# CARAVELS: a case study using peer-to-peer transaction, voluntary demand response, and energy storage optimization

Rafael Silva[1], Rita Costa[1], Luis Gomes[1], and Zita Vale[1]

[1] GECAD – Research Group on Intelligent Engineering and Computing for Advanced Innovation and Development, LASI – Intelligent Systems Associate Laboratory, ISEP, Polytechnic of Porto, Portugal
{rdpds, ricdc, log, pnf, zav}@isep.ipp.pt

*Abstract*—**With the appearance of energy communities and the proposal of novel energy management models for these communities, it is imperative the proposal of novel deployment infrastructures that could accommodate the needs of energy communities. The deployment of models in pilot sites and energy communities is needed to enable the correct testing and validation of models. Therefore, this work presents a distributed energy community infrastructure based on software containers, known as CARAVELS. The system is designed to support energy communities by enabling the deployment, testing, and validation of novel energy management models. The platform was tested in a virtual energy community consisting of 7 real buildings/members and 20 simulated buildings/members, demonstrating its capability to deploy and run energy models in real-world settings. The results suggest that CARAVELS is effective for managing energy in energy communities.**

*Index Terms*--**Energy community, Peer-to-peer energy transaction, Sustainability, Voluntary demand response**

## I. Introduction

The definition and dissemination of energy communities, as well as citizen energy communities, as presented in DL 2019/944 of the European Parliament and of the Council [1], enabled the group of energy consumers to create collaborative and cooperative communities where they drive to pursue stable environments promoting the sustainability of the system [2]. Energy communities have been largely studied and deployed, mainly as pilot sites, in a large set of countries [3], [4]. However, as it combines multiple disciplinary domains and requires stable infrastructure together with proper energy models, there has been no defined solution that could model, manage, and operate an energy community.

Since the adoption of smart grids, consumers, prosumers, and producers, are desired to be active participants in the electric grid [5]. Energy communities could be ideal environments to deploy participation models to increase the users' conscience and engagement towards the use and source of the energy they consume. The collaborative and cooperative environment of energy communities enables the deployment of energy models, such as demand response (DR) and peer-to-peer (p2p) energy transactions, to promote the interactions among community's members. However, the motivation of the members must be taken into account as they can have significant implications in the user participation [6].

Being representative of multiple members, energy communities are distributed systems with multiple stakeholders that must be modeled into the system [7]. Therefore, distributed solutions must be considered when proposing a solution for energy communities. These solutions could be centralized, distributed, or both, depending on the model and functionality adopted. For example, the energy community can assume a distributed resource management solution while providing a centralized p2p energy transaction model.

Data privacy, ownership, and trustworthy systems are also concerns that must be addressed in energy communities [8]. The existence of multiple stakeholders, where each has one or more data sources, needs to be accounted for to guarantee the protection of the users' data. Therefore, data treatment, storage, and transmission are critical topics that must be considered. In [9], a decentralized solution was proposed to address data concerns in energy communities, giving control and decision power to the data owners.

The use of artificial intelligence (AI) models and their integration in energy communities are common topics addressed in scientific studies [10]. The capability of such models to interpret data, extract profiles, and knowledge are very important when developing energy management models that could negatively impact users. Therefore, several energy-related models are proposed using AI-based models or supported by AI-based models. An example of this is the use of energy management models that operate on a day-ahead basis, which require proper energy forecasting models, usually based on machine learning algorithms [11], [12].

This paper provides a better understanding of the proposed solution CARAVELS. This solution is a distributed solution based on software containers and Internet of Things (IoT) devices, enabling and promoting retrofitting. CARAVELS deploys a Kubernetes node in each community member which will then aggregate the information of multiple IoT devices. This guarantees the local use of data without compromising the privacy of the user. This distributed solution enables the efficient deployment of models in the energy community. The energy community operator can manage the infrastructure and models in a centralized approach.

In this paper, three energy models will be presented and demonstrated: a p2p auction-based model for energy transactions among members, a voluntary-basis DR program designed for energy communities, and an energy storage system (ESS) management model for the management of ESS at community level. These models were integrated into CARAVELS and deployed in a real case study. Despite having real interactions and communications, there are no physical connections between the members, which makes it a virtual energy community. This community is composed of 7 real members and 20 simulated members.

