

Effects of dynamic-demand-control appliances on the power grid frequencyE. B. Tchawou Tchuisseu,¹ D. Gomila,¹ D. Brunner,^{1,2} and P. Colet¹¹*Instituto de Física Interdisciplinar y Sistemas Complejos, IFISC (CSIC-UIB), Campus Universitat Illes Balears, E-07122 Palma de Mallorca, Spain*²*Département Optique Institut FEMTO-ST UMR 6174, Université Bourgogne Franche-Comté CNRS, 25030 Besançon cedex, France*

(Received 28 March 2017; revised manuscript received 5 July 2017; published 2 August 2017)

Power grid frequency control is a demanding task requiring expensive idle power plants to adapt the supply to the fluctuating demand. An alternative approach is controlling the demand side in such a way that certain appliances modify their operation to adapt to the power availability. This is especially important to achieve a high penetration of renewable energy sources. A number of methods to manage the demand side have been proposed. In this work we focus on dynamic demand control (DDC), where smart appliances can delay their switchings depending on the frequency of the system. We introduce a simple model to study the effects of DDC on the frequency of the power grid. The model includes the power plant equations, a stochastic model for the demand that reproduces, adjusting a single parameter, the statistical properties of frequency fluctuations measured experimentally, and a generic DDC protocol. We find that DDC can reduce small and medium-size fluctuations but it can also increase the probability of observing large frequency peaks due to the necessity of recovering pending task. We also conclude that a deployment of DDC around 30–40% already allows a significant reduction of the fluctuations while keeping the number of pending tasks low.

DOI: [10.1103/PhysRevE.96.022302](https://doi.org/10.1103/PhysRevE.96.022302)**I. INTRODUCTION**

Since the mid-1990s, several studies have been carried out on how to control, stabilize, and improve the efficiency of the electricity power grid [1–9]. Electric fluctuations have two main contributions. On the side of demand, households and commercial users attached to a grid do not have ideal constant or periodic needs. A significant amount of devices are switched on and off at an arbitrary time over the day. On the side of generation, although the majority of the production is provided by nuclear and fossil power plants, renewable energy sources account already for a significant fraction of the production (over 20% in Spain, for instance) and there is an overall growing interest in their use. Renewables are also constantly subjected to fluctuations at different time scales, for instance, turbulence of the wind on windmills, clouds passing over photovoltaic panels, etc. The fluctuations on the demand side combined with the fluctuations of the production due to renewable sources constantly unbalance production and demand, creating fluctuations of the frequency of the system. We note that in some instances a group of households may decide to switch on a device at the same time, due to, for example, a TV program break. Indeed, the largest “TV Pickup” recorded in the UK was about 2.8 GW at the end of the 1990 World Cup semi-final between England and West Germany [10]. Here, we are concerned with random fluctuations. Large fluctuations triggered by external factors will not be considered in this work.

In general, electricity generated by a power plant and consumed in households and commercial users is characterized by three main properties: voltage, current, and frequency. Despite some proposals [7,11], there is practically no electricity storage capacity in the power grid; therefore the power generated at any time has to match exactly the power consumed by all the loads attached to the grid [4]. Thus, the current is expected to vary according to the consumption. However, voltage and frequency must be kept within acceptable ranges by the system operator as established by the legislation of each country. For instance, in most parts of

Europe, the grid is designed to run at 50 ± 0.5 Hz and tension variations below 7% [12,13]. Eventual mismatches between supply and demand modify the frequency of the grid. In fact, if at any time power demand exceeds supply, the frequency falls (associated to turbine rotation slowdown). Frequency is then a good proxy to monitor the stability of the power grid.

Based on the mismatch between supply and demand, due to the fact that turbine governors cannot follow fluctuations in demand fast enough, several approaches to control the demand side of the grid fluctuations have been proposed. One is dynamic demand control (DDC), which provides frequency regulation by controlling the demand side of the grid, reducing the load when the grid is under stress and increasing it when there is a surplus generation [4,14]. While stabilizing the power grid, DDC aims at making it more resilient against power outages, as well as saving spinning reserve availability. DDC can be implemented as an external control or integrated within the electric appliances [4,15]. DDC algorithms can work in rechargeable devices such as laptops or mobile phones or with appliances for thermal applications like air conditioners, refrigerators, electric boilers, etc. These devices sometimes can delay or advance their operation minutes or even hours without disrupting the user comfort.

Besides DDC, other works proposing demand-side control are the following: In 1979 Schweppe *et al.* in Ref. [16] proposed the concept of “frequency adaptive power energy re-scheduler” (FAPER). They proposed the idea of individual load control through responding to frequency. Econnect Ltd. developed a “Distributed Intelligent Load Controller” using additional loads to handle excess generation and load shedding to handle shortages based on frequency sensing using fuzzy logic [17]. The UK firm ResponsiveLoad Ltd. has developed a frequency-dependent load controller that uses various frequency limits to affect the probability of switching [18]. Recently, a dynamic demand response where the price of electricity is directly linked to the frequency has been considered [19].

One of the issues of using DDC applied to domestic appliances is their synchronization. The algorithm contin-

uously monitors the electrical frequency, shutting off the appliance load when the electrical frequency drops below a lower threshold and keeping it in the off state until the frequency goes above an upper threshold [12]. If all DDC-controlled appliances operate with the same threshold, then the simultaneous responses can lead to oscillatory instabilities in the frequency [4,12]. To address this issue, randomization of the action of each appliance has been suggested [4,12,20,21].

