Large Electric Load Fluctuations in Energy-Efficient Buildings and how to Suppress them with Demand Side Management

Tommaso Coletta*, Robin Delabays*†, Laurent Pagnier*‡ and Philippe Jacquod*
* School of Engineering, University of Applied Sciences of Western Switzerland HES-SO CH-1951 Sion, Switzerland
† Section of Mathematics, University of Geneva, CH-1211 Geneva, Switzerland
‡ Institute of Theoretical Physics, EPF Lausanne, CH-1015 Lausanne, Switzerland
email: {tommaso.coletta, robin.delabays, laurent.pagnier, philippe.jacquod}@hevs.ch

***Abstract*—Demand side management (DSM) is known for generating synchronized behaviors of aggregated loads that can lead to large power fluctuations [1]. In contrast to this well-studied occurrence, we report here on the emergence of novel synchronized behaviors of thermostatically-controlled electric heating systems in buildings with good thermal insulation and important solar radiation gains without DSM. To suppress the resulting large load fluctuations on the distribution grid we propose a centralized DSM algorithm that smoothens the total load curve – including electric heating and all other domestic appliances – of the cluster of dwellings it pilots. Setting up the baseline load is based on weather forecasts for a receding time-horizon covering the next 24 hours, while control actions are based on a priority list which is constructed from the current status of the dwellings. We show numerically that our DSM control scheme can be generically used to modify load curves of domestic households to achieve diverse goals such as minimizing electricity costs, peak shaving and valley filling.**

***Keywords*—Demand side management, thermostatically controllable loads, direct load control, residual load**

I. Introduction.

Under the energy transition, the penetration of nondispatchable electricity productions is steadily increasing, which results in large uncontrolled fluctuations in power generation. Enforcing the balance between power demand and supply becomes a challenge, which one standardly tries to meet with electrical energy storage. An alternative to storage is demand side management (DSM) where consumption is modified to balance production [2]. In parallel to changes in production, the energy transition further aims at making human activities more energy-efficient. Buildings and households are one of the main targets as they represent 30 to 40% of the total energy consumed in western countries. Energy-efficiency is improved via increased electrification of heating and cooling. This generates new opportunities for DSM, because thermostatically controlled loads (TCL) such as electric heaters, AC coolers and water boilers are characterized by a significant usage flexibility [3]. Also, buildings have a sizeable thermal inertia that allows to delay or anticipate electric heating operation. Anticipated operation allows to store electric power as thermal energy, while delayed operation releases part of that thermal energy in the form of reduced power demand. A variety of DSM schemes based on TCL have been proposed for minimizing electricity costs [4], to provide active power reserves [5], [6], [7] or ancillary services such as primary voltage control [8] and primary frequency control [9], [10], [1].

We have found that electric heating systems in energy-efficient buildings submitted to homogeneous weather conditions undergo a synchronization transition where a large fraction of them switches on and off in unison – a fact that has not been recognized so far. Fig. 1 shows the aggregated electric heating load of a collection of 1000 energy-efficient individual houses. Atmospheric conditions – shown in Panel a) – are obtained from a historical time series corresponding to sunny winter days in February in the city of Paris, France. Details of the calculation are discussed below. During four consecutive days the weather is clear, solar heat gains through windows are important and contribute significantly to the heating of the buildings. One sees in panel b) that this leads to the switching off of most of the heating systems at around noon. Solar gains next vanish as the sun sets, which leads to the switching on of most heating systems. We show below that this synchronous operation persists over periods of sunny weather and is damped only after few consecutive days with reduced solar radiation. It arises systematically in energy-efficient buildings with large solar energy gains and good thermal insulation, during sunny winter days. Fig. 1 b) further

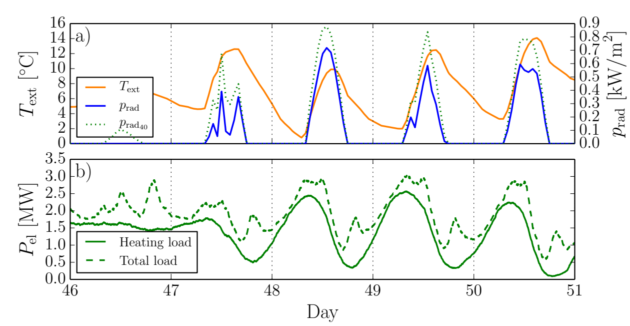

Fig. 1. Uncontrolled operation of 1000 aggregated households over one typical winter week in Paris, France. a) External temperature and solar radiation data. b) Full line: load of the electric heating systems (heat pumps with coefficient of performance $\epsilon_{\rm COP}$ = 3). Dashed: total load, including heating systems and domestic appliances. Vertical dotted lines indicate midnight.

