ASYNCHRONOUS LOCAL VOLTAGE CONTROL IN POWER DISTRIBUTION NETWORKS

Hao Zhu[†] *and Na Li* [‡]

[†] Department of ECE, University of Illinois, Urbana, IL, 61801

[‡] School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138

ABSTRACT

High penetration of distributed energy resources presents significant challenges and provides emerging opportunities for voltage regulation in power distribution systems. Advanced power-electronics technology makes it possible to control the reactive power output from these resources, in order to maintain a desirable voltage profile. This paper develops a local control framework to account for limits on reactive power resources using the gradient projection optimization method. Requiring only local voltage measurements, the proposed design does not suffer from the stability issues of (de-)centralized approaches caused by communication delays and noises. Our local voltage design is shown to be robust to potential asynchronous control updates among distributed resources in a "plug-and-play" distribution network.

1. INTRODUCTION

Recent advances in power distribution networks have led to a growing interest in voltage regulation using distributed energy resources (DERs) [12]. Highly variable renewable generation and abrupt electric vehicle charging could result in unexpected network-wide voltage fluctuations, at time-scales much faster than the traditional control devices for voltage regulation; see e.g., [6, 9, 12]. At the same time, the power-electronics interface of DERs allows for managing its reactive power output, serving as the emerging resources for fast responding to voltage fluctuations away from the rated values.

With the full network information available centrally, the voltage control problem can be cast as an optimal power flow (OPF) one that minimizes the network voltage mismatch error [5]. Several distributed optimization algorithms have also been proposed to solve this centralized problem, relying on information exchanges among neighboring buses [3,9,10,17]. More recently, a stochastic-approximation approach has been adopted in [7] to handle high system variability and measurement noises. Nonetheless, all these (de-)centralized approaches would require high-quality communication of the measurement and control signals, which is not yet available for every distribution system. Since these optimization-based control methods are designed in an open-loop fashion, potential communication delays or noises would challenge their optimality and even stability for real-time implementations.

To tackle this, voltage control strategies have been designed by using only locally available information such as voltage magnitude measurements [8, 12, 15, 16]. With no centralized coordination, the local control approach could be challenged by the system stability concern when implemented in real time. One popular design is the droop control, in which the reactive power output from DERs scales linearly with the instantaneous local voltage mismatch. As shown in [4], the droop scaling factor has to be small enough to ensure system stability, which could result in insufficient utilization of reactive power resources. A general framework of iterative local voltage control strategies has been recently developed in [18], which can generalize the droop design. Motivated by a surrogate voltage mismatch error proposed in [4], the gradient-projection (GP) method [2, Sec. 3.3] has been proposed for the resource-constrained voltage regulation problem, which turns out to decouple into totally local updates. Convergence analysis for the GP method can be readily applied to establish the stability conditions for the proposed control design.

The goal of this paper is to generalize the GP-based voltage control framework to the scenario of asynchronous local control updates. Due to the heterogeneity of various DER resources, it is very likely that individual nodes are not synchronized in executing their local updates in a "plug-and-play" distribution network. Motivated by this, the stability analysis for asynchronous local voltage control would be of interest. Under the framework of classical asynchronous distributed optimization [2, Ch. 6-7], the choice of stepsize has to be more conservative in order to account for potential delay in information exchange among peer processors. Smaller stepsize could result in slower convergence speed for asynchronous iterative algorithms. Interestingly, our GP-based local control method does not suffer from this stepsize issue. This is because the voltage measurement input in our control design always provides the instantaneous feedback of the power network-wide information. Thanks to this feature, the asynchronous local control updates enjoys the same condition to the synchronous case, on the stepsize choice for convergence.

2. MODELING AND PROBLEM STATEMENT

Consider a distribution network with $(N+1)$ *buses* collected in the set $\mathcal{N} := \{0, 1, \ldots, N\}$, and *line segments* represented

by the set $\mathcal{L} := \{(i,j)\} \subset \mathcal{N} \times \mathcal{N}$; see Fig. 1 for a radial network illustration. We will consider tree-topology networks, where the number of lines $|\mathcal{L}| = (N+1) - 1 = N$.

