Compositional Neural Certificates for Networked Dynamical Systems

Songyuan Zhang SZHANG21@MIT.EDU
Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA, USA.

Yumeng Xiu YXIU2@ANDREW.CMU.EDU
Department of Mechanical Engineering, Carnegie Mellon University, Pittsburgh, PA, USA.

Guannan Qu GQU@ANDREW.CMU.EDU
Department of Electrical and Computer Engineering, Carnegie Mellon University, Pittsburgh, PA, USA.

Chuchu Fan CHUCHU@MIT.EDU
Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA, USA.

Editors: N. Matni, M. Morari, G. J. Pappas

Abstract

Developing stable controllers for large-scale networked dynamical systems is crucial but has long been challenging due to two key obstacles: certifiability and scalability. In this paper, we present a general framework to solve these challenges using compositional neural certificates based on ISS (Input-to-State Stability) Lyapunov functions. Specifically, we treat a large networked dynamical system as an interconnection of smaller subsystems and develop methods that can find each subsystem a decentralized controller and an ISS Lyapunov function; the latter can be collectively composed to prove the global stability of the system. To ensure the scalability of our approach, we develop generalizable and robust ISS Lyapunov functions where a single function can be used across different subsystems and the certificates we produced for small systems can be generalized to be used on large systems with similar structures. We encode both ISS Lyapunov functions and controllers as neural networks and propose a novel training methodology to handle the logic in ISS Lyapunov conditions that encodes the interconnection with neighboring subsystems. We demonstrate our approach in systems including Platoon, Drone formation control, and Power systems. Experimental results show that our framework can reduce the tracking error up to 75% compared with RL algorithms when applied to large-scale networked systems. [1]

Keywords: Neural Certificates, ISS Lyapunov Functions, Networked Dynamical Systems

1. Introduction

Large-scale networked dynamical systems play an important role across a wide spectrum of real-world applications, including power grids (Zhao et al., 2014), vehicle platoons (Stankovic et al., 2000), drone swarms (Tedrake, 2022), transportation networks (Varaiya, 2013), etc. The control, and in particular stabilization, of such networked systems has long been recognized as a challenging problem as the dimension of the state and input spaces of such networked systems is usually very high, and existing methods often suffer from the "curse-of-dimensionality" (Powell, 2007).

Classical approaches for stabilization of dynamical systems include LQR (Linear Quadratic Regulator) for linear systems (Khalil et al., 1996; Dullerud and Paganini, 2013). For nonlinear systems, certificates like Lyapunov functions can be used to guide the search for a stabilizing controller

1. Project website: https://mit-realm.github.io/neuriss-website/. The appendix can be found on the project website.

and certify the stability of the closed-loop system (Slotine and Li, 1991). However, certificates are usually constructed on a case-by-case basis. While there exist approaches like SOS (Sum-Of-Squares) that can construct certificates for general classes of nonlinear systems (Parrilo, 2000), they are not scalable to large-scale networked dynamical systems, since the number of polynomial coefficients in an SOS program grows exponentially w.r.t. the dimension of system (Parrilo, 2000).

A recent line of work parameterizes control certificates (*e.g.* Lyapunov functions, barrier functions) and controllers as neural networks (NNs) and learns them jointly from data (Chang et al., 2019; Jin et al., 2020; Chow et al., 2018). They have been successfully applied to nonlinear systems and achieved good performance on complex control tasks (Richards et al., 2018; Manek and Kolter, 2019) thanks to the representation power of NNs. However, the existing works mainly focus on single-agent systems with relatively small state space ($\leq$ 10 dimensions) or multi-agent systems without coupled dynamics (Chang et al., 2019; Qin et al., 2021b). Applying these approaches to large-scale networked systems can be challenging due to the exponential growth of the sample complexity and the hardness of training NNs with large input spaces. Despite the challenge, many networked dynamics often contain sparse network structures that can be exploited to help training, and the question we try to answer in this paper is: *can we exploit network structure to learn neural certificates and stabilize large scale networked systems in a scalable and effective manner?*

To answer the question, we view the large networked dynamical system as a group of smaller subsystems interconnected through a graph. Instead of learning a single certificate for the entire system, we find a decentralized ISS (Input-to-State Stability) Lyapunov function (Sontag, 2013; Liu et al., 2011; Jiang and Liu, 2018) and a decentralized controller for each subsystem. Although ISS Lyapunov functions have been known for decades, it is not straightforward to adapt them as neural certificates for the stabilization of large networked systems due to the following reasons: 1) Existing ISS Lyapunov theory requires checking a condition involving global information of the networked system (Liu et al., 2011) thus is not entirely decentralized; 2) Existing ISS Lyapunov theory requires finding different ISS Lyapunov functions for each subsystem and therefore is computationally expensive for systems with many subsystems; 3) Each subsystem, as well as the corresponding ISS Lyapunov function, are intertwined with neighboring subsystems and therefore cannot be learned straightforwardly as Lyapunov functions for a single system such as that in Chang et al. (2019).

To tackle these challenges, we propose **Neur**al **ISS** Lyapunov functions (NeurISS) that make the following **contributions:** 1) We show that the ISS Lyapunov functions only need to satisfy a *local* condition involving local information from neighboring subsystems in order to collectively constitute a compositional certificate to certify the stability of the entire dynamical system (Lemma 1); 2a) We prove that under certain conditions, the compositional certificate for a small networked system can be generalized to be used on a more extensive system that has a similar structure without re-training (Lemma 2), which improves the scalability of the proposed approach as one can reduce a large training task to a smaller training task with a smaller network size; 2b) We extend the notion of the ISS Lyapunov function to robust ISS Lyapunov function for control-affine systems (Lemma 3) so that similar subsystems that have different parameters can share the same ISS Lyapunov functions, which not only reduces the number of ISS Lyapunov functions we need to learn for large-scale networked systems but also improves the robustness of the learned results against model uncertainties. 3) Furthermore, we develop a novel approach to encode the ISS logic condition that intertwines neighboring subsystems into the training loss function (Section 4).

