

University of Groningen

Energy Smart Buildings

Setz, Brian; Haghshenas, Kawsar; Aiello, Marco

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version

Early version, also known as pre-print

Publication date:

2022

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Setz, B., Haghshenas, K., & Aiello, M. (2022). *Energy Smart Buildings: Parallel Uniform Cost-Search with Energy Storage and Generation*. arXiv.

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Energy Smart Buildings: Parallel Uniform Cost-Search with Energy Storage and Generation

Brian Setz*, Kawsar Haghshenas†, and Marco Aiello‡
Service Computing Department, University of Stuttgart
Stuttgart, Germany
Email: {firstname}.{lastname}@iaas.uni-stuttgart.de

Abstract—The amalgamation of Internet of Things and the smart grid enables the energy optimal scheduling of appliances based on user needs and dynamic energy prices. Additionally, progress in local storage technology calls for exploiting additional sources of flexibility. In this paper, we propose a scheduling approach for building operation management, considering factors such as energy storage, local energy generation, and dynamic energy prices. In addition, we propose a new optimization strategy to discover the optimal scheduling of devices. Our approach utilizes parallel uniform cost-search to explore the complex search space and to find the optimal schedule within a user-acceptable amount of time. The evaluation utilizes real-world data for the devices, and the price signals, while the architecture is designed following a micro-service approach, enabling modularity and loose-coupling. The evaluation shows that including local energy storage as part of the optimization problem further reduces overall costs by up to 22.64% when compared to schedules without energy storage. Parallel uniform cost-search decreases the time to find the optimal schedule by a factor of 4.7 with respect to the traditional uniform cost-search algorithm.

Index Terms—Smart Buildings, Internet of Things, Device Scheduling, Demand Response, Sustainability

I. INTRODUCTION

In 2022, gas and electricity prices have surged in Europe as fossil fuel prices rise, emphasising the need for solutions to reduce energy use [1]. Buildings are a major source of energy usage and a natural place to look for cutting consumptions and costs. Building operations management encompasses the controlling and monitoring of buildings, and can assist in the optimization of buildings. Building management benefits greatly from two emerging paradigms: the smart power grid and the digital interface to everyday objects known as the Internet of Things (IoT). The smart grid aims to integrate the behaviours and actions of all connected actors, be it consumers, producers, or prosumers by leveraging on information and communication technologies. Local energy generation and storage, economic efficiency, and sustainability are key concepts associated with the smart grid [2]. Whereas the IoT paradigm enables a wide variety of devices and machines to interact with one another programmatically and autonomously [3]. This enables the vision of smart buildings, which are able to use devices to sense and control, and in turn

influence the energy consumption of the building by reacting to different signals while satisfying user needs. The dynamic real time price of energy is a prototypical signal which the smart grid is envisioned to provide to end users [4]. This signal is influenced by weather condition, fossil fuel prices, and current demand [5], [6]. End users receive an incentive to actively participate by changing and rescheduling their load patterns based on current prices, while not being forced to [7], [8]. This differs from demand response schemes where the users leave the ultimate control of their consumption to the energy provider [9].

Smart buildings are able to sense their context of use and control systems and devices based on user requirements and external signals [10]. As a reaction to the dynamic price signal, the building can operate on subsystems by actuating on them, and by postponing their operation to times where the signal is expected to be more favourable. For example, the scheduled time of heating for a meeting room can be advanced or postponed, if the energy prices are expected to increase considerably during the scheduled heating period. However, deciding when to schedule a load in order to minimise costs is a non trivial task, especially when local energy generation and storage add uncertainty to the overall set of possible decisions to make.

We have previously proposed an architecture and optimization strategy to control office appliances based on user needs and dynamic prices [11]. Since then, we have witnessed a widespread adoption of IoT solutions to monitor and control appliances and building systems, IoT and data science approaches to building context and human activity recognition, as well as rapid developments in electricity storage technologies. This provides both challenges and opportunities for building energy optimization.

One of the major challenges are that more uncertainty is present in the system and that the complexity of monitoring and decision making increases. The opportunity comes from the major flexibility and controllability of the system. Therefore, we approach the problem considering the current state of the art and resort to improved techniques for the optimal scheduling of office buildings. Specifically, we present an approach to find the optimal load schedule that minimises costs while considering (1) price signals from the grid and prosumers, (2) local energy generation, (3) local energy storage, and (4) the scheduling constraints of typical office

The presented research is partially funded by the Dutch Research Council (NWO) in the framework of the Indo-Dutch Science Industry Collaboration programme with project NextGenSmart DC (629.002.102).

devices. Based on this information, an optimization problem is formulated. The resulting search space is large, as for each time slot there are many permutations of possible actions that can be taken. The contributions made and presented here include:

- A general optimal device scheduling algorithm for smart grid enabled buildings that considers: photovoltaic panels, micro wind turbines, and battery energy storage. Our evaluation shows a reduction in costs by up to 22.64%;
- A parallel uniform-cost search algorithm to find the optimal schedule that minimizes overall costs while ensuring performance when navigating complex search spaces. We observe a decrease in search time by a factor of 4.7 compared to traditional uniform-cost search;
- A micro-service architecture for the integration of building management systems with independent cloud services, enabling a loosely coupled, modular systems that can be adjusted for any building; and,
- An evaluation using real data for the scheduled devices and for the prediction of renewable energy generation. The device data was collected from an office environment over a period of eight months through plug load monitoring, with measurements taken every 10 seconds.

The rest of the paper is organised as follows. In Section II, we review related work and place the present contribution in the context of it. The problem definition, proposed approach, models, scheduling policies, and the architecture of the various components of the systems are described in Section III. This is followed by the evaluation with real data, in Section IV. The evaluation considers both the economic and performance aspects. The discussion of the results, the limitations, as well as concluding remarks are summarized in Section V.

II. RELATED WORK

A common approach to the scheduling of loads is using Mixed-Integer Linear Programming (MILP) models. Paterakis et al. present a method for optimal scheduling of household appliances [12]. The scheduling is performed under hourly pricing and peak power limiting demand response strategies. The authors utilize a MILP model for optimal scheduling, distinguishing between thermostatically and non-thermostatically controllable devices. The results show that the load factor can be improved significantly, while the economic costs slightly increase. Duman et al. also demonstrate the use of a MILP model for scheduling the operation of shiftable loads [13]. The distinction is made between time-shiftable loads, power-shiftable loads, and thermostatically controlled loads. Their simulation shows that daily costs can be reduced by 53.2%. A similar method has been proposed by Nan et al. for optimal demand response scheduling in residential communities [14]. A clear distinction is made between interruptible loads, shiftable loads, and distributed generation. The authors define a MILP problem to be solved in order to find a solution that reduces peak load and peak-valley differences and, in turn, the overall cost of power. The end user costs are reduced by up to 1.52%.

While MILP appears to be the predominant approach, alternatives have been proposed. Amer et al. propose a home energy management system centred around a multi-objective optimization problem [15]. A distinction is made between shiftable loads, non-shiftable loads, and active loads. Energy storage and PV energy resources are also considered. The objective is to balance the benefits for the end-user and the Distribution System Operator (DSO). The results show a reduction in energy cost of 31%, and a reduction in demand peak by 18%. Lu et al. present a different approach for demand response management using reinforcement learning and artificial neural networks [16]. While a distinction between different types of loads is made, energy storage and renewable generation are not considered. The obtained cost savings are between 7.3% and 72.3%. The authors demonstrate that their approach outperforms MILP as the number of iterations increases. The method proposed by Li et al. utilises deep-reinforcement learning for the optimal scheduling of home appliances [17]. The appliances are split into three categories: deferrable devices, regulatable devices, and critical devices. The approach reduces the overall electricity costs by 31.6%. An approach that is more closely related to the work that is presented in this paper, is the work of Fioretto et al. [18]. They define a constraint optimization problem, which is solved by a distributed algorithm that divides the problem into individual sub-problems to overcome the complexity of the problem. The savings compared to the baseline use case are over 50%. However, in our work the problem is not divided into sub-problems, as this would lead to non-optimal global solutions.

