POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

Fuel Cell Modeling for Microgrid Control Applications

Tesi di Laurea Magistrale in
Automation and Control Engineering

Author: **Francesco Berruti**

Student ID: 249009
Advisor: Prof. Giampaolo Manzolini
Co-advisors: Riccardo Simonetti
Academic Year: 2025-2026

Abstract

The increasing penetration of intermittent renewable power sources necessitates the development of robust long-term storage solutions to ensure grid stability and energy security. While lithium-ion batteries are highly effective for short-term storage, long-term storage requires different solutions. Hydrogen is one of the most promising among them, due to its high energy density and scalability. This thesis focuses on modeling a hydrogen storage system, specifically the fuel cell, and integrating it into a broader control architecture based on a hierarchical Model Predictive Control (MPC).

First, an experimental campaign was conducted to characterize the fuel cell and refine thermal, voltage, and efficiency models. These revealed significant deviations from nominal manufacturer specifications, identifying a gross efficiency of 48% (excluding auxiliaries), and a peak net efficiency of 32 %. The campaign was also challenging as the fuel cell had hardware limitations, which led to frequent flooding events and caused many problems during startup.

To mitigate these issues, the real-time MPC layer was modified to incorporate logic capable of handling this unpredictable behavior, using a battery buffer to compensate for the consequences of these faults.

The algorithm was then tested experimentally, with necessary simplifications due to laboratory constraints.

Subsequently, it was validated through a year-long simulation of the Ginostra microgrid, a small town on Stromboli island. Two scenarios were analyzed: one with a smaller tank and one that enabled fully renewable generation. The comparison between the two scenarios demonstrated the effectiveness of the algorithm and highlighted the challenges related to seasonality in the context of intermittent power generation, as capacity requirements increase significantly to achieve complete energy autonomy.

Keywords: Microgrid, Hydrogen Storage, Model Predictive Control, Seasonal Storage, Fuel Cell Modeling, Energy Management System.

Abstract in lingua italiana

La crescente penetrazione di fonti rinnovabili intermittenti necessita di soluzioni robuste per l'accumulo di energia nel lungo periodo, al fine di garantire la stabilità della rete e la sicurezza energetica. Mentre le batterie agli ioni di litio sono altamente efficaci per l'accumulo a breve termine, quello a lungo termine richiede soluzioni diverse. L'idrogeno rappresenta uno dei candidati più promettenti, grazie alla sua elevata densità energetica e scalabilità. Questa tesi si concentra sulla modellazione di un sistema di accumulo ad idrogeno, in particolare la cella a combustibile, e sulla sua integrazione in una più ampia architettura di controllo, basata sul Model Predictive Control (MPC).

Innanzitutto è stata condotta una campagna sperimentale, per caratterizzare la cella a combustione e affinare i modelli termici, della tensione e dell'efficienza. Questa ha rivelato una deviazione significativa dalle caratteristiche nominali indicate dal produttore, identificando un'efficienza lorda del 48% (esclusi gli ausiliari), ed un'efficienza netta massima del 32%. La campagna ha inoltre evidenziato limitazioni hardware, che spesso causavano allagamenti e instabilità durante l'avvio.

Per mitigare questi problemi, il livello MPC real-time è stato modificato per incorporare una logica in grado di gestire questi comportamenti imprevedibili, utilizzando una batteria tampone per compensare le conseguenze di questi guasti.

L'algoritmo è stato poi testato sperimentalmente, anche se con delle semplificazioni dovute a limitazioni nel laboratorio.

Successivamente, è stato validato attraverso una simulazione di un anno della microrete di Ginostra (isola di Stromboli). Sono stati analizzati due scenari: uno con un serbatoio più piccolo, e uno con un serbatoio di capacità sufficiente da permettere una generazione completamente rinnovabile. Il confronto tra questi due scenari ha dimostrato l'efficacia dell'algoritmo, ed ha evidenziato le sfide legate alla stagionalità nel contesto della generazione intermittente, poiché i requisiti di accumulo crescono significativamente per raggiungere la completa autonomia energetica.

Parole chiave: Microreti, Accumulo a Idrogeno, Controllo Predittivo (MPC), Accumulo Stagionale, Modellazione Celle a Combustibile, Sistemi di Gestione dell'Energia.

Contents

Abstract **i**

Abstract in lingua italiana **iii**

Contents **v**

1 Introduction **1**

2 State of the art **7**

3 Modeling **9**

3.1 Models 9

3.1.1 Fuel cell system and auxiliaries 10

3.1.2 Thermal modeling 11

3.1.3 Current voltage model 14

3.2 Overview of the MG^2Lab Control Architecture 16

3.2.1 First layer 17

3.2.2 Second layer 21

3.2.3 Innovations added to implement the fuel cell 21

4 Experimental implementation **25**

4.1 The physical system 25

4.2 Experimental campaign 26

4.2.1 Characterization of thermal model 27

4.2.2 Characterization of voltage model 32

4.2.3 Efficiency curves 33

4.3 Experimental deployment 35

4.3.1 Description of the experiment 35

4.3.2 Practical implementation Challenges and workarounds 36

5 Algorithm Validation for Seasonal Storage: The Ginostra Microgrid Case **39**
5.1 Introduction 39
5.2 The Ginostra Challenge: Isolation and Seasonality 40
5.3 Simulation framework 42
5.4 Assumptions 44

6 Results **45**
6.1 Case study 45
6.1.1 Planned scenario 45
6.1.2 Ideal fully renewable scenario 48
6.2 Results of experimental implementation 51
6.3 Second layer results in simulation 55
6.3.1 Simple battery buffer 55
6.3.2 Taking advantage of the two layers 57

7 Conclusions and future developments **59**
7.1 Future developments 60

Bibliography **61**

List of Figures **67**

List of Tables **71**

List of Symbols **73**

Appendix **77**
7.2 Fault characterization 77
7.2.1 DC/DC faults 77
7.2.2 Flooding 77

1 | Introduction

There is broad scientific consensus that climate change will have devastating impacts on humanity if no action is taken [1]. In this context, it is crucial to address electricity generation, one of the biggest sources of greenhouse gas emissions: in 2024 electricity and heat production accounted for nearly 30% of total CO_2 emissions, the largest single-sector source [2].

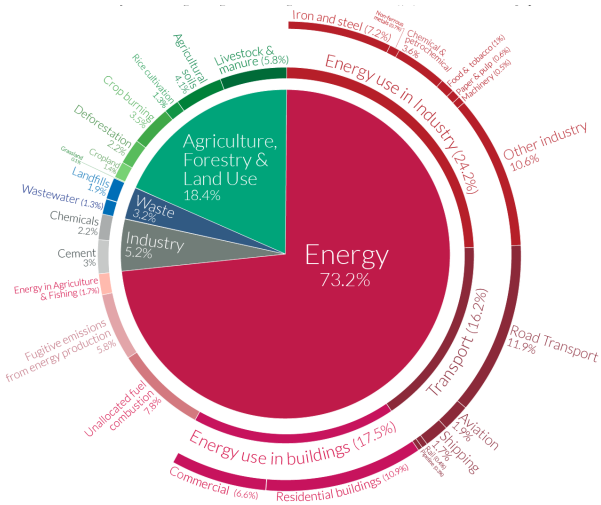

Figure 1.1: Share of global greenhouse emissions by industry

However, recent trends are encouraging as in 2024 around 80% of the growth in power generation came from renewable energy [3], indicating that renewables represent the most viable path to a more sustainable energy grid. The EU has set very ambitious goals, aiming to become the first climate-neutral continent. Most importantly, the current goal is to reduce greenhouse gas emissions by at least 55% by 2030 compared to 1990, by transitioning to renewable energy sources (RES) and improving energy efficiency across all sectors. The subsequent goal is to reach net zero emissions by 2050.

This very fast development poses many challenges in grid management. Most renewable energy sources (excluding nuclear) are intermittent and not directly dispatchable: unlike thermal power plants, they do not contribute to the grid's inertia, and any weather change can have an immediate impact on power generation. For this reason, grid operators have to increasingly rely on advanced control strategies, supported by increasingly accurate load and weather forecasting, and they are starting to investigate sector coupling techniques, where transport, heating and electricity generation can be integrated to increase grid stability.

For all these reasons it is increasingly important to develop technologies to increase grid stability, and this is where storage comes into play. The idea is to counter this volatility with a buffer capable of absorbing excess generation and releasing it when production falls

short. This capacity to smooth fluctuations in generation and demand makes storage a cornerstone of a resilient renewable-based energy system.
A distinction must be made in energy storage for these purposes:

- Short-term energy storage: it is typically sized for seconds to hours, and it is dominated by flywheels and Li-ion batteries. The former are mainly used for frequency regulation, as they are not expected to store any significant energy. The latter are used for longer time frames due to their very high efficiency and very low self-discharge, but cannot be easily scaled due to their high cost per kWh and capacity degradation, compared to other storage solutions. Lithium-ion battery energy storage increases with volume, which means it's a function of its constituent materials. Most of these are Critical Raw Materials(CRM), such as lithium, cobalt and nickel. Scaling battery storage means scaling the use of these materials linearly, which means they are not economically competitive as storage grows.

- Long-term energy storage: it is the main focus of this thesis, sized for days to months it is designed to deal with problems connected to seasonality, whether that is caused by intermittent generation or intermittent loads. This is always useful, but strictly needed when penetration of renewable energy sources is over 80% [4–7]. It is typically done with Power-to-x technologies, which transform electrical power into an energy medium and then back to electricity. Power-to-gas is one of the most promising, with hydrogen as the leading candidate. Hydrogen requires platinum (another CRM) in the fuel cells, but does not require it to scale in energy storage. There is a decoupling between power and energy, which allows for mass storage at a lower cost. This is a very important factor when comparing different scales of storage.

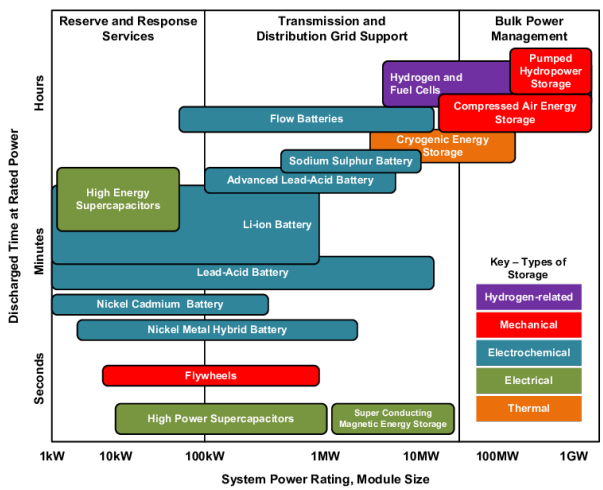

Figure 1.2: Comparison of different technologies at different storage size

The main form of long-term energy storage has traditionally been hydroelectric storage through the use of dams [8]: when energy is in excess, water is pumped into the reservoir, and when it's needed it is run through a turbine to generate electricity. Although effective and simple, this method is clearly limited in its applications, as the specific requirements to build a dam are very strict.
Many more methods are being considered to address the problem in a more versatile and distributed way, and one of them is hydrogen. The process of electrolysis involves the use of electricity to separate two molecules of water into its components:

$$2H_2O + \text{Electricity} \rightarrow 2H_2 + O_2 \tag{1.1}$$

The reverse reaction is instead oxidation, where hydrogen and oxygen react to create water and electricity.

$$2H_2 + O_2 \rightarrow 2H_2O + \text{Electricity} + \text{Heat} \tag{1.2}$$

The high energy content of hydrogen per unit mass and its ability to be compressed or liquefied make it a promising medium for energy storage, as this process can be easily repeated over time.
However hydrogen has a major drawback, which is efficiency. While Li-ion batteries have an efficiency of about $80\% - 95\%$[9] the hydrogen electrolyzer-fuel cell cycle has an efficiency of about $30\% - 50\%$ [10]. For this reason it is not economically viable unless the storage size and seasonal variation justify the energy losses.
Another form of long-term energy storage is thermal storage, but at the current stage it is typically limited to the scale of days, not entire months. It works by heating a medium, like molten salt, rocks or water, and then reusing that heat to generate electricity when needed via heat engines and turbines.

One factor to consider when comparing hydrogen to other long term energy storage solutions is the possibility of delivering it through pipelines or using it in different industries if more convenient, such as transportation or manufacturing. [11] [12].
Currently, most of the hydrogen is produced on site through the use of fossil fuels when necessary, like for the production of ammonia and many other chemicals. Hydrogen could be transported using the already existing natural gas pipeline network, and brought where necessary if it is in excess of the power grid's needs. This versatility makes hydrogen not only an energy storage system, but also a commodity that can be exchanged in real time depending on the markets.

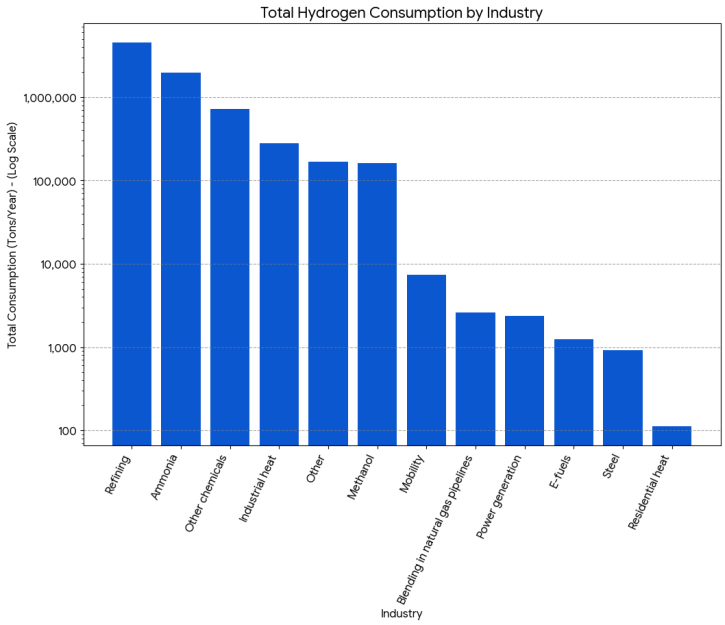

Figure 1.3: Hydrogen consumption by sector from the European hydrogen observatory

Another important factor in this rapid change is a phenomenon unique in the history of the electric grid: the introduction of prosumers. These are entities that not only draw power from the grid, but also feed their excess generation into the grid. The decentralization of both generation and installation means that the flow of power is no longer unidirectional; grids must handle generation coming from thousands of different sources, each requiring balancing, metering and compensation. A high share of prosumers has caused grid overloads in the past, leading grid operators to impose negative prices to prevent overloads [13–17].

Traditional power systems were designed around large, centralized plants and passive consumers. The growing penetration of renewable generation has fundamentally changed this paradigm, requiring flexible, distributed architectures capable of local balancing and autonomous operation. Microgrids represent a key building block of this emerging smart grid structure, as they enable local energy generation and consumption, reducing losses and increasing reliability.

Formally, a microgrid is defined by the U.S. Department of Energy as "a group of interconnected loads and distributed energy resources within clearly defined electrical boundaries that acts as a single controllable entity with respect to the grid. A microgrid can connect and disconnect from the grid to enable it to operate in both grid-connected or island-mode"[18]. The complexity of the control problem grows exponentially with the size of the electric grid considered, which is why microgrids are especially interesting for experimenting with new technologies.

The interest in microgrids goes beyond simple experimentation: the simplicity of control has pushed many studies to focus on splitting the macrogrids into multiple microgrids that can sometimes change dynamically, to simplify the management of the distributed energy sources and reduce the burden on a centralized controller [19, 20].
These setups are mostly experimental for now, but islanded microgrids are present in many remote places, where the traditional grid cannot reach. Italy is prone to many geographical challenges that impede a single, unique grid, with 19 small inhabited islands that currently mostly rely on diesel[21], and many other mountain communities. Reducing their dependence on fossil fuel has impacts not only directly on the pollution produced locally during generation, but also on all the logistics necessary to bring fuel to these remote locations. This has been accomplished before with a combination of hydrogen storage, batteries, and solar generation, as will be discussed later in this thesis.

The control of these systems is challenging, as the load must be met at all times, while optimizing for the constraints introduced by each generator, and the respective efficiency curves. This is achieved with the use of forecasting methods that are ever-improving, both for load and generation. Controller decisions must account for them, while optimizing for multiple scenarios, as every forecast has unmodeled dynamics. On top of this most control strategies implement multiple layers, working on different time scales to guarantee optimal decisions once the uncertainties realized.

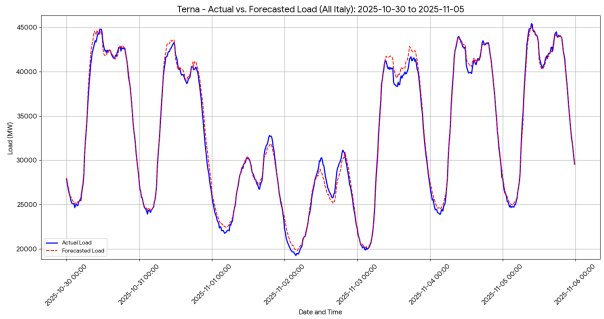

Figure 1.4: Italian load as forecasted by Terna and measured across a week in October.

Despite the ever improving quality of simulation tools it is always necessary to have experimental setups, to capture any unexpected or unmodeled dynamics. They are a crucial intermediary between theoretical research and practical implementation.
For these reasons it is important to laboratories like the one in the MG^2Lab. This is an innovative microgrid testing ground, which integrates Li-ion batteries and hydrogen storage with solar photovoltaic generators and a back-to-back (B2B) converter to simulate a load, all while being connected the main grid, in case of necessity. The MG^2Lab provides an ideal environment to investigate the role of hydrogen-based storage in microgrid control. The results obtained here can inform future strategies for large-scale deployment

of hybrid renewable systems and contribute to the broader development of smart grid architectures. The purpose of this thesis is to model and integrate into a microgrid control algorithm a hydrogen storage system, focusing on both its dynamic behavior and its role in grid balancing. Specifically, an $H_2Vector$ PET-10 fuel cell will be used to gather all the necessary experimental curves, and run all the tests in the experimental microgrid of $MG2LAB$, located in the Milan Bovisa campus. Some limitations of the fuel cell also gave an opportunity to tackle some deeper control algorithms for the dynamical behavior of the fuel cell's activity. This thesis is organized as follows:

- **Chapter 2** presents the state-of-the-art in microgrid control, analyzing both data-driven and model-based paradigms to identify the specific research gap this work addresses.
- **Chapter 3** details the models that will be implemented in the control system, or used to evaluate overall performance
- **Chapter 4** fits and validates these models using experimental data, and introduces the experimental setup for the validation of the algorithm over a short period
- **Chapter 5** introduces the case study used for validation of the algorithm over a longer simulated time span
- **Chapter 6** shows the results of the implementation of all control algorithms, both in simulation and on the experimental setup.
- **Chapter 7** concludes the thesis by summarizing the key contributions and suggesting directions for future research.

