Grant Agreement: 768936

D3.3: REPORT - INTEGRATED IMPROVED INNOVATIONS IN ENERPIPE NETWORK

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 768936.

PROJECT DOCUMENTATION SHEET	
Project Acronym	**TEMPO**
Project Full Title	Temperature Optimisation for Low Temperature District Heating across Europe
Grant Agreement	**768936**
Call Identifier	H2020-EE-2017-RIA-IA
Topic	EE-04-2016-2017 – New heating and cooling solutions using low grade sources of thermal energy
Type of action	IA Innovation Action
Project Duration	54 months (October 2017 – March 2022)
Coordinator	Vlaamse Instelling voor Technologisch Onderzoek NV (BE) – VITO
Consortium partners	Noda Intelligent Systems AB (SE) – NODA Austrian Institute of technology (AT) – AIT Thermaflex international holding (NL) - THF Steinbeis innovation (DE) – Solites Enerpipe (DE) – Enerpipe Vattenfall (DE) – Vattenfall A2A Calore & servizi (IT) - A2A Hogskolan i Halmstad (SE) – HU Euroheat & Power (BE) – EHP
Website	www.tempo-dhc.eu
Disclaimer	The sole responsibility for the content of this document lies with the authors. It does not necessarily reflect the opinion of the funding authorities. The funding authorities are not responsible for any use that may be made of the information contained herein.

DELIVERABLE DOCUMENTATION SHEET	
Number	**D3.3**
Title	Report integrated improved innovations in ENERPIPE network
Related WP	WP3 (Enerpipe Demonstrator)
Related Task	Task 3.3 (Implementation of the TEMPO innovations)
Lead Beneficiary	ENERPIPE
Author(s)	Davy Geysen (VITO) – davy.geysen@vito.be Chris Hermans (VITO) – chris.hermans@vito.be Andreas Fiegl (ENERPIPE) – andreas.fiegl@enerpipe.de Manfred Abt (ENERPIPE) – manfred.abt@enerpipe.de Christian Johanson (Noda) – cj@noda.se Jens Brage (Noda) – jens.brage@noda.se Aurelien Bres (AIT) – aurelien.bres@ait.ac.at Thomas Schmidt (Solites) - schmidt@solites.de
Contributor(s)	Günther Beck – gunther.beck@noda.se
Reviewer(s)	Dirk Vanhoudt (VITO) - dirk.vanhoudt@vito.be
Nature	R (Report)
Dissemination level	PU (Public)
Due Date	October 2021 (M37)
Submission date	29 October 2021 (M37)
Status	Final

QUALITY CONTROL ASSESSMENT SHEET			
Issue	**Date**	**Comment**	**Author**
V0.1	11/10/2021	First draft	Davy Geysen (VITO) WP Leader
V0.2	15/10/2021	Input by partners	All involved partners
V0.3	22/10/2021	Second draft	All involved partners
V0.4	27/10/2021	Peer review	Dirk Vanhoudt (VITO)
V1.0	29/10/2021	Quality review	Johan Desmedt (VITO) - Project coordinator
V1.1	29/10/2021	Submission to EC	Johan Desmedt (VITO) - Project coordinator

EXECUTIVE SUMMARY

This deliverable "Report integrated improved innovations in Enerpipe network" is developed within WP3 of the H2020 TEMPO project. TEMPO is the acronym for: Temperature Optimisation for Low Temperature District Heating and focusses on the development, demonstration and deployment of innovations for low temperature district heating (DH) networks across Europe. TEMPO aims at reducing DH network system temperatures to achieve improved network efficiency, costs competitiveness and capability of integrating sustainable energy sources like renewable and residual heat.

This document will give a brief description of the different innovations together with a detailed explanation of the technical and economic aspects of implementing them at the demo site. Furthermore, we will also discuss the quality control measures that were put into place and the commissioning and initial operation of the demo site.

The ENERPIPE demo is located in the town of Windsbach, 35 km south west of Nuremberg (Germany). The municipality of Windsbach is developing a new residential housing project in a rural area. Several technological innovations will be implemented on the demo site to analyse their impact while operating on a real district heating network. The following technological innovations are implemented at the demo site:

- Supervision ICT platform
- Visualization tools for expert and non-expert users
- Smart DH controller
- Decentralised buffers
- Optimisation of the building installations

The houses will be heated by means of a District Heating network and the project consists of 3 phases:

- Phase 1
  - 59 single family houses
  - 3 multi family houses
- Phase 2
  - Additional 57 single family houses
- Phase 3
  - Additional 20 single family houses

Only Phase 1 and 2 are included as part of the H2020 TEMPO project. Phase 3 is not part of the H2020 TEMPO project. As of October 2021, 74 dwellings are connected to the heating network. There still are some houses to be built, so a connection ratio of approximately 75% may be expected.

Figure 1: The Enerpipe demo in the municipality of Windsbach

The network was designed for approximately 90% connection ratio, with a then-estimated total heat demand of 1 GWh per year.

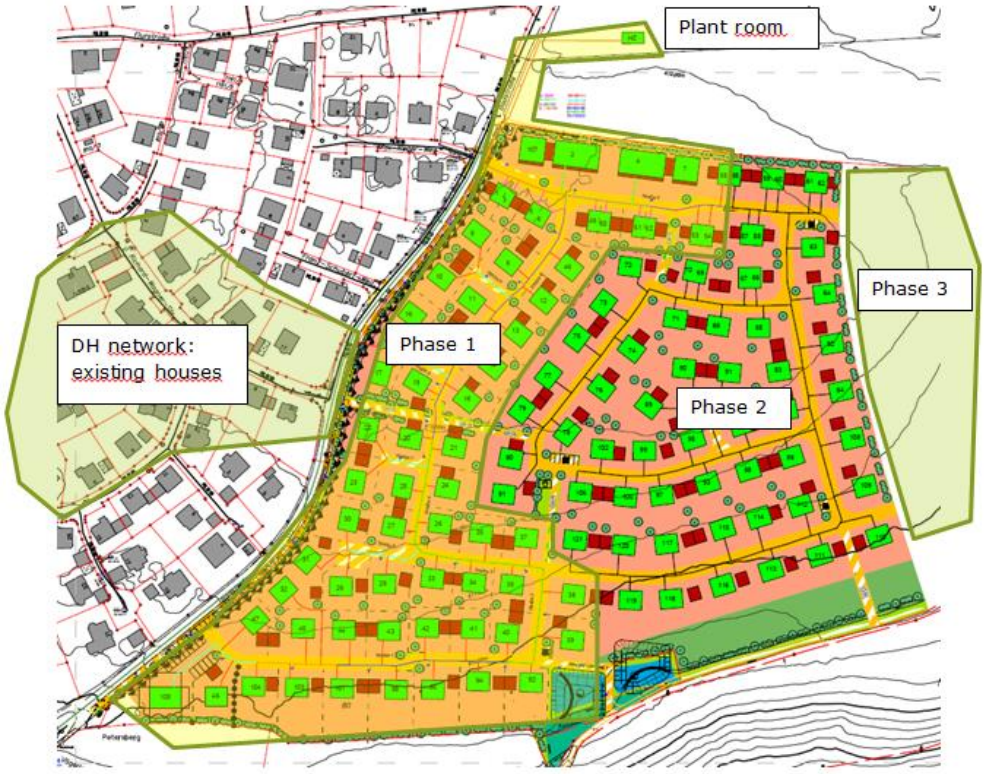

Figure 2: District heating scheme of the Windsbach project

In order to determine the necessary heat network pipe diameters, a temperature difference of 30K between flow and return has been assumed, since the houses that would be built should typically be thoroughly insulated and mostly equipped with floor heating or other

low-temperature heating systems. The district heating network water is being transported by means of PE-Xa tubes (mainly twin-tubes).

The heat stems from a biogas plant ("satellite") some 800m away from the heating central, and two biogas-fueled cogeneration (or CHP, for "combined heat and power) plants inside said heating central. They provide 200 and 400kW of electrical peak power, with a corresponding 207 and 505kW of thermal output. During summer months, the satellite provides enough waste heat to cover the whole network's heat demand. During winter, the cogeneration plants run intermittently, meeting the network's heat demand while storing heat in two 20m³ buffer storage tanks inside the heating central. In case of a CHP plant failure, there's a backup gas burner that allows burning the biogas for heat generation.

Smart DHN controller

By applying the theoretical algorithms that were developed in WP1 into the Windsbach demo of WP3, we encountered many practical and technical challenges. At first, the optimization algorithm needs to be controlling a system which is already operated with a conventional controller, e.g. the Schneid controller in Windsbach. The behavior of this controller was different than the assumptions made in WP1, e.g. the backup controller did not safeguard end-user comfort if the smart controller was sending out non-optimal control signals. Due to this, a new optimal control approach, purely data-driven, was defined, implemented and deployed. In order to get the calculated control signal of this approach applied on a distributed buffer, four different platforms are used. The signal gets calculated on a VITO server, where it sent to the EnergyView platform. In the EnergyView platform another piece of software is running to export the control signals out of EnergyView, at the defined times, and send them to the Schneid system. The Schneid system in the end sends the charging signals to the buffers. This workflow implies that there are four different places where errors can occur and each of the systems has its own expert and responsible person. In order to debug the whole system, different people are needed. As a result, in the majority of the cases it takes longer to locate the problem than to solve it. In that sense, it would be better if the conventional controller is able to receive external control signals and apply them directly to the devices in the system (if comfort is not jeopardized). This would rule out two of the four systems, doubling the debugging efficiency. In an ideal world, the installed system together with its operating software, would be designed in cooperation with the company responsible for deploying the smart controller. However, in the majority of the cases this is not feasible as the DHN and its controller are already in operation at the time the smart controller is defined. Therefore, the 20%-80% rule certainly applies when deploying a control algorithm in practice, e.g. 20% of the time is needed to develop the algorithm and 80% is needed to get it up and running in the field.

Decentralized buffers

Installing decentralized buffers in each dwelling allowed for a comparatively slim design of the district heating network. The first two phases of the development area are completed to a large degree, and there seem to be no signs for supply shortages. The network flow temperature could be reduced significantly, from 80°C in a baseline-scenario to 65-72°C. However, the original concept behind the heating network with the centralized buffers aimed for a significantly lower return temperature. There are a couple of issues increasing said return temperature, e.g. badly manufactured short return lances, heating systems

with elevated temperatures inside several dwellings, and a rather remote multi-family house requiring an elevated temperature level for its apartment stations, sending back high temperature. The same multi-family house requires high network flow temperatures, that could otherwise be further reduced. Nevertheless, the return temperature would be even higher in a network with secondary instead of primary storage buffers.Several consumers with flawless performance (including low return temperatures) show that the initial assumptions might be met (or at least approached) in a network with reasonably designed and executed building installations.

TABLE OF CONTENTS

GLOSSARY / LIST OF ACRONYMS

ACRONYM	DEFINITION
DHS	District Heating System
DHN	District Heating Network
DH	District Heating
DHW-station	Domestic Hot Water - station
CHP	Combined heat and power
SoC	State of charge (temperatures in a decentralized buffer)

1 IMPLEMENTED TEMPO INNOVATIONS

The following paragraphs illustrate a general description of these innovations and how the innovations are implemented in this specific demonstration site. The TEMPO solution package of the Enerpipe demonstrator consists of the following innovations:

- Supervision ICT platform
- Visualization tools for expert and non-expert users
- Smart DH controller
- Decentralised buffers
- Optimisation of the building installations

1.1 SUPERVISION ICT PLATFORM

BASELINE VERSUS TEMPO IMPLEMENTATION (STATE OF THE ART)

In traditional district heating systems, primarily data for billing is saved and analysed, and then only analysed sufficiently enough to perform the billing itself. The TEMPO supervision ICT platform bases the pre-processing and data analysis on a more operational perspective. In substations with suboptimal behaviour, the water from the DH network is cooled insufficiently and as a result the return temperature to the DH network will be higher than needed. The supervision ICT platform serves to detect such faults and deviations causing suboptimal behaviour in consumer substations. An important part of that part is the pre-processing phase. The analysis is based on data. However, data is often corrupted and often slightly different than desirable, e.g., it is desirable to treat DHW separately from heating, but often only one heat meter is installed. And we often wish for more flow meters and temperature sensors. Hence, it is necessary to make educated guesses of what's going on and what data to trust. This is especially true in the Windsbach case, with distributed local storage that creates a user profile not commonly seen in district heating systems.

The primary focus of the first version of the ICT platform was to create a robust data management platform, including the functionality to make it accessible by project partners. Such a set-up is also required for future potential scalability and replicability. Included in this application was an API (application programming interface) wrapper for easy access by project partners. This is based on standard Rest API functionality and includes encryption and proper user authentication.

IMPROVEMENTS TO STATE OF THE ART IMPLEMENTATION

For the second version of the ICT platform, the focus has been to implement tools for improved data management. This included improved functionality for extraction when dealing with necessary cross-referencing and error handling, transformation of sensors names and units of measure including removal of repeated values and outliers based on standard deviation, resampling on desired time resolution, interpolation over shorter periods of time to supply missing values and computation of synthetic values. These are all tools that support the general usability of the data management system.

An important addition to the ICT platform was the ability for script management built-into the system. This makes it possible to manage simple, but common, task by a scripting language built into the ICT system. An example of this is the manage virtual sensors, e.g.

calculating the mean value of a set of physical measurements. Having the ability for such script management greatly enhances the usability of an ICT data management platform, and supports scalability since such scripts can be re-used to create the required behaviour on the different levels of the data management process.

FUNCTIONALITY

The following figure shows a screenshot from the script management implementation in the EnergyView ICT system. This is based from the Windsbach demo. In this particular case, the script systems is used to transform time series control plans to current individual control signals. The script manager is based on Tengo, which is a small and dynamic script language for the Go programming language. This particular script has made it possible for other project actors to simplify their implementations for active control in the Windsbach demo.

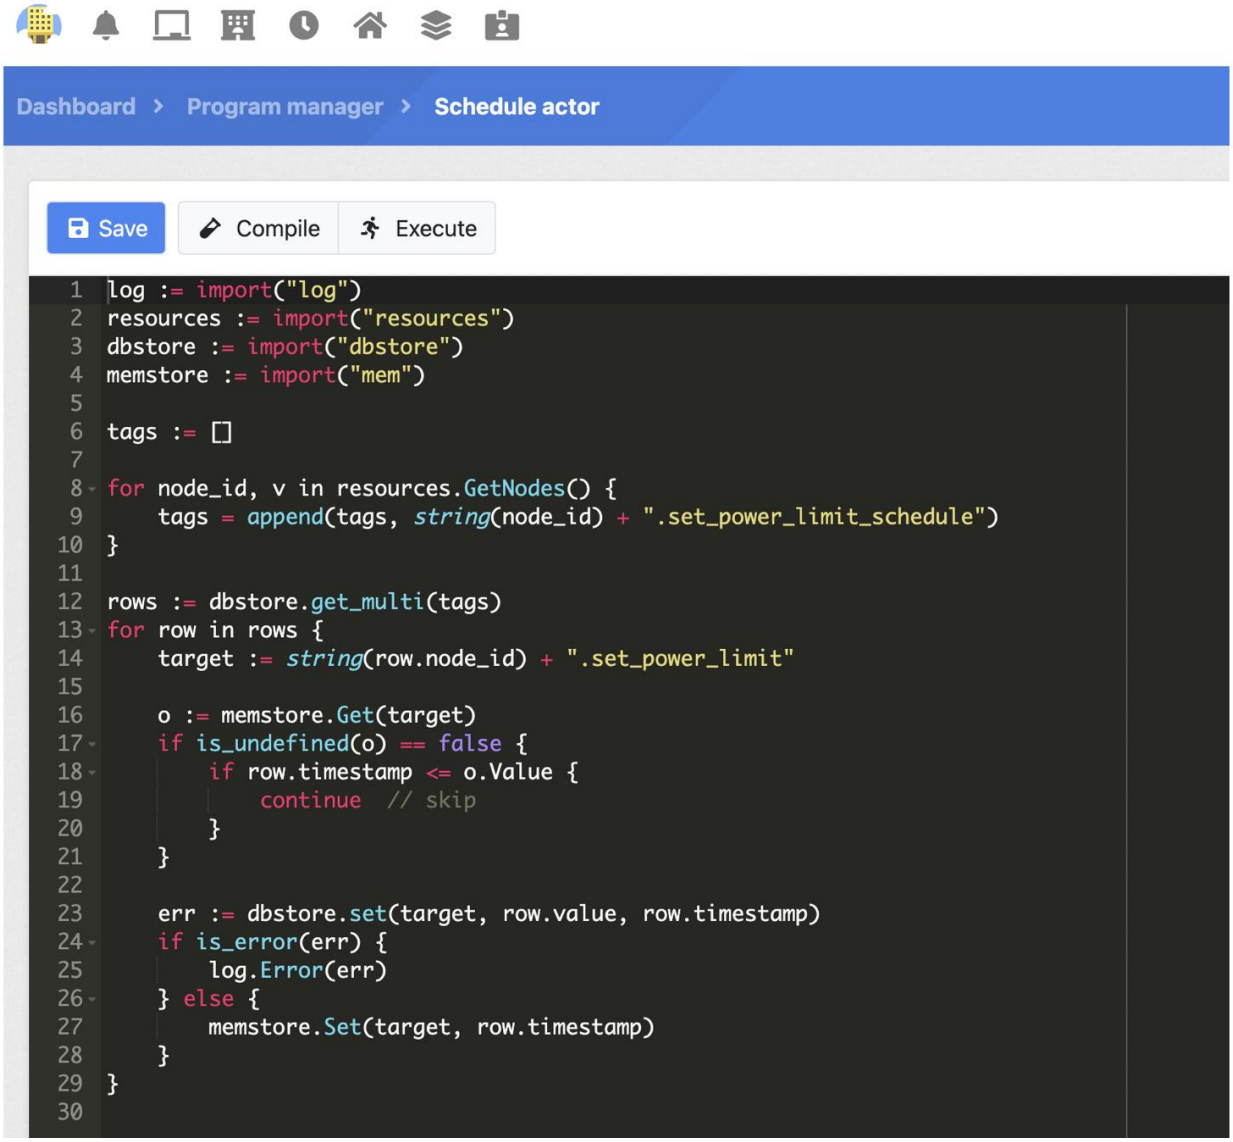

Figure 3: Screenshot of script management implementation in NODA's EnergyView

The following figure shows a screenshot of the API management user console in the EnergyView ICT system. This is also from the deployed system in Windsbach, and this is currently being used by other project partners in order to perform the actual control testing. The ICT system manages different levels of security from the lowest level (i.e. Nobody) to a selection of more advanced access levels up to full system administrator. To ensure secure communication, the system used authentication for both manual and programmatical access. In this screenshot the API key is blocked by a red square.

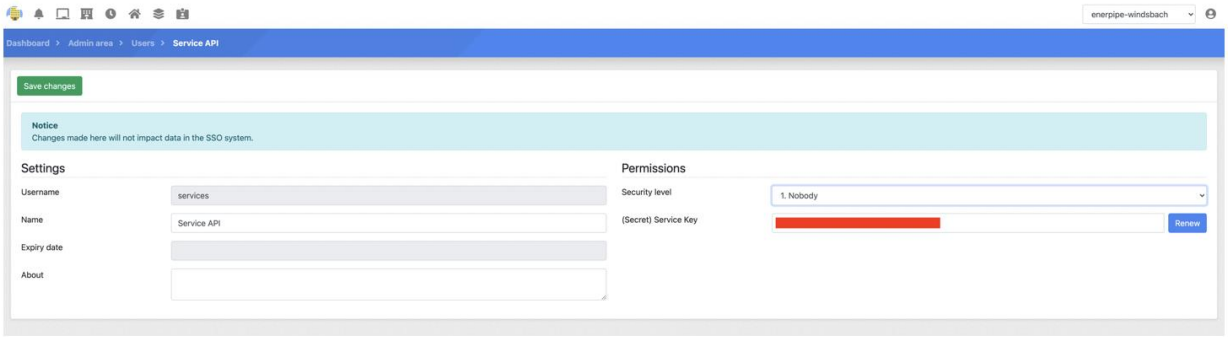

Figure 4: Screenshot of API management user console

COSTS OF THE SUPERVISION ICT PLATFORM

Data access for the ICT platform has been implemented through an integration with the existing Schneid control system. This integration is used as the core communication channel for all the system development effort in Windsbach. The overall ICT platform has been developed with the whole project in mind. Therefore, it is hard to exactly separate what part directly relates to the Windsbach demonstrator. However, considering a relevant share of the overall work brings the estimate to about EUR 50 000 relating to this demonstrator for the above-described effort. This includes the sub-contracting for the Schneid integration with Beck & Partner.

