

Smart Grid Dispatch Strategy for ON/OFF Demand-Side Devices

Benjamin Biegel

Palle Andersen
Jakob Stoustrup

Tom S. Pedersen

Kirsten Mølgaard Nielsen
Lars Henrik Hansen

Abstract—We consider an aggregator managing a portfolio of runtime and downtime constrained ON/OFF demand-side devices. The devices are able to shift consumption in time within certain energy limitations. We show how the aggregator can manage the portfolio of devices to collectively provide upward and downward regulation. Two control strategies are presented enabling the portfolio to provide regulating power while respecting the runtime, downtime, and energy constraints of the devices. The first strategy is a predictive controller requiring complete device information; this controller is able to utilize the full flexibility of the portfolio but can only handle a small number of devices. The second strategy is an agile controller requiring less device information; this controller is able to handle a large number of devices but not able to utilize the full flexibility of the portfolio.

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]. Many actions have been taken from a political point to increase the penetration of renewables: in the US, almost all states have renewable portfolio standards or goals that ensure a certain percentage of renewables [2]. Similarly, the commission of the European Community has set a target of 20 % renewables by 2020 [3], while China has doubled its wind power production every year since 2004 [4]. In Denmark, the 2020 goals are 35 % sustainable energy over all energy sectors and 50 % wind power in electrical energy sector [5].

A major challenge arises when replacing central power plants with renewable energy sources: the central power plants do not only deliver power but also provide ancillary services to ensure a reliable and secure electrical power system. This includes frequency stability support, power balancing, voltage control, etc. When the conventional power plants are replaced with renewables such as wind turbines and photovoltaics, the ability to provide ancillary services in the classical sense disappears; the renewable energy sources will often fully utilize the available power and thus not be able to provide balancing ancillary services. Moreover, many renewable sources are characterized by highly fluctuating

power generation: they can suddenly increase or decrease production depending on weather conditions. These rapid production changes are not always predictable and can therefore have severe consequences for grid stability [6].

It is therefore evident that in a grid with high penetration of renewables, the need for balancing ancillary services will increase [7], [8]. As conventional power plants are pushed out gradually, alternative sources of ancillary services must be found. One of the approaches to obtaining alternative ancillary services is the *smart grid* concept, where demand-side devices with flexible power consumption take part in the balancing effort [9], [10]. The basic idea is to let an *aggregator* manage a portfolio of flexible demand-side devices and utilize the accumulated flexibility in the unbundled electricity markets on equal terms with conventional generators [11].

In this work, we consider the class of flexible consumption devices that only can be switched either ON or OFF possibly with minimum runtime and minimum downtime constraints. This covers a large range of different devices, for example thermal devices such as heat pumps, refrigeration and freezer systems, etc. We present two different direct load control strategies for enabling these devices to provide ancillary services: a predictive and an agile controller. The predictive controller requires full knowledge of all device parameters and provides an upper performance bound. This controller is, however, only able to handle a limited number of devices due to the computational burden. On the other hand, the agile controller is able to handle many devices and requires only little knowledge of the device parameters at the expense of not being able to utilize the full flexibility.

The paper structure is as follows. In Sec. II, the system architecture is presented. Following, in Sec. III, it is described how flexible ON/OFF consumers are able to deliver regulating reserves. In Sec. IV and Sec. V, the predictive and agile control strategies are presented. Numerical examples demonstrating these strategies are presented in Sec. VI. Finally, Sec. VII concludes the work.

II. SYSTEM ARCHITECTURE

In this section, we describe the overall relation between consumers, the aggregator, and the electricity markets.

A. The Aggregator as a Player in the Electricity Markets

We consider an unbundled liberalized electricity market 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 speaking,

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.

in an unbundled electricity market, TSOs do not own production units and must therefore procure ancillary services in the electricity markets to ensure system stability [12].

The aggregator is a legal entity able to enter into flexibility contracts with consumers. These contracts allow the aggregator to manage the consumers' flexible consumption; hereby, the aggregator is able to utilize the accumulated consumer flexibility to participate in the electricity markets. The flexible devices are managed by the aggregator through a technical unit often referred to as a virtual power plant (VPP). This setup is illustrated in Fig. 1 and inspired by [11]. In this work, we consider an aggregator utilizing the consumer flexibility to participate in the regulating power markets.

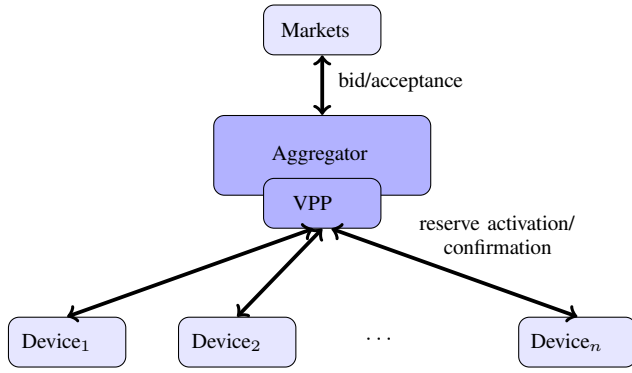


Fig. 1. Aggregator bidding in the electricity markets by managing n flexible ON/OFF devices through a VPP.

