Schweizerische Eidgenossenschaft
Confédération suisse
Confederazione Svizzera
Confederaziun svizra

Department of the Environment,
Transport, Energy and Communication DETEC

Swiss Federal Office of Energy SFOE
Energy Research

Final report

MuLDeR

Multi-Level Demand Response

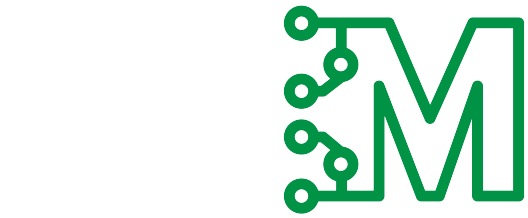

M U L D E R

University of Applied Sciences and Arts
of Southern Switzerland

SUPSI

Date: 20 December 2020

Place: Bern

Publisher:
Swiss Federal Office of Energy SFOE
Research Programme, Grid
CH-3003 Bern
www.bfe.admin.ch
energieforschung@bfe.admin.ch

Author:
Lorenzo Nespoli, ISAAC, SUPSI (lorenzo.nespoli@supsi.ch)
Federico Rosato, ISAAC, SUPSI (federico.rosato@supsi.ch)
Davide Strepparava, ISAAC, SUPSI (davide.strepparava@supsi.ch)
Vasco Medici, ISAAC, SUPSI (vasco.medici@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/501705

Summary

A comprehensive approach towards the actuation of flexibility, taking into account its effects at all grid levels and for all actors involved, is crucial for a successful implementation of demand response (DR). The goal of the MuLDeR project is to develop grid-aware mechanisms for the activation of demand response. Different methods to coordinate flexible demand and storage have been investigated. In particular, a multi-level hierarchical distributed control scheme making use of trusted aggregators at different levels of the grid has been studied in detail. The possibility of doing without aggregators using a gossip-based coordination mechanism based on a peer-to-peer network has also been explored. Mechanisms that allow offering DR services to all actors involved and at all grid levels while making sure that the control actions are not violating grid constraints throughout the entire network and that power quality is guaranteed have been studied. The applicability of the proposed solutions has been evaluated in simulation and is being tested in a pilot project, in the region of Lugano (https://lic.energy).

Résumé

Une approche globale de l'activation de la flexibilité, prenant en compte ses effets à tous les niveaux du réseau et pour tous les acteurs impliqués, est cruciale pour une mise en œuvre réussie de la réponse à la demande (DR). L'objectif du projet MuLDeR est de développer des mécanismes conscients de l'état du reseau pour l'activation de la réponse à la demande. Différentes méthodes de coordination de la demande flexible et du stockage ont été étudiées. En particulier, un schéma de contrôle distribué hiérarchique à plusieurs niveaux utilisant des agrégateurs fiables à différents niveaux du réseau a été étudié en détail. La possibilité de se passer d'agrégateurs en utilisant un mécanisme de coordination basé sur les ragots et sur un réseau d'égal à égal a également été étudiée. Des mécanismes permettant d'offrir des services de RD à tous les acteurs concernés et à tous les niveaux du réseau tout en s'assurant que les actions de contrôle ne violent pas les contraintes du réseau dans son ensemble et que la qualité de l'électricité est garantie ont été étudiés. L'applicabilité des solutions proposées a été évaluée en simulation et est actuellement testée dans le cadre d'un projet pilote, dans la région de Lugano (https://lic.energy).

Sommario

Riassunto. Un approccio globale all'attuazione della flessibilità, che tenga conto dei suoi effetti a tutti i livelli della rete e per tutti gli attori coinvolti, è fondamentale per un'implementazione di successo del demand response (DR). L'obiettivo del progetto MuLDeR è quello di sviluppare meccanismi grid-aware per l'attivazione del DR. Sono stati studiati diversi metodi per coordinare la domanda flessibile e lo stoccaggio. In particolare, è stato studiato in dettaglio uno schema di controllo distribuito gerarchico multilivello che fa uso di aggregatori a diversi livelli della rete. È stata anche esplorata la possibilità di fare a meno degli aggregatori utilizzando un meccanismo di coordinamento basato sul gossip e su una rete peer-to-peer. Sono stati studiati meccanismi che permettono di offrire servizi di DR a tutti gli attori coinvolti e a tutti i livelli della rete, assicurandosi che le azioni di controllo non violino i vincoli della rete in tutta la rete e che la qualità dell'energia sia garantita. L'applicabilità delle soluzioni proposte è stata valutata in simulazione ed è in fase di test in un progetto pilota, nella regione di Lugano (https://lic.energy).

Contents

List of abbreviations

ADMM	Alternating Direction Method of Multipliers
AMM	Automated Market Making
AT	Auditable Tariff app
BAU	Business As Usual
D-OPF	Distributed Optimal Power Flow
DSM	Demand Side Management
DSO	Distribution System Operator
EPNE	Ex-Post Nash Equilibrium
GBM	Gradient Boosted Models
GDI	Grid Dependence Index
GHI	Global Horizontal Irradiance
HW	Holt Winters
IC	Incentive Compatible
IR	Individual Rationality
ISO	Independent System Operator
KPI	Key Performance Indicator
LGB	LightGBM
LIC	Lugaggia Innovation Community
LV	Low Voltage
MAE	Mean Absolute Error
MPC	Model Predictive Control
NE	Nash Equilibrium
NILM	Non-Intrusive Load Monitoring
NWP	Numerical Weather Predictions
OPF	Optimal Power Flow

P2P	Peer to Peer
PCC	Point of Common Coupling
PDF	Probability Density Function
pFB	preconditioned FOrward Backward
PV	Photovoltaic
RES	Renewable Energy Sources
SCC	Self Consumption Community
SOC	State Of Charge
SPB	Sharing Problem Based
VCG	Vickrey-Clarke-Groves
VNE	Variational Nash Equilibrium
VSC	Voltage Sensitivity Coefficients

1 Introduction

It is widely recognized that in the future the intelligent activation of demand response (DR) will contribute to a more reliable power system and higher price stability on the power markets and aid the transition to a new paradigm of distributed generation and storage, by minimizing the need for new infrastructural investments. With the penetration of renewable energy sources (RES) and storage, the demand side is becoming increasingly capable of providing flexibility services and contributing to a reliable power system and price stability on power markets. The coordination needed to provide flexibility services at different grid levels implies substantial load correlation and openly challenges the traditional assumption of "statistical smoothing" under which distribution system operators (DSOs) have designed their grids; let alone the effects of the ever-increasing penetration of photovoltaic (PV) systems at the low voltage level. As flexible generation and consumption capacity will be more and more fragmented and distributed, to better exploit it and maximize its economic profitability, many flexibility owners will be required to coordinate with each other, when responding to demand response (DR) signals. This requires a comprehensive approach towards the actuation of flexibility, taking into account the effect of DR at all grid levels. In particular, one needs to ensure that DR does not create congestions or voltage violations at any point of the distribution grid.
Proper regulatory mechanisms defining rules for remunerations and compensations need to be set up to allow the integration of DR services in the wholesale market. Such practices need to address the problem of conflicting interests between market participants. Smart algorithms need to be devised to work in this framework. A market design that rewards flexibility needs to be set up. The highly stochastic nature of RES generation calls for a market that is able to operate in near real-time.
The goal of the MuLDeR project is to develop grid-aware mechanisms for the activation of demand response in different voltage levels of the distribution grid. In order to achieve this goal, new incentive schemes for the remuneration of flexibility and new technical solutions for the coordination of the controllable users' devices must be created. Distributed optimization, which consists in collectively solving an optimization problem by decomposing it into simpler sub-problems being solved by distributed agents, poses itself as an ideal solution for a harmonized actuation of DR, as it brings together the optimality of a centralized solution and the scalability of a decentralized one. On the other hand, single market participants' interests can be conflicting, and it is imperative to design a market that rewards flexibility actuation fairly, discourages cheating and preserve privacy. The method developed in Mulder relies on distributed control and game theory. It counts on a hierarchical structure of aggregators, reflecting the voltage level separations of the electrical grid. In this approach, upper levels of the hierarchy have only access to aggregated information from the underlying levels, allowing scalability and privacy preservation. Moreover, this scheme considers only communication regarding power consumption/production forecasts for the agent and does not imply sharing private information about one's own system parameters, not even to the trusted party. The energy market that we designed has the following four core properties:

1. the market rules generate a unique (generalized) Nash equilibrium (NE);
2. the market rules ensure that no one is better off leaving the designed market, in expectation (Individual Rationality - IR).
3. the market incentivizes the agents telling the truth (Incentive Compatible - IC);
4. deviation of the actual behaviour from the one established at consensus is discouraged and/or can be efficiently and correctly identified as either bona fide or fraudulent, and the consequent actions taken are appropriate and fair.
5. the market will take into account grid constraints, making agent pay fees dependent on the state of the grid (estimated with the convex formulation), while avoiding unfair treatment of any agent due to their particular conditions. Switzerland is embracing a causal principle on the price formation for end users, as stated in the recent modification to the Federal Electricity Supply Act [1]. This means that the electrical bills "should reflect costs caused by end users", which is in line with this market property.

6. the NE can be obtained in reasonable computation time, with a protocol connecting different voltage levels.

Under the first three conditions, the market switches from competitive to cooperative and end users won't have any incentive in deviating from the actions computed by the algorithm deployed on their smart devices. The fourth point will ex-post identify deviations from the declared actions and use them to punish malicious agents. Methods that do not require aggregators' presence have also been evaluated and are discussed in this report. Our conclusion is that, although the class of problems considered (sharing problems) can be decentralised through gossip protocols, the number of messages required before convergence (and thus the overall computation time) is higher than for a parallel solution scheme. Moreover, the presence of coupling constraints and the fact that the goal is a function of the actions of all users require an additional cryptographic layer to preserve the privacy of private information.
Besides the definition of price formation scheme and market mechanism, a coordination protocol is also needed in order to converge to the market equilibrium in a multilevel hierarchical setting. This requires to define how information flows through the different levels of the network and which control algorithm is employed. Since we desire to solve the market without disclosing internal parameters of the end user's private devices nor constraints, we chose to adopt an iterative solution based on distributed control. For the hierarchical formulation, we designed a forward-backward communication protocol, in which agents from different levels of the grid send their forecasted power production or consumption to an aggregator in the upper level of the grid. Then, the aggregators send back a coordination signal to the agents, based on the design of the energy market. The process is carried out in an iterative way until convergence is reached.
A simulation framework was designed to benchmark the proposed market formulation and coordination mechanisms, both from the grid and economic point of view. The results are presented and discussed in detail here.
The market design and control strategy are currently being tested in the pilot project Lugaggia Innovation Community (LIC, lic.energy). LIC is a self-consumption community that optimises and automates the use of local solar energy among 19 prosumers in the same district, in combination with a public battery system. In LIC, smart meters have been equipped with computing units capable of solving the distributed optimisation problem designed in this project, and a privacy-preserving method has been implemented to calculate the redistribution of money. Tests are currently underway. The pilot configuration and preliminary results are presented are the end of this report.
The report is structured as follows: first a literature review on market formulation and distributed control is presented in section 2. Then, the energy market formulation is presented in section 3. This is followed by the design of the control algorithms that optimise the economic rewards within the designed market, in section 4. Section 5 focuses on extending the market formulation to the multi-level setting. The decentralization of the coordination mechanism via gossip protocol and the preservation of privacy are discussed in sections 6 and 7, respectively. Sections 8 and 9 cope with the evaluation of the market design and control algorithms in simulation. Section 10 deals with the implementation of market and control in the pilot project LIC. Finally, section 11 contains the conclusions.

2 Literature review on energy market formulation and distributed control

The mathematical problem of coordinating distributed agents in order to achieve a desired global effect is pretty general, and different approaches can be found in the literature for its application to the electrical grid. We will generally refer to this class of problem as distributed optimal power flow (D-OPF). This can be described as a networked optimization problem, in which each agent can only communicate with their neighbors, which are defined by a communication graph [2]. When dealing with D-OPF, grid constraints must be decomposed among the different agents, and we refer to them as coupling constraints. Various approaches exist to solve them in a distributed way [3, 4, 5]. In [6] a non-exhaustive review on D-OPF methods can be found. A non-comprehensive review on architectures for distributed

MPC can be found in [7]. In [8] several papers on distributed MPC, including deterministic and robust approaches, are cataloged based on process commonalities, control architecture and theoretical properties. When applied to power systems, most of the studies are focused on a single level decomposition [9, 10, 11], where agents can communicate with a coordinator. Fewer works are focused on completely decentralized protocols. Notably, in [12], a method to solve the D-OPF through the alternating direction method of multipliers (ADMM) on arbitrary graphs is presented, which is referred by the authors as proximal message passing. Methods found in the D-OPF literature for coordination of distributed agents are usually based on some decomposition techniques, such as ADMM, primal-dual Douglas-Rachford splitting [9], which is equivalent to ADMM (cfr. [13] §4.5), Dantzig-Wolfe decomposition [14, 15], proximal minimization [16].

Despite the abundant literature on distributed control, the ongoing paradigm shift from centralized to decentralized generation poses new challenges. When power generators belong to a single entity, which is usually the assumption in D-OPF, the redistribution of revenues among competing entities does not have to be taken into account in the problem decomposition by the independent system operator (ISO). This is not the case for demand side management (DSM) in general, in which prosumers' actions are motivated by their own utility. In this setting, proper regulatory mechanisms defining rules for remunerations and compensations need to be set up to allow the integration of DSM services in the wholesale market. Such rules need to address the problem of conflicting interests between market participants.

Additionally, a growing number of publications [17, 18, 19] highlights the possibility for some agents to alter the distributed optimization mechanism in their favour, to the detriment of the welfare of the other agents, by just lying about their predictions. This setting requires to treat the problem decomposition in a game theoretic way, shifting the focus from distributed control to market design.

Different authors have designed market rules using a game theoretic framework, describing the problem as a non-cooperative game. The underling concept in game theory and mechanism design is that, if the right amount and type of incentives are given, rational prosumers will decide to cooperate.

Most of the game theoretic works on en Vickrey-Clarke-Groves (VCG) mechanisms [20, 21, 22] . The VCG has many teoretical theoretical properties, among which being weakly budget-balanced and strategy-proof. The strategy-proofness is achieved by the mechanism paying each agent the sum of the value of all other agents. In this way, the social choice coincides with the interest of each agent, but at the same time it makes the mechanism extremely expensive for the aggregator, which should pay each agent the total amount of money that it would have paid in a trusted setting. This problem is alleviated using the Clarke pivot rule, with a resulting total value $c_{tot,i}$ for the i_{th} agent equal to:

$$c'_{tot,i}(x) = \max_{u \in \mathcal{U}_i} \sum_{i=1}^{N} c_{tot,i}(u) - \max_{u_{-i} \in \mathcal{U}_{-i}} \sum_{j \neq i} c_{tot,j}(u) \tag{1}$$

It is clear that since the second term does not directly nor *indirectly* depend on any decision of i, the only way to influence its total value is influencing the first term. Since maximizing the first term is equal to maximizing the social welfare, the mechanism implements social welfare maximization and is strategy-proof. Anyway, the requirement that the second term has to be independent from agent i implies that N optimization problems must be solved, each of which is performed without considering a given agent, which is a major drawback when the total number of agents is big. Since the computational cost of solving a single problem is expected to scale linearly with the number of users, the overall computational cost will scale as $\mathcal{O}(n^2)$. Moreover, when trying to distribute strategy-proof mechanisms as VCG, dominant strategy equilibrium can not be guaranteed anymore [23]. This effect is also known as the cost of decentralization. The strongest notion of equilibrium when decentralizing a mechanism is known as ex-post Nash equilibrium. Informally, this equilibrium guarantees that the suggested strategy $\sigma^*(\theta_i)$ is the best strategy when all the other agents follow it. Formally:

$$c_{tot,i}\left(g(\sigma^*(\theta_i), \sigma^*_{-i}(\theta_{-i}))\right) \geq c_{tot,i}\left(g(\sigma'(\theta_i), \sigma^*_{-i}(\theta_{-i}))\right) \tag{2}$$

Authoritative work on decentralized mechanism have been carried out in the contest of algorithmic game theory [23]. The ex-post implementations of VCG has been theoretically studied in [24], where design principles are presented. Notable work has been done in [25] trying to reuse parts of the redundant

computation of distributed VCGs mechanisms, in order to reduce the computational time. However, the works do not provide any guarantees on the computational time reduction, which is problem-specific. Moreover, including shared constraints in a VCG mechanism is not straightforward. A way of considering coupling constraints in games comes from the seminal work of [26] on n-person non cooperative games. In this work, conditions for the existence and uniqueness of the NE are given in terms of convexity of the considered game map. For instance, this work has been used in [27], applied to the problem of bidirectional energy trading between residential loads ad electric vehicles. The concept has been lately reconsidered, showing the existence of a new class of equilibrium refinement, called Variational NE (VNE) [28], which guarantees the same marginal prices for all the users at equilibrium. In [29] conditions are given for the uniqueness of a VNE in terms of monotonicity of the game map, in presence of coupling constraints. A similar description using monotone operators is provided in [30]. Since VNEs can incorporate coupling constraints and its distributed solution can be achieved using strategies which are well known in the field of distributed control, with well studied convergence properties and which scales better compared to VCGs, we chose to adopt them as a base for the formulation of the multilevel energy market.

3 Energy Market Formulation

The final market formulation we developed is agnostic with respect to the underlying business model, and can be adapted to any case fitting in the following setting:

- A group of end users can control some deferrable loads (e.g. electric boilers, heat pumps, etc...) or an electrochemical storage which can be used in order to shift production/consumption patterns
- An Independent System Operator (ISO) exploits a defined business model which uses end users' flexibility (e.g. committing to a predefined day ahead power profile as a service to a balance responsible party, providing regulation energy reserve, etc...) and redistribute part of the capital gain to flexible end users as a reward to their flexibility.

Under these assumptions, the problem can be mathematically formulated as a variant of the so called sharing problem [31]:

$$\begin{aligned} &\underset{u \in \mathcal{U}}{\text{argmin}}\, e(S_{\emptyset} u) + \sum_{i=1}^{N} c(u_i) \\ &s.t. : A_{\lambda} u \leq b \end{aligned} \tag{3}$$

where u_i are the actions associated with the agent i, $\mathcal{U} = \prod_{i=1}^{N} \mathcal{U}_i$ is the Cartesian product of the flexible users feasible sets, $e(u)$ is a system level objective which defines the business model, $c(u_i)$ are the costs of each flexible user in the business as usual case, and $u = [u_1^T, ..u_N^T] = [u_i]_{i=1}^N$ is the vector of the concatenated actions of all the flexible users. Here A_{λ} is a constraint matrix, taking into account the linear influence of the end users' powers on the problem constraints, and b encodes box constraints limits on the power and voltages at specified grid's nodes.

A notable business model is the one of energy communities, or self consumption community (SCC), for which the billing is done at the point of common coupling (PCC) with the main grid. This means that users inside the energy community always pay a lower energy bill if there is a heterogeneous mix of generation and production inside the energy community. In the case of energy communities, the function $e(u)$ is the surplus that the agent community has in paying the energy at the point of common coupling with the electrical grid:

$$e(x) = c\left(\sum_{i=1}^{N} u_i\right) - \sum_{i=1}^{N} c(u_i) \tag{4}$$

where $u_i \in \mathbb{R}^{\mathrm{T}}$ is the vector of total power of the ith agent, $c(\cdot)$ is the energy cost function defined as:

$$c(p_t) = \begin{cases} p_{b,t}p_t, & \text{if} \quad p_t \geq 0 \\ p_{s,t}p_t, & \text{otherwise} \end{cases} \tag{5}$$

where $p_{b,t}$ and $p_{s,t}$ are the buying and selling tariffs, respectively, at time t and p is the power at the households' electric main. In order to induce agents to follow the proposed mechanism, we must ensure that the energy tariff they pay participating in the market is always lower than the one they pay in the base case. This is always true when we are not taking into account grid constraints, since $e(u)$ as defined in (4) is always negative, when $p_{b,t} \geq p_{s,t}$, as usual in energy tariffs. However, if the agents are located in a grid with big voltage oscillations, the Lagrangian dual variables (which we can interpret as punishment prices) could be such that the cost paid by the agents is higher than $\alpha_i e(u)$.

In paper [32] we propose a single level coordination mechanism for problem (3), and addressed the following problems:

- Existence and uniqueness of a generalized-VNE for this class of sharing problems. We prove that decomposing (3) using different repartition weights for the surplus, induces a game with unique generalized variational equilibrium, which can be reached jointly minimizing the utility function of the agents, given by:

$$c_{tot,i} = c(u_i) + \alpha_i e(u) + \lambda_i^T u_i \tag{6}$$

$$= \alpha_i c\left(\sum_{i=1}^{N} u_i\right) + (1-\alpha_i)c(u_i) + \lambda_i^T u_i \tag{7}$$

  where α is a repartition coefficient for prosumer i, $\lambda_i \in \mathbb{R}^{2T}$ is a vector of Lagrangian multipliers associated with the i_{th} agent.

- Enforcement of individual rationality. Lagrangian multipliers in equation 6 are not bounded a-priori; this could lead to a situation in which the prosumers pay more than in the base case. To ensure IR, we encode it in the optimization scheme. At each iteration, for each time step in the horizon, we increment the Lagrangians only if the following condition holds:

$$\alpha_i e(u_t^k) + \lambda_{i,t}^T u_{i,t} \leq 0 \qquad \forall i \in N, \quad \forall t \in T \tag{8}$$

  where a negative value means that the prosumer is gaining a reward. This could obviously result in the impossibility to satisfy the coupling constraints. We can give the following straightforward economic interpretation to this mechanism: each agent would opt-out from the game as soon as the energy tariffs become unfavorable with respect to the base case. Condition (8) prevent this from happening. In the presence of bad power quality, the DSO could provide favorable energy tariffs to prosumers participating in the mechanism, ensuring that condition (8) is met with high probability.

- Comparison with the preconditioned forward backward (pFB) algorithm. We compared the proposed algorithm with the pFB. [30] states that all the projected-gradient algorithm class, when applied to aggregative games, can be interpreted as pFB. In figure 1, an example of solution is shown in terms of a group of prosumer's battery state of charge, aggregated power profiles and voltage at the PCC. Results for the game associated to 3 are reported, using the proposed formulation and the pFB algorithm. In this case, grid constraints, represented by the dashed red lines, are respected. This could not always be the case, for two reasons. The first is that the prosumers' flexibility could be in any case not enough to keep power and voltages inside the constraints set. For instance, this could depend on the total number of agents which are willing to participate to the proposed energy sharing community, which could be low with respect to the total number of users. Secondly, as anticipated, the hard cap on the Lagrangian multipliers enforcing individual rationality could prevent prosumers from using all their batteries' flexibility for grid regulation. The comparison was done by means of iterations before convergence, as depicted in figure 2, and showed that

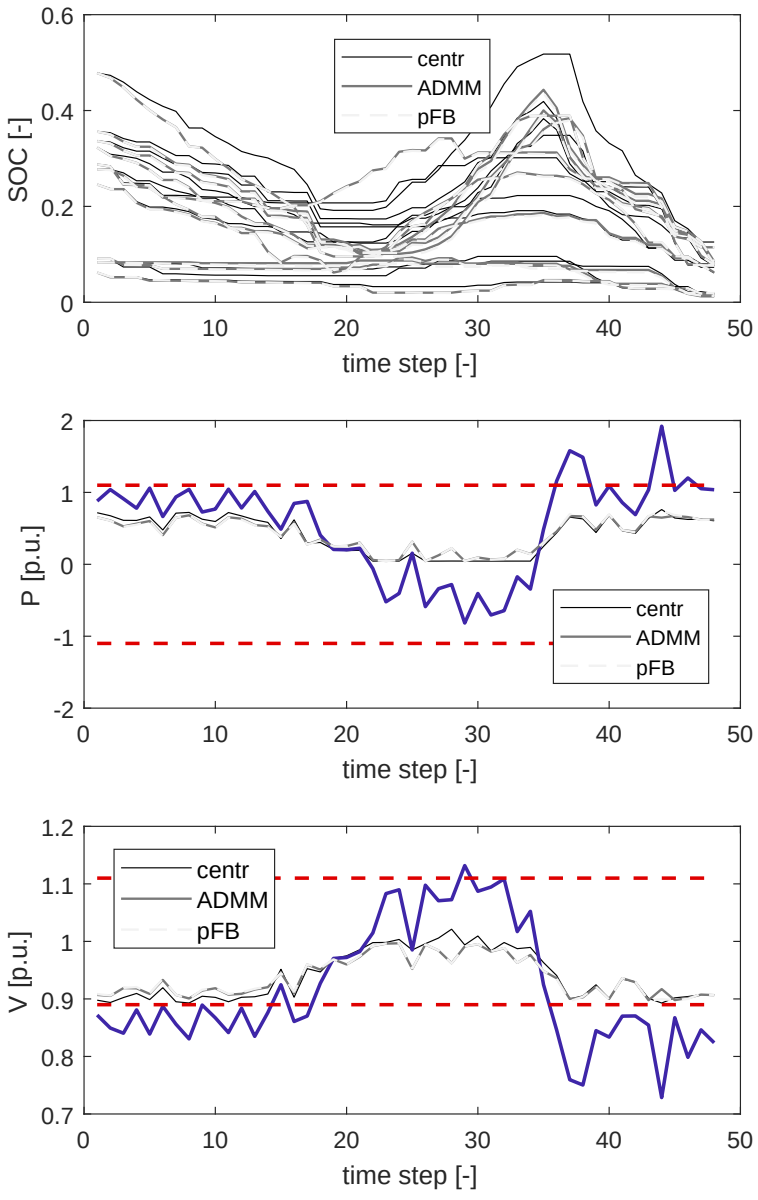

Figure 1: Time series example, N = 10. Blue: forecasted profiles. Red: constraints. Grays: solutions of the centralized and decentralized approaches. Top: state of charge for each battery. Middle: power profiles. Bottom: voltage profiles.

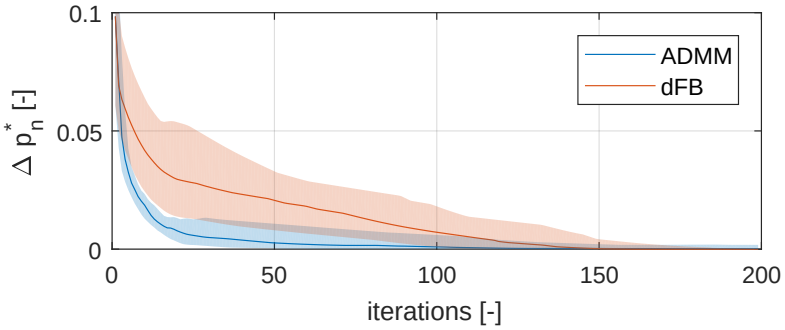

Figure 2: Normalized optimal value p^*. The thick lines denote the median, while the shaded areas are the 25% and 75% quantiles.

the ADMM based formulation requires on average fewer iterations to converge, which is desirable, since for each iteration step, the overhead and delays associated with a round of messages between the agents and the coordinator (or between the agents in the decentralized version) must be taken into account.

3.1 Redistribution coefficients and credible threats

Modeling the market as a non-cooperative game in theory prevents us from the possibility that some users could hack their smart devices in order to modify the decision algorithm controlling the deferrable loads. In fact, this would be an irrational action, since the cost specified in 6 induces a unique NE, where all the agents costs are minimized. However, some users could try to steer the NE in their favor. Since the final cost for the users depends on the actions of all the SCC members, could be possible for a user to influence the NE, misreporting the forecast of its own energy profile. The α_i coefficient is intended to promote a truthful report of the users' power forecasts. The α_i in (6) redistribute the SCC's economic surplus (4), so that the user is interested in maximizing it. In [33] we proposed the following formula for α_i:

$$\alpha_{i,t} = a_{i,t} \frac{\sum_{k=t-\tau}^{t} |x_{i,k}|}{\sum_{k=t-\tau}^{t} \sum_{i=1}^{N} |x_{i,k}|} \tag{9}$$

where τ is a characteristic period (e.g. one week). This is basically a normalized moving average of the consumption (or production) of the i_{th} agent. In other words, the discount experienced by the SCC member is proportional to the produced or consumed kWh. For the numerical experiments carried out in the following sections, we have adopted a different definition which promote fairness: the α_i coefficients are computed using moving averages of the change in the surplus due to the action of the i_{th} agent. In other words, called $e(x_{-i})$ the surplus computed without the action of the i_{th} agent, the α_i are defined as:

$$\alpha_{i,t} = a_{i,t} \frac{\sum_{k=t-\tau}^{t} \tilde{\alpha}_{i,k}}{\sum_{k=t-\tau}^{t} \sum_{i=1}^{N} \tilde{\alpha}_{i,k}} \tag{10}$$

$$\tilde{\alpha}_i = e(x_i) - e(x_{-i}) \tag{11}$$

This value is then reduced, multiplying it by a_i, which takes into account the accuracy of the next step ahead declared actions:

$$a_{i,t} = 1 - \frac{|p_{i,t} - p^*_{i,t}|}{|p_{i,t}|} \tag{12}$$

where $p^*_{i,t+1}$ is the optimal power profile declared by the i_{th} agent at time t for time t+1, after convergence of the distributed algorithm, while $p_{i,t+1}$ is the true value, realized at $t+1$. Note that this mechanism favor those users inside the SCC who do not have stochastic load nor production, that is, those members who

participate in the SCC only to perform price arbitrage, e.g. a user with only an electric battery. This is reasonable, since this kind of users provide a service to all the other member, actively maximize SCC's self consumption.

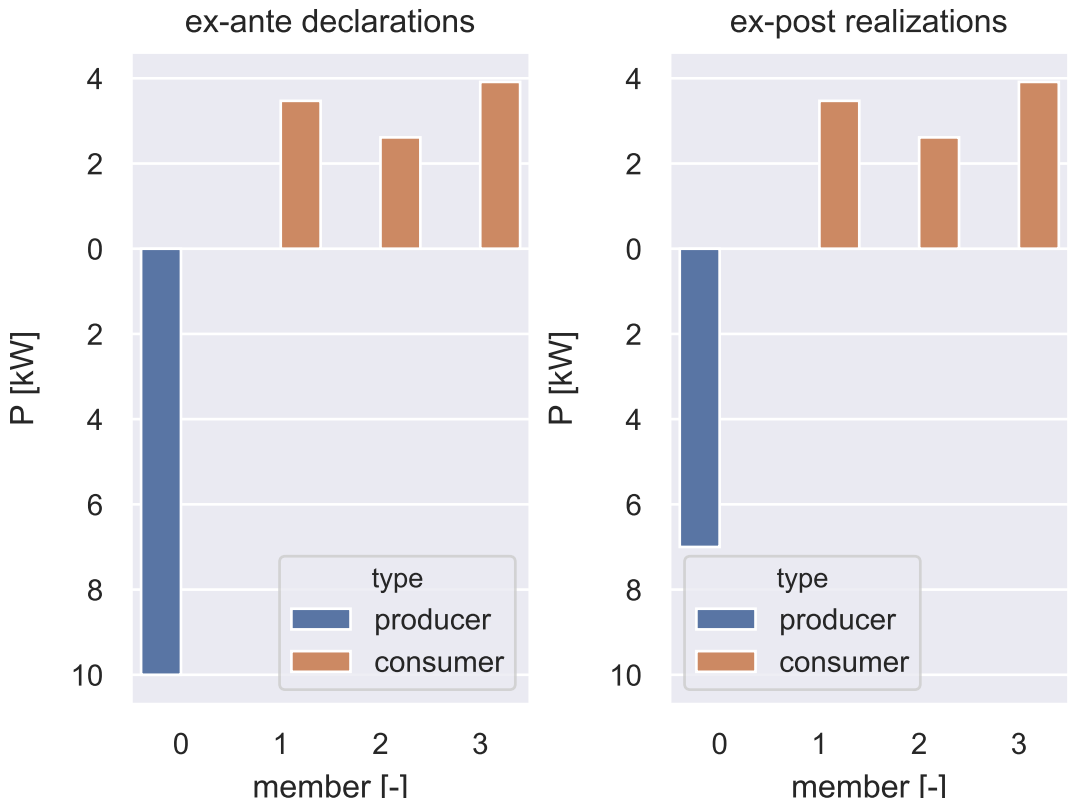

Figure 3: Example of EPNE manipulation. The only energy producer (blue) on purpose overshoot its own forecasted energy production during the coordination phase (left panel). The consumers increase their consumption plan in order to consume the cheap energy produce inside the SCC. Ex-post, when the decisions are implemented (right panel), while the consumers get stick to their declared plans, the producer fail to deliver all the promise energy. This behavior has no consequences for the producer, which sell all of his energy inside the SCC at a higher price, but it lower the SCC welfare.

