



REEL Demo – Romande Energie ELectric network in local balance Demonstrator

Deliverable: 3b Assessment of the performance of the
decentralized battery control strategies, based on 1 year
data

Demo site: Chapelle

Developed by

Vasco Medici, ISAAC, SUPSI

Lorenzo Nespoli, ISAAC, SUPSI

Davide Strepparava, ISAAC, SUPSI

Roman Rudel, ISAAC, SUPSI

[Mendrisio, 31.03.2022]



University of Applied Sciences and Arts
of Southern Switzerland

SUPSI

Date: 31. March 2022

Place: Bern

Publisher:

Swiss Federal Office of Energy SFOE
Research Programme, Energy - Economy - Society (EES)
CH-3003 Bern
www.bfe.admin.ch
energieforschung@bfe.admin.ch

Agent:

University of Applied Sciences and Arts of Southern Switzerland (SUPSI)
Department for Environment Constructions and Design
Institute for Applied Sustainability to the Built Environment
Via Flora Ruchat-Roncati 15
CH-6850 Mendrisio
<https://www.supsi.ch/isaac>

Author:

Vasco Medici, ISAAC, SUPSI (vasco.medici@supsi.ch)
Lorenzo Nespoli, ISAAC, SUPSI (lorenzo.nespoli@supsi.ch)
Davide Strepparava, ISAAC, SUPSI (davide.strepparava@supsi.ch)
Roman Rudel, ISAAC, SUPSI (roman.rudel@supsi.ch)

SFOE head of domain: Dr. Michael Moser, michael.moser@bfe.admin.ch

SFOE programme manager: Dr. Michael Moser, michael.moser@bfe.admin.ch

SFOE contract number: SI/501523

The author of this report bears the entire responsibility for the content and for the conclusions drawn therefrom.



Contents

1	Introduction	4
2	Pilot setup	4
3	Results	7
3.1	Evaluation of control performance	7
3.1.1	Multi-level distributed model predictive control algorithm description	7
3.1.2	Challenges in Battery control	8
3.1.3	Control performance	8
3.1.4	Peak reduction analysis	13
3.2	Evaluation of forecasting performance	18
3.2.1	Forecasting algorithm	18
4	Discussion	21
4.1	Limitations of current approach and future work	21
5	References	22



1 Introduction

SUPSI developed demand-side management (DSM) algorithms aimed at controlling several flexible resources at different levels in the grid, taking into account grid constraints. To this end, a distributed multilevel predictive control system was developed. The controller relies on a machine learning-based forecast developed in parallel. In this report, we analyze the performance of the forecasting and control algorithms applied to the management of two battery energy storage systems in the Chapelle-sur-Moudon case study.

This report is structured as follows:

- In section 2, we present the test setup in the Chapelle-sur-Moudon pilot.
- In section 3, we analyze the results, in particular:
 - In section 3.1, we analyze the control part. We describe the algorithms used and analyze their effect in the pilot.
 - In section 3.2, We deal with the load curve prediction part, briefly describing the forecasting algorithms and evaluating the performance during the test phase.
- In section 4, we discuss the limitations of the chosen approach in the available configuration and propose future developments that could improve its performance.

2 Pilot setup

The ReEL DSM demo site consists of a section of LV grid in Chapelle-Sur-Moudon with roughly 300kWp of PV installed. A total of 7 nodes are monitored, as shown in figures 1 and 2. A total of 3 batteries are installed in the grid, two of which can be directly controlled:

- Node 107: a 300 kWh district-level battery with charging and discharging power of [50, -200] kW located at the PCC. It is intended to be used for peak shaving to smooth the consumption peak and reduce the production peaks generated by the PV power plants.
- Node 106: a 40 kWh battery with charging and discharging power of [10, -20] kW is owned by an end-user and is operated to maximize self-consumption and minimize its billing costs.

Nodes were monitored at 1-second resolution by GridEye devices, which provided active and reactive power, voltage and current, per phase. The measurement history is not queryable, so data must be constantly fetched from the GridEye server every second. This is not an ideal setup, as, if for any reason, the server is not reachable, data get lost. But the system was very reliable, and very few data were lost.

The batteries are managed by Aurora's Grid, which provided us with an API that allows us to query the state of the batteries and set active power setpoints. Although GridEye offered a sampling resolution of 1 second, after some testing and discussions with Aurora's Grid engineers, we decided to set the battery data sampling and control resolution to 1 minute since the system could not handle faster paces. However, this is suitable for the application we had in mind, which was never meant to perform real-time control, as the distributed control mechanism has an iterative convergence mechanism, making it inherently slow.

An Influxdb database has been created, in which all measurement data, both from GridEye and from the batteries are stored. The database is used both to collect the data necessary to train the forecasting algorithms, and to analyze the performance of the control.

A Grafana web interface was used to visually assess data and control quality. An example of data visualization for three days of test is shown in figure 3.

