Predicting depreciation of flexibility in thermal buffer clusters in the PowerMatcher smart grid

Bachelorproject

N.A. van der Veen, s2041928, n.a.van.der.Veen@student.rug.nl,
Supervisors: P.A. MacDougall* and M.A. Wiering†

July 26, 2013

Abstract

An increase of less controllable supply and demand on the energy market asks for effective market balancing. The PowerMatcher is a market-based coordination mechanism for smart grid systems in which the balance is held by shifting the price. Sometimes the PowerMatcher is not able to hold the balance. It has a depreciation of the flexibility in supply or demand. We simulated the behavior of thermal clusters reacting on: 1) a request for demand or supply, 2) changes in heat demand of the households. We found a few variables that can be used as predictors for the depreciation of flexibility. We used these variables for creating models that can predict flexibility depreciation. We could predict depreciation of flexibility half an hour before with an accuracy of 76% and a hitrate of 73%. These kinds of predictions could help grid operators to prevent violation of security of supply or safety constraints.

1 Introduction

1.1 Smart grids: balancing supply and demand by using flexiblity

In today's power grid, supply follows the demand [3]. The grid was designed to operate vertically [14]: central generation, transmission via high voltage grid to substations and distribution to consumers via medium and low voltage power lines. Since the world is becoming more conscious of the necessity for clean and sustainable energy sources, the use of renewables, for example wind [5] and solar [13], is increasing. Renewables as solar and wind sources are non-flexible and are not able to follow the demand. The supply from these renewables at one moment is rather different from another moment due to weather changes. In today's power grid, a certain amount of the produced solar and wind energy would need to be curtailed or exported [17]. These renewables are not always produced centrally, therefore system operators are facing another challenge: regulating supply penetration on decentral level [15].

The supply side of electricity systems will become larger, more diverse, more decentralized and less controllable. The demand side will also be larger [1], especially on the household level. People are using more electrical devices and the use of electrical vehicles is increasing [18]. The household demand is in the current energy market seen as a non-flexible factor [3], but it may provide some flexibility to the supply and demand balance. For example a fridge could be turned off as long as the fridge keeps a climate that is acceptable according to food standards. Utilizing the flexibility of end consumers and producers, can reduce investments and operational costs needed for a future reliable electricity system. In section 1.3 flexibility is defined and further explained.

There is a strong consensus [10] for promoting the necessity of a new electricity grid that is able to balance supply and demand better: the Smart Grid. In the Smart Grid new information and communication technologies are used to revolutionize the

*TNO
†University of Groningen, Department of Artificial Intelligence

electrical power grid [4]. There are worldwide multiple smart grid projects and billions are invested in these projects [6]. There are several visions on the aim, design and implementation of smart grids, but in all visions the flexibility of consumers and producers is used to match the demand and supply at a certain point in time better.

We researched the flexibility in demand and supply in the PowerMatcher, a decentralized market-based smart grid using agent-based software to balance supply and demand. The PowerMatcher is further explained in section 1.2.

One of the strengths of an agent-based implementation is its autonomous and flexible behavior [19]. A drawback of an agent-based system is that it is not easy to monitor and control, because the emergent behavior of the system is inherently unpredictable [7]. A grid operator is responsible to maintain the comfort requirements of users and needs to import or export energy to the cluster if the cluster could not balance itself. Therefore they have to be able to predict the behavior of the cluster and to take action. Especially in uncommon situations, a grid operator needs to be prepared for balancing problems. We developed and tested a few methods for forecasting depreciation of flexibility using actual values of some variables in the cluster.

Our goal was to find a method that can forecast depreciation of flexibility in an acceptable time course. In power grid systems, an hour would be a preferred time, but forecasting methods that predict the problem half an hour before are also interesting. Today's most rapid power plants have a start-up time of 30 minutes [2].

We could predict depreciation of flexibility half an hour before with an accuracy of 76% and a hitrate of 73% using price fluctuations. We discovered the potential of variables as fill level of the total system and percentage of devices turned on and found some effects that could indicate depreciation of flexibility. The models we made to predict depreciation of flexibility based on these effects were not as succesful as the price fluctuation indicator. The effects were not enough present in the simulation data we used for testing the model performance.

1.2 The PowerMatcher

The PowerMatcher is a multi-agent based distributed software system that coordinates the energy supply and demand of houses and small offices in the network. Figure 1 shows the hierarchical structure of a Powermatcher cluster.

Within a PowerMatcher cluster, agents are organized into a logical tree. Devices in houses and small offices are represented by the local device agents in the leaves. To represent a clusters of devices, concentrator agents are used. The concentrator agent is able to do (local) congestion management. The root of the tree is formed by the auctioneer agent, a unique agent that handles the price forming by searching for the equilibrium price. Sometimes an objective agent is added to the tree. The function of this agent is to mimic the business logic for this specific cluster by adding demand or supply [9]. We use this agent in our experiments to force a cluster to follow a sine wave demand or supply request pattern.

Due to aggregation in a binary tree, the complexity of the market algorithm is $O(\log \alpha)$. Where α is the number of device agents [8].

1.2.1 Agent descriptions

Local device agents The local device agents are control agents that aim to reach the goals of the local device in an economical optimal way. These agents communicate with the cluster by buying or selling electricity. They give their bids (via concentrator agents) to the auctioneer and receive price updates from the auctioneer. The local device agents determine how large the amount of power is that the agent will produce or consume, based on its own latest bid and the current electricity price on the market.

Auctioneer agent The auctioneer agent performs the price-forming process. The auctioneer collects the current bids of all agents and searches for the equilibrium price. It returns a price update if the equilibrium is changed.

Concentrator agents Concentrator agents are representatives of a cluster of local device agents. These agents collect the bids of the agents and communicate this to the auctioneer. In the opposite direction, they receive price updates from the auctioneer and communicate this update to local device agents.

