Les Cahiers du GERAD ISSN: 0711–2440

Flexibility product for water heater aggregators on electricity markets

M. Pied, M. F. Anjos,
R. P. Malhamé

G–2019–43

June 2019

La publication de ces rapports de recherche est rendue possible grâce au soutien de HEC Montréal, Polytechnique Montréal, Université McGill, Université du Québec à Montréal, ainsi que du Fonds de recherche du Québec – Nature et technologies.

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2019
– Bibliothèque et Archives Canada, 2019

The publication of these research reports is made possible thanks to the support of HEC Montréal, Polytechnique Montréal, McGill University, Université du Québec à Montréal, as well as the Fonds de recherche du Québec – Nature et technologies.

Legal deposit – Bibliothèque et Archives nationales du Québec, 2019
– Library and Archives Canada, 2019

GERAD HEC Montréal
3000, chemin de la Côte-Sainte-Catherine
Montréal (Québec) Canada H3T 2A7

Tél. : 514 340-6053
Téléc. : 514 340-5665
info@gerad.ca
www.gerad.ca

Flexibility product for water heater aggregators on electricity markets

Marie Pied [a,b]
Miguel F. Anjos [a,c,d]
Roland P. Malhamé [a,b]

[a] *GERAD, Montréal (Québec), Canada, H3T 2A7*

[b] *Department of Electrical Engineering, Polytechnique Montréal (Québec) Canada, H3T 1J4*

[c] *Department of Mathematics and Industrial Engineering, Polytechnique Montréal (Québec) Canada, H3C 3A7*

[d] *School of Mathematics, University of Edinburgh, Edinburgh EH9 3FD, United Kingdom*

marie.pied@polymtl.ca
roland.malhame@polymtl.ca
anjos@stanfordalumni.org

June 2019
Les Cahiers du GERAD
G–2019–43

Abstract: Intermittent renewable energy, such as solar and wind, brings uncertainty into the grid. To increase their contribution into the energy mix, load management solutions are necessary to correct the resulting typical mismatches between generation and demand. This can be achieved rather effectively with thermostatic loads such as space heaters or water heaters by considering them as means of storage. This article proposes a mean field game-based controller to provide load flexibility into the grid using a multi-layer water heater model. A uniform local state feedback law is used to track the temperature trajectory specified by an aggregator for the group of controlled devices. The law is computed via a near fixed-point algorithm. A scheduling problem for the desired mean water heater target temperatures over a time horizon is formulated to find the maximum flexibility available from the group of loads while maintaining the typical post-control load oscillations within predefined bounds over a fixed time period. The solution of the scheduling problem is obtained by solving a linear optimization problem with upper and lower bounds on the power drawn by the group to converge to an acceptable mean temperature schedule.

Keywords: Smart Grid, thermal storage, mean field games, optimal control

Acknowledgments: This research was supported by the NSERC Energy Storage Technology Network (NESTNet) and by a GERAD scholarship.

Nomenclature

Parameters

$x_{l,t}$	Temperature of layer l at time t
$\bar{u}_{l,t}$	Energy input by heating at layer l at time t
$\dot{m}_t^{in}$	Rate of water extraction at time t
$\dot{Q}_l$	Presence/absence of a heating element in layer l
x_{env}	Temperature of the surroundings
x_{in}	Inlet fluid temperature
x^{low}	Lower comfort temperature
x^{high}	Upper comfort temperature
M_l	Fluid mass in layer l
A_l	Lateral surface of layer l
C_{pf}	Specific heat of the fluid
U	Loss coefficient between tank and surroundings
$\dot{V}_j^{mix}$	Flow for extraction of type j
N_wh	Number of water heater considered
Δ_t	Time discretization step
T_{start}	Time at which control of water heaters starts
T_{end}	Time at which control of water heaters stops
T_1	Time separating in two the control horizon, on each separate interval the objective is slightly different. See Figure 2 for details.
d_t	Uncontrolled demand at time t
p_t^b	Base aggregate water heater power demand at time t
x_{mix}	Temperature desired by the customer
$\dot{V}_j^{mix}$	Rate of extraction for events of type j
$\zeta_{j,\infty}(t)$	Quasi steady-state probability of state j of the water extraction Markov Chain associated to the infinitesimal generator L_t defined by $\begin{cases} \zeta_\infty(t)L_t = 0 \\ \sum_{j\in\Theta} \zeta_{j,\infty}(t) = 1 \end{cases}$
ρ	Water density
V	Water heater volume
A	Water heater surface area
$C_{direction}$	Indicator for increase (+1) or decrease (-1) in power consumption
$r_{rebound}$	The acceptable range for the post-control rebound
c_d	a integer coefficient for the bisection alike method.

Decision variables

e_t	Energy stored in the water heaters at time t
φ_t	Energy to inject in the water heaters at time t

1 Introduction

Many jurisdictions have adopted energy transition policies that focus on increasing the use of intermittent renewable energy sources in the electricity mix. However, a massive introduction of such sources brings new challenges linked to the instability they can potentially bring to the electricity grid because their power generation is highly variable during the day. To increase the integration of power produced by these sources, it is essential to have sufficient means to ensure the balance between generation and demand. This balance is often achieved by high marginal cost generation such as gas-fired plants (when renewable power is insufficient) or by curtailing renewable generation (when too much of it is available). Load management is thus a promising means to support greater integration of renewables.

Load management or demand response consists in controlling flexible loads to compensate for the fluctuation of generation. Much research has been done to quantify the demand response potential for peak reduction and load shifting, see e.g. [1, 2]. Projects such as PowerShift Atlantic [3] have demonstrated that this is both technically possible and economically promising. Industrial customers can already participate in demand response programs in Ontario [4].

In order to make residential load management programs possible, a new entity needs to enter into play: the aggregator. Indeed, on its own, a residential building or house has only a negligible impact on the balance of the grid, but when aggregated, the group can have a significant impact. The role of the aggregator is thus to offer a flexibility product on the electricity market while incentivizing consumer participation [5,6]. In particular, resildential customers can participate via the energy storage potential of various devices such as batteries, water heaters, and space-heaters. However, few studies address the important question of the nature of the offers that aggregators of residential loads can reasonably bid in markets [7].

