



Balancing supply and demand in the presence of renewable generation via demand response for electric water heaters

Adham I. Tammam¹ · Miguel F. Anjos^{1,2} · Michel Gendreau³

Published online: 31 March 2020

© The Author(s) 2020

Abstract

With the increasing presence of renewable energy sources in the electrical power grid, demand response via thermostatic appliances such as electric water heaters is a promising way to compensate for the significant variability in renewable power generation. We propose a multistage stochastic optimization model that computes the optimal day-ahead target profile of the mean thermal energy contained in a large population of heaters, given various possible wind power production and uncontrollable load scenarios. This optimal profile is calculated to make the variable net demand as even as possible.

Keywords Demand response · Electric water heaters · Stochastic optimization · Renewable power generation

List of symbols

Indices

t	Time steps
n	Nodes of the scenario tree
i, j	States of the Markov chain / types of water extraction

Parameters

K	Thermal conductivity per unit length of EWHs
-----	--

✉ Miguel F. Anjos
anjos@stanfordalumni.org

Adham I. Tammam
adham.ismail@polymtl.ca

Michel Gendreau
michel.gendreau@cirrelt.net

¹ GERAD and Department of Mathematics and Industrial Engineering, Polytechnique Montreal, C.P. 6079, Succ. Centre-Ville, Montreal, QC H3C 3A7, Canada

² School of Mathematics, University of Edinburgh, James Clerk Maxwell Building, The King's Buildings, Peter Guthrie Tait Road, Edinburgh EH9 3FD, United Kingdom

³ CIRRELT and Department of Mathematics and Industrial Engineering, Polytechnique Montreal, C.P. 6079, Succ. Centre-Ville, Montreal, QC H3C 3A7, Canada

A	Total surface area of all the EWHs
C^p	Hot-water specific heat
ρ	Water density
V	Total volume of hot water in the EWHs
N^{ewh}	Number of EWHs
T^L	Inlet water temperature
T^{env}	Environment temperature
$\Theta = \{1, 2, \dots, \mathcal{J}\}$	Set of water extraction values for the Markov chain θ_t
$\Lambda = \{\lambda_{ij}, i, j, = 1, \dots, \mathcal{J}\}$	Infinitesimal generator for the Markov chain θ_t
p_i	Probability of occurrence of each state of the Markov chain
T^{mix}	Desired temperature for the end-user of hot water
\dot{V}_i^{mix}	Desired flow of water for the end-user for extraction of type i
e^{min}	Minimum amount of energy that can be stored in the EWHs
e^{max}	Maximum amount of energy that can be stored in the EWHs
T^{min}	Lower limit of the temperature of the population of EWHs
T^{max}	Upper limit of the temperature of the population of EWHs
N	Number of nodes in the scenario tree
T	Number of time periods
Δt	Discrete time step
$\{\xi_t\}_{t \in T}$	Scenario tree
$d(\omega_n)$	Value of the uncontrollable demand at node n of the scenario tree
$r(\omega_n)$	Wind power production at node n of the scenario tree
\hat{n}	Parent node of node n in the scenario tree
\mathbb{P}_n	Marginal probability of the occurrence of node n

Decision variables

e_t	Energy stored in the system at time step t
$x(e_t)$	Quantity of energy injected into the EHWs at time step t
p_n	Net demand curve for the mean field controller

Other variables

$\ell_1(e_t)$	Heat transfer by conduction
$\ell(e_t)$	Energy loss due to hot water extraction and heat transfer by conduction
$\theta_t, t \geq 0$	Continuous-time Markov chain of water extractions
$f(e_t)$	Density function of e_t

1 Introduction

Compared to thermal resources, renewable energy sources are expensive in terms of equipment, installation, and maintenance. Their increasing use in electric grids is mainly due to the desire to reduce greenhouse gas emissions from burning fossil fuels. The variability of renewables creates challenges for system operators who must ensure a balance between supply (constrained by ramping limits) and demand. A promising tool to provide means to cope with these challenges is the provision of flexibility by the demand from the loads in the system. In particular, the load curve of thermostatically controlled appliances (TCAs)

such as electric water heaters (EWHs), space heaters, and the batteries of connected electric vehicles, can in principle be reshaped while respecting the end-user comfort constraints.

The work in this paper was carried out as part of the smartDESC (smart Distributed Energy Storage Controller) project (Sirois et al. 2017). The objective of this project was to develop and validate a scalable methodology to harness the energy potential of very large numbers of small TCAs distributed throughout an electrical grid. The resulting methodology provides a tool that can be used for traditional peak shaving as well as to reduce the impacts of the fluctuations of intermittent renewable energy sources, particularly solar and wind.

The smartDESC project focused on controlling the load of EWHs, which can store energy for considerable periods because of their high thermal inertia. Their demand peak coincides with the peak of the total demand, so a significant reduction of the load curve peak could potentially be achieved (Sepulveda et al. 2010). According to Natural Resources Canada, the power consumption of EWHs in Canada can reach 21.7% of the total demand (Energy Use Data Handbook 2004). With renewable generation capacity growing quickly in Canada, notably with wind capacity doubling by 2040 (Canada's Energy Future 2017), EWHs as a balancing resource is only going to grow in importance.

More generally, the demand-side management of TCAs is a promising way to counter-balance the variability of renewables, and it has been the focus of many studies. Dynamic programming models have been developed to minimize the peak load given a deterministic demand (Bischke and Sella 1985; Cohen and Wang 1988; Zhang and Li 2013). A control algorithm (Kondoh et al. 2011) has been developed to allow TCAs to follow regulation signals in order to stabilize a network supplied with renewable resources. A fuzzy logic control strategy (Nehrir and LaMeres 2000; LaMeres et al. 1999) can be used for TCA load shifting from peak to off-peak periods.