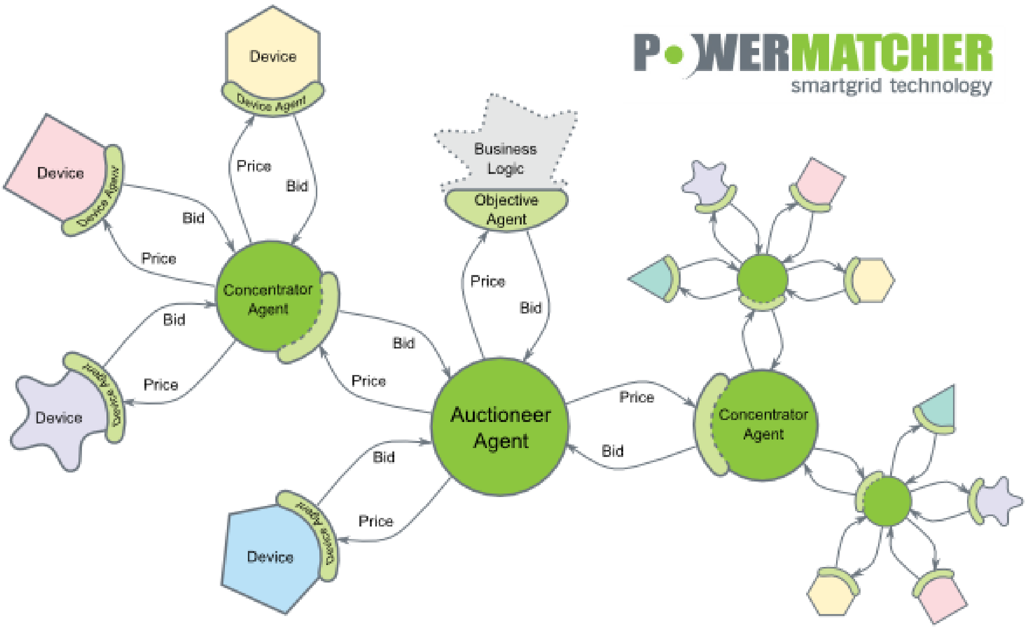

Figure 1: Schematic overview of the PowerMatcher concept

Objective agents When the objective agent is absent, a cluster strives for an equal supply and demand within the cluster itself. Depending on the specific application, the goal of the cluster might be different. This goal can be realized by implementing an objective agent. The objective agent adds demand or supply to the cluster, so the balance of the cluster will shift to the purpose of the cluster. A purpose of a cluster could be: producing a certain amount of energy at night and consuming another amount at day.

The PowerMatcher technology has been implemented in a number of small field tests scaling up to 22 households [9] and large simulations (about 3000 households) [11]. In [11], simulations show that the PowerMatcher smart grid offers a huge potential in utilizing flexibility in demand and supply of households to accommodate mass integration (60%-90%) of green electricity. In field tests over a longer period peak load reduction of 15% was achieved [9]. Before the end of 2013, more than 1000 households will be powerMatcher-equipped [8].

1.2.2 Bids and prices

The local device agents give bids to the auctioneer. With these bids, the agents communicate their flexibility to the auctioneer. A heat-pump is a local device agent that turns electrical energy into thermal energy. A heat-pump has a minimal and a maximal temperature between which the thermal buffer needs to maintain the temperature.

We represent in figure 2 the 'space' between the minimal and maximal temperature as a buffer. If the minimal temperature is reached, then we have situation A in which the heat pump has to turn on whatever the price is. If the maximal temperature is reached, then we have situation C: the heat pump wants to turn off or stay off, independent of the price. If the temperature is somewhere between the minimal and maximal temperature, then we have situation B. In this situation the heat pump has flexibility. If the price is low, then the heat pump turns on, otherwise it turns off.

Most devices have a must-run situation. For example a washing machine cannot be turned off when it is on. Also a heat pump has a minimal run-time. It is not efficient to turn off and on the heat pump every minute.

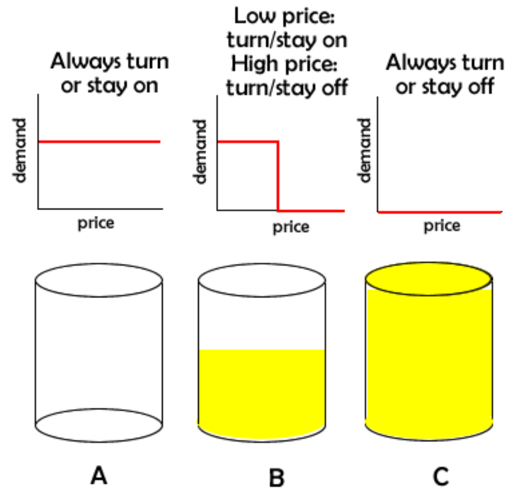

Figure 2: The fill levels of a heat pumps

1.3 Ramp up and ramp down flexibility of household clusters in smart grids

A device has flexibility if it is capable of shifting its production or consumption of energy in time within the boundaries of end-user comfort requirements [16]. A household device provides a relatively low demand or supply, so their flexibility is not interesting to the stakeholders in the energy grid. When we connect a lot of these devices into a cluster, we can use the sum of the flexibility of the devices. The flexibility of this cluster is measured as the amount of power increase or decrease, with respect to its current power consumption or production, that can be sustained for a given period of time [16].

For example, if you could turn off 1000 fridges at a certain time t, those were operating at 700 Watt, then the ramp down flexibility is $700 * 1000$ Watt. If these fridges are turned off and could be turned on, then the ramp up flexibility is 700 kW. This kind of clusters could be used as a Virtual Power Plant, because they can deliver some supply (or less demand) to the grid during a certain time. These clusters could also be used as a flexible large-scale consumer. How the cluster could be used is dependent of the composition and characteristics of the cluster.