For bus i, let V_i denote its voltage magnitude, while p_i and q_i denote the real and reactive power injection, all in per unit (p.u.). Bus 0 typically corresponds to the substation bus, assumed to be of fixed voltage; i.e., $V_0 = 1$. For each (i,j), let r_{ij} and x_{ij} denote its line resistance and reactance, and P_{ij} and Q_{ij} the real and reactive line flow from bus i to j, respectively. In addition, the subset $\mathcal{N}_j \subset \mathcal{N}$ has all bus j's neighboring buses that are further down from bus 0. Assuming negligible line flow losses and $V_j \approx 1$ for every j, one can simplify the model flow using the so-termed LinDistFlow equations [1], given for each $(i,j) \in \mathcal{L}$ as

$$P_{ij} - \textstyle\sum_{k\in\mathcal{N}_j} P_{jk} = -p_j, \tag{1a}$$

$$Q_{ij} - \textstyle\sum_{k\in\mathcal{N}_j} Q_{jk} = -q_j, \tag{1b}$$

$$V_i - V_j = r_{ij}P_{ij} + x_{ij}Q_{ij}. \tag{1c}$$

Clearly, (1a)-(1b) represent the real and reactive power flow balance at bus j, while (1c) relates the line voltage drop to the line flow. The LinDistFlow model linearly relates voltage V_j at every bus to the network power injection.

To better represent the linear relation, we will develop the matrix-vector form of the LinDistFlow model. To this end, let matrix $\mathbf{M}^o$ of size $(N+1) \times N$ denote the graph incidence matrix for $(\mathcal{N}, \mathcal{L})$; see e.g., [13, pg. 6]. Its ℓ-th column would correspond to a line segment $(i,j) \in \mathcal{L}$, with all entries zero except for the i-th and j-th ones being ± 1. Let vector $\mathbf{m}_0^T$ be the first row of $\mathbf{M}^o$, with the rest in the $N \times N$ $\mathbf{M}$; i.e., $\mathbf{M}^o = [\mathbf{m}_0\ \mathbf{M}^T]^T$. Since $(\mathcal{N}, \mathcal{L})$ is a connected tree, the rank of $\mathbf{M}^o$ equal to $(N+1)-1 = N$ [13]. Hence, the square matrix $\mathbf{M}$ is of full rank N and invertible. Upon defining the voltage vector $\mathbf{V} := [V_1, \ldots, V_N]^T$ and similarly other power related vectors, one can stack the LinDistFlow equations in (1a)-(1c) for every $(i,j) \in \mathcal{L}$ into

$$-\mathbf{M}\mathbf{P} = -\mathbf{p}, \tag{2a}$$

$$-\mathbf{M}\mathbf{Q} = -\mathbf{q}, \tag{2b}$$

$$[\mathbf{m}_0\ \mathbf{M}^T][V_0\ \mathbf{V}^T]^T = \mathbf{D}_r\mathbf{P} + \mathbf{D}_x\mathbf{Q} \tag{2c}$$

where $\mathbf{D}_r$ is the $N \times N$ diagonal matrix with the ℓ-th diagonal entry equal to r_{ij}; and similarly the diagonal $\mathbf{D}_x$ captures all reactance x_{ij}'s. Solving for $\mathbf{P}$ and $\mathbf{Q}$ and substituting $V_0 = 1$ into (2c) give rise to

$$\mathbf{V} = \mathbf{R}\mathbf{p} + \mathbf{X}\mathbf{q} - \mathbf{M}^{-T}\mathbf{m}_0 \tag{3}$$

where $\mathbf{R} := \mathbf{M}^{-T}\mathbf{D}_r\mathbf{M}^{-1}$ and $\mathbf{X} := \mathbf{M}^{-T}\mathbf{D}_x\mathbf{M}^{-1}$ are the network resistance and reactive matrices, respectively. As shown by [4, 18], both $\mathbf{R}$ and $\mathbf{X}$ are symmetric positive definite (PD).

The goal of voltage regulation is to attain a desired voltage profile $\mathbf{V} \to \boldsymbol{\mu}$ by optimizing the injected reactive power in

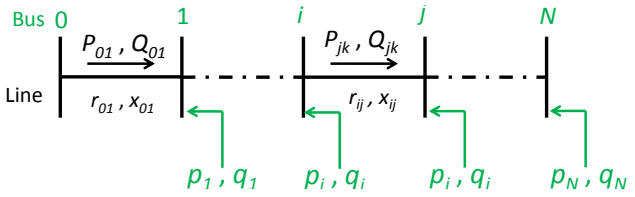

Fig. 1. A radial distribution system with bus and line associated variables.