978-1-5090-3358-4/16/$31.00 ©2016 IEEE

shows that oscillations clearly affect the total load, obtained by summing the consumptions of the electric heating and of the domestic appliances. Because they are dephased with respect to solar radiation, these oscillations cannot be directly compensated by local photovoltaic productions. One unwanted, and so far overlooked consequence of the energy transition is thus that increasing energy efficiency in buildings eventually leads to a synchronous behavior of large sets of electric heating systems that may hamper the operation of electric power distribution systems.

In this manuscript, we propose a centralized DSM coordination scheme for non-disruptive peak shaving and valley filling of the load curve of aggregated households. By non-disruptive, we mean with no significant impact on end-use performance. Our method is based on a new concept which we call *residual consumption*, and which in the present context refers to the sum of all non-flexible consumptions (mostly those of domestic appliances) from which non-flexible delocalized productions (mostly photovoltaics) are subtracted. Clearly, the residual consumption measures the amount of electricity that the considered group of households has to request from its electricity provider at a given time.

Our aim is therefore to predict moments of high and low residual consumption in a time window spanning, for instance, the next 24 hours. Then, in order to provide a service to the distribution system operator (DSO) by smoothing the overall load curve, the flexible electric heating systems are turned on at times of lesser residual consumption and turned off when it is larger. This is done by piloting the loads to try and match a predicted optimal load curve.

The paper is organized as follows: Section II describes the model used to account for heat exchanges and temperature control in buildings. Section III presents the details of the implementation of our control algorithm and compares it to conventional thermostat regulation. The simulation parameters and the results obtained are presented in Section IV. A brief conclusion is given in Section V.

II. Electric Heating Modeling

We consider a standard model for thermostatically-constrained electric heating [11], [12]. The state of each building is described by a single average internal temperature T whose time-evolution is governed by

$$C\,\frac{dT}{dt} = \kappa\;\left[T_{\rm ext}(t) - T\right] + P_{\rm h}(t) + P_{\rm rad}(t)\,. \tag{1}$$

The left-hand side gives the change in energy stored in the building's thermal mass. The constant C [Wh/K] is a thermal inertia and depends on the volume and the material the building is made of. The first term on the right-hand side represents the heat exchange with the exterior. The effective thermal conductivity κ [W/K] is determined by the building's insulation – a better insulation means a smaller κ. The terms $P_{\rm h}$ [W] and $P_{\rm rad}$ [W] in (1), are respectively the thermal power of the heating system and the solar radiation gain. The latter is obtained from meteorological solar radiation data, the window surface S and a transmission coefficient g,

$$P_{\rm rad}(t) = p_{\rm rad}(t)\;S\;\;g\,. \tag{2}$$

The heating power provided by the heating system $P_{\rm h}(t)$ is a discrete valued function, which either vanishes when the heating is off or takes on its nominal value, $P_{\rm h}(t) = P_{\rm h}^{\rm n}$, when the heating is on. Focusing on energy-efficient buildings, the heating systems we consider are heat pumps with a constant coefficient of performance (COP) $\epsilon_{COP} = 3$. This is reasonable for modern ground source heat pumps.

III. Thermostatic Control vs. Demand Side Management.

A. Purely Thermostatic Control

A thermostat keeps a building's temperature within a certain range. In a conventional control scheme, heating is turned on as soon as the internal temperature of the building reaches a minimal temperature fixed by the user. It then keeps operating until the temperature reaches a maximal setpoint, at which time it is turned off. The building's temperature then falls down, leading to a new cycle. The minimal and maximal temperatures define a comfort temperature interval within which the building's temperature should remain at all times, $T(t) \in [T_{\rm ref} - \Delta, T_{\rm ref} + \Delta]$, $\forall t$. It is defined by a reference temperature $T_{\rm ref}$ and a tolerance temperature interval 2Δ.

Solar radiation significantly contributes to heating. In order to maintain the temperature within the comfort bracket, it must therefore be controlled, e.g. by shutting window blinds as soon as the temperature exceeds a maximal set point. This strongly suppresses the solar gains and is modeled in (2) by a reduction of g.

B. DSM Coordination Algorithm

As an alternative to the thermostat regulation protocol just outlined, we present a simple control scheme whose aim is to shift the electric consumption required for heating when it is more beneficial. Here "beneficial" may mean "financially beneficial" (minimizing electricity costs), "beneficial from the point of view of the DSO" (providing ancillary service, smoothing the load curve and so forth) or "beneficial for self-consumption" (allowing to consume locally as much as possible of a local production).

