detection, Feature Extraction and Multi-Label Classification in NILM Signals*, in IEEE Sensors Journal, vol. 22, no. 1, pp. 501-509, 1 Jan.1, 2022, doi: 10.1109/JSEN.2021.3127322.
 - [82] S. Mari, G. Bucci, F. Ciancetta, E. Fiorucci and A. Fioravanti, *An Embedded Deep Learning NILM System: A Year-Long Field Study in Real Houses*, in IEEE Transactions on Instrumentation and Measurement, vol. 72, pp. 1-15, 2023, Art no. 2531215, doi: 10.1109/TIM.2023.3328085
 - [83] Y. Liu, L. Bai, J. Ma, W. Wang and W. Ouyang, *Self-Supervised Feature Learning for Appliance Recognition in Nonintrusive Load Monitoring*, in IEEE Transactions on Industrial Informatics, vol. 20, no. 2, pp. 1698-1710, Feb. 2024, doi: 10.1109/TII.2023.3280445.
 - [84] G. Baglietto, S. Massucco, F. Silvestro, A. Vinci and F. Conte, *Optimization Model for Demand Flexibility Based on a Non-Intrusive Load Disaggregation Tool*, in IEEE Transactions on Industry Applications, vol. 62, no. 1, pp. 436-445, Jan.-Feb. 2026, doi: 10.1109/TIA.2025.3590337.
 - [85] G. Baglietto, S. Massucco, F. Silvestro, A. Vinci and F. Conte, *A Non-Intrusive Load Disaggregation Tool based on Smart Meter Data for Residential Buildings*, 2023 IEEEIC / I&CPS Europe, Madrid, Spain, 2023, pp. 1-6, doi: 10.1109/IEEEIC/ICPSEurope57605.2023.10194690.
 - [86] CEI TS 13-82:2020-04, *Electricity metering equipment - Communication with end-user devices. Part 1: Use cases*
 - [87] CEI TS 13-83:2020-04, *Electricity metering equipment - Communication with end-user devices. Part 2 : Data model and application layer*
 - [88] J. Y. Park, X. Yang, C. Miller, P. Arjunan, Z. Nagy, *Apples or oranges? Identification of fundamental load shape profiles for benchmarking buildings using a large and diverse dataset*, Applied Energy, vol. 236, pp. 1280-1295, 2019, ISSN 0306-2619, doi: 10.1016/j.apenergy.2018.12.025.
 - [89] S. Bañales, R. Dormido, N. Duro. *Smart Meters Time Series*

- Clustering for Demand Response Applications in the Context of High Penetration of Renewable Energy Resources.*, Energies 2021, vol. 14, 3458. doi: 10.3390/en14123458
- [90] P. J. Rousseeuw, *Silhouettes: A graphical aid to the interpretation and validation of cluster analysis*, Journal of Computational and Applied Mathematics, vol. 20, 1987, pp. 53-65, ISSN 0377-0427, doi: 10.1016/0377-0427(87)90125-7.
- [91] T. Calinski and J. Harabasz, *A dendrite method for cluster analysis*, Communications in Statistics - Theory and Methods, vol. 3, no. 1, pp. 1–27, 1974, doi: 10.1080/03610927408827101.
- [92] D. L. Davies and D. W. Bouldin, *A Cluster Separation Measure*, in IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. PAMI-1, no. 2, pp. 224-227, April 1979, doi: 10.1109/TPAMI.1979.4766909.
- [93] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, E. Duchesnay, *Scikit-learn: Machine Learning in Python*, Journal of Machine Learning Research, vol. 12, pp. 2825–2830, 2011.
- [94] F. Conte, F. Silvestro, A. Vinci, and A. R. Di Fazio, *Optimal coordination and discount allocation in residential renewable energy communities with smart home appliances*, in Proc. IEEE Belgrade PowerTech, 2023, pp. 1–8, doi: 10.1109/PowerTech55446.2023.10202822.
- [95] M. Saviozzi, S. Massucco, and F. Silvestro, *Implementation of advanced functionalities for Distribution Management Systems: Load forecasting and modeling through Artificial Neural Networks ensembles*, Electric Power Systems Research, vol. 167, pp. 230–239, Feb. 2019, doi: 10.1016/j.epsr.2018.10.036.
- [96] S. Makridakis, E. Spiliotis and V. Assimakopoulos, *Statistical and Machine Learning forecasting methods: Concerns and ways forward*, PLoS One. 2018 Mar 27;13(3):e0194889. doi: 10.1371/journal.pone.0194889.
- [97] A. Mystakidis, P. Koukaras, N. Tsalikidis, D. Ioannidis and C. Tjortjis, *Energy Forecasting: A Comprehensive Review of*

