

Satisfiability of Elastic Demand in the Smart Grid

Jean-Yves Le Boudec
EPFL

Lausanne, Switzerland

Email: jean-yves.leboudec@epfl.ch

Dan-Cristian Tomozei
EPFL

Lausanne, Switzerland

Email: dan-cristian.tomozei@epfl.ch

Abstract—We study a stochastic model of electricity production and consumption where appliances are adaptive and adjust their consumption to the available production, by delaying their demand and possibly using batteries. The model incorporates production volatility due to renewables, ramp-up and ramp-down time, uncertainty about actual demand versus planned production, delayed and evaporated demand due to adaptation to insufficient supply. We study whether threshold policies stabilize the system. The proofs use Markov chain theory on general state space.

Keywords-Dynamical Systems, Smart Grids, Elastic Demand, Macroscopic Model, Stability

I. INTRODUCTION

Recent results on modelling the future electricity markets [11] suggest that they may lead to highly undesirable equilibria for consumers, producers or both. A central reason for such an outcome might be the combination of volatility in supply and demand, the delays required for any unplanned capacity increase, and the inflexibility of demand which leads to high dis-utility costs of blackouts. Further, the use of renewable energy sources such as wind and solar increases volatility and worsens these effects [8].

Flexible load is advocated in [6] as a mechanism to reduce ramp-up requirements and adapt to the volatility of electricity supply that is typical of renewable sources. A deployments report [4] shows the feasibility of delaying air conditioners using signals from the distributor. Adaptive appliances combined with simple, distributed adaptation algorithms were advocated in [5], [7]; they are assumed to reduce, or delay, their demand when the grid is not able to satisfy them. Some examples might be: e-cars, which may have some flexibility regarding the time and the rate at which their batteries can be loaded; heating systems or air conditioners, which can delay their demand if instructed to; hybrid appliances which use alternative sources in replacement for the energy that the grid cannot supply. If the alternative energy source is a battery, then it will need to be replenished at a later point in time, which will eventually lead to later demand.

The presence of adaptive appliances may help address the volatility of renewable energy supply, however, backlogged demand is likely to be merely delayed, rather than canceled; this introduces a feedback loop into the global system of

consumers and producers. Potentially, one might increase the backlogged demand to a point where future demand becomes excessive. In other words, one key question is whether it is possible to stabilize the system. This is the question we address in this paper.

To address this fundamental question, we consider a macroscopic model, inspired by the model in [8]. We assume that electricity supply follows a two-step allocation process: first, in a forecast step (*day ahead market*) demand and renewable supply are forecast, and the total supply is planned; second, (*real time market*) the actual, volatile demand and renewable supply are matched as possible. We assume that the rate at which supply can be varied in the real time step is subject to ramp-up and ramp-down constraints. Indeed, it is shown in [3] that it is an essential feature of the real time market. We modify the model in [8] and assume that the whole demand is adaptive. While this is clearly an exaggerate assumption, we do it for simplicity and as a first step, leaving the combination of adaptive and non-adaptive demand to a later research. We are interested in simple, distributed algorithms, as suggested in [5], therefore, we assume that the suppliers cannot directly observe the backlogged demand; in contrast, they see only the effective instantaneous demand; at any point in time where the supply cannot match the effective demand, the backlogged demand increases.

Our model is macroscopic, so we do not model in detail the mechanism by which appliances adapt to the available capacity; several possible directions for achieving this are described in [5]. However we do consider two essential parameters of the adaptation process. First, the *delay* $1/\lambda$ is the average delay after which frustrated demand is expressed again. Second, the *evaporation* μ is the fraction of backlogged demand that disappears and will not be resubmitted per time unit. The inverse delay λ is clearly positive; in contrast, as discussed in Section II, it is reasonable to assume that some adaptive appliances naturally lead to a positive evaporation (this is the case for a simple model of heating systems), but it is not excluded that inefficiencies in some appliances lead to negative evaporation.

Within these modelling assumptions, the electricity suppliers are confronted with a scheduling issue: how much capacity should be bought in the real time market to match

the adaptive demand. The effect of adaptation is to increase the latent demand, due to backlogged demand returning into the system. This is the mechanism by which the system might become unstable. We consider a threshold based mechanism as in [8]. It consists in targeting some fixed supply reserve at any point in time; the target reserve might not be met, due to volatility of renewable supply and of demand, and due to the ramp-up and ramp-down constraints.