Lee and Wilkins (1983) have proposed a deterministic linear optimization model that decides the number of EWHs to which a control scheme selected from a predefined set should be applied in order to reduce the peak load. In Laurent et al. (1995), a column generation approach is applied to a load management problem where the objective is to minimize the maximum peak of a known load profile by choosing from a large set of admissible interruption scenarios established in advance. A metaheuristic algorithm based on particle swarm optimization to manage the power consumption of EWHs has been proposed (Rosario et al. 2011; Sepulveda et al. 2010). Directly controlling the power consumption of EWHs could reduce power losses in an electric grid as in Salehfar and Wehbe (2001).

The aforementioned studies consider a deterministic setting, but in reality the load demand and renewable supply are uncertain. Malik and Havel (2014) proposed centralized direct load control of EWHs to reduce the peak imports and exports in the Czech electricity market. Their approach takes into account the stochastic nature of load demand and renewable production. The dispatch of the EWHs is decided through a two-stage stochastic optimization program, where the first stage computes the overall EWH load, and the second stage adjusts the dispatch according to the actual power supply and demand.

The geographical distribution of the EHWs, and of distributed energy resources in general, is an important practical concern. Multiple challenges arise in the control of a large number of small but diverse storage devices spread over a wide area. A control schedule applied indiscriminately can reduce their natural load diversity, inducing the payback phenomenon that may create new peak loads (Laurent et al. 1995). Moreover, carefully controlling a large population of storage devices requires sophisticated mathematical models and significant computational power.

For these reasons, the methodology proposed in smartDESC combines a stochastic optimization model and a mean field model. More specifically, we consider a two-phase approach

to the control problem. In the first phase, we optimally schedule the day-ahead load of a homogeneous aggregated model of the EWH population. The objective is to even out the net load as much as possible. There are two stochastic parameters: (1) the uncontrollable load (the total load after excluding the controllable EWH demand) and (2) the renewable supply. We use a stochastic optimization model to compute an optimal power profile (OPP) for this aggregated model. This profile is then translated into a temperature profile that specifies a series of hourly setpoints that the hot water as an aggregate is required to reach, in order to achieve the OPP. In the second phase, a local control module sends instructions to individual EWHs to ensure that the mean thermal energy of the EWH population follows the OPP as closely as possible. This local controller (Kizilkale and Malhame 2014) is based on mean field theory; we refer to it as a mean field controller (MFC). This paper is concerned with the first phase. We refer the reader to (Kizilkale and Malhame 2014; Kizilkale and Malhamé 2013) for more information on the second phase.

The contribution of this paper is firmly set in the context described above. We propose a multistage stochastic optimization model called the *Scheduler* that, given the current state of the EHWs, as well as information on the total demand and the wind production for the next T time periods, computes the OPP that minimizes the mean variation of the net demand over those T time periods. The optimal solution is then sent to the MFC to be applied via instructions to each individual EHW. A key point here is that the signals sent to the EWHs are in terms of temperature, which is a proxy for energy, but the desired effect is quantified in terms of power shifted from certain periods of time to others to achieve the reduction of variation of the net demand.

It is important to note that in the context of smartDESC, the stochastic optimization problem is solved on a rolling horizon basis. In other words, at each time step, the Scheduler receives updates on the energy capacity of the EHWs, the demand, and the wind production, and builds a new scenario tree for T periods using this updated information. It then computes a new OPP over the new scenario tree, and again sends the optimal solution to the MFC. We operate in this way because at the start of each day, or whenever an unexpected event causes reality to deviate significantly from the forecast, it is essential to generate a new OPP to prevent too much error accumulation at the level of the MFC. The stochastic optimization model is thus solved repeatedly to calculate the temperature profile that best matches the new reality, and updated instructions are computed and sent by the MFC to the EHWs.

The rest of this paper is organized as follows. In Sect. 2 we introduce the aggregated EWH model and show how to maintain OPP feasibility for the MFC. In Sect. 3 we present the Scheduler. In Sect. 4 we discuss the scenario generation. In Sect. 5 we present a case study and the computational results, and Sect. 6 provides concluding remarks.

2 Aggregated EWH model

Our optimization formulation is designed to work in conjunction with the mean field model (Kizilkale and Malhame 2014), where each EWH is modeled by assuming that the reservoir is made up of n fully mixed equal volume layers. The cold water inlet is in the bottom layer, and hot water is drawn from the top layer. Our formulation uses an aggregated EWH model by considering a group of homogeneous EWHs as one large thermal battery. We thus model the EWHs as a single-layer large reservoir with a controllable mean water temperature given a sufficiently good approximation of the energy that the EWH population is capable of absorbing as well as of the losses (mostly due to hot water draw events). Thermal energy conservation for the aggregated EWH model, also called the system dynamics, is expressed as:

$$e_{t+1} = e_t + x(e_t) - \ell(e_t). \tag{1}$$

Here e_t is the stored energy (i.e., the system state) at time step t ; $x(e_t)$ is the decision variable that represents the quantity of energy injected into the reservoir, which depends on the current system state; and $\ell(e_t)$ is the system loss due to hot water extraction and heat transfer by conduction.

The heat transfer by conduction, denoted $\ell_1(e_t)$, is

$$\ell_1(e_t) = KA \left(\frac{e_t}{C^p \rho V} + N^{\text{ewh}}(T^L - T^{\text{env}}) \right) t \tag{2}$$

where we have the following quantities are that exogenous inputs: K is the thermal conductivity per unit length of EWHs, A is the total surface area of all the EWHs, C^p is the hot-water specific heat, ρ is the water density, V is the total volume of hot water in the EWHs, N^{ewh} is the number of EWHs in the model, T^L is the inlet water temperature, and T^{env} is the environment temperature.

