

Congestion Avoidance in Low-Voltage Networks by using the Advanced Metering Infrastructure

Benoît Vinot, Florent Cadoux, Nicolas Gast, Rodolphe Heliot, Victor Guin

► **To cite this version:**

Benoît Vinot, Florent Cadoux, Nicolas Gast, Rodolphe Heliot, Victor Guin. Congestion Avoidance in Low-Voltage Networks by using the Advanced Metering Infrastructure. 2018. hal-01784386

HAL Id: hal-01784386

<https://hal.inria.fr/hal-01784386>

Submitted on 4 May 2018

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.

Congestion Avoidance in Low-Voltage Networks by using the Advanced Metering Infrastructure

Benoît Vinot
Schneider-Electric
Grenoble Cedex 9, France
benoit.vinot@schneider-electric.
com

Florent Cadoux*
Univ. Grenoble Alpes
CNRS, Grenoble INP, G2ELab
Grenoble, France
florent.cadoux@grenoble-inp.fr

Nicolas Gast
Univ. Grenoble Alpes
CNRS, Inria, Grenoble INP, LIG
Grenoble, France
nicolas.gast@inria.fr

Rodolphe Héliot
Schneider-Electric
Grenoble Cedex 9, France
rodolphe.heliot@schneider-electric.
com

Victor Gouin
Univ. Grenoble Alpes
CNRS, Grenoble INP, G2ELab
Grenoble, France
victor.gouin@grenoble-inp.fr

ABSTRACT

Smart Grids aim at leveraging communication and computation capabilities to better operate and control electrical networks. The idea of controlling the output of decentralized photovoltaic (PV) generators so as to avoid current and/or voltage issues in the distribution grid is, in particular, thoroughly investigated by the research community. The simplest and most commonly considered control methods consist in using a feed-back controller based on a local voltage measurement at PV generator level. The main advantages of such methods is that they can be implemented at low cost and require no specific information about the network on which they are deployed. In this paper, we consider an alternative, that uses the advanced metering infrastructure (AMI) as the basis for PV generation control. Our contributions are as follows. First, we cast the problem of PV generation curtailment in the general framework of control theory, which sheds some light on the structure of previously proposed controllers. Second, we show how the advanced metering infrastructure may be used on one hand to infer some knowledge about the underlying network (learning phase), and on the other hand, to observe some of the disturbances that apply to our system so that they can be taken into account in the control phase. While doing so, we take great care of accounting for the inherent limitations of the advanced metering infrastructure in terms of measurements and communication, which is a major difference between this paper and previous work on the topic. Third, by means of numerical simulations, we compare our proposed feed-forward controller with two other controller structures: open-loop, and feed-back. We demonstrate that our feed-forward controller — that requires no prior knowledge of the underlying electrical network — brings significant performance improvements as it can effectively suppress over-voltage and over-current while requiring low power curtailment.

*Grenoble INP: Institute of Engineering Univ. Grenoble Alpes

CCS CONCEPTS

• **Hardware** → **Smart grid**; *Power networks*; • **Mathematics of computing** → *Linear programming*; • **Networks** → *Network control algorithms*;

ACM Reference Format:

Benoît Vinot, Florent Cadoux, Nicolas Gast, Rodolphe Héliot, and Victor Gouin. 2018. Congestion Avoidance in Low-Voltage Networks by using the Advanced Metering Infrastructure. In *Proceedings of ACM e-Energy*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 INTRODUCTION

Advanced metering infrastructures (AMIs) are currently deployed in many countries [1]. In addition to streamlining the billing process, the rationale is that these infrastructures will help mitigating the grid congestions caused by the rise of distributed generation, improve grid observability, and provide new customer services. For example when local generation is high and demand is low, distribution grids may experience overloads in transformers and possibly in lines, and voltage excursions outside of the allowed voltage range. Among these issues, voltage excursions and in particular over-voltages are the major problem of embedded generation on low-voltage distribution networks [2].

The most standard solution to solve this kind of issues is grid reinforcement, which involves replacing or adding lines and/or transformers. Grid reinforcement comes at a relatively high cost, and is not always the best option [3]. As a consequence, many alternate solutions have been suggested by the industry and the research community, emphasising the need for more *active* distribution networks. One of these solutions is to equip medium-to-low-voltage (MV/LV) transformers with an on-load tap-changer (OLTC) [3, 4, 5] or some other type of voltage-regulating device. This solution does involve some additional hardware, yet may not be as expensive as “fully” reinforcing the grid.

Another solution is to leverage the so-called flexibility of loads and generators, namely, their ability to modify their (active or reactive) power consumption or output whenever

this is needed to mitigate grid congestions. This approach is appealing since it could potentially eradicate congestions without CAPital EXpenditure (CAPEX) cost for the Distribution System Operator (DSO). Various methods have thus been proposed along those lines, involving loads and/or generators, active or reactive power, and local or centralised control. One of the most notorious of these numerous options is probably to locally control the reactive power output of PV inverters based on local measurements — *a.k.a.* $Q(U)$ controller — in order to flatten the voltage profile of the feeder [6, 7], but other approaches have received some attention as well.

In this paper, we focus on controlling the *active* power output of PV inverters as in, for instance, [8]. The drawback of this choice is that PV producers may suffer from a loss of profit proportional to the amount of PV energy curtailed. There are, however, strong arguments in favour of active power control:

- it does not involve additional CAPEX cost contrary to grid reinforcement and OLTC-like technologies;
- in a low voltage (LV) network, the effect of active power on voltage is usually several times stronger than the effect of reactive power;
- active power curtailment makes it possible to deal with current constraints (overloads), not only with voltage problems — contrary to reactive power control and OLTC-like technologies;
- finally, active power curtailment will occur only on the rare occasions where solar irradiance is maximal while local consumption is minimal, so that the loss-of-profit of PV producers will be limited.

