

Bridging Transient and Steady-State Performance in Voltage Control: A Reinforcement Learning Approach with Safe Gradient Flow

Jie Feng¹ Wenqi Cui² Jorge Cortés³ Yuanyuan Shi¹

Abstract—Deep reinforcement learning approaches are becoming appealing for the design of nonlinear controllers for voltage control problems, but the lack of stability guarantees hinders their real-world deployment. This paper constructs a decentralized RL-based controller for inverter-based real-time voltage control in distribution systems. It features two components: a transient control policy and a steady-state performance optimizer. The transient policy is parameterized as a neural network, and the steady-state optimizer represents the gradient of the long-term operating cost function. The two parts are synthesized through a safe gradient flow framework, which prevents the violation of reactive power capacity constraints. We prove that if the output of the transient controller is bounded and monotonically decreasing with respect to its input, then the closed-loop system is asymptotically stable and converges to the optimal steady-state solution. We demonstrate the effectiveness of our method by conducting experiments with IEEE 13-bus and 123-bus distribution system test feeders.

I. INTRODUCTION

Voltage safety is one of the primary concerns of power system operation, which requires the voltage magnitude to stay in an acceptable range under all working conditions [1]. In recent years, the integration of distributed energy resources (DERs) such as roof-top solar and electric vehicles has led to rapid and unpredictable fluctuations in the load and generation profiles of the distribution systems, thus leading to challenges in real-time voltage control for distribution grids.

There have been tremendous efforts made to overcome this challenge. Much of the attention has focused on optimizing the steady-state cost for voltage control [2]–[6], which refers to the operation cost after the system voltage has settled to equilibrium. However, as the system is subject to more frequent disturbances from load and generation fluctuations, optimizing the transient performance (i.e., how to optimally stabilize voltage after disturbances) becomes of equal importance.

The transient performance for the voltage control problem involves minimizing the voltage recovery time, at the minimum control effort. However, optimizing the transient cost for voltage control is a challenging task [7], as both the cost functions and system dynamics can be nonlinear. This is further complicated by the lack of exact model

knowledge and limited communications in the distribution grid. Recently, reinforcement learning (RL) has emerged as a powerful approach for addressing model-free nonlinear control problems. There has been considerable interest in developing RL-based controllers for optimizing the transient performance of voltage control problems. We refer readers to a recent survey [8].

Recent research has revealed that RL with a monotone policy network can ensure transient stability for voltage control [9]–[11]. However, these works do not offer guarantees regarding the optimality of the steady state. Steady-state requirements are difficult to enforce in RL since training can only occur over a finite horizon. Motivated by the challenges, the question we want to address in this paper is,

Can RL be structured to optimize both transient and steady-state performance for voltage control?

The key idea underlying our approach is the synthesis of a neural-network-based transient control policy and a steady-state optimizer (represented by the gradient of the cost function) in a safe gradient flow framework [12]. This enables us to coordinate these two sub-controllers to optimize both transient and steady-state performance while respecting the reactive power constraint and guaranteeing closed-loop stability. We summarize our main contributions as follows:

- We design a decentralized RL-based controller that optimizes both transient and steady-state performance for the distribution system voltage control;
- We prove that the proposed controller design guarantees both transient stability and steady-state optimality, for strictly convex objective functions (in Theorem 1);
- We demonstrate the effectiveness of the proposed method with extensive numerical experiments. Our method can reduce over 30% transient cost compared to controllers that only optimize the steady-state or transient performance, and guarantee optimal steady-state cost.

The remaining parts of this paper are organized as follows. Section II presents the distribution system voltage control problem and steady-state optimization solution as preliminaries. Section III presents the proposed transient and steady-state reinforcement learning algorithm. Its stability and optimality guarantees are rigorously proved in Section IV. Section V evaluates the results in IEEE test feeders and provides a discussion, followed by concluding remarks in Section VI.

II. MODEL & PRELIMINARIES

In this section, we review the distribution system power flow model and introduce the voltage control problem.

*The authors are supported by NSF ECCS-2200692, ECCS-2153937, and ECCS-1947050 awards and the Jacobs School Early Career Faculty Development Award.

¹Jie Feng and Yuanyuan Shi are with the Department of Electrical and Computer Engineering, University of California San Diego, jif005@ucsd.edu, yyshi@eng.ucsd.edu.

²Wenqi Cui is with the Department of Electrical and Computer Engineering, University of Washington, Seattle, wenqicui@uw.edu.

³Jorge Cortés is with the Department of Mechanical and Aerospace Engineering, University of California San Diego, cortes@ucsd.edu.