2 State of the art

Research on microgrid control has evolved rapidly in the last decade, driven by the increasing penetration of distributed renewable resources.
Modern approaches aim to achieve optimal performance despite uncertainties, combining control theory, optimization, and artificial intelligence. The literature generally converges toward hierarchical structures that balance computational efficiency and scalability, ensuring both local stability and global coordination. The main idea behind these algorithms is to have different levels of control, depending on the dynamics each layer has to capture and control, from day-ahead weather forecasting to real-time balancing of the grid.
Within this framework, numerous strategies have been proposed, ranging from purely data-driven to fully model-based algorithms.
Data-driven techniques have been applied before to the control of the microgrid operation, even though they are not the main interest of this thesis they warrant an explanation.
Tools like artificial neural networks are not structurally capable of guaranteeing constraints, such as the power balance, and are therefore not well suited for this purpose. However, they can work with other algorithms to provide forecasts of weather or load, crucial for most algorithms [22]. In spite of this, it is a highly researched topic, with focus on easing the computational load, as it increases very quickly with the network size [23]. Another important factor is the low adaptability these algorithms have to a change in the grid's composition, but this could be addressed through transfer learning [24].
Among the data driven techniques the work of Stefano Leonori et al [25]. is one of the most unique, utilizing an interpretable fuzzy inference system synthesized through a genetic algorithm, but it is strictly experimental.
Another frontier of data-based algorithms is Q-learning, where the grid management problem is formulated as a Markov decision process: states are the commitment status of each generator, the power requested from them, the load, the SOC of the batteries, and the price of electricity [26]. Particularly noteworthy is the work of Hao Zhou et al., who developed a Q-learning multi-agent algorithm where the producers and the users are agents, who aim to maximize their revenue [27].
Despite their flexibility, data driven methods are not able to give the formal guarantees

that a system like this requires. Many of these methods are not structurally able to enforce hard-coded constraints, which is a strong limitation.
This directs research, and this thesis, mostly towards model-based algorithms. These are simpler to understand and manage, as they leverage a white-box model of the system, to ensure all physical constraints are respected by design. Hybrids exist, an example is model-based convolutional neural networks (CNNs) [28], but they are still highly experimental.
Most model-based algorithms use some form of an optimization problem, where a model of the system is leveraged to minimize an economic cost function.
To this end, nonlinear models are implemented in the literature, due to their precise modeling capabilities [29]. They are much more flexible and accurate than linearized models. Usually these are however considered only after the binary commitment status variables have been decided, as the complexity of nonlinear mixed integer programs is currently too high for practical implementation. This means they can only be implemented in the real-time balancing of the system; however, in this case speed is crucial, and computational efficiency must be thoroughly considered.
For all these reasons the MG^2Lab focuses on control methods based on linearized or piecewise linear models [30–32].
This introduces some complications in modeling the systems properly, as there are very few components that are actually linear in a microgrid, but at the current stage there is no realistic path to solve nonlinear mixed-integer programs at the required speed. It is possible to use piecewise characteristic curves (such as the diesel generators), but this approach is mostly avoided due to the increasing computational complexity.

3 Modeling

3.1. Models

This chapter is meant to introduce the system-level models necessary for integration into the Model Predictive Control (MPC) framework. Due to the constraints set by a linear optimization solved in real time the complex dynamic models are actually not directly implementable, but they provide a basis for the necessary linear approximation. Other than the efficiency model, two other models are important to characterize the fuel cell behavior:

- A thermal model is important for the control problem, as the reaction is eased by temperature, and this limits the dynamic response of the fuel cell.
- A voltage-current model, to understand all of the real working points beyond the nominal declared behavior.

The chapter first provides an overview of the physical system and its auxiliaries, then details the simplified thermal, current-voltage, and efficiency models used by the control algorithm.

3.1.1. Fuel cell system and auxiliaries

The fuel cell system considered is the $H_2Vector$ PET-10, an integrated unit with a nominal output of 10.6 kW. Before modeling the stack itself, the Balance of Plant (BOP) components—the auxiliaries—must be defined, as they represent significant, nonlinear consumption.

- Constant consumption: control panels and air fans to cool the cell and feed the air needed for the reaction. Together, they always consume about 2kW
- Variable consumption: cooling system. Depending on the heat that needs to be dissipated, it either implements an on/off logic based on cooling water temperature or is simply left on, and the flow is controlled through valve opened by a P controller. This consumes about 1.5kW.

However, as it is possible to note in figure 3.1, even at a constant power auxiliary behavior and consumption change overtime. This means that it is impossible to statically map auxiliary power consumption to power consumed. Since it is not linearly related to any measured value, it was decided to simply model it as a constant 3.5 kW loss. This assumption is not significant, as power consumption changes very little with power, and even when the on off logic is implemented it is generally inefficient to run the fuel cell at less than 3.5 kW, as it switches from providing to requiring power. This means that the fuel cell should, optimally, always be run at its highest output, where net efficiency is highest.

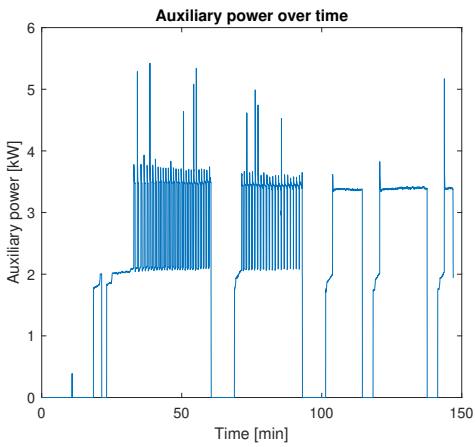

Figure 3.1: Power consumed by auxiliaries over time, in a test ran at constant power. It is possible to note the switch in temperature control strategy. Every drop in power is fuel cell downtime

3.1.2. Thermal modeling

Fuel cells require very careful thermal modeling, as temperature is crucial in controlling the reaction and its efficiency. Pushing too much current at low temperatures can damage the fuel cell overtime, shortening its lifespan. For this reason, many 1D and 2D temperature models have been developed[33–35].
However, the main interest of this thesis is not to control the reaction happening inside the fuel cell, but rather to control the power setpoints to follow, assuming the lower-level control to be managed internally. For this reason, only a lumped model will be considered. The producer enforces the following rule:

$$I_{max}(T) = \begin{cases} 2T_{FC}, & \text{if } \; T_{fc} \leq 55\,^{\circ}\text{C} \\ 230 & \text{else} \end{cases}$$

A thermal model was developed from experimental data to precisely characterize this constraint. The fuel cell stack is enclosed in a cabinet and cooled by a combination of forced air and an auxiliary water-cooling system. Both the cabinet air temperature and the cooling water temperature were measured and included in the model identification process. Over longer time scales (on the order of several hours), the auxiliary cooling unit itself exhibits thermal inertia; its temperature was therefore incorporated in the model to improve prediction accuracy during optimization, although these slow dynamics are largely outside the control horizon of interest. While estimating the model the cooling water temperature and the fuel cell temperature were assumed to be approximately equal.

$$\begin{cases} T_{FC}^{t+1} = T_{FC}^{t} + T_s[L_1(T_{cab}^{t} - T_{FC}^{t}) + L_2(T_{aux}^{t} - T_{FC}^{t}) \cdot aux_{on}^{t} + L_3 P_{fc}^{t}] \\ T_{cab}^{t+1} = T_{cab}^{t} + T_s[L_4(T_{FC}^{t} - T_{cab}^{t}) + L_5(T_{amb}^{t} - T_{cab}^{t})], \\ T_{aux}^{t+1} = T_{aux}^{t} + Ts[L_6(T_{FC}^{t} - T_{aux}^{t}) \cdot aux_{on}^{t} + L_7(T_{amb}^{t} - T_{aux}^{t})]. \end{cases}$$

Here T_{FC}, T_{cab},T_{aux} are respectively the temperature of the fuel cell, the temperature of the cabinet and the temperature of the auxiliary unit, while aux_{on} is a boolean parameter, that was set to zero or one depending on whether the cooling auxiliaries were off or on. T_s is the sampling time of the system. P_{fc} is the power generated by the fuel cell, the only energy input of the system. The auxiliary's behavior changed with power and time, as mentioned earlier.
However, the thermal model cannot be used directly in the optimization problem, as it is strictly nonlinear due to the on/off behavior of the auxiliaries. Nevertheless it can be used to derive a linear approximation.

Firstly, since the optimization problem is formulated using a power balance equation, a linear relationship between current and power had to be considered to reformulate this constraint.
The linearization was applied at low power, as this is where the constraint must be enforced; higher power levels can only be reached when temperature is above 55 °C, 110A corresponds to approximately 6kW. The constraint can therefore be rewritten as

$$P_k \leq 2fT_k \tag{3.1}$$

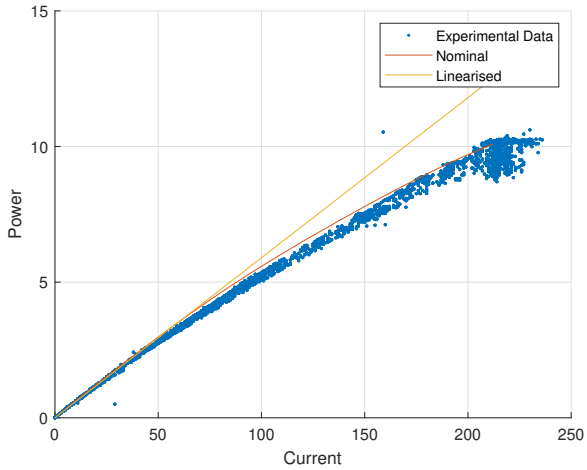

Figure 3.2: Linear approximation of the current/power relationship

After this issue was resolved, the problem had to be reformulated to account for the nonlinearities of the temperature model. A simple ramp-up constraint could have been implemented; however, this method lacked memory, preventing the fuel cell to start as already warm and thus bypass the initial thermal transient. This constraint would be enforced at all times, regardless of the state of the fuel cell, slowing down the dynamic response of the system.
For this reason, a first-order dynamic system was introduced to approximate the constraint. This was introduced alongside a second always-active constraint, $P_{fc} \leq 10.6$.
The parameters were identified using a prediction error minimization, by comparing it with the thermal model developed earlier. The model parameters were identified using a prediction-error minimization approach, by comparing the dynamic model with the previously developed thermal model under maximum allowable power. The estimation minimized the deviation between the two constraints over time.

The resulting formulation can be expressed in linear form by assuming that the derivative remains constant over each sampling interval:

$$
\begin{aligned}
P_k &\leq 2f(T'_k + T_{amb}) \\
T'_{k+1} &= T'_k + T_s \dot{T}'_k \\
\dot{T}'_k &= A' T'_k + B' P^k_{fc} \\
A' &\leq 0 \\
B' &\geq 0
\end{aligned}
\tag{3.2}
$$

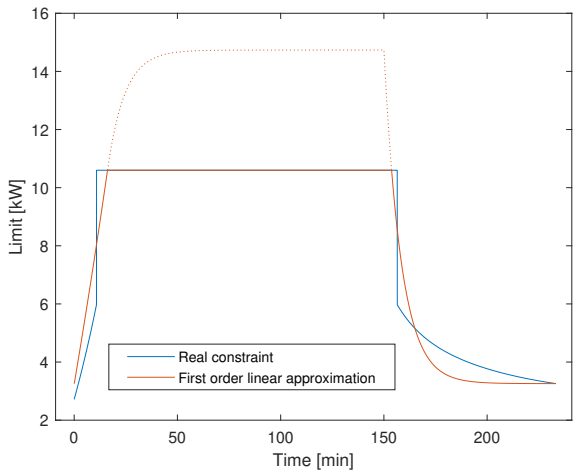

Figure 3.3: Comparison between the first order transfer function constraint and the real one, both when the power applied is the highest allowed by the constraint and when the fuel cell is turned off

This fictitious temperature is designed to reach a higher steady-state value, and it is replaced by the fixed constraint when the latter becomes stricter.
This is a rough approximation and is not intended to replace the internal thermal control enforced by the fuel cell. Instead, it helps reduce the computational burden of the second layer, which has to balance the grid despite the model limitations. The steady state also changes, depending on the input power, and can get below 10.6 kW. This is not, however, a big issue, as the fuel cell always has some inertia, and therefore gradual setpoint changes are preferred.
The main limitation of this method is the lack of measured feedback, as it has no direct connection with the measured temperature, and it is a completely synthetic constraint. While the model remains stable, it reduces some of the advantages of the MPC framework.

3.1.3. Current voltage model

To further characterize the stack, an analysis of the voltage current characteristic was carried out, and a model was estimated to compare it to the nominal one. By extrapolating a model from the data is was possible to characterize the entire polarization curve, not just the points gathered experimentally.
A detailed description of the model can be found in [36].
The main output voltage is the sum of the elements reported in table 3.1:

Table 3.1: Voltage Component Descriptions

Voltage Component	Symbol	Description
Reversible Voltage	V_{rev}	Theoretical maximum voltage from chemical potential
Activation Voltage	V_{act}	Voltage needed to overcome reaction energy barriers; dominates at low current
Ohmic Losses	V_{ohm}	Linear loss from inherent resistance in electrodes and membrane.
Concentration Losses	V_{conc}	Loss from mass transport limitations (reactants in, water out); dominates at high current.

Therefore, the overall model equation is:

$$V_{stack} = V_{rev} - V_{act} - V_{conc} - V_{ohm} \tag{3.3}$$

For the derivation of each equation, refer to [36]. This can be modeled as an equivalent circuit model:

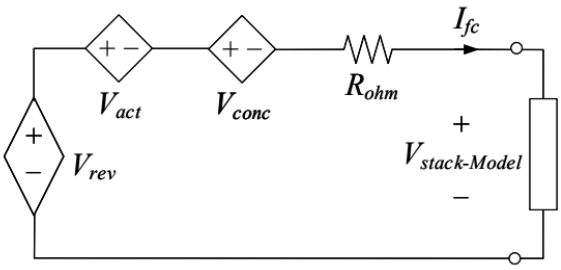

Figure 3.4: Equivalent circuit of the Fuel cell stack

Where:

$$V_{rev} = N_s[1.229 - 4.3 \times 10^{-5} T_{fc,K} \ln(\frac{\mathrm{P_{H_2}P_{O_2}}}{\mathrm{P_{H_2O}}})]$$
$$V_{act} = N_s\{(a_0 + a_1 T_{fc,K}) + (b_0 + b_1 T_{fc,K}) \ln(I_{fc})\}$$
$$V_{conc} = N_s(m_0 + m_1 T_{fc,K}) e^{nI_{fc,k}}$$

Here N_s is the number of cells in the stack, $P_{[\ ]}$ is the partial pressure of species [ ] (in this case 1 for hydrogen, 0.21 for oxygen and 1 for water), $T_{fc,K}$ is the fuel cell temperature in Kelvin, I_{fc} is the current, while $a_0, a_1, b_0, b_1, R_{ohm}, m_0, m_1, n$ are the modeling parameters. Note that the choice of the partial pressures implies an assumption of purely liquid output water and not humidified hydrogen, which is not always correct. However, the difference in voltage introduced by this assumption is negligible, as it is scaled by a term in the order of 10^{-5}.

This model will not be implemented into the MPC framework, as it is not necessary and strictly nonlinear, but it is useful to characterize the fuel cell and compare it with nominal data. It allows to extrapolate a more general characterization from the experimental data.

3.2. Overview of the MG^2Lab Control Architecture

The laboratory implements a two-layer control structure, with two separate controllers operating at different time scales, as described in [30–32]

- First layer: A MILP problem is solved with a horizon of 24 hours and time intervals of 15 minutes. Since it relies on forecasts of both weather and load it requires a stochastic component, optimizing across multiple scenarios. The resulting problem is computationally expensive, and is only suitable for high-level day-ahead scheduling. Owing to its long prediction horizon it is able to plan ahead and decide resource allocation throughout the day, while minimizing economic cost.
- Second layer: Operating with a much finer time resolution, a second optimization problem is solved, with a time horizon of 15 minutes and a resolution of 1 minute. This preserves the commitment variables decided by the first layer, and balances the power in real time, based on the realized values of uncertain variables. Since it maintains the commitment statuses, all the variables are continuous and linear, making the problem simpler and faster to solve. The higher temporal resolution also allows for more accurate models, capturing faster dynamics not noticeable in the first layer. The main focus is to balance the grid and respond to unforeseen disturbances.

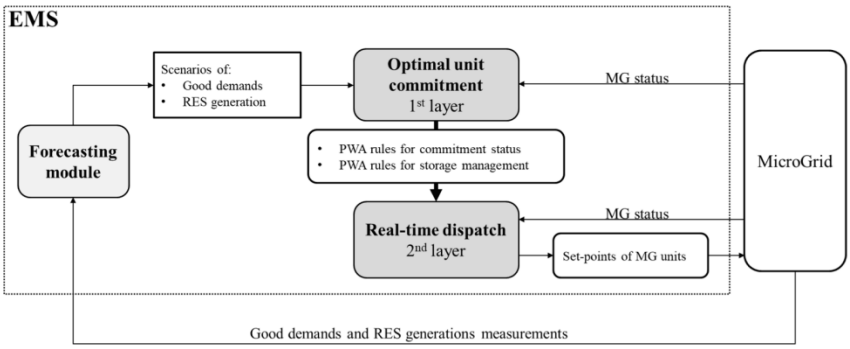

Figure 3.5: Hierarchical structure of the EMS and information flow.

3.2.1. First layer

Since the first layer is focused on optimizing the resources throughout the day, it is more sophisticated and computationally intensive than the second layer, and requires some careful characterization. This section summarizes its main formulation for context; more detailed information can be found in [30–32]

The optimization variables are the power outputs (or consumptions) of each component, expressed as a fraction of their rated power, and their corresponding binary commitment statuses:

$$\begin{aligned} x_i^t &\in [0,1] \\ z_i^t &\in \{0,1\} \\ \\ z_i^t &\geq x_i^t \\ x_i^t &\geq P_i^{min} z_i^t \\ \\ \forall i &\in \{fc, f, batt, el\} \\ \forall t &\in [0, t^{horizon}] \end{aligned} \tag{3.4}$$

Here x_i^t denotes the normalized power output of component i at time t, and z_i^t its binary commitment variable. The last two constraints guarantee logical consistency between the variables: a component can only output (or consume) power if its commitment variable is set to one, and vice versa.

The fundamental constraint enforces the instantaneous power balance across the system:

$$P_f^t + P_{solar}^t + P_{fc}^t = P_{batt}^t + P_{el}^t + P_{load}^t + P_{grid}^t \tag{3.5}$$

Here, P_f is the power generated from fossil sources, P_{fc} is the power generated by the fuel cell, P_{batt} is the battery power (positive if charging negative if discharging), P_{el} is the electrolyzer consumption, P_{load} is the total consumption and P_{grid} is the power exchanged with the external grid.

This global constraint couples all components, while each device is subject to its own operational limits.

Energy storage systems. For both the battery and the hydrogen storage unit, the state of charge (SOC) is constrained between its minimum and maximum limits. The evolution of the SOC over time is modeled as:

$$\begin{gathered} SOC_{i,\min} \leq SOC_i^t \leq SOC_{i,\max} \\ SOC_i^{t+1} = SOC_i^t + c_i P_i^t, \\ \dot{SOC}_i^t \leq \dot{SOC}_{i,\min}, \\ \forall t \in [0, t^{horizon}],\ \forall i \in \{\text{Hstorage}, \text{batt}\}. \end{gathered} \tag{3.6}$$

Generators The power output of each generator is limited by its dynamic ramping constraints:

$$\begin{gathered} P_i^{t+1} \leq a_i P_i^t + b_i, \\ \forall i \in \{fc, el, f\},\ \forall t \in [0, t^{horizon}]. \end{gathered} \tag{3.7}$$

These constraints are coupled with a startup cost, to ensure that the physical generators are not switched on and off, and the startups are limited. This ensures that any unmodeled startup dynamic present in the physical generators affects the result as little as possible, and the generators are able to follow the given setpoints accurately.

Cost The objective function is formulated to minimize total operational cost. Each component either produces or consumes a resource, to which an economic cost is assigned according to system priorities, which forms the objective to be minimized. Additional terms account for maintenance and operational costs.

