Final report

1.1 Project details

Project title	Forecasting and Aggregation of SME Heat Pumps
Project identification (program abbrev. and file)	64016-0053
Name of the programme which has funded the project	EUDP
Project managing company/institution (name and address)	Energy Cool Navervej 10, 7000 Fredericia
Project partners	EWII, DTU
CVR (central business register)	33067135
Date for submission	29/3-2019

1.2 Short description of project objective and results

The purpose of CASHPump project is to develop and demonstrate a solution that can enable use of HP installed in small and medium-sized enterprises (SMEs) to deliver balancing power to the power market. The project will focus on development and demonstration of:

1. Forecasting models of heat needs based on few sensor inputs and calculation of activation potential for aggregator
2. Integration in aggregator software, which pools HP together with other flexible assets and enables control of activations
3. System configuration with control hardware, sensors and a buffer, which is commercially scalable both for new installations and as a retrofit to existing installations
4. Drafting and development of standard contracts for aggregator services for future commercialization

A solution for item 1-3 has been demonstrated with good results.

Formålet med CASHPump projektet er at udvikle og demonstrere en løsning som muliggør leverance af systemydelser til elmarkedet fra varmepumper installeret i små og mellemstore virksomheder. Projektet fokuserer på udvikling og demonstration af:

1. Model til forecast af varmebehov baseret på få sensorinput og beregning af potentialet for aktivering for en aggregator
2. Integration i aggregator software, som kan pulje VP-aktivering med andre fleksible enheder og muliggør aktivering
3. Systemkonfiguration med styrings hardware, sensorer og buffertank, som er kommercielt skalerbart både til nye installationer og retrofit på eksisterende
4. Skabelon for standard aftalen imellem aggregator og VP ejer til fremtidig kommercialisering

En løsning for punkt 1 til 3 blev demonstreret med gode resultater.

1.3 Executive summary

The main objective for the project was to develop and demonstrate a commercial scalable solution for allowing HPs installed in SMEs to deliver balancing power to the power market. Such a solution would also allow for optimizing the operation of the HP with respect to other cases such a varying spot prices and tariffs. Delivering balancing power to the power mar-

kets requires the ability to control the HP and the ability to forecast when and how much power can be activated – in this project turning off (delaying) expected consumption and hence reducing grid load is considered an activation. The key to forecasting how much power can be activated is a forecast of the expected consumption for a given period and a forecast of the expected indoor temperature development when the HP is turned off.
In the project it has been demonstrated that the forecasting can be performed with high accuracy using only one temperature sensor and power meter connected only to the HP though adding a heat meter will improve accuracy. Two types of HPs have been tested in the project and both required a serial modbus connection for controlling the settings of the HP. The setup used for the demonstrations in the project is simple. It consists of an interface box connecting to the heatpump, the power meter and the heat meter using modbus. A modem is used for connecting to Energy Cloud. We find that the obtained results provide evidence that with a simple and scalable setup it is possible to aggregate heat pumps in a manner suitable for providing regulating power.
In the project the interface box is based on a Schneider PLC but any hardware platform capable of communicating via modbus and modem can be used. Indoor temperature is measured using an IC-meter device directly connected via its own modem – this is very flexible and easy to install but a more cost-effective solution may be required for a large scale roll out.
During the project it has become evident, that challenge exists in relation to collecting sufficient training data for the forecast algorithms while also providing a high number of activations. In general, how to select training data and validate that a set of training data provides sufficient accurate forecasting is an area for future research and testing.

1.4 Project objectives

The objective for the project was to develop and demonstrate a setup that enables medium sized HPs (10 – 50 kW electrical power) to participate in the balancing power marked in a commercially viable way. It is assumed that the individual HPs (customers) will access the balancing power market through an aggregator but no assumptions are made about the aggregator's role in relation to other actors in the power market (e.i. the aggregator could be a balance responsible party, power retailer or independent third party). Participation in the balancing power market requires accurate forecasting of the load that will be delivered, measurement of the actual consumption. For the solution to be commercially viable, the hardware setup in addition to the HP must be cheap both with regards to hardware cost and installation cost, and it must be robust with respect to maintenance.
Forecasting the available power requires a forecast of how long the HP can be turned of starting from a specified point in time, until the indoor temperature reaches a lower threshold where the HP must be turned on again. It also requires a forecast of the expected electrical power consumption from the HP during the period – and in more general terms a forecast of the electrical power consumption from the HP that can be used as a baseline. Such forecasting has been demonstrated before using a range of (sensor) input that makes the solution unsuitable for commercial roll-out. Hence performing a similar forecasting with fewer inputs and still sufficiently accurate is the key objective.
As the project focuses on reaching a commercially viable solution factors such as installation cost and maintenance has been considered.
To demonstrate the solution two test sites (buildings) were identified where heat-meter and indoor temperature sensors were installed. These were used for collecting training data for forecasting of indoor temperature development. All testing during the first heating season in the project were performed by manually switching heating off in the buildings and manually turning it on again – making it impractical to perform a high number of tests.
Installation of heat pumps in the buildings were not a part of the project but supervised by the project to ensure that the installed heat pumps had the required interfaces for operation. Installation in the Strib library test site was delayed and operational for the second heating season in the project. Therefore, the project period was extended to allow for a two heating seasons as planned. In December 2018 (during the additional heating season) Middelfart municipality decided to renovate Strib library cutting the test period short. However, at this point a series of test had been performed during November and December providing sufficient test data. At the Gudsøgaard test site, a heat pump was not installed (for reasons out-

side the scope of the project), so this test site has only been used for testing the forecasting of indoor temperature.

One planned outcome, a template for commercial agreement between aggregator and heat pump owner, has not been produced due to changes in organization and focus of one of the project partners (see section 1.5.7).

The key result of the project is a simple hardware setup that has been demonstrated to enable forecasting of activation potential for heat pumps, and the solution has been integrated into the Energy Cloud aggregator software. This provides a setup that is commercially viable for large scale roll-out, though some cost reductions should be implemented before such roll-out as discussed in section 1.7. Also, some challenges have been identified in relation to combining high numbers of activations (which is desired from a commercial perspective) with training data for the forecasting algorithms.

1.5 Project results and dissemination of results

1.5.1 Screening and selecting buildings

A key focus for the project is the practical aspects of using flexibility from HPs. One aspect of this is the building and the installed heating system in the building – whether it is suited for installation of a HP and whether it is possible to monitor the building in a reliable way using only one or a few temperature sensors.

Some buildings were visited but found too complex for the project because measuring indoor temperature in one or a few rooms would not be representative and because controlling the heating from one place (only using the heat pump – not individual thermostats) would not be possible. The buildings also had several heating sources, something which is often the case in buildings where new parts have been added several times. Only controlling the HP would lead to "imbalance" in the indoor temperature in the building resulting in much wasted energy and poor indoor climate.

During the screening it became clear that aiming for a simple and low-cost system like the one envisioned in the project is only relevant in some types of buildings. The approach and functionality developed in this project can be extended to a more complex setup, but much cost and energy would go into developing interfaces to control the different heating systems and such a task would most likely only be relevant during an upgrade or renovation of the total heating system in such buildings.

After screening several types of building we ended up choosing 2 test sites. The first was Strib Library which are owned by the municipality (Middelfart). One of the main reasons, was that the building only where used by people for short time, so if the temperature varied because of the test we have to make in the project only a few people would be inconvenient compared to a nursery and a School with a lot of people. The Second place where at 2 rental properties at Gudsøgaard in Fredericia.

We made an agreement with both Strib library and the people in the rental properties that the indoor temperature would be between 19-24 degrees during the project period.

1.5.2 Installation setup

The objective of the project was to define a setup that is both simple and sufficiently cheap. Focus was on minimizing the number of sensors and keeping installation cost low. Therefore, the suggested sensor setup consists of one power meter, one heat meter and one indoor temperature sensor. To easy installation and flexibility during the test, IC-meter was selected as indoor temperature sensor. It has a build in modem and communicates with IC-meter cloud where data it made available through a REST-API for integration in other systems such as Energy Cloud. The two other sensors are typically installed at the same location as the HP and therefore simple to connect to the same communication box. For this purpose, Energy Cool's Remote Service Hub (RSH) communication box was used. The RSH has a PLC with modbus connection and a modem for communication with Energy Cloud. An modbus interface was added to the HP, so all local communication was via modbus.

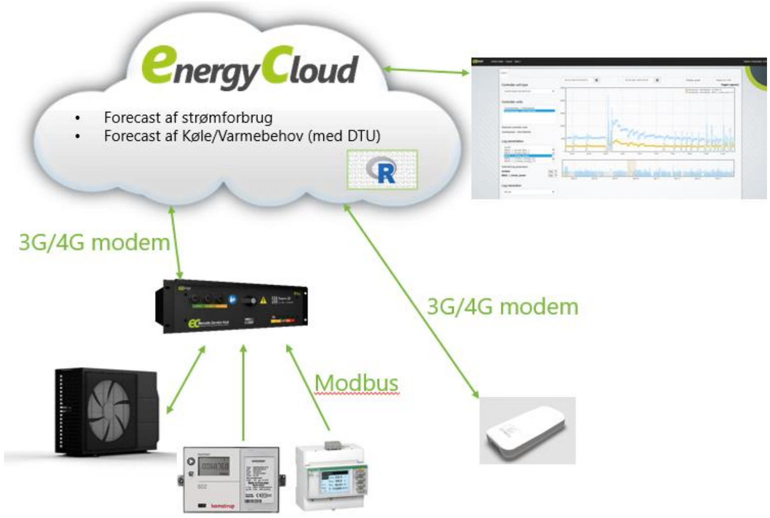
EnergyCloud
Forecast af strømforbrug
Forecast af Køle/Varmebehov (med DTU)
3G/4G modem
3G/4G modem
Modbus