This paper proposes a flexibility product that a water heater aggregator could offer. water heaters have significant potential because they are already present in large numbers in many residential contexts. However, to achieve this, the aggregator needs a control strategy for the consumption of large numbers of water heaters, and this large-scale coordination involves various challenges. The recent SMARTDesc project [8] explored the possibility of using water heaters for load management and developed a decentralized control architecture to manage them using a relatively recent theoretical development, namely mean field game theory [9, 10].

Mean field game theory is at the intersection of statistical mechanics, game theory and optimization. This theory is devoted to the analysis of games with a large numbers of players who have negligible individual impact but collectively create a stable mass effect [11, 12]. The corresponding controls have several advantages, including decentralization and communication parsimony. In this paper we adopt the mean field control strategy as it seems ideally suited for the large-scale control problem at hand.

There are two main contributions in this paper. The first one is the extension of a near fixed-point algorithm, first proposed for space-heaters in [13], to the case of water heaters. It is a way of deriving a decentralized control law of individual water heaters that yield an overall behavior consistent with the aggregator's target. The second one is the development of a scheduler that produces mean target temperatures for the group of controlled water heaters, inspired by [9]. These mean target temperatures correspond to the maximum load decraese or increase that the aggregator can offer while preserving customer safety and comfort, and meeting post-control power oscillation amplitude constraints upon restoration of ordinary thermostatic control.

This paper is structured as follows. In Sections 2 and 3 we summarize the water heaters model and the control used for this work and developed in [10]. Section 4 presents the near fixed-point algorithm, and Section 5 described the optimization problem to evaluate the possible flexibility offer. Section 6 discusses the simulation results obtained for the proposed strategy. Section 7 concludes the paper.

2 Water heater model

To describe the dynamics of the temperature of a water heater, we model the tank using n equal volume layers with uniform temperature as shown in Figure 1. This model reflects the stratification of temperature in a typical tank.

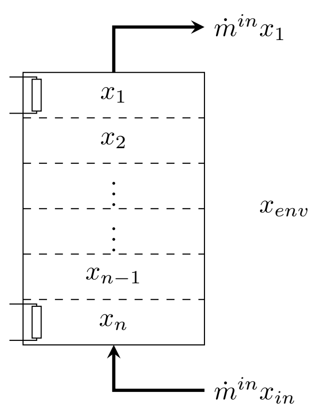

Figure 1: Water heater model

The thermal dynamics are described by the energy balance in each layer:

$$M_l C_{pf} \frac{dx_{l,t}}{dt} = UA_l(x_{env} - x_{l,t}) + \dot{Q}_l \bar{u}_{l,t} + \dot{m}_t^{in} C_{pf}(x_{(l+1),t} - x_{l,t}),\ l \neq n, \tag{1a}$$

$$M_l C_{pf} \frac{dx_{l,t}}{dt} = UA_l(x_{env} - x_{l,t}) + \dot{Q}_l \bar{u}_{l,t} + \dot{m}_t^{in} C_{pf}(x_{in,t} - x_{l,t}),\ l = n. \tag{1b}$$

Note that $\dot{m}_t^{in}$ is modeled as a piece-wise constant process with the extraction and transition rate evolving according to a continuous time Markov chain,with states θ_t taking values in $\Theta = \{1, 2, \dots, p\}$, and an infinitesimal generator denoted $L_t = (L_{q,k})_{q\in\Theta,k\in\Theta}$ as introduced in [14]. Each state j corresponds to a type of water event (no event, dishwashing, showering, etc.) and is associated to a specific flow $\dot{V}_j^{mix}$ so that $\dot{m}_t^{in} = \dot{V}_{\theta_t}^{mix}$. Note also that we consider only two heating elements, one at the top of the tank, and one at the bottom. Equations (1) can be written in linear form for all $l \in \{1, \dots, n\}$ with $x_t = (x_{1,t}, \dots, x_{n,t})^T$, $\bar{u}_t = (\bar{u}_{1,t}, \dots, \bar{u}_{n,t})^T$, and $\bar{A}(\theta_t)$, B, $\bar{c}(\theta_t)$ the matrices resulting from this change of notation. The thermal dynamics thus take the form

$$\frac{dx_t}{dt} = A(\theta_t)x_t + B\bar{u}_t + \bar{c}(\theta_t) \tag{2}$$

In the optimal control strategy detailed in Section 3, in order to keep the customers comfortable, we do not penalize the effort to maintain the water heater at its temperature at the start of the control horizon, but rather the effort to deviate from it when aiming for a different temperature. The thermal effort to remain on average at the initial temperature is thus obtained for free in our formulation, and the thermal dynamics can be written as

$$\frac{dx_t}{dt} = A(\theta_t)x_t + Bu_t + c(\theta_t) \tag{3}$$

where $A(\theta_t)$ and $c(\theta_t)$ are modified from (2) to account for the free effort.

3 Control strategy

The control strategy we implement was introduced in [10]. It is based on so called mean field game theory, a theoretical development that occurred in the past 15 years [11,12]. We consider a homogeneous group of N water heaters, in particular with identical layer structure and water extraction statistics. Their mean temperature is to follow a target temperature of y. As customary in game theory, a cost function is attributed to a generic individual water heater i as follows:

$$\begin{aligned} J_i^N(u^i, j, t) =& E\left[\int_t^T \left[(Hx_{i,\tau} - z)^2 q_\tau^y\right] d\tau | \theta_t = j\right] \\ &+ E\left[\int_t^T \left[(Hx_{i,\tau} - Hx_{i,0})^2 q^{x_0}\right] d\tau | \theta_t = j\right] \\ &+ E\left[\int_t^T \left[||u_{i,\tau}||_R^2\right] d\tau | \theta_t = j\right] \\ &+ E\left[(Hx_{i,T} - z)^2 q_T^y | \theta_t = j\right] \\ &+ E\left[(Hx_{i,T} - Hx_{i,0})^2 q^{x_0} | \theta_t = j\right] \end{aligned} \tag{4}$$

where:

- $q_t^y = \left|\lambda \int\limits_0^t (H\bar{x}_\tau - y)d\tau\right|$;
- z is set to x^{low} if the objective is to decrease the mean aggregate temperature, and to x^{high} if the objective is to increase that temperature;
- $\bar{x}_t = \sum\limits_{i=1}^{N} \frac{1}{N} x_{i,t}$ is the vector of mean temperatures of the water heaters;
- $H = \begin{pmatrix} \frac{1}{n} & \frac{1}{n} & \cdots & \frac{1}{n} \end{pmatrix}$;
- $||u_{i,t}||_R^2 = (u_{i,t})^T R u_{i,t}$.

