



Pareto local search for a multi-objective demand response problem in residential areas with heat pumps and electric vehicles

Thomas Dengiz ^{a,*}, Andrea Raith ^b, Max Kleinebrahm ^a, Jonathan Vogl ^a, Wolf Fichtner ^a

^a Karlsruhe Institute of Technology (KIT), Institute for Industrial Production, Karlsruhe, Germany

^b University of Auckland, Department of Engineering Science, Auckland, New Zealand

ARTICLE INFO

Dataset link: [\[GitHub\]](#), [\[KITOpen\]](#)

Keywords:

Multi-objective optimization
Demand response
Pareto local search
Reinforcement learning
Residential area

ABSTRACT

In future energy systems characterized by significant shares of fluctuating renewable energy sources, there is a need for a fundamental change in electricity consumption. The energy system must adapt to the intermittent generation of renewable energy sources. This can be achieved by using flexible electrical loads, such as heat pumps and electric vehicles, with efficient control methods. In this paper, we introduce the Pareto local search method PALSS which employs heuristic search operations to solve the multi-objective demand response problem in residential areas, resulting in superior performance to existing approaches. Flexible electrical loads are shifted with the objective of minimizing the electricity cost and peak load while maintaining the inhabitants' comfort in favorable ranges. Furthermore, we extend PALSS by incorporating reinforcement learning into the search operations in the approach RELAPALSS. For the evaluation, we employ the dichotomous method to obtain solutions that are guaranteed to be Pareto-optimal, serving as benchmarks. The results demonstrate that PALSS outperforms state-of-the-art multi-objective evolutionary algorithms by 16% (18% for RELAPALSS) regarding the performance indicator Generational Distance and by 128% (130% for RELAPALSS) for the indicator Hypervolume. The runtime for PALSS and RELAPALSS is reduced by up to 92% compared to exact methods.

1. Introduction

Future energy systems require flexible electricity loads to cope with the intermittent electricity generation by renewable energy sources. In residential areas, electric heating systems (like heat pumps) coupled to thermal storage, and electric vehicles (EV), are especially suitable for demand response by adjusting their electricity consumption based on an external signal [1,2]. In 2022, space heating accounted for approximately 63.5% of the total energy usage in EU households and about 15% for domestic hot water (DHW) [3]. Due to their high efficiency, heat pumps are especially well-suited for demand response strategies. They can utilize existing infrastructure like the building's thermal mass and hot water tanks to shift electricity usage over time while maintaining occupant comfort. Advanced measurement technologies, like smart meters, combined with intelligent monitoring and control strategies, are transforming the traditional electricity grid into a smart grid. Intelligent strategies for centrally controlling smart buildings can lower electricity costs, enhance the utilization of renewable energy, and contribute to grid stabilization [4].

A price signal tied to electricity market prices is typically used to control electricity consumption [5]. The problem with using a price

or availability signal for renewable energy sources is that all buildings tend to react to the lowest price at the same time [6]. This leads to new peak loads in the electricity grid that might harm the transformers or the transmission lines. This issue becomes increasingly significant with the growing residential demand driven by the electrification of the mobility and heating sectors, particularly through the use of EVs and heat pumps [7].

Thus, the resulting optimization problem for controlling flexible devices in a residential area has multiple conflicting goals. Next to the minimization of the electricity cost, the reduction of the peak load in the local grid is an important objective. Further, considering the inhabitants' thermal comfort is crucial for shifting the operation of the electrical heating devices. Since the multi-objective scheduling problem is NP-hard [8], exact methods for finding optimal solutions are computationally intensive and therefore impractical for real-world scenarios involving residential areas with multiple buildings [9]. Consequently, heuristic methods are needed to solve the multi-objective optimization problem for the intelligent control of flexible devices. While heuristics cannot guarantee to find the globally optimal solution, they have shown

* Corresponding author.

E-mail address: thomas.dengiz@kit.edu (T. Dengiz).

<https://doi.org/10.1016/j.energy.2025.138063>

Received 29 October 2024; Received in revised form 12 August 2025; Accepted 14 August 2025

Available online 21 August 2025

0360-5442/© 2025 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

Nomenclature**Acronyms**

BT	Building type
DHW	Domestic hot water
EV	Electric Vehicle
GD	Generational Distance
HV	Hypervolume
MILP	Mixed Integer Linear Programming
NSGA-III	Non-dominated Sorting Genetic Algorithm 3
NSGA-II	Non-dominated Sorting Genetic Algorithm 2
PALSS	Pareto local search for load shifting
PLS	Pareto local search
RELAPALSS	Reinforcement learning assisted Pareto local search
RL	Reinforcement Learning
RVEA	Reference Vector Guided Evolutionary Algorithm
SPEA-II	Strength Pareto Evolutionary Algorithm

Indices

B	total number of buildings
b	building
i	iteration
t	time slot
Z	total number of time slots

Parameters

η_b	charging efficiency of the EV
ρ^{Concrete}	density of concrete
$\text{COP}_{t,b}$	coefficient of performance of the heat pump
$a_{t,b}$	availability parameter of the EV
c^{Concrete}	specific heat capacity of concrete
C_b^{EV}	energy capacity of the EV's battery
d^{DHW}	factor for controlling the allowed deviation from the initial DHW volume
$d^{\text{Temperature}}$	factor for controlling the allowed deviation from the initial temperature
k_b	maximum number of heat pump starts
P_b^{HP}	maximum electrical power of the heat pump
p_t	time-variable electricity price
$P_{t,b}^{\text{el}}$	electrical power of the inflexible devices
$P_{t,b}^{\text{EVDrive}}$	power consumption of the EV while driving
$P_{t,b}^{\text{EVMax}}$	maximum power of the EV charging station
$Q_{t,b}^{\text{DemandDHW}}$	demand for DHW
$Q_{t,b}^{\text{DemandSH}}$	demand for space heating
$Q_{t,b}^{\text{LossesDHW}}$	energy losses of the hot water tank
$Q_{t,b}^{\text{LossesSH}}$	energy losses of the heating system
T_b^{max}	maximum building temperature
T_b^{min}	minimum building temperature
T^{DHW}	temperature of the domestic hot water

V_b^{DHWmax}	maximum volume of the hot water tank
V_b^{DHWmin}	minimum volume of the hot water tank
V_b^{UFH}	volume of the underfloor heating system
mod^{min}	minimum modulation degree of the heat pump

Variables

$\text{SOC}_{t,b}$	state of charge of the EV
$h_{t,b}^{\text{runDHW}}$	binary variable indicating if the heat pump is running for DHW
$h_{t,b}^{\text{runSH}}$	binary variable indicating if the heat pump is running for space heating
$h_{t,b}^{\text{switchedOff}}$	binary variable indicating if the heat pump is switched off
P^{Peak}	electrical peak load
$P_i^{\text{PeakShift\%}}$	share of shifted load using the Peak-Shift-Operator
$P_i^{\text{PriceShift\%}}$	share of shifted load using the Price-Shift-Operator
$P_{t,b}^{\text{EV}}$	charging power of the EV
$Q_{t,b}^{\text{DHW}}$	heating energy of the heat pump for DHW
$Q_{t,b}^{\text{SH}}$	heating energy of the heat pump for space heating
$T_{t,b}^{\text{Building}}$	temperature of the building
$V_{t,b}^{\text{DHW}}$	usable volume for DHW of the hot water tank
$x_{t,b}$	modulation degree of the heat pump for space heating
$y_{t,b}$	modulation degree of the heat pump for DHW

very good performance with strongly reduced computational times in numerous demand response studies [10].

Compared to single-objective optimization, multi-objective optimization offers several key advantages [11,12]. The outcome of an optimization problem in a multi-dimensional objective space is a Pareto set – a collection of optimal trade-off solutions (Pareto-optimal solutions) – which provides greater decision-making flexibility. In multi-objective optimization, the weights of the objectives do not need to be specified beforehand, allowing the decision-maker to explore trade-offs and understand objective conflicts before expressing preferences. In contrast, the weighted-sum approach used in single-objective optimization requires the decision maker to specify objective weights prior to optimization, which can introduce arbitrary or biased results depending on the chosen weights. Moreover, multi-objective optimization methods offer a more comprehensive understanding of the interactions between conflicting objectives, revealing the sensitivity of these objectives to various decision variables. This level of insight is generally unattainable with single-objective optimization. Additionally, it is known that the weighted-sum approach cannot find all Pareto-optimal solutions, as discussed in Section 4.3, thereby limiting the exploration of the solution space in the objective domain. Exact methods for solving multi-objective optimization problems are generally more computationally demanding than those for single-objective problems, further justifying the use of heuristic approaches in applications with realistic problem sizes.

Heuristic local search methods based on Pareto-optimality have been successfully used in several applications of multi-objective optimization [13,14]. However, the literature lacks multi-objective heuristics that incorporate demand response-specific search operations, particularly for exploiting the flexibility of thermal storage in combination

Table 1

Comparison of relevant papers from the literature.

	RL with domain knowledge	Multi-objective optimization	Pareto local search	Compared to metaheuristics	Compared to exact methods
Pinto et al. 2021 [15]	X	(✓)	X	X	X
Peirelinck et al. 2024 [16]	✓	(✓)	X	X	X
Pinto et al. 2021 [17]	✓	(✓)	X	X	X
Shakouri et al. 2017 [18]	X	(✓)	X	X	✓
Tang et al. 2022 [19]	X	(✓)	X	X	X
Dutra et al. 2019 [20]	X	(✓)	X	X	✓
Khalid et al. 2018 [21]	X	(✓)	(✓)	✓	X
Wynn et al. 2022 [22]	X	✓	X	✓	X
Terlouw et al. 2019 [23]	X	✓	X	X	(✓)
Song et al. 2022 [24]	X	(✓)	X	X	(✓)
Wang et al. 2020 [25]	X	✓	X	✓	X
Kazmi et al. 2019 [26]	✓	X	X	X	X
Wei et al. 2015 [27]	X	✓	X	✓	X
Chiu et al. 2020 [28]	X	✓	✓	✓	X
Present work	✓	✓	✓	✓	✓

with heat pumps. To address this gap, we propose a Pareto-based local search method to solve the multi-objective demand response optimization problem for a residential area equipped with thermal storage and EVs. The goal is to minimize the electricity costs while simultaneously reducing the peak load and maintaining the thermal comfort of the inhabitants in acceptable ranges.

The remainder of the paper is organized as follows: In Section 2, we describe the relevant literature and explain the main contribution of this paper. We define the multi-objective optimization problem in Section 3 and introduce our novel Pareto local search approach and an exact method from the literature for solving the problem in Section 4. The results of our experiments are presented in Section 5. The paper ends with a conclusion and outlook in Section 6.

