

Thermal Storage Power Balancing with Model Predictive Control

Rasmus Halvgaard, Niels K. Poulsen, Henrik Madsen and John B. Jørgensen

Abstract—The method described in this paper balances power production and consumption with a large number of thermal loads. Linear controllers are used for the loads to track a temperature set point, while Model Predictive Control (MPC) and model estimation of the load behavior are used for coordination. The total power consumption of all loads is controlled indirectly through a real-time price. The MPC incorporates forecasts of the power production and disturbances that influence the loads, e.g. time-varying weather forecasts, in order to react ahead of time. A simulation scenario demonstrates that the method allows for the integration of flexible thermal loads in a smart energy system in which consumption follows the changing production.

I. INTRODUCTION

Integration of large amounts of renewable energy sources in the power system, such as wind and solar energy, introduces large fluctuations in power production. This type of green energy must be either stored or consumed right away. Consuming all of it as it is produced requires a very flexible and controllable power consumption. Thermal loads, in particular, consume power and often have flexible operating temperatures and thermal storage capacity. Examples of controllable electric thermal loads are heat pumps in buildings [1], auxiliary heating in solar collector storage tanks [2], and commercial and domestic refrigeration systems [3]. In a smart energy system these loads can potentially offer flexibility if they are pooled together into a large-scale system with large power consumption. With the right control scheme this large-scale system of flexible thermal loads can help balance changing power production levels by adjusting the consumption of the loads accordingly [4]. However, an incentive to help balance the power and a method for coordinating must be established.

In this paper an indirect control strategy is proposed where a control signal, referred to as a control price, communicates the need for balancing. The control price is linearly linked to the temperature set points and therefore indirectly influences the total power consumption of a group of thermal loads. This group is often referred to as an *aggregation* of loads, and all loads are connected to an *aggregator* [5]. The aggregator broadcasts the current control price, which is translated by each load individually into a local temperature set point to be followed. Based on a model of the aggregated consumption response to the control price, closed-loop feedback is provided at the aggregator level by measuring the total

power consumption. In this way the aggregator is able to balance instantaneous power or track an amount of power already bought from a market [6]. In this paper the aggregator controller is based on Model Predictive Control (MPC) [7]. The MPC uses an estimated low-order autoregressive (ARX) model for real-time power balancing. Moreover, an integrator model is added to eliminate model and forecast errors and to achieve offset-free tracking.

For control, the MPC needs a model of the aggregated thermal loads. This model should predict the effect of a price change and calculate a single control price, which is broadcast to all loads. The aggregator model is estimated from the price response and may be very small compared to a centralized model that includes detailed information about all loads.

Compared to a centralized direct control strategy, the decentralized indirect method described in this paper reduces the aggregator problem complexity considerably. The computation efforts are decreased dramatically and the need for two-way communication is eliminated. The relationship between control price and set point in this paper was inspired by [8]. A similar concept of balancing is found in [9], where simple hysteresis control is used. In [10] an indirect price strategy based on bilevel programming and a large centralized model is proposed to minimize power imbalances accounting for the load's response to the price signal. An example of a centralized direct control strategy can be found in [11]. Note that most centralized formulations can be solved more efficiently through decomposition of the optimization problem into smaller subproblems. However, two-way communication is still needed for coordination and as the number of loads increase a decentralized approach is needed. A completely decentralized approach, where optimization variables are exchanged between loads as dynamic prices, is considered in [12].

This paper is organized as follows. In Section II we formulate an aggregated model of a large-scale system of thermal loads. Section III describes the MPC that controls the aggregated loads. In Section IV the control method is demonstrated through simulation. The *control price* concept is discussed further in Section V, while Section VI provides conclusions.

II. MODELING

First we model the dynamics of the thermal loads and their closed loop behavior with Linear Quadratic (LQ) controllers. Then connection to the aggregator and the estimated aggregated model is described. For notational simplicity,

*This work was partly supported by the DSF iPower (Strategic Platform for Innovation and Research in Intelligent Power) project.