B. The Regulating Power Market

The suppliers can submit bids for upward regulation (increased production or reduced consumption) or downward regulation (decreased production or increased consumption) in the regulating power market. In the delivery hour, the TSO will activate the submitted bids if needs for upward or downward regulation occur.

The focus of this work is a dispatch strategy for a portfolio of devices activated for a given regulating power delivery. This means that we do not consider flexibility estimation, bidding strategies or similar in this work; rather, we describe how the portfolio should be managed to deliver regulating power once activated.

C. Demand-Side Devices as Power Reserves

Through a VPP, the aggregator manages a portfolio of runtime/downtime constrained ON/OFF devices with flexible power consumption. This covers a large class of devices; for example, thermal devices with large time constants such as electrically heated houses, refrigerations systems, water heaters, etc. [13]. The power consumption of these devices is not continuously adjustable; rather, the devices are either turned ON or OFF.

In order for consumption devices to provide ancillary services, they must be separated from and independent of

ordinary consumption and must be approved by a TSO as consumption that can be used as regulation reserves [14]. The hourly energy consumption of the portfolio must equal the energy bought at the spot-markets as long as the portfolio is not activated for power reserves. Upon activation, the hourly energy consumption of the portfolio must be adjusted accordingly.

In this work, we assume that the portfolio of devices under the jurisdiction of the aggregator indeed is approved by a TSO. Moreover, we assume that the necessary two-way communication between the aggregator's VPP and the ON/OFF devices exists as illustrated in Fig. 1.

III. REGULATING RESERVES VIA ON/OFF DEVICES

In this section we describe how a portfolio of ON/OFF devices collectively can deliver regulating power.

A. ON/OFF Consumption Devices

The VPP manages a portfolio of n flexible ON/OFF consumption devices. We assume that these devices can be modeled as energy storages with a time-varying drain. Denote the energy levels of the devices $x(k) \in \mathbf{R}^n$, the nominal device power ratings $p(k) \in \mathbf{R}^n$, and the drain rates $v(k) \in \mathbf{R}^n$, where k is the sample number using a sampling time T_s . We model device i as

$$x_i(k+1) = x_i(k) + T_s (p_i u_i(k) - v_i(k)) \quad (1)$$

$$x_i(1) = x_i^0 \quad (2)$$

where $x^0 \in \mathbf{R}^n$ represents the initial states of the devices and

$$u(k) \in \{0, 1\}^n \quad (3)$$

describes the state of each device: $u_i(k) = 1$ if device i is ON and $u_i(k) = 0$ if device i is OFF. The storage capacities are limited in size, thus we have

$$0 \preceq x(k) \preceq \bar{x} \quad (4)$$

where $\bar{x} \in \mathbf{R}^n$ describes the devices' energy limits and \preceq represents componentwise inequality. The interpretation of these limitations depends on the type of device. For space heating systems, space cooling systems, water heating systems, etc., the limits could represent an allowable temperature band [13].

The ON/OFF devices are furthermore characterized by *minimum runtime constraints* and *minimum downtime constraints* describing that once a device is turned ON, it must remain ON for a certain amount of time; similarly, that once a device is turned OFF, it must remain OFF for a certain amount of time. We use $\bar{r}, \underline{r} \in \mathbf{Z}_+^n$ to describe the runtime and downtime limits by letting \bar{r}_i be the minimum number of samples device i must remain ON once turned ON and by letting \underline{r}_i be the minimum number of samples device i must remain OFF once turned OFF:

$$u_i(k) - u_i(k-1) = 1 \implies u_i(k+1) = 1, \dots, u_i(k + \bar{r}_i - 1) = 1 \quad (5)$$

$$u_i(k) - u_i(k-1) = -1 \implies u_i(k+1) = 0, \dots, u_i(k + \underline{r}_i - 1) = 0 \quad (6)$$

where (5) describes the runtime constraint while (6) describes the downtime constraint. This type of constraints occur in many ON/OFF devices such as thermal systems where rapid switching of the compressor can damage the device or reduce lifetime significantly; likewise, rapid switching of for example heat pumps, will deteriorate performance.

B. Provisions of Regulating Reserves

The portfolio of ON/OFF devices is separated from and independent of regular consumption and is approved by the TSO as being able to deliver regulating reserves. The portfolio must therefore consume the electricity purchased at the spot-markets for each hour of the day. If the portfolio is activated for upward regulation, the consumption must be decreased accordingly the given hour; similarly, if activated for downward regulation, the consumption must be increased accordingly the given hour.

For simplicity, we make two assumptions that do not correspond to the regulating power markets. First, we assume that regulating power deliveries always are activated for exactly one delivery hour. In reality, however, the activation can be done for a shorter period and also within a delivery hour. Second, we assume that regulating power can be delivered in any manner, as long as the correct volume of energy is provided within the delivery hour. In reality, however, the regulating power must be provided at constant power.