In this work we study a generic algorithm of DDC to study the effects of this technique on the power grid fluctuations. To do so, we first introduce a simple stochastic model for the demand such that when coupled to a standard power plant model it is capable of reproducing the statistical properties of the frequency fluctuations measured experimentally. We then introduce DDC on the appliances and analyze its effect on the frequency fluctuations as well as its efficiency in stabilizing the electrical power grid. We find that DDC can significantly reduce small and medium-size fluctuations; however, the recovery of pending tasks may increase the probability of large frequency peaks. Finally, we consider that only a fraction of the load of the system is controlled by smart appliances and study how this fraction affects the overall performance of the power grid.

The paper is organized as follows. In Sec. II we describe the standard power plant model to be considered in this paper. In Sec. III we introduce the stochastic model for demand which is calibrated in Sec. IV on the basis of experimental measurements of the frequency fluctuations. In Sec. V we introduce the model for dynamic demand control. In Sec. VI we study the effects of DDC on frequency fluctuations and determine the most suitable operation parameter values. In Sec. VII we analyze the how the effectiveness of DDC depends on the number of smart devices in the grid. Finally, in Sec. VIII, we give some concluding remarks.

II. POWER PLANT MODEL

A conventional power plant is, roughly speaking, constituted by a generator and a governor. The generator is the responsible to produce electricity from a fuel or a renewable energy source, and the governor is the specific control method used to match the power of the generator to the demand. The generator is typically composed by a mechanical part, often a turbine fixed on a ferromagnetic rotor which rotates between the stator winding and an electrical part formed by coils wound around the stator. Both parts are coupled magnetically. By applying the Newton law on the turbine, the well-known swing equation describing the dynamics of the generator can be derived [21]:

$$\frac{d\omega}{dt} = \frac{\omega}{2HP_G}(P_m - P_e), \quad (1)$$

where P_G and H are the nominal capacity and the inertia constant of the generator, P_m is the mechanical power generated by the turbine (or other means), and P_e is the total power of the electric current passing through the coils around the stator. Throughout this work we take $P_G = 37320$ MW based on the UK national power grid [4]. We note, however, that this value is arbitrary in this work and it fixes only the scale of P_m and P_e .

The total electric load P_e can be divided in two parts: a non-frequency-sensitive load and frequency-sensitive load, such that

$$P_e(\omega, t) = \left(1 + D \frac{\omega - \omega_R}{\omega_R}\right) P(t), \quad (2)$$

where ω_R is the grid reference frequency, D is the fraction of the load which is frequency sensitive, such as electrical motors, and $P(t)$ is the total power at $\omega = \omega_R$.

In Eq. (1), if the electrical power exceeds the input mechanical power, for instance, as a consequence of a sudden load increase, the frequency of the system decreases. Conversely, if the input mechanical power exceeds the load, the frequency increases. The governor is responsible to restore the frequency to its reference value, and it does so in two steps. Under a supply-load unbalance, the primary regulation acts within tens of seconds increasing (or decreasing) the mechanical power to halt the decline (or rise) in frequency. The secondary regulation, which acts within tens of minutes, incorporates spinning reserve to the generation in order to restore the frequency to its reference value ω_R [4,21]. Secondary regulation is conditioned to the availability of sufficient spinning reserved. Primary and secondary regulations can be respectively modeled by the following equations:

$$\frac{dP_m}{dt} = \frac{1}{\tau_g} \left[P_s - P_m - \frac{P_G}{R\omega_R} (\omega - \omega_R) \right], \quad (3)$$

$$\frac{dP_s}{dt} = -\frac{K}{\omega_R} (\omega - \omega_R). \quad (4)$$

Here R is the governor speed regulation, P_s is the spinning reserve power used at a given time, K is the gain of the secondary controller, and τ_g is the time constant of the turbine.

Figure 1 illustrates the response of the system after a sudden load increase that takes place at time $t = 1$ minute. Immediately after the load increase the frequency decreases. The first response of the power plant due to the governor, Eq. (3), acts to stop the frequency decline, and within tens of second the frequency is momentarily stabilized at a value below ω_R , ($\omega = 49.93$ Hz). After stabilizing the decline, secondary regulation, Eq. (4) returns the frequency to its




FIG. 1. Frequency response after a sudden increase of the demand. Initially, the system is operating at a constant load of 36000 MW. At $t = 1$ min a sudden increase of load of 1320 MW is applied. The reference frequency of the system is $\omega_R = 50$ Hz and the power of the generator $P_G = 37320$ MW. Other parameters (see text) are $\tau_g = 0.78$ s, $H = 4$ s, $D = 0.026$, $R = 0.04$, and $K = 6600$ MW/s.

reference value $\omega_R = 50$ Hz within about 12 min, which corresponds to a realistic power plants response time [4].

III. STOCHASTIC DEMAND MODEL

We propose a very simple model able to reproduce the main statistical properties of real demand fluctuations by fitting a single parameter. This will allow us to study general effects of applying DDC to the power grid in the next section. To model a load consisting of generic domestic appliances such as refrigerators, freezers, air conditioners, electric heaters, dishwashers, chargeable portable devices, etc, that can switch on and off at any time, we consider N devices or bunches of aggregated devices. The total load demand will be

$$P(t) = \sum_{j=1}^N P_j(t), \quad (5)$$

where $P_j(t)$ is the load of device (or group of devices) j at time t .

