A macroscopic view of demand-side grid regulation through fluid queueing models and $\mathcal{H}_2$ control.

Federico Bliman, Fernando Paganini, *Fellow, IEEE*, and Andres Ferragut

***Abstract*—We consider a demand aggregator handling a large number of deferrable loads through a simple interface, who must align their overall consumption with an operator-provided reference. A macroscopic model of the load aggregate is developed as a two-state stochastic differential equation, with a scalar input controlling deferred service, and where deadlines are enforced. Through linearization the feedback control design is cast within the framework of $\mathcal{H}_2$ control, with an objective that embeds frequency domain characteristics of the tracked reference. The solution yields a distributed implementation with mild communication requirements. We tested it in a simulation environment using real-world grid frequency regulation signals, achieving high performance with the relevant industry metric.**

***Index Terms*—Frequency regulation, demand response, deferrable loads, optimal control.**

I. Introduction

The matching of supply and demand in electric power networks is a longstanding issue, addressed by mechanisms at multiple time-scales: from long-term delivery contracts, through hour-ahead spot markets, to real-time balancing at the scale of seconds. Our focus is on the latter *frequency regulation* ancillary service, which traditionally has been viewed as generation-side control to follow fluctuations in demand, implemented in two stages [1]: *primary* regulation or "droop control" carried out by governors at rotating machines, which maintain equilibrium at the expense of deviations from nominal frequency; and *secondary* regulation or "Automatic Generation Control" (AGC), orchestrated by the System Operator (SO) to restore nominal frequency and inter-area flows.

This traditional picture is being challenged on different fronts [14]. On one hand, penetration of renewable generation of a volatile nature (wind, solar) means that demand is not the only source of fluctuations, thus increasing pressure on this service. In compensation, balancing from battery storage is becoming more available [8], and also in a smarter grid, controlled loads can do their share of regulation. Indeed, due to their relatively small inertia, loads are in a position to be among the fastest responders for load balancing.

These new trends are leading to changes in the operation of the regulation market. For instance, the PJM operator [20] has in recent years diversified its AGC, providing two alternative regulation signals for providers to track: the standard ("RegA") signal aimed at slower responders who can provide a steady-state power deviation, and a dynamic signal with zero mean ("RegD") aimed at storage devices which can provide power quickly but not sustain it over time. Precisely this second type of response can be provided on the demand side, when the loads involved cannot be curtailed but are *deferrable*: this is the situation considered in the present paper.

Of the possible loads to be summoned, we are interested in the scenario of a large number of small (e.g. domestic scale) loads which collectively provide the ancillary service, controlled by a load aggregator entity (e.g. [9], [11]) to follow an operator-provided directive. Related work on this idea is reviewed below. Our primary objective is to regulate the total power consumption in a manner that avoids complexity at the load end, and also avoids maintaining individual load information by the aggregator. Rather, the control should be based on aggregate population information and communicate to loads through a streamlined interface.

The modeling strategy of this paper starts in Section II with a Markov model tracking deferrable and non-deferrable populations, and where control is a scalar variable expressing service deferral. A *macroscopic* model is then presented in Section III, in the form of a stochastic differential equation; this model is suitably linearized for control design purposes.

The design of a feedback controller based on population measurements to track an operator reference is cast in Section IV in the $\mathcal{H}_2$-control framework, which is well-suited to model the relevant bandwidth of fast regulation signals. In Section V we discuss a decentralized implementation of such a control scheme, and provide simulations at the microscopic load level to validate its ability to track real-life regulation signals from PJM. Conclusions are given in Section VI. Preliminary results leading up to this paper were presented in [2], [4].

A. Related Work

The possibility of implementing the frequency regulation ancillary service by controlling aggregates of loads has been investigated recently by many authors. In many cases the proposal involves a microscopic description, where discrete loads are individually identified and some kind of large-scale online optimization is performed to dispatch them [23], [24], [26]. To avoid such load micro-management on the part of the aggregator, a macroscopic approach is required.

An aggregated model was pursued in a series of papers on thermostatically controlled loads (TCLs). [15] uses a state-space model representing populations in different temperature ranges, as a basis for model predictive control. In [12], [13] the thermal storage is given an equivalent battery representation within certain bounds, used to determine the ex ante regulation

The authors are with the Mathematical Analysis in Telecommunications and Energy (MATE) Research Group, Universidad ORT Uruguay, Montevideo, Uruguay (e-mail: {paganini,ferragut}@ort.edu.uy). Research partially supported by ANII-Uruguay.

commitments; the model is not, however, directly applied in the stack priority scheme used for control. An extension of the battery point of view for more generic loads is given in [19], with the observation that load deferability behaves in essence like storage, in particular not allowing for steady-state biases in demand. Bounds are given on the equivalent battery, with some consideration to the randomness of the load process. In our work, we invoke stochastic queueing more comprehensively as a basis for dynamic modeling and control.

Another series of references consider the stochastic control of ON-OFF loads to provide frequency regulation. In [7], [25] a decentralized frequency measurement is the threshold for stochastic ON-OFF transitions, thereby implementing load-side droop control. A proposal for secondary AGC with ON-OFF pool-pump loads is given in [18]; the mathematics covers broad territory, but the essence is a linear state-space model representing occupation probabilities of a Markov model for the pump state, controlled by a parameter from the balancing authority. Our proposal of controlling populations through a scalar deferral action parameter bears some philosophical resemblance, but details appear quite different.

Finally, we mention that another recent stream of literature [17], [27] has looked at load-side frequency regulation coupled with the grid dynamics, proposing new signals by which the SO may better carry out the balancing task. This issue will not be pursued here, we will assume the SO communicates power directives as is current practice [20], and the task of our load aggregate is to track such reference signal.

II. Load deferral in terms of controlled queues

We consider an aggregator entity that manages a large quantity of loads, with a known *mean* consumption power p^*; we assume this quantity has been procured in advance. In real-time, discrete loads materialize as a sequence of requests for service with inherent randomness; a passive system where loads are served immediately upon arrival would result in a stochastic aggregate power consumption, varying around the mean. If service can be partially deferred, then a controlled schedule could remove these variations and, furthermore, possibly arrange for the power consumption to follow a desired reference around the mean, thus providing a balancing service.

We start with a stochastic model of the load process. Assuming stationarity of demand, a natural model is a Poisson process of requests for energy, arriving at rate λ; for simplicity we assume they all have the same nominal power p_0,[1] but may differ in the required energy. This would be the situation of, for instance, a fleet of electric vehicles with homogeneous ratings, but different charging requirements upon arrival.

