Robust Online Voltage Control with an Unknown Grid Topology

Christopher Yeh
California Institute of Technology
Pasadena, CA, USA
cyeh@caltech.edu

Jing Yu
California Institute of Technology
Pasadena, CA, USA
jing@caltech.edu

Yuanyuan Shi
UC San Diego
San Diego, CA, USA
yyshi@eng.ucsd.edu

Adam Wierman
California Institute of Technology
Pasadena, CA, USA
adamw@caltech.edu

ABSTRACT

Voltage control generally requires accurate information about the grid's topology in order to guarantee network stability. However, accurate topology identification is a challenging problem for existing methods, especially as the grid is subject to increasingly frequent reconfiguration due to the adoption of renewable energy. Further, running existing control mechanisms with incorrect network information may lead to unstable control. In this work, we combine a nested convex body chasing algorithm with a robust predictive controller to achieve provably finite-time convergence to safe voltage limits in the online setting where the network topology is initially unknown. Specifically, the online controller does not know the true network topology and line parameters, but instead must learn them over time by narrowing down the set of network topologies and line parameters that are consistent with its observations and adjusting reactive power generation accordingly to keep voltages within desired safety limits. We demonstrate the effectiveness of the approach using a case study, which shows that in practical settings the controller is indeed able to narrow the set of consistent topologies quickly enough to make control decisions that ensure stability.

CCS CONCEPTS

• **Hardware** → **Smart grid**; **Power networks**; • **Theory of computation** → **Online learning algorithms**; • **Computer systems organization** → *Sensor networks*; • **Computing methodologies** → *Control methods.*

KEYWORDS

voltage control, distribution grid, convex body chasing

ACM Reference Format:
Christopher Yeh, Jing Yu, Yuanyuan Shi, and Adam Wierman. 2022. Robust Online Voltage Control with an Unknown Grid Topology. In *The Thirteenth ACM International Conference on Future Energy Systems (e-Energy '22), June 28–July 1, 2022, Virtual Event, USA.* ACM, New York, NY, USA, 11 pages. https://doi.org/10.1145/3538637.3538853

e-Energy '22, June 28–July 1, 2022, Virtual Event, USA

ACM ISBN 978-1-4503-9397-3/22/06.
https://doi.org/10.1145/3538637.3538853

1 INTRODUCTION

Operators of electricity distribution grids must maintain voltages at each bus within certain operating limits, as deviations from such limits may damage electrical equipment and cause power outages [1, 17]. This "voltage control" or "voltage regulation" problem has been well-studied in the literature, *e.g.*, [24, 27, 37] and the references therein. Voltage control algorithms aim to guarantee grid stability while minimizing the costs associated with control inputs. Typically, the algorithms assume *exact knowledge* of the underlying grid topology.

However, the exact grid topology and line parameters are often not known, and inexact knowledge may lead to stability problems for voltage control algorithms [20, 25]. This problem is exacerbated by the increasing integration of distributed energy resources (DERs), such as photovoltaic and storage devices. Especially in distribution grids, where DERs often do not belong to the electricity utility, the grid operator may lack up-to-date information about the grid topology [9, 22]. While a grid operator can install sensors to help identify the current network topology, unless such sensors are densely deployed, uncertainty about the topology remains; and so cost is prohibitive. Furthermore, parts of the grid may undergo frequent reconfiguration either due to load balancing [3] or unplanned maintenance. Thus, distribution grid operators cannot expect to operate with perfect topology information and the design of voltage control robust to unknown grid topology is crucial.

To date, only a limited number of voltage control mechanisms have been studied in the case when the grid topology is uncertain. One common design is to learn a voltage controller via deep reinforcement learning (DRL), *e.g.*, [11, 15, 31, 33, 34] and references within. However, such approaches have neither performance nor voltage stability guarantees. Thus, they are not suitable for safety-critical infrastructure. Two recent works [7, 30] propose a model-free DRL approach for voltage control with stability guarantees. The main tool being used in [7, 30] is Lyapunov stability theory, from which a structural constraint for stable controllers is derived, and policy optimization with the constraint is performed. In contrast, our framework jointly learns the system model (consistent to data) and stable controller, in an online fashion.

Another common approach for voltage control when grid topology is uncertain, is to use a two-stage model-based approach: first, estimate the network topology, a.k.a., system identification, using structured neural networks; and second, apply an existing model

predictive control with the identified model, *e.g.*, [6]. However, accurate topology identification is a challenging problem [29] and existing methods require voltage measurements over hundreds of time steps [8, 22], after which uncertainty still remains. This is problematic because, during the time where system identification is being performed, the network is not able to respond effectively to disturbances, since using an incorrect model may lead to unstable control. Given the timescale of grid topology changes in practice, a different approach is needed.

1.1 Contributions

We propose a new approach for voltage control over an uncertain grid topology that does not perform system identification and voltage control separately. Instead, our approach robustly learns to stabilize voltage within the desired limits directly, *without any prior knowledge of the topology and without needing to precisely learn the topology*.

Our approach takes ideas from online nested convex body chasing (CBC) [2] and robust predictive control and combines them using a new learning framework [18] to apply them to voltage control for the first time. Intuitively, we use a nested CBC algorithm in order to track the set of topologies that are consistent with the observed voltage measurements—as more measurements are taken the set of consistent topologies shrinks (and so the sets are nested). As these measurements are taken, a form of robust predictive control is used for voltage control, where the robustness guarantee is used to ensure the uncertainty about the topology can be handled. Our main result (Theorem 1) provides a finite error stability bound for the overall controller, which is summarized in Algorithm 1. This represents the first voltage control algorithm that is provably robust to uncertainty about network topology.

In addition to providing theoretical guarantees, we demonstrate the effectiveness of our proposed approach using a case study of a 56-bus distribution grid from the Southern California Edison (SCE) utility [13]. In this setting, we give the controller no prior information about the topology of the grid, yet the controller quickly narrows down the set of topologies and line parameters that are consistent with its observations and adjusts reactive power generation to keep voltages within desired safety limits when faced with disturbance. In fact, our controller's performance nearly matches that of controllers which assume perfect knowledge of the topology, even when given only partial observations of bus voltages.

1.2 Related Work

The problem of voltage control has a long history with many important contributions [4, 6, 14, 21, 26, 28, 32, 33, 36, 37] (and the references within). Classic voltage regulation devices such as tap-changing transformers [14, 28] are effective in dealing with *slow* voltage variations. However, with fast time variations introduced by renewables, a growing body of literature has focused on inverter-based controllers that can quickly respond by adjusting their active and reactive power set-points [4, 21, 26, 32, 37]. Most of these works cast voltage control as an optimization problem and then propose different centralized or decentralized algorithms depending on the communication infrastructure. Critically, all of these voltage control methods assume that the underlying grid topology is *known*.

Some recent works consider voltage control with unknown network topology and parameters. These works either use a two-stage model-based approach of first performing system identification and then optimizing over the identified model [6], or an entirely model-free to learn a controller via deep reinforcement learning [11, 15, 31, 33, 34] which has no performance or voltage stability guarantees. In contrast, our work considers model-based approach that jointly learns the system model and a controller.

An important tool for voltage control is model predictive control (MPC), which has been investigated in a number of works, *e.g.*, [6, 16, 23]. Of particular relevance to this paper is work on robust MPC algorithms for voltage control such as [23]. While many proposals have emerged, to this point provable bounds for robust MPC algorithms have typically been elusive. A key part of our proposed framework is the development of a robust control algorithm for voltage control with a provable robustness guarantee. This is summarized in Theorem 3.

The standard approach for handling uncertainty about network topology is to estimate the topology using a form of system identification. There is a growing literature of such approaches, e.g., [8, 9, 19, 20, 22, 25]. A prominent approach is to use graphical models for topology reconstruction [8], via maximum likelihood methods while enforcing other structural restrictions like low-rank and sparsity. Our approach in this paper is novel because system identification is not performed separately from control. Instead of seeking to estimate the topology, the controller itself is learned directly without the intermediate step of system identification.

2 MODEL

We study voltage control on an unknown grid topology. We consider a radial power distribution circuit represented as a network $G = (\mathcal{N}, \mathcal{E})$, where $\mathcal{N}$ is the set of buses (nodes) and $\mathcal{E} \subset \mathcal{N} \times \mathcal{N}$ is the set of lines (edges). The buses are numbered $\mathcal{N} = \{0, 1, 2, \ldots, n\}$, where bus 0 is the substation and other buses are branch buses. Let $\mathcal{C} \subseteq \mathcal{N}$ denote the subset of buses with controllable reactive power injection. Because the network is radial (tree-structured) and rooted at the substation (bus 0), there is a unique path $\mathcal{P}_i$ from bus 0 to any other bus i.

For each line $(i, j) \in \mathcal{E}$, its complex impedance is $r_{ij} + \mathbf{i}x_{ij}$, with real-valued line parameters $r_{ij}, x_{ij} > 0$ (units Ω). Define the following matrices $R^\star, X^\star \in \mathbb{R}^{n \times n}$, which are computed from the network topology and line parameters. Assuming that G is a connected graph (*i.e.*, no bus is disconnected from the substation), then $R^\star, X^\star$ are positive definite and have strictly positive entries [12].

$$R^\star_{ij} := 2 \sum_{(h,k) \in \mathcal{P}_i \cap \mathcal{P}_j} r_{hk}, \quad X^\star_{ij} := 2 \sum_{(h,k) \in \mathcal{P}_i \cap \mathcal{P}_j} x_{hk}, \quad i, j \in \{1, \ldots, n\} \tag{1}$$

Let $v \in \mathbb{R}^n$ denote the squared voltage magnitude at the buses, excluding the substation. Let $p + \mathbf{i}q$ denote the complex power injection at the buses, where $p \in \mathbb{R}^n$ (units W) is the net active power injection, and $q \in \mathbb{R}^n$ (units Var) is the net reactive power injection. We assume that the active power injection is exogenous, but that reactive power at each bus can be decomposed as $q = q^c + q^e$, where q^c is the "controllable" component and q^e is the "exogenous" (*i.e.*, uncontrollable) component.