2. Related work and contribution

2.1. Related work

Table 1 lists the relevant studies from the literature. Pinto et al. [15] use (Reinforcement Learning) RL to optimize four buildings with heat pumps and domestic hot water (DHW). They aim to minimize the energy costs and the peak load. Peirelinck et al. [16] include domain knowledge into their RL application to control the DHW heating of buildings to minimize the electricity cost by having a tariff that depends on the peak load. Pinto et al. [17] combine a Long-Short-Term-Memory (LSTM) with RL to control heat pumps. They use the LSTM to learn the building dynamics for evaluation from a building simulation tool. Shakouri et al. [18] define a multi-objective optimization problem for minimizing the costs and peak load in an intelligent grid by shifting the electricity consumption of household devices. They solve the problem with Goal programming. In [19], Tang et al. use RL to operate an integrated energy system with electricity, gas, and heating networks. Their objectives also include the minimization of costs and peak load. Dutra et al. [20] introduce a decentralized heuristic for demand response of an aggregator. The aggregator controls multiple flexible loads such that the costs and the peak load are reduced. Khalid et al. 2018 [21] define a novel local search method based on two metaheuristics (genetic algorithm and hybrid bacterial foraging). Their flexible devices include air conditioners (AC), water heaters, washing machines, and dishwashers.

Wynn et al. [22] use a multi-objective version of the metaheuristic gray wolf optimization applied for scheduling flexible resources in a microgrid. Next to minimizing the costs and the peak load, they consider consumers' dissatisfaction in their optimization problem. Terlouw et al. [23] define a multi-objective optimization problem for an aggregator with different battery technologies for a whole community. The objectives comprise minimizing costs, CO₂-emissions, and the peak

load. In the study [24], Song et al. solve a day-ahead battery scheduling problem for residential buildings and include the goal of costs, peak load, and consumer satisfaction in the objective function. Wang et al. [25] use the multi-objective genetic algorithm NSGA-II for the optimization of a combined cooling, heating, and power system. The goals are the reduction of CO₂-emissions, costs, energy usage, and the enhancement of grid integration. Kazmi et al. [26] train an RL agent for thermostatically controlled loads (heat pump, DHW). They include domain knowledge into RL by adding additional training features derived from thermodynamic laws. The goal is to minimize energy consumption in residential areas. Wei et al. [27] use multi-objective particle swarm optimization to optimize the operation of a heating-ventilation and air conditioning (HVAC) system. The goals include minimizing the energy consumption, CO₂-emissions and the temperature of buildings while maximizing the indoor air quality. Chiu et al. [28] introduce a Pareto local search approach based on evolutionary algorithms for demand response in a residential area. The Pareto local search controls the AC, EV, and shiftable devices to minimize the energy costs and the peak load.

While every study considers multiple objectives, most use the weighted sum approach to transform the multi-dimensional objective space into a one-dimensional space (ticks in brackets). Only [22,23,25,27,28] optimize in a multi-objective space and generate trade-off (Pareto-optimal) solutions. However, none of those studies compares their introduced approach to an exact method for multi-objective optimization that systematically generates the true Pareto-front containing all Pareto-optimal solutions. Chiu et al. [28] introduce a Pareto local search approach aimed at approximating the true Pareto front using local search methods. However, their approach relies on the local search operations of evolutionary algorithms, which do not incorporate problem-specific heuristics tailored to the demand response problem with variable electricity tariffs. Moreover, their method does not utilize thermal storage to exploit the flexibility of heat pumps. While several studies use RL [15–17,19,26], only [16,17,26] include domain knowledge into the decision making of their trained agent.

2.2. Contribution

To the best of our knowledge, this study is the first to introduce a Pareto local search method with problem-specific operations tailored to shift flexible electricity loads from heat pumps and EVs in residential areas, within a multi-dimensional objective space. Approaches for multi-objective optimization problems aim to find optimal trade-off solutions for decision-making. This feature of heuristic control approaches for demand response in residential areas has not been adequately explored in the literature. We also extend our approach by integrating RL into the decision-making and evaluating the novel

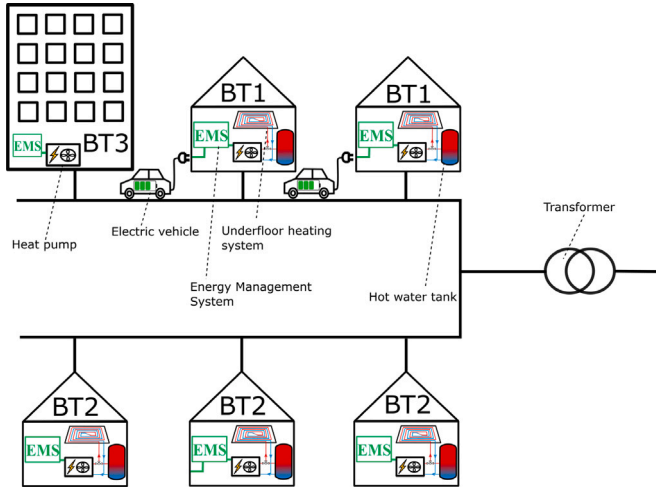


Fig. 1. Residential area with three different building types.

approaches. Our approach is the only one to utilize RL combined with domain knowledge of demand response operations to optimize within a multi-dimensional objective space. Moreover, we have made the code and data used in this study publicly available to facilitate reproducibility and support future research.

In summary, our paper has the following three unique contributions:

- We introduce a novel approach for Pareto local search for demand response in residential areas with heat pumps, thermal storage and EVs that utilizes heuristic search operations. The heuristic search operations we propose for this domain are new and designed specifically for multi-objective demand response problems.
- We include a model-free RL approach into the Pareto local search's decision-making in a multi-dimensional objective space and thus include domain knowledge into the training of the RL agent. To the best of our knowledge, this is the first approach that combines RL with domain-specific knowledge for multi-objective demand response optimization.
- For the evaluation, we solve the multi-objective optimization problem for a residential area using four state-of-the-art meta-heuristics and the *Dichotomous method*, an exact method that can guarantee to find the whole Pareto-front. To our knowledge, no prior work on multi-objective demand response has included exact methods for benchmarking. Further, we compare them to an existing conventional control approach that is used most often nowadays in buildings. Notably, our evaluation is extensive, comparing our methods to a greater number and a wider variety of algorithms than the studies found in the literature.

3. Multi-objective demand response optimization problem for the residential area

In this section, we formulate the multi-objective demand response optimization problem for the residential area, which consists of three building types (BT) visualized in Fig. 1. BT1 is a single-family building that uses a heat pump for space heating and domestic hot water. An underfloor heating system and a hot water tank serve as thermal storage. Further, the inhabitants charge the EV at home. BT2 is similar to BT1, but does not use an EV. BT3 is a multi-family building that uses a heat pump for space heating but not for DHW supply (DHW requires higher temperatures which reduces the efficiency of heat pumps).

The primary goals for a centralized controller are to minimize the total electricity costs with a time-variable electricity tariff p_t and to

minimize the peak load P^{Peak} in the residential area.

$$\min(\text{Cost}, P^{Peak}) \quad (1)$$

subject to:

$$T_b^{\min} \leq T_{t,b}^{\text{Building}} \leq T_b^{\max} \quad \forall t, b \quad (2)$$

$$V_b^{\text{DHWmin}} \leq V_{t,b}^{\text{DHW}} \leq V_b^{\text{DHWmax}} \quad \forall t, b \quad (3)$$

$$T_{Z,b}^{\text{Building}} \geq T_{1,b}^{\text{Building}} \cdot d^{\text{Temperature}} \quad \forall b \quad (4)$$

$$V_{Z,b}^{\text{DHW}} \geq V_{1,b}^{\text{DHW}} \cdot d^{\text{DHW}} \quad \forall b \quad (5)$$

$$T_{t,b}^{\text{Building}} = T_{t-1,b}^{\text{Building}} + \frac{Q_{t,b}^{\text{SH}} - Q_{t,b}^{\text{DemandSH}} - Q_{t,b}^{\text{LossesSH}}}{V_b^{\text{UFH}} \cdot \rho^{\text{Concrete}} \cdot c^{\text{Concrete}}} \quad \forall t \neq 1, \forall b \quad (6)$$

$$V_{t,b}^{\text{DHW}} = V_{t-1,b}^{\text{DHW}} + \frac{Q_{t,b}^{\text{DHW}} - Q_{t,b}^{\text{DemandDHW}} - Q_{t,b}^{\text{LossesDHW}}}{T^{\text{DHW}} \cdot \rho^{\text{Water}} \cdot c^{\text{Water}}} \quad \forall t \neq 1, \forall b \quad (7)$$

$$Q_{t,b}^{\text{SH}} = x_{t,b} \cdot P_b^{\text{HP}} \cdot \text{COP}_{t,b} \cdot \Delta t \quad \forall t, b \quad (8)$$

$$Q_{t,b}^{\text{DHW}} = y_{t,b} \cdot P_b^{\text{HP}} \cdot \text{COP}_{t,b} \cdot \Delta t \quad \forall t, b \quad (9)$$

$$h_{t-1,b}^{\text{runSH}} + h_{t-1,b}^{\text{runDHW}} \leq h_{t,b}^{\text{runSH}} + h_{t,b}^{\text{runDHW}} + h_{t,b}^{\text{switchedOff}} \quad \forall t \neq 1, \forall b \quad (10)$$

$$h_{t,b}^{\text{runSH}} + h_{t,b}^{\text{runDHW}} \leq 1 \quad \forall t, b \quad (11)$$

$$\sum_{t=1}^Z h_{t,b}^{\text{switchedOff}} \leq k_b \quad \forall b \quad (12)$$

$$x_{t,b} \leq h_{t,b}^{\text{runSH}} \quad \forall t, b \quad (13)$$

$$x_{t,b} \geq h_{t,b}^{\text{runSH}} \cdot \text{mod}^{\min} \quad \forall t, b \quad (14)$$

$$y_{t,b} \leq h_{t,b}^{\text{runDHW}} \quad \forall t, b \quad (15)$$

$$y_{t,b} \geq h_{t,b}^{\text{runDHW}} \cdot \text{mod}^{\min} \quad \forall t, b \quad (16)$$

$$P_{t,b}^{\text{EV}} \leq P_{t,b}^{\text{EVMax}} \cdot a_{t,b} \quad \forall t, b \quad (17)$$

$$0 \leq \text{SOC}_{t,b} \leq 1 \quad \forall t, b \quad (18)$$

$$\text{SOC}_{Z,b} \geq \text{SOC}_{1,b} \cdot d^{\text{SOC}} \quad \forall b \quad (19)$$

$$\text{SOC}_{t,b} = \text{SOC}_{t-1,b} + \frac{P_{t,b}^{\text{EV}} \cdot \eta_b \cdot \Delta t - P_{t,b}^{\text{EVDrive}} \cdot \Delta t}{C_b^{\text{EV}}}, \quad \forall t \neq 1, \forall b \quad (20)$$

$$\text{Cost} = \sum_{t=1}^Z \sum_{b=1}^B ((x_{t,b} + y_{t,b}) \cdot P_b^{\text{HP}} + P_{t,b}^{\text{EV}} + P_{t,b}^{\text{el}}) \cdot \Delta t \cdot p_t \quad (21)$$

$$P^{\text{Peak}} = \max_t \left\{ \sum_{b=1}^B ((x_{t,b} + y_{t,b}) \cdot P_b^{\text{HP}} + P_{t,b}^{\text{EV}} + P_{t,b}^{\text{el}}) \right\} \quad (22)$$

$$x_{t,b} \in [0, 1], \quad y_{t,b} \in [0, 1], \quad h_{t,b}^{\text{runSH}} \in [0, 1], \quad h_{t,b}^{\text{runDHW}} \in [0, 1], \quad h_{t,b}^{\text{switchedOff}} \in [0, 1] \quad (23)$$

Eq. (2) makes sure that the temperature $T_{t,b}^{\text{Building}}$ for every building b is always between a lower temperature limit T_b^{\min} and an upper temperature limit T_b^{\max} during every time slot t . In analogy, Eq. (3) forces the volume of the hot water tank $V_{t,b}^{\text{DHW}}$ to be between the corresponding minimum V_b^{DHWmin} and maximum volume V_b^{DHWmax} . To ensure that the values of the building temperature and the usable volume of the hot water tank are not too far away from their initial values, we add Eq. (4) and (5). These equations include the two parameters $d^{\text{Temperature}}$ and d^{DHW} whose values are between 0 and 1. Thus, we can control the deviation from the allowed initial values.

