Deliverable 6.456
Integrated Prototype Release Documentation

Revision : 1.0
Due date : 31 May 2015
Actual submission date : 31 May 2015
Lead Coordinator : UCC4C

Deliverable Administration & Summary			
No & name	**D6.456- Integrated Prototype Release Documentation**		
Status	Final Release	**Due Date**	2015-05-31
Author(s)	H. Simonis, UCC4C A. Schülke, M. Schmidt, NEC		
Editor	H. Simonis, UCC4C		
DoW	Task 6.4: Documentation UCC-NUIC (4C), NEC and Cartif document the algorithms and tools developed. The documentation must describe the choices made in selecting the final architecture, a detailed description of the algorithms and strategies used, and performance evaluation results (compared to the status-quo) obtained from monitoring tools during operation. (result: D6.456; TL: UCCNUIC (Group 4C))		
Comments			

Document history

V	Date	Author	Description
0.1	24 Mar 2015	UCC4C, NEC, UTRC-I	Draft outline
0.2	07-05-2015	NEC	Input related to Frankfurt demonstration case
0.3	15-05-2015	UCC4C	Input related to UCC Campus demonstration case
0.4	20-05-2015	NEC	Adaptation of input to Frankfurt site, add references
0.5	21-05-2015	UCC4C	Added material about UCC Campus
1.0	26-05-2015	UCC4C, NEC	Final polishing

Submitted Document History

V	Date	Author	Description
1.0	31-05-2015	UCC4C, NEC	Final version 1.0

Table of contents

1. Executive Summary

This deliverable D6.456 provides the documentation of the tools and algorithms used in WP6 at the Load Balancing use cases at the use cases UCC Campus' CHP operation (Ucd[1] UCC 6.1) and Grass Heating Load Balancing for the Commerzbank Arena (Ucd HSG 6.1m). Additionally, performance evaluation results (compared to the status-quo) are summarized.

The outlining for Ucd UCC 6.1 will report on the documentation of the optimization algorithm to provide operational schedules for the UCC CHP to enable energy cost reduction, as described in D6.1 and necessary adaptations.

For the Ucd HSG 6.1m, the control strategies as used for optimization of the grass heating operation reaching energy savings, and load balancing control for embedding grass heating into daily operation of the entire heating concept of the Commerzbank Arena Frankfurt are documented. Extensions beyond the descriptions given in D6.4 are outlined.

[1] Ucd: Use case description

2. INTRODUCTION

2.1 Purpose

The main scope of this deliverable D6.456 (Task 6.4) is the documentation of the tools and algorithms used at the Load Balancing use cases:

Use Case Ucd UCC 6.1: optimization algorithm to provide operational schedules for the UCC CHP to enable energy cost reduction.

Use Case Ucd HSG 6.1m: control strategies for optimization of the grass heating operation reaching energy savings, and load balancing control for embedding grass heating into daily operation of the entire heating concept of the Commerzbank Arena Frankfurt.

2.2 The D6.456 partners and their contributions

This task has been performed jointly by UCC-NUIC (4C) and NEC. The major contributions are:

- **UCC-NUIC (4C)** coordinated the documentation of the Ucd UCC 6.1, which had been jointly investigated by UCC-4C and UTRC-I during this project.

- **NEC** provided the documentation for the Ucd HSG6.1m which had been investigated together with HSG Zander in this project.

2.3 Differences to Previous Document Versions

This is the first version of the D6.456 document.

3. Final Architecture

The CAMPUS21 system is designed and built as an ICT-integration platform to integrate pre-existing, independently functioning building management subsystems into a holistically managed monitoring and control platform.

The core target of WP6 is the utilization of energy and context data in order to develop Load Balancing methods in an integrative manner utilizing the different data sources and operational information available through the CAMPUS21 data system. Figure 1 provides the final architectural view of the CAMPUS21 high-level view as introduced in deliverable D2.4 [1] and adapted through the development of the use cases documented in D4.456 [2].

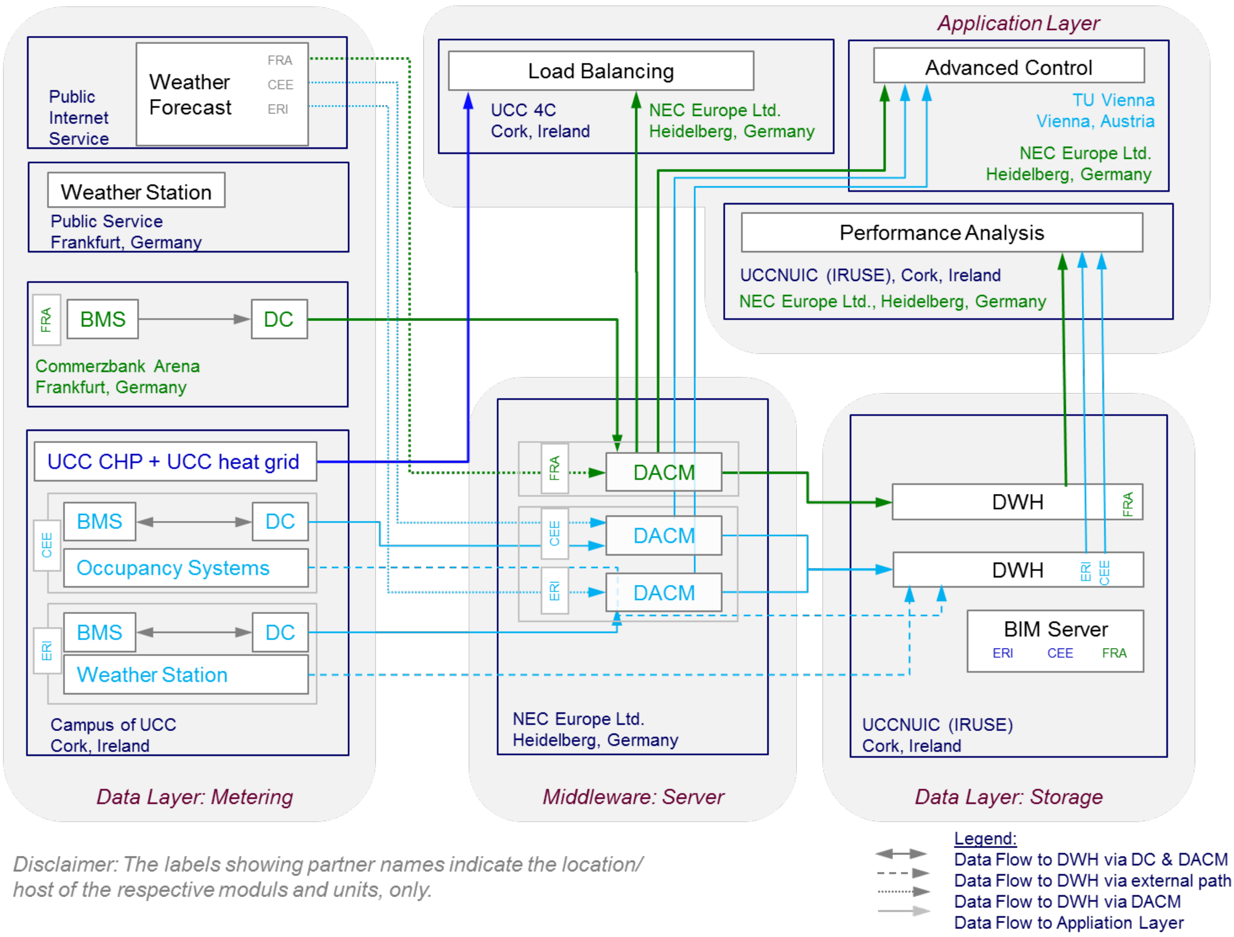

Figure 1: Final CAMPUS21 overall architecture, shown with hosting locations and data information flow

For WP6 and this D6.456, the data flow and module host/ locations for Ucd UCC6.1 are represented in dark-blue, and for Ucd HSG 6.1m in green.

The middleware deployment details are provided in deliverable D4.456 [2].

3.1 Ucd UCC 6.1: CHP Optimization on UCC Campus

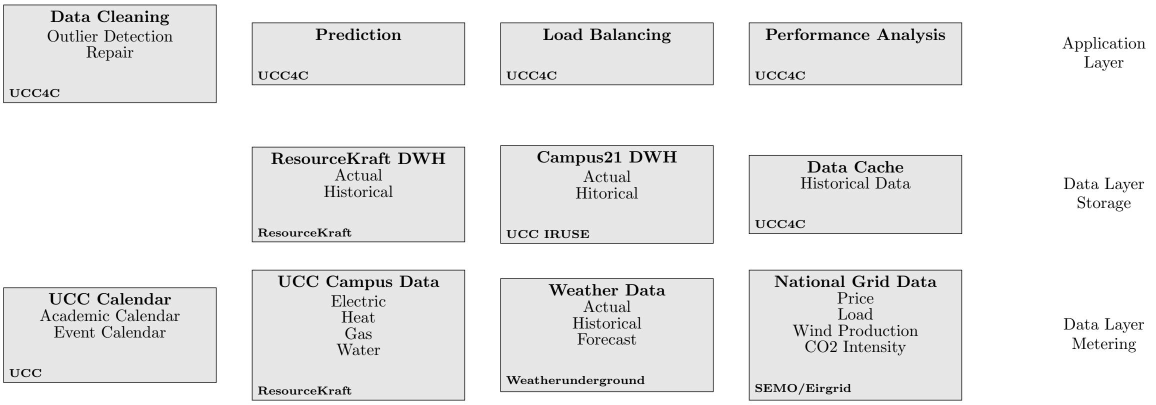

Figure 2: Final Campus21 Architecture for UCC Campus Experiments

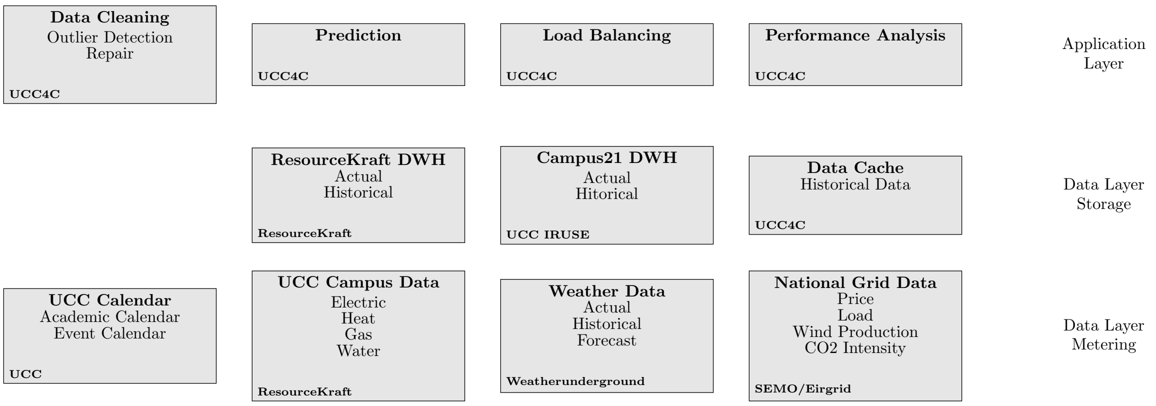

Figure 2 shows the elements of the overall architecture relevant to the UCC Campus experiments. The metering layer consists of four main sources. The UCC calendar provides the academic calendar, defining the teaching, exam and holiday periods for the academic year, and the event calendar, which indicates special events on the campus. The existing ResourceKraft sensor system is used to collect historical and actual data about electricity and heat consumption on the campus, together with gas use and water consumption data. These data form the basis of the forecasts for electricity and heat consumption on the campus. The weather data comes from weatherunderground, who provide historical, actual and forecast data for the Cork airport. This data feeds into the forecast of the grid CO2 intensity. The last set of data is provided by SEMO and Eirgrid, the Irish electricity market operator and the national grid operator, respectively. This data covers historical and actual data for grid electricity price, national grid load, total wind production and CO2 intensity. We also use the existing forecasts for price, national grid load and wind production. There currently is no national forecast for CO2 intensity, which we therefore have to produce ourselves.

The second layer in the architecture model is the data storage layer, which for the UCC Campus covers three elements. We interact with the ResourceKraft system through their data warehouse, which provides the campus measurement data in a consistent, normalized format. Data is then stored in the Campus21 data warehouse, for long term storage and analysis. Finally, we locally cache historical data on the load balancing system to minimize network load and latency.

The third layer is the application layer, where the following operations are performed. As a first step, a data cleaning procedure uses outlier detection and specific repair functions to remove improper values in the data streams, replacing them with more likely values. The second step performs prediction of demand and resource costs based on actual and historical data. These predictions are then used in the load balancing routine itself to generate a schedule for the future. The actual values are evaluated in the performance analysis to provide key performance indicators.

