

A Virtual Battery Model for Packetized Energy Management*

Luis A. Duffaut Espinosa Adil Khurram Mads R. Almassalkhi†

Abstract—The goal of this paper is to develop a low-order model to represent the coordination of distributed energy resources based on concepts that make the internet work and is known as packetized energy management (PEM). The low-order model includes energy as a state variable together with dynamic opt-out constraints and internal packet request feedback, which in principle turns the model into a PEM virtual battery (PEM-VB) model. The paper focuses on a homogeneous aggregation of electric water heaters (EWHs) under PEM. It is shown that the bottom-up logic of the PEM-VB makes the system observable mainly due to the convenient feeding back of the number of packet requests through the communication channel established between the PEM coordinator and devices enrolled in the scheme. Without such extra information, the system loses the ability to observe the system's stored energy state. A procedure for computing the maximum and minimum energy bounds for the PEM-VB is developed. Moreover, the PEM-VB is explicitly used as the underlying model for an extended Kalman filter observer formulation with the purpose of estimating the energy stored in a simulated ensemble of agent-based EWHs under PEM. The use of PEM-VB is demonstrated in pre-positioning of flexible resources depending upon load forecasts. Finally, conclusions and future directions are provided.

Index Terms—Demand dispatch, distributed energy resources, packetized energy management, virtual battery

I. INTRODUCTION

The overarching goal in power systems operations is to deliver energy in an efficient, reliable, and economical manner. To achieve this, power system operators employ hierarchical primary, secondary, and tertiary frequency regulation schemes that uphold the century-old operating paradigm of *supply follows demand*. As electrification and decarbonization policies are pursued, the levels of variable, renewable generation will increase, which will require that power system operator think beyond *supply follows demand* to achieve the desired power systems objectives. This means one needs to consider the potential flexibility provided by, for instance, internet-enabled, connected, and responsive loads, which are part of the broad class of behind-the-meter (BTM) distributed energy resources (DERs) [1]. By effectively coordinating DERs and regulating their aggregate response, it is possible to turn the operating paradigm on its head and have (responsive) demand *follow* (variable) supply. In fact, it has been shown that BTM DERs can represent an inexpensive source of flexibility for the grid [2]–[4]. Thus, this paper focuses on a DER coordination scheme recently developed by the authors and their collaborators known as packetized energy management. This scheme enables information to be exchanged between the DERs and a coordinator, which informs the device to transition its operational state (e.g., consume, standby, and/or discharge), so that the coordinator can provide the appropriate control signals to control aggregate responsive net-demand (total consumption minus supply) as a source of grid flexibility. This control of the aggregate net-demand

is often with respect to tracking a power reference signal, such as that provided by a market-facing aggregator or independent system operator (ISO) and is similar to automatic generation control (AGC) and other ancillary services.

A natural question that arises when managing a fleet of DERs as a source of flexible demand is *what power reference signals can and cannot be tracked?* To tackle this question, it is important to have a notion of the current energy state or state of charge (SoC) of the aggregation. If that is the case, then the ensemble of DERs can be considered to be a virtual battery, e.g., [5]–[9]. In [5], a direct load control scheme is presented from which power/energy limits on the aggregate response of thermostatically controlled loads (TCLs) are defined as a function of ambient conditions. A formal virtual battery (VB) model is then developed for TCLs in [6] that couples the power response and fleet's SoC in a low-order, dynamic model. Similar methods are then adapted to develop VB models for large-scale fleets of TCLs and pool pumps in [8] that are utilized in a feedforward scheme together with myopic load coordination scheme. In [9], machine learning methods are used along with a direct load control scheme to parameterize a low-order VB model that works well for small-scale aggregations of less than 300 electric water heaters (EWHs) and air conditioners (ACs). The temporal flexibility or slack inherent to deferrable DERs, such as electric vehicles (EVs), are characterized as a VB in [7], where online control policies are developed to dispatch the DERs within local and VB constraints.

However, the parametrization of VBs often require DER coordination schemes that assume full or partial knowledge of the individual's DER energy or human comfort states. This information overhead may be perceived by the DER owner as violating their privacy. For example, with the charging profile of an electric vehicle owner, one can extract specific driving patterns such as miles driven by the owner in a day. Furthermore, constantly streaming information from large-scale fleets of DERs to a centralized coordinator limits scalability of the schemes due to communication rates and costs.

In this work, a low-order VB model is developed and parametrized based on a light-weight communication infrastructure that underpins the bottom-up, event-based DER coordination approach called packetized energy management [10]–[12]. PEM is illustrated in Fig. 1. In this scheme, information about the energy state of the DER ensemble is transmitted indirectly by energy request pings received at the load coordinator. These requests are received from each DER on an asynchronous, event-based timescale, which reduces communication overhead and enables a scalable implementation [13]. While PEM is applicable to DERs in general [11], the focus of this work is on EWHs.

The main contribution of this paper is the development of a predictive, low-order model of the aggregation of EWHs under PEM that is capable of providing bounds and estimates of the stored energy of the ensemble, while ensuring that local quality of service (QoS) constraints are satisfied. The PEM-VB model is shown to accurately estimate the SoC of a fleet of EWHs using an extended Kalman filter (EKF). Furthermore, an optimal

*This work was supported by the U.S. Department of Energy's ARPA-E award DE-AR0000694 and NSF grant CMMI-1839387.

The authors are with the Department of Electrical and Biomedical Engineering, University of Vermont, Burlington, VT 05405 USA

†M. Almassalkhi is co-founder of startup company Packetized Energy, which is commercializing aspects related to PEM.




Fig. 1: Closed-loop feedback system for PEM with P_{ref} provided by the grid or market operator and the aggregate net-load P_{dem} measured by the coordinator.

control problem (OCP) is presented that can be used to modify the aggregate power consumption of EWHs and ensure that the end-user remains unaffected by the coordination. The predictive PEM-VB model allows system operators to optimize use of flexible demand to counter the variability inherent to renewable generation.

