A JAX-Accelerated Simulation Framework for Multi-Agent Energy Management in Energy Communities

Demonstration Track

Hicham Azmani
Vrije Universiteit Brussel
Brussels, Belgium
hicham.azmani@vub.be

Andries Rosseau
Vrije Universiteit Brussel
Brussels, Belgium
andries.rosseau@vub.be

Marjon Blondeel
Vrije Universiteit Brussel
Brussels, Belgium
marjon.blondeel@vub.be

Ann Nowé
Vrije Universiteit Brussel
Brussels, Belgium
ann.nowe@vub.be

ABSTRACT

The global push toward renewable energy has accelerated the formation of energy communities, where households produce and share electricity locally to reduce grid strain and promote sustainability. Regulators and researchers are still exploring optimal energy exchange mechanisms. However, communities often decide how they will exchange energy themselves, even though many lack the information and tools needed to make informed decisions. To address these challenges, we introduce a JAX-accelerated simulation-based framework that allows researchers to prototype and evaluate diverse energy exchange models under realistic conditions. On top of this framework, we present an interactive demonstration targeted at legislators, citizens, and other non-technical stakeholders, offering an intuitive introduction to foundational concepts and a hands-on environment for experimenting with different community setups and exchange mechanisms. The project website is available here. Beyond technology, our work is grounded in an interdisciplinary project integrating legal analysis, social science research, and public engagement via citizen jury sessions. By bridging these domains, we aim to empower communities and decision-makers to make more informed, equitable choices in transitioning to sustainable energy systems.

KEYWORDS

Multi-Agent Systems; Mechanism Design; Energy Management

ACM Reference Format:
Hicham Azmani, Andries Rosseau, Marjon Blondeel, and Ann Nowé. 2025. A JAX-Accelerated Simulation Framework for Multi-Agent Energy Management in Energy Communities: Demonstration Track. In *Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Detroit, Michigan, USA, May 19 – 23, 2025*, IFAAMAS, 3 pages.

Proc. of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2025), Y. Vorobeychik, S. Das, A. Nowé (eds.), May 19 – 23, 2025, Detroit, Michigan, USA. © 2025 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org).

1 INTRODUCTION

The global transition to renewable energy sources and decentralized energy production is driven by the urgent need to address climate change. This transition is characterized by different initiatives worldwide, such as the Community Choice Aggregation (CCA) [11] in the U.S., microgrid programs in China [7] and Australia [6], and energy communities in the European Union [8]. Localized energy communities, where users produce and share local renewable electricity, accelerate the shift toward clean energy sources while maintaining grid stability and resilience. Despite legal frameworks and regulatory directives promoting decentralization, the exact methods of exchanging energy locally are frequently left in the hands of the communities, many of whom lack the information needed to make fully informed decisions [8]. To address this, we have developed an interactive and user-friendly demonstration that aims to simplify these complex issues, providing citizens with insights to make informed decisions about their energy exchanges.

The widespread adoption of solar panels and smart meters enable new opportunities for energy communities. Across Europe, renewable energy overproduction often creates surpluses that cannot be injected into the grid, sometimes even forcing producers to pay for surplus energy. This highlights the urgent need for demand flexibility and local balancing strategies [12]. Energy communities can provide this flexibility by exchanging surplus energy locally, limiting the strain on the existing grid infrastructure and energy markets. However, coordinating such exchanges is complex, as renewable energy production and consumption patterns can vary dramatically, often on timescales as short as 15 minutes, making manual oversight impractical [10]. Fixed prices are likewise inadequate, as they do not reflect the actual cost of energy production [12]. By employing AI-driven agents that continuously monitor and adapt to real-time conditions, communities can allocate resources more efficiently and adjust pricing fairly.

To deal with this complexity, we have developed a JAX hardware-accelerated simulation-based framework. This framework provides mechanism design researchers with the tools to test and refine energy exchange models under diverse conditions. Additionally, our interactive and user-friendly demonstration is accessible to legislators, citizens, and individuals without coding experience. Within the framework, each household in an energy community is represented by an AI agent, allowing for controlled experimentation with diverse and heterogeneous agent populations and energy exchange mechanisms. Our work is rooted in an interdisciplinary project integrating in-depth legal analysis of relevant regulations and norms with social science research on citizens' values and principles.

2 TECHNOLOGY

2.1 Simulation-Based Framework

Our simulation-based framework leverages JAX [3], a high-performance numerical computing library for Python, to enable rapid prototyping, training, and evaluation of AI-driven agents in diverse energy communities. Additionally, our framework aligns with well-established APIs and standards from the multi-agent reinforcement learning (MARL) community [14], benefiting from novel algorithms and advances in the broader MARL community while contributing to a growing ecosystem of JAX-accelerated MARL environments [2, 4, 5, 9, 14]. Finally, as outlined in the following subsections, our standardized APIs further simplify the introduction of new agent models and exchange mechanisms, encouraging ongoing research into efficient and fair local energy-sharing strategies.

