Les Cahiers du GERAD ISSN: 0711–2440

Smart distributed energy storage controller (smartDESC)

F. Malandra, A. C. Kizilkale,
F. Sirois, B. Sansò, M. F. Anjos,
M. Bernier, M. Gendreau, R P. Malhamé

G–2020–15

February 2020

Citation suggérée : F. Malandra, A. C. Kizilkale, F. Sirois, B. Sansò, M. F. Anjos, M. Bernier, M. Gendreau, R P. Malhamé (Février 2020). Smart distributed energy storage controller (smartDESC), Rapport technique, Les Cahiers du GERAD G–2020–15, GERAD, HEC Montréal, Canada.

Avant de citer ce rapport technique, veuillez visiter notre site Web (https://www.gerad.ca/fr/papers/G-2020-15) afin de mettre à jour vos données de référence, s'il a été publié dans une revue scientifique.

Suggested citation: F. Malandra, A. C. Kizilkale, F. Sirois, B. Sansò, M. F. Anjos, M. Bernier, M. Gendreau, R P. Malhamé (February 2020). Smart distributed energy storage controller (smartDESC), Technical report, Les Cahiers du GERAD G–2020–15, GERAD, HEC Montréal, Canada.

Before citing this technical report, please visit our website (https://www.gerad.ca/en/papers/G-2020-15) to update your reference data, if it has been published in a scientific journal.

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

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

Dépôt légal – Bibliothèque et Archives nationales du Québec, 2020
– Bibliothèque et Archives Canada, 2020

Legal deposit – Bibliothèque et Archives nationales du Québec, 2020
– Library and Archives Canada, 2020

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

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

Smart distributed energy storage controller (smartDESC)

Filippo Malandra [b]

Arman C. Kizilkale [c]

Frédéric Sirois [d]

Brunilde Sansó [a,d]

Miguel F. Anjos [a,e]

Michel Bernier [f]

Michel Gendreau [a,g]

Roland P. Malhamé [a,d]

[a] *GERAD, Montréal (Québec), Canada*

[b] *University at Buffalo, Department of Electrical Engineering, Buffalo, NY, U.S.A.*

[c] *Automat.AI, Montréal (Québec), Canada*

[d] *Department of Department of Electrical Engineering, Polytechnique Montréal (Québec) Canada*

[e] *University of Edinburgh, School of Mathematics, Edinburgh, Scotland*

[f] *Department of Department of Mechanical Engineering, Polytechnique Montréal (Québec) Canada*

[g] *Department of Department of Mathematical and Industrial Engineering, Polytechnique Montréal (Québec) Canada*

February 2020
Les Cahiers du GERAD
G–2020–15

Abstract: While one can exploit the storage properties and thus the deferability or anticipation potential of many classes of power system loads (such as thermal loads) as an increasingly needed tool to mitigate renewable sources variability, the challenge to do so in an optimal and coherent manner is immense. This is in view of the sheer number and dynamic diversity of the loads that can be involved in any large-scale application of such an approach. The smartDESC concept is a control architecture that was developed for this purpose. It builds on the more pervasive communication means currently available (such as Advanced Metering Infrastructures), as well as the mathematical tools of (i) aggregate load modeling, (ii) renewable energy forecasting, (iii) optimization theory, deterministic or stochastic, and (iv) the recent developments in control of large-scale systems based on game theory and called mean-field (MF) control theory, which allow producing a scalable yet optimal approach to the decentralized control of large pools of loads. This paper presents the building blocks of the smartDESC architecture, together with an associated simulator and simulation results.

Acknowledgments: This project has been primarily funded by the Canadian federal government's ecoENERGY Innovation Initiative (ecoEII), between 2012 and 2016 (project RENE–034). The authors gratefully acknowledge this financial contribution, as well as the in-kind contributions of many other partners, namely: Laboratoire sur les Technologies de l'Énergie (LTE) of IREQ (Institut de Recherche d'Hydro-Québec), WPred Inc., Artelys, SG2B, and Coopérative St-Jean-Baptiste (a small electricity distributor in Montérégie, QC, Canada). The authors also address a special thanks to all researchers and students who took part to this project, namely (in alphabetical order): D. Beauvais, B. Bourdel, J. Coulombe, S. Fan, F. Li, R. Losseau, K. Ratelle, M. Sauvé, J. Solis and A. I. Tammam.

1 Introduction

Since the seventies, load management via direct load control and its predominantly price induced variant, demand response, have been considered as tools for shaping the load demand in power systems so as to achieve peak load shaving and valley filling [21, 28]. Such measures help defer generation and capacity expansion, and operate generators in the vicinity of their most efficient operating point. However, the seriousness with which demand response (or load dispatch) are being considered is a relatively recent phenomenon. It is mainly caused by the increasing share of intermittent renewable energy sources (such as wind and solar) in the energy mixes of power producers worldwide [3, 7, 5, 2]. The ensuing variability and diminished predictability of generation has indeed put strong pressures on the ability of independent system operators to maintain grid stability and insure reliable power delivery and transmission to consumers in their power pool [6, 1].