R. Halvgaard, N. K. Poulsen, H. Madsen, and J. B. Jørgensen are with Department of Applied Mathematics and Computer Science, Technical University of Denmark, Matematiktorvet, 2800 Kgs. Lyngby, Denmark {rhal, nkpo, hmad, jbjjo}@dtu.dk

the discrete time step subscript k has been omitted in the following while the superscript $+$ denotes $k + 1$.

A. Thermal load

Each thermal load is modeled by the discrete-time state space model

$$x^+ = Ax + Bu + Ed + w \quad (1a)$$

$$y = Cx \quad (1b)$$

y is the output temperature, u is the power consumption, d and $w \in N(0, \sigma_w^2)$ are disturbances influencing the states x . The disturbance d could be an outdoor temperature acting on a building, or solar radiation, while w is the unmodeled process noise. For the method described in this paper we design LQ controllers to track a temperature set point r . When choosing linear controllers, the aggregated model is also linear and allows for a linear MPC at the aggregator level. The unconstrained LQ controller should be able to track the set point with no offset by rejecting any disturbances. Offset-free integral control is achieved by augmenting the state vector x with an integrating state \tilde{x} , such that $x_i = [x^T \ \tilde{x}^T]^T$. From now on the different loads are subscripted i to denote the L different loads. With integral control the i th load is then modeled by

$$x_i^+ = A_i x_i + B_i u_i + F_i r_i + E_i d_i + G_i w_i \quad (2a)$$

$$y_i = C_i x_i \quad (2b)$$

with

$$A_i = \begin{bmatrix} A & 0 \\ -C & I \end{bmatrix} \quad B_i = \begin{bmatrix} B \\ 0 \end{bmatrix} \quad C_i = [C \ 0] \quad (3a)$$

$$F_i = \begin{bmatrix} 0 \\ I \end{bmatrix} \quad E_i = \begin{bmatrix} E \\ 0 \end{bmatrix} \quad G_i = \begin{bmatrix} I \\ 0 \end{bmatrix} \quad (3b)$$

The following linear control law is applied to track the temperature set point r_i

$$u_i = -K_i x_i \quad K_i = [K \ -\tilde{K}] \quad (4)$$

A stationary control gain K_i has been designed for each load with the weights $Q_i \geq 0$ and $R_i \geq 0$ on the states x_i and control action u_i , respectively. All loads are assumed stable and controllable. The assumption of full-state feedback is justified by the use of SISO models later in the numerical example in Section IV. Alternatively, a Kalman filter could be applied to estimate any unmeasured states. The controller weights should be tuned separately for each load to trade off long settling times for temperature overshoot and power consumption.

As the aggregator objective is to manipulate power consumption indirectly through the set point, the relationship between u_i and r_i must be modeled. In our case a linear expression for power consumption is readily available from the control law (4). Inserting (4) in (2) gives us the closed-loop model

$$x_i^+ = (A_i - B_i K_i) x_i + F_i f_i(p) + E_i d_i + G_i w_i \quad (5a)$$

$$z_i = u_i = -K_i x_i \quad (5b)$$




Fig. 1. Load i with system $l_i : (A_i, B_i, C_i, E_i)$ and LQ integral controller.




Fig. 2. Function $f(p)$ from control price to temperature set point

The aggregator measures the power consumption and not the temperature. Therefore, power consumption is now defined as the model output z_i from the i th load. y_i still indicates the temperature output. In (5) the temperature set point r_i has been replaced by a function f_i with the aggregator control price p as argument. The control price is a scalar that is broadcast to all loads, reflecting the need for balancing. A block diagram of the controlled load is shown in Fig. 1. Here it is seen how the control price is added as input to the closed-loop model. Each load must map the control price to an individual temperature set point. This mapping is done by the affine function $f_i(p)$ defined for each load

$$r_i = f_i(p) = -\frac{\bar{r}_i - \underline{r}_i}{\bar{p} - \underline{p}}(p - \underline{p}) + \bar{r}_i \quad (6)$$