TEST RESULTS

The functionality of the ICT platform has been evaluated in relation to the primary usage in the Windsbach, which is the active control of distributed storage buffers. This includes basic data fault management as part of the pre-processing stage. It also includes the API access functionality which at the core of enabling the end-to-end functionality of the control system. This has been tested using unit testing according to standard software engineering practices.

1.2 VISUALIZATION TOOLS FOR EXPERT AND NON-EXPERT USERS

BASELINE VERSUS TEMPO IMPLEMENTATION (STATE OF THE ART)

The primary visualization tool for the work in Windsbach is the EnergyView system. This system is both a data management system (i.e. the basis for the ICT platform) and a web-based visualization tool for human users. Therefore, it is to a large extent co-existing with the functionality described in the previous section about the ICT platform. In other words, EnergyView is basically a front-end to the ICT platform and they share the same database structure and architecture. Most configuration of the ICT system is also done through the same front-end system, which makes the system interlinked.

The visualization front-end provides a way to really examine and evaluate the operational behavior of all the data being injected into the system regardless of where this data comes from or the original format of that data. In this sense, the EnergyView system provides an important addition to the usability of these type of systems. The following figure provides an example live data from the Windsbach demo, and shows how data from different sources co-exist in the same system.

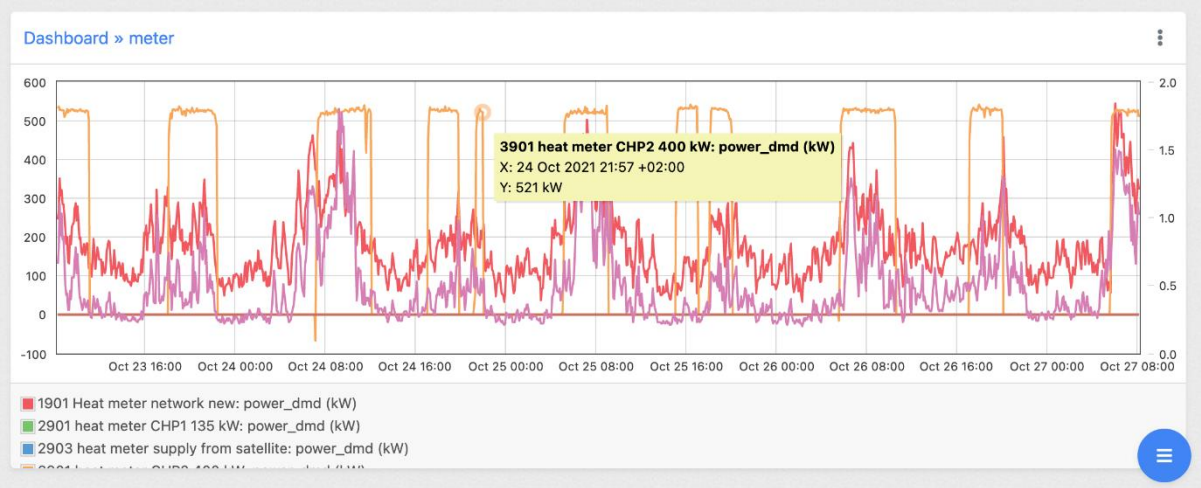

Figure 5: Screenshot of live data from the Windsbach demo site

On addition that has been made for the first version is an improved alarm system. This has made the system more usable in practice by providing a way to better notice and understand possible faults in the data management process. A complete alarm management pipeline has been implemented. This includes different levels of alarm notifications such as open, acknowledged, shelved, expired and closed. The following figure shows the alarm system running in the Windsbach demo. At the time of screenshot, there were no alarms in the system.

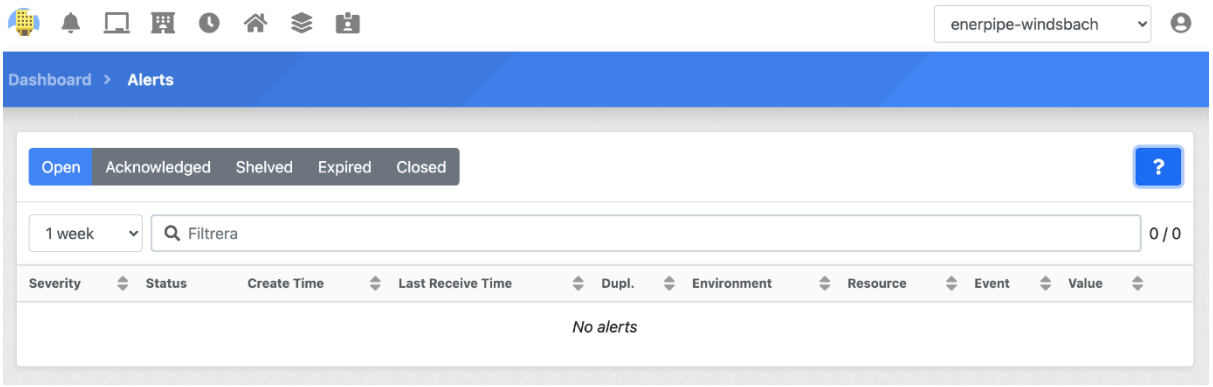

Figure 6: Screenshot of alarm system running in the Windsbach demo site

IMPROVEMENTS TO STATE OF THE ART IMPLEMENTATION

To further the usability of the visualization system the widget system has been further developed. Widgets can be seen as small applications that can be used to set up a visual dashboard which is directly related to the needs and wants of the user. This process makes it possible to adapt the visualization tool to either expert or non-experts, since it can be configured on what data to show and how to show it.

The widget system now handles a dedicated widget app as a front-end for the alarm and alerts system. It also provides an app for audit log, which provides insight into important changes to the configuration or operational behavior of the system. A widget dedicated for taking notes has also been added. This seemingly simply widget is actually based on markup language, which means that it can embed webpages and links and thereby provide a powerful tool to create customer-specific user experiences. There is also a widget to display geographical information about buildings. This obviously require compliance with GDPR, so it is not extensively used within the project itself. However, it is an important addition for post-project scalability. The widget system is continuously being extended and improved. The following picture shows a screenshot of the widget configuration page.

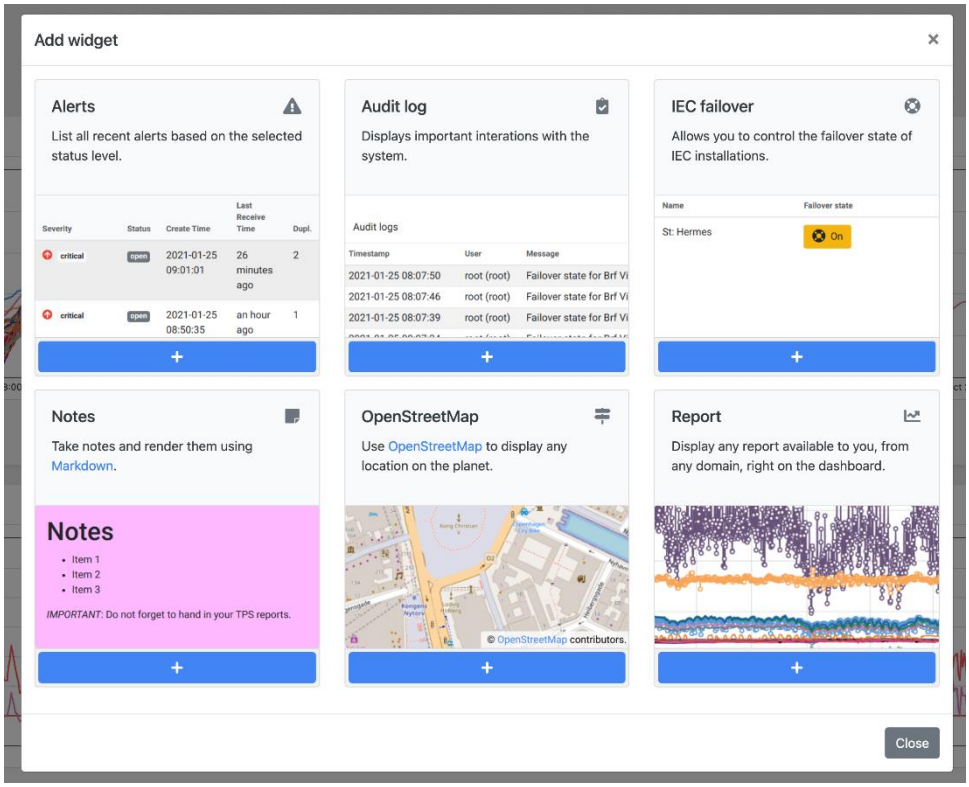

Figure 7: Screenshot of the widget system in the Windsbach demo site

FUNCTIONALITY

An important addition to the visualization tool is the solution integration functionality. This makes it possible to easily integrate data from different sources and in different formats. This functionality is used to ease the integration with the underlying Schneid control system used in the substations for each individual building. The figure below shows an overview of the solution integration that are now implemented in the Windsbach demo.

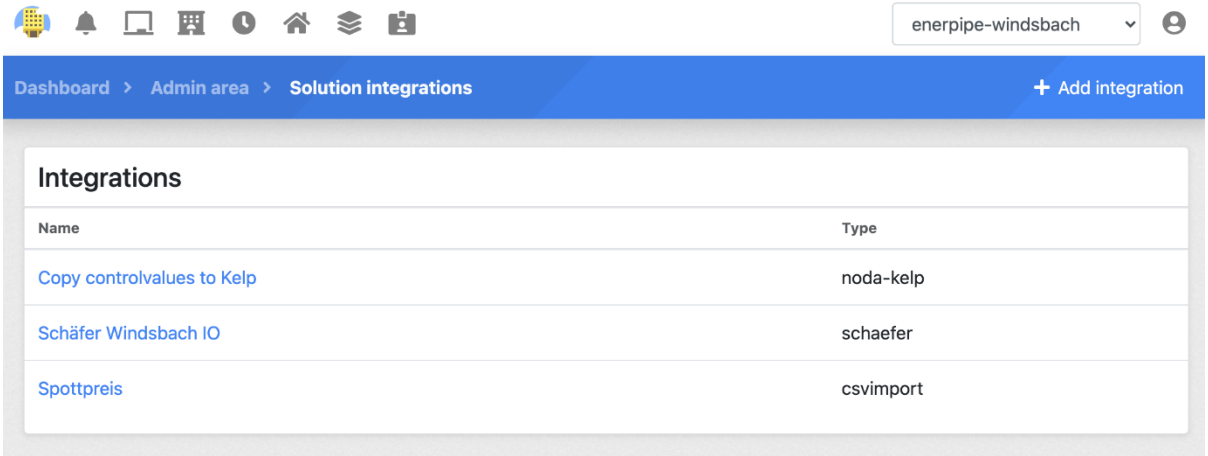

Figure 8: Screenshot of the solution integration in the Windsbach demo site

This has, for example, made it possible to easily interpret imported files in CSV format and automatically inject such data into the system. In the Windsbach demo, this has made it possible to programmatically manage spot prices for electricity as input for the optimization models. In this case, the functionality for solution integration is used to set up rules based on regular expressions to parse the content of those CSV files. The screenshot below shows how this is currently implemented in the Windsbach demo.

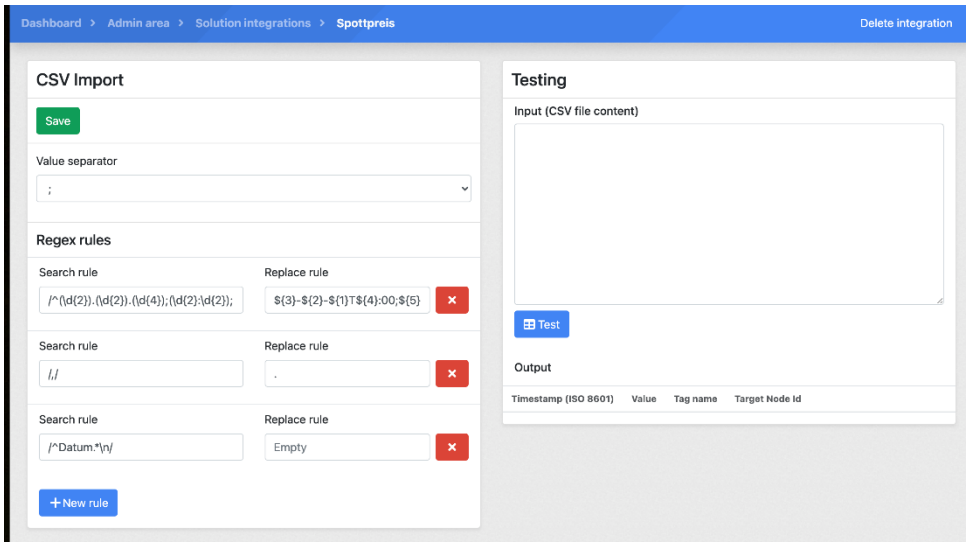

Figure 9: Screenshot of the import solution in the Windsbach demo site

COSTS OF THE VISUALIZATION TOOLS

The development of the EnergyView front-end is spread out over the whole project, and also intertwined with the development of the ICT data management system. This makes it hard to quantify the development cost specifically for the Windsbach demo. However, to deploy the system anew will require the integration and the configuration of the dashboard and reports. Both these are now easily done based on the innovative functionality added to the system. This ensures post-project scalability and replicability in general. The development effort done within the project is not required if deploying the system a second time.

TEST RESULTS

The EnergyView front-end development has been performed according to industry standard software development standards. This includes unit testing and deployment testing before operational deployment. The EnergyView front-end has been used throughout the project as a visualization tool, by several project partners. It is therefore currently considered thoroughly tested in an operational environment.

KPI COMPARISON

Table 4: Overview of efficiency gains due to the individual TEMPO innovations in HT solution package for urban areas

LT solution package for urban areas	Efficiency gain during peak load	Efficiency gain during off-peak load
Decentralised buffers	6 %	6 %
Visualisation tools	3 %	3 %
Smart DH network controller	7 %	15 %
Optimisation of building installation	7,5 % (static) + 10 % (dynamic)	7,5 % (static) + 10 % (dynamic)
Total estimate	**About 25 %**	**About 30 %**

The efficiency gains by implementing visualization tools is based on previous research studies showing that operators having more knowledge about their system and readily availability of relevant data also have the ability to manage their system in a more optimal way. To what extent this has happened so far is under evaluation within the project, and must be verified in relation to other on-going energy efficiency measures.

IMPACT OF THE RESULTS ON THE OVERALL PROPOSED SOLUTION

The ICT platform and the EnergyView front-end both contribute to the implementation and on-going work of the project. They provide a way for project partners to see what is going on with the algorithms. Smart algorithms cannot exist by themselves in some vacuum. They have to exist in some framework, and the EnergyView system provides that framework. Therefore, EnergyView has a significant impact on the overall project performance and post-project scalability and usability potential, although it is not part of the underlying algorithms as such.

1.3 DECENTRALISED BUFFERS

BASELINE VERSUS TEMPO IMPLEMENTATION (STATE OF THE ART)

In conventional district heating systems, only the heating central is equipped with a buffer. In the tempo innovation, every dwelling is equipped with a buffer tank to store some heat. This allows for lower heat demand peaks in those dwellings, since short-time peaks like running a bath will be supplied from the buffer, at least partially.

Nowadays, dwellings that already even have a buffer tank usually store secondary water after it has been heated up by the district heating water. In the TEMPO project, however, the buffers contain primary water (the same water that is inside the district heating network). This allows for lower heat network flow temperatures, since the domestic hot water station mounted on the buffer tank will have access to the higher primary temperature, as compared to the lower secondary temperature. Using a buffer storage tank inside every dwelling allows stopping the heating network's recirculation pump. This reduces electricity consumption and heat losses, but requires adequate superior control. Furthermore, the TEMPO innovation is a plug-and-play set of equipment that may be installed very quickly, as the installer does not need to install and mount every single component individually in each dwelling. The 250l buffer tank comes with mounted domestic hot water station, control valves, heating pump, mixing valve, expansion vessel and much more. Its dimensions are very compact and do not exceed those of a common fridge (approx. 61 x 61 x 205cm).

The picture above shows the first supplied version of the decentralized buffer unit

Alternatively, for larger homes, a few units of a 600l-variant have been installed, with a rather similar scheme. However, its dimensions are larger.

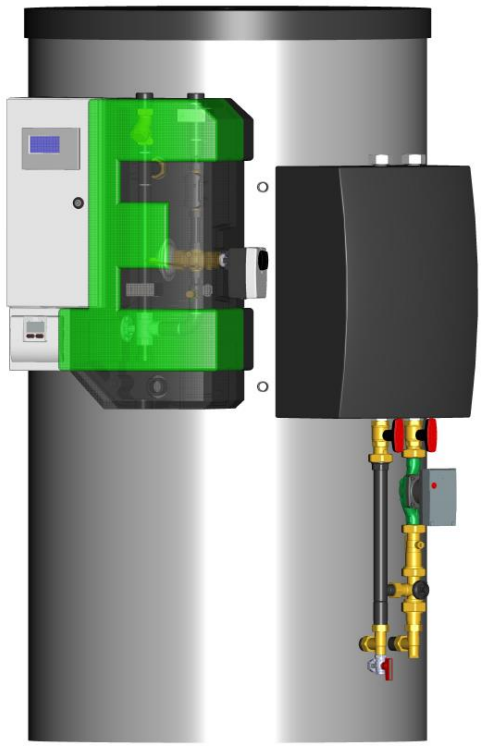

Figure 10: Figure o f larger buffer tank

IMPROVEMENTS TO STATE OF THE ART IMPLEMENTATION

During the course of the project, the innovation was constantly further developed. Especially hydraulically, a number of aspects have been improved:

- accessibility of the components is now much better, resulting in shorter installation times,
- the domestic hot water station has been properly redesigned and is now equipped with its own heat quantity measurement,
- domestic hot water circulation lance has been introduced into the buffer tank, to improve buffer stratification in buffers with domestic hot water circulation,
- hydraulics section may now be installed independently from the buffer tank – faster assembly.