We model the hot water extraction process as a time process on a finite state space that satisfies the Markov property. Specifically, we adopted the model in Kizilkale and Malhame (2014) where extraction is modeled as a continuous-time Markov chain. It is denoted θ_t , $t \geq 0$, and takes values in $\Theta = \{1, 2, \dots, \mathfrak{J}\}$, with the identical infinitesimal generator $\Lambda = \{\lambda_{ij}, i, j, = 1, \dots, \mathfrak{J}\}$, where each state consists of different drawn water volumes depending on the type of event such as shower or hand washing. In a continuous time Markov chain, the useful information that we use to calculate the probability of a given state is the distribution of the waiting time at every state. The infinitesimal generator parameters, also called transition rate parameters or matrix, describe the rate of movement between states. Hence, λ_{ij} describes the rate of transition to state j from state i . This transition rate is then used to calculate the probability of occurrence of each state, denoted p_i and defined as

$$p_i = \frac{\Pi_i}{\sum_{k=0}^K \Pi_k} \tag{3}$$

$$\Pi_0 = 1, \quad \Pi_i = \frac{\lambda_{0,1}\lambda_{1,2} \dots \lambda_{i-1,i}}{\lambda_{1,0}\lambda_{2,1} \dots \lambda_{i,i-1}} \text{ for } i \geq 1.$$

Given this, we can aggregate the losses due to extraction by considering the expected flow of drawn hot water for each type of event i as follows:

$$\ell_2 = \rho C^p (T^{\text{mix}} - T^L) \sum_{i=1}^{\mathfrak{J}} N^{\text{ewh}} p_i \dot{V}_i^{\text{mix}}. \tag{4}$$

Here we assume that the end-user mixes hot and cold water together to obtain the desired temperature T^{mix} and the desired flow \dot{V}_k^{mix} , with the flow depending on the type of extraction i .

The water temperature must be bounded below to prevent bacterial contamination (especially from *Legionella pneumophila*, for which the growth potential is almost zero above 46°C (Lacroix 1999)) and bounded above for end-user safety. The zone between these two bounds is the comfort zone that we represent for the aggregate of the EHWs using the constraints:

$$e^{\text{max}} = N^{\text{ewh}} \rho V C^p (T^{\text{max}} - T^L) \quad \forall t, s \tag{5}$$

$$e^{\text{min}} = N^{\text{ewh}} \rho V C^p (T^{\text{min}} - T^L) \quad \forall t, s$$

Note that these expressions are also bounds on the energy that can be stored in the set of EWHs. In other words, the maximum quantity of energy the system is able to absorb, $x(e_t)$,




Fig. 1 Population distribution near the lower bound

depends on its mean current state, e_t , because as this state approaches the upper bound, it is able to accept less energy. Moreover, the EWHs will consume a minimum quantity of energy to prevent the system in aggregate from going below the lower limit of the comfort zone.

To bound the aggregated power consumption of the EWHs, we considered the MFC module of Kizilkale and Malhame (2014) as a black box and used it to calculate the minimum and maximum electric power that the EWH population can consume for all reachable values of e_t . Note that not all members of the EWH population reach the same energy level. They are distributed with a certain variance and skewness, where this distribution is not normal because of the comfort zone constraint that trims the tails of the probability density function. Furthermore, the variance and skewness depend on the control. The variance tells us how far the extreme energy states are from the average. We figure out that when the population average is next to the boundaries, the population is more squeezed, in contrary to the case when the population's mean is far from the boundaries. The former gives us less energy state diversity. The skewness is also important because it shows that the distribution of the EWHs is not symmetric around the mean (i.e., the probability of having EWHs with energy state higher than the average is greater than the probability of having EWHs with less energy). Therefore, we expect that asking the EWH population to reduce its mean temperature is an easier job than increasing it, because the pool size of EWHs which temperature is higher than the mean is larger, hence more flexibility to reduce the mean with minimum individual disturbance in this case. We expect also that with higher diversity, less violation of individual comfort zones will occur.

We conducted a study of the distribution of the EWH population state around a finite set of system states e_t in order to randomly generate initial EWH states following this precalculated probability density function. For every discrete value of a target energy level e_t , we begin the simulation assuming that the energy state of the EWH population follows a normal distribution with mean μ and variance 1. The tails of the distribution are truncated because we have lower and upper bounds on the inner temperature of the hot water inside the EWH tank, but as shown in Figs. 1, 2 and 3, the area under the probability density function near the bounds is sufficiently small to argue that ignoring the truncation is a fair assumption. The MFC is asked to control the EWH population so that its mean energy state reaches e_t . When the population mean converges to e_t , we calculate its variance and skewness around e_t ; we denote this density function by $f(e_t)$. Then, for every e_t , the state of the EWH population is initialized so that its state distribution follows $f(e_t)$. The MFC is next asked to ensure that the mean population temperature is between its lower and upper limits, T^{\min} and T^{\max} . The




Fig. 2 Population distribution near the upper bound




Fig. 3 Population distribution in the middle of the comfort zone

lower and upper aggregated power consumption bounds can then be calculated. We use this simulation to initialize the energy state of the EHW population, to reflect as much as possible a reasonable initial state of the population, before applying our target control trajectory.

Figures 1, 2 and 3 show the distribution of the thermal energy of the population as its mean moves towards the lower bound, the upper bound, and the middle of the comfort zone. We can see that the thermal energy distribution has a clear negative skew due to the hot water draw events. These result in thermal losses and cause a large portion of the EWHs to have a temperature below the mean population temperature. Furthermore, the distribution variance shrinks as the mean approaches the bounds, which affects the population's diversity.

Figures 4, 5 and 6 are concerned with bounding the power consumption of the EWHs. Figure 4 shows the simulated results of the minimum and maximum power consumption for 200 EWHs. The feasible region of $x(e_t)$ is the region between the two monotonically nonincreasing functions. We will need to integrate these bounds in our model. To integrate the upper bound, we apply a linear regression, as shown in Fig. 5. For the lower bound, we use a convex quadratic regression, as shown in Fig. 6; the resulting quadratic function is then outer-approximated by a piecewise linear function formed using supporting hyperplanes.