Under the linearized Simplified DistFlow model [21],

$$v = R^\star p + X^\star q + v^0 \mathbf{1}_n = X^\star q^c + v^{\text{par}}, \tag{2}$$

where $v^{\text{par}} = R^\star p + X^\star q^e + v^0 \mathbf{1}_n \in \mathbb{R}^n$ ("par" stands for "partial") represents the effect of the exogenous quantities on voltage, and v^0 is a known, fixed constant representing the squared voltage magnitude at the substation.

We can model this as a discrete-time linear system

$$v(t+1) = X^\star q^c(t) + v^{\text{par}}(t). \tag{3}$$

Substituting $u(t) = q^c(t) - q^c(t-1)$ (change in controllable reactive power injection) and $w(t) = v^{\text{par}}(t) - v^{\text{par}}(t-1)$ (change in exogenous noise) yields the linear dynamical system

$$v(t+1) = v(t) + X^\star u(t) + w(t). \tag{4}$$

The voltage control problem [13] is to drive the squared voltage magnitudes of each bus from an initial state $v(1) \in \mathbb{R}^n$ into a given multi-dimensional interval $[\underline{v}, \overline{v}] \subset \mathbb{R}^n$; it is possible that $v(1)$ does not start within the interval due to some large initial disturbance. For all $t \geq 2$, the voltage control algorithm aims to maintain $v(t)$ within $[\underline{v}, \overline{v}]$, ideally keeping $v(t)$ as close as possible to a "nominal" value $v^{\text{nom}} \in [\underline{v}, \overline{v}]$, typically $v^{\text{nom}} = (\underline{v} + \overline{v})/2$. The cost for deviating from v^{nom} is measured by $\|v(t) - v^{\text{nom}}\|_{P_v}^2$ for some positive semidefinite matrix P_v, where $\|x\|_A^2 := x^\top A x$.

At each time step, buses may change their reactive power injection in order to regulate the voltage close to v^{nom}. The reactive power injection must remain within a given bound $q^c(t) \in [\underline{q}, \overline{q}] \subset \mathbb{R}^n$, and we assume $q^c(0)$ indeed starts within $[\underline{q}, \overline{q}]$. Because buses not in C do not have any ability to control the reactive power injection: $\forall i \notin C.\ \underline{q}_i = \overline{q}_i = 0$. In our model, we do not place any hard "ramp constraints" on $u(t)$. However, we impose a quadratic ramping cost $\|u(t)\|_{P_u}^2$ where P_u is a positive semidefinite matrix.

To drive voltage back to the desired interval, and minimize the aforementioned voltage violation cost $\|v(t) - v^{\text{nom}}\|_{P_v}^2$ and control cost $\|u(t)\|_{P_u}^2$, one needs the exact system dynamics (3) for choosing the optimal reactive power injections $q^c(1), q^c(2), \ldots$. However, in distribution systems, the exact network parameters are often unknown or hard to estimate.

In this paper, we work with the voltage control problem on an *unknown* grid topology. We assume that the true $X^\star$ lies within a known compact set $\mathcal{X} \subset \mathbb{S}_+^n \cap \mathbb{R}_+^{n \times n}$ and that we only have access to the real-time voltage measurement $v(t)$ at each bus. We perform voltage control while learning the system model at the same time. ($\mathbb{S}_+^n$ is the set of $n \times n$ positive semidefinite matrices, and $\mathbb{R}_+^{n \times n}$ is the set of $n \times n$ matrices with nonnegative entries.)

3 ROBUST ONLINE VOLTAGE CONTROL

In this section we introduce our robust online voltage control algorithm and its performance bound (Theorem 1), which is the main result of this paper.

3.1 Algorithm

The structure of the algorithm is summarized in Figure 1 and detailed in Algorithm 1. As the figure shows, the algorithm consists of two main components, a consistent model chasing algorithm SEL and a robust control oracle Π, which are then combined by adapting

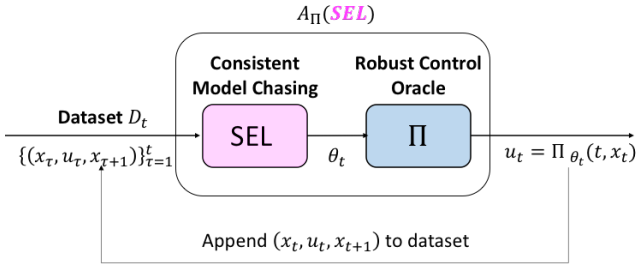

Figure 1: Online robust control framework

ideas from [18]. These two crucial components are detailed in steps (2) and (3) in Algorithm 1, respectively.

The model chasing algorithm SEL performs nested CBC, which is the online problem of choosing a sequence of points within sequentially nested convex sets, with the aim of minimizing the sum of distances between the chosen points [2]. In our setting, the nested convex sets are the consistent sets of possible model parameters, described in Section 4. We use a simple projection-based algorithm that is more computationally efficient than the Steiner point-based approaches, which have state-of-the-art competitive ratio for nested CBC but are not computationally tractable for high-dimensional settings like voltage control.

The robust control oracle Π we use is novel and is developed specifically to provide a provable robustness guarantee (Theorem 3). This robustness guarantee is necessary for the analysis which integrates SEL with Π to provide the finite mistake guarantee of the overall algorithm. See Section 4 for details. Note that other choices for both of these components are possible, as long as they provide the guarantees needed in the analysis in Section 4.

Intuitively, SEL and Π are combined in a way such that Π outputs an action that causes a voltage limit violation, SEL always reduces the uncertainty about the model by a minimum amount. SEL ensures that our current model estimate $\hat{X}_t$ is consistent with observed data collected so far. This model improvement means that Π cannot take too many "bad" actions before the system uncertainty is small.

One important detail in Algorithm 1 is the inclusion of the slack variable ξ. If no slack variable is included (equivalently, with $\xi = 0$ fixed), the optimization problem in Π is guaranteed (Theorem 3) to be feasible and keep the voltage within limits only when the current model estimate $\hat{X}_t$ is close enough to the true model. However, such a guarantee does not necessarily hold when ξ is allowed to be nonzero. On the other hand, since the current model estimate $\hat{X}_t$ may in general be far from the true model, the optimization problem Π without a slack variable may not be feasible; that is, there may not be a control action that keeps the predicted voltage within limits under an incorrect model.

Thus, to be precise, for our finite mistake guarantee (Theorem 1) to hold, the optimization problem for the robust control oracle Π should first be solved without the slack variable. This ensures that if $\hat{X}_t$ is sufficiently close enough to the true model, then the algorithm will not make a mistake. In the case that Π is infeasible, then it should be solved again with a slack variable, which ensures feasibility. However, solving Π twice is unnecessary in practice, and so we have written Algorithm 1 to reflect its practical implementation.

Algorithm 1 Online Robust Voltage Controller

Inputs

- fixed squared voltage magnitude at substation: $v^0 \in \mathbb{R}^n$
- desired nominal squared voltage magnitude: $v^{\text{nom}} \in \mathbb{R}^n$
- limits on the squared voltage magnitude: $[\underline{v}, \overline{v}] \subset \mathbb{R}^n$
- limits on the reactive power injection: $[\underline{q}, \overline{q}] \subset \mathbb{R}^n$
- initial state: $v(1), q^c(0) \in \mathbb{R}^n$
- state and action cost matrices: $P_v, P_u \in \mathbb{S}^n_+$
- compact convex uncertainty set for the model parameter: $\mathcal{X} \subset \mathbb{S}^n_+ \cap \mathbb{R}^{n \times n}_+$
- compact convex uncertainty set for exogenous voltage quantities: $\mathcal{V}^{\text{par}} \subset \mathbb{R}^n$
- upper bound for noise: $\eta > 0$
- weight for slack variable: $\beta > 0$

Procedure

(1) Initialize an empty trajectory $D_0 = [\,]$. Set $t = 1$.
(2) If $t = 1$, initialize estimate of model parameters $\hat{X}_1 \in \mathcal{X}$. Otherwise, query the model chasing algorithm for a new parameter estimate: $\hat{X}_t \leftarrow \texttt{SEL}[D_{t-1}]$.

$$\texttt{SEL}[D_t]: \min_{\hat{X}_t \in \mathbb{S}^n} \|\hat{X}_t - \hat{X}_{t-1}\|^2_{\triangle} \tag{5a}$$

$$\text{s.t. } \hat{X}_t \in \mathcal{X} \tag{5b}$$

$$\forall (v_i, v_{i+1}, u_i, q^c_i) \in D_t:$$

$$-\eta \mathbf{1} \le v_{i+1} - v_i - \hat{X}_t u_i \le \eta \mathbf{1} \tag{5c}$$

$$v_{i+1} - \hat{X}_t q^c_i \in \mathcal{V}^{\text{par}} \tag{5d}$$

(3) Query the robust control oracle for the next control action: $u(t) \leftarrow \Pi_{\hat{X}_t}(v(t))$.

$$\Pi_{\hat{X}_t}: \min_{u \in \mathbb{R}^n, \xi \in \mathbb{R}_+} \left\|\hat{v}' - v^{\text{nom}}\right\|^2_{P_v} + \|u\|^2_{P_u} + \beta \xi^2 \tag{6a}$$

$$\text{s.t. } \underline{q} \le q^c(t-1) + u \le \overline{q} \tag{6b}$$

$$\hat{v}' = v(t) + \hat{X}_t u \tag{6c}$$

$$k = \eta + \rho(\epsilon) \|u\|_2 \tag{6d}$$

$$\underline{v} + (k - \xi)\mathbf{1} \le \hat{v}' \le \overline{v} - (k - \xi)\mathbf{1} \tag{6e}$$

(4) Update $q^c(t) \leftarrow q^c(t-1) + u(t)$. Apply the control action $u(t)$. Observe the system transition to $v(t+1) = v(t) + X^\star u(t) + w(t)$.
(5) Append $(v(t), v(t+1), u(t), q^c(t))$ to the trajectory D_{t-1},

$$D_t = \left[(v(i), v(i+1), u(i), q^c(i))\right]_{i=1}^{t}.$$

(6) Increment $t \leftarrow t + 1$. Repeat from Step (2).

