Designing effective frequency response patterns for flexible thermostatic loads

Vincenzo Trovato, Simon H. Tindemans and Goran Strbac
Department of Electrical and Electronic Engineering – Imperial College London
London, United Kingdom

***Abstract*—Future power systems will have to integrate large amounts of wind and solar generation to drastically reduce CO_2 emissions. Achieving this goal comes at the cost of a reduced level of the system inertia and an increased need for fast response services. Previous research has shown the effectiveness of thermostatically controlled loads (TCLs) providing frequency response, and the ability to accurately control the aggregate power consumption of TCLs. In this paper, we explore the design space of frequency response patterns of flexible TCLs. Two distinct frequency response implementations are presented. The first makes the TCLs' power consumption a linear function of system frequency and/or its rate of change; in the second, TCLs respond to a frequency event tracking a pre-programmed reference power profile. Computer simulations illustrate strengths and weaknesses of the proposed implementations in the context of the GB 2020 Gone Green scenario.**

Keywords — Load management, Frequency response, Thermostatically controlled loads, Demand response.

I. Introduction

The penetration of renewable energy sources (RES) in electricity systems worldwide is quickly increasing in an attempt to reduce the greenhouse gas emissions produced by conventional generators. However, beside the environmental benefits introduced, this profound change in the traditional generation portfolio is a concern for system operators; since most of the RES (e.g. wind and solar generation) are mechanically decoupled from the AC network, they do not naturally contribute to the system inertial response [1]. With a reduced level of inertia, the absolute value of the rate of change of frequency (RoCoF) after a sudden infeed loss will increase considerably [2]. A larger RoCoF may trigger the generators' RoCoF-sensitive protection schemes; the consequent cascading disconnection process may lead, in the worst case, to a blackout [3].

During the first instants after a generator failure, the RoCoF is primarily a function of the system inertia and the amplitude of the generation-demand imbalance, before the generators' governor response is deployed. A look at the future GB generation scenario reveals the severity of the risk faced by National Grid (NG), the GB system operator. The installed wind capacity is expected to reach 25 GW by the year 2020 and the maximum infeed generation loss grows from 1.32 GW up to 1.8GW [4]. Moreover, if the frequency drops more rapidly, conventional generators may not be fast enough at limiting the deviation and keeping it within the security limit by means of their governor responses. The resulting frequency nadir would potentially activate the costly Load Frequency Demand Disconnection (LFDD) scheme [4]. As a result, the lack of system inertia may cause violations of the GB Security and Quality of Supply Standard (GB-SQSS) [5]. Hence, a significant increase of the amount of primary frequency response is required to contain the fast transient frequency evolution within acceptable limits. Previous works investigated the ability for thermostatically controlled loads (TCLs), such as refrigerators and air conditioners, to modulate their aggregate power consumption in order to provide short-term frequency response [6, 7]. However both these methods revealed drawbacks. The first showed the tendency for devices to synchronize their cooling cycles; the second solves the synchronization at the expense of a limited speed and magnitude of the power reduction available for the network.

In this paper we build upon our recent results that enable accurate decentralized and non-disruptive control of aggregate power consumption of distributed TCLs [8]. In particular we compare the performances of two distinct implementations of the hybrid stochastic-threshold controller developed in [8]. The first makes the TCLs response a linear function of the system frequency and/or its RoCoF at all times; four combinations are considered. The second implementation considers the tracking of a pre-programmed reference power profile triggered by the initial frequency deviation after a generation outage. We test the effectiveness of the two controller implementations on a basic but intuitive power system model; this model mimics the dynamics of the future GB network under the 2020 Gone Green scenario [4]. Computer simulations illustrate the increased ability of the power system to integrate renewable generation when TCLs are controlled with the strategies proposed. Furthermore, the positive contribution of demand side response is studied not only over the short-term time scale of primary response (few seconds) but also over the longer time window of the secondary response (several minutes) [9].

II. Aggregation and Control of TCLs

Recently, we proposed in [8] a novel approach to accurately control the aggregate power consumption $P(t)$ [W] of a large population of heterogeneous thermostatic appliances by means of a decentralized deterministic-stochastic controller. The capabilities of this controller are summarized by the constraints

$$P(t) = P_0 \cdot \Pi(t) \tag{1}$$

$$\Pi_{min} \leq \Pi(t) \leq \Pi_{max} \tag{2}$$

$$T_{min} \leq \bar{T}(t) \leq T_{max} \tag{3}$$

The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 283012.