A. Branch Flow Model

We consider the linearized branch flow model [3] in a tree-structured distribution network for theoretical analysis. The system is defined as $\mathcal{G} = (\mathcal{N}_0, \mathcal{E})$, consisting of a set of nodes $\mathcal{N}_0 = \{0, 1, \dots, n\}$ and an edge set \mathcal{E} . Node 0 is known as the substation, and $\mathcal{N} = \mathcal{N}_0 / \{0\}$ denotes the set of nodes excluding the substation node. Each node $i \in \mathcal{N}$ is associated with an active power injection p_i and a reactive power injection q_i . Let v_i be the squared voltage magnitude, and let p, q and v denote $\{p_i, q_i, v_i\}_{i \in \mathcal{N}}$ stacked into a vector. The variables satisfy the following equations, $\forall i \in \mathcal{N}$,

$$p_i = -P_{ji} + \sum_{k:(i,k) \in \mathcal{E}} P_{ik}, \quad q_i = -Q_{ji} + \sum_{k:(i,k) \in \mathcal{E}} Q_{ik}, \quad (1a)$$

$$v_i = v_j - 2(r_{ji}P_{ji} + x_{ji}Q_{ji}), \quad (1b)$$

where j is the parent node of i in the distribution network, P_{ji} and Q_{ji} represent the active power and reactive power flow on line (j, i) , and r_{ji} and x_{ji} are the line resistance and reactance. (1) can be written in the vector form,

$$v = Rp + Xq + v_0 \mathbf{1} = Xq + v^{env}, \quad (2)$$

where $v^{env} = Rp + v_0 \mathbf{1}$ is the non-controllable part. $R = [R_{ij}]_{n \times n}$, $X = [X_{ij}]_{n \times n}$ are defined as $R_{ij} := 2 \sum_{(h,k) \in \mathcal{P}_i \cap \mathcal{P}_j} r_{hk}$, $X_{ij} := 2 \sum_{(h,k) \in \mathcal{P}_i \cap \mathcal{P}_j} x_{hk}$. Here, \mathcal{P}_i is the set of lines on the unique path from bus 0 to bus i , and v_0 is the squared voltage magnitude at the substation bus. R and X are positive definite matrices and all elements are positive [9].

We make the following assumptions that are well-justified for voltage control on distribution networks [3], [9].

Assumption 1. *The system models (i.e., matrices R and X in (2)) are time-invariant, and the controllers are installed in every bus without real-time communication.*

Assumption 2. *There is a timescale separation between the voltage dynamics and the dynamics of inverters, so that the controlled inverter injects instantaneously the exact value of reactive power computed by the controller.*

B. Voltage Control Problem

The optimal voltage control problem at the steady state is,

$$\min_q F(q) = C(q) + \frac{1}{2} q^\top X q + q^\top \Delta \bar{v} \quad (3a)$$

$$\text{s.t. } \underline{q} \leq q \leq \bar{q} \quad (3b)$$

where $C(q)$ is the control cost and $\Delta \bar{v} = v^{env} - v^{nom}$. We define the reactive power safety set as $\mathcal{S}_q = \{q \in \mathbb{R}^n | q \leq \bar{q}\}$. For the per-unit system, we define $v^{nom} = 1$ p.u. Using (2), the objective function can be rewritten as $F(q) = C(q) + \frac{1}{2} (v - v^{nom})^\top X^{-1} (v - v^{nom}) - \frac{1}{2} \Delta \bar{v}^\top X^{-1} \Delta \bar{v}$. Since the last term is a constant, the objective function finds an optimal trade-off between minimizing the control cost $C(q)$ and the voltage deviation $\frac{1}{2} (v - v^{nom})^\top X^{-1} (v - v^{nom})$. Following [5], we consider $C(q) = \sum_i C_i(q_i)$, $C_i(q_i) = \frac{\eta_i}{2\bar{s}_i} q_i^2$, where $\eta_i, \bar{s}_i > 0$ represent the cost of reactive power of bus i and its apparent power capacity. Compactly, $C(q) = \frac{1}{2} q^\top C_q q$, where C_q is the diagonal matrix $\text{diag}\{\frac{\eta_i}{\bar{s}_i}\}_{i \in \mathcal{N}}$.

Note that the objective function (3a) can be equivalently written as the sum of cost at all nodes by (2),

$$F(q) = \sum_i^n \left(C_i(q_i) + \frac{1}{2} (q_i(v_i + v_i^{env} - 2v_i^{nom})) \right).$$

The gradient of the objective function ∇F is,

$$\nabla F = C_q q + X q + \Delta \bar{v} \stackrel{\textcircled{1}}{=} C_q q + v - v^{nom}. \quad (4)$$

where $\textcircled{1}$ follows from (2) and the definition of $\Delta \bar{v} := v^{env} - v^{nom}$. We write $\nabla F_i = \frac{\eta_i}{\bar{s}_i} q_i + v_i - v_i^{nom}$. The decomposable structure of the objective function and the gradient enables decentralized training and deployment of a controller. We make the assumption that the optimal solution of (3) is unique and the corresponding voltage v^* lies in the safe voltage range.

Assumption 3. *The steady-state optimization problem (3) is strictly convex, the optimal solution (v^*, q^*) satisfies $v^* \in \mathcal{S}_v = \{v \in \mathbb{R}^n : v_i \leq v_i \leq \bar{v}_i\}$, $q^* \in \mathcal{S}_q$ where v_i, \bar{v}_i are upper and lower bounds of desired system voltage magnitudes.*

To solve (3), [2] introduces a projected gradient method

$$q_i(t+1) = [q_i(t) - \gamma \nabla F_i]_{\bar{q}_i}^{\underline{q}_i}, \quad (5)$$

where $[\cdot]_b^a$ denotes the projection onto $[a, b]$, and \underline{q}_i and \bar{q}_i are the lower and upper bound of reactive power capacity. If the stepsize γ satisfies $\gamma < \frac{2}{\lambda_{\max}(\nabla^2 C(q) + X)}$, where λ_{\max} denotes the largest eigenvalue, $v(t)$ and $q(t)$ under the controller (5) converge to (v^*, q^*) – the optimal solution of (3). However, this approach does not account for optimizing transient performance, which is critical when the system is subject to rapid voltage deviations due to renewable integration and EV charging (reflected in changes in v^{env}). This limitation motivates our design of a controller that jointly optimizes both steady-state and transient performance.

