FlexMyHeat

D3.1 – Simulation results for flexibility potential on Belgian electricity grid

Table of Contents

Table of Figures

List of Tables

I. INTRODUCTION

The FlexMyHeat project aims at understanding the role that heat pumps and decentralized storage solutions will play in 2030 and 2050 as a source of flexibility for the national electricity system.

The extra need for electricity by shifting from fossil fuel based heating systems to heat pumps in the upcoming years will increase peak loads in the Belgian electricity grids (on top of increased peak loads caused by other domains that get electrified such as mobility and industry) and will thus lead to challenges with respect to the energy security of supply, the net balance and (on a more local level) to congestion of the grid infrastructure.

However, by properly controlling these heat pumps in combination with local storage solutions, unlocking the available flexibility, this challenge can be turned into an opportunity for the grid, contributing to the national and regional balance of the Belgian electricity system.

The goal of FlexMyHeat is to quantitatively analyze the impact and value of the increased deployment of heat pumps and decentralized electrical/thermal storage in 2030 and 2050 on the Belgian electricity system, including proposed control/coordination strategies at (a combination of) various timescales, ranging from day-ahead markets to imbalance markets.

This quantitative assessment is performed for different scenarios:

- *Business-as-usual:* considering the heat pumps and possibly associated local storage as independent devices, only optimized for local objectives, i.e., maximizing PV self-consumption. Thus, no dynamic interaction from the grid side to exploit their flexibility.
- *Individual smart control:* optimized control of the flexibility opportunities offered by the heat pump or storage devices individually, so assuming that any other devices are only optimized for PV self-consumption maximization.
- *Integrated smart control:* combined optimization of both the heat pump and storage devices for local and market objectives

D1.1 focused on the business-as-usual scenarios while D2.1 focused on the individual smart control assessment of battery systems. In this deliverable, we present the results of the individual smart control of heat pumps and thermal storage systems, and the results for the integrated smart control scenarios.

This deliverable is structured as follows:

- **Section II** describes the creation of temperature setpoint profiles based on data from an apartment building. These profiles are used as input for the RL based control algorithms to assess that the user comfort is guaranteed.
- **Section III** describes our MCTS based control algorithm for heat pumps to minimize energy costs, reduce peak loads while guaranteeing user comfort.
- **Section IV** describes our MCTS based methodology for controlling heat pumps in combination with thermal storage.
- **Section V** describes our methodology for the combined control of battery storage, heat pump and thermal storage.
- **Section VI** describes the results for the 3 analyzed scenarios: smart heat pump control, combined control of heat pump and thermal storage, and integrated control of battery storage, heat pump and thermal storage.
- Finally, in **Section VII**, we provide our conclusions and describe possible next steps.

II. Indoor Temperature Setpoint

To control heat pumps and thermal storage, the temperature setpoint is used as an input to ensure user comfort. The temperature setpoint depends on user behavior and habits, and it can vary throughout the day. In this section, we analyze indoor temperature data from several apartments to obtain their daily temperature setpoint profiles.

We used residential indoor temperature data collected from a 41-apartment building in Nivelles, Wallonia. More details about the dataset are provided in Deliverable D.1.1. We chose four different apartments for the study, with the following features: A001 – 2-bedroom, 87.74 m²; B103 – 1-bedroom, 58.35 m²; C204 – 2-bedroom, 83.02 m²; and A401 – 3-bedroom, 152.92 m².

Setpoint profiles are obtained based on the extracted indoor air temperature (TEx) and supplied air temperature (TSu) profiles for each apartment over the entire year. As an example, TEx and TSu profiles for a specific month are shown in Figure 1.