Fig. 4 Simulated bounds

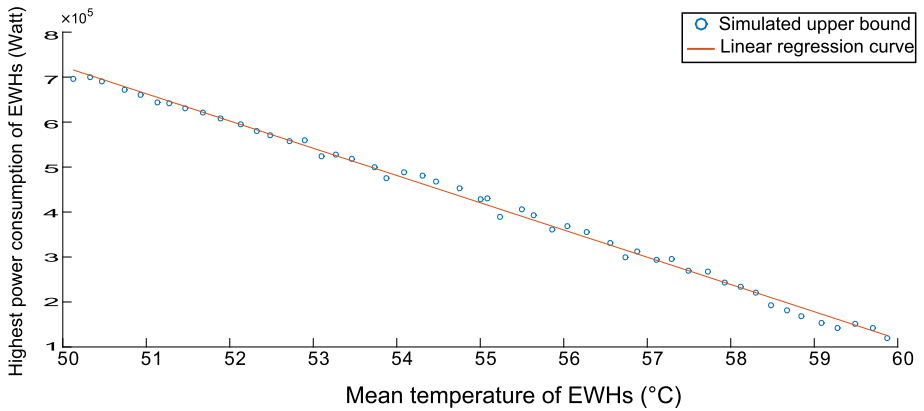


Fig. 5 Linear regression for the upper bound

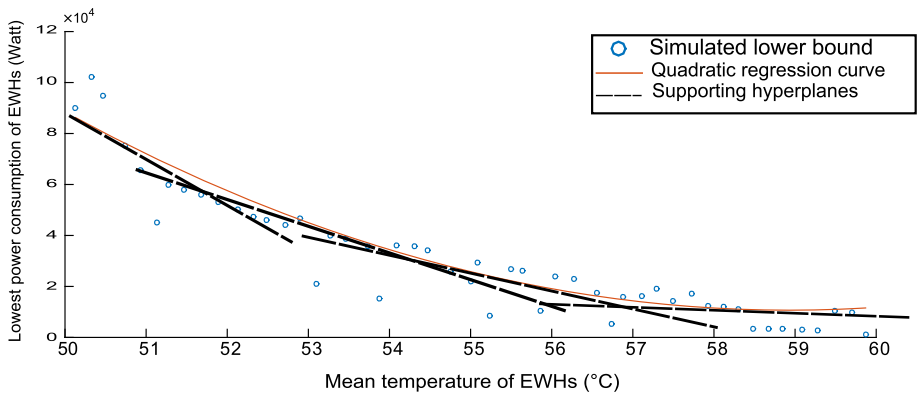


Fig. 6 Quadratic regression for the lower bound

3 Stochastic model

Several optimization techniques have been developed to deal with uncertain parameters, namely stochastic dynamic programming (Bertsekas 2007), robust optimization (Ben-Tal and Nemirovski 2000; Bertsimas et al. 2011; Bertsimas and Sim 2004; Ferreira et al. 2012), chance-constrained optimization (Dorini et al. 2013), and stochastic optimization with recourse (Birge and Louveaux 2011). See (Sahinidis 2004) for a review of optimization under uncertainty in general, and (Pisciella et al. 2015) for a detailed description of different optimization techniques with different levels of granularity for the description of the problem of capacity expansion in power systems.

Uncertain parameters are usually modeled by either distributions or stochastic processes. The former are used when the decision is to be made over a single stage, and the latter when a series of decisions must be taken over multiple stages (Kaut and Wallace 2003). In our case, decisions are taken for hourly time steps in the planning horizon, and the two stochastic processes are the uncontrollable demand and the wind power production.

Stochastic optimization with recourse is our chosen approach. In certain special cases the model is solvable directly with continuous distributions, but most methods require discrete distributions of finite cardinality. The continuous random process must therefore be approximated by a discrete finite set of outcomes in the form of a *scenario tree* that represents the diffusion of stochastic information into the future. The tree represents the multiple stages of the observation of the possible outcomes of the random variables in time. Decisions are taken at different stages depending on the available data at the given stage, and regardless of future observations that are considered uncertain. As more observations are revealed, recourse decisions are made taking the revealed information into account while being consistent with the decisions made in previous stages.

Figure 7 illustrates a scenario tree with three stages, where the root node is the value of the discrete stochastic process $\{\xi_t\}$ at $t = 0$ and is considered deterministic (i.e., has a probability of occurrence equal to 1). Two possible outcomes at the next stage $t = 1$ are represented by two nodes with values ω_1 and ω_2 . Each of these can lead to two possible realizations of the random process, (ω_3, ω_4) and (ω_5, ω_6) respectively, with their conditional probabilities shown on the arcs. The number of stages in a scenario tree does not necessarily reflect the number of time steps in the optimization problem but rather the number of times that new observations of the random process are made. The larger the scenario tree, the better the representation of the uncertain stochastic process, but also the larger the optimization problem.

3.1 Mathematical model

The objective of the Scheduler is to even out the net demand curve. This curve is denoted p_n and is equal to the demand remaining (in MW) after the renewable production has been dispatched. We compute the optimal electric power consumption (in MWh) of the controllable EWH load of EWHs, denoted $x_n(e_n)$, so that valleys are filled and peaks are lowered while respecting the end-user comfort. The net demand is

$$p_n = d(\omega_n) - r(\omega_n) + \frac{x_n(e_n)}{\Delta t} \quad (6)$$

where $d(\omega_n)$ and $r(\omega_n)$ are respectively the observed values of the uncontrollable demand (in MW) and the wind power production (in MW) at node $n \in 1, \dots, N$ of the scenario tree




Fig. 7 Three-stage scenario tree

with N nodes, and Δt is the discrete time step. To respect the nonanticipativity conditions, the decision $x_{\hat{n}}(e_{\hat{n}})$ is taken at the parent node of node n in the scenario tree, denoted \hat{n} , before we observe the realization of $d(\omega_n)$ and $r(\omega_n)$. At the root node d and r are deterministic (observed) parameters, where d is the actual total demand, and r is the actual wind power.

