




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 and Cleantech

Final report

LIC - Lugaggia Innovation Community





University of Applied Sciences and Arts
of Southern Switzerland

SUPSI

Date: 8.8.2022

Place: Bern

Publisher:

Swiss Federal Office of Energy SFOE
Research Programme, Pilot and Demonstration Programme
CH-3003 Bern
www.bfe.admin.ch
energieforschung@bfe.admin.ch

Co-financed by:

Fondo Energie Rinnovabili (FER)
Repubblica e Cantone Ticino
<https://www4.ti.ch/generale/piano-energetico-cantonale/fondo-energie-rinnovabili-fer/fondo-energie-rinnovabili-fer/>

Ente e Agenzia Regionale per lo Sviluppo del Luganese
Via Cantonale 10, 6942 Savosa
<https://ersl.ch/>

Author:

Vasco Medici, ISAAC, SUPSI (vasco.medici@supsi.ch)
Lorenzo Nespoli, ISAAC, SUPSI (lorenzo.nespoli@supsi.ch)
Davide Strepparava, ISAAC, SUPSI (davide.strepparava@supsi.ch)
Francesca Cellina, ISAAC, SUPSI (francesca.cellina@supsi.ch)
Matteo Salani, IDSIA, SUPSI (matteo.salani@supsi.ch)
Marco Derboni, IDSIA, SUPSI (marco.derboni@supsi.ch)
Davide Rivola, Hive Power SA (davide.rivola@hivepower.tech)
Daniele Farrace, AEM SA (dfarrace@aemsa.ch)
Roman Rudel, ISAAC, SUPSI (roman.rudel@supsi.ch)
SFOE head of domain: Karin Söderström, karin.soederstroem@bfe.admin.ch
SFOE programme manager: Dr. Michael Moser, michael.moser@bfe.admin.ch
SFOE contract number: SI/501840

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



Summary

The municipality of Capriasca and the local DSO Azienda Elettrica di Massagno (AEM SA) have promoted the creation of a self-consumption community in the commune of Lugaggia. The Lugaggia Innovation Community (LIC) consists of 18 houses, five of which have roof-mounted photovoltaic systems with a total rated power of 45 kWp, a kindergarten with a 27 kWp photovoltaic array, and a neighborhood battery with a capacity of 60 kWh. For AEM SA, the purpose of the trial was to test and verify its ability to provide new energy services to its customers by leveraging two innovative technical solutions provided by Swiss companies Optimatik AG and Hive Power SA. The first solution consists of a centralized energy management platform that uses the existing smart meter infrastructure for sensing and actuation. The second solution implements a decentralized control approach and a secure, privacy-preserving internal billing system secured by blockchain technology. This report summarizes the results of field experimentation, the stakeholder inquiry, and analysis conducted in simulation, aimed at verifying the technical and economic feasibility of the proposed solutions, and provides insights into the evolution of energy communities in Switzerland.

Zusammenfassung

Die Gemeinde Capriasca und der lokale Stromversorger Azienda Elettrica di Massagno (AEM SA) haben gemeinsam einen Zusammenschluss zur Eigenversorgung in der Gemeinde Lugaggia gegründet. Die Lugaggia Innovation Community (LIC) besteht aus 18 Häusern, von denen fünf über Photovoltaikanlagen auf dem Dach mit einer Gesamtleistung von 45 kWp verfügen, einem Kindergarten mit einer Photovoltaikanlage von 27 kWp und einer Gemeinschaftsbatterie mit einer Kapazität von 60 kWh. AEM SA bezweckte mit diesem Projekt, Angebote neuer Energiedienstleistungen zu prüfen, indem es zwei innovative technische Lösungen nutzte, die von den Schweizer Unternehmen Optimatik AG und Hive Power SA entwickelt wurden. Die erste Lösung besteht aus einer zentralen Plattform für das Energiemanagement der LIC, welche die bestehende Smart-Meter-Infrastruktur für die Erfassung der Daten und Steuerung der Wärmepumpen, Boiler und Batterie nutzt. Die zweite Lösung beruht auf einem dezentralen Steuerungsansatz und einem sicheren, den Datenschutz garantierendem internes Abrechnungssystem, das mittels der Blockchain-Technologie abgesichert ist. Der vorliegende Bericht fasst die Ergebnisse der Feldversuche, der Befragung der Stakeholder und der Simulationsanalysen zusammen. Mit Studie überprüft die technische und wirtschaftliche Machbarkeit der vorgeschlagenen Lösungen und zeigt Möglichkeiten für die Entwicklung von Zusammenschlüssen zum Eigenverbrauch in der Schweiz auf.

Résumé

La municipalité de Capriasca et le GRD local Azienda Elettrica di Massagno (AEM SA) ont encouragé la création d'une communauté d'autoconsommation dans la commune de Lugaggia. La communauté d'innovation de Lugaggia (LIC) se compose de 18 maisons, dont cinq sont équipées de systèmes photovoltaïques montés sur le toit d'une puissance nominale totale de 45 kWp, d'une école enfantine dotée d'une installation photovoltaïque de 27 kWp et d'une batterie de quartier d'une capacité de 60 kWh. Pour AEM SA, l'objectif était de tester et de vérifier sa capacité à fournir de nouveaux services énergétiques à ses clients en s'appuyant sur deux solutions techniques innovantes fournies par les sociétés suisses Optimatik AG et Hive Power SA. La première solution consiste en une plateforme de gestion centralisée de l'énergie qui utilise l'infrastructure existante des compteurs intelligents pour la détection et l'actionnement. La seconde solution met en œuvre une approche de contrôle décentralisée et un système de facturation interne sécurisé, préservant la vie privée en utilisant la technologie blockchain. Ce rapport résume les résultats de l'expérimentation sur le terrain, l'enquête auprès des parties prenantes



et l'analyse menée en simulation, visant à vérifier la faisabilité technique et économique des solutions proposées, et donne un aperçu de l'évolution des communautés énergétiques en Suisse

Sommario

Il comune di Capriasca e l'azienda elettrica di Massagno (AEM SA) hanno promosso la creazione di una comunità di autoconsumo nella frazione di Lugaggia. La Lugaggia Innovation Community (LIC) è composta da 18 case, di cui 5 sono dotate di impianti fotovoltaici installati sul tetto, per una potenza nominale totale di 45 kWp, una scuola dell'infanzia, con un impianto fotovoltaico da 27 kWp ed una batteria di quartiere con una capacità di 60 kWh. Per il gestore di rete AEM, lo scopo della sperimentazione era quello di testare e verificare la propria capacità di fornire nuovi servizi energetici ai propri clienti, facendo leva su due soluzioni tecniche innovative fornite dalle società svizzere Optimatik AG e Hive Power SA. La prima soluzione consiste in una piattaforma centralizzata di gestione dell'energia, che utilizza l'infrastruttura esistente dei contatori intelligenti per il rilevamento e l'attivazione. La seconda soluzione implementa un approccio di controllo decentralizzato ed un sistema di fatturazione interna sicuro ed attento alla privacy, garantito dalla tecnologia blockchain. Questo rapporto riassume i risultati della sperimentazione sul campo, dell'indagine sugli stakeholder e delle analisi svolte in simulazione, volte a verificare la fattibilità tecnica ed economica delle soluzioni proposte, e fornisce degli spunti di discussione sull'evoluzione delle comunità energetiche in Svizzera.



Contents

List of abbreviations	9
1 Introduction	10
1.1 Background information and current situation	10
1.2 Purpose of the project	11
2 Lab environment	11
2.1 Hardware setup	12
2.1.1 Landis+Gyr E450	12
2.1.2 Strato	12
2.1.3 NUC	12
2.2 Software setup	13
2.2.1 Meter optical reader	14
2.2.2 InfluxDB	14
2.2.3 Grafana	14
3 Pilot environment	15
3.1 Stakeholders onboarding	15
3.1.1 Legal structure of the community	15
3.1.2 Stakeholders' acknowledgment	16
3.1.3 Contracts submission to the stakeholders	16
3.2 Field configuration	16
3.3 Pricing scheme	16
3.4 Web portal	20
3.5 Centralized data management framework	20
3.6 Decentralized data management framework	21
3.6.1 Hardware setup	22
3.6.2 Software setup	23
4 Centralized management through smart meters	23
4.1 Optiflex description	24
4.2 Non-Intrusive Load monitoring	24
4.3 Estimation of global irradiance	24
4.4 Demand estimation	27
4.5 Scheduler	28
4.6 Scalability of the centralized solution	29
4.6.1 Computational time	29
4.6.2 Data requirements	29
4.7 Extensions for near real-time DSM	30
4.8 Report on self-consumption via battery control	31
4.9 Report on peak shaving via DSM	35
5 Decentralized management	38



5.1	Algorithms description	38
5.1.1	Battery control algorithm	38
5.1.2	Boiler algorithm	40
5.1.3	Coordination	41
5.2	Report on battery control	42
5.3	Report on boiler control	43
5.3.1	Local disaggregation	44
5.3.2	Control performance	45
5.4	Decentralized management through blockchain-enabled smart meters	49
5.4.1	1 st layer application: asset manager (AM)	50
5.4.2	2 nd layer solutions: state channels	52
5.4.3	2 nd layer solutions: sidechains	53
5.4.4	Cosmos and Tendermint	53
5.4.5	Applications deployed in LIC: <i>Metering</i>	56
5.4.6	Applications deployed in LIC: <i>Auditable Tariffs (AT)</i>	57
5.4.7	Applications deployed in LIC: <i>Pre-paid scenario (PS)</i>	60
5.4.8	Applications deployed in LIC: <i>PARITY Market (PM)</i>	60
5.5	Application sustainability	64
5.5.1	Hardware requirements	64
5.5.2	Energy consumption and greenhouse gas emission evaluation	65
6	Techno-economic performance analysis	66
6.1	PV and distributed batteries	66
6.2	District battery	72
6.3	Flexible loads	78
6.4	Coordination	78
6.4.1	Batteries	79
6.4.2	Water heaters	80
6.5	Grid support potential	81
6.6	Delay investments in grid refurbishment through storage and DSM	81
6.6.1	Numerical results	83
7	Stakeholders, policy and legal evaluation	86
7.1	Methodology	86
7.1.1	Questionnaire administration and response rate	87
7.1.2	Limitations of the research design	87
7.2	Characteristics of LIC households and houses	88
7.3	Level of engagement and satisfaction with the LIC community (RQ1)	88
7.3.1	Attitudes towards self-consumption communities	88
7.3.2	Level of engagement and satisfaction with LIC web portal	90
7.4	Factors affecting attitudes towards the LIC community (RQ2)	93
7.4.1	Community identity	94
7.4.2	Social norms	94
7.4.3	General trust	95
7.4.4	Attitude towards renewable energies	95
7.5	Perceived effects of membership to the LIC community (RQ3)	96
7.6	Effect on attitudes towards renewable energies and environmental behaviour (RQ4)	98



7.7 Policy and regulatory framework	102
7.8 The evolution of the LIC self-consumption community since its start	103
8 Conclusions and recommendations	104
8.1 The perspective of the pilot owner	104
8.2 Centralized vs decentralized control	104
8.2.1 Thermal loads	104
8.2.2 Battery	105
8.3 The Evolution of Policy and Regulatory Framework	106
8.4 The present and the future of LIC	106
9 Publications	108
10 References	109



List of abbreviations

AEM Azienda elettrica di Massagno SA

AM Asset manager

AMM Automated market making

AT Auditable tariff

BaU Business as usual

BFE Bundesamt für Energie

DSO Distribution system operator

EC Energy community

EVM Ethereum virtual machine

GDPR General data protection regulation

GHI Global horizontal irradiance

HES head-end system

KPI Key performance indicator

LCOE Levelized cost of electricity

LEM Local energy market

LIC Lugaggia Innovation Community

MAPE Mean absolute percent error

MPC Model predictive control

NILM Non-intrusive load monitoring

P2P Peer to peer

PLC Power line communication

PM Parity Market

PS Prepaid scenario

PV Photovoltaic

RMSE Root mean squared error

RQ Research question



SCC	Self-consumption community
SLA	Service level agreement
SOC	State of charge
TMY	Typical meteorological year
WACC	Weighted average cost of capital
ZEV	Zusammenschluss zum Eigenverbrauch



1 Introduction

1.1 Background information and current situation

In Switzerland, solar generation is reaching considerable levels of penetration, which is expected to rise further thanks to a wide range of environmental, social, technical and economic drivers (not least thanks to the “Energy Strategy 2050”). From this growth, a range of issues related to the operation and economic impacts of photovoltaic (PV) generation on the power grid arises. From a technical point of view, the stochastic nature of solar production causes operational challenges. Among them, the unbalance between production and consumption, overvoltage and overload of grid components are the most common ones. From the economic point of view, the increase in self-generation tends to reduce the turnover of distribution system operators (DSO), considering that the grid component of the electricity tariff is usually a function of the consumed energy. On the other hand, due to the above-mentioned technical problems, the investments in the network infrastructure are expected to increase. As a consequence, in Switzerland grid tariffs are likely to increase and with them the social disparities between people who can afford a PV plant and those who do not.

To restore fairness, one solution would be to redesign the grid tariffs. As the grid is dimensioned as a function of the maximum power demand (or supply) and not of the total energy consumed, a more appropriate design of a grid tariff could result in having at least part of the costs proportional to power (or its peak like it is already usually done for big consumers). Following the SUPSI experience with decentralized energy management achieved through the project GridSense, the novel tariff schemes are resulting from the partners’ expertise from previous and ongoing projects:

- NEMoGrid (ERA-Net Smart Grids Plus grant agreement No 646039) regarding the design and the evaluation of new business models, favouring the grid-integration of decentralized energy resources
- MuLDeR (Research Program Grids BFE), regarding the design of mechanisms that allow offering demand-response services to all actors involved and at all grid levels, while making sure that the control actions are not violating grid constraints and guaranteeing the power quality.

However, this does not solve the problem of the increasing grid infrastructural costs. Among the technical measures for mastering this challenge, the intelligent management of the flexibility available at the demand side is recognized as a promising approach to relieving the network stress. In Switzerland, most of the grid issues related to PV penetration are located in the low voltage levels of distribution grids, close to the demand side. To avoid creating unbalances in the distribution grid and to improve the grid energy efficiency, demand side management, alongside with local storage, could be used to realign consumption and production. The first and most direct way to align consumption with production is the optimization of self-consumption, directly at the point of PV generation. At a household level, this can be achieved by shifting flexible loads like electric boilers, heat pumps and electric vehicle chargers, as much as possible to periods of abundant solar generation. Additionally, electric storage can be installed locally. Smart self-consumption solutions can also steer consumption towards the periods of maximum generation, in order to lower the peaks of injected power.

The new energy ordinance, in force since January 1, 2018, allows the establishment of self-consumption communities (SCC), in which energy can be exchanged internally between their members, without being subject to grid tariffs and grid taxes. SCCs are an efficient solution for the reduction of grid issues, even if the local grid still has to be maintained with some costs. By increasing the number of loads having access to local generation (and possibly storage) raises the self-consumption potential, allows for a finer control of power injection and consumption and lowers the Levelized Cost of Energy (LCOE) of PV and storage, making them more attractive.

The new law also introduces other important changes and opportunities. The DSO is allowed to install intelligent control and regulation systems at final consumers or producers. And, if a smart meter is installed, the DSO can propose new tariffs as an alternative to the standard ones, which are based



on energy for at least 70%. Therefore, the DSO could propose a tariff to a SCC which incentivizes peak shaving and will help him keeping the grid under nominal operating conditions. It could also offer flexibility control services to the SCC.

1.2 Purpose of the project

The municipality of Capriasca installed a 30 kWp PV plant in the village of Lugaggia on the roof of the local kindergarten. The building is located on the edge of a residential area, mainly consisting of single-family houses. The self-consumption potential of the kindergarten is limited because most of the production takes place during school summer holidays when the local consumption is low. AEM, the DSO serving the area, therefore promoted the creation of a SCC named Lugaggia Innovation Community (LIC), connecting together the kindergarten and ten nearby houses. Differently from other research projects, such as "Quartierstrom", the energy exchange inside the community is compliant with existing laws regulating the Self Consumption Communities. By creating the SCC, AEM aims at testing and verifying its capability to provide new energy services to its customers, by leveraging on two novel technical solutions provided by the Swiss companies Optimatik and Hive Power:

- The first solution consists of a centralized energy management platform, which uses the existing smart meter infrastructure for sensing and actuation
- The second solution implements a decentralized control approach secured by blockchain technology and requires the installation of computing and controlling unit, connected to the smart meters via DLMS interface.

To further increase the flexibility in the SCC, AEM installed a district-level storage system.

The project aims to:

1. Evaluate the needs and requirements to the realization of LIC in a real environment. The project aims to provide recommendations how to allow and facilitate the replicability and scalability of peer-to-peer self-consumption communities. In particular, with respect to the needs of the public interest, to ensure fair treatment of all stakeholders (especially in areal situations) and implement measures for the correct use of energy resources (avoid sporadic tips, excessive consumption, control of equipment with the mandatory announcement, etc.)
2. Assess blockchain as a decentralized billing management method introduced by the utility
3. Compare centralized vs decentralized load management methods from the DSO point of view (grid costs), energy consumption and economic point of view
4. Help to assess the local flexibility potential and the different ways in which it could be exploited from a technical point of view
5. Evaluate the degree of knowledge or acceptance among the community stakeholders to be willing to participate in these new self-consumption communities; a living lab to test users' acceptance was set up.

2 Lab environment

Prior to field installation of the hardware needed for decentralized load control and also to begin to familiarize with the meters installed by AEM, a laboratory test environment was set up. The hardware and software of the test environment are described below. The same combination of hardware and software was then used in the pilot project.



2.1 Hardware setup

2.1.1 Landis+Gyr E450

E450, produced by Landis+Gyr¹, is an industrial and widely used device in the field of residential smart metering. It has been exhaustively tested in the laboratory due to its installation in the prosumer constituting the LIC demonstrator. E450 provides a notable collection of signals that can be accessed for monitoring. Besides, two relays of the meter can be remotely actuated. It is important to remark that E450 provides an optical serial interface that can be used to continuously read data and send command using the open DLMS protocol².



Figure 1: E450 smart meter

2.1.2 Strato

The Strato Pi CM board³, shown in Figure 2, is based on a Raspberry CM platform⁴. The Raspberry CM combines the computational power and 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 restarted the operating system without problems. Table 2 summarizes the main features of the Strato devices.

2.1.3 NUC

Strato are embedded devices that can be used in a wide variety of applications. However, they are based on ARM architecture [1] and, consequently, they can not manage specific applications that require too

¹<https://www.landisgyr.it/product/landisgyr-e450/>

²<https://www.dlms.com/>

³<https://www.sferalabs.cc/strato-pi/>

⁴<https://www.raspberrypi.org/products/compute-module-3>



CPU	4 ARMv8 64-bit 1.2GHz
RAM	1 GB
DISK	32 GB
USB ports	2
Connectivity	1 Ethernet port
OS	Raspbian GNU/Linux 9 (Stretch)

Table 2: Strato main features



Figure 2: Strato device

significant hardware resources. In order to manage these peculiar cases, the NUC board ⁵ was taken into account. It is an embedded solution provided by Intel®, more powerful and robust than the Strato. For example, in the LIC pilot for details) the NUC acts like an aggregator in the community running specific control applications that cannot be managed by the Strato devices. Table 2 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 3: NUC main features

2.2 Software setup

The present section describes the main software installed on the Strato and NUC devices. It is important to remark how these applications are not the only ones deployed on the boards. Fundamentally, they are needed by the applications described in sections 3.5 and 3.6 to operate correctly. Thus, they can be considered as a software layer between the smart meters and the developed applications, which can be installed locally on the Strato or remotely on other machines. Figure 3 reports the interactions between the main software running on a Strato, which are described in the following of this section, and a custom application, like the ones described in sections 3.5 and 3.6.

⁵<https://www.intel.com/content/www/us/en/products/details/nuc.html>




Figure 3: Main applications running on Strato

2.2.1 Meter optical reader

Meter Optical Reader is a custom Python application aiming to continuously provide a bidirectional interface with the smart meter. The hardware interface between the Strato and the meter is provided by a USB optical reader. Thanks to it, the software can gather electrical signals (e.g. active and reactive power, current, voltage, etc.) from the smart meter and send command to its relays. The data acquired by *Meter Optical Reader* are saved locally in an InfluxDB database⁶, which provides a local data source that can be used by the applications described chapters 3.5 and 3.6. In addition, *Meter Optical Reader* is configured to save the datasets also on a remote InfluxDB database server. This dual data saving is due to be able to manage both a centralized control algorithm based on a central database and a decentralized approach, please refer to sections 3.5 and 3.6.

The software manages to collect an entire dataset, comprehensive of 27 signals, with a remarkable time resolution of approximately 5 seconds. Regarding the commands, *Meter Optical Reader* provides a REST API that can be used by authorized applications to actuate the meter relays. The data exchanging is based on the DLMS protocol [2] and exploits the open-source Gurux library [3]. Table 4 reports the signals currently collected and the commands with the related DLMS codes.

2.2.2 InfluxDB

As explained in Section 2.2.1, *Meter Optical Reader* needs a local database server in order to locally save the collected measurements. This data management is provided by the installation of an InfluxDB server on each Strato. It is a time-series database optimized for the management of time-series datasets. Besides, InfluxDB can easily run on boards based on ARM architecture like Raspberry and, consequently, also Strato. Thus, it is used by *Meter Optical Reader* to store the collected data. However, the resource requirements of InfluxDB on Strato were exhaustively checked and no problem was encountered in terms of CPU, RAM and disk usage.

2.2.3 Grafana

In order to monitor the data stored in InfluxDB by *Meter Optical Reader*, an installation of Grafana⁷ has been installed on each Strato. Similarly to InfluxDB, its behaviour was tested to control if its usage was heavy for the operating system. As in the case of InfluxDB, no problems were encountered.

⁶<https://www.influxdata.com/>

⁷<https://grafana.com>



V1	'1.0.32.7.0.255'
V2	'1.0.52.7.0.255'
V3	'1.0.72.7.0.255'
I1	'1.0.31.7.0.255'
I2	'1.0.51.7.0.255'
I3	'1.0.71.7.0.255'
IN	'1.0.91.7.0.255'
ITot	'1.0.90.7.0.255'
freq	'1.0.14.7.0.255'
PImp	'1.0.1.7.0.255'
PExp	'1.0.2.7.0.255'
QImp	'1.0.3.7.0.255'
QExp	'1.0.4.7.0.255'
S	'1.0.9.7.0.255'
PF	'1.0.13.7.0.255'
AUL1UL2	'1.0.81.7.10.255'
AUL1UL3	'1.0.81.7.20.255'
AUL2UL3	'1.0.81.7.21.255'
AUL1IL1	'1.0.81.7.40.255'
AUL2IL2	'1.0.81.7.51.255'
AUL3IL3	'1.0.81.7.62.255'
EPImp	'1.1.1.8.0.255'
EPExp	'1.1.2.8.0.255'
EQImp	'1.1.3.8.0.255'
EQExp	'1.1.4.8.0.255'
Relay1	'0.1.96.3.10.255'
Relay2	'0.2.96.3.10.255'

Table 4: Signals acquired by *Meter Optical Reader*

3 Pilot environment

3.1 Stakeholders onboarding

We deployed a 3 steps procedure for setting up the LIC Community:

- Legal structure of the community
- Stakeholders' information
- Contracts submission to the stakeholders

3.1.1 Legal structure of the community

The legal structure has been setup as a simple partnership managed by AEM, the local DSO, which is still owning the low voltage cable supplying the district, although it has been put out of the grid solidarity. The LIC simple partnership signed a supply contract and a grid contract with AEM on one hand, and with all the end users/producers inside the district on the other one. At the end of the pilot it is under consideration to create a benevolent association ("Eingetragener Verein") among the buildings' owners and AEM.



3.1.2 Stakeholders' acknowledgment

The stakeholders' acknowledgment happened in two intermediary steps. First of all, we did send to all the LIC's end users a leaflet explaining LIC's goals, its way of working and highlighting end users' rights and constraints. Secondly, we did organise a general meeting gathering all the stakeholders in order to have the time to present and discuss LIC activities and end users' implication. This meeting was well attended and enjoyed the support of the Municipality (which is also involved in LIC through the Kindergarten), which played an important role for boosting trust among the participants.

3.1.3 Contracts submission to the stakeholders

Finally, we did meet each candidate in a 1:1 individual meeting, which led to the contract signing. It is important to emphasize that all the LIC end users decided to opt in and sign the contracts. In the future, for setting up new "self-consumption districts" inside AEM supply territory, based on LIC experience and results dissemination among AEM's end users, we will need to simplify this procedure which has been quite time consuming (an investment which is hard to be paid back through the Community income).