Let us remark that $x_{i,t}$ and $\bar{x}_t$ are vectors whose dimension is the number of layers in the tank. Thus Hx_t^i is the mean temperature of water heater i.

This formulation of the cost function is unusual in that the cost coefficient q_t^y generating the pressure to go toward z (first term of (4)) is an *integral cost* depending on the deviation from the target. This means that the pressure (either to store energy or to decrease energy power consumption) continues to build as long as the mean temperature has not reached the target temperature. This temperature change is partially countered by the second term of (4) which penalizes deviations from the agent's initial temperature. Thus, each agent reaches its own specific steady-state with a mean temperature somewhere between the initial temperature and temperature z, while the overall mean temperature for the set of water heaters reaches the target y. This happens while minimizing relative temperature changes in each water heater. Futhermore, those water heaters that can contribute the most are subject to the highest pressure, and contribute accordingly when computing their best response policy. The third term of (4), limits the contribution of each water heater to the global effort, in order to favor local customer comfort. The last two terms represent the final cost.

When the number of controlled water heaters is very large, the laws of large numbers dictate that the aggregate mean temperature vector $\bar{x}_t$ converges to a deterministic (yet a priori unknown) trajectory. Because that trajectory no longer depends on the actions of individual agents, (4) can be viewed as an isolated agent, leading to a classical optimal tracking problem. Viewed as a tracking problem for a linear quadratic regulator, this problem can be solved through a system of Riccati equations with variables Π_t^j and offset variables $s_{i,t}^j$ [10].

This system can be used to compute the control we need to apply to each individual agent in order to achieve the common goal of reaching the target temperature. The system depends on the unknown q_t^y and to obtain it, we need to consider that *individuals optimally responding to the assumed q_t^y must collectively produce a mean temperature response $\bar{x}_t$ such it replicates the assumed q_t^y*. When this condition is fullfilled, one can claim that the Nash equilibrium of the game has been reached. The above argument implies that we need to find the fixed point of the following system:

$$rClq_t^y = \left|\lambda \int_0^t (H\bar{x}_\tau - y)d\tau\right| \tag{5a}$$

$$-\frac{d\Pi_t^j}{dt} = \Pi_t^j A^j + {A^j}^T \Pi_t^j - \Pi_t^j B R^{-1} B^T \Pi_t^j + \sum_{k\in\theta} L_{j,k}\Pi_t^k + (q_t^y + q^{x_0})H^T H,\ j \in \Theta \tag{5b}$$

$$\Pi_T^j = (q_T^y + q^{x_0})H^T H,\ j \in \Theta \tag{5c}$$

$$-\frac{ds_t^j}{dt} = (A^j - BR^{-1}B^T\Pi_t^j)^T s_t^j + \Pi_t^j c^j - (q_t^y z + q^{x_0} H\bar{x_0})H^T + \sum_{k\in\theta} L_{j,k} s_t^k,\ j \in \Theta \tag{5d}$$

$$s_T^j = -(q_t^y z + q^{x_0} H\bar{x_0})H^T,\ j \in \Theta \tag{5e}$$

$$\begin{aligned}
\frac{d\bar{x}_t^j}{dt} &= (A^j - BR^{-1}B^T\Pi_t^j)\bar{x}_t^j + \sum_{k\in\theta} L_{k,j}\bar{x}_t^k \quad (5f)\\
&\quad +\zeta_{j,t}c^j - \zeta_{j,t}BR^{-1}B^T s_t^j,\ j\in\Theta\\
\bar{x}_t &= \sum_{j\in\theta}\bar{x}_t^j \quad (5g)\\
\bar{x}_t^j &= E(\mathbb{1}(\theta_t = j)\bar{x}_t),\ j\in\Theta \quad (5h)\\
\text{where } \zeta_t &= [\zeta_{1,t},\cdots,\zeta_{p,t}] \text{ is defined by: } \frac{\partial\zeta_t}{\partial t} = \zeta_t L^T \quad (5i)
\end{aligned}$$

Computing the fixed point of (5a) corresponds to finding the global strategy of the mean field game. Then q_t^y can be used to find the individual control law of each agent. The algorithm used to find the fixed point is described in the next section.

4 Near fixed-point algorithm

While it is established in [10] that under some technical conditions a fixed points always exists, it may not always be desirable, i.e, associated with a bounded q_t^y as t goes to infinity (or equivalently, a mean water heater temperature which converges to the target y). We instead look for a desirable near fixed-point such that the trajectory converges to y when $t\to\infty$, which means that the cost coefficient trajectory $q_t^y(\lambda)$ must converge to some q_∞^y satisfying the steady state equation of (5a) with $\bar{x}_\infty = y$. As the convergence to a fixed point with an iterative algorithm highly depends on the choice of the integration coefficient λ, we modify our approach relative to [10], inspired by the near fixed point calculations in [13], to rely on the solution of a suitable optimization problem.