Let l be the index of the operation hour and let the electricity purchased in the electricity markets for hour l be denoted $e_{\text{spot}}(l) \in \mathbf{R}$. Further, let $e_{\text{reg}}(l) \in \mathbf{R}$ denote the activated regulating power delivery in time period l and define $e_{\text{reg}}(l)$ as positive for upward regulation and negative for downward regulation in production terms. The energy reference $e_{\text{ref}}(l) \in \mathbf{R}$ for the portfolio is hereby given by

$$e_{\text{ref}}(l) = e_{\text{spot}}(l) - e_{\text{reg}}(l) \quad (7)$$

meaning that the portfolio of ON/OFF devices must consume the energy $e_{\text{ref}}(l)$ in hour l .

C. Regulating Power via ON/OFF Devices

As described, it is assumed that the power consumption within hour l can be chosen in any way as long as the energy reference $e_{\text{ref}}(l)$ is tracked according to the requirements. The portfolio is operated at a sampling rate T_s which is in the magnitude of minutes and thereby faster than the one-hour energy periods. The total power consumption of the portfolio at time sample k is denoted $p_{\text{out}}(k) \in \mathbf{R}$ and given by

$$p_{\text{out}}(k) = \mathbf{1}^T p(k) \quad (8)$$

where $\mathbf{1}$ is a vector with all components one. The hourly energy consumption $e_{\text{out}}(l) \in \mathbf{R}$ of the portfolio is found by integrating the portfolio power consumption $p_{\text{out}}(k)$ over each hour l :

$$e_{\text{out}}(l) = T_s \sum_{k=k_1(l)}^{k_2(l)} p_{\text{out}}(k) \quad (9)$$

where $k_1(l)$ and $k_2(l)$ indicate the first and last sample of the power consumption within hour l :

$$k_1(l) = \frac{3600}{T_s}(l-1) + 1, \quad k_2(l) = \frac{3600}{T_s}l \quad (10)$$

as $\frac{3600}{T_s}$ corresponds to the number of samples within one delivery hour. As the portfolio operates as a regulating reserve provider, it must be assured that the difference between the hourly energy consumption reference $e_{\text{ref}}(l)$ and the hourly energy consumption $e_{\text{out}}(l)$ is sufficiently small, hence we must minimize the tracking error $e_{\text{error}}(l) \in \mathbf{R}$ given by

$$e_{\text{error}}(l) = |e_{\text{ref}}(l) - e_{\text{out}}(l)|. \quad (11)$$

D. Summary of ON/OFF Device Characteristics

To visualize some of the concepts introduced in this section, we conclude with a small example. Consider a portfolio of 20 ON/OFF devices. The parameters of the portfolio are not important for this example but can be found later in (21). We assume that each device is operated by a local hysteresis controller on the form

$$u_i(k) = \begin{cases} 1 & \text{if } x_i(k) \leq 0 \\ 0 & \text{if } x_i(k) \geq \bar{x}_i \\ u_i(k-1) & \text{otherwise.} \end{cases} \quad (12)$$

The hourly power consumption $p_{\text{ref}}(k)$ and the energy consumption $e_{\text{out}}(l)$ are presented in Fig. 2 along with the energy levels of 5 of the devices, for a 10 hour period. For comparison, the figure also shows the nominal power consumption, given directly by summation of the drain rates $v(k)$, and the nominal energy consumption, given by the accumulated drain per hour. The nominal energy consumption could correspond to the expected energy consumption and therefore the electricity we have purchased at the spot-market.

An important point can be made from the energy delivery plot: large deviations between purchased electricity and actual consumption can occur due to the stochastic behavior of the ON/OFF devices. This is, however, not acceptable as a provider of regulating reserves. Therefore, a controller must manage the switching of the ON/OFF devices to assure that we indeed consume the purchased electricity. Further, this controller must adjust the consumption when activated for upward or downward regulation. Such controllers are developed in the following two sections.

IV. PREDICTIVE CONTROLLER SYNTHESIS

In this section, we design a predictive controller to manage the portfolio of runtime/downtime constrained ON/OFF devices. The controller relies on perfect information of the future load $v_i(k)$, the power rating p_i , and the capacity \bar{x}_i of all devices for a given horizon of L hours.

The predictive controller is not to be seen as an implementable strategy as it is not realistic to achieve such perfect information several hours ahead; further, the algorithm will show to be computationally heavy and thus only applicable for a limited number of devices. On the contrary, the predictive controller serves as an *upper performance bound*: it uses perfect information of the future conditions and finds a