-
- Techniques and Technologies*, Energies, 2024; 17(7):1662. doi: 10.3390/en17071662
- [98] M. Gupta, A. Arya, U. Varshney, J. Mittal, and A. Tomar, *A review of PV power forecasting using machine learning techniques*, Progress in Engineering Science, vol. 2, no. 1, p. 100058, Jan. 2025, doi: 10.1016/j.pes.2025.100058.
- [99] T. Hong, S. Fan, *Probabilistic electric load forecasting: A tutorial review*, International Journal of Forecasting, vol. 32, no. 3, 2016, pp. 914–938. doi: 10.1016/j.ijforecast.2015.11.011
- [100] J. Liu, X. He, K. Li, and W. Xue, *A Review of State-of-the-Art AI and Data-Driven Techniques for Load Forecasting*, Energies, vol. 18, no. 16, pp. 4408–4408, Aug. 2025, doi: 10.3390/en18164408.
- [101] R. Chandrasekaran and S. K. Paramasivan, *Advances in Deep Learning Techniques for Short-term Energy Load Forecasting Applications: A Review*, Archives of Computational Methods in Engineering, vol. 32, no. 2, pp. 663–692, Jun. 2024, doi: 10.1007/s11831-024-10155-x.
- [102] M. Hasan et al., *A state-of-the-art comparative review of load forecasting methods: Characteristics, perspectives, and applications*, Energy Conversion and Management: X, vol. 26, p. 100922, Feb. 2025, doi: 10.1016/j.ecmx.2025.100922.
- [103] Y. Eren and İ. Küçükdemiral, *A comprehensive review on deep learning approaches for short-term load forecasting*, Renewable and Sustainable Energy Reviews, vol. 189, pp. 114031–114031, Jan. 2024, doi: 10.1016/j.rser.2023.114031.
- [104] M. Alhussein, K. Aurangzeb and S. I. Haider, *Hybrid CNN-LSTM Model for Short-Term Individual Household Load Forecasting*, IEEE Access, vol. 8, pp. 180544–180557, 2020. doi: 10.1109/access.2020.3028281.
- [105] T.-Y. Kim and S.-B. Cho, *Predicting residential energy consumption using CNN-LSTM neural networks*, Energy, vol. 182, pp. 72–81, Sep. 2019, doi: 10.1016/j.energy.2019.05.230.
- [106] M. Abadi et al., *TensorFlow: Large-scale machine learning on heterogeneous systems*, 2015, Software available from