Note also that $\alpha_{i,t}$ are always positive, but that they do not sum to 1, due to the introduction of $a_{i,t}$. This means that the mechanism generates a revenue stream for the SCC manager, which can be interpreted as an insurance against the discrepancy of the declared actions and their realization. In the game theoretic literature, this is known as a weakly budget balanced mechanism [34], that is, agents pay a greater then zero tax to a central entity running the mechanism. While this is not a good property in general, and in most of the game theoretic literature tries to achieve budget balanced mechanisms, to see why introducing $a_{i,t}$ in our case is beneficial, we should consider that:

- In order to guarantee IR, we cannot arbitrary punish the SCC members or owners as a function of their declared or predicted actions
- Not punishing them for wrong prediction will possibly lead agents to be too risk averse, since the cost of doing an optimistic (from the point of view of users' costs) prediction is externalized.
- Due to Myerson-Satterthwaite impossibility theorem, no mechanism is capable of achieving individual rationality, efficiency, and budget balance at the same time for general valuation functions

[35][36].

- The strongest notion of equilibrium attainable through a distributed mechanism is ex-post NE (EPNE) [37][23], which informally states that is guaranteed agents will minimize their costs if all the other agents follow a common protocol (intended to maximize the welfare). When this is not the case, a self-serving agent could try to steer the EPNE declaring false intentions for his next actions. This can be seen through a simple example. Consider a mechanism in which energy producers get paid more if their energy is consumed inside the SCC, and consumer pays less if they consume energy from the SCC. In this case, an agent owning only a PV could overshoot its forecast for the next step ahead. This will lead consumers with flexible loads to increase their consumption, making advantage of the cheap energy supplied by the PV owner. In the actuation phase, after the coordination step, all the energy produced by the PV owner will be almost surely be consumed inside the SCC, while the exceeding demand needs of some of the prosumers will be covered by (the more expansive) energy drown by the main grid. In this way the PV owner effectively steered the mechanism in his favor, externalizing his costs due to the uncertainty of its own predictions to the detriment of all the other consumers. This example is illustrated in Fig. 3.

4 Control Algorithms

In the following the algorithms used to control the distributed batteries and electric boilers are described. For sake of simplicity we describe the algorithms in the case in which the objective function is the economic cost in the business as usual. This formulation is also used in the implicit coordination case, where the prices are dynamically changed and each agent uses its own device to reduce its overall costs, without any kind of communication. Batteries and boilers are controlled through a model predictive control (MPC) approach: at each timestep of the simulation, the controller solves an optimization problem using consumption and production forecasts for the next day-ahead. Once the optimal solution has been found, the algorithms actuate only the first control action, and the procedure is repeated. We later extend it to the sharing problem with coupling constraints case presented in [38].

4.1 Battery control algorithm

The battery controller is supposed to be interfaced with the battery energy management system, returning an estimation of the battery's state of charge and injected and withdrawn power, into and from the battery. In this setting, the battery can be considered as a one state fully observed system and applying the MPC is straightforward. The formulation of the battery control algorithm for the implicit coordination is based on the work published in [33], and has been further improved to decrease the overall computational time, exploiting a new formulation for enforcing mutual exclusivity in charging and discharging operations. We report it in the following. Called $u = [p_{ch}^T, p_{ds}^T]^T \in \mathbb{R}^{2T}$ the vector of concatenated decision variables for the control horizon T, where p_{ch} and p_{ds} are the battery charging and discharging power, respectively, $\tilde{u} = [p_{ch}, p_{ds}] \in \mathbb{R}^{T \times 2}$ being the same vector reshaped in a 2 columns matrix, $\hat{p} \in \mathbb{R}^T$ being the forecasted power at household's main for the next contro horizon, $y \in \mathbb{R}^T$, $s_{ch} \in \mathbb{R}^T$, $s_{ds} \in \mathbb{R}^T$

being three auxiliary variables, we seek to solve the following problem:

$$u^*, y^* = \underset{u,y}{\operatorname{argmin}} \sum_t^T y_t + \|s_{ch}\|^2 + \|s_{ds}\|^2 \tag{13}$$

$$x_{t+1} = Ax_t + B\tilde{u}^T \tag{14}$$

$$y \succcurlyeq p_b \left(\tilde{u}[1,-1]^T + \hat{p}\right) \tag{15}$$

$$y \succcurlyeq p_s \left(\tilde{u}[1,-1]^T + \hat{p}\right) \tag{16}$$

$$x \in [x_{\min}, x_{\max}] \quad u \in [u_{\min}, u_{\max}] \tag{17}$$

$$s_{ch}, s_{ds} \succcurlyeq 0 \tag{18}$$

$$s_{ch} \succcurlyeq -\hat{p} \qquad s_{ds} \succcurlyeq \hat{p} \tag{19}$$

$$u \preccurlyeq [s_{ch}, s_{ds}] \tag{20}$$

where $\succcurlyeq$ stands for $\succcurlyeq_{R_+}$, indicating element-wise inequalities, $p_b \in \mathbb{R}^T$ and $p_s \in \mathbb{R}^T$ are the business as usual buying and selling prices. We start analyzing the objective function (13) term-wise. The first summation in (13) represents the total cost of the agent in the business as usual case. For prosumers, the cost function can be either positive or negative, depending on the overall power at their household's main and can be expressed as in equation (5). The cost can be thought of as the maximum over two

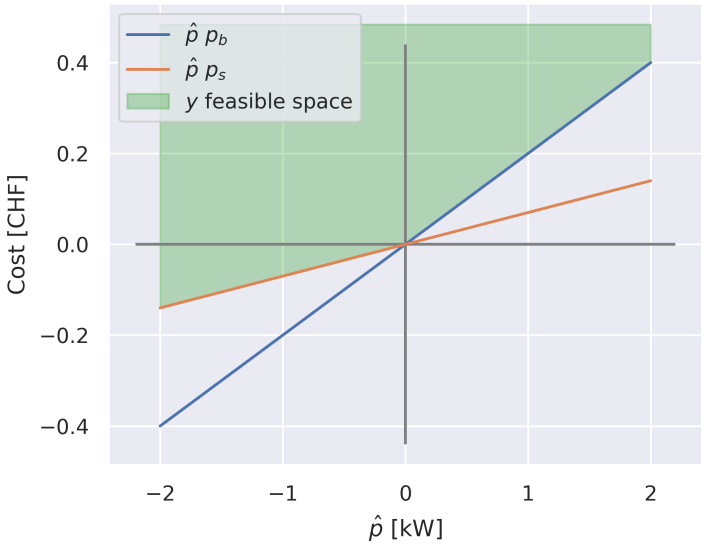

Figure 4: Visual explanation of the scope of the y variable. When linearly penalized, y is pushed to its feasible space's lower borders, collapsing on the cost function $c(p)$ in (5)

affine functions (the first and second line of equation (5), respectively). Equations (15),(16) constraint y to live in the epigraph of the maximum of these two affine functions. Minimizing y then guarantees that its value at the optimum, y^*, will lie on the epigraph's lower boundary (and will thus represents the prosumer's total costs), as shown in figure 4. Equation (14) describes the battery's dynamics. $A \in \mathbb{R}_+$ and $B \in \mathbb{R}_+^{1\times 2}$ are the discrete dynamics matrices obtained by the continuous one through exact discretization [39]:

$$\begin{aligned} A &= e^{A_c dt} \\ B &= A_c^{-1}\left(A_d - I\right) B_c \end{aligned} \tag{21}$$

where $A_c = \frac{1}{\eta_{sd}}$ and $B_c = [\eta_{ch}, \frac{1}{\eta_{ds}}]$, and η_{sd}, η_{ch} and η_{ds} are the caracteristic self-discharge constant, charge and discharge efficiencies, respectively. Since B_c defines an asymmetric behaviour in charging

and discharging (even with equal charging/discharging coefficients), solving the battery scheduling requires to use two different variables for the charging and discharging powers, p_{ch} and p_{ds}. When considering grid constraints, the battery can try to dissipate energy through round-trip efficiency to help respect negative grid constraints (when there is an excessive PV generation), so that in this case we need explicit binary complementary constraints for enforcing mutual exclusivity (the battery cannot charge and discharge at the same time). This can be obtained in three ways: explicitly modeling the bi-linear constraint $p_{ch}p_{dc} = 0$, introducing a binary variable and model it through big M formulation, or trying to restrict their feasible space. The first way will make the problem non-linear, while the second will turn it into a MIQP introducing a binary variable; as both options will increase the computational time, we introduced a new formulation exploiting the third way. Charging and discharging powers are effectively separated using the auxiliary variables s_{ch} and s_{ds}. The feasible space of s_{ds} is constrained to be the epigraph of the maximum between 0 and the forecasted power at the main. As shown in figure 5 for the case of s_{ds}, the equations (18) and (19) constraint these auxiliary variables to live in the positive half-plane and to be higher than the power profile at main (or its negative value for s_{ch}). When s_{ch} is quadratically punished, it will shrink on the lower boundary of the epigraph, (orange line in the second panel of figure 5). Its optimal value can then be used to define the feasible regions of the battery charging power, as done by equation (20). The same reasoning done in figure 5 for the discharging power can be applied to define the feasible regions for the battery's charging power; this will result in two disjoint feasible sets for the charging and discharging powers.

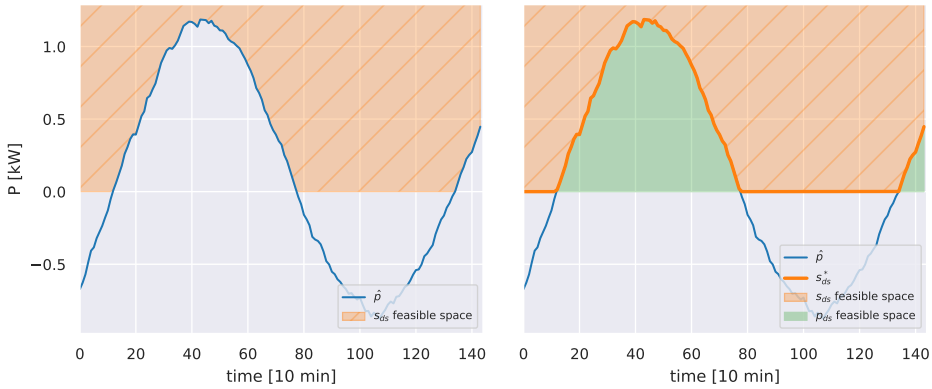

Figure 5: Visual explanation of the change in the feasible space for the discharging power.

4.2 Boiler algorithm

For the electric boilers, we cannot realistically assume them to be a fully observable systems. In fact, this assumption will require to have several sensors indicating their internal temperatures at different heights of the boilers. In a realistic setting, existing electric boilers has no more than two temperature sensors, used by their internal hysteresis controllers, and this information cannot typically be read from an external controller. Furthermore, consider the following simplistic one state model for the boiler's thermal dynamics:

$$cM\frac{\partial T}{\partial t} = c\Gamma\left(T_{i,t} - T_{o,t}\right) - U\left(T - T_{ext}\right) + P_{el,t} \tag{22}$$

Despite its simplicity, this model requires to know the incoming/outgoing water flux Γ, which means that a fluximeter must be installed. This is not possible but in pilot projects, since installation costs of these sensors will completely cancel out the economic benefit of an avoided grid refurbishment.

As such, we assume that we can only exploit the electric power measurements for controlling electric boilers. Furthermore, we expect to be able to only turn off the boiler through a relay, and not forcing it on (due to safety reasons, since we do not any feedback). Given these constraints, the electric boiler's nominal power and energy needs are estimated using historical data of their power consumption. Then, the algorithm decides when to force off the boiler such that the boiler can always satisfy its energy needs inside 3-hours slots. We based our algorithm on the work published in [40]. The algorithm is summarized in the following points:

- The nominal power of the boiler, P_{nom} is estimated from historical power data.
- The energy needs of the boiler are forecasted using a LightGBM[1] model taking as input past data of the boiler's power profile, as well as weather predictions for the next 24 hours. Furthermore, forcing the boiler off could result in an energy rebound effect. This can be corrected by passing to the forecaster also historical values of the control action as a categorical binary variable (since we want to forecast the energy needs of the uncontrolled boiler, this approximately counteracts our action on the system).
- The algorithm decides when to force off the boiler such that the boiler can always satisfy its energy needs inside 3-hours slots. For example, if a consumption of 2 kWh is forecasted between 18h-21h, and the estimated nominal power is of 4 kW, the boiler can be forced off at most 2h30min during this period.

Even if the boiler cannot actively be forced on, the internal control of the boiler, which is usually an hysteresis based on one or two temperature sensors, will automatically turn it on if its internal temperature is too low. The mathematical formulation is the following:

$$u^*, y^* = \underset{u,y}{\operatorname{argmin}} \sum_t^T y_t - \sum_t^T \min(\gamma, 0) \tag{23}$$

$$\text{s.t.} \quad y \succcurlyeq p_b \left(\hat{p}_b(1-u) + \hat{p}\right) \tag{24}$$

$$y \succcurlyeq p_s \left(\hat{p}_b(1-u) + \hat{p}\right) \tag{25}$$

$$S\left[(1-u)p_{nom} - \hat{p}_b\right] \succcurlyeq \gamma \tag{26}$$

$$\sum_{t=1}^{T-1} |\Delta u| \succcurlyeq n_{ch} \tag{27}$$

where y has the same role as in the battery optimization problem, representing the total costs for the prosumer, γ is a slack variable which relax the energy invariance constraint (26). Here S is a summation matrix which sum the energy in the pre-defined time slots (3 hours). Equation (27) further prevents the boiler for being turned on and off more than n_{ch} times in a control horizon.

4.3 Sharing problem

The presented problem formulations for the battery and the boiler, (13) and (23) respectively, minimize the end users' business as usual costs. These can be adapted to solve the decomposed sharing problem (3) simply modifying the part of the objective function representing the end users' costs and altering the feasible space for the charging and discharging powers. In particular, using the expression presented in section 3 for the total costs for the agent, the economic cost of the agent becomes:

$$c_{tot,i} = c(u_i) + \alpha_i e(u) \tag{28}$$

$$= \alpha_i c\left(\sum_{i=1}^N u_i\right) + (1-\alpha_i)c(u_i) \tag{29}$$

[1] https://lightgbm.readthedocs.io/

This cost function must be augmented with the Lagrangian multipliers coming from the decompostion of problem (3):

$$c_{tot,i} = \alpha_i c\left(\sum_{i=1}^{N} u_i\right) + (1-\alpha_i)c(u_i) + \lambda_i^T u \tag{30}$$

$\lambda_i = [\lambda_{i,ch}^T, \lambda_{i,ds}^T]^T \in \mathbb{R}^{2T}$ being the vector of Lagrangian multipliers associated to the i_{th} agent. Briefly speaking λ_i is a filtration of the overall λ induced by A_λ in (3). More details about this filtration process can be found in [33, 41]. In order to allow batteries to charge and discharge when system-level constraints are violated, we must further modify the feasible space of p_{ch} and p_{ds} in (13). In particular, equations (19) and (20) becomes:

$$s_{ch} \succcurlyeq -\hat{p} + M\lambda_{i,ch} \tag{31}$$
$$s_{ds} \succcurlyeq -\hat{p} + M\lambda_{i,ds} \tag{32}$$
$$\tag{33}$$

Here M is a big constant, that we set to $1e6$ in the simulations. An example of change in the feasible space of p_{ds} is shown in figure 6.

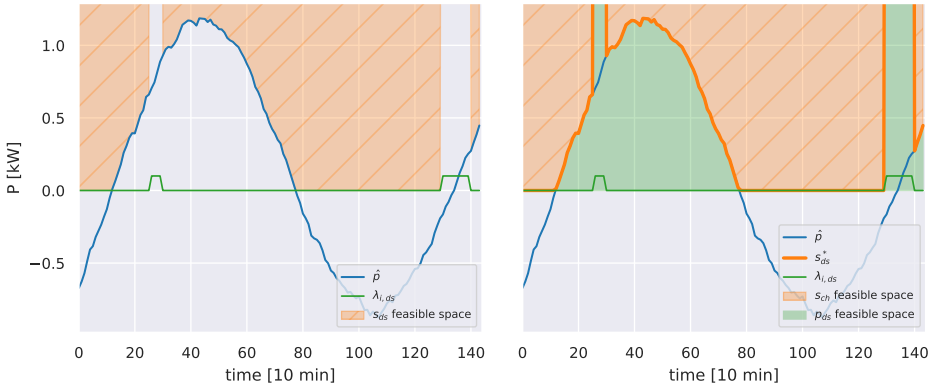

Figure 6: Visual explanation of the change in the feasible space for the discharging power, when s_{ds} is modified taking into account the Lagrangian associated to the battery's charging power $\lambda_{i,ds}$.

Finally, using the preconditioned forward-backward formulation, agents perform a gradient descent step in the direction of the negative gradient of the system level cost. This can be formulated as the minimization of the linearization of the system level cost around the previous state, plus a quadratic punishment on the action at the previous iteration; more details on this equivalence can be found in [33]. Replacing the agent cost with the auxiliary variable y as in (13) and (23), the final objective function (for the battery) then becomes:

$$\alpha_i \nabla c\left(\sum_{i=1}^{N} u_{i,pre}\right)^T u_i + (1-\alpha_i)\sum_{i=1}^{T} y + \lambda_i^T u + \rho_d \|u - u_{pre}\|^2 + \|s_{ch}\|^2 + \|s_{ds}\|^2 \tag{34}$$

where $u_{i,pre}$ are the agents actions at the previous iteration.

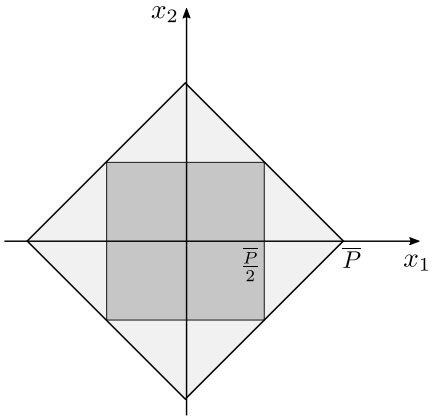

Figure 7: Feasible sets for the space of two prosumers' actions, for the constraint $|x_1 + x_2| < \overline{P}$ in the case of individual policies (dark gray) and in the case of communication (light gray). Communication enlarges the feasible set, thus potentially improving the solution.

5 Multilevel hierarchical control

The flexibility of prosumers can be exploited providing multiple services to the electrical grid, for example secondary control and energy management. Due to the electrical grid interconnections, in order to provide these services safely, an optimization method taking into account their effects on multiple voltage levels should be used. Moreover, coordinating groups of prosumers is always better than having prosumers to follow individual policies, in terms of aggregated effect. For example, let us consider the case in which we want to minimize a given objective function in terms of the aggregated power profile of a group of prosumers connected at the same LV transformer. Additionally, suppose that we require the power at the PCC to be symmetrically bounded, that is, $Sx \in \left[-\overline{P}, \overline{P}\right]$, where with $S = \mathbb{1}_N^T \otimes I_T$ we denoted the time summation matrix of all the prosumers' actions. In order to guarantee that this constraint is never violated using individual policies, we must replace it with the set of prosumers' constraints $u_i \in \left[-\overline{P}/N, \overline{P}/N\right]$. On the other hand, the feasible space of prosumers' actions would increase when allowing communication, thus leading to a better solution, as shown in figure 7. Formally:

Lemma 5.1. *The n-dimensional hypercube intersecting each axis in $\{-\overline{P}/N, \overline{P}/N\}$ is always included in the convex hull defined by the points $\left[e_i(-\overline{P}, \overline{P})\right]_{i=1}^N$, where e_i is the versor of the ith dimension.*

Proof. The maximum distance of the center from the n-dimensional hypercube defined in 5.1, equal to half of its diagonal, is $\overline{P}/\sqrt{N}$, which equal to the distance of the hyperplanes defined in 5.1 form the center, in each orthant. □

This means that constraining the aggregate instead of the single users will enlarge the feasible set $\mathcal{X}$, always resulting in a (not strictly) better solution.

In the literature, different methods have been proposed to coordinate groups of prosumers by means of an aggregator. Most of them refer to a single level hierarchy, as in [11, 15]. Both [10, 42] consider distributed MPC for a single level hierarchy, without grid constraints. In [43] a three level hierarchy has been proposed for the control of reactive power from PV inverters. The adopted sequential strategy does not explicitly take into account grid constraints and is solved through particle swarm optimization heuristic.

In paper [44], we present a multilevel hierarchical algorithm for the coordination of prosumers located in different voltage levels of the electrical grid. The hierarchical structure of the grid is described by means of a rooted tree. An example of rooted node structure is shown in figure 8. The coordination of prosumers is made possible by aggregators located at branching nodes of the tree. It is not necessary

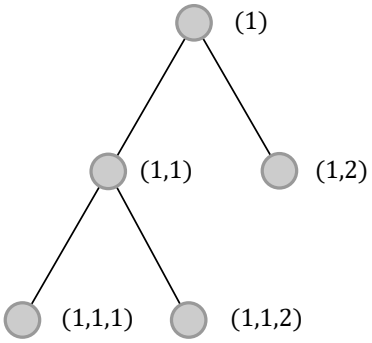

Figure 8: Example of rooted tree hierarchical structure

that the levels of the hierarchical structure coincide with the grid voltage levels. For instance, the first level of aggregation could be done at building level, while the second level could be placed at the LV cabinets. Here we briefly introduce the some notations needed to describe the hierarchical structure, which are useful to explain the coordination mechanism.

Definition 5.1 (Node sets)**.**

1. Descendants of node $A = (d_1, ..d_L)$. $\mathcal{D}(A) = \left\{A_j = (j_1, ..j_{L_j}) : L_j > L, (j_1, .., j_L) = A\right\}$
2. Leaf node $\mathcal{L}(\tau) = \left\{A \in \tau : \mathcal{D}(A) = \emptyset\right\}$
3. Nodes in level l. $\mathcal{N}_l(\tau) = \left\{A \in \tau : L = l\right\}$
4. Branching nodes. $\mathcal{B}(\tau) = \tau \setminus \mathcal{L}(\tau)$
5. Anchestors of node A. $\mathcal{A}(A) = \left\{A_j \in \tau : A \in \mathcal{D}(A_j)\right\}$

Formally, given N users with a controllable loads, located in a hierarchical structure, we solve the hierarchical counterpart of problem 3:

$$\begin{aligned} &\underset{u \in \mathcal{U}}{\text{argmin}}\, e(S_\emptyset u) + \sum_{i=1}^{N} c(u_i) \\ &s.t. : S_B u \leq v_B \quad \forall\, B \in \mathcal{B}(\tau) \end{aligned} \tag{35}$$

where $S_\emptyset = \mathbb{1}_N^T \otimes \mathbb{I}_T$ is a summation matrix summing all the users' actions u. Here S_B are similar summation matrices, each of which encodes a constraint in the set of branching points of the hierarchical structure, $\mathcal{B}(\tau)$. They are defined as:

$$S_B = \left[M_{B,A_j}\right], \quad M_{B,A_j} = \begin{cases} a_{B,j}\mathbb{I}_\mathrm{T}, & \text{if} \quad A_j \in \mathcal{D}(B) \\ 0_T, & \text{otherwise} \end{cases} \tag{36}$$

where $\mathbf{0}_T$ and $\mathbb{I}_\mathrm{T}$ are the zero and identity matrix of size T respectively, $a_{B,j}$ is a weight associated with the j_{th} descendant of branching node B. For instance, for constraining power at a give point of the grid, neglecting power loss, the coefficients $a_{B,j}$ are equal to 1 for the nodes directly connected to branch B and 0 for the other ones.

Despite having a simple formulation, problem (35) is very flexible, allowing having multiple objectives in different branches, or appliances providing possibly competing services to the grid, like primary regulation and dispatchable operations. In paper [44] we showed that the problem (35) can be solved using

a forward-backward communication protocol, with an iterative round of message passing spanning the hierarchical structure from the leaf node to the root node and vice versa. The mathematical formulation of the problem is the following:

$$u_i^{k+1} = \underset{u_i}{\text{argmin}}\, f_{c,i}(u_i) + \sum_{B \in \mathcal{A}(A_i)} \frac{1}{2\rho} \| \tag{37}$$

$$= (S_B x^k - y_B^k)/N_B - a_{B,i} u_i^k + u_i + \lambda_B^k \|_2^2 \tag{38}$$

$$y_\emptyset^{k+1} = \mathbf{\Pi}_{\mathcal{Y}_\emptyset}(\mathbf{prox}_{\rho e}(S_\emptyset x^{k+1} \lambda_\emptyset^k)) \tag{39}$$

$$y_B^{k+1} = \mathbf{\Pi}_{\mathcal{Y}_B}(S_B x^{k+1} + \lambda_B^k) \tag{40}$$

$$\lambda_B^{k+1} = \lambda_B^k + \frac{\rho}{N_B}(S_B x^{k+1} - y_B^{k+1}) \tag{41}$$

where N_B is the number of descendants of branch B, y are auxiliary variables used to decompose the problem and to treat the coumpling constraints, $\mathbf{\Pi}_{\mathcal{Y}_B}$ stands for the projection operator over the box constrain set defined by $\mathcal{Y}_B$ and $\mathbf{prox}_{\rho e}$ is the proximal operator for function e, with parameter ρ and $\mathcal{A}(A_i)$ are the ancestors of node A_i. Since the root node update involves the minimization of system-level objective function e, equation 39 projects its proximal minimization into the root node constraint set $\mathcal{Y}_\emptyset$, similarly to proximal gradient methods, as the forward-backward splitting algorithm [45]. As can be seen from equation (38), each agent needs some information from its ancestors in order to solve its optimization problem. This information can be provided by their respective parent nodes. This allows the algorithm to be solved in a forward-backward passage. In the forward passage each branch B sends its reference signal r_B and the one received by its parent to his children, which propagate it downwards through the hierarchy. At the same time, prosumers in leaf nodes solve their optimization problem as soon as they receive their overall reference signal r_i. In the backward passage agents send their solutions to their parents, which collect them and send the aggregated solution upward. Note that r_B contains only aggregated information from branch B, which ensures privacy among prosumers.

6 Assessment of gossip algorithms for decentralization

6.1 Introduction

The decomposed mechanisms presented in the previous sections can be decentralized, that is, run without the help of a central aggregator. In order to remove the need of an aggregator, there are two main options:

1. All the computations normally carried out by the aggregator, that is, equations (39), (40) and (41), must be deferred to a leader among the agents;
2. Since all the computations done by an aggregator node can be reduced to inexpensive algebraic calculations, all the agents could process them in parallel, in a similar fashion to what is done by the Ethereum Virtual Machine.

While the first option is less expensive computationally and from a communication complexity point of view, the high redundancy provided by the second one makes it particularly appealing for a distributed computation. The decentralization of the algorithm requires also to adopt some cryptographic or obfuscation protocol (like differential privacy) to prevent a disclosure of private information among the users. In fact, while we have regarded aggregators as trusted third parties, we consider the users as “semi-honest” agents, also known as the honest-but-curious setting [46]. This means that we assume that users are willing to cooperatively solve the problem, but they are also possibly interested in knowing the single agents private information such as energy consumption, production or their final bill. One way

to achieve a decentralized control is is by using the gossip protocol to exchange messages among the agents. The gossip protocols are well known in the telcom sector, as can be used to pass data on communication networks with an arbitrary topology. In this context, they can be used to estimate the signals that all the agents need to know in order to solve their own problem (38). In [47], the authors suggest to use the gossip protocol to asynchronously solve an unconstrained aggregative game using a first order method. For this class of games and solution scheme, the only thing needed by the agents to effectively coordinate is an esteem of the sum of their actions. This estimation can be done using gossip. In this chapter, starting from this baseline algorithm, we will examine its feasibility in terms of simulation speed and analyze and possible solutions to the security and privacy issues.

6.2 Baseline gossip method

The mathematical formalization of the problem we ideally want to solve is as follows: We have N agents, identified by an index i, each owning a private, real-valued data point, V_i. The objective is to develop a distributed protocol for the computation of the sum of the agents' private values without a significant privacy sacrifice, that is, preventing the possibility of an agent being in general able to reconstruct the private value of another agent.

As described in [47], the baseline gossip algorithm works as follows. The agents initially set $v_{0,i}$ to their private value, then repeatedly organize in pairs in a series of events k and compute the average of their $v_{k,i}$. Each agent sets a clock that ticks according to a Poisson process with parameter λ, where λ, in a naturally symmetric case, is agreed-upon and shared among all agents. When the clock of agent i tick, he randomly contacts a neighbor j. They then exchange their current estimates and compute a new intermediate estimate as the simple average of their current estimates. The intermediate estimates of all agents converge to a common value. Convergence is checked by the agents periodically: when the difference between the E values of the agents is under a certain commonly agreed threshold, the agents broadcast a convergence message. When all the agents are converged, the iterations stop. The value of the sum of the agents' private value is then the final, agreed upon value E multiplied by the number of agents.

6.2.1 Experimental convergence

Useful theoretical convergence properties are hard to obtain, since they involve the use of supermartingales. We will here show in Figure 9 the results of a simulation of the gossip protocol. The algorithm will be executed on randomly generated connected graphs whose nodes represent the agents. Each agent is assigned an uniformely random number from 0 to 1. Each gossip iteration consists of 2 messages (each of the two peers sending their $v_{k,i}$ to the other). The number of agents and the $1-\ell^0$ density metric of the adjacency matrix (where ℓ^0 is a simple measure of sparsity, the proportion of zero elements on the total) are varied. The mean and standard deviation σ of the number of messages exchanged before attaining convergence are shown in the figure. The data is obtained from 100 attempts for each combination and the relative tolerance for convergence is 10^{-3}. Notwithstanding the relatively loose convergence tolerance, the number of messages exchanged, even in this baseline version, proved to be high in relation to the simplicity of the task to obtain.

6.3 Security

As previously mentioned, in order to attain real-world applicability, security and anonymity features have to be taken into account. In this context, this proved challenging, requiring modifications that increased the communication overhead, further aggravating the overall speed problems.