3.2 Assumptions

Before presenting the main results, we introduce several key assumptions that underlie our analysis and discuss why they are both needed and practical.

The first assumption, stated below, is standard and bounds the noise in the dynamics.

Assumption 1. *The change in noise is bounded as*

$$\forall t: \|w(t)\|_\infty \le \eta, \tag{7}$$

where $w(t) = v^{par}(t) - v^{par}(t-1)$, $\eta \in [0, \overline{\eta}]$ *is a known constant and* $\overline{\eta} = \min_{i=1,\dots,n}(\overline{v}_i - \underline{v}_i)/2$.

Physically, this bound represents an assumption that the active and exogenous reactive power injection does not vary dramatically between time steps, as can be seen by expanding the definition of $w(t)$:

$$\begin{aligned} w(t) &= v^{\text{par}}(t) - v^{\text{par}}(t-1) \\ &= R^\star(p(t) - p(t-1)) + X^\star(q^e(t) - q^e(t-1)). \end{aligned}$$

For example, if the net active and exogenous reactive power injection is the same at time steps t and $t-1$, then $w(t) = 0$.

The requirement that $\eta \le \overline{\eta}$ ensures that the space between $\underline{v} + \eta\mathbf{1}$ and $\overline{v} - \eta\mathbf{1}$ is nonempty so that, for any $v(t)$ inside this space, the addition of noise $w(t)$ does not push the voltage outside of the target voltage limits. On the other hand, if this space were empty, then for any $v(t)$ and $u(t)$, there would exist some $w(t)$ satisfying $\|w(t)\|_\infty \le \eta$ such that $v(t+1)$ exceeds the desired voltage limits. Because we seek a voltage controller that is robust to change in exogenous noise, we must assume that the space is nonempty.

Our second assumption provides a bound on the uncertainty about the network topology, such as the maximum connectivity and impedances.

Assumption 2. *The true model* $X^\star$ *lies within a known compact, convex set* $\mathcal{X} \subset \mathbb{S}^n_+ \cap \mathbb{R}^{n \times n}_+$. *Let* $\text{diam}(\mathcal{X})$ *denote the diameter of* $\mathcal{X}$*:*

$$\text{diam}(\mathcal{X}) = \sup_{X_1, X_2 \in \mathcal{X}} \|X_1 - X_2\|_{\triangle}.$$

Assumption 2 ensures that the unknown true model parameters $X^\star$ belong to a compact, convex set $\mathcal{X}$, which is a minimal assumption necessary for proving an analytic guarantee. This compact set forms the starting point of our consistent model chasing algorithm SEL.

We equip the space $\mathbb{S}^n$ of $n \times n$ symmetric matrices with a norm $\|\cdot\|_{\triangle}$ defined as

$$\|A\|_{\triangle} = \|\text{upper-triangle}(A)\|_2 = \sqrt{\sum_{i=1}^{n} \sum_{j=i}^{n} A^2_{ij}}.$$

We use this choice of norm in order to isometrically map the space of symmetric matrices to Euclidean space, thereby enabling us to take advantage of known results on nested convex body chasing within Euclidean space.

Finally, our third assumption is about the existence of feasible control actions for the robust control oracle. This assumption can be interpreted as either a bound on the noise, or a requirement that the controllable reactive power injection be flexible enough to satisfy the demand of any noise.

Assumption 3. *There exists a compact, convex set* $\mathcal{V}^{par} \subset \mathbb{R}^n$ *such that* $\forall t \ge 0: \; v^{par}(t) \in \mathcal{V}^{par}$. *Furthermore, for some known* $\epsilon > 0$,

$$\forall v^{par} \in \mathcal{V}^{par}, X \in \mathcal{X}.$$

$$\exists q^c \in [\underline{q}, \overline{q}] \text{ s.t. } Xq^c + v^{par} \in [\underline{v} + (\eta + \epsilon)\mathbf{1}, \overline{v} - (\eta + \epsilon)\mathbf{1}].$$

Intuitively, the η padding is required for robustness to the noise $w(t)$, while the ϵ padding is required for robustness to model uncertainty (*i.e.*, uncertainty about $X^\star$).

3.3 Main result

We can now state our main result, which is a finite-error bound for Algorithm 1.

Theorem 1 (Main Result). *Under Assumptions 1 to 3, Algorithm 1 ensures that the voltage limits will be violated at most*

$$\frac{4\gamma}{\rho(\epsilon)}\operatorname{diam}(\mathcal{X}) + 1 \tag{8}$$

times, where $\rho(\epsilon) = \frac{\epsilon}{2\left\|\overline{q}-\underline{q}\right\|_2}$, $\gamma = \pi(m-1)m^{m/2}$, *and* $m = n(n+1)/2$.

To the best of our knowledge, this result is the first provable stability bound for voltage control in a setting where the network topology is unknown. It highlights that the Algorithm 1 can ensure stability even after *unknown* changes to the network topology, *e.g.*, due to maintenance, failures, etc., without the need to perform system identification.

Intuitively, this result guarantees that the model chasing algorithm SEL will learn a "good enough" model for control quickly. When the robust controller Π makes a mistake, the model chasing algorithm will learn from that mistake and significantly reduce the set of consistent models. Because the initial set of consistent models is bounded, and this set shrinks a significant amount after each mistake, the total number of mistakes is bounded.

To interpret the error bound (8) in Theorem 1, we notice that it is proportional to the diameter $\operatorname{diam}(\mathcal{X})$ of the parameter space $\mathcal{X}$ and the competitive ratio γ of the consistent model chasing algorithm, and inverse proportional to the oracle robustness margin ρ. Note that the dependence on m of the consistent model chasing competitive ratio γ is very conservative. Because of computational tractability concerns, our implementation of SEL uses a projection-based algorithm rather than the state-of-the-art Steiner point method [2, 5]. For the algorithm in [5], $\gamma_{\text{Steiner}} = m/2$. As our case studies show, in practice the projection-based algorithm used in SEL performs much better than the worst-case bound.

We outline a proof of Theorem 1 in the next section, but before doing so we want to highlight one piece of that proof that is of independent interest. In particular, a major step in the proof is to provide a feasibility guarantee for the robust control oracle component Π of the algorithm, which is done in Theorem 3.

4 PROOFS

We now prove our main result Theorem 1. Our proof builds on and adapts the approach of [18], which outlines a general framework for integrating model chasing and robust control. To explain the general framework, we first consider a discrete-time nonlinear dynamical system

$$x_{t+1} = f_*(x_t, u_t) + w_t, \qquad x_0 \text{ given}, \qquad (f_*, \mathbf{w}) \in \mathcal{F},$$

where $x \in \mathcal{S} \subseteq \mathbb{R}^n$ is the system state and $u \in \mathcal{U} \subseteq \mathbb{R}^m$ is the control input. The unknown function f_* and disturbance sequence $\mathbf{w} \in \ell^\infty(\mathbb{Z}_+; \mathbb{R}^n)$ belong to an uncertainty set $\mathcal{F}$, and the disturbance is bounded as $\|\mathbf{w}\|_\infty \le \eta$. Assume that $\mathcal{F}$ has a *compact parametrization* $(\mathbb{T}, \mathsf{K}, d)$, where $\mathbb{T} : \mathsf{K} \to \wp(\mathcal{F})$ is a mapping from the parameter space K to a set of functions and disturbances such that $\mathcal{F} \subseteq \bigcup_{\theta\in\mathsf{K}} \mathbb{T}[\theta]$. $\wp(\mathcal{F})$ denotes the powerset of $\mathcal{F}$. Let d denote a metric on K, so (K, d) is a compact metric space.

The control objective is specified as a sequence of indicator "goal" functions $\mathcal{G} = (\mathcal{G}_0, \mathcal{G}_1, \dots)$. Each $\mathcal{G}_t : \mathcal{X} \times \mathcal{U} \to \{0, 1\}$ encodes a desired condition per time step t:

$$\mathcal{G}_t(x_t, u_t) = \mathbf{1}[x_t, u_t \text{ violate desired condition at time } t].$$

The main result of [18] specifies a set of sufficient conditions for a finite-mistake guarantee. These conditions decouple online robust control into separate online learning and robust control components.