Innovations added to implement the fuel cell

The optimal behavior would be to set a priority order, where solar energy is used whenever available, if it is not enough then batteries should intervene and then the fuel cell, only if the batteries are not available anymore. If this is also depleted or not available then the diesel generation should turn on. Vice versa, when there is extra energy, it should go to the batteries first, and any energy that is in excess to that should go to the electrolyzer. This should be an emergent property if the optimization problem.
The main idea behind the implemented cost function is to assign an economic value to every resource, so as to ensure that the set priorities are respected. To determine the value given to the hydrogen, it is important to consider the cost associated to fuel consumption:

$$C_t^{ud} = W_0 P\ m_f \tag{3.8}$$

where P is the price of fuel in €/l and m_f is the fuel intake in liters per sampling time, in this case 15 minutes, and W_0 is a weighting coefficient, useful to shape the relative economic contribution.
In our case hydrogen does not have a direct economic value, as it is produced exclusively from renewable sources and the study does not consider the possibility of selling it or using it for other purposes. For this reason, an artificial value has been assigned to it, using an equivalent-energy valuation.

$$C_t^{ud} = W_1 \; PM_H \frac{LHV_H}{LHV_f} \tag{3.9}$$

The price is still the one of fuel in €/l, M_H is the total stored volume of hydrogen in liters, and LHV_H and LHV_f are respectively the lower heating values of hydrogen and the considered fuel in kWh/l. W_1 is simply a weight sized so that $W_1 \ll W_0$.
The economic value of hydrogen is the price of the amount of fuel that would provide the same energy, scaled by an empirically chosen factor, to ensure hydrogen is always preferred to fuel.
The same can be done with the batteries, but with a lower multiplier $W_2 < W_1$ so that they are preferred to hydrogen. However, the choice of this weight requires careful tuning: if W_2 is set too low the system may avoid charging the batteries in favor of hydrogen storage, which is energetically inefficient. In the extreme case, the controller may even choose to discharge the batteries to fill the hydrogen storage tank. Assigning the weights directly to energy simplifies this trade-off, as the cost function directly accounts for the corresponding efficiency losses. For energetic consistency it is sufficient that:

$$W_1 \cdot \frac{\eta_{el}}{\eta_{batt}^{char}} < W_2 \leq W_1 \cdot \frac{\eta_{batt}^{disch}}{\eta_{fc}} \ll W_0 \tag{3.10}$$

where $\eta_{[\;]}$ is the efficiency of component [ ], with a distinction between charging and discharging efficiency for the battery.
This balance ensures that it is always better to both charge and discharge the batteries: the lower bound ensures batteries are not ignored to power the electrolyzer, while the higher bound ensures batteries are discharged first.

Symbol	Meaning	Typical value
W_0	Fuel cost weight	1.00
W_1	Hydrogen storage value weight	0.15
W_2	Battery storage value weight	0.10

Table 3.2: Weighting factors used in the cost function.

As an example,table 3.3 and 3.4 respectively report the value increment or cost that the optimization problem computes, considering it has 10kWh to store away or generate from the different sources. they keep efficiencies into account, and considers a 1.5€/l diesel cost. Auxiliaries are not considered in this computation, as their importance depends on power. Regardless, since the auxiliaries of the fuel cell are much higher, the incentive to use hydrogen is lower, not changing the results.

Storage method	**Value[€]**
Battery	0.13
Hydrogen	0.11

Table 3.3: Value of storing 10 kWh

Energy source	**Cost[€]**
Diesel	3.5[a]
Battery	0.15
Hydrogen	0.438

[a] Assuming a 0.4 efficiency from the diesel generator

Table 3.4: Cost of generating 10 kWh.

The remaining essential elements of the model are the efficiency curve, approximated as linear, and the relationships defining power, hydrogen flow, and hydrogen storage.

$$P_{fc}^{t} = x_{fc}^{t} P_{fc}^{max} - 3.5 \tag{3.11}$$

$$H_{flow,fc}^{t} = q x_{fc}^{t} P_{fc}^{max} \tag{3.12}$$

$$SOC_{Hstorage}^{t+1} = SOC_{Hstorage}^{t} + \frac{H_{flow,el}^{t} - H_{flow,fc}^{t}}{H_{storage}^{max}} T_s \tag{3.13}$$

In the case of a non-islanded microgrid, where the power deficit is compensated by the network, the cost is moved from equivalent fuel to kWh directly. This implies the removal of the scaling factor $\frac{1}{LHV_f}$, and the price becomes the one of electric current, in €/kWh.

3.2.2. Second layer

Once the commitment variables are determined by the first layer, the second layer optimizes the power flows in real time while maintaining the operational status previously assigned to each generator. This is enabled by the batteries, which act as a buffer to absorb deviations as uncertainties are realized. The optimization problem retains the same structure as the first layer, with an added cost term penalizing deviations from the planned states of charge of the battery and hydrogen storage system. The models used here are generally more detailed, capturing faster dynamics.

3.2.3. Innovations added to implement the fuel cell

The main difference is that at this level the temperature constraint can be addressed, as the time step is refined enough to capture it accurately. The control strategy implemented to manage this constraint could also compensate for a secondary issue: hardware and software imperfections in the specific fuel cell used. Power outages and fluctuations on the order of several kW frequently occurred during startup due to faults in the onboard DC/DC converter, and the unit often experienced flooding, requiring periodic shutdowns for drainage. More detailed explanations of such issues can be found in the appendix.
These disturbances had to be mitigated by a control layer, since the grid could not be expected to handle such unpredictable behavior, nor could they be easily forecast, especially within a linear framework.

Battery buffer The first control strategy tested introduced a power mismatch in the MPC, representing the difference between the power expected from the fuel cell by the grid and the power actually delivered. This allowed for a battery pack to be installed in parallel to the fuel cell, compensating for all its deficiencies. The battery would then be recharged when possible. For safety reasons the fuel cell was already equipped with a supercapacitor, but this device could provide high power at very low energy capacity.

New variables were defined, to account for the difference between the actual power delivered by the fuel cell and the one expected by the grid, and for the separate behavior of the auxiliary battery pack:

$$
\begin{aligned}
P_{fc}^{exp,t} &= x_{fc}^{exp,t} P_{fc}^{max} - 3.5 \\
P_{batt}^{em,t} &= x_{batt}^{em,t} P_{batt}^{em,max} \\
0 \leq SOC_{batt}^{em,t} &\leq 1 \\
SOC_{batt}^{em,t+1} &= SOC_{batt}^{em,t} + \frac{P_{batt}}{E_{batt}^{max}} T_s \\
P_{fc}^{t} &= x_{fc}^{t} P_{fc}^{max} - 3.5 \\
P_{fc}^{t} - P_{batt}^{em} &= P_{fc}^{exp,t} \\
x_{fc}^{t}, x_{fc}^{exp,t}, x_{batt}^{t} &\in [0,1] \\
z_i \geq x_i,\ \forall i &\in [fc, batt^{em}] \\
\forall t &\in [0, t^{horizon}]
\end{aligned}
\tag{3.14}
$$

The only important thing to note about the commitment variable constraints is that the fuel cell's binary status is shared by P_{fc}^{exp}. In other words if the power is expected to come from the fuel cell the unit must be switched on. Otherwise the formulation is equivalent to allowing the grid to directly draw power from the auxiliary battery.
The reverse, however, does not necessarily hold: the system may decide to recharge the auxiliary battery when it is convenient to do so, even if the grid does not request fuel cell operation. These new variables, however, are only implemented in the second layer, that has no authority over the commitment variables.
The power balance constraint was then modified to reflect the introduction of the new variables:

$$
P_f^t + P_{solar}^t + P_{fc}^{exp,t} = P_{load}^t + P_{el}^t + P_{batt}^t \tag{3.15}
$$

The structure is identical to the previous formulation, with two key differences:

- P_{batt} No longer represents all the batteries anymore, as a portion is reserved to support the fuel cell
- P_{fc}^{exp} is the power expected from the fuel cell, rather than the power actually delivered.

The cost was changed in the following way, in addition to the base optimization objective:

$$
\begin{aligned}
&\min_{U}\ C \\
&C = C_N - W_3 P \frac{E_{batt}^{em,t}}{LHV_f}
\end{aligned}
$$

Here C_N is the cost of the base optimization objective, $E_{batt}^{em,t}$ is the energy stored in the auxiliary battery attached to the fuel cell in kWh, LHV_f is the lower heating value of the fossil fuel considered in kWh/l, and W_3 is a scaling value.
Following the same logic as before, it must hold that:

$$W_0 \gg W_3 > \max\left(W_1 \cdot \frac{\eta_{batt}^{disch}}{\eta_{fc}} \quad , \quad \frac{W_1}{\eta_{batt}^{char} \cdot \eta_{fc}}\right) = \frac{W_1}{\eta_{batt}^{char} \cdot \eta_{fc}} \tag{3.16}$$

This guarantees that, regardless of fuel cell efficiency, it is never economically optimal to draw energy from the battery, thus ensuring it will be discharged only if necessary. The $max()$ operator ensures that the fuel cell is also used to recharge the battery. If the efficiency of the fuel cell is not modeled as constant, the minimum value is picked.
The optimization variables are the standard ones, once x_{fc}^{exp} and $x_{batt}^{em,t}$ are factored in.

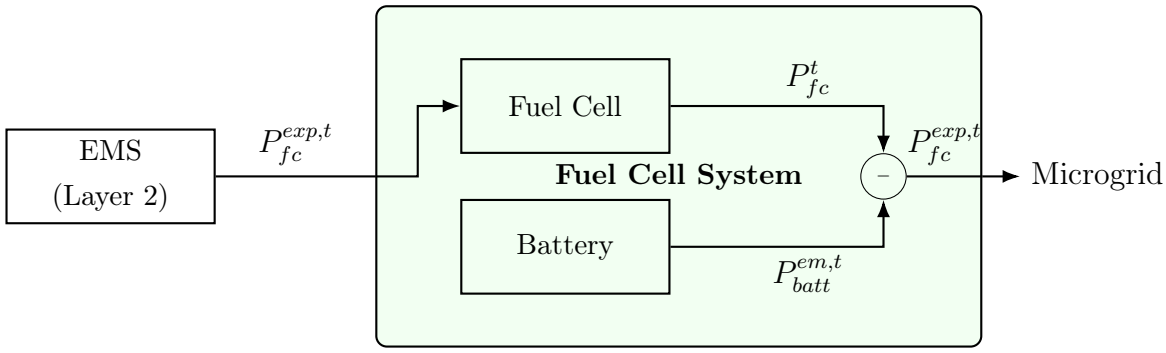

Figure 3.6: Conceptual schematic of the second-layer 'battery buffer' strategy. The Energy Management System (EMS) issues a power reference $P_{fc}^{exp,t}$ to the fuel cell system. This demand is met by the combined action of the actual fuel cell (P_{fc}^{t}) and the dedicated auxiliary battery ($P_{batt}^{em,t}$), which actively compensates for internal fluctuations, startup delays, and flooding-induced power drops.

Taking advantage of the two layers architecture The MPC strategy works well under most operating conditions; however it cannot guarantee the battery will stay charged. Its commitment variable is decided by the first layer, and it charges the battery only when excess generation is available beyond the grid demand. In the event of frequent floodings, the energy required by the system might increase significantly.
This happens because the second layer acts reactively: it cannot decide to turn components on or off; it can only react within the constraints defined by the first layer. This reduces computational effort, but it also limits the controller's ability to respond once the uncertainties are realized and the faster dynamics are considered.
To reduce the energy needed at startup the output of the first layer is automatically

adjusted to set the commitment variable of the fuel cell to 1 one time step before it is actually supposed to turn on. As a result, when the grid requires the fuel cell to provide power it will already be at operating temperature, and the battery can be recharged with any surplus energy. Once the battery is fully charged, the remaining available power can be used to drive he fuel cell's auxiliaries, allowing it to stay warm regardless of the size or state of charge of the batteries. This approach allows a smaller battery pack to be installed alongside the fuel cell, or enables the existing battery pack to be maintained at a higher state of charge in case of future unexpected disturbances. However, hydrogen consumption will increase, as the fuel cell must stay on -and power its auxiliaries- for 15 minutes longer than expected. This effect should be taken into account with a startup penalty in the first layer optimization. Notably, the MPC algorithm naturally selects this strategy given a window long enough, since the energy stored in the battery has a higher weight in the cost function. However, this only holds true if it is able to choose its own binary variables, which is not an option with the fast, real-time control layer.

4 | Experimental implementation

4.1. The physical system

The system considered is a PET-10 fuel cell energy system, manufactured by $H_2Vector$. It has a nominal power of 10.6 kW, with a nominal maximum hydrogen flow rate of 120Nl/min. The system integrates a NEDSTACK FCS 10-XXL PEM fuel cell stack, composed of 75 fuel cells in series, providing an open circuit voltage of 72.8V and a maximum current of 230A.

Figure 4.1: The PET-10 fuel cell energy system

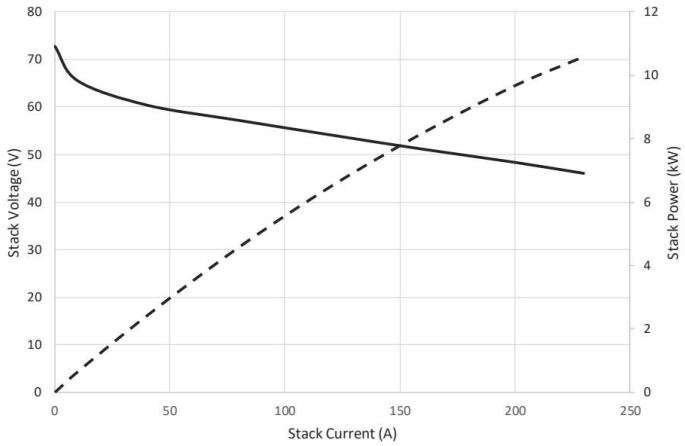

Figure 4.2: Nominal Power-Voltage-Current of the fuel cell stack

It is a unit intended for medium scale residential applications, and its losses can be used for central heating.

The system has auxiliary units for cooling, venting, and other functions, which account for a maximum consumption of 3.5 kW. This means the maximum output power it can

provide continuously is 7.1 kW.
This was not included in the efficiency computations, as the behavior of the auxiliaries has been estimated experimentally, and will be discussed later. The nominal hydrogen consumption was estimated using Faraday's law.

$$\dot{m}_{\mathrm{H}_2} = \frac{I_{fc}}{2F} \tag{4.1}$$

where I is the current, F is the Faraday constant $F = 96485\ \frac{C}{mol}$ and $\dot{m}_{\mathrm{H}_2}$ is the molar flow rate in $\frac{mol}{s}$.

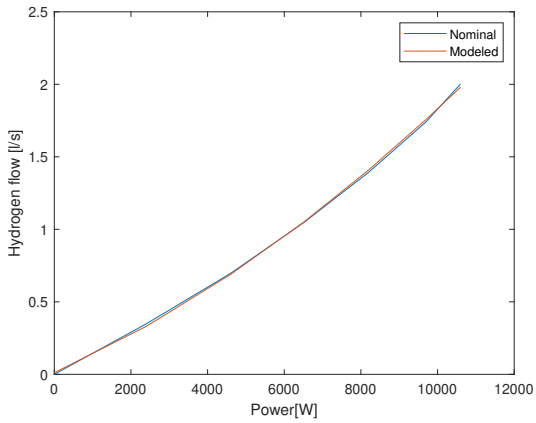

Figure 4.3: Nominal hydrogen consumption

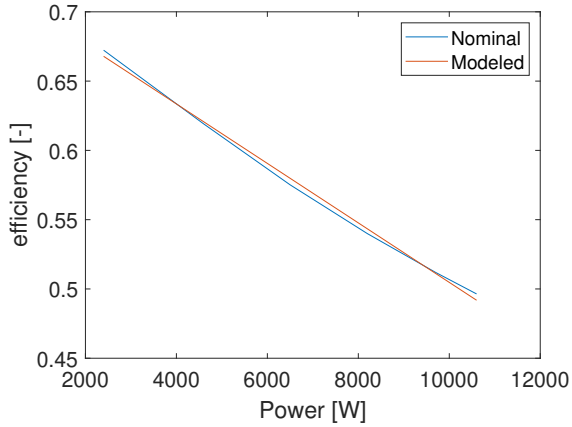

Figure 4.4: Nominal efficiency

Results show that the nominal consumption curve is quadratic ($R^2 \approx 0.9995$) , meaning that the efficiency is approximately linear($R^2 \approx 0.9966$). This will be later compared with the experimental one, which is significantly lower.

4.2. Experimental campaign

The fuel cell had many implementation issues. For this reason many of the trials were aimed at identifying and troubleshooting these issues, rather than characterizing the different models.
Throughout the entire campaign the fuel cell would fail to drain, and for this reason it often flooded and had to turn off to get rid of the accumulated water. This made the identification of many parameters, especially for the thermal model, more complicated. On top of this the on-board DC/DC was badly programmed, causing a power fluctuations of many kW on startup. A more precise characterization with examples is in the appendix.

As a consequence many of the tests are segmented, with power fluctuations hard to simulate or keep into account, and not really important to characterize the nominal behavior. These issues were solved right before the experimental implementation of the algorithm, which left little time for nominal characterization.
Despite all of these imperfections it was possible to accumulate enough data (1500 minutes) to characterize the fuel cell, even if a relevant part of it was either at high power or in downtime, causing a possible imbalance in the results.
All communication with the fuel cell was managed through the programmable logic controller (PLC) supervising the entire microgrid. It handled communication and data acquisition, making sure that all the data was properly logged for later analysis. During the experiments all the power produced was fed into the main grid of campus Bovisa.

4.2.1. Characterization of thermal model

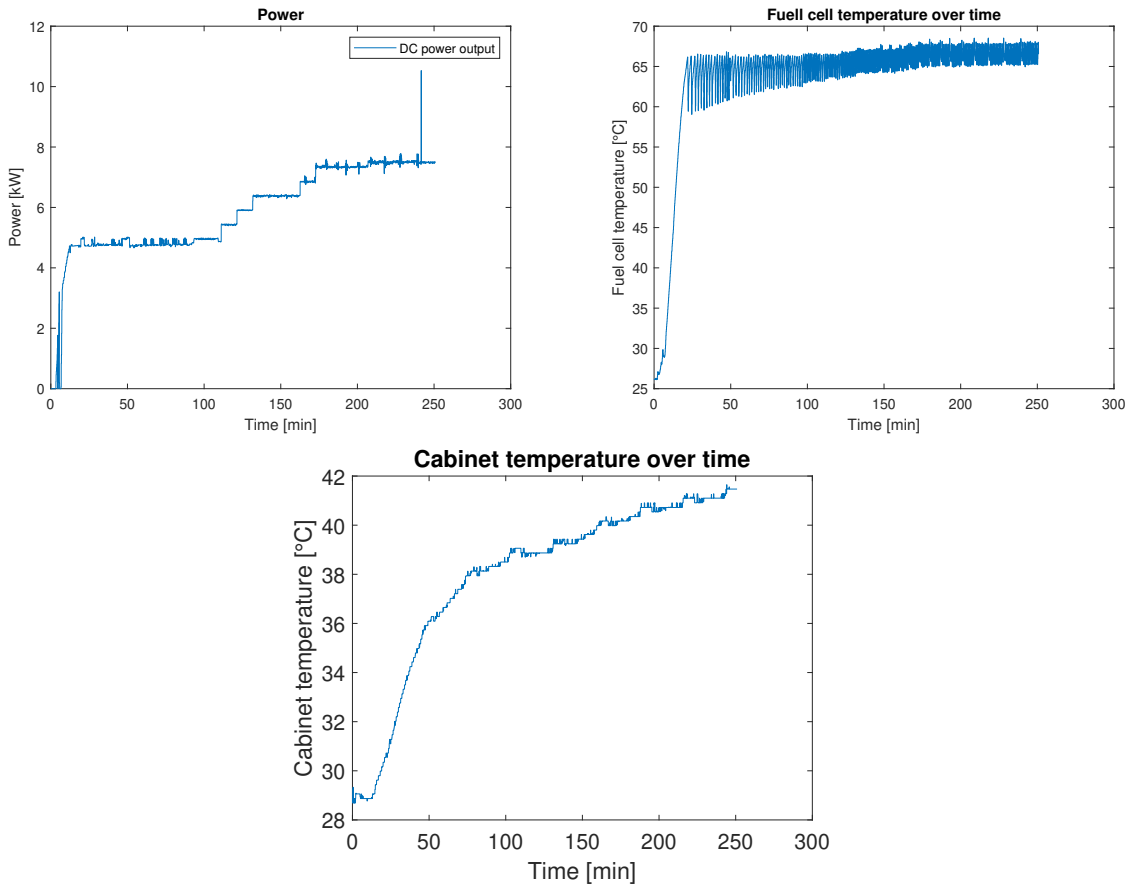

Figure 4.5: Power and temperature curve of the test used to characterize the fuel cell

The test used to characterize the thermal model was the one in 4.5.
The setpoints to follow were empirically set to:

Setpoint[kW]	Time [minutes]
4.7	111
5.4	10
5.9	10
6.4	30
6.8	10
7.5	80

Table 4.1: Setpoint of the test considered

The reason why this test was chosen to characterize the thermal model was the clear temperature curve over time, showing not only the transient of the fuel cell of the cabinet reaching steady state but also the cooling auxiliary. This can be inferred by the fuel cell's temperature, which reaches a constant maximum after 150 minutes, as the effectiveness of the cooling system reaches a steady state. Moreover, the cooling strategy was maintained the same throughout the entire test, which simplifies the model, and no floodings occurred (one of three such tests out of sixteen).
By using a prediction error minimization algorithm the parameters L_1 to L_6 were estimated, and an overall thermal model was estimated. A sequential estimation method was adopted to improve convergence: first L_4 and L_5 were estimated using the measured fuel cell temperature as a fixed input. Next, the cabinet temperature was fixed to its measured value, and L_1 and L_3 were estimated from the initial transient when auxiliaries were off. With these initial estimates the remaining parameters were computed on the entire dataset.
This procedure was repeated iteratively until convergence. The resulting values are reported in table 4.2

Table 4.2: Parameter values for L_1 through L_7.

