

Stability Constrained Reinforcement Learning for Decentralized Real-Time Voltage Control

Jie Feng Student Member, IEEE, Yuanyuan Shi Member, IEEE, Guannan Qu Member, IEEE, Steven H. Low Fellow, IEEE, Anima Anandkumar Fellow, IEEE, and Adam Wierman Member, IEEE

Abstract—Deep reinforcement learning has been recognized as a promising tool to address the challenges in real-time control of power systems. However, its deployment in real-world power systems has been hindered by a lack of explicit stability and safety guarantees. In this paper, we propose a stability-constrained reinforcement learning (RL) method for real-time implementation of voltage control, that guarantees system stability both during policy learning and deployment of the learned policy. The key idea underlying our approach is an explicitly constructed Lyapunov function that leads to a sufficient structural condition for stabilizing policies, i.e., monotonically decreasing policies guarantee stability. We incorporate this structural constraint with RL, by parameterizing each local voltage controller using a monotone neural network, thus ensuring the stability constraint is satisfied by design. We demonstrate the effectiveness of our approach in both single-phase and three-phase IEEE test feeders, where the proposed method can reduce the transient control cost by more than 26.7% and shorten the voltage recovery time by 23.6% on average compared to the widely used linear policy, while always achieving voltage stability. In contrast, standard RL methods often fail to achieve voltage stability.

Index Terms—Voltage control, Reinforcement learning, Lyapunov stability

I. Introduction

The voltage control problem is one of the most critical problems in the control of power network. The primary purpose of voltage control is to maintain the voltage magnitude within an acceptable range under all possible working conditions. Due to the recent proliferation of distributed energy resources (DERs) such as solar and electric vehicles, voltage deviations are becoming increasingly complex and unpredictable. As a result, conventional voltage regulation methods based on on-load tap changing transformers, capacitor banks, and voltage regulators [1], [2] may fail to respond to the rapid and possibly large fluctuations. DERs can adjust the reactive power output based on the real-time voltage measurement to achieve fast and flexible voltage stabilization [3].

To coordinate the inverter-connected resources for real-time voltage control, one key challenge is to design control rules that can stabilize the system at scale with limited

information. Despite the progress, most of the existing work has only been able to optimize the steady-state cost, i.e. the cost of the operating point after the voltage converges (see [4]–[7] and references within). However, as the system is subject to more frequent load and generation fluctuations, optimizing the transient performance becomes of equal importance. Once a voltage violation happens, an important goal is to bring the voltage profile back to the safety region as soon as possible, at the minimum control costs.

Optimizing or even analyzing the transient cost of voltage control has long been challenging as this is a nonlinear control problem [8]. The challenge is further complicated by the fact that exact model of a distribution system is often unknown due to frequent system reconfigurations [9] and limited communication infrastructure. Reinforcement Learning (RL) has emerged as a promising tool to address this challenge. One intriguing benefit of RL methods is their model-free characteristic, which means no prior knowledge of the system models is required. Further, with the introduction of neural networks to RL, deep reinforcement learning has great expressive power and has shown impressive results in learning nonlinear controllers with good transient performance.

Despite the promising attempts, one difficulty in applying RL for voltage control is the lack of stability guarantee [10], [11]. Even if the learned policy may appear “stable” on a training data set, it is not guaranteed to be stable in unseen test cases and stability requires explicit characterization. Motivated by this challenge, the question we address in this paper is:

Can RL be applied for voltage control with a provable stability guarantee?

The key idea underlying our approach is that, with a judiciously chosen Lyapunov function, strict monotonicity of the policy is sufficient to guarantee voltage stability (Theorem 1). Given that monotonicity is a model-free constraint, it is practical to design a stabilizing RL controller without model knowledge. To enforce this structural constraint, we propose a decentralized controller (Stable-DDPG, Algorithm 1) which integrates the stability constraint with a popular RL framework deep deterministic policy gradient (DDPG) [12] through monotone policy network design. The proposed method enables us to leverage the power of deep RL to improve the transient performance of voltage control without knowing the underlying model

The authors are supported by NSF grant ECCS-2200692.

Jie Feng and Yuanyuan Shi are with the Department of Electrical and Computer Engineering, University of California San Diego jif005,yus047@ucsd.edu.

Guannan Qu is with the Department of Electrical and Computer Engineering, Carnegie Mellon University.

Steven H. Low, Anima Anandkumar, and Adam Wierman are with the Computing and Mathematical Sciences Department, Caltech.

parameters. We conduct extensive numerical case studies in both single-phase and three-phase IEEE test feeders to demonstrate the effectiveness of the proposed Stable-DDPG with both simulated voltage disturbances and real-world data. The trained Stable-DDPG can compute control actions efficiently (within 1 ms), which facilitates real-time implementation of neural network-based voltage controllers. This paper extends the result of our previous conference version [8] in the following aspects:

- We extend the stability analysis from continuous-time to discrete-time systems to better accommodate the discrete-time nature of inverter-based controllers.
- We construct a new discrete-time Lyapunov function and derive the structural constraints for stabilizing controller in Theorem 1. The discrete-time stability constraint requires the policy to be monotonically decreasing, and lower bounded by a value related to the sampling time. The clear relationship between stability and sampling time can assist the practical implementation of the proposed voltage controller with a finite sampling time. As the sampling time $\Delta T \rightarrow 0$, the stability condition reduces to the continuous-time stability condition.
- We test the proposed approach through extensive numerical studies in IEEE single-phase and three-phase systems with simulated and real-world data.

A. Related work

a) Steady-state cost optimization: Existing literature in optimizing the steady-state cost of voltage control can be roughly classified into two categories based on [6]: feed-forward optimization methods and feedback control methods. A typical example of feed-forward optimization methods is Optimal Power Flow (OPF) based methods [4], [13]–[16], where control actions are calculated by solving an optimization problem to minimize the power loss subject to voltage constraints. These algorithms assume knowledge of both the system models and the disturbance (e.g., load or renewable generations). Additionally, the computational cost of solving the OPF problem makes it difficult to respond to rapidly varying voltage profiles. On the other hand, feedback control methods do not assume to know the system model or the disturbance explicitly but take measurements of voltage magnitudes to decide the reactive power setpoints. In terms of time scale, the feedback controllers could work on a faster time scale as it does not require solving an optimization problem at each time step to decide the control actions. A popular feedback controller is the droop control, which is adopted by the IEEE 1547 standard [3]. However, as shown in [17], basic droop control can lead to instability if the controller gains are selected improperly. With more sophisticated structures, feedback controllers could achieve promising performance [6]. Regardless of the progress, optimizing the transient cost using feedback control methods remains challenging as the power flow equations are nonlinear, and the transient performance cost function can also be non-convex. The difficulty will be further exacerbated when the

controllers are required to be optimized in a decentralized manner.

b) Transient cost optimization: There has been tremendous interest in using RL for transient performance optimization in voltage control [18]–[30]. Given different communication conditions, existing RL methods fall into three categories, centralized, distributed, and decentralized controllers. Centralized controllers mean that the agent has access to global operating conditions, which leads to a powerful controller [21]. However, sophisticated communication networks are demanded and the agent has to deal with high-dimensional information. In the distributed setting, the network is first partitioned into small regions, and each region is assigned with an RL agent [19], [22], [26]. The agent has full observation of nodes located within the region. Decentralized controllers [8], [18], [31] are trained only with local measurements, and thus require no communication among peers, reducing the local computational burden. Please refer to a recent review [32] for more comprehensive overview. Despite the promise of RL for optimizing the transient performance, a widely-recognized issue is that RL lacks a provable stability guarantee, which is the main problem we want to tackle in this paper.

c) Lyapunov approaches in RL: Using Lyapunov functions in RL was first introduced by [33], where an agent learns to control the system by switching among the base controllers. These controllers are designed by using a specific Lyapunov function such that any switching policy is stable for the system. However, this work does not discuss how to find a candidate Lyapunov function in general, except for a case-by-case construction. A set of recent works including [34], [35] have attempted to address this challenge by jointly learning a Lyapunov function and a stabilizing policy. [34] uses linear programming to find the Lyapunov function, and [35] parameterizes the Lyapunov function as a neural network. To find a valid Lyapunov function and the corresponding controller, stability conditions are incorporated as a soft penalty during training and verified after training. In the context of these works, our contribution can be viewed as explicitly constructing a Lyapunov function for the voltage control problem to guide policy learning, rather than learning Lyapunov functions. Closest in spirit to our paper is [31], which proposes a stable RL approach for frequency control. However, their approach only applies to the frequency control problem, while our method works for voltage control which requires a different Lyapunov function design. Interestingly, both our work and prior work [31] arrive at a similar stability condition, that is strict policy monotonicity guarantees system stability.

II. Model & Preliminaries

In this section, we review distribution system power flow models for both single-phase and three-phase grids.