We demonstrate NeurISS using three examples - Power systems, Platoon, and Drone formation control, and show that NeurISS can find certifiably stable controllers for networks of size up to

100 subsystems. Compared with centralized neural certificate approaches, NeurISS reaches similar results in small-scale systems, and can generalize to large-scale systems that centralized approaches cannot scale up to. Compared with LQR, NeurISS can deal with strong coupled networked systems like the microgrids, and reaches smaller tracking errors on both small and large-scale systems. Compared with RL (PPO, LYPPO, MAPPO), our algorithm achieves similar or smaller tracking errors in small systems, and can hugely reduce the tracking errors in large systems (up to 75%).

Related Work. Safe machine learning, neural certificates, and reinforcement learning all have rich literature. Due to space limits, we only mention the most related works.

Neural Certificates. Mostly related is the line of work on learning neural certificates. This line of work focuses on searching for a controller together with a certificate that guarantees the soundness of the controller. Such neural certificates include Lyapunov-like functions for stability guarantees (Chang et al., 2019; Jin et al., 2020; Richards et al., 2018; Manek and Kolter, 2019; Abate et al., 2020; Dawson et al., 2021; Gaby et al., 2021), barrier functions for safety guarantees (Jin et al., 2020; Qin et al., 2021b; Xiao et al., 2021; Peruffo et al., 2021; Srinivasan et al., 2020), contraction metrics for tracking guarantees (Sun et al., 2020; Chou et al., 2021), etc. Through learning proof of the correctness of the controllers, these approaches address the concerns about the safety, stability, and reliability of the controllers on a large variety of tasks, including precision quadrotor flight through turbulence (Sun et al., 2020), walking under model uncertainties (Castañeda et al., 2021), tracking with high-dimensional dynamics (Chou et al., 2021), and safe decentralized control of multi-agent systems (Qin et al., 2021b; Meng et al., 2021). Compared to these works, we learn ISS Lyapunov functions, which are decentralized and scalable to large-scale networked systems. We will compare the proposed approach with such neural certificate approaches in Section 5.

ISS Lyapunov Function. The concept of ISS and ISS Lyapunov function is long established in control theory (Sontag, 2013). ISS Lyapunov function for networked dynamical systems was proposed in Jiang et al. (1996) for a two subsystem case and generalized in Liu et al. (2011, 2012) for multiple subsystems (cf Jiang and Liu (2018) for a review). Compared to these works, our paper builds upon the ISS Lyapunov concept to learn neural certificates for networked systems.

Reinforcement Learning (RL). RL is a popular paradigm in the learning-to-control community with various approaches like (Deep) Q networks (Mnih et al., 2013) , policy optimization (Schulman et al., 2015, 2017) and multi-agent versions of them (Yu et al., 2021) (cf. Sutton and Barto (2018) for a review). However, standard RL is reward-driven and does not formally guarantee stability. Recently, there has been research on learning certificates in the RL process (Berkenkamp et al., 2017; Chow et al., 2018; Cheng et al., 2019; Han et al., 2020; Chang and Gao, 2021; Zhao et al., 2021; Qin et al., 2021a), but none of them considers decentralized compositional certificates for networked systems. As a result, they are not scalable to large-scale networked systems as they lack the ability to deal with the sheer dimensions of the state space and the exponential growth of sample complexity. We will compare our approach with popular RL algorithms in Section 5.

2. Problem Setting

In this paper, we consider the following networked dynamical system involving n subsystems $\mathcal{N} = \{1, 2, \ldots, n\}$. The dynamics of each subsystem are given by

$$\dot{x}_i = f_i(x_i, x_{i_1}, x_{i_2}, \ldots, x_{i_{n_i}}, u_i) = f_i(x_i, x_{\mathcal{N}_i}, u_i) \tag{1}$$

where $x_i \in \mathbb{R}^{d_i}$, $u_i \in \mathbb{R}^{p_i}$ is the state and control inputs of each subsystem i, and $x_{\mathcal{N}_i} = (x_{i_1}, x_{i_2}, \ldots, x_{i_{n_i}})$ is used to denote the states of the neighbors of subsystem i, $\mathcal{N}_i = \{i_1, i_2, \ldots, i_{n_i}\}$ (not including i itself) which affect the dynamics of the subsystem i. We use $x = (x_1, \ldots, x_n)$, $u = (u_1, \ldots, u_n)$, and $f = (f_1, \ldots, f_n)$ to denote the vector of states, actions, and dynamics across all subsystems (*i.e.* the overall system), respectively. We also denote $d = \sum_{i=1}^n d_i$ and $p = \sum_{i=1}^n p_i$ to the dimension of x and u, respectively. Our goal is to design a controller $u = \pi(x)$ such that the closed-loop system is asymptotically stable around a goal set $\mathcal{X}^{\text{goal}}$, formally defined as follows.

Definition 1 *Consider a goal set $\mathcal{X}^{\text{goal}} := \mathcal{X}_1^{\text{goal}} \times \cdots \times \mathcal{X}_n^{\text{goal}}$ where each $\mathcal{X}_i^{\text{goal}}$ is a closed convex subset of $\mathbb{R}^{d_i}$. The closed-loop system is globally asymptotically stable about $\mathcal{X}_i^{\text{goal}}$ if for any initial state $x(0)$, the trajectory $x(t)$ satisfies $\lim_{t\to\infty} \text{dist}(x(t), \mathcal{X}^{\text{goal}}) = 0$, where $\text{dist}(x, \mathcal{X}^{\text{goal}}) := \inf_{y\in\mathcal{X}^{\text{goal}}} \|x - y\|$ is the distance between the point x and the set $\mathcal{X}^{\text{goal}}$.*

Example 1 *A simple networked system is the truck Platoon system with $n+2$ trucks, where the 0-th (leading) truck can drive freely within the speed and acceleration limits, and the $(n+1)$-th (last) truck will be driven in a way so that the total length of the platoon is roughly kept as a pre-defined constant. We assume other trucks are controllable but can only measure the distance to the two trucks directly in front of and behind themselves. We want to control these trucks so that the trucks in the whole platoon are spread evenly. Specifically, for truck $i \in \{1, 2, ..., n\}$, the states are given as $x_i = [p_i^f, p_i^b, v_i]^\top$ and the neighboring subsystems are trucks $\mathcal{N}_i = \{i-1, i+1\}$, where p_i^f is the distance between the i-th truck and the $(i-1)$-th truck, p_i^b is the distance between the i-th truck and the $(i+1)$-th truck, and v_i is the velocity of the i-th truck. The control input is the acceleration of the i-th truck. Therefore, the dynamics of truck i is $\dot{x}_i = [v_{i-1} - v_i, v_i - v_{i+1}, a_i]^\top$. The goal set for each truck i is uniquely defined as the set of states satisfying $p_i^f = p_i^b$.*

Notations. Function $\alpha : [0, \infty) \to [0, \infty)$ is said to be class-$\mathcal{K}$ if α is continuous, strictly increasing, and $\alpha(0) = 0$. Class-$\mathcal{K}$ function α is said to be class-$\mathcal{K}_\infty$ if $\lim_{a\to+\infty} \alpha(a) = +\infty$.