Let τ_i denote the nominal service time of the i-th load, i.e. the time (energy$/p_0$) it would take to serve it at full nominal power. If load i is deferrable, there is an additional spare time or *laxity* L_i it can tolerate and still meet its deadline. For modeling purposes, we assume that τ_i, L_i are independent, exponential random variables for each load, with respective means τ, L, and also independent of the arrival process.

[1]This assumption simplifies the mathematics. In Section V we remark on its removal.

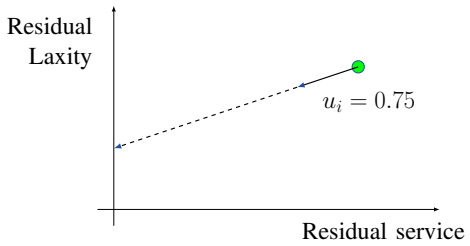

Fig. 1. Service-laxity trajectory under service level u_i

The deferral action is defined by the *service level*, a scalar variable $u_i \in (0, 1]$ specifying the fraction of nominal power at which a load is served.[2] After an interval of time dt, a load served at power $u_i p_0$ will have attained $u_i dt$ of its service time requirement, but also will have consumed $dt - u_i dt = (1 - u_i)dt$ of its spare time. Fig. 1 shows the trajectory in (residual service-time/laxity) space when u_i is held constant.

A trajectory reaching the vertical axis completes service and leaves the system. If the horizontal axis is crossed first, laxity expires and the load misses its deadline.

While applying an individualized service level u_i could provide a very fine grained control of power and deadlines, we wish to avoid this level of micro-management by the aggregator; instead, we will work with simple global commands for the aggregate of loads. We discuss two options below.

A. A single class aggregate control

The simplest choice would be to apply a common service level u for the entire population of loads. Assuming initially u is held constant, the service time is an exponential random variable of mean τ_i/u, and therefore the number of active loads behaves as the following $M/M/\infty$ queue [22], a birth and death process with the following transition rates:

- $n \mapsto n + 1$ with rate λ, reflecting a Poisson arrival.
- $n \mapsto n - 1$ with rate $\frac{nu}{\tau}$, reflecting the completion of one service among n independent $\exp(\frac{u}{\tau})$ random variables.

The preceding Markov chain has a well-known Poisson$\left(\frac{\lambda\tau}{u}\right)$ steady state distribution, which provides useful information about the system in steady state. In particular:

- The mean power consumption is $E[p] = p_0 u E[n] = p_0 \lambda \tau$, independent of u. Deferring service has no effect on mean consumption, it must match the mean demand.
- The variance is $Var(p) = p_0^2 u^2 Var(n) = p_0^2 \lambda \tau u$. Thus deferring service (reducing u) smoothes the power profile.

However we note a limitation of the preceding control strategy: if we apply an indiscriminate service reduction to all loads, some will inevitably miss their deadlines; in particular, for any fixed $u < 1$ there is a positive probability

$$P\left[\frac{\tau_i}{u} > \frac{L_i}{1-u}\right] = \frac{(1-u)/L}{u/\tau + (1-u)/L} \qquad (1)$$

that the trajectory in Fig. 1 crosses the horizontal axis. This issue will be avoided by a slightly more complicated policy.

[2]We assume such modulation in power consumption is possible, which is the case for electric vehicles. If loads must be ON or OFF, a similar aggregate effect could be obtained by serving a fraction of loads at nominal power.

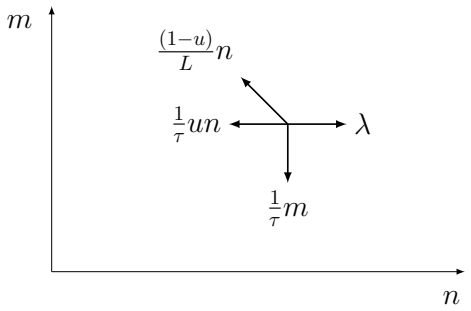

Fig. 2. Two-class Markov state diagram with transition rates

B. Two-class control policy enforcing deadlines

The simplest control strategy that avoids missing deadlines is for a load which runs out of laxity to start consuming at full power. Thus loads will be divided in two classes: $n(t)$ will be the population of loads that at time t still have remaining laxity, and thus are able to obey a command u of deferred service; the remainder of active loads $m(t)$ with expired laxity will be served at full power until their completion.

With these notations, returning to the original stochastic assumptions, the behavior of the population variables over time is a continuous-time Markov chain with state (n, m) and transition rates depicted in Figure 2:

1. $(n, m) \mapsto (n+1, m)$ is a Poisson arrival, at rate $w_1 = \lambda$.
2. $(n, m) \mapsto (n-1, m)$ represents a load from the n queue completing service. The rate $w_2 = \frac{nu}{\tau}$ corresponds to a service level u as in the previous section.
3. $(n, m) \mapsto (n-1, m+1)$ represents the transition between the n and m queues due to expiration of laxity, with rate $w_3 = \frac{n(1-u)}{L}$ (see below).
4. $(n, m) \mapsto (n, m-1)$ represents a departure from the m queue. The transition rate $w_4 = \frac{m}{\tau}$ reflects that service here is at full power.

We expand briefly on the justification of w_2, w_3. Given u, each of the n deferrable loads may complete service after an $\exp(u/\tau)$-distributed time, or run out of laxity after an $\exp((1-u)/L)$-distributed time. So a departure of this queue occurs at the minimum of $2n$ (n of each kind, independent) exponentials, thus having an $\exp(nu/\tau+n(1-u)/L)$ distribution; this justifies the total departure rate w_2+w_3. Multiplying it by the probability (1) that laxity runs out first, yields w_3; the complementary thinning gives w_2.

