

A Two-Stage Distributed Architecture for Voltage Control in Power Distribution Systems

Brett A. Robbins, *Student Member, IEEE*, Christoforos N. Hadjicostis, *Senior Member, IEEE*, and Alejandro D. Domínguez-García, *Member, IEEE*

Abstract—In this paper, we propose an architecture for voltage regulation in distribution networks that relies on controlling reactive power injections provided by distributed energy resources (DERs). A local controller on each bus of the network monitors the bus voltage and, whenever there is a voltage violation, it uses locally available information to estimate the amount of reactive power that needs to be injected into the bus in order to correct the violation. If the DERs connected to the bus can collectively provide the reactive power estimated by the local controller, they are instructed to do so. Otherwise, the local controller initiates a request for additional reactive power support from other controllers at neighboring buses through a distributed algorithm that relies on a local exchange of information among neighboring controllers. We show that the proposed architecture helps prevent voltage violations and shapes the voltage profile in radial distribution networks, even in the presence of considerable penetration of variable generation and loads. We present several case studies involving 8-, 13-, and 123-bus distribution systems to illustrate the operation of the architecture.

I. INTRODUCTION

DRIVEN by the US DOE *Smart Grid* initiative, and its European counterpart *Electricity Networks of the Future*, power distribution systems are undergoing radical transformations in structure and functionality. These transformations are enabled by the integration of (i) advanced communication and control, (ii) variable generation, e.g., photovoltaics (PVs), and (iii) storage-capable loads, e.g., plug-in hybrid electric vehicles (PHEVs). These generation and storage resources are commonly referred to as distributed energy resources (DERs).

To date, the relatively small penetration of DERs in distribution systems has allowed regulations pertaining to their control to be limited to (i) maintaining a constant power factor, (ii) following scheduled dispatches from an operator, and (iii) disconnecting from the grid when a fault occurs [1]. However, it has been reported (see, e.g., [2], [3]) that increased penetration of DERs in distribution systems is likely to cause voltage problems, thus requiring additional control mechanisms. This is due to the fact that, unlike transmission systems, typical line reactance to resistance ratios in distribution systems are such that bus voltages are much more sensitive to changes

in active power injections [4], [5]. In this regard, PV-based electricity generation can be highly variable and ramp up on the order of 15% of its capacity per minute across a network with intermittent cloud cover [2]. Additionally, on a clear day, a high penetration of PV installations has the potential to cause voltage rise and over-voltages from a reversal of active power flow originating from net-positive power injections. With respect to this, the University of Illinois solar decathlon house—the *Gable Home* [6]—and the *Equinox* house [7] are examples of residential PV installations capable of producing an amount of power significantly larger than its average load during peak hours of the day. Similarly, the additional power demand introduced by massive charging of PHEVs can potentially cause unacceptable voltage drops [8].

Voltage violations are traditionally handled by automatically-controlled tap-changing under load (TCUL) transformers, set voltage regulators (SVR), and manually-controlled fixed/switched capacitors [2], [9], [10]. However, existing equipment is not inherently designed to handle the variability introduced by DERs, and the lifetime of these components (e.g., the switches and tap changers) could be dramatically reduced due to the increased number of operations that they may undergo [2]. A potential solution to this problem lies on the utilization of DERs to provide reactive power support for voltage regulation through the proper control of the power electronics that interface them with the grid [11]–[13]. In this paper, we pursue this idea and propose a two-stage voltage control architecture that relies on controlling and coordinating the reactive power provided by DERs. In this regard, the commercial products described in [13], [14] are examples of existing rooftop and pole-mount PV solutions capable of providing reactive power support; these products have wireless communication capabilities that allow them to be controlled through cellular, Ethernet, or WiMax backhaul networks.

Our proposed architecture consists of two stages. In the first stage, the voltages at certain buses in the network are monitored by a local controller. If the local controller at a particular bus senses that its voltage is above or below certain thresholds imposed by performance specifications (e.g., $\pm 5\%$ around a nominal value [1]), it will first estimate the amount of reactive power that should be injected into the bus to clear the voltage violation. This estimate is obtained by using the sensitivity of the bus voltage magnitude to changes in reactive power injections in the same bus. Then, if the DERs directly connected to the bus can provide the estimated reactive

B. A. Robbins and A. D. Domínguez-García are with the Department of Electrical and Computer Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, 61801. E-mail: {robbs3, aledan}@ILLINOIS.EDU.

C. N. Hadjicostis is with the Department of Electrical and Computer Engineering, University of Cyprus, Nicosia, Cyprus. E-mail: chadjic@UCY.AC.CY.

This work was supported in part by the National Science Foundation (NSF) under grant ECCS-1135598 and Career Award ECCS-CAR-0954420; and by the Consortium for Electric Reliability Technology Solutions (CERTS).

power, they will be instructed to do so; otherwise, they will output their maximum/minimum capacity. In the second stage, the difference in reactive power between the local controller estimate and what the DERs connected to the bus can provide will be requested from other buses that have additional capacity. Through a distributed algorithm that only requires an exchange of information among neighboring controllers (e.g., through wireless or power line communications), each local controller calculates its fair contribution to meet the additional request.

Other (centralized) solutions to the voltage control problem in distribution networks have been proposed in [15]–[17]. For example, the authors in [15] propose an optimal multi-agent scheme that provides reactive power support in distribution feeders, and assumes that DERs have two-way communication with a single controller, either directly or through other DERs. In order to correct limit violations, agents are assigned to be managers or contractors that bid on reactive power contributions based on bus sensitivities. In [16], the authors partition the system buses into groups (agents) and solve local optimal power flow problems through a hierarchical chain of command structure. The authors in [17] maintain a database of limited historical observations and the corresponding solutions to a nonlinear optimal power flow problem so as to avoid computing its solution each time. In all the solutions described above, it is necessary to overlay a communication network connecting a central controller with each DER and requires the controller to know the available DERs at all times. By contrast, our control architecture offers many potential benefits over a centralized one: (i) it is more economical because it does not require a significant communication infrastructure overlay, (ii) computational requirements for the local controller are relatively low, and (iii) local information is sufficient to control the DERs.

It is worth noting that there are also several decentralized or distributed strategies that address the voltage regulation problem in distribution networks [18], [19]. The authors of these works propose a switching control scheme where the DERs are operated with a constant power factor while bus voltages are within specifications, and, whenever there is a voltage violation, their reactive power is adjusted so that the system returns to the desired operating conditions. This is similar to the first stage of the proposed control scheme, but lacks the additional support provided by the second stage.

In subsequent developments, we assume that the local controller on each node has aggregate information on the total reactive power capacity available (upper and lower limits) from DERs connected to its bus. This collective upper (lower) capacity limit is determined by the sum of the individual DER upper (lower) limits, which the local controller needs to obtain. In practice, this can be accomplished in a variety of ways. One possibility is to have the local controller directly communicate with each individual DER. In such case, each DER can directly report its individual capacity limits, which can vary over time depending on the specific operating conditions of the DER. An alternative to the above approach can be implemented using a variation of the algorithm used in the second stage, which can

be modified to coordinate the DERs on each bus in a distributed fashion; for further details, the reader is referred to [20], [21].

The remainder of this paper is organized as follows. Section II provides the distribution system model used for control design purposes, and the communication network model used to describe the exchange of information between local controllers. The proposed two-stage control architecture is presented in Section III, while Section IV illustrates its operation on an 8-bus network. Section V presents a case study involving the IEEE 123-bus system with a high penetration of PVs and PHEVs. Section VI discusses extensions to unbalanced three-phase systems. Concluding remarks are presented in Section VII.

II. SYSTEM MODEL

In this section, we develop a power distribution system model, which is used in Section III to design the voltage control system; this model describes the evolution of bus voltage magnitudes as active and reactive power injections change over time. Additionally, we introduce the network communication model that describes the exchange of information between the local controllers that are geographically dispersed throughout the electrical network.

