A Comparison of Algorithms for Controlling DSRs in a Control by Price Context Using Hardware-in-the-loop Simulation

Fabrizio Sossan, Henrik Bindner

***Abstract*—With future increasing of electric energy production from fluctuating sources, the need for regulating power will rise and conventional power plants - that today provide all power system ancillary services - could not have the capability and the flexibility of providing it. Demand Side Resource, DSRs, are electric loads whose power consumption can be shifted without having a big impact on the primary services they are supplying and they are suitable for being controlled according the needs of regulating power in the electric power system. In this paper the performances and the aggregate responses provided by three algorithms for controlling electric space heating through a broadcasted *price signal* are compared. The algorithms have been tested in a software platform with a population of buildings using a *hardware-in-the-loop* approach that allows to feedback into the simulation the thermal response of a real office building; the experimental results of using a model predictive controller for heating a real building in a variable price context are also presented.**

This study is part of the Flexpower project whose aim is investigating the possibility of creating an electric market for regulating power with a big participation of DSRs and small scale generation units.

***Index Terms*—control by price, Demand Side Resources, Smart grids.**

I. Introduction

Demand side resources, or DSRs, are electric loads that provide services that are naturally coupled to some kind of storage; this allows to control, schedule or shift their power consumption without having a big direct impact on the quality of the primary services they are providing to the users. Examples of demand side resources are space or water heating, electric vehicles or also fully deferrable load such as washers or dishwashers. Because of their flexibility, DSRs are suitable to be controlled in order to contribute to power system services with respect to their constraints, physical limits and local settings [1]. DERs exploitation is based on the consideration that the contribution from the single unit is small but the aggregate response of a big number of devices might be relevant. Demand side resources can be directly controlled (for example by a power reference signal), they can react in order to response to a deviation of the grid frequency or, a time shift in their electric power usage can be achieved using a price signal for the electric energy [2][3]: this economic incentive should induce the demand side user to consume more power when the energy is cheap (and in the case to store it) and to reduce the consumption when the price for the energy is high. Users response to price signal is spontaneous and based on local comfort or operational preferences. Using price signal, is a convenient form of control since the decision is computed locally while in the case of direct control the information (local conditions for example) should be propagated from each devices up to some aggregator which should send a control signal to drive each of them. The critical part of a control-by-price approach is the price signal itself because, for producing it, the response of the distributed demand side resources should be known since it is required in the market process[4]. In Flexpower project, the aim is to use control by price approach for supporting the amount of regulating power needed by the power system [5].

In this paper the performances of three algorithms with price responsiveness capability for controlling domestic electric space heating through a broadcasted price signal are compared. The proposed algorithms have different level of requirements: two of them work basing their decision on historical data of the price signal; the third one is a model predictive controller (*MPC*) and it uses both prices and real weather forecast.

In this comparison, the price signal is artificial and it is built with the aim of highlighting the pros and cons of each control algorithm. The algorithms have been compared adopting both a final user point of view, evaluating the deviation from the optimal comfort level (the indoor temperature set-point in this case) and the total cost for the energy used, and the power system one; from the power system perspective, it would be useful having a devices with a good price responsiveness capability, able to reduce or to decrease the power consumption steadily for a long time and predictable in the behavior.

The control algorithms have been tested using a Java simulation platform with a population of houses. For introducing more realism into the simulation, a *hardware-in-the-loop* approach has been used: the thermal dynamics of a real building have been introduced in the simulation and they are used for perturbing in real time the behaviors of all the houses inside the population.

Detailed descriptions about the simulation platform, the hardware in the loop feedback, the thermal models, the control algorithms and how their performances have been compared are given in Section II.

Section III is for presenting and discussing the results of the simulations: both the aggregate response of the population of buildings and the experimental result of the model predictive controller applied on a real house are reported. A comparison of the performances obtained with the different control algorithms is then presented.

Section IV is for conclusions.