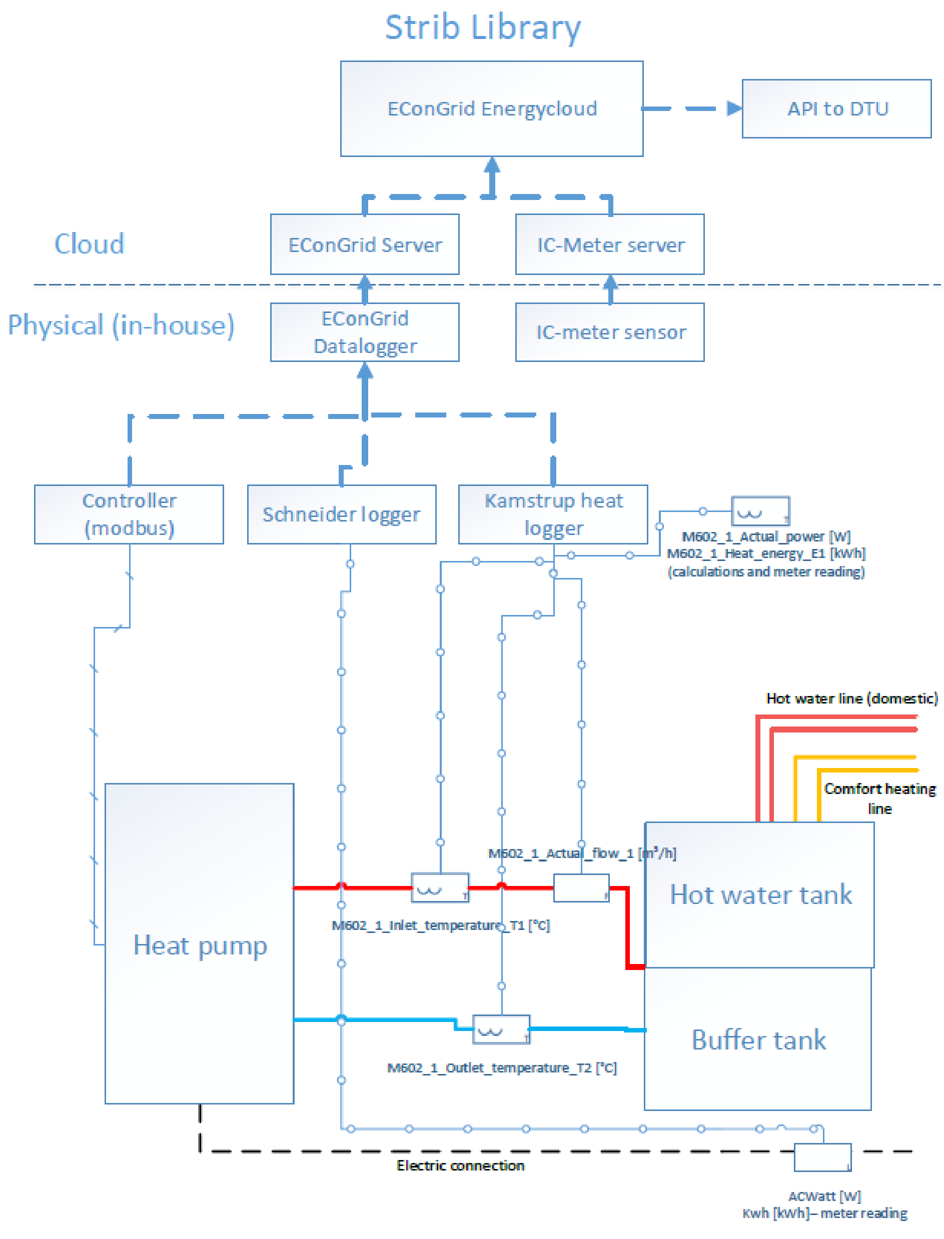
Strib Library
EConGrid Energycloud
API to DTU
Cloud
EConGrid Server
IC-Meter server
Physical (in-house)
EConGrid Datalogger
IC-meter sensor
Controller (modbus)
Schneider logger
Kamstrup heat logger
M602_1_Actual_power [W]
M602_1_Heat_energy_E1 [kWh]
(calculations and meter reading)
Hot water line (domestic)
Comfort heating line
M602_1_Actual_flow_1 [m³/h]
Hot water tank
Heat pump
M602_1_Inlet_temperature_T1 [°C]
Buffer tank
M602_1_Outlet_temperature_T2 [°C]
Electric connection
ACWatt [W]
Kwh [kWh]– meter reading

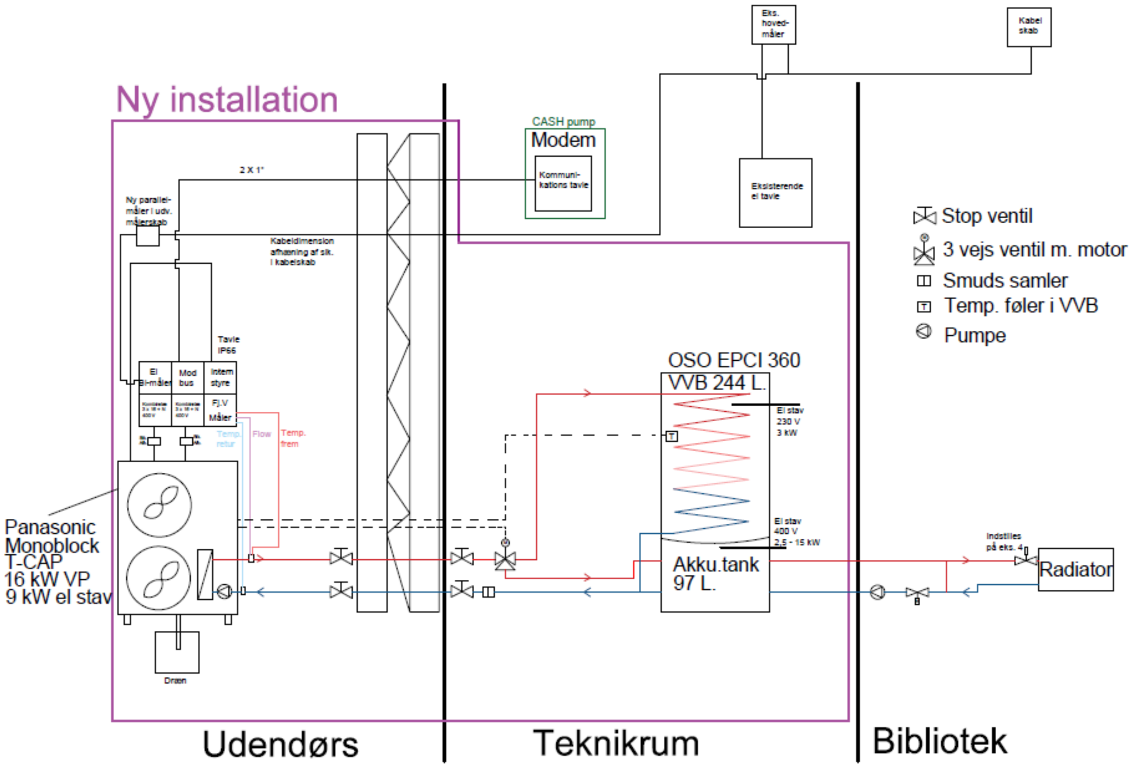

At both test sites heat meters and temperature loggers where installed as fast as possible, so the project could start having data from the existing heating systems. At Strib there was a gasboiler installed and at Gudsøgaard it was 2 oil boilers. The data gave us knowhow about how much energy the buildings need compared to the outdoor temperature. Before the heat pump was installed, we tried to test the buildings heating capacity by manual short down of the gas and oilboiler at site. That gave us knowledge about the buildings reaction when we turned off the heating for the building which we used to train the algorithm DTU have developed during the project.

None of the 2 testcases had a heat pump installed at the beginning of the project. EWII Energy offered the municipality at Middelfart a new concept called "Heat pump at subscription" where EWII as utility company install, own and operate a heat pump at the customer. The customer payed for the heating the produced by the heat pump similar to district heating. The heat pump at Strib where installed the 5 dec 2017, and after that date it was possible to turn the heatpump on/off remote via the interface box controlled by Energy Cloud.

The choice of Panasonic T-Cap monoblock heat pump to Strib library, where because it has a modbus communications opportunity and it was frequency regulated between 3-16kW. Furthermore, the T-Cap has the advantaged compered to many other heat pumps that it can deliver the capacity also when it is -12° outside, where the most others heat pumps loose capacity. This unique technology gave us the opportunity to make a 1:1 replacement with the gasboiler which had a capacity at 25kW at -12° outdoor temperature. That capacity was achieved using the Panasonic T-Cap 16kW Heat pump + 9 kW electricity boiler at -12°C.

At Gudsøgaard we unfortunately not where able to change the heating system from oilboilers to heat pumps because of future plans for the site, which made it difficult to change to the right solutions during the project period.

