Electricity Market Optimization of Heat Pump Portfolio

Benjamin Biegel Palle Andersen Tom S. Pedersen Kirsten Mølgaard Nielsen
Jakob Stoustrup Lars Henrik Hansen

***Abstract*— We consider a portfolio of domestic heat pumps controlled by an aggregator. The aggregator is able to adjust the consumption of the heat pumps without affecting the comfort in the houses and uses this ability to shift the main consumption to hours with low electricity prices. Further, the aggregator is able to place upward and downward regulating bids in the regulating power market based on the consumption flexibility. A simulation is carried out based on data from a Danish domestic heat pump project, historical spot prices, regulating power prices, and spot price predictions. The simulations show that electricity price reductions of $18-20$ % can be achieved compared to the heat pumps currently in operation.**

I. Introduction

With an increasing focus on climate-related issues and rising fossil fuel prices, the penetration of renewable energy sources is likely to increase in the foreseeable future throughout the developed world [1]. The Danish electric power system, which is the focus of this work, is a particularly interesting case with a wind energy penetration of 30 % today and an expected 2020 penetration of 49.5 % [2]. When conventional power plants are outdone with renewables such as wind turbines and photovoltaics, the ability to provide balancing services in the classical sense disappears, as the renewable energy sources often will fully utilize the available power. It is therefore evident that in a grid with high penetration of renewables, alternative sources of balancing services must be established. One of the approaches to obtaining such services is the *smart grid* concept, where demand-side devices with flexible power consumption take part in the balancing effort [3]. The basic idea is to let an *aggregator* manage and optimize a portfolio of flexible demand-side devices on behalf of the balancing responsible party (BRP) for this consumption. This allows the balancing responsible to utilize the accumulated flexibility in the unbundled electricity markets on equal terms with conventional generators [4]. In the future Danish electricity system it is expected that domestic heat pumps will play an important role as flexible consumption: already now, around 27.000 heat pumps are installed in Denmark [5], and potentially $205,000$ households can benefit from replacing oil-fired boilers with heat pumps in the coming years [6]. It is therefore most relevant to consider how to aggregate and control this flexibility towards the electricity markets.

Control of smaller flexible consumers to support grid stability has been discussed as early as the 1980s [7]. Since, the topic of demand-side management has received much attention from a research perspective including control of heat pumps [8], [9]. In particular, optimization of heat pumps has received much attention in Denmark the last few years. See, e.g., [10], [11], [12], [13], [14]. These works consider how the operation of heat pumps can be optimized to support grid stability and how to lower the operational electricity costs by performing spot price optimization of the consumption. None of the works do, however, account for the structure of the Nordic system, which consists of a day-ahead spot market and an intra day balancing market. As an example of this, several works use the electricity spot price as a *price signal* that the aggregator will face without any planning phase. In this work we move closer to the real electricity market by including both a day-ahead planning phase and an intra-day balancing market. The aggregator will purchase electricity based on spot price *predictions* at the day-ahead market and will intra-hour adjust the operation of the portfolio according to the experienced load and possibly place bids of upward and downward regulation in the intra-day market. Finally, the intra-day imbalances will be settled as balancing power according to the regulations.

The structure of the paper is as follows. First in Sec. II we model the heat pump portfolio and in Sec. III we describe the electricity markets. In Sec. IV we develop a control strategy that takes the spot prices into account and a control strategy that further is able to bid into the regulating power market. Following in Sec. V we show two simulation examples of the developed control strategies and finally in Sec. VI we conclude the work.

II. Lumped Portfolio Model

A. The Heat Pump Project "Styr Din Varmepumpe"

Several large projects dealing with flexible consumption are currently ongoing in Denmark. The project "Styr din varmepumpe" (meaning: *control your heat pump*) deals specifically with understanding domestic heat pumps and how they can be operated depending on the electricity markets [15]. In this project, around 200 heat pumps installed in Danish homes have been subsequently equipped with various measurement devices such that power consumption, flows, temperatures etc. can be measured. The data is collected

B. Biegel, P. Andersen, T. S. Pedersen, K. M. Nielsen and J. Stoustrup are with the Department of Electronic Systems, Automation and Control, Aalborg University, Denmark; e-mail: {bbi,pa,tom,kmn,jakob}@es.aau.dk.

L. H. Hansen is with DONG Energy A/S (largest Danish power producer), Copenhagen, Denmark; e-mail: LARHA@dongenergy.dk.

The work is completed as a part of two projects: the *iPower* project supported by the Danish government via the DSR-SPIR program 10-095378 and the *READY* project supported by PSO funds administered by Energinet.dk via the ForskEl project program 2012-1-10757.

via Internet connections and can be used for modeling and analysis. This data forms the background for this work.

B. System Architecture

The starting point of this work is the Nordic unbundled liberalized electricity system architecture. In this setup, the Transmission System Operators (TSOs) are responsible for secure and reliable system operation and must consequently ensure balance between production and consumption. Generally, in an unbundled electricity system, TSOs do not own production units and must therefore procure ancillary services in the electricity markets to ensure system stability.

The aggregator is a legal entity able to enter into flexibility contracts with consumers allowing the aggregator to manage the consumers' flexible consumption; in return, the consumers will achieve some type of compensation. This enables the aggregator to utilize the accumulated flexibility to participate in the electricity markets through the consumers' BRP. The flexible devices are managed by the aggregator through a technical unit often referred to as a Virtual Power Plant (VPP).

C. Aggregated House Model

It is desired to have a simple model that describes the accumulated flexibility of all the houses in the portfolio rather than a model for each of the houses. There are two reasons for this. The first reason is that it can be computationally difficulty to perform optimization across thousands of heat pumps each described by its own model. The second reason is that it is difficult to predict the future behavior of a single house due to the many unpredictable disturbances affecting a house: fluctuating sunshine, opening/closing of doors and windows, the use of wood stove etc. For a lumped heat pump model, however, these local disturbances will, however, smooth out as the number of houses in the portfolio increases.