A. Branch Flow Model for Single-phase Grids

We consider the branch flow model [1] in a radial distribution network. Consider a distribution grid $\mathcal{N} = (\mathcal{P}_0, \mathcal{G})$, consisting of a set of $\mathcal{P}_0 = \{0, 1, \dots, n\}$ nodes and an edge set $\mathcal{G}$. In the graph, node 0 is known as the substation, and all the other nodes are buses that correspond to residential areas. We also use $\mathcal{P} = \mathcal{P}_0 / \{0\}$ to denote the set of nodes excluding the substation node. Each node $i \in \mathcal{P}$ is associated with an active power injection p_i and a reactive power injection q_i . Let V_i be the complex voltage and $v_i = \|V_i\|^2$ is the squared voltage magnitude. We use notation $\mathbf{p}, \mathbf{q}$ and $\mathbf{v}$ to denote the p_i, q_i, v_i stacked into a vector. $\mathbf{p}, \mathbf{q}$ and $\mathbf{v}$ satisfy the following equations, $\mathcal{D}_j \in \mathcal{P}, i = \text{parent}(j)$,

$$p_j = P_{ij} - r_{ij} l_{ij} - \sum_{k: (j,k) \in \mathcal{E}} P_{jk}, \quad (1a)$$

$$q_j = Q_{ij} - x_{ij} l_{ij} - \sum_{k: (j,k) \in \mathcal{E}} Q_{jk}, \quad (1b)$$

$$v_j = v_i - 2(r_{ij} P_{ij} + x_{ij} Q_{ij}) + (r_{ij}^2 + x_{ij}^2) l_{ij}, (i, j) \in \mathcal{G} \quad (1c)$$

where $l_{ij} = \frac{P_{ij}^2 + Q_{ij}^2}{v_i}$ is the squared current, P_{ij} and Q_{ij} represent the active power and reactive power flow on line (i, j) , and r_{ij} and x_{ij} are the line resistance and reactance. Equation (1a) and (1b) represent the real and reactive power conservation at node j , and (1c) represents the voltage drop from node i to node j .

Following [36], if the higher order power loss term can be ignored by setting $l_{ij} = 0$, we obtain the following linear approximation model,

$$p_j = -P_{ij} + \sum_{k: (j,k) \in \mathcal{T}} P_{jk}, \quad q_j = -Q_{ij} + \sum_{k: (j,k) \in \mathcal{T}} Q_{jk}, \quad (2a)$$

$$v_j = v_i - 2(r_{ij} P_{ij} + x_{ij} Q_{ij}), (i, j) \in \mathcal{E} \quad (2b)$$

We can rearrange the above equations into the vector form,

$$\mathbf{v} = R\mathbf{p} + X\mathbf{q} + v_0 \mathbf{1} = X\mathbf{q} + \mathbf{v}^{env}. \quad (3)$$

where matrix $R = [R_{ij}]_{n \times n}, X = [X_{ij}]_{n \times n}$ are given as follows, $R_{ij} := 2 \sum_{(h,k) \in \mathcal{P}_i \cap \mathcal{P}_j} r_{hk}, X_{ij} := 2 \sum_{(h,k) \in \mathcal{P}_i \cap \mathcal{P}_j} x_{hk}$ where $\mathcal{X}_i \prec E$ is the set of lines on the unique path from bus 0 to bus i . Here we follow [17] to separate the voltage magnitude $\mathbf{v}$ into two parts: the controllable part $X\mathbf{q}$ that can be adjusted via adjusting reactive power injection $\mathbf{q}$ through the inverter-based control devices, and the non-controllable part $\mathbf{v}^{env} = R\mathbf{p} + v_0 \mathbf{1}$ that is decided by the load and PV active power $\mathbf{p}$. Matrix X and R satisfy the following property, which is crucial for the stable control design.

Proposition 1 ([17] Lemma 1). Suppose $x_{ij}, r_{ij} > 0$ for all (i, j) . Then, X and R are positive definite matrices.

B. Multi-phase Grid Modeling

We now introduce an abridged version of the branch flow model in three-phase distribution systems. For simplicity, it is first assumed that all buses are served by all three

phases, so we can use 3-dimensional vectors to represent system variables. With slight abuse of notation, $\mathbf{P}_{ij}$ is a 3-dimensional vector such that $\mathbf{P}_{ij} = [P_{ij}^a, P_{ij}^b, P_{ij}^c]^\top$. $\mathbf{S}_{ij}, \mathbf{Q}_{ij}$ are defined in the same way. The vectors of power injections and complex voltages are denoted by $\mathbf{s}_i$ and $\mathbf{v}_i$, respectively. $\mathbf{Z}_{ij} \in \mathbb{S}^3$ is the phase impedance matrix for line (i, j) , where $\mathbf{Z}_{ij} = \mathbf{R}_{ij} + j\mathbf{X}_{ij}$.

We further assume that the phase voltages of arbitrary bus i are approximately balanced with absolute value $\hat{v}_i$, then $\mathbf{v}_i$ can be estimated by $\hat{v}_i \alpha$, where $\alpha = [1 \ a \ a^2]^\top, a = e^{-j2\pi/3}$. Define $\hat{\mathbf{Z}}_{ij} = \text{diag}(\alpha^H) \mathbf{Z}_{ij} \text{diag}(\alpha)$, following (2), the linear approximate three-phase model is,

$$\mathbf{s}_j = \mathbf{S}_{ij} + \sum_{k: (j,k) \in \mathcal{E}} \mathbf{S}_{jk} \quad (4a)$$

$$\mathbf{v}_i - \mathbf{v}_j = 2Re[\hat{\mathbf{Z}}_{ij}^H \mathbf{S}_{ij}] \quad (4b)$$

Notice that the vector variables can be arranged either by bus or by phase. For example, the voltage magnitude can be rearranged by phase as $\tilde{\mathbf{v}} = [\tilde{\mathbf{v}}_a, \tilde{\mathbf{v}}_b, \tilde{\mathbf{v}}_c]$, where $\tilde{\mathbf{v}}_a = [v_1^a, v_2^a, \dots, v_n^a]$, $\tilde{\mathbf{v}}_b$ and $\tilde{\mathbf{v}}_c$ share the similar definition. Recall that $\mathbf{v} = [\mathbf{v}_i, i \in \mathcal{P}]$, which is ordered by bus. With a permutation matrix T_v , the transformation between two formats can be represented by $\mathbf{v} = T_v \tilde{\mathbf{v}}$. The three-phase branch flow model can then be arranged to a compact form, the same as the single-phase model,

$$\mathbf{v} = R\mathbf{p} + X\mathbf{q} + v_0 \mathbf{1}_{3N} = X\mathbf{q} + \mathbf{v}^{env}. \quad (5)$$

For a detailed mathematical derivation, please refer to [37]. Notice that the single-phase and three-phase system dynamics share the same linear approximation model (3), $\mathbf{v} = X\mathbf{q} + \mathbf{v}^{env}$, allowing us to derive the Lyapunov equation and stability conditions based on the same analytical model.

Assumption 1. Assume every matrix $\mathbf{X}_{ij} = \frac{1}{2} \begin{bmatrix} 2x_{ij}^{aa} & x_{ij}^{ab} & x_{ij}^{ac} \\ x_{ij}^{ab} & 2x_{ij}^{bb} & x_{ij}^{bc} \\ x_{ij}^{ac} & x_{ij}^{bc} & 2x_{ij}^{cc} \end{bmatrix}$ is strictly diagonally dominant with positive diagonal entries for all edges $(i, j) \in \mathcal{G}$.

It is denoted in [37], due to the structure of distribution lines, the diagonal dominance conditions in Assumption 1 are generally satisfied for multi-phase grids. According to Corollary 1 in [37], if Assumption 1 holds, X is positive definite for three-phase distribution system.

III. Voltage Control Problem Formulation

The voltage control problem can be modeled as a control problem in a quasi-dynamical system with state $\mathbf{v}$ and controller $\mathbf{q}$. Given the current voltage measurement $\mathbf{v}(t)$ and other available information, the controller determines a new reactive power injection $\mathbf{q}(t+1)$. The new $\mathbf{q}(t+1)$ will result in a new voltage profile $\mathbf{v}(t+1)$. We envision that the reactive power loop is embedded in an inverter control loop and operates at very fast timescales [21], and denote the change rate of reactive power injection as, $\dot{q}_i(t) := u_i(t)$. Using the zero-order hold on the inputs

and a sample time of ΔT , we get the closed-loop voltage control dynamics as,

$$\mathbf{v}(t+1) = X\mathbf{q}(t+1) + \mathbf{v}^{env}, \quad (6a)$$

$$q_i(t+1) = q_i(t) + \Delta T \pm u_i(v_i(t)), \mathcal{D} / \mathcal{P} \quad (6b)$$

where $\mathbf{u} = (u_1, \dots, u_n)$ is the decentralized voltage controller. Note that (6b) represents the class of incremental voltage controller. As shown in [17], a decentralized controller only depends on the current time step information, i.e., $q_i(t) = u_i(v_i(t))$ is not possible to stabilize the voltage $\mathbf{v}$ under arbitrary disturbance, while the incremental voltage controller guarantees the existence of stabilizing controllers. This motivates our focus on incremental voltage controllers.

A. Voltage Stability

Voltage stability is defined as the ability of the system voltage trajectory to return to an acceptable range after arbitrary disturbance. See Definition 1 below.

