

Stability-Constrained Learning for Frequency Regulation in Power Grids With Variable Inertia

Jie Feng¹, *Graduate Student Member, IEEE*, Manasa Muralidharan², *Graduate Student Member, IEEE*, Rodrigo Henriquez-Auba³, *Member, IEEE*, Patricia Hidalgo-Gonzalez⁴, and Yuanyuan Shi¹

Abstract—The increasing penetration of converter-based renewable generation has resulted in faster frequency dynamics, and low and variable inertia. As a result, there is a need for frequency control methods that are able to stabilize a disturbance in the power system at timescales comparable to the fast converter dynamics. This letter proposes a combined linear and neural network controller for inverter-based primary frequency control that is stable at time-varying levels of inertia. We model the time-variance in inertia via a switched affine hybrid system model. We derive stability certificates for the proposed controller via a quadratic candidate Lyapunov function. We test the proposed control on a 12-bus 3-area test network, and compare its performance with a base case linear controller, optimized linear controller, and finite-horizon Linear Quadratic Regulator (LQR). Our proposed controller achieves faster mean settling time and over 50% reduction in average control cost across 100 inertia scenarios compared to the optimized linear controller. Unlike LQR which requires complete knowledge of the inertia trajectories and system dynamics over the entire control time horizon, our proposed controller is real-time tractable, and achieves comparable performance to LQR.

Index Terms—Power systems, data-driven control, time-varying systems.

I. INTRODUCTION

RENEWABLE energy has experienced a rapid growth, and is predicted to constitute one-third of the global generation mix by 2025 [1]. Unlike conventional generators, renewables are interfaced with the grid via power

electronic inverters, which lack rotational inertia. In synchronous generator-dominated grids, aggregated rotational inertia slows down the system frequency response in the event of an imbalance between power supply and demand [2]. This allows controllers sufficient response time to restore frequency to its nominal value. In contrast, renewable-dominated grids are characterized by low and time-varying inertia, and fast frequency dynamics [3]. This creates a need for real-time tractable, and fast-acting controllers (in the order of milliseconds) [4] that are capable of stabilizing a disturbance without impacting frequency-dependent load shedding and other protection schemes [5].

Several works have proposed virtual inertia and damping allocation as a solution to address the frequency stability challenges in low inertia systems [6], [7], [8]. The proposed approaches vary in the order of their frequency dynamics models and in the metrics used for characterizing controller performance (see [9] for a comprehensive overview). Nevertheless, it's challenging to extend the controllers designed for constant inertia to a time-varying system, given that even exponentially stable subsystems can become unstable with time-varying parameters [10]. Only a few papers have considered frequency control for variable inertia systems [11], [12], [13], [14], [15]. The work in [12] proposes a robust controller that optimizes the worst case system performance via a $\mathcal{H}_\infty$ loop shaping controller that adapts to time-varying frequency and damping in low-inertia systems. The work in [13] explicitly considers the temporal variation in inertia by modeling it as multiplicative and additive noise in the linearized stochastic swing dynamics model, and characterizes controller performance via the $\mathcal{H}_2$ system norm. The work in [11] models the time-varying frequency dynamics as a switched affine hybrid system, with the system switching through different modes representing different levels of inertia. Using this framework, the authors solve a receding-horizon model predictive control problem for dynamic virtual inertia placement. The work in [14] proposes a data-driven virtual inertia controller that learns an optimal linear control gain from the finite-horizon LQR solution that stabilizes the switched system in all inertia modes. The work in [15] expands on [11], [14] to prove the existence of such a stabilizing time-invariant linear controller.

There has been a lot of recent interest in implementing machine learning to address frequency regulation challenges in the power grid (see [16] for a review). However, the key challenge is that standard learning techniques do not provide stability guarantees. Particularly, a neural network controller with a low training loss may actually lead to

Manuscript received 8 March 2024; revised 10 May 2024; accepted 27 May 2024. Date of publication 31 May 2024; date of current version 17 June 2024. This work was supported in part by the UC-National Laboratory In-Residence Graduate Fellowship under Grant L24GF7923; in part by NSF under Grant ECCS-2200692; and in part by the Jacobs School Early Career Faculty Development Award 2023. Recommended by Senior Editor C. Prieur. (*Corresponding author: Jie Feng.*)

Jie Feng and Yuanyuan Shi are with the Department of Electrical Engineering, University of California San Diego, San Diego, CA 92121 USA (e-mail: jason222hh@gmail.com).

Manasa Muralidharan is with the Department of Mechanical and Aerospace Engineering and the Center for Energy Research, University of California San Diego, San Diego, CA 92121 USA.

Rodrigo Henriquez-Auba is with the Grid Planning and Analysis Center, National Renewable Energy Laboratory, Golden, CO 80401 USA.

Patricia Hidalgo-Gonzalez is with the Department of Electrical Engineering, the Department of Mechanical and Aerospace Engineering, and the Center for Energy Research, University of California San Diego, San Diego, CA 92121 USA.