Several works have suggested the use of a first order model to describe the energy level, or temperature, in a house; see, e.g., [12], [14], [16]. Such a model can be formulated as

$$\dot{T}_i(t) = \frac{1}{R_i C_i}\left(T_{\text{a},i}(t) - T_i(t)\right) + \frac{1}{C_i}\left(u_i(t) + v_i(t) + w_i(t)\right) \tag{1}$$

for house number i where the constants $R_i, C_i \in \mathbf{R}_+$ are the thermal resistance and the heat capacity of the house, respectively, while $T_i(t) \in \mathbf{R}$ is the indoor temperature and $T_{\text{a},i} \in \mathbf{R}$ is the outdoor (ambient) temperature affecting the house. The input $u_i(t)$ is the electrical equivalent of the stored thermal energy, $v_i(t) \in \mathbf{R}$ represents a deterministic daily load pattern on the house, i.e. $v_i(t) = v_i(t + 24 \text{ hours})$ while $w_i(t) \in \mathbf{R}$ is an exogenous disturbance. Note that this model covers houses with electrical heating, but also with a transformation of parameters a house with a heat pump with a given COP by letting $C_i = C_{\text{house}}/\text{COP},\ R_i = R_{\text{house}}\text{COP}$.

In this work we assume that we can describe the entire heat pump portfolio by a first order model. This is clearly a rough assumption: if the individual houses are described by first order models as (1), the order of the lumped model will be the total number of houses unless some houses have identical parameters C_i, R_i. It may seem crude to make a lumped first order model, as the houses definitely will have different thermal resistance and heat capacity; however, the parameters will be in the same order of magnitude as all the houses are standard-sized Danish houses. Further, it must be stressed that the purpose of the aggregated model is not to accurately describe the houses' states; rather, the purpose is to have a model suitable for rough planning of the future electricity consumption. The benefit of actually having a very accurate model will also be very limited as the flexibility optimization depends on several parameters that we do not know accurately such as future temperatures and spot prices. Finally, attempts on individual household modeling on inhabited houses show that the disturbances often are so great that the actual house dynamics cannot be observed. Further argumentation and real life demonstrations motivating the use of a lumped heat pump model can be found in [17].

This leads us to the following model description of the entire portfolio. Let $T(t) \in \mathbf{R}$ be the average indoor temperature, $T_\text{a}(t) \in \mathbf{R}$ the average outdoor temperature, $u(t) \in \mathbf{R}$ the average heat pump power input, $v(t) \in \mathbf{R}$ the average daily load profile, and $w(t) \in \mathbf{R}$ the average disturbance. The aggregated model can then be described as

$$\dot{T}(t) = \frac{1}{RC}\left(T_\text{a}(t) - T(t)\right) + \frac{1}{C}\left(u(t) + v(t) + w(t)\right) \tag{2}$$

where the constants $R, C \in \mathbf{R}_+$ are the parameters of the aggregated model. As mentioned, a benefit of this model is that the outdoor temperature $T_\text{a}(t)$, the daily load profile $v(t)$, and the exogenous input $w(t)$ to an extend will smooth out as the number of houses increase. Note that we in this work only consider scheduling of the operation of the accumulated system represented by (2); we do not discuss how to control the individual devices but assume that an underlying dispatch algorithm distributes power to the individual houses in order to be able to let local control loops reject individual disturbance patterns. For details on how this can be achieved, see for example [18], [19].

D. Thermal Flexibility for House Heating

Figure 1 shows indoor temperature measurements from four of the houses that are at part of the heat pump project over a one-month period. The heat pumps operate using the default heat pump controller. The figure shows that the indoor temperature varies several degrees for all the houses over the period, which indicates the foundation for this work: that people are used to and comfortable with indoor temperatures varying a couple of degrees, hence the indoor temperature in a house does not have to be fixed at a given temperature setpoint. This motivates a formulation where the indoor temperature is allowed to vary within a given interval for each house $\underline{T}_i \leq T_i \leq \overline{T}_i$. This gives the following requirement to the aggregated model:

$$\underline{T} \leq T(t) \leq \overline{T} \tag{3}$$

where $\underline{T}, \overline{T} \in \mathbf{R}$ describe the average temperature limits.

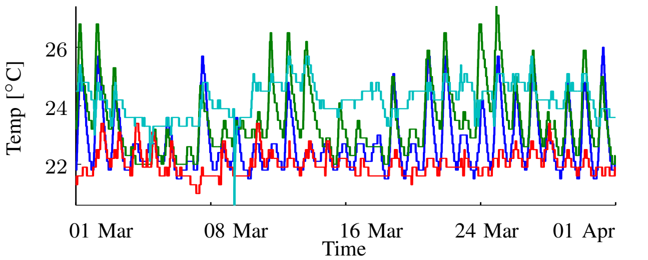

Fig. 1. One month's indoor temperature measurements for four houses during March 2012.

Finally, the power consumption of a heat pump is limited $\underline{u}_i \le u_i \le \overline{u}_i$ which for the aggregated model implies that

$$\underline{u} \le u(t) \le \overline{u} \tag{4}$$

where $\underline{u}, \overline{u} \in \mathbf{R}$ describe the average power limits.

Further note that honoring the temperature and power limits on the aggregated system as described by (3) and (4) will not guarantee that the individual device constraints are honored; this is the task of the dispatcher which is not described in this work.

E. Model Estimation

The purpose of the heat pump portfolio is to optimize the flexibility towards spot prices and the intra-day markets. As these are hourly markets in the Nordic system we discretize the portfolio model with a sampling time of 1 hour and obtain

$$T(k+1) = aT(k) + (1-a)T_{\mathrm{a}}(k) + b\left(u(k) + v(k) + w(k)\right) \tag{5}$$

where $a, b \in \mathbf{R}$, which depend on R, C and are found by discretization, and k is used to indicate the sample number.