978-1-4799-7993-6/15/$31.00 ©2015 IEEE

The first equation shows that the instantaneous aggregate power consumption of a population of TCLs is given by the average steady state population consumption P_0 [W], modulated by a reference signal $\Pi(t)$. The steady state is described by $\Pi(t) = 1$. It is demonstrated in [8] heterogeneous appliances can achieve this target in a decentralized manner if they share the same reference signal. The control framework allows for either an *off-line* broadcast or an *on-line* local computation of $\Pi(t)$; in the first case TCLs follow a pre-programmed profile, in the second one, they compute the reference profile depending on the actual evolution of external variables (e.g. system frequency). However, the possible reference power profiles are bounded by (2); we refer the reader to section II-B for the mathematical expression of Π_{min} and Π_{max} and to [8] for further details and the derivation of these thresholds. The final feature (3) is a temperature constraint, reflecting the need for TCLs to perform their primary function. The population average temperature $\bar{T}(t)$ [°C] – and the temperature of individual appliances – is kept within an acceptable interval $[T_{min}, T_{max}]$ at all times. This constraints is equivalent to an integral constraint on $\Pi(t)$ (as will be shown in equation 6).

The constraints (2) and (3) are sufficient conditions, so a reference signal that satisfies them is guaranteed to be implementable by individual devices. As a result, there is no need to perform device-level simulations to prove the feasibility of a specific response. Finally, the collective response does not require a central real-time communication infrastructure (costly and not immune to disturbances).

A. Short derivation of aggregate TCL modeling and control

Below we present a summary description of the TCL aggregate thermal model. Let us consider first a single generic thermostatically controlled device a; its regular operation consists in the alternation of two operating conditions. The first is a *passive* heating/cooling phase during which the power drained is essentially nil; instead, the specular condition is *active* in the sense that the power consumption is constantly at the nominal value P_{on}^a [W] (maximum). Hence, the appliances will perform a state switching from passive to active mode when the controlled temperature reaches the upper/lower temperature limit, T_{max}/T_{min}; a state jump in the opposite direction (active → passive) is triggered by reaching the lower/upper threshold, T_{min}/T_{max}. These appliances are usually described by variations on the following linear first order dynamic model [7]:

$$\frac{dT^a(t)}{dt} = \begin{cases} -\frac{1}{\tau^a}\left(T^a(t) - T_{off}^a\right) & if\ s^a(t) = 0 \\ -\frac{1}{\tau^a}\left(T^a(t) - T_{on}^a\right) & if\ s^a(t) = 1 \end{cases} \tag{4}$$

It illustrates the evolution of the temperature $T^a(t)$ of an appliance a in relation with the binary thermostat's cooling state $s^a(t)$, that can thus assume values 1 if the TCL is *on* or 0 if *off*. Moreover, the thermal time constant of the model is τ^a [s]; T_{off}^a [°C] and T_{on}^a[°C], respectively, denote the ambient temperature and the asymptotic temperature achieved by a TCLs that it is always maintained in the *on* state. T_{on}^a incorporates a physical model of heat exchange. Considering a population of N>>1 TCLs, equation (5) follows from the linearity of (4);

$$\frac{1}{N}\sum_a \frac{dT^a(t)}{dt} = -\frac{1}{N}\sum_a \left[\frac{T^a(t) - T_{off}^a + s^a(t)\left(T_{off}^a - T_{on}^a\right)}{\tau^a}\right] \tag{5}$$

This property indicates that the mean temperature $\bar{T}(t) = \frac{1}{N}\sum_a T^a(t)$ of a large population of TCLs evolves as (6), in case of TCLs with identical parameters τ, T_{on}, T_{off} (these parameters no longer need the superscript a)

$$\frac{d\bar{T}(t)}{dt} = -\frac{1}{\tau}\left[\bar{T}(t) - T_{off} + \pi(t)\cdot\left(T_{off} - T_{on}\right)\right]. \tag{6}$$

Hence, $\pi(t) = \frac{1}{N}\sum_a s^a(t) = P(t)/(N \cdot P_{on}^a)\ \epsilon\ [0,1]$ is the aggregate power consumption $P(t)$ relative to the maximum achievable level. A comparison with (4) shows that the average temperature of a cluster of TCLs actually retraces the temperature evolution of a single large device with a variable cooling rate. The definition of the steady state relative power consumption π_o directly follows from (6)

$$\pi_o = \frac{T_{off} - \bar{T}_o}{T_{off} - T_{on}}, \tag{7}$$

with $\bar{T}_0$ the average steady state temperature.

