

# Integrating Power-to-Heat Services in Geographically Distributed Multi-Energy Systems: A Case Study from the ERIGrid 2.0 Project

Giuseppe Silano<sup>1,2</sup>, Evangelos Rikos<sup>3</sup>, Vetrivel Rajkumar<sup>4</sup>, Oliver Gehrke<sup>5</sup>, Tesfaye Amare Zerihun<sup>6</sup>, Carmine Rodio<sup>1</sup>, and Riccardo Lazzari<sup>1</sup>

**Abstract**—This paper investigates the integration and validation of multi-energy systems within the H2020 ERIGrid 2.0 project, focusing on the deployment of the JaNDER software middleware and universal API (uAPI) to establish a robust, high-data-rate, and low-latency communication link between Research Infrastructures (RIs). The middleware facilitates seamless integration of RIs through specifically designed transport layers, while the uAPI provides a simplified and standardized interface to ease deployment. A motivating case study explores the provision of power-to-heat services in a local multi-energy district, involving laboratories in Denmark, Greece, Italy, the Netherlands, and Norway, and analyzing their impact on electrical and thermal networks. This paper not only demonstrates the practical application of Geographically Distributed Simulations and Hardware-in-the-Loop technologies but also highlights their effectiveness in enhancing system flexibility and managing grid dynamics under various operational scenarios.

## I. INTRODUCTION

In the evolving landscape of multi-energy systems, the drive towards integrating diverse energy sources and optimizing their operations remains a pressing concern [1], [2]. This integration is crucial as it encompasses various dimensions of energy systems including power, heat, and transport, all undergoing significant transformations due to decentralization, decarbonization, and digitalization [3], [4]. Traditionally, research in this domain has been skewed towards optimizing the planning and operational strategies within isolated engineering frameworks, often overlooking the intricate interconnections and interdependencies that exist among different energy networks [5], [6].

Recent advancements have triggered a wave of innovation aimed at addressing these complexities, with significant emphasis on the development of new tools and methodologies that span across different engineering domains [7], [8]. Notably, the concept of Hardware-in-the-Loop (HIL) and its derivative, Geographically Distributed Simulations (GDS),

have emerged as pivotal in enhancing the validation and deployment of novel multi-energy systems [9], [10]. These methodologies facilitate accelerated validation processes and enable the harnessing of distributed resources and expertise, thereby broadening the scope of what can be achieved within the constraints of individual RIs.

This paper explores the intersection of technological advancements and their practical applications, particularly emphasizing the use of HIL and GDS in multi-energy systems. Specifically, within the framework of the H2020 ERIGrid 2.0 project<sup>1</sup>, it examines the deployment of the JaNDER [11], [12] software middleware and universal API (uAPI) [8], [13] for establishing a secure, reliable, high-data-rate, and low-latency communication between Research Infrastructures (RIs). A motivating case study focuses on providing ancillary services through power-to-heat strategies in a local multi-energy district, examining their impact on electrical and thermal networks. This initiative aims to enhance power system flexibility as requested by the Transmission and Distribution System Operators, utilizing a combination of electric and thermal storage systems, alongside demand response strategies from controllable loads such as heat pumps and electric boilers [14]. Additionally, the heating system is designed to offer flexibility to the electrical system, e.g. for congestion management and power balancing services. Consequently, the key contributions of this paper are:

- Providing a detailed description of the software setup used for implementing communication and conducting geographically distributed experiments among RIs using the JaNDER open-source middleware and the uAPI. The presented case study serves as a practical guide that other researchers can use, modify, or integrate to test their algorithms and understand different approaches within their RIs. This enables the study of performance and stability in multi-energy systems.
- Demonstrating the feasibility of providing power-to-heat services in a local multi-energy district and assessing their impacts on electrical and thermal networks. Furthermore, it highlights the potential of such services to manage congestions and provide balancing power.

The remainder of this paper is organized as follows. Section II discusses the motivating case study. Section III details the software framework for RI communication and experiment management. The experimental results of GDS are presented in Section IV. Section V concludes the paper.

