Energieforschungsprogramm

Publizierbarer Endbericht

Programmsteuerung:
Klima- und Energiefonds

Programmabwicklung:
Österreichische Forschungsförderungsgesellschaft mbH (FFG)

Endbericht
erstellt am

26.07.2026

Projekttitel:
EnergyDec
Decision-Making and Optimisation for Distributed Energy Management

Projektnummer: **FO999888428**

Ausschreibung	7. Ausschreibung Energieforschungsprogramm
Projektstart	01.01.2022
Projektende	31.12.2025
Gesamtprojektdauer (in Monaten)	48 Monate
ProjektnehmerIn (Institution)	Wirtschaftsuniversität Wien Forschungsinstitut für Kryptoökonomie
AnsprechpartnerIn	Univ.Prof. i.R. Dipl.-Ing. Dr. Hermann Kaindl
Postadresse	Perspektivstraße 4, 1020 Wien
Telefon	-
Fax	-
E-mail	hermann.kaindl@wu.ac.at
Website	https://www.wu.ac.at/cryptoeconomics

EnergyDec
Decision-Making and Optimisation for Distributed Energy Management

AutorInnen: Julian Binz, Aleksey Bratukhin, Benedikt Herold, Adél Hartvig-Kelemen, Hermann Kaindl, Roland Kuras, Thomas Reisinger, Thilo Sauter, Albert Treytl, Stefan Wilker

1 Table of Content

2 Introduction

Energy management systems primarily focus on optimising individual buildings. The goal is, therefore, to optimise energy consumption at a given location based on its technical capabilities and level of self-generation.

The core innovation of the project is an integrated approach to optimisation with cross-cutting distributed decision-making, which was developed and evaluated for the energy management of multiple buildings with a wide variety of structures. The goal of the project was to maximise the utilization of renewable energy from a holistic perspective across multiple sites through distributed decision-making, based on a multitude of such factors.

For example, a school with large PV roof areas and daytime demand is very different from a residential quarter with heat pumps and evening peaks. Coordinating such sites across the public grid creates additional flexibility, increases renewable utilization, and reduces simultaneous peaks that would otherwise burden the network.

To this end, the project developed a new approach to distributed decision-making that can incorporate combined top-down and bottom-up decisions as well as peer-to-peer decisions at the same level. This approach was innovatively integrated with optimisation for individual buildings for cross-site decision-making across multiple diverse buildings.

The methods employed encompass modeling, simulation and optimisation. They have been integrated with our new approach to distributed decision-making.

This report is organised in the following manner. First, we present the content on our decision-making architecture, the BIFROST simulation environment, coupling BIFROST with the Community Controller (the optimiser from TU Wien), optimisation using, a new grey-box model for heat demand, and our running example for final evaluation of decision-making. Then we show our results on the grey-box model for heat demand, geographically distributed properties, and the final evaluation. Based on that, we draw our conclusion. Finally, we give an outlook and propose recommendations.

3 Presentation of Content

First, we present our new decision-making architecture. Since we investigated this approach in simulations, we briefly present the tool used for our simulations and its integration with the optimiser used. This simulation environment as given was not able to make optimisations and decisions on the basis of temperatures inside buildings. Hence, we developed a model for heat demand of buildings. For the final evaluation of decision-making in our project, we describe the example used in detail.

3.1 Decision-Making Architecture

In complex energy systems, decisions should be made through a combination of top-down and bottom-up approaches, as well as through peer-to-peer interactions at the same level. To this end, we developed a new decision-making architecture, as illustrated in Figure 1. In fact, this figure shows only an excerpt of two Decision-Making Units (DMUs) that are connected for both control flow and data flow. One DMU is higher in a hierarchy and the other lower. This excerpt concerns only this binary relationship between these two DMUs. Each of them has inputs and outputs via interfaces with other software. The Upper DMU can communicate its decision (II) via output and/or forward it to the Lower DMU. The Upper DMU can also specify a problem in the form of *Objectives*, *Preferences* und *Constraints*, and instruct the Lower DMU to perform a task (I).

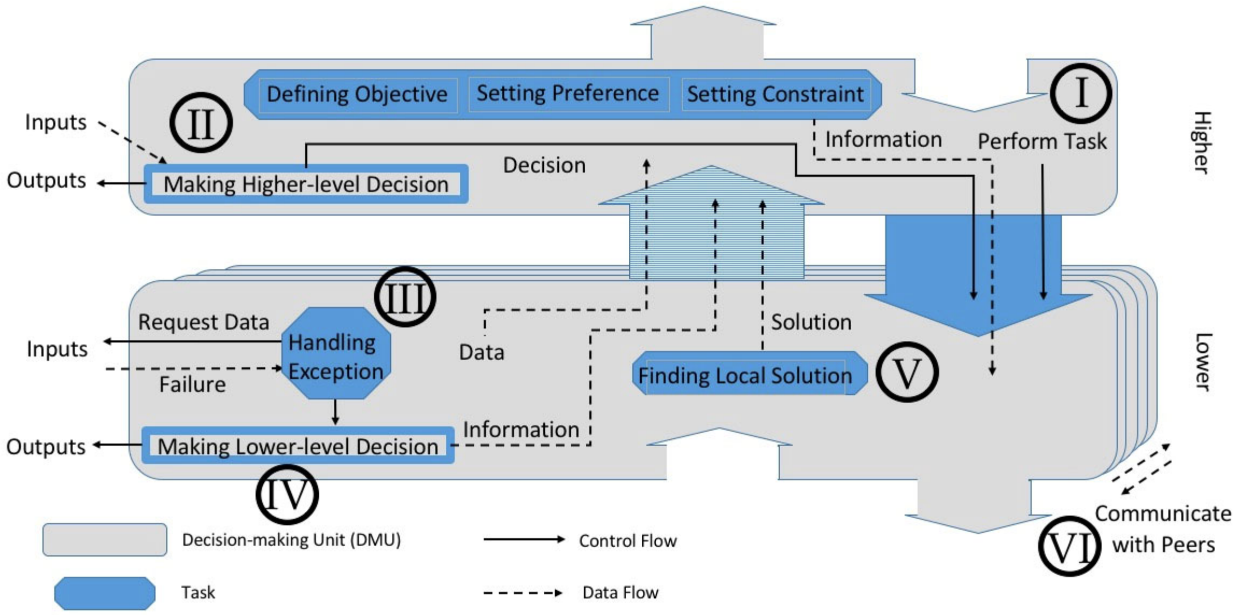

Figure 1: Excerpt from two Decision-Making Units (DMUs)

For such a task, the lower DMU must find a local solution (V) and communicate it upward. In addition, the lower DMU can handle an exception (III) or an error. In this regard, it should make a lower-level decision (IV) and communicate it via the output of other software to the higher DMU. In general, there will be several Lower DMUs of this type for each Upper DMU. These Lower DMUs are considered peers that can communicate with one another (VI) like autonomous agents.

