A Real-World District Community Platform as a Cyber-Physical-Social Infrastructure Systems in the Energy Domain

Johannes Galenzowski
Simon Waczowicz
Karlsruhe Institute of Technology
Germany
johannes.galenzowski@kit.edu
simon.waczowicz@kit.edu

Stefan Meisenbacher
Ralf Mikut
Karlsruhe Institute of Technology
Germany
stefan.meisenbacher@kit.edu
ralf.mikut@kit.edu

Veit Hagenmeyer
Karlsruhe Institute of Technology
Germany
veit.hagenmeyer@kit.edu

ABSTRACT

The transition of urban energy supply into a regenerative but also digitalized smart ecosystem leads to an unprecedented necessity to rethink the energy supply infrastructure in city districts. Recent research strongly suggests that it is highly beneficial to consider platform solutions holistically as a cyber-physical-social infrastructure system. However, there are few detailed descriptions of the practical implementation of such a system. We aim to close this gap by presenting the findings from the real-world implementation of a district community platform as a cyber-physical-social infrastructure system. In doing so, we address the energy systems, the relevant (time series) data and services as well as the stakeholders involved.

CCS CONCEPTS

• **Hardware → Smart grid**; **Energy metering**; **Sensor devices and platforms**.

KEYWORDS

city district, cyber-physical-social infrastructure system, urban energy transition, energy management, district community platform

ACM Reference Format:
Johannes Galenzowski, Simon Waczowicz, Stefan Meisenbacher, Ralf Mikut, and Veit Hagenmeyer. 2023. A Real-World District Community Platform as a Cyber-Physical- Social Infrastructure Systems in the Energy Domain. In *The 10th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '23), November 15–16, 2023, Istanbul, Turkey*. ACM, New York, NY, USA, 8 pages. https://doi.org/10.1145/3600100.3626347

1 INTRODUCTION

Reducing greenhouse gas emissions to tackle climate change is one of the most significant challenges for people and governments around the world [17]. Almost 70 % of the global population is expected to live in towns and cities by 2050 [11]. Further, 37 % of the world's greenhouse gas emissions are related to buildings [19] and are directly linked to energy consumption. Worldwide, the residential sector contributes 20 %, the service sector 14 % and the transport sector 35 % to the primary energy consumption [10]. These sectors all come together in city districts, so districts play a crucial role in transforming the energy sector.

Terminology. The present work focuses on how an interplay of physical, cyber, and social systems in city districts can add value for the energy system. The interaction of networks of physical and computational components, so-called cyber-physical systems (CPSs), is well known in literature [1]. The main application classes of CPSs are smart healthcare, smart manufacturing, and smart cities [1]. Within smart cities or smart city districts, human behavior plays a crucial role, which is why Cassandras [3] refers to them as cyber-physical-social systems (CPSSs). In the context of energy, the CPS is often referred to as "Smart Grid" [21].

Because we focus on the energy system as an infrastructure, our CPSS is a cyber-physical-social infrastructure system (CPSIS). However, the distinction between CPSS and CPSIS is usually not made in the literature. For example, works like Zhang et al. [23] that highlight the importance of the energy system in the domain of CPSS also refer to a CPSS and not a CPSIS. Other works like [4, 12, 20] describe a similar system to the one we present but call it a renewable energy community (REC). Thus, the district community platform (DCP) we present can be understood as a CPSIS or a REC, as well as a CPSS depending on the domain and research community (see Figure 1). However, it is crucial to distinguish a DCP from an information and communications technology platform (ICTP), since an ICTP is only the cyber system that is part of the DCP. This is in line with the definition of Aguida et al. [1], that a "Smart City" is a CPSS that utilizes information and communications technology (ICT) to provide services to achieve social goals. Gong et al. [8] point out that Internet of Things (IoT) is an enabler to control devices of a CPSIS and is, besides cloud services, a crucial part of the ICTP.

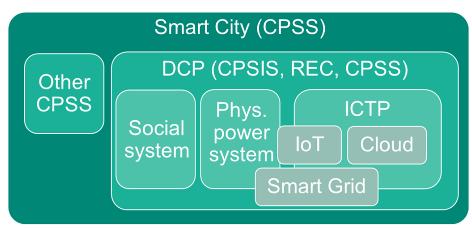

Figure 1: Relation of the terminology cyber-physical-social infrastructure system (CPSIS), cyber-physical-social system (CPSS), district community platform (DCP), renewable energy community (REC), information and communications technology platform (ICTP) and Internet of Things (IoT)

BuildSys '23, November 15–16, 2023, Istanbul, Turkey

ACM ISBN 979-8-4007-0230-3/23/11.
https://doi.org/10.1145/3600100.3626347

Related work. As stated in the terminology above, different perspectives on DCPs exist, which is also apparent in related work. Gong et al. [8] propose a framework with a CPSS perspective for modeling smart buildings, integrating them into a large-scale real-world virtual power plants (VPPs) with 130 buildings in Shanghai. The grouping of the buildings into blocks is to a large extent comparable to a district approach. However, the authors focus more on describing the framework and theoretical concepts. Details on the real-world implementation regarding stakeholders, the data used or a description of the services of their ICTP is missing.

Other works like Cavallaro et al. [4] extensively focus on the social benefits of RECs. However, the considered example of a REC comprises only one photovoltaic (PV) plant and 20 families. A description of a real-world large-scale district or a detailed discussion of the real-world implementation is missing.

In comparison, other works like Vernay et al. [20] or Krug et al. [12] focus on legal and regulatory aspects and business models. They describe interesting policy concepts and market perspectives of RECs. However, they lack to describe the technical aspects or detailed presentation of real-world DCP examples.