tensorflow.org

- [107] F. Chollet et al., *Keras*, 2015, <https://keras.io>
- [108] F. Hofmann, *Linopy: Linear optimization with n-dimensional labeled variables*, Journal of Open Source Software, vol. 8(84), n. 4823, 2023. doi: 10.21105/joss.04823.
- [109] Q. Huangfu and J. Hall, *Parallelizing the dual revised simplex method*, Math. Prog. Comp., vol. 10, p. 119–142, 2018. doi: 10.1007/s12532-017-0130-5.
- [110] R. C. Martin, *Clean Architecture*. Prentice Hall, 2017.
- [111] M. Gaglione, M. Ricci, B. Massa, G. Baglietto and A. Vinci. *Lightness - Engaging Citizens Community in the Future of Energy*. In Conference Proceedings of BEHAVE 2023, 7th European Conference on Behaviour Change and Energy Efficiency, Maastricht, The Netherlands, 28-29 November 2023.
- [112] M. Preziosi, *PESTLE Analysis & benchmarking of CEC implementations with stakeholder mapping*, Deliverable D2.1 of the Lightness Project, 2021. Available online: https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=0801_66e5dc57e5c1&appId=PPGMS
- [113] N. Nikolova-Deme, *Regulatory Impact Analysis for CECs and flexibility services*, Deliverable D2.2 of the Lightness Project, 2021. Available online: https://ec.europa.eu/research/participants/documents/downloadPublic?documentIds=0801_66e5e3d594ef&appId=PPGMS
- [114] Etra I+D, A. Marqués, M. Serrano and L. Alacreu, *Landscape analysis for energy platforms/consumer applications*. FINAL STUDY REPORT, 2023. Available at: <https://digitalstrategy.ec.europa.eu/en/library/empowering-consumers-leveraging-digital-technologyfacilitate-voluntary-energy-reductions>
- [115] European Commission, European Climate, Infrastructure and Environment Executive Agency, B. Alpagut, X. Zhang, A. Gabaldon, *Digitalization in urban energy systems : outlook 2025, 2030 and 2040*. Publications Office of the European Union, 2022.

<https://data.europa.eu/doi/10.2926/070630>

- [116] IEA. *Empowering people – the role of local energy communities in clean energy transitions*. IEA, Paris, 2023. Available at: <https://www.iea.org/commentaries/empowering-people-the-role-of-local-energy-communities-in-clean-energy-transitions>, License: CC BY 4.0
- [117] G. Baglietto, F. Belziti, B. Massa and A. Vinci, *Validation Analysis of Demonstration Scenarios*, Deliverable D7.2 of the Hestia Project, 2024. Available online: <https://cordis.europa.eu/project/id/957823/results>
- [118] Efficiency Valuation Organization, *International Performance Measurement and Verification Protocol (IPMVP) – Core Concepts*. March 2022, available at: <https://evoworld.org/en/library/download-protocol-documents-mainmenu-en/ipmvp-core-concepts/2025-2022-new-ipmvp-core-concepts-in-english>
- [119] A. S. Alrobaie and M. Krarti, *Measurement and Verification Building Energy Prediction (MVBEP): An interpretable data-driven model development and analysis framework*, Energy and buildings, vol. 295, pp. 113321–113321, Sep. 2023, doi: 10.1016/j.enbuild.2023.113321.
- [120] E. Olabarrieta, A. Vinci and A. Pérez, *Multi-actor energy performance contracting*, Deliverable D2.2 of the Neon Project, 2022. Available online: <https://neonproject.eu/publication-results/project-deliverables/>
- [121] European Union, *Directive 2012/27/EU of the European Parliament and of the Council of 25 October 2012 on energy efficiency, amending Directives 2009/125/EC and 2010/30/EU and repealing Directives 2004/8/EC and 2006/32/EC*. Official Journal of the European Union, L 315, 2012
- [122] D. Tzani, V. Stavrakas, M. Santini, S. Thomas, J. Rosenow, and A. Flamos, *Pioneering a performance-based future for energy efficiency: Lessons learnt from a comparative review analysis of pay-for-performance programmes*, Renewable and Sustainable Energy Reviews, vol. 158, p. 112162, Apr. 2022, doi:

10.1016/j.rser.2022.112162.

- [123] T. Shang, K. Zhang, P. Liu, and Z.-W. Chen, *A review of energy performance contracting business models: Status and recommendation*, Sustainable Cities and Society, vol. 34, pp. 203–210, Oct. 2017, doi: 10.1016/j.scs.2017.06.018.
- [124] J. Szinai, M. Borgeson and E. Levin, *Putting Your Money Where Your Meter Is: A Study of Pay-for-Performance Energy Efficiency Programs in the United States*, Natural Resources Defense Council: New York, NY, USA, 2017; Available online: <https://www.nrdc.org/>
- [125] W. Tushar, T. K. Saha, C. Yuen, D. Smith, and H. V. Poor, *Peer-to-Peer Trading in Electricity Networks: An Overview*, IEEE Transactions on Smart Grid, pp. 1–1, 2020, doi: 10.1109/tsg.2020.2969657.
- [126] D. Kirli et al., *Smart contracts in energy systems: A systematic review of fundamental approaches and implementations*, Renewable and Sustainable Energy Reviews, vol. 158, p. 112013, Apr. 2022, doi: 10.1016/j.rser.2021.112013.
- [127] X. Wang, Y. Liu, R. Ma, Y. Su, and T. Ma, *Blockchain enabled smart community for bilateral energy transaction*, International Journal of Electrical Power & Energy Systems, vol. 148, p. 108997, Jun. 2023, doi: 10.1016/j.ijepes.2023.108997.