Let $\mathcal{S}_\lambda(\bar{x}_t(\lambda))$ be the solution of (5a) for mean temperature trajectory $\bar{x}_t$ and coefficient λ in the definition of q_t^y. We want to select the trajectory that is closest to a fixed point within a family of mean trajectories $\bar{x}_t(\lambda)$ that possess the correct steady-state behavior. This family, first introduced in [13] for space heaters, is constructed as follows. Let $N_q > n_q > 1$ and $t_0 > 0$. We solve system (5a) with the cost coefficient

$$q_t^y = \begin{cases} n_q q_\infty^y & \text{if } t\in[0,t_0]\\ q_\infty^y & \text{if } t\geq t_0\end{cases}$$

to obtain $\bar{x}_{1,t}$ and with

$$q_t^y = \begin{cases} N_q q_\infty^y & \text{if } t\in[0,t_0]\\ q_\infty^y & \text{if } t\geq t_0\end{cases}$$

to obtain $\bar{x}_{2,t}$. Although they may not be fixed points, these two trajectories satisfy the correct steady-state behavior and constitute the bounds of the family. The associated lambdas are respectively

$$\lambda_1 = \frac{q_\infty^y}{\left|\int_0^\infty (H\bar{x}_\tau^1 - y)d\tau\right|} \text{ and } \lambda_2 = \frac{q_\infty^y}{\left|\int_0^\infty (H\bar{x}_\tau^2 - y)d\tau\right|}.$$

The family is then defined as

$$\mathcal{F}(f) = \{\bar{x}_t(\lambda)|\lambda = \frac{q_\infty^y}{\left|\int_0^\infty (H\bar{x}(\lambda)_\tau - y)d\tau\right|},$$

$$\bar{x}_t(\lambda) = f\bar{x}_{1,t} + (1-f)\bar{x}_{2,t},$$

$$f\in[0,1], \lambda\in[\min(\lambda_1,\lambda_2),\max(\lambda_1,\lambda_2)]\}.$$

Within this family, we select the trajectory that is the closest to a fixed point using the following optimization problem where the optimiwed variable are on n_q, N_q, t_0 and f:

$$\begin{aligned}
\min \quad & a_1 \overbrace{||\bar{x}_t(\lambda) - \mathcal{S}_\lambda(\bar{x}_t(\lambda))||_{\mathcal{L}_2}}^{\text{so as to be a fixed point}} + a_2 \overbrace{(\mathcal{S}_\lambda(\bar{x}_t(\lambda))(T) - y)^2}^{\text{convergence to the target}} \\
\text{s.t.} \quad & \bar{x}_{1,t}, \lambda_1, \bar{x}_{2,t}, \lambda_2 \text{ computed as described above} \\
& \bar{x}_t(\lambda) \in \mathcal{F}(f) \\
& N_q \in [1, N_q^{max}] \\
& n_q \in [1, n_q^{max}] \\
& t_0 \in [0, t_0^{max}]
\end{aligned}$$

where N_q^{max},n_q^{max} and t_0^{max} are chosen arbitrarily, (4,2,5) in our study.

The choice of the length of the control horizon T and of n_q and N_q allows us to impose somewhat the speed at which we wish the aggregate control to operate. The resulting optimal trajectory $\bar{x}_t$ can either i) be sent by the aggregator to all water heaters so that they implement locally their optimal control policy, or ii), if local computational capacity allows it, the optimization can be carried out locally by each water heater with only the aggregate mean temperature vector communicated to them at the start of the control horizon.

We tried several solvers in Julia to solve this problem: `BlackBoxOptim.jl` [15] (denoted BBO) with a differential evolution strategy (Adaptive DE/rand/1/bin with radius limited sampling), LBFGS and Gradient Descent (denoted GD) from the package `Optim.jl` [16], and finally just an iteration of the family by iterating the variables in their scope. We tested them with several initial and target temperatures; a pair of initial and target temperatures is called an instance. We tested the solvers on instances with differences between initial and target temperatures of $\{0.1, 0.2, 0.3, 0.4, 0.5, 0.6, 0.7, 0.8, 0.9\}$. We used an `Intel(R) Core(TM) i7-2600 CPU @ 3.40GHz` with 8 processors. The results are reported in Table 1.

The *Distance* indicator is computed as the mean of $||\bar{x}_t(\lambda_{opt}) - \mathcal{S}^\infty_{\lambda_{opt}}(\bar{x}_t(\lambda_{opt}))||_1$ over all the instances, where $\mathcal{S}^\infty_{\lambda_{opt}}(\bar{x}_t(\lambda_{opt}))$ is the solution of (5a) after we have iterated several times starting with $\bar{x}_t(\lambda_{opt})$ for $\bar{x}_t$ and λ_{opt} for λ until we converge to a fixed point of the system with a fixed accuracy $\epsilon > ||\bar{x}_t(\lambda_{opt}) - \mathcal{S}_{\lambda_{opt}}(\bar{x}_\tau(\lambda_{opt}))||_1$. If there is no convergence to a fixed point after a large number of iterations the *distance* indicator will be very large. The *time* indicator is the mean CPU time over all instances.

Table 1: Comparison of four optimization solvers

Solver	BBO	LBFGS	GD	Iteration
Nunmber of unsolved instances (out of 549)	0	15	8	0
Distance	0.0232	0.0294	0.0261	0.0118
Time (s)	33.98	33.23	32.86	70.03

All four methods display good convergence as the *Distance* indicator remains relatively small. Looking at the unsolved instances, the LBFGS and Gradient Descent methods perform poorly because they failed many times whereas the other two methods solved all the instances. Finally, between BBO and Iteration, we select BBO because it takes half the time on average.

5 Flexibility product

In this section we describe how to determine the maximum increase/decrease of power consumption that the aggregator is able to achieve with a given set of water heaters. This information is then used in the approach from Section 4 to determine the target temperatures.