Definition 1 (Voltage stability). The closed loop system is stable if for any $\mathbf{v}^{env}$ and $\mathbf{v}(0)$, $\mathbf{v}(t)$ converges to the set $S_v = \{\mathbf{v} / \mathbb{R}^n : \underline{v}_i \leq v_i \leq \bar{v}_i\}$ in the sense that $\lim_{t \rightarrow \infty} \text{dist}(\mathbf{v}(t), S_v) = 0$ and the distance is defined as $\text{dist}(\mathbf{v}(t), S_v) = \min_{\mathbf{v}' \in S_v} \|\mathbf{v}(t) - \mathbf{v}'\|$.




Fig. 1: Voltage stability of bus i.

With high penetration of DERs, rapid changes of load and renewable generation often happen in a fast time scale, thus it is important to ensure the designed controller meets the stability condition. With the requirement for voltage stability, the optimal voltage control problem can be formulated as,

$$\min_{\theta} J(\theta) = \sum_{t=0}^T \gamma^t \sum_{i=1}^n c_i(v_i(t), u_i(t)) \quad (7a)$$

$$\text{s.t. } \mathbf{v}(t+1) = X\mathbf{q}(t+1) + \mathbf{v}^{env}, \quad (7b)$$

$$q_i(t+1) = q_i(t) + \Delta T \pm u_i(t) \quad (7c)$$

$$u_i(t) = g_{\theta_i}(v_i(t)) \quad (7d)$$

$$\text{Voltage stability holds.} \quad (7e)$$

The goal of the voltage control problem is to reduce the total cost (7a) for time steps t from 0 to T , which consists of two parts: the cost of voltage deviation and the cost of control actions. One can choose different cost functions (e.g., one-norm, two-norm, or infinity-norm), depending on the system performance metrics and control devices. Our stability-constrained RL framework can accommodate




Fig. 2: Control System Architecture.

different types of cost functions mentioned above. In particular, in our experiment, we use $c_i(v_i(t), u_i(t)) = \eta_1 \max(v_i(t) - \bar{v}_i, 0) + \min(v_i(t) - \underline{v}_i, 0)^2 + \eta_2 |u_i(t)|$. Here η_1, η_2 are coefficients that balance the cost of action with respect to the voltage deviation. Voltage dynamics are represented by equations (7b)-(7c). (7d) specifies the decentralized policy structure $u_i(t) = g_{\theta_i}(v_i(t))$ only depends on local voltage measurement $v_i(t)$. Here θ_i is the policy parameter for the local policy at node i , and $\theta = (\theta_i)_{i \in \mathcal{N}}$ is the collection of the local policy parameters.

Transient cost vs. stationary cost. Our problem formulation in (7) is different from some of those in the literature, e.g., [5], [6], [38], [39], in the sense that the existing works typically consider the cost in steady-state, meaning the cost is evaluated at the fixed point or stationary point of the system. In contrast, our work considers the transient cost after a voltage disturbance, which is also an important metric for the performance of voltage control. An important future direction is to unify these two perspectives and design policies that can optimize both the transient and stationary costs.

B. Solving Voltage Control Problem via RL

In order to solve the optimal voltage control problem in (7), one needs the exact system dynamics, i.e., X . However, for distribution systems, the exact network parameters are often unknown or hard to estimate in real systems [7]. RL provides a powerful paradigm for solving (7), by training a policy that maps the state to action via interacting with the environment, so as to minimize the loss function defined as (7a). There are many RL algorithms to solve the policy minimization problem (7), and in this paper, we focus on the class of RL algorithms called policy optimization. We define the state space of each local controller as the nodal voltage deviation, represented by $v_i / \mathbb{R}$ (single-phase) or $v_i / \mathbb{R}^3$ (three-phase). The action space is defined as the range of potential reactive power changes, represented by $u_i / \mathbb{R}$ (single-phase) or $u_i / \mathbb{R}^3$ (three-phase).

Generally speaking, we parameterize each of the controllers, i.e., $u_i(t) = g_{\theta_i}(v_i(t))$ as a neural network with weights θ_i . The procedure is to run gradient methods on the policy parameter θ_i with learning rate α_i , $\theta_i \rightarrow \theta_i - \alpha_i J(\theta_i)$. As we are dealing with deterministic policies and continuous state space, one of the most popular choices is the Deep Deterministic Policy Gradient (DDPG) [12], where the policy gradient $J(\theta_i)$ is approximated by

$$\frac{1}{N} \sum_{j \in B} u_i \hat{Q}_{\phi_i}(v_i, u_i) \big|_{v_i=v_i[j], u_i=-g_{\theta_i}(v_i[j])} - \theta_i g_{\theta_i}(v_i) \big|_{v_i[j]}, \quad (8)$$

where $g_{\theta_i}(v_i)$ is the actor network, and $\{v_i[j], u_i[j]\}_{j \in B}$ are a batch of samples with batch size $\|B\| = N$ sampled from the replay buffer which stores historical state-action transition tuples of bus i . Here $\hat{Q}_{\phi_i}(v_i, u_i)$ is the value network (a.k.a critic network) that can be learned via temporal difference learning,

$$\min_{\phi_i} L(\phi_i) = E_{(v_i, u_i, c_i, v'_i)} [Q_{\phi_i}(v_i, u_i) - (c_i + \gamma Q_{\phi_i}(v'_i, g_{\theta_i}(v'_i)))] \quad (9)$$

where v'_i is system voltage after taking action u_i and realization of v_i^{env} . For more details of DDPG, readers may refer to [12].

In standard DDPG, stability is not an explicit requirement, it plays the role of implicit regularization since instability leads to high costs. However, the lack of an explicit stability requirement can lead to several issues. During the training phase, the policy may become unstable, causing the training process to terminate. Even after a policy is trained, there is no formal guarantee that the closed loop system is stable, which hinders the learned policy's deployment in real-world power systems where there is a very strong emphasis on stability. Next, we will introduce our framework that guarantees stability in policy learning.

IV. Main Results

We now introduce our stability-constrained RL framework for voltage control. We demonstrate that the voltage stability constraint can be translated into a monotonicity constraint on the policy, that can be satisfied by a careful design of monotone neural networks.

A. Voltage Stability Condition

In order to explicitly constrain stability for RL, we constrain the search space of policy in a subset of stabilizing controllers from Lyapunov stability theory. In particular, we use a generalization of Lyapunov's direct method, known as LaSalle's Invariance theorem for deriving the stability condition.

Proposition 2 (LaSalle's theorem for discrete-time system [40]). For dynamical system $x(t+1) = f(x(t))$, suppose $V : \mathbb{R}^n \rightarrow \mathbb{R}$ is a continuously differentiable function satisfying $V(x) < 0$ and $V(f(x)) \leq V(x)$ for $x \in \mathbb{R}^n$. Let E be the set of all points in $\mathbb{R}^n$ where $V(f(x)) = V(x) = 0$, and let M be the largest invariant

set in E . If there exists $a \in \mathbb{R}^+$ such that the level set $L_a := \{x : V(x) \geq a\}$ is bounded, then for any $x(0) \in L_a$ we have $\text{dist}(x(t), M) \rightarrow 0$ as $t \rightarrow \infty$. Further, if V is radially unbounded, i.e. $V(x) \rightarrow \infty$ as $\|x\| \rightarrow \infty$, then, for any $x(0) \in \mathbb{R}^n$, we have $\text{dist}(x(t), M) \rightarrow 0$ as $t \rightarrow \infty$.

The key to ensure stability is to find a controller $\mathbf{u} = g_{\theta}(\mathbf{v})$ and a Lyapunov function V , such that the stability conditions in Proposition 2 can be satisfied. For the voltage control problem defined by (7), $\mathbf{v}(t+1) = \mathbf{v}(t) + I_{\Delta T} X \mathbf{u}(t)$, where $I_{\Delta T}$ is a diagonal matrix with diagonal entries equal to ΔT . Since the control input $\mathbf{u} = g_{\theta}(\mathbf{v})$ depends on state $\mathbf{v}$, the closed-loop system dynamics can be written as $\mathbf{v}(t+1) = \mathbf{v}(t) + I_{\Delta T} X g_{\theta}(\mathbf{v}) := f_u(\mathbf{v}(t))$. We consider the following Lyapunov function,

$$V(\mathbf{v}) = (\mathbf{v} - f_u(\mathbf{v}))^T X^{-1} (\mathbf{v} - f_u(\mathbf{v})) \quad (10)$$

where X is the network reactance matrix defined in (3), that is a positive definite matrix for both single-phase and three-phase distribution grids. V is positive definite and is radially unbounded if $\|\mathbf{v}\| \rightarrow \infty$ as $\|\mathbf{v}\| \rightarrow \infty$. From LaSalle's theorem in Proposition 2, if $V(f_u(\mathbf{v})) \leq V(\mathbf{v})$ and $V(f_u(\mathbf{v})) = V(\mathbf{v}) = 0$ only when $\mathbf{v} \in S_v$, where S_v is the voltage safety set defined as $S_v = \{\mathbf{v} \in \mathbb{R}^n : v_i \geq \bar{v}_i\}$. Then we have for every initial voltage profile $\mathbf{v}(0) \in \mathbb{R}^n$, $\mathbf{v}(t)$ converges to the largest invariant set in S_v . Furthermore, suppose for all i , the control action satisfies $u_i = 0$ for $v_i \in [\underline{v}_i, \bar{v}_i]$. Then S_v itself is an invariant set.

