



HAL
open science

Scalable and Privacy-Preserving Admission Control for Smart Grids

Giovanni Neglia, Giuseppe Di Bella, Laura Giarré, Ilenia Tinnirello

► **To cite this version:**

Giovanni Neglia, Giuseppe Di Bella, Laura Giarré, Ilenia Tinnirello. Scalable and Privacy-Preserving Admission Control for Smart Grids. [Research Report] RR-8769, Inria Sophia Antipolis. 2015. hal-01193294

HAL Id: hal-01193294

<https://inria.hal.science/hal-01193294>

Submitted on 4 Sep 2015

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Scalable and Privacy-Preserving Admission Control for Smart Grids

Giovanni Neglia, Giuseppe Di Bella,
Laura Giarré, Ilenia Tinnirello

**RESEARCH
REPORT**

N° 8769

September 2015

Project-Team Maestro



Scalable and Privacy-Preserving Admission Control for Smart Grids

Giovanni Neglia*, Giuseppe Di Bella†,
Laura Giarré‡, Ilenia Tinnirello‡

Project-Team Maestro

Research Report n° 8769 — September 2015 — 26 pages

Abstract: Energy demand and production need to be constantly matched in the power grid. The traditional paradigm to continuously adapt the production to the demand is challenged by the increasing penetration of more variable and less predictable energy sources, like solar photovoltaics and wind power. An alternative approach is the so called *direct control* of some inherently flexible electric loads to shape the demand. Direct control of deferrable loads presents analogies with flow admission control in telecommunication networks: a request for network resources (bandwidth or energy) can be delayed on the basis of the current network status in order to guarantee some performance metrics. In this report we go beyond such an analogy, showing that usual teletraffic tools can be effectively used to control energy loads. In particular we propose a family of control schemes which can be easily tuned to achieve the desired trade-off among resources' usage, control overhead and privacy leakage.

Key-words: smart grids, direct load control, admission control, privacy, large deviation

* Inria Sophia-Antipolis Méditerranée Email: giovanni.neglia@inria.fr

† Swarm Joint Open Lab, TELECOM ITALIA, Email: giuseppe.dibella@telecomitalia.it.

‡ Università degli Studi di Palermo, Email: {firstname.lastname}@unipa.it

**RESEARCH CENTRE
SOPHIA ANTIPOLIS – MÉDITERRANÉE**

2004 route des Lucioles - BP 93
06902 Sophia Antipolis Cedex

Contrôle d'admission passant à l'échelle et préservant la vie privée pour les réseaux intelligents

Résumé : Demande et production d'énergie doivent être constamment équilibrées dans le réseau électrique. Le paradigme traditionnel qui consiste à adapter dynamiquement la production à la demande est contesté par la montée en puissance de sources d'énergie plus variables et moins prévisibles, comme l'énergie solaire photovoltaïque et l'énergie éolienne. Une approche alternative est le contrôle direct de certaines charges électriques intrinsèquement flexibles dans le but de modeler la demande. Le contrôle direct des charges reportables présente des analogies avec le contrôle d'admission des flux de données dans les réseaux de télécommunication : une demande pour les ressources du réseau (bande passante ou énergie) peut être différée selon l'état actuel du réseau afin de garantir certaines mesures de performance. Dans ce rapport, nous allons au-delà d'une telle analogie et montrons que les outils de télétrafic habituels peuvent être efficacement utilisés pour contrôler les charges électriques. En particulier, nous proposons une famille de schémas de contrôle qui peuvent être facilement configurés pour atteindre le compromis souhaité entre utilisation de ressources, coûts du contrôle et révélation d'informations privées.

Mots-clés : Réseaux intelligents, contrôle direct des charges électriques, contrôle d'admission, grandes déviations

1 Introduction

Direct Load Control (DLC) allows energy utilities to control electric loads at the customers' premises. In the past DLC was used in critical situations to prevent blackouts by shutting down these loads. More recently, an extensive use of DLC has been advocated as a way to shape energy demand peaks or provide other ancillary services [1, 2, 3].

In this report we consider a scenario where DLC functionalities are deployed at a large set of small deferrable energy loads, like appliances at residential users. As we are going to discuss, the required additional intelligence and communication capabilities may be introduced through *smart plugs*, without the need to replace older "dumb" appliances. For such a scenario, we propose a simple control mechanism that guarantees that the instantaneous power demand exceeds a given bound with probability smaller than ϵ . Only a stochastic characterization of the power demand from each class of appliances is required. This mechanism combines two different operation paradigms. In the first one, appliances need to ask a controller the permission to start, and the controller will limit the number of simultaneously active appliances to $n(t)$. In the second one, an activation probability function— $p(t)$ —is broadcast periodically to all the appliances; appliances do not notify the controller but they start with probability $p(t)$ and postpone their decision to the time $t + T$ with probability $1 - p(t)$. The first operation paradigm requires more communication exchanges between the appliances and the controller and reveals more information about the customers' habits. The second paradigm works in an open-loop fashion on the basis of historical aggregated data and, then, does not disclose any private information. At the same time the lack of an exact knowledge of the current number of active appliances generates a lower average utilization of the resources in order to satisfy the constraint. We simply combine the two paradigms by means of a probability q : when an appliance wants to start operating, it will ask the permission to the controller with probability q , and it will decide autonomously using the function $p(t)$ with probability $1 - q$. The parameter q can then be chosen in order to achieve the wished trade-off among resources' usage, control overhead and privacy leakage. In fact, by increasing q we gradually i) reduce privacy by exposing more the energy profile of each user, ii) increase control overhead because the controller needs to directly interact with a large number of appliances, iii) increase the efficiency. Fig. 1 qualitatively depicts the effect of the parameter q . The detailed description of our system is in Sec. 3, and our analysis (Sec. 4) and experiments (Sec. 6) allow to quantify these trade-offs.

In our mechanism, the control policy is determined by the two functions $p(t)$ and $n(t)$. In Sec. 4, we show how techniques developed for flow admission control in packet networks can be used to determine a stationary control policy (i.e. $p(t) = p$ and $n(t) = n$), when the appliances' activation rate is assumed to be time-invariant. Our report contributes then to show how teletraffic engineering tools can be advantageously used also in the context of future smart power grids. In Sec. 5, we derive time-variant control policies for the more realistic case when the appliance activation rate i) is time-varying and ii) needs to be estimated. Our main goal in the present report is not to design a full-fledged solution for DLC, but rather to show how teletraffic tools (such as those used for Call Admission Control in telecommunication networks) can be effectively used to achieve a large scale deployment of the DLC approach to shape energy peaks.

2 Related Work

Direct load control is a specific mechanism of demand side management, that allows electric utilities to turn specific users' appliances on and off during peak demand periods and critical events. Most of the current large scale DLC programs work on thermostatic loads [1, 2], such

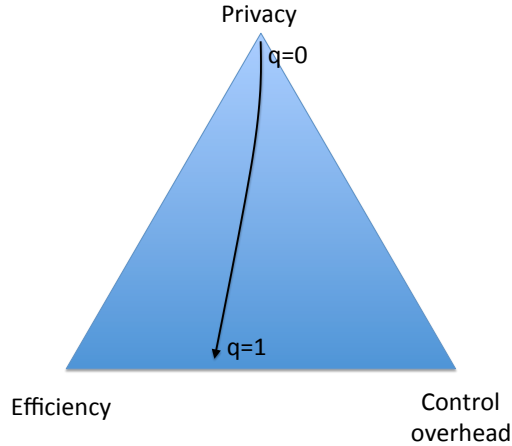


Figure 1: Different trade-offs among resources' usage, control overhead and privacy leakage achievable by tuning the parameter q .

as air conditioners and heating systems, because they allow a fine-tuning regulation of power demand. Alternative approaches consider electric vehicles or other battery-empowered appliances [4] which, beside acting as adaptive loads, can also reinject energy in the grid. as adaptive loads to be used for regulation. Multiple load typologies, including interruptible or deferrable loads such as the pool pumps, have been also considered for responding to different frequency components of the regulation signal [5]. Our focus in this report is on deferrable loads and simple control policies requiring limited intelligence in the network.

Direct load control programs have been activated in pilot projects for both residential and industrial customers. For residential customers, some utilities also offer commercial services [3] based on the deployment of radio-controlled switches on air-conditioning units or electric water heaters at the customers' premises. For industrial customers, the control mechanisms are very simple and often based on voice dispatch (i.e. a telephone network) [6]. Alternative communication infrastructures are currently considered for future large scale deployments [7]. While power grid operators envision the utilization of advanced metering infrastructure, mainly based on power line communications (PLC) and GPRS technologies, for transporting DLC signals, Internet connections available in most households may enable more flexible programs managed by energy suppliers or third parties [8]. Other solutions are considering the use of PLC from energy meters to zone data concentrators, which can access the Internet by means of a public telecommunication network [9]. Our solution can benefit from data concentrators' broadcast functionalities to control a large number of appliances, but it is otherwise essentially architecture-agnostic.

The specific DLC control policy can be designed on the basis of dynamic programming optimization [10], fuzzy logic-based decisions [11], or other profit maximization schemes [12]. In this report we consider simpler approaches rooted in teletraffic engineering which intrinsically work better for large scale systems. Indeed, loads on the electrical grid are multiplexed at different aggregation levels (distribution transformers, primary station, prize zone), similarly to traffic from data sources multiplexed at a router of a different hierarchical level. The use of teletraffic tools is not completely novel. Network calculus has been recently exploited to size energy batteries [13] and transformers [14]. Queuing theory is also used in [15] for sizing the population of customers subscribing a DLC program under a given maximum tolerable delay for activating the controlled appliances. In this report we consider a different application of teletraffic tools to smart grids, highlighting the analogy between DLC and call admission control.