2.1.1 Energy Profiles. Our framework provides diverse data sources and generation methods to study energy exchange mechanisms that are robust, fair, and effective across a spectrum of real-world energy community configurations. Firstly, we integrate existing benchmark datasets such as CityLearn [10, 16], which are widely used in the MARL community, and offer realistic electricity consumption and production patterns. However, since predefined scenarios may not always capture the specific conditions or challenges of a particular community, we incorporate flexible data generation tools such as PVLib [1] and ANTgen [13]. PVLib enables the creation of synthetic solar production profiles, allowing users to simulate anything from urban neighborhoods with limited rooftop space to rural communities where each household can host its own photovoltaic system. Similarly, ANTgen offers an approach for modeling household consumption, varying occupant schedules, and appliance usage to reflect diverse lifestyles and building types. Our framework also supports user-provided data.

2.1.2 Agent Modelling. Each household is represented as an AI-driven agent whose decisions depend on the underlying production and consumption patterns. To capture the complexity and diversity of real-world energy communities, our framework supports heterogeneous agent populations. In terms of sophistication, some agents rely on simple, rule-based strategies, while others employ reinforcement learning methods—such as Q-Learning [17] or Proximal Policy Optimization (PPO) [15], to dynamically adapt their behavior over time. These differences in decision-making complexity allow us to investigate how easily less sophisticated agents can be disadvantaged or outperformed under different energy exchange mechanisms. In addition to their decision-making strategies, agents can also differ in their technological resources. While some may have solar installations, battery storage systems, or both, others rely entirely on external sources. Finally, agents can pursue different goals, from maximizing individual profits to achieving greater self-consumption or promoting community-wide stability and sustainability.

2.1.3 Exchange Mechanisms. Our framework includes multiple exchange mechanisms that define how agents trade energy. To measure the added value of such energy community-driven exchanges, we include a simple baseline scenario in which no local exchange occurs. This reference point highlights the differences in cost, efficiency, and fairness that emerge once agents start trading energy among themselves. Building on this baseline, we offer a double auction market and an agent-based pricing system. The double auction mechanism, a well-known approach in electricity markets, requires participants to submit bids to buy and offers to sell energy, with a central clearing process matching them to determine final trade volumes and prices. In contrast, the agent-based pricing mechanism employs a specialized reinforcement learning agent to dynamically adjust the community's energy price in real-time, responding to fluctuations in supply and demand. Additionally, our framework features a standardized API for integrating new mechanisms to support further mechanism design research.

2.2 Demonstration

While the simulation-based framework targets mechanism design researchers, the accompanying demonstration provides a user-friendly interface targeting a broader audience, including citizens, community organizers, and policymakers. It simplifies the underlying concepts, enabling non-experts to understand local energy exchanges, explore the implications of different setups, and make informed decisions about their own contexts. The demo is a web app containing several parts, the first introduces key concepts using interactive visualizations and examples. Secondly, they can test several scenarios, starting at the household level, adjusting parameters such as occupancy patterns, daily routines, or seasonal variations. As they change these inputs, the plots update in real time, illustrating how energy consumption and production fluctuate with lifestyle choices or seasonalities. Building on these insights, the demonstration then focuses on entire energy communities, where multiple households pool their resources and interact through local energy exchange mechanisms. Users can create their own energy communities by choosing the exchange mechanisms and populating them with agents of differing capabilities, objectives, and sophistication. They can use predefined datasets, generate synthetic profiles, or upload their own data to ground exploration in locally relevant and realistic contexts. As they modify configurations and rerun simulations, users observe how the chosen rules shape the outcomes of cost, efficiency, and fairness. By presenting complex ideas incrementally and allowing participants to interact with the system in real-time, the demonstration offers a hands-on, accessible platform.

ACKNOWLEDGMENTS

This work was supported by the Brussels Capital Region and the Flemish Government under the Innoviris "COOMEP" project and the 'Onderzoeksprogramma Artificile Intelligentie (AI) Vlaanderen' program.

REFERENCES

[1] Kevin S Anderson, Clifford W Hansen, William F Holmgren, Adam R Jensen, Mark A Mikofski, and Anton Driesse. 2023. pvlib python: 2023 project update. *Journal of Open Source Software* 8, 92 (2023), 5994.