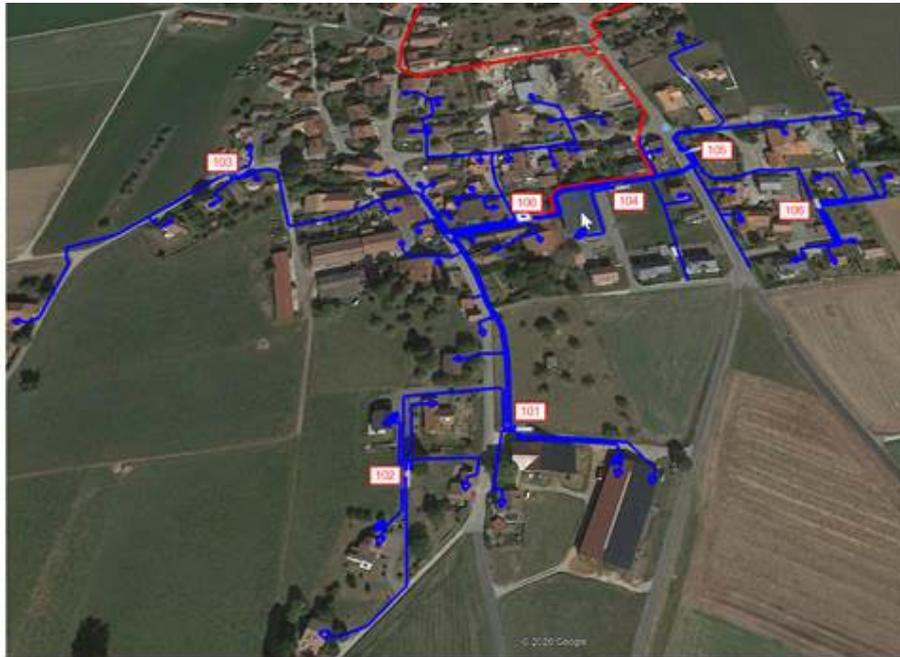


Figure 1: The ReEL demo site aerial view

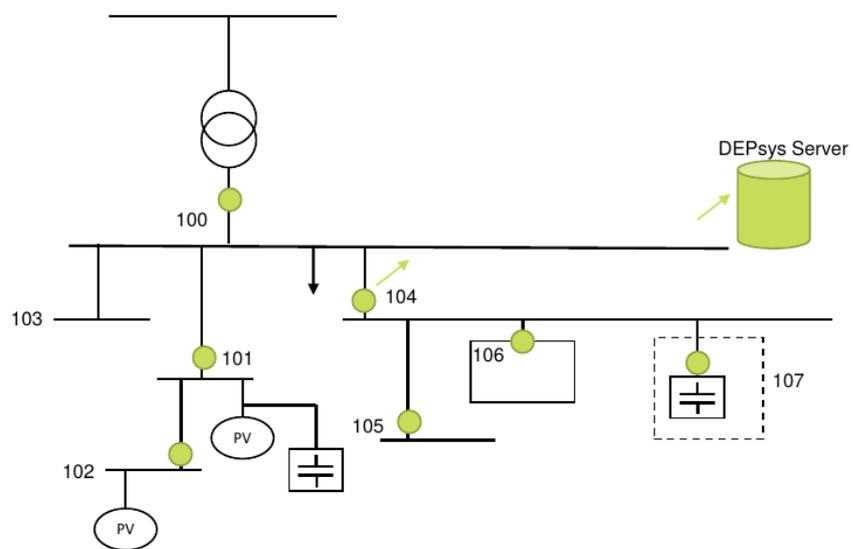


Figure 2: The ReEL demo site single line diagram

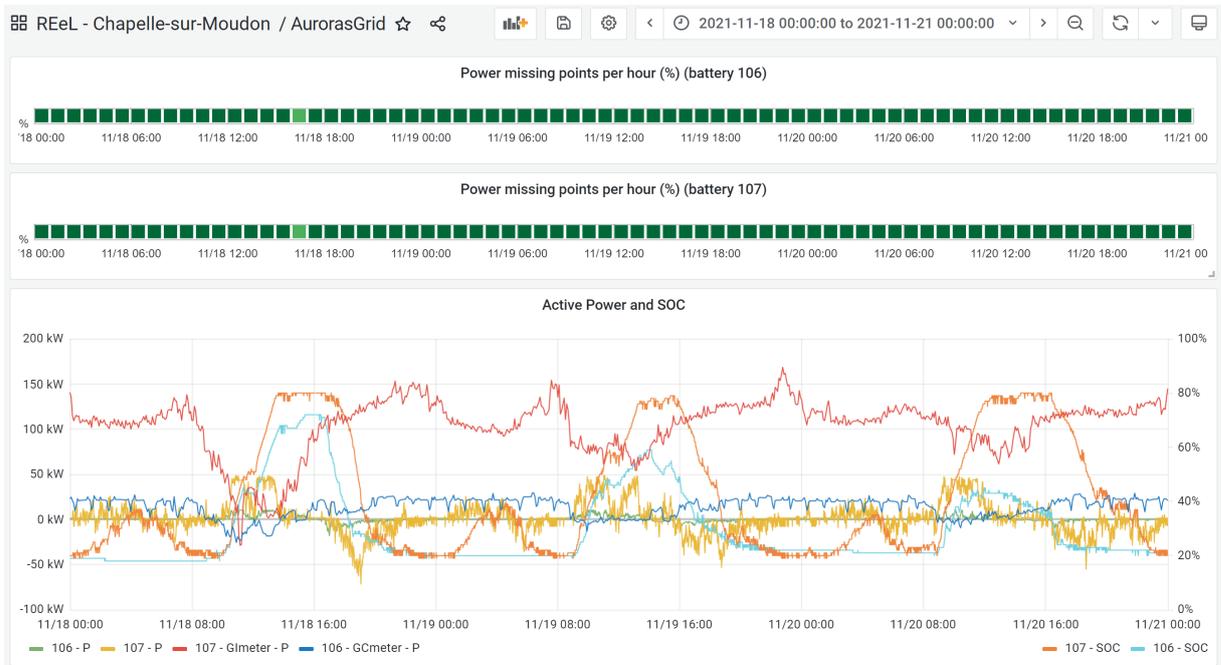


Figure 3: Pilot data visualization



3 Results

The testing phase lasted 2 months, from October 24th to December 20th 2021. A previous testing phase was planned in spring 2021, but at the time the API interface with the batteries proved to be unreliable. Since then Aurora's grid has upgraded battery hardware and software, greatly improving the reliability of the system. The first couple of weeks of successful testing phase were used to fine tune the algorithms and test the battery control interface. In the following period, unfortunately, the PV production was limited and lower than the consumption of local loads, so most of it was 100% self-consumed. Under these circumstances, we were unable to extensively test our self-consumption optimization algorithms. Figure 4 shows the time course of power at the node 106 and state of charge for the small battery that tries to optimize the self-consumption of the node, during the testing period. It can be noted that most of the time all the production is self-consumed without the need for battery intervention, and that when the battery is charged to absorb the excess PV power, it rarely reaches a state of charge (SOC) above 50%. In the following two subsections, we present the results of the test campaign. First, we focus on the performance evaluation of the forecasting algorithms, which is of paramount importance as we use predictive control methods. Then we present the results of battery control.

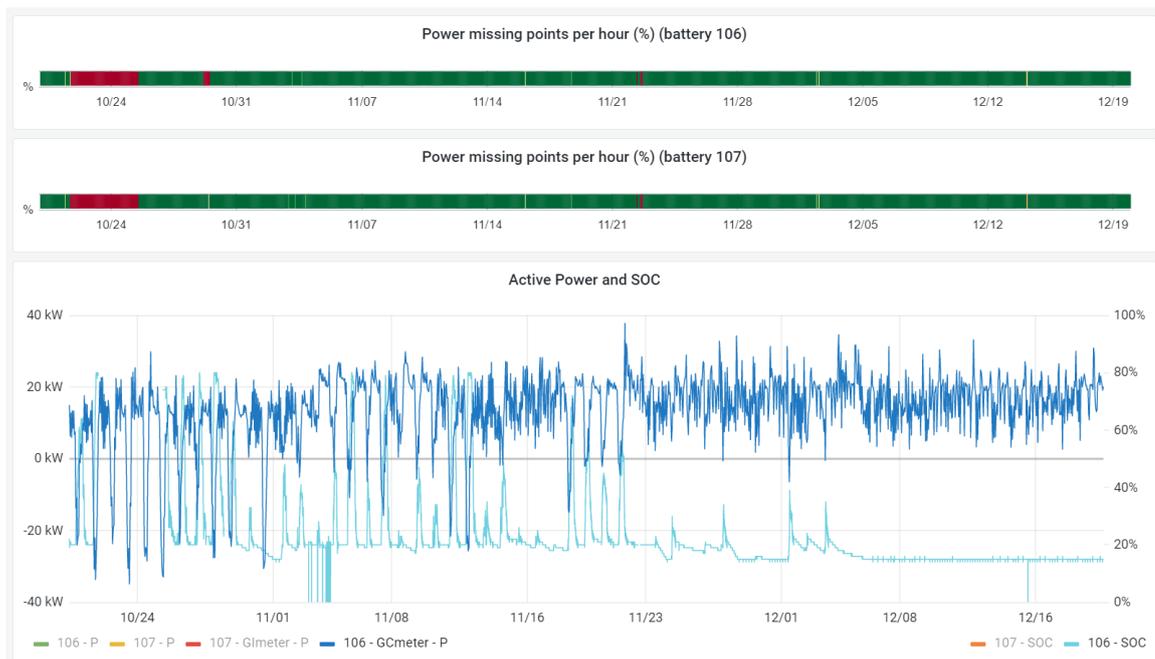


Figure 4: Time course of power at node 106 and SOC of the small battery.