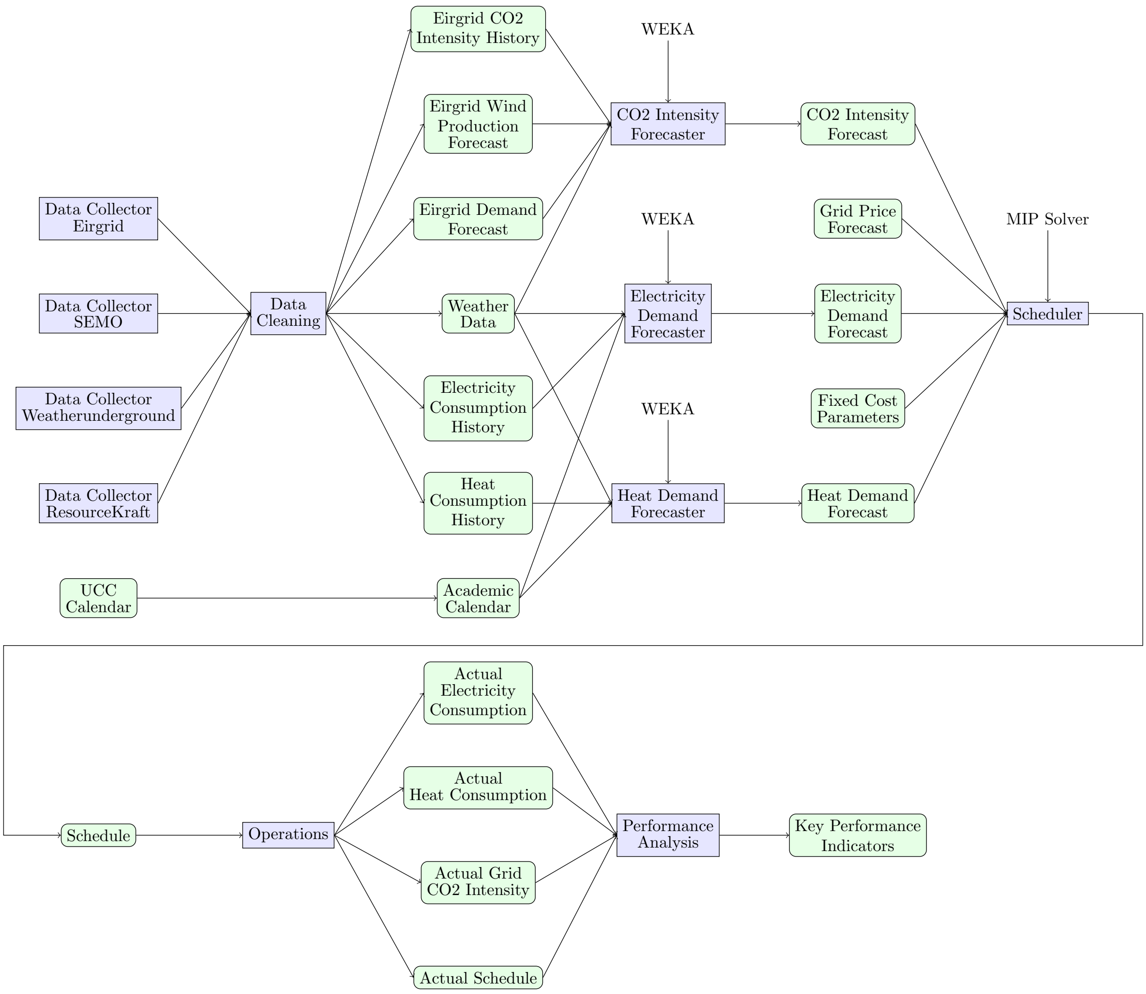

Figure 3: Final Deployment Diagram for UCC Campus

Figure 3 shows the deployment diagram for the UCC Campus experiments. There are six main steps in the operation.

- The first step performs the data collection from the different feeds and the import to the storage layer. Each data collector uses specific code to connect to the feed, extract the required data and convert it to the internal storage format. The feeds use different formats (JSON, XML, xls, csv) that are handled through common Java libraries.

- The next step performs data cleaning, detecting missing and probably incorrect data, replacing those values with estimated values. Different repair strategies have been implemented, and are used according to the data feed and its properties.
- In the next step, the required forecasts for the system are prepared. This is using feature based regression models, developed for each forecast type. We use the Weka tool as the underlying software tool to produce the forecasts.
- In the following step, the load balancing system computes a schedule for the next 24 hours based on the forecasts computed by our tools and fore- casts/data loaded from other sources. The optimization tool used is a MIP (mixed integer programming) solver.
- In the fifth step, the generated schedule is passed to operations for implementation. It is then run on the actual device. The controlled device is the CHP plant on the campus, for which the mode is set according to the schedule. Its local control adjusts for changes in demand as required.
- In the last step, results monitoring operations are collected with the same data collection tools, and the performance of the system is analysed. This provides a set of key performance indicators.

3.2 Ucd HSG 6.1m: Heat Demand Load Balancing at Commerzbank Arena Frankfurt

The following figure represents the final system architecture view, in place for the Commerzbank Arena experiments performed in WP6. Data analysis, optimization and evaluation are realized by the modules of the application layer.

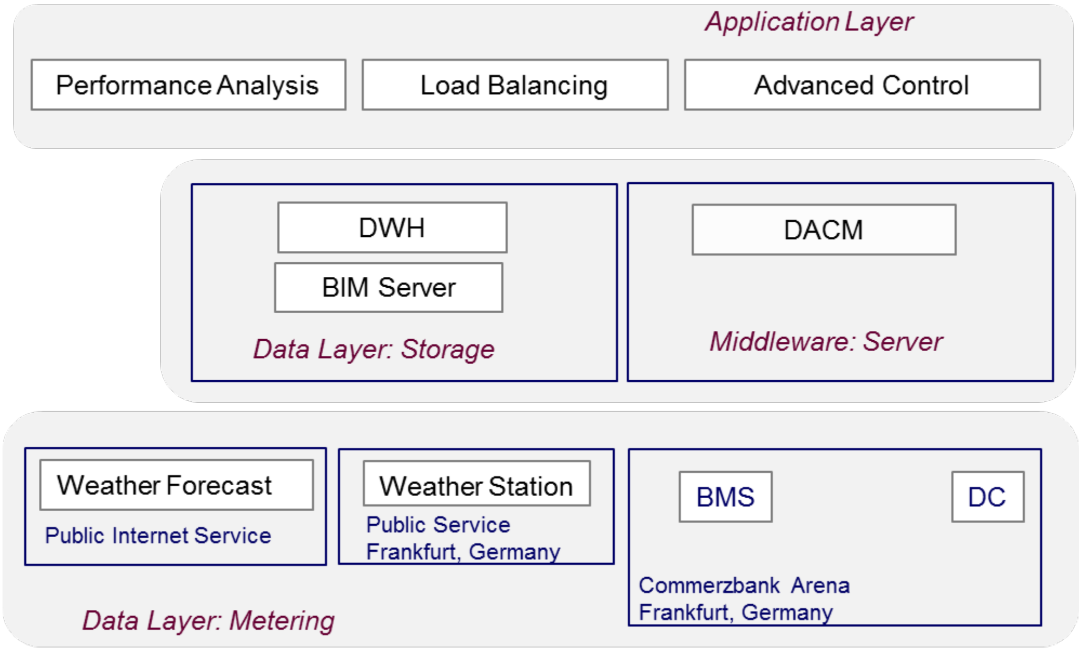

Figure 4: Final CAMPUS21 System Architecture for experiement realization at the Commerzbank Arena Frankfurt

The assessment of the existing control strategies had been investigated using historical data records stored at the Commerzbank Arena BMS local data storage. These records have only limited information, relevant to ensure correct operation.

With the implementation of a detailed smart metering concept within the Commerzbank Arena, detailed information of the consumption of heating and electrical energy, but also the material flows could be collected. Data aggregation deployed over the DC-DACM middleware connecting with local BMS-DC gateways, to the CAMPUS21 Data Warehouse as well as to external context information enabled a holistic data aggregation in real time (realized in 10 minute data aggregation intervals).

For WP6, the applicaton layer – realized for the use case at the Commerzbank Arena - contained 3 modules to execute the experiements, data analysis, control actuations, and the overall evaluation:

- Load Balancing: Real Time data analytics of operational and context information; enforcement of basic control mechanisms;
- Advanced control: Data analytics for advanced scenarios involving operation parameter forecast based on analysis of real-time operational and context parameters as well as weather forecast; application of machine learning technologies (deep belief networks for regression) for operational parameter forecast.
- Performance Analysis: Calculation of the KPI parameters UPT, OPT; and calculation of energy and cost savings in comparison to previous winter season

4. CONTROL ALGORITHM CHOICES

4.1 Ucd UCC 6.1: CHP Optimization on UCC Campus

The detailed description of the model used is contained in Deliverable D6.2. This model was used in the first set of experiments in April 2014. Some changes to the model were requested by Dalkia to make the generated schedule more easily acceptable to the plant operator. These changes were described in Deliverable D6.4, and were tested in a second set of experiments on the plant in November 2014.

We recall the main assumptions for the selected approach:

During the project duration, the heating and electricity supply for the UCC campus was split between two operators. Dalkia owned and operated the CHP plant, providing electricity and heat to the college, and UCC B&E operated steam boilers to supplement the heating when required. The conditions and prices for the energy supply were specified in a contract between ESB, the initial owner and operator of the CHP, and the university from 1999. This contract was taken over by Dalkia when they acquired the plant from ESB. The electricity tariff for the college is a fixed price day/night tariff, which does not consider the real-time national market price, and which also does not depend on whether the CHP is running or not. The heat provided by the CHP plant is paid for with a lump sum every year. The contract stipulates a 15-hour per day running period for the CHP, independent of heating demand in the college. During summer, only the Food Sciences building consumes steam for process needs, any remaining heat is dumped by coolers in the boiler house.

When the CHP is not running, or when demand exceeds capacity of the plant, Dalkia imports electricity from the grid, at market price. As this is sold on at a fixed price, Dalkia may earn or lose money depending on the grid price at the moment.

In the Irish market, the actual price is not known at time of consumption, but only with a delay of 4 days (EP2 price from SEMO). A 24-hour ahead price forecast is given daily by the market operator, but the potential error in the price is quite large. As we will see in the detailed analysis, any later forecast (EA2, WD1) does not significantly improve on the error of the initial (EA) forecast. This means that the price information does not improve as we get closer to time of implementation.

The hardware of the control system of the CHP plant was designed not to allow export of electricity to the grid, any changes of that system would require new hardware and potentially updates to the grid connector itself. This means that the CHP plant can't be used to export electricity at peak price periods, It can only satisfy the internal demand of the core campus buildings.

Dalkia can also not easily extend the running time of the CHP plant, as their contract with the gas supplier imposes penalties if demand exceeds the usual consumption per day (15 hours for two machines operating at full power). If only one plant is operating, longer total run-times are possible.

The current contract between UCC and Dalkia is quite favorable for UCC, but imposes inefficiencies in operating the plant. Any change to minimize total cost or total energy use would reduce overall cost, but increase the cost for UCC, while Dalkia would benefit from these changes.

The CO2 intensity of the Irish grid electricity varies significantly with the amount of wind power generated, and other factors. But on most days, the highest O2 intensity is reached in the early morning hours. At these times, the internal CO2 balance of the CHP plant would reduce overall CO2 emissions. But this is impractical, as electricity demand on the campus at that time is minimal.

The operating modes of the CHP plant can be changed during the daytime. As there is no operator on-site during the night, operation during the night is not possible. The on/off times in the morning and evening can be changed to match requirements of the college.

After the first set of experiments, Dalkia expressed concerns about running the plant at intermediate output levels. According to their experience, this leads to increased maintenance cost and downtimes due to faults. Their strong preference was to either run one or both plants at maximal output, or to follow the electricity demand on the campus exactly. This leads to a drastic simplification of the optimization model, as only very few discrete output levels need to be considered.

4.2 Ucd HSG 6.1m: Heat Demand Load Balancing at Commerzbank Arena Frankfurt

The core of the Use Case HSGZ 6.1 addresses the limited heating capacities at the Frankfurt Commerzbank Arena. Centred on the grass field operation as major commercial preference, the system operation ensures the service for the grass heating demand. If operated in parallel to other systems, service delivery and thermal comfort conditions in other areas like the offices, conference centre, meeting rooms are sacrificed.

According to the operational staff, the problem occurs during the heating season (late Oct – early April), when the heating demand exceeds the capacity of the gas boiler hydronic system. The main demand during this period relates to the grass heating system of about 50% of the peak output of the gas boilers accounting for approximately 40% of annual gas consumption.

The development of the control algorithm took a step-wise approach:

(1) Identification of the energy efficiency potential for the grass heating operation (based on historic data of past years, and first winter season with real-time measurements) documented in [3]
(2) Implement automatic control of grass heating operation to increase the operation performance documented in [4] (and [5])
(3) Balance grass heating demand embedded into stadium operation without sacrificing thermal comfort conditions in other areas of the arena documented in [5]

The following subsection will provide the summary of experimental decisions and reasoning for final advanced control decision.

4.2.1 Traditional control – Identification of efficiency potential

The first analysis of the grass heating characteristics focussed on the understanding of the effect of the different control schemes in winter 2013/2014. The currently deployed grass field operation is built upon two distinct control schemes which can be described as:

- Control 1: Heat Demand logic patterns with either
  - (1.A) Threshold of $T_{external}$ at ~6° C, or
  - (1.B) Time schedule, e.g. 12 hr interval
- Control 2: Setpoint $T_{root,set}$ for grass heating supply temperature

Concluding over several past winter season data sets, the following terms are important to define dynamic setting of the supply temperature setpoint $T_{G,set}$ and heat demand request:

- T_{root} – current value and its trend, its heating and cooling speed on heat demand on/off
- $T_{external}$ – forecast within different time frames related to heating cycles and heating/cooling speed, and
- the difference of T_{root} - $T_{external}$ to estimate the heat demand to equalize the heat loss over the grass field

Aiming for optimization of operation performance and a time-scheduled load balancing concept, we studied in D6.3 [3] in more detail the aspects of:

- Impacting the actual heating time based on $T_{external}$ forecast and metering
- Exploiting T_{root} deviation dynamics for various heat demand request periods
- Apply T_{root} range from historical data into control dynamics
- Consider the heat loss- heat request relationship to drive the time pattern

Results:

We had found two opportunities of improving the grass heating:

1. The current grass heating control system is not driven by T_{root} but governed by $T_{external}$ instead, resulting potentially wasteful heating.
2. The operational staff interference with the grass heating system limits the ability to predict its energy consumption based on $T_{external}$ - in particular when compared to fully automated heating control systems deployed in the sports facility such as the static heating system.

These initial studies indicate a possible reduction of the hourly energy demand of about 0.1 – 0.2 MWh per hour, when the operation aims to limit the heat loss over the grass field. This heat loss limitation could be enforced to restrict the temperature difference T_{root} - $T_{external}$ below a range of about 10-12 K.

Additionally, the insights obtained from the thermal behaviour of the field could allow to schedule a pre-heating of the grass field followed by a complete grass heating switch-off for reducing the grass heating's thermal demand for up to several hours.

The details of these studies can be found in deliverable D6.3 [3].

To evaluate these findings, the main experimental season in winter 2014-2015 enforced a consequent energy efficient grass operation first, and then applied step-wise dynamic control schemes.