List of Publications

During the PhD, the following scientific papers were co-authored:

Journal papers

1. **G. Baglietto**, S. Maccarini, A. Traverso and P. Bruttini, *Techno-economic comparison of supercritical CO₂, steam and Organic Rankine cycles for Waste Heat Recovery applications*. ASME, Journal of Engineering for Gas Turbines and Power, April 2023; 145(4): 041021.
2. G. Yiasoumas, L. Berbakov, V. Janev, A. Asmundo, E. Olabarrieta, A. Vinci, **G. Baglietto** and G.E. Georghiou, *Key Aspects and Challenges in the Implementation of Energy Communities*. Energies 2023, 16, 4703. <https://doi.org/10.3390/en16124703>
3. **G. Baglietto**, S. Massucco, F. Silvestro, A. Vinci and F. Conte, *Optimization Model for Demand Flexibility Based on a Non-Intrusive Load Disaggregation Tool*. IEEE Transactions on Industry Applications, 2025. doi: 10.1109/TIA.2025.3590337.

Conference papers

4. **G. Baglietto**, S. Massucco, F. Silvestro, A. Vinci and F. Conte, *A Non-Intrusive Load Disaggregation Tool based on Smart Meter Data for Residential Buildings*. In Proceedings of the 23rd IEEE

International Conference on Environment and Electrical Engineering, Madrid, Spain, 6–9 June 2023.

5. M. Gaglione, M. Ricci, B. Massa, **G. Baglietto** and A. Vinci, *Lightness - Engaging Citizens Community in the Future of Energy*. In Conference Proceedings of BEHAVE 2023, 7th European Conference on Behaviour Change and Energy Efficiency, Maastricht, The Netherlands, 28-29 November 2023.

Ongoing papers

6. **G. Baglietto**, A. Bagnasco, L. Belussi, D. Menegazzo, G. Mosaico, A. Pievatolo, M. Saviozzi and A. Vinci, *Algorithms for Energy management of a Zero Energy Building within a Renewable Energy Community*.

List of Projects

During the PhD, active contributions were made to the following projects:

1. Within the **ANTES** project, funded under the cross-border cooperation program INTERREG-ALCOTRA, the activity focused on evaluating the energy impact of greenhouse cultivation, correlating electrical energy consumption with plant biomass production to provide a quantitative basis for efficiency interventions. To this end, a cloud-based real-time monitoring architecture was designed and implemented to precisely track the power consumption of the energy assets in an experimental greenhouse (geothermal heat pump, basal heating, LED lighting, and automated ventilation). The system involved the installation of IoT measurement devices at the edge level for the secure acquisition and transmission of telemetry via the MQTT protocol. At the cloud level, data was processed by a broker, stored within a time-series database (InfluxDB), and made accessible through customized dashboards (Grafana), providing an essential diagnostic tool for the real-time analysis of energy performance and the evaluation of the economic sustainability of agricultural operations.
2. Within the research project belonging to the RAISE ecosystem (Spoke 3: Environmental Protection and Care), a comprehensive monitoring architecture and a control system were implemented at

the **ZEB laboratory** of the ITC-CNR (as described in section 4.4 of this thesis). A field interaction system was designed to acquire real-time data from heterogeneous sources, including environmental sensors, a photovoltaic plant, energy storage systems, and heat pumps, managing data flow through InfluxDB, Grafana, and an MQTT broker. Furthermore, an advanced control system was integrated to optimize controllable resources—specifically batteries and heat pumps—aiming to reduce operational costs and maximize economic benefits, such as those related to energy communities. To achieve this, a MPC strategy based on a MILP formulation was designed. This approach leverages load and generation forecasts generated by ANNs and incorporates user thermal comfort constraints evaluated through a simplified thermal model fed by weather forecasts and field measurements.