3.1 Evaluation of control performance

3.1.1 Multi-level distributed model predictive control algorithm description

The multilevel hierarchical algorithm that we developed for the coordination of prosumers located in different voltage levels of the electrical grid is presented in [1]. The hierarchical structure of the grid is described by means of a rooted tree. At the top level the objective is peak shaving and valley filling,



while and in level below, the local objective of the small battery is self-consumption optimization, as summarized in table 1. Since the two objectives could be in conflict with each other, we adopted a slight

	PV	Controllable batteries	Objective	Control type
ReEI	200 kWp	40 kWh [-20, 10] kW 300 kWh [-200, 50] kW	Cost reduction Peak shaving	distributed, lexicographic

Table 1: Technical characteristics of the demo site

modification of the algorithm presented in [1]: at each iteration of the coordination process, the smaller battery performs a lexicographic optimization in which it will at first optimize for its own costs and in a second moment will try to synchronize with the higher level of the hierarchy (in this case constituted of just another agent), to perform peak shaving. In the ReEI demo, the privately owned battery (the small one) has no economic reason to synchronize with the district level battery to perform peak shaving. In fact, the objective of the private battery is to increase its own self-consumption. However, typically, several equivalent solutions for the charging and discharging operations exist, which achieve the same results in terms of self-consumption. A win-win solution is to use a lexicographic approach for the small battery: at first, an optimal scheduling for the small battery, which maximizes its owner's self-consumption, is obtained. This optimal scheduling generates a cost and a final state of charge of the battery. These can be used as constraints during the coordination needed for peak shaving. A more detailed description of the algorithm can be found in the deliverable *3d3 Design and test of distributed DSM algorithms that use communication and new forecasting models*.

3.1.2 Challenges in Battery control

The SOC reported by the battery management system (BMS) was highly inaccurate and very power-dependent (figure 5). We were unable to obtain information on how the estimation of SOC is done in the BMS, but our observations lead to the conclusion that SOC is estimated by the battery management system (BMS), almost only on the basis of the closed source voltage, which causes that when power is fed into the battery, the estimate of the state of charge suddenly increases and, conversely, when power is drawn from it, the SOC decreases. This is particularly true when the power set points have a steep ramp. We decided not to develop our own SOC estimator, as this was out of the scope of the project. Instead, to minimize the oscillation in the SOC estimate, we implemented a ramp constraint on the set point power sent to the batteries. The maximum rate of change in absolute power was set to 2kW per minute and 10kW per minute for the small and the big battery, respectively.

Another minor practical problem we encountered during the test phase is the standby power of the 300kWh battery, which appeared to be around 1.4kW. During the period in which the small battery was not operated because the PV was completely self-consumed, it required 500Wh per day in order to stay above the safe limit of 15% SOC. This is 1.25% of the total 40kWh capacity per day, which is a surprisingly high number.

3.1.3 Control performance

Batteries were operated within a safe range between 20% and 80% SOC. Those numbers were communicated to us by Aurora's Grid. Figure 6 shows a summary of battery control during the testing period. It shows the total active power at both nodes with and without batteries and the battery power as a function of the hour of the day and the day of the week. One can notice that the small battery was charged mainly from the excess PV power and mostly discharged during the evening peaks. The outlier 5kW charging events that spread throughout the day are emergency charging events, which were automatically triggered in our software when the SOC of the battery would fall below 15%.

The big battery instead was allowed to charge from the grid, and therefore we can see that it tends to absorb energy in periods in which the power at node 100 is low or negative and reinject energy when the consumption at the node is high, during the morning and evening peaks. Since at node 106, PV produc-

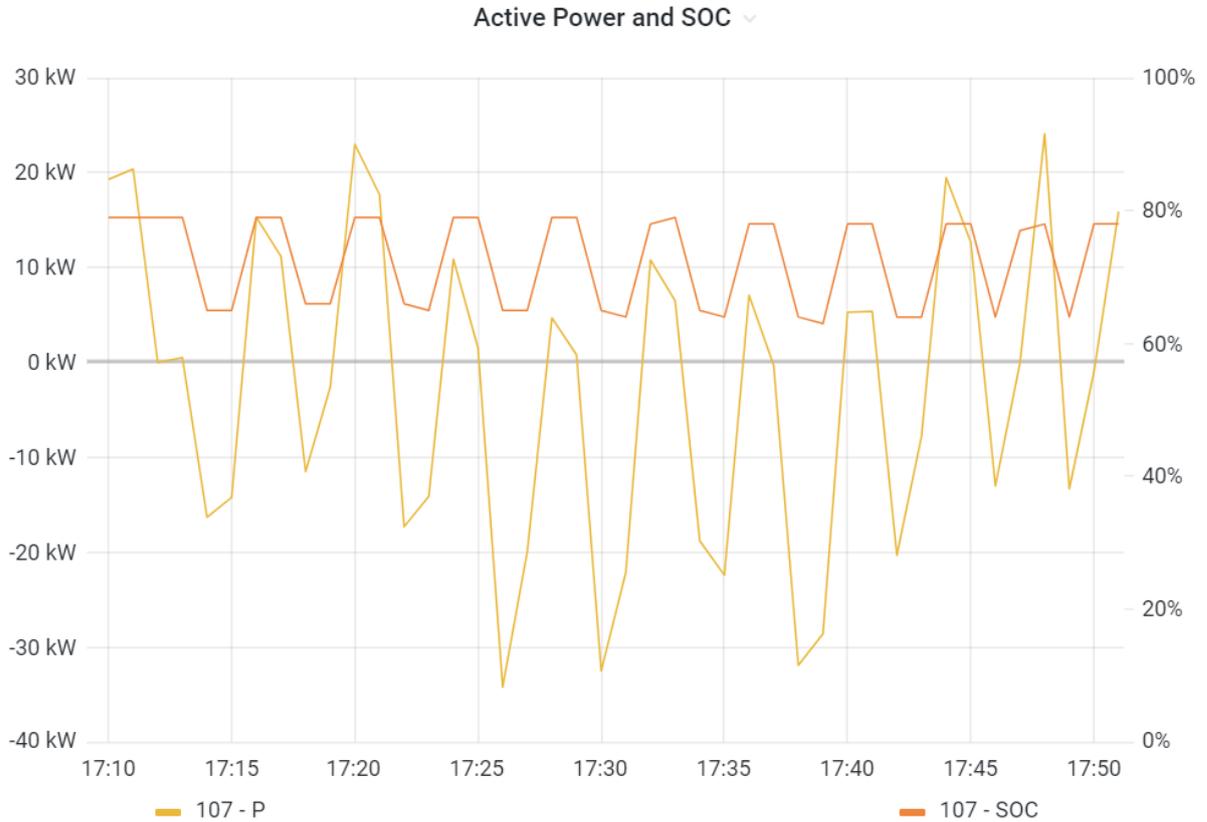


Figure 5: Oscillations in SOC readings as a function of the active power requested.

tion was not sufficient for the vast majority of the time, the advantage given by coordinating the batteries with each other could not be evaluated on the empirical data. However, we have demonstrated the advantage of coordination in the deliverable *3Overall-Cross-site Comparison of the performance of different DSM strategies. Investigation of the possible conflicts (LIC and Chapelle-sur-Moudon) and sub-optimality issues*, in simulation. The algorithm employed is the same, and it worked as specified. We focus here on analyzing the impact of control at the coupling point of the pilot network (node 100). The main objective of the control was to reduce power fluctuations at node 100, through an objective function that quadratically punished them.

To better assess control quality, we decided to calculate what would have happened under ideal conditions. We simulated ideal batteries with the same specifications as those used in the project and used perfect forecasts to control them. The batteries were modelled in simulation with a simple 1-state model, using constant charge and discharge efficiencies of 95% and a self-discharge of 3% per month. Simulating ideal batteries allows us to better assess the performance of our control, against a best-case scenario. We also decided to evaluate the effect of imposing a ramp constraint on the battery power set point. Since we added it to overcome the problem in the SOC readings we received from the battery, we wanted to assess whether this would hinder the control under ideal conditions.