4.2.2. Automatic control for increasing grass heating operation performance

As learnt from our studies, the impact of the control mode depends strongly on the environmental context (like climate or ambient temperature, irrigation, wind, etc.). Within the special case of grass pitch heating, the ***driving point for correct operation is the thermal inertia characteristics*** of the grass pitch. This however is strongly dependent on environmental context which needs an integrated evaluation to choose the right control option to serve the operational goals efficiently. It also very much depends on the operational settings how the impact correlates with the context. Traditionally, the grass field operations are controlled via the external temperature, and some human experience

The core idea is to ***modify the control strategy from outside-temperature driven control to root-temperature driven control***. For this, the control the grass heating energy consumption will restrict the root temperature into a small lower temperature band. The real target of the grass heating are in creating optimal grass growing conditions, reflected in the grass root temperature mainly. The external temperature itself is an indicator impacting the grass root conditions, but not only. The conditions are governed by the thermal characteristics and its responsiveness to contextual modifications.

The following measures had been taken:

- For winter 2015, the operational rule sets the grass heating operation to be performed only if the air temperature is below 7^{o}C.
- We started with a basic set of control modes to specify their thermal impact and the grass system behaviour.

The following series of experiments have been performed, are ongoing or prepared in next step, as shown in Figure 5 in relation to the external air temperature (in the reporting period September 2014 – March 2015).

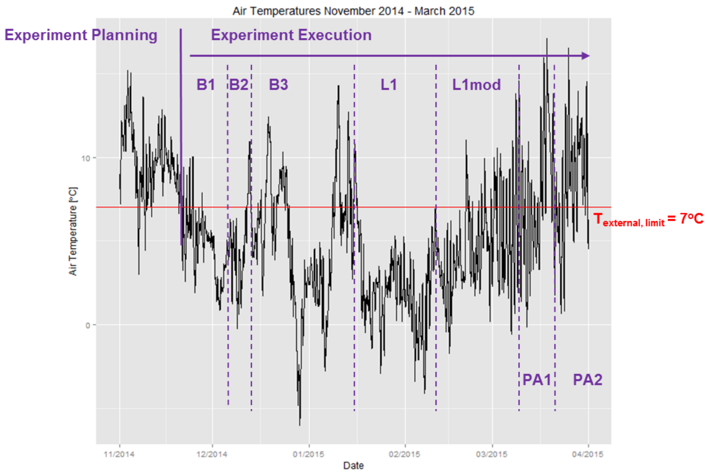

Figure 5: Series of experiments performed during reporting period Sept 2014 - Jan 2015 (incl planning phase), showing experiment condition for $T_{external}$

We experimented with setups of so-called basic experiments labelled as B1, B2 and B3 in order to test control efficiency. The simple control actions are summarized in

Table 1: Experiment plan for basic experiments [5]

No.	Experiment	Time Period	Set points	Condition for control
B1	"static"	24 Nov 2014 – 04 Dec 2014	3 nights each with $T_{G, set}$ = {18°C, 20°C, 22°C } Heat Demand: ON/OFF Time: 6pm – 6am	Heat Demand = ON, if T_{root} < 12.0°C Heat Demand = OFF, if T_{root} > 14.0°C
B2	"variable"	04 Dec 2014 – 11 Dec 2014	$T_{G, set} \in [12\,^{o}C, 22\,^{o}C]$ Heat Demand: ON Time: 6pm – 6am	Set $T_{G, set} = 22\,^{o}C$, if $T_{root} < 12.0^{o}C$ Set $T_{G, set} = 12\,^{o}C$, if $T_{root} > 14.0^{o}C$ Set $T_{G, set} = T_{G, set} + 0.5\,^{o}C$, if $T_{root} < 12.5^{o}C \land T_{G, set} < 22.0^{o}C$ Set $T_{G, set} = T_{G, set} - 0.5\,^{o}C$, if $T_{root} > 13.5^{o}C \land T_{G, set} > 12.0^{o}C$
B3	"variable" + pre-heating	11 Dec 2014 – 16 Jan 2015	$T_{G, set} \in [12\,^{o}C, 22\,^{o}C]$ Heat Demand: ON Time: 6pm – 6am	Set $T_{G, set} = 22\,^{o}C$, if $T_{root} < 12.0^{o}C$ Set $T_{G, set} = 12\,^{o}C$, if $T_{root} > 14.0^{o}C$ T_{supply} Modification rules before 4am: Use B2 logic T_{supply} Modification rules 4am – 6am: Set $T_{G, set} = T_{G, set} + 2.0\,^{o}C$, if $T_{root} < 13.5^{o}C \land T_{G, set} < 20.0^{o}C$ Set $T_{G, set} = T_{G, set} - 2.0\,^{o}C$, if $T_{root} > 13.75^{o}C \land T_{G, set} > 14.0^{o}C$

Result:

Systematic Experiments (November 2014 – January 2015) has been the core of the activities to prepare for the load balancing experiments. With the control logics enforced under B1 – B3, the impact of the coordination on variable supply temperature and on/off control have been quantified under the limitation for night heating operation. The prime aim is to avoid the under- and over-performance of the system.

The details are outlined in deliverable D6.4 [5]. A summary of results is provided in the following figure:

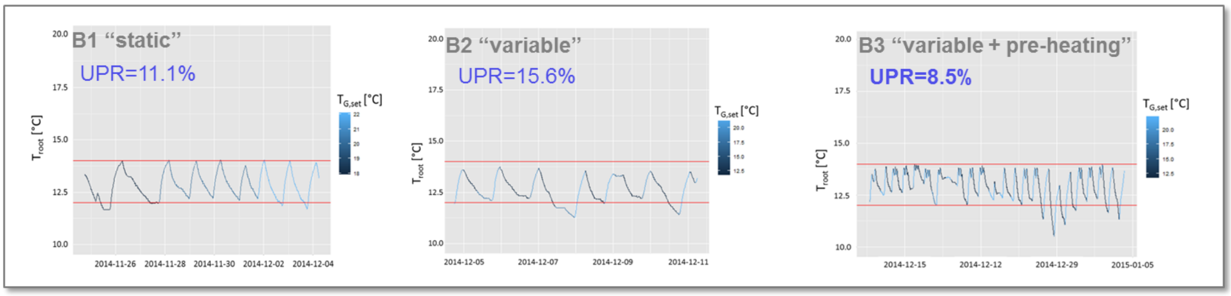

Figure 6: comparison of system performance for the basic experiments B1- B3

The “variable” control mode B2 provided about ~ 25% - 35% energy savings, HDD-normalized.

4.2.3 Load Balancing of Grass operation in context to full arena heating demand

Exploiting the results from the basic experiments operated at night operation (isolated to grass heating scope only), we decided target for an advanced control model based on dynamically predicting the thermal behaviour when shifting to the day-operation embedded into the full arena heating demand.

Load Balancing has been performed in two control modes:

- L1: with a pre-defined heating option into the daily scheme and with B3 night heating + pre-heating
- L1mod: modified version of L1 with a more aggressive day time heating exploiting higher $T_{external}$

The originally planned L2 that would reduce heat demand of non-critical systems in favour of the Grass Heating System under certain conditions was not conducted after consultation with operational staff due to the relatively low impact of non-critical systems when compared to the Grass Heating System dimensions.

The experiments L1 and L1mod build upon the basic experiments and extend the heating phase into daytime.

During daytime the heating capacity constraints can cause underperformance of some of the attached thermal consumption sub-systems. This happens when the thermal output of the stadium’s gas boilers is insufficient for the overall heating demand, resulting in a dropping temperature of the main thermal supply circuit denoted $T_{supply;overall}$. The lower threshold value of $T_{supply;overall} > 80^{o}C$ was defined.

As the heating demand of offices and other arena areas peaks at office hour start, logic B3 was chosen to be used during night time. This choice finishes the night operation with a respective pre-heat cycle, which results in higher T_{root} and with the option to deactivate the grass heating more easily. For daytime grass heating control operation, logic B2 was chosen when $T_{supply;overall}$ exceeds 80°C.

After a successful conclusion of L1, given the little prospect of success of the planned L2 experiments, it was decided to experiment with a more aggressive version of L1 denoted L1mod. Again, during night B3 was chosen to be used. During daytime heating, for colder days when $T_{external} < 5$°C the $T_{supply;overall}$ threshold was lowered to 75°C below which the Grass heating system was switched off. If $T_{external} \geq 5$°C, the $T_{supply;overall}$ threshold was lowered to 70°C below which the Grass heating system was switched off. In both cases For daytime grass heating control operation the logic B2 was used if the $T_{supply;overall}$ threshold was not violated.

The Load Balancing experiment L1 and L1mod conditions are detailed in [5] and summarized in Table 2.

After additional discussions with the operational staff, **<u>the experiment L2</u>** was not considered to promise significant energetic savings beyond L1 due to the relatively large Grass Heating system demand. Thus it was decided to utilize the precious time window remaining with the more promising L1mod.

Table 2: Experiment plan for load balancing experiments [5]

No.	Experiment	Time Period	Set points	Condition for control
L1	"day"	16 Jan 2015 – 11 Feb 2015	$T_{supply} \in$ [12 °C, 22 °C] Heat Demand: ON Time: full day	Between 6 pm – 6am: Logic as in B3 Day time heating: Logic as in B2 with additional check on overall heating supply temperature Set T_{supply} = 12°C if $T_{supply;overall} < 80$°C If $T_{supply;overall} \geq$ 80°C use Logic of B2
L1mod	"aggressive"	11 Feb 2015 – 11 Mar 2015	$T_{supply} \in$ [12 °C, 22 °C] Heat Demand: ON Time: full day	Between 6 pm – 6am: Logic as in B3 Day time heating: If $T_{external}$<5°C: Set T_{supply} = 12°C if $T_{supply;overall}$ < 75 °C If $T_{supply;overall} \geq$ 75°C use Logic of B2 If $T_{external}$≥5°C: Set T_{supply} = 12°C

				if $T_{supply;overall} < 70^oC$ If $T_{supply;overall} \geq 70^oC$ use Logic of B2
L2	"load"	cancelled	$T_{supply} \in [12\ ^oC,\ 22\ ^oC]$ Heat Demand: ON Time: full day	Between 6 pm – 6am: Logic as in B3 Day time heating: Set $T_{supply} = 12^oC$ if $T_{supply;overall} < 70\ ^oC$ (tbc) If $T_{supply;overall} \geq 80^oC$ use Logic of B2 If $70^oC \leq T_{supply;overall} < 80^oC$ reduce heating in low priority heating systems

Result:

Systematic and extensive Experiments (January – March 2015) were conducted. The control logics B2 and B3 were reused.

The details are outlined in deliverable [5]. A summary of results is provided in the following figure:

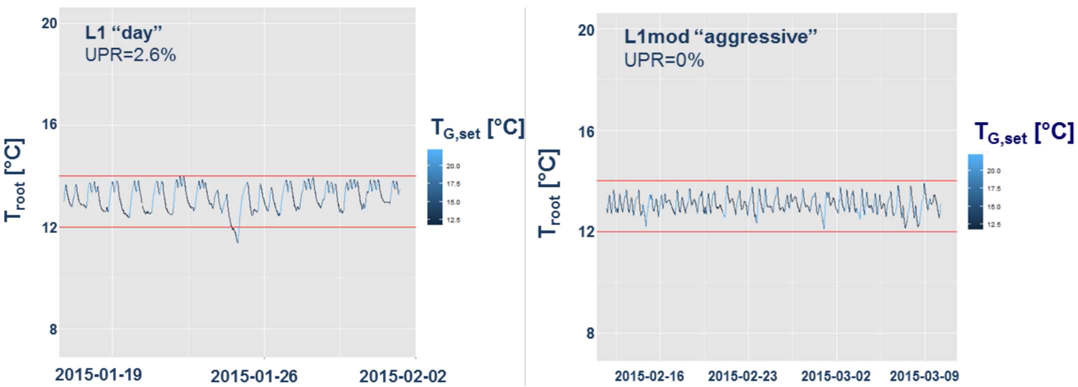

Figure 7: comparison of system performance for the basic experiments L1 and L1mod

Not surprisingly, using daytime for heating reduced UPR compared to the basic experiments. While L1 exhibited similar normalized energy demand to B3, L1mod used significantly less energy – caused in part due to the higher air temperatures (despite the HDD normalization we applied). However, even when considering only the colder February subset of the experiment L1mod experiment phase, its performance is better than L1. Detailed analysis of the experiments is documented in [5].

4.2.4 Preheating with advanced forecasting of grass root temperature trends

Since discussions with the operational staff indicated that the experiments of load balancing reached the limits of the potential energy savings and time left for experiments until the end of the heating season at the end of March, it was decided to investigate the potential of applying enhanced methods for forecasting T_{root} trends in the context of pre-heating.

With the expected rising of $T_{external}$ it was decided to apply these experiments during night time heating only to avoid influence from daytime heating concepts.

Several studies with the standard model of linear regression and the advanced non-linear regression method using Deep Belief Networks were conducted to assess the T_{root} forecast performance. Ultimately, two sets of **experiments PA1 and PA2** with Deep Belief Networks.

During night time (before 4 am) the **experiment PA1** forecast a first 6 hour T_{root} trend using the recent 6 hour histories of T_{root} and $T_{external}$ under the assumption no heating was in effect. Then, combining this 6 hour T_{root} trend with weather forecast information for another 6 hour T_{root} forecast led to a 12 hour prediction of T_{root}. If this forecast predicted a violation of the minimum T_{root} (12°C) threshold, T_{root} was forecast for 6 hours under different T_{Gset} settings. Then the heating scenario which did not violate the maximum T_{root} (14°C) threshold was chosen. In the last two hours of the night (4 am – 6 am) the same logic was applied, but the maximum T_{root} threshold violations were only considered if they occurred within the first 3 hours of heating – this in turn led to a pre-heating effect similar to B3.

During night time (before 4 am) the **experiment PA2** behaved as PA1 with the modification that directly a 12 hour T_{root} trend was forecast using a single DBN and a larger history of information.

