Project RENE-034: Managing Energy Storage Capacities Dispersed in an Electrical Grid to Reduce the Effects of Renewable Energy Source Variability – PUBLIC REPORT

Frédéric Sirois, Benoît Bourdel et Roland Malhamé

Département de Génie électrique, Polytechnique Montréal, Montréal (QC), Canada

Last update: July 22nd, 2017

Other contributors with the proponent:

Arman C. Kizilkale, J. Coulombe, S. Fan, F. Li, R. Losseau, F. Malandra, K. Ratelle, J. Solis, A. I. Tammam, M. Anjos, M. Bernier, M. Gendreau, B. Sanso

Partners:

IREQ – Laboratoire sur les Technologies de l'Énergie (LTE), WPred inc., Artelys, SG2B, Coopérative St-Jean-Baptiste

TABLE OF CONTENTS

LIST OF FIGURES

Chapitre 1

Executive Summary

Project RENE-034 received funding between 2012 and 2016 as part of the Canadian federal government's ecoENERGY Innovation Initiative (ecoEII), in the "R & D Program" category. The project proponent was Ecole Polytechnique de Montréal, in Montreal, Canada. The project was carried out in cooperation with several Canadian partners, namely the Laboratoire sur les Technologies de l'Énergie (LTE) at IREQ (Institut de Recherche d'Hydro-Québec), WPred Inc., Artelys, SG2B, as well as the Coopérative St-Jean-Baptiste (a small electricity distributor in Montérégie).

The objective of the project, which the proponent renamed smartDESC (smart Distributed Energy Storage Controller), was to provide a *scalable* approach to harness the potential of the millions of dispersed energy storage associated devices in an electrical grid, such as electric water heaters, space heaters, electric vehicle batteries, etc. as a tool to help shave peak system loads or smooth the fluctuations of intermittent renewable energy sources (mainly wind or solar). The scientific underpinning of the smartDESC project is so-called Mean Field Games theory, a recent large scale systems control theoretic development whereby control objectives are achieved by means of a game formulation ; this turns individual devices into decision making agents thereby achieving the crucial property of *decentralization* of the control. Furthermore, the large scale nature of the control problem (millions of contributing devices) is turned into an advantage in that it makes it possible to capitalize on the *predictability* arising from the law of large numbers. As a result, the massive information exchanges between the local level and a more general coordination level typically needed for such control architectures can be significantly reduced (reduced communications requirements). The function of the coordination level is to compile statistics about the controlled devices group dynamics for model building purposes and, based on system level information, to compute and communicate the desired collective behaviour to local controllers, in the form for instance of a short term target mean energy level for the population of controlled devices, as well as its desirable but feasible target evolution over the next hours.

Thus summarizing, the major advantage of the approach developed here is that it drastically reduces telecommunications requirements on electric distribution networks, which often have only a very limited bandwidth to work with (for example that of a network of smart meters), while devices make their own local decisions. Moreover, once the millions of controlled storage sites can be organized to behave collectively as directed, they can be used as a giant battery to help smooth renewable sources fluctuations, and raise their permissible penetration levels. Generating adequate targets requires weather forecasts based short-term predictions of expected renewable generation, together with that of the projected electricity consumption of customers connected to the electrical grid to which the storage sites belong. The information is subsequently processed in a centralized location and used to generate the desired collective behaviour that is transmitted to the local controllers through a telecommunications network.

All of the components needed to complete a simulation based proof of concept for the mean field control of a population of electric water heaters within an electrical distribution network were developed as part of this project. This primarily entailed developing models of (i) electric water heaters dynamics, (ii) the associated random hot water demand stochastic processes, and (iii) the telecommunications network involved (mesh network of smart meters). Mathematical developments were also required in terms of the control algorithms themselves, as well as to optimize the overall collective behaviour of the storage elements based on a specific objective, e.g., flattening the electricity consumption curve, maximizing wind penetration, etc. All of these developments ultimately took the form of computer codes in various environments and languages (Java, C++, Fortran, etc.). An overall simulator had to be developed, based on the JADE multi-agent platform, to ensure the interoperability of all of these codes in a common computational space.

Finally, an extension received in the last four months of the project also made it possible to build a physical platform to demonstrate the feasibility of a hardware implementation (physical controller operating on a real water heater). This physical “node” was set up so as to be able to operate as an element (agent) of the simulation and, therefore, in this sense is also part of the simulation based proof of concept.

To date, the mean field control architecture proposed as part of the smartDESC project has been shown, through simulated case studies, to successfully meet the desired objectives. Indeed, the overall simulator made it possible to establish the benefits expected from harnessing the collective smoothing potential of the storage sites while keeping low bandwidth requirements. The project can thus be considered as successfully completed.

Chapitre 2

Project Description and Execution

2.1 Project objective

The main objective of the smartDESC project was to develop a decentralized control architecture for the energy storage capacities found in an electrical grid, but geographically dispersed among its customers (e.g., electric water heaters, air-conditioned spaces, hybrid electric vehicle batteries, etc.) or elsewhere in the grid. The proposed architecture was intended to offset the undesirable effects of the temporal variability inherent in both the load on the grid and the renewable component of electricity generation (wind, solar energy). It goes without saying that this storage site management scheme had to be designed under the paramount constraint of *preserving customer comfort*. At the end of the smartDESC project, the potential of the global control architecture including 1) ability to flatten the electricity consumption curve ; 2) ability to achieve a smooth load recovery (e.g., following a blackout) ; and 3) ability to help smooth the mismatch between wind generation and grid electricity demand, was to be demonstrated through various case studies inspired by actual network data.

To accomplish these aims, the project was divided into several development stages, as follows : (i) modelling the individual behaviour of the loads in question (primarily water heaters for the proof of concept) and their aggregate behaviour ; (ii) generating, through mathematical programming, the desirable mean behaviour for these loads in order to offset sources of variability in the grid's generation or load ; (iii) setting up decentralized control using mean field theory in order to achieve the desired collective behaviour ; and (iv) verifying the approach using an overall simulator to : 1) estimate the additional gains made possible in terms of renewable energy penetration ; and 2) test the ability of the architecture to respond to unforeseen events.

The financial aspect (time dependent energy pricing) was not part of the smartDESC approach, but it would be possible to include it in the future should the need arise.

2.2 Necessity of the project

The need to generate more electricity from renewable sources is unanimously recognized internationally. It is also an important part of Canada's energy planning, given the country's geographic expanse and the associated huge potential for development of the renewable energy industry. However, apart from hydroelectricity, which is a very reliable and flexible energy source due to the large amount of energy stored in water reservoirs, the major difficulty associated with integrating intermittent renewable energy sources (sun, wind) into electrical grids is their variability over time. There is never any guarantee that the power required by consumption centres can be provided exclusively by such energy sources at any given time. The convenience of being able to store produced renewable energy when it is available is therefore a priority in any energy strategy based on intermittent renewable energy sources.

Because large-scale electrical power storage is still difficult considering the high level of consumption, one possible approach is to fully decentralize energy storage by making the greatest possible use of the storage potential that already exists at customer sites, especially the thermal energy stored in an electric water heater or even in buildings themselves (furniture, walls, etc., also accumulate thermal energy) [Hughes, 2010]. However, the number of storage sites to control quickly becomes huge, requiring an unconventional control architecture to make the most of this storage potential inherent in electrical grid customers. The challenge is even greater if this storage management is required to be completely transparent in terms of customer comfort.

Most conventional approaches to load management and, by extension, to the management of distributed energy storage in an electrical grid, are so-called "centralized" approaches, in the sense that load control depends on a signal from a mainframe that has to be updated regularly based on intensive feedback that provides the state of the controlled loads to the mainframe. This is achieved at the cost of substantial telecommunications bandwidth, thus rendering impractical the application of such controls to very large systems. If bandwidth is limited, then the rate at which information is refreshed generally has to be compromised, which has a particular impact on 1) customer comfort (increased unavailability of hot water, etc.) ; and/or 2) overall system performance in terms of the integration of renewables (less wind penetration, etc.).

Against this backdrop, there was a need to explore options for controlling millions of storage sites distributed in an electrical grid without increasing telecommunications requirements disproportionately, and while ensuring satisfactory customer comfort at all times. The smartDESC project proved that, when applied to this specific problem, mean field game/control theory is one possible solution.

2.3 Identification and solicitation of partners

When the project was launched in 2012, the proponent had identified a certain number of partners, the main one being Hydro-Québec, which at the time was working to develop a domestic load management strategy to help the company reduce its wintertime peak in electricity consumption. More specifically, the expertise of the Laboratoire sur les Technologies de l'Énergie (LTE) at IREQ (Institut de Recherche d'Hydro-Québec) in modelling energy storage elements proved to be a major asset to the smartDESC project. Discussions between the LTE researchers and the smartDESC project team sped up the model development and validation phase considerably, especially for the water heater models and the models of the process according to which residential customers draw hot water. Other discussions with the LTE, IREQ and Hydro-Québec Distribution also helped to expand the potential scope of the architecture proposed as part of smartDESC and especially to clearly define the case studies used to complete the simulation based proof of concept.

Another project partner is WPred Inc., a company that produces wind power forecasts within a given geographic area based on publicly available weather data from Environment Canada. There is also Artelys, which developed an electricity demand forecast module for the smartDESC team based on historical consumption data from previous years. Finally, the Coopérative St-Jean-Baptiste [Coop, 2016], a small private electricity distributor located in Montérégie, provided complimentary electricity consumption anonymized data for its network of 6,000 customers over a two-year period, enabling the smartDESC case studies to be based on real (and realistic) data. The combined contributions of WPred, Artelys and the Coop therefore made it possible to develop case studies involving wind power fluctuations balancing on an electrical grid. The latter was the most important challenge for the completion of smartDESC's intended proof of concept, i.e., demonstrating the possibility of integrating more renewable energy into the electrical grid by means of decentralized control based on mean field theory.

Finally, toward the end of the project, SG2B (Smart Grid to Business), which provides smart-grid consulting services, offered the project team its assistance with planning the building of a physical water-heater controller, as well as a real water-heater test platform controlled via a mean field controller and connected to the overall simulator developed as part of the project.

Large manufacturers of electrical equipment were also approached before and during the project, but because the goal of the smartDESC project was to arrive at a simulation based proof of concept and not to develop commercial equipment, the contribution from these companies was limited to general recommendations aimed at better planning for a potential technological development phase.

During the project, the proponent had no difficulty finding suitable partners to cover the range of expertise needed to complete the mean-field-control proof of concept.

2.4 Difficulties encountered

The project was conducted in a university setting, which meant involving a certain number of master's and doctoral students, who are sometimes difficult to recruit. However, in this case, the recruitment process at the beginning of the project went very well, as did the recruitment of two research associates, one dedicated to background theoretical work on mean field theory and the other to the implementation of the overall simulator. Aside from the fact that it took longer than initially anticipated to decide on the architecture of the overall simulator, the project went as initially planned and the results are exactly where they were expected to be after the 42 months of the project.

Chapitre 3

Summary of Developments

3.1 General overview

The core smartDESC architecture is based on a mean field control architecture, itself based on the theory of mean field games, introduced in 2006 [1], as the limit of non-cooperative games with many players [Huang, Malhamé and Caines 2006, Lasry and Lions, 2007]. The most attractive feature of mean field game theory is the considerable simplification of interactions between players or equivalently "agents". Indeed, agents determine their optimal strategy by considering the evolution of the *mass of other agents* as a whole rather than all individual behaviours (that is to say, each of the other agent' actions considered individually).

In our particular power system context, and for the purposes of the smartDESC project, water heater/controller pairs are considered as the agents. Using mean field game theory, it is possible to achieve optimal control of a large population of these agents by strictly minimizing communications through local calculations that anticipate the behaviour of the group. The decentralized implementation of the control also safeguards the possibility of a device temporarily opting out of the control at any given time if the computed control is anticipated as leading to customer discomfort or compromised safety. The control was only limited to water heaters in order to focus as many project resources as possible on the proof of concept. In no way does this limit the application of the mean field control to other types of storage elements, such as building heating systems, electric vehicle batteries, etc.

In order to demonstrate the feasibility and advantages of such large-scale decentralized control, the smartDESC project was broken down into a certain number of tasks, the main ones relating to the development of different computer modules, including :

- a module digitally simulating water-heater dynamics [Solis, 2015],
- an electric water heater aggregate electricity consumption profile optimization module,
- a local water heater control module based on mean field game theory,
- a module simulating a mesh telecommunications network based on smart meters like those installed by Hydro-Québec (and many other electric companies),
- a multi-agent simulation module based on JADE (Java Agent DEvelopment Framework).

It was the multi-agent simulation module that allowed all of the other modules to be integrated into a unique environment. Once this integration was complete, case studies could be set up that made it possible to draw some conclusions about the performance of the mean field control based architecture.

Finally, in the last few months of the project, an actual electric water heater controlled by a physical prototype of a mean field controller was set up to demonstrate the practical feasibility of the concept.

This chapter summarizes the main achievements associated with the tasks listed above. The next chapter presents the results of the case studies.

3.2 Module development and integration into the simulator

Every subsection set out here begins with an introduction to the theoretical principles associated with the module in question and ends with practical information about the implementation of the module in the simulator.

1. Professor Roland Malhamé, one of the two main investigators on this project, is one of its pioneers.

3.2.1 Module digitally simulating water-heater physics

Theoretical principles

The water heater (or hot water tank) is a ubiquitous electric household appliance. Its main function is to maintain a supply of hot water so that it is continuously available for household users. It is therefore an excellent way of storing thermal energy and especially, if done correctly, a way of storing that energy at a strategic time in relation to energy supply and demand on the electrical grid.

There are many types of water heaters (gas, instantaneous, thermodynamic), but the most common is the conventional electric water heater. The operating principle is simple : one or two electric elements are immersed in a tank of water. The water is heated locally by the Joule effect by passing an electric current through one or both of the elements. The temperature is then equalized by convection movements between the cold and hot water due to the difference in their density.

Buoyancy forces hot water to sit on top of cold water, which means that the temperature gradient is opposite to the gravitational field. Therefore, in order to respect this dynamic, the cold water inlet is located at the bottom of the water heater, whereas the hot water outlet is at the top.

When hot water is drawn, cold water is instantly added at the bottom of the water heater so that the volume of water in the water heater remains constant. Stratification of the water in the water heater can be observed because cold water is denser than hot water. The lower, cold part of the water heater is then heated by the heating element at the bottom of the water heater.

The temperature of water heaters is also a critical factor in the health safety of their users. While water that is too hot may cause burns, water that is not hot enough promotes the growth of bacteria inside the water heater. The temperature therefore has to be kept within well-defined limits.

Conventional electric water heaters are controlled thermostatically. Thermostatic controls are defined by two variables : the set temperature (usually around 60 °C) and a tolerance zone (*deadband*, usually ±3 °C). So, when the water temperature reaches the set value less the tolerance zone ($60 - 3 = 57$ °C) in the vicinity of a heating element, the element comes on at its rated output (a few kW). When the heated water reaches the set temperature plus the tolerance zone ($60 + 3 = 63$ °C), the heating element is no longer supplied with electricity. This is the classic phenomenon of hysteresis found in domestic thermostatic controls. In conventional thermostatic mode, the top element, which is closer to the hot water outlet, takes absolute priority over the bottom element to ensure that hot water is always available to the user.