<sup>1</sup>G. Silano, R. Lazzari, and C. Rodio are with the Ricerca sul Sistema Energetico, Milan, Italy, (emails: {name.surname}@rse-web.it).

<sup>2</sup>G. Silano is with the Czech Technical University in Prague, Czech Republic, (email: giuseppe.silano@fel.cvut.cz).

<sup>3</sup>E. Rikos is with the Centre For Renewable Energy Sources and Saving, Athens, Greece, (email: vrikos@ces.gr).

<sup>4</sup>V. Rajkumar is with the Delft University of Technology, Delft, The Netherlands, (email: v.subramaniamrajkumar@tudelft.nl).

<sup>5</sup>O. Gehrke is with the Technical University of Denmark, Roskilde, Denmark, (email: olge@dtu.dk).

<sup>6</sup>T. A. Zerihun is with the SINTEF Energy AS, Trondheim, Norway (email: tesfaye.zerihun@sintef.no).

This work received funding under the European Community's Horizon 2020 Program (H2020/2014-2020) in project “ERIGrid 2.0” (Grant Agreement No. 870620).

<sup>1</sup><https://erigrd2.eu/>

## II. MOTIVATING CASE STUDY

The case study considered here explores the provision of critical services to the electrical grid, with particular emphasis on congestion management – including the constraints associated with electrical import and export – and regulating power provision. The focal point of this case study is the assessment of a Centralized Supervisory Controller (CSC) within a geographically distributed local multi-energy system. This analysis is crucial for understanding the operational dynamics of system components interconnected across various regions.

The experimental setup encompasses a tripartite system: an electrical grid with energy sources, a thermal network, and an advanced control system, all graphically represented in Figure 1. The *thermal infrastructure* (red dashed box) is divided into two subsystems. The first is a District Heating Network (DHN) located in Denmark at the Technical University of Denmark (DTU), interfaced with a Combined Heat and Power (CHP) unit based in Italy at the Ricerca sul Sistema Energetico (RSE), connected through an electrical coupling unit. The second subsystem consists of a thermal load (L1) powered by an Electrical Heat Pump (EHP), both situated in Greece at the Center For Renewable Energy Sources (CRES). The *electrical infrastructure* (blue dashed box) includes a distribution grid and a few controllable units. The distribution grid adopts the CIGRE LV-distribution benchmark [15], characterized by a 0.4 kV, 50 Hz low-voltage system with resources working as controllable current sources, allowing precise control over active ( $P$ ) and reactive power ( $Q$ ) through the calculation of current magnitude and phase relative to the bus voltage. The schematic representation of this network is detailed in Figure 3. This network is emulated through an RTDS real-time simulator<sup>2</sup> located in the Netherlands at the Technische Universiteit Delft (TUD), with two physically different nodes virtually interfaced in a GDS setup. To these nodes, PCC2 and PCC4, are connected the simulated Battery Energy Storage System (BESS) of Stifelsen for industriell og teknisk forskning (SINTEF), and the CHP with the EHP, respectively. The *control system* delves with accurately determining setpoints for these controllable entities to ensure the delivery of required services.

Regarding the hardware setup of the experiment, it includes several key components. The EHP (CRES) has nominal power 16 kW and capacity 18 kV A, with controllable variables including its operating state (ON/OFF). Various measurements are recorded, such as active, reactive, and apparent power, energy consumption, voltage, frequency, and indoor temperature at multiple points. Interfacing with other RIs involves data collection via an eWon 4001 datalogger and control through a Raspberry Pi 4 Model B. The DHN (DTU) consists of two 440 m long double pipes connecting two buildings as depicted in Figure 2. Building 1 houses a controllable electrical heat source and heat consumer, while Building 2 functions as a thermal substation interconnecting the pipes, creating a total pipe length of 880 m between

the source and sink. The heat source includes nine electrical flow heaters totaling 22.5 kW, feeding into a 200 L accumulator tank. To address the size mismatch between the CHP plant at RSE (46 kW to 81 kW) and the DTU heat source (0 kW to 22.5 kW), a linear mapping function offsets and scales the thermal power setpoints received through the coupling interface, ensuring the full controllable range of the CHP plant is utilized. The three-phase distribution grid (TUD) operates at 0.4 kV, 50 Hz and supplies five loads (see Figure 3). The BESS managed by SINTEF includes a converter and a power amplifier, along with measurement points including State-of-Charge (SoC) and instantaneous power. Lastly, the CHP (RSE) plant utilizes a natural gas internal combustion engine, bidirectional converters, and power amplifiers within a three-phase low-voltage grid setup.

