



POLITECNICO
MILANO 1863

SCUOLA DI INGEGNERIA INDUSTRIALE
E DELL'INFORMAZIONE

EXECUTIVE SUMMARY OF THE THESIS

Fuel Cell Modeling for Microgrid Control Applications

LAUREA MAGISTRALE IN AUTOMATION AND CONTROL ENGINEERING - INGEGNERIA
DELL'AUTOMAZIONE

Author: FRANCESCO BERRUTI

Advisor: PROF. GIAMPAOLO MANZOLINI

Co-advisor: RICCARDO SIMONETTI

Academic year: 2025-2026

1. Introduction

The shift of the power grid towards intermittent renewable power generation is the main driver behind research in the control of electrical grids. In this context storage systems are critical, as they are the main tool to mitigate the challenges of non dispatchable generators.

Currently, Li-ion batteries are the dominant technology for short-term storage, but long-term storage is less commercially mature. Hydrogen is one of the most interesting tools considered for this purpose, thanks to its high energy density and versatility, allowing the coupling of different sectors. This thesis develops and integrates a linearized model of a hydrogen storage system, with a particular focus on the fuel cell, into a hierarchical Model Predictive Control (MPC) algorithm.

2. Fuel Cell models

The fuel cell was characterized using three models:

- **Thermal model:** A thermal model is essential to characterize a fuel cell, as temperature strongly influences the maximum speed of the reaction, and therefore maxi-

mum power output. Generating high power at low temperatures highly reduces the fuel cell's longevity. This model accounts for the nonlinear behavior arising from the cooling control rules implemented by the producer, but a linear approximation was used to improve control, albeit at the cost of losing the benefits of feedback.

- **Voltage model:** Although not necessary for control, it was important to characterize the fuel cell further, and understand its operating points.
- **Efficiency curve:** An efficiency model was essential to characterize consumption beyond the nominal consumption, to accurately predict fuel consumption.

3. Experimental campaign

An experimental campaign was conducted to calibrate all of the models. This was complicated by many operational faults, as the fuel cell had multiple setup issues. Flooding was frequent and startup was often disrupted by intermittent power generation due to an issue with the DC/DC converter. However, these issues provided an opportunity to tackle more advanced control strategies in the real-time layer,

isolating these issues from the main grid. The fuel cell was repaired right before the experimental implementation of the algorithm, and for this reason this advanced fault-handling logic was not tested experimentally.

Thermal model A thermal model was crucial, as the producer enforced a rule where the maximum current could not be more than twice the temperature in Celsius. This limit was applied until the fuel cell reached 55 °C, after which the limit was removed. Beyond the nonlinearity of the constraint itself, a temperature model was necessary to understand the limitations imposed by the manufacturer’s rule.

Two temperatures are measured in the considered fuel cell: the temperature of the cooling water (estimated to be about the same as the fuel cell temperature) and the temperature of the fuel cell enclosure. To improve model estimation one more state was introduced into the system, which is the temperature of the cooling auxiliary. This state, while not relevant for control, was important for the parameter estimation. Over a timescale of hours, it would warm up, making cooling less effective and changing the behavior of the water temperature.

An iterative prediction error minimization procedure was used to estimate all model parameters.

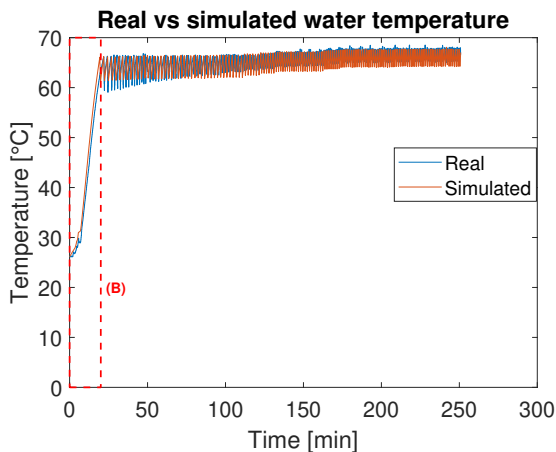


Figure 1: Comparison between simulated and real fuel cell temperature behavior, in the nonlinear model

However, the fuel cell’s internal controller implemented a nonlinear on/off rule, which renders this model incompatible for direct implementa-

tion in the linear optimization problem the controller solves in real time. For this reason, a linear model was developed, based on the nonlinear one, with the sole purpose of enforcing the constraint. The connection with the physical system was therefore lost, which resulted in the loss of a very important feedback, although with no risk of divergence. The underlying concept involves simulating a first order system, whose behavior closely approximated the constraint itself, when paired with a second constraint always bounding the power to be less than the maximum one. This rendered the set of constraints completely linear, at the cost of decoupling the constraint from the real measured variable, and introducing a less accurate approximation.