When $(\underline{p}, \bar{p}) = (-1, 1)$ and $(\underline{r}_i, \bar{r}_i) = (b_i - a_i, b_i + a_i)$, (6) reduces to

$$f_i(p) = -a_i p + b_i \quad (7)$$

When the price is constrained the function $f_i(p)$ also constrains the temperature set point to a certain interval defined by a_i and b_i . This mapping is illustrated in Fig. 2 and is key to understanding the role of the control price. Note that for cooling systems the sign on the slope a will be chosen opposite of (7).

B. Aggregated model

We can put all the closed-loop models from (5) together to get a large linear model of all L loads. This augmented state space model subscripted a is then

$$x_a^+ = A_a x_a + B_a(p) + E_a d + G_a w_a \quad (8a)$$

$$z_a = C_a x_a \quad (8b)$$

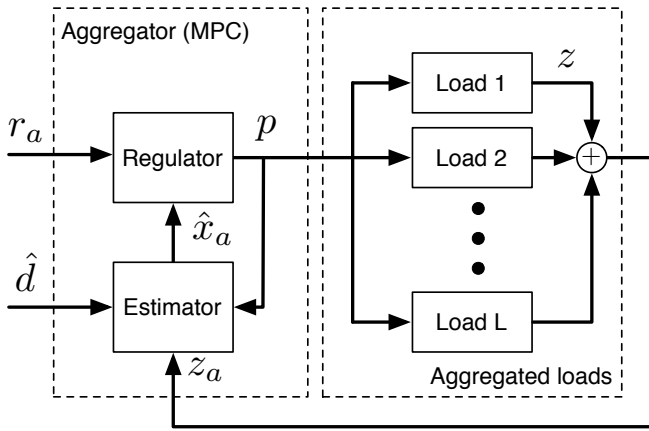


Fig. 3. System overview of aggregator and loads.

with $B_a(p) = F_a f_a(p)$ and

$$\begin{aligned}
 A_a &= \text{blkdiag}(A_1 - B_1 K_1, A_2 - B_2 K_2, \dots, A_L - B_L K_L) \\
 F_a &= \text{blkdiag}(F_1, F_2, \dots, F_L) \\
 G_a &= \text{blkdiag}(G_1, G_2, \dots, G_L) \\
 f_a(p) &= [f_1(p) \quad f_2(p) \quad \dots \quad f_L(p)]^T \\
 x_a &= [x_1^T \quad x_2^T \quad \dots \quad x_L^T]^T \\
 w_a &= [w_1^T \quad w_2^T \quad \dots \quad w_L^T]^T \\
 E_a &= [E_1^T \quad E_2^T \quad \dots \quad E_L^T]^T \\
 C_a &= -[K_1^T \quad K_2^T \quad \dots \quad K_L^T] \cdot \mathbf{1}
 \end{aligned}$$

The derived closed-loop model (8) of all loads describes the aggregated response from price to power consumption. The desired SISO model with $z_a \in \mathbb{R}^1$ and $p \in \mathbb{R}^1$ is here formed by the output matrix C_a , which sums all power consumption contributions $z_a = \sum_{i=1}^L z_i$. The disturbance, $d \in \mathbb{R}^1$, is assumed from now on to be a scalar influencing all loads. It is assumed that the main component of this disturbance can be forecast, while the remaining tracking errors from disturbances are assumed to be eliminated by the MPC integral controller.

As the number of loads increase, so does model complexity and, ultimately, controller computation time. However, this high dimensional model can be well approximated by a lower-order model [13], [14]. In our method, we reduce the model by estimating a low-order AR model from the simulated response

$$A_{ar}(q^{-1})\hat{z}_a = B_{ar}(q^{-1})p + E_{ar}(q^{-1})\hat{d} + \eta \quad (9)$$

η is the unmodeled disturbances. The model (9) can be used for model estimation and control. It is assumed to be accurate enough to enable the aggregator MPC to eliminate model mismatch errors through an observer and stable closed-loop feedback. Forecasts of the load disturbance \hat{d} is also added to the model with the term E_{ar} .