Practical aspects

The water heater was modelled as a Java code. Such modelling makes it possible to simulate the operation of a water heater in a reliable and realistic way, as described below. It is based on the water heater segmentation into N horizontal layers, each of which has a uniform temperature and exchanges heat with the layers above and below, the walls of the water heater, and any heating elements. This water heater model is illustrated in Figure 3.1. An algorithm also approximates convection movements to prevent the lower layers from getting hotter than the top layers, which would not be natural.

In the simulation, the drawing of water is generated based on a so-called Markov chain. It is a two-state chain :

- a low state (0) equivalent to no water being drawn : the flow rate is zero,
- a high state (1) equivalent to water being drawn : the flow rate is then equal to the average output of a water draw event, i.e., 1.31 ℓ/min for the purposes of this study.

The Markov chains were generated based on data from actual measurements taken by Hydro-Québec. They are combined with aggregate water-draw parameters. These are Markov chains with time-varying parameters, i.e., the values of the rates of transition between the two water-draw states change depending on the time of day, to reproduce the statistics of the real water extraction processes as faithfully as possible.

Two versions of Markov chains are used : one for weekdays (Monday to Friday) and one for weekends (Saturday and Sunday). The inherently random nature of Markov chains allows for some variability in individual water-draw scenarios. This variety of scenarios would exist even if the parameters of the Markov chains used were common to the entire population of water heaters controlled.

Under certain conditions, this water heater model yields linear equations that describe the water heater's thermal operation. The linearity of the equations is a prerequisite for an analytical solution to the mean field (control equations (see section 3.2.3).

The equations pertaining to the temperature of the water heater also make it possible to link the water heater's mean temperature to its energy state. The water heater's energy state, expressed in kWh, reflects the amount of energy injected into the water heater to raise it from its “cold” state (water heater filled with cold water) to a given

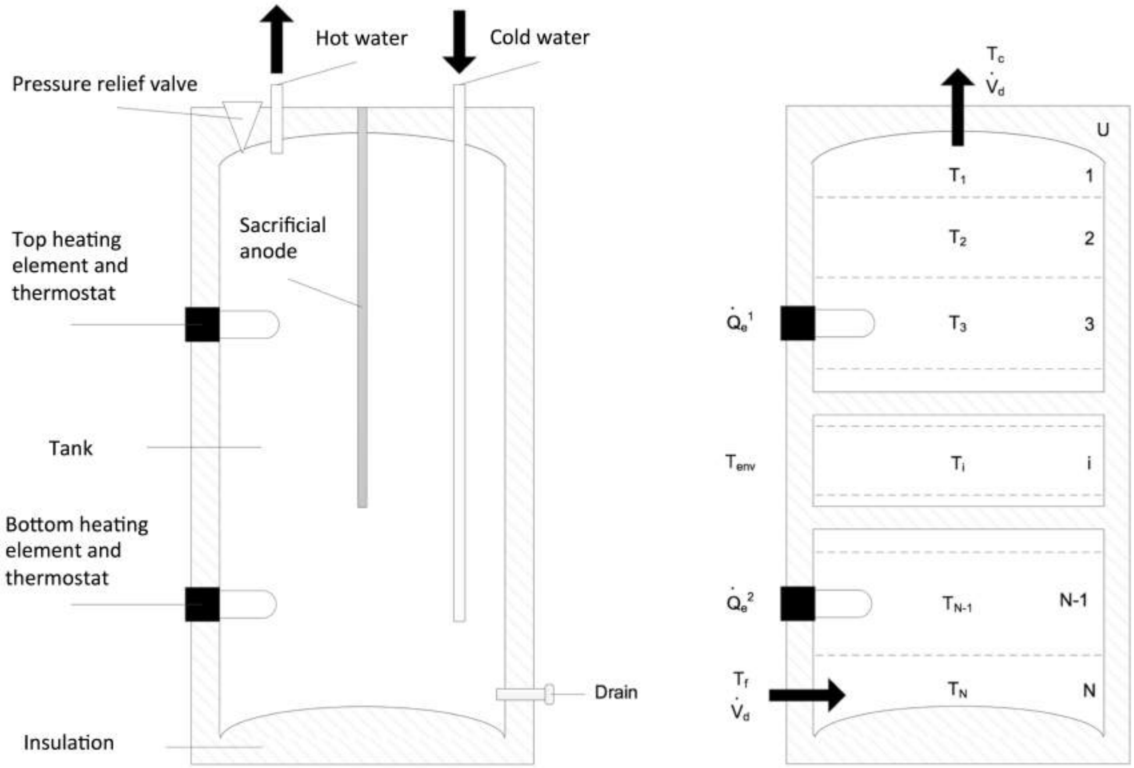

FIGURE 3.1 – Diagram of a water heater and its N layer model.

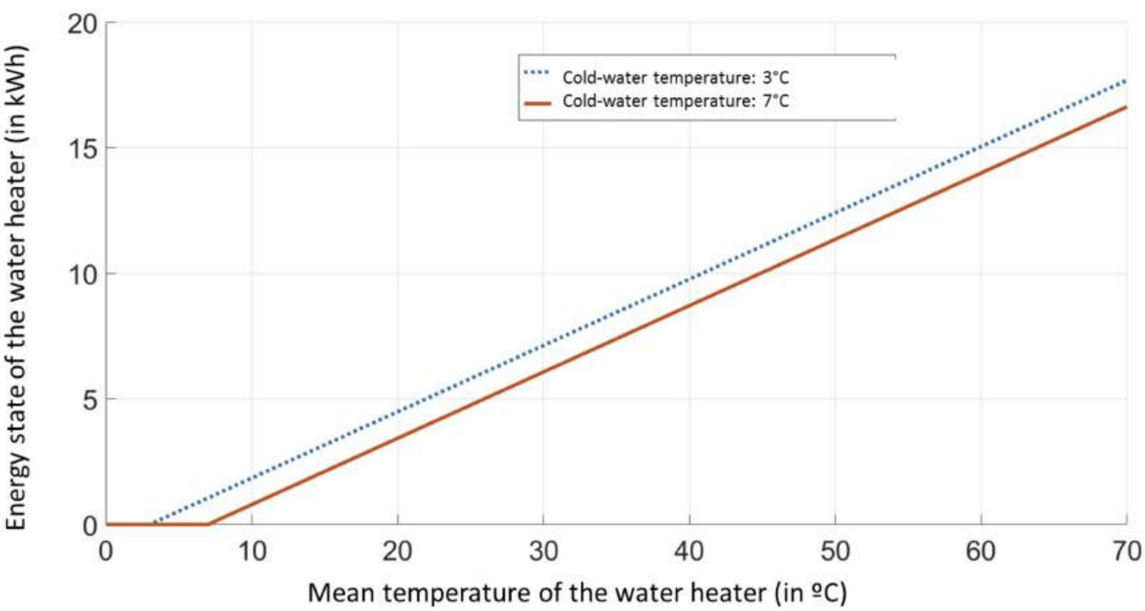

FIGURE 3.2 – Energy state of a 60-gallon water heater as a function of its mean internal temperature and of the cold-water temperature.

mean temperature. In other words, it is the amount of energy stored in the water heater. Figure 3.2 depicts this relationship between the water heater's mean temperature and its energy state for two different input cold water temperatures [2].

The water heater simulation module operates jointly and interactively with the "house" module, which is described in section 4.1 and models complete residential electricity consumption, and with the mean field control module, which is described in section 3.2.3. The control module makes it possible to determine the power that will be stored in the water heater during the next time interval in the simulation.

3.2.2 Optimization module for the aggregated electricity consumption profile

Theoretical principles

It is critical to continuously balance electrical power generation and consumption in order to maintain a stable electrical grid (constant voltage and frequency). It is therefore preferable for the load profile to be as "smooth" and stable as possible so that the level of generation does not constantly have to be changed, which is difficult with thermal or nuclear power plants. In that case, so-called "peak load" power plants, which are often more costly and polluting than so-called "base load" power plants, have to be started up whenever consumption peaks. Conversely, during consumption troughs, base load power plants have to be regulated outside of their optimal operating point. This reduces their performance and increases production costs. Even for hydro-electric power plants which can be quite flexible in their operation, adjustments that are too frequent tend to wear out the valves that control the flow of water. A smooth electricity demand profile is therefore the most economical scenario for an electrical grid operator with a "conventional" generation mix.

The growing penetration of intermittent renewable (and therefore uncontrollable) energies, such as solar photovoltaic and wind, creates additional stresses by adding generation side variability to demand side variability, thus making the load/production balancing problem even more difficult than with traditional grid generation mixes. Under these circumstances, having some control on the load side (aggregate load profile) clearly becomes an asset for the grid operator as it constitutes an additional degree of freedom to cope with this variability. The advantage of such a control is both technical, since it enables the penetration of a greater proportion of renewable energy while limiting stability problems of the grid, and economical, since it helps limiting electricity demand when energy is more costly to generate. At the moment, these advantages are perceived primarily by electricity distribution companies, and not generally by their customers, unless the former offer sufficiently more favorable rates to participating customers.

That said, the aim of the smartDESC project is not to generate energy savings at the customer end, but rather to achieve temporal control over energy storage through customers' water heaters in a way that is completely transparent to them, while substantially benefiting electricity distributors in terms of the operating costs and flexibility of their grid. In other words, the aim of the optimization called for here is to shift customers' water heater-related consumption in time, without changing their consumption habits or net energy balance. In this sense, the optimizer's role is to provide sequences of energy storage targets that help level as much as possible the net load that is perceived by the conventional inertial portion of the grid generation mix, whether or not there is renewable generation on the grid. In the former case, improvements are mainly due to wind/solar power balancing, while in the latter, improvements are on the side of the global electricity demand itself.

Practical aspects

In any event, the algorithm outputs 24 mean temperatures (corresponding to target energy levels stored in the complete water heater population), i.e., a target mean temperature for each of the next 24 hours. The mean temperature of the entire population is to reach the intended target at the end of every hour in order to achieve the desired optimization.

In fact, to achieve its purpose, the algorithm applies its optimization to a hypothetical water heater that represents the aggregate of all of the water heaters in the controlled population. The input to this virtual "mega-water heater" is the sum of demands to draw hot water from the individual water heaters, assuming that this aggregate behaviour is known, for example, following a measurement campaign carried out in advance. In the future, one could consider gathering the statistics of local hot water extraction rates based on observations of the load through local controllers installed at customers' locations. The algorithm also uses a certain number of parameters as inputs, such as the physical features of the water heaters, the number of houses controlled, the mean temperature of the water heaters at the beginning of the control horizon (which can be determined based on the

2. The temperature of the cold water depends on the season. It is around 3°C in winter and 10°C in summer

last target temperature from the previous day) and the non-controllable consumption projections for the coming day.

Using all of these data, the optimizer is able to generate the 24 target mean temperatures that are to be transmitted to all of the controlled water heaters (see Figure 3.3).

The algorithm can also operate in “stochastic” mode. Given that the main source of variability in stochastic modelling is the long-term forecasting of weather conditions, the algorithm will then use N weather scenarios as inputs, with as many different probabilities of occurrence, instead of a single, supposedly known and perfect, scenario. This especially affects the wind power expected in the 24 hours covered by the optimization. Furthermore, if a more refined model had been considered for the non-controllable fraction of the load, the temperature projections would also have a significant impact on this via the buildings' electric heating/air conditioning systems.

The stochastic optimizer therefore factors all of these scenarios into its calculation by weighting each by its probability of occurrence. At the end of optimization, the algorithm yields a single set of 24 mean temperatures that considers the uncertainty of occurrence of each of the proposed N scenarios. Although the resulting solution is less perfect than the one obtained for a single, fully known future scenario, it should provide greater immunity against potential discrepancies between projections and reality.

In terms of connection to the overall simulator, the optimization program acts as a black box with well-defined inputs/outputs. In fact, the code is written in C++ and then converted into a Linux shared library (.so) that is readable by the multi-agent simulator (coded in JAVA) through the multi-language JNA (Java Native Access) interface.

Finally, it should be noted that at the start of each day, or whenever deemed necessary, such as when an unexpected event causes reality to deviate significantly from projection, it becomes important to generate a new control to prevent too much error accumulation. In this case, the optimization algorithm is once again invoked with the various above-mentioned parameters to calculate a trajectory that better matches the new reality. The code then provides 24 new temperature set points, which are transmitted to all connected water heaters. These set points are received and processed by the mean field (local) controller, which will now be described in greater detail. The optimizer therefore operates on a rolling horizon and updates its calculations every hour when it receives more current information.

3.2.3 Local mean field control module

Theoretical principles

As indicated in section 3.2.1, as of 2016, the power injected into most water heaters is controlled exclusively by thermostat, meaning that the water heater consumes power from the electrical grid as soon as it begins to cool down, irrespective of other consumption on the same grid. The mean field controller developed as part of the

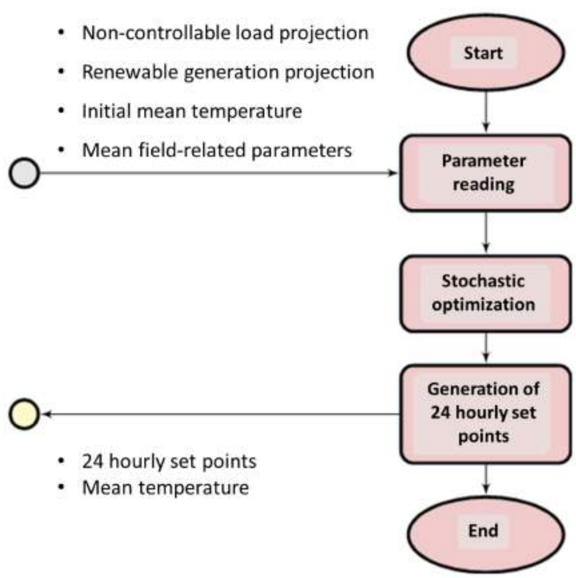

FIGURE 3.3 – Block diagram representation of the load optimization algorithm.

smartDESC project is designed to replace the conventional thermostatic control. It has to be able to receive and interpret the general control transmitted by the optimization algorithm described in section 3.2.2, and to apply it locally without the need for higher-level communication, until the next time the information is refreshed.

Going back to the explanations provided in the background to the project, the objective of smartDESC is to show that every water heater can be supplied with an individual control that uses only minimal bandwidth in terms of telecommunications, while maintaining the best possible user comfort. So, the "trajectory" calculated and sent to the local controllers by the load profile optimizer is common to all of the water heaters, but is interpreted locally by the mean field controller in a way that is tailored to the individual water heater based on the local environment. This way, by following the principles of mean field game theory [Huang et al., 2006], we ensure that, for a sufficiently large (in the statistical sense) set of controlled individuals, the trajectory of the mean temperature of the fleet of water heaters converges toward the group trajectory provided by the optimizer.