Nevertheless, we have been positively surprised by the cooperative attitude of the end users (as previously said all of them signed the contracts and are participating to the pilot project) and their attitude to start thinking as a community and less as single users. For instance, a building owner decided to put its PV program on hold because he would like to be sure not to impact negatively on the self-consumption balance inside the community. Similarly, other end users offered us their rooftop, if the construction of additional PV capacity would help in increasing resilience and autonomy of the LIC.

3.2 Field configuration

The LIC pilot⁸ is located in Lugaggia, a small village near Lugano. In a part of the municipality, a self-consumption community has been created with the collaboration of AEM⁹, the local DSO. The self-consumption community consists of 18 residential houses and a kindergarten. Figure 4 shows the pilot map.

LIC community is composed by the two blue polygons reported in Figure 4, representing the kindergarten at the center top and the LIC houses at the bottom left. Three houses, represented by a sun icon in Figure 4, are equipped with photovoltaic systems installed on the roof, for a total nominal power of 33 kWp. Besides, a 27 kWp photovoltaic plant and a battery with a capacity of 60 kWh are installed in the kindergarten.

3.3 Pricing scheme

In the decentralized setting, the end users are organized in an energy community (EC). The goal of the community is to maximize its welfare, by reducing the costs for the consumers and increasing the revenues of producers. At the same time, we want to formulate a market, meeting the following conditions:

- A fair redistribution of money among the market players, according to their contribution to the market's intended outcome, must be guaranteed
- The market must induce a variance reduction in the aggregated power profile
- The market must be compatible with the current legal energy billing framework, considering both produced and consumed energy in a given timeslot

⁸<https://lic.energy/>

⁹<http://aemsa.ch>




Figure 4: Lugaggia pilot

- The market must induce an increase of self-consumption at community level, while steering the overall power profile at will of third parties.

In order to achieve these points, we set up an automated market making (AMM) mechanism [4, 5]; this is defined by a set of simple and interpretable price formation 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
- 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
- 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
- The difference between the community buying and selling prices covers the cost to setup, operate and maintain the community infrastructure.



These AMM rules can be mathematically expressed as:

$$p_b = [E_c p_b^{BaU} - \min(E_c, E_p) (p_b^{BaU} p_b^{P2P})] / E_c \quad (1)$$

$$p_s = [E_p p_s^{BaU} - \min(E_c, E_p) (p_s^{P2P} - p_s^{BaU})] / E_p \quad (2)$$

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. Peers clearly profit from the difference in price between BaU and community, but the third party 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 third party interest to maximize self-consumption (no conflicting interests between peers and community admin). The AMM mechanism dictate the price formation inside the community. The prices as a function of the consumed and produced energy inside the EC can be seen in Figure 5. It can be shown that these prices generate convex costs as a function of the agents' actions, and thus are amenable to be jointly optimized in a distributed way, as in [6, 7].




Figure 5: Buying (upper surface) and selling price (lower surface) p_b and p_s as a function of the produced and consumed energy inside the EC.

Let's consider the cost of the i th agent of consuming energy inside the EC. If the agent changes its consumption e_i , it directly influences p_b since its consumption is included in E_c . Due to the presence of the min operator, we must split the p_b expression in two cases to study its convexity, namely the case in which E_p is bigger or smaller than E_c . In the first case, the EC is a net energy producer, and the i th agent's total costs can be expressed as (simplifying the above expression for p_b):

$$c_i = p_b e_i = p_b^{P2P} e_i \quad (3)$$

Which is linear; in the second case, the EC is a net energy consumer, and the expression of p_b reduces to:

$$c_i = p_b e_i = p_b^{BaU} e_i - E_{p,0} \frac{p_b^{BaU} - p_b^{P2P}}{E_{c,0} + e_i} \quad (4)$$



which is convex in e_i . Here $E_{c,0}$ is the initial consumed energy of the EC, before the influence of the i th agent; similarly, $E_{p,0}$ is the produced energy before the production of the i th agent (in this case fixed at 0). Since the i th agent can switch the EC from being a net energy importer to being a net energy exporter, the two expression must be combined to study the overall convexity of the i th agent costs.

Figure 6 shows the combination of the two expression. We fixed $E_{c,0}$ and $E_{p,0}$ to 1 and 5 kWh respectively, and spanned the consumption e_i of the i th agent from 0 to 10 kWh, so that the EC passes from a net energy producer to a net energy consumer in the graph. The blue line shows the linear price which generates in the case the EC was a net energy producer, the orange one the one if the EC was a net energy consumer, and the dashed green line is the true cost. As the true cost is the maximum of two convex expressions, it is also convex due to convexity rule of composite convex functions. A similar reasoning can be done for the case in which the agent is a producer, and we can reach to the same conclusion. Thus, the costs function is convex with respect to the agents' actions.




Figure 6: buying cost (in green) as a composition of the two sub-cases in which the EC is a net energy importer (orange) or net energy exporter (blue). The shape of the lines depend on the fixed values of $E_{c,0}$ and $E_{p,0}$, but the final expression is always convex.

The selling price of locally produced energy inside the community was set by AEM at 9cts/kWh. Then, AEM conducted a cost analysis of the distribution network for the community's internal network and, for purchase prices and configured the pricing scheme using the parameters p_b^{BaU} and p_b^{P2P} shown in Table 5:



	p_b^{BaU}	p_b^{P2P}
Grid tariff	7.3	4.5
Energy	8.1	9
Taxes	5.6	0
Admin costs	0	2.5
Total	21	16

Table 5: Buying prices (in cts/kWh) in the business as usual (BaU) and self-consumption community (P2P).

Administration costs were calculated so as to redistribute all income from community establishment to community users, but without incurring financial losses. Administration costs also cover for the losses inside the community's internal distribution grid. The amortization of the battery costs was not considered while choosing the internal prices for the energy community, that is, the battery is considered to be a sunk cost undertaken by the community administrator. As demonstrated in section 6.2 (table 14), under the considered revenue stream and current battery prices (just due to self-consumption optimization), the installation of an electric battery is not economically sustainable, that is, in order to cover for its costs the administrator should increase internal prices above the external ones; at this point wouldn't make sense for the users to join the energy community.

Summarizing, the pricing scheme is configured with the parameters for p_b^{BaU} , p_s^{BaU} , p_b^{P2P} and p_s^{P2P} shown in Table 6.

	outside the community	inside the community
buying	$p_b^{BaU} = 21$	$p_b^{P2P} = 16$
selling	$p_s^{BaU} = 6$	$p_s^{P2P} = 9$

Table 6: Energy prices (in cts/kWh) applied in the evaluation of economic impact.

3.4 Web portal

Hive Power has developed and made available to users a web interface (Figure 7) on which they can view:

- The energy consumption and injection at your home's connection point
- The energy consumption and injection at the virtual community connection point, which is the sum of all the meters' measurements
- The power of the battery and its state of charge
- The trend of prices in the domestic energy market.

The web portal allows users to be aware of their consumption and production and gives them a tool they can use to shift other flexible loads at times of the day when energy prices are lower. The web portal has been integrated into the social engagement part of the work, presented in Section 7.

3.5 Centralized data management framework

The Optiflex project, described in Section 4, aims to develop a centralized control solution fully based on the existing smart metering infrastructure, without the need for additional hardware. Consequently, a solution allowing to retrieve the data from the smart meters and send commands to them was developed.




Figure 7: Web portal

Figure 8 depicts the data flow of the centralized system, which is based on two main components: the smart metering head-end system (HES) and the Big Data platform Kibid provided by Kisters AG¹⁰. The HES provides the low-level interface with the Landis+Gyr smart meters. Normally, the smart meters are read using PLC via a data concentrator, which is represented by the light blue rectangle in Figure 8. The meters that required to be read at a higher frequency (e.g. the coupling point of the community) are directly connected to the HES using fiber.

The data that are continuously sent to the HES by the meters are stored in a time-series database within Kibid. It is important to remark that this centralized solution was not developed only for the LIC pilot but for the entire area managed by AEM in two projects funded by Innosuisse (26756.1 PFES-ES and 43383.1 IP-EE). Currently, data concentrators collect data of about 9'000 smart meters. Both HES and Kibid are installed on-premise on a set of machines to assure a high level of robustness and redundancy. The data exchange between HES and Kibid shown in Figure 8 is provided using the IEC61968-9 protocol [8]. Third-party applications can access the datasets stored in Kibid using libraries developed by Kisters and based on REST APIs. Each access requires a successful login of the third-party application.

Thank to the Kibid interface, various measurements can be downloaded for each meter. Table 7 reports the one-minute resolution signals currently downloaded using the Kibid interface and the related DLMS codes.

3.6 Decentralized data management framework

The solution described in Section 3.5 is suitable for a centralized approach. On the other hand, the applications described in Section 3.6 require a specific decentralized setup in terms both of deployed hardware and of data collection software. An example of the needing of this peculiar setup is the fact that the presented framework is not able to guarantee measurements with a time resolution smaller than a minute. Besides, the data collection is not performed in real-time and some

¹⁰<https://energie.kisters.de/loesungen-produkte/kibid-big-data-data-analytics/>




Figure 8: Data flow between HES, Kibid and the smart meters

P_import	"0.2.3.4.1.1.12.0.0.0.0.0.0.0.0.38.0"
P_export	"0.2.3.4.19.1.12.0.0.0.0.0.0.0.0.38.0"
I_L1	"0.2.3.0.0.1.4.0.0.0.0.0.0.128.0.5.0"
I_L2	"0.2.3.0.0.1.4.0.0.0.0.0.0.64.0.5.0"
I_L3	"0.2.3.0.0.1.4.0.0.0.0.0.0.32.0.5.0"
PF_L1	"0.2.3.12.0.1.38.0.0.0.0.0.0.128.0.0.0"
PF_L2	"0.2.3.12.0.1.38.0.0.0.0.0.0.64.0.0.0"
PF_L3	"0.2.3.12.0.1.38.0.0.0.0.0.0.32.0.0.0"
Q_export	"0.2.3.4.19.1.12.0.0.0.0.0.0.0.0.63.0"
Q_import	"0.2.3.4.1.1.12.0.0.0.0.0.0.0.0.63.0"
V_L1	"0.2.3.0.0.1.54.0.0.0.0.0.128.0.29.0"
V_L2	"0.2.3.0.0.1.54.0.0.0.0.0.0.64.0.29.0"
V_L3	"0.2.3.0.0.1.54.0.0.0.0.0.0.32.0.29.0"

Table 7: Signals acquired through Kibid interface (1 minute data resolution)

delays can occur in the data collection. This problem can affect also the actuation of the smart meter relays, which cannot be performed directly but has to be addressed by the Kibid and HES systems. For the aforementioned reasons, an additional decentralized framework was implemented, comprehensive of both hardware and software components that are described in the following sections.

3.6.1 Hardware setup

In order to have a direct interface with the smart meters, the devices tested in laboratory (see Section 3.6.1 and 2.2 for details) were deployed in the cabinets of the LIC houses.

Basically, each LIC end-user corresponds to a node, which is associated to a single point of delivery and equipped with a smart meter. Thus, a Strato device, described in Section 2.1.2, is connected to each of the smart meters via an optical USB port. Besides, in the basement of the kindergarten, close to the central battery, a NUC machine (please refer to Section 2.1.3 for details) has been installed. This machine acts as the aggregator of the community and as an interface with the battery, as described in Section 2.1.3.

Figure 9 shows a Strato installation in a LIC 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.




Figure 9: Strato setup in a cabinet

3.6.2 Software setup

Section 3.6.1 describes the nodes installed in the LIC pilot where the applications explained in sections 2.2, 3.5 and 3.6 run. In addition to them, to have secure and easy connections among the nodes, a VPN has been deployed in the LIC pilot. Each Strato and NUC has its own VPN certificate providing a static IP address. The VPN is ruled by an OpenVPN¹¹ instance installed in a SUPSI server.

It is important to remark how this tool has extreme importance in the maintenance of the nodes installed in the pilot, such as the Strato and the NUC devices. Indeed, with OpenVPN each node can be easily and safely accessed to perform various operations, including the upgrading of the running applications, the checking of the operating system status and the interactions between the nodes.

In addition, an application acting as an interface with the community battery is continuously running on the NUC device. The interaction with the battery is guaranteed by a fast queue protocol defined by Kisters.

4 Centralized management through smart meters

The Optiflex project has been funded by two consecutive Innosuisse projects (26756.1 PFES-ES and 43383.1 IP-EE), we refer the reader to the report documentation of the projects that can be accessed using the ARAMIS platform <https://www.aramis.admin.ch/>. Some developments of the Optiflex platform have been designed specifically for LIC and, after a brief description of the platform, we focus on these developments.

¹¹<https://openvpn.net>



4.1 Optiflex description

Four organizations were involved in this project: AEM as the customer and technical requirement provider, SUPSI (responsible for algorithm development), Optimatik as the project center (including technical consulting), and KISTERS (software company from the energy industry, responsible for the development/transfer of the previous results from OptiFlex into their standard product FlexManager).

A simplified representation of the Optiflex suite of modules and their interaction is reported in "Figure 10".




Figure 10: Schematic representation of Optiflex.

The metering infrastructure feeds metering data into the system, data is processed to understand loads' behavior with the Non-Intrusive Load Monitoring (NILM) component according to past weather data. Load behavior is then used to train a demand model able to forecast future needs according to future weather forecasts. Finally, grid loads are steered by a Demand Side Management (DSM) component and the process is continuously iterated.

4.2 Non-Intrusive Load monitoring

NILM is necessary when electrical loads are not directly monitored but an estimate of their use is needed. Commonly, the smart-metering infrastructure provides the necessary measurements at the Point of Common Coupling (PCC) of each household, which are composed of active and reactive power measurements as well as voltages and currents. The purpose of NILM is to detect the activation of major controllable loads (normally heat pumps and domestic water heaters) that can be steered via the actuators installed at the smart meter and separate their power footprint from the other loads that are considered as uncontrollable.

The main issue to implement effective disaggregation algorithms is the sampling frequency. While the majority of the approaches rely on sampling rates in the range from 0.1 Hz to 1 Hz, in Optiflex we deal with much lower sampling frequencies in the order of one sample per minute.

In this context, our purpose is to control a pre-defined set of flexibilities, in particular heat pumps and electric heaters, for which we know the nominal power and the relay status with a 5 minutes resolution, in addition to the meter power profile.

Given these data, the disaggregation algorithm aims to detect whether a flexibility has been absorbing power or not during a given time interval.

4.3 Estimation of global irradiance

Specific development of Optiflex has been conceived for LIC to deal with the presence of non-monitored PV installations. The PV power production can be absorbed by local electrical loads in what is commonly referred to as *self-consumption*. In such cases, the metering infrastructure records the net power

consumption only. Figure 11 illustrates the case where some loads activate during the period when PV is producing power.



Figure 11: Meter power profile in presence of a non-monitored PV installation

When the PV installations are not monitored, the methodologies for NILM must be enriched in order to account for PV production. In particular, one could exploit global horizontal irradiance (GHI) data coming from weather services and estimate the PV power production based on each installation's nominal data (nominal ac and dc power). In order to do so, software libraries are normally used. One popular software library is PvLib, a community supported tool that provides a set of functions and classes for simulating the performance of photovoltaic energy systems [9].

In our case, accurate weather data is not available. Therefore we reconstruct an approximation of the global irradiance data by exploiting the metering data of all installations that are close to one another. The rationale behind this approach is that not all major loads are absorbing power simultaneously and the effect of PV installations is directly visible in the metering data.

Let H be the set of metering infrastructures reading neighborhood households equipped with PV installations, $p_{t,h}$ identifies the power reading at time $t \in T$ of household $h \in H$. Let GC_t be the global irradiance in clear sky conditions at time $t \in T$, computed using PvLib. Let $\hat{p}_{t,h}$ be an upper bound on the power production of the PV installation of household h computed with PvLib using the known nominal data of the installation and the global irradiance with clear sky conditions and $r_{t,h}$ be the ratio between the upper bound and the negative portion of the power reading limited by an upper bound \hat{r} (equal to 1.1 in our tests to allow for some additional freedom) and \hat{r}_t be the maximum of such ratios for $t \in T$:

$$r_{t,h} = \min\{\hat{r}, \min\{0, p_{t,h}\} / \hat{p}_{t,h}\} \quad \forall t \in T, h \in H \quad (5)$$

$$r_t = \max_{h \in H} \{r_{t,h}\} \quad \forall t \in T \quad (6)$$

Using r_t we approximate the global irradiance \tilde{I}_t as a fraction of the clear sky global irradiance GC_t and estimate the PV installation power production $\tilde{p}_{t,h}$ of household h with PvLib.

$$\tilde{I}_t = GC_t \cdot r_t \quad \forall t \in T \quad (7)$$

We recall that five photovoltaic installations are present in the LIC area. The peak powers of the installations are 30.0, 14.0, 10.0, 10.0 and 9.0 kW for a total of 73 kWp.

The rooftop of the Kindergarten, see Figure 4, is also equipped with a pyranometer for measuring solar irradiance gathering data with one minute resolution. We use this data for validation purposes only. Figure 12 shows the global irradiance for a period of two months.

We gather real data from the LIC area for the period comprised from 04.04.2020 to 03.06.2020. We collect one data point per minute. For space reasons, we summarize results as weekly averages. Table 8 reports the test weeks and the average daily measured GHI in the first two columns. Then the remaining columns report on the Weekly Average Root Mean Square Error (RMSE) and the Weekly Average Mean Absolute Percentage Error (MAPE) of the estimated GHI with respect to the measured GHI. In our tests, we use time discretization of the set T of 5 minutes leading to $|T| = 288$ time steps. Columns




Figure 12: GHI measured on the Kindergarten's rooftop, in W/m^2

marked with > 0 and $> AVG$ report values when the time steps are filtered by considering those with measured GHI larger than 0 and larger than the daily average, respectively. Indeed, several time steps are characterized with a measured GHI = 0 (during night time) and the GHI is easily estimated. This is confirmed by results in Table 8 observing that that the values in the column RMSE(>0) are always larger than those in column RMSE. Instead, by definition, MAPE is computed on time steps where the measured GHI is strictly positive. When observing time steps where the measured GHI is larger than the daily average, that is on time steps which should correspond to large PV power production, we do not observe a clear trade off in the quality of the estimation. Considering daily results, we report that MAPE($>AVG$) varies from 4.6% to 41.9%. The plot of these two cases are shown in pictures 13a and 13b. In Figure 13a we can appreciate the correct estimation of a sudden drop of the GHI around 9AM. We see that the worst cases occur when the daily avg GHI is very low (cloudy days) and the PV production is indeed less relevant. If we limit the investigation to sunny or scattered days (e.g., with GHI $> 100 W/m^2$) we observe that the worst MAPE is 21.2% (see Figure 14).

Week	Daily Avg GHI	RMSE	RMSE (>0)	RMSE ($>AVG$)	MAPE (%)	MAPE ($>AVG, %$)
Week_1	241.8	61.5	87.7	86.8	15.4	13.6
Week_2	245.8	42.7	59.4	62.2	11.0	9.6
Week_3	191.1	33.1	45.5	46.5	11.4	15.2
Week_4	159.0	39.2	55.4	59.7	12.2	12.1
Week_5	273.4	37.8	48.9	50.9	8.2	6.7
Week_6	178.8	44.7	61.5	68.4	12.1	17.1
Week_7	248.3	39.2	52.1	56.6	10.5	9.3
Week_8	288.3	44.1	59.6	62.5	12.7	9.2
Week_9	257.6	50.3	67.2	74.8	12.8	11.0

Table 8: Weekly average RMSE and MAPE comparing estimated GHI and measured GHI, from 04.04.2020 to 03.06.2020

While the estimation quality may seem particularly bad on some days, the outcome of PV production estimation is sufficient for the purposes of non-intrusive load monitoring. We were not able to perform the validation on the real data (i.e. coming from the field tests in LIC) because the PV production is not monitored nor the physical consumption of the single devices, there is only one meter for the entire building. Instead, we tested the method of GHI estimation in simulation and results are reported in the publication [10] we observe that despite the relative large errors in GHI estimation the error in disaggregation is 2.25% in average with the highest error being 9.01%




Figure 13: Measured and estimated GHI




Figure 14: Measured and estimated GHI on 29.05.2020, MAPE = 21.2%

4.4 Demand estimation

With demand estimation, we predict the time-dependent needs of a flexibility to be connected to the grid, that is the amount of time a flexibility must be allowed to drain power from the grid in different moments of the day. More formally, we aim to predict the future power usages of a flexibility $y_t, y_{t+1}, \dots, y_{N-1}$ using the output of the disaggregation algorithm that estimates past power usages $y_{t-m}, y_{t-m+1}, \dots, y_{t-1}$, where M is the number of historical data used to forecast, and N is the number of future values being forecast.

There are several forecasting methods available, namely: moving average, seasonal method with error correction, autoregressive model, autoregressive integrated moving average model, function fitting neural network, and nonlinear autoregressive neural network. In the current implementation of Optiflex, we devised a software architecture capable of using different techniques with minor implementation effort. We currently provide a demand estimation algorithm that hybridizes a seasonal method with a simple classification method.

Seasonal method: we assume that demand has daily-seasonal patterns. Thus, this method predicts the future demand values by computing statistical distributions for the values at the same time point in previous days. Let m be the period of the seasonality $l = \frac{M}{m}$ be the number of available seasonal data, and the future values.



Table 9: Classification rules

	Daily T avg \geq Monthly T avg	Daily T avg $<$ Monthly T avg
Daily I avg \geq Monthly I avg	Hot – Sunny day	Cold – Sunny day
Daily I avg $<$ Monthly I avg	Hot – Cloudy day	Cold – Cloudy day

As an example, average values can be then computed by:

$$\hat{y}'_{t+i} = \frac{\sum_{j=1}^l y_{t+i-jm}}{l}, \quad i = 0, 1, \dots, N-1 \quad (8)$$

Classification method: the output of the disaggregation is a data sample that is used to calibrate a seasonal model, we hybridize the seasonal method with a simple classification method, thus obtaining multiple seasonal models, one for each class. Fundamentally:

- We define a set of classes $C \cup \{c_0\}$, where c_0 is a base class
- We classify the sampling data as belonging to two classes c_j and c_0
- We update the seasonal data of classes c_j and c_0
- We predict the class of the future power usages and predict using the appropriate seasonal model.

The method is generic (classes can be defined in several ways) but in the current implementation of Optiflex we use four classes + the base class (class 0).

The classification method is based on aggregated daily weather data: average daily temperature, average daily irradiance. To classify samples, we use static data associated with the location of the pilot site: average monthly temperature, and average monthly irradiance. The four classes intuitively correspond to “Hot - Sunny”, “Hot - Cloudy”, “Cold - Sunny”, and “Cold - Cloudy” days.

Classification is then performed as in “Table 9” where letter I stands for Irradiance and letter T stands for Temperature.

4.5 Scheduler

The scheduling algorithm simultaneously considers the entire set of controllable flexibilities. It implements a Model Predictive Control scheme, in summary, it considers the control actions over a future period of time called planning horizon and it actually implements only a smaller portion of the control actions called control horizon and the process repeats when the control horizon is elapsed. The planning horizon is defined as a discretized time interval T divided in timeslots (24h, discretized in 288 timeslots of 5 minutes each in the current implementation of Optiflex) and the control horizon is 1 timeslot. The control actions are therefore discretized and assumed constant during each timeslot. Control actions can be categorized as follows:

- Binary control actions: related to the control of a power switch. 1: flexibility connected to the grid, 0: otherwise
- Continuous control action: this is normally related to a power setpoint of the flexibility (commonly associated with storages or chargers), values are bound to operational constraints of the flexibility.

The scheduler has to respect some constraints operating the flexibilities. In particular for binary flexibilities, the amount of time the flexibility is connected to the grid must be sufficient to ensure that the flexibility can satisfy the energy demand. Furthermore the flexibility should not change its state too frequently and too many times during period T in order to preserve the lifetime of the physical switch.



For continuous flexibilities, the scheduler must maintain the state of the flexibility within bounds and reach the desired state at time t .

The scheduler considers the behavior of the rest of the grid in order to account for the uncontrollable portion of the power the model considers an aggregated signal. The software architecture of Optiflex, as done for demand estimation, allows to use modular implementations of the scheduler. For moderate sized pilot sites (hundreds of flexibilities) an exact approach based on a MILP formulation is used. For larger pilot sites (thousands of flexibilities) a fast optimizer based on local search heuristics is used.