Reynolds et al. [16] propose the "Computational Urban Sustainability Platform" as a CPSIS that is quite similar to our DCP. However, at the time the authors published the paper, the platform was still under development. The concept outlines of forecasts and optimization are given, but no concrete service functions or required time series from the productive usage are described. Moreover, even though the authors also implicitly consider social aspects, they did not mention a CPSS but rather framed it as an ICTP.

Lastly, meta studies like Wang et al. [21] show that CPSS also become increasingly relevant for research coming from the "Smart Grid" community. The authors discuss aspects of the distribution grid as well as the increasing electric vehicle (EV) integration that are crucial for DCPs. However, as a meta study and due to its grid focus, it lacks the aspect of districts and real-world implementation.

Main contributions. The related work analyzed above reveals a lack of real-world projects that test CPSISs in the energy domain in city districts and discuss results in the multidimensional context of a CPSIS. Therefore, the main contributions of this paper are the following:

(1) In-depth presentation of a real-world productive DCP with a specialized focus on technical devices, system boundaries, and hierarchy in city districts.
(2) Detailed presentation on how the multidimensionality of a CPSIS is reflected in the practice of a DCP.
(3) Discussion on required (time series) data and services for a large-scale CPSIS in the energy sector.
(4) Illustration of the widespread involvement of social factors through the large number of stakeholders.

Paper structure. The remainder of the paper is organized as follows: section 2 presents the real-world district and the physical assets of our DCP. In section 3, we discuss social aspects and their relation to the ICT and the data basis. The following section 4 provides details on the services required by the DCP and the involved stakeholders. Section 5 concludes and gives an outlook for future work.

2 THE CITY DISTRICT: ENERGY DOMAIN AND PHYSICAL ASSETS OF THE DCP

To give insight into our DCP, the physical, cyber, and social components as well as the interfaces and connecting networks are outlined in the following. Participants in a DCP are connected via energy flows (electricity, district heating), material flows (gas, electric vehicles), and information flows (from physical to cyber systems and vice versa as well as between individual participants).

2.1 Structure of the district

Since our DCP is intended to provide services on the energy side, it is sensible to first describe the district at the highest level according to the requirements of the energy sector.

In the energy sector, a relevant point for zoning urban districts from a legal and business model perspective are the connection points from the private to the public grid [7], which are referred to as grid connection points (GCPs) in the following. GCPs can aggregate several buildings or just a single one (see Figure 2).

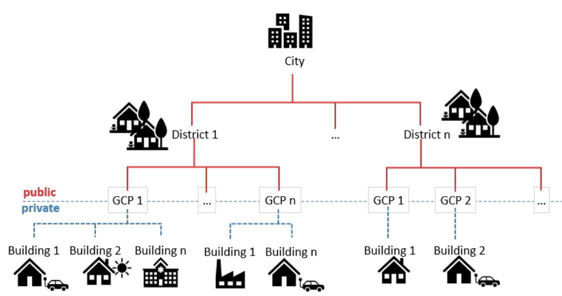

Figure 2: Concept of GCPs as boundary [7]

Our real-world district lies in the city of Karlsruhe (Germany) and consists of six GCPs. The district extends over roughly 460 m×270 m respectively 124 200 m^2. Figure 3 gives an overview of the energy systems in this district. The electrical systems vary greatly between

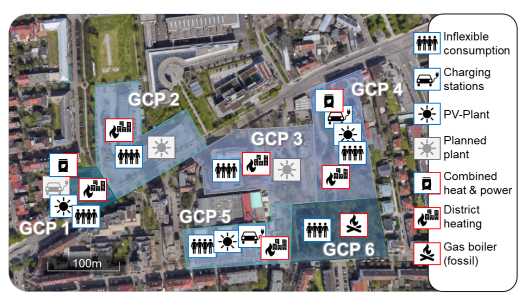

Figure 3: Geographical layout of the physical components of the DCP in the real-world district in Karlsruhe, Germany (Project: Smart East, https://smart-east-ka.de/)

the GCPs. In the heating sector, however, the systems are very similar: five of the six GCPs are supplied via district heating, and only one with natural gas. All of them have an inflexible electrical base load, although they show significant differences in the load curve due to the heterogeneous utilization pattern of the different buildings. Local generation consists of PV plants and combined heat and power (CHP) units. In terms of controllable power, charging infrastructures (CIs) are the easiest to control and the largest flexible plants. Therefore, further flexible plants such as CHP or heating, ventilation, and air conditioning (HVAC) were not considered in the first version of our DCP. If the flexible plants would also feed in, they would be called prosumer instead of consumer, while the overall concept remains unchanged.

Our goal is to be able to capture all energy flows into or out of the district in the form of time series, as a dynamic quantity, during online operation and for later processing. Therefore, in the following, it must be explained how the interface from the physical to the cyber system respectively information system can take place.

2.2 Relevant measuring points

To find the relevant places to deploy measurement devices, the energy flows in the system must be understood in more detail. We identified the power flows for a single GCP as shown in Figure 4.