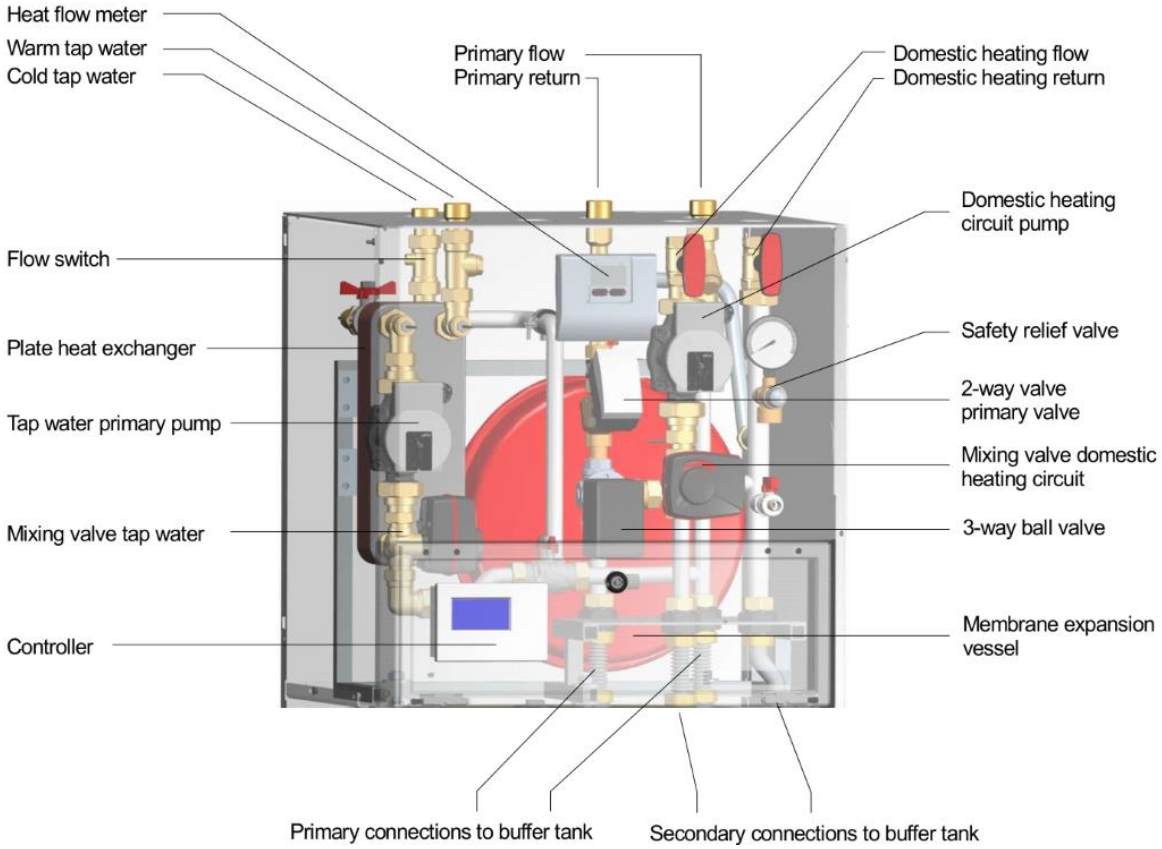

Figure 11: Schematic overview of decentralised buffer

FUNCTIONALITY

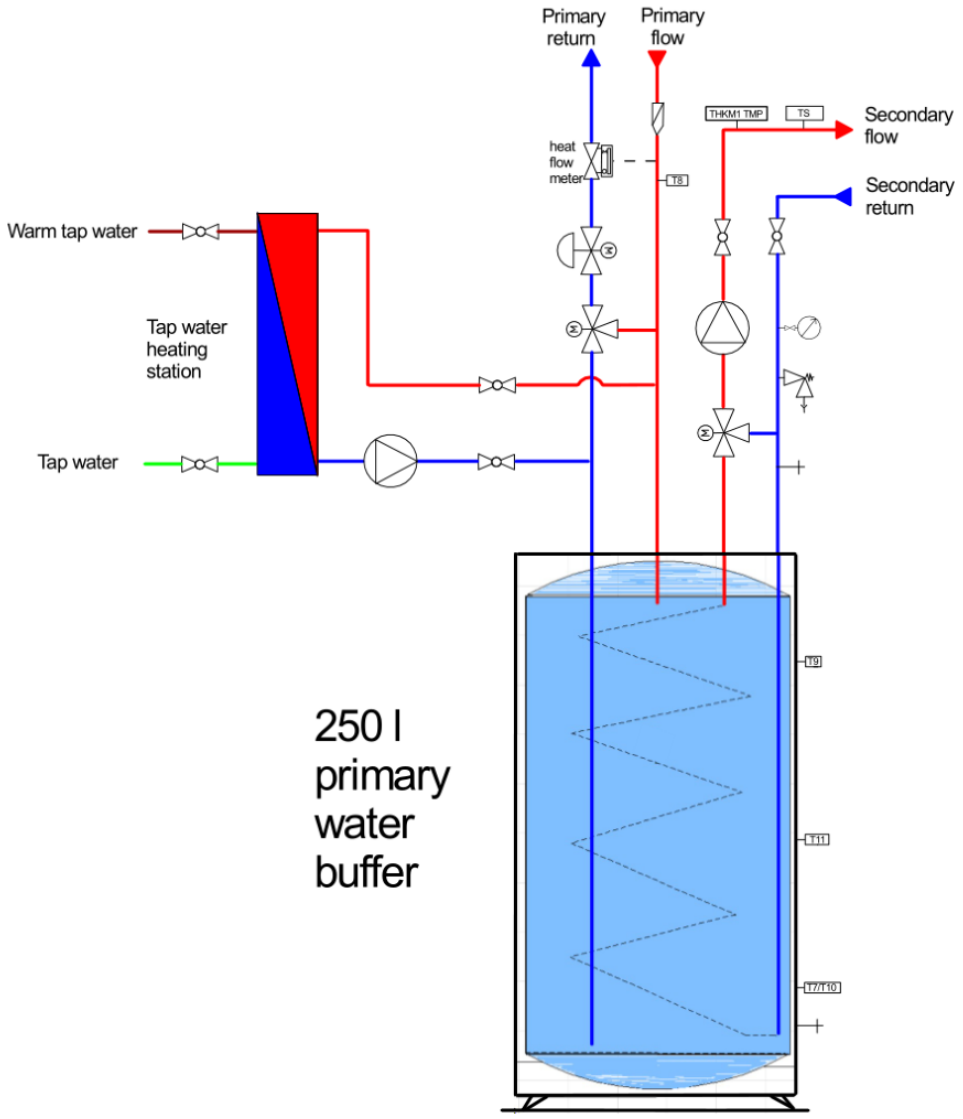

Figure 12: Schematic overview of the connections of the decentralised buffer

One major advantage of the system is that the hot water for the domestic hot water station (for "warm tap water", as it is called here) comes directly from the primary flow, or the buffer top. In each case, the temperature loss of the heat exchanger between primary and secondary side is avoided. This allows for lower heating network temperatures, lowering heat losses across the network.

COSTS OF THE DECENTRALIZED BUFFERS

Various cost factors have to be taken into account, i.e. piping, heat transfer stations (consumer installation), and heating central components cost.

Due to the smaller peak load of dwellings with a decentralized buffer, the pipes may be executed in smaller diameters, and consequently more often as twin pipes instead of two separate pipes. This means less material cost, less installation cost, and less heat losses.

The price difference in pipe material alone sums up to roughly 28.000€ of otherwise 552.000€. Instead of ~3540m of pipes, only ~3330m have to be installed. Heat losses of roughly 18 MWh may be avoided with standard operation parameters.

On the other hand, the installation of standard heat transfer stations is simpler and cheaper than that of the combined unit with buffer (~4.790€ vs. ~5.060€), even when considering the required additional building installation. This difference should become smaller when the production and sales process picks up momentum.

Since peak load is reduced significantly in the heating network with decentralised buffers, and less central storage capacity is required due to the fragmentation of the total heating network water volume, significantly smaller equipment may be installed inside the heating central, possibly even allowing for an overall smaller heating central. Additionally, with smaller main pumps and peak volume flow, pump energy cost may be reduced. Especially during summer nights, the pumps may even be turned off completely when no heat is required. In Windsbach, roughly 50.000€ of the heating central's installation cost could be saved, not even taking into account a smaller heating central.

In summary, while the higher costs for the buffer tank units are roughly compensated by lower heating grid installation costs, the heating central's initial as well as operational cost make a significant difference, rendering a network with decentralized buffers more economical by far.

TEST RESULTS

In the following, rather than showing test results of the buffers in an artificial test environment, some exemplary decentralised buffer storage stations of the Windsbach network and their behaviour during operation are shown.

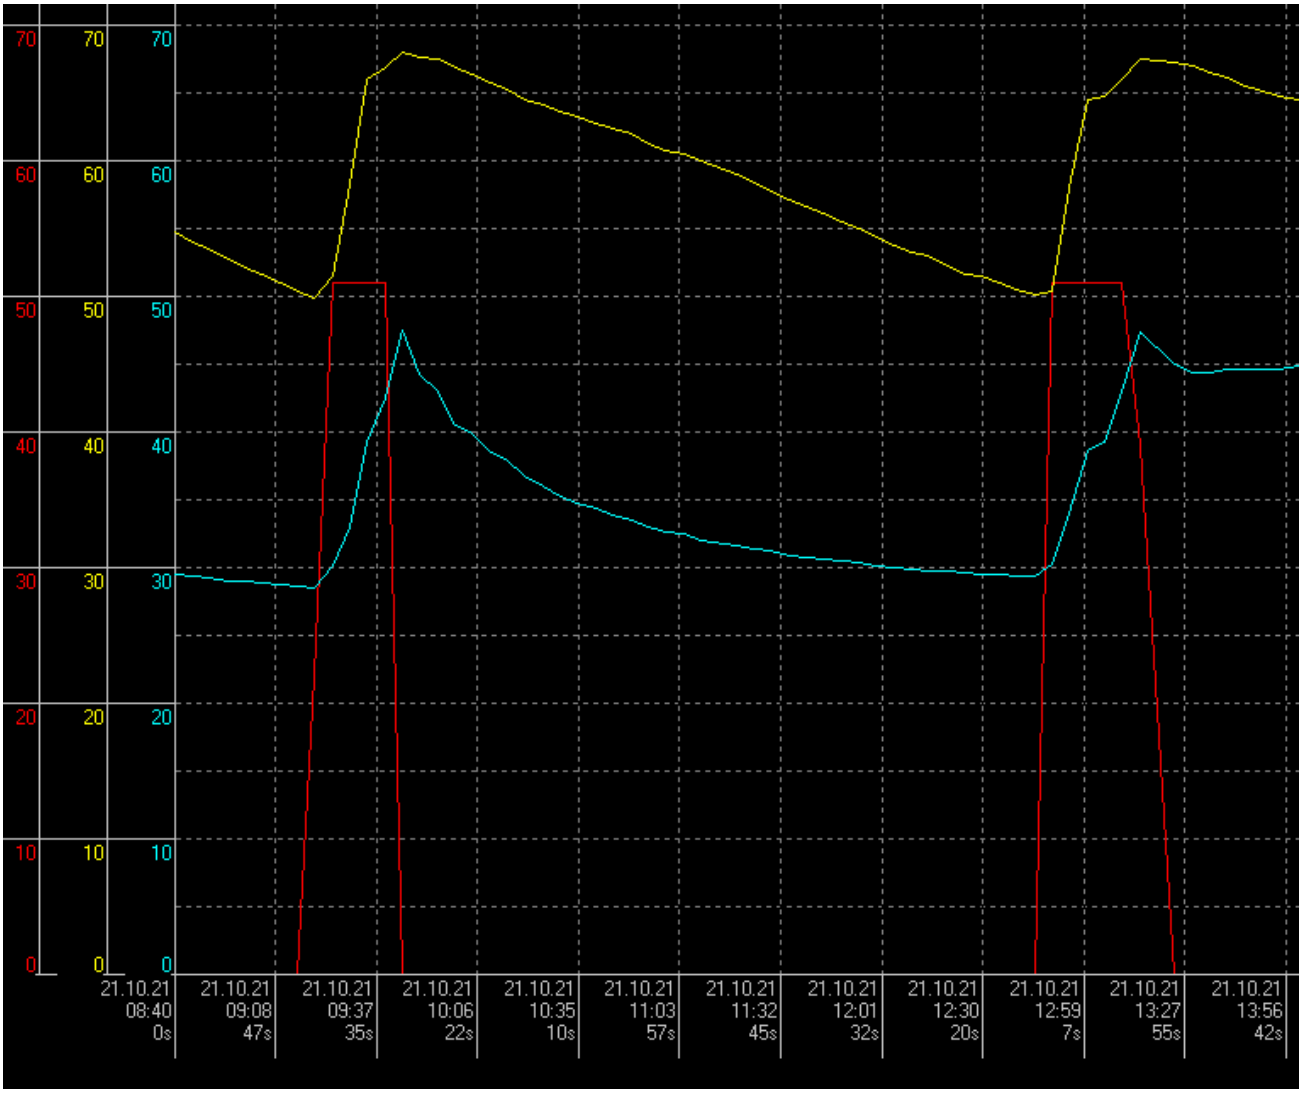

Figure 13: Test results of the decentralised buffer - the degree of opening of the primary valve (red, %), the buffer top temperature (yellow, °C) and buffer bottom temperature (blue, °C).

The temperatures are especially meaningful as long as the valve is open. The maximum district heating water flow for typical single-family homes is limited to 500l/h, which is why the valve will not open to 100%.

Charging starts when both the upper buffer temperature is below 50°C and the lower buffer temperature is below 40°C. Then bottom water will be sent back to the heating network. Following the blue curve while the valve is open, it becomes clear that the return temperature of the water is significantly below 40°C for a large portion of the charging process. It stops when the bottom temperature reaches 45°C (in winter).

KPI COMPARISON

By using decentralized buffers, peak loads in the network could be reduced, thus smaller pipe diameters and consequently more twin tubes with less heat losses could be used. At the same time, the heating network flow temperature could be lowered to between 65°C in summer and 72°C in winter. The lower flow temperature is acceptable for the heating

network thanks to avoiding the temperature loss of a heat exchanger between primary water and domestic hot water stations. Thus, calculated network heat losses could be reduced from ~270kWh/a for a scenario without decentralized buffers to ~249kWh/a. These 21.000kW/h of diffrence equal roughly 7,8% of avoided losses.

Table 4: Overview of efficiency gains due to the individual TEMPO innovations in HT solution package for urban areas

LT solution package for urban areas	Efficiency gain during peak load	Efficiency gain during off-peak load
Decentralised buffers	6 %	6 %
Visualisation tools	3 %	3 %
Smart DH network controller	7 %	15 %
Optimisation of building installation	7,5 % (static) + 10 % (dynamic)	7,5 % (static) + 10 % (dynamic)
Total estimate	**About 25 %**	**About 30 %**

1.4 SMART DH CONTROLLER

BASELINE VERSUS TEMPO IMPLEMENTATION (STATE OF THE ART)

Decentralized buffer control

Decentralized buffers, as explained in the previous section, are not yet widely present in DHNs although they can reduce the network investment cost and network peak loads [1]. Currently limited research exists on the topic of smart control of decentralized buffer charging in DHNs. In [2] for example, the authors focus on smart control of distributed thermal energy storage charging with the objective of increasing the economical profits of operating a CHP. However in TEMPO the objective of charging the decentralized buffers is to further decrease the peak load of the network in order to reduce the use of the gas peak load boilers. In the end network investment costs can also be reduced as smaller backup installations can be installed in new DHNs.

Currently the **baseline** control of decentralized buffers is based purely on the different temperature measurements in the storage. In the Enerpipe demo the buffers contain 3 temperature sensors: one at the top, one in the middle and one at the bottom. The baseline charging behaviour is as follows:

- If (T_{top} < ($Setpoint_{top}$ +$Hysteresis_{top}$) **and** T_{bottom} < ($Setpoint_{bottom}$ + $Hysteresis_{bottom}$))
  - Start charging
- If (T_{bottom} > $Setpoint_{bottom}$)
  - Stop charging

The default setpoint and hysteresis settings are shown in **Table 1**.

Table 1: default setpoints and hysteresis settings

outdoor temperature	setpoint bottom buffer temp	setpoint top buffer temp
°C	°C	°C
+20	55	60
+5	45	60
-10	45	60
Hysteresis	-5 K	-10 K

Concretely this means that when the outside temperature is 20 °C the buffer will be charged according to the following schema:

- Start charging
  - **T_{top} < 50 °C** (60 °C – 10 °C) AND

- T_{bottom} **< 50 °C** (55 °C – 5 °C)
- Stop charging
  - T_{bottom} **> 55 °C**

The buffer is charged at a maximum charging power of 28 kW which slowly decreases when the bottom temperature is reaching its setpoint minus 1K. In this case the controller closes the primary valve steadily to limit the overshoot of the bottom setpoint. As the buffers are only controlled based on their own temperatures, they do not contribute to the optimal operation of the total DHN. In order to use the flexibility of the distributed thermal energy storages for optimizing the overall DHN operation, a new control algorithm was developed in WP1 of TEMPO which is discussed in detail in the Coordinated Buffer Charging section of [3]. In this report however we will focus on how this algorithm was applied on the demo site and the technical challenges we faced.

The **first version of the control algorithm** was developed to exploit the total flexibility available in the decentralized buffers. This could be achieved by sending an optimal charging profile to each individual buffer. To keep the solution scalable, an optimal aggregated profile is calculated taking into account the aggregated flexibility of all the buffers, afterwards this profile is disaggregated over the different individual buffers using the Alternating Direction Method of Multipliers (ADMM), which is a specific type of dual decomposition. It is an algorithm that solves convex optimization problems by breaking them into smaller pieces, each of which are then easier to handle [4]. One of the requirements to implement this approach is the ability of the buffers to follow a power consumption profile, with a resolution greater or equal to 10 minutes. This is the format of the disaggregated power profile calculated by the ADMM algorithm. However when the integration of the communication system in Windsbach was ready (February 2019), we started to do buffer response tests to see if the real system behaves as expected. This is important because one of the most important features in smart control is that the backup controller needs to remain in operation to the end user's comfort. In Windsbach this comes down to the inhabitant having access to hot water at all times. However, when analysing the results of one of these response tests it was seen that comfort violations could occur when sending power charging profiles to the individual buffers, see Figure 14.

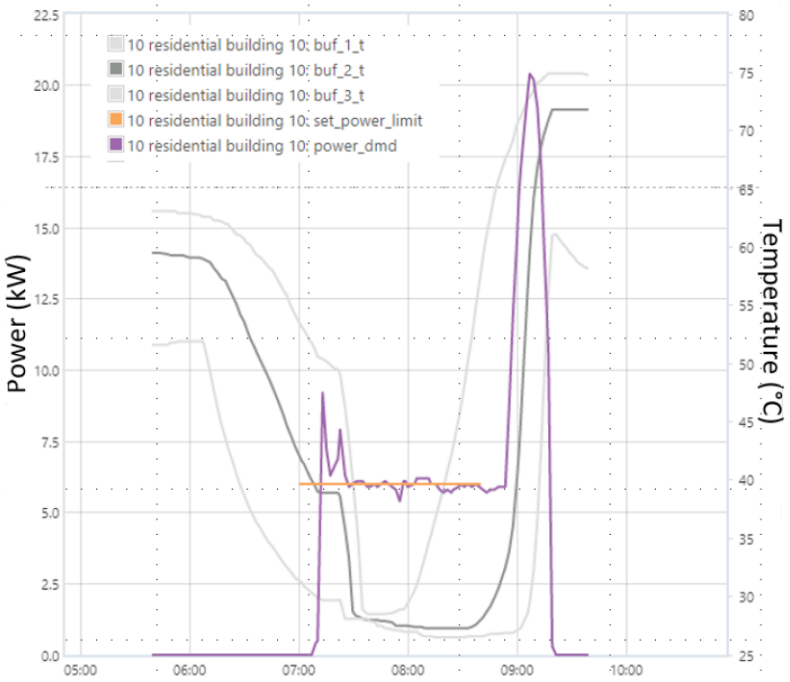

Figure 14: buffer response test with possible comfort violation

The light grey lines represent the top and bottom temperatures of the buffer, the orange line shows the control signal that was applied and the purple line shows the real charging profile that was followed. The goal of this response test was to charge the buffer slowly with 6 kW during 2 hours (7 am to 9 am). However at 7:30 am there is a large temperature drop in the buffer, probably caused by somebody taking a shower. Despite this drop in temperature the buffer keeps on charging with 6 kW and is not able to charge the buffer sufficiently while the control signal is applied. This can be seen by the charging peak at 9 am when the response test ended. It was expected that the backup controller would make sure that the charging power of the buffer would increase to full power around 7:30 am to maintain the comfort. The cause of this behaviour was found in the fact that **the bypass valve of the buffer is also constrained by our control signal**, e.g. the maximum charging power through the bypass valve is also limited to 6 kW when the charging power of the buffer is set to 6 kW. This effect had a major impact on our control algorithm because it cannot be deployed in a system in which the backup control is not able to kick in. As a result, a new way of controlling the buffer had to be devised.