Figure 7 shows the distribution of the active power at the coupling point of the test network (node 100) during the period in which the algorithms were activated. The *no_bat* case represents the power measured at the node 100 after the battery actions were subtracted from it. Losses were not taken into account. We can see that our algorithms effectively reduce the power excursions at node 100 and flatten the profile. However, we can also notice that the distribution's right tail is much better in the ideal case. This happens primarily because the forecasts are not perfect. The bias in the forecasts probably plays a crucial role. We often observed that the battery was depleted too early, leading to a peak in



consumption. The results would probably be much better if we managed to reduce the bias.

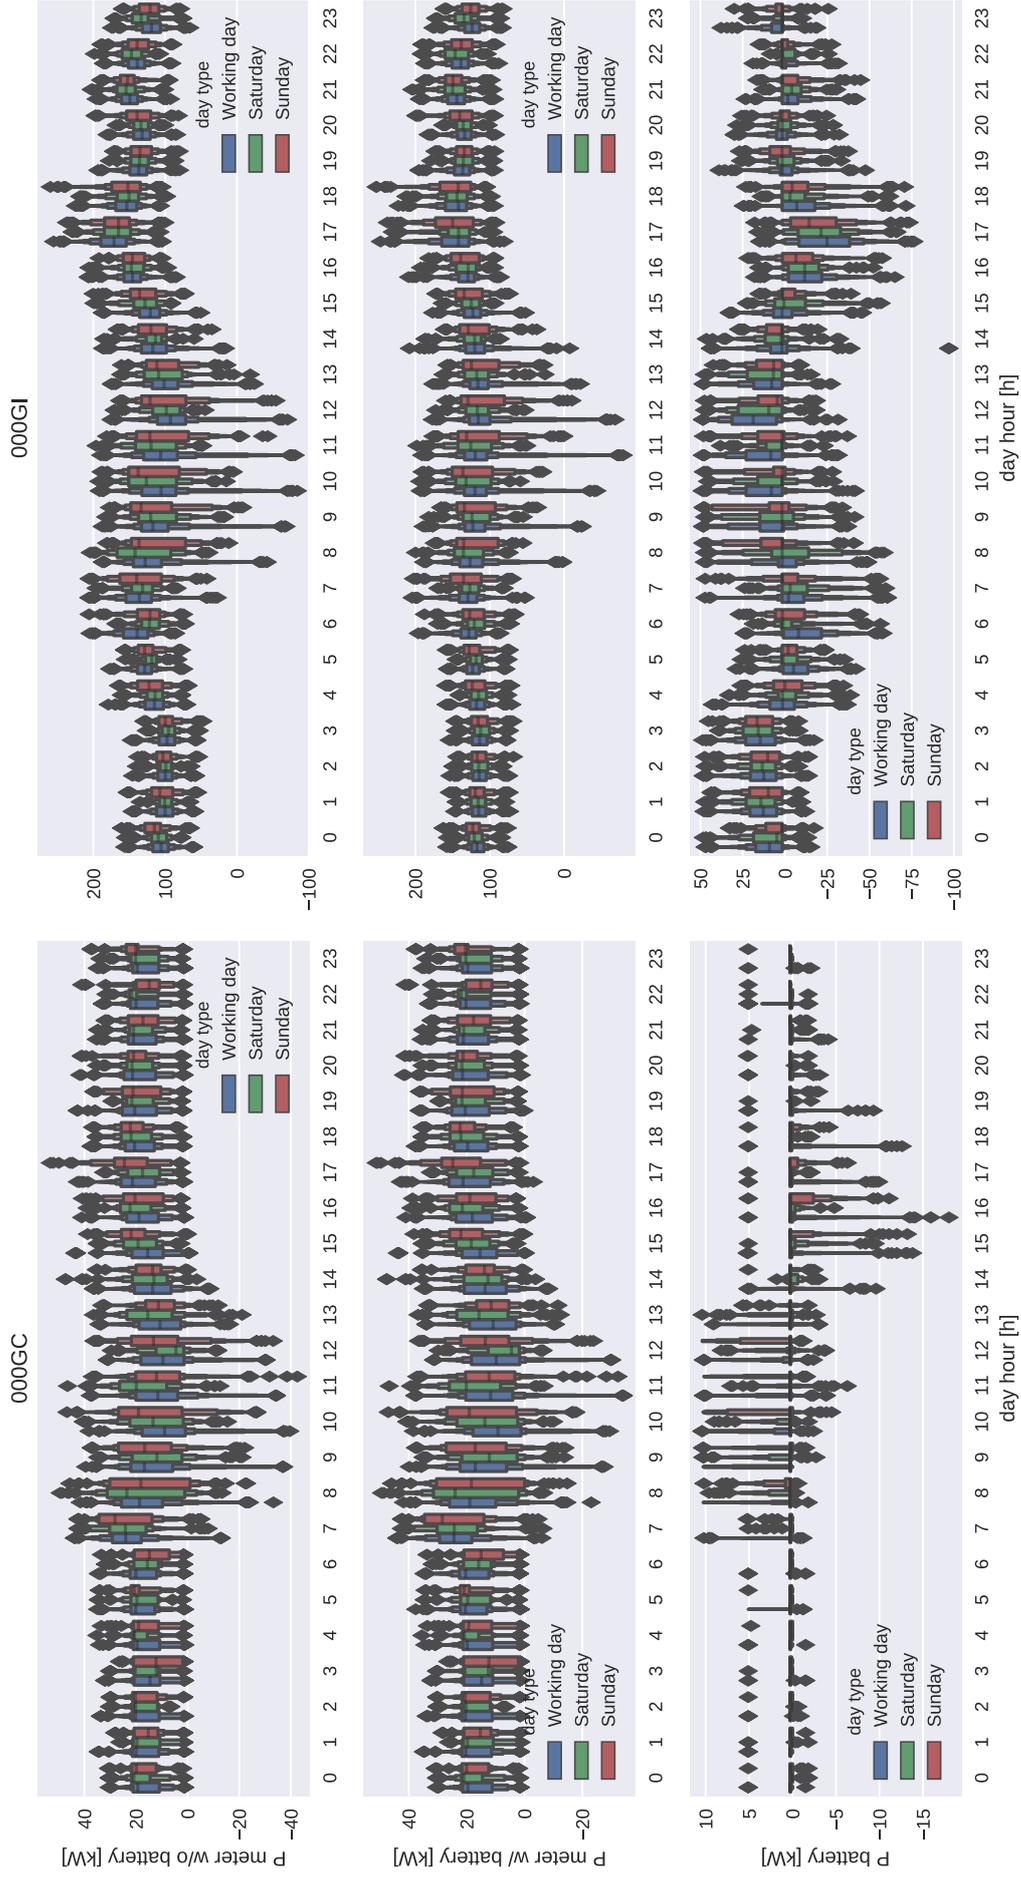


Figure 6: Batteries action as a function of the week and hour of the day, during the testing period.