4.6 Scalability of the centralized solution

4.6.1 Computational time

For what concerns scalability tests, we performed *in silico* tests and *in production* tests. In simulation, we performed tests with up to 5'000 flexibilities. Results confirm that the developed algorithms can handle the optimization of such a big-scale test case with a moderate impact on computational time (order of seconds, see Figure 15).

In a production environment, we were able to test with up to 647 meters and 156 flexibilities in total, as at that time it was the largest possible test set with real data. The computation time was in line with the values obtained in simulation for this number of meters and flexibility. Details of the computing environment:

1. Intel(R) Core(TM) i7-4770K CPU @ 3.50GHz CPU:4, RAM: 16GB
2. Operating System: Linux Ubuntu 18.04
3. All algorithms are implemented in Python language with interpreter Python 3.6
4. All algorithms are mono-threaded in the current implementation of Optiflex.

The Optiflex algorithms that are computationally more expensive are the disaggregator and the scheduler. The disaggregation time scales linearly with the number of flexibilities, as each one of them is disaggregated singularly, while the heuristic scheduler scales with $O(n^2)$. Consequently, by significantly increasing the number of flexibilities scheduled, one could run into computational time problems. For example, maintaining the current configuration, scheduling 100'000 flexibilities would take about 25 minutes. However, this problem could be easily solved by dividing flexibilities into smaller groups according to their location in the distribution grid.

4.6.2 Data requirements

The algorithm uses data with a 5-minute time resolution. The disaggregator runs once a day, and each time one day of data is disaggregated. Therefore 576 datapoints are used per flexibility (288 for active power P and 288 for reactive power Q of the meter). The scheduler runs every 5-minutes and generates a 288-points signal for each of the controlled flexibilities. These are not very large amounts of data, and, even with a high number of flexibilities, they can be easily handled by the Big Data platform Kibid.




Figure 15: Computational time of scheduler.

4.7 Extensions for near real-time DSM

A typical deployment of Optiflex, like the one illustrated in Figure 10, is characterized by a data collection process that does not need any additional device or a dedicated collection infrastructure. The downside is that it is not suited for real-time applications.

Anyway, the LIC pilot features a community battery installed with the purpose of maximizing the self-consumption of the community which requires near real-time control. In this project, we extended the Optiflex framework to include a fast communication channel via TCP protocol to allow the NUC to write data in Kibid with a frequency of 1 sample per minute. Data collection requires the acquisition of the net-power metered at the Point of Common Coupling of the LIC community with the rest of the grid and the power metering of the community battery. The software architecture is unchanged as different software modules can work simultaneously on the same data storage platform.

The control algorithm is purely reactive. It aims at maximizing self-consumption as a first objective




Figure 16: Schematic representation of Optiflex extension to deal with near real-time control.




Figure 17: Samples of battery control from 28.6.2020 to 8.7.2020

and performing peak-shaving as a secondary objective. The algorithm tracks the net power profile and charges the battery in case the net power is below a given threshold (set to 0 in the deployment, i.e., the community is producing more than what is consuming). The algorithm decides to inject power into the grid when the net power is above a certain threshold. The threshold is updated dynamically: it tracks the 75th percentile of the net power in the last days, therefore the algorithm injects in the high portion of power consumption of the community reducing peaks. If the battery is still full approaching the morning, the threshold adjusts automatically so that the battery gets almost empty for the next charging cycle. Discharging power is also capped at 20kW to prevent the fast aging of the battery.

4.8 Report on self-consumption via battery control

In order to compute the fraction of self consumption, the net photovoltaic production must be computed. As mentioned in section 4.3, photovoltaic installations of the LIC community are non-monitored. Therefore we exploited our GHI estimator to estimate PV production.

The self consumption ratio (S_r) is the amount of PV production locally consumed and it is computed as follows:

$$S_r = \frac{\int^T P_{pv} - P_{export}}{\int^T P_{pv}}$$

The period in which we were able to perform both PV estimation and direct control of the community battery starts from September 2020. Figures 18 to 22 illustrate different periods in which community battery was operated with power limits set to 20 and 50kW in both charging and discharging operations.

We observe that in September 2020, fig. 18, the battery operates in a limited way with some days almost idle. Anyway, the self consumption amounts to 98.54%. This is an indication that the installed PV power is relatively small with respect to the community consumption. This is confirmed by Figure 19 in which the battery is not operated but the self consumption still amounts to 91.85%.

The period in which we did not control the battery and the self consumption reached one of the lowest values is August 2021 (see fig 20 and a detail in fig. 21) in which the self consumption ratio is as high as 83.35%

The best self-consumption performances were reached in May 2021, see fig. 22, where battery was operated without limits and self consumption ratio reached 99.95%, that is basically all photovoltaic




Figure 18: LIC with Battery control, limited power to 20kW, from 9.9.2020 to 25.9.2020, self consumption 98.54%

production was used by the community or stored in the community battery.



Figure 19: LIC without Battery control from 21.3.2021 to 16.4.2021, self consumption 91.85%



Figure 20: LIC without Battery control from 9.8.2021 to 29.8.2021, self consumption 83.35%



Figure 21: LIC without Battery control detail 14 and 15 August 2021




Figure 22: LIC with Battery control, limited power up to 50kW, from 17.5.2021 to 16.6.2021, self consumption 99.95%



4.9 Report on peak shaving via DSM

During the period in which we applied centralized DSM, we were able to control only up to 6 resistive boilers and just one heat pump. At the time of the test for the centralized approach, only these devices were correctly connected to the switches of the smart meters. The control actions were successfully computed, sent to devices and actuated by the smart-metering infrastructure. The state of the relays was monitored, and we were able to verify that actuation took place and devices were correctly switched off by the system.

These limited control actions did not result in a sensible shift of the peak consumption.

Anyway, some quantitative analysis can be performed observing when these actions occur with respect to the overall network load. We report the number of energized flexibilities, i.e. the control actions that allow loads from absorbing power from the grid, with respect to the PCC power. We observe that a shift in load occur when flexibilities are prevented to absorb power during times of high consumption.

In Table 10 for different periods of 2021 and for the 75th, 90th, and 95th percentiles, we report:

- “Number of samples”: The number of time slots in which PCC power was above the percentile (in our setting one time slot lasts 5 minutes)
- “One flexi ON”: The number of time slots in which at least one flexibility was allowed to absorb power
- “Number flexi ON / demand”: The ratio of energized flexibilities with respect to the total demand, i.e. the total number of time slots in which the flexibilities are expected to be energized
- “One flexi ON / number of samples”: The ratio between the number of time slots in which at least one flexibility was allowed to absorb power over the number of time slots above the percentile.

These tests are also illustrated in figures 23-29.

We observe that in April 2021, late June 2021 and August 2021, DSM limited the activation of flexibilities to about half of the timeslots when the PCC power was high. Notably, this is more clear for the 95th percentile. The effect of DSM is less apparent for May 2021 and early June 2021. The demand shift is also clear from the fact that the total demand scheduled for the periods is always less than the related complemented percentile. We remark that the PCC power is of course unknown to the system when actuation decisions are taken. DSM is performed using a forecast estimate of the uncontrollable load + photovoltaic production.

In order to further evaluate the potential of peak shaving with DSM control we run simulation with a calibrated model of the LIC area with the help of a low voltage simulation framework called OPTISIM and developed within SUPSI. In the simulation, we controlled 10 heat pumps and 12 electric heaters totalling 26kW nominal power. The simulation considered a 30kWp PV installation on the kindergarten and other 3 PV plants totalling 32.76kWp.

In “Figure 30” we report the overall active power of the test-set. It ranges approximately from -20kW to 80kW. In the top sub-plot we report the overall power when all loads are not controlled and can freely absorb power when necessary. In the bottom sub-plot we report the overall power when all flexibilities are




Figure 23: Energized Flexibilities vs PCC power, April 2021



Figure 24: Energized Flexibilities vs PCC power, May 2021




Figure 25: Energized Flexibilities vs PCC power, June 2021, 8 days




Figure 26: Energized Flexibilities vs PCC power, June 2021, 22 days




Figure 27: Energized Flexibilities vs PCC power, June 2021, 3 days detail




Figure 28: Energized Flexibilities vs PCC power, July 2021



Figure 29: Energized Flexibilities vs PCC power, August 2021



Period: 2021-04-01 - 2021-04-30 (30 days)			
Percentile	75°	90°	95°
Number of samples	2048	820	410
One flexi ON	1251	484	229
Number flexi ON / demand	0.2	0.08	0.04
One flexi ON / number of samples	0.61	0.59	0.56

Period: 2021-05-01 - 2021-05-31 (31 days)			
Percentile	75°	90°	95°
Number of samples	2139	856	428
One flexi ON	1556	608	299
Number flexi ON / demand	0.24	0.09	0.04
One flexi ON / number of samples	0.73	0.71	0.7

Period: 2021-06-01 - 2021-06-08 (8 days)			
Percentile	75°	90°	95°
Number of samples	504	202	101
One flexi ON	329	131	68
Number flexi ON / demand	0.21	0.08	0.04
One flexi ON / number of samples	0.65	0.65	0.67

Period: 2021-06-08 - 2021-06-30 (22 days)			
Percentile	75°	90°	95°
Number of samples	1561	625	313
One flexi ON	820	320	141
Number flexi ON / demand	0.2	0.07	0.03
One flexi ON / number of samples	0.53	0.51	0.45

Period: 2021-07-01 - 2021-07-31 (31 days)			
Percentile	75°	90°	95°
Number of samples	1642	657	329
One flexi ON	996	420	190
Number flexi ON / demand	0.24	0.1	0.04
One flexi ON / number of samples	0.61	0.64	0.58

Period: 2021-08-01 - 2021-08-31 (31 days)			
Percentile	75°	90°	95°
Number of samples	2156	862	432
One flexi ON	1152	448	200
Number flexi ON / demand	0.21	0.08	0.04
One flexi ON / number of samples	0.53	0.52	0.46

Table 10: Energized flexibilities with respect to PCC power percentiles

controlled by Optiflex. We observe a pattern for the uncontrolled case where loads tend to accumulate during morning and evening hours forming power peaks.

We then observe that Optiflex is capable to prevent the formation of such high peaks by spreading loads along the day. We report that the overall energy provided on a daily basis for the test-set does not differ between the two settings, that is all loads are absorbing the same amount of energy. From a preliminary analysis of the test-set we observe KPI peak reductions of 30-40% with cases of peak reduction of up to 50%.



Figure 30: Simulation of LIC Pilot, January, detail.

5 Decentralized management

5.1 Algorithms description

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.

5.1.1 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 [11], 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 $u = [p_{ch}^T, p_{ds}^T]^T \in R^{2T}$ 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, $\tilde{u} = [p_{ch}, p_{ds}] \in R^{T \times 2}$ being the same vector reshaped in a 2 columns matrix, $\hat{p} \in R^T$ being the forecast power at household's main for the next control horizon, $y \in R^T$, $s_{ch} \in R^T$, $s_{ds} \in R^T$



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

$$u^*, y^* = \underset{u, y}{\operatorname{argmin}} \sum_t^T y_t + \|s_{ch}\|^2 + \|s_{ds}\|^2 \quad (9)$$

$$x_{t+1} = Ax_t + B\tilde{u}^T \quad (10)$$

$$y p_b (\tilde{u}[1, -1]^T + \hat{p}) \quad (11)$$

$$y p_s (\tilde{u}[1, -1]^T + \hat{p}) \quad (12)$$

$$x \in [x_{\min}, x_{\max}] \quad u \in [u_{\min}, u_{\max}] \quad (13)$$

$$s_{ch}, s_{ds} \geq 0 \quad (14)$$

$$s_{ch} - \hat{p} \quad s_{ds} \hat{p} \quad (15)$$

$$u [s_{ch}, s_{ds}] \quad (16)$$

where \hat{p} stands for R_+ , indicating element-wise inequalities, $p_b \in R^T$ and $p_s \in R^T$ are the business as usual buying and selling prices. We start analyzing the objective function (9) term-wise. The first summation in (9) 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 (17):

$$c(p_t) = \begin{cases} p_{b,t} p_t, & \text{if } p_t \geq 0 \\ p_{s,t} p_t, & \text{otherwise} \end{cases} \quad (17)$$

where p_b and p_s are the prices generated by the price scheme 1.




Figure 31: 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 (17)

The cost can be thought of as the maximum over two affine functions (the first and second line of equation (17), respectively). Equations (11),(12) 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 31. Equation (10) describes the battery's dynamics. $A \in R_+$ and $B \in R_+^{1 \times 2}$ are the discrete dynamics



matrices obtained by the continuous one through exact discretization [12]:

$$\begin{aligned} A &= e^{A_c dt} \\ B &= A_c^{-1} (A_d - I) B_c \end{aligned} \quad (18)$$

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 characteristic 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 32 for the case of s_{ds} , the equations (14) and (15) 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 32). Its optimal value can then be used to define the feasible regions of the battery charging power, as done by equation (16). The same reasoning done in Figure 32 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 32: Visual explanation of the change in the feasible space for the discharging power.

5.1.2 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:



$$cM \frac{\partial T}{\partial t} = c\Gamma (T_{i,t} - T_{o,t}) - U (T - T_{ext}) + P_{el,t} \quad (19)$$

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 have 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 [13]. 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 LightGBM¹² 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:

$$u^*, y^* = \underset{u, y}{\operatorname{argmin}} \sum_t^T y_t - \sum_t^T \min(\gamma, 0) \quad (20)$$

$$\text{s.t.} \quad y p_b (\hat{p}_b (1 - u) + \hat{p}) \quad (21)$$

$$y p_s (\hat{p}_b (1 - u) + \hat{p}) \quad (22)$$

$$S [(1 - u) p_{nom} - \hat{p}_b] \gamma \quad (23)$$

$$\sum_{t=1}^{T-1} |\Delta u| n_{ch} \quad (24)$$

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 relaxes the energy invariance constraint (23). Here S is a summation matrix which sums the energy in the pre-defined time slots (3 hours). Equation (24) further prevents the boiler from being turned on and off more than n_{ch} times in a control horizon.

5.1.3 Coordination

The presented problem formulations for the battery and the boiler, (9) and (20) respectively, minimize the end users' business as usual costs. These can be adapted to jointly minimize the cost of the energy community. In particular, called $e(x) = c \left(\sum_{i=1}^N u_i \right) - \sum_{i=1}^N c(u_i)$ the economic surplus generated by

¹²<https://lightgbm.readthedocs.io/>



being in an energy community, if we allocate an α_i fraction of it to each user, the total economic cost of the agent becomes:

$$c_{tot,i} = c(u_i) + \alpha_i e(u) \quad (25)$$

$$= \alpha_i c \left(\sum_{i=1}^N u_i \right) + (1 - \alpha_i) c(u_i) \quad (26)$$

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 [11]. Replacing the agent cost with the auxiliary variable y as in (9) and (20), the final objective function (for the battery) then becomes:

$$\alpha_i \nabla c \left(\sum_{i=1}^N u_{i,pre} \right)^T u_i + (1 - \alpha_i) \sum_{i=1}^T y + \lambda_i^T u + \rho_d \|u - u_{pre}\|^2 + \|s_{ch}\|^2 + \|s_{ds}\|^2 \quad (27)$$

where $u_{i,pre}$ are the agents actions at the previous iteration. Minimizing the aforementioned objective function leads to a weighted (with coefficient α_i) Nash Equilibrium.

5.2 Report on battery control

The battery has had several problems. One module has been replaced, and the battery was not controllable for quite some time. Also, because the community is already self-consuming a lot of energy in winter, in this period of the year, the battery is not particularly useful when it comes to self-consumption optimization and was left idle (Figure 33). Even when it worked, sometimes it had problems estimating the state of charge. For example, it would not report a SOC of more than 80%, when it was clear that it was 100% charged by monitoring its voltage. These problems, unfortunately, often coincided with the periods when it was scheduled to test the decentralized algorithms. However, the results obtained during the short periods of decentralized battery operation confirmed the results of the simulations. The




Figure 33: Time course of power at the community's coupling point in 2021.

battery control algorithm differs from the centralized algorithm in that it is predictive and not reactive. In contrast to the centralized algorithm, a forecast of the power trend at the community coupling point is made, and, based on it, battery actions are planned. As in the centralized case, the battery is in charge



of both optimizing the community's self-consumption and reducing power fluctuations by peak shaving and valley filling.

Figure 34 shows an example of battery operation using the Hive Power algorithm. It can be seen that most of the time, the battery is able to flatten the consumption profile of the community, thus reducing its peak powers and consequently its losses. It can be seen, however, that when the forecast significantly underestimates consumption, one gets into situations where the battery is drained entirely ahead of time. An example can be spotted by looking at Figure 34 in the morning of July 22nd at 08:00. The battery aggressively discharges during the evening of July 21, effectively lowering the peak consumption for that day; however the battery hit the lower bound of the state of charge overnight, finding itself empty in the morning of 22, when a high peak occurs. At present, battery control is purely deterministic; a possible evolution of it would be to extend it to a probabilistic (robust or stochastic) formulation. This is, however, very complicated to set up in the context of distributed control and has not yet been done by Hive Power.




Figure 34: Example of battery control, when actuated using Hive Power's algorithm.

5.3 Report on boiler control

For decentralized load control, up to 5 boilers were controlled. This equals to one unit less compared to the control undertaken during the centralized algorithm tests, since one of the boiler had technical actuation issues. Unfortunately, a reading problem occurred at the community coupling point during the load control testing phase. For reasons that are still unknown, the meter installed by AEM at the coupling point began to send plausible but wrong data. Another testing phase was foreseen for the second part of 2021, but the battery broke, and we preferred to continue the test of the centralized load control solution. This report will show the results of decentralized load control on three boilers during the testing phase between November 2020 and March 2021. As also shown by simulations, the cost reduction potential given by controlling thermal loads is limited. Nevertheless, in the period in which the community meter was working, the results proved to be in line with what was expected. Controlling the boilers compared to controlling the battery poses greater challenges. First of all, you do not have absolute control; rather, you can only force the load off, but you cannot force it on. This, in the control phase, makes estimating consumption complicated. Another problem that has proven to be complicated to solve is the need to disaggregate consumption from the house meter measurements; this task was harder w.r.t. the centralized setting, due to computational restrictions posed by the distributed hardware. The centralized solution for the disaggregation was formulated as a mixed integer quadratic program, and the GUROBI solver was used to solve it; unfortunately, it was not possible to install GUROBI on the ARM architecture. The choice was made not to install a separate meter for thermal loads for financial reasons. Reading data through an optical interface allows data to be acquired from the meter with a sampling rate of about one measurement every 5 seconds. This sampling rate is still too slow to



use high-frequency NILM algorithms, and the advantage over the 1-minute resolution of the centralized solution is almost negligible. As for the centralized solution, the relay status readout is available and used. The next section describes the disaggregation algorithm that runs on the Strato PC installed on each meter of the community.

5.3.1 Local disaggregation

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 than one minute.

This setting differs from the centralized case mainly on the hardware used to perform the disaggregation: in the centralized case we were not restricted by the computational power of the distributed hardware (the Stratos), which allowed to use optimization-based algorithms. In this case we had to opt for a computationally less intensive algorithm. 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 heuristic is used to merge the derivatives
- Identify clusters in the jumps using an unsupervised algorithm (Bayesian Gaussian Mixture Model [14]). Fig. 35 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 separated 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 $x_{gr} + \text{centroid}$ is close to the signal




Figure 35: 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).

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 algorithm 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. 36.

5.3.2 Control performance

During the testing phase, the distributed algorithms that controlled both the battery and the boilers were activated. Since when the distributed control was successfully activated, the self-consumption of the community at its coupling point was naturally 100%, the battery was idle most of the time. Unfortunately, it was not possible to carry out a second phase of distributed load and battery control testing, as the battery broke down several times and could not be repaired in time. Figure 37 shows an example of controlling a water heater for a period of one week. It can be seen that the activation of the relay that allows the water heater to consume coincides with periods when consumption is low and periods when the internal energy purchase price is low. The boiler is also allowed to stay on during periods of high price and consumption. There are several reasons for this. First of all, it is always vital to ensure the comfort of users, and since the performance of the disaggregation algorithms was not particularly good, a solution that overestimated the energy requirements of the boilers was chosen. Second, the quality of the forecast influences the performance since these are predictive and not reactive algorithms.

Figures 38 shows how long the boiler was allowed to stay on in percent of the time, as a function of the quantile of the power at the coupling point, detrended by the average daily power at the coupling point. Figure 39 shows how long the boiler was allowed to stay on in percent of the time, as a function of the internal buying price in the community, defined by 1.




Figure 36: 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).

It can be seen how boilers tend to be allowed to turn on longer when energy prices are low. It can also be seen how similarly boilers are forced off longer in periods when community consumption is high. This corresponds, of course, to the desired behavior. It is also important to note that although this effect does not appear to be very pronounced, a flat non-increasing curve is already a positive result. In a situation where the control of the boilers was leased to a random scheduler, to an increase in the consumption quantile at the coupling point would have corresponded an increase in the percentage in which the boilers are left on since they actively and significantly contribute to increasing community consumption.




Figure 37: Example of boiler control. PCC: Power at the coupling point of the community. Relay: Status of the relay (0: forced off, 1: allowed on)




Figure 38: Relative time the boilers were allowed to charge as a function of the power at the coupling point.




Figure 39: Relative time the boilers were allowed to charge as a function of the internal price in the community.



5.4 Decentralized management through blockchain-enabled smart meters

An energy community like LIC has an intrinsic decentralized structure, being roughly constituted by a collection of smart meters, able to acquire, process and store data about consumption/production and by the related prosumers. For that reason, a decentralized EC management based on blockchain technology constitutes an interesting solution. Unfortunately, the usage of blockchain technology in the energy sector, and in general in IoT applications, still encounters barriers that can limit its adoption. Among them, currently the main challenges are related to privacy management and scalability issues. As regards the privacy aspect, which is basically the proper handling of private data (e.g. energy consumption) on an immutable ledger like a blockchain, different strategies can be adopted in order to be compliant with the different national and international regulations. For example, GDPR [15], adopted in the EU, poses significant barriers to store private data on blockchain networks.

As mentioned above, scalability is certainly one of the most significant challenges for blockchain adoption, especially within IoT applications, which is unfortunately the case of solutions tailored for the ECs management. Under a generic point of view, scalability on a blockchain is determined by the number of transactions it can complete in a second [16]. To scale up effectively, a blockchain would need to perform thousands of transactions per second (TPS). Yet, the Bitcoin blockchain can only handle about 5 to 7 TPS while the Ethereum can handle 10 to 25 TPS [17]. This is a sharp contrast to systems such as VISA that handle 1500 to 2000 TPS [18]. However, straightforward scaling up of the blockchain interferes with its integral security and decentralization attributes. In simple terms, scaling up the blockchain means enabling it to process transactions faster. To achieve this, the same system would either need to simplify its security measures or decentralization capacity.

Scalability, security and decentralization are the three main pillars that affect the mass adoption of blockchain technology. Decentralization and security are its key selling points. Yet, the market cannot meet its current and future demand without sustainable scalability. A strong structural correlation exists between these three pillars. This situation is commonly referred to as the 'Blockchain Scalability Trilemma' as described by Vitalik Buterin, the founder of Ethereum [19] [20]. According to Buterin, you cannot improve all the pillars of an individual blockchain system together. The system limits you to improve only one or two pillars at the expense of the third. Developers found that altering the parameters of one pillar would in turn compromise the other two pillars of the system. The three scenarios below illustrate the blockchain scalability trilemma:

- Improving Decentralization and Security protocols reduces Scalability options
- Improving Scalability and Security protocols reduces Decentralization options
- Improving Scalability and Decentralization protocols reduces Security options.

Several solutions have been proposed to increase the solutions scalability in order to have transactions both faster to be performed and more economically sustainable. This fact would improve adoption of blockchain in a wide variety of applications, among them the data management in the energy sector.

A promising approach to front the scalability problem in a blockchain is basically to move the issue to a second layer, the main blockchain itself being the first one. Currently, the available 2nd layer solutions are essentially based on two technologies: state channels and sidechains. In both the cases, the most of the user interactions is moved out of the blockchain, on the second layer, where fast, safe and cheap transactions between participants can take place.

The following of this chapter describes the main available second-layer blockchain solutions and developed applications that exploited the LIC pilot. Section 5.4.1 details the AssetManager application, which is based, not needing significant scalability requirements, on a 1st layer approach; instead the rest of the chapter relates to second-layer solutions. A first part refers to the two aforementioned 2nd layer solutions in sections 5.4.2 and 5.4.3, while 5.4.4 introduces the technological framework used in LIC. Then sections 5.4.5, 5.4.6, 5.4.7 and 5.4.8 describe in detail the 2nd layer applications developed and deployed in LIC.