III. AGGREGATOR CONTROLLER

The MPC is well suited for control at the aggregator level due to the following reasons. It handles capacity constraints

indirectly through a limit on the price. It rejects disturbances and is able to track the power consumption reference r_a with a small error, since r_a is known ahead of time so the MPC can react in advance. In practice, r_a could be a time-varying forecast of wind power production, and at every time step the MPC takes continuously updated forecasts into account.

A. Aggregator objective

The method presented in this paper will indirectly change the power consumption of all thermal loads through a price that is linearly related to the temperature set points. In this way the aggregator puts a price on heating or cooling, and indirectly on electricity as well. The set points will be set at a high temperature set point when the price is low, and at a low temperature when the price is high. However, the interval within which the temperature set point is allowed to vary and is up to the individual load, e.g. it could be the temperature comfort interval in a building heated by a heat pump [1]. The temperature interval could even be set at zero by setting $a_i = 0$ in (7), but then the aggregator would have no flexibility to exploit. Note that the same method holds for refrigeration systems. However, in this case the set points should set at a low temperature when the price is low. Controlling the loads through a price requires a model of the price response as well as models of the thermal load behavior. The loads are connected to the aggregator through a control price p as shown in Fig. 3. The total power consumption of the loads z_a is measured by the aggregator that estimates an aggregated model and provides closed-loop feedback with an MPC for tracking the power consumption reference r_a .

B. Offset-free ARX MPC

We assume the model (8) to be estimated from data as an ARX model on the form (9). To obtain offset-free tracking we replace the unmodeled term η by an integrator model [15]

$$\eta = \frac{1 - \alpha q^{-1}}{1 - q^{-1}} e \quad (10)$$

α is a tuning parameter [15]. The observer error $e = z_a - \hat{z}_a$ is obtained from measurements of the aggregated response z_a (8). Adding the integrator model (10) to (9) yields the controller model in ARMAX form

$$A_{c,ar} = (1 - q^{-1})A_{ar} \quad E_{c,ar} = (1 - q^{-1})E_{ar} \quad (11a)$$

$$B_{c,ar} = (1 - q^{-1})B_{ar} \quad C_{c,ar} = 1 - \alpha q^{-1} \quad (11b)$$

The final controller model used as a predictor is obtained by realizing (11) as a discrete state-space model in innovation form

$$\hat{x}_a^+ = A_c \hat{x}_a + B_c p + E_c \hat{d} + K_c e \quad (12a)$$

$$\hat{z}_a = C_c \hat{x}_a \quad (12b)$$

This is the one-step predictor. For predicting j -steps ahead the term $K_c e$ is omitted.

Algorithm 1 MPC algorithm

For every time step find the optimal control price p_k^*
 Require $\mathcal{P} = (\hat{x}_{a,k}, p_{k-1}, \{\hat{d}_{k+j}\}_{j=0}^{N-1}, \{r_{a,k+j}\}_{j=0}^{N-1})$
 $z_{a,k} = C_a x_{a,k}$ {Measure}
 $e_k = z_{a,k} - C_c \hat{x}_{a,k}$ {Estimate error}
 $\{p_{k+j}\}_{j=0}^{N-1} = \mu(\mathcal{P})$ {Solve QP}
 $x_{a,k+1} = A_a x_{a,k} + F_a f_a(p_k) + E_a d_k + G_a w_{a,k}$ {Actuate}
 $\hat{x}_{a,k+1} = A_c \hat{x}_{a,k} + B_c p_k + E_c \hat{d}_k + K_c e_k$ {Predict}
return p_k^*