For the sake of simplicity, in this work we consider that $P_j(t)$ can only take the values 0 (off) or P_0 (on) [22]. We consider that appliances in the off state switch on with a rate p , while running devices switch off with rate q . This creates a fluctuating demand with statistical properties that depend on the parameters p and q . Throughout this work, we will consider that the rates are constant and identical, $p = q$, such that the average power remains constant. Time varying rates $p(t)$, $q(t)$ following daily demand patterns will be considered elsewhere.

This simple model corresponds exactly to a Markov process for a system composed of N particles, each one making transitions between two states (on, off) with rates p and q [23]. If there is no interaction among devices, then this problem can be solved exactly, and for the case $p = q$ the average power demand is $\langle P \rangle = NP_0/2$, and, in the stationary regime, the size of the fluctuations is proportional to \sqrt{N} , with variance $\sigma_P = \sqrt{N}P_0/2$. As a matter of fact, the variance of the fluctuations at all times is given by

$$\sigma_P^2(t) = NP_0^2[\mu(t) - \mu(t)^2], \quad (6)$$

where the probability $\mu(t)$ of finding a device on is given by

$$\mu(t) = \frac{1}{2}(1 - e^{-2pt}) + \mu(0)e^{-2pt}. \quad (7)$$

Figure 2 shows the evolution of the probability of finding a device on as obtained from Eq. (7) for different values of p to illustrate the difference in the characteristic time $\tau = \frac{1}{2p}$ to reach the stationary state. The probability computed from a numerical simulation averaging over a large number of noise realizations is shown for $p = 6.55 \times 10^{-4} \text{ s}^{-1}$ in blue.

We now focus on the fluctuations in the total load predicted by the model. Time scales play a very important role, as the features of the fluctuations in Fig. 3 are quite different depending on the observation time scale. For short time scales, the above stochastic process is essentially a random walk, as the on/off switchings of a device at each time correspond to the characteristic step forward or backward of a random walk. As a matter of fact, from (6) one can show that for short times the variance of the fluctuations grows as $\sigma_P(t) \propto \sqrt{t}$, characteristic of a random walk. Figure 3(a) shows a time series of the demand P for $N = 1000$ devices of power




FIG. 2. Probability $\mu(t)$ of finding a device on as obtained from Eq. (7) for different values of the switching on probability p . The corresponding characteristic times to reach the stationary states are $\tau = 125$ min (red line), 12.5 min (green line), and $\tau = 1.25$ min (black line). Here $\mu(0) = 0$.

$P_0 = 132$ MW with a high time resolution, where the discrete jumps can be clearly appreciated. We take $P_0 = 132$ MW such that for the chosen power plant parameters and $N = 1000$ the system is under moderate stress (medium-large frequency fluctuations) in order to better check the effectiveness of DDC. Then, P_0 has to be understood as the total aggregated power of a large group of appliances. We note that the same variance of the fluctuations can be obtained considering a larger number of devices with smaller power if scaled appropriately.

As the number of devices is finite, the random walk is bounded, and fluctuations can not grow indefinitely. For long times ($t \gg \tau$), the variance saturates to $\sigma_P = \sqrt{N}P_0/2$. At these large time scales the fluctuations look more like white Gaussian noise with the latter standard deviation [Fig. 3(b)]. These features become clear looking at the power spectrum of the total load $P_e(\omega) = \int e^{i\omega t} P_e(t) dt$, as shown in Fig. 4. For observations at very large time scales (low frequencies) the spectrum is flat, characteristic of white noise, while for short time scales (high frequencies) the spectrum displays a power-law shape with a slope of -2 , characteristic of a random walk [23].




FIG. 3. Demand fluctuations at different time scales for $N = 1000$ devices of power $P_0 = 132$ MW. The switching probability is $p = 6.55 \times 10^{-4} \text{ s}^{-1}$.




FIG. 4. Power spectrum of the electric load produced by the demand model for $p = 6.55 \times 10^{-5} \text{ s}^{-1}$ (red), $p = 6.55 \times 10^{-4} \text{ s}^{-1}$ (blue), and $p = 6.55 \times 10^{-3} \text{ s}^{-1}$ (green). Averages over 100 noise realizations are shown. Dashed line has slope -2 for comparison.

IV. COMPARISON OF THE STOCHASTIC DEMAND MODEL WITH REAL MEASUREMENTS

We next compare the spectrum of the frequency fluctuations obtained from numerical simulations of the model with the spectrum of frequency fluctuations measured from a wall-plug outlet. This allows us to estimate a realistic value of the probability p .