III. JOINT TRANSIENT AND STEADY STATE PERFORMANCE OPTIMIZATION

In this section, we first introduce the joint transient and steady-state optimization problem. Then, we propose a *Transient and Steady-state Reinforcement Learning (TASRL)* framework to solve it. We provide transient stability and steady-state optimality guarantees in the next section.

From the system operator perspective, we wish to achieve two main objectives: 1) *transient stability and performance*: fast convergence of system voltage to the desired operating range \mathcal{S}_v (e.g., $\pm 5\%$ around v^{nom}) after a disturbance; 2) *steady-state performance*: maintaining the system operation at the most economical point. Thus, the joint transient and steady-state optimization problem is formulated as,

$$\min_{\theta} J(\theta) = \int_{t=0}^{t_f} \gamma \sum_{i=1}^n c_i(v_i(t), q_i(t)) \quad (6a)$$

$$\text{s.t. } v(t) = X q(t) + v^{env}, \quad (6b)$$

$$\dot{q}(t) = f_{\theta}(q(t), v(t)), \quad (6c)$$

$$v^* = \lim_{t \rightarrow t_f} v(t), q^* = \lim_{t \rightarrow t_f} q(t), \quad (6d)$$

$$q^* \text{ is the optimal solution for (3),} \quad (6e)$$

$$q(t) \in \mathcal{S}_q, \forall t \text{ and } v^* \in \mathcal{S}_v. \quad (6f)$$

where γ is the discount factor and c_i is the cost function at node i , for which we choose $c_i(v_i, q_i) = C_i(q_i) + \frac{1}{2}q_i(v_i + v_i^{emv} - 2v_i^{nom})$. $f_\theta(\cdot, \cdot)$ is the controller to be optimized. Here, $t_f \in \mathbb{R}_{>0}$ is the (possibly unbounded) stabilization time. During the transient period $[0, t_f]$, the goal is to recover the voltage quickly under limited reactive power resources while minimizing the control effort. In addition, we want the system to converge to the steady-state optima $\lim_{t \rightarrow t_f} v(t) \rightarrow v^*$, $\lim_{t \rightarrow t_f} q(t) \rightarrow q^*$, and such that the equilibrium point (v^*, q^*) solves the steady-state optimization (3).

Controller synthesis inspired by safe gradient flow: Here, we introduce the proposed decentralized controller design. Each controller measures the local voltage magnitude and computes the local reactive power injection without real-time communication. To jointly optimize the transient and steady-state performance, let's start by considering a direct combination of a transient policy $\pi_\theta(v)$, parameterized by a neural network, and the gradient $\nabla F(q)$, cf. (4),

$$\dot{q} = f_\theta(q, v) := -\nabla F(q) + \pi_\theta(v), \quad (7)$$

with $q(0) = 0$. The main issue with the controller (7) is that the resulting trajectory of $q(t)$ may not satisfy the reactive power constraint $\underline{q} \leq q(t) \leq \bar{q}$ at all times. To enforce it, we build on the safe gradient flow framework introduced in [12]. This design employs a control barrier function $g(q)$ to ensure that a given dynamics never leaves a safe set $\mathcal{S}_q = \{q \in \mathbb{R}^n | g(q) \leq 0\}$, where $g(q)$ is defined by

$$g(q) = [I \quad -I]^\top q + [-\bar{q} \quad \underline{q}]^\top. \quad (8)$$

If the dynamics $\dot{q} = f_\theta(q, v)$ satisfies

$$\frac{\partial g}{\partial q} f_\theta(q, v) \leq -\alpha g(q), \quad (9)$$

then, by Nagumo's theorem [13], $q(t)$ must stay inside the safe region \mathcal{S}_q for all t . Here, the hyperparameter $\alpha > 0$ indicates the degree of conservatism regarding the reactive power constraints, where the larger α is, the more flexibility is allowed when $q(t)$ is not reaching the constraints. To ensure that the controller in (7) satisfies the safety constraints, the safe gradient flow framework [12] prescribes modifying it minimally according to the following control barrier function quadratic program (CBF-QP) safety filter [14],

$$f_\theta(q, v) = \arg \min_{\xi \in \mathbb{R}^n} \frac{1}{2} \|\xi - (-\nabla F(q) + \pi_\theta(v))\|_2^2 \quad (10a)$$

$$\text{s.t. } \frac{\partial g}{\partial q} \xi \leq -\alpha g(q) \quad (10b)$$

Plugging in $g(q)$ gives,

$$f_\theta(q, v) = \arg \min_{\xi \in \mathbb{R}^n} \frac{1}{2} \|\xi + \nabla F(q) - \pi_\theta(v)\|_2^2 \quad (11a)$$

$$\text{s.t. } \alpha(\underline{q} - q) \leq \xi \leq \alpha(\bar{q} - q) \quad (11b)$$