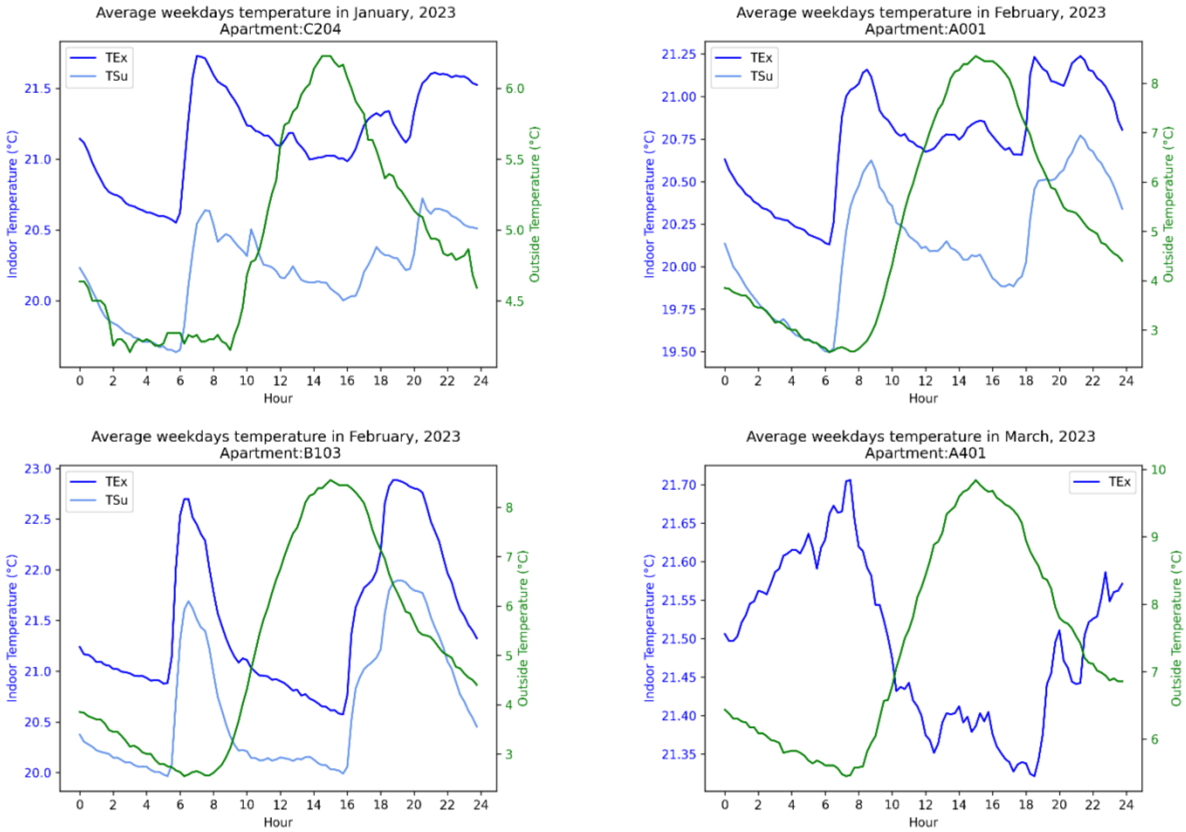

Figure 1 Temperature profiles for selected apartments for a chosen month

Since hours with a sudden rise or drop in temperature indicate a change in the setpoint, these profiles can be used to derive the setpoint temperature profiles for each apartment, as illustrated in Figure 2. All profiles include two major peaks: a morning peak, which corresponds to the hours people spend preparing before work, and an evening peak, which corresponds to the hours after they arrive home.

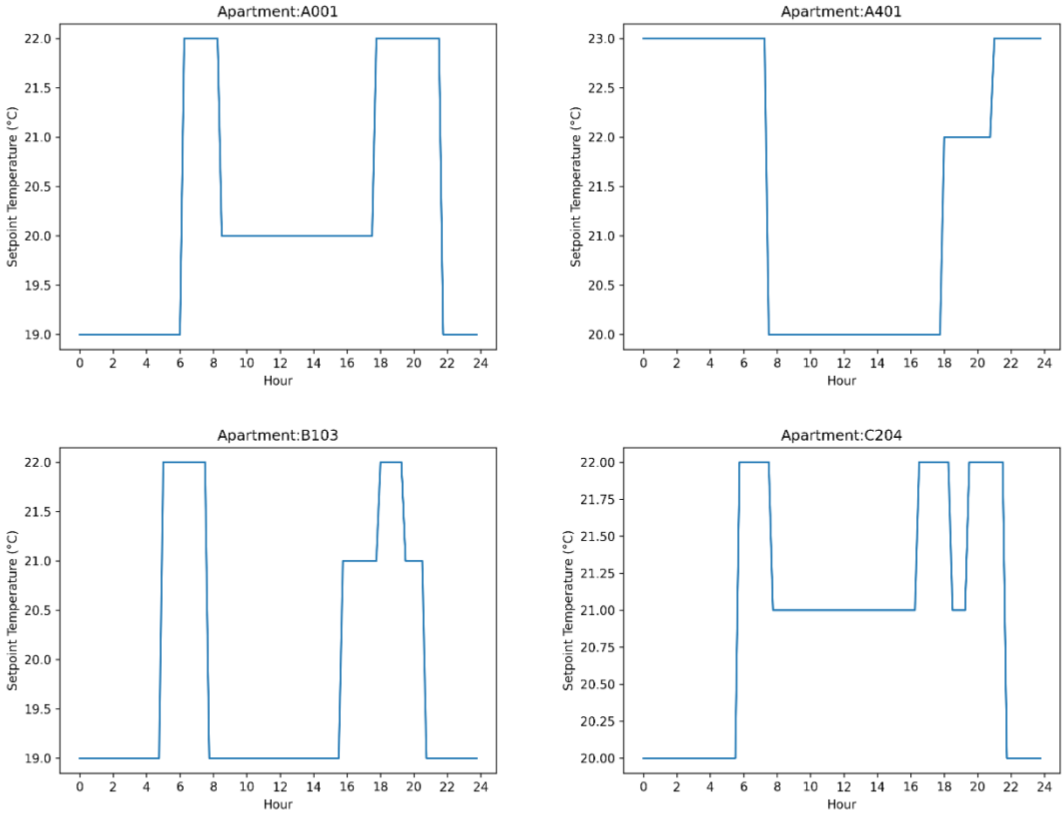

Figure 2 Setpoint temperature profile for each apartment

III. Heat Pump Control Algorithm

As discussed in previous deliverables, a simple rule-based controller cannot effectively control heat pumps. In this section, we introduce our smart controllers for heat pumps, designed to minimize energy costs and reduce negative impacts on peak power demand while maintaining user thermal comfort.

III.1.1. MDP Formulation

We model the sequential control problem of heating a building as a Markov Decision Process (MDP) to minimize the energy cost while staying close to the desired temperature set by the users [2]. The problem is partially observable because some variables (e.g., internal heat gains of the building) remain hidden from the control agent. We define the observable building state at each time step as follows

$$s_t = (t, T_{a,t}, T_{r,t}, T_{r\_set,t}, T_{m,t}, \pi_t, P_t^{PV}, P_t^{load}, P_t^{Bat}, a_{t-h:t-1}^{phys})$$

where t is the hour of day, $T_{a,t}$, $T_{r,t}$, and $T_{r\_set,t}$ are the ambient, room, and user setpoint temperatures at time t, respectively, $T_{m,t}$ is the (estimated) mass temperature of the building at time t, π_t indicates the electricity price at time t, P_t^{PV} and P_t^{load} represent the PV generation of the household and the non-flexible load consumption, respectively. $a_{t-h:t-1}^{phys}$ represent the energy consumed by the heating system in the previous h time steps.

The agent can take 2 possible actions as follows

$$a_t \in A, \qquad A = \{0, P_{max}\}$$

where P_{max} is the maximum electric power of the heat pump. The agent takes an action every 5 minutes. To meet heat pump constraints, we use a backup controller to override actions that violate these constraints. More specifically, the minimum on-time and off-time for all heat pumps are 15 minutes and 10 minutes, respectively.

the reward function is formulated as follows.

$$r_t = \begin{cases} -P_t^{agg}\, \pi_t^{buy} + PF, & P_t^{agg} > 0 \\ -P_t^{agg}\, \pi_t^{inj} + PF, & P_t^{agg} \leq 0 \end{cases}$$

$$P_t^{agg} = P_t^{load} + a_t \Delta t - P_t^{PV}$$

$$PF = -(T_{r\_set,t} - T_{r,t+1})^+ c_1 - (T_{r,t+1} - T_{r\_set,t})^+ c_2$$

where the first term in r_t calculates the energy cost, while the second term ensures that user constraints are satisfied. c_1 and c_2 are hyperparameters used to balance the energy cost objective with the user constraints. In our experiments, the user thermal comfort optimization will use asymmetrical settings ($c_1 > c_2$), because we want to avoid excessive penalization for preheating the room.