Each mean field controller would, in practice, be a "smart" box connected to the heating elements in the water heater it controls and to a telecommunications network. Based on the information the optimizer receives from above and the water heater current thermal state, its role is to generate a feedback law on the water heater local state vector as a function of time for each of the two heating elements. In other words, the mean field controller translates, in real time, the group mean temperature set point expressed as 24 hourly temperatures into a locally calculated sequence of feedback laws to be applied to each heating element in a controlled water heater.

The mean field control's ability to generate an individual control that is locally adjusted to a water heater state and dynamic parameters is one of the strong points of the concept. In fact, it is important to understand that the central controller does not give the local controllers an individual set point, but only very brief information about the desired mean behaviour for the group. The individual set point is generated locally and adjusted to each water heater's specific situation. Achieving such decentralized control without any unwanted side effects is no mean feat.

There is also a strong element of fairness (or equity) in the mean field control algorithm. In fact, the individual cost structure at play is such that the coldest water heaters contribute less than the hottest water heaters to peak reduction efforts (and, conversely, to filling consumption troughs). Note also that the relative energy state of the water heaters in relation to one another varies constantly during the simulations (see section 4.2.2). So, the hottest water heaters in the studied population may be among the coldest water heaters a few hours later. Over the long term, a certain balance is maintained in contributions to the common objective.

To achieve its purpose, the mean field has to maintain some variability in the mean temperature of each individual water heater compared with the mean temperature of the group. This is critical to avoiding excessively large and inefficient temperature shifts in the water heaters, which would always be constrained to move toward the desired mean temperature for the group.

The mean field control algorithm that has been implemented can currently operate only with piecewise constant trajectories. The consequence of this is that jumps in the target mean temperature every new hour may result in small, grid-wide consumption peaks until the spike has passed (a few minutes). Formulae allowing for a smooth trajectory in segments that are linear and continuous from one interval to the next would have to be developed in order to address this issue (future work). In the meantime, to mitigate the problem, the mean field algorithm performs a linear interpolation of the hourly control into an interval control that changes every 15 minutes. This makes it possible to smooth the power control somewhat by reducing the magnitude of jumps between two intervals (four small jumps per hour instead of one large hourly jump).

Practical aspects

Like the load optimization algorithm, the mean field controller module takes the form of a Linux shared library written in C++ and interfaced with the simulator through JNA.

Every time the load optimization algorithm generates a new 24-temperature control, the control is sent to all of the controllers so that they update the optimal control law to be followed for the water heater coupled to them.

At every time step in the simulation, the water heater queries the mean field controller about the power it should inject into each heating element. The controller decides the total electrical power the water heater will have to consume based on information about the desired trajectory of the group, the vector of temperatures in the different strata of the water heater, and whether or not water is drawn (see Figure 3.4).

Finally, the controllers transmit their mean temperature up the chain (i.e. to the mainframe) randomly and according to a Poisson distribution (uniform probability over time). This random sampling enables the optimization algorithm to keep its population sample current and, therefore, to correct its estimated statistical distribution of the entire population. Based on the experiments conducted as part of the project, it seems that an average of four heat recoveries per water heater per day have to be considered in order to maintain comprehensive information that is statistically representative of reality for a water heater population of only a few thousand.

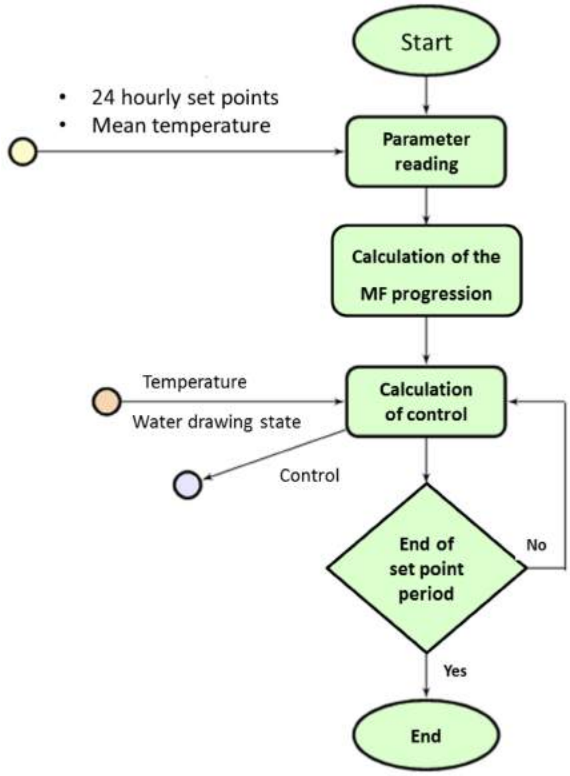

FIGURE 3.4 – Block diagram representation of the mean field control algorithm.

To date, the computing time associated with the mean field algorithm is 1 to 10 seconds per local controller on a 3 GHz processor. In fact, this is much faster than is actually necessary, since the local controller's processor has very little to do once the calculation is complete. In other words, it would not be necessary to have a very powerful microcontroller in order to implement the concept.

3.2.4 Module simulating a "mesh" telecommunications network

Theoretical principles

The load control simulations would not be entirely realistic if they did not include an aspect relating specifically to the telecommunications network and to the transfer of information between the different actors in the system.

The module responsible for performing the telecommunications network simulation is based on the operation of a "mesh/grid" network. This was done in order to reproduce the operation of the electricity meters installed by some North American electric companies, including Hydro-Québec. This establishes that each house acts as a node in the telecommunications grid and is responsible for receiving and transmitting information from one direct neighbour to another, until the message reaches its final destination.

Practical aspects

In the simulator, the telecommunications network is the nerve center of the relationship between the different agents. At every time step, the agents send the telecommunications module the messages they need to transmit, as well as the corresponding recipients. The telecommunications module then performs a simulation separate from the energy simulation, with a much shorter time step (suitable for the telecommunications world, i.e. in the range of a millisecond instead of minutes), in order to simulate the transfer of messages from their starting nodes to their destination nodes.

The telecommunications simulation factors in the phenomena of interference, packet loss and distance between the different houses, as well as the packet routing, transmission and retransmission policies. When a message reaches its recipient in the telecommunications network simulation, the telecommunications module then transmits the message to the recipient agent in the main (energy) simulation. In other words, two simulations are always running jointly, i.e. :

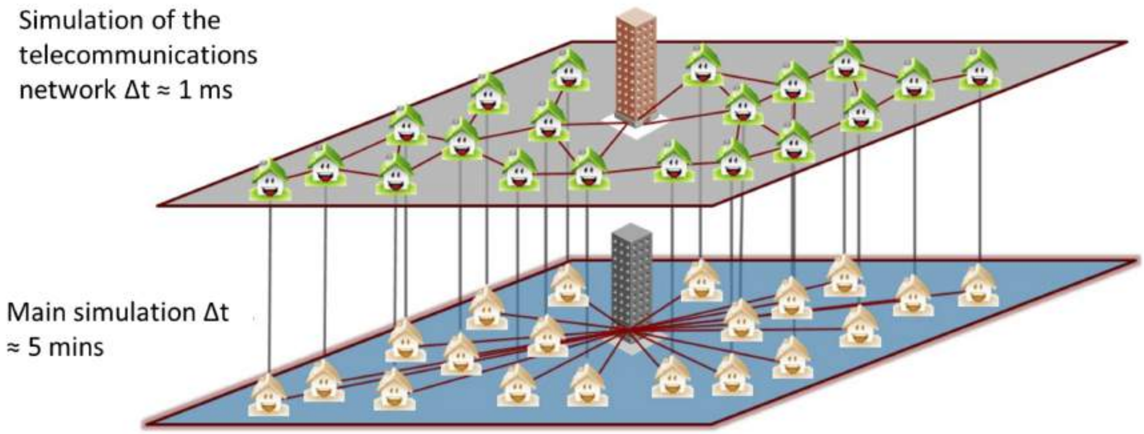

FIGURE 3.5 – Schematic representation of both levels of simulation carried out simultaneously, i.e., the energy simulation (bottom) and telecommunications simulation (top).

1. The main simulation, with a time step of a few minutes, which deals primarily with energy aspects related to the heating of the water heaters and to load control,
2. The telecommunications simulation, with a much faster time step (in the range of a millisecond), which deals only with phenomena related to the transfer of messages between actors in the main simulation. All messages pass through this simulation, making it possible to generate the period between transmission and receipt.

Figure 3.5 depicts this duality between the two simulations, which are simultaneous, synchronous and consistent. The two networks have different topologies of connection between agents. Indeed, while the topology of the main simulation is calculated based on electricity distribution lines, the topology of the telecommunications simulation is based on the architecture of the telecommunications network, the nodes of which are at the geographic location of the houses, where the smart meters are installed.

3.2.5 Multi-agent simulation module (overall simulator)

This final module, called the overall simulator for the purposes of the smartDESC project, is much more generic than the specialized modules described above. It is actually a code based on the JADE (Java Agent DEvelopment Framework) environment, which was designed to implement multi-agent systems. In this case, the agents are the specialized modules described above, repeated enough times to model the desired system. These agents are the actors in the simulation, and they use the mechanisms provided by JADE to interact with one another.

This section introduces the overall operation of the simulator comprising all of the modules described above. Figure 3.6 depicts the simulator's software architecture. It clearly shows all of the different modules, as well as the connections between them.

The "Mainframe" agent is at the top of the figure ; it centralizes both :

— load optimization operations, managed by a C++ shared library (in green),
— the telecommunications network (in pink), which transfers messages between the mainframe and the water heaters,
— the participating water heaters mean temperature estimation function, which retrieves messages from the water heaters and transmits them to the optimization algorithm,
— the simulation's backup and log function.

Then, there are the House/Water Heater pairs, with multiplicity N ranging from 1 to 10 000 [3]. The "House" agent plays two specific roles :

— generating the non-controllable load based on a data file,
— transmitting this load and the water heater's load to the mainframe. This transmission of electrical information does not go through the telecommunications network. It uses a direct link through the mechanisms provided by JADE.

The "Water Heater" agent includes the following modules :

3. Based on 10,000 houses, the number of simultaneous connections to certain configuration files exceeds the maximum value authorized by the Linux operating system. Coupling with a database system should increase this value.

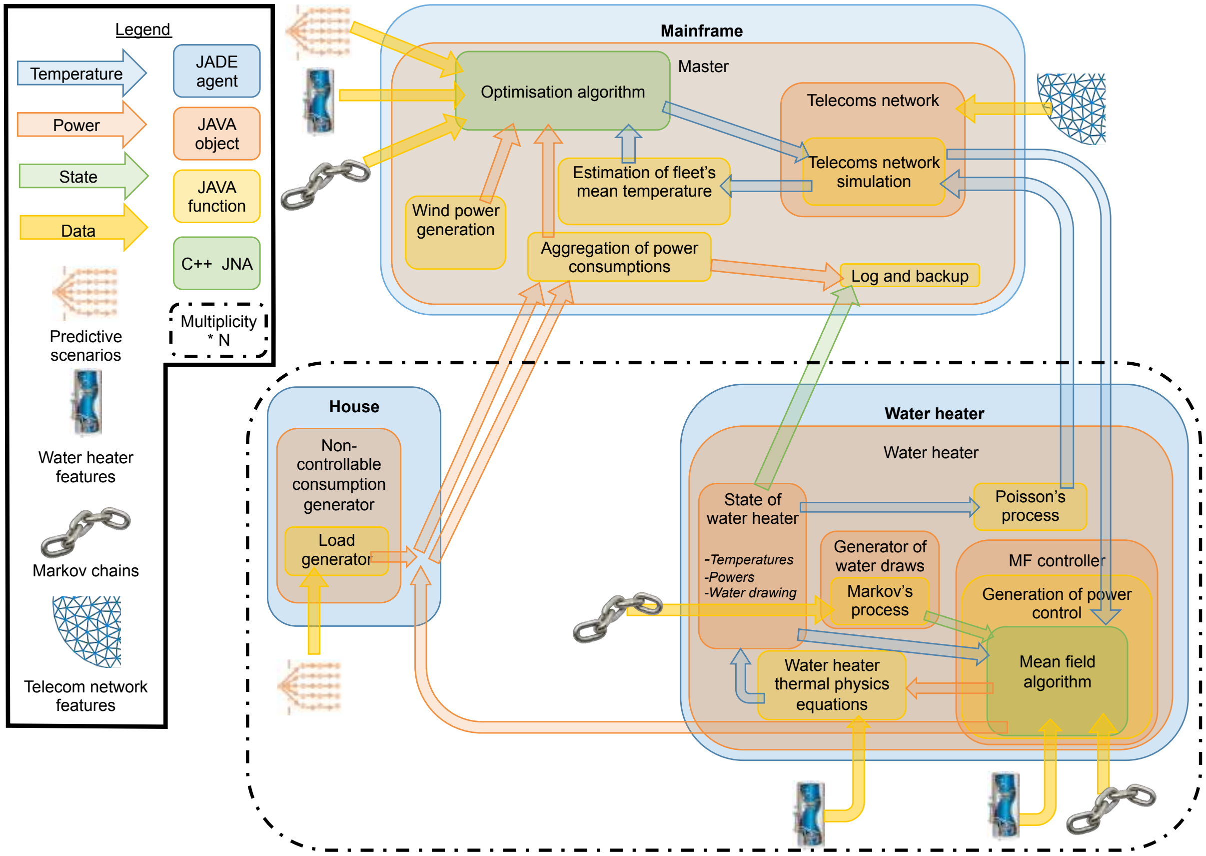

FIGURE 3.6 – Diagram of the overall simulator and of the interactions between the different agents and modules within the agents.

— the mean field algorithm, which receives the 24 set temperatures and the mean temperature of the controlled water heaters from the mainframe at the beginning of the optimization horizon, and which generates a power to be injected into the water heater based on the state of the latter and presence or absence of hot water extraction,
— the algorithm for thermal simulation of the water heater, which retrieves this power and generates a new thermal state of the water heater,
— the water draw sequences generator, based on a Markov process whose parameters are variable over time.

Despite the generic programming of the JADE environment and the implementation quality of the JAVA code needed to manage time simulations under these circumstances, it took considerable work to get all of the above modules operating consistently. On the whole, however, once the loop is closed, the simulator operates as expected : the mean field algorithm serves its purpose by performing the controls generated by the optimization algorithm, and the optimization algorithm uses the mean field algorithm to generate its control.

Although the initial aim was to bring the simulator to a level of maturity where the source code could be shared publicly, this objective could not be achieved. It is worth mentioning that the vision for the simulator changed significantly between the start of the project in 2012 and the end in 2016, as did the underlying libraries and computer technologies. Instead of focusing massive effort on updating code at the end of the project, it was considered more important to spend that time building a physical platform consisting of a water heater and mean field controller (see next section), in order to make the concept more tangible to the industry and encourage the launch of a phase of *demonstration* of the concept on a fleet of real water heaters. Still, the code for the simulator could be shared upon request, without any guarantee of technical support since the project team has now disbanded.

3.3 Development of a physical platform

A physical platform was set up in the last four months of the project. It consists of an actual water heater, a physical prototype of a mean field controller and various instrumentation devices (for flow measurement ; voltage, current and electrical power measurement ; measurement of the water heater's internal temperature over 10 layers, etc.). The purpose of this physical platform was to :

1. confirm that the thermal model of the water heater is realistic,
2. establish the feasibility of physically producing the local controller based on the mean field control concept,
3. stimulate industry interest in a mean field control demonstration phase using a real fleet of water heaters.