## III. DISTRIBUTED RI SETUP

This section details the software framework developed to facilitate communication among the RIs engaged in the experimental setups of the motivating case study. At the core of this framework is the deployment of uAPI [8], [13], powered by the JaNDER middleware [11], [12]. The uAPI serves as a transport-independent abstraction layer, enabling the use of various software middleware for multi-RI experiments. This eliminates the need to implement individual laboratory interfaces and provides common core functionality, such as accessing a list of available signals, RI statuses, and more. JaNDER complements the uAPI by enabling secure and efficient data exchange between RIs via an HTTPS-secured internet connection. Its primary functionality includes replicating infrastructure data across network nodes, ensuring that data captured in local database is simultaneously mirrored in a cloud-based database. This setup guarantees data consistency and redundancy without altering or interpreting the data's structure.

A schematic representation of the JaNDER's software architecture can be found in Figure 4, and the source code used for the experiments is publicly available on GitHub<sup>3</sup>. These resources provide a practical guide for other researchers to replicate, modify, or integrate the setup for testing their algorithms and studying different approaches within their RIs. This framework enables the study of performance and stability in multi-energy systems, offering a robust platform for real-time experimentation and data analysis. Further details about the open-source middleware are discussed in [11], [12] but are omitted here for brevity.

The operational deployment of the JaNDER middleware and uAPI requires a Linux-based Operating System (OS), such as Ubuntu. This choice promotes the use of open-source tools and avoids dependency on specific proprietary licenses. The setup process involves several key steps. Firstly, a Linux-based OS is installed on a dedicated or secondary machine separate from the RI's main data monitoring and control infrastructure. Following the OS installation, Docker<sup>4</sup>