The key question now reduces to how can we design the controller $\mathbf{u} = g_{\theta}(\mathbf{v})$ such that the closed-loop system satisfying these two properties:

- 1) $V(f_u(\mathbf{v})) \leq V(\mathbf{v})$ for $\mathbf{v} \in S_v$
- 2) $V(f_u(\mathbf{v})) = V(\mathbf{v}) = 0$ for $\mathbf{v} \in S_v$

Theorem 1 presents a sufficient structural condition for the above properties to hold, thus guaranteeing voltage stability.

Theorem 1 (Voltage stability condition). Suppose for all bus i , $g_{\theta_i}(\cdot)$ is a continuously differentiable function satisfying $u_i = g_{\theta_i}(v_i) = 0$ for $v_i \in [\underline{v}_i, \bar{v}_i]$. Further, each $\frac{\partial g_{\theta_i}}{\partial v_i}$ satisfies equation (11) on $(\underline{v}_i, \bar{v}_i]$ and $[\bar{v}_i, \infty)$

$$\frac{2}{\Delta T} X^{-1} \leftarrow \frac{\partial \mathbf{u}}{\partial \mathbf{v}} \leftarrow 0 \quad (11)$$

and $\lim_{|v_i| \rightarrow \infty} \|g_{\theta_i}(v_i)\| = \infty$. Then, the voltage stability defined in Definition 1 holds.

Equation (11) shows that when the sampling time $\Delta T \rightarrow 0$, the stability condition reduces to the continuous time stability condition $\frac{\partial \mathbf{u}}{\partial \mathbf{v}} \leftarrow 0$ as first shown in [8]. As the length of the sampling time increases, $\frac{\partial \mathbf{u}}{\partial \mathbf{v}}$ also needs to be lower bounded by $\frac{2}{\Delta T} X^{-1}$ and upper bounded by 0. As the typical sampling frequency of real-world inverters is in kHz scale [41], the left-hand side of (11) is naturally satisfied in most cases. Therefore, we focus on the right-

Algorithm 1 Stable-DDPG Learning Process

Initialization: Initial Q network $Q_{\phi_i}(v_i, u_i)$, monotone policy network $g_{\theta_i}(v_i)$ with parameters ϕ_i, θ_i ; empty replay buffers $\mathcal{E}_i$ for all buses $i = 1, \dots, n$.

- 1: for $j = 0$ to N_{ep} do
- 2: Randomly generate initial states $\mathbf{v}(0)$ with voltage violation for all nodes
- 3: for $t = 0$ to N_{step} do
- 4: Observe state $v_i(t)$, compute action based on current policy $u_i = g_{\theta_i}(v_i(t))$, $\mathcal{D}$
- 5: Execute the joint action $\mathbf{q}(t+1) = \mathbf{q}(t) + \Delta T \mathbf{u}(t)$, transit to next state $\mathbf{v}(t+1)$
- 6: Store $(v_i(t), u_i(t), c_i(t), v_i(t+1))$ in $\mathcal{E}_i$, $\mathcal{D}$
- 7: end for
- 8: if $\text{len}(\mathcal{E}_i) > \text{batch size}$ then
- 9: for $i = 1, \dots, n$ do
- 10: Randomly sample N state-transition data pairs from replay buffer D_i , $B_i = \{(v_i, u_i, c_i, v'_i)\}_{i=1}^N$.
- 11: Update the policy network by Eq (8)
- 12: Update the Q-function network Eq (9)
- 13: end for
- 14: end if
- 15: end for

hand side condition, $\frac{\partial \mathbf{u}}{\partial \mathbf{v}} \leftarrow 0$. Because of the decentralized characteristic, $\frac{\partial \mathbf{u}}{\partial \mathbf{v}}$ is a diagonal matrix.

$$\frac{\partial \mathbf{u}}{\partial \mathbf{v}} = \begin{bmatrix} \frac{\partial g_{\theta_1}}{\partial v_1} & \text{+++} & 0 \\ & \ddots & \\ 0 & \text{+++} & \frac{\partial g_{\theta_n}}{\partial v_n} \end{bmatrix} \quad (12)$$

Thus, if each g_{θ_i} is strictly monotonically increasing, i.e., $\frac{\partial g_{\theta_i}(v_i)}{\partial v_i} > 0$, $\mathcal{D}$, the voltage stability condition $\frac{\partial \mathbf{u}}{\partial \mathbf{v}} \leftarrow 0$ will be met. We note that a similar stability condition for the discrete-time voltage dynamics has been shown in [42], while our condition ensures globally asymptotic stability rather than the local stability guarantee in [42].

B. Stability-Constrained RL Algorithm

Combining the structural constraints for stabilizing controllers in Theorem 1 and the DDPG algorithm for solving voltage control in Section III-B, we now present the design of Stable-DDPG algorithm. The proposed stability-constrained policy learning algorithm is summarized in Algorithm 1.

As we notice in Algorithm 1, the general algorithm flow of Stable-DDPG is the same as DDPG, and the only difference is in the policy network parameterization. Since Theorem 1 restricts the class of stabilizing decentralized controllers to be strictly monotonically decreasing. Thus, we need to incorporate this structural condition into the policy design. Essentially, any monotone functions can be used for parameterizing the policy function, e.g., a linear policy $u_i = k_i v_i$, $\mathcal{D}v_i < \underline{v}_i$ or $v_i > \bar{v}_i$ with k_i is positive; and $u_i = 0$ for $\underline{v}_i \geq v_i \geq \bar{v}_i$. To leverage the superior expressiveness of neural networks, we represent




Fig. 3: Three-phase Monotone Voltage Controller, where each policy is parameterized as a Stacked ReLU Monotone Network.

$u_i = g_{\theta_i}(v_i)$ as a monotone neural network. There are several existing designs for the monotone neural networks in literature, e.g. [31], [43], [44]. In this paper, we follow the monotonic neural network design in [31, Lemma 3], which guarantees universal approximation of all single-input-single-output monotonic increasing functions [31, Theorem 2]. This design uses a single hidden layer neural network with d hidden units and ReLU activation, which is defined below.

Corollary 1. (Stacked ReLU Monotone Network [31, Lemma 3]) The stacked ReLU function constructed by Eq (13) is monotonic increasing for $x > 0$ and zero when $x \geq 0$.

$$\xi^+(x; w^+, b^+) = (w^+)^{\mathcal{E}} \text{ReLU}(\mathbf{1}x + b^+) \quad (13a)$$

$$\sum_{l=1}^{d'} w_l^+ > 0, \forall d^{\mathcal{E}} = 1, \dots, d, b_1^+ = 0, b_l^+ \leq b_{l-1}^+, \forall l = 2, \dots, d \quad (13b)$$

The stacked ReLU function constructed by Eq (14) is monotonic increasing for $x < 0$ and zero when $x \leq 0$.

$$\xi^-(x; w^-, b^-) = (w^-)^{\mathcal{E}} \text{ReLU}(-\mathbf{1}x + b^-) \quad (14a)$$

$$\sum_{l=1}^{d'} w_l^- < 0, \forall d^{\mathcal{E}} = 1, \dots, d, b_1^- = 0, b_l^- \leq b_{l-1}^-, \forall l = 2, \dots, d \quad (14b)$$

1) Single-phase Monotone Voltage Controller: Following the stability constraint (11), we set the single-phase voltage controller to be monotonically increasing with voltage controller. To incorporate the dead-band within range $v_i \in [\underline{v}_i, \bar{v}_i]$, we parameterize the controller at bus i as $g_{\theta_i}(v_i) = [\xi_{\theta_i}^+(v_i - \bar{v}_i) + \xi_{\theta_i}^-(v_i - \underline{v}_i)]$, where $\xi_{\theta_i}^+(v_i - \bar{v}_i) : \mathbb{R} \rightarrow \mathbb{R}$ is monotonically increasing for $v_i > \bar{v}_i$ and zero when $v_i \leq \bar{v}_i$, and $\xi_{\theta_i}^-(v_i - \underline{v}_i) : \mathbb{R} \rightarrow \mathbb{R}$ is monotonically increasing for $v_i < \underline{v}_i$ and zero otherwise. Because $u_i(t) = g_{\theta_i}(v_i(t))$, $\frac{\partial \mathbf{u}}{\partial \mathbf{v}} \leftarrow 0$ is satisfied.

2) Three-phase Monotone Voltage Controller: For the three-phase voltage controller, we generalize the single-input-single-output monotone policy network to three-dimensional input and output by deploying a single-phase controller for each phase. As demonstrated in Figure 3, we disentangle the AC voltage observations per phase and treat each phase as a single-phase input. In this way, $\frac{\partial u_i}{\partial v_i} =$

$\begin{bmatrix} \frac{\partial u_i^a}{\partial v_i^a} & 0 & 0 \\ 0 & \frac{\partial u_i^b}{\partial v_i^b} & 0 \\ 0 & 0 & \frac{\partial u_i^c}{\partial v_i^c} \end{bmatrix}$ is a diagonal matrix with negative entries, thus the stability condition $\frac{\partial \mathbf{u}}{\partial \mathbf{v}} \leftarrow 0$ is satisfied.

We conclude this section with two remarks.

Remark 1. The simplified Distflow models in (3) and (5) are introduced for theoretic analysis only. The original nonlinear dynamics are deployed in numerical experiments, and our proposed method stabilizes the systems in all test scenarios.