A microcontroller inside the controller itself manages all aspects of control, receiving as inputs information (emulated) from the mainframe, as well as local measurements of internal water heater temperatures. The microcontroller can return information for the mainframe as an output as needed, especially the water heater's measured (or estimated) mean temperature. For testing purposes, all information is exchanged through an Ethernet or WiFi telecommunications network. No "mesh counter" network has been emulated on the physical platform to date. Note that the code for the mean field controller used in the simulator has been fully loaded onto the microcontroller.

The external components of the controller include a solenoid valve on the platform for simulating the drawing of water in a manner that exactly follows the Markov chain model programmed into the simulator, with a constant flow rate while water is being drawn. An electronic circuit for modulating the electrical power transmitted to the water heater elements was also developed so that the practical application of the mean field control follows the approach taken in the simulation. Note that the technological choices that led to the construction of the platform are not discussed in detail here because some aspects are subject to intellectual property protection. However, apart from a special sensor that had to be developed to measure the water heater's internal temperature in 10 levels, all development used commercially available equipment. Figure 3.7 shows a few pictures of the physical platform just before it was commissioned.

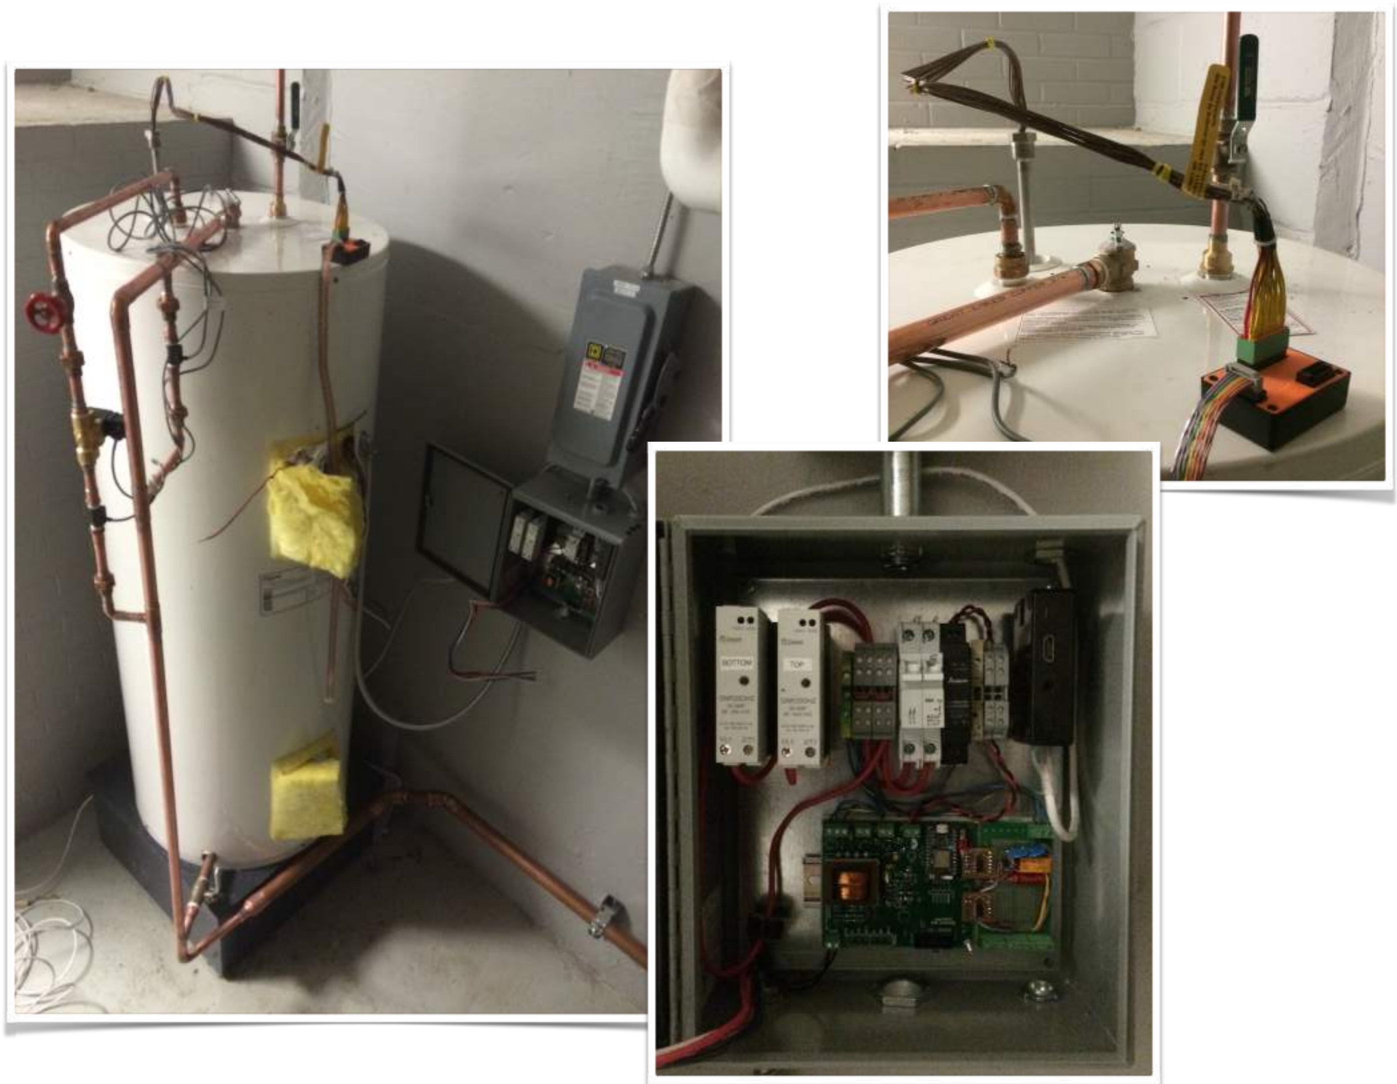

FIGURE 3.7 – Photos of the physical platform set up at Polytechnique Montréal to test the hardware implementation of the mean field controller on a real water heater. The platform is fully equipped with instruments, and the drawing of water is controlled by a solenoid valve.

Chapitre 4

Summary of Results

4.1 Set-up of case studies

4.1.1 Working assumptions

All of the case studies developed as part of this project were based on the electricity consumption data provided by Coop St-Jean-Baptiste. These data were used to identify the rate at which electrical power is consumed by a pool of 6,000 customers over a period of two years. The data were preprocessed to remove customers not representative of the residential sample, especially farms, businesses and a few industries.

Once this processing was completed, the proportion of the consumption associated with water heaters, called the "controllable load" here, had to be separated from the rest of the consumption, called "non-controllable load", which includes all other electric household appliances [1]. This disaggregation of the domestic load was done using the hot water drawing statistics provided by Hydro-Québec and the water heater simulator, all through a special routine written in Java in the overall simulator. Although imperfect, this *a posteriori* load separation made it possible to carry out the proof of concept outlined in this chapter.

Houses are actually supplied with electricity through an electrical distribution network (medium-voltage grid, e.g., 12 to 25 kV), itself powered by a high-voltage grid (>60 kV) via an electrical transformer station. In fact, houses are not connected directly to the medium-voltage grid, but rather to a local low-voltage grid (120/240 V in North America). For the purposes of this project, the decision was made to model the individual behaviour of at least one hundred and no more than a few thousand houses in order to 1) ensure some statistical diversity of the overall load, which is required in order for the mean field control to operate properly; and 2) stay within the capacity limits of the overall simulator, which is capped at about 10 000 simultaneous agents for operation on a single machine. This number of houses corresponds approximately to the electrical load that one or more distribution lines from a single substation has to supply during winter peak periods. In order to simplify the model and focus development efforts on the core innovation, the decision was made not to specifically model the distribution network's electrical behaviour. Because the network's architecture is generally radial and the network itself is relatively short in urban and semi-urban areas, it was deemed acceptable at this stage of development to disregard voltage drops and losses in the network, and to consider the total load delivered to customers to be equal to the load measured upstream from the lines in question. The reactive power on the network has therefore been disregarded for the purposes of this study. It would be easy to add a network electrical parameter calculation module to the overall simulator. The proponent already has this simulator, which simply needs to be properly interfaced with the simulator to be used in the case studies.

Finally, it is worth noting that all of the power results below are standardized by the number of houses for ease of results comparison between the different case studies. All of the power graphs provided in this chapter therefore show the mean power consumed per house, expressed in kW/house.

4.1.2 Practical aspects

In the simulator, each "House" agent is an independent and autonomous actor in the simulation connected to a single agent called the "Mainframe" (or "Master") in Figure 3.6, in addition to being linked to a "Water Heater"

1. Note that the terms "load," "power" and "electricity consumption or demand" are strictly equivalent for the purposes of this report.

agent, as explained above. One function also independently and randomly generates the house's non-controllable load.

Every simulated water heater transmits its electricity consumption to the associated house at every time step in the simulation. The house also generates its own non-controllable load and transmits the total load consumed (non-controllable + controllable) to the mainframe, which aggregates the consumption of all of the houses. This process emulates the power measurement that would normally happen at a substation, since we opted not to specifically simulate the electrical grid. In a physical implementation of the mean field control, this information about individual consumption would not be sent to the mainframe ; only the aggregate measurement from the substation would be transmitted. For optimal emulation of this process, the mainframe in the simulation works only with the aggregate load value, representing the total electricity consumed by the population of houses.

The population's non-controllable load is obtained using real data provided by Coop St-Jean-Baptiste. In this case study, we use the week of 1st january 2014 illustrated in Figure 4.1, which represents the consumption peak for winter 2014-2015. This way, we create difficult conditions for ourselves in order to substantiate our levelling ability.

4.2 Case 1 : load levelling and deterministic optimization

4.2.1 General considerations

The purpose of the first case study outlined here is to level the overall load seen by the electrical grid using thermal energy storage management in a population of water heaters. The grid here is made up exclusively of residential loads (houses), as explained in the previous section. The control circuit used is the mean field control, and the optimization process is deterministic. A single non-controllable consumption scenario is therefore assumed for the next 24 hours when calculating the target mean temperatures.

The simulations are organized according to the following logic :

— Fifteen minutes before midnight, the overall optimizer generates the target mean temperatures for the next 24 hours and transmits them to all of the houses via the telecommunications system. The mean temperature of the water-heater fleet is also transmitted to the houses,

— At midnight, each water heater mean field controller analyzes the data received and sequentially generates the structure of the state feedback control laws for the next 24 hours,

— At each time step of the simulation (five minutes here), each water heater mean field controller determines the power to be injected into the water heater [2] based on the calculated laws, the temperature of each layer

2. Note : In developing the mean field control, it was assumed that the power of the heating elements could be modulated between 0 and 100% of the maximum power of a conventional heating element.

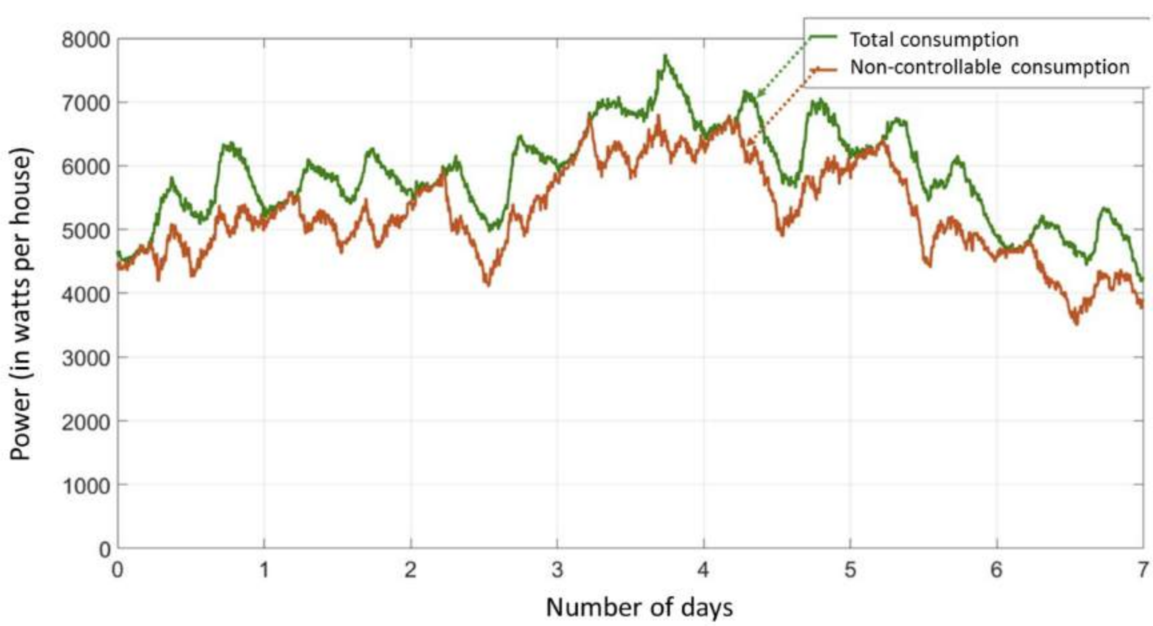

Figure 4.1 – Non-controllable consumption for the week of December 30, 2013, to January 5, 2015.

of the water heater (estimated based on only two measurements) and the local hot water draw state,
— At approximately four random times per day, the water heater local controller transmits the water heater mean temperature to the mainframe via the telecommunications system. This makes it possible to update the estimated mean temperature of the population, an information needed by water heater controllers at the start of the next 24 hours calculations cycle.
— All of the modules described in the previous chapter are operational during the simulation for the case study.

4.2.2 Results

Overall behaviour of the mean field control

Figure 4.2 hows the different power curves of interest after the mean field control has been applied to a population of 400 houses fitted with electric water heaters in case study 1. It includes the following curves :
— in blue : controllable consumption, which represents the consumption of water heaters fitted with a mean field controller,
— in purple : the optimal consumption curve, i.e. the target of the deterministic optimizer based on future consumption projections,
— in yellow : the total consumption of houses with water heaters fitted with mean field controllers (here, all 400 houses,
— in green (as a basis for comparison) : the total consumption of houses in the *absence* of any control (so, conventional water heaters with thermostatic controllers).

The mean field control clearly demonstrates its ability to reduce both consumption peaks and troughs, and therefore to achieve good load levelling. The yellow curve, which is the result of the simulation with mean field control, is close to the purple optimal curve. Operation by plateau for the load is observed throughout these seven days.

The "microvibrations" observed in the simulated total curve and optimal curve are due to the fact that optimization occurs hour by hour, while the power varies every five minutes. From this perspective, the purpose of the mean field is not to regulate these fluctuations, but rather to manage power over a longer time scale (from a few hours to a few days).

So, with the mean field control, the winter peak in 2014, which happens on the 3rd day, drops from 7,800 W to 7000 W which represents an 800 W absolute, or 10% relative, reduction.

It is important to point out that water heaters have a limited available energy storage margin. In fact, the temperature of water heaters has to be kept between 50 and 60°C. As a result, the optimizer cannot completely level the load on a 24-hour horizon : instead, it returns one optimization per level that follows the general consumption trend for the day.