Parameter	Value
L_1	0.0024
L_2	0.0219
L_3	0.0294
L_4	$2 \cdot 10^{-4}$
L_5	$4 \cdot 10^{-4}$
L_6	0.0084
L_7	0.0402

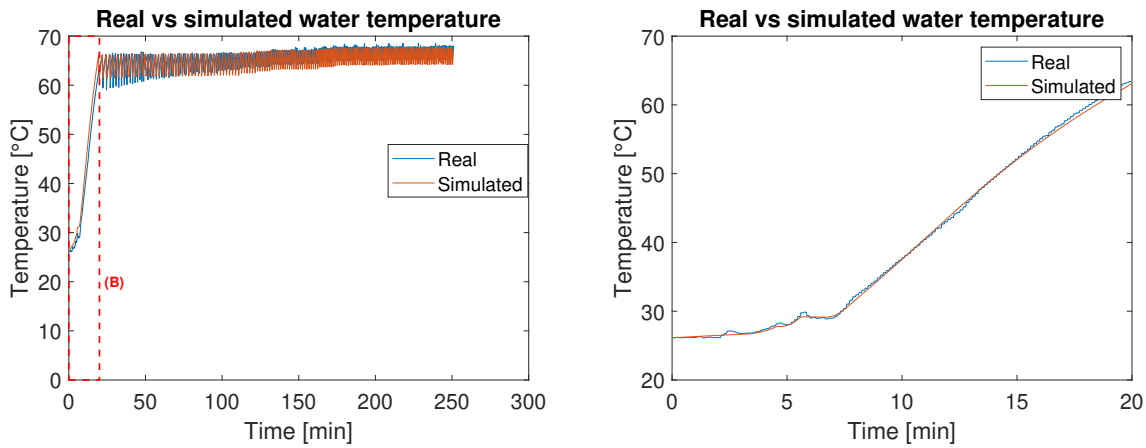

Figure 4.6: Comparison between real and simulated water temperature, with zoom on the initial transient (B)

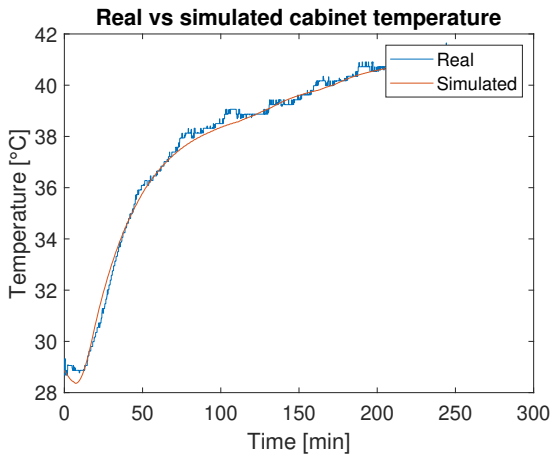

Figure 4.7: Comparison between real and simulated cabinet temperature

Temperature	R^2
Fuel cell (transient)	0.99
Fuel cell(entire test)	0.93
Cabinet(entire test)	0.97

Table 4.3: Accuracy of results on the test they were fit on

The results accurately capture both transient and steady-state behavior. The model was used to study the overall fuel-cell behavior, given the limited availability of experimental data.

The focus was on fuel cell temperature transient, as that is the one that is most interesting for control, and that is reflected in the results. To ensure it was not a case of overfitting the same metrics were measured across different tests:

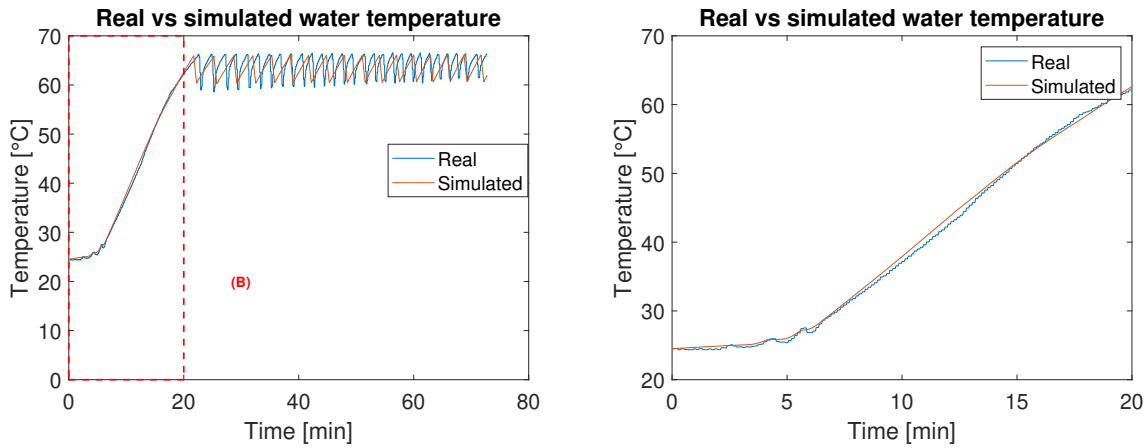

Figure 4.8: Comparison between real and simulated water temperature, with zoom on the initial transient (B) on a different test

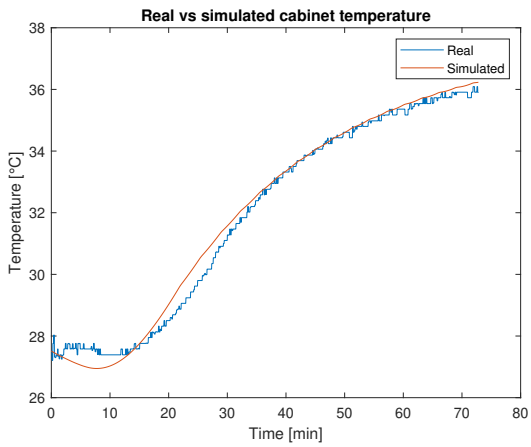

Figure 4.9: Comparison between real and simulated cabinet temperature on a different test

Temperature	R^2
Fuel cell (transient)	0.9970
Fuel cell(entire test)	0.967
Cabinet(entire test)	0.9886

Table 4.4: Accuracy of results on a different test

4.2.2. Characterization of voltage model

This model is not dynamic, but depends directly on measured values, point by point. Despite its high non linearity this eases the parameter identification procedure, as not only multiple tests can be chained together with no issues, but faults are also less of a problem, once cleaned of clear outliers caused by them. The same method as the thermal model was used, but it was not necessary to use any iterative procedure, thanks to the static nature of this problem.
All of the tests were used for the identification of the parameters of the voltage model, to use as many combinations of voltage, current and temperature as possible.

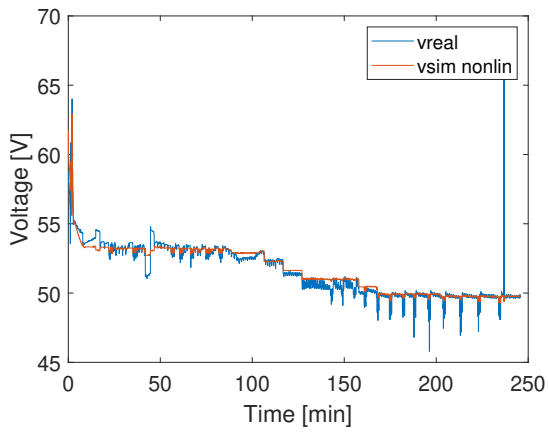

Figure 4.10: Results over one test

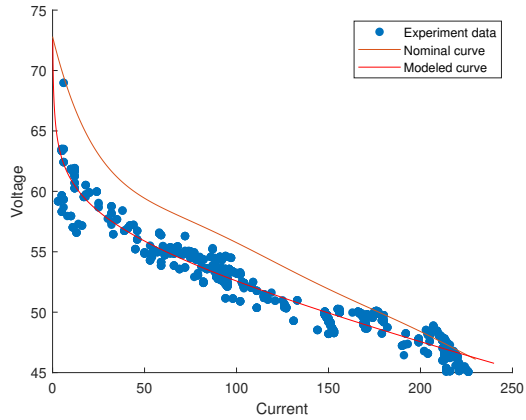

Figure 4.11: Comparison with experimental data and nominal V-I curve at 62 °C

Over the 1500 minutes of data available the measure R^2 is 0.95. The voltage is steadily below the nominal one, matching the efficiency results, since a higher current is necessary to achieve the same power. Note that at high power the two curves seem to merge, as is reflected in the efficiency computation. That is also the area where the fuel cell will be used the most, as the consumption of auxiliary components is only loosely related to power.
It is also clear that the relation between temperature and voltage is very small, as both a_1 and b_1 are approximately zero. Only m_1 carries some contribution, as it is related to the

Table 4.5: Parameter results

Item	Value
a_0	0.3322
a_1	-8e-06
b_0	0.0309
b_1	-5e-07
R	4.5644e-04
m_0	1.1299
m_1	-0.0034
n	-0.9205

concentration losses: viscosity decreases with temperature and it is an important factor in the transportation of gasses to the reaction site.

4.2.3. Efficiency curves

The entire dataset could be used, for the same reason as the voltage model, despite being affected more from faults.
Across all the experiments the following curves were gathered:

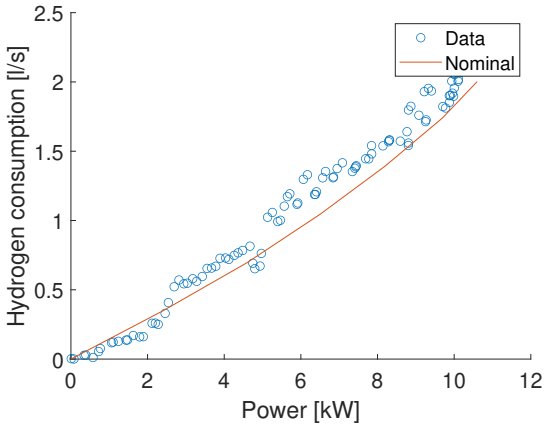

Figure 4.12: Hydrogen consumption compared to the nominal one

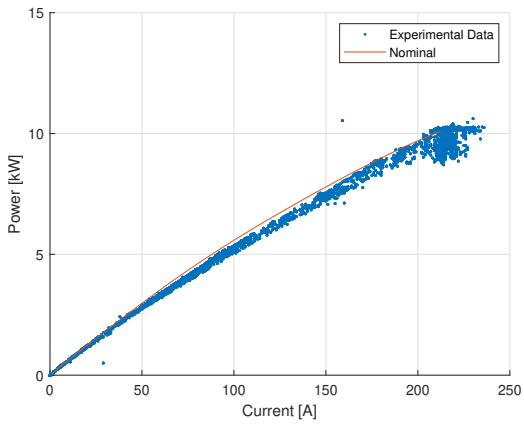

Figure 4.13: Current to power compared to nominal

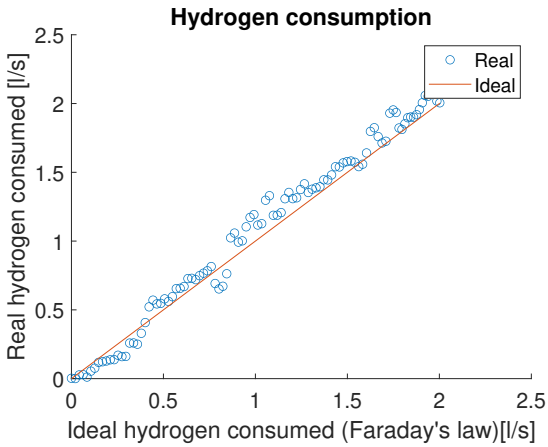

Figure 4.14: Hydrogen consumption compared to the theoretical minimum, computed through Faraday's law. Note that the data at very low powers seems to imply a better than ideal consumption, showing the dataset's limitations

The measured relation from power to hydrogen consumption was approximately linear [1], as confirmed by the linear regression in figure 4.15. The linear model performs worse on

[1] $R^2 = 0.9776$

lower powers, but because of the auxiliaries any power below 3.5 kW is not operationally convenient in most situations, while a power below 2 kW is never favorable, as the consumption is always higher than the production.
Results show that a constant efficiency of 0.48 is a good approximation, about equal to the nominal efficiency at the highest power level. However, it is possible that this is a consequence of the dataset that was available, and further characterization at lower power is required. This might reveal a piecewise approximation is more accurate, as the nominal data seems to show.

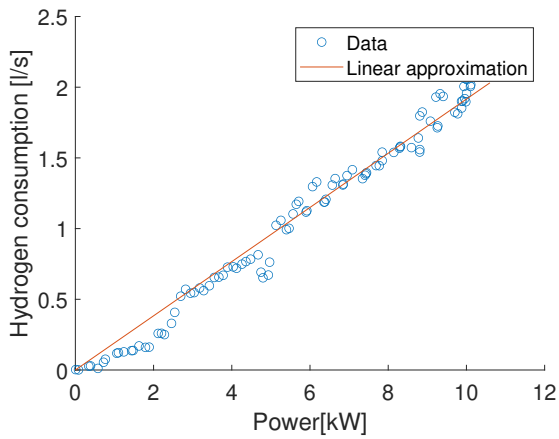

Figure 4.15: Linearized approximation of hydrogen consumption curve

The maximum efficiency, net of auxiliaries, is therefore around 32 %.

4.3. Experimental deployment

With the physical models defined, the first step was an experimental validation of the algorithm, to understand the consequences of the choices taken in modeling the components. The algorithm was deployed on the central Programmable Logic Controller (PLC) of the laboratory, with the objective of studying its behavior in a real-life setting. This PLC is responsible for all the high-level decision-making, data logging and the dispatch of all power setpoints to the individual assets (generators, fuel cell ecc..).

Table 4.6: Experimental Setup Parameters (MG^2LAB)

Component	Parameter	Value
Battery pack	Capacity	20 kWh[a]
	Self-discharge	1 kW
Hydrogen tank	Capacity	30 Nm^3
	Pressure	30 bar
PEM Electrolyzer	Max Power	55 kW
	Auxiliary Power	5 kW (when on)
	Efficiency	0.6 (constant)
PEM Fuel Cell	Max Power	10.6 kW
	Auxiliary Power	3.5 kW(when on)
	Efficiency	0.48 (constant)
B2B Converter	Max Load possible (Simulated)	110 kW
Grid	Connection	Yes

[a] Actual installed capacity is 140 kWh in two 70 kWh battery packs, but only 20 kWh were considered in this experiment

4.3.1. Description of the experiment

Ideally, the full hierarchical MPC architecture described in Chapter 3 would have been deployed to control the experimental microgrid. However, strict integration constraints regarding the laboratory's existing infrastructure and safety protocols precluded the full deployment of the optimization layer in the real-time environment.
Instead, a simplified control logic was implemented. This configuration retained the proposed fuel cell model, but utilized a reduced cost function and relied on the laboratory's legacy models for all peripheral components. These models, especially the batteries, imposed significant limitations, resulting in a power dispatch that reflects the rigid constraints imposed by the existing framework, rather than the economic optimality intended

by the full architecture.
Nevertheless, the experiment serves as a critical validation of the actuation fidelity of the fuel cell, and the strength of the developed model. It should be viewed as a test of the model used in control in a real-world deployment, rather than a demonstration of the higher-level economic optimization capabilities of the architecture.
The experiment consisted in using the B2B converter to simulate the load of a residential profile, shown in figure 4.16. It was sized to be the simulation of a small condominium, so to ensure the fuel cell had to turn on at night. The batteries were sized accordingly to this goal, and reduced to 20 kWh.

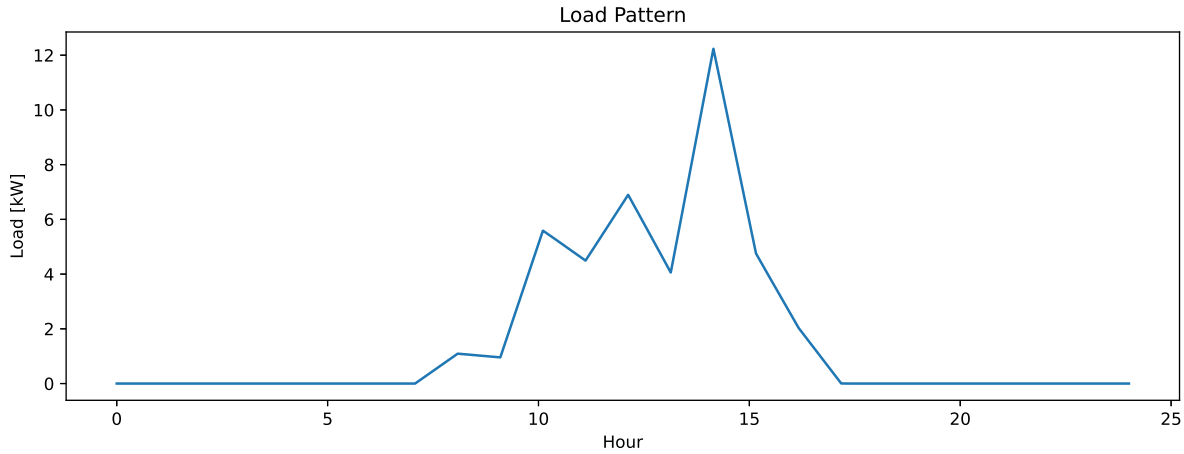

Figure 4.16: Load used for the real-work experiment

4.3.2. Practical implementation Challenges and workarounds

The transition from a purely theoretical simulation to an hardware implementation faced some challenges related to hardware non-idealities, and characteristics of the specific hardware set up. The following sections will detail these challenges, and the solutions implemented to overcome them.