3. Compositional Neural Certificates

3.1. Decentralized Controller and ISS Lyapunov Functions for Networked Systems

Lyapunov functions are widely used to guarantee the stability of dynamical systems (see Appendix A for an introduction). A common paradigm for stabilizing a dynamical system is to jointly search for a controller $u = \pi(x)$ and a Lyapunov function $V(x)$. However, this approach is not scalable for large-scale networked dynamical systems due to the sheer dimension of the state space, and the controller of the form $u = \pi(x)$, which requires global information of the entire network.

To address the issues, our framework NeurISS includes two key components: decentralized controllers and compositional certificates. We consider the class of *decentralized controllers* $u_i = \pi_i(x_i)$, which only needs local information within the small subsystem. Further, we consider *compositional certificates*, that is, instead of finding a single Lyapunov function $V(x)$ for the whole system, we find one Lyapunov function $V_i(x_i)$ for each subsystem i. The individual Lyapunov functions only depend on the subsystem state, which is much smaller in dimension. Further, based on Liu et al. (2011), we provide the following Lemma 1 which shows that when the individual Lyapunov functions V_i satisfy an ISS-style condition, they will certify the stability of the entire dynamical system. Since it is the collection of the ISS Lyapunov functions $\{V_i\}_{i=1}^n$ that certify

the stability of the entire dynamics, we also call such ISS Lyapunov functions $\{V_i\}_{i=1}^n$ as a "compositional" certificate to distinguish them from typical certificates that only contain one Lyapunov function for the entire system. A proof of Lemma 1 is given in Appendix B.

Lemma 1 *Suppose each subsystem has a decentralized controller $u_i = \pi_i(x_i)$ and a continuously differentiable function $V_i(x_i)$. Suppose: (1) For each i, there exists $\mathcal{K}_\infty$ functions $\underline{\alpha}_i, \bar{\alpha}_i$ such that $\underline{\alpha}_i(\mathrm{dist}(x_i, \mathcal{X}_i^{\mathrm{goal}})) \le V_i(x_i) \le \bar{\alpha}_i(\mathrm{dist}(x_i, \mathcal{X}_i^{\mathrm{goal}}))$; (2) For each i, there exists $\alpha_i > 0$ and class-$\mathcal{K}$ functions $\chi_{ij}, j \in \mathcal{N}_i$ satisfying $\chi_{ij}(a) < a, \forall a > 0$, such that $\forall x_i, x_{\mathcal{N}_i}$,*

$$V_i(x_i) \ge \max_{j \in \mathcal{N}_i} \chi_{ij}(V_j(x_j)) \quad \Rightarrow \quad [\nabla V_i(x_i)]^\top f_i(x_i, x_{\mathcal{N}_i}, \pi_i(x_i)) \le -\alpha_i\, V_i(x_i). \tag{2}$$

Then, the closed-loop system under controllers $\pi_1, \ldots, \pi_n$ is globally asymptotically stable around $\mathcal{X}^{\mathrm{goal}}$. Such functions $V_i(x_i), i = 1, \ldots, n$ are called ISS Lyapunov functions.

We note that Lemma 1 is a variant of the result in Liu et al. (2011), in that we explicitly consider the network structure in the dynamics (1). As a result, in the ISS implication condition (2), we need to test $V_i(x_i)$ versus the max of $V_j(x_j)$ over only the neighbors $\mathcal{N}_i$, as opposed to the entire network as in Liu et al. (2011). This effectively makes (2) a condition that can be checked locally at each subsystem. One benefit of the local structure in the implication condition (2) is that it allows us to use certificates from smaller networks to compose certificates for larger networks that consist of blocks of the smaller networks. We will discuss this in detail in Section 3.2. Moreover, our results can also be robustified so a single ISS Lyapunov function can be used across different subsystems and handle uncertain parameters in the dynamics. We will explain this in detail in Section 3.3.

3.2. Network Generalizability

As discussed in Section 3.1, the condition (2) only involves f_i, V_i, and the Lyapunov functions of neighbors $\{V_j\}_{j \in \mathcal{N}_i}$. With such a local architecture, we present the following Lemma 2 that shows the decentralized controllers π_i and ISS Lyapunov functions V_i for a small system can be "ported over" to a larger dynamical system that has a similar symmetric structure to the smaller dynamical system. The proof of Lemma 2 is postponed to Appendix B.

Lemma 2 *Consider a networked dynamical system with node set $\mathcal{N}$, neighborhood sets $\mathcal{N}_i$, and dynamics functions f_i, and suppose there exist decentralized controllers π_i such that the closed-loop dynamical system admits a compositional certificate V_i that satisfies the conditions in Lemma 1 with parameters χ_{ij}, α_i. Suppose there is another dynamical system with node set $\tilde{\mathcal{N}}$, neighborhood sets $\tilde{\mathcal{N}}_j$, and dynamics functions $\tilde{f}_i$. Suppose for each $j \in \tilde{\mathcal{N}}$, there exists a one-to-one map $\tau_j : \{j\} \cup \tilde{\mathcal{N}}_j \to \mathcal{N}$ such that $\tau_j(\tilde{\mathcal{N}}_j) = \mathcal{N}_{\tau_j(j)}$, and $\tilde{f}_j = f_{\tau_j(j)}$. Further, suppose $\forall j, j' \in \tilde{\mathcal{N}}$, $\forall \ell \in \tilde{\mathcal{N}}_j \cap \tilde{\mathcal{N}}_{j'}$, we have $V_{\tau_j(\ell)} = V_{\tau_{j'}(\ell)}$. Then, $\tilde{\pi}_j = \pi_{\tau_j(j)}$ is a stabilizing controller for the new system with compositional certificate $\tilde{V}_j = V_{\tau_j(j)}$.*