B. Additional characteristics of the control framework

The reference power profile is defined as $\Pi(t) = \pi(t)/\pi_0$, so that that $\Pi(t) = 1$ corresponds to a steady state condition for the power consumption. The design of the reference profile $\Pi(t)$ is subject to limitations in (2); in [8] we derived these constant thresholds as function of the temperatures of the TCL model. The mathematical expressions of these quantities is:

$$\Pi_{min} = \frac{\left(T_{off} - T_{max}\right)\left(\bar{T}_o - T_{min}\right)}{\left(T_{max} - T_{min}\right)\left(T_{off} - \bar{T}_o\right)} \tag{8a}$$

$$\Pi_{max} = \frac{\left(T_{max} - \bar{T}_o\right)\left(T_{max} - T_{on}\right) + \left(T_{off} - T_{max}\right)\left(T_{max} - T_{min}\right)}{\left(T_{max} - T_{min}\right)\left(T_{off} - \bar{T}_o\right)} \tag{8b}$$

Note that the control strategy developed in [8] is applied to an heterogeneous population of TCLs. In this paper we assume, for simplicity, the presence of identical TCLs, but extensions to heterogeneous device populations are straightforward.

III. Designing Controller Implementations

We now proceed with the descriptions of the two controller implementations. Their aim is to maintain the grid frequency above the minimum security threshold f_{min} and its rate of change above $RoCoF_{min}$. It is worth pointing out that this analysis will focus on *initial* demand-response power reductions and subsequent recovery process; these actions are performed to limit the frequency drop (negative value) consequent to a generator outage. Cases where the system frequency suddenly becomes larger than the nominal value (total generation larger than total system load) are therefore not considered in this paper.

A. The frequency linear controller

The control architecture introduced in previous section allows making a simple choice for tailoring demand side response profiles. Hence, we propose

$$\Pi_{lin} = 1 + K_1 \Delta f + K_2 \frac{d\Delta f}{dt} \tag{9}$$

This controller has the desirable property that, after the demand response activation, the power consumption of the TCL population is always a linear function of the system frequency (modulated by K_1) and its rate of change (K_2).

Initial proposals for temperature-deadband controllers (e.g. [6]) could not guarantee such a linear relation between power consumption and frequency. Such controllers induce an initial power reduction that is approximately proportional to the initial frequency drop, but when frequency recovers, the TCLs that have been switched off do not automatically switch back on. Instead, they do it when their temperatures reach the constant upper thresholds.

Considering (9), four sub-implementations of the linear controller are proposed:

$$A: K_2 = 0; K_1 | \Pi_{lin}(f_{min}) = \Pi_{sust} \tag{10a}$$

$$B: K_2 = 0; K_1 | \Pi_{lin}(f_{min}) = \Pi_{min} \tag{10b}$$

$$C: K_1 = 0; K_2 | \Pi_{lin}(RoCoF_{min}) = \Pi_{min} \tag{10c}$$

$$D: K_1 | \Pi_{lin}(f_{min}) = \Pi_{min}, K_2 | \Pi_{lin}(RoCoF_{min}) = \Pi_{min} \tag{10d}$$

In case A, the value of K_1 is chosen such that when the frequency achieves the minimum security threshold f_{min}, the reference power level is

$$\Pi_{sust} = 1 - \left(\frac{T_{max} - \bar{T}_o}{T_{off} - \bar{T}_o} \right); \tag{11}$$

This represents the lower *sustainable* power limit, which can be sustained indefinitely without violating the temperature constraint (3) [8]. This *conservative* assumption permits to ignore the dynamics of frequency recovery; even if frequency is constantly maintained at f_{min} the appliances will still be able to provide their support indefinitely. Case B is similar, but it provides greater frequency support because the devices achieve the is the minimum accessible instantaneous power level Π_{min} in (8a). This response cannot be maintained indefinitely, but this poses no problems in practice given the short duration of frequency support. Both cases A and B are insensitive to the RoCoF ($K_2 = 0$).