5.4.1 1st layer application: asset manager (AM)

In the context of fully automating a community via blockchain, first of all, the community's users and assets need to be mapped within it. For this, SUPSI has developed a prototype asset management application called asset manager (AM). The main element of AM is the prosumers community, i.e. a group of users that decide to manage together a local energy market. Examples can be a district with a unique DSO (Distribution System Operator), a condominium or an energy community like LIC. Each user can be a producer and consequently, generate energy to sell to other users or externally to the DSO. Besides, some assets related to energy production and storage too expensive to be owned by a single user can be maintained by the entire community defining a governance approach. AM implements the automatic management of an assets collection for a community, taking into account both the asset governance and its operations (e.g. the revenues generation and distribution). Practically, AM provides the community governance using Aragon platform [21]. Aragon provides an open-source platform, named *AragonOS* [22], where different Ethereum applications can work together, similarly to processes in an operating system. Besides, using *AragonOS* a significant collection of applications already developed and tested by Aragon is available. Among them, the *Voting* application [23] provides the functionalities needed to manage the governance via a collection of votings, which can be configured and customized (e.g. majority to reach, quorum, etc.). Thus, the most meaningful operations related to an asset are handled by the governance community using *Voting* and reported in the following list:

- *Asset creation*: a new asset has to be bought and installed
- *Shares initial distribution*: after an asset creation, the shares related to the owners have to be distributed
- *Asset deactivation*: an asset has a problem and is not able to generate revenues
- *Asset activation*: an asset can generate revenues again (e.g. it has been repaired).

Fig. 40 shows the interactions between AM and *Voting* applications via *AragonOS*. It is important to remark how unique AM smart contract that works with *AragonOS* is *AssetManager*.




Figure 40: Interactions between *Voting* and *AssetManager* using *AragonOS*. The green circles refer to smart contracts provided by Aragon.

When community votes in favour of asset creation, a new instance of *Asset* and *SHT* are deployed on the blockchain. In details, an asset is defined in the smart contracts with the following features:

- a numerical *id*, used by the DAO to identify the asset
- a *name* and a *description*, useful to briefly describe the asset



- a trusted *oracle*, practically a wallet allowed to perform significant operations such as the distribution of the revenues
- the address of its *SHT* contract for the management of the shares, just deployed during the asset creation
- the address of *RVT* token, which will be used to distribute the revenues.

In order to correctly implement all the operations, such as the revenue distribution, during its life an asset can be in one of the following states:

- ACTIVE: when shares can be exchanged, and the shareowners can not claim any revenues
- FROZEN: when shares cannot be exchanged, and the shareowners can claim revenues
- INACTIVE: when shares cannot be exchanged, and the shareowners can not claim revenues.

After the creation, an asset is in ACTIVE state, i.e. it is working and creating revenues (e.g. a PV plant is producing energy).

When an asset is created, the related *SHT* is deployed. Besides, 1000 *SHT* are minted and assigned to the asset. Thus initially, all shares are owned by the *SHT* smart contract. To complete the asset setup, the community has to vote to decide how the shares will be distributed among the users. The community can determine the shares amounts of every user. Once a user has received the number of its shares, it can exchange them according to ERC-20 standard [24] when the asset is ACTIVE.

As mentioned above, three states are available for an asset. Fig. 41 depicts the transitions with the related Solidity functions. It is important to remark that `activateAsset()` and `deactivateAsset()`, pertinent to ACTIVE \leftrightarrow INACTIVE transitions, have to be performed by the community with a vote. Instead, `setRevenue()` and `defrostAsset()` can be performed only by the trusted aforementioned oracle.




Figure 41: State transitions of an asset

When the asset state is *ACTIVE* a revenue can be set by the asset trusted oracle using the `setRevenue()` functions. After this transaction, the asset state becomes *FROZEN*, and a certain amount of *RVT* tokens are staked in the *Asset* smart contract. Besides, a new *RevenueClaiming* contract is deployed on the blockchain. It contains all the information related to the revenue (e.g. the *RVT* amount, the shareowners having already claimed the income, etc.). When the state is *FROZEN*, no shares can be transferred



to avoid cheatings due to possible double-spending. Besides, each shareowner is allowed to claim its revenue portion, i.e. a part of RVT staking proportional to owned shares, using the `claim()` function. After a certain amount of time, reasonably long to assure an easy claim to the shareowners, the asset state becomes ACTIVE again, thanks to `defrostAsset()` function performed by the trusted oracle. Fig. 42 shows the timeline of a revenue life: the initial setting, the claims when the asset is FROZEN, and the defrosting. Besides, when an asset is FROZEN no new revenues can be created. Consequently, the life of revenue currently claimable has to end with a `defrostAsset()` transaction before having a new one.




Figure 42: Timeline of a revenue life

If an asset is unable to create new revenues (e.g. a PV plant has an electrical problem and needs to be repaired to produce energy again), the community has to vote to set the state to INACTIVE using the function `deactivateAsset()`. When an asset is INACTIVE, its shares are locked, and no new revenues can be generated. If and when the asset will be able again to generate new revenues (i.e. a PV plant has been repaired), the community can vote to reactivate it using the function `activateAsset()`.

5.4.2 2nd layer solutions: state channels

State channels technology provides secure P2P channels that can be used by two nodes connected to a blockchain to exchange data avoiding to do transactions on the main network. The only transactions performed on-chain are the opening and closing transactions, which are needed to establishes and to settle the state channel. The main process downside is that it requires the full availability of the two involved participants; otherwise the security of the exchanged data, typically tokens, can be compromised. As a consequence, state channels are substantially useful in cases where nodes exchange many state updates over a long period, to mitigate the costs of creating and settling a channel. For this reason, probably the payment channel is the typical use case where this technology is mostly used. The following list reports promising projects that provide state channels implementation:

- Lightning Network [25]
- Raiden Network [26]
- Liquidity [27]
- Celer [28]

The first three projects are among the first to provide *state channels* platforms. More precisely, they are projected to provide payment channels in order to transfer tokens in the channels. Instead, Celer project aims to manage exchanging of any type of data, not only tokens.

State channels is undoubtedly a meaningful 2nd layer technology, but it appears not to be functional for the needed usage. Indeed, in an energy community like LIC it would be advantageous to provide a light and cheap blockchain usable by all the users instead of having a huge quantity of channels connecting couples of nodes. For example, the information related to nodes consumption and production could



be accessible by all the nodes with different level of allowances, not only by the two nodes creating a channel. Indeed, for what concerns our discussion, the unique use case where a *state channels* solution like the one provided by Celer project [28] could be exploited, is an energy community composed only by two nodes.

For the reasons mentioned above, *state channels* technology was not considered as suitable for the management of 2nd layer blockchain solutions in an energy community.

5.4.3 2nd layer solutions: sidechains

Sidechains are separate blockchain networks, able to interact with the main chain. They have their own consensus mechanism, level of security, and tokens and can be public or private. When sidechain security is compromised, the damage does not affect the other connected networks. Moreover, two sidechains connected together can transfer any data. For example, tokens can be exchanged at a pre-determined 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.

As opposed to state channels, transactions that occur on a sidechain are not private between the participants of a transaction. They are published on the sidechain network and thus visible to anyone who has access to the ledger. Moreover, the chain nodes do not need to be always available.

The main drawback of this technology is that setting up a sidechain is a significant effort, as it means to build the entire platform from scratch.

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 meters and used to store data about nodes and interact with other main blockchains.

For these reasons, the *sidechains* approach was chosen to be implemented. The project selected for the implementation of sidechain solutions in LIC was Cosmos [29], one of the most notable platforms currently available. Cosmos was chosen because it provides significant warranties in terms of documentation, applications to develop sidechains and strategy to create networks of chains.

5.4.4 Cosmos and Tendermint

Cosmos [29] provides a complete environment to create custom side chains, which can be used as meaningful solutions for decentralized management of energy communities. Under a generic point of view, a Cosmos network is a collection of independent and interconnected custom blockchains, which run in parallel. Each blockchain is powered by Tendermint consensus algorithm, which provides a byzantine-fault-tolerant mechanism [30].

Theoretically, the creation of a blockchain is based on the development of its three main layers: *Networking*, *Consensus* and *Application*. In a blockchain ecosystem based on smart-contracts like Ethereum, an application layer is provided, named Ethereum Virtual Machine (EVM), and is used to deploy on the chain smart contracts with custom functionalities. This approach has many advantages, but it is based on a monolithic platform where the three layers cannot work separately.

The solution proposed by Cosmos is more modular: it is a solution that packages the networking and consensus layers of a blockchain into a generic engine, allowing developers to focus on application development. The communication with the *Application* layer is provided by the ABCI protocol. Figure 43 shows the structure of a typical Tendermint application.

As explained above, Tendermint is substantially a consensus/networking algorithm that can be easily configured and be used together with Cosmos SDK to create a custom side chain to manage a EC. Moreover, the custom sidechain can be integrated inside the Cosmos Network using the IBC protocol and interacting with other blockchains, typically exchanging data, as shown in Figure 44. The next paragraphs describe in details all the main elements that constitute the Cosmos framework.

It is important to remark how the decentralized approach of Cosmos can be beneficial in the management of energy communities. Indeed, some chains can be demanded for specific purposes, such as




Figure 43: Tendermint application structure




Figure 44: Tendermint blockchains in Cosmos network connected via IBC protocol

the consumption and production measurements, which are implicitly customized for the single community. Instead, other chains can be used for more general tasks (e.g. tokens management, community governance, etc.).

Consensus layer: *Consensus* layer is the part of a Tendermint application to whom the management of the consensus in the blockchain is demanded. Currently, two mechanisms are implemented and usable in Tendermint: *Proof-of-Authority* (PoA) [31] and *Proof-of-Stake* (PoS) [32]. These approaches are very sustainable under the computational point of view, unlike mechanisms based on *Proof-of-Work* [33], still broadly used in blockchain-based solutions.

It is remarkable to note that the sustainability of PoA and PoS has two meaningful consequences. The former is that embedded devices such as the smart meters have sufficient resources to run the Tendermint applications. Moreover, the energy consumed by these two mechanisms is negligible if compared to PoW.

Both PoA and PoS were positively tested on the LIC pilot, which is basically an energy community composed of 20 nodes. Thus, on the pilot a sidechain of 20 nodes was deployed. In a first step a chain based on PoA, simpler to configure than PoS, was implemented and deployed since June 2020 until the end of the year. Instead, since the beginning of 2021 PoS was positively tested in a new sidechain, which



is still running in the pilot nodes and will be used for new developments. It is important to remark how the consensus mechanism choice entails no consequences for the integration in the Cosmos platform. Indeed, the blockchains shown in Fig. 44 can communicate via IBC also if they use different consensus algorithms.

The consensus in a Tendermint network is effectively managed by a subset of the nodes, named *validators*. They are the unique nodes allowed both to create new blocks and to decide the transactions to insert in a block. Consequently, no new blocks can be created if all the validators are not connected to the network (e.g. for a power outage). In order to minimize the probability that this event occurs, the validators have to be chosen carefully depending on the installation of the nodes. For example, a good candidate to be a validator should have both a stable Internet connection and a safe power supply system.

Networking layer: *Networking* layer manages the connections between the nodes of Tendermint application taking care to properly transfer all over the network the information stored in the sidechain via transactions.

The links between the nodes are managed by a P2P protocol, which is secured by the usage of private and public keys for the identification. The protocol used is TCP. Upon a nodes establishes a successful TCP connection with a peer, two handshakes are performed to empower the chain security: the former for the encrypted authentication, the latter in order to check the Tendermint versioning. Thus, this layer provides a fast, secure and configurable platform that manages the network connections in a Tendermint application.

It is important to note how Tendermint provides the feature to configure persistent peers. These are intended to be trusted peers that can be very helpful to define a stable P2P structure. To empower the stability in the community side chain, in the pilot where a side chain has been deployed it was decided to configure a collection of persistent peers able to connect all the nodes in the network safely. This decision was taken considering how meters are usually installed in cabinets where the data connectivity can be affected by local interference. Consequently, a P2P structure taking into account the different installations can be useful to have a more stable network.

Application layer The previous paragraphs describe how the consensus and networking tasks, mandatory in any blockchain-based solution, are already provided by Tendermint code. It is possible to configure the software (e.g. setting PoA or PoS, defining the persistent peers, etc.), but no changes to the code are needed. Instead, the *Application* layer is the custom application part that has to be developed. Regarding the used technologies, all the code was developed in Go language [34] because Cosmos and Tendermint platforms are based on this programming language. Consequently, all the official documentation and the available examples are based on this programming language and so the development of custom Go applications was strongly facilitated.

The main aims of a custom *Application* layer are substantially two: the former is to provide communication with the other layers via ABCI, the latter is to interact with off-chain applications that try to perform transactions and queries via a REST API. Figure 45 shows the interactions between the aforementioned components in a chain of two nodes.

Figure 45 shows how ABCI manages the interactions between the custom applications with the other layers. Basically, it consists of a collection of Go methods that are suitable to properly operates the chain. For example, *CheckTx()* function is used to check if a transaction has to be rejected or not, performing the filtering feature mentioned above.

The second aim of the *Application* layer is to provide an interface for off-chain software that has to exchange information with the sidechain. This interaction is provided by a REST interface that allows the performing of query and transactions. The usage of a REST API has the significant consequence that the off-chain applications can be developed in any programming language, not necessarily in Go. Acting this way, specific functionalities, which can be more easily developed in other languages such as Python, are included in the off-chain part, demanding to the *Application* layer only the functionalities directly related to the chain interactions. An example can be the data gathering from a smart meter that has to be periodically stored in the chain and can be easily performed by dedicated Python modules




Figure 45: *Application* layer interactions in a chain constituted by two nodes

provided by the manufacturer of the meters. The REST provides two API services: the former for the transactions, which can result in a status change of the chain, the latter related to queries that are used to read data from the chain.

Under a logical point of view, it is simple to note how the *Application* layer can act as a custom filter and an adapter between the off-chain parts and the side chain, managed by Tendermint consensus/networking layers. The first task is fundamentally used to check the nodes allowance when they attempt to access to the chain with a transaction or a query. Typically, the nodes in a chain can have different access rights, which can be managed by the *Application* layer. In addition to the filtering task, this layer can be used to transform the information to store in the blockchain when custom transactions occur. Figure 46 shows an example of how filtering and adapting functionalities can operate.




Figure 46: *Application* layer filtering and adapting functionalities

5.4.5 Applications deployed in LIC: *Metering*

While in the previous sections a generic description of blockchain-based solutions suitable to be used in ECs has been reported, with a significant spotlight on the description of the Cosmos framework, in the present and following sections the focus is on the sidechain applications developed and deployed in LIC. Each of them is a decentralized solution based on Cosmos sidechains, which run on the LIC embedded devices described in sections 2 and 3.6.

The *Metering* application was the first to be developed and, as a consequence, its functionalities are quite simple. Basically, every quarter of hour it saves data about average consumption and production



of the last 15 minutes. The datasets are saved in plain text and each meter is allowed to read data related to other nodes. Consequently, every 15 minutes the instances of *Metering* running in parallel on the nodes perform new transactions and save the related data on the chain. Figure 47 shows the energy consumption and production stored in the LIC sidechain. The green nodes correspond to producer, whereas the red ones to the consumers.




Figure 47: LIC data saved on sidechain by *Metering* application

5.4.6 Applications deployed in LIC: Auditable Tariffs (AT)

AT application provides an efficient approach to make the payment verifiable and the aggregator auditable by the users, preserving their privacy. In *AT* each node saves encrypted data about consumption and production on the Lugaggia chain. The data encoding is based on homomorphic encryption, which is locally provided by the *TrustedSum* Go application running on the Strato device (see Sections 2.1.2, 3.6.1 and 2.2 for details). *TrustedSum* code is released as open-source on the Gitlab platform¹³. 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 verify the payment validity, which is function of the community production.

AT algorithm is based on the trusted voting system described in [35]. Figure 48 shows the problem: Alice, Bob, Carol, Dave and Eve can vote in an election, but keep their votes secret. Besides, none of them trust Trent to count the votes. In [36] the authors shows how, with a proper blockchain-based approach, there is no need for a trust infrastructure and voters privacy is preserved. In addition, a zero-knowledge proof approach based on the Schnorr proof [37] and the Fiat-Shamir heuristic [38] can effectively perform a decentralized voting system. Figure 49 reports the operations sequence related to the trusted voting algorithm. As first step, each voters i registers its voting key with the following equation:

$$key_i = g^x \pmod{p} \quad (28)$$

where g and p are prime numbers with a size of 1024 bits chosen equal for all the participants, while x is a secret value chosen at random by the i th agent. Once the voters are registered, their voting keys are published on the sidechain. Next each voter calculates a y_i value as follows:

¹³<https://gitlab.com/supsi-dacd-isaac/trusted-sum>



$$y_i = \prod_{j=1}^{i-1} g^{x_j} / \prod_{j=i+1}^n g^{x_j} \quad (29)$$

Then, each voter can compute and share a vote v_i and its associated auditable hash:

$$\begin{aligned} v_i &= g^{x_i y_i} g^{v_i} \pmod{p} \\ \text{hash}_i &= H(g^{x_i y_i} g^{v_i}) \pmod{p} \end{aligned} \quad (30)$$

Now, it can be shown that if we define y_i as in (29), the following equivalence holds [39]:

$$\prod_{i=1}^n g^{x_i y_i} \pmod{p} = 1 \quad (31)$$

Thanks to (31), all the voters can retrieve the sum of all the votes without knowing the other's voters values with:

$$\begin{aligned} \text{tally} &= \prod_{i=1}^n g^{x_i y_i} g^{v_i} \pmod{p} \\ &= g^{\sum v_i} \pmod{p} \end{aligned} \quad (32)$$

Known $g^{\sum v_i} \pmod{p}$, it is finally possible for the voters to retrieve $\sum v_i$ through a grid search.




Figure 48: Trusted vote approach




Figure 49: Trusted vote sequence



It is easy to note how the aforementioned voting system shown in Figures 48 and 49 can be easily applied in the *AT* case. Indeed, *AT* can be seen as a "weighted" voting system, where the weight of each user depends on its energy consumption. Thus, the voting system can be used to implement a trusted certification algorithm of the community energy consumption and, as a consequence, of the applied tariffs. Under an operational point of view, every 15 minute *AT* performs the sequence described in Algorithm 1 interacting with *TrustedSum* application on each Strato device installed in LIC pilot (see Sections 2.1.2, 3.6.1 and 2.2 for details).

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.

Figure 50 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).




Figure 50: AT data flow

The data flow shown in Figure 50 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 and described in [35]. For this reason, a custom Go application was developed to run locally on the embedded devices the operations explained in Algorithm 1. On a Strato *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 of the nodes number, being the most of the resources used for the final decryption of the encoded sum. Consequently, the application scalability is guaranteed towards the community dimension.

5.4.7 Applications deployed in LIC: *Pre-paid scenario (PS)*

Similarly to *Metering*, *PS* saves information about production and consumption on the chain. *PS* is released as open-source on the Github platform¹⁴ ¹⁵. The main difference compared to *Metering* is that the energy data are completely tokenized using the official library provided by Cosmos. According to energy produced or consumed by a meter named *M* in a defined period, *M* mints or burns a corresponding amount of tokens.

To have *PS* properly functioning, periodically (e.g. monthly) the community administrator saves on the chain the settings that will regulate the application. They are reported in Table 11.

<i>convfactor_cons</i>	Consumed energy [Wh] to tokens to burn
<i>convfactor_prod</i>	Produced energy [Wh] to tokens to mint

Table 11: Main parameters used in PS application

Every 15 minutes, *PS* tokenizes the consumed and produced energies related to the last quarter of an hour. Consequently, the energy values are transformed in tokens related both to the energy consumption/productions and the power peak. Algorithm 2 reports in details how the energy is tokenized by *PS*.

Algorithm 2: PS operations sequence on a generic meter *M*

Get data about energy produced (*Ep*) and consumed *Ec* in the last 15 minutes

QUERY: Read PS parameters

QUERY: Read *M* tokens balance

newBalance = *balance*

bonus = *convfactor_prod* * *Ep*

penalty = *convfactor_cons* * *Ec*

newBalance = *newBalance* + *bonus* - *penalty*

TRANSACTION: Set the new token balance of *M*

5.4.8 Applications deployed in LIC: *PARITY Market (PM)*

During 2021 the activities of blockchain-based solutions for LIC mainly refer to the development and the deployment of *PARITY* decentralized application. *PARITY* (Prosumer AwaRe, Transactive Markets for Valorization of Distributed flexibility enabled by Smart Energy Contracts) is a H2020 project (864319 - 2019-2023 and LIC is participating in it).

PARITY aims to go beyond the traditional “top-down” grid management practices by delivering a unique local flexibility market platform through the seamless integration of IoT and blockchain technologies. By delivering a market for automated flexibility exchange based on smart contracts blockchain, *PARITY* will facilitate efficient and transparent local flexibility transactions and reward flexibility in a cost-reflective and symmetric manner, through price signals of higher spatio-temporal granularity based on real-time grid operational constraints and available DER flexibility. The *PARITY* solution will be demonstrated in 4 pilot sites around the EU (ES, CH with LIC, SE and GR) to validate its effectiveness across climatic, cultural and techno-regulatory conditions. The trials of the *PARITY* algorithms, among them the blockchain application that will be described in the following, will be performed in 2022.

Like all the blockchain solutions developed in LIC, the *PARITY* application, in the following named *PM*, is based on sidechains technology and has two main tasks. The former is the management and solution of

¹⁴<https://github.com/supsi-dacd-isaac/cosmos-apps>

¹⁵<https://github.com/supsi-dacd-isaac/cosmos-apps-handlers>



Local Energy Markets (LEM), played every quarter of an hour by some prosumers that are part of the EC and, consequently, nodes in the sidechain. The periodic solution of LEM, based on an Automated Market Mechanism (AMM) described in the following, entails a correspondent movement of tokens between the nodes. The latter is the management and periodic checking of generic Service Level Agreement (SLA) between a single prosumer and the sidechain administrator related to a set of KPIs (e.g. an internal temperature threshold that has never to be overcome). If a node is not compliant with its SLAs, then it has a penalty in terms of burnt tokens. In addition to the aforementioned main tasks, PM has also to store data about the forecast of energy production and consumption.

Before proceeding with the description in detail of the tasks described above, the roles that can be played by the different nodes in PM are introduced. This aspect is strictly connected to the access to the data managed by the sidechain and is crucial to understand how the application operates. In PM three different roles are defined and can be played by a node:

- **DSO**, responsible for managing on the sidechain the list of the prosumers belonging to the EC and to save every 15 minutes information about the state of EC grid for the following quarter of an hour. The periodic storing of this information is extremely important, being necessary to solve the market related to each: temporal slot of 15 minutes.
- **Aggregator**, which has to periodically define the LEM features (e.g. the nodes allowed to play the market) and SLA settings in accordance with the nodes (e.g. the temperature threshold related to a specific SLA). Besides, it solves sets of LEMs acting together with the prosumers.
- **Oracle/Prosumer**, which basically saves data about its energy consumption and production on the sidechain when it plays LEM and solves sets of LEMs together with the Aggregator. In addition, it stores on the sidechain the needed information related to the SLA with which the prosumer is involved.

It is important to note how there can be only one DSO and Aggregator in a sidechain. In accordance with the role played, a single node is allowed to modify only specific parts of the sidechain performing the proper transactions. The read-only access to the sidechain, based on queries, works similarly.

LEM running and solving: In PM every LEM has a duration of 15 minutes and can be divided into the following three phases:

- Phase 1: The LEM initial setting is provided by DSO and Aggregator nodes before the start of the related quarter of the hour
- Phase 2: The LEM settling is managed by the oracles participating to the LEM and the Aggregator nodes after the market end
- Phase 3: The LEMs solving is provided by the Markets Solver instances of the oracles participating to the LEM and of the Aggregator, as depicted in Figure 1. In accordance with the energy produced and consumed by each oracle, its token balance is updated. The algorithm that explains how the LEM solution works is described at the end of the paragraph.

It is important to note that the third phase has not to be necessarily performed every 15 minutes as the other. Indeed, more LEMs can be grouped and solved in sequence in order to minimize the number of transactions on the sidechain. Thus, it is useless to solve each market every quarter of an hour: it is more meaningful to have longer resolutions, e.g. a monthly basis. Figure 51 shows an example with the running of four LEMs in the period [12:00-13:00] of a generic day and their solution that, for simplicity reasons, is performed on an hourly basis. In this example, a sidechain constituted of 6 nodes with 4 oracles is considered. Orange transactions are performed by DSO node, green ones by Aggregator node, whereas the other colours refer to Oracles nodes.