A. Power Distribution Network

Consider a distribution system with $n + 1$ buses indexed by $i = 0, 1, \dots, n$. At time instants $r = 0, 1, \dots$, the voltage magnitude and angle of bus i are denoted by $V_i[r]$ and $\theta_i[r]$, respectively. We assume that bus $i = 0$ is the feeder and will be treated as an infinite bus, therefore $V_0[r]$ and $\theta_0[r]$ remain constant for all r . The remaining n buses are considered to be PQ buses. Let $V[r] = [V_1[r], V_2[r], \dots, V_n[r]]^T$ denote the vector of bus voltage magnitudes and $\theta[r] = [\theta_1[r], \theta_2[r], \dots, \theta_n[r]]^T$ denote the vector of bus voltage angles (both $V_0[r]$ and $\theta_0[r]$ are omitted). At time instant r , let $P_i[r]$ and $Q_i[r]$ be the active and reactive power injections in bus i , respectively; the corresponding active and reactive power injections vectors at the PQ buses are denoted by $P[r] = [P_1[r], P_2[r], \dots, P_n[r]]^T$ and $Q[r] = [Q_1[r], Q_2[r], \dots, Q_n[r]]^T$. We define $\Delta V[r] = V[r+1] - V[r]$ and $\Delta\theta[r] = \theta[r+1] - \theta[r]$ as the vectors describing small variations in voltage magnitudes and angles between times r and $r + 1$; while variations in active and reactive power injections at PQ buses are defined as $\Delta P[r] = P[r+1] - P[r]$ and $\Delta Q[r] = Q[r+1] - Q[r]$. Then,

$$\begin{bmatrix} \Delta P[r] \\ \Delta Q[r] \end{bmatrix} = \begin{bmatrix} H & N \\ K & L \end{bmatrix} \begin{bmatrix} \Delta\theta[r] \\ \Delta V[r] \end{bmatrix}, \quad (1)$$

where

$$H = \left[\frac{\partial P_i}{\partial \theta_j} \right], N = \left[\frac{\partial P_i}{\partial V_j} \right], K = \left[\frac{\partial Q_i}{\partial \theta_j} \right], L = \left[\frac{\partial Q_i}{\partial V_j} \right].$$

[Note that all of the partial derivatives defining the entries of H , N , K , and L are evaluated at $V[r]$, $\theta[r]$, $P[r]$, $Q[r]$, and therefore H , N , K , and L are functions of r ; however, in the remainder, we suppress the argument for ease of notation.]

A standard assumption used in the analysis of transmission systems is that the entries of H , L are much larger than the

entries of N , K . This effectively decouples (1) so that variations in active power injections primarily affect bus voltage angles, whereas variations in reactive power injections directly affect bus voltage magnitudes. This is a consequence of the fact that the per unit reactance to the per unit resistance ratio of transmission lines, commonly referred to as “ x/r ratio”, is large [22]. In the case of a distribution system, this assumption is not valid since the transmission line x/r ratios are much lower [4]. As a result, in distribution systems, bus voltages are much more sensitive to changes in active power than typically observed in transmission systems.

In our setting, $\Delta P[r]$ describes the changes in active power injections that arise from DERs, e.g., PV rooftop installations and PHEVs, and represents an external “disturbance” over which we do not have control. These injections will have a noticeable impact on the network voltage profile. To mitigate the effect of $\Delta P[r]$ on system voltages, we assume that we have control over $\Delta Q[r]$. We are interested in the effect that uncontrolled variations in active power and controlled variations in reactive power injections have on bus voltage magnitudes. Assuming that H is invertible, then it follows from (1) that $\Delta\theta[r] = -H^{-1}N\Delta V[r] + H^{-1}\Delta P[r]$, and

$$\begin{aligned} \Delta V[r] &= (L - KH^{-1}N)^{-1}(\Delta Q[r] - KH^{-1}\Delta P[r]) \\ &= S\Delta Q[r] + w[r], \end{aligned} \quad (2)$$

where $S \equiv (L - KH^{-1}N)^{-1}$ is assumed to be invertible, and $w[r] \equiv -(L - KH^{-1}N)^{-1}KH^{-1}\Delta P[r]$ captures the effect of uncontrolled variations in active power injections on bus voltage magnitude. Now, by unwrapping (2), the recurrence relation that describes how the bus voltage magnitudes evolve with time is given by

$$V[r+1] = V[r] + S\Delta Q[r] + w[r]. \quad (3)$$

Although we did not make it explicit, the matrix S is in general a function r . On the other hand, the variations of S with r are relatively small for a wide range of operating conditions [22]. In this regard, for the test systems used in the case studies of Sections IV–VI, we verified that, for a wide range of operating conditions, the entries of S typically remain within 3% of their average value.

B. Network Communication

It is assumed that certain buses of the electrical network have a local controller that can monitor the bus voltage and make local control decisions based on the exchange of information with a subset of other controllers. Some of these controllers may be located at buses that are directly connected to the bus of the given controller, but, in general, the exchange of information between the n controllers can be arbitrary. It is convenient to capture this exchange of information between local controllers by a *directed graph* $\mathcal{G}_d = \{\mathcal{V}, \mathcal{E}\}$, where $\mathcal{V} = \{1, 2, \dots, n\}$ represents the set of vertices (nodes, which represent the controllers), and $\mathcal{E} \subseteq \mathcal{V} \times \mathcal{V}$ represents the set of directed edges, i.e., $(j, i) \in \mathcal{E}$ when node j can receive




Figure 1: Timeline for the first and second stages.

information from node i . By convention, we assume that self-loops are not contained in \mathcal{E} . All of the nodes that can send information to node j are said to be the in-neighbors of node j and are represented by the set $\mathcal{N}_j^- = \{i \in \mathcal{V} : (j, i) \in \mathcal{E}\}$. The cardinality of \mathcal{N}_j^- is called the in-degree of node j and is denoted by \mathcal{D}_j^- . The nodes that can receive information from node j are called its out-neighbors and are represented by the set $\mathcal{N}_j^+ = \{l \in \mathcal{V} : (l, j) \in \mathcal{E}\}$; the out-degree of node j is \mathcal{D}_j^+ . A directed graph is considered *strongly connected* if any two vertices $i, j \in \mathcal{V}$, $i \neq j$, can be joined by a path that starts at node i and ends at node j .

III. TWO-STAGE VOLTAGE CONTROL ARCHITECTURE

In this section, we first provide an overview of the proposed two-stage voltage control architecture. Then we formulate each stage that comprises the architecture and analyze their stability and convergence properties.

A. Overview

Figure 1 shows the timeline for the operation of the two stages. Without loss of generality, assume that at each bus j of the network, there is a local controller that monitors the bus voltage magnitude V_j . At fixed time instants $r = \lfloor k/k_0 \rfloor$, $k = 0, 1, \dots$, for some sufficiently large k_0 (to be precisely defined later), each local controller executes the first stage. In this stage, if the local controller on bus j detects a voltage violation, i.e., V_j is outside specifications, it estimates the amount of reactive power that, if injected (or consumed) in bus j , will correct the violation; then, if the DERs connected to bus j collectively have the capacity to inject (or consume) the estimated reactive power request, the controller instructs these DERs to do so. [In Section I we discussed two alternatives for how the local controller connected to bus j can, a priori, determine the DER collective upper/lower capacity limits.] Otherwise, the local controller will execute the second stage by initiating a request for an additional amount of reactive power to be injected in other buses of the network to help clear its voltage violation. This request is relayed by the local controller to other neighboring local controllers through a distributed iterative algorithm that ensures that nodes with available reactive power capacity will provide additional support.

The second stage exploits the sensitivity of V_j to reactive power injections in neighboring buses and has a net effect of a globally homogeneous raising (or lowering) of all the bus voltages. This effect is similar to the one that results from adjusting the taps on a TCUL transformer or a SVR

in the sense that it affects all bus voltages; however, there are some differences. In particular, the action of a TCUL, or a SVR, will *uniformly* raise (or lower) the voltage across the network. In the two-stage control architecture, the reactive power injections in bus j will primarily affect the voltage at this bus and the voltages at buses downstream of it (in a radial system), with the effect attenuating as we move towards the feeder. It is important to note that the proposed voltage control architecture is not intended to replace current voltage control systems, but rather to supplement their action while (i) minimizing their usage by handling faster voltage variations due to changes in renewable-based power injections, and (ii) having them intervene only during extreme circumstances rather than minor, possibly temporary, violations.

B. First Stage Control

At time instants $r = 0, 1, \dots$, the local controller located on bus j measures the voltage $V_j[r]$. If it detects a voltage violation (i.e., $V_j[r] \notin [\underline{V}_j, \overline{V}_j]$, where \overline{V}_j and \underline{V}_j are bus j 's upper and lower voltage limits, respectively), then the local controller will estimate the amount of additional reactive power $\rho_j[r]$ needed to clear the violation. In the first stage, we assume that the controller at bus j does not have voltage information for other buses available. Therefore, the estimate of $\rho_j[r]$ is given by

$$\rho_j[r] = \begin{cases} \frac{\alpha}{s_{jj}}(\underline{V}_j - V_j[r]), & V_j[r] < \underline{V}_j, \\ 0, & \underline{V}_j \leq V_j[r] \leq \overline{V}_j, \\ \frac{\alpha}{s_{jj}}(\overline{V}_j - V_j[r]), & \overline{V}_j < V_j[r], \end{cases} \quad (4)$$

where $\alpha > 0$ is some constant (to be made precise later) and $s_{jj} = \partial V_j / \partial Q_j$. Thus, if the DERs connected to each bus j collectively have the capacity (with respect to their previous demand) to provide $\rho_j[r]$, then the net change in reactive power injections for all of the nodes in the network is given by $\Delta Q[r] \equiv \rho[r]$, where $\rho[r] = [\rho_1[r], \rho_2[r], \dots, \rho_n[r]]^T$ (assuming that the reactive power consumed by loads does not substantially change). Next, we provide a condition on the value that α must take to ensure the stability of the closed-loop system that results from applying $\Delta Q[r] \equiv \rho[r]$ to (3).