In contrast, case C uses only the RoCoF as a control signal ($K_1 = 0$). This means the TCL contribution is higher during the very first instants of the frequency transient period as the RoCoF is minimum (negative value), but the contribution vanishes when frequency is at the nadir; afterwards, the RoCoF becomes a positive quantity (still constrained by (8b)) and thus the TCL power consumption is higher than the steady state level, i.e. $\Pi(t) > 1$. The implication is that the energy borrowed from TCLs is (partially) paid back already by the end of the primary response service.

Finally, case D tailors the power response depending on both frequency and RoCoF, combining elements from B and C. This version provides rapid and substantial support as the control action is maximized from the first instants after the frequency drop (RoCoF term) and continuing until the frequency reaches the nadir (frequency term). In all cases, the on-line controller enforces that the reference power consumption $\Pi(t)$ respects constraint (2) at all times. The temperature limits in (3) are also enforced – and can be monitored by integrating (6). However, the latter are not usually binding on the time scale of frequency restoration.

B. The pre-programmed controller

The second application of the hybrid stochastic-threshold controller considers the tracking of pre-programmed power curve in response to a generator outage. This response may be triggered by the consequent frequency drop, but is not affected by the frequency signal after it commences; this additional feature could be implemented without significant changes. Nevertheless, the programmed response should depend on the system state and the size of the population of frequency-responsive loads. For instance, the tailored response might be able to avoid over-reaction that might result in a positive frequency excursion when the TCL power reduction is larger than the initial generation-demand imbalance. In order to estimate the actual power imbalance, TCLs need to calculate the rate of change of frequency and recognize the level of system inertia H. As the approximate system inertia may be predicted with reasonable accuracy, this information could be broadcast by means of smart meters, along with an indication of the total amount of frequency-responsive load on the system. The communication delays involved in this channel would not be critical. By having awareness of these quantities, the amplitude of the power imbalance ΔP can be inferred by:

$$\Delta P = 2H \frac{df(t)}{dt} \tag{12}$$

Note that (12) holds only in the first seconds after the frequency event. Depending on the detected amplitude of ΔP and the collective response capability, the appropriate pre-programmed response would be triggered for the appliance.

The TSO or the demand response aggregator has the freedom to design complex response $\Pi_{pp}(t)$ according to system commercial/technical requirement. The only hard limitations are enforced by constraints (3) and (2). For the simulations carried out we assume for simplicity the perfect evaluation of ΔP; moreover we focus on the impact of the maximum power support from TCLs in order to face the maximum expected infeed generation loss [4].

C. Demand side response activation

An important aspect is the criterion for demand response activation; both methods exploit the grid frequency signal and its rate of change to trigger the TCLs frequency response. Fig. 1 shows the activation scheme.

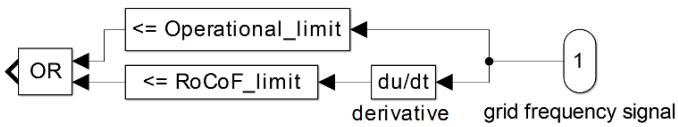

Fig. 1. Schematic of the demand side response activation.

Demand response is triggered by either one of two conditions that correspond to an emergency state for network frequency control. In the first case, frequency drops below the

operational limit [5]; in the second case, the RoCoF (negative quantity if frequency drops) falls below the minimum accepted value [5].In practice, such a design could be implemented in a fully decentralized way.

IV. Case Study

We consider two cases in accordance to National Grid's 2020 Gone Green Scenario [10]. The high wind scenario assumes 20 GW of wind power supplying the system; in the low wind scenario, the wind output is limited to 5 GW. For each scenario, we vary the system demand from 30 to 55 GW in 5 GW steps. Moreover, we assume that the wind farms do not provide inertial support or governor response [1]. The operational limit equals 49.8 Hz [5], while the minimum RoCoF is set to -0.5 Hz/s. It is worth pointing out that this value differs from the current NG setting (-0.125 Hz/s [5]) and it is in agreement with NG suggestions for future low carbon scenarios [4].

A. Power system model

We make use of the standard linear frequency response model described in [11]; this model considers the presence of generators able to provide primary response only and other units that can supply both services. The model is then extended (see Fig.2) to integrate the demand response support and the reserve contribution provided by further generators.