Digital Object Identifier 10.1109/LCSYS.2024.3408068

2475-1456 © 2024 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.

See <https://www.ieee.org/publications/rights/index.html> for more information.

system instabilities (i.e., unbounded frequency deviations) when implemented during testing as observed in [17], [18]. Stability is critical for power system operation because it can lead to catastrophic consequences, e.g., blackouts [19]. To take stability into account for a nonlinear time-varying system with uncertainty, [20] proposes a quadratic constraint based on restricting the partial gradients of the control policy to a bounded (safety) set to guarantee global asymptotic stability. Another class of papers are based on integrating Lyapunov stability constraints into neural-network-based design [18], [21], [22], [23], [24], [25], [26], [27], [28], and enforce strict monotonicity of the policy for stability [18], [21], [23], [24], [26], [27], [28]. In particular, [21], [24], [28] propose a nonlinear controller for primary frequency control in lossless power networks considering nonlinear frequency dynamics. For the proposed candidate Lyapunov function, local asymptotic stability is guaranteed as long as the controller is a monotonically increasing function that passes through the origin. The works [18], [23], [25], [26], [27] propose nonlinear controllers for voltage control considering the linearized distribution grid power flow model, and show that an incremental control with monotonically decreasing instantaneous control functions guarantee stability. For problems where it is hard to analytically derive a Lyapunov function, [17], [22], [29] propose methods for learning a candidate Lyapunov function jointly with a controller, with both parameterized as neural networks. To the best of the authors' knowledge, no existing work has explored learning-based frequency control that guarantees stability under *time-varying* frequency dynamics.

This letter leverages the power of deep learning to optimize the performance of inverter-based frequency control under time-varying inertia while maintaining stability guarantees. The contributions of this letter are three-fold:

- 1) a fast-acting, data-driven controller that considers frequency dynamics with time-varying inertia,
- 2) that satisfies Lyapunov stability conditions for a switched affine hybrid system (i.e., in all inertia modes and switching sequences), and
- 3) that performs comparably to finite-horizon LQR in stabilizing a frequency disturbance in milliseconds (real-time computationally tractable).

The remainder of this letter is organized as follows. In Section II, we introduce the switched affine hybrid system model of frequency dynamics for a system with time-varying inertia, and the finite-horizon LQR formulation for frequency control. In Section III, we expand on the notion of Lyapunov stability for a switched affine hybrid system in all inertia modes, and present our proposed controller. We derive Lyapunov stability constraints for the proposed controller, and propose a method to integrate the derived stability constraints into a learning algorithm. We implement the proposed stability-constrained learning algorithm on a modified Kundur 12-bus 3-area test system and present our simulation's results in Section IV. Lastly, in Section V we conclude and outline future work.

II. MODEL AND PROBLEM FORMULATION

A. Frequency Dynamics as a Hybrid Switching System

Consider a power network described by an undirected graph, $\mathcal{G} = \{\mathcal{N}, \mathcal{E}\}$, where $\mathcal{N} = \{1, \dots, n\}$ represents the buses (nodes) of the power network and $\mathcal{E}$ represents the edges (transmission lines). Assuming identical unit voltage magnitudes, and purely inductive lines, a small-signal approximation

of the swing equation [30] gives us the following linearized dynamics, $\forall i \in \mathcal{N}$:

$$m_i \ddot{\theta}_i(t) + d_i \dot{\theta}_i(t) = p_{in,i}(t) - \sum_{j \in \mathcal{N}} b_{ij}(\theta_i(t) - \theta_j(t)), \quad (1)$$

where θ_i is the voltage phase angle at node i , m_i is the inertia coefficient at node i , d_i is a lumped parameter representing the droop control or frequency damping coefficient at node i , $p_{in,i}$ represents the power input at node i , and b_{ij} is the susceptance of the transmission line between nodes i and j .

Shifts towards greater integration of renewable energy sources introduce variability in the system's inertia. This variability is attributed to the varying ratios of power generation from renewables (and their associated power electronic converters) and conventional generators across different times of the day. We use [14] to model these new dynamics as a switched affine hybrid system to capture the time-variance of the system and put the frequency dynamics in (1) in the compact vector form:

$$\begin{bmatrix} \dot{\theta}(t) \\ \dot{\omega}(t) \end{bmatrix} = \underbrace{\begin{bmatrix} 0 & I \\ -M_q^{-1}L & -M_q^{-1}D \end{bmatrix}}_{A_{q(t)}} \underbrace{\begin{bmatrix} \theta(t) \\ \omega(t) \end{bmatrix}}_{x(t)} + \underbrace{\begin{bmatrix} 0 \\ M_q^{-1} \end{bmatrix}}_{B_{q(t)}} \underbrace{p_{in}(t)}_{u(t)}, \quad (2)$$