(3). The flat voltage profile is a popular choice, corresponding to setting $\boldsymbol{\mu} = \mathbf{1}$. Since the injected reactive power $\mathbf{q} = \mathbf{q}^g - \mathbf{q}^c$, with $\mathbf{q}^g$ denoting the reactive power supplied by DERs and $\mathbf{q}^c$ the uncontrollable part corresponding to load consumption, the LinDistFlow model (3) becomes

$$\mathbf{V} = \mathbf{R}\mathbf{p} + \mathbf{X}\mathbf{q}^g - \mathbf{X}\mathbf{q}^c - \mathbf{M}^{-T}\mathbf{m}_0 = \mathbf{X}\mathbf{q}^g + \bar{\mathbf{V}} \tag{4}$$

where $\bar{\mathbf{V}} := \mathbf{R}\mathbf{p} - \mathbf{X}\mathbf{q}^c - \mathbf{M}^{-T}\mathbf{m}_0$ is the voltage profile when the input $\mathbf{q}^g = \mathbf{0}$. Under given $\mathbf{p}$ and $\mathbf{q}^c$, the problem becomes to minimize the voltage mismatch error by solving

$$\mathbf{q}^\dagger = \arg\min_{\mathbf{q}} \ \frac{1}{2}\|\mathbf{V} - \boldsymbol{\mu}\|^2 = \frac{1}{2}\|\mathbf{X}\mathbf{q} + \bar{\mathbf{V}} - \boldsymbol{\mu}\|^2 \tag{5a}$$

$$\text{subject to} \quad \underline{\mathbf{q}} \le \mathbf{q} \le \bar{\mathbf{q}} \tag{5b}$$

where $\|\cdot\|$ is the Euclidean norm operator, and the constraint (5b) accounts for the limits of reactive power resources at every bus. This is a box-constrained quadratic program (QP) and could be easily solved if a centralized controller can instantaneously acquire the full network information and feedback the decisions to remote DERs. Without requiring a centralized controller, this paper will develop a solution approach using only local information by leveraging the gradient-projection (GP) method.

Remark 1. Albeit an approximation, the linear model (4) works well for distribution networks with power losses and non-flat voltage. Otherwise, it can be thought of as the sensitivity analysis under small system changes. Moreover, it has been shown in [18] that (4) would hold for meshed networks that not of tree topology. Interestingly, the inverse of $\mathbf{X}$ given by $\mathbf{B} := \mathbf{X}^{-1} = \mathbf{M}\mathbf{D}_x^{-1}\mathbf{M}^T$ is also a PD matrix, and it is actually the power network *Bbus matrix* used in the fast decoupled power flow (FDPF) analysis; see e.g., [14, Sec. 6.16]. As shown in [18], the local control design based on (4) can be applied to meshed networks and even three-phase distribution networks.

3. GP-BASED LOCAL VOLTAGE CONTROL

To develop local control schemes, consider a surrogate problem to (5) by weighting the voltage mismatch norm using $\mathbf{B} := \mathbf{X}^{-1}$, as given by

$$\mathbf{q}^\star = \arg\min_{\underline{\mathbf{q}} \le \mathbf{q} \le \bar{\mathbf{q}}} \ f(\mathbf{q}) := \frac{1}{2}(\mathbf{X}\mathbf{q} - \tilde{\mathbf{V}})^T\mathbf{B}(\mathbf{X}\mathbf{q} - \tilde{\mathbf{V}}) \tag{6}$$

where the voltage difference $\tilde{\mathbf{V}} := \boldsymbol{\mu} - \bar{\mathbf{V}}$. The surrogate problem is still convex, because its Hessian $\mathbf{X}$ is a PD matrix. Because of the box constraint, the weighted error norm in (6) would attain a different solution compared to the original unweighted problem (5). If every bus has unlimited reactive power resource, the box constraint would become inactive and the problem reduces to an unconstrained one. Accordingly, both error norm objectives would achieve the minimum at zero with equal optimal solutions as $\mathbf{q}^\dagger = \mathbf{q}^\star = \mathbf{X}^{-1}(\boldsymbol{\mu} - \bar{\mathbf{V}})$. This implies that under abundant reactive power resources, the optimal solution to (6) has the potential to closely approximate the minimum of the (5), suggesting promising performance could be achieved by solving the surrogate problem instead. This weighted norm has been proposed in [4] for reverse engineering the classical droop control design.