Remark 2. The stability criteria defined by Theorem 1 is a sufficient condition for asymptotic stability, which does not give any explicit guarantee to stabilize the systems in finite steps. To achieve exponential stability, the Lyapunov condition should be strengthened as $V(\mathbf{v}_{t+1}) - V(\mathbf{v}_t) \geq cV(\mathbf{v}_t)$, $0 < c < 1$ [40]. In this case, the stability condition is given by $\frac{1+\sqrt{1-c}}{\Delta T} X^{-1} < \frac{\partial \mathbf{u}}{\partial \mathbf{v}} < \frac{1-\sqrt{1-c}}{\Delta T} X^{-1}$. Given that ΔT is small, we can often find a constant c such that the trained policy satisfies the above inequality. As a result, the system can be input-to-state stable [45] with the proposed controller.

V. Case Study

We demonstrate the effectiveness of the proposed Stable-DDPG approach (Algorithm 1) on both single-phase and three-phase IEEE distribution test systems. Source code and data are available at <https://github.com/JieFeng-cse/Stable-DDPG-for-voltage-control>.

A. Experimental Setup

For single-phase test feeders, we use the IEEE 13-bus feeder and IEEE 123-bus feeder as the test cases, which are modified from three-phase models in [46]. For three-phase test feeders, we test on the unbalanced IEEE 13-bus feeder and IEEE 123-bus feeder [46]. Simulations for single-phase systems are implemented with pandapower [47], and the three-phase system is obtained by OpenDSS. We simulate different voltage disturbance scenarios: 1) High voltages: the PV generators are generating a large amount of power, this corresponds to the daytime scenario in California where there is abundant sunshine that can result in high voltage issues. 2) Low voltages: the system is serving heavy loads without PV generation. It corresponds to late afternoon or night when there is low/no solar generation but still a significant load. For each scenario, we randomly vary the active power injections thus obtaining different degrees of voltage violations, i.e., 5% to 15% of the nominal value. We set $\Delta T = 1s$ for all numerical experiments and verify the stability condition for all the trained policies. All experiments are conducted with an AMD 5800X CPU and an Nvidia 1080Ti GPU.

Baselines: We test the proposed stable-DDPG approach (Algorithm 1), against the following baseline algorithms. Details about the algorithm implementations are provided in Appendix B.




Fig. 4: Schematic diagram of the IEEE 13-bus system and IEEE 123-bus system.

TABLE I: Performance of linear, DDPG, and Stable-DDPG on 500 voltage violation scenarios for IEEE 13-bus system.

Method	Voltage recovery steps		Reactive power (Mvar)	
	Mean	Std	Mean	Std
MPC	4.55	8.90	7.62	16.40
Linear	5.31	3.19	8.22	10.72
Stable-DDPG	4.47	2.43	6.75	8.08
DDPG	6.61	20.67	30.20	120.24
DDPG*	2.31	1.18	3.65	3.21

Note: DDPG* denotes the performance of the DDPG policy in the subset of testing cases when it was able to stabilize the voltage.

a) Linear policy with deadband: $u_i(v_i) = \epsilon_i([v_i \bar{v}_i]^+ [v_i \ v_i]^+)$ (where $[x]^+ = \max(x, 0)$), and the new reactive power injection is $q_i(t) = q_i(t-1) + \Delta T u_i(v_i)$. This linear controller has been widely used in the power system control community [17]. With each $0 < \epsilon_i \geq \frac{2\sigma_{\min}(X)}{\sigma_{\max}^2(X)}$, linear policy guarantees stability but may lead to suboptimal control cost. We optimize the linear controller in an RL framework, where ϵ_i is the learnable parameter to obtain the best-performing linear policy for comparison.

b) Standard DDPG algorithm: is suggested for voltage control in [23], [26]. Standard DDPG minimizes the control cost without an explicit stability guarantee.

c) DDPG*: We denote the subset of results where the standard DDPG policy is able to maintain voltage stability as DDPG*.

B. Single-phase Simulation Results

1) 13-bus system: The nominal voltage magnitude at each bus except substation is 4.16 kV. The safe operation range S_v is defined as $\leq 5\%$ of the nominal value, that is $[3.952\text{kV}, 4.368\text{kV}]$. The overall training time for the Stable-DDPG algorithm is 71s, and the DDPG algorithm can be trained around 85s. Stable-DDPG can quickly learn to stabilize the system with relatively low cost. It also has smaller variance compared to standard DDPG.

Model Predict Control. We assume perfect knowledge of matrix X for the IEEE 13-bus system. Considering a finite look-ahead time window $H = 30$, which equals the episodic length of the RL training, the centralized Model Predict Control (MPC) algorithm can be formulated as follows.

Method	MPC	Linear	Stable-DDPG	DDPG
Time (ms)	449.83	0.16	0.37	0.17

TABLE II: Computational time comparison.

$$\begin{aligned} \underset{\hat{\mathbf{u}}(t), \hat{\mathbf{v}}(t), \hat{\mathbf{q}}(t)}{\operatorname{argmin}} \quad & \sum_{k=0}^H \sum_{i=1}^n c_i(\hat{v}_i(t+k), \hat{u}_i(t+k)) \quad (15a) \\ \text{subject to} \quad & \hat{\mathbf{q}}(t+k+1) = \hat{\mathbf{q}}(t+k) + \Delta T \hat{\mathbf{u}}(t+k), \quad (15b) \\ & \hat{\mathbf{v}}(t+k+1) = \mathbf{X} \hat{\mathbf{q}}(t+k+1) + \mathbf{v}^{env}, \quad (15c) \\ & \hat{\mathbf{v}}(t) = \mathbf{v}(t), \hat{\mathbf{q}}(t) = \mathbf{q}(t) \quad (15d) \\ & \underline{v} \leq \hat{\mathbf{v}}(t+H) \leq \bar{v} \quad (15e) \end{aligned}$$

For a fair comparison, the cost function $c_i(\hat{v}_i(t+k), \hat{u}_i(t+k))$ of MPC is chosen to be the same as the cost function used in RL training. At each time step, the finite-horizon optimal control problem (15) is solved to obtain the control sequence. We write $\hat{\mathbf{u}}^*(t), \hat{\mathbf{v}}^*(t), \hat{\mathbf{q}}^*(t)$ as the optimal control sequence and the corresponding voltage trajectory. The control action is later selected as $\mathbf{u}(t) = \hat{\mathbf{u}}^*(t)$. We use this centralized MPC algorithm as a baseline for the IEEE 13-bus system.

Control Performance. We compare the performance of the proposed Stable-DDPG method against linear policy, standard DDPG, and MPC policies on 500 different voltage violation scenarios. Table I shows the results. Notably, Stable-DDPG outperforms the centralized MPC algorithm even when the exact linearized system dynamics model is known for the MPC. In this case, the linearized model provides a reasonable approximation with some approximation error. As a result, our proposed Stable-DDPG algorithm, which interacts with the nonlinear power flow simulator for policy training can outperform the centralized MPC method. It is also worth mentioning that the computational time of the proposed Stable-DDPG (0.37ms) is on the same scale as the Linear controller (0.16ms) while significantly smaller than the MPC (449.83ms) as shown in Table II. Stable-DDPG can support a control frequency of up to 2 kHz, which enables real-time decentralized voltage control.

Figure 5 demonstrates the percentage of voltage instability cases in the 500 testing scenarios. If the controller is able to bring back the voltage of all controlled buses to [3.952kV, 4.368kV], the trajectory will be marked as “stable”. Otherwise, we record the final voltage magnitudes of the controlled buses and categorized based on the violation magnitude. Our method achieves voltage stability in all scenarios, whereas DDPG may lead to voltage instability even in this simple setting, with the final voltage beyond the $\leq 5\%$ range for about 4% of the test scenarios.

Test with Real-world Data. Finally, we test the proposed method using real-world data from DOE [6]. We compare the voltage dynamics without voltage control and when Stable-DDPG is used. We simulate a massive solar penetration scenario where all buses are associated with




Fig. 5: Voltage stability for single-phase 13 bus test system. The left plot is the voltage violation for each bus, the right plot is the largest violation bus.

PV and voltage controllers. The voltage control results are given in Figure 6. There are severe voltage violations without control, due to the high volatility in load and PV generation. In contrast, Stable-DDPG quickly brings the voltage into the stable operation range, which further demonstrates its applicability in power system voltage control. For the 13-bus network in Fig 4, with a control frequency larger than 0.82 Hz, both sides of the stability constraint $\frac{2}{\Delta T} X^{-1} \leftarrow \frac{\partial \mathbf{u}}{\partial \mathbf{v}} \leftarrow 0$ will hold.




Fig. 6: Stable-DDPG test with real-world load and PV dataset. The left plot is the PV and aggregated load. The right two plots are the voltage without control and with Stable-DDPG, where colored curves show voltage at different buses.

2) 123-bus system: We further test the controller performance in the IEEE 123-bus test feeder, which has 14 PV generators and controllers randomly selected to be placed at Buses 10, 11, 16, 20, 33, 36, 48, 59, 61, 66, 75, 83, 92, and 104. The system diagram is shown in Figure 4 (right). The nominal voltage magnitude at each bus except substation is 4.16 kV, and the acceptable range of operation is $\leq 5\%$ of the nominal value which is [3.952kV, 4.368kV].