AC/DC inverter nonlinear behavior

The first thing to consider was a problem with power conversion. The relation between the DC power generated by the fuel cell and the AC power read on the AC side had values diverging in a nonlinear fashion. This had to be implemented using a piecewise approximation. This issue will not be considered in the case study, as it is related to this specific set up, and not to the inner functioning of the fuel cell, and this implies an additional drop in overall system efficiency.

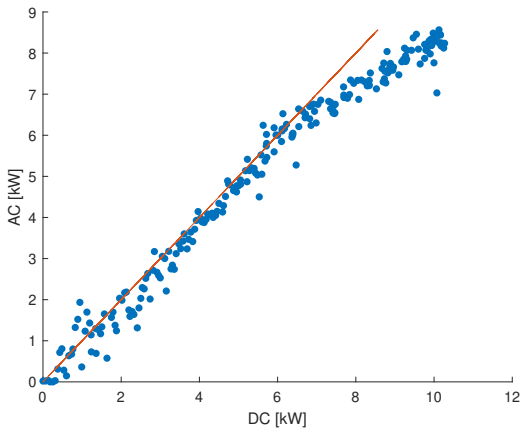

Figure 4.17: Nonlinear diverging curve between DC power and AC power

Setpoint translation

The fuel cell could not reliably receive power setpoints, and it was impossible to directly control the current of the fuel cell. However, it was possible to control the current after the first DC/DC converter onboard of the fuel cell. Empirically, it was found that this aimed at maintaining a voltage of about 44V, therefore the setpoint was found with a simple division of the wanted power over 44.

Temperature transient

This simple control logic did not have any temperature model, for this reason the first power setpoints were limited automatically, so to avoid overloading the fuel cell. This had to be done in real time, and was not considered in the optimization problem.

5 | Algorithm Validation for Seasonal Storage: The Ginostra Microgrid Case

5.1. Introduction

Once the algorithm was validated experimentally, its long-term effectiveness was assessed through a simulated case study. This aspect is particularly important for hydrogen storage, since its primary advantages lie in large-scale, long-term energy storage. Over a shorter period, such as daily operation, its benefits are less apparent. The results will demonstrate that a properly designed control algorithm can substantially enhance microgrid operation through the efficient use of hydrogen.
This will be done over two different scenarios, one where the hydrogen storage system is not scaled properly to guarantee full seasonal storage, and one where it is. Results will show that increasing only hydrogen storage will lead to full self reliance.
Islands provide an excellent example of microgrid applications, as Italy has 19 inhabited islands disconnected from the main grid due to geographical constraints[21]. As a result, each island has its own isolated grid, typically managed by a local utility. At present, they are mostly powered by a small number of diesel generators due to the challenges of renewable energy integration. However, there are plans to expand renewable energy penetration on the islands [37].

5.2. The Ginostra Challenge: Isolation and Seasonality

Figure 5.1: Air view of Stromboli island

Ginostra is a small village belonging to the municipality of Lipari (ME) and located on the island of Stromboli. It has approximately 40 permanent residents, and an annual consumption of 172 MWh[38]. Stromboli itself is isolated from the national grid, but due to its remote position Ginostra is also disconnected from all the other settlements on the island. As a result it has to manage its own independent power system.

The village experiences pronounced seasonality, which results in a highly variable electrical demand. Currently, it is powered by a 100 kW photovoltaic plant coupled with diesel generators.

Fuel is delivered from the port to the generators by helicopter, which is both inefficient and costly, but necessary due to its geographical location. This further justifies the adoption of alternative energy technologies. Thanks to its Mediterranean climate, the site has high solar potential throughout the year, which can be exploited to support energy generation.

In 2018 the EU-funded project REMOTE was launched, in collaboration with Politecnico di Torino. It was aimed at experimenting with cutting-edge technologies and studying the behavior of microgrids across Europe, initially with four demonstration sites:

– Two in Italy (Ginostra, in the south, and Ambornetti in Castelmagno, Piedmont, in the north),

– One in the village of Agkistro, on the border between Greece and Bulgary,

– One in Frøya, located in the Froan archipelago, in the north of Norway.

The project's first phase concluded in 2023, and new demonstration sites are now under development around the globe. Unfortunately after the planning phases the two Italian sites were dismissed, in favor of one on Canary island, Spain [39].
In the case of Ginostra, the design specifications were as reported in [40], and shown in Table 5.1.

Table 5.1: Sizing of the planned project

Component	Value
Thermal power installed [kW]	96[a]
PV power installed [kW]	170[b]
Battery capacity installed [kWh]	600
Hydrogen storage[kg]	100
Electrolyzers installed power [kW]	50
Fuel cells installed power [kW]	50
Renewables penetration [%]	95

[a] Implemented with two 48 kW generators
[b] An upgrade of a 100 kW photovoltaic plant installed by Enel in 2004.

This configuration is sufficient to decrease the use of diesel generators by 95%, according to the project's initial estimates.
Beyond its well-defined component sizing, this network was selected for two main reasons:

- Seasonality: long-term storage is particularly useful when there are significant seasonal variations, either in demand or in generation. [12]. The high touristic load of the island leads to strong fluctuations in energy consumption, with monthly energy consumption nearly triplicating from winter to summer[38], and this forces an oversizing of the energy grid if not properly managed.

- Moderate size: the upfront investment required for hydrogen systems is usually justified only in grids of sufficient scale, as the economic benefits in residential applications are often limited [41], and the cost per kW of small plants is significantly higher [42]. However, very large networks introduce additional complexity, and are not compatible with the experimental data that were gathered. This network is big enough for the investment to be realistic, but still small enough to retain detailed control over individual generators.

5.3. Simulation framework

The case study was implemented in a simulation environment designed to replicate the behavior of the aforementioned microgrid. It was run over a year with an hourly definition, using Pyomo to compute the optimal decisions defined by the optimization problem presented in chapter 3. Given the large temporal scope, only the results from the first layer are shown, as all the dynamics captured by the second layer are largely irrelevant at this scale.

To run the simulation, both load and solar generation profiles were required. Since this data is not publicly available, it was estimated using *HOMER* (Hybrid Optimization of Multiple Energy Resources), a widely used software for modeling and optimizing microgrids and hybrid renewable energy systems. Using this software a seasonally variable profile was generated. This profile was configured to grow from June to August, reaching a peak twice that of April and three times that of January, before decreasing through September and October to represent the seasonality.

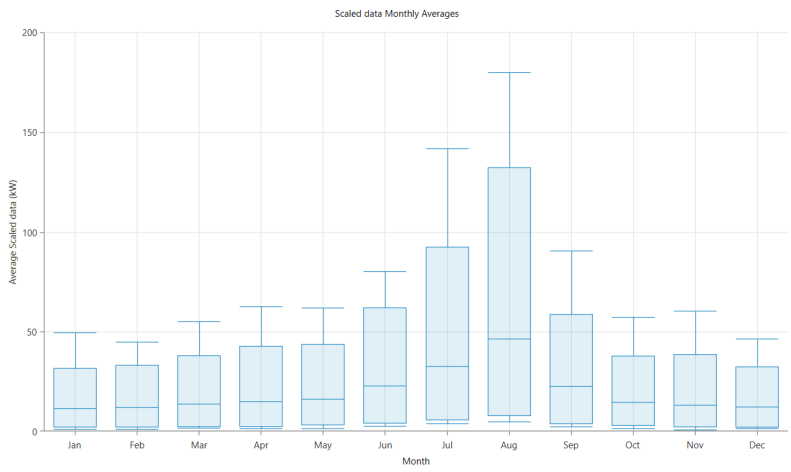

Figure 5.2: Monthly load profile, the horizontal bars represent the monthly median load, the blue rectangles span from the lowest to the highest daily averages, and the vertical bars indicate the instantaneous monthly minimum and maximum values.

The fuel cell and electrolyzer models used in the simulation correspond to the laboratory prototypes, with multiple units connected in parallel when higher power output was required. This is a suboptimal way to implement this system, as efficiency generally improves in machines of higher power.

Then a second scenario is considered, where the efficiency of the fuel cell is taken from

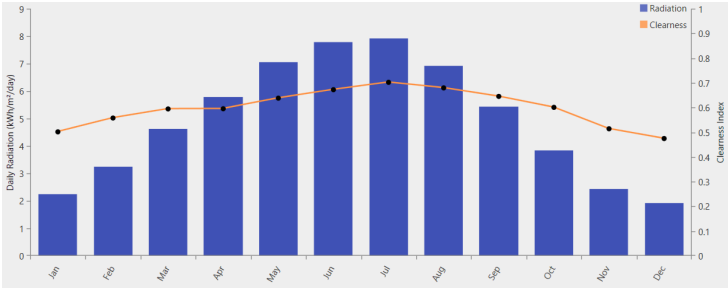

Figure 5.3: Solar profile throughout the year

literature (0.45 net of auxiliaries[43]), and the hydrogen storage was rescaled to achieve fully renewable power generation, resulting in the system defined in table 5.2

Table 5.2: Sizing of the plant

Component	**Value**
Thermal power installed [kW]	96[a]
PV power installed [kW]	170
Battery capacity installed [kWh]	600
Hydrogen storage [kg]	689
Electrolyzers installed power [kW]	50
Fuel cells installed power [kW]	50
Renewables penetration [%]	100

[a] Implemented with two 48 kW generators

Since the original case expected compression at 250 bar this was also implemented in the simulations, modeled as a 7.5% higher consumption for the electrolyzer [44]. However, for this second case the volume is high enough to consider compressing it up to 700 bar. The results indicate the energy to do so is available, without installing new generation. However, this carries many complications, including higher storage costs and additional safety precautions.

Pressure [bar]	**Volume [m^3]**
250	33
450	23.2
700	17.4

Table 5.3: Volume of stored hydrogen at different pressures

5.4. Assumptions

The following assumptions were implemented, to avoid excessive computational complexity:

- Converter losses are neglected
- The diesel generator's efficiency curve has not been estimated from experimental data, but a standard one was used.
- The battery auxiliaries are assumed to be the same as the ones available to the laboratory set in parallel. Each pack is 70 kW, and has 1kW of self consumption when turned on.

6 | Results

6.1. Case study

6.1.1. Planned scenario

This is the previously outlined scenario (the one planned for realization) but using the efficiency curves obtained from the laboratory setup. This corresponds to a hydrogen storage system round trip efficiency of 28% excluding the auxiliaries, generally low compared to values presented in literature[45, 46].

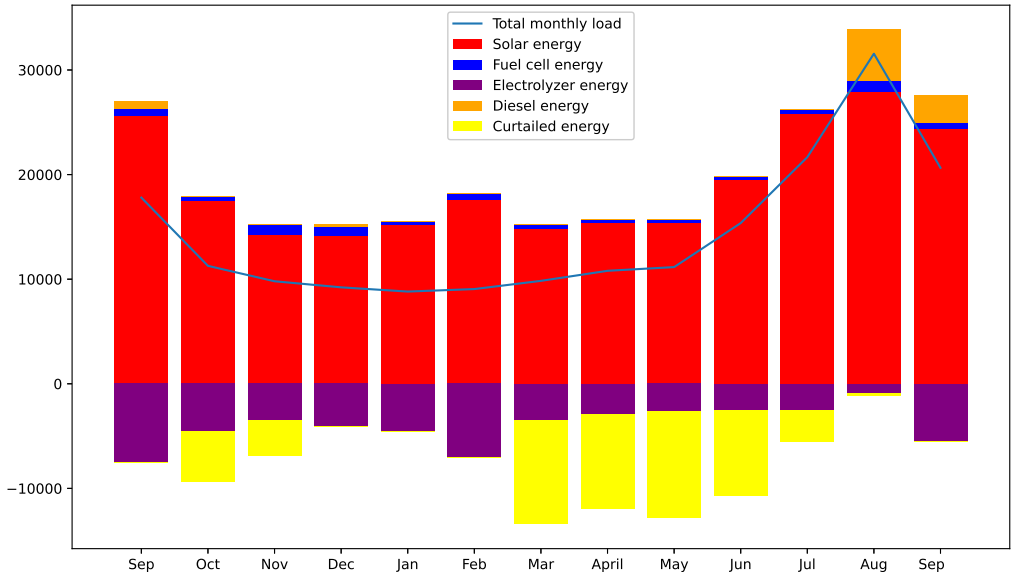

Figure 6.1: Energy balance of the microgrid across the year

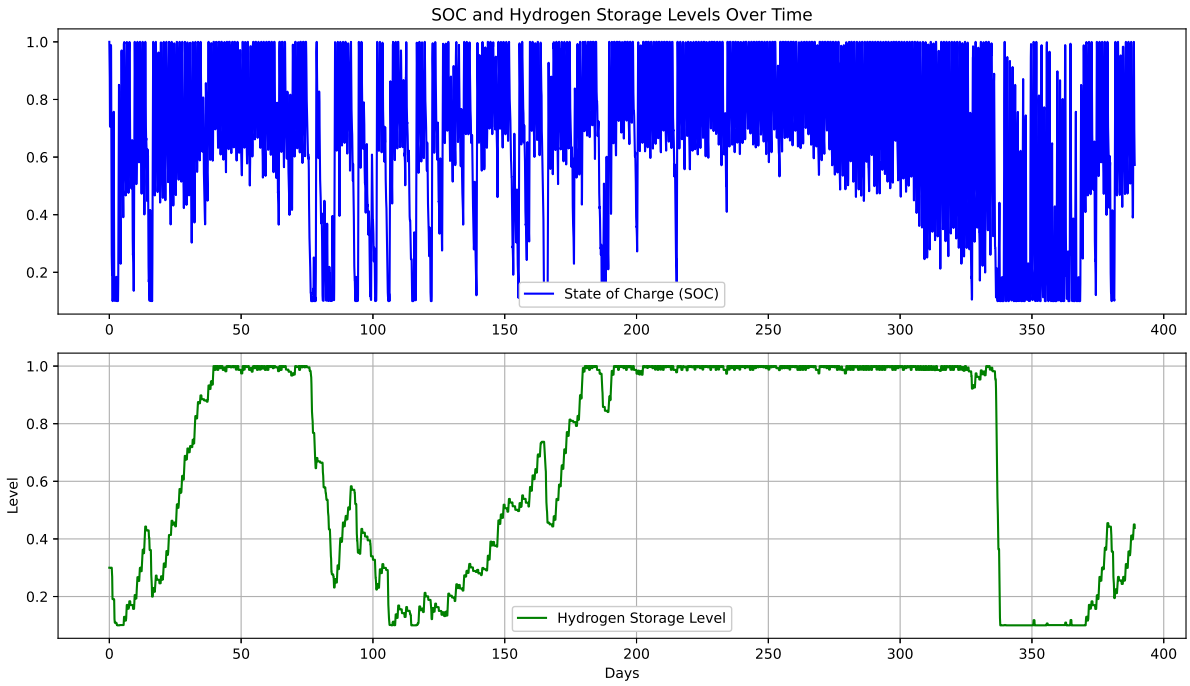

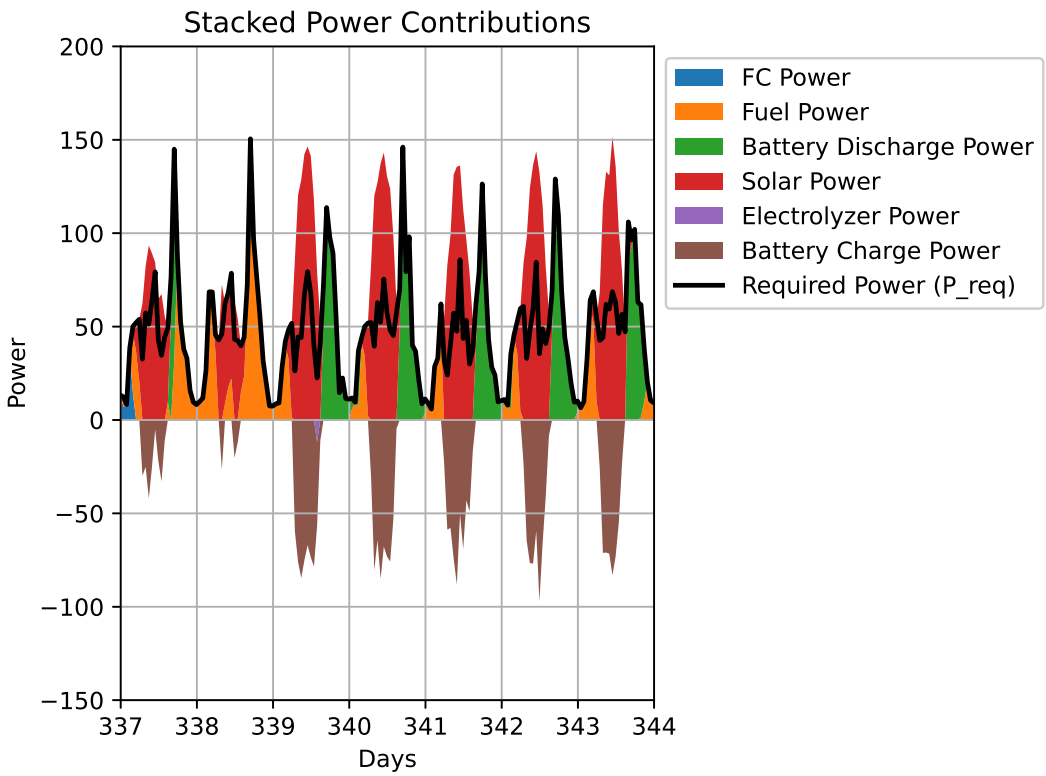

Figure 6.2: Example of a week in August

Hydrogen storage is only ever discharged twice: once in winter when solar power is lowest and once in August, when load is highest. It is big enough to mostly satisfy winter demand (with diesel generation playing a very minor role in December), but it is not sized to cover the peak load demand. It is worth noting that in all months the system generates more power than it is strictly required. This phenomenon is most pronounced in the months when the hydrogen storage tank is empty and energy can be stored for future use. On top of this there are losses caused by the batteries. They are not represented here as they are charged and discharged multiple times a month, accumulating an integral of about zero over 30 days, but have a round trip efficiency of about 85% and auxiliaries of a relevant size are needed for cooling and control.

In this setup, the system still achieves a renewable generation of 95% (or 1568 liters of diesel consumed), but it is also clear that the contribution of hydrogen is very limited, if we were to remove it the renewable generation would only decrease to 93 %(or 2177 liters of diesel consumed) making the inclusion of hydrogen definitely not worth the cost, as diesel generation is still required, and so are the logistics necessary to maintain it.

However, an important assumption was made, multiple small fuel cells in parallel were used to utilize the experimental data. This is not how it is generally done and it is very inefficient. Moreover, the fuel cell that is available in the laboratory has been modeled with an overall low efficiency, and the weight of all the balance of plant (BOP) auxiliaries is exceptionally high. Using the same fuel cell in this scenario does not improve results significantly (1% more renewable penetration) because of the low influence of hydrogen, but it allows for the system to continue working thorough the entire winter without turning on the diesel generators, requiring it only during the summer at peak load.

6.1.2. Ideal fully renewable scenario

In this scenario the fuel cell is modeled using higher efficiency curve (0.45 net of auxiliaries) taken from literature [43], and the hydrogen storage was rescaled to achieve fully renewable power generation, resulting in the system defined in the next pages.