1) *Choice of α for stability:* Note from (4) that, by defining

$$V_j^{ref}[r] = \begin{cases} \underline{V}_j, & V_j[r] < \underline{V}_j, \\ V_j[r], & \underline{V}_j \leq V_j[r] \leq \overline{V}_j, \\ \overline{V}_j, & \overline{V}_j < V_j[r], \end{cases} \quad (5)$$

we can write

$$\Delta Q[r] \equiv \rho[r] = \alpha D (V^{ref}[r] - V[r]), \quad (6)$$

where $V^{ref}[r] = [V_1^{ref}[r], V_2^{ref}[r], \dots, V_n^{ref}[r]]^T$, $\alpha > 0$, and D is a diagonal matrix with $d_{ii} = 1/s_{ii}$. Then, by substituting (6) into (3), we get

$$V[r+1] = (I - \alpha SD) V[r] + \alpha SD V^{ref}[r] + w[r], \quad (7)$$

with $V^{ref}[r]$ as defined in (5), from where it is easy to see that $V^{ref}[r]$ is bounded for all r . Also, from the definition of $w[r]$ in (3), it is obvious that $w[r]$ is also bounded for all r .

Then, since the system (7) is a linear time-invariant system driven by bounded inputs $V^{ref}[r]$ and $w[r]$, ensuring the stability of this system, i.e., that $V[r]$ remains bounded for all r , is equivalent to ensuring that the system is internally stable (see, e.g., [23]), i.e., the eigenvalues of $(I - \alpha SD)$ must lie within the unit circle. This can be accomplished by choosing α such that

$$\alpha < \alpha_c = \min_i \left\{ \frac{2 \operatorname{Re}\{\lambda_i\}}{|\lambda_i|^2} \right\}. \quad (8)$$

where $\lambda_i = \operatorname{Re}\{\lambda_i\} + j \operatorname{Im}\{\lambda_i\}$ denotes the i^{th} eigenvalue of SD . The derivation of (8) is included in Appendix B.

Remark 1: The developments above assume a fixed electrical network configuration that results in a single S ; however, the control system should be able to adapt to configuration changes (potentially resulting in different S 's) and ensure that the system in (7) is stable for all possible configurations. In this regard, we envision that the local controller at each bus j could have a database, calculated off-line, with the value of s_{jj} and α corresponding to each network configuration. Then, upon a change in configuration, the local controllers would be notified, and they would update α and the s_{jj} 's accordingly. \square

2) *Practical considerations for implementation:* The action of (6) on the system dynamics as defined by (3), which results in (7), is equivalent to those of a discrete-time integrator. Thus, whenever there is a voltage violation at bus j , the action of the first stage controller will asymptotically drive the voltage V_j to either \overline{V}_j or to \underline{V}_j depending on the nature of the violation. In practice, it is desirable that the first stage controller stops iterating after a finite number of steps r_0 . In order to achieve this, in (4), we replace \overline{V}_j by $\overline{V}_j - \varepsilon_1$ and \underline{V}_j by $\underline{V}_j + \varepsilon_1$ for some $\varepsilon_1 > 0$ small. The result is such that, for any $\varepsilon_1 > 0$, there is some finite r_0 such that $V_j[r] \in [\underline{V}_j, \overline{V}_j]$ for all $r \geq r_0$ and all j ; thus the controller action stops after r_0 steps.

In all our numerical experiments, we verified that letting $\alpha = 1/n$ (where n is the number of buses in the network with the feeder omitted) be the default gain satisfies the condition in (8). This choice of α may not necessarily result in the shortest settling time; however, it helps prevent overshoots if several local controllers are acting simultaneously. In this regard, we have observed that S is in general a full matrix whose entries are usually on the same order of magnitude, except for those associated with buses that are furthest apart, i.e., buses at the end of a sublateral have a low voltage sensitivity with respect to injections into buses near the feeder, and vice versa. This implies that injections of reactive power at any bus affect (to a more or less extent) all other bus voltages throughout the network. Furthermore, for typical distribution system line parameter values, the column and row sums of SD are typically well below $1.5n$, so scaling the contribution of each bus by $1/n$ ensures that the spectral radius of SD is less than 2 [24]. This accounts for the worst-case scenario in which all the buses are subject to either under-voltage or over-voltage violations. Finally, it is worth noting that the local controllers can easily estimate n using a distributed algorithm similar to the one to be described in Section III-C.

C. Second Stage Control

If the reactive power estimate $\rho_j[r]$ is within the limits that the DERs connected to bus j can provide for all j , then the second stage is not required at step r . Otherwise, the second stage compensates for the capacity constraint violations from the first stage by adjusting every node's contributions to globally raise, or lower, bus voltages across the network. This is accomplished through a distributed algorithm that relies on a local exchange of information among neighboring controllers. We assume that the graph that describes the exchange of information between local controllers is strongly connected, but not necessarily complete, i.e., it is not necessarily the case that each node can communicate directly with every other node in the graph. [If the graph was complete, then we could easily use a centralized dispatch strategy.]

Let $\underline{q}_j, \bar{q}_j$, where $\underline{q}_j \leq 0 \leq \bar{q}_j$, be the total lower and upper limits on the amount of reactive power that DERs at bus j can provide. Let $q_j[r]$ be the amount¹ of reactive power provided by the DERs connected to bus j at instant r . Then, assuming that $\underline{q}_j \leq q_j[r] \leq \bar{q}_j, \forall j$, the total estimated reactive power to be provided by node j is

$$\hat{q}_j[r+1] = q_j[r] + \rho_j[r], \quad (9)$$

where $\rho_j[r]$ is the estimate from (4). Then, if $\underline{q}_j \leq \hat{q}_j[r+1] \leq \bar{q}_j, \forall j$, the second stage is not necessary, i.e., every node j can provide itself the amount of reactive power estimated in the first stage. Otherwise, whenever $\hat{q}_j[r+1] \geq \bar{q}_j$ or $\hat{q}_j[r+1] \leq \underline{q}_j$, for some j (which means that at least one node cannot correct its voltage violation by itself), the buses that have additional capacity will calculate the amount of reactive power they need to provide as an attempt to raise the voltage in the network through the distributed iterative algorithm described next.

1) *Second stage algorithm:* Let $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ be a strongly connected directed graph describing the exchange of information between local controllers. Each node $j \in \mathcal{V}$ maintains three auxiliary variables $\mu_j[k]$, $\bar{\nu}_j[k]$, and $\underline{\nu}_j[k]$, and updates them to $\mu_j[k+1]$, $\bar{\nu}_j[k+1]$, and $\underline{\nu}_j[k+1]$, respectively, via a weighted linear combination of their previous $\mu_j[k]$, $\bar{\nu}_j[k]$, and $\underline{\nu}_j[k]$, respectively, and those of its in-neighbors, i.e., $\{\mu_i[k] \mid i \in \mathcal{N}_j^-\}$, $\{\bar{\nu}_i[k] \mid i \in \mathcal{N}_j^-\}$, and $\{\underline{\nu}_i[k] \mid i \in \mathcal{N}_j^-\}$, respectively; specifically,

$$\mu_j[k+1] = \sum_{i \in \{\mathcal{N}_j^-\} \cup \{j\}} \frac{1}{1 + \mathcal{D}_i^+} \mu_i[k], \quad (10)$$

$$\bar{\nu}_j[k+1] = \sum_{i \in \{\mathcal{N}_j^-\} \cup \{j\}} \frac{1}{1 + \mathcal{D}_i^+} \bar{\nu}_i[k], \quad (11)$$

$$\underline{\nu}_j[k+1] = \sum_{i \in \{\mathcal{N}_j^-\} \cup \{j\}} \frac{1}{1 + \mathcal{D}_i^+} \underline{\nu}_i[k], \quad (12)$$

¹In (3), we defined $\Delta Q_j[r] = Q_j[r+1] - Q_j[r]$, where $Q_j[r]$ is the total reactive power injection (with appropriate sign) in bus j that arises from both DERs and loads, i.e., $Q_j[r] = q_j[r] + q_j^L[r]$, where $q_j^L[r]$ denotes the (uncontrolled) reactive power injection arising, e.g., from loads, which is assumed to remain constant, i.e., $\Delta Q[r] = Q[r+1] - Q[r] = q[r+1] - q[r]$.