Example 2 *For the Platoon system, using Lemma 2, we can prove that a stabilizing controller for a 5-truck system $\pi_i, i = 1, \ldots, 5$ can be generalized to any system with $n > 5$. Let $\tilde{\mathcal{N}}$, $\tilde{\mathcal{N}}_i$, $\tilde{f}_i$ be the subsystems, neighborhood, and dynamics functions of the Platoon system with n trucks, and $\mathcal{N}$, $\mathcal{N}_i$, f_i be those of the system with 5 trucks. Note that in the Platoon system we have $\mathcal{N}_j = \{j-1, j+1\}$ and the same for $\tilde{\mathcal{N}}_j$. We define the one-to-one mapping τ_j in the following way: For the first truck,*

let $\tau_1(0) = 0, \tau_1(1) = 1, \tau_1(2) = 2$. *For the last truck, let* $\tau_n(n-1) = 4, \tau_n(n) = 5, \tau_n(n+1) = 6$. *For other trucks* $j = 2, 3, \ldots, n-1$, *let* $\tau_j(j-1) = 2, \tau_j(j) = 3, \tau_j(j+1) = 4$. *Further, let* $V_2 = V_3 = V_4$ *and* $\pi_2 = \pi_3 = \pi_4$ *in the system with* 5 *trucks. In this way, we can check that* $\forall j, j' \in \tilde{\mathcal{N}}$, $\forall \ell \in \tilde{\mathcal{N}}_j \cap \tilde{\mathcal{N}}_{j'}$, *we have* $V_{\tau_j(\ell)} = V_{\tau_{j'}(\ell)}$. *Then following Lemma 2, we can conclude that* $\tilde{\pi}_1 = \pi_1$, $\tilde{\pi}_n = \pi_5$, $\tilde{\pi}_j = \pi_2 = \pi_3 = \pi_4, j = 2, 3, \ldots, n-1$ *are stabilizing controllers for the new system with certificates* $\tilde{V}_1 = V_1, \tilde{V}_n = V_5$, *and* $\tilde{V}_j = V_2 = V_3 = V_4, j = 2, 3, \ldots, n-1$.

3.3. Robust ISS Lyapunov Functions

Many subsystems in a networked system are very similar in terms of dynamics and network structure but may have different parameters. Furthermore, system dynamics may have model uncertainties and unknown parameters (Dawson et al., 2021). For instance, in the Platoon example, trucks may have different weights, and the leading truck can have unknown velocity and acceleration, but the platoon system should be stabilized for any driving style of the leading truck. Having a robust version of ISS Lyapunov functions that can work for a set of different subsystems with different parameters can significantly reduce the number of ISS Lyapunov functions we need to find for a large networked system and also improve the robustness of the resulting controller. To tackle this, we show that the robust ISS Lyapunov functions can be established for control-affine systems taking the form: $\dot{x}_i = h_i(x_i, x_{\mathcal{N}_i}; \beta) + g_i(x_i, x_{\mathcal{N}_i}; \beta)u_i$, where $\beta \in \mathcal{B}$ is the parameter of the dynamics that models uncertainties. Such an assumption is not restrictive and can cover a large range of physical systems, *e.g.* , systems following the manipulator function (Tedrake, 2022). Under this assumption, we further introduce robust ISS Lyapunov functions to guarantee the global asymptotic stability of systems with uncertainties. The proof of Lemma 3 is postponed to Appendix B.

Lemma 3 (Robust ISS Lyapunov Functions) *Given a networked dynamical system with control-affine dynamics with bounded parametric uncertainty* $\beta \in \mathcal{B}$, *where* $\mathcal{B}$ *is the convex hull of parameters* $\beta_1, \beta_2, \ldots, \beta_{n_\beta}$. *If there exists ISS Lyapunov functions* V_i *satisfying the conditions in Lemma 1 for each* $\beta_j, j \in \{1, 2, \ldots, n_\beta\}$, *the dynamics* h_i *and* g_i *are affine with respect to* β, *then the closed-loop system is globally asymptotically stable with any* $\beta \in \mathcal{B}$.

Example 3 *The robust ISS Lyapunov functions can be directly used to the Platoon system.* $v_0 = v_{n+1}$ *is a parameter of this system, which is bounded between* $v_{\min}$ *and* $v_{\max}$ *by assumption. Following Lemma 3, if we can find ISS Lyapunov functions for both* $v_{\min}$ *and* $v_{\max}$, *we can ensure our system is stable with any velocity of the leading truck.*

4. Learning Compositional Certificates and Controllers

Based on the compositional certificate developed in Section 3, we now focus on jointly learning the individual ISS Lyapunov functions and the decentralized controllers. We note that while there are many existing approaches to learn neural certificates (Dawson et al., 2021; Gaby et al., 2021), they can not be directly applied here because we have a unique imply condition (2) that can not be handled by the existing approaches. We will introduce a novel approach to incorporate the imply condition as specially designed loss terms. To proceed, we start with formally defining the parameterization of the decentralized controllers and the ISS Lyapunov functions.

Controllers and ISS Lyapunov Functions Parameterization. We focus on decentralized controllers, in which the control u_i of subsystem i only depends on the subsystem state x_i, i.e.

$u_i = \pi_i(x_i; \theta_i)$. Here π_i is an NN with θ_i as the parameters. We parameterize Lyapunov function $V_i(x_i)$ as $V_i(x_i; S_i, \omega_i, \nu_i) = x_i^\top S_i^\top S_i x_i + p_i(x_i; \omega_i)^\top p_i(x_i; \omega_i) + q_i(x_i; \nu_i)$. where $S_i \in \mathbb{R}^{d_i \times d_i}$ is a matrix of parameters, $p_i(x_i; \omega_i)$ is an NN with weights ω_i, and $q_i(x_i; \nu_i)$ is another NN with weights ν_i and ReLU as the output activation function, which is only applied to the output of the NN. The first term in $V_i(x_i)$ is a quadratic term to capture the linear part of the non-linear dynamics. The second term is a sum-of-squares term to capture the polynomial part of the dynamics. Finally, the third term is used to model the residues. Using this form, the ISS Lyapunov function satisfies $V_i(x_i) \geq 0$ by construction.