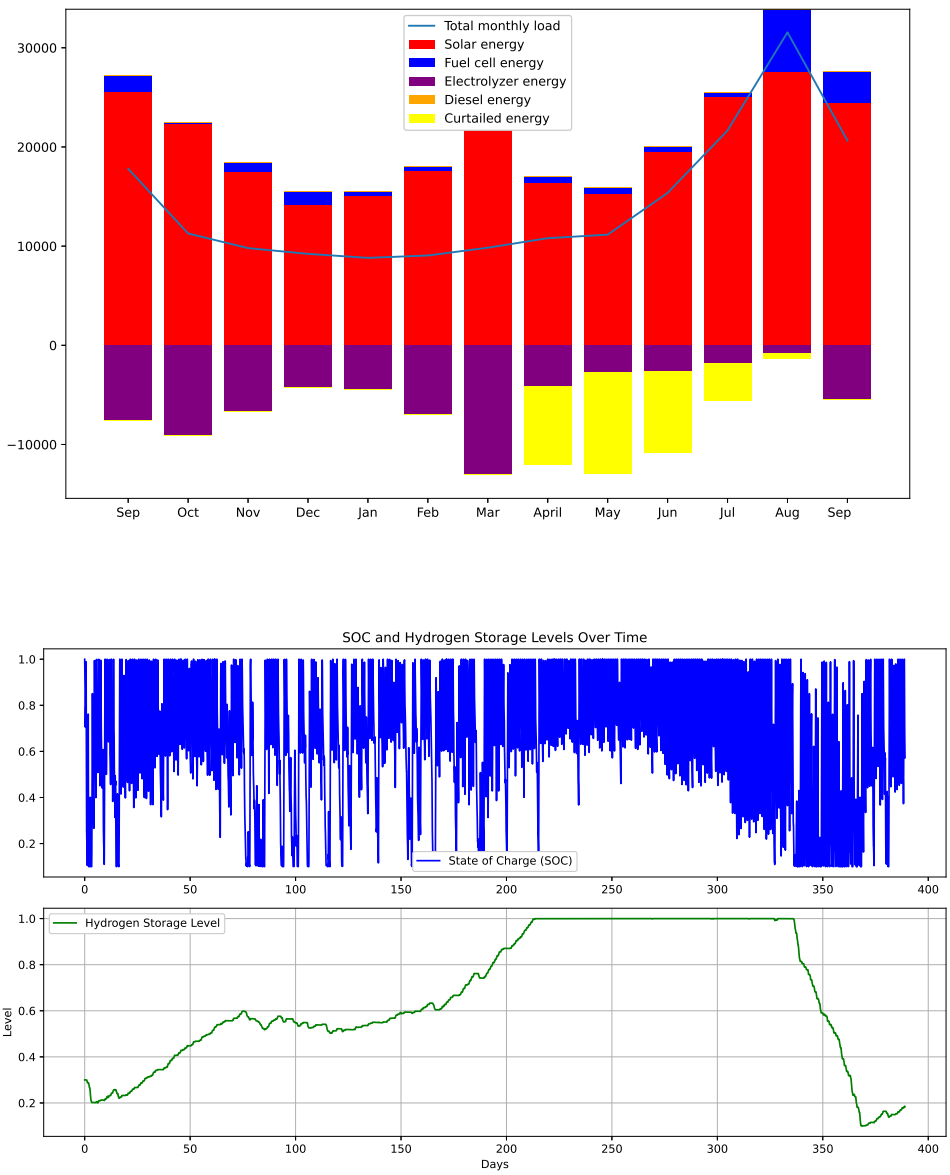

Figure 6.3: Results of the simulation, showing the energy balance(top) and the SOC of batteries and hydrogen tank during the year, in the case of fully renewable generation

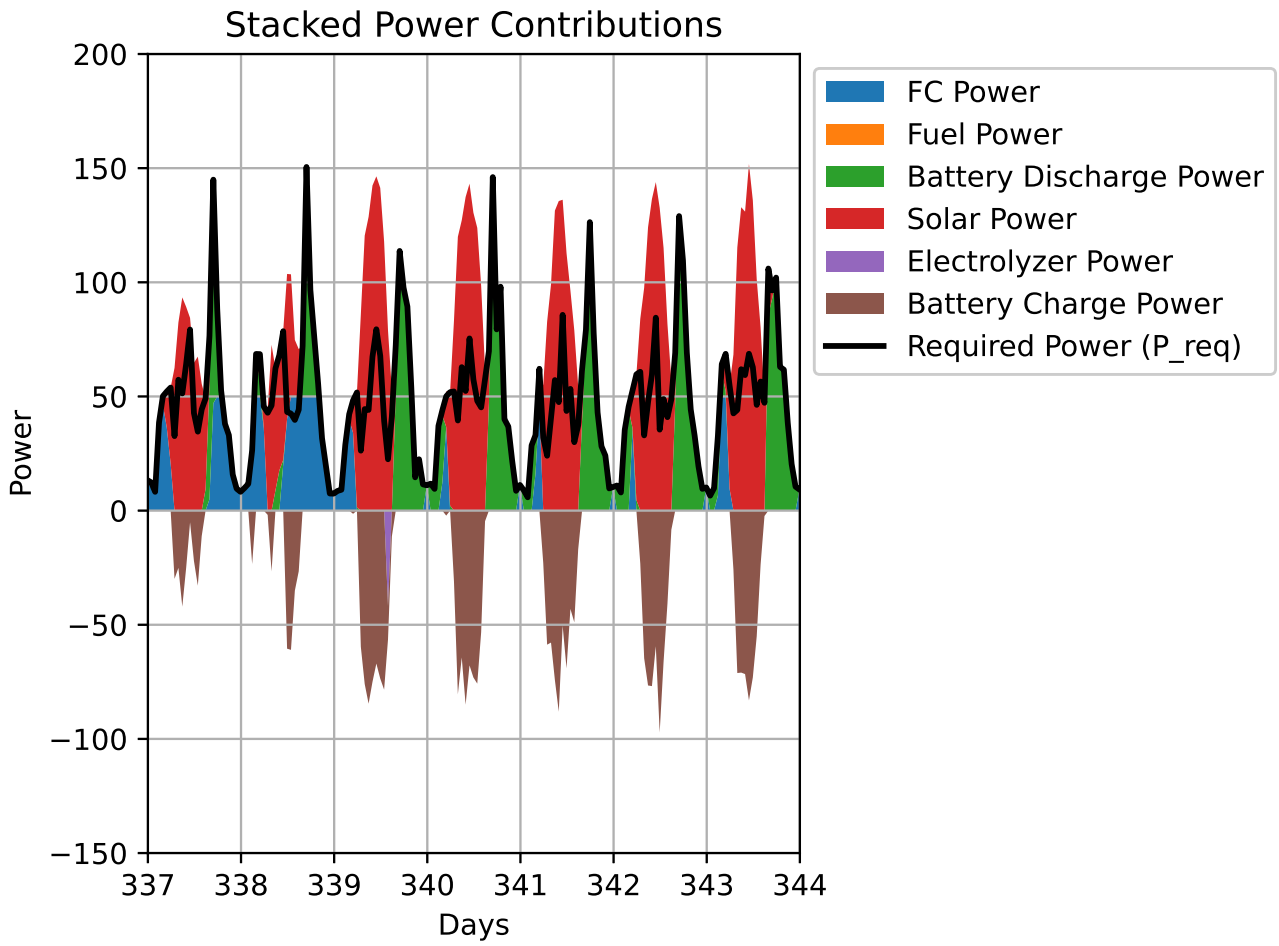

Figure 6.4: Example of a week in August

In this case the storage capacity is sufficient to handle the seasonal variation,eliminating the need for Diesel generators operation.
Notably, curtailment still presents a significant share of total losses. This suggests it would be technically feasible to decrease the installed photovoltaic capacity, while increasing the hydrogen storage further, and achieve the same renewable share. However, determining this optimal balance constitutes an economic sizing problem, which lies beyond the scope of this thesis. This also indicates that further compressing the hydrogen would be feasible, with the only consequence of moving up the moment storage fills up. The complications connected with high storage pressure must be considered.
It is important to note that the the fuel cell power output is not always enough to cover the instantaneous load. Batteries are therefore required to compensate for short-term power gaps as well as for for energy balancing, as illustrated in the example week provided. For this reason, it may be useful to reserve a properly sized battery pack to intervene only

during such conditions, using a modified cost function like the one that was defined for the second layer of the optimization problem.
Overall, this configuration proves that the technology that was planned was already enough to achieve full renewable generation. Hydrogen storage tanks are the only components not to require rare earth elements in the entire systems, and can be added in parallel to the ones that would have been installed. The seasonal storage works as expected, slowly storing energy during the low season, and using it when it is necessary.
However, this comes at the price of oversizing both PV and hydrogen storage capacity. This suggests that this design should not be considered as driven by economic reason, but instead by the goal of autarky and emissions reduction. It highlights the fundamental problem with an energy grid fully powered by intermittent renewable energy sources: ensuring security of supply requires a surplus of capital-intensive assets (generation or storage), most of which run at low utilization during much of the year. This is especially true in places with very high seasonality, such as Ginostra. While load seasonality has always led to oversizing (even with dispatchable generators) generation seasonality is a new paradigm that will have to be faced in the upcoming future, as renewable generation finds its place into the energy mix.

Table 6.1: Comparison of Ginostra Case Study Scenarios

Parameter	**Scenario 1: Planned**	**Scenario 2: Ideal**
FC Efficiency (Net)	32% (from lab data)[a]	45% (from literature)[b]
PV Power	170 kW	170 kW
H2 Storage	100 kg	689 kg
Result: Renewable Share	**95%**	**100%**
Result: Diesel Used	**1200 liters**	**0 liters**

[a] Peak efficiency, net of auxiliaries.

[b] Net constant efficiency, including auxiliaries.

6.2. Results of experimental implementation

The interpretation of these results requires the consideration of three critical system constraints:

- The battery auxiliaries were not implemented as on/off in the legacy model. For this reason the legacy model imposes a continuous discharge of 3.14 kW, regardless of state of charge or connection to the network. This represents a substantial parameter discrepancy, as even when connected to grid the battery's self discharge is 1 kW. The lack of an on/off binary variable meant that a significant amount of the energy has to go into powering the batteries to avoid total discharge, particularly during overnight periods.
- Since the grid is connected to the Bovisa campus, and can output up to 110 kW of power, there are some strict safety requirements. As such, the generation cannot vary too fast, and has to be easily predictable. For this reason there is a tracking penalty: the second layer must track the grid power exchange profile scheduled by the first layer as precisely as possible. This means that, even if the energy the first layer sets aside is not necessary, the second layer will still have to adhere to the planned import/export profile.
- The price of electricity is time varying, which creates opportunity for arbitrage.The absence of a discount factor in the cost function allows for future energy savings, incentivizing multi-day storage strategies.

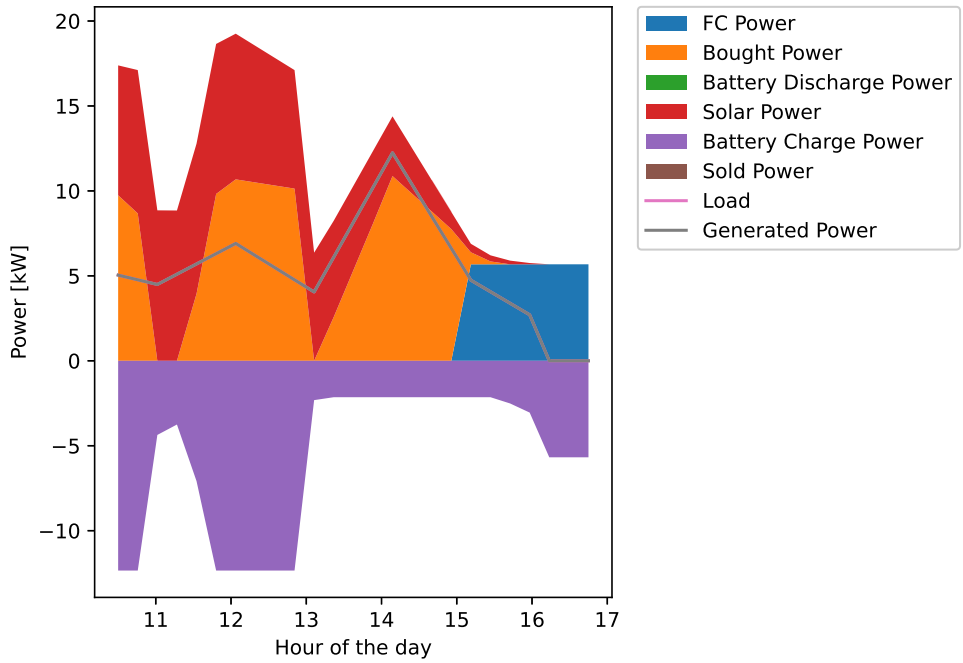

Figure 6.5: Initial first layer output at the start of the experiment

Figure 6.5 shows the importance of correctly setting the value of hydrogen and batteries and the importance of updating the power dispatch profile as the prediction horizon advances. Since hydrogen is not valued, at the end of the day it is used to recharge the battery. Neither the battery state of charge nor the hydrogen state of charge appear in the cost function; thus, storage levels only influence the objective function if they are strictly required within the optimization window. However, as the receding horizon progresses, the controller revises this strategy, in favor of maintaining the hydrogen for later use, as the forecasted load of the following day enters the horizon of the optimization problem. The use of the fuel cell is delayed, likely attributable to the electricity price.
Retrospective analysis using the full algorithm demonstrates the system is capable of fully sustained island-mode operation, eliminating the need for grid reliance entirely

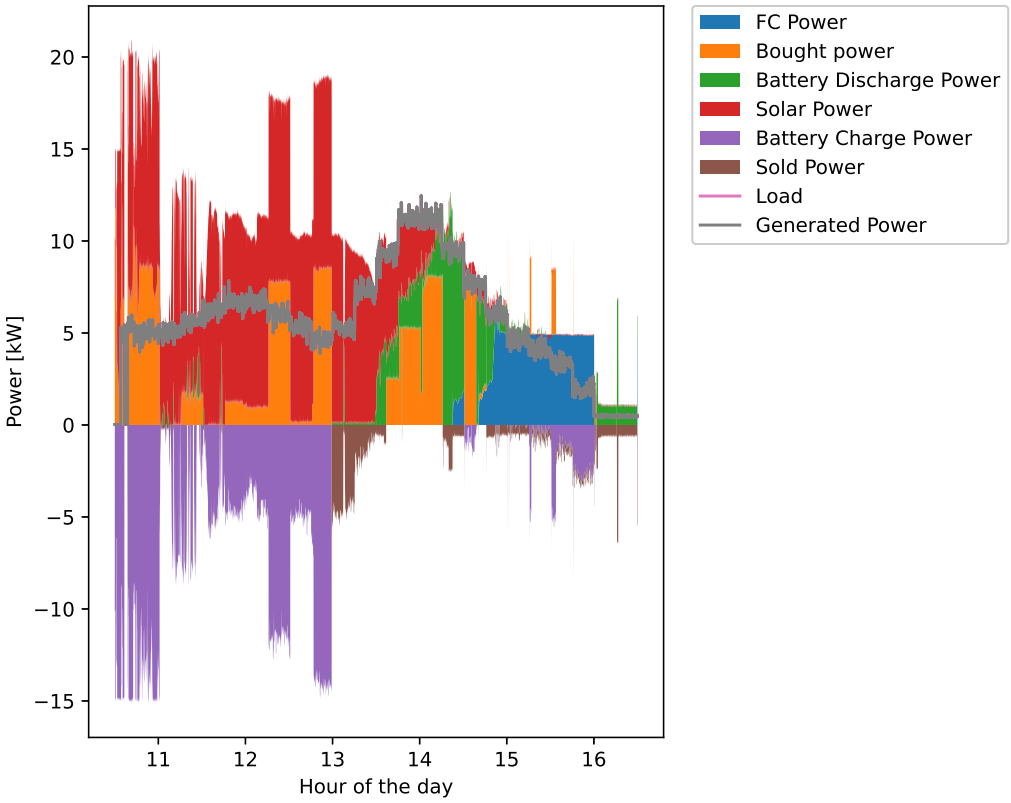

Figure 6.6: Second layer results

The first-layer dispatch schedule exhibited minimal variation via receding horizon updates. The most relevant are the decision to switch off the fuel cell earlier, and a reduction in scheduled grid import as the expected self-discharge did not materialize. The importance of the second layer is underlined here, as it was able to utilize the batteries to compensate for all high frequency fluctuations of solar generation shown in figure 6.7. The fuel cell experienced a fault on startup, but was then manually restored to nominal operation and was able to maintain a full power profile. The data explicitly captures the temperature transient at the start, not considered by the first layer.

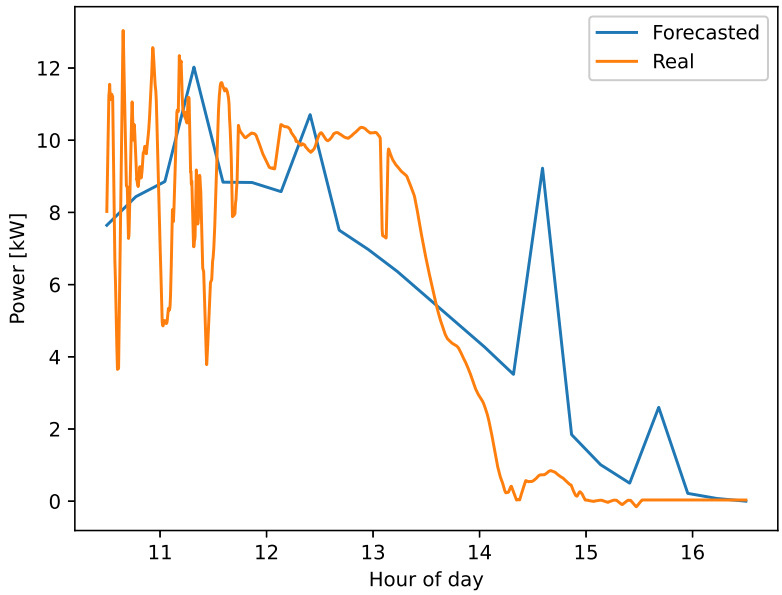

Figure 6.7: Actual solar generation against forecasted solar generation.

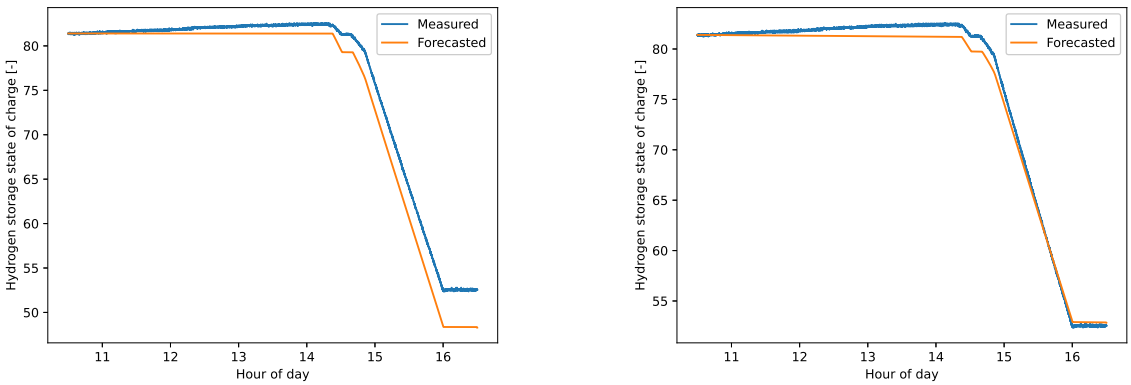

Figure 6.8: Validation of the hydrogen consumption model. Left: Modeled with constant auxiliary consumption. Right: Incorporating the measurement of the real auxiliary consumption

The results in figure 6.8 validate the efficacy of the efficiency model, while highlighting the significant impact that the assumption of constant auxiliary consumption carries. This discrepancy further demonstrates the need for a real-time control layer, capable of compensating the power in real time.