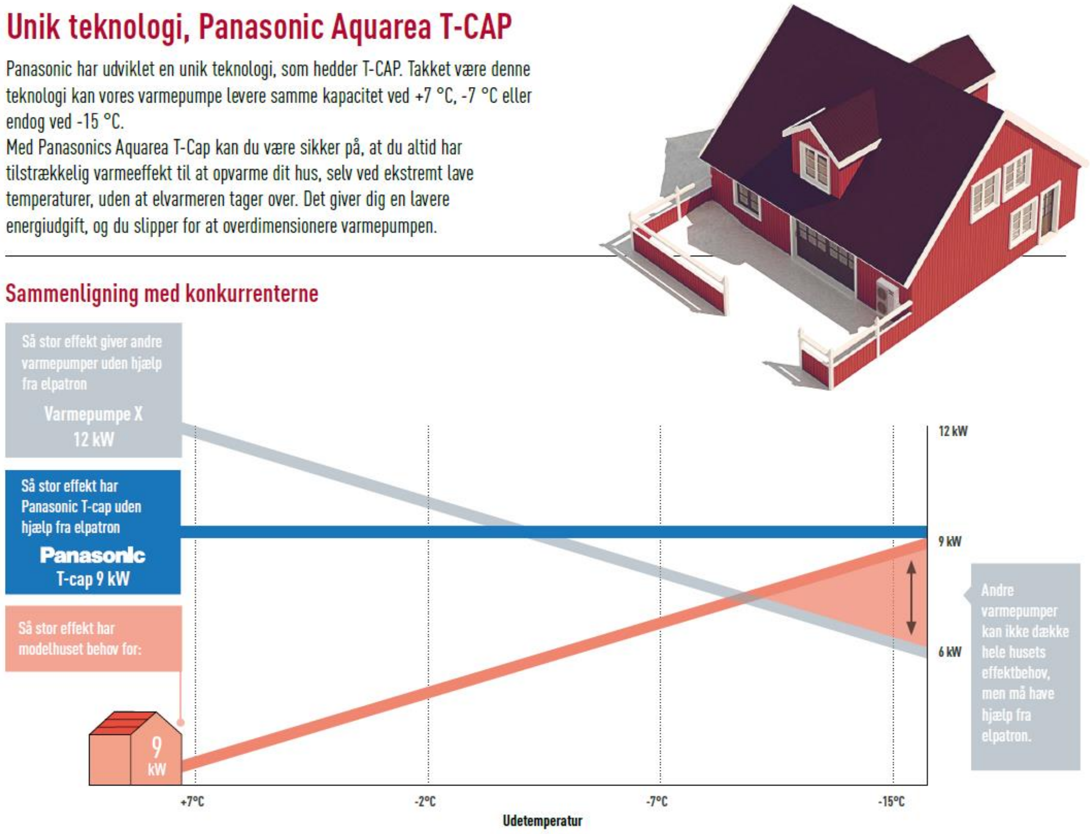

1.5.3 *Forecasting activation potential*

Forecasting the activation potential for a HP in a building requires two independent steps:

- Step one is a forecast of the electricity consumed by the HP under normal operation. This is refered to as the baseline consumption. Since activations here are considered as a shutdown of the HP, the activation potential for a HP is the forecasted consumption for a given period of time under regular operating conditions, i.e. usually with a default thermostatic control.
- Step two is to calculate the (maximum) period of time the HP can be shutdown without violating are required indoor temperature lower level, hence a forecast of the indoor temperature when the HP is shutdown is needed.

In this section the forecasting techniques applied and an analysis of the forecasting performance for the test case is presented.

First the results of forecasting of electrical power (or load) for the heat pump is presented, and thereafter for the indoor temperature forecasting.

1.5.4 *Electrical load forecasting*

In this section the results of forecasting the electrical load of the heat pump using the R package "Online Forecast" (see attached document "Description of online forecast package") are presented. The model is a based on a recursive least squares estimation technique, as the one applied in (ref load forecast paper). First the data is presented and then the results.

Load forecast data

The data used for load forecasting consists hourly values of:

- Numerical Weather Predictions (NWPs) of external air temperature and solar radiation. Each hour the most recent hourly NWPs 10 hours ahead are used. The weather forecasts are from openweathermap and have 3 hour values, which are interpolated to 1 hour values.

- Observed load, i.e. the average electrical power consumed by the heat pump.

The data set over the entire period is plotted as hourly values:

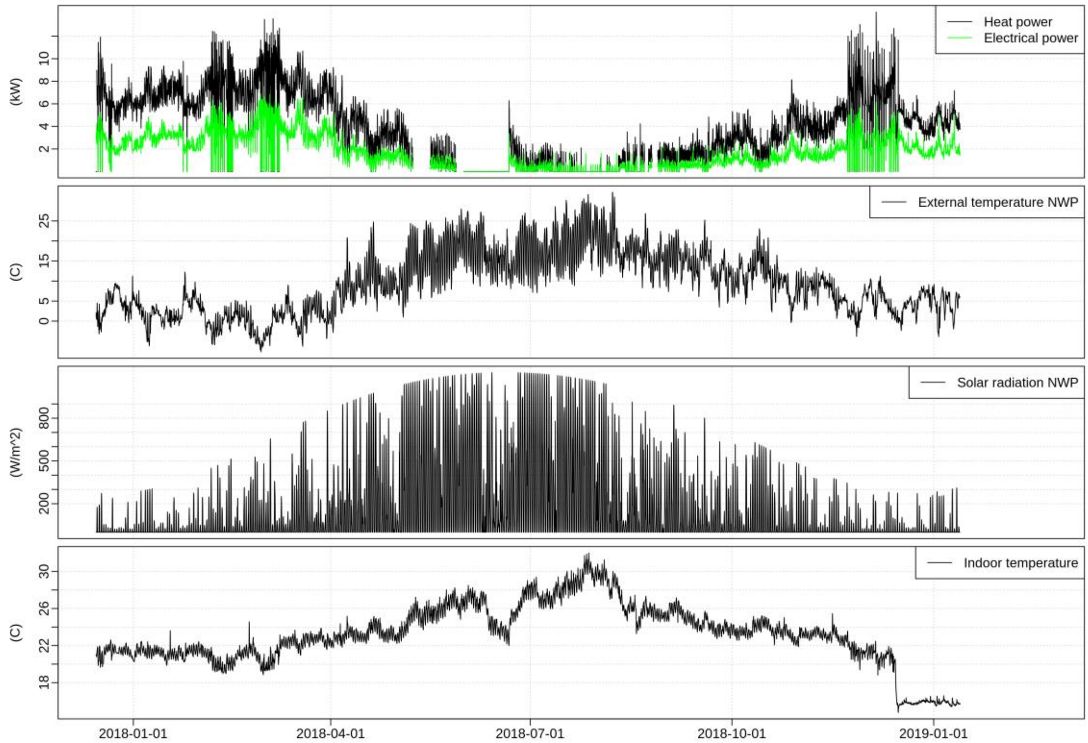

The activation periods in March and December 2018 can easily be seen. Further, it can be seen that the indoor temperature is around 22-23 C in the heating period, with some lower values during the activation periods.

The data from the period from start November to mid December, which will be used more intensively in the analysis, is shown below:

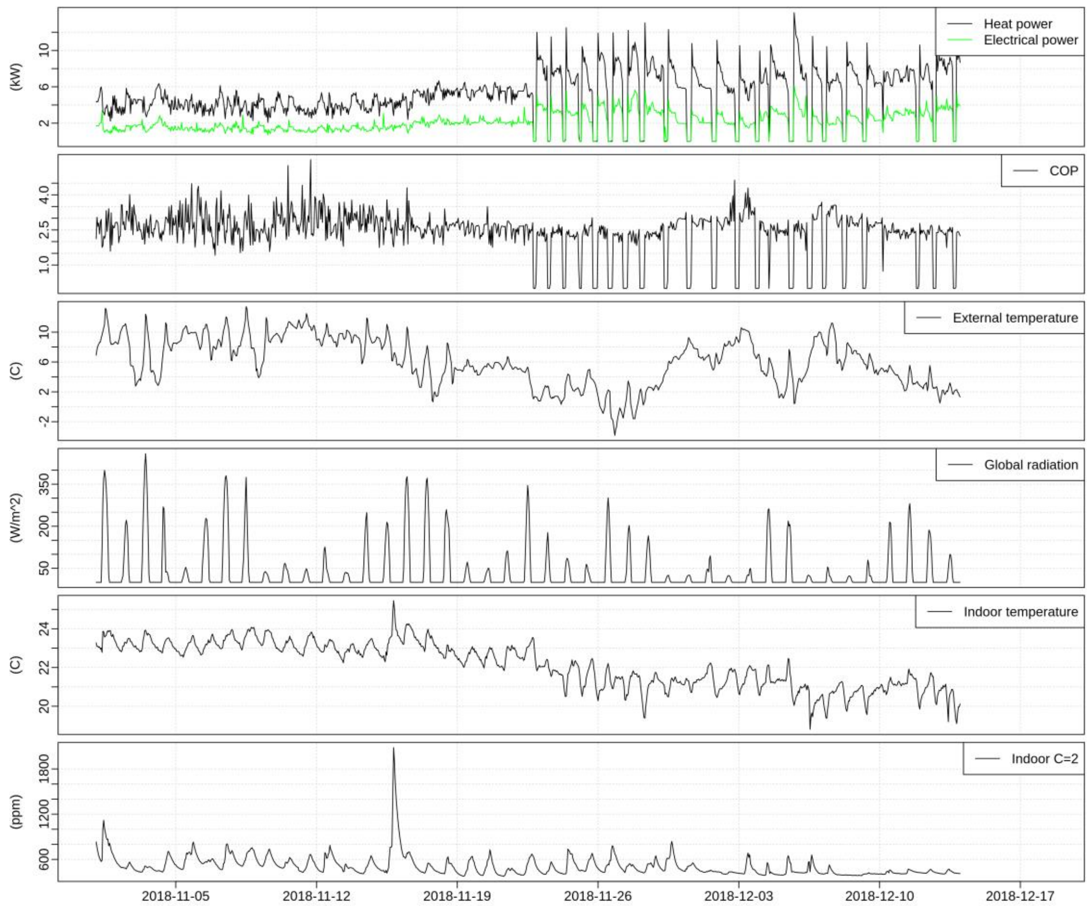

The following is noticed about this period:

- First part is of with no activations, hence just the default thermostatic control of the heating in the building.
- Second part is with activations, where the heat pump is shut down in shorter periods once every day at different times of day.
- The external temperature starts at levels around 5-10 degrees and drops to a colder period down around -2 to 2.
- Solar radiation varies, such that there are sunny days (mostly in the beginning) and cloudy days.
- The indoor temperature levels drop in the activation period, this can both be due to colder weather, as well as the shut-downs.
- The CO2 levels indicate activity in the building most days, with one very high peak, most probable this an event where many people were in the building at the same time.