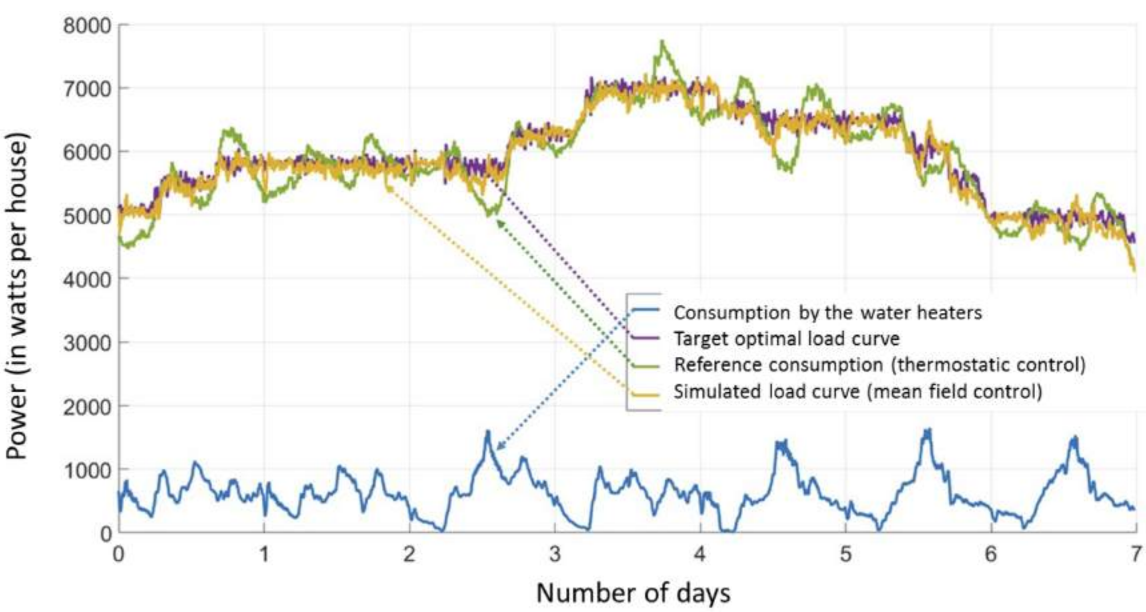

FIGURE 4.2 – Result of the mean field control for case study 1, applied to a population of 400 houses.

Mean temperature and power of the water heater population

Figure 4.3 depicts the outcome of the simulation in terms of water heater temperature. The red line is the target mean temperature calculated by the optimizer, and the blue is the simulated mean temperature for the population of 400 houses all fitted with a mean field controller. A point is depicted every five minutes along each curve.

This figure illustrates the mean field controllers' excellent ability to follow the mean temperature target when the mean temperature follows an increasing trajectory ("loading" period of the water heater fleet). However, the two curves uncouple when the target mean temperature drops below 52°C. This uncoupling can be explained by two factors :

This figure illustrates the mean field controllers' excellent ability to follow the mean temperature target when the mean temperature follows an increasing trajectory ("loading" period of the water heater fleet). However, the two curves uncouple when the target mean temperature drops below 52°C. This uncoupling can be explained by two factors :

- The central optimizer always considers a normal distribution of the water heater population around the mean temperature. However, this approximation is valid only in the domain $[52, 58]$°C. Near the 50 and 60°C limits, the local controller (mean field or not) does not allow the temperature to rise above 60°C or drop below 50°C for customer safety and comfort reasons, resulting in a distortion of the presumed normal distribution, which then crashes at these limits, thereby biasing the optimizer's projections,
- The water heaters' heating and cooling mechanisms are not the same. During heating, power is injected into the water heater, which is always possible. Conversely, cooling requires customers to draw hot water, which is a stochastic process. The controller can therefore do nothing to accelerate the unloading of the water heaters if its pace is slower than the optimizer projected.

Figure 4.4 depicts the mean power consumed by the water heater population when the population is subject to the mean field control. It is interesting to consider this figure alongside Figure 4.3. It is worth noting the close

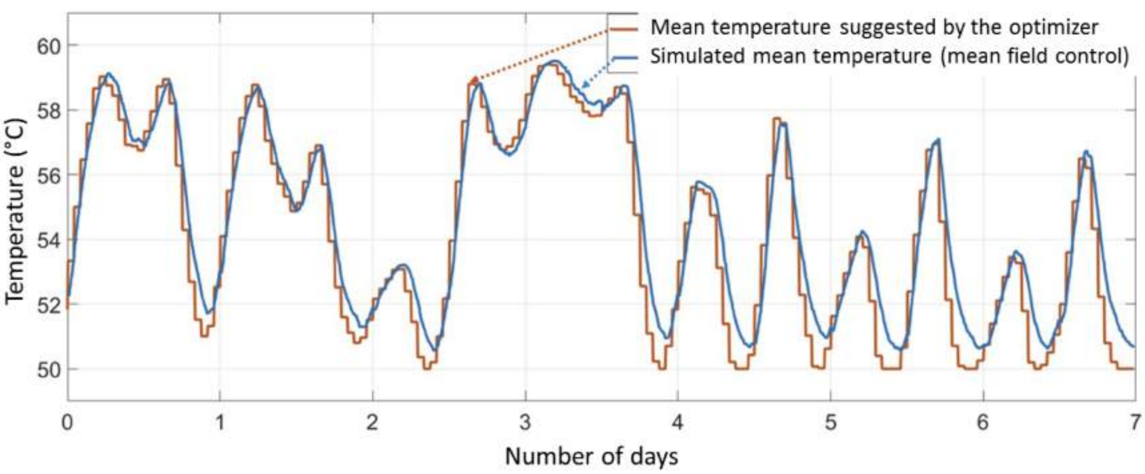

FIGURE 4.3 – Changes in the mean temperature of the water heater population with mean field control over 400 houses.

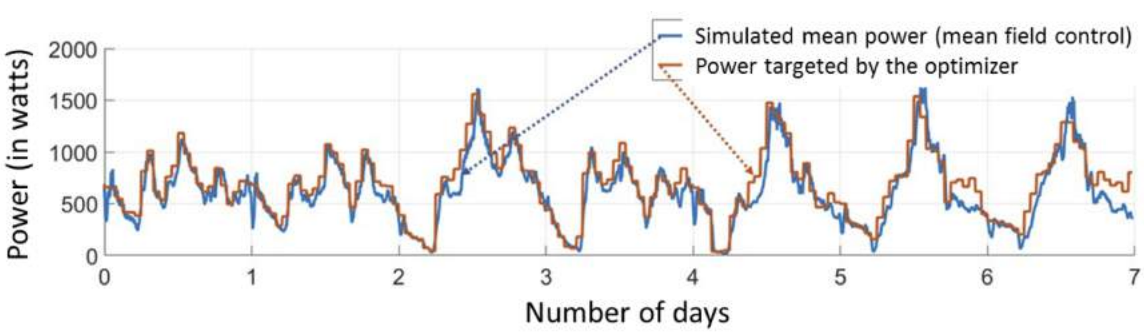

FIGURE 4.4 – Mean power consumed by the water heater population with mean field control over 400 houses.

proximity of the optimal power (red) and simulated power (blue) curves, which shows that the local control does its job very well. It is important to remember that this controller receives no power set point, but only the trajectory of the mean temperature of the water heater population.

Notice, though, that a few instances of uncoupling, for example at the end of the 5^{th} day, are related to the gap between the desired mean temperatures and the mean temperatures achieved by the control (effect of the limits, as discussed above). Some solutions are currently being explored to improve the control and/or optimization when a significant proportion of the population reaches the temperature limits. The next section provides more detail about the temperature distribution within the water heater population.

Mean temperature distribution : equity between water heaters

As explained in section 3.2.3, one intended purpose of the mean field control is to customize a common control for every water heater based on local constraints and need, while maintaining some concept of "equity" in the controlled fleet. In fact, the mean field control seeks to minimize the disruption it might cause to the water heaters statistical mean temperature distribution by avoiding, for example, bringing the entire population to exactly the same temperature. In other words, the goal is to maintain significant variance around the target mean temperature at all times in order to avoid the effects of water heater synchronization sometimes observed in a more conventional control (e.g., by load shedding). Figure 4.5 depicts the population's thermal distribution during the seven days of mean field control. Note that the population's variance remains significant, i.e., the mean temperature difference between the hottest and coldest water heaters remains large at all times. The curve for the coldest 10% (in dark blue) also never falls below the 48°C bar for the mean field algorithm used here, and the number of dissatisfied customers is basically nil over the eight days of the simulation, i.e. no water heater supplied hot water at a temperature below 45°C.

Figure 4.5 also highlights the competition between two opposing "forces" :

— the first tends to reduce variance (by bringing the temperature curves closer together) : it represents the fleet heating phases. When the fleet temperature increases, the control that is given tightens distribution around the target mean temperature,

— the second force tends to increase variance : it represents the fleet's cooling phases (unloading). As previously discussed, the temperature of the water heaters is achieved mainly through the drawing of hot water, which is a random process. This process therefore naturally helps to increase the population temperature variance.

The mean field control is capable of managing competition between these opposing forces with the balance needed to maintain the desired variance in the population and, therefore, to maintain constant flexibility in terms of energy storage for the electrical grid.

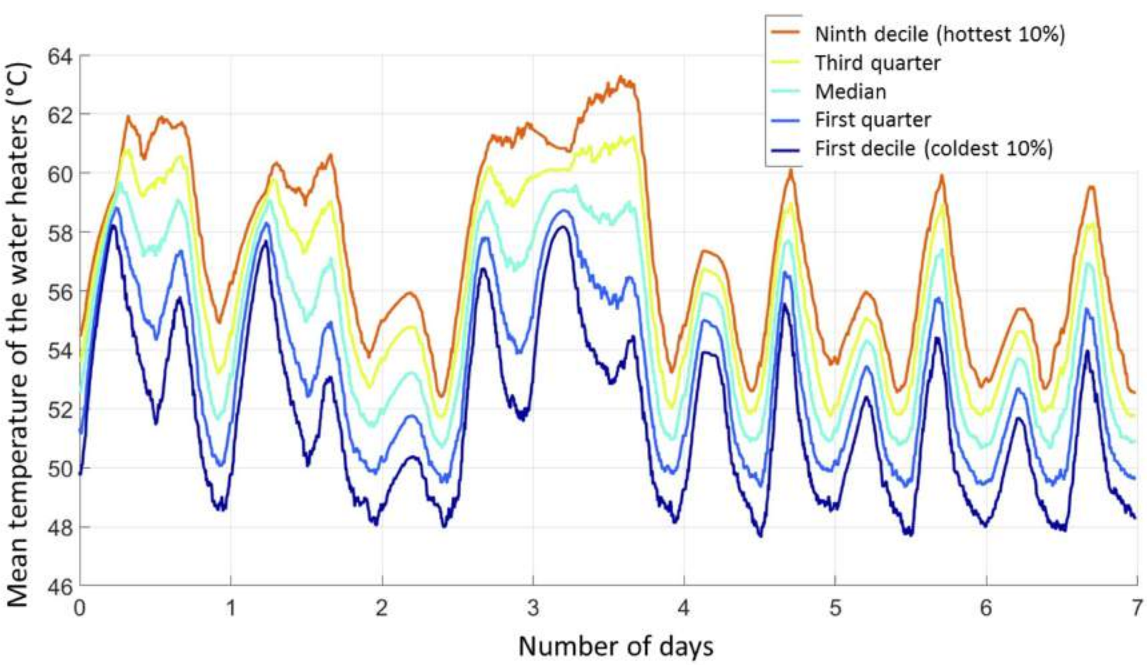

FIGURE 4.5 – Mean temperature distribution of a fleet of 400 water heaters with a deterministic mean field control.

Finally, Figure 4.6 depicts the individual mean temperature trajectories of three water heaters selected at random from a population of 400 water heaters with a mean field control. The first and last deciles in Figure 4.3 are also plotted. This figure clearly illustrates the equitable principle underlying the mean field control. In fact, looking more specifically at water heater 3 (in green), it appears that at time $t = 2.7$ days, it is among the coldest water heaters, then at time $t = 4.5$ days, it is among the hottest. The same observation can be made for the other water heaters.

Note that despite a slowly changing thermal mean and distribution, each water heater's individual position within this distribution changes on a very regular basis : it is not always the same individuals in the population that contribute to the overall energy increase or decrease goal. The water heaters behaviour within their group is analogous to several behaviours found in nature, especially in a school of fish or flock of starlings in flight (see Figure 4.7). Indeed, each starling's position within the flock varies constantly, randomly and quickly, but the trajectory and size of the flock themselves remain stable. The speed and movements of the flock as a whole are consistent, as though the flock were a single entity with a will of its own.

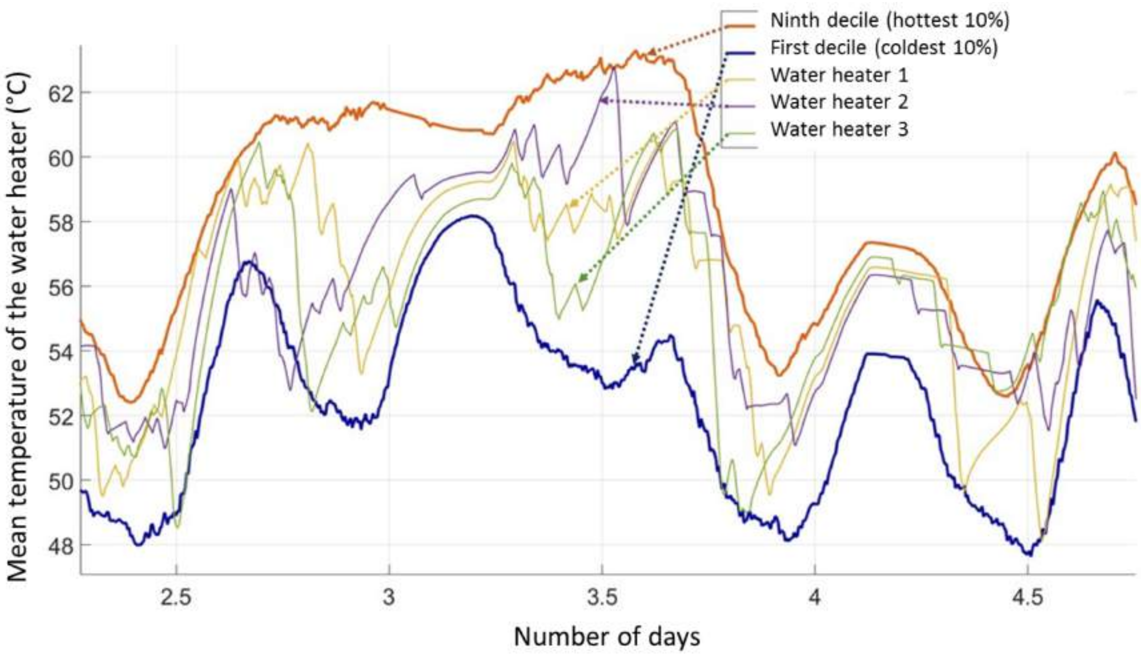

FIGURE 4.6 – Individual mean temperature trajectories of three water heaters selected at random from a population of 400 water heaters with a mean field control.