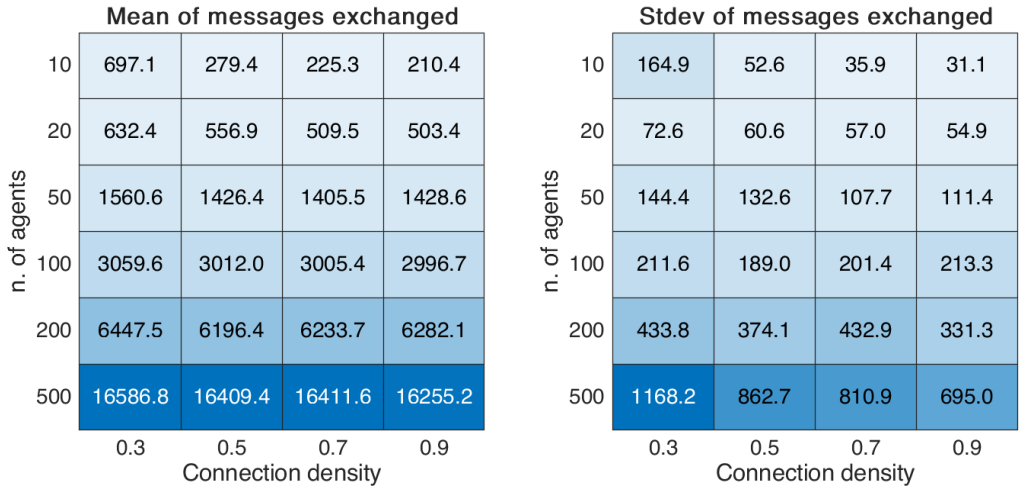

Figure 9: Experimental investigation of messages exchanged for convergence of the gossip algorithm

6.3.1 Asymmetric disclosure

This approach implies fully trusted agents. The only way of exchanging the private values naively is to do it in a sequential way, opening an opportunity for the first receiver to asymmetrically avoid disclosure.

In order to tackle this problem, we propose the following modification: The secret exchange phase between two partners is carried out using a simultaneous secret exchange protocol, also called gradual secret release protocol. Efficient protocols for this task exist, such as the classic ones proposed in [48]. This guarantees that two agents do not have to trust each other in sending their private values in sequential messages. Using this technique, though, first of all commands a computational effort that does not exist in the baseline, due to the need of computing the mathematical primitives needed by the obfuscation and decoding. Secondly, the protocols involve sending several messages per secret shared, implying a communication overhead that changes the order of magnitude of messages exchanged.

6.3.2 Clock spoofing

A second problem we notice on the baseline algorithm is that it involves revealing the agent's own initial private value, that is the sensitive one, to one other agent: the one with which is paired first. In particular, a malicious agent could pretend its clock ticked very rapidly after the start of the round and contact a neighbor J of his choice, obtaining with very high probability the initial value V_j of that agent. We can modify the protocol in the following way to address this problem: The secret exchange will happen over a secure channel. The first message the sending agent A will emit will be an anonymous proposal for initiating a secure key exchange protocol over an insecure network, such as the Diffie–Hellman key exchange protocol. The first receiving agent B will then randomly choose either to accept the key exchange and send back the following key exchange message, or to relay the message to a random neighbor C. This random neighbor will therefore see the message as if it originated from B, and has no way to tell the origin with certainty. C performs the same actions. When C, or any other following agent, decides to answer the key exchange protocol, the answers are relayed back by the same chain of neighbors. When the keys are exchanged, the gradual secret revealing happens inside this secure channel and it is therefore fully secured against eavesdropper. The following encrypted communication can happen both using the original chain of agents as a tunnel, or through distributed relayed broadcast of anonymous encrypted messages.

6.4 Implementation attempts

Implementation attempts for the techniques above proved slow to converge. Even the naive version of the gossip algorithm requires a significant amount of messages compared to aggregator-based methods, for example, simple leader selection followed by disclosure of the private values to the leader. This is due to the slow "diffusion" of the values throughout the network Since the naive gossip algorithm is, as we saw, sensible to attacks, the amount of overhead was deemed disproportionate to the benefits.

The two techniques described above allow for a high standard of security and anonymity, but imply significant further communication overhead - especially the gradual secret exchange, that requires a large amount of messages instead of just two containing the private values, with a severe impact on the wall clock time of convergence of the method as a whole. If one wants only to avoid just the second problem - maliciously establishing a frequent communication with a particular neighbor in a short time, in order to gain statistical insight on the profile of his secret values - the secret tunnel method can be used alone, independently from the simultaneous secret release, but the execution still requires further wall-clock time compared to the naive case.

It should be noted that the secret tunnel method allows, in addition to anonymization, connections between agents that might not be directly neighbors, promoting a longer reach of diffusion of the secrets in the original method. By taking into account only the "core" iterations, this promotes faster convergence of the pure gossip protocol, but this advantage does not nearly offset the communication overhead.

7 Privacy-preserving decentralized control mechanisms

As seen in the previous chapter, despite the feasibility of the approach using a gossip protocol, decentralizing the algorithm using it results in higher convergence times compared to a distributed approach (with an aggregator). Theoretical convergence properties are hard to obtain and, in any case, the experimental number of messages exchanged was high. Secondly, in order to preserve privacy, this method must be coupled with some cryptographic or differential privacy protocol, further worsening the overhead. For these reasons, instead of using a gossip communication scheme, we opted for a second way. Since we don't have any real topology restriction for the communication of the messages between agents (as in networked optimization), we can adopt an homomorphic encryption scheme for the cryptographic summation of private data. One option is to use the secure sum protocol [49], which is based on threshold homomorphic encryption [50] and random shares. One downside of this protocol is that agents can collude in order to alter the actual sum [51], however the protocol has been extended to work for an honest majority [52], further loosing the assumption of honest-but-curious agents. The secure sum protocol can be informally described as follows: given a group of N agents with private values $[v_1, ...v_N]$, a master agent, m generates a random number, r, drawing from a uniform distribution in the range $[0, s]$, where s is the upper bound of the expected sum of all the agent values. Agent m then passes $V_m = (r + v_m) \bmod s$ to the next agent in the group. Since r is randomly chosen from a uniform distribution, also V_m will be uniformly randomly distributed on $[0, s]$. Then, each agent pass the message summing its own value, thresholding it with s. In the last passage, the i_{th} agent passes his message back to agent m, which can then subtract r and retrieve the correct sum Σ.

Even in its simpler version, the secure sum protocol requests N rounds of sequential (ring) communication, since the sum can be decrypted only by the agent instantiating the message passing. Assuming all the N participants use the same ordered ring, and with a smart message gathering this can be lowered to $2N$. However, taking into account its extension to work with an honest majority requires non trivial message coordination and increases the number of messages to be sent sequentially. A simpler communication protocol with a stronger guarantee on privacy can be obtained adopting the voting method described in [53], lately implemented in an ethereum-based voting protocol [54] . The authors propose a voting system based on threshold homomorphic encryption which the strongest possible guarantee against collusion: the protocol guarantees privacy up to N-1 colluding agents [55]. Moreover, the protocol uses broadcast communication and has only two rounds, which is the best round efficiency for

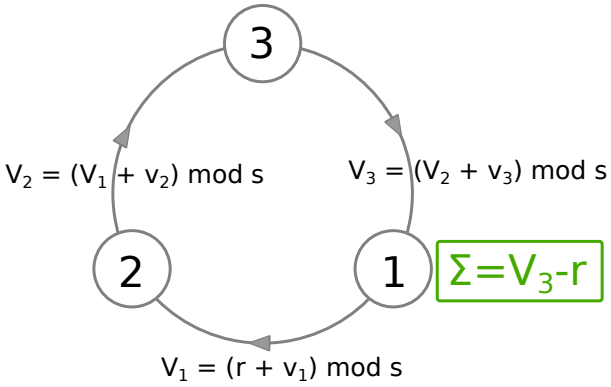

Figure 10: Schematics of the secure sum protocol. Agent 1 draws the random number r and at the end of the protocol he can retrieve the total sum Σ. No agents acquires any knowledge about other agents' private values during the protocol.

multi-party secure computation protocols [56]. Even if the presented method is thought to be used to sum boolean values (yes/no votes), we found it suitable to be applied to the sum of continuous values as well. The following list describes the steps constituting this methodology.

1. A list of eligible voters $(P_1, P_2, ..., P_n)$ is established
2. P_i selects a random value x_i as private voting key
3. P_i calculates and publishes its registration key

$$key_i = g^x \mod p \tag{42}$$

4. Each voter P_i calculates y_i:

$$y_i = \frac{\prod_{j=1}^{i-1} g^{x_j}}{\prod_{j=i+1}^{n} g^{x_j}} \tag{43}$$

   Equation (43) ensures the following equivalence:

$$\prod_{i=1}^{n} g^{x_j y_j} \mod p = 1 \tag{44}$$

5. P_i provides a verifiable hash of its vote v_i

$$Verification_i = hash(g^{x_i y_i}) \mod p \tag{45}$$

6. P_i calculates and publishes its encrypted vote PV_i:

$$PV_i = g^{x_i y_i} g^{v_i} \mod p \tag{46}$$

7. Then any voter can perform the following calculation:

$$\prod_{i=1}^{n} g^{x_i y_i} g^{v_i} \mod p \tag{47}$$

8. Using Equation (47) a voter can determine $g^{\sum v_i}$ and then, after an exhaustive research, obtain the final result $\sum v_i$.

Figure 11 shows the aforementioned election process in an implementation using an Ethereum blockchain to store the needed data. In the LIC community described in Section 10.1, this method has been developed and deployed, and its practical implementation is reported in details in Section 10.2.

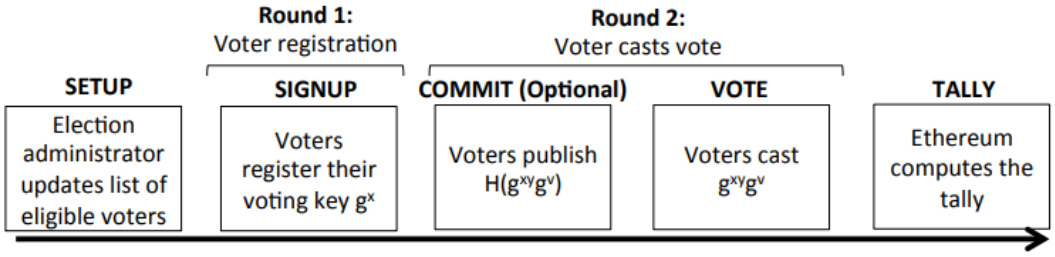

Figure 11: Election process, figure from [54]

8 Simulation environment

In order to allow for easy experimentation with the agent algorithms, a comprehensive simulation framework was prepared. The simulation software is called OPTISIM and provides the possibility to simulate algorithmic agents interacting with the grid. A scheme for the simulator is provided in Figure 12. The agents can be simulated in parallel, since communication in OPTISIM is based on message-exchanging through a central message broker. The main driving script runs the simulation and, at every step, launches the parallel simulation of the agents with a message containing information about the electrical step just simulated (such as voltage at the node of the agent, etc.), if needed for deciding the subsequent course of action. Single agents can comprise an arbitrary number of individual physical components such as building thermal envelopes, heating terminals, boilers, heat pumps, batteries, etc. Each of these elements has a corresponding physical model code, configured with the specific parameters of the agents' instance. Both the model and the agents can send, through the central message broker, records of their internal state to a time-series database (InfluxDB), allowing for storage and subsequent analysis of the simulation results. Both the electrical quantities and the agents' internal evolving state are, therefore, recorded. Furthermore, the same time-series database stores useful information available to the agents, such as meteorological data, temperature and irradiance.

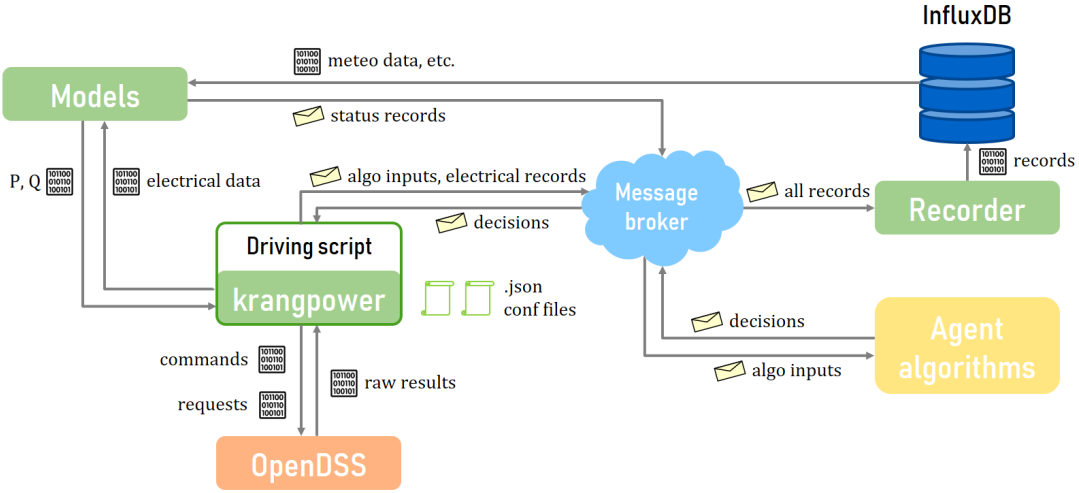

Figure 12: Conceptual scheme of the simulator

8.1 Buildings and appliances models

8.1.1 Heating system and control logic

In order to obtain a representative dataset for Switzerland, we used the STASCH6 standard [57] and its variants as a reference for the heating system and the control logic. The STASCH6 standard comprehends 3 main components: an heatpump (HP), a water tank used as an energy buffer, and a heating element delivering heat to the building. The HP control logic is based on two temperature sensors placed at different heights of the water tank, while the circulation pump connecting the tank with the building's heating element is controlled by an hysteresis on the temperature measure by a sensor placed inside the house.
We describe the control logic in a sequential way, following the heating components of the system. The first decision is taken by the building central controller, which decides its working mode, that is, if the building needs to be cooled or heated, based on a moving average of the historical data of the external temperature:

$$\begin{cases} wm_t & = -1 & \textbf{if} \quad T_{ma,t} > T_{max,ma} \\ wm_t & = 1 & \textbf{if} \quad T_{ma,t} < T_{min,ma} \\ wm_t & = 0 & otherwise \end{cases} \tag{48}$$

where the working mode wm_t is negative when the building requires to be cooled, positive when heating is required, and 0 when no actions are needed.$T_{max,ma}$ and $T_{min,ma}$ represent the maximum and minimum values of the external temperature's moving average, which is based on the past 7 days. The actual activation of the heating element is controlled by the hysteresis on the internal temperature of the building, T_z. If the working mode is positive, this is given by:

$$\begin{cases} s_{hy,t} & = 1 & \textbf{if} \quad (\, T_z < T_{min,hy} - \Delta T/2) \\ & & \textbf{or} \quad (T_z < T_{min,hy} + \Delta T/2 \quad \textbf{and} \quad s_{hy,t-1}) \\ s_{hy} & = 0 & otherwise \end{cases} \tag{49}$$

where $s_{hy,t}$ is the state of the hysteresis at time t, 1 meaning that the circulation pump of the heating element must be activated, and ΔT was chosen to be equal to $1°C$. For completeness, we report also the control logic when the building is in cooling mode:

$$\begin{cases} s_{hy,t} & = 1 & \textbf{if} \quad (\, T_z > T_{max,hy} + \Delta T/2) \\ & & \textbf{or} \quad (T_z > T_{max,hy} - \Delta T/2 \quad \textbf{and} \quad s_{hy,t-1}) \\ s_{hy} & = 0 & otherwise \end{cases} \tag{50}$$

The incoming water temperature in the heating element is then modulated linearly through a 3-way valve between a maximum and minimum value, based on the external temperature, both in the heating and cooling modes. When operative, the heating element requests hot or cold water to the water tank, which control logic is based on two temperature sensors located in two different layers. When the building is in heating mode, the control logic is a simple hysteresis based on the temperature of the sensor in the uppermost layer, which is identical to the one in (49). When in cooling mode, the control logic is the following:

$$\begin{cases} s_{hy,t} & = -1 & \textbf{if} \quad (\, T_{up} > T^c_{max} + \Delta T/2) \\ & & \textbf{or} \quad T_{low} > T^c_{max} + \Delta T/2 \\ s_{hy,t} & = 0 & \textbf{if} \quad (\, T_{low} < T^c_{min}) \ \textbf{or} \ (T_{up} < T^c_{max} - \Delta T/2) \\ s_{hy,t} & = s_{hy,t-1} & otherwise \end{cases} \tag{51}$$

where T_{up} and T_{low} are the temperature measured by the upper and lower sensors, respectively, and T^c_{min} and T^c_{max} are the minimum and maximum desired temperatures of the water in the tank while in cooling mode.
The value of $s_{hy,t}$ is then communicated to the HP. In the case in which the HP is also used for the domestic hot water (DHW), the DHW tank is always served with priority by the HP.

8.1.2 Building model

We modeled the building thermal dynamics with a simple one state RC equivalent model, as done in [58]. The main reason for this choice is that it is hard to generalize RC models with higher number of states, since no values can be found in the literature for the needed parameters. Estimating an RC model from data requires different measurements of temperatures, internal and solar gains, at a resolution of at least 10 minutes. This kind of datasets are extremely hard to find, and limited to only a few, often undwelled, cases. These equivalent RC circuit parameters could, in theory, be estimated starting from first principles, however recently proposed studies show that this can give worse results then estimating a model from data [59]. The second reason is that, while a higher order model leads in general to smaller one step ahead residuals compared to a lower order model, the loss of accuracy passing from a one state model to an higher order one when considering a longer period of simulation is much lower [60]. Last, when considering RC models for buildings with a number of states higher than 3, the chances of overfitting are high, and additional measurements such as the heat fluxes between thermal zones are required to guarantee observability. Alternatively, pseudo-random binary sequences can be applied to the heating systems in order to excite the system in a wide range of frequencies [61], while being uncorrelated with other exogenous inputs. This technique induces high changes in internal temperature of the building and cannot clearly be applied to occupied buildings.
The final model is the following;

$$C\frac{\partial T_z}{\partial t} = \frac{T_{ext} - T_z}{R} + kQ_h + A_{eq}I_s \tag{52}$$

where T_{ext} is the the external temperature, R is the equivalent thermal resistance for the building, C is the thermal capacitance, k is a parameter weighting the estimated power coming from the heating system Q_h, I_s is the incoming solar radiation and A_{eq} is the estimated equivalent window area. In order to obtain representative simulations, R, C, k and A_{eq} were estimated from statistical data for Swiss households. For the simulation regarding the LIC pilot, R was directly estimated from data and C and A_{eq} where estimated from the buildings' equivalent area.

8.1.3 Floor heating

Considering a fixed and uniform temperature for the ground and the building internal temperature at each time-step and stationary conditions, we can retrieve the analytical expression of the temperature profile along the pipe, through the energy balance on an infinitesimal element of the pipe. This can be expressed as:

$$\frac{\partial cT_x}{\partial t} = \Phi_x - \Phi_{x+\partial x} + \dot{q}_{up} + \dot{q}_{down} \tag{53}$$

where c is the heat capacity in J/K, x is the distance from the pipe entrance, T_x is the temperature of the water inside the pipe at x, Φ are enthalpy flows at the entrance and exit of the considered infinitesimal volume, $\dot{q}_{up}$ and $\dot{q}_{down}$ are the heating powers from the building and from the ground. Expressing the latter through equivalent resistance taking into account convective and conductive effects, the balance in steady state can be rewritten as:

$$\frac{\dot{m}c_p}{\rho^*}\frac{\partial T_x}{\partial x} = \frac{R_{down}T_z + R_{up}T_g}{R_{down} + R_{up}} - T_x = T^a - T_x \tag{54}$$

where T^a is the asymptotic temperature and where:

$$R_{down} = \frac{1}{h_{in}w} + \frac{1}{h_{u,eq}w} + R_u \tag{55}$$

$$R_{up} = \frac{1}{h_{in}w} + R_g \tag{56}$$

$$\rho^* = \frac{R_{up} + R_{down}}{R_{up}R_{down}} \tag{57}$$

where w is the diameter of the tube, h_{in} is the internal coefficient of heat transfer, which can be retrieved using available empirical relation for fully developed flow with fixed temperature at the boundary conditions [62], $h_{u,eq}$ is the heat transfer coefficient between the floor and the building air including both the effect for natural convection and radiation. The values of $h_{u,eq}$ can be found in the literature [63],[64]. The value of the thermal resistances R_u and R_g, towards the floor and the ground, can be found in the literature as well. We can reformulate (54), making it adimensional through a change of variable:

$$\frac{\partial \Theta}{\partial \mathcal{X}} = -\Theta \tag{58}$$

from which solution we can retrieve the temperature profile of the water inside the pipe:

$$T_x = T^a + (T_0 - T^a)e^{\frac{-xp^*}{\dot{m}c_p}} \tag{59}$$

where T_0 is the temperature of the water at the pipe inlet. We can use (59) to retrieve the heating power flowing into the building, integrating $\dot{q}_{up}(x)$ along the pipe.

$$\dot{Q}_{up} = \int_0^L \dot{q}_{up}(x)\mathrm{d}x = \int_0^L \frac{T(x) - T_z}{R_{up}}\mathrm{d}x \tag{60}$$

where L is the length of the serpentine. Integrating, we obtain

$$\dot{Q}_{up} = \frac{(T^a - T_z)L - (T_L - T_0)\frac{\dot{m}c_p}{\rho^*}}{R_{up}} \tag{61}$$

where T_L is the temperature of the water at the outlet of the serpentine. Note that the equation (61) tends to $(T_L - T_0)\dot{m}c_p$ when R_{down} increase and R_{up} is kept fixed.
The nominal mass flow of the heating system and the length of the serpentine are found as the solution of the following optimization problem:

$$\underset{L,\dot{m}}{\text{argmin}} \left(\dot{Q}_{up}(L) - \dot{Q}_{nom}\right)^2 + 10^{-3}\left(\dot{m} - \dot{m}_{nom}\right)^2 \tag{62}$$

where $\dot{m}_{nom}$ is a reference mass flow, equal to $0.1\,[kg/s]$ and $\dot{Q}_{nom}$ is the power required to keep the building internal temperature constant under reference conditions (we used an external temperature of -4°C and a desired internal temperature of 20 °C):

$$\dot{Q}_{nom} = \frac{\Delta T_{ref}}{R} \tag{63}$$

where R is the resistance of an equivalent RC circuit describing the heating dynamics of the building.

8.1.4 Water tanks and boilers

The water tank connected with the floor heating, which is used as a buffer by the heat pump, and the boiler for the DHW, are modeled as a N-states fully-mixed stratified tanks. Despite not being able to model buoyancy driven effects such as heat plumes and transient de-stratification, this kind of models are suitable for 1D simulations and control [65].

The dynamic equation describing the evolution of the temperature of the tank's layers is the following:

$$C\frac{\partial T_i}{\partial t} = \dot{Q}^u_{buo,i} + \dot{Q}^d_{buo,i} + \dot{Q}_{h,i} + \dot{Q}_{loss,i} + \dot{Q}^u_{cond,i} + \dot{Q}^d_{cond,i} + c_p\dot{m}(T_{i-1} - T_i) \tag{64}$$

where T_i is the temperature of the i_{th} layer, Q^u_{buo},Q^d_{buo},Q^u_{cond},Q^u_{cond} are the thermal powers due to buoyancy and conduction, from the lower and upper layer, respectively. The last term represents the enthalpy flow due to mass exchange, while C is the thermal capacity of the layer, in $[J/K]$ and $Q_{h,i}$ is the thermal

power due to an electric resistance (for the boiler) or an heat exchange (for the heating system buffer). The expression for the above thermal power are the following:

$$\dot{Q}^{u}_{buo,i} = k \max(T_{i+1} - T_i, 0)N, \quad 0 \quad for \quad i = N \tag{65}$$

$$\dot{Q}^{d}_{buo,i} = k \max(T_{i-1} - T_i, 0)N, \quad 0 \quad for \quad i = 1 \tag{66}$$

$$\dot{Q}^{u}_{cond,i} = u_{amb}(T_{i+1} - T_i), \quad 0 \quad for \quad i = N \tag{67}$$

$$\dot{Q}^{d}_{cond,i} = u_{amb}(T_{i-1} - T_i), \quad 0 \quad for \quad i = 1 \tag{68}$$

$$\dot{Q}_{loss,i} = u_{amb}(T_{ext} - T_{i)} \tag{69}$$

$$\dot{Q}_{h,i} = \dot{Q}_{tot}/n_h \quad if \quad i \in \mathcal{I} \tag{70}$$

$$\tag{71}$$

where N is the number of layers, u_{amb} is the equivalent thermal loss coefficient with the ambient and $\mathcal{I}$ is the set of the n_h layers heated by the heat exchange (or electric resistance). The buoyancy model is the one proposed in the IDEAS library [66]. Detailed description of the parameters for the boiler model can be found in [40].

8.1.5 Heat pump model

The heat pump is modeled by means of interpolated tables, in which heating and electrical power are available as a function of the evaporator and the condenser temperatures. The tables were taken from the energy simulation software Polysun (Vela Solaris AG, Winterthur, Switzerland). When the heat pump produces heat for both the heating system and the domestic hot water, its control logic prioritizes the latter, meaning that the buffer is heated as long as the DHW tank temperature sensor reaches the upper bound of its hysteresis control.

8.1.6 PV model

Residential PV power plants were modeled using the Sandia National Laboratories PV Collaborative Toolbox [67], which is based on the 1985 Grover Hughes' Engineering Astronomy course at Sandia National Laboratories, using typical inverters and polycrystalline modules data.

8.2 Electrical simulation

The electrical power flow simulation guided by the main script is carried out with Krangpower, an internally developed python library based on OpenDSS. It allows accessing the different functionalities of OpenDSS within the OPTISIM framework by providing modern interfaces, such as structured information retrieval, dynamic querying, graphing, and other such functionality, in order to provide a systematized interface and avoid the continuous need for scripts custom-tailored to the particular simulation.

The configuration of the simulation can get quite complicated, since full information must be provided about the grid, the agents, the models they contain with all the possible parameters. The file type used is json, due to the inherent hierarchical structure, useful for specifying the agents, then their content, and then their parameter. Krangpower uses a specifically structured json file that contains information about the grid, essentially a list of elements (for example, transformers and lines) with their parameters and topological information about the buses they are connected to.

As explained in the scheme, the algorithms that govern the agents run separately from the other components, and they interface themselves with the rest of the system by, again, standardized messaging. The algorithms implement the ideas explained in the theoretical chapters. In order to allow experimentation with algorithms that involve inter-agent communication, an infrastructural package was created, called Gossipy, that allows agents to be grouped in "chats" where they can address each other and where they can receive messages coming from outside (for example, the aforementioned messages from the

driving script that trigger the start of the elaboration by the agents). In this framework, agents actions are programmed as "reactions" to the arrival of a certain type of message. One or more aggregators can be added as additional entities, not representing any real element in the grid, that can receive and send messages. The various steps and iterations in the procedures described in the previous chapters thus are represented as explicit message exchanges between the aggregator and the agents themselves, or among the agents in the P2P case.

9 Algorithm evaluation

9.1 Simulation configuration

The grid of the neighborhood that participates in the LIC project (described in 10.1) has been mapped and its components simulated using the tools described in 8. The loads and PV plants characteristics were reconstructed based on the data originally provided by the local DSO. In the pilot project the total installed PV power is 64kWp. With respect to the original configuration, the PV power was increased by an additional 33kWp, in order to simulate a likely future situation, justifying the use of local storage.

9.1.1 Control configurations

The control algorithms developed in the project were tested in two main configurations:

1. Distributed batteries. In this first configuration, the four households of the LIC project that possess a PV plant have been equipped with a 15kWh/7kW battery, for a total of 60kWh/28kW. We assumed full control and observability of the batteries. The model of the battery as well as the optimization problems which they solve are described in detail in 4.1.
2. Water heaters. In the second configuration, 12 existing water heaters for a total of 61kW power and 4450l water storage were controlled. In this case, we assumed that the boilers could only be switched off, but not forced on, which is the typical case with ripple control. We assumed that the boilers internal temperature was not observable. Therefore, in order not to undermine user comfort, the control algorithms have been designed to slightly overestimate the energy need of the boilers and to make sure that this energy need is satisfied at all-time, by allowing the boilers to turn on for a sufficiently long amount of time in the control horizon. Moreover, the algorithms guarantee at least a minimum switch-on time of 2h every 12h. The model of the water heaters, as well as the optimization problems which they solve, are described in detail in 4.2.

9.1.2 Grid constraints

One of the characteristics of the explicit coordination mechanisms described in section 3 is that it directly takes into account grid constraints, which weigh on the objective function of each agent, proportionally to their influence on them. This means that each agent is rewarded if it contributes positively to the respect of the constraint, while it is punished if, with its actions, he contributes to the violation of the constraint. In our simulations, we have introduced a constraint on the maximum power measured at the main coupling point of the energy community. This type of constraint is particularly interesting because it can serve two purposes. First, it can be based on an actual technical limitation, i.e. a limit on the maximum power of a transformer or the ampacity of the line serving the community. But it can also be used to limit the monthly peak of the community towards the DSO. For example in the case of LIC, the energy community, being a big consumer with a total yearly consumption above 100MWh, needs to pay a power tariff of 7.5 CHF/kW to the DSO, based on its monthly peak. Therefore by actively reducing the maximum peak of the community, the agents directly contribute to lowering the total costs for the community. Two limits for

the maximum positive and negative power at the coupling point were chosen based on an estimate of how much could ideally be steered using the available batteries if one had perfect forecasts. The limits have been chosen based on a baseline simulation of the energy community, without batteries. The limits have been selected to be as low as possible, given that the following two criteria are respected:

1. The maximum daily energy exceeding the positive and negative limits must be smaller than the total energy storable in the batteries. This assumption means that in the worst day, all the batteries should have been empty (respectively full) to fulfill negative (respectively positive) grid constraints.
2. The difference between the minimum limit and the quantile 0.01 of the power at the coupling point and between the quantile 0.99 of the power at the coupling point and the maximum limit must be smaller than the maximum charge and discharge power of the batteries, respectively.

The power limits selected using the above mentioned criteria for the simulation of the LIC community are summarised in Table 2 and shown in detail for the month of July in figure 13. We stress that these limits where chosen automatically based on yearly simulation in which no devices where controlled.

month	negative limit (kW)	positive limit (kW)
1	-4.82	55.38
2	-23.09	51.62
3	-34.23	45.18
4	-39.32	33.76
5	-43.87	22.38
6	-48.25	18.48
7	-43.67	19.02
8	-42.69	19.19
9	-41.17	20.27
10	-14.34	33.64
11	-8.34	45.50
12	-1.54	52.45

Table 2: Selected power limits at the coupling point of the community as a function of the month for the simulated LIC community.