C. Model Predictive Control

We use the tracking MPC formulation from [16] and minimize the residual by solving the following optimization problem at every time step k

$$\begin{aligned}
 & \text{minimize} && \frac{1}{2} \sum_{j=0}^{N-1} \|\hat{z}_{a,k+1+j} - r_{a,k+1+j}\|_2^2 + \lambda \|\Delta p_{k+j}\|_2^2 \\
 & \text{subject to} && (12) \\
 & && \hat{x}_{a,k+1+j} = A_c \hat{x}_{a,k+j} + B_c p_{k+j} + E_c \hat{d}_{k+j} \\
 & && \hat{z}_{a,k+j} = C_c \hat{x}_{a,k+j} \\
 & && -1 \leq p_{k+j} \leq 1 \\
 & && \Delta p_{\min} \leq \Delta p_{k+j} \leq \Delta p_{\max}
 \end{aligned} \tag{13}$$

The optimal control price $\{p_k^*\}_{j=0}^{N-1}$ is found over the prediction horizon $j = 0, 1, \dots, N-1$. The control price minimizes the deviations from the power consumption reference based on model predictions of the aggregated thermal loads. The first control price p_0^* is broadcast to all loads and the process is repeated at the next time step. Only the optimal control price at the current time step is implemented, e.g. the current price, and consequently closed-loop feedback is obtained. This is often referred to as the receding horizon principle. A regularization term is also added to the objective with penalty λ on the price rate Δp to enforce stability.

Algorithm 1 shows the closed-loop MPC algorithm that runs at every time step [16]. The MPC control law $p = \mu(\mathcal{P})$ is evaluated by solving (13), and real-time computation is enabled from the low-order aggregated controller model.

IV. NUMERICAL EXAMPLE

We model the individual load with a first-order transfer function $G(s)$ from power consumption u to temperature y

$$G(s) = \frac{c}{\tau s + 1} \tag{14}$$

τ is the time constant and c is the gain. The same model is used to model the disturbance response from d to y . We discretize with a zero-order hold and sampling period $T_s = 1$. As an example we set up a portfolio of $L = 10$ loads, each modeled with (14) and the parameters from Table I. We select different price scalings in a , but use the same temperature interval bias b . The tuning weight R is set rather high to minimize control action and is not tuned separately for each load.

i	τ	c	a	b	Q	R
1	55.4	1.32	1.00	21	1	2
2	68.3	2.94	1.22	21	1	2
3	27.4	2.91	1.44	21	1	2
4	58.6	1.97	1.67	21	1	2
5	53.2	2.60	1.89	21	1	2
6	36.9	1.28	2.11	21	1	2
7	45.7	1.84	2.33	21	1	2
8	53.4	2.83	2.56	21	1	2
9	85.8	2.58	2.78	21	1	2
10	77.7	2.92	3.00	21	1	2

TABLE I
PARAMETERS FOR NUMERICAL EXAMPLE WITH $L = 10$

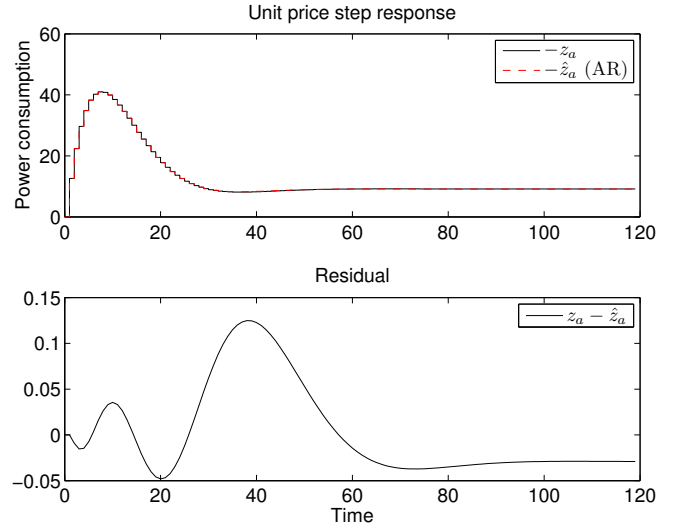


Fig. 4. Power consumption response and estimated response to unit price step (upper) and their residual (lower).