<sup>2</sup><https://www.rtds.com/>

<sup>3</sup><https://github.com/ERIGrid2/JRA-3.1-JaNDER-API>

<sup>4</sup><https://www.docker.com/>

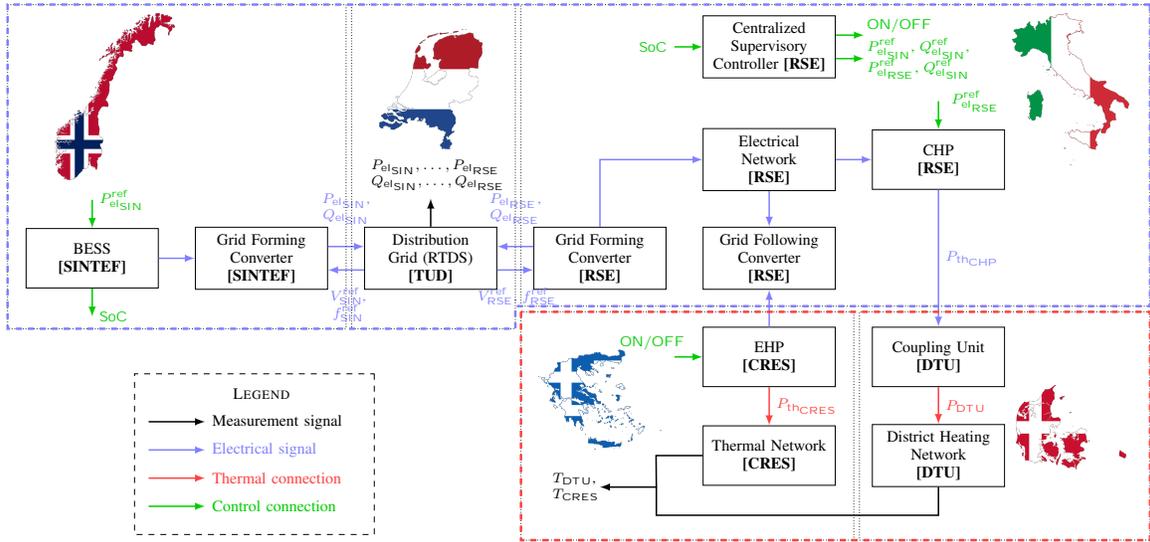


Fig. 1: Schematic representation of the multi-energy district case study, with the electrical and thermal subsystems highlighted by blue and red dashed boxes, respectively.

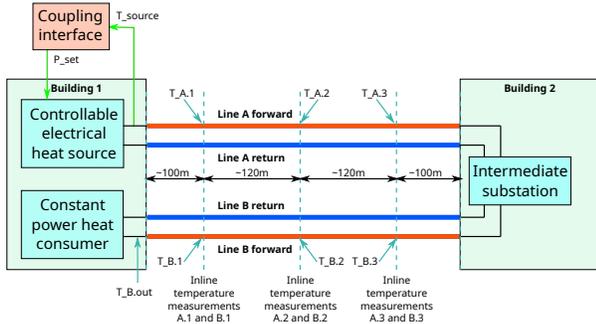


Fig. 2: Schematic representation of the DHN at DTU, with approximate measurement locations marked

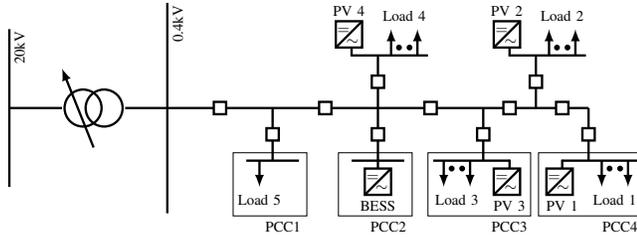


Fig. 3: Schematic representation of the CIGRE LV-distribution benchmark grid along with the PCC.

is used to containerize applications, ensuring secure, isolated, and portable environments. This streamlines the software deployment process and minimizes configuration requirements. Secure communication between RIs is facilitated by generating private and public certificates. These certificates provide robust authentication and encryption mechanisms [8], [11]. Finally, environment configuration involves updating configuration files to include namespaces for the RIs, such as RSE, TUD, and CRES, ensuring correct identification within the network.

Once the development environment is configured, the

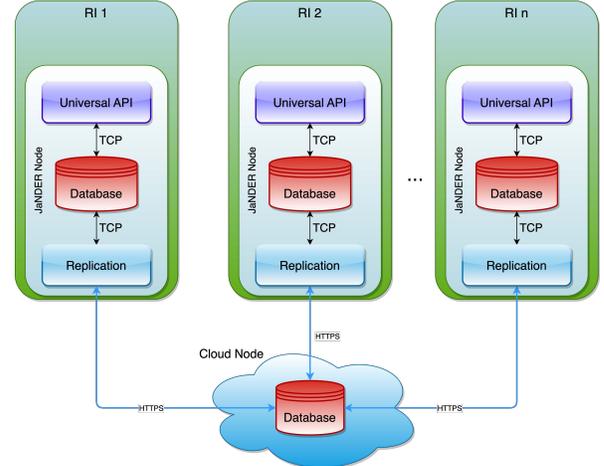


Fig. 4: Diagram of the JaNDER middleware, featuring RIs each equipped with the uAPI and local database instances, which are connected to a shared database in the cloud node.

Docker container must be compiled. At this stage, communication between the RIs is established, facilitating data exchange via the uAPI. The uAPI serves as an abstraction layer that enables integration with SCADA systems by defining a set of REST functions [13]. The implementation of these functions supports GET/SET operations, which are essential for real-time monitoring and control of the multi-energy system. For instance, the uAPI defines endpoints for data retrieval (GET) and data updates (SET). The REST functions allow the SCADA system to request current data states or send control commands to the various components of the multi-energy system. Detailed documentation and code examples to implement these REST functionalities are available in the uAPI repository<sup>5</sup>.