It was found that the only way to avoid limiting the power through the bypass valve is to purely decide on the charging time of the buffer and not on the power, as it can only be charged at full power. This implies that part of the total buffer flexibility is not available to the smart controller because of the inability of the buffer to follow a specific charging profile. The new way of controlling the buffers consists of sending a forced loading signal to each buffer individually, leading to a charging cycle at maximum power. Similar to the baseline buffer charging, some constraints also apply on this forced loading scenario. A buffer will react to a forced loading signal only when the following requirements are met:

- Buffer available for forced loading
  - $T_{bottom} < (Setpoint_{bottom} + Hysteresis_{bottom})$
- Stop forced loading
  - $T_{bottom} > Setpoint_{bottom}$

Specifically, the only difference with the baseline scenario is the fact that the buffer will react to a forced loading signal even if the temperature at the top of the buffer is still above its setpoint. As the baseline control is always active as it serves as the backup controller, the flexibility is limited as follows in case of an outdoor temperature of 5 °C:

- Forced loading possible when:
  - **T_{bottom} < 40 °C** (45 °C – 5 °C)
- Baseline backup controller will take over when:
  - **T_{top} < 50 °C** (60 °C – 10 °C) AND
  - **T_{bottom} < 40 °C** (45 °C – 5 °C)

Therefore the available flexibility window for forcing the buffer to start charging is limited to the time it takes for the top temperature of the buffer to go below the top setpoint + hysteresis (when T_{bottom} < $Setpoint_{bottom}$ + $Hysteresis_{bottom}$). However the measurements show that by the time the bottom temperature is less than 40 °C **it only takes 30 to 60 minutes** before the top temperature drops below 50 °C at which point **the backup controller will start a full charging cycle**. To increase this flexibility window up to about 90 minutes, the bottom hysteresis was set to 0, which advances the point at which the forced charging signal can be applied. The downside of **removing the hysteresis** is the fact that the buffer would only **charge a limited amount of energy (± 2 kWh)** if it would be started immediately when T_{bottom} < 40 °C because the charging will stop when T_{bottom} > 40 °C. Version 2 of the algorithm, which is discussed in the following subsection, uses this forced charging signal to optimize the buffer charging.

CHP optimization

Initially, the purpose of the smart controller was focused solely on charging the decentralized buffers but during consultations with the DHN owner it became clear that it would be fairly straightforward to add control of the central CHPs which produce the vast majority of the heat in the network. All hardware was already in place to control the CHPs. The owner was receiving the day ahead price tariff at which the produced electricity could be sold and a control box was installed which is able to control the CHP based on a given production plan.

The baseline control of a CHP in a district heating network is mostly based on a simple rule system taking into account the heat load of the network and the SoC of the central buffers, if available. In the Windsbach demo however the baseline control is different as the network owner already implements a type of manual smart control in which he checks the day-ahead price tariff and, based on his experience, defines a plan for the CHP to operate during high electricity prices taking into account the expected heat load of the DHN. The owner executes this task twice a day updating the plan for the next 24 hours. As explained in [3] and [5] a mixed-integer linear programming (MILP) approach was used to automate this cost optimized operation of the CHP, sending a mail twice a day to the network owner with the optimal plan. Version 1 calculated an optimal plan for the only commissioned CHP available in the network with a thermal capacity of 215 kW. These mailings have started in June 2020, however in the winter season of 2020-2021 the total heat load of the network exceeded the CHP capacity causing the CHP to operate at full power 100% of the time. In case of excess heat load, it was provided by the gas peak boilers. When the CHP is operational 100% of the time, no flexibility is available and thus no mails are sent.

IMPROVEMENTS TO STATE OF THE ART IMPLEMENTATION

Decentralized buffer control

As discussed above, the new version of the control algorithm only calculates the charging time of a buffer, which is a major difference to the complete charging profile with varying power that was calculated in version 1. In order to switch to this new smart control scenario, a completely data driven approach was implemented. This Reinforcement Learning (RL) approach is explained in detail in the Coordinated Buffer Charging section of [5]. Data-driven RL control algorithms have already been applied in charging thermal storages in buildings but mostly in the context of an electric grid. In [6] the authors state that RL control is favourable over baseline control in terms of cost savings, when used in an electrical grid, but it does not reach the level of optimal predictive control (which was the concept of version 1 in TEMPO). A recent research paper [7], which also focuses on electricity consumption, describes a deep RL algorithm which was able to reach a decrease of peak demand up to 12% and achieving an average peak-to-average ratio by 6%. In version 2 of TEMPO a similar RL approach is deployed in a DHN to lower the thermal peaks in the network.

In this approach an agent will take into account the State Of Charge (SoC) of each buffer and combine these states into one big state space. Based on this state space it will decide which of the buffers in the DHN should start charging in the next 10 minutes. Afterwards the agent will be rewarded/penalized for the decisions that were taken. If the actions have caused a high peak demand, they will be penalized. If the actions have caused the backup controller to kick in, they will also be penalized. These two penalty terms will make sure that the controller will not charge a buffer too early or too late. The most important feature of this approach is that the agent learns from each decision that it has made, which makes the controller self-learning. The longer the controller will be operational in the system, the better it learns the responses of the distributed buffers.

Testing of this controller started in June 2020, first new buffer response tests were executed to guarantee that this type of control could not lead to comfort loss. The results of these tests were satisfactory and the tests of the new algorithm started in November 2020. In this period it became clear that the buffers were not responding well to the control signals that were sent from our controller. It was a difficult task to pinpoint why the control signals applied correctly at some times and not at other times. In the end a bug in one of the two underlying communication platforms was found. These two platforms, EnergyView and Schneid, are used to send the control signals from the algorithm all the way to the installed decentralized buffers. Where Schneid is the conventional controller, which is running as the main operation software in the plantroom of the DHN, and EnergyView is the API and visualization tool that is used on top of the Schneid software to enable the integration of the smart controller. The bug consisted of the optimal charging profile not being sent from EnergyView to the Schneid system at the desired times but only when the schedule from the smart controller was sent to the EnergyView API. This resulted in unpredictable charging behaviour of the buffers. By February 2021 this bug was fixed and the testing restarted. Yet another problem surfaced as the buffers were now forced to start a charging cycle but they were reacting very slow to this signal with an average delay of 20 minutes. A live response test, in which Enerpipe, Noda and VITO were involved, was executed to follow the charging cycle of a buffer all the way from sending the control signal to EnergyView, from EnergyView to the Schneid system and further on to the buffer. It

showed that a slow opening charging valve led to the delay in the buffer response. In consultation with Enerpipe more aggressive valve settings were defined in March 2021. The valve opening times were changed from 15 seconds to 5 seconds and from 20 seconds to 10 seconds. After this change, the maximum charging power was reached within 10 minutes after sending the control signal.

When all the technical issues were solved, the control algorithm could be tested but unfortunately the heating season was almost over. The RL peak shaving algorithm was still tested in March 2021 but it was applied on 20 buffers only. The disadvantage of this RL approach is the fact that the state space grows exponentially when adding extra buffers to the system. In order to make the solution scalable, adaptations need to be made to the algorithm to make sure the state space and thus calculation time does not explode. These adaptations were made over the summer period of 2021 and are now ready to be tested during the 2021-2022 heating season.

CHP optimization

The improved version of the CHP optimization consisted of changing the algorithm to enable the integration of the newly commissioned CHP with an electrical power of 400 kW. It is expected that this CHP will be able to cover the total heat load of the network during the heating season of 2021-2022. Since half October 2021 the mails are adapted to include the plan for the 400 kW CHP.

COSTS OF THE SMART DH CONTROLLER

The costs of the smart DHN controller consist only of development costs of the algorithm:

- VITO
  - Developing the coordinated buffer charging algorithms (work done in WP1)
    - Estimated: 70 k€
  - Adapting the coordinated buffer charging algorithms to the real life system in Windsbach with its specific measurements and control options (WP3)
    - Estimated: 40 k€

TEST RESULTS

Decentralized buffer control

The test results of all the different tests that have been executed up until now are summarized in Table 2. The most important results and their impact has been discussed in the above sections.

Table 2: Test results smart distributed buffer charging

start	end	type	Action	Details
2019-03-16	**2019-07-22**	baseline	Baseline spring/summer 2019	
2019-07-23	**2019-08-08**	fault	Outlier period	Test period in which Enerpipe was experimenting to increase the capacity of the decentralised buffers. However, this led to overloading of the buffers, causing too high return temperatures
2019-08-09	**2020-01-15**	baseline	Baseline autumn/winter 2019	
2020-01-16	**2021-01-16**	fault		Period with high impact of buffer faults, filtered out of analysis
2020-01-28	**2020-01-28**	baseline	Measurements logged in EnergyView	Windsbach successfully connected to EnergyView
2020-02-11	**2020-02-11**	test	First buffer response test	See appendix A – First buffer response test (11/02/2020)
2020-03-19	**2020-03-19**	test	First test of setting set_power_limit through EnergyVille API	set_power_limit successfully set but no reaction from buffer
2020-03-31	**2020-03-31**	test	1st charging profile tests (10 min: 15 kW, 10 min: 5 kW, 10 min: 10 kW)	See appendix A – Charging profile test (31/03/2020)
2020-04-09	**2020-04-09**	test	2nd charging profile tests (20min: 10 kW, 20min: 3 kW, 20min: 6kW	See appendix A – Charging profile test (09/04/2020)

2020-04-15	**2020-04-15**	test	3rd charging profile tests (20min: 20 kW)	See appendix A – Charging profile test (15/04/2020)
2020-04-17	**2020-04-17**	test	4th buffer response tests (3 hours: 6 kW) for buffer 10 (COMFORT VIOLATION DUE TO LIMITING BYPASS VALVE)	See appendix A – Charging profile test (17/04/2020)
2020-05-22	**2020-05-25**	test	6th buffer response test, first test with new control signal, always charge at max power limit. Not possible to define charging power anymore. (10min: max power")	See appendix A – New control scenario test (25/05/2020)
2020-06-04	**2020-06-06**	test	7th buffer response test, second test with new control signal (20min: max power)	See appendix A – New control scenario test (06/06/2020)
2020-11-25	**2020-12-22**	test	First reinforcement learning algorithm test	Results not representative due to bug in EnergyView API, control signals were not sent according to the given schedule
2021-01-04	**2021-01-11**	test	Second reinforcement learning algorithm test	Results not representative due to bug in EnergyView API, control signals were not sent according to the given schedule
2021-01-11	**2021-01-15**	fault	No data, communication lost	
2021-01-15	**2021-01-22**	test	Third reinforcement learning algorithm test	Results not representative due to bug in EnergyView API, control signals were not sent according to the given schedule
2021-02-03	**2021-02-08**	test	Test EnergyView API bugfix, set_power_limit_schedule	Bug fix worked as expected, set_power_limit_schedule is copied to set_power_limit according to the given schedule

2021-02-19	**2021-03-11**	test	First reinforcement learning algorithm test with fixed API	Slow buffer response to control signals (+- 20 minutes delay)
2021-03-17	**2021-03-17**	test	Live test low lever buffer control to identify slow buffer response	See appendix A – Aggressive valve setting test (17/03/2021)
2021-03-18	**23/03/2021**	test	All 20 buffers are configured with new aggressive valve setting and RL algorithm can schedule already the next 5 minutes	
2021-03-26	**2021-03-27**	test	All 20 buffers are configured with new aggressive valve setting and RL algorithm can schedule already the next 5 minutes	Some unexpected results were seen as buffers were only controlled at multiples of hours, whereas they could be controlled at multiples of 10 minutes. This behaviour is analysed and bug fixes were made over Summer 2021
2021-10-01	…	test	Tests peak shaving algorithm extension from 20 buffers to all buffers	

CHP optimization

As discussed above, the DHN owner is responsible for deciding on the production plan that the CHP needs to follow. We will advise the owner with our optimal plan by sending an updated plan for the next 24 hours twice a day (at 8h45 and 16h45 CET). Figure 15 shows the advised optimal plan in blue, together with the followed plan in green. The followed plan deviates from the optimal plan around 20h00 at which the CHP stops operating earlier than expected. This can also be seen in the deviation of the real state of charge of the central buffer unit (red dots) and the estimated SoC of our optimal plan (orange).

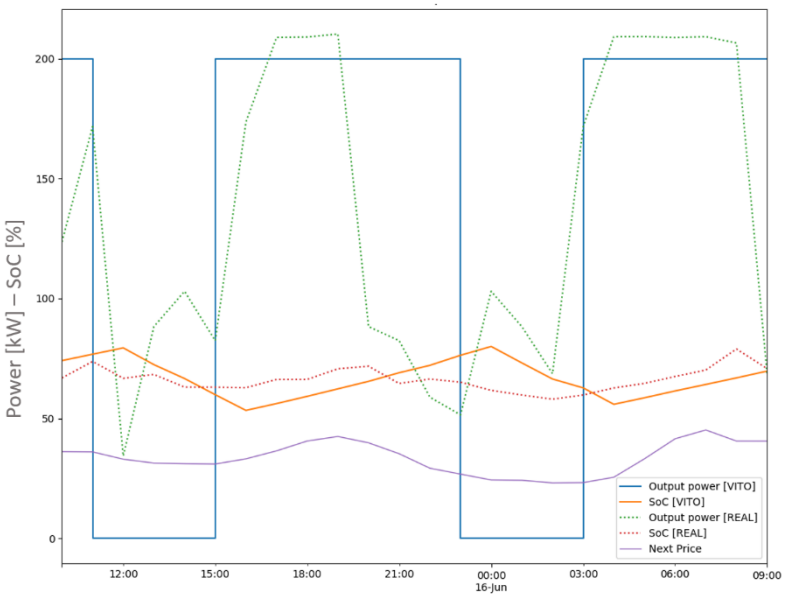

Figure 15: example of CHP optimization (blue lines is the optimal plan and green is the followed plan)

Looking at the period of July 2020 – October 2020, **the average electricity price** at which the produced electricity could be sold was **36,1 €/MWh**. If the **optimal plan** was followed correctly throughout this period, **the average selling price** of the produced electricity was **43,5 €/MWh**. Looking at the spread of the price in Figure 16, it shows that 75% of the times the electricity price is lower than 43,5 €/MWh.

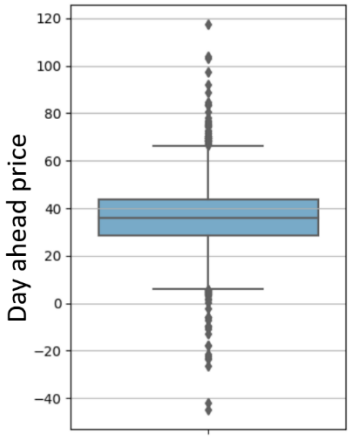

Figure 16: spread of the day-ahead electricity prices (July 2020 - October 2020)

However, due to the increasing numbers of connected buildings to the DHN, the heat load of the network exceeded the capacity of the CHP from November 2020 onwards. Therefore the CHP was running constantly, backed up with the peak boilers. As a result, there was no longer any flexibility available to optimize the CHPs operation.

During summer 2021 the heat load of the network, consisting of DHW only, was fully provided by the satellite site and the CHP was switched off most of the time. Since October 2021 the second, and bigger (400 kW_{th}), CHP is commissioned and it is expected that it will be able to deliver all the heat in the network at all times, apart from some extreme weather conditions. VITO changed the optimization algorithm to take into account this new CHP and since 14 October 2021, mails are being sent containing the optimal plan for the

400 kW_{th} CHP. In addition to that, the network owner now also feeds back to VITO the actual plan that is defined in the system every day. Due to the rising energy prices, it is expected that the potential profit increase of selling the produced electricity will be higher during this heating season. In the end, the optimizer not only increases the revenues of the electricity sales, it also unburdens the DHN owner as the plan is automatically calculated and even submitted if desired. Furthermore it reduces the maintenance costs of the CHP by minimizing the number of ON/OFF cycles.

KPI comparison

Due to the numerous technical challenges that were faced during the different buffer test scenarios, it is not possible to do a correct KPI comparison at this moment. The controller is focused on flattening the total heat load profile of the network, therefore it is important that the algorithm would be able to run for at least 3 weeks during a heating season to analyse its impact on the system. This will be done during the upcoming heating season of 2021 – 2022. The controller will be able to increase the efficiency in a network by shaving/shifting the heat load peaks to off-peak times which ensures that the less efficient gas peak boilers do not have to be used. An efficiency gain of 7% during peak load times is to be expected case the CHP is not over dimensioned, otherwise the gas peak load boilers are obsolete. However, the controller will not impact the efficiency in the network in off-peak times and thus the efficiency gain of 15% will not be reached in these off-peak ranges. The results and impact on the KPIs of the smart controller will be discussed in D5.3.

1.5 OPTIMISATION OF THE BUILDING INSTALLATIONS

BASELINE VERSUS TEMPO IMPLEMENTATION (STATE OF THE ART)

There does not appear to be a standardized way of dealing with the optimization of building installations in district heating systems. According to a recent survey with Swedish utilities [8], a variety of methods are used to detect anomalous substations, including monthly checks of customer billing data and the analyses based on return temperature, flow and overflow. The same study highlighted the importance of having physical access to the customers' installations, which is done either by signing service agreements or by including yearly inspections in the district heating pricing. The most advanced approaches proposed in the state of the art (e.g. [9]) are usually based on the evaluation of automatic meter readings, for instance with clustering and threshold methods. These methods often make it possible to identify suboptimal behaviour, but not to assess or predict its causes, or to provide indications with regards to further diagnosis and resolution steps.

Compared to the state of the art, one distinctive feature of the TEMPO innovation was a general investigation of situations leading to high return temperatures and their modelling in coupled building and system simulations. This allowed the training of fault detection and diagnosis algorithms (based on simulation data) on the one hand, and the development of a guide for technical audits of building installations on the other hand.

IMPROVEMENTS TO STATE OF THE ART IMPLEMENTATION

Several improvements were made in version 2 as compared to version 1, including continuous improvements in terms of simulation models, calibration of simulations and adjustments to the practical guide, but the main improvements were as follows:

- Fault detection and diagnosis with new algorithms and for shorter periods.
- Interactive guide for technical audit of building installations: Whereas in version 1 the practical guide for technical audit of building installations was a static (PDF) document, in version 2 it has become a more interactive guide in the form of a web page. This web page makes the exploration of the guide easier by providing a variety of links.
- Connecting fault predictions with the guide: Having a web version of the guide for technical audit of building installations made it possible to establish a connection between two outcomes of the innovation which were not yet connected in version 1: the fault predictions made for given data and the guide. In version 2, the fault predictions resulting from the application of the simulation-trained fault detection and diagnosis algorithms to actual monitoring data are summarized in reports. These reports refer to fault types and metrics for which links to the guide are provided, enabling a combination of the (automated) data-driven insights with the (manual) audit process.

FUNCTIONALITY

The innovation "optimisation of the building installations" combines the two related functions illustrated in Figure 17, which are the simulation-based fault detection and diagnosis for secondary systems and the guide for auditing building installations.

Simulation-based fault detection and diagnosis is based on data-driven algorithms which have been trained on simulation data (from detailed simulations carried out in Modelica) and are applied on monitoring data from the individual substations. The predictions made by the algorithms for a given period are summarized in an HTML document (Figure 17, left) along with metrics calculated for the period. This document also includes links to relevant sections of the guide for auditing building installations (Figure 17, right).