Measurements were performed using basic custom electronics implemented on a Raspberry Pi 2. The Raspberry Pi 2 device is a small, single board and general-purpose computer with a 900-MHz quad-core ARM-compatible CPU. The electrical grid voltage signal was obtained from a standard wall plug and was subsequently scaled to the voltage range of an analog digital converter IC (ADC, Texas Instruments ADC7816). The analog-to-digital converter was connected to the Raspberry Pi via a serial bus link. The grid signal was sampled using one of the CPUs cores with a resolution of 12 bits at a rate of 62.5 KSamples/s. The power grid waveform was recorded in real-time during a temporal window of 10 s which corresponds to ≈ 200 periods of the grid signal. Following completion of the 10-s sampling window, the sampling core passed the data to the CPUs second core for data analysis. As such, the device was capable to continuously sample and process the input waveform without loss of data. By measuring the temporal positions of the grid signal's zero crossings, the second CPU core obtained an average frequency for each 10-s window.

Before its utilization as a grid frequency spectrum analyzer, we extensively calibrated the stability and resolution of our device, using a Keysight 33120A signal generator as a calibration source. Figure 5 shows two example spectra obtained during calibration. Data shown in Fig. 5(a) correspond to the spectral distribution during a ~ 10 -hour calibration test using a 50-Hz signal. The obtained stability was excellent, showing no indications of drifts while reaching a spectral resolution

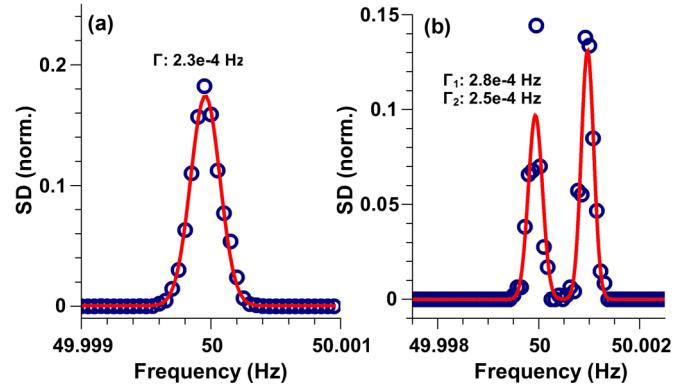


FIG. 5. Calibration of the spectrum analyzer. Power spectrum for a constant frequency input during ~ 10 h (a) and for a bimodal spectrum with a frequency separation of 1 mHz (b).

of $\sim 2.5 \times 10^{-4} \text{ Hz}$. In a second calibration, we use a signal which rapidly switches from 50 Hz to 50.001 Hz. Figure 5(b) shows the obtained spectra, demonstrating that the device is capable of clearly detecting frequency changes below 1 mHz.

Figure 6(a) shows the experimental results for the frequency obtained measuring the output of a power outlet. The frequency is observed to fluctuate in a range of $\pm 0.2 \text{ Hz}$ around the 50-Hz reference value. Figure 6(c) shows the power spectrum of the frequency. At low frequencies the spectrum is basically flat while it shows a power-law decay for large frequencies.

For comparison, Figs. 6(b) and 6(d) show the frequency time trace and power spectrum obtained from a simulation




FIG. 6. Time series (a) and power spectrum (c) of the frequency measured at an electrical outlet at IFISC. Panels (b) and (d) show the equivalent results from a numerical simulation with $p = 6.55 \times 10^{-4} \text{ s}^{-1}$, $N = 1000$, $P_0 = 132 \text{ MW}$ and power plant parameters as in Fig. 1.




FIG. 7. Frequency time series and power spectrum for a numerical simulation with $p = 6.55 \times 10^{-5} \text{ s}^{-1}$ [(a) and (c)] and with $p = 6.55 \times 10^{-3} \text{ s}^{-1}$ [(b) and (d)]. Other parameters as in Fig. 6.

of Eqs. (1)–(3) with the stochastic demand model described in Sec. III. Stochastic demand fluctuations as those shown in Fig. 3 translate to fluctuations of the frequency [Fig. 6(b)]. The power spectrum of the frequency fluctuations displays a plateau for low frequencies and decay as a power law for high frequencies [Fig. 6(d)], reproducing the characteristics of the power spectrum of the demand fluctuations (Fig. 4). Therefore, changing p in the stochastic model shifts the characteristic time scales of the frequency fluctuations as well, which allows us to adjust the value of p in order to fit the experimental data. Figure 7 shows the result generated by the model for two very different values of p . A shift in the frequency where the power spectrum starts to decay can be clearly appreciated. We finally take $p = 6.55 \times 10^{-4} \text{ s}^{-1}$ as a good value that reproduces the power spectrum of the frequency fluctuations measured experimentally [Fig. 6(a)]. For this probability of switching, on average, about six of the 1000 devices receive and order to switch on or off every 10 s.

We find that the adjusted probability p corresponds to a characteristic time for the saturation of the fluctuations around 12 min. Frequency fluctuations have flat spectrum as white noise for time scales slower than this characteristic time. For faster time scales frequency variations behave as a random walk.

The experimental data display also a superimposed periodicity, as revealed by the peak in the power spectrum around $1 \times 10^{-3} \text{ Hz}$, corresponding to oscillations with a period of approximately 15 min, clearly observable in the data, and which are not introduced in the model.

V. DYNAMIC DEMAND CONTROL

We next include a fraction $\gamma_1 \equiv n/N$ of smart devices operating under a DDC algorithm which delays the device switching on or off if the instant frequency is beyond a given tolerance. Pending tasks are recovered later in periods of

favorable frequency conditions in order to ensure that the average energy consumption at the end of the day is the same as in absence of smart devices.