<sup>5</sup><https://github.com/ERIGrid2/JRA-3.1-api>

TABLE I: Signals exchanged among RIs, including their symbols and operational ranges.

Sym.	Unit	Min	Max	Sym.	Unit	Min	Max
$P_{elSIN}$	kW	-40	40	$Q_{elSIN}$	kVAr	-5	5
$P_{thCHP}$	kW	46	81	$P_{elSIN}^{ref}$	kW	-40	40
$\bar{P}_{DTU}$	kW	0	25	SoC	%	0	100
$V_{SIN}^{ref}$	V	150	400	$f_{SIN}^{ref}$	Hz	48	52
$P_{elRSE}$	kW	-100	100	$Q_{elRSE}$	kVAr	-50	50
$V_{RSE}^{ref}$	V	150	400	$f_{RSE}^{ref}$	Hz	48	52
$P_{thCRES}$	kW	0	30	$T_{DTU}$	°C	0	100

#### IV. EXPERIMENTAL RESULTS

To validate the effectiveness of the proposed experimental setup and demonstrate the feasibility of providing ancillary services through power-to-heat strategies in a local multi-energy district, a sector-coupling experimental demonstration inspired by [7] was conducted.

The performed experiments demonstrated the impact of providing flexibility services on both the electrical and thermal networks. The amount of flexibility requested by the system operator is achieved through the control of electric and thermal units, such as storage systems, heat pumps, thermal loads, and electric boilers. The distributed laboratory setup, as shown in Figure 1, utilizes the uAPI for data exchange. Given the “slower” dynamics typical of electro-thermal experiments, which operate on the scale of seconds, the data-exchange rate is set between 1 Hz to 2 Hz. Two distinct working scenarios were delineated:

- **Case 1 – Overvoltage Scenario:** this scenario addresses an overvoltage condition triggered by high Photovoltaic (PV) generation, either PV1 and PV3 located at PCC4 and PCC3, respectively, as depicted in Figure 3. This condition is coupled with low electricity demand. It is noteworthy that the DHN does not influence the outcomes in this scenario.
- **Case 2 – Undervoltage Scenario:** this scenario examines an undervoltage condition resulting from reduced or nonexistent PV generation and elevated electricity demand. It mirrors the overvoltage scenario (Case 1) with the critical distinction of employing negative threshold values for the operational parameters of the units. This condition necessitates the engagement of both the CHP unit and thermal storage to ensure grid stability.

This comprehensive case study aims to shed light on the operational efficacy of a CSC within a distributed multi-energy framework, highlighting its capacity to manage grid congestions and provide regulating power under a spectrum of operational conditions.

Figure 5 illustrates the outcomes of the overvoltage scenario aimed at addressing high PV generation coupled with low electricity consumption, leading to an overvoltage condition. To prepare for this test, TUD developed profiles for the PVs and simulated loads<sup>6</sup>, inducing a significant voltage

<sup>6</sup>The PVs and simulated loads profiles are not included here for brevity and as they are not crucial to the purpose of the case study which primarily focuses on showcasing the viability of deployment of the JaNDER software middleware and uAPI.

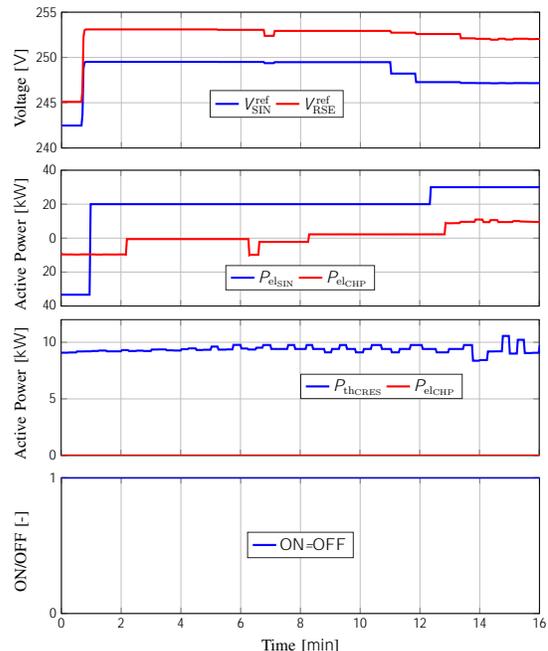


Fig. 5: Reference voltage sent from the distribution grid (TUD) to the grid forming converters (SINTEF and RSE), along with the active power generated by the BESS and CHP and enable signals of the EHP in the overvoltage scenario.