The paper is organized as follows. Section II summarizes the PEM methodology and mechanism. In Section III, the low order model for PEM is developed as well as its energy bounds and conditions for local observability are provided. In Section IV, numerical simulations are presented including an extended Kalman filter formulation using the PEM-VB as the underlying model for an ensemble of EWHs that are simulated individually. Furthermore, an OCP formulation is presented that can be used in day ahead planning. The conclusions are given in the final section.

II. PRELIMINARIES ON PEM

Packetized energy management for diverse DERs has been presented in the author's earlier work [10]. State-bin transition models with high-dimensional state space were also developed along with control schemes discussed in [14]–[16]. This paper, however, focuses on a new, low-order, energy-based model for a fleet of PEM-enabled EWHs.

When applied to an ensemble of EWHs, PEM utilizes a probabilistic scheme based on the local dynamic state related to the energy content of each individual EWH. For instance, an electric water heater stores thermal energy that is proportional to the temperature difference between the hot water in the tank and ambient conditions. A notion of state of charge can then be associated with EWHs, which allows one to define the quality of service based on how close the energy level or temperature is to a customer-defined set point. For example, an EWH QoS is satisfied, if its temperature remains within a user's predefined temperature deadband. PEM's bottom-up approach (detailed in [10], [14]–[16]) is summarized as follows:

- i. An EWH estimates or measures its local SoC.
- ii. If the SoC is within a predefined range of comfort, the EWH, based on its SoC, probabilistically requests to consume energy from the grid at a fixed rate (e.g., 4kW) and for a pre-specified epoch (e.g., 5 minutes) to beget an energy packet (e.g., 0.33 kWh). If the SoC is too low, the EWH automatically opts out of PEM and charges. This mode is called OPT-OUT and is added to guarantee QoS. Once SoC is returned within limits, EWH opts back into PEM.
- iii. If a request is received, the coordinator either accepts or denies the EWH's packet request based on grid or market conditions. If the request is denied, go to i. If the request is accepted, consume the energy packet and then go to i.

Based on the previous description, an EWH can be in either one of three states: ON, OFF or OPT-OUT. Furthermore,

randomization is injected to the request rule based on the local SoC, which limits synchronization and promotes equitable access to the grid. Fig. 1 illustrates the closed-loop system under PEM [10].

The temperature dynamics of the n -th EWH is given by the simplified equation,

$$T_n[k+1] = T_n[k] + \Delta t \left(\frac{\eta_n P_{\text{rate},n}}{c_p \rho L_n} \phi_n[k] - \frac{\Delta T_n[k]}{\tau} - \frac{Q_n[k]}{c_p \rho L_n} \right), \quad (1)$$

where $c_p = 4.186 \frac{\text{kJ}}{\text{kg}^\circ\text{C}}$ is the specific heat constant, T_{amb} is the ambient temperature, $\tau = 150 \times 3600$ seconds is the standing loss time constant to ambient temperature and $\rho = 0.990 \frac{\text{kg}}{\text{L}}$ is the density of water when close to 50°C . For the n -th EWH, $P_{\text{rate},n}$ is the energy transfer rate, L_n is liters in the container tank, $\Delta T_n[k] := T_n[k] - T_{\text{amb}}$, T_n^{set} is the setpoint of the thermodynamic switch, η_n is the efficiency of the system, $Q_n[k]$ is the heat loss from the tank due to water usage and $\phi_n[k] \in \{1, 0\}$ is the state of the thermodynamic switch. Observe the three terms multiplying Δt in the right hand side of (1) correspond to the added power from the grid, standing losses, and end-user power consumption, respectively. The water usage profiles, Q_n are generated according to the procedure outlined in [10].

The rate governing how EWHs request packets is predefined by a so-called *probability of request* function. Such a function provides the probability at which an energy packet is requested by the n -th EWH. If the corresponding temperature of the EWH is $T_n[k]$ and the desired set-point is $T_n^{\text{set}} \in (\underline{T}_n, \bar{T}_n)$ over time k (for discretization time-step Δt), then the probability of request is given by a cumulative distribution function over the range of admissible dynamic states. For example, the exponential distribution gives

$$P_{\text{req}}(T_n[k]) := 1 - e^{-\mu(T_n[k])\Delta t}, \quad (2)$$

where $\mu(T_n[k]) > 0$ is a variable rate parameter dependent on the local dynamic state. For energy packet requests, it follows that

$$\mu(T_n[k]) = \begin{cases} 0, & \text{if } T_n[k] \geq \bar{T}_n \\ m_R \left(\frac{\bar{T}_n - T_n[k]}{\bar{T}_n - T_n^{\text{set}}} \right) \left(\frac{T_n^{\text{set}} - T_n}{T_n^{\text{set}} - \underline{T}_n} \right), & \text{if } T_n[k] \in (\underline{T}_n, \bar{T}_n) \\ \infty, & \text{if } T_n[k] \leq \underline{T}_n \end{cases}$$

where $m_R > 0$ [Hz] is a design parameter that defines the mean time-to-request (MTTR). Moreover, the energy at time k stored in the tank of the n -th EWH is obtained from

$$E_n[k] = c_p \rho L_n \Delta T_n[k]. \quad (3)$$

Energy can be translated into SoC, $\zeta_n[k] = \frac{E_n[k] - E_{\text{min},n}}{E_{\text{max},n} - E_{\text{min},n}}$ where $\zeta_n \in [0, 1]$ and $E_{\text{max},n}, E_{\text{min},n}$ are maximum and minimum limits on energy respectively. The procedure to obtain limits on energy are presented in the next section.

III. VIRTUAL BATTERY MODEL FOR EWHs

In this section, a notion of average energy or SoC of an aggregation of EWHs is developed. By doing so, a fleet of EWHs can be abstracted into a PEM *virtual battery*. Furthermore, charging and discharging characteristics of this PEM-VB are obtained based on the total power consumption of the fleet and the requests received by the coordinator. The desired low order virtual battery model then consists of three states; (i) average energy of the fleet, (ii) total number of EWHs consuming power and (ON) (iii) total number of EWHs in OPT-OUT. Total number of EWHs in the ON and OPT-OUT modes are required to capture the charging characteristics of PEM-VB. The discharge characteristics are governed by the end-use process and is discussed next.