Our contribution is to show that if evaporation is positive, then any such threshold policy does stabilize the system. In contrast, if evaporation is negative, then there exists no threshold policy which stabilizes the system. The case where evaporation is exactly equal to 0 remains unsolved.

Thus, our results suggest that evaporation plays a central role. Simple adaptation mechanisms as described in this paper might work if evaporation is positive (as one may perhaps generally expect), but will not work if evaporation is negative, i.e. if the fact that demand is backlogged implies that a higher fraction of demand returns into the system. This suggests that future research be done in order to gain a deeper understanding of evaporation, whether it can truly be assumed to be positive, and if not, how to control it.

We use discrete time, for tractability. We use the theory of Markov chains on general state spaces in [9]. In Section II we describe the assumptions and the model, and relate our model to prior work. In Section III we study the stability of the system under threshold policies. We conclude in Section IV.

II. MODEL AND ASSUMPTIONS

A. Assumptions and Notation

We use a discrete model, where $t \in \mathbb{N}$ represents the time elapsed since the beginning of this day. The time unit represents the time scale at which scheduling decisions are done, and is of the order of the second.

The supply is made of two parts: the planned supply $G^f(t)$, forecast in the day-ahead market, and the actual supply $G^a(t)$, which may differ, due to two causes. First, the forecasted supply may not be met, due to fluctuations, for example in wind and sunshine. Second, the suppliers attempt to match the demand by adding (or subtracting) some supply, bought in the real time market. We assume that this latter term is limited by ramp-up and ramp-down constraints. We model the actual supply as

$$G^a(t) = G(t-1) + G^f(t) + M(t), \quad (1)$$

where $M(t)$ is the random deviation from the planned supply due to renewables, $G(t-1)$ is the supply decision in the real time market. We view $G^a(t)$ as deterministic and given(it was computed yesterday in the day-ahead market), $M(t)$ as an exogenous stochastic process, imposed by nature, and $G(t-1)$ as a control variable.

We call $D^a(t)$ the “natural” demand. It is the total electricity demand that would exist if the supply would

be sufficient. In addition, there is, at every time t , the *backlogged demand* $B(t)$, which results from adaptation: $B(t)$ is the demand that is expressed at time t due to a previous demand being backlogged. The total effective demand, or expressed demand, is

$$E^a(t) = D^a(t) + B(t). \quad (2)$$

We model the effect of demand adaptation as follows.

$$B(t) = \lambda Z(t), \quad (3)$$

$$Z(t+1) = Z(t) - B(t) - \mu Z(t) + F(t), \quad (4)$$

$$F(t) = [E^a(t) - G^a(t)]^+. \quad (5)$$

We used the convenient notation $(a)^+ := \max(0, a)$.

In the above equations, $F(t)$ is the *frustrated* demand, i.e. the denied satisfaction at time t . Eq. (5) expresses that, through adaptation, the demand that is served is equal to the minimum of the actual demand and the supply. The variable $Z(t)$ is the *latent backlogged* demand; it is the demand that was delayed, and might later be expressed. It is incremented by the frustrated demand.

The frustrated demand is expressed with some delay; we model this with Eq. (3), where λ^{-1} is the average delay, in time slots.

The evolution of latent backlogged demand is expressed by Eq. (4). The expressed demand $B(t)$ is removed from the backlog (some of it may return to the backlog, by means of Eq. (5) at some later time). The remaining backlog may evaporate at a rate μ , which captures the effect on total demand of delaying some demand. Delaying a demand may indeed result in a decreased backlogged demand, in which case the evaporation factor is positive. For example this occurs if we delay heating in a building with a heating system that has a constant energy efficiency; such a heating system will request more energy in the future, but the integral of the energy consumed over time is less whenever some heating requests are delayed. In this case, positive evaporation comes at the expense of a (hopefully slight) decrease in comfort (measured by room temperature). In other cases, though, we may not exclude that evaporation be negative. This may occur for example with heat pumps [1].