We define the temperature difference equation in Eq. (6). The temperature of each building b at time t is equal to its temperature from the previous time slot plus the energetic difference divided by the volume of the underfloor heating system V^{UFH} , the density ρ^{Concrete} and the specific heat capacity of concrete c^{Concrete} . The heat pump's heating energy for space heating $Q_{t,b}^{\text{SH}}$ increases the temperature of

the building while the demand for space heating $Q_{t,b}^{\text{DemandSH}}$ and the losses of the heating system itself $Q_{t,b}^{\text{LossesSH}}$ decrease it. The demand for space heating comprises a time series incorporating the building's transmission and ventilation losses, as well as internal and solar gains. We describe in Section 5.1 how we generated the time series. In analogy, Eq. (7) defines the energetic difference equation for the volume of the hot water tank. The heating energy flows for DHW from the heat pump $Q_{t,b}^{\text{DHW}}$ increase the usable volume of DHW while the demand for DHW $Q_{t,b}^{\text{DemandDHW}}$ and the hot water tank's standing losses $Q_{t,b}^{\text{LossesDHW}}$ decrease it. As hot water tanks use water as storage medium, the energetic difference is divided by the temperature T^{DHW} , density ρ^{Water} and specific heat capacity c^{Water} of water.

To calculate the heat pump's heating energy for space heating, Eq. (8) multiplies the modulation degree of the heat pump for space heating $x_{t,b}$ with the heat pump's maximum electrical power P_b^{HP} , the coefficient of performance $\text{COP}_{t,b}$ and the time resolution Δt . Likewise, Eq. (9) uses the heat pump's modulation degree for heating the hot water tank $y_{t,b}$ to derive the heating flows to the hot water tank.

Eq. (10) introduces the binary variables $h_{t,b}^{\text{unSH}}$ and $h_{t,b}^{\text{unDHW}}$ that indicate whether the heat pump is running for space heating or DHW at time t . Further, the binary variable $h_{t,b}^{\text{switchedOff}}$ has the value 1 if the heat pump is being switched off at time slot t . As the heat pump cannot heat the building (space heating) and the hot water tank simultaneously, we add Eq. (11). To ensure that the maximum number of heat pump starts does not exceed k_b , Eq. (12) is added. Eq. (13) to Eq. (16) force the heat pump's modulation degree to be either 0 or higher than the minimum modulation degree mod^{min} .

The charging power for the electric vehicle $P_{t,b}^{\text{EV}}$ should not exceed the maximum power of the charging station $P_{t,b}^{\text{EVMax}}$. Moreover, the EV can only be charged at a particular time slot t if its binary availability variable $a_{t,b}$ indicates that the EV is at home and thus available for charging. Hence, we introduce Eq. (17) to incorporate these constraints. Eq. (18) ensures that the state of charge of the electric vehicle's battery $\text{SOC}_{t,b}$ is always between 0 and 1, while Eq. (19) hinders it to be too far away from its initial value at the end of the optimization horizon Z . The difference equation for the $\text{SOC}_{t,b}$ is defined in Eq. (20). The charging power of the EV, adjusted by the charging efficiency η_b , increases the SOC while the required electrical power while driving the EV $P_{t,b}^{\text{EVDrive}}$ reduces it. The parameter C_b^{EV} quantifies the energetic capacity of the battery.

Eq. (21) defines the costs and, thus, the first objective. The total costs are calculated by summing the electrical power consumption of the heat pump, the EV, and the inflexible household appliances $P_{t,b}^{\text{el}}$. This is multiplied by the time resolution and the current price p_t of a time-variable electricity tariff. For the costs, we sum up over all time slots and buildings.

To determine the peak load as the second objective, we use Eq. (22), which takes the maximum value of the combined electrical load of all buildings over all time slots. To model the maximum function, we introduce a set of linear constraints requiring P^{Peak} to be greater than or equal to the load at each time slot in the implementation. Since P^{Peak} is minimized in the objective, it will exactly equal the maximum load across all time slots. Thus, despite involving a maximum function, the problem formulation remains linear.

While the defined optimization problem is valid for the whole residential area, it must be noted that BT2 and BT3 do not have an EV, and BT3 only uses the heat pump for space heating. For these buildings, the corresponding equations are neglected. The optimization problem is a mixed-integer linear problem (MILP).

4. Methods for solving the multi-objective optimization problem

In this section, we introduce our novel algorithms for shifting electrical loads in a residential area and thereby find Pareto-optimal solutions for the minimization of the peak load and the electricity costs.

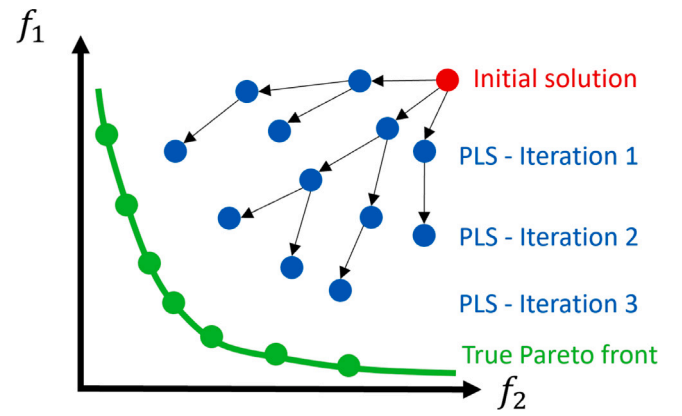


Fig. 2. Solution space of a minimization problem with Pareto local search.

A solution is Pareto-optimal when there exists no other solution that improves at least one objective without worsening the others.

In Section 4.1, we explain the Pareto local search approach for demand side management *PALSS*, which constitutes the main contribution of this paper. Next, we describe *RELAPALSS* in Section 4.2, which uses reinforcement learning to assist *PALSS*. In Section 4.3, we explain the *Dichotomous method*, which is an exact method for multi-objective optimization that we use to obtain a real Pareto-front of the problem that consists of so-called supported Pareto-optimal solutions. We use a conventional control approach for a lower benchmark in our evaluations that we describe in Section 4.4.

4.1. Pareto local search for load shifting (*PALSS*)

4.1.1. Pareto local search

Pareto local search (PLS) aims to find a good approximation of the true Pareto-front (that is not known before) in a reasonable time. PLS extends the iterative improvement algorithm, called local search, from single-objective to multi-objective problems for the systematic exploration of the solution space [14]. It navigates the neighborhood of solutions within its archive, incorporating new non-dominated solutions while removing dominated ones to refine the archive. Fig. 2 illustrates the outcomes in the solution space of three PLS iterations for the minimization of two objectives. In many applications, the search starts from a naive solution of the problem that is easy to find.

Fig. 3 depicts the schematic workflow of our developed method *PALSS*. The search starts with a conventional solution (see Section 4.4).

PALSS generates the population of the first generation by using local search operators (see Section 4.1.2) to generate k solutions. Afterward, the same local search operations generate m new solutions from each of the k existing solutions and store each of the $k \cdot m$ solutions in a temporary archive. Next, *PALSS* uses the Pareto sorting approach (see Section 4.1.3) to filter the solutions in the temporary archive and reduce them again to k solutions. These solutions are Pareto-optimal (based on the solutions in the archive) and form the population of the next iteration. This procedure is repeated for n iterations (in case the algorithm meets no other stopping criterion, such as a time limit or predefined target values of the different goals), and the algorithm outputs the approximated Pareto front.

4.1.2. Local search operators

PALSS makes use of two shifting operations to adjust the load profiles of individual buildings: the *Price-Shift-Operator* and the *Peak-Shift-Operator*. Both use as input the resulting electrical load profile of the residential area from a solution and the time-variable electricity tariff. The *Price-Shift-Operator* shifts loads from time slots with high electricity prices to time slots with low prices. Fig. 4 illustrates this

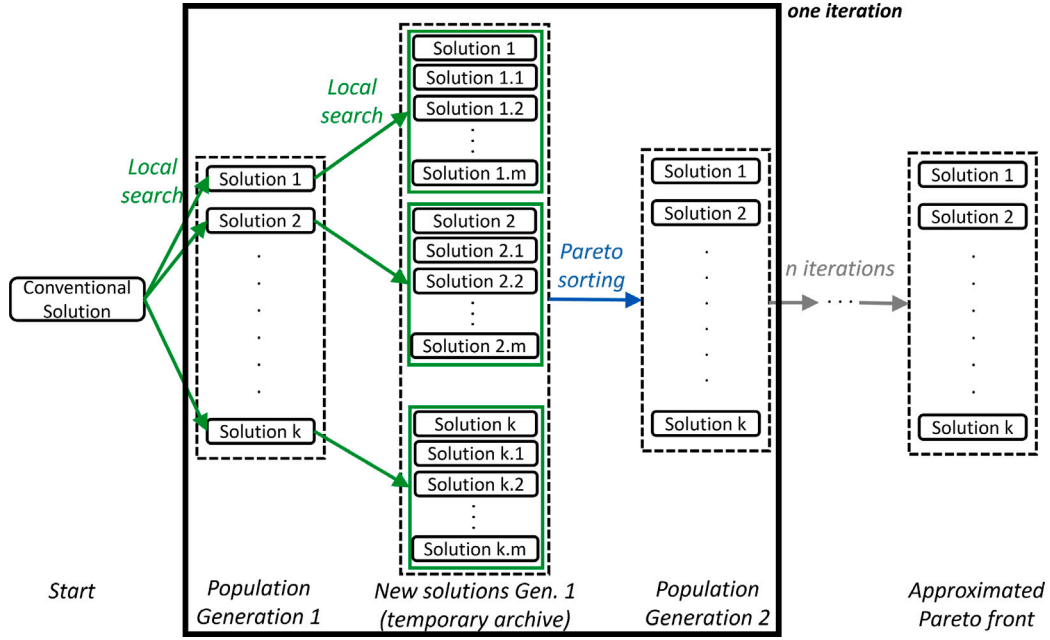


Fig. 3. Schema of PALSS.