III.1.2. Physics-informed Neural Network

To simulate the environment dynamics, we use a model that captures the thermal behavior of a building. For this purpose, we employ a physics-informed neural network forecaster, shown in Figure 3, to estimate the next building states of the system. The architecture is explained in detail in [3]. The model uses an encoder-based neural network to project the most recent building states into a compact hidden state, which corresponds to the building's mass temperature in a simple RC model of its thermodynamic behavior. The decoder network then uses this hidden state, along with observable states, to predict the next building state.

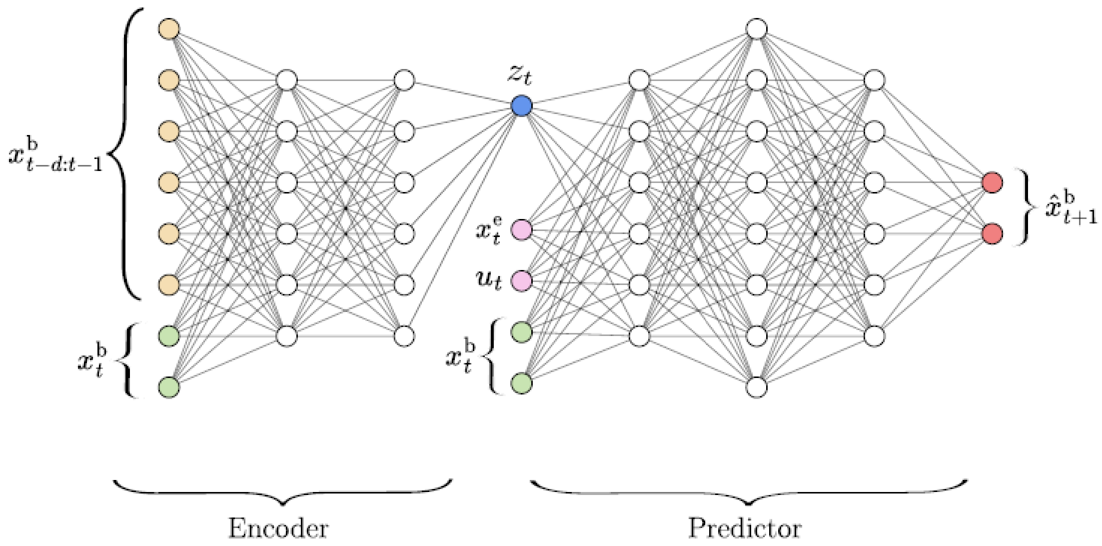

Figure 3 Architecture of the used physics-informed neural network

The physics-informed network was trained offline using historical data from buildings collected in February and March 2023. Figure 4 demonstrates the performance of the trained network on data from May 2023, showing that the model successfully estimated the next indoor temperature.

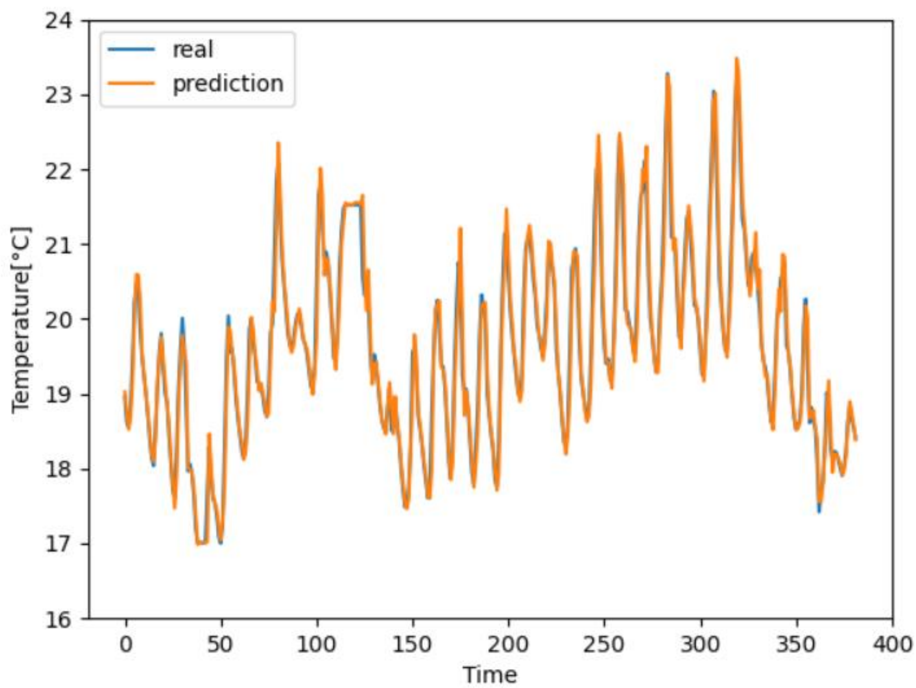

Figure 4 The performance of the physics-informed model trained on building 1 on test days.

III.1.3. Monte Carlo Tree Search

We use Monte Carlo tree search (MCTS) approach to control heat pump and thermal storage. The heat pump action is obtained at each time step using MCTS. Figure 5 illustrates an overview of MCTS. In MCTS, the possible scenarios that the current state may evolve into are represented in a tree structure, modeling subsequent states of the environment as nodes and the actions governing transitions between them as edges. The general structure of the MCTS algorithm comprises four sequential phases that are repeated iteratively, until an acceptable solution is obtained:

1) Selection: select a node (i.e., a system state) to further explore actions for.
2) Expansion: expand the tree by adding new node(s), i.e., roll out possible actions from that state.
3) Simulation: evaluate the node value by performing Monte Carlo simulations.
4) Backpropagation: propagate the information acquired back to the root node. After iterating through these phases ("searching"), the tree is used at inference time to select the actions based on the node values to maximize them.