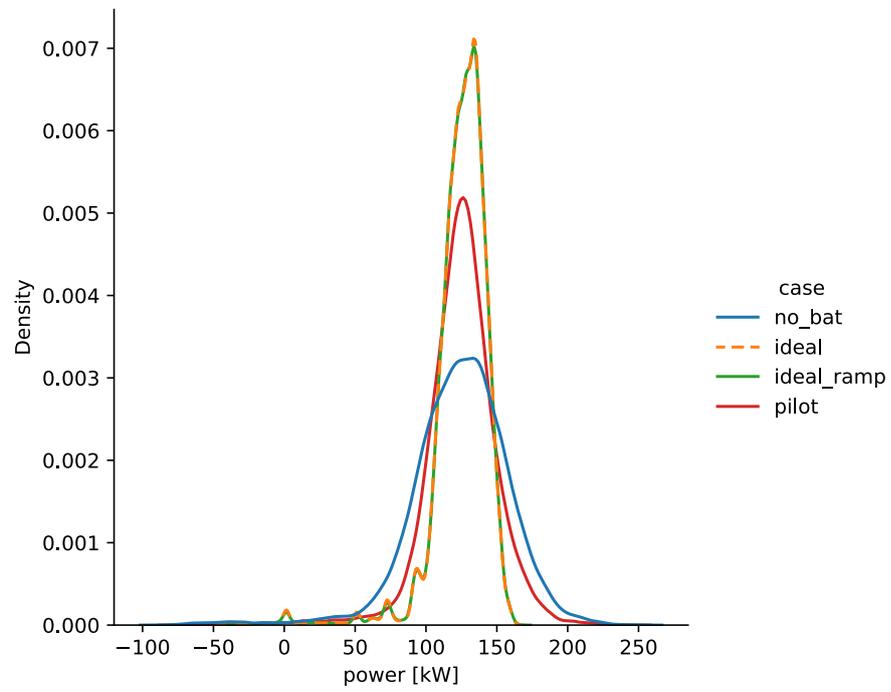


Figure 7: Power distribution during the test period, sampled with 1-minute resolution.



3.1.4 Peak reduction analysis

We analyzed the performance of the algorithms for peak shaving. In this case, we evaluated the reduction and the ideal potential reduction in daily peaks with a sampling of 1-minute (or control sampling time) and 15-minute peaks (the usual peak invoicing sampling time in Switzerland). For each day of the test phase, we calculated the maximum power consumed, after averaging the values according to the sampling time.

Figure 8 and 9 show the time-series of daily peaks during the testing period in which both batteries were actively controlled, with a sampling time of 1 minute and 15 minutes, respectively. Here one can clearly see that the ramping constraint imposed on the battery affects the high-resolution peak reduction capabilities. This effect disappears when the sampling resolution is 15 minutes.

The empirical cumulative density function of the peak reduction is shown in figures 10 and 11 for the sampling resolutions of 1 and 15 minutes, respectively. The average relative and absolute daily peak reduction is summarized in table 2. The algorithm reached slightly above 1/4 of the maximum possible peak reduction in 1-min resolution, and slightly above 1/3 of the maximum reduction in 15-min resolution. For a more detailed analysis of the cost, refer to deliverable 3f - *Assessment of investment costs of controllable batteries and comparison with grid refurbishment*. However, it is important to notice that the peak is billed on the monthly and not on the daily maximum, which makes the task of optimizing peak costs very difficult, since it only takes the algorithm to make a mistake once a month to undo the effect of days in which the controller actually helped lower the peak. In this case, it might be interesting for the DSO to revise the peak billing method to make the service more attractive, which in any case helps to reduce losses on the network and extend the life of components on the network.

controller	1 min relative reduction [%]	1 min absolute reduction [kW]	15 min relative reduction [%]	15 min absolute reduction [kW]
ideal	35.59	76.12	29.75	58.18
ideal_ramp	32.36	69.30	29.75	58.17
pilot	8.41	17.86	11.02	21.28

Table 2: Peak reduction during test period

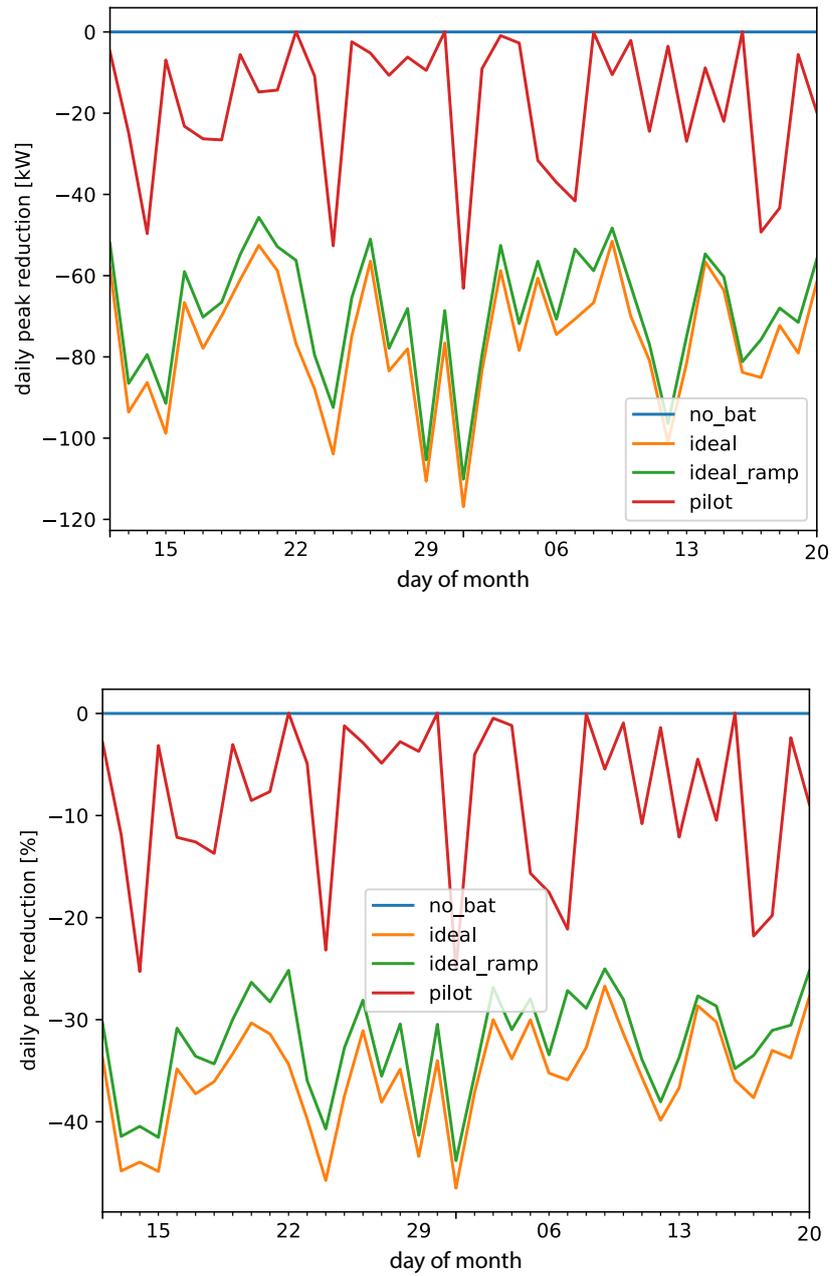


Figure 8: Time-series of peak in the period between 12.11.21 and 21.12.21, sampled with 1-minute resolution.