Evening out the net demand curve involves minimizing the absolute value of the difference in the net demands at two consecutive nodes in the scenario tree multiplied by their probability of occurrence:

$$\min \sum_{n=0}^N \mathbb{P}_n z_n \tag{7}$$

where \mathbb{P}_n is the marginal probability of the occurrence of node n , independent from its parent node \hat{n} , and z_n is the absolute value of the difference between the net demands at node n and its parent node \hat{n} . We can model this absolute value using the following linear formulation:

$$z_n \geq p_n - p_{\hat{n}}, \quad z_n \geq p_{\hat{n}} - p_n. \tag{8}$$

The resulting stochastic optimization model is:

$$\begin{aligned} \min_{e, x, z \in \mathbb{R}^n} \quad & \sum_{n=0}^N \mathbb{P}_n z_n \\ \text{s.t.} \quad & z_n \geq p_n - p_{\hat{n}} \quad \forall n \\ & z_n \geq p_{\hat{n}} - p_n \quad \forall n \\ & p_n = d(\omega_n) - r(\omega_n) + \frac{x_{\hat{n}}(e_{\hat{n}})}{\Delta t} \quad \forall n \\ & e_n = e_{\hat{n}} + x_{\hat{n}}(e_{\hat{n}}) - \ell(e_{\hat{n}}) \quad \forall n \\ & \ell_1(e_n) = KA \left(\frac{e_n}{C^p \rho V} + N^{\text{ewh}}(T^L - T^{\text{env}}) \right) \Delta t \quad \forall n \\ & \ell_2 = \rho C^p (T^{\text{mix}} - T^L) \sum_{i=1}^J N^{\text{ewh}} p_i \dot{V}_i^{\text{mix}} \end{aligned}$$

$$\begin{aligned}
\ell(e_{\bar{n}}) &= \ell_1(e_{\bar{n}}) + \ell_2 \quad \forall n \\
e_o &= N\rho V C^p (T_{init} - T^L) \\
e_n &\leq N^{\text{ewh}} \rho V C^p (T_{\max} - T^L) \quad \forall n \\
e_n &\geq N^{\text{ewh}} \rho V C^p (T_{\min} - T^L) \quad \forall n \\
x_n(e_n) &\leq A_1(e_n) + B_1 \quad \forall n \\
x_n(e_n) &\geq \underline{Q}(e^i) + \hat{Q}(e^i)(e_n - e^i) \quad \forall n, i \\
x_n(e_n) &\geq 0 \quad \forall n \\
e_n &\geq 0 \quad \forall n
\end{aligned}$$

Here we have included the bounds on $x_n(e_n)$ from the regression approximation, and also the last two inequalities that require the non-negativity of the energy injected and of the stored energy.

4 Scenario generation

In this section we describe the generation of the multistage scenario tree for our stochastic optimization problem. Artelys, a company specializing in optimization, decision support, and modeling, cooperated with us by developing a load forecast model trained over a set of historical data for power consumption from 2012 to 2014. The data were provided by the supervisory control and data acquisition (SCADA) section of the Coopérative Régionale d'Électricité de Saint-Jean-Baptiste de Rouville (CoopSJB). The CoopSJB data are collected from five distribution substations in Mont Saint-Hilaire, a suburb of Montreal, for 6819 houses. They are normalized to obtain the mean power consumption per house. The forecast model takes as inputs the hourly wind speed and temperature of day d_k and the hourly forecast for day d_{k+1} , and it outputs the hourly load demand forecast for d_{k+1} . Multiple load demand curves are generated using wind speed and temperature ensemble forecasts provided by Environment Canada, which releases each day a set of 22 forecasts for the next 144 h. The uncontrollable demand component is obtained by computing an estimate of the hourly EWH consumption using the data from CoopSJB, and subtracting this estimate from the total demand. Wind power scenarios are generated from the wind speed ensemble forecasts using the approach proposed in Tammam et al. (2015). Every uncontrollable load demand scenario is coupled with its corresponding wind power scenario resulting from the same wind speed forecast.

We construct a collection of 22 scenarios with 24 nodes per scenario for the 24-h horizon in the following way. Each node has two (hourly) values: uncontrollable demand and wind power production. We first construct a so-called “comb” scenario tree with what amounts to two stages: the first stage corresponds to the root, and the second stage is composed of 22 possible branches from the root. Each of those branches covers all the remaining time periods in a deterministic fashion. It is common to use this kind of two-stage structure to describe complex systems evolving over several periods: Except for the first period, all periods are lumped into the second stage. In the context here, we chose to use this deterministic-equivalent approach because it is compatible with the application of this model jointly with the MFC. Note that while it is well known that this approach does not generally scale well for large scenario trees, we do not experience this scaling problem because we restrict ourselves to “comb” scenario trees.

However, as a consequence, this tree structure does not capture future uncertainty well because it does not possess the non-anticipativity property. To overcome this shortcoming, we apply the forward tree construction algorithm of Heitsch and Römisch (2007) to obtain a tree with only 22 leaves but that is non-anticipative. While this scenario tree represents only a small fraction of the possible outcomes over a 24-h period, it still forces decisions to account for uncertainty.

5 Case study and results

We performed a case study in the context of the smartDESC (smart Distributed Energy Storage Controller) project. The goal was to obtain a proof of concept for a hierarchical control architecture managing the power consumption of dispersed energy storage devices in the electric grid. The controller aimed to mitigate the variability caused mainly by the increasing penetration of renewable energy resources.