As in [8], we assume that the natural demand can be forecast with some error, so that

$$D^a(t) = D(t) + D^f(t), \quad (6)$$

where the forecasted demand $D^f(t)$ is deterministic and $D(t)$, the deviation from the forecast, is modelled as an exogenous stochastic process. We assume that the day ahead forecast is done with some fixed safety margin r_0 , so that

$$G^f(t) = D^f(t) + r_0. \quad (7)$$

B. The Stochastic Model

We model the fluctuations in demand $D(t)$ and renewable supply $M(t)$ as stochastic processes such that their difference $M(t) - D(t)$ is an ARIMA(0, 1, 0) Gaussian process, i.e.

$$(M(t+1) - D(t+1)) - (M(t) - D(t)) = N(t+1) \quad (8)$$

where $N(t)$ is white Gaussian noise, with variance σ^2 . This is the discrete time equivalent of Brownian motion, as in [8].

Let $R(t)$ be the reserve, i.e. the difference between the actual production and the expressed demand, defined by

$$R(t) = G^a(t) - E^a(t) = G(t-1) - \lambda Z(t) + r_0 + M(t) - D(t), \quad (9)$$

and let $H(t)$ be the increment in supply bought on the real time market, i.e.

$$H(t) = G(t) - G(t-1). \quad (10)$$

Putting all the above equations together, we obtain the system equations:

$$\begin{aligned} R(t+1) &= R(t) + H(t) + N(t+1) - \lambda(Z(t+1) - Z(t)), \\ Z(t+1) &= (1 - \lambda - \mu)Z(t) + \mathbb{1}_{\{R(t) < 0\}} |R(t)|, \end{aligned}$$

Thus we can describe our system by a two-dimensional stochastic process $X(t) = (R(t), Z(t))$, with $t \in \mathbb{N}$.

The sequence $H(t)$ is the control sequence. It is constrained by the *ramp-up and ramp-down constraints*:

$$-\xi \leq H(t) \leq \zeta, \quad (11)$$

where $\xi > 0$ and $\zeta > 0$ are some positive constants.

We assume a simple, threshold based control, which attempts to make the reserve equal to some threshold value $r^* > \zeta$; therefore

$$H(t) = \max(-\xi, \min(\zeta, r^* - R(t))). \quad (12)$$

In summary, we have as model the stochastic sequence $X = (X(t))_{t \in \mathbb{N}}$ defined by

$$\begin{aligned} R(t+1) &= R(t) - \lambda \mathbb{1}_{\{R(t) < 0\}} |R(t)| + \lambda(\lambda + \mu)Z(t) \\ &\quad + (\zeta \wedge (r^* - R(t))) \vee (-\xi) + N(t+1), \end{aligned} \quad (13)$$

$$Z(t+1) = \mathbb{1}_{\{R(t) < 0\}} |R(t)| + (1 - \lambda - \mu)Z(t), \quad (14)$$

where N is an iid white Gaussian noise sequence of variance σ^2 . Note that X is a Markov chain on the state space $\mathbf{S} = \mathbb{R} \times \mathbb{R}^+$.

C. Related Models

Let us now discuss some similar models which have been considered in the literature.

In [9], the authors consider the so-called Linear State Space model (LSS), which introduces an n -dimensional stochastic process $X = \{X_k\}_k$, with $X_k \in \mathbb{R}^n$. For matrices $F \in \mathbb{R}^{n \times n}$, $G \in \mathbb{R}^{n \times p}$, and for a sequence of i.i.d. random

variables of finite mean taking values in \mathbb{R}^p , the process evolves as

$$X_k = FX_{k-1} + GW_k, \quad k \geq 1. \quad (15)$$

Our model (13)-(14) is in fact a superposition of three such LSS models, depending on the current state of the Markov chain. The challenge of showing that our model is stable comes from the fact that in the part of the state space in which $R(t) < 0$, the corresponding LSS does not satisfy the stability condition (LSS5) of [9] (which requires that the eigenvalues of F be contained in the open unit disk of \mathbb{C}).

In [10] a slightly different model for capturing elasticity of demand is proposed. More specifically, the authors consider a scenario in which a deterministically *bounded* amount of demand arrives at each time step, while the supplier decides whether to buy an additional amount of energy from an external source at a certain cost. Unsatisfied demand is backlogged. A threshold policy is analyzed and found to be stable (the size of the backlog is found to be deterministically bounded) and optimal. Pricing decisions are also explored. The main differences with the present work are the following:

- Additional parameters which model delay and loss of backlogged demand (i.e. λ and μ) enrich our model's expressivity.
- We consider potentially unbounded demand, modeled as a 0-mean Gaussian random variable, which makes proving stability more challenging.
- No results on pricing and cost-optimality are included in the present work.