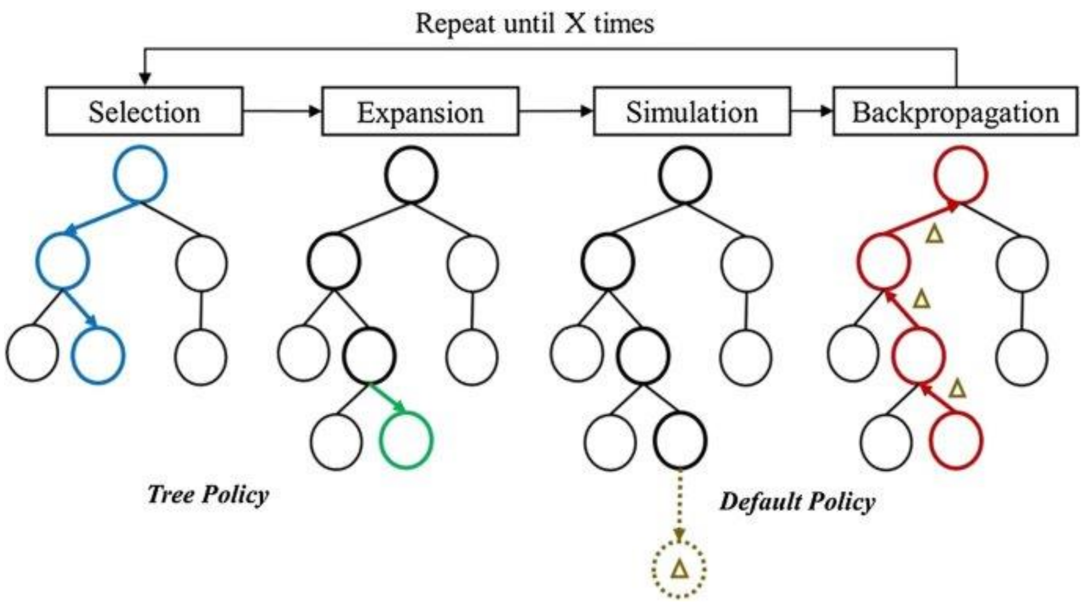

Figure 5 Monte Carlo tree search framework overview

IV. Thermal Storage and Heat Pump Control Algorithm

Thermal storage provides an opportunity to store heating energy for later use in space heating, reducing the need to activate the heat pump. This section describes our methodology for effectively integrating thermal storage with heat pumps.

IV.1.1. MDP Formulation

The control problem is the same as the one explained in Section IV. The observable building state at each time step is formulated as follows

$$s_t = (t, T_{a,t}, T_{r,t}, T_{r\_set,t}, T_{m,t}, \pi_t, P_t^{PV}, P_t^{load}, P_t^{Bat}, SOC_t^{TS}, a_{t-h:t-1}^{phys})$$

where SOC_t^{TS} indicates thermal storage state of charge at time t.

The agent can take 4 possible pairs of actions as follows

$$a_t = (a_t^{HP}, a_t^{TS}) \in A, \qquad A = \{(0,0), (P_{max}^{HP}, 0), (0, -P_{max}^{TS}), (P_{max}^{HP}, P_{max}^{TS})\}$$

where P_{max}^{HP} and P_{max}^{TS} represent the maximum thermal power of the heat pump and thermal storage, respectively. Note that the positive thermal storage action means that it is charged and the negative one means that it is discharged. Similar to Section IV, a backup controller is used to override actions and prevent violations of the heat pump and thermal storage constraints. The heat pump constraints are explained in Section IV, and the thermal storage SoC is limited to 0%, 50% (one cell fully charged), or 100% (both cells fully charged), prioritizing longevity through a 50-50 cell capacity distribution.

The reward function is defined similar to Section IV.

Thermal storage is controlled using the trained physics-informed neural network and MCTS, similar to the approach used for heat pump control.

V. Multi-asset Smart Control

To unlock the full potential of all flexible assets, we control them simultaneously. For this purpose, heat pumps and thermal storage are controlled using MCTS, as explained in Section IV, while the battery is controlled using RL, described in detail in Section D.2.1. Since the time resolution of battery control is one hour, the RL agent determines the battery action every hour. The heat pumps and thermal storage are controlled every 5 minutes, taking into account the battery's action and the household's net consumption.

VI. RESULTS

In this section, the performance of the proposed data-driven controllers for heat pumps, thermal storage, and multiple assets is evaluated. A similar case study as detailed in Deliverable D.1.1 was used, involving 120 households equipped with PV installations across five different locations in Belgium. To study the impact of user comfort on flexibility, we consider three different temperature deadbands for the user temperature setpoint: users do not allow any deviation from the setpoint, allow at most a 0.5 ℃ deviation, or allow at most a 1 ℃ deviation. We investigate the case study under two different price schemes: bi-hourly tariff and dynamic day-ahead prices. In bi-hourly tariff there is two different tariffs for peak and off-peak hours that are calculated based on average day-ahead prices of that hours on that month.

VI.1. Heat Pump Smart Control Results

Table 1 presents the average annual electricity costs of households with heat pumps under different control strategies. Our proposed controller reduced the electricity cost by 13.22% and 10.23% compared to the business-as-usual controller. As expected, allowing a wider temperature deadband provides greater flexibility and cost reductions.

Table 1 Overview of the average annual electricity bill in 2030 for heat pump control

Price Scheme	RBC	Smart Control	
		Flexibility 0.5 ℃	Flexibility 1 ℃
Day-ahead	1170.75 €	1087.89 €	1015.9 €
Bi-hourly tariff	1118.1 €	1058.7 €	1003.7 €

Figure 6 illustrates the average monthly electricity bill of households based on day-ahead prices. The results show that the proposed method consistently reduces the energy costs of households.