One of the striking characteristics of so-called smart grids is an ability to rely on both an increasingly pervasive communication network and an improved electricity distribution network, to shift the responsibility of balancing electricity demand with power generation from being solely a generation side task, to one shared increasingly by both producers and customers. Furthermore, the role of consumers is gradually changing in that they can also contribute to power generation or delivery, mostly through rooftop solar panels, electric batteries in dwellings or electric vehicles [4, 10, 18]. This drastically modified electric grid landscape has produced both a need to anticipate the new dynamics via adequate modeling tools, and a need to develop a new control theoretical framework capable of handling a very large number of heterogeneous control points and the presence of a large number of agents or decision makers. In that context, ideas of hierarchical control architectures are beginning to emerge. On top of the hierarchical control architecture sits a coordinator level [19, 27, 20]. The coordinator can be either the power system operator or a third party called aggregator, who by means of financial incentives, has enlisted the commitment of a large pool of customer loads. The loads must be coordinated so as to deliver an aggregate load behavior which is desirable from the power system point of view. The key to the scalability in these architectures is that the coordinator level must work with a collection of low order macro models of the various homogeneous load subclasses in order to identify optimal feasible aggregate load trajectories at a reasonable computational cost. Subsequently, global requirements must translate into microscopic individual load control actions. In [19, 27], the microscopic control actions are still dictated centrally. By contrast, in [20], the devices switch probabilistically according to a coordinator signal strength, and convergence to aggregate objectives is achieved via both control and reliance on the law of large numbers. The current approach is closer to [20], except that in producing the microscopic device-wise actions to attain the global objectives, the recent theory of mean-field (MF) games [13] is applied. Devices are attributed an individual cost function which encapsulates conflicting objectives: local customer comfort and safety, but with compulsory global objectives enforced on the mean behavior of the population of devices [8, 9]. This has been developed within a three year project called Smart Distributed Energy Storage Controller (smartDESC).

Storage is essential for mitigating the intermittent character of pervasive renewable energy sources, such as wind and solar. One largely underutilized form of storage, but whose potential is increasingly considered, is the distributed storage capacity naturally present in the grid such as (i) Electric Water Heaters (EWHs) and electric space heaters in households and commercial or industrial buildings [27], (ii) the anticipated large numbers of batteries in electric vehicle parks to be [13], (iii) the thermal energy storage added in a number of large buildings for load management purposes [24], (iv) the thermal energy associated with heating/cooling loads in general [19], and (v) highly flexible loads such as swimming pool pumps [20]. Challenges concerning this type of capacity are multi-fold, not the least of which is the sheer number of individually small but diversified contributors involved in most power systems (in the millions) which need to be mathematically modeled, reached through a pervasive communication system, monitored for safety and comfort, and optimally coordinated for best performance. In what follows, an overview of the components of the smartDESC architecture is presented.

The smartDESC concept consists in laying out a scalable computation and communication architecture for turning a large pool of participating energy storage capable loads into a virtual battery under the control of a single *aggregator*. While in our presentation, the *aggregator* is considered to be the power system operator, the proposed architecture could work equally well for a third-party aggregator offering load dispatch services on the energy markets. The focus of the analysis is on EWH type of loads. The proposed smartDESC architecture is displayed in Figure 1.

At the top left sits a *coordinator*: its function is to produce piecewise-constant "optimal" targets for the mean energy content per device in the aggregate, or equivalently, mean water temperature, over successive 30-minute periods. At the top right, a node represents the renewable generation forecast. At the bottom of the figure, we can find the collection of controlled devices (here the EWHs). In the smartDESC architecture, EWHs must manage to collectively meet the dictated aggregate mean energy targets (i) with minimal information exchange with the *coordinator* (decentralization), (ii) with no impact on users' comfort, and (iii) while keeping the devices in their safe zone of operation. In the middle, we can observe the *communication* module, which permits the required data flows. In what follows, we further detail the operation of these modules.

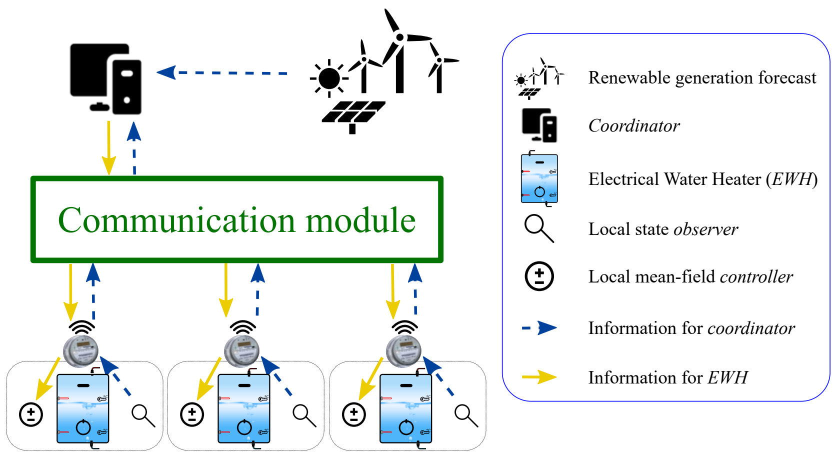

Figure 1: Architecture of the smartDESC concept.

2 The coordinator module

In order to enable the desired optimization function of the *coordinator*, i.e. that of coordinating devices to help mitigate the generation variability due to intermittency of the renewable energy sources, two conditions must be met: the *coordinator* must have a reasonably reliable, yet sufficiently simple model of the aggregate storage potential of the Electric Water Heaters (EWHs); (ii) the *coordinator* must be fed with forecasts of the renewable energy to be expected over the next 24 hours.