After this first introductory section, section II presents the proposed solution, considering the main components required for the deployment in an energy community. Section III presents the case study and results obtained during the implementation and use of the proposed solution in a virtual energy community composed of real physical members together with simulated members. Finally, section IV presents the main conclusions of this work.

## II. CARAVELS

The solution presented in the paper is known as CARAVELS and has been developed in multiple research and development projects. This paper will focus on how the integration of multiple energy models can be achieved in this solution and how they interact among them to provide a complete solution for energy communities.

Figure 1 shows the architecture overview of CARAVELS, where on its base are visible IoT devices that provide monitoring and controlling abilities to the Kubernetes' nodes available in the system. The number of nodes will depend on the number of members and stakeholders available in the deployed system. For instance, it is possible to have additional computation resources as nodes, to provide more computational capability for the system. This is particularly important when using single-board computers to represent members within a community, as they have limited computational capability. In these cases, it is recommended to have additional servers, available as nodes, to allow members to move their containers. This enables nodes with limited capability to move containers with AI-based models so they can be (re)trained.

The containers in the same Kubernetes node communicate, among them, using Message Queuing Telemetry Transport (MQTT) protocol to create shared data streams with multiple subscribers and publishers. Communication between nodes uses HyperText Transfer Protocol (HTTP) communications inside the Kubernetes network. To enable remote communication as well as ensuring privacy and security, CARAVELS uses a Virtual Private Network (VPN) across its infrastructure nodes where all the nodes can interact even though they do not share the same local network. Kubernetes is used to create a common infrastructure to manage and deploy services as containers.

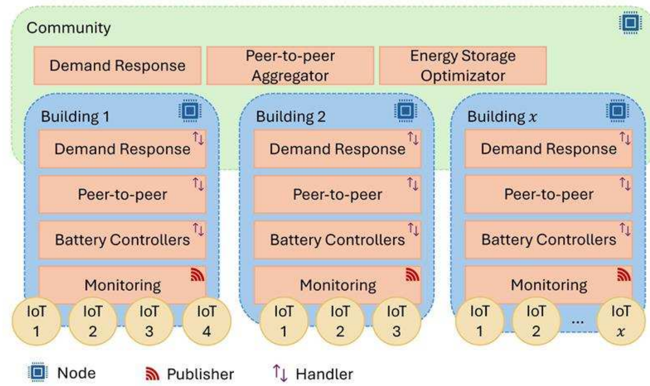


Figure 1. CARAVELS architecture overview

Each node of CARAVELS provides a web-based graphical interface where users can access monitor, manage, and control their infrastructure. Figure 2 shows the dashboard of a community's member where it is possible to see the real-time consumption, generation, and flexibility, a previous short history of these values, and a six-hour forecast of the same values. On the left of the page, it is visible the menu available where the user can navigate to the: dashboard, energy status, IoT devices list, building's room, battery resources, accessing and shared tokens, and the user's profile.

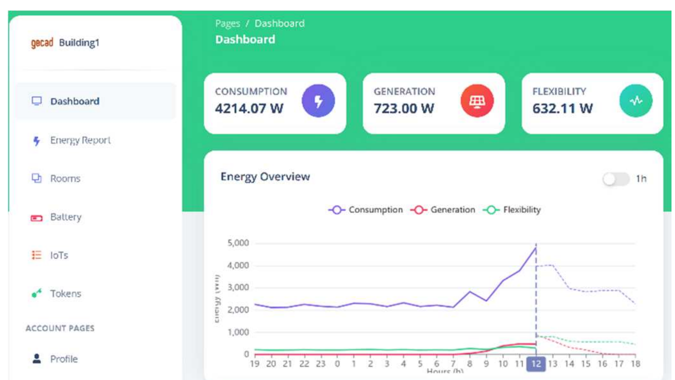


Figure 2. Web application on the phone (print screen from the application)

### A. IoT-Based Real-Time Data Acquisition Subsystem

As previously stated, CARAVELS makes use of IoT devices, that are easy to be acquired and installed, to enable remote monitoring and/or control of sensors and resources. Also, this allows the user to integrate already existing IoT devices into CARAVELS. To enable this integration, CARAVELS provides multiple container applications for multiple communication protocols, such as HTTP, MQTT, and Advanced Message Queuing Protocol (AMPQ). However, having a container-based infrastructure enables CARAVELS to evolve over time and provide additional communication protocols without compromising the already implemented solutions. Nonetheless, IoT devices, when integrated, must provide viable communication with one of the CARAVELS protocols, otherwise, it will not be possible to be integrated. The integration also enables the continuous use of the IoT devices with previous solutions, such as Google Home and the respective manufacturer mobile application.