The online learning component requires a consistent model chasing algorithm SEL, which takes as input the current observed trajectory $D_t = [(x_i, x_{i+1}, u_i)]_{i=1}^t$ and outputs an estimated parameter $\theta_t \in \mathsf{K}$. The estimated parameter θ_t must be *consistent* with D_t.

Definition 1 (Consistent Parameter). *We say* $\theta \in \mathsf{K}$ *is consistent with* D_t *if there exists* $(f, \mathbf{w}) \in \mathbb{T}[\theta]$ *such that*

$$\forall (x_t, x_{t+1}, u_t) \in D_t : \; x_{t+1} = f(x_t, u_t; \theta) + w_t.$$

Let P_t denote the set of all parameters consistent with D_t. We say SEL is γ-competitive if $\sum_{t=1}^\infty d(\theta_t, \theta_{t-1}) \le \gamma \max_{\theta\in\mathsf{K}} d(P_\infty, \theta)$ holds for a fixed constant $\gamma > 0$, which we call the *competitive ratio*.

The robust control component requires a control oracle Π, which given the current state x_t and a parameter θ_t, outputs a control action $u_t = \Pi_{\theta_t}(x_t)$ that is robust for all systems that are close to θ_t. In particular, we call a control oracle ρ-*robust* for control objective $\mathcal{G}$, if all trajectories in $S^\Pi[\rho; \theta]$ achieve $\mathcal{G}$ after finitely many mistakes, and $S^\Pi[\rho; \theta]$ is defined as the set of all possible trajectories generated by $\Pi_{\hat\theta}$ for all $\hat\theta$ such that $d(\theta, \hat\theta) \le \rho$.

$$S^\Pi[\rho;\theta] = \left\{ \begin{array}{l|l} D_\infty = [(x_t, x_{t+1}, u_t)]_{t=1}^\infty : & (f, \mathbf{w}) \in \mathbb{T}[\theta], \\ \quad u_t = \Pi_{\hat\theta}(x_t) & d(\hat\theta, \theta) \le \rho \\ \quad x_{t+1} = f(x_t, u_t) + w_t & \end{array} \right\} \tag{9}$$

Due to the page limit, we refer readers to [18] for a more detailed discussion of consistent model chasing algorithms and ρ-robust control oracles. As a summary, if SEL chases consistent models and Π is a robust oracle for $\mathcal{G}$, then the resulting $A_\Pi(\mathsf{SEL})$ algorithm achieves a finite mistake guarantee, which is stated in the following.

Theorem 2. *[18, Theorem 2.5] Assume that* SEL *chases consistent models and* Π *is a robust oracle for objective* $\mathcal{G}$. *Then for any starting point* x_0 *and trajectory* $[(x_t, u_t)]_{t=0}^\infty$ *generated by* $\mathcal{A}_\Pi(\mathsf{SEL})$ *(illustrated in Figure 1), the following mistake guarantees hold: (i) If* Π *is robust, then* $\sum_{t=0}^\infty \mathcal{G}_t(x_t, u_t) < \infty$; *(ii) If* Π *is uniformly* ρ*-robust and* SEL *is* γ*-competitive, then*

$$\sum_{t=0}^{\infty} \mathcal{G}_t(x_t, u_t) < \max\left\{1, M_\rho^\Pi\right\}\left(\frac{2\gamma}{\rho}\operatorname{diam}(\mathsf{K}) + 1\right)$$

where M_ρ^Π *denotes the worst case total mistakes of the* ρ*-robust control oracle* Π.

To apply Theorem 2 to prove Theorem 1 we need to prove that (i) the proposed algorithm Equation (5) chases consistent models and has a bounded competitive ratio, and (ii) the proposed robust algorithm in Equation (11) is a ρ-robust control oracle, for bounded

disturbance in the system topology. In particular, the correspondence of the definitions are as follows. We have $\theta = X$, and

$$\begin{aligned} \mathsf{K} &= \mathcal{X}, \quad v(1), q^c(0) \text{ given} \\ d(X, X') &= \|X - X'\|_\triangle \\ \mathbb{T}[X] &= \{(f, \mathbf{w}) \mid f(v, u) = v + Xu,\ \|\mathbf{w}\|_\infty \le \eta\} \\ \mathcal{F} &= \bigcup_{X \in \mathcal{X}} T[X] \\ \mathcal{G}_t &= \mathbf{1}[v(t) \in [\underline{v}, \overline{v}]]. \end{aligned}$$

We begin by proving that the set defined from (5b), (5c), (5d) in Algorithm 1 is consistent with the trajectory D_T.

Lemma 1 (SEL is consistent). *Suppose D_T is a trajectory of voltage measurements and control actions taken up to time T:*

$$D_T = [v(t), v(t+1), u(t), q^c(t)]_{t=1}^T.$$

The set

$$P_T := \left\{ \hat{X} \in \mathcal{X} \;\middle|\; \begin{array}{l} \forall \left[v(t), v(t+1), u(t), q^c(t)\right] \in D_T : \\ \quad \left\|v(t+1) - v(t) - \hat{X}u(t)\right\|_\infty \le \eta, \\ \quad v(t+1) - \hat{X}q^c(t) \in \mathcal{V}^{par} \end{array} \right\} \tag{10}$$

is a consistent set for D_T, i.e., $\hat{X} \in \mathcal{X}$ is consistent (Definition 1) if and only if $\hat{X} \in P_T$.

Proof. Consider any $\hat{X} \in P_T$. For $t \in \{1, \dots, T\}$, define

$$\hat{w}(t) := v(t+1) - v(t) - \hat{X}u(t).$$

Then, $\hat{w}(t)$ clearly satisfies Assumption 1. Moreover, let

$$v^{\text{par}}(t) := v(t+1) - \hat{X}q^c(t).$$

Then $v^{\text{par}}(t)$ is the corresponding admissible noise that satisfies Assumption 3.

Conversely, if $\hat{X} \notin P_T$, then either $\left\|v(t+1) - v(t) - \hat{X}u(t)\right\|_\infty > \eta$ for some t (and therefore $\|\hat{w}(t)\|_\infty > \eta$) or $v^{\text{par}}(t)$ violates Assumption 3. Thus, P_T contains exactly all $X \in \mathcal{X}$ that are consistent with D_T. □

Observe that P_T is a closed, bounded, and convex set. Furthermore, it is non-empty, since $X^\star \in P_T$. Intuitively, P_T is the smallest set containing all parameters that could generate the observed trajectory $\{v(t)\}_{t=1}^{T+1}$ along with a corresponding admissible sequence of noise compatible with Assumptions 1 to 3.

Now that we have defined the consistent sets P_t, we can express SEL equivalently as solving $\min_{\hat{X}_t \in \mathbb{S}^n} \left\|\hat{X}_t - \hat{X}_{t-1}\right\|_\triangle^2$ s.t. $\hat{X}_t \in P_t$. This is a nested convex body chasing problem, where we aim to minimize the movement distance $\left\|\hat{X}_t - \hat{X}_{t-1}\right\|_\triangle$ between nested convex sets $P_t \subseteq P_{t-1}$. By leveraging known results about nested convex body chasing algorithms [2, 5], we can prove that the model chasing algorithm SEL described by (5) has a bounded competitive ratio. This is formalized in the following lemma.

Lemma 2 (SEL is competitive). *For any compact convex space $\mathsf{K} \subset \mathbb{S}^n$, the greedy projection algorithm for consistent model chasing (CMC) in Equation (5) achieves a competitive ratio*

$$\gamma = \pi(m-1)m^{m/2}$$

where $m = n(n+1)/2$.

Proof. The normed vector space $(\mathbb{S}^n, \|\cdot\|_\triangle)$ is isometrically isomorphic to the Euclidean space $(\mathbb{R}^m, \|\cdot\|_2)$ with $m = n(n+1)/2$. The mapping between the two spaces is the vectorization of the upper-triangle of the symmetric matrix:

$$A \in \mathbb{S}^n \longleftrightarrow [A_{1,1:n}, A_{2,2:n}, \dots, A_{n,n}]^\top \in \mathbb{R}^m.$$

Note that $\forall d \in \mathbb{N}: \ (\omega_d/\omega_{d-1}) \le \pi$, where ω_d denotes the surface area of the d-sphere. Then, by [2, Theorem 1.3], the greedy projection algorithm achieves a competitive ratio of at most $\pi(m-1)m^{m/2}$. □

Finally, we show that our controller Π is ρ-robust. In particular, we prove that $\Pi_{\hat{X}}$ makes no mistakes ($M_\rho^\Pi = 0$) given a consistent model $\hat{X} \in P_t$.

Theorem 3 (Π is ρ-robust). *Under Assumptions 1 to 3, suppose $\hat{X} \in P_t$, where P_t is given in (10) for $t \ge 1$. Define $\rho(\epsilon) = \epsilon / \left(2\|\overline{q} - \underline{q}\|_2\right)$. Then, the following optimization problem is feasible,*

$$\min_{u \in \mathbb{R}^n} \quad \left\|\hat{v}' - v^{nom}\right\|_{P_v}^2 + \|u\|_{P_u}^2 \tag{11a}$$

$$s.t. \quad \underline{q} \le q^c(t-1) + u \le \overline{q} \tag{11b}$$

$$\hat{v}' = v(t) + \hat{X}u \tag{11c}$$

$$k = \eta + \rho(\epsilon)\|u\|_2 \tag{11d}$$

$$\underline{v} + k\mathbf{1} \le \hat{v}' \le \overline{v} - k\mathbf{1}. \tag{11e}$$

Further, the solution of (11), $u(t)$, guarantees voltage stability for all $X \in \mathcal{X}$ such that $\left\|X - \hat{X}\right\|_\triangle \le \rho(\epsilon)$. That is, $v(t) + Xu(t) + w(t) \in [\underline{v}, \overline{v}]$ for all $w(t)$ satisfying Assumption 1.