The controller (11) is our proposed controller to solve the joint optimization problem in (6). We envision that (11) finds the safe control action closest to $-\nabla F(q) + \pi_\theta(v)$ while

Algorithm 1 Transient and Steady-state Reinforcement Learning (TASRL) for Distribution Grid Voltage Control

Ensure: policy networks $\pi_{i, \theta_i}(v_i)$ with parameters θ_i ; hyperparameter α ; sampling time h ; replay buffers $\mathcal{D}_i, \forall i \in \mathcal{N}$.

for $j = 0$ to N_{ep} **do**

 Randomly generate initial states $v(0)$

for $t = 0$ to N_{step} **do**

 For each agent $i \in \mathcal{N}$

 Measure the current state $v_i(t)$

 Compute the control action (reactive power adjustment) $f_{i, \theta_i}(q_i(t), v_i(t)) = [\pi_{i, \theta_i}(v_i(t)) - \nabla F_i]_{\alpha(\underline{q}_i - q_i(t))}^{\alpha(\bar{q}_i - q_i(t))}$

 Execute $q_i(t+1) = q_i(t) + h f_{i, \theta_i}(q_i(t), v_i(t))$

 Transit to next state $v_i(t+1)$, receive cost $c_i(t)$

 Store $\{v_i(t), q_i(t), f_{i, \theta_i}(q_i(t), v_i(t)), -c_i(t), v_i(t+1)\}$ in \mathcal{D}_i

 Update policy network θ_i

end for

end for

ensuring the reactive power constraints are never violated. If sufficient reactive power capacity exists, $f_\theta(q, v) = -\nabla F(q) + \pi_\theta(v)$, otherwise, the action is projected to ensure reactive power capacity constraints are met. As $\alpha \rightarrow \infty$, (11) reduces to a projection of $-\nabla F(q(t)) + \pi_\theta(v(t))$ onto $[\underline{q}, \bar{q}]$.

Proposition 1. *The optimal solution to (11) is given by,*

$$f_\theta(q, v) = [\pi_\theta(v) - \nabla F(q)]_{\alpha(\underline{q} - q)}^{\alpha(\bar{q} - q)} \quad (12)$$

where $[\cdot]_b^a$ denotes the projection onto $[a, b]$.

Proposition 1 is a direct result of solving the QP [15, Chapter 8] formulated by Equation (11). We summarize the proposed controller in Algorithm 1. As observed in Algorithm 1, the controller computation and execution are decentralized. The training process of TASRL follows the standard policy optimization RL algorithms. Each local transient policy $\pi_{i, \theta_i}(v_i)$ can be parameterized as neural networks (with requirements specified in Section IV Def. 1) and trained together with ∇F to optimize the transient performance. The proposed TASRL framework is general and can be integrated with most policy optimization methods, including DDPG [16], PPO [17], and TRPO [18].

IV. TRANSIENT STABILITY AND STEADY-STATE OPTIMALITY GUARANTEES

In this section, we establish the closed-loop stability and optimal steady-state performance properties of Algorithm 1. The guarantees rely on certain structural constraints of the transient policy in the next definition.

Definition 1 (Stable decentralized transient policy). *A set of local policy $\{\pi_{i, \theta_i}, \forall i \in \mathcal{N}\}$ is a stable transient policy if it satisfies the following conditions for each bus $i \in \mathcal{N}$:*

- 1) $\pi_{i, \theta_i}(v_i)$ is a continuously differentiable function satisfying $\pi_{i, \theta_i}(v_i) = 0$ for $v_i \in [\underline{v}_i, \bar{v}_i]$;
- 2) $\pi_{i, \theta_i}(v_i)$ is monotonically decreasing for $v_i \in (-\infty, \underline{v}_i) \cup (\bar{v}_i, \infty)$;

- 3) $\pi_{i,\theta_i}(v_i)$ is bounded i.e. $c\alpha(q'_i - q_i(t)) \leq \pi_{i,\theta_i}(v_i(t)) \leq c\alpha(\bar{q}'_i - q_i(t))$, where $q'_i = q_i(1 - \varepsilon)$, $\bar{q}'_i = \bar{q}_i(1 - \varepsilon)$, and $\varepsilon \in (0, 1)$ $c \in [0, 1)$.

We write $\pi_{i,\theta_i}(v_i) = \pi_{i,\theta_i}(v_i) - \pi_{i,\theta_i}(v_i^*)$ as $\pi_{i,\theta_i}(v_i^*) = 0$. For $v_i \neq v_i^*$, we define $K_{ii}(v_i) = \frac{\pi_{i,\theta_i}(v_i) - \pi_{i,\theta_i}(v_i^*)}{v_i - v_i^*}$ and $K(v) := -\text{diag}(K_{11}(v_1), K_{22}(v_2), \dots, K_{nn}(v_n))$. By the monotonically decreasing condition in Definition 1, when $v_i \in (-\infty, \underline{v}_i) \cup (\bar{v}_i, \infty)$, $K_{ii}(v_i) < 0$. When $v_i \in [\underline{v}_i, \bar{v}_i]$ and $v_i \neq v_i^*$, $K_{ii}(v_i) = 0$. We define $K_{ii}(v_i) = 0$ if $v_i = v_i^*$. As a result, for every v , we can write $\pi_\theta(v) = -K(v)(v - v^*)$. Define $\sigma_{\max}(\cdot)$ and $\sigma_{\min}(\cdot)$ as the largest and smallest singular value of a matrix. The following result establishes the theoretical guarantees.