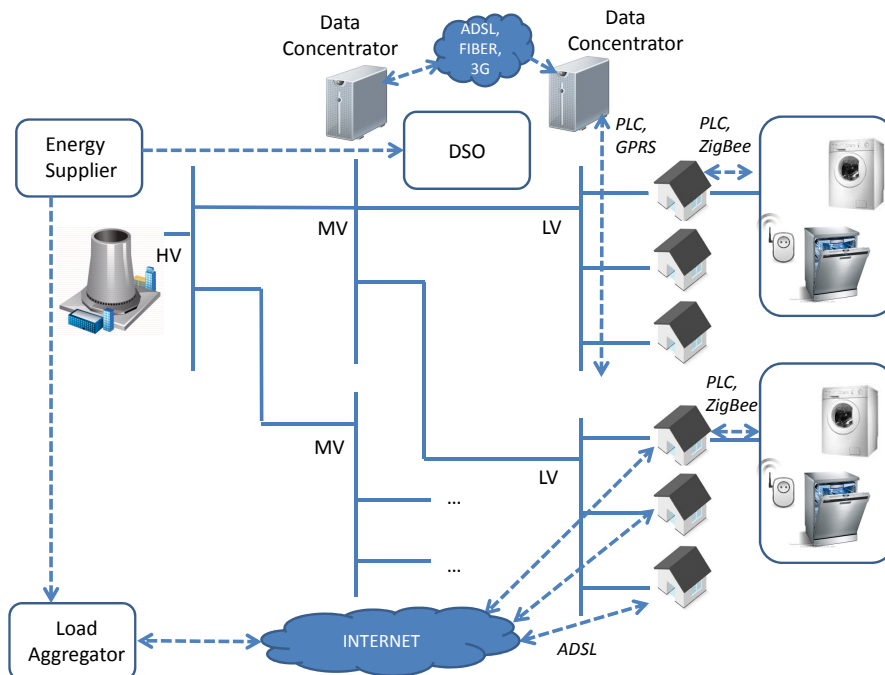


Figure 2: Reference scenario: communication infrastructure and control system for DLC managed by a DSO and/or a load aggregator.

3 System Description

We consider the problem of peak shaving for an energy supplier, which wants to reduce its customers' consumption during the time of the day when energy costs are higher. To achieve these goals, the energy suppliers may interact with the distribution system operators (DSO), and/or with novel intermediate figures called load aggregators [16] by means of a communication infrastructure transporting the control messages. Both possibilities are illustrated in Fig. 2 where dashed lines indicate control messages flows. The solution based on the DSO can take advantage of broadcast functionalities deployed at the data concentrators. In fact broadcasting of low-rate control signals can be very efficiently performed by PLCs in the low-voltage (LV) distribution grid [17]. The aggregator can exploit standard Internet connectivity to interact with appliances at customers' premises. The two approaches can coexist. Our solution can operate in both the scenarios, so we will talk generically about the load controller to denote the entity which drives the appliances.

The energy supplier specifies an high-level command for the controller in terms of maximum tolerable probability to exceed a given power demand in the controlled area. The load aggregator is responsible to meet this requirement by deciding about the activation of deferrable loads, such as dish-washers and laundry-machines, in the controlled area. On the basis of the demand prediction, the requirements are mapped into a control signal that is broadcast to all the controlled households for deciding about the admission or the deferral of load activation requests. Each household is equipped with a gateway able to receive the control signal from the controller and to interact with the domestic appliances by means of local area technologies (such as ZigBee or




Figure 3: Household control model in terms of activation probability $p(t)$ and ratio of forwarded queries q .

PLC). Smart appliances can be natively equipped with a programmable interface able to communicate with the gateway, while dumb appliances can be controlled by means of smart plugs [18].

In our control system, the signal broadcast by the load controller represents the probability $p(t)$ to accept or defer a load activation request during the intervals of the day. At each household, the gateway can autonomously process a novel activation request on the basis of this function: the activation request is accepted with probability $p(t)$ and deferred of a fixed time interval T with probability $1 - p(t)$. This operation applies to a fraction $1 - q$ of all the requests. The others are directly forwarded to the load controller and the gateway waits for an ACK from the load controller in order to accept the request. The controller maintains a cap $n(t)$ on the number of appliances active at a given time instant (among those it is aware of, i.e. those whose activation request has been sent to the controller). The ACK can then be delayed until some previously activated appliance does not stop operating. This mechanism allows the load controller to actually estimate the time-varying power demand, but also to have a tighter control on the aggregate power demand and then to achieve a more effective resource utilization in the controlled area.

Fig. 3 summarizes the actuator model at the household: different control operations can be programmed by the load controller by specifying q and $p(t)$. For example, when q is set to 1, all the activation requests are forwarded to the load controller by means of unicast transmissions. This implies a better control on the aggregated power demand, because the number of active appliances is known and the only source of randomness is due to the appliance consumptions. On the contrary, when q is very small, most of communications are unidirectional (from the load controller to the households), while decisions can be taken locally with minimal delays. However, the aggregate power consumption is affected by two sources of randomness (the one related to the number of appliance active, and the other one related to the appliance consumption) which require a lower admission rate for not exceeding the power constraint.

4 Control in a Stationary Setting

Our goal is to design a direct load control system able to impose a certain threshold on the overall power absorbed from a set of appliances with known statistical properties. In the following, we will consider each appliance power consumption as a stochastic process, characterized by a random variable X with probability density function equal to $f_X(x)$.

Let $m(t)$ be the number of active appliances at time t , then the total absorbed power is $P(t) = \sum_{i_1}^{m(t)} X_i$. First we consider that the energy utility wants to guarantee that $P(t)$ does

not exceed the threshold \bar{P} with probability larger than ϵ , i.e.:

$$\Pr\left(\sum_{i=1}^{m(t)} X_i > \bar{P}\right) < \epsilon. \quad (1)$$

Then, in Sec. 4.4, we show how similar results can be derived for a different type of constraints on $P(t)$.

4.1 Classic large deviation results for $q = 1$

If $q = 1$, i.e. if all the queries are forwarded to the controller, our problem is to determine the maximum value n , such that

$$\Pr\left(\sum_{i=1}^n X_i > \bar{P}\right) < \epsilon. \quad (2)$$

The same problem has been considered in the context of Call Admission Control (CAC) in telecommunication networks. The interpretation of the quantities is different: the purpose is to determine the maximum number n of homogeneous data flows, each with instantaneous rate X_i in order to guarantee that the aggregate rate on a link exceeds the value \bar{P} (e.g. the link capacity) with probability at most ϵ .

Different approaches have been proposed to determine n . Here, we introduce a simple one based on basic large deviation results (see e.g. [19, Ch. 6]), but more sophisticated ones could also be applied [20]. Our main goal is hereafter to show how teletraffic tools can in general be reused in this different context, rather than develop the solution in all its details.

Let $M_X(s) = \ln \mathbb{E}[e^{sX}]$ be the cumulant generating function. From the Chernoff bound and basic properties of the cumulant generating function, it follows easily that:

$$\ln \Pr(X_1 + \dots + X_n > nc) \leq n \inf_{s \geq 0} [M_X(s) - sc] \quad (3)$$

Moreover, Cramer's theorem proves that this bound is asymptotically tight:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln \Pr(X_1 + \dots + X_n > nc) = \inf_{s \geq 0} [M_X(s) - sc] \quad (4)$$

Practically speaking, when n is large, Cramer's theorem can be used to approximate the probability that the sum of n independent random variables exceeds a bound (\bar{P} in our case) as follows:

$$\ln \Pr(X_1 + \dots + X_n > \bar{P}) \approx \left(\inf_{s \geq 0} [nM_X(s) - s\bar{P}] \right),$$

and then n can be determined as the largest integer such that

$$\left(\inf_{s \geq 0} [nM_X(s) - s\bar{P}] \right) < \ln \epsilon.$$

The set of values $\{0, 1, \dots, n\}$ is called the *acceptance region* for the admission controller.

The approach can be easily generalized to a finite number of appliances' classes and the concept of *effective bandwidth* can be defined for each class in order to easily derive a subset of the acceptance region [19, Ch. 6].

4.2 Extension to $q < 1$

In this section we show how the results above can be extended to the scenario where a subset of the appliances do not perform any query to the controller, but start autonomously their operation. For the moment we assume that such appliances turn on according to a homogeneous Poisson process with rate λ_c . In the following Sec. 5 we explain how this rate can be determined taking into account the retrial mechanism described in Sec. 3 and how one can deal with the actual time-variant activation process. For the moment, we just anticipate that λ_c is a linear function of $1 - q$ and of the instantaneous probability of activation p : $\lambda_c \propto (1 - q)p$. In order to keep the equations simple, we assume that all the appliances have the same activation time equal to D , but the results can be easily extended to the case when D is a random variable.

Let N be the random number of appliances starting autonomously. Under the assumptions indicated above, N is distributed as a Poisson random variable with parameter $E[N] = \Lambda_c = \lambda_c D$. The condition in Eq. (2), then becomes

$$\Pr \left(\sum_{i=1}^n X_i + \sum_{j=n+1}^{n+N} X_j > \bar{P} \right) < \epsilon, \quad (5)$$

where $(X_i)_{i \in \mathbb{N}}$ is a sequence of independent and identically distributed random variables.