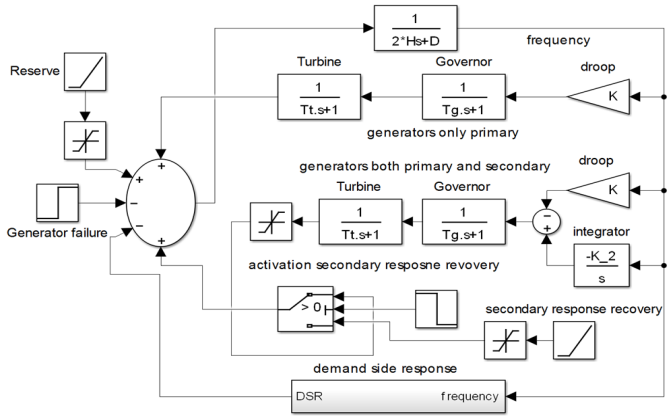

Fig. 2. Linear system frequency response model.

Furthermore it is worth pointing out that the provision of secondary response from generators is maintained for 30 minutes after the failure [9]; afterwards, the power provided reduces linearly (within 15 minutes) and it is replaced by reserve units. The parameters of this model are chosen to qualitatively match the frequency dynamics of the GB system; In particular, the time constants for the governor and the turbine are t_g=1s and t_t=10s. The constant of inertia for the synchronous generators is 4.5s, while the values for the gains are K=3.33 and K_2=0.025. All these parameters are multiplied by the amount of synchronous generation [MW] and divided by the nominal frequency [Hz]; the load damping effect D=1 is in units of MW/Hz as it is multiplied for the total system demand and divided by the nominal frequency. A step function simulates at the time failure $t_{fail} = 10s$ a sudden generator outage; the amplitude of the generation loss is set at 1.8 GW that nowadays represents the maximum expected infeed generation loss for the GB power system [4]. The software used for simulations is Matlab-Simulink [12].

B. Thermostatically controlled loads

In this paper, we consider domestic refrigerators with a built-in freezer compartment. A single thermostat is assumed to govern the freezer temperature evolution; by controlling the freezers' temperature dynamics, the temperatures of the fridge compartment are maintained within safe limits. The parameters of the first order thermal model (4) are taken from [13]; in particular, T_{max}=-14°C, T_{min}=-21°C, T_{on}=-151°C, T_{off}=20°C, τ=5h and P_{on}^{a}=180W. The cluster of TCLs is assumed to count 45 million units [13] so that their population level steady state power consumption is 1.77 GW. The appliance power response with the frequency linear controller is exclusively driven by the frequency deviation and RoCoF evolution. For the second implementation, the reference power profile is designed off line and we take as a reference response profile the blue solid curve in Fig.3. It is worth pointing out that this paper does not aim to determine the optimal shape of the pre-programmed response profile $\Pi_{pp}(t)$; the optimal allocation of frequency response services could be also driven by further technical and economic issues.

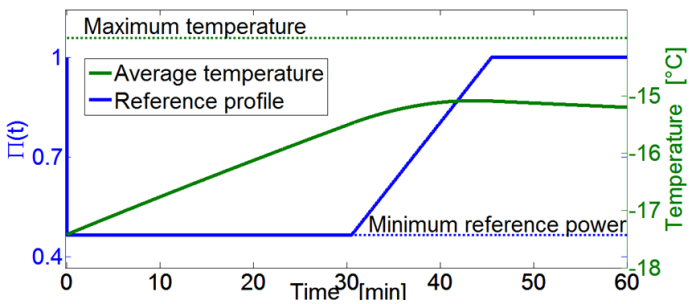

Fig. 3. Reference power profile (blue solid) and associated average temperature evolution (green solid).

The selected reference profile $\Pi_{pp}(t)$ can be followed by TCLs as it respects condition (2) and the associated average temperature evolution does not violate (3). Note that $\Pi_{pp}(t)$ provides maximum support for primary and secondary response: once the frequency drop triggers the demand reduction, it achieves the lower boundary of equation (2) (see (8a) also) in one second. This reduced power level is maintained for 30s (primary response) and for a further 30 minutes (secondary response) [9]. Afterwards, the profile recovers the pre-fault value by means of a positive ramp (15 min). Payback patterns are not considered in the short-term scope of this analysis. However, the stochastic-threshold controller developed in [8] is able to boost the power consumption in order to speed up the recovery of the average temperature. Even in this case the optimal shape and amount of extra-power absorbed would be the outcome of several technical and economic aspects.