In [10] the above Markov chain was analyzed for the case of a fixed $u = u^*$; in particular it is shown that it admits a stationary distribution of product form

$$\pi(n, m) = e^{-\rho_n - \rho_m} \frac{\rho_n^n}{n!} \frac{\rho_m^m}{m!}, \quad n, m \in \mathbb{N}; \tag{2}$$

namely, n and m in steady state are independent random variables with Poisson distribution. The parameters are $\rho_n = \lambda/\nu$, $\rho_m = \lambda\tau(1-u^*)/\nu L$, where we introduce the notation

$$\nu := \frac{u^*}{\tau} + \frac{(1-u^*)}{L}. \tag{3}$$

Here, our main interest is in a *controlled* service deferral $u(t)$, as a means to provide frequency regulation service. We state our high-level objective as follows:

Problem 1: Given a power reference signal $p^* + r(t)$ from the system operator, design a control system for the deferral action signal $u(t)$ such that the aggregate power consumption

$$p(t) = p_0 \left(n(t)u(t) + m(t)\right) \tag{4}$$

tracks the reference; i.e. the tracking error $e(t) = p^* + r(t) - p(t)$ has small variance $E[e(t)^2]$.

For control design toward this objective, the Markov chain model is unwieldy. Assuming populations are moderately large, we turn instead to differential equation models that treat load populations as continuous variables.

III. Macroscopic dynamic model

A first model of this nature was pursued in [2] for the one-class control of Section II-A, replacing the Markov chain by the ordinary differential equation (fluid flow) model[3]:

$$\dot{n}(t) = \lambda - \frac{1}{\tau} n(t)u(t). \tag{5}$$

In [2] we showed that this simple model leads to a reference tracking controller with good performance. However due to the requirement of enforcing deadlines, from now on we will focus exclusively on the two-class control of Section II-B.

A. Stochastic Differential Equation Model

In the fluid-flow version of the Markov chain of Figure 2, n and m become continuous variables; an ODE analogous to (5) can be written. However since part of the objective is to control variability, we require a finer model that preserves a macroscopic view of the randomness. This model takes the form of *stochastic* differential equation (SDE), as follows:

$$\begin{aligned} dn = &\left(\lambda - \frac{nu}{\tau} - \frac{n(1-u)}{L}\right) dt \\ &+\sqrt{\lambda}\, dW_1 - \sqrt{nu/\tau}\, dW_2 - \sqrt{n(1-u)/L}\, dW_3, \end{aligned} \tag{6a}$$

$$\begin{aligned} dm = &\left(\frac{n(1-u)}{L} - \frac{m}{\tau}\right) dt \\ &+ \sqrt{n(1-u)/L}\, dW_3 - \sqrt{m/\tau}\, dW_4. \end{aligned} \tag{6b}$$

The model is written in stochastic calculus notation. Each of the four transitions in the Markov chain, with rate w_i is replaced by two terms: the drift $w_i\, dt$ plus the noise $\sqrt{w_i}\, dW_i$, where $W_i(t)$ are independent Wiener processes, $i = 1, 2, 3, 4$. For mathematical justification of the model and its relationship through scaling with the Markov chain, we refer to [16].

The above state-space dynamics has $u(t)$ as control input, and the main output of interest is the aggregate power consumption (4) of both types of loads in the population.

[3]Formally, the solutions to (5) are limits of the scaled stochastic processes $\frac{1}{k} n^{(k)}(t)$ as $k \to \infty$, where $n^{(k)}$ is the Markov chain state under scaled arrival parameter $k\lambda$, and suitably scaled initial condition. For details see [22].

A first step in the analysis is to find the equilibrium of this model for the case of a fixed control $u(t) \equiv u^*$. Setting to zero the drift terms in (6) we have the equilibrium point

$$n^* = \frac{\lambda}{\nu}, \qquad m^* = \frac{\lambda\tau(1-u^*)}{\nu L}, \tag{7}$$

with ν from (3). The equilibrium power of the cluster is then found to be

$$p^* = p_0(n^*u^* + m^*) = \lambda p_0 \tau. \tag{8}$$

Note that p^* does not depend on u^*, and matches the mean exogenous demand for power (load arrival rate times mean energy). Deferrable loads must be served sooner or later, and cannot provide a steady-state deviation in consumed power.

B. Linearized dynamics

The dynamics (6) is a nonlinear, stochastic differential equation. For purposes of control design, we will consider instead a *linearization* of the dynamics around the operating point (u^*, n^*, m^*, p^*), using $\tilde{u} = u - u^*$, $\tilde{n} = n - n^*$, etc. to denote incremental quantities.

We will replace the drift terms, which have equilibrium value 0, by their first-order approximation

$$\lambda - \frac{nu}{\tau} - \frac{n(1-u)}{L} \approx -\underbrace{\left[\frac{u^*}{\tau} + \frac{1-u^*}{L}\right]}_{\nu}\tilde{n} + \left[\frac{n^*}{L} - \frac{n^*}{\tau}\right]\tilde{u};$$

$$\frac{n(1-u)}{L} - \frac{m}{\tau} \approx \frac{1-u^*}{L}\tilde{n} - \frac{1}{\tau}\tilde{m} - \frac{n^*}{L}\tilde{u}.$$

The noise terms $\sqrt{w_i}\, dW_i$ will be replaced by $\sqrt{w_i^*}\, dW_i$; this makes the noise enter the dynamics linearly, and is justified since the equilibrium value $\sqrt{w_i^*} > 0$, so locally around equilibrium, this constant will dominate first and higher order terms. The proposed linearized dynamics are:

$$\overbrace{\begin{bmatrix} d\tilde{n} \\ d\tilde{m} \end{bmatrix}}^{dX} = \overbrace{\begin{bmatrix} -\nu & 0 \\ \frac{1-u^*}{L} & -\frac{1}{\tau} \end{bmatrix}}^{A} \overbrace{\begin{bmatrix} \tilde{n} \\ \tilde{m} \end{bmatrix}}^{X} dt + B_1 dW + \overbrace{\begin{bmatrix} \frac{n^*}{L} - \frac{n^*}{\tau} \\ -\frac{n^*}{L} \end{bmatrix}}^{B_2} \tilde{u} dt \tag{9a}$$

$$\tilde{p} = \overbrace{\begin{bmatrix} p_0 u^* & p_0 \end{bmatrix}}^{C} \begin{bmatrix} \tilde{n} \\ \tilde{m} \end{bmatrix} + \overbrace{p_0 n^*}^{D_2} \tilde{u}, \tag{9b}$$

where $dW = (dW_1, dW_2, dW_3, dW_4)^T$ and

$$B_1 = \begin{bmatrix} \sqrt{\lambda} & -\sqrt{\frac{n^*u^*}{\tau}} & -\sqrt{\frac{n^*(1-u^*)}{L}} & 0 \\ 0 & 0 & \sqrt{\frac{n^*(1-u^*)}{L}} & -\sqrt{\frac{m^*}{\tau}} \end{bmatrix}. \tag{9c}$$

Note that (9b) is the linearization of (4). Clearly, the above dynamics is a local approximation, whose validity hinges on the fact that the state stays close to its equilibrium value, itself dependent on the feedback control to be designed. For this reason, as in any control design based on a linearized model, experimental validations will be required.