We defined that depreciation of flexibility (F) is the case if:

$$F^{+} < P_{des\_out} \lor F^{-} > P_{des\_out} \tag{1.1}$$

F^{+} and F^{-} are the flexibility to ramp up and ramp down. P_{des_out} is the desired output of the cluster. This is zero, if the cluster is not asked to bring some supply or demand to the grid. If the cluster is asked to produce or consume power, then this is the desired output.

2 Methods

2.1 Simulation tool and agent models

We used the PowerMatcher Simulation Tool for the experiments. This tool has been used before in [11] and [12]. The simulation tool simulates the behavior of a PowerMatcher controlled smart grid. The demand or supply of the device models are based on data from dutch households. The simulation behavior is therefore representing the behavior of a PowerMatcher controlled smart grid acting in the real world.

We used two types of household device agents in the PowerMatcher simulations: heat pump agents and micro-CHP agents. Heat pumps are devices that transfer heat by using electrical energy. Micro-CHPs generate heat and power using fuel power. If micro-CHPs are turned on for heating, they produce power as well. If households request for heating, than both micro-CHPs and heat pumps are turned on. Heat pumps are consumers of electrical energy and micro-CHPs are producers.

The objective agent is the third agent we use in the PowerMatcher Simulation Tool. The objective agent adds demand or supply from external sources to the cluster, so the balance of the cluster will shift to the purpose of the cluster. A system operator can use the objective agent if he would like to use the PowerMatcher cluster as a Virtual Power Plant.

2.2 Experiment setup

The main objective of the simulation setup is to simulate a balanced cluster that has sometimes depreciation of flexibility. We chose to use 1000 micro-CHPs and 1000 heat pumps that have the same specifications and the same capacity. The specifications of the heat pumps and micro-CHPs are found in table 1. We made the decision for 2000 devices because the simulation run-time would be in a reasonable time. It is considered to be large enough to be representative for the behavior of thermal buffer clusters in the real word.

The heat demand from the 2000 households is generated by a pattern generator that generates heat demand patterns based on real data. If people come home late afternoon, we see that the heat demand peak is growing. Average, the behavior of the households that have a micro-CHPs and households that have a heat pump are the same during the day. Since the characteristics of the heat pumps and micro-CHPs are the same, the sum of power supply of the micro-CHPs and the sum of power demand of the heat pumps are near equal. The cluster consisting of 2000 devices is able to balance its own demand and supply.

The output of an ideal balanced cluster is an output request for demand of supply equal to zero Watt. In smart grids, we have the possiblity to control the output of a cluster in a certain way. In a PowerMatcher controlled smart grid we can do this by using an objective agent. We choose to use a remarkable pattern (a sine wave) for the controlled profile request, so we can easily trace the influence of the pattern on the cluster behavior.

During the simulation, every virtual minute there is an interaction between the device agents, concentrator agents and auctioneer agent. Therefore, every minute there is a new price determined by the auctioneer based on the bids of the devices agents. Every minute, the devices update their willingness to bid based on the actual price and their state. The actual ramp up and ramp down flexibility values are available on auctioneer level when the auctioneer received the actual bids of the devices agent.

2.2.1 Following a profile

The sine wave that gives the demand or supply of the objective agent has to be in proportion with the cluster. It cannot be too small, otherwise we could not trace its influence on the cluster. It cannot be too large either, because we are not interested in situations in which the cluster has too much depreciation of flexibility. We used a sine wave with a period of 1/5 of a day. The amplitude was found by sending different constant controlled profiles and find then a constant demand that can be followed by the cluster during 1/10 of a day. We took the value of this controlled profile as amplitude.

We can use this rule-of-thumb in situations where the influence of devices with a negative nominal electric power (so they are on the supply side) is (almost) equal to the influence of devices with a positive nominal electric power. If there is no balance between these two groups in the cluster, then the objective agent has to compensate for the difference. In this experiment, we used only clusters that are balanced, as you can see in table 1. Therefore we could use this rule-of-thumb without compensating for imbalances.

The objective agent also needs to compensate if the volume and/or temperature boundaries of the thermal devices are too small to deal with the heat demand. Especially in winter months or in other periods with a high amount of heat demand, this could be the case. So even a cluster that in fact is balanced, like the cluster we use, can not always follow a sine wave. It even cannot balance itself. Therefore the devices we use in our clusters as shown in table 1 are chosen so they can deal with heat demand profiles that are normal for average dutch households in the month March.

	$P_{nom,th}$ (W) *	$P_{nom,el}$ (W) †	V(L)	T_{min} ($^{\circ}C$)	T_{max}($^{\circ}C$)	minimal runtime(sec)
micro-CHP	4000	1000	120	40	70	900
Heat Pump	4000	1000	120	40	70	900

Table 1: Specification of used devices in the simulations

2.3 Procedure

We simulated the behavior of our cluster with heat pumps and micro-CHPs and an objective agent asking for an appropriate profile during the month March. There is some change in heat demand during the day. Sometimes, the cluster is not able to deal with these fluctuations in heat demand and there is depreciation of flexibility. Those events we call 'red events': the ramp up flexibility is lower than the profile asked by the objective agent (the sine wave) or the ramp down flexibility is higher than the asked profile (see also the definition of depreciation of flexibility in section 1.3).

Before 'red events' are the case, there is probably a period in which the cluster is healthy, but some variables maybe could indicate that the probability for a coming 'red event' is high. We decided to cut the simulation in time windows and use a 'sliding window', to go along these windows and determine the state of the cluster in that window. It is our goal to make a model that could indicate the period before the 'red events' (the 'yellow event') and could give an alarm.