Table 2

Probabilities of *PALSS* for selecting the highest price hours in each iteration using the *Price-Shift-Operator* (Probabilities for the lowest price hours are equivalent to those for the highest price hours).

Iteration	Highest hour 1	Highest hour 2	Highest hour 3	Highest hour 4	Highest hour 5
1	0.410	0.328	0.123	0.082	0.057
2	0.393	0.311	0.139	0.098	0.059
3	0.377	0.295	0.156	0.11	0.062
4	0.361	0.279	0.172	0.12	0.068
5	0.344	0.262	0.189	0.135	0.07

operation. Flexible electrical devices (heat pumps, EVs) try to reduce their load during time slots with high prices if possible. In compensation for that, the devices try to increase their electricity consumption during low-price periods.

For the selection of the affected time slots and the amount of shifted loads, we use a heuristic approach. Table 2 shows the probabilities for selecting the hours with the highest, second highest, third highest prices, etc., during each iteration. The probability of selecting the time slot with the highest and second highest prices is very high initially and slightly decreases in every new iteration. In contrast, the probability of selecting the hour with the third, fourth, and fifth highest price to reduce the load increases in each iteration. The probabilities for selecting the time slot with the lowest value to shift the load to are determined analogously. Using selection probabilities enhances the exploration of the search space. We determined these probabilities after running a large number of experiments. The concrete probabilities depend on the specific use case and scenario. However, the general idea of having high probabilities at the beginning for the highest prices which decrease over time, and analogously, having lower probabilities for the lower prices that slightly increase in every iteration, is useful for all probabilistic demand response heuristics.

Next to selecting the time slots for load shifting, the *Price-Shift-Operator* determines the amount of load that is shifted between the time slots. For this purpose, it uses the following uniform distribution in iteration i :

$$\Delta P_i^{\text{PriceShift}\%} \sim U(20 - i, 40 - 2i), \quad i \in \{1, \dots, 5\} \quad (24)$$



Fig. 4. Price-Shift-Operator.

According to Eq. (24), the percentage of shifted loads $p_i^{\text{load}\%}$ in iteration i is randomly determined by using a uniform distribution with diminishing upper and lower boundaries in each iteration. Those values yielded the best results in our experiments. These percentages and the maximum number of iterations (which is five in our case) also depend on the specific use case and scenario. The central idea is that in early iterations, higher amounts of loads are shifted while the amount decreases with increasing iteration counter. Using random values ensures diversity in the temporary archive.

The *Peak-Shift-Operator* is similar to the *Price-Shift-Operator* as Fig. 5 illustrates. It always chooses the time slot with the highest load in the whole residential area and shifts a certain amount to time slots with low prices. The probabilities for selecting the time slots with the lowest prices are equivalent to the ones of the *Price-Shift-Operator* in Table 2. The algorithm likewise randomly determines the amount of shifted load

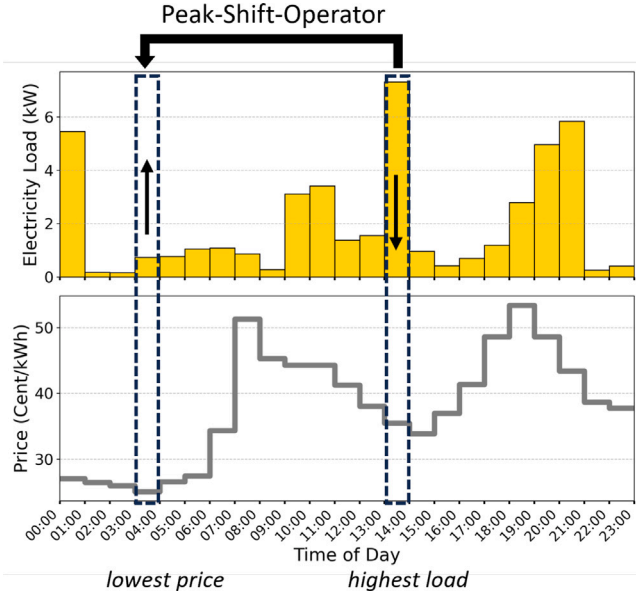


Fig. 5. Peak-Shift-Operator.

$P_i^{\text{PeakShift\%}}$ by using a uniform distribution with decreasing boundaries in each iteration i as shown in Eq. (25).

$$\Delta P_i^{\text{PeakShift\%}} \sim U(10 - i, 25 - i), \quad i \in \{1, \dots, 5\} \quad (25)$$

The boundaries can be flexible. We determined the specific values by running multiple experiments and choosing the combination that yielded the best results. PALSS randomly chooses to either apply the *Peak-Shift-Operator* or the *Price-Shift-Operator* to generate one new candidate solution. This is done several times for every solution in the generation to generate a temporary archive with new candidate solutions.

4.1.3. Pareto sorting

The number of the new candidate solutions in the temporary archive is higher than the population size. Thus, the *Pareto sorting* eliminates several solutions from it. First, the algorithm determines the Pareto-optimal solutions amongst the candidate solutions of the prior generation. Further, all solutions that violate any constraints are eliminated. A score is calculated by using Eq. (26) for all remaining solutions.

$$\text{Score} = \frac{\text{Cost}}{\text{Cost}_{\text{Conventional}}} + \frac{P^{\text{Peak}}}{P^{\text{Peak}}_{\text{Conventional}}} \quad (26)$$

The equation divides the cost of a solution by the cost of the conventional solution. Likewise, the peak power is divided by the peak of the conventional solution. A lower score indicates a better fulfillment of the two objectives as we deal with a minimization problem.

If the number of Pareto-optimal solutions in the temporary archive of generation i is lower than the population size, all Pareto-optimal solutions are transferred to the next generation $i + 1$. The remaining positions in generation $i + 1$ are filled by successively adding the dominated solutions of the previous generation's i temporal archive with the lowest Score value. If the number of Pareto-optimal solutions exceeds the population size, the algorithm only adds the solutions with the lowest (best) Score values to the new generation.

4.2. Reinforcement learning assisted pareto local search (RELAPALSS)

Instead of heuristically choosing the time slots and the amount of shifted electrical loads within the *Peak-Shift-Operator* and the *Price-Shift-Operator*, we also investigate the use of model-free reinforcement

learning algorithms to learn the parameters for the local search operators. We use the *Proximal Policy Optimization Algorithm* [29] as it yielded promising results in many different applications and tends to be sample-efficient. We train it on 12 days for one configuration of the residential area with 20 buildings.

To derive the state for the *Price-Shift-Operator*, the RELAPALSS algorithm uses the five time slots with the highest and lowest electricity prices. It then calculates the share of electrical load from the overall load profile of the residential area (as percentages) for each of those time slots. This results in a ten dimensional state space. This state vector provides the agent with information about the load distribution during the most expensive and cheapest time slots, enabling it to learn shifting strategies based on price signals. For the *Peak-Shift-Operator*, the calculation is similar. The difference is that only the five time slots with the lowest electricity prices are considered, making the state space five-dimensional. This operator focuses solely on reducing the peak load by shifting demand from the peak period to the five lowest-price time slots.

The action space for the *Price-Shift-Operator* is a three-dimensional vector. The first component determines the source time slot for shifting the electricity load. The agent chooses the source time slot out of the five time slots with the highest prices indicated by the state space vector. The second component contains the target time slot to shift the load to (time slots with lowest prices), while the third quantifies the amount of shifted electrical loads.

The action space is merely two-dimensional for the *Peak-Shift-Operator*. The first dimension determines the target time slot to shift the load to, and the second component decides about the amount of shifted electrical loads.

During the training phase the RELAPALSS algorithm calculates the reward after every application of the local search operators for every building. If the *Price-Shift-Operator* is chosen as the local search operator, the reward quantifies the percentage cost reduction of the individual building multiplied by a constant weighting factor. For the *Peak-Shift-Operator*, the reduced peak load of the whole residential area is used as a reward for the *Proximal Policy Optimization Algorithm*. Throughout training, the agent interacts with the simulation environment by selecting and applying both types of local search operators. Over time, it learns to choose operator parameters that are likely to improve the overall objective.

4.3. Dichotomous method

One approach to solving multi-objective optimization problems is scalarization. Here, a weighted sum scalarization is applied to the multi-objective problem with objective functions in Eq. (1) to obtain a *weighted sum objective*

$$\min \lambda \text{Cost} + (1 - \lambda) P^{\text{Peak}}, \quad (27)$$

where $\lambda \in (0, 1)$. The multi-objective load shift problem with objective Eq. (27) is now a single-objective MILP. For every choice of $\lambda \in (0, 1)$ it is guaranteed that a *supported* Pareto-optimal solution is obtained (and all such solutions can be found) [30]. The so-called dichotomous approach is an effective strategy for exploring all necessary weights in order to guarantee that all supported Pareto efficient solutions are found [e.g. 31]. Initially the two lexicographically optimal solutions of the problem are computed which prioritize minimizing a single objective (Cost or P^{Peak}) and among all optimal solutions of this first problem then minimize the other objective, to obtain the initial two points P_1, P_2 on the Pareto-front in Fig. 6. The dichotomous method iterates by updating the weight λ and solving a weighted sum problem for each pair of neighboring points until no new points are found — intuitively the next point is the one furthest from the line connecting the two neighboring points in Fig. 6. The process terminates with a so-called complete set of extreme supported Pareto-optimal solutions. It should be noted that, unlike a multi-objective continuous linear

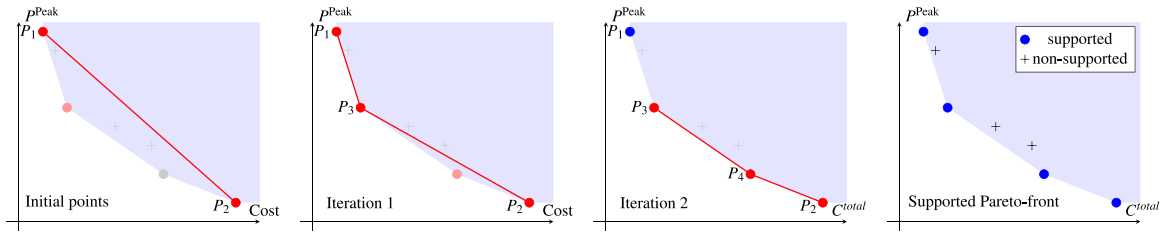


Fig. 6. Illustration of dichotomous method and (non-)supported Pareto-optimal points.

programme, a multi-objective MILP also has other *non-supported* Pareto-optimal solutions (lying in the interior of the convex hull of supported Pareto-optimal solutions, see bottom-right of Fig. 6) that cannot be obtained by the dichotomous method. Hence, the dichotomous method obtains a subset of the true Pareto-front. All obtained solutions are guaranteed to be Pareto-optimal, even if the dichotomous method is terminated early due to an imposed time limit.