Each DMU can implement any type of automated decision-making, ranging from stored deterministic decisions to methods drawn from operations research, decision theory, cognitive science, or artificial intelligence (e.g., utility-based, see [4]). Whatever is already being used effectively can be integrated here, as can new methods and approaches. Our architecture combines its decisions as explained above.

3.2 BIFROST Simulation Environment

BIFROST[1] is the simulation backbone used in EnergyDec to build and evaluate realistic settlements, communities, and distributed infrastructure scenarios [2]. Its strength is the combination of a visual settlement model, a configurable simulation engine, and interfaces through which external modules can interact with the data model during runtime. This makes it suitable not only for isolated building simulations, but also for multi-instance scenarios in which several communities are coupled through distributed optimisation (Figure 2).

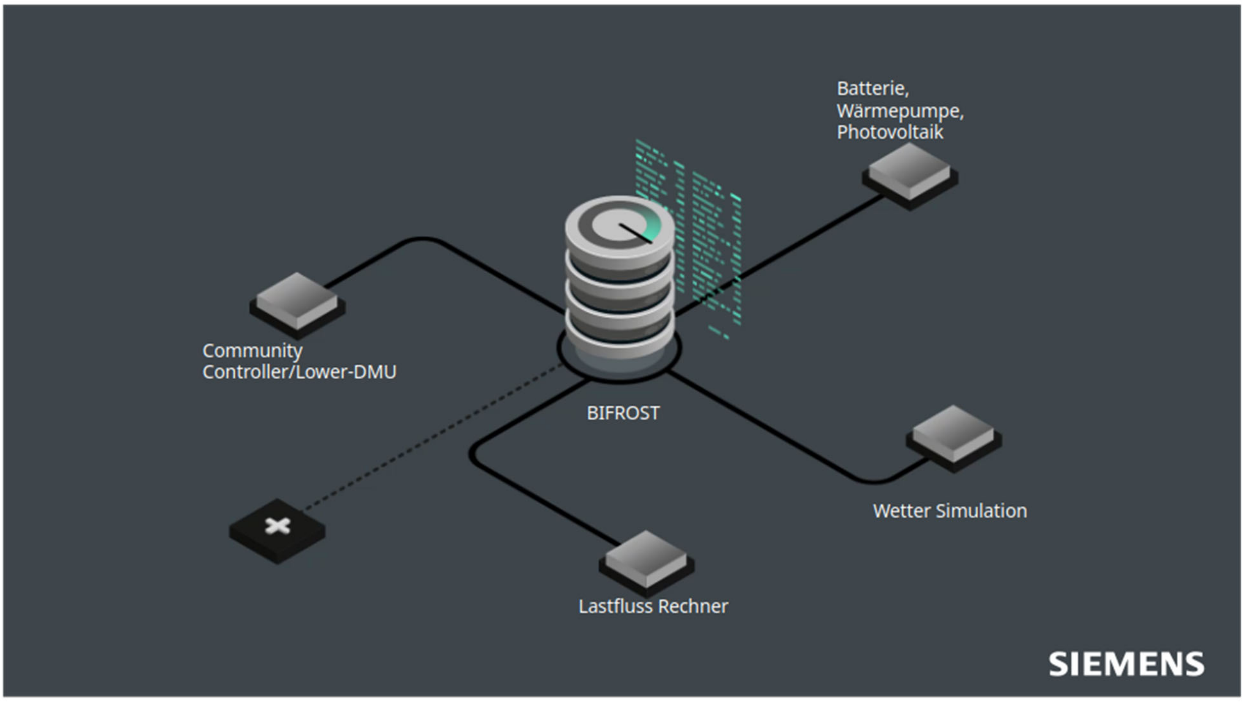

Figure 2: General BIFROST core concept

Within the native BIFROST interface, simulation outputs can be inspected as time-series plots for individual components. For the EnergyDec workflow, however, additional post-processing and cross-scenario analysis were required. Hence, BIFROST serves as the source of exported metadata, settlement descriptions, hardware configuration data, and generated time series, while an additional data-processing layer ingests these exports and exposes them in a dedicated graphical user interface. This enables flexible evaluation of time series, combination of curves, KPI-oriented analysis, and rapid comparison of simulation outcomes across use cases.

For the project use cases, each participating site can be represented by its own BIFROST instance (Figure 3). These instances are linked to the Upper-DMU, which synchronizes otherwise asynchronous simulations by waiting for the current simulation step from all parallel instances before triggering the next

[1] "Narrative simulation for Smart Energy scenarios."
https://BIFROST.siemens.com/en

distributed optimisation step. In that way, local site simulations are reassembled into a coherent distributed experiment in which the Upper-DMU can process transformer loads, compute coordination targets, and send updated control parameters back to the Lower-DMUs.

Figure 3: General example view of a BIFROST settlement and simulation result plots

3.3 Optimisation Using the Community Controller

The optimisation for a Lower-DMU uses a pre-existing *Community Controller* originally developed in the cFlex project (FFG 871657) for the optimal use of flexibility assets such as battery storage, controllable loads, photovoltaic generation, and heat pumps (Figure 4). For its used in the EnergyDec project, the pre-existing optimisation core was extended by a dedicated interface to DMUs and by an optional human-in-the-loop decision layer.

At the operational level, the Community Controller solves a local optimisation problem that balances three main objectives: self-optimisation of the community, reduction of steep power gradients, and smoothing of peak loads. Via a REST-based interface, it collects the flexibility information available for the next optimisation horizon and computes schedules and control actions according to the selected objective weighting.

An important implementation advantage has been that the Community Controller was already integrated into BIFROST and can directly use its simulations and weather forecasts. Since BIFROST serves as the virtual testbed for the evaluation, this prior integration enabled a seamless coupling between simulation, local optimisation, and hierarchical coordination.