To decide afterwards when the model should give the alarms, we labeled the windows based on the available flexibility in the window and on the label given to the windows before. If there is depreciation of flexibility in one or more data points we give the label 'red event'. If there is no depreciation of flexibility in the window, but an hour later there is a red event, we label the event as 'yellow event'. If there is no depreciation of flexibility and in the near future there is no red event, we give the label 'green event'. Also we give the label 'black event' a few minutes after a red event.

We used 2 simulated weeks (20160 data points) as training set. First we found out what indicators are useful for prediction. After that, we decided rules for our prediction models based on the observations we did in the training set. Then we tested the models on the test set of 201500 (10 times 2 simulated weeks) data points.

3 Results

3.1 Finding indicators

Figure 3 shows the flexibility of the cluster consisting of 1000 heat pumps and 1000 microCHPs that is forced to follow the sine wave demand/supply pattern of the objective agent with amplitude= 200.000 W and period = 0.2 day. Between $t = 0.5$ and $t = 0.85$, there is depreciation of ramp up flexibility. Before and after this depreciation, the cluster flexibility is following the requested sine wave profile. We defined this kind of behavior 'healthy': there is no flexibility depreciation.

We see in figure 3 that after $t = 0.5$, there is depreciation of flexibility and the cluster is not able to follow the sine wave anymore. The depreciation of flexibility is an effect of changes in the heat demand of the end-users. These changes in heat demand cause changes in some variables that are influencing the flexibility. We assume that these variables are following also the sine wave input in an approximately constant way. If the effect of the changes is traceable in the flexibility (the output), then the cause of the depreciation is probably also traceable in the patterns of the cluster variables.

We selected three variables as possible indicators for depreciation of flexibility:

- Average fill level of the devices with positive nominal power ($\overline{P}_{+}$) and average fill level of devices with negative nominal power $\overline{P}_{-}$.
- Average amount of flexible devices that are turned on. We also make here a distinction between devices with negative and devices with positive nominal power.
- The price given by the auctioneer.

Our choice for the first two variables are based on observations as shown in figure 4. If you compare figure 4 with figure 3, then you can see that the two variables are following the sine wave pattern. Before the cluster turns in a flexibility depreciation

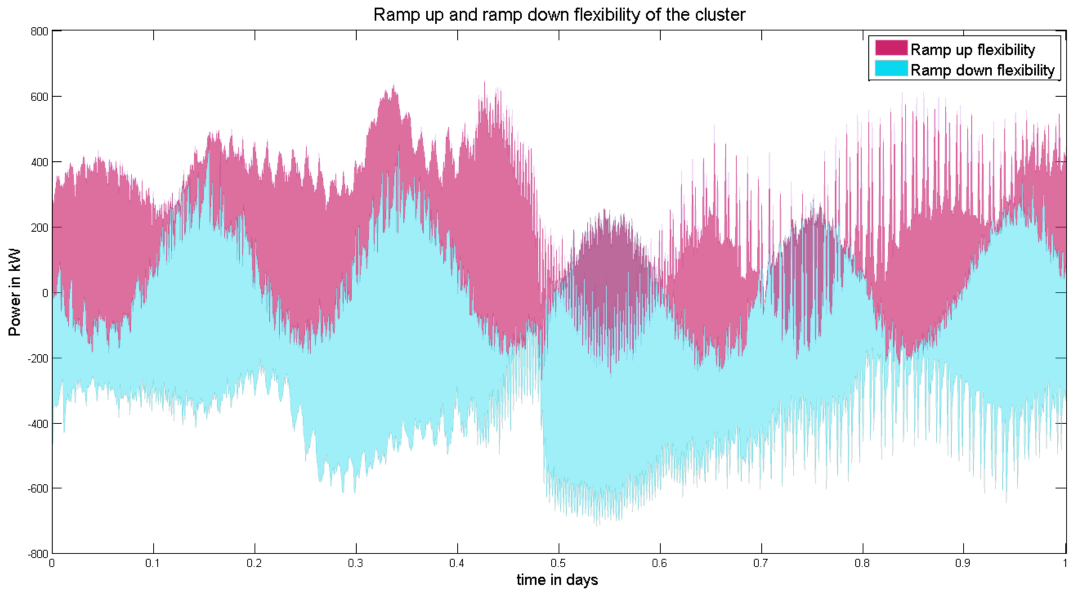

Figure 3: Ramp up and ramp down flexibility of a cluster consisting of 1000 heat pumps and 1000 microCHPs that is forced to follow a sine wave demand/supply pattern with amplitude=200.000 W and period = 0.2 day. The heat demand pattern is based on the heat demand patterns of normal households on March 1st. The boundary between the surfaces that represent the flexibility is the output of the cluster.

situation, we are able to see that the pattern of the two variables is changing. If we have a method that could monitor this behavior (not following the sine wave anymore), as humans are able to do, we could use these variables to predict depreciation of flexibility an hour before the actual depreciation.

The third variable (price level) follows also a sine wave, but that sine wave is not in phase with the sine wave input. The price follows on the sine wave input with a shift of $1/4$ phase to the right as shown in figure 5. Remarkable is the fast changing price after $t = 0.45$. These price fluctuations are followed by depreciation of flexibility. Therefore, these fluctuations could maybe be used as an indicator for a decrease of flexibility.

We used the three selected variables to indicate the 'yellow events' and so for predicting the 'red events'. Therefore we looked at the correlation between the input sine wave (the controlled profile) and these variables. We have chosen a few derived indicators that we would like to use for prediction:

- Correlation between input sine wave and difference in average amount of devices that are turned on between devices with positive (heat pumps in our experiment) and negative (micro-CHPs in our experiment) nominal electric power.
- Correlation between input sine wave and difference in average fill level between devices with positive (heat pumps in our experiment) and negative (micro-CHPs in our experiment) nominal electric power.
- The correlation between input sine wave and the price determined by the auctioneer.