The constraint of maintaining each house in the appropriate temperature interval provides restricted flexibility to shift load. In order to increase the overall flexibility we aggregate the thermostatically controlled heating systems of a large district and focus on their collective behavior. We first discretize time into intervals Δt ($t_i = i\,\Delta t$) and encode the information of when it is desirable to consume in a predetermined target consumption profile $\{P_i\}$, generated by a central controller. This profile is chosen, for instance, according to one of the specific objectives listed above. It is based on weather forecasts for the solar radiation $\{P_{{\rm rad},i}\}$ and the outside temperature $\{T_{{\rm ext},i}\}$, $i = 1, \ldots N$, in a time-window corresponding to a receding time-horizon $N\Delta t$. How to effectively construct $\{P_i\}$ is discussed below. Here we always consider a time-horizon of 24 hours. At each time step, loads adapt their instantaneous consumption, so that the aggregated consumption $\{E_i\}$ minimizes the deviations from $\{P_i\}$ as measured by the variance

$$F\left(\{P_i\},\{E_i\}\right) = \sum_{i=1}^{N}\left(E_i - P_i\right)^2\,. \tag{3}$$

Minimizing F in Eq. (3) is easy if the total load is ideally flexible, i.e. the total consumption can be distributed over any time interval of the day, and if it can take any value between zero and the maximum total power. Under this assumption, $\{E_i\}$ is allowed to vary freely, with the only constraint that the total daily consumption $W_{\rm tot}$ is fixed,

$$\sum_{i=1}^{N} E_i \, \Delta t = W_{\rm tot} \, . \tag{4}$$

Then an analytical closed form solution to the minimization of F under the constraint (4) is readily obtained using the method of Lagrange multipliers,

$$\tilde{E}_i = \frac{W_{\rm tot}}{N \Delta t} - \frac{1}{N} \sum_{l=1}^{N} P_l + P_i \, . \tag{5}$$

The existence of an analytical solution indicates that the problem is well posed. However is Eq. (5) of much use, given that individual TCLs are not ideally flexible ? We next argue that the aggregation of sufficiently many TCLs increases load flexibility overall and renders the problem closer to that of ideally flexible loads. First, with $M \gg 1$ TCLs, one may legitimately hope that at any time there are enough individual loads ready to turn on or off – the aggregated load can be increased or decreased at any time to adapt to DSM goals. Second, the aggregated consumption increases roughly in steps of $P^{\rm n}_{\rm tot}/M$, the aggregated nominal power $P^{\rm n}_{\rm tot}$ divided by the number M of individual loads, so that the relative power consumption becomes effectively a continuous variable for $M \gg 1$. This is why we make use of the enhanced overall flexibility brought about by aggregation. We stress however, that in the following, individual TCLs are never assumed fully flexible.

The control algorithm works as follows. At every time step, the central controller estimates the thermal energy consumption $W_{\rm tot}$ of the district for the receding time window covering the next 24 hours. This is made based on weather (temperature and solar radiation) forecasts and on the current energy content of the buildings. Let i_0 denote the considered time step. From (1), the estimate for $W_{\rm tot}$ at i_0 is obtained as

$$\begin{aligned} W_{{\rm tot},i_0} = & \sum_k \sum_{i=i_0}^{i_0+N} \left[\kappa^{(k)} \left(-T_{{\rm ext},i} + T^{(k)}_{\rm ref} \right) - p_{{\rm rad},i} \, S^{(k)} \, g \right] \, \Delta t \\ & + \sum_k C^{(k)} \left[T^{(k)}_{\rm ref} - T^{(k)}_{i_0} \right] , \end{aligned} \tag{6}$$

where the sum over k runs over the whole district, with quantities corresponding to individual dwellings labeled by superscripts (k). The first term on the right-hand side of (6) corresponds to the energy needed to maintain every building at its reference temperature, given the external temperature forecast, minus the energy provided by the forecasted solar radiation. The second term additionally takes into account the excess or deficit in stored thermal energy at the time of the estimate, compared to the energy stored at the reference temperature.