One year's data from 130 heat pump heated houses is used to identify the parameters a, b via quadratic fitting. For details on such parameter estimation, see for example [14]. Figure 2 shows the power added to the portfolio of houses from the heat pump u, the daily load profile v, the exogenous input w, and the energy that drains out due to the lower ambient temperature, which we denote d for *drain*. The figure shows averages for the entire portfolio over a two-month period. The figure illustrates that the average heat pump power u throughout the period is in the order of $1.0-2.5$ kW. Further it can be seen that the load v varies daily between 500 and 700 W describing the average profile of heat added by people in the house, electronics, wood stove, etc. Finally, the unpredictable load w has a contribution in the magnitude ± 500 W caused by the disturbances that cannot be captured by the daily load profile. The parameters of the model reveal a time constant of 33 hours.

III. Electricity Markets

In this section we briefly describe the electricity markets that the aggregator faces.

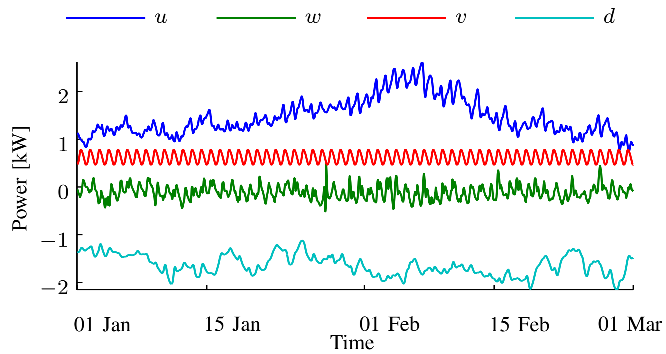

Fig. 2. Energy transfer to portfolio of houses from heat pump u, daily load v, exogenous load w, and drain d.

A. Electricity Spot Market

In the Nordic system, electricity is bought and sold at the electricity spot market. Every day before 12 p.m. (noon), buyers and sellers of electricity can place bids into the electricity spot market specifying what volume they will buy/sell depending on the price of electricity for each hour of the following day. The hourly spot prices for the following day will be set to the intersection between the aggregated supply and demand curves. All electricity traded in each hour will be settled at this spot price; further, the spot price determines the volumes of traded electricity. As the spot prices are unknown at the time when electricity is purchased, the aggregator must rely on a spot price prognosis when purchasing electricity day-ahead.

Let the spot prices at hour k be denote $\pi(k)$ and let $\tilde{\pi}(k)$ denote the prediction of $\pi(k)$ available day-ahead before gate closure. To illustrate this, assume that the current time is between 11 a.m. and 12 p.m. (last hour before gate closure); further, let this correspond to sample $k = 12$. At this time we know the spot prices for the current day $\pi(1), \ldots, \pi(24)$, but we do not know the spot prices the following day (the day-ahead) $\pi(25), \ldots, \pi(48)$, which are not announced until $k = 13$ (i.e. 1 p.m.). We do, however, have spot price predictions for the following day, $\tilde{\pi}(25), \ldots, \tilde{\pi}(48)$. This is illustrated in Fig. 3. The figure further illustrates what is generally the case, namely that the predictions are able to capture the shape of the actual spot price realization.

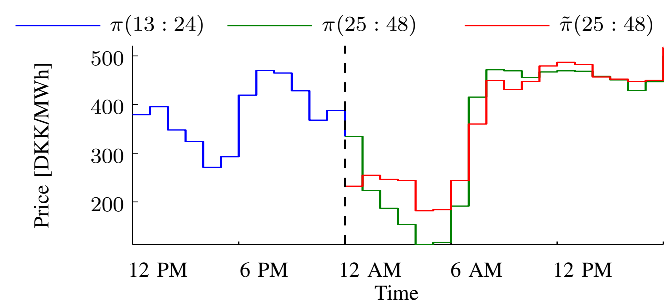

Fig. 3. Spot prices π and predictions $\tilde{\pi}$ on January 9 and 10, 2011.

B. Balancing Power and Regulating Power

During the day, the portfolio of consumers will consume a given volume of power $u(k)$ for each of the hours of the day. If the actual consumption $u(k)$ deviates from the electricity purchased day-ahead in the spot market, the aggregator will by definition trade the difference as *balancing power* with the TSO. The price of balancing power is the regulating power price (RP price) which we denote $\pi_{\text{RP}}(k)$. The total cost $J_{\text{el}}(k)$ of electricity in hour k thus depends on the electricity purchased at the spot market, denoted $u_{\text{spot}}(k)$, and the electricity actually consumed $u(k)$:

$$J_{\text{el}}(k) = u_{\text{spot}}(k)\pi(k) + (u(k) - u_{\text{spot}}(k))\pi_{\text{RP}}(k) \quad (6)$$

assuming that all consumers contributing to $u(k)$ are hourly metered and settled.

To counteract for the imbalances caused by electricity traders, the TSO activates regulating power from the regulating power market. Providers of regulating power can place bids in the regulating power market up to 45 minutes before the hour of operation, specifying the price at which they are willing to increase or decrease production or consumption. The TSO will activate the required volume of regulating power and will select the bids in merit order after price. The RP price will be set as (defined by) the bidding price of the most expensive regulating power bid activated in a delivery hour. If the direction of regulation is upward, the RP price will be greater than or equal to the spot price; similarly, if the direction of regulation is downward, the RP price will be less than or equal to the spot price. The RP price will be used to settle all the provisions of regulating power in that given hour. Further, it will be used to settle all imbalances according to the power balancing settlement procedures. Note that the RP price is not published until after the hour in question.