Large deviation results hold for large systems, for example when both the number of random variables (n) and the threshold to be exceeded (nc) diverge as in Eq. (4). In Eq. (5) we need to let both the first and the second addend scale in the same way, otherwise one of them would become negligible in comparison to the other. We assume then that the mean Λ_c of the number of appliances starting autonomously scales linearly with n : $\Lambda_c = n\Lambda_c^0 = n(1 - q)pD\lambda_c^0$, where we also took into account λ_c 's dependence on p and $1 - q$. It is then possible to prove the following result:

Theorem 1. *Let $(X_i)_{i \in \mathbb{N}}$ be a sequence of independent and identically distributed random variables with cumulant generating function $M_X(s)$, and N be a Poisson random variable with mean $n\Lambda_c^0$. We assume that $M_X(s)$ is defined for any s and $\Pr\{X > c\} > 0$. It holds:*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln \Pr \left(\sum_{i=1}^n X_i + \sum_{j=n+1}^{n+N} X_j > nc \right) = \inf_{s \geq 0} \left[M_X(s) + \Lambda_c^0 \left(e^{M_X(s)} - 1 \right) - sc \right].$$

Proof. Let us define $\{Y_{i,k}, i = 1, \dots, n, k \in \mathbb{N}\}$ to be a set of independent random variables distributed as X and $\{N_i, i = 1, \dots, n\}$ to be a set of independent Poisson random variables with mean Λ_c^0 . The sum of the variables N_i has the same distribution of N and then:

$$\sum_{i=1}^n \sum_{k=1}^{N_i} Y_{i,k} \stackrel{d}{=} \sum_{j=n+1}^{n+N} X_j.$$

Let $Z_i \triangleq X_i + \sum_{k=1}^{N_i} Y_{i,k}$, for $i = 1, \dots, n$. It follows that

$$\sum_{i=1}^n Z_i \stackrel{d}{=} \sum_{i=1}^n X_i + \sum_{j=n+1}^{n+N} X_j.$$

All the variables Z_i are independent and identically distributed with cumulant generating function $M_Z(s) = M_X(s) + \Lambda_c^0 (e^{M_X(s)} - 1)$. The thesis follows by applying directly Cramer's theorem to $\sum_{i=1}^n Z_i$.¹ \square

¹ An anonymous reviewer from IEEE CDC 2015 suggested this simpler proof. Our original proof is in Appendix A.




Figure 4: Two-state Markov process as a simplified model of the appliance power consumption.

In a finite-size system, this result is used to approximate the probability that the total power consumption exceeds the bound \bar{P} as follows:

$$\ln \Pr \left(\sum_{i=1}^n X_i + \sum_{j=n+1}^{n+N} X_j > \bar{P} \right) \approx \inf_{s \geq 0} \left[nM_X(s) - \lambda_c D \left(e^{M_X(s)} - 1 \right) - s\bar{P} \right].$$

and then the acceptance region is made by all the pairs of values (p, n) such that:

$$\inf_{s \geq 0} \left[nM_X(s) + \lambda_c D \left(e^{M_X(s)} - 1 \right) - s\bar{P} \right] \leq \ln \epsilon. \quad (6)$$

In the next subsection, we characterize such region for a simple case.

4.3 A numerical example

In order to illustrate the admission control rule determined by the results above, we consider a toy example where there is a single class of appliances, which would like to activate according to a Poisson process with rate λ . Then, there will be a Poisson process with rate $q\lambda$ to the query-response system, and a spontaneous activation process of appliances with rate $p(1-q)\lambda$ (we ignore retrials for the moment). We also assume that the appliance's consumption profile can be modeled as the two-state Markov process in Fig. 4. The appliance stays in the high consumption state on average 10 minutes and consumes 1.5kW, it stays in the low consumption state on average 20 minutes and consumes 500W. Literature suggests that Markov processes can indeed accurately model appliances' consumption [21].

For this appliance, the cumulant generating function can be easily calculated as $M_X(s) = \ln(1/3e^{s1.5} + 2/3e^{s0.5})$. Then inequality (6) can be used to determine the acceptable pairs (p, n) with $\lambda_c = p(1-q)\lambda$:

$$\inf_{s \geq 0} \left[nM_X(s) + p(1-q)\lambda D \left(e^{M_X(s)} - 1 \right) - s\bar{P} \right] \leq \ln \epsilon. \quad (7)$$

Fig. 5 shows the frontiers of 10 different acceptance regions calculated for $\lambda = 12$, $D = 90$ minutes, $q = 0.1, 0.2, \dots, 1$, $\epsilon = 0.1$ and $\bar{P} = 0.8\lambda D \mathbb{E}[X]$ (the 80% of the average absorbed power in absence of the control). Each curve² is made by the pairs (p, n) for which the constraint in (6) is satisfied with an equality. The acceptance region is then made by all the points that are

² Obviously n can only assume integer values, but here we are considering a continuous relaxation. In practice, the values can then be lower-rounded.

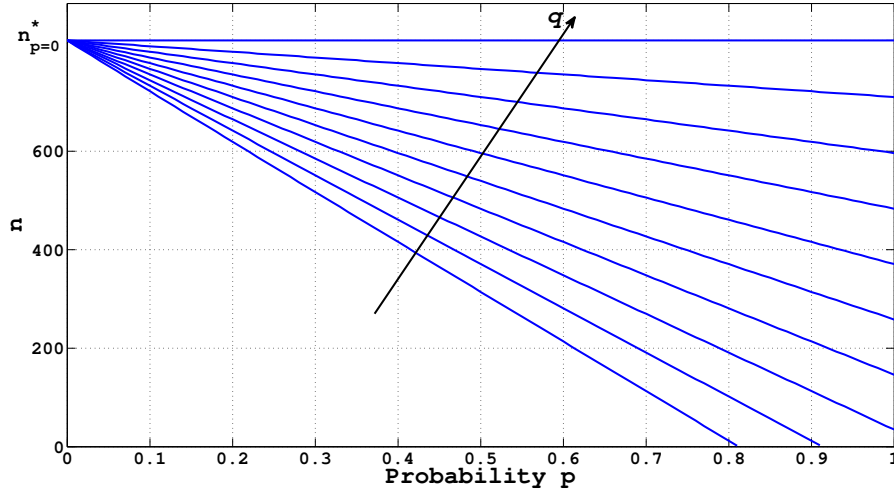


Figure 5: Acceptance regions of the pairs (p, n) for different values of the querying probability q .

below this curve. This is evident analytically because the left term in (7) is decreasing in n and p , but it is also intuitively clear because of the physics of the system: for a given acceptable pair (p, n) , if we limit more p or n ($p' < p$ or $n' < n$), then less appliances will be admitted in the system and the constraint will be even more satisfied (both (p', n) and (p, n') are acceptable). We observe the expected trade-off between n and p : on the frontier as p increases, n needs to decrease. The more appliances we decide to admit through the query-response system, the more we have to limit through p the expected number of appliances which will activate autonomously.

Another interesting remark about the frontiers in Fig. 5 is that they look almost linear. This can clearly significantly simplify the determination of the acceptance region and leads to important conclusions about the efficiency of resources' usage. For this reason, we investigate deeper such behavior. The term $e^{M_X(s)}$ grows very fast for $s \geq 0$ and because of this the point of minimum for the left-hand side of Eq. (7) is close to 0, specially for larger $p(1-q)$. It is then possible to consider a first order Taylor approximation ($e^{M_X(s)} - 1 \approx M_X(s)$) and the frontier can be approximated as follows:

$$\inf_{s \geq 0} [(n + (1-q)pD\lambda) M_X(s) - s\bar{P}] = \ln \epsilon,$$

For a given value of q , the frontier is made by pairs (p, n) such that:

$$n + (1-q)pD\lambda = \text{constant}. \quad (8)$$

To completely characterize this linear relation, it is enough to observe that for $p = 0$ the maximum value for n (let us denote it as $n_{p=0}^*$) does not depend on q , i.e. all the frontiers pass by the point $(0, n_{p=0}^*)$ and they are described by the following linear equation parameterized in q :

$$n + (1-q)pD\lambda = n_{p=0}^*. \quad (9)$$

If there are n appliances activated through the query-response mechanism, the expected number of appliances in the system is $n + (1-q)pD\lambda$ and then the expected power consumption is $E[P] = (n + (1-q)pD\lambda)E[X]$. This means that for a given q value all the points in the




Figure 6: Analogy between the energy bucket on the left and a work-conserving queuing system on the right.

frontier described by Eq. (9) have the same resources' usage $E[P] = n_{p=0}^* E[X]$ while satisfying the constraint. Moreover, for any value of q we can achieve the same efficiency by selecting (n, p) on the corresponding frontier. Although Eq. (9) stems from a linear approximation, it suggests that the efficiency (in terms of average utilization) of our control mechanism is not very sensitive neither to the pair (n, p) (as long as it is a point of the frontier), nor to the parameter q . This is confirmed by our numerical analysis: for example we observed less than 5% reduction of the maximum expected load changing q from $q = 1$ to $q = 0$ despite the fact that for $q = 0$ the control needs to deal also with the variability of the number of appliances admitted. However, the importance of direct queries in practical applications is not only due to efficiency reasons, but also to the possibility to accurately estimate the appliance activation rate as discussed in Sec. 5. Appendix B further discusses the relation between n, p, q and average utilization.

