



## Fully distributed peer-to-peer optimal voltage control with minimal model requirements

Ortmann, Lukas; Prostejovsky, Alexander; Heussen, Kai; Bolognani, Saverio

*Published in:*  
Electric Power Systems Research

*Link to article, DOI:*  
[10.1016/j.epsr.2020.106717](https://doi.org/10.1016/j.epsr.2020.106717)

*Publication date:*  
2020

*Document Version*  
Peer reviewed version

[Link back to DTU Orbit](#)

*Citation (APA):*  
Ortmann, L., Prostejovsky, A., Heussen, K., & Bolognani, S. (2020). Fully distributed peer-to-peer optimal voltage control with minimal model requirements. *Electric Power Systems Research*, 189, Article 106717. <https://doi.org/10.1016/j.epsr.2020.106717>

---

### General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

# Fully Distributed Peer-to-Peer Optimal Voltage Control with Minimal Model Requirements

Lukas Ortmann  
Saverio Bolognani  
Automatic Control Laboratory  
ETH Zurich  
Zurich, Switzerland

Alexander Prostejovsky  
Kai Heussen  
Department of Electrical Engineering  
Technical University of Denmark (DTU)  
Roskilde, Denmark

**Abstract**—This paper addresses the problem of voltage regulation in a power distribution grid using the reactive power injections of grid-connected power inverters. We first discuss how purely local voltage control schemes cannot regulate the voltages within a desired range under all circumstances and may even yield detrimental control decisions. Communication and, through that, coordination are therefore needed. On the other hand, short-range peer-to-peer communication and knowledge of electric distances between neighbouring controllers are sufficient for this task. We implement such a peer-to-peer controller and test it on a 400 V distribution feeder with asynchronous communication channels, confirming its viability on real-life systems. Finally, we analyze the scalability of this approach with respect to the number of agents on the feeder that participate in the voltage regulation task.

**Index Terms**—distributed control, distributed optimization, power distribution grids, reactive power, Volt/VAR control.

## I. INTRODUCTION

Future power distribution grids are expected to host a significant portion of the total generation capacity, for the most part from renewable energy sources like solar and micro-wind installations. Meanwhile, the deployment of a distributed electric mobility infrastructure will substantially increase the loading of this infrastructure. This transition will inevitably affect the operating regime of distribution feeders, and will increase the risk of both overvoltage and undervoltage contingencies. On the other hand, microgenerators and electric vehicle charging stations will offer unprecedented voltage control flexibility via their power inverters, offering a finely distributed network of reactive power compensators.

For the control of these reactive power compensators, a multitude of decentralized Volt/VAR feedback control strategies have been proposed (e.g., Volt/VAR droop control; cf. [1]

for a literature review) and ultimately incorporated in many grid codes and standards [2]–[4]. These strategies rely on the control architecture schematically represented in Figure 1, where each power inverter independently regulates its reactive power injection based on the voltage measurement performed at its point of connection, typically via a static update map

$$q_h(t+1) = f_h(v_h(t)).$$

The update map  $f_h$  is usually the outcome of heuristic design procedures. In most cases the design is completely model-free (no grid information is used), although computational design approaches have also been proposed [5].

Fully decentralized feedback control solutions present multiple advantages, such as:

- high robustness, given by the absence of a single point of failure;
- economical deployment and retrofitting (plug-and-play);
- minimal actuation time delays, due to the absence of any communication;
- modularity and interoperability, as individual inverters do not coordinate their action;
- scalability and computational simplicity.

However, purely decentralized control strategies fail to ensure feasible voltages, even if such a feasible solution exists, as recently proven in [1]. Conversely, centralized feedback Volt/VAR solutions are guaranteed to drive the system to a feasible voltage profile, using the same measurements collected in the decentralized setting (i.e., only voltage magnitude




Fig. 1: Schematic representation of the control architecture employed by fully decentralized Volt/VAR feedback strategies, e.g. [2]–[4].

This research has been performed using the ERIGrid Research Infrastructure and is part of a project that has received funding from the European Union's Horizon 2020 Research and Innovation Program under the Grant Agreement No. 654113. The support of the European Research Infrastructure ERIGrid and its partner Technical University of Denmark is very much appreciated. This paper reflects only the authors' view and the EU Commission is not responsible for any use that may be made of the information it contains.

The research leading to this work was supported in part by the Swiss Federal Office of Energy grant #SI/501708 UNICORN.

Corresponding author: Lukas Ortmann, email: ortmannl@ethz.ch.

measurements of the inverters) but processing them in a centralized manner. We refer to [6], [7] for a recent review of *feedback optimization* methods that can be employed to design these centralized feedback Volt/VAr strategies, and to [8] for an experimental validation that demonstrates a remarkable robustness against model uncertainty. The disadvantages of centralized feedback optimization are that a communication channel between a central computational unit and all the power inverters is required, and a global model of the grid needs to be known at this central location.

This paper is motivated by a fundamental question: is it possible to achieve optimal Volt/VAr regulation without collecting all measurement and all model information in a centralized location? A limited number of recent works contributed towards an answer to this question by proposing feedback control strategies that are extremely parsimonious in terms of information that inverters need to communicate:

- A distributed solution for the voltage regulation and loss minimization problem is proposed in [9], allowing asynchronous communication between agents (but relying on both angle and magnitude measurements).
- In [10], power inverters are controlled by individual automata that communicate a “distress signal” only when their regulation problem becomes infeasible; however, this strategy is not guaranteed to converge to the optimal regulation.
- In a similar spirit, [11] proposes a distributed strategy in which inverters communicate only when triggered by local voltage violation rules; an all-to-all communication channel is however assumed.
- A primal-dual method that requires only communication between neighboring inverters is proposed in [12];
- The authors of [13] demonstrate how coordination between inverters can be achieved by only transmitting a few bits of information;
- a distributed dual ascent method is employed in [14], allowing for delayed communication between inverters;
- finally, [1] proposes a distributed synchronous dual ascent method with a nested quadratic program.