The proposed DDC algorithm is implemented on top of the demand model explained in the previous section. It works as follows: At every time step devices can randomly turn on or off or remain in the same state they were before. However, for smart devices, before committing any change of state, the DDC measures the grid frequency ω and the change is only committed if ω is within a suitable range. Smart devices in the off state that randomly would switch on effectively do so only if the frequency is above a minimum level $\omega > \omega_R - \epsilon$. When a switch-on is prevented by the DDC, the missing consumed energy is accounted to be used at a later time, frequency conditions permitting. Similarly, smart devices in the on state that randomly would switch off effectively do so only if the frequency is below a maximum level $\omega < \omega_R + \epsilon$, and when a switching-off is skipped, the extra consumed energy is accounted to be saved later. The objective is that, in the long run, every smart device has used the same total energy as if it were not smart.

In the following we refer to the extra (saved) energy consumption generated by the DDC control as pending tasks. Energy-consuming pending tasks, namely pending tasks that require switching on a device to recover from a previous instance in which the device could not turn on when it was required, are recovered only if frequency is above a threshold $\omega > \omega_r + \epsilon_1$. Similarly, pending tasks saving energy are recovered only if $\omega < \omega_r - \epsilon_1$. To avoid the simultaneous switching of all devices with pending tasks when these thresholds are crossed, each device starts recovering pending tasks with probability γ . The randomization of the response of appliances is known to avoid oscillations created by the synchronizations of smart devices [4,12,20,21].

Thus the overall DDC algorithm consists then of two distinct parts: decision on committing actions and recovery of pending tasks, and it has only three parameters, namely ϵ for the allowed range to commit actions and ϵ_1 and γ for the recovery of pending tasks.

VI. EFFECTS OF DDC ON THE POWER GRID FREQUENCY

We next analyze the role of the different DDC parameters. Typically, it will be convenient to set ϵ below the statutory limit for frequency variations ($\pm 0.2 \text{ Hz}$). Throughout this paper we take $\epsilon = 0.05 \text{ Hz}$. In this section we consider that DDC is applied to all devices ($n = N$), while in Sec. VII we will address the effect of applying DDC only to a fraction of them.

We first consider a very large value of ϵ_1 so that pending tasks are not recovered in practice. In this situation, as shown in Fig. 8, the control efficiency is very good and the frequency stays almost always within the tolerance range $\omega_R - \epsilon < \omega < \omega_R + \epsilon$. The counterpart is, however, that depending on the system conditions smart devices will consume more (or less) energy than what they were supposed to consume to perform whatever task they were designed to do, accumulating pending tasks indefinitely, which is unrealistic.

We define the pending tasks Q_i of smart device i as the absolute value of the energy that this device has consumed in




FIG. 8. Time series of (a) the demand P_e , (b) frequency ω , and (c) pending tasks Q with DDC applied to all devices. We have considered $\epsilon = 0.05$ Hz, $\epsilon_1 = 0.1$ Hz, and $\gamma = 1$. Other parameters as in Fig. 6

excess or in shortage with respect to the reference case of no applying any DDC control. Pending tasks per smart device are given by $Q = \frac{1}{n} \sum_i Q_i$.

To effectively recover pending tasks we have to decrease ϵ_1 . Decreasing the value of ϵ_1 and setting recovery probability γ to 1 leads to an effective recovery of all pending task as soon as the frequency crosses the threshold. Although this avoids the accumulation of pending tasks, it leads to large frequency fluctuations (even larger than without DDC). All devices with energy demanding pending tasks turn on simultaneously when the frequency crosses $\omega_R + \epsilon_1$ from below. Similarly, all devices with pending tasks saving energy will simultaneously turn off when the frequency crosses $\omega_R - \epsilon_1$ from above.

In order to avoid this synchronized switching, it is required that $0 < \gamma < 1$ such that pending tasks are recovered progressively, avoiding sudden demand peaks. To determine suitable values of ϵ_1 and γ , we have explored this two-parameter space performing numerical simulations and computing the variance of the frequency fluctuations σ_ω^2 and number of pending tasks averaged over noise realizations $\langle Q \rangle$. Figure 9 shows the results for changing ϵ_1 for a fixed value of $\gamma = 1.2 \times 10^{-3}$. For $\epsilon_1 < -\epsilon$, pending tasks are recovered practically immediately, washing out the effect of DDC. As ϵ_1 is increased, DDC starts acting, progressively reducing frequency fluctuations. We consider only positive values of ϵ_1 as recovering pending task under unfavorable frequency conditions is not recommendable. For ϵ_1 small the average number of pending tasks reaches a stationary value. This stationary value increases slowly with ϵ_1 . A qualitative change in the number of pending tasks occurs at $\epsilon_1 \approx 0.06$ Hz, slightly above the value of the tolerance $\epsilon = 0.05$ Hz of the frequency control. At this value pending tasks increase very sharply. For ϵ_1 above this value, the number of pending tasks does not reach an stationary value, and it diverges for large times. The plateau shown in Fig. 9(b) corresponds to the value of the pending tasks after a finite time, $t = 135$ h,




FIG. 9. Frequency variance (a) and pending tasks at time $t = 135$ h averaged over 100 realizations of the noise (b). We have considered $\gamma = 1.2 \times 10^{-3}$. Other parameters as in Fig. 8.

used for the numerical simulations. Physically, frequency fluctuations are quite reduced and very rarely reach this large values of ϵ_1 , precluding the recovering of pending tasks. In this circumstance the recovery probability γ is practically irrelevant and the dynamics is similar to that shown in Fig. 8 for $\epsilon_1 = 0.1$ and $\gamma = 1$.