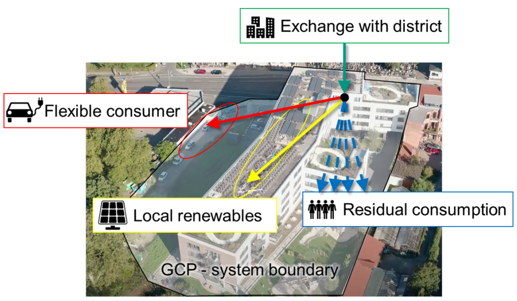

Figure 4: Energy flows of a GCP (arrows showing convention for positive flow direction)

For this example, we find over 30 electricity meters for residual consumers (blue, dotted lines). Of these, only four were found to be connected for remote readout. Upgrading all missing ones for remote readout would mean significant financial costs and pose problems for data privacy. However, we found that measuring individual residual consumers is unnecessary for the relevant DCP functionality (see subsection 3.2). Only the sum of the residual consumption $P_{\text{sum,res.cons.}}$ is required for each GCP. Consequently, the following time series are relevant:

- Exchange with district (P_{exch})
- Local renewable generation ($P_{\text{loc,gen}}$)
- Flexible consumer ($P_{\text{flex,cons}}$)
- Sum of residual consumption ($P_{\text{sum,res.cons.}}$)

Still, the sum of the residual consumption ($P_{\text{sum,res.cons.}}$) is usually also difficult to measure. For example, if local renewables or flexible consumers are connected to the same sub-distribution as the inflexible consumers. In this case, there is simply no busbar, where current sensors could be placed that only includes residual consumption. Our solution is to circumvent the measurement by calculating the residual consumption mathematically from the other load profiles according to the following equation:

$$P_{\text{sum,res.cons.}} = P_{\text{exch.}} - P_{\text{loc.gen.}} - P_{\text{flex. cons.}} \tag{1}$$

Flexible consumers usually already have integrated measuring devices for their control and billing purposes. This further saves the need for installing new metering devices. In some cases, the local generation plants also have a sufficiently easy-to-implement interface for measurement data acquisition. In summary, in the ideal case, only a single new meter is needed instead of more than 35 for all individual producers and consumers. In the standard case, two measuring devices can be assumed if the flexible consumers record their measured values themselves and the residual consumption is calculated. Therefore, we recommend installing meters at the following points: Exchange with district (P_{exch}) and local renewable generation ($P_{\text{loc,gen}}$).

3 THE ALL-PERVADING SOCIAL ASPECT

It must be noted that a DCP involves social interaction at various dimensions. The social aspects are comprised of which stakeholder interacts with the CPS at which lifecycle phase for which use case (see Figure 5). In particular, this paper aims to illustrate how each technical implementation, service, and design decision can take into account aspects such as human constraints, privacy, productivity enhancement, and user convenience. DCPs aim to bring together

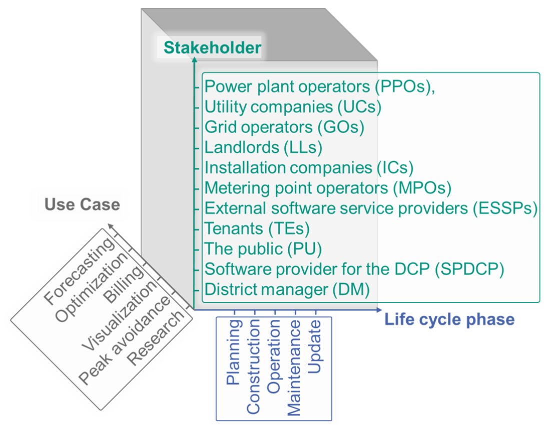

Figure 5: The linking of the social component of the DCP with all the technical aspects is multidimensional

climate protection, digitalization, business models and participation. Relevant stakeholders are: power plant operators (PPOs), utility companies (UCs), grid operators (GOs), landlords (LLs), installation companies (ICs), metering point operators (MPOs), external software service providers (ESSPs), tenants (TEs), the public (PU) as well as a software provider for the DCP (SPDCP) and a district manager (DM). A distinction must be made between different forms of stakeholder involvement. The DCP consists of mechanics (hardware), electrics, and software. Moreover, participation can take place at any life cycle phase: planning, construction, operation, maintenance, and update. Therefore, the human-centered aspect is a major focus when discussing the different parts of the DCP in the following.

3.1 IoT and Information flow as social challenge

As depicted in section 2, the district consists of six GCPs and stretches over an area of 460 m × 270 m. To connect all the facilities shown in Figure 3 to the ICTP of the DCP, a widespread communication network is needed.

As a limitation, it has to be said that the majority of the meters are installed in basements, where there is no mobile phone reception and therefore a wireless solution via a local radio network as well as LTE was ruled out from the outset. Another reason that played a role in the exclusion of LTE is that we also want to do research on our software architecture. We have therefore installed meters that are suitable for high-frequency measuring of the network states presented in subsection 3.2 and built a software architecture that allows direct access to significantly more than 2000 Modbus registers with measured values per meter at any time during operation. If one does not need to carry out any further configuration during operation and only a limited data set is to be delivered from the meter to an endpoint of a metering point operator in commercial-scaled operation, LTE can again become significantly more interesting.

Since they belong to different owners, the GCPs do not have data cabling that directly connects them. However, all GCPs are essentially connected to the Internet. Therefore, it is necessary to consider either constructing a new local network that directly connects all the energy plants or to fulfill the data exchange via the Internet. Due to the high costs, the temporal effort and the legal problems if one wanted to lay a private data line across public areas, the construction of a dedicated physical district network was ruled out. With much smaller effort, all energy-related systems were connected to the Internet of the individual tenants or landlords. The only serious disadvantage of using the Internet is the reduced reliability. As a result, no real-time control can be implemented via the DCP's cyber network without a local fallback option.

Several stakeholders and devices are involved in the data transmission from the energy systems to the cloud part of the platform, as shown in Figure 6. With Figure 6 we aim at describing general possible failure points. These possible failure points cannot be eliminated. But an awareness of them can help greatly in troubleshooting and highlight the importance of cooperation between the stakeholders.