The continuous time model used in [2] seeks to capture the presence of two types of energy sources, primary and ancillary, the latter being less desirable (i.e. more costly) than the former, both subject to (different) ramp-up constraints. A threshold policy is again discussed in the context of rigid demand, which is simply dropped if not enough energy is available. The analyzed Markov chain is a two-dimensional process having on the first coordinate the quantity of energy used from the ancillary source in order to satisfy as much demand as possible, and on the second coordinate the reserve (i.e. energy surplus).

III. SYSTEM STABILITY UNDER A THRESHOLD POLICY

For presentation ease, we consider like in [8] the case $\xi = \infty$. In other words, we relax the ramp-down constraint, i.e. the surplus energy can be disposed of easily. In the extended version of the paper [1] we show that a finite value of ξ does not alter the results. Thus, (11) becomes:

$$H(t) \leq \zeta, \quad (16)$$

and the threshold policy (12) writes as

$$H(t) = \min(\zeta, r^* - R(t)). \quad (17)$$

Let us study how the system stability of (13)-(14) depends on the evaporation parameter μ .

Define the following three domains:

$$\begin{aligned} D_1 &= (-\infty, 0) \times \mathbb{R}_+, \\ D_2 &= [0, r^* - \zeta) \times \mathbb{R}_+, \\ D_3 &= [r^* - \zeta, +\infty) \times \mathbb{R}_+. \end{aligned}$$

Then, denoting

$$\begin{aligned} X(t) &= \begin{bmatrix} R(t) \\ Z(t) \end{bmatrix}, N_0(t) = \begin{bmatrix} N(t) \\ 0 \end{bmatrix}, \zeta_0 = \begin{bmatrix} \zeta \\ 0 \end{bmatrix}, \\ r_0^* &= \begin{bmatrix} r^* \\ 0 \end{bmatrix}, A_1 = \begin{bmatrix} 1 + \lambda & \lambda(\lambda + \mu) \\ -1 & 1 - \lambda - \mu \end{bmatrix}, \\ A_2 &= \begin{bmatrix} 1 & \lambda(\lambda + \mu) \\ 0 & 1 - \lambda - \mu \end{bmatrix}, A_3 = \begin{bmatrix} 0 & \lambda(\lambda + \mu) \\ 0 & 1 - \lambda - \mu \end{bmatrix}, \end{aligned}$$

the process (13)-(14) rewrites in matrix form:

$$X(t+1) = N_0(t+1) + \begin{cases} A_1 X(t) + \zeta_0, & X(t) \in D_1, \\ A_2 X(t) + \zeta_0, & X(t) \in D_2, \\ A_3 X(t) + r_0^*, & X(t) \in D_3. \end{cases} \quad (18)$$

The main reason for which the analysis of system stability is challenging is the fact that both A_1 and A_2 admit 1 as an eigenvalue.

The first result of this section can be stated in the form of the following

Theorem 1. *If $\mu > 0$, the Markov chain (13),(14) is positive Harris and ergodic. For any initial distribution ρ , the chain converges to its unique invariant probability measure π in total variation norm, i.e. denoting the transition probability by \mathbb{P} ,*

$$\left\| \int_{\mathcal{S}} \rho(dx) \mathbb{P}^n(x, \cdot) - \pi(\cdot) \right\| \rightarrow_{n \rightarrow \infty} 0.$$

Recall that the total variation norm of a signed measure ν is defined as

$$\|\nu\| := \sup_{f: |f| \leq 1} |\nu(f)|.$$

The proof uses the theory of general state space Markov chains. The following lemmas are instrumental in proving the result. For brevity, we only provide proof outlines, while the complete proofs can be found in the extended version [1].

Lemma 1. *If $1 - \lambda - \mu < 1$, then there exists a measure φ on \mathcal{S} such that the Markov chain (13),(14) is φ -irreducible.*

Proof: (Outline) Fix some finite closed interval I and some $a > 0$, and consider the set $E = I \times [0, a]$.