rise at the nodes connected to SINTEF ( $V_{SIN}^{ref}$ ) and RSE ( $V_{RSE}^{ref}$ ). During the test, SINTEF and RSE, integrated into the simulation model (see Figure 1), active power consumption ( $P_{elSIN}$  and  $P_{elRSE}$ ) to mitigate voltage deviations. SINTEF employed a simplified control approach, adjusting the active power setpoint of the battery in a stepwise manner, as depicted in the lower left plot of Figure 5. For instance, significant power consumption was activated if the voltage rise exceeded 5% and deactivated if it dropped below 0%, ensuring stability during the test. Similarly, RSE employed the same hysteresis approach to control the electricity consumption of CRES’ heat pump, reducing the voltage rise at the corresponding node, as shown in the plots on the right side of Figure 5. It is important to note that the 5% threshold is indicative and may vary based on the simulated profiles, highlighting the necessity for accurate estimations via offline simulations. Furthermore, in this scenario, where an increase in electric load is necessary, neither the thermal power of the combined heat and power (CHP) unit nor the thermal capacity of the heat network (DTU) can be utilized to mitigate the overvoltage condition.

In this context, RSE could still communicate the thermal ( $P_{thCHP}$ ) setpoint value to DTU, but it would need to be set to zero as it does not play any role. Whereas, Figure 6 illustrates the reference ( $P_{elSIN}^{ref}$ ) and actual ( $P_{elSIN}$ ) active power profiles, as well as the relative increase in the SoC of the BESS, at the SINTEF facility. Additionally, it shows the reactive power profiles of SINTEF ( $Q_{elSIN}$ ) and RSE ( $Q_{elRSE}$ ), along with the reference frequencies ( $f_{SIN}^{ref}$  and  $f_{RSE}^{ref}$ ).

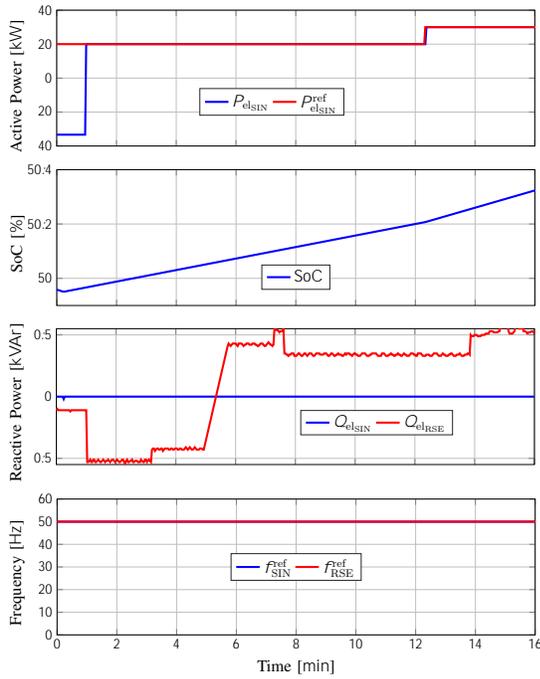


Fig. 6: Reference and actual active power profiles at the SINTEF facility, along with the SoC of the BESS, during the overvoltage scenario experiment. Additionally, reactive power and frequency data are included.