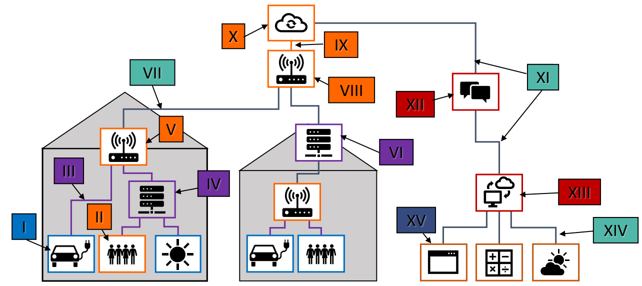

I. (PPO) Network adapter or web server of flexible consumer
II. (MPO) Meter of the MPO
III. (LL,TE) Local Ethernet cable
IV. (LL,TE, MPO) Local switch
V. (MPO) VPN gateway
VI. (LL, TE, UC) Local router on internet side of VPN gateway
VII. (UC) Internet building to metering-point operator
VIII. (MPO) VPN endpoint device (network card, software)
IX. (MPO) Network of metering-point operator
X. (MPO) Data collection and processing backend of MPO
XI. (UC, SPDCP, MPO) Internet link to communication broker
XII. (SPDCP,DM) Hardware/software of communication broker
XIII. (SPDCP, DM) Platform
XIV. (UC, ESSP, SPDCP, DM) Link to external services
XV. (ESSP) Servers and cables of external service provider

Figure 6: Possible failure points of data links, including responsible stakeholder: power plant operator (PPO), metering point operator (MPO), landlord (LL), tenant (TE), utility company (UC), software provider for the DCP (SPDCP), district manager (DM) and external software service provider (ESSP)

3.2 The presupposed data basis for all use cases

In CPSIS, data is the linking element between hardware, software services and the interaction between all stakeholders. A special feature of the energy industry compared to other research areas in urban neighborhoods is that, in addition to static data and relational data of users and facilities, temporally related data series (time series) play a crucial role. Fundamental to discuss is the required temporal resolution and whether they are needed online or only in hindsight.

Since in CPSISs, technical implementations should always serve the social goals, it is necessary to derive the required time series and resolutions from the use cases, the DCP faces. Therefore, in the following, we discuss the use case axis in Figure 5.

Data basis for use case 1: Forecasting

The temporal resolution of the forecasts depends on the one hand on the temporal resolution of the optimization algorithm for which the forecast is required, and on the other hand on the available resolution of any additional input data. Market-based optimizations are usually based on the spot market prices, which are currently provided in Germany with a temporal resolution of 15 min. Additional input data are for example weather data. The DWD provides measured weather data with a resolution of 10 min and forecasts with a resolution of 1 h. When considering all these facts combined, it becomes clear that a temporal resolution of 15 min is optimal for forecasts in the current market and data environment. Certain time series like weather for the next day or the sum of residual consumption of the last day are required for cyclical automated online execution of the machine learning models used for forecasting. Other time series, like the historical weather or the local renewable generation or sum of residual consumption over a larger period, are only recorded for offline retraining of the machine learning models. Hence, the requirement is to record $P_{\text{loc,gen}}$, $P_{\text{flex,cons}}$ and P_{exch} as well as to calculate $P_{\text{sum,res.cons.}}$ for the largest available period, with at least 15 min resolution. The forecasts are divided into load forecasts, that depend solely on historical load patterns (use case 1.1), and PV forecasts, that are weather dependent (use case 1.2). For a detailed description of the services, see section 4.

Data basis for use case 2: Optimization

The functioning of the optimization is explained in detail in subsection 4.4. As discussed in the forecasting use case, it is a market-based optimization that uses spot market prices (c_{sp}), information on the local generation ($P_{\text{loc,gen}}$) and consumption ($P_{\text{sum,res.cons.}}$) and the state of the flexible consumer. For CI the relevant state of the flexible consumer is the energy charged in the current charging event (e_{charged}).

Table 1: Required time series for all use cases of the DCP. The grid status data set consists of three-phase and total active, reactive and apparent power, three-phase voltage, current and their total harmonic distortion (THD), grid frequency and $\cos(\varphi)$

Use case	Processing	Parameter	Unit	Available resolution	Required resolution	Required period
1.1.) Load forecast	online	$P_{\text{sum,res.cons.}}$	W	1 min	15 min	last day
	offline	$P_{\text{sum,res.cons.}}$	W	1 min	15 min	since construction
1.2.) PV forecast	online	global radiation	$\mathrm{W\,m^{-2}}$	1 h	15 min	next day
		air temperature	°C	1 h	15 min	next day
		wind speed	$\mathrm{m\,s^{-1}}$	1 h	15 min	next day
	offline	global radiation	$\mathrm{W\,m^{-2}}$	10 min	15 min	since construction
		air temperature	°C	10 min	15 min	since construction
		wind speed	$\mathrm{m\,s^{-1}}$	10 min	15 min	since construction
		$P_{\text{loc,gen}}$	W	1 min	15 min	since construction
2.) Optimization	online	c_{sp}	€/MWh	15 min	15 min	next day
		$P_{\text{loc,gen}}$	W	15 min	15 min	next day
		$P_{\text{sum,res.cons.}}$	W	15 min	15 min	next day
		e_{charged}	W h	event	event	recent
3.) Billing	offline	c_{el}	€/MWh	15 min	15 min	last year
		$P_{\text{loc,gen}}$	W	15 min	15 min	last year
		$P_{\text{flex,cons}}$	W	15 min	15 min	last year
4.) Peak avoidance	online	P_{exch}	W	1 s	1 s	recent, current 15 min
5.) Research	offline	P_{exch}	W	1 s	1 s	since construction
		$P_{\text{flex,cons}}$	W	1 s	1 s	since construction
		$P_{\text{loc,gen}}$	W	1 s	1 s	since construction
		grid status data set	various	1 s	1 s	during measuring campaigns
6.) Visualization	online	all above	all above	1 s to 1 h	1 min	construction - next day

Data basis for use case 3: Billing

To set the monetary profit incentives for the flexible consumers properly, they have to be billed depending on the electricity price and the local renewable generation. In billing concepts such as "Mieterstrom" (landlord-to-tenant electricity), the consumers' consumption $P_{\text{flex,cons}}$ and the renewable generation $P_{\text{loc,gen}}$ as well as the electricity price c_{el} are countervailed on a 15 min basis. The price in the landlord-to-tenant electricity business model consists of a fixed price for local renewable generation and a power purchase price for the residual energy c_{el}, usually but not necessarily the spot price. The shares of locally generated and purchased energy are then multiplied with their corresponding prices for each 15 min.