Consider the measure φ_E defined as follows: for any Borel set A , $\varphi_E(A) := \nu(A \cap E)$, where ν denotes the Lebesgue measure on \mathbb{R}^2 .

We show that our chain is φ_E -irreducible, that is, if $A \subset \mathcal{S}$ is such that $\varphi_E(A) > 0$, then for all $x \in \mathcal{S}$, there is a strictly




Figure 1. Typical trajectory

positive probability that the time τ_A of return in A is finite: $L(x, A) = \mathbb{P}(\tau_A < +\infty) > 0$.

We consider any measurable set $B \subset E$. By Proposition 4.2.1 (ii) from [9], it suffices to show that there exists a finite $T > 0$, such that the probability of hitting B in T steps starting from any point $x \in \mathcal{S}$ is lower-bounded by a factor $\alpha(B, T, x) > 0$ (which might depend on B , T and x) times the irreducibility measure φ_E of B :

$$\mathbb{P}^T(x, B) \geq \alpha \varphi_E(B).$$

Indeed, if $\varphi_E(B) > 0$, then necessarily $\mathbb{P}^T(x, B) > 0$, and hence, $L(x, B) > 0$, since T is finite.

The key element in the proof is the observation that, at each time step, the R coordinate positions the Markov chain in region D_3 with a strictly positive probability due to the Gaussian sequence $(N(t))_t$. Furthermore, in this region the Z coordinate decreases geometrically at rate $1 - \lambda - \mu$. This property enables us to exhibit trajectories having a finite number of steps and leading from any point x to the set B . Such a trajectory is shown in Figure 1. ■

Lemma 1 shows that for any compact set of strictly positive Borel measure of the state space, its hitting time starting from any point is finite with strictly positive probability.

A set C is said to be ν_T -small for some non-trivial measure ν_T and a positive integer T , if for all $x \in C$, the probability of reaching any measurable B in T steps is lower-bounded as $\mathbb{P}^T(x, B) \geq \nu_T(B)$.

Furthermore, a set C is *petite* if there exists a distribution h on the positive integers and a non-trivial measure ν_h , such that for any $x \in C$, and for any Borel set B , the transition kernel of the *sampled chain* has the following property:

$$K_h(x, B) := \sum_{t \geq 0} h(t) \mathbb{P}^t(x, B) > \nu_h(B).$$

A ν_T -small set is implicitly δ_T -petite, where δ_T is the Dirac distribution.

Lemma 2. For any set $C = J \times [0, b]$, with J a finite closed interval and $b > 0$, there exists $T_0 > 0$ and non-trivial measures $(\nu_T)_{T \geq T_0}$ such that C is ν_T -small for all $T \geq T_0$.

Proof: (Outline) In Lemma 1 we essentially proved that we can reach any bounded Lebesgue-measurable set B of positive measure from any state x in a finite number of steps. In order to give an upper bound on the required number of steps, we defined the set $E = I \times [0, a]$ and we introduced the measure φ_E , which is defined as the Lebesgue measure of the set obtained via intersection with E .

In order to prove smallness, we need to eliminate the dependence of the number of steps T on the specific starting point $x \in C$ and on the destination B . Since we rely on the deterministic geometric decrease of the coordinate Z in region D_3 , a choice for B which is such that all the points in $B \cap E$ have a Z -coordinate less than a certain δ , automatically leads to a logarithmic number of steps in δ (thus, arbitrarily large, for δ close to 0).

Instead, we pick a finite closed interval I and constants $\Delta > \delta > 0$, and we define the set $A := I \times [\delta, \Delta]$. Then it can be shown that there exists a minimum number of steps T_0 depending on A and C , and an $\alpha(T) > 0$ which does not depend on B , or the starting point x , but which does depend on T , such that for measures

$$\nu_T := \alpha(T)\varphi_A,$$

the set C is ν_T -small, for all $T \geq T_0$. ■

We give two direct consequences of the two above lemmas.

Since any set $C = J \times [0, b]$ is both ν_T - and ν_{T+1} -small for some $T > T_0$, it follows that

Corollary 1. The Markov chain (13),(14) is aperiodic.

Additionally,

Corollary 2. Any compact subset of the state space is petite.