Finally, in our previous work, we demonstrated that the optimal scheduling of the operation of devices based on user-defined policies can result into significant savings [11]. From the economic perspective, the obtained savings are 35% on average. Furthermore, energy savings of 10% are also observed. The occupant satisfaction study in our previous work shows that comfort and satisfaction is preserved while our system operates the devices. In this work, we further tackle more complex scheduling problems by including real-time predictions of renewable energy generation, and the possibility for energy storage. To counteract the increased complexity, a novel algorithm is proposed that parallelises the solution of to the scheduling problem. The distinct subsystems are encapsulated in a micro-service architecture that promotes modularity and scalability.

The approach taken in our work also recognises the categorizations of shiftable, non-shiftable, and active loads as done in several existing studies. The difference is that these properties are encapsulated in fine grained policies which are assigned to individual devices. We consider energy storage to be a property of a device, and therefore it is also an assignable policy. Furthermore, we guarantee to find the optimal solution to the scheduling problem, as all possible and valid permutations of the state space are explored. In addition, the approach proposed in this work is distributed in nature, both the novel parallel uniform-cost search algorithm and the underlying micro-service architecture can be distributed

across multiple machines. Additionally, we use real-world and real-time data from multiple sources to predict the renewable energy generation and model the day-ahead prices. To the best of our knowledge, our proposed approach is the first to combine these aspects into a singular approach.

III. APPROACH

First we define the energy optimization problem in terms of the relationships between devices, their desired behaviour, and the available energy sources. This is followed by the models of the energy sources, in the present case: wind turbines, photovoltaic panels, and prosumers on the same local grid. Energy storage in the form of batteries is also modelled as part of the optimization problem. The last part of the problem definition relates to the user policies associated to the office devices. Such policies drive the way in which each one can be scheduled. Finally, we describe the algorithm to find the optimal schedule, and provide an overview of the proposed architecture.

A. Problem Definition

Given the devices to be scheduled, the scheduling policy for every device, and the available energy sources, the goal is to find an optimal day-ahead schedule that minimizes the costs while ensuring all of the constraints as defined by the scheduling policies are satisfied. The following problem definition is based and extends the definitions from our previous work [11].

A day is divided into time slots with a fixed duration. For every time slot t with length $t.duration$, $ES(t) = \{es_i\}$ represents the available energy sources at time slot t . An energy source consists of a tuple $es_i = \langle cost, energy \rangle$, where cost is the price per kWh at time t for energy source i , and energy is the available energy that the given energy source can provide at time t . The energy sources include both renewable energy sources and neighbouring prosumers.

The set of all building devices that are amenable to scheduling is D . For each device $d_i \in D$ there exists a tuple $d_i = \langle d_{id}, S_i \rangle$ where d_{id} is the identifier of the device, and S_i is the set of available states in which the device can be. Every state $s_{ij} \in S$ consists of the tuple $\langle s_{id}, power \rangle$ where s_{id} is the identifier of the state, and $power$ is the power consumption in that particular state.

Each device d_i has associated one or more scheduling policies that determine the scheduling behaviour for the device. The set P contains all of the policies that are associated with devices. Each policy $p_i \in P$ is represented by the triple $p_i = \langle p_{id}, type, parameters \rangle$ where p_{id} is the identifier of the policy, $type$ is the policy type, and $parameters$ are the specific parameters for that type of policy.

A schedule $X = \{x_{ti}\}$ is a set of values per device per time slot, where each value $x_{ti} \in S_i$ denotes the state that a device i assumes at time t . Then we have:

Definition 1 (Office optimization problem). *Find a schedule $X = \{x_{ti}\}, \forall t \in T, \forall i \in D$ that is optimal. An schedule is optimal iff*

$$\sum_{t \in T} cost(t, e_t) \rightarrow min, \forall p \in P : satisfied(p, X) \quad (1)$$

where

$$e_t = \sum_{d_i \in D} x_{ti}.power * t.duration \quad (2)$$

In other words, the schedule is optimal when the devices are scheduled in such a way that the policy constraints for each device are satisfied and the total costs of operation are minimal. To find the optimal solution requires exploration of the state space. The state space consists of all possible permutations of device states for each time slot. A schematic overview of the approach we propose is given in Figure 1.

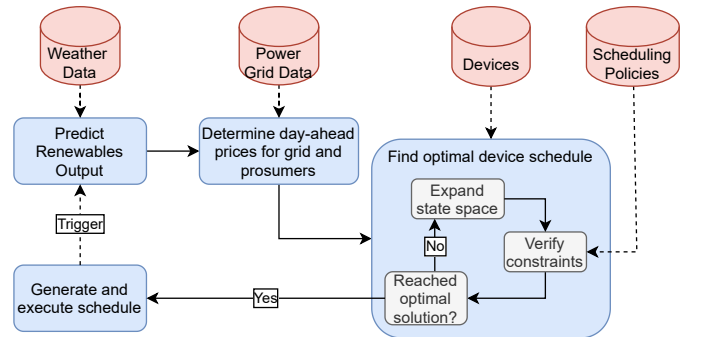


Fig. 1: A high-level overview of the approach proposed in this work

The process of finding an optimal schedule is started by a trigger, for example, the beginning of a new day. First, the energy generation of the renewable energy sources is estimated based on weather data and individual models for each renewable energy source. Next, the cost of energy is determined for each of the prosumers as well as the DSO. All of this data is passed to our scheduling algorithm which continuously expands the state space, verifies constraints, and checks if the optimal solution is found. The complexity of this state space depends on the set of devices, the assigned scheduling policy for each device, and the scheduling horizon. Once the optimal solution is found, a schedule is generated in JSON format, which can then be executed by the building management system. The process is repeated whenever a trigger takes place.

B. Energy Sources: Wind Turbines

One of the energy sources that is considered is local (on-site) energy generation using wind turbines. The following mathematical model is used to represent the wind turbine's power output and is based on the work of Xia et al. [19]:

$$P_w = \frac{1}{2} \rho C_p A v_w^3 \quad (3)$$

where:

P_w	power output (W);
ρ	air density (kg/m ³);
C_p	power coefficient;
A	wind turbine swept area (m ²);
v_w	wind speed (m/s).

The wind turbine swept area and the power coefficient are dependent on the particular wind turbine model. As per the Betz Limit, the theoretical maximum power coefficient is 0.59, though in practice the power coefficient is generally in the 0.35-0.45 range [20].

To determine the air density, first the saturation vapour pressure must be calculated. We resort to the work of Herman Wobus who has provided an approximating polynomial to calculate the vapour pressure [21]:

$$p_s = e_{so}/C^8$$

$$C = c_0 + T_d(c_1 + T_d(c_2 + T_d(c_3 + T_d(c_4 + T_d(c_5 + T_d(c_6 + T_d(c_7 + T_d(c_8 + T_d c_9)))))))))) \quad (4)$$

where:

p_s	saturation vapour pressure (hPa);
e_{so}	known constant (6.1078);
$c_0 - c_9$	known constants;
T_d	dew point (°C).

Once the saturation vapour pressure is known, the air density can be calculated according to the following equation:

$$p_d = p - p_s$$

$$\rho = \frac{100p_d}{R_{da}T_{ok}} + \frac{100p_s}{R_{wv}T_{ok}} \quad (5)$$

where:

p_d	dry air pressure (hPa);
p	air pressure (hPa);
p_s	saturation vapour pressure (hPa);
ρ	air density (kg/m ³);
R_{da}	dry air gas constant (J/kgK);
R_{wv}	water vapour gas constant (J/kgK);
T_{ok}	outside air temperature (K).

And finally, the cut-in and cut-out air speeds of the wind turbine have to be taken into account. If the air speed is above (cut-out) or below (cut-in) a certain value, the output of the turbine will be zero:

$$P_{turbine} = \begin{cases} P_w & \text{if } v_{in} \leq v_w \leq v_{out} \\ 0 & \text{otherwise} \end{cases} \quad (6)$$

where:

$P_{turbine}$	turbine power output (W);
v_{in}	cut-in air speed (m/s);
v_{out}	cut-out air speed (m/s).

The cut-in and cut-out air speed are dependent on the wind turbine model.

C. Energy Sources: Photovoltaic Panels

Photovoltaic panels are a common local generation means. We use the following formulation to model the photovoltaic panel's power output based on the work of Chenet al. [22]:

$$P_{photovoltaic} = \eta SI(1 - 0.005(T_{oc} - 25)) \quad (7)$$

where:

$P_{photovoltaic}$	photovoltaic power output. (W);
η	panel efficiency (%);
S	array area (m ²);
I	solar irradiation (W/m ²);
T_{oc}	outside air temperature (°C).