where \mathcal{D}_i^+ is the out-degree of node i . Each node j sets its initial conditions in (10)–(12) respectively to

$$\mu_j[0] = \begin{cases} \hat{q}_j[r+1] - \bar{q}_j, & \hat{q}_j[r+1] > \bar{q}_j, \\ \hat{q}_j[r+1] - \underline{q}_j, & \hat{q}_j[r+1] < \underline{q}_j, \\ 0, & \text{otherwise,} \end{cases} \quad (13)$$

$$\bar{\nu}_j[0] = \begin{cases} \bar{q}_j - \hat{q}_j[r+1], & \hat{q}_j[r+1] < \bar{q}_j, \\ 0, & \hat{q}_j[r+1] \geq \bar{q}_j, \end{cases} \quad (14)$$

$$\underline{\nu}_j[0] = \begin{cases} \underline{q}_j - \hat{q}_j[r+1], & \hat{q}_j[r+1] > \underline{q}_j, \\ 0, & \hat{q}_j[r+1] \leq \underline{q}_j. \end{cases} \quad (15)$$

Then, at every step k , for each j such that $\underline{\nu}_j[k] \neq 0$ or $\bar{\nu}_j[k] \neq 0$, the corresponding local controller computes

$$\eta_j[k] = \begin{cases} \frac{\mu_j[k]}{\underline{\nu}_j[k]} \underline{\nu}_j[0], & \mu_j[k] < 0, \\ \frac{\mu_j[k]}{\bar{\nu}_j[k]} \bar{\nu}_j[0], & \mu_j[k] \geq 0, \end{cases} \quad (16)$$

where $\eta_j[k]$ will asymptotically converge to

$$\eta_j = \begin{cases} \frac{\sum_{i=1}^n \mu_i[0]}{\sum_{i=1}^n \underline{\nu}_i[0]} \underline{\nu}_j[0], & \lim_{k \rightarrow \infty} \mu_j[k] < 0, \\ \frac{\sum_{i=1}^n \mu_i[0]}{\sum_{i=1}^n \bar{\nu}_i[0]} \bar{\nu}_j[0], & \lim_{k \rightarrow \infty} \mu_j[k] \geq 0; \end{cases} \quad (17)$$

the derivation of this result can be found in Appendix C.

From (17), it is obvious that $\sum_{j=1}^n \eta_j = \sum_{j=1}^n \mu_j[0]$, i.e., the total amount of reactive power (including both positive and negative contributions) that constrained nodes cannot provide remains asymptotically constant. Also, $\eta_j \leq \bar{\nu}_j[0]$ if $\lim_{k \rightarrow \infty} \mu_j[k] \geq 0$ and $\eta_j > \underline{\nu}_j[0]$ if $\lim_{k \rightarrow \infty} \mu_j[k] < 0$. Finally, bus j adjusts its reactive power contribution to

$$q_j[r+1] = \begin{cases} \bar{q}_j, & \hat{q}_j[r+1] + \eta_j > \bar{q}_j, \\ \underline{q}_j, & \hat{q}_j[r+1] + \eta_j < \underline{q}_j, \\ \hat{q}_j[r+1] + \eta_j, & \text{otherwise.} \end{cases} \quad (18)$$

From the above development, it follows that the additional reactive power to be requested in the second stage will either be a net injection or a net consumption. In this regard, it is reasonable to assume that a distribution network will typically experience one type of voltage violation at any given moment. Simultaneous over- and under-voltage violations would imply that the distribution system lines have substantial losses, which is unlikely in real system; however, we assume that this is a possibility. Thus, each node j maintains $\bar{\nu}_j[k]$ and $\underline{\nu}_j[k]$, and computes the appropriate solution in (17) once $\mu_j[k]$ converges.

2) *Effect of network connectivity on convergence speed:* For any strongly connected graph $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$ describing the exchange of information between local controllers, the steady-state solution of the distributed algorithm in (17) is independent of \mathcal{G} . However, for a given size of the vertex set \mathcal{V} , the connectivity between the nodes as described by \mathcal{E} , which determines the weights in (10)–(12), affects the convergence speed of the distributed algorithm. In this regard, by letting $\mu[k] =$

$[\mu_1[k], \mu_2[k], \dots, \mu_n[k]]^T$, $\bar{v}[k] = [\bar{v}_1[k], \bar{v}_2[k], \dots, \bar{v}_n[k]]^T$, and $\underline{v}[k] = [\underline{v}_1[k], \underline{v}_2[k], \dots, \underline{v}_n[k]]^T$, and defining a matrix $P = [p_{ji}]$, with

$$p_{ji} = \begin{cases} \frac{1}{1+\mathcal{D}_j^+}, & j = i, \\ \frac{1}{1+\mathcal{D}_i^+}, & j \neq i, (j, i) \in \mathcal{E}, \\ 0, & j \neq i, (j, i) \notin \mathcal{E}, \end{cases} \quad (19)$$

then, we can write the iterations in (10)–(12) in matrix form (as in (20)–(22) in Appendix C). The matrix P is column stochastic and has the same sparsity structure (except for the diagonal entries) as the adjacency matrix of \mathcal{G} . Furthermore, since \mathcal{G} is strongly connected, the matrix P is primitive [24]. Now, it is well-known (see, e.g., [24]) that, for some accuracy level ε_0 , the second largest eigenvalue modulus $|\lambda_2|$ of P determines the number of iterations k_0 after which $\|\mu[k] - \mu\|_\infty \leq \varepsilon_0$ and $\|\bar{v}[k] - \bar{v}\|_\infty \leq \varepsilon_0$ ($\|\underline{v}[k] - \underline{v}\|_\infty \leq \varepsilon_0$), $\forall k \geq k_0$.

In general, the more connected the graph is the faster the algorithm converges, i.e., the smaller k_0 is; however, in order to make a quantitative statement, it is necessary to check the value of $|\lambda_2|$. On the other hand, there are results in the spectral graph theory literature (see, e.g., [25]) that establish upper bounds on $|\lambda_2|$ in terms of the number of nodes and the diameter of \mathcal{G} , e.g., the maximum shortest path between any two nodes. In order to determine the execution time of the algorithm, we need to fix the time τ_0 for completing each iteration. In this regard, we have shown in [26] that using commercial off-the-shelf hardware we can complete an iteration step τ_0 within 10 to 40 ms; should the hardware be designed specifically for this application, the iteration step time τ_0 might be decreased even further. Then, given τ_0 and k_0 , the actual time it takes for the distributed algorithm to converge is $\tau_1 = k_0\tau_0$, which in turn determines the minimum time between actions of the first stage controller. Note that in the case studies that follow, we ensure that τ_1 is much smaller than the reaction time² for TCULs.

IV. EXAMPLE: 8-BUS DISTRIBUTION SYSTEM

Consider the 8-bus system shown in Fig. 2; operational requirements specify that bus voltage magnitudes must lie within $\pm 5\%$ of 1 p.u. at all times. Line impedance data, system loading data, and the aggregated reactive power capacity limits of the DERs on each bus can be found in Appendix D, from which the matrices S and D in (7) can be obtained. Then, given S and D , and following the notation in (8), we obtain that $\alpha_c = 0.4$, thus, for the first stage, the gain α of each local controller needs to be chosen so that $\alpha < \alpha_c$.

A. Second Stage Implementation

We consider the two communication network topologies displayed in Fig. 3. Topology 1 (Fig. 3(a)) mirrors the physical network with undirected communication links between the

²This partially explains why the effect of TCULs was not included in the case studies, since we naturally wanted to focus on the effectiveness of the proposed control scheme; nevertheless, future work will consider the coupling of TCULs with the proposed control scheme.





Figure 2: 8-bus system: electrical network graph



(a) Topology 1.

(b) Topology 2.


Figure 3: 8-bus system: communication network graphs.

controllers of any two buses that are electrically connected by a line. In Topology 2 (Fig. 3(b)), we add bidirectional communication links between nodes 1 & 7 and nodes 5 & 7, as well as make the links between nodes 1 & 2 and nodes 4 & 5 directional (directed from 1 to 2, and from 4 to 5, respectively). For each topology, the matrix P , with entries as defined in (19), is given in Appendix D. We assume that each iteration of the distributed algorithm requires $\tau_0 = 10$ ms, which is consistent with the experimental findings in [26] (see also discussion in Section III-C2). For an accuracy level of $\varepsilon_0 = 10^{-3}$, the algorithm needs $k_0 = 39$ iterations to converge when implemented over Topology 1, and $k_0 = 17$ iterations when implemented over Topology 2. Thus, the time that it takes for the algorithm to converge is 0.39 s for Topology 1 and 0.17 s for Topology 2.


B. System Response for Different Scenarios

Next, we illustrate the operation of the voltage control architecture for both over-voltage and under-voltage violations. Additionally, for the under-voltage violation, and assuming different values of α and ε_1 , we study the time it takes for the control system to correct this violation.