Also, in this paper, we consider the problem of controlling the output of PV generators in a low-voltage network *by leveraging the advanced metering infrastructure* (AMI). The AMI capabilities are used at two different levels. First, we do not assume any knowledge of the details of the underlying electrical network (topology, characteristics of the lines, etc), but we *infer* from AMI data the information that we need for control purposes. Second, we view the AMI as a measurement and communication platform on which a centralised controller is deployed. This leads us to formulate the problem of optimal energy curtailment as an optimisation problem in which the network constraints are learnt from past historic AMI data, and in which some of the parameters (termed “disturbances” below) that enter the optimisation problem are obtained from the most recent available AMI measurements. We end up with a feed-forward controller whose input is the latest data measured by the AMI, and that sends back to each producer a maximal generation “quota”.

We then evaluate the performance of our controller by means of numerical simulations. These experiments use network data from the “Low Voltage Network Solutions” project [9] and simulate the exact three-phased load-flow equations. We compare our proposed feed-forward controller with popular alternatives such as open-loop controllers and pure local feed-back controllers like $P(U)$ or $Q(U)$. For comparison

purpose with our AMI-based feed-forward controller, we also introduce an AMI-based feed-back controller.

Our main conclusions are the following. First, we demonstrate that the relative lateness and inaccuracy of AMI-measured data, and the long duration of control time-steps, which are inherent to the limitations of the AMI infrastructure, have a completely different impact on AMI-based feed-back and feed-forward controllers. Namely:

- in spite of these two impediments, our AMI-based feed-forward controller performs well,
- whereas AMI-based feed-back control significantly suffers from communication limitations and performs very poorly.

Second, in spite of AMI limitations, our centralised feed-forward controller performs better, overall, than local controllers (that do not use the AMI and are thus oblivious to communication problems). Indeed, it is the only controller that is able to suppress most of the over-voltage or over-power constraints, and appears as an appealing method to mitigate congestions created by PV generators without reinforcement and with limited solar energy curtailment.

Road-map. This paper is organised as follows. In Section 2, we review some of the related work. In Section 3, we first formulate the mathematical problem that we aim to solve and discuss design choices that pertain to the structure of the controller. In Section 4, we detail several different controller implementations that are then compared by simulation in Section 5. Finally, Section 6 gathers some conclusions from our study and suggests some directions for future research.

2 RELATED WORK

In this section, we review some of the approaches that were previously suggested in the literature for the PV control problem. Since the controller proposed in this paper is feed-forward, the literature review below focuses on previous works that also used a feed-forward logic (although the “feed-forward” terminology is often not explicitly used by the authors themselves). On the contrary, we do not review the abundant literature about pure feed-back controllers. See Figure 1 for a quick over-view of the differences between open-loop, feed-back and feed-forward, and Section 3.3 for more details.

In [10], the authors propose a centralised controller structure based on a linear program (LP) that sets the value of controllable devices such as flexible PV generation. It is worth noting that the power output of non-flexible devices, such as loads, enters as a parameter in the LP formulation. In our words, this is a feed-forward controller in which the disturbance (here, the power output of non-flexible devices) is measured instantaneously and with perfect accuracy — an optimistic assumption that is questioned in the present paper.

In [11], five different control methods are presented. The so-called “zero current injection” method may also be cast in the realm of feed-forward methods, where the consumption of each “prosumer” (customer with both load and generation)

is measured locally and local generation is controlled so as to keep the net power injection below a certain threshold (which is 0 in [11]). Contrary to the method presented in [10], the “zero current injection” method only requires local (domestic-level) communication and is thus decentralised from the grid viewpoint. The other four methods presented in [11] are purely feed-back.

The pros and cons of local (domestic-level) controllers as in [11], and centralised (grid-level) controllers as in [10], in the context of PV generation control, is discussed in [7]. The obvious advantage of centralised control is the benefit of coordination, and its obvious drawback is the cost and complexity of communication. What is less obvious is the impact on the *speed* of the controller when choosing centralised or decentralised control: indeed, if the communication channel is relatively slow, centralised control might itself be too slow to be practical. This is particularly relevant in the context of PV generation control: indeed, the only communication channel generally available today to utilities is the AMI, which usually has limited communication capabilities. This fact is well acknowledged in [7], which states: “An important design decision *that weighs on the speed and quality of communications* required is whether the control should be centralised or distributed (*ie.*, local)”. In spirit, [7] is thus relatively close to the work presented in this paper, but the design choices made by the authors are different and almost opposite: [7] uses reactive power control, whereas we advocate active power control; the controller structure is local in [7] and centralised in the present paper; and the controller logic is relatively straightforward in [7], with the advantage of simplicity, whereas we use a more elaborate logic that requires some knowledge of the underlying grid (this knowledge being here the result of a learning phase, see Section 4.4.1).

The authors of [12] make a step further by proposing a hybrid controller that uses both the feed-back and the feed-forward logic simultaneously. Indeed, the controller makes use of both:

- active and reactive power values, which are disturbances as far as non-flexible devices are concerned,
- and voltage values, which are outputs.

The controller structure presented in [12] is quite involved, and uses in particular a state estimator instead of direct measurements of observable variables and disturbances, but the idea of hybridising feed-back and feed-forward control could be applied to simpler controllers as well.

All in all, some papers from the literature — such as those presented above — did already propose feed-forward controllers for the PV generation control problem, and with several flavours (centralised or decentralised, purely feed-forward or hybrid). But the inherent limitations of the AMI, seen as a measurement and communication platform for control purpose, and their impact on the efficiency of various controller structures, have received little attention so far. This is the topic addressed by the present paper.

3 MODEL

3.1 Description of the system

3.1.1 Network, loads and generators. We consider a three-phase (AC) low voltage network connected to an upstream medium voltage (MV) network. We assume that the voltage at the point of coupling with the MV network is fixed, and the loads are all modelled as constant (complex) power, single-phase loads. We assume that all PV generators are controllable, whereas loads are not. Finally, we consider over-voltage constraints and over-power flowing through the MV/LV transformer. These assumptions are only meant to simplify the mathematical formulation of the equations that will appear in the sequel of this paper: our arguments can be generalised to encompass, for instance, situations where the load is controllable as well; or to take into account under-voltage, or other current constraints.

We will study the network over a certain period of time. For each load, a load curve is thus defined, as well as a curve of “producible power” for each PV generator. The curve of producible power is defined as the power that would be produced by the considered generator if no active power control was used. In our simulations below, both types of curves are defined with a fixed time-step of 2 minutes.