An example of the regulating power price is illustrated in Fig. 4 where we compare the hourly spot price π with the regulating power price π_{RP}. The figure illustrates that the

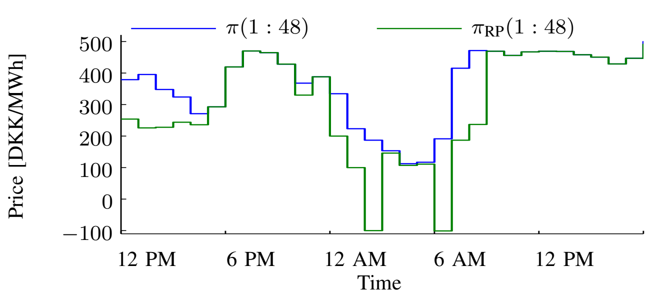

Fig. 4. An example with relatively large differences between the spot prices π and regulating power prices π_{RP}. January 9 and 10, 2011.

regulating power price is lower than the spot price in the first hours of the figure indicating that the system is in downward regulation. If we have a higher consumption in these hours than the purchased electricity at the spot market, we will buy this additional electricity (balancing power) from the TSO at RP price which is beneficial as the RP price is low; on the contrary, if we have a lower consumption in these hours, we will sell the excess electricity to the TSO at RP price which is disadvantageous as the RP price is low. As the RP price is not published until after the hour in question, it is not possible to adjust the consumption corresponding to the RP price.

IV. Controller Synthesis

In this section we describe two strategies of operating the flexibility of the heat pumps: a strategy that optimizes the flexibility according to spot prices and a strategy that also bids into the regulating power market.

A. Spot Price Optimization

The overall idea in the spot price optimization control is that the aggregator achieves a lower operational cost of the portfolio of heat pumps by shifting consumption into hours of low electricity prices. This optimization is non-trivial due to the fact that the spot prices only apply for the electricity purchased at the spot market – if the consumption of the portfolio deviates from the purchased volumes, the difference will be settled using the regulating power price according to (6). For this reason, the optimization is divided into a day-ahead optimization that determines the volumes that will be bought at the electricity markets and an intra-day optimization that operates the portfolio hour by hour. This is elaborated in the following.

1) Day-ahead Optimization: The key idea in the day-ahead optimization is to purchase the electricity needed for the following day based on spot price and outdoor temperature predictions such that the cost of operating the portfolio is minimized. To formally describe this, we assume that the current hour k corresponds to the last hour before gate closure, i.e., the hour between 11 a.m. and 12 p.m. Define $\mathcal{K} = \{k+13, \dots, k+36\}$; hence $\mathcal{K}$ will correspond to a set of the 24 hours of the following day which we have to purchase electricity for. The overall objective is to minimize the electricity cost for operation the following day which is given by $\sum_{\kappa \in \mathcal{K}} \pi(\kappa)u(\kappa)$ where $\pi(\kappa)$ is the hourly spot price and $u(\kappa)$ is the hourly consumption. Again we remind that the spot prices $\pi(\kappa)$ are not available day-ahead; hence, spot price predictions $\tilde{\pi}(k)$ must be utilized to minimize the objective.

Further, we want the portfolio to be at a temperature setpoint T_{sp} in steady state instead of converging to either of the temperature limits $\overline{T}, \underline{T}$. This is achieved by minimizing the norm of a state $x(k) \in \mathbf{R}$ that corresponds to the integrated temperature error as described by:

$$x(k+1) = x(k) + T(k) - T_{\text{sp}} \quad (7)$$

In this work it is chosen to minimize the integrated temperature error in a two-norm sense as described in the following.

To perform this optimization, it is necessary to collect predictions of both the temperature and integrated temperature tracking error for the first hour of the following day $\tilde{T}(k+13), \tilde{x}(k+13)$ which are 12 hours into the future; besides, spot price and outdoor temperature predictions for each hour of the following day $\tilde{\pi}(\kappa), \tilde{T}_{\text{a}}(\kappa), \kappa \in \mathcal{K}$ must be collected. Based on this data we can formulate the optimization as follows:

$$\begin{aligned} \text{minimize} \quad & \sum_{\kappa \in \mathcal{K}} \left(\tilde{\pi}(\kappa) u(\kappa) + k_I x^2(\kappa) \right) \\ \text{subject to} \quad & T(\kappa+1) = aT(\kappa) + (1-a)\tilde{T}_{\text{a}}(\kappa) + \\ & \qquad\qquad b\left(u(\kappa) + v(\kappa)\right), \quad \kappa \in \mathcal{K} \\ & x(\kappa+1) = x(\kappa) + T(\kappa) - T_{\text{sp}}, \quad \kappa \in \mathcal{K} \\ & u(\kappa) \in \mathcal{U}, \quad T(\kappa) \in \mathcal{T}, \quad \kappa \in \mathcal{K} \\ & T(k+13) = \tilde{T}(k+13) \\ & x(k+13) = \tilde{x}(k+13) \end{aligned} \tag{8}$$

where the variables are $u(\kappa), T(\kappa), x(\kappa), \kappa \in \mathcal{K}$ and $k_{\text{I}} \in \mathbf{R}$ is a trade-off parameter. The data to the problem is the predicted spot prices and outdoor temperatures $\tilde{\pi}(\kappa), \tilde{T}_{\text{a}}(\kappa), \kappa \in \mathcal{K}$, the daily load profile $v(\kappa), \kappa \in \mathcal{K}$, and the predicted temperature and integrated error in the first hour of the following day $\tilde{T}(k+13), \tilde{x}(k+13)$. The sets $\mathcal{T}, \mathcal{U}$ represent the power and temperature limitations as described by (3) and (4). The solution $u^\star_{\text{spot}}(\kappa), \kappa \in \mathcal{K}$ are the volumes of electricity we will purchase for the following day.

The reason for choosing a horizon of 24 hours is that 24 hour forecasts of the following day's spot prices is a standard product that can be purchased from forecasting providers. As the time constant of the aggregated houses is in the magnitude of 33 hours, a longer horizon could be beneficial; however, such long spot price predictions are not available.