We note that we have not provided direct comparisons with other works in the literature. There are two reasons for this. First, most of the other approaches mentioned in Sect. 1 are deterministic. Second, many of the approaches in the literature are centralized methods that address the whole process, unlike ours that is (a) designed specifically to work together with a request dispatch method that is completely separate, and for our purposes is a “black box” (the MFC), and (b) centralized methods are highly impractical because of the high communication requirements, and indeed the architecture of SmartDESC (Sirois et al. 2017) was motivated by the desire to avoid this.

5.1 Results

The aggregated model used in this case study represents a population of 200 EWHs; but the curves in this section are averaged over 200 houses. The planning horizon is 24h with hourly time steps. We use a rolling horizon: at every time step t , a new scenario tree, denoted $\{\xi_t\}_{t \in T}$, is created based on the available wind speed and demand forecasts. We retain only the optimal solution $x_{0,t}^*$ at the root node $\xi_{0,t}$ of the tree, where the root node contains the values of total demand and wind production observed at the previous time step $t - 1$.

We scale the wind power production scenarios so that the average of their maximum values over the planning horizon is equal to 10% of the average value of the maximum uncontrollable demand over the horizon. We did this to test the load curve reduction the system can achieve when the maximum wind power production is on average equal to 10% of the uncontrollable demand.

As previously mentioned, the root node of the scenario tree $\{\xi_t\}_{t \in T}$ contains the actual realization of the total demand and wind power. Therefore, we must construct a deterministic observation of these two parameters over the planning horizon. We present three cases in which the mean uncontrollable demand forecast is considered as the observed uncontrollable demand, and this demand is coupled with three different wind power observations: minimum, average, and maximum power. This is depicted in Figs. 8, 9 and 10, where the upper curve is the uncontrolled demand, the lower curve is the net demand, and the region in between represents the magnitude of the wind power. We see that, in spite of the wind fluctuations, our approach can absorb this intermittence to ensure the balance between supply and demand.

Figures 11, 12 and 13 depict the load reduction curves with low, average, and high wind penetration. There are six curves for each of these cases of wind penetration:




Fig. 8 Case study with minimum wind power




Fig. 9 Case study with average wind power

- The EHW Target Power is the target given by the Scheduler to the MFC.
- The EHW Simulated Power is the level of consumption that the MFC is able to achieve (given the EHW target power).
- The Total Target Power is the resulting total power that the Scheduler expects to achieve by proposing the EHW target power.
- The Total Simulated Power is the total power that the MFC is able to achieve.
- Finally, the Total Power With Thermostatic EWH is the power level before we intervene, i.e., it is the power curve before our balancing is applied.

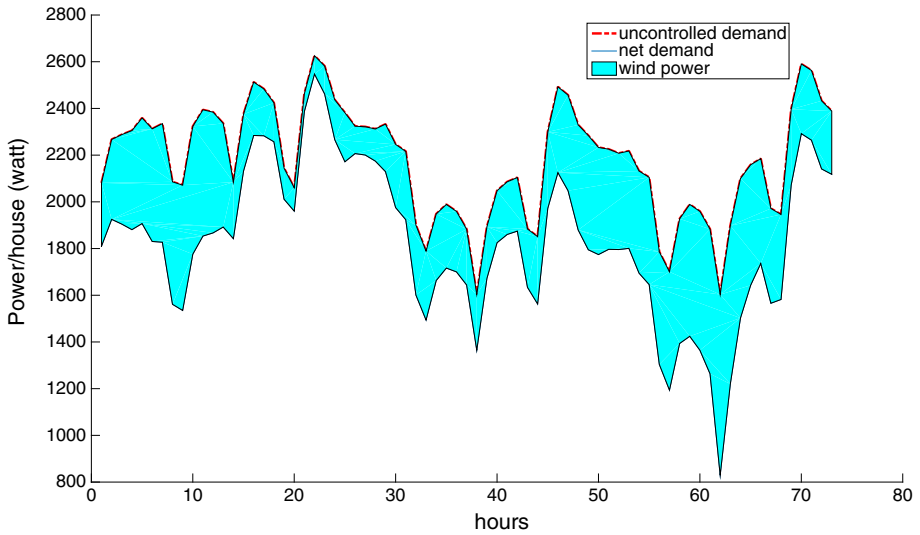


Fig. 10 Case study with maximum wind power




Fig. 11 Load reduction with low wind power

From these figures, we can make the following observations:

- The target power consumption profile generated by the Scheduler is generally feasible with respect to the MFC, i.e., the EWH simulated power curve and the EWH target power curve almost overlap.
- The total simulated power generally fluctuates less than the total power before we intervene, i.e., the green curve achieves significant curtailment of the peaks of the blue curve, and moreover the green curve also fills several of the low valleys of the blue curve.