We experimented with setups of so-called basic experiments labelled as *PA1* and *PA2* in order to test control efficiency. The control actions are summarized in Table 3.

Table 3: Experiment plan for advanced pre-heating experiments [5]

No.	Experiment	Time Period	Set points	Condition for control
PA1	"DBN6"	11 Mar 2015 – 18 Mar 2015	$T_{G,set} \in$ [12 °C, 22 °C] Heat Demand: ON/OFF Time: night	Between 6 pm – 4am: Forecast 12 hour T_{root} trend without heating. If violation of $T_{root,min}$: Forecast 6 hour T_{root} trends in different $T_{G,set}$ scenarios. Choose largest $T_{G,set}$ not violating $T_{root,max}$ Between 4 am – 6am: Same as before, but use 3 hour trend for TG,set determination
PA2	"DBN12"	20 Mar 2015 –	$T_{G,set} \in$ [12 °C, 22 °C]	Between 6 pm – 6am: Same as PA1 but use

		31 Mar 2015	Heat Demand: ON/OFF Time: night	for 12 hour cooling trend a single DBN.

Result:

Systematic Experiments (March 2015) were conducted. Due to the higher external temperatures in March, the reported energy savings are drastic despite HDD normalization and should be considered with caution - the details are outlined in deliverable [5]. UPR was lower than with the other night time heating experiments (B1-B3). A summary of UPR results is provided in the following figure:

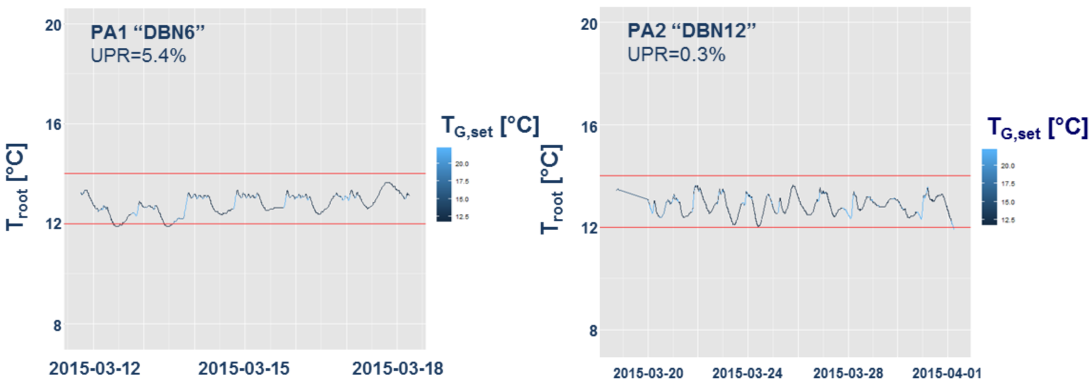

Figure 8: comparison of system performance for the basic experiments PA1 and PA2

5. PERFORMANCE EVALUATION RESULTS

5.1 Ucd UCC 6.1: CHP Optimization on UCC Campus

In this section we evaluate the quality of the different components of the application layer of the UCC Campus use case. We begin with the data cleaning operation, then describe the predictors we are using and check their quality over a period of multiple years, and then look at the scheduling results of the load-balancing tool.

Data Cleaning

The data cleaning operation takes the raw input data, which may be missing some measurement values, rejects data outside the acceptable range of measurements, and then repairs the values with different strategies. Table 13 in the appendix shows the values ranges for the ResourceKraft sensors that are used. The values are expressed per time interval of 15 minutes, in the unit used inside the load balancing application. Note that the ResourceKraft DHW does not seem to apply any repair methods itself. In the Campus21 architecture, neither the storage component nor the middleware perform validity checks, so that these have to be applied at the application layer.

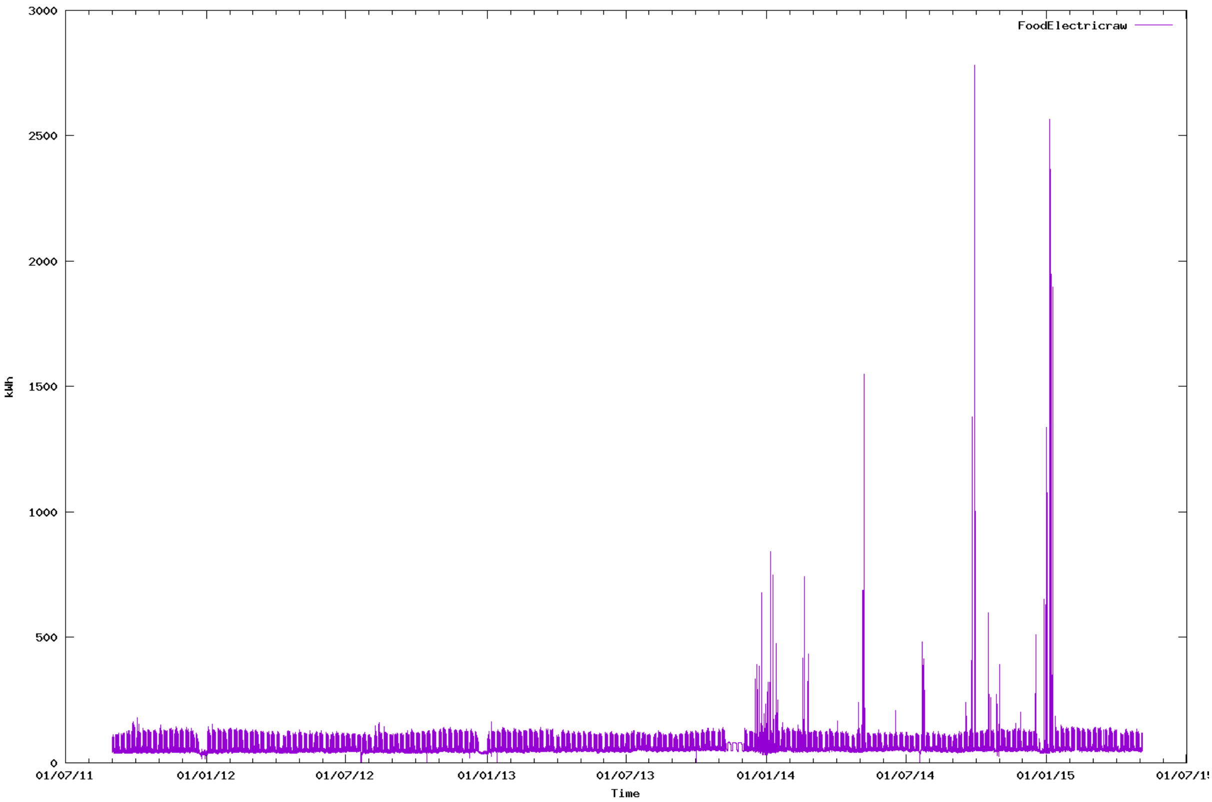

Figure 9: Food Sciences Building Electricity Consumption Raw Data (11-2011-05/2015)

Figure 9 shows a plot of the electricity consumption for the Food Sciences building from November 2011 to May 2015. Note, that from the beginning of 2014 spurious high measurements are reported. A check with UCC B&E confirmed that these spikes do not represent real values, but are due to sensor failure. The data cleaning operation repairs the values as shown in Figure 10.

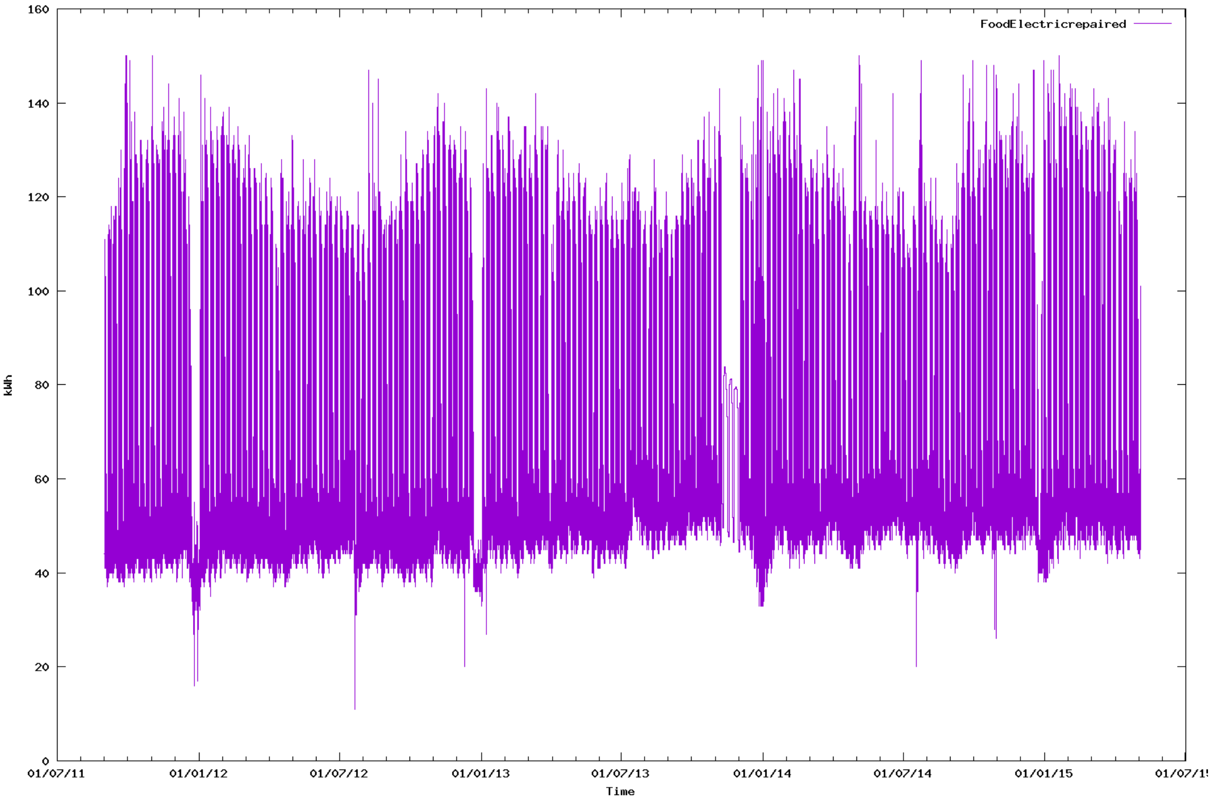

Figure 10: Food Sciences Building Electricity Consumption Repaired Data (11/2011-05/2015)

A decrease of the electricity consumption during Summer and a significant drop in consumption over the Christmas break are clearly visible.

Forecast Quality

We now consider the quality of the forecasts for the different elements of the system. As we are planning 24 hours ahead, it is important that the quality of the forecasts is high enough to generate meaningful schedules

Quality Indicators

We use the following functions to evaluate the quality of forecasts, using n data points with actual values a_i and forecasted values f_i.

$$\text{MAE} = \frac{\sum_{i=1}^{n} |f_i - a_i|}{n} \tag{5.1}$$

$$\text{MAE Over} = \frac{\sum_{i=1}^{n} \max(0, f_i - a_i)}{n} \tag{5.2}$$

$$\text{MAE Under} = \frac{\sum_{i=1}^{n} \max(0, a_i - f_i)}{n} \tag{5.3}$$

$$\text{MAPE} = 100\frac{\sum_{i=1}^{n} |\frac{f_i - a_i}{a_i}|}{n} \tag{5.4}$$

$$\text{SMAPE} = 100\frac{\sum_{i=1}^{n} |f_i - a_i|}{\sum_{i=1}^{n}(f_i + a_i)} \tag{5.5}$$

$$\text{MSE} = \frac{\sum_{i=1}^{n}(f_i - a_i)^2}{n} \tag{5.6}$$

$$\text{SQRT(MSE)} = \sqrt{\frac{\sum_{i=1}^{n}(f_i - a_i)^2}{n}} \tag{5.7}$$

Campus Electricity Demand

Table 4 compares different baseline forecasts for the electricity consumption on the campus with our feature based predictor. We report mean absolute error (MAE), mean absolute error below exact value (MAE Under), mean absolute error above exact value (MAE Over), mean absolute percentage error (MAPE), symmetric mean absolute percentage error (SMAPE), mean squared error (MSE) and the square root of the mean squared error (SQRT(MSE)) as defined in the previous Section. As baseline comparison we use the data from one hour ago (Hour), from one day ago at the same time (Day), from one week ago (Week), from 52 weeks ago (52 Weeks), a mixed forecast (Mixed) and our feature base forecast. The values are reported for the time period from beginning of 2013 to the 5th of May, 2015.

Table 4: Comparison of UCC Campus Electricity Demand Forecasts (1/1/2013 - 5/5/2015)

Method	MAE	MAE Under	MAE Over	MAPE	SMAPE	MSE	SQRT(MSE)
Hour	26.98	13.50	13.48	7.02	3.68	1808.67	42.53
Day	50.44	25.26	25.18	12.98	6.88	9468.78	97.31
Week	26.67	13.85	12.82	7.53	3.64	3405.86	58.36
Week52	22.13	7.26	14.87	6.28	2.99	1595.11	39.94
Mixed	17.72	8.42	9.30	5.07	2.41	920.19	30.33
Forecast	15.72	8.08	7.64	4.63	2.15	645.95	25.42

The mixed forecast uses the following decision tree:

- If the current day is a bank holiday, use the data from the same holiday last year (i.e. for Easter Monday 2015 use data from Easter Monday 2014), even if the dates are different.
- If the day one week ago was not a public holiday, and the university calendar type is the same as for the current day, use the value from a week ago.

- If the corresponding day one year ago was not a public holiday, use the data from that day.
- Otherwise go back one week at a time, until you find a day that was not a public holiday.

The corresponding day one year ago is either the day 52 weeks ago, if it is not a holiday and the academic types match, or the day of the same academic type and academic week number a year ago, if it was not a public holiday.

We can see that the value from one week ago at the same time is a better predictor than the value from one hour ago, or from one day ago. The value one year ago is better still, indicating that the electricity demand does not change much with weather conditions, but is largely determined by the academic calendar. The mixed forecast avoids problems with forecasts for (or before and after) bank holidays, but our feature based predictor is even better, without requiring a detailed understanding of the underlying mechanism.