V. Results

The effectiveness of the proposed control methods is illustrated in Fig.4 for both the low (4.A) and high (4.B) wind scenarios. Without implementing any demand control scheme, the frequency nadirs (violet) would drop below the security limit, (grey) f_{min}= 49.2 Hz, in most low-demand scenarios,

especially in the high wind case. The results do not significantly improve by making use of versions A (yellow) and C (green) of the frequency linear controller. The former only enables a poor maximum power reduction ($\Pi_{min} < \Pi_{sust}$), whereas the latter contributes mostly in the first instants, dropping to zero the frequency nadir.

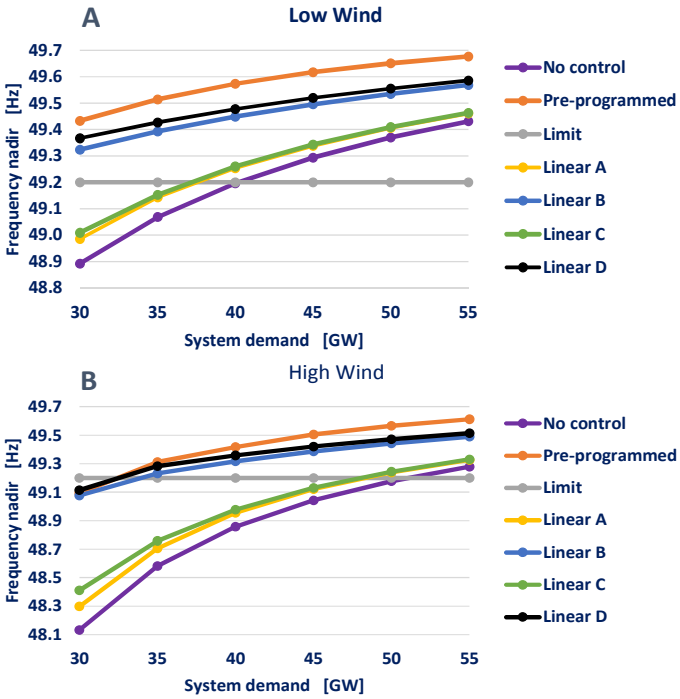

Fig. 4 Frequency nadirs for low wind (A) and high (B) wind penetration.

Conversely, the other linear controllers (case B (blue) and D (black)) and the pre-programmed implementation (orange) provide promising overall results. The associated frequency nadirs are above the security limit, implying that the available wind capacity can be used to supply the system demand without being curtailed. Only in one severe scenario (30 GW demand and 20 GW of wind) are the TCLs unable to contain the frequency drop within the security limit; a deeper penetration of TCLs (e.g. commercial refrigeration) could easily cover this frequency response shortage. We now compare in more detail the three most effective solutions. Version D of the linear controller performs slightly better than version B; the power reduction is quicker as it exploits the RoCoF evolution that is maximum just after the frequency drop. On the other hand, the pre-programmed controller generally results in the smallest frequency deviations. For a few severe cases in the high wind scenario, the performance of the pre-programmed controller and the linear one (case D) are equivalent; frequency falls vary rapidly (high RoCoF), so the initial power reduction with the linear controller resembles the fast reduction of the pre-programmed power profile.

For the remainder of the paper, we will only consider the scenario with 45 GW of demand, 20 GW of wind and the remaining difference supplied by conventional generators. Moreover, we will focus only on the most effective strategies, the pre-programmed controller (red) and Case D of the frequency linear controller (black).

A. *Frequency evolution*

The first instants (50s after the disturbance) of the system frequency evolution is shown in Fig. 5. With both responsive demand strategies, the frequency nadir is held at a significantly higher value than the security limit of 49.2 Hz. The frequency evolution with the linear controller is more damped; this effect is consequence of the positive value of the RoCoF after that frequency achieves the nadir. In fact, during this phase, the second term of (9), $K_2 \cdot \frac{d\Delta f}{dt}$, opposes the first term, thus reducing the TCLs' response. This counteractive behavior terminates quickly as the RoCoF tends to zero after the primary response provision, while the frequency deviation is a negative quantity at all times.