In order to comply with the first condition, a large EWH with water mass equal to the total water mass in the aggregated devices is used for each homogeneous subpopulation of devices. In addition, while the power of the heating element in the aggregate device is bounded above by the sum of maximal heating powers, and below by zero, in practice, only a subset of that range can be achieved by the aggregate. This range will depend on the current temperature distribution within the devices. An algorithm going back and forth between aggregate modeling, optimization, the *coordinator* module, and simulation of impact on devices is developed to correctly estimate the maximal and minimal heating rate bounds [12]; finally, the physical parameters of the aggregated EWH, including inlet water temperature, heat loss rates, and overall water extraction statistics in the controlled pool of devices must be periodically updated.

All the above functions must be realized by a *data analysis* submodule within the *coordinator* module, as shown in Figure 2. Based on the above information, the *coordinator* module calls the *optimization* submodule, which performs a deterministic or stochastic optimization over an adequately chosen time horizon. The *optimization* submodule also generates the optimal feasible mean target temperature profile for each homogeneous subgroup of EWHs. The *optimization* module can perform deterministic or stochastic optimization. In the former, classical peak shaving objectives are employed whereas, in the latter, forecasts on the renewable generation are considered. The optimization problem must be solved sufficiently quickly to allow reacting, at least in part, to unpredicted changes relative to any of the forecast quantities (the spinning reserves being the ultimate line of defense against the thoroughly intractable components of this variability). Additional details on the *optimization* submodule can be found in [26]. Ultimately, the 24-hour mean temperature profile targets for each homogeneous subgroup of controlled EWHs, together with the associated initial mean water temperatures, are transmitted to the *communication* module to be broadcasted to the corresponding subgroups of devices.

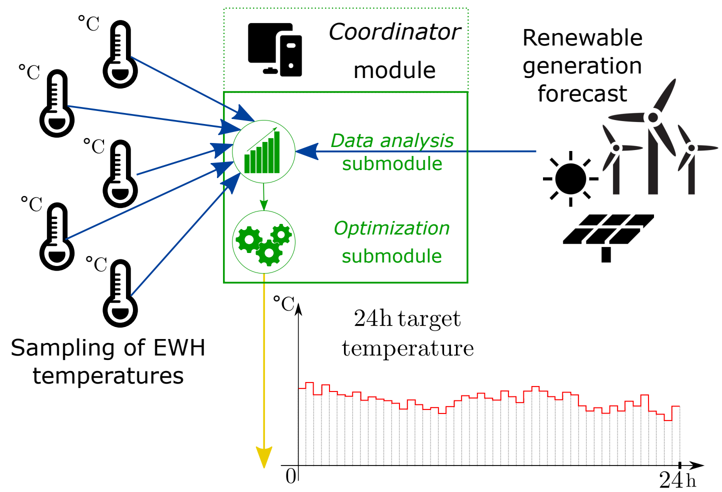

Figure 2: Schematic of the coordinator module.

3 The communication module

Telecommunications play a key role in each smart grid application: the bidirectional flow of information is a key component of the intelligence in smart grids. It is therefore important to account for and model how data flow throughout the smartDESC network. In Figure 3, the implementation of the *communication* module is depicted. In the top left of this figure, there is a simplified architecture of the proposed smartDESC simulator with one *coordinator* and 3 EWHs. As can be seen, each EWH is composed of a local state *observer*, a local *MF controller* and a *smart meter*, which is used to wirelessly transmit/receive the smartDESC packets.

Essentially, smartDESC requires two types of communication: (i) downlink, from the *coordinator* to all the local *MF controllers* of the collection of devices; (ii) uplink, in the opposite direction (see Figure 3). The *communication* module handles the transmission of 2-way information between the *coordinator* module and the collection of devices in a timely manner. The uplink direction is represented in Figure 3 by dashed blue arrows, while the downlink direction is represented in the same figure by solid yellow arrows.

Indeed, the *coordinator* needs to periodically send the target mean temperatures to all the local *MF controllers* as segregated into homogeneous subgroups. Also, it must provide to each subgroup updated information about the mean temperature of the subgroup at the beginning of each control period. Based on this information, the MF optimal state feedback policy for the particular subgroup to which an EWH belongs is locally calculated and locally applied to each device *as long as it does not violate predefined comfort and safety constraints*. In addition, in order to generate an anonymized temperature distribution estimation of the whole EWH population, each *MF controller* sends back to the *coordinator*, at low frequency and randomly distributed times, its current mean temperature with a time stamp; a monthly update of the EWH local water draw statistical mathematical model is also sent back to the *coordinator*. The overall data is consolidated at the level of the *data analysis* submodule of the *coordinator* to maintain updated versions of aggregated models.

The overall exchange rate of information is very low (on average below 10 messages per day per EWH) and the packet size is very small (in the order of few hundreds of bytes). Among the available communication architectures, it was proposed to use a particular type of Advanced Metering Infrastructure (AMI) that uses unlicensed radio frequencies to create a mesh topology between the smart meters and the power utility Metering Data Management System. This architecture is called RF-mesh and was originally deployed for reading remotely the residential and commercial power meters, but it is also suitable for applications without high communication requirements, such as those of smartDESC. The RF-mesh system is a middle layer that allows the communication between the local *MF controllers* and the *coordinator*. To prove the concept before suggesting a field deployment, a dedicated *communication* module was implemented and integrated into a smartDESC simulator. The role of the *communication* module is to model the network delay introduced by the RF-mesh system. Further details can be found in [15, 16, 14, 17].