where $\theta(t) \in \mathbb{R}^n$ and $\omega(t) \in \mathbb{R}^n$ respectively denote the angle and frequency deviations from their nominal values at all nodes at time t , $x \in \mathbb{R}^{2n}$ corresponds to the stacked state vector of angle and frequency deviations at all nodes at time t , $u = p_{in}(t) \in \mathbb{R}^n$ is the control action at time t , $D = \text{diag}(d_i) \in \mathbb{R}^{n \times n}$ is a diagonal matrix containing the droop/damping coefficients at all nodes, and $L \in \mathbb{R}^{n \times n}$ is the Laplacian of the grid. Further, M_q represents the inertia matrix in mode $q \in \{1, \dots, p\}$ and each $M_q = \text{diag}(m_{i,q}) \in \mathbb{R}^{n \times n}$ is a diagonal matrix with $m_{i,q}$ denoting inertia constant at node i . Each operational mode q in (2) corresponds to a specific inertia value h_q (in seconds), based on the mix of online generators and converters at time t . In particular, inertia of < 2 s represents a renewable-dominated node, $2 - 4$ s represents a hydroelectric generation-dominated node, and $4 - 10$ s represents a thermal generation-dominated node [2]. The inertia coefficients are calculated as $m_{q,i} = \frac{2h_q S_{rated,i}}{\omega_s}$, where $\omega_s = 50$ Hz and $S_{rated,i}$ is the net power rating at node i . Note that x , u , and q are all time-dependent variables; we omit the t in their notation for brevity.

B. Training Set Generation

To train a policy that minimizes frequency deviations and control costs, we generate a training set for the learning-based controller with the following finite-horizon LQR:

$$\min_{u,x} \int_{t=0}^T (x^\top Q x + u^\top R u) dt, \quad (3a)$$

$$\text{s.t. } \dot{x} = A_q x + B_q u \quad \forall t \in [0, T], \quad (3b)$$

$$x(0) = x^{(0)}, \quad (3c)$$

where $x^{(0)}$ is the initial state, states x and control actions u are the decision variables. The objective function includes quadratic costs for frequency deviation and control action over the optimization time horizon, where $Q \geq 0$ and $R > 0$ can be tuned according to the desired control objectives. The finite-horizon LQR has complete knowledge of the dynamics, including the modes, over the entire time horizon. Problem (3) is a Quadratic Programming problem solvable with CVX [31].

III. METHODOLOGY

In this section, we begin by deriving a Lyapunov function for the swing dynamics with time-varying inertia described in (2). Then, we derive the stability condition under this Lyapunov function and present our proposed controller. Finally, we introduce the training method for the proposed algorithm. Figure 1 shows a diagram of the proposed policy.

A. Candidate Lyapunov Function for the Switched System

The first-stage objective is to find a Lyapunov function for the switched affine hybrid system. We propose to *jointly* identify a Lyapunov function $V(x) = x^T P x$ and a *linear* feedback controller $u = Kx$ that can stabilize all the modes. In the next part, we will show how to improve the performance of the linear controller via a neural network residual policy while maintaining the stability guarantee.

Formally, we aim to find a controller $u = Kx$ and a common Lyapunov function $V(x) = x^T P x$, $P \succ 0$ such that the following Lyapunov stability condition is satisfied,

$$(A_q + B_q K)^T P + P(A_q + B_q K) \prec 0, \quad \forall q \in \{1, 2, \dots, p\}, \quad (4)$$

Given that the Lyapunov function $V(x) = x^T P x$ is positive definite since $P \succ 0$ and satisfies the Lie derivative condition (4), the linear feedback controller $u = Kx$ guarantees the stability of the switched affine hybrid system.

However, a notable challenge is that the constraint (4) is *nonconvex* in the decision variables K and P due to the products of bi-linear terms. Thus, we adopt a change of variables following [32, Ch. 7]. Defining $X = P^{-1}$ and $Y = KX$, we formulate the joint Lyapunov function and stable control identification feasibility problem as follows,

$$\min_{X, Y} C, \quad (5a)$$

$$\text{s.t. } A_q X + X A_q^T + B_q Y + Y^T B_q^T \prec 0, \quad \forall q \in \{1, \dots, p\}, \quad (5b)$$

$$X \succ 0, \quad (5c)$$

where C is a constant. This problem is convex and feasible as [15] analytically guarantees it. Suppose (X_s, Y_s) are solutions of (5). Then the feasible control gain K_s is derived as $K_s = Y_s(X_s)^{-1}$ and the Lyapunov function is defined as

$$V = x^T P_s x, \quad P_s = (X_s)^{-1}. \quad (6)$$

Therefore, no matter how many different modes the switched affine system has, we can always find a linear controller and a common Lyapunov function [15].