A. End-use model of hot water extraction

Local end-user events of the n -th EWH correspond to hot water usage that result in heat loss from the tank, denoted by Q_n in (1). To obtain a low-order model, the aggregate statistics of the fleet must be required. For this purpose, it was shown in [16] that these user-driven events can be modeled by the so-called *Poisson rectangular pulses* (PRPs) [17]. To clarify notation, the subscript n is omitted hereafter and focus is placed on a single EWH.

Assume that there exists an appropriate probability space $(\Omega, \mathcal{P}, \mathcal{F})$, where Ω is the set of events, $\mathcal{F}$ a filtration, and P the probability measure of elements in $\mathcal{F}$. The stochastic differential model for a PRP, $w(t)$ is given as [17],

$$dw(t) = (v(t) - w(t))dN_1(t) - w(t)dN_2(t), \quad (4)$$

where N_1 (N_2) is an independent, stationary Poisson point process with constant rate parameter λ_1 (λ_2), representing the initiation (conclusion) of a random end-user event and $v(t)$ is a random variable independent of N_1 and N_2 that describes the intensity of the end-user event. For an EWH, $v(t)$ describes the power used for increasing the tank temperature and is considered here exponentially distributed with mean λ .

Denote the expected value of the random process w as $\bar{w}(t) := \mathbb{E}[w(t)]$. Due to the independence of the processes ΔN_1 , ΔN_2 and $v(t)$ in time, one can compute the expected end-user event as

$$\frac{d\bar{w}(t)}{dt} = (\bar{v}(t) - \bar{w}(t))\lambda_1 + \bar{w}(t)\lambda_2. \quad (5)$$

The solution of (5) when $w(0) = 0$ is

$$\bar{w}(t) = \mathbb{E}[v] \frac{\lambda_1}{\lambda_1 + \lambda_2} (1 - \exp(-(\lambda_1 + \lambda_2)t))$$

The expected event reaches steady state as t goes to infinity. Hence, the mean of end-user event in steady state is

$$\bar{w}_{\text{sst}} := \lim_{t \rightarrow \infty} \bar{w}(t) = \frac{\mathbb{E}[v]\lambda_1}{\lambda_1 + \lambda_2}. \quad (6)$$

The next theorem describes the probability distribution of these events as the number of devices increases. A reasonable assumption is that the end-user event for each EWH are independent and identically distributed (i.i.d.) random processes. One can then obtain the steady state statistics of the aggregation of the process w .

Theorem 1: The aggregation of individual end-user events, w , is distributed in steady state as $\mathcal{N}(\bar{w}_{\text{sst}}, \sigma_w/\sqrt{N_e})$, where N_e is the total number of end-user event processes and $\bar{w}_{\text{sst}}$ and σ_w are the corresponding expected value and standard deviation of the process w in steady state.

Proof: See [15]. ■

Example 1: If $v \sim \exp(\lambda)$, then $\sigma_w = \lambda\sqrt{2p-p^2}$ and $\bar{w}_{\text{sst}} = \lambda p$, where $p := \frac{\lambda_1}{\lambda_1 + \lambda_2}$. The average of 2,000 water usage profiles generated from (4) with $\lambda = 2.1$ liters per minute, $\lambda_1 = 1/3600 \text{ sec}^{-1}$ and $\lambda_2 = 1/800 \text{ sec}^{-1}$ results in the mean and standard deviation of 0.3868 and 0.0382 respectively, whereas for the one generated using Theorem 1 gives a mean of 0.3818 and standard deviation of 0.0369. □

This example shows that the average effect produced by the fleet of EWHs driven by i.i.d. PRPs ($w(t)$) is the same as the aggregate driven by the i.i.d. process $\gamma(t)$ distributed according to $\mathcal{N}(\bar{w}_{\text{sst}}, \sigma_w)$. Finally, in case of a fleet of EWHs, the hot water usage varies depending upon the time of the day. For example, the water usage is higher in the morning due to people taking showers than the afternoon, but remains relatively constant within a period of 1 to 2 hours. Therefore, a full day can be divided into periods

of relatively constant water usage. Historical data is then used to obtain the statistics for each period as discussed in [18].

With the aggregate statistics of the hot-water usage now known, the PEM-VB is formally derived next.

B. Virtual Battery Model

The simple average of an homogeneous population of N EWHs temperatures is $T_{\text{avg}} = \sum_{n=1}^N T_n/N$. From (1) and provided that Q is the discrete-time equivalent of γ , we have,

$$T_{\text{avg}}[k+1] = \left(1 - \frac{\Delta t}{\tau}\right) T_{\text{avg}}[k] + \frac{\Delta t T_{\text{amb}}}{\tau} - \frac{\Delta t \sum_{i=1}^N (Q_i[k])}{N c_p \rho L} + \frac{\eta \Delta t P_{\text{rate}} \sum_{i=1}^N \phi_i[k]}{N c_p \rho L}.$$

From Section III-A and the law of large numbers, one has that $\sum_{i=1}^N Q_i[k]/N$ becomes $\mu_Q := \bar{w}_{\text{sst}}$, which results in

$$T_{\text{avg}}[k+1] = \left(1 - \frac{\Delta t}{\tau}\right) T_{\text{avg}}[k] + \frac{T_{\text{amb}} \Delta t}{\tau} - \frac{\Delta t}{c_p \rho L} (\mu_Q) + \frac{\eta \Delta t P_{\text{rate}} (N_{\text{on}}[k] + N_{\text{opt}}[k])}{c_p \rho L N}, \quad (7)$$

where $\sum_{i=1}^N \phi_i[k]$ is the total number of EWHs that are either in ON state (N_{on}) and in OPT-OUT state (N_{opt}), that is,

$$\sum_{i=1}^N \phi_i[k] = N_{\text{on}}[k] + N_{\text{opt}}[k]. \quad (8)$$

Each EWH then requests probabilistically based on $P_{\text{req}}(T_{\text{avg}}[k])$, which means that the total number of requests received at the coordinator during the interval $[k-1, k]$ is $x_r[k] := P_{\text{req}}(T_{\text{avg}}[k])(N - N_{\text{on}}[k] - N_{\text{opt}}[k])$. Define $\beta[k]$ as the proportion of the $x_r[k]$ requests that are accepted by the coordinator at time k and $\beta^-[k]$, the proportion of EWHs in the ON state that complete their packet at time k . The dynamics of the number of ON EWHs is then identified as

$$N_{\text{on}}[k+1] = N_{\text{on}}[k] - \beta^-[k] N_{\text{on}}[k] + \beta[k] x_r[k]. \quad (9)$$