Transactions 1 and 2 must be performed before the beginning of the period related to the LEM. Essentially, these transactions record in the chain the LEM main features, like the parameters needed to




Figure 51: LEM operation sequence

calculate the market energy prices. More specifically, transaction 1 stores on the sidechain the grid state signal and is performed by the DSO node. Instead, transaction 2 saves the other main LEM features and is run by the Aggregator node. Once the first LEM (LEM_Q1) is finished, transactions 3, 4, 5 and 6 are performed by the oracles to store on the sidechain their energy production/consumption during LEM_Q1. In addition, after the oracles' transactions, the Aggregator Node closes the LEM (i.e. transaction 7). Finally, after the end of the four LEMs (LEM_Q1 – LEMQ4), the nodes related to them (i.e. the Aggregator Toolset and the oracles), perform transactions to update their token balances. In case the production of an oracle exceeds the consumption in the hour [12:00-12:15], then its related prosumer has a token reward, whose amount depends on the energy prices calculated in the four LEM. Instead, in case of an exceeding consumption, the prosumer has a token penalty.

Considering the solution of a single LEM, e.g. LEM_Q1 in Figure 51, it is performed by PM application taking into account the price scheme introduced in section 3.3.

SLA running and solving: The management of the SLA/KPI is similar to the procedure aforementioned described for LEM. Indeed, the operations can be divided into the following three phases:

- Phase 1: The SLA initial setting provided by the Aggregator node before the starting of the period related to the specific SLA. Basically, SLA is a container of KPIs active in a specific period. Then, each KPI is defined by standard features, like the list of the oracles that must be compliant with it or the token penalty that has to be applied if necessary
- Phase 2: The SLA settling, managed by the oracles that must be compliant with the KPI
- Phase 3: The SLA solving, provided by the oracles participating to the KPI.

Similarly to the LEM case, the final phase is not performed when each SLA is ended. Specifically, many SLAs are grouped and solved together to minimize the number of transactions on the sidechain. There are two main differences between SLA/KPI and LEM managements. The former relates to the temporal aspect, whereas all the LEM in PARITY have a fixed life of 15 minutes, each SLA can have a different period of activation. The latter difference regards how the token balances are updated when a set of SLA is solved. Indeed, only penalties are applied in case of an oracle not compliant with a KPI, and there are no rewards like in the case of the producers participating in a LEM. Figure 52 shows an example with the running of two SLAs with an hourly duration in the period [12:00-14:00] of a generic day and their solution that, for simplicity reasons, is performed every two hours. In the example, a sidechain



constituted of 6 nodes with 4 oracles is considered, and each SLA has one KPI. Green transactions are performed by Aggregator node, whereas the other colours refer to oracle nodes.




Figure 52: SLA/KPI operation sequence

Transactions 1 and 2 are performed before the beginning of the period related to the SLA of the first hour, in the example in Figure 52 the period [12:00-13:00]. These transactions save the SLA/KPI features on the sidechain. More specifically, transaction 1 stores on the sidechain the SLA metadata, for example the activation period. Instead, transaction 2 saves the features related to the KPI. The two transactions are performed by the Aggregator node. Once the first SLA is finished, the nodes involved with the KPI save the related data in the sidechain (transactions 3, 4, 5 and 6). Finally, after the end of the SLA (two hours time window), the related oracles perform transactions if their token balances have to be updated. In the case shown in Figure 52 only two oracles are not compliant with at least a KPI and, consequently, they have to run the transactions 7 and 8.

Forecast management: The forecast storage on the sidechain is implemented basically to provide the related future energy prices for a single oracle/prosumer, considering equations 1. This information is periodically updated with the latest forecasts by the Aggregator node and the oracles every quarter of an hour. In PARITY project, one-day ahead forecasts are taken into account; as a consequence, the saved forecasts are composed of 96 elements. Figure 53 depicts how the oracles and the aggregator nodes have to interact with BA performing periodic transactions to update every 15 minutes the forecasts of the next 96 slots. Green transactions are performed by the Aggregator node, whereas the other colours refer to oracles ones.




Figure 53: Forecast management on the sidechain



Privacy management: The compliance with General Data Protection Regulation (GDPR) [15] is mandatory for the PARITY applications. It is well-known how data privacy has to be carefully taken into account in the development of blockchain solutions, being this technology intrinsically decentralized. From an operative point of view, PM preserves privacy by exploiting the pseudonymization strategy [40]. Thus, the private data that must be saved on the sidechain in order to properly run the energy market (e.g. the serial number of the smart meter owned by the prosumer) are properly anonymized and, consequently, PM is compliant with GDPR.

PM deployment and preliminary tests: The deployments and tests of PM started in the second half of 2021 and it is expected to continue until the end of 2022. LEM and SLA are continuously created, settled and solved on the sidechain and the latest forecast updated, similarly as shown in figures 51, 52 and 53. Figure 54 reports the structure of the Cosmos 5.4.4 sidechain currently active. The light blue circles represent nodes without specific permissions, instead the orange ones are the chain validators, i.e. the unique elements in the network allowed to create new blocks. Under the point of view of the roles of the nodes, *nuc01* is both the DSO and the Aggregator node; instead the other nodes are oracles related to a specific prosumer.




Figure 54: Sidechain structure

5.5 Application sustainability

We evaluated the sustainability of the blockchain-based local market implemented in LIC, in terms of resources allocations, power consumption and greenhouse gas emissions. The evaluation is based on the PARITY application (see Section 5.4.8), which is the one that requires the highest hardware resources and, as a consequence, highest energy consumption of all the sidechain applications tested in the LIC pilot and presented in Sections 5.4.5, 5.4.6, 5.4.7 and 5.4.8.

5.5.1 Hardware requirements

The application runs on the Strato devices, described in Section 2.1.2. For the analysis of the resource allocation, we monitored the following 4 KPIs:



- CPU loading
- RAM usage
- Disk occupancy
- Network usage

Table 12 reports the results obtained after 3 weeks of continuous running of LEM and SLA/KPI on the Lugaggia sidechain as previously explained.

CPU loading	< 10%
RAM usage	~ 210MB
Disk occupancy	~ 70MB/day
Network usage (upload)	~ 4MB/hour
Network usage (download)	~ 6MB/hour

Table 12: Hardware resources usage of PM application in LIC sidechain

While CPU, RAM and network usages do not have a significant impact on the Strato devices, the values reported in Table 12 show that the disk usage is the KPI that has to be taken into account carefully. For what concerns the deployment of PM application in LIC, the choice to use disks with a capacity of 32 GB on the Strato devices is sufficient to run various LEM and SLA/KPI during the tests in the pilots to fully validate the PARITY strategies. Regarding the connectivity, in LIC each Strato communicates using a 4G mobile dongle; no issues specific to the sidechain have been experienced with this type of connection, which can be easily affected by problems like interference and obstacles that can weaken the signal.

5.5.2 Energy consumption and greenhouse gas emission evaluation

The energy consumption and the carbon intensity of the PM application were estimated considering a CPU usage of 10%, as reported in Table 12 and a maximum power consumption of 6 W, according to the documentation of Raspberry Pi-3+¹⁶, the board used by Strato devices. The power consumption of the meter's optical interface was considered negligible, while for the 4G router that provides the Internet connection to the stratos, we conservatively assumed that it operated at full power 100% of the time. Therefore, this estimate can be considered a worst-case scenario.

Regarding the conversion from energy to CO₂eq emissions, we took into account the factor provided by electricitymaps¹⁷ for Switzerland during 2021. Table 13 reports the parameters described above and the obtained values for a sidechain node.

Application CPU usage	10	%
Raspberry PI3+ maximum consumption	6.0	W
Application power consumption	0.6	W
4G router power consumption	3.5	W
Yearly energy consumption	35.9	kWh
CO ₂ eq emissions conversion factor ¹⁸	111.7	g/kWh
CO ₂ eq emissions	4.01	kg

Table 13: Yearly energy consumption and CO₂eq emissions of PM application, per node.

The total yearly consumption per node under the worst conditions is 35.9kWh with a CO₂eq emissions of 4.01kg (most likely lower). Without considering the consumption of the 4G router, which is unfortunately

¹⁶<https://www.raspberrypi.com/documentation/computers/raspberry-pi.htmltypical-power-requirements>

¹⁷<https://app.electricitymaps.com>



difficult to estimate, the values would be 5.25kWh and 587gCO₂eq. The emission values are in any case relatively low thanks to the design choice not to use the Proof-of-Work consensus mechanism [33], which typically requires significantly higher CPU utilization and, consequently, significantly higher power consumption. Please refer to Sections 5.4.3 and 5.4.4 for more details about the selected blockchain solution.

Regarding the entire sidechain scalability, the tests performed in LIC showed a linear scale of consumption depending on the number of nodes constituting the sidechain.

It is also correct to say that grey energy related to the production, transportation and installation of the stratos, which certainly has an impact on the life cycle analysis of the blockchain solution, has not been considered. Unfortunately, this part of the analysis is beyond the competence of the project team.

6 Techno-economic performance analysis

In this section, we will evaluate the techno-economic performance of establishing a prosumer self-consumption community, depending on the presence of PV and batteries and the activation of flexible loads such as boilers and heat pumps. The goal is to identify which solutions make it attractive to set up a prosumer community. We also analyze the effect that the penetration of PV and batteries has on the gains, respectively savings, of users who already had a PV system.

The work was carried out in simulation. For this, a digital twin of the LIC community was constructed. The thermal model of the buildings was modeled with an RC circuit with one state and was identified from the data monitored by AEM's smart meters in this way:

1. R was calculated from the building energy signature, which, as a function of the average daily outdoor temperature, returns the energy consumed by the building. The slope of the energy signature allows the thermal resistance of the building to be derived
2. The thermal inertia C was estimated from the year of construction and building type.

Non-thermal electric loads consumption and domestic hot water usage profiles were generated using the load profile generator application ¹⁹.

The adopted pricing scheme is the one presented in 3.3.

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 scheme. 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.

6.1 PV and distributed batteries

We evaluated the effect of PV and home batteries in the LIC digital twin. For PV production, the database of the suitability of roofs for use of solar energy of the Federal Office of Energy ²⁰ was used, and the most suitable roof pitches were identified for each building. Using Meteonorm software ²¹, a typical meteorological year (TMY) was generated for the Lugaggia location, and the pvlib library²² was used to generate representative PV production profiles for each pitch of the roofs in the LIC community. Different PV penetration and battery scenarios were generated. PV was gradually increased until all buildings possessed some. The choice was made not to fill the roofs completely but to size the systems to achieve a certain level of self-consumption. This is a more realistic setting and similar to the current sizing practice for rooftop-mounted PV power plants w.r.t. installing as many kW as possible, as the

¹⁹<https://www.loadprofilegenerator.de>

²⁰www.sonnendach.ch

²¹<https://meteonorm.com/>

²²<https://github.com/pvlib/pvlib-python>

latter strategy is usually not economically optimal. In the case of a PV system without a battery, the self-consumption target was 40%; in the case of a PV system with a battery, the self-consumption target was raised to 80%. Pitches with a higher production potential were filled first. Of course, it is possible that the self-consumption target may not be reached because the roof was completely filled. Battery were sized using the empirical rule suggested by EnergieSchweiz in [41], which states that:

$$\text{Energy battery [kWh]} = \text{PV nominal power [kW]} \cdot 1.5 \quad (33)$$

Figure 56 shows the different scenarios that were generated. They are not evenly spaced, as buildings are added one after the other in a random way, and each of them has a fixed-size PV plant on the roof. Battery control algorithms lexicographically optimize self-consumption and power fluctuations (peak shaving) at the individual building level. A predictive algorithm is needed to perform peak shaving. The control algorithm is based on MPC with a 24-hour horizon. A machine learning-based predictor is used to predict the energy production and demand of houses in the control horizon. One year of operation

Figure 55: PV and battery penetration scenarios




Figure 56: PV and battery penetration scenarios

of the community was simulated and the pricing scheme presented in 3.3 was applied to calculate how much each prosumer would profit from being a member of the community in the different scenarios. The results are shown in Figure 57. We can see that when PV penetration is low, PV system owners (prosumers) have an advantage over consumers and benefit more from being members of a community. This is due to the scarcity of self-generated energy, which means that virtually all the energy produced by PV systems is sold within the community. On the other hand, since there is little self-produced energy available, consumers can buy little of it and consequently take less advantage of being community members than prosumers. As the PV installed within the community increases, the relative savings of producers and consumers level off. The abundance of self-generated energy means that there will more often be an opportunity for consumers to consume energy at a lower price than the power company, while prosumers will not always be able to sell their excess PV production within the community at a higher price, thus reducing their margin.




Figure 57: Relative savings due to community membership as a function of PV and battery penetration.




Figure 58: Total yearly savings in a community settings, with respect to a business as usual case with the same PV and battery penetration. The labels mean pv:[total installed PV power in kW]_bat:[total installed battery capacity in kWh]

Another interesting observation is that battery owners have more significant relative savings than their counterparts who own PV systems without batteries. However, it is essential to emphasize that this is a relative saving. In fact, if one looks at the absolute savings of users given by belonging to the community, this is instead greater in the case where users do not own a battery (Figure 58). This makes sense since the locally installed battery optimizes the self-consumption of the individual home, making it more self-sufficient and reducing the amount of energy its owner needs to buy from the outside. It can be concluded that, for those who have a battery that can be purchased off the shelf today (and therefore only optimizes their self-consumption at building level), it is still attractive to be part of a self-

consumption community with residential users, but in absolute terms the annual savings are less than for a user without a battery.

Let's analyze the total savings of simulated community members. In Figure 59, it can be seen that this peaks around a penetration of about 61.8kW of PV and then decreases as PV and battery penetration increases. This point represents the penetration for which there is generally a greater benefit in being part of a community than not being part of it and does not necessarily represent the economic optimum for community members. In fact, this is given by the sum of the bill savings given by purchasing a PV system and perhaps a battery, the initial investment in them, and only ultimately, by being part of a community.




Figure 59: Total savings due to community membership as a function of PV and battery penetration.

We delved into the question of how much community membership is profitable for PV systems owners. Figure 61 shows the annual bill savings for PV and battery owners in the community compared with the case where they did not have PV and batteries and were not part of a community. Figure 62, on the other hand, shows the relative additional savings on the annual bill given solely by being in a community versus not being in a community, defined as:

$$\frac{c_{SCCPV} - c_{BaUPV}}{c_{BaUPV} - c_{BaUNoPV}} \quad (34)$$

where c_{SCCPV} is the yearly electricity bill in a self-consumption community (SCC) setting with PV (and battery) installed, and c_{BaUPV} and $c_{BaUNoPV}$ is the bill in a business as usual (BaU) scenario with, respectively without PV installation.

It is pretty clear that being a member of a community is always a benefit. It is guaranteed by law and by the tariff scheme adopted. It can be seen, however, that as PV penetration in the community increases, this advantage, in terms of annual savings, decreases slightly, as we already pointed out previously. It can also be seen that for a battery owner, the benefit in terms of additional bill savings is significantly less than for someone who instead decides to install only a PV system without storage. Based on these assumptions, a property owner could most likely choose not to invest in a storage system, which is still very expensive and hardly profitable, but to become a community member, because this would significantly affect his savings without additional expenses.




Figure 60: Yearly balance of simulated users in the case in which they are not members of the community. Negative values: costs, positive values: earnings.




Figure 61: Total yearly savings in a community settings, with respect to a business as usual case in which no PV or batteries were installed.




Figure 62: Relative additional annual savings given by community membership, compared to the annual savings given by purchasing PV (and battery) in a business-as-usual scenario.

6.2 District battery

An alternative setup to having distributed batteries is to have a centralized one for the entire neighborhood, as is the case in the LIC pilot. Figure 64 shows the average savings for users in the LIC self-consumption community in the years 2020 and 2021.




Figure 63: Relative annual savings given by community membership for the LIC community members in the years 2020 and 2021.

Since the battery was not operated continuously and was broken for several months, these data are not representative of a situation in which the battery is available most of the time. For this reason, we decided to estimate the maximum potential savings the battery can give in the real LIC pilot configuration. The real battery power was subtracted from the power measured at the grid connection point of the community, and savings were calculated in the case in which the battery was absent. Losses were not considered during this operation, which adds a small error. Then the actions of an ideal battery of the same size as the one installed in LIC were calculated using the control algorithms to which we fed perfect forecasts. We evaluated the potential savings in the case in which the battery provided only a self-consumption optimization service and in the case in which it provided a self-consumption optimization service and a reduction in maximum peak consumption. Since AEM has a peak tariff, this makes a difference. In any case, it can be seen how the potential for user savings is relatively low. This is because the energy injected into the grid by the community is relatively little (see Figure 33). For this reason, the cost-benefit analysis was also conducted in simulation. The algorithms deployed in the LIC pilot are precisely the same as those run in simulation. Therefore, one can expect the same results in a real pilot, and the results obtained when the battery was operated confirm it.

To evaluate the financial performance of the installation of a battery by the community manager, it is crucial to note that there are two possible types of configurations when it comes to integrating the battery into the domestic market.

1. The battery takes an active part in the market. This means that the community members can sell and buy their energy to and from the battery at the p_s^{P2P} and p_b^{P2P} prices




Figure 64: Relative annual savings given by community membership for the LIC community members in the year 2021 with ideal battery. SC: self-consumption, SC+PS: self-consumption and peak shaving.

2. The battery does not take an active part in the market. This means the community members cannot sell and buy their energy to and from the battery. In this case, the battery is used by the administrator exclusively for self-consumption optimization at the community coupling point.

The choice of the market configuration strongly depends on who invested in the battery. Even if the community members did not invest in the battery, option 1. was adopted in the pilot project case, and the community members fully profit from the battery's presence.

We decided to simulate a community in which all dwellings are provided with a PV system sized on 40% self-consumption at the individual house level. This results in a total of 114.3kWp of installed PV. The battery size was then varied from 15 to 240kWh in several growing steps. The battery control algorithm lexicographically optimizes community self-consumption and power fluctuations (peak shaving) at the coupling point of the community. Like in the previous case the control algorithm is based on MPC with a 24-hour horizon. A machine learning-based predictor is used to predict the energy production and demand at the coupling point of the community in the control horizon.

As can be seen in Figure 65, community users take advantage of the presence of the battery in the shared configuration. As the size of the battery increases, the savings also increase. However, it can be seen that beyond 180kWh, the battery does not allow users to save more. Figures 66 and 67 show the annual earnings of the simulated community administrator in the case where the battery participates and does not participate in the market, respectively.

To test whether the investment in a battery is profitable, we estimated a battery life of 10 and 15 years (as battery life depends both on cycling and calendar ageing) and operation and maintenance costs of \$60 per year [42], assuming them as constant with the battery size in the considered range. We then calculated the maximum price per kWh that a battery could cost to break even on the investment after 10 years of operation, which corresponds to the cost per kWh of the battery that results in a net present value (NPV) of zero after 10 years. Lower costs would result in profits for those who invested in the battery, while higher costs would result in losses. A weighted average cost of capital (WACC) of 2% was




Figure 65: Relative additional annual savings given by community membership as a function of district battery size, in the case in which the battery participate into the market.

considered in the calculations. The results are shown in Figure 68.

One can see how a 60kWh battery, like the one installed in LIC, would have to cost less than 300CHF/kWh to be profitable, even when all community members have PV. The battery in the LIC project costed about 750CHF/kWh, making it decidedly unprofitable, also in light of the lower PV penetration.

If the battery life is extended to 15 years, the maximum price is higher, but still below what can be actually found on the market. The break-even prices are summarized in Table 14.

Battery capacity [kWh]	break-even price for a life of 10 years [CHF/kWh]	break-even price for a life of 15 years [CHF/kWh]
15	313.8	448.8
36	303.8	434.5
60	285.7	408.7
120	235.8	337.4
180	184.5	263.9
240	140.3	200.7

Table 14: Break-even prices of the district battery as a function of its size and lifespan.




Figure 66: Admin yearly profits (top) and share of the yearly total savings with respect to the BaU scenario (bottom), as a function of district battery size, in the case in which the battery participates in the market.




Figure 67: Admin yearly profits (top) and share of the yearly total savings with respect to the BaU scenario (bottom), as a function of district battery size, in the case in which the battery does not participate in the market.




Figure 68: Battery profitability analysis. Top: Additional yearly profit given by the installation of a battery in the community. Bottom: Maximum battery cost in CHF/kWh, to return the investment within 10 years.

6.3 Flexible loads

The evaluation of the optimization potential given by controlling thermal loads was also done in simulation. Preliminary simulations showed that the optimization potential given by controlling loads is less than that of controlling one or more batteries. Field experiments confirmed our conclusions. Simulation allows the effects of flexibility actuation to be evaluated under controlled conditions, which cannot be done in the field.

Again, a community in which all buildings are equipped with a PV system was chosen to be simulated. Three cases were evaluated:

1. No control of thermal loads
2. Control of water heaters
3. Control of water heaters and heat pumps

One year of operation was simulated. The control algorithms optimized users' costs, knowing they were community members. Due to its iterative nature, no distributed control approach was used in the case, as it would have been too slow to simulate a full year. However, the shorter simulation results presented in [43] showed only limited advantage of coordinated over uncoordinated control. Also in this case an MPC control strategy based on realistic forecasts was used.

Figure 69 shows the balance at the end of the simulated year for the community members.




Figure 69: Yearly balance of simulated community members as a function of thermal loads control. Negative values: costs, positive values: earnings.

One can notice how controlling the thermal loads slightly improves the financial performance of the community members. This is especially true when heat pumps are controlled.

6.4 Coordination

Different coordination mechanisms for batteries were tested in simulation and deployed in the field, a more in depth analysis of the simulation results can be found in [43]. Since the coordination required more computationally expensive algorithms, the simulation were carried out for a single month - the monthly of July - and for a single PV configuration - LIC actual real configuration.



In 6.4.1, the economical impact of the batteries is considered. The central battery is evaluated both while operating alone and in coordination with the water heaters, while the distributed batteries will be considered as a separate case. In 6.4.2 the same is done for the water heaters. In both sections, the "controlled" cases are evaluated against a baseline case with no batteries present and water heaters always on. This baseline case, in turn, is evaluated financially both for the "business as usual" case (bau) and for the case in which agents are part of an self-consumption community (scc).

6.4.1 Batteries

The results are summarized in Table 15. Here follows a brief explanation of the various scenarios and control strategies found in the table.

The central battery was considered both operating alone or in coordination with the boilers (cases *cb only*/*cb+boilers*). Two different control algorithms for governing the simulated distributed batteries were tested:

1. **Explicit coordination with grid constraints:** in explicit mode, the batteries can address each other with messages; grid constraints are enforced by introducing a steep fee when they are about to be violated. This ensures the respect of the limits, but potentially generates a net cash flow from the agents to the administrator that cannot be mitigated for theoretical reasons (case *db expl lims*)
2. **Implicit coordination** In implicit mode the batteries forecast the internal price signals and aim to reduce their own overall cost (case *db impl*).

It emerges from the results that the mere fact of being a member of a community generates considerable savings for end-users. In the cases with the central battery on the kindergarten, the PCC obtains significant economical advantage due to being the admin and owner of the battery. As expected, the introduction of distributed batteries, by encouraging self-consumption both locally and at the community level, improves the financial situation of the end-users and also constitutes an advantage for the community administrator over the baseline cases. The explicit coordination method in which agents communicate with each other is the most effective from a financial point of view.

The results with individual data points are also plotted in Figure 70. In the case in which distributed batteries are present, their 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. In the case in which the central battery is present, on average the end users profit slightly less than in the case in which the batteries are distributed, but the administrator of the community profits more (see Table 15). This is fair because in this case the administrator of the community has to return the investment in the battery and bare the associated economic risk.

	user costs	PCC costs	PCC balance
baseline bau	1476.1	0.0	0.0
baseline scc	993.1	600.4	-392.6
cb only	985.6	334.7	-650.9
cb+boilers	969.9	339.5	-630.3
db expl lims	821.7	360.3	-461.3
db impl	847.5	387.2	-460.3
db no coordination	927.1	530.7	-396.3

Table 15: Total costs (in CHF) for the end users and the administrator of the community using the LIC pricing scheme for the case of battery control. The keywords describing the control types are explained above.




Figure 70: Difference in the monthly costs using LIC pricing scheme, with respect to the *baseline bau* case, for the simulated month of July.