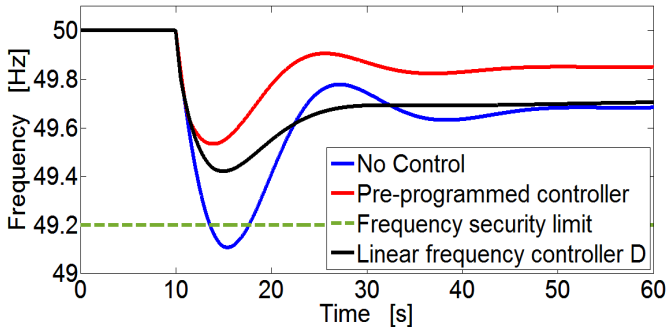

Fig. 5. Initial evolution of system frequency with pre-programmed controller (red), frequency linear controller D (black) and without (blue) demand support.

We now increase the time window to assess the benefits of demand side response also in the secondary response/reserve time interval. The complete frequency evolution until the restoration of the steady state condition is shown in Fig. 6.

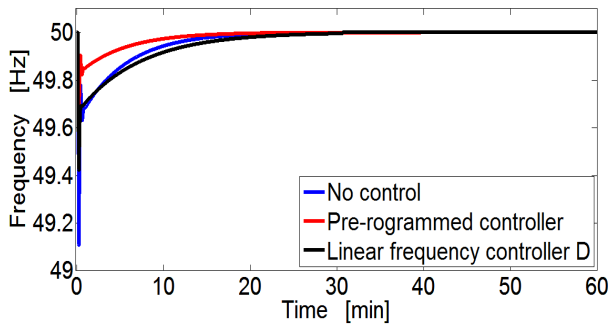

Fig. 6. System Freqeuncy evolution and restoration with pre-programmed controller (red), frequency linear controller D (black) and without (blue) demand support.

The provision of secondary response with the pre-programmed strategy (red) speeds up the restoration of frequency to acceptable values as it maintains the power at the minimum level for 30 min. The linear controller (black) instead imposes a slightly slower frequency recovery compared to the reference case (blue).

B. *Impact of TCL support on response and reserve services provided by conventional generators*

The second case study highlights the ability of demand response to replace frequency services provided by conventional generators operated part-loaded – or by costly

rapid start units. The fast response provided by generators only in charge of primary control drastically reduces (see Fig. 7) as a result of the TCL contribution (red and black curves), compared to the reference case without demand (blue).

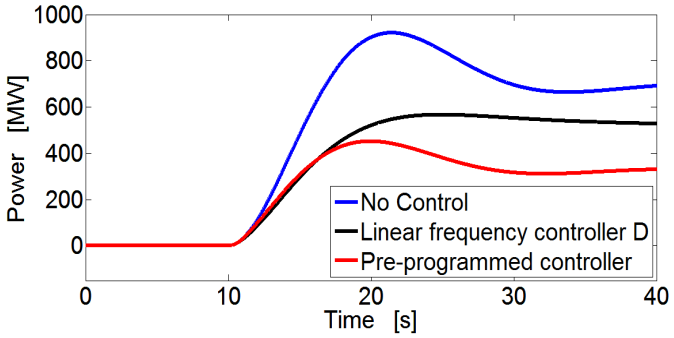

Fig. 7. Power supplied by generators that provide only primary response with pre-programmed controller (red solid), frequency linear controller (black solid) and without demand support (blue solid).

This result has a significant impact on system scheduling; the reduced need for part-loaded generators implies an ability to de-commitment expensive units. The amount of power supplied by generators providing both primary and secondary response (Fig. 8) is significantly reduced only with the pre-programmed controller (dash red) as it sustains the maximum TCL response for 30 min.

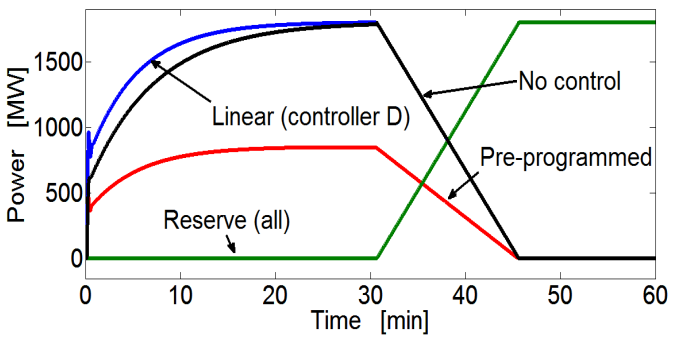

Fig. 8. Power provided by generators in charge of primary and secondary response with pre-programmed controller (red), frequency linear contorller (black) and without demand support (blue). Reserve with and without demand response (green) provided by additional generators.