Control Performance. Compared with IEEE 13-bus system, the IEEE 123-bus system is more sophisticated. As a result, the computation cost for simulation is higher. The policy training time for the Stable-DDPG is 1450.08s and 1300.14s for the DDPG. Table III compares the voltage recovery time and reactive power consumption of the trained controllers. Although DDPG performs slightly better when it can successfully stabilize the system (denoted as DDPG*), the lack of stability guarantee can lead to oscillations and instability, thus resulting higher overall costs. As shown in Figure 7, the DDPG voltage controller without considering stability can lead to voltage instability, while the proposed Stable-DDPG controller

TABLE III: Performance of linear, DDPG, and Stable-DDPG on 500 voltage violation scenarios with IEEE 123-bus case.

Method	Voltage recovery steps		Reactive power (Mvar)	
	Mean	Std	Mean	Std
Linear	41.30	20.30	1529.62	1302.60
Stable-DDPG	32.35	15.40	1178.77	992.70
DDPG	73.91	36.72	4515.33	2822.96
DDPG*	29.11	22.10	1148.24	1357.08

Note: DDPG* denotes the performance of the DDPG policy in the subset of testing cases when it was able to stabilize the voltage.




Fig. 7: Stable-DDPG and DDPG were tested on a low voltage scenario simulation. The left plot is the voltage trajectories, and the right plot is the reactive power injection.

shows good performance for the same test scenario. Figure 8 shows that our proposed Stable-DDPG stabilizes the system voltage in all test scenarios within 100 steps. In contrast, for DDPG, about 10% of buses' voltages are still beyond the $\leq 5\%$ range after a maximal control period (Fig. 8 Left), which accounts for approximately 63% of test scenarios (Fig. 8 Right). This further highlights the necessity of explicitly considering stability in learning-based controllers.

C. Three-phase Simulation Results

We now evaluate Stable-DDPG in three-phase systems. All simulations are built with the OpenDSS public models [48].

1) 13-bus system: To stabilize all the nodes of the network, we installed a PV generator and controller in every node except the substation node. The nominal voltage magnitude and the acceptable range are the same as in the single-phase experiment. Table IV summarizes the performance of different controllers. Our proposed




Fig. 8: Voltage stability for single-phase 123 bus test system. The left plot is the voltage violation for each bus, the right plot is the largest violation bus.

TABLE IV: Performance of linear, DDPG, and Stable-DDPG on 500 voltage violation scenarios with three-phase IEEE 13-bus test case.

Method	Voltage recovery steps		Reactive power (Mvar)	
	Mean	Std	Mean	Std
Linear	19.75	9.10	46.55	37.76
Stable-DDPG	14.61	3.74	29.94	16.07
DDPG	73.32	42.58	118.44	74.01
DDPG*	5.39	1.99	18.42	11.05

Note: DDPG* denotes the performance of the DDPG policy in the subset of testing cases when it was able to stabilize the voltage.

TABLE V: Performance of linear, DDPG, and Stable-DDPG on 500 scenarios with three-phase IEEE 123-bus test case.

Method	Voltage recovery steps		Reactive power (Mvar)	
	Mean	Std	Mean	Std
Linear	18.18	4.54	439.99	310.23
Stable-DDPG	12.70	4.99	318.31	273.22
DDPG	59.82	46.46	4715.57	3993.85
DDPG*	6.12	0.96	126.77	35.68

Note: DDPG* denotes the performance of the DDPG policy in the subset of testing cases when it was able to stabilize the voltage.

method achieves the best overall performance with a fast response and less reactive power consumption compared to the baseline linear policy and DDPG policy. While the DDPG algorithm has an impressive voltage recovery time and control cost if it successfully stabilizes the system (DDPG*), the percentage of stabilizing test cases is only around 34%. About 16.5% of buses' voltages fail to recover to the nominal range that spans 66% of 500 test scenarios, whereas Stable-DDPG achieves voltage stability in all scenarios. Furthermore, compared to the optimized linear policy, our method can save about 26.0% in time and 35.7% in reactive power consumption.

2) 123-bus system: Finally, we evaluate the proposed model with the unbalanced three-phase IEEE 123-bus system. The PV generator and controllers are installed in the same location as the single-phase IEEE 123-bus system. We summarize the control performance of different methods with the three-phase IEEE 123-bus system in Table V. According to the results, the average recovery time of the Stable-DDPG controller is 30% quicker compared to the optimal linear controller. Moreover, the reactive power consumption of the Stable-DDPG is 27.7% less than the optimal linear controller. Due to the absence of a stability guarantee, with the DDPG controller, 57.2% of the 500 test scenarios have at least one bus that fails to recover within 100 steps, leading to a significantly longer response time and a considerable increase in reactive power consumption.

D. Further Discussion

The above results also reveal an important trade-off between stability and the expressiveness of neural networks. DDPG algorithm with standard neural network policy obtains the best transient performance if it was

able to stabilize the system (see the performance of DDPG*). However, without a stability guarantee, the DDPG controller can lead to unstable working conditions, thus incurring overall high costs compared to both optimized linear policy and Stable-DDPG policy. With the monotone policy network, Stable-DDPG maintains the voltage magnitude in all test scenarios at the cost of a less flexible neural network parameterization. The linear policy can be regarded as an extreme example of a restricted neural net with only one learnable parameter, its slope, and thus might get sub-optimal performance compared to the monotone neural network with more learnable parameters.


VI. Conclusion and Future Works

In this work, we propose a stability-constrained reinforcement learning framework that formally guarantees the stability of RL for distribution system voltage control. The key technique that underpins the proposed approach is to use the Lyapunov stability theory and enforce the stability condition via monotone policy network design. We demonstrate the performance of the proposed method in IEEE single-phase and three-phase test systems. In terms of future work directions, one limitation of the proposed decentralized Stable-DDPG controller is that it can only guarantee voltage stability for the controlled buses. It is an interesting future direction to consider communications between neighboring nodes and design distributed controllers to ensure stability guarantees for buses without control. It is also a valuable future direction to unify the proposed approach in optimizing the transient cost of voltage control with steady-state cost optimization to obtain the best of both worlds. Additionally, a challenging and important task is to extend the monotone neural network design to multi-input multi-out monotone neural networks for the three-phase voltage controllers.

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] K. Turitsyn, P. Sülç, 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, June 2011.
- [3] "Ieee standard for interconnection and interoperability of distributed energy resources with associated electric power systems interfaces," *IEEE Std 1547-2018 (Revision of IEEE Std 1547-2003)*, pp. 1–138, 2018.
- [4] M. Farivar, R. Neal, C. Clarke, and S. Low, "Optimal inverter VAR control in distribution systems with high PV penetration," in *IEEE Power and Energy Society General Meeting*, San Diego, CA, July 2012.
- [5] H. Zhu and H. J. Liu, "Fast local voltage control under limited reactive power: Optimality and stability analysis," *IEEE Trans. Power Syst.*, vol. 31, no. 5, pp. 3794–3803, 2016.
- [6] G. Qu and N. Li, "Optimal distributed feedback voltage control under limited reactive power," *IEEE Trans. Power Syst.*, vol. 35, no. 1, pp. 315–331, 2019.
- [7] Y. Chen, Y. Shi, and B. Zhang, "Data-driven optimal voltage regulation using input convex neural networks," *Electric Power Systems Research*, vol. 189, p. 106741, 2020.
- [8] Y. Shi, G. Qu, S. Low, A. Anandkumar, and A. Wierman, "Stability constrained reinforcement learning for real-time voltage control," in *2022 American Control Conference (ACC)*. IEEE, 2022, pp. 2715–2721.
- [9] Y. Weng, Y. Liao, and R. Rajagopal, "Distributed energy resources topology identification via graphical modeling," *IEEE Transactions on Power Systems*, vol. 32, no. 4, pp. 2682–2694, 2016.
- [10] "Global survey of regulatory approaches for power quality and reliability," *Electric Power Research Institute*, Palo Alto, CA, Tech. Rep., 2005.
- [11] H. Haes Alhelou, M. E. Hamedani-Golshan, T. C. Njenda, and P. Siano, "A survey on power system blackout and cascading events: Research motivations and challenges," *Energies*, vol. 12, no. 4, p. 682, 2019.
- [12] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.
- [13] Y. Xu, Z. Y. Dong, R. Zhang, and D. J. Hill, "Multi-timescale coordinated voltage/var control of high renewable-penetrated distribution systems," *IEEE Transactions on Power Systems*, vol. 32, no. 6, pp. 4398–4408, 2017.
- [14] P. Sülç, S. Backhaus, and M. Chertkov, "Optimal distributed control of reactive power via the alternating direction method of multipliers," *IEEE Transactions on Energy Conversion*, vol. 29, no. 4, pp. 968–977, 2014.
- [15] W. Zheng, W. Wu, B. Zhang, H. Sun, and Y. Liu, "A fully distributed reactive power optimization and control method for active distribution networks," *IEEE Transactions on Smart Grid*, vol. 7, no. 2, pp. 1021–1033, 2016.
- [16] Z. Tang, D. J. Hill, and T. Liu, "Distributed coordinated reactive power control for voltage regulation in distribution networks," *IEEE Transactions on Smart Grid*, vol. 12, no. 1, pp. 312–323, 2021.
- [17] N. Li, G. Qu, and M. Dahleh, "Real-time decentralized voltage control in distribution networks," in *2014 52nd Annual Allerton Conference on Communication, Control, and Computing (Allerton)*. IEEE, 2014, pp. 582–588.
- [18] H. Liu and W. Wu, "Online multi-agent reinforcement learning for decentralized inverter-based volt-var control," *IEEE Transactions on Smart Grid*, vol. 12, no. 4, pp. 2980–2990, 2021.
- [19] J. Wang, W. Xu, Y. Gu, W. Song, and T. C. Green, "Multi-agent reinforcement learning for active voltage control on power distribution networks," in *Advances in Neural Information Processing Systems*, vol. 34, 2021.
- [20] Y. Gao, W. Wang, and N. Yu, "Consensus multi-agent reinforcement learning for volt-var control in power distribution networks," *IEEE Trans. Smart Grid*, pp. 1–1, 2021.
- [21] X. Sun and J. Qiu, "Two-stage volt/var control in active distribution networks with multi-agent deep reinforcement learning method," *IEEE Trans. Smart Grid*, pp. 1–1, 2021.
- [22] Y. Zhang, X. Wang, J. Wang, and Y. Zhang, "Deep reinforcement learning based volt-var optimization in smart distribution systems," *IEEE Trans. Smart Grid*, vol. 12, no. 1, pp. 361–371, 2021.
- [23] P. Kou, D. Liang, C. Wang, Z. Wu, and L. Gao, "Safe deep reinforcement learning-based constrained optimal control scheme for active distribution networks," *Applied Energy*, vol. 264, p. 114772, 2020.
- [24] Y. Chen, Y. Shi, D. Arnold, and S. Peisert, "Saver: Safe learning-based controller for real-time voltage regulation," *arXiv preprint arXiv:2111.15152*, 2021.
- [25] Q. Yang, G. Wang, A. Sadeghi, G. B. Giannakis, and J. Sun, "Two-timescale voltage control in distribution grids using deep reinforcement learning," *IEEE Transactions on Smart Grid*, vol. 11, no. 3, pp. 2313–2323, 2020.
- [26] S. Wang, J. Duan, D. Shi, C. Xu, H. Li, R. Diao, and Z. Wang, "A data-driven multi-agent autonomous voltage control framework using deep reinforcement learning," *IEEE Trans. Power Syst.*, 2020.
- [27] W. Wang, N. Yu, Y. Gao, and J. Shi, "Safe off-policy deep reinforcement learning algorithm for volt-var control in power distribution systems," *IEEE Transactions on Smart Grid*, vol. 11, no. 4, pp. 3008–3018, 2020.
- [28] D. Cao, W. Hu, J. Zhao, Q. Huang, Z. Chen, and F. Blaabjerg, "A multi-agent deep reinforcement learning based voltage reg-