The fact that $P_{\text{avg}}(T_{\text{avg}})$ is nonlinear makes (9) nonlinear. Also, (9) assumes that even though some EWHs have opted out of the PEM scheme, they still notify the coordinator whenever they leave and re-enter the scheme. Thus, we have

$$N_{\text{opt}}[k+1] = N_{\text{opt}}[k] + \xi_{\text{opt-out}}[k] - \xi_{\text{opt-in}}[k], \quad (10)$$

where $\xi_{\text{opt-out}}$ and $\xi_{\text{opt-in}}$ are the number of devices opting out and opting in PEM during Δt , respectively. Thus, under the ON and OPT-OUT dynamics, the total number of devices, N , remains constant, as expected.

For each Δt , the control input $\beta[k]$ is determined from

$$\beta[k] = (P_{\text{ref}}[k] - P_{\text{dem}}[k]) (P_{\text{rate}} x_r[k])^{-1}, \quad (11)$$

where, $P_{\text{dem}}[k] := P_{\text{rate}} (N_{\text{on}}[k] + N_{\text{opt}}[k])$, whereas the input $\beta^-[k]$ is determined by introducing a packet-duration-timer state that runs at the coordinator. The timer accumulates the number of requests accepted at each timestep and moves them forward in time, deterministically, until the packet is consumed, which transitions the EWH's mode from ON to OFF. The packet duration timer dynamics is given by

$$\tau[k+1] = M\tau[k] + e_1 \beta[k] x_r[k], \quad (12)$$

where $\tau \in \mathbb{R}^{n_p}$, n_p is the number of timer steps given by $n_p = \lceil t_p / \Delta t \rceil$, t_p is the packet duration, $e_i \in \mathbb{R}^{n_p}$ is an elementary vector whose i -th component is 1, $M \in \mathbb{R}^{n_p \times n_p}$ is such that its first lower diagonal comprised of ones and zero everywhere

else. The packet expiration rate is obtained from the last state of the timer as $\beta^-[k] = e_{n_p}^\top \tau[k]$. Finally, converting temperature to energy using (3) gives the desired three state $(E, N_{\text{on}}, N_{\text{opt}})$ and 4-input $(\beta, \beta^-, \xi_{\text{opt-out}}, \xi_{\text{opt-in}})$ PEM-VB.

The PEM-VB achieves its maximum energy limit and is “fully charged” when all requests are accepted ($\beta = 1$). Similarly, the PEM-VB becomes fully discharged when all requests are rejected ($\beta = 0$). Figs. 2-3 show charge and discharge cycles of the PEM-VB for 1000 EWHs having packet duration of 5 minutes, rated power 4.5kW and ambient conditions $T_{\text{amb}} = 14^\circ\text{C}$. This shows that the usual notions of state of charge can be associated to a fleet of EWHs operating under PEM. Note that as PEM-VB discharges, the number of OPT-OUTs (N_{opt}) increases and is an indication of decreasing QoS. Information on EWHs in OPT-OUT mode is obtained at the coordinator from measurements of $(\xi_{\text{opt-in}}, \xi_{\text{opt-out}})$.




Fig. 2: SoC during a 30 hour dis/charging simulation showing a first-order response. Highlighted region shows PEM-VB charging.




Fig. 3: N_{on} , N_{off} and N_{opt} during a 30 hour charging/discharging simulation. The total number of devices is always preserved.

C. Virtual battery energy limits

The limits of the PEM-VB energy are computed by analyzing the steady-state conditions for salient inputs β and β^- .

1) *Upper energy limit:* From (3), (7) and (9),

$$E_{\text{max}} = \lim_{k \rightarrow \infty} c_p \rho L (T_{\text{avg}}[k] - T_{\text{amb}}) |_{\beta=1, \beta^-=1/n_p},$$

provides the energy upper limit. In the authors' earlier work [12], it was shown that $\beta^- \approx 1/n_p$ when all requests are accepted ($\beta = 1$). Also, from the fact that $N_{\text{opt}}[\infty] = 0$ and $N_{\text{on}}[k+1] = N_{\text{on}}[k]$ when $\beta = 1$, one has that

$$N_{\text{on}}[\infty] = \frac{\beta P_{\text{req}}(T_{\text{avg}}^*) N}{\beta^- + \beta P_{\text{req}}(T_{\text{avg}}^*)}.$$

Replacing $N_{\text{on}}[\infty]$ into (7) gives $P_{\text{req}}(T_{\text{avg}}^*) = 0.00828$ since $(1 - \Delta t/\tau) \approx 1$, which corresponds to $N_{\text{on}}[\infty] = 142$. $N_{\text{on}}[\infty]$ here matches the value provided by the simulation in Fig. 3 and

amounts to $T_{\text{avg}} = 54.23^\circ\text{C}$. Using (3), the maximum energy value is $E_{\text{max}} = 11.58 \text{ MWh}$. To determine the available energy from the ensemble, lower energy bound is required.

2) *Lower energy bound:* The lower limit of the PEM-VB is achieved at,

$$E_{\text{min}} = \lim_{k \rightarrow \infty} c_p \rho L (T_{\text{avg}}[k] - T_{\text{amb}}) |_{\beta=0, \beta^-=0}.$$

It should be noted that there are no EWHs in the ON state ($N_{\text{on}}[\infty] = 0$) when $\beta = 0$, that is all requests are denied. Therefore, EWHs transition between OFF and OPT-OUT, however, $N_{\text{opt}}[\infty]$ and hence the number of EWHs in OFF state are constant meaning that $\xi_{\text{opt-out}}$ and $\xi_{\text{opt-in}}$ are equal, as shown in Fig. 3. This type of behavior corresponds to a 2-state Markov chain in steady state and its statistics are then used to determine $N_{\text{opt}}[\infty]$ as shown next.