3.1.2 Measurements and communications. The upstream MV voltage, the active load and the solar irradiance are measured, and considered as observed disturbances over which we have no control. Voltage magnitude at customer nodes, which is the output of the system, is also measured, whereas reactive power is not.

Active power and voltage measurements, as well as telecommunications, are handled by the AMI. In our simulations below, we consider that voltage and energy measurements are taken over a duration of 10 minutes [13], then communicated to the central controller. Conversely, setpoints that are sent from the centralised controller are applied for a duration of 10 minutes. To the best of our knowledge, this update rate (once every 10 minutes) is in the correct order of magnitude of what could be achieved in practice given the current performance of existing metering technologies. Because the state of the electrical network evolves much faster (with 2 minute time-steps in our simulation below), centralised control intrinsically suffers from using relatively obsolete data.

This important issue was acknowledged, to some extent, in [10], which states that a margin of safety (using a -7% voltage limit instead of the actual -10% limit) should be used in practice to make up for “any unexpected short-term variations in demand”. This relevant argument however remains purely theoretical, and neither the approach presented in [10] nor — as far as we know — other previous related research, actually attempted to explicitly capture such “short-term variations” in the simulations. The present paper is thus the first in which these limitations of the AMI are taken into account.

3.1.3 Controller. The central controller gathers measurements and sends back quotas at fixed time intervals, here every 10 minutes. For a given PV generator g , we denote by p_g^{\max} its power rating (*i.e.* the maximal power that g can produce). We denote by $\alpha_g(t) \in [0, p_g^{\max}]$ the generation quota sent to the generator g at time t . In vector notation, the vector p^{\max} fixes an upper bound on vector $\alpha(t)$ at all times t . We wish to design a controller that achieves two goals:

- first, the controller ensures that the allowable voltage and power will not be excessively violated (over-voltages and over-powers through the transformer);
- second, the controller strives to maximise the amount of generation quotas (within the limit $\alpha(t) \leq p^{\max}$) sent to PV generators, in order to maximise the amount of PV energy actually produced.

The first item above imposes a set of constraints, while the second is in the form of an objective function; mathematically, we will thus formulate the logic of our controller as an optimisation problem.

3.2 Parameters, variables and constraints

We now formulate the problem in mathematical notations. Let \mathbf{L} be the set of loads, and \mathbf{G} the set of generators. For $l \in \mathbf{L}$ and for $g \in \mathbf{G}$, we denote by $p_l \geq 0$ and $p_g \leq 0$ respectively the active power consumed and produced by the considered customer. The active power $p_g(t)$, generated by generator g , is the minimum between the producible power $\hat{p}_g(t)$ (which is determined by the current value of solar irradiance), and the power quota $\alpha_g(t)$ sent by the central controller. As both values are negative, this rewrites as:

$$p_g(t) = \max(\hat{p}_g(t), \alpha_g(t)). \quad (1)$$

We also introduce the phase-to-neutral voltage magnitude, which we denote by u_l ($l \in \mathbf{L}$) for loads, and u_g ($g \in \mathbf{G}$) for generators. The active power flowing through the transformer of the given electrical LV network is denoted by p_{trans} .

For convenience, we define the vector of all active power outputs

$$p(t) := ((p_l(t))_{l \in \mathbf{L}}, (p_g(t))_{g \in \mathbf{G}}); \quad (2)$$

and the vector of all phase-to-neutral voltage magnitudes

$$u(t) := ((u_l(t))_{l \in \mathbf{L}}, (u_g(t))_{g \in \mathbf{G}}),$$

as well as the vector of power generation quotas

$$\alpha(t) := (\alpha_g(t))_{g \in \mathbf{G}}.$$

Note that control is here unilateral: variable $p_g(t)$ may only be upper-bounded, not fixed. For simplicity, the voltage at the point of common coupling with the upstream MV network is set to a fixed value, which reduces the list of (meaningful) observed disturbances to active power load on one hand, and producible power on the other hand. Our approach would remain the same if a non-constant time-series was used for the upstream voltage.

Active and reactive power consumption (or generation) at all customer nodes, along with the network structure and

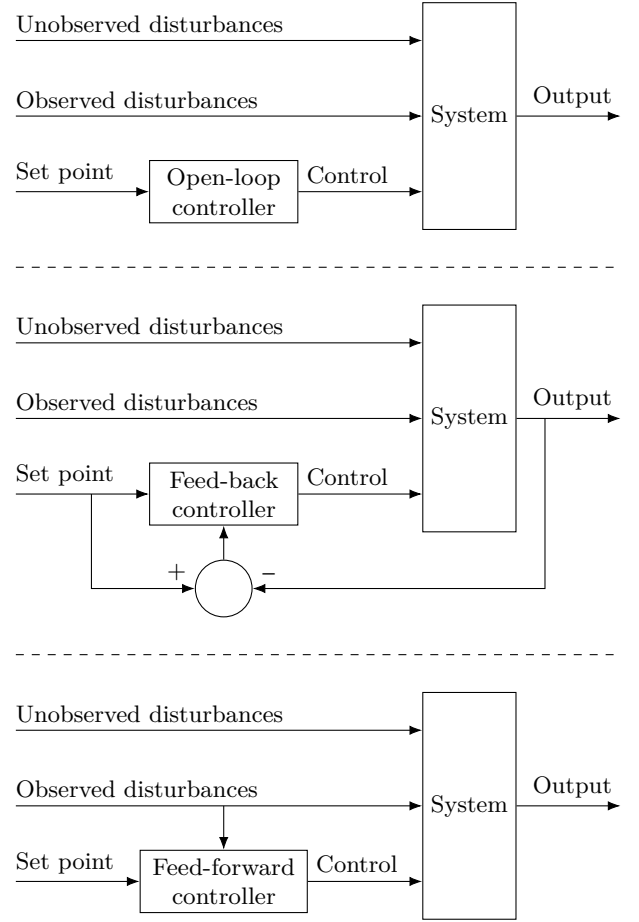


Figure 1: open-loop, feed-back and feed-forward control.

upstream MV voltage, determine the complex voltages everywhere in the network through the (three-phase, unbalanced) load flow equations. These equations are standard [14] and are omitted here for the sake of concision.