1) *Feeder over-voltage*: We consider a scenario in which bus 1 is subject to an over-voltage violation; this could potentially arise if there is a severe over-voltage in the transmission network and the transformers at the substation could no longer adjust their taps to lower the bus voltages. In this case, we assume the feeder's voltage is $1.08\angle 0$ p.u., which results in $V_1 = 1.0642$ p.u. We set the gain of each controller to be $\alpha = 0.3 < \alpha_c = 0.4$. When the second stage is implemented over the network in Fig. 3(a), the evolution of the voltage profile in the network is displayed in Fig. 4(a) for the first 25 s. The corresponding evolution of reactive power injections that result from the combined action of first and second stages is displayed in Fig. 4(b), whereas the individual responses of both stages are displayed in Fig. 5 for the first 3 s. In this case, the first stage controller in bus 1 tries to fix the voltage violation by



(a) Voltage response.



(b) Reactive power support.

Figure 4: 8-bus system: response to over-voltage.





(a) First stage control \hat{q} .(b) Second stage control η .

Figure 5: 8-bus system: first and second stage response to over-voltage.

demanding local consumption of reactive power; as it can be seen in Fig. 4(b), this results in the reactive power provided by bus 1 reaching its minimum capacity within 1 s. Then, the action of the second stage makes the local controllers in the other buses react in order to bring down the voltage magnitude of bus 1, thus reactive power starts being consumed in these nodes. This causes an under voltage violation in buses 5 and 7. However, the violation is corrected by subsequent actions of the controllers, and the voltage in all buses is restored to values within $\pm 5\%$ of 1 p.u. after approximately 15 s.



(a) Voltage response.



(b) Reactive power support.

Figure 6: 8-bus system: system response to under-voltage.


(a) η for Topology 1.(b) η for Topology 2.

Figure 7: 8-bus system: distributed algorithm response to under-voltage.

2) *Sublateral under-voltage*: In this case, the feeder voltage is set to $0.98\angle 0$ p.u., which results in under-voltages on buses 2 through 7, with the lowest voltage magnitude of 0.8830 p.u. on bus 5. For the first stage, we set the gain of each local controller to be $\alpha = 0.22$. Figure 6(a) shows the evolution of the voltage profile in all nodes when the second stage is implemented over the network in Fig. 3(b). In this case, we can see that it takes about 2.1 s to bring the bus voltage magnitudes within the levels specified by operational requirements ($\pm 5\%$ of 1 p.u.). Figure 6(b) shows, as time evolves, the reactive power injections on each of the buses that result from the combined actions of first and second stages. In this figure, we can see that the local controllers on buses 2 through 7 swiftly begin to demand reactive power so as to raise their bus voltages. On the other hand, since bus 1 is initially within its voltage limits, the reactive power injection on this bus does not start until other buses reach their reactive power capacity limits. Then, the second stage starts demanding reactive power from any bus that has available capacity, which includes bus 1; thus, the evolution of q_1 is entirely determined by the action of the second stage.

Table I: 8-BUS SYSTEM: SETTLING TIME FOR DIFFERENT PARAMETER VALUES.

Accuracy ε_1	Controller Gain α	Topology	
		Time 1 [s]	Time 2 [s]
10^{-2}	$1/n$	1.17	0.51
	0.22	0.39	0.17
	α_c	0.39	0.17
10^{-3}	$1/n$	5.46	2.38
	0.22	2.34	1.02
	α_c	0.39	0.17
10^{-4}	$1/n$	9.36	4.08
	0.22	4.68	2.04
	α_c	0.39	0.17

For both communication topologies, Table I compares the settling times for different values of ε_1 and α . From the results in this table and Fig. 7, it is clear that the system response is always faster when the second stage is implemented using the communication topology in Fig 3(b), where Fig. 7 compares the first two intervals of the second stage for Figs. 3(a) and 3(b), respectively. Additionally, the time it takes for the system to correct a violation substantially decreases when α is chosen to be closer to its critical value.

V. CASE STUDY: 123-BUS DISTRIBUTION SYSTEM

In this section, we illustrate the performance of the two-stage control architecture in a 123-bus system with a significant penetration of PV installations and loading from PHEVs.

A. Modified IEEE 123-bus Distribution System

The system for this case study was adapted from the IEEE 123-bus test system [27]. The one-line diagram is displayed in Fig. 8, where the 85 load buses are denoted by triangular-shaped markers. The original IEEE 123-bus system was modeled as a






Figure 8: 123-bus system: electrical network graph.



(a) Seven day load curve.



(b) PV generation and PHEV active load.

Figure 9: 123-bus system: seven-day loading, PV generation, and PHEV charging profile.


three phase unbalanced network, but was simplified to be analyzed as a single phase network (for the system data, see [28]); however, as discussed in Section VI, the results can be extended to unbalanced systems. The communication network used to implement the second stage mirrors the physical network with the additional communication links that are represented as red dashed lines in Fig. 8. This results in $k_0 = 600$ iterations compared to $k_0 = 5750$ without the additional links. As in the 8-bus system, we assume that $\tau_0 = 10$ ms, so it takes $\tau_1 = k_0\tau_0 = 6$ s for the algorithm to converge. Operational requirements constrain voltage magnitudes to lie within $\pm 5\%$ of 1 p.u.

We held the taps constant for all of the TCUL transformers and automatic SVRs for the following reasons: (i) cases are structured such that additional injections/loads will induce over- and under-voltages, (ii) it is assumed that the ability of the nodes to react to voltage violations is faster than that of the TCUL transformer (in fact, we ensure this by the particular choices of k_0 and τ_0 above), and (iii) the purpose of the proposed control method is to minimize the execution of traditional voltage control strategies, so we are interested in studying the ability of our control architecture to correct violations³ without additional help from the TCULs.


B. Seven-Day Scenario

For a time frame of seven days, we analyze the impact of loading, variable PV generation, and PHEV charging on the system voltage profile when (i) there is no voltage control, and (ii) the two-stage voltage control system is implemented.

³Certainly, it is worthwhile for future work to carefully study the coupling between the proposed control mechanism and TCULs, however, this is beyond the scope of this article.



(a) Uncontrolled voltage response.




(b) Controlled voltage response.

Figure 10: 123-bus system: voltage profile evolution with and without control.


1) Loading, PV generation, and PHEV charging profiles:

In order to develop a realistic loading profile, we utilized historical data about consumer demand gathered from the PJM interconnection [29], which offers 18 years of metered data supplied by their electric distribution companies. Figure 9(a) shows the average hourly diversified demand for the PJM Mid-Atlantic Region from the last week of 2011 normalized so that the peak value of the curve is 1. Then, in order to properly scale the profile in Fig. 9(a), we assumed that all of the loads in our case study are residential, so that we can compare the total energy used on each load bus with the average residential energy consumption to determine the number of consumers located at that bus. A typical American household consumed, on average, 953 kWh per month in 2010 [30]. Since weekly behaviors are cyclic, we estimated that the average weekly energy usage is approximately 219 kWh.


For generating realistic PV generation profiles, we used data acquired in 2007 by the University of Nevada, Las Vegas [31]; this dataset has a highest resolution of 1 min. For generating uncontrolled PHEV charging profiles, we used an hourly load curve for charging obtained from the National Renewable Energy Laboratory (NREL) [32]. Similar to the loading profile in Fig. 9(a), Fig. 9(b) plots the normalized active power injections for the PV systems and the uncontrolled load curve for the PHEVs. In order to properly scale the PV generation profile, we consider a scenario where there is a 50% penetration of PV installations whose capacities are half that of the *Gable Home's* [6], i.e., 4.5 kW installed capacity with a 5 kVA inverter and ± 2.18 kVAR of reactive power provision capacity provided by the inverter, which is slightly overrated. Similarly, in order to



(a) PV active power injections.



(b) Uncontrolled voltage response.



(c) Controlled voltage response.

Figure 11: 123-bus system: voltage response with and without control for a one-hour period with highly-variable PV.

properly scale the PHEV charging profile, we assume there is a 40% penetration of PHEVs whose charging peak is at 1.3 kW with inverters sized at 1.72 kVA and ± 1.13 kVAR [32].