In order not model the load during regular operation, i.e. how it would be in with no shut-down, the periods with activations are removed, as well as the period from 1st May to 1st of October. This leads to the load from the periods shown below:

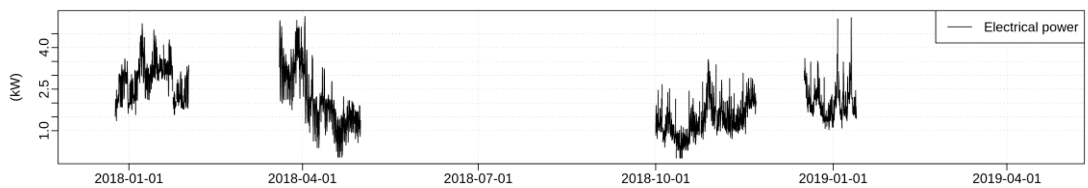

A training period is needed before each forecast test period and in order to always calculate the results on exactly the same data. 55 days which have a sufficiently long previous period train period, have been selected, such that the forecasts evaluation scores (RMSE and bias) are only calculated on the following data

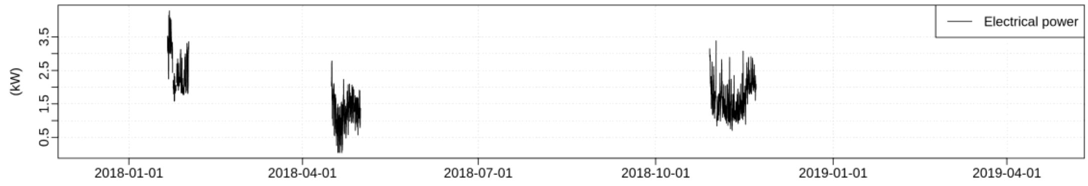

Note that all evaluation is generated by training only on the past period and calculating the evaluation scores on the selected day.

Load forecasting

Forecasts of the load calculated up to 10 hours ahead is shown for a short period in the plot below. The forecasts are updated every hour, but only the forecasts calculated every 15 hours is included in the plot for clarity. The grey vertical lines marks the issue time (the time the forecast is calculated) and the red lines marks the forecasts ahead in time from each of these points:

In order to get more insights in how well the forecasts are following the observed load, then all the one and ten-step forecasts of are extracted (since the forecasts are calculated each hour, then a series of hourly values of a given horisont can easily be extracted from the results):

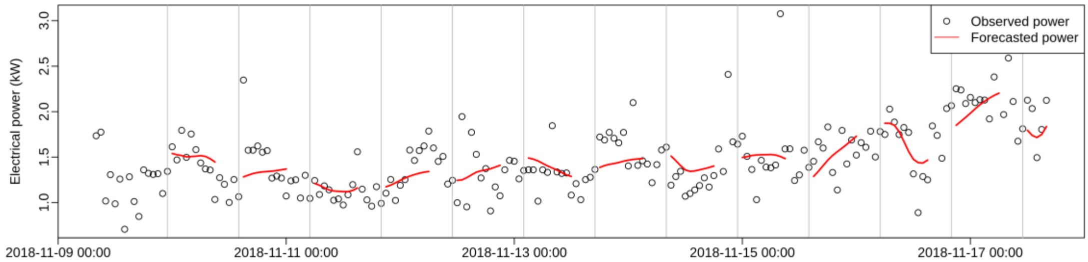

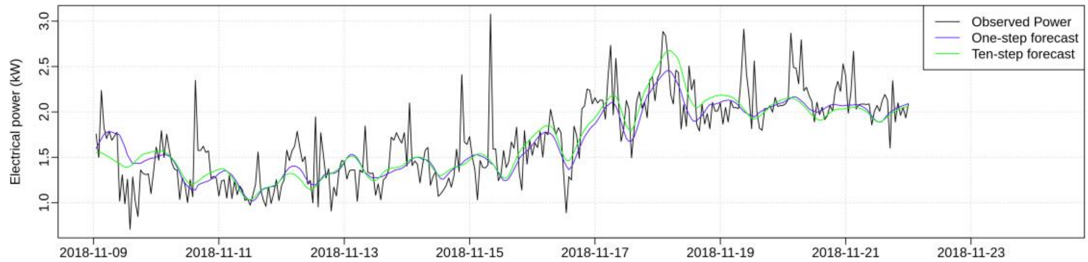

It is found that the forecasts follow very well the mean observed load. Of course, variation occur and in particular some features seems to be not modelled by the forecast model:

- Some peaks lasting only a single hour are not modelled, they occur randomly and unpredictable.
- Some periods with 3-4 hours of consistently higher load are found, e.g. at 2018-10-11 and 2018-13-11. A further investigation of these, plotting 5 min. values in these periods and the input series, did not reveal any clear patterns or relationships, neither in time (some of these are at night time and others at day time). It is thus left as some unknown effect, which could be further investigated, maybe it is a particular effect on this system.

Here we present the main results from this analysis, which is an analysis of the forecasting accuracy as a function of:

- the length of the training period (“trainlen” in days).

- the length of a gab between the training period (the period on which the model is trained) ("gablen"} in days).

The periods are illustrated using the figure below:

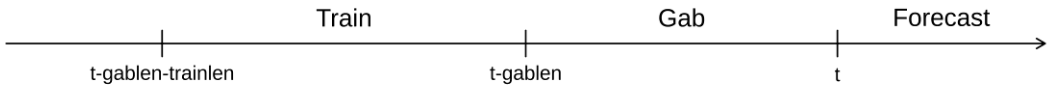

where "t" is the issue time of the forecasts (i.e. the time at which the forecast ahead in time is calculated).

In the plot below the two following forecasting evaluation criterias (scores) are plotted:

- The Root Mean Square Error (RMSE) of the forecast error averaged over all horizons

$$\frac{1}{10}\sum_{k=1}^{10}\sqrt{\frac{1}{n}\sum_{t=1}^{n}\varepsilon_{t+k|t}^2}$$

where the k-step forecast error is $\varepsilon_{t+k|t} = P_{t+k|t} - \widehat{P}_{t+k|t}$, where P is electrical power, k is the forecast horizon, t is the time index which counts all the n hours in the period and k is the forecast horizon in hours. The forecasts up of 10 hours ahead is included. And in in total $55 \cdot 24 = 1320$ hours are included in the period.

- The bias of the forecast error averaged over all horizons

$$\frac{1}{10}\sum_{k=1}^{10}\frac{1}{n}\sum_{i=1}^{n}\varepsilon_{t+k|t}$$

The x-axis is the "gablen" in days, see the illustration above, and the different colors marks the different training period lengths:

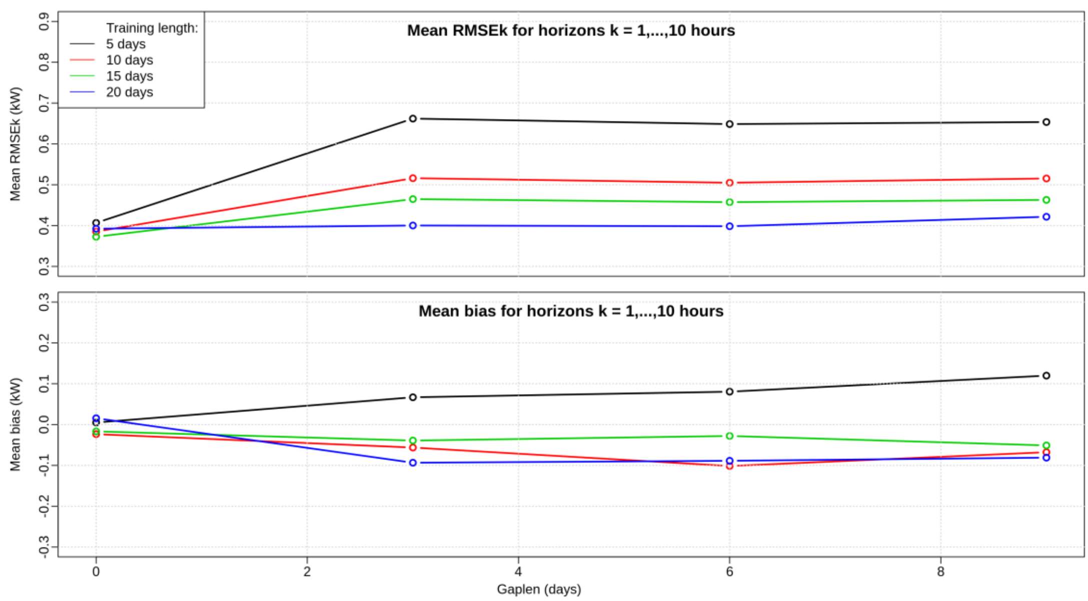

Considering the forecast accuracy in measured by RMSE it is found that the developed forecasts models will require some training period before the forecast accuracy converge. It it found that 5 days is too short a period, although the performance with no gap is more or less independent of the training length, however with a gab 20 days are needed to keep a stable forecasting accuracy.

Forecast accuracy for aggregated loads

A few considerations about the forecast accuracy for upscaling by pooling multiple similar

systems has been carried out. Roughly the standard deviation of a sum of n_{sys} independent random variables scales by the square root of the number of systems, i.e.

$$RMSE_{\mathrm{sum}} \approx \sqrt{\mathrm{Var}\left(\sum_{i=1}^{n_{sys}} \varepsilon_{i,t+k|t}\right)} = \sqrt{\sum_{i=1}^{n_{sys}} \mathrm{Var}\left(\varepsilon_{i,t+k|t}\right)} = \sqrt{n_{sys}\mathrm{Var}\left(\varepsilon_{t+k|t}\right)} = \sqrt{n_{sys}}RMSE$$

The average power in the forecast test periods is around 2 kW, hence a rough calculation for

$n_{sys} = 500$, which would realize 1 MW of flexibility, of the RMSE of the forecast of the aggregated electrical load will be

$$RMSE_{\mathrm{sum}} = 22.4 \cdot 0.4 = 9 \mathrm{~kW}$$

hence a 95% confidence interval would be roughly $\pm$ 20 kW from the 1000 kW. This is result is only valid if the forecast errors are independent, which they will not be, since:

- The forecasts rely of weather forecasts, which indeed are not perfect, hence their errors will probagate throught the forecasting model
- Some events leading to forecast errors, e.g. Christmas evening, will cause the forecast errors to be correlated.