4.4 A different constraint

Before describing how the above control policy can actually be used in a case when the activation pattern is time-varying and unknown, we want to briefly show that teletraffic tools can also be used to deal with different constraints. In particular, the constraint $\Pr(P(t) > \bar{P}) \leq \epsilon$ ignores the time duration during which the constraint \bar{P} is violated. It may be significant for the energy utility to target the following different type of constraint.

Let us define an *energy bucket* as follows: the bucket can store an energy capacity equal to B , it is filled with constant rate \bar{P} and emptied with the variable rate $P(t)$. Let $L(t)$ denote the energy stored in the bucket at time t . Fig. 6 describes on the left side the energy-buffer. While inspired by the token-bucket used to shape data traffic, this energy-bucket is only used to express a constraint. If one requires that $L(t) > 0, \forall t$ then $P(t)$ can occasionally exceed the constraint \bar{P} by δP , but only for the time interval $[t, t + L(t)/\delta P]$. The ratio behind such bound is that smaller violations would be tolerated for a longer time. Given the uncertainty on the energy demand, the control policy cannot guarantee the deterministic bound described above, but rather its probabilistic version:

$$\Pr(L(t) = 0) < \epsilon, \forall t. \quad (10)$$

The possibility to use teletraffic tools (and in particular large deviation results) also to deal with the constraint in (10) derives from the following equivalence: the energy-bucket on the left of Fig. 6 is equivalent to the queue in the same figure on the right, where the input and the output process have been reversed, i.e. a traffic with rate $P(t)$ is offered to a queue with service

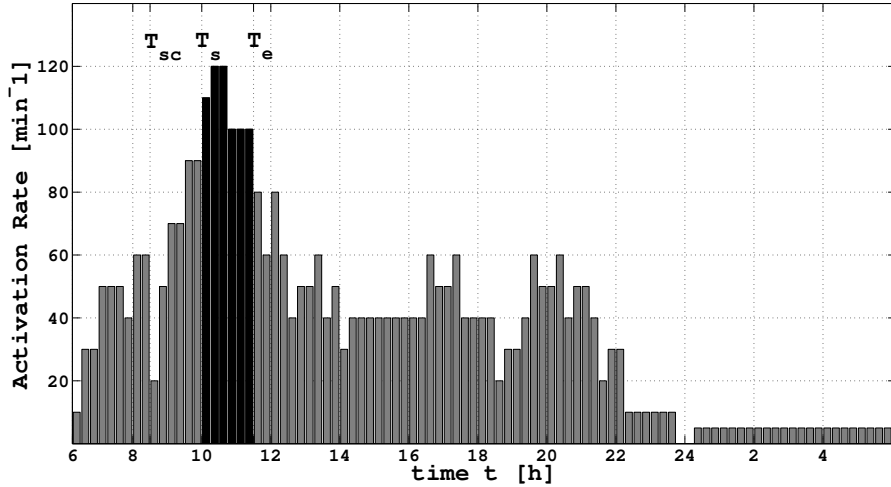


Figure 7: Aggregate activation rate for 30000 washing machines.

rate \bar{P} . Specifically, if $L(0) = B$ and $Q(0) = 0$, then $L(t) = Q(t)$ for every $t \geq 0$. This result is quite evident, but we have not found it in the literature, so we prove it in Appendix C.

The important consequence for our purposes is that controlling the activation of electric appliances so that the aggregate consumption rate satisfies constraint (10) is equivalent to admit data flows to a queue so that the probability to lose packets because of buffer overflow is smaller than ϵ . In fact (under large deviation approximation):

$$\Pr(Q(t) > B) \approx \Pr(Q(t) = B) = \Pr(L(t) = 0).$$

Similar results for the acceptance region hold [19, Ch. 6]:

$$\Pr(L(t) = 0) \approx \exp\left(\sup_{t \geq 0} \inf_{s \geq 0} [stn\alpha(s, t) - s(B + \bar{P}t)]\right),$$

where $\alpha(s, t) = \frac{1}{st} \log \mathbb{E} [e^{sX(0, t)}]$, with $X(0, t)$ being all the random energy consumption of one appliance during $[0, t]$.

This example has further shown the potential relevance of applying teletraffic tools to load control in smart grids.

5 Control in a Time-Variant Setting

In this section, we show how the analytical results derived in the previous section can be practically used when the appliances' activation process is not stationary, its rate is unknown, and probabilistically controlled appliances retry to activate some time later.

Indeed, the usage of electric appliances exhibits a strong time-of-the-day effect. For example Fig. 7 shows the activation rate of washing machines over 15 minutes intervals, as derived from data in [22]. We need then to take into account the effect of a time variant $\lambda(t)$, which is generally unknown even if historical data may be available. Moreover, because of the retrial mechanism from probabilistically controlled appliances, the actual activation rate results time-varying even when the spontaneous activation rate λ is constant.

Because of the time variability of the process to control, the control policy needs to be time-varying too, i.e. in general we will have $n(t)$ and $p(t)$. There are different possible choices on how to jointly adapt the two control actions, which lead to different performance in terms of resources' usage (see Appendix B for the effect of n and p on the efficiency), communication requirements (e.g. if p is constant, it does not need to be transmitted periodically) and fairness between the two groups of appliances—those controlled probabilistically and those controlled through query-response—(e.g. in terms of delay before the activation). For simplicity, in this paper, we decided to consider $n(t) = \text{const} = n$ constant and then to compensate for process changes by dynamically tuning $p(t)$. The control is time-slotted with time intervals of length T_c . Without loss of generality we consider that the control starts at time $t = 0$, and we denote by p_k the value of the control action during the k -th time slot, i.e. $p(t) = p_k$ for $t \in [(k-1)T_c, kT_c)$. For the sake of simplicity we will also assume that the retrial delay T is equal to T_c and the activation time of the appliance is a multiple of T_c ($D = dT_c$), even if the three parameters are in general independent.

Fig. 8 shows the model of the whole system we are going to describe in the following sections.

5.1 The plant

As it is usual in control theory, we call plant the combination of the process under control and the control actuator. The input is the spontaneous activation process with rate λ_k during the k -th time slot, i.e. the process of the activation instants in absence of any form of control (we omit time dependence in the figure). The control system assumes $\lambda(t)$ to be constant during a control slot, while this is not necessarily the case. In Sec. 6 we evaluate the effects of such an approximation. We keep assuming that the point process of all the activation time instants can be correctly modeled by a (non-homogeneous) Poisson process with rate $\lambda(t)$, because it originates from the superposition of many independent individual choices (users deciding to turn on their appliances). The initial requests are randomly split in two independent Poisson processes with rate respectively $q\lambda(t)$ and $(1-q)\lambda(t)$. The rate of appliances that will consider to activate autonomously in slot k is $\lambda_{eq,k}$. It holds:

$$\lambda_{eq,k} = (1-q)\lambda_k + \lambda_{eq,k-1}(1-p_{k-1}), \quad (11)$$

where the first term is due to all the appliances that are considering to activate themselves for the first time during the k -th time slot, and the second is due to those that have already considered this decision in the $(k-1)$ -th slot and have postponed it to the current one. The actual activation rate of probabilistically controlled appliances³ is

$$\lambda_{c,k} = p_k \lambda_{eq,k}.$$

Finally, the actual number of active appliances $m_{pc,k}$ admitted through the probabilistic control and active at time kT_c is equal to those arrived in the interval $[kT_c - D, kT_c]$. This is a Poisson random variable with expected value

$$\mathbb{E}[m_{pc,k}] = \sum_{h=k-d+1}^k \lambda_{c,h}.$$

5.2 The controller

The upper part of the figure describes the controller. This directly receives the activation queries with rate $q\lambda_k$ and it manages the activations as in a $M/D/n$ queue by guaranteeing that the

³ It is possible to show that also the point process of such activations is a Poisson process [23].

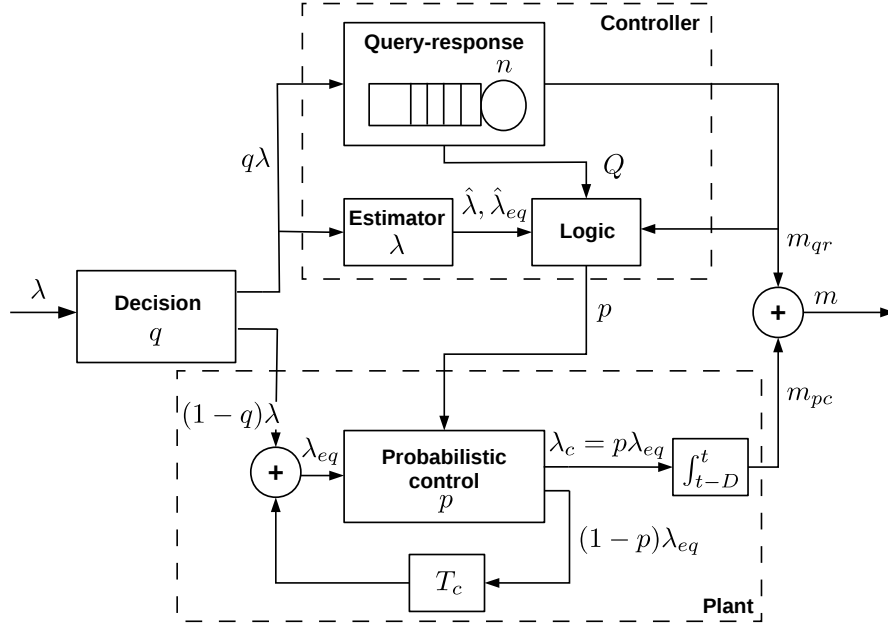


Figure 8: System block diagram.

number of appliances active at a given time does not exceed n . Let $m_{qr,t}$ denote the number of appliances active at time t .

