

A Robust, Efficient Predictive Safety Filter

Wenceslao Shaw Cortez, Jan Drgona, Draguna Vrabie, Mahantesh Halappanavar

Abstract—In this paper, we propose a novel predictive safety filter that is robust to bounded perturbations and is implemented in an even-triggered fashion to reduce online computation. The proposed safety filter extends upon existing work to reject disturbances for discrete-time, time-varying nonlinear systems with time-varying constraints. The safety filter is based on novel concepts of robust, discrete-time barrier functions and can be used to filter any control law. Here, we use the safety filter in conjunction with Differentiable Predictive Control (DPC) as a promising offline learning-based policy optimization method. The approach is demonstrated on a two-tank system, building, and single-integrator example.

I. INTRODUCTION

Control barrier functions are a useful tool for ensuring constraint satisfaction of nonlinear, dynamical systems [1]. This method has received increasing attention recently due to its modularity and applicability to learning-based techniques that generally do not have guarantees of constraint satisfaction, i.e., safety [2], [3]. Barrier functions are also advantageous for being robust to perturbations and provide asymptotic stability to the safe set.

One criticism of barrier function methods, however, is the difficulty in constructing the functions themselves. [4] provides a way to construct barrier functions using maximal output admissible sets, but is dependent on a known stabilizing control law. Verification methods have been developed for uncertain system dynamics [5], but only probabilistic guarantees of safety are provided. Some synthesis methods are restricted to specific types of systems [6] or are dependent on sampling methods or sum-of-squares techniques, which do not scale well with system size [7]. Other approaches require expert-provided trajectories in order to learn the barrier function [8], [9]. Most of these approaches require significant offline data and computation, may be subject to conservatism, and are generally restricted to time-invariant systems. An alternative approach considered here is to use a prediction horizon to relax these restrictions.

Many existing barrier function methods implement the safety-critical control as a nonlinear program, which is effectively a 1-step look-ahead model predictive control (MPC)

[†] W. Shaw Cortez, J. Drgona, D. Vrabie, and M. Halappanavar are with the Data Science and Machine Intelligence Department at the Pacific Northwest National Laboratory, Richland, WA, USA. E-Mail: {w.shawcortez, jan.drgona, draguna.vrabie, mahantesh.halappanavar}@pnnl.gov

This research was partially supported by the U.S. Department of Energy, through the Office of Advanced Scientific Computing Research's "Data-Driven Decision Control for Complex Systems (DnC2S)" project, and through the Energy Efficiency and Renewable Energy, Building Technologies Office under the "Advancing Market-Ready Building Energy Management by Cost-Effective Differentiable Predictive Control" projects. PNNL is a multi-program national laboratory operated for the U.S. Department of Energy (DOE) by Battelle Memorial Institute under Contract No. DE-AC05-76RL0-1830.




Fig. 1: Schematics of the proposed event-triggered robust predictive safety filter combined with learning-based policy.

problem. The ability to predict system behavior can relax conservatism in the 1-step look-ahead approach. There exist methods that combine the finite horizon, MPC setup with control barrier functions [10]–[14] as well as those that use robust MPC methods [15]–[19]. Here we focus on the discrete-time safety addressed in [12]–[19], which unfortunately do not provide hard guarantees of safety for perturbed, nonlinear, time-varying systems with time-varying constraints. In [12], a predictive safety filter is developed for time-invariant, nonlinear systems that is always guaranteed to be feasible and ensures asymptotic stability to the safe set. In [13], safety and asymptotic stability are addressed by combining MPC with control barrier functions for nonlinear, time-invariant systems. In [14], a disturbance observer is combined with an MPC-based barrier function control law to guarantee the safety of perturbed nonlinear, time-invariant systems with measurement noise. Other types of predictive safety filters have been developed for handling perturbations, but are restricted to time-invariant systems and require global bounds on the disturbance [15], are restricted to linear time-invariant systems [16], or focus on probabilistic safety [17]–[19]. We note that there exist many robust, nonlinear MPC methods, but many time-varying implementations are focused on linear systems [20], [21] or adaptive methods [22], [23] wherein there is control over how the time-varying components update. Here we are interested in the development of a robust, predictive safety filter for time-varying systems with time-varying constraints. This sole focus on safety (instead of safety and performance common to MPC) allows for less restrictive conditions to be satisfied, can yield additional properties such as convergence to the safe region [12], and as will be shown here can reduce the amount of computation needed online.

As mentioned previously, predictive safety encompasses the 1-step safety filters common to barrier function methods. Existing discrete-time barrier function methods have been ap-

plied to bipedal robotics [24] and adaptive high-order systems [25], for example, but there is a limited range of methods for handling bounded perturbations. [26] presents a robust barrier function that uses GP estimate of disturbance, but no hard guarantees of safety are provided. [27] presents barrier functions that are robust to stochastic perturbations. [28] use barrier functions for multi-agent partially observable Markov decision processes for probabilistic safety guarantees. [29] uses MPC in combination with barrier functions, but no formal guarantees of safety are provided. [30] combines high-order barrier functions in an MPC framework. [13] combines nonlinear MPC with barrier functions, focusing on enhancing feasibility. [31] uses MPC with barrier functions for obstacle avoidance. Apart from [26], [27] [12], and [14], none of the previous methodologies address robustness to perturbations. The methods from [26] [27] only address probabilistic safety, whereas the asymptotic stability results from [12] are effective for vanishing perturbations, but not for bounded, non-vanishing perturbations which are common in practice. In the continuous-time domain, there exist many solutions that are robust to bounded perturbations [32], [33], but this has been lacking in the discrete-time domain. Furthermore, many of the aforementioned approaches require constant on-line computation of optimization problems, which can be restrictive for computationally limited edge devices. A new safety filter is needed that is robust, addresses time-varying systems/constraints, and is only solved if the nominal control will violate constraints.

In this work, we extend the predictive safety filter [12] to address bounded perturbations as well as time-varying systems and constraints. Furthermore, we address computation demand via an event-triggering scheme. This event triggering only requires solving the safety filter optimization problem if the nominal control performs an unsafe action. We also extend the existing results in the literature to address robust discrete-time barrier functions with a 1-step safety filter that reduces computation and simplifies the process of barrier function verification. The resulting methodology is guaranteed to be safe and robust. Figure 1 illustrates the proposed method. The main contributions of this work are: (1) Robust predictive safety filter for time-varying systems and constraints, (2) 1-step robust safety filter for discrete-time systems, and (3) event-triggering methods for reducing online computations.

The paper is outlined as follows. In Section II we introduce the notation, relevant preliminaries, and the problem formulation. In Section III, we introduce the robust barrier function formulation and the robust predictive safety filter. In Section IV, we extend the predictive safety filter with an event-triggering scheme and develop a 1-step robust safety filter for robust discrete-time barrier functions. In Section V, we apply the proposed method to a two-tank system, a building system, and a single integrator.

II. BACKGROUND

A. Notation and Preliminaries

Let $\mathbb{N}$ be the set of natural numbers including 0. The notation $\mathcal{L}_f^k$ is used to denote the Lipschitz constant of a Lipschitz function, $\mathbf{f}(\mathbf{x}, k)$, with respect to the second argument.

Similarly, $\mathcal{L}_f^x$ is used to denote the Lipschitz constant of $\mathbf{f}$ with respect to the first argument. The notation $\mathcal{L}_f$ will be used to denote a bound with respect to the function $\mathbf{f}$ and will be explicitly defined in the text.

We say a function $z : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}$ is locally Lipschitz continuous on $\mathcal{D} \subset \mathbb{R}^n$ if z satisfies the following Lipschitz condition:

$$\|z(\mathbf{x}, k) - z(\mathbf{y}, k)\| \leq \mathcal{L}_z^x \|\mathbf{x} - \mathbf{y}\|$$

for all $\mathbf{x}, \mathbf{y} \in \mathcal{D}$, $k \in \mathbb{N}$. Unless otherwise stated, z can be discontinuous with respect to k .

We say a function $z : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}$ is 1-step bounded on $\mathbb{N}$ with $\mathcal{L}_z^k$ if z satisfies the following condition:

$$\|z(\mathbf{x}, k+1) - z(\mathbf{x}, k)\| \leq \mathcal{L}_z^k \quad (1)$$

for all $\mathbf{x} \in \mathbb{R}^n$, $k \in \mathbb{N}$.

B. Problem Formulation

Consider the following discrete-time dynamical system:

$$\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, k) + \mathbf{d}(\mathbf{x}_k, \mathbf{u}_k, k), \quad \mathbf{x}_0 \in \mathbb{R}^n, k \in \mathbb{N} \quad (2)$$

where $\mathbf{f} : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{N} \rightarrow \mathbb{R}^n$ is a function defining the known model, $\mathbf{u} \in \mathcal{U}$ is the control input for the compact set $\mathcal{U} \subset \mathbb{R}^m$, and $\mathbf{d} : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{N} \rightarrow \mathcal{D} \subset \mathbb{R}^n$ is a function defining the unknown disturbance.

Next, we define the safe region. Consider a function $b : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}$. This function b is used to define general system constraints that should be satisfied, whose related safe set is defined as follows:

$$\mathcal{X}(k) := \{\mathbf{x} \in \mathbb{R}^n : b(\mathbf{x}, k) \geq 0\} \quad (3)$$

The problem addressed here is to design a control law that implements a nominal control law, denoted $\mathbf{u}_{nom} : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathcal{U}$, while always making sure the system remains in the safe set $\mathcal{X}(k)$. Here the nominal control could be any previously defined control law or learning-based control that has no safety guarantees associated with it. The purpose of the proposed control is to only override the nominal control to ensure safe action.

Problem 1. Consider the system (2) with unknown disturbance. Given a safe region $\mathcal{X}(k)$ from (3) defined for all $k \in \mathbb{N}$, and any nominal control $\mathbf{u}_{nom} : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathcal{U}$, design a control law $\mathbf{u}_k : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathcal{U}$ such that $\mathbf{x}_k \in \mathcal{X}(k)$ for all $k \in \mathbb{N}$.

III. ROBUST PREDICTIVE SAFETY FILTER

A. Robust Barrier Functions

In this section, we present the proposed predictive safety filter method to address Problem 1. The approach considered here is based on control barrier functions for discrete-time systems.