We are now ready to give the following

Proof: (of Theorem 1)

We prove that the function $H : \mathbb{R} \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$,

$$x = \begin{bmatrix} r \\ z \end{bmatrix} \mapsto H(x) = (r + \lambda z)^2 + (r + (\lambda + \mu)z)^2 \quad (19)$$

is a Lyapunov function for the system (13)-(14). It can be shown that H is unbounded off petite sets, that is for any $n < \infty$ the sublevel set $C_H(n) := \{y : H(y) \leq n\}$ is small. Furthermore, there exist constants $a, b, c > 0$ and a set $C = [-a, a] \times [0, b]$ such that the drift

$$DH(x) := \mathbb{E}_x H(X(1)) - H(x)$$

satisfies

$$DH(x) \leq -1 + c\mathbb{1}_C.$$

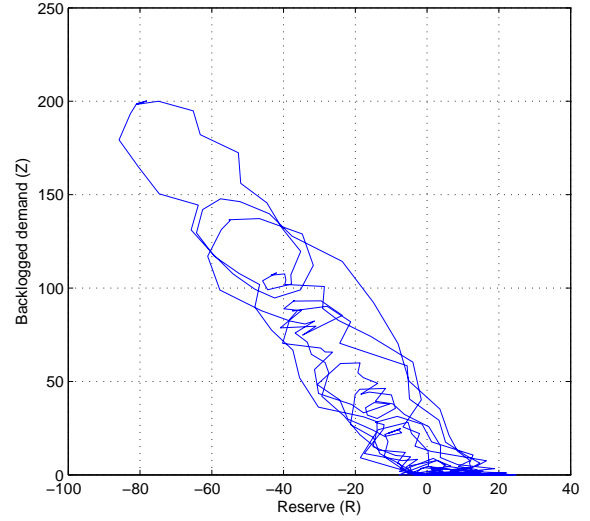


Figure 2. 500 iterations of the Markov process (13)-(14) for $\zeta = 1, r^* = 10, \sigma = 5, \lambda = 0.3, \mu = 0.1$

Separate cases need to be considered for each of the three regions D_1, D_2 and D_3 .

A consequence of this result is that, by Theorem 9.1.8 of [9] and by Corollary 2, which shows that the set C is small, the chain is Harris recurrent. Furthermore, by Theorem 10.0.1 of [9], it admits a unique invariant measure π . Finally, by Theorem 13.0.1 of [9] and by Corollary 1, we get the finiteness of π and we conclude. ■

This first result signifies that for positive evaporation $\mu > 0$, the simple threshold policy (17) stabilizes the system. A typical simulated trajectory is shown in Figure 2. Most points are found around the state $(r^*, 0)$, with some excursions in domain D_1 due to the variability of the demand.

The second result of this section concerns the case for which we have negative evaporation $\mu < 0$. It is stated as the following

Theorem 2. If $\mu < 0$, the Markov chain (13)-(14) is non-positive.

Proof: (Outline) Notice that if $\mu \leq -\lambda$, then the Z coordinate of the Markov chain cannot decrease. Hence the chain is trivially unstable (it is not even φ -irreducible).

In the case $-\lambda < \mu < 0$, we turn again to [9] for proving the claim. In this case the Markov chain (13)-(14) is φ -irreducible. We need to find a function H which satisfies the hypothesis of Theorem 11.5.2 from [9] to show non-positivity.

Define $H : \mathbb{R} \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$,

$$H(r, z) = \begin{cases} \log\left(\frac{r+(\lambda+\mu)z}{\mu}\right) & \text{if } r + (\lambda + \mu)z \leq \mu, \\ 0 & \text{otherwise.} \end{cases} \quad (20)$$

Then, it can be shown that H has finite increments in any point of the state space, namely that

$$\sup_{x \in \mathbf{S}} \mathbb{E}_x |H(X(1)) - H(x)| < +\infty. \quad (21)$$

In order to show this property of H , we need to consider points x in all three domains D_1 , D_2 and D_3 , while distinguishing between the points $x = (r, z)$ in D_1 for which $r + (\lambda + \mu)z \leq \mu$ holds, and those for which it does not.