However, it is absolutely certain that the forecast error will decrease drastically when many systems are aggregated together, and it will be possible to accurately model this uncertainty, also as a dependence of e.g. the weather forecasts and calendar events. This should be studied when data from many systems are available.

1.5.5 Indoor temperature forecasting during activation events

As explained in the beginning of the section, it is needed to forecast the indoor temperature during activation events in order to estimate when the indoor temperature fall below a user defined comfort temperature level.

The forecasting of indoor temperature during activation events were the first forecast developed because this did not require the heat pumps to be installed. Different methods were tested at both test sites. The results shown below are the final forecasts validated at Strib library during the 2018 heating season.

Data for indoor temperature forecast

The data used for the study consists of:

- Observations every 5 minute of:
  - Indoor temperature
  - Electrical and heating power
  - Indoor CO2
  - Indoor sound level
- Weather forecasts linearly interpolated from 3 hour averaged to 5 minute values of:
  - External temperature

- Global radiation
- Wind speed

The used data is plotted below. It is the period of activations between 2018-11-22 and 2018-12-14:

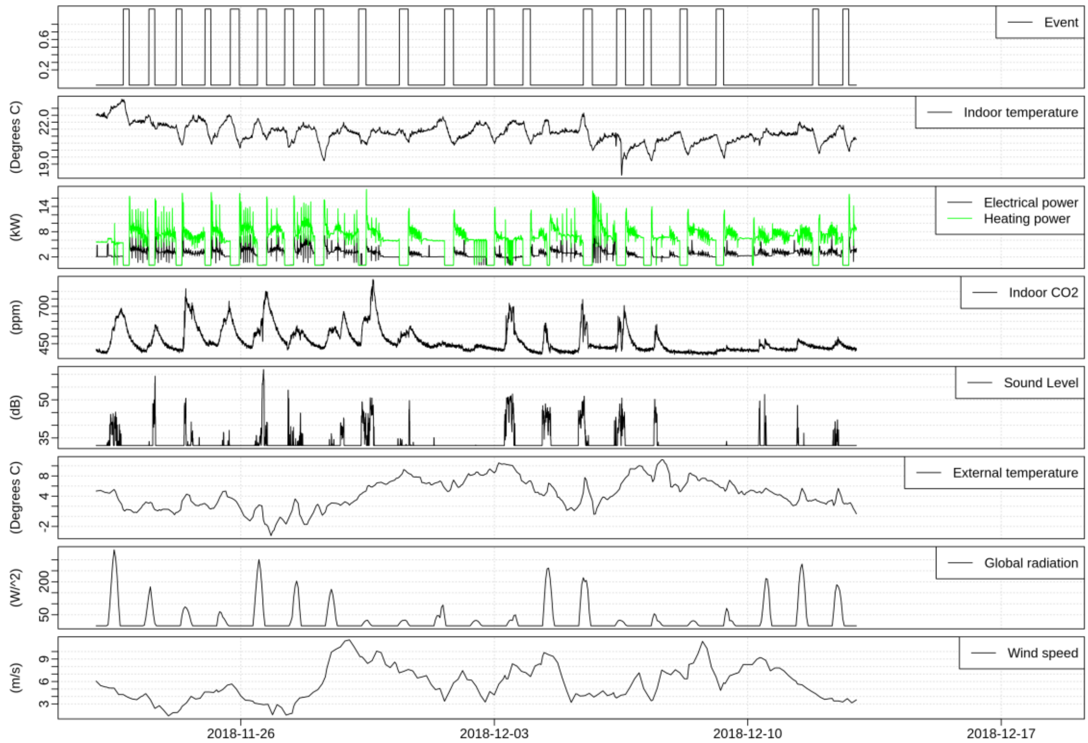

We find the following interesting from this plot:

- Clearly the indoor temperature decrease when the heat pump is shut down.
- The heat pump COP heating power and electrical power follow each other (COP around 2.5).
- From the CO2 and sound level it is seen that there are two weekends with little indoor activity (1. to 2. and 8. to 9. Dec.) and also after the last weekend.
- All climate variables vary relatively much.

We find interest in one particular event on the 6th Dec., where the indoor temperature drops very fast. A plot of a shorter period around that day is:

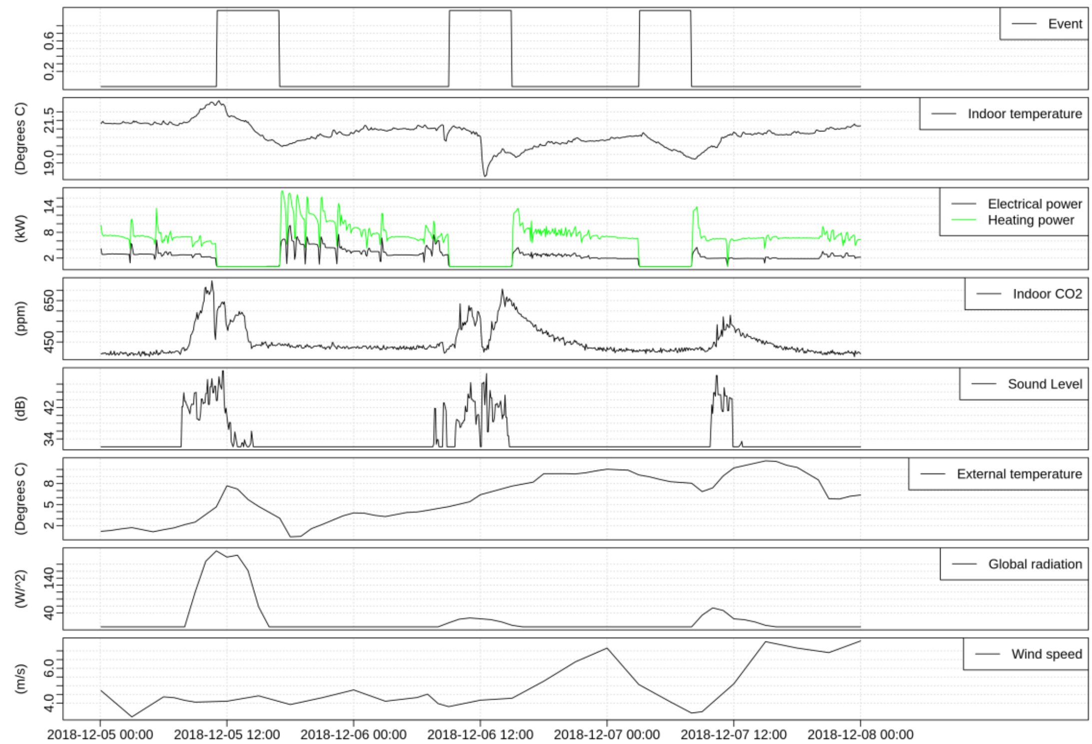

From which we find that a ventilation must have occurred, probably from window opening, since both the indoor temperature and the CO_2 level drops very rapidly. From the experience gained when applying the models it is clear that this particular event does influence the results (not the forecasts themselves, but on the performance measures) it is chosen to remove it from the data set. In further studies more detailed analysis of the impact of such events should be carried out.

Indoor temperature forecast examples

The first 7 events are reserved for initial training. The forecasts calculated using the model for the remaining events are plotted below:

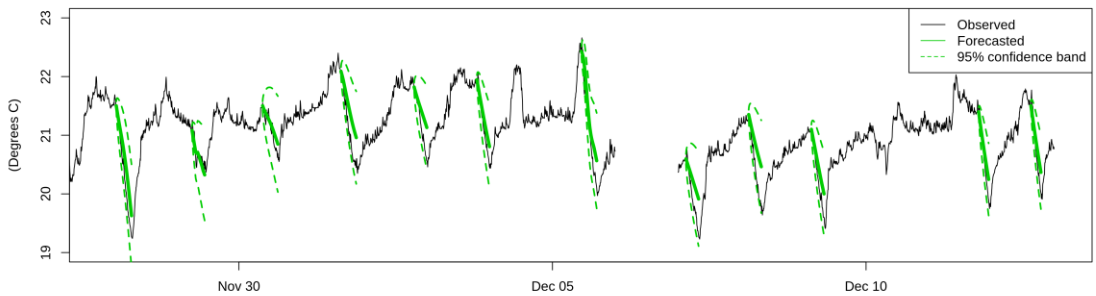

In addition to the green line marking the point forecast, which is an estimate of the mean value of the indoor temperature (i.e. a single point value for each step ahead in time), the dashed lines in the figure marks the estimated 95% confidence interval ahead in time.