A. Estimated controller model

Fig. 4 depicts the unit price step response of the chosen model. The response fits an AR(2,2) model quite well, except for a small stationary error, even though the response is the sum of $L = 10$ different first-order models. For the chosen numerical parameters in Table I we find the following model

$$\hat{A}_{ar}(q^{-1}) = 1 - 1.756q^{-1} + 0.7798q^{-2} \tag{15a}$$

$$\hat{B}_{ar}(q^{-1}) = -12.62q^{-1} + 12.40q^{-2} \tag{15b}$$

$$\hat{E}_{ar}(q^{-1}) = -2.286q^{-1} + 2.038q^{-2} \tag{15c}$$

Also a response from the disturbance to power consumption was used to estimate the polynomial $\hat{E}_a(q^{-1})$. The final controller model is obtained by adding the integrator model as in (11).

In this numerical example some process noise was added to simulate unmodeled disturbances with $w_a \in N(0, 0.01)$. We tune the controller with the parameters to $\alpha = 0.7$, $\lambda = 10^3$ and $(\Delta p_{\min}, \Delta p_{\max}) = (-1, 1)$. These parameters matter considerably when the plant is subject to stochastics [17].

B. Simulation results

A simulation of the MPC price control with the estimated models is shown in Fig. 5 and 6. The upper plot of Fig. 5 shows the aggregated response and how the power reference is tracked by the aggregator MPC with no offset errors. Since prediction of the reference r_a is available to the MPC, control prices which indirectly change power consumption are broadcast ahead of time in order to minimize the residual. The residual is plotted below, along with the control price. After 55 time steps the control price constraint is active at -1 . At this point the power reference is very high and the aggregator demands all available power from the thermal loads. Note that the control price is not constant when the total reference power consumption is constant, i.e. tracking a constant power requires a ramping of the price due to the dynamics of the loads.

There is a small stationary offset error during ramping of the reference. A double integrator can be added to (10) to eliminate the error. However, this requires the LQ load controllers to increase their order as well. This means the response will be more sensitive to noise and the tuning parameter α becomes extremely important [17]. In this work we accept the ramp offset error and use a single integrator.

The bottom plot of Fig. 6 shows the temperatures of the loads. Some loads are more flexible than others and allow a wider temperature interval, i.e. $b \pm a$ with a large a_i , indicated by the various dashed lines at different levels. Consequently, a more flexible load will have a more varying temperature. However, the temperature is still ensured to lie within the predefined interval, $b \pm a$, due to the constrained control price. The temperature interval can be adjusted for each load by the scaling a_i and can even be time varying. Naturally, the temperature does not depend exclusively on the control price; it also depends on the dynamics, i.e. the time constant, of the load and its controller tuning. The power consumption of each load is shown in the upper plot of the figure. As intended, power consumption mainly occurs when the price is low, as becomes evident when comparing to the price in Fig. 5. The stationary power consumption, when $p = 0$ and $d = 0$, varies from load to load as observed in Fig. 6 because of the different initial levels of u . In our example the combined stationary power consumption of all the loads when disregarding the disturbance is

$$z_a^0 = \sum_{i=1}^L \frac{b_i}{c_i} \quad (16)$$

In Fig. 5 power consumption was plotted around zero as the deviation from this stationary consumption z_a^0 . From (16) it can be seen that the stationary power consumption depends on the number of loads L , their temperature settings (b, a) , their efficiency c , and disturbance d . The methods accounts for local disturbances by forecasting a global disturbance \hat{d} that acts on all loads. Any remaining sources of error will be eliminated by the MPC. A disturbance has also been used in the simulation shown in Fig. 5. After thirty time steps the disturbance kicks in, e.g. a change in outdoor temperature




Fig. 5. Simulation of the aggregator tracking a power consumption r_a by controlling an aggregation of thermal loads. Total power consumption z_a is plotted around zero as the deviation from the stationary consumption z_a^0 . The normalized residual is plotted below along with the control price p . As intended, load consumption is highest when the price is low. The disturbance is forecast dh and eliminated by the MPC. The disturbance shown here is scaled and does not match the units of the y-axis.