6.4.2 Water heaters

Water heater control also leads to improved financial figures, although the effect is not as pronounced as with battery control. The results are summarized in Table 16 and Figure 71. 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.

	user costs	PCC costs	PCC balance
baseline bau	1476.1	0.0	0.0
baseline scc	993.1	600.4	-392.6
boilers expl lims	971.3	588.2	-383.1
boilers impl lims	969.6	574.3	-395.2

Table 16: Total costs (in CHF) for the end users and the administrator of the community using the implicit costs redistribution mechanism for the case of boiler control.




Figure 71: Difference in the monthly costs using LIC pricing scheme, with respect to the *baseline bau* case, for the simulated month of July. Water heaters control case.

6.5 Grid support potential

The central battery can be used, not only to maximize community self-consumption, but also to reduce its impact on the power grid. The control algorithms used in LIC actively seek to reduce power withdrawal and injection into the grid at the community's coupling point. This not only reduces losses within the community but also optimizes against power tariffs. The disadvantage of power tariffs nowadays is that they are based on the maximum monthly peak, which makes the control task extremely challenging. It is enough to make a mistake in predicting consumption and emptying the battery too much one day per month to nullify the efforts, which instead resulted in effective peak reduction on other days.

The battery, however, effectively contributes to stabilizing the power at the coupling point and consequently reducing voltage swings and losses within the community, as shown in Figures 72, 73, and 74, which report the values extracted from the load flow calculation performed during the simulations. Losses reduction could be further improved, if one would site the district battery differently. At the moment the battery is installed in the kindergarten, which is at the end of a long line.

6.6 Delay investments in grid refurbishment through storage and DSM

In this section we estimate the reduction in cable aging given by the reduction in the current flowing through it, due to battery actions and active control of thermal loads. We quantify by how much cable life is extended and grid refurbishment delayed, and calculate how much this saves on grid refurbishment.

Cable degradation is mainly due to the degradation of the cable insulation. The main aging mechanisms are degradation due to high temperatures and due to electrical stress. The authors in [44] propose a




Figure 72: Empirical PDF of the active power at the community's coupling point, as a function of district battery capacity.




Figure 73: Empirical PDF of the mean L-N voltage inside the community, as a function of district battery capacity.

multi-stress model, which can be expressed as:

$$\frac{L}{L_o} = \frac{L_1}{L_0} \frac{L_2}{L_o} \dots \frac{L_N}{L_0} G(S_1, S_2, \dots S_N) \quad (35)$$

where L is the life of the cable under multiple stresses and $L_1 \dots L_N$ are single stress lives for stress $S_1 \dots S_N$ and G is a correction function taking into account that life under multiple stresses is usually higher than that derived by simple multiplicative laws. Since the single-stress electrical and temperature effects on cable life can be modeled by respectively modeled with an inverse power law

$$\frac{L_E}{L_0} = \left[\frac{E}{E_o} \right]^{-n} \quad (36)$$

and with a ratio of Arrhenius laws

$$\frac{L_T}{L_0} = \exp(-E_a \theta / R) \quad (37)$$

where E_a is the activation energy associated to the insulation material, R is the gas constant and $\theta =$




Figure 74: Yearly losses inside the community, as a function of the district battery capacity.

$\frac{T - T_{ref}}{TT_{ref}}$ where T_{ref} is the reference temperature. The authors propose the following aging law:

$$\frac{L}{L_0} = \left[\frac{E}{E_o} \right]^{-(n-b\theta)} \exp(-B\theta) \quad (38)$$

However, the interaction between electrical and thermal aging is usually disregarded (see for example the International Atomic Energy Agency (IAEA) Cable Ageing technical report [45] and [46]). Furthermore, authors in [47] compared different aging models for cable insulations, and reported only a small difference between models considering and disregarding the electrical field and temperature interaction. Therefore, we just considered temperature-induced degradation in the present study.

Insulator temperature computation Insulator temperature can be computed as a function of ambient temperature and the instantaneous cable's power. If the exact temperature at insulator-conductor interface is needed, an equivalent resistance coefficient must be used, which can be estimated starting from geometrical considerations on the cable stratigraphy and type of insulator used, as shown for example in [48]. In this study we have used a simpler method, as suggested in [49], where the cable operating temperature is found by using the Joules loss formula and knowing the maximum operating power and its corresponding temperature:

$$T(t) = T_a(t) + (T_{max} - T_{ref}) \left(\frac{P(t)}{P_{max}} \right)^2 \quad (39)$$

where $T(t)$ and $T_a(t)$ are the cable operating temperature and the ambient temperature at the current time, T_{max} and T_{ref} are the maximum operating cable temperature and reference temperature (20 °C), while $P(t)$ and P_{max} are the current and maximum cable power. Equation 39 is then plugged into 37 to retrieve the power-dependant instantaneous degradation.

We didn't estimate the impact on the aging of the secondary substation, since we just measured the power at the LIC's PCC. Estimating the aging of the transformer would have required the whole aggregated power, which includes other distribution cabinets, which are not part of LIC.

6.6.1 Numerical results

The presented methods to estimate the acceleration aging factors for the cables have been applied to the simulations of controlled batteries, boilers and HPs in LIC, using realistic forecasts for the power at the PCC and the internal community prices. The simulations refer to a one year period. For the battery,

	whole year	T > T ref	cost reduction [CHF]
battery: 15 kWh	2.22%	1.92%	1076
battery: 30 kWh	3.92%	3.37%	1866
battery: 60 kWh	6.60%	5.64%	3056
battery: 120 kWh	9.84%	8.49%	4479
battery: 180 kWh	10.87%	9.30%	4872
battery: 240 kWh	10.93%	9.41%	4923
battery: 0kWh, Boil:12, HP:0	1.63%	1.33%	748
battery: 0kWh, Boil:12, HP: 13	3.83%	2.39%	1337

Table 17: First column: relative life extension considering all the timesteps. Second column: relative life extension considering only times in which the cable's temperature is higher than the reference. Third column: avoided costs.

the boilers and the HPs, we simulated a lexicographic control in which the battery optimizes at first for energy costs of the community, and in a second optimization performs peak shaving using the optimal costs retrieved in the first simulation as constraints. Each agent has two forecasts, one for the positive and one for the negative part of the 15 minutes energy at PCC, E_c , E_p . These are used to estimate future internal prices for the SCC, using equations 1. Note that this setting is equivalent to the implicit control strategy used for the simulations in section 6.4.1, but with an additional lexicographic step.

We simulated the control for increasing sizes of the central battery, ranging from 0 kWh to 240, in steps of increasing size. For all the simulations we considered the maximum hosting capacity for PV, under the assumption that PV system sized on 40% self-consumption at the individual house level. This equal to 114 kWp. In figure 75 the key quantities for the computation of the cable degradation are shown for a period of three days. The top plot shows the power at the PCC, the blue line indicating what would have happened without the batteries in place. The red line shows the resulting power under the real battery operations in the considered period. The batteries try to perform peak shaving flattening the overall consumption. During the first day we can see the batteries charge from the installed PV (when the PCC is injecting back into the public grid) and then soon discharges themselves; this is likely due to forecasting errors overestimating the future consumption. The middle plot shows the cable's temperature obtained by equation (39), while the last plot shows the degradation factor computed from equation (37). It can be seen how even small deviations in the cable temperatures can result in high changes in the degradation factor, as expected. From the plot is clear that the battery operated in a peak-shaving mode has a high impact in reducing the cable temperature and, as a consequence, to increase the life of the cable. The effect of the real batteries on the aging fall in between the baseline case and the operation of ideal batteries with perfect knowledge on the future. The results on the total monitoring period are reported in table 17, as increase in operational life, that is, called L_{base} the life of the base case (without battery, boilers and HP optimization):

$$\frac{L}{L_0} / \frac{L_{base}}{L_0} - 1 = \frac{L}{L_{base}} - 1 \quad (40)$$

The first column reports the relative improvement of lifespan over the base case, considering all the monitoring period. The second column considers just the periods for which the cable's temperature was above the reference operating temperature, which is the temperature for which the standard life of the cable is computed, and under which no life shortening can be considered, which in this case was fixed to 15 °C. The third column reports the avoided costs computed as:

$$C_{avoided} = C_{tot} \frac{L - L_{base}}{L} \quad (41)$$

We can use the estimation over cables and transformer degradation reduction from table 17 to infer a total cost reduction induced by the battery, boilers and HPs operations. Considering a similar effect over the operating temperatures for all the lifespan of the cable, we can estimate an increase of lifespan in the range of 1.3%-9.4% for the cables. A realistic estimation of the main cables of the considered grid




Figure 75: Example of time series for the computation of cable degradation. Top: instantaneous power for the PCC cable. Middle: corresponding cable temperature computed using equation (39). Bottom: degradation factor for the considered cases, computed using equation (37)

(including installation costs), based on a consultation with AEM DSO, is about 57250 CHF (20250 CHF cables and 37 kCHF installation). Under these assumptions, we can estimate a cost reduction between 748 CHF and 4923 CHF, if considering degradation to start after the reference operating temperatures. Note that assuming that all the cables degrade at the same pace of the main one results in a conservative estimation of the avoided costs, since the normalized power of the main cable at each timestep is the usually less than the average normalized powers of the connected cables.



7 Stakeholders, policy and legal evaluation

This work package focused on exploring the perspective of the households engaged in the Lugaggia Innovation Community (LIC), in order to inform later processes aimed at the replication of the LIC material and regulatory solutions to other contexts and to support their large-scale diffusion.

Particularly, WP5 explored the knowledge, attitude and perceptions of the household members of the LIC community, with the aim of analysing their level of engagement with the community, the direct and indirect benefits of the LIC model, as well as possible negative outcomes and barriers to its implementation. More in details, our analysis tackled the following set of research questions (RQs):

- RQ1:** Which is the level of engagement and satisfaction with the LIC community, its governance structure, and its web portal tool by LIC members?
- RQ2:** Do community identity factors, social norms, knowledge of the energy system, and attitudes towards renewable energies affect attitude towards the LIC community and its level of engagement with it?
- RQ3:** Which are the perceived direct effects of membership to the LIC community, regarding both energy availability and its cost?
- RQ4:** Are attitudes towards renewable energies and household reported energy and environmental behaviour affected by membership to the LIC community?

To enrich our understanding on the implications of the LIC community from the perspective of its members, we also looked for possible differences — if any — between general LIC household members and the sub-sample of LIC households that are equipped with a PV power plant. Coherently with the recent literature, we refer to them as “prosumers” [50], while we use the term “consumers” for general LIC members.

7.1 Methodology

To address our research questions, we opted for a “before-after” analysis based on panel data collected via a survey questionnaire. We designed a three-wave survey investigation, aimed at collecting *ex-ante*, *in itinere* and *ex-post* information on attitudes, perceptions towards self-consumption communities (and, more specifically, towards the LIC community), as well as on its impacts.

The *ex-ante* survey was launched right before the beginning of the LIC pilot activities on the field, once LIC household members had been identified and had confirmed their engagement in the community, but before they could experience anything related to community membership (October 2019). The *in itinere* survey was launched exactly one year after, when pilot activities were halfway (October 2020), while the *ex-post* survey was launched at the end of the LIC pilot activities (October 2021).

In all the three cases, the same questions were offered, with only minor modifications aimed at guaranteeing the right tenses were used, such as for instance using the past for events that had already happened during pilot activities on the field. Doing so, data collected during the three survey waves can be directly compared, and a longitudinal assessment of any changes having occurred throughout the pilot activities can be performed.

Furthermore, the *ex-ante* survey included basic information on socio-economic parameters, while the *in itinere* and *ex-post* surveys also included questions aimed at assessing the LIC web portal and the LIC billing process, besides questions aimed at detecting possible changes related with the household composition (number of individuals permanently living in the house), the installation of electricity-related plants and devices (PV power plant, electric battery storage, electric vehicle, plug-in hybrid vehicle), or the energy retrofitting of the home (change in the heating system, thermal insulation of the building’s roof, windows, or facade). The latter piece of information allowed us to identify possible critical changes preventing before-after comparisons for the specific household.



7.1.1 Questionnaire administration and response rate

Part of the survey questions dealt with the entire household and part of them dealt with specific attitudes, behaviours or values of the specific survey respondent within the household. To guarantee consistency as much as possible, we asked each household to identify a survey respondent (“responsible for the household”) and to guarantee she/he would remain the same at each wave. We cannot guarantee this has always been the case, but we are confident this has happened, since the invitation to answer the survey was always sent via email to a specific email address (the email of the respondent), which was provided by the household at registration and signing of the contract to join the LIC community.

The questionnaires were offered online and delivered via the open source software tool for online surveys “Limesurvey”. Each survey wave was announced by a news post in the LIC newsletter, which was periodically sent via email to all LIC members, followed by a customised message sent to each LIC member. The latter directly offered a customised link with a token for unique identification of the respondent, automatically generated by the Limesurvey tool, aimed at guaranteeing pseudonymisation of the responses. Each survey questionnaire was kept open for six weeks and automatic reminder emails were sent by Limesurvey every two weeks, to those who had not answered yet.

The *ex-ante* survey was launched before the beginning of the LIC pilot activities on the field, on October, 13 2019. By the start of December 2019, 17 of the 19 LIC users had responded ((89% of the total). The *in itinere* survey was launched exactly one year after. By the start of December 2020, 18 of the 19 LIC users had responded (again, 89% of the total). Finally, the *ex-post* survey was launched on early December 2021, and by mid January 2022, again 15 of the 19 household had responded (79% of the total). When merging the three datasets, however, it appears that only 14 of the 19 LIC users (74% of the total) have answered to all the three questionnaires. Here we provide a description of participants to LIC activities based on the data collected via the *ex-ante* survey, and then address our research questions by analysing changes over time. For this purpose, we consider differences emerging throughout the whole LIC project, thus we directly compare *ex-ante* and *ex-post* responses to the same questions, by using the maximum number of responses available for both survey waves (responses by 14 LIC members). When instead we analyse questions related with a single survey wave, we use the maximum available number of responses. For each comparison, tests of statistical significance were performed. Since the sample is very small, we cannot rely on asymptotic normality of the variables’ distribution. Therefore, we always performed visual inspection of the values of the variables, to check for normality, coupled with statistical normality tests such as the Shapiro-Wilk test. Whenever normality was found, we performed paired t-tests. When this was not the case, we used the non-parametric Wilcoxon signed-rank test.

To collect the respondent’s opinion, for many factors under analysis we used a 5-point Likert-scale, such as “1– Completely disagree” and “5– Completely agree”. When identifying mean values of the responses to such questions by the whole LIC sample, we made the following assumptions: mean values lower than 2.5 correspond to the “disagree part” of the scale; mean values between 2.5 and 3.5 are neutral/moderate; mean values larger than 3.5 correspond to the “agree part” of the scale.

7.1.2 Limitations of the research design

We would like to stress that, with such a research design, which does not include a control group, no causal analysis can be performed. Namely, detection of possible changes in the LIC members’ responses over time cannot be causally attributed to participation in the LIC community. We performed in fact an exploratory study, aimed at identifying relevant topics or situations that deem further in-depth analysis based on a strict policy-analysis scheme. Furthermore, we are well aware that, even in the cases in which statistical significance was found, any generalization of the obtained results should be carefully made, mainly due to the small sample size under this investigation. Despite such key limitations, the current analysis provides a preliminary contribution to the knowledge of self-consumption communities and their members, opening-up lines of investigation for future activities.



7.2 Characteristics of LIC households and houses

Before dealing with our specific research questions, we summarize the profile of the LIC members, based on the socio-economic information we collected via the *ex-ante* survey.

The person responsible for the household (in terms of the energy contract) tends to be male (12 respondents out 17), older than 35 years-old. Level of education and profession are highly variable among the sample. LIC households are composed by a minimum of two to a maximum of five people. Only four of the households included in the LIC community include underage people.

All houses are owned by the LIC members; the large majority of them are independent houses, with only two semi-detached houses. Furthermore, the large majority of LIC households has already installed some kind of passive energy systems, such as walls and windows insulation. Five of the houses have also installed active solar energy systems (PV power plant), but none of them at the start of the project had a battery storage system. Interestingly, electric vehicles (EVs) are also available in the LIC sample: two out of the 17 responding households in fact own an electric vehicle.

7.3 Level of engagement and satisfaction with the LIC community (RQ1)

We started by exploring the level of engagement and satisfaction with the LIC community by its members. For this purpose, survey questionnaires included batteries of 5-point Likert questions about the attitudes towards self-consumption community projects, the way self-consumption communities are created and managed (governance aspects), as well as support tools such as the LIC web portal.

7.3.1 Attitudes towards self-consumption communities

The questions part of this series are inspired by [51] and [52], with minor modifications to fit the LIC case. All respondents show high appreciation of the idea of self-consumption communities, in general, and of LIC community in particular: they are highly in favour of such communities and are proud to be part of LIC (Table 18). Over time, data show a tendency towards the decrease of such an appreciation—however with values remaining higher than the 3.5 threshold we indicated above—and in any case statistical significance of the observed differences is only found for the general appraisal about “being in favour” of self-consumption communities.

The feeling of being part of a community with other people and households is instead less pronounced: there is an intermediate agreement about “having a lot in common with other LIC members” and “feeling attached to them”. Furthermore, also in this case a decreasing trend is observed over time, particularly for the feeling of attachment to other members, whose decrease over time is statistically significant at the 5% level.

It also seems that, at the start of pilot activities, membership to LIC was a topic of conversation with other members of their circles, though over time the novelty of the topic tended to decrease, and LIC members were less prone to “disseminating” the LIC initiative to their contacts (though the decrease is not statistically significant).

Analysis of the responses of the prosumers (n=4), not reported here, shows slightly higher values. However, due to the very low sub-sample size, there is no statistical difference at any significance level, compared with the mean values observed for all the LIC members. Similarly, no statistically significant differences emerge in the evolution of the responses by the prosumers, throughout the evolution of the pilot activities.

We also explored the reasons why LIC members accepted to participate in the LIC self-consumption community (Table 19): they mostly valued the fact that the community is part of a research project and that it is promoted and managed by their electricity provider AEM. These two elements favoured their trust in the project, even more than the fact that LIC was promoted by their municipality and that they were guaranteed by AEM that they would not have been requested to pay higher electricity bills than



Table 18: Attitudes towards self-consumption community (SCC) projects.

1 - I completely disagree; 5 - I completely agree

n = 14	Ex-ante (T1)		In itinere (T2)		Ex-post (T3)		Wilcoxon test (T3-T1)	
	Mean	SD	Mean	SD	Mean	SD	V	p-value
I am in favour of energy SCCs	4.9	0.4	4.8	0.4	4.4	0.9	10	0.0947*
I feel a sense of pride in being part of the Lugaggia SCC	4.5	0.9	4.5	0.8	4.3	0.8	15	0.3741
I have a lot in common with other members of the Lugaggia SCC	3.5	0.8	3.3	0.5	3.0	0.9	25	0.0651*
Being a member of the Lugaggia SCC is an important part of who I am	3.1	1.2	3.1	1.1	2.9	1.1	18.5	0.4879
I feel attached to the other members of the Lugaggia SCC	3.1	0.8	3.1	0.9	2.6	0.9	15	0.0477**
I am proud to be part of a SCC	4.3	1.0	4.1	0.8	3.8	1.1	6	0.1814
I like talking about the Lugaggia SCC in presence of others	3.9	0.9	3.9	1.0	3.7	0.9	3	0.3458

Statistical significance: * 10%, ** 5%, *** 1%

they would have payed under the “traditional” tariff scheme by AEM. Despite levels of agreement on intention to join self-consumption communities generally decrease between the *ex-ante* and the *ex-post* survey, no statistically significant differences are found.

Furthermore, we explored more general preferences regarding hypothetical participation in other self-consumption communities (Table 20). In the latter case, we considered different governance structures of hypothetical self-consumption communities, in terms of their initiator (promoter) and manager, and asked LIC members their intention to join such hypothetical self-consumption communities. Their preference goes to communities that are both promoted and managed by public institutions, such as the canton or consortia of municipalities, followed in order of preferences by their municipality alone, private companies and, in the last position, other citizens. In the small sample of LIC members, it seems therefore that the “cooperative” approach that characterises many examples of totally bottom-up renewable energy communities [52], is not perceived as highly appealing: impartial and independent public institutions are largely favoured.

The same questions were made in the *ex ante* and *ex post* survey. Despite levels of agreement on reasons for joining LIC and intention to join hypothetical self-consumption communities generally decrease between the *ex-ante* and the *ex-post* survey, no statistically significant differences emerge between the two periods. Furthermore, and highly relevant, responses show that if guarantees on the maximum price of electricity currently offered by AEM were lacking, the intention to join a self-consumption community would largely decrease (from 4.21 points to 2.50 points in a 5-point Likert scale). This indicates that, according to the small sample of LIC members, for self-consumption communities to be actually appealing, they should be managed not only with the aim of maximising the self-consumption rate of PV electricity production: monetary aspects have to be included in the objective functions to be maximised within the community’s electricity management algorithms.

Finally, responses by prosumers (not reported here but available upon request) indicate slightly higher values for all the questionnaire items. Particularly, they show a slightly higher agreement for joining a SCC also if guarantees on electricity prices are lacking (average response equal to 2.75 points out of



5, compared to the average response by LIC members equal to 2.50 points). Also, all prosumers highly value the fact the LIC is part of a research project (average response equal to 5.00): this suggests that prosumers perceive themselves as innovators and pioneers, and thus are more inclined, interested, and open to research activities than “average households”.

Table 19: Conditions for participation in the LIC self-consumption community (n=14).

1 - I completely disagree; 5 - I completely agree

I accepted to join the LIC SCC because...	Ex-ante (T1)		In itinere (T2)		Ex-post (T3)		Wilcoxon test (T3-T1)	
	Mean	SD	Mean	SD	Mean	SD	V	p-value
It is promoted by my municipality	4.14	0.95	—	—	3.64	1.15	4.5	0.2740
It is promoted by my electricity provider	4.36	1.15	—	—	4.00	0.88	15	0.3951
It is managed by my electricity provider	4.36	1.15	—	—	4.00	0.88	15	0.3951
It is a research project	4.86	0.36	—	—	3.93	1.00	36	0.0154
I had guarantees on the maximum price of electricity	4.21	0.80	—	—	3.93	1.00	41	0.4929

Statistical significance: * 10%, ** 5%, *** 1%

Table 20: Conditions for participation in a hypothetical self-consumption community (n=14).

1 - I completely disagree; 5 - I completely agree

I would join a SCC...	Ex-ante (T1)		In itinere (T2)		Ex-post (T3)		Wilcoxon test (T3-T1)	
	Mean	SD	Mean	SD	Mean	SD	V	p-value
Promoted by private citizens (neighbours)	2.86	1.46	—	—	2.36	0.93	15.5	0.3387
Managed by private citizens (neighbours)	2.29	1.14	—	—	2.14	0.95	28	0.5328
Promoted by private companies	3.21	1.25	—	—	2.86	1.17	31	0.3346
Managed by private companies	3.29	1.27	—	—	2.86	1.10	54.5	0.2264
Managed by my municipality	4.07	1.21	—	—	3.14	1.23	61	0.0842
Promoted by a public institution (consortium of municipalities, canton, other)	4.21	1.05	—	—	3.57	1.16	24	0.1033
Managed by a public institution (consortium of municipalities, canton, other)	4.21	1.05	—	—	3.43	1.09	32	0.0544
Even without guarantees on the maximum price of electricity	2.50	1.34	—	—	2.57	1.45	16	0.8211

Statistical significance: * 10%, ** 5%, *** 1%

7.3.2 Level of engagement and satisfaction with LIC web portal

The *in itinere* and *ex post* surveys also provided us with the opportunity to evaluate the level of engagement by the LIC members with the LIC web portal. The web portal was in fact aimed at increasing the awareness of the LIC members about the electricity autarky level of the self-consumption community and the interactions between LIC members and the grid in terms of exchanges of electricity.



How much was the web portal used and how was it assessed by the LIC members? About one third of LIC members did not access their account in the web portal (see Figure 76). This was in some cases due to an explicit lack of interest in the web portal, while in most cases LIC members refer of problems of communication about its meaning or how to access it (forgotten credentials, unknown link), or about the lack of time to access it (see Figure 77). Even though the web portal was presented in at least two e-newsletter posts and personal credentials to access it were distributed to each LIC member via a dedicated email message (check this is correct), this feedback shows that additional communication and/or targeted support would have been useful to favour LIC members to explore the web portal at least once, thus allowing them to fully grasp its potentials.

Have you ever visited your personal page on the LIC webportal?




Figure 76: Number of LIC members who accessed the web portal (*in itinere* and *ex post* surveys).