Campus Heat Demand

Table 5 shows the corresponding analysis of the heat forecast, predicting the sum of all heat consumption values used. The quality of the forecast is much worse than the electricity demand forecast, but still better than any of the available fixed forecast values.

Table 5: Comparison of UCC Campus Heat Demand Forecasts (1/1/2013-5/5/2015)

Method	MAE	MAE Under	MAE Over	MAPE	SMAPE	MSE	SQRT(MSE)
Hour	23.94	11.98	11.96	NaN	14.35	4040.78	63.57
Day	36.13	18.08	18.05	NaN	21.65	6083.61	78.00
Week	32.23	16.34	15.89	NaN	19.36	4806.90	69.33
Week52	62.72	50.02	12.70	NaN	48.41	11613.07	107.76
Mixed	37.45	23.28	14.17	NaN	23.74	6147.18	78.40
Forecast	29.06	13.24	15.83	NaN	17.15	3058.08	55.30

Note that the MAPE value is not defined for the heat consumption, as the actual value is zero for many time periods during nights, and during the Summer shutdown period.

National Grid CO2 Intensity

As input to the optimization routine we need a prediction of the CO2 intensity of the grid electricity, as this influences the objective function when we minimize CO2 emissions. The results of our forecast compared to some baseline cases are shown in Table 6. An important factor influencing the CO2 is the wind energy production in the country, which largely follows the weather and does not show a regular day/week pattern. This means that the forecast error is much smaller than the error of the base line comparisons.

Table 6: Comparison of National Grid CO2 Intensity Forecasts (1/1/2013-5/5/2015)

Method	MAE	MAE Under	MAE Over	MAPE	SMAPE	MSE	SQRT(MSE)
Hour	20.64	10.40	10.24	4.51	2.27	901.99	30.03
Day	58.62	29.39	29.23	13.40	6.44	6151.30	78.43
Week	75.64	38.05	37.59	17.31	8.31	9555.08	97.75
Week52	86.51	33.61	52.89	20.19	9.30	11824.60	108.74

Mixed	76.97	38.44	38.53	17.66	8.45	9823.81	99.12
Forecast	41.01	20.96	20.06	9.34	4.51	2843.94	53.33

National Grid System Load

In order to evaluate the quality of our forecasting methodology, we also wanted to compare our forecast against an existing forecast for national system load, i.e. electricity demand for the Irish market. SEMO produces different versions of a forecast for the electricity demand that is published at the same time as the price forecasts. Table 7 compares these forecasts with our own model for a 24-hour ahead load forecast.

Table 7: Comparison of National Electricity Load Forecasts (1/1/2013-5/5/2015)

Method	MAE	MAE Under	MAE Over	MAPE	SMAPE	MSE	SQRT(MSE)
Hour	201.27	100.63	100.64	5.45	2.65	72073.55	268.47
Day	249.41	124.90	124.51	6.64	3.28	128270.14	358.15
Week	198.82	100.79	98.03	5.34	2.62	76439.20	276.48
Week52	191.88	101.17	90.71	5.18	2.53	64640.78	254.25
Mixed	193.15	98.68	94.47	5.20	2.54	65949.52	256.81
SEMO EA	232.90	8.23	224.67	6.47	2.98	83891.75	289.64
Forecast	113.43	60.21	53.22	3.05	1.50	40250.89	200.63

Note that the forecast error in the SEMO Forecast is quite asymmetrical, it very rarely overestimates consumption, but often underestimates the load significantly. This is probably due to the methodology used. Our forecast is more balanced, and overall has about half the error of the SEMO Forecast.

Existing Grid Electricity Price Forecast

We also evaluate the different, existing forecasts for the grid electricity price published by the market operator SEMO. Table 8 shows the statistical analysis based on data from January 2013 to May 2015. This compares the first 24 hours ahead price forecast EA with later forecasts, against the definite price EP2 that is published after the event. The EA forecast is typically available at 10:00 with a 24 hour forecast starting at 06:00 on the next day, the EA2 is available later on the day. The WD1 forecast is published at noon on the day of operations, while the EP1 is an ex-post price published on the day after operations. The MAPE value is not defined for these forecasts, as the actual price is zero (or even negative) for some time periods.

Table 8: Comparison of SEMO Price Forecasts (1/1/2013-5/5/2015)

Method	MAE	MAE Under	MAE Over	MAPE	SMAPE	MSE	SQRT(MSE)
SEMO EA	15.47	8.76	6.71	n/a	13.06	963.56	31.04
SEMO EA2	15.64	8.59	7.05	n/a	13.15	972.06	31.18
SEMO WD1	15.41	8.10	7.31	n/a	12.88	910.83	30.18
SEMO EP1	9.54	4.79	4.75	n/a	7.92	515.90	22.71

We can see that the MAE, SMAPE and MSE values for the EA2 forecast are worse than for the earlier EA forecast. This indicates that we are better off using the EA forecast only in producing our schedule. Even the WD1 forecast (during execution of the plan) is not much better than the EA forecast. Even the initial ex-post values in the EP1 series have a significant error attached to them.

Experimental Results

We performed tests of the full load balancing system on the plant installed at UCC during two weeks in 2014, supported by Dalkia, the CHP plant operator, and Buildings & Estates (B&E) from the university side. We first describe the experimental protocol followed during these experiments, and then recall key results for the two test weeks. The material is covered in more detail in Deliverables D6.3 and D6.4. At the end, we investigate the impact of using forecasting on the schedule quality.

Experimental Protocol

The protocol to run the experiments was taking the following form: The schedule was produced for one day at a time, during the day before the event.

At some time on the day before the run, we triggered the data collection from all internal and external sources. During the tests we varied the collection time in the interval from 09:00 to 15:00, to confirm our understanding of the impact of that data collection cutoff point. The SEMO SMP price forecast, although published in theory at 06:00 every morning, on some days did not become available until 10:00. In the second test week data feed for our tool included input from B&E in the form of an event calendar that helps to understand special demand patterns on the campus.

Once all data was collected, the analysis and forecasting tools were run to clean the data and produce forecasts for all relevant parameters for the next day.

Given the forecasts, the scheduler produced multiple variants of the schedule based on changes of the objective function and the constraints considered. Most of these schedules are generated for comparison only, only one of the schedules was transferred to operations.

The provisional, generated schedule was transferred electronically to Dalkia and in report form to both Dalkia and Buildings and Estates for comment. A short production conference call was called to allow for discussion and potential changes based on the feedback provided. On two days of the second test week, the proposed schedule was modified to take preferences from Dalkia into account, and a new schedule was transmitted. The objective was to have an agreed schedule in place at the latest on 17:00 on the day before, to define the schedule of the CHP plant from 00:00 to 24:00 on the next day.

Dalkia then implemented the schedule on the actual devices during the day. This involves both on-site personnel, and their off-site operations center.

Data about the actual performance on the day was collected and evaluated during and at the end of each day, so that performance analysis could be completed.

As the real-time electricity price is only known four days after the event, the exact actual cost of the schedules could only be computed the following week.

First Week of Experiments, April 2014

The first tests of the system were performed in April 2014, in the week before Easter. The Friday of that week is Good Friday, which is not a public holiday in Ireland, but on which the operation on the campus is quite limited. This was a good test to check for the ability of the demand forecast to deal with this exceptional case. Figure 11 shows the predicted versus actual electricity use on the campus, indicating that the semi-holiday was recognized well.

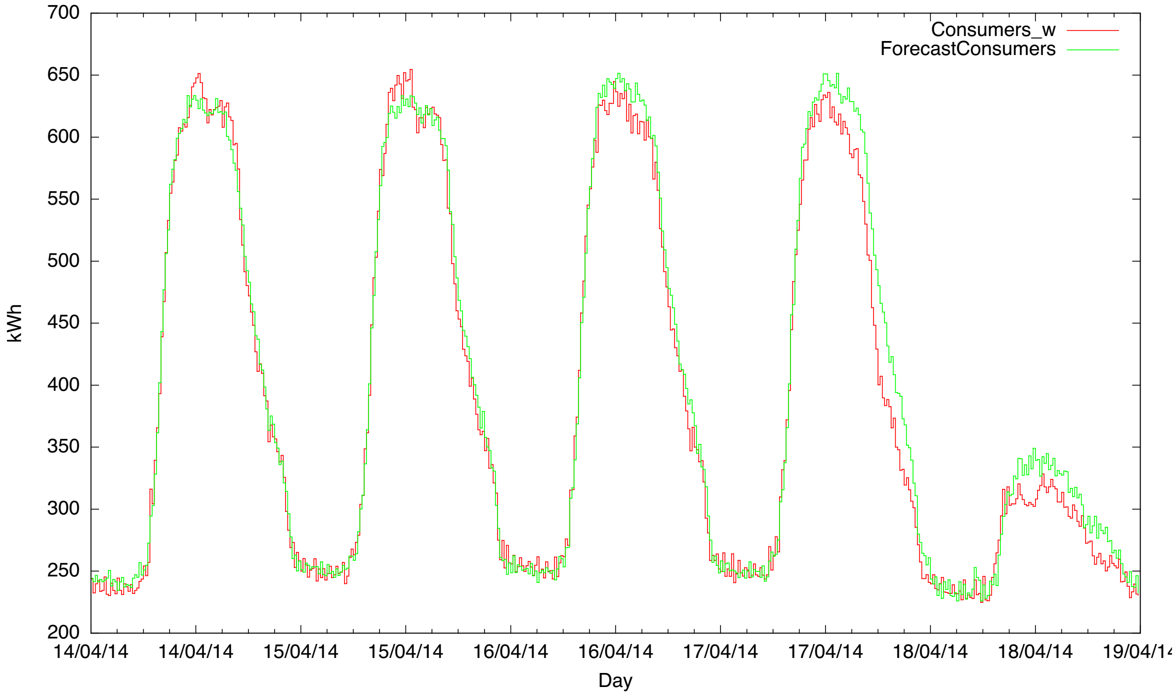

Figure 11: Campus Electricity Demand, Actual and Forecast, Week 14-18 Apr, 2014

Figure 12 shows the proposed schedules for the CHP plant, as delivered to Dalkia each day. The schedules for Monday, Tuesday, Wednesday and Friday minimize operational cost, while respecting operational constraints. The schedule for Thursday runs a continuous schedule for both plants, and was used to calibrate the model of the internal power consumption of the CHP plant in these modes.

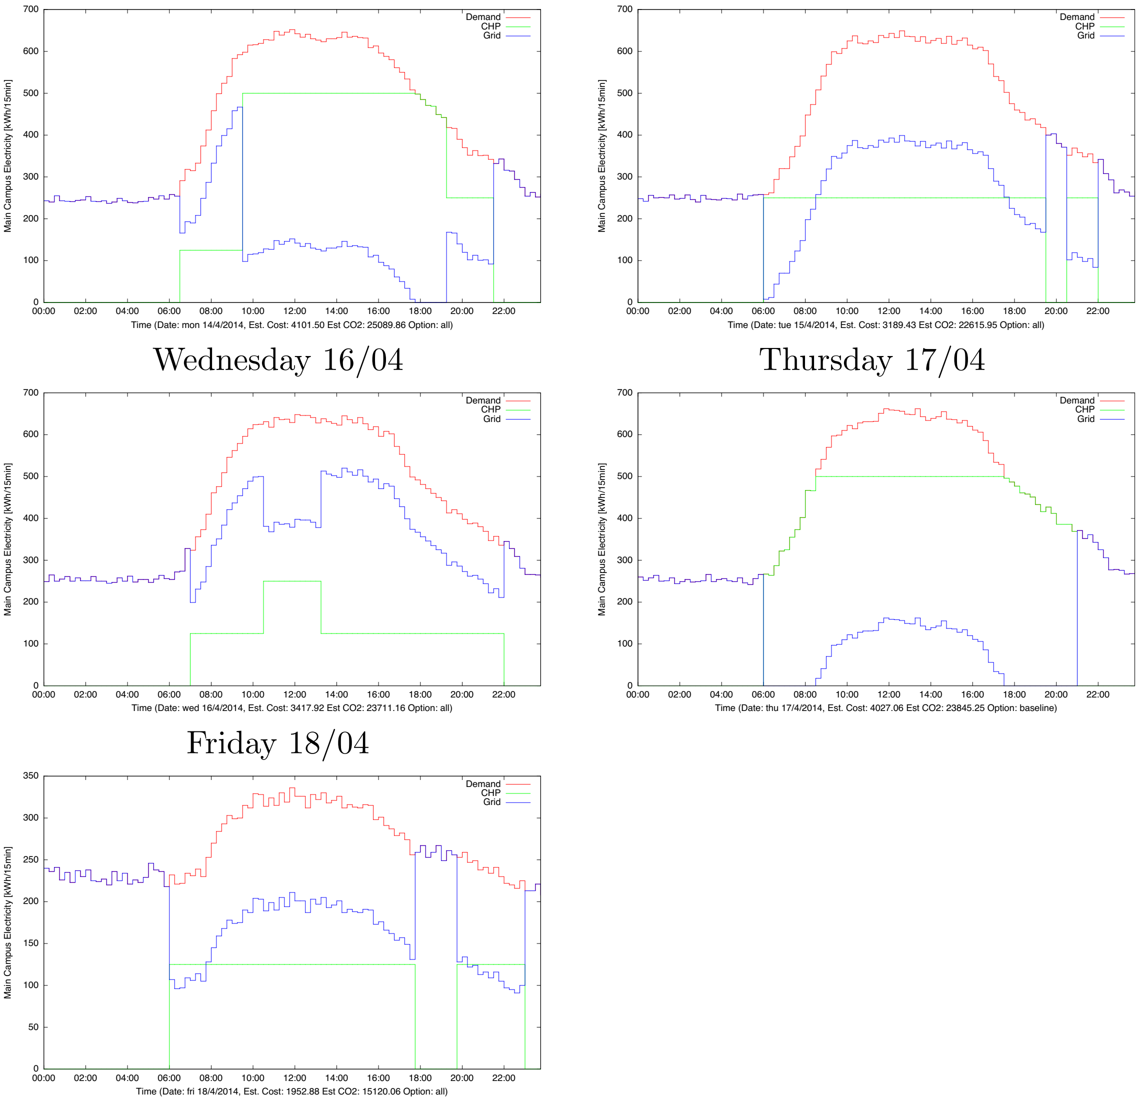

Figure 12: Proposed Schedules for First Test Week, April 14-18, 2014

The recorded output of the CHP plant is shown in Figure 13, which differs from the planned schedule mainly due to unforeseen downtimes, repairs and testing.