Let $h : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}$ denote a barrier function and let the corresponding time-varying set be defined by:

$$\mathcal{C}(k) = \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}, k) \geq 0\} \quad (4)$$

The main idea behind the barrier function is that if h satisfies certain dynamic conditions for (2), then for any $\mathbf{x}_k \in \mathcal{C}(k)$, for any time $k \in \mathbb{N}$, there exists a control law to ensure $\mathcal{C}(k)$ is forward invariant, i.e., if $\mathbf{x}_0 \in \mathcal{C}(0)$, then $\mathbf{x}_k \in \mathcal{C}(k)$ for all $k \in \mathbb{N}$. Thus if $\mathcal{C}(k) \subset \mathcal{X}(k)$ for all $k \in \mathbb{N}$ and $\mathbf{x}_0 \in \mathcal{C}(0)$, then it follows that the safety condition from Problem 1 is satisfied: $\mathbf{x}_k \in \mathcal{X}(k)$ for all $k \in \mathbb{N}$. The key then is to determine what dynamic conditions h needs to satisfy in order to guarantee the safety condition in the presence of unknown disturbances. We will also extend this notion using prediction so that the dynamic condition only needs to hold at a horizon $N \in \mathbb{N}$, $N \geq 1$.

Before presenting the dynamic condition to address unknown disturbances we make the following realistic assumptions on (2) and h :

Assumption 1. *The system (2) satisfies the following:*

- 1) *The function $\mathbf{f} : \mathbb{R}^n \times \mathbb{R}^m \times \mathbb{N} \rightarrow \mathbb{R}^n$ is locally Lipschitz continuous on a set $\mathcal{W} \subset \mathbb{R}^n$, for which $\mathcal{W} \supset \mathcal{X}(k) \forall k \in \mathbb{N}$, with Lipschitz constant $\mathcal{L}_f^x \in \mathbb{R}_{\geq 0}$.*
- 2) *For the unknown disturbance in (2), there exists a bound $\mathcal{L}_d \in \mathbb{R}_{\geq 0}$, $\mathcal{L}_d < \infty$, satisfying:*

$$\mathcal{L}_d := \sup_{\mathbf{x} \in \mathcal{X}(k), \mathbf{u} \in \mathcal{U}, k \in \mathbb{N}} \|\mathbf{d}(\mathbf{x}, \mathbf{u}, k)\| \quad (5)$$

Assumption 2. *For $N \in \mathbb{N}$, $N \geq 1$, let $\mathcal{Y} \subset \mathbb{R}^n$ satisfy:*

$$\mathcal{Y} \supset \bigcup_{\mathbf{x} \in \mathcal{X}(k), k \in \mathbb{N}} \left\{ \mathbf{y} \in \mathbb{R}^n : \|\mathbf{y} - \mathbf{x}\| \leq \mathcal{L}_d \sum_{j=0}^{N-1} (\mathcal{L}_f^x)^j \right\}$$

The following Lipschitz conditions hold:

- 1) *The function $b : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}$ is locally Lipschitz continuous on $\mathcal{Y}$ with Lipschitz constant $\mathcal{L}_b^x \in \mathbb{R}_{\geq 0}$.*
- 2) *The function $h : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}$ is locally Lipschitz continuous on $\mathcal{Y}$, with Lipschitz constant $\mathcal{L}_h^x \in \mathbb{R}_{\geq 0}$.*

Remark 1. *Many dynamical systems satisfy the Lipschitz condition in Assumption 1 including robotics and buildings [6], [34], which is less restrictive than the smoothness requirement in similar MPC settings [21]. Also, the disturbance condition is a common assumption [15] and can be satisfied directly if a) $\mathbf{d}$ is continuous on its domain, $\mathcal{C}(k)$ is compact for all time, and $\mathbf{d}$ is bounded in time, or b) $\mathcal{D}$ is compact. Many existing time-invariant barrier function methods assume $\mathcal{C}(k)$ is compact which is a special case of this assumption [12]. Furthermore, Assumption 2 restricts only how h and b change with respect to $\mathbf{x}$. This assumption still allows for discontinuous jumps in time, which is reflected in many building applications. The Lipschitz condition in Assumption 2 takes into account the effect of the disturbance on b and h . The condition says that set $\mathcal{Y}$ must contain the set $\mathcal{X}(k)$ in addition to a ball, dependent on the disturbance bound, surrounding each point in $\mathcal{X}(k)$. Finally, there exist methods to compute Lipschitz constants for general systems [35] as well as neural-network models used in system identification [36], [37].*

Now we define the associated safety condition in a *predictive* context. For a horizon N , the associated dynamic condition is:

$$\delta \bar{h}_N(\mathbf{x}_k, \mathbf{u}_k, k) := h(\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, k), k+1) - \mathcal{L}_h^x \mathcal{L}_d (\mathcal{L}_f^x)^{N-1} \geq 0 \quad (6)$$

The condition in (6) requires that at the end of the N -step horizon, the predicted system needs to be sufficiently far into the safe region to take into account the worst possible disturbance. We define the *N -step robust, discrete-time barrier function* as follows:

Definition 1. *Consider the system (2) and the function $h : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}$ with $\mathcal{C}(k)$ defined by (4) and*

$$\bar{\mathcal{K}}(\mathbf{x}, k) := \{\mathbf{u} \in \mathcal{U} : \delta \bar{h}_N(\mathbf{x}, \mathbf{u}, k) \geq 0, \text{ if } \mathbf{x} \in \mathcal{C}(k)\}. \quad (7)$$

We say h is an N -step robust, discrete-time barrier function, if $\bar{\mathcal{K}}(\mathbf{x}, k) \neq \emptyset$, $\forall \mathbf{x} \in \mathcal{C}(k)$, $\forall k \in \mathbb{N}$.

Remark 2. *The motivation behind (6) is based on recent results for discrete-time barriers [4]. Consider the conventional time-invariant barrier function where $N = 1$ and no disturbance exists such that (6) becomes: $h(\mathbf{x}_{k+1}) \geq 0$. This is a necessary and sufficient condition for ensuring forward invariance of $\mathcal{C}$ as shown in [4]. This is notably different from other discrete-time barriers [25], [28] which adopt the continuous-time formulation: $h(\mathbf{x}_{k+1}) - h(\mathbf{x}_k) \geq -\gamma(h(\mathbf{x}_k))$ which rearranges to: $h(\mathbf{x}_{k+1}) \geq h(\mathbf{x}_k) - \gamma(h(\mathbf{x}_k))$ for an extended class- $\mathcal{K}$ function $\gamma : \mathbb{R} \rightarrow \mathbb{R}$ satisfying $\gamma(r) \leq r$. Although it seems that γ allows for freedom in the design of h , in the discrete-time case this is not true. Consider the simple choice of a scalar γ for which the condition $\gamma(r) \leq r$ becomes $\gamma \leq 1$, the difference condition becomes $h(\mathbf{x}_{k+1}) \geq (1 - \gamma)(h(\mathbf{x}_k))$. For any $\gamma < 1$, this requires $h(\mathbf{x}_{k+1})$ to be strictly greater than zero which is overly restrictive for ensuring safety. The choice of $\gamma = 1$ of course results in the same condition from [4]. The proposed predictive form in (6) extends the recent development from [4] to account for prediction and robustness in the barrier function.*

The *N -step robust, discrete-time barrier function* will be used to ensure there always exists a control N steps in the future to enforce safety. Before presenting the safety filter, we present a result on how to ensure future state trajectories are safe for an N -step look-ahead. To do so, we first define the prediction of a future state as:

$$\hat{\mathbf{x}}_{N|k} = \hat{\mathbf{f}}(\mathbf{x}_k, \{\mathbf{u}_{l|k}\}_{l=0}^{N-1}, k) := \mathbf{f}(\dots \mathbf{f}(\mathbf{x}_k, \mathbf{u}_{0|k}, k), \dots \mathbf{u}_{N-1|k}, k + N - 1) \quad (8)$$

where the notation $\{\mathbf{u}_{l|k}\}_{l=0}^N := \{\mathbf{u}_0, \mathbf{u}_1, \dots, \mathbf{u}_N\}$ is used to denote the set of inputs $\mathbf{u}_{l+k}$ starting from time $l = 0$ up to $l = N$. The notation $\mathbf{x}_{N|k}$ is the state $\mathbf{x}_{k+N}$ at time $k + N$ beginning from time k at which point $\hat{\mathbf{x}}_{0|k} = \mathbf{x}_k$.

Next, we define the following term to accommodate future predictions with respect to $\mathcal{X}(k)$:

$$\delta \hat{b}_{N|k}(\mathbf{x}_k, \{\mathbf{u}_{l|k}\}_{l=0}^{N-1}, k) := b(\hat{\mathbf{f}}(\mathbf{x}_k, \{\mathbf{u}_{l|k}\}_{l=0}^{N-1}, k), k+N) - \mathcal{L}_b^x \mathcal{L}_d \sum_{j=0}^{N-1} (\mathcal{L}_f^x)^j \geq 0 \quad (9)$$

Finally, we define the following set of safety-admissible control actions as:

$$\hat{\mathcal{K}}_{N|k}(\mathbf{x}_k, k) := \{\mathbf{u}_{i-1+k} \in \mathcal{U}, i \in [1, N] \subset \mathbb{N} : \delta \hat{b}_{i|k}(\mathbf{x}_k, \{\mathbf{u}_{l|k}\}_{l=0}^{i-1}, k) \geq 0\} \quad (10)$$

In the following proposition, we show that control inputs that lie inside the predictive set of safety admissible actions from (10) ensure open-loop safety of the system over the finite horizon.

Proposition 1. Consider the system (2) for which Assumption 1 holds. Given a function $b : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}$ satisfying Assumption 2.1, let $\hat{\mathcal{K}}_{N|k}$, $\delta \hat{b}_{N|k}$, and $\mathcal{X}(k)$ be defined by (10), (9), and (3) respectively. For a given $\mathbf{x}_{k_0} \in \mathcal{X}(k_0)$, and $k_0, N \in \mathbb{N}$, $N \geq 1$, if $\{\mathbf{u}_{l|k_0}\}_{l=0}^{N-1} \in \hat{\mathcal{K}}_{N|k_0}(\mathbf{x}_{k_0}, k_0)$ is applied to (2) in open-loop from time k_0 to $N-1$, then $\mathbf{x}_k \in \mathcal{X}(k)$ for all $k \in [k_0, k_0 + N] \subset \mathbb{N}$.

Proof. Let $i \in [1, N] \subset \mathbb{N}$. The proof follows by showing that the error between the unperturbed prediction $\hat{\mathbf{x}}_{i|k}$ (defined by (8)) and perturbed prediction $\mathbf{x}_{i|k}$ (defined by (2)) is appropriately bounded and that this bound is taken into account via $\delta \hat{b}_{i|k}$. By definition of $\delta \hat{b}_{N|k}$ and $\hat{\mathcal{K}}_{N|k}$, $\hat{\mathbf{x}}_{i|k} \in \mathcal{X}(k+i) \subset \mathcal{W}$ for all $i \in [1, N]$, i.e., the future trajectory of the unperturbed system is safe, the dynamics are locally Lipschitz continuous on this trajectory, and the bound on the disturbance holds. This cannot yet be said for the unperturbed trajectory.