To the best of the authors’ knowledge, none of these distributed solutions has been implemented and tested on a real grid with physically distributed computations.

In this work, we provide a proof-of-concept demonstration of how Volt/VAr regulation can be achieved via a distributed feedback control law, namely under the specifications that:

- each inverter can only establish asynchronous peer-to-peer communication with its neighboring inverters;
- each inverter only maintains model information regarding its grid neighborhood;
- no central coordination unit is present.

The reported experiment also validates other important features of this distributed solution such as its robustness against noisy measurements, its real-time computational feasibility, and the viability of algorithm distribution in a peer-to-peer setting with no master algorithm synchronization.

Finally, we investigated scalability of the proposed approach via a series of numerical experiments.

## II. DISTRIBUTED VOLTAGE CONTROL

In this section we report the procedure proposed in [1] to design a distributed controller for the Volt/VAr regulation problem. Although a synchronous communication channel was assumed in [1], it provides the key idea on how to achieve optimal coordination via only short-range exchange of information.

### A. Feedback Optimization Controller

The controller is derived from the optimization problem

$$\begin{aligned} \min \quad & \frac{1}{2}q^T M q \\ \text{subject to} \quad & v_{\min} \leq v_h(q, w) \leq v_{\max} \quad \forall h \\ & q_{\min} \leq q_h \leq q_{\max} \quad \forall h. \end{aligned} \quad (1)$$

where the matrix  $M$  is a square, symmetric and positive definite design parameter and  $v$  and  $q$  are the vectors we obtain by stacking the voltages  $v_h$  and reactive power set-points  $q_h$  of the different inverters, respectively. The function  $v_h(q, w)$  is the steady-state map of the nonlinear power flow equations that defines voltages  $v_h$  as a function of both reactive powers  $q$  and external influences  $w$  (e.g., active and reactive demands, active generation).

Active power injections are not a decision variable in (1) for the following reason. Controlling the active power of devices comes with an economic cost, whereas the usage of reactive power is free (neglecting the active power losses generated by the reactive power flows). Therefore, it is typically preferred to use the reactive power capabilities in the network to their full extent before controlling active power injections.

To solve (1) we introduce the dual multipliers  $\lambda_{h,\min}$  and  $\lambda_{h,\max}$  for the voltage constraints of every inverter  $h$ . Stacking them gives us the vector  $\lambda = \begin{bmatrix} \lambda_{\min} \\ \lambda_{\max} \end{bmatrix}$  with which we form the Lagrangian  $\mathbf{L}(q, \lambda)$  by dualizing the voltage constraints:

$$\begin{aligned} \mathbf{L}(q, \lambda) = \frac{1}{2}q^T M q + \sum_h \lambda_{h,\min}(v_{\min} - v_h(q, w)) \\ + \sum_h \lambda_{h,\max}(v_h(q, w) - v_{\max}). \end{aligned} \quad (2)$$

We thus define the equivalent dual optimization problem

$$\begin{aligned} \max_{\lambda \geq 0} \min_q \quad & \mathbf{L}(q, \lambda) \\ \text{subject to} \quad & q_{h,\min} \leq q_h \leq q_{h,\max} \quad \forall h. \end{aligned} \quad (3)$$

The optimization problems (1) and (3) have the same solution (Strong Duality Theorem, [15, Proposition 5.3.2]). We adopt an iterative dual ascent update on  $\lambda$  to compute the solution of (3), obtaining

$$\begin{aligned} \lambda_{\min}(t+1) &= [\lambda_{\min}(t) + \alpha \nabla_{\lambda_{\min}} \mathbf{L}(q(t), \lambda(t))]_{\geq 0} \\ &= [\lambda_{\min}(t) + \alpha(v_{\min} - v(q(t), w))]_{\geq 0} \\ \lambda_{\max}(t+1) &= [\lambda_{\max}(t) + \alpha \nabla_{\lambda_{\max}} \mathbf{L}(q(t), \lambda(t))]_{\geq 0} \\ &= [\lambda_{\max}(t) + \alpha(v(q(t), w) - v_{\max})]_{\geq 0}. \end{aligned} \quad (4)$$

As we can see every inverter integrates its own voltage violation with a gain of  $\alpha$ . This corresponds to the integral part of a PI-controller and can be done locally, by using

feedback from the physical system through voltage magnitude measurements  $v(t)$  of the inverters, rather than via a numerical evaluation of  $v(q(t), w)$ . To find the optimal reactive power set-points we use the newly calculated  $\lambda(t + 1)$  and solve

$$\begin{aligned} q(t + 1) = & \arg \min_q \mathbf{L}(q, \lambda(t + 1)) \\ \text{subject to } & q_{h,\min} \leq q_h \leq q_{h,\max} \quad \forall h. \end{aligned} \quad (5)$$

Towards this goal, we introduce the approximation

$$\frac{\partial v(q, w)}{\partial q} \approx X \quad (6)$$

where  $X$  is the reduced bus reactance matrix that can be derived from the grid topology and the cable data. The sensitivity described by  $X$  is similar to power transfer distribution factors for active power generation on the transmission level. Under no-load conditions and the assumption of negligible cable resistances this approximation is accurate, because the nonlinearity of the power flow equations is mild near this operating point [16]. In our application the system can be heavily loaded and the cable resistances are high. It was shown in [8] that feedback optimization is sufficiently robust against this model mismatch.