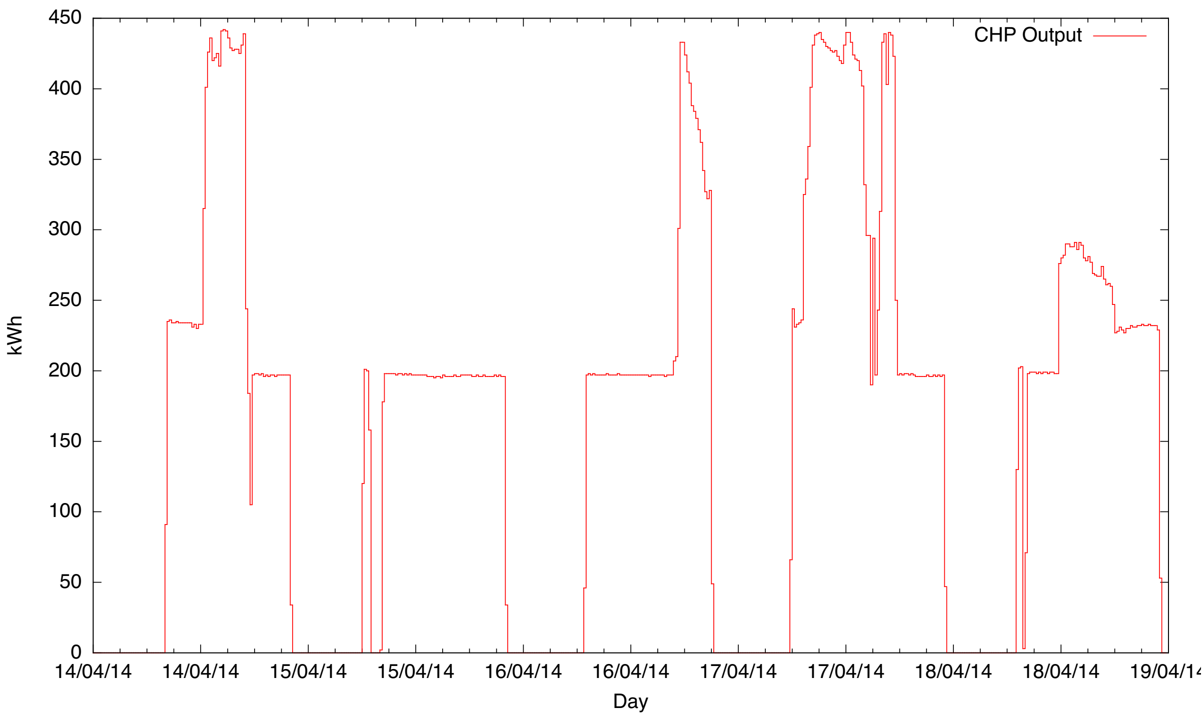

Figure 13: Actual Schedule for First Week of Experiments

The tests confirmed the validity of the experimental protocol, but also highlighted a number of additional constraints of the operator:

- Dalkia expressed a strong preference to run the plants only at maximal power levels, with either both or only a single plant operating. Intermediate power level are only acceptable when following the electric demand on the campus exactly, typically ramping up in the morning and ramping down in the evening.
- In order to meet the contractual obligations towards the college, Dalkia wants to operate the plant for at least 15 hours per day.
- At the same time, running for more than 15 hours at full power on both plants pushes the gas consumption into the penalty band of the contract with the gas supplier. This would incur significant extra monthly cost, and should be avoided.
- Operation between 23:00 and 05:00 on the next morning is only possible in exceptional cases, as there is no human operator on duty for this period.

Second Week of Experiments, November 2014

For the second week of experiments in November 2014 we modified the scheduler to take the extra constraints expressed by Dalkia into account. We introduced two new constraint settings, **FullPower** and **HalfPower**, to control the use of the extra constraints. In the **FullPower** mode, the plant has to run for 15 hours, but not during 23:00 to 05:00. If the plant is shut down, it has to stay off for at least 2 hours. If the plant is turned on, it must run for at least one hour. The plant can be off, run at full power on both engines, or follow the electricity demand on the campus exactly. In the **HalfPower** mode, the plant can also be run on full power with either a single or both engines, while keeping all other constraints. We compare this against a **Baseline** mode where the plant is turned on at 07:00, runs for 15 hours at maximal output, and then is turned off, and against a **Pure** mode, where no additional constraints apply.

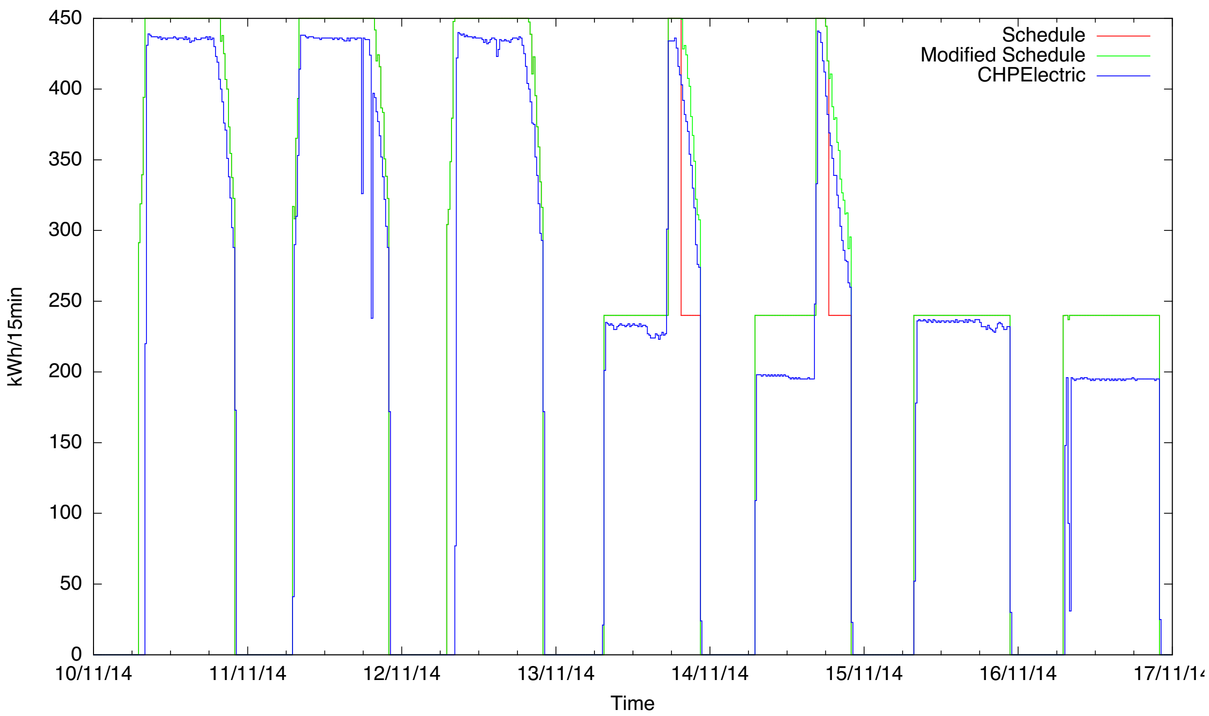

Figure 14: Schedule for Second Week of Experiments

Figure 14 shows the proposed schedules for the week, including the weekends, as well as the modified schedule after discussion with Dalkia, and the recorded output of the CHP plant for the week of operation. In the production conference call foreseen in the experimental protocol Dalkia suggested to extend the run of the second engine during Thursday and Friday evenings, a safeguard against real-time electricity price spikes not predicted by the SEMO price forecast. The **FullPower** mode was tested on Monday to Wednesday; the **HalfPower** mode was used on Thursday to Sunday. Table 9 compares the computed costs of the different scheduling variants, highlighting the selected schedule in yellow, and the actual results in green.

Table 9: Predicted vs. Actual Cost

Day	Baseline	FullPower	HalfPower	Pure	Actual	Percentage
Monday	4306.48	4288.56	3786.16	3433.50	4444.49	3.2
Tuesday	4476.24	4412.03	3978.43	3851.18	4480.01	0.08
Wednesday	4349.66	4323.38	3880.55	3693.83	4317.22	0.75
Thursday	4347.25	4323.01	3780.70	3468.73	4867.42	28.7
Friday	4126.39	4123.01	3767.26	3593.37	3948.71	4.8
Saturday	2963.69	2938.76	2606.41	2246.23	2905.04	11.5
Sunday	2337.24	2326.98	2186.23	1559.94	2408.36	10.2

The computed and actual values are quite similar, with the exception of Thursday and the weekend. On Thursday the uncorrected measurement of the Main Restaurant heat meter increased the "actual" consumption values significantly. For the weekend, the demand forecasts were less accurate, leading to a larger error in the predicted cost.

Figure 15 shows the evolution of the total cost over the week, in 15 minute intervals. We see the total cost, the share of the cost paid by UCC, and the share paid by Dalkia. The cost value for Dalkia is negative at most time periods, indicating their operating margin, and only increases significantly when importing electricity from the grid at high prices.

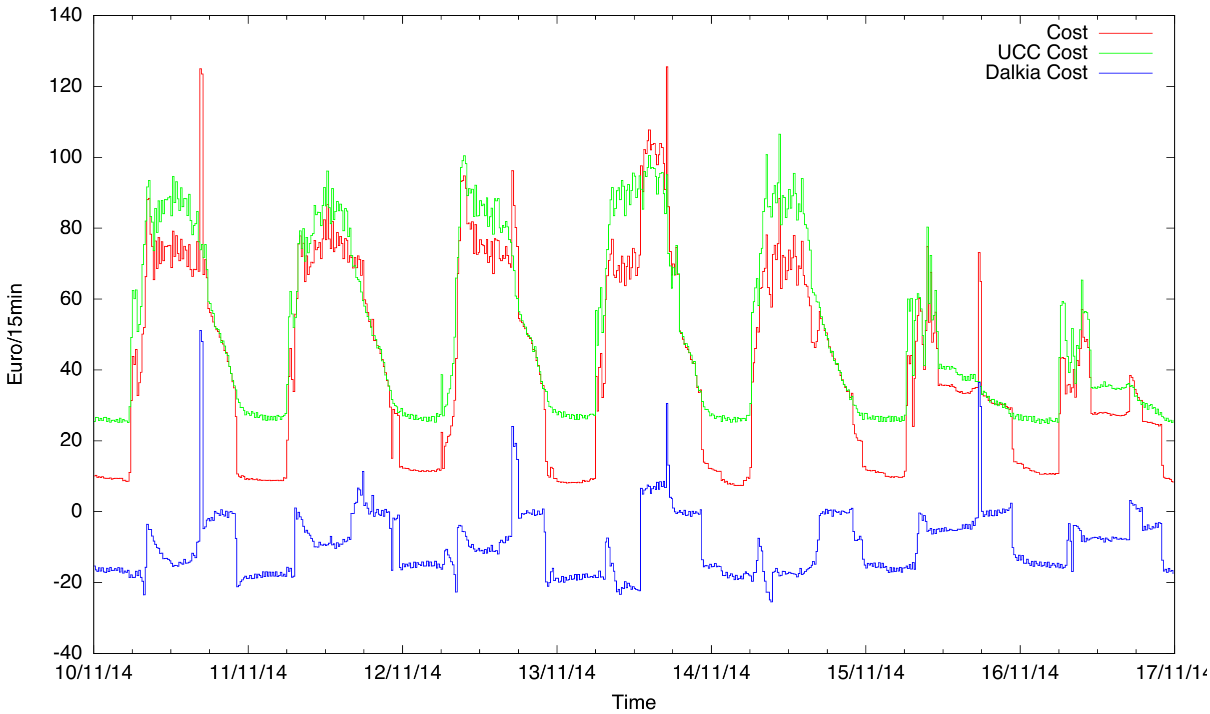

Figure 15: Total Cost Over Time

Figure 16 shows the different elements for the UCC cost component. The majority of the cost is incurred by the electricity usage (paid for at a fixed price to Dalkia), the gas cost for operating the steam boiler is a much smaller cost factor.

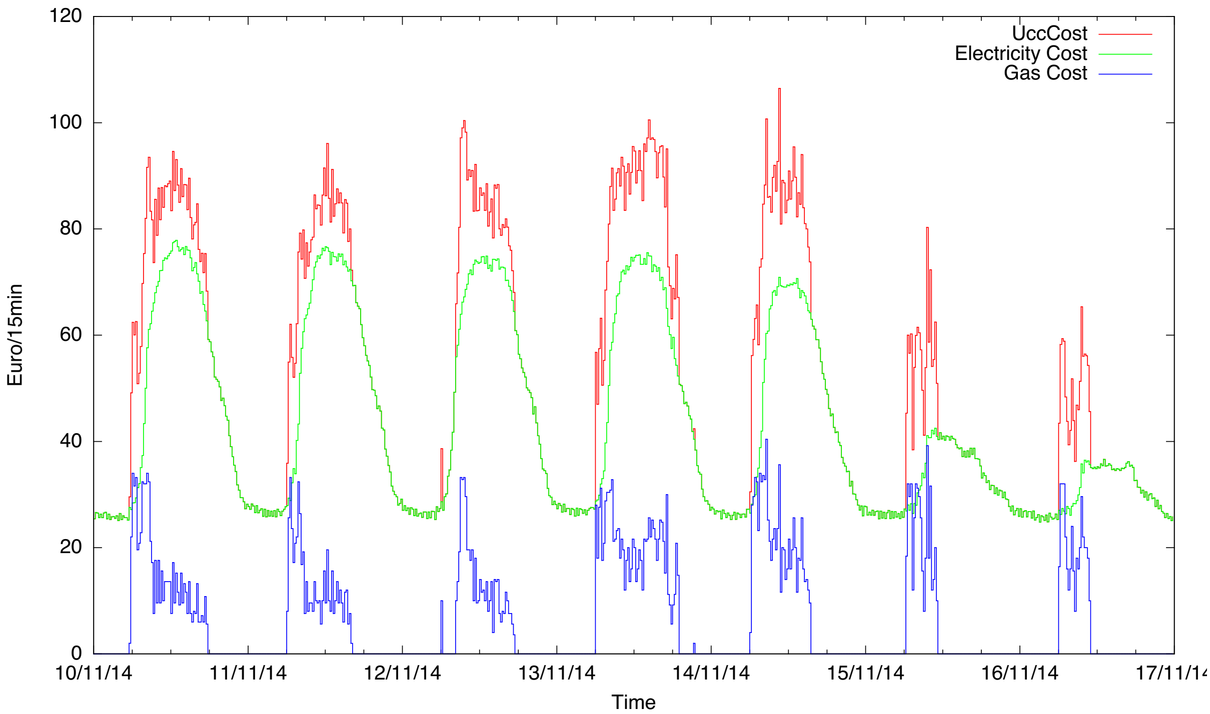

Figure 16: Elements of UCC Cost over Time

Figure 17 shows the elements of cost for Dalkia. The green curve is the amount paid by UCC for the electricity consumed, charged at a fixed price. The blue curve is the cost of operating the CHP plant, while the magenta curve is the cost of importing electricity from the grid at market price.