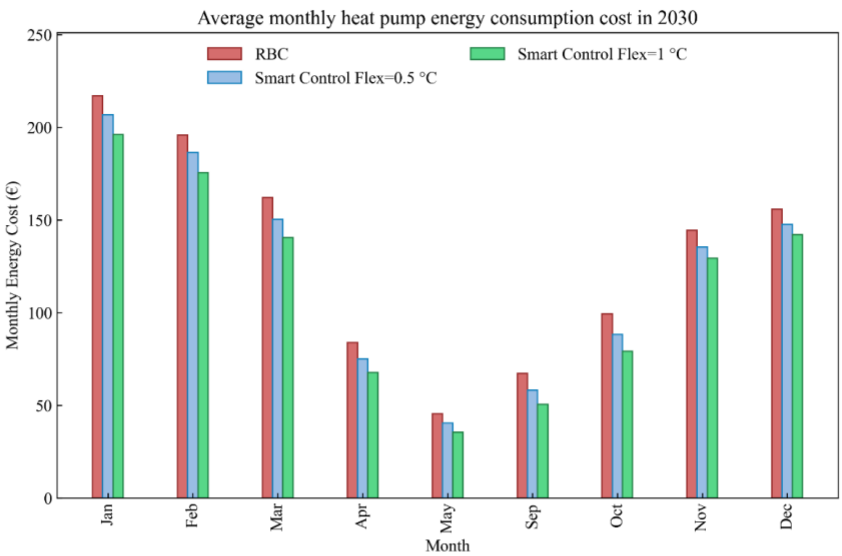

Figure 6 Average household monthly electricity bill in 2030

The average daily peak power of households with heat pumps is shown in Figure 7. In the base scenario, the households have no flexible asset. On average, the smart controller reduces peak power by 9.4% and 13.3% for 0.5 ℃ and 1 ℃ deadbands, respectively, compared to the rule-based controller.

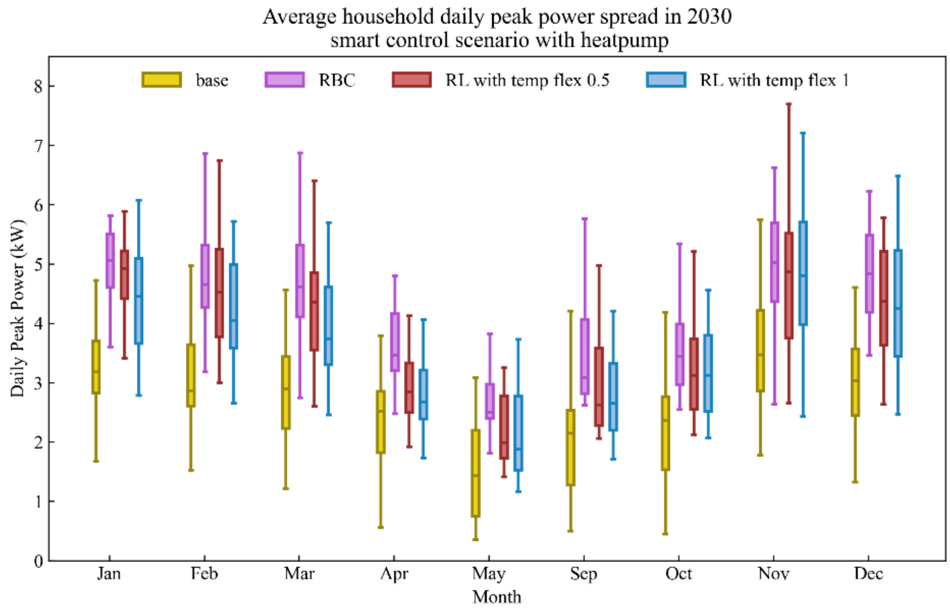

Figure 7 Average daily peak power spread for households with heat pumps

To better interpret the control logic of the proposed method, heat pump actions for an example day, are demonstrated in Figure 8. In this case, the user temperature flexibility is 1℃. The backbone of the smart strategy is preheating, where the controller starts the heat pump earlier to preheat the building when prices are low. Since the temperature setpoint of the house increases at 5:00 and 15:45, the agent turns on the heat pump about two hours earlier to heat the household at a lower cost while ensuring user satisfaction. However, between 16:30 and 18:30, when the price is high, the agent must use the heat pump due to the high temperature setpoint. In this case, using thermal storage can help avoid operating the heat pump during the evening.

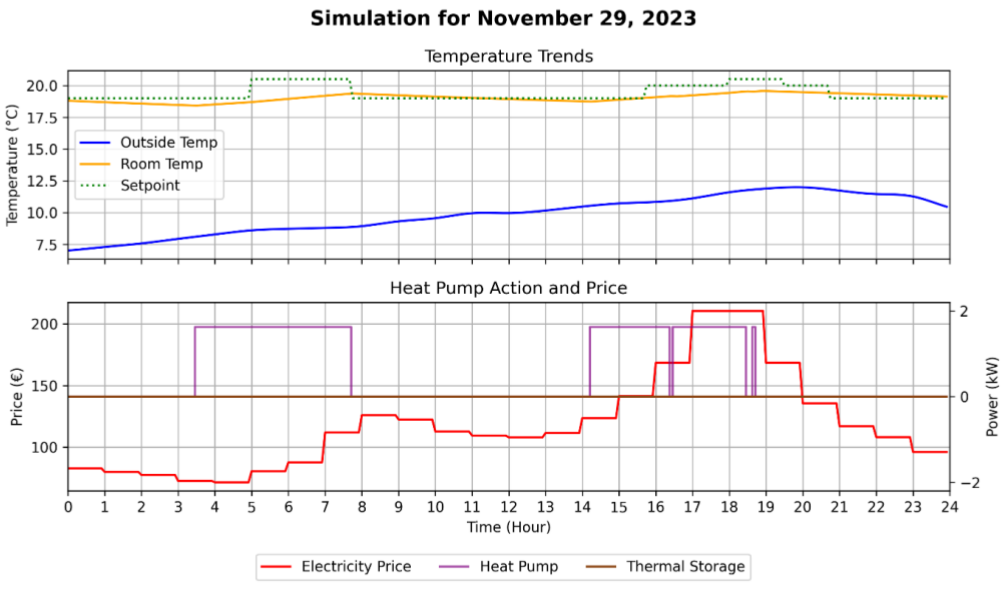

Figure 8 heat pump control actions for an example day

VI.2. Thermal Storage and Heat Pump Smart Control Results

the average annual electricity costs of households with heat pumps and thermal storage under different control strategies is shown in Table 2. Controlling the heat pump and thermal storage simultaneously could decrease the electricity cost by 29.78% and 26.05% compared to the business-as-usual controller. Comparing the results with those in Table 1 indicates that controlling the thermal storage reduced the cost by 19.1% and 17.6% under different pricing schemes.