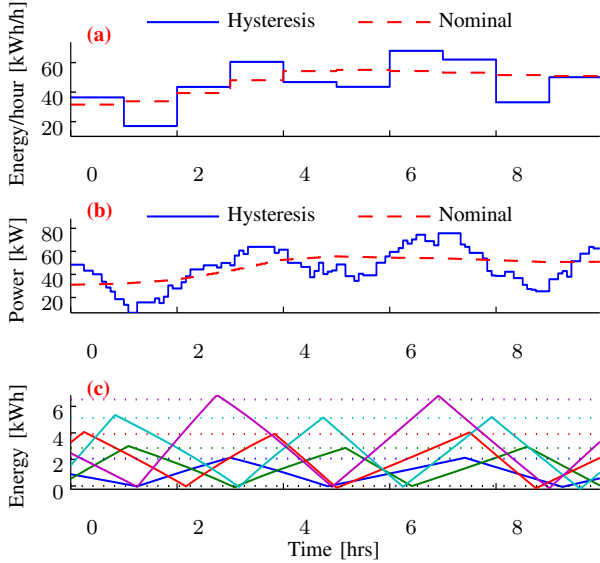


Fig. 2. Behavior of portfolio controlled with local hysteresis controllers compared to the nominal (predictable) behavior. Subplot (a): hourly energy consumption; subplot (b): power consumption; subplot (c): energy levels $x(k)$ for 5 of the 20 devices and corresponding energy limitations \bar{x} .

control strategy for the portfolio within the control horizon of L hours if feasible. This upper bound allows us to evaluate the performance of the agile controller which is presented in next section.

A. Optimization of ON/OFF Devices

The objective of the predictive control strategy is to determine the ON/OFF pattern of each device in the portfolio such that the hourly energy consumption of the portfolio tracks the energy reference while the devices honor runtime, downtime, and energy constraints. Define \mathcal{I} as the set of all devices, \mathcal{L} as the set of the L delivery hours, and \mathcal{K} as the set of time samples from the beginning of the first delivery hour to the end of the last delivery hour:

$$\mathcal{I} = \{1, \dots, n\}, \quad \mathcal{L} = \{1, \dots, L\}, \quad \mathcal{K} = \{1, \dots, K\}, \quad (13)$$

where $K = 3600L/T_s$ is the number of time samples within the horizon of L delivery hours.

Based on the previously introduced model, we can summarize the constraints and roughly formulate the predictive controller as follows.

$$\begin{aligned} & \text{minimize} && \sum_{l \in \mathcal{L}} e_{\text{error}}(l) \\ & \text{subject to} && \text{Eqs. 1-3, 5-6, } i \in \mathcal{I}, k \in \mathcal{K} \\ & && \text{Eqs. 4, 8, } k \in \mathcal{K} \\ & && \text{Eqs. 9, } l \in \mathcal{L} \end{aligned} \quad (14)$$

where the variables are $u(k), k \in \mathcal{K}$. Denote a solution to the optimization problem $u^*(k), k \in \mathcal{K}$. Note that $v(k), k \in \mathcal{K}$ is data to the problem meaning that perfect drain rate predictions are required to solve the problem. The solution $u^*(1), \dots, u^*(K)$ will describe how the devices can be switched ON and OFF such that the energy reference

$e_{\text{ref}}(1), \dots, e_{\text{ref}}(L)$ is tracked within the smallest average deviation, while runtime, downtime, and energy constraints are honored.

Notice that in this work we simply use Problem (14) in a static manner, i.e. we perform an open loop optimization over the whole horizon. This is done as the solution only is used as an upper performance bound based on perfect portfolio knowledge as previously described. The optimization problem could, however, be implemented in a receding horizon fashion where we optimize over a given horizon, apply the first element of the solution $u^*(1)$ to the plant, and then reoptimize the following sample after new information is obtained [15].

B. Binary Linear Optimization Problem

Problem (14) is a mixed integer linear optimization problem: dynamics (1), (2), state limitations (4), and conversion from power to energy (8), (9) are linear constraints. Further, the runtime and downtime constraints (5), (6) can be rewritten into linear constraints, see for example [16], [17]; similarly, the energy constraint (11) can be rewritten into linear constraints, see for example [18]. Finally, the ON/OFF constraint (3) makes the optimization problem binary (mixed integer). This mixed integer linear optimization problem resembles a unit commitment problem [19]. Generally speaking, this type of program is hard and can only be solved for a smaller number of devices and for shorter time horizons when using commercial optimization tools. For a larger number of devices, alternative methods are needed. As it is desired to be able to aggregate and control thousands of devices, alternative control strategies are needed. Therefore, an agile strategy is presented in the following section, relying on fast sorting algorithms rather than mixed integer optimization.

V. AGILE CONTROLLER SYNTHESIS

In this section, we present an agile controller that is able to overcome some of the limitations of the predictive strategy. By agile is meant a controller that seeks to maximize the agility of the portfolio by utilizing the least agile devices first [20]. In this context, an agile device corresponds to a device that is able to change state but does not demand a state change within a short horizon. Three major advantages are that the agile controller is able to:

- 1) handle a portfolio with a large numbers of devices
- 2) operate with little knowledge of the device parameters
- 3) handle devices that autonomously switch state.

These three features are necessary in a real life scenario where it is desired to aggregate thousands of small devices and where specific knowledge of each single device is difficult to assess and expensive to communicate. In the following, we describe an agile controller that satisfies the above three features.