Observe that Equation (11) corresponds to Equation (6) in Algorithm 1 with the slack variable set to zero. We note that the robustness margin ρ decreases as $[\underline{q}, \overline{q}]$ increase. The intuitive reason is that the voltage is more sensitive to changes in $\hat{X}$ when the range of possible u's expands. Therefore, a fixed voltage buffer of ϵ in constraints (6d) and (11e) affords less robustness to changes in $\hat{X}$ as $[\underline{q}, \overline{q}]$ gets larger.

Proof of Theorem 3. First, we prove that the optimization problem (11) is feasible. Define a new variable $q^c := q^c(t-1) + u$. Let

$$v^{\text{par}}(t-1) := v(t) - \hat{X}q^c(t-1) \tag{12}$$

$$\hat{v}'(u) := v(t) + \hat{X}u = v(t) + \hat{X}[q^c - q^c(t-1)] \tag{13}$$

$$= \hat{X}q^c + v^{\text{par}}(t-1) \tag{14}$$

$$v'(u) := v(t) + Xu = \hat{v}'(u) + (X - \hat{X})u. \tag{15}$$

Here, v^{par} is the conjectured admissible noise when we assume the underlying parameter is $\hat{X}$. Recall (4) and we can interpret $\hat{v}'$ as the one-step prediction of voltage under the selected consistent model $\hat{X}$ given a control action u and the current voltage $v(t)$, without the disturbance term. Similarly, v' is the disturbance-free one-step prediction of voltage under a different model X such that $\left\|X - \hat{X}\right\|_\triangle \le \rho(\epsilon)$ given the same control action u and the current voltage $v(t)$.

Since $\hat{X}$ is consistent with P_{t-1}, $v^{\text{par}}(t-1) \in \mathcal{V}^{\text{par}}$. Therefore, by Assumption 3, there exists $q^c \in [\underline{q}, \overline{q}]$ such that

$$\underline{v} + (\eta + \epsilon)\mathbf{1} \le \hat{v}'(u) \le \overline{v} - (\eta + \epsilon)\mathbf{1}. \tag{16}$$

Then the u corresponding to such q^c, *i.e.*, $u = q^c - q^c(t-1)$, clearly satisfies constraint (11b). By Lemma 3, for all X satisfying $\|X - \hat{X}\|_\triangle \le \rho(\epsilon)$,

$$-\rho(\epsilon)\|u\|_2 \mathbf{1} \le (X - \hat{X})u \le \rho(\epsilon)\|u\|_2 \mathbf{1}. \tag{17}$$

Adding (17) with (16) and using (15) yields

$$\underline{v} + (\eta + \epsilon - \rho(\epsilon)\|u\|_2)\mathbf{1} \le v'(u) \le \overline{v} - (\eta + \epsilon - \rho(\epsilon)\|u\|_2)\mathbf{1}. \tag{18}$$

By choosing

$$\rho(\epsilon) = \frac{\epsilon}{2\left\|\overline{q} - \underline{q}\right\|_2},$$

and observing that u must satisfy $\|u\|_2 \le \left\|\overline{q} - \underline{q}\right\|_2$ since $u = q^c - q^c(t-1)$, we have

$$\rho(\epsilon)\|u\|_2 \le \frac{\epsilon}{2} \le \epsilon - \rho(\epsilon)\|u\|_2.$$

Therefore, we can use this relation to upper and lower bound (18) and arrive at

$$\underline{v} + (\eta + \rho(\epsilon)\|u\|_2)\mathbf{1} \le v'(u) \le \overline{v} - (\eta + \rho(\epsilon)\|u\|_2)\mathbf{1} \tag{19}$$

This holds for any one-step noiseless prediction $v'(u)$ from X such that $\left\|X - \hat{X}\right\|_\triangle \le \rho(\epsilon)$. Since $\hat{X}$ trivially lies in this set, $\hat{v}'$ also satisfies (19), which in turn shows that u satisfies constraint (11e). Thus, the optimization problem is feasible.

Next, we show that every solution u from (11) generated with $\hat{X}$ guarantees that $v(t) + Xu \in [\underline{v}, \overline{v}]$ for all X such that $\left\|X - \hat{X}\right\|_\triangle \le \rho(\epsilon)$. Subtracting (17) from constraint (11e) and using Equation (15) yields

$$\underline{v} + \eta\mathbf{1} \le v'(u) \le \overline{v} - \eta\mathbf{1}.$$

By Assumption 1, any admissible disturbance $w(t)$ is bounded as $-\eta\mathbf{1} \le w(t) \le \eta\mathbf{1}$. This means that

$$\underline{v} \le v(t+1) = v'(u) + w(t) \le \overline{v},$$

which shows the control action computed using (11) guarantees voltage stability for all $X \in \mathcal{X}$ such that $\left\|X - \hat{X}\right\|_\triangle \le \rho(\epsilon)$. □

Lemma 3. *For all $A \in \mathbb{S}^n$, $b \in \mathbb{R}^n$, and $\alpha \in \mathbb{R}_+$,*

$$\|A\|_\triangle \le \alpha \quad \Longrightarrow \quad -\alpha\|b\|_2 \mathbf{1} \le Ab \le \alpha\|b\|_2 \mathbf{1}.$$

Proof. Let A_i denote the i^{th} row of A. By symmetry of A,

$$\|A_i\|_2^2 = \sum_{j=1}^{n} A_{i,j}^2 = \sum_{k=1}^{i-1} A_{k,i}^2 + \sum_{j=i}^{n} A_{i,j}^2$$
$$\le \sum_{k=1}^{n}\sum_{j=k}^{n} A_{k,j}^2 = \|A\|_\triangle^2 \le \alpha^2,$$

so $\|A_i\|_2 \le \alpha$. Then

$$-\alpha\|b\|_2 \le -\|A_i\|_2\|b\|_2 \le (Ab)_i \le \|A_i\|_2\|b\|_2 \le \alpha\|b\|_2.$$

This holds for all $i \in \{1, \ldots, n\}$, which yields the desired result. □

Finally, combining Theorem 3 with Lemma 2 and applying Theorem 2 completes the proof of Theorem 1.

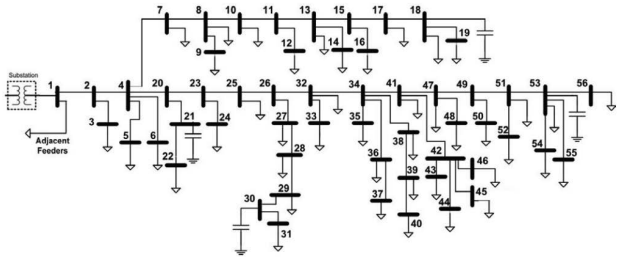

Figure 2: Schematic diagram of SCE 56 bus distribution system.

5 CASE STUDY

We demonstrate the effectiveness of Algorithm 1 using a case study based on a single-phase 56-bus network from the Southern California Edison (SCE) utility (Figure 2). We use a setup that mimics what has been used previously in the voltage control literature. In particular, the detailed line parameters r_{ij} and x_{ij} that we use for this network are the same as those in Table 1 of [13]. In our experiments, we use the linear power model in Equation (2) to solve for voltages; we do not use non-linear power flow models.

5.1 Experimental Setup

We use real-world data [26] collected from the unmodified network, with power injection at buses in $C = \{2, 4, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 19, 20, 23, 25, 26, 32\}$ adjusted by scaled actual generation from a nearby photovoltaic system. Exogenous active and reactive power injection measurements are taken at each bus at 6-second intervals over a 24-hour period. We assume that controllers with reactive power injection capacity are set up at every node. Figure 3a plots these values for several buses to illustrate the setting considered.

The network parameters used in our experiments are as follows:

- nominal squared voltage magnitude at the substation $v^0 = v^{\text{nom}} = (12\text{kV})^2$
- squared voltage magnitude limits $[\underline{v}, \overline{v}] = [0.95\text{pu}, 1.05\text{pu}] = [(11.4\text{kV})^2, (12.6\text{kV})^2]$
- reactive power injection limits $[\underline{q}, \overline{q}] = [-0.24\text{MVar}, 0.24\text{MVar}]$
- state and input cost matrices $P_{\overline{v}} = 0.1I$, $P_u = 10I$
- initial state $v(1) = Rp(0) + Xq^e(0) + v^0\mathbf{1}$, $q^c(0) = \mathbf{0}$

Note that, in comparison to previous papers in the voltage control literature, our reactive power injection limits $[\underline{q}, \overline{q}]$ are slightly more generous than the ±0.2 MVar used in, e.g., [26]. We choose ±0.24 MVar because even a controller with perfect knowledge of the future would need reactive power injection capabilities of at least ±0.238 MVar in order to maintain $v(t) \in [\underline{v}, \overline{v}]$ (if $\underline{q} = -\overline{q}$).

In order to truly satisfy the requirement in Assumption 3 $v(t) \in [\underline{v} + (\eta + \epsilon), \overline{v} - (\eta + \epsilon)]$ with $\epsilon = 0.1$, the reactive power injection capabilities need to be at least ±0.455 MVar. As we show in our experiments with only ±0.24 MVar range of control, Assumption 3 does not need to be fully satisfied in order for our method to still provide strong empirical results.