The IoT-based approach of CARAVELS enables its deployment in a significant number of pilots without the need for new equipment installation. While providing a viable solution for monitoring and control. The use of IoT devices also enables CARAVELS to be used outside energy communities, such as in smart buildings and smart communities where energy is important but where more domains are needed to create intelligent models for application.

### B. Energy Forecast Models

Currently, CARAVELS provides multiple energy forecast models for energy consumption, energy generation, and energy flexibility. The models integrated were based on past published works, such as [13], [14] and [14]. These AI-based models for energy forecasting provide predictions for day-ahead and hour-ahead considering consumption, generation, and flexibility. A total of six energy forecast models were integrated and deployed in each community member and in the community operator. Each of the models has a dedicated container responsible for querying the needed data and providing the forecasted value, or values, to the database manager container that will store the value(s) in the database so every other container can access it.

The real members of the virtual energy community used in this paper use single-board computers, such as a Raspberry Pi 4, to host CARAVELS. Thus, the members' computation resources are not sufficient to train the energy forecast models. To overcome this problem, the training was executed prior to the deployment of the model in the member's Kubernetes node.

### C. Peer-to-Peer Energy Trading

To enable energy sharing among the community members, the auction-based p2p model proposed in [15] was integrated into CARAVELS for day-ahead p2p transactions. The model is supported by the energy forecast models of consumption and generation. These day-ahead forecasts provide 24 values for hourly energy forecast which is then converted, by the p2p container, to the energy that the member will try to buy or sell in the p2p model.

Each community member has access to the energy market price and sets the p2p price, per kWh, according to the price of the market. The price is set by the user by choosing a percentual variation of the market price. The bids from each member are sent to the energy community operator via HTTP communication. This bid includes the amount of energy in kWh and the price in EUR/kWh.

The p2p model results in direct transactions among peers that are later communicated to each member. Participation in the p2p auction-based model is optional; however, once a member chooses to participate, they are obligated to fulfill the results of the auction. For instance, if a member sells 0.5 kWh, then it must supply 0.5 kWh the following day. This may involve purchasing the energy from their retailer in order to fulfill their commitment in the p2p market.

### D. Voluntary-Basis Demand Response Program

To assist the energy community operator, a voluntary DR program designed for energy communities, that was published in [16], was integrated in CARAVELS. This voluntary DR program enables the community operator to balance consumption and generation to make the community independent from external energy suppliers.

The integration of the DR program was more complex than the p2p model, as it involved multiple processes in different time periods and required coordination between all members. Despite this complexity, the DR program is also supported by the energy forecast models deployed in the nodes of the community operator and the community members. The DR program requires a day-ahead forecast of energy consumption, generation, and flexibility.

The program starts within the community operator by forecasting energy values for the next and identifying periods where the DR event(s) could be launched. A period, $p$, is viable for a DR event if the two constraints are true

$$consumption_p > generation_p \tag{1}$$

$$consumption_{p<} - flexibility_p \leq generation_p \tag{2}$$

where $consumption$, $generation$, and $flexibility$ represent the sum of values of the energy community (i.e., the sum of each individual member's value).

After identifying the viable DR event periods, the community operator enquires the community members to analyze the opportunity and communicate the flexibility available to that period. If a member wishes to participate in the DR event, they need to respond. The member must specify the amount of flexibility it has for (i) reduction and (ii) shifting.

The community operator will aggregate all the member's proposals and rank the participants using the following four metrics:

- total number of participations vs. the provided flexibility;
- participation rate vs. the provided flexibility;
- flexibility by reductions for the future event vs. the effort rate;
- cost of shifting vs. the effort rate.

The ranking process utilizes clustering algorithms to group participants and subsequently assigns scores for each metric. After applying these metrics, a fairness mechanism is implemented to ensure equal opportunities for all members. This mechanism is necessary because the metrics evaluate historical participation, which could disadvantage those who have not previously participated.