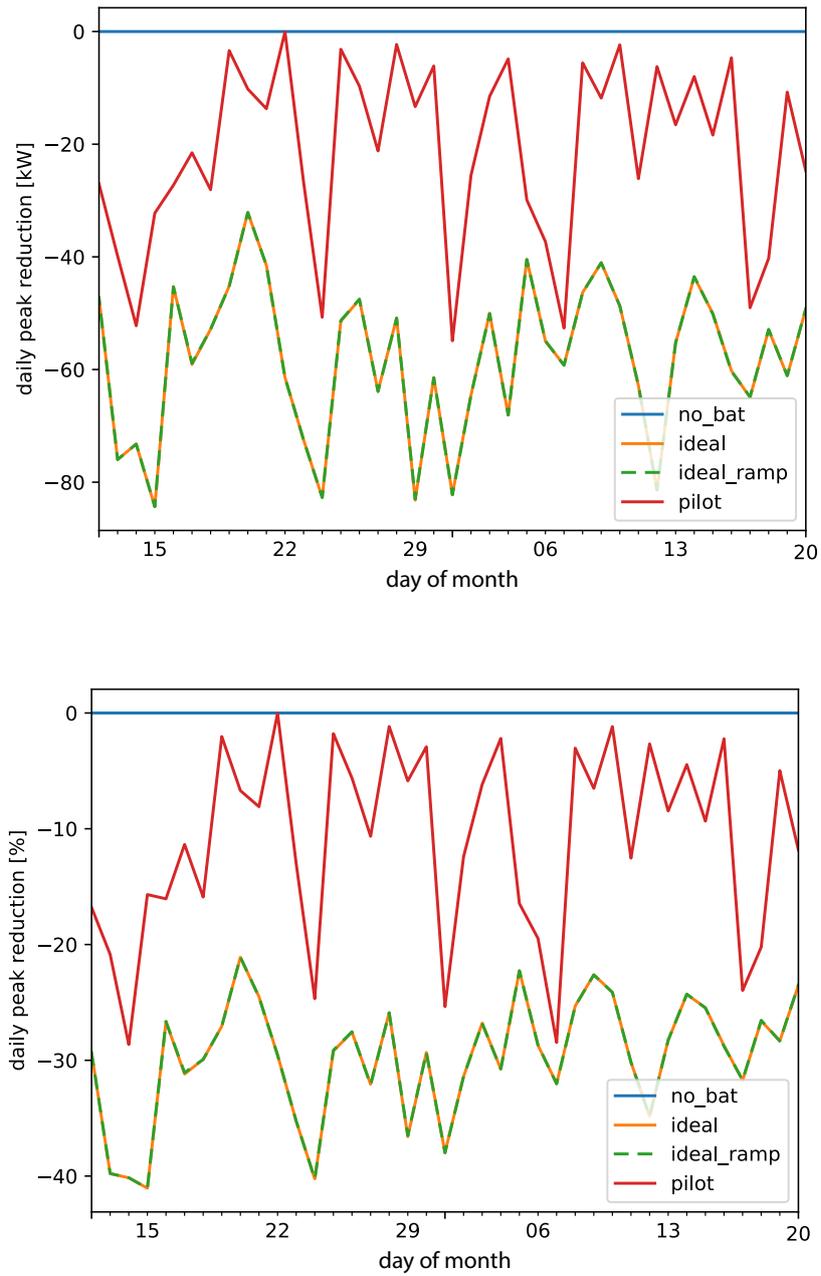


Figure 9: Time-series of peak in the period between 12.11.21 and 21.12.21, sampled with 15-minute resolution.

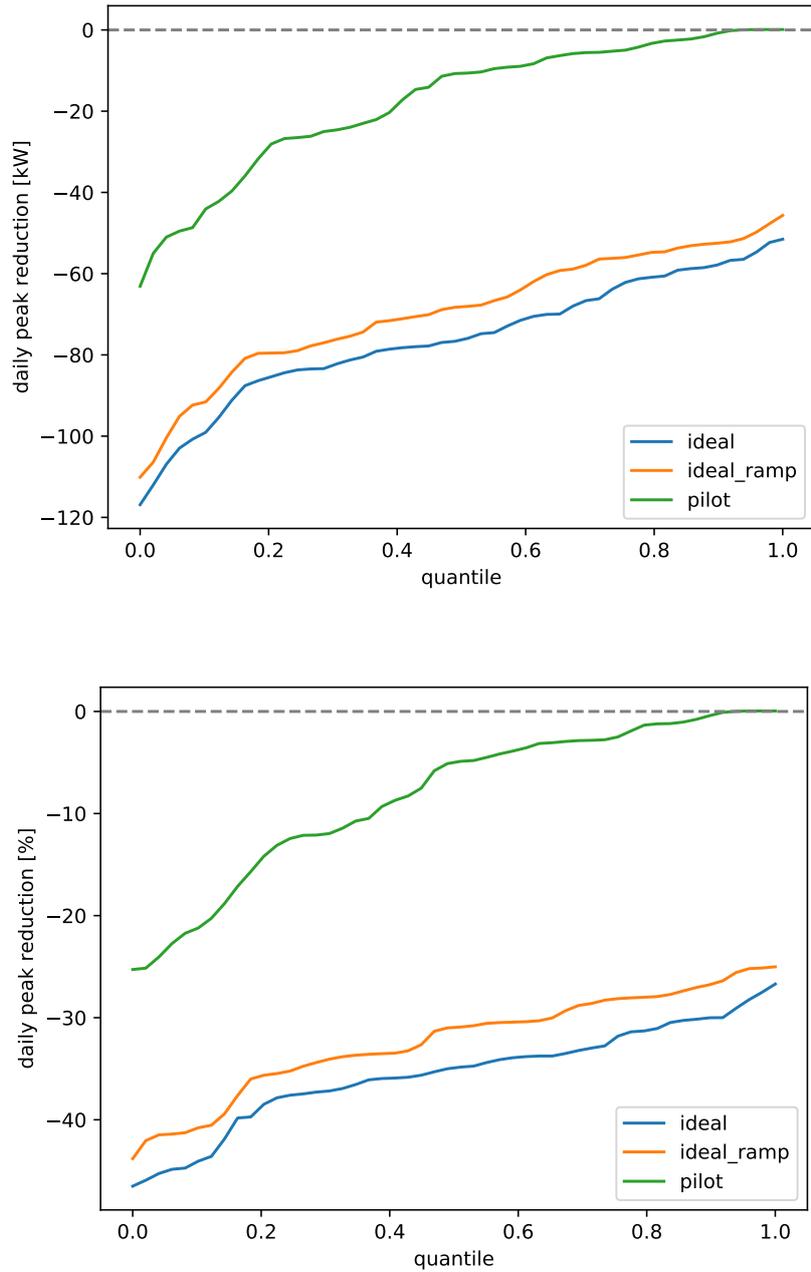


Figure 10: Empirical CDF of peak in the period between 12.11.21 and 21.12.21, sampled with 1-minute resolution.