4.4. Conventional control

In our application, the conventional control uses a hysteresis-based two-point control approach for the heat pumps for DHW similar to the one described in [32]. The heat pump starts heating the DHW tank with full power when its usable hot water volume reaches a lower limit. After the volume reaches an upper limit, the heating stops.

For space heating, the heat pump's modulation degree is continuously adjusted to have a constant temperature if possible. Otherwise, the conventional control also uses a small hysteresis to control the deviations from the desired temperature. This control approach is typically used for inverter heat pumps, which can not only be switched on and off but are also capable of modulating their power output [33].

Further, the conventional control approach uses uncontrolled charging [34] for the EV and thus charges the EV with full power if it is connected to the charging station until the battery is fully charged.

5. Results

5.1. Scenarios for the analysis

Table 3 lists the relevant parameters for the residential area which serves as the test environment, along with the corresponding sources. We investigate three configurations of the residential area: one with ten buildings (4 BT1, 4 BT2, 2 BT3), one with 20 buildings (8 BT1, 8 BT2, 4 BT3), and one with 30 buildings (12 BT1, 12 BT2, 6 BT3). For the electricity price, we use the hourly data of the Day-Ahead market in Germany from 2021, which we took from the *ENTSO-E Transparency Platform* [35]. We scaled the prices up so that they are in line with the average electricity price for residential customers in Germany. We use the software tool *DHWcalc* [36] to generate the DHW demand profiles. The time resolution Δt for the optimization and the simulations is 30 min and the optimization horizon is one day.

We further use the CREST model [47] in combination with German time use data [48] to simulate occupancy behavior which builds the basis for the consistent simulation of electricity and space heating demand. By considering local weather conditions [49], occupancy profiles, and information on household device equipment [50], electrical appliance demand profiles are generated. The calculation of the space heating demand is based on a 5R1C model following DIN EN ISO 13790, incorporating internal gains derived from the household simulation and thermal building parameters taken from the TABULA residential building typology [51].

The buildings are all located in the federal state of *Schleswig-Holstein* in the north of Germany. We use data collected in the project *iZEUS* [52] for the EV driving patterns and their availability.

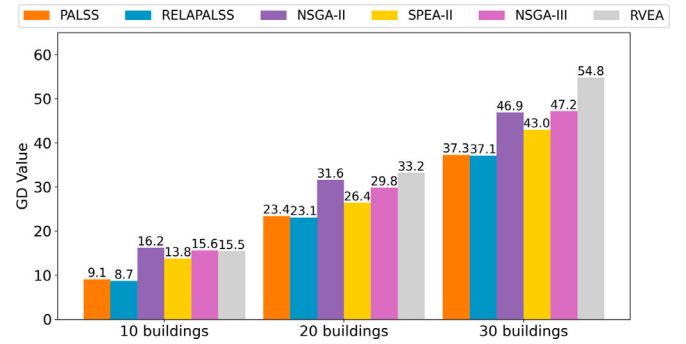


Fig. 7. Results for the Generational Distance (GD) averaged over 18 days and 5 runs.

5.2. Evaluation

We compare our developed algorithms *PALSS* and *RELAPALSS* to the well-established multi-objective evolutionary algorithms *Non-dominated Sorting Genetic Algorithm II (NSGA-II)* [53] and *Strength Pareto Evolutionary Algorithm (SPEA-II)* [54]. Additionally, we include newer evolutionary algorithms into our analysis like *Non-dominated Sorting Genetic Algorithm III (NSGA-III)* [55] and *Reference Vector Guided Evolutionary Algorithm (RVEA)* [56]. We restrict the runtime for each algorithm to solve the optimization problem for one day to 10 min. For the evaluation, we randomly chose 18 days from the heating period in Germany (October to March).

Fig. 7 shows the results of our analysis for the Generational Distance (GD) averaged over 18 days and 5 runs per method. Thus, each bar represents the average value of 90 runs. GD evaluates the quality of a set of solutions to approximate the true Pareto-front. It quantifies the average distance between the solutions and the points on the true Pareto-front [57]. A smaller GD value indicates a better approximation. For all three configurations of the residential area, *PALSS* and *RELAPALSS* yield significantly better results compared to the four metaheuristics *NSGA-II*, *NSGA-III*, *RVEA* and *SPEA-II*. Using RL for selecting the relevant parameters of *PALSS* leads to a small additional improvement as in all categories *RELAPALSS* leads to lower GD values on average. *SPEA-II* yields better approximations compared to the other evolutionary algorithms. *PALSS* leads to improvements of about 16.1% compared to *SPEA-II* while *RELAPALSS* is 17.2% better. Generally, the GD values increase with an increasing number of residential buildings. This is because the problem complexity strongly increases with a higher number of buildings, and we execute every algorithm for only ten minutes independently of the problem size. Thus, only a smaller fraction of the search space can be explored using our developed approaches and the metaheuristics.

All methods lead to crucial improvements compared to the conventional control approach, which, on average, achieved GD values of 22.0 for 10 buildings, 40.4 for 20 buildings, and 60.0 for 30 buildings.

To evaluate the coverage of the objective space by the approximated Pareto-front, we use the hypervolume (HV) indicator [58]. HV quantifies the hypervolume between an approximated Pareto-front and a

Table 3
Parameters of the residential area.

Parameter	Value	Source	Comment
Heated area of the buildings BT1, BT2	140 m ²	[37]	Average German single family house
Heated area of the building BT3	12 · 75 m ²	[38]	Residential area example (Germany)
Concrete width (underfloor heating system)	7 cm	[39]	DIN 18560 (screeds in buildings)
Density of concrete	2400 $\frac{\text{kg}}{\text{m}^3}$	[9]	European standards for concrete EN 206-1
Heat capacity of concrete	1000 $\frac{\text{kJ}}{\text{kg} \cdot \text{K}}$	[9]	European standards for concrete EN 206-1
Supply temperature of the heating system	30 °C	[9]	
Temperature range for space heating	21–23 °C		Assumption for optimal comfort
DHW tank volume	200 l	[40]	For 4 inhabitants
Losses of space heating	45 W, 220 W		Assumed 45 W (BT1, BT2), 220 W (BT3)
Losses of DHW tank	35 W	[41]	2nd highest EU efficiency class (814/2013)
Electrical power of the heat pump BT1, BT2	3 kW	[42]	Sized to outside temperature in Germany
Electrical power of the heat pump BT3	15 kW	[42]	Sized to outside temperature in Germany
Minimal modulation of the heat pump	20%	[43]	Model: Bosch Compress 6800i AW
COP of air-source heat pump for $\Delta T = 28$ °C	4.8	[43]	Model: Bosch Compress 6800i AW
COP of air-source heat pump for $\Delta T = 33$ °C	3.9	[43]	Model: Bosch Compress 6800i AW
Battery capacity of the EV	60 kWh	[44]	Model: Opel Ampera-e
Charging efficiency of the EV	89%	[9]	Model: Opel Ampera-e
Maximal charging power for home charging	4.6 kW	[45]	Wallbox: KEBA KeContact P30
Average length of rides for the EV	45 km	[46]	Inspired by the German Mobility Study

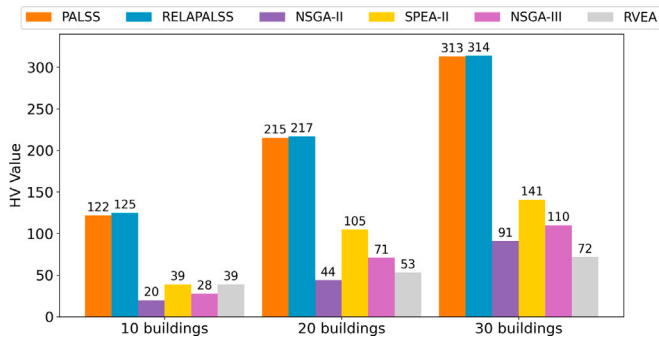


Fig. 8. Results for the Hypervolume (HV) averaged over 18 days and 5 runs.

reference point. Higher values indicate a better coverage. We choose the conventional control's resulting costs and peak load as the reference point in the two-dimensional objective space. Fig. 8 depicts the results for HV of our experiments averaged over 18 days and 5 runs. Analogous to GD, the results for all configurations of the residential area reveal that our developed methods *PALSS* and *RELAPALSS* strongly outperform the four metaheuristics *NSGA-II*, *NSGA-III*, *RVEA* and *SPEA-II*. On average, *PALSS* achieves a 128.1% improvement over the best-performing metaheuristic, *SPEA-II*, while *RELAPALSS* surpasses *SPEA-II* by 130.2%. In this case, the HV values increase with increasing number of buildings. A larger problem size increases the improvements over the naively generated conventional solution. Thus, based on the reference point, the approximated Pareto-fronts enclose a larger volume. However, the distance to the real Pareto-front becomes larger with higher problem complexity. The average HV values for the dichotomous method are 303.9 for 10 buildings, 933.7 for 20 buildings, and 1687.2 for 30 buildings.

Fig. 9 illustrates the found Pareto-fronts of all six methods for different days and configurations of the residential area. It can be seen that the distance between the approximated Pareto-fronts of the various methods and the true Pareto-front, obtained by using the dichotomous method, becomes more significant with increasing problem size. As the wholesale prices on the electricity market can also be negative, meaning that you get money when consuming electricity, the resulting costs of day 11 are also negative. The Pareto-fronts for the different days also reveal that cost minimization and peak load minimization are conflicting objectives since minimizing one goal without worsening the other is impossible.

We tested different parameters for each method and used the ones that yielded the best results. For all six methods *PALSS*, *RELAPALSS*,

NSGA-II, *NSGA-III*, *RVEA* and *SPEA-II*, the population size was 20. The number of new solutions for each member in the population (offsprings) was 3 for the methods *PALSS* and *RELAPALSS* while for the two metaheuristics, the number was 10. All four metaheuristics used the Simulated Binary Crossover (SBX) and started from the solution of the conventional control approach. While the runtime for these four methods was chosen to be 10 min, the execution time of the dichotomous method depends on the problem size. For a residential area with 10 buildings, the dichotomous method required an average of 42 min per day; for 20 buildings, 94 min; and for a problem size of 30 buildings, 122 min. Thus the runtime was up to 92% higher compared to ones of the heuristics with the chosen configuration. We use the Gurobi solver to solve the resulting single-objective optimization problems of the dichotomous method [59]. All experiments were carried out on an Intel Xeon E-1650 v2, 3.5 GHz, 6 Cores. We set the optimality gap for exactly solving the resulting single-objective optimization problems to 0.1% and the maximum runtime for a single problem to 10 min.