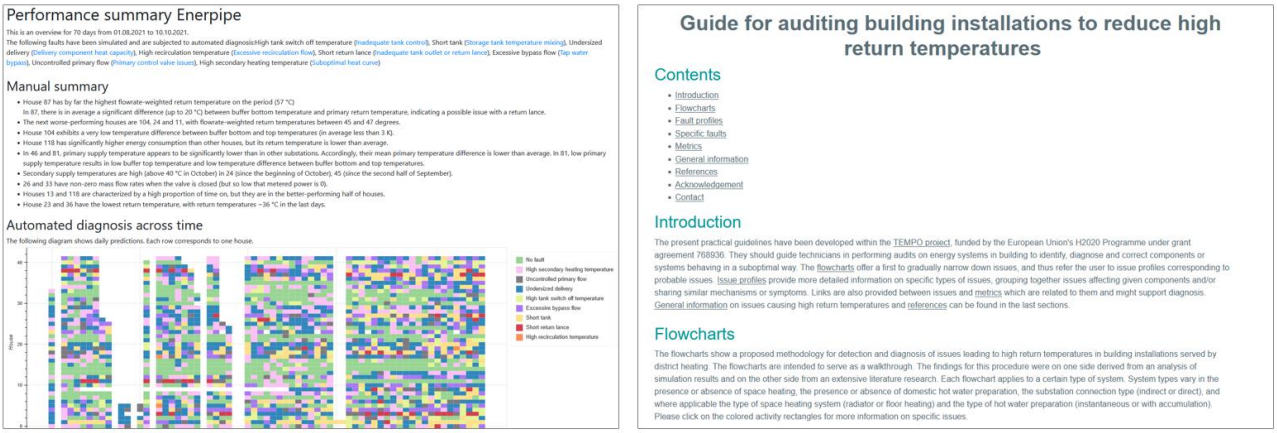

Figure 17: Screenshots of a fault detection and diagnosis summary (left) and of the interactive guide for auditing building installations (right).

COSTS OF THE INNOVATION

The costs of the innovation "optimization of the building installations" within the TEMPO project amounted to approximately 300 hours[1]. Given the unusual substation type used in the Enerpipe demonstrator, the preparation of a detailed simulation model including controls proved to be time-consuming. As summarized in Table 3, it can be expected that the effort incurred during the TEMPO project could be greatly reduced if the innovation is applied in a similar network. Not only does the learning curve mean that activities could be carried out more efficiently in further networks, but also the effort for simulation model development would not have to be repeated, as long as the installed systems remain the same.

Table 3: Estimated efforts for the innovation "optimization of the building installations" for the Enerpipe demonstration network in Windsbach within the TEMPO project and for a similar network (reusing TEMPO developments).

Activity	Effort in TEMPO in hours	Estimated effort in further similar network in hours
Development of detailed simulation model	90	0
Calibration and reruns of simulation model	60	20
Training fault detection and diagnosis algorithms	30	15

[1] Approximately the effort of AIT in WP3 (2 person months)

Inference with the fault detection and diagnosis algorithms	40	15
Adapting the guide to insights from the demonstrator	40	5
Result evaluation and communication	40	15
Sum	300	70

TEST RESULTS

For the innovation to be evaluated in terms of accuracy of detection and diagnosis accuracy, true fault labels are required. Since true fault labels require audits in the building installations, they could only be obtained in limited numbers. The most clear-cut faults which occurred during the project duration were the faults related to excessively short return lances, which occurred especially in the second half of 2020 and were solved for the most part in January 2021. In this case, the simulation-based fault detection and diagnosis algorithms were to a large extent successful at detecting malfunction. As shown in the next section, identifying and remediating faults in building installations can result in very significant performance improvements.

KPI COMPARISON

The summary of key performance indicators before and after fixing the issue of short return lances provided in Table 4 for five building installations shows that the performance improvement after diagnosing and fixing a fault can be substantial, both in terms of energy demand (30 % reduction) and in terms of mean return temperatures (more than 10 K lowering). These high numbers are mitigated by the fact that not all installations are affected by faults, and that other fault types generally have more subtle impacts. Finally, it is difficult to ascribe improvements to the innovation alone, as faults were also discovered by more traditional means including manual data inspection.

Table 4: Key performance indicators before (from beginning of recording in 2019-2020, depending on house, to January 15th 2021) and after (January 19th 2021 to October 15th 2021) fixing issue of short return lances in five houses.

House	58	70	71	73	109	Average
Performance before fix - mean power in kW	2.48	1.71	1.78	3.11	2.62	2.34
Performance after fix - mean power in kW	1.45	1.40	1.42	1.65	2.14	1.61
Performance difference in %	-41.8	-17.7	-20.2	-46.9	-18.3	-31.1
Performance before fix - mean return temperature in °C	64.50	53.67	45.63	44.78	56.26	52.97
Performance after fix - mean return temperature in °C	40.16	47.09	38.51	40.34	37.36	40.69
Performance difference in K	-24.3	-6.6	-7.1	-4.4	-18.9	-12.3

Table 4: Overview of efficiency gains due to the individual TEMPO innovations in HT solution package for urban areas

LT solution package for urban areas	Efficiency gain during peak load	Efficiency gain during off-peak load
Decentralised buffers	6 %	6 %
Visualisation tools	3 %	3 %
Smart DH network controller	7 %	15 %
Optimisation of building installation	7,5 % (static) + 10 % (dynamic)	7,5 % (static) + 10 % (dynamic)
Total estimate	**About 25 %**	**About 30 %**

IMPACT OF THE RESULTS ON THE OVERALL PROPOSED SOLUTION

By contributing to the elimination of faults in building installation, the innovation "optimization of the building installations" contributes to the goals of lowering return temperature and increasing energy efficiency. It also appears to be complementary to the other innovations, as faults in building installations compromise the ability to apply smart controls and other advanced strategies.

2 CONCLUSIONS

Smart DHN controller

By applying the theoretical algorithms that were developed in WP1 into the Windsbach demo of WP3, we encountered many practical and technical challenges. At first, the optimization algorithm needs to be controlling a system which is already operated with a conventional controller, e.g. the Schneid controller in Windsbach. The behavior of this controller was different than the assumptions made in WP1, e.g. the backup controller did not safeguard end-user comfort if the smart controller was sending out non-optimal control signals. Due to this, a new optimal control approach, purely data-driven, was defined, implemented and deployed. In order to get the calculated control signal of this approach applied on a distributed buffer, four different platforms are used. The signal gets calculated on a VITO server, where it sent to the EnergyView platform. In the EnergyView platform another piece of software is running to export the control signals out of EnergyView, at the defined times, and send them to the Schneid system. The Schneid system in the end sends the charging signals to the buffers. This workflow implies that there are four different places where errors can occur and each of the systems has its own expert and responsible person. In order to debug the whole system, different people are needed. As a result, in the majority of the cases it takes longer to locate the problem than to solve it. In that sense, it would be better if the conventional controller is able to receive external control signals and apply them directly to the devices in the system (if comfort is not jeopardized). This would rule out two of the four systems, doubling the debugging efficiency. In an ideal world, the installed system together with its operating software, would be designed in cooperation with the company responsible for deploying the smart controller. However, in the majority of the cases this is not feasible as the DHN and its controller are already in operation at the time the smart controller is defined. Therefore, the 20%-80% rule certainly applies when deploying a control algorithm in practice, e.g. 20% of the time is needed to develop the algorithm and 80% is needed to get it up and running in the field.

Decentralized buffers

Installing decentralized buffers in each dwelling allowed for a comparatively slim design of the district heating network. The first two phases of the development area are completed to a large degree, and there seem to be no signs for supply shortages. The network flow temperature could be reduced significantly, from 80°C in a baseline-scenario to 65-72°C. However, the original concept behind the heating network with the centralized buffers aimed for a significantly lower return temperature. There are a couple of issues increasing said return temperature, e.g. badly manufactured short return lances, heating systems with elevated temperatures inside several dwellings, and a rather remote multi-family house requiring an elevated temperature level for its apartment stations, sending back high temperature. The same multi-family house requires high network flow temperatures, that could otherwise be further reduced. Nevertheless, the return temperature would be even higher in a network with secondary instead of primary storage buffers.Several consumers with flawless performance (including low return temperatures) show that the initial assumptions might be met (or at least approached) in a network with reasonably designed and executed building installations.

REFERENCES

[1] J. M. Jebamalai, K. Marlein and J. Laverge, "Influence of centralized and distributed thermal energy storage on district heating network design," *Energy,* vol. 202, no. 0360-5442, p. 117689, 2020.

[2] D. Vanhoudt, . B. Claessens , R. Salenbien and J. Desmedt, "The use of distributed thermal storage in district heating grids for demand side management," *ArXiv,* vol. abs/1702.06005, 2017.

[3] D. Vanhoudt, A. Soares, G. Suryanarayana, D. Geysen and Van Oevelen Tijs, "D1.1 Report innovation installation guidelines 1st version," TEMPO, 2019.

[4] S. Boyd, N. Parikh, E. Chu, B. Peleato and J. Echstein, "Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers," *Foundations and Trends in Machine Learning,* vol. 3, no. 1935-8237, pp. 1-122, 2011.

[5] D. Vanhoudt , C. Johansson, R.-R. Schmidt, M. Abt, D. Geysen, C. Hermans, L. Scapino and T. Van Oevelen, "Report innovation installation guidelines 2nd version," VITO, 2021.

[6] G. P. Henze and J. Schoenmann, "Evaluation of Reinforcement Learning Control for Thermal Energy Storage Systems," *HVAC&R Research,* vol. 9, no. 3, pp. 259-275, 2003.

[7] G. Pinto, M. S. Piscitelli, J. R. Vazquez-Canteli, N. Zoltán and A. Capozzoli, "Coordinated energy management for a cluster of buildings through deep reinforcement learning," *Energy,* vol. 229, no. 0360-5442, p. 120725, 2021.

[8] S. Månsson, P.-O. J. Kallioniemi, M. Thern, T. Van Oevelen and K. Sernhed, "Faults in district heating customer installations and ways to approach," *Energy,* 2019.

[9] H. Gadd and S. Werner, "Fault detection in district heating substations," *Applied Energy,* 2015.

[10] B. Windholz, D. Basciotti and C. Marguerite, "Deliverable 5.1: Generic monitoring principle and equipment for TEMPO," TEMPO, 2018.

[11] T. Nuytten, B. Claessens, K. Paredis, J. Van Bael and D. Six, "Flexibility of a combined heat and power system with thermal energy storage for district heating," *Applied Energy,* vol. 104, no. 0306-2619, pp. 583-591, 2013.

APPENDIX A – BUFFER CHARGING TEST RESULTS

1. First buffer response test (11/02/2020)

Following the completion of the transfer of measurement data from Windsbach to NODA Sweden, a test to confirm the stations‘ reaction to switching impulses was run.

On the weekend before it was tested if the stations react when their are NOT activated with the proper controller setting. These tests showed that switching has no effect if the settings in the stations are set only to „volume limit“ (as one would expect). For the actual tests these settings were corrected to their proper value (code „3“ for combined volume- and load-limit).

The tests were carried out on Monday, Feb. 10th 2020, by Guenther Beck and were supervised from the grid side by Andreas Fiegl (Enerpipe).

The time slot was 13:20 to 16:30.

Method of switching

The switching impulses were simulated from a separate software tool that ran on the server in Windsbach.

Each run used an input file that contained the station number, the starting time and end time for each switch, as well as a (fictional) offset [K] and one of the DHW-options (automatic, forced loading, blocked loading).

The buildings were switched in a sequence of 1 building each minute. Each duration was 13 minutes, so that a wave of buildings formed with max. 14 active at the same time.

The correct and timely switching as well as the automatic reset after expiry of the duration was observed.

Four runs were made with different settings:

1) an offset was set that led to 0% of load limit in combination with forced loading:
    a. the values were correctly transferred to Winmiocs and then to the stations
    b. the stations did not react on the dhw-setting, because the load limit was 0% and this has priority
    c. stations that had a load at the time of switching, were gradually reduced to 0%
2) a (negative) offset was set that led to 100% of load limit in combination with forced loading:
    a. the values were correctly transferred to Winmiocs and then to the stations
    b. the stations did correctly react on the dhw-setting, because the load limit was 100%
3) an offset was set that led to random limits of 20-40% of load limit in combination with forced loading:
    a. the values were correctly transferred to Winmiocs and then to the stations
    b. the stations did react on the dhw-setting, but the load was kept below the limit (as it should be)

   c. stations that had a higher load at the time of switching, were gradually reduced to the limit

4) an offset was set that led to 100% of load limit in combination with load blocking:
   a. the values were correctly transferred to Winmiocs and then to the stations
   b. it was not easy to see if the blockade actually had an effect or if the lack of buffer loading was due to the fact that buffers were still full from the previous operation. But in any case no violation of the load blockade could be observed. Ideally such a test should take place when there is a typical demand for loading, e.g. between 5:00 and 7:00 am.

Results / Summary

The results confirmed that the data chain to the stations works as expected and that the stations react to the settings as expected.

The run 2) triggered a peak load at node 1901 (total grid) with

- peak load of 361 kW (285 kW as 15 min avg.)
- peak volume flow of 10.296 l/h (8796 l/h as 15 min avg.)
- lowest return temperature primary of 44°C (45°C as 15 min avg.)

The peak of run 2) can be seen in the graph below.

Run 1) and run 4) can be seen as lows in the demand, that are stronger than the lows of the previous days.

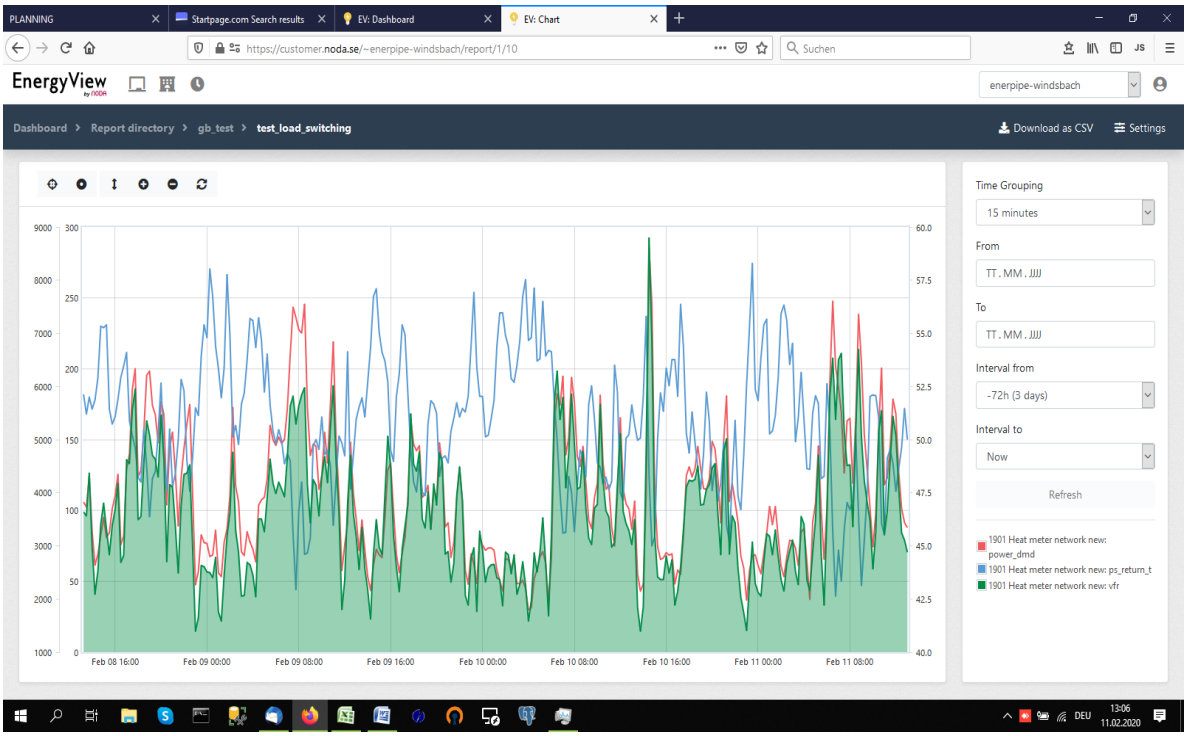

NOTE: the report can be found in EnergyView under: Dashboard/Report directory/gb_test/test_load_switching

Recommendation for data format of switching impulses

It seems to be the most straight way to send directly the two components that can be switched:

- the % of load limit or its equivalent in absolute kW
- the desired DHW program to switch

In this case it must be observed that the load limit has priority in case there is a conflict between the two settings.

2. Charging profile test (31/03/2020)

During the charging profile tests, specific predefined charging profiles are sent to the buffers when the following requirements are met:

- T_{top} < 56 °C
- T_{bottom}< 40 °C

In all the following charging profile tests, 2 graphs of each buffer will be shown:

1. Graph with original values from EnergyView (no resampling)
    a. Set_power_limit: the power limit that has been set by VITO [kW]
    b. Power_dmd: the measured power demand [kW]
    c. Avg_temp: the average of the 2 or 3 measured temperatures per buffer [deg C]
2. Graph with resampled values from Energyview (10min resampling using means)
    a. Set_power_limit: the power limit that has been set by VITO [kW]
    b. Power_dmd: the measured power demand [kW]
    c. Avg_temp: the average of the 2 or 3 measured temperatures per buffer [deg C], this is NOT resampled because it doesn't clutter the graph

In this test the following charging profile was sent to the buffers:

- 0 - 10 min: 15 kW
- 10 – 20 min: 5 kW
- 20 - 30 min: 10 kW

Analysis based on these values show that:

1. Group 1:
    a. None of the sites was able to follow the profile in a +- correct way, especially the first 15kW was never reached (in most cases not a lot happened when setting the 15 kW power limit)
    b. Site 11 and site 18 didn't react
    c. Temperatures in all buffers was quite high +50 deg C (except for site 11) but this cannot be the reason why they weren't loading at 15 kW because in other groups buffers with these temperatures did.
2. Group 2:
    a. None of the sites was able to follow the profile in a +- correct way, especially the first 15kW was never reached (in most cases not a lot happened when setting the 15 kW power limit). Site 23 followed the first limit but this is probably because it was charging at that moment (to be checked).
    b. Site 22, site 29, site 30 didn't react. Site 22 did something but the profile looks constant.
3. Group 3:

   a. Site 38 seems to follow the profile quite well (even the first 15kW), it had a low temperature (<40 deg C). To be checked if this site started loading just before our signals came
   b. Site 35 did not react.
4. Group 4:
   a. None of the sites was able to follow the profile in a +- correct way, especially the first 15kW was never reached (in most cases not a lot happened when setting the 15 kW power limit)
   b. Site 94, site 104 and site 24 didn't react
   c. Although site 48 and site 50 had low temperatures, they didn't react to the first 15 kW setting
   d. Site 105 seems to react faster than other sites but it has a 70kW connection capacity vs the normal 30kW connection capacity.