Figure 2: Comparison between linear and nonlinear constraint

Voltage model A voltage model was estimated through parameter identification, starting from the equivalent circuit model in [1]. This highly nonlinear model is not a dynamical one, meaning the estimation was more straightforward, with no need for iterative techniques. The results show a behavior significantly worse than the nominal voltage curve, with a voltage always below that indicated by the manufacturer. The real voltage however converged toward the nominal voltage at higher power. The efficiency curve matches these results, showing an overall poor performance compared to the expected one. This might be related to the many faults that characterized the entire experimental campaign, or simply to the low quality of the overall system, beyond the stack itself. This finding underscores the need to gather experimental data,

rather than the ideal specifications of the manufacturer.




Figure 3: Modeled voltage-current characteristics compared to the nominal one at 62 °C

Efficiency model The efficiency model was straightforward and based on a linear approximation, which yielded an R^2 of 0.97 over the entire dataset.




Figure 4: Linear approximation of consumption, compared with experimental data

4. Additions to the optimization problem

For details regarding the entire architecture, refer to [3].

First layer The cost function was modified to give an economic value to hydrogen, which is, in reality, a cost-free resource, as it is generated only by renewable sources.

The optimal behavior would be to set a priority

order: solar energy is used whenever available. If solar generation is not sufficient, then batteries intervene and the fuel cell is used as a last resort. Vice versa, when there is extra energy, it should go to the batteries first, and any energy in excess of that should go to the electrolyzer. This is accomplished by giving an economic value to the energy stored in each of the resources. Ignoring additional scaling factors that appear in the cost function, if we call W_0 the weight assigned to energy coming from fuel, W_1 the value assigned to hydrogen storage, W_2 the weight assigned to the batteries, then it is only necessary that:

$$W_1 \cdot \frac{\eta_{el}}{\eta_{batt}^{char}} < W_2 \leq W_1 \cdot \frac{\eta_{batt}^{disch}}{\eta_{fc}} \ll W_0 \quad (1)$$

where $\eta_{[]}$ is the efficiency of component [], with a distinction between charging and discharging efficiency for the battery. This ensures that the value of the energy stored in hydrogen is higher than the value of energy stored in batteries. However, the lower bound prevents the controller from discharging the batteries (or not charging them) to increase hydrogen storage. Then all the essential components and commitment variables were added:

$$\begin{aligned} x_{fc}^t &\in [0, 1] \\ z_{fc}^t &\in \{0, 1\} \\ z_{fc}^t &\geq x_{fc}^t \\ x_{fc}^t &\geq P_{fc}^{min} z_{fc}^t \\ P_{fc}^t &= x_{fc}^t P_{fc}^{max} \end{aligned} \quad (2)$$

$$H_{storage}^{t+1} = H_{storage}^t + \frac{H_{flow,el} - H_{flow,fc}}{H_{storage}^{max}}$$

$$H_{storage}^t \in [H_{storage}^{min}, H_{storage}^{max}]$$

$$\forall t \in [0, t^{horizon}]$$

Second layer The main structure of the second layer remained similar to that of the first layer, but a further control logic was implemented, to handle the very frequent faults and the thermal constraint.

The power actually coming from the fuel cell was decoupled from the power expected from the grid, using a battery as a buffer. This buffer is not visible by the first layer, and is only accessible by the fuel cell. Its energy has a high value in the cost function, ensuring that the fuel cell

is used to recharge it.

This already handles faults, as long as the fuel cell has the sufficient capacity to recharge the battery. However, this is not guaranteed, and to handle the thermal constraint the energy gap to be refilled is not negligible. For this reason it is prudent to automatically turn on the fuel cell a time step (15 minutes) earlier, so to ensure the fuel cell is warm when the grid requires it to be on, and the power required from the battery for startup is minimal (during warm up auxiliaries might consume more than the fuel cell can provide). This might however increase hydrogen consumption: the fuel cell must stay on for longer, but, depending on the frequency of the faults, it might discharge and recharge the battery less, increasing overall efficiency.