We note that, actually, for small values of γ and intermediate values of ϵ_1 , the recovery of pending task helps in controlling frequency fluctuations, because additional devices switch on (off) only for high (low) frequencies, helping reducing fluctuations even further. This is the reason why the variance of the frequency fluctuations can be even lower than in the case without recovery of pending tasks. As a matter of fact, the variance of the frequency fluctuations has a minimum around $\epsilon_1 = 0.037$ Hz.

Considering the variance of the frequency fluctuations σ_ω^2 and the number of pending tasks, from Fig. 9 one would conclude that the optimum value for $\gamma = 1.2 \times 10^{-3}$ is $\epsilon_1 = 0.037$ Hz, but this is not the whole story. Looking in detail at trajectories of the frequency for $\epsilon_1 = 0.037$ Hz, one observes that, despite the fact that the variance is lower than for larger values of ϵ_1 , there are extreme events in which the frequency takes very large or very small values, outside the statutory limits. These events are rare but pose a great risk to the system since they could trigger a failure or blackout. Figure 10 shows trajectories of the frequency in simulations with different values of ϵ_1 . These extreme events are caused by the random synchronization of pending tasks recovery, and their probability increases decreasing ϵ_1 .

Figure 11 shows the estimated cumulative rank probability R of finding a fluctuation larger than $\Delta\omega = |\omega - \omega_R|$ calculated as

$$R = 1 - \frac{i}{m-1}, \quad (8)$$

where i is the rank of the frequency fluctuation of size $\Delta\omega$ in a very long time series with m samples. The black dots shows




FIG. 10. Time series of the frequency for different values of ϵ_1 : (a) $\epsilon_1 = 0.022$ Hz, (b) $\epsilon_1 = 0.05$ Hz, and (c) $\epsilon_1 = 0.06$ Hz. Other parameters as in Fig. 8.

the case of random switching on and off without DDC for comparison. For the parameter considered, fluctuations above $\Delta\omega = 0.2$ Hz are practically nonexistent. When applying DDC with a large ϵ_1 , for instance, $\epsilon_1 = 0.1$ Hz, shown in yellow triangles in Fig. 11, the frequency fluctuations are largely reduced as reflected in the narrowing of the width of the probability distribution. However, since the frequency variations rarely go above $\Delta\omega = 0.1$, pending tasks are very rarely recovered and they keep accumulating. For $\epsilon_1 = 0.037$ Hz, corresponding to the minimum of the variance of the frequency fluctuations, shown in green stars in Fig. 11, the average number of pending tasks is low and small- to medium-size fluctuations are largely suppressed; however, we observe a power-law tail of the probability distribution, indicating that large fluctuations have non-negligible probabilities to occur.

A compromise is to choose $\epsilon_1 = 0.06$ Hz (gray circles in Fig. 11). For this value, we obtain a fairly low value for the variance of the fluctuations and average number of pending task, and a reasonably low probability of large events to occur, although the distribution still has power-law tails, signaling that the problem of large fluctuations due to recovery of accumulated pending task does not disappear completely. This




FIG. 11. Cumulative rank probability distribution of finding a fluctuation larger than $\Delta\omega = |\omega - \omega_R|$. Parameters as in Fig. 8.




FIG. 12. Frequency variance (a) and pending tasks per smart device at time $t = 135$ h averaged over 100 realizations of the noise (b). We have considered $\epsilon_1 = 0.06$ Hz. Other parameters as in Fig. 8.

phenomenon is similar to what it is observed for the occurrence of large blackouts in power grids due to risk-adverse policies: Strict control policies suppress small or mid-size blackouts but increase the probability of very large ones [24]. This is a signature of a system operating close to a critical condition [25]. As a matter of fact, in Fig. 11 we observe a qualitative change in the shape of the probability distributions for $\Delta\omega > \epsilon_1$, signaling the recovery of accumulated pending tasks as the cause for the long tail of the distribution.

We now focus on the effect of varying the value of the recovery probability γ . Figure 12 shows the variance of the frequency fluctuations and the pending tasks at time $t = 135$ h averaged over 100 realizations of the noise as a function of γ for a fixed value of $\epsilon_1 = 0.06$ Hz. For low values of γ the probability of recovering pending task is low, and therefore the control of the frequency fluctuations is very efficient at the expense of accumulating an increasing number of pending task. In fact, for $\gamma = 0$, no task is recovered and $\langle Q \rangle$ diverges, increasing the time. For γ not infinitesimally small, the amount of pending tasks evolves in time to a stationary value. As γ is increased, the stationary value for $\langle Q \rangle$ decreases rapidly without detrimental of the frequency variance. If γ is too large (beyond the range of the figure), then many devices start recovering pending tasks simultaneously, causing synchronization peaks in the demand and increasing frequency fluctuations again. We find, then, that good values for the threshold and probability of recovering pending tasks are $\epsilon_1 = 0.06$ Hz and $\gamma = 1.2 \times 10^{-3}$, respectively. These values have been determined for a constant fixed grid load. As a matter of fact, the optimum value of ϵ_1 depends on the load of the grid. The smaller the load, the smaller the value of ϵ_1 in order to avoid the accumulation of pending tasks; otherwise, if the load is low and ϵ_1 too high, then frequency fluctuations never reach this threshold and pending tasks cannot be recovered. The precise dependence of ϵ_1 on the system load and the effect of time-varying loads will be investigated elsewhere.