Let z_{opt} and z_{off} be the two states (OPT-OUT and OFF) of the Markov chain, then its dynamics are given by,

$$\begin{pmatrix} z_{\text{off}}[k+1] \\ z_{\text{opt}}[k+1] \end{pmatrix} = \begin{pmatrix} 1-p_1 & p_2 \\ p_1 & 1-p_2 \end{pmatrix} \begin{pmatrix} z_{\text{off}}[k] \\ z_{\text{opt}}[k] \end{pmatrix}. \quad (13)$$

Clearly, this model permits a non-trivial unique stationary distribution for probabilities p_1 and p_2 . The stationary distribution provides the *averaged occupancy* of each state, which is the percentage of EWHs that, on average, are either in z_{off} or z_{opt} . That is, if $\pi = (\pi_{\text{off}}, \pi_{\text{opt}})^\top$ denotes such stationary distribution, then $\pi_{\text{off}} = \frac{p_2}{p_1 + p_2}$ and $\pi_{\text{opt}} = \frac{p_1}{p_1 + p_2}$ provides occupancy of z_{off} and z_{opt} respectively. Now, the transition probabilities p_1, p_2 are determined from mean sojourn time of the Markov chain. Let t_{opt} be the number of time steps an EWH spends in OPT-OUT, on average, before transitioning to OFF, then t_{opt} is the *mean sojourn time* of z_{opt} . Similarly, t_{off} is the mean sojourn time of z_{off} .

Consider first t_{opt} , which physically represent the time taken by an EWH that starts in OPT-OUT at $\underline{T}$ to reach T_{pem} (the pre-defined temperature at which EWH re-enter the PEM scheme). Further, denote with $\psi(z_{\text{opt}})$, the expected number of time steps needed to reach state z_{off} given that one starts in z_{opt} and $\psi(z_{\text{off}})$ if one were to start in state z_{off} . Forcing the state z_{off} to be absorbing, it follows that $\psi(z_{\text{off}}) = 0$ and $\psi(z_{\text{opt}}) = 1 + (1 - p_2)\psi(z_{\text{opt}})$, which provides $\psi(z_{\text{opt}}) = 1/p_2$. Thus, p_2 describes how many time steps on average EWH stays in OPT-OUT, and the actual expected time spent in OPT-OUT is trivially $t_{\text{opt}} = \psi(z_{\text{opt}})\Delta t = \Delta t/p_2$. Similarly, one can obtain the expected time spent in OFF as $t_{\text{off}} = \Delta t/p_1$. The sojourn times $t_{\text{opt}}, t_{\text{off}}$ can be obtained by solving the continuous version of (1), $\dot{T} = AT + D$ that has the solution

$$T(t^*) = (e^{At}(D + AT_0) - D)/A, \quad (14)$$

where $A = -1/\tau$, $D = T_{\text{amb}}/\tau - \mu_Q/(c_p \rho L) + \phi \eta P_{\text{rate}}/(c_p \rho L)$, $\phi \in \{0, 1\}$ and T_0 is the initial temperature.

For the simulation in Fig. 3, $t_{\text{opt}} = 9.27 \text{ min}$ is obtained from (14) when $T_0 = \underline{T}$ and $\phi = 1$. Similarly, $t_{\text{off}} = 57.4 \text{ min}$ for $T_0 = T_{\text{pem}}$, $\phi = 0$. Finally $N_{\text{opt}}[\infty] = \pi_{\text{opt}} N = 139$ that matches the total number of OPT-OUTs observed in Fig. 3, that is 140 which amounts to modeling error of $< 1\%$ for a population of 1000 EWHs. One can now solve for T_{avg} in (7) to obtain minimum temperature that the PEM-VB can achieve: 49.8°C . Applying (3), the minimum energy of the PEM-VB is $E_{\text{min}} = 10.30 \text{ MWh}$.

The stored energy on the PEM-VB is $E_{\text{max}} - E_{\text{min}} = 1.28 \text{ MWh}$, which is close to the 1.275 MWh obtained from the simulation. Thus, the PEM-VB model captures energy bounds well. Next, we examine how to estimate the PEM-VB states from measured outputs.

D. Local Observability

This section focuses on observability of the PEM-VB model. Specifically, it is shown that the energy of the system is locally observable exactly due to the internal feedback from the requests, which signifies the importance of the bottom up approach and the value of the overhead created by the bi-directional, but sparse communications between EWHs and the coordinator.

The low-order PEM-VB model enables analysis of *strong local observability*. This is a consequence of the implicit function theorem [19] and mimics the process used for linear control systems. Consider the system

$$x[k+1] = f(x[k], u[k]), \quad y[k] = h(x[k]), \quad (15)$$

where f and h are smooth functions. The time increments of the output can be computed from (15) as

$$\begin{pmatrix} y[k] \\ y[k+1] \\ \vdots \\ y[k+n-1] \end{pmatrix} = \begin{pmatrix} h(x[k]) \\ h(f(x[k], u[k])) \\ h(f(f(x[k], u[k]), u[k+1])) \\ \vdots \\ h(\mathcal{L}^{n-1} f(x[k], u[k])) \end{pmatrix}, \quad (16)$$

where $\mathcal{L}^{n-1} f$ is the n -th iterative substitution of f . Then (16) is written as $\mathcal{Y}_k = \mathcal{H}(x[k], u[k])$. If $\mathcal{H}$ is invertible, then the system is observable and one can solve locally for $x[k]$, which requires that the Jacobian of $\mathcal{H}$, say $O_{\mathcal{H}}$, is invertible. Hence, if $\det(O_{\mathcal{H}}) \neq 0$ then (15) is called *strong locally observable*.