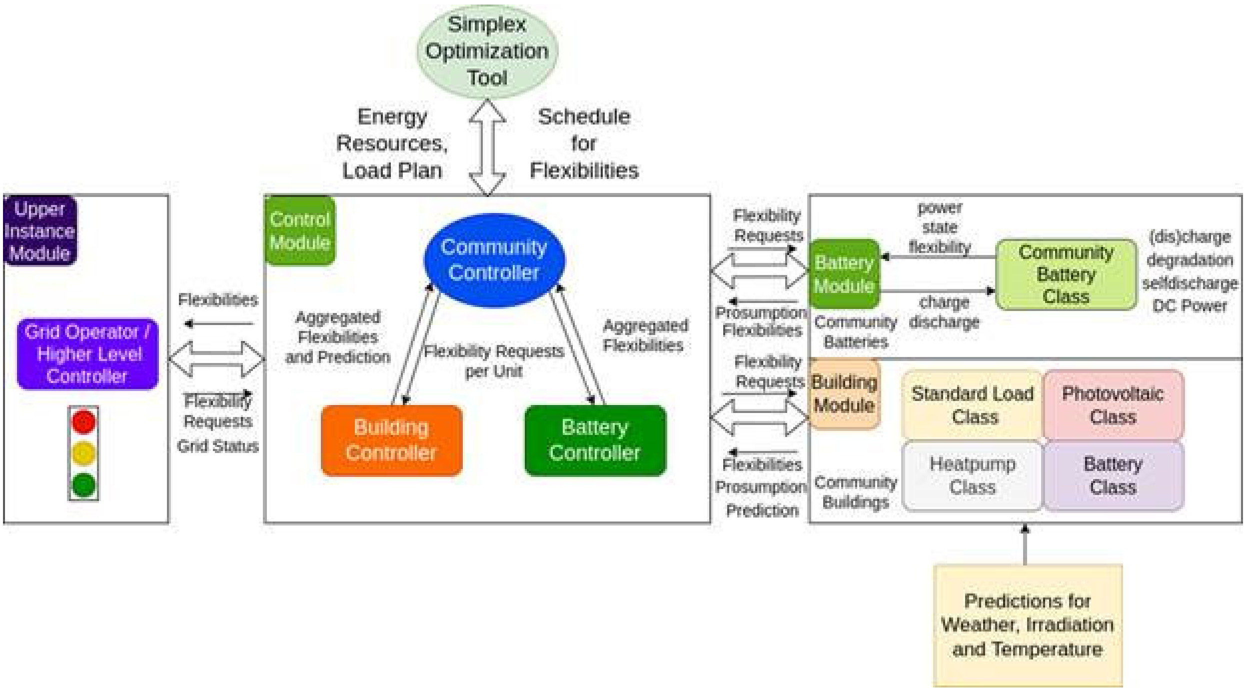

Figure 4: Conceptual representation of the Community Controller originating from cFlex

The Community Controller receives locally available forecasts and flexibility descriptions, accepts scenario parameterization from the Lower-DMU interface layer, and produces technically feasible schedules for the connected devices. The Lower-DMU as a whole is responsible for orchestrating this process, mapping candidate parameter sets to Community Controller inputs, triggering the corresponding optimisation runs, and passing the selected result forward for execution.
For the implementation developed during the project, the Community Controller exposes an interface that allows the Lower-DMU to adjust optimisation setpoints and related parameters at each simulation step (Figure 5). These higher-level signals influence the local optimisation, but do not replace it: local forecasts, device constraints, and community-specific operating conditions remain part of the optimisation process. This is especially relevant in situations such as grid congestion, where targeted avoidance of additional feed-in or load peaks may be required.

.

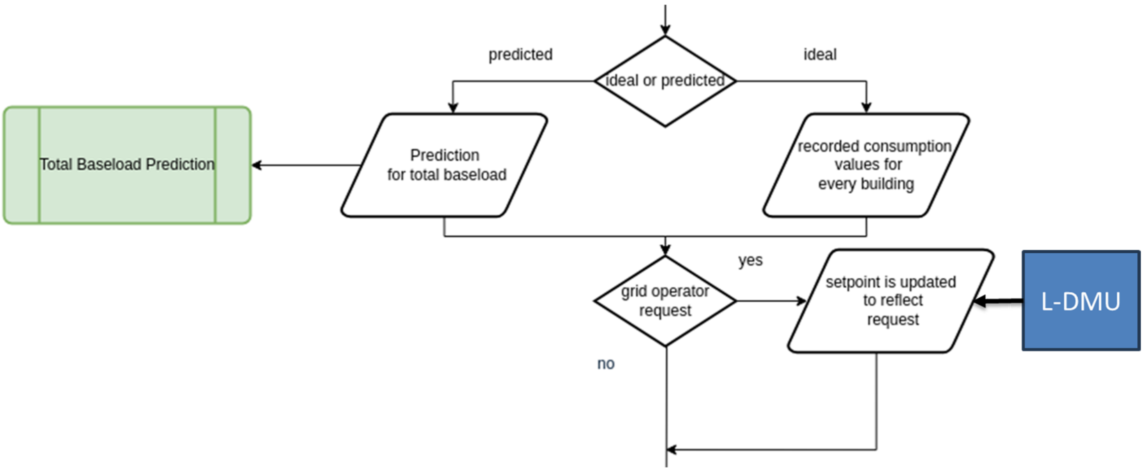

Figure 5: Flow chart of Community Controller and L-DMU interface for setpoints

Actually, multiple candidate parameter sets can be processed in this way. More precisely, the Lower-DMU can invoke the Community Controller *n* times, with *n* different parameter sets, and thereby generate *n* alternative optimisation schedules for the following 24-hour horizon (Figure 6).

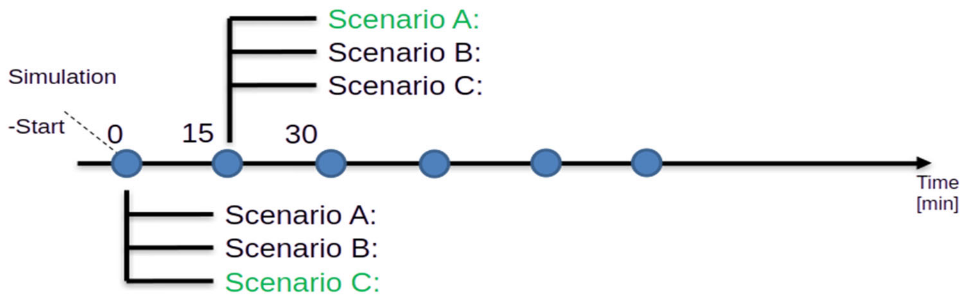

Figure 6: Overview of the simulation process with multiple scenarios per step

The Community Controller supports three weighted objective dimensions: reduction of peak load, reduction of gradients, and tracking of a transformer-level setpoint. By adjusting these weights, the Mixed-Integer Linear Programming optimisation can be steered toward desired operational behaviour. In its default mode, the controller generates a rolling 24-hour schedule at every 15-minute simulation step based on its own forecasts [1].

The different scenarios can be viewed as looking-ahead into different possible “worlds”. The results can then be used to decide on which of these “worlds” is the desired one to be achieved and to take the corresponding actions in the sense of, e.g., reducing the desired room temperature in order to save energy.

3.4 Grey-Box Model for Heat Demand

To enable more accurate simulation of buildings with regard to energy management and, in particular, heat demand, various approaches were implemented for the development of a grey-box model. The goal was to achieve better results for EnergyDec through more accurate simulation of building energy management and to enable a more robust evaluation. Figure 7 schematically shows which factors influencing heat demand must be taken into account.

3.4.1 Static Heat Demand Model

A static heat demand model was developed, implemented, and evaluated to simulate the annual heat demand of households. The goal was to map a realistic consumption profile based on weather- and building-specific parameters. The model was integrated into Bifrost as a Python module and run in parallel with the existing heat demand model.

The characteristics of the static model include heat losses through building components (conduction, radiation, convection), solar gains, hot water demand, electrical loads, human body heat, and ventilation losses. In addition, dynamic model components that take into account factors such as the thermal capacities of building components and temperature variations were also included in the model.