The total electric consumption of the neighborhood is the sum of a flexible and a non-flexible contribution. The flexible consumption can be shifted in time, at least partially. The non-flexible consumption corresponds to all consumptions that either cannot be shifted, or whose flexibility is not used. In our case, the flexible consumption is that of the electric heating system, while the non-flexible consumption is the sum of the consumptions of all other domestic appliances. The district's need for electric power is further quantified by the residual consumption R_i, which, as discussed above, is the sum of all non-flexible consumptions minus all non-flexible local productions. Our goal is to construct a coordination algorithm shifting the flexible consumption in such a way that the sum of the residual consumption and of the flexible consumption is smoother than without control. Thus, we want the electric heating system to function mostly in the valleys of the residual consumption. To achieve this, we chose the target consumption profile [the P_i's in (3)], as

$$P_i = \frac{W_{\rm tot}}{N \, \Delta t} - \epsilon_{\rm COP} \, R_i \, . \tag{7}$$

From (5) the optimal thermal consumption profile is given by

$$\tilde{E}_{i_0+i} = \left(\frac{W_{{\rm tot},i_0}}{N \, \Delta t} + \frac{1}{N} \sum_{l=1}^{N} \epsilon_{\rm COP} \, R_{i_0+l} \right) - \epsilon_{\rm COP} \, R_{i_0+i} \, . \tag{8}$$

Note that in (7) and (8) the residual consumption is an electric power and therefore it must be multiplied by the COP to obtain the corresponding thermal power. Obviously, the total electric consumption $\tilde{E}_{i_0+i}/\epsilon_{\rm COP} + R_{i_0+i}$ gives a smoothly varying function of time. It is given by the sum of the forecasted heating energy required for the next 24h plus the residual consumption for the next 24h, distributed uniformly over the N time intervals.

With estimates for $W_{\rm tot}$ and $\{R_i\}$ in hand, the central controller generates the optimal consumption profile $\{\tilde{E}_i\}$ in the receding time window $i = i_0, \ldots i_0 + N$, according to (8). The controller's task is then to adapt the district's actual load to $\{\tilde{E}_i\}$, under the end-use constraint that each bulding's inside temperature lies at any time within the comfort interval, $T^{(k)}_i \in [T^{(k)}_{\rm ref} - \Delta^{(k)}, T^{(k)}_{\rm ref} + \Delta^{(k)}]$. In our approach, this end-use constraint has priority over everything else. The procedure goes as follows. The central controller receives from each building: i) the temperature $T^{(k)}_{i_0}$ and the comfort temperature-interval, and ii) the nominal power $P^{{\rm n},(k)}_{\rm h}$ and the state $s^{(k)}_{i_0}$ of each heating system. A priority list is then constructed by ranking the houses according to a priority index

$$\eta^{(k)}_{i_0} \equiv \frac{T^{(k)}_{i_0} - \left(T^{(k)}_{{\rm ref},i_0} - \Delta^{(k)} \right)}{2 \, \Delta^{(k)}} \in [0, 1] \, , \tag{9}$$

encoding how close the $k^{\rm th}$ house is to the minimal ($\eta^{(k)}_{i_0}$ small) or to the maximal temperature ($\eta^{(k)}_{i_0}$ closer to one).

The controller next computes the district's instantaneous consumption,

$$E^{\rm inst}_{i_0} = \sum_k P^{{\rm n},(k)}_{\rm h} \, s^{(k)}_{i_0} \, , \tag{10}$$

which is the heating load that the district would have with purely thermostatic control. The controller then compares $E^{\rm inst}_i$ to $\tilde{E}_i$. If $E^{\rm inst}_i < \tilde{E}_i$, it instructs the coldest buildings, starting with those with lowest $\eta^{(k)}_{i_0}$, to turn on their heating systems until the actual load reaches $\tilde{E}_i$. If $E^{\rm inst}_i > \tilde{E}_i$ on the other hand, it instructs the warmest buildings, starting with those with highest $\eta^{(k)}_{i_0}$, to turn off their heating system. Implemented

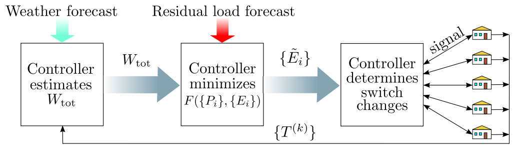

Fig. 2. Schematic representation of the control loop.

blindly, this procedure could lead either to violations of the temperature constraint $T_i^{(k)} \in [T_{\rm ref}^{(k)} - \Delta^{(k)}, T_{\rm ref}^{(k)} + \Delta^{(k)}]$ or to frequent switchings of the heating systems (the number of which determines the lifetime of a heat pump) or both. To avoid that, the heating devices react to off (resp. on) commands only if the building's temperatures belong to $T \in [T_{\rm ref} - 0.9\ \Delta, T_{\rm ref} + \Delta]$ (resp. $T \in [T_{\rm ref} - \Delta, T_{\rm ref} + 0.9\ \Delta]$) – when temperatures are outside these intervals, thermostat rules discussed in Sec. III-A apply. Thus our control algorithm accounts both for the "willingness" [quantified in terms of the priority index (9)] and the "availability" (corresponding to which pumps are available to change their operating state) of the TCLs to participate in the DSM coordination.