We model the set of water heaters as one large aggregated water heater with a unique layer and we compute an energy balance to find the maximum amount of energy that can injected in this large water heater during a given Δ_t. The stored energy can then be converted into a target temperature for the mean field control using the equation

$$y_t = \frac{e_t}{N_{wh}\rho C_{pf}} + x_L.$$

Two conditions need to be satisfied: i) the water heater power demand must be as close to constant as possible during the control horizon, and ii) the rebound after the horizon must lie within an acceptable range, noted $r_{rebound}$. The rebound is the drop/peak in the power demand after the water heaters revert to the thermostatic mode relative to the demand in the no-control scenario. The different time interval and objectives of our flexibility are summaries in Figure 2. To find the maximum achievable flexibility under these two conditions, we developed an algorithm formed of 4 blocks. The first block is a *scheduler* which compute a temperature schedule depending on bounds in the injected energy φ_t, $\varphi^{min} = (\varphi^{min}_{T_{start}}, \varphi^{min}_{T_{start}+\Delta_t}, \cdots, \varphi^{min}_{T_{end}})$ and $\varphi^{max} = (\varphi^{max}_{T_{start}}, \varphi^{max}_{T_{start}+\Delta_t}, \cdots, \varphi^{max}_{T_{end}})$. Initial values of this vectors are $\varphi^{min}_{init} = (0, 0, \cdots, 0)$ and $\varphi^{max}_{init} = (\dot{Q}\Delta_t N_{wh}, \cdots, \dot{Q}\Delta_t N_{wh})$. The second block is a *Simulator* which perform a Monte Carlo simulations of $N_w h$ water heaters under the temperature schedule just computed, the control described in Section 3 and the dynamics described in Section 2. The third one is an *Updater* which actualizes the value of the bounds φ^{min}, φ^{max} depending on whether the rebound constraint is satisfy or not. The last one is the *Convergence test* based on the results of which we go back to the first block or we exit the algorithm. The flow chart of the algorithm can be found in Figure 3.

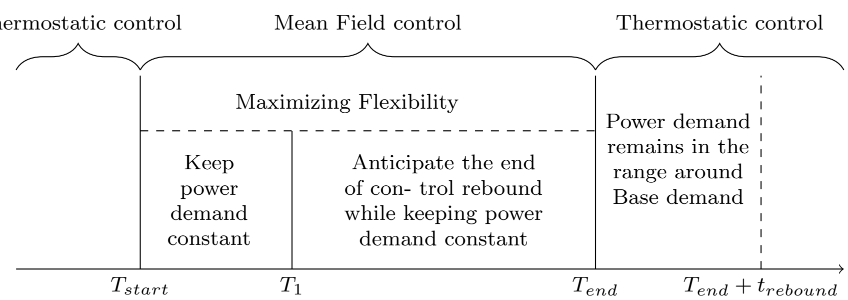

Figure 2: Definition of the goals depending on time intervals

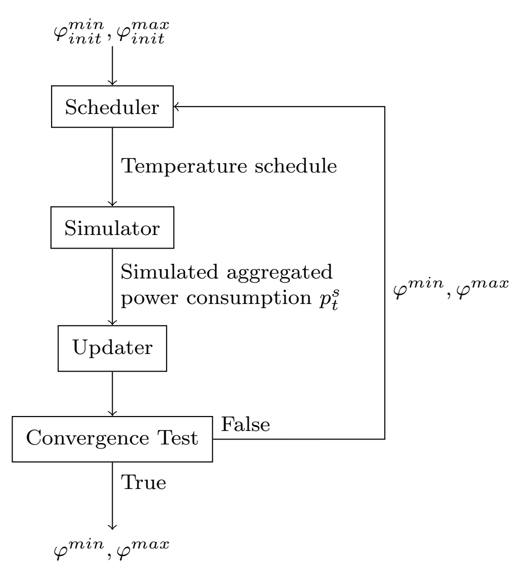

Figure 3: Description of the algorithm to implement the rebound constraint

5.1 Scheduler

The scheduler is an optimization problem, the output of which is a temperature schedule for the mean field controller.

Objective function to be maximized

$$\begin{aligned} Flexibility(\varphi_t, z_t) = a_1 &\sum_{t \in]T_{start}, T_{end}]} C_{direction}(p_t^b - p_t) \\ &- a_2 \sum_{t \in]T_{start}, T_1[} z_t - a_3 \sum_{t \in [T_1, T_{end}]} z_t \end{aligned}$$

The first term maximizes the increase/decrease in power demand by maximizing the distance from the base power demand to the power demand at optimality, while the second and third terms minimize the change in power demand to keep it close to constant on $]T_{start}, T_1[$ and $[T_1, T_{end}]$.

Constraints

- The residential power consumption at each time step is the sum of the uncontrollable part d_t and the controllable part, namely the amount of energy injected in the water heater:

$$p_t = d_t + \frac{\varphi_t}{\Delta_t}$$

- The change in power consumption z_t must be greater than or equal to $|p_t - p_{t-1}|$:

$$\begin{aligned} z_t &\geq p_t - p_{t-1} \\ z_t &\geq p_{t-1} - p_t \end{aligned}$$

- The energy stored in the water heater at time t is equal to the amount of energy stored at $t-1$ plus the energy injected minus the losses.

$$\begin{aligned} e_t &= e_{t-1} + \varphi_t - l(e_{t-1}) \\ l(e_t) &= l_c(e_t) + l_e(t) \end{aligned}$$

  where :

$$\begin{aligned} l_c(e_t) &= UA\left(\frac{e_t}{C_{pf}\rho V} + N_{wh}(x_{in} - x_{env})\right) \\ l_e(t) &= \rho C_{pf}(x_{mix} - x_{in})V^{extract}(t) \\ V^{extract}(t) &= \sum_{j \in \Theta} N_{wh}\zeta_{j,\infty}(t)\dot{V}_j^{mix}\Delta_t \end{aligned}$$

  $l_c(e_t)$ correspond to the loss from heat transfer by conduction with the environnement and $l_e(t)$ is the energy loss due to the water extraction, $V^{extract}(t)$ representing the expectation over the extraction type of volume of water drawn for the aggregated water heater.
- The energy stored in the water heater is linked to the temperature:

$$\begin{aligned} e_{T_{start}} &= N_{wh}\rho V C_{pf}(x_{init} - x_{in}) \\ e_t &\geq N_{wh}\rho V C_{pf}(x_{low} - x_{in}) \\ e_t &\leq N_{wh}\rho V C_{pf}(x_{high} - x_{in}) \end{aligned}$$

  This give us bounds on the energy we are able to store, depending on the initial amount of energy stored.