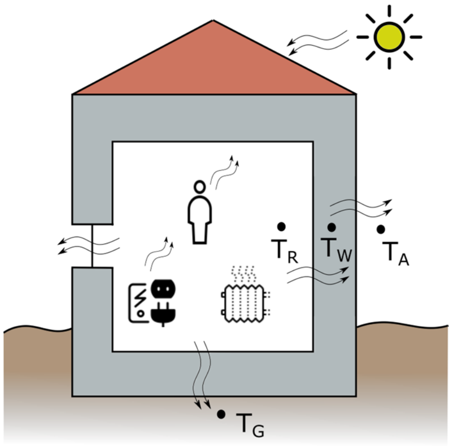

Figure 7: Heat Demand Model for Building

For the evaluation, the model was simulated over the course of a year under various weather conditions and building parameters. Four households with different insulation standards and living areas (200–500 m²) were analysed. The resulting average annual consumption was approximately 84.3 kWh/m² (heating and cooling energy) or 75.5 kWh/m² (heating energy only), placing it in the upper range of expected values. The analysis was performed using a custom-developed script to evaluate the simulation data, as the BSX web interface does not provide exportable individual data. The model was implemented as a standalone Python module and can be run via the Bifrost instance with room temperature specifications.

3.4.2 Dynamic Heat Demand Model

The static heat demand model was expanded to include a dynamic model that accounts for the difference to the target temperature and how that difference changes over time. The model was implemented as a standalone Python module and can be used in parallel with the existing model. No adjustments to the Community Controller were necessary, as the model uses electrical power ratings that distinguish between heating and cooling demands.

The dynamic model includes an integrated linear heat pump model with three different efficiency levels for heating, domestic hot water, and cooling. These efficiency levels are based on typical flow temperatures (35°C for heating, 55°C for domestic hot water) and depend on the outdoor temperature. A switching logic has been implemented to switch between heating and cooling modes, which selects the operating mode based on the average energy demand for the next 24 hours and taking hysteresis into account. An inertia logic prevents frequent switching, which corresponds to real-world behavior.

To ensure realistic indoor temperatures, threshold values were introduced that simulate ventilation when temperatures are too high or too low. These are currently set as fixed values, as fresh air requirements and temperature equalization impose different constraints, which necessitates more sophisticated modeling.

Another focus was the detailed modeling of the thermal capacities of room and wall masses. The set temperature, actual heat input/output, and separate temperatures for the room and wall were taken into account. The calculation logic was divided into four methods and moved to a separate module, enabling independent testing with Pytest. Only the electrical power is transmitted for communication with the Community Controller; differentiation between heating and cooling modes is handled internally within the model.

The evaluation of the dynamic model yielded realistic consumption values: The average annual consumption was 40.3 kWh/m²—significantly lower and more realistic than in the static model. The daily profiles reflect weather dependence, ventilation behavior, and heating capacity limits. The simulation time for a household with 96 forecasts was only 2–3 ms.

Overall, a robust, weather- and usage-dependent dynamic heat demand model with realistic load forecasting, flexible operating mode selection, and good integration into the existing simulation infrastructure was successfully implemented.

3.5 Running Example for Final Evaluation of Decision-Making Integrated with Optimisation

For presenting the work in the last period, we use the following running example, for which many experiments in the simulation environment have been made, which included a new implementation of automated decision making.

Figure 8 shows an instantiation of our generic decision-making architecture with a hierarchical structure. H-DMU is a “higher” decision-making unit with the primary purpose of coordination its assigned “lower” decision-making units. In this instantiation, there are three of them.

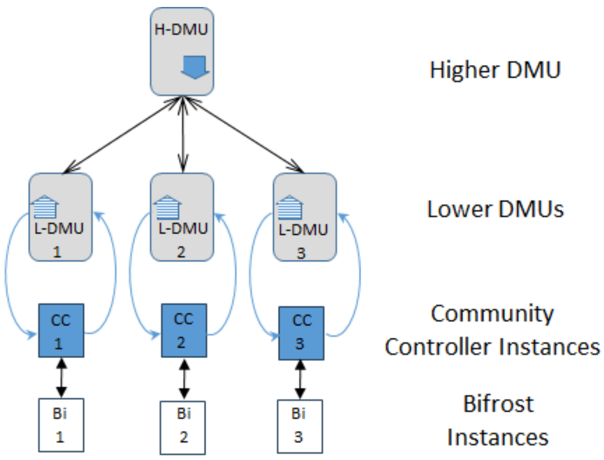

Figure 8: Instantiation of the Architecture and Connection with the Community Controller and BIFROST

Each of these L-DMUs represents some real estate. Note, that for different scenarios of the use case, these can be different real estates. In particular, these may be located close to each other, or they can be in distant locations. Having one in the Western part of Austria and another one, say, in Vienna, even different weather conditions can be expected, which gives rise to interesting cases of decision making. Figure 8 also illustrates the connection with the *Community Controller* (CC), which is the tool employed for elaborate and sophisticated *local optimisations*. In fact, each of the three L_DMUs has its own CC instance assigned to it.

The CC instances actually optimise in a simulation environment, which is in our project Bifrost. More precisely, as illustrated in Figure 8 as well, each CC instance has its own Bifrost instance assigned to it. Each L-DMU operates at the level of a single energy community. Its objective is to optimise local electricity consumption and production over a 24-hour horizon with a 15-minute time granularity. The L-DMU processes detailed forecasts of consumption, generation, flexibility ranges, and local constraints to determine how the community can shift or modulate its prosumption (combined production and consumption). The outcome is a set of *flexibility envelopes*, expressed as upward and downward ranges, indicating how much the community can increase or decrease its net power exchange with the grid at each time step. These envelopes serve as the actionable interface to the upper level. Rather than exchanging detailed appliance-level data, communities expose only how far their prosumption can be adjusted at each time step. This abstraction preserves privacy, reduces data volume, and enables scalable decision-making across many communities.

The H-DMU coordinates the three communities represented by the L_DMUs simultaneously. Using the aggregated flexibility envelopes provided by all L-DMUs, it performs *system-level optimisation* across the same 24-hour horizon. Its role is to balance global objectives that exceed the scope of individual communities. Three optimisation criteria are taken into account:

1. **Load Balancing:** smoothing the aggregated load profile to reduce peaks and mitigate grid stress.
2. **Energy Costs:** shifting flexible loads to periods with lower electricity prices to minimize total operational cost.
3. **CO_2 Emissions:** leveraging periods with lower CO_2 intensity in the electricity mix to reduce environmental impact.

The H-DMU continuously evaluates trade-offs among these criteria, using the available flexibility ranges to dispatch coordinated setpoints to each community. This ensures that system-level goals are achieved without violating local constraints.

In more detail, the H-DMU aggregates all incoming predictions from the L-DMUs assigned to it. Using these data, the H-DMU computes the optimal joint energy profile across communities, considering the selected optimisation criterion (load balancing, energy costs, or CO_2 minimization). This is implemented via a *model predictive control* (MPC) formulation solved with convex optimisation. The output consists of *prosumption targets* (setpoints) for each community for the 24-hour horizon (Shape: 96 × n, where $n = 3$ is the number of Community Controllers).