2) Intra-Day Optimization: Day-ahead we purchase electricity at the spot market based on predictions of load on the heat pump as described above. Intra-day we decide how to actually operate the portfolio. This intra-day operation may differ from the plan made day-ahead, as the houses will experience loads that differ from the predictions. Different strategies can be chosen for the intra-day operation. One strategy is to track the electricity we have purchased day-ahead as closely as possible to avoid trading balancing power with the TSO at possibly unfavorable prices. Another option, which we choose in this work, is to simply consider the known spot prices as predictions of the regulating power price. The reason for choosing this approach is that the upward and downward regulating power prices on average only differs around 10 % from the spot price; further, we are only penalized when our imbalance is in the same direction as the overall imbalance – else we are rewarded according to the (6). This indicates that deviations from the purchased electricity typically are not associated with a large penalty.

Assume that we are at hour k and let H describe the number of hours into the future that the spot prices are known. Further let $\mathcal{H} = \{k, \ldots, k+H-1\}$ denote the set of future hours where the spot price is known. To perform intra-day optimization we must collect the current temperature and current integrated error $T(k), x(k)$; further, outdoor temperature predictions and known spot prices must be collected $\tilde{T}_{\text{a}}(\kappa), \pi(\kappa), \kappa \in \mathcal{H}$. The object of this problem is again to minimize the cost of operating the portfolio and the integrated error subject to the temperature bands. By using the known spot prices as predictions of the regulating power price, the intra-day optimization problem becomes:

$$\begin{aligned} \text{minimize} \quad & \sum_{\kappa \in \mathcal{H}} \left(\pi(\kappa) u(\kappa) + k_I x^2(\kappa) \right) \\ \text{subject to} \quad & T(\kappa+1) = aT(\kappa) + (1-a)\tilde{T}_{\text{a}}(\kappa) + \\ & \qquad\qquad b\left(u(\kappa) + v(\kappa)\right), \quad \kappa \in \mathcal{H} \\ & x(\kappa+1) = x(\kappa) + T(\kappa) - T_{\text{sp}}, \quad \kappa \in \mathcal{H} \\ & u(\kappa) \in \mathcal{U}, \quad T(\kappa) \in \mathcal{T}, \quad \kappa \in \mathcal{H} \end{aligned} \tag{9}$$

where the variables are $u(\kappa), T(\kappa), x(\kappa), \kappa \in \mathcal{H}$. The data to the problem is the known spot prices and outdoor temperature predictions $\pi(\kappa), \tilde{T}_{\text{a}}(\kappa), \kappa \in \mathcal{H}$, the daily load profile $v(\kappa), \kappa \in \mathcal{H}$, and the current temperature and integrated error $T(k), x(k)$. We denote the solution $u^\star_{\text{intra}}(\kappa)$.

The first element of the solution $u^\star_{\text{intra}}(k)$ is now applied meaning that the VPP will regulate the portfolio to collectively consume the electricity $u^\star_{\text{intra}}(k)$ within the current hour. In this work we do not discuss how the power $u^\star_{\text{intra}}(k)$ is dispatched among the individual heat pumps – we only state that the heat pump portfolio collectively should consume $u^\star_{\text{intra}}(k)$ within hour k.

Note that other strategies could be implemented instead; for example, the day-ahead optimization could be merged with the intra-day optimization when planning how to purchase electricity day-ahead.

3) Algorithm: We are now able to describe the algorithm for spot price optimization. This is presented in Algorithm 1.

Algorithm 1: Spot Price Optimization Algorithm

---

for $k = 1, 2, \ldots$ **do**
  Collect the published spot prices $\pi(\kappa)$ and temperature predictions $\tilde{T}_{\text{a}}(\kappa)$ for the horizon $\kappa \in \mathcal{H}$ and collect current state of the heat pump portfolio $T(k), x(k)$;
  Optimize intra-day operation of the portfolio by solving Problem 9 to obtain $u^\star_{\text{intra}}(\kappa), \kappa \in \mathcal{H}$;
  Let VPP steer portfolio's consumption to $u^\star_{\text{intra}}(k)$;
  if *Current hour is between 11 a.m. and 12 p.m.* **then**
    Collect predictions of spot prices and temperatures for the following day $\tilde{\pi}(\kappa), \tilde{T}_{\text{a}}(\kappa), \kappa \in \mathcal{K}$;
    Collect state predictions $\tilde{T}(k+13), \tilde{x}(k+13)$ (available from the latest solution of (9));
    Optimize spot trades by solving Problem 8 to achieve $u^\star_{\text{spot}}(\kappa), \kappa \in \mathcal{K}$;
    Purchase electricity $u^\star_{\text{spot}}(\kappa), \kappa \in \mathcal{K}$ for the following day.
  end
end

B. Spot Price and Regulating Power Optimization

In this section we present an extension to the spot price optimization strategy by letting the aggregator place bids of regulating power via the BRP into the regulating power market.

1) Consumer as Provider of Regulating Power: A BRP for flexible consumption participating in the regulating power market must submit operational schedules for the portfolio's planned consumption and is allowed to update this schedule if it is discovered that the schedule cannot be followed [20]. The regulations do not specify the deadline for updating the operational schedules. In this work we assume that the aggregator is allowed to update the operational schedule for consumption up to 45 minutes before the hour of operation; hereafter, the aggregator must commit to the planned consumption for the given hour.

2) Bidding in the Regulating Power Market: We enable the aggregator via the BRP to bid in the regulating power market by expanding Algorithm 1. Let $u_{\text{sch}}(k)$ denote the scheduled power consumption at hour k. The aggregator cannot update this volume after hour $k-1$, i.e., the hour before operation. Further, let $u_{\text{act}}(k)$ denote the volume of regulating power the aggregator is activated to deliver in hour k; hence, the portfolio *must* consume the power $u_{\text{reg}}(k) = u_{\text{sch}}(k) + u_{\text{act}}(k)$ in hour k. Note that we use the convention that $u_{\text{act}}(k) < 0$ corresponds to upward regulation and $u_{\text{act}}(k) > 0$ corresponds to downward regulation.