The array area and panel efficiency depend on the configuration and on the model of the photovoltaic panels. The solar irradiation and outside air temperature depend on the location and weather. This model assumes that the panels utilize a two-axis maximum power point tracker. As a result, the surface of the panel is always perpendicular to the sunlight. Therefore, the direct normal solar irradiance is used in the calculation of the panel's power output.

D. Energy Sources: Prosumers and DSO

The final energy sources considered in this work are the energy provided by neighbouring prosumers and by the DSO. These are both seen as energy providers independently of the underlying technology used. For instance the prosumer could also have photovoltaic panels or storage systems, while the DSO also has access to the energy coming from the transmission infrastructure. For both, we need to define the cost of the energy, and the amount of energy available during a given time slot. For prosumers, the available energy (kWh) is sampled from a normal distribution:

$$X \sim \mathcal{U}_{[a,b]} \quad (8)$$

whereas the cost of the energy (€/kWh) is sampled from a Gaussian distribution:

$$X \sim \mathcal{N}(\mu, \sigma^2) \quad (9)$$

As typical in practice, we assume that the DSO can always fulfil the demand of the system. Though, the cost of energy from the DSO will typically be higher than those of prosumers, especially in times of high demand. As for the cost of energy supplied by the DSO, this is modelled after the day-ahead energy market prices for DSOs, which are then scaled to values more representative of real-world end-user prices.

E. Energy Storage: Battery

Energy storage provides for important flexibility in complex smart energy systems. It can take many shapes and sizes, from fly wheels, to rechargeable batteries, to hydroelectric dams. In this work, residential energy storage in the form of household batteries is considered as it is the most easy and likely to be deployed in an office building. The battery should adhere to the following constraints, based on the definitions by Chen et

al. [22]. If $C(t)$ is the charge of the battery at time t , P_t^c the power charged to the battery, P_t^d the power discharged from the battery, and η is the charging efficiency, then

$$\begin{aligned} \text{Charge: } C(t+1) &= C(t) + \frac{\Delta t P_t^c}{\eta} \\ \text{Discharge: } C(t+1) &= C(t) - \frac{\Delta t P_t^d}{\eta} \end{aligned} \quad (10)$$

under the maximum charge and discharge limits:

$$\begin{aligned} 0 &\leq P_t^c \leq P_t^{c,max} \\ 0 &\leq P_t^d \leq P_t^{d,max} \end{aligned} \quad (11)$$

energy storage limits:

$$C_{min} \leq C(t) \leq C_{max} \quad (12)$$

and the start and end limits:

$$C(t_{start}) = C(t) = C(t_{end}) \quad (13)$$

where $P_t^{c,max}$ and $P_t^{d,max}$ are the maximum charge rate and maximum discharge rate respectively, C_{min} and C_{max} are the minimum and maximum energy that the battery can store, and $C(t_{start})$ and $C(t_{end})$ represent the charge of the battery at the beginning and end of the scheduling period.

F. Scheduling Policies

The way devices operate varies depending on their main function and user needs. This influences the way they are scheduled. To represent this, we have introduced several policies to represent usage [11]. The policies are then translated into scheduling constraints when solving the optimization problem. The following scheduling policies are available:

- **Total Policy:** defines the total amount of time a device d_i should be in state s_{ij} . For example, for any device with a battery that needs to be recharged, this is the time it takes to fully charge it. The amount of time is fixed, though it need not be continuous.
- **Continuous Policy:** similar to the *total* policy, it defines the total amount of time a device d_i should be in state s_{ij} . The key difference being that once the device is in state s_{ij} , it should remain in this state for the given amount of time without interruptions.
- **Repeat Policy:** defines the amount of time a device d_i should be in state s_{ij} with a certain periodicity. For example, a freezer that should be scheduled repeatedly every hour to maintain the appropriate temperature.
- **Multiple Policy:** combines the *continuous* and *repeat* policies. This policy is for devices that need to be in state s_{ij} for a certain number of times, but should not be interrupted before a certain deadline once they enter that state. For example, a printer can have multiple jobs to complete, but cannot be interrupted while executing the job.
- **Strict Policy:** defines a schedule that is known ahead of time, determining exactly when device d_i should be

in state s_{ij} , for every available time slot. For example, a safety light that should be on from dusk till dawn. It would not be acceptable to only turn the light on for 15 minutes every hour.

- **Pattern Policy:** similar to the *strict* policy, it defines ahead of time what the expected usage of a device is. However, this policy does not offer the possibility of the device to be controlled, it simply provides the expected usage pattern.
- **Sleep Policy:** defines a no-op policy, during the entire specified period, device d_i should be in state s_{ij} and should not be operated upon. For example, this is useful when a device should be turned off for a period of time in order to save energy.
- **Battery Policy:** indicates that a device d_i is able to act as short-term energy storage, allowing for charging and discharging of energy. Examples of devices that can have the battery policy include stationary home energy storage devices, as well as connected electric vehicles.

The formal definitions for each policy, except for the *battery* policy which we introduce for the purpose of the present work, can be found in [23].

G. Parallel Uniform-cost Search

The goal of this work is to find the optimal solution to the previously defined problem. Optimal in this case, refers to the solution that minimizes the overall costs. The cost of the solution can be defined in multiple ways, for example, as the total cost of electricity, or as the carbon footprint associated with the energy consumption. To solve the problem optimally, the state space has to be explored. The state space can be interpreted as a weighted graph, where each node is a permutation of the state space, edges represent the transformation from one state to another and have an associated cost, the depth of the graph is equal to the number of timeslots. Uniform-cost search is an uninformed search algorithm used for traversing graphs, finding the path with the lowest cost from the root node to the destination node [24]. This work presents a novel parallel uniform-cost search algorithm, which finds the optimal device schedule significantly faster than regular uniform-cost search.

The proposed algorithm is given in Algorithm 1. The search is initialized by creating an empty node and adding it to the global priority queue. Each node contains a permutation of device states, and the cost associated with being in this state. The priority queue ensures that the node with the lowest cost is the first node in the queue. Each thread also maintains a local priority queue. When the local queue is empty, the thread takes a node from the global queue and starts expanding the state space while populating its local queue. These newly expanded states are added back into the priority queue, assuming none of the policies are violated and they pass the heuristic checks. This is repeated until one of the threads finds a solution to the scheduling problem. Once a possible solution is found, the local queues are merged with the global queue to verify that there are no nodes with a lower cost in any of the local queues. If this is true, then the optimal solution is returned.

Two heuristics are applied to reduce the complexity of the search space. The first heuristic check is to determine whether the newly expanded states violate any policies. If a policy is violated, the state is invalid and should not be explored further. For example, a device with a *sleep* policy from 01:00 AM till 06:00 AM should not be scheduled during this given period. If the device is scheduled during this period in the newly expanded state, then the state violates the policy and therefore the child node is pruned from the search space. The second heuristic checks if there are partial solutions that have already been explored and that have the same outcome (cost and number of time slots) as the newly expanded state. If the outcome of the partial solution is a duplicate, it is discarded. In case the outcome is the same, but the cost is lower, the original node with the same outcome is replaced by the newly created child node. In case the outcome of both nodes are identical, the child node is pruned from the search space.

H. Architecture

The cost optimization of smart grid enabled buildings requires the availability of a wide variety of data sources. For example, the reliable prediction of the renewable generation for the next day is based on the weather forecast. Weather services vary considerably in terms of the predicted features (radiation, temperature, wind, humidity, etc.), the spatial granularity, and the temporal resolution. The microservice architecture that we design and propose enables the decoupling of the system's components, where each loosely coupled service adheres to the single responsibility principle [25]. Furthermore, this architectural design promotes modularity, allowing services to be easily replaced by alternative services, even at runtime [26]. An example is replacing the weather forecast of provider A, by the forecasting service of provider B, because provider B provides better local forecasts at an increased resolution without affecting the overall operation of the system. The proposed microservice architecture is shown in Figure 2. It consists of nine distinct microservices, identified by the rounded rectangles. The arrows indicate the interdependencies between them. Four of these services depend on external components, which can be cloud services (e.g. DarkSky for weather data), or data repositories (e.g. MongoDB for storing device data). In case a different weather forecast supplier should be used, it is a matter of replacing the weather forecast service with a different one. Communication between services is exclusively done via the REST architectural style, as each REST service exposes a REST API that accepts HTTP requests. The Content-Type of each HTTP response is *application/json*. The deployment of the microservices is handled by means of containerization, where each service is hosted within individual Docker container. What follows next is a description of each microservice.