Thanks to the separable structure of the box constraint, (6) can be easily solved using the gradient-projection (GP) method, a generic constrained optimization solver; see e.g., [2, Sec. 3.3]. Upon forming the gradient direction

$$\nabla f(\mathbf{q}) := \mathbf{X}\mathbf{q} - \tilde{\mathbf{V}}, \tag{7}$$

the GP iteration works by projecting the gradient update, as

$$\mathbf{q}(t+1) = \mathbb{P}[\mathbf{q}(t) - \epsilon \nabla f(\mathbf{q}(t))]\,,\ \forall t \geq 0 \tag{8}$$

where the stepsize $\epsilon > 0$, and the operator $\mathbb{P}$ projects any input to the box $[\underline{\mathbf{q}}, \bar{\mathbf{q}}]$. The projection is extremely efficient through entry-wise thresholding.

Because of the projection operation, $\mathbf{q}(t) \in [\underline{\mathbf{q}}, \bar{\mathbf{q}}]$ always holds for any $t > 0$. It is feasible to use the latest GP iterate as the network reactive power input by setting $\mathbf{q}^g = \mathbf{q}(t)$ at every iteration t. Under The linear relation in (4) suggests that the gradient per iteration t

$$\nabla f(\mathbf{q}(t)) = \mathbf{X}\mathbf{q}(t) - \tilde{\mathbf{V}} = \mathbf{V}(t) - \boldsymbol{\mu}. \tag{9}$$

Clearly, its j-th entry $\nabla_j f(\mathbf{q}(t)) = V_j(t) - \mu_j$ only depends on the local bus voltage, and does not require information on the full vector $\mathbf{q}(t)$. Hence, the GP iteration (8) can be completely decoupled into local updates, as given by

$$q_j(t+1) = \mathbb{P}_j\left[q_j(t) - \epsilon(V_j(t) - \mu_j)\right]\ \forall j,\ t \tag{10}$$

where $\mathbb{P}_j$ is the local projection operator at bus j to the interval $[\underline{q}_j, \bar{q}_j]$. The proposed local control updates using (10) are equivalent to the centralized counterpart (8).

Proposition 1. *Under constant $\bar{\mathbf{V}}$, the local update* (10) *converges to the optimum $\mathbf{q}^\star$ if $0 < \epsilon < 2/K$, where K is the spectral norm of matrix $\mathbf{X}$.*

This proposition follows directly from the convergence analysis for general GP algorithms in [2, Sec. 3.3], which relies on the Lipschitz continuity condition $\|\nabla f(\mathbf{q}) - \nabla f(\mathbf{q}')\| \leq K\|\mathbf{q} - \mathbf{q}'\|$, for any $\mathbf{q}$ and $\mathbf{q}'$. The bound $\epsilon < 2/K$ would ensure sufficient descent in the objective at every iteration. To meet the Lipschitz continuity condition, K should be no greater than the spectral norm of the Hessian $\mathbf{X}$, which equals to its largest eigenvalue since $\mathbf{X}$ is symmetric and PD.

Remark 2. As a special case of the first-order algorithms, the GP method typically attains a linear rate of convergence [2, Sec. 3.3]. The convergence speed of GP algorithms would be determined by the condition number of the Hessian matrix. To accelerate the convergence, one can scale the stepsize to be d_j per bus j using a diagonal $\mathbf{D} := \mathrm{diag}(d_1, \ldots, d_N)$. As shown in [18], this approach effectively scales the Hessian to be $[\mathbf{D}^{1/2}\mathbf{X}\mathbf{D}^{1/2}]$. To improve the condition number, we can select $\mathbf{D} = [\mathrm{diag}(\mathbf{X})]^{-1}$ in order to approximate the Newton's method. This diagonal scaling does not affect the separability of the projection operator. Hence, the GP updates with diagonally scaled Hessian would still converge, even for the ensuing analysis of asynchronous GP updates.