- The amount of energy that can be injected into the water heater is bounded by φ^{min} and φ^{max} at each time step. $B_{down}(t)$ or $B_{up}(t)$ in (6a), are active depending on the value of $C_{direction}$ and prevents overshooting/undershooting the initial temperature as we act to decrease/increase the power consumption. φ_t is bounded above/below appropriately:

$$\varphi_t \geq \varphi_t^{min} \tag{6a}$$
$$\varphi_t \leq \varphi_t^{max} \tag{6b}$$
$$\varphi_t \geq B_{down}(t-1) \tag{6c}$$
$$\varphi_t \leq B_{up}(t-1) \tag{6d}$$

where :

$$\begin{aligned}
B_{down}(t) =& N_{wh}\rho V C_{pf}(C_+ x_{low} + C_- x_{init} - x_{in}) - e_t + l_t \\
B_{up}(t) =& N_{wh}\rho V C_{pf}(C_- x_{high} + C_+ x_{init} - x_{in}) - e_t + l_t \\
C_+ =& \frac{(1 + C_{direction})}{2} \\
C_- =& \frac{(1 - C_{direction})}{2}
\end{aligned}$$

5.2 Simulator

The *Simulator* performs the simulation of the N_{wh} water heaters individually under the temperature schedule just computed by the *Scheduler* through a Monte-Carlo simulation on the horizon time $[T_{start} - 2h, T_{end} + t_{rebound}]$. We perform this simulation because the power demand during the post T_{end} rebound phase, needed to evaluate the rebound value, depends on the temperature distribution within the water heaters, and this distribution cannot be obtained from an aggregate single water heater model.

5.3 Updater

The *Updater* takes as input the aggregated power demand computed by the *Simulator* p_t^s, p_t^b and $r_{rebound}$.

First, the *Updater* performs a rebound constraint test :

1. If $|p_t^s - p_t^b| \geq r_{rebound} p_t^b$, the rebound constraint is not satisfied
2. If $|p_t^s - p_t^b| \leq r_{rebound} p_t^b$, the rebound constraint is satisfied

Then, it updates the bounds φ_t^{max} and φ_t^{min} using a technique similar to the one used in the bisection method.

- If we are in case 1, the bounds are too permissive. The bounds are updated the following way :
  - If $C_{direction} = 1$: $\varphi_{prev}^{max} = \varphi^{max}$ and $\varphi_{T_{end}}^{max} = \dfrac{(c_d - 1)\varphi_{T_{end}}^{max} + \varphi_{prev,T_{end}}^{min}}{c_d}$
  - If $C_{direction} = -1$: $\varphi_{prev}^{min} = \varphi^{min}$ and $\varphi_{T_{end}}^{min} = \dfrac{\varphi_{prev,T_{end}}^{max} + (c_d - 1)\varphi_{T_{end}}^{min}}{c_d}$
- If we are in case 2, the bounds may not be permissive enough. The bounds are updated the following way :
  - If $C_{direction} = 1$: $\varphi_{prev}^{min} = \varphi^{max}$ and $\varphi_{T_{end}}^{max} = \dfrac{(c_d - 1)\varphi_{T_{end}}^{max} + \varphi_{prev,T_{end}}^{max}}{c_d}$
  - If $C_{direction} = -1$: $\varphi_{prev}^{max} = \varphi^{min}$ and $\varphi_{T_{end}}^{min} = \dfrac{\varphi_{T_{end}}^{min} + (c_d - 1)\varphi_{prev,T_{end}}^{min}}{c_d}$

φ_{prev}^{min} and φ_{prev}^{max} are initialized with φ_{init}^{min} and φ_{init}^{max}.

We can observe that acting on $\varphi^{max}_{T_{end}}, \varphi^{min}_{T_{end}}$ will lead to a corresponding change in the temperature schedule. The last target of the schedule will follow the same trend as $\varphi^{max}_{T_{end}}, \varphi^{min}_{T_{end}}$, when we decrease $\varphi^{max}_{T_{end}}$ it will result in a decrease of the last target temperature for example. Combining this to the last term of the objective function, which imposes that the power consumption remain as constant as possible on $[T_1, T_{end}]$, we will observe in the next section, that the target will gradually increase or decrease on $[T_1, T_{end}]$ to prepare water heaters to shift to thermostatic control and anticipate the post-control rebound.

5.4 Convergence test

We exit the algorithm if $C_+|\varphi^{max}_{prev,T_{end}} - \varphi^{max}_{T_{end}}| + C_-|\varphi^{min}_{prev,T_{end}} - \varphi^{min}_{T_{end}}|$ is small enough or if we reached the maximum permitted number of iterations.

6 Case study

We use in this case study the setup of the SMARTDesc project [8]. We consider 500 identical water heaters with a two-layer tank. The infinitesimal generator L_t of the Markov chain modeling water extraction is piecewise constant every 2h during the day with values taken from [17]. The Markov chain has 2 states $\theta_t \in \{0, 1\}$ that represent the absence or presence of water extraction, and $\dot{m}_t = \dot{V}^{mix}\theta_t$ where $\dot{V}^{mix} = 2.62\ell/\text{min}$ is the extraction flow. The parameters are provided in Table 2.

Table 2: Parameters value for simulations

$\dot{Q}_l$	$4500\ J/s$	A	$2.55\ m^2$
x_{env}	$25°C$	M_l	$136.5\ kg$
x_{in}	$15°C$	C_{pf}	$4190\ J/(kgK)$
x^{low}	$50°C$	U	$28.38 J/(m^2 K min)$
x^{high}	$60°C$	$\dot{V}^{mix}_j$	$2.62\ l/min$
x_{mix}	$38°C$	q^{x_0}	$8000\ h^{-1}$
V	$273\ \ell$	R	$\begin{pmatrix} 0.025 & 0 \\ 0 & 0.025 \end{pmatrix} h^{-1}$

We ran the scheduler described in Section 5 with $T_1 = T_{start} + 2h$ and $T_{end} = T_{start} + 4h$, i.e., the water heaters are controlled for 4 hours, maximizing flexibility during the first 2 h ($C_{direction} = -1$), and anticipating the end of control rebound during the other 2 h ($C_{direction} = -1$). The results are reported in Figure 4 and Figure 5 respectively.