B. Proposed Controller

Having identified a stable linear controller for the hybrid system and the corresponding Lyapunov function, we now introduce our proposed algorithm. To guarantee stability and improve performance, we parameterize the proposed controller as a *combination* of a stabilizing linear feedback controller and a nonlinear residual $\pi_\psi(\cdot) : \mathbb{R}^{2n} \rightarrow \mathbb{R}^n$ (Fig. 1). The nonlinear residual is parameterized as a neural network with parameter ψ , and its inputs are the current states x . We consider a controller defined as follows:

$$u = Kx + \pi_\psi(x). \quad (7)$$

By combining the linear controller with a neural network residual, we have a dual benefit: it can inherit the guarantee




Fig. 1. The diagram of the proposed controller depicts a combination of a linear controller $Kx(t)$ and a neural network residual $\pi_\psi(\cdot)$, constrained to actions that satisfy the Lyapunov stability conditions specified in (10).

of the linear control while gaining additional flexibility for performance optimization with a neural network.

To ensure the stability of the closed-loop system we define $A_{cl,q} = A_q + B_q K$ and we utilize the Lyapunov function from (6) to derive the algebraic stability constraint, $\forall x \neq 0$:

$$\dot{V} = x^T (A_{cl,q}^T P + P A_{cl,q}) x + 2x^T P B_q \pi_\psi(x) < 0. \quad (8)$$

We introduce short-handed notations $V_{\text{Lin}}(x, q) := x^T (A_{cl,q}^T P + P A_{cl,q}) x$ and $g(x, q) = 2B_q^T P x$. Thus, (8) can be compactly written as,

$$V_{\text{Lin}}(x, q) + g^T(x, q) \pi_\psi(x) < 0, \quad \forall x \neq 0. \quad (9)$$

The switching system is asymptotically stable with the proposed policy if (9) holds for all time t . We summarize the main result as the following Theorem.

Theorem 1: Consider a controller defined as $u_\psi(x) = Kx + \Pi[\pi_\psi(x)]$, where $u = Kx$ is a linear controller that can stabilize the switching system (2), and $\pi_\psi(x)$ is a neural network controller parameterized by ψ . For all $x \neq 0$, define the projection operation $\Pi[\pi_\psi(x)]$ as follows (10),

$$\lambda^* = \left\lceil \frac{g^T(x, q) \pi_\psi(x) + V_{\text{Lin}}(x, q) + \epsilon}{g^T(x, q) g(x, q)} \right\rceil^+, \quad (10a)$$

$$\Pi[\pi_\psi(x)] = \pi_\psi(x) - \lambda^* g(x, q), \quad (10b)$$

where $\epsilon > 0$ is a sufficiently small constant. For $x = 0$, the projection is $\Pi[\pi_\psi(0)] = 0$. Then the closed-loop system (2) is asymptotically stable with respect to the origin for arbitrary switching signal $q(t) : \mathbb{R}^+ \mapsto \{1, \dots, p\}$.

Proof: Consider the following convex programming,

$$\Pi[\pi_\psi(x)] = \underset{\xi}{\operatorname{argmin}} \frac{1}{2} \|\xi - \pi_\psi(x)\|^2, \quad (11a)$$

$$\text{s.t. } V_{\text{Lin}}(x, q) + g^T(x, q) \xi \leq -\epsilon. \quad (11b)$$

Given that K is a stable linear controller for the system, with the common Lyapunov function $V = x^T P x$, thus for any $q \in \{1, \dots, p\}$ $V_{\text{Lin}}(x, q) < 0, \forall x \neq 0$. Therefore, the project problem (11) is always feasible with a solution $\xi = 0$. Consider the Lagrangian of (11)

$$\mathbb{L}(\xi, \lambda) = \frac{1}{2} \|\xi - \pi_\psi(x)\|^2 + \lambda (V_{\text{Lin}}(x, q) + g^T(x, q) \xi + \epsilon),$$

The optimal primal and dual solutions (ξ^*, λ^*) satisfy the Karush-Kuhn-Tucker (KKT) conditions,

$$\nabla_\xi \mathbb{L} = \xi^* - \pi_\psi(x) + \lambda^* g(x, q) = 0, \quad (12a)$$

$$\lambda^* (V_{\text{Lin}}(x, q) + g^T(x, q) \xi^* + \epsilon) = 0, \quad (12b)$$

$$V_{\text{Lin}}(x, q) + g^T(x, q) \xi \leq -\epsilon, \lambda^* \geq 0. \quad (12c)$$

Thus, (10b) is a direct result from (12a). Substituting (10b) in (12b) gives (10a). As a result, (10) provides an optimal solution of (11). This solution guarantees that $V_{\text{Lin}}(x, q) + g^\top(x, q)\Pi[\pi_\psi(x)] < 0, \forall x \neq 0$ holds for all time t , which leads to asymptotic stability with respect to the origin. ■

Therefore, the final designed policy reads as follows,

$$u_\psi(x) = Kx + \Pi[\pi_\psi(x)], \quad (13)$$

where $\Pi[\pi_\psi(x)]$ is given by (10).