Nevertheless, we will gain confidence on the procedure by showing that in open loop (no active control of u), the second moment predictions in steady state of this linearized model model match exactly those of the nonlinear SDE (6), and even those of the original Markov chain.

C. Variance calculation for fixed input

We consider here the case where there is no active control, $\tilde{u} \equiv 0$ (i.e. we define a constant deferral action u^*). We will compute the steady-state covariance matrix of the state in (9a), given (see e.g. [6]) by the solution Q to the Lyapunov equation

$$AQ + QA^T + B_1 B_1^T = 0. \tag{10}$$

Solving this equation for the values for A, B_1 in (9) yields

$$Q = \frac{\lambda}{\nu}\begin{bmatrix} 1 & 0 \\ 0 & \frac{\tau}{L}(1-u^*) \end{bmatrix}. \tag{11}$$

Note that this result is consistent with the Markov model; indeed, the stationary distribution with independent Poisson random variables (2) has the covariance Q above.

In fact a third method is available for Markov chains like ours where the transition rates are affine on the state: closed-form dynamic equations for the moments (mean and covariance) can be written (see [5] in the context of chemical reactions); they also give the same result in steady state.

The resulting variance of the output p is:

$$E\left[(\tilde{p})^2\right] = CQC^T = p^* p_0 \left[1 - \frac{1}{\frac{1}{1-u^*} + \frac{\tau}{Lu^*}}\right]. \tag{12}$$

Note that the choice of u^* affects the variance. Even with no real-time control, there is a variance reduction obtained through fixed deferral, its optimum occurring at

$$u^*_{opt} = \frac{\sqrt{\tau}}{\sqrt{L} + \sqrt{\tau}}. \tag{13}$$

In the remainder of the paper we will add active control of $\tilde{u}(t)$, to track an exogenous reference. In that case we do not claim the linearized model has an exact evaluation of variance; it is a local approximation which must be validated.

D. Plant transfer function

Before proceeding with control design it is instructive to find the (open loop) transfer function of our linear plant, from the control input $\tilde{u}$ to the output $\tilde{p}$, in the Laplace domain. This is found from standard calculations to be

$$T_{up}(s) = C(sI - A)^{-1}B_2 + D_2 = \frac{p_0\lambda}{\nu}\frac{s}{s+\nu}. \tag{14}$$

We make the following observations:

- A transfer function of first order indicates a non-minimal realization; indeed, the state is not fully controllable from the input u; in [3] the appropriate Kalman decomposition is performed to exhibit the uncontrollable state. The full-state model must still be used for control design, however, since the full state is excited by noise, and observable.
- The zero at $s = 0$ shows that we can exert no control in "DC" frequency, i.e. no steady-state balancing power. Again, this also happens with a storage device.
- On the other hand, response is firm at high frequency, indicating that our system is capable of very fast control. The boundary between both regimes is the natural pole ν from (3), which corresponds to a time constant of the order of mean service times and laxities. Our system will be able to provide balancing power when this requirement is faster than these load intrinsic times.

IV. $\mathcal{H}_2$ CONTROL FOR REFERENCE TRACKING

In this section we employ the linear state-space model (9) of the aggregate deferrable load dynamics to design a feedback controller for the input $u(t)$. Our objective, stated in Problem 1, is to offer a balancing service to the grid by having the aggregate power consumption follow a reference signal provided by the system operator (SO).

Specifically, a provider of this ancillary service must commit to varying its power consumption up to a fraction θ of its nominal power p^*, in response to a real-time signal $\rho(t) \in [-1, 1]$ that it receives every few seconds from the SO. Upon receiving this signal the load should ideally become

$$p(t) = p^*(1 + \theta\rho(t)) = p^* + \overbrace{\theta p^* \rho(t)}^{r(t)}. \tag{15}$$

A. Maximum offered regulation

A regulation provider is rewarded by the maximum deviation θp^* it is able to offer, thus our convenience is to make θ as large as possible. The maximum theoretical value is $\theta = 1$, which would imply varying the power in the range $[0, 2p^*]$.

In our system of deferrable loads this value is not achievable, because consumed power must lie within the bounds

$$p_0 m(t) \le p(t) \le p_0[n(t) + m(t)];$$

in particular the lower bound is always positive since we have chosen not to defer the loads $m(t)$ with expired laxity, and the upper bound is constrained by loads currently present. Both bounds are time-varying, but we can get an estimate of the achievable margin by applying the equilibrium values.

In particular, imposing that the committed minimal power $p^*(1-\theta)$ is above $p_0 m^*$ and recalling $p^* = p_0(n^*u^* + m^*)$ leads to the bound

$$\theta \le \frac{n^*u^*}{n^*u^* + m^*} = \frac{Lu^*}{Lu^* + \tau(1-u^*)}.$$

Similarly, the upper bound $p^*(1+\theta) \le p_0[n^* + m^*]$ gives

$$\theta \le \frac{n^*(1-u^*)}{n^*u^* + m^*} = \frac{L(1-u^*)}{Lu^* + \tau(1-u^*)}.$$

The upper bounds on θ are, respectively, increasing and decreasing in u^*, and they become equal in $u^* = \frac{1}{2}$; so this value provides the maximum (symmetric) regulation capability, namely $\theta_{\max} = \frac{L}{L+\tau}$. We will use this choice of u^* in what follows; note that it need not coincide with the value from (13) providing minimal open-loop power variability.

B. Regulation signal characterization

Having decided on the amplitude of reference signals we are offering to track, the next key requirement for a good tracking control design is to characterize their spectral content.