This was implemented in the optimization problem by adding the following variables and constraints:

$$\begin{aligned} P_{fc}^{exp,t} &= x_{fc}^{exp,t} P_{fc}^{max} \\ P_{fc}^t - P_{batt}^{em,t} &= P_{fc}^{exp,t} \\ \forall t \in [0, t^{horizon}] \end{aligned} \quad (3)$$

The expected fuel cell power is the one that appears in the grid balancing equation, so that from the grid's perspective the fuel cell is an ideal generator.

5. Experimental implementation

Once the models were calibrated, it was necessary to test the algorithm in a real world scenario. Ideally, the algorithm would have been directly tested on the experimental grid. Due to laboratory constraints, this was not possible, and an heavily simplified control logic was used. This used a simpler cost function (simple minimization of electricity bought during the day) and only implemented the developed fuel cell model. Other hardware dynamics connected to the laboratory's specific setup had to be considered, and influenced the optimization results. This does not serve as a demonstration of the economic optimization capabilities of the algorithm, rather as a proof of concept of the actuation fidelity of the fuel cell and the capability of the architecture to balance the grid in a real-world scenario. Additionally, it serves as additional model validation.

A standard residential load of a small condo-

minium was used, temporally compressed to align it with the operational hours of the laboratory.




Figure 5: Load profile used for the experiment

It was also necessary to address an issue strictly related to the laboratory setup, a divergence between the AC and DC side. This was modeled with a piecewise linear approximation.




Figure 6: Divergence between AC and DC power output

Notably, this reduced the maximum output of the fuel cell (net of auxiliaries) further, down to little above 5 kW. The battery size was limited to 20 kWh, to ensure fuel cell utilization.

6. Case study

The Ginostra microgrid was chosen for the long-term simulation. This is an isolated village on the island of Stromboli, which manages its own microgrid due to significant geographical constraints. It has about 40 year-round residents, but its population more than doubles in summer, causing strong seasonal trends. This is also reflected in energy consumption, which grows from 10 MWh in January to 30 MWh in August [2]. This very strong seasonality further complicates the use of non dispatchable generation, and

presents a significant challenge that requires significant storage, making it an ideal, manageably-sized case study. Relevantly, a hydrogen storage system was already sized and planned for construction, which was later canceled in favor of other options.

7. Results

7.1. Case study

Scenario planned for execution The sizing for the planned system was used, substituting the fuel cell with multiple units of the prototype present in the laboratory.

This is not the most efficient implementation, as efficiency drops with lower power systems, and in the second scenario a different setup is considered.




Figure 7: Energy balance month by month in the scenario planned for execution




Figure 8: Average state of charge of battery packs and hydrogen storage over time in the scenario planned for execution

As expected, hydrogen is used in the winter months to decrease diesel consumption to almost zero, but the size of the storage tank is not sufficient to power the grid throughout August.

Ideal renewable scenario Here, a more efficient fuel cell model from the literature is considered, and the storage tank is resized to allow for fully renewable generation.




Figure 9: Energy balance month by month in the fully renewable scenario

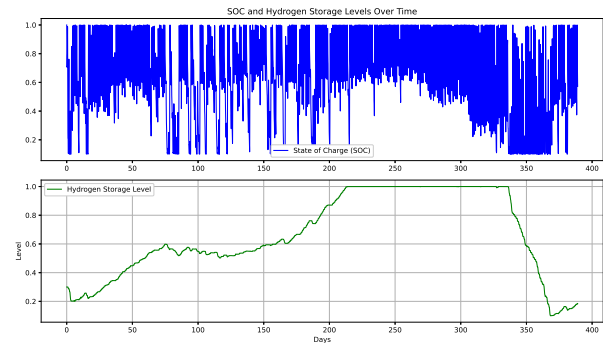


Figure 10: Average state of charge of battery packs and hydrogen storage over time in the fully renewable scenario

Now diesel generation is fully replaced both during winter and summer, managing to take advantage of the power that was curtailed in early summer.

7.2. Experimental results




Figure 11: Second layer results

Beyond the suboptimal behavior of the simpler control logic, the capability of the second layer to dynamically balance the grid despite solar generation unpredictable behavior is clear. The fuel cell experienced a startup fault, but after manual intervention it was able to turn back on and function at full power. It is possible to see the temperature transient, and how the batteries are able to balance the power imbalance.