Day Ahead Prices Service: requires an Energy Identification Code to identify the energy market that is used, and a date for which the day ahead prices should be retrieved. This service uses the cloud API provided by the European Network of Transmission System Operators transparency

Algorithm 1 Parallel Uniform-cost Search

```

1: procedure UCS_PARALLEL(problem, threads)
2:   initialNode  $\leftarrow$  new(node)
3:   globalQueue  $\leftarrow$  push(globalQueue, initialNode)
4:   globalVisited  $\leftarrow$   $\emptyset$ 
5:   solutionFound  $\leftarrow$  False
6:   for i  $\leftarrow$  1 to threads do
7:     startNode  $\leftarrow$  pop(globalQueue)
8:     localQueue  $\leftarrow$  push(localQueue, startNode)
9:     localVisited  $\leftarrow$   $\emptyset$ 
10:    while not solutionFound do
11:      node  $\leftarrow$  pop(localQueue)
12:      if goalReached(problem, node.state) then
13:        synchronizeThreads(threads)
14:        if peek(globalQueue).cost > node.cost
15:          then
16:            solutionFound  $\leftarrow$  True
17:            return solution(node)
18:          end if
19:        end if
20:        localVisited  $\leftarrow$  insert(localVisited, node)
21:        for newState in expand(problem, node.state)
22:          do
23:            c  $\leftarrow$  createChild(problem, node, newState)
24:            if policiesViolated(c.state) then
25:              continue
26:            else if c.state not in localQueue or visited
27:              then
28:                localQueue  $\leftarrow$  push(localQueue, c)
29:                else if c.state in localQueue with higher
30:                cost then
31:                  localQueue  $\leftarrow$  replace(localQueue,
32:                  c.state, c)
33:                else
34:                  continue
35:                end if
36:              end for
37:              if length(localQueue)  $\leq$  threshold then
38:                Merged localQueue into globalQueue
39:                Merged localVisited into globalVisited
40:              end if
41:            end while
42:          end for
43:        failure
44:      end procedure

```

platform [27] in order to obtain realistic and real-time day ahead prices. The services returns the hourly day ahead prices for the specified date.

Air Density Service: requires the temperature, air pressure, and dew point in order to calculate the air density. This service implements the Herman Wobus vapour pressure polynomial to calculate the saturation vapour pressure at the dew point temperature, which can then be used

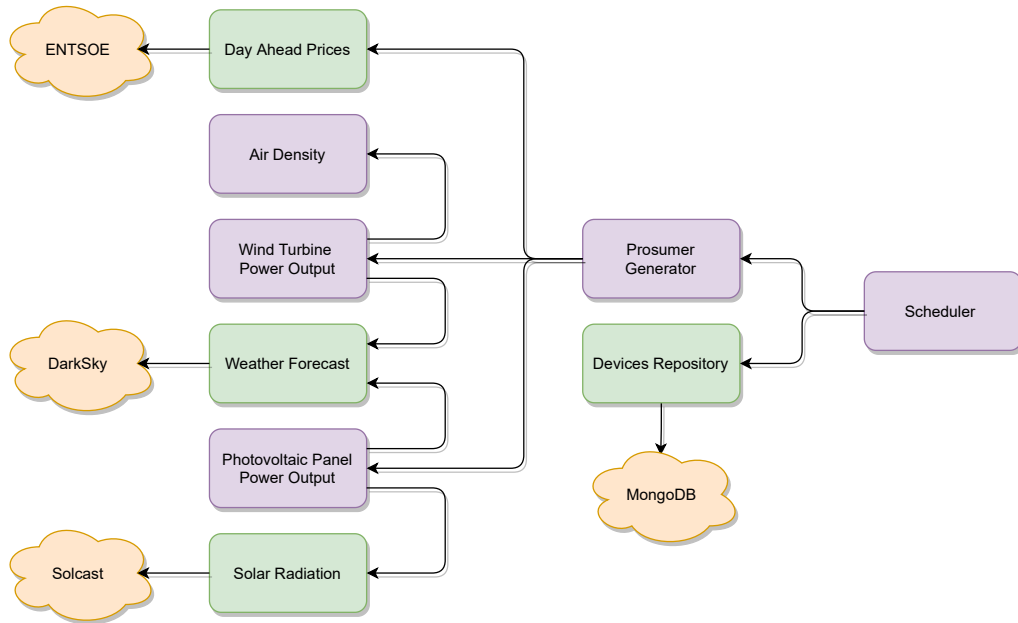


Fig. 2: Microservice architecture. The green coloured shapes represent services that can easily be exchanged to support a different data source. The purple shapes are the core services. The orange shapes are external services. The arrows represent the interdependencies.

to calculate the pressure of dry air. Finally, the service returns the air density for the given parameters.

Wind Turbine Power Output Service: requires the turbine's blade radius, wind speed, air density, and power coefficient to calculate the turbine's power output. This service also requires the cut-in and cut-out wind speed of the turbine to be provided. Based on these parameters, the turbine power output is given.

Weather Forecast Service: requires the latitude and longitude of the location for which the weather forecast should be obtained. This service returns the temperature, dew point, humidity, air pressure, and wind speed. The weather service used in this work is DarkSky [28] as it provides high temporal resolution forecasts for our desired location.

Photovoltaic Panel Power Output Service: requires the direct normal solar irradiance, temperature, area of the panel, and the efficiency of the panel. The service assumes that a maximum power point tracker is used. Based on the parameters provided to the service, the expected power output is returned.

Solar Radiation Service: requires the latitude and longitude of the location for which the solar radiation should be collected. The solar irradiance data is provided by the cloud API from Solcast [29]. This service returns the direct normal irradiance, diffused horizontal irradiance, and global horizontal irradiance.

Prosumer Generator Service: requires the location (latitude and longitude) for which the data for simulated prosumers should be generated, as well as the specifications of


the available PV and/or wind turbine resources. This service generates a list of energy sources which includes prosumers representing the neighbouring buildings connected to the smart grid. Each prosumer is able to deliver a certain amount of energy for a certain cost during a certain period of time. Added to this list of prosumers are the available renewable resources of the building in question, as well as an energy source representing the (traditional) grid in case the energy provided by the renewables and prosumers is insufficient to satisfy the demand.

Device Repository Service: returns the list of devices which should be scheduled, their power consumption in different states, as well as the policy for each device. The devices, their power states and their policies are stored in the document-oriented database MongoDB.


Scheduler Service: requires a list of energy sources (prosumers), a list of devices to schedule, the initial charge of the battery, as well as the algorithm that should be used. The modularity allows for selecting among versions of the uniform grid-search problem algorithm. The service returns the day ahead schedule of the devices with a 15 minute granularity, this predetermines the power state that each of the devices should be in for every 15 minute interval. Some meta-data on the performance of the algorithm is also included in the data that is returned by this service.

IV. EVALUATION

To evaluate the quality and possibilities of the proposed architecture and scheduling approach, we consider offices.



(a) Histogram of power data



(b) Clustering of power data

Fig. 3: Display device data

The specific setup consists of individual offices and a shared kitchen area, each with numerous devices. Smart plugs are used to collect power data from the devices situated in the offices and the kitchen area. This data is used to define the devices' profiles required by the scheduler. The renewables energy sources and energy storage are modelled after commercially available products. Combined with real-world weather data, the power output of the renewables is computed. The device profiles and energy sources are then used as input for the scheduler. Three different cases are considered; for each case we evaluate it with and without the availability of energy storage. For each generated schedule, the economic cost, the energy demand profile, and the battery charge are analysed.

A. Devices

In this work, 8 different devices are scheduled: a coffee machine, a fridge, a laptop, a microwave, a printer, a display, a thin client, and a battery based energy storage device. Two power states for each controllable device are considered. These power states are defined on the basis of the historical power data of the devices. The historical data was collected over a period of eight months, with measurements taken every 10 seconds. The K-means clustering algorithm [30] is applied to specify the boundary between different power states based on the power values and the density of the points.

As an example, Figure 3a shows the historical power data and its histogram for the display device. As is seen from the histogram, as well as examined by the K-means inertia plot and elbow method [31], the optimum number of clusters for the data set is two. Figure 3b shows the clustering result highlighting the centroids of the clusters. The boundary between the two clusters is placed between the two centroids. Table I shows the cluster centroids for different devices with

K-means clustering algorithm. S0 and S1 in this table, stand for two different power states.