2 things stand out:

1. No site followed the profile (but maybe some would have been able if they started with a lower temperature)
2. Site 11, 18, (22), 24, 29, 30, 35, 94 ,104 didn't react

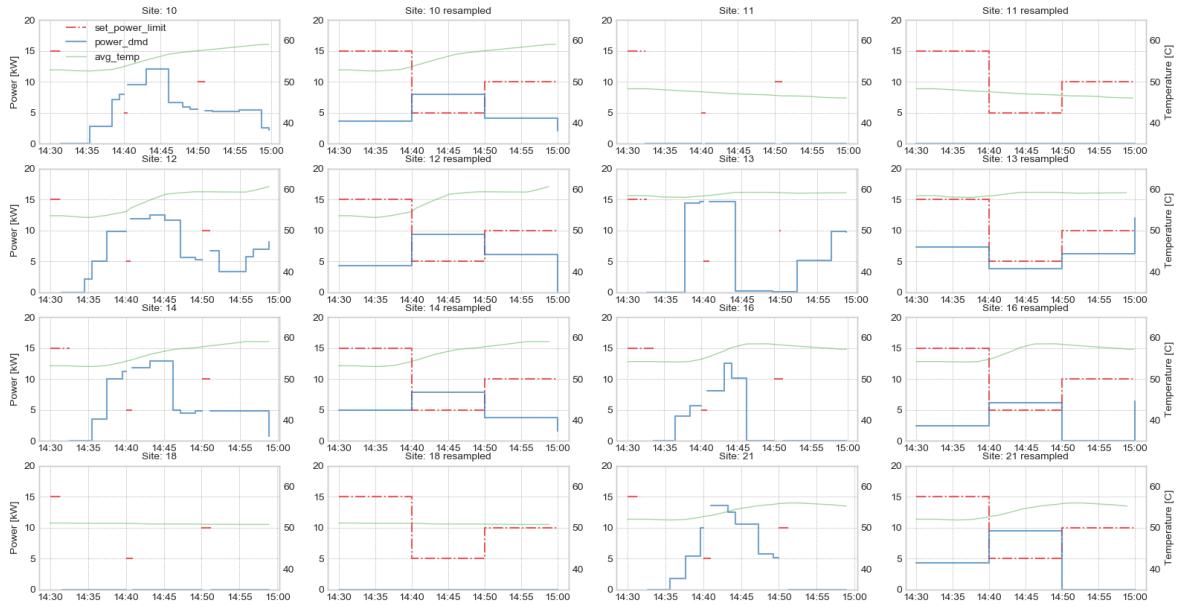

Figure 18: overview of test in group 1 (10, 11, 12, 13, 14, 16, 18, 21)

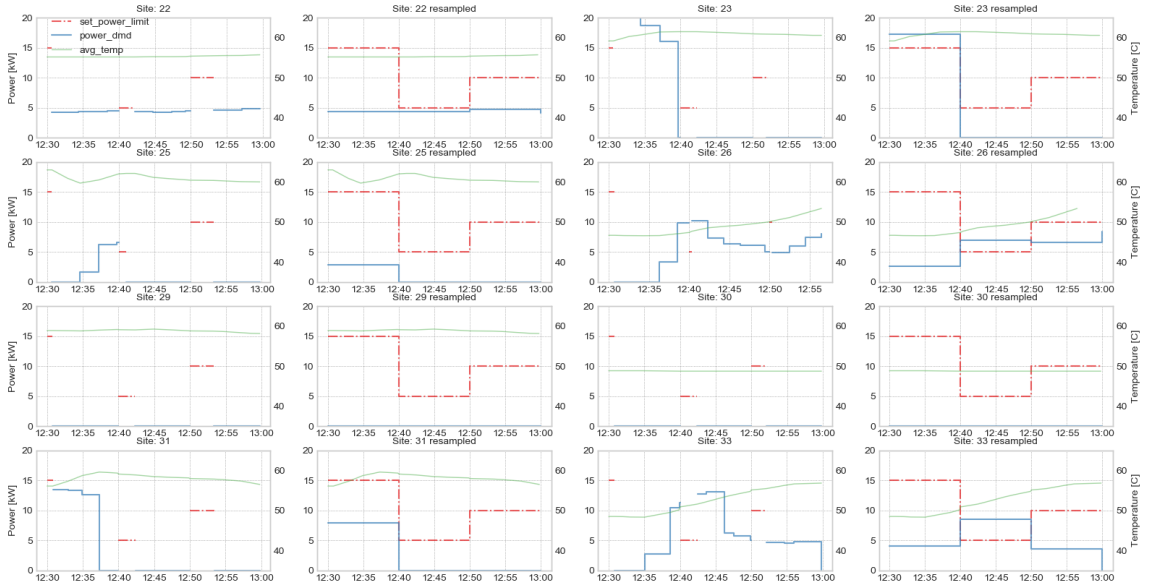

Figure 19: overview of charging profile test in group 2 (22, 23, 25, 26, 29, 30, 31, 33)

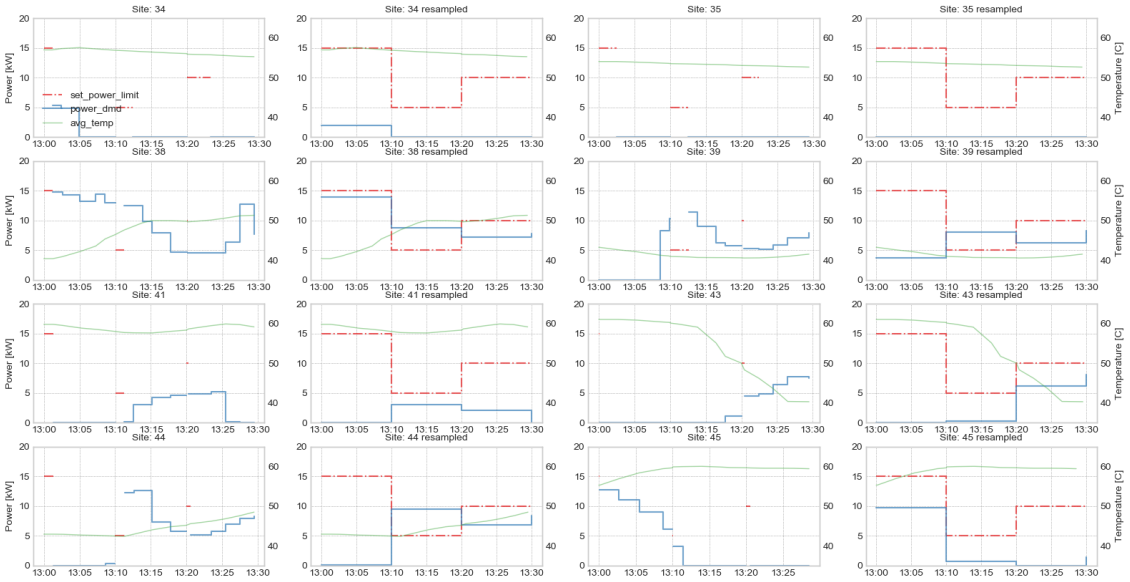

Figure 20: overview of charging profile test in group 3 (34, 35, 38, 39, 41,43, 44, 45)

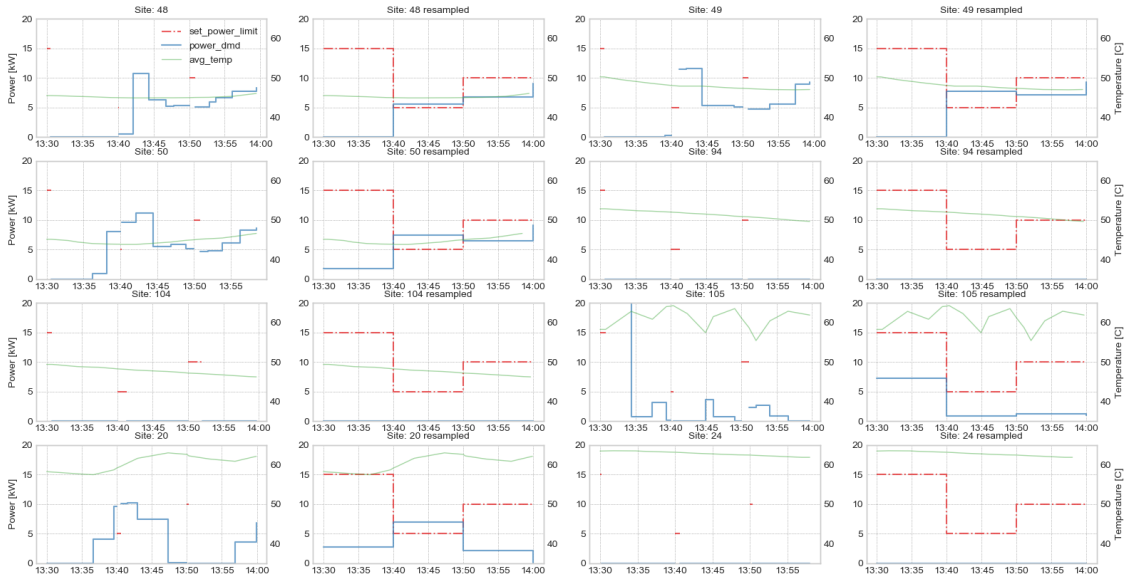

Figure 21: overview of charging profile test in group 4 (48, 49, 50, 94, 104, 105, 20, 24)

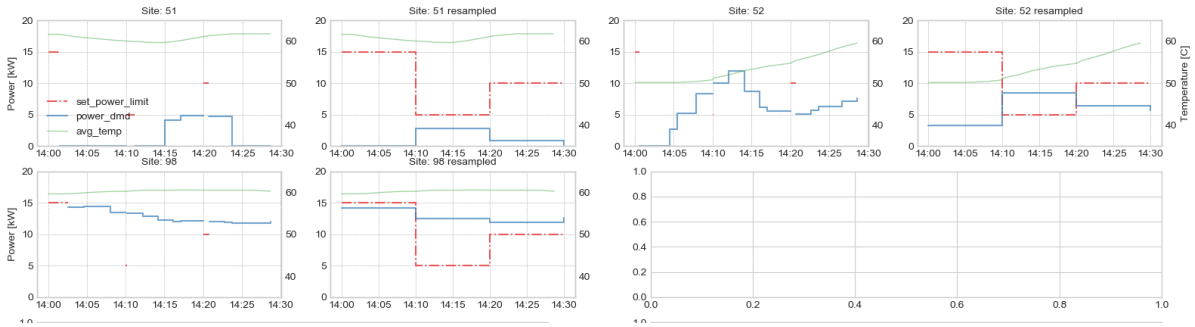

Figure 22: overview of charging profile test in group 5 (51, 52, 98)

3. Charging profile test (09/04/2020)

In this test the following charging profile was sent to the buffers:

- 0 - 20 min: 10 kW
- 20 – 40 min: 3 kW
- 40 - 60 min: 6 kW

Buffers 29, 104,105 and 98 were not tested as they never reached the temperature limits of top_temp < 56 deg and bottom_temp < 40 deg.

All tested buffers reacted to the control signals (in on or the other a way), some were more precise than others. Most buffers reduced their load when the top temperature was around 70 deg but there were a couple which reduced their load already after the first 20 minutes even though their top temperature was significantly lower than 70 deg:

- 23, 26, 34 ,(39), 45, 94

There is also a problem with buffer 13 for which the top and middle sensor are switched (buf_2_t is the top temperature and buf_1_t is the middle temperature).

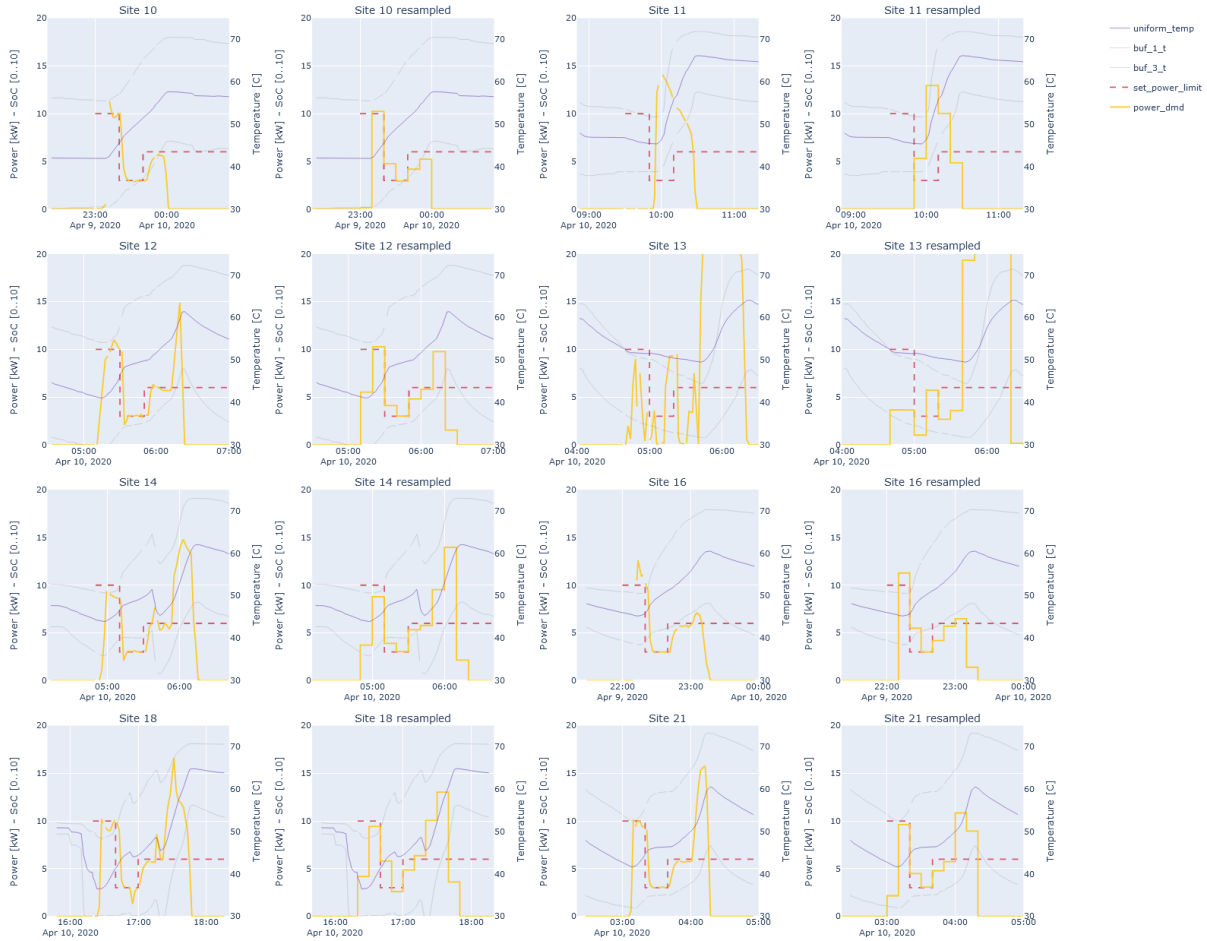

Figure 23: overview of 2nd charging profile test in group 1 (10, 11, 12, 13, 14, 16, 18, 21)

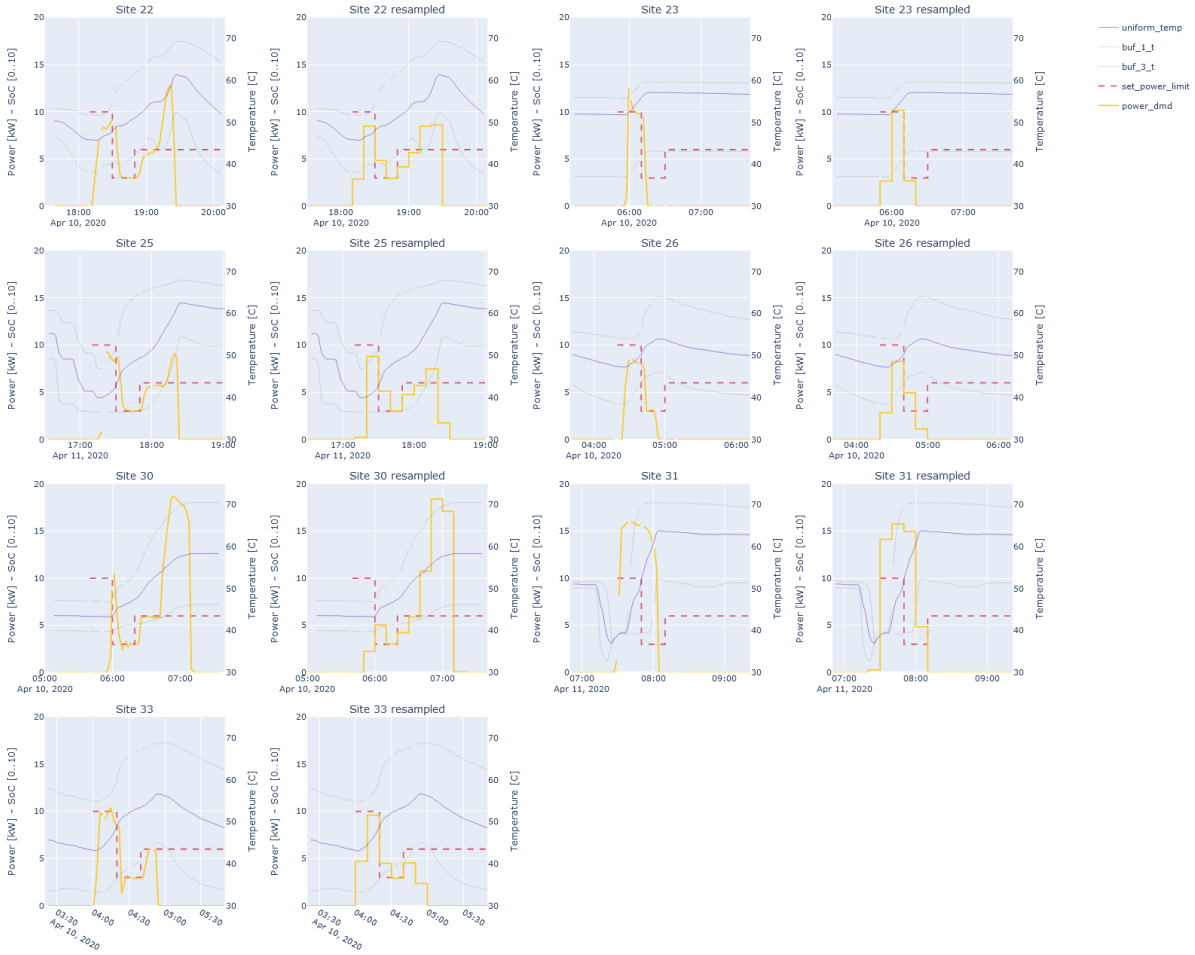

Figure 24: overview of 2nd charging profile test in group 2 (22, 23, 25, 26, 30, 31, 33)