3.2 Two algorithms

We chose two approaches to use the indicators for predicting 'red events' by indicating 'pre-red events' (the 'yellow events'). The idea is that the algorithm gives an alarm if it finds that one or more variables are satisfying 'yellow event' conditions.

The first approach is based on the idea that the input sine wave induces sine wave patterns in other variables if the cluster behaves well. We saw in the

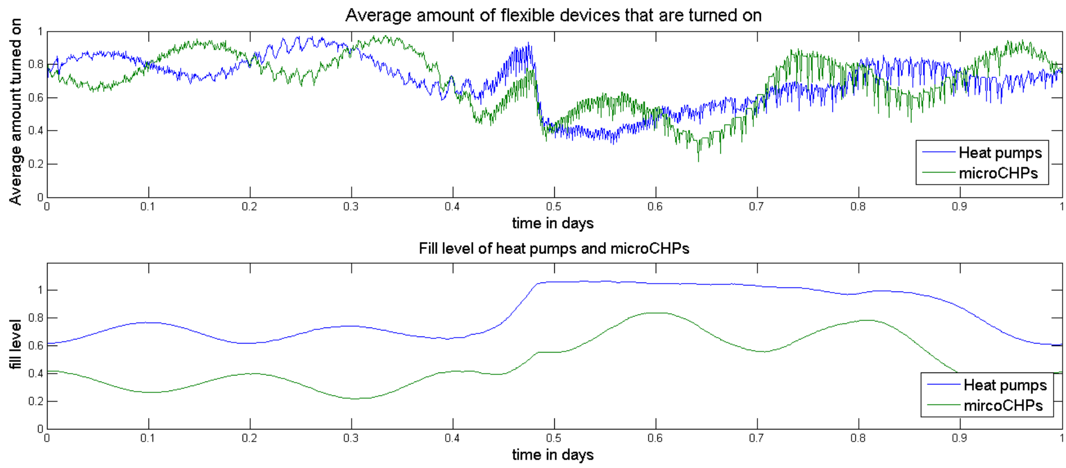

Figure 4: Fill level of the thermal buffers (below) and the average amount of flexible devices that are turned on (above) for the cluster in figure 3

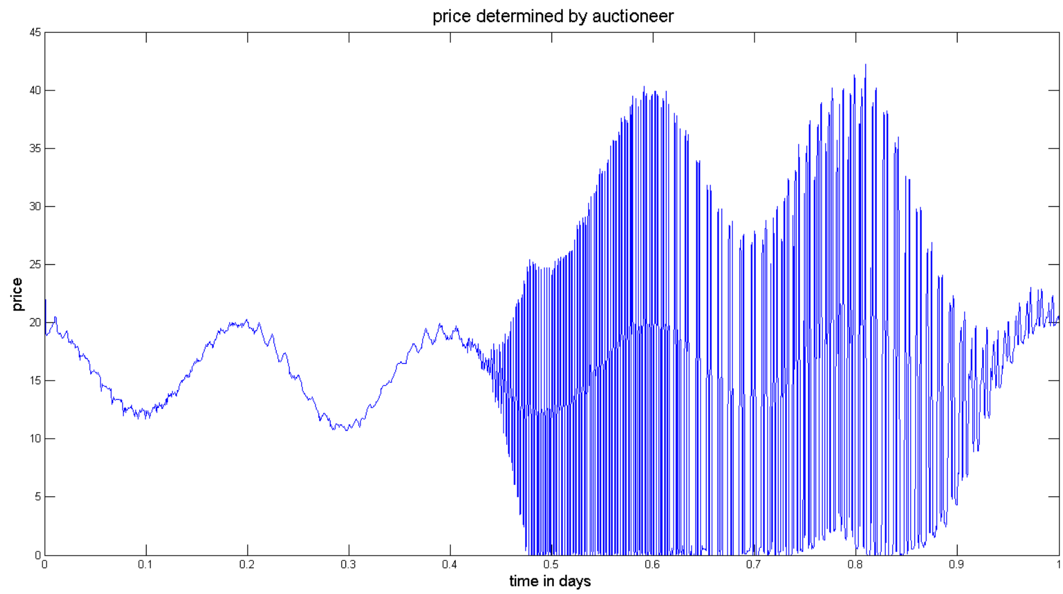

Figure 5: Price level during the simulation of the cluster in figure 3

test data (see the example in figures 3, 4 and 5) that in a 'yellow event' period, the indicator variables are no longer following the sine wave.

The algorithm draws for n data points in a sliding window the most optimal linear regression model using least-square fit. The independent variable is the input sine wave (the request from the objective agent) and the dependent variable is one of the indicator values. We made for the test set graphs to find out what information about the residuals is useful. This information from the test set was used to calibrate the model that gives the alarms for red events. The algorithm is described in the code 3.1.

Algorithm 3.1 Algorithm 1: This is the final algorithm used on the test set. In the training phase, we used also other rules in the latest if-loop.

```
for every window n do
    make linear regression model M for sine wave input and variable of interest
    log the residual on M for every data point x in window n
    take the mean for all x
    if mean x > threshold th AND mean residuals for p previous steps < step p+1 then
        give alarm
    end if
end for
```

Algorithm 3.2 Algorithm 2. The accepted residual R could for example be defined as: max over all windows n in model M of mean residual of window

```
for every window n do
    label n as red, yellow, green or black
    if n.label == green then
        if model M not available then
            update model M with data points in window n
        else
            log residuals r for data points x in window n
            if Mean residuals r < accepted residual R then
                update model M with data points in window n
            else
                give alarm
            end if
        end if
    end if
end for
```

The second approach is based on the idea that a 'healthy' cluster is following a certain profile that is in line with the sine wave input. The algorithm learns a model of this profile by learning from examples of 'green events'. The sliding window approach makes online learning possible: The model is updated by the data from a new sliding window on condition that the window is not a red event and on the condition that the data fit in the model. If the data do not fit in the model, an alarm is given. Algorithm 2 is described in code 3.2.