9.1.3 Coordination mechanisms

In section 3 and 4.3, we have presented a market design with an associated distributed control approach that allows maximizing money savings by leveraging the fact of being grouped in an energy community. One of the advantages of the presented model is that it allows explicitly enforcing grid constraints. This control strategy requires active coordination between the various players involved, who communicate their predictions to each other and iteratively solve an optimization problem. Therefore, we will call this mechanism "explicit coordination".
An alternative coordination mechanism is possible without the need for communication between users but only based on a price signal. A default rule that allows calculating the price of electricity sold at any time is defined. Users can then react according to this price signal. We will call this mechanism "implicit coordination". The "implicit coordination mechanism" is in place in the pilot project LIC (section 10.1). Various pricing schemes can be adopted, in this case we opted for a simple and easy to explain one. We define the price formation mechanism using extremely simple and interpretable rules:

- The energy consumed from the external grid shall be paid for as if the consumer were not part of the community.
- The energy consumed from inside the community is paid for at a total price lower than the standard tariff of the energy supplier and DSO, with a discount proportional to the ratio of the total produced and consumed energy.

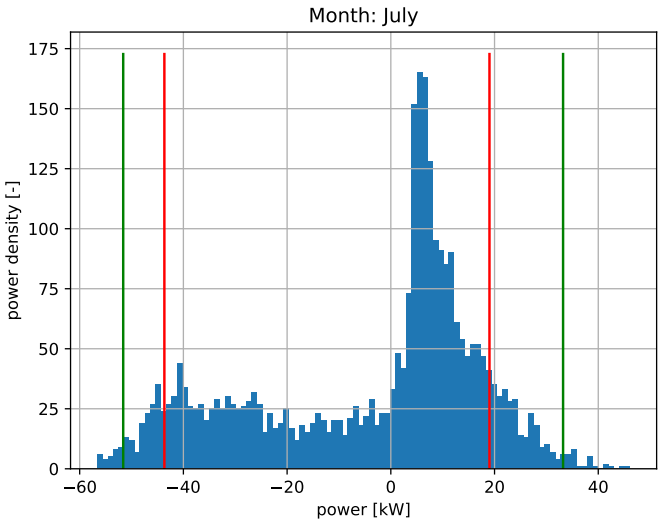

Figure 13: Power limits for the simulated month of July. Blue bars: histogram of the active power at the coupling point of the community. Green vertical lines: quantiles 0.01 and 0.99. Red bars: selected negative and positive power limits (-43.67kW and 19.02kW, respectively).

- The energy injected into the external grid shall be remunerated as if the consumer were not part of the community.
- The energy injected, which is consumed inside the community is remunerated at a price higher than the standard tariff of the energy supplier, with a discount proportional to the ratio of the total consumed and produced energy.

This simple set of rules push flexible users to increase the overall self-consumption. The energy price is calculated using an automated market making (AMM) mechanism, which follows these principles:

- The self-consumed energy is equally split among the community members proportionally to their consumption and production.
- The instantaneous buying and selling prices are dynamic, but for a given time slot they are the same for everyone.

We define the price functions of our AMM mechanism starting from the previously introduced rules, which can be expressed formally as:

$$\begin{aligned} p_b &= \left(E_c p_b^{BAU} - \min(E_c, E_p)(p_b^{BAU} - p_b^{P2P})\right)/E_c \\ p_s &= \left(E_p p_s^{BAU} - \min(E_c, E_p)(p_s^{P2P} - p_s^{BAU})\right)/E_p \end{aligned} \tag{72}$$

where p_b and p_s are the buying and selling prices generated by the AMM, E_c and E_p are the sum of the energy consumed and produced inside the energy community, while p_b^{BAU}, p_s^{BAU}, p_b^{P2P} and p_s^{P2P} are the buying and selling prices in the Business as Usual (BAU) case and inside the energy community. In such pricing configuration, peers clearly profit from the difference in price between BAU and P2P, but the community administrator also earns money, when energy is self-consumed inside the community. It is important to notice that P2P tariff is applied only to the energy produced by the members of the community, as a consequence it is also in the administrator interest to maximize self-consumption (no conflicting interests between peers and community admin). We can develop some intuition on how these prices reduces the variance of the aggregated power profile and maximize self- consumption plotting them as a function of the grid dependence index, defined as:

$$GDI = (E_p - E_c)/(E_p + E_c) \tag{73}$$

The GDI defines how much the community is dependent on the main grid, which provides an infinite reservoir of negative or positive energy. When the GDI is equal to 1, no one is consuming inside the community, while a GDI of -1 indicates that no one is producing. As shown in Fig. 14, the selling price for a net energy producer increases as the GDI moves from 1 to 0, then reaching a plateau. The same is seen for the buying price for a net consumer, decreasing while the GDI shifts from –1 to 0, and getting constant thereafter. This means that the community maximizes his welfare when the GDI is 0, that is when the buying price is minimized and the selling price is maximized for the agents. This means that the community maximizes its welfare when the self-consumption is maximized.

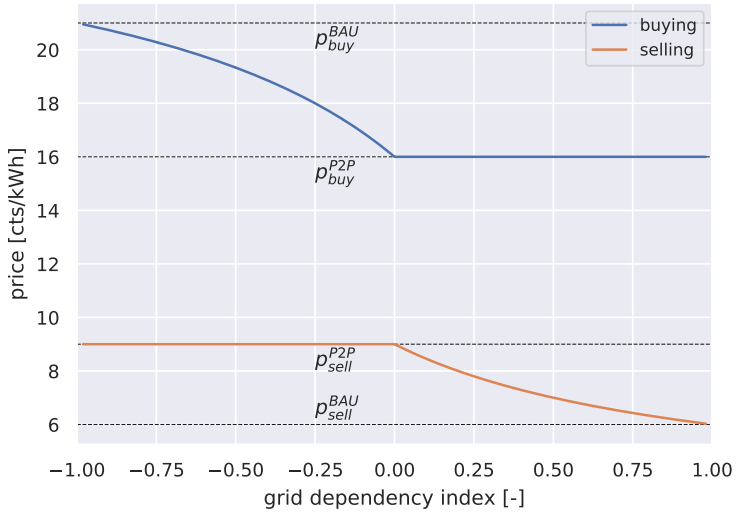

Figure 14: Buying and selling prices as a function of the GDI.

In the following sections the performance of the implicit and explicit coordination schemes are evaluated in simulation for both control configurations (batteries and water heaters), from a techno-economical point of view.

9.2 Simulation results

We will focus our analysis on the simulated month of July. Yearly simulations will be available in future reports, as the computational time required exceeded the time at our disposal. Simulations step was set to 5 minutes. The computational time required by a simulation step varies considerably depending on the coordination type, the forecasting method used and the control model. It can take from about 1.5 seconds in case of implicit coordination of batteries to about 30 seconds in the case of explicit coordination of boilers. Explicit coordination requires a much higher computational time since it is solved using an iterative approach. In this series of detailed simulations, the iterative distributed problem is solved using the pFB method, described in section 3. Parameter tuning for the pFB algorithm has proved to be essential and delicate. There is, in fact, a tradeoff between optimal conversion speed to the solution and stability. After parameter tuning, we limited the maximum number of iterations to 50, as we evaluated that, on average, the coordination converged sufficiently well after this amount of iterations.
The optimization problems are always solved using a model predictive control MPC approach over a horizon of 24 hours. In order to make the simulations as realistic as possible, we have assumed that the controller does not have perfect forecasts of any future variable. On the contrary, we have trained different types of forecasting algorithms and evaluated their performance, the results of which are presented in section 9.2.3.
The local DSO doesn't allow the injection of energy from batteries into the grid, as any injection into the

grid is remunerated as being produced by green energy sources. In the context of an energy community, there is no law regulating the injection of energy into the community's internal grid. Nevertheless, in the case of distributed batteries, there is a possibility that the energy injected into the internal grid will flow from the PCC into the main grid, especially if the batteries do not coordinate. To tackle this problem, we have tested two variants of control algorithms: A variant in which the batteries cannot inject energy into the internal grid, except in case they have to contribute to the respect of grid constraints (control model explained in detail in section 4.3), and a more permissive variant, in which batteries are allowed to inject energy into the internal grid, with the risk of leaking into the main grid from the PCC. The following scenarios were tested:

1. Baseline (no batteries, water heaters always enabled)
2. Batteries:
   (a) Batteries not coordinating: In this case the batteries are operated in local (at house level) self-consumption optimization mode. They have no incentives in coordinating, nor in shaving peaks. A delayed charging strategy is in place, to preserve battery life.
   (b) Batteries in implicit coordination mode without the possibility to inject energy into the grid: In this case the batteries will try to optimize the household energy costs, using the pricing scheme presented in section 9.1.3.
   (c) Batteries in implicit coordination mode, with the possibility to inject energy into the grid
   (d) Batteries in explicit coordination mode without grid constraints enforcement, without the possibility to inject energy into the grid. In this case the batteries coordinate with each other, and the surplus money derived from the fact of being members of in an energy community, with the method described in 4.3.
   (e) Batteries in explicit coordination mode with grid constraints enforcement, without the possibility to inject energy into the grid.
   (f) Batteries in explicit coordination mode with grid constraints enforcement, with the possibility to inject energy into the grid.
3. Water heaters:
   (a) Water heaters in implicit coordination mode.
   (b) Water heaters in explicit coordination mode with grid constraints enforcement.

9.2.1 Evaluation of impact on the grid

Batteries The impact of the control actions on the grid has been evaluated by quantifying the effect of the control actions on the power at the point of common coupling (PCC) of the community, the total losses in its internal grid and the voltage at its nodes. The total losses considered are the ohmic losses (dissipated for Joule's effect) over the internal grid lines, that in this case are equal to the difference between the power measured at the PCC and the sum of the powers of the meters in the grid. The time-series of the power at the PCC for the first three days of July is shown in figure 15. Figure 16 shows the probability density function (PDF) of the power at the PCC, and figure17 allows to further analyse the effects of batteries on the active power at PCC, by displaying statistics of the total controlled power as a function of the power at the PCC.

It is clear that the mere presence of batteries in the grid helps to minimize power and voltage fluctuations. It is equally clear, however, that without proper coordination of their actions, the benefit is minimal. In general, allowing batteries to inject energy into the internal grid even when grid constraints are not violated improves performance across all KPIs. However, one can notice in figure 15 and figure 17 that implicit coordination through price signals can cause significant power fluctuations, which with excessive penetration of storage systems could lead to instability. At certain times of the day, the energy injection of batteries into the internal grid is excessive and results in an energy leak into the main grid.

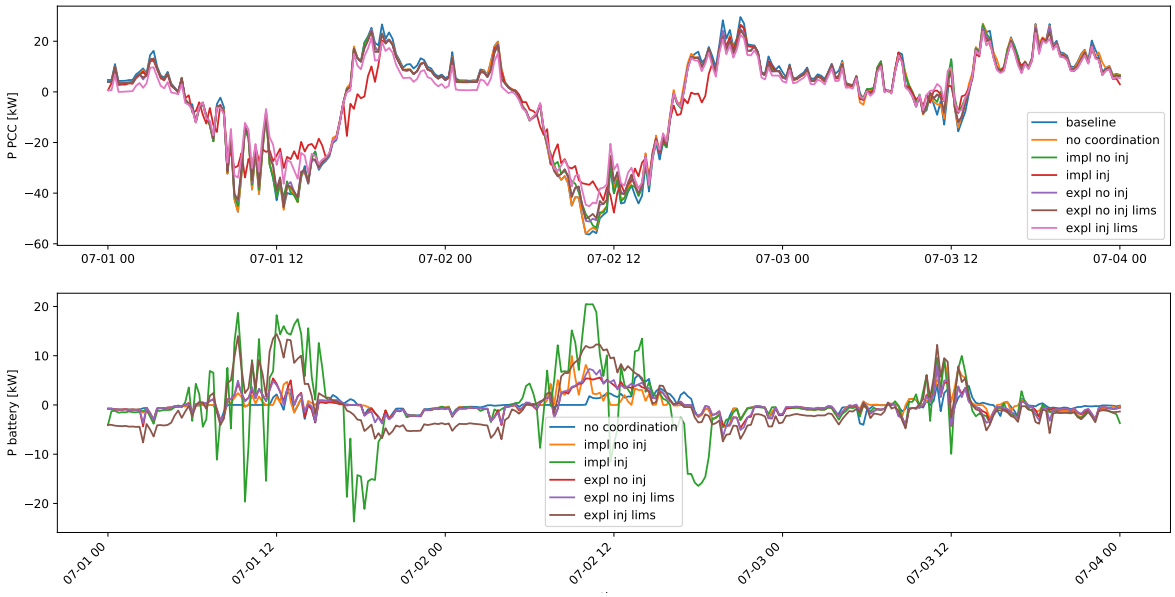

Figure 15: Time-series of the active power at the point of common coupling of the community and the total power of the distributed batteries for the first 3 days of July. Negative power corresponds to injection into the grid, positive power corresponds to consumption. Top: Active power at PCC. Bottom: active power of batteries. *baseline*: baseline simulation with no batteries, *no coordination*: batteries optimizing local self-consumption only, *impl no inj*: implicit coordination, no injection into the grid allowed, *impl inj*: implicit coordination, injection allowed, *expl no inj*: explicit coordination, no injection allowed, *expl no inj lims*: explicit coordination, no injection allowed, power limits at PCC enforced, *expl inj lims*: explicit coordination, injection allowed, power limits at PCC enforced.

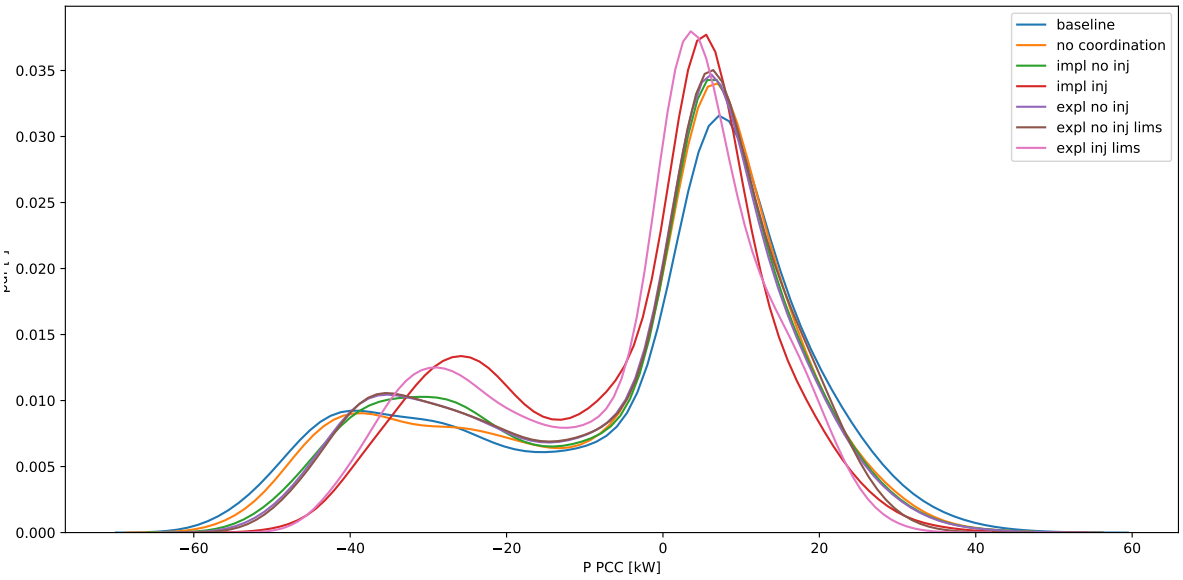

Figure 16: Probability density function (PDF) of the power at the PCC of the community.

On can also notice that the synchronism between battery actions is greater when the power at the PCC

is negative (i.e. when the community is injecting into the main grid). While loads of individual homes are only minimally correlated with each other, production from photovoltaic systems is almost perfectly correlated. This, combined with the fact that all agents use the same methods to forecast PV production, generates a higher synchronization of the batteries when they are charging.
Explicitly enforcing the grid constraints, as expected, helps to reduce the positive and negative peaks at the PCC. Nevertheless, it can be noted that, due to the fact that the control is based on imperfect forecasts, these limits are still violated. Future work will concentrate on the development of robust and stochastic distributed control strategies that will take into account the uncertainty in the forecasts.
Allowing the batteries to discharge into the grid at any time, and not only when constraints violations are expected, leads to better performance. This is because the actions of the batteries would have been beneficial even without enforcing grid constraints.

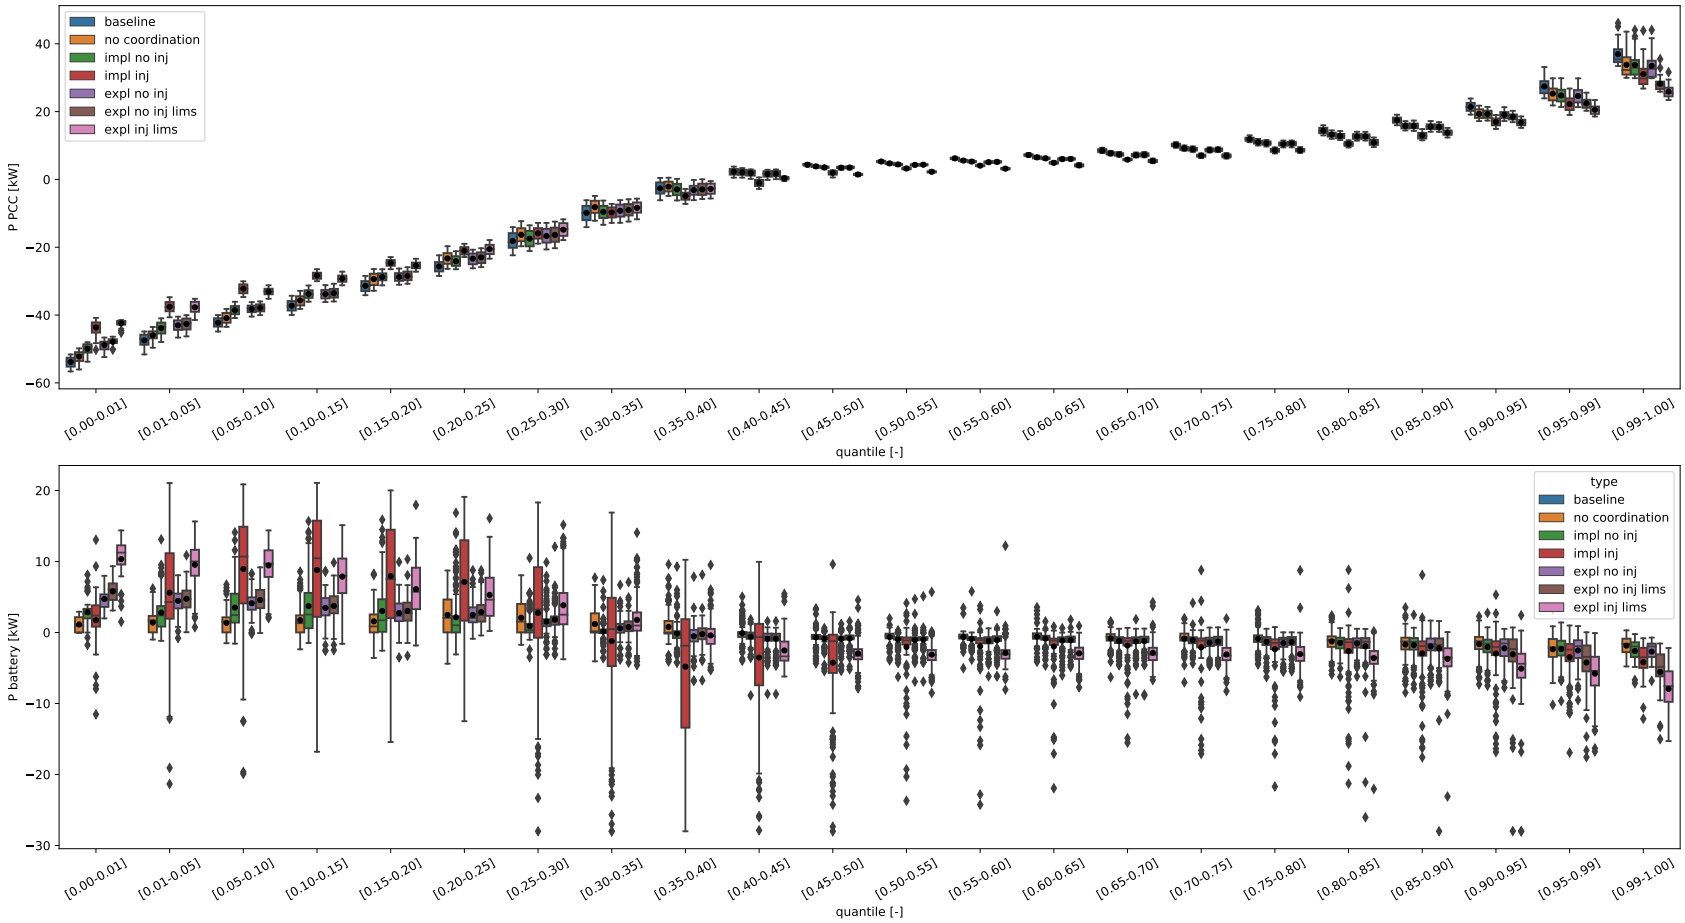

Figure 17: Effect of batteries control on the power at the point of common coupling (PCC) of the community. Top: boxplots of the power at the PCC per inter-quantile range. Bottom: boxplot of the total power of the batteries per inter-quantile range of the power at the PCC.

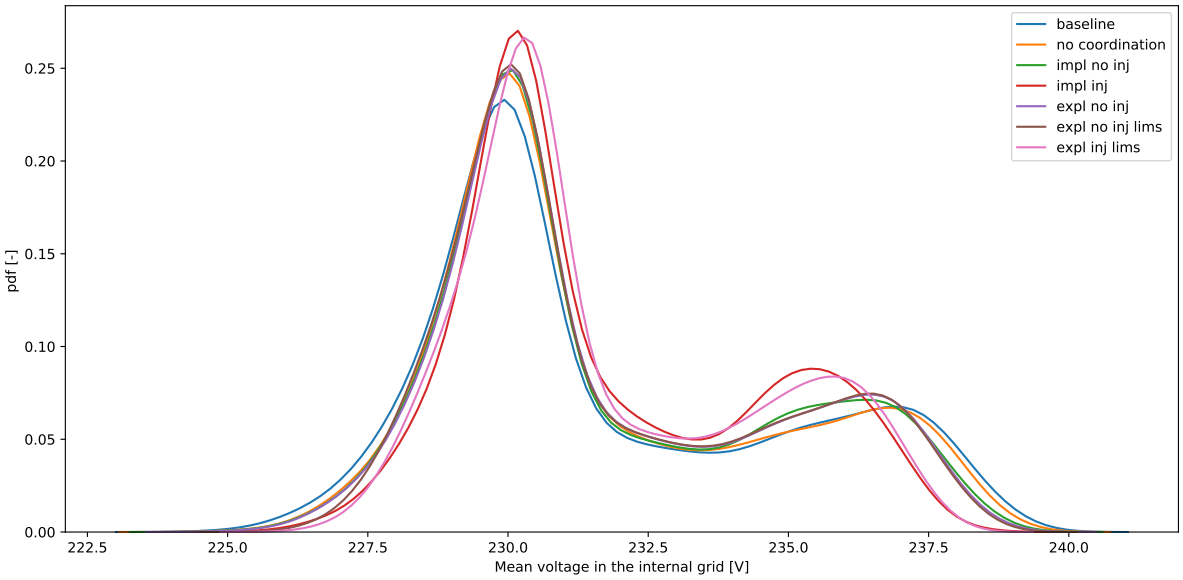

Figure 18: Probability density function (PDF) of the mean L-N voltage inside the community.

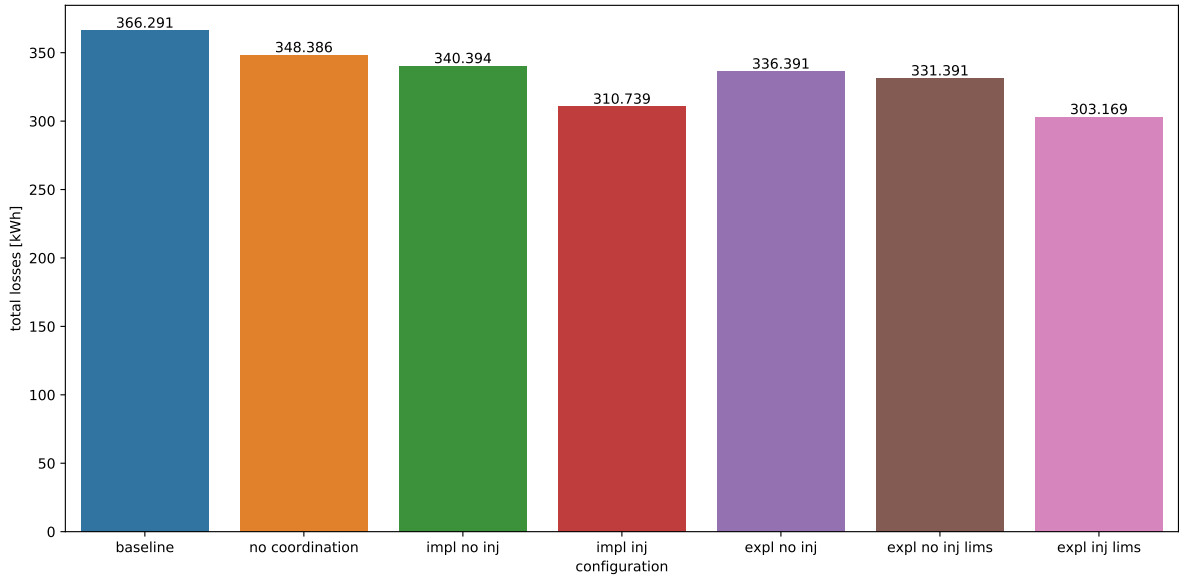

Figure 19: Total ohmic losses inside the community for the month of July.

In the case of explicit coordination, the strategy adopted has the advantage of converging towards a single Nash equilibrium, which guarantees a fair solution for all participants in the cooperative game. However, one might further increase the community surplus if one did not take into account that this might disadvantage some individual participants in the game. To assess the difference between the maximum attainable surplus, and the attainable surplus in the condition where a Nash equilibrium is reached, a simulation was run in which batteries were forced to ignore their internal costs and maximise the community surplus, only. This is achieved by setting $\alpha = 1$ in equation 7. The effect of this strategy on the grid can be seen in figures 20, 21, 22 and 23. One can notice the effect of maximizing the surplus of the community on the grid parameters is not significantly different than when single users' costs are taken into account.

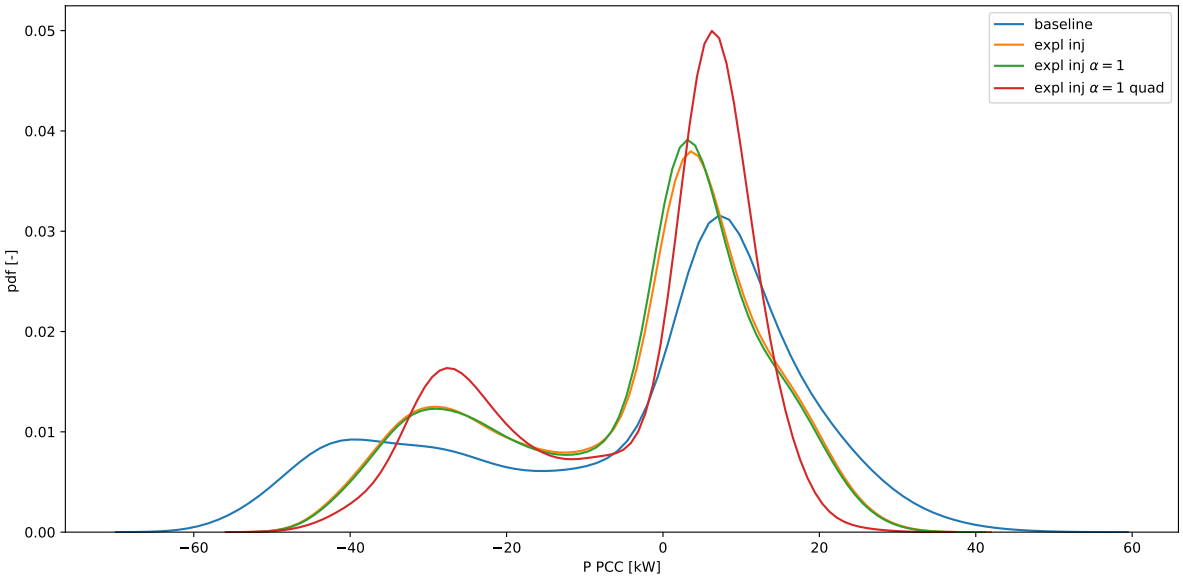

Figure 20: Probability density function (PDF) of the power at the PCC of the community. *baseline*: baseline simulation with no batteries, *expl inj*: explicit coordination, injection allowed, power limits at PCC enforced, *expl inj* $\alpha = 1$: explicit coordination, injection allowed, power limits at PCC enforced using Lagrangian dual variables, $\alpha = 1$, *expl inj* $\alpha = 1$ *quad*: explicit coordination, injection allowed, power limits at PCC enforced using quadratic punishment in the global objective function, $\alpha = 1$.

Another way to limit grid stress is to punish quadratically the excursions from a nominal value. This can be achieved by adding a quadratic punishment term to the objection function 3. Such strategy was applied to limit the power excursions at the coupling point. In this case the batteries try to flatten the profile as much as possible regardless of whether the grid constraints are violated or not. In this case, α was still set to 1, so that the batteries would only optimize against the community surplus. The effect on the grid of such strategy is shown in figures 19, 20, 21, 22 and 23. One can notice that, in this case the effect on the grid KPIs is more pronounced. The PDFs of power and mean voltage are narrower, and the losses, as expected, are lower. However, setting $\alpha = 1$ and adding a quadratic punishment in the objective function could theoretically worsen the financial balance of end-users. This is discussed in detail in section 9.2.2.

Water heaters Water heater control also leads to improved grid KPIs, although the effect is not as pronounced as with battery control. The effect of water heaters control on the power at the PCC of the community is shown in figure 24. It can be seen that the control of the boilers allows reducing the peaks of energy consumption significantly. At the same time, the effect on the reduction of power injection is minimal. This is mainly because the controllability of boilers is not as refined as in the case of batteries and presents inherent asymmetries. In fact, boilers can be forced off, but cannot be switched

on on command if their internal controller based on hysteresis does not consider it necessary to charge. This, combined with the fact that it is not possible to force the boilers off for too long without risking compromising the comfort of the end-users, means that often when boilers are left free to turn on, they don't actually do so. This leads to an imperfect reduction of the negative peaks, which would require the boilers to turn on to be smoothed out. The effect the difference between implicit and explicit coordination is less pronounced. Raising the aggressiveness of the boiler control algorithms by lowering the turn-on time should help increase control performance against the grid KPIs. However, there is a risk of violating user comfort. In this case, the comfort settings are very conservative, as the same settings are used in the LIC pilot project. Additional simulations are planned with more aggressive settings.