This approximation makes  $v(q, w)$  linearly dependent on  $q$ , and we can approximate (5) with a convex quadratic optimization problem (QP). This QP involves the decision variables of all DERs and can be solved by collecting all the necessary information (the multipliers  $\lambda(t + 1)$  and the parameters  $X$ ) in a central control unit [8]. In the following we use the idea proposed in [1] to show how (5) can also be solved in a distributed manner, without centralized computation or centralized model knowledge.

### B. Distributing the Controller

To solve the subproblem (5) in a distributed manner we perform  $K$  iterative steps, which will have to be executed between the times  $t$  and  $t + 1$ . To denote these iterative steps we introduce a new iteration counter  $\tau$ . We also introduce the dual multipliers  $\mu_{h,\min}$  and  $\mu_{h,\max}$  for the reactive power constraints of every inverter  $h$ , which we stack in the vector  $\mu = [\mu_{\min}^{\min}, \mu_{\max}^{\max}]$ . By dualizing the reactive power constraints, we define the Lagrangian

$$\begin{aligned} \mathbf{L}_\lambda(q, \mu) = & \mathbf{L}(q, \lambda) + \sum_h \mu_{h,\min}(q_{\min} - q_h) \\ & + \sum_h \mu_{h,\max}(q_h - q_{\max}) \end{aligned} \quad (7)$$

and the following optimization problem:

$$\max_{\mu \geq 0} \min_q \mathbf{L}_\lambda(q, \mu). \quad (8)$$

The optimization problems (5) and (8) have the same solution (Strong Duality Theorem, [15, Proposition 5.3.2]). Similarly as




Fig. 2: Schematic representation of neighbor-to-neighbor communication, where we adopt the definition of neighbors from [9]: two inverters are neighbors if the electrical path connecting them does not pass through any other bus where a controlled inverter is connected.

before, we solve this optimization problem via gradient ascent iterations on  $\mu$  with step size  $\gamma$ :

$$\begin{aligned} \mu_{\min}(\tau + 1) &= [\mu_{\min}(\tau) + \gamma \nabla_{\mu_{\min}} \mathbf{L}_\lambda(\hat{q}(\tau), \mu(\tau))]_{\geq 0} \\ &= [\mu_{\min}(\tau) + \gamma(q_{\min} - \hat{q}(\tau))]_{\geq 0} \\ \mu_{\max}(\tau + 1) &= [\mu_{\max}(\tau) + \gamma \nabla_{\mu_{\max}} \mathbf{L}_\lambda(\hat{q}(\tau), \mu(\tau))]_{\geq 0} \\ &= [\mu_{\max}(\tau) + \gamma(\hat{q}(\tau) - q_{\max})]_{\geq 0} \end{aligned} \quad (9)$$

where

$$\hat{q}(\tau) = \arg \min_q \mathbf{L}_\lambda(q, \mu(\tau)).$$

Observe, that the update of  $\mu_{\min}$  and  $\mu_{\max}$  can be done locally by every inverter by integrating the constraint violation of the virtual quantity  $\hat{q}(\tau)$ . In order to compute the unconstrained minimizer  $\hat{q}(\tau)$ , we take the derivative  $\nabla_q \mathbf{L}_\lambda(q, \mu)$  and obtain

$$\nabla_q \mathbf{L}_\lambda(q, \mu) = Mq + \frac{\partial v}{\partial q}(\lambda_{\max} - \lambda_{\min}) + \mu_{\max} - \mu_{\min}. \quad (10)$$

As stated before, we approximate the derivative  $\partial v / \partial q$  with  $X$  and set (10) to 0. We then solve for  $q$  and obtain

$$\begin{aligned} \hat{q}(\tau) &= -M^{-1}X(\lambda_{\max} - \lambda_{\min}) \\ &+ M^{-1}[\mu_{\max}(\tau) - \mu_{\min}(\tau)]. \end{aligned} \quad (11)$$

Equation (11) reveals that all the communication requirements of the proposed iterative algorithm are encoded in the sparsity of the matrices  $M^{-1}$  and  $M^{-1}X$ . In fact, off-diagonal non-zero elements of these two matrices determine components of  $\lambda$  and  $\mu$  that need to be communicated between inverters in order to compute  $\hat{q}(\tau)$ .

In order to maximize the sparsity of both these matrices, we exploit the structure inherited from the physical system. We inherit the formal definition of neighboring inverters from [9], see Figure 2. Neighbors according to this definition can be conveniently discovered via correlation analysis of the voltage measurements, even without central supervision (see [17], [18] and references therein).

Due to the sparsity of the power flow equations,  $G = X^{-1}$  is a sparse matrix: namely,  $G_{ij}$  of  $G$  is non-zero only if the buses  $i$  and  $j$  are neighbors, and  $G_{ij}$  depends only on the electrical impedance of the path between  $i$  and  $j$ .

Based on this observation, we choose  $M = X$  which yields

$$\hat{q}(\tau) = \lambda_{\min} - \lambda_{\max} + G(\mu_{\min}(\tau) - \mu_{\max}(\tau)). \quad (12)$$

Therefore inverter  $i$  only needs to gather  $\mu_{j,\min}$  and  $\mu_{j,\max}$  from their neighbors  $j$  to calculate  $\hat{q}_i(\tau)$ . Note that  $M = X$

is possibly not the only choice that allows to distribute the algorithm, if one accepts to use a descent direction in the gradient steps which is not the steepest one [6, Section III.F].