We can however claim that for some $i \in [1, N] \subset \mathbb{N}$, $\mathbf{x}_{i|k} \in \mathcal{X}(k+i)$. This proof is straightforward by simply choosing $i = 1$ and noting that $\hat{\mathbf{x}}_{0|k_0} = \mathbf{x}_{0|k_0} = \mathbf{x}_{k_0} \in \mathcal{X}(k_0) \subset \mathcal{W}$ so the locally Lipschitz property and disturbance bound of Assumption 1 hold for $\mathbf{x}_{k_0}$, which along with the triangle inequality yields: $\|\hat{\mathbf{x}}_{1|k_0} - \mathbf{x}_{1|k_0}\| = \|\mathbf{f}(\hat{\mathbf{x}}_{0|k_0}, \mathbf{u}_{0|k_0}, k_0) - \mathbf{f}(\mathbf{x}_{0|k_0}, \mathbf{u}_{0|k_0}, k_0) - \mathbf{d}(\mathbf{x}_{0|k_0}, \mathbf{u}_{0|k_0}, k_0)\| \leq \mathcal{L}_d + \mathcal{L}_f^x \|\hat{\mathbf{x}}_{0|k_0} - \mathbf{x}_{0|k_0}\|$. From Assumption 2, the local Lipschitz property of b applies to $\mathbf{x}_{1|k_0}$ and $\hat{\mathbf{x}}_{1|k_0}$, such that the following holds: $|b(\hat{\mathbf{x}}_{1|k_0}, k_0 + 1) - b(\mathbf{x}_{1|k_0}, k_0 + 1)| \leq \mathcal{L}_b^x \|\hat{\mathbf{x}}_{1|k_0} - \mathbf{x}_{1|k_0}\|$. Thus from the previous relations we have: $|b(\hat{\mathbf{x}}_{1|k_0}, k_0 + 1) - b(\mathbf{x}_{1|k_0}, k_0 + 1)| \leq \mathcal{L}_b^x \mathcal{L}_d$. This implies that: $b(\mathbf{x}_{1|k_0}, k_0 + 1) \geq b(\hat{\mathbf{x}}_{1|k_0}, k_0 + 1) - \mathcal{L}_b^x \mathcal{L}_d$. By assumption, since $\{\mathbf{u}_{l|k_0}\}_{l=0}^{N-1} \in \hat{\mathcal{K}}_{N|k_0}(\mathbf{x}_{k_0}, k_0)$, $\mathbf{u}_{0|k_0}$ satisfies: $\delta \hat{b}_{1|k_0}(\mathbf{x}_{k_0}, \{\mathbf{u}_{l|k_0}\}_{l=0}^0, k_0) \geq 0$ for which $b(\hat{\mathbf{x}}_{1|k_0}, k_0 + 1) \geq \mathcal{L}_b^x \mathcal{L}_d$. This yields: $b(\mathbf{x}_{1|k_0}, k_0 + 1) \geq b(\hat{\mathbf{x}}_{1|k_0}, k_0 + 1) - \mathcal{L}_b^x \mathcal{L}_d \geq 0$ such that $\mathbf{x}_{1|k_0} = \mathbf{x}_{k_0+1} \in \mathcal{X}(k_0 + 1)$.

To complete the proof, we repeat the above argument recursively for $i = 2, \dots, N$. Note that for a given i , $\mathbf{x}_{i-1|k_0} \in \mathcal{X}(k_0 + i - 1)$ so that the Lipschitz property of $\mathbf{f}$ always holds. Next we define the following relations for any $i \in 1, \dots, N$. Using the Lipschitz property of $\mathbf{f}$ and the bound $\mathcal{L}_d$ along with the triangle inequality yields: $\|\hat{\mathbf{x}}_{i|k_0} - \mathbf{x}_{i|k_0}\| =$

$\|\mathbf{f}(\hat{\mathbf{x}}_{i-1|k_0}, \mathbf{u}_{i-1|k_0}, k_0 + i - 1) - \mathbf{f}(\mathbf{x}_{i-1|k_0}, \mathbf{u}_{i-1|k_0}, k_0 + i - 1) - \mathbf{d}(\mathbf{x}_{i-1|k_0}, \mathbf{u}_{i-1|k_0}, k_0 + i - 1)\| \leq \mathcal{L}_d + \mathcal{L}_f^x \|\hat{\mathbf{x}}_{i-1|k_0} - \mathbf{x}_{i-1|k_0}\|$. We repeatedly apply the above process until we reach $i = 1$. This can be written as follows for $F(y) = \mathcal{L}_d + \mathcal{L}_f^x y$ and $F_i(y) = \underbrace{F \circ F \circ \dots \circ F}_{i \text{ times}}(y)$: $\|\hat{\mathbf{x}}_{i|k} - \mathbf{x}_{i|k}\| \leq \underbrace{F \circ F \circ \dots \circ F}_{i \text{ times}}(\|\hat{\mathbf{x}}_{0|k} - \mathbf{x}_{0|k}\|) = F_i(\|\hat{\mathbf{x}}_{0|k} - \mathbf{x}_{0|k}\|)$. Now since $\hat{\mathbf{x}}_{0|k_0} = \mathbf{x}_{0|k_0} = \mathbf{x}_{k_0}$, it is clear from the previous relation that $\|\hat{\mathbf{x}}_{i|k} - \mathbf{x}_{i|k}\| \leq F_i(0)$ for which by evaluation $F_i(0) = \mathcal{L}_d \sum_{j=0}^{i-1} (\mathcal{L}_f^x)^j$.

Now it is clear that from the bound defined by $F_i(0)$, the Lipschitz property of b applies to each $\mathbf{x}_{i|k_0}$ so that the following holds: $|b(\hat{\mathbf{x}}_{i|k_0}, k_0 + i) - b(\mathbf{x}_{i|k_0}, k_0 + i)| \leq \mathcal{L}_b^x \|\hat{\mathbf{x}}_{i|k_0} - \mathbf{x}_{i|k_0}\|$. Thus from the previous relations we have: $|b(\hat{\mathbf{x}}_{i|k_0}, k_0 + i) - b(\mathbf{x}_{i|k_0}, k_0 + i)| \leq \mathcal{L}_b^x \mathcal{L}_d \sum_{j=0}^{i-1} (\mathcal{L}_f^x)^j$. This implies that:

$$b(\mathbf{x}_{i|k_0}, k_0 + i) \geq b(\hat{\mathbf{x}}_{i|k_0}, k_0 + i) - \mathcal{L}_b^x \mathcal{L}_d \sum_{j=0}^{i-1} (\mathcal{L}_f^x)^j \quad (11)$$

By assumption, since $\{\mathbf{u}_{l|k_0}\}_{l=0}^{N-1} \in \hat{\mathcal{K}}_{N|k_0}(\mathbf{x}_{k_0}, k_0)$, $\{\mathbf{u}_{l|k_0}\}_{l=0}^{i-1}$ satisfies: $\delta \hat{b}_{i|k_0}(\mathbf{x}_{k_0}, \{\mathbf{u}_{l|k_0}\}_{l=0}^{i-1}, k_0) \geq 0$ for which $b(\hat{\mathbf{x}}_{i|k_0}, k_0 + i) \geq \mathcal{L}_b^x \mathcal{L}_d \sum_{j=0}^{i-1} (\mathcal{L}_f^x)^j$ (see (8), (9)). Combined with (11) yields: $b(\mathbf{x}_{i|k_0}, k_0 + i) \geq b(\hat{\mathbf{x}}_{i|k_0}, k_0 + i) - \mathcal{L}_b^x \mathcal{L}_d \sum_{j=0}^{i-1} (\mathcal{L}_f^x)^j \geq 0$ such that $\mathbf{x}_{i|k_0} = \mathbf{x}_{k_0+i} \in \mathcal{X}(k_0 + i)$. Application of this for all $i \in [1, N] \subset \mathbb{N}$ with $k = k_0 + i$ completes the proof. $\square$

It is important to emphasize that Proposition 1 does *not* require that b is a barrier function in any sense. This result is a sufficient condition for ensuring that in the finite horizon, the safety conditions will be met. The *N-step robust, discrete-time barrier function* will then ensure that the open-loop control from Proposition 1 can be applied for all future time to enforce safety. This will be presented in the next subsection.

B. Predictive Control

Here we present the robust predictive safety filter control. Consider h as an *N-step robust, discrete-time barrier function*. The robust predictive safety filter is defined by the following **safety filter problem**, respectively:

$$\mathcal{P}_{sf}(\mathbf{x}_k, k) = \underset{\mathbf{u}_{l|k}}{\operatorname{argmin}} \quad \|\mathbf{u}_{nom}(\mathbf{x}_k, k) - \mathbf{u}_{l|k}\| \quad (12a)$$

$$\text{s.t. } \forall l = 0, \dots, N-1 :$$

$$\hat{\mathbf{x}}_{0|k} = \mathbf{x}_k, \quad (12b)$$

$$\hat{\mathbf{x}}_{l+1|k} = \mathbf{f}(\hat{\mathbf{x}}_{l|k}, \mathbf{u}_{l|k}, l + k), \quad (12c)$$

$$\mathbf{u}_{l|k} \in \mathcal{U}, \quad (12d)$$

$$b(\hat{\mathbf{x}}_{l|k}, l + k) \geq \mathcal{L}_b^x \mathcal{L}_d \sum_{j=0}^{l-1} (\mathcal{L}_f^x)^j, \quad (12e)$$

$$h(\hat{\mathbf{x}}_{N|k}, k + N) \geq \mathcal{L}_h^x \mathcal{L}_d (\mathcal{L}_f^x)^{N-1} \quad (12f)$$

$$\{\mathbf{u}_{l|k}^*\}_{l=0}^N = \mathcal{P}_{sf}(\mathbf{x}_k, k) \quad (13)$$

Remark 3. *There are several differences between the proposed robust predictive safety filter and that of [12]. First, the proposed safety filter only requires the solution of one optimization problem whereas [12] requires an additional feasibility problem to be solved. Second, the robustness-related terms, i.e., $\mathcal{L}_b^x, \mathcal{L}_d^x, \mathcal{L}_h^x$, and $\mathcal{L}_f^x$ ensure any bounded disturbance cannot push the system outside of the safe set, which is not ensured in [12].*

The barrier function h acts as a terminal constraint and must be an N -step robust, discrete-time barrier function for robustness guarantees to hold. The h function must be constructed to satisfy the conditions of Definition 1, which will most likely result in a conservative terminal set. Thus for more aggressive, but safe performance, a designer would construct h conservatively, but then enlarge the prediction horizon, N . The constraints defined by b can then be satisfied without checking that b itself is a barrier function in any sense. The trade-off is that as N increases, more computation is required to check that safety is guaranteed, while the system state is able to deviate further from the conservative bound defined by h . For a time-invariant h , the linearization approach from [12] can be used to construct h to satisfy Definition 1.