The constraints that must be taken into account derive from the upper voltage limit:

$$u(t) \leq u^{\max}, \quad \forall t, \quad (3)$$

where u^{\max} is the maximum allowed phase-to-neutral voltage magnitude (*eg.* 7% or 10% above the nominal value of voltage) and the maximum allowed active power at the transformer level:

$$|p_{\text{trans}}(t)| \leq p_{\text{trans}}^{\max}. \quad (4)$$

Note that in theory, we should also consider a lower voltage limit constraint $u^{\min} \leq u(t)$. In practice, however, this constraint is never violated because the congestion in our experiments are caused by over-production and not over-consumption. This is why in the remainder of the paper, we only consider upper voltage limit.

3.3 Controller structures: generalities

Let us consider an abstract system whose state evolves under the effect of two kinds of inputs, namely the *control* variable and the (observed or unobserved) *disturbances*, and which produces some output. Generally speaking, the goal of control is usually to get the system output to follow a desired trajectory; as noted above however, no such “desired trajectory” is predefined in the context of PV curtailment: the aim of control is here better expressed in the language of optimisation. The control objectives being defined, three broad categories of controllers may be used to achieve them, as illustrated by Figure 1:

- open-loop control, which consists in setting the value of the control variable regardless of the value of both disturbances and system output;
- feed-back control, which consists in observing the output of the system and choosing the value of the control variable accordingly;
- and feed-forward control, which consists in measuring the value of the (observed) disturbances and to choose the value of the control variable that, *according to our model of the system*, is the best response to these disturbances.

The above-mentioned notion of “model of the system” is important: while tuning the design parameters of the simplest controllers (for instance a simple PID feed-back controller) may be achieved based only on minimal information about the system, more elaborate controllers may on the contrary need a richer model of the controlled system. The various ways to devise such a rich model for our PV generation control problem are discussed in Section 4.4.1 below. The next subsection will clarify what “control variable”, “disturbance” and “output” mean in our context. Also note that the feed-back and feed-forward methods may be combined (and are often indeed combined, in practice); in this case, the controller uses measurements of both the current disturbances and the current output.

3.4 Application to the PV generator output control problem

The system that we are willing to control is a low-voltage network involving flexible PV generators. From the electrical viewpoint, the system is considered to be in quasi-steady state. Our control variable is a “generation quota” that we will send to flexible generators, that is to say, a maximum power output that each flexible generator will not be allowed to exceed. Keeping in mind our general control theory framework, we consider that the state of the system is determined on the one hand by this control variable, and on the other hand by some “disturbances”, among which the following are considered observable:

- voltage at root node,
- current active power consumption of loads (and current active power output of non-flexible generators, if any);
- and current solar irradiance.

In order to make practical use of the value of these disturbances in a feed-forward controller, we need to measure them. In this paper, we make the following assumptions.

- First, we assume that the voltage magnitude at root node is observed thanks to some dedicated sensor inside the MV/LV substation.
- We assume that the active power consumption of loads may be measured *through the AMI*, while acknowledging the inherent limitations, in terms of measurement and communication, of this infrastructure.
- We assume that solar irradiance may be either measured directly (in case equipping MV/LV substations with some kind of light sensor would be seen as impractical) or estimated thanks to recent measurements of the output of the non-curtailed PV generators.

These disturbances, in particular the value of non-flexible load and solar irradiance, may change very quickly compared to the speed of our controller; this issue is discussed in the present paper.

All other potential disturbances are considered as unobserved. This applies in particular to the reactive power consumption of loads. Arguably, this data could also be measured through the AMI just like active power, and thus be considered as observable. However the reactive consumption of loads is much smaller than their active consumption [15], and in addition, the effect on voltage of 1 kVAr of reactive power in an LV network is itself much smaller than the effect of 1 kW of active power (due to the high R/X ratio of LV lines), so that we dispense from taking into account reactive power as an observed disturbance. Finally, the output of the system is the set of (phase-to-neutral) voltage magnitudes at all customer nodes in the network, and the value of power that flows through the transformer. Just like the power consumption of loads, we assume that voltage at customer nodes may be observed thanks to the AMI. The power flow through the transformer is assumed to be measured thanks to a dedicated sensor.

We now have cast our PV curtailment problem into the general framework of control theory, except for one thing: the aim of control was not defined yet. As already discussed, the main point is here, first to eradicate over-voltages, and second to produce as much solar energy as possible. Secondary objectives, such as fairness between producers, may be introduced, but these are not considered in the present paper.

3.5 Controller structures for PV generation curtailment

Having formulated the PV curtailment problem in the language of control theory in Section 3.4, and following our previous discussion about controller structures in Section 3.3, we may control PV output using:

- an open-loop controller, that would consist in permanently limiting the output of PV panels (for instance, no more than 80% of the peak power of the PV modules should be produced at any time). This option is quite brutal and could potentially lead to curtail generation

even in the absence of any grid constraint; but it is simple and reliable.

- A feed-back controller, that would read measurements of voltage at customer nodes and power flow at transformer level, and react to over-voltages by somehow adjusting the generation quota of flexible generators.
- Or a feed-forward controller, that would adjust the generation quotas based on measurements of solar irradiance, voltage at root node and current power consumption of loads.

In this paper, we focus on controllers that leverage the AMI, and argue that feed-forward control is much better suited than feed-back control in this context.

3.6 Which model of the system should our feed-forward controller embed?

In practice, two options are available: either we consider that the network data (topology, line impedances, phase connection of customers, *etc*) is known; or we attempt to learn the model from past measurements. In practice, reliable network data at LV level is scarce; on the contrary, under our assumption that the control system is implemented on top of the AMI, historical data will be plentiful. As a consequence, we assume that the feed-forward controller feeds past measurements of (active) power and (phase-to-neutral) voltage into a learning algorithm in order to infer how to respect network constraints. In our experiments below, we use a simple linear regression (see the discussion around Equation (7)). As this simple model gives us satisfactory results, we do not discuss further this topic in this paper and left the search for the best learning algorithm for future work.

4 CONTROL STRATEGIES

In this section, we now describe the different controllers that we will compare numerically in Section 5.