Each device needs to have a policy assigned to it in order for the scheduler to understand the constraints under which to schedule it. An overview of each device and the assigned policy is shown in Table II. The coffee machine and the fridge require the *repeat* policy, as these devices need to be repeatedly turned on to maintain their temperature. The laptop uses the *total* policy as it needs to be charged for a total amount of time. The microwave follows a predefined *pattern* based on the usage of this device by the users; it is used more during lunch time. The printer has multiple printing jobs that need to be performed, therefore it has a *multiple* policy. The display and its associated thin client are only on during office hours, and therefore require a *strict* policy. From the perspective of the scheduler, the battery energy storage is just another device that can be scheduled, though this device does not only consume energy, but can also return energy at a later point in time. The energy storage is assigned the *battery* policy.

The total daily energy demand for these devices is approximately 15.6 kWh. The majority of this demand originates from the coffee machine, the printer, and the microwave. The coffee machine needs to repeatedly reheat water to maintain the required temperature, the printer has numerous large printing jobs that need to be completed, and the microwave is used heavily during lunch breaks.

B. Energy Sources and Storage

For local energy generation, the building has a single micro wind turbine, and a photovoltaic panel installation. The building also has battery storage. The geographical location of the building affects the energy sources, as the renewable energy generation is influenced by the weather, and the prosumers are constructed based on the day-ahead energy of the energy market. We choose as location Stuttgart, Germany.

The wind turbine is modelled after a Sonkyo Energy Windspot 1.5, for which the technical specifications are shown in Table III. The power coefficient is derived from the specifications as the manufacturer does not provide it explicitly. It should be noted that typically the specifications are optimistic, and that in practice the power coefficient will be lower. However, for the purpose of this evaluation the differences between theoretical and practical performance are not of major importance. The cost of locally generated wind energy is fixed at 0.08 €/kWh.

TABLE I: Centroids of the power states for different devices

Device	Power state S0 (watt)	Power state S1 (watt)
Coffee Machine	22.48	2555.84
Fridge	0.05	50.8
Laptop	14.12	76.61
Microwave	4.3	1567.71
Printer	11.13	654.59
Display	0.63	36.36
Thin Client	0.07	10.52

TABLE II: Device Policies

Device	Policy
Coffee Machine	REPEAT
Fridge	REPEAT
Laptop	TOTAL
Microwave	PATTERN
Printer	MULTIPLE
Display	STRICT
Thin Client	STRICT
Battery Energy Storage	BATTERY

For the photovoltaic panels, it is assumed that the building is equipped with two-axis maximum power point tracking solar panel installation. The specifications of the panels, as shown in Table IV, are based on the HiTech Solar 250Wp Black 60 cells panel. The installation consists of 6 panels, with a total of 1.5 kW peak power, and an area of 9.9 m². The cost is set to 0.06 €/kWh.

As for the battery-based energy storage, the assumption is that the building has a single battery storage installation. The assumed energy storage is based on AlphaESS SMILE3. The technical specifications of the battery energy storage are given in Table V.

Both the prosumers and the grid are modelled using the day ahead prices of the German electricity market, as published on the ENTSO-E Transparency Platform. To model the grid, the hourly cost of energy is to be equal to the hourly day-ahead prices as obtained from ENTSO-E. For the purpose of this evaluation, we assume that the grid is always able to satisfy the energy demand of the building, in case the renewables and prosumers are not able to. The obtained grid prices are normalised between €0.40/kWh and €0.60/kWh. The prosumers are also modelled based on the hourly day-ahead prices. It is assumed there are 10 different prosumers available at all time. To introduce some variability, the hourly energy cost for each prosumer is sampled from Equation 9, where $\mu = \frac{\text{grid_price}}{1.5}$ and $\sigma = 0.025$. The available energy for each prosumer is sampled from Equation 8, where $a = 0$ and $b = 1$.

C. Results: Economic

Three distinct cases are considered in order to evaluate the economic benefits of scheduling devices with and without energy storage in the presence of locally generated renewable energy. Each case has a unique price signal, modelled on the day ahead prices of the German electricity market. Furthermore, each case also has a unique renewable energy profile which is dependent on the weather forecast for that

TABLE III: Windspot 1.5 - Specifications

Rotor Swept Area	12.88 m ²
Rated Power	1.5 kW
Rated Speed	12 m/s
Cut In Speed	3 m/s
Survival Speed	60 m/s
Power Coefficient	0.11

TABLE IV: HiTech Solar - Specifications

Dimensions	1x1.65 m
Panel Area	1.65 m ²
Cells	60
Technology	Monocrystalline
Maximum Power	250 Wp
Panel Efficiency	15.30 %

TABLE V: SMILE3 - Specifications

Capacity	2.8 kWh
Charging/Discharging Current	60 A
Charging/Discharging Power	3000 W
Depth of Discharge	95 %

day. The renewable energy profile and the price signal for Case A (February 5th, 2022) are shown in Figures 4a and 4b, respectively. Case A is a cloudy and extremely windy day, especially after 14:00. The price signal for case A shows that the price is relatively low until around 11:00. The prosumer signal is the average price signal of the 10 available prosumers. Case B (February 6th, 2022), as shown in Figures 4c and 4d, is a windy day with some PV generation throughout the day. And finally, Case C (February 8th, 2022) is a sunny day with wind speeds below the cut-in wind speed of the turbine, as can be seen in Figure 4e. The price signal is low for the majority of the day, until approximately 16:00, as shown in Figure 4f.

The energy sources and their prices, as well as the list of devices, form the input of the scheduler. In total, six schedules are generated. Two schedules for each of the three cases, one with energy storage, and one without energy storage. The values of energy sources do not change for each individual case, the only change is the inclusion and exclusion of the energy storage. The scheduling horizon is set to 24 hours, the length of each scheduling period is set to 15 minutes, and the energy storage is empty at the start of the day.

Figure 5a, 5b, and 5c show the cumulative charge of the energy storage device (battery), the cumulative cost, and the cost per period, for case A. The cumulative charge is only applicable to the schedule with energy storage. The remaining two figures visualise the results for both the schedule with energy storage, and the schedule without energy storage. It is apparent from Figure 5a that the optimal schedule, i.e. the schedule with the lowest cost, does not charge or discharge the battery. The reason for this behaviour is the availability of cheap locally generated renewable energy (wind) throughout the day. The final cumulative cost is identical for both schedules, though the scheduling of devices is slightly different as can be seen in Figure 5c, where the cost per period differs slightly from 18:00 onwards. This difference can be explained by the parallel nature of the algorithm used by the scheduler, it is possible that the search space is explored in a different order, though this clearly does not affect the goal of minimizing total cost. The peaks in the cost per period graph are caused by the fridge and coffee machine repeatedly turning on and off with a certain periodicity.