For this purpose we turn to a particular family of real-life regulation signals $\rho(t)$ taken from PJM [21], a regional transmission operator in the US. We performed a spectral density estimation based on these PJM signals using MATLAB's signal identification toolbox. They are found to be band limited, with cutoff frequency $\omega_r \approx 1.65 \times 10^{-2}\, rad/s$,

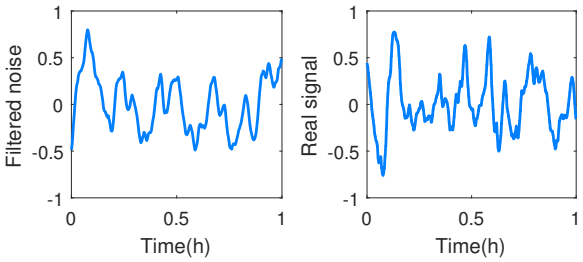

Fig. 3. Artificial regulation signal from filtered white noise, in comparison with a real regulation signal from PJM.

after which they present a roll-off of $40\, db/dec$, indicating a second-order filtering. A closer inspection shows a resonance in the cutoff frequency with a damping factor of $\zeta \approx 0.4$. We therefore approximated the practical signals as generated by white noise through the frequency weighting filter

$$R_\rho(s) = \frac{\kappa_r \omega_r^2}{s^2 + 2\zeta\omega_r s + \omega_r^2}, \tag{16}$$

where $\kappa_r \approx 3$ was chosen to match the mean signal power.

In Fig. 3 we can see a 1- hour simulation of filtered white noise along with a real regulation signal, with a qualitatively similar behavior. Thus, we have confidence that a controller designed to optimally track such filtered white noise will have good performance with the practical regulation signals.

C. $\mathcal{H}_2$-optimal control

We are now ready to tackle the feedback control design problem. A first block diagram is given in Fig. 4, where:

- The plant P is our deferrable load system, with state-space realization (9). Its inputs are our control signal $\tilde{u}(t)$ and the noise, represented in the diagram by $v(t)$ of dimension 4 (in classical notation, the "white noise" vector v stands informally for the "derivative" of W).
- The dynamic model of our synthetic command signal is included, as white noise v_ρ going through the second-order filter $R_\rho(s)$ in (16). The state-variables of this filter are $\rho, \dot{\rho}$; we will assume the controller has access to both. Although in practice only $\rho(t)$ is directly available, it has enough over-sampling that a simple estimate $\dot{\rho}(t) \approx \frac{\rho(t) - \rho(t-T_s)}{T_s}$ has enough accuracy for control purposes. Concretely in the PJM signals with bandwidth $f_r \approx 2.6.10^{-3} Hz$, the sampling rate is $f_s = 0.25 Hz$.
- The controller K is assumed then to have access to the complete (4th order) dynamic state $(\tilde{n}, \tilde{m}, \rho, \dot{\rho})$. With this information it acts on the service level $\tilde{u}(t)$.
- The main performance objective is for the tracking error

$$e(t) = r(t) - \tilde{p}(t) = p^*(1 + \theta\rho(t)) - p(t) \tag{17}$$

to be small, measured by its variance as in Problem 1. For a well-posed problem we must also penalize the control signal $\tilde{u}$, which also helps keep the system in the linear regime. Introducing a tradeoff parameter k_1, our cost function will be

$$J_2 := E[(k_1 e)^2 + (\tilde{u})^2]. \tag{18}$$

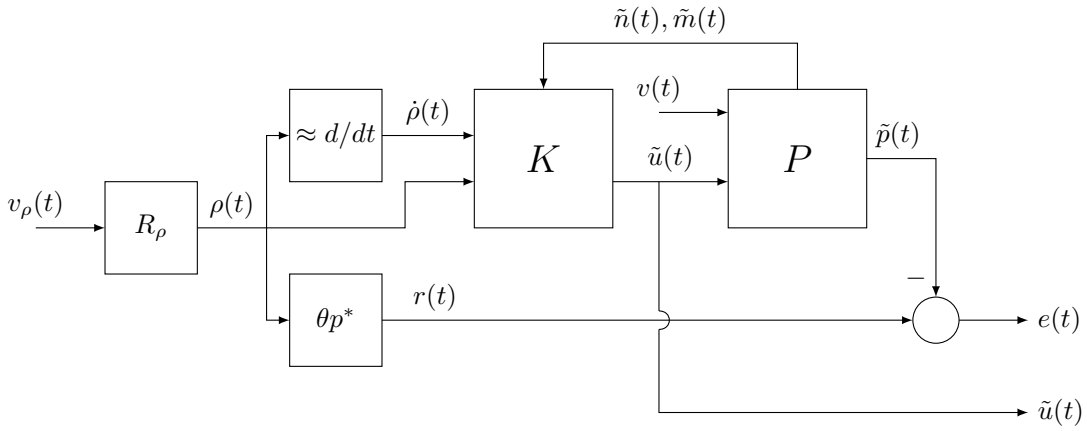

Fig. 4. Block diagram of the control system, including the model for reference signals

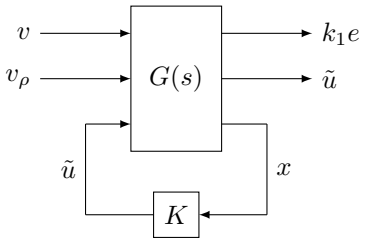

Fig. 5. Setup for feedback control design

By using such a quadratic cost for a system driven by white noise, we are in the standard setup of $\mathcal{H}_2$-optimal control. This is exhibited by rearranging the blocks into the "generalized plant" model of Figure 5; in this setup, the design objective is to find a controller K that internally stabilizes the feedback loop and minimizes the norm

$$\|T(s)\|_{\mathcal{H}_2}^2 = \frac{1}{2\pi}\int_{-\infty}^{\infty} T(j\omega)^* T(j\omega) d\omega$$

of the closed-loop transfer function $T(s)$ between exogenous noises (v, v_ρ) and the vector of penalized variables $(k_1 e, \tilde{u})$. For background on this classical problem we refer to [28].