4.1 Open-loop

The open-loop controller is the simplest of all. It is parameterised by a value $\theta \in [0, 1]$ which indicates how much of its nominal power a generator is allowed to produce. In terms of vector α , we apply the following control by defining $(\alpha_g(t))_{g \in \mathbf{G}}$:

$$\forall g \in \mathbf{G}, \alpha_g(t) := \theta p_g^{\max}$$

Applying the “open-loop 100%” strategy means that PV generation is never curtailed, which amounts to applying no control. The “open-loop 0%” strategy means that PV generators are not allowed to produce at all. The “open-loop 75%” strategy means that the PV are allowed to produce up to 75% of their nominal power. Note that in practice this strategy, that curtails 25% of the available *power*, actually curtails less than 5% of the solar *energy* (see Figure 6) because PV generators rarely produce at their maximal power.

4.2 Inverter-level feed-back control

The pure feed-back methods that we consider are the well-known $P(U)$ and $Q(U)$ controls. The $P(U)$ control for the generators defines the following $(\alpha_g(t))_{g \in \mathbf{G}}$:

$$\forall g \in \mathbf{G}, \alpha_g(t) := \beta(u_g(t)) p_g^{\max}, \quad (5)$$

where β is a function defined in the Figure 2.

The $Q(U)$ control for the generators defines their reactive power output:

$$\forall g \in \mathbf{G}, q_g(t) := \gamma(u_g(t)) q_g^{\max}, \quad (6)$$

where γ is a function defined in the Figure 2.

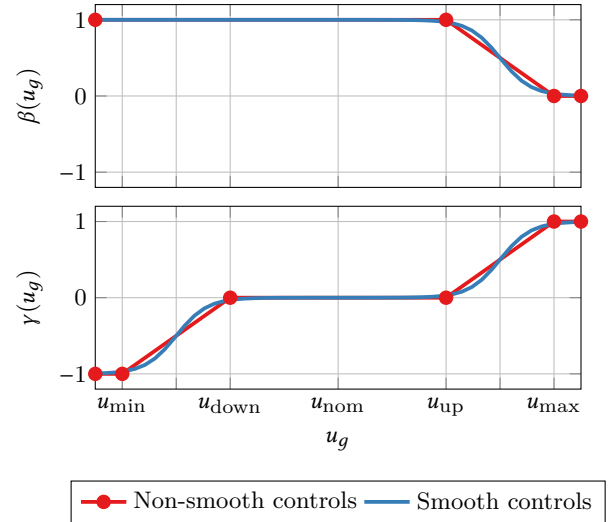


Figure 2: Local control functions used for the feed-back controls. Smooth functions for $P(U)$ and $Q(U)$. Non-smooth functions for the “AMI-based feed-back” control.

In the numerical experiments, we implement these controls directly at the level of our home-made load flow solver by using a standard Newton-Raphson method. Because this method applies to smooth functions, for the numerical experiments, we consider a smooth version (sigmoid function) of the functions β and γ , as depicted in Figure 2.

4.3 AMI-based feed-back control

For the sake of comparison with our AMI-based feed-forward controller, we also implemented an AMI-based feed-back controller. The only difference with inverter-level feedback control is that there is now significant delay between the measurement of voltage (through the AMI) and the application of the control setpoint. Generally speaking, from the viewpoint of control, delays are detrimental and promote in particular instability; this detrimental effect will indeed be observed in our simulations below.

Because function β does not enter the Newton-Raphson solver anymore, its non-smooth version is used, as depicted in Figure 2. The control becomes

$$\forall g \in \mathbf{G}, \alpha_g(t) := \beta(u_g(t)) p_g^{\max}$$

where $u_g(t)$ denotes the latest, and already relatively obsolete, AMI voltage measurement (whereas in Equation (5), $u_g(t)$ was the value of the voltage computed directly inside the load flow algorithm, without any delay).

4.4 Feed-forward

Here we present our feed-forward controller.

4.4.1 Surrogate model of load-flow equations. As mentioned above, DSOs often lack reliable information about their LV networks. As a consequence, we advocate a data-driven approach that consists in using AMI measurements to infer the network model that will be embedded in our centralised controller. In this paper, we assume that a (linear) surrogate model of the following form is available to the central controller:

$$u(t) \approx Ap(t) + b, \quad (7)$$

where matrix A and vector b are determined by a learning algorithm fed with historical smart metering data, and is thus available in practice to the distribution system operator. We also assume that an equivalent (linear) model (C, d) is also available for the active power flowing through the transformer:

$$p_{\text{trans.}}(t) \approx Cp(t) + d. \quad (8)$$

In our simulations below, the learning phase was performed using a simple linear regression approach.

4.4.2 Formulation. We are now ready to design our feed-forward controller. We assume that, before taking our decision, we are given a forecast of the consumption for the next time step that we denote by $(\tilde{p}_l)_{l \in \mathbf{L}}$ and a forecast of the PV generation that we denote by $(\tilde{p}_g)_{g \in \mathbf{G}}$. In practice, in the numerical simulation, for the load we use a persistence forecast: \tilde{p}_l is equal to the consumption during the last time step at this node. For the production, we set \tilde{p}_g as the average between the production during the last time step and p_g^{\max} .

We then choose the quotas α according to the following formulation:

$$\alpha^*(t) = \underset{p_g}{\operatorname{argmax}} - \sum_{g \in \mathbf{G}} p_g \quad (9)$$

subject to

$$\begin{aligned} Ap + b &\leq u^{\max}, \\ -p_{\text{trans.}}^{\max} &\leq Cp + d \leq p_{\text{trans.}}^{\max}, \\ p_l &= \tilde{p}_l \quad (\forall l \in \mathbf{L}), \\ \tilde{p}_g(t) &\leq p_g \leq 0 \quad (\forall g \in \mathbf{G}), \end{aligned}$$

where p denotes the vector of active power as in Equation (2).

The rationale behind the optimisation problem Equation (9) is that we try to maximise the quota allocated to PV generators while respecting the constraints given by the surrogate linear models Equation (7) and Equation (8) and given a forecast value \tilde{p} for the consumption and production during the next time steps. Recall that in our notations, a load p_l is a positive value while a production p_g is a negative value.