The schedules for case B clearly exhibit different behaviour

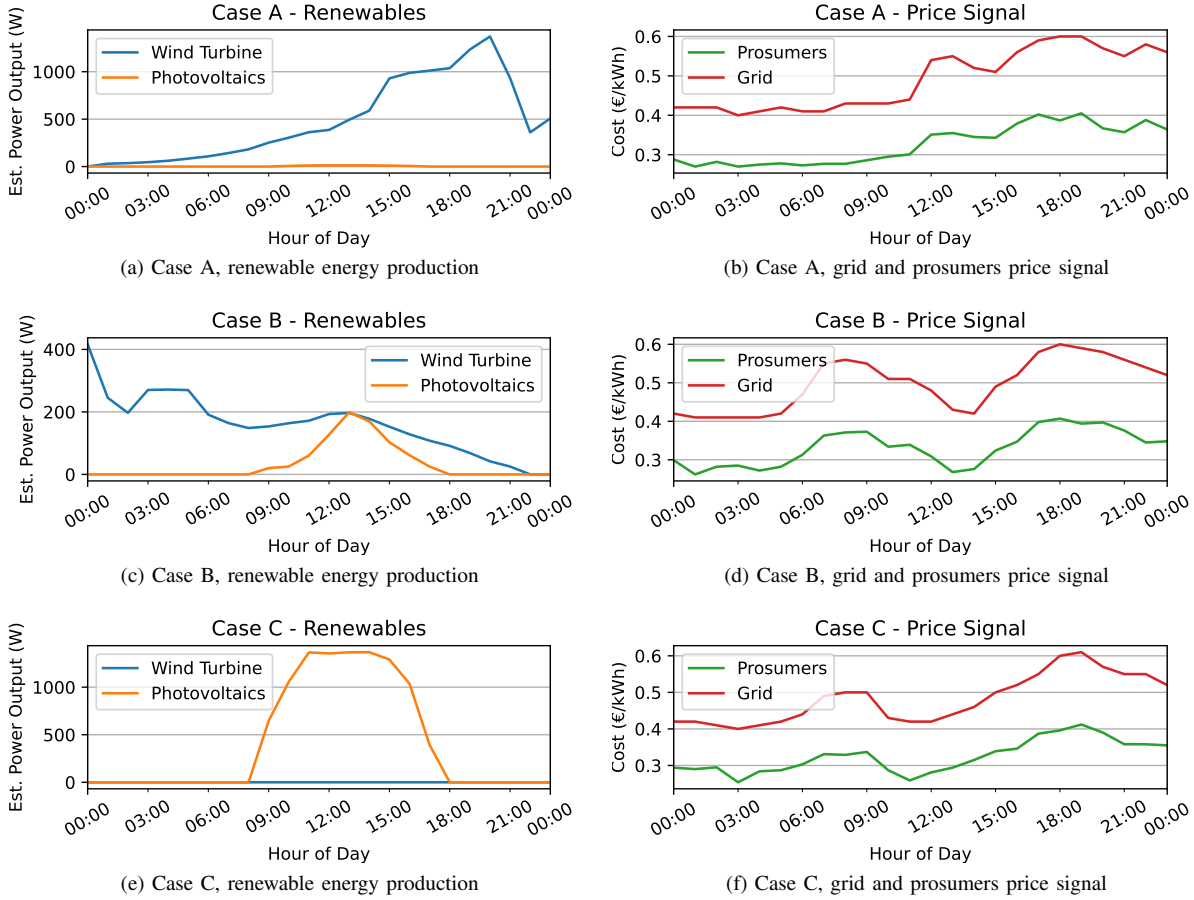


Fig. 4: The renewable energy production and price signals for each of the three considered cases

when compared to case A. Figure 5d demonstrates this, the battery is charged in the beginning of the day, and is not discharged until the end of the day. When looking at the availability of locally generated renewable energy, and the price signal, it becomes clear why case B behaves in such way. Until approximately 19:00 there is sufficient renewable energy available to meet most of the demand. Simultaneously, the price signal remains high after 17:00. Therefore, the battery is charged in the early hours of the day, as cheap locally generated renewable energy is available, and is discharged after 19:00, as the renewable energy generation is not sufficient, and the price signal is high.

For case C, as visible in Figure 5g, the charge of the battery fluctuates between 00:00 and 09:00, after which it is fully charged, and remains fully charged until 19:00. From Figure 5h one notices that the cumulative costs of the schedule which includes the battery are significantly lower. These savings are obtained after 19:00, when the price signal is high and there is no renewable energy available, as visible in Figure 5i.

An overview of the economic costs associated to each of the 6 schedules is shown in Table VI. The price signal and availability of locally generated renewable energy greatly af-

fect the potential economic benefits of utilizing energy storage. Regardless of these factors, the scheduler is always able to find the optimal schedule, whether this includes utilizing energy storage or not.

D. Results: Performance

When optimizing energy it is important that the system is fast enough to be used in the building and that it does not consume more resources than it saves. To study this we perform the analysis on a Dell PowerEdge R7425 server with two AMD EPYC 7551 2.0GHz 32-core processors, 512GB of RAM, and read-optimized SSDs with a throughput of 560MB/s. The bottleneck of this test bed is the relatively low clock speed of the AMD EPYC 7551, which limits the

TABLE VI: Economic Benefits Overview

Case	Battery?	Cost (€)	Cost Difference
A	No	1.46	-
A	Yes	1.46	-
B	No	1.84	-
B	Yes	1.69	-€0.15 (8.15%)
C	No	2.12	-
C	Yes	1.64	-€0.48 (22.64%)

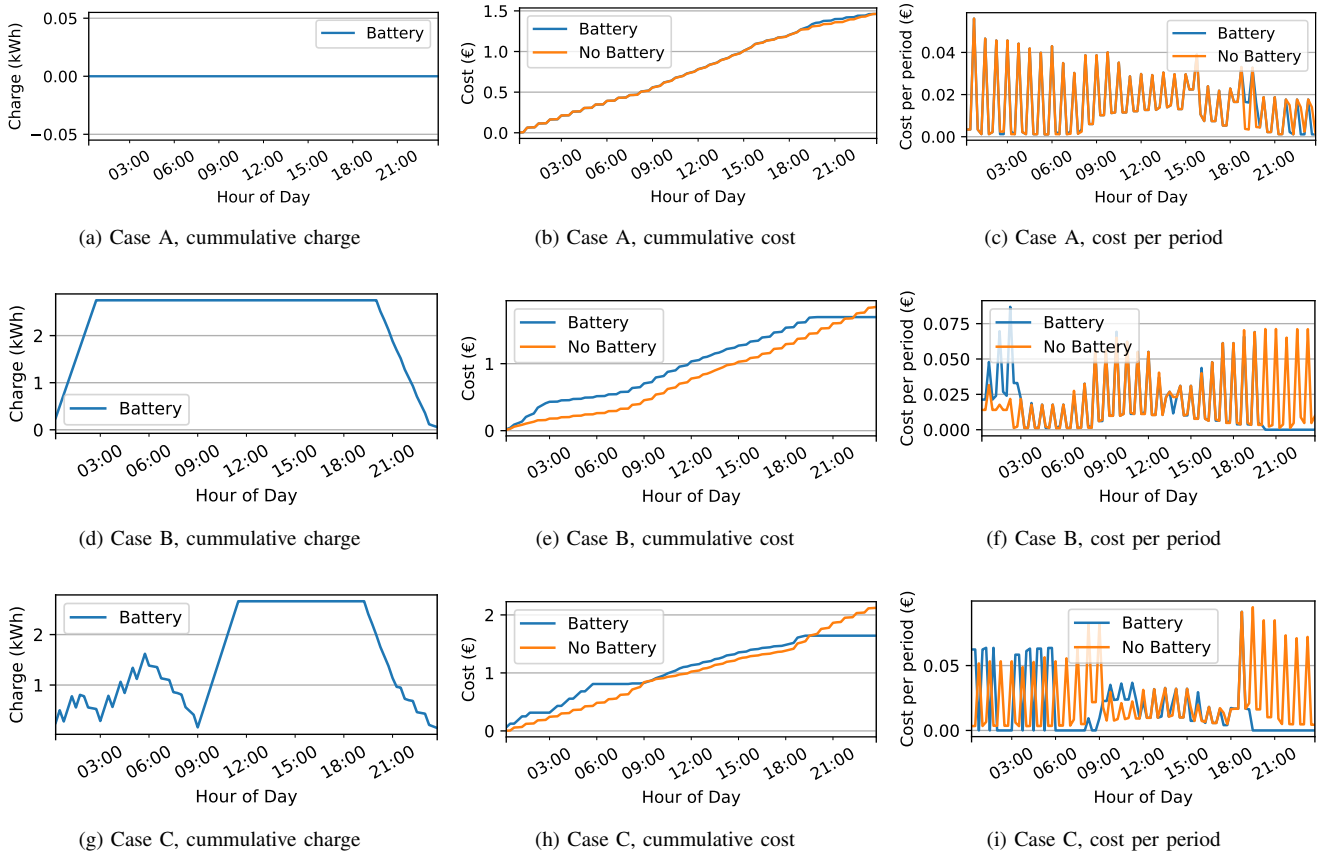


Fig. 5: The cumulative battery charge, cumulative cost, and cost per period associated with the generated schedules

number of operations per second when executing the uniform grid search.

To evaluate the performance of the algorithms, a set of test cases is generated by varying numerous parameters. These parameters include: the number of devices, scheduling horizon, number of threads, number of iterations, and algorithm variation. The number of devices ranges from 1 to 12. The scheduling horizon ranges from 2 hours to 24 hours, in 1 hour increments. The number of threads ranges from 1 to 16. The number of iterations is set to 10, and indicates how many times the same test case is run. And finally, the algorithm variations include the original algorithm from our previous work [11], a variation of the original algorithm that minimizes memory operations, and the parallel algorithm proposed in this work. The original algorithm was modified to support battery device policy and the multitude of microservices. The result is 27.600 distinct test cases.