Theorem 1 (Transient Stability and Steady-State Optimality). *Suppose Assumption 3 holds, π_θ is a stable decentralized transient policy according to Definition 1, and $2\sigma_{\max}(K(v)) \leq \sigma_{\min}(C_q X^{-1} + I)$, $\forall v \in \mathbb{R}^n$, then with a sufficiently large α , the closed-loop system is asymptotically stable with controller (11). In addition, $\lim_{t \rightarrow t_f} v(t) = v^*$, $\lim_{t \rightarrow t_f} q(t) = q^*$, q^* is the global minimizer of optimization problem (3), $v^* \in \mathcal{S}_v$, and $q(t) \in \mathcal{S}_q, \forall t \geq 0$.*

Theorem 1 shows that with a Lipschitz-like bound on the transient policy $\pi_{i,\theta_i}, \forall i \in \mathcal{N}$, the proposed controller in (11) obtains both transient stability and steady-state optimality while respecting the reactive power capacity constraint at all times. We present its proof below.

Lemma 1. *Suppose $2\sigma_{\max}(K(v)) \leq \sigma_{\min}(C_q X^{-1} + I)$, $\forall v \in \mathbb{R}^n$ and let $\langle \cdot, \cdot \rangle$ be the dot product, then $\|-\nabla F(q)\|^2 + 2\langle \pi_\theta(v), -\nabla F(q) \rangle \geq 0, \forall v \in \mathbb{R}^n, q \in \mathcal{S}_q$.*

Proof. Following (4) and (2), $\nabla F(q) = \nabla F(q) - \nabla F(q^*) = (C_q X^{-1} + I)(v - v^*)$. Denote $A = C_q X^{-1} + I$, we have

$$\begin{aligned} & \|-\nabla F(q)\|^2 + 2\langle \pi_\theta(v), -\nabla F(q) \rangle \\ &= (v - v^*)^\top \left[A^\top A + K(v)A + A^\top K(v) \right] (v - v^*). \end{aligned}$$

To ensure $\|-\nabla F(q)\|^2 + 2\langle \pi_\theta(v), -\nabla F(q) \rangle \geq 0$, it suffices to have $A^\top A + K(v)A + A^\top K(v) \succeq 0$.

$$A^\top A + K(v)A + A^\top K(v) \succeq 0 \quad (13a)$$

$$\iff [A + K(v)]^\top [A + K(v)] \succeq K(v)^\top K(v) \quad (13b)$$

$$\iff \sigma_{\min}(A + K(v)) \geq \sigma_{\max}(K(v)) \quad (13c)$$

By [19, Proposition 9.6.8], $\sigma_{\min}(A + K(v)) \geq \sigma_{\min}(A) - \sigma_{\max}(K(v))$. Thus, $\sigma_{\min}(A) \geq 2\sigma_{\max}(K(v))$ is a sufficient condition for $\|-\nabla F(q)\|^2 + 2\langle \pi_\theta(v), -\nabla F(q) \rangle \geq 0$. \square

A. Proof of Theorem 1

Proof of Theorem 1. By design, the proposed controller (11) guarantees $q(t) \in \mathcal{S}_q, \forall t \geq 0$ by the CBF-QP safety filter. We will work with the equivalent controller form in (10) throughout the proof since the inequality constraints are in a more compact form. The Lagrangian of (10) is,

$$L(\xi, \omega; q) = \frac{1}{2} \|\xi + \nabla F(q) - \pi_\theta(v)\|_2^2 + \omega^\top \left(\frac{\partial g(q)}{\partial q} \xi + \alpha g(q) \right),$$

where ω is the nonnegative Lagrange multiplier for (10b). The Karash-Kuhn-Tucker (KKT) conditions [20, Chapter 3] of (10) are:

$$\xi + \nabla F(q) - \pi_\theta(v) + \frac{\partial g(q)}{\partial q} \omega = 0 \quad (14a)$$

$$\omega \geq 0, \frac{\partial g(q)}{\partial q} \xi + \alpha g(q) \leq 0 \quad (14b)$$

$$\omega^\top \left(\frac{\partial g(q)}{\partial q} \xi + \alpha g(q) \right) = 0 \quad (14c)$$

Because (10) is strongly convex with respect to ξ , the existence of a (ξ, ω) satisfying (14) is sufficient for $\xi = f_\theta(q)$. To verify the feasibility of the KKT conditions, we apply the Mangasarian-Fromovitz Constraint Qualification (MFCQ) condition [20, Chapter 3], which requires a $\xi \in \mathbb{R}^n$ s.t.

$$\nabla g_i(q)^\top \xi < 0 \quad \forall i \in I_0(q) = \{1 \leq i \leq 2n | g_i(q) = 0\}$$

where $I_0(q)$ is the active constraint set. Given the specific structure of $g(q), \forall q \in \mathcal{S}_q$, there always exists an ξ such that the MFCQ is satisfied. By Lemma 4.5 of [12], the existence of a solution (ξ, ω) satisfying (14) is guaranteed. We next characterize the stability properties of our proposed algorithm with the solution (ξ, ω) of the KKT conditions (14).