The total power consumption considered is that of the 500 homes and represents the sum of the uncontrollable demand, d_t, and that of the 500 water heaters simulated independently with distinct initial conditions and extraction trajectories. This corresponds to Monte Carlo simulations of 500 water heaters with the parameters stated above in Table 2. The base power demand with which our simulations are compared is the total power consumption of 500 homes. Power demand data is taken from the public demand data report of the Independent Electricity System Operator (IESO) of Ontario [18]. We used the data from January 30, 2019. This data represents the full power demand of Ontario; to obtain the demand of 500 homes we have applied a reduction coefficient, 10^{-4}, to this overall demand to obtain only the power consumption of 500 homes.

We consider the two following cases :

- Case 1.a : Active control between 7a.m and 11a.m. First a 2 hours period of power reduction relative to base power demand, then 2 hours to anticipate the rebound; after 11a.m mean field control ends and one reverts to the classical thermostatic control of water heaters.
- Case 1.b : Active control between 2p.m and 4p.m. First a 2 hours period of power increase relative to base power demand, then 2 hours to anticipate the rebound; after 6p.m mean field control ends and one reverts to the classical thermostatic control of water heaters.

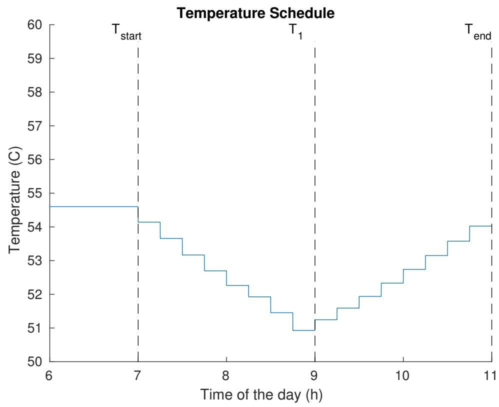

(a) Mean field target schedule temperature

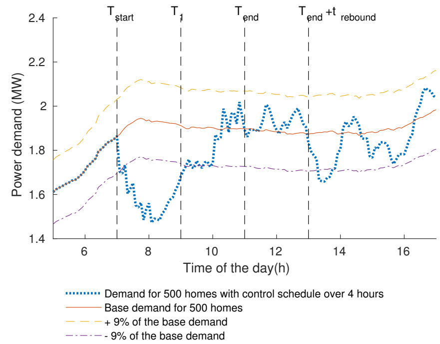

(b) Comparison of baseline power demand and the power demand under the schedule in Figure 4a

Figure 4: Case 1.a - Decrease of the power demand

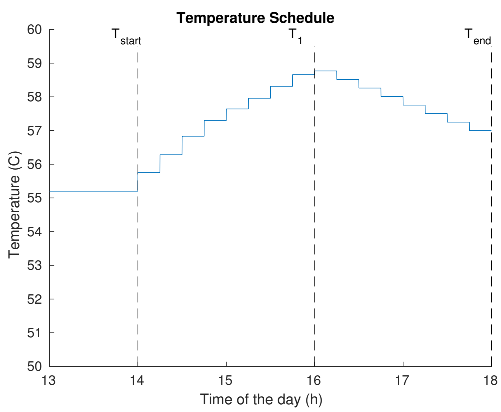

(a) Mean field target schedule temperature

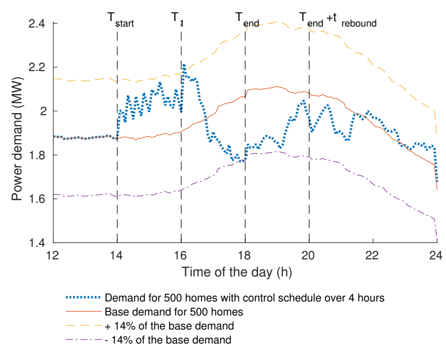

(b) Comparison of baseline power demand and the power demand under the schedule in Figure 5a

Figure 5: Case 1.b - Increase of the power demand

In Figure 4 and Figure 5, we display the output of our optimization for the two cases. Results are very similar. Looking at the target temperature schedules (Figure 4a and 5a), on the first part of the control period (before T_1) the strategy is to have a target temperature that decreases (resp. increases) up to T_1. After T_1 the target starts to increase (resp. decrease) again in order to heat (resp. cool) the water heaters gradually to anticipate the rebound and have the distribution of water heaters temperature not too close to the comfort temperatures. In Figure 4b and 5b, we observe that while achieving the objective of decreasing the power demand (resp. increasing the power demand), the rebound stays within a 9% range (resp. 14% range) of the base power consumption. This rebounds are the minimum we can achieve for the given parameters (T_{start}, T_{end}, $t_{rebound}$, ...).

In case 2, we can see also that in the schedule we reach the upper comfort limit temperature. So if the initial temperature was lower, we anticipate that we could have had more flexibility, as we could have increased the temperature more. An initial phase were we decrease the temperature of the water heaters, to anticipate our need of flexibility, could then be beneficial. Reciprocal situation can appear while decreasing the power demand.

Economic discussion

To assess the economic viability of such an aggregator we consider the case of Ontario with the Ontario time-of-use (TOU) and the canadian dollars currency. TOU prices are used to compare how much consumers would have paid in case where no flexibility control is applied on the water heaters, versus that when it is applied [19]. The difference will indicate the minimum financial incentives to be given to participating customers: 0 if they pay less when we apply control on the water heaters, and the net cost of the flexibility action if customers pay more when we apply control. We then need to evaluate how much the aggregator can sell their flexibility on the market. For this we rely on the *Auction clearing price* of the Demand Response Auction in Ontario [18]. In this auction, bidders commit to provide flexibility every business day of a six month period. Given this two quantities, we can evaluate the possible profit for the aggregator for each flexibility product. Results for 1.a and 1.b are summarized in Table 3. Then we perform a Net Present Value study for a water heater aggregator. We consider that the initial investment will be a participation to the purchase of new water heaters for the 500 homes, with a cost of 425$ for a new water heater. Water heaters are said to have a 10 to 15 years lifetime. Then the net cash flows each month will be the profit obtained with the combination of the two flexibility offers 1.a and 1.b every day and we consider a discount rate of 6% per year, or 0.49% per month [20]. Results of this study can be found in Table 4. With a participation of 40% from the aggregator, the NPV becomes positive in 4.5 years, and with a 50 %, it is in almost 6 years. This type of investment seems to be profitable with a participation rate below 40 to 50% taking into account the lifetime of a water heater.