Due to the high dimensionality of the resulting data, it is not possible to visualize all of the results in their entirety. Therefore, we focus on a number of specific cases which characterize the general performance of the approach. Figure 6a shows the mean performance when scheduling 8 devices with a varying scheduling horizon going from 2 to 24 hours. This case is selected because the scheduling

problem is reasonably complex. The time it takes to find the optimal schedule (scheduling duration) is measured in minutes. It is immediately clear that the memory-optimized and parallel algorithms perform considerably better than the original algorithm. The more complex the search space (the broader the scheduling horizon), the more pronounced the difference becomes. At the 24 hour scheduling horizon, the original algorithm (30 minutes) is 2.9 times slower compared to memory-optimized (10.2 minutes), and 4.7 times slower compared to parallel (6.2 minutes) with 4 threads.

Figure 6b shows the same results as the previous figure, except the original algorithm is excluded, and the results for the varying thread count parameters are added. This highlights the performance differences between number of threads. The memory-optimized and parallel (1 thread) algorithms have identical performance (10.2 minutes at 24:00 hours scheduling horizon), as is expected due to the fact that the parallel algorithm is derived from the memory-optimized one. Increasing the thread count from 1 to 2, improves the performance by a factor of 1.4 (7.2 minutes at 24:00). When increasing the thread count to 3 or 4, the performance improves by a factor of 1.6 (6.2 minutes at 24:00) for both cases when compared to the memory-optimized results. This is due to the implementation of the concurrent priority queue. As this implementation




Fig. 6: Performance of the variations of the uniform grid search algorithm

requires synchronization when adding and removing nodes to the queue, acquiring the lock becomes the bottleneck.

In Figure 6c, the case of 4 devices and a 06:00 hour scheduling horizon is considered. This case is selected because complexity-wise it is the opposite of the previous one: the search space is small. When the search space is small, the overhead of synchronizing 4 threads becomes larger than the time it takes for the single threaded variations of the algorithm to find the solution. This is clearly visible at 03:00, where the parallel algorithms (with 4 threads) requires 15 milliseconds to complete, while parallel (with 1 thread) requires only 2 milliseconds. In practice, the search space is typically more complex, and a difference of 13 milliseconds is negligible in the broader perspective of the system's execution time.

In the three preceding cases the number of devices is fixed. Figure 6d (logarithmic y-axis) demonstrates the effect of increasing the device count, and therefore increasing the search space. The complexity does not change significantly when adding the 4th and 5th devices, this is due to the policy that is assigned to these devices. Some policies (such as *pattern*, *strict*, and *idle*) only minimally increase the search space.

When ignoring the 4th and 5th devices, as the number of devices increases, the scheduling duration increases exponentially.

To clearly illustrate the bottleneck problem caused by the locking mechanism of the queue, the number of devices is reduced to 6, the thread is set to range from 1 to 16, and the scheduling horizon is fixed at 24:00 hours. Figure 7a visualizes the mean scheduling duration in milliseconds with varying thread counts. The original algorithm (red dashed line) is included for reference. For this particular scheduling problem, it is clear that 3 threads provides the fastest scheduling duration. As the thread count grows beyond 3, the synchronization mechanism in the queue becomes the bottleneck. In addition, Figure 7b shows the standard deviation for each thread count obtained after 10 iterations.

V. DISCUSSION, LIMITATIONS & OUTLOOK

In our previous work we demonstrated the concept of scheduling devices against a price signal in order to decrease energy cost in an office building [11]. The present work expands that concept and improves upon it in multiple ways.

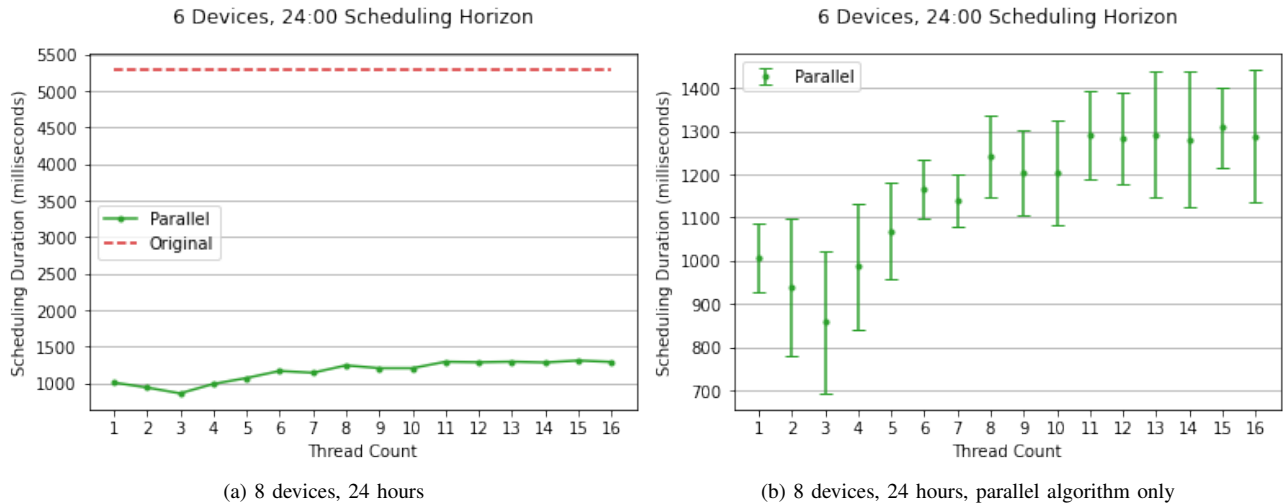


Fig. 7: Thread counts vs. performance

First of all, by adding real time modelling of renewable energy generation, as well as the inclusion of energy storage. Second, the modelling of energy sources is supported by a loosely coupled, modular microservice architecture which enables the system to adapt to different cloud service providers. Furthermore, real world data is used for the scheduled devices, and for the prediction of renewable energy generation. And finally, significant performance enhancements are made by introducing a parallel uniform cost-search algorithm.

From the economic perspective, the inclusion of energy storage coupled with device scheduling can potentially lead to significant cost savings. However, as shown by the presented results, this depends greatly on the price signal and on the availability of locally generated renewable energy. For example, when sufficient amounts of energy are generated throughout the day, the battery is not scheduled to charge, as charging is not free and the number of charge cycles of a battery is limited. Therefore, in the worst case, the economic benefit of adding energy storage is null on a daily basis. Nevertheless, in several cases conditions are beneficial for energy storage, for example when price signals fluctuate and the locally generated renewable energy is not sufficient to meet the demand, then there are significant cost savings to be made. The results show cost savings of up to 8.15% per day, going as high as 22.64% under the more favourable conditions. While the evaluation focuses on the economic perspective, the algorithm is agnostic with regards to the costs. In other words, instead of providing a list of energy sources and their cost expressed in euros per kWh, one could provide instead the cost expressed as grams of CO₂ equivalent emissions per power unit (gCO₂eq per kWh) to improve the sustainability rather than the economic costs.

From the performance perspective, the optimization and parallelization of the uniform cost-search algorithm result in significant performance gains. The original algorithm requires

up to 30 minutes to schedule 8 devices with a scheduling horizon of 24 hours. The optimizations alone reduce this time to 10 minutes. Further gains are made by parallelizing the uniform-cost search algorithm, completing the scheduling in 6.2 minutes with 4 threads. It is also possible to spread the load over multiple machines using distributed agents, where each agent solves a part of the problem. However, accessing the shared resources requires synchronization, and therefore needs to be executed sequentially. This sequentially executed section of the algorithm becomes the bottleneck. The bottleneck is a manifestation of Amdahl's law [32], which in the context of parallel programming describes that the maximum possible speed-up is disproportionately limited by the sequential section of the application.

Finding the optimal solution to a scheduling problem is a time consuming task. In the present work, the scheduler operates on 15 minute intervals with a maximum scheduling horizon of 24 hours. Decreasing the size of the intervals, expanding the scheduling horizon, and increasing the number of scheduled devices all exponentially increase the complexity of the problem. To counteract these limitations, the scheduling problem can be divided into smaller problems, which are solved individually and later combined into one schedule. While these local schedules are optimal, the global schedule is likely not to be.