A. Agile Controller Structure

The agile controller consists of two parts: a feedback controller and a dispatcher, see Fig. 3. Each device in




Fig. 3. Illustration of the agile strategy: a controller integrates the power error p_{error} between the power reference p_{ref} and the measured power consumption p_{out} to determine a control signal p_{ctrl} to the dispatcher. The dispatcher translates the control signal p_{ctrl} to ON/OFF signals u .

the portfolio operates by hysteresis control corresponding to (12) such that each device autonomously switches state if it reaches its energy limits. The feedback controller and dispatcher work on top of this: the power consumption of the portfolio p_{out} is measured and subtracted from a power reference p_{ref} resulting in a power error p_{error} which is the input to the feedback controller. The controller determines the control signal p_{ctrl} and feeds this signal to the dispatcher which translates p_{ctrl} to ON/OFF signals as described by u .

To emphasize the simplicity and robustness of the agile controller, we assume that the available information is very limited as described by the following:

- 1) the individual drain rates $v_i(k)$ are unknown,
- 2) the individual power ratings p_i are unknown; only the mean power rating $\tilde{p} = \frac{1}{n} \mathbf{1}^T p$ and the real-time total power consumption $p_{\text{out}}(k)$ are known,
- 3) the individual device states $x(k)$ and energy limits \bar{x}_i are unknown; only the real-time *state of charge* $s_i(k) = x_i(k)/\bar{x}_i$ is known.

Hereby, each device is only required to communicate the state of charge $s_i(k) \in \mathbf{R}$ to the VPP, which significantly reduces the communication flow. Further, the VPP tasks are simplified as it is sufficient to estimate the mean power rating $\tilde{p} \in \mathbf{R}$ instead of individual power ratings and drain rates. Note, however, that this relaxation requires that the power ratings, drain rates, and energy limits of the devices in the portfolio are within the same order of magnitude, i.e., this is not intended for a portfolio mixing for example large-scale CHP heating elements and small domestic heat pumps.

B. Energy Reference and Power Reference

The controller must ensure that the energy reference $e_{\text{ref}}(l)$ is tracked for each delivery hour l . This is done by translating the energy reference $e_{\text{ref}}(l)$ to a power reference $p_{\text{ref}}(k)$. The sampling rate of the power reference is in the magnitude of minutes and thus faster than the hourly sampling time of the energy reference. A freedom lies in the translation from energy reference to power reference. In this work, this freedom is utilized to make the power reference smooth over time such that fast power reference jumps are avoided. In this work we construct a filter that minimizes the two-norm of the change in power from sample to sample; however, other methods can be chosen. Later, in Fig. 5, this smoothing is seen when comparing the energy reference in subplot (a) with the power reference in subplot (b).

C. Feedback Controller

The feedback controller measures the power consumption of the portfolio $p_{\text{out}}(k)$ and compares this with the power

reference $p_{\text{ref}}(k)$ to determine the power error $p_{\text{error}}(k) \in \mathbf{R}$:

$$p_{\text{error}}(k) = p_{\text{ref}}(k) - p_{\text{out}}(k). \quad (15)$$

The feedback controller is implemented as a pure integral controller as our main objective is to follow the hourly energy reference $e_{\text{ref}}(l)$ which is exactly the integrated power. Further, the integral action will provide the necessary robustness to cope with the incomplete knowledge of the portfolio. The control signal $p_{\text{ctrl}}(k)$ is therefore simply found as

$$p_{\text{ctrl}}(k) = p_{\text{ctrl}}(k-1) + k_I p_{\text{error}}(k) \quad (16)$$

where $k_I \in \mathbf{R}$ is the integral gain.

D. Agile Dispatcher

The dispatcher translates the control signal $p_{\text{ctrl}}(k)$ into an ON/OFF signal $u(k)$ to the devices. The basic idea in the dispatcher is to *maximize the agility of the portfolio* meaning that the least agile devices should be activated first.

1) *Feasible Devices*: First, it is necessary to examine the subset of devices $\mathcal{I}_{\text{up}}(k) \subseteq \mathcal{I}$ able to provide upward regulation at time k and the subset of devices $\mathcal{I}_{\text{down}}(k) \subseteq \mathcal{I}$ able to provide downward regulation at time k . For a device to provide upward regulation at sample k , it must currently be in state ON and be able to switch to state OFF which requires that it has been ON for at least $\bar{\tau}_i$ samples. Similar argumentation can be made for a device to be able to provide downward regulation. Define the counters $c_i(k) \in \mathbf{Z}_+^n$ as

$$c_i(k) = \begin{cases} c_i(k-1) + 1 & \text{if } u_i(k) = u_i(k-1) \\ 1 & \text{otherwise} \end{cases} \quad (17)$$

such that $c_i(k)$ is the number of samples that device i has been in its current state $u_i(k)$. Then the sets

$$\mathcal{I}_{\text{up}}(k) = \{i \in \mathcal{I} | u_i(k-1) = 1, c_i(k-1) \geq \bar{\tau}\} \quad (18)$$

$$\mathcal{I}_{\text{down}}(k) = \{i \in \mathcal{I} | u_i(k-1) = 0, c_i(k-1) \geq \underline{\tau}\} \quad (19)$$

will describe the devices feasible for upward and downward regulation, respectively.