With a proper choice of the gain  $\gamma$  (for which we refer to Section V) the alternate execution of (9) and (12) is guaranteed to converge to the solution to (8). We assume that the number of iterations  $K$  is chosen sufficiently large so that, after  $K$  iterations,  $\hat{q}$  is accepted as the solution to (5) and determines the next set-point  $q(t+1)$ . The effect of this approximation is also studied in Section V.

The resulting control algorithm consists in a main loop, reported hereafter as Algorithm 1, and a nested iterative procedure, Algorithm 2. Communication between agents only happens as part of Algorithm 2, when the dual multipliers  $\mu$  of the reactive power constraints need to be communicated with neighbors (steps 7–8). All other steps are basic numerical operations that each inverter performs locally. The resulting control architecture is represented in Figure 3.

Note that the implementation of our controller inherits the theoretical guarantees provided in [1], including Proposition 6 that guarantees asymptotic optimality (under the linearity condition (6) and assuming that (5) is solved exactly).

---

**Algorithm 1** Feedback optimization controller

---

- 1: Initialize:  $\lambda_{h,\min}$  and  $\lambda_{h,\max}$  with 0
  - 2: **loop**
  - 3:   Locally measure the voltage magnitude  $v_h$
  - 4:   Locally update  $\lambda_{h,\min}$  and  $\lambda_{h,\max}$  via (4)
  - 5:   Jointly compute  $q_h$  via Algorithm 2
  - 6:   Locally apply the new set-point  $q_h$
  - 7:   Wait until next system interrupt
  - 8:   ▷ System interrupts generated every  $T$  seconds
- 

---

**Algorithm 2** Distributed QP solver

---

- 1: **if** Algorithm 2 was never executed previously **then**
  - 2:   Initialize:  $\mu_{h,\min}$ ,  $\mu_{h,\max}$  and  $\hat{q}_h$  with 0
  - 3: **else**
  - 4:   Keep previous values to warm start
  - 5: counter = 0                                   ▷ Iteration counter
  - 6: **repeat**
  - 7:   Locally update  $\mu_{h,\min}$  and  $\mu_{h,\max}$  via (9)
  - 8:   Send  $\mu_{h,\min}$  and  $\mu_{h,\max}$  to neighbors
  - 9:   Receive  $\mu_{i,\min}$  and  $\mu_{i,\max}$  from all neighbors
  - 10:   Locally, compute  $\hat{q}_h$  via (12)
  - 11:   counter = counter + 1
  - 12: **until** counter ==  $K$
  - 13: Return the solution  $q_h = \hat{q}_h$
- 

### III. EXPERIMENTAL SETUP

The experiment has been implemented in the SYSLAB facility located on the Risø campus of the Technical University of Denmark. The setup consists of a 400 V three-phase electric grid connecting a variety of DERs (solar panels, wind turbines, a flow battery, a diesel generator, controllable loads, among




Fig. 3: Control architecture. Measurement and actuation is performed locally by each controller. Only the dual multipliers  $\mu_{\min}$  and  $\mu_{\max}$  need to be communicated to neighboring peers.

others). Each device has an associated computer node running a distributed monitoring and control platform.

#### A. Algorithm Implementation and Deployment

An existing distributed optimization framework developed at DTU [19] was adopted to implement the proposed distributed optimization controllers over an asynchronous communication channel. Each computer node implements Algorithm 1 in major fixed time intervals of  $T = 10$  seconds, based on their individual clock. This is therefore the rate at which measurements are collected (line 3) and the system is actuated (line 6). The choice of such a long interval is due to hardware constraints given by the laboratory setup. A more frequent actuation is often possible. However, the actuation interval should be long enough for the system to settle and reach its steady state. The frequency at which the system can be actuated will always be significantly lower than the rate at which inverters can communicate (see Section V for a discussion on the implications on the algorithm scalability).

Algorithm 2 is executed in  $K$  iterations. Lines 7 and 8 of this algorithm require communication between neighbours, where the communication time is variable, dependent on uncontrollable influences. Coherency of the algorithm, and thereby a synchronous advancement of the algorithm steps, is achieved by letting individual nodes remain idle until data has been received from all neighbours (line 8 of Algorithm 2). This way, the synchronous Algorithm 2 is transparently implemented on an asynchronous communication channel, which has better scaling properties than a synchronous one in such a setup [19]. ZeroMQ [20] is used as the underlying messaging library with TCP transport, facilitating reliable data delivery. The code comprising the distributed framework and algorithm is deployed to each of the active SYSLAB node computers and operates as a local process.

#### B. Test Case and Experiment Design

The topology and operational set-points are designed to produce a voltage drop at the beginning and an overvoltage at the end of the feeder. Without proper reactive power control, the feeder’s ability to host renewable energy infeed is limited and generation would need to be curtailed. The setup consists of the flow battery, two photovoltaic arrays (PV), an adjustable resistive load, and a utility grid connection (PCC). This test system is illustrated in Figure 4, and Figure 5 presents the corresponding implementation on the SYSLAB topology view.




Fig. 4: Sketch of the voltage profile, the distribution feeder and the peer-to-peer communication. The colors of the voltage profile and the diagram match the colors in the topology in Figure 5.