Why haven't you visited the LIC webportal?




Figure 77: Reasons for not accessing the LIC web portal (*in itinere* and *ex post* surveys).

Through the *in itinere* survey we also explored the evaluation of the LIC web portal by the LIC members who had accessed it. Table 21 summarizes the outcome. The portal was appreciated as a useful (average evaluation 4.57 out of 5), educative (4.43), interesting (4.43), easy to use (4.57), and easy to access (4.43) tool. On the other hand, it was judged as not particularly appealing from the graphical

How frequently do you visit the LIC web portal?




Figure 78: Frequency of access to the LIC web portal (*in itinere* and *ex post* surveys).

point of view (2.71) and not sufficiently intuitive and intelligible (2.30). As a possible enrichment of the offered features, a LIC member suggested it should also offer comparisons with the “normal situation”, namely the traditional electricity provision by the AEM GDO, and include a feature showing the overall amount of energy that the whole LIC community has avoided to feed or request from the grid.

Responses to the same questions by LIC prosumers confirm higher familiarity with monitoring tools aimed at providing quantitative feedback on energy consumption and production, as well as their interest for such kind of tools and the related data. They all attributed the web portal the maximum score (5.00 points), regarding how informative/educative, useful, and interesting it is — which is about 0.5 points higher than LIC members as a whole. Also, they judged the web portal to be more intuitive than the average LIC members did (with a difference of about 0.8 points). It is likely that, since the installation of their own PV plant, prosumers were used to regularly monitor the amount of electricity produced by the plant, as well as the amount of electricity self-consumption, in order to both assess the financial return to their investment, and also to increase their self-esteem and perception of self-efficacy in contributing to the energy transition. The LIC web portal is therefore a particularly valuable tool especially to prosumers.

Table 21: Level of satisfaction with the LIC web portal by the LIC members who at least accessed it once. Evaluation collected during the *in itinere* survey (n=8).

The LIC web portal is...		In itinere (T2)	
		Mean	SD
1- Difficult to access	5 - Easy to access	4.43	1.13
1 - Difficult to use	5 - Easy to use	4.57	1.13
1 - Graphically appealing	5 - Graphically not appealing	2.71	1.60
1 - Intuitive	5 - Unintelligible	2.43	1.51
1 - Poorly informational	5 - Educative	4.43	0.98
1 - Useless	5 - Useful	4.57	0.79
1 - Boring	5 - Interesting	4.43	0.98
1 - Below my expectations	5 - Above my expectations	3.57	0.79

Finally, we measured the effects of the LIC web portal in terms of possible increases in awareness about the household's energy consumption and energy literacy, as well as its possible contribution to the increase in self-consumption rates, by changing electricity consumption behaviours at home. Responses by LIC members show that the web portal has increased awareness on knowledge of the consumption patterns of one's own household (on average, 3.71 points out of 5.00), though its contribution to energy



literacy and increase of knowledge of the energy system in general is quite limited (3.14 points). LIC members also show a slight intention to change their household electricity consumption patterns, with the aim of favouring the exploitation of higher shares of solar energy. In Section 7.6 we will come back to this point and check if such a stated behaviour intention is actually confirmed by the daily energy consumption behaviours they declare to put into practice.

Also in this case, responses by prosumers show a larger intention to favour high self-consumption rates (5.00 points out of 5.00, instead of 3.71). This again confirms that prosumers joining LIC have higher intrinsic motivation and interest towards energy topics than the average LIC members. The same conclusions can be drawn by considering that prosumers are less interested than LIC members as a whole in getting information about saved money and carbon emissions (respectively, 3.50 and 3.00 points out of 5.00, compared with 4.14 and 4.14 by average LIC members): to the prosumers, the web portal is interesting *per se*, since it provides feedback on their intrinsic motivation to support the energy transition, and it does not need to provide additional monetary or climate feedback to trigger them to action.

Table 22: Assessment of the LIC web portal. Evaluation collected during the *in itinere* survey (n=8).
1 - I completely disagree; 5 - I completely agree

How much do you agree with the following sentences?	In itinere (T2)	
	Mean	SD
Thanks to the LIC web portal, my awareness of the energy consumption of my household is higher	3.71	1.11
Thanks to the LIC web portal, I learnt something more on the way the world of energy works	3.14	0.69
I check data on the LIC web portal and try to move my energy consumption in the hours when production of solar energy is higher	3.71	0.95
On the LIC web portal I would like to monitor my energy consumption data expressed in monetary units (Swiss francs)	4.14	0.69
On the LIC web portal I would like to monitor how much CO_2 emissions have been saved	4.14	0.90

7.4 Factors affecting attitudes towards the LIC community (RQ2)

This research question aims at investigating which factors might influence individual attitudes towards self-consumption communities and, ultimately, the decision to join one of them. To get a preliminary understanding of the relevant factors, we explored scientific literature about closely connected community-based renewable energy projects [53],[54], [55], [51], [56], [57], [58], defined as “formal or informal citizen-led initiatives which propose collaborative solutions on a local basis to facilitate the development of sustainable energy technologies and practices, producing local benefits” ([52], p. 613).

The literature analysis suggests that the following factors may influence individual attitude towards energy self-consumption communities:

- *Community identity*, that is “the feeling of attachment to the community, taking pride in the community and having friends within the community” ([59], p. 797, cited in [57])
- *Social norms*, namely the perception by an individual that a behaviour must (or must not) be performed, regardless of its outcome for the individual [60]. Note that, following [57], these can also be influenced by community identity: higher community identities imply higher social norms
- *General Trust*, which, following [51] (p. 2657), we interpret as “the nature and quality of the relationships between people and organisations within a community”
- *Attitude towards renewable energies*, namely the level of support to renewable energy sources [52], which in turn depends on the individual’s pro-environmental self-identity (the perception of



oneself as an individual who cares for the environment, according to the conceptualisation by [52]).

We thus explored each of them, with a dedicated set of questions in all the three survey questionnaires. Here we report responses to the *ex ante* and *ex post* surveys and compare them to verify if any changes occurred during the LIC activities.

7.4.1 Community identity

To assess the level of community identity, we use the construct by [57], reported in Table ???. LIC Members show a moderate sense of belonging to the local community where they live in, which slightly decreases over time. At the start of LIC activities, in fact, they tend to talk about their community as a "great place to live", while at the end of LIC activities such an agreement decreases, and such a change is significant at the 10% significance level.

Overall, though we have not quantitatively computed correlations due to the limited LIC sample size, the moderate feeling of belonging to their community seems to show that "community identity" was not among the key factors driving the LIC members' positive attitudes towards self-consumption communities.

Table 23: Sense of belonging to the local community where LIC members live in (n=14).

1 - I completely disagree; 5 - I completely agree

	Ex-ante (T1) Mean	SD	In itinere (T2) Mean	SD	Ex-post (T3) Mean	SD	Wilcoxon test (T3-T1) V	p-value
I feel strongly attached to the community I live in (people and places)	3.43	0.76	3.64	0.74	3.5	0.76	20	0.7897
There are many people in my community whom I think of as good friends	3.43	0.94	3.36	0.93	3	1.04	29	0.1236
I often talk about my community as being a great place to live	4.36	0.74	4.5	0.65	3.86	0.77	30	0.0969*

Statistical significance: * 10%, ** 5%, *** 1%

7.4.2 Social norms

Also social norms were measured based on the construct proposed by [57], which focuses on energy consumption and renewable energies. Responses, reported in Table 24, show that the social context of LIC members did not lead them to feel social pressure to save energy or use energy from renewable sources. Nevertheless, LIC members indicated high social acceptance for their participation in local energy projects. Interestingly, prosumers indicate very close answers to those of LIC members as a whole. It seems therefore that in this case both electricity prosumers and regular consumers share exactly the same social context: it appears prosumers are still "pioneers", which constitute the exception in their social context, and are not yet part of a niche of other prosumers in which using renewable electricity and saving energy has already settled as a dominant social norm.

Considering the evolution over time, social norms, measured in terms of peer pressure and expectations by peers, did not change significantly between the first and last questionnaire waves. The picture emerging from the responses is that social norms about energy saving and use of renewable energies are still relatively low and not really perceived by the LIC members (both prosumers and consumers) – even though their social circles support participation in energy projects. This therefore suggests that social norms did not play a key role neither in driving the LIC members' positive attitudes towards self-consumption communities nor in triggering their engagement within the LIC community.



Table 24: How social norms about energy consumption and renewable energies affect LIC members (n=14).

1 - I completely disagree; 5 - I completely agree

	Ex-ante (T1)		In itinere (T2)		Ex-post (T3)		Wilcoxon test (T3-T1)	
	Mean	SD	Mean	SD	Mean	SD	V	p-value
Many of my peers use energy generated from renewable energy sources	2.93	0.47	3.07	0.73	3.14	0.53	4	0.4076
Saving energy is expected by my peers	2.93	0.83	3.00	0.78	2.86	0.66	8.5	0.8902
People I care about would approve my participation in local energy projects	4.14	0.86	4.00	0.88	3.50	1.16	50.5	0.1220

Statistical significance: * 10%, ** 5%, *** 1%

7.4.3 General trust

To assess trust, we refer to the question proposed by the 2010 “Environment” module by the International Social Survey Programme (ISSP) [61], also used by [57]. To assess the general attitude towards the concept of “trust”, the question investigates whether the respondent believes that most people can generally be trusted, or that care is always needed when dealing with other people. All LIC members showed a neutral position with respect to trust, which also remained constant over time and across the three survey waves. Prosumers showed slightly higher values, indicating a marginally higher openness to trusting people in general, though their responses remain within the range of neutral positions.

The obtained responses suggest that general trust in other people has not played a key role in either shaping attitudes towards the LIC community or favouring higher engagement by the LIC members.

Table 25: General attitude towards the concept of “trust” by LIC members (n=14).

1 - I completely disagree; 5 - I completely agree

	Ex-ante (T1)		In itinere (T2)		Ex-post (T3)		Wilcoxon test (T3-T1)	
	Mean	SD	Mean	SD	Mean	SD	V	p-value
Generally speaking, would you say that most people can be trusted or that you can't be too careful in dealing with people?	3.29	1.07	3.29	1.20	3.21	1.05	31.5	0.7055

Statistical significance: * 10%, ** 5%, *** 1%

7.4.4 Attitude towards renewable energies

To assess attitudes towards renewable energies, we refer to the constructs by [52] and customise them in order to mostly focus on PV renewables and thus specifically refer to the LIC case.

On average, LIC members showed a very positive attitude towards renewable energies, as well as PV, agreeing that more renewable energy projects should be developed [M=4.8, SD=0.6], that more photovoltaic plants must be developed [M=4.4, SD=0.9], and that photovoltaic plants offer an answer to the climate crisis [M=4.3, SD=0.7]. As expected, prosumers showed even higher scores, confirming a definitely positive attitude towards renewables and PV.

It appears, therefore, that a positive attitude towards renewable energies characterises all the LIC members. Though the current research design does not allow us to identify causal relationships, whose



identification would require a counter-factual control group, this suggests that LIC members' attitude towards renewable energies could be a good predictor of their positive attitudes towards self-consumption communities and of their engagement with the LIC community.

Table 26: Attitudes towards renewable energies (n=14).

1 - I completely disagree; 5 - I completely agree

	Ex-ante (T1) Mean	SD	In itinere (T2) Mean	SD	Ex-post (T3) Mean	SD	Wilcoxon test (T3-T1) V	p-value
More renewable energy projects should be developed	4.8	0.6	4.9	0.4	4.6	0.6	7.5	0.4237
More photovoltaic plants must be developed	4.4	0.9	4.5	0.7	4.5	0.9	6	0.7656
Photovoltaic plants have a critical landscape impact	2.5	0.9	2.3	1.0	2.3	1.1	35	0.374
Photovoltaic plants offer an answer to the climate crisis	4.3	0.7	4.1	1.0	3.6	1.2	47	0.4416

Statistical significance: * 10%, ** 5%, *** 1%

7.5 Perceived effects of membership to the LIC community (RQ3)

The three survey waves allow us to investigate the perceived effects of the LIC community on the availability of energy and its overall cost: do LIC members think that, after joining the LIC community, the quantity and quality of electricity available to them has changed? If so, how do they think it has changed? And what do they think about the cost of electricity? Do they think it remained equal or decreased, compared with the past, when their electricity was provided by the AEM DSO? We exclude the possibility of an increase in electricity costs compared with the past, since the LIC members were all guaranteed they would not have experienced any increase in the cost of electricity compared with the electricity bills they would have paid as customers of AEM. The contract signed between each LIC member and the LIC community in fact states that, on varying their electricity consumption, any possible differences between the electricity bills issued by the LIC self-consumption community and the bills they would have received by AEM if they had remained AEM customers, would have been covered by AEM — and this only holds for increases in billings: any decrease in billings after having joined LIC, would have been in favour of the LIC members themselves.

The set of questions we used to investigate the perceived effects of LIC membership is reported in Table 27. Responses to questions in the *ex ante* survey reflect the prior expectations by LIC members, before they start interacting with the LIC community; *in itinere* and *ex post* responses to the same questions reflect instead their perceptions based on the actual LIC experience: by comparing them, we can therefore get an indication of how much expectations differed from reality.

Ex ante responses show that LIC members had neutral expectations about the cost of energy under the self-consumption community: on average they thought their electricity bills would not be less expensive than in the past [$M=3.86$, $SD=0.77$]. Also, they believed they would not experience more water shortages [$M=1.71$, $SD=0.91$ for statement "My household might experience hot water shortages"] or higher black-out risk than in the past [$M=1.79$, $SD=0.97$ for statement "My household will run the risk of more black-outs or electricity shortages than in the past"], and would be able to heat their home as much as they liked, just like in the past [$M=4.00$, $SD=1.04$]. In this case, prior perceptions by prosumers were aligned to the answers by LIC members as a whole.

By comparing responses over time, it seems that, thanks to the real-life experience of the LIC community, the perception of its positive impacts got strengthened: the fear of hot water shortages or black-outs is lower, and the perception that homes can be heated as much as LIC members would like to, is higher. In all cases, however, the differences are not statistically significant.

To the opposite, instead, in the *ex post* survey the perception to consume renewable energies has



Table 27: Perceived effects of membership to LIC community (n=14).

1 - I completely disagree; 5 - I completely agree

Joining the Lugaggia SCC...	Ex-ante (T1)		In itinere (T2)		Ex-post (T3)		Wilcoxon test (T3-T1)	
	Mean	SD	Mean	SD	Mean	SD	V	p-value
The electricity bill of my household will be/were less expensive than in the past	3.86	0.77	3.36	1.01	2.79	1.12	70	0.01425**
My household might experience/experienced hot water shortages	1.71	0.91	1.50	0.76	1.93	1.14	15	0.7193
My household will mostly consume/consumed renewable-based electricity	3.71	0.99	3.21	1.05	3.21	1.37	30	0.0969*
My household will run/ran the risk of more black-outs or electricity shortages than in the past	1.79	0.97	2.00	1.18	2.00	1.11	12	0.7921
My household will consume/consumed more renewable-based electricity than in the past	4.43	0.76	4.29	0.91	3.93	0.92	15	0.0533*
I will be/was able to heat my home when and how much I will like to/liked to, exactly as in the past	4.00	1.04	4.36	0.93	4.43	0.76	13	0.2760

Statistical significance: * 10%, ** 5%, *** 1%

decreased, and in this case the difference is statistically significant. The prior expectation was that most of the LIC members' electricity demand would have been covered by PV electricity, while the third survey shows LIC members believe a lower part of their electricity demand to be covered by PV. To explain such a change, we make the hypothesis that a "reality-check" via the LIC web portal data has led LIC members to think that a relevant share of their electricity consumption was still covered by the electricity mix offered by the grid. Alternatively, the decrease in the reported share of renewables might be due to the fact that LIC membership has become a "routine": the novelty effect might have decreased, and LIC members might have partly forgotten about the origins of the electricity they were consuming. Differently than average LIC members, however, such a decrease is not observed among the prosumers: their perception of mostly consuming renewable-based electricity, and to do so more than in the past, remains high also in the third survey.

Furthermore, in the *in itinere* and *ex post* surveys the perception that electricity bills were less expensive than in the past has decreased as well, and the difference between the first and the third survey is statistically significant at the 5% level. This means that the amounts of the actual bills received by the LIC members were actually comparable with previous bills —or, at least, that they were perceived as comparable—and not lower than them. The same change is also observed among prosumers, in this case with statistical significance at the 10% level (which is also due to the very small sample size of prosumers, equal to 4 households).

Overall, the collected data tend to show that average LIC members perceive no relevant differences between their situation as a "traditional" AEM customer, before joining the LIC community, and their situation as a LIC member: neither regarding the amount of electricity available to address their needs, nor regarding the cost of electricity, nor even regarding the composition of the electricity mix they consume. Therefore, it seems that transition to the LIC community has occurred without any negative impacts, compared with the previous situation. This is also highly correlated with the fact that attitudes towards the LIC community remain high also in the *ex post* survey (see Table 18).



7.6 Effect on attitudes towards renewable energies and environmental behaviour (RQ4)

We finally investigated whether, after participation in the LIC community, any changes in the attitude towards renewables and/or their reported energy and environmental behaviour could be observed. If so, we could suppose that LIC membership has increased the energy and environmental awareness by its members, thus potentially triggering a change in their reported behaviour. Also in this case, we remark that caution would be needed, since the lack of a counter-factual control group would not allow us to assume causal relationships: any observed change in behaviour could be due to unobserved factors other than membership to LIC community.

Against this background, we first checked possible changes over time in the attitudes towards renewables and, more specifically, photovoltaic energies. As shown in Table 18, such attitudes did not significantly change between the three survey questionnaires. This is also because LIC members already had a very positive attitude towards renewable energies, as well as PV, since the very beginning (they highly agree that more renewable energy projects and photovoltaic plants need to be developed, that photovoltaic plants offer an answer to the climate crisis, and that the latter have no critical landscape impact).

Anyway, we observe a statistically significant increase (at the 5% level) in the level of agreement by the LIC members to the statement that “LIC would benefit if its members would install new PV plants”: before the start of LIC, the level of agreement was intermediate, equal to 3.36, while at the end of LIC it had increased by about 0.5 points out of 5 (see Table 28). This could be regarded as a consequence of the increase in awareness of the potential of PV and its contribution to a future low-carbon and self-sufficient energy system.

Table 28: Perceived benefits by membership to self-consumption communities (n=14).
1 - I completely disagree; 5 - I completely agree

The Lugaggia SCC would benefit if...	Ex-ante (T1)		In itinere (T2)		Ex-post (T3)		Wilcoxon test (T3-T1)	
	Mean	SD	Mean	SD	Mean	SD	V	p-value
Its members install new photovoltaic plants	3.36	1.01	3.57	1.28	3.86	1.17	70	0.0143**
Its members install more battery storage capacity	3.14	1.35	3.36	1.39	3.79	1.31	15	0.7193
Its members install new electric heating systems (e.g. heat pumps)					3.86	1.03		

Statistical significance: * 10%, ** 5%, *** 1%

Regarding changes in reported behaviour, we explored both the LIC members’ “electricity behaviour” at home, namely all the practices that they perform within the household and that have an impact in terms of electricity consumption, and their environmental behaviour, namely key practices that households perform and that have an impact on the environment. In both cases, we looked at the evolution of such reported practices over time.

Regarding electricity consumption practices, we investigated the frequency at which LIC members perform energy saving behaviours, such as washing laundry at low temperatures, turning standby off, cooking with lids on the pots, or also turning down the thermostat to reduce heating demand — a relevant electricity consuming behaviour for those households that are heated via heat pumps. The full list of investigated behaviours, inspired by the work of [62], is reported in Table 29, where the frequency of performing each behaviour is reported on a 1 (never) to 5 (always) scale. Please, note some of the behaviours are expressed as energy saving behaviours, some other are expressed as energy wasting behaviours. Responses have to be properly assessed in each case and, when needed, re-coded, if one wants to get a coherent indication about the frequency of energy-saving behaviours.

Overall, average LIC members reported intermediate values for most of the practices we investigated:



they perform the energy-saving behaviours neither too frequently nor too rarely. Only exceptions are the use of dishwasher and washing machine, which are mostly used only when fully loaded, and the habits of turning the lights off when leaving the room, turning TV off if no-one is watching it, and closing the water when brushing teeth. Furthermore, no relevant differences emerge by only considering responses by the prosumers.

Table 29: Reported electricity behaviour by LIC members (n=14).
1 - Never; 5 - always

How often do you perform the following activities?	Ex-ante (T1)		In itinere (T2)		Ex-post (T3)		V	Wilcoxon test (T3-T1) p-value
	Mean	SD	Mean	SD	Mean	SD		
Turn down/off heating before leaving for winter holidays	2.29	1.07	2.43	1.09	2.57	1.22	9	0.4403
Wash laundry at lower temperatures (30 or 40 °C)	2.71	0.47	2.93	0.83	2.71	0.91	7.5	1.0000
Turn off standby on appliances	2.36	1.15	2.79	1.05	2.64	1.15	1.5	0.2652
Use the mashing machine in the evening/night	2.29	0.83	2.29	0.83	2.29	0.91	10.5	1.0000
Adjust room temperature according to room's usage (e.g. heat/cool less in unused rooms)	2.86	0.86	2.86	0.95	2.86	1.03	10.5	1.0000
Cook covering the pots with lids	2.93	0.83	2.79	0.89	3.07	0.92	3.5	0.7103
Use the dishwasher when it is not fully loaded	1.50	0.52	1.71	1.07	1.43	0.65	9	0.7656
Let the water run while brushing teeth	1.50	0.76	1.50	0.65	1.79	0.80	0	0.0719*
Use the dishwasher in the evening/night	2.07	0.73	2.71	0.73	2.21	0.70	5	0.5716
TV is on and no-one is watching	1.29	0.47	1.29	0.61	1.50	0.94	5	0.5716
Turn off the light when leaving a room	3.93	0.27	3.71	0.47	3.71	0.61	3	0.3711
Use the washing machine when it is not fully loaded	1.79	0.80	1.71	0.91	1.50	0.52	21	0.2402
When buying an electrical appliance, consciously pay attention to its energy consumption	2.86	0.95	3.07	0.92	3.14	1.03	4	0.4076

By comparing the responses over time, no relevant changes emerge: for some questions, a slight increase in the frequency of energy-saving practices is found, while for some others a decrease in the frequency of energy-saving practices emerges. In all cases, however, the observed changes are not statistically significant, apart for the behaviour regarding brushing one's teeth, which however goes in the opposite direction than expected: responses show a slight increase, significant at the 10% level, in the average frequency of brushing one's teeth with the water running. If warm or hot water is used, this would imply an increase in energy demand —and specifically of electricity, in case of heat pumps or electric (water) heating systems. The limited amount of such an increase, however, does not allow to think of rebound effects, such as those for instance observed by [63] or [64] in domestic contexts after the introduction of solar PV plants. Future research could benefit by the quantitative analysis of the overall electricity consumption by LIC members before and after the activation of the LIC community, in comparison with a similar control group of households, via a panel data regression model. This would allow to identify whether possible rebound effects (or, to the opposite, energy saving effects) were



actually induced by LIC membership.

Finally, we also note that LIC membership has not triggered a change in the patterns of use of key domestic appliances, such as the dish-washer and the washing machine: we would have in fact expected to observe a decrease of their use during the night, justified by the desire to increase the rates of self-consumption of photovoltaic electricity produced during the day and exchanged via the LIC community. Reported behaviour, instead, indicates the opposite trend — though, again, observed differences are not statistically significant. LIC members in fact reported that they are now using more often the dishwasher at night, while use of the washing machine during the night is declared to have remained constant.

We finally explored the broader environmental behaviour by the LIC members. The aim in this case was to spot possible “spillover effects”, namely changes in practices that have an environmental impact, but are outside the direct realm of electricity consumption in households. If an increase in behaviours and practices that are beneficial to the environment is observed between the *ex ante* and *ex post* survey, such a change could be related with the LIC community — though, as already indicated, we could not simply conclude it was *caused* by membership to the LIC community.

The questions we considered to identify presence of spillover effects are taken by the 2010 IPSS Environment module by the International Social Survey Programme [61], and refer to recycling, consumer choices for organic products, reducing car and plane use, saving water, and adopting a vegetarian diet (see Table 30).

Table 30: Reported environmental behaviour by LIC members (n=14).