Estimators. The controller does not know the state of the plant (e.g. how many appliances are taking the decision to activate autonomously). It needs then to estimate the rates λ and λ_{eq} . An estimate for quantity x is denoted as \hat{x} . In particular two different set of estimates will be useful: $\hat{\lambda}_k^p$ and $\hat{\lambda}_{eq,k}^p$ will estimate the sequences until the current slot k , while $\hat{\lambda}_{k+1}^f$ and $\hat{\lambda}_{eq,k+1}^f$ will be used as a prediction for the slot $k+1$ in order to determine p_{k+1} . In this paper we consider simple estimators for these quantities, but we show that they work reasonably well in Sec. 6. Given $N_{r,k}$ the number of queries received during the k -th slot, the Maximum Likelihood Estimator for λ_k is simply $\hat{\lambda}_k^p = N_{r,k}/(qT_c)$. Clearly, this estimate could be improved if some a-priori is available (e.g. from historical data). The prediction for slot $k+1$ is $\hat{\lambda}_{k+1}^f = \hat{\lambda}_k^p$.

The estimates $\hat{\lambda}_{eq,k}^p$ and $\hat{\lambda}_{eq,k+1}^f$ are simply obtained as a function respectively of $(\hat{\lambda}_k^p, \hat{\lambda}_{eq,k-1}^p)$ and of $(\hat{\lambda}_{k+1}^f, \hat{\lambda}_{eq,k}^p)$ using the same relation that links $\lambda_{eq,k}$ to λ_k , i.e.

$$\hat{\lambda}_{eq,k}^p = (1-q)\hat{\lambda}_k^p + \hat{\lambda}_{eq,k-1}^p(1-p_{k-1}).$$

$$\hat{\lambda}_{eq,k+1}^f = (1-q)\hat{\lambda}_{k+1}^f + \hat{\lambda}_{eq,k}^p(1-p_k).$$

Control logic. The controller determines p_{k+1} on the basis of the acceptance region derived as described in Sec. 4. For simplicity we consider a linearized frontier, but the following operation could be adapted to the case when the frontier is described by a more generic function. Moreover, in Sec. 4.3 we have shown that the linear approximation works very well. Let $n^{(q)}(p)$ be the frontier parametric equation for a given q value. We denote the minimum value of n on the frontier by $n_{\min}^{(q)}$ (it is 0 if the frontier intersects the segment $[0,1]$). Moreover let $p_{\max}^{(q)} \triangleq \max\{p | n^{(q)}(p) \geq n_{\min}^{(q)}\}$, i.e. the largest probability value in the frontier. For the sake of fairness

between the probabilistically controlled appliances and the querying ones, we would like the system to work in a point with coordinates (p, n) approximatively in the middle of the frontier. Under the linear approximation, and considering the maximum arrival rate in the control period, the corresponding value of n is $n_{\min}^{(q)} + \lambda_{\max}(1-q)p_{\max}^{(q)}/2$. It has to be remarked that the expected number of appliances at time t that have been activated through the query-response mechanism can be at most $\int_{t-D}^t \lambda(\tau)q d\tau$ (see discussion in Appendix B). Then we have set n as follows

$$n = \min \left(0.9 \max_{t \in [T_{sc}, T_e]} \left\{ \int_{t-D}^t \lambda(\tau)q d\tau \right\}, n_{\min}^{(q)} + \lambda_{\max}(1-q) \frac{p_{\max}^{(q)}}{2} \right). \quad (12)$$

We then select the control probability p in order to keep the total expected number of appliances constant, i.e.

$$\Lambda_{k+1} + n = \text{const}, \quad (13)$$

where Λ_{k+1} is the expected number of active probabilistically controlled appliances at the end of the $k+1$ -th slot, and the constant is equal to $n_{\min}^{(q)} + \lambda_{\max}(1-q)p_{\max}^{(q)}$.⁴ In the stationary case without retrials, it was simply $\Lambda_{k+1} = \lambda_c D = \lambda(1-q)pD$. Here, we can express Λ_{k+1} as the sum of two terms, one (denoted as $\Lambda_{hist,k+1}$) due to all the probabilistically controlled appliances already in the system at the begin of slot $k+1$ (or equivalently at the end of slot k), the other due to the estimated number of appliances which will activate during slot $k+1$. We can estimate $\Lambda_{hist,k+1}$ as follows

$$\hat{\Lambda}_{hist,k+1} = \sum_{h=0}^{d-2} p_{k-h} \hat{\lambda}_{eq,k-h}^p T_c,$$

and then

$$\hat{\Lambda}_{k+1} = \hat{\Lambda}_{hist,k+1} + \hat{\lambda}_{eq,k+1}^f T_c p_{k+1}. \quad (14)$$

Finally, p_{k+1} can be iteratively derived from Eq. (13) and Eq. (14) as:

$$p_{k+1} = \frac{n_q^* - n - \hat{\Lambda}_{hist,k+1}}{\hat{\lambda}_{eq,k+1}^f T_c}.$$

This may lead to a too prudential strategy, because the number $m_{qr,k+1}$ of active appliances in the query-response queue can be significantly smaller than n , specially at the begin of the control period (the queue fills initially with rate λq) and then the configuration above would lead to a severe underutilization. A solution is to predict the number of active appliances in the query-response queue during the residual control period (i.e. $\hat{m}_{qr,h}$ for $h > k$), on the basis of the queue status at time k and expected activation requests. Indeed, the activity state of the enqueued appliances can be deterministically calculated from their activation/de-activation time instants, while the number of future arrivals can be predicted from the current estimate of the arrival rate $\hat{\lambda}_{k+1}^f$. In particular let $a_{qr,h}$ be the total number of number of appliances in the query-response system (active or waiting) at slot h . This can be estimated as sum of two terms. The first term is the number of appliances currently (at slot k) in the query-response system which will be in the system at slot h . It can be immediately calculated from the current status of the queue (number of active appliances, the activation time instants, and number of appliances in the queue), we denote it as $a_{qr,h}^p$. The other term ($\hat{a}_{qr,h}^f$) requires the prediction of

⁴ The reader may wonder why we have introduced $n_{\min}^{(q)}$ and $p_{\max}^{(q)}$ and in particular why the linearized frontier is not expressed simply by $\Lambda_{k+1} + n = n_{p=0}^*$. The answer to these questions is in Appendix B.




Figure 9: Instantaneous absorbed power (with and without the applied control), and the computed activation probability $p(t)$.

future arrivals from now until h and can be estimated as $\hat{a}_{qr,h}^f = \lambda_k^f(h-k)T_c$. The estimate for $m_{qr,h}$ is then $\hat{m}_{qr,h} = \min(n, a_{qr,h}^p + \hat{a}_{qr,h}^f)$.

Then, p_{k+1} can be calculated replacing n in Eq. (13) with the most pessimistic forecast until the end of the control period, i.e. with $\max\{\hat{m}_{qr,h}, h > k\}$. This is the approach we adopted.

6 Numerical Results

In this section we show the performance of the system described in Sec. 5 in a realistic setting. In particular we consider 30000 washing machines under control. The instantaneous power consumption of a washing machine is assumed to follow the simple model in Fig. 4 and the activation time is $D = 90$ minutes. The spontaneous activation rate $\lambda(t)$ is the one in Fig. 7, derived from experimental data in [22]. The largest expected power demand P_{\max} is at around time 11.00am. We assume the energy supplier sets the constraint as $\Pr(P > \bar{P} = 0.8P_{\max}) < 0.1$ in the interval $[T_s, T_e] = [10.00\text{am}, 11.30\text{am}]$. In order to be sure to satisfy the constraint at $t = 10.00\text{am}$, the control needs to start at $T_{sc} = 8.30\text{am}$.

Fig. 9 plots the evolution of the power demand with and without control for $q = 0.5$ and $T_c = 15$ minutes together with the probability signal $p(t)$. We observe that the controller does not actually affect the system ($p(t) = 1$) until $t = 10.00\text{am}$ and that it actually manages to maintain the absorbed power below \bar{P} for the whole duration of the control interval. Observe also how power consumption significantly increases after T_e . This is due to the fact that a severe constraint has been imposed for a long time interval. Moreover, the power increase can be made smooth by gradually increasing \bar{P} after T_e .

The following figures 10 and 11 show the utilization—evaluated as ratio between the average power demand and \bar{P} —and the probability to exceed the bound at time $t = 10.26\text{am}$ estimated over 1000 simulations for different values of $q \in [0, 1]$ and for $T_c = 1, 5, 15, 30$ minutes. The time instant of observation falls in the interval where there is the largest activation rate after a period when the rate has been almost constantly increasing. It is then a particularly critical instant for




Figure 10: Resources utilization at 10.26am.

the control system. All confidence intervals in the figures have 95% confidence level.

The upper plots of figures 10 and 11 correspond to the ideal case when the controller has perfect estimation of the average request rate in the next timeslot, i.e. $\hat{\lambda}_{k+1}^f = \lambda_{k+1}$. We observe that, as we expected (see Appendix B), utilization increases as q increases, but no more than 3%. The probability to exceed the bound roughly reflects the same trends, with higher probabilities corresponding to higher utilization. In any case the probability values are well below $\epsilon = 0.1$. The curves almost overlap for all the values of T_c but $T_c = 30$ minutes. This is due to the fact that the actual activation request rate is constant over 15 minutes time intervals, then for $T_c = 1, 5, 15$ minutes, the knowledge of the average rate in the next slot corresponds to the knowledge of the actual rate. For $T_c = 15$ minutes, there is some mismatch for some estimates, but it does not affect much the operation, while for $T_c = 30$ minutes, the average arrival rate is a bad predictor for the actual arrival rate.