FIG. 13. Time series of the demand P_e for an increasing fraction of smart devices: (a) $\gamma_1 = 0$, (b) $\gamma_1 = 0.01$, (c) $\gamma_1 = 0.5$, and (d) $\gamma_1 = 1$. We have considered $\gamma = 1.2 \times 10^{-3}$. Other parameters as in Fig. 12.

VII. FRACTION OF SMART DEVICES ON THE GRID

We now focus on the case in which only a fraction γ_1 of devices are smart. So far we have considered the two extreme cases: no smart devices ($\gamma_1 = 0$) and all the load smart ($\gamma_1 = 1$). A more realistic case would be a grid where only part of the load is smart, while the rest keeps switching on and off according to random demand. In order to study how




FIG. 14. Time series of the frequency ω for an increasing fraction of smart devices: (a) $\gamma_1 = 0$, (b) $\gamma_1 = 0.01$, (c) $\gamma_1 = 0.5$, and (d) $\gamma_1 = 1$. Parameters as in Fig. 13.




FIG. 15. Time series of the pending tasks per smart device for an increasing fraction of smart devices: (a) $\gamma_1 = 0$, (b) $\gamma_1 = 0.01$, (c) $\gamma_1 = 0.5$, and (d) $\gamma_1 = 1$. Parameters as in Fig. 13.

the performance of the grid depends on the number of smart devices, we have performed a series of simulations with the same stochastic realization but varying the fraction of smart devices γ_1 . Figures 13, 14, and 15 show time traces of the demand P_e , frequency ω , and pending tasks per smart device Q for an increasing number of smart devices in the system.




FIG. 16. Frequency variance (a) and number of pending tasks per smart device at time $t = 135$ h averaged over 100 realizations of the noise (b). Parameters as in Fig. 13.




FIG. 17. Cumulative rank probability distribution of finding a fluctuation larger than $\Delta\omega = |\omega - \omega_R|$ for $\gamma_1 = 0.3$. Other parameters as in Fig. 13.

Figure 16 shows the dependence of the frequency variance σ_ω^2 and the number pending tasks per smart device at time $t = 135$ h averaged over 100 realizations of the noise $\langle Q \rangle$ on the fraction of smart devices γ_1 . We observe that the frequency variance decreases very fast increasing the fraction of smart devices. The variance saturates at $\gamma_1 = 0.2$. Increasing the fraction of smart devices above this value does not significantly reduce the frequency variance. The number of pending tasks per smart device decreases as well with the fraction of smart devices and saturates at a larger value $\gamma_1 \approx 0.5$. Thus for $0.2 < \gamma_1 < 0.5$, while increasing the number of smart devices has a little effect on the overall frequency fluctuations, it does reduce the average pending tasks. The frequency of occurrence of large frequency fluctuations does not change significantly for moderate reductions of the fraction of smart devices. In Fig. 17, we show the rank probability of the frequency fluctuations for a penetration of smart devices of 30% ($\gamma_1 = 0.3$), which does not present qualitative changes with respect to the 100% case.

Since the deployment of smart devices suppose an additional cost and extra complexity for the appliances, and the global performance of the grid does not significantly increase for γ_1 above 0.2–0.4, we conclude that aiming for a 30% of the total load being smart would be a reasonable objective in terms of cost-benefits for society.

VIII. CONCLUSIONS

We have proposed a simple model to study the effects of dynamic demand control on the frequency of the power

grid. Our model makes use of the well-established equations for a power plant with primary and secondary regulation and introduces a simple stochastic model for the power demand. We have shown that the model can reproduce the statistical properties of real measurements of the frequency fluctuations adjusting a single parameter, namely the switching probability of the devices.

The model also allows for the application of a generic DDC protocol to a fraction of devices in order to study its effect on the dynamics. The generic DDC protocol consists of two parts: control, by which on or off orders on devices are actually committed only if the frequency is within a suitable range, and recovery of pending tasks only when the frequency value is appropriate and performed randomly to avoid instabilities generated by simultaneous switching of all devices with accumulated pending tasks.

We have found that DDC can significantly reduce the variance of the fluctuations by delaying the switching of smart devices and recovering the pending tasks later. However, the recovery of pending tasks modifies qualitatively the probability distribution of the frequency fluctuations, introducing large tails with a power-law shape. Therefore, depending on parameters, while DDC can reduce small or medium-size fluctuations, it can also increase the probability of observing large frequency fluctuations with respect to the case without control. This rare events can potentially trigger a failure of the system and strategies to avoid them have to be addressed. We have identified the most suitable parameter range for practical operation. Finally, we have also found that there is no need to apply DDC to all devices in order to achieve significant effects. Frequency fluctuations can be effectively reduced already with 20% of devices being smart and a ratio of 30–40% allows to reduce the fluctuations while keeping the pending tasks per device low.