Data basis for use case 4: Peak avoidance

One operational goal of a DCP is to minimize the power peaks at the GCP, meaning the peaks of the measured P_{exch} profile. However, unlike in theoretical research on grid support, the question of cost avoidance is usually the driving factor in real-world implementations. In the real world, costs are incurred in the form of peak power prices. These are dependent on the highest 15 min average value of P_{exch} within a month or a year [7]. If P_{exch} is measured with a higher resolution than 15 min, countermeasures can be taken immediately within a critical 15 min interval by adjusting the flexible loads $P_{\text{flex,cons}}$. In general, the faster the reaction, the more efficient the countermeasures. Resolutions from 1 s to 1 min are conceivable. The most suitable resolution, should be investigated in further research. As presented in section 1 no fail-safe online control can be implemented safely via a cloud application of a DCPs over the internet. Therefore, it should be emphasized that peak avoidance through DCP aims at economic benefits. This means that the physical power limits of the transformers or other components are never reached. In the event of a failure, the system would remain in a safe state.

Data basis for use case 5: Research

It is difficult to estimate which data will become relevant for future research in the field of districts. The research question that needs to be answered will determine the data required. The research subject could be, for example:

- Investigate the advantages of using higher resolution data for all existing use cases in a DCP.
- Gain insights into consumption and generation behavior that reveal potentials for planning, maintenance, savings, or efficiency increase.
- Analyze the state of the power grid and changes imposed by larger renewable generators or flexible consumers.
- Acquire data for previously unknown research questions.

These subjects can be broadly divided into improving the existing application and investigating novel issues. For the questions on existing applications, the time series already identified and listed for the other use cases should in any case be recorded and saved with the highest 1 s resolution. For the other questions, an extended data set on the power grid status was defined. Our specific infrastructure allows direct access to the meters via a protected VPN tunnel to record the unique grid status data set in the district. The grid status data set is currently not yet recorded in everyday operation but can be activated for certain interesting time intervals for all 17 meters instrumented by us as part of measurement campaigns. Our grid status data set consists of 26 values listed in Table 1.

Data basis for use case 6: Visualization

To be able to adequately recognize charging processes, special features in inflexible base load consumption, and weather effects in renewable generation, a higher resolution than 15 min is required. We consider 1 min to be sufficient for visualization.

4 SERVICES OF THE DCP

The necessary services are based on the goals and requirements of the stakeholders. For more information on the stakeholders, see section 3. In this section, the focus is purely on the software services of the DCP. The main services are depicted in Figure 7. In the following, the features of the services are briefly explained.

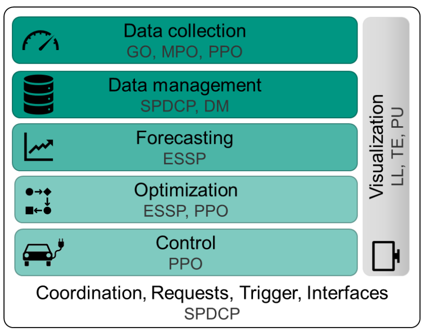

Figure 7: Software services of our DCP related to grid operator (GO), metering point operator (MPO), power plant operator (PPO), software provider for the DCP (SPDCP), district manager (DM), external software service provider (ESSP), landlord (LL), tenant (TE) and the public (PU)

4.1 Data collection

The data collection service is most integrated with the physical infrastructure. It ensures that the data is collected at the points described in subsection 2.2. To accomplish this, it uses the communication network shown in Figure 6. Various depth of integration are possible. The data collection service can be implemented anywhere from the MPO or PPO to the SPDCP. In our case, it is carried out by the MPO as a cloud service that converts Modbus TCP/IP and HTTPS as well as SFTP (depending on the meter) from the meters to MQTT in our data standard. We also implemented a data link to enable the import of measurements from BEMCom [22].

4.2 Data management

The ideal standards for describing data depend on the requirements of the people interacting with the data, the communication protocols used, and the domain in which the data exists. Since numerous stakeholders are involved in DCPs as shown in section 3, the requirement arises to utilize data structures that are easy to understand and describable in a standardized manner. The main communication protocols used to exchange data between stakeholders are MQTT and REST. Other protocols are implemented towards external service providers like MPO or ESSP only as connecting elements. As a domain, "energy" is not a sufficiently precise description. The specific domain is energy management in the context of users, markets, and interaction with the grid at a common GCP. The requirements can then be listed in the form:

- Easily human-readable
- Standardized description
- Suited for MQTT and REST
- Fitted for data exchange between services
- Describe all relevant aspects of energy management

Based on these requirements, it can be decided, for example, whether XML or JSON would be more suitable. In the context of data transfer via MQTT and REST, JSON is better human-readable and easier to parse (directly into other data objects or classes). Another point in favor of JSON is that the data transferred per single message is small in quantity and simple in structure. This follows from the domain (energy management) with a high number of single measurement- or set-points. Individual services internally may, of course, have other data structures and types, for example, XML in visualization or YAML or database formats in data storage.