An immediate result of Lemma 1 is

$$\begin{aligned} \|\pi_\theta(v) - \nabla F(q)\|^2 &= \langle \pi_\theta(v) - \nabla F(q), \pi_\theta(v) - \nabla F(q) \rangle \\ &= \|\pi_\theta(v)\|^2 + \|-\nabla F(q)\|^2 + 2\langle \pi_\theta(v), -\nabla F(q) \rangle \geq \|\pi_\theta(v)\|^2 \end{aligned}$$

For every bus i s.t. $\pi_{i,\theta_i}(v_i) \neq 0$, we have $\|\pi_{i,\theta_i}(v_i) - \nabla F_i(q)\|_{\alpha(q_i - \bar{q}_i)}^{\alpha(\bar{q}_i - q_i)} \geq \|\pi_{i,\theta_i}(v_i) - \nabla F_i(q)\|_{\alpha\varepsilon(q_i)}^{\alpha\varepsilon(\bar{q}_i)}$. Given that $-\nabla F_i(q)$ is bounded on both sides, there exists a finite α such that $\|\pi_{i,\theta_i}(v_i) - \nabla F_i(q)\|_{\alpha(q_i - \bar{q}_i)}^{\alpha(\bar{q}_i - q_i)} = \pi_\theta(v_i) - \nabla F_i(q)$. A similar reasoning holds if $\pi_{i,\theta_i}(v_i) = 0$. Therefore, with a sufficiently large α , we have $\|f_\theta(q, v)\| = \|\pi_\theta(v) - \nabla F(q)\|_{\alpha(q - \bar{q})}^{\alpha(q - \bar{q})} \geq \|\pi_\theta(v)\|$. Using $F(q)$ as a Lyapunov-like function, from Eq (14a), we have $\nabla F(q) = -\frac{\partial g(q)}{\partial q} \omega - f_\theta(q, v) + \pi_\theta(v)$, thus

$$\begin{aligned} L_{f_\theta} F(q(t)) &= f_\theta(q(t), v(t))^\top \nabla F(q(t)) \\ &= f_\theta(q(t), v(t))^\top \left(-\frac{\partial g(q(t))}{\partial q(t)} \omega - f_\theta(q(t), v(t)) + \pi_\theta(v(t)) \right) \\ &= -\|f_\theta(q(t), v(t))\|^2 + f_\theta^\top(q(t), v(t)) \pi_\theta(v(t)) + \alpha \omega^\top g(q(t)) \\ &\leq -\|f_\theta(q(t), v(t))\|^2 + \|f_\theta(q(t), v(t))\| \|\pi_\theta(v(t))\| + \alpha \omega^\top g(q(t)) \\ &\leq 0 \end{aligned} \quad (15)$$

where the second equality follows Eq (14c). Given that $g(q(t)) \leq 0$, ω is a nonnegative dual variable, $\alpha > 0$, $\alpha \omega^\top g(q(t)) \leq 0$ holds, which leads to the final inequality.

Furthermore, $L_{f_\theta} F(q) = 0$ if and only if $f_\theta(q^*, v^*) = 0$. Then there exists $(0, \omega^*)$, which is the solution of (14). Plug $(0, \omega^*)$ into (14), it is reduced to

$$\nabla F(q^*) - \pi_\theta(v^*) + \frac{\partial g(q^*)}{\partial q} \omega^* = 0 \quad (16a)$$

$$\omega^* \geq 0, \alpha g(q^*) \leq 0 \quad (16b)$$

$$(\omega^*)^\top (\alpha g(q^*)) = 0 \quad (16c)$$

By Assumption 3 where $v^* \in \mathcal{S}_v$, we have $\pi_\theta(v^*) = 0$. Given that $\alpha > 0$, (v^*, q^*) is the optimal solution of (3). Due to the strict convexity of $F(q)$, q^* is the unique global minimizer. By Lyapunov Stability Theory [21, Chapter 3], we conclude that the closed loop system is asymptotically stable with respect to the global minimizer. \square

B. Stable transient policy design

We now present the neural network design that meets the stable decentralized transient policy $\pi_{i,\theta_i}(v_i)$ in Definition 1.

Conditions 1) and 2): To ensure conditions 1) and 2), we adopt the structure in [22] for each bus i . The stacked ReLU function constructed by (17) is monotonically decreasing for $v_i - \bar{v} > 0$ and zero when $v_i - \bar{v} \leq 0$.

$$\xi^+(v_i - \bar{v}; w^+, b^+) = (w^+)^T \text{ReLU}(\mathbf{1}(v_i - \bar{v}) + b^+), \quad (17a)$$

$$\sum_{l=1}^{d'} w_l^+ < 0, \forall d' = 1, \dots, d, b_1^+ = 0, b_l^+ \leq b_{l-1}^+, \quad (17b)$$

$\forall l = 2, \dots, d$. The stacked ReLU function constructed by (18) is monotonically decreasing for $v_i - \underline{v} < 0$ and zero otherwise.