On average, the training time for *RELAPALSS* was 4 h, and we trained it on 12 different days, using the conventional solution as the starting solution. It should be noted that we trained the RL model solely on a residential area consisting of 20 buildings. We subsequently applied this trained model to other configurations of the residential area, such as those with 10 and 30 buildings, without retraining. The results reveal that the trained model is thus scalable and can be applied to different sizes of the residential area.

We implemented the simulations and the optimization in the programming language *Python*. For the MILP we use the package *Pyomo* [60], for the evolutionary algorithms *Pymoo* [61], for the RL *Stable-Baselines3* *stable-baselines3* [62] and *gymnasium* [63]. All methods use a super-ordinate additional controller that adjusts their actions if a constraint violation is about to occur. The code repository containing the commented code and an executable notebook is openly available at [64]. The input data and result profiles are provided at [65].

Fig. 10 in the appendix plots exemplary profiles for one day (Day 332) resulting from the different control approaches and input profiles (DHW, Space heating, number of EVs, outside temperature, and price). The profiles result from a residential area with 20 buildings (8 BT1, 8 BT2, 4 BT3).

5.3. Critical appraisal

We made some simplifications in our study to carry out the experiments. To solve the optimization problem optimally, we assumed a perfect forecast of future demands. Thus, we did not incorporate any uncertainties. To cope with uncertainties of energy consumption, forecasts of energy demand, for example, using artificial neural networks,

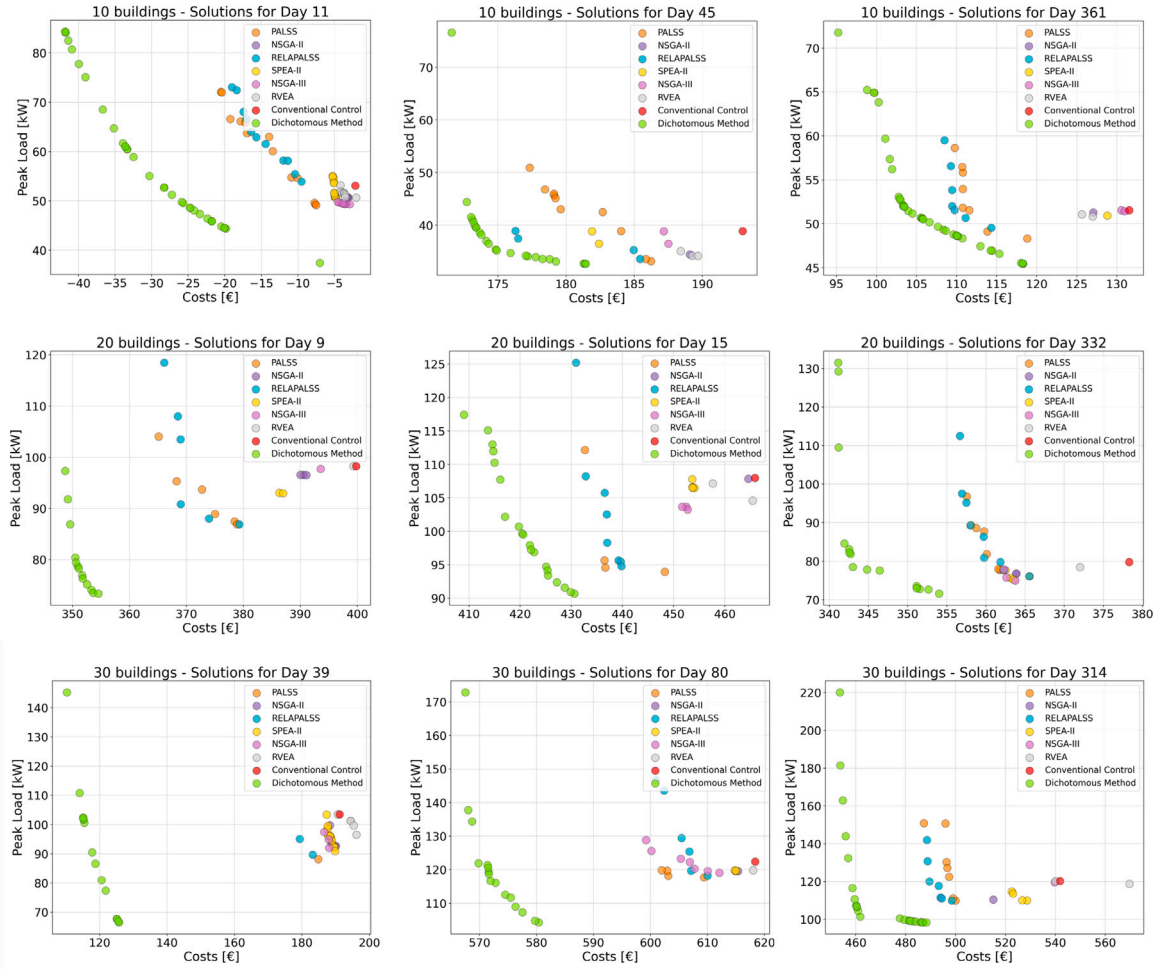


Fig. 9. Pareto fronts of the different methods.

are necessary [66]. Further, the building model for the optimization assumes one uniform temperature in the whole building and does not consider heat flows between different rooms or internal heat gains. Moreover, we use a centralized control approach to control the flexible loads. While this typically leads to the best results, it might infringe on the privacy of the inhabitants [67]. It should be noted that the control mechanisms we introduce require an aggregated schedule for the residential area, which can be obtained either through data-driven forecasting or via decentralized load profile sharing, as proposed by Dengiz et al. [67]. Thus, they can be combined with decentralized scheduling-based control approaches. Hence, we are convinced that our approaches can be used in many simulations and real-world applications. We also set a maximum runtime for the single-objective optimization problems that the dichotomous method repeatedly solves. Thus, it could happen that the exact methods did not solve the optimization problem to optimality. Because of this, what we refer to as the true Pareto-front that we use for evaluation, is itself only an approximation. However, this approximation is very close to the theoretical true front. It should be noted that the introduced approaches are heuristic in nature and therefore cannot guarantee globally optimal solutions. However, they are capable of finding high-quality solutions efficiently, even for large-scale and complex problem instances.

6. Conclusion

In this paper, we introduce two novel approaches for the multi-objective optimization of the conflicting goals to minimize the electricity cost and the peak load in a residential area. The approaches

intelligently control flexible electrical devices like heat pumps and electric vehicles such that the electricity consumption of the residential area reacts to an external electricity price signal. The two approaches *Pareto local search for shifting loads (PALSS)* and *Reinforcement learning assisted Pareto local for shifting loads (RELAPALSS)* define a concept for Pareto local search with heuristic neighborhood searching operations. *RELAPALSS* additionally uses reinforcement learning to adjust the parameters of the local search operations.

Our simulations on a residential area with multiple building types and varying sizes demonstrate the effectiveness of our approach. The key findings are:

- PALSS and RELAPALSS strongly outperform state-of-the-art evolutionary algorithms (NSGA-II, NSGA-III, RVEA, SPEA-II) in the two key evaluation metrics for multi-objective optimization: Generational Distance (up to 18% improvement) and Hypervolume (up to 130% improvement).
- The runtime of PALSS and RELAPALSS is reduced by up to 92% compared to the exact Dichotomous method.
- Including reinforcement learning in RELAPALSS leads to further improvements in solution quality of about 2%.
- The proposed methods significantly reduce both electricity costs and peak loads compared to conventional control approaches and metaheuristics. Thus, they can contribute to better utilizing volatile renewable energy sources in residential areas.

In future work, we intend to analyze the use of imitation learning in combination with domain knowledge to improve our local search