The GP-based control design requires each bus to measure its local voltage magnitude, which can be implemented with minimal hardware requirements. The computation is also efficient since it only involves scalar operations. In addition, even though the analytical results rely on the linearized power flow model, the proposed local voltage control (10) can be easily implemented in any realistic distribution networks such as three-phase systems. Albeit its simple design and robust features, the local update could suffer from performance degradation due to the surrogate error norm objective in (6), compared to a centralized approach with full network information available. In the future, we would be interested in quantifying this effect for given reactive power limits.

4. ASYNCHRONOUS VOLTAGE CONTROL

Asynchronous voltage control is motivated by the heterogeneity of reactive power resources. Allowing each of them to perform local update at different rates would facilitate the "plug-and-play" functionality for DER integration. Since implementing (10) only requires local voltage information, the nodes that can compute faster do not need to wait for other slower ones. This strategy allows the faster nodes to execute more iterations for a given time, leading to quicker response to local voltage violations.

To this end, let the set $\mathcal{T}_j$ collect all the iterations at which bus j executes its local update. The asynchronous counterpart of (10) is given by

$$q_j(t+1) = q_j(t) + s_j(t),\ \forall j,\ t \tag{11}$$

where the difference

$$s_j(t) := \begin{cases} \mathbb{P}_j\left[q_j(t) - \epsilon(V_j(t) - \mu_j)\right] - q_j(t), & t \in \mathcal{T}_j \\ 0, & t \notin \mathcal{T}_j \end{cases} \tag{12}$$

To establish the convergence condition, we will constrain every bus to update sufficiently often. This is similar to the

bounded update delay assumption in the analysis of partially asynchronous algorithms in [2, Ch. 7]. Specifically, there exists a positive integer T such that for every bus j and every $t \geq 0$, at least one of the elements of the set $\{t, t+1, \ldots, t+T-1\}$ belongs to $\mathcal{T}_j$. Therefore, every bus performs a local update at least once within every T iterations.

The general analysis of [2, Ch. 7] also considers the scenario of *bounded information delay* in computing the local update. Due to potential communication delays among distributed processors, local updates may not be performed based on the latest system state information. The analysis in [2, Ch. 7] assumes that the local information used for computing the gradient direction could be outdated by at most T iterations. Under both types of delay, the sufficient convergence conditions for general asynchronous GP algorithms in [2, Sec. 7.5] state that (11)-(12) will converge if the stepsize $0 < \epsilon < 1/[K(1+T+NT)]$. This upper bound on ϵ depends on the time delay constant T, based on the slowest processor. Hence, the stepsize could be much smaller than $2/K$ in Proposition 1, potentially leading to much slower convergence than the synchronous case.

Interestingly, our asynchronous update (11)-(12) does not suffer from the information delay. Since $\mathbf{V}(t) - \boldsymbol{\mu} = \mathbf{X}\mathbf{q}(t) - \tilde{\mathbf{V}} = \nabla f(\mathbf{q}(t))$ always holds, $s_j(t)$ is computed at every $t \in \mathcal{T}_j$ with the latest $\mathbf{q}(t)$. In other words, the network power flow model guarantees that local voltage $V_j(t)$ contains the up-to-date information on the gradient direction. This is different from the implementation of most parallel and distributed algorithms, where local updates require information sent by peer processors. Thanks to this feature of local voltage control, the asynchronous version enjoys the same convergence condition to the synchronous case, as follows.

Proposition 2. *Under the bounded update delay assumption, the asynchronous update in* (11)-(12) *converges to the optimum* $\mathbf{q}^\star$ *if* $0 < \epsilon < 2/K$.