We fix $\epsilon = 0.1$ in our experiments, corresponding to a robustness margin $\rho = \epsilon/(2\|\overline{q} - \underline{q}\|_2) \approx 0.014$. We set $\eta = 8.65$, which upper-bounds the true maximum change in exogenous noise observed in

our data:

$$\max_t \left\| R(p(t) - p(t-1)) + X(q^e(t) - q^e(t-1)) \right\|_\infty \approx 8.625.$$

For the robust controller Π, we use a slack variable weight $\beta = 100$ in the cost function, and we set $\mathcal{V}^{\text{par}} = [\underline{v^{\text{par}}}, \overline{v^{\text{par}}}]$ to be the rectangle around the true noise:

$$\forall i \in \{1, \dots, n = 55\}: \quad \underline{v^{\text{par}}}_i = \min_t v_i^{\text{par}}(t), \quad \overline{v^{\text{par}}}_i = \max_t v_i^{\text{par}}(t).$$

In our implementation of the consistent model chasing algorithm, we make a few changes from the procedure described in Algorithm 1. In order to keep the consistent model chasing optimization problem (5) computationally tractable, we do not use the full trajectory D as in constraints (5c)-(5d). Instead, we include the 20 latest observations and a set of 80 more random samples $(v(t), v(t+1), u(t), q^c(t)) \sim D$. This provides a computationally tractable approximation of the uncertainty set.

Our experiments aim to understand the performance of the proposed online robust controller under different levels of uncertainty and the effect of different initializations of $\hat{X}$. In particular, we consider uncertainty sets of the form

$$\mathcal{X}_\alpha = \left\{ \hat{X} \in \mathbb{S}_+^n \cap \mathbb{R}_+^{n \times n} : \ \|\hat{X} - X^\star\|_\triangle \le \alpha \left\|X^\star\right\|_\triangle \right\} \tag{20}$$

where larger α corresponds to a larger uncertainty set. Note that $\text{diam}(\mathcal{X}_\alpha) = 2\alpha \left\|X^\star\right\|_\triangle$.

We initialize $\hat{X}$ by adding noise to the true $X^\star$ in two ways. First, we multiply the line impedance coefficients x_{ij} by a random scaling factor sampled from $\text{Uniform}[1-\sigma, 1+\sigma]$ for some $\sigma \in [0, 1]$. Second, we randomly permute the bus ordering, so that $\hat{X}$ corresponds to a permuted grid topology. Finally, we project $\hat{X}$ into the uncertainty set $\mathcal{X}_\alpha$.

We consider 3 settings:

(1) moderate uncertainty and error ($\alpha = 0.5$), with $\sigma = 0.5$
(2) large uncertainty and error ($\alpha = 1.0$), with $\sigma = 1.0$
(3) large uncertainty ($\alpha = 1.0$) with $\sigma = 1.0$, but with moderate error (the initial $\hat{X}$ is projected into $\mathcal{X}_{0.5}$)

By comparing case (3) to cases (1) and (2), we distinguish the impact of the uncertainty set size and the error of the initial guess for $\hat{X}$, each of which presents different challenges for Algorithm 1.

We compare our experiments against four baselines, illustrated in Figure 3: (a) the case of no controller, (b) the case where our robust controller is used with perfect knowledge of the network topology, (c) a model-agnostic decentralized controller from [21, Section IV], and (d) another model-agnostic decentralized controller from [21, Section V]. The figure highlights that, without a controller, buses 19 and 31 violate the upper and lower voltage limits, respectively, by a significant margin. In contrast, the robust controller given the true $X^\star$ keeps the voltage within the limits for all buses, as expected. The decentralized model-agnostic controllers (c) and (d), which are supposed to be robust to the underlying topology, do not perform well, as their theoretical convergence guarantees [21] only hold for fixed v^{par} and for unbounded reactive power injection limits $(\underline{q}, \overline{q})$.

5.2 Experimental Results

Our experimental results focus on demonstrating the ability of Algorithm 1 to stabilize the system without knowledge of the network topology. To highlight the performance of the algorithm, we consider settings with moderate and large amounts of uncertainty in Figures 4 and 5. Importantly, Algorithm 1 stabilizes the system without performing system identification. In fact, our results highlight that the algorithm still has significant uncertainty about the topology at the end of the experiments, while also providing nearly the same stabilization performance as the robust controller does with complete information about the topology.

Our simulations use the linearized system dynamics Equation (3), and the convex optimization problems for SEL and Π are solved with CVXPY [10] using the MOSEK solver. Code for our simulations will be made available on GitHub.

Experiment 1: Moderate Uncertainty. Our first set of experiments runs the robust online voltage controller (Algorithm 1, denoted Π+SEL) in a scenario where the norm of the initial parameter estimate is within 50% of the true parameter norm. Figure 4 shows that the proposed controller can consistently maintain a voltage profile within the nominal operation range. The solid orange line and shaded orange region represent the mean and ±1 standard deviation of voltages from the proposed controller, over 4 random choices of $\hat{X}_1$. In contrast, the robust controller Π on its own (without the consistent model chasing algorithm SEL to update the model) exhibits large voltage oscillations shown in blue and fails to stabilize the voltage. This is a consequence of Π choosing actions based on incorrect information about the network topology. Using the wrong knowledge is even worse than not having a controller (Figure 3a) in this case. This demonstrates a clear need for learning the topology rather than purely replying on a robust voltage controller in the case of uncertainty about the network topology.

Figure 4c shows the evolution of the convex model chasing algorithm results across different choices of $\hat{X}_1$. Notably, even though we only approximate the true consistent set through a small random sample, the distance between the learned model $\hat{X}_t$ and true model $X^\star$ decreases monotonically over time. However, the uncertainty does not converge to zero, illustrating that Algorithm 1 does not perform complete system identification and instead learns just enough about the topology in order to stabilize the system. This is a key novelty of our approach, and enables the algorithm to quickly adapt to uncertain network conditions.

Finally, we compare the performance of Algorithm 1 with that of the robust controller that has perfect knowledge of the network topology (Figure 3b). The performance of Algorithm 1 is comparable despite its lack of knowledge of the topology. One reason for the near-optimal performance of Algorithm 1 in this case is that the robust controller Π turns out to be even more robust than our theoretical results suggest. In particular, empirically it is robust to $\hat{X}$ up to a distance $0.4 \left\|X^\star\right\|_\triangle \approx 66$ away from the true $X^\star$, even though our theoretical results only guarantee its robustness for distances up to $\rho \approx 0.014$.

Experiment 2: Large Uncertainty. Next, we test the proposed robust online voltage controller in a more challenging setting where there is a large amount of uncertainty. Here, the initial $\hat{X}$ is generated from impedance values with up to 100% error from the true values. Consistent with the moderate uncertainty case, our method manages to maintain voltage stability across all buses. Even though the initial model estimation can be very different compared to the

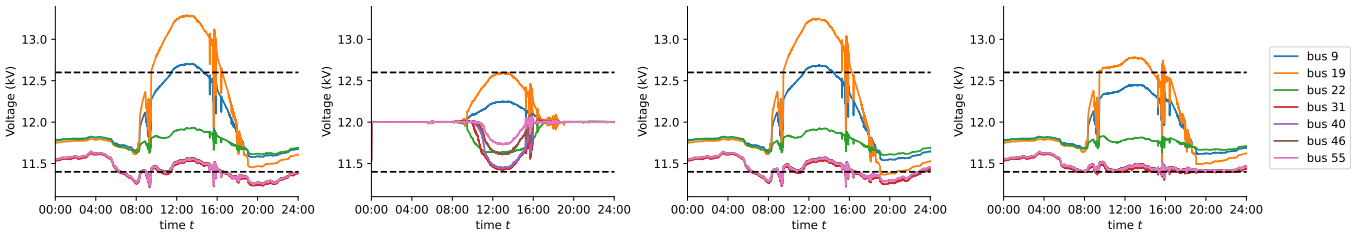

Figure 3: Voltage profile of SCE 56 bus distribution system with PV generators. (a) without control (b) robust controller Π given the true $X^\star$ (c) model-agnostic controller from [21, Section IV] (d) model-agnostic controller from [21, Section V].