The lower plots in figures 10 and 11 show the same metrics when the simple estimators described in Sec. 5 are used. In this case, we expect an increase of utilization, due to the fact that the controller will usually underestimate λ_{k+1} for $t < 10.00$ am (because the arrival rate keeps increasing), so it will select a too high probability p_{k+1} allowing the activation of a number of appliances larger than the correct value. This error has larger consequences for small q , when a larger percentage of appliances is activated through the probabilistic control. For T_c different from 30 minutes, there is only a slight increase in the utilization, but still it has a significant effect on the overload probability. In particular for $q < 0.1$ the constraint is no more satisfied. For $T_c = 30$ minutes, the controller uses the average rate measured in [9.30am, 10.00am] to estimate the arrival rate during the interval [10.00am, 10.30am] with about a 20% of relative error. This justifies the bad performance achieved with this setting. We observe that, even without using better estimators, one could simply counteract the estimation errors, for example by reducing p_{k+1} by a given factor corresponding to the maximum variability of the arrival rate from a control slot to the following one.

We also show the curves for the time-average utilization (Fig. 12 and Fig. 13), and overload probability (Fig. 14 and Fig. 15) over the whole control interval (for both the cases of exact and estimated λ_{k+1}).

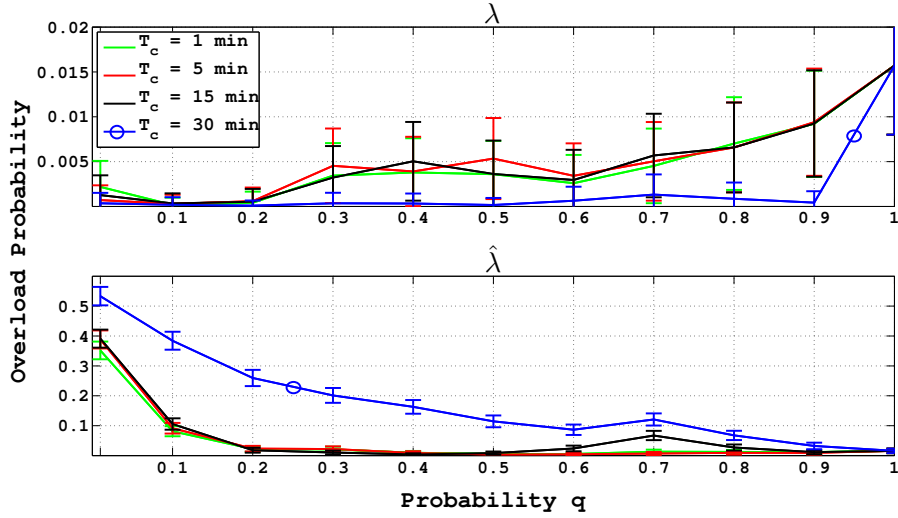


Figure 11: Probability to exceed \bar{P} at 10.26am.

A final aspect to evaluate is how the communication overhead at the controller depends on the request arrival rate, the control interval $[T_{sc}, T_e]$ and the parameters q and T_c . Here we refer to the scenario when the controller is deployed at the DSO and broadcast functionalities are available at data concentrators (as discussed in Sec. 3). Let η be the number of data concentrators. Each query forwarded to the controller during the interval $[T_{sc}, T_e]$ generates an ACK message and during this interval the controller needs to send one message every T_c to each of the η concentrators. Then, the expected number of messages at the controller is:

$$2 \int_{T_{sc}}^{T_e} q\lambda(t)dt + \frac{T_e - T_{sc}}{T_c} \eta \quad (15)$$

An alternative solution is to rely only on queries. In this case the expected number of queries is $2 \int_{T_{sc}}^{T_e} \lambda(t)dt + \eta$, because all the queries will be forwarded to the controller and one message needs to be sent to all the concentrators (and then broadcast to all the appliances) to specify the interval of control.⁵ Fig. 16 plots the expected number of messages obtained from Eq. (15) for different values of the number of concentrators η and $T_c = 15$ minutes, as well as for the solution relying only on queries (the horizontal line). Different values of T_c would simply shift up or down the costs without changing the slope. The intersection between one of the oblique lines and the horizontal one determines the q value for which the two approaches have the same communication overhead. For smaller values of q our solution is more advantageous in this regard, and obviously in terms of privacy.

7 Conclusion

In this report we propose a DLC scheme for smart grids that can work with a large number of dumb appliances, and can be configured for providing the desired trade-off between resource utilization, communication overhead and privacy. The control system has been designed by using

⁵ If the interval is not specified the appliances would need to query the controller all the day along, with an expected number of messages $2 \int_{\text{one day}} \lambda(t)dt + \eta$.

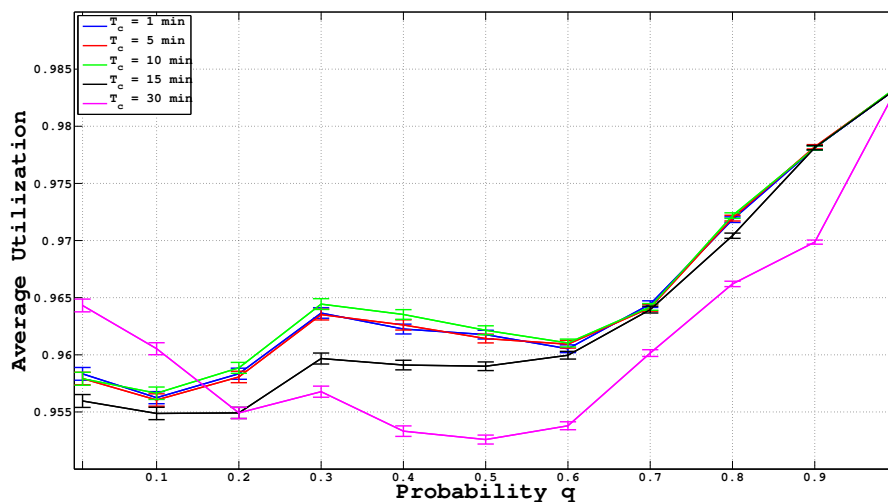


Figure 12: Average resources utilization over the entire control interval $[T_s, T_e]$ with known λ .

well-established teletraffic tools for admission control, which work well for large scale systems, in order to provide probabilistic guarantees on the power demand in the controlled zone.

The main idea is controlling a class of electric appliances by combining a centralized query-response system with a probabilistic system (periodically programmed by the central controller) able to take local decisions. The highest is the ratio between the activation requests that are processed locally and the ones that are forwarded to the controller, the lowest are the information leakage on user habits and the communication overheads, at the expense of control efficiency (in terms of resource utilization and probability to exceed the required power demand).

References

- [1] L. Mathieu, S. Koch, and D. Callaway, “State estimation and control of electric loads to manage real-time energy imbalance,” *Power Systems, IEEE Transactions on*, vol. 28, no. 1, pp. 430–440, 2013.
- [2] H. Hao, Y. Lin, A. Kowli, P. Barooah, and S. Meyn, “Ancillary service to the grid through control of fans in commercial building hvac systems,” *IEEE Transactions on Smart Grid*, 2013, under review.
- [3] DLC programs of Indiana Utilities. [Online]. Available: <http://www.in.gov/oucc/2391.htm>
- [4] W. Kempton and S. E. Letendre, “Electric vehicles as a new power source for electric utilities,” *Transportation Research Part D: Transport and Environment*, vol. 2, no. 3, pp. 157 – 175, 1997. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1361920997000011>
- [5] S. P. Meyn, P. Barooah, A. Busić, and J. Ehren, “Ancillary service to the grid from deferrable loads: The case for intelligent pool pumps in florida,” in *IEEE CDC*, 2013, pp. 6946–6953.
- [6] D. Callaway and I. Hiskens, “Achieving controllability of electric loads,” *Proceedings of the IEEE*, vol. 99, no. 1, pp. 184–199, Jan 2011.




Figure 13: Average resources utilization over the entire control interval $[T_s, T_e]$ with the estimation of λ .

- [7] H. Farhangi, “The path of the smart grid,” *Power and Energy Magazine, IEEE*, vol. 8, no. 1, pp. 18–28, January 2010.
- [8] E. Ancillotti, R. Bruno, and M. Conti, “The role of communication systems in smart grids: Architectures, technical solutions and research challenges,” *Computer Communications*, vol. 36, no. 17–18, pp. 1665–1697, 2013.
- [9] “Demand-response technology roadmap.” [Online]. Available: http://innovation.edf.com/fichiers/fckeditor/Commun/Innovation/conference/Dherbecourt-Technology_Roadmap_DR_28th_May_2009_V1_1_va.pdf
- [10] Y. Y. Hsu and C. C. Su, “Dispatch of direct load control using dynamic programming,” *IEEE Transactions on Power Systems*, vol. 6, no. 3, pp. 1056–1061, 1991.
- [11] K. Bhattacharyya and M. L. Crow, “A fuzzy logic based approach to direct load control,” *IEEE Transactions on Power Systems*, vol. 11, no. 2, pp. 708–714, 1996.
- [12] G. B. Sheble and K. H. Ng, “Direct load control—a profit-based load management using linear programming,” *IEEE Transactions on Power Systems*, vol. 13, no. 2, pp. 668–694, 1998.
- [13] Y. Ghiassi-Farrokhfal, S. Keshav, and C. Rosenberg, “Toward a realistic performance analysis of storage systems in smart grids,” *Smart Grid, IEEE Transactions on*, vol. PP, no. 99, pp. 1–1, 2014.
- [14] O. Ardakanian, S. Keshav, and C. Rosenberg, “On the use of teletraffic theory in power distribution systems,” in *Future Energy Systems: Where Energy, Computing and Communication Meet (e-Energy), 2012 Third International Conference on*, May 2012, pp. 1–10.
- [15] G. Di Bella, L. Giarré, M. Ippolito, A. Jean-Marie, G. Neglia, and I. Tinnirello, “Modeling Energy Demand Aggregator for Residential Users,” in *IEEE CDC*, 2013, pp. 6280–6285.
- [16] N. Ruiz, I. Cobelo, and O. J., “A Direct Load Control Model for Virtual Power Plant Management,” *IEEE Transactions on Power Systems*, vol. 24, no. 2, pp. 959–966, 2009.