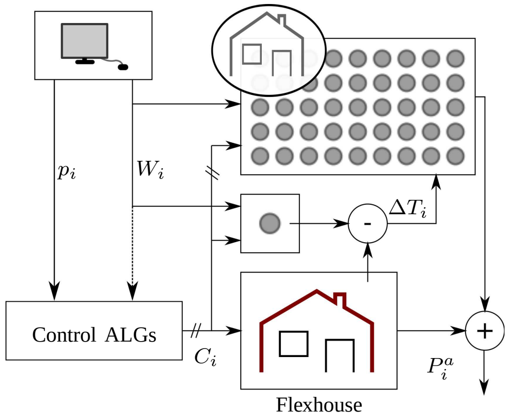

Fig. 1. The experimental setup used for the proposed simulations with the FlexHouse building as the *hardware-in-the-loop* for reproducing the real environmental conditions. Weather forecast W_i and the price p_i signal are used.

II. Methods

The software tool used for performing the simulations is a Java based dynamic simulator which allows to implement generic models and control algorithms. The simulation platform allows for simulating a large number of units and for having different controllers implemented concurrently; the functionalities of the simulation platform have been integrated with the distributed power system facility SYSLAB[6] in order to perform a real time hardware in the loop simulation introducing the thermal dynamics of a real building, FlexHouse [7]. FlexHouse is a small office building which is heated by ten 1kW electric heaters that can be controlled.

The diagram of the simulation scenario is shown in Figure 1. The box with the small gray circles represents the population of houses that is implemented in the simulator. The disturbation signal ΔT_i, which is obtained comparing the indoor temperature of FlexHouse with its model implemented in the simulator, perturbs the behavior of all the simulated buildings in order to reproduce the effect of the real environment and operating conditions (uncertainties in the modeling and in weather forecast). Of course environmental local conditions like wind, outside temperature and solar radiation, act on FlexHouse.

The output of each implemented control algorithm is a signal that drives the electric heaters. In the case of FlexHouse, the same control signal is applied both to the real building and to the model inside the simulation platform in order to compare the real behavior with the simulated one.

The output of the simulation is the temperature profile and the power usage of each building: the sum of them gives the aggregate electric power consumption. Also the experimental data regarding FlexHouse activity are made available by the simulation platform.

In this experimental setup, the price signal p_i is simulated and it is built for trying to demonstrate which algorithm performs better within this framework. In accordance with Flexpower project, the price for the electric energy is updated every five minutes.

Weather forecast W_i for the area is delivered by a FTP service; the weather forecast is used both by the model predictive controller and for evolving the state of thermal models. Forecast is released every day and with a time resolution of one hour: smaller time steps are required for computing the models evolution so the same resolution of the price signal, five minutes, is achieved simply using linear interpolation.

Despite local conditions measurements, such as outdoor temperature or solar radiation, are available in SYSLAB, they are not used by the implemented control algorithms because it would not be realistic in an future diffusion of the control by price since it would increase the complexity and cost of the system for the many sensors to add around.

A. *Thermal models*

Reference [8] present several linear thermal models for FlexHouse. They are reported in increasing order of complexity (states number) and they are built using a grey-box approach where the parameters are computed using a maximum likelihood estimator with real measurements from FlexHouse. The one selected for this implementation is the simplest one, a 1-state linear model. This model is used both for describing FlexHouse dynamics and the population of buildings; for this last purpose, the parameters of the model have been slightly varied following a normal distribution for taking into account small differences in size and thermal conductivity features.

The population of the buildings is composed by four identical groups of fifty houses each. Each group of buildings is driven by a different type of controller (presented later) and the temperature evolution of each model is perturbed at each step by the measured input discussed before.