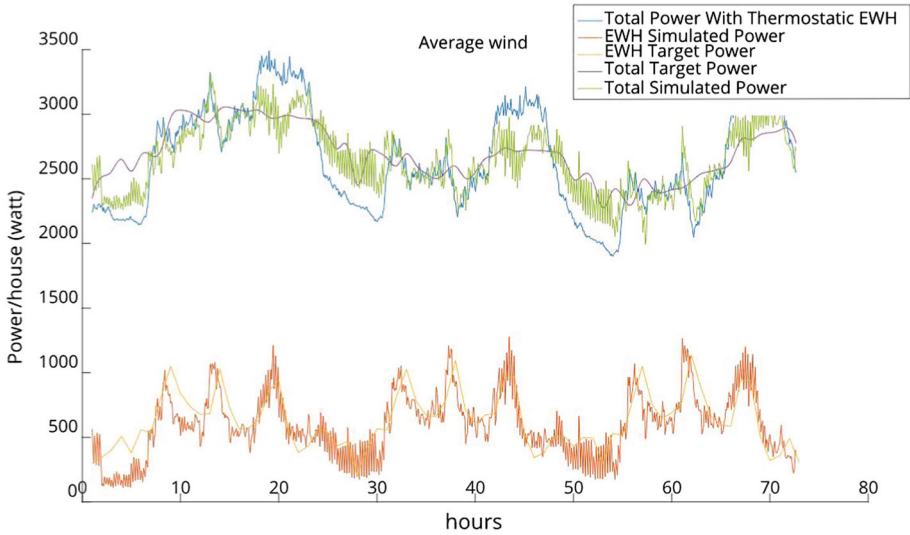


Fig. 12 Load reduction with average wind power




Fig. 13 Load reduction with high wind power

The outcomes observed in Figs. 11 to 13 can also be quantified by the results in Table 1. This table shows the peak reduction resulting from the direct control of the power consumption of the EWH population as a percentage of the load peak in the thermostatic control mode (third column); the reduction in demand variance as a percentage of the demand with the thermostatic control of EWHs (fourth column); and the computational time required to solve the rolling-horizon problem (last column). The reduction in demand variance is an estimate of how well the net demand curve is evened out.

Table 1 Comparison of three cases of wind blow for 10% and 20% of wind power penetration over three days

Wind penetration (%)	Wind	Peak reduction (%)	Net demand variation reduction (%)	Time (s)
10	High	6.68	46.40	21.535
	Average	7.85	49.42	21.471
	Low	7.84	50.82	21.454
20	High	6.51	32.82	22.674
	Average	8.57	42.60	21.849
	Low	8.46	45.16	22.544

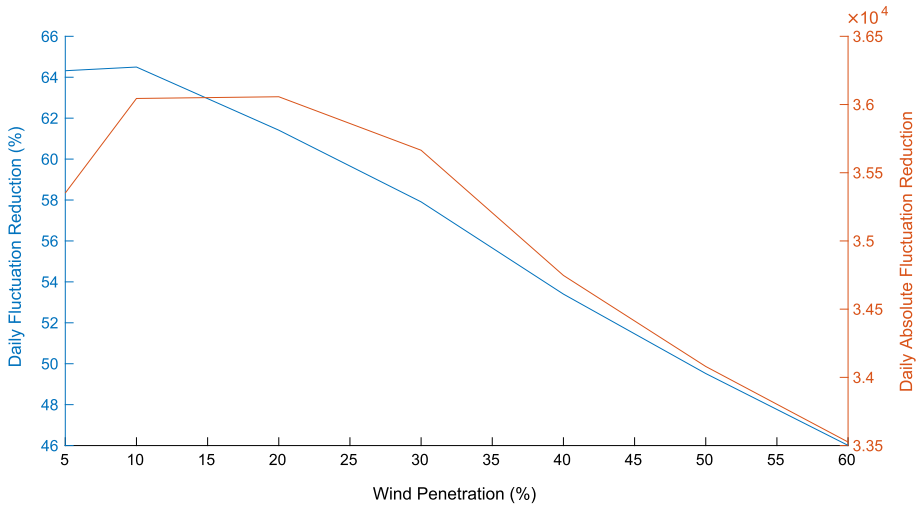


Fig. 14 Reduction of daily fluctuation versus wind penetration

Clearly, as wind power is added to the grid, the variation in the net demand increases, and it is more challenging to manage the EWH load to reduce the demand fluctuation. This can be seen by comparing the net demand variation reductions for 10% and 20% wind penetration.

Two other observations are important. First, Fig. 14 shows that when the wind power penetration to the grid increases, the reduction of the net demand load curve fluctuation is smaller. This is because of the limited ability of EWHs to absorb the wind power variability. On the other hand, Fig. 15 shows how this increase in wind power permits the operator to reduce the peak of the load curve, although eventually this gain in peak reduction decreases as the penetration of renewables increases. These results were obtained with a participation of 100% of the EWH population in the demand response program.

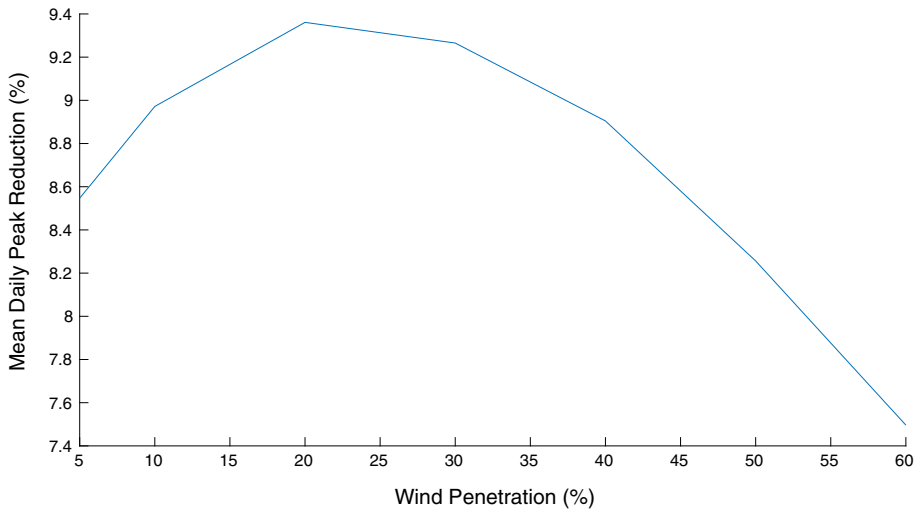


Fig. 15 Mean daily peak reduction versus wind penetration

6 Conclusion

We have proposed a multistage stochastic optimization model for load reduction in the presence of renewable resources attached to the grid, by means of the storage capacity of residential EWHs. This model is a part of the smartDESC project. The project provides a hierarchical control architecture that manages dispersed devices locally and more efficiently while achieving the global goals of peak reduction and net demand flattening. The model shows the impact of renewable resources on the variability of the net demand curve. This gives the system operator information about how much renewable power it can afford while maintaining a stable and serviceable demand.