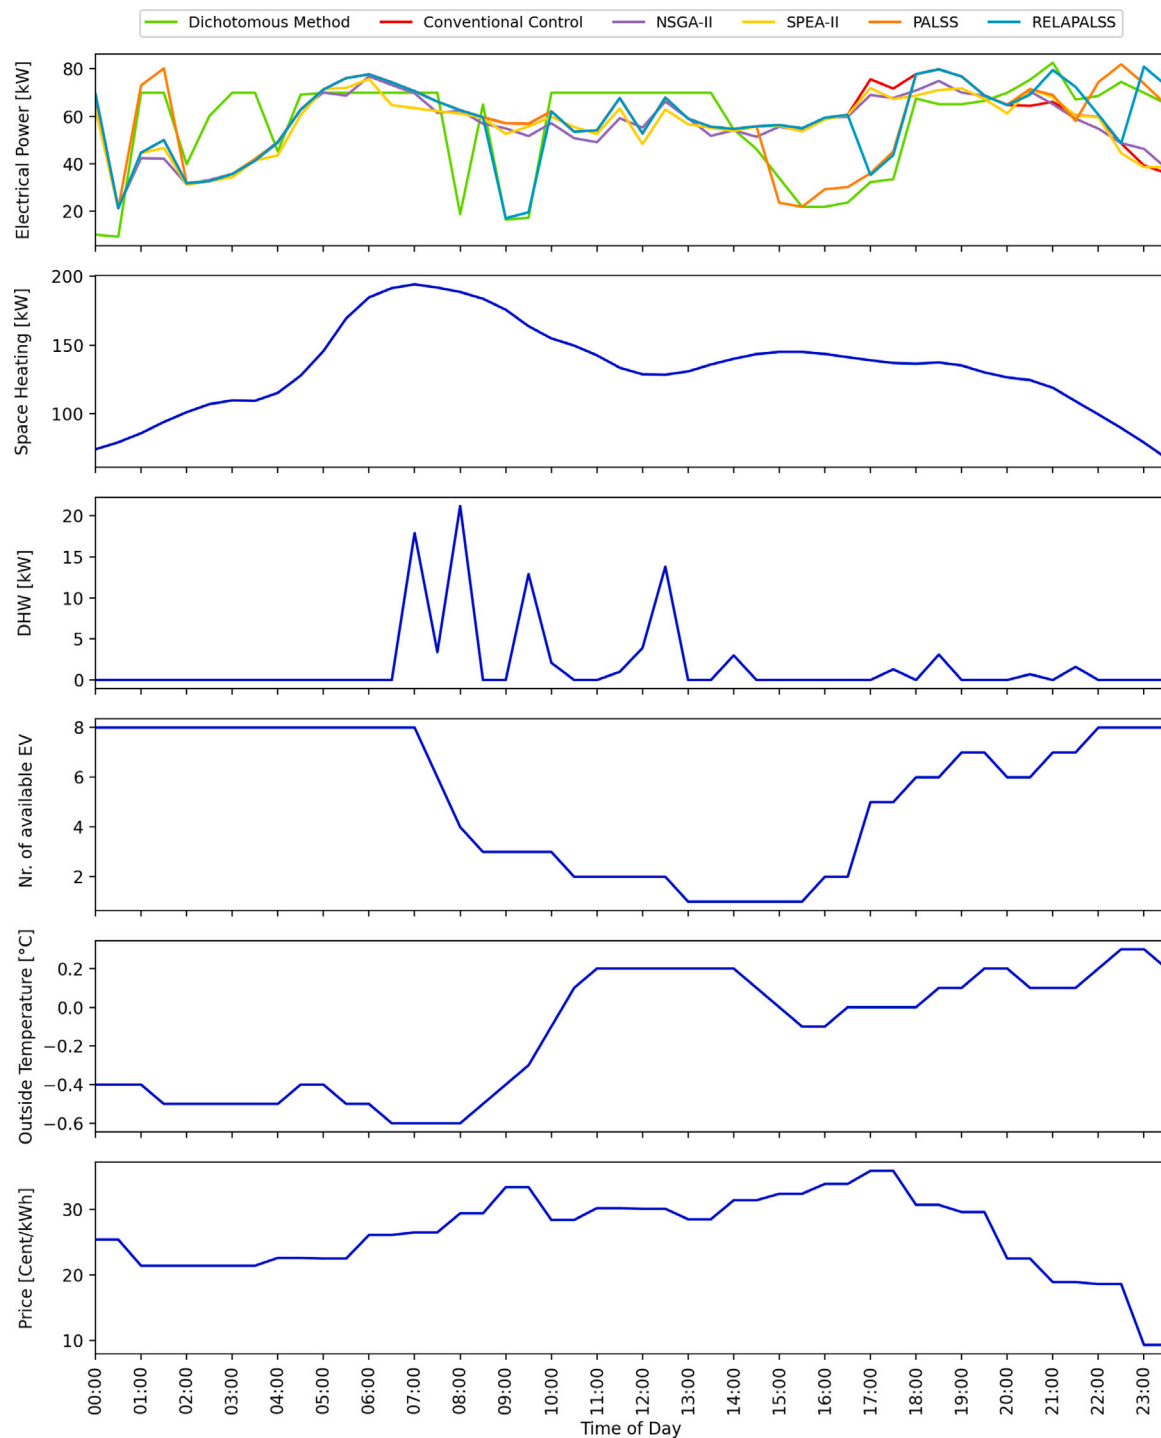


Fig. 10. Exemplary profiles for one day in November for a residential area with 20 buildings.

approaches further. Moreover, we plan to integrate other flexibility options like batteries or additional electric heating elements in gas heating systems into our model. Also, combining our developed methods with decentralized coordination approaches constitutes a possible future task.

CRediT authorship contribution statement

Thomas Dengiz: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Andrea Raith:** Writing – original draft, Methodology, Investigation, Conceptualization. **Max Kleinebrahm:** Writing – original draft,

Validation, Resources, Data curation. **Jonathan Vogl:** Writing – review & editing, Validation, Software. **Wolf Fichtner:** Writing – review & editing, Validation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by the project AsimutE (Autoconsommation et Stockage Intelligents pour une Meilleure Utilisation de l'Énergie) from the European Territorial Cooperation program Interreg, Germany. The authors also want to thank the Karlsruhe House of Young Scientists (KHYS) for enabling a research stay, which made the necessary collaborations for this paper possible.

Appendix

Exemplary profiles for one day in November for a residential area with 20 buildings are presented in Fig. 10 (to enhance visual clarity, only the two best-performing metaheuristics for that day are included in the illustration).

Data availability

Both the code and data are openly accessible. The repository containing the commented code and an executable notebook is available at [GitHub]. The input data and result profiles can be accessed at [KITOpen].

References

- [1] Shao S, Pipattanasomporn M, Rahman S. Demand response as a load shaping tool in an intelligent grid with electric vehicles. *IEEE Trans Smart Grid* 2011;2(4):624–31. <http://dx.doi.org/10.1109/TSG.2011.2164583>.
- [2] Patteeuw D, Bruninx K, Arteconi A, Delarue E, D'haeseleer W, Helsen L. Integrated modeling of active demand response with electric heating systems coupled to thermal energy storage systems. *Appl Energy* 2015;151:306–19. <http://dx.doi.org/10.1016/j.apenergy.2015.04.014>.
- [3] Eurostat. Energy use in EU households in 2022 lowest since 2016. 2024, URL <https://ec.europa.eu/eurostat/web/products-eurostat-news/w/ddn-20240605-2>, [Accessed 23 September 2024].
- [4] Dileep G. A survey on smart grid technologies and applications. *Renew Energy* 2020;146:2589–625. <http://dx.doi.org/10.1016/j.renene.2019.08.092>, URL <https://www.sciencedirect.com/science/article/pii/S0960148119312790>.
- [5] Shariatzadeh F, Mandal P, Srivastava AK. Demand response for sustainable energy systems: A review, application and implementation strategy. *Renew Sustain Energy Rev* 2015;45:343–50. <http://dx.doi.org/10.1016/j.rser.2015.01.062>.
- [6] Kühnrich M, Stute J, Klingler A-L. Impacts of avalanche effects of price-optimized electric vehicle charging - does demand response make it worse? *Energy Strat Rev* 2021;34:100608. <http://dx.doi.org/10.1016/j.esr.2020.100608>.
- [7] Misconel S, Zimmermann F, Mikurda J, Möst D, Kunze R, Gnann T, et al. Model coupling and comparison on optimal load shifting of battery electric vehicles and heat pumps focusing on generation adequacy. *Energy* 2024;305:132266. <http://dx.doi.org/10.1016/j.energy.2024.132266>, URL <https://www.sciencedirect.com/science/article/pii/S0360544224020401>.
- [8] Ullman J. NP-complete scheduling problems. *J Comput System Sci* 1975;10(3):384–93. [http://dx.doi.org/10.1016/S0022-0000\(75\)80008-0](http://dx.doi.org/10.1016/S0022-0000(75)80008-0).
- [9] Dengiz T, Jochem P. Decentralized optimization approaches for using the load flexibility of electric heating devices. *Energy* 2020;193:116651. <http://dx.doi.org/10.1016/j.energy.2019.116651>, URL <https://www.sciencedirect.com/science/article/pii/S0360544219323461>.
- [10] Dengiz T. Optimization approaches for exploiting the load flexibility of electric heating devices in smart grids [Ph.D. thesis], Karlsruher Institut für Technologie (KIT); 2021. <http://dx.doi.org/10.5445/IR/1000131495>.
- [11] Zheng F, Jin K, Xu Y, Liu M. Unrelated parallel machine scheduling with processing cost, machine eligibility and order splitting. *Comput Ind Eng* 2022;171:108483. <http://dx.doi.org/10.1016/j.cie.2022.108483>, URL <https://www.sciencedirect.com/science/article/pii/S0360835222005058>.
- [12] Branke J. In: *Multiobjective optimization: Interactive and evolutionary approaches*, vol. 5252, Springer Science & Business Media; 2008.
- [13] Luo Q, Wu G, Ji B, Wang L, Suganthan PN. Hybrid multi-objective optimization approach with Pareto local search for collaborative truck-drone routing problems considering flexible time windows. *IEEE Trans Intell Transp Syst* 2022;23(8):13011–25. <http://dx.doi.org/10.1109/TITS.2021.3119080>.
- [14] Paquete L, Chiarandini M, Stützle T. Pareto local optimum sets in the biobjective traveling salesman problem: An experimental study. In: Gandibleux X, Sevaux M, Sörensen K, T'kindt V, editors. *Metaheuristics for multiobjective optimisation*. Berlin, Heidelberg: Springer Berlin Heidelberg; 2004, p. 177–99. http://dx.doi.org/10.1007/978-3-642-17144-4_7.
- [15] Pinto G, Piscitelli MS, Vázquez-Canteli JR, Nagy Z, Capozzoli A. Coordinated energy management for a cluster of buildings through deep reinforcement learning. *Energy* 2021;229:120725. <http://dx.doi.org/10.1016/j.energy.2021.120725>.
- [16] Peirelinck T, Hermans C, Spiessens F, Deconinck G. Combined peak reduction and self-consumption using proximal policy optimisation. *Energy AI* 2024;16:100323. <http://dx.doi.org/10.1016/j.egyai.2023.100323>.
- [17] Pinto G, Deltetto D, Capozzoli A. Data-driven district energy management with surrogate models and deep reinforcement learning. *Appl Energy* 2021;304:117642. <http://dx.doi.org/10.1016/j.apenergy.2021.117642>.
- [18] Shakouri G. H, Kazemi A. Multi-objective cost-load optimization for demand side management of a residential area in smart grids. *Sustain Cities Soc* 2017;32:171–80. <http://dx.doi.org/10.1016/j.scs.2017.03.018>.
- [19] Tang H, Li Y, Li D, Wang H. Dynamic optimization of on-grid integrated energy system considering peak-shaving demand via learning methods. *IEEE Trans Electr Electron Eng* 2022;17(10):1409–19. <http://dx.doi.org/10.1002/tee.23651>.
- [20] de Souza Dutra MD, Anjos MF, Digabel SL. A framework for peak shaving through the coordination of smart homes. In: 2019 IEEE PES innovative smart grid technologies conference - Latin America. 2019, p. 1–6. <http://dx.doi.org/10.1109/ISGT-LA.2019.8895476>.
- [21] Khalid A, Javaid N, Guizani M, Alhussein M, Aurangzeb K, Ilahe M. Towards dynamic coordination among home appliances using multi-objective energy optimization for demand side management in smart buildings. *IEEE Access* 2018;6:19509–29. <http://dx.doi.org/10.1109/ACCESS.2018.2791546>.
- [22] Wynn SLL, Pinthurat W, Marungsri B. Multi-objective optimization for peak shaving with demand response under renewable generation uncertainty. *Energies* 2022;15(23). <http://dx.doi.org/10.3390/en15238989>.
- [23] Terlouw T, AlSkaif T, Bauer C, van Sark W. Multi-objective optimization of energy arbitrage in community energy storage systems using different battery technologies. *Appl Energy* 2019;239:356–72. <http://dx.doi.org/10.1016/j.apenergy.2019.01.227>.
- [24] Song Z, Guan X, Cheng M. Multi-objective optimization strategy for home energy management system including PV and battery energy storage. *Energy Rep* 2022;8:5396–411. <http://dx.doi.org/10.1016/j.egy.2022.04.023>.
- [25] Wang J, Liu Y, Ren F, Lu S. Multi-objective optimization and selection of hybrid combined cooling, heating and power systems considering operational flexibility. *Energy* 2020;197:117313. <http://dx.doi.org/10.1016/j.energy.2020.117313>.
- [26] Kazmi H, Suykens J, Balint A, Driesen J. Multi-agent reinforcement learning for modeling and control of thermostatically controlled loads. *Appl Energy* 2019;238:1022–35. <http://dx.doi.org/10.1016/j.apenergy.2019.01.140>.
- [27] Wei X, Kusiak A, Li M, Tang F, Zeng Y. Multi-objective optimization of the HVAC (heating, ventilation, and air conditioning) system performance. *Energy* 2015;83:294–306. <http://dx.doi.org/10.1016/j.energy.2015.02.024>.
- [28] Chiu W-Y, Hsieh J-T, Chen C-M. Pareto optimal demand response based on energy costs and load factor in smart grid. *IEEE Trans Ind Inform* 2020;16(3):1811–22. <http://dx.doi.org/10.1109/TII.2019.2928520>.
- [29] Schulman J, Wolski F, Dhariwal P, Radford A, Klimov O. Proximal policy optimization algorithms. 2017. <http://dx.doi.org/10.48550/arXiv.1707.06347>, arXiv:1707.06347.
- [30] Ehrgott M. *Multicriteria optimization*. 2nd ed. Springer; 2005.
- [31] Cohon J. *Multiobjective programming and planning*. New York: Academic Press; 1978.
- [32] Boait P, Dixon D, Fan D, Stafford A. Production efficiency of hot water for domestic use. *Energy Build* 2012;54:160–8. <http://dx.doi.org/10.1016/j.enbuild.2012.07.011>.
- [33] Gomes A, Antunes CH, Martinho J. A physically-based model for simulating inverter type air conditioners/heat pumps. *Energy* 2013;50:110–9. <http://dx.doi.org/10.1016/j.energy.2012.11.047>.
- [34] Qian K, Zhou C, Allan M, Yuan Y. Modeling of load demand due to EV battery charging in distribution systems. *IEEE Trans Power Syst* 2011;26(2):802–10. <http://dx.doi.org/10.1109/TPWRS.2010.2057456>.
- [35] Entso-E Transparency Platform. Day-ahead prices. 2023, URL <https://transparency.entsoe.eu/>, [Accessed 24 April 2024].
- [36] Jordan U, Vajen K. Dhwcalc: Program to generate domestic hot water profiles with statistical means for user defined conditions. In: Proceedings of the ISES solar world congress. Orlando, FL, USA; 2005, URL <https://www.uni-kassel.de/maschinenbau/en/institute/thermische-energietechnik/fachgebiete/solar-und-anlagentechnik/downloads>, [Accessed 24 April 2024].
- [37] Bausparkasse Schwäbisch Hall AG. Ein einfamilienhaus bauen: So wird der traum wirklichkeit. 2024, URL <https://www.schwaebisch-hall.de/ratgeber/neubau-und-anbau/zusammen-leben/das-ideale-familienhaus.html>, [Accessed 24 April 2024].
- [38] ZimmerMeisterHaus Service & Dienstleistungen GmbH. Mehrfamilienhaus öpfingen. 2024, URL <https://www.zmh.com/wohnbau-gewerbearbeit/beispiele-wohn-gewerbearbeit/portrait/mehrfamilienhaus-oepfingen.html>, [Accessed 24 April 2024].
- [39] German Institute for Standardization. DIN 18560-2:2022-08 estriche im bauwesen - teil 2: Estriche und heizestriche auf dämmschichten. 2022, URL <https://www.beuth.de/de/norm/din-18560-2/356233036>, [Accessed 24 April 2024].
- [40] ökoloco GmbH. Warmwasserspeicher: Vor- und nachteile, größe und hersteller. 2024, URL <https://oekoloco.de/heizungen/heizungskomponenten/warmwasserspeicher/>, [Accessed 24 April 2024].