ACKNOWLEDGMENTS

We acknowledge helpful discussions with B. A. Carreras and financial support from Agencia Estatal de Investigación (AEI, Spain) and Fondo Europeo de Desarrollo Regional under Project No. ESoTECoS Grants No. FIS2015-63628-C2-1-R (AEI/FEDER,UE) and No. FIS2015-63628-C2-2-R (AEI/FEDER,UE). E.B.T.-T. also acknowledges the fellowship FIS2015-63628-CZ-Z-R under the FPI program of MINEICO, Spain.

[1] R. H. Lasseter, Microgrids, in *Proc. IEEE Power Eng. Soc. Winter Meet.* (IEEE, New York, NY, 2002), pp. 305–308.
 [2] R. H. Lasseter, CERTS microgrid, in *Proc. 2007 IEEE Intl. Conf. Syst. of Syst. Eng.* (IEEE, San Antonio, TX, 2007).
 [3] M. H. Nehrir, B. J. LaMeres, and V. Gerez, A customer-interactive electric water heater demand-side management strategy using fuzzy logic. In *Proc. IEEE Power Eng. Soc. Winter Meet.* (IEEE, Los Alamitos, CA, 1999), pp. 433–436.
 [4] J. A. Short, D. G. Infield, and L. L. Freris, *IEEE Trans. Power Syst.* **22**, 1284 (2007).

[5] A. Brooks, E. Lu, D. Reicher, C. Spirakis, and B. Wehl, *IEEE Power Energy Mag.* **8**, 20 (2010).
 [6] N. Lu and T. Nguyen, Grid Friendly TM appliances-Load-side solution for congestion management, in *Proc. IEEE Transm. Distrib. Conf. Exhib.* (IEEE, Dallas, TX, 2006), pp. 1269–1273.
 [7] M. Takagi, K. Yakaji, and H. Yamamoto, Power system stabilization by charging power management of plug-in hybrid electric vehicles with LFC signal, in *Vehicle Power and Propulsion Conference 2009* (IEEE, Los Alamitos, CA, 2009), pp. 822–826.

- [8] A. Molderink, V. Bakker, M. Bosman Johann, L. Hurink and G. J. M. Smit, *IEEE Trans. Smart Grid* **1**, 109 (2010).
- [9] A. M. Kosek, G. T. Costanzo, H. W. Bindner, and O. Gehrke, An overview of demand side management control schemes for buildings in smart grids, in *2013 IEEE International Conference on Smart Energy Grid Engineering* (IEEE, Oshawa, ON, 2013), pp. 1–9.
- [10] https://en.wikipedia.org/wiki/TV_pickup; Can you have a big ‘switch off’?, BBC News available at <http://news.bbc.co.uk/1/hi/magazine/6981356.stm>.
- [11] D. Shively, Energy storage methods for renewable energy integration and grid support, in *Proc. IEEE Energy 2030 Conf.* (IEEE, Atlanta, GA, 2008), pp. 1–6.
- [12] Parliament of the United Kingdom. Electricity Act (1989).
- [13] National Grid. Electricity—Real Time Operational Data. [<http://www.nationalgrid.com/uk/Electricity/Data/Realtime/>] (2011).
- [14] M. R. V. Moghadam, R. T. B. Ma, and R. Zhang, Distributed frequency control via demand responses in smart grids, in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process* (IEEE, Vancouver, BC, Canada, 2013), pp. 5233–5237.
- [15] S. Hild, S. Leavy, C. Grf, and B. Sorazu, *IEEE Trans. Smart Grid* **5**, 328 (2014).
- [16] F. Schweppe, R. Tabors, J. Kirtley, H. Outhred, F. Pickel, and A. Cox, *IEEE Trans. Power App. Syst.* **99**, 1151 (1980).
- [17] K. Pandiaraj, P. Taylor, and N. Jenkins, *IEEE Trans. Energy Convers.* **16**, 14 (2001).
- [18] D. Hirst, Responsive Load Ltd., U.K. Patent GB2361118.
- [19] B. Schäfer, M. Matthiae, M. Timme, and D. wittthaut, *New J. Phys.* **17**, 015002 (2015).
- [20] A. Mohsenian-Rad, V. W. S. Wong, J. Jatskevich, R. Schober, and A. Leon-Garcia, *IEEE Trans. Smart Grid* **1**, 320 (2010).
- [21] H. Saadat, *Power System Analysis*, 2nd ed. (McGraw–Hill, New York, 2002).
- [22] The role of diversity in the power of each group will be considered elsewhere.
- [23] R. Toral and P. Colet, *Stochastic Numerical Methods* (Wiley-VCH, Weinheim, 2014).
- [24] U. Bhatt, D. Newman, B. Carreras, V. Lynch, and I. Dobson, Understanding the effect of risk aversion on risk, in *Proc. 38th Annual Hawaii Int. Conf. on System Sciences* (IEEE, Big Island, HI, 2005), p. 64b; D. Newman, B. Carreras, V. Lynch, and I. Dobson, Evaluating the effect of upgrade, control and development strategies on robustness and failure risk of the power transmission grid, in *Proc. 41st Annual Hawaii Int. Conf. on System Sciences* (IEEE, Waikoloa, HI, 2008), pp. 182–188.
- [25] B. Ph. van Milligen, B. A. Carreras, and D. E. Newman, *Chaos* **26**, 033109 (2016).