2) *Least Agile Device First*: The dispatcher is given the control signal p_{ctrl} and must determine if some of the devices in the portfolio must be switched from ON to OFF or vice versa. The agile dispatch strategy is to choose among the devices available for upward (downward) regulation the device closest to its upper (lower) bound. This strategy can be interpreted in different ways. One interpretation is that this is the strategy that will operate the devices as close as possible to the nominal hysteresis control strategy previously presented. Another interpretation is that this strategy maximizes the agility of the portfolio by always selecting the least agile device, see for example [20]. Finally, this strategy can

be interpreted as resembling the scheduling algorithm known as “least laxity first”, where the process with the smallest process slack time is activated first [21].

3) *Dispatch Algorithm*: Under the assumption that each device has a nominal power consumption given by \tilde{p} , the dispatcher expects the power output of the portfolio to equal $\tilde{p}\mathbf{1}^T u(k)$. Therefore, the dispatcher will choose to switch the state of $|n_{\text{sw}}(k)|$ devices at time k :

$$n_{\text{sw}}(k) = \text{round} \left(p_{\text{ctrl}}(k) / \tilde{p} - \mathbf{1}^T u_{\text{meas}}(k) \right) \quad (20)$$

where $u_{\text{meas}}(k) \in \mathbf{R}^n$ is the measured ON/OFF-state of the n devices at time k and $\text{round}(\cdot)$ is the “round to nearest integer” function. Note that it is necessary to measure the ON/OFF-states of the devices at time k as some devices may have reached the limitations and autonomously switched state according to the local hysteresis control. The dispatcher will switch $\max(0, n_{\text{sw}}(k))$ devices from OFF to ON and $\max(0, -n_{\text{sw}}(k))$ devices from ON to OFF at time k . Hereby the expected power output $\tilde{p}\mathbf{1}^T u(k)$ will get as close as possible to the control signal $p_{\text{ctrl}}(k)$.

In order to maximize the agility of the portfolio, we simply activate the device closest to its limit first. When $n_{\text{sw}}(k) < 0$, we need to decrease consumption and switch the $-n_{\text{sw}}(k)$ devices with the highest state of charge from ON to OFF; similarly, when $n_{\text{sw}}(k) > 0$, we need to increase consumption and therefore switch the $n_{\text{sw}}(k)$ devices with the lowest state of charge from OFF to ON. This way of finding $u(k)$ is described in Algorithm 1. The algorithm simply states that if $n_{\text{sw}}(k) < 0$, upward regulation is provided by selecting the $-n_{\text{sw}}(k)$ devices with highest state of charge from \mathcal{I}_{up} (if non-empty) and switching the state of these devices from ON to OFF, and vice versa for downward regulation.

Algorithm 1: Agile Dispatch Algorithm.

```

Initialize  $u(k) := u_{\text{meas}}(k)$ ;
collect control signal  $p_{\text{ctrl}}(k)$  and find  $n_{\text{sw}}(k)$  by (20);
for  $j = 1, \dots, |n_{\text{sw}}(k)|$  do
  update  $\mathcal{I}_{\text{up}}(k), \mathcal{I}_{\text{down}}(k)$  based on (18) and (19);
  if  $n_{\text{sw}}(k) > 0$  and  $\mathcal{I}_{\text{down}} \neq \emptyset$  then
    find the least agile device that can provide
    downward regulation:  $i := \text{argmin}_{i \in \mathcal{I}_{\text{down}}} s_i$ ;
    switch device ON:  $u_i(k) := 1$  ;
  else if  $n_{\text{sw}}(k) < 0$  and  $\mathcal{I}_{\text{up}} \neq \emptyset$  then
    find the least agile device that can provide
    upward regulation:  $i := \text{argmax}_{i \in \mathcal{I}_{\text{up}}} s_i$ ;
    switch device OFF:  $u_i(k) := 0$  ;
  end
end
apply  $u(k)$  to the portfolio;

```

E. Agile Controller Algorithm

We are now ready to describe the full algorithm of the agile controller, see Algorithm 2. As mentioned, the algorithm can be visualized as in Fig. 3.

Algorithm 2: Agile Controller Algorithm.

```

Initialize Determine the energy reference  $e_{\text{ref}}(l)$  by (7)
and convert to a smooth power reference  $p_{\text{ref}}(k)$ ;
for  $k = 1, \dots, 3600L/T_s$  do
  if  $e_{\text{reg}}(l)$  is changed by system operator then
    Update energy reference  $e_{\text{ref}}(l)$  by (7) and
    convert to a smooth power reference  $p_{\text{ref}}(k)$ ;
  end
  Measure current power consumption  $p_{\text{out}}(k)$  and
  determine  $p_{\text{error}}(k)$  according to (15);
  Obtain  $p_{\text{ctrl}}(k)$  by integration according to (16);
  Translate  $p_{\text{ctrl}}(k)$  to  $u(k)$  according to Algorithm 1;
  Dispatch the ON/OFF signals  $u(k)$  to the portfolio;
end

```

VI. NUMERICAL EXAMPLES

In this section, two numerical examples are considered: a small-scale example where the predictive strategy and the agile strategy are compared and a large-scale example that only the agile controller is able to handle. A sampling time $T_s = 5$ minutes is used.