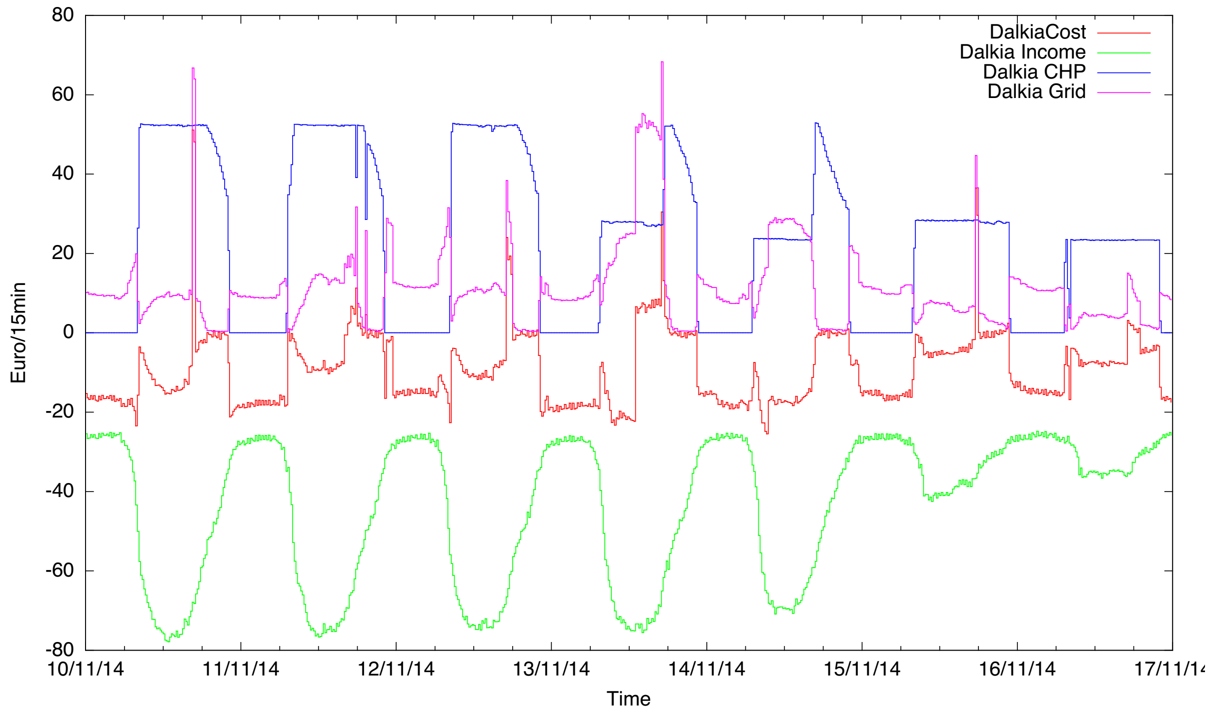

Figure 17: Elements of Dalkia Cost over Time

An important point affecting the ability to perform more experiments on the actual plant is demonstrated in Table 10, which compares the computed costs for the different scheduling variants for Monday, November 10th, 2014. The **Pure** model has a much lower overall cost (Total) compared to the **Baseline** (a saving of 873 €, or more than 20%), but the cost for UCC increases by 3750 €, while the cost for Dalkia turns into a profit. For the other modes, the potential savings are smaller, but the imbalanced cost change is similar. This indicates that the current contractual situation is highly favorable to UCC, and any changes will lead to higher cost for the college. At the same time, this fact limits the opportunity for performing experiments on the actual plant, as the cost to the college becomes prohibitive.

Table 10: Computed Scheduling Variants for Monday, November 10th, 2014

Model Variant	CO2	Total	CHP	Boiler	Grid	UCC	Dalkia
BASELINE	26734.87	4306.48	3383.25	317.93	605.30	1955.17	2351.31
PURE	26359.18	3433.50	262.80	1256.69	1914.00	5799.14	-2365.64
FULLPOWER	26560.66	4288.56	3219.17	394.78	674.61	2198.34	2090.22
HALFFULLPOWER	25982.54	3786.16	1823.53	683.52	1279.11	3763.85	22.31

Impact of Forecast on Schedule Quality

The two weeks of experiments with the actual plant showed that the experimental protocol is feasible, and that results obtained by scheduling with forecasted data produces results that are close to actual values. One unanswered question is what overhead is caused by using the forecast data to generate the schedule. For this we compare a hypothetical schedule produced with ex-post information to the schedules compared with the forecasted data, and both to a fixed baseline schedule. The process is shown in Figure 18. After the event, we can run the optimization module (on the left hand side) with ex-post data to obtain a solution based on perfect knowledge, i.e. exact electricity prices, exact heat and electricity demand, and CO2 intensity. We use that solution to find the effect of using forecast data in the decision making. This follows the process on the right hand side. We first produce the schedule with forecasted data, and then evaluate its actual cost based on the ex-post data. This cost is then comparable to the value produced on the left.

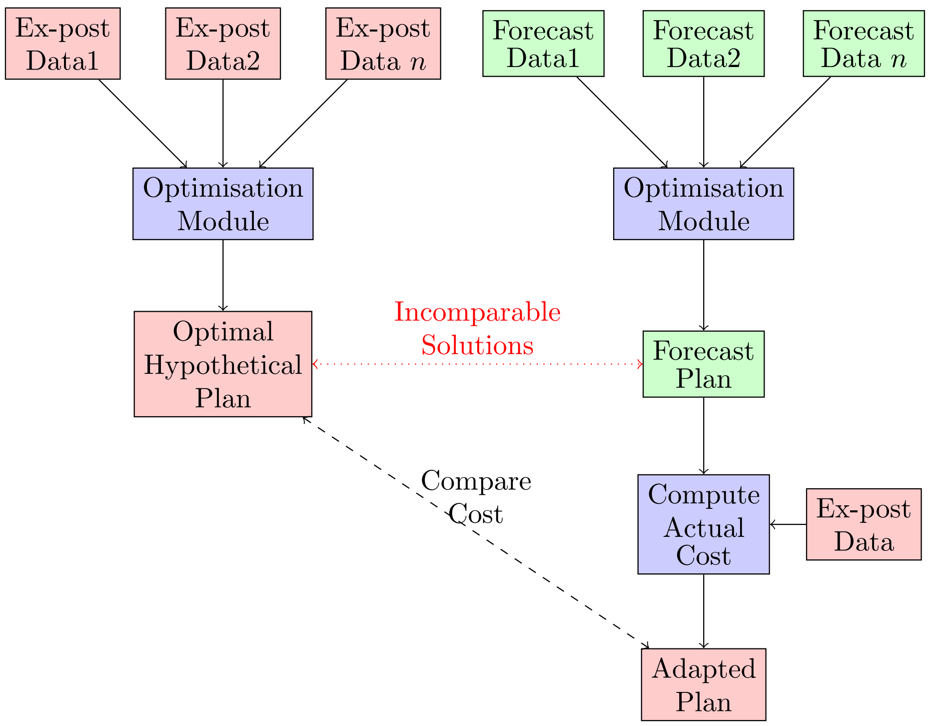

Figure 18: Evaluation of Optimization Process

For our analysis we use the daily forecasts for a period from 1/1/2013 to 30/4/2015, and perform the load balancing optimisation for each day with both the forecasted data and the recorded, actual values. For most days, the results are very similar to those obtained in our experiments although both seasonal and weekly effects change the overall cost of the schedule. Figure 19 shows the daily cost for the **Baseline**, the **FullPower** and **HalfPower** variants, and the **Pure** model. The **Baseline** cost is the highest value on most days, closely followed by the **FullPower** model variant. The **Pure** model achieves good savings on most days, and the **HalfPower** result lies between the two extremes. On a few, exceptional days, the generated schedules perform poorly, due to large price spikes at unexpected times of the day.

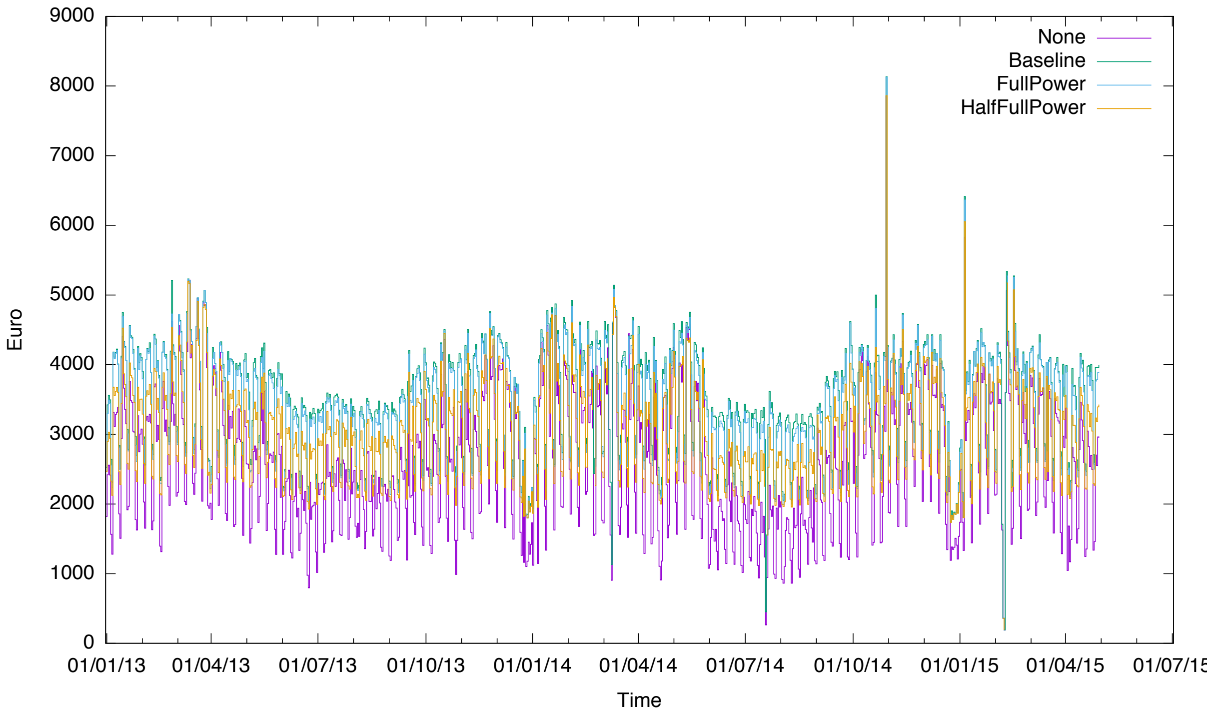

Figure 19: Daily Cost for Model Variants, Schedule Based on Forecast, Evaluated on Observed Values

Figure 20 shows the cost reduction (in percent) for the result of the **HalfPower** model over the **Baseline** result. The average reduction is around 10%, but the value varies from day to day, and is negative on some days.

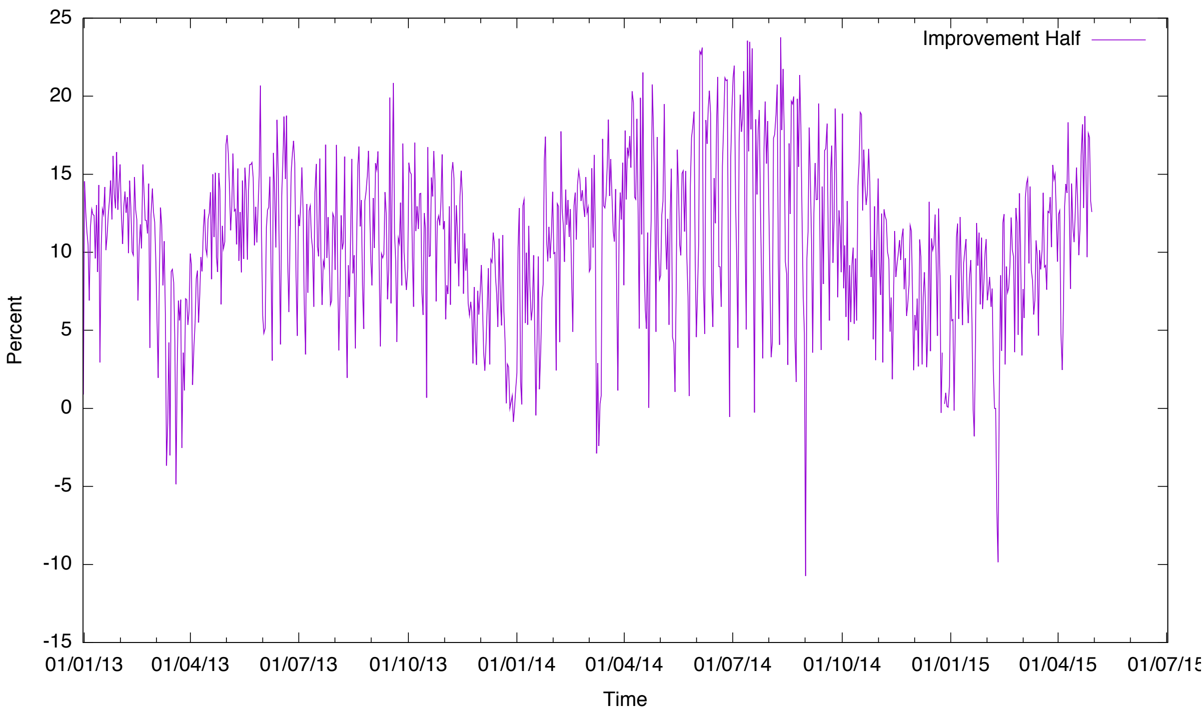

Figure 20: Improvement (in Percent) of HalfPower Model over Baseline

Table 11 summarizes the results. For the period from 1/1/2013 to 30/4/2015, we compare the cost of the schedules created either with perfect knowledge or using the forecasted values against the baseline operation.