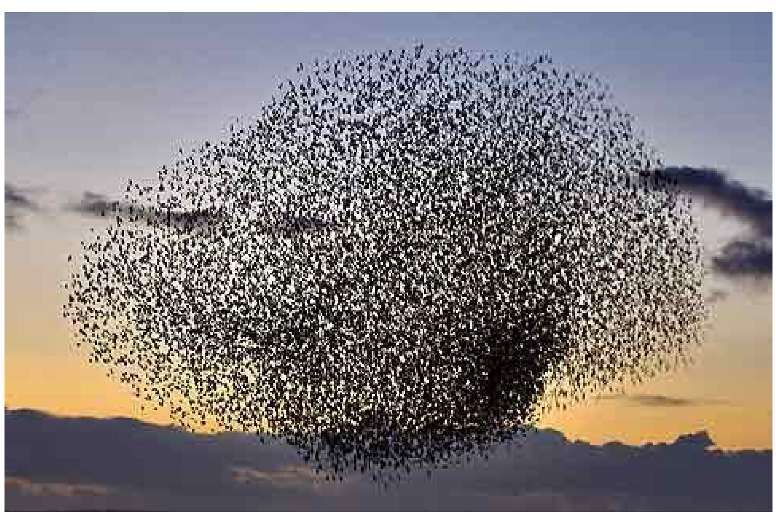

FIGURE 4.7 – Flock of starlings *(source : unknown).*

TABLE 4.1 – Summary of the results associated with the telecommunications network's performance for a mean field simulation of 400 houses in a grid with a total of 3,300 meters, with hourly updating of the control from the mainframe.

Mean ascending delay	19 s
Maximum ascending delay	38 s
Mean descending delay	303 s
Maximum descending delay	414 s
Mean network node occupancy	0.13 %
Maximum network node occupancy	22 %

Impact of the mean field control on the telecommunications network

As previously mentioned (see section 3.2.4), messages relating to the control of the water heaters and to the occasional transmission of their mean temperature (or of their energy state) to the mainframe are carried over a telecommunications network. This telecommunications network has all of the features of the one developed and used by Hydro-Québec for its smart meters : it is a "mesh" network with routers and collectors, with throughputs of 9.6 kbps between meters and routers, and of 19.2 kbps between routers and collectors.

The telecommunications network plays a dual role. It :

1. makes the simulations produced even more realistic since it reproduces the period between the time messages are sent and received,
2. quantifies the rate at which the meters are used to perform the mean field control.

In this manner, statistics can be produced at the end of the simulations. In the case of the above simulation, for a mean field control with 400 houses that have water heaters controlled by mean field out of a total of 3,300 house [3], the results obtained are summarized in Table 4.1.

The "ascending delay" represents the time it takes for a message sent by the water heater controller to reach the mainframe. The "descending delay" represents the time it takes for broadcasts sent from the mainframe to reach all of the controllers. The broadcast represents the transmission of the common control (24 target mean temperatures + mean temperature of the population) from the mainframe to all of the water heaters.

Broadcast duration is particularly slow compared with ascending delay, with a difference factor of 10. The reason for this is the lack of a broadcast protocol natively integrated into Hydro-Québec meters. As a result, broadcasts can be made only through individual point-to-point transmissions, which are therefore particularly costly in terms of bandwidth. It would be possible to reduce this duration significantly by updating the meters' communications protocol (if the regulations allow, of course).

The main conclusion that can be drawn from this exercise is that the occupancy of network nodes is particularly low, which aligns with the goal of the mean field control, i.e., to achieve a common objective through a very decentralized control requiring very little communication between actors. The maximum observed occupancy of 22% reflects a strategic node in the mesh that could potentially be reduced through a message "re-routing" process.

3. 400 of the 3,300 simulated meters generate and receive messages related to the mean field control. The other meters participate in the telecommunications network by repeating the messages, but do not send any messages related to the water heater in their associated house. These water heaters are not actually controlled.

4.3 Case 2 : load recovery following a blackout

A *black-out* is widespread power failure affecting a neighbourhood, city or entire region for a period ranging from a few minutes to several days. The reasons for blackouts vary but often have to do with energy demand exceeding available transmission or generating capacity, or with a major infrastructure failure, for example, following a natural disaster [Pourbeik et al. 2006]. Blackouts in interconnected electrical grids also have a domino effect, and a blackout in one area may spread to neighbouring areas [Pourbeik et al., 2006].

The analysis in this case study is confined to observing the load recovery speed of a water heater population following a power outage lasting about one hour and 20 minutes and involving the same distribution line as in case study 1. The point at which power is restored to the grid following an outage is particularly critical, because much of the water heater population is waiting for heating, which creates a temporary consumption peak until the population resumes its usual diversity of energy states. In this study, it is assumed that only the water heaters have synchronized load recovery dynamics, although other components of the non-controllable load (electric heating, refrigerators, etc.) would actually magnify the severity of the peak observed during load recovery.

Load recovery in a population of water heaters fitted with conventional thermostatic controls is affected by a severe payback phenomenon *payback* (green curve in Figure 4.8). This figure also includes the non-controllable consumption curve (orange), the optimal curve provided by the deterministic optimizer (in purple), and the total consumption curve when the water heaters are fitted with a mean field controller. The first conclusion is that the mean field control is capable of drastically reducing the fleet's load recovery peak. There is a simple explanation for this : while thermostatic controllers try to supply water heaters with the power lost during the outage as soon as the power comes back on (using the full rated power of the water heaters), the mean field controller offsets this loss more slowly, all the while maintaining the diversity of power demands within the population. Indeed, the controller does not attempt to make up for the energy loss during blackout immediately, but rather over a longer period. This is reflected in slightly higher power demand than the water heaters undisturbed optimal demand trajectory, but the demand quickly *returns* to the optimum level after about one hour.

It is important to note that no special instruction is sent to the mean field controller during or after the outage : it naturally manages load recovery in a nearly ideal way. In other words, the mean field controller is "aware" of the heating loss and therefore of the resulting impact on the mean temperature of the entire population. When power is restored to the grid, the mean field controllers simply continue to operate based on the last control broadcast (at midnight the same day). This ability to maintain the initial course of action is critical because power outages also generally disrupt the telecommunications network.

Note also that in the event of an outage lasting a few hours, the lack of power to the water heaters results in

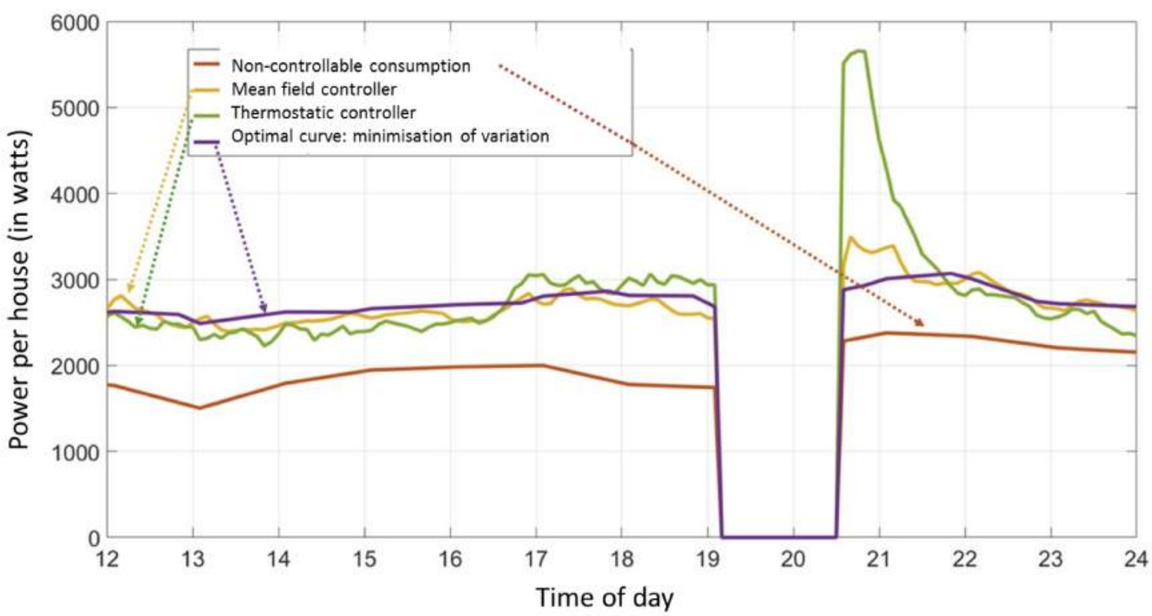

Figure 4.8 – Profiles of the mean power consumed per house in connection with a power failure between 7 : 10 p.m. and 8 : 30 p.m. for a fleet of 400 water heaters fitted with i) thermostatic (green) and ii) mean field (yellow) controllers.

a drop of a few degrees in the mean temperature, which in most cases is imperceptible to users. For example, in the simulation for Figure 4.8, the customer dissatisfaction rate is nil in both cases (thermostatic and mean field). However, the power peak is radically different (major reduction in load recovery from 2,500 W to 500 W per house, i.e. 80%). The impact on electric companies is a quicker and safer return to service in areas affected by a blackout

One bias of this study is that the parameters for the random water drawing process are considered to remain the same during power failures, i.e., consumers' hot water consumption habits do not change in the event of a power failure. This is certainly not the case in practice, but that jeopardize the ability of the mean field control to manage load recovery seamlessly.

4.4 Case 3 : wind power balancing and stochastic optimization

4.4.1 General considerations

This final case study is the ultimate objective of the smartDESC project : to provide distributed energy storage capacity in the electrical grid to facilitate the penetration of renewable energies, in this case wind energy. The fluctuations of this production are strongly correlated with that of weather. We are therefore bringing the inherent uncertainty of weather forecasts to bear on the storage energy planning problem : this is the role of the stochastic optimizer described in section 3.2.2.

4.4.2 Stochastic projection generation

In addition to the mean temperature of the population (estimated by sampling), the stochastic optimizer requires as inputs N non-controllable consumption projections and M wind generation scenarios. The scenarios are generated as follows :

1. WPred, a company specializing in wind and solar generation projections, provides 22 equiprobable weather scenarios for the next 24 hours. These projections are made using public data provided by Environment Canada. The scenarios include the following information : wind speed and direction from 10 to 150 metres above ground level, as well as the air temperature and atmospheric pressure. [4],
2. This information makes it possible to estimate the wind generation of a wind farm sized to produce 10% of the electricity consumed by the study population. The estimate is made using the wind turbine power curves provided by the manufacturers,
3. These scenarios also support a non-controllable load projection module provided by Artelys. This module can use a weather forecast and a load history for a given power line to generate a non-controllable consumption projection that is a function of temperature, wind speed, time and day of the week. The 22 weather scenarios therefore provide 22 non-controllable consumption scenarios.

Figure 4.9 illustrates the above process. Figure 4.10 and Figure 4.11 show the different wind generation and non-controllable consumption scenarios for an autumn day. The probability of each of these scenarios occurring is the same in this case at $1/22 = 4.545\%$.

The following assertions can be made based on these curves :

— Based on the macrodata used, the discrepancies between the different wind generation scenarios become significant beyond a six-hour horizon,
— The wind scenarios develop very differently and may intersect,
— The electricity consumption projection is much less diverse : the rate of consumption is the same for all of the scenarios,
— The electricity consumption scenarios generally do not intersect. Over a period of about 10 hours, they could even be defined by a single scenario. A more precise description of the domestic loads would probably make it possible to achieve greater variability between the different scenarios, but this was not explored here.

These figures show that the complexity and variability of the problem has to do mainly with the wind generation projections. The consumption scenarios are considered to be basically identical (as a first approximation). The wind scenarios are all very different and their variations are not correlated. It is virtually impossible to identify a general trend among the wind generation profiles.

Throughout the following simulations, it is assumed that only one of these scenarios is fully realized. In other words, it is assumed that the above projections are sufficiently accurate for one of the 22 proposed scenarios to describe reality. The stochastic optimizer obviously has no prior information about the real scenario.

4. The air's temperature and pressure directly affect its density, and the output of wind turbines depends on the density of the fluid that sets them rotating.

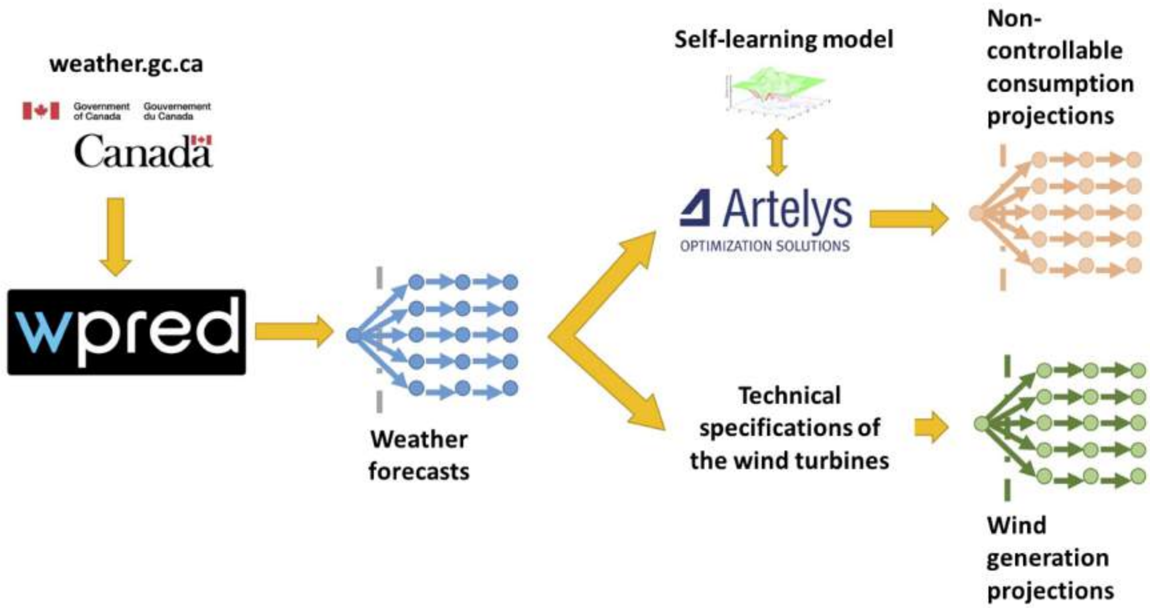

FIGURE 4.9 – Generation chain for consumption and wind energy production projections.

The scenarios are generated every hour for a period of 24 hours, enabling optimization to be implemented on a *rolling horizon*. These scenarios are transmitted to the mainframe, which generates the resulting optimal control.

4.4.3 Stochastic optimization

Using all of the scenarios as an input enables the stochastic optimizer to generate a unique control (24 target mean temperatures for the next 24 hours), which factors in the diversity of future developments. This control is updated every hour to take into account the latest developments in wind generation and non-controllable consumption.