C. Analysis

The following assumption is made regarding the terminal constraint function h :

Assumption 3. *The function $h : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}$ with corresponding sets $\mathcal{C}(k)$ and $\bar{\mathcal{K}}$ defined by (4) and (7), respectively, is an N -step robust, discrete-time barrier function.*

In the following theorem, we show that the proposed control is always feasible and ensures the states remain inside of the safe set $\mathcal{X}(k)$ for all time in the presence of disturbances. Let the robust safety set and robust terminal safe set be respectively defined by:

$$\mathcal{X}^l(k) = \{\mathbf{x} \in \mathbb{R}^n : b(\mathbf{x}, k) \geq \mathcal{L}_b^x \mathcal{L}_d^x \sum_{j=0}^{l-1} (\mathcal{L}_f^x)^j\},$$

$$\forall l \in \{0, \dots, N\} \quad (14)$$

$$\mathcal{X}_f^N(k) = \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}, k) \geq \mathcal{L}_h^x \mathcal{L}_d^x (\mathcal{L}_f^x)^{N-1}\} \quad (15)$$

The following assumption is required to ensure the robust safety sets are non-empty and that if the final state in the horizon lies in terminal robust safety set, then it also lies in the robust safety set, i.e., $\mathcal{X}_f^N(k+N) \subset \mathcal{X}^N(k+N)$, $\forall k \in \mathbb{N}$.

Assumption 4. *The following conditions hold:*

- 1) $\mathcal{X}^l(k+l) \neq \emptyset$, $\forall l \in \{0, \dots, N\}$, $\forall k \in \mathbb{N}$
- 2) $\mathcal{X}_f^N(k) \neq \emptyset$, $\forall k \in \mathbb{N}$
- 3) $\mathcal{X}_f^N(k+N) \subset \mathcal{X}^N(k+N)$, $\forall k \in \mathbb{N}$.

Remark 4. *The last condition in Assumption 4 requires that the robust safety region at prediction time N must contain the level set defined by the N -step robust, discrete-time barrier function. The idea here is that we treat $\mathcal{X}_f^N$ as the terminal set, and so if the system reaches this terminal set at time N , then*

it will satisfy the system constraints defined by $\mathcal{X}$ while taking into account the effects of the perturbation. This condition can be satisfied by shrinking $\mathcal{C}(k)$ appropriately.

The following theorem presents one of the main results and states that the proposed robust safety filter ensures the states remain inside of the safe region $\mathcal{X}(k)$ and feasibility of the proposed control for all time. The main differences between this result and that of [12] is that here we address robustness to perturbations and extend the approach to time-varying systems with time-varying constraints. The key in this proof is in handling the disturbances in all future predictions of the system by exploiting Proposition 1 and Definition 1.

Theorem 1. *Consider the system (2) satisfying Assumption 1 and suppose a nominal control law $\mathbf{u}_{nom} : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathcal{U}$ is given. Given functions $b : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}$ and $h : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}$ for which Assumption 2 holds, let $\bar{\mathcal{K}}$, $\delta \hat{\mathbf{b}}_{N|k}$, and $\mathcal{X}(k)$ be defined by (10), (9), and (3), respectively. Suppose Assumptions 3 and 4 hold. For (2) in closed-loop with (12), (13), and for any given $k_0 \in \mathbb{N}$, $\mathbf{x}_{k_0} \in \mathcal{X}(k_0)$ for which $\mathcal{P}_{sf}(\mathbf{x}_{k_0}, k_0)$ is feasible, the following statements hold:*

- 1) *The control defined by (12), (13) is feasible for all $k \geq k_0$.*
- 2) *$\mathbf{x}_k \in \mathcal{X}(k)$ for all $k \geq k_0$, $k \in \mathbb{N}$.*

Proof. Proposition 1 ensures that the open-loop application of $\{\mathbf{u}_{l|k}^*\}_{l=0}^{N-1}$ to (2) is safe, i.e., $\mathbf{u}_k = \mathbf{u}_{0|k}^*$, $\mathbf{u}_{k+1} = \mathbf{u}_{1|k}^*$, ..., $\mathbf{u}_{k+N-1} = \mathbf{u}_{N-1|k}^*$, ensures that $\mathbf{x}_\kappa \in \mathcal{X}(\kappa)$ for all $\kappa \in \{k, \dots, k+N\}$. Thus since $\mathcal{P}_{sf}(\mathbf{x}_{k_0}, k_0)$ is feasible, we thus know that $\mathbf{x}_{k_0+1} \in \mathcal{X}(k_0+1)$.

In the closed-loop response, we want to show that $\mathcal{P}_{sf}$ is always feasible so that we can repeatedly apply Proposition 1 to ensure safety holds for all future time. In order to show that feasibility of $\mathcal{P}_{sf}(\mathbf{x}_k, k)$ implies feasibility of $\mathcal{P}_{sf}(\mathbf{x}_{k+1}, k+1)$, we need to show that there exists a feasible control action at $k+1$.

We will denote the resulting state trajectory associated with $\{\mathbf{u}_{l|k}^*\}_{l=0}^{N-1}$ applied to the unperturbed system (12c) as $\{\hat{\mathbf{x}}_{l|k}\}_{l=0}^N$, i.e., this is the predicted system trajectory with the optimal control at time k . The feasible control action chosen is defined as follows:

$$\mathbf{u}_{l|k+1}^+ \in \begin{cases} \{\mathbf{u}_{l+1|k}^*\}, \forall l \in \{0, \dots, N-2\}, \\ \bar{\mathcal{K}}(\hat{\mathbf{x}}_{N-1|k+1}, k+N), \text{ for } l = N-1 \\ \quad \text{if } \hat{\mathbf{x}}_{N-1|k+1} \in \mathcal{C}(k+N), \\ \{\mathbf{0}\}, \text{ for } l = N-1 \text{ if } \hat{\mathbf{x}}_{N-1|k+1} \notin \mathcal{C}(k+N) \end{cases} \quad (16)$$

for which the chosen control action yields an unperturbed trajectory $\{\hat{\mathbf{x}}_{l|k+1}\}_{l=0}^{N-1}$, where $\hat{\mathbf{x}}_{l|k+1}$ is defined by (12c) starting at time $k+1$ with $\hat{\mathbf{x}}_{0|k+1} = \mathbf{x}_{k+1}$. Note that $\{\hat{\mathbf{x}}_{l|k+1}\}_{l=0}^{N-1}$ may not be equal to $\{\hat{\mathbf{x}}_{l|k}\}_{l=1}^N$ because $\hat{\mathbf{x}}_{0|k+1} = \mathbf{x}_{k+1}$, i.e., the initial condition in the prediction of the unperturbed state at time $k+1$ is the result of the *true* system dynamics.

The last step here is to show that the control action (16) yields $\hat{\mathbf{x}}_{l|k+1}$ that satisfy $b(\hat{\mathbf{x}}_{l|k+1}, k+l+1) \geq \mathcal{L}_b^x \mathcal{L}_d^x \sum_{j=0}^{l-1} (\mathcal{L}_f^x)^j$ for all $l \in \{0, \dots, N-1\}$ and $h(\hat{\mathbf{x}}_{N|k+1}, k+N+1) \geq \mathcal{L}_h^x \mathcal{L}_d^x (\mathcal{L}_f^x)^{N-1}$.

We first follow a similar approach to the proof of Proposition 1 and define the error between the unperturbed trajectory

at k , $\hat{\mathbf{x}}_{i+1|k}$, and the resulting trajectory from $\mathbf{u}_{i|k+1}^+$, $\hat{\mathbf{x}}_{i|k+1}$, for all $i \in \{0, \dots, N-1\}$ using the Lipschitz property of $\mathbf{f}$. Consider the following relation for $i \in \{1, \dots, N-1\}$:

$$\begin{aligned} \|\hat{\mathbf{x}}_{i|k+1} - \hat{\mathbf{x}}_{i+1|k}\| &= \|\mathbf{f}(\hat{\mathbf{x}}_{i-1|k+1}, \mathbf{u}_{i-1|k+1}^+, k+i) \\ &\quad - \mathbf{f}(\hat{\mathbf{x}}_{i|k}, \mathbf{u}_{i-1|k+1}^+, k+i)\| \\ &\leq \mathcal{L}_f^x \|\hat{\mathbf{x}}_{i-1|k+1} - \hat{\mathbf{x}}_{i|k}\| \end{aligned} \quad (17)$$

For each $i \in \{1, \dots, N-1\}$ we can repeat the application of (17) until $i = 1$, which yields: $\|\hat{\mathbf{x}}_{i|k+1} - \hat{\mathbf{x}}_{i+1|k}\| \leq (\mathcal{L}_f^x)^i \|\hat{\mathbf{x}}_{0|k+1} - \hat{\mathbf{x}}_{1|k}\|$. Now recall that $\hat{\mathbf{x}}_{0|k+1} = \mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_{0|k}^*, k) + \mathbf{d}(\mathbf{x}_k, \mathbf{u}_{0|k}^*, k)$ and $\hat{\mathbf{x}}_{1|k} = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_{0|k}^*, k)$. Substitution into the above inequality along with (5) yields: $\|\hat{\mathbf{x}}_{i|k+1} - \hat{\mathbf{x}}_{i+1|k}\| \leq (\mathcal{L}_f^x)^i \|\hat{\mathbf{x}}_{0|k+1} - \hat{\mathbf{x}}_{1|k}\| \leq (\mathcal{L}_f^x)^i \|\mathbf{f}(\mathbf{x}_k, \mathbf{u}_{0|k}^*, k) + \mathbf{d}(\mathbf{x}_k, \mathbf{u}_{0|k}^*, k) - \mathbf{f}(\mathbf{x}_k, \mathbf{u}_{0|k}^*, k)\| \leq (\mathcal{L}_f^x)^i \mathcal{L}_d$. Applying the Lipschitz property of b yields: $|b(\hat{\mathbf{x}}_{i|k+1}, k+i+1) - b(\hat{\mathbf{x}}_{i+1|k}, k+i+1)| \leq \mathcal{L}_b^x \mathcal{L}_d (\mathcal{L}_f^x)^i$, which implies that, for all $i \in \{0, \dots, N-1\}$:

$$b(\hat{\mathbf{x}}_{i|k+1}, k+i+1) \geq b(\hat{\mathbf{x}}_{i+1|k}, k+i+1) - \mathcal{L}_b^x \mathcal{L}_d (\mathcal{L}_f^x)^i \quad (18)$$

Recall that if $\mathcal{P}_{sf}(\mathbf{x}_k, k)$ is feasible, for $i \in \{0, \dots, N-2\}$ $b(\hat{\mathbf{x}}_{i+1|k}, k+i+1) \geq \mathcal{L}_b^x \mathcal{L}_d \sum_{j=0}^i (\mathcal{L}_f^x)^j$, which yields: $b(\hat{\mathbf{x}}_{i|k+1}, k+i+1) \geq b(\hat{\mathbf{x}}_{i+1|k}, k+i+1) - \mathcal{L}_b^x \mathcal{L}_d (\mathcal{L}_f^x)^i \geq \mathcal{L}_b^x \mathcal{L}_d \sum_{j=0}^i (\mathcal{L}_f^x)^j - \mathcal{L}_b^x \mathcal{L}_d (\mathcal{L}_f^x)^i = \mathcal{L}_b^x \mathcal{L}_d \sum_{j=0}^{i-1} (\mathcal{L}_f^x)^j$. Thus the chosen $\{\mathbf{u}_{l|k+1}^+\}_{l=0}^{N-3}$ and resulting $\{\hat{\mathbf{x}}_{l|k+1}\}_{l=0}^{N-2}$ satisfies (12e) for $l \in \{0, \dots, N-2\}$.