For a better description of the thermal behavior of a building, a two states model could be used for representing the transients both of the air and of the building envelope. The reason why a simpler 1-state model has been chosen is that FlexHouse only disposes of the indoor air temperature information so this is the only measure available for initializing the model (a state estimator could be considered in the case of linear models with larger order). The hardware in the loop approach provides to the simulated population of house an even more realistic footprint to the thermal transient. The problem still remains for the MPC, where the 1-state model could not give a good prediction on the future indoor temperature during transient (the steady state response is the same) resulting in a worse performance of the controller.

B. *Control algorithms*

The control algorithms for whom the performance is compared have different levels of complexity and requirements.

The simplest one is the traditional thermostatic controller and it has been included mainly for having a comparison with the current situation.

The second algorithm is a simple extension of the traditional thermostat where a price response capability has been added: in the temperature hysteresis cycle, the controller can choose to switch on the heating for storing thermal energy if the proposed price for the electric energy is cheaper than the prices paid in average (in the past) or otherwise to shutdown it.

The third control algorithm is proposed by [9] and [10]; again, the decision process is based upon considerations on the historical prices. At each time i, when a new price for electric energy arrives, the temperature set-point for the thermostatic controller is modified from the optimal one using the quantity ΔT_i defined in the Equation 1:

$$\Delta T_i = -k \; \hat{p}_i \tag{1}$$

$$\hat{p}_i = \frac{p_i - \overline{p}_i}{\sigma_i} \tag{2}$$

where p_i, $\overline{p}_i$ and σ_i in Equation 2 are respectively the current price, the moving average and standard deviation of the price history for the last 24 hours. Coefficient k is a positive constant and sets the price responsiveness capability of the controller. $\hat{p}_i$ is called *relative price*. The result achieved by this controller is to produce a temperature deviation ΔT proportional to the relative price and in general it is negative when the price is greater than the average of the old ones: so the new temperature set-point will be smaller than the previous one and vice-versa.

The fourth and last controller uses a thermal model for computing the future temperature states of the house and it minimizes a cost expression in the form of Equation 3 with finding an appropriate heating power profile; the target of the algorithm is to reduce future deviations from the temperature set-point paying as little as possible for the electric energy.

$$J = k_1 q(T(t_f)) + \int_0^{t_f} k_2 q(T(t)) + r(u(t))\overline{p}(t) \; dt \tag{3}$$

The symbols T and u, both time dependent, are for indoor temperature and heating power respectively. The integration interval starts from the current instant, 0, until the length of optimization horizon t_f. The functions $q(T)$ and $r(u)$ are numerical weighting barrier functions that assume high values when the independent variable approaches non admissible values. Their shape is shown in Figure 2. Function $\overline{p}(t)$ is the price signal that, multiplied by the electric power used by the heaters and integrated in time, gives a energy cost. Since the integral looks ahead in time, $\overline{p}(t)$ is a price forecast series. As said before, the price signal is artificial and, in this approach, it is assumed to know the price forecast without error. The factor k_2 determines the price responsiveness capability of the controller. The goal of the optimization process is finding the optimal control law $u^o(t)$ that is able to minimize the cost expression J of Equation 3. For finding the minimizing control $u^o(t)$, a numeric algorithm based on the gradient descent method[11] has been used and it is applied each time a new price for the electric energy is released (five minutes). The indoor temperature T should be treated as stochastic variable but in this approach the problem formulation is left deterministic since the receding horizon configuration brings a natural feedback into the system because the algorithm retrieves a new indoor temperature reading every five minutes, together with the price signal and weather forecast. A similar approach for a MPC but using linear programming is in [12] and [13].

Both end user and power system point of views have been taken into account in the comparison of the performances of the different algorithms. The following three indicators have been evaluated.