To be able to describe data cleanly, a common schema is needed. In computer science, this is usually classes of the respective programming languages. But there are also such schemas regarding data standards for data transmission. These are XML schema definitions (XSDs) for XML and JSON Schema for JSON. The existing EEBUS standard [6] is ideally suited for the domain. The EEBUS standard provides sufficient data models for all relevant use cases in the domain described above. These are described in a XSD schema. Unfortunately, a translation into a JSON schema and thus the option to validate JSON messages is missing. There are a variety of data standards for the power system. However, besides a clear focus on the relevant domain, these often lack easy implementability in common programming languages (respectively prefabricated easily importable classes and packages). According to the authors, the application of a data standard primarily fails due to the availability of simple implementations in common programming languages (Java, Python, C++). A simple JSON schema can serve as a template and even be converted into classes of the respective programming languages. In any case, sent messages can be validated against the schema in all languages. Our recommendations are: JSON schema, based on EEBUS SPINE, and Pydantic models (see subsection 4.3).

4.3 Forecasting services

The forecasting services we develop are based on FastAPI and pyWATTS [9]. The data transfer when calling the API is standardized according to the state of the technology via Pydantic models. For further details on data formats also see subsection 4.2. The services are deployed via GitLab CI, JFrog Artifactory and with the help of Argo CD and Vault by HashiCorp. General functions are the creation of a new forecast data object (PV plant or inflexible consumer), training with historical data, deleting or modifying a forecast data object and obtaining a forecast for a specific object. The time series relevant for training (offline) and those relevant for obtaining forecasts (online) have already been presented in Table 1.

PV forecasting service - AutoPV. As listed in Table 1, the PV forecast can be generated purely based on weather data. Actual measurement data and metadata are only required when creating or training the plant. The identification of a plant across creation, training, and execution is undertaken via a UUID (version 4). When creating a new plant, its kWp power must be specified as metadata. During training, the PV mounting configuration is estimated based on the kWp power and measurement data. After training, both the estimated mounting configuration and the model of the PV plant are stored for making predictions afterward. A high-quality estimate of the mounting configuration requires the location of the PV plant and weather data corresponding to this location. To ensure these two requirements, the location of the plant is also stored in the plant model in the form of longitude, latitude, and altitude. For a DCP, however, this information is usually optional, in contrast to, for example, a generic nationwide forecasting service. The same applies to variables such as a scaling factor or period length. The reason is that these variables are typically identical for all forecasts inside a single DCP. In our case, we use weather data from the DWD from the station "Rheinstetten" with the MOSMIX ID "10731" and the observation ID "4177". For detailed insights and evaluations of AutoPV, we refer to [13].

Load forecasting service. The load forecast model uses equidistant time series, in our case, 15 min as shown in Table 1. During training, a period of several days or months is needed, ideally a whole year. During execution, the most recent 24 h (i.e., the last 96 15 min values) are passed to the service and the service delivers the next 96 values as a response. For the underlying model, it is not relevant whether the periods are 15 min, 1 h or an other duration. However, since models trained with a certain period can only meaningfully forecast time series with the same period, the scaling factor and time period are again part of the submitted data model. The load forecast is calculated via pyWATTS [9] and scikit-learn [15] using linear regression.

4.4 Optimization service and control service

The core function of the DCP is the control of the flexible consumers based on the time series presented in Table 1. The information flow in which time series are used is shown in Figure 8. In general, a new optimization is triggered by the arrival of new information. In the simplest case, this is an updated forecast or the arrival or departure of vehicles at the charging stations. The next most frequent reason

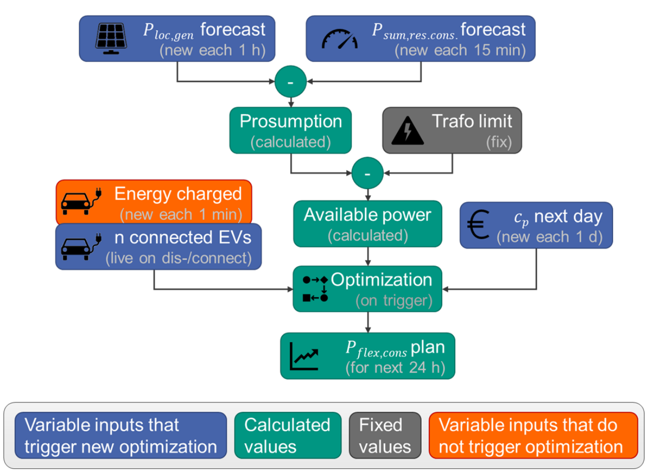

Figure 8: Inputs and calculation steps for optimization

for a re-optimization, besides the arrival or departure of a vehicle, is a new load forecast. Since the load forecast is always based on the 96 last measured 15 min power values, a new forecast can be generated as soon as a new value is available for a further 15 min period. However, the change in input for the machine learning model is minimal, as 95 values remain unchanged. The update frequency of the PV forecast is determined by the availability of new weather forecasts. When using DWD MOSMIX [5] the forecast is updated each hour around 25 min past the full hour. The least frequently updated variable is the price, with one update per day. For EPEX SPOT clearing starts at 12:00 CET and data is available 1 h - 2 h later depending on data provider. In stationary operation with a constant number of EVs, the optimization is consequently only renewed every 15 min with some variables only updated once per day. Numerous works exist on specific algorithms for optimization (an extensive overview can be found in Shewale et al. [18] or Panda et al. [14]) and therefore no further details will be given here. To import the price time series, we recommend an external API like awattar [2]. The functionality of the load and PV forecast services, which were developed by us, are briefly presented in the subsection 4.3.

The control is performed directly by the charging stations and vehicles. The values are transmitted as set points through an Open Charge Point Protocol (OCPP) backend from the ICTP to the CI.