3. Within the research project belonging to the RAISE ecosystem, entitled “**A-ISoIE** (Management Solutions for Electrical Energy Distribution Systems with Centralized and Distributed Artificial Intelligence)”, the functionalities of a Distribution Management System were extended through Artificial Intelligence-based operational algorithms. A substantial part of the activity was dedicated to consolidating test scenarios, with a specific focus on RECs and the flexibility market. In particular, the functions and logical schemes necessary for developing the software components were identified. An optimization procedure was designed to minimize the energy costs of a test substation configured as a prosumer, coordinating flexible loads (such as HVAC systems and batteries) and implementing local control through edge computing logic. Furthermore, a control algorithm and a rule-based decision-making process for the management of controllable assets were designed to fulfil the simulated distributor's requests under active participation in the flexibility market.
4. Within the framework of the **LIGHTNESS** project, funded by the EU under the Horizon 2020 programme, the socio-technical dimensions of citizen engagement within energy communities have been analysed. The activity, in collaboration with Axpo Energy Solutions Italia, focused on evaluating the multidimensional barriers to community deployment and the integration of user-centric digital dashboards (as described in section 7.1 of this thesis).

In particular, the crucial role of reflexive learning mechanisms in upskilling end-users to participate in demand-side flexibility services was investigated, leading to a direct scientific contribution via a conference paper (“Lightness - Engaging Citizens Community in the Future of Energy” presented at the BEHAVE 2023 conference in Maastricht). The analysis concluded that while digital applications are highly scalable, physical and interactive tools remain a fundamental technological necessity for engaging residents, especially in environments lacking prior social bonds.

5. Within the framework of the **HESTIA** project, funded by the EU under the Horizon 2020 programme, a scalable architecture for next-generation Demand Response services has been evaluated. The activity, in collaboration with Axpo Energy Solutions Italia, focused on treating energy supply and demand side operations in a holistic, multi-carrier manner to unlock the flexibility of the residential sector (as described in section 7.2 of this thesis). In particular, rigorous data-driven M&V frameworks, based on the International Performance Measurement and Verification Protocol (IPMVP), were structured to quantify the residential flexibility actually unlocked across heterogeneous European pilot sites. This evaluation resulted in a direct technical contribution via a deliverable concerning the validation analysis of the demonstration scenarios.
6. Within the framework of the **NEON** project, funded by the EU under the Horizon 2020 programme, the economic and contractual architectures required to ensure the long-term viability of CECs have been analysed. The activity, in collaboration with Axpo Energy Solutions Italia, focused on delivering advanced, cross-sectoral energy services through the adaptation of Energy Performance Contracting and Pay-for-Performance schemes to the residential sector (as described in section 7.3 of this thesis). In particular, a Smart Grid Architecture Model (SGAM) compliant collaborative platform was implemented, utilizing distributed ledger technologies (blockchain) and automated smart contracts to securely manage data and facilitate Peer-to-Peer trading. These models and services were empirically validated across four European demonstration pilot sites, resulting in a direct scientific and technical contribution via a journal article (“Key Aspects and

Challenges in the Implementation of Energy Communities” published in *Energies*) and a deliverable concerning multi-actor energy performance contracting.