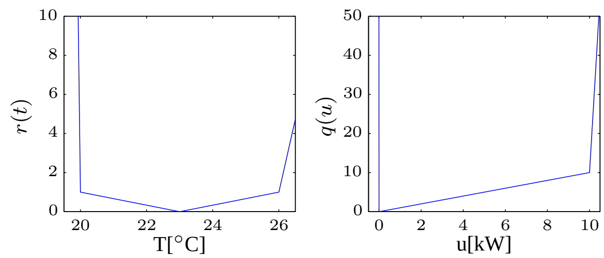

Fig. 2. The weighting functions used in the cost expression of the MPC.

1) temperature comfort penalty function: it is computed as the integral in time of the aboslute value of the temperature deviation ΔT from the set-point (chosen as $23°C$ for all the population of simulated buildings and FlexHouse). States of under-temperature have been weighted 1.2 times more than over-temperature conditions; a second temperature penalty function is considered and it is the integral in time of the temperature deviation when it is out of the admissible comfort interval (defined as $23 \pm 3°C$).

2) total cost of energy: it is simply the integral in time of the electric power used for heating multiplied by the price signal;

3) price responsiveness capability: it is defined as the variation of heating power produced by a control algorithm with respect to price variation that is dP/dp with P power consumption.

The first two indicators take into account the user comfort and economic benefit, while the last one is important from power system perspective since it gives a measurement of the control capability that each algorithm adds, in this case, to domestic electric heating. The second indicator, the total cost for the energy, gives also a measurement of the capability of the algorithm to move electric power consumption in instants of time with lower energy prices.

The number of on-off cycles produced by the algorithms has not been take into account in this analysis; anyway this is a relevant aspect if electromechanical devices for powering heaters are used, since their life is heavily affected by the number of switching cycles.

III. Results and Discussion

As introduced before, the simulation here presented is obtained applying the heating power setting computed by the predictive controller to both a real building (FlexHouse) and its model. The difference between the two responses is then used for perturbing the state of all the other buildings inside a real time simulation, where the behavior of 200 houses (divided in four identical groups of 50 buildings with each group driven by one of the four algorithms discussed before) is reproduced. The temperature set-point for all the buildings is $23°C$ and the allowed offset is $\pm 3°C$. An artificial new price is broadcasted every five minutes together with real weather forecast. The price signal is a sine wave with decreasing frequency. Plots in Figure 3 and 4 are split in two parts for

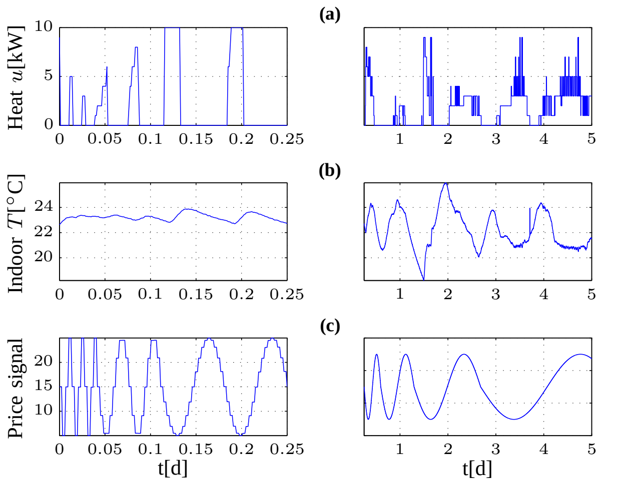

Fig. 3. The experimental data from FlexHouse with the model predictive controller. Plot *a* shows the heating power, *b* is the indoor temperature and *c* is the artificial price signal.