The state-space realization of the generalized plant G incorporates a realization for $R_\rho(s)$, yielding an augmented state $x = [\tilde{n},\, \tilde{m}, \rho, \dot{\rho}]^T$, and is given as follows:

$$G(s) = \left[\begin{array}{c|cc} \hat{A} & \hat{B}_1 & \hat{B}_2 \\ \hline \hat{C}_1 & 0 & \hat{D}_{12} \\ I & 0 & 0, \end{array}\right], \tag{19}$$

in terms of the block matrices

$$\hat{A} = \begin{bmatrix} A & 0 \\ 0 & A_{22}^r \end{bmatrix},\ \hat{B}_1 = \begin{bmatrix} B_1 & 0 \\ 0 & B_{12}^r \end{bmatrix},\ \hat{B}_2 = \begin{bmatrix} B_2 \\ 0 \end{bmatrix},$$

$$\text{where}\quad A_{22}^r = \begin{bmatrix} 0 & 1 \\ -\omega_r^2 & -2\zeta\omega_r \end{bmatrix},\quad B_{12}^r = \begin{bmatrix} 0 \\ \kappa_r \omega_r^2 \end{bmatrix};$$

$$\hat{C}_1 = \frac{k_1}{k_2}\begin{bmatrix} -p_0 u^* & -p_0 & \theta p^* & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \hat{D}_{12} = \frac{1}{k_2}\begin{bmatrix} -k_1 p_0 n^* \\ 1 \end{bmatrix}.$$

The penalized output corresponds to the cost in (18), except we divided by the constant $k_2 = (1 + (k_1 p_0 n^*)^2)^{\frac{1}{2}}$, which provides the simplifying normalization $\hat{D}_{12}^* \hat{D}_{12} = 1$. Under these conditions, it is shown in [28] that the optimal $\mathcal{H}_2$ controller is a *static* state-feedback law

$$\tilde{u} = -Fx = -(\hat{B}_2^* X + \hat{D}_{12}^* \hat{C}_1)x, \tag{20}$$

where X is the stabilizing solution to the Algebraic Riccati Equation

$$\begin{aligned}(\hat{A}^* - \hat{C}_1^* \hat{D}_{12} \hat{B}_2^*)X + X(\hat{A}^* - \hat{B}_2 \hat{D}_{12}^* \hat{C}_1) \\ -X\hat{B}_2\hat{B}_2^* X + \hat{C}_1^*(I - \hat{D}_{12}\hat{D}_{12}^*)\hat{C}_1 = 0.\end{aligned} \tag{21}$$

The expression for the optimal cost from [28], undoing the normalization is:

$$J_2 = k_2^2\ \mathrm{trace}(\hat{B}_1 X \hat{B}_1^T). \tag{22}$$

The design knob in the problem is the parameter k_1; as it becomes larger, the tracking error takes priority in the objective, resulting in better performance. This has a limit, however: as $\tilde{u}$ is penalized less in relative terms, it will become larger and nonlinear effects come into play, most seriously saturation of $u(t) \in (0, 1]$, which breaks the feedback and deteriorates performance. To keep units normalized, it is convenient to specify k_1 as a multiple of $\frac{1}{p_0 n^*}$.

V. Implementation and performance

We have designed an optimal controller to be run by a demand aggregator for the purpose of tracking a regulation signal received from the SO. We now describe a distributed implementation of this controller, with particular focus on the required communication between aggregator and loads.

The information required by the aggregator is the full-state $(n, m, \rho, \dot{\rho})$; it receives $\rho(t)$ with a high sampling rate from the SO, and as argued above $\dot{\rho}(t)$ is easily estimated. In regard to the population states $n(t)$ and $m(t)$, these can be easily tracked with minimal communication with the loads: the aggregator only needs to be notified when each load arrives, runs out of laxity or leaves the system. The signal $u(t)$ can be broadcast periodically by the aggregator to loads, for instance upon reception of a new sample of the regulation directive.

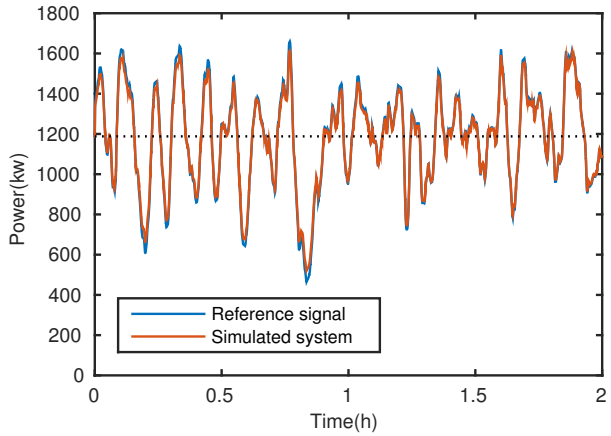

Fig. 6. Tracking a regulation signal with deferrable, tunable power loads.

Loads can keep track of their own laxity and switch to nominal power upon its expiration, so this part of the actuation is fully decentralized, and guarantees deadlines are met. In terms of applying the service level $u(t)$ (fraction of nominal power) to the $n(t)$ loads with remaining laxity, the simplest situation is for loads (e.g. electric vehicles) with tunable power, which can follow $u(t)p_0$ almost instantaneously. We focus on this case first, below we discuss other situations.

A. Simulation results

To validate our approach we employ a detailed simulation of the deferrable load system at the microscopic level. This is a discrete event simulator, in which fresh loads arrive at random times, each one with a random service time and laxity. The reference input is a regulation test signal, preloaded from PJM [21]. At each event, the simulator updates the population variables, and the deferral signal u in (20). The control signal is used to determine the processing rate $p_0 u$ for loads, from where the departure or laxity expiration events are endogenously calculated. We record state, control and power output variables. This setup is thus reproducing, as closely a possible, a real experimental situation.

Our first result shows that the method works extremely well when the number of active loads is in the order of 300, a situation that could apply to an EV garage. The parameters $\tau = 3$ h, $L = 6$ h, $p_0 = 6.6$ kW are chosen with this case in mind. Taking an arrival rate of $\lambda = 1$ load/min and $u^* = 0.5$, the equilibrium point is $n^* = 240$, $m^* = 60$, $p^* = 1188KW$. The regulation amplitude offered is $\theta_{max} = L/(L+\tau) = 0.66$. We chose a value $k_1 = \frac{3}{p_0 n^*}$ for the tradeoff parameter.

We ran a 12 h simulation, two hours of which are shown in Fig. 6. We show the PJM target value $p^*(1 + \theta\rho(t))$ together with the power output $p(t)$ of our $\mathcal{H}_2$ system, which qualitatively achieves very close tracking.