This explains why the maximisation of $-\sum_g p_g$ corresponds to maximising the quota of allocated to the PV.

5 NUMERICAL EXPERIMENTS

5.1 Dataset

The data necessary to run the simulations presented in this section is extracted from the database of the ‘‘Low Voltage Network Solutions’’ [9] that was originally introduced to the research community by the university of Manchester. It contains the topology of 25 electrical distribution networks. A single network includes several feeders connected to the same MV/LV substation. The largest feeder has a few hundred clients, and the smallest less than 10. This database also provides several load and generation profiles.

We extracted 21 feeders from this data set, and used the load and PV generation profiles provided. We attributed one of these profiles to each customer on the 21 selected feeders, while respecting the following ratio: 50% of customers were pure consumers, and 50% were ‘‘prosumers’’ (*i.e.* had both load and generation). All customers are single-phased.

5.2 Learning versus control phases

The tests were divided in two steps :

Learning step: we use winter load profiles for consumers; this is to avoid biasing the experiment by using the same dataset for training and for control (the control phase will be carried out with summer profiles). The load profiles peak at approximately 6 kW with a power factor of 0.95. Generators peak at 3 kW with unity power factor. With this data, we simulate a duration of one day, with 2-minute time-steps. This first batch of simulations is performed without any control, and provides a database of historical measurements of (active) power and (phase-to-neutral) voltage on each network. We then apply a simple linear regression to this dataset in order to compute matrices and vectors A, b, C, d of Equations (7) and (8).

Control step: we use summer load profiles, still peaking at 6 kW. The PV profiles now also peak at 6 kW. We re-run the simulation using 2-minute time-steps, and with 10-minute control time-steps. The calculation is run again over 24 hours, once for each of the control strategies below.

5.3 Control strategies

The control strategies are the following.

Open-loops: five open-loops strategies are compared, corresponding to five curtailment levels: generation is curtailed when it exceeds respectively 0%, 25%, 50%, 75% and 100% of the installed capacity. Note that the open-loop controller with a 100% limit does not curtail any generation at all, while the open-loop controller with a 0% limit prevents generation entirely. See Section 4.1.

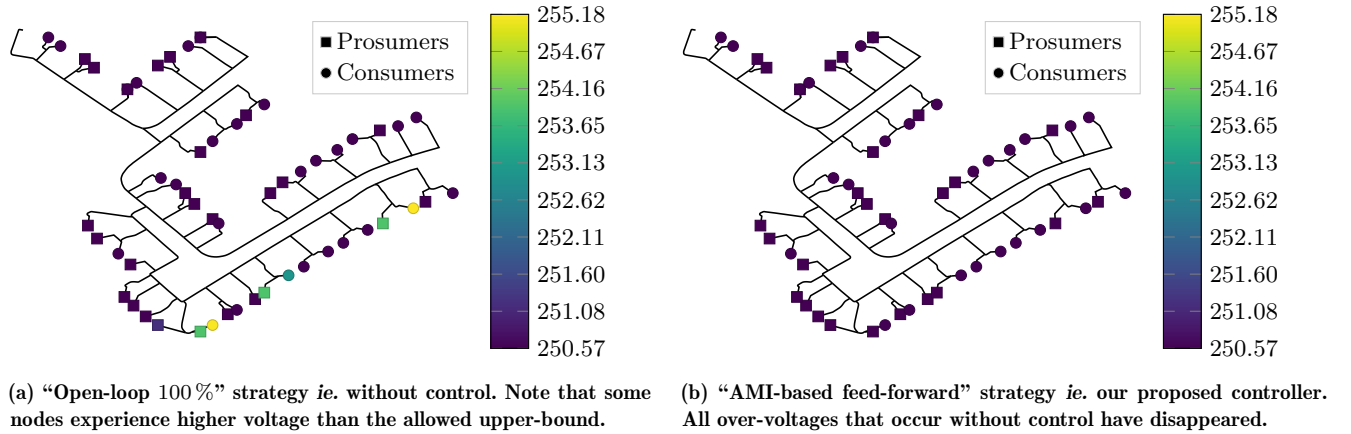


Figure 3: Mean over-voltages for the feeder 1 of the network 1 of [9].

P(U) and Q(U): these are the popular inverter-level control strategies that are described in Section 4.2, and implemented directly in the load flow solver for simulations.

AMI-based Feed-back: this non-standard control strategy is similar to the $P(U)$ control, except that it is implemented at AMI level (not at inverter level) and therefore suffers from delay in the measurements. We present this strategy to compare it with our feed-forward AMI-based controller below. See Section 4.3.

Feed-forward: this is our proposed strategy, which is described in Section 4.4.

The local control functions of Figure 2 were parameterised as follows: $u_{\text{nom}} = 400/\sqrt{3}$, $u_{\text{min}} = 0.915 u_{\text{nom}}$, $u_{\text{max}} = 1.085 u_{\text{nom}}$, $u_{\text{down}} = 0.9575 u_{\text{nom}}$ and $u_{\text{up}} = 1.0425 u_{\text{nom}}$.

5.4 Technical constraints

The nominal phase-to-neutral voltage magnitude for the simulations is $400/\sqrt{3} \approx 230.94$ V. Voltage at the secondary winding of the transformer is fixed to this nominal value, and the maximal allowed voltage magnitude is +8.5 % above nominal value. Said otherwise, over-voltage constraints begin at 250.57 V.

Regarding the limit on the maximum power flowing at the MV/LV substation through the transformer, we first record the maximum value that occurs during the “test step” (without any control action), and then consider that the capacity of the transformer is equal to 90 % of this value.

5.5 Results: depth of voltage violations

Figure 4 shows the results of our simulations in terms of violations of the upper voltage limit, for each of the control strategies. The value displayed is an average over all 21 feeders and over all simulation time-steps. The vertical bars represent the standard deviations. Note that these bars are not meant to represent the confidence intervals on the mean. Rather, they show the diversity of situations of the various networks.




Figure 4: over-voltages per control strategy. The coloured bar shows the mean value of over-voltages across feeders and time steps. The black line shows standard deviation.