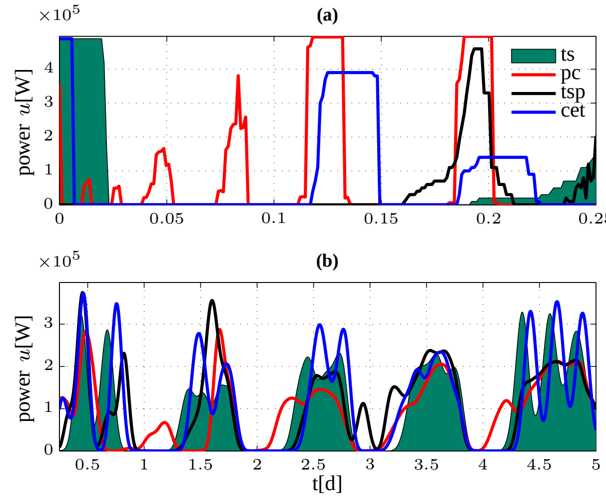

Fig. 4. A comparison of the output of the four different controllers. The two plots are consecutive and with different time scale. Data in the plot *b* have been filtered with a low pass filter for a better visualization of the main frequency components.

convenience of visualization; the former part is for the high frequency components of the price signal, the latter is for the consecutive rest of time where price signal presents slower variations. The optimization horizon length of the predictive controller is five hours and it has been chosen according the thermal time constant of FlexHouse; as said before, MPC uses real weather forecast and the price signal forecast is supposed known without error.

Time $t = 0$ in the plots refers to 18:16 pm of October the 13^{rd} 2011 and the simulation lasts for 5.2 days.

Since the price signal is artificial and it was created for enabling the comparison between algorithms, its unit of measure on the plots is not specified and it is intended as a generic monetary value for unit of energy.

The firsts plots in Figure 3 show the experimental results of the application of the model predictive controller to FlexHouse. The attempt of the controller to move the electric energy consumption when the price is low is visible. With the decreasing of the frequency of the price signal, the controller starts to lose the capability of shifting the energy usage and the power is dictated by the indoor temperature because the house does not have enough thermal inertia to maintain the indoor temperature in a acceptable range.

Figure 4 shows the aggregate response of the four different groups of buildings; each of them is driven by a different controller according to the plot legend: *ts* stays for thermostatic controller, *tsp* is for second algorithm discussed in Section II, *cet* is the third and *pc* is the model predictive controller. The data in the bottom graph have been elaborated with a low pass filter for a better visualization. As Figure 4 shows, the output produced by each controller with price responsiveness capability is very different from the green area that is the thermostat action; this means that they are all exploiting the flexibility offered by the price signal but in different ways. Again, with the increasing of the period of the price signal, all the controllers tend to show a behavior more and more similar to the thermostatic controller action because the thermal inertia is not enough for both being able respond to slow price variation and maintaining the indoor temperature in an acceptable comfort level.

Figure 5 shows the average temperature profiles for the four different groups of buildings. Plot *a* shows the instant deviation from the indoor temperature set-point ($23°C$). Plot *b* shows the accumulated error computed as the integral of the absolute value of the deviation, where the states of under temperature have been weighted 20% more than states of over temperature. Plot *c* shows the accumulated absolute value of the error when the temperature is not in the admissible range $23 \pm 3°C$. Plot *b* shows that the thermostatic controller is the more precise in keeping the temperature close to the optimal set-point; plot *c* shows that the model predictive controller is the one that performs best in keeping the temperature of the buildings in the defined acceptable range. It is worth to notice the behavior of the two thermostatic based controllers (black and green profiles): in fact they always should be able to maintain the temperature in the thermostatic interval but, in the days of the simulation, the weather was cold ($7/12°C$) and sunny in the Copenhagen area and during the day hours the sun was able to warm the buildings in a significative way and more than the allowed range. This explain why the thermostatic based controllers have this accumulated error different than zero; in this case the model predictive controller takes advantage of the weather forecast and it is able to perform even better than them. The control algorithm that corresponds to the blue line does not show a good performance in plot *c* because the way the temperature offset is computed (Equation 1) can easily lead to temperature references quite different from the real desired optimal one.