Table 2 Overview of the average annual electricity bill in 2030 for heat pump and thermal storage control

Price Scheme	RBC	Smart Control	
		Flexibility 0.5 ℃	Flexibility 1 ℃
Day-ahead	1170.75 €	850.04 €	822.08 €
Bi-hourly tariff	1118.1 €	828.3 €	826.8 €

The average monthly electricity bill of households exposed to day-ahead prices is demonstrated in Figure 9. The results indicate a significant and consistent reduction in their monthly costs, demonstrating that the proposed smart control is not biased toward any specific season or month and is a generalizable method.

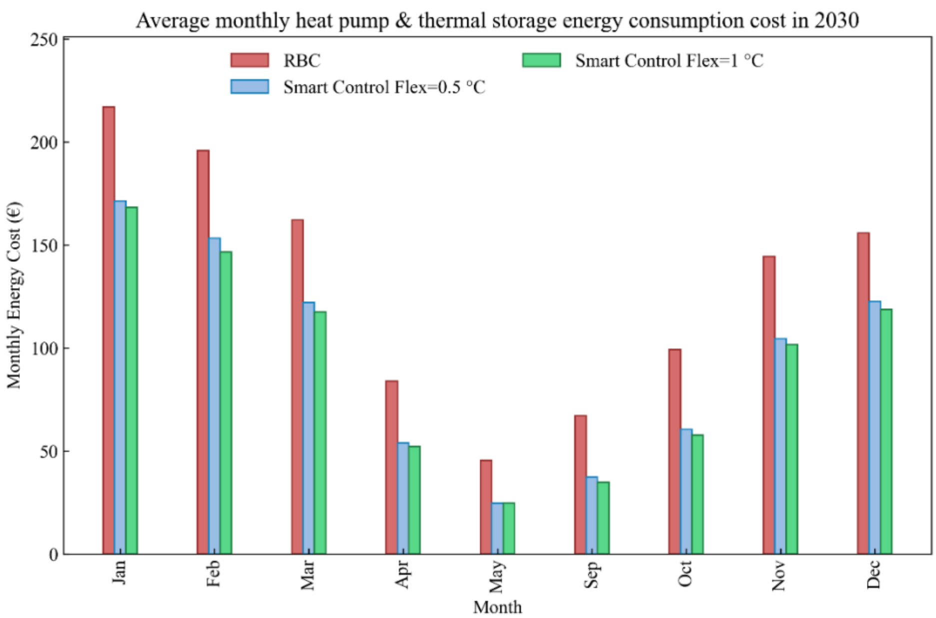

Figure 9 Average household monthly electricity bill in 2030 with heat pumps and thermal storage

Figure 10 illustrates the average daily peak power of households equipped with heat pumps and thermal storage. On average, the smart controller lowers peak power by 22.7% and 24.3% for deadbands of 0.5 ℃ and 1 ℃, respectively, compared to the rule-based controller. Moreover, compared to the scenario with only heat pumps shown in Figure 7, controlling thermal storage could further reduce the average peak power by 14.68% and 12.7% for deadbands of 0.5 ℃ and 1 ℃, respectively.

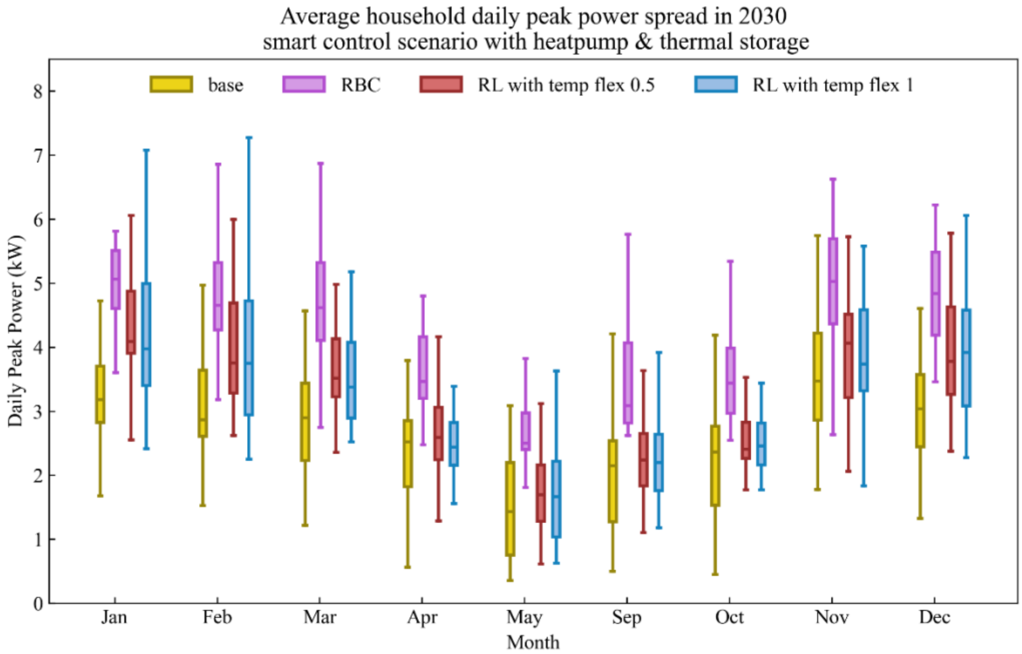

Figure 10 Average daily peak power spread for households with heat pumps and thermal storage

In this scenario as shown in Figure 11, the smart controller follows the same control logic described in Figure 8. Thermal storage is charged when the price is low and discharged when the price is high and heating is needed such as between 4:30 and 5:30.