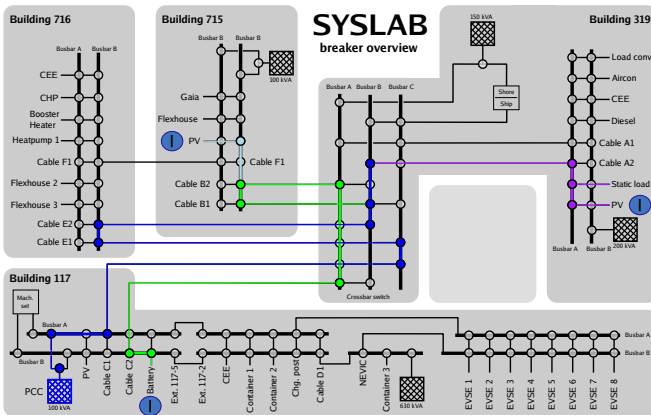


Fig. 5: SYSLAB infrastructure with the used topology. The colors match the colors in the diagram and in the voltage profile in Figure 4. Inverters participating in the control algorithm are marked with an I.

The active power injection  $p_3$  of the battery is interpreted as a renewable power source, which is not to be curtailed; its active power infeed is set to  $p_3 = 10$  kW. The static load is set to an active power consumption of 15 kW ( $p_1 = -15$  kW) which is larger than the local production, therefore causing a positive active power flow from the substation. PVs are fluctuating power sources. Therefore, to facilitate repeatability of the experiments and to allow for a comparison between different controllers, the PVs are curtailed to not inject active power ( $p_2 = 0$  kW). The different nodes are connected via cables with non-negligible resistance, see Table I. Due to a weak link (resistive) cable connecting the battery to the grid, the battery encounters an overvoltage when the reactive power injection is zero. Both PVs and the battery can measure their voltage magnitudes, and their reactive power injections can be controlled. The PV inverters have a reactive power range of  $\pm 6$  kVAr and the battery can be actuated with  $\pm 8$  kVAr. The PVs and the battery can communicate with their neighbors, while the load is uncontrolled and unmeasured. The voltage limits are defined to be 0.95 p.u. and 1.05 p.u.

TABLE I: Parameters of the cables between busbars/devices.

Cable	Length [m]	Cross section [mm <sup>2</sup> ]	R [Ω]	X [Ω]
C1	700	240	0.085	0.054
E1	450	240	0.055	0.035
E2	450	240	0.055	0.035
A2	25	95	0.0078	0.002
PV1	83	16	0.095	0.007
Static Load	11	95	0.002	0.001
B1	350	95	0.11	0.027
PV2	8	6	0.025	0.0008
B2	350	95	0.11	0.027
C2	700	240	0.085	0.054
Battery	100	2.5	0.774	0.012

## IV. EXPERIMENTAL RESULTS

In this section, we first demonstrate the suboptimal performance of a decentralized (purely local) controller on the proposed system, by implementing the droop control recommended by the IEEE standard [2]. We then execute the proposed distributed controller, evaluate its control performance, examine the nested execution of Algorithm 2, hint at a windup phenomenon in case of problem infeasibility, and analyze the trade-off between control performance and communication complexity.

### A. Local Control

The droop controller that we implement is the one proposed in [2] and similar to the ones suggested in the recent grid codes [3], [4]. Every inverter implements the following piecewise linear control law.

$$q_h = \begin{cases} q_{\max} & v_h < v_1 \\ q_{\max} \frac{v_2 - v_h}{v_2 - v_1} & v_1 \leq v_h \leq v_2 \\ 0 & v_2 \leq v_h \leq v_3 \\ q_{\min} \frac{v_h - v_3}{v_4 - v_3} & v_3 \leq v_h \leq v_4 \\ q_{\min} & v_4 < v_h \end{cases}$$

where  $v_h$  is the measured voltage magnitude,  $q_h$  is the calculated reactive power injection,  $q_{\min}$  and  $q_{\max}$  are the minimum and maximum reactive power injection. We tune the droop curve to  $v_1 = 0.95$  p.u.,  $v_2 = 0.99$  p.u.,  $v_3 = 1.01$  p.u. and  $v_4 = 1.05$  p.u.

The resulting performance of the controller is reported in Figure 6. When the control is activated at minute 3, only the controller at the battery detects a voltage violation and immediately lowers its reactive power injection to the minimum. However, this is not sufficient to regulate the voltage to the desired voltage range. The PV systems do not detect an overvoltage and therefore do not draw reactive power. Without introducing coordination between the inverters, the persistent overvoltage at the battery cannot be prevented. Therefore, all local control strategies fail in this setup, as established from a theoretical perspective in [1].

Figure 6 also shows that PV1 injects reactive power around minute 4 of the experiment. This worsens the overvoltage at






Fig. 6: Performance of local droop control (IEEE 1547 standard).




Fig. 7: Performance of the distributed voltage controller,  $\alpha = 100$ .

the battery, which shows that local control decisions can in some cases be even detrimental.

### B. Controller Evaluation

Figure 7 shows the performance of the distributed voltage controller with a gain of  $\alpha = 100$ ,  $K = 100$  communication steps to distributively solve the QP, an ascent step length of  $\gamma = 0.005$  and with matrices  $M = X$  and  $G = X^{-1}$ :

$$X = \begin{bmatrix} 0.10 & 0.09 & 0.09 \\ 0.09 & 0.11 & 0.11 \\ 0.09 & 0.11 & 0.16 \end{bmatrix}, G = \begin{bmatrix} 48.3 & -40.7 & 0 \\ -40.7 & 61.8 & -18.7 \\ 0 & -18.7 & 19.1 \end{bmatrix}.$$

Cable data have been used to compute the matrix  $G$ , although the necessary parameters could also be estimated (see [21] for an experimental demonstration on the same network). Notice that, as expected, the matrix  $G$  has the sparsity pattern induced by the topology of the distribution grid (zero elements in the positions corresponding to non-neighbors). The system is initialized with zero reactive power flow.<sup>1</sup> The controller is activated after 3 minutes and drives all voltages to the desired range. After 11 minutes the active power of the battery, which produces the overvoltage, is brought to 0 kW. The algorithm promptly responds by bringing the reactive power injections of all the power inverters to 0 kVAR.