The energy consumption associated with the schedule generation is not included in the economic evaluation. As is shown in the performance evaluation, the time it takes to solve the scheduling problems is in the order of several minutes. The system only needs to be available during this time period. The execution of the schedule can be performed by a simple embedded device. As per the performance evaluation, the scheduling of 8 devices requires at most 6.2 minutes when using 4 parallel threads. The 32-core AMD EPYC 7551 CPU is responsible for the majority of the system's power consump-

tion. Assuming each of the 4 parallel threads utilises 100% of their allocated CPU core, the overall CPU utilization is 12.5%. According to the SPECpower_ssj2008 benchmarks [33], this would equate to a power consumption of 100W. Therefore, generating the schedule requires 0.01 kWh. We can conclude that the cost of the energy consumption of the system itself is insignificant with respect to the overall economic savings.

From a system architectural point of view, we proposed a microservice architecture that encapsulates the three critical components in several microservices. This enables the system to leverage the real time data provided by weather services and the ENTSO-E transparency platform to generate realistic price signals on the fly. Furthermore, a performance-oriented version of the uniform cost-search algorithm is proposed, which parallelizes the search for the optimal schedule. The performance evaluation shows that this parallel algorithm decreases the time to find the optimal schedule by a factor of 4.7 compared to the previous approach. And finally, this work evaluates the economic benefits of including local energy storage in the scheduling problem, demonstrating that the inclusion of energy storage can further reduce costs by up to 22.64%, depending on the price signal and availability of renewable energy.

ACKNOWLEDGEMENTS

We thank Alexander Lazovik, Viktoriya Degeler, and Andrea Pagani for discussion about the approach and architecture presented here as an extension of our joint previous work [11].

REFERENCES

- [1] A. Ari, N. Arregui, S. Black, O. Celasun, D. Iakova, A. Mineshima, V. Mylonas, I. Parry, I. Teodoru, and K. Zhunussova, "Surging energy prices in europe in the aftermath of the war: How to support the vulnerable and speed up the transition away from fossil fuels," 2022.
- [2] P. E. Teixeira Martins, M. Oleskovicz, and A. L. da Silva Pessoa, "A survey on smart grids: concerns, advances, and trends," in *2019 IEEE PES Innovative Smart Grid Technologies Conference - Latin America (ISGT Latin America)*, 2019, pp. 1–6.
- [3] I. Lee and K. Lee, "The Internet of Things (IoT): Applications, investments, and challenges for enterprises," *Business Horizons*, vol. 58, no. 4, pp. 431–440, 2015.
- [4] H. Chen, P. Chou, S. Duri, H. Lei, and J. Reason, "The design and implementation of a smart building control system," in *2009 IEEE International Conference on e-Business Engineering*, 2009, pp. 255–262.
- [5] Q. Wang, M. Liu, and R. Jain, "Dynamic pricing of power in smart-grid networks," in *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, 2012, pp. 1099–1104.
- [6] G. A. Pagani and M. Aiello, "Generating realistic dynamic prices and services for the smart grid," *IEEE Systems Journal*, vol. 9, no. 1, pp. 191–198, 2015.
- [7] A. Gholian, H. Mohsenian-Rad, and Y. Hua, "Optimal industrial load control in smart grid," *IEEE Transactions on Smart Grid*, vol. 7, no. 5, pp. 2305–2316, 2016.
- [8] L. Fiorini and M. Aiello, "Energy management for user's thermal and power needs: A survey," *Energy Reports*, vol. 5, pp. 1048–1076, 2019.
- [9] S. J. Darby, "Load management at home: advantages and drawbacks of some 'active demand side' options," *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, vol. 227, no. 1, pp. 9–17, 2013.
- [10] F. Nizamic, T. A. Nguyen, A. Lazovik, and M. Aiello, "Greenmind - an architecture and realization for energy smart buildings," in *Proceedings of the 2014 conference ICT for Sustainability*. The Netherlands: Atlantis Press, Aug. 2014, pp. 20–29.
- [11] I. Georgievski, V. Degeler, G. A. Pagani, T. A. Nguyen, A. Lazovik, and M. Aiello, "Optimizing energy costs for offices connected to the smart grid," *IEEE Transactions on Smart Grid*, vol. 3, no. 4, pp. 2273–2285, Dec.
- [12] N. G. Paterakis, O. Erdinç, A. G. Bakirtzis, and J. P. S. Catalão, "Optimal household appliances scheduling under day-ahead pricing and load-shaping demand response strategies," *IEEE Transactions on Industrial Informatics*, vol. 11, no. 6, pp. 1509–1519, 2015.
- [13] A. C. Duman, H. S. Erden, Ömer Gönül, and Önder Güler, "A home energy management system with an integrated smart thermostat for demand response in smart grids," *Sustainable Cities and Society*, vol. 65, p. 102639, 2021.
- [14] S. Nan, M. Zhou, and G. Li, "Optimal residential community demand response scheduling in smart grid," *Applied Energy*, vol. 210, pp. 1280–1289, 2018.
- [15] A. Amer, K. Shaban, A. Gaouda, and A. Massoud, "Home energy management system embedded with a multi-objective demand response optimization model to benefit customers and operators," *Energies*, vol. 14, no. 2, 2021.
- [16] R. Lu, S. H. Hong, and M. Yu, "Demand response for home energy management using reinforcement learning and artificial neural network," *IEEE Transactions on Smart Grid*, vol. 10, no. 6, pp. 6629–6639, 2019.
- [17] H. Li, Z. Wan, and H. He, "Real-time residential demand response," *IEEE Transactions on Smart Grid*, vol. 11, no. 5, pp. 4144–4154, 2020.
- [18] F. Fioretto, W. Yeoh, and E. Pontelli, "A multiagent system approach to scheduling devices in smart homes," in *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, ser. AAMAS '17. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems, 2017, p. 981–989.
- [19] Y. Xia, K. Ahmed, and B. Williams, "Wind turbine power coefficient analysis of a new maximum power point tracking technique," *IEEE Transactions On Industrial Electronics*, vol. 60, no. 3, pp. 1122–1132, Mar. 2013.
- [20] The Royal Academy of Engineering, "Wind turbine power calculations," *Mechanical and Electrical Engineering / Power Industry*, 2010.
- [21] C. Gerald and P. Wheatley, *Applied Numerical Analysis*, 7th ed. United Kingdom: Pearson College Div, 2003.
- [22] S. X. Chen, H. B. Gooi, and M. Q. Wang, "Sizing of energy storage for microgrids," *IEEE Transactions on Smart Grid*, vol. 3, no. 1, pp. 142–151, 2012.
- [23] V. Degeler, "Dynamic rule-based reasoning in smart environments," Ph.D. dissertation, University of Groningen, 2014.
- [24] S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, 3rd ed. United States: Prentice Hall, 2010.
- [25] K. Bakshi, "Microservices-based software architecture and approaches," in *2017 IEEE Aerospace Conference*, 2017, pp. 1–8.
- [26] M. Kalske, N. Mäkitalo, and T. Mikkonen, "Challenges when moving from monolith to microservice architecture," in *Current Trends in Web Engineering*, I. Garrigós and M. Wimmer, Eds. Cham: Springer International Publishing, 2018, pp. 32–47.
- [27] ENTSO-E, "Data and information on wholesale energy generation, transmission, and consumption." [Online]. Available: <https://transparency.entsoe.eu/>
- [28] DarkSky, "Hyperlocal weather information." [Online]. Available: <https://darksky.net>
- [29] Solcast, "Global solar irradiance data and pv system power output data." [Online]. Available: <https://solcast.com/>
- [30] J. MacQueen *et al.*, "Some methods for classification and analysis of multivariate observations," in *Proceedings of the fifth Berkeley symposium on mathematical statistics and probability*, vol. 1, no. 14. Oakland, CA, USA, 1967, pp. 281–297.
- [31] A. Ng, "Clustering with the k-means algorithm," *Machine Learning*, 2012.
- [32] G. M. Amdahl, "Computer architecture and amdahl's law," *IEEE Solid-State Circuits Society Newsletter*, vol. 12, no. 3, pp. 4–9, 2007.
- [33] SPEC Power Benchmark, "SPECpower_ssj@ 2008." [Online]. Available: https://www.spec.org/power_ssj2008/