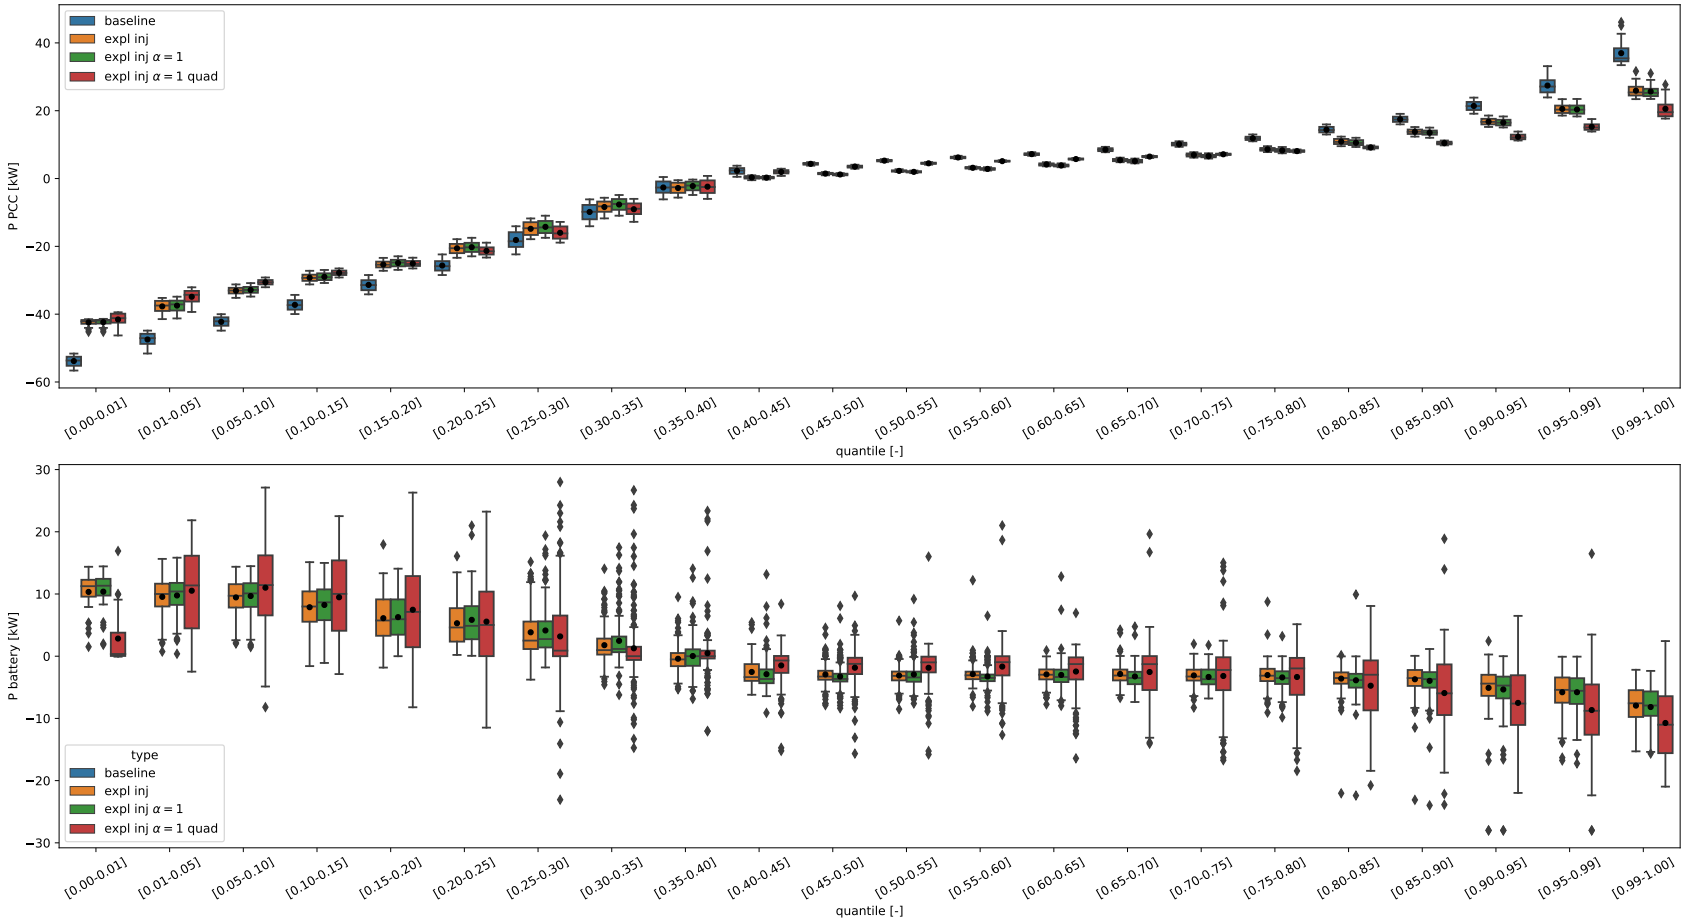

Figure 21: Effect of batteries control on the power at the point of common coupling (PCC) of the community. Top: boxplots of the power at the PCC per inter-quantile range. Bottom: boxplot of the total power of the batteries per inter-quantile range of the power at the PCC. The high variance of the batteries' power in the bottom plot starting from quantile 0.3 up to quantile 0.7 of the PCC's power, is due to the fact that the PCC's power is close to 0 in this range.

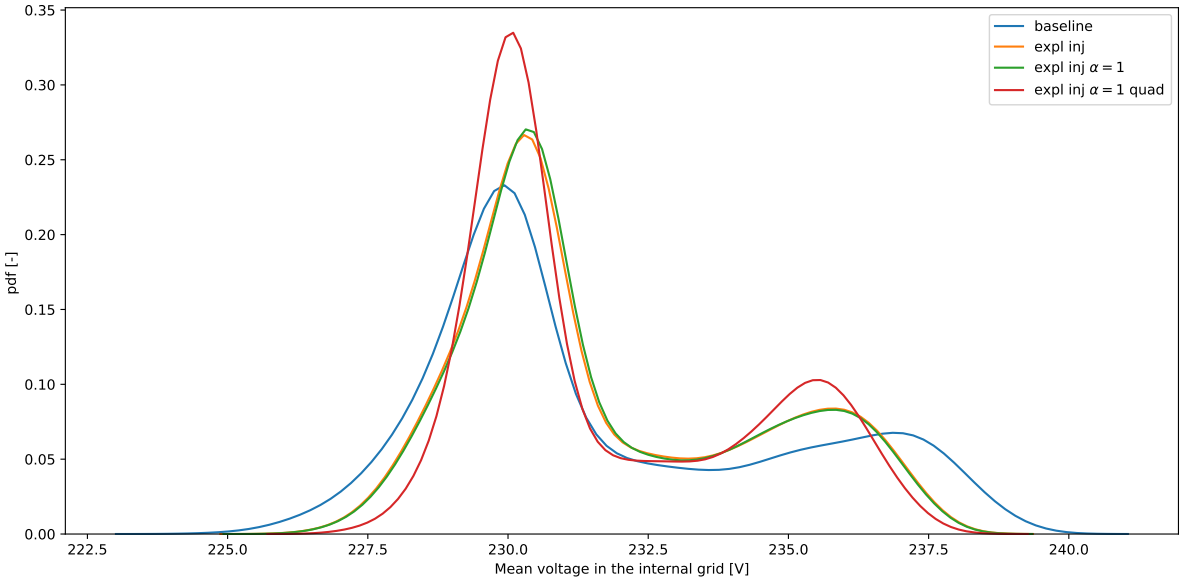

Figure 22: Probability density function (PDF) of the mean L-N voltage inside the community.

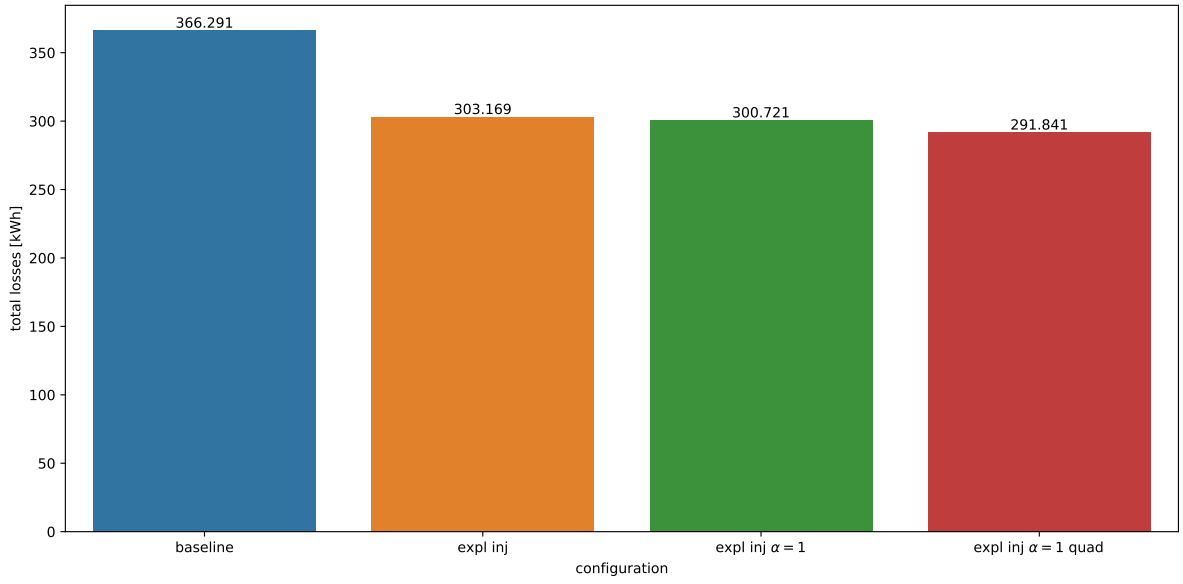

Figure 23: Total losses inside the community for the month of July.

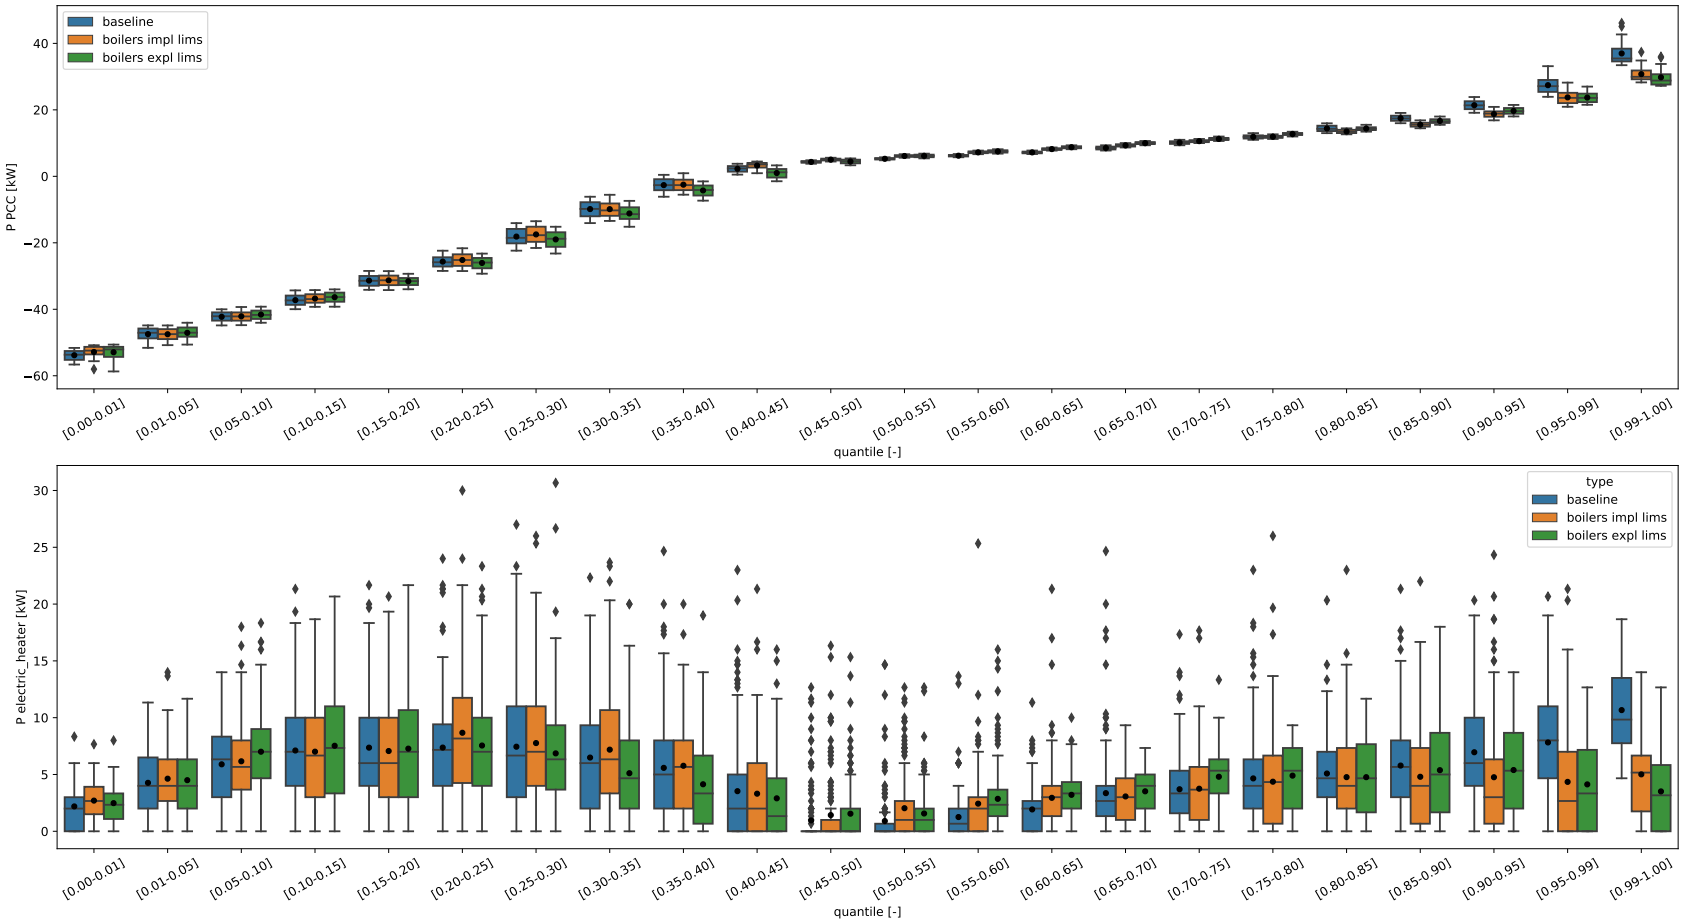

Figure 24: Effect of water heaters control on the power at the point of common coupling (PCC) of the community. Top: boxplots of the power at the PCC per inter-quantile range. Bottom: boxplot of the total power of the batteries per inter-quantile range of the power at the PCC.

9.2.2 Evaluation of economic impact

In the previous section, the different control strategies for batteries and water heaters have been evaluated against grid KPIs. We now focus on the financial performance for both the end-users and the administrator of the energy community. Two pricing schemes have been evaluated. The sharing-problem-based (SPB) pricing scheme against which the explicit control algorithms optimize, presented in section 4.3 and the LIC pricing scheme presented in section 9.1.3. In the first case, the community's surplus (equation 4) is equally split among the end-users and the community manager. The half of the surplus accruing to end-users is then divided according to their contribution to the surplus according to the rule defined in equations 10-12. Users who contributed the most to the surplus receive more. The redistributed surplus from now on is named bonus.
In the second case, the energy prices applied inside the LIC community are used. In this case, the prices for selling to and buying from the community are fixed and summarized in table 3. The self-consumed energy is equally split among the community members proportionally to their consumption and production, which results in buying and selling prices that are dynamic, but that for a given 15min slot are the same for everyone.

	outside the community	inside the community (LIC case)
buying	21	16
selling	6	9

Table 3: Energy prices (in cts/kWh) applied in the evaluation of economic impact.

In both cases, the community administrator pays the bill at the coupling point, where the DSO's prices are applied and gets paid by the end-users according to the above-mentioned pricing schemes. The difference between the administrator costs and revenues is used to pay for the internal grid, cover the administrative costs, and ideally make some profit.
Both, the effect of batteries and water heaters control is evaluated in the next paragraphs.

Batteries The different control strategies presented in section 9.2.1 are evaluated. The total costs for the end-users and the administrator, for the SPB and LIC pricing schemes are shown in tables 4 and 5, respectively. The difference in the monthly costs with respect of the business as usual case is shown in figures 25 and 26, for the SPB and LIC pricing schemes, respectively. The business as usual pricing scheme is that in which the users are not members of an energy community. The baseline_scc case represents the case in which the users are members of a community, but no batteries are present.
It can be seen that the mere fact of being a member of a community generates considerable savings for end-users. And, as expected, the introduction of batteries, by encouraging self-consumption both locally and at the community level, further improves the financial situation, for the end-users and for the community administrator. The explicit coordination method in which agents communicate with each other is the most effective from a financial point of view. In both pricing schemes, the battery owners are among those who are profiting the most from being inside a community. This is desirable, as they have to amortise the cost of the battery and fair since their actions contribute the most to increase the community's self-consumption. At the moment, the preliminary results (one month simulations) are not enough to estimate a payback time for the installation of a BES system. Based on the simulations done for the month of July, considering the explicit no injection case reported in table 5 and the installation cost of a residential BES of 515 CHF/kWh (as per the Tesla Powerwall 2), we can esteem a lower bound for the average payback time for the end user would be of 8.9 years for a 10 kWh BES system.
In section 9.2.1, we also proposed two alternative models, in which the batteries only optimised the community's own consumption, regardless of the effect this might have on their individual costs: models *expl inj* $\alpha = 1$ and *expl inj* $\alpha = 1$ *quad*. Both of the above models are not guaranteed to reach a generalised Nash equilibrium and consequently have the potential to reduce savings (if not cause additional costs) to individual users. Their financial figures are shown in tables 6, 7 and figures 27, 28. As could already be seen by analysing performance against grid parameters (section 9.2.1), the model in which alpha is set equal to 1 differs little from that in which it is not, also with respect to financial performance. As

	user costs bau	bonus	user costs tot	PCC costs	PCC balance
baseline bau	1476.3	0	1476.3	0	0
baseline scc	1476.3	-453	1023.3	600.5	-422.9
expl inj lims	1376	-519.6	856.4	360.4	-496
expl no inj	1386.3	-456.4	929.9	500.5	-429.4
expl no inj lims	1387.7	-465	922.8	484.2	-438.7
impl inj	1401.2	-519	882.2	387.2	-494.9
impl no inj	1400.5	-457.8	942.7	512.2	-430.5
no coordination	1412.2	-454.7	957.4	530.8	-426.6

Table 4: Total costs (in CHF) for the endusers and the administrator of the community using the explicit costs redistribution mechanism for the case of battery control. *baseline*: baseline simulation with no batteries, *no coordination*: batteries optimizing local self-consumption only, *impl no inj*: implicit coordination, no injection into the grid allowed, *impl inj*: implicit coordination, injection allowed, *expl no inj*: explicit coordination, no injection allowed, *expl no inj lims*: explicit coordination, no injection allowed, power limits at PCC enforced, *expl inj lims*: explicit coordination, injection allowed, power limits at PCC enforced.

expected, the community administrator's financial figures improve when $\alpha = 1$ since the batteries only optimise for the community surplus, which directly translates to higher revenues for the administrator. One can also see that the total costs for the end-users are slightly lower in the $\alpha = 1$ case in both pricing schemes. However, the average costs for battery owners are slightly higher (figures 27 and 28), which is consistent with the fact that they are sacrificing part of their personal revenues in order to increase the general welfare of the community. These effects are even bigger in the case in which the batteries are contributing to the reduction of quadratic punishment at the coupling point. In this case, the administrator is earning more money, but this is not the case for the community members, especially the battery owners. This is because there is actually no financial incentive to mitigate power excursions at the coupling point and, therefore, the battery actions are not sufficiently rewarded. In this analysis, we did not include remuneration schemes to reward users for helping to comply with network constraints. For the case in which constraints are enforced using upwardly trimmed Lagrangian multipliers, putting a hard cap on them corresponding to the price the grid operator is willing to pay to enforce the respect of the grid constraint is a fairly straightforward way to implement such mechanism, as the actual costs are directly reflected in the optimisation function of each individual user. Such a mechanism could also be applied to perform peak shaving actively. In a big enough community, the administrator will need to pay peak costs at the coupling point. Peak mitigation could be implemented exactly in the same way in which grid constraints are enforced, using Lagrangian multipliers. This concept is easily explained using a concrete example. The current monthly community peak to date is 50kW. In this case, the administrator will set a constraint on the maximum power at the coupling point to 50kW and reward those who allow the power to remain below the set limit by placing a hard cap on Lagrangians that corresponds to a reward significantly higher than the average cost of energy. For example, the administrator could reward the end-users with 0.1CHF/kW, which for a 15-minute step would correspond to 0.4CHF/kWh. It is not trivial, however, to understand if the administrator will save money by doing so. This is due to the highly non-linear nature of the existing billing system, which is based on the maximum monthly peak. In other words, if for 15 minutes in a month something goes wrong (e.g. a battery goes offline, or the forecasts are extremely bad), all the work done by the batteries, and remunerated by the administrator, is wasted. This topic is being currently investigated and will be discussed in future reports in the context of the LIC project.

	user costs	PCC costs	PCC balance
baseline bau	1476.3	0	0
baseline scc	993.1	600.5	-392.7
expl inj lims	821.7	360.4	-461.4
expl no inj	899.5	500.4	-399.0
expl no inj lims	891.8	484.2	-407.7
impl inj	847.6	387.2	-460.3
impl no inj	912.2	512.2	-400.0
no coordination	927.1	530.8	-396.3

Table 5: Total costs (in CHF) for the endusers and the administrator of the community using the implicit costs redistribution mechanism for the case of battery control

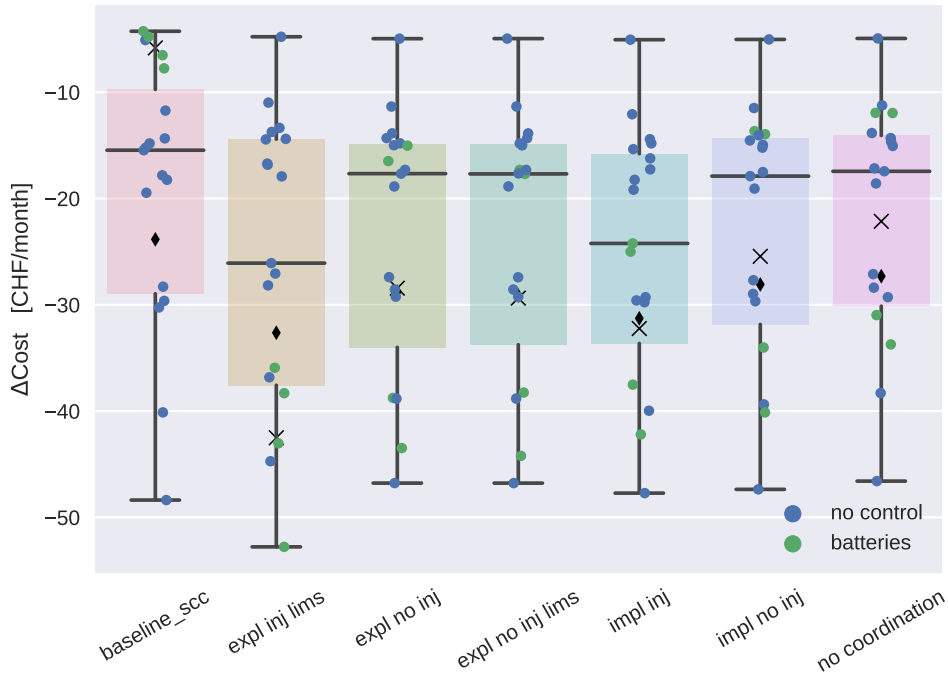

Figure 25: Boxplots of the difference in the monthly costs using SPB pricing scheme with respect of the business as usual case (in which the end-users are not part of an energy community), for the simulated month of July. The box shows the quartiles of the dataset while the whiskers extend to show the rest of the distribution, except for points that are determined to be outliers. Points that extend beyond 1.5 times the inter-quartile range, i.e. $v < q25 - 1.5(q75 - q25)$ or $v > q75 + 1.5(q75 - q25)$ are considered outliers. Black diamond: average price reduction for all the users. Black X: average price reduction for users with batteries.

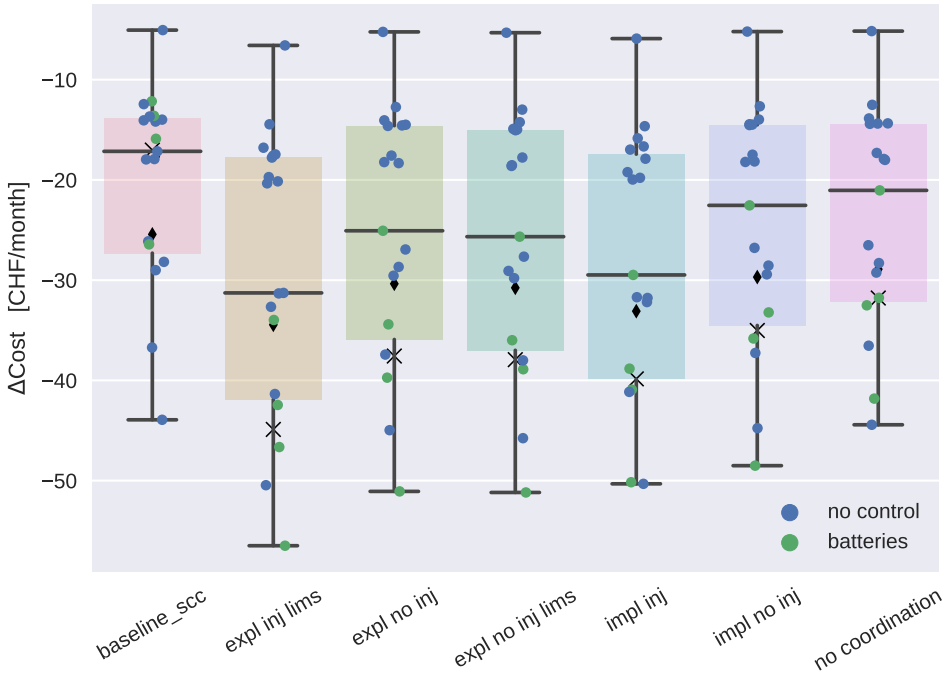

Figure 26: Boxplots of the difference in the monthly costs using LIC pricing scheme with respect of the business as usual case (in which the end-users are not part of an energy community), for the simulated month of July. Black diamonds: average price reduction for all the users. Black Xs: average price reduction for users with batteries.

	user costs bau	bonus	user costs tot	PCC costs	PCC balance
baseline bau	1476.3	0	1476.3	0	0
baseline scc	1476.3	-453.0	1023.3	600.5	-422.9
expl inj	1376.0	-519.6	856.4	360.4	-496.0
expl inj $\alpha = 1$	1391.2	-536.3	854.8	341.9	-512.9
expl inj $\alpha = 1$ quad	1438.8	-563.5	875.3	333.9	-541.4

Table 6: Total costs (in CHF) for the endusers and the administrator of the community using the explicit costs redistribution mechanism for the case of battery control, setting $\alpha = 1$ and using quadratic punishment.

	user costs	PCC costs	PCC balance
baseline bau	1476.3	0	0
baseline scc	993.1	600.5	-392.7
expl inj	821.7	360.4	-461.4
expl inj $\alpha = 1$	819.1	341.9	-477.1
expl inj $\alpha = 1$ quad	837.8	333.9	-503.9

Table 7: Total costs (in CHF) for the endusers and the administrator of the community using the implicit costs redistribution mechanism for the case of battery control, setting $\alpha = 1$ and using quadratic punishment.

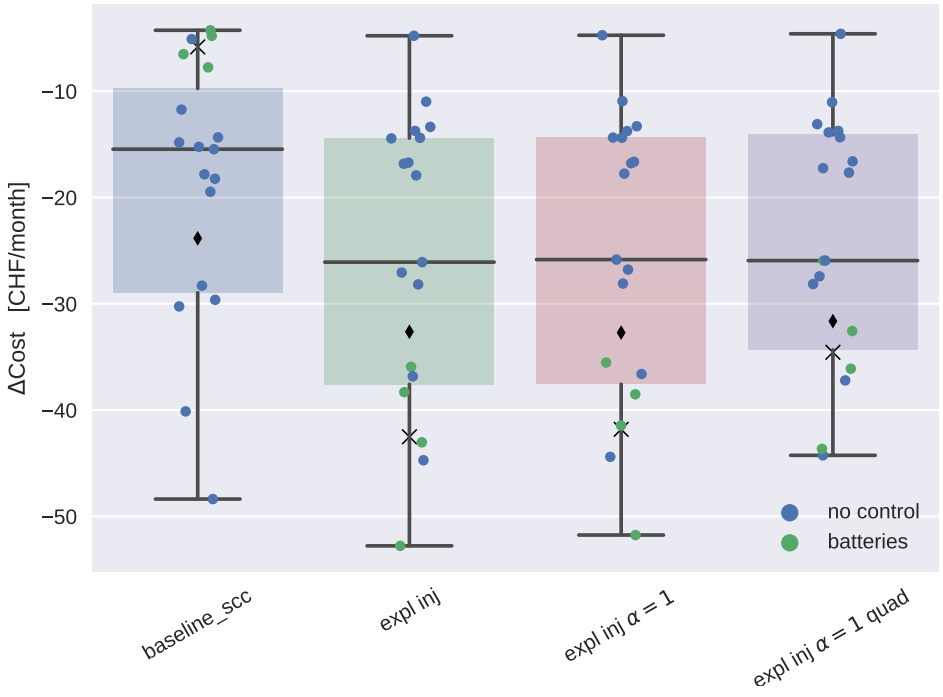

Figure 27: Difference in the monthly costs using SPB pricing scheme, with respect of the business as usual case, in which the end-users are not part of an energy community, for the simulated month of July. Batteries control case, setting $\alpha = 1$ and using quadratic punishment.

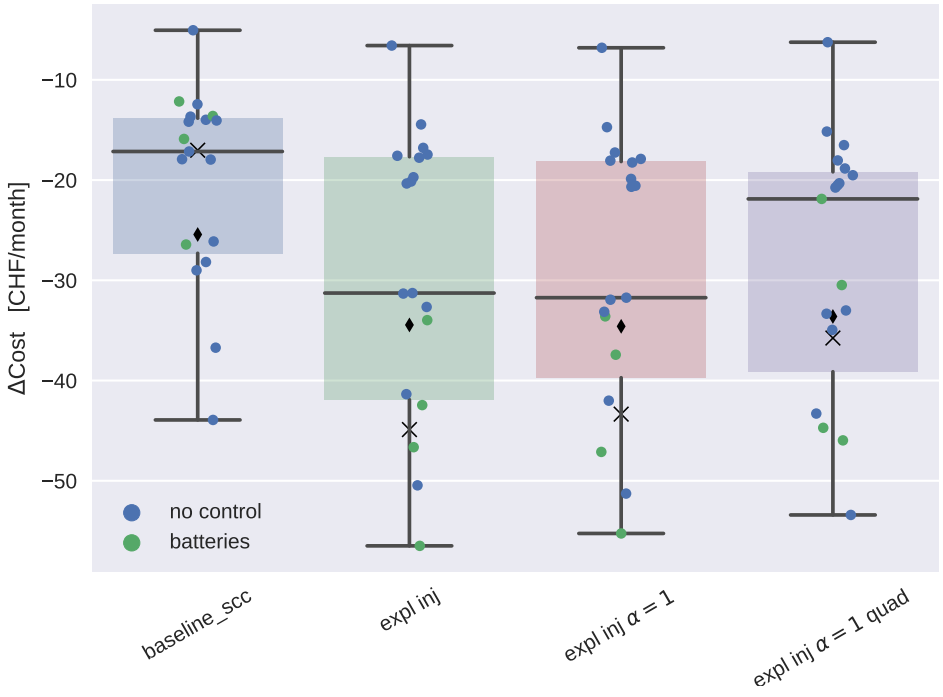

Figure 28: Difference in the monthly costs using LIC pricing scheme, with respect of the business as usual case, in which the end-users are not part of an energy community, for the simulated month of July. Batteries control case, setting $\alpha = 1$ and using quadratic punishment.