Remark 1: To analyze the impact of linearization error on stability, consider the following nonlinear dynamics,

$$\dot{x} = A_q x + B_q u + \eta(x),$$

where $\eta(x)$ is the linearization error depending on state. With the same Lyapunov function $V = x^\top P x$, the stability condition for the nonlinear dynamics reads as follows,

$$\dot{V} = x^\top \underbrace{\left(A_{cl,q}^\top P + P A_{cl,q} \right)}_{V_{\text{Lin}}(x,q) + g^\top(x,q)\xi \leq -\epsilon} x + 2x^\top P B_q \pi_\psi(x) + 2x^\top P \eta(x).$$

Therefore, our stability constraint in (11b) guarantees stability of the nonlinear dynamics in a local region around the origin when the linearization error is small.

C. Stability-Constrained Learning for Frequency Regulation

In Section III-A, we introduced feasible solutions for the feedback control gain K and the Lyapunov function obtained by solving (5). However, the direct solution of (5) often leads to suboptimal outcomes. Indeed, it has been observed in previous studies [24] that neural network-based controllers can reduce the control cost by 30% and shorten the frequency recovery time by a third compared to a linear control policy. To address this, we propose an iterative approach in this section to jointly optimize the neural network residual controller, the linear feedback control gain, and the Lyapunov function. The proposed training method bypasses the need for hand-tuning, as for the BMI algorithm [15] or the feature selection required for a regression-based linear controller [14].

With the training set (x_{lqr}, u_{lqr}) generated through finite-horizon LQR by solving (3), we optimize the proposed controller to mimic the behavior of the LQR controller. Considering the distinct parameterization of the linear controller and the neural network, we decompose the optimization of the proposed controller into two sub-problems: (1) optimizing the nonlinear residual $\pi_\psi(x)$ and (2) optimizing the linear feedback controller Kx and the Lyapunov stability certificate $V(x)$. The optimization problem for ψ is

$$\min_{\psi} \mathcal{L}_\psi = \|u_\psi(x_{lqr}) - u_{lqr}\|, \quad (14)$$

where u_ψ is defined in (13) and $\|\cdot\|$ is the Frobenius norm.

Then, we fix ψ and optimize the linear controller and Lyapunov function. This sub-problem is formulated as an optimization problem with Lyapunov stability constraints. We use the Cholesky decomposition [33] of P to enforce the positive definiteness of the Lyapunov function, i.e., $P = QQ^\top$, where Q is a lower triangular matrix with positive diagonal entries. The stability constraint (4) is formulated as a soft penalty and verified post-training. We parameterize the

Algorithm 1 Stability-Constrained Learning

Ensure: NN $\pi_\psi(x)$; K, Q ; dataset (x_{lqr}, u_{lqr}) ; feasible solution K_s, P_s ; epoch number N_{ep} , batch number N_b , batch size N_s , training step for linear controller N_l ; constants c_1, c_2, c_3 , learning rates η_1, η_2 .

- 1: Initialize K, Q with K_s, P_s ;
- 2: **for** $i = 0$ to N_{ep} **do**
- 3: **for** $j = 0$ to N_b **do**
- 4: Randomly sample N_s pairs from (x_{lqr}, u_{lqr}) ;
- 5: Update ψ by $\psi = \psi - \eta_1 \nabla_\psi \mathcal{L}_\psi$ in (14);
- 6: **for** step = 0 to N_l **do**
- 7: Update: $K = K - \eta_2 \nabla_K \mathcal{L}_{(K,Q)}, Q = Q - \eta_2 \nabla_Q \mathcal{L}_{(K,Q)}$ in (15).
- 8: **end for**
- 9: **end for**
- 10: **end for**

Cholesky decomposition matrix Q and the linear controller gain matrix K as the learnable parameters.

$$\min_{K,Q} \mathcal{L}_{(K,Q)} = c_1 \|K\| + c_2 \|Kx_{lqr}\| + \|u_\psi(x_{lqr}) - u_{lqr}\|$$

$$+ c_3 \sum_{i=0}^{2n} \sum_q \max\left(0, \text{eig}_i(A_{cl,q}^\top P + P A_{cl,q})\right),$$

$$P = QQ^\top, \quad Q \text{ is lower triangular}, \quad (15a)$$

$$Q_{ii} \in \mathbb{R}, Q_{ii} > 0, \forall i \in \{1, \dots, n\}, \quad (15b)$$

where ψ is fixed, only K and Q are optimized. The coefficients c_1, c_2, c_3 are objective weights. The second row for $\mathcal{L}_{(K,Q)}$ is a summation of all positive eigenvalues of $A_{cl,q}^\top P + P A_{cl,q} < 0, \forall q \in \{1, \dots, p\}$ to penalize violations of the Lyapunov stability constraint (4). We deploy a warm start using the feasible solution of (5) for K and Q , where Q is decomposed from P_s , such that we start with a stabilizing but suboptimal solution, improving learning efficiency. We minimize the norm of the control gain $\|K\|$ and the linear control action $\|Kx_{lqr}\|$ to avoid the large overshoot observed in the feasible solution. With $\pi_\psi(x) = 0$, the same procedure solely optimizes the linear controller with its Lyapunov function, termed *Linear-opt*.