To satisfy (12e) for $l = N-1$, we note that since $\mathcal{X}_f^N(k+N) \subset \mathcal{X}^N(k+N)$ from Assumption 4, $h(\hat{\mathbf{x}}_{N|k}, k+N) \geq \mathcal{L}_h^x \mathcal{L}_d (\mathcal{L}_f^x)^{N-1} \implies b(\hat{\mathbf{x}}_{N|k}, k+N) \geq \mathcal{L}_b^x \mathcal{L}_d \sum_{j=0}^{N-1} (\mathcal{L}_f^x)^j$. Substitution into (18) for $i = N-1$ yields: $b(\hat{\mathbf{x}}_{N-1|k+1}, k+N) \geq b(\hat{\mathbf{x}}_{N|k}, k+N) - \mathcal{L}_b^x \mathcal{L}_d (\mathcal{L}_f^x)^{N-1} \geq \mathcal{L}_b^x \mathcal{L}_d \sum_{j=0}^{N-1} (\mathcal{L}_f^x)^j - \mathcal{L}_b^x \mathcal{L}_d (\mathcal{L}_f^x)^{N-1} = \mathcal{L}_b^x \mathcal{L}_d \sum_{j=0}^{N-2} (\mathcal{L}_f^x)^j$. Thus the chosen $\mathbf{u}_{N-2|k+1}^+$ and resulting $\hat{\mathbf{x}}_{N-1|k+1}$ satisfies (12e) for $l = N-1$.

To satisfy (12f), we substitute h with associated Lipschitz constant $\mathcal{L}_h^x$ (note the same analysis applies to h) into (18) for $i = N-1$, which yields: $h(\hat{\mathbf{x}}_{N-1|k+1}, k+N) \geq h(\hat{\mathbf{x}}_{N|k}, k+N) - \mathcal{L}_h^x \mathcal{L}_d (\mathcal{L}_f^x)^{N-1} \geq \mathcal{L}_h^x \mathcal{L}_d (\mathcal{L}_f^x)^{N-1} - \mathcal{L}_h^x \mathcal{L}_d (\mathcal{L}_f^x)^{N-1} \geq 0$. Thus $\hat{\mathbf{x}}_{N-1|k+1} \in \mathcal{C}(k+N)$ and $\bar{\mathcal{K}}(\hat{\mathbf{x}}_{N-1|k+1}, k+N) \neq \emptyset$ since h is an N -step robust, discrete-time barrier function. Thus from (16), $\mathbf{u}_{N-1|k+1}^+ \in \bar{\mathcal{K}}(\hat{\mathbf{x}}_{N-1|k+1}, k+N)$. By definition, this implies that $h(\hat{\mathbf{x}}_{N|k+1}, k+N+1) = h(\mathbf{f}(\hat{\mathbf{x}}_{N-1|k+1}, \mathbf{u}_{N-1|k+1}^+, k+N), k+N+1) \geq \mathcal{L}_h^x \mathcal{L}_d (\mathcal{L}_f^x)^{N-1}$. Thus the chosen control $\mathbf{u}_{N-1|k+1}^+$ and resulting $\hat{\mathbf{x}}_{N|k+1}$ satisfies (12f).

We have thus constructed a feasible solution of $\{\mathbf{u}_{l|k+1}^+\}_{l=0}^{N-1}$ for which all the constraints of $\mathcal{P}_{sf}(\mathbf{x}_{k+1}, k+1)$ are satisfied. This implies that at the next time step, there exists a control such that $\mathcal{P}_{sf}$ is feasible and $\mathbf{x}_{k+1} \in \mathcal{X}(k+1)$. Since $\mathcal{P}_{sf}(\mathbf{x}_{k_0}, k_0)$ is feasible, this process can be repeated by induction for all $k \geq k_0$ such that the control is feasible for all $k \geq k_0$ and $\mathbf{x}_k \in \mathcal{X}(k)$ for all $k \geq k_0$, which concludes the proof. $\square$

Theorem 1 ensures safety for all time. There is no explicit restriction on the size of N in order to ensure safety. Adjusting

N simply provides a trade-off between computational complexity and conservatism. Smaller values of N result in smaller optimization problems to be solved and smaller robustness margins, however this will yield a more conservative control law since the system state will be forced to remain close to the robust terminal safe set. On the other hand, larger values of N result in larger optimization problems, but allow for states to deviate from the robust terminal safe set so long as at time N , the system predictions can be returned to the robust terminal safe set. Also, note that as N increases the robustness margins increase, which may violate Assumption 4.

Remark 5 (Multiple constraints). *Note that in $\mathcal{P}_{sf}$, only one constraint function, b , is enforced throughout the trajectory for readability. In practice, multiple $b_i : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}$ could be implemented in $\mathcal{P}_{sf}$ so long as each b_i satisfies the same requirements as b . To be clear, the following sets would be used in place of $\mathcal{X}(k)$ and $\mathcal{X}^l(k)$ for $i \in \{1, \dots, M\}$: $\mathcal{X}_i(k) := \{\mathbf{x} \in \mathbb{R}^n : b_i(\mathbf{x}, k) \geq 0\}$, $\mathcal{X}(k) = \cap_{i=1}^M \mathcal{X}_i(k)$, $\mathcal{X}_i^l(k) := \{\mathbf{x} \in \mathbb{R}^n : b_i(\mathbf{x}, k) \geq \mathcal{L}_b^x \mathcal{L}_d \sum_{j=0}^{l-1} (\mathcal{L}_f^x)^j\}$, $\mathcal{X}^l(k) = \cap_{i=1}^M \mathcal{X}_i^l(k)$.*

IV. EXTENSIONS AND IMPLEMENTATION

Here we make modifications to the robust predictive safety filter to reduce online computation. We first extend the robust predictive safety filter using an even-triggered scheme. We then use the results of the predictive safety filter to provide an efficient 1-step safety filter for existing barrier function implementations.

A. Event-triggered control

We propose an event-triggered control policy that implements the robust predictive safety filter as long as the system is safe and implements the predictive control policy (13) otherwise to ensure safety. This control policy is defined in Algorithm 1.

Algorithm 1 Event-Triggered, Robust Safety Filter

- 1: Given: $\mathbf{x}_k, k$.
 - 2: Initialize $\mathbf{x}_{0|k} = \mathbf{x}_k$.
 - 3: **for** $l \in \{0, \dots, N-1\}$ **do**
 - 4: Compute $\mathbf{u}_{nom_{l|k}} = \mathbf{u}_{nom}(\mathbf{x}_{l|k}, k+l)$.
 - 5: Compute $\hat{\mathbf{x}}_{l+1|k} = \mathbf{f}(\hat{\mathbf{x}}_{l|k}, \mathbf{u}_{nom_{l|k}}, k+l)$.
 - 6: **end for**
 - 7: Define $\{\mathbf{u}_{nom_{l|k}}\}_{l=0}^{N-1}$.
 - 8: Define the roll-out $\{\hat{\mathbf{x}}_{l|k}\}_{l=0}^N$.
 - 9: **if** $\hat{\mathbf{x}}_{l|k} \in \mathcal{X}^l(k)$ for $l \in \{0, \dots, N-1\}$ and $\hat{\mathbf{x}}_{N|k} \in \mathcal{X}_f^N(k)$ **then**
 - 10: **return** $\mathbf{u}_{nom_{0|k}}$.
 - 11: **else**
 - 12: Solve (13) for $\{\mathbf{u}_{l|k}^*\}_{l=0}^{N-1}$.
 - 13: **return** $\mathbf{u}_{0|k}^*$.
 - 14: **end if**
-

Corollary 1. *Suppose the conditions of Theorem 1 hold. Given a nominal control law $\mathbf{u}_{nom} : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathcal{U}$, suppose the system (2) is in closed-loop with the control from Algorithm*

1. If $\mathcal{P}_{sf}(\mathbf{x}_{k_0}, k_0)$ is feasible for any $\mathbf{x}_{k_0} \in \mathcal{X}(k_0)$, $k_0 \in \mathbb{N}$, then $\mathbf{x}_k \in \mathcal{X}(k) \forall k \geq k_0, k \in \mathbb{N}$.

Proof. Follows directly from the proof of Theorem 1. $\square$

Remark 6. The predictive safety filter is solely focused on enforcing safety. This allows for more modularity in its use with existing control laws over MPC methods that solve safety and performance in one control law. By using prediction only for safety, the event-triggering approach in Algorithm 1 is straightforward and can be used on compute-limited systems. This is advantageous when coupled with learning-based controllers which require low onboard computational resources to implement, but do not guarantee safety for general systems.

B. Efficient 1-step Robust Safety Filter

The robust predictive safety filter is dependent on the existence of an N -step robust, discrete-time barrier function, which can be difficult to synthesize or verify. In special cases, we can check for the robust, discrete-time barrier function condition without checking $\hat{\mathcal{K}}$ over all $\mathcal{C}(k)$. For such cases we restrict our attention to 1-step safety filters typically seen in the literature. Here we show that under stronger conditions, we can simplify the event-triggering conditions and the design of h .