3.3 Results of the algorithms

3.3.1 Results of algorithm 1

We tested the performance of the algorithm by comparing the given alarms by the labeling (afterwards) of sliding windows based on flexibility. The sliding windows are labeled based on the flexibility (red event, yellow event, green event, black event). If the flexibility is a 'green event', but in less than 72 minutes (4 sliding windows) later we notice a red event, the event is labeled as a 'yellow event'.

Finding model parameters from training data. We found the mean of the residuals from the price model as the most useful indicator for the approach implemented in sliding window algorithm 1. Figure 6 shows the mean residuals for every sliding window during a simulation. The residuals from the price model are increasing if a red event is going to happen. We use this fact to give the alarms together with the use of a residual threshold.

We found that a threshold of mean price residual = 14 gives the best results on the training set (93% hitrate). An alarm is then given for a sliding window if the mean price residual is 14 or higher and the mean price residual of the previous sliding window is lower than this sliding window.

Test results of best algorithm. We tested the best algorithm for the price indicator on the test set of 201500 data points. We got an accuracy rate of 76% and a hitrate of 73% (61% for early events) of the red events that need to be indicated (see table 2). Red events that need to be indicated are windows labeled as red event after a yellow period. If the previous window was also a red event, then the event is counted as the same red event. Successful alarm points are alarms given during yellow events. The total amount of alarm points are the alarm points given during green and yellow events. If an event is red or black, there is never an alarm given.

50% of all data points are not red events. If we choose to select 50% of these points to be an alarm point, we would assign $0.5 * 0.5 * 201500$ alarm points. Since every red event has a yellow period of 4 data point, $0.5 * 4 * 377$ yellow events will be classified as alarm point. The hitrate of a random alarm assignment (probability of being assigned as an alarm point is 50%) would be 15%.

Results for other indicators. We found that the other indicators are less useful for predicting red events. For example, figure 7 shows the residual plot for the sliding windows based on the percentage of heat pumps that are turned on. Whatever threshold we choose, the accuracy is not good, because the residuals of a lot of the green and red events are in this range. We see also that the residuals are mostly increasing before a red event, but this is also the case for some green events, so the accuracy will remain low.

3.3.2 Results of algorithm 2

The idea of algorithm 2 was to make a model that follows the correlation between indicators and sine wave input. The algorithm learns from the previous green events. Next data points are tested by the model: If their deviation from the model is too large, then they are indicated as pre-red events. Since we saw in our data set examples like figure 4, we had reason to think that this information would bring us a good indicator. Nevertheless, this kind of information is not general enough. The accuracy for average heat pumps turned on, average micro-CHPs turned on, fill level for heat pumps and fill level for CHPs were below 30%.

4 Discussion

Images of indicators like figure 4 indicate that the fill level and average amount of devices turned on could be useful for prediction. Nevertheless the algorithm we built could not use these indicators with success. How could we explain this? Since we use in our simulation randomly generated heat demand patterns, there is a possibility that those patterns are responsible for the promising data curves. Therefore it is useful to research the influence of these heat demand patterns.

Also it is useful to research the influence of the initial fill levels. If the initial fill level is relatively high, the fill level will be higher in the next time steps, if flexibility is asked from the devices. High fill levels are causing inflexible behavior of the cluster. In the used simulations it was not possible to control the initial fill levels, so the initial fill levels were different. The behavior of the clusters is similar, but if the initial fill levels are relatively close to the maximum or minimum, there is more depreciation of flexibility in the cluster.

Figure 8 and figure 9 show the fill levels and average amount of devices turned on for two different simulations with the same characteristics and simulated on the same day. Only initial fill levels and some random processes are different. The initial fill level of the heat pumps in figure 8 starts around 70% and the initial fill level of the micro-CHPs around 10%. We see that when the heat demand of the households increases, the cluster is not able to turn on more heat pumps, because the buffers are almost 100% filled. Micro-CHPs want to turn on, because more heat demand is asked, but because there is no much demand flexibility from the heat pumps, this is not possible.

The amount of red events was around 50% in the test set, what is a bit high. We saw that the amount and length of red events was increasing during the 2 week simulation. This is an interesting effect that needs further research.

A drawback of the simulations we used is that we could not control all parameters. There are simulation settings possible in which this is possible, but due to technical problems, we were not able to use these settings. Future research has to be done with these settings, to research further the interesting features discovered in this research.