$$\xi^-(v_i - \underline{v}; w^-, b^-) = (w^-)^T \text{ReLU}(-\mathbf{1}(v_i - \underline{v}) + b^-), \quad (18a)$$

$$\sum_{l=1}^{d'} w_l^- > 0, \forall d' = 1, \dots, d, b_1^- = 0, b_l^- \leq b_{l-1}^-, \forall l = 2, \dots, d. \quad (18b)$$

Condition 3) This condition requires the output of $\pi_{i,\theta_i}(v_i)$ to be bounded. We use tanh activation function to scale the output as a percentage while preserving its sign. Then the percentage is multiplied by the absolute value of the bounds. The local policy network of bus i is defined as,

$$\begin{aligned} \pi_{i,\theta_i}(v_i) = & c\alpha(q_i - \bar{q}_i) \tanh(\xi^+(v_i - \bar{v}; w^+, b^+)) \\ & + c\alpha(\bar{q}_i - q_i) \tanh(\xi^-(v_i - \underline{v}; w^-, b^-)). \end{aligned} \quad (19)$$

The monotonicity of tanh ensures that 1) and 2) are satisfied.

V. EXPERIMENTS

In this section, we demonstrate the effectiveness of the proposed method in two IEEE distribution test systems.

A. Experiment Setup

We evaluate our approach on the IEEE 13-bus and 123-bus test feeders [23]. The nominal voltage magnitude for both environments at each bus except the substation is 4.16 kV, and the acceptable range of operation is $\pm 5\%$ of the nominal value, which is [3.952kV, 4.368kV]. Though our theoretical analysis is based on the linearized system model in (2), all experiments are run using Pandapower [24] as the nonlinear power flow simulator to evaluate the algorithm performance. We simulate 5% to 15% high voltage and low voltage violation scenarios. Two baselines are tested:

- 1) Stable-DDPG [9] with Safety Filter: The Stable-DDPG [9] optimizes the transient performance with a stability guarantee. To enforce reactive power safety, we incorporate the CBF-QP safety filter in (10) by replacing the obj (10a) with $\arg \min_{\xi \in \mathbb{R}^n} \frac{1}{2} \|\xi - \pi_\theta(v)\|_2^2$.




Fig. 1: Left: Schematic diagram of the 13 bus system with three PV generators and voltage controllers located at nodes 2, 7, and 9. Right: Schematic diagram of 123 bus system, with 14 PV generators and voltage controllers located at nodes 10, 11, 16, 20, 33, 36, 48, 59, 61, 66, 75, 83, 92, and 104.

TABLE I: Performance of 500 scenarios for 13-bus system.

Method	Transient Performance		Steady State
	Recovery Time (s)	Transient Cost	Objective $F(\mathbf{q})$
Stable-DDPG [9]	10.18	-5.61	-0.09
Safe gradient flow [12]	3.04	-6.74	-0.11
TASRL	2.60	-6.76	-0.11

Note: Smaller value means better performance for all three columns.

- 2) Safe gradient flow [12]: Safe gradient flow optimizes the steady-state performance with reactive power safety.

B. Results

1) *IEEE 13-bus*: IEEE 13-bus system is a standard radial distribution system depicted in Figure 1 (Left), where three PV stations and controllers are located at buses 2, 7, and 9.

Table I compares the transient and steady-state performance of 500 different voltage violation scenarios. Clearly, TASRL achieves the best performance for both transient and steady states. For this test case, the magnitude of the gradient is relatively large. As a result, the transient performance of the safe gradient flow is close to TASRL, and the Stable-DDPG is underperforming. In terms of steady-state performance, both the safe gradient flow and the TASRL achieve the best result as $v(t)$ and $q(t)$ converge to the steady-state optima (v^*, q^*) .

2) *IEEE 123-bus*: Figure 1 (Right) demonstrates the IEEE 123-bus distribution test feeder, which has 14 PV generators and controllers randomly placed in the network. We summarize the performance of our method and the baselines in Table II. The average response time for TASRL is 12.08 steps, which saved 77% of time compared to the safe gradient flow, and 30% compared to the Stable-DDPG. Both the TASRL and safe gradient flow obtain optimal steady-state cost. Interestingly, compared to the results of IEEE 13-bus, the gap between the steady-state performance of Stable-DDPG and the other two methods is larger. This indicates that optimizing the steady-state performance becomes increasingly crucial as the system complexity increases.

TABLE II: Performance of 500 scenarios for 123-bus system.

Method	Transient Performance		Steady State
	Recovery Time (s)	Transient Cost	Objective $F(\mathbf{q})$
Stable-DDPG	17.36	-303.44	-4.90
Safe gradient flow	52.43	-254.72	-5.95
TASRL	12.08	-333.03	-5.95

Note: Smaller value means better performance for all three columns.

Figure 2 shows an example control trajectory of the proposed approach and the baselines at bus 20 and bus 66. Although both the Stable-DDPG and the proposed method restore voltage quickly, Stable-DDPG uses more reactive power at bus 66 and less at bus 20, leading to a suboptimal solution for $F(q)$. On the other hand, the safe gradient flow and the proposed method converge to the same steady state, while the safe gradient flow's convergence is slower.




Fig. 2: Control trajectory of IEEE 123-bus system. The voltage trajectory is shown in the left plot, the middle plot displays the reactive power usage, and the right plot shows the objective function's trajectory.