[2] ASME 2024. *JaxLife: An Open-Ended Agentic Simulator*. Artificial Life Conference Proceedings, Vol. ALIFE 2024: Proceedings of the 2024 Artificial Life Conference. ASME. https://doi.org/10.1162/isal_a_00770 arXiv:https://direct.mit.edu/isal/proceedings-pdf/isal2024/36/47/2461075/isal_a_00770.pdf

[3] James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. 2018. *JAX: composable transformations of Python+NumPy programs*. Google. http://github.com/jax-ml/jax

[4] Ruan de Kock, Omayma Mahjoub, Sasha Abramowitz, Wiem Khlifi, Callum Rhys Tilbury, Claude Formanek, Andries P. Smit, and Arnu Pretorius. 2023. Mava: a research library for distributed multi-agent reinforcement learning in JAX. *arXiv preprint arXiv:2107.01460* (2023). https://arxiv.org/pdf/2107.01460.pdf

[5] Benjamin Ellis, Jonathan Cook, Skander Moalla, Mikayel Samvelyan, Mingfei Sun, Anuj Mahajan, Jakob N. Foerster, and Shimon Whiteson. 2023. SMACv2: an improved benchmark for cooperative multi-agent reinforcement learning. In *Proceedings of the 37th International Conference on Neural Information Processing Systems* (New Orleans, LA, USA) *(NIPS '23)*. Curran Associates Inc., Red Hook, NY, USA, Article 1634, 27 pages.

[6] MA Farrelly and S Tawfik. 2020. Engaging in disruption: A review of emerging microgrids in Victoria, Australia. *Renewable and Sustainable Energy Reviews* 117 (2020), 109491.

[7] Wei Feng, Ming Jin, Xu Liu, Yi Bao, Chris Marnay, Cheng Yao, and Jiancheng Yu. 2018. A review of microgrid development in the United States–A decade of progress on policies, demonstrations, controls, and software tools. *Applied energy* 228 (2018), 1656–1668.

[8] Jiabo He. 2019. The development and utilization of microgrid technologies in China. *Energy Sources, Part A: Recovery, Utilization, and Environmental Effects* 41, 13 (2019), 1535–1556.

[9] Michael Matthews, Michael Beukman, Benjamin Ellis, Mikayel Samvelyan, Matthew Jackson, Samuel Coward, and Jakob Foerster. 2024. Craftax: a lightning-fast benchmark for open-ended reinforcement learning. In *Proceedings of the 41st International Conference on Machine Learning* (Vienna, Austria) *(ICML'24)*. JMLR.org, Article 1428, 34 pages.

[10] Kingsley Nweye, Kathryn Kaspar, Giacomo Buscemi, Tiago Fonseca, Giuseppe Pinto, Dipanjan Ghose, Satvik Duddukuru, Pavani Pratapa, Han Li, Javad Mohammadi, et al. 2024. CityLearn v2: Energy-flexible, resilient, occupant-centric, and carbon-aware management of grid-interactive communities. *CoRR* abs/2405.03848 (2024). https://doi.org/10.48550/ARXIV.2405.03848 arXiv:2405.03848

[11] Eric O'Shaughnessy, Jenny Heeter, Julien Gattaciecca, Jenny Sauer, Kelly Trumbull, and Emily Chen. 2019. Empowered communities: The rise of community choice aggregation in the United States. *Energy Policy* 132 (2019), 1110–1119.

[12] Oleksandr Prokhorov and Dina Dreisbach. 2022. The impact of renewables on the incidents of negative prices in the energy spot markets. *Energy Policy* 167 (2022), 113073.

[13] Andreas Reinhardt and Christoph Klemenjak. 2020. How does Load Disaggregation Performance Depend on Data Characteristics?: Insights from a Benchmarking Study. In *e-Energy '20: The Eleventh ACM International Conference on Future Energy Systems, Virtual Event, Australia, June 22-26, 2020*. ACM, Australia, 167–177. https://doi.org/10.1145/3396851.3397691

[14] Alexander Rutherford, Benjamin Ellis, Matteo Gallici, Jonathan Cook, Andrei Lupu, Garð ar Ingvarsson Juto, Timon Willi, Ravi Hammond, Akbir Khan, Christian Schroeder de Witt, Alexandra Souly, Saptarashmi Bandyopadhyay, Mikayel Samvelyan, Minqi Jiang, Robert Lange, Shimon Whiteson, Bruno Lacerda, Nick Hawes, Tim Rocktäschel, Chris Lu, and Jakob Foerster. 2024. JaxMARL: Multi-Agent RL Environments and Algorithms in JAX. In *Advances in Neural Information Processing Systems* (Vancouver, Canada), A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (Eds.), Vol. 37. Curran Associates, Inc., 50925–50951. https://proceedings.neurips.cc/paper_files/paper/2024/file/5aee125f052c90e326dcf6f380df94f6-Paper-Datasets_and_Benchmarks_Track.pdf

[15] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347* (2017).

[16] José R. Vázquez-Canteli, Sourav Dey, Gregor Henze, and Zoltán Nagy. 2020. CityLearn: Standardizing Research in Multi-Agent Reinforcement Learning for Demand Response and Urban Energy Management. *CoRR* abs/2012.10504 (2020). arXiv:2012.10504 https://arxiv.org/abs/2012.10504

[17] Christopher JCH Watkins and Peter Dayan. 1992. Q-learning. *Machine learning* 8 (1992), 279–292.