Water heaters Water heater control also leads to improved financial figures, although the effect is not as pronounced as with battery control, as shown in tables 8, 9 and figures 29, 30. As already mentioned in section 9.2.1, the lack of controllability of the water heaters (they can be forced off, but cannot be switched on on command) and the stringent constraints imposed on user comfort are limiting the effect of boilers control. Raising the boiler control algorithms' aggressiveness by lowering the turn-on time should help increase control performance against the financial KPIs, with the risk of violating user comfort. Additional simulations are planned with more aggressive settings.

	user costs bau	bonus	user cost tot	PCC costs	PCC balance
baseline scc	1476.3	-453.0	1023.3	600.487641	-422.7
expl lims	1442.3	-441.4	1000.8	588.231841	-412.6
impl lims	1454.0	-454.0	999.9	574.382804	-425.5

Table 8: Total costs (in CHF) for the endusers and the administrator of the community using the explicit costs redistribution mechanism for the case of boiler control.

	user costs	PCC costs	PCC balance
baseline scc	993.1	600.5	-392.7
boilers expl lims	971.4	588.2	-383.2
boilers impl lims	969.7	574.4	-395.3

Table 9: Total costs (in CHF) for the endusers and the administrator of the community using the implicit costs redistribution mechanism for the case of boiler control.

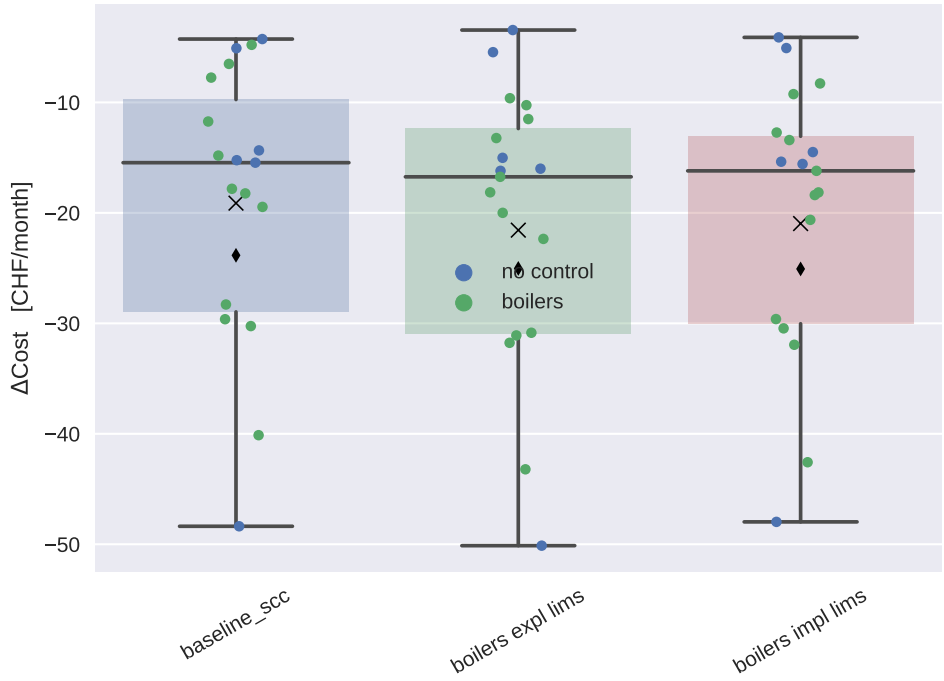

Figure 29: Difference in the monthly costs using SPB pricing scheme, with respect of the business as usual case, in which the end-users are not part of an energy community, for the simulated month of July. Water heaters control case.

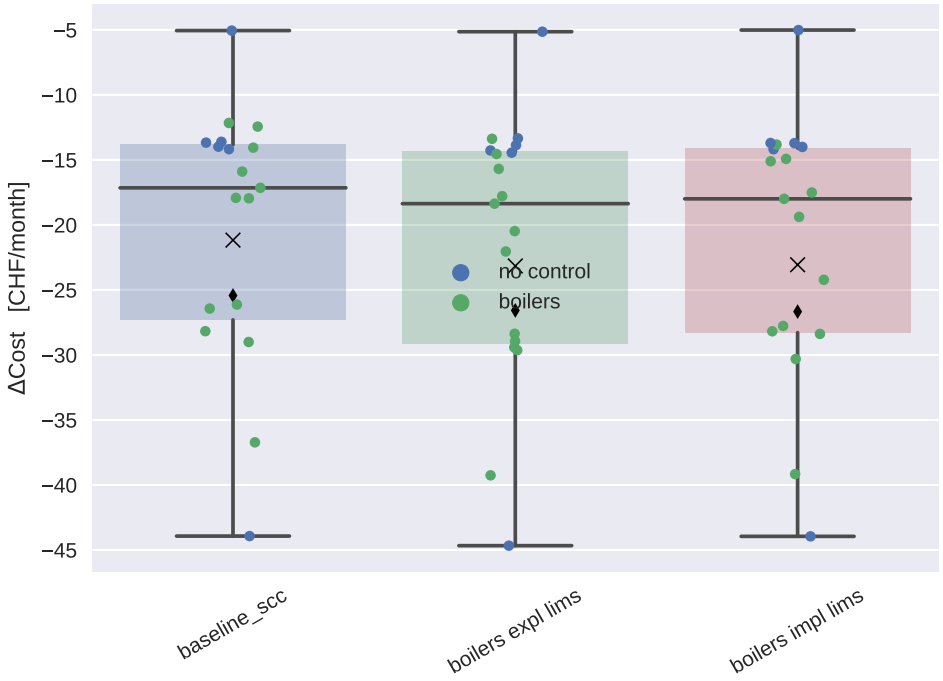

Figure 30: Difference in the monthly costs using LIC pricing scheme, with respect of the business as usual case, in which the end-users are not part of an energy community, for the simulated month of July. Water heaters control case.

9.2.3 Effect of forecasting algorithms on the performance

Residential power profiles are characterized by high variance and right-skewed data. This is due to the non-synchronous activation of loads and the presence of devices with a predominant power consumption (e.g. heat pumps and electric boilers) w.r.t. the other appliances. The activation of these loads doesn't usually have a strict seasonality effect; for example, heat pumps are usually characterized by several turn on events during the day, but the activation time seldom coincide with the one of the previous day. This makes the power profile hard to forecast.

In the following we describe the forecasting models that we have used for the closed loop comparison. We focused on methods which 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 [68] and [69], and focused on the Holt-Winters (HW) 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. For these latter models, 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 $Y \in \mathbb{R}^{(\mathrm{T-h-e})\times \mathrm{h}}$, and the Hankel matrix of the past regressors, $X_p \in \mathbb{R}^{(\mathrm{T-h-e})\times \mathrm{n_x e}}$, where n_x is the number of regressors. Verbosely, X_p and Y can be written as:

$$X_p = \begin{bmatrix} x_{1,t-e} & x_{1,t-e+1} & \dots & x_{1,t} & x_{2,t-e} & \dots & x_{n_x,t} \\ & & & \dots & & & \\ x_{1,t-e+1} & x_{1,t-e+2} & \dots & x_{1,t+1} & x_{2,t-e+1} & \dots & x_{n_x,t+1} \\ x_{1,T-2h} & x_{1,T-2h+1} & \dots & x_{1,T-h} & x_{2,T-2h} & \dots & x_{n_x,T-h} \end{bmatrix} \tag{74}$$

$$Y = \begin{bmatrix} y_{t+1} & y_{t+2} & \dots & y_{1,t+h} \\ & & & \dots \\ y_{T-h+1} & y_{T-h+2} & \dots & y_T \end{bmatrix} \tag{75}$$

where $x_{1,t}$ stands for the first regressor at time t. In hour case, we fixed $h = 288$, corresponding to a prediction horizon of 24 hours ahead with a 5 minutes sampling time. The past regressors matrix X_p is then augmented with categorical time features, e.g. day of week, and numerical weather prediction (NWP) variables, to obtain the final regressors matrix X.

Holt-Winters model with double seasonality The Holt-Winters (HW) model [70] is a special class of the exponential smoothing [71], which consists of three smoothing equations, such that the final prediction is a combination of the level a, trend b and seasonality s. We tested different flavors of the HW families and based on performance, we adopted a double seasonality additive HW:

$$\begin{aligned} \hat{y}_{t+h} &= (a_t + hb_t) + s_{1,t-p(1)+1+(h-1)\setminus p_1} + s_{2,t-p_2+1+(h-1)\setminus p_2} \\ a_t &= \alpha(y_t - s_{1,t-p_1} - s_{2,t-p_2}) + (1-\alpha)(a_{t-1} + b_{t-1}) \\ b_t &= \beta(a_t - a_{t-1}) + (1-\beta)b_{t-1} \\ s_{1,t} &= \gamma_1(y_t - a_t - s(2, t-p_2)) + (1-\gamma_1)s_{1,t-p_1} \\ s_{2,t} &= \gamma_2(y_t - a_t - s(1, t-p_1)) + (1-\gamma_1)s_{2,t-p_2} \end{aligned} \tag{76}$$

where α, β, γ_1 and γ_2 are parameters to be learned from data, while $p_1 = 96$ and $p_2 = 672$ are the periods of the seasonalities, and $\setminus$ is the modulo operator. The values for p_1 and p_2 correspond to a daily and weekly period. The model (76), and HW in general, do not include exogenous inputs. Since quantities like external temperature and irradiance are important explanatory variables in load forecasting, we included them with an a-priori linear detrend, such that the new target is $y = y - X\beta_d$, where X is a three column matrix containing GHI, T and the unit vector (for the intercept), and β_d is the vector of linear coefficients. Usually, a single set of α, β, γ_1 and γ_2 values is fitted, and the prediction of each step ahead is obtained applying equations 76 recursively, as usually done for state-space systems. To increase the accuracy of the method, we instead fitted 5 different models: the first two models have the only purpose of predicting the first and second step ahead, respectively. The third model predicts the steps from the third to the fifth, the fourth model predicts the steps up to the 20th and the last one predict the resto of the steps up to 288. Each model has its own set of parameters, which where fitted through

random search with a budget of 3000 samples. Since the trend of the time series was negligible, we fixed the β parameters to 0.

Single GBM A single gradient boosted model was fitted to predict all the 288 steps ahead. In order to obtain this, we stacked 288 copies of the features matrix (74), each one decorated with an additional column containing the number of the step ahead to predict. The target was replaced with the opportune reshape of matrix (9.2.3), $\tilde{Y} \in \mathbb{R}^{(T-h-e)h}$.

Independent GBMs In this approach, we fitted 288 models, each one taking as input the same matrix (74), but predicting different columns of (9.2.3). This allows the model to be more expressive (having much more parameters w.r.t. the single model strategy; as we will see this is especially beneficial for the first steps predictions. The drawback is an increase of computational time.

Independent GBMs with Huber loss In figure 31, 1 month power profiles of 4 buildings are aggregated in boxen plots based on the hour of the day, and plotted along with observations (black dots). As can be seen, the distribution at each hour of the day is severely skewed, approaching a bi-modal distribution for night hours. In the literature, prediction of unbalance data is usually threaded with the use of tweedie loss [72]; however, this approach is useful only in the case of extremely unbalance and zero inflated data, which is not the case as we can see in figure 31. In an attempt to make the model more robust against outliers, we fitted independent GBMs using Huber loss with standard parameters.

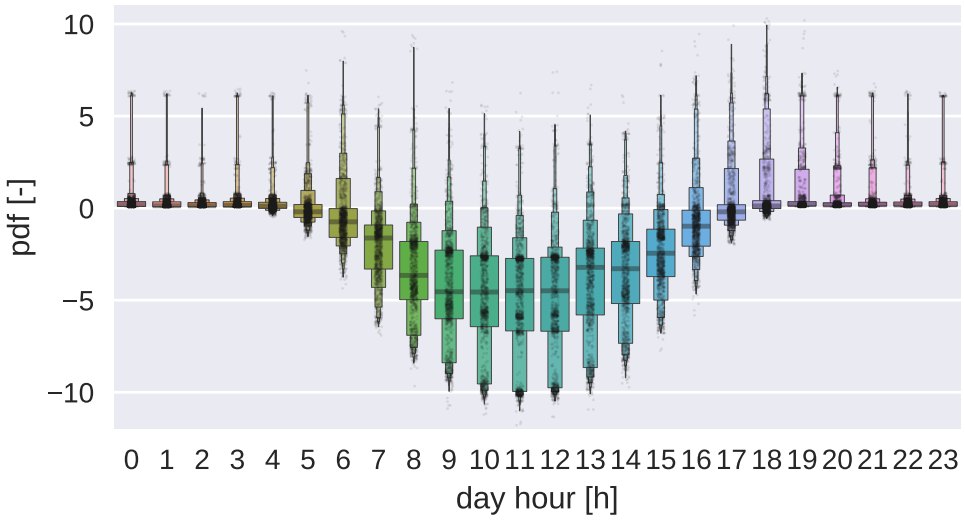

Figure 31: Boxen plots and observations (black dots) of the power distributions of 4 controlled buildings. The observations are groped by hour of the day. The w analyzed buildings has roof-mounted PV plants. All day hours present skewed, but different, distributions.

Hybrid GBMs This final model is an hybrid approach between the Single GBM and the independent GBMs approach. We fitted 5 independent models for the first steps ahead, which we expected to be more important in terms of controllers' performances, and a single GBM for the rest of the steps.

9.2.4 Forecasting accuracy

The forecasting accuracy was tested on a 1 month simulation. Figure 32 shows an example of day ahead predictions for the household's power profile of one of the controlled agents. It can be noticed that the HW model is able to approximately predict the high power peaks due to the presence of an electric boiler, capturing the seasonality of the time series. At the same time, since the turn on events of the boiler are not exactly periodic, the other models tend to underestimate the power profile, actually treating the turn on events as outliers.

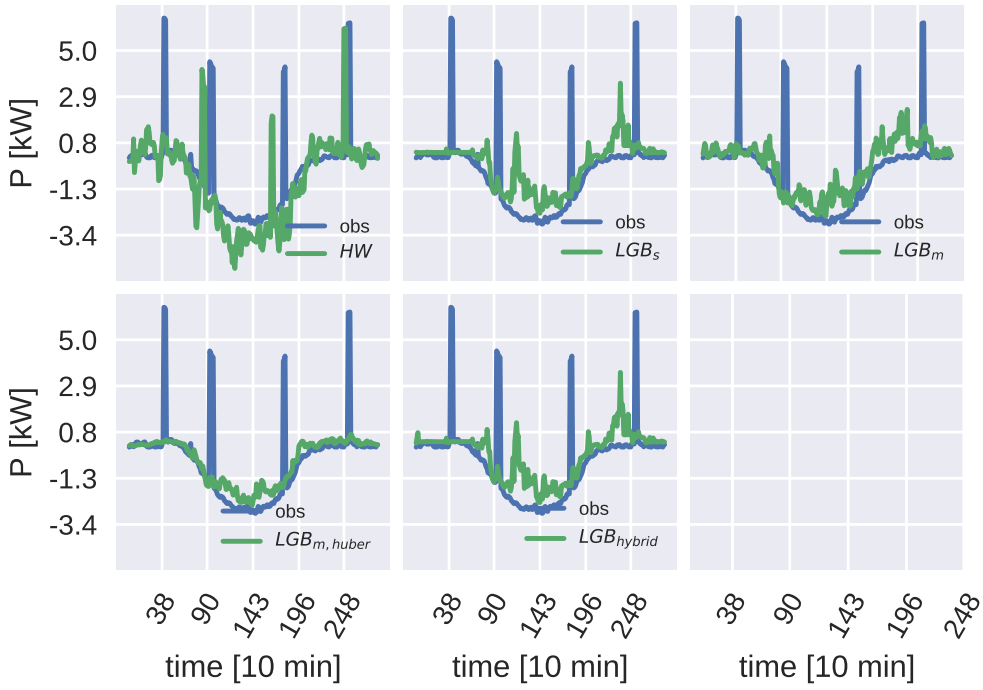

Figure 32: Example of predictions for an household main's power profile, for the tested forecasters.

More insight can be given by plotting the expected mean absolute error (MAE) of the different forecasters as a function of the day hour (vertical axis) and 5 minutes ahead time of the prediction (horizontal axis), as done in figure 33. For all the methods, we can see that the combination of step ahead and step of the day close to the antidiagonal present the lowest values in terms of MAE. This means that in a time window of a few hours centred around midnight, the predictability of the signal is high. While the non-parametric models are in general better than the HW, the approach using multiple independent models trained with Huber loss do not improve upon the others; this means that treating the turn on of big electric appliances as outliers doesn't improve the overall MAE. The effect of having several models for the initial steps of the HW based forecaster is clearly visible. This is also visible for the hybrid LGB model, and is better seen in figure 34. Here the MAE for the different forecasters is plotted as a function of step ahead (left) and the time of the prediction (right). The best forecaster in predicting the first step ahead is the HW based forecaster. The independent GBM models and the hybrid approach achieved similar accuracy for the first step, while using the Huber loss deteriorated the performances. In general the hybrid approach combines the first step accuracy of the independent models forecaster with the more stable accuracy of the single GBM forecaster for successive steps ahead. Plotting the MAE as a function of the prediction time highlight the more stable predictions of the non-parametric models w.r.t. the HW model, which is less accurate when predicting the next day ahead during post meridian hours.

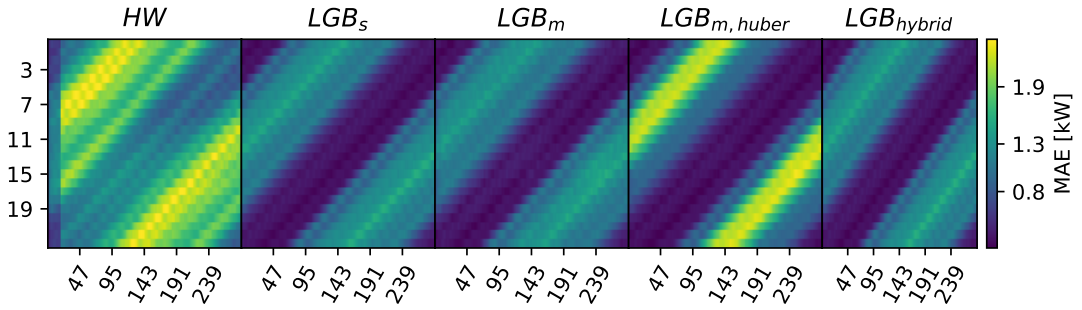

Figure 33: Results in terms of MAE for the households with controlled batteries. The results where mediated using bins of day hour (vertical axis) and 5 minutes ahead time of the prediction (horizontal axis).

9.2.5 Closed loop performances

Figure 36 and 35 show the differences in monthly costs with respect to the business as usual, in the case of the explicit and implicit control strategies, for different forecasters. Additionally, we tested the HW and LGB hybrid models in the case in which the battery's constraint for which it cannot charge nor inject in the main grid are removed. In this last case, the battery complementarity constraint on charging and discharging operations is modeled with an integer variable and a big M formulation. For both the business models the change in costs w.r.t. the base case is always positive (meaning a lower cost w.r.t. the base case), as expected since the methods guarantee a reduction of costs for the end-users by construction. The reduction of costs is higher for households with batteries since those are the only controllable devices in this test. The outlier in both plots is the kindergarten, in which a 27 kWp photovoltaic system and a battery with a capacity of 60 kWh are installed. For the implicit coordination, figure 36, no substantial difference is seen among the forecasters, in terms of distribution means. For the explicit coordination case, the cost reduction for the HW is slightly higher in expectation. A more meaningful comparison can be done plotting the differences for the HW and LGB hybrid forecasters, in the case in which injection and charging into e from the grid are allowed for batteries, which is shown in figure 37 and in figure 38 for the implicit and explicit coordination cases, respectively. Both for the implicit and the explicit coordination cases, two households with a boiler and a battery worsen their performance when using HW forecaster, w.r.t. the case in which the battery can't inject or charge in and from the grid. Since the relaxed problem allows the battery to perform arbitrage, expanding the feasible space for charging and discharging operations, this can be due to the low accuracy of the HW forecaster for the higher steps-ahead. This hypothesis is strengthened by the fact that the LGB hybrid forecaster, which has a higher accuracy w.r.t. the HW model for higher steps-ahead, shows a consistent decrease of costs for the end-users with a battery. While for the implicit coordination the LGB hybrid model guarantee a reduction in monthly costs for all the participants, this is not the case for the explicit coordination, in which agents without a battery see a slight increase of costs w.r.t. the case in which batteries cannot interface with the main grid. This is due to the redistribution model for the explicit coordination presented in section 3.1, which incentives those agents which had a greater impact in reducing the surplus function of the SCC.

Effect on power at PCC Figure 39 shows the effect of the different forecasting algorithms on the distribution at the PCC. The HW model is surprisingly good at coordinating batteries when the uncontrolled power at the PCC is in its exreme quantiles. The second best performer in terms of effectiveness in coordination when the PCC power is at extreme values is the hybrid GBM. This suggest that the accuracy on the prediction of the first step ahead is the most important factor when considering grid constraints. Figure 40 shows the pdf of the power at the PCC, which gives a more complete overview on the different performances of the forecasting models. Once again, it's clear that the best models at reducing the variance of the power at the PCC are the HW and the hybrid GBM models.

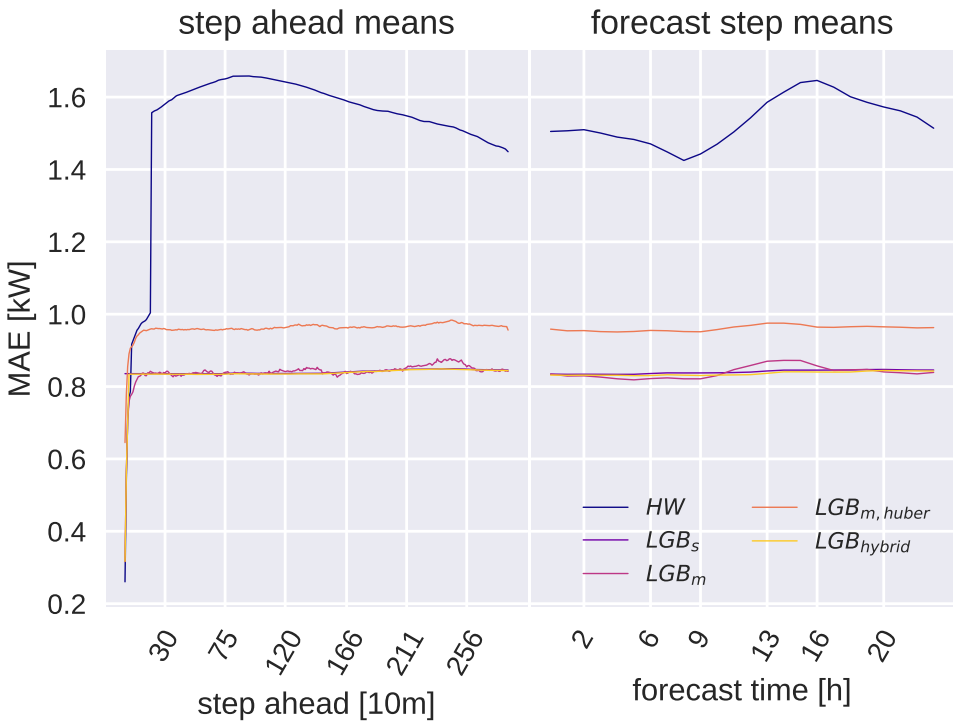

Figure 34: Mean results in terms of MAE for the households with controlled batteries. Left: results mediated on the step ahead. Right: results mediated on the time of the day on which the forecast was performed.

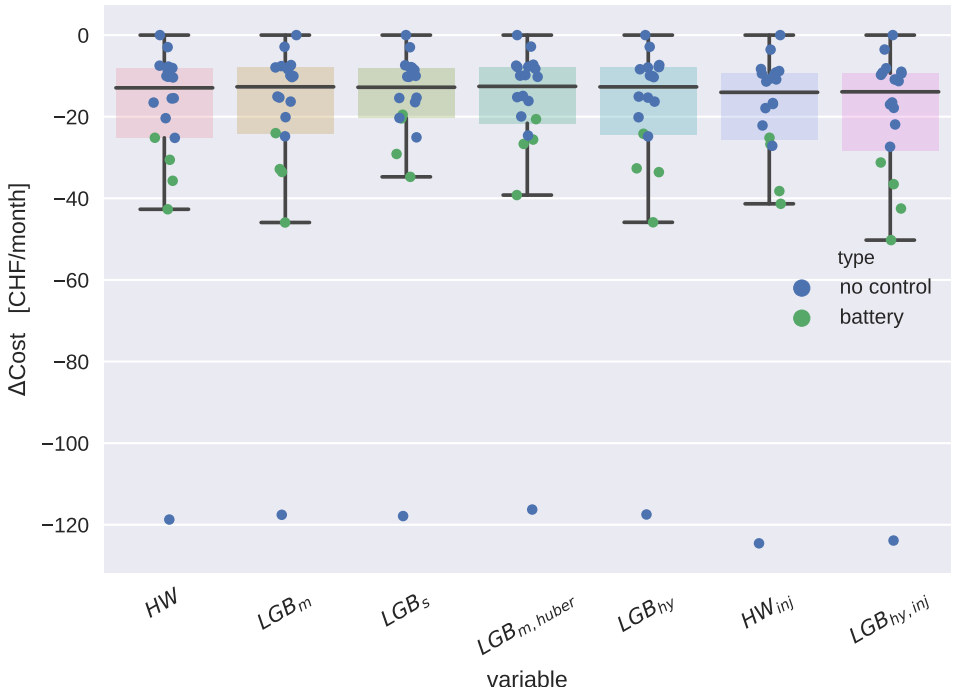

Figure 35: Difference in monthly costs w.r.t. the BaU, for the implicit coordination.

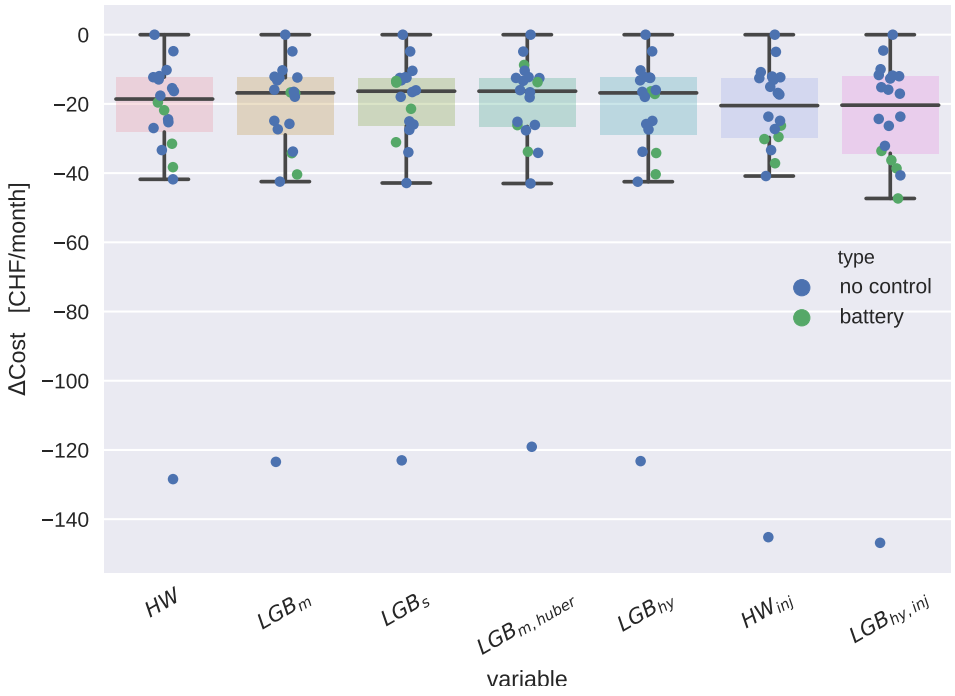

Figure 36: Difference in monthly costs w.r.t. the BaU, for the explicit coordination.

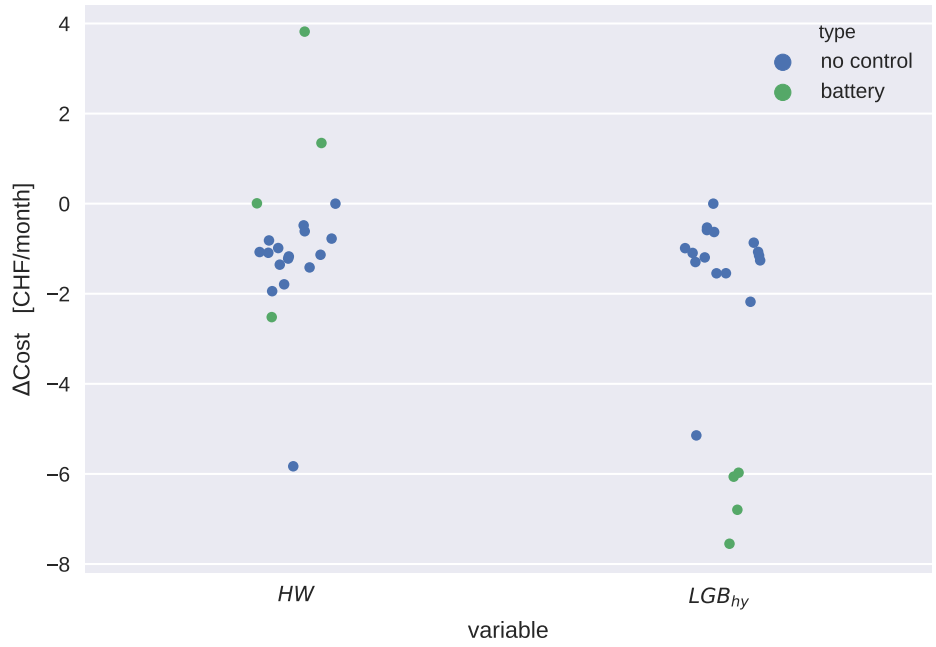

Figure 37: Difference in monthly costs w.r.t. the BaU with no allowed injections, with implicit coordination.