Sharing ISS Lyapunov Across Subsystems. Following Lemma 2 and Lemma 3, to reduce the number of neural networks and to make the ISS Lyapunov functions learned in small-scale networked system generalizable to large-scale systems, we use the following weight sharing technique. We let the subsystems share the same ISS Lyapunov functions if their dynamics are similar. In addition, we can also let similar subsystems share the same controller. In this way, the number of trainable parameters can be reduced, and we can easily apply the controllers trained in small-scale systems to large-scale systems. For example, in the previous Platoon system, we can let the j-th truck, $j = 2, 3, \ldots, n-1$, share the same ISS Lyapunov function and the same controller.

Gain Function Parameterization. We use linear functions to model the gain functions χ_{ij} in Equation (2). We let $\chi_{ij}(x) = \chi_i(x; k_i) = \mathrm{Sigmoid}(k_i)x, \forall j$, where k_i is a trainable parameter, x is the scalar input of the gain functions, which is always the output of the ISS Lyapunov functions. In this way, the condition $\chi_{ij}(x) < x, \forall x > 0$ is satisfied by construction.

Loss Functions. A key challenge in learning ISS Lyapunov functions is how to ensure condition (2) is satisfied. We now propose a methodology that promotes (2). Let Boolean $A_i, B_i \in \{0, 1\}$ be[2]

$$A_i = V_i(x_i) \geq \max_{j \in \mathcal{N}_i} \chi_i(V_j(x_j)), \quad B_i = [\nabla V_i(x_i)]^\top f_i(x_i, x_{\mathcal{N}_i}, \pi_i(x_i)) \leq -\alpha_i V_i(x_i). \tag{3}$$

Then condition (2) can be written as $A_i \Rightarrow B_i$, which is the same as $\neg A_i \vee B_i$, or $\max\{\neg A_i, B_i\}$. However, this kind of formulation is not trainable for neural networks because Boolean variables are not differentiable. To settle this problem, we introduce the following losses:

$$\mathcal{L}_{A_i} = \mathrm{ReLU}\left(V_i(x_i) - \max_{j \in \mathcal{N}_i} \chi_i(V_j(x_j)) + \epsilon_A\right), \tag{4}$$

$$\mathcal{L}_{B_i} = \mathrm{ReLU}\Big([\nabla V_i(x_i)]^\top f_i(x_i, x_{\mathcal{N}_i}, \pi_i(x_i)) + \alpha_i V_i(x_i) + \epsilon_B\Big), \tag{5}$$

where ϵ_A and ϵ_B are small parameters that encourages strict satisfactions and generalization abilities (Dawson et al., 2021). $\mathcal{L}_{A_i}$ and $\mathcal{L}_{B_i}$ can address the problem introduced by Boolean variables A_i and B_i, but they introduce a new problem that $\max\{\neg A_i, B_i\}$ cannot be written as $\max\{\mathcal{L}_{A_i}, \mathcal{L}_{B_i}\}$ since the two losses are not comparable. To address this issue, we minimize the loss $\mu_{A_i}\mathcal{L}_{A_i} + \mu_{B_i}\mathcal{L}_{B_i}$ instead, where μ_{A_i} and μ_{B_i} are two hyper-parameters for balancing the two losses. Note that in practice, since we often simulate the dynamical systems in a discrete way, we can use two ways to calculate $\nabla V_i(x_i)$ in $\mathcal{L}_{B_i}$. First, we can directly calculate the gradient of $V_i(x_i)$ w.r.t. x_i. For the second method, we can just do a one-step simulation $x_i \to x_i^{\mathrm{next}}$, and approximate $[\nabla V_i(x_i)]^\top f_i(x_i, x_{\mathcal{N}_i}, \pi_i(x_i))$ with $(V_i(x_i^{\mathrm{next}}) - V_i(x_i))/\Delta t$, where Δt is the simulation time step.

To ensure the condition $V_i(x_i) = 0$ for $x_i \in \mathcal{X}_i^{\mathrm{goal}}$ is satisfied, we introduce another loss term $\frac{1}{|\widehat{\mathcal{X}}_i^{\mathrm{goal}}|} \sum_{x_i^{\mathrm{goal}} \in \widehat{\mathcal{X}}_i^{\mathrm{goal}}} |V_i(x_i^{\mathrm{goal}})|$, where $\widehat{\mathcal{X}}_i^{\mathrm{goal}}$ is a randomly sampled set of states from $\mathcal{X}_i^{\mathrm{goal}}$. In

2. For notational simplicity, from now on we omit all the notations of parameters in the function approximators.

addition, we add $\|\pi_i(x_i) - u_i^{\text{nominal}}\|^2$ to the loss, where u_i^{nominal} is the control signal calculated by some nominal controller. We use the Droop controller and LQR controller in our experiments. We add the nominal controller so that the learned controller can explore the "informed region" near the nominal control signal rather than randomly, in order to accelerate the training. We do not need the nominal controller to be stable or optimal, and the learned controller behaves much better than the nominal controller as shown in Section 5. The final loss function used in training is

$$\mathcal{L} = \sum_{i=1}^{n} \left[\frac{1}{|\widehat{\mathcal{X}}_i^{\text{goal}}|} \sum_{x_i^{\text{goal}} \in \widehat{\mathcal{X}}_i^{\text{goal}}} |V_i(x_i^{\text{goal}})| + \mu_{A_i}\mathcal{L}_{A_i} + \mu_{B_i}\mathcal{L}_{B_i} + \mu_{\text{ctrl}}\|\pi_i(x_i) - u_i^{\text{nominal}}\|^2 \right], \tag{6}$$

where μ_{ctrl} is a tuning parameter, and the training parameters are $\theta, S, \omega, \nu, k$.