The nonlinear and linear policies are updated iteratively using gradient descent until convergence or the maximum step limit is reached, as outlined in Algorithm 1. It's worth noting that while this iterative algorithm does not guarantee global optimality, the numerical experiments show that the trained controller (13) achieves performance comparable to the LQR controller with stability guarantees.

IV. SIMULATIONS

In this section, we demonstrate the effectiveness of the proposed algorithm via numerical experiments. We compare against the LQR, the linear controller, i.e., (*Linear*), and the optimized linear controller, i.e., (*Linear-opt*).

A. Experimental Setup

We implement the proposed controller in the modified Kundur 12-bus 3-region network [34] (cf. Fig. 2), with discretized network dynamics using zero-order hold.

The inverter-based generation resources are deployed for frequency control. We consider 9 inertia modes, with the inertia in each mode $h_q \in \{0.5, 1, 1.5, 2, 2.5, 3, 3.5, 5, 9\}$ s. For simplicity, we assume the same normalized inertia coefficient for all nodes in a given mode q , i.e., $m_{q,i} = m_q = h_q, \forall i$. The droop coefficient is 0.5 at all nodes and the control frequency is 1 kHz. We solve the finite-horizon LQR




Fig. 2. 230kV/100 MVA Kundur 12-bus 3-area test network with $0.0001 + 0.001$ p.u. line impedance.

TABLE I
PERFORMANCE OF LQR, BASE CASE (LINEAR), OPTIMIZED LINEAR (LINEAR-OPT), AND PROPOSED CONTROLLER ACROSS 100 SCENARIOS

Method	Settling time (ms)		Overshoot (Hz)		Avg Cost	
	Mean	Std	Mean	Std	Mean	Std
LQR	122.8	77.1	0.097	0.033	99.0	76.18
Proposed	147.7	84.8	0.096	0.035	107.1	82.39
Linear	554.8	24.2	0.269	0.021	4497.4	130.5
Linear-opt	605.2	48.8	0.083	0.034	244.7	92.97

problem (3) using $Q = \text{diag}(0, 5e^4 I)$ and $R = I$ and collect 200 trajectories, assuming both frequency and actions in per unit (p.u.). The nonlinear residual is modeled using a three-layer fully connected neural network with 300 and 400 hidden units, respectively. The network takes the state x as input and outputs the control actions (real power set-points).

Online inertia modes are estimated by a neural network pre-trained with the collected trajectories and used by (10) for the projection, where the input to the neural network is two consecutive frequency observations $\omega(t), \omega(t+1)$ and the corresponding action $u(t)$. The estimator achieves 90.64% accuracy for inertia mode classification. We optimize the proposed controller by Algorithm 1 with $c_1 = 0.1, c_2 = 0.01, c_3 = 500, \eta_1 = 0.001, \eta_2 = 0.01, N_{ep} = N_b = 300, N_s = 256, N_l = 5$, and compare its performance against LQR, *Linear* and *Linear-opt* in 100 distinct frequency deviation scenarios, each with a time horizon of 1 s (1000 steps) and a random initial frequency deviation at each bus sampled from a uniform distribution $\mathcal{U}_{[-0.3, 0.3]}$ Hz. The scenarios commence in random operational modes $q(0) \in \{1, \dots, 9\}$. Subsequently, based on a uniform distribution, the inertia of the system either remains constant, increases or decreases every 0.1 s.

We utilize three key metrics to evaluate our controller: (i) **settling time**, defined as the average duration for the controller to reduce frequency deviations to under 0.01 Hz; (ii) **overshoot**, which is the average maximum frequency deviation observed in each scenario; and (iii) **average cost** (Avg Cost) $\frac{1}{N} \sum_{i=0}^{N=100} \sum_{t=0}^{T=1000} x_i^\top Q x_i + u_i^\top R u_i$.

B. Results

In 100 scenarios (cf. Table I), the proposed controller matches LQR's performance and outperforms *Linear* and *Linear-opt* in settling time and control cost due to the flexibility of the nonlinear residual.

Figures 3 and 4 illustrate state and control trajectories for one inertia switching scenario, where bus 9 has an initial frequency deviation of -0.3 Hz and other buses have random deviations. From Figure 3, all controllers stabilize bus 9 despite the inertia change. Specifically, *Linear* induces a




Fig. 3. State and control trajectories for LQR, the proposed controller, *Linear*, and *Linear-opt* at bus 9 with a zoomed-in view for 0-0.2 s. The background color represents the inertia modes.




Fig. 4. State and control trajectories of all buses with the proposed controller with a zoomed-in view for 0-0.1 s.