The mainframe needs the following information to carry out its stochastic optimization :

- The non-controllable consumption scenarios (twenty-two 24-hour scenarios),
- The wind generation scenarios (twenty-two 24-hour scenarios),
- The actual non-controllable consumption (a single scenario chosen from among the 22 possible scenarios),
- The actual wind generation (a single scenario chosen from among the 22 possible scenarios),
- The mean temperature of the fleet of water heaters (obtained using the mean temperatures transmitted randomly by the local controllers),
- The parameters of the water drawing processes for the next 24 hours,
- The characteristics and operating conditions of the controlled water heaters (power, capacity, water inlet temperature, etc.).

All of this information enables the mainframe to generate the mean temperature trajectory associated with these scenarios. This information is refreshed and transmitted to the controllers every hour.

In this study, non-controllable consumption and wind generation are two similar energy “flows”. Indeed, they are both non-controllable and constitute data exogenous to the simulation. These two “flows” are consequently aggregated in the results set out below into a single flow called **non-controllable load**. It is therefore important to understand that this “non-controllable load” is equivalent to the houses’ **non-controllable consumption less wind generation** (renewable power generation mismatch with the non-controllable load).

Figure 4.12 illustrates the process by which the non-controllable load is obtained from wind generation and non-controllable consumption. The wind generation and non-controllable load curves are specific scenarios chosen randomly from among the 22 scenarios generated and, for the purposes of the simulation, are considered to be “the real scenarios”, i.e. “those that actually occur”. The stochastic optimizer is not aware of this choice and receives the 22 scenarios with related probabilities of occurrence (which are equal for the purposes of this study).

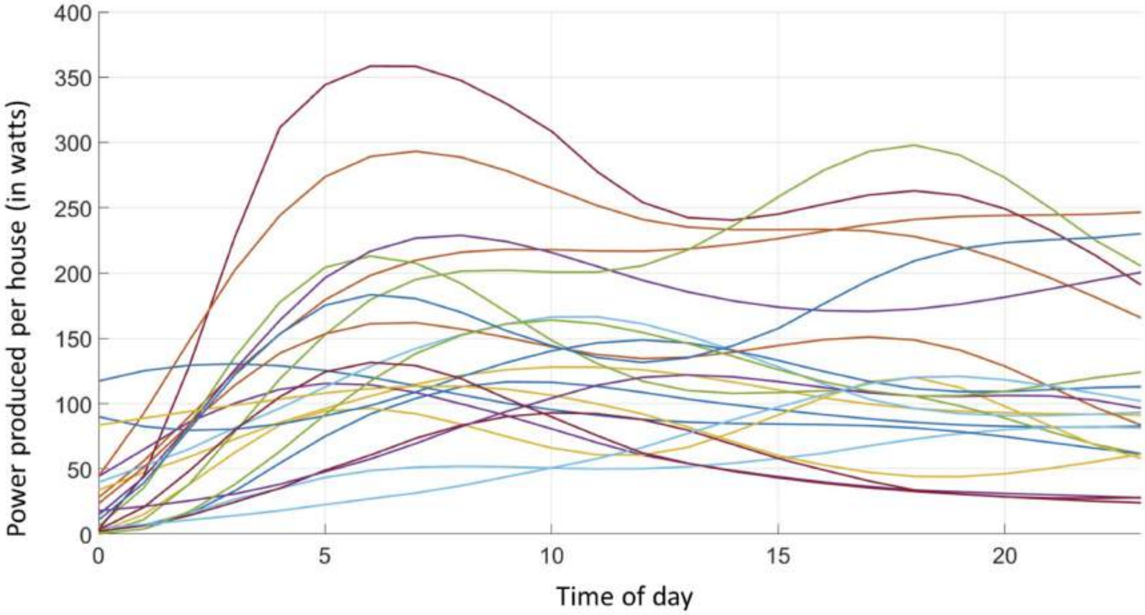

FIGURE 4.10 – Wind generation projection scenarios (22) for the day of September 2, 2016, in suburban Montréal.

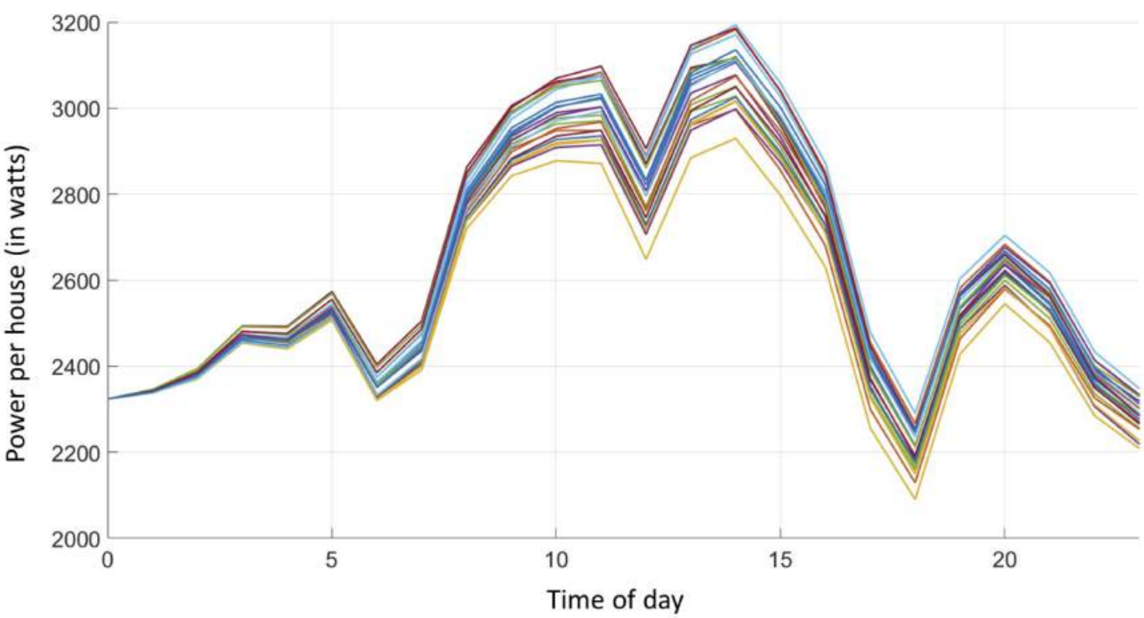

FIGURE 4.11 – Consumption projection scenarios (22) for the day of September 2, 2016, in suburban Montréal.

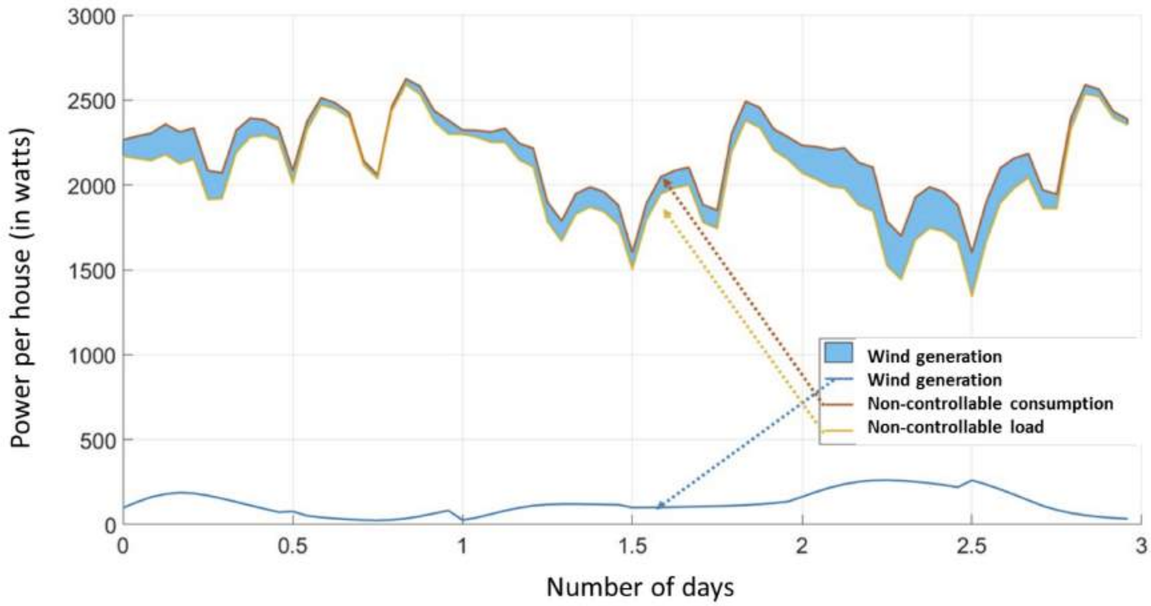

FIGURE 4.12 – Non-controllable consumption and wind generation scenario used for case study 3.

Result of the stochastic optimization for 10% installed wind power

The simulation takes place over three days in September 2015 for a Montreal suburb.

Figure 4.13 summarizes the results in three curves :

— the non-controllable load curve (orange) : described above, it is equivalent to non-controllable consumption less wind generation,

— the simulated total load (seen by the electricity distribution company) : it is equivalent to the mean consumption of the houses fitted with mean field controllers less wind generation,

— the optimal load : it is the result of the optimization performed by the stochastic optimizer using the load variation limiting method (levelling of the power seen by the electrical grid). It also includes wind generation.

All of these curves include wind generation in their trace, meaning that actual wind generation (from the chosen scenario) is subtracted from the houses' consumption. The goal here is to show that it is possible to absorb a wind farm's production locally, i.e. to achieve *wind power balancing*. The weather data that make it possible to produce the generation and consumption scenarios are the same : either a substation or small city is used.

Figure 4.13 clearly shows the way mean field control makes it possible to absorb produced wind energy locally. The electrical grid is treated as a system separate from the wind farm, and the power consumption it sees corresponds to the yellow curve in Figure 4.13. The variations and unpredictable nature of wind energy are therefore managed locally and, as a result, are nearly invisible to the rest of the grid. In this simulation in particular, the wind energy produced is equivalent to 5% of the total energy consumed, and the wind power produced may total 10% of the houses total electricity consumption.

It is possible to compare the load of the fleet fitted with mean field controllers to the same fleet (with the same wind generation) fitted with thermostatic controllers. Figure 4.14 shows both the total thermostatic load (in **green** and the total mean field load (in **yellow**). It turns out that mean field control results in better levelling of the total load (including wind generation) than thermostatic control, as already demonstrated in our previous case studies. Peaks are reduced and consumption troughs are filled. Although the total mean field load curve does not follow the optimal trajectory (in **purple**) perfectly because of possible discrepancies between target energy levels and the actual levels realized, the results are still very good compared to conventional thermostatic control.

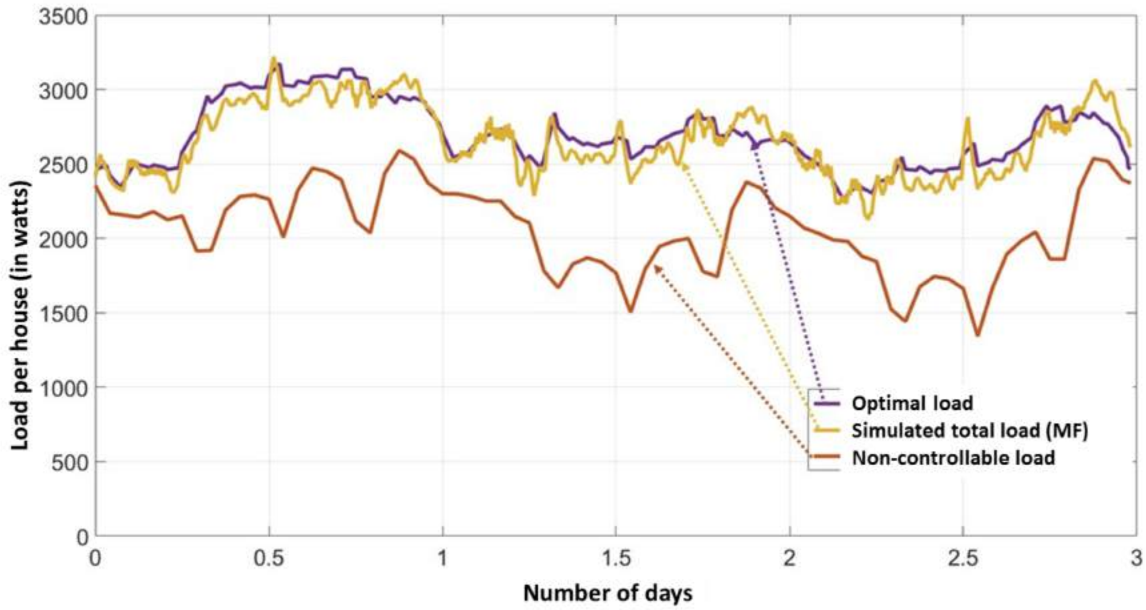

FIGURE 4.13 – Results of a three-day stochastic simulation for 600 houses with a penetration rate of 10% wind power compared with the grid's peak power.

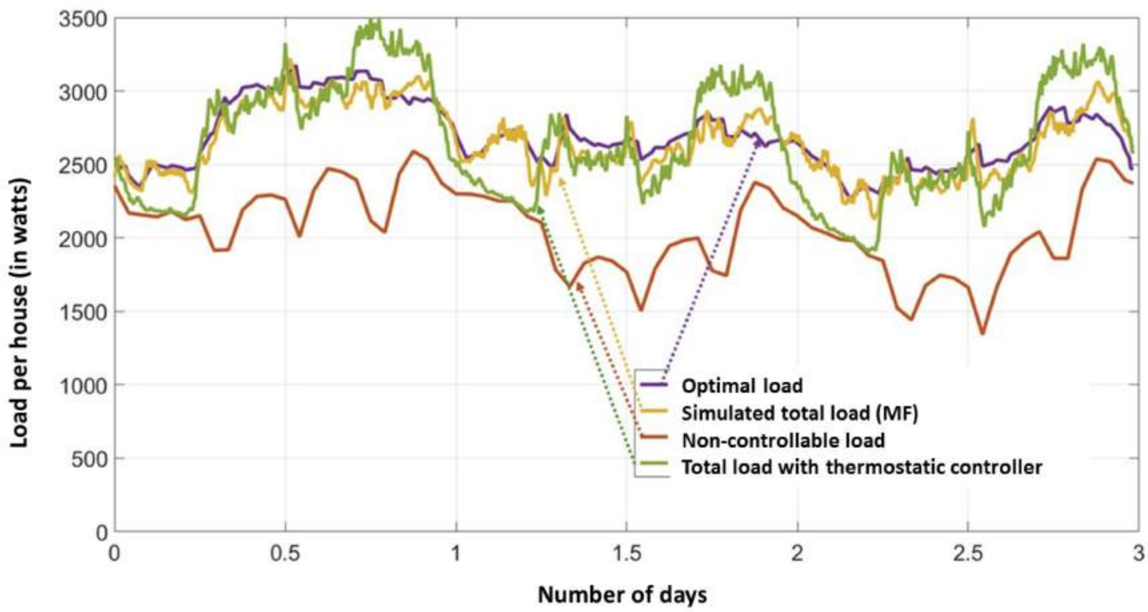

FIGURE 4.14 – Comparison between mean field and thermostatic control for a population of 600 houses with a penetration rate of 10% wind power compared with the grid's peak power.