The PEM-VB states are $x[k] = (E[k], N_{\text{on}}[k], N_{\text{opt}}[k])^\top$, its inputs $u[k] = (\beta[k], \beta^-[k], \xi_{\text{opt-out}}, \xi_{\text{opt-in}})$, and the components of $f = (f_1, f_2, f_3)^\top$ and h in (15) are

$$f_1(x, u) = a_1 x_1 - \Delta t \left(\frac{\eta P_{\text{rate}}(x_2 + x_3)}{N} - \mu_Q \right), \quad (17a)$$

$$f_2(x, u) = (1 - u_2 - u_1 P_{\text{req}}(x_1)) x_2 - u_1 P_{\text{req}}(x_1) (x_3 - N), \quad (17b)$$

$$f_3(x, u) = x_3 + u_3 - u_4, \quad (17c)$$

where $a_1 := (1 - \Delta t / \tau)$ and

$$h(x) = (P_{\text{rate}}(x_2 + x_3), P_{\text{req}}(x_1)(N - x_2 - x_3))^\top \quad (18)$$

Example 2: Computing the Jacobian $O_{\mathcal{H}}$ for f and h in (17) and (18), respectively, gives $\text{rank}(O_{\mathcal{H}}) = 3$ for all inputs, except if $\beta[k] + \beta^-[k] = 0 \forall k$, which occurs regularly as shown in the numerical simulations provided in the upcoming section (e.g., see Fig. 4) and yields $\text{rank}(O_{\mathcal{H}}) = 2$. In this case, rank is lost since no EWHs are in the ON state ($x_2 = 0, \forall k$). However, the dynamics (9) for x_2 then cancel and the EWHs cycle between OFF and OPT-OUT states. Hence, it is still possible to estimate x_1 since devices in the OFF-state ensure condition $x_r > 0$ and the coordinator has information from EWHs that temporarily opt-in and opt-out of PEM ($\xi_{\text{opt-in/out}}$).

However, when information on x_r is unavailable to the coordinator, $O_{\mathcal{H}}$ losses rank when $\beta[k] + \beta^-[k] = 0$. Specifically, when $\beta = 0$ one has that $\text{rank}(O_{\mathcal{H}}) = 2$. The situation is even worse when $\beta = \beta^- = 0$ providing $\text{rank}(O_{\mathcal{H}}) = 1$. The latter case makes the system unable to estimate the energy of the PEM-VB. In other words, while many VB approaches assume full information and control of every DER state that is generally not possible in practice due to privacy and communication constraints. In that context, PEM's anonymous, light-weight, and probabilistic request mechanism satisfies both communication and privacy limitations due to its asynchronous, internet-like protocols. $\square$

IV. NUMERICAL RESULTS

First, the tracking limits of the PEM-VB are considered. Consider a fleet of 1000 EWHs having packet duration of 5 minutes, rated power of 4.5kW and $\Delta t = 15s$. The first numerical simulation aims to illustrate the behavior of the PEM-VB when tracking a reference power signal. The blue curve at the top plot of Fig. 4 was selected. The signal comprised of a scaled and zero mean AGC signal from [20] that has been placed around two power levels (300 and 1000 kW) for the purpose of letting the system reach its energy limits while tracking. Observe that the SoC plot saturates reaching either 0% or 100% exactly when the PEM-VB can no longer track the reference. This is clearly reflected in the saturation of β and β^- . Furthermore, in the intervals [8,11] and [21,24], the PEM-VB reaches 0% SoC and is unable to track the reference. In this case the inputs $\beta = \beta^- = 0$, the system is cycling between OFF and OPT-OUT states and the dynamics of the ON states becomes irrelevant. Contrarily, in intervals [1,4], [14,18] and [27,30], $\beta = 1$ and $\beta^- = 1/n_p$ the system reaches 100% SoC and cannot track the reference due to the lack of flexibility in the PEM-VB. Staying around the set point (50%) allows the PEM-VB to track the reference and also corresponds to the nominal power of the system [12].




Fig. 4: Tracking when reaching upper and lower energy boundaries. System is unable to track once energy limits are reached.




Fig. 5: Tracking an AGC signal from ISO-NE.

A. Estimation of SoC using PEM-VB

The next simulation involves the estimation of the SoC of an ensemble of EWHs via a simple extended Kalman filter (EKF) formulation. Here an agent-based simulation of a fleet of 1000 EWHs with packet duration of 5 minutes, packet power of 4.5kW, and $\Delta t = 15$ seconds is performed (i.e., 1000 individual models based on (1) are simulated and aggregated). Here the standard EKF formulation is used on (17) assuming additive Gaussian noise for the state and measurement update equations.

The estimation of the SoC is now performed online from the measurements of demand power, P_{dem} , and the number

of requests, x_r from a fleet of 1000 EWHs operating under PEM. The fleet is tracking a scaled AGC signal centered around 500 kW and the result is shown in Fig. 5. It is found that the root-mean-square error (RMSE) in the SoC estimation using the PEM-VB as the model for the EKF estimator is less than 2%.

The next simulation involves comparing the estimation procedure using the EKF formulation with and without the information of number of requests. A simple step-down reference signal for 150 EWHs having packet duration of 5 minutes and packet power of 2kW suffices to illustrate the two cases. Fig. 6 shows the result when requests are available to the coordinator. Observe that the reference signal is tracked, as expected, and the EKF using the PEM-VB model reproduces demand power and the number of requests accurately while estimating SoC with approximately 2% RMSE. In the simulation, the Gaussian noises R_1 and R_2 have been tuned to improve the estimation.

On the other hand, the case where no requests are fed back into the estimation produces Fig. 7. Even though demand power is estimated correctly, the EKF is not able to estimate the system's SoC accurately. In the simulation, only the PEM scheme accepts information about the number of requests to perform the tracking of the reference. But the EKF does not receive such information resulting in more than 10% RMSE. Furthermore, the simulation was limited to 150 minutes since the error increases after that time.




Fig. 6: Estimating SoC while tracking a reference with x_r .




Fig. 7: Estimating SoC while tracking a reference without x_r .