A. Small-Scale Example

In this example, the portfolio consists $n = 20$ ON/OFF devices with parameters

$$\begin{aligned}
 p_i &\in [2, 9], & v_i(k) &\in [0, p_i], & [\text{kW}], \\
 \bar{x}_i &\in [1, 7], & x_i^0 &\in [1, \bar{x}_i], & [\text{kWh}], \\
 \bar{r}_i &= 6, & \mathcal{L}_i &= 6, & [\text{samples}].
 \end{aligned} \quad (21)$$

The parameters are selected such that the time to fully charge a device and to fully discharge a device are uniformly distributed in the interval 1 to 4 hours. Further, the load vectors $v(k)$ are chosen such that the total load curve $\mathbf{1}^T v(k)$ has the typical consumption shape with a morning and an evening peak as is visible in subplot (a) of Fig. 5. The runtime and downtime constraints are identical and equal to 6 samples corresponding to 30 minutes. These parameters are the same as used in the hysteresis controller case presented in Fig. 2.

A horizon of 10 hours is considered. The energy reference $e_{\text{ref}}(l)$ is set equal to the nominal energy consumption in the first and last 4 hours of the horizon. In hour 5 and 6, the energy consumption is set such that the maximum possible energy consumption is moved from hour 6 to hour 5; the volume of energy we can move is found via the predictive controller. The reference and the behavior of both the predictive and the agile controller are illustrated in Fig. 4.

This numerical example shows a number of interesting results. In subplot (a) we notice, that the agile controller is able to track the same reference as the predictive controller except for the two hours where load is shifted. In these two hours, the predictive controller is able to move at most 30 kWh while the agile controller is able to move at most 19 kWh corresponding to 63 % of the maximum possible.

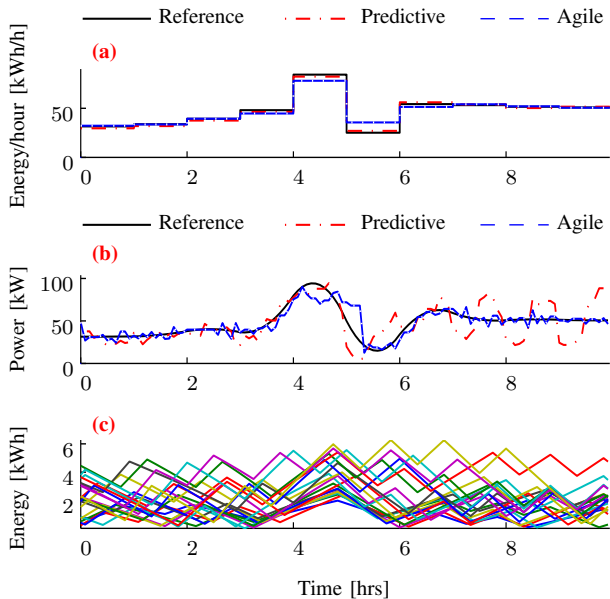


Fig. 4. Comparison of the behavior of the predictive controller and the agile controller tracking the hourly energy reference $e_{\text{ref}}(l)$. Subplot (c) shows the energy levels of the devices in the predictive controller case.

Subplot (b) of Fig. 4 shows the power reference found by smoothing the energy reference; further, this plot shows the power consumption in case of the agile controller and in case of the predictive controller. This plot shows an important difference between the two control strategies: the agile controller seeks to track this power reference, while the predictive controller does not consider the power reference; instead it directly considers the energy reference which in this case causes a fluctuating power consumption.

Finally, subplot (c) of Fig. 4 shows the energy levels of the 20 devices in the case of the predictive method. This plot illustrates the fundamental idea of moving consumption in time: almost all devices are ON in hour 5 to increase consumption lifting the energy levels of all devices; following, in hour 6, many of the devices are switched OFF again.

B. Large-Scale Example

We consider a portfolio of $n = 10,000$ ON/OFF devices with parameter distributions and runtime/downtime limitations similar to the previous example. A horizon of 24 hours is used. The predictive controller is not able to handle a portfolio of this size, therefore we only consider the agile controller. We consider an energy reference equal to the nominal power consumption; however, we shift a total of 14 MWh of consumption from the afternoon peak hours to the off-peak hours in the evening as depicted in subplot (a) of Fig. 5. The agile controller is able to track the reference with an error less than 0.5 MWh/h throughout all 24 hours.

In subplot (b) of Fig. 5, the nominal power consumption and the power reference are showed. Subplot (c) show the energy levels of 100 of the 10,000 devices in the portfolio. This figure shows the behavior of the controller: the overall energy levels in the devices are reduced in the afternoon peak hours to assure that the consumption is decreased as

required; following, the energy levels are restored when the energy consumption reference is increased in the evening.

Finally, subplot (d) shows the number of devices n_{avail} able to perform upward regulation and downward regulation:

$$n_{\text{avail}}(k) = \text{card}(\mathcal{I}_{\text{up}}(k)) + \text{card}(\mathcal{I}_{\text{down}}(k)) \quad (22)$$

where $\text{card}(\mathcal{X})$ denotes the cardinality of \mathcal{X} , i.e., the number of elements in \mathcal{X} . The plot shows that throughout the delivery period, there are between 2,000 and 7,500 available devices. Further, the plot shows that after the consumption of the devices is reduced at hour $l = 14$, the number of available devices decreases as the overall energy level must be kept low until the consumption of the devices is increased at hour $l = 19$.