For a more in-depth analysis of the control behavior we provide the data in Figure 8 for a controller with  $\alpha = 50$  and  $K = 50$ . We report both the electrical quantities  $v$  and  $q$  and the controllers' internal variables  $\lambda_{\max}$  and  $\mu_{\min}$  ( $\lambda_{\min}$  and  $\mu_{\max}$  remain zero in this experiment). Once the controller is activated at 3 minutes, the voltage violation leads to a growing  $\lambda_{3,\max}$  at the battery. As this integral variable grows, the battery starts drawing reactive power. Once the reactive power  $q_3$  of the battery reaches the battery's reactive power limit, the corresponding multiplier  $\mu_{3,\min}$  starts growing. At

each iteration of Algorithm 2, this value is communicated to PV2. Ultimately, PV2 starts drawing reactive power as well (thus participating to the voltage regulation task). Once the reactive power limit of PV2 is reached, its  $\mu_{2,\min}$  value becomes positive and PV1 starts to draw reactive power. As long as there remains an overvoltage at the battery, the battery keeps integrating its  $\lambda_{3,\max}$ , which leads to a larger reactive power demand by the inverter that is closer to battery and is not yet saturated. Finally, the voltage converges to the voltage constraint. Once that point is reached the system has settled (not fully represented in Figure 8). Three remarks are due.

- There is no central clock signal and the different inverters time their iterations of Algorithm 1 independently. Measurements are therefore not perfectly synchronous. We do not observe any detrimental effect in the experiment.
- Each controller gathers raw voltage measurements. No filtering or state estimation is performed (which, in general, would require a system model and further exchange of information). The control performance seems to be unaffected by measurement noise and quantization.
- With a smaller actuation interval (smaller  $T$ ) and therefore more frequent control actuations the settling time of the algorithm can be reduced.

### C. Convergence of Algorithm 2

In Figure 9, we can see how the internal variables  $\hat{q}$  of the three inverters converge during the execution of Algorithm 2. The algorithm is started as soon as the multipliers  $\lambda_{\max}$  are updated with the measured voltage violation. Agents update their internal variable  $\hat{q}_h$  (orange dots in Figure 9) and their multipliers  $\mu_{h,\min}, \mu_{h,\max}$  (not represented) while communicating with their neighbors at each iteration. After  $K$  iterations, the internal value  $\hat{q}$  is used to actuate the system by updating the reactive power set-points for the inverters (blue line). A few remarks are due:

<sup>1</sup>Due to an inaccuracy of the sensor used by the internal reactive power controller of the battery, we can observe a small tracking error. The reported measurements in the figures are from accurate sensors.




Fig. 8: Behavior of the distributed controller with control gain  $\alpha = 50$ . Upper panels: electrical quantities  $v, q$ . Lower panel: dual multipliers  $\lambda_{\max}$  and  $\mu_{\min}$  (evaluated at the last step of Algorithm 2).

- due to the warm start of the algorithm and the relatively small changes in  $\lambda_{\max}$ , the initialization of  $\hat{q}$  is already close to the final (optimal) value;
- $K = 40$  iterations suffice for the convergence of Algorithm 2 in this experiment (see Section V for further discussion on the effect of early termination of Algorithm 2);
- the time needed to complete Algorithm 2 is significantly shorter than the sampling rate of Algorithm 1 (10 s).

#### D. Controller Windup

Figure 10 illustrates the behavior of the proposed scheme when the Volt/VAR regulation problem is temporarily unfeasible. A persistent overvoltage at the battery leads to a constantly growing  $\lambda_{3,\max}$ . All inverters are drawing their maximum reactive power, which confirms that the voltage cannot be regulated: there does not exist a feasible reactive power input such that all voltages are within the voltage limits.

Once we remove the cause of the overvoltage (at approx. 4 minutes) and the voltage drops, the inverters do not adjust their




Fig. 9: Convergence of one execution of Algorithm 2. Reactive power set-points (solid) and updates of the internal variable  $\hat{q}$  (dotted).




Fig. 10: Controller windup due to a persistent overvoltage.

reactive power injection, but remain saturated at their limit value for several minutes. This phenomenon corresponds to the windup behavior that is often observed in integral controllers. Here, the integrator is  $\lambda_{3,\max}$  of the battery.

One solution to this windup problem is to stop the integration of the voltage violation once all inverters have saturated. While this is an easy modification for a centralized controller (see [8]), a more sophisticated anti-windup scheme is needed in a distributed setup, where no single agent is aware of the infeasibility of the optimization problem. The design of an effective distributed anti-windup scheme is an interesting and open problem *per se*.

#### E. Control Performance vs Communication Complexity

The ability of performing optimal voltage control without global communication comes at a price. As detailed in Section II-B, in order to obtain an iterative update that



TABLE II: Comparison between the steady-state of the distributed algorithm and the maximal-fairness set-points that minimize (13).