Let $\mathcal{A}(k) \subset \mathbb{R}^n$ be a set encompassing the boundary of $\mathcal{C}(k)$, defined by $\partial\mathcal{C}(k) = \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}, k) = 0\}$, which is defined by:

$$\mathcal{A}(k) = \{\mathbf{x} \in \mathbb{R}^n : h(\mathbf{x}, k) \in [0, a]\} \quad (19)$$

for some $a \in \mathbb{R}_{>0}$. The set $\mathcal{A}(k)$ defines the region around the controlled safe set boundary, which is the only region where we care about enforcing the safe set condition, i.e., $h \geq 0$. In this respect, we define the set of admissible control inputs for which $h \geq 0$ holds if $\mathbf{x} \in \mathcal{A}(k)$:

$$\mathcal{K}(\mathbf{x}, k) := \{\mathbf{u} \in \mathcal{U} : \delta \bar{h}_{N=1}(\mathbf{x}, \mathbf{u}, k) \geq 0, \text{ if } \mathbf{x} \in \mathcal{A}(k)\} \quad (20)$$

We note that in this case, for any k and $\mathbf{x} \in \mathcal{C}(k) \setminus \mathcal{A}(k)$, then $\mathcal{K}(\mathbf{x}, k) = \mathcal{U}$ and so we need only check $\mathcal{A}(k)$ to determine if h satisfies the conditions of Definition 1 for $N = 1$.

First, in the next theorem we ensure that for a 1-step robust, discrete-time barrier function, forward invariance holds for any $\mathbf{u} \in \hat{\mathcal{K}}$.

Theorem 2. Consider the system (2) satisfying Assumption 1 and let $N = 1$. Given functions $b : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}$ and $h : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathbb{R}$, let $\hat{\mathcal{K}}$ and $\mathcal{X}(k)$ be defined by (10) and (3), respectively. Suppose Assumptions 2.2 and 3 hold and $\mathcal{C}(k) \subset \mathcal{X}(k) \forall k \in \mathbb{N}$. If $\mathbf{x}_{k_0} \in \mathcal{C}(k_0) \subset \mathcal{X}(k_0)$ for some $k_0 \in \mathbb{N}$ and $\mathbf{u}_k \in \hat{\mathcal{K}}(\mathbf{x}_k, k)$ is implemented in closed-loop with (2), then $\mathbf{x}_k \in \mathcal{C}(k) \subset \mathcal{X}(k)$ for all $k \geq k_0, k \in \mathbb{N}$.

Proof. Due to the restriction to $N = 1$, this proof is slightly different to that of Theorem 1. The main difference is that here we do not require any predictive component related to b . For clarity, we present a simplified proof based on the similar approach from Theorem 1. From the proof of Proposition 1, it is clear that for any $\mathbf{x}_k \in \mathcal{C}(k)$, the local Lipschitz property of h

holds for $\mathbf{x}_{k+1}$. Thus by the Lipschitz property, the following holds for any $k \in \mathbb{N}$: $|h(\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, k) + \mathbf{d}(\mathbf{x}_k, \mathbf{u}_k, k), k+1) - h(\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, k), k+1)| \leq \mathcal{L}_h^x \|\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, k) + \mathbf{d}(\mathbf{x}_k, \mathbf{u}_k, k) - \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, k)\| = \mathcal{L}_h^x \|\mathbf{d}(\mathbf{x}_k, \mathbf{u}_k, k)\| \leq \mathcal{L}_h^x \mathcal{L}_d$. It follows then that $h(\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, k), k+1) - \mathcal{L}_h^x \mathcal{L}_d \leq h(\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, k) + \mathbf{d}(\mathbf{x}_k, \mathbf{u}_k, k), k+1) = h(\mathbf{x}_{k+1}, k+1)$. Thus from (6), $h(\mathbf{x}_{k+1}, k+1) \geq h(\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, k), k+1) - \mathcal{L}_h^x \mathcal{L}_d \geq 0$. Thus $\mathbf{x}_{k+1} \in \mathcal{C}(k+1)$. Now it is clear that since $\mathbf{x}_{k_0} \in \mathcal{C}(k_0)$, then by induction $\mathbf{x}_k \in \mathcal{C}(k) \subset \mathcal{X}(k)$ for all $k \geq k_0$. $\square$

Next, we show that under stronger conditions, we can enforce safety by only implementing a control in $\mathcal{A}(k)$.

Theorem 3. Suppose the conditions of Theorem 2 hold with $\mathcal{K}(\mathbf{x}, k)$ defined by (20) in place of $\hat{\mathcal{K}}$, $\mathcal{A}(k)$ defined by (19), and h is 1-step bounded on $\mathbb{N}$ with $\mathcal{L}_h^k \in \mathbb{R}_{\geq 0}$. Further suppose there exists $\mathcal{L}_f \in \mathbb{R}_{>0}$ such that the following holds for all $\mathbf{x} \in \mathcal{Y}$, $\mathbf{u} \in \mathcal{U}$, $k \in \mathbb{N}$:

$$\|\mathbf{f}(\mathbf{x}, \mathbf{u}, k) - \mathbf{x}\| \leq \mathcal{L}_f, \quad (21)$$

for which $a \geq \mathcal{L}_h^x(\mathcal{L}_f + \mathcal{L}_d) + \mathcal{L}_h^k$, and the resulting $\mathcal{A}(k) \neq \emptyset \forall k \in \mathbb{N}$. If $\mathbf{x}_{k_0} \in \mathcal{C}(k_0) \subset \mathcal{X}(k_0)$ for some $k_0 \in \mathbb{N}$, and $\mathbf{u}_k \in \mathcal{K}(\mathbf{x}_k, k)$ is implemented in closed-loop with (2), then $\mathbf{x}_k \in \mathcal{C}(k) \subset \mathcal{X}(k)$ for all $k \geq k_0, k \in \mathbb{N}$.

Proof. Since the conditions of Theorem 2 hold on $\mathcal{A}(k) \subset \mathcal{C}(k)$, we need only consider the case when $\mathbf{x}_k \in \mathcal{C}(k) \setminus \mathcal{A}(k)$ (i.e. $h(\mathbf{x}_k, k) > a$). Let $\bar{\mathbf{f}}_k = \mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, k) + \mathbf{d}(\mathbf{x}_k, \mathbf{u}_k, k)$ to simplify the notation, for which (21) and the triangle inequality yields: $\|\bar{\mathbf{f}}_k - \mathbf{x}_k\| = \|\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, k) + \mathbf{d}(\mathbf{x}_k, \mathbf{u}_k, k) - \mathbf{x}_k\| \leq \|\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, k) - \mathbf{x}_k\| + \|\mathbf{d}(\mathbf{x}_k, \mathbf{u}_k, k)\| \leq \mathcal{L}_f + \mathcal{L}_d$. Using the following conditions: $|h(\mathbf{y}, k) - h(\mathbf{x}, k)| \leq \mathcal{L}_h^x \|\mathbf{y} - \mathbf{x}\|$, $|h(\mathbf{x}, k+1) - h(\mathbf{x}, k)| \leq \mathcal{L}_h^k$ and the triangle inequality, we arrive at the following inequality: $|h(\bar{\mathbf{f}}_k, k+1) - h(\mathbf{x}_k, k)| = |h(\bar{\mathbf{f}}_k, k+1) - h(\bar{\mathbf{f}}_k, k) + h(\bar{\mathbf{f}}_k, k) - h(\mathbf{x}_k, k)| \leq |h(\bar{\mathbf{f}}_k, k) - h(\mathbf{x}_k, k)| + |h(\bar{\mathbf{f}}_k, k+1) - h(\bar{\mathbf{f}}_k, k)| \leq \mathcal{L}_h^x \|\bar{\mathbf{f}}_k - \mathbf{x}_k\| + \mathcal{L}_h^k \leq \mathcal{L}_h^x(\mathcal{L}_f + \mathcal{L}_d) + \mathcal{L}_h^k \leq a$. From this relation, it follows that $h(\mathbf{x}_{k_0+1}, k_0+1) \geq h(\mathbf{x}_{k_0}, k_0) - a > 0$. Since $\mathbf{x}_{k_0} \in \mathcal{C}(k_0)$, then by induction $\mathbf{x}_k \in \mathcal{C}(k)$ for all $k \geq k_0$. $\square$

Remark 7. The condition (21) may seem highly restrictive at first, however many systems in fact satisfy this condition. For example, if $\mathcal{Y}$, i.e., the set that contains $\mathcal{X}(k)$ for all time, is compact and $\mathbf{f}$ is time-invariant and a continuous function on its domain then in fact (21) must hold, since $\mathbf{x}$ and $\mathbf{f}$ are both bounded on the respective sets. Continuity of $\mathbf{f}$ is a common assumption for many discrete-time systems [12] and in many instances the safe set is compact [6], [12]. In such cases, single integrators and mechanical/robotic systems [6] satisfy (21). The 1-step bound condition only restricts how much h can change with respect to time, between time steps. This condition is needed to not require checking the barrier condition in $\mathcal{C}(k) \setminus \mathcal{A}(k)$.

Remark 8. Theorem 3 provides a new condition of forward invariance for discrete-time systems that does not require checking the condition (6) for all states in $\mathcal{C}(k)$. This extends the standard results of forward invariance from [38] (for which Theorem 3 provides a sufficient condition) by requiring stronger conditions of h and the system dynamics. Theorem

3 provides a similar notion of forward invariance used in continuous-time systems, for which one need-only check how the system behaves near the constraint boundary to ensure safety.

The predictive safety control can be implemented in a 1-step look-ahead fashion, typical of many existing barrier function methods. In this case, the proposed control can be simplified into the following optimization problem:

$$\mathcal{P}_{sf}(\mathbf{x}_k, k) = \underset{\mathbf{u}_k}{\operatorname{argmin}} \quad \|\mathbf{u}_{nom}(\mathbf{x}_k, k) - \mathbf{u}_k\| \quad (22a)$$

$$\text{s.t.} \quad (22b)$$

$$\hat{\mathbf{x}}_{0|k} = \mathbf{x}_k, \quad (22c)$$

$$\hat{\mathbf{x}}_{1|k} = \mathbf{f}(\hat{\mathbf{x}}_{0|k}, \mathbf{u}_k, k), \quad (22d)$$

$$\mathbf{u}_k \in \mathcal{U}, \quad (22e)$$

$$h(\hat{\mathbf{x}}_{1|k}, k+1) \geq \mathcal{L}_h^x \mathcal{L}_d \quad (22f)$$

In the special case where the conditions of Theorem 3 hold, the simpler Algorithm 2 can be implemented. Note that Algorithm 2 is not the same as Algorithm 1 with $N = 1$. In Algorithm 2, no system dynamics roll-out is required. The only condition that needs to be checked is if $h(\mathbf{x}_k, k) > a$.