Figure 6 shows the cost of the energy for the different control algorithms. Plot *a* is the instant energy bill and plot *b* represents the average on the population amount of money that the final user would save (in a variable price context of course) applying the different algorithms with respect to

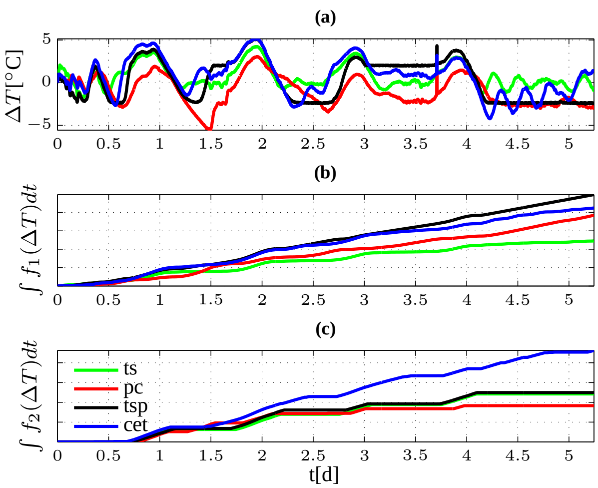

Fig. 5. Plot *a* shows the temperature deviation from the set-point (23°C) for the different algorithms. Plot *a* is the accumulated error (with states of under temperature have been weighted 1.2 times more). Plot *c* is the accumulated error for the temperature deviation from the comfort area that is 23 ± 3°C.

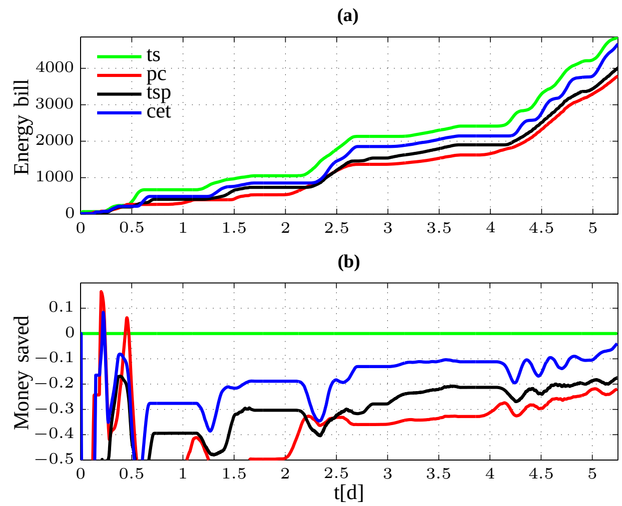

Fig. 6. Plot *a* shows the instant values of the electric bill. Plot *b* shows the money saved with respect to the traditional thermostatic controller in a price signal context.

using traditional thermostatic controller. Plots in Figure 6 show that the model predictive controller is the algorithm that performs best together with the very simple modified thermostatic controller. It is worth to notice that, despite weather forecast are real, the price signal is assumed known and so in the case of the energy bill the performance of the model predictive controller could be worse if wrong forecast are provided. Anyway, with good price forecast the MPC controller is the algorithm supposed to still give the best result because it is the only one that, using forecasts, takes advantage of looking into the future for a better utilisation of the thermal energy stored in the buildings. So, for example, if the price forecast are indicating an increment in the future energy price, MPC algorithm is the only one that is able to start to store thermal energy at some point (with switching the heating on) in order to avoid to pay more money for the same energy amount later.

Figure 7 shows the price responsiveness for the proposed algorithms. Price responsiveness is defined as $\frac{dP}{dp}$, that is the variation of electric power usage achieved by an algorithm given a certain variation of the price of the energy. Each point on the plot is the mean of a distribution of values that are the different responses that the algorithm gives to such variation of price. The standard deviation σ for each distribution is reported on each mean value with the vertical bar on the plot: how the bars show, the aggregate response is not really predictable and may vary in a very significative way because it is function of a wide number of parameters, such as time, outside temperature, wind speed, solar radiation, thermal features of the buildings, thermal history and so on. Anyway the contour lines in the plot of Figure 7, that are the fourth grade polynomials obtained applying least square method to the mean values of the distributions, show a clear tendency we could expect about the behavior of the algorithms that is that the controllers with price responsiveness capability are able to reduce the power consumption when there is a positive variation of the price and vice-versa; as the flat green lines shows, the thermostatic controller does not offer price responsiveness at all. The algorithms that show the best price responsiveness capabilities are the model predictive controller and the one named *cet*.