	$\arg \min q^T X q$	$\arg \min J_{\text{fair}}$	difference
$J_{\text{fair}}$	2.12	1.98	6.9%
PV1 $q_1$	-2.06 kVAr	-3.76 kVAr	0.28 [p.u.]
PV2 $q_2$	-6 kVAr	-4.6 kVAr	0.23 [p.u.]
Battery $q_3$	-8 kVAr	-8 kVAr	0 [p.u.]

only requires neighbor-to-neighbor communication we had to constrain the choice of the quadratic cost parameter  $M$  in (1).

We showed that  $M = X$  is a valid choice,  $X$  being the grid susceptance matrix. As discussed in [9], the minimization of  $q^T X q$  is connected to the minimization of power losses caused by reactive power flows (under the assumption of homogeneous X/R ratio). Moreover, as discussed in [1], the cost  $q^T X q$  can then be rewritten as  $(Xq)^T G X q$ , where  $Xq$  is the first order approximation of the voltage drop caused by reactive power injection. Therefore, because  $G$  has the structure of a Laplacian,  $q^T X q$  promotes equal voltage drops in the network.

In general, however, a network operator may be interested in minimizing a different cost function, e.g.

$$J_{\text{fair}} := \sum_i (q_i / q_i^{\max})^2 \quad (13)$$

which promotes proportional fairness in the use of the reactive power capacity of each inverter. The difference in the reactive power set-points and in the resulting cost is reported in Table II. Given the inexpensive nature of reactive power, these differences are in most cases acceptable.

## V. SCALABILITY

In this section we investigate how the performance of the proposed feedback scheme scales with the number of nodes. In order to perform this analysis, we consider a fictitious scenario which is identical to the one described in Section III-B, but where  $N'$  extra “dummy” nodes have been added on the line connecting PV2 to the Battery (see Figure 11). These nodes are equally spaced and have zero reactive power capability. They can communicate one to the next one, so that the entire communication graph becomes a line of  $N = N' + 3$  nodes. Their presence therefore affect the execution of the algorithm without affecting the optimal solution of the problem.

We compare two implementations of our method: in the first case (that we denote as  $K \gg 1$ ), we allow an arbitrarily large number  $K$  of communication steps between each actuation step (namely, we allow communication until convergence up to a tolerance of 10 VAr); in the second case, we only allow one communication step for each actuation step ( $K = 1$ ).

This second case closely resembles what was proposed [9]. More generally, we consider it as a prototype for the other distributed methods available in the literature, where communication and actuation are always interleaved one-to-one. These include for example the primal-dual methods proposed in [13] or in [12]. The other methods reviewed in the Introduction also share the same interleaving between communication and actuation. We will see in this section how




Fig. 11: Diagram of the electrical topology and of the communication graph used in the numerical analysis of the algorithm scalability.

this appears to be a design choice that limits performance as the network grows in size. In the comparison, it is important to keep in mind that iterations that only require computation and communication can be executed much faster than iterations that require actuation of the system and measurement. We therefore counted and reported them separately, as *communication steps* and *actuation steps*.

We executed the algorithm with  $K \gg 1$  and  $K = 1$  for networks of different sizes, and these are the main findings.

### Ease of tuning:

- If  $K$  is large, then the tuning of the inner optimization gain  $\gamma$  becomes very simple; Figure 12 shows how a large  $K$  gives a plateau of valid choices for  $\gamma$ .
- An optimal value of  $\gamma$  as a function of the grid parameters has been suggested in [9, Corollary 2] when  $K$  is unbounded, and seems to be an excellent choice also for  $K$  finite but sufficiently large. In the specific case of Figure 12, the recommended  $\gamma$  is  $1/(2\sigma(G)) = 0.004$ .
- For large  $K$ , tuning  $\alpha$  becomes significantly simpler. All the executions of  $K \gg 1$  in Table III use the same parameter  $\alpha = 100$  and attain similar time to convergence, while  $\alpha$  needed to be substantially re-tuned when  $K = 1$ .

### Controller performance:

- By allowing many communication steps, the number of actuation steps required for convergence is significantly reduced (see Table III).
- The number of communication steps required for convergence of the nested subproblem increases with the size of the network (although not exponentially). Remember that these steps only require communication and computation, and we showed that they may also be performed asynchronously (in contrast to the actuation steps, which need to be synchronous). By using an asynchronous implementation of the inner loop data exchange, the time needed for a single communication step is only determined by the communication speed between two neighbouring nodes [19].
- The performance of the controller degrades gracefully if an upper bound on the communication steps is imposed (see Table IV).




Fig. 12: Number of actuation steps required for convergence as a function of  $\gamma$  and of the number of communication steps  $K$ . Network with 3 inverters and  $\alpha = 100$ . Each curve stops, at the right end, at the largest  $\gamma$  that does not cause instability.

TABLE III: Actuation steps required for convergence. For  $K \gg 1$ , parameters are constant and are  $\alpha = 100$ ,  $\gamma = 1/(2\sigma(G))$ . For  $K = 1$ ,  $\alpha$  has been optimized for every instance.

nodes	$K \gg 1$		$K = 1$
	communication steps	actuation steps	actuation steps
3	35	46	258 ( $\alpha = 40$ )
7	566	46	2975 ( $\alpha = 3.1$ )
10	1417	47	5972 ( $\alpha = 1.6$ )
30	14515	51	>30000
100	19848	61	>30000

These findings indicate how the decomposition of the iterative optimization scheme into an iteration that requires actuation of the grid (and therefore cannot be executed too frequently) and a nested sequence of communication steps is fundamental for the overall scalability of the solution.