We computed the empirical RMS values of the tracking error e from (17) and the command signal u, and the resulting empirical estimate for the cost, $\hat{J}_2$. We obtain

$$\|e\|_{RMS} = 17.4kW; \qquad \|\tilde{u}\|_{RMS} = 0.1661;$$
$$\hat{J}_2 = k_1^2 \|e\|_{RMS}^2 + \|\tilde{u}\|_{RMS}^2 = 0.0287.$$

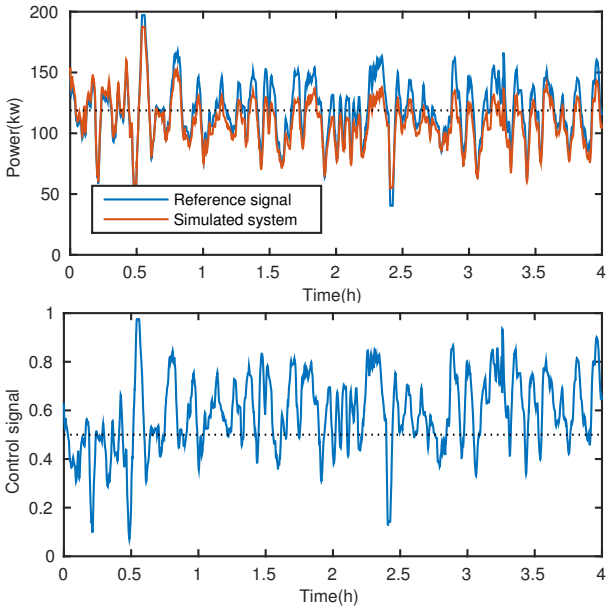

Fig. 7. System operation with a low load population.

A first point is model validation; we compare the cost with the theoretical value from (22), which gives $J_2 = 0.0304$. This close match (6% error) indicates that in this regime, the linearized model is approximating well the system dynamics.

A second comment is that the RMS tracking error is only 1.5 % of the average power, quantifying the excellent regulation service provided. For additional evaluation we computed a performance score used by PJM to rank regulation resources [20]. This score is calculated comparing the reference signal with the actual response form the system and is the average of three components: correlation, delay and precision; all of them measured in a scale from 0 to 1. A score of 0.75 is required for participation in the market, and values above 0.9 are considered excellent. The score corresponding to our control system was found to be 0.978, indicating a highly satisfactory performance.

To test the limits of our procedure, we try operating the system with an order of magnitude fewer loads, reducing the arrival rate to $\lambda = 0.1$ load/min, with the other parameters unchanged. Our mean load counts are now $n^* = 24, m^* = 6$. Figure 7 plots 4 hours of simulation with the same test signal, showing the power tracking and also the control signal u.

A degradation of tracking performance is easily observed, the RMS error is now 10% of the reference signal. The variation of $u(t)$ increases ($\|u\|_{RMS} = 0.2$, 40% of u^*); from the figure we see that it oscillates fully within its range, at certain times hitting its limits. In this scenario it is expected that nonlinearities will come into play: indeed the comparison of the empirical and theoretical $\mathcal{H}_2$ cost now gives a greater discrepancy:

$$\hat{J}_2 = 0.083, \qquad J_2 = 0.1162;$$

the linearized model is overestimating variances by 40%. Still, even in this taxing regime the PJM score is at 0.89, within acceptable limits for frequency regulation service.

B. ON/OFF and heterogeneous loads

We briefly comment here on the application of the same control system to other scenarios of interest. Due to space limitations, we refer to the thesis [3] for extensive details.

A first extension is to consider loads which are not tunable in power, but rather can only be ON at power p_0 or OFF, always with laxity in their service time. For instance water heaters, pool pumps, etc. may fall in this category. In that case we can emulate the power consumption of the aggregate through randomization: upon receipt of $u(t) \in [0,1]$ a load with laxity turns itself ON with this probability; the law of large numbers implies that the aggregate consumption will be close to $u(t)n(t)p_0$. In fact, to accumulate a balancing power of sufficient entity with these smaller loads, an aggregator *must* handle large numbers. Simulations reported in [3], [4] verify the very close reference tracking achieved by this method.

The randomization approach can also provide another extension, required for such broader load categories: allowing for *heterogeneous* nominal power. To handle this diversity with a macroscopic view, we design our controller for the *mean* rating p_0, and rely again on the law of large numbers to guarantee that switching ON independently with probability u loads of the deferrable class, their aggregate consumption approximates unp_0; this was also tested successfully in [3].

One limitation of randomization is that if ON/OFF decisions are made at every step, a large number of service interruptions results, which may not be acceptable. The challenge is to mitigate this issue without a centralized management of individual loads by the aggregator. In [3] we developed and tested a heuristic method to enforce a limit $\#_i^{maxint}$ on the number of interruptions, in a decentralized way: every time a load turns ON, it stays so for a minimum time $T_i = \tau_i/(\#_{maxint} + 1)$. Under this behavior, the aggregator cannot count on its current value of u being instantaneously in force; what it can do is track the real fraction $u_r := \frac{n_{ON}}{n}$ of loads which are ON, and issue to responsive loads a directive meant to align u_r with u. We remit details to [3]; simulations there show that performance degrades severely if no interruptions are allowed, but with an average of the order of 2 interruptions per load, a very high PJM score can be achieved in this way.

VI. Conclusions

We have shown a methodology for an aggregator of deferrable loads to provide frequency regulation to the electric grid. A simple, scalar directive from the aggregator specifies the level of service deferral, and is controlled in feedback to track a directive from the SO.

The design of this feedback is reduced to the setting of $\mathcal{H}_2$ control, through a macroscopic dynamic model in the form of a stochastic differential equation, suitably linearized. Simulations validate the excellent performance achievable by this system for a moderate scale of aggregation.

The control is directly applicable to loads with tunable power, such as EVs. For ON/OFF loads, the main implementation challenge is to handle constraints on service interruptions. A heuristic approach was outlined in [3], a more formal treatment is a topic of future research.

References

[1] A. R. Bergen and V. Vittal, *Power systems analysis*. Prentice Hall, 2000.

[2] F. Bliman, A. Ferragut, and F. Paganini, "Controlling aggregates of deferrable loads for power system regulation," in *Proc. of the American Control Conference*, 2015.

[3] F. Bliman, "Frequency regulation in electric power systems using deferrable loads," Master's thesis, Universidad ORT Uruguay, 2016.

[4] F. Bliman, F. Paganini, and A. Ferragut, "Grid frequency regulation with deferrable loads-an H2-optimal control approach." in *IEEE Conference on Control Applications (CCA)*, 2016, pp. 273–278.

[5] C. Briat and M. Khammash, "Computer control of gene expression: Robust setpoint tracking of protein mean and variance using integral feedback," in *51st IEEE Conf. on Decision and Control (CDC)*, 2012.

[6] J. Burl, *Linear Optimal Control: $\mathcal{H}_2$ and $\mathcal{H}_\infty$ methods*. Menlo Park, CA: Addison Wesley, 1999.