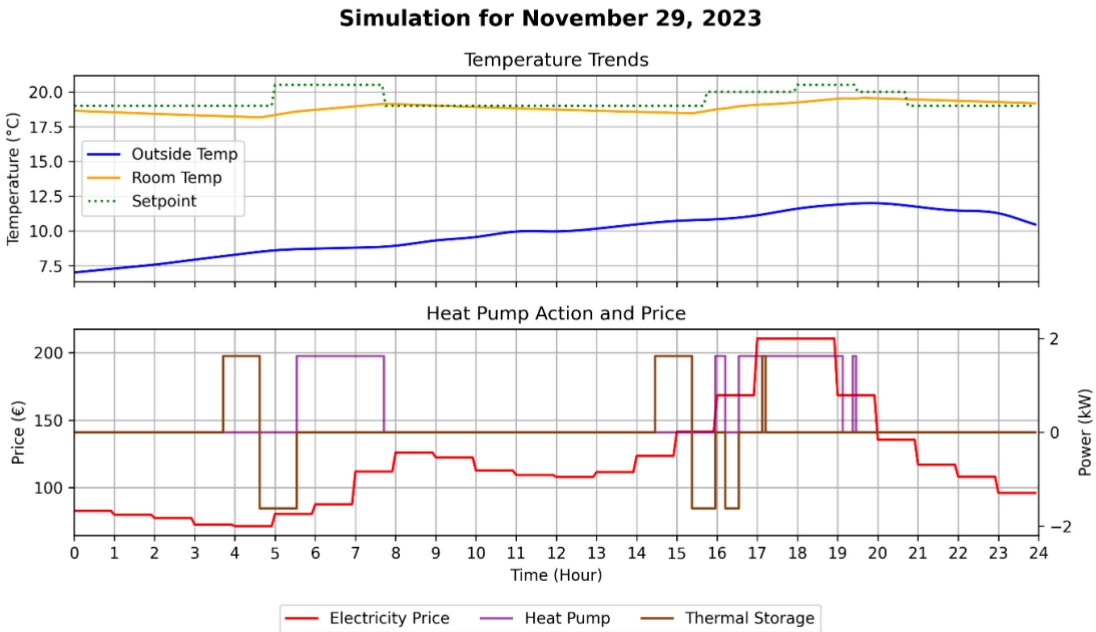

Figure 11 heat pump and thermal storage control actions for an example day

VI.3. Multi-asset Smart Control Results

VI.3.1. Prioritize electric battery over thermal storage

The average annual electricity costs of households with all flexible assets under different control strategies is shown in Table 3. Based on the results, controlling all flexible assets can reduce the average electricity cost of households by 35.1% and 29.9% compared to the scenario in which the heat pumps strictly follow the temperature setpoint.

Table 3 Overview of the average annual electricity bill in 2030 for battery, heat pump and thermal storage control

Price Scheme	RBC	Smart Control	
		Flexibility 0.5 ℃	Flexibility 1 ℃
Day-ahead	1170.75 €	784.23 €	760.12 €
Bi-hourly tariff	1118.1 €	791.33 €	783.92 €

Figure 12 shows the average monthly electricity bill of households exposed to day-ahead prices. The results indicate that during April, May, September, and October, when heating is not needed, the energy cost reduction is primarily due to smart battery control. However, during the other months, the smart control of both the heat pump and thermal storage plays the major role in reducing costs.

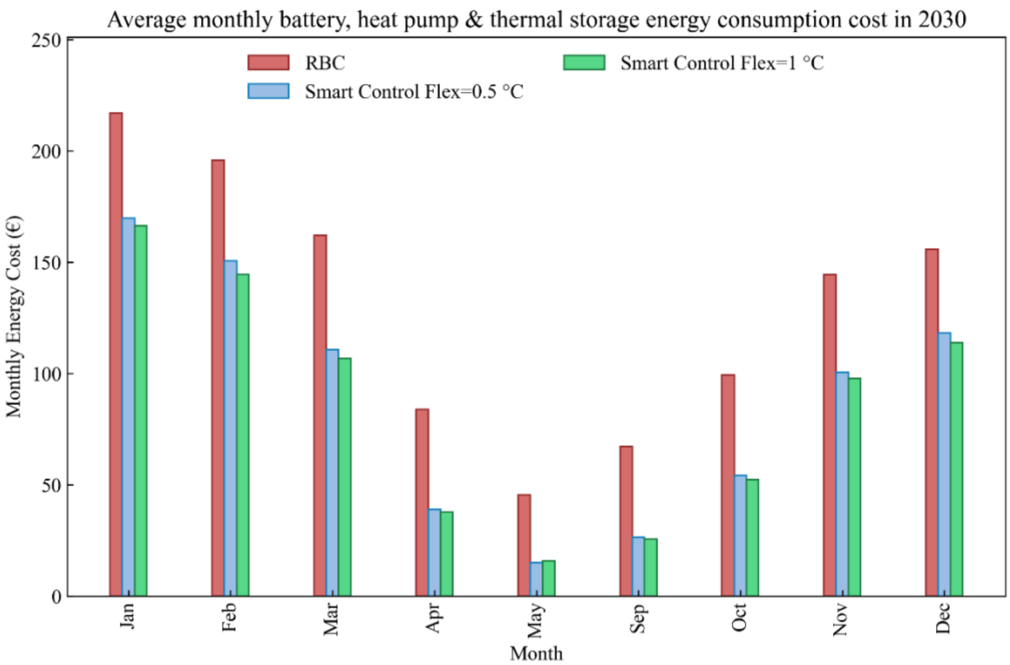

Figure 12 Average household monthly electricity bill in 2030 with battery, heat pumps and thermal storage

Figure 13 shows the average daily peak power of households with battery, heat pumps and thermal storage. Smart control of all flexible assets results in a peak power reduction of 24.8%

and 26% for deadbands of 0.5 ℃ and 1 ℃, respectively, compared to the business-as-usual control of heat pumps.