In the lower part of Figure 3, a representation of the RF-Mesh network is depicted. As shown in this figure, the smart meters included in the smartDESC architecture represent only a fraction of all the smart meters simulated in the RF-Mesh *communication* module: this is due to the fact that the number of smart meters in a typical RF-Mesh system exceeds the number of EWHs analyzed in smartDESC. Even though a large number of smart meters included in the *communication* module are not used to carry smartDESC data, they are important to accurately model a realistic RF-Mesh system and obtain meaningful results on the network performance (e.g., collision probability, resource utilization, and delay).

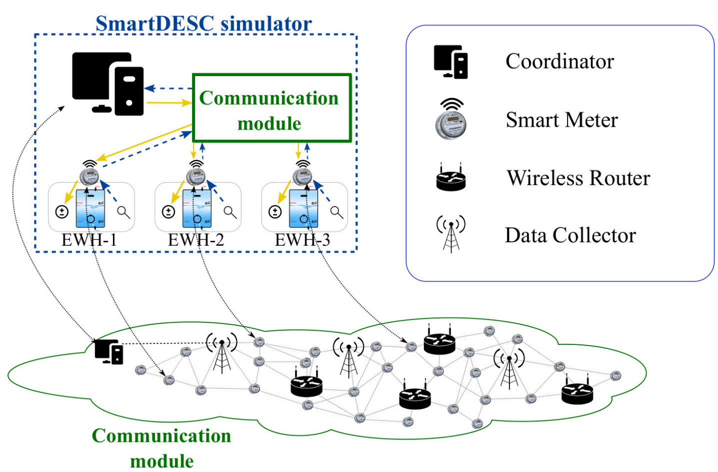

Figure 3: Schematic of the communication module.

4 The local controller

Upon receiving a mean target temperature signal from the *coordinator* module, which is specific to its particular homogeneous subgroup, the local *MF controller* must synthesize a local state feedback heating control law consistent with that signal. While the control law is common to the whole subgroup, it is computed locally so as to minimize the overall communications requirements of the architecture. This computation is enabled by a sequence of building blocks, as shown in the schematic in Figure 4: (i) an EWH stochastic state space model: it characterizes the internal dynamics of the EWH when subjected to random hot water extraction processes; (ii) a local state *observer*: it is needed because the EWH energy state, including in particular the binary hot water demand component of the state vector, cannot be directly measured, but must rather be inferred from a reduced set of measurements performed on the device; (iii) a local *MF controller*: it is the heart of the computation leading to an EWH control law which, while applied based on local information only, leads to a collective behavior of the homogeneous subgroup *consistent* with the targets dictated by the *coordinator*. We now provide further details on these building blocks.

4.1 EWH stochastic state space model

There is a trade-off between the modeling accuracy of EWHs and the simplicity of calculating aggregate dynamics for a large population. First of all, it is no surprise that the actual temperature dynamics in an EWH tank, which is subject to thermal stratification, is nonlinear. However, aggregation of nonlinear dynamics poses further complexity on the optimization module as well as for local control synthesis; therefore, a linear model is adopted. After some adjustments to reflect the nonlinear convection effects have been included in the model, the linear model has been observed to provide acceptable accuracy in comparison with simulations performed with TRNSYS, a widely accepted thermal system simulation program. Linearity in dynamics is a key component for this work. The linear model is depicted in Figure 4. In this specific case, there are 10 water stratification layers, one local *MF controller* connected to two heating elements, and one local *observer* module, which receives the temperature from two sensors. Cold water flows from the bottom layer to the top layer when water draw events occur. Furthermore, the statistics of hot water extraction events are modeled as a binary state Markov chain with switching rates that vary depending on the hour of the day [11].

4.2 Local state observer

As seen in Figure 4, the thermal dynamics of an EWH is represented through a 10-dimensional state, representing the average temperature of different water stratification layers in the EWH. In this particular case, the heating elements and the temperature sensors are located on the fourth and the ninth layers, respectively. As previously mentioned, the local *MF controller* uses a local vector state feedback law, which needs two pieces of information: (i) the mean temperature of all EWH model strata $T_i, \forall i = 0, \ldots, 9$, and (ii) the boolean value θ, which indicates whether there is water drawn from the tank at a particular instant or not. The objective of the *observer* module is to estimate these values through the readings of two temperature sensors placed near existing heating elements of water heaters (see Figure 4). Further details on the modelling of the local state *observer* can be found in [23]. Let us denote the temperature estimates as $\hat{T}_i$, and the binary hot water demand state estimate as $\hat{\hat{\theta}}$.