[7] M. Donnelly, D. Harvey, R. Munson, and D. Trudnowski, "Frequency and stability control using decentralized intelligent loads: Benefits and pitfalls," in *IEEE PES General Meeting*, 2010, pp. 1–6.

[8] B. Dunn, H. Kamath, and J.-M. Tarascon, "Electrical energy storage for the grid: a battery of choices," *Science*, vol. 334, no. 6058, pp. 928–935, 2011.

[9] EnergyPool, http://www.energy-pool.eu/en/home/.

[10] A. Ferragut and F. Paganini, "Queueing analysis of service deferrals for load management in power systems," in *53rd Annual Allerton Conference*, 2015, pp. 91–98.

[11] GoodEnergy, http://www.goodenergy.com.

[12] H. Hao, B. M. Sanandaji, K. Poolla, and T. L. Vincent, "Frequency Regulation from Flexible Loads: Potential, Economics, and Implementation," in *Proc. of the American Control Conference*, 2014.

[13] ——, "Aggregate flexibility of thermostatically controlled loads," *IEEE Transactions on Power Systems*, vol. 30, no. 1, pp. 189–198, 2015.

[14] A. Ipakchi and F. Albuyeh, "Grid of the Future," *IEEE power & energy magazine*, vol. 7, pp. 52–62, 2009.

[15] S. Koch, J. Mathieu, and D. Callaway, "Modeling and Control of Aggregated Heterogeneous Thermostatically Controlled Loads for Ancillary Services," in *Power Systems Computation Conference*, 2011.

[16] T. G. Kurtz, "Strong approximation theorems for density dependent Markov chains," *Stochastic Processes and their Applications*, vol. 6, no. 3, pp. 223–240, 1978.

[17] E. Mallada, C. Zhao, and S. Low, "Optimal load-side control for frequency regulation in smart grids," in *52nd Annual Allerton Conference*, 2014, pp. 731–738.

[18] S. P. Meyn, P. Barooah, A. Bušić, Y. Chen, and J. Ehren, "Ancillary service to the grid using intelligent deferrable loads," *IEEE Transactions on Automatic Control*, vol. 60, no. 11, pp. 2847–2862, 2015.

[19] A. Nayyar, J. Taylor, A. Subramanian, K. Poolla, and P. Varaiya, "Aggregate Flexibility of a Collection of Loads," in *Proc. of the 52nd IEEE Conference on Decision and Control*, 2013.

[20] PJM, "PJM Manual 12: Balancing Operations," http://www.pjm.com/~/media/documents/manuals/m12.ashx.

[21] ——, "Regulation Self Test Signals," http://www.pjm.com/markets-and-operations/ancillary-services.aspx.

[22] P. Robert, *Stochastic networks and queues*. Berlin: Springer-Verlag, 2003.

[23] A. Subramanian, M. Garcia, D. Callaway, K. Poolla, and P. Varaiya, "Real-Time Scheduling of Distributed Resources," *IEEE Transactions on Smart Grid*, vol. 4, pp. 2122–2130, 2013.

[24] K. Trangbæk, M. Petersen, J. Bendtsen, and J. Stoustrup, "Exact power constraints in smart grid control," in *50th IEEE Conference on Decision and Control*, 2011, pp. 6907–6912.

[25] T. L. Vincent, K. Poolla, S. Mohagheghi, and E. Bitar, "Stability guarantees for primary frequency control with randomized flexible loads," in *American Control Conference (ACC)*, 2016, pp. 2328–2333.

[26] H. Zarkoob, S. Keshav, and C. Rosenberg, "Optimal contracts for providing load-side frequency regulation service using fleets of electric vehicles," *Journal of Power Sources*, vol. 241, pp. 94–111, 2013.

[27] C. Zhao, U. Topcu, N. Li, and S. Low, "Design and stability of load-side primary frequency control in power systems," *IEEE Transactions on Automatic Control*, vol. 59, no. 5, pp. 1177–1189, 2014.

[28] K. Zhou, J. Doyle, and K. Glover, *Robust and Optimal Control*. Prentice Hall, 1996.

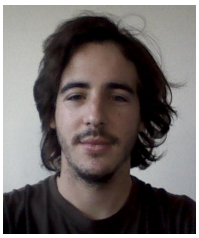

Federico Bliman received his Electrical Engineering degree from Universidad de la República, Montevideo, Uruguay, in 2011, and his Master degree in Electrical Engineering from Universidad ORT, Uruguay, in 2016. From 2014 to 2016 he was part of the Mathematical Analysis in Telecommunications and Energy (MATE) research group where he worked on his thesis. Now he moved to the countryside where he is doing field investigation on sustainable living.

Fernando Paganini (M'90–SM'05–F'14) received his Electrical Engineering and Mathematics degrees from Universidad de la República, Montevideo, Uruguay, in 1990, and his M.S. and PhD degrees in Electrical Engineering from the California Institute of Technology, Pasadena, in 1992 and 1996 respectively. His PhD thesis received the 1996 Wilts Prize and the 1996 Clauser Prize at Caltech. From 1996 to 1997 he was a postdoctoral associate at MIT. Between 1997 and 2005 he was on the faculty the Electrical Engineering Department at UCLA, reaching the rank of Associate Professor. Since 2005 he is Professor of Electrical and Telecommunications Engineering at Universidad ORT Uruguay. Dr. Paganini has received the 1995 O. Hugo Schuck Best Paper Award, the 1999 NSF CAREER Award, the 1999 Packard Fellowship, the 2004 George S. Axelby Best Paper Award, and the 2010 Elsevier Scopus Prize. He is a member of the Uruguayan National Academy of Sciences, and the Latin American Academy of Sciences. He is a Fellow of the IEEE. His research interests are control and networks.

Andrés Ferragut obtained his PhD degree in Electrical Engineering (Telecommunications) from Universidad de la República, Uruguay (2011). He previously held teaching and research positions at the Mathematics and Electrical Engineering Depts. at Univ. de la República (2000 to 2009). He also spent two years (2004-2006) as a research intern at Télécom Paris, France. Since 2011 he is Associate Professor at Universidad ORT Uruguay, working in the Mathematical Analysis in Telecommunications and Energy (MATE) research group. In 2012, Dr. Ferragut was awarded the annual Uruguayan Engineering Academy prize for the best doctoral thesis in Electrical Engineering. Since 2017, he is also Editor for IEEE/ACM Transactions on Networking. His research interests are stochastic processes and queueing theory applied to the mathematical modeling of networks.