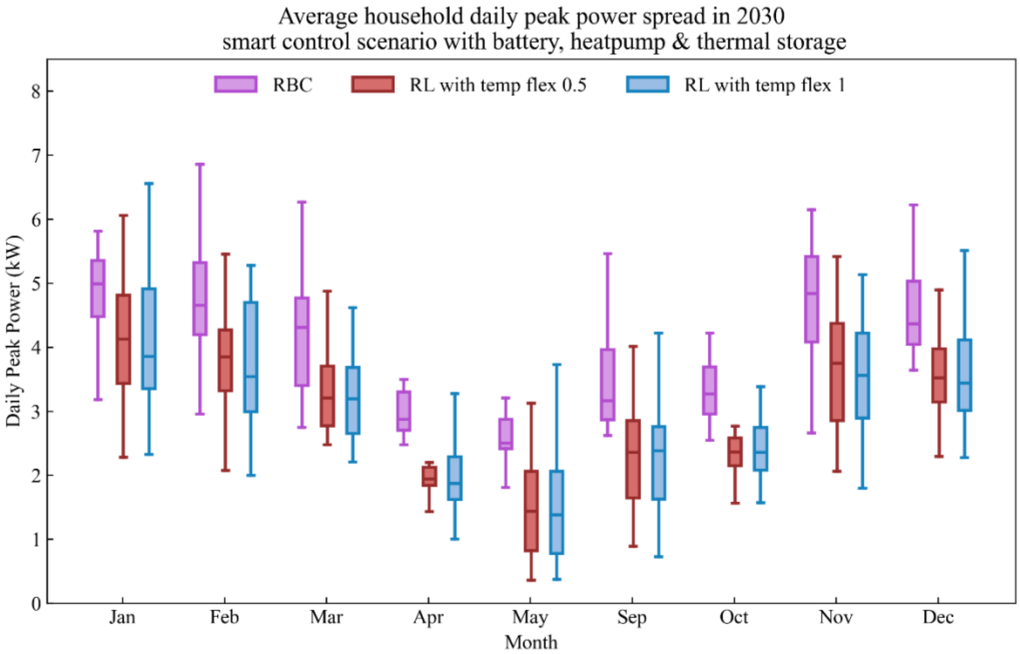

Figure 13 Average daily peak power spread for households with battery, heat pumps and thermal storage

VI.3.2. Prioritize thermal storage over electric battery

Table 4 indicates the average annual electricity costs of households for control strategies that prioritize thermal storage charging over electric battery charging. The results show approximately a 1% cost reduction compared to control strategies that prioritize electric battery charging over thermal storage charging. This is because the battery can more effectively adapt to the operational profile of the heat pump and thermal storage system.

Table 4 Overview of the average annual electricity bill in 2030 for battery, heat pump and thermal storage control

Price Scheme	RBC	Smart Control	
		Flexibility 0.5 ℃	Flexibility 1 ℃
Day-ahead	1170.75 €	775.92 €	751.34 €
Bi-hourly tariff	1118.1 €	782.36 €	774.84 €

"The average daily peak power of households under control strategies that prioritize thermal storage charging over electric battery charging is illustrated in Figure 14 Average daily peak power spread for households with battery, heat pumps and thermal storage. Coordinated control of all flexible assets leads to an approximate 20% reduction in peak power demand compared to the conventional control of heat pumps. However, compared to the results in Section VI.3.1, these control strategies achieve a slightly smaller reduction in peak power. This is because, on sunny days, both thermal storage and batteries are charged to maximize the use of PV, which leads to higher peak power.

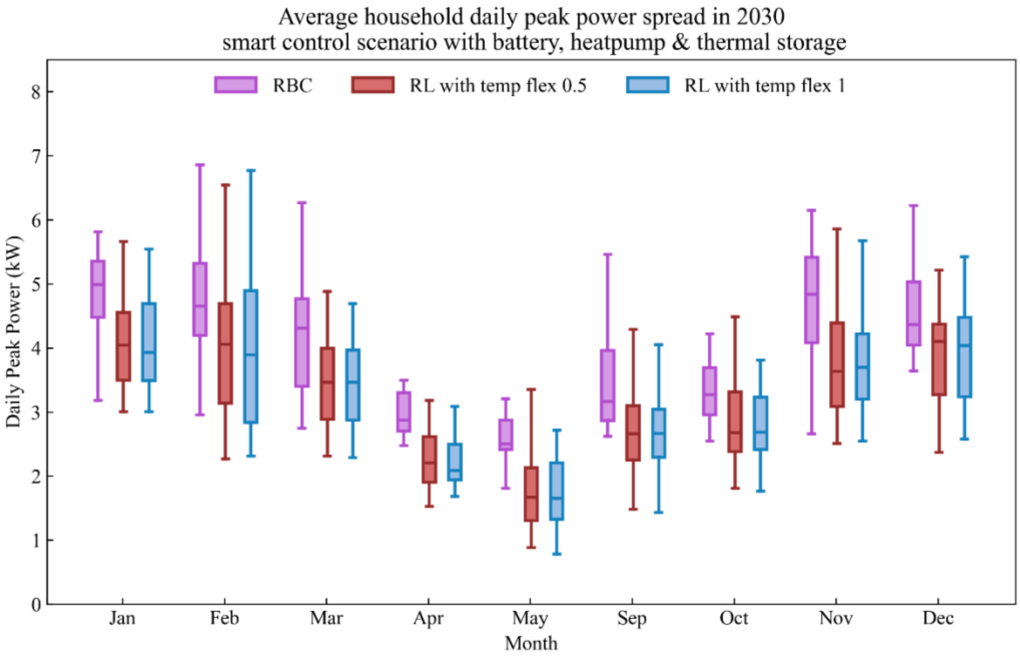

Figure 14 Average daily peak power spread for households with battery, heat pumps and thermal storage

VII. Conclusion and Next Steps

The results show that simultaneously controlling all flexible assets — including the battery, heat pump, and thermal storage — can significantly reduce the impact of heat pumps on residential loads. The smart control of these assets can reduce the average peak power by 24.8% and 26% for deadbands of 0.5 ℃ and 1 ℃, respectively, compared to the business-as-usual control of heat pumps. The smart control of all assets results in a 35.1% reduction in the electricity bill compared to a scenario in which heat pumps strictly follow the user-defined temperature setpoint. The results reveal that the cost reduction during warm months is primarily due to smart battery control, while during cold months, it is mainly achieved because of the control of the heat pump and thermal storage.

VIII. References and Internet Links

[1] G. Gokhale, B. Claessens, and C. Develder, "Sample efficient reinforcement learning for building control: Leveraging physics informed latent representations," in *Proceedings of the 15th ACM International Conference on Future and Sustainable Energy Systems*, 2024, pp. 496–497.

[2] F. Pavirani, G. Gokhale, B. Claessens, and C. Develder, "Demand response for residential building heating: Effective Monte Carlo Tree Search control based on physics-informed neural networks," *Energy Build*, vol. 311, p. 114161, 2024.

[3] G. Gokhale, B. Claessens, and C. Develder, "Physics informed neural networks for control oriented thermal modeling of buildings," *Appl Energy*, vol. 314, p. 118852, 2022.