To get a more detailed insight into the forecasts two events are plotted:

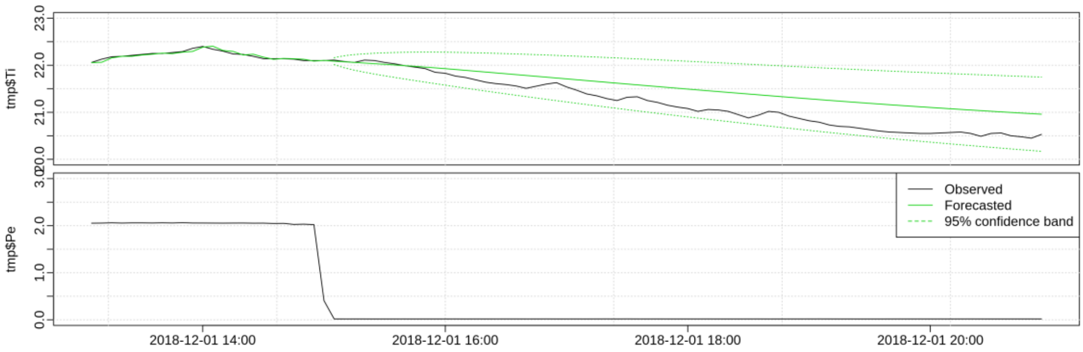

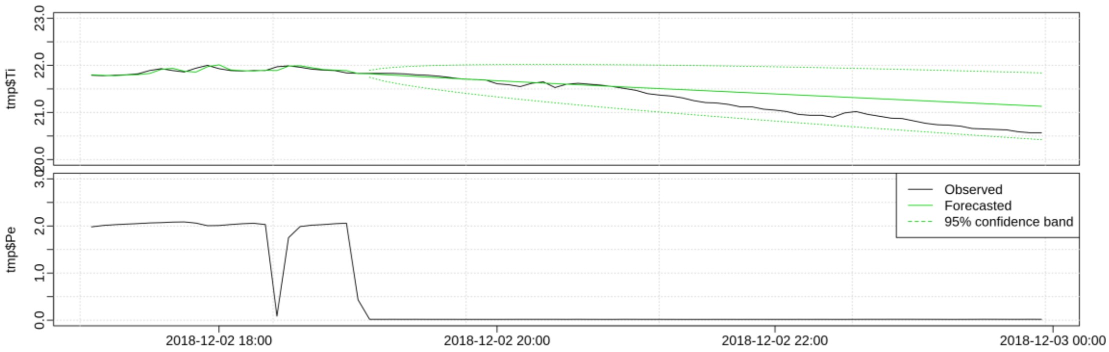

In the first it can be seen how the forecast follow quite well the decreasing indoor temperature, after 6 hours the realized error is aound 0.3 degrees, whereas in the second event the indoor temperature does decrease more rapidly than the forecast and after 5 hours the difference is around 0.8 degrees. Further, the confidence band in the second event didn't catch the observed. Generally, these plots reveal, what will clear from the analysis below, that the forecasts are biased, such that the observed temperature, in average, drop faster than the model predicts.

Indoor temperature forecast model

The applied indoor temperature forecas model is a first order ARMAX model. As an example, the model with power and external temperature as input can be written as

$$T_{i,t+1} = H_p(B)\,P_t + H_e(B)\,T_{e,t} + H_\varepsilon(B)\,\varepsilon_{t+1}$$

where t is the time index (i.e. one step in t is 5 min.), P_t is the power (either heat or electrical) and $T_{e,t}$ is the external temperature weather forecast. B is the back-shift operator.

The one-step ahead error ε_{t+1} for $t = 1, \ldots, n-1$ is assumed to be i.i.d.

The transfer functions in the model above are

$$H_p(B) = \frac{\omega_p}{1+\phi B}\ ,\ H_e(B) = \frac{\omega_e}{1+\phi B}\ ,\ H_\varepsilon(B) = \frac{\omega_\varepsilon}{1+\phi B}$$

hence this model has 4 parameters $(\omega_p, \omega_p, \omega_\varepsilon, \phi)$, which are estimated using the R package MARIMA.

The model can be extended to include more of the available inputs (observed values and weather forecasts) by adding new term with similar structure as the power and external temperature in the model above.

In order to select the best model the following two properties of the model were varied:

- Model inputs. Inputs were included increasing the model complexity and number of parameters.
- Training period length. The parameters were fitted on a training period with which include the period of a given number of previous events.

The forecast performance measured by average RMSE over all events and horizons (up to 47 steps ahead, which is nearly 4 hours ahead) is plotted below as a function of the training period length in number of previous events included. Further, a line for each of the different models of increasing complexity is included. These results are with the *heating power* as input power P_t. In the legend the included inputs for the model are indicated.

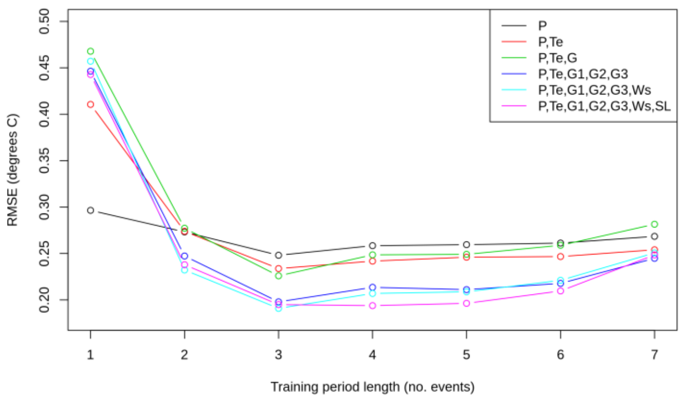

It is apparent that the forecast accuracy is improved until the training period length is increased to 3 (i.e. since there were appr. one activation event per day, this amounts to training periods of length 3 days, which include 3 shutdowns). However, for most of the models, except the full model including all inputs, the RMSE increase again thereafter. Hence, the conclusion is, that using these models it is profitable not to having too long a training period. For the full model, which does achieve the lowest RMSE, the RMSE doesn't increase before the training period length is 6 events.

Comparison of different sensor configurations

An alternative to the heating power as input to the mode, is using the electrical power to the heat pump. In the application cases we are considering, an electrical power meter must always be installed, however, the heat meter is only used for the forecasting of indoor temperature. It would decrease the installation cost, if the electrical power meter is sufficient to achieve the required accuracy. Note, in the case where the heat pump is operated by a provider of heating as a service, the heat meter is required for billing purposes, and hence does not represent an additional cost in this case.

The forecast performance measured in average over the events as a function of the horizon k is plotted below, with the *electrical power* for the heat pump as input and a training period length of 4 events, is plotted:

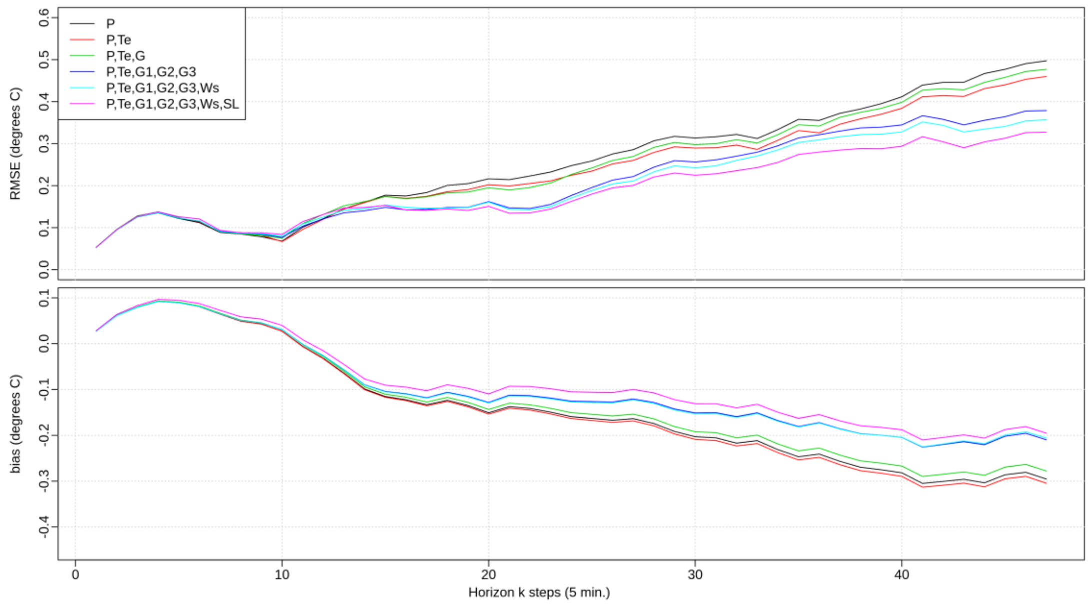

Clearly, the accuracy decrease (increasing RMSE) and the absolute bias level increase (negative bias means that the forecast is too high in average). These two are related in the way that an increase in absolute bias leads to an increase in RMSE, but an increase in RMSE doesn't mean an increase in bias.

In order to understand the impact of using the electrical power as model input, compared to using the heating power, the same plots as above are generated. If the model is sufficiently accurate using the electrical power, the heat meter is not required in the hardware setup.

The plots below are the results using the electrical power as input P:

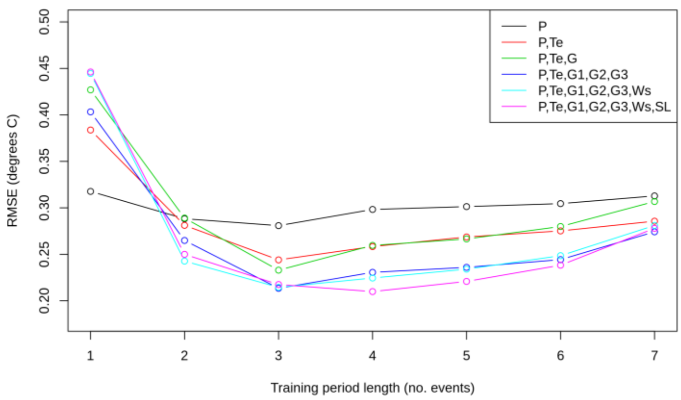

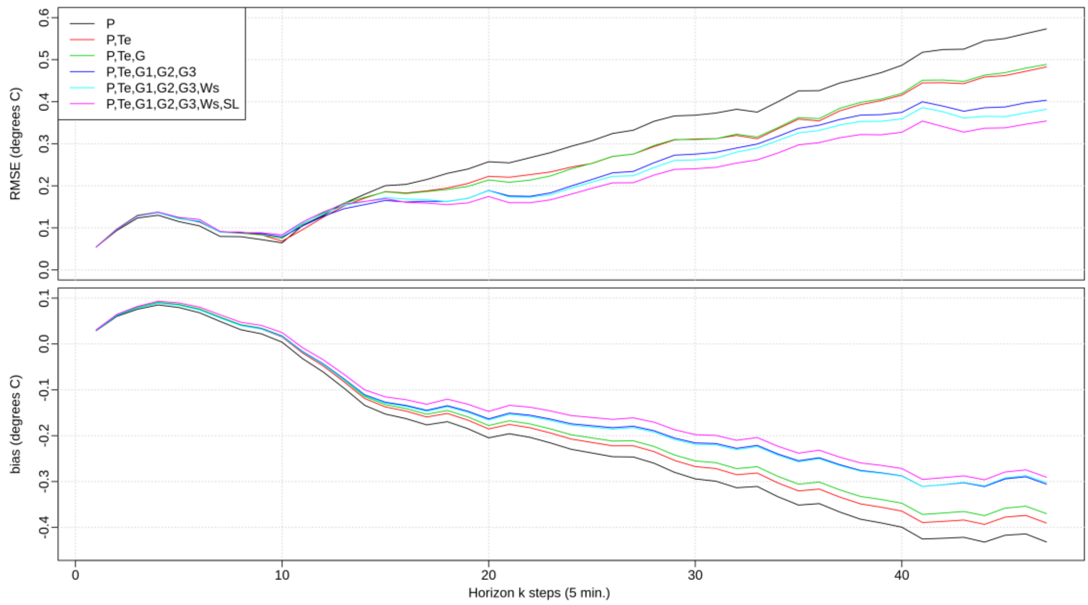

Comparing the RMSE of the full model only a very little increase is found from using the heating power as input, namely from around 0.33 to 0.35. As for the bias the increase is relatively higher, from -0.2 to -0.3.

It is clear, that using the electrical power does have an impact on the forecast performance, and especially the bias in the models will be an important aspect to improve in future work, since bias will propagate through an aggregation and lead to bias in upscaling with many systems.

Conclusions from indoor temperature forecasting:

- Some small loss in accuracy using electrical power instead of heating power as input in the models is found, mostly resulting in an increase in bias, however this can be dealt with, such that that should be only minor differences in forecasting accuracy between using the heat and electrical load as input.
- Bias should be dealt with in future works. Maybe continuous models (better for stiff systems) could be applied, optimization with regards to longer horizons or non-parametric bias correction model could be developed.
- More accurate weather forecasts will probably lead to higher forecasting accuracy.
- The current results indicate that adaptive models are needed. Using recursive models with parameter varying smoothly over time instead of the current fixed time window would be of interest.
- Studies with more buildings and automatic model selection are needed in the future. Upscaling forecast accuracy should be studied, bias is potentially an issue which must be dealt with, but could be corrected for in an aggregated model.

1.5.6 Software integration

All data is stored and managed in Energy Cool's Energy Cloud. The first step in the project was to integrate the data collected through the interface box on the sites. Initially, this was metering data from heat meters and later in the project a power meter and the heat pump was added. Indoor climate data was collected through an API at the IC-meter cloud service and weather forecasts were integrated from Open Weather Map (see https://openweathermap.org/).

The data was made available to the project partners through a REST API and through the existing Energy Cloud web interface. The web interface was used for manual inspection of data and monitoring of equipment performance on the sites. For development and test of the forecast functionality, data from the REST API was used in the forecasting tools developed in

R. Later when the developed forecast functionality was ready, an interface was developed to integrate the R functionality into Energy Cloud.

Integration of data and control of the different types of HP was also done through the on-site interface box. As an add on to the Panasonic HP, an IntesisBox was used, adding a Modbus interface to the HP. This basically allows to set a number of registers in the HP and through these, control the operational mode and different temperature thresholds. This interface was used both for controlling activations in the HP and for optimizing the operational setup. Commands for activating the HP (but setting operational mode using the Modbus registers) were implemented in the interface box and made available for Energy Cloud. All test activations were made by activating these commands in Energy Cloud, either automatically through a scheduler or manually through a web interface in Energy Cloud. The REST API was also extended to enable activations to be scheduled by the project partners.

In the operational scenario the schedule for activating the HP is made by Energy Cloud by aggregating the flexibility from the HPs together with flexibility from other assets. The overall concept for aggregating large sets of different types of assets in Energy Cloud is shown in the figure below.

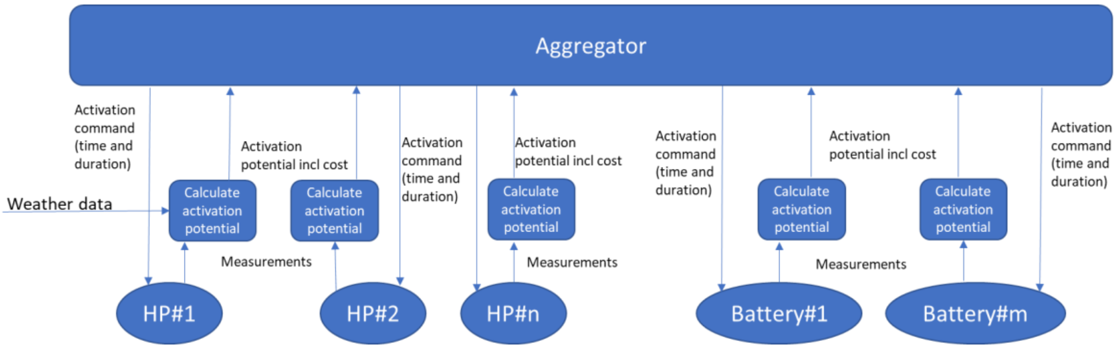

The key for combining flexibility from different types of assets in the Energy Cloud aggregator is that all calculated activation potentials are of the same type. Calculation of the activation potential depends of the type of asset and the physical characteristics. In the project, this was developed for HP using the forecasting described above. The aggregation algorithm of Energy Cloud was also extended to handle a weight or cost of an activation potential to allow for prioritizing between different types of assets with different associated costs. (Note, a cost could also be related to different types of contracts with different owners of assets, however, this had not been relevant at an earlier stage in the development of Energy Cloud.)
The cost used was a simple linear cost in the volume (watt hours) for an activation. The cost could also reflect the difference between the electricity price for the hour(s) where consumption is moved away and the cost of the hour(s) with increased consumption. Hence the cost could in principle be negative if the electricity price at the time when the HP catches up is lower than in the time period where it was forecasted to run in. For the algorithm this makes no difference since it selects the combination of activation potentials that results in the lowest cost (including negative costs).
One could image other types of costs than linear. For batteries one would expect a maximum utilization during a period of time to ensure that batteries are not worn out earlier that expected due to cycling as a result of activations. In such a case, the cost will also depend on whether the current number of cycle/activations performed on the battery exceeds the planned number of cycles or not.
The sites Energy Cool had planned to use for combining activations of HPs with activation of tele-sites (where activation potential is based on back-up batteries) were discontinued during the duration of this project (for reasons unrelated to the CASHPump project). Therefore, a combined activation with physical assets has not been performed. However, the aggregator scheduling activation potentials is the same that had earlier been used for activating back-up batteries in tele-sites. Should both types of assets be available at the same time again, the aggregator will make combined activations.

The planned integration of HP data and control into Energy Cloud has been achieved and the calculation of activation potential has been made.

1.5.7 Dissemination

During the project period we have communicated the knowhow and experience we have build. We have made a homepage where people could read about the project
http://econgrid.com/index.php/cashpump/

Furthermore, we have sent out press release when the project started. It was also posted a series of short posts at LinkedIN to raise focus on heat pumps and the project. (see Annex A)

We presented the project at the experience at "Klimalabmøde" at Middelfart municipality the 30-10-2017. This was a large event with more than 100 participants and former minister and EU commissioner Connie Hedegaard as key note speaker.
A closing workshop was held at EWII on the 19-03-2019. The program was aimed at participants owning buildings relevant for medium sized heat pumps, heat pump distributers and installers.
At both events municipalities' members, supplies and consultant where represented so they have heard about the opportunities and benefits that can be obtained by controlling heat pumps with simple equipment.

The results of the project will be presented in several papers. The following are under development:

- The load forecasting modelling techniques have been implemented in an R (ref) package, which works and have been tested in this and other projects. Some documentation are needed for the particular code, however a report describing the models in detail is finished (see attached document) and some vignettes (kind of tutorials with code examples for R packages). The report is planned to be submitted to the useR journal (ref).
- An article based on the results in this section will be published at a journal, probably Applied Energy or similar.

During the project EWII has gained a lot of knowhow about how heat pumps cloud be used in the electricity marked with flexibility. Furthermore, we have obtained knowledge about how to operated heat pumps efficiently based on the detailed data the project have delivered. That has provided EWII the opportunity to offer the new concept "heatpumps at subscription" to customer in hole Denmark which is align with our strategy for the future.

Two master students from SDU in Odense have completed their masters project at EWII and have used data from the project. They have also participated in some project meetings to learn more about heat pump operation and have presented the findings from their master project at a project meeting.

During the project period the project partner EConGrid was closed and all activities including the CASHPump project were taken over by Energy Cool (EConGrid was a daughter company of Energy Cool). A change in strategy and focus is one result of Energy Cool taking over EConGrid. The SW platform Energy cloud has continued to be operational both in the CASHPump project and commercially, but focus has shifted from aggregation of flexible energy to monitoring and optimizing sites with technical installations such as HPs and cooling equipment (Energy Cool's primary business). Energy Cool aims at using Energy Cloud, including the forecasting functionality developed in CASHPump, to optimize HPs as part of a solution for companies delivering heat as a service. This has not resulted in contracts yet, but is one element in the business plan for continuing to develop and operate Energy Cloud.

1.6 Utilization of project results

The project objectives were defined to support two overall purposes after the project

1) Improve the business case for owning/operating HP by adding additional value through payments for flexibility – and hence support the roll-out of HPs.
2) Demonstrate the technology for adding another type of asset to aggregators of small to medium assets such as EConGrid.