Figure 7 presents the outcomes of the undervoltage scenario, where the voltage at the coupling points of SINTEF and RSE suddenly drops below the reference value (240 V) due to low or zero PV generation or high consumption in the grid (see Figure 3). The procedure for this scenario parallels that of the overvoltage scenario, employing a negative threshold value of 5% for activation and 0% for deactivation of units. Initially, the CRES EHP operates to ensure its deactivation during the experiment, as shown in the lower right corner of Figure 7. Hence, the scenario necessitates additional generation by the CHP, requiring the utilization of both the CHP of RSE and the thermal network (DTU).

To restore the voltage to the reference value (or close to it), RSE generated active power ( $P_{elRSE}$ ) by disabling the EHP ( $P_{thCRES}$  goes to zero) and activating the CHP, while the BESS stopped charging ( $P_{elSIN}$  goes to zero). This recovery process, depicted in Figure 8, occurs within minutes due to the rapid dynamics involved. Notably, the increase in voltage is not directly correlated with active power but instead requires reactive power ( $Q_{elSIN}$  and  $Q_{elRSE}$ ), as evident in the plots (see Figure 8). The system utilizes thermal energy to provide services, with insufficient generation of reactive power resulting in the need to generate active power. This supports the fault ride-through capability, enabling the RIS to remain connected despite voltage fluctuations and provide voltage support and thermal heating services. Figure 9 demonstrates a decrease in load due to the BESS stopping its charging, illustrating the impact of the electrical grid on the thermal network. The battery was maintained at 50% capacity, as observed in the lower left corner of Figure 9.

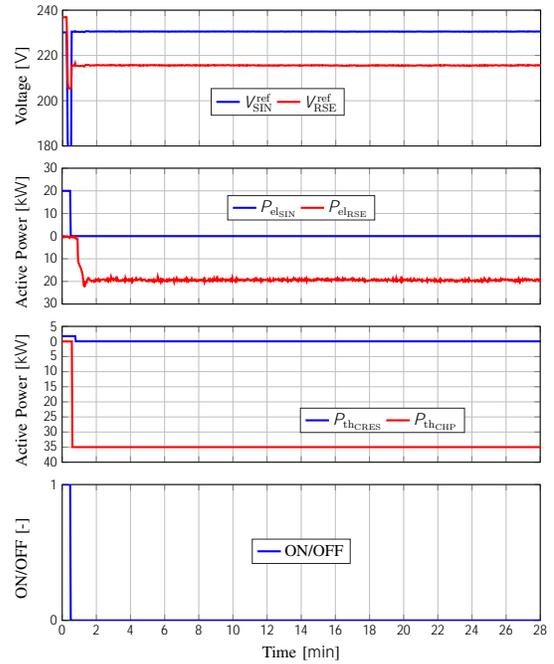


Fig. 7: Reference voltage sent from the distribution grid (TUD) to the converters (SINTEF and RSE), along with the active power generated by the BESS and CHP, and the enable signals of the EHP, in the undervoltage scenario.

The control strategy was simplified, primarily focusing on adjusting active power without considering other parameters like temperatures. If any anomalies arise during the experiments, involved partners could locally override setpoints and deactivate equipment. While both experiments could be combined into one, conducting them separately was more time-efficient, considering the real-time nature of the tests.

It is important to note that one of the challenges of multi-domain experiments is the difference in the time scales at which relevant phenomena occur. Thermal systems react orders of magnitude more slowly than electrical systems. This is illustrated by the thermal response observed during the experiment, as depicted in Figure 10. The bottom plot shows the output of the controllable heat source over a three-hour experiment. The heat source tracks the remote CHP plant's output with an offset and scaling factor to match their controllable ranges, limited by a resolution of 2.5 kW due to the load steps. The dashed vertical line indicates the end of the coupling experiment, after which no further data is exchanged, and the heat source is turned off. Post-coupling, all dynamic processes occur solely within the thermal system. The upper plot shows dynamic processes continuing for about twelve hours due to the small pump used, which caused low mass flow rates. However, even with a larger pump, the thermal response would still be slower.

Figure 10 only shows the forward line's response; similar dynamics occur in the return line. Towards the end, cold return water from the heat consumer lowers the temperature at the network start ("Line A in"), as the heat source has been off for hours and the accumulator tank's energy is depleted.