## VI. CONCLUSION

We implemented for the first time a fully distributed peer-to-peer Volt/VAr controller on a real low-voltage distribution network. The controller at each inverter only uses local voltage measurements and the required model knowledge is only the electrical distance to its neighbors. No filtering or centralized estimation is needed, and the controller is able to drive the system to an optimal point where all voltage and reactive power constraints are satisfied. Each inverter is allowed to communicate only with its neighbors in the electric topology. We also showed that the performance of such a distributed strategy scales nicely with the size of the grid, as long as the communication rate is substantially higher than the measurement/actuation rate. Moreover, we highlight some directions for future investigation, such as optimizing the trade-off between communication complexity and performance, detecting problem infeasibility, and analyzing finite-time convergence of the nested algorithm.

TABLE IV: Actuation steps required for convergence, for different values of  $K$  (communication steps) in a network of 30 nodes.

$K$	100	300	1000	3000	10000	30000
actuation steps	724	211	67	39	50	51

## VII. ACKNOWLEDGMENT

The authors are grateful to Dimitrios Semitsoglou-Tsiapos for his support in adapting his distributed framework to the control algorithm at hand.

## REFERENCES

- [1] S. Bolognani, R. Carli, G. Cavraro, and S. Zampieri, "On the need for communication for voltage regulation of power distribution grids," *IEEE Trans. Control Netw. Syst.*, vol. 6, no. 3, pp. 1111–1123, 2019.
- [2] IEEE 1547-2018, "Standard for interconnection and interoperability of distributed energy resources with associated electric power systems interfaces," Tech. Rep., 2018.
- [3] "VDE-AR-N 4105: Generators connected to the LV distribution network – Technical requirements for the connection to and parallel operation with low-voltage distribution networks," 2018.
- [4] "Commission Regulation (EU) 2016/631 of 14 April 2016 establishing a network code on requirements for grid connection of generators," *OJ*, vol. L112, pp. 1–68, 2016.
- [5] S. Karagiannopoulos, P. Aristidou, and G. Hug, "Data-driven local control design for active distribution grids using off-line optimal power flow and machine learning techniques," *IEEE Trans. Smart Grid*, 2019.
- [6] F. Dörfler, S. Bolognani, J. W. Simpson-Porco, and S. Grammatico, "Distributed control and optimization for autonomous power grids," in *Proc. European Control Conference (ECC)*, Jun. 2019.
- [7] A. Bernstein and E. Dall'Anese, "Real-time feedback-based optimization of distribution grids: A unified approach," *IEEE Trans. Control Netw. Syst.*, vol. 6, no. 3, pp. 1197–1209, Sep. 2019.
- [8] L. Ortmann, A. Hauswirth, I. Caduff, F. Dörfler, and S. Bolognani, "Experimental validation of feedback optimization in power distribution grids," *Electric Power Systems Research*, 2020, to appear.
- [9] S. Bolognani, R. Carli, G. Cavraro, and S. Zampieri, "Distributed reactive power feedback control for voltage regulation and loss minimization," *IEEE Trans. Autom. Control*, vol. 60, no. 4, Apr. 2015.
- [10] F. Olivier, P. Aristidou, D. Ernst, and T. Van Cutsem, "Active management of low-voltage networks for mitigating overvoltages due to photovoltaic units," *IEEE Trans. Smart Grid*, vol. 7, no. 2, pp. 926–936, Mar. 2016.
- [11] S. Magnússon, C. Fischione, and N. Li, "Optimal voltage control using event triggered communication," in *Proc. 10th ACM Intl. Conf. Future Energy Systems*, 2019, pp. 343–354.
- [12] G. Qu and N. Li, "Optimal distributed feedback voltage control under limited reactive power," *IEEE Trans. Power Syst.*, vol. 35, no. 1, pp. 315–331, Jan. 2020.
- [13] S. Magnússon, G. Qu, C. Fischione, and N. Li, "Voltage control using limited communication," *IEEE Trans. Control Netw. Syst.*, vol. 6, no. 3, pp. 993–1003, Sep. 2019.
- [14] S. Magnússon, G. Qu, and N. Li, "Distributed optimal voltage control with asynchronous and delayed communication," *IEEE Trans. Smart Grid*, 2020, in press.
- [15] D. P. Bertsekas, *Nonlinear programming*, 2nd ed. Belmont (MA): Athena Scientific, 1999.
- [16] S. Bolognani and F. Dörfler, "Fast power system analysis via implicit linearization of the power flow manifold," in *53rd Allerton Conf.*, 2015.
- [17] S. Bolognani, "Grid topology identification via distributed statistical hypothesis testing," in *Big Data Application in Power Systems*, R. Arghandeh and Y. Zhou, Eds. Elsevier, 2018, pp. 281–301.
- [18] D. Deka and M. Chertkov, "Topology learning in radial distribution grids," in *Big Data Application in Power Systems*, R. Arghandeh and Y. Zhou, Eds. Elsevier, 2018, pp. 261–279.
- [19] D. Semitsoglou-Tsiapos, "Evaluation of distributed algorithms for electric power systems applications," Master's thesis, Technical University of Denmark (DTU), 2019. [Online]. Available: <https://findit.dtu.dk/en/catalog/2452696264>
- [20] P. Hintjens, *ZeroMQ: Messaging for many applications*. " O'Reilly Media, Inc.", 2013.
- [21] A. M. Prostejovsky, O. Gehrke, A. M. Kosek, T. Strasser, and H. W. Bindner, "Distribution line parameter estimation under consideration of measurement tolerances," *IEEE Transactions on Industrial Informatics*, vol. 12, no. 2, pp. 726–735, 2016.