After evaluating the members eligible for the DR event, the community operator sends out invitations. This invitation will be based on a predetermined order of members' scores and does not require a response from the invited members. Members have the option to accept or decline the invitation from the community operator. If a member declines, the community operator will invite the next member on the ranked list.

During the DR event, the community operator is for responsible for monitoring the event to ensure the energy balance is maintained. Throughout the one-hour event, the community operator will use a monitoring container that checks the community's balance every 10 minutes. If the balance does not meet the desired criterea (i.e., the community consumption is higher than its generation) the community operator will invite new members to participate in real-time.

The integration of the DR program requires the deployment of multiple containers across all members and the community operator. All the containers share the member's/node's database and infrastructure.

### *E. Community-Level Energy Storage Optimization*

Additionally to the p2p model and the DR program, it was integrated an ESS model to manage a set of batteries deployed in the energy community, owned by the community's members but centrally managed by the community operator. The deterministic ESS management model schedules the charging and discharging of each battery unit considering the community demand and energy prices. The model integrated was published in [17].

To provide an optimized schedule of batteries, the integrated model uses energy forecast models for consumption and generation in the day-ahead. The schedule is provided for the next 24 hours and the charging and discharging actions are programmed in a container in CARAVELS to force the planned actions.

The integration of this model was done using a single container deployed in the community operator. However, communications with members were open in order to allow the remote control over the members' ESS. The control of the ESS is done using the HTTP protocol available in the battery inverter. The ESS is modeled in CARAVELS as an IoT device.

## III. Case Study

To validate the integrated models, a case study was conducted within a virtual energy community. The composition of the community consists of 27 nodes representing community members, of which 5 consumers and 22 prosumers. There are 5 real residential households, each equipped with anonymized IoT systems to monitor the consumption of the house, 2 real office buildings with photovoltaic generation and 173 IoT devices, and 20 simulated houses with different consumption and generation profiles. Furthermore, the virtual energy community also possesses 6 batteries, of which 3 with 3.6 kWh capacity, and 3 with 2.4 kWh capacity, these batteries are deployed in the office buildings.

The community members representing real buildings use the single-board computer Raspberry Pi 4 as a computation resource where CARAVELS is installed. The simulated buildings use virtual nodes in Kubernetes to create the deployment environment needed to accommodate all the containers. Each member's node has multiple containers running the services and models deployed in that node. Consumers have similar containers to each other, differing in the IoT communication protocols. The prosumers have the same containers as consumers but with the addition of the containers to monitor and forecast energy generation. A total of 28 Kubernetes nodes are used, one for each member and one for the community operator.

### *A. Peer-to-peer Energy Trading*

The p2p auction-based model is executed in day-ahead for the full 24 periods of the day (i.e., one per hour). Figure 3 shows the p2p matrix where it is possible to see all the transactions that will be made among members. In CARAVELS, this information is only shown in the community operator interface and it can be seen as energy values, in kWh, or economic values, in EUR. At the community's member-side, the information regarding p2p is also visible but only transactions to or from that member are shown. Transactions among other members are not publicly displayed to the member.

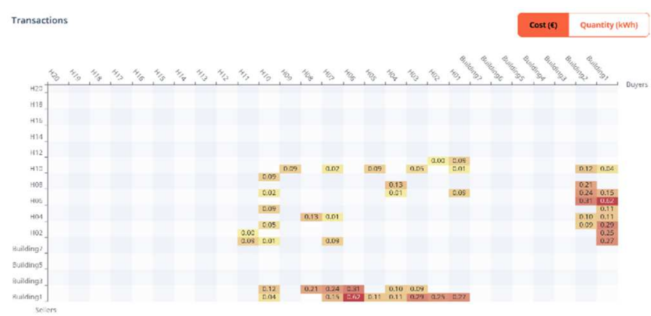


Figure 3. p2p energy transaction matrix between members (print screen from the application)

The use of p2p shows the ability of CARAVELS to provide energy forecast models, namely for consumption and generation. To enable the p2p model, CARAVELS also made available the energy market prices for the next day for members to define their bids.