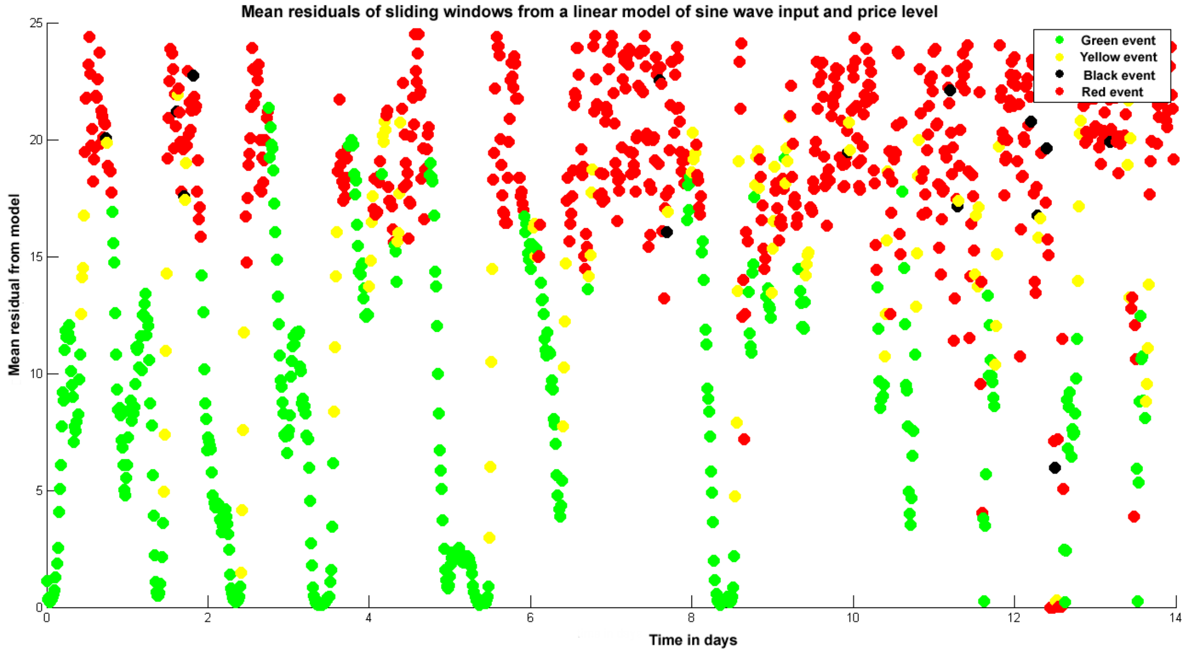

Figure 6

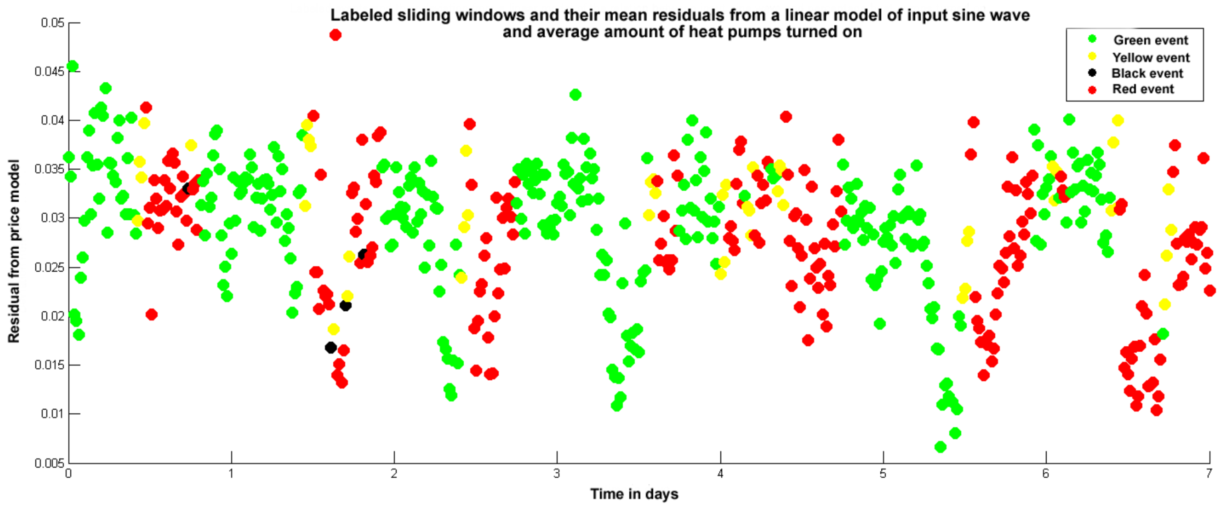

Figure 7

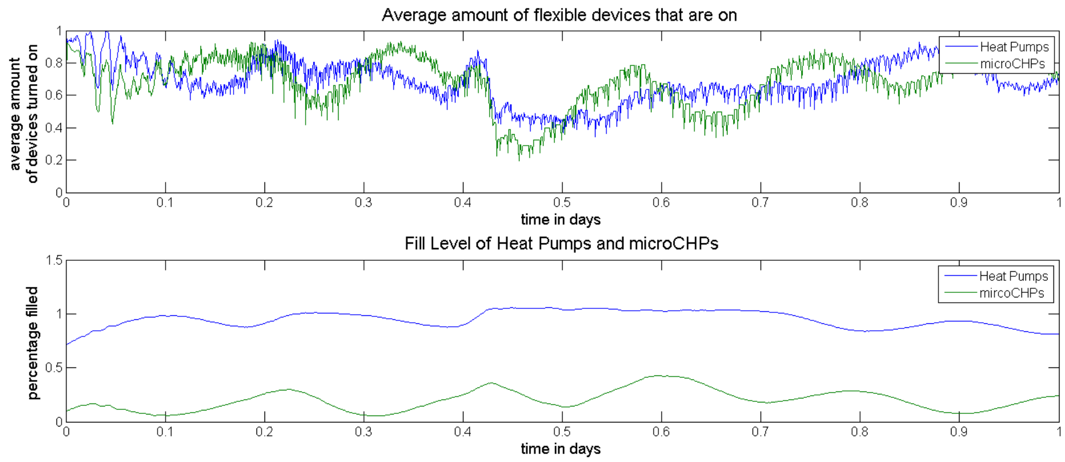

Figure 8: Fill level and percentage devices turned on for simulation 1, march 2nd