VII. CONCLUSION

In this work we showed how a portfolio of runtime and downtime constrained ON/OFF devices with flexible power consumption can be managed to collectively provide a delivery of regulating power. We described how to track a regulating power reference based on a predictive controller requiring perfect information of the device parameters. The predictive strategy was able to fully utilize the flexibility of the devices and thereby provide the largest possible amount of regulating reserves. Following, we described how to track the regulating power reference based on an agile control strategy relying only on estimates of the device parameters. The agile controller was able to track an energy reference even for a large number of devices and with very limited knowledge of the portfolio parameters; however, it was not able to utilize the flexibility to the limits as the predictive controller.

REFERENCES

- [1] Department of Energy, "Energy efficiency and renewable energy," U.S. Government, Department of Energy, Tech. Rep., 2008, DOE/GO-102008-2567.
- [2] K. S. Cory and B. G. Swezey, "Renewable portfolio standards in the states: Balancing goals and implementation strategies," National Renewable Energy Laboratory, Tech. Rep., 2007, NREL/TP-670-41409.
- [3] Commission of the European Communities, "A european strategy for sustainable, competitive and secure energy," 2006, COM(2006) 105 final.
- [4] Y. Yan, C. Xia, Z. Song, and T. Shi, "Assessing the growth and future prospect of wind power in china," in *Electrical and Control Engineering (ICECE), 2010 International Conference on*, june 2010, pp. 3391–3395.
- [5] 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.
- [6] A. S. Chuang and C. Schwaegerl, "Ancillary services for renewable integration," in *Proceedings of Integration of Wide-Scale Renewable Resources Into the power Delivery System*, Jul. 2009.
- [7] P. Nyeng, "System integration of distributed energy resources - ict, ancillary services and markets," Ph.D. dissertation, The Technical University of Denmark, 2010.
- [8] P. E. Sørensen and M. Togeby and T. Ackermann and D. K. Chandrashekhara and J. P. F. Horstmann and H. Jóhannsson and A. H. Nielsen and P. Nyeng and T. W. Rasmussen and Z. Xu, "Ecogrid.dk phase 1 wp4 report: New measures for integration of large scale renewable energy," Forskningscenter Risø, Tech. Rep., 2008, DOE/GO-102008-2567.




Fig. 5. 24 hour simulation of a portfolio of 10,000 ON/OFF devices. Subplot (a) show the energy reference, the nominal consumption, and the response of the agile controller; similarly for the power in the subplot (b). Subplot (c) show the energy levels of 100 of the devices. Finally, subplot (d) shows the number of devices n_{avail} available for either upward or downward regulation.

- [9] I. A. Hiskens, "Load as a controllable resource for dynamic security enhancement," in *Proceedings of the IEEE Power Engineering Society General Meeting*, June 2006.
- [10] K. Moslehi and R. Kumar, "A reliability perspective of the smart grid," *IEEE Transactions on Smart Grid*, vol. 1, no. 1, pp. 57–64, 2010.
- [11] B. Biegel, P. Andersen, J. Stoustrup, L. H. Hansen, and D. V. Tackie, "Information modeling for direct control of distributed energy resources," in *Submitted for publication*, 2012.
- [12] Energinet.dk, "Regulation C2: The balancing market and balance settlement," Energinet.dk, Tech. Rep., 2011.
- [13] T. S. Pedersen, P. Andersen, K. Nielsen, H. L. Stærmosse, 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.
- [14] Energinet.dk, "Regulation C3: Handling of notifications and schedules," Energinet.dk, Tech. Rep., 2011.
- [15] J. Mattingley, Y. Wang, and S. Boyd, "Receding horizon control," *Control Systems, IEEE*, vol. 31, no. 3, pp. 52–65, 2011.
- [16] M. Carrion and J. Arroyo, "A computationally efficient mixed-integer linear formulation for the thermal unit commitment problem," *Power Systems, IEEE Transactions on*, vol. 21, no. 3, pp. 1371–1378, aug. 2006.
- [17] J. Arroyo and A. Conejo, "Optimal response of a thermal unit to an electricity spot market," *Power Systems, IEEE Transactions on*, vol. 15, no. 3, pp. 1098–1104, aug. 2000.
- [18] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
- [19] X. Guan, Q. Zhai, and A. Papalexopoulos, "Optimization based methods for unit commitment: Lagrangian relaxation versus general mixed integer programming," in *Power Engineering Society General Meeting, 2003, IEEE*, vol. 2, july 2003, p. 4 vol. 2666.
- [20] 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.
- [21] H. Aydin, R. Melhem, D. Mosse, and P. Mejia-Alvarez, "Optimal reward-based scheduling for periodic real-time tasks," *Computers, IEEE Transactions on*, vol. 50, no. 2, pp. 111–130, feb. 2001.