Training Procedure. During training, we draw samples by randomly sampling states in the state space and in the goal set. We first initialize the controller by minimizing the loss $\|\pi_i(x_i) - u_i^{\text{nominal}}\|^2$, and then fix the controller to initialize the ISS Lyapunov function by minimizing the loss $\mathcal{L}_{B_i}$. After the initialization, we minimize loss (6) to train the controllers, the ISS Lyapunov functions, and the gain functions jointly. The contour plots of the learned robust ISS Lyapunov functions are provided in Appendix C.2.5.

5. Experiments

We demonstrate NeurISS in 3 environments including Power system, Platoon, and Drone, aiming to answer the following questions: How does NeurISS compare with other algorithms in the case of stabilizing networked systems? Can NeurISS perform similarly or surpass the centralized controllers in small-scale networked systems? Can NeurISS scale up to large-scale networked systems? We provide implementation details, introductions to the systems, and more results in the appendix.

Baselines. We compare NeurISS with both centralized and decentralized baselines. For centralized ones, we compare with the state-of-the-art RL algorithm PPO (Schulman et al., 2017), the RL-with-Lyapunov-critic algorithm LYPPO (Chang and Gao, 2021), and the centralized Neural CLF controller (NCLF) (Dawson et al., 2021). For decentralized ones, we compare with the classical LQR (Kwakernaak et al., 1974) controller and the multi-agent RL algorithm MAPPO (Yu et al., 2021). We hand-craft reward functions based on the common way of designing reward functions for tracking problems for the RL algorithms. For LQR, since the agents only have local observations, we calculate the goal point for the LQR controller based on local observations in each time step.

5.1. Environment Descriptions

Power Systems. We consider two control problems in power systems. Firstly, we consider a networked microgrid system introduced in Huang et al. (2021), where there is an interconnection of 5 microgrids. Each microgrid i has two states $x_i = (\delta_i, E_i)$ where δ_i is the voltage phase angle and E_i is the voltage magnitude. The goal is to design controllers so that δ_i, E_i can converge to their reference values $\delta_i^{\text{ref}}, E_i^{\text{ref}}$. Secondly, we consider a distribution grid voltage control problem (Shi et al., 2022) which we name GridVoltage8. The goal is to drive the distribution grid voltage to the nominal value 1.0. Due to space limits, more details of the two systems are deferred to Appendix C.1. Since the dynamics of the power systems are not separable, we use a droop controller as one of the baselines and the nominal controller for NeurISS and NCLF instead of LQR.

Environment	Microgrid5	GridVoltage8	Platoon5	Drone2x2
NeurISS	2027.25 ± 18.15	$\mathbf{2483.70} \pm 1.68$	$\mathbf{2054.89} \pm 95.84$	$\mathbf{1713.73} \pm 13.35$
LQR	——	——	1894.64 ± 5.20	1209.15 ± 5.27
Droop	1431.30 ± 55.03	2251.70 ± 0.00	——	——
PPO	1970.89 ± 43.97	1086.72 ± 1.13	1489.32 ± 125.24	1489.32 ± 125.24
LYPPO	$\mathbf{2234.58} \pm 34.92$	1086.27 ± 2.12	707.08 ± 300.61	41.18 ± 6.05
MAPPO	1934.72 ± 149.66	1086.06 ± 0.27	1880.80 ± 155.49	1549.06 ± 271.32
NCLF	1887.36 ± 26.98	1078.19 ± 660.76	1979.02 ± 8.03	871.87 ± 46.02

Table 1: The expected reward of NeurISS and the baselines in the small-scale environments

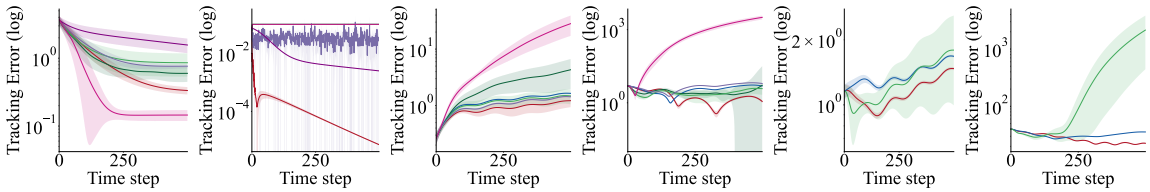

Figure 1: The tracking error in log scale w.r.t. time step of NeurISS and the baselines in the small-scale environments (*a-d*) and the large-scale environments (*e-f*). The line shows the mean tracking error while the shaded region shows the standard deviation.

Platoon. The Platoon system has been introduced in the examples before. We use the LQR controller as the nominal controller of NeurISS and NCLF. Because of the robustness and generalizability of NeurISS we let the controller and the ISS Lyapunov functions of the first and the n-th truck share the same weights, while other trucks also share the same weights. We let the controllers of MAPPO share the weights in the same way for a fair comparison. In testing, we let the leading truck's acceleration follow a sin-like curve with clips, which is hard to track. In the small-scale training and testing, we use $n = 5$ trucks. In the large-scale testing, we use $n = 100$ trucks.

Drone. We design the Planar Drone Formation Control environment to further demonstrate the capability of NeurISS in complex networked systems. In this environment, at the beginning of the simulations, the planar drones (Tedrake, 2022) stay on the ground. As the simulations start, we want the drones to form a 2D mesh grid while tracking a given trajectory. The states of the drones are modeled as a 2-D platoon system, which is given by $x_i = [p_i^l, p_i^r, p_i^u, p_i^d, \theta_i, v_i^x, v_i^y, \omega_i]^\top$, where $p_i^l, p_i^r, p_i^u, p_i^d$ are the distances from drone i to the left, right, up, down drones, θ_i is the angle between the drone and the horizontal line, v_i^x and v_i^y are velocities and ω_i is the angular velocity. The control inputs of each drone are the forces generated by the two propellers. We use the LQR controller as the nominal controller for NeurISS and NCLF, and let the controllers in NeurISS and MAPPO share the same weights. Because of the robustness, we also let the ISS Lyapunov functions in NeurISS share the same weights, so we only need 1 ISS Lyapunov function. In testing, we set the target trajectory to follow a horizontal line with a sin-like acceleration. In the small-scale training and testing, we use $2 \times 2 = 4$ drones. In large-scale testing, we use $10 \times 10 = 100$ drones.