With computing the average of the standard deviation of the populations along the x-axis, it is possible to get an idea of which algorithm would result in the best predictability of the price response for this simulation: the controller with the less standard deviation is *cet*, then *tsp* and finally the model predictive controller.

An interesting point of the price responsiveness of Figure 7 is to see if increasing the number of buildings of one order of magnitude could improve its shape in the sense of having smaller values for the standard deviations and, so, increasing the predictability of the behaviors of the buildings; in fact these curves could be used for estimating the prediction of power usage variation given a electric energy price and they could be used for bidding in the electric market for computing a price signal.

IV. Conclusion

In this paper the performances of three different algorithms suitable for controlling domestic electric space heating through price signal are compared. Control algorithms have been compared analyzing the aggregate response of a population of buildings inside a simulation platform using a *hardware-in-the-loop* approach: a temperature feedback signal from a real office building perturbs in real time the behaviors of all the buildings simulated in the platform in order to take into account environmental uncertainties and so pretend more realism from the simulation. Besides, the experimental results of using a model predictive controller on a real house have been presented. Real weather forecast for the area are used for the model predictive controller. The price signal used in the simulation is artificial and created for enabling the comparison between the different algorithms.

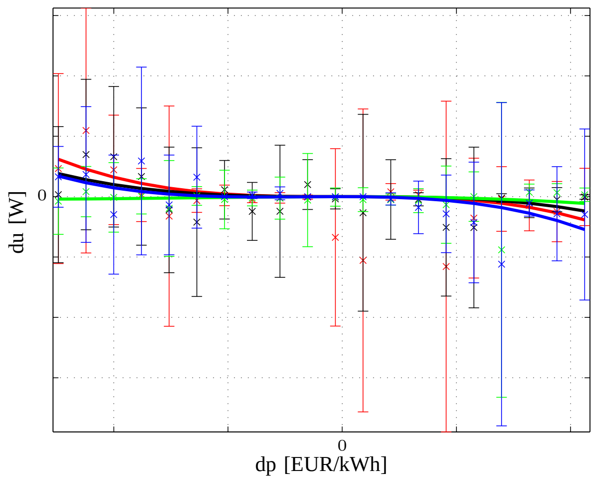

Fig. 7. The price responsiveness for the four different algorithms, defined as the variation in the power consumption as a function of a variation in the price; vertical bars show the standard deviations for the different responses.

Control algorithms are compared taking into account the comfort level they can achieve (measured as the deviation from the optimal indoor temperature), the total cost for the energy they require and the price responsiveness capability, defined as the change in the electric power usage when a variation in the electric energy price is proposed.

All the control algorithms here presented show a positive capability of moving the energy consumption from moments where the energy is expensive to cheaper energy time frames (Figure 7) and maintaining a reasonable comfort level in the buildings. All the algorithms start to lose effectiveness when the frequency of the price signal decreases; this because the thermal inertia of the buildings is not big enough to maintain the indoor temperature in the allowed range.

Anyway the algorithm that shows the best global performance is the model predictive controller; in the case of the indoor temperature comfort level, it behaves even better than the thermostatic based controllers (Figure 5.c) because, taking advantage of weather forecast and the thermal model, can predict the effects of solar radiation (that in the days of the experimental simulation was important) and take appropriate counteractions. MPC is also able to produce the cheapest energy bill if compared with the other two algorithms. In the case of the proposed simulation, we assume perfect foresight for the price signal and this would not be the case in a real application where errors in forecasting are expected. Anyway in presence of good forecast for the price signal, the model predictive controller should still be the best one also in terms of energy cost since it is the only one that can look into the future and, for example, start to store energy if it detects a future increase in the price; this allows a better management of the energy stored in the buildings.