B. Pre-positioning flexible load w.r.t. day ahead forecasts

In this section, the predictive capability of the PEM-VB model is shown. For this purpose, consider the net load in Vermont, shown in Fig. 8, during two consecutive days in the summer of 2017 [21]. The net load is fulfilled by a mixture of conventional generation and variable distributed renewable generation. As shown in Fig. 8, a difference of over 150MW was observed due to June 6 being

cloudy resulting in lower distributed generation. On the other hand, June 7 was forecasted to be sunny and the conventional generation had to ramp up/down to account for this variability, which is costly. It is, therefore, desirable to minimize the deviation of conventional generation from a given set-point. Provided the day ahead forecasts of the net load P_f are available, the nominal power consumption (P_{nom}) of EWHs can be modified (due to their flexibility) using an optimal control problem to design a control input P_{ref} . This input, when fed to the PEM system of Fig. 1 minimizes the ramp up/down of generators. The PEM-VB model has the desired low order structure suitable for the OCP formulation as shown next.

The optimal control problem uses the PEM-VB model with the state space given by states $x[k] = (x_1[k], x_2[k])^\top$, and mapping $f(x[k], u[k]) = (f_1(\cdot), f_2(\cdot))^\top$ given in (17a), (17b). The control input $u[k] := P_{ref}[k]$ is considered as the optimization variable. Furthermore, it is assumed here that under nominal conditions N_{opt} is small enough and can be omitted from the PEM-VB model. Another reason for this is that in this work ξ_{opt-in} and $\xi_{opt-out}$ are considered as inputs, which in a predictive framework they require to be modeled as part of the dynamic equations. This is outside the scope of the paper and left for future work. However, in our simulations, it is observed that if the energy remains within $[20, 80]\%$ then the number of opt-out DERs is small enough ($\leq 1.5\%$ on average) and can be ignored for the purpose of this day-ahead reference planning problem. Let x_1^{max} , x_1^{min} be the energy limits of the PEM-VB obtained in Section III, then the reduced energy bounds $\underline{x}_1 > x_1^{min}$ and $\bar{x}_1 < x_1^{max}$ are used instead.

Consider now a scaled version of the net-load day-ahead forecast P_f , observed on June 7, as shown at the top plot of Fig. 9 (a). Let g be the contribution of conventional generation and $l := \sum_{i=1}^K P_f[k]$ be the average net-load. The objective is to keep the conventional generation (g) as close as possible to l by modifying P_{nom} of EWHs. Given an initial state of PEM-VB $x_0 \in \mathbb{R}^2$, an OCP similar to the one in [8], is defined for minimizing the following cost function,

$$\begin{aligned} \chi(P_{ref}[k], g[k], x[k]) &= \sum_{k=1}^{K+1} (c_1(\Delta P_{dev}[k-1])^2 + c_2(\Delta g_{dev}[k])^2 + c_3(\Delta x_1[k])^2) + \\ &\quad \sum_{k=1}^{K+1} (c_4(\Delta P_{ref}[k-1])^2 + c_5(\Delta g[k])^2), \end{aligned} \quad (19)$$

where, $c_i > 0, i = 1, \dots, 5$ are weights that have been tuned (here $c = (c_1, c_2, c_3, c_4, c_5) = (1, 100, 100, 1, 100)$), so that the ramp up/down costs of $g[k]$ are higher than those of EWHs, $\Delta P_{dev}[k] := P_{ref}[k] - P_{nom}[k]$ and $\Delta g_{dev}[k] := g[k] - l$ penalize the deviation of EWHs and conventional generation from P_{nom} and l , respectively. Similarly, $\Delta x_1[k] := x_1[k] - [1 \ 0]x_0$ discourages deviation of EWH energy from the nominal. The remaining two terms, $\Delta P_{ref} := P_{ref}[k] - P_{ref}[k-1]$, $\Delta g[k] = g[k] - g[k-1]$ are ramp rate limits. The resulting optimization problem is given by,

$$\min_{P_{ref}[k], g[k], x[k]} \chi(P_{ref}[k], g[k], x[k]) \quad (20a)$$

$$\text{s.t. } x[k+1] = f(x[k], P_{ref}[k]) \text{ and (12),}$$

$$P_{ref}[k] \geq P_{rate}x_2[k], \quad (20b)$$

$$P_{ref}[k] \leq P_{rate}(P_{req}(x_1[k])(N - x_2[k]) + x_2[k]), \quad (20c)$$

$$P_f[k] = \Delta P_{dev}[k] + g[k], \quad (20d)$$

$$x \leq x[k] \leq \bar{x}, \forall k = 1, \dots, K+1, \quad (20e)$$

$$x[0] = x_0, x_1[K+1] = [10]x_0, \quad (20f)$$

where, $\underline{x}, \bar{x}$ are the bounds on the state, the constraints (20b), (20c) on P_{ref} are obtained from $0 \leq \beta[k] \leq 1$ and (11), (20d) is for energy conservation, (20e) represent bounds on PEM-VB energy limits, (20f) is the initial state x_0 and terminal (sustainability) constraint. The timer states ($\tau[k]$) have been included in the OCP so that the expiration rate $\beta^-[k]$ can always be calculated from the timer states and hence it does not constitute an input to PEM-VB. Purpose of the timer is to keep track of expiring packets. Time step of the OCP has been set equal to the packet size, so only a single timer state, $n_p = 1$, is required.

The optimization problem is solved using IPOPT in Julia with JuMP on a MacBook Pro with a 2.2 GHz processor and 16 GB memory. The solution was obtained within 7 seconds. The result of solving the optimization problem (20) for a fleet of 15,000 EWHs is shown in Fig. 9. Ramp up/down of the generation $g[k]$, shown in blue in Fig. 9 (a), is minimal. To minimize ramping of generators, the optimal control problem anticipates the increase in distributed generation and preheats the EWH fleet. This can be seen from Fig. 9 (b) where P_{ref} and P_{nom} have been plotted. The evolution of the EWHs SoC is shown in Fig. 9 (c). PEM-VB captures the average SoC or energy of the fleet, shown in blue in Fig. 9 (c). Notice that between hrs 16-20 the fleet is discharging to account for the decrease in distributed generation. This pushes the fleet to the lower end of the energy limit and as a result, opt-outs increases. To capture this behavior, a model is required and will be explored in future publications. Furthermore, it should be noted that although opt-out was not included explicitly in the OCP, the EWHs can still track the day ahead reference P_{ref} and minimize generator deviation from the given set point.