Figure 12: Hydrogen consumption computed a posteriori using the efficiency model, compared to the real one. Using the assumption of constant auxiliary consumption (left), using the measured consumption (right)

7.3. Simulated second layer

As the second layer was not implemented in the real system, it was evaluated through simulation to understand the behavior when it is necessary for the fuel cell to follow a specific setpoint. The battery was set to be only 90% full at the start of the simulation, in order to show the two different behaviors regarding interaction with batteries.

Without warm up First, the results of the modified optimization problem were simulated.

The battery is used on startup to fill the gap created by the thermal constraint. Later, when the fuel cell is warm it uses excess power to recharge the battery. This is only possible because the fuel cell is not requested to run at full power. This operating point, however, is sub-optimal, as the constant auxiliary load means that the fuel cell operates most efficiently when at full power, and the optimizer tries to avoid it if possible.

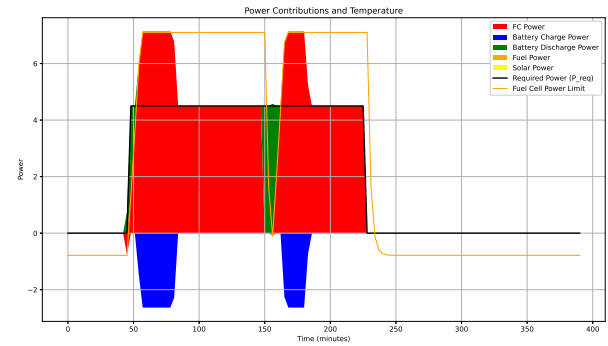


Figure 13: Step response of the system, and response to a flooding, without warm up

Including the warm up logic Adding the warm up timestep decreases the use of the battery, and gives the fuel cell a chance to recharge the battery while it warms up.

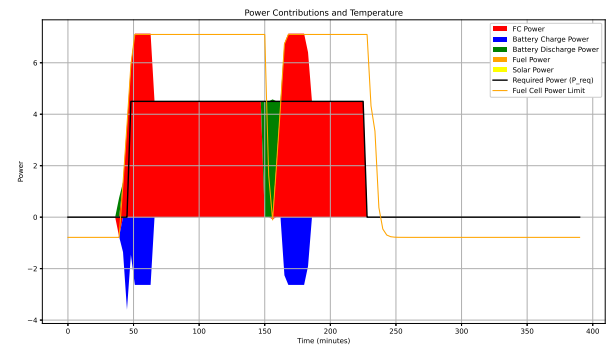


Figure 14: Step response of the system with warm up, and response to a flooding with warm up

8. Conclusions

This thesis focused on the modeling and integration of a fuel cell system into a Model Predictive hierarchical control. An experimental campaign was conducted, despite multiple hardware failures, and the fuel cell was characterized. This revealed significant deviations from the

nominal behavior, highlighting an especially low efficiency, about 32 %, and a voltage profile consistently lower than the manufacturer's data. Despite strict laboratory constraints that necessitated a simpler control logic, the experimental implementation successfully validated the actuation fidelity of the fuel cell system. The results demonstrated the robustness of the architecture to unforeseen generation. By dynamically modulating the battery buffer, the controller successfully shielded the grid from high-frequency solar fluctuations, ensuring that the net power exchange remained stable and compliant with safety profiles.

Furthermore, the simulation of the Ginostra microgrid confirmed the critical role of hydrogen in addressing seasonal imbalances, but also the challenges that this carries.

Future Developments To further advance this research, future work should focus on:

Full MPC Deployment: Implementing the complete economic optimization layer in the real-time environment to validate the cost-saving potential demonstrated in simulations.

Hardware Correction: Addressing the specific hardware faults identified during the campaign, particularly the fuel cell's drainage system failures and the DC/DC converter instabilities.

Advanced Logic Validation: Experimentally validating the simulated "warm-up" strategy, which pre-activates the fuel cell to manage thermal constraints and minimize battery depletion during startup.

based on stochastic model predictive control for off-grid microgrids. *Advances in Applied Energy*, 2:100028, 05 2021.

References

- [1] Purushottam Kumar and Sanjib Ganguly. Digital twin approach for pem fuel cell degradation prognostics using equivalent electrical circuit model and polarization curves. 2025.
- [2] 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.
- [3] Simone Polimeni, Lorenzo Meraldi, Luca Moretti, S. Leva, and Giampaolo Manzolini. Development and experimental validation of hierarchical energy management system