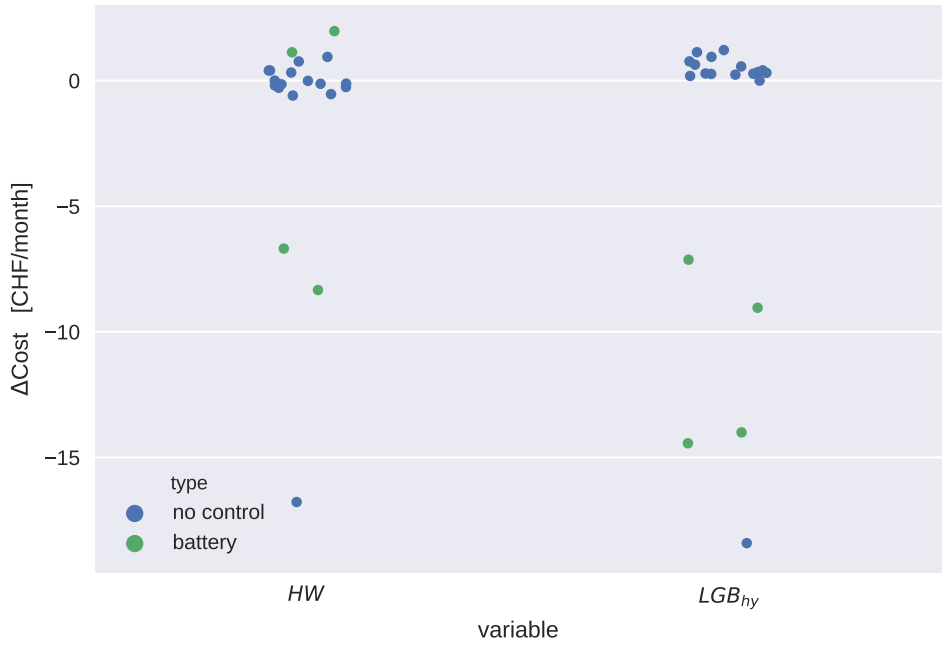

Figure 38: Difference in monthly costs w.r.t. the BaU with no allowed injections, with explicit coordination.

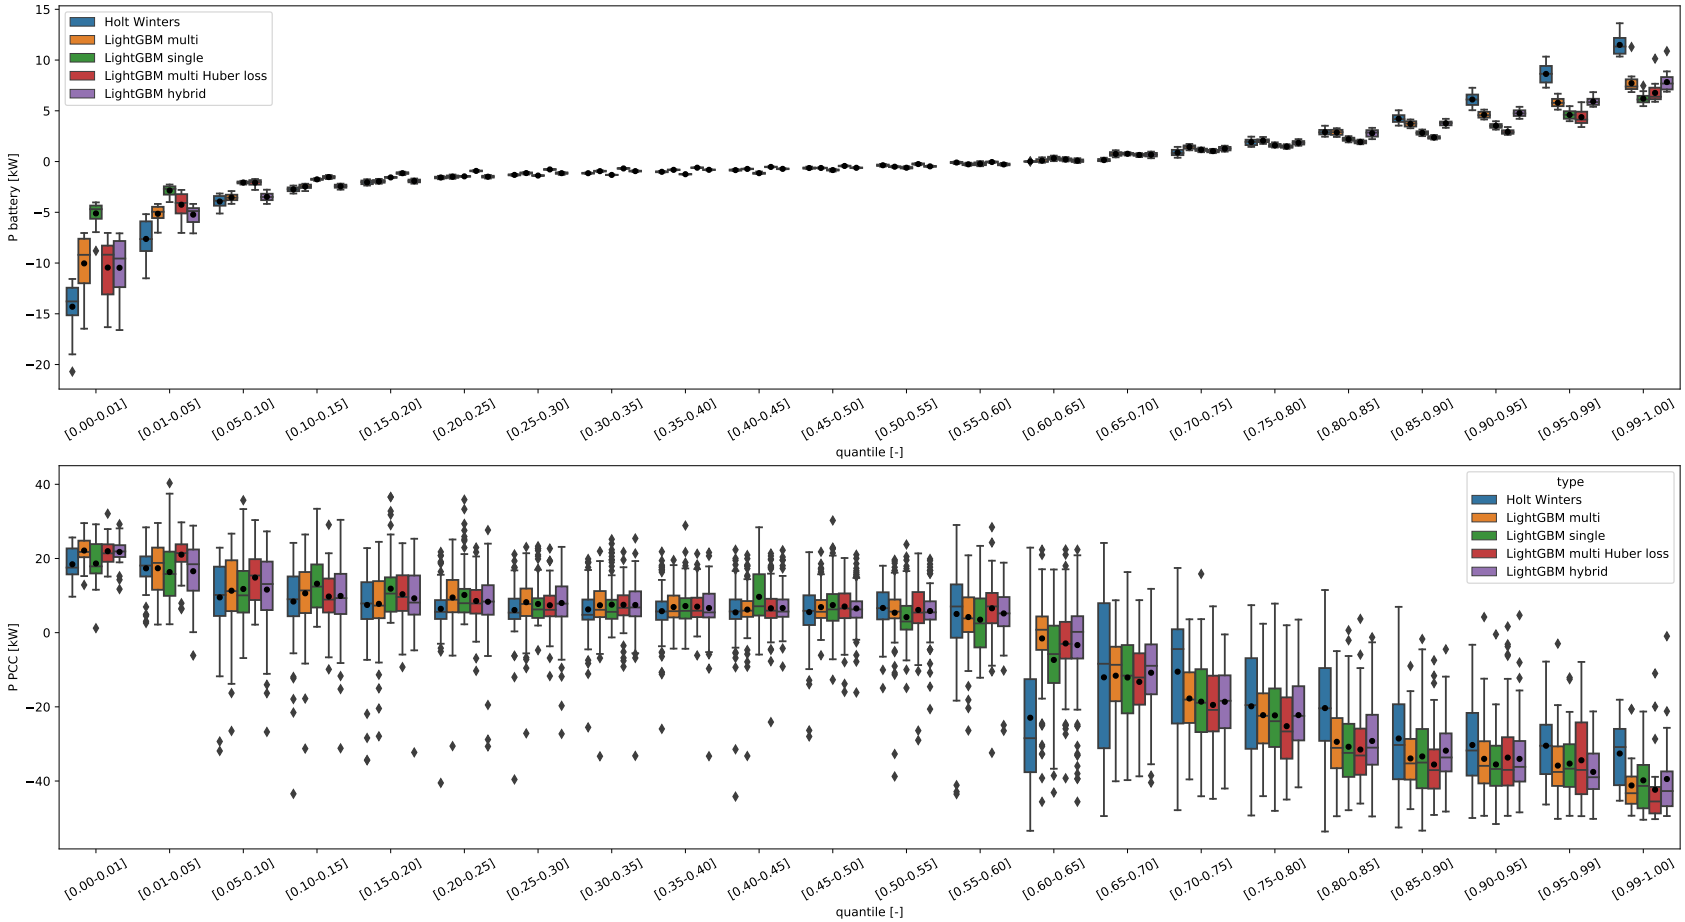

Figure 39: Effect of the different forecasting algorithms coupled with battery control on the power at the point of common coupling (PCC) of the community. Top: boxplots of the power at the PCC per inter-quantile range. Bottom: boxplot of the total power of the batteries per inter-quantile range of the power at the PCC.

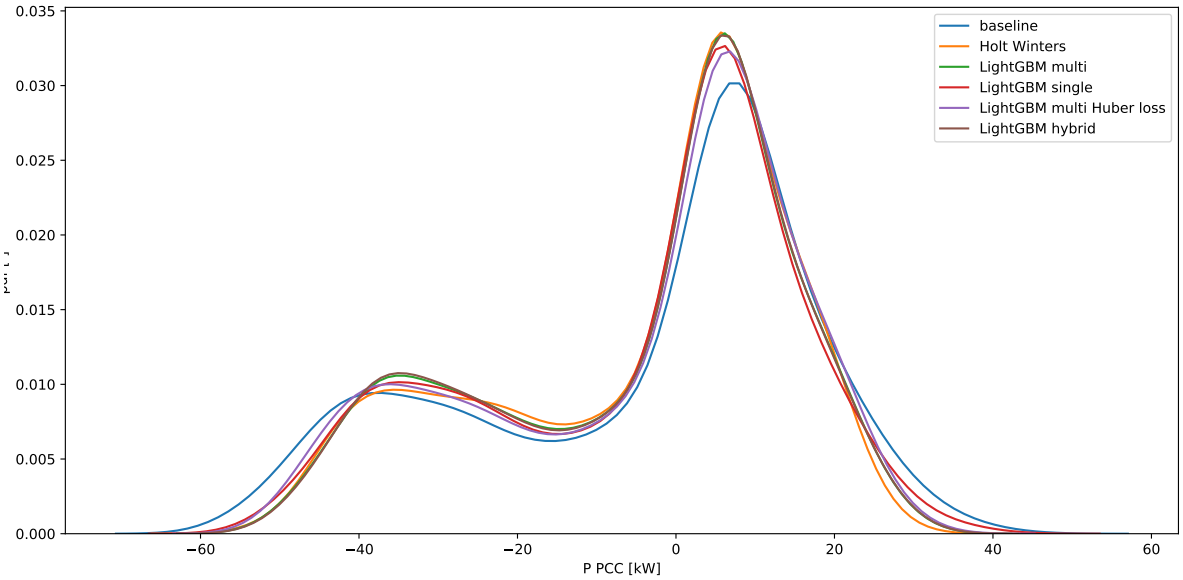

Figure 40: Probability density function (PDF) of the power at the PCC using diffrent forecasting algorithms.

9.3 Multi level algorithm

The algorithm to solve the multilevel coordination problem with grid constraints presented in 5 has been tested on several hierarchical configurations and distributions of smart agents represented by controllable batteries. Since the multilevel setting requires an higher amount of steps to converge with respect to the single level case, we only tested it without electrical simulation coupling, in order to test its correctness and characterize it in terms of computational time. A campaign of detailed simulations is undergoing and the results will be presented in future reports. An example of solution in terms of aggregated quantities and state of charge of controllable batteries for a 4 levels hierarchy is shown in figure 41. In order to test the algorithm also in the case of voltage violations, reasonable voltage sensitivity coefficients were assigned at random to the agents.

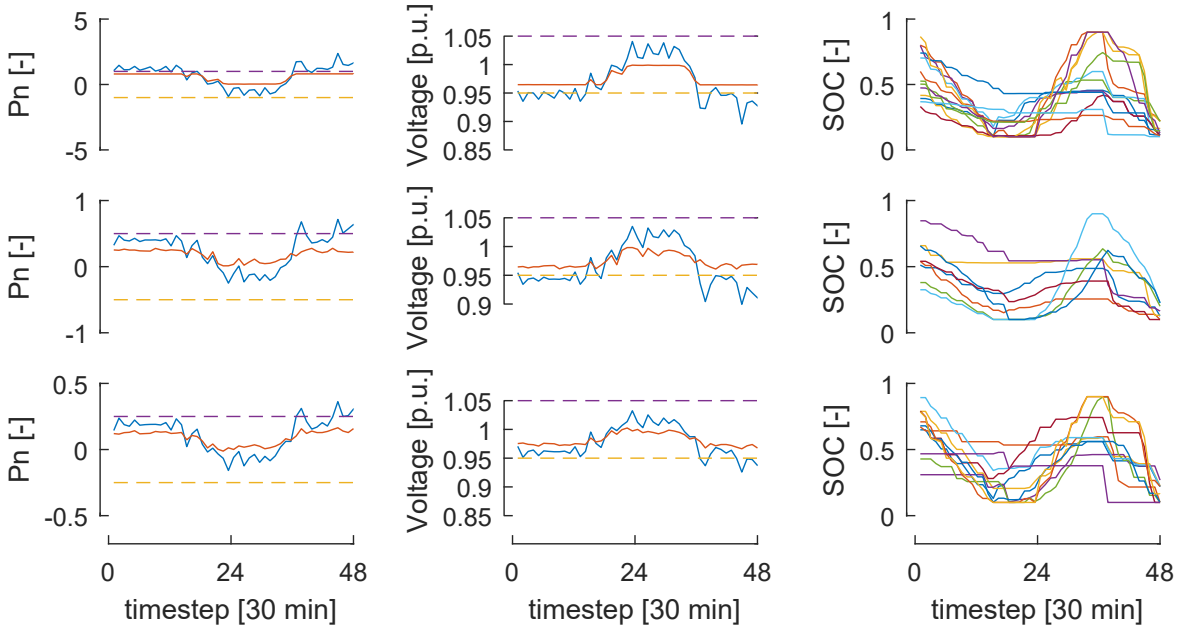

Figure 41: Example case. The considered hierarchical structure has 4 levels, and a single branching node in the first 3 levels. For the first two columns, the first row refers to the root node, second and third rows to the second and third level. First column represents the aggregated power profiles. Blue line: no battery actions. Red line: optimized power profiles. Dashed lines: power constraints. The second column represent voltage profiles. Blue line: no battery actions. Red line optimized voltage profiles. voltages and the state of charge (SOC) of the batteries are shown. The third column represents state of charge of prosumers' batteries in the second, third and fourth levels.

A trivial way of decomposing the problem would consist in reaching convergence in the lower level, then sending the results to the upper level of the hierarchy. In this way, the whole number of iteration before convergence in the lower level would represent a single iteration in the upper level. This solution concept would clearly result in an exponential computational time with respect to the number of considered levels, namely:

$$N_{tot} \sim tnN_i^L \prod_{l=1}^{L} N_b^l \sim N_i^L N_b^{\frac{L+L^2}{2}} \tag{77}$$

where t is the computational time for each agent for solving its local problem, n is the number of agents per level, N_i is the number of iterations before convergence for a single level, L is the number of levels, and N_b is the number of branches per level. Instead of following this strategy, we decompose the monolithic formulation of the problem to obtain nearly linear convergence, with respect to the number of

levels. A study on iterations needed for convergence is shown in figure 43, for hierarchical structures with different number of levels and controlled agents. This was obtained creating random hierarchical structure with a maximum of 4 aggregation levels and of 20 agents per terminal aggregation nodes. The near-linear convergence with the number of levels is shown in figure 42, where the computational time before convergence from 500 simulations is shown in terms of boxplots for each considered number of levels.

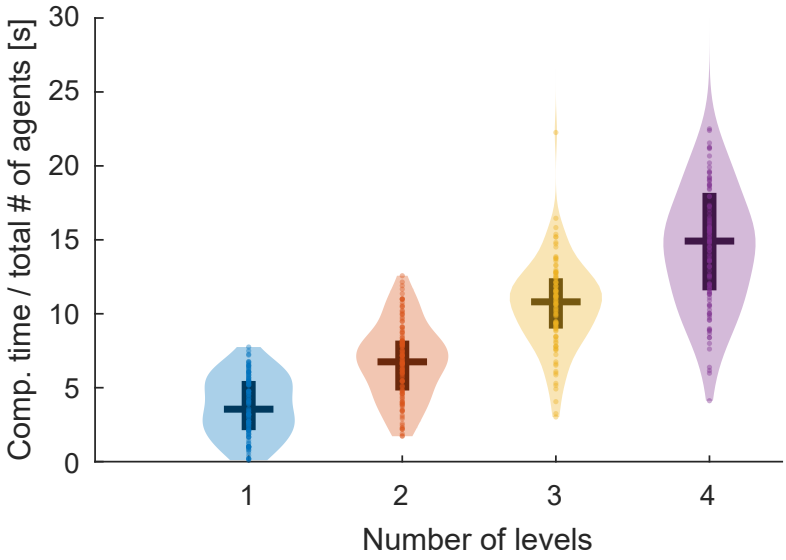

Figure 42: Estimated pdfs of the computational time divided by the total number of agents, as a function of the number of levels in the hierarchy. The vertical bar is the interquartile range, the horizontal line is the median.

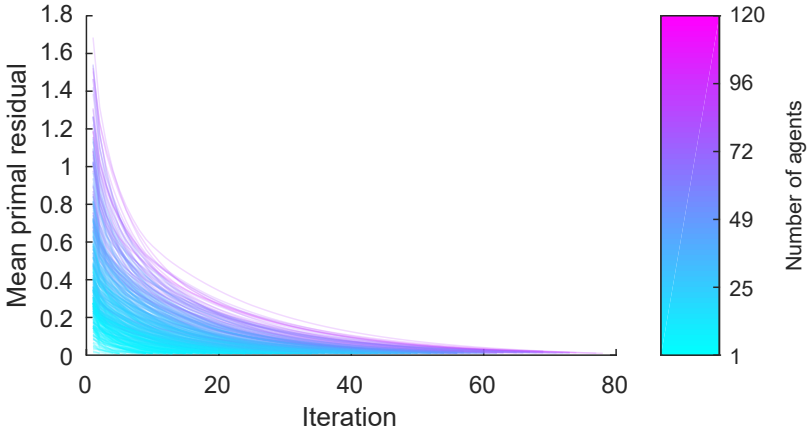

Figure 43: Mean overall primal residual for all the simulations, plotted versus the number of iterations before convergence, colored against the number of agents.

10 Pilot implementation

10.1 Pilot description

The LIC pilot[2] is located in Lugaggia, a small village near Lugano. In a part of the municipality, an energy community has been created with the collaboration of AEM[3], the local DSO. The energy community consists of 18 residential houses and a kindergarten. Four houses are equipped with photovoltaic systems installed on the roof, for a total nominal power of 37 kWp. A 27 kWp photovoltaic system and a battery with a capacity of 60 kWh are installed in the kindergarten. The figure 44 shows the pilot of Lugaggia with the nodes, corresponding to the different participants.

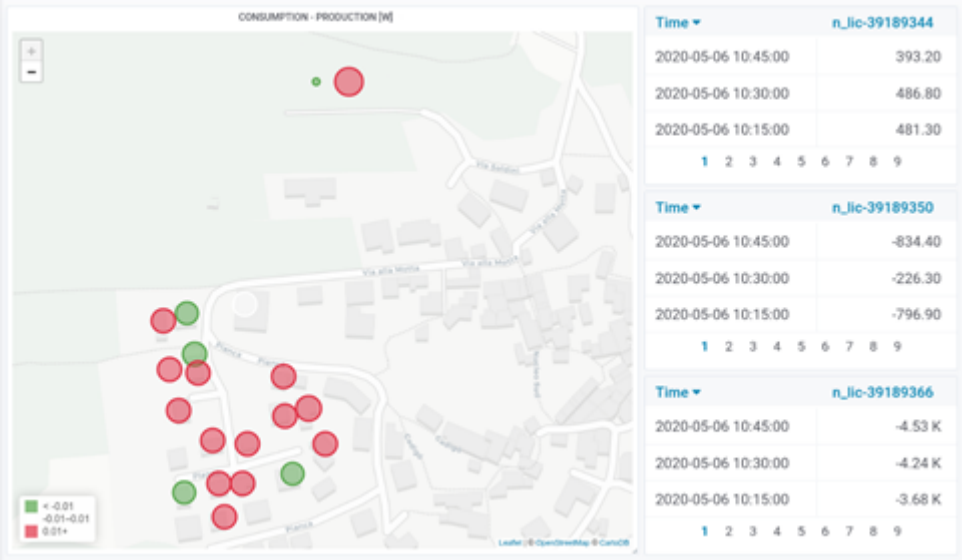

Figure 44: Lugaggia pilot

Each end-user corresponds to a node, associated with a single point of delivery, which is equipped with a smart meter. Besides, a Strato Pi CM[4], a compact industrial PC based on the Raspberry Pi Compute Module, is connected to each of the smart meters via an optical USB port. Each Strato acquires production/consumption data with a time resolution of 5 seconds and controls the two relays onboard the smart meter. In the basement of the kindergarten, close to the central battery, a NUC[5] machine has been installed. This machine acts as the aggregator of the community and as an interface with the battery.

10.1.1 Strato

The Strato Pi CM board is based on a Raspberry CM platform[6]. The Raspberry CM combines the computational power an easiness to use of the Raspberry Pi, i.e. a complete Linux operating system based on an ARM v8 platform, with the high reliability and service continuity of an industrial PC. This is achieved, in particular, thanks to the absence of an SD card, which is substituted with a much more robust internal eMMC Flash, and thanks to the presence of a hardware watchdog. Lab tests were conducted to verify the stability of the system against sudden power outages. All the devices always

[2] https://lic.energy/
[3] http://aemsa.ch
[4] https://www.sferalabs.cc/strato-pi-cm/
[5] https://www.intel.com/content/www/us/en/products/boards-kits/nuc.html
[6] https://www.raspberrypi.org/products/compute-module-3

restarted the operating systems without problems. Table 10 summarizes the main features of the Strato devices installed in the Lugaggia pilot.

CPU	4 ARMv8 64-bit 1.2GHz
RAM	1 GB
DISK	32 GB
OS	Raspbian GNU/Linux 9 (Stretch)

Table 10: Strato main features

Figure 45: Strato setup in a cabinet

Figure 45 shows a Strato installation in a Lugaggia cabinet. The Strato is in the red rectangle, and it is connected to the smart meter (green rectangle) via an optical USB reader (violet circle). The Internet connectivity is guaranteed by a USB dongle (blue rectangle), that provides a 4G data mobile connection. On the Strato, a Python application gathering electrical signals (e.g. active and reactive power, current, voltage, etc.) from the smart meter is continuously running.

10.1.2 NUC

Installed in the kindergarten where power outages are avoided, NUC is an embedded solution provided by Intel®, more powerful and robust than the Strato. A NUC was deployed in the pilot to have a machine more powerful than the Strato to run eventual tasks requiring significant computational resources. Table 10 summarizes its main features.

CPU	8 Intel(R) Core(TM) i7-8559U CPU @ 2.70GHz
RAM	16 GB
DISK	500 GB
OS	Ubuntu 18.04.4 LTS

Table 11: NUC main features

10.2 Auditable tariff

In an energy community as LIC, the members capability to audit the market tariffs applied by the DSO preserving the prosumers privacy is remarkably useful. Thus, a dedicated application, in the following named AT, was developed and deployed on the Strato devices. The basic idea behind AT is that the tariff is applied by the DSO according to the consumption of the entire community every 15 minutes. As a consequence, AT has to calculate the energy consumed by the whole community and provide it to the LIC members preserving the privacy, i.e. a member is not allowed to know the consumption of the single community prosumers. Currently, an instance of AT continuously runs on each Strato devices every quarter of hour. Fundamentally, the application is based on the following technologies:

- Decentralization based on sidechain usagefor the data storage
- Homomorphic encryption for privacy management

10.2.1 Decentralized data storage

AT provides a sidechain platform where the energy data can be properly stored. More specifically, Sidechains are separate blockchain networks, able to interact with the main chain. They have their own consensus mechanism, level of security and tokens. When sidechain security is compromised, the damage does not affect the other connected networks. Moreover, connected sidechains can transfer data. For example, tokens can be exchanged at a predetermined rate between the main chain and the sidechain. Basically, the main chain should provide the security of the entire ecosystem, while the transactions outsourced to the sidechain can sacrifice decentralization in return for scalability and velocity.
This solution appears to be extremely suitable to be used in the data management of energy communities. Indeed, custom chains could also be tailored for embedded devices such as the Strato and used to store data about nodes and interact with other main blockchains.
The sidechain running on LIC Strato is based on Tendermint[7] framework application, a consensus/networking algorithm that can be easily configured and used to create a custom sidechain. Figure 46 shows the structure of the sidechain network. Each node corresponds to a Strato and, consequently, to a prosumer. In addition, the orange nodes are the unique nodes allowed to create new blocks on the sidechain. The arrows show the P2P connections among the nodes.

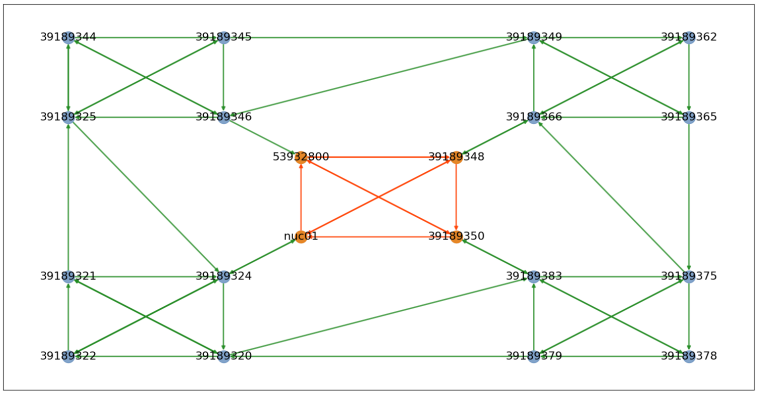

Figure 46: LIC network structure

As a consequence, using AT each node can save data about its energy consumption on the sidechain. Moreover, thank the data decentralization implicit in any blockchain technology, each node can access

[7] https://tendermint.com/

energy data about the other LIC members properly querying its local copy of the chain. As a consequence, a prosumer can calculate the LIC consumption and audit the tariff, but it violates the users privacy. The following Section 10.2.2 describes how it is possible both to audit the tariff and to preserve privacy.

10.2.2 Privacy management

The decentralized approach described in Paragraph 10.2.1 provides an ideal platform to store energy data. Unfortunately, it does not guarantee privacy preservation without proper data encryption. Thus, the AT instances running on LIC nodes save encrypted data about consumption on LIC chain. The data encoding is based on homomorphic encryption, which is locally provided by a custom application running locally on each Strato. Thanks to this type of data encoding, a node is not able to decrypt the information related to the other meters, i.e. privacy is preserved. Moreover, it can calculate the consumption/production of the entire community, and consequently verifies the tariff validity, which is a function of the community consumption.
Algorithm 1 briefly shows the sequence of the transactions performed by on the LIC nodes every quarter of an hour.

Algorithm 1 AT operations sequence

minute T=0,15,30,45
T+1: Registration
Each node saves on-chain the registration string and off-chain locally the related key
T+2: Encoding
Using the registrations saved by all the nodes at minute *T+1* and its key (saved off-chain), each node encodes the average consumption and production of the previous quarter of hour (e.g. T=15: data about [00:00-14:59] will be encoded). The encoded value is saved on-chain by all the nodes.
$\geq$**T+3: Sum and decoding**
Using the encoded values saved on chain at minute *T+2* each node can calculate the encrypted sum of the entire community and then, applying the homomorphic decoding, obtains the plain text value.

Fig. 47 reports the data flow showing how AT works in a simplified chain of three nodes. They save on the chain encoded values, as explained in 1 (orange arrows). Then a generic node (*N* can be *A*, *B* or *C*) is able to sum the encoded values and then decode the results obtaining the plain text total consumption or production (light blue arrow).

The data flow shown in Fig. 47 requires significant resources in terms of computational power and memory on the Strato boards. This is mainly due to the mathematical operations used by the homomorphic encryption. More precisely, on a Strato device AT needs less than a minute to decode and find the plain test solution, instead on the much more powerful NUC it converges to the solution approximately in 5 seconds. It is meaningful to note how this amount of time is independent from the nodes number, being the great most of the resources used for the final decryption of the encoded sum. Consequently, the application scalability is guaranteed towards the community dimension.

10.2.3 Application sustainability

The AT application runs in the Strato devices, which compose the Tendermint sidechain in the LIC pilot. To analyze their sustainability, it has been properly tested both in a laboratory and on the pilot. More precisely, on LIC Strato the sidechain runs since June 2020 managing the transactions performed by AT every 15 minutes. Table 12 summarizes the information about the transactions actuated on the chain and the results in terms of resources usage.

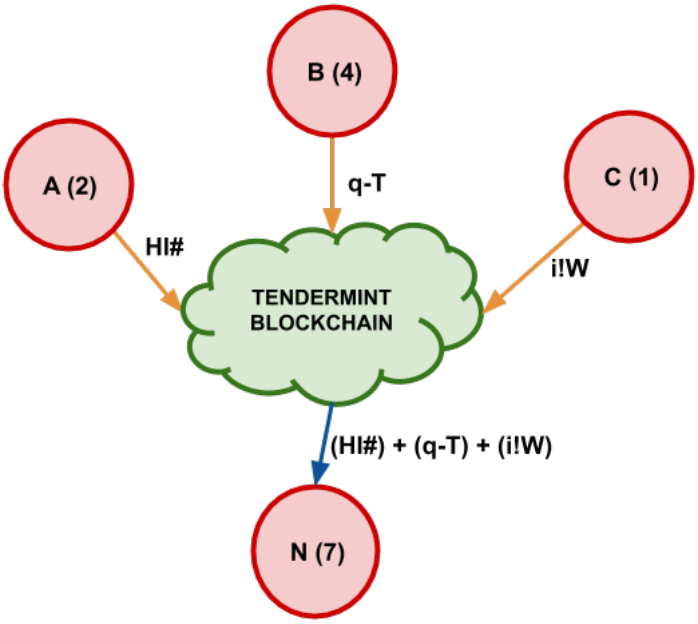

Figure 47: AT data flow

Nodes composing the chain	20
Time resolution	15 minutes
Transactions/15m	3
Transactions/1d for a node	288
Transactions/1d	5760
Load average (15 minutes)	0.3
CPU usage	5%
Memory usage	77 MB
Disk usage per day (community)	55 MB
Disk usage per day (single node)	2.7 MB
Disk usage per year (community)	20 GB
Network usage per day (download)	118 MB
Maximum reached data speed (download)	20.5 KB/s
Network usage per day (upload)	125 MB
Maximum reached data speed (upload)	19.7 KB/s

Table 12: Summary about tests performed on LIC nodes

The results reported in Table 12 shows how the sustainability of the Tendermint chain is guaranteed. As regards the computational power and memory usage, Strato devices can easily manage the applications. Instead, an aspect that has to be taken into account with special importance is disk usage. Indeed, 4 MB of space is daily used in the chain to store the data related to one node. Thus, the network cardinality has to be considered to assure safe data storing, especially if the number of transactions is much higher than in AT. However, the availability of solutions with more and more capacious disks continues to increase. Consequently, also this hardware constraint can be reasonably managed with a proper preliminary analysis.
Regarding the network usage, AT daily needs approximately 125 MB of data for the download, 121 MB for the upload. Furthermore, the data speed of the applications never overcomes 20 KB/s. These network requirements can be easily provided by standard connectivity services (e.g. wired connections, wi-fi). Moreover, also typical mobile subscriptions currently available on the market can be used for the data exchanging in the network as in LIC.

10.3 Unsupervised Non Intrusive Load Monitoring

The LIC's controllable loads are not directly monitored, as is usually done in demand side management projects. This poses challenges in the control of these devices, as planning the control of the devices requires to forecast their energy needs, that is, to know their disaggregated power profiles; on the other hand, not needing to install dedicated sensors for each controlled appliance, significantly lowers the technology cost of the DSM solution, as this only relies on an already installed sensor: the smart meter.

Many non-intrusive load monitoring algorithms (NILM) have been proposed in the literature; however, many of them cannot be applied out of the box in this case, as we are considering the following setting:

- unsupervised learning: we do not possess the ground truth for the output of the disaggregation. We are in the most general case in which we do not know: number of appliances, types of appliances, their nominal power, nor their characteristic power profiles.
- we are working with 10s sampling time, but we would like to use an algorithm able to disaggregate with lower sampling times (down to 10 minutes), as smart meters usually send data with a granularity of minutes.
- the algorithm must be able to run on the Strato's ARM architecture, which is also used to solve the optimal control problem: both training and test must run in less then one minute.