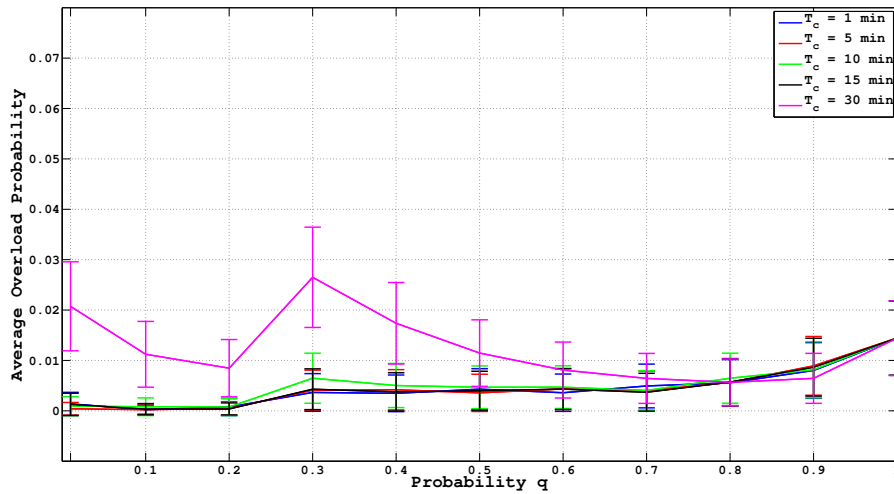


Figure 14: Average probability to exceed \bar{P} over the entire control interval $[T_s, T_e]$ with known λ .

- [17] “An overview to new plc applications and new challenges - maximum integrated products.” [Online]. Available: http://www.ieee-isplc.org/2011/Razazian_talk.pdf
- [18] T. Ganu, D. P. Seetharam, V. Arya, R. Kunnath, J. Hazra, S. A. Husain, L. C. De Silva, and S. Kalyanaraman, “nplug: A smart plug for alleviating peak loads,” in *Proceedings of the 3rd International Conference on Future Energy Systems: Where Energy, Computing and Communication Meet*, ser. e-Energy '12. New York, NY, USA: ACM, 2012, pp. 30:1–30:10.
- [19] F. Kelly and E. Yudovina, *Stochastic Networks*. Cambridge University Press, 2014.
- [20] C. Chang, *Performance Guarantees in Communication Networks*, ser. Performance Guarantees in Communication Networks. Springer London, 2000.
- [21] J. Widén and E. Wäckelgård, “A high-resolution stochastic model of domestic activity patterns and electricity demand,” *Applied Energy*, vol. 87, no. 6, pp. 1880 – 1892, 2010.
- [22] R. Miceli, “Sustainable development and energy saving laboratory.” DIEET - University of Palermo, Tech. Rep., 2007.
- [23] G. Neglia, G. Di Bella, L. Giarré, and I. Tinnirello, “Unidirectional Direct Load Control through Smart Plugs,” in *IEEE 53rd Annual Conference on Decision and Control (CDC 2014)*, Los Angeles, United States, Dec. 2014. [Online]. Available: <https://hal.inria.fr/hal-01094640>
- [24] A. Ganesh, N. O’Connell, and D. Wischik, *Big queues*, ser. Lecture notes in mathematics. Berlin, Heidelberg, Paris: Springer, 2004.
- [25] J. Tao, “Large deviation estimate in ruin theory,” 1997, actuarial Research Clearing House. [Online]. Available: <https://www.soa.org/Library/Research/Actuarial-Research-Clearing-House/1990-99/1997/Arch-1/arch97v120.aspx>

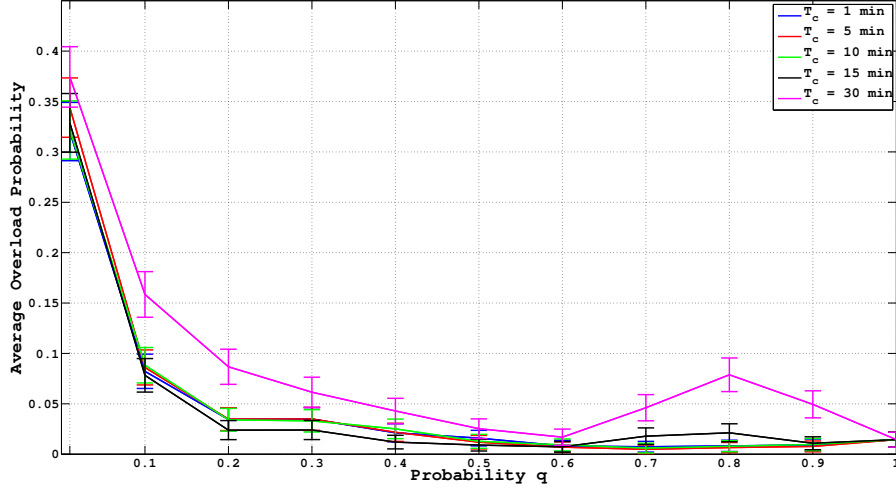


Figure 15: Average probability to exceed \bar{P} over the entire control interval $[T_s, T_e]$ with the estimation of λ .

A Alternative Proof of Theorem 1

Let $Y_n = \sum_{i=1}^n X_i$ and $W_n = \sum_{j=n+1}^{n+N} X_j$ where N is a Poisson r.v. with parameter $n\Lambda_0$ (and then W_n is a compound Poisson r.v.). It holds $M_{Y_n}(s) = nM_X(s)$ and $M_{W_n}(s) = n\Lambda_c^0(e^{M_X(s)} - 1)$. Define also $Z_n = Y_n + W_n$ and let $f_1(s) = M_X(s)$ and $f_2(s) = \Lambda_c^0(e^{M_X(s)} - 1)$.

We apply Chernoff bound to the variable $Z_n - nc$, we obtain:

$$\begin{aligned}
 \ln \Pr(Z_n - nc > 0) &\leq \inf_{s \geq 0} \ln \mathbb{E} \left[e^{s(Z_n - nc)} \right] \\
 &= \inf_{s \geq 0} [\ln \mathbb{E} [e^{sY_n}] + \ln \mathbb{E} [e^{sW_n}] \\
 &\quad - nsc] \\
 &= \inf_{s \geq 0} [nM_X(s) + n\Lambda_c^0 (e^{M_X(s)} - 1) \\
 &\quad - nsc] \\
 &= n \inf_{s \geq 0} [f_1(s) + f_2(s) - sc].
 \end{aligned}$$

$M_X(s)$ is defined for any s and then derivable (see [24]).

Given that $M_{Z_n}(s) = n(M_X(s) + \Lambda_c^0(e^{M_X(s)} - 1))$, also $M_{Z_n}(s)$ is defined and derivable for any s . $M'_{Z_n}(0) = \mathbb{E}[Z_n] = n(\mathbb{E}[X](1 + \Lambda_c^0)) < nc$ then $M_{Z_n}(s) - c$ is decreasing for $s = 0$. Moreover $M_{Z_n}(s) \geq n\Pr(X > c)e^{sc}$ and then $M_{Z_n}(s) - nsc$ diverges when s diverges. It follows that it exists a finite $s^* > 0$ at which $M_{Z_n}(s) - nsc$ is minimized ($M_{Z_n}(s)$ is also convex, then s^* is unique) and $M'_{Z_n}(s^*) - nc = 0$, i.e. $f'_1(s^*) + f'_2(s^*) = c$.

We can then consider $c_1 = f'_1(s^*)$ and $c_2 = f'_2(s^*)$ and then s^* is a point of minimum for the functions $f_1(s) - sc_1$ and $f_2(s) - sc_2$ and $c_1 + c_2 = c$. Then it holds:

$$\inf_{s \geq 0} [f_1(s) - sc_1] + \inf_{s \geq 0} [f_2(s) - sc_2] = \inf_{s \geq 0} [f_1(s) + f_2(s) - sc].$$

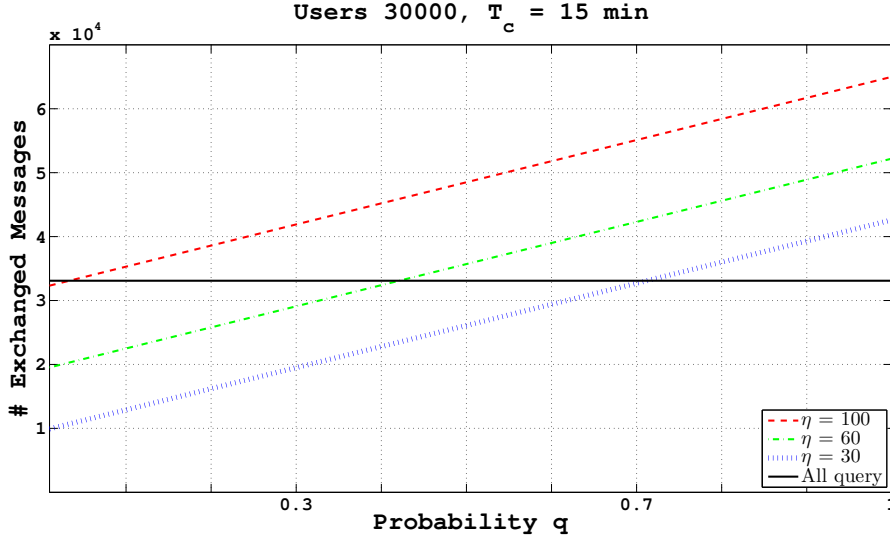


Figure 16: Expected number of exchanged messages for different numbers of data concentrators.