2) *Base case (no voltage control)*: We fix the feeder voltage to its upper limit of $1.05 \angle 0$ p.u. and, for a week-long period, simulate the evolution of the voltage profile in the network. The results for buses 2, 55, 71, 81, 89, 106, 112, 115, and 123 are plotted in Fig. 10(a), where $\max\{V_2, V_{89}\}$ and $\min\{V_{55}, V_{71}, V_{81}, V_{106}, V_{112}, V_{115}, V_{123}\}$ represent the upper and lower bounds for all bus voltage magnitudes at any given time, respectively. In this figure, we can see that the maximum voltage (outside operational requirements) among all the buses in the network, which takes the value of 1.0568 p.u., is reached by bus 89, whereas the minimum voltage (outside operational requirements) among all the buses in the network, which takes the value of 0.9132 p.u., is reached by bus 123. Note that the voltage on bus 89 exceeds that of the feeder; this results from the fact that many of the customers in this network are assumed to have PV installations capable of creating net active power injections similar to those of the *Gable Home* and the *Equinox House*. As a result, there is a power flow reversal at midday when the load is at its lowest and the PV generation is at its

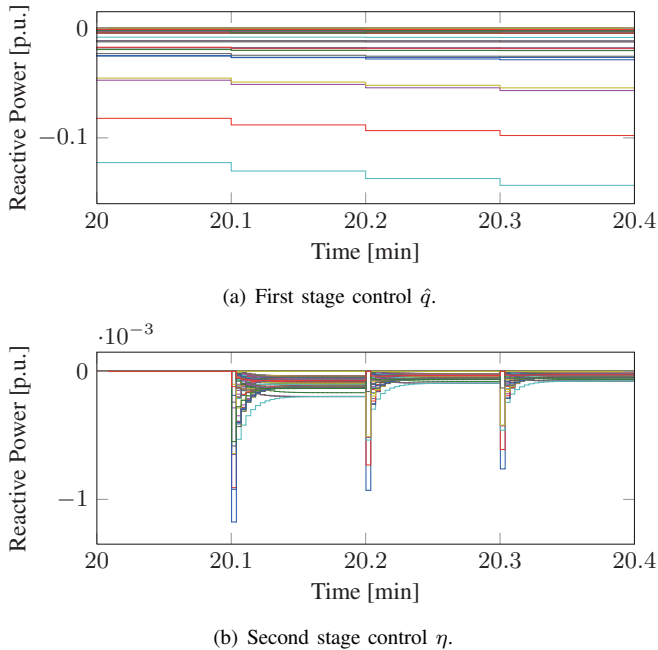


Figure 12: 123-bus system: first and second stage response.

highest. Hence, the voltage drop across a few lines are towards the feeder, rather than away from it. Under-voltages are due to the significant loading introduced by the PHEVs.

3) *System response with two-stage voltage control:* In accordance with the base case, and considering an accuracy level of $\varepsilon_1 = 10^{-3}$, Fig. 10(b) displays the evolution of the voltage in buses 2, 55, 71, 81, 89, 106, 112, 115, and 123. Among all the buses in the network, bus 89 reaches the highest voltage of 1.0498 p.u. and bus 123 observes the lowest voltage of 0.9504 p.u., which are both within operational requirements. Thus, the two-stage voltage control is able to maintain all bus voltages within their performance requirements of $\pm 5\%$ of 1 p.u.

C. One-Hour High-Variability Scenario

Next, we assess the effectiveness of the control system to handle high-variability scenarios in PV generation. In order to do so, we consider the highly-variable PV profile of Fig. 11(a) extracted from the dataset in [31]; which corresponds to a one-hour time frame between 11:30am to 12:30pm.

As before, the feeder voltage is fixed at $1.05 \angle 0$ p.u. Since we are operating at midday, the loads are held at 82% of their maximum values and loading from PHEVs is neglected. The uncontrolled and controlled voltage responses are displayed in Figs. 11(b) and 11(c), respectively. Without control, 33 buses experience an over-voltage and 57 buses will trigger the control actions for the first stage. With voltage control, several buses experience an over-voltage at 11 min into the simulation with the highest voltage reading on bus 89 of 1.0546 p.u. This is promptly corrected with most of the buses within tolerance after 12 s (two control cycles) and bus 89 requiring 36 s. There is a second instance at 20 min where bus 89 experiences an over-voltage of 1.0514 p.u., which is corrected in 12 s. Finally, Fig. 12 shows the responses for both stages.

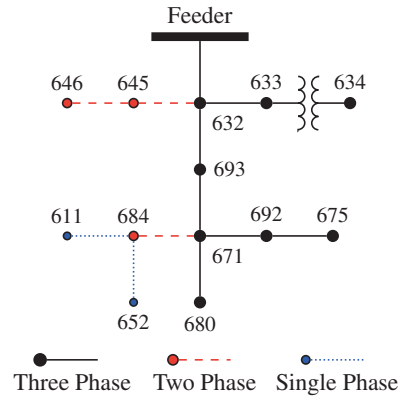


Figure 13: 13-bus system: electrical network graph.

VI. EXTENSION TO UNBALANCED THREE-PHASE SYSTEMS

The ideas presented in Section III and demonstrated in Sections IV–V can be extended to three-phase unbalanced systems by implementing independent two-stage voltage control architectures on each phase. We assume that there are single-phase DERs on each phase capable of providing reactive power support; this assumption is realistic as there are commercially available residential-PV solutions that are also capable of providing reactive power [14].

To illustrate the idea, we use the IEEE 13-bus test system, displayed in Fig. 13, which is unbalanced and contains sublaterals that are single, two, and three phase. The line parameters can be found in [27] and the loads are listed in Appendix E. Similar to the case studies in Section V, we modified the system by removing the SVR located between the feeder and bus 632. We added bus 693 to represent the distributed load defined along the line, and we increased the loads on phases A and B to ensure that under-voltages would appear on each phase. Additionally, we turned the switch between buses 671 and 692 into a transmission line that shares the same length, conductor type, and configuration as the line between buses 692 and 675. We assume that the DERs on each bus have aggregated reactive power capacities between 0.2 and 0.3 p.u.

The communication network mirrors the one-line diagram, with the exception that phases A and B have 10 nodes and phase C has 12 nodes that participate due to the single and two phase interconnections with $k_0 = 133$ and $\tau_0 = 10$ ms. The feeder is modeled as an infinite bus with a fixed voltage magnitude of 0.975 p.u. and is assumed to be balanced. This causes under-voltages on 7 buses in phase A, 1 bus in phase B, and 6 buses in phase C. Figure 14 shows the voltage response and reactive power support per bus on each phase. Figure 14(e) reveals the coupling between phases: phase B voltages drop as a result of the reactive power injected on phases A and C.

Although we have not included the evolution of the second stage, its impact is evident whenever any particular bus reaches its capacity. Observe buses 646 and 634 on phase B in Figs. 14(b) and 14(e). Both buses begin to inject reactive power to raise their voltages, but they quickly reach their capacities. Once bus 646 exhausts its resources, the buses with available capacity begin compensating, as a result of the second stage.

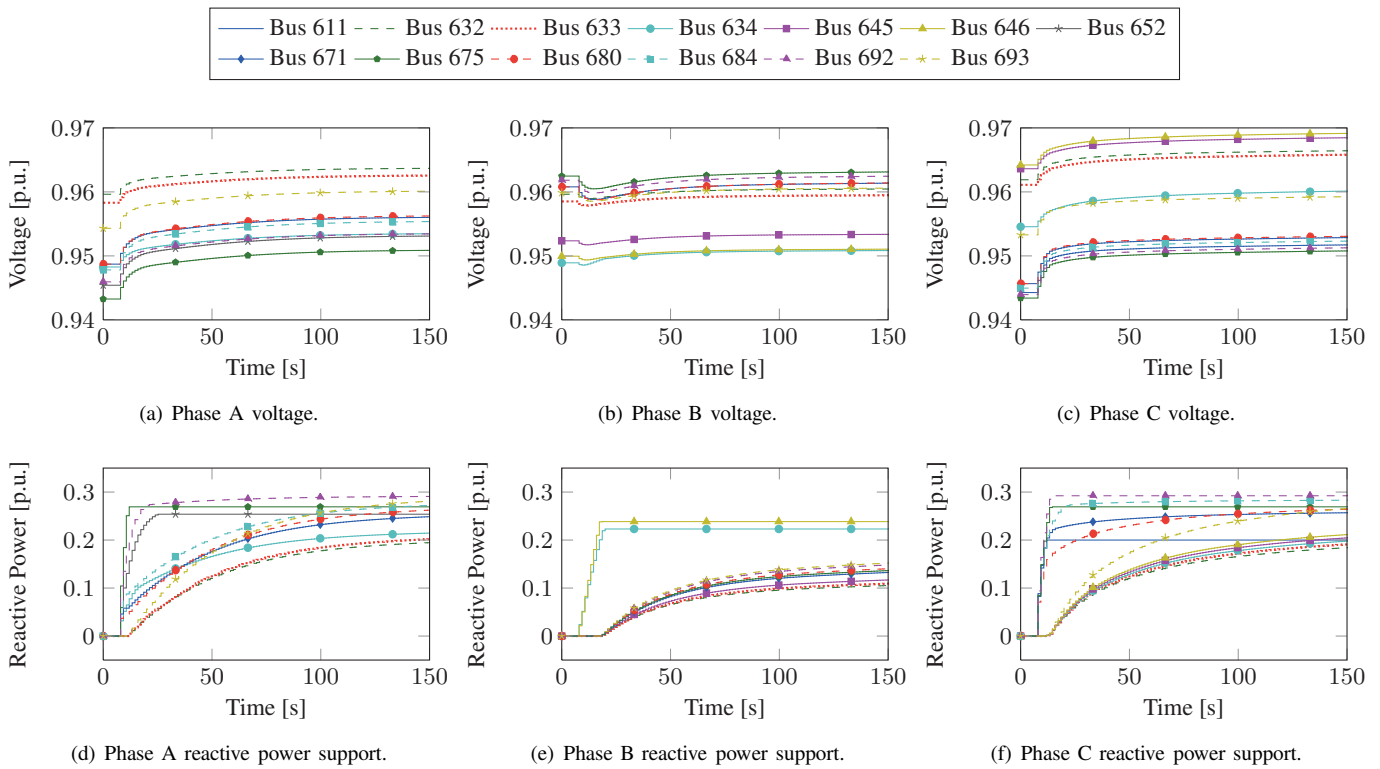


Figure 14: Three-phase implementation of the two-stage architecture.