Result of the stochastic optimization for 20% installed wind power

As a final case, in order to further push the limits of the mean field control, the wind farm's output was doubled to 20% of the electrical grid's peak power. As a result, a 500 W peak in wind generation (per house) is observed on the 3rd day in Figure 4.15. Comparing this with the maximum non-controllable load ($\approx$ 2500 W) yields a ratio of 20%.

Figure 4.16 depicts the results of the stochastic optimization. It includes :

— In **orange**, the non-controllable load, which is equivalent to the houses' non-controllable consumption less wind generation,

— In **purple**, the consumption resulting from the stochastic optimization. This curve also corresponds to the total load if the water heaters follow exactly the mean temperature trajectory determined by the stochastic optimizer,

— In **yellow**, the simulated total load for a fleet of water heaters fitted with mean field controllers. There are still deviations from the optimal load due to “crashing” of the water heaters’ thermal distribution at the temperature limits (50°C and 60°C),

— In **green**, the total load for a fleet of water heaters fitted with thermostatic controllers and, therefore, without consumption optimization.

All of these curves include wind generation in the total load, i.e., wind generation was subtracted from total consumption to obtain the total load curves seen by the electrical grid.

Once again, the ability of mean field control to help absorb wind fluctuations is clearly demonstrated despite the high proportion in this case (20% of the grid’s peak power). In fact, the variations in the **yellow** curve are quite small compared with the optimal curve, whereas there are significant power spikes in the **green** curve. For the rest of the electrical grid (i.e. all except the population of controlled houses and the wind farm), the load of the body of houses and wind turbines varies only slightly over time.

4.5 Tests on the physical platform

As the very last stage of the project, tests were performed on the physical platform described in section 3.3 in order to i) calibrate the thermal model used in the simulator, and ii) demonstrate that the local controller required by the mean field control could be implemented in a practical way using existing commercial equipment.

Initially, simple heating and cooling cycles were applied to the water heater to test the thermal model. We were able to calibrate all of the model’s constants after a few iterations. Then, the thermostatic mode was successfully tested and, finally, the local controller was tested for a month, as described below.

The methodology used to test the local controller in connection with a mean field control generated by simulation was as follows : a mean field simulation was first carried out on a large population of water heaters using the overall simulator, and the state of all of the simulated water heaters was saved as a function of time (temperature profile, water drawing, heating state). Next, the temperature of the physical water heater was initialized to the same state as a first simulated water heater (the initial state was always taken overnight so that the temperature would be consistent throughout the volume of the water heater as a result of little water being drawn during this period). Then, the same water drawing conditions as in the simulated case were applied to the platform for 24 hours. This way, it was possible to emulate as many simulated days as there are water heaters in the simulation, each with a different hot water extraction history. This process is obviously very time consuming (24 hours of real time per water heater, plus about one hour per initialization), so it is not possible to physically emulate hundreds of scenarios unless the platform is left running for several months.

In this case, it had been possible to emulate 30 water-heater days on the platform at the time of writing this report. Although there seems to be every indication that the physical implementation of the local controller works well, a larger number of results would be needed in order to conclude that the group mean temperature follows closely the trajectory dictated by the mainframe, in keeping with the law of large numbers on which all of mean field control is based. This report contains no figure relating to the results obtained on the platform, because the results have not been analyzed in detail to date. However, one thing is certain : the mean field controller has not yet been optimized for an industrial roll-out, and the work needed to get there will have to be done as part of a subsequent industrialization phase.

Finally, it is important to stress the fact that the entire physical infrastructure required to achieve mean field control over a large population of water heaters cannot actually be tested using the current platform. In fact, although the simulator current structure makes it possible to reproduce the entire data aggregation aspect quite realistically, in practice, thousands of sensors would return information via the local controllers and an as-yet-undefined telecommunications network.

The effect of transmission failures or the failure of any other equipment in such a network is yet to be quantified, but more importantly the effect of likely discrepancies in diversified hot water drawing statistics in a real network, when compared to that estimated and used for mean field control calculations, remains a main cause for potential concern. Indeed, the law of large numbers works (as clearly shown in simulation), and physical implementation was demonstrated via the platform, but the correct statistical parameters are still required in order to complete the calculations. The issue of gradual on-line learning of the statistical water drawing models, as well as the classification of consumers into homogeneous classes for the purpose of well-targeted control, should therefore become important research topics in the future.

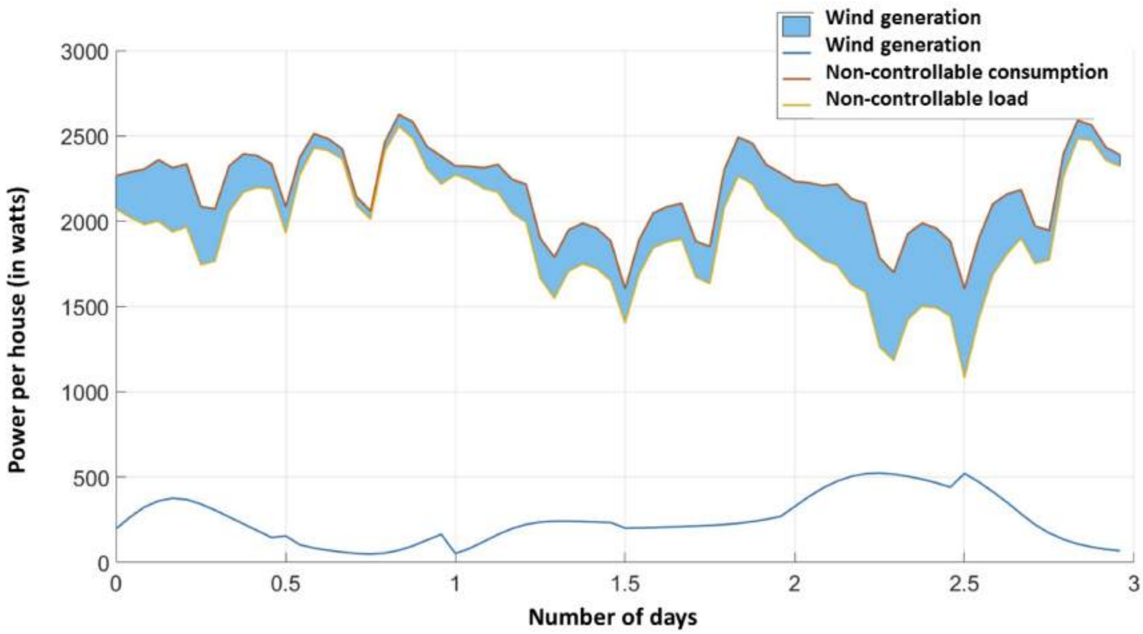

FIGURE 4.15 – Non-controllable consumption and wind generation scenario used for a case study with 20% installed wind power compared with the electrical grid's peak power.

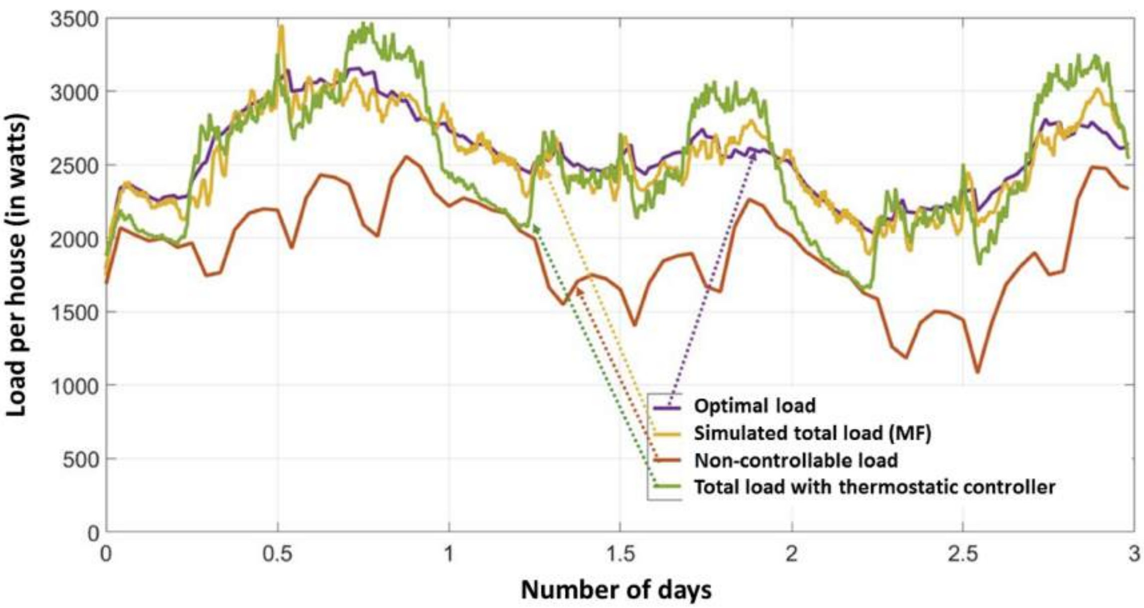

FIGURE 4.16 – Comparison between mean field and thermostatic control for a population of 600 houses with wind generation totalling 20% of the electrical grid's peak power.

Chapitre 5

Conclusion

5.1 Project benefits and results

The smartDESC project (RENE-034), carried out as part of the Canadian federal government's ecoENERGY Innovation Initiative (ecoEII), has demonstrated the effectiveness of mean field control theory for managing a large number of energy storage elements distributed in an electrical grid, in order to smooth electricity demand over time or, even more importantly in terms of program objectives, to achieve wind power balancing in an electrical grid. In the latter case, this opens the door to a potential increase in wind energy penetration in Canadian electrical grids. This report has shown that 20% wind penetration could be fully absorbed through the proper management of thermal energy storage in household water heaters fitted with mean field controllers. This percentage is in addition to any other proportion of wind power already installed in the grid, which potentially pushes the penetration rate envelope beyond conventional thresholds, without having to compensate the variability due to this extra wind power through additional power plants.

This project also showed that the mean field control helps mitigate the peak problem created by load recovery in electrical grids, i.e. a power peak that occurs after the electricity supply is restored to a grid following a blackout. The mean field control also requires very little bandwidth in terms of telecommunications. As a result, a mean field control applied to an electrical grid could make do with communication via a network of very low-bandwidth smart meters and would take up less than 1% of their bandwidth.

One interesting spin-off of this project was the development of a multi-agent simulator that makes it possible to couple a set of disparate processes in a synchronized time simulation. This simulator software structure (based on the JADE environment) is available upon request, but no warranty or technical support can be offered for its operation since the project team has now disbanded.

Finally, a physical platform with a real water heater controlled by a mean field controller was built. The platform made it possible to confirm that the thermal model of the water heater used in the simulator could be reliably tuned. It was also possible to confirm that the local controller required for the mean field control was physically feasible without having to use sophisticated components : a simple microcontroller, some Ethernet or WiFi telecommunications equipment and basic instruments are sufficient.

5.2 Next steps for R&D in the field

Generally speaking, the proof of concept of the mean field control was successfully completed by simulation.

A physical platform was also built toward the end of the project to confirm empirically that the local controller required by the mean field control can be built using commercially available technologies, and that the local controller behaves as the simulation suggests. At this stage, the concept could be implemented almost as-is in a pilot project, provided that the following enabling elements are gathered :

1. a manufacturer to produce the mean field controllers,
2. an electricity distribution company that would be prepared to install the necessary equipment in its grid at both the centralized and local level (on customers' premises),
3. depending on the circumstances, an aggregation service provider to manage the data gathering and flow between the customer and the mainframe,
4. a proponent to oversee project management and coordination.

Any improvements made through research carried out alongside a potential pilot project could be incorporated into the pilot project through a firmware update to the local controllers, and also in terms of continuous improvement of the optimization algorithms in the mainframe.

The presence of the initial promoter (Polytechnique Montreal) would still be required for technology transfer during the development phase prior to reaching a first prototype. In addition, a number of still unresolved research issues, if addressed, would help improve the performance of the mean field controllers, in particular when applied to electric water heaters. In particular, one could mention :

1. the development of a learning algorithm to identify the stochastic processes of hot water extraction to improve the accuracy and performance of local controllers,
2. the lessening the customer's comfort constraints on the accuracy of overall mean target temperature tracking in electric water heaters (boundary effects),
3. the generalization of mean field control approaches to other types of loads and energy storage capable elements, including electric heating and air conditioning loads [Moffet et al., 2012]), electric vehicles [Tuffner and Kintner-Meyer, 2011], etc.,
4. the generalization of the optimization algorithms to account for multiple classes of controlled loads,
5. etc.

Research improvements that could be realized in the process of a potential future prototyping process could be integrated in the form of updates to the *firmware* of local controllers, and also in terms of continuous improvement to the macroscopic optimization algorithms in the central computer.

Bibliographie

Larry Hughes. Meeting residential space heating demand with wind-generated electricity. *Renewable Energy*, 35(8) : 1765–1772, 2010.

Coopérative régionale d'électricité de St-Jean-Baptiste-de-Rouville, 2016. URL : `http://www.coopsjb.com/`. [Online ; accessed March 8^{th} 2016].

Jean-Michel Lasry and Pierre-Louis Lions. Mean field games. *Japanese Journal of Mathematics*, 2(1) : 229–260, 2007.

Jérôme Solis. Développement d'un estimateur d'état énergétique d'un chauffe-eau pour un contrôle par champ moyen. *Master's thesis*, Polytechnique Montréal, 2015.

Minyi Huang, Roland P Malhamé, and Peter E Caines. 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.

Pouyan Pourbeik, Prabha S Kundur, and Carson W Taylor. The anatomy of a power grid blackout. *IEEE Power and Energy Magazine*, 4(5) : 22–29, 2006.

Marc-André Moffet, Frédéric Sirois, and David Beauvais. Case studies : Balancing wind generation using electric thermal storage and electric water heaters. *CanmetENERGY*, Report no. 2012 067 RP CAS 411 SGZONE : 32 p., 2012.

Frank K Tuffner and Michael CW Kintner-Meyer. Using electric vehicles to meet balancing requirements associated with wind power. *Pacific Northwest National Laboratory (PNNL)*, Report no. PNNL-20501 : 48 p., 2011.

Roland Malhamé and Chee-Yee Chong. Electric load model synthesis by diffusion approximation of a high-order hybrid-state stochastic system. *IEEE Transactions on Automatic Control*, 30(9) : 854–860, 1985.

Arman C. Kizilkale, Frédéric Sirois *et al.* Tammam, M. Anjos, M. Bernier, M. Gendreau, B. Sanso, and R. P. Malhamé. smartDESC : smart distributed energy storage controller. *to be submitted to IEEE Power and Energy Magazine*, 2017. URL to preliminary version : `http://www.cim.mcgill.ca/%7Earman/preprints/2016PEM.pdf`.

Arman C. Kizilkale and Roland P. Malhamé. *Collective target tracking mean field control for Markovian jump-driven models of electric water heating loads*, chapter 20 in : *Control of Complex Systems : Theory and Applications*, edited by K. G. Vamvoudakis and S. Jagannathan, Butterworth-Heinemann (Elsevier), p. 559–584, 2016.