5.2. Results

The results show that compared with the baselines, NeurISS can achieve comparable or better rewards and tracking errors in small-scale environments, and significantly higher rewards and lower tracking errors in large-scale environments, which demonstrate its efficacy and generalizability.

Small-scale Experiments. In Table 1, we show the expected rewards and standard deviations of NeurISS and the baselines, and in Figure 1 (a-d) we show the tracking error w.r.t. the simulation time steps. We can observe that in the small-scale system Microgrid5 (10 dimensions), NeurISS achieves the second highest expected reward, and second lowest tracking error, while LYPPO behaves the best. In GridVoltage8 (8 dimensions) and larger systems Platoon5 and Drone2x2 (15 and 24 dimensions), NeurISS achieves the highest expected rewards and the lowest tracking error. Note that in Platoon5 and Drone2x2, we do not have full knowledge of the tracking trajectory, and the trajectory changes fast, so the tracking error cannot converge to 0. NeurISS has this performance because of its ability to learn decentralized controllers *jointly* with the ISS Lyapunov functions as certificates. Compared with NCLF, NeurISS performs better because it is hard to find a global CLF for networked systems. PPO and MAPPO achieve lower rewards than NeurISS. They are policy gradient methods to approximate the solution of the Bellman equation, so there is no certificate of their stability. LYPPO, although outperforms NeurISS in very small systems (Microgrid), its performance drops a lot in larger systems (Platoon and Drone). This is because in small-scale systems, with the guidance of CLF, RL can achieve the goal very quickly to maximize the cumulative reward, but NeurISS only seeks to reach the goal without targeting on the convergence speed. Therefore, NeurISS convergences slower than LYPPO. However, in larger scale systems, because of the hardness of finding a correct global CLF, LYPPO receives the wrong guidance by the wrong CLF, and thus behaves much worse than NeurISS and even PPO and MAPPO. For the nominal controllers, LQR is designed for linear systems and the Droop controllers are hand-tuned. Therefore, their performance is hard to guarantee in complex nonlinear networked systems.

Large-scale Experiments. One key advantage of the proposed framework is network generalizability (Section 3.2), where the decentralized controllers and ISS Lyapunov functions trained in small networked systems can be directly applied to large networked systems without further training, while the centralized controllers need a really long time to be trained on large-scale systems (Approximately 250 hours for Drone10x10). We test the 3 decentralized approaches, NeurISS, MAPPO, and LQR in large-scale Platoon and Drone systems, Platoon100 and Drone10x10, with 100 trucks and $10 \times 10 = 100$ drones. We show the tracking errors w.r.t. the simulation time steps in Figure 1 (e-f). We observe that NeurISS has the smallest tracking errors in both environments, with large gaps to others, which shows that NeurISS has the strongest scalability.

6. Conclusion

In this paper, we propose a neural compositional certificate framework for stabilizing large-scale networked dynamical systems. Limitations of the approach include: 1) the approach requires the knowledge of the dynamical system functions f_i; 2) the network generalizability result Lemma 2 requires a strong symmetric condition; 3) the robust ISS Lyapunov result Lemma 3 assumes the dynamical system is control affine; 4) the approach only learns a compositional certificate using finite samples but does not verify it, so in some sense, the ISS Lyapunov functions we learn are only candidate ISS Lyapunov functions. These limitations are all interesting future directions.

Acknowledgments

The Defense Science and Technology Agency in Singapore and the C3.ai Digital Transformation Institute provided funds to assist the authors with their research. Guannan Qu is also supported by NSF Grant 2154171. However, this article solely reflects the opinions and conclusions of its authors and not DSTA Singapore, the Singapore Government, or C3.ai Digital Transformation Institute.

References

Alessandro Abate, Daniele Ahmed, Mirco Giacobbe, and Andrea Peruffo. Formal synthesis of lyapunov neural networks. *IEEE Control Systems Letters*, 5(3):773–778, 2020.

Felix Berkenkamp, Matteo Turchetta, Angela Schoellig, and Andreas Krause. Safe model-based reinforcement learning with stability guarantees. *Advances in neural information processing systems*, 30, 2017.

Fernando Castañeda, Jason J Choi, Bike Zhang, Claire J Tomlin, and Koushil Sreenath. Gaussian process-based min-norm stabilizing controller for control-affine systems with uncertain input effects and dynamics. In *2021 American Control Conference (ACC)*, pages 3683–3690. IEEE, 2021.

Ya-Chien Chang and Sicun Gao. Stabilizing neural control using self-learned almost lyapunov critics. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1803–1809. IEEE, 2021.

Ya-Chien Chang, Nima Roohi, and Sicun Gao. Neural Lyapunov Control. In *Advances in Neural Information Processing Systems*, volume 32, pages 3245–3254, 2019.

Richard Cheng, Gábor Orosz, Richard M Murray, and Joel W Burdick. End-to-end safe reinforcement learning through barrier functions for safety-critical continuous control tasks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3387–3395, 2019.

Glen Chou, Necmiye Ozay, and Dmitry Berenson. Model error propagation via learned contraction metrics for safe feedback motion planning of unknown systems. In *2021 60th IEEE Conference on Decision and Control (CDC)*, pages 3576–3583. IEEE, 2021.

Yinlam Chow, Ofir Nachum, Edgar Duenez-Guzman, and Mohammad Ghavamzadeh. A lyapunov-based approach to safe reinforcement learning. *Advances in neural information processing systems*, 2018.

Charles Dawson, Zengyi Qin, Sicun Gao, and Chuchu Fan. Safe nonlinear control using robust neural lyapunov-barrier functions. *Conference on Robot Learning*, 2021.

Geir E Dullerud and Fernando Paganini. *A course in robust control theory: a convex approach*, volume 36. Springer Science & Business Media, 2013.

Nathan Gaby, Fumin Zhang, and Xiaojing Ye. Lyapunov-net: A deep neural network architecture for lyapunov function approximation. *arXiv preprint arXiv:2109.13359*, 2021.

Minghao Han, Lixian Zhang, Jun Wang, and Wei Pan. Actor-critic reinforcement learning for control with stability guarantee. *IEEE Robotics and Automation Letters*, 5(4):6217–6224, 2020.