The bidding strategy must ensure that we only deviate from the scheduled consumption $u_{\text{sch}}(k)$ if this is economically favorable for the aggregator. A number of different bidding strategies can be imagined. In this work we do not seek to predict the future RP prices, which is a very difficult task, but instead implement a simple strategy that examines the marginal cost associated with being activated in the following hour for a given delivery of regulating power. This marginal cost is then used as our bid in the regulating power market. As the Nordic regulating power market has a minimum bid size of 10 MW, we simply examine the marginal cost of delivering any feasible multiple of 10 MW. As an example, for a portfolio with limits $\underline{u} = 0$ MW, $\overline{u} = 40$ MW with a scheduled consumption of $u_{\text{sch}}(k) = 10$ MW for the next hour, we will examine the marginal cost of delivering 10 MW upward regulation and $10, 20$, and 30 MW downward regulation. Following, we use these marginal costs as bids for the four regulating power deliveries.

Described more formally, we determine the cost of activating a regulating power activation of $u_{\text{act}}(k)$ by solving

$$\begin{array}{ll} \text{minimize} & \displaystyle\sum_{\kappa\in\mathcal{H}} \left(\pi(\kappa)u(\kappa) + k_I x^2(\kappa)\right) - T(H)\frac{\tilde{\pi}_{\text{avg}}}{b} \\ \text{subject to} & T(\kappa+1) = aT(\kappa) + (1-a)\tilde{T}_{\text{a}}(\kappa) + \\ & \qquad\qquad b\left(u(\kappa) + v(\kappa)\right), \quad \kappa \in \mathcal{H} \\ & x(\kappa+1) = x(\kappa) + T(\kappa) - T_{\text{sp}}, \quad \kappa \in \mathcal{H} \\ & u(\kappa) \in \mathcal{U}, \quad T(\kappa) \in \mathcal{T}, \quad \kappa \in \mathcal{H} \\ & u(k) = u_{\text{sch}}(k) + u_{\text{act}}(k). \end{array} \tag{10}$$

The variables and the data to the problem have all been previously described (see Problem 8) except $\tilde{\pi}_{\text{avg}}$ which is the predicted average spot price for the following day; hence the term $T(H)\frac{\tilde{\pi}_{\text{avg}}}{b}$ is a way to appraise the energy stored in the portfolio at the end of the horizon. The optimal value is denoted $J^\star(u_{\text{act}}(k))$ and describes the cost if we choose to bid and are activated for regulating power of volume $u_{\text{act}}(k)$.

The regulating power bid $\pi_{\text{bid}}(u_{\text{act}}(k))$ for providing the regulating power delivery $u_{\text{act}}(k)$ can be found as the RP price where the cost of not providing regulating power $J^\star(0)$ equals the cost of being activated for a delivery $u_{\text{act}}(k)$ given by $J^\star(u_{\text{act}}(k))$ plus the portfolio's imbalance cost $u_{\text{act}}(k)(\pi_{\text{RP}}(k) - \pi(k))$. We can therefore find the regulating power bid $\pi_{\text{bid}}(u_{\text{act}}(k))$ associated with a regulation power delivery $u_{\text{act}}(k)$ by solving

$$J^\star(0) = J^\star(u_{\text{act}}) + u_{\text{act}}(\pi_{\text{bid}}(u_{\text{act}}(k)) - \pi(k)). \tag{11}$$

We illustrate the equation with a small example: assume the cost of operating the portfolio with no activation is $J^\star(0) = 1000$ DKK while the cost of delivering 10 MW of downward regulation ($u_{\text{act}} = 10$ MW) is $J^\star(10) = 1.200$ DKK and assume a spot price of 200 DKK/MW. In this case, our downward regulating power bid is $\pi_{\text{bid}}(10) = 180$ DKK/MW according to (11) as we will break even at this price while we will profit if the regulating power price becomes even lower. If the cost of operating the portfolio to provide 10 MW of upward regulation ($u_{\text{act}} = -10$ MW) is $J^\star(-10) = 1.200$ DKK, the regulating power bid will be $\pi_{\text{bid}}(-10) = 220$ DKK/MW according to (11).

3) Algorithm: We can now describe the algorithm of operating the portfolio to both perform spot price optimization and also bid into the regulating power market. This is presented in Algorithm 2.

Algorithm 2: Regulating Power Algorithm

```
for k = 1, 2, ... do
    Collect data and perform intra-day optimization
    similar to Algorithm 1;
    Determine regulating power bids π_bid(u_act(k)) for
    all feasible multiples of 10 MW based on
    Problem (10) and (11) and place bids in market;
    if Current hour is between 11 a.m. and 12 p.m. then
        Purchase electricity similar to Algorithm 1;
    end
    if Activated for delivery u_act(k) then
        Let VPP steer portfolio's consumption to
        u(k) = u_sch + u_act(k);
    else
        Let VPP steer portfolio's consumption to
        u(k) = u_sch(k);
    end
    update scheduled consumption for next hour
    u_sch(k + 1);
end
```

V. Numerical Simulations

We perform two simulations to examine the presented control algorithms and use data from the "Styr din varmepumpe" project as a benchmark. Hereby we get a benchmark that corresponds to heat pumps operating using their default controllers which are only concerned with honoring a temperature setpoint and do not take spot prices into account. In both cases we consider a portfolio of $10,000$ heat pumps with heat capacity and drain rate as estimated in Sec. II-E and a nominal power consumption of 4 kW; further, an allowable temperature band of $\pm$ 2 °C around a setpoint of 21.5 °C is assumed. A sampling time of 5 minutes is used. The utilized data are:

- Spot price data from Nord Pool,
- Spot price predictions provided by [21],
- Outdoor temperature and daily loads from the "Styr din varmepumpe" project,
- Outdoor temperature predictions from the Danish Meteorology Institute.

We perform simulations for a full year and assume a liquid market where we do not affect the market prices[1].