TABLE II
PERFORMANCE OF LQR (LQR-C) AND PROPOSED CONTROLLER (PROPOSED-C) WITH HARDWARE CONSTRAINTS

Method	Settling time (ms)		Overshoot (Hz)		Avg Cost	
	Mean	Std	Mean	Std	Mean	Std
LQR-C	143.7	84.7	0.114	0.039	106.2	81.1
Proposed-C	152.2	79.1	0.124	0.043	141.7	112.3

relatively large overshoot and large control actions, while *Linear-opt* reduces these but has a longer frequency restoration time (≈ 1 s). Moreover, the control action of both linear controllers might fail to converge to the optimal solution once the frequency is restored. The proposed controller, similar to LQR, achieves fast frequency recovery within 0.1 s and low cost across all buses (cf. Figures 3 and 4), demonstrating the efficiency of our algorithm.

C. Controller Constraints

We explore hardware constraints and nonlinear power flow dynamics' impact. Integrating action constraints, $u \leq \bar{u}$, into both LQR formulation and projection (10), we train the controller with LQR solutions that incorporate these constraints. Table II outlines the constrained controller's performance, where $|u_i| \leq 0.5$ p.u. for all $i \in \mathcal{N}$. Despite a slight increase in control cost and settling time compared to LQR, our controller still achieves fast frequency recovery.

V. CONCLUSION

We propose a stability-constrained data-driven controller for frequency regulation with time-varying inertia. Our method integrates a linear controller with a neural network-based nonlinear residual, where the linear controller can stabilize the switching system with a joint Lyapunov function. During

training, the linear controller, the corresponding Lyapunov function, and the nonlinear residual are optimized iteratively in an end-to-end manner. The stability of the closed-loop system is further enforced by projecting the nonlinear residual to guarantee the Lyapunov condition. Thanks to the nonlinear residual, the policy can approximate the LQR solution and achieve a comparable performance. Although we successfully identified a valid Lyapunov function and a linear feedback controller, jointly optimizing these with stability and optimality guarantees remains challenging and is a key area for future research. Future work includes adaptation to nonlinear inverter dynamics, delay-aware controller design, decentralizing the controller to reduce communication needs, and improving robustness to parameter measurement errors and variability.