From the point of view of the infrastructure, model predictive controller does not have more hardware requirements than other controllers since they all need to receive a price signal and some kind of logic to implement the algorithm; once the communication and hardware requirements for enabling the transmission of price signal are set up, weather forecast and price forecast are services that can be built on the top of the system without having impact on the complexity of the hardware architecture. The problem of how to compute good price forecast still remains since the aggregate response of DSRs is not perfectly predictable, time variant and dependent on many variables such as local conditions.

Acknowledgment

This work is supported by Flexpower project[5] (ForskEl 2010-1-10486).

References

[1] Z. Xu et al., "Design and Modelling of Thermostatically Controlled Loads as Frequency Controlled Reserve", Power Engineering Society General Meeting IEEE, 2007.
[2] Schweppe, F.C. et al., "Homeostatic Utility Control," Power Apparatus and Systems, IEEE Transactions on , vol.PAS-99, no.3, 1980.
[3] Alvarado, F.L.; , "Is system control entirely by price feasible?", Proceedings of the 36th Annual Hawaii International Conference, 2003.
[4] N. Lu et al., "Simulating price responsive distributed resources," Power Systems Conference and Exposition, 2004. IEEE PES, 2004.
[5] http://www.ea-energianalyse.dk/projects-english/1027_flexpower_market_design.html
[6] O. Gehrke, "Self-Organising Distributed Control of a Distributed Energy System with a High Penetration of Renewable Energy", PhD thesis, 2010.
[7] Yi Zong et al., "Active load management in an intelligent building using model predictive control strategy," PowerTech, IEEE Trondheim, 2011
[8] P. Bacher et al., *Models for Energy Performance Analysis*, IMM Technical Report, 2010.
[9] D. Hammerstrom et al., Pacific Nortwest GridWise Testbed demonstration Projects, Part I. Olympic Penninsula Project National Technical Iformation Service, U.S. Dept. Commerce, 2007.
[10] P. Nyeng et al., "Information and Communications Systems for Control-by-Price of Distributed Energy Resources and Flexible Demand," Smart Grid, IEEE Transactions on , vol.2, no.2, pp.334-341, 2011.
[11] R.F. Stengel, Optimal control and estimation, Dover publications, 1993.
[12] Yi Zong et al., "Model Predictive Control Strategy for a Load Management Research Facility in the Distributed Power System with High Wind Penetration -Towards a Danish Power System with 50% Wind Penetration," Power and Energy Engineering Conference (APPEEC), 2011.
[13] T.G. Hovgaard et al., "Energy Efficient Refrigeration and Flexible Power Consumption in a Smart Grid", Risø International Energy Conference 2011.

Fabrizio Sossan was born on 9th April 1985 in Genova, Italy. He received his master in Robotic and Control Engineering from University of Genova in 2009. From December 2010 he is a PhD student at IES Risø DTU, Denmark; he deals with integration of distributed energy resources in electric power systems with a high share of renewable energies. Email: faso@risoe.dtu.dk

Henrik Bindner was born on 30 June 1964 in Copenhagen, Denmark. He received his master in electrical engineering from the Technical University of Denmark in 1988. Since 1990 he has been with Risø DTU National Laboratory for Sustainable Energy in the Wind Energy Division, currently as a Senior Scientist. He has mainly been working on integration of wind energy into power system. The work has included analysis, design and control of small island systems as well as technologies and techniques for integration of wind in large systems. Currently, he is researching how distributed energy resources can be applied to increase penetration of wind energy. Email: hwbi@risoe.dtu.dk.
Risø DTU, Frederiksborgvej 399, Roskilde 4000, Denmark, http://risoe.dtu.dk/