Acknowledgements This work was funded by an ECO Energy Innovation Initiative grant from Natural Resources Canada and by the NSERC Energy Storage Technologies (NEST) Strategic Network.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ben-Tal, A., & Nemirovski, A. (2000). Robust solutions of linear programming problems contaminated with uncertain data. *Mathematical Programming*, 88(3), 411–424.
- Bertsekas, D. P. (2007). *Dynamic programming and optimal control* (4th ed.). Belmont: Athena Scientific.
- Bertsimas, D., Brown, D. B., & Caramanis, C. (2011). Theory and applications of robust optimization. *SIAM Review*, 53(3), 464–501.
- Bertsimas, D., & Sim, M. (2004). The price of robustness. *Operations Research*, 52(1), 35–53.
- Birge, J., & Louveaux, F. (2011). *Introduction to stochastic programming*., Springer series in operations research and financial engineering Berlin: Springer.

- Bischke, R., & Sella, R. A. (1985). Design and controlled use of water heater load management. *IEEE Transactions on Power Apparatus and Systems*, PAS-104(6), 1290–1293.
- Canada's Energy Future 2017: Energy Supply and Demand Projections to 2040. Natural Energy Board of Canada, Cat. No. NE2-12/2017E-PDF, 2017.
- Cohen, A. I., & Wang, C. (1988). An optimization method for load management scheduling. *IEEE Transactions on Power Systems*, 3(2), 612–618.
- Dorini, G., Pinson, P., & Madsen, H. (2013). Chance-constrained optimization of demand response to price signals. *IEEE Transactions on Smart Grid*, 4(4), 2072–2080.
- Energy Use Data Handbook. Natural Resources Canada, 2004.
- Ferreira, R., Barroso, L., & Carvalho, M. (2012). Demand response models with correlated price data: A robust optimization approach. *Applied Energy*, 96, 133–149.
- Heitsch, H., & Römisch, W. (2007). Scenario tree modeling for multistage stochastic programs. *Mathematical Programming*, 118(2), 371–406.
- Kaut, M., & Wallace, S. W. (2003). *Evaluation of scenario-generation methods for stochastic programming* (pp. 14–2003). Stochastic Programming E-Print Series: In *World Wide Web*.
- Kizilkale, A. C., & Malhamé, R. P. (2013). Mean field based control of power system dispersed energy storage devices for peak load relief. In *IEEE 52nd Annual conference on decision and control (CDC)*, 4971–4976.
- Kizilkale, A., & Malhamé, R. (2014). Collective target tracking mean field control for Markovian jump-driven models of electric water heating loads. Groupe d'Etudes et de Recherche en Analyse des Décisions (GERAD), Technical Report.
- Kondoh, J., Lu, N., & Hammerstrom, D. (2011). An evaluation of the water heater load potential for providing regulation service. *IEEE Transactions on Power Systems*, 26(3), 1309–1316.
- Lacroix, M. (1999). Electric water heater designs for load shifting and control of bacterial contamination. *Energy Conversion & Management*, 40, 1313–1340.
- LaMeres, B., Nehrir, M., & Gerez, V. (1999). Controlling the average residential electric water heater power demand using fuzzy logic. *Electric Power Systems Research*, 52(3), 267–271.
- Laurent, J.-C., Desaulniers, G., Malhamé, R., & Soumis, F. (1995). A column generation method for optimal load management via control of electric water heaters. *IEEE Transactions on Power Systems*, 10(3), 1389–1400.
- Lee, S. H., & Wilkins, C. L. (1983). A practical approach to appliance load control analysis: A water heater case study. *IEEE Transactions on Power Apparatus and Systems*, PAS-102(4), 1007–1013.
- Malik, O., & Havel, P. (2014). Active demand-side management system to facilitate integration of renewable energy sources (RES) in low-voltage distribution networks. *IEEE Transactions on Sustainable Energy*, 5(2), 673–681.
- Nehrir, M., & LaMeres, B. (2000). A multiple-block fuzzy logic-based electric water heater demand-side management strategy for leveling distribution feeder demand profile. *Electric Power Systems Research*, 56(3), 225–230.
- Pisciella, P., Bertocchi, M., & Vespucci, M. T. (2015). Models for optimization of power systems. *Numerical analysis and optimization* (pp. 289–317). Cham: Springer.
- Rosario, P. A., Morsi, W., & Chang, L. (2011). Application of PSO to optimize the operation of electric water heaters for reserve provision. In *Electrical power and energy conference (EPEC)*. IEEE, 164–169.
- Sahinidis, N. V. (2004). Optimization under uncertainty: State-of-the-art and opportunities. *Computers & Chemical Engineering*, 28(6–7), 971–983.
- Salehfar, H., & Wehbe, A. (2001). Direct control of residential water heater loads to reduce power system distribution losses. In *Power engineering society winter meeting*. IEEE, 3, 1455–1460.
- Sepulveda, A., Paull, L., Morsi, W., Li, H., Diduch, C., & Chang, L. (2010). A novel demand side management program using water heaters and particle swarm optimization. In *Electric power and energy conference (EPEC)*. IEEE.
- Sirois, F., Bourdel, B., & Malhamé, R. (2017). Project RENE-034: Managing energy storage capacities dispersed in an electrical grid to reduce the effects of renewable energy source variability. Public Report. http://www.profs.polymtl.ca/f.sirois/smartDESC_public_report.pdf
- Tammam, A. I., Watters, C. S., Anjos, M. F., & Gendreau, M. (2015). *A methodology for ensemble wind power scenarios generation from numerical weather predictions (accepted for IEEE PES GM 2016 proceedings)*. GERAD, Technical Report.
- Zhang, L., & Li, Y. (2013). Optimal energy management of wind-battery hybrid power system with two-scale dynamic programming. *IEEE Transactions on Sustainable Energy*, 4(3), 765–773.