Based on optimisations according to these three criteria, a founded and rational **decision** has to be made on a related strategy to be followed by the system comprised by these DMUs, which H-DMU communicates then to its three assigned L-DMUs. The idea is to decide for the strategy corresponding to the criterion with the overall best result, where the other criteria also achieve “good” results.

For implementing this idea, we use a Weighted Sum Approach with a single objective defined as follows for combining these three optimisation criteria:

$$\min_{x} \quad w_1 \cdot f_1(x) + w_2 \cdot f_2(x) + w_3 \cdot f_3(x)$$

Where:

$$w_1 + w_2 + w_3 = 1$$

We set these weights based on policy, user preference, and regulatory priorities.

Let us explain how this decision is being made using the following example as summarized below:

Weights	0.5	0.3	0.2	
Optimisation Criterion	Evaluation Function	Weighted Sum		
	Balance	Price	CO_2	
Balancing	**1**	5	10	**4**
Price	5	**2**	7	4.5
CO_2	7	8	**2**	6.3
Decision for	Balancing			

These data result from the following optimisation runs and the decision making based on them:

Run 1: Optimisation with the target Balancing

Scores: Balancing = 1, Price = 5, CO_2 = 10

Weighted score = 4

1*0.5+5*0.3+10*0.2 = 4

Run 2: Optimisation with the target Price

Weighted score = 4.5

Run 3: Optimisation with the target CO_2

Weighted score = 6.3

Decision is made for “Balancing”.

This weighted aggregation of the objective values allows the H-DMU to identify the strategy that provides the **best overall trade-off between economic, environmental, and grid stability objectives**.

The strategy decided for is communicated by the H-DMU to its three assigned L-DMUs. This is done via sending the selected target trajectory for each CC instance to its corresponding L-DMU. These targets represent the proposed setpoints that should be implemented locally unless they violate important community-level constraints or preferences.

4 Results and Conclusions

First, we present the results of our evaluation of the heat demand model. Then we show results of our approach to implement and use geographically distributed properties. With our simulation environment, we investigated how it can be used for finding an optimal battery configuration. Finally, we present the results of our final evaluation of decision-making as integrated with optmisation.

4.1 Evaluation of the Grey-Box Model for Heat Demand

To evaluate the developed grey-box model (GBM), a comparison was made with the WPuQ dataset, which was created as part of the Wind-Solar-Heat Pump District project. The aim of this project was to analyze the behavior and effects of operating multiple heat pumps within a local power grid. The dataset includes detailed electrical load profiles from 38 single-family homes and their heat pumps in Hamelin, Lower Saxony, as well as aggregated measurements at the substation level. Since the WPuQ dataset reflects heating and hot water demands, the GBM was considered exclusively in heating mode. The results show that the GBM reflects realistic behavior (Figure 9): The simulated annual consumption of the settlement is 162,528 kWh compared to 170,602 kWh in the measured data, which corresponds to a relative error of less than 5%. Given the simplified building modeling and limited data availability, this represents a good result to be used as a basis for the simulations in our project.

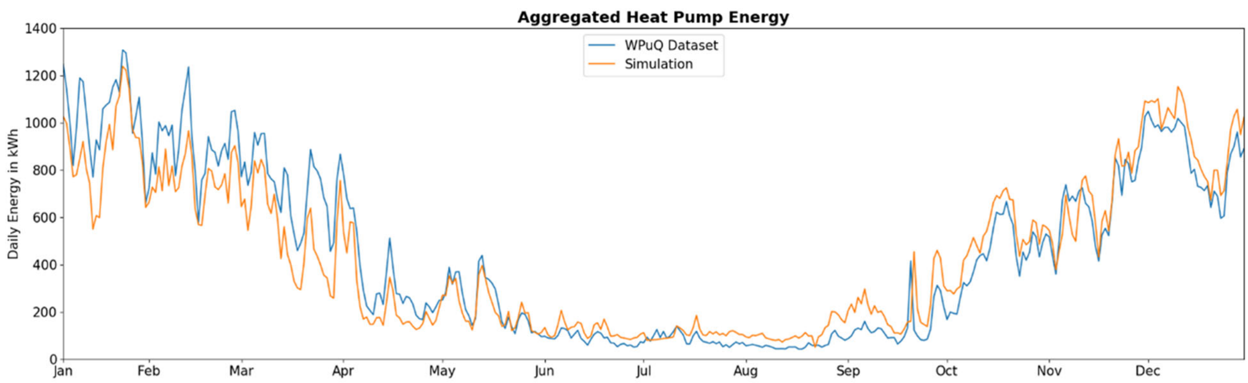

Figure 9: Aggregated Heat Pump Energy

In addition, two settlements with identical layouts were simulated using location-specific weather data from East Tyrol and the Tyrolean Oberland. Based on these results, the models and algorithms in the DMU were optimised. The simulation results include both the electrical power supplied to the respective grid stations in the two settlements and the optimised schedule for energy transfer between the locations, calculated by the Upper DMU (Figure 10). The results indicate existing potential for targeted energy exchange between geographically distributed settlements. In particular, they demonstrate how the coordinated use of thermal flexibility can lead to an improved energy balance for the energy community.

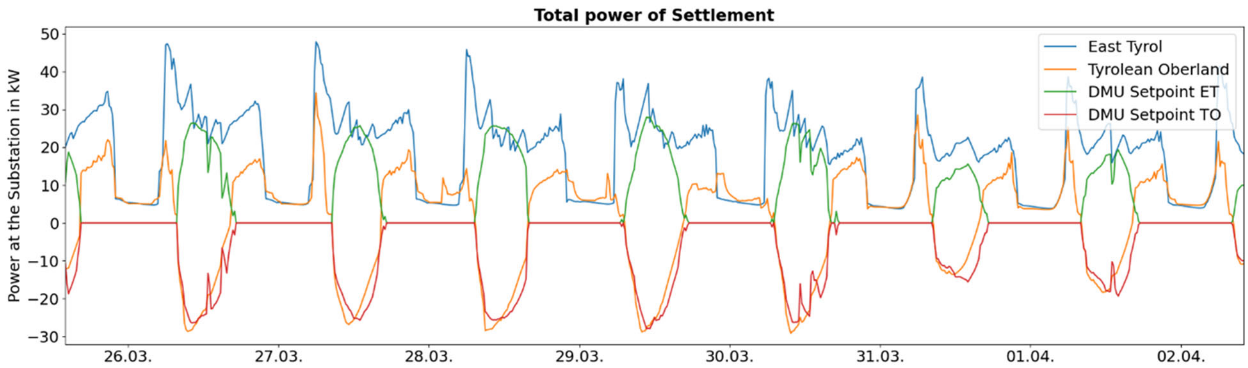

Figure 10: Total Power Of Settlement

4.2 Geographically Distributed Properties

In BIFROST, only a single scenario—particularly with regard to weather conditions—can be simulated at any given point in the simulation. In EnergyDec, however, it was necessary to simulate different weather conditions simultaneously—for example, in western and eastern Austria—so that distributed decision-making and, consequently, distributed energy management can be analyzed in such situations.