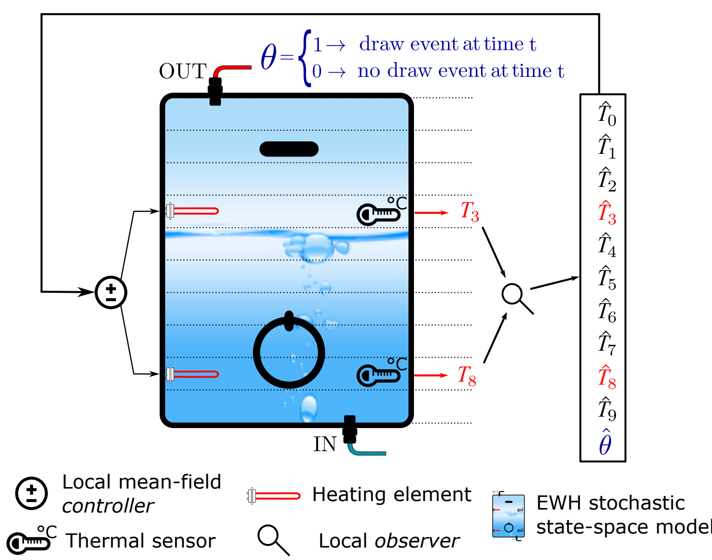

Figure 4: Schematic of the local controller of an EWH.

4.3 Local MF controller: local actions with global results

A significant challenge of a large-scale control of dispersed energy storage in power systems is the presence of literally millions of control points, each contributing a small amount to the overall load and generation smoothing effort. A classical direct control view of this problem, as has been the case in traditional load management programs, would quickly run into the limits of large (and expensive) data communication requirements, and more importantly a very large computational burden for the synthesis of the required control signals. Instead, the current hybrid control architecture that was adopted enables one to address this issue: it is centralized at the system level for optimal scheduling, as described previously, but decentralized at the implementation level to allow scalability and increase the level of customer acceptance of the controls (since enforcement of security and comfort constraints becomes local). Such an organizing architectural principle is possible thanks to a state-of-the-art development in control theory known as *MF control*. Roughly described, MF control relies on the statistical predictability of large numbers of similar controlled plants or devices to produce optimal controls (i) with significantly reduced communication requirements, and (ii) using distributed computation to apply the required local control signals. Simply put, MF control is the main enabler for smartDESC.

Let us further detail what MF control is and how it can help us in achieving the goal of best steering millions of individual energy storage elements so that their aggregated behavior follows an optimal target trajectory specified at the systems level (further details can be found in [8, 9]). As stated before, a linear dynamics has been adopted for EWHs. A careful choice of a quadratic cost function induces a linear feedback controller, with coefficients calculated backwards in time thanks to so-called Riccati equations. The coefficients are parametrized by θ, a boolean process. The state 0 represents *no water draw* event, whereas 1 represents *water draw*. A scheme of the operation of the local *MF controller* for say EWH_{ij}, i.e. the i^{th} EWH in the j^{th} subgroup, is reported in Figure 5. Note that the *MF controller* receives the 24-hour target profile for the EWH optimal mean temperature of EWHs in the j^{th} subgroup, as computed by the optimization module. It also receives the estimated mean temperature in the j^{th} subgroup at the start of the control horizon. This information is then used by the local *MF controller* to compute a local state feedback law through a fixed-point calculation (see [8] for details). This control law is applied in the feedback loop (see the lower part of Figure 5). The local *MF controller* receives the estimate of the temperature $\hat{T}_i$ of all the layers i in the EWH as well as $\hat{\theta}$. This data is used to set the value of the instantaneous power of the two heating elements of the EWH using an optimal feedback control law. The updated temperature is then captured by the two thermal sensors (i.e., T_3 and T_8). The *observer* employs these two values to update the temperature vector estimates (i.e., $\hat{T}_i$) as well as water draw state estimates (i.e., $\hat{\theta}$). The remarkable fact is that, thanks to the law of large numbers, a locally computed and implemented control law leads at the aggregate level to the mean behavior dictated by the *coordinator*, and thus to a system-wide optimal behavior. Note however that for this to happen, the data concerning the EWH physical parameters in homogeneous subgroups must be regularly updated (for instance, the water inlet temperature changes over the seasons). In addition, the statistical parameters of the water usage statistics must also be updated periodically. These parameter updates can be carried out locally and sent back to the optimization module - at low rates - through the communication module.

5 The smartDESC simulator

The smartDESC simulator was developed to design and evaluate load control, renewable penetration and demand response algorithms **in a realistic setting**. The whole environment setup is implemented using Jade (Java Agent Development Framework), which provides a versatile platform for parallel computing. The smartDESC framework provides a layer of abstraction facilitating the design of custom models and provides base classes, conventions and guidelines for their development. It also provides an interface for the user and the environment in which components can run in a coherent manner.