Fig. 8: Netload in Vermont during two consecutive days in summer.




Fig. 9: Optimal dispatch of EWHs. (a) Net load forecasts (P_f) and setpoint of generators (g), (b) nominal power consumption, P_{nom} and the input P_{ref} , (c) evolution of EWHs SoC distribution. PEM-VB captures the average SoC or energy of the fleet. The Kullback-Leibler divergence between initial and final distributions is 0.061.

V. CONCLUSION AND FUTURE WORK

This paper developed a low-order model for a homogeneous ensemble of EWHs operating under PEM. The PEM-VB, is

capable of predicting the SoC and energy bounds. It was shown that the model is strong locally observable and represents a bottom-up scheme when the number of requests is available to the coordinator. The estimation of the PEM-VB states has been illustrated with an EKF implementation in which the SoC of an ensemble of individually simulated EWHs is recovered with an accuracy of 2% RMSE. Furthermore, it was shown that the PEM-VB model can also be used in an optimal control problem to account for the variability in renewables.

REFERENCES

- [1] M. Golden, A. Scheer, and C. Best, "Decarbonization of electricity requires market-based demand flexibility," *The Electricity Journal*, vol. 32, no. 7, p. 106621, 2019.
- [2] J. L. Mathieu, M. Dyson, and D. S. Callaway, "Using residential electric loads for fast demand response: The potential resource and revenues, the costs, and policy recommendations," in *American Council for an Energy-Efficient Economy*, 2012, pp. 189–201.
- [3] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent, "Potentials and economics of residential thermal loads providing regulation reserve," *Energy Policy*, vol. 79, pp. 115 – 126, 2015.
- [4] N. J. Cammardella, R. W. Moye, Y. Chen, and S. P. Meyn, "An energy storage cost comparison: Li-ion batteries vs distributed load control," *Clemson University Power Systems Conference*, pp. 1–6, 2019.
- [5] J. L. Mathieu, M. Kamgarpour, J. Lygeros, and D. S. Callaway, "Energy arbitrage with thermostatically controlled loads," in *European Control Conference*, 2013, pp. 2519–2526.
- [6] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent, "Aggregate flexibility of thermostatically controlled loads," *IEEE Transactions on Power Systems*, vol. 30, no. 1, pp. 189–198, 2015.
- [7] D. Madjidian, M. Roozbehani, and M. A. Dahleh, "Energy storage from aggregate deferrable demand: Fundamental trade-offs and scheduling policies," *IEEE Transactions on Power Systems*, vol. 33, no. 4, pp. 3573–3586, 2018.
- [8] N. Cammardella, J. Mathias, M. Kiener, A. Bušić, and S. Meyn, "Balancing california's grid without batteries," in *IEEE Conference on Decision and Control*, 2018, pp. 7314–7321.
- [9] I. Chakraborty, S. P. Nandanoori, and S. Kundu, "Virtual battery parameter identification using transfer learning based stacked autoencoder," in *IEEE International Conference on Machine Learning and Applications*, 2018, pp. 1269–1274.
- [10] M. Almassalkhi, J. Frolik, and P. Hines, "Packetized energy management: Asynchronous and anonymous coordination of thermostatically controlled loads," in *2017 American Control Conference*, 2017, pp. 1431–1437.
- [11] M. Almassalkhi, L. A. Duffaut Espinosa, P. D. Hines, J. Frolik, S. Paudyal, and M. Amini, *Asynchronous Coordination of Distributed Energy Resources with Packetized Energy Management*. New York, NY: Springer New York, 2018, pp. 333–361.
- [12] L. A. Duffaut Espinosa, M. Almassalkhi, P. Hines, and J. Frolik, "System properties of packetized energy management for aggregated diverse resources," *Power Systems Computation Conference*, pp. 1–7, 2018.
- [13] K. Desrochers, A. Khurram, M. Amini, A. Giroux, F. Wallace, J. Slinkman, V. Hines, M. Almassalkhi, and P. Hines, "Real-world, full-scale validation of power balancing services from Packetized virtual batteries," in *Conference on Innovative Smart Grid Technologies North America*, 2019, pp. 1–5.
- [14] L. A. Duffaut Espinosa and M. Almassalkhi, "A packetized energy management macromodel with quality of service guarantees for demand-side resources," *IEEE Transactions on Power Systems*, to appear.
- [15] L. A. Duffaut Espinosa, M. Almassalkhi, and A. Khurram, "Reference-tracking control policies for packetized coordination of diverse der populations," *IEEE Transactions on Control Systems Technology*, submitted.
- [16] L. A. Duffaut Espinosa, M. Almassalkhi, P. Hines, and J. Frolik, "Aggregate modeling and coordination of diverse energy resources under packetized energy management," in *56th IEEE Conference on Decision and Control*, 2017, pp. 1394–1400.
- [17] S. Buchberger and L. Wu, "Model for instantaneous residential water demands," *Journal of Hydraulic Engineering*, vol. 121, no. 3, pp. 232–246, 1995.
- [18] S. Wong, W. Muneer, M. S. Nazir, and A. Prieur, *Designing, operating, and simulating electric water heater populations for the smart grid*, 2014.
- [19] W. Rudin, *Principles of Mathematical Analysis*, ser. International series in pure and applied mathematics. McGraw-Hill, 1976.
- [20] ISO New England, "Simulated automatic generator control (AGC) setpoint data: Energy neutral," <https://www.iso-ne.com/isoexpress/web/reports/grid/-tree/simulated-agc>, Last Accessed: 2017-10-18.
- [21] —, "Energy load and demand reports," ATM Forum Contribution 94-0735R1, Aug. 1994. [Online]. Available: <https://www.iso-ne.com/isoexpress/web/reports/load-and-demand>