Therefore, we adapted the simulation environment in BIFROST so that users can quickly and easily switch between different weather and geographic configurations. BIFROST's weather module requires historical weather data for the respective location in order to account for weather conditions during the simulation. Weather data was downloaded via Open-Meteo, which supplies all meteorological variables required by the model.

The Community Controller's weather forecast module requires the geographic coordinates of the respective location as well as the minimum and maximum temperatures from historical data, and operates according to the clear-sky principle. This means that, for the purposes of cloud forecasting, a cloudless sky is always assumed initially. If the current weather differs from this, a corresponding adjustment is made. Once the necessary weather configurations have been defined for each relevant location, the appropriate configuration for weather and location can be selected in the user interface at the start of a simulation.

The technical implementation involves running two BIFROST simulations with different weather configurations, with a new BIFROST module simulating the energy exchange between the two properties. Each of these properties is assigned a DMU, with both being subordinate to a higher-level DMU. This allows for the coordinated exchange of energy among multiple properties, with our primary focus in EnergyDec being the avoidance of CO_2 emissions.

To assess the potential for shifts resulting from the geographically distributed approach, two BIFROST simulations with identical configurations but different weather conditions were compared. For this purpose, weather data from Berndorf (Lower Austria) and Bludenz (Vorarlberg) provided by Open-Meteo were used as examples, and the defined configuration of a property was applied. The rated capacities of the heat pumps were sized in BIFROST so that the load profile was of the same order of magnitude as the dataset from AMPEERS ENERGY used in earlier simulations. To specifically compare the potential for shifts between distributed properties and to account only for the influence of weather, identical hardware configurations were used for both properties (Table 1).

Table 1: Hardware configuration of the distributed properties.

	Gymnasium	Property 1	Property 2	Energy community
PV system	300 kWp	36 kWp	47 kWp	
Heat pump	-	48 kW	63 kW	
Electric storage				100 kWh

To this end, a week was selected in which weather conditions varied as much as possible and were ideal for the PV system. Since the efficiency of PV panels is higher at lower temperatures, this parameter was also taken into account alongside solar radiation and cloud cover. Under the ideal assumption that there are no transmission costs or losses, but only the current demand is met, and the forecasts match the historical data, this resulted in a shifting potential (green area) of 1.28 MWh for the data from the first week of March 2021 (Figure 11).

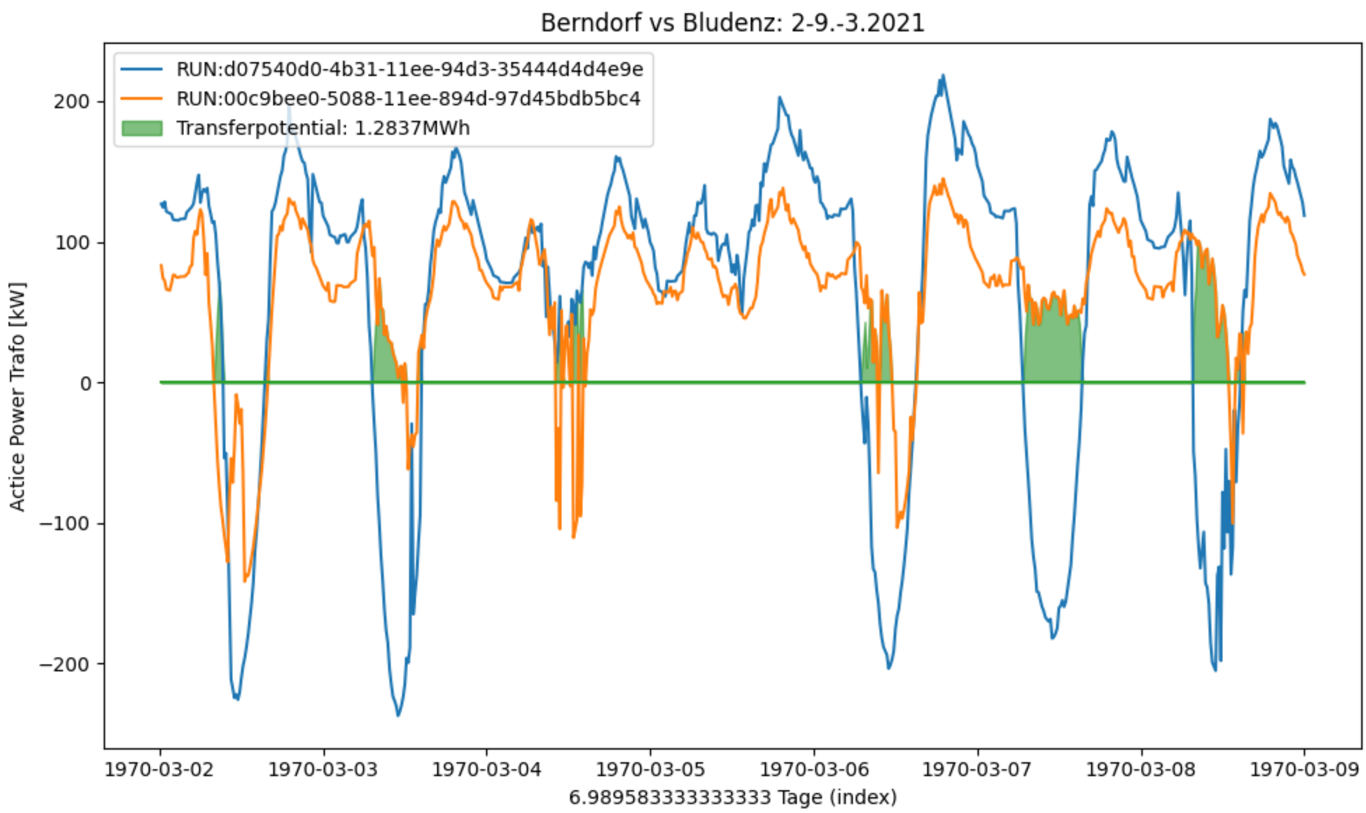

Figure11: potential for shifts between two geographically distributed properties.

4.3 Finding an Optimal Battery Configuration

With our simulation environment, we investigated how it can be used for finding an optimal battery configuration. The main objective of integrating battery storage into a PV system is to increase self-consumption of locally generated electricity. Since the price of grid electricity can be higher than the remuneration received for feeding electricity into the grid, households benefit economically when they consume their own PV electricity instead of purchasing electricity from the grid [3]. The goal of this simulation was to determine gains that a battery installation brings in the energy arbitrage. That allows us to estimate the optimal configuration of the battery size in terms of Return-of-Investment (ROI).

The average energy price rate of change determines the scores in the cost optimisation since the algorithm tends to trade the energy according to the prices. Table 2 shows the annual rate of change in electricity prices over the last five years. 2022 is an anomaly due to the high prices in general and price fluctuations in particular, due to the energy crisis caused by the political events in Europe. Therefore, this year is not representative of the general trend. It can be assumed that due to increasing share of renewables in the energy basket, the slow trend of increasing fluctuations will continue and 11.5 hourly rate of change is a good indication for analysis.