A. Simulation 1: Spot Price Optimization

Algorithm 1 is utilized to operate the portfolio for spot price optimization for a full year. The resulting average temperature, power consumption, and costs are illustrated in Table I. In Fig. 5 the operation over 5 days is presented to illustrate the behavior of this controller. The top subplot shows the spot price predictions (red) and realizations (blue). The second subplot shows the power consumption of the heat pumps in the "Styr din varmepumpe' project (green) upscaled from the 130 available measurements to $10,000$ heat pumps. In the same subplot we show the power consumption when the portfolio is operated by the controller developed in this work (purple). Finally, the lower subplot shows the resulting average indoor temperature with the spot price controller operating the portfolio (purple) compared to the observed data for that period (green).

Together, the three subplots show the main result of the spot optimizing controller: that the developed controller is able to shift the main consumption to hours of low spot prices while keeping the temperature fluctuations in the same magnitude as the houses currently experience. It is important to notice that the aggregated portfolio is idealized as no delays, ramping constraints, etc. are included. This becomes evident in the idealized on/off characterized power consumption of the portfolio as illustrated in Fig. 5. Hence, the performance in the simulations will be higher than what we can expect by implementing the strategy.

Finally, notice that the spot price optimization will cause the natural smoothing of the heat pump portfolio consumption to disappear which may distribution grid congestion issues. This problem is, however, outside the scope of this work.

B. Simulation 2: Regulating Power Optimization

Algorithm 2 is utilized to operate the portfolio both for spot price optimization and for providing regulating power.

[1]It is difficult to predict how the market volatility will evolve the following years: increasing penetration of renewables and increasing oil prices suggests higher and more fluctuating prices while increasing volumes of flexibility and new transmission cables suggest the opposite.

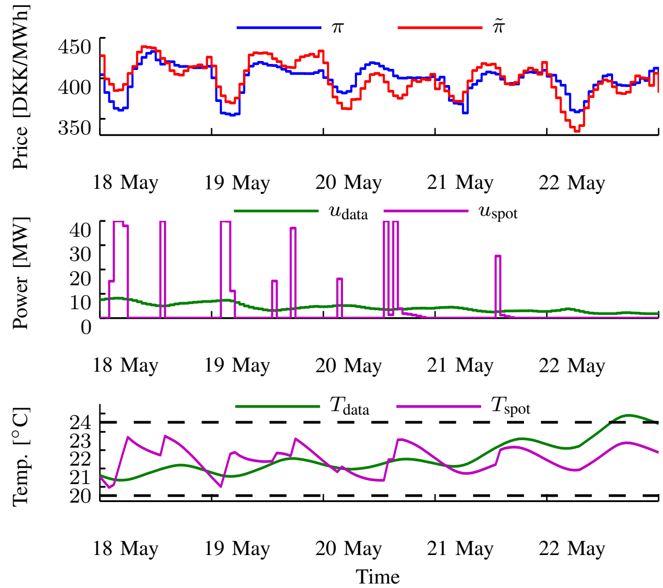

Fig. 5. Simulated heat pump portfolio optimized towards spot prices (purple) compared to upscaled real measurements (green).

		Data	Spot.	Reg.
Avg. temp.	[°C]	21.5	21.6	21.6
Avg. pwr.	[W]	732	737	744
Avg. spot.	[DKK/MWh]	356	282	270
Total cost per hp.	[DKK]	2.285	1.819	1.759
Savings	[%]	0	17.8	19.9

TABLE I

PERFORMANCE COMPARISON OF MEASUREMENTS AND THE TWO CONTROL STRATEGIES DEVELOPED IN THIS WORK.

Again, the end results are presented in Table I while a 5-day closeup is presented in Fig. 6. The top subplot shows the spot price realizations (blue) and the regulating power realizations (red). Following in the second subplot, the activations of regulating power is shown (yellow) along with the resulting consumption (purple) for the regulating power controller. It is observed, that the portfolio is activated for upward regulation in the cases where the RP price is significantly high; similarly the portfolio is activated for downward regulation when the RP price is significantly low. Note the portfolio is not able to perform upward regulation (decrease consumption) in the start of May 21st where the highest regulating power price is observed as the consumption already is scheduled to be zero and cannot be decreased further. Finally, the third plot again shows the resulting temperatures indicating that the fluctuations in the case of the regulating power controller is in the same order of magnitude as the observed indoor temperatures in the same period.

C. Comparison

In Table I we compare the two controllers with the data observed the same year. The first row shows that the average temperature based on measurements (data) is 21.5 °C, which therefore is used as a setpoint for the two controllers resulting

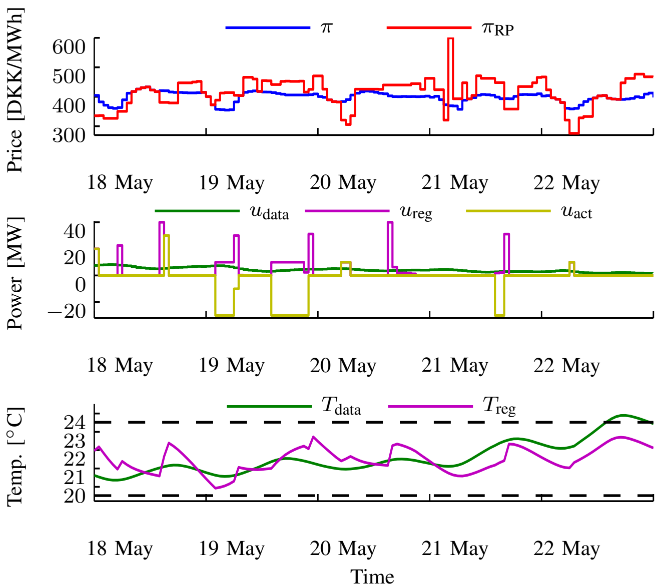

Fig. 6. Simulated heat pump portfolio optimized towards spot prices and the regulating power market (purple) compared to upscaled real measurements (green).

in almost identical average temperature. The next row shows the average power consumption which is measured to be 732 W while the two control strategies require a slightly higher power consumption. The average spot price based on the data is 356 DKK/MWh which is close to the yearly average spot price of 357 DKK/MWh – this is a result of the smooth power consumption of the heat pumps. By comparison, the spot price optimizing controller is able to lower this around 18 % while the controller that also bids into the regulating power market is able to save around 20 %. We observe that the annual savings per heat pump is in the magnitude of 470 DKK for spot price optimization but only additionally 60 DKK when also providing regulating power. We remind the reader that the simulated results are based on a somewhat idealized model; hence it should be expected that the savings when implementing this in real life will be lower. As described, actual spot price predictions are utilized for the simulation. By applying the actual spot prices, i.e. perfect predictions, we gain additionally 5 percentage points illustrating that the spot price predictions are of reasonable quality.