To our knowledge no commercial aggregators of small to medium sized assets exists in Denmark at the moment. Therefore, item two will not be used immediately after the project. However, a practical setup has been demonstrated and way to forecast and hence also baseline expected consumption has been demonstrated. The setup is sufficiently simple to be the basis for a commercial scale roll-out. The interface box should be adjusted for cost optimization purposes before initiating a commercial scale roll-out but this is considered basic changes without impact on the overall setup or functionality.

Some other areas still need to be addressed before a commercial scale roll out can be initiated. A system delivering regulating power from aggregated HPs (and possibly other types of assets) shall be technical approved by the TSO, Energinet. As the delivery will go through a Balance Responsible Party (BRP), things such a data communication and status signals shall also be agreed with the BRP. A key element of the approval is agreeing on the volume of the delivery and hence the baseline for the HP(s). The forecasting demonstrated in this project should be used as baseline and is shown to be accurate with sufficient training data. However, there is a trade-off between collecting data baselining/forecasting and delivering regulating power. How to handle this trade-off in the best way and set requirements for training data for baseline or in other ways validate/approve baselines needs to be settled before actual deliveries can start.

The developed solution can also be used to optimize HP operation with respect to spot prices and differentiated tariffs. In most cases the difference between the hourly spot price is not enough to cover the investment of interface hardware. However, significant savings can be achieved in areas with differentiated tariffs such as the one implemented by DSO Radius. Minimizing the HP operation during peak load times not only saves the customer/operator cost but also supports the overall goal of implementing peak load tariffs.

Combining this with the benefits of monitoring a HP installation and optimizing it could be a sound business case especially for companies operating HP. Delivering heat produced from HPs as a service to business customers is currently being explored by several companies, e.g. EWII, Best Green, Sustain Solutions. For these companies, operating and optimizing HP according to power cost using the forecasting demonstrated in the project, will be a benefit to the overall business case. An improved business case will increase the number of buildings "suited" for HP and hence support the general roll-out of HP.

The results and experience obtained during the project will be part of the continued work at EWII both in terms of software solutions and transferred to the business through the people who has participated in the project.

The results in the project are based on a continued development of results on forecasting indoor temperature and energy consumption at DTU. The new results obtained will also be used as basis and input for new projects.

1.7 Project conclusion and perspective

The project set out to test and demonstrate a commercial and scalable way of adding revenue from the regulating power market to owners of SME heat pumps. This required the ability to forecast the activation potential of a heat pump based on a cheap hardware setup and integration into an aggregator software solution, that can provide the accurate forecasting of activation potential needed to participate in the regulating power market.

Such a solution has been successfully tested in a real life case at Strib library and integrated into the aggregation software platform Energy Cloud.

Provided with the right set of training data, the forecasting delivers accurate results of both electrical consumption forecast and indoor temperature development forecast. However, the accuracy is depended of the training data, and this raises the question of how to specify or validate a set of training data to be used as input to the baseline required for participating in the regulating power market. This is one of the questions that needs to be addressed in co-operation with Energinet. One aspect of this question is also the fact, that deliveries of regulating power leaves periods of time (during the activation and the time period after when the heat pump increases indoor temperature back to the "normal" temperature) where data for training cannot be collected. Hence, there is a trade-off between the number of possible activations and the amount of training data available.

Given the cost of hardware and installation compared to the possible financial benefits obtainable in the regulating power market (especially assuming an independent aggregator and balance responsible party also taking part of the revenue), the overall business case is not viable at the moment. In areas where variable tariffs are implemented, the same methods can be used to optimize operation of the heat pump in relation to the tariffs. This adds a substantial revenue (saving). If the heat pump is owned and operate by a provider of heating as a service, most of the hardware and installation cost is required for that business case. In this case, the solution tested in the project can be an important addition to the total business case.

One final area for future research and development is the requirements for the building and heating system. The demonstrations in the project were performed on a building with a simple structure (one main building with few rooms) and simple heating system. For a larger scale roll-out the solution should be applicable in more complex buildings and this is also an area for future research and test.

Annex A
Bilag Pressemeddelse

Intelligent styring af varmepumper som en del af den grønne omstilling

Varmepumper, der kan justeres til et komfortabelt varmeniveau ved hjælp af intelligent software, skal sikre at små og mellemstore virksomheder skubber til den grønne omstilling.
Hvor lang tid kan man lukke helt ned for varmen fra en varmepumpe, og fortsat opnå et komfortabelt varmeniveau i små og mellemstore virksomheder? Det er udgangspunktet for et nyt samarbejde mellem softwareudvikler EConGrid, Danmarks Tekniske Universitet, DTU Compute og energirådgivningen i EWII.
Målet er, at man via intelligent måling og styring af varmepumpens varmeproduktion i små og mellemstore virksomheder, kan kontrollere den varme, som udledes i virksomhedernes lokaler. Samtidig vil der kunne opnås en økonomisk besparelse ved at strømmen, som anvendes til varmepumpen, kommer fra eller leveres til det såkaldte regulerkraftmarked. Det er Energinet.dk som systemansvarlig, der styrer regulerkraftmarkedet og balancerer produktion og forbrug af strøm til elnettet.
"Projektet bygger bl.a. på vores eksisterende viden om varmepumper, vores erfaringer med at anvende prognoser til at forudsige varmebehovet, og den viden vi har om fleksibel varmeproduktion", fortæller Martin Vesterbæk, der er senior projektleder i energirådgivningen hos EWII, og som også er én af kræfterne bag dette projekt.
Projektet trækker også på erfaringer og viden fra DTU Compute i forbindelse med CITIES-forskningsprojektet, som netop indebærer forskning i relation til Smart Grids og integration af vedvarende energi.
Lige nu bliver den intelligente software, som er udviklet af EConGrid, testet flere forskellige steder i Trekantområdet. Det er bl.a. erhvervslejemålet Gudsøgaard i Fredericia og Strib Bibliotek i Middelfart Kommune, der stiller sig til rådighed som testvirksomheder. Her er opsat intelligente følere og målere, der opsamler data, som de tre partnere skal blive klogere på.
Vindenergi er et ekstra plus
Hensigten er, at rykke et skridt nærmere uafhængighed af fossile brændsler, da varmepumper bruger strøm, som blandt andet genereres af vindmøller. Derudover håber projektets aktører at få erfaringer med fremtidig udvikling af regulerkraftmarkedet til udnyttelse af både nye og eksisterende varmepumper hos virksomhedsejere.
"Løsningen her er en del af Smart Grid – det intelligente elsystem, der er en af de grundlæggende forudsætninger for, at vi kan være en del af den grønne omstilling", siger Martin Vesterbæk.

BOKS: CASH-Pump hedder projektet
Projektets navn er "Commercial forecasting and Aggregation of SME Heat Pump" – men forkortet CASH-pump. Projektet transformerer varmepumpeejere i små og mellemstore virksomheder fra passive til proaktive forbrugere ved at designe en forretningsmodel og en teknisk løsning, der gør det let og økonomisk fordelagtigt at gå på regulerkraftmarkedet med deres varmepumper.

CASHPump
Commercial Forecasting and Aggregation of SMV Heat Pumps

BOKS: Hvad er regulerkraftmarkedet?
Regulerkraftmarkedet har til formål at opretholde balancen i elmarkedet og den overordnede stabilitet i elsystemet igennem balancering af strømproduktion og strømforbrug. Regulerkraft er en strømreserve, hvor en given strømpulje enten opreguleres eller nedreguleres.

BOKS: Støtte til projektet
Projektet CASH-pump er støttet af EUDP, der er Energiteknologisk Udviklings- og Demonstrationsprogram. Ordningen støtter ny teknologi inden for energi, som kan bidrage til at indfri Danmarks målsætningen inden for energi og klima.

Projektpartnere:

EUDP   EConGrid   DTU   EWII

Bilag linkedIN

Opslag 1 – september

Hvor lang tid kan man lukke helt ned for varmen fra en varmepumpe, og fortsat opnå et komfortabelt varmeniveau? Det har jeg sammen med DTU Compute og softwareudviklere fra EConGrid sat mig for at finde ud af. Det er et spændende projekt som bl.a. bygger på erfaringer om varmepumper og tidligere forskningsprojekter om Smart Grid og integration af vedvarende energi.

Short link: https://goo.gl/RBTJdx
Link: https://ewii.com/om-ewii/nyheder-og-presse/Intelligent-styring-af-varmepumper-som-en-del-af-den-gr%C3%B8nne-om-stilling?utm_source=linkedin&utm_term=september&utm_content=martin%20vesterbaek

Opslag 2 – oktober

Varmepumper, der kan justeres til et komfortabelt varmeniveau ved hjælp af intelligent software, skal være med til at sikre at små og mellemstore virksomheder skubber til den grønne omstilling – det er en del af min hver dag sammen med DTU Compute og softwareudviklere fra EConGrid.

Short link: https://goo.gl/sSn62u
Link: https://ewii.com/om-ewii/nyheder-og-presse/Intelligent-styring-af-varmepumper-som-en-del-af-den-gr%C3%B8nne-omstilling?utm_source=linkedin&utm_term=oktober&utm_content=martin%20vesterbaek

Opslag 3 – november

Skal vi skubbe til den grønne omstilling og Smart Grid?
Lige nu tester vi sammen med DTU Compute og softwareudviklere fra EConGrid intelligent styring af varmepumper hos erhvervslejemålet Gudsøgaard i Fredericia og Strib Bibliotek i Middelfart Kommune. Her samler vi brugbar data op, som giver os indblik i, om varmepumper kan sikre et komfortabelt varmeniveau med styring der både giver komfort og regulerekraft i el-nettet.

Short link: https://goo.gl/hcYg1X
Link: https://ewii.com/om-ewii/nyheder-og-presse/Intelligent-styring-af-varmepumper-som-en-del-af-den-gr%C3%B8nne-om-stilling?utm_source=linkedin&utm_term=november&utm_content=martin%20vesterbaek