Fig. 6. Load output temperatures y_i (lower) and their temperature intervals $b_i \pm a_i$ (dashed lines). Also their power consumptions u_i are plotted (upper).

which changes the power consumption. By forecasting the disturbance, the tracking error can be greatly improved. A forecast that is close to the real disturbance is implemented and this is why almost no deviations are seen at the disturbance transitions after 30 and 70 time steps. A single disturbance acting on all loads can be justified when the loads are geographically close to each other and the disturbance considered is the outdoor temperature. Solar radiation has a more local impact on buildings but can also be forecast for a larger area [18].

V. DISCUSSION

The control price used in this method should *not* be interpreted as the final billing price for each load. The control price helps the aggregator meet its balancing objective, but does not create an incentive for the loads to choose high temperature intervals, a_i . High temperature intervals increase the flexibility and thereby also the regulating power. Moreover, it lowers temperature variations for all aggregated loads. Consequently, all loads are forced to help the aggregator reach its tracking goals regardless of their own optimal strategy. This method is not optimal for every load in terms of energy savings, but will ensure that total aggregated power consumption follows production, e.g. from wind, to the benefit of the overall energy system.

A negative control price should not be considered a subsidy. However, the sign on p merely says whether the aggregator needs up or down regulation. It is evident from the simulations that if the price is negative a majority of the time, often the loads with the largest temperature interval will be at very low temperatures, thus saving a lot of power. If the price is mostly positive, the loads with low temperature intervals will save power. The opposite is true if we consider cooling rather than heating. In this case, the final electricity cost for each load should not depend exclusively on the control price. There must be a clear incentive to provide a large temperature interval, since this will enable more power at the aggregator level and less discomfort for all loads. Final billing could be calculated on the basis of consumption $u_{i,k}$, and temperature interval $2a_i$, defining how much load i allows the temperature to vary. Also the heat capacity of the load, e.g. the time constant τ_i , could play a role if it was measured. Instead of billing for power consumption using the control price, we suggest putting a price on flexibility, i.e. the temperature interval a_i which, in practice, could be time varying.

As a consequence of using linear unconstrained controllers for the loads, no actuator saturation was considered. If actuator constraints are involved, the price response will not be linear, and clipping of the power will be observed. As a result the response in Fig. 2 might look more sigmoidal and bend at the price limits ± 1 . One way to prevent this problem is to restrain loads from setting a_i too high compared to its capacity and the expected disturbances. Another way is to include an adaptive model of the price response. Note that time-varying linear models can be easily implemented in the MPC algorithm by changing the coefficients of the controller model (15).

VI. CONCLUSIONS

The method described in this paper enables a linear MPC, based on a low-order SISO ARX model, to balance power production with consumption of a considerable number of thermal loads in real-time. The method requires linear temperature set point controllers to control the loads as well as model estimation of the price response at the aggregator level. The aggregator MPC controls total power consumption of all loads indirectly through a broadcast real-time price, i.e. one-way communication. It also handles the load temperature

constraints through price constraints. The MPC incorporates forecasts of disturbances and power production, e.g. time-varying wind power forecasts, in order to react ahead of time. Added integral control eliminates model and forecast errors, while feedback is provided by measuring total load power consumption. Individual loads can set their own desired upper and lower temperature bounds. The method was demonstrated through simulation and allows for integration of flexible thermal loads a smart energy system in which consumption follows a changing production.