6.3. Second layer results in simulation

6.3.1. Simple battery buffer

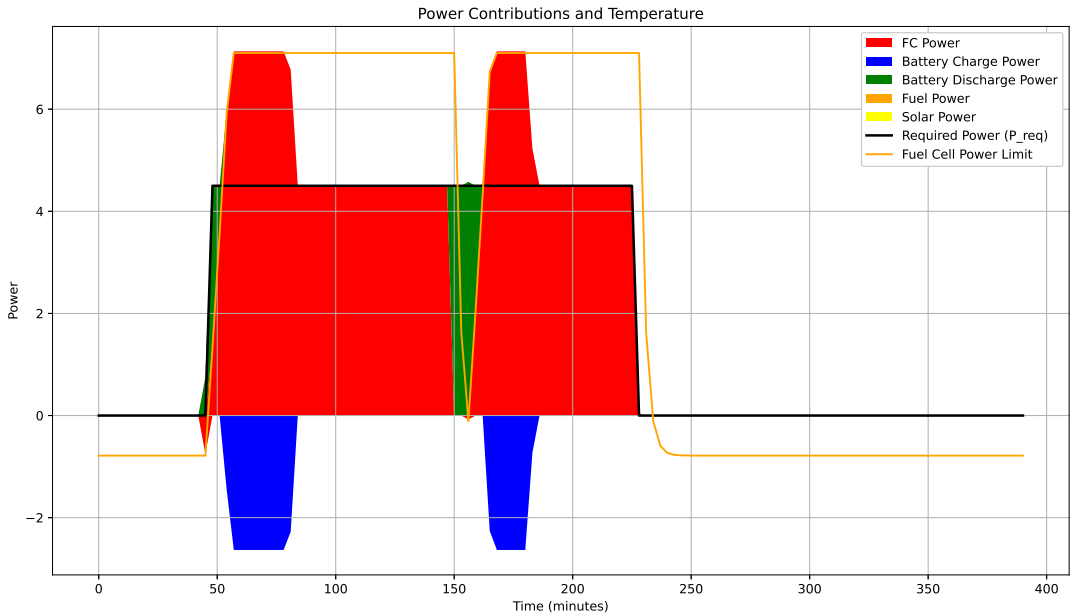

Figure 6.9: Behavior during start up and flooding, power plot

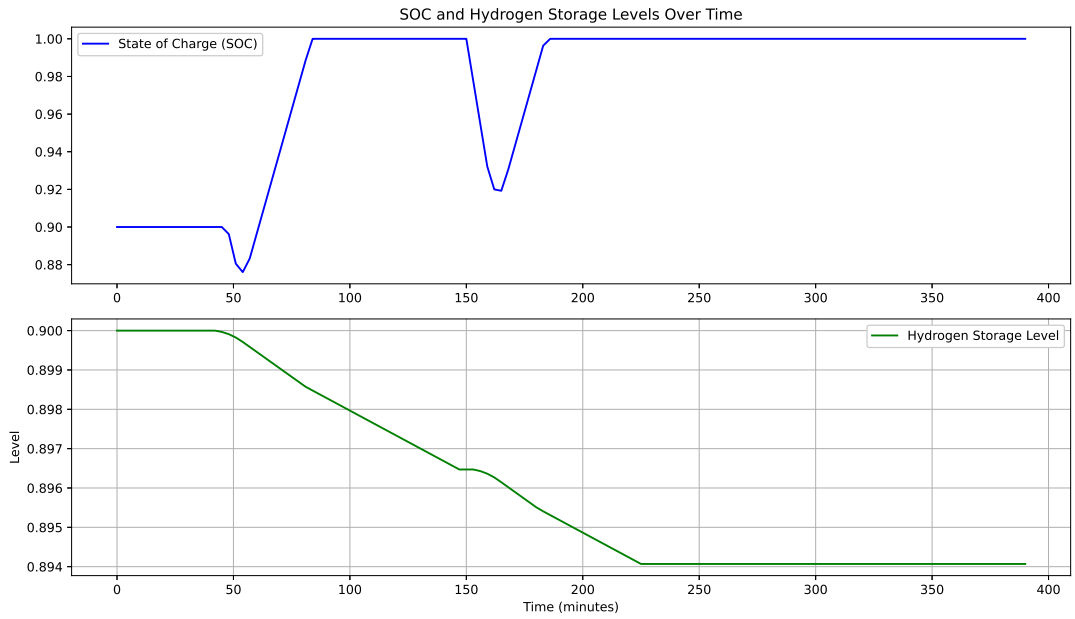

Figure 6.10: Behavior during start up and flooding, states of charge

The battery intervenes whenever necessary, both during the fuel cell warm-up and during simulated flooding. During startup, the available output power is lower than the auxiliary power demand, requiring additional power from the battery. In practice, the cooling is normally off during startup, so such a large imbalance would not occur; however, that cannot be represented in this model.
During flooding, the power demand decreases significantly, as the cell uses most of the downtime to drain. The results show the fuel cell recharging the battery whenever possible, but this may not always be feasible, for instance during continuous operation close to saturation. This issue is addressed in the following section, particularly with regard to warming up.

In this simulation, a 10kWh battery was used to study the response; however similar performances could be achieved with a much smaller battery, as the minimum observed SOC is around 92%. The component size needs to be estimated experimentally, depending on the likelihood of flooding and on the average power demand placed on the fuel cell by the grid. With this strategy the fuel cell effectively behaves as an ideal generator from the grid's perspective. This is especially useful in the context of the two-layer architecture: in MG^2LAB the first layer works on a 15-minute control sampling time, which is too coarse to capture the temperature constraint.

6.3.2. Taking advantage of the two layers

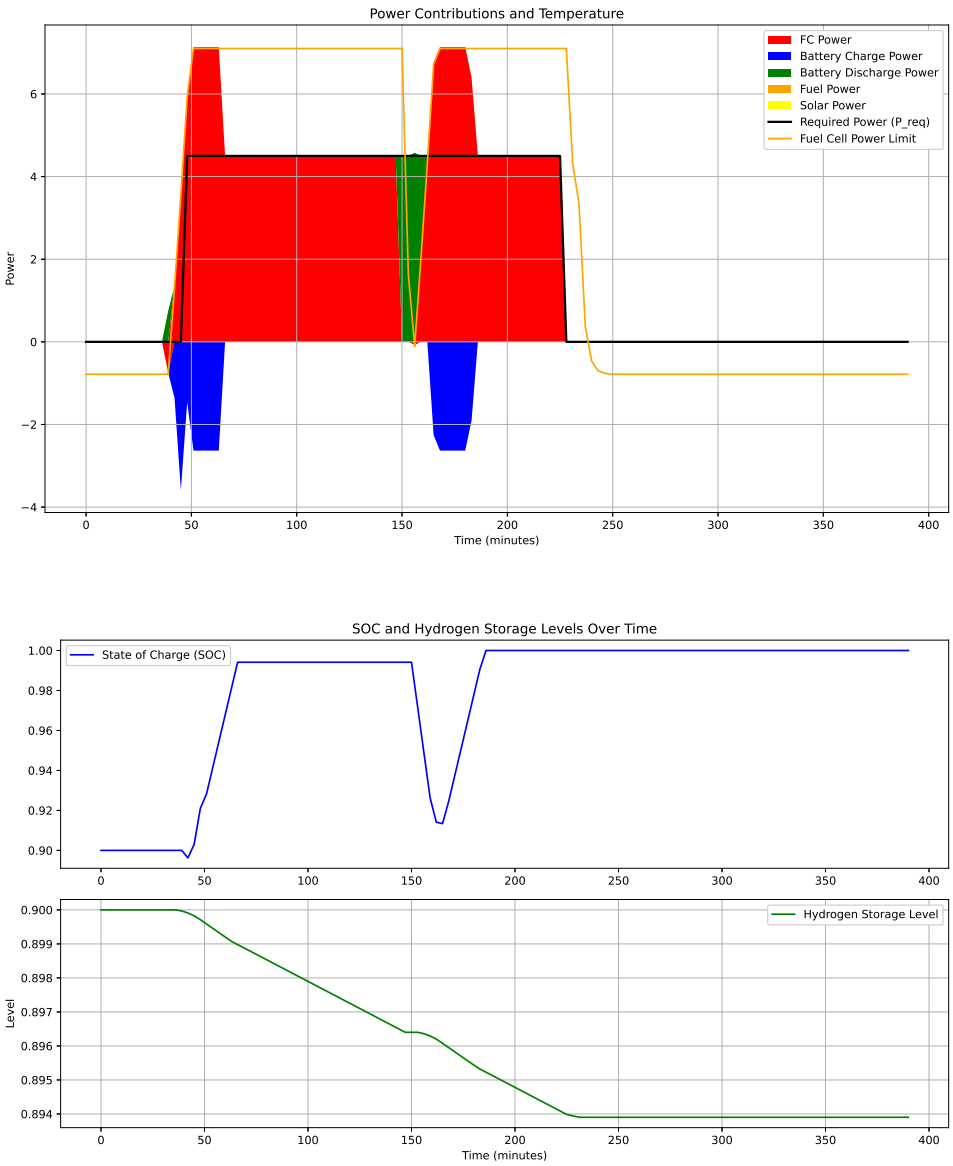

Figure 6.11: Behavior during startup and flooding in simulation: power balance(top) and state of charge of batteries and hydrogen (bottom), with a 10kWh battery, enough for 1 hour of outage

Allowing the fuel cell to be activated 15 minutes before its scheduled operation gives it time to warm up and recharge the batteries if needed, as there is no limit to how much power it can use to increase the SOC in that time frame, other than its own saturation limit.

However, it must, at the very least, to keep its own auxiliaries for 15 minutes, so it will

use more hydrogen.
Due to the specific laboratory setup, this strategy was never tested experimentally, as the flooding issue was temporarily resolved by other means, and introducing such a modification in the second layer was deemed too technically complex for handling only the initial transient.
Overall, the integration of the auxiliary battery and the proposed control logic significantly improves the robustness of the microgrid operation, effectively decoupling the fuel cell's transient behavior from the grid. Although the extended preheating strategy was not experimentally validated, the simulation results demonstrate its potential to mitigate startup instabilities and maintain system balance under uncertainty. These developments strengthen the adaptability of the two-layer control framework, ensuring reliable operation even in the presence of hardware imperfections and unmodeled disturbances. These responses were added to show the response when the fuel cell was the only way to make the problem feasible. However, this is not usually the case, and the second layer can balance the rest of the plant in real time, it usually does not require the battery. The weight given to the battery must be tuned, depending on the reliance the network has on the fuel cell.

7 | Conclusions and future developments

A model for the fuel cell considered was developed, based on the experimental data and compared it against the nominal manufacturer specifications.
The experimental performance was found to be significantly lower than the nominal values, with a gross efficiency, excluding auxiliaries, of 48% and a voltage-current characteristic consistently below the declared curve.
The resulting model was then integrated into an existing hierarchical two-layer Model Predictive Control algorithm. A dedicated cost function was formulated, and the system constraints were reformulated to keep the problem linear.
The algorithm was then implemented experimentally, although with some necessary simplifications due to laboratory constraints.
The same algorithm was subsequently tested over a year-long simulation, using firstly an already sized microgrid, then with a manually sized one to show the potential of reaching fully renewable generation. The control algorithm behaved as expected, storing hydrogen when it was able to and then using it to generate power only when the alternative was diesel.
This thesis demonstrates the potential of hydrogen-based electric storage, which enables long-term energy buffering at scales where it would be impractical and economically unjustified using batteries alone. Imperfections of the hardware are also addressed, with a dedicated control algorithm to handle floodings and black outs. This highlights the importance of a two layer control structure, as generators have complex dynamics that cannot be captured in a linear model, and have to be addressed in real time. This is especially highlighted in situations like this, where the hardware is especially faulty for reasons that could not be addressed directly. As renewable penetration and electrification of all sectors grows, both load and generation will become more complicated to foresee, and will require more complex models. Addressing imperfections like the ones shown here is a starting point to handle these issues.
It's important to note that the algorithm was generalized, despite of these issues, and can

be applied independently of the fuel cell considered.

7.1. Future developments

Improving the data This study had severe hardware limitations, as the results show. The experimental data lacked and was affected by significant hardware faults. Most of the data was gathered at high powers, possibly skewing the results. Future developments should focus on increasing the dataset to improve the models, and ensure their accuracy.

Further validation of the temperature model An experimental campaign needs to be ran, to validate the temperature model, as it could not be done in this context. However, the lack of feedback is a strong limitation, and the development of a model more closely related to a measured one could improve control.

Further yearly studies For what regards a yearly study, the same methodology could be extended to include hydrogen demand, as would occur in the case of an industrial microgrid. Unfortunately the size of such systems was not compatible with the available experimental data, and was therefore outside the scope of this study.

General further validation The most important next step is a more comprehensive experimental validation of this algorithm, since only its most fundamental components could be implemented in the real system.

Bibliography

[1] World Meteorological Organization. Wmo global annual to decadal climate update (2025–2029). Report, World Meteorological Organization (WMO), May 26 2025. URL https://wmo.int/files/wmo-global-annual-decadal-climate-update-2025-2029.

[2] Hannah Ritchie, Max Röser, and Pablo Rosado. Greenhouse gas emissions by sector. *Our World in Data*, 2024.

[3] International Energy Agency. Electricity 2025. Technical report, International Energy Agency, 2025. URL https://iea.blob.core.windows.net/assets/0f028d5f-26b1-47ca-ad2a-5ca3103d070a/Electricity2025.pdf.

[4] Herib Blanco and André Faaij. A review at the role of storage in energy systems with a focus on power to gas and long-term storage. *Renewable and Sustainable Energy Reviews*, 81:1049–1086, 2018. doi: 10.1016/j.rser.2017.07.062. URL https://www.sciencedirect.com/science/article/pii/S1364032117311310.

[5] Alexander N.C. Edington. The role of long duration energy storage in decarbonizing power systems. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA, 2019. URL https://dspace.mit.edu/handle/1721.1/122095.

[6] Bethany A. Frew, Sarah Becker, Michael J. Dvorak, Gorm B. Andresen, and Mark Z. Jacobson. Flexibility mechanisms and pathways to a highly renewable us electricity future. *Energy*, 101:65–78, 2016. ISSN 0360-5442. doi: https://doi.org/10.1016/j.energy.2016.01.079. URL https://www.sciencedirect.com/science/article/pii/S0360544216300032.

[7] Stefan Weitemeyer, David Kleinhans, Thomas Vogt, and Carsten Agert. Integration of renewable energy sources in future power systems: The role of storage. *Renewable Energy*, 75(C):14–20, None 2015. doi: 10.1016/j.renene.2014.09.028. URL https://ideas.repec.org/a/eee/renene/v75y2015icp14-20.html.

[8] Terna. Sezione "accumulii" — generazione, dati di accumulo energetico. https://dati.terna.it/generazione#accumuli, 2025.

[9] A. G. Olabi, Joseph Onwubiko, Nikhil Singh, T. Al-Ansari, Abdelkareem M.A., Hegazy Rezk, and Hussein M. Maghrabie. Energy storage and its role in renewable energy systems. *Energy*, 220:119693, 2021. ISSN 0360-5442. doi: 10.1016/j.energy.2020.119693.

[10] Tom Headley and Marwan Khzouz. Chapter 11: Hydrogen energy storage. In *Energy Storage Handbook*. Sandia National Laboratories, 2021.

[11] Tay Son Le, Tuan Ngoc Nguyen, Dac-Khuong Bui, and Tuan Duc Ngo. Optimal sizing of renewable energy storage: A techno-economic analysis of hydrogen, battery and hybrid systems considering degradation and seasonal storage. *Applied Energy*, 336:120817, April 2023. ISSN 0306-2619. doi: 10.1016/j.apenergy.2023.120817. URL `http://dx.doi.org/10.1016/j.apenergy.2023.120817`.

[12] Caterina Cara, Paolo Marocco, Riccardo Novo, Matti Koivisto, Massimo Santarelli, and Giuliana Mattiazzo. Modeling the long-term evolution of the italian power sector: The role of renewable resources and energy storage facilities. *International Journal of Hydrogen Energy*, 59:1183–1195, 2024. doi: 10.1016/j.ijhydene.2024.01.358. URL `https://doi.org/10.1016/j.ijhydene.2024.01.358`.

[13] Financial Times. Negative european energy prices hit record level, 2024. URL `https://www.ft.com/content/1f94d0b4-c839-40a2-9c8d-782c00384154`.

[14] Mohit Chhabra. Feasting on the glut: Putting renewable curtailment in context, December 2024. URL `https://www.nrdc.org/bio/mohit-chhabra/feasting-glut-putting-renewable-curtailment-context`.

[15] Andreas Horsch and Benjamin Aust. Negative market prices on power exchanges: Evidence and policy implications from germany. *The Electricity Journal*, 33, 01 2020. doi: 10.1016/j.tej.2020.106716.

[16] François Benhmad and Jacques Percebois. Photovoltaic and wind power feed-in impact on electricity prices: The case of germany. *Energy Policy*, 119:317–326, 2018. ISSN 0301-4215. doi: https://doi.org/10.1016/j.enpol.2018.04.042. URL `https://www.sciencedirect.com/science/article/pii/S0301421518302581`.

[17] Marek Pavlík, František Kurimský, and Kamil Ševc. Renewable energy and price stability: An analysis of volatility and market shifts in the european electricity sector (2015–2025). *Applied Sciences*, 15(12), 2025. ISSN 2076-3417. doi: 10.3390/app15126397. URL `https://www.mdpi.com/2076-3417/15/12/6397`.

[18] U.S. Department of Energy. Doe microgrid workshop report. Technical report, U.S. Department of Energy, 8 2012.

[19] Andrey Bernstein, David Fobes, Vaibhav Donde, Matthew Reno, and Line Roald. Dynagrid: Dynamic microgrids for large-scale der integration and electrification. National Renewable Energy Laboratory (NREL), Golden, CO (United States), 09 2022. URL `https://www.osti.gov/biblio/1891199`.

[20] Arman Amini Badr, Amin Safari, and Sajad Najafi Ravadanegh. Power system islanding considering effect of microgrid by milp formulation. *Electric Power Systems Research*, 216:109093, 2023. ISSN 0378-7796. doi: https://doi.org/10.1016/j.epsr.2022.109093. URL `https://www.sciencedirect.com/science/article/pii/S0378779622011427`.

[21] Smart Island. Levanzo – replicabilità delle isole egadi, 2025. URL `http://www.smartisland.eu/replicabilita/egadi/levanzo.html`.

[22] Raniyah Wazirali, Elnaz Yaghoubi, Mohammed Shadi S. Abujazar, Rami Ahmad, and Amir Hossein Vakili. State-of-the-art review on energy and load forecasting in microgrids using artificial neural networks, machine learning, and deep learning techniques. *Electric Power Systems Research*, 225:109792, 2023. ISSN 0378-7796. doi: https://doi.org/10.1016/j.epsr.2023.109792. URL `https://www.sciencedirect.com/science/article/pii/S0378779623006818`.

[23] A. V. Waghmare, V. P. Singh, T. Varshney, and P. Sanjeevikumar. A systematic review of reinforcement learning-based control for microgrids: trends, challenges, and emerging algorithms. *Discover Applied Sciences*, 7(939), 2025. doi: 10.1007/s42452-025-07529-6. URL `https://doi.org/10.1007/s42452-025-07529-6`.

[24] Sijia Li, Frede Blaabjerg, and Amjad Anvari-Moghaddam. A transferable drl-based intelligent secondary frequency control for islanded microgrids. *Electronics*, 14(14), 2025. ISSN 2079-9292. doi: 10.3390/electronics14142826. URL `https://www.mdpi.com/2079-9292/14/14/2826`.

[25] Stefano Leonori, Enrico De Santis, Antonello Rizzi, and F.M. Frattale Mascioli. Optimization of a microgrid energy management system based on a fuzzy logic controller. In *IECON 2016 - 42nd Annual Conference of the IEEE Industrial Electronics Society*, pages 6615–6620, 2016. doi: 10.1109/IECON.2016.7793965.