Once the switching commands are sent, the controller lets the district evolve according to (1) during one time step Δt. The new building temperatures are then transmitted back to the central controller, which closes the control loop [see Fig. 2]. The next step is then initiated with a new estimate of $W_{\rm tot}$ and of $\{R_i\}$ for the next 24h, a new priority list is constructed and the procedure just outlined is implemented again. In our receding horizon approach, new estimates of $W_{\rm tot}$ and $\{R_i\}$ and a new optimization of the consumption profile are performed at each time step. This allows to correct forecast errors rapidly.

IV. Simulation Parameters and Results.

A. Building parameters

In all our simulations, we consider a district consisting of 1000 individual houses, modeled by 1000 copies of (1) discretized in time steps $\Delta t = 1$ min. The parameters we use in these differential equations are different for each house and distributed as follows. The thermal capacity and thermal conductivity are uniformly distributed in the intervals

$$C \in [13, 27]\ \text{kWh/K}\,, \quad \kappa \in [200, 400]\ \text{W/K}\,, \qquad (11)$$

which are a common ranges for detached European houses with a floor surface ranging from 100 to 200 m^2. These thermal conductivity values include both recent energy efficient houses ($\kappa \in [200, 300]$ W/K) as well as older houses as commonly built in the early 90's ($\kappa \in [300, 400]$ W/K) [7].

We calibrate the nominal heating power of each heat pump as $P_{\rm h}^{\rm n} = \kappa \cdot 30\,\text{K}$, so that the pumps can provide sufficient heating power to maintain the houses at $T = 20°$C against an external temperature of $T_{\rm ext} = -10°$C. For this choice, and given (11), the nominal thermal powers of the heating systems are uniformly distributed in the interval

$$P_{\rm h}^{\rm n} \in [6, 12]\ \text{kW}\,. \qquad (12)$$

The houses in our district may have different orientations and are thus subject to different solar gains. This is modeled by distributing their effective south facing window surface uniformly in the interval

$$S \in [5, 15]\ \text{m}^2\,. \qquad (13)$$

Furthermore, the transmission coefficient of the windows is either $g = 0.6$ or $g = 0.12$ when the blinds are up or down respectively. Blinds in different dwellings are closed at the setpoint $T_{\rm blinds}^{(k)}$, uniformly distributed in the interval $[T_{\rm ref}^{(k)} + \Delta^{(k)}/2, T_{\rm ref}^{(k)} + \Delta^{(k)}]$, and are opened back at $T_{\rm ref}^{(k)}$. The comfort interval parameters $T_{\rm ref}$ and Δ reflect end-user habits and preferences. They are uniformly distributed as

$$T_{\rm ref} \in [21, 23]\ °\text{C}\,, \quad \Delta = 1.5\ °\text{C}\,. \qquad (14)$$

B. Non-flexible domestic appliances

The electrical load of the district we consider is the sum of the consumption of domestic appliances and of the electric heating systems. DSM coordinates only the latter as domestic appliances are considered non-flexible in this work. We generate load profiles from domestic appliances with the software BEHAVSIM which is based on recorded consumptions of different household appliances [13]. The resulting load for all domestic appliances in our district is shown in Fig. 3 a). The non-flexible load displays a characteristic double peak structure with peaks at noon and in the early evening. A smaller peak is also present in the early morning of working weekdays. There is a $\sim$ 1.3 kW daily excursion per house between the peak consumption of 1.5 kW and the base load of 0.2 kW.

In investigations of districts with their own, local PV production, we try to additionally consume locally as much of the PV production as possible. The quantity of interest in this case is the residual consumption. For a district with 1000 houses and 10^4 m^2 of solar panels it is shown in Fig. 5 a). Compared to the consumption of domestic appliances shown in Fig. 3 a) we see that, on sunny days, the PV production erases the noon consumption peak. The residual consumption even becomes negative at times of relatively low consumption and large PV production, indicating a need to either store or export part of the local PV production.

C. Meteorological data

Time series for the external temperature $T_{\rm ext}(t)$ and solar radiation data $p_{\rm rad}(t)$ are obtained from the software Meteonorm, which gives either true (recorded) or interpolated data[1]. Below, we discuss sequentially different situations of typical winter conditions in Paris, France, both with and without PV production that needs to be self-consumed.