VI. Conclusion

In this work we showed how the consumption of a portfolio of heat pumps could be optimized towards spot price predictions day-ahead and adjusted intra-day to ensure comfort. Simulations were presented showing that savings in terms of reduced electricity costs in the magnitude of 18 % could be achieved compared to conventional heat pump operation. The controller was further extended to also bid into the regulating power market increasing the savings up to around 20 %. The savings 18 – 20 % correspond to around 500 DKK/year indicating that the equipment and installation costs must be very small to justify this type of optimization. Both controllers were designed based on the current regulations in the Nordic electricity market.

References

[1] Department of Energy, "Energy efficiency and renewable energy," U.S. Government, Department of Energy, Tech. Rep., 2008, DOE/GO-102008-2567.

[2] Danish Ministry for Climate, Energy and Buildings (Klima,- energi - og bygningsministeriet), "The energy- and climate goals of the danish government and results of the energy agreement 2020 (regeringens energi- og klimapolitiske mål – og resultaterne af energiaftalen i 2020)," 2012.

[3] I. A. Hiskens, "Load as a controllable resource for dynamic security enhancement," in *Proceedings of the IEEE Power Engineering Society General Meeting*, June 2006.

[4] Energinet.dk and Danish Energy Association, "Smart grid in Denmark 2.0," 2012, Report.

[5] Danish Energy Department, "Pilotprojekt for individuel rådgivning af boligejere med oliefyr vedrørende installation af varmepumpe," March 2012, Report.

[6] COWI (made for the Danish Energy Department), "Fremtidig metode til opgørelse af bestanden af varmepumper til opvarmning af helårshuse i danmark," November 2011, Report.

[7] F. Schweppe, R. Tabors, J. Kirtley, H. Outhred, F. Pickel, and A. Cox, "Homeostatic utility control," *Power Apparatus and Systems, IEEE Transactions on*, vol. PAS-99, no. 3, pp. 1151 –1163, may 1980.

[8] A. Arteconi, N. Hewitt, and F. Polonara, "Domestic demand-side management (dsm): Role of heat pumps and thermal energy storage (tes) systems," *Applied Thermal Engineering*, vol. 51, no. 12, pp. 155 – 165, 2013.

[9] C. Molitor, F. Ponci, A. Monti, D. Cali, and D. Muller, "Consumer benefits of electricity-price-driven heat pump operation in future smart grids," in *Smart Measurements for Future Grids (SMFG), 2011 IEEE International Conference on*, nov. 2011, pp. 75 –78.

[10] K. Hedegaard, B. V. Mathiesen, H. Lund, and P. Heiselberg, "Wind power integration using individual heat pumps – analysis of different heat storage options," *Energy*, vol. 47, no. 1, pp. 284 – 293, 2012.

[11] R. Halvgaard, N. Poulsen, H. Madsen, and J. Jorgensen, "Economic model predictive control for building climate control in a smart grid," in *Innovative Smart Grid Technologies (ISGT), 2012 IEEE PES*, jan. 2012, pp. 1–6.

[12] T. S. Pedersen, P. Andersen, K. Nielsen, H. L. Stærmose, and P. D. Pedersen, "Using heat pump energy storages in the power grid," in *IEEE Multi-Conference on Systems and Control, Denver, CO, USA*, Denver, CO, Sep. 2011.

[13] F. Tahersima, J. Stoustrup, S. Meybodi, and H. Rasmussen, "Contribution of domestic heating systems to smart grid control," in *Decision and Control and European Control Conference (CDC-ECC), 2011 50th IEEE Conference on*, dec. 2011, pp. 3677 –3681.

[14] M. U. Kajgaard, J. Mogensen, A. Wittendorf, A. T. Veress, and B. Biegel, "Model predictive control of domestic heat pump," in *American Control Conference, Washington, USA*, Jun. 2013.

[15] Energinet.dk (Danish TSO), "Fra vindkraft til varmepumper," July 2012, Report. [Online]. Available: http://www.styrdinvarmepumpe.dk

[16] P. P. Madsen, "Energy storage in ordinary houses. a smart grid approach," in *In proceedings of IFAC Conference on Power Plant and Power Systems Control, Toulouse, France*, Aug. 2012.

[17] B. Biegel, M. B. Madsen, P. Andersen, J. Stoustrup, and L. H. Hansen, "Lumped thermal household model," in *In proceedings The 4th European Innovative Smart Grid Technologies (ISGT) Conference, Copenhagen, Denmark*, October 2013.

[18] M. Petersen, L. H. Hansen, and T. Mølbak, "Exploring the value of flexibility: A smart grid discussion," in *IFAC Power Plant and Power Systems Control, Toulouse, France*, Aug. 2012.

[19] B. Biegel, P. Andersen, T. S. Pedersen, K. M. Nielsen, J. Stoustrup, and L. H. Hansen, "Smart grid dispatch strategy for on/off demand-side devices," in *In proceedings of IEEE European Control Conference, Zurich, Switzerland*, jul. 2013.

[20] Energinet.dk, "Ancillary services to be delivered in denmark – tender condition," Energinet.dk, Tech. Rep., 2011.

[21] T. Jónsson, "Forecasting and decision-making in electricity markets with focus on wind energy," Ph.D. dissertation, The Technical University of Denmark, 2010.