First note that large violations of the upper-voltage limit occur when no control action is taken (“Open-loop 100 %”), and these issues fully disappear when generation is curtailed below 25 % (“Open-loop 0 % and 25 %”). In other words: our data set exhibits a situation where active power control makes sense. As an illustration, Figure 3 shows over-voltages in the first feeder of the first network, without control (“Open-loop 100 %”) and with our feed-forward strategy. This figure shows that voltage constraints are localised on a few network nodes and can be mitigated. Also note that voltage issues naturally tend to decrease when the generation quota is reduced, which appears clearly when comparing the results obtained with various open-loops strategies.

The results obtained with the various control strategies are the following. First, the well-known $P(U)$ control performs very well on the upper-voltage constraint, while the $Q(U)$ control is less efficient; this is because voltage violations are here relatively strong, the reactive capacity of generators is limited, and the R/X ratio of the lines in our dataset is larger than 1 (around 3 for most lines). Second, the AMI-based feed-back control performs poorly due to the control delay introduced by the AMI, which impairs the stability of the controller. Finally, our proposed feed-forward strategy performs well in reducing the over-voltages, although not as well as the $P(U)$ control; this is again an effect of the limitations of the AMI.

5.6 Results: depth of transformer overloads

The second constraint of interest is the upper bound on the (reverse) power flow through the transformer. Figure 5 shows the average over all 21 feeders, and over all time-steps, of power violations (counted negatively because power *leaves* the low-voltage network). Standard deviation is also displayed.




Figure 5: depth of transformer overloads, per control strategy. The coloured bar shows the mean value of overloads across feeders and time steps. The black line shows standard deviation.

Note that the open-loop strategy that curtails power over 75% of the generator capacity is sufficient to mitigate transformer overloads. Unsurprisingly, the *local* controls $P(U)$ and $Q(U)$ fail in solving this *global* constraint. The AMI-based feed-back control fails for the same reason. On the contrary, observe that our proposed AMI-based feed-forward strategy performs well, although not perfectly – still because of the limitations of the AMI.

Table 1: Percentage of AMI measurements that exceed the allowed limits.

Strategy	Over-voltages [%]	Over-powers [%]
Open-loop 100%	2.45	4.53
Open-loop 75%	1.98	0.00
Open-loop 50%	1.02	0.00
Open-loop 25%	0.00	0.00
Open-loop 0%	0.00	0.00
Feed-back	1.12	3.58
Feed-forward	0.10	0.06
$P(U)$	0.00	3.21
$Q(U)$	0.56	4.33

5.7 Results: frequency of violations

Table 1 shows another aspect of voltage and power violations: their frequency of occurrence, rather than their depth. Even on this very sunny summer day with maximal PV generation, violations are quite rare: indeed, congestions only occur in a short time slot around noon, which supports the idea that reinforcement (*permanently* upgrading the network to solve a *very temporary* constraint) might not be the best option.

5.8 Results: energy curtailed

We now turn to assessing the amount of lost energy that is incurred by each of our various control strategies. This indicator is plotted on Figure 6, as a percentage of the total energy that would be produced if no control action was taken (“Open-loop 100%” strategy).




Figure 6: curtailed energy, per strategy. The coloured bar shows the mean value of lost energy across feeders and time steps. The black line shows standard deviation. Standard deviation is zero for open-loop strategies.

As expected, the open-loop strategies curtail more energy when the associated threshold is lowered; for instance, observe that curtailing generation above 50 % of generator capacity leads to wasting slightly more than 20 % of the available solar energy, while curtailing generation above 75 % of the capacity wastes only a few percent of this energy. By definition, the $Q(U)$ control does not curtail any energy at all. The three remaining strategies — AMI-based feed-back, feed-forward and $P(U)$ — are very comparable in this respect: in our experiments, they curtailed nearly the same amount of solar energy.

5.9 Results: trade-off between constraint violations and loss-of-energy

Figures 7 and 8 show the efficiency of each control strategy in terms of constraint mitigation, along with its cost in terms of curtailed energy. On both figures, good controllers are the ones in the lower-left corner, that is to say, controllers that mitigate constraints well at modest cost.




Figure 7: over-voltages versus curtailment (“Ol.” stands for “Open-loop”, “Fb.” for “Feed-back” and “Ff.” for “Feed-forward”).

We observe that our proposed feed-forward controller is the only one that lies in the lower-left corner on both graphs; indeed, the $P(U)$ strategy does slightly better than ours as far as voltage constraints are concerned (Figure 7) but fails in terms of mitigating overloads (Figure 8). On the contrary, the open-loop strategy at 75 % level gives similar performance than ours in terms of overloads, but much worse in terms of over-voltages. Also note that the open-loop strategy must be tuned manually in order to obtain good results in terms of overloads, as shown by Figure 8 where open-loop strategies with levels other than 75 % perform poorly. On the contrary, our AMI-based feed-forward method does not need any parameter tuning, which is a significant advantage of this method.




Figure 8: transformer overloads versus curtailment (“Ol.” stands for “Open-loop”, “Fb.” for “Feed-back” and “Ff.” for “Feed-forward”).

6 CONCLUSION

In this paper, we studied different control methods to eliminate congestions (over-voltage and over-power) caused by distributed PV generation: some local, and some centralised and based on the advanced metering infrastructure (AMI). We discussed how the AMI could be used for control purpose on two levels: firstly to infer knowledge of the underlying network structure, and secondly to populate the values of the parameters that enter the optimisation problem that is solved inside the centralised controller. These arguments led us to the design of a feed-forward controller, that we compared with other controller structures — including some local controllers that do not leverage the AMI. Since our model captures the inherent measurements and communication capabilities of the AMI, it allowed us to assess for the first time the impact of these limitations on various controller structures, and to conclude about the respective value of local and AMI-based controllers. Our simulations showed that an AMI-based feed-forward controller is very well-suited to mitigate voltage and current constraints in a low voltage network, while an AMI-based feed-back controller yielded disappointing results. We also showed that our (centralised) AMI-based feed-forward controller was able to solve issues that are not solved by purely local feed-back controller.