### *B. Demand Response Event*

The integration of the DR program was done to observe the capacity of CARAVELS to execute multiple processes in different types considering multiple communications, announcements, and invitations. Figure 4 shows the DR event monitoring chart where the community energy balance was monitored during the one-hour event. As can be seen, the energy balance was above zero in two correction periods, i.e., at 40 and 50 minutes after the beginning of the event. At these moments, additional members were invited to participate in the event and to force a decrease in the energy balance of the community. In Figure 4, the corrections are identified with different colors.

The execution of DR in the energy community was only possible by the integration of energy forecast models, namely consumption, generation, and flexibility. Also, the entire infrastructure of Kubernetes was tested during this event as communication among members was mandatory to enable the launch of the program.


Figure 4. DR event monitoring (print screen from the application)

### C. Energy Storage Optimization

The last model that was integrated enabled a centralized scheduler optimization of ESS at the community level. To enable this centralized management, CARAVELS used a sharing resource functionality that enables the owner of a resource to share that resource, for monitoring and/or control, with another member. Therefore, the batteries, located on the office buildings, were shared by those two members with the community operator.

The ESS optimization model also uses the energy consumption and generation forecasts at the community level and the energy market price for the next 24 hours. The optimization is executed every hour to adjust the battery management considering the current context. Figure 5 shows the state of charge of the batteries, individually, for the past periods and for the future periods that were scheduled by the optimizer. The planned actions of charging and discharging of the batteries are scheduled in CARAVELS. In this case study, only planned charging and discharging actions were sent to the batteries.

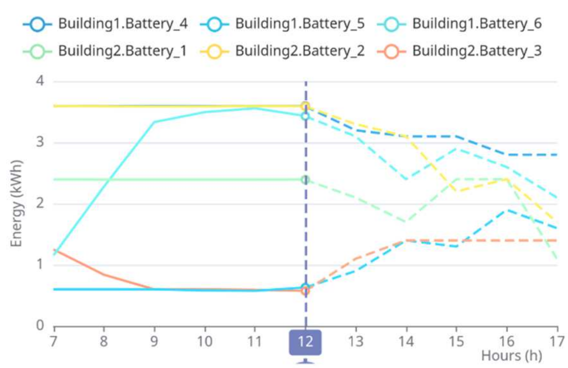


Figure 5. State of charge profile of ESS considering past periods and planned actions (print screen from the application)

## IV. Conclusion

The dissemination of energy communities has been a reality in a significant number of countries, by implementing pilot sites and real legalized energy communities. However, there is a significant gap between the energy management models deployed and the models proposed in the literature. Therefore, new infrastructures that enable the fast deployment of models in order to assess their quality without compromising the infrastructure system and the community resources are needed to fasten the deployment of novel models that can contribute to the community and improve their sustainability and members' participation.

In this work, CARAVELS was used to integrate three different management models supported by energy forecast models, which were also integrated. The integrated models were extracted from previous publications.

The case study, using a virtual energy community of 27 members, integrated energy forecasting models for consumption, generation, and flexibility, a peer-to-peer auction-based energy transaction model, a voluntary-basis demand response program, and a centralized energy storage system optimization. The results demonstrate the capabilities of CARAVELS to integrate such models and provide the needed information for the models and for the users, members and the community operator. These models can contribute to each other and operate independently to manage the community and its resources. This demonstrates how a distributed system can support evolving energy communities, facilitating a broader implementation of energy communities that benefit both community and building-level services.

## Acknowledgment

The authors acknowledge the work facilities and equipment provided by GECAD research center (UIDB/00760/2020), DOI:10.54499/UIDB/00760/2020 to the project team.

## References

[1] "DIRECTIVE (EU) 2019/944 OF THE EUROPEAN PARLIAMENT AND OF THE COUNCIL." Accessed: Feb. 07, 2025. [Online]. Available: https://eur-lex.europa.eu/legal-content/EN/TXT/HTML/?uri=CELEX:32019L0944

[2] L. Gruber, U. Bachhiesl, and S. Wogrin, "The current state of research on energy communities," Elektrotechnik und Informationstechnik, vol. 138, no. 8, pp. 515–524, Dec. 2021, doi: 10.1007/S00502-021-00943-9/FIGURES/6.

[3] S. Heuninckx, G. te Boveldt, C. Macharis, and T. Coosemans, "Stakeholder objectives for joining an energy community: Flemish case studies," Energy Policy, vol. 162, p. 112808, Mar. 2022, doi: 10.1016/J.ENPOL.2022.112808.