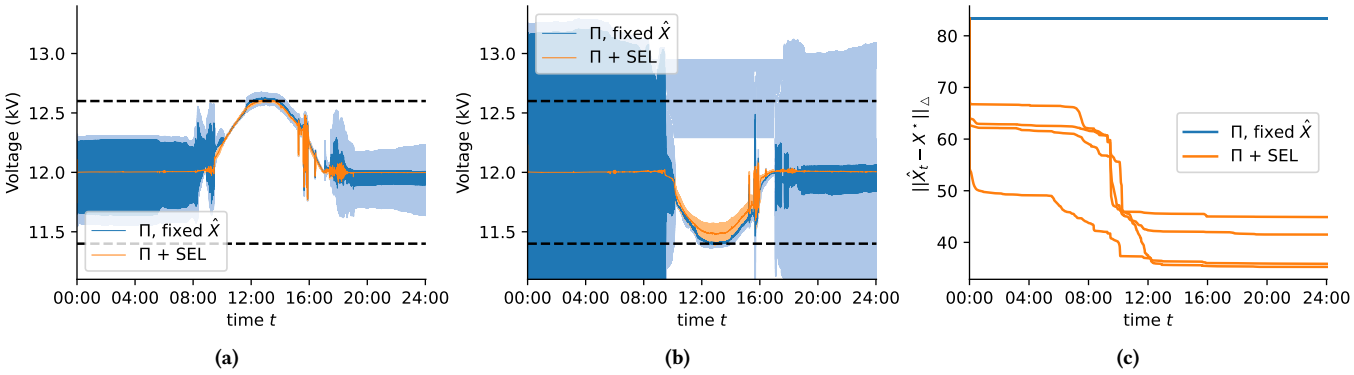

Figure 4: Results for Algorithm 1 on the SCE 56-bus distribution system with moderate-sized uncertainty set ($\alpha = 0.5$) and moderate initial error ($\|\hat{X} - X^\star\|_\triangle = 0.5\left\|X^\star\right\|_\triangle$). (a) Voltage profile of bus 19, where dark blue and dark orange lines plot the mean voltage for the robust controller Π with fixed $\hat{X}$ vs. Π paired with SEL, respectively, across 4 random initializations of $\hat{X}$; shaded light blue and light orange regions indicate ± 1 standard deviation. The dark blue line looks like a blob because Π over-corrects its mistakes when given incorrect network parameters. (b) Same as (a), but for bus 31. (c) Model uncertainty decreases as SEL learns over time, for the same 4 choices of $\hat{X}$ from (a) and (b). The algorithm keeps the voltage within limits even though the model estimates $\hat{X}_t$ are imperfect, demonstrating that complete system identification is not necessary.

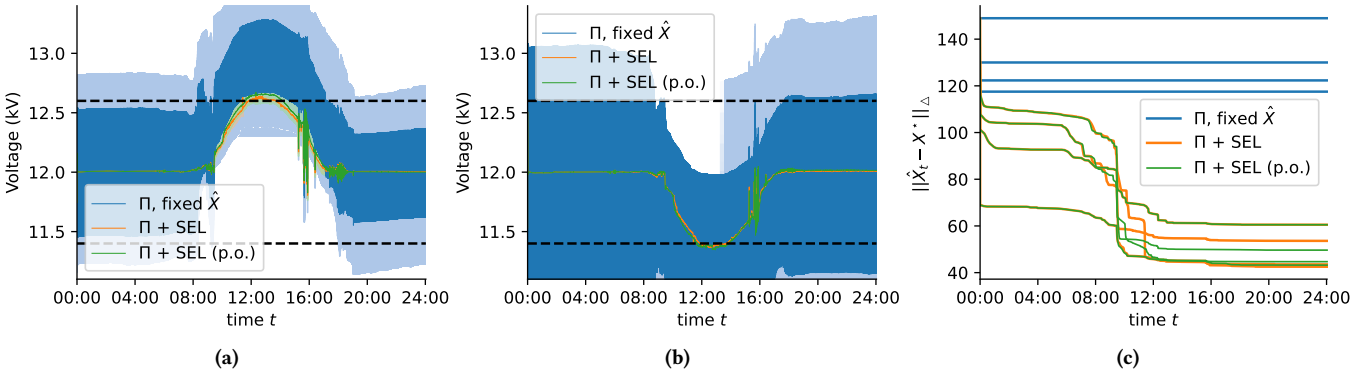

Figure 5: Same as Figure 4, except the uncertainty set and initial error are both larger ($\alpha = 1.0$, $\|\hat{X} - X^\star\|_\triangle = \left\|X^\star\right\|_\triangle$.) The robust controller Π with fixed $\hat{X}$ performs significantly worse than in Figure 4, while Π paired with SEL performs only slightly worse. We also experiment with partial observability of voltages (denoted "p.o.") by withholding observations of voltages from certain buses, including buses 19 and 31 whose voltage profiles are plotted in (a) and (b). As expected, our controller performs worse with partially observations compared to full observations, but only marginally so. In (c), it is visible that the lack of voltage observation at these 7 nodes does not significantly affect the ability of SEL to learn consistent models.

ground truth system (*i.e.*, $\left\|\hat{X} - X^\star\right\|_\triangle > 100$), Algorithm 1 quickly learns from the mistakes and refines the model estimation. Notice in Figure 5c how the model error drops quickly during the voltage violation periods between time 10:00 and 12:00. However, the robust controller on its own (without SEL) fails to control the voltage in the large uncertainty scenario: the voltage profiles at buses 19 (Figure 5a) and 31 (Figure 5b) deviate significantly from the nominal value. This may lead to unsafe operating conditions that violate regulatory requirements, with potentially catastrophic consequences, *e.g.*, blackouts. In addition, the robust controller on its own is quite sensitive to the initial choice of $\hat{X}$, with a much larger standard deviation than Algorithm 1. Similar to the previous experiment, the model error shown in Figure 5c decreases monotonically with a nonzero final estimation error.

In this large uncertainty setting, we also experiment with partial observability of voltages (denoted "p.o." in Figure 5) by withholding observations of voltages from buses $i \in \{9, 19, 22, 31, 40, 46, 55\}$. That is, we still permit the robust controller Π to control the reactive power injection at these buses, but SEL does not use the voltages $v_i(t)$ from these buses as part of the constraints for the consistent sets. As expected, our controller performs worse in this partially-observed setting, but only marginally so. As shown in Figure 5c, the lack of voltage observation at these 7 nodes does not significantly affect the ability of SEL to learn consistent models.

Experiment 3: Moderate Initial Error, Large Uncertainty. Finally, we test our robust online voltage controller with a large uncertainty set ($\alpha = 1.0$) but moderate initial error ($\hat{X}_1$ projected into $\mathcal{X}_{0.5}$). We found the plots of voltage profiles to be nearly indistinguishable from the moderate uncertainty/moderate error setting (Figure 4), so we have omitted inclusion of those plots. This empirical observation matches the intuitive idea that for large uncertainty settings, the observed trajectory data is more informative in the definition of consistent sets than the initial uncertainty set $\mathcal{X}$. This explains why our method performs similarly in both the medium uncertainty/medium error and large uncertainty/large error settings.

6 CONCLUSION

This paper provides the first controller that establishes a finite-mistake guarantee for voltage control in a setting with an unknown grid topology. We showed that traditional voltage control algorithms that rely on knowing network information may cause grid instability when given incorrect information about the network topology; furthermore, decentralized network-agnostic control algorithms may also fail when subject to realistic noise and constraints on control inputs. In contrast, our proposed algorithm is able to learn the grid topology in an online fashion and provably converge to a stable controller. Further, simulated experiments on a 56-bus distribution grid demonstrate the effectiveness of our algorithm in a practical scenario.

As the current algorithm is highly centralized, future works may consider more decentralized approaches to topology-robust voltage control in order to enable faster real-time control. Ideas from works such as [35] can potentially be adapted. Further studies may also explore loosening the radial topology assumption to accommodate a wider range of distribution grids. This would be a challenging, but important, extension. Finally, an interesting algorithmic extension is to consider computationally efficient convex body chasing algorithms with better competitive ratios. Existing methods based on Steiner point [2, 5] achieve nearly-optimal competitive ratio but are computationally inefficient in high dimension settings such as voltage control, so designing efficient approximate Steiner point algorithms could potentially lead to significant performance improvements.

ACKNOWLEDGMENTS

We thank Dimitar Ho for useful conversations as well as the anonymous reviewers for their careful reading of this paper and insightful suggestions. This work was supported by NSF grants CNS-2146814, CPS-2136197, CNS-2106403, NGSDI-2105648, and ECCS-2200692.

REFERENCES

[1] 2005. *Global Survey of Regulatory Approaches for Power Quality and Reliability.* Technical Report 1008589. EPRI, Palo Alto, CA, USA. 146 pages. https://www.epri.com/research/products/1008589

[2] C. J. Argue, Sébastien Bubeck, Michael B. Cohen, Anupam Gupta, and Yin Tat Lee. 2019. A Nearly-Linear Bound for Chasing Nested Convex Bodies. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms* (San Diego, CA, USA) *(SODA '19)*. Society for Industrial and Applied Mathematics, USA, 117–122. https://doi.org/10.1137/1.9781611975482.8

[3] M.E. Baran and F.F. Wu. 1989. Network reconfiguration in distribution systems for loss reduction and load balancing. *IEEE Transactions on Power Delivery* 4, 2 (1989), 1401–1407. https://doi.org/10.1109/61.25627

[4] Saverio Bolognani, Ruggero Carli, Guido Cavraro, and Sandro Zampieri. 2015. Distributed Reactive Power Feedback Control for Voltage Regulation and Loss Minimization. *IEEE Trans. Automat. Control* 60, 4 (April 2015), 966–981. https://doi.org/10.1109/TAC.2014.2363931

[5] Sébastien Bubeck, Bo'az Klartag, Yin Tat Lee, Yuanzhi Li, and Mark Sellke. 2020. Chasing nested convex bodies nearly optimally. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms* (Salt Lake City, Utah). Society for Industrial and Applied Mathematics, USA, 1496–1508. https://doi.org/10.1137/1.9781611975994.91

[6] Yize Chen, Yuanyuan Shi, and Baosen Zhang. 2020. Data-Driven Optimal Voltage Regulation Using Input Convex Neural Networks. *Electric Power Systems Research* 189 (Dec. 2020), 106741. https://doi.org/10.1016/j.epsr.2020.106741

[7] Wenqi Cui, Jiayi Li, and Baosen Zhang. 2021. Decentralized Safe Reinforcement Learning for Voltage Control. *arXiv preprint arXiv:2110.01126* (Oct. 2021), 1–7. https://doi.org/10.48550/arXiv.2110.01126