1 - Never; 5 - Always

How often do you perform the following activities?	Ex-ante (T1)		In itinere (T2)		Ex-post (T3)		Wilcoxon test (T3-T1)	
	Mean	SD	Mean	SD	Mean	SD	V	p-value
Sort glass or tins or plastic or newspapers and so on for recycling	3.93	0.27	4.00	0.00	4.00	0.00	0	1.0000
Buy fruit and vegetables grown without pesticides or chemicals	3.00	0.55	3.14	0.53	2.86	0.53	14	0.4840
Cut back on driving a car for environmental reasons	2.00	0.55	1.86	0.86	1.93	0.83	12	0.8241
Reduce the energy you use at home for environmental reasons	2.50	0.76	2.29	0.61	2.50	0.76	5	1.0000
Save water for environmental reasons	2.57	0.76	2.29	0.61	2.57	0.51	5	1.0000
Avoid buying certain products for environmental reasons	2.71	0.73	2.54	0.78	2.64	0.63	9	0.7656
Eat meat during main meals (lunches and dinners)	2.21	0.70	2.23	0.60	2.14	0.36	4	0.7728
Travel by car or motorbike	3.21	0.70	2.86	0.66	2.93	0.83	12.5	0.2031
Renounce to a flight for environmental reasons	1.71	0.91	2.00	0.88	2.21	0.89	8	0.1758
Use car- or bike-sharing	1.07	0.27	1.36	0.50	1.14	0.36	0	1.0000

According to the *ex ante* survey, average LIC members report intermediate values for most of the environmental-friendly behaviours we investigated: just like their reported electricity behaviours, they perform environmental-friendly behaviours neither too frequently nor too rarely — with a tendency to report reducing car or plane use less frequently than they report reducing meat consumption, or buying organic food or environmental-fair products, or even sorting their waste. Also in this case, no relevant differences emerge by considering the responses by prosumers alone.

Furthermore, a comparison over time shows no statistically significant change among such behaviours.




Therefore, even though we observe slight fluctuations in the responses between the *ex ante* and *ex post* surveys, we cannot conclude that a change has actually happened for reasons other than by chance. We conclude, therefore, that participation in the LIC community seems not to have induced relevant positive spill-over effects in environmental behaviours outside the domain of home energy consumption.

Finally, again following the IPSS Environment module, we also explored membership to environmental protection groups, or support to their activities, or activism for environmental issues. The responses by LIC members, reported in Figure 79, show that, between the first and the last survey, an increase in environmental-related action has occurred among LIC members. The number of LIC members who are also members of an environmental protection association increased, just like the number of those who signed a petition and the number of those who have financially supported environmental associations. Also in this case, such a change cannot be causally attributed to membership to the LIC community. Future research could, however, explore whether there is any connection between the observed increase in support to environmental activism and LIC membership and, if so, how such a support relates to the observed lack of personal, direct engagement in environmental behaviour change.



Ex ante survey (n=14)



Ex post survey (n=14)




Figure 79: Reported support to environmental topics by LIC members (*ex ante* and *ex post* surveys).

7.7 Policy and regulatory framework

Swiss law currently allows to set up a self-consumption community (legally a Zusammenschluss zum Eigenverbrauch, ZEV) like LIC, only if one is directly connected with their neighbors without using DSO's



public grid. In the case of LIC, the portion of the network has been formally unbundled from the rest of the grid by AEM, and its costs are covered by the internal tariff (see Section 3.3, Table 5). In the actual context, the LIC solution is not easily replicable for two main reasons:

1. All the users need to participate in the community. In the LIC case, the municipality and the DSO were involved in the project, which made it easier to convince all the users of the selected distribution cabinet to participate. In the case of a less known and trusted private company, it would be difficult to set up an energy community in a neighborhood, as some users would likely refuse to participate. Rewiring them to ensure they have a direct connection with the DSO's grid would most likely cost too much and hinder the entire project.
2. The DSO is not required to rent parts of its network to third parties, so not everyone who wants to set up in the community in Switzerland is entitled to the same treatment.

Unbundling parts of the grid to lease them to third parties also raises issues of equity. In Switzerland, grid costs are socialized, and all users of the same DSO with less than 50MWh per year are subject to the same grid tariffs. Hypothetically, if the DSO decided on a case-by-case basis, it might be more willing to rent or sell expensive network pieces with long lines serving low-consumption users. In this case, the costs per kWh of this particular portion of the network are higher than average, which would allow the DSO to reduce tariffs. On the other hand, it could be more reluctant to lease or sell short lines serving high-consumption users because this would cause it to raise its grid tariffs.

In addition, in the current regulatory setting, it is very difficult to manage the dynamic evolution of the community. Adding or removing users is very difficult and requires interventions on the network, which discourages the operation of community expansion. AEM received multiple applications from other Lugaggia residents living near the community but had to decline them because they were not connected to the proper distribution box.

7.8 The evolution of the LIC self-consumption community since its start

Although this is not particularly noticeable in the evolution of responses to the three surveys, in practice LIC users have been particularly active in the past two years. Since the beginning of the project:

- A new 8kWp PV plant was installed
- Two members have inquired with AEM with the intention of installing a PV plant
- Two users have replaced their oil-based heating system with heat pumps.

It is likely that at least some of these changes were motivated by the additional benefits of being in a community. It seems that, in practice, belonging to a community has made users more attentive to their consumption and the topic of energy transition in general.



8 Conclusions and recommendations

8.1 The perspective of the pilot owner

LIC was able to demonstrate the potential of self-consumption communities in conjunction with smart algorithms to reduce grid injection. These findings are particularly important for DSOs, which will face an increasing share of renewable electricity feed-in to the grid in the near future due to residential distributed generation. In fact, energy communities can help reduce and defer expensive grid interventions at a relatively low cost, being the first action to take for a more efficient distribution. However, while it was shown that the district battery plays an important role in increasing self-consumption, unfortunately it was found that battery systems are not yet financially sustainable if used to maximize self-consumption. Based on these results, AEM is convinced of the potential of energy communities and will establish more self-consumption communities, in particular to identify sustainable business models for battery applications, better understand the convergence between electricity and heat, and leverage smart grid infrastructure.

8.2 Centralized vs decentralized control

Both proposed solutions were shown to be able to control the battery and thermal loads. We will analyze the control of thermal loads and battery separately.

8.2.1 Thermal loads

Both solutions encountered several difficulties related to the control of thermal loads. Mostly, there were underestimated difficulties in disaggregating the loads from the readings at the main house meter. Another problem was related to the inaccuracy of the flexibility metadata. Very often, the nameplate data of the loads (flexibility connected to the meter's relay, nominal power, ...) provided to the algorithms were incorrect. It was often not clear what kind of load was being controlled; for example in the case of a heat pump that also produces hot water, it was not always clear how it reacted to the control signal. More and more heat pumps are equipped with an inverter and can modulate their power, which makes them particularly difficult to disaggregate. In any case, the simulation work also emphasized how controlling thermal loads leads to relatively small savings for community users, which makes it clear that the system costs per controlled unit must be low in order to return on the investment, since the benefits are also low. The decentralized solution is pretty expensive, in terms of:

- **Hardware costs:** In this case a PC based on the raspberry Pi CM connected to the smart meter using an optical interface. In total more than 300CHF, but this could be drawn down considerably if one decided to industrialize the product
- **Operating costs:** In order to communicate with each other and coordinate, and to run the blockchain market, the single units need broadband data communication, which in the case of the pilot project was solved by connecting the PCs to the 4G network
- **Installation costs:** The hardware needs to be installed close to the smart meter. The installation is relatively simple, but installation costs are particularly high in Switzerland.

The centralized solution, instead, does not require any installation cost at the end users' premises, assuming that the loads are already connected to the ripple control, as it is in the case of AEM and many other DSOs. The costs of the centralized solution are limited to the costs of the IT infrastructure for data acquisition and control. Since some of these costs are fixed and independent of the number of controlled flexibilities, the centralized solution becomes more and more economically attractive as the number of controlled loads increases. In the case of AEM, the solution is applied to all flexibilities on the network, which were previously connected to the company block, which considerably draws down the cost per



unit. However, it is not a simple solution to deploy. The integration of the centralized control algorithm with the smart metering infrastructure is required, which makes the solution not always applicable and dependent on the availability of data from the meters and the ability to set the schedules of their relays remotely and effectively. The amount of data transmitted to the centralized control algorithm is large, and the DSO's communications network must be able to support it. If future smart meters had computing capabilities similar to those of the computers used in the pilot project, this would greatly reduce data traffic. Complex operations requiring quite a lot of data (such as disaggregation) could be carried out on board, and a decentralized control solution would again become attractive.

Another evolution that would certainly positively impact the quality of control would be to increase the controllability of thermal loads, particularly heat pumps. At the moment, the tested systems can only force them off, but thanks to the thermal storage offered by the buffer (or the building itself²³), heat pumps have far greater flexibility than can be achieved with the current system. On this point, the SmartGridReady initiative²⁴, which aims to develop standard interfaces with flexible loads that allow two-way communication and consequently finer control, should be welcomed. Of course, it would mean that locally there would be a need to directly connect to the heat pumps, and the decentralized solution is the best suited for that.

8.2.2 Battery

As for the battery, the issue is different. Both solutions require additional hardware installation; in fact, even in the case of centralized control, it was necessary to install a meter capable of communicating the power reading to the community coupling point in pseudo real-time. And, since the battery can cost several thousands of francs and needs to be installed and connected to a broadband network anyway, the extra cost of installing a local PC to control it impacts much less the overall costs of the solution per unit installed. The centralized solution is based on a simpler reactive algorithm, while the decentralized solution uses a predictive and deterministic MPC-based control algorithm. The second solution showed better performance in peak shaving. However, it is difficult to quantify how much improvement potential there is in a general case, as it depends greatly on the quality of the forecast. In general, the larger the community, the easier it is to predict its production and load, the better the MPC control would perform.

The analysis of battery installation's financial feasibility showed that battery installation is still not economically advantageous for the community administrator under the price conditions applied in the LIC community (see Section 6.2). However, it was pointed out in Section 6.6 how installing a battery helps extend the life of network components, especially cables and transformers. When the battery is operated to reduce the oscillations quadratically at the community's coupling point, the life of cables can be extended. Should DSOs propose rates on power not solely related to the maximum monthly peak but also to the ability to flatten the consumption and production profile over the entire month, this could make the installation of a battery more profitable. A suggestion would be to extend the local energy market by adding a tariff component directly dependent on the power at the coupling point. This is the most effective way to further reduce the variance of the aggregated power profile and thus reducing voltage fluctuations in the local grid without crafting ad-hoc energy tariffs for the single user, as shown in Equation 42.

$$p_P = \beta E \quad (42)$$

Where β is power tariff parameter and E is the total energy produced/consumed by all the market participants or the total energy at the point of the grid determining the grid state, e.g. a LV/MV transformer. The price p_P , in cts/kWh, varies every 15 minutes as a function of E . With this model, if the community is consuming a lot, the users who contribute to the peak with their load will have to pay proportionally to it. On the contrary, if, at the same time, other users inject power into the community and therefore contribute to reducing the consumption peak, they will be paid for their contribution.

²³<https://www.aramis.admin.ch/Grunddaten/?ProjectID=38727>

²⁴<https://smartgridready.ch/>



8.3 The Evolution of Policy and Regulatory Framework

As we discussed in section 7.7, the current law in Switzerland does not lend itself to encouraging the establishment of self-consumption communities between existing buildings connected by the public grid, although this, from the point of view of optimizing self-consumption, would make sense. Therefore, most self-consumption communities are currently formed within individual buildings (apartment buildings) or in newly constructed neighborhoods. Self-consumption communities remain, in our view, an effective tool to encourage PV penetration while reducing investments in the grid because they encourage the installation of PV that is consumed locally, and does not have to be transported to the distribution grid. However, we find that a virtual community model, such as the Italian one, would be easier to implement and would serve the same purpose. In Italy, it is sufficient to reside under the same secondary substation to form self-consumption communities. The incentive calculation is based on virtually self-consumed energy and disbursed by the state. This system also works in a liberalized market, in which every user has the right to choose their own energy supplier, and in the case where, in order to receive incentives, the PV plant owner is forced to sell excess energy to a third party. In a virtual community setting, it would be easier to create energy communities and to see them evolve in time by adding or removing users.

8.4 The present and the future of LIC

In Section 7, we have seen that users in the community are generally satisfied. Although they initially expected more of their consumption to be covered by nearby PV plants, they still have saved an average of 5% on their bills, with peaks of 15% for owners of large PV systems with little self-consumption (Figure 64). A new PV plant was installed, and new ones will likely be installed in the future. This kind of evolution, with the expansion of decentralized generation within a community, is indeed welcomed and easier in a neighborhood community setting, like in LIC. Some community members have replaced their oil-fired heating systems with heat pumps and thus take advantage of the community's local rates to heat their homes, as well. There was a realization that even by providing users with a web portal allowing them to view their consumption and production and identifying the times of day when energy prices were cheapest, they only partially took advantage of the opportunity. People did not visit the site often (some never), which makes us assume that they did not actively shift their consumption to times when PV production was high and, consequently, local prices low. This is in line with what has been observed in the literature and in previous pilot projects carried out by SUPSI in Switzerland. In our view, automation of load control is preferable to having end users shift loads manually. It is also likely that in a future context of much higher end-user energy prices than today, users will certainly be more careful about their consumption, which may push them to move their load more actively during periods when energy prices in the community are lower. At present, loads that a normal user is unable to shift manually such as boilers and heat pumps have been automatically controlled. But, in the future, other household appliances will also be connected to the Internet and consequently controllable remotely. And above all, electric cars will come, which with their high consumption and great flexibility lend themselves particularly well to automatic control.

As more control is deployed in distribution networks as a consequence of the increasing penetration of renewables, the LV grids are becoming increasingly complex systems. This setting does not lend itself particularly well to the use of centralized solutions. Large-scale optimization should be broken down into smaller-scale distributed problems managed by decentralized optimization solutions. However, even if decentralized optimization frameworks are computationally scalable, they require complete trust among all parties, lack protection against cyber attacks, and can put consumer privacy at risk. This fosters a promising application field for distributed smart contract-enabled blockchain framework establishing consensus-based security and a transparent energy monetization and cost-reflective energy services framework. In LIC, we successfully tested the implementation of the local P2P market, but such a framework should offer a wide spectrum of services, like the subscription to different energy management services, guarantee of origin certificate credits, and real-time asset valuation. Blockchain could constitute a unifying framework for real-time energy management, utility analytics and grid situational awareness, distributed intelligence and predictive control, followed by automated secure settlements



under a truly transactive energy framework.

The experience of the LIC will not end with this project; the self-consumption community will continue to exist, and will serve as a pilot project in two European projects. In the first project (H2020 PARITY), a blockchain-based framework for offering flexibility services to DSOs will be investigated. In the second project (Horizon Europe FEDECOM), the possibility of federating different energy communities with each other to offer services to DSOs and TSOs will be investigated.



9 Publications

1. Salani, Matteo, Marco Derboni, Davide Rivola, Vasco Medici, Lorenzo Nespoli, Federico Rosato, and Andrea E. Rizzoli. "Non intrusive load monitoring for demand side management." *Energy Informatics* 3, no. 1 (2020): 1-12.
2. Strepparava, Davide, Lorenzo Nespoli, Evgenia Kapassa, Marios Touloupou, Leonidas Katelaris, and Vasco Medici. "Deployment and analysis of a blockchain-based local energy market." *Energy Reports* 8 (2022): 99-113.
3. Strepparava, Davide, Federico Rosato, Lorenzo Nespoli, and Vasco Medici. "Privacy and Auditability in the Local Energy Market of an Energy Community with Homomorphic Encryption." *Energies* 15, no. 15 (2022): 5386.



10 References

References

- [1] “https://en.wikipedia.org/wiki/ARM\textunderscore_architecture.”
- [2] “<https://www.dlms.com/>.”
- [3] “<https://github.com/Gurux/Gurux.DLMS.Python>.”
- [4] G. Angeris, H.-T. Kao, R. Chiang, C. Noyes, and T. Chitra, “An analysis of Uniswap markets,” pp. 1–25, 2019. [Online]. Available: <http://arxiv.org/abs/1911.03380>
- [5] G. Angeris and T. Chitra, “Improved Price Oracles: Constant Function Market Makers,” *SSRN Electronic Journal*, pp. 1–29, 2020.
- [6] 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.
- [7] G. Belgioioso and S. Grammatico, “Projected-gradient algorithms for Generalized Equilibrium seeking in Aggregative Games are preconditioned Forward-Backward methods,” *arXiv*, 2018.
- [8] “https://en.wikipedia.org/wiki/IEC_61968.”
- [9] W. F. Holmgren, C. W. Hansen, and M. A. Mikofski, “pvlib python: a python package for modeling solar energy systems,” *Journal of Open Source Software*, vol. 3, no. 29, p. 884, 2018. [Online]. Available: <https://www.theoj.org/joss-papers/joss.00884/10.21105.joss.00884.pdf>
- [10] M. Salani, M. Derboni, D. Rivola, V. Medici, L. Nespoli, F. Rosato, and A. E. Rizzoli, “Non intrusive load monitoring for demand side management,” *Energy Informatics*, vol. 3, no. 1, pp. 1–12, 2020.
- [11] 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.
- [12] L. S. Shieh, H. Wang, and R. E. Yates, “Discrete-continuous model conversion,” *Topics in Catalysis*, 1980.
- [13] 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.
- [14] C. M. Bishop, *Pattern Recognition and Machine Learning*, 2013, vol. 53, no. 9.
- [15] “Gdpr <https://gdpr-info.eu/>.”
- [16] “<https://blockchainhub.net/blog/blog/blockchain-scalability-statchannels-sidechains/>.”
- [17] “<https://alephzero.org/blog/what-is-the-fastest-blockchain-and-why-analysis-of-43-blockchains/>.”
- [18] “<https://academy.binance.com/en/glossary/transactions-per-second-tps>.”
- [19] “<https://medium.com/@akash\textunderscore13214/the-scalability-trilemma-in-blockchain-75fb57f646df>.”
- [20] “<https://www.bitorb.com/campus/what-is-the-scalability-trilemma>.”
- [21] “Aragon, <https://aragon.org>.”
- [22] “AragonOS, <https://github.com/aragon/aragonOS>.”



- [23] "Voting app, <https://github.com/aragon/aragon-apps/tree/master/apps/voting>."
- [24] "ERC20, <https://github.com/ethereum/eipsi/issues/20>."
- [25] "Lightning network <https://lightning.network/>."
- [26] "Raiden network <https://raiden.network/>."
- [27] "Liquidity <https://liquidity.network/>."
- [28] "Celer <https://www.celer.network/>."
- [29] "Cosmos network <https://cosmos.network/>."
- [30] "Byzantine fault tolerance https://en.wikipedia.org/wiki/byzantine_fault."
- [31] "Proof of authority https://en.wikipedia.org/wiki/proof_of_authority."
- [32] "Proof of stake https://en.wikipedia.org/wiki/proof_of_stake."
- [33] "Proof of work https://en.wikipedia.org/wiki/proof_of_work."
- [34] "Go programming language <https://golang.org/>."
- [35] "<https://asecuritysite.com/encryption/go\textunderscore priv>."
- [36] P. McCorry S. F. Shahandashti and F. Hao, "A smart contract for board- room voting with maximum voter privacy," *International Conference on Financial Cryptography and Data Security*, pp. 357–375, 2017.
- [37] C.-P. Schnorr, "Efficient signature generation by smart card," *International Conference on Financial Cryptography and Data Security*, vol. 4, no. 3, pp. 161–174, 1991.
- [38] A. S. A. Fiat, "How to prove yourself: Practical solutions to identification and signature problems," *Conference on the Theory and Application of Cryptographic Techniques*, vol. 4, no. 3, pp. 186–194, 1986.
- [39] 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.
- [40] "Pseudonymization <https://en.wikipedia.org/wiki/pseudonymization>."
- [41] EnergieSchweiz, "Stationäre Batteriespeicher in Gebäuden," Tech. Rep., 2018. [Online]. Available: <https://pubdb.bfe.admin.ch/de/publication/download/9430>
- [42] T. Brinsmead, P. Graham, J. Hayward, E. Ratnam, and L. Reedman, "Future energy storage trends: An assessment of the economic viability, potential uptake and impacts of electrical energy storage on the nem 2015–2035," *Report Prepared for the Australian Energy Market Commission*, 2015.
- [43] Tech. Rep.
- [44] G. Montanari and L. Simoni, "Aging phenomenology and modeling," *IEEE Transactions on Electrical Insulation*, vol. 28, no. 5, pp. 755–776, Oct. 1993, conference Name: IEEE Transactions on Electrical Insulation.
- [45] "Assessing and Managing Cable Ageing in Nuclear Power Plants," Sep. 2016, publisher: IAEA. [Online]. Available: <https://www.iaea.org/publications/8753/assessing-and-managing-cable-ageing-in-nuclear-power-plants>
- [46] M. H. Kim, H. J. Seo, S. K. Lee, and M. C. Lee, "Influence of Thermal Aging on the Combustion Characteristics of Cables in Nuclear Power Plants," p. 17, 2021.



- [47] V. N. Pugach, D. A. Polyakov, K. I. Nikitin, I. L. Zakharov, and N. N. Petukhova, "Analysis of Temperature Conditions Influence on Cables Insulation Operation Life," *Journal of Physics: Conference Series*, vol. 1260, no. 5, p. 052027, Aug. 2019. [Online]. Available: <https://iopscience.iop.org/article/10.1088/1742-6596/1260/5/052027>
- [48] D. D. Evans, "Calculated operating temperatures of thermally insulated electric cables," p. 48.
- [49] S. Sachan, R. Wen, Y. Xiang, L. Yao, and C. Zhou, "A Stochastic Electrothermal Degradation Model of Power Cables," p. 10, 2015.
- [50] Y. Parag and B. K. Sovacool, "Electricity market design for the prosumer era," *Nature energy*, vol. 1, no. 4, pp. 1–6, 2016.
- [51] G. Walker, P. Devine-Wright, S. Hunter, H. High, and B. Evans, "Trust and community: Exploring the meanings, contexts and dynamics of community renewable energy," *Energy policy*, vol. 38, no. 6, pp. 2655–2663, 2010.
- [52] T. Bauwens and P. Devine-Wright, "Positive energies? An empirical study of community energy participation and attitudes to renewable energy," *Energy Policy*, vol. 118, pp. 612–625, 2018.
- [53] J. C. Rogers, E. A. Simmons, I. Convery, and A. Weatherall, "Public perceptions of opportunities for community-based renewable energy projects," *Energy policy*, vol. 36, no. 11, pp. 4217–4226, 2008.
- [54] G. Walker, "What are the barriers and incentives for community-owned means of energy production and use?" *Energy policy*, vol. 36, no. 12, pp. 4401–4405, 2008.
- [55] G. Walker and P. Devine-Wright, "Community renewable energy: What should it mean?" *Energy policy*, vol. 36, no. 2, pp. 497–500, 2008.
- [56] T. Van Der Schoor and B. Scholtens, "Power to the people: Local community initiatives and the transition to sustainable energy," *Renewable and sustainable energy reviews*, vol. 43, pp. 666–675, 2015.
- [57] B. J. Kalkbrenner and J. Roosen, "Citizens' willingness to participate in local renewable energy projects: The role of community and trust in germany," *Energy Research & Social Science*, vol. 13, pp. 60–70, 2016.
- [58] V. Brummer, "Community energy—benefits and barriers: A comparative literature review of community energy in the UK, germany and the USA, the benefits it provides for society and the barriers it faces," *Renewable and Sustainable Energy Reviews*, vol. 94, pp. 187–196, 2018.
- [59] M. Van Vugt, "Central, individual, or collective control? Social dilemma strategies for natural resource management," *American Behavioral Scientist*, vol. 45, no. 5, pp. 783–800, 2002.
- [60] M. Hechter and K.-D. Opp, "Social norms," 2001.
- [61] A. Franzen and D. Vogl, "Two decades of measuring environmental attitudes: A comparative analysis of 33 countries," *Global Environmental Change*, vol. 23, no. 5, pp. 1001–1008, 2013.
- [62] B. Sütterlin, T. A. Brunner, and M. Siegrist, "Who puts the most energy into energy conservation? a segmentation of energy consumers based on energy-related behavioral characteristics," *Energy Policy*, vol. 39, no. 12, pp. 8137–8152, 2011.
- [63] G. Deng and P. Newton, "Assessing the impact of solar pv on domestic electricity consumption: Exploring the prospect of rebound effects," *Energy Policy*, vol. 110, pp. 313–324, 2017.
- [64] Y. Qiu, M. E. Kahn, and B. Xing, "Quantifying the rebound effects of residential solar panel adoption," *Journal of environmental economics and management*, vol. 96, pp. 310–341, 2019.