Table 2: Average rate of change in electricity prices in Austria

Year	Hourly rate of change of electricity prices (day-ahead market)
2021	8.38
2022	20.05
2023	10.22
2024	11.95
2025	11.5

The simulations were performed on the 2021 year data with average fluctuations of 8.38. In particular, the simulations were performed on January (3.6), August (6.3) and November (14.4) data (Table 3). The simulations highlighted that the size of the community is not a factor in the monetary gains from the arbitrage and the main factor is the charge rate of the battery that determines how fast the arbitrage can be performed.

Table 3: Battery gains vs price rate of change ratios

	Price rate of change	Battery gains Euros/kW	Ratio of battery/price
August	6.3	2.1	3
November	14.4	4.32	3.33
January	3.6	2.11	3.32

Figure 12 shows the monthly profit per kW of battery capacity depending on the charge rate. As shown in the figure, the profit flattens if the battery can be fully charged or discharged within the hour. That is because the granularity of prices is 1 hour. The chart also shows that the monetary gains are proportional to the prices rate of change and therefore, we can apply the proportions to the current trends of the price fluctuations to estimate the possible monetary gains.

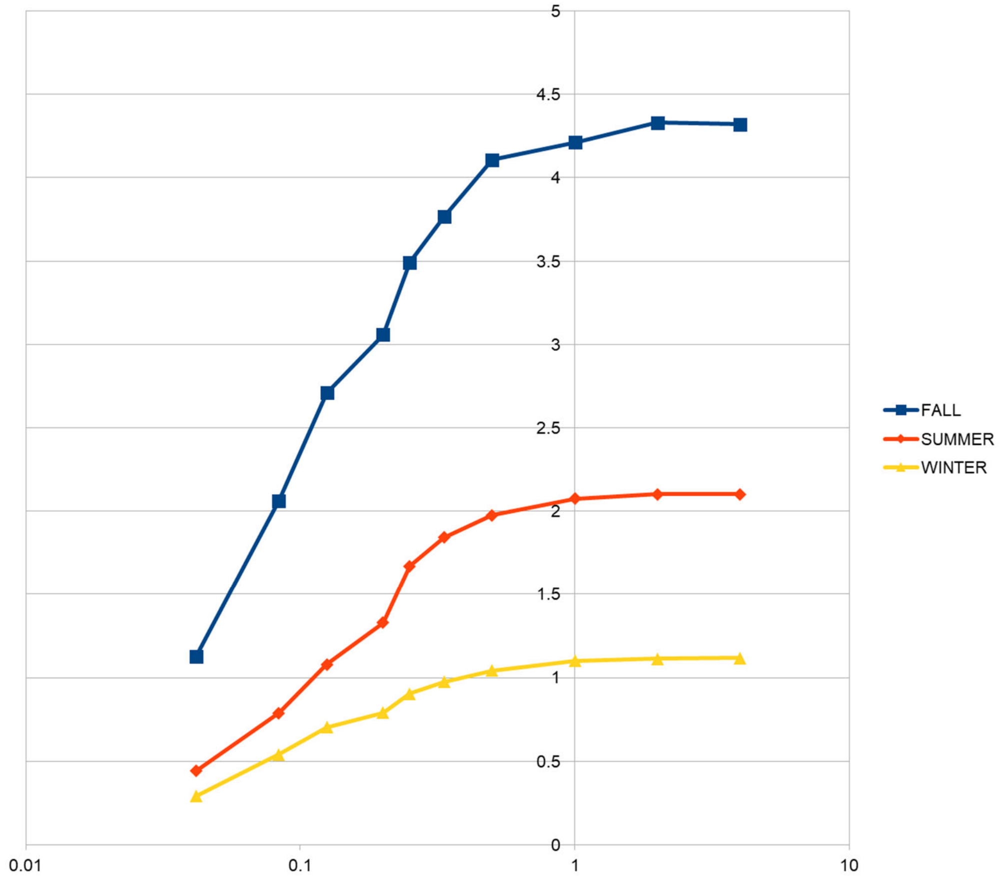

Figure 12: Monthly profit per kW of battery capacity depending on the charge rate

The simulation shows that the ratio between the battery gains per kW and the rate of the change of prices remain stable with an average ratio of 3.18. That allowed us to estimate the possible gains for the whole year. For 2025, estimated gains are 3.61 euros per 1kW of battery capacity per month and 43.32 euros per year.

Considering that the average price for battery storage are as low as 125 euros per 1kW of installed capacity, the Return of Investment is less than 3 years. Furthermore, the prices of batteries continue falling, which will make using batteries even more profitable.

4.4 Results of Final Evaluation of Decision-Making Integrated with Optimisation

According to the use case of our running example as described above, the evaluation is based on a comprehensive set of **480 simulation scenarios**, where the variability comes from varying environmental conditions, system configurations, and coordination levels between optimisation layers. To capture **seasonal differences in energy production and consumption patterns**, the scenarios are divided into **summer and winter conditions**. In total, **300 scenarios represent summer operation**, characterized by high photovoltaic (PV) production and moderate demand, while **180 scenarios represent winter operation**, where PV generation is significantly lower and energy demand—particularly for heating—is higher.

In addition to seasonal variation, the scenarios differ in the **level of knowledge and control that the Higher Decision Making Unit (H-DMU) has over the Lower Decision Making Units (L-DMUs)**. These variations allow the evaluation of different coordination architectures ranging from fully centralized to more decentralized decision making. For the summer cases, the scenarios are evenly distributed across three coordination settings:

4. **Full knowledge and control:** The H-DMU has complete visibility of local system states and optimisation variables and can directly influence or override L-DMU optimisation results.
5. **Partial knowledge/control:** The H-DMU receives summarized or aggregated flexibility information from L-DMUs and can influence system-level decisions but does not fully control local optimisation processes. The main difference to full control is the knowledge over the battery state that reduces optimisation scores due to inability of the L-DMU to fulfil the goals set by the H-DMU because of the battery limitations unknown to the H-DMU.
6. **Limited knowledge/control:** L-DMUs perform local optimisations autonomously and communicate only limited results or flexibility envelopes to the H-DMU. The main difference between the partial and limited control is H-DMU knowledge of the battery state. In the latter, if the battery is used for charging, it cannot be used for charging in the next step.

All simulation runs have been executed in **True Prediction** mode, where weather and load predictions match the actual values, yielding a prediction error of zero. Prediction accuracy itself has been validated in prior projects (cFlex and SONDER) and is not the subject of this evaluation.

All three Bifrost instances share the same settlement configuration and differ only in their weather profiles. Weather is the primary driver of PV and Heat pump based flexibility, and this single-factor variation allows controlled analysis.

For the sake of comparison, simulations with **No Optimisation** have also been run — Prosumption without control; batteries unused.