Compared to classical control strategies (open-loop and local feed-back), our AMI-based feed-forward method is competitive: from the performance perspective, we demonstrated its ability to solve both types of constraints (not only voltage constraints, but also overloads) while limiting curtailment levels to a minimum; and from the implementation perspective, this is achieved with limited communication needs on the infrastructure side. Also, the method does not require any parameter tuning, nor prior detailed modelling of the underlying electric network — since it uses inference from AMI data instead. This is key in the context of electric distribution systems, where reliable network data is scarce.

This paper also suggests some directions for future research, such as the issue of fairness between clients of the same network, and the problem of adequately combining the feed-forward and feed-back control logics. On the first topic, we simply observe that a fairness-aware feed-forward controller could be easily derived from (9) by adding constraints or by modifying the objective function. On the second topic, our research suggests that the best combination of feed-back and feed-forward control for PV generation curtailment may be of the following form: an AMI-based feed-forward controller, combined with local feed-back controllers. This particular hybrid method was however not investigated in the present paper, where we focused on pure control strategies. These two topics are left open for future research.

REFERENCES

- [1] Office of Industries. *Global Market for Smart Electricity Meters: Government Policies Driving Strong Growth*. ID-037. Unites-States: US International Trade Commission, June 2014. URL: https://www.usitc.gov/publications/332/id-037smart_meters_final.pdf (visited on 01/17/2018).
- [2] C. L. Masters. “Voltage rise: the big issue when connecting embedded generation to long 11 kV overhead lines”. In: *Power Engineering Journal* 16.1 (Feb. 2002), pp. 5–12. ISSN: 0950-3366. DOI: 10.1049/pe:20020101.
- [3] Bianca Barth, Giorgia Concas, Eduardo Binda Zane, Olivier Franz, Pablo Frias, Roland Hermes, Ricardo Lama, Holger Loew, Carlos Mateo, Manoël Rekinge, Pablo Michele Sonvilla, and Michel Vandenberg. *PV Grid-Final Project Report*. Berlin, Aug. 2014. URL: <http://www.pvgrid.eu/home.html> (visited on 12/21/2017).
- [4] N. Daratha, B. Das, and J. Sharma. “Coordination Between OLTC and SVC for Voltage Regulation in Unbalanced Distribution System Distributed Generation”. In: *IEEE Transactions on Power Systems* 29.1 (Jan. 2014), pp. 289–299. ISSN: 0885-8950. DOI: 10.1109/TPWRS.2013.2280022.
- [5] Kalle Rauma. “Industrial aspects of voltage management and hosting capacity of photovoltaic power generation in low voltage networks”. Theses. Université Grenoble Alpes, Mar. 2016. URL: <https://tel.archives-ouvertes.fr/tel-01345202> (visited on 12/21/2017).
- [6] E. Demirok, P. C. González, K. H. B. Frederiksen, D. Sera, P. Rodriguez, and R. Teodorescu. “Local Reactive Power Control Methods for Overvoltage Prevention of Distributed Solar Inverters in Low-Voltage Grids”. In: *IEEE Journal of Photovoltaics* 1.2 (Oct. 2011), pp. 174–182. ISSN: 2156-3381. DOI: 10.1109/JPHOTOV.2011.2174821.
- [7] K. Turitsyn, P. Sulc, S. Backhaus, and M. Chertkov. “Options for Control of Reactive Power by Distributed Photovoltaic Generators”. In: *Proceedings of the IEEE* 99.6 (June 2011), pp. 1063–1073. ISSN: 0018-9219. DOI: 10.1109/JPROC.2011.2116750.
- [8] Q. Zhou and J. W. Bialek. “Generation curtailment to manage voltage constraints in distribution networks”. In: *Transmission Distribution IET Generation* 1.3 (May 2007), pp. 492–498. ISSN: 1751-8687. DOI: 10.1049/iet-gtd:20060246.
- [9] Alejandro Navarro Espinosa. “Dissemination Document “Low Voltage Networks Models and Low Carbon Technology Profiles””. In: (2015).
- [10] P. Richardson, D. Flynn, and A. Keane. “Optimal charging of electric vehicles in low-voltage distribution systems”. In: *2012 IEEE Power and Energy Society General Meeting*. 2012 IEEE Power and Energy Society General Meeting. July 2012, pp. 1–1. DOI: 10.1109/PESGM.2012.6344912.
- [11] J. Seuss, M. J. Reno, M. Lave, R. J. Broderick, and S. Grijalva. “Advanced inverter controls to dispatch distributed PV systems”. In: *2016 IEEE 43rd Photovoltaic Specialists Conference (PVSC)*. 2016 IEEE 43rd Photovoltaic Specialists Conference (PVSC). June 2016, pp. 1387–1392. DOI: 10.1109/PVSC.2016.7749842.
- [12] Fabio Bignucolo, Roberto Caldon, and Valter Prandoni. “Radial MV networks voltage regulation with distribution management system coordinated controller”. In: *Electric Power Systems Research* 78.4 (Apr. 1, 2008), pp. 634–645. ISSN: 0378-7796. DOI: 10.1016/j.epsr.2007.05.007. URL: <http://www.sciencedirect.com/science/article/pii/S0378779607001125> (visited on 12/21/2017).
- [13] Légifrance. *Arrêté du 4 janvier 2012 pris en application de l'article 4 du décret n° 2010-1022 du 31 août 2010 relatif aux dispositifs de comptage sur les réseaux publics d'électricité*. Aug. 31, 2010. URL: <https://www.legifrance.gouv.fr/affichTexte.do?cidTexte=JORFTEXT000025126353&categorieLien=id> (visited on 01/24/2018).
- [14] William H. Kersting. *Distribution System Modeling and Analysis, Third Edition*. CRC Press, Jan. 24, 2012. 459 pp. ISBN: 978-1-4398-5622-2.
- [15] Ian Richardson, Murray Thomson, David Infield, and Conor Clifford. “Domestic electricity use: A high-resolution energy demand model”. In: *Energy and Buildings* 42.10 (Oct. 1, 2010), pp. 1878–1887. ISSN: 0378-7788. DOI: 10.1016/j.enbuild.2010.05.023. URL: <http://www.sciencedirect.com/science/article/pii/S0378778810001854> (visited on 01/19/2018).