In addition, the power ramp rate to provide secondary response decreases, facilitating the replacement of fast but costly generators with other units. The frequency linear controller cannot guarantee similar benefits as the TCL participation in the secondary response is poor, despite the significant primary response contribution. The reserve supply after 30 min does not change; however, the generators that now do not provide secondary response anymore can facilitate the fulfillment of the reserve requirement.

CONCLUSIONS

The paper explored the range of frequency response contributions that can be delivered by smart thermal appliances with freely designable response curves, and how they affect the system's frequency performance. Two qualitatively different implementations, the frequency linear controller and the pre-programmed controller, were introduced and analyzed in the context of a low carbon generation scenario. Different variations of the frequency linear controller were tested. The versions of the linear controller with a maximum reduction equal to Π_{sust} and the version that is only sensitive to the RoCoF were found to provide poor frequency support. The variations that are able to reduce the TCLs' consumption up to Π_{min} in response to frequency deviations, either with or without additional sensitivity to the RoCoF, provide substantial support to the system's primary response. The alternative strategy, the pre-programmed controller, has the benefit of providing both substantial support and flexibility. The system operator or demand response aggregator has the freedom to design complex responses in accordance with the technical and commercial requirements of the network. Moreover, this implementation enlarges the range of applications for TCLs; beyond the frequency services, TCLs could perform energy arbitrage using the same pre-programmed control framework. Further work will study the impact of delays in the activation and deployment of the demand reduction, and – for the preprogrammed controller - the impact of mistakes in system identification.

REFERENCES

[1] G.Lalor, M. Mullane and M. O'Malley, "Frequency control and wind turbine technologies," *IEEE Trans. Power Syst.,* vol. 20, n. 4, pp. 1905-1913, Nov 2005 .

[2] PPA Energy, "Rate of Change of Frequency (ROCOF) Review of TSO and Generator Submissions - Final Report Submitted to: Commission for Energy Regulation," May 2013. [Online]. Available: http://www.cer.ie/docs/000260/ppa-rocof-final-report.pdf.

[3] A. Dyśko, I. Abdulhadi, X. Li and C. Booth, "Assessment of Risks Resulting from the Adjustment of ROCOF Based Loss of Mains Protection Settings," June 2013. [Online]. Available: http://www.nationalgrid.com/NR/rdonlyres/D3F18F81-BFE8-4BA1-8B82-CCD6CD0A0A4F/62018/GC0035IndustryConsultationv10.pdf.

[4] National Grid, "Electricity ten year statement," November 2013. [Online]. Available: http://www2.nationalgrid.com/UK/Industry-information/Future-of-Energy/Electricity-ten-year-statement.

[5] National Grid, "GB Security and quality of supply standard," 2012. [Online]. Available: http://www2.nationalgrid.com/UK/Industry-information/Electricity-codes/System-Security-and-Quality-of-Supply.

[6] J. Short, D. Infield, and L.L. Freris, "Stabilization of grid frequency through dynamic demand control," *IEEE Trans. Power Syst,* n. 3, 2007.

[7] D. Angeli and A. Kountouriotis, "A Stochastic Approach to "Dynamic-Demand" Refrigerator Control," *IEEE Transactions on Control Systems Technology,* vol. 20, n. 3, pp. 581 - 592, 2012.

[8] S.Tindemans, V.Trovato and G.Strbac, "Decentralised control of thermostatic loads for flexible demand response," *IEEE Trans. Control Syst. Technol. - in press,* 2014.

[9] National Grid , "Mandatory Frequency Response," [Online]. Available: http://www2.nationalgrid.com/uk/services/balancing-services/frequency-response/mandatory-frequency-response/.

[10] National Grid, "Frequency response technical sub-group report," Nov 2011. [Online]. Available: http://www.nationalgrid.com.uk/.

[11] P. Kundur, Power system stability and control, London: McGraw-Hill, 1994.

[12] Mathworks, "Matlab, the language of thechnical computing," [Online]. Available: http://www.mathworks.co.uk/products/matlab/.

[13] V.Trovato, S. Tindemans and G. Strbac, "Security constrained economic Dispatch with Flexible Thermostatically Controlled Loads," in *5th IEEE PES ISGT Conference*, Istanbul, 2014.