Table 3: Profit results for different cases

Case number	Duration	Percentage of rebound	Flexibility provided (kW)	Profit (CAD)
1.a	4	9%	212	49.7
1.b	4	14%	44	10.3

Table 4: Results of the profitability study

Discount rate	Participation (%)	Profitability time (month)	NPV(10 y)	NPV(15 y)
0.49%	100	177	-49 196$	2 992 $
0.49%	50	70	57 053 $	109 242 $
0.49%	40	54	78 303 $	130 492 $
0.49%	30	39	99 553 $	151 742 $
0.49%	0	0	163 303 $	215 492 $

7 Conclusion

In this paper, we have proposed a flexibility product for a water heater aggregator. The control strategy used for the water heaters is a mean field approach that is suited for controlling large scale groups. We have adapted a recently developed strategy [13] for solving the mean field control problem for space-heaters to the case of water heaters. More precisely, we used an optimization approach to find a trajectory for the water heaters to follow as an approximate, but guaranteed to meet target, mean field control strategy. Then, to assess the possible flexibility that the group of water-heaters could provide during a fixed interval, we used a linear program that works as a target scheduler for the mean field controller. The output is a schedule that permits to maximize the flexibility provided. To incorporate an additional constraint that ensures that the rebound, at the end of the control period, in the power consumption remain limited compared to the base power consumption, a bisection like method algorithm has been added to iterate the scheduler and find eligible schedules.

In future work, it would be interesting to consider some stochastic optimization for the scheduler to take into account the uncertainty on both the production of renewable sources and on the water extraction, to have a more robust schedule. It could also be of interest to develop a specific solver to compute the near fixed point using the form of the optimization problem to achieve greater efficiency. Also, more detailed economic modelling could give more accurate information on the profitability of such aggregator.

References

[1] A. Mahmood, A. R. Butt, U. Mussadiq, R. Nawaz, R. Zafar, and S. Razzaq, Energy sharing and management for prosumers in smart grid with integration of storage system, in 2017 5th International Istanbul Smart Grid and Cities Congress and Fair (ICSG), 2017, pp. 153–156.

[2] S. M. Basnet, H. Aburub, and W. Jewell, Effect of demand response on residential energy efficiency with direct load control and dynamic price control, in 2015 IEEE Eindhoven PowerTech, 2015, pp. 1–6.

[3] PSA, PowerShift Atlantic Final Report, 2015.

[4] IESO, Demand Response Auction in Ontario, 2019, [Accessed: February 2019]. [Online]. Available: http://www.ieso.ca/en/sector-participants/market-operations/markets-and-related-programs/demand-response-auction

[5] L. Gkatzikis, I. Koutsopoulos, and T. Salonidis, The role of aggregators in smart grid demand response markets, IEEE Journal on Selected Areas in Communications, 31(7):1247–1257, 2013.

[6] C. Eid, P. Codani, Y. Chen, Y. Perez, and R. Hakvoort, Aggregation of demand side flexibility in a smart grid: A review for european market design, in 2015 12th International Conference on the European Energy Market (EEM), 2015, pp. 1–5.

[7] A. Mohsenian-Rad, V. W. S. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, Autonomous demand-side management based on game-theoretic energy consumption scheduling for the future smart grid, IEEE Transactions on Smart Grid, 1(3):320–331, 2010.

[8] F. Sirois, B. Bourdel, and R. P. Malhamé, Project RENE-034 : managing energy storage capacities dispersed in an electrical grid to reduce the effects of renewable energy source variability, 2017.

[9] A. I. Tammam, Lissage optimal de la charge électrique en présence de sources d'énergies renouvelables via le pilotage de la consommation des chauffe-eau, Ph.D. dissertation, École Polytechnique de Montréal, 2016.

[10] A. Kizilkale and R. Malhamé, Chapter twenty - collective target tracking mean field control for Markovian jump-driven models of electric water heating loads, in Control of Complex Systems, K. G. Vamvoudakis and S. Jagannathan, Eds. Butterworth-Heinemann, 2016, pp. 559–584.

[11] J. Lasry and P. Lions, Jeux à champ moyen. ii–horizon fini et contrôle optimal, Comptes Rendus Mathématique, 34(10):679–684, 2006.

[12] M. Huang, R. P. Malhamé, P. E. Caines et al., Large population stochastic dynamic games: closed-loop mckean-vlasov systems and the nash certainty equivalence principle, Communications in Information & Systems, 6(3):221–252, 2006.

[13] A. C. Kizilkale, R. Salhab, and R. P. Malhamé, An integral control formulation of mean field game based large scale coordination of loads in smart grids, Automatica, vol. 100, pp. 312–322, 2019.

[14] R. Malhamé, A jump-driven markovian electric load model, Advances in Applied Probability, 22(3):564–586, 1990.

[15] R. Feldt, Blackboxoptim.jl, https://github.com/robertfeldt/BlackBoxOptim.jl, 2018.

[16] P. K. Mogensen and A. N. Riseth, Optim: A mathematical optimization package for Julia, Journal of Open Source Software, 3(24):615, 2018.

[17] B. Bourdel, Gestion de stockage d'énergie thermique d'un arc de chauffe-eaux par une commande à champ moyen, Master's thesis, École Polytechnique de Montréal, 2016.

[18] IESO, Power Data - Data Directory, 2019, [Accessed: February 2019]. [Online]. Available: http://www.ieso.ca/en/Power-Data/Data-Directory

[19] Ontario Hydro, Ontario TOU prices, 2019, [Accessed: February 2019]. [Online]. Available: http://www.ontario-hydro.com/current-rates

[20] Jean-René Tagne Kuelah, L'évaluation du rendement d'une dépense publique, 2016, [Accessed: June 2019]. [Online]. Available: https://www.economie.gouv.qc.ca/fileadmin/contenu/publications/administratives/rapports/rendement_depense_publique.pdf