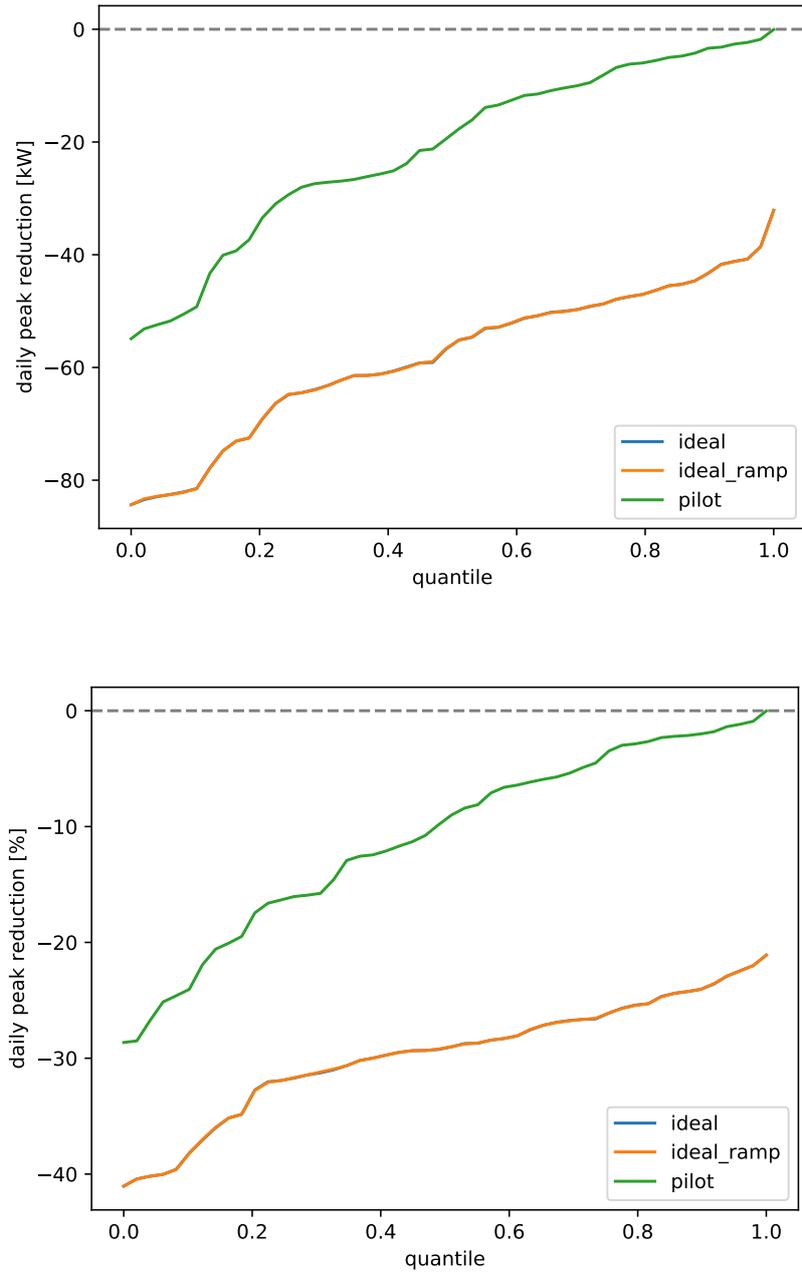


Figure 11: Empirical CDF of peak in the period between 12.11.21 and 21.12.21, sampled with 15-minute resolution.



3.2 Evaluation of forecasting performance

3.2.1 Forecasting algorithm

In the previous deliverable *3d3 Design and test of distributed DSM algorithms that use communication and new forecasting models comparison CsM e LIC*, we analyzed the performance of different forecasting models and have used them to assess the closed-loop performance of storage and load control. We focused on methods that have already proved to be accurate in forecasting 24 hours ahead residential power profiles. In particular, we tried to improve the performance of the methods that were tested in [2] and [3], and focused on the Holt-Winters method and on different forecasting techniques exploiting gradient boosted models (GBM), a family of competition-winning, general-purpose, non-parametric regressors, which exploit sequential model fitting and gradient descent to minimize a specific loss function.

In the pilot phase we decided to use a forecaster based on the Light Gradient Boosting Machine (LightGBM), a general-purpose, non-parametric regressor, which exploit sequential model fitting and gradient descent to minimize a specific loss function. We applied a preliminary causal embedding of the explanatory variables, in order to capture seasonal effects. Starting from the original time series $s \in \mathcal{S}$, a predictors (or regressors) matrix X and a target matrix Y are obtained. Given a dataset with T observations, a prediction horizon of h steps ahead, and an history embedding of e steps, we obtain the Hankel matrix of targets $Y \in \mathbb{R}^{(T-h-e) \times h}$, and the Hankel matrix of the past regressors, $X_p \in \mathbb{R}^{(T-h-e) \times n_x e}$, where n_x is the number of regressors. Verbosely, X_p and Y can be written as:

$$X_p = \begin{bmatrix} x_{1,t-e} & x_{1,t-e+1} & \dots & x_{1,t} & x_{2,t-e} & \dots & x_{n_x,t} \\ x_{1,t-e+1} & x_{1,t-e+2} & \dots & x_{1,t+1} & x_{2,t-e+1} & \dots & x_{n_x,t+1} \\ x_{1,T-2h} & x_{1,T-2h+1} & \dots & x_{1,T-h} & x_{2,T-2h} & \dots & x_{n_x,T-h} \end{bmatrix} \quad (1)$$

$$Y = \begin{bmatrix} y_{t+1} & y_{t+2} & \dots & y_{1,t+h} \\ y_{T-h+1} & y_{T-h+2} & \dots & y_T \end{bmatrix} \quad (2)$$

where $x_{1,t}$ stands for the first regressor at time t .

In our case, the battery controller runs every 1 minute, solving an optimization problem using model predictive control over a receding horizon of 24 hours. The optimization strategy implements a variable step length: the first 5 steps have a 1 minute duration, the following 2 have a 5 minutes duration, and the rest 15, as shown in figure 12. This way, we reduced the number of steps h from 1440 to 102. The past regressors matrix X_p is then augmented with categorical time features, e.g. day of week, and NWP variables, to obtain the final regressors matrix X .

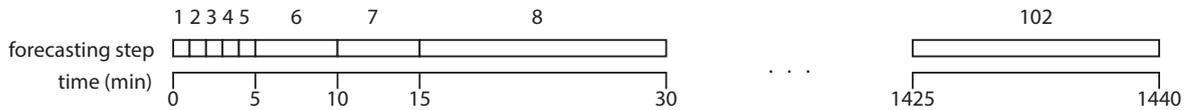


Figure 12: Forecasting steps and their duration.

In order to realize the control we depicted in section 3.1, we need to forecast the time-course of power at two points:

- At node 106, to optimize local self-consumption
- At node 100, to perform peak shaving and valley filling

The forecasting algorithms, described in 3.2.1, were trained based on one year of historical grid data, measured by the GridEye devices installed at the nodes 100 and 106. Unfortunately, we didn't get access to the battery data for the entire period, when available, the battery actions were subtracted from the power measured at the nodes.



The accuracy of the forecasting varied as a function of the forecasting step and of the hour of the day. Figure 13 shows the absolute error made by the forecasting algorithm during the test period. Not surprisingly, the closer in time the forecasts are to the time at which they were produced, the better they are. It can also be noted that the time of day when forecasts are most difficult to drill down is around noon, which is due to the fact that PV is not easy to predict, even if numerical weather predictions are fed into our forecasters. The same data are shown in Figure 14, where the mean absolute error (MAE) is calculated as a function of time of day and number of steps ahead.