REFERENCES

- [1] R. Halvgaard, N. K. Poulsen, H. Madsen, and J. B. Jørgensen, "Economic Model Predictive Control for Building Climate Control in a Smart Grid," in *Proceedings of IEEE PES Innovative Smart Grid Technologies (ISGT)*, 2012.
- [2] R. Halvgaard, P. Bacher, B. Perers, E. Andersen, S. Furbo, J. B. Jørgensen, N. K. Poulsen, and H. Madsen, "Model Predictive Control for a Smart Solar Tank based on Weather and Consumption Forecasts," in *Proceedings of 2012 International Conference on Solar Heating and Cooling for Buildings and Industry (SHC)*, 2012.
- [3] T. G. Hovgaard, L. F. S. Larsen, K. Edlund, and J. B. Jørgensen, "Model predictive control technologies for efficient and flexible power consumption in refrigeration systems," *Energy*, vol. 44, no. 1, pp. 105–116, 2012.
- [4] F. Oldewurtel, A. Ulbig, A. Parisio, G. Andersson, and M. Morari, "Reducing peak electricity demand in building climate control using real-time pricing and model predictive control," in *49th IEEE Conference on Decision and Control (CDC)*, dec 2010, pp. 1927–1932.
- [5] M. Braun and P. Strauss, "A review on aggregation approaches of controllable distributed energy units in electrical power systems," *International Journal of Distributed Energy Resources*, vol. 4, pp. 297–319, 2008.
- [6] A. Subramanian, M. Garcia, A. Dominguez-Garcia, D. Callaway, K. Poolla, and P. Varaiya, "Real-time Scheduling of Deferrable Electric Loads," in *2012 American Control Conference (ACC), Montreal, Canada, June 27-29, 2012*.
- [7] R. Scattolini, "Architectures for distributed and hierarchical Model Predictive Control – A review," *Journal of Process Control*, vol. 19, pp. 723–731, 2009.
- [8] O. Corradi and H. Ochsenfeld, "Integration of fluctuating energy by electricity price control," Master's thesis, Tech. Univ. Denmark, 2011.
- [9] P. Nyeng, J. Østergaard, M. Tøgeby, and J. Hethy, "Design and Implementation of Frequency-responsive Thermostat Control," in *45th International Universities Power Engineering Conference (UPEC)*, 2010, pp. 1–6.
- [10] M. Zugno, J. M. Morales, P. Pinson, and H. Madsen, "Modeling Demand Response in Electricity Retail Markets as a Stackelberg Game," in *Proceedings of IAEE conference*, 2012.
- [11] T. G. Hovgaard, K. Edlund, and J. B. Jørgensen, "The Potential of Economic MPC for Power Management," in *49th IEEE Conference on Decision and Control (CDC)*, 2010, pp. 7533–7538.
- [12] M. Kraning, E. Chu, J. Lavaei, and S. Boyd, "Message Passing for Dynamic Network Energy Management," [*To appear in*] *Foundations and Trends in Optimization*, 2013.
- [13] J. B. Jørgensen and S. B. Jørgensen, "MPC-relevant prediction-error identification," in *Proceedings of ACC*, 2007, pp. 128–133.
- [14] P. Bacher and H. Madsen, "Identifying suitable models for the heat dynamics of buildings," *Energy and Buildings*, vol. 43, no. 7, pp. 1511–1522, 2011.
- [15] J. K. Huusom, N. K. Poulsen, S. B. Jørgensen, and J. B. Jørgensen, "Tuning SISO offset-free Model Predictive Control based on ARX models," *J. Process Control*, 2012.
- [16] J. B. Jørgensen, J. K. Huusom, and J. B. Rawlings, "Finite Horizon MPC for Systems in Innovation Form," in *Proceedings of CDC*, 2011, pp. 1896–1903.
- [17] J. K. Huusom, N. K. Poulsen, S. B. Jørgensen, and J. B. Jørgensen, "ARX-Model based Model and Predictive Control and with Offset-Free and Tracking," in *Proceedings of the 20th European Symposium on Computer Aided Process Engineering (ESCAPE)*, 2010.
- [18] P. Bacher, H. Madsen, and B. Perers, "Short-term solar collector power forecasting," in *In proceedings of ISES Solar World Conference*, 2011.