D. District without local production

For a district without local production, the residual consumption is given by the load of the domestic appliances. An example of such residual consumption is shown in Fig. 3 a). It corresponds to 1000 individual houses in January in Paris, France, for which temperature and solar radiation time series are shown in Fig. 3 b). Without PV production, the solar radiation influences our results only via the radiation gains defined in (1) and (2).

[1]For more details we refer the reader to *meteonorm.com*.

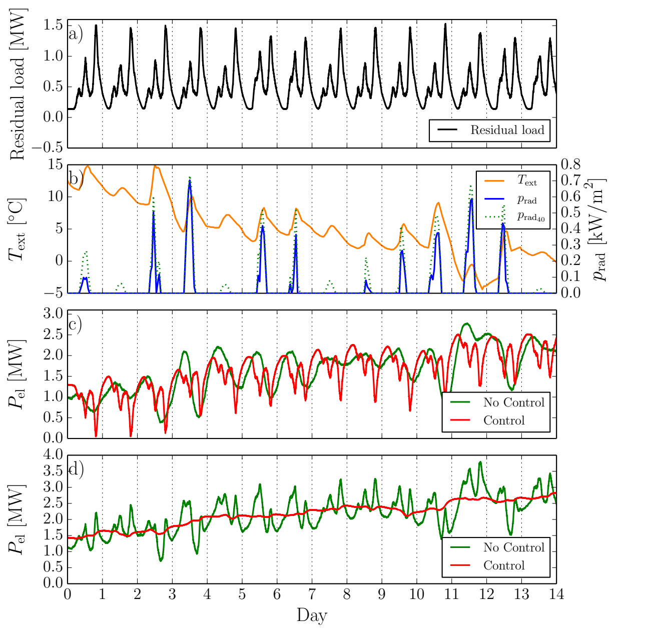

Fig. 3. Time series for electrical consumptions and weather conditions for two January weeks in Paris, France. Consumptions are aggregated for a district of 1000 houses and vertical dotted lines correspond to midnight. a) Residual load from all domestic appliances (there is no PV production). b) External temperature and solar radiation data. c) Total electrical consumption of the heating systems. d) Total load given by the sum of the heating consumption and of the residual load.

Results for the electric heating and total loads, with and without DSM coordination are shown in panels c) and d) of Fig. 3. It is clearly seen in Fig. 3 c) that, without DSM coordination, the heat pumps tend to synchronize and the heating consumption correlates with the external weather conditions. The trend is especially visible on clear, cold days, when solar radiation provides enough heating power [the last term in (1)] so that the need for electric heating is lower from midday until the late afternoon. The temperature drops fast, however, when the sun sets, at which time a large fraction of heat pumps switches on and, because the night is cold, they stay on until the next day. If that day is also clear, a new cycle starts and the pumps are almost perfectly synchronized. This behavior is observed almost throughout the month (though only results for the first two weeks are reported here), with fluctuations being damped, but not suppressed, on cloudy days with little or no solar radiation.

Numerical simulations performed using different weather conditions (not shown) corroborate our finding that, without control, large load fluctuations already emerge after a single sunny day and reach their maximal amplitude of 1.5-2.0 kW/house after two consecutive clear days. We conclude that the onset of synchronization of heat pumps is fast in energy-efficient buildings, and mostly depends on solar radiation, external temperature being a subdominant factor.

The synchronization of the heat pumps is also reflected in the total load shown in Fig. 3 d), which exhibits daily fluctuations reaching 2 MW for our set of 1000 houses. Striking in Fig. 3 d) is that the total load on the distribution network exhibits sharp ramp-down of the order of 1 MW in just two hours or less during the sunniest days of the month. Such abrupt variations need to be smoothed to avoid potential service disruptions. Achieving this is a priori a hard task: it requires to operate a significant fraction of heat pumps at times of higher solar gains and to turn them off during cold night hours.

Fig. 3 c) and d) show how the DSM coordination algorithm proposed manages to smooth the total load of our district. The heating consumption in Fig. 3 c) is clearly anticorrelated with the load from the domestic appliances. The result is that the total load, shown in Fig. 3 d), becomes remarkably smooth. With DSM, the amplitude of daily fluctuations is suppressed by a factor five or more compared to thermostatic control.