Algorithm 2 1-Step Event-Triggered, Robust Safety Filter

- 1: Given: $\mathbf{x}_k, k$.
 - 2: **if** $h(\mathbf{x}_k, k) > a$ **then**
 - 3: **return** $\mathbf{u}_k = \mathbf{u}_{nom}(\mathbf{x}_k, k)$.
 - 4: **else**
 - 5: Compute (13) with $\mathcal{P}_{sf}$ from (22) for $\mathbf{u}_k^*$.
 - 6: **return** $\mathbf{u}_k = \mathbf{u}_k^*$.
 - 7: **end if**
-

Corollary 2. Suppose the conditions of Theorem 3 hold. Given a nominal control law $\mathbf{u}_{nom} : \mathbb{R}^n \times \mathbb{N} \rightarrow \mathcal{U}$, suppose the system (2) is in closed-loop with the control from Algorithm 2. Then if $\mathbf{x}_{k_0} \in \mathcal{C}(k_0) \subset \mathcal{X}(k_0)$ for any $k_0 \in \mathbb{N}$, then $\mathbf{x}_k \in \mathcal{C}(k) \subset \mathcal{X}(k)$, $\forall k \geq k_0, k \in \mathbb{N}$.

Proof. Follows from Theorem 3. $\square$

V. NUMERICAL EXAMPLES

In this section, the proposed controllers from Algorithm 1 and 2 are implemented on several systems. The nominal control used was the learning-based differentiable predictive control (DPC) from [39], [40]. The DPC method in all presented examples was designed and trained using the Neuromancer scientific machine learning library [41]. All code was developed in Python using the CasADi optimization module [42] to solve $\mathcal{P}_{sf}$. For each example, the DPC control policy was parametrized by a multi-layer perceptron (MLP) of 2 layers with 32 internal states each and a Gaussian Error Linear Unit (GELU) activation. The MLP was implemented with a sigmoid scale method to ensure input constraints were satisfied.

A. Two-Tank Example

We consider the nonlinear, discretized two-tank system : $\bar{f}_1(\mathbf{x}, \mathbf{u}) := x_1 + \Delta t (c_1(1 - u_2)u_1 - c_2\sqrt{x_1})$, $\bar{f}_2(\mathbf{x}, \mathbf{u}) := x_2 + \Delta t (c_1u_1u_2 + c_2\sqrt{x_1} - c_2\sqrt{x_2})$ where $x_1, x_2 \in \mathbb{R}_{>0}$ are the height of the liquid in tank 1 and 2, respectively, $c_1 = 0.8$ is the inlet valve coefficient, $c_2 = 0.4$ is the outlet valve coefficient, and $u_1, u_2 \in \mathbb{R}_{\geq 0}$ are the pump and valve control terms, respectively. The sampling time used was $\Delta t = 0.1$ s. The objective is to keep the liquid between desired levels in each tank i.e., $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^2 : 0.2 \leq x_i \leq 1, i \in \{1, 2\}\}$ for $\mathbf{x} = [x_1, x_2]^T$, while respecting the input constraints defined by: $\mathcal{U} = \{\mathbf{u} : 0 \leq u_i \leq 1.0, \forall i \in \{1, 2\}\}$ and tracking a piece-wise constant reference trajectory $\mathbf{r} : \mathbb{N} \rightarrow \mathbb{R}^n$.


Note that some optimization modules, such as CasADi, have difficulty when non-smooth dynamics are included in the optimization. However, thanks to the robust nature of the proposed control, we re-write the two tank system with an approximation of $\sqrt{\cdot}$ as: $f_1(\mathbf{x}, \mathbf{u}) = x_1 + \Delta t (c_1(1 - u_2)u_1 - c_2\psi(x_1))$, $f_2(\mathbf{x}, \mathbf{u}) = x_2 + \Delta t (c_1u_1u_2 + c_2\psi(x_1) - c_2\psi(x_2))$, where $\psi : \mathbb{R} \rightarrow \mathbb{R}$ is a 7th order polynomial fit of $\sqrt{x}$ on $[0.2, 1]$, which yields:

$$\begin{aligned} x_{1,k+1} &= f_1(\mathbf{x}_k, \mathbf{u}_k) + (\bar{f}_1(\mathbf{x}_k, \mathbf{u}_k) - f_1(\mathbf{x}_k, \mathbf{u}_k)) + w_1(k) \\ x_{2,k+1} &= \underbrace{f_2(\mathbf{x}_k, \mathbf{u}_k)}_{\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k)} + \underbrace{(\bar{f}_2(\mathbf{x}_k, \mathbf{u}_k) - f_2(\mathbf{x}_k, \mathbf{u}_k)) + w_2(k)}_{\mathbf{d}(\mathbf{x}_k, \mathbf{u}_k, k)} \end{aligned}$$


where $w_1(k) = w_2(k) = \bar{w} \sin(k)$ are additional perturbations acting on the system for $\bar{w} \in \mathbb{R}_{>0}$. We define $h(\mathbf{x}) = \varepsilon - (x_1 - x_{r_1})^2 + \rho(x_2 - x_{r_2})^2$, $\mathbf{x}_r = [x_{r_1} = 0.63, x_{r_2} = 0.63]^T$, $b_1(\mathbf{x}) = x_{max} - x_1$, $b_2(\mathbf{x}) = x_1 - x_{min}$, $b_3(\mathbf{x}) = x_{max} - x_2$, $b_4(\mathbf{x}) = x_2 - x_{min}$, where $x_{min} = 0.2$, $x_{max} = 1.0$, $\rho = 2.69$, and b_1, b_2, b_3, b_4 are used as discussed in Remark 5 and note that these all have the same Lipschitz constant $\mathcal{L}_b^x$. The function h is shown to be an N -step robust, discrete-time barrier function in Appendix VII-B for $N = 20$. The parameters used for the first simulation of the two-tank are as follows: $\bar{w} = 0.00001$, $\mathcal{L}_f^x = 1.205$, $\mathcal{L}_b^x = 1.0$, $\mathcal{L}_h^x = 1.331$, $\mathcal{L}_d = 0.0000542$, $\varepsilon = 0.12$ and $\mathbf{x}_r = [0.63, 0.63]^T$. Note that $\mathcal{L}_d$ was determined by combining $\bar{w}$, the additional perturbation, with the computed bounded error between $\psi(x)$ and $\sqrt{x}$ on $[0.2, 1.0]$. The DPC control law for the two-tank system was designed to track the time-varying reference and penalty functions were used to address the system constraints.

In the first set of simulations, the DPC control, denoted ‘DPC’ was implemented alone in one scenario and the proposed control (1), denoted ‘DPC+SF’, was implemented in the second scenario with $N = 20$. For the chosen N , the conditions of Assumption 4 hold. The results of the simulation are shown in Figure 2.


The plots in Figure 2 show that the ‘DPC’ control is able to track the reference trajectory reasonably well, but violates the system constraints. The constraint violations occur at time steps $k = 24, 324$ and 493 . The ‘DPC+SF’ control, on the other hand, implements the ‘DPC’ control as close as possible, but then overrides the ‘DPC’ control (see Figure 2c) to ensure the states remain inside their constraint bounds (see Figure 2a). Furthermore, the event-triggering from Algorithm 1 is shown in Figure 2b for which the optimization problem from the



(a) State trajectories with constraint bounds (black, solid lines).



(b) Events triggered (1 is triggered, 0 is not triggered) for each simulation scenario. Note the 'DPC' scenario is not shown as no safety filter is implemented.



(c) Input trajectories with constraint bounds (black, solid lines).

Fig. 2: (Two-tank example) Comparison of trajectories for 'DPC' and 'DPC+SF' implementations with the proposed robust safety filter on the perturbed two-tank system with $\bar{w} = 0.00001$.

safety filter is only solved when the system states move close to the constraint boundary. These results show the advantage of the proposed methodology from Algorithm 1 in enforcing system constraints without requiring the optimization problems to be solved at every instant when implemented online.

Although these results are promising, there is a trade-off between the size of the optimization problems from the safety filter and the magnitude of the disturbance that can be tolerated by the system. In other words, as N increases the robustness margins from (14) and (15) increase, which restricts

the magnitude of $\mathcal{L}_d$ that can be tolerated, and vice versa. To show this, we implemented the proposed control with a higher perturbation, $\bar{w} = 0.001$ for which $\mathcal{L}_d = 0.00145$, and resulted in a smaller prediction horizon of $N = 6$. The results of this simulation are shown in Figure 3. The plots show more conservative behavior in the 'DPC+SF' case compared to the previous simulation, especially between $k = 14$ and $k = 150$ and $k = 301$ and $k = 451$. Within both of these time intervals, the shorter prediction horizon requires the system to stay closer to $\mathcal{C}$ to satisfy the safety conditions. Note however that the proposed control is able to reject the larger perturbation here and keep the system safe.

Finally, to showcase the approach compared to the existing predictive safety filter of [12], we implement that predictive safety filter on the previous example with $\bar{w} = 0.001$. The same parameters used for the safety filter from Algorithm 1 were used for the safety filter of [12] for a direct comparison in addition to $\Delta = 0.00000005$ (see [12]). The results of the safety filter implementation are shown in Figure 4, which shows similar performance to the proposed control in Figure 3, except for the state constraint violation that occurs between $k = 476$ and $k = 53$ (see Figure 4b). Furthermore, the control from [12] does not include an event-triggering and so must be solved at all instances in time. As a result, the proposed control was 66.0% faster in online implementation compared to that of [12] (see Figure 3b). These results show that the proposed control is beneficial for handling bounded perturbations to guarantee hard safety constraints at all times, while reducing the online computations required.


B. Building Example

Here we apply the proposed method to a single zone building temperature control problem from [43]. This is a nonlinear, time-varying system that incorporates an estimate of the ambient disturbance affecting the building, defined as follows: $\mathbf{x}_{k+1} = \underbrace{A\mathbf{x}_k + B\mathbf{u}_k + \hat{\mathbf{w}}(k)}_{\mathbf{f}(\mathbf{x}_k, \mathbf{u}_k, k)} + \underbrace{(\mathbf{w}(k) - \hat{\mathbf{w}}(k))}_{\mathbf{d}(k)}$, with

$$A = \begin{bmatrix} 0.995 & 0.0017 & 0.0 & 0.0031 \\ 0.0007 & 0.996 & 0.0003 & 0.0031 \\ 0.0 & 0.0003 & 0.983 & 0.0 \\ 0.202 & 0.488 & 0.01 & 0.257 \end{bmatrix}, B = \begin{bmatrix} 0.00000176 \\ 0.00000176 \\ 0.0 \\ 0.000506 \end{bmatrix}$$

where $\mathbf{w} : \mathbb{N} \rightarrow \mathbb{R}^n$ is the perturbation of the environment affecting the system, and $\hat{\mathbf{w}} : \mathbb{N} \rightarrow \mathbb{R}^n$ is the estimate of $\mathbf{w}$, which is commonly approximated using annual statistical information of the building and surrounding region (see Figure 5). The control $\mathbf{u}$ is the heat flow of the HVAC system with input constraints set: $\mathcal{U} = \{\mathbf{u} \in \mathbb{R} : 0 \leq u \leq 5000\}$, and the number of system states is $n = 6$. Note that these dynamics were developed for a sampling rate of 0.1 s. For this example, $\mathcal{L}_d = 0.005$ and $\mathbf{d}(k) = -\mathcal{L}_d \mathbf{1}_n$ was used to demonstrate the worst-case disturbance that is constantly lowering the temperature of the building.