Tong Huang, Sicun Gao, and Le Xie. A neural lyapunov approach to transient stability assessment of power electronics-interfaced networked microgrids. *IEEE Transactions on Smart Grid*, 13(1): 106–118, 2021.

Zhong-Ping Jiang and Tengfei Liu. Small-gain theory for stability and control of dynamical networks: A survey. *Annual Reviews in Control*, 46:58–79, 2018.

Zhong-Ping Jiang, Iven MY Mareels, and Yuan Wang. A lyapunov formulation of the nonlinear small-gain theorem for interconnected iss systems. *Automatica*, 32(8):1211–1215, 1996.

Wanxin Jin, Zhaoran Wang, Zhuoran Yang, and Shaoshuai Mou. Neural certificates for safe control policies. *arXiv preprint arXiv:2006.08465*, 2020.

IS Khalil, JC Doyle, and K Glover. *Robust and optimal control*. Prentice hall, 1996.

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Huibert Kwakernaak, Raphael Sivan, and Bjor N D Tyreus. Linear optimal control systems. 1974.

Tengfei Liu, David J Hill, and Zhong-Ping Jiang. Lyapunov formulation of iss cyclic-small-gain in continuous-time dynamical networks. *Automatica*, 47(9):2088–2093, 2011.

Tengfei Liu, Zhong-Ping Jiang, and David J Hill. Lyapunov formulation of the iss cyclic-small-gain theorem for hybrid dynamical networks. *Nonlinear Analysis: Hybrid Systems*, 6(4):988–1001, 2012.

Steven H Low. Convex relaxation of optimal power flow—part i: Formulations and equivalence. *IEEE Transactions on Control of Network Systems*, 1(1):15–27, 2014.

Gaurav Manek and J Zico Kolter. Learning stable deep dynamics models. *Advances in neural information processing systems*, 2019.

Yue Meng, Zengyi Qin, and Chuchu Fan. Reactive and safe road user simulations using neural barrier certificate. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2021.

Takeru Miyato, Toshiki Kataoka, Masanori Koyama, and Yuichi Yoshida. Spectral normalization for generative adversarial networks. *arXiv preprint arXiv:1802.05957*, 2018.

Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

Pablo A Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. California Institute of Technology, 2000.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019.

Andrea Peruffo, Daniele Ahmed, and Alessandro Abate. Automated and formal synthesis of neural barrier certificates for dynamical models. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*, pages 370–388. Springer, 2021.

Warren B Powell. *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. John Wiley & Sons, 2007.

Zengyi Qin, Yuxiao Chen, and Chuchu Fan. Density constrained reinforcement learning. In *International Conference on Machine Learning*, pages 8682–8692. PMLR, 2021a.

Zengyi Qin, Kaiqing Zhang, Yuxiao Chen, Jingkai Chen, and Chuchu Fan. Learning safe multi-agent control with decentralized neural barrier certificates. In *International Conference on Learning Representations*, 2021b.

Antonin Raffin, Ashley Hill, Adam Gleave, Anssi Kanervisto, Maximilian Ernestus, and Noah Dormann. Stable-baselines3: Reliable reinforcement learning implementations. *Journal of Machine Learning Research*, 22(268):1–8, 2021. URL http://jmlr.org/papers/v22/20-1364.html.

Spencer M Richards, Felix Berkenkamp, and Andreas Krause. The lyapunov neural network: Adaptive stability certification for safe learning of dynamical systems. In *Conference on Robot Learning*, pages 466–476. PMLR, 2018.

John Schulman, Sergey Levine, Pieter Abbeel, Michael Jordan, and Philipp Moritz. Trust region policy optimization. In *International conference on machine learning*, pages 1889–1897. PMLR, 2015.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Yuanyuan Shi, Guannan Qu, Steven Low, Anima Anandkumar, and Adam Wierman. Stability constrained reinforcement learning for real-time voltage control. In *2022 American Control Conference (ACC)*, pages 2715–2721. IEEE, 2022.

Jean-Jacques E Slotine and Weiping Li. *Applied nonlinear control*. Prentice Hall, 1991.

Eduardo D Sontag. *Mathematical control theory: deterministic finite dimensional systems*, volume 6. Springer Science & Business Media, 2013.

Mohit Srinivasan, Amogh Dabholkar, Samuel Coogan, and Patricio A Vela. Synthesis of control barrier functions using a supervised machine learning approach. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7139–7145. IEEE, 2020.

Srdjan S Stankovic, Milorad J Stanojevic, and Dragoslav D Siljak. Decentralized overlapping control of a platoon of vehicles. *IEEE Transactions on Control Systems Technology*, 8(5):816–832, 2000.

Dawei Sun, Susmit Jha, and Chuchu Fan. Learning certified control using contraction metric. In *Conference on Robot Learning*, 2020.

Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

Russ Tedrake. *Underactuated Robotics*. 2022. URL http://underactuated.mit.edu.

L. Thurner, A. Scheidler, F. Schäfer, J. Menke, J. Dollichon, F. Meier, S. Meinecke, and M. Braun. pandapower — an open-source python tool for convenient modeling, analysis, and optimization of electric power systems. *IEEE Transactions on Power Systems*, 33(6):6510–6521, Nov 2018. ISSN 0885-8950. doi: 10.1109/TPWRS.2018.2829021.

Pravin Varaiya. Max pressure control of a network of signalized intersections. *Transportation Research Part C: Emerging Technologies*, 36:177–195, 2013.

Wei Xiao, Ramin Hasani, Xiao Li, and Daniela Rus. Barriernet: A safety-guaranteed layer for neural networks. *arXiv preprint arXiv:2111.11277*, 2021.

Chao Yu, Akash Velu, Eugene Vinitsky, Yu Wang, Alexandre Bayen, and Yi Wu. The surprising effectiveness of ppo in cooperative, multi-agent games. *arXiv preprint arXiv:2103.01955*, 2021.

Changhong Zhao, Ufuk Topcu, Na Li, and Steven Low. Design and stability of load-side primary frequency control in power systems. *IEEE Transactions on Automatic Control*, 59(5):1177–1189, 2014.

Weiye Zhao, Tairan He, and Changliu Liu. Model-free safe control for zero-violation reinforcement learning. In *5th Annual Conference on Robot Learning*, 2021.