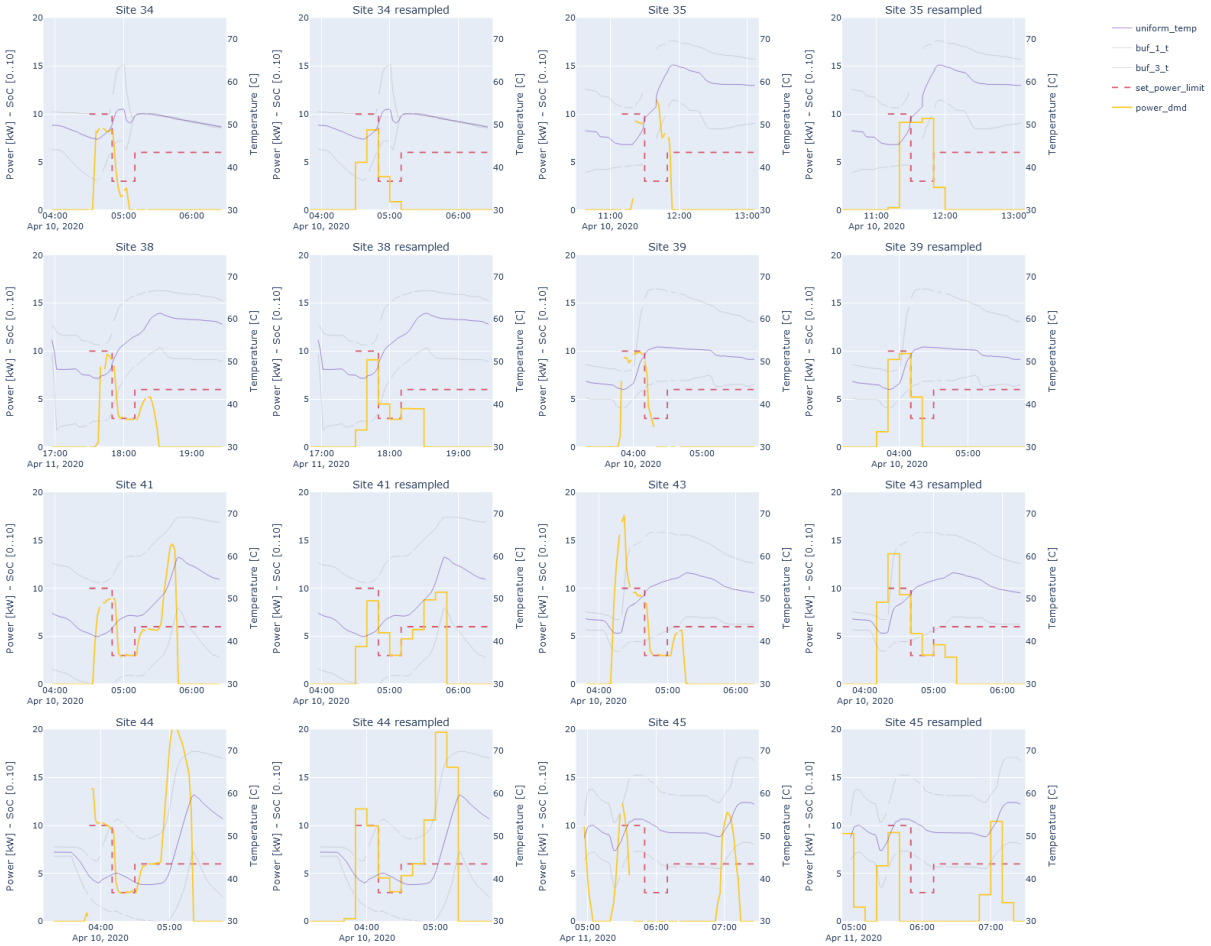

Figure 25: overview of 2nd charging profile test in group 3 (34, 35, 38, 39, 41, 43, 44, 45)

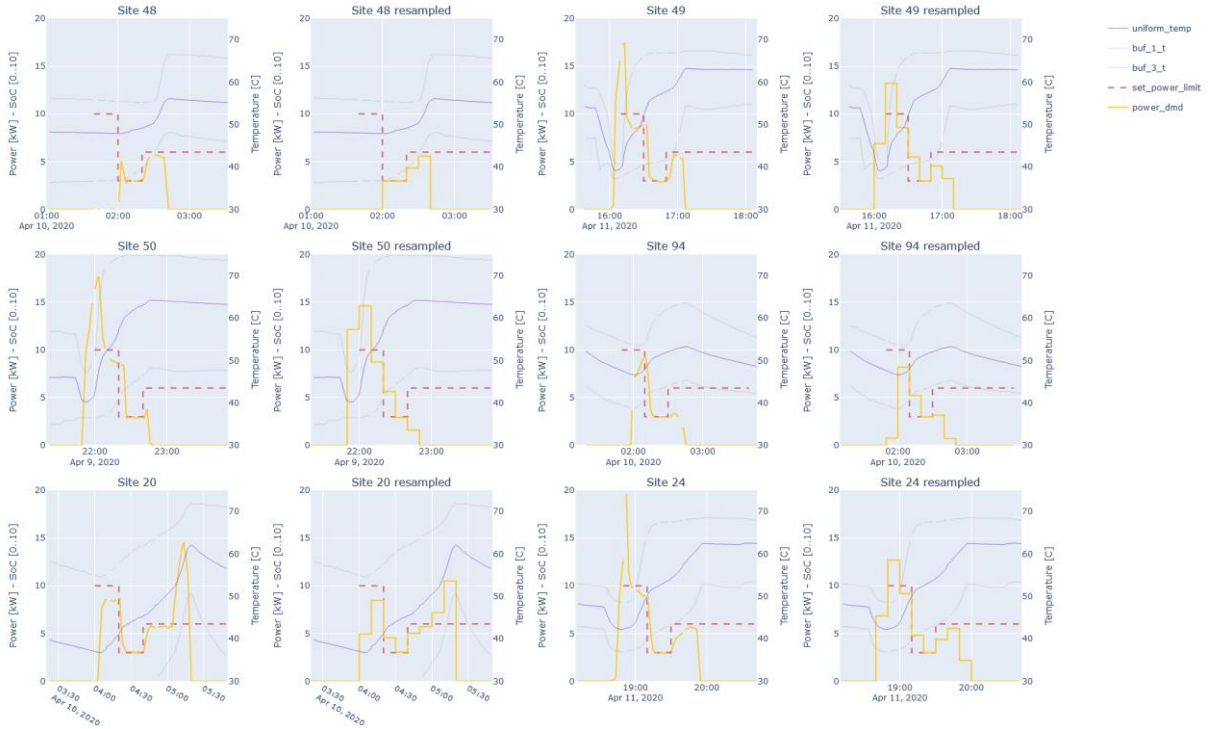

Figure 26: overview of 2nd charging profile test in group 4 (48, 49, 50, 94, 20, 24)

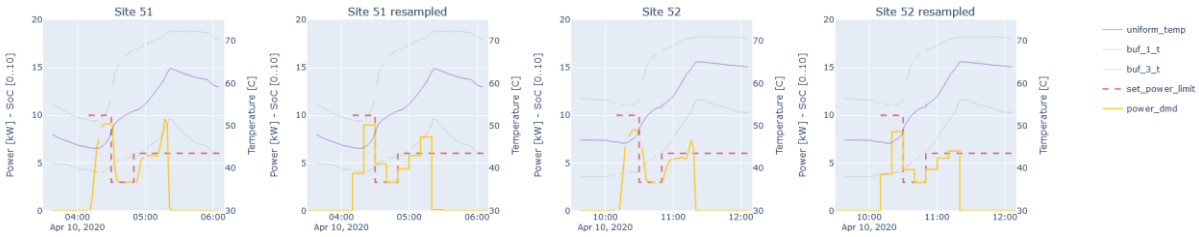

Figure 27: overview of 2nd charging profile test in group 5 (51, 52)

4. Charging profile test (15/04/2020)

In this test the following charging profile was sent to the buffers:

- 0 - 20 min: 20 kW

Based on these results a technical meeting was planned to discuss how we can change the buffer setpoints to improve the buffer response results. With all the tests done up until now, it is clear that the normal operation of the buffer still interferes too much with our desired control. For example, if we want to partly charge the buffer but the bottom temperature is not above 45 degrees when we stop charging, basline charging takes over and keeps on charging until this setpoint is reached.

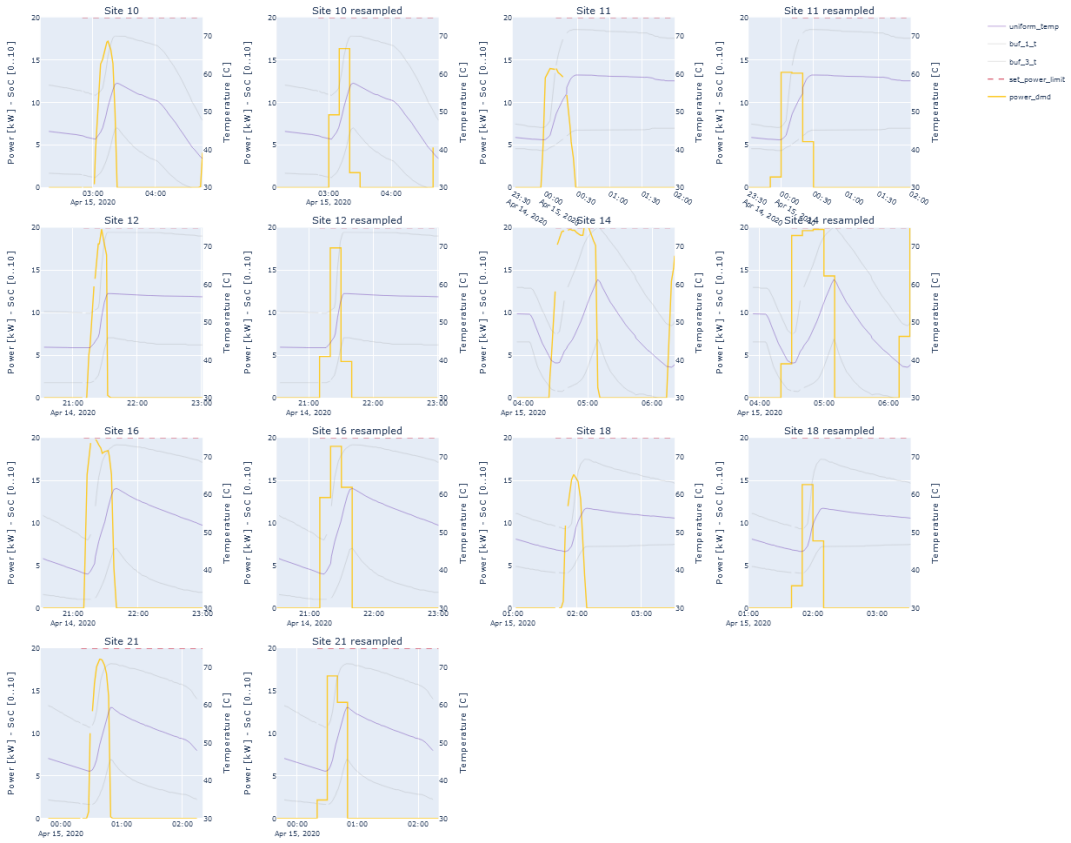

Figure 28: overview 3rd charging profile test group 1 (10, 11, 12, 14, 16, 18, 21)

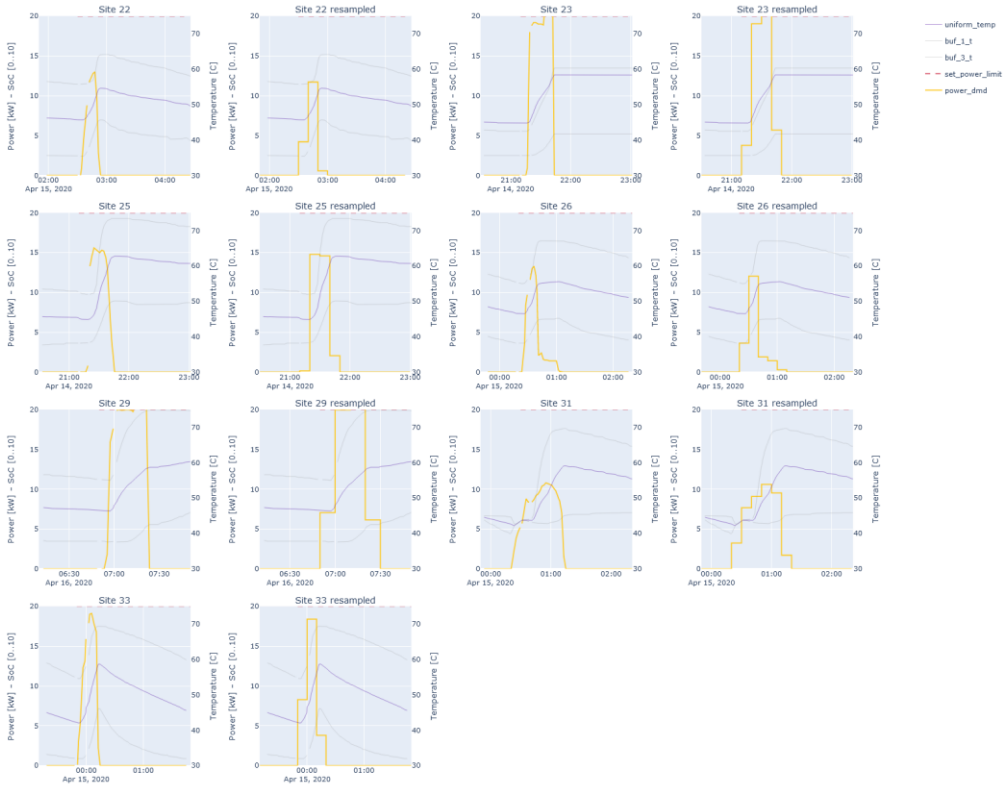

Figure 29: overview 3rd charging profile test group 2 (22, 23, 25, 26, 29, 31, 33)

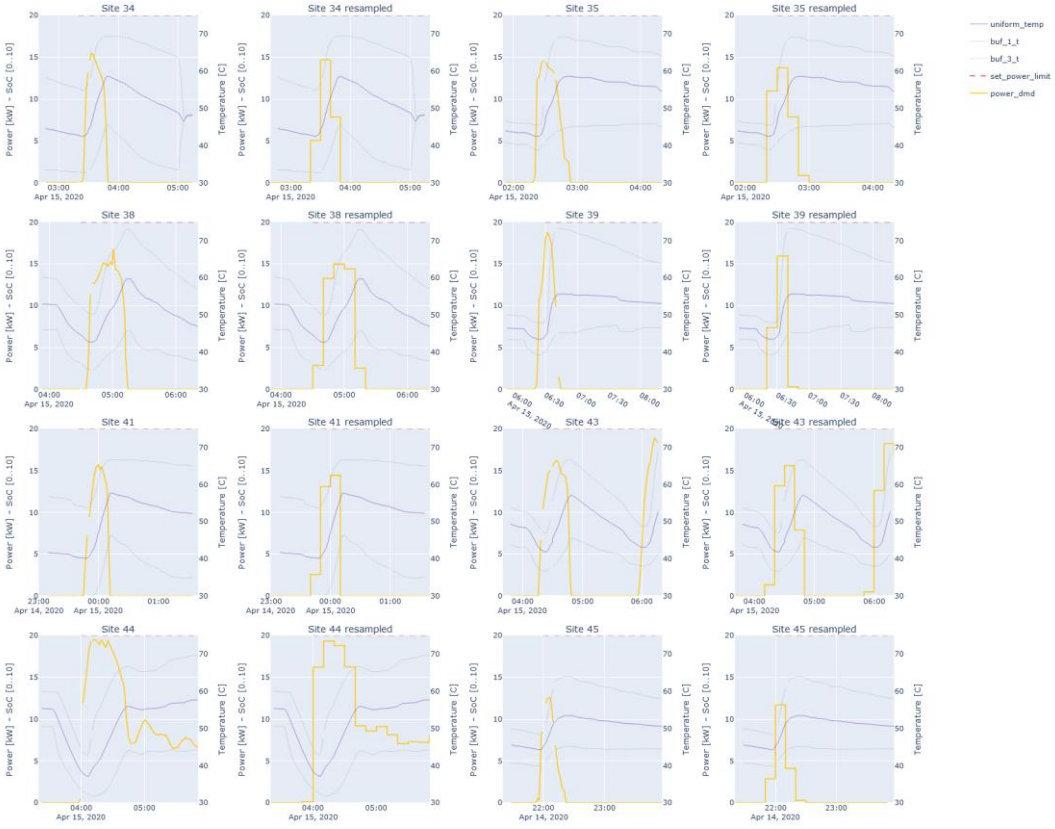

Figure 30: overview 3rd charging profile test group 3 (34, 35, 38, 39, 41, 43, 44, 45)

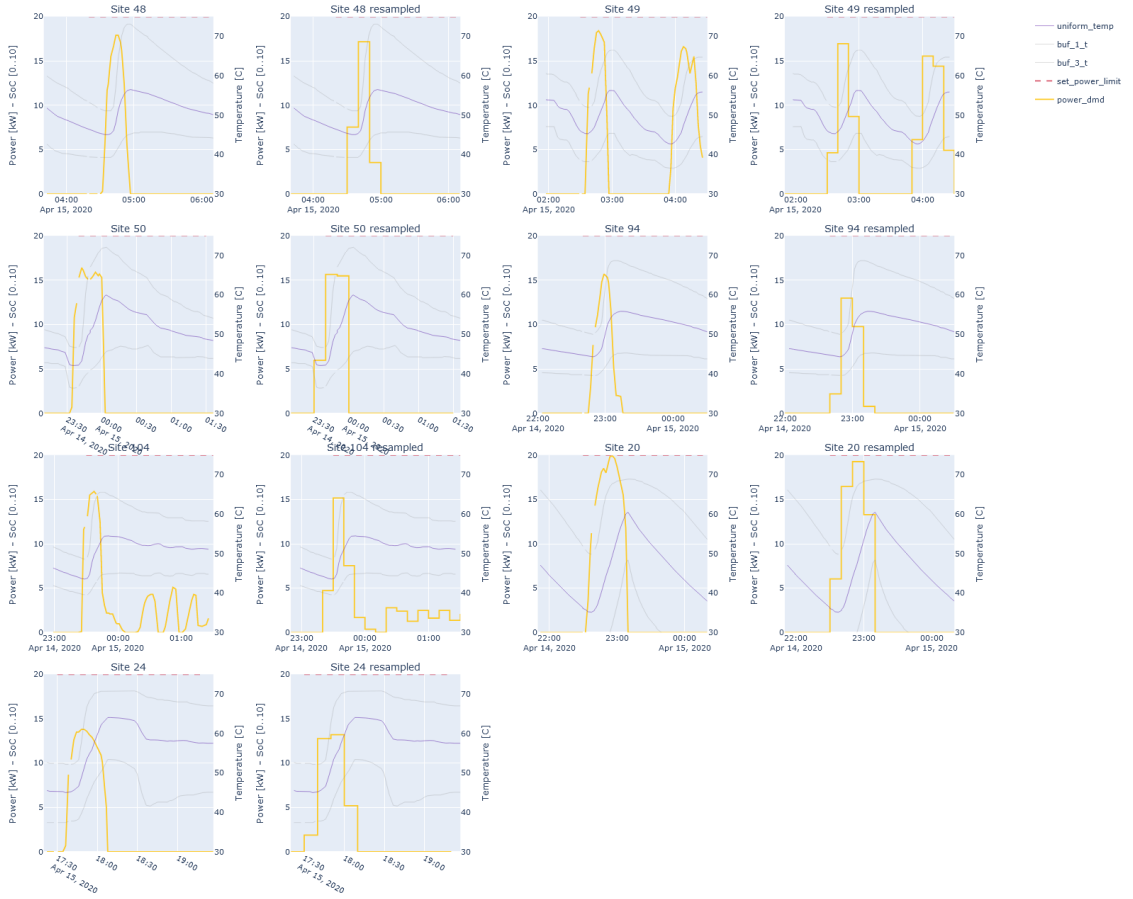

Figure 31: overview 3rd charging profile test group 4 (48, 49, 50, 94, 104, 20, 24)

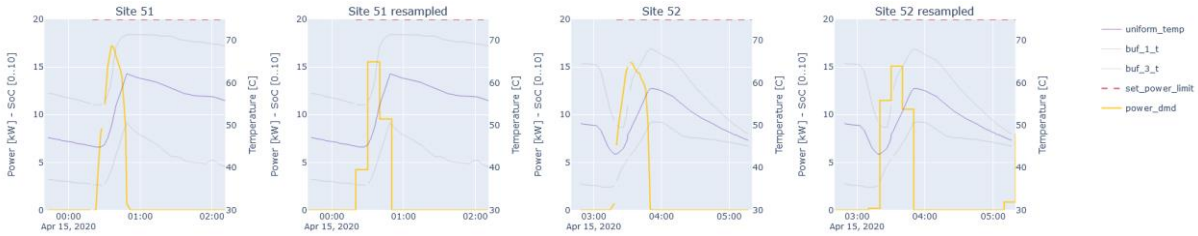

Figure 32: overview 3rd charging profile test group 5 (51, 52)

5. Charging profile test (17/04/2020)

In this test the following charging profile was sent to the buffers:

- 0 - 180 min: 6 kW

The results of this test showed a possible comfort violation for buffer 10, which is discussed in the report based on **Figure 14**. Based on the results of this test, it was decided to investigate other options to control the distributed buffers.