The interaction between the two stages is noticeable in phases A and C when buses near the violations reach their capacities as well. Although each phase was controlled independently from the other two, the controllers were able to correct all of the bus voltages within an accuracy of $\varepsilon_1 = 10^{-3}$ in 25.27 s, 7.98 s, and 10.64 s for phases A, phase B, and phase C, respectively.

VII. CONCLUDING REMARKS

In this paper, we proposed a two-stage control architecture for voltage control in distribution networks with substantial penetration of variable generation. In order to illustrate the operation of the proposed architecture, we developed two case studies involving an 8-bus system and the IEEE 123-bus test system; both case studies assumed balanced operation. Then, we extended these ideas to unbalanced three-phase systems and provided results for a case study involving the IEEE 13-bus test system.

As already mentioned in the introduction, the main difference between the proposed control scheme and existing decentralized or distributed strategies that address the voltage regulation problem in distribution networks (e.g., [18], [19]) is that the authors of these works propose a switching control scheme where the DERs are operated with a constant power factor while bus voltages are within specifications. Otherwise, whenever there is a voltage violation, their reactive power is adjusted so that the system returns to the desired operating conditions. This is similar to the first stage in our architecture, but lacks the additional support provided by the second stage.

Future work will carefully study the coupling between the proposed control mechanism and TCULs. We also intend to design an optimal high level system control that maintains the benefits of a low level distributed strategy for the integration of a large number of devices with minimal communication requirements.

APPENDIX

A. Notation

S	Sensitivity matrix $\partial V/\partial Q$
ρ_j	Reactive power increment to be injected in node j as estimated by the first stage
\mathcal{N}_j^-	Set of nodes that can send information to node j
D_j^+	Number of nodes that node j can send information to
k_0, ε_0	Number of steps k required for the distributed algorithm to converge within some tolerance ε_0
r_0, ε_1	Number of steps r needed to correct the limit violation within some accuracy level ε_1
τ_0	Time required to perform a single iteration step
τ_1	Time required by both stages to correct a violation
\hat{q}_j	Node j 's first stage reactive power support estimate
q_j	Node j 's actual reactive power support
$\bar{q}_j, \underline{q}_j$	Node j 's upper/lower reactive power constraints
α	Gain used by the local controller in the first stage
μ_j	Node j 's reactive power request to other nodes
$\bar{v}_j, \underline{v}_j$	Node j 's reported available capacities for second stage
η_j	The solution for node j from the distributed algorithm

B. Stage One Stability Analysis

Consider the system in (7); in order to ensure its stability, we need to ensure that $|\sigma_i(I - \alpha SD)| < 1$, $\forall i$, where $\sigma_i(I - \alpha SD)$ denotes the i^{th} eigenvalue of $I - \alpha SD$. Let $\lambda_i = \text{Re}\{\lambda_i\} + j \text{Im}\{\lambda_i\}$ denote the i^{th} eigenvalue of the matrix SD . Then, it follows that every $\sigma_i(I - \alpha SD)$ is given by $\sigma_i = 1 - \alpha \lambda_i$. Therefore, $\sigma_i = (1 - \alpha \text{Re}\{\lambda_i\}) + j \alpha \text{Im}\{\lambda_i\}$, and $|\sigma_i|^2 = 1 - 2\alpha \text{Re}\{\lambda_i\} + \alpha^2 |\lambda_i|^2$. Then, $|\sigma_i|^2 < 1$, if $\alpha < 2 \text{Re}\{\lambda_i\} / |\lambda_i|^2$. Thus, $|\sigma_i(I - \alpha SD)| < 1$, $\forall i$, if $\alpha < \alpha_c = \min_i \left\{ \frac{2 \text{Re}\{\lambda_i\}}{|\lambda_i|^2} \right\}$.

C. Stage Two Convergence Analysis

Let $\mu[k] = [\mu_1[k], \mu_2[k], \dots, \mu_n[k]]^T$, $\bar{v}[k] = [\bar{v}_1[k], \bar{v}_2[k], \dots, \bar{v}_n[k]]^T$, and $\underline{v}[k] = [\underline{v}_1[k], \underline{v}_2[k], \dots, \underline{v}_n[k]]^T$. Then the iterations in (10)–(12) can be rewritten in matrix form as

$$\mu[k+1] = P\mu[k], \quad (20)$$

$$\bar{v}[k+1] = P\bar{v}[k], \quad (21)$$

$$\underline{v}[k+1] = P\underline{v}[k], \quad (22)$$

where $P \in \mathbb{R}^{n \times n}$ is a primitive column stochastic matrix. Primitivity follows since P is (i) a nonnegative matrix, (ii) the assumption that the graph of P is strongly connected implies that P is irreducible, and (iii) P is aperiodic since it contains at least one $p_{jj} > 0$ [24]. This ensures that (20)–(22) converge to the unique solutions $\mu = \lim_{k \rightarrow \infty} \mu[k] = (\sum_{i=1}^n \mu_i[0])\pi$, $\bar{v} = \lim_{k \rightarrow \infty} \bar{v}[k] = (\sum_{i=1}^n \bar{v}_i[0])\pi$, and $\underline{v} = \lim_{k \rightarrow \infty} \underline{v}[k] = (\sum_{i=1}^n \underline{v}_i[0])\pi$, where $\pi = [\pi_1, \pi_2, \dots, \pi_n]^T$ is the unique solution of $\pi = P\pi$ satisfying $\sum_{i=1}^n \pi_i = 1$, and $\pi_i > 0$, $\forall i$, (see, e.g., [24]). If $\sum_{i=1}^n \mu_i[0] \geq 0$, then $\lim_{k \rightarrow \infty} \mu_j[k] \geq 0 \forall j$, and therefore

$$\lim_{k \rightarrow \infty} \frac{\mu_j[k]}{\bar{v}_k[k]} = \frac{\sum_{i=1}^n \mu_i[0]}{\sum_{i=1}^n \bar{v}_i[0]}, \quad (23)$$

from where it follows that

$$\eta_j = \lim_{k \rightarrow \infty} \eta_j[k] = \frac{\sum_{i=1}^n \mu_i[0]}{\sum_{i=1}^n \bar{v}_i[0]} \bar{v}_j[0], \quad \forall j. \quad (24)$$

A similar reasoning can be used to establish that, whenever $\sum_{i=1}^n \mu_i[0] \leq 0$, $\eta_j = (\sum_{i=1}^n \mu_i[0] / \sum_{i=1}^n \underline{v}_i[0]) \underline{v}_j[0]$, $\forall j$.

D. Data for 8-bus System Example

For the topology shown in Fig. 3(a), the transition matrix P that results from the set of weights that define the distributed algorithm is given

$$P = \begin{bmatrix} 1/2 & 1/3 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1/2 & 1/3 & 1/4 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1/3 & 1/4 & 1/3 & 0 & 1/3 & 0 & 0 \\ 0 & 0 & 1/4 & 1/3 & 1/2 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1/3 & 1/2 & 0 & 0 & 0 \\ 0 & 0 & 1/4 & 0 & 0 & 1/3 & 1/2 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1/3 & 1/2 & 0 \end{bmatrix}, \quad (25)$$

whereas for the topology shown in Fig. 3(b) is given by

$$P = \begin{bmatrix} 1/3 & 0 & 0 & 0 & 0 & 0 & 1/4 \\ 1/3 & 1/2 & 1/4 & 0 & 0 & 0 & 0 \\ 0 & 1/2 & 1/4 & 1/3 & 0 & 1/3 & 0 \\ 0 & 0 & 1/4 & 1/3 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1/3 & 1/2 & 0 & 1/4 \\ 0 & 0 & 1/4 & 0 & 0 & 1/3 & 1/4 \\ 1/3 & 0 & 0 & 0 & 1/2 & 1/3 & 1/4 \end{bmatrix}. \quad (26)$$

The load and DER capacity data for the example in Section IV are listed in Table II. Note that the capacities are denoted by *under* and *over* for the under-voltage and over-voltage examples, respectively. Line parameters are listed in Table III.

Table II: 8-BUS LOADING DATA.