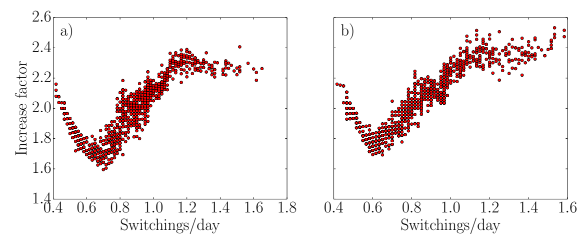

Fig. 4. Increase in the number of heat pump switchings resulting from DSM, a) no PV production, b) 10 m^2/house of PV panels. Horizontal axis: average daily number of heat pump switchings without DSM. Vertical axis: multiplicative factor of the number of switchings with DSM.

Any DSM protocol such as ours should have almost no negative impact on end-use performance. For the results presented in Fig. 3 c) and d), we have that $T_i^{(k)} \in [T_{\text{ref}}^{(k)} - \Delta^{(k)}, T_{\text{ref}}^{(k)} + \Delta^{(k)}]$ for all houses at any time of the simulation. Additionally, we find that the total consumption with DSM is few percent smaller than without DSM. Finally, Fig. 4 a) shows that, with DSM, the number of switchings of the heat pumps increases only by a factor of 1.6 to 2.4, which is tolerable.

For other weather conditions (not shown), we find that our DSM protocol performs as well as it does for Paris. It is non-disruptive in the sense discussed above, and it reduces the daily excursions of the total load below 0.25 MW for 1000 buildings.

E. Non-flexible domestic appliances and local PV production

We next consider a district with a local PV production corresponding to a total of 10^4 m^2 of PV panels for our 1000 houses. We generate the PV production time series by multiplying the solar radiation data on a south facing, 40° degrees inclined surface by the total surface of PV panels, times an efficiency coefficient of 15%. Our goal is now not only to smoothen the load curve but to do so while simultaneously consuming as much as possible of the PV production. This is not a trivial task – solar panels produce the most when solar gains are maximal, i.e. when electric heating systems would rather switch off. We show that our DSM protocol can still achieve this goal, at least for not too large PV penetration.

The residual load, shown in Fig. 5 a), is now given by the domestic appliances' consumption minus the PV production. Meteorological conditions are the same as in the previous example [see Fig. 3 b)]. The aggregated heating consumption of the 1000 heat pumps is presented in Fig. 5 b) for both coordinated and uncoordinated operation. In the uncoordinated

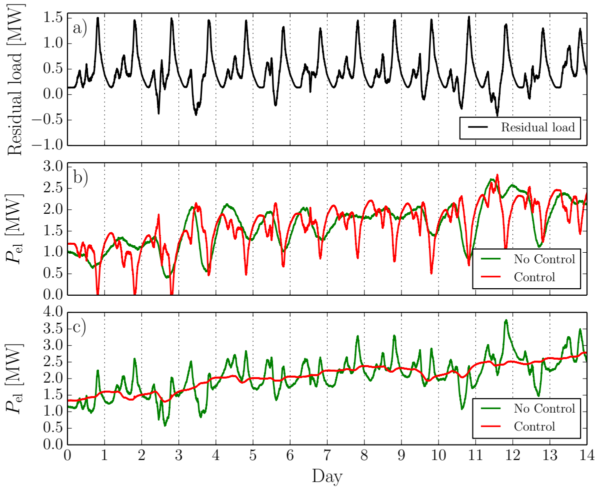

Fig. 5. Time series for electrical consumptions for two January weeks in Paris, France. Consumptions are aggregated for a district of 1000 houses and vertical dotted lines correspond to midnight. a) Residual load from all domestic appliances minus a local PV production from 10^4 m^2 of solar panels. b) Total electrical consumption of the heating systems. c) Total load given by the sum of the heating consumption and of the residual load.

case the results are the same as those presented in Fig. 3 c) – they do not depend on PV production. In contrast, the heating consumption in the coordinated case is now higher at noon, to compensate the PV production. Fig. 5 c) finally presents the total load on the distribution grid, which is about as smooth as without PV production. We conclude that our algorithm is able to smooth the load of a relatively large district while simultaneously absorbing a local PV production, even with large solar radiation power – despite the fact that the latter simultaneously reduces the need for heating and increases PV production.

We note that, with PV production, our coordination protocol still meets our requirement of being non-disruptive: we observe no violation of the comfort temperature interval, nor a dramatic increase in the number of heat pump switches [see Fig. 4 b)]. We even find a small, though not statistically significant, reduction of the total consumption with DSM coordination compared to thermostatic control.