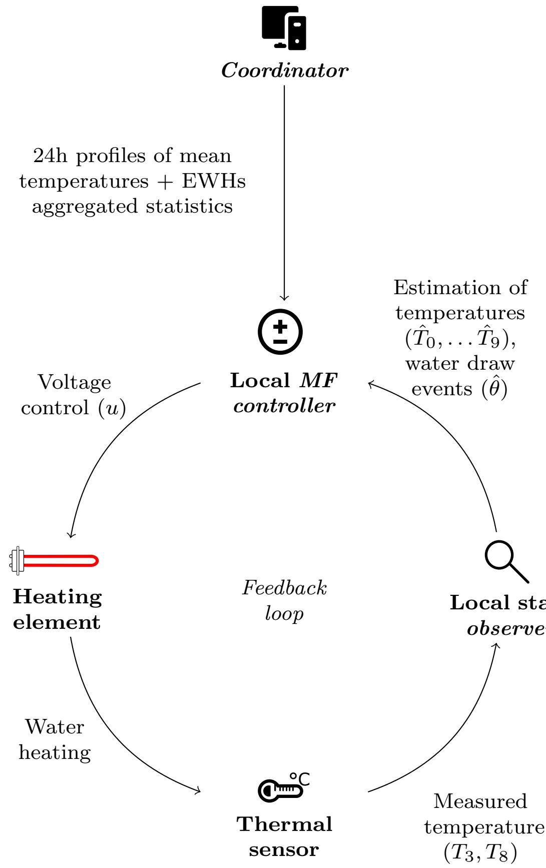

Figure 5: Schematic of the feedback loop in the local controller.

Two scenarios simulation results are presented here. The case presented is based on a population of 400 smartDESC-controlled EWHs, constituting the experimental group. The comparison benchmark is a system of same size, but with traditional thermostatically controlled EWHs and setpoints at 56°C and 61°C, constituting the control group. These 2 temperatures represent the boundaries of the customer comfort zone: when the temperature is below 56°C a traditional thermostatic controller activates the heating elements, which are subsequently turned off when the temperature reaches 61°C. In order to visualize the full effect of the local smartDESC *MF controllers*, the thermostats have been disabled for the experimental group: by doing this, the activation of the heating elements is solely decided by local *MF controllers*. Another note is that the *observer* module is also disabled; we basically assume that the exact temperature profile of the tank, and the instantaneous water draw state are available to the *coordinator* at each time instant. This gives a theoretical upper bound on the performance that can be achieved with this approach. The first simulation result is displayed in Figure 6. At time 0, the optimization module is provided with the seven-day profile of (i) total power consumption (green curve), and (ii) consumption due to thermostatically controlled EWHs.

The optimization module receives the uncontrolled load forecast, which corresponds to the total consumption minus the thermostatically controlled EWH consumption, together with the 7-day power consumption profile for the uncontrolled EWH population in a Business-as-usual (BaU) scenario. The optimization module also calculates the red curve on the figure, which represents the target power profile for EWHs. The purple curve represents the total target power, i.e., the uncontrolled load plus the EWH target power. Simulation results are also shown: the blue curve represents the power consumption of controlled EWHs while the orange curve corresponds to the aggregate power consumption of the whole system.

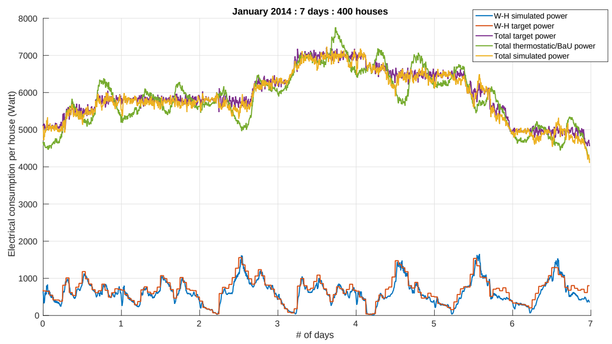

Figure 6: Simulation results for a deterministic scenario and classical peak shaving objectives.

Simulation results indicate that the mean power profile of the EWHs (blue) satisfactorily tracks the target (red), and that the total simulated power consumption (orange) is close to the target power (purple). The total simulated power curve (yellow) also shows a reduced variability and considerable peak shaving, when compared to the total consumed power in the reference scenario (green).

In the second simulation, a more complex scenario was considered, assuming the *coordinator* is provided with *random* forecasts of wind production and of uncontrolled electric load demand. Additional details on the generation of wind power scenarios from numerical weather predictions can be found in [25] The simulation results in this scenario are displayed in Figure 7. In this figure, the upper border of the cyan area represents the uncontrolled demand, whereas the cyan area itself shows the potential time-varying range of wind power produced. The bottom border of the cyan area represents the net uncontrolled demand. Given the forecast, the *coordinator* computes a target profile for the EWH population (red curve). Summing up the uncontrolled demand with the EWH target power, one obtains the purple curve, representing the total target power. The green curve represents, as in the previous case, the total consumed power when regular thermostatically-controlled EWHs are employed. The behavior of smartDESC controlled EWHs is subsequently simulated: the blue curve represents the simulated power consumption of the 400 EWHs and the yellow curve represents the total simulated power. Load balancing characteristic of the yellow curve can easily be seen in the figure, in comparison with the green curve. Of particular interest is the valley in the early parts of day 2, where most of the valley filling is due to the peak of wind production.

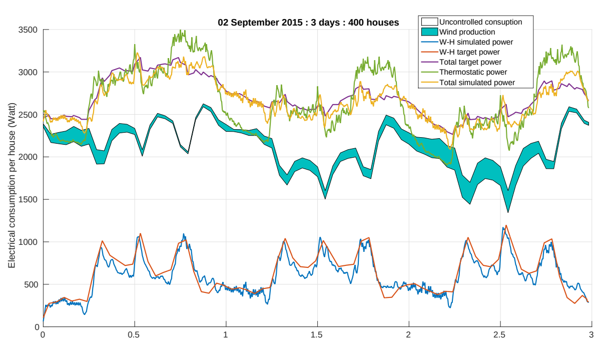

Figure 7: Simulation results in a more complex scenario with forecasts on wind production.