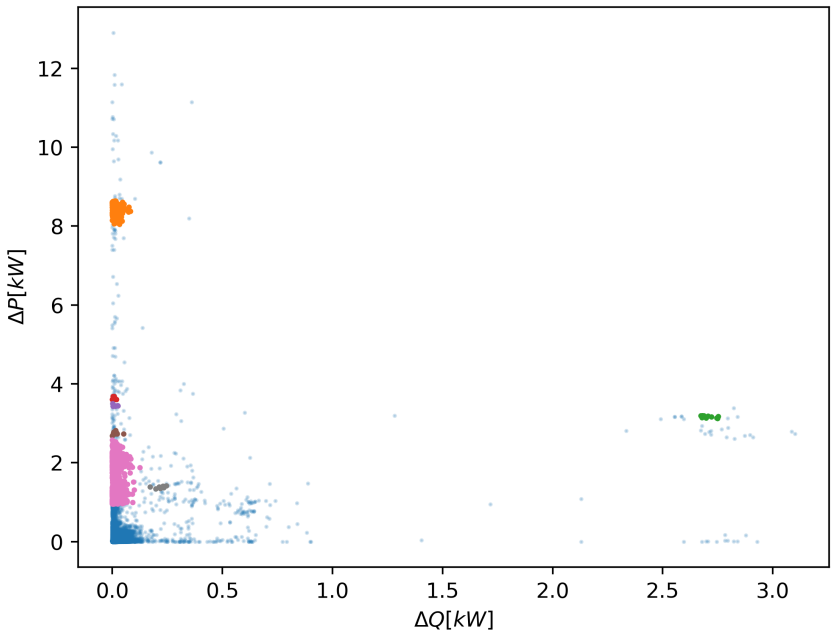

Figure 48: Unsupervised identification of clusters in the ΔQ-ΔP plane. The clustering algorithm correctly identified a boiler, in orange (ΔP of 8 with no associated change in reactive power) and an heat pump, in green (ΔP of 3.7 and ΔQ of 2.6, respectively).

Given all these constraints, we have developed an in-house disaggregation algorithm. The algorithm sequentially disaggregates a power signal of length N, $x = \{(x_i)\}_{i \in N} = \{(p_i, q_i)\}_{i \in N}$, given its first order (corrected) derivatives of active and reactive powers (p and q). The disaggregation is based on the following steps:

- Merge the first order derivatives of p,q. This is necessary to identify groups of jumps. Without this step, similar jumps will create lines in a scatter plot, due to the fact that jumps are often intra-

sampling. This means that a turn on/off of a load can span two timesteps, getting fragmented. An euristic is used to merge the derivatives.

- identify clusters in the jumps using an unsupervised algorithm (Bayesian Gaussian Mixture Model [73]). Fig. 48 shows the identification of clusters in the ΔQ-ΔP plane. The blue points are identified as not belonging to any clusters, while the other colors indicate points in a region of high density, significantly separed from other points. The clusters are then ranked with a pseudo-density measures and reordered based on their significance for disaggregation (the values of the centroids - higher jumps corresponding to more significance)
- clusters showing strong temporal correlations are merged together, obtaining meaningful groupings (a load can have a small turn on jump, a steady slope followed by a bigger turn off jump).
- derivatives are mapped to the most likely cluster. For each point in time associated with a cluster, we do the following. We start identifying a "ground" value, x_{gr}, for each jump in the time series belonging to the current cluster. Then, for each jump, the disaggregation follows 3 criteria for identifying the presence of the current load:

  1. if x_{gr} + centroid is close to the signal
  2. if subtracting x_{gr} + centroid to the signal don't make it too negative
  3. if subtracting x_{gr} + centroid to the signal reduces the signal variance

  if one of the three conditions is met, the algo keeps travelling forward (if the current jump was positive) or backward (if was negative). An example of the resulting disaggregated power profiles is shown in fig. 49.

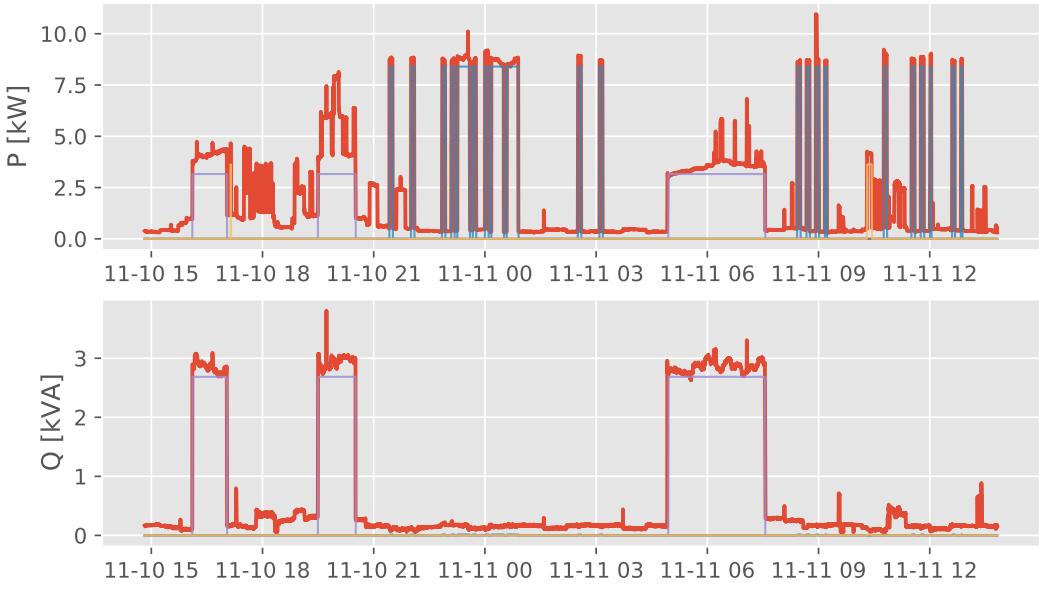

Figure 49: Example of disaggregated time series on test data. Upper plot: active power. Lower plot: reactive power. Three clusters are visible, in particular the boiler (only active power) and the heat pump (both active and reactive power).

10.4 Control: preliminary results

Algorithm tests in the LIC pilot project are underway. The thermal loads of users are controlled using the Strato devices, which control the relays of the smart meters via the optical interface. The algorithms described in Section 4 are being tested.
Preliminary tests have been successfully performed to demonstrate real-time control of the water heaters and the battery. An example is shown in figure 50. The implicit and explicit coordination mechanisms are not being tested. Figure 51 shows 2 days of control of one of the boilers in LIC, using the implicit coordination algorithm, on which the limits on the total number of daily actuation commands of the relay described in equation 27 has not yet been implemented. In the same figure one can also notice that the boiler is correctly disaggregated from the main power of the house. Unfortunately, due to low PV production, the battery remains idle in the winter months as the community never injects into the grid.

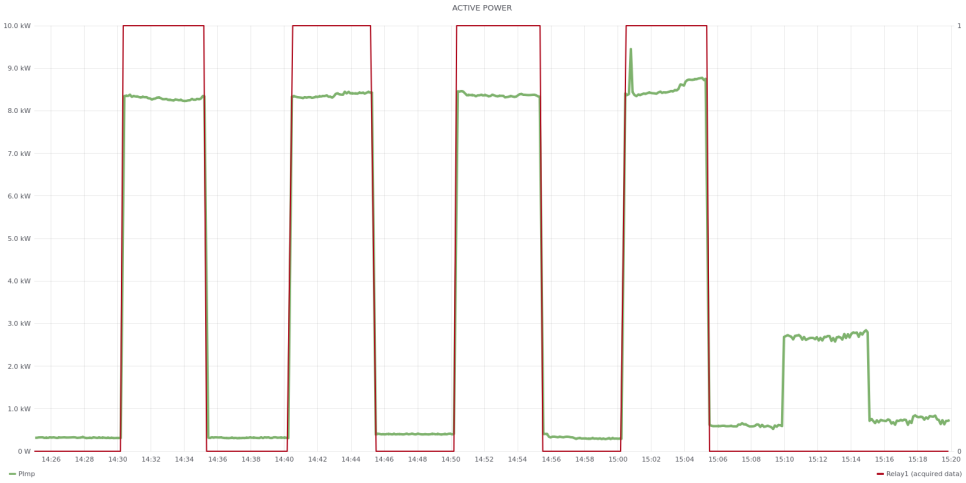

Figure 50: Control of a 8 kW boiler

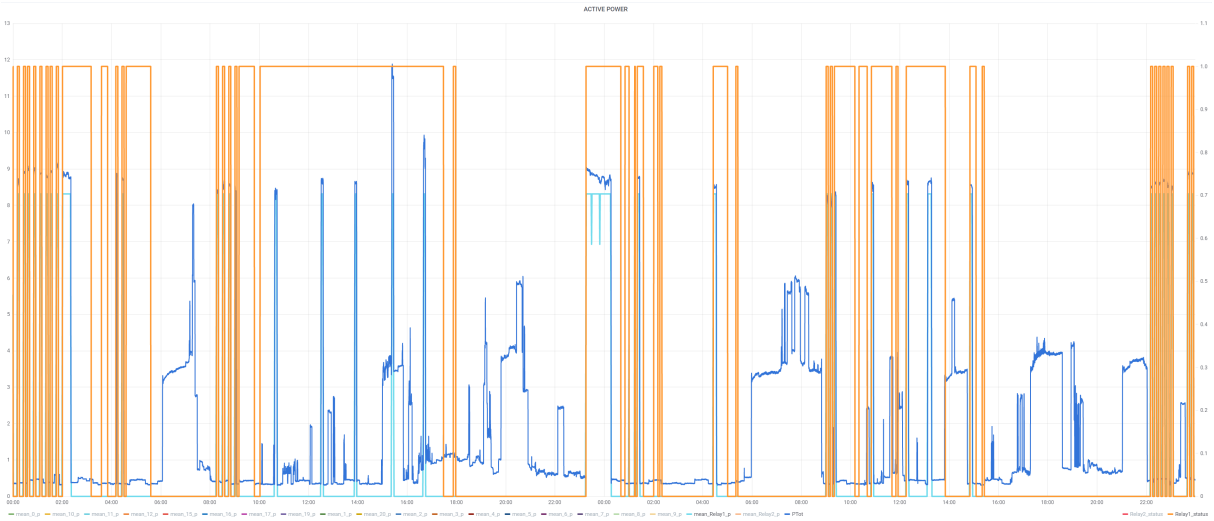

Figure 51: Example of water heater control. Orange line: relay status (0: forced off, 1: allowed on). Blue line: total active power of the smart meter. Cyan line: disaggregated water heater active power.

11 Conclusions

Self-consumption communities have been identified as the most interesting solution supported by the current legal framework, for building a hierarchical market taking into account grid constraints. The energy market proposed in section 3 generates a cost that depends on the consumption of all the energy community participants. That is, it is not directly proportional to the energy consumed by the end-user, and it gets lower when the self-consumption inside the energy community increases. Switzerland is embracing a causal principle on the price formation for end-users, as stated in the recent modification to the Federal Electricity Supply Act [1]. This means that the electrical bills "should reflect costs caused by end-users". However, the Electricity Supply Ordinance [74] states that DSO must guarantee to the end-users that at least 70% of their bills are directly proportional to their energy use; at the same time, they can offer opt-in tariffs in which this percentage is reduced. Under these constraints, the tariff proposed in section 3 can be potentially applied in Switzerland.

The energy market has been mathematically described as a sharing problem, in which agents have to share energy under coupling constraints representing electrical grid safety limits. The monolithic problem has been decomposed in an equivalent distributed formulation; in this setting, the market model has been analyzed from a game-theoretic perspective, using the framework of non-cooperative games. In section 3.1, solutions have been proposed to mitigate malicious actions from self-serving users, introducing repartition coefficients which are proportional to the role of each agent in reducing the system-level objective.

A hierarchical distributed formulation for multilevel control based on ADMM is presented in section 5. This formulation can be easily coupled with the non-cooperative energy market formulation presented in 3, and the following considerations on decentralized, parallel, and privacy-preserving formulations apply for both the single level and the multilevel energy markets.

In section 6, gossip protocols have been analyzed as a communication scheme for obtaining fully decentralized algorithms. Even if the considered class of problems (sharing problems) can be decentralized through gossip protocols, the number of messages required before convergence (and thus the overall computational time) is higher than a parallel solution scheme. Moreover, the presence of coupling constraints and the fact that the objective is a function of all users' actions require an additional cryptographic layer to preserve the privacy of private information. Given these considerations, two different homomorphic-based solutions were tested, which can be used to build a decentralized and fully-parallelized solution scheme, namely, the secure sum protocol and a modification of a secure voting protocol. The latter was finally chosen and implemented in the LIC pilot since it provides the strongest theoretically obtainable guarantees against agent collusions.

The energy market formulation and the optimal coordination schemes they generate has been tested against an "implicit" benchmark coordination scheme in which agents optimize for energy prices formed by an automated market making (AMM) mechanism. The rules of the AMM are simple to explain and the energy prices it generates can be used by agents to optimize their own energy consumption, providing an off-the-shelf distributed (indirect) coordination mechanism. The drawback is the lack of mutual information between agents' decisions and that grid constraints are not explicitly taken into account.

In section 9.2, we evaluated the performances of the algorithms on a simulation environment replicating the LIC pilot project. Both the explicit and implicit coordination schemes effectively reduced the overall variance of the power profile at the PCC; as expected the explicit coordination scheme performed better in terms of violation of power constraints at PCC, when power constraints are enforced. It is reasonable to expect an even wider spread between the two coordination mechanisms in terms of variance at the PCC, with an increasing number of controllable agents. We additionally tested variants of the explicit control mechanism presented in section 3, where the repartition coefficients for agents are unitary. This

case is equivalent to a centralized control where the only objective is the minimization of the costs at the PCC. Results showed that power distribution at PCC is only slightly affected by this modification, indicating that the distributed non-cooperative formulation of the sharing problem is already highly effective in minimizing the overall costs at the PCC. An additional variant of the control problem in which a small quadratic punishment was added shown the possibility of further reducing the overall costs while performing peak shaving. This additional cost reduction obtained punishing extreme charging/discharging operations is likely due to the mitigation of erroneous decisions introduced by imperfect forecasts. Furthermore, the effect of removing the constraint for which distributed batteries cannot discharge into the main grid was tested; this problem relaxation was shown to be beneficial for bot grid constraints and economic results. The economic analysis showed the same rank in results among different coordination schemes, w.r.t. costs at the PCC: allowing injection from distributed batteries and imposing a quadratic punishment on the overall PCC power profiles further reduces the costs at PCC. On the other hand, the quadratic punishment has a negative impact on the single agent's energy costs, when they are computed with the rules presented in 3, as expected.

In subsection 9.2.3 the effect of different forecasters was assessed on closed-loop performances for both the implicit and explicit coordination schemes when controlling batteries. The analysis showed how the accuracy for the first steps ahead is the most important factor for the reduction of the agents' costs. On the other hand, when charging/discharging from/into the grid is allowed, the ability to forecast step ahead further in time becomes increasingly important.

In section 9.3 the multilevel implementation presented in section 5 has been tested in a simplified setting, showing quasi-linear convergence time with the number of considered levels. More detailed closed-loop simulations for the multilevel setting, with a stronger coupling with an exact model for the electrical grid, are planned as an addendum to this report.

Finally, in section 10.1, the LIC pilot project implementation is described. End users with controllable devices have been equipped with local computing modules, which are connected to the smart meters and are able to solve the distributed problems required for the clearing of the energy market, exploiting the homomorphic encryption protocol presented in section 7 for keeping their power consumption profile private; at the same time, they have enough computational power to run the ad-hoc non-intrusive load monitoring disaggregation described in section 10.3, required to control thermal loads, such as heat pumps and electric boilers. The implementation of the billing scheme through the use of Tendermint is described.

12 Publications

L. Nespoli and V. Medici, "Constrained hierarchical networked optimization for energy markets," inProceedings - 2018 IEEE PES Innovative Smart Grid Technologies Conference Europe, ISGT-Europe 2018, 2018

L. Nespoli, M. Salani, and V. Medici, "A rational decentralized generalized Nash equilibrium seekingfor energy markets," in2018 International Conference on Smart Energy Systems and Technologies,SEST 2018 - Proceedings, 2018

L. Nespoli, "Model Based Forecasting for Demand Response Strategies," Thesis, 2019 https://doi.org/10.5075/EPFL-THESIS-9415

13 References

[1] "LAEI 734.7, https://www.admin.ch/opc/it/classified-compilation/20042411/index.html."

[2] B. Yang and M. Johansson, "Distributed optimization and games: A tutorial overview," in *Networked Control Systems*, 2010, vol. 406, pp. 109–148.

[3] B. HomChaudhuri and M. Kumar, "A newton based distributed optimization method with local interactions for large-scale networked optimization problems," pp. 4336–4341, 2014.

[4] E. Ghadimi, I. Shames, and M. Johansson, "Multi-step gradient methods for networked optimization," *IEEE Transactions on Signal Processing*, vol. 61, no. 21, pp. 5417–5429, 2013.

[5] B. Johansson, M. Rabi, and M. Johansson, "A randomized incremental subgradient method for distribued optimization in netowrked systems," *SIAM Journal on Optimization*, vol. 20, no. 3, pp. 1157–1170, 2009.

[6] D. K. Molzahn, F. Dorfler, H. Sandberg, S. H. Low, S. Chakrabarti, R. Baldick, and J. Lavaei, "A Survey of Distributed Optimization and Control Algorithms for Electric Power Systems," *IEEE Transactions on Smart Grid*, vol. 3053, no. c, p. 1, 2017.

[7] R. Scattolini, "Architectures for distributed and hierarchical Model Predictive Control - A review," *Journal of Process Control*, vol. 19, no. 5, pp. 723–731, 2009. [Online]. Available: http://dx.doi.org/10.1016/j.jprocont.2009.02.003

[8] R. Negenborn and J. Maestre, "Distributed Model Predictive Control: An overview of features and research opportunities," *Proceedings of the 11th IEEE International Conference on Networking, Sensing and Control*, no. April, pp. 530–535, 2014. [Online]. Available: http://ieeexplore.ieee.org/document/6819682/

[9] R. Halvgaard, L. Vandenberghe, N. K. Poulsen, H. Madsen, and J. B. Jørgensen, "Distributed Model Predictive Control for Smart Energy Systems," *IEEE Transactions on Smart Grid*, vol. 11, no. 4, pp. 1–9, 2014.

[10] P. Braun, T. Faulwasser, L. Gr, C. M. Kellett, S. R. Weller, and K. Worthmann, "Hierarchical Distributed ADMM for Predictive Control with Applications in Power Networks," *IFAC Journal of Systems and Control*, vol. 3, pp. 10–22, 2018.

[11] M. Diekerhof, F. Peterssen, and A. Monti, "Hierarchical Distributed Robust Optimization for Demand Response Services," *IEEE Transactions on Smart Grid*, vol. 3053, no. c, p. 1, 2017.

[12] M. Kraning, E. Chu, and S. Boyd, "Dynamic Network Energy Management via Proximal Message Passing," *Foundations and Trends in Optimization*, vol. 1, no. 2, pp. 70–122, 2013.

[13] N. Parikh and S. Boyd, "Proximal Algorithms," *Foundations and Trends®in Optimization*, vol. 1, no. 3, pp. 123–231, 2013.

[14] T. G. Hovgaard, K. Edlund, and J. B. Jørgensen, "The Potential of Economic MPC for Power Management," pp. 7533–7538, 2010.

[15] P. McNamara and S. McLoone, "Hierarchical Demand Response for Peak Decomposition," *IEEE Transactions on Smart Grid*, vol. 6, no. 6, pp. 2807–2815, 2015.

[16] K. Margellos, A. Falsone, S. Garatti, and M. Prandini, "Distributed Constrained Optimization and Consensus in Uncertain Networks via Proximal Minimization," *IEEE Transactions on Automatic Control*, vol. 9286, no. c, pp. 1–16, 2017.

[17] P. Scott and S. Thiebaux, "Identification of Manipulation in Receding Horizon Electricity Markets," *IEEE Transactions on Smart Grid*, vol. 3053, no. c, pp. 1–12, 2017.

[18] S. Weerakkody, X. Liu, S. H. Son, and B. Sinopoli, “A Graph-Theoretic Characterization of Perfect Attackability for Secure Design of Distributed Control Systems,” *IEEE Transactions on Control of Network Systems*, vol. 4, no. 1, pp. 60–70, 2017.

[19] P. Velarde, J. M. Maestre, H. Ishii, and R. R. Negenborn, “Vulnerabilities in Lagrange-Based DMPC in the Context of Cyber-Security,” *2017 IEEE International Conference on Autonomic Computing (ICAC)*, pp. 215–220, 2017. [Online]. Available: http://ieeexplore.ieee.org/document/8005352/

[20] L. Makowski and J. M. Ostroy, “Vickrey-Clarke-Groves mechanisms and perfect competition,” *Journal of Economic Theory*, vol. 42, no. 2, pp. 244–261, 1987.

[21] E. H. Clarke, “Multipart pricing of public goods,” *Public Choice*, vol. 11, no. 1, pp. 17–33, 1971.

[22] W. Vickrey, “COUNTERSPECULATION, AUCTIONS, AND COMPETITIVE SEALED TENDERS,” *The Journal of Finance*, vol. 16, no. 1, pp. 8–37, 1961.

[23] N. Noam, T. Roughgarden, E. Tardos, and T. Wexler, *Algorithmic Game Theory*. Cambridge University Press, 2007.

[24] D. C. Parkes and J. Shneidman, “Distributed Implementations of Vickrey-Clarke-Groves Mechanisms,” *Third Joint Conference on Autonomous and Multiagent Systems*, pp. 261–268, 2004. [Online]. Available: http://dash.harvard.edu/bitstream/handle/1/4054438/Parkes_Distributed.pdf?sequence=2

[25] A. Petcu, B. Faltings, and D. C. Parkes, “M-DPOP: Faithful distributed implementation of efficient social choice problems,” *Journal of Artificial Intelligence Research*, vol. 32, pp. 705–755, 2008.

[26] J. B. Rosen, “Existence and Uniqueness of Equilibrium Points for Concave N-Person Games,” *Econometrica*, vol. 33, no. 3, pp. 520–534, 1965.

[27] B.-g. Kim, S. Member, S. Ren, M. V. D. Schaar, J.-w. Lee, and S. Member, “Bidirectional Energy Trading and Residential Load Scheduling with Electric Vehicles in the Smart Grid,” *IEEE Journal on selected areas in communication*, vol. 31, no. 7, pp. 1219–1234, 2013.

[28] A. A. Kulkarni and U. V. Shanbhag, “On the variational equilibrium as a refinement of the generalized Nash equilibrium,” *Automatica*, 2012.

[29] D. Paccagnan, M. Kamgarpour, B. Gentile, F. Parise, J. Lygeros, and D. Paccagnan, “Distributed computation of Nash Equilibria in aggregative games with coupling constraints,” in *2016 IEEE 55th Conference on Decision and Control (CDC)*, Las Vegas, USA, 2016, pp. 6123–6128.

[30] G. Belgioioso and S. Grammatico, “Projected-gradient algorithms for Generalized Equilibrium seeking in Aggregative Games are preconditioned Forward-Backward methods,” *arXiv*, 2018.

[31] M. Hong, Z. Q. Luo, and M. Razaviyayn, “Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems,” *SIAM Journal on Optimization*, vol. 26, pp. 337–364, 2016.

[32] L. Nespoli, M. Salani, and V. Medici, “A rational decentralized generalized Nash equilibrium seeking for energy markets,” in *2018 International Conference on Smart Energy Systems and Technologies, SEST 2018 - Proceedings*, 2018.

[33] ——, “A rational decentralized generalized Nash equilibrium seeking for energy markets,” in *2018 International Conference on Smart Energy Systems and Technologies, SEST 2018 - Proceedings*, 2018.

[34] Y. Shoham and K. Leyton-Brown, *Multiagent systems: Algorithmic, Game-Theoretic, and logical foundations*, 2008.

[35] R. Cavallo, “Optimal Decision-Making With Minimal Waste: Strategyproof Redistribution of VCG Payments,” *Aamas-06*, pp. 603–607, 2006. [Online]. Available: cavallo-aamas-06.pdf

[36] R. B. Myerson and M. A. Satterthwaite, “Efficient mechanisms for bilateral trading,” *Journal of Economic Theory*, 1983.

[37] T. Tanaka, F. Farokhi, and C. Langbort, “Faithful Implementations of Distributed Algorithms and Control Laws,” vol. 4, no. 2, pp. 191–201, 2013. [Online]. Available: http://arxiv.org/abs/1309.4372%0Ahttp://dx.doi.org/10.1109/TCNS.2015.2489318

[38] L. Nespoli, M. Salani, and V. Medici, “A rational decentralized generalized Nash equilibrium seeking for energy markets,” in *2018 International Conference on Smart Energy Systems and Technologies, SEST 2018 - Proceedings*. Institute of Electrical and Electronics Engineers Inc., 10 2018.

[39] L. S. Shieh, H. Wang, and R. E. Yates, “Discrete-continuous model conversion,” *Topics in Catalysis*, 1980.

[40] L. Nespoli, A. Giusti, N. Vermes, M. Derboni, A. Rizzoli, L. Gambardella, and V. Medici, “Distributed demand side management using electric boilers,” *Computer Science - Research and Development*, vol. 32, no. 1-2, 2017.

[41] 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.

[42] K. Baker, J. Guo, and S. Member, “Distributed MPC for Efficient Coordination of Storage and Renewable Energy Sources Across Control Areas,” *IEEE Transactions on Smart Grid*, vol. 7, no. 2, pp. 992–1001, 2016.

[43] X. Su, M. A. S. Masoum, and P. J. Wolfs, “Multi-Objective Hierarchical Control of Unbalanced Distribution Networks to Accommodate More Renewable Connections in the Smart Grid Era,” *IEEE Transactions on Power Systems*, vol. 31, no. 5, pp. 3924–3936, 2016.

[44] 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.

[45] L. Stella, A. Themelis, and P. Patrinos, “Forward – backward quasi-Newton methods for nonsmooth optimization problems,” *Computational Optimization and Applications*, vol. 67, no. 3, pp. 443–487, 2017.

[46] A. Gascón, P. Schoppmann, B. Balle, M. Raykova, J. Doerner, S. Zahur, and D. Evans, “Privacy-Preserving Distributed Linear Regression on High-Dimensional Data,” *Proceedings on Privacy Enhancing Technologies*, vol. 2017, no. 4, pp. 345–364, 2017.

[47] J. Koshal, A. Nedić, and U. V. Shanbhag, “Distributed Algorithms for Aggregative Games on Graphs,” *Operation Research*, pp. 680–704, 2016. [Online]. Available: http://arxiv.org/abs/1605.00267

[48] T. Okamoto and K. Ohta, “How to Simultaneously Exchange Secrets by General Assumptions,” Tech. Rep.

[49] R. Mendes and J. P. Vilela, “Privacy-Preserving Data Mining: Methods, Metrics, and Applications,” *IEEE Access*, vol. 5, pp. 10 562–10 582, 2017.

[50] R. Cramer, I. Damgård, and J. B. Nielsen, “Multiparty computation from threshold homomorphic encryption,” *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 2045, pp. 280–300, 2001.

[51] S. Goryczka and L. Xiong, “A Comprehensive Comparison of Multiparty Secure Additions with Differential Privacy,” *IEEE Transactions on Dependable and Secure Computing*, vol. 14, no. 5, pp. 463–477, 2017.

[52] C. Clifton, M. Kantarcioglu, J. Vaidya, X. Lin, and M. Y. Zhu, “Tools for privacy preserving distributed data mining,” *ACM SIGKDD Explorations Newsletter*, vol. 4, no. 2, pp. 28–34, 2002.

[53] F. Hao, P. Y. Ryan, and P. Zieliński, “Anonymous voting by two-round public discussion,” *IET Information Security*, vol. 4, no. 2, pp. 62–67, 2010.

[54] P. McCorry, S. F. Shahandashti, and F. Hao, “A smart contract for boardroom voting with maximum voter privacy,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2017.

[55] F. Hao and P. Zieliński, “A 2-round anonymous Veto protocol,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2009.

[56] R. Gennaro, Y. Ishai, E. Kushilevitz, and T. Rabin, “On 2-round secure multiparty computation,” in *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2002.

[57] T. Afjei, U. Schonhardt, C. Wemhöner, M. Erb, H. R. Gabathuler, H. Mayer, G. Zweifel, M. Achermann, R. von Euw, and U. Stöckli, “Standardschaltungen für Kleinwärmepumpenanlagen Teil 2: Grundlagen und Computersimulationen. Schlussbericht,” Tech. Rep., 2002.

[58] L. Girardin, “A GIS-based Methodology for the Evaluation of Integrated Energy Systems in Urban Area,” p. 218, 2012. [Online]. Available: http://infoscience.epfl.ch/record/170535

[59] K. Arendt, M. Jradi, H. R. Shaker, C. T. Veje, and S. Denmark, “Comparative Analysis of white- , gray- and black-BOX models for thermal simulation of indoor environment: teaching building case study,” in *Building Performance Modeling Conference*, 2018, pp. 173–180.

[60] G. Reynders, J. Diriken, and D. Saelens, “Quality of grey-box models and identified parameters as function of the accuracy of input and observation signals,” *Energy and Buildings*, vol. 82, pp. 263–274, 2014. [Online]. Available: http://dx.doi.org/10.1016/j.enbuild.2014.07.025

[61] P. Bacher and H. Madsen, “Identifying suitable models for the heat dynamics of buildings,” *Energy and Buildings*, vol. 43, no. 7, pp. 1511–1522, 7 2011. [Online]. Available: http://linkinghub.elsevier.com/retrieve/pii/S0378778811000491

[62] F. P. Incropera, D. P. DeWitt, T. L. Bergman, and A. S. Lavine, *Fundamentals of Heat and Mass Transfer*, 2007.

[63] T. Cholewa, M. Rosiński, Z. Spik, M. R. Dudzińska, and A. Siuta-Olcha, “On the heat transfer coefficients between heated/cooled radiant floor and room,” *Energy and Buildings*, vol. 66, pp. 599–606, 2013.

[64] P. Wallenten, “Heat Transfer Coefficients in a Full Scale Room With and Without Furniture,” *Lund Institute of Thechnology*, pp. 1–8, 1999.

[65] T. Schütz, H. Harb, R. Streblow, and D. Müller, “Comparison of models for thermal energy storage units and heat pumps in mixed integer linear programming,” 2015.

[66] R. D. Coninck, R. Baetens, D. Saelens, A. Woyte, L. Helsen, A. Mechanics, and B. P. Section, “Rule-based demand side management of domestic hot water production with heat pumps in zero energy neighbourhoods,” *Journal of Building Performance Simulation*, 2013.

[67] J. S. Stein, “The photovoltaic Performance Modeling Collaborative (PVPMC),” *Conference Record of the IEEE Photovoltaic Specialists Conference*, pp. 3048–3052, 2012.

[68] L. Nespoli, “MODEL BASED FORECASTING FOR DEMAND RESPONSE STRATEGIES,” Ph.D. dissertation, 2019.

[69] 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

[70] C. C. Holt, “Forecasting seasonals and trends by exponentially weighted moving averages,” *International Journal of Forecasting*, 2004.

[71] E. S. Gardner, “Exponential smoothing: The state of the art,” *Journal of Forecasting*, 1985.

[72] H. Zhou, W. Qian, and Y. Yang, “Tweedie gradient boosting for extremely unbalanced zero-inflated data,” *Communications in Statistics: Simulation and Computation*, 2020.

[73] C. M. Bishop, *Pattern Recognition and Machine Learning*, 2013, vol. 53, no. 9.

[74] “OAEI 734.71, https://www.admin.ch/opc/it/classified-compilation/20071266/index.html.”