Bus	P_L	Q_L	\bar{q}^{over}	\underline{q}^{over}	\bar{q}^{under}	\underline{q}^{under}
1	0.70	0.10	1.000	-1.000	0.20	-0.20
2	0.85	0.25	0.700	-0.700	0.20	-0.20
3	0.60	0.15	0.625	-0.625	0.30	-0.20
4	1.25	0.50	0.500	-0.500	0.50	-0.50
5	0.90	0.30	0.425	-0.425	0.35	-0.35
6	0.10	0.10	0.650	-0.650	0.40	-0.40
7	1.00	0.35	0.625	-0.625	0.20	-0.20

Table III: 8-BUS LINE DATA.

From Bus	To Bus	R [p.u.]	X [p.u.]	B [p.u.]
1	2	0.0010	0.0077	0.0158
2	3	0.0029	0.0145	0.0275
3	4	0.0015	0.0083	0.0142
4	5	0.0035	0.0153	0.0322
3	6	0.0015	0.0065	0.0134
6	7	0.0011	0.0091	0.0188

E. 13-bus Three Phase Distribution Network

The load and DER capacity data for the example in Section VI are listed in Table IV (note that $\underline{q} = -\bar{q}$); line parameter data can be found in [27].

Table IV: 13-BUS LOADING DATA.

Bus	S_A [p.u.]	S_B [p.u.]	S_C [p.u.]	\bar{q}_A	\bar{q}_B	\bar{q}_C
611	-	-	$1.70 - j0.2$	-	-	0.20
632	0	0	0	0.21	0.21	0.21
633	0	0	0	0.22	0.22	0.22
634	$1.92 + j1.32$	$1.68 + j1.35$	$1.20 + j0.90$	0.22	0.22	0.22
645	-	$2.38 + j1.88$	0	-	0.23	0.23
646	-	$3.22 + j1.98$	0	-	0.24	0.24
652	$1.54 + j1.03$	-	-	0.25	-	-
671	$4.62 + j2.64$	$5.39 + j3.30$	$3.85 + j2.20$	0.26	0.26	0.26
675	$5.82 + j0.28$	$0.95 - j1.10$	$2.90 + j0.12$	0.27	0.27	0.27
680	0	0	0	0.28	0.28	0.28
684	0	-	0	0.28	-	0.28
692	0	0	$1.70 + j1.51$	0.29	0.29	0.29
693	$0.17 + j0.10$	$0.66 + j0.38$	$1.17 + j0.68$	0.30	0.30	0.30

REFERENCES

- [1] "IEEE application guide for IEEE Std 1547, IEEE standard for interconnecting distributed resources with electric power systems," *IEEE Std 1547.2-2008*, pp. 1–207, April 2009.
- [2] K. Turitsyn, P. Sulc, S. Backhaus, and M. Chertkov, "Distributed control of reactive power flow in a radial distribution circuit with high photovoltaic penetration," in *Proc. of IEEE Power and Energy Society General Meeting, 2010*, Minneapolis, MN, July 2010, pp. 1–6.
- [3] P. Carvalho, P. Correia, and L. Ferreira, "Distributed reactive power generation control for voltage rise mitigation in distribution networks," *IEEE Trans. Power Sys.*, vol. 23, no. 2, pp. 766–772, May 2008.
- [4] W. H. Kersting, *Distribution System Modeling and Analysis*. New York, NY: CRC Press, 2001.
- [5] A. Keane, L. Ochoa, E. Vittal, C. Dent, and G. Harrison, "Enhanced utilization of voltage control resources with distributed generation," *IEEE Trans. Power Sys.*, vol. 26, no. 1, pp. 252–260, Feb. 2011.
- [6] S. V. Dhople, J. L. Ehlmann, C. J. Murray, S. T. Cady, and P. Chapman, "Engineering systems in the Gable home: A passive, net-zero, solar-powered house for the U. S. Department of Energy's 2009 Solar Decathlon," in *Proc. of Power and Energy Conference at Illinois, 2010*, Feb. 2010, pp. 58–62.
- [7] Newell Instruments. (2012, January) Equinox house. [Online]. Available: <http://newellinstruments.com/equinox>
- [8] C. Guille and G. Gross, "A conceptual framework for the vehicle-to-grid (V2G) implementation," *Energy Policy*, vol. 37, no. 11, pp. 4379–4390, Nov. 2009.
- [9] D. Schooley, Written communication from ComEd, March 2011.
- [10] P. W. Sauer and M. A. Pai, "A comparison of discrete vs continuous dynamic models of tap-changing-under-load transformers," in *Proc. of Bulk Power System Voltage Phenomena III: Voltage Stability, Security, and Control*, Davos, Switzerland, Aug. 1994.
- [11] G. Joos, B. T. Ooi, D. McGillis, F. D. Galiana, and R. Marceau, "The potential of distributed generation to provide ancillary services," in *Proc. of 2000 IEEE Power Engineering Society Summer Meeting*, vol. 3, July 2000, pp. 1762–1767.
- [12] A. D. Domínguez-García, C. N. Hadjicostis, P. T. Krein, and S. T. Cady, "Small inverter-interfaced distributed energy resources for reactive power support," in *Proc. of the IEEE Applied Power Electronics Conference and Exposition*, Fort Worth, TX, Feb. 2011, pp. 1616–1621.
- [13] Petra Solar. (2009) SunWave Communications System. South Plainfield, NJ. [Online]. Available: <http://www.petrasolar.com/>
- [14] SolarBridge Technologies. (2009) Pantheon Microinverter. South Plainfield, NJ. [Online]. Available: <http://www.petrasolar.com/>
- [15] M. E. Baran and I. M. El-Markabi, "A multiagent-based dispatching scheme for distributed generators for voltage support on distribution feeders," *IEEE Trans. Power Sys.*, vol. 22, no. 1, pp. 52–59, Feb. 2007.
- [16] K. Rogers, R. Klump, H. Khurana, A. Aquino-Lugo, and T. Overbye, "An authenticated control framework for distributed voltage support on the smart grid," *IEEE Trans. Smart Grid*, vol. 1, no. 1, pp. 40–47, June 2010.
- [17] D. Villacci, G. Bontempi, and A. Vaccaro, "An adaptive local learning-based methodology for voltage regulation in distribution networks with dispersed generation," *IEEE Trans. Power Sys.*, vol. 21, no. 3, pp. 1131–1140, Aug. 2006.
- [18] P. Vovos, A. Kiprakis, A. Wallace, and G. Harrison, "Centralized and distributed voltage control: Impact on distributed generation penetration," *IEEE Trans. Power Sys.*, vol. 22, no. 1, pp. 476–483, Feb. 2007.
- [19] T. Sansawatt, J. O'Donnell, L. F. Ochoa, and G. P. Harrison, "Decentralised voltage control for active distribution networks," in *Proceedings of the 44th International Universities Power Engineering Conference (UPEC)*, Sept. 2009, pp. 1–5.
- [20] A. D. Domínguez-García and C. N. Hadjicostis, "Coordination and control of distributed energy resources for provision of ancillary services," in *Proc. of IEEE International Conference on Smart Grid Communications*, Gaithersburg, MD, Oct. 2010, pp. 537–542.
- [21] B. A. Robbins, A. D. Domínguez-García, and C. N. Hadjicostis, "Control of distributed energy resources for reactive power support," in *Proc. of North American Power Symposium*, Boston, MA, August 2011.
- [22] A. Bergen and V. Vittal, *Power System Analysis*. Upper Saddle River, NJ: Prentice Hall, 2000.
- [23] E. D. Sontag, "On the input-to-state stability property," *Systems & Control Letters*, vol. 24, pp. 351–359, 1995.
- [24] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 23rd ed. New York, NY: Cambridge University Press, 2010.
- [25] C. Godsil and G. Royle, *Algebraic Graph Theory*. New York, NY: Springer, 2001.
- [26] S. T. Cady, A. D. Domínguez-García, and C. N. Hadjicostis, "Robust implementation of distributed algorithms for control of distributed energy resources," in *Proc. of North American Power Symposium*, Boston, MA, Aug. 2011.
- [27] IEEE PES. (2010, Sept.) Distribution test feeders. [Online]. Available: <http://www.ewh.ieee.org/soc/pes/dsacom/testfeeders/index.html>
- [28] B. A. Robbins and A. D. Domínguez-García, "Distributed algorithms for voltage control in electrical networks," Power Systems Engineering Research Center, PSERC Publication 11-07, September 2011.
- [29] PJM. (2011, March) Historical metered load data. [Online]. Available: <http://www.pjm.com/markets-and-operations/>
- [30] Energy Information Agency. (2012, January). [Online]. Available: <http://www.eia.gov/electricity/data.cfm>
- [31] University of Nevada. (2011, March) Measurement and instrumentation data center. [Online]. Available: <http://www.nrel.gov/midc/unlv/>
- [32] "Costs and emissions associated with plug-in hybrid electric vehicle charging in the Xcel Energy Colorado service territory," National Renewable Energy Laboratory, Tech. Rep. TP-640-41410, Jan. 2007.