After having shown the positive effects of using smartDESC-controlled EWHs on the aggregated behavior of the system, we now focus on population dynamics. Figure 8 shows how the temperature profile of EWHs evolves in time. It is seen that the system maintains variability in temperature at all times but the transitions are smooth. The absence of sudden changes in the temperature distribution is important because it shows that the use of *MF*-based control does not cause steep changes in the tank mean temperature, which would constitute an undesirable situation. Fairness is also important: one of the objectives of the control architecture is that the burden of lowered or increased mean temperature should be evenly distributed across the population. In order to put the fairness characteristic of the control algorithm to test, we report in Figure 9 the temperature curves of three randomly selected EWHs (yellow, green and purple curves). The orange and blue curves represent the top and bottom deciles, respectively. One can see that individuals switch position quite often. We did not observe the unfortunate scenarios with some EWHs predominantly within the cold temperature or some others always within the hot temperature zone. The two scenarios would cause the undesirable consequences of a reduction of customer comfort or a higher cost of heating, respectively.

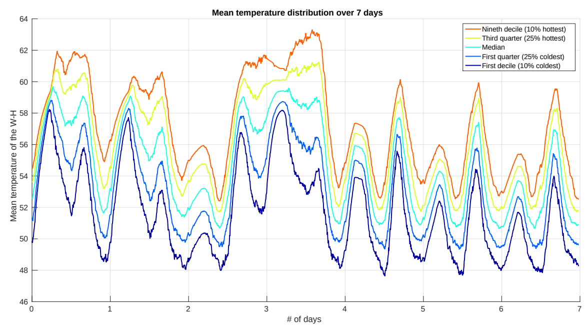

Figure 8: Percentile evolution of the mean temperature of the EWHs.

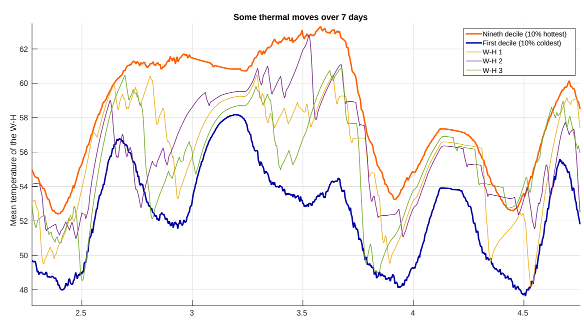

Figure 9: Temperature profile of three EWHs and percentile evolution.

In terms of communications, numerical results show the delay in the downlink direction to be significantly higher than the uplink delay. This is because the transmission times of uplink packets, randomly generated by each local *MF controller*, are usually well spread over time. On the other hand, the *coordinator* transmits the same target temperature to all the local *MF controllers* at the same time, in a broadcast fashion. This mechanism causes network congestion, mainly due to the limited size of node buffers. Moreover, the required utilization of the communication channels is very low (the smart meters transmit packets only about 0.13% of the time on average).

6 Conclusion

The smartDESC concept offers an integrated architecture to turn the energy storage potential of distributed existing electrical devices into a reliable and responsive asset for load and generation leveling. The foundation of smartDESC lays in a recent theoretical development, i.e. *MF* control theory, which is used to manage the macro-micro separation of optimization at the utility level and control at the local level. It constitutes a novel paradigm that could find applications in many different areas of engineering. The smartDESC architecture has been implemented in a simulator and tested on realistic case studies scenarios involving an homogeneous population of EWHs in a grid with renewable penetration. The results have shown that it works as expected, and no substantial issues of scalability are expected to arise. It is worth mentioning that a physical local *MF controller* installed on a real EWH has also been implemented and tested jointly with the simulated EWHs (not shown in this paper), and the concept also proved to work on the physical device. The smartDESC concept is therefore ready to evolve towards a hardware testing stage. More details about the smartDESC project can be found in the public report available online [22].

For further reading

[1] Thomas Ackermann, Enrico Maria Carlini, Bernhard Ernst, Frank Groome, Antje Orths, Jon O'Sullivan, Miguel de la Torre Rodriguez, and Vera Silva. Integrating variable renewables in europe: Current status and recent extreme events. IEEE Power and Energy Magazine, 13(6):67–77, 2015.

[2] Dana-Alexandra Ciupăgeanu, Gheorghe Lăzăroiu, and Linda Barelli. Wind energy integration: Variability analysis and power system impact assessment. Energy, 185:1183–1196, 2019.

[3] California Energy Commission et al. California renewable energy overview and programs. California Energy Commission. Accessed February 18th, 2013.

[4] Georges El Rahi, Walid Saad, Arnold Glass, Narayan B Mandayam, and H Vincent Poor. Prospect theory for prosumer-centric energy trading in the smart grid. In 2016 IEEE Power & Energy Society Innovative Smart Grid Technologies Conference (ISGT), pages 1–5. IEEE, 2016.

[5] J Gillis. A tricky transition from fossil fuel: Denmark aims for 100 percent renewable energy. New York Times, 2014.

[6] Lion Hirth and Inka Ziegenhagen. Control power and variable renewables. In 2013 10th International Conference on the European Energy Market (EEM), pages 1–8. IEEE, 2013.