- [41] Umweltbundesamt. Ökodesign-richtlinie und energieverbrauchskennzeichnung - warmwasserbereiter und warmwasserspeicher. 2013, URL https://www.umweltbundesamt.de/sites/default/files/medien/376/dokumente/oekodesignrichtlinie_und_energieverbrauchskennzeichnung_warmwasserbereiter.pdf, [Accessed 24 April 2024].
- [42] Bundesverband Wärmepumpe (BWP) eV. Klimakarte. 2024, URL <https://www.waermepumpe.de/normen-technik/klimakarte/>, [Accessed 24 April 2024].
- [43] Bosch Thermotechnik GmbH. Luft-wasser-wärmepumpe compress 6800i AW. 2024, URL <https://www.bosch-homecomfort.com/de/de/ocs/wohngebäude/compress-6800i-aw-19312695-p/>, [Accessed 24 April 2024].
- [44] Allgemeiner Deutscher Automobil-Club eV(ADAC). Opel ampera-e first edition (07/17 - 01/18). 2018, URL <https://www.adac.de/rund-ums-fahrzeug/autokatalog/marken-modelle/opel/ampera/2generation/275080/>, [Accessed 24 April 2024].
- [45] Allgemeiner Deutscher Automobil-Club eV(ADAC). KEBA KeContact P30. 2024, URL <https://www.adac.de/rund-ums-fahrzeug/elektromobilität/tests/wallboxen/keba-kecontact-p30-id-4180/>, [Accessed 24 April 2024].
- [46] Federal ministry of transport and digital infrastructure. Mobilität in deutschland - MiD. 2018, URL https://bmdv.bund.de/SharedDocs/DE/Anlage/G/mid-ergebnisbericht.pdf?__blob=publicationFile, [Accessed 24 April 2024].
- [47] Richardson I, Thomson M, Infield D, Clifford C. Domestic electricity use: A high-resolution energy demand model. *Energy Build* 2010;42(10):1878–87. <http://dx.doi.org/10.1016/j.enbuild.2010.05.023>.
- [48] Statistisches Bundesamt (Destatis). Einkommen, konsum und lebens bedingungen - zeitverwendung. 2024, URL https://www.destatis.de/DE/Themen/Gesellschaft-Umwelt/Einkommen-Konsum-Lebensbedingungen/Zeitverwendung/_inhalt.html#sprg234984, [Accessed 24 April 2024].
- [49] Copernicus Climate Change Service. ERA5: Fifth generation of ECMWF atmospheric reanalyses of the global climate - datasets. 2017, URL <https://cds.climate.copernicus.eu/#!/search?text=ERA5&type=dataset>.
- [50] BDEW Bundesverband der Energie- und Wasserwirtschaft eV. Stromverbrauch im haushalt 2016. 2016, <https://www.bdew.de/media/documents/BDEWStromverbrauchimHaushaltJuli2016Charts.pdf>, [Accessed 17 March 2016].
- [51] Institut Wohnen und Umwelt GmbH. Deutsche wohngebäudetypologie. Beispielhafte maßnahmen zur verbesserung der energieeffizienz von typischen wohngebäuden – zweite erweiterte auflage -. 2015, URL https://www.episcope.eu/downloads/public/docs/brochure/DE_TABULA_TypologyBrochure_IWU.pdf, [Accessed 24 April 2024].
- [52] Paetz A-G, Kaschub T, Jochem P, Fichtner W. Load-shifting potentials in households including electric mobility - a comparison of user behaviour with modelling results. In: 2013 10th international conference on the European energy market. 2013, p. 1–7. <http://dx.doi.org/10.1109/EEM.2013.6607324>.
- [53] Deb K, Pratap A, Agarwal S, Meyarivan T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans Evol Comput* 2002;6(2):182–97. <http://dx.doi.org/10.1109/4235.996017>.
- [54] Zitzler E, Laumanns M, Thiele L. SPEA2: Improving the strength Pareto evolutionary algorithm. TIK report 103, ETH Zurich; 2001, URL <https://www.research-collection.ethz.ch/handle/20.500.11850/145755>.
- [55] Deb K, Jain H. An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, part I: Solving problems with box constraints. *IEEE Trans Evol Comput* 2014;18(4):577–601. <http://dx.doi.org/10.1109/TEVC.2013.2281535>.
- [56] Cheng R, Jin Y, Olhofer M, Sendhoff B. A reference vector guided evolutionary algorithm for many-objective optimization. *IEEE Trans Evol Comput* 2016;20(5):773–91. <http://dx.doi.org/10.1109/TEVC.2016.2519378>.
- [57] Ishibuchi H, Masuda H, Tanigaki Y, Nojima Y. Modified distance calculation in generational distance and inverted generational distance. In: Gaspar-Cunha A, Henggeler Antunes C, Coello CC, editors. *Evolutionary multi-criterion optimization*. Cham: Springer International Publishing; 2015, p. 110–25. http://dx.doi.org/10.1007/978-3-319-15892-1_8.
- [58] Bader J, Zitzler E. HypE: An Algorithm for Fast Hypervolume-Based Many-Objective Optimization. *Evol Comput* 2011;19(1):45–76. <http://dx.doi.org/10.1162/EVCO.a.00009>.
- [59] Gurobi Optimization L. Gurobi Optimizer Reference Manual. 2023, URL <https://www.gurobi.com>.
- [60] Hart WE, Watson J-P, Woodruff DL. Pyomo: modeling and solving mathematical programs in python. *Math Program Comput* 2011;3(3):219–60. <http://dx.doi.org/10.1007/s12532-011-0026-8>.
- [61] Blank J, Deb K. Pymoo: Multi-objective optimization in python. *IEEE Access* 2020;8:89497–509. <http://dx.doi.org/10.1109/ACCESS.2020.2990567>.
- [62] Raffin A, Hill A, Gleave A, Kanervisto A, Ernestus M, Dormann N. Stable-Baselines3: Reliable reinforcement learning implementations. *J Mach Learn Res* 2021;22(268):1–8, URL <http://jmlr.org/papers/v22/20-1364.html>.
- [63] Towers M, Terry JK, Kwiatkowski A, Balis JU, Cola Gd, Deleu T, et al. *Gymnasium*. Zenodo; 2023, <http://dx.doi.org/10.5281/zenodo.8127026>, URL <https://zenodo.org/record/8127025>.
- [64] Dengiz T, Vogl J. MultiObjective_LocalSearch_DSM. 2024, GitHub repository, https://github.com/thomasdengiz/MultiObjective_LocalSearch_DSM.
- [65] Dengiz T, Kleinebrahm M. Input data and result data for the paper “Pareto local search for a multi-objective demand response problem in residential areas with heat pumps and electric vehicles”. 2024, Repository KITopen, <http://dx.doi.org/10.35097/ZxeqNfKvVIQcjQAt>.
- [66] Neto AH, Fiorelli FAS. Comparison between detailed model simulation and artificial neural network for forecasting building energy consumption. *Energy Build* 2008;40(12):2169–76. <http://dx.doi.org/10.1016/j.enbuild.2008.06.013>, URL <https://www.sciencedirect.com/science/article/pii/S037878808001448>.
- [67] Dengiz T, Jochem P, Fichtner W. Demand response through decentralized optimization in residential areas with wind and photovoltaics. *Energy* 2021;223:119984. <http://dx.doi.org/10.1016/j.energy.2021.119984>.