7. Within the research project **SPARKIE** (Smart Platform for Aggregated Renewable Energy Communities), funded under the PE2-NEST program, a smart platform for the management of RECs was developed. This activity, carried out during a research mobility period at Welld Sagl (Lugano, Switzerland), constitutes the core of Chapter 6 of this thesis. The work initially focused on defining the software requirements for energy monitoring and the management of both individual users and communities as a whole. Subsequently, advanced forecasting and optimization services were integrated into the platform using Python. Specifically, a photovoltaic generation forecasting algorithm was implemented, utilizing a decision tree approach to select between clear-sky models and ANNs, alongside load forecasting models based on CNN-LSTM architectures. Furthermore, an energy management optimization procedure was designed using MILP. This procedure processes generation and load forecasts to compute the optimized load profiles for controllable assets, such as batteries and generation plants, ultimately defining the operational setpoints for direct field control to either maximize self-consumption or provide aggregated flexibility services to DSOs.
8. Within the research project belonging to the RAISE ecosystem (Spoke 4: Smart and Sustainable Ports), titled “Innovative demand-responsive energy system for port environment”, the implementation of a Virtual Aggregation Platform for a Port Renewable Energy Community (PREC) has been developed. The activity focused on the deployment of an advanced EMS integrating a PV plant, a BESS and smart multiservice charging pillars at the **Marina di Finale** pilot site. In particular, an IoT-based telemetry architecture utilizing time-series databases and web-based dashboards was established to enable real-time data acquisition and, then, PV and load forecasting. This infrastructure allows for the continuous execution of an optimization algorithm designed to actively manage the microgrid’s demand-side flexibility by dispatching dynamic set-points to field devices, ensuring optimal energy flows and reducing operational costs.

List of attended Courses and Conferences

During the PhD, the following courses and conferences were attended:

In presence courses and conferences

1. Participation in the PhD course “Introduction to Data-driven Modeling and Machine Learning” at the University of Genoa (Total: 10 hours).
2. Seminar regarding the Electricity Market held by Terna (Total: 4 hours).
3. Participation in the PhD course “Artificial Intelligence (AI) for Energy Systems” at the Politecnico di Milano (Total: 40 hours).
4. Participation in the “23rd IEEE International Conference on Environment and Electrical Engineering / 7th I&CPS Industrial and Commercial Power Systems Europe”, held from June 6 to 9, 2023, at Universidad Carlos III de Madrid, Spain.
5. Participation in the BEHAVE 2023 conference, the 7th European Conference on Behaviour Change for Energy Efficiency, from November 28 to 29 in Maastricht, Netherlands.

Online courses

1. Participation in FIRE online courses: “RES Generation Path: REC,

PV and storage, SSPC, Biomass Generators”, “Formulas and numbers for energy conversion”, “Statistical tools in energy uses and economic-financial analysis under uncertainty conditions”, “Energy Audit in the industrial sector pursuant to UNI CEI EN 16247 3” and “Fundamentals of Energy Management” (Total: 68 hours).

2. Participation in AiCARR online courses: “REC and self-consumption: current status and future prospects”, subsequent update “Renewable Energy Communities. Decree 414/23 and GSE Application Rules”, and “Quantification and valorization of energy savings” (Total: 15 hours).
3. Participation in the Schneider Electric online course “Integrated Digital Academy: Energy Efficiency and Iproject” (Total: 25 hours).
4. Participation in the Bachelor’s degree course at the University of Genoa in Electrical Engineering “Electric Power Systems” – Prof. Marco Invernizzi (Total: 120 hours).
5. Participation in the iKn-Italy course on the Electricity Market and TIDE reform (Total: 10.5 hours).
6. Participation in the course “Smart Grids: evolution of distribution networks and energy transition” provided by CEI (Total: 8 hours).
7. Participation in the course “IoT concepts, architectural principles and maritime applications” organized by DITEN-Unige in collaboration with Åbo Akademi University (Finland) within the LeaderSHIP project (Total: 8 hours).

International and Industrial Experiences

During the PhD, the following industrial and international research experiences were carried out:

1. Six-month industrial internship at IESolutions Soluzioni Intelligenti per l'Energia s.r.l. (formerly a spin-off of the University of Genoa) under the supervision of Dr. Andrea Vinci. The collaboration focused on the research, design, and deployment of industrial energy monitoring systems.
2. Six-month international internship at WellD Sagl in Lugano, Switzerland, under the supervision of Mr. Matteo Codogno. The activity involved collaborating on the SPARKIE project within NEST - Network 4 Energy Sustainable Transition - Spoke 8, aimed at developing an advanced software architecture for the monitoring and management of RECs.