- ulation using coordinated pv inverters,” IEEE Transactions on Power Systems, vol. 35, no. 5, pp. 4120–4123, 2020.
- [29] H. Liu and W. Wu, “Two-stage deep reinforcement learning for inverter-based volt-var control in active distribution networks,” IEEE Transactions on Smart Grid, vol. 12, no. 3, pp. 2037–2047, 2021.
- [30] C. Yeh, J. Yu, Y. Shi, and A. Wierman, “Robust online voltage control with an unknown grid topology,” in Proceedings of the Thirteenth ACM International Conference on Future Energy Systems, 2022, pp. 240–250.
- [31] W. Cui, Y. Jiang, and B. Zhang, “Reinforcement learning for optimal primary frequency control: A lyapunov approach,” IEEE Transactions on Power Systems, vol. 38, no. 2, pp. 1676–1688, 2022.
- [32] X. Chen, G. Qu, Y. Tang, S. Low, and N. Li, “Reinforcement learning for selective key applications in power systems: Recent advances and future challenges,” IEEE Transactions on Smart Grid, vol. 13, no. 4, pp. 2935–2958, 2022.
- [33] T. J. Perkins and A. G. Barto, “Lyapunov design for safe reinforcement learning,” Journal of Machine Learning Research, vol. 3, no. Dec, pp. 803–832, 2002.
- [34] Y. Chow, O. Nachum, E. Duenez-Guzman, and M. Ghavamzadeh, “A lyapunov-based approach to safe reinforcement learning,” Advances in Neural Information Processing Systems, 2018.
- [35] Y.-C. Chang, N. Roohi, and S. Gao, “Neural lyapunov control,” Advances in Neural Information Processing Systems, 2019.
- [36] M. E. Baran and F. F. Wu, “Network reconfiguration in distribution systems for loss reduction and load balancing,” IEEE Power Engineering Review, vol. 9, no. 4, pp. 101–102, 1989.
- [37] V. Kekatos, L. Zhang, G. B. Giannakis, and R. Baldick, “Voltage regulation algorithms for multiphase power distribution grids,” IEEE Transactions on Power Systems, vol. 31, no. 5, pp. 3913–3923, 2016.
- [38] S. Bolognani, R. Carli, G. Cavarero, and S. Zampieri, “Distributed reactive power feedback control for voltage regulation and loss minimization,” IEEE Trans. Autom. Control, vol. 60, no. 4, pp. 966–981, April 2015.
- [39] Z. Tang, D. J. Hill, and T. Liu, “Fast distributed reactive power control for voltage regulation in distribution networks,” IEEE Trans. Power Syst., vol. 34, no. 1, pp. 802–805, 2019.
- [40] N. Bof, R. Carli, and L. Schenato, “Lyapunov theory for discrete time systems,” arXiv preprint arXiv:1809.05289, 2018.
- [41] M. Blachuta, R. Bieda, and R. Grygiel, “Sampling rate and performance of dc/ac inverters with digital pid control—a case study,” Energies, vol. 14, no. 16, 2021.
- [42] W. Cui, J. Li, and B. Zhang, “Decentralized safe reinforcement learning for voltage control,” arXiv preprint arXiv:2110.01126, 2021.
- [43] A. Wehenkel and G. Louppe, “Unconstrained monotonic neural networks,” Advances in Neural Information Processing Systems, vol. 32, 2019.
- [44] W. Cui, Y. Jiang, B. Zhang, and Y. Shi, “Structured neural-pi control with end-to-end stability and output tracking guarantees,” Advances in Neural Information Processing Systems, 2023.
- [45] H. K. Khalil, Nonlinear systems, vol. 3.
- [46] K. P. Schneider, B. A. Mather, B. C. Pal, C.-W. Ten, G. J. Shirek, H. Zhu, J. C. Fuller, J. L. R. Pereira, L. F. Ochoa, L. R. de Araujo, R. C. Dugan, S. Matthias, S. Paudyal, T. E. McDermott, and W. Kersting, “Analytic considerations and design basis for the ieee distribution test feeders,” IEEE Transactions on Power Systems, vol. 33, no. 3, pp. 3181–3188, 2018.
- [47] L. Thurner, A. Scheidler, J. Dollichon, F. Schäfer, J.-H. Menke, F. Meier, S. Meinecke et al., “pandapower - convenient power system modelling and analysis based on pypower and pandas,” Tech. Rep., 2016.
- [48] “Opends.” [Online]. Available: <https://sourceforge.net/p/electricdss/code/HEAD/tree/trunk/Distrib/IEEETestCases/>
- [49] S. Janković and M. Merkle, “A mean value theorem for systems of integrals,” Journal of mathematical analysis and applications, vol. 342, no. 1, pp. 334–339, 2008.




Jie Feng (Student member, IEEE) received the B.E. degree in Automation from Zhejiang University, Hangzhou, China, in 2021. He is currently pursuing his Ph.D. degree in Electrical and Computer Engineering at the University of California, San Diego. His research interests focus on stability-constrained machine learning for power system control.



Yuanyuan Shi (Member, IEEE) is an Assistant Professor of Electrical and Computer Engineering at the University of California, San Diego. She received her Ph.D. in Electrical Engineering, masters in Electrical Engineering and Statistics, all from the University of Washington, in 2020. From 2020 to 2021, she was a Postdoctoral Scholar with the Department of Computing and Mathematical Sciences, Caltech. Her research interests include machine learning, dynamical systems, and control, with applications to sustainable power and energy systems.



Guannan Qu (Member, IEEE) received his B.S. degree in electrical engineering from Tsinghua University, Beijing, China, in 2014, and his Ph.D. degree in applied mathematics from Harvard University, Cambridge, MA, USA, in 2019. He is currently an Assistant Professor with the Electrical and Computer Engineering Department, Carnegie Mellon University, Pittsburgh, PA, USA. From 2019 to 2021, he was a Postdoctoral Scholar with the Department of Computing and Mathematical Sciences, Caltech, Pasadena, CA, USA. His research interests include control, optimization, and machine/reinforcement learning with applications to power systems, multi-agent systems, Internet of things.