4.5 Visualization service

The front-end of the DCP is shown in Figure 9. First, it shows a general overview of the area, which reveals details about the facilities when zooming in further. Clicking on the items displays measurement values in 1 min and 15 min resolution for all meters shown on the map. In other tabs, individual customers can log in, charging processes can be displayed and charging stations can be reserved in advance.

5 CONCLUSION

We present the real-world, productively operated district community platform (DCP) of an urban district in Karlsruhe's Oststadt. In doing so, we look at the connected systems, the relevant measuring points, and the information and communications technology (ICT). We address the multidimensional social aspects. In particular, we work out in which technical issues people are necessarily involved as a social component. This concerns especially not only the tenants or the users of the system but also all the companies responsible for the communication and physical infrastructure, as well as all the information and communications technology platform (ICTP) services. We show, which time series and services are necessary for the operation and which data standards we recommend for uniform communication. In particular, as depicted in Figure 9, we demonstrate how a central ICTP can bundle the requirements and social goals and give an accessible interface for all stakeholders.

Further work will include analyzing the data collected for research purposes. This includes heat consumption time series. It will be investigated which other facilities can be integrated into the DCP, for example, whether it would make sense to switch from district heating to heat pumps in the long term.

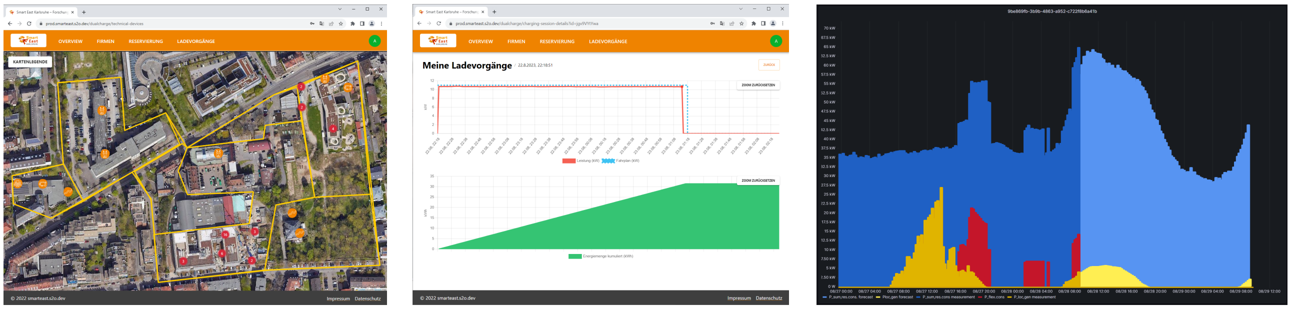

(a) Landing page: Overview of district with currently 48 devices / linked time series

(b) Charging process with a $P_{\mathbf{flex,cons}}$ fixed limit of 11 kW

(c) Example GCP with $P_{\mathbf{sum,res.cons.}}$ (blue), $P_{\mathbf{loc,gen}}$ (yellow) and $P_{\mathbf{flex,cons}}$ (red) including 24 h **forecast (pale)**

Figure 9: Results from the real-world DCP operation

ACKNOWLEDGMENTS

This work was supported by the 'Smart East' (funding reference L75 21113) and Helmholtz 'ESD' project. The DCP was a cooperative effort of the partners Seven2one, InnoCharge GmbH, FZI Research Center for Information Technology and Stadtwerke Karlsruhe GmbH. We also thank all associated partners for their support, and the Institute for Automation and Applied Informatics for the provision of algorithms and IT infrastructure.

AUTHOR CONTRIBUTIONS

Author contributions according to Contributor Roles Taxonomy (CRediT, https://credit.niso.org/): *Conceptualization:* J.G., S.W.; *Methodology:* J.G.; *Formal analysis:* J.G.; *Investigation:* J.G.; *Writing - original draft:* J.G., S.W.; *Writing - review and editing:* J.G., S.W., S.M., R.M., V.H.; *Visualization:* J.G.; *Supervision:* S.W., R.M., V.H.; *Project administration:* S.W.; *Funding acquisition:* V.H.

REFERENCES

[1] Mohamed Anis Aguida, Samir Ouchani, and Mourad Benmalek. 2020. A Review on Cyber-Physical Systems: Models and Architectures. In *2020 IEEE 29th International Conference on Enabling Technologies: Infrastructure for Collaborative Enterprises (WETICE)*. IEEE, Bayonne, France, 275–278. https://doi.org/10.1109/WETICE49692.2020.00060

[2] aWATTar Deutschland GmbH. 2023. aWATTar Service API. https://www.awattar.de/services/api

[3] Christos G. Cassandras. 2016. Smart Cities as Cyber-Physical Social Systems. *Engineering* 2, 2 (June 2016), 156–158. https://doi.org/10.1016/J.ENG.2016.02.012

[4] Enzaemilia Cavallaro, Maria Rosaria Sessa, and Ornella Malandrino. 2023. Renewable Energy Communities in the Energy Transition Context. *International Journal of Energy Economics and Policy* 13, 3 (May 2023), 408–417. https://doi.org/10.32479/ijeep.14230

[5] Deutscher Wetterdienst. 2023. Wetter und Klima - Deutscher Wetterdienst - Leistungen - Model Output Statistics-MIX (MOSMIX). https://www.dwd.de/DE/leistungen/met_verfahren_mosmix/met_verfahren_mosmix.html

[6] EEBus e.V. 2023. home. https://www.eebus.org/

[7] Johannes Galenzowski, Simon Waczowicz, and Veit Hagenmeyer. 2023. A new Data-Driven Approach for Comparative Assessment of Baseline Load Profiles Supporting the Planning of Future Charging Infrastructure. In *Companion Proceedings of the 14th ACM International Conference on Future Energy Systems (e-Energy '23)*. 8. https://doi.org/10.1145/3599733.3600245