V. Conclusion.

Heat pumps operated by purely thermostatic control in energy-efficient buildings quickly synchronize during clear winter days due to strong solar heat gains. We found that typically 60% of the heating systems turn on and off simultaneously, which leads to large load modulations and sharp ramp-ups and -downs of more than 1 MW in just two hours for 1000 houses. This undesirable synchronization is likely to become more pronounced as buildings become more and more energy efficient.

We have constructed a non-disruptive DSM algorithm that strongly suppresses these large load fluctuations and demonstrated its load-shifting potential over long periods of time. We furthermore showed that the same control scheme is able to additionally absorb a local PV production, even though the latter is produced at the same time solar gains are maximal and heat pumps have a natural tendency to switch off. We note that our algorithm is robust in that it works well in different geographical regions, regardless of meteorological conditions.

Our proof-of-principle algorithm relies on a two-way communication protocol between a central controller and local load controllers. We see this as the main shortcoming in our approach and future works should attempt to adapt this idealized control scheme to one-way communication.

Acknowledgment

This work has been supported by the Swiss NSF under an AP Energy Grant and by the HES-SO under the I1 funding scheme. The authors thank G. Basso, P. Ferrez, P.-O. Moix, G.-A. Morand and P. Roduit for useful discussions and especially D. Gabioud for ideas that initiated this project.

References

[1] K. Dehghanpour and S. Afsharnia, “Electrical demand side contribution to frequency control in power systems: a review on technical aspects,” *Renewable and Sustainable Energy Reviews*, vol. 41, pp. 1267 – 1276, 2015.

[2] D. Callaway and I. Hiskens, “Achieving controllability of electric loads,” in *Proceedings of the IEEE*, vol. 99, no. 1, Jan 2011, pp. 184–199.

[3] J. L. Mathieu, M. Dyson, and D. S. Callaway, “Using residential electric loads for fast demand response: The potential resource and revenues, the costs, and policy recommendations,” in *Proceedings of the ACEEE Summer Study on Buildings, Pacific Grove, CA*, Aug 2012, pp. 189–203.

[4] M. Kamgarpour, C. Ellen, S. Soudjani, S. Gerwinn, J. Mathieu, N. Müllner, A. Abate, D. Callaway, M. Franzle, and J. Lygeros, “Modeling options for demand side participation of thermostatically controlled loads,” in *Bulk Power System Dynamics and Control - IX Optimization, Security and Control of the Emerging Power Grid, IREP Symposium*, Aug 2013, pp. 1–15.

[5] S. Koch, D. Meier, M. Zima, M. Wiederkehr, and G. Andersson, “An active coordination approach for thermal household appliances; local communication and calculation tasks in the household,” in *PowerTech, IEEE Bucharest*, June 2009, pp. 1–8.

[6] N. Lu and Y. Zhang, “Design considerations of a centralized load controller using thermostatically controlled appliances for continuous regulation reserves,” *IEEE Transactions on Smart Grid*, vol. 4, no. 2, pp. 914–921, June 2013.

[7] G. Maître, G. Basso, C. Steiner, D. Gabioud, and P. Roduit, “Distributed grid storage by ordinary house heating variations: A swiss case study,” in *Euromicro Conference on Digital System Design*, Aug 2015, pp. 241–249.

[8] K. Christakou, D.-C. Tomozei, J.-Y. Le Boudec, and M. Paolone, “Gecn: Primary voltage control for active distribution networks via real-time demand-response,” *IEEE Transactions on Smart Grid*, vol. 5, no. 2, pp. 622–631, March 2014.

[9] Y. Mu, J. Wu, J. Ekanayake, N. Jenkins, and H. Jia, “Primary frequency response from electric vehicles in the great britain power system,” *IEEE Transactions on Smart Grid*, vol. 4, no. 2, pp. 1142–1150, June 2013.

[10] K. Samarakoon, J. Ekanayake, and N. Jenkins, “Investigation of domestic load control to provide primary frequency response using smart meters,” *IEEE Transactions on Smart Grid*, vol. 3, no. 1, pp. 282–292, 2012.

[11] S. Ihara and F. Schweppe, “Physically based modeling of cold load pickup,” *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-100, no. 9, pp. 4142–4150, Sept 1981.

[12] S. Kundu, N. Sinitsyn, S. Backhaus, and I. Hiskens, “Modeling and control of thermostatically controlled loads,” in *17th Power Systems Computation Conference (PSCC), Stockholm, Sweden*, August 2011.

[13] P. Ferrez and P. Roduit, “Non-intrusive appliance load curve disaggregation for service development,” in *IEEE International Energy Conference (ENERGYCON)*, May 2014, pp. 813–820.