Steven Low (Fellow, IEEE) is the F. J. Gilloon Professor of the Department of Computing & Mathematical Sciences and the Department of Electrical Engineering at Caltech. Before that, he was with AT&T Bell Laboratories, Murray Hill, NJ, and the University of Melbourne, Australia. He has held honorary/chaired professorship in Australia, China and Taiwan. He was a co-recipient of IEEE best paper awards, an awardee of the IEEE INFOCOM Achievement Award and the ACM SIGMETRICS Test of Time Award, and is a Fellow of IEEE, ACM, and CSEE. He was well-known for work on Internet congestion control and semidefinite relaxation of optimal power flow problems in smart grid. His research on networks has been accelerating more than 1TB of Internet traffic every second since 2014. His research on smart grid is providing large-scale electric vehicle charging to workplaces. He received his B.S. from Cornell and Ph.D. from Berkeley, both in EE.



Anima Anandkumar (Fellow, IEEE) works on AI algorithms and its applications to many domains in scientific areas. She is a fellow of the IEEE and ACM, and is part of the World Economic Forum's Expert Network. She has received several awards including the Guggenheim and Alfred P. Sloan fellowships, the NSF Career award, and best paper awards at venues such as Neural Information Processing and the ACM Gordon Bell Special Prize for HPC-Based COVID-19 Research. She recently presented her work on AI+Science to the White House Science Council. She received her B. Tech from the Indian Institute of Technology Madras and her Ph.D. from Cornell University and did her postdoctoral research at MIT. She was principal scientist at Amazon Web Services, and is now senior director of AI research at NVIDIA, and Bren named professor at Caltech.



Adam Wierman (Member, IEEE) Adam Wierman is a Professor in the Department of Computing and Mathematical Sciences (CMS) at the California Institute of Technology. He is the director of the Information Science and Technology (IST) initiative and served as Executive Officer (a.k.a. Department Chair) of CMS from 2015-2020. Additionally, he serves on the Advisory Board of the Linde Institute of Economic and Management Sciences and previously served on the Advisory Board of the “Sunlight to Everything” initiative of the Resnick Institute for Sustainability. He received his Ph.D., M.Sc., and B.Sc. in Computer Science from Carnegie Mellon University in 2007, 2004, and 2001, respectively, and has been a faculty at Caltech since 2007.

Appendix

Appendix A: Proof of Theorem 1

Proof of Theorem 1. Recall the closed-loop voltage dynamics $\mathbf{v}(t+1) = f_u(\mathbf{v}(t))$ with $\mathbf{u}(t) = g(\mathbf{v}(t))$. Let define $h(\mathbf{v}(t)) = \mathbf{v}(t) - f_u(\mathbf{v}(t))$. The Lyapunov function could be expressed compactly as follows,

$$V(\mathbf{v}(t)) = h(\mathbf{v}(t))^T X^{-1} h(\mathbf{v}(t)) \quad (16)$$

We further write $h(\mathbf{v}_{t+1})$ ¹ in terms of $h(\mathbf{v}_t)$ as follows,

$$h(\mathbf{v}_{t+1}) = h(\mathbf{v}_t) + \int_0^1 \frac{\partial h}{\partial \mathbf{v}}(\mathbf{v}_t + t(\mathbf{v}_{t+1} - \mathbf{v}_t))(\mathbf{v}_{t+1} - \mathbf{v}_t) dt$$

From Kowalewski’s Mean Value Theorem (Theorem 1 in [49]), we have $h(\mathbf{v}_{t+1}) = h(\mathbf{v}_t) + J_h(\mathbf{v}_{t+1} - \mathbf{v}_t)$, where $J_h = \sum_{i=1}^n \lambda_i \frac{\partial h}{\partial \mathbf{v}}(\mathbf{v}_t + k_i(\mathbf{v}_{t+1} - \mathbf{v}_t))$ for $k_i \in [0, 1]$, $\lambda_i \in [0, 1]$ for all i and $\sum_{i=1}^n \lambda_i = 1$. Note that, $\mathbf{v}_{t+1} - \mathbf{v}_t = f_u(\mathbf{v}_t) - \mathbf{v}_t = -h(\mathbf{v}_t)$. Thus, we get

$$h(\mathbf{v}_{t+1}) = (I - J_h)h(\mathbf{v}_t) \quad (17)$$

Therefore,

$$\begin{aligned} V(\mathbf{v}_{t+1}) &= h(\mathbf{v}_{t+1})^T X^{-1} h(\mathbf{v}_{t+1}) \\ &= h(\mathbf{v}_t)^T (I - J_h)^T X^{-1} (I - J_h) h(\mathbf{v}_t) \end{aligned} \quad (18)$$

We denote $G(\mathbf{v}, \theta) = \frac{\partial f_u}{\partial \mathbf{v}} + \frac{\partial f_u}{\partial \mathbf{u}} \frac{\partial \mathbf{u}}{\partial \mathbf{v}}$ as the Jacobian of the closed-loop voltage dynamics. and we then define $J_G = \sum_{i=1}^n \lambda_i G(\mathbf{v}_t + k_i(\mathbf{v}_{t+1} - \mathbf{v}_t), \theta)$, where k_i and λ_i follow the definition of J_h . From the definition of J_h and $h(\mathbf{v}_t)$, we have $J_G = I - J_h$. Thus we get

$$V(\mathbf{v}_{t+1}) - V(\mathbf{v}_t) = h(\mathbf{v}_t)^T (J_G^T X^{-1} J_G - X^{-1}) h(\mathbf{v}_t) \quad (19)$$

With Jensen’s inequality, $\mathcal{D}x / \mathbb{R}^n$, we further have

$$\begin{aligned} x^T J_G^T X^{-1} J_G x &= \mathbb{E} [X^{-1/2} J_G x]^2 \\ &= \mathbb{E} \left[\sum_{i=1}^n \lambda_i X^{-1/2} G(\mathbf{v}_t + k_i(\mathbf{v}_{t+1} - \mathbf{v}_t), \theta) x \right]^2 \\ &\geq \sum_{i=1}^n \lambda_i \mathbb{E} [X^{-1/2} G(\mathbf{v}_t + k_i(\mathbf{v}_{t+1} - \mathbf{v}_t), \theta) x]^2 \\ &= \sum_{i=1}^n \lambda_i x^T G(\mathbf{v}_t + k_i(\mathbf{v}_{t+1} - \mathbf{v}_t), \theta)^T \nabla \\ &\quad X^{-1} G(\mathbf{v}_t + k_i(\mathbf{v}_{t+1} - \mathbf{v}_t), \theta) x, \end{aligned} \quad (20)$$

¹We use shorthand $\mathbf{v}_{t+1}$ instead of $\mathbf{v}(t+1)$ to simplify the notation throughout the proof.

Therefore, with $G(\mathbf{v}, \theta)^T X^{-1} G(\mathbf{v}, \theta) \preceq X^{-1} \prec 0$, $\mathcal{D}x / \mathbb{R}^n$, we have $V(\mathbf{v}_{t+1}) - V(\mathbf{v}_t) < 0$ as long as $h(\mathbf{v}_t) \neq 0$, which means the Lyapunov function is decreasing along the system trajectory. Lastly, recall that $g_{i,\theta_i}(v_i) = 0$ for $v_i \in [\underline{v}_i, \bar{v}_i]$, so $V(\mathbf{v}_{t+1}) - V(\mathbf{v}_t) = 0$ implies that $\mathbf{v}_t \in S_v$.

Given that $G(\mathbf{v}, \theta) = I + I_{\Delta T} X \frac{\partial \mathbf{u}}{\partial \mathbf{v}}$, the stability condition becomes

$$(I + I_{\Delta T} X \frac{\partial \mathbf{u}}{\partial \mathbf{v}})^T X^{-1} (I + I_{\Delta T} X \frac{\partial \mathbf{u}}{\partial \mathbf{v}}) \preceq X^{-1} \prec 0,$$

Because of the decentralized characteristic, $\frac{\partial \mathbf{u}}{\partial \mathbf{v}}$ is a diagonal matrix. Expanding the multiplication terms, we get the stability condition as

$$\frac{2}{\Delta T} X^{-1} \prec \frac{\partial \mathbf{u}}{\partial \mathbf{v}} \prec 0 \quad (21)$$

By LaSalle’s Invariance Principle and the fact that $\lim_{\mathbf{v} \rightarrow \infty} \|g_\theta(\mathbf{v})\| = \infty$, the stability constraint is summarized in Theorem 1. $\square$

Appendix B: Experimental Details

We use Pytorch to build all RL models. Table VI show the hyperparameters for of the methods. The linear policy only has one parameter which is the slope, and it is optimized with the same RL framework. The Stable-DDPG requires monotonicity of the policy network, which leads to a specially designed one-layer monotone neural network. The Q network of all three baselines and the policy network of the DDPG are designed as three-layer fully connected neural networks, the numbers of hidden units are listed in the following table. More details about the simulation setup and model hyperparameters for all the testing cases can be found in <https://github.com/JieFengcse/Stable-DDPG-for-voltage-control>.

Hyper-parameters	DDPG	Stable-DDPG	Linear
Policy network	100-100	100	1
Q network	100-100	100-100	100-100
Discount factor (λ)	0.99	0.99	0.99
Q network learning rate	$2e^{-4}$	$2e^{-4}$	$2e^{-4}$
Maximum replay buffer size	1000000	1000000	1000000
Target Q network update ratio	$1e^{-2}$	$1e^{-2}$	$1e^{-2}$
Batch size	256	256	256
Activation function	ReLU	ReLU	ReLU

TABLE VI: Algorithm Hyperparameters