[26] Hang Shuai, Xiaomeng Ai, Jiakun Fang, Wei Yao, and Jinyu Wen. Double deep q-learning based real-time optimization strategy for microgrids, 2021. URL `https://arxiv.org/abs/2107.12545`.

[27] Hao Zhou and Melike Erol-Kantarci. Decentralized microgrid energy management: A multi-agent correlated q-learning approach. In *2020 IEEE International Conference on Communications, Control, and Computing Technologies for Smart Grids (SmartGridComm)*, page 1–7. IEEE, November 2020. doi: 10.1109/smartgridcomm47815.2020.9302992. URL `http://dx.doi.org/10.1109/SmartGridComm47815.2020.9302992`.

[28] Xiaoyu Ge and Javad Khazaei. Physics-informed convolutional neural network for microgrid economic dispatch. *Sustainable Energy, Grids and Networks*, 40:101525, 2024. doi: 10.1016/j.segan.2024.101525. URL `https://doi.org/10.1016/j.segan.2024.101525`.

[29] Armin Nurkanović, Amer Mešanović, Andrea Zanelli, Gianluca Frison, Jonathan Frey, Sebastian Albrecht, and Moritz Diehl. Real-time nonlinear model predictive control for microgrid operation. In *2020 American Control Conference (ACC)*, pages 4989–4995, 2020. doi: 10.23919/ACC45564.2020.9147816.

[30] Simone Polimeni, Lorenzo Meraldi, Luca Moretti, S. Leva, and Giampaolo Manzolini. Development and experimental validation of hierarchical energy management system based on stochastic model predictive control for off-grid microgrids. *Advances in Applied Energy*, 2:100028, 05 2021. doi: 10.1016/j.adapen.2021.100028.

[31] Simone Polimeni, Luca Moretti, Giampaolo Manzolini, S. Leva, Lorenzo Meraldi, and Pietro Raboni. Numerical and experimental testing of predictive ems algorithms for pv-bess residential microgrid. pages 1–6, 06 2019. doi: 10.1109/PTC.2019.8810548.

[32] Simone Polimeni, Luca Moretti, Emanuele Martelli, Sonia Leva, and Giampaolo Manzolini. A novel stochastic model for flexible unit commitment of off-grid microgrids. *Applied Energy*, 331:120228, 2023. ISSN 0306-2619. doi: https://doi.org/10.1016/j.apenergy.2022.120228. URL `https://www.sciencedirect.com/science/article/pii/S0306261922014854`.

[33] Adam Z. Weber and John Newman. Coupled thermal and water management in polymer electrolyte fuel cells. *Journal of The Electrochemical Society*, 153(12):A2205, oct 2006. doi: 10.1149/1.2352039. URL `https://doi.org/10.1149/1.2352039`.

[34] Manish Khandelwal, sungho Lee, and M.M. Mench. One-dimensional thermal model of cold-start in a polymer electrolyte fuel cell stack. *Journal of Power Sources*, 172: 816–830, 10 2007. doi: 10.1016/j.jpowsour.2007.05.028.

[35] Tao-Feng Cao, Hong Lin, Li Chen, Ya-Ling He, and Wen-Quan Tao. Numerical

investigation of the coupled water and thermal management in pem fuel cell. *Applied Energy*, 112:1115–1125, 2013. doi: 10.1016/j.apenergy.2013.02.031.

[36] Purushottam Kumar and Sanjib Ganguly. Digital twin approach for pem fuel cell degradation prognostics using equivalent electrical circuit model and polarization curves. 2025.

[37] Ministry of the Environment and Energy Security (Italy). National integrated plan for energy and climate (pniec) – draft update 2021-2030. `https://commission.europa.eu/system/files/2023-07/ITALY%20-%20DRAFT%20UPDATED%20NECP%202021%202030%20%281%29.pdf`, June 2023.

[38] Paolo Marocco, Massimo Santarelli, and Andrea Lanzini. An milp approach for the optimal design of a stand-alone hybrid pv-battery-hydrogen system: A case study of ginostra (italy). *Applied Energy*, 295:117070, 2021. doi: 10.1016/j.apenergy.2021.117070. URL `https://doi.org/10.1016/j.apenergy.2021.117070`.

[39] Politecnico di Torino. A new remote demo in the canary islands, 5 2021. URL `https://archivio-poliflash.polito.it/en/research_innovation/a_new_remote_demo_in_the_canary_islands`.

[40] M. Gandiglio, P. Marocco, I. Bianco, D. Lovera, G.A. Blengini, and M. Santarelli. Life cycle assessment of a renewable energy system with hydrogen-battery storage for a remote off-grid community. *International Journal of Hydrogen Energy*, 47(77):32822–32834, 2022. ISSN 0360-3199. doi: https://doi.org/10.1016/j.ijhydene.2022.07.199. URL `https://www.sciencedirect.com/science/article/pii/S0360319922032505`.

[41] Pietro Lubello, Mattia Pasqui, Alessandro Mati, and Carlo Carcasci. Assessment of hydrogen-based long term electrical energy storage in residential energy systems. *Smart Energy*, 8:100088, 2022. ISSN 2666-9552. doi: https://doi.org/10.1016/j.segy.2022.100088. URL `https://www.sciencedirect.com/science/article/pii/S2666955222000260`.

[42] Anita H. Reksten, Magnus S. Thomassen, Steffen Møller-Holst, and Kyrre Sundseth. Projecting the future cost of pem and alkaline water electrolysers; a capex model including electrolyser plant size and technology development. *International Journal of Hydrogen Energy*, 47(90):38106–38113, 2022. ISSN 0360-3199. doi: https://doi.org/10.1016/j.ijhydene.2022.08.306. URL `https://www.sciencedirect.com/science/article/pii/S0360319922040253`.

[43] Elena Crespi, Giulio Guandalini, Sönke Gößling, and Stefano Campanari. Modelling

and optimization of a flexible hydrogen-fueled pressurized pemfc power plant for grid balancing purposes. *International Journal of Hydrogen Energy*, 46(24):13190–13205, 2021. ISSN 0360-3199. doi: 10.1016/j.ijhydene.2021.01.085.

[44] European Commission. Demonstration of innovative solutions for high-capacity, reliable, flexible, and sustainable hydrogen compression technologies in commercial applications. EU Funding & Tenders Portal, Jan 2024. URL `https://ec.europa.eu/info/funding-tenders/opportunities/portal/screen/opportunities/topic-details/horizon-jti-cleanh2-2024-02-04`.

[45] Naef A. A. Qasem and Gubran A. Q. Abdulrahman. A recent comprehensive review of fuel cells: History, types, and applications. *International Journal of Energy Research*, 2024. doi: 10.1155/2024/7271748. URL `https://doi.org/10.1155/2024/7271748`.

[46] Xiaohua Wang, Andrew G. Star, and Rajesh K. Ahluwalia. Performance of polymer electrolyte membrane water electrolysis systems: Configuration, stack materials, turndown and efficiency. *Energies*, 16(13), 2023. ISSN 1996-1073. doi: 10.3390/en16134964. URL `https://www.mdpi.com/1996-1073/16/13/4964`.

List of Figures

1.1 Share of global greenhouse emissions by industry 1
1.2 Comparison of different technologies at different storage size 2
1.3 Hydrogen consumption by sector from the European hydrogen observatory 4
1.4 Italian load as forecasted by Terna and measured across a week in October. 5

3.1 Power consumed by auxiliaries over time, in a test ran at constant power. It is possible to note the switch in temperature control strategy. Every drop in power is fuel cell downtime 10
3.2 Linear approximation of the current/power relationship 12
3.3 Comparison between the first order transfer function constraint and the real one, both when the power applied is the highest allowed by the constraint and when the fuel cell is turned off 13
3.4 Equivalent circuit of the Fuel cell stack 14
3.5 Hierarchical structure of the EMS and information flow. 16
3.6 Conceptual schematic of the second-layer 'battery buffer' strategy. The Energy Management System (EMS) issues a power reference $P_{fc}^{exp,t}$ to the fuel cell system. This demand is met by the combined action of the actual fuel cell (P_{fc}^{t}) and the dedicated auxiliary battery ($P_{batt}^{em,t}$), which actively compensates for internal fluctuations, startup delays, and flooding-induced power drops. 23

4.1 The PET-10 fuel cell energy system 25
4.2 Nominal Power-Voltage-Current of the fuel cell stack 25
4.3 Nominal hydrogen consumption 26
4.4 Nominal efficiency 26
4.5 Power and temperature curve of the test used to characterize the fuel cell . 27
4.6 Comparison between real and simulated water temperature, with zoom on the initial transient (B) 29
4.7 Comparison between real and simulated cabinet temperature 30

4.8 Comparison between real and simulated water temperature, with zoom on the initial transient (B) on a different test 31
4.9 Comparison between real and simulated cabinet temperature on a different test 31
4.10 Results over one test 32
4.11 Comparison with experimental data and nominal V-I curve at 62 °C 32
4.12 Hydrogen consumption compared to the nominal one 33
4.13 Current to power compared to nominal 33
4.14 Hydrogen consumption compared to the theoretical minimum, computed through Faraday's law. Note that the data at very low powers seems to imply a better than ideal consumption, showing the dataset's limitations . 33
4.15 Linearized approximation of hydrogen consumption curve 34
4.16 Load used for the real-work experiment 36
4.17 Nonlinear diverging curve between DC power and AC power 37

5.1 Air view of Stromboli island 40
5.2 Monthly load profile, the horizontal bars represent the monthly median load, the blue rectangles span from the lowest to the highest daily averages, and the vertical bars indicate the instantaneous monthly minimum and maximum values. 42
5.3 Solar profile throughout the year 43

6.1 Energy balance of the microgrid across the year 45
6.2 Example of a week in August 46
6.3 Results of the simulation, showing the energy balance(top) and the SOC of batteries and hydrogen tank during the year, in the case of fully renewable generation 48
6.4 Example of a week in August 49
6.5 Initial first layer output at the start of the experiment 52
6.6 Second layer results 53
6.7 Actual solar generation against forecasted solar generation. 54
6.8 Validation of the hydrogen consumption model. Left: Modeled with constant auxiliary consumption. Right: Incorporating the measurement of the real auxiliary consumption 54
6.9 Behavior during start up and flooding, power plot 55
6.10 Behavior during start up and flooding, states of charge 55

6.11 Behavior during startup and flooding in simulation: power balance(top) and state of charge of batteries and hydrogen (bottom), with a 10kWh battery, enough for 1 hour of outage 57

7.1 Example of a fault of the DC/DC converter 77
7.2 Example of a typical test, affected by flooding 78

List of Tables

3.1 Voltage Component Descriptions 14
3.2 Weighting factors used in the cost function. 19
3.3 Value of storing 10 kWh 20
3.4 Cost of generating 10 kWh. 20

4.1 Setpoint of the test considered 28
4.2 Parameter values for L_1 through L_7. 29
4.3 Accuracy of results on the test they were fit on 30
4.4 Accuracy of results on a different test 31
4.5 Parameter results 32
4.6 Experimental Setup Parameters (MG^2LAB) 35

5.1 Sizing of the planned project 41
5.2 Sizing of the plant 43
5.3 Volume of stored hydrogen at different pressures 43

6.1 Comparison of Ginostra Case Study Scenarios 50

List of Symbols

Variable	Description	SI unit
$\boldsymbol{A'}$	System coefficient of fictitious temperature T'	1/s
$\boldsymbol{a_0, a_1}$	Voltage model parameters (activation loss)	-
$\boldsymbol{aux_{on}}$	Boolean parameter (1 if cooling auxiliaries are on)	-
$\boldsymbol{B'}$	Input coefficient of fictitious temperature T'	kW/°C/s
$\boldsymbol{b_0, b_1}$	Voltage model parameters (activation loss)	-
$\boldsymbol{c_i}$	Linear coefficient of discharge of component i	-
$\boldsymbol{C}$	Total cost function	€
$\boldsymbol{C_N}$	Base optimization cost	€
$\boldsymbol{C_t^{ud}}$	Operational cost at time t	€
$\boldsymbol{E_{batt}^{em,t}}$	Energy stored in the emergency battery at time t	kWh
$\boldsymbol{E_{batt}^{max}}$	Maximum energy capacity of the emergency battery	kWh
$\boldsymbol{F}$	Faraday constant	C/mol
$\boldsymbol{f}$	Linear coefficient from current to power	kW/A
$\boldsymbol{H_{flow,el}^t}$	Hydrogen flow rate of the electrolyzer at time t	l/s
$\boldsymbol{H_{flow,fc}^t}$	Hydrogen flow rate of the fuel cell at time t	l/s
$\boldsymbol{H_{storage}^{max}}$	Maximum hydrogen storage capacity	l
$\boldsymbol{I}$	Current	A
$\boldsymbol{I_{fc}}$	Fuel cell current	A
$\boldsymbol{I_{max}(T)}$	Maximum allowable current as a function of temperature	A
$\boldsymbol{L_1 - L_7}$	Coefficients of thermal model	1/s
$\boldsymbol{LHV_f}$	Lower heating value of fossil fuel	kWh/l
$\boldsymbol{LHV_H}$	Lower heating value of hydrogen	kWh/l
$\boldsymbol{M_H}$	Total stored volume of hydrogen	l
$\boldsymbol{\dot{m}_{H_2}}$	Molar flow rate of H2	mol/s
$\boldsymbol{m_0, m_1}$	Voltage model parameters (concentration loss)	-

Variable	Description	SI unit
m_f	Fuel intake	l
N_s	Number of cells in the stack	-
n	Voltage model parameter (concentration loss)	-
P	Price of fuel	€/l
P_i^{min}	Minimum power output for component i	kW
P_i^t	Power generated or requested by instrument i at time t	kW
P_{batt}^{em}	Power of the emergency battery	kW
$P_{batt}^{em,max}$	Maximum power of the emergency battery	kW
P_{fc}^{exp}	Expected power output from fuel cell (grid perspective)	kW
P_{fc}^{max}	Maximum fuel cell power output	kW
P_{H_2}	Partial pressure of hydrogen	atm
P_{H_2O}	Partial pressure of water	atm
P_{O_2}	Partial pressure of oxygen	atm
q	Linear coefficient relating fuel cell power to hydrogen flow	l/s/kW
R_{ohm}	Ohmic resistance of the stack	Ω
SOC_i^t	State of Charge of component i at time t	-
$SOC_{i,min}$	Minimum State of Charge limit	-
$SOC_{i,max}$	Maximum State of Charge limit	-
$SOC_{batt}^{em,t}$	State of Charge of the emergency battery at time t	-
$S\dot{O}C_i^t$	Rate of change of State of Charge at time t	1/s
$S\dot{O}C_{i,min}$	Minimum rate of change of State of Charge	1/s
$t^{horizon}$	Optimization time horizon	s
T_k'	Fictitious temperature introduced to model thermal constraint	°C
$\dot{T}_k'$	Time derivative of the fictitious temperature	°C/s
T_s	Sampling time	s
T_{amb}	Ambient temperature	°C
T_{aux}	Temperature of the auxiliary cooling unit	°C
T_{cab}	Temperature of the cabinet the fuel cell is in	°C
T_{FC}	Fuel cell temperature	°C
$T_{fc,K}$	Fuel cell temperature in Kelvin	K
U	Vector of optimization variables	-
V_{act}	Activation voltage loss	V

Variable	Description	SI unit
V_{conc}	Concentration voltage loss	V
V_{ohm}	Ohmic voltage loss	V
V_{rev}	Reversible (Nernst) voltage	V
V_{stack}	Total voltage output of the fuel cell stack	V
W_0	Weighting coefficient for fossil fuel cost	-
W_1	Weight given to hydrogen energy in cost function	-
W_2	Weight given to battery energy in cost function	-
W_3	Weighting value for emergency battery cost	-
x_i^t	Normalized power output of component i at time t	-
x_{batt}^{em}	Normalized power of the emergency battery	-
x_{fc}^{exp}	Normalized expected power from fuel cell	-
z_i^t	Binary commitment variable for component i at time t	-
$\eta_{[]}$	Efficiency of component []	-
η_{batt}^{disch}	Efficiency of battery discharge	-
η_{batt}^{char}	Efficiency of battery charge	-

Appendix

7.2. Fault characterization

7.2.1. DC/DC faults

These are characterized by strong current fluctuations in the initial transient, and are cause by a mismatch between the power requested by the inverter and the maximum power that the fuel cell can provide. These have been present in almost all of the experiments. For the experiment a fix was implemented, using additional communication directly with the inverter.

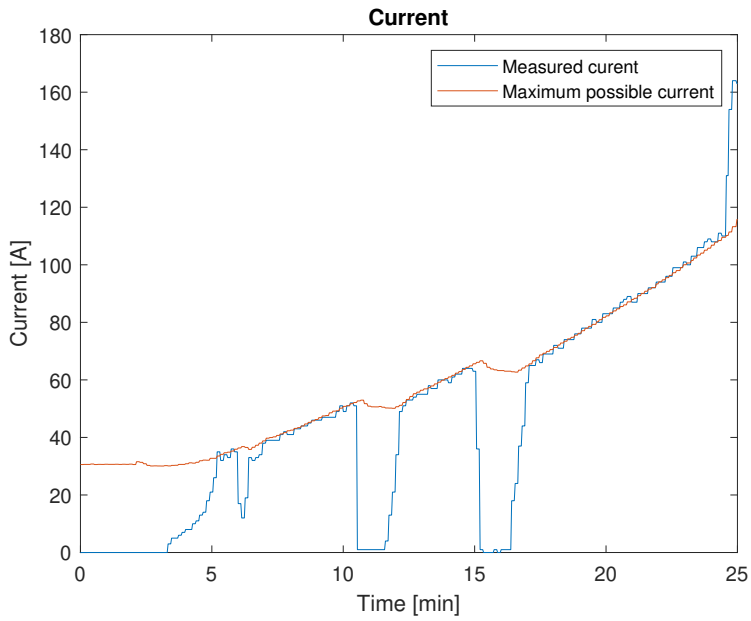

Figure 7.1: Example of a fault of the DC/DC converter

7.2.2. Flooding

These are characterized by a sudden drop in power, that could last several minutes. They were caused by a malfunctioning valve, that was fixed right before the experiment.

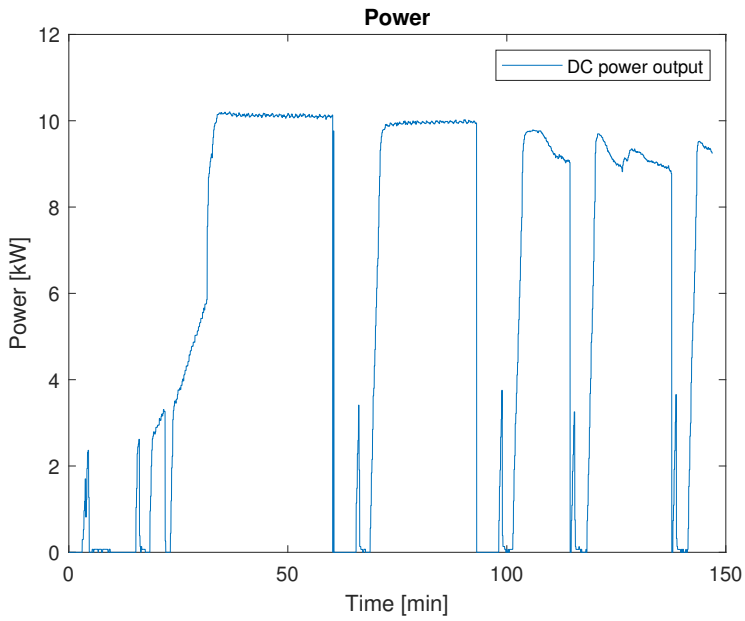

Figure 7.2: Example of a typical test, affected by flooding