Table 11: Daily Cost with Actual and Forecasted Values

Baseline	Pure	Full	Half	Computed Value
3470.70	2600.26	3414.57	3068.68	Daily Cost with Perfect Knowledge [€/day]
n/a	870.43	56.13	402.01	Savings over Baseline [€/day]
n/a	25.08	1.62	11.58	Percentage Savings [%]
n/a	2690.25	3450.17	3102.95	Daily Cost with Forecast Data [€/day]
n/a	780.26	20.34	367.57	Savings over Baseline [€/day]
n/a	22.48	0.59	10.59	Percentage Savings [%]
n/a	2.60	1.03	0.99	Loss due to Forecast [%]

With the hypothetical perfect knowledge, the **Pure** model achieves a cost reduction of 25%. This is reduced to 11.5% for the **HalfPower** model, where the additional constraints remove some of the opportunity for cost reduction. The **FullPower** model is only 1.6% better than the **Baseline**, indicating that the additional constraints do not allow enough changes from the **Baseline** schedule. If we consider the schedules built from the forecasted data, we see than the improvements are quite similar. The cost of using the forecast is 2.6% for the **Pure** model, and less than one percent for the **HalfPower** model. Even though the error in the forecast can be quite high, the resulting schedule is not much more expensive than the result obtained with perfect knowledge.

These results justify the approach taken in our system. The loss due to the forecast error is small compared to the savings possible by using optimization, at the same time the overhead of operational constraints imposed by the operator is significantly higher than the impact of the forecast error as well.

The results also show that the existing contract structure of an owner/operator running part of the complete system imposes many restrictions due to long-term contracts that do not take current energy prices into account.

5.2 Ucd HSG 6.1m: Heat Demand Load Balancing at Commerzbank Arena Frankfurt

We refer to [5] for more detailed discussion of the findings.

In CAMPUS21, Load Balancing has been performed in two control modes:

- L1: with a pre-defined heating option into the daily scheme and with B3 night heating and pre-heating
- L1mod: a more aggressive approach than L1.

The results for the **L1 experiments** in comparison with the night time heating B1-B3, PA1 and PA2 experiments can be comprehended with following aspects for increased performance and energy savings:

- Performance of load balancing experiments: UPR=0% (in experimental mode).

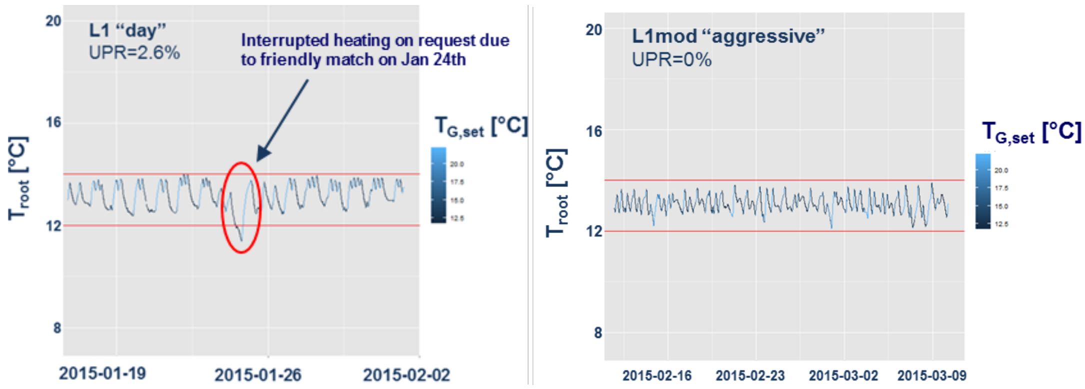

All night time experiments exhibited worse UPR when heating logic was not interrupted.

- Energy Efficiency and consumption savings:

For the February subset of the L1mod experiment, a reduction of ~62.5 % could be calculated for the median HDD normalized daily grass heating consumption compared to the baseline winter of 2013/2014. This resulted in median HDD normalized daily savings of 23.9% for the gas consumption reported for the entire arena as fuel for the heating system. At the same time, no underperformance was detected during the L1mod experiment.

For the entire experiment phase (from B1-PA2) – compared to last winter – the median HDD normalized daily grass heating consumption was reduced from 2797 kWh/DD to 1065 kWh/DD – a reduction of ~62%. At the same time, the median HDD normalized daily gas consumption of the entire arena was reduced from 9640 kWh/DD to 7120 kWh/DD – a reduction of ~26%.

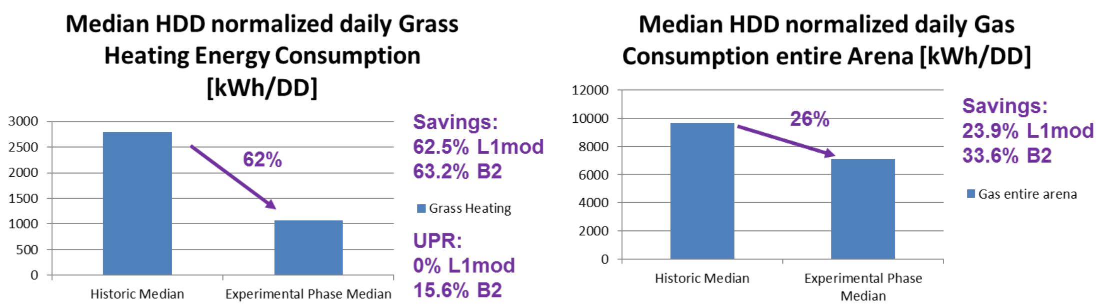

In summary, the CAMPUS21 experimental phase with grass heating results in the following findings:

- Automatic control of complex systems under severe resource constraints can enhance daily operation and best practices
- Basic control schemes achieve high energy savings and are easily explainable to operational staff, but suffer potentially from under performance
- Pre-Heating scenarios can address resource constraint situations and mitigate under performance risks – at the cost of higher energy use. Experiments B1 - B3 demonstrate this.
    - More advanced forecasting methods PA1 and PA2 are expected to improve the energetic behavior of the pre-heating approach.
- The daytime load balancing scheme achieves highly competitive energy savings, comparable with the best basic experiment schemes while avoiding under performance.

6. References

[1] CAMPUS-21, "Deliverable 2.4 Installation of CAMPUS21 System Platform ECOService", revision 2, November 2013.

[2] CAMPUS-21, "Deliverable 4.456 Integrated Prototype Release Documentation", revision 1, May 2015.

[3] CAMPUS-21, "Deliverable 6.3 Evaluated Load Balancing Tools", revision 3, October 2014.

[4] CAMPUS-21, "Deliverable 6.4 Validated and Adapted Load balancing Algorithms & Tools", revision 1, February 2015.

[5] CAMPUS-21, "Deliverable 6.4 Validated and Adapted Load balancing Algorithms & Tools", revision 2, May 2015.

7. APPENDIX

7.1 Ucd UCC 6.1: CHP Optimization on UCC Campus

The following table gives an updates list of the ResourceKraft data points that have been used for the analysis.

Table 12: Updated Data Points from ResourceKraft

DataPoint	Field	Measures	Unit
/ds:6173	ENERGY/RESOURCE USAGE	Food Science Steam Meter (ENERGY USAGE)	J
/ds:6171	ENERGY/RESOURCE USAGE	Boole Steam Meter (ENERGY USAGE)	J
/ds:1380	ELECTRICITY/RESOURE USAGE	Kane Incomer Total (ELECTRICITY USAGE)	kWh
/ds:1358	ELECTRICITY/RESOURE USAGE	Boole Q-2 Main Meter (ELECTRICITY USAGE)	kWh
/ds:1365	ELECTRICITY/RESOURE USAGE	Boole Q+4 Main Meter (ELECTRICITY USAGE)	kWh
/ds:1393	ELECTRICITY/RESOURE USAGE	Main Electricity (incl Med Bldg) (ELECTRICITY USAGE)	kWh
/ds:1394	ELECTRICITY/RESOURE USAGE	Main CHP Elec Meter (ELECTRICITY USAGE)	kWh
/ds:1386	ELECTRICITY/RESOURE USAGE	O Rahilly Building (ORB) Main Meter (ELECTRICITY USAGE)	kWh
/ds:1350	ELECTRICITY/RESOURE USAGE	Pagoda Main Meter (ELECTRICITY USAGE)	kWh
/ds:2260	ELECTRICITY/RESOURE USAGE	Site Electricity Total Main Campus (ELECTRICITY USAGE)	kWh
/ds:2262	ELECTRICITY/RESOURE USAGE	ESB Day Incomer (ELECTRICITY USAGE)	kWh
/ds:2231	GAS/RESOURCE USAGE	Main Campus Boiler House Gas (GAS USAGE)	m3
/ds:6170	WATER/RESOURCE USAGE	CHP Steam Meter (WATER USAGE)	l
/ds:6134	WATER/RESOURCE USAGE	CHP Hot Water (WATER USAGE)	l
/ds:4535	WATER/RESOURCE USAGE	Make up Feed Tank (WATER USAGE)	l
/ds:1370	ELECTRICITY/RESOURE USAGE	Main Meter Food (ELECTRICITY USAGE)	kWh
/ds:1377	ELECTRICITY/RESOURE USAGE	Pharmacy Electricity (ELECTRICITY USAGE)	kWh
/ds:1392	ELECTRICITY/RESOURE USAGE	Main Meter Bio (ELECTRICITY USAGE)	kWh
/ds:2425	ENERGY/RESOURCE USAGE	Main Restaurant Heat Meter (ELECTRICITY USAGE)	J
/ds:2426	ENERGY/RESOURCE USAGE	Windle Heat Meter (ELECTRICITY USAGE)	J
/ds:2428	ENERGY/RESOURCE USAGE	Electrical Engineering Heat Meter (ELECTRICITY USAGE)	J

/ds:2429	ENERGY/RESOURCE USAGE	Civil Engineering Heat Meter (ELECTRICITY USAGE)	J
/ds:2430	ENERGY/RESOURCE USAGE	Quadrangle Heat Meter (ELECTRICITY USAGE)	J
/ds:2427	ENERGY/RESOURCE USAGE	Kane Heat Meter (ELECTRICITY USAGE)	J

The next table shows the specified values ranges for the different data points from the ResourceKraft system. The values have been defined manually by observing the values of multiple years.

Table 13: Acceptable Value Ranges for ResourceKraft Sensors

Data Point	Measures	Min	Max
/ds:6173	Food Science Steam Meter (ENERGY USAGE)	0 J	1e9 J
/ds:6171	Boole Steam Meter (ENERGY USAGE)	0 J	4e8 J
/ds:1380	Kane Incomer Total (ELECTRICITY USAGE)	40 kWh	200 kWh
/ds:1358	Boole Q-2 Main Meter (ELECTRICITY USAGE)	5 kWh	120 kWh
/ds:1365	Boole Q+4 Main Meter (ELECTRICITY USAGE)	3 kWh	80 kWh
/ds:1393	Main Electricity (incl Med Bldg) (ELECTRICITY USAGE)	0.5 kWh	30 kWh
/ds:1394	Main CHP Elec Meter (ELECTRICITY USAGE)	40 kWh	200 kWh
/ds:1386	O Rahilly Building (ORB) Main Meter (ELECTRICITY USAGE)	2 kWh	100 kWh
/ds:1350	Pagoda Main Meter (ELECTRICITY USAGE)	10 kWh	150 kWh
/ds:2260	Site Electricity Total Main Campus (ELECTRICITY USAGE)	kWh	kWh
/ds:2262	ESB Day Incomer (ELECTRICITY USAGE)	kWh	kWh
/ds:2231	Main Campus Boiler House Gas (GAS USAGE)	0 $\mathbf{m}^3$	250 $\mathbf{m}^3$
/ds:6170	CHP Steam Meter (WATER USAGE)	1	1
/ds:6134	CHP Hot Water (WATER USAGE)	1	1
/ds:4535	Make up Feed Tank (WATER USAGE)	1	1
/ds:1370	Main Meter Food (ELECTRICITY USAGE)	10 kWh	150 kWh
/ds:1377	Pharmacy Electricity (ELECTRICITY USAGE)	10 kWh	150 kWh
/ds:1392	Main Meter Bio (ELECTRICITY USAGE)	25 kWh	100 kWh
/ds:2425	Main Restaurant Heat Meter (ELECTRICITY USAGE)	0 kWh	200

			kWh
/ds:2426	Windle Heat Meter (ELECTRICITY USAGE)	0 kWh	200 kWh
/ds:2428	Electrical Engineering Heat Meter (ELECTRICITY USAGE)	0 kWh	200 kWh
/ds:2429	Civil Engineering Heat Meter (ELECTRICITY USAGE)	0 kWh	200 kWh
/ds:2430	Quadrangle Heat Meter (ELECTRICITY USAGE)	0 kWh	350 kWh
/ds:2427	Kane Heat Meter (ELECTRICITY USAGE)	0 kWh	600 kWh