6. New control scenario test (25/05/2020)

Goal

Test the new buffer control approach which only uses a start charging signal without defining a power limit to make sure that the comfort boundaries are never violated. The latter could when a power limit was applied because this power limit was also applied to the DHW bypass valve. This could lead to cold water in case of hot water demand.

The main goal of the test is to see if this control solution works as expected and if there is a correlation between the state of charge of the buffer and the instantaneous power it starts charging with. For optimal control we need to be able to predict the charging power of the buffer when we give it the signal to start charging.

How

The test sequence is as follows:

- Every 5 minutes top and bottom temperature of all buffers, that have not yet been tested, are checked
    - If buffer temperatures are dropping and **bottom temperature < 40** and **top temperature < 56**
        - Set **set_power_limit = 100 for the following 10 minutes** (we can only set control signals starting from a multiple of 10 minutes)
        - Setting the set_power_limit >= 20 will force the buffer charging
    - Add the id of the tested buffer to the tested_buffer list
- The test stops automatically when all buffers have been tested

Results

The below list shows the results per tested buffer (250 l):

- **SoC [%}:** the SoC at the start of the control
- **Power [kW]:** max charging power during control and 10 minutes after control
- **Kwh_1:** charged energy during first 10 minutes of the charging
- **Kwh_2:** charged energy during second 10 minutes of the charging
- **Kwh_total:** total charged energy during 20 minutes
- **delay [min]:** delay in start control signal and real power increase

	[%]	[kW]	[kWh]	[kWh]	[kWh]	[min]
site_id	**soc**	**power**	**kwh_1**	**kwh_2**	**kwh_total**	**delay**
10	37.50	12.9	0.00	2.03	2.03	13.27
11	37.25	13.6	1.13	2.12	3.25	3.07
12	43.08	13.0	0.62	2.15	2.77	8.27
14	**36.50**	**20.6**	**2.68**	**3.35**	**6.04**	**-11.50**
18	32.25	15.3	0.04	2.41	2.45	8.27
21	44.67	13.9	0.07	2.06	2.13	8.27
25	31.25	15.3	0.06	2.37	2.43	8.43

26	38.83	12.1	0.31	1.79	2.10	8.60
33	42.92	12.1	0.72	0.97	1.69	8.33
35	**33.25**	**10.7**	**1.52**	**0.59**	**2.11**	**-6.73**
43	**22.08**	**16.2**	**2.41**	**2.31**	**4.73**	**-1.57**
44	36.92	9.6	0.00	1.58	1.58	13.42
49	**33.00**	**17.7**	**2.66**	**2.52**	**5.17**	**-1.88**
94	38.12	13.9	0.64	2.19	2.83	8.27
20	21.08	0.0	0.00	0.00	0.00	0.00
51	28.75	15.1	0.60	2.22	2.82	8.50
52	41.29	15.1	0.47	2.09	2.56	8.43
98	38.75	12.0	0.44	1.98	2.42	8.27

The last column shows the delay (in minutes) between setting the buffer to forced charging and the actual change in charging power. The bold rows show a negative delay, which indicates that the buffer started charging before the control signal was set.

Figure 33 shows a scatter plot of the maximum charging power of a buffer versus the SoC of the buffer at the start of the charging. The buffers with a negative delay have been left out because for them the charging power is not related to the control signal.

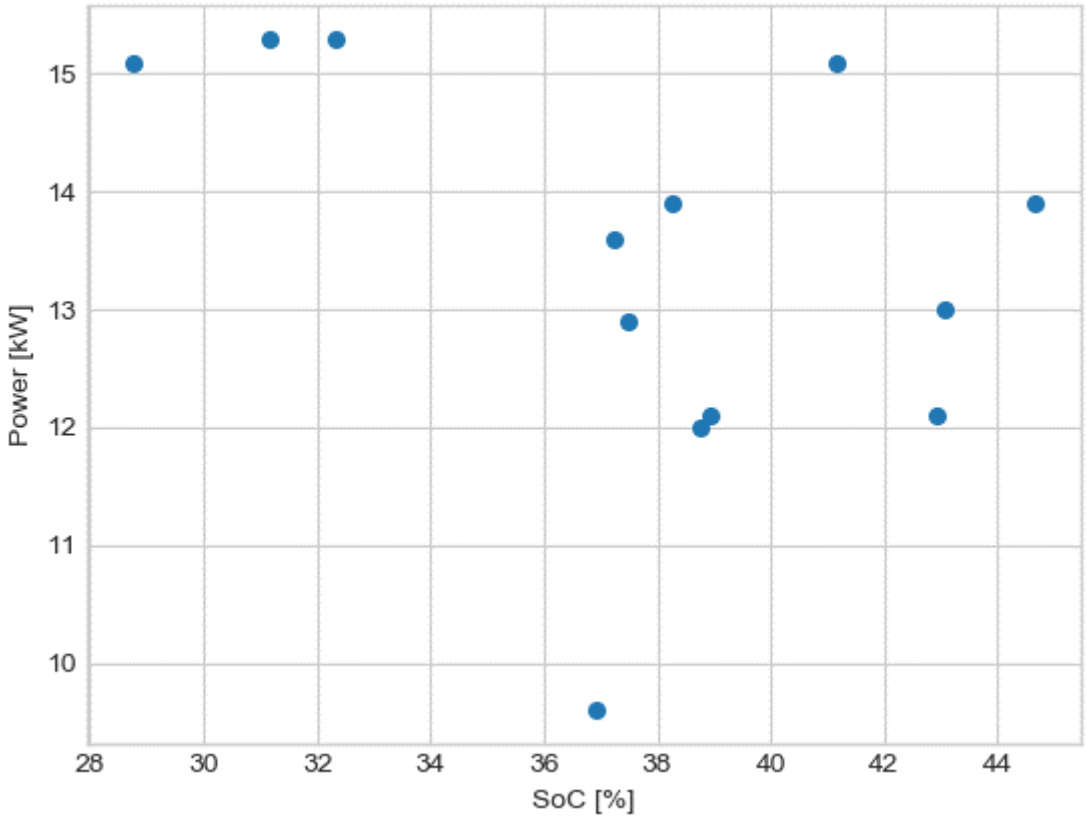

Figure 33: Maximum charging power versus SoC of the buffer

The above table and figure show that:

- The **maximum charging power** is between **12 and 15.3 kW** except for an outlier of site 44 with a max of 9.6 kW (as seen in figure 2, site 44 also reaches 13 kW of charging power but only 48 minutes after the control signal has been applied and the table only takes into account 10 minutes after the control signal)
- Buffers with an **SoC < 35%** charge at a power of **+- 15 kW**
- Buffers with an **SoC > 35%** charge at a power **in between 12 and 14 kW** except for site 52 which charges at 15.1 kW
- The average delay between setting the control signal and the increase of the charging power is around 8 minutes. This means that this can add up to a delay of 18 minutes as we can only set the control signal for the next multiple of 10 minutes.
- The buffers with a negative delay which started charging automatically before the control signal was set, typically charge with a higher maximum power: 20.6 kW, 16.2 kW and 17.7 kW. They also charge more energy +- 5 kWh as opposed to +- 2.5 kWh for the ones that started with the control signal.

The results for the 600l buffers look similar, except their max charging power is higher =- 20 kW with a SoC around 50%:

[%]	**[kW]**	**[kWh]**	**[kWh]**	**[kWh]**	**[min]**

site_id	soc	power	delay	kwh_1	kwh_2	kwh_total
13	50.25	19.6	1.52	3.08	4.61	8.43
23	44.00	21.4	1.38	3.57	4.95	8.48

Conclusions

- 250 liter buffers
  - **SoC < 35%** charge at a power of **+- 15 kW**
  - **SoC > 35%** charge at a power **in between 12 and 14 kW**
  - Average delay of charging 8 minutes
  - Average charged energy = 2.5 kW
- 600 liter buffers
  - 40 % < SoC 50 % charge at a power of +- 20 kW
  - Average delay of charging 8.5 minutes
  - Average charged energy 4.5 kWh (but charges longer than 20 minutes after control signal has been set so effective charged energy is higher)
- Buffers that started charging before control signal was set (250l)
  - Charge at higher powers (16.2 kW to 17.7 kW)
  - Charge more energy, average = 5 kWh

7. New control scenario test (06/06/2020)

Goal

Test the new buffer control approach which only uses a start charging signal without defining a power limit to make sure that the comfort boundaries are never violated. The latter could happen in previous response tests where a power limit was applied because this power limit was also applied to the DHW bypass. This could lead to cold water in case of hot water demand.

The main goal of the test is to see if this control solution works as expected and if there is a correlation between the state of charge of the buffer and the instantaneous power it starts charging with. For optimal control we need to be able to predict the charging power of the buffer when we give it the signal to start charging.

How

The test sequence is as follows:

- Every 5 minutes top and bottom temperature of all buffers, that have not yet been tested, are checked
  - If buffer temperatures are dropping and **bottom temperature < 40** and **top temperature < 56**
    - Set **set_power_limit = 100 for the following 20 minutes** (we can only set control signals starting from a multiple of 10 minutes)
    - Setting the set_power_limit >= 20 will force the buffer charging

- Add the id of the tested buffer to the tested_buffer list, the test stops automatically when all buffers have been tested

Results

The below list shows the results per tested buffer (250 l):

- **SoC [%]:** the SoC at the start of the control
- **Power [kW]:** max charging power during control and 10 minutes after control
- **Kwh_1:** charged energy during first 10 minutes of the charging
- **Kwh_2:** charged energy during second 10 minutes of the charging
- **Kwh_total:** total charged energy during 20 minutes
- **delay [min]:** delay in start control signal and real power increase

	[%]	[kW]	[kWh]	[kWh]	[kWh]	[min]
site_id	**soc**	**power**	**kwh_1**	**kwh_2**	**kwh_total**	**delay**
10	33.92	15.8	0.00	2.28	4.67	13.63
11	39.50	10.2	1.06	0.43	1.49	-15.92
12	38.25	15.9	0.13	2.64	3.71	8.68
14	32.58	16.2	0.10	2.61	4.96	9.15
16	38.75	15.4	0.01	2.24	3.44	8.82
18	**27.58**	**23.5**	**1.84**	**3.10**	**5.75**	**-0.98**
21	45.67	13.6	0.14	2.06	3.48	8.98
22	**21.92**	**21.4**	**1.78**	**2.54**	**5.25**	**-6.02**
25	**39.17**	**10.2**	**1.05**	**1.37**	**2.42**	**-6.38**
26	41.83	13.6	0.02	2.27	3.47	8.65
31	**18.75**	**14.7**	**1.43**	**2.34**	**5.74**	**-5.92**
33	**27.17**	**16.6**	**1.64**	**0.00**	**1.64**	**-16.28**
34	**24.25**	**8.5**	**0.85**	**1.16**	**2.55**	**-6.10**
35	**27.58**	**15.8**	**1.62**	**2.09**	**3.71**	**-6.30**
38	41.75	10.1	0.15	1.67	3.08	8.95
39	47.17	9.8	0.10	1.63	2.16	8.97
41	**34.33**	**14.1**	**1.37**	**0.97**	**2.34**	**-16.05**
43	36.75	14.7	0.13	1.98	4.20	8.75
44	38.00	11.1	1.08	1.42	2.96	4.13

	[%]	[kW]	[kWh]	[kWh]	[kWh]	[min]
site_id	**soc**	**power**	**kwh_1**	**kwh_2**	**kwh_total**	**delay**
45	39.75	11.5	0.01	1.87	3.07	8.87
48	39.62	0.0	0.00	0.00	0.00	0.00
50	**31.42**	**12.7**	**1.21**	**2.06**	**4.01**	**-11.23**
94	**25.00**	**14.4**	**1.37**	**2.10**	**4.20**	**-0.93**
20	-16.58	0.0	0.00	0.00	0.00	0.00
24	34.50	13.2	0.01	2.04	3.99	8.75
51	42.75	17.2	0.18	2.74	4.60	8.72
52	42.33	13.9	0.15	2.24	4.10	8.78
98	**37.75**	**11.3**	**1.08**	**1.71**	**2.79**	**-5.90**

The last column shows the delay (in minutes) between setting the buffer to forced charging and the actual change in charging power. The bold rows show a negative delay, which indicates that the buffer started charging before the control signal was set.

Buffer 20 is not active at this moment and buffer 48 did not react to the control signal, both are left out of the analysis.

Figure 34 shows a scatter plot of the maximum charging power of a buffer versus the SoC of the buffer at the start of the charging. The buffers with a negative delay have been left out because for them the charging power is not related to the control signal.

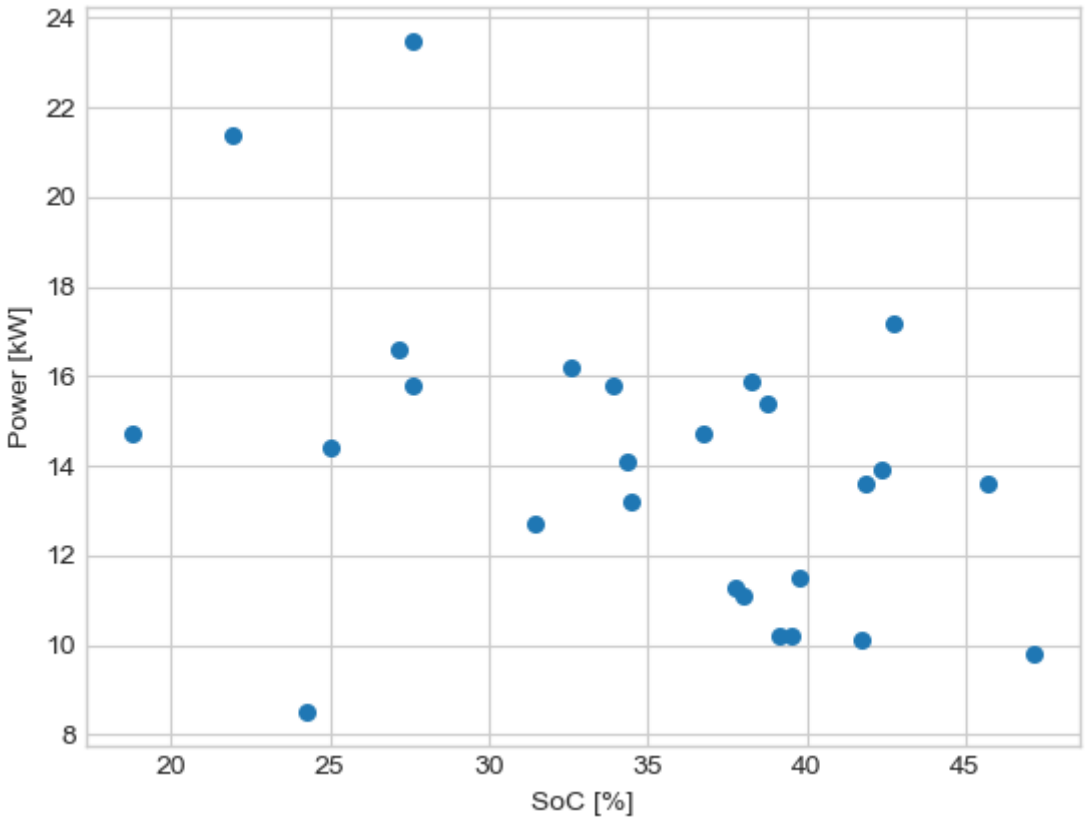

Figure 34: Maximum charging power versus SoC of the buffer

The above table and figure show that:

- The **maximum charging power** is between **12 and 18 kW** except for 2 outliers above 20 kW but they started charging automatically (negative delay for 18 and 22)
- Buffers with an **SoC < 35%** charge at a power **in between 12 kW and 17 kW**
- Buffers with an **SoC > 35%** charge at a power **in between 10 kW and 16 kW**
- The average delay between setting the control signal and the increase of the charging power is around 8,8 minutes. This means that this can add up to a delay of 18,8 minutes as we can only set the control signal for the next multiple of 10 minutes.
- The buffers with a negative delay which started charging automatically before the control signal was set, typically charge with a higher maximum power. They charge a similar amount of energy +- 3.5 kWh as opposed to +- 3,7 kWh for the ones that started with the control signal.

The results for the 600l buffers look similar, except their max charging power is higher +- 20 kW with a SoC around 40 to 50% and 24 kW at 30%:

	[%]	[kW]	[kWh]	[kWh]	[kWh]	[min]
site_id	**soc**	**power**	**kwh_1**	**kwh_2**	**kwh_total**	**delay**
13	56.83	19.5	0.85	2.88	4.90	3.95
23	30.75	24.5	0.02	3.83	7.90	9.15

29	40.25	21.5	0.24	3.39	7.12	8.90

Conclusions

- 250 liter buffers
  - **SoC < 35%** charge at a power of **in between 12 and 17 kW**
  - **SoC > 35%** charge at a power **in between 10 and 16 kW**
  - Average delay of charging 8,8 minutes
  - Average charged energy = 3,7 kW
- 600 liter buffers
  - Soc = 40 % to SoC 50 % charge at a power of +- 20 kW
  - Average delay of charging 7,33 minutes
  - Average charged energy 3.5 kWh (but charges longer than 20 minutes after control signal has been set so effective charged energy is higher)

8. Aggressive valve setting test (17/03/2021)

In the test of the new control option it was found that there is a significant delay between the time the control signal is applied and the time a buffer starts charring. This was mainly due to the slow opening of the primary valve which leads to a low flow. In this test VITO and Enerpipe tested buffer 49 with a more aggressive valve setting leading to charging behaviour after only 4 minutes, as can be seen in Figure 35. The red line shows the control signal that was sent, sending a value > 25 will induce a force charging cycle) and it can be seen that only 4 minutes after sending this control signal, the buffer starts charging at max power (green line).

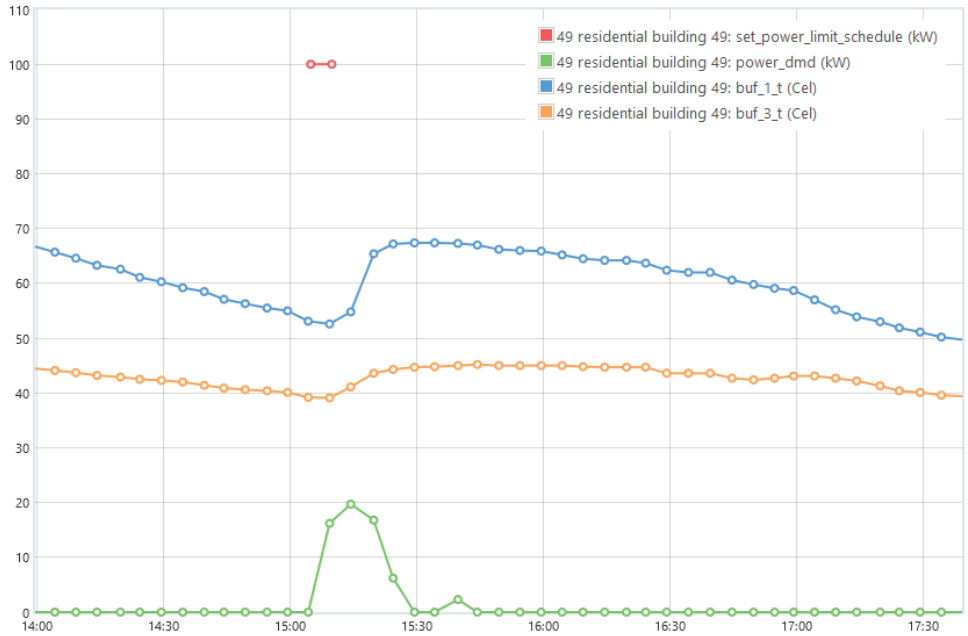

Figure 35: test aggresive valve settings at buffer 49