[8] Kai Gong, Jianlin Yang, Xu Wang, Chuanwen Jiang, Zhan Xiong, Ming Zhang, Mingxing Guo, Ran Lv, Su Wang, and Shenxi Zhang. 2022. Comprehensive review of modeling, structure, and integration techniques of smart buildings in the cyber-physical-social system. *Frontiers in Energy* 16, 1 (Feb. 2022), 74–94. https://doi.org/10.1007/s11708-021-0792-6

[9] Benedikt Heidrich, Andreas Bartschat, Marian Turowski, Oliver Neumann, Kaleb Phipps, Stefan Meisenbacher, Kai Schmieder, Nicole Ludwig, Ralf Mikut, and Veit Hagenmeyer. 2021. pyWATTS: Python Workflow Automation Tool for Time Series. https://doi.org/10.48550/arXiv.2106.10157 arXiv:2106.10157 [cs].

[10] International Energy Agency (IEA). 2021. *Key World Energy Statistics 2021*. Technical Report. Paris. https://iea.blob.core.windows.net/assets/52f66a88-0b63-4ad2-94a5-29d36e864b82/KeyWorldEnergyStatistics2021.pdf

[11] International Energy Agency (IEA). 2022. *World Energy Outlook 2022*. Technical Report. Paris. https://iea.blob.core.windows.net/assets/830fe099-5530-48f2-a7c1-11f35d510983/WorldEnergyOutlook2022.pdf

[12] Michael Krug, Maria Rosaria Di Nucci, Matteo Caldera, and Elena De Luca. 2022. Mainstreaming Community Energy: Is the Renewable Energy Directive a Driver for Renewable Energy Communities in Germany and Italy? *Sustainability* 14, 12 (Jan. 2022), 7181. https://doi.org/10.3390/su14127181 Number: 12.

[13] Stefan Meisenbacher, Benedikt Heidrich, Tim Martin, Ralf Mikut, and Veit Hagenmeyer. 2023. AutoPV: Automated photovoltaic forecasts with limited information using an ensemble of pre-trained models. In *Proceedings of the 14th ACM International Conference on Future Energy Systems*. 386–414. https://doi.org/10.1145/3575813.3597348 arXiv:2212.06797 [cs].

[14] Subhasis Panda, Sarthak Mohanty, Pravat Kumar Rout, Binod Kumar Sahu, Mohit Bajaj, Hossam M. Zawbaa, and Salah Kamel. 2022. Residential Demand Side Management model, optimization and future perspective: A review. *Energy Reports* 8 (Nov. 2022), 3727–3766. https://doi.org/10.1016/j.egyr.2022.02.300

[15] Fabian Pedregosa, Gael Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine Learning in Python. *Journal of Machine Learning Research* 12 (2011), 2825–2830.

[16] Jonathan Reynolds, Yacine Rezgui, and Jean-Laurent Hippolyte. 2017. Upscaling energy control from building to districts: Current limitations and future perspectives. *Sustainable Cities and Society* 35 (Nov. 2017), 816–829. https://doi.org/10.1016/j.scs.2017.05.012

[17] Filipe Duarte Santos, Paulo Lopes Ferreira, and Jiesper Strandsbjerg Tristan Pedersen. 2022. The Climate Change Challenge: A Review of the Barriers and Solutions to Deliver a Paris Solution. *Climate* 10, 5 (May 2022), 75. https://doi.org/10.3390/cli10050075 Number: 5.

[18] Amit Shewale, Anil Mokhade, Nitesh Funde, and Neeraj Dhanraj Bokde. 2022. A Survey of Efficient Demand-Side Management Techniques for the Residential Appliance Scheduling Problem in Smart Homes. *Energies* 15, 8 (Jan. 2022), 2863. https://doi.org/10.3390/en15082863 Number: 8.

[19] United Nations Environment Programme. 2022. *Global Status Report for Buildings and Construction: Towards a Zero-emission, Efficient and Resilient Buildings and Construction Sector*. Technical Report. Nairobi. https://globalabc.org/sites/default/files/2022-11/FULL%20REPORT_2022%20Buildings-GSR_0.pdf

[20] Anne-Lorène Vernay, Carine Sebi, and Fabrice Arroyo. 2023. Energy community business models and their impact on the energy transition: Lessons learnt from France. *Energy Policy* 175 (April 2023), 113473. https://doi.org/10.1016/j.enpol.2023.113473

[21] Yi Wang, Chien-Fei Chen, Peng-Yong Kong, Husheng Li, and Qingsong Wen. 2023. A Cyber–Physical–Social Perspective on Future Smart Distribution Systems. *Proc. IEEE* 111, 7 (July 2023), 694–724. https://doi.org/10.1109/JPROC.2022.3192535

[22] David Wölfle, Manuel Lösch, and Hartmut Schmeck. 2022. BEMCom: a framework for the efficient creation of hardware abstraction layers for building energy management. *ACM SIGEnergy Energy Informatics Review* 2, 1 (Feb. 2022), 20–25. https://doi.org/10.1145/3527579.3527581

[23] Jun Jason Zhang, Fei-Yue Wang, Xiao Wang, Gang Xiong, Fenghua Zhu, Yisheng Lv, Jiachen Hou, Shuangshuang Han, Yong Yuan, Qingchun Lu, and Yishi Lee. 2018. Cyber-Physical-Social Systems: The State of the Art and Perspectives. *IEEE Transactions on Computational Social Systems* 5, 3 (Sept. 2018), 829–840. https://doi.org/10.1109/TCSS.2018.2861224 Conference Name: IEEE Transactions on Computational Social Systems.