C. Effect of design parameter of safe gradient flow

We illustrate the effect of hyperparameter α in Figure 3. Smaller α results in a smoother voltage trajectory, as the controller becomes more conservative for the reactive power capacity constraints. The right plot shows the reactive power injection, which is not significantly affected by α when it is far from the capacity limit due to the presence of a transient performance optimizer (bus 16). However, when approaching the capacity limit, smaller α will slow down the rate of change of reactive power injection (bus 66).




Fig. 3: Control trajectory of IEEE 123-bus system with our proposed method using different α . For simplicity, we did not plot $q_{16} = -21.6$.

VI. CONCLUSIONS

We proposed the TASRL framework to optimize transient and steady-state performance simultaneously for voltage control and established formal guarantees for it. The main insight underlying our approach is that, by synthesizing a stable transient policy and a steady-state optimizer within a safe gradient flow framework, the performance of different time scales can be optimized end-to-end. Our proposed method was tested on both IEEE 13-bus and 123-bus systems. The results demonstrate that TASRL not only converges to the steady-state optimal solution but also exhibits superior transient performance compared to existing methods. Future work will (i) extend the theoretical analysis to include the controller dynamics and nonlinear system models and (ii)

generalize our design to handle time-varying and dynamic loads while maintaining the stability guarantees.

REFERENCES

- [1] M. E. Baran and F. F. Wu, "Optimal capacitor placement on radial distribution systems," *IEEE Transactions on Power Delivery*, vol. 4, no. 1, pp. 725–734, 1989.
- [2] M. Farivar, X. Zhu, and L. Chen, "Local voltage control in distribution systems: An incremental control algorithm," in *IEEE International Conference on Smart Grid Communications*, 2015, pp. 732–737.
- [3] N. Li, G. Qu, and M. Dahleh, "Real-time decentralized voltage control in distribution networks," in *52nd Annual Allerton Conference on Communication, Control, and Computing*, 2014, pp. 582–588.
- [4] H. Zhu and H. J. Liu, "Fast local voltage control under limited reactive power: Optimality and stability analysis," *IEEE Transactions on Power Systems*, vol. 31, no. 5, pp. 3794–3803, 2016.
- [5] G. Qu and N. Li, "Optimal distributed feedback voltage control under limited reactive power," *IEEE Transactions on Power Systems*, vol. 35, no. 1, pp. 315–331, 2019.
- [6] Z. Yuan, G. Cavraro, M. K. Singh, and J. Cortés, "Learning provably stable local volt/var controllers for efficient network operation," *IEEE Transactions on Power Systems*, 2022, submitted.
- [7] J. Feng, Y. Shi, G. Qu, S. H. Low, A. Anandkumar, and A. Wierman, "Stability constrained reinforcement learning for decentralized real-time voltage control," arXiv 2209.07669, 2022.
- [8] 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.
- [9] Y. Shi, G. Qu, S. Low, A. Anandkumar, and A. Wierman, "Stability constrained reinforcement learning for real-time voltage control," in *American Control Conference*. IEEE, 2022.
- [10] J. Feng, Y. Shi, G. Qu, S. H. Low, A. Anandkumar, and A. Wierman, "Stability constrained reinforcement learning for real-time voltage control in distribution systems," arXiv preprint arXiv:2209.07669, 2022.
- [11] W. Cui, J. Li, and B. Zhang, "Decentralized safe reinforcement learning for inverter-based voltage control," *Electric Power Systems Research*, vol. 211, p. 108609, 2022.
- [12] A. Allibhoy and J. Cortés, "Control barrier function based design of gradient flows for constrained nonlinear programming," arXiv preprint arXiv:2204.01930, 2022.
- [13] F. Blanchini and S. Miani, *Set-Theoretic methods in control*. Boston, MA: Birkhäuser, 2007.
- [14] T. Gurriet, A. Singletary, J. Reher, L. Ciarletta, E. Feron, and A. Ames, "Towards a framework for realizable safety critical control through active set invariance," in *ACM/IEEE International Conference on Cyber-Physical Systems*. IEEE, 2018, pp. 98–106.
- [15] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge University Press, 2004.
- [16] 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.
- [17] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," arXiv:1707.06347, 2017.
- [18] J. Schulman, S. Levine, P. Abbeel, M. Jordan, and P. Moritz, "Trust region policy optimization," in *International Conference on Machine Learning*. PMLR, 2015, pp. 1889–1897.
- [19] D. S. Bernstein, *Matrix mathematics: theory, facts, and formulas*, 2nd ed. Princeton University Press, 2009.
- [20] D. Bertsekas, *Nonlinear programming*. Athena Scientific, 1999.
- [21] H. K. Khalil, *Nonlinear systems; 3rd ed*. Prentice-Hall, 2002.
- [22] W. Cui, Y. Jiang, B. Zhang, and Y. Shi, "Structured neural-PI control for networked systems: Stability and steady-state optimality guarantees," arXiv preprint arXiv:2206.00261, 2022.
- [23] 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.
- [24] 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.
- [25] K. Turitsyn, P. Sulc, S. Backhaus, and M. Chertkov, "Options for control of reactive power by distributed photovoltaic generators," *Proceedings of the IEEE*, vol. 99, no. 6, pp. 1063–1073, 2011.