[8] Deepjyoti Deka, Scott Backhaus, and Michael Chertkov. 2016. Estimating distribution grid topologies: A graphical learning based approach. In *2016 Power Systems Computation Conference (PSCC)*. IEEE, Genoa, Italy, 1–7. https://doi.org/10.1109/PSCC.2016.7541005

[9] Deepjyoti Deka, Scott Backhaus, and Michael Chertkov. 2018. Structure Learning in Power Distribution Networks. *IEEE Transactions on Control of Network Systems* 5, 3 (Sept. 2018), 1061–1074. https://doi.org/10.1109/TCNS.2017.2673546

[10] Steven Diamond and Stephen Boyd. 2016. CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research* 17, 83 (2016), 1–5. http://jmlr.org/papers/v17/15-408.html

[11] Jiajun Duan, Di Shi, Ruisheng Diao, Haifeng Li, Zhiwei Wang, Bei Zhang, Desong Bian, and Zhehan Yi. 2020. Deep-reinforcement-learning-based autonomous voltage control for power grid operations. *IEEE Transactions on Power Systems* 35, 1 (Jan. 2020), 814–817. https://doi.org/10.1109/TPWRS.2019.2941134

[12] Masoud Farivar, Lijun Chen, and Steven Low. 2013. Equilibrium and dynamics of local voltage control in distribution systems. In *Proceedings of the IEEE Conference on Decision and Control* (Firenze, Italy). IEEE, USA, 4329–4334. https://doi.org/10.1109/CDC.2013.6760555

[13] Masoud Farivar, Russell Neal, Christopher Clarke, and Steven Low. 2012. Optimal inverter VAR control in distribution systems with high PV penetration. In *2012 IEEE Power and Energy Society General Meeting* (San Diego, CA, USA). IEEE, USA, 1–7. https://doi.org/10.1109/PESGM.2012.6345736

[14] Chao Gao and Miles A Redfern. 2010. A review of voltage control techniques of networks with distributed generations using on-load tap changer transformers. In *45th International Universities Power Engineering Conference UPEC2010* (Cardiff, UK). IEEE, 1–6.

[15] Yuanqi Gao, Wei Wang, and Nanpeng Yu. 2021. Consensus Multi-Agent Reinforcement Learning for Volt-VAR Control in Power Distribution Networks. *IEEE Transactions on Smart Grid* 12, 4 (2021), 3594–3604. https://doi.org/10.1109/TSG.

2021.3058996

[16] Yifei Guo, Qiuwei Wu, Houlei Gao, Xinyu Chen, Jacob Østergaard, and Huanhai Xin. 2019. MPC-Based Coordinated Voltage Regulation for Distribution Networks With Distributed Generation and Energy Storage System. *IEEE Transactions on Sustainable Energy* 10, 4 (Oct. 2019), 1731–1739. https://doi.org/10.1109/TSTE.2018.2869932

[17] Hassan Haes Alhelou, Mohamad Esmail Hamedani-Golshan, Takawira Cuthbert Njenda, and Pierluigi Siano. 2019. A survey on power system blackout and cascading events: Research motivations and challenges. *Energies* 12, 4 (2019), 682. https://doi.org/10.3390/en12040682

[18] Dimitar Ho, Hoang Le, John Doyle, and Yisong Yue. 2021. Online Robust Control of Nonlinear Systems with Large Uncertainty. In *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics (AISTATS)* (San Diego, CA, USA), Vol. 130. PMLR, 3475–3483.

[19] Vassilis Kekatos, Georgios B. Giannakis, and Ross Baldick. 2014. Grid topology identification using electricity prices. In *2014 IEEE PES General Meeting | Conference & Exposition* (National Harbor, MD, USA). IEEE, 1–5. https://doi.org/10.1109/PESGM.2014.6939474

[20] Haoran Li, Yang Weng, Yizheng Liao, Brian Keel, and Kenneth E. Brown. 2019. Robust Hidden Topology Identification in Distribution Systems. *arXiv preprint arXiv:1902.01365* (Feb. 2019). https://doi.org/10.48550/arXiv.1902.01365

[21] Na Li, Guannan Qu, and Munther Dahleh. 2014. Real-time decentralized voltage control in distribution networks. In *2014 52nd Annual Allerton Conference on Communication, Control, and Computing, Allerton 2014* (Monticello, IL, USA). IEEE, 582–588. https://doi.org/10.1109/ALLERTON.2014.7028508

[22] Yizheng Liao, Yang Weng, Meng Wu, and Ram Rajagopal. 2015. Distribution grid topology reconstruction: An information theoretic approach. In *2015 North American Power Symposium (NAPS)* (Charlotte, NC, USA). IEEE, 1–6. https://doi.org/10.1109/NAPS.2015.7335248

[23] Salish Maharjan, Ashwin M. Khambadkone, and Jimmy Chih Hsien Peng. 2021. Robust constrained model predictive voltage control in active distribution networks. *IEEE Transactions on Sustainable Energy* 12, 1 (June 2021), 400–411. https://doi.org/10.1109/TSTE.2020.3001115

[24] Daniel K Molzahn, Florian Dörfler, Henrik Sandberg, Steven H Low, Sambuddha Chakrabarti, Ross Baldick, and Javad Lavaei. 2017. A survey of distributed optimization and control algorithms for electric power systems. *IEEE Transactions on Smart Grid* 8, 6 (2017), 2941–2962. https://doi.org/10.1109/TSG.2017.2720471

[25] Seiun Park, Deepjyoti Deka, and Michael Chcrtkov. 2018. Exact topology and parameter estimation in distribution grids with minimal observability. In *2018 Power Systems Computation Conference (PSCC)*. IEEE, 1–6. https://doi.org/10.23919/PSCC.2018.8442881

[26] Guannan Qu and Na Li. 2020. Optimal Distributed Feedback Voltage Control under Limited Reactive Power. *IEEE Transactions on Power Systems* 35, 1 (Jan. 2020), 315–331. https://doi.org/10.1109/TPWRS.2019.2931685

[27] Yann G. Rebours, Daniel S. Kirschen, Marc Trotignon, and Sbastien Rossignol. 2007. A Survey of Frequency and Voltage Control Ancillary Services—Part I: Technical Features. *IEEE Transactions on Power Systems* 22, 1 (2007), 350–357. https://doi.org/10.1109/TPWRS.2006.888963

[28] Tomonobu Senjyu, Yoshitaka Miyazato, Atsushi Yona, Naomitsu Urasaki, and Toshihisa Funabashi. 2008. Optimal Distribution Voltage Control and Coordination With Distributed Generation. *IEEE Transactions on Power Delivery* 23, 2 (2008), 1236–1242. https://doi.org/10.1109/TPWRD.2007.908816

[29] Y. Sharon, A. M. Annaswamy, A. L. Motto, and A. Chakraborty. 2012. Topology identification in distribution network with limited measurements. In *2012 IEEE PES Innovative Smart Grid Technologies (ISGT)*. 1–6. https://doi.org/10.1109/ISGT.2012.6175638

[30] Yuanyuan Shi, Guannan Qu, Steven Low, Anima Anandkumar, and Adam Wierman. 2021. Stability Constrained Reinforcement Learning for Real-Time Voltage Control. *arXiv preprint arXiv:2109.14854* (Sept. 2021). https://doi.org/10.48550/arXiv.2109.14854

[31] Xianzhuo Sun and Jing Qiu. 2021. Two-Stage Volt/Var Control in Active Distribution Networks With Multi-Agent Deep Reinforcement Learning Method. *IEEE Transactions on Smart Grid* 12, 4 (2021), 2903–2912. https://doi.org/10.1109/TSG.2021.3052998

[32] Zhiyuan Tang, David J. Hill, and Tao Liu. 2019. Fast Distributed Reactive Power Control for Voltage Regulation in Distribution Networks. *IEEE Transactions on Power Systems* 34, 1 (2019), 802–805. https://doi.org/10.1109/TPWRD.2018.2868158

[33] Shengyi Wang, Jiajun Duan, Di Shi, Chunlei Xu, Haifeng Li, Ruisheng Diao, and Zhiwei Wang. 2020. A Data-Driven Multi-Agent Autonomous Voltage Control Framework Using Deep Reinforcement Learning. *IEEE Transactions on Power Systems* 35, 6 (2020), 4644–4654. https://doi.org/10.1109/TPWRS.2020.2990179

[34] Hanchen Xu, Alejandro D. Domínguez-García, and Peter W. Sauer. 2020. Optimal Tap Setting of Voltage Regulation Transformers Using Batch Reinforcement Learning. *IEEE Transactions on Power Systems* 35, 3 (May 2020), 1990–2001. https://doi.org/10.1109/TPWRS.2019.2948132

[35] Jing Yu, Dimitar Ho, and Adam Wierman. 2022. Online Stabilization of Unknown Networked Systems with Communication Constraints. *arXiv preprint arXiv:2203.02630* (March 2022). https://doi.org/10.48550/arXiv.2203.02630

[36] Baosen Zhang, Albert YS Lam, Alejandro D Domínguez-García, and David Tse. 2014. An optimal and distributed method for voltage regulation in power distribution systems. *IEEE Transactions on Power Systems* 30, 4 (2014), 1714–1726. https://doi.org/10.1109/TPWRS.2014.2347281

[37] Hao Zhu and Hao Jan Liu. 2015. Fast local voltage control under limited reactive power: Optimality and stability analysis. *IEEE Transactions on Power Systems* 31, 5 (2015), 3794–3803. https://doi.org/10.1109/TPWRS.2015.2504419