Proof: By projecting any q_j' to $[\underline{q}_j, \bar{q}_j]$, it holds that $[\mathbb{P}_j(q_j') - q_j][\mathbb{P}_j(q_j') - q_j'] \leq 0$ for any $q_j \in [\underline{q}_j, \bar{q}_j]$ [2, Sec. 3.3]. This implies that at every iteration $t \in \mathcal{T}_j$

$$s_j(t)[s_j(t) + \epsilon \nabla_j f(\mathbf{q}(t))] = \epsilon s_j(t) \nabla_j f(\mathbf{q}(t)) + s_j^2(t) \leq 0. \tag{13}$$

With the Lipschitz continuity of $f(\cdot)$, the descent lemma in [2, Prop. A.32] entails that for every $t \geq 0$

$$\begin{aligned} f(\mathbf{q}(t+1)) &= f(\mathbf{q}(t) + \mathbf{s}(t)) \\ &\leq f(\mathbf{q}(t)) + \mathbf{s}(t)^T \nabla f(\mathbf{q}(t)) + (K/2)\|\mathbf{s}(t)\|^2 \\ &\leq f(\mathbf{q}(t)) - \left(\frac{1}{\epsilon} - \frac{K}{2}\right) \|\mathbf{s}(t)\|^2. \quad [\text{cf.}(13)] \end{aligned}$$

If $C := 1/\epsilon - K/2 > 0$ or equivalently $\epsilon < K/2$, summing up this inequality over all iterations yields

$$\textstyle\sum_{\tau=0}^{\infty} \|\mathbf{s}(\tau)\|^2 \leq \frac{1}{C} f(\mathbf{q}(0)) < \infty,$$

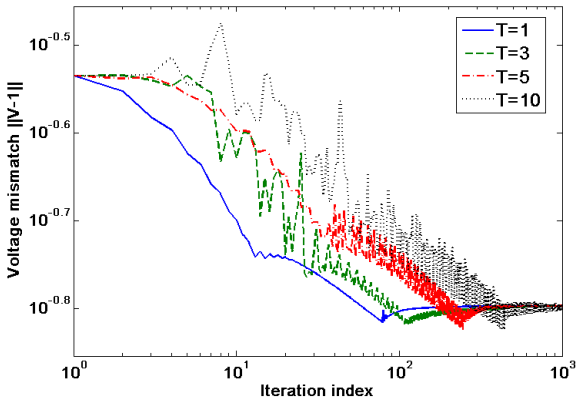

Fig. 2. Iterative voltage mismatch error norm for various choices of delay constant T.

which implies that $\lim_{t\to\infty} s_j(t) = 0$ at every bus j, satisfying the convergence condition for GP algorithms. □

The asynchronous GP algorithm also attains a linear rate of convergence under some proper conditions of $f(\cdot)$ [11]. However, its convergence speed can become slower due to the delay in executing updates. In general, optimal stepsize for asynchronous methods is challenging to characterize.

5. NUMERICAL TESTS

To validate the convergence analysis for the asynchronous control design, we have tested it on a synthetic radial network as in Fig. 1 with $N = 10$. A uniform line reactance model is assumed such that $x_{ij} = 1$ for every $(i,j) \in \mathcal{L}$. The baseline voltage profile $\tilde{\mathbf{V}}$ is randomly generated using a Gaussian distribution $\mathcal{N}(\mathbf{1}, \sigma^2\mathbf{I})$ with $\sigma = 0.1$, and the desired flat voltage profile is set to be $\boldsymbol{\mu} = \mathbf{1}$. The reactive power limits are set to be $[-0.2,\ 0.2]$ at every bus. The stepsize at bus j is set to be ϵd_j, where the scaling matrix $\mathbf{D} = [\text{diag}(\mathbf{X})]^{-1}$. Under this setting, the maximum and minimum eigenvalues of the effective Hessian matrix are $\lambda_{\max} = 7.2572$ and $\lambda_{\min} = 0.0334$, respectively. Hence, the bound on stepsize $\epsilon < 0.2756$, while we set $\epsilon = 2/(\lambda_{\max} + \lambda_{\min}) \approx 0.274$ as the choice for best convergence rate under the synchronous scenario.

For given delay constant T, bus j performs the local update once every T_j iterations, with integer T_j drawn from a uniform distribution in $[1,\ T]$. Fig. 2 plots the iterative voltage mismatch error $\|\mathbf{V} - \mathbf{1}\|$ for various choices of T. Although the error is not monotonically decreasing, it converges at a linear rate under all scenarios. The case of $T = 1$ corresponds to the synchronous scenarios. Hence, the asynchronous control would attain the same optimal solution as the synchronous case. Fig. 2 also corroborates the choice of ϵ in Proposition 2, as compared to the upper bound $1/[K(1+T+NT)] = 0.0025$ for $T = 10$. We are interested to further investigate the choice of stepsize for improving convergence rate of the asynchronous method.