[4] A. Parreño-Rodriguez, A. P. Ramallo-González, M. Chinchilla-Sánchez, and A. Molina-García, "Community energy solutions for addressing energy poverty: A local case study in Spain," Energy Build, vol. 296, p. 113418, Oct. 2023, doi: 10.1016/J.ENBUILD.2023.113418.

[5] G. Dileep, "A survey on smart grid technologies and applications," Renew Energy, vol. 146, pp. 2589–2625, Feb. 2020, doi: 10.1016/J.RENENE.2019.08.092.

[6] A. Bourazeri and J. Pitt, "Collective attention and active consumer participation in community energy systems," Int J Hum Comput Stud, vol. 119, pp. 1–11, Nov. 2018, doi: 10.1016/J.IJHCS.2018.06.001.

[7] J. García-Martínez, J. L. Reyes-Patiño, L. B. López-Sosa, and L. F. Fuentes-Cortés, "Anticipating alliances of stakeholders in the optimal design of community energy systems," Sustainable Energy Technologies and Assessments, vol. 54, p. 102880, Dec. 2022, doi: 10.1016/J.SETA.2022.102880.

[8] S. Heuninckx, M. Meitern, G. te Boveldt, and T. Coosemans, "Practical problems before privacy concerns: How European energy community initiatives struggle with data collection," Energy Res Soc Sci, vol. 98, p. 103040, Apr. 2023, doi: 10.1016/J.ERSS.2023.103040.

[9] N. Teixeira, R. Barreto, L. Gomes, P. Faria, and Z. Vale, "A Trustworthy Building Energy Management System to Enable Direct IoT Devices' Participation in Demand Response Programs," Electronics 2022, Vol. 11, Page 897, vol. 11, no. 6, p. 897, Mar. 2022, doi: 10.3390/ELECTRONICS11060897.

[10] S. S. Binyamin, S. A. Ben Slama, and B. Zafar, “Artificial intelligence-powered energy community management for developing renewable energy systems in smart homes,” Energy Strategy Reviews, vol. 51, p. 101288, Jan. 2024, doi: 10.1016/J.ESR.2023.101288.

[11] D. Putz, M. Gumhalter, and H. Auer, “The true value of a forecast: Assessing the impact of accuracy on local energy communities,” Sustainable Energy, Grids and Networks, vol. 33, p. 100983, Mar. 2023, doi: 10.1016/J.SEGAN.2022.100983.

[12] H. Firoozi, H. Khajeh, and H. Laaksonen, “Flexibility Forecast at Local Energy Community Level,” 2021, doi: 10.1109/ISGTEUROPE52324.2021.9640214.

[13] L. Gigoni et al., “Day-Ahead Hourly Forecasting of Power Generation from Photovoltaic Plants,” IEEE Trans Sustain Energy, vol. 9, no. 2, pp. 831–842, Apr. 2018, doi: 10.1109/TSTE.2017.2762435.

[14] M. H. Alobaidi, F. Chebana, and M. A. Meguid, “Robust ensemble learning framework for day-ahead forecasting of household based energy consumption,” Appl Energy, vol. 212, pp. 997–1012, Feb. 2018, doi: 10.1016/J.APENERGY.2017.12.054.

[15] H. Pereira, L. Gomes, and Z. Vale, “Peer-to-peer energy trading optimization in energy communities using multi-agent deep reinforcement learning,” Energy Informatics, vol. 5, no. 4, pp. 1–18, Dec. 2022, doi: 10.1186/S42162-022-00235-2/FIGURES/11.

[16] R. Barreto, L. Gomes, and Z. Vale, “Smart Solution for Energy Communities: Integrating Demand Response and Unsupervised Learning Evaluation Metrics,” 2024 IEEE 22nd Mediterranean Electrotechnical Conference, MELECON 2024, pp. 144–149, 2024, doi: 10.1109/MELECON56669.2024.10608742.

[17] R. Costa, R. Silva, R. Faia, L. Gomes, P. Faria, and Z. Vale, “Empowering energy management in smart buildings: A comprehensive study on distributed energy storage systems for Sustainable consumption,” Energy Build, vol. 324, p. 114953, Dec. 2024, doi: 10.1016/J.ENBUILD.2024.114953.