The objective is to minimize the energy usage, i.e., the control input, while remaining within time-varying constraints defined by: $b_1(\mathbf{x}, k) = -C\mathbf{x} + \bar{r}(k)$, $b_2(\mathbf{x}, k) = C\mathbf{x} - \underline{r}(k)$,



(a) State trajectories with constraint bounds (black, solid lines).




(b) Events triggered (1 is triggered, 0 is not triggered) for each simulation scenario. Note the 'DPC' scenario is not shown as no safety filter is implemented.




(c) State trajectories (zoomed-in) with constraint bounds (black, solid lines). Shows proposed safety filter satisfying system constraints.

Fig. 3: (Two-tank example) Comparison of trajectories for 'DPC' and 'DPC+SF' implementations with the proposed robust safety filter on the perturbed two-tank system with $\bar{w} = 0.001$.



(a) State trajectories with constraint bounds (black, solid lines).



(b) State trajectories (zoomed-in) with constraint bounds (black, solid lines). Original safety filter violates safety constraints.

Fig. 4: (Two-tank example) Results for predictive safety filter of [12] for 'DPC' and 'DPC+SF' implementations on the perturbed two-tank system with $\bar{w} = 0.001$.

where $C = [0, 0, 0, 1]$ and $\bar{r}, \underline{r} : \mathbb{N} \rightarrow \mathbb{R}$ define the time-varying comfort constraint bounds. The constraint functions were implemented as per Remark 5.

Here we take a different approach to the previous example and show how the proposed control could be implemented from a practical perspective. First, it is reasonable to assume that the building HVAC system is capable of keeping the building room temperature within a temperature range despite reasonable fluctuations in the ambient temperature. Here we use a temperature range from 19.55 C to 20.45 C. We also assume that if the output of the system, $C\mathbf{x}$, which represents the temperature of the single zone, is bounded, then the states are also bounded. Thus we effectively assume the following function is a *N-step robust, discrete-time barrier function*: $h(\mathbf{x}) = \varepsilon - \|C\mathbf{x} - y_r\|_2^2$, where $\varepsilon = 0.9$, $y_r = 20$ C is the reference temperature, and $N = 6$. Note that neither




Fig. 5: Estimate of environmental perturbations on the building, $\hat{w}(k)$.

the state constraint sets: $\mathcal{X}(k)$ nor the terminal constraint set $\mathcal{C}$ are compact, however compactness is not required for the results of Theorem 1. We define the robustness margins used: $\mathcal{L}_f^x = 0.9998$, $\mathcal{L}_b^x = 1.0$, $\mathcal{L}_h^x = 100$. The margins for $\mathcal{L}_f^x$ and $\mathcal{L}_b^x$ were computed using standard methods of computing Lipschitz constants [35]. For $\mathcal{L}_h^x$, since h is quadratic, we use the assumption that the states remain bounded if the temperature is bounded in the terminal set to make a conservative estimate of the bound of $\|x\|$ for all $x \in \mathcal{C}$.

From Figure 6, it is clear that the sets from Assumption 4 are non-empty and that the robust safety set with respect to the terminal barrier function, i.e., $\mathcal{X}_f^N(k+N)$, is a subset of the robust safety set with respect to the state constraints, i.e., $\mathcal{X}^N(k+N)$. Thus the conditions of Assumption 4 are satisfied. The DPC control law was designed using a regulation loss term to minimize the control input of the system. In addition, a smoothness term was added to prevent sharp changes in the control input, and a penalty loss was used to address state constraints (see [39]).

The proposed control from Algorithm 1 is compared to using the DPC control alone. We refer to the control from Algorithm 1 as ‘DPC+SF’, and DPC as simply ‘DPC’. Both controllers were run with and without the perturbation term $d(k)$. The results are shown in Figure 7.

Figure 7a shows the comparisons of the various different simulations. The unperturbed cases show that the ‘DPC+SF’ and ‘DPC’ controllers overlap throughout the entire simulation in both the temperature and input trajectories. This indicates that in the absence of uncertain disturbances, DPC is able to respect the safety constraints. If the safety constraints were ever to be violated, the safety filter would have altered the DPC control. This result is expected as DPC is known for (probabilistically) ensuring constraint satisfaction in the absence of perturbations. This is not the case with the perturbed simulations. The trajectories associated with the dynamics subject to $d(k)$ show that the safety filter (see the ‘DPC+SF’ trajectories) must alter the original DPC control to remain





Fig. 6: (Building example) Comfort constraints (black, solid lines) along with terminal constraint (blue, solid lines). The sets from Assumption 4 are depicted as: $\mathcal{X}^l(k+l)$ for $l \in \{0, \dots, N\}$ is the region between the gray, dashed lines depicted at intervals of 50 time steps, $\mathcal{X}_f^N(k)$ is the region between the cyan, dotted lines.

within the comfort constraints. The ‘DPC+SF’ remains within both the comfort constraints and input constraints. The ‘DPC’ control on the other hand violates the comfort constraints (see Figure 7a). These results show that the proposed control, ‘DPC+SF’ is able to robustly handle uncertain perturbations, while remaining point-wise, minimally close to the DPC control law. Note that in this example, the objective was to reduce the amount of energy used, while staying within the comfort constraints, which resulted in trajectories close to the system constraints. As a result, the simulation resulted in events triggering at every time step. This highlights an important aspect of the event-triggering component in that if the nominal control aims to keep the system close to the constraint bounds, the safety filter will need to be solved at all time steps.


C. Simple Example (1-step Safety Filter)

In this simple example, we apply Algorithm 2 to the perturbed single integrator discretized using the standard Euler method: $x_{k+1} = \underbrace{x_k + \Delta t(u_k)}_{f(x_k, u_k, k)} + \underbrace{\Delta t w(k)}_{d(x_k, u_k, k)}$, where $\Delta t = 0.01$ s is the sampling time and $w(k) = 0.02 \sin(0.5k)$ is a perturbation on the system. The input constraint set is: $\mathcal{U} = \{u \in \mathbb{R} : |u| \leq 10\}$. The barrier function used here is: $h(x, k) = \varepsilon - (x - x_r(k))^2$, for $x_r(k) = 0.5 \sin(0.05k)$. The objective is to keep the system inside the time-varying tube defined by $\mathcal{C}(k) \subseteq \mathcal{X}(k)$, similar to the continuous-time version from [6].


To apply Algorithm 2, we need to ensure that h is a 1-step robust, discrete-time barrier function and that the conditions of Theorem 3 hold. The Lipschitz constants and bounds from Theorem 3 are computed by using standard methods for which $\mathcal{L}_h^x = 0.894$, $\mathcal{L}_h^k = 0.072$, $\mathcal{L}_d = 0.02$, $\mathcal{L}_f = 0.11$, $\mathcal{L}_f^x =$



(a) Temperature trajectories with comfort constraints (black, solid lines).



(b) Temperature trajectories (zoomed-in) with comfort constraints (black, solid lines).



(c) Input trajectories with constraint bounds (black, solid lines).

Fig. 7: (Building example) Comparison of trajectories for ‘DPC+SF’ with perturbation, ‘DPC’ with perturbation, ‘DPC+SF’ without perturbation, and ‘DPC’ without perturbation.





Fig. 8: (Simple example) The barrier function safe set, $\mathcal{C}(k)$, is the region between the blue, solid lines. The set $\mathcal{X}_f^N(k)$ is the region between the cyan, dotted lines. The set $\mathcal{A}(k)$ is the region between the pink solid lines and the blue solid lines.

1.0. See Appendix VII-A to check that h is a *1-step robust, discrete-time barrier function*. We chose $a = \mathcal{L}_h^x(\mathcal{L}_f + \mathcal{L}_d) + \mathcal{L}_h^k = 0.180$. Since $0 < a < \varepsilon$ and $a > 0$, $\mathcal{A}(k) \neq \emptyset$ and $\mathcal{A}(k) \subset \mathcal{C}(k)$. All the conditions of Theorem 3 have been met. Figure 8 shows the sets associated with this example. In this example, the DPC control was used to reduce the amount of online computations needed to satisfy the objective. DPC was tasked with tracking the center of the tube to avoid triggering the event in Algorithm 2.


Algorithm 2 was implemented for this example along with a ‘no control’ case denoted ‘ $u_{nom} = 0$ ’. This yields three different controllers. First is the ‘ $u_{nom} = 0$ ’ control in which $u(x, k) = u_{nom}(x, k) = 0$ is implemented in closed-loop with the system. Second is the ‘ $u_{nom} = 0 + SF$ ’ in which Algorithm 2 is implemented with $u_{nom}(x, k) = 0$. Third is the ‘DPC + SF’ in which Algorithm 2 is implemented as stated. Figure 9 shows the results of these three scenarios.

In Figure 9, several concepts are demonstrated. First, it is clear that the ‘ $u_{nom} = 0$ ’ case is not able to stay within the tube as it quickly leaves the safe set. The same control implemented in Algorithm 2, i.e., ‘ $u_{nom} = 0 + SF$ ’, matches the ‘ $u_{nom} = 0$ ’ control until the system attempts to leave the safe region and the safety filter acts to keep the system within the safe set. This demonstrates the basic concept that the safety filter can ensure constraint satisfaction.


Second, the results show the impact of combining the DPC control with the SF in Algorithm 2. In the ‘ $u_{nom} = 0 + SF$ ’ case, the safety filter gets triggered frequently which requires the system to continually solve the nonlinear program online. This is shown in Figure 9b. Note that the events only occur when the state enters $\mathcal{A}(k)$. In the ‘DPC + SF’ case, the safety filter is only solved within the first few time steps as the initial condition lies inside of $\mathcal{A}(k)$. However the DPC control is able to keep the system centered in the safe set which results in no further events triggering for the duration of the simulation.



(a) State trajectories with barrier function safe set (blue, solid lines).



(b) Events triggered (1 is triggered, 0 is not triggered) for each simulation scenario. Note the case where $u_{nom} = 0$ is not shown as no safety filter is implemented.



(c) Input trajectories with constraint bounds (black, solid lines).

Fig. 9: (Simple example) Comparison of trajectories for ‘unom=0 + SF’ (red, solid curve), ‘unom=0’ (orange, dashed curve), ‘DPC+SF’ (magenta, dotted curve).

This shows that combining DPC, i.e., a learning-based control, with the safety filter allows for significant reduction in online computations, while still providing guarantees of safety in the presence of perturbations.

VI. CONCLUSION

In this