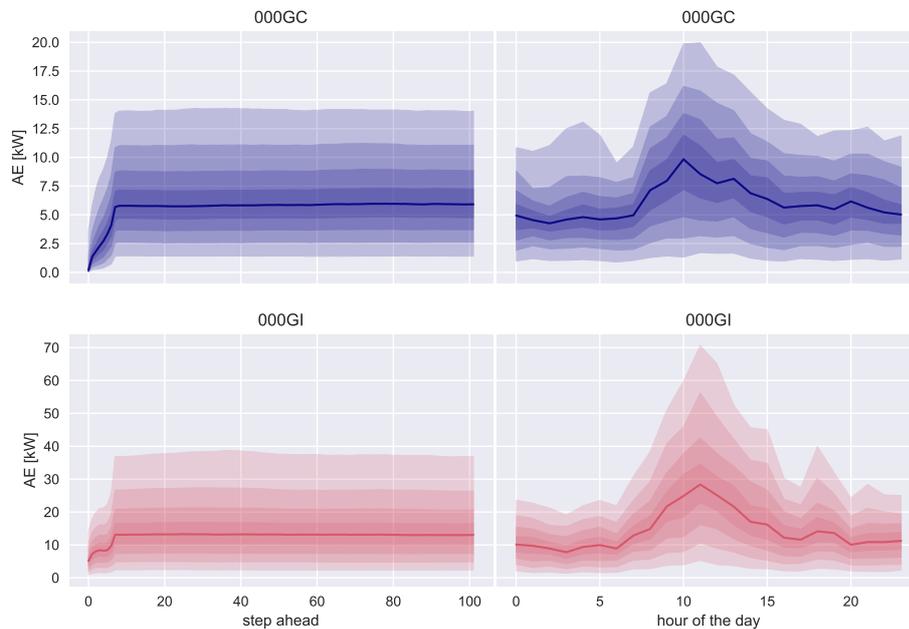


Figure 13: Distribution of absolute error for the forecasters at the nodes 106 (000GC) and 100 (000GI). Left: error distribution as a function of the step ahead. Right: error distribution as a function of the time of the day.

Since our forecasts are used to control a battery, it is important to analyze the forecast error not only in terms of MAE or RMSE, but it is equally important to understand whether the forecaster is biased in underestimating or overestimating the power trend at the node. Overestimating PV production, for example, would lead to a control strategy that does not fully charge the battery, underestimating the consumption would lead to a premature battery depletion. Figure 15 shows that in both cases the forecaster displays a significant positive bias. In our case, the residuals were calculated as $e_t = P_{observed_t} - P_{forecast_t}$, which means that a positive bias results in an underestimation of the total consumption at the forecast nodes. This has a consequence on the control, as we will show in section 3.1.4. We hypothesize that this bias is due to the fact that we could not subtract battery actions from the training data, which led to an underestimation of grid consumption since, at least at some times, it was masked by battery actions. However, it is possible that there was an evolution in the type and number of loads present at the nodes.

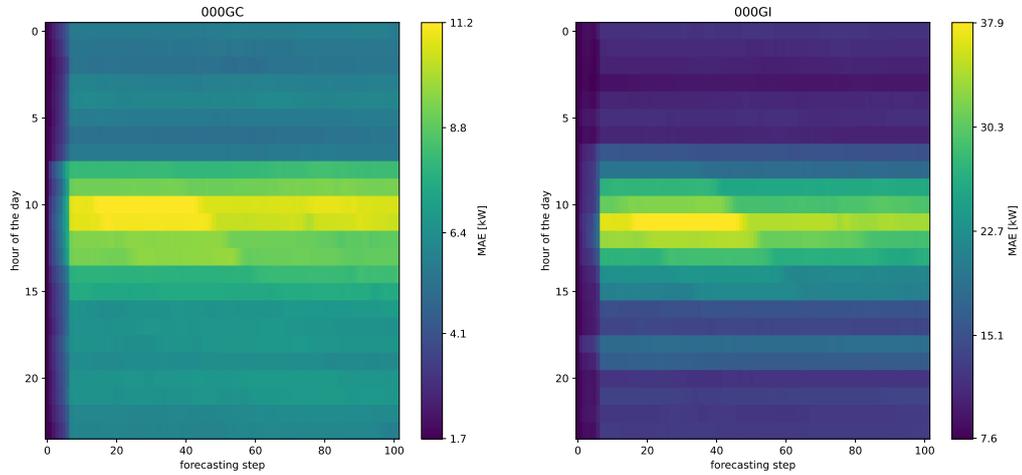


Figure 14: MAE of the forecasters as a function of forecasting step and hour of the day. Left: node 106 (000GC), Right: node 100 (000GI).

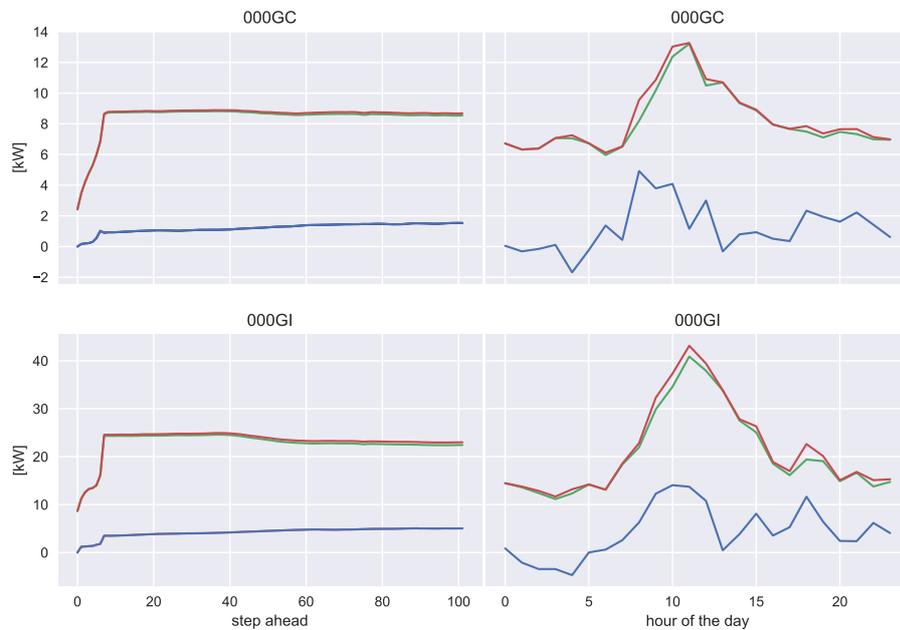


Figure 15: Decomposition of root mean squared error into bias and standard deviation for the forecast at the nodes 106 (000GC) and 100 (000GI). blue lines: bias, green line: standard deviation, red line: RMSE. Left: as a function of the step ahead. Right: as a function of the time of the day.



4 Discussion

4.1 Limitations of current approach and future work

Controllers were initially designed with many flexible loads and batteries distributed throughout the grid in mind. This was based on what was developed in FURIES and for the Gridsense commercial product. We focused on the development of algorithms that would work at different voltage levels in the grid, taking into account grid constraints. Unfortunately, only two batteries could be controlled in this case, and for a limited time.

One of the limitations of the chosen approach is that it relies on deterministic optimization. For the specific problem of peak shaving, it would be particularly useful to use an approach that accounts for uncertainty in predictions of energy production and consumption, i.e., robust or stochastic control. A possible extension of our formulation, which would certainly help to achieve better results in peak shaving, would be the implementation of multilevel stochastic control.

Multilevel stochastic control requires to model the joint probability distribution of the power at all the levels, and to take joint decisions, for example by constructing a multivariate stochastic tree with a number of dimensions equal to the number of agents. As the initial number of agents was thought to be high, this wasn't considered feasible and we opted instead for a deterministic formulation. But in this case, in which only two batteries are controlled, it would be a sound approach.



5 References

- [1] L. Nespoli and V. Medici, "Constrained hierarchical networked optimization for energy markets," in *Proceedings - 2018 IEEE PES Innovative Smart Grid Technologies Conference Europe, ISGT-Europe 2018*, 2018.
- [2] L. Nespoli, "MODEL BASED FORECASTING FOR DEMAND RESPONSE STRATEGIES," Ph.D. dissertation, 2019.
- [3] L. Nespoli, V. Medici, K. Lopatichki, and F. Sossan, "Hierarchical demand forecasting benchmark for the distribution grid," *Electric Power Systems Research*, vol. 189, no. August, p. 106755, 2020. [Online]. Available: <https://doi.org/10.1016/j.epsr.2020.106755>