Observe now that:

$$\begin{aligned} \ln \Pr(Z_n > nc) &= \ln \Pr(Y_n + W_n > nc) \\ &\geq \ln \Pr(Y_n > nc_1) + \ln \Pr(W_n > nc_2). \end{aligned}$$

Because of Cramer's theorem, it holds:

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \ln \Pr(Y_n > nc_1) &= \inf_{s \geq 0} [M_X(s) - sc_1] \\ &= \inf_{s \geq 0} [f_1(s) - sc_1] \end{aligned}$$

A similar large deviation result for a compound Poisson process is in [25] and leads to:

$$\begin{aligned} \lim_{n \rightarrow \infty} \frac{1}{n} \ln \Pr(W_n > nc_2) &= \inf_{s \geq 0} [\Lambda_c^0(e^{M_X(s)} - 1) - sc_2] \\ &= \inf_{s \geq 0} [f_2(s) - sc_2]. \end{aligned}$$

Both the upper and the lower bound for $1/n \ln \Pr(Z_n > nc)$ converge to $\inf_{s \geq 0} [f_1(s) + f_2(s) - sc]$, then it follows:

$$\lim_{n \rightarrow \infty} \frac{1}{n} \ln \Pr(Z_n > nc) = \inf_{s \geq 0} [M_X(s) + \Lambda_c^0(e^{M_X(s)} - 1)].$$

■

B On the Frontiers

If we choose a control setting (p, n) and there are actually n appliances activated in the system through the query-response mechanism, then the expected power consumption is

$$E[P] = (n + p(1 - q)D\lambda)E[X]. \quad (16)$$

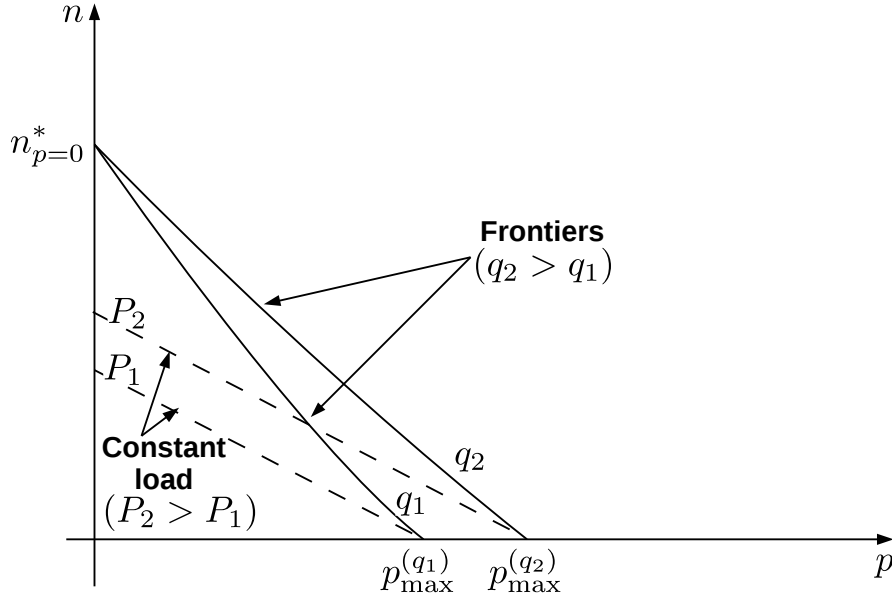


Figure 17: Qualitative behavior of the frontiers. The segments are the geometric loci with constant expected power consumption.

We say that a control setting uses the resources more efficiently (shortly, it is more efficient) than another, if it can sustain a larger expected power demand without violating the bound. The linear approximation for evaluating the frontiers of the acceptance regions given in Eq. (9) leads to interesting consequences about efficiency, because it implies that all the pairs (p, n) on the frontiers have the same efficiency independently from q . This result is counter-intuitive: we could expect that the query-response control offers a better handle on the total power consumption, because the exact number of working appliances is known, and only the randomness of individual power absorption determines the variability of power consumption. On the other hand, the probabilistic control has another source of variability, i.e. the number of appliances actually admitted (that is a Poisson random variable). In this section we are going to show that this is the case, and the independency is simply and artefact of the linear approximation.

The first observation is that the frontiers are not segments but in any case they are convex curves. Indeed, it can be shown from Eq. (7) that the acceptance region is concave because its complement is the intersection of a family of half-planes and then convex. Moreover, for $p = 0$ the maximum number of appliances that can be admitted does not depend on q , then all the frontiers intersect the y -axis in the same point $(0, n_{p=0}^*)$. Fig. 17 shows then a qualitative representation of what the frontiers look like. The loci with constant expected power consumption are characterized by

$$n + p(1 - q)D\lambda = \text{const.}$$

We compare the slope of such segments $-(1 - q)D\lambda$ with the slopes of the tangents to the frontiers. Given a pair (p, n) on a frontier, let $s^* \geq 0$ be a value at which the infimum of $nM_X(s) + p(1 - q)\lambda D(e^{M_X(s)} - 1)$ is attained. The tangent to the frontier in (p, n) has slope:

$$-(1 - q)\lambda D \frac{e^{M_X(s^*)} - 1}{M_X(s^*)} \leq -(1 - q)\lambda D, \quad (17)$$

because for any x it holds $e^x \geq x + 1$. Taking into account these considerations, we have plotted in Fig. 17 two loci with constant power demand (dashed lines). The figure confirms then our intuition: for a given frontier (a given q value), the larger n , the higher the efficiency. Despite all these considerations, as Fig. 17 shows and Eq. (9) justifies analytically, this effect is negligible.

There is another aspect to be mentioned. We have assumed above that by selecting a pair (p, n) we can actually have n appliances activated through the query-response system, but this may not be the case if q is small. In particular if $q\lambda D \ll n$, the number of appliances activated through the query-response system will be smaller than n most of the time. Observe in fact that $q\lambda D$ is the expected number of requests to the controller during an interval n . If such number is much smaller than n , the cap will be ineffective most of the time. Practically speaking the maximum value of n to be considered for a given q is $q\lambda D$. This remark is reflected in our way to set n , see Eq. (12).

We conclude this appendix with some additional consequences of Eq. (17) and the convex shape of the frontiers on the use of the linear approximation. If we consider a segment with slope $-(1 - q)D\lambda$ passing by a point $(p_{\max}^{(q)}, 0)$, the segment is guaranteed to lie inside the acceptable region. On the contrary, a segment with the same slope passing by $(0, n_{p=0}^*)$, would lie outside the acceptable region (even if very close to it). This is the reason why in Sec. 5 we have introduced $p_{\max}^{(q)}$ to configure the system instead to rely on Eq. (9).

C Equivalence between the Energy-Bucket and the Queue

We show the equivalence for generic time-variant input/output rates. In particular we replace \bar{P} with $T(t)$.

We write the equations for the two dynamics:

$$\frac{dL(t)}{dt} = \begin{cases} (T(t) - P(t))^+ & L(t) = 0 \\ T(t) - P(t) & 0 \leq L(t) \leq B \\ (P(t) - T(t))^+ & L(t) = B \end{cases} \quad (18)$$

$$\frac{dQ(t)}{dt} = \begin{cases} (P(t) - T(t))^+ & Q(t) = 0 \\ P(t) - T(t) & 0 \leq Q(t) \leq B \\ (T(t) - P(t))^+ & Q(t) = B \end{cases} \quad (19)$$

If $L(0) = L_0$ and $B(0) = B - L_0$, it can be easily checked that it will be $\frac{dQ(t)}{dt} = -\frac{dL(t)}{dt}$ for any $t \geq 0$. Then if $L(t)$ is the solution of (18) with $L(0) = L_0$, $Q(t) = B - L(t)$ is the solution of (19) with $Q(0) = B - L_0$. ■

Contents

1	Introduction	3
2	Related Work	3
3	System Description	5
4	Control in a Stationary Setting	6
4.1	Classic large deviation results for $q = 1$	7
4.2	Extension to $q < 1$	8
4.3	A numerical example	9
4.4	A different constraint	11

5	Control in a Time-Variant Setting	12
5.1	The plant	13
5.2	The controller	13
6	Numerical Results	16
7	Conclusion	18
A	Alternative Proof of Theorem 1	22
B	On the Frontiers	23
C	Equivalence between the Energy-Bucket and the Queue	25



**RESEARCH CENTRE
SOPHIA ANTIPOLIS – MÉDITERRANÉE**

2004 route des Lucioles - BP 93
06902 Sophia Antipolis Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399