The L-DMUs prioritize following the setpoints communicated from the H-DMU. The peak shaving weight is reduced to 20, allowing the H-DMU's optimisation criteria to dominate.

Parameter	Value
Peak Shaving Weight	20
Setpoint Weight	1920
Consumption	383.28 MWh/year
Installed PV	235.15 kWp
Battery (Fenecon Industrial L)	450 kWh

For each scenario, **three optimisation criteria** have been applied, as explained above:

1. **Cost optimisation** – minimizing total energy procurement and operational costs.
2. **CO_2 optimisation** – minimizing emissions associated with energy consumption.
3. **Load balancing / grid friendliness** – reducing peak demand and improving grid stability.

Each objective is first optimised individually to determine the best achievable outcome for that criterion. The **Higher Decision Making Unit (H-DMU)** then makes a decision as explained above.

We also compared the **distributed** decision-making approach with **local** peak shaving as a reference to analyze the impact of the hierarchical Upper/Lower DMU concept by comparing it against predominantly local optimisation.

Local Peak Shaving (Reference)
The L-DMUs retain maximum weight on local peak shaving, representing the default configuration without distributed coordination. Setpoint weight (H-DMU influence) is minimal.

Parameter	Value
Peak Shaving Weight	1920
Setpoint Weight	20
Consumption	383.28 MWh/year
Installed PV	235.15 kWp
Battery (Fenecon Industrial L)	450 kWh

Results
For each of the optimisation criteria used by the H-DMU individually, the runs with *Local Peak Shaving* show measurable improvements: *balancing* achieves 36.32% grid-load reduction, *cost* saves 1,249 €/year, and *CO_2* reduces emissions by 272 kg/year.

When switching from local optimisation to distributed coordination (setpoint mode), *balancing* performance drops from 36.32% to 1.11%, while *cost* and *CO_2* savings increase dramatically (6,073.9 € and 4,323.9 kg, respectively). This quantifies the trade-off between local flexibility and global coordination.

Summary

Aspect	Finding
Balancing optimisation	Up to 36% grid-load reduction (local peak shaving)
Cost optimisation	Up to 6,074 €/year savings (distributed decision making)
CO_2 optimisation	Up to 4,324 kg/year reduction (distributed decision making)
Local vs. distributed trade-off	Balancing drops from 36% to 1% when shifting to distributed decision making.

The hierarchical decision-making concept developed in this project demonstrates that distributed energy communities can be effectively managed for load balancing, cost minimization, and CO_2 reduction using a coordinated Upper/Lower DMU architecture. The configurable weight mechanism allows operators to define the trade-off between local autonomy and distributed coordination depending on the specific use case.

4.5 Conclusion

The core innovation of the project is an integrated approach to optimisation with distributed decision-making, which maximises the utilisation of renewable energy across multiple sites from a holistic perspective.

In particular, we found the potential benefit of energy management across different geographic locations even higher than expected, especially regarding savings of CO_2 and costs. However, these savings have to be balanced with peak shaving.

Our approach can also help to reduce investments in battery electric storage systems and increase the return on investment.

5 Outlook and Recommendations

The exploitation of the project results is planned primarily through integration into professional workflows rather than direct mass-market software deployment. We consider our approach as complementary to existing planning tools on the market, hence filling a gap in advanced system optimisation and analysis.

Building on our approach and its results, we recommend to pursue and further develop *optimising* distributed systems of buildings and properties, which *founded decisions* for energy management can be based upon. Our generic decision-making architecture can be instantiated in many ways for these purposes.

We recommend to provide *more funding for studies* and the implementation of *distributed energy management*, since we found large potential for its practical benefit.

In particular, we recommend providing *regulatory incentives for distributing generation and demand assets across different geographic locations*, which can benefit from different weather conditions (such as solar surplus at one site coinciding with cloud cover at another).

We recommend *subsidies to prioritise the installation of heat pumps and thermal storage*. These lead to high reductions of CO_2 emissions especially during winter, and they can also support cooling of buildings in summer, which will become more and more relevant in coming years.

We also recommend to set up *dynamic energy communities and energy management systems* across buildings to increase direct consumption of renewable energy when available, which will reduce the need for investments in energy storage.

We recommend to regulatory bodies *enforcement of vendor-neutral control standards and open communication protocols* across all flexibility assets, including storage systems, heat pumps, and EV chargers. This will greatly facilitate their optimised use in practice.

We also recommend to regulatory bodies the *creation of standardised legal pathways for sharing anonymised generation and consumption data*. The availability of such data is essential for reliable predictive power models.

We recommend providing standardized, machine-readable access to already published Distribution System Operator (DSO) data, such as published *coverage areas for energy community boundaries* via open APIs. Standardizing these data streams is essential for allowing third-party optimisation frameworks to align energy community setups accordingly.

We also recommend adopting the decision-making approach developed in this project to resolve *optimisation conflicts* resulting from DSOs enforcing dynamic power limits with the upcoming *operational envelopes*. In addition, successful deployment requires more than technical retrofitting: broad *user acceptance* hinges on automated comfort preservation, fair financial incentives, and transparent prosumer control.

Strategically, we recommend *close involvement of interest groups* in the processes and *harmonisation of the legal framework*.

6 References

[1] N. Efkarpidis et al., Proactive Scheduling of Mixed Energy Resources at Different Grid Levels, in IEEE Transactions on Sustainable Energy, vol. 15, no. 2, pp. 952-963, April 2024, doi: 10.1109/TSTE.2023.3320055
[2] R. Mosshammer, K. Diwold, A. Einfalt, J. Schwarz and B. Zehrfeldt, BIFROST: A Smart City Planning and Simulation Tool. In: W. Karwowski, T. Ahram (eds) Intelligent Human Systems Integration,

Proceedings of Advances in Intelligent Systems and Computing, vol 903., 2019, Springer, Cham., San Diego, CA, USA, doi:10.1007/978-3-030-11051-2_33
[3] M. Naumann, Karl, R. C., Truong, C. N., Jossen, A., and Hesse, H. C. (2015). Lithium-ion battery cost analysis in PV-household application. Energy Procedia, 73, 37-47.
[4] A. Rosenfeld and S. Kraus, Predicting Human Decision-Making: From Prediction to Action, Morgan & Claypool, 2018.

7 Contact Data

Project leader: Univ.Prof. i.R. Dipl.-Ing. Dr. Hermann Kaindl
WU Vienna
Research Institute for Cryptoeconomics
Perspektivstraße 4, 1020 Wien, hermann.kaindl@wu.ac.at;
https://www.wu.ac.at/cryptoeconomics

List of further project partners:

Dipl.-Ing. Stefan Wilker
TU Wien
Institute of Computer Technology

Dipl.-Ing. Albert Treytl
University for Continuing Education Krems
Department for Integrated Sensor Systems

Dr. Roland Kuras
PowerSolution Energieberatung GmbH

Sebastian Gerlich
Bundesimmobiliengesellschaft m.b.H

Carsten Röskes
AMPEERS ENERGY GmbH