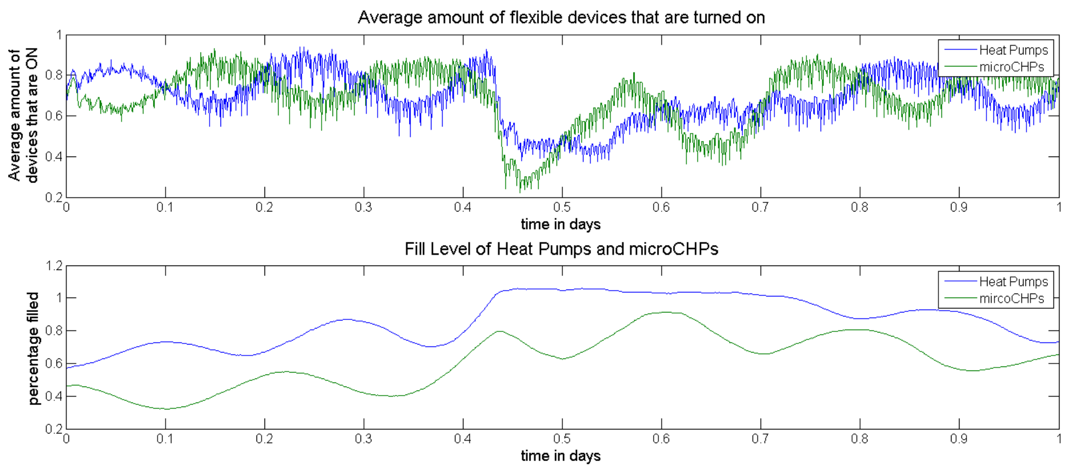

Figure 9: Fill level and percentage devices turned on for simulation 2, march 2nd

Red events that need to be indicated	377
Indicated red events	275
Early indicated red events‡	230
Hitrate (%)	73
Hitrate early indicators (%)	61
Successful alarm points	503
Alarm points	658
Accuracy (%)	76

Table 2: Results of algorithm 1 over 201500 data points, indicator = price level. Red events that need to be indicated are windows labeled as red event after a yellow period. If the previous window was also a red event, than the event is counted as the same red event.

5 Acknowledgement

A special thank you goes to those who contributed next to my supervisors to this thesis: dr. Koek Kok* for his valuable comments and ideas on the experiment setup and interpretation of the results. Prof. dr. ir. Marco Aiello§ for his helpful information about smart grids.

References

[1] International Energy Agency. *Key World Energy Statistics 2012*. IEA, 2012.

[2] Lothar Balling (Siemens AG). Fast cycling and rapid start-up: new generation of plants achieves impressive results. *Modern Power Systems*, pages 35–41, january 2011.

[3] John Casazza. *Understanding electric power systems: an overview of the technology and the marketplace*. John Wiley & Sons, 2003.

[4] Janaka Ekanayake, Nick Jenkins, Kithsiri Liyanage, Jianzhong Wu, and Akihiko Yokoyama. *Smart grid: technology and applications*. John Wiley & Sons, 2012.

[5] Lauha Fried, Steve Sawyer, Shruti Shukla, and Liming Qiao (eds.). *Global Wind 2012 Report*. Global Wind Energy Council, 2013.

[6] Vincenzo Giordano, Flavia Gangale, Gianluca Fulli, Manuel Sánchez Jiménez, Ijeoma Onyeji, Alexandru Colta, Ioulia Papaioannou, Anna Mengolini, Corina Alecu, Tauno Ojala, et al. Smart grid projects in europe. *JRC Reference Reports*, 8, 2011.

[7] Nicholas R. Jennings. On agent-based software engineering. *Artificial Intelligence*, 117(2):277 – 296, 2000.

[8] J.K. Kok. *The PowerMatcher, Smart Coordination for the Smart Electricity Grid*. PhD thesis, Vrije Universiteit Amsterdam, 2013.

[9] J.K. Kok, .B Roossien, P.A. MacDougall, O.P. Pruissen, G. Venekamp, I.G. Kamphuis, J.A.W. Laarakkers, and C.J. Warmer. Dynamic pricing by scalable energy management systems - field experiences and simulation results using powermatcher. In *IEEE Power and Energy Society General Meeting*, pages 1 –8, july 2012.

[10] Ruofei Ma, Hsiao-Hwa Chen, Yu-Ren Huang, and Weixiao Meng. Smart grid communication: Its challenges and opportunities. *IEEE Transactions on Smart Grids*, 4(1):36–46, 2013.

[11] P. MacDougall, C. Warmer, and Koen Kok. Mitigation of wind power fluctuations by intelligent response of demand and distributed generation. In *2nd IEEE PES International Conference and Exhibition on Innovative Smart Grid Technologies (ISGT Europe)*, pages 1 – 6, dec 2011.

[12] Pamela MacDougall, Koen Kok, Cor Warmer, and Bart Roossien. Quantifiying flexibility for smart grid services. *IEEE Transactions on Smart Grids*, 2013.

§University of Groningen, Johann Bernoulli Inst. for Math. and CompSc.

[13] Gatan Masson, Marie Latour, Manol Rekinger, Ioannis-Thomas Theologitis, Myrtho Papoutsi, and Craig Winneker (ed.). *Global Market Outlook for Photovoltaics for 2013-2017*. European Photovoltaic Industry Association, may 2013.

[14] James Momoh. *Smart grid: fundamentals of design and analysis*. The Institute of Electrical and Electronics Engineers, John Wiley & Sons, 2010.

[15] S. Prousch, C. Breuer, and A. Moser. Optimization of decentralized energy supply systems. In *7th International Conference on the European Energy Market (EEM)*, pages 1–6, 2010.

[16] Bart Roossien. Mathematical quantification of near real-time flexibility for smart grids. Technical report, Flexiness D8.1, 2010.

[17] Tennet. Eindrapportage systeemintegratie windvermogen. Technical Report MR 05-373, Tennet, november 2005.

[18] Tali Trigg and Paul (eds.) Telleen. Global EV outlook, understanding the electric vehicle landscape to 2020. *International Energy Agency*, 2013.

[19] Michael Wooldridge and Nicholas Jennings. Intelligent agents: Theory and practice. *Knowledge engineering review*, 10(2):115–152, 1995.