[7] Noah Johnson and Diana Kennis. Feed-in tariffs for solar powered cities. Available at http://www.worldwatch.org/node/5430.

[8] Arman C Kizilkale and Roland P Malhamé. Collective target tracking mean field control for markovian jump-driven models of electric water heating loads. In Control of Complex Systems, pages 559–584. Elsevier, 2016.

[9] Arman C Kizilkale, Rabih Salhab, and Roland P Malhamé. An integral control formulation of mean field game based large scale coordination of loads in smart grids. Automatica, 100:312–322, 2019.

[10] Marko Kovač, Gašper Stegnar, Fouad Al-Mansour, Stane Merše, and Andrej Pečjak. Assessing solar potential and battery instalment for self-sufficient buildings with simplified model. Energy, 173:1182–1195, 2019.

[11] JC Laurent and RP Malhamé. A physically-based computer model of aggregate electric water heating loads. IEEE transactions on power systems, 9(3):1209–1217, 1994.

[12] Romain Losseau. Modélisation agrégée de chauffe-eau électrique commandé par champ moyen pour la géstion des charges dans un réseau. Master's thesis, École Polytechnique de Montréal, 2016. Available at https://publications.polymtl.ca/2175/.

[13] Zhongjing Ma, Duncan S Callaway, and Ian A Hiskens. Decentralized charging control of large populations of plug-in electric vehicles. IEEE Transactions on Control Systems Technology, 21(1):67–78, 2013.

[14] F. Malandra and B. Sansò. Analytical performance analysis of a large-scale RF-Mesh smart meter communication system. In 2015 IEEE Power Energy Society Innovative Smart Grid Technologies Conference (ISGT), pages 1–5. IEEE, November 2015.

[15] F. Malandra and B. Sansò. A simulation framework for network performance evaluation of large-scale RF-mesh AMIs. Simulation Modelling Practice and Theory, 75:165–181, 2017.

[16] Filippo Malandra. A Framework for the Performance Analysis and Simulation of RF-Mesh Advanced Metering Infrastructures for Smart Grid Applications. PhD thesis, École Polytechnique de Montréal, 2016. https://publications.polymtl.ca/2422/.

[17] Filippo Malandra and Brunilde Sanso. PeRF-Mesh: A performance analysis tool for large scale RF-Mesh-based smart meter networks with FHSS. In 2015 IEEE International Conference on Smart Grid Communications (SmartGridComm), pages 792–797. IEEE, 2015.

[18] Hannah Mareike Marczinkowski and Poul Alberg Østergaard. Residential versus communal combination of photovoltaic and battery in smart energy systems. Energy, 152:466–475, 2018.

[19] Johanna L Mathieu, Maryam Kamgarpour, John Lygeros, and Duncan S Callaway. Energy arbitrage with thermostatically controlled loads. In 2013 European Control Conference (ECC), pages 2519–2526. IEEE, 2013.

[20] Sean Meyn, Prabir Barooah, Ana Bušić, and Jordan Ehren. Ancillary service to the grid from deferrable loads: The case for intelligent pool pumps in florida. In 52nd IEEE Conference on Decision and Control, pages 6946–6953. IEEE, 2013.

[21] Bridger M Mitchell, Willard G Manning, and Jan Paul Acton. Electricity pricing and load management: Foreign experience and California opportunities. Rand, 1977.

[22] Frédéric Sirois, Benoit Bourdel, and Roland P Malhamé. Management of distributed energy storage capacity scattered in electric power systems for damping the variability of renewable energy sources – PUBLIC REPORT of project RENE-034. Natural Resources Canada, July 2017. Available at https://www.nrcan.gc.ca/energy/funding/current-funding-programs/eii/16102.

[23] Jérôme Solis. Développement d'un estimateur d'état énergétique d'un chauffe-eau pour un contrôle par champ moyen. Master's thesis, École Polytechnique de Montréal, 2015. Available at https://publications.polymtl.ca/2014/.

[24] Fatemeh Tahersima, Jakob Stoustrup, Soroush Afkhami Meybodi, and Henrik Rasmussen. Contribution of domestic heating systems to smart grid control. In 2011 50th IEEE Conference on Decision and Control and European Control Conference, pages 3677–3681. IEEE, 2011.

[25] A. I. Tammam, C. S. Watters, M. F. Anjos, and M. Gendreau. A methodology for ensemble wind power scenarios generation from numerical weather predictions. In 2016 IEEE Power and Energy Society General Meeting (PESGM), pages 1–5, July 2016.

[26] Adham I. Tammam. Lissage optimal de la charge électrique en présence de sources d'énergies renouvelables via le pilotage de la consommation des chauffe-eau. PhD thesis, École Polytechnique de Montréal, 2016. Available at https://publications.polymtl.ca/2254/.

[27] Evangelos Vrettos, Stephan Koch, and Göran Andersson. Load frequency control by aggregations of thermally stratified electric water heaters. In 2012 3rd IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe), pages 1–8. IEEE, 2012.

[28] Xiaojing Xu, Chien-fei Chen, Xiaojuan Zhu, and Qinran Hu. Promoting acceptance of direct load control programs in the united states: Financial incentive versus control option. Energy, 147:1278–1287, 2018.