REFERENCES

- [1] *Electricity Market Report 2023*, Int. Energy Agency, Paris, France, 2023.
- [2] P. Kundur, *Power System Stability and Control*. New York, NY, USA: McGraw-Hill Inc., 2002.
- [3] A. Ulbig, T. S. Borsche, and G. Andersson, "Impact of low rotational inertia on power system stability and operation," *IFAC Proc. Vol.*, vol. 47, no. 3, pp. 7290–7297, 2014.
- [4] A. Ulbig, T. Rinke, S. Chatzivasileiadis, and G. Andersson, "Predictive control for real-time frequency regulation and rotational inertia provision in power systems," in *Proc. 52nd IEEE Conf. Decis. Control*, 2013, pp. 2946–2953.
- [5] U. Markovic, O. Stanojev, P. Aristidou, E. Vrettos, D. Callaway, and G. Hug, "Understanding small-signal stability of low-inertia systems," *IEEE Trans. Power Syst.*, vol. 36, no. 5, pp. 3997–4017, Sep. 2021.
- [6] D. Gross, S. Bolognani, B. Poolla, and F. Dörfler, "Increasing the resilience of low-inertia power systems by virtual inertia and damping," in *Proc. IREP Symp.*, 2017, p. 64.
- [7] B. K. Poolla, D. Groß, and F. Dörfler, "Placement and implementation of grid-forming and grid-following virtual inertia and fast frequency response," *IEEE Trans. Power Syst.*, vol. 34, no. 4, pp. 3035–3046, Jul. 2019.
- [8] A. Venkatraman, U. Markovic, D. Shchetinin, E. Vrettos, P. Aristidou, and G. Hug, "Improving dynamic performance of low-inertia systems through eigensensitivity optimization," *IEEE Trans. Power Syst.*, vol. 36, no. 5, pp. 4075–4088, Sep. 2021.
- [9] F. Milano, F. Dörfler, G. Hug, D. J. Hill, and G. Verbič, "Foundations and challenges of low-inertia systems (invited paper)," in *Proc. Power Syst. Comput. Conf. (PSCC)*, 2018, pp. 1–25.
- [10] L. Vu, D. Chatterjee, and D. Liberzon, "Input-to-state stability of switched systems and switching adaptive control," *Automatica*, vol. 43, no. 4, pp. 639–646, 2007.
- [11] P. Hidalgo-Gonzalez, D. S. Callaway, R. Dobbe, R. Henriquez-Auba, and C. J. Tomlin, "Frequency regulation in hybrid power dynamics with variable and low inertia due to renewable energy," in *Proc. IEEE Conf. Decis. Control*, 2018, pp. 1592–1597.
- [12] G. S. Misyris, S. Chatzivasileiadis, and T. Weckesser, "Robust frequency control for varying inertia power systems," in *Proc. IEEE PES Innovat. Smart Grid Technol. Conf. Eur.*, 2018, pp. 1–6.
- [13] Y. Guo and T. H. Summers, "A performance and stability analysis of low-inertia power grids with stochastic system inertia," in *Proc. Am. Control Conf.*, 2019, pp. 1965–1970.
- [14] P. Hidalgo-Gonzalez, R. Henriquez-Auba, D. S. Callaway, and C. J. Tomlin, "Frequency regulation using data-driven controllers in power grids with variable inertia due to renewable energy," in *Proc. IEEE Power Energy Soc. Gener. Meet.*, 2019, pp. 1–5.
- [15] P. Srivastava, P. Hidalgo-Gonzalez, and J. Cortés, "Learning constant-gain stabilizing controllers for frequency regulation under variable inertia," *IEEE Control Syst. Lett.*, vol. 6, pp. 3056–3061, 2022.
- [16] 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 Trans. Smart Grid*, vol. 13, no. 4, pp. 2935–2958, Jul. 2022.
- [17] W. Cui and B. Zhang, "Lyapunov-regularized reinforcement learning for power system transient stability," *IEEE Control Syst. Lett.*, vol. 6, pp. 974–979, 2022.
- [18] Y. Shi, G. Qu, S. Low, A. Anandkumar, and A. Wierman, "Stability constrained reinforcement learning for real-time voltage control," in *Proc. Am. Control Conf.*, 2022, pp. 2715–2721.
- [19] 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.
- [20] M. Jin and J. Lavaei, "Stability-certified reinforcement learning: A control-theoretic perspective," *IEEE Access*, vol. 8, pp. 229086–229100, 2020.
- [21] Y. Jiang, W. Cui, B. Zhang, and J. Cortés, "Stable reinforcement learning for optimal frequency control: A distributed averaging-based integral approach," *IEEE Open J. Control Syst.*, vol. 1, pp. 194–209, 2022.
- [22] T. Huang, S. Gao, and L. Xie, "A neural Lyapunov approach to transient stability assessment of power electronics-interfaced networked microgrids," *IEEE Trans. Smart Grid*, vol. 13, no. 1, pp. 106–118, Jan. 2022.
- [23] J. Feng, Y. Shi, G. Qu, S. H. Low, A. Anandkumar, and A. Wierman, "Stability constrained reinforcement learning for Decentralized real-time voltage control," *IEEE Trans. Control Netw. Syst.*, early access, Dec. 1, 2023, doi: [10.1109/TCNS.2023.3338240](https://doi.org/10.1109/TCNS.2023.3338240).
- [24] W. Cui, Y. Jiang, and B. Zhang, "Reinforcement learning for optimal primary frequency control: A Lyapunov approach," *IEEE Trans. Power Syst.*, vol. 38, no. 2, pp. 1676–1688, Mar. 2023.
- [25] J. Feng, W. Cui, J. Cortés, and Y. Shi, "Bridging transient and steady-state performance in voltage control: A reinforcement learning approach with safe gradient flow," *IEEE Control Syst. Lett.*, vol. 7, pp. 2845–2850, 2023.
- [26] Z. Yuan, G. Cavarro, M. K. Singh, and J. Cortés, "Learning provably stable local volt/Var controllers for efficient network operation," *IEEE Trans. Power Syst.*, vol. 39, no. 1, pp. 2066–2079, Jan. 2023.
- [27] Z. Yuan, G. Cavarro, and J. Cortés, "Constraints on OPF surrogates for learning stable local volt/Var controllers," *IEEE Control Syst. Lett.*, vol. 7, pp. 2533–2538, 2023.
- [28] W. Cui, Y. Jiang, B. Zhang, and Y. Shi, "Structured neural-PI control with end-to-end stability and output tracking guarantees," in *Proc. 37th Conf. Neural Inf. Process. Syst.*, 2024, pp. 1–24.
- [29] Y.-C. Chang, N. Roohi, and S. Gao, "Neural Lyapunov control," in *Proc. 33rd Conf. Neural Inf. Process. Syst.*, 2022, pp. 1–10.
- [30] B. K. Poolla, S. Bolognani, and F. Dörfler, "Optimal placement of virtual inertia in power grids," *IEEE Trans. Autom. Control*, vol. 62, no. 12, pp. 6209–6220, Dec. 2017.
- [31] A. Agrawal, R. Verschuere, S. Diamond, and S. Boyd, "A rewriting system for convex optimization problems," *J. Control Decis.*, vol. 5, no. 1, pp. 42–60, 2018.
- [32] S. Boyd, L. El Ghaoui, E. Feron, and V. Balakrishnan, *Linear Matrix Inequalities in System and Control Theory*. Philadelphia, PA, USA: Soc. Ind. Appl. Math., 1994.
- [33] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes in C (2nd ed.): The Art of Scientific Computing*. Cambridge, MA, USA: Cambridge Univ. Press, 1992.
- [34] T. S. Borsche, T. Liu, and D. J. Hill, "Effects of rotational inertia on power system damping and frequency transients," in *Proc. IEEE Conf. Decis. Control*, 2015, pp. 5940–5946.