6. REFERENCES

[1] M. E. Baran and F. F. Wu, "Optimal capacitor placement on radial distribution systems," *IEEE Trans. Power Delivery*, vol. 4, no. 1, pp. 725-734, 1989.

[2] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and Distributed Computation: Numerical Methods*, Belmont, M.A. : Athena Scientific, 1997.

[3] E. Dall'Anese, H. Zhu, and G. B. Giannakis, "Distributed optimal power flow for smart microgrids," *IEEE Trans. Smart Grid*, vol. 4, no. 3, pp. 1464-1475, Dec. 2013.

[4] M. Farivar, C. Lijun, S. Low, "Equilibrium and dynamics of local voltage control in distribution systems," in Proc. *52nd IEEE Annual Conf. on Decision and Control*, pp. 4329-4334, Dec. 2013.

[5] M. Farivar, R. Neal, C. Clarke, and S. Low, "Optimal inverter VAR control in distribution systems with high PV penetration," in *Proc. of IEEE Power and Energy Society General Meeting*, San Diego, CA, June 2012.

[6] P. Jahangiri and D. C. Aliprantis, "Distributed Volt/VAr Control by PV Inverters," *IEEE Transactions on Power Systems*, vol. 28, no. 3, pp. 3429-3439, Aug. 2013.

[7] V. Kekatos, G. Wang, A.-J. Conejo, and G. B. Giannakis, "Stochastic Reactive Power Management in Microgrids with Renewables," *IEEE Trans. on Power Systems*, 2015 (early access).

[8] N. Li, G. Qu, and M. Dahleh, "Real-time decentralized voltage control in distribution networks," *52nd Annual Allerton Conf. on Communication, Control, and Computing*, pp. 582-588, 2014.

[9] B. A. Robbins, C. N. Hadjicostis, and A. D. Dominguez-Garcia, "A Two-Stage Distributed Architecture for Voltage Control in Power Distribution Systems," *IEEE Trans. on Power Systems*, vol. 28, no. 2, pp. 1470-1482, May 2013.

[10] P. Sulc, S. Backhaus, and M. Chertkov, "Optimal distributed control of reactive power via the alternating direction method of multipliers," *IEEE Trans. Energy Conversion*, vol. 29, no. 4, pp. 968-977, 2014.

[11] P. Tseng, "On the Rate of Convergence of a Partially Asynchronous Gradient Projection Algorithm," *SIAM Journal of Optimization*, vol. 1, no. 4, pp. 603-619,, 1991.

[12] K. Turitsyn, P. Sulc, S. Backhaus, and M. Chertkov, "Options for control of reactive power by distributed photovoltaic generators," *Proc. of the IEEE*, vol. 99, no. 6, pp. 1063-1073, 2011.

[13] D. B. West, *Introduction to graph theory*, 2nd ed., Upper Saddle River: Prentice hall, 2001.

[14] A. J. Wood, B. F. Wollenberg, and G. B. Sheble, *Power Generation, Operation, and Control*, 3rd ed., New York, NY: John Wiley and Sons, 2013.

[15] H.-G. Yeh, D. F. Gayme, and S. H. Low, "Adaptive VAR control for distribution circuits with photovoltaic generators," *IEEE Trans. on Power Systems*, vol. 27, no. 3, pp. 1656-1663, 2012.

[16] B. Zhang, A. Dominguez-Garcia, and D. Tse, "A Local Control Approach to Voltage Regulation in Distribution Networks," in *Proc. North America Power Symposium*, pp. 1-7, 2013.

[17] B. Zhang, A. Y.S. Lam, A. D. Dominguez-Garcia, and D. Tse, "An Optimal and Distributed Method for Voltage Regulation in Power Distribution Systems," *IEEE Trans. on Power Systems*, vol. 30, no. 4, pp. 1714-1726, July 2015.

[18] H. Zhu and H.-J. Liu, "Fast Local Voltage Control Under Limited Reactive Power: Optimality and Stability Analysis," *IEEE Trans. on Power Systems*, 2015 (accepted).