Moreover, it can be shown that there exists a constant $v_0 > 1$ which is such that for all $x = (r, z) \in \mathbf{S}$ satisfying $r + (\lambda + \mu)z \leq \mu v_0$, the drift of H evaluated in such points is positive:

$$\mathcal{D}H(x) := \mathbb{E}_x H(X(1)) - H(x) > 0. \quad (22)$$

Define the set $C := \{x = (r, z) : r + (\lambda + \mu)z \geq \mu v_0\} \cap \mathbf{S}$. Then any x in the complementary $\bar{C} := C^c \cap \mathbf{S}$ is such that $H(x) > \sup_{x_1 \in C} H(x_1)$ and $\mathcal{D}H(x) > 0$. Using (21) and (22) we can apply Theorem 11.5.2 from [9] and conclude. ■

Let us sum up these results. We have shown that:

- If a positive fraction of latent jobs disappear during each time slot ($\mu > 0$), then *any* threshold policy stabilizes the system.
- On the other hand, if delaying any job results in an increase of its requirement/workload by a positive fraction ($\mu < 0$), then *there exists no threshold policy* that stabilizes the system.

The critical case for which $\mu = 0$ remains to be analyzed.

IV. CONCLUSION

In this paper, we considered a macroscopic model for electricity production and consumption. We assumed that the allocation process is done in two steps: a first step, in which demand and renewable supply are forecast (the day-ahead market) and a second step, in which real-time demand and real-time supply are matched as closely as possible (the real-time market), under ramp-up and ramp-down constraints regarding the rate at which real-time supply can be varied. We further assumed that all demand is adaptive and that backlogged demand cannot be observed by the supplier.

We introduced two parameters:

- the average delay $1/\lambda$ after which frustrated demand is expressed again, and
- the evaporation μ , the fraction of backlogged demand that disappears.

We showed that, if the evaporation μ is positive, then a threshold policy, targeting a fixed supply reserve r^* at

any time (under the ramp-up and ramp-down constraints), manages to stabilize the system for *any* value of r^* .

We also showed that, in case of negative evaporation μ , there exists no threshold policy which stabilizes the system.

Further research is needed to understand and control the phenomenon of negative evaporation.

REFERENCES

- [1] J.-Y. L. Boudec and D.-C. Tomozei. Satisfiability of elastic demand in the smart grid. *CoRR*, abs/1011.5606, 2010.
- [2] I.-K. Cho and S. P. Meyn. Optimization and the price of anarchy in a dynamic newsboy model. In *Stochastic Networks, invited session at the INFORMS Annual Meeting*, pages 13–16, 2005.
- [3] I.-K. Cho and S. P. Meyn. Efficiency and marginal cost pricing in dynamic competitive markets. To appear in *J. Theor. Economics*, 2006.
- [4] J. Eto, J. Nelson-Hoffman, C. Torres, S. Hirth, B. Yinger, J. Kueck, B. Kirby, C. Bernier, R. Wright, A. Barat, et al. Demand response spinning reserve demonstration. 2007.
- [5] S. Keshav and C. Rosenberg. How Internet Concepts and Technologies Can Help Green and Smarten the Electrical Grid. In *Proc. ACM SIGCOMM Green Networking Workshop*, August 2010.
- [6] J. Kueck, B. Kirby, T. Laughner, and K. Morris. Spinning Reserve from Responsive Load. Technical report, Oak Ridge National Laboratory (ORNL), 2009.
- [7] N. F. Maxemchuk. Communicating appliances. 2nd IMDEA Networks Annual International Workshop Energy Efficiency and Networking (Madrid), May-June 2010.
- [8] S. Meyn, M. Negrete-Pincetic, G. Wang, A. Kowli, and E. Shafieepoorfard. The value of volatile resources in electricity markets. Submitted to the *49th Conf. on Dec. and Control*, 2010.
- [9] S. P. Meyn and R. L. Tweedie. *Markov Chains and Stochastic Stability (2nd edition)*. Cambridge University Press, 2009.
- [10] M. J. Neely, A. S. Tehrani, and A. G. Dimakis. Efficient algorithms for renewable energy allocation to delay tolerant consumers. Technical Report arXiv:1006.4649, Jun 2010.
- [11] G. Wang, A. Kowli, M. Negrete-Pincetic, E. Shafieepoorfard, and S. Meyn. A control theorist's perspective on dynamic competitive equilibria in electricity markets. In *Proc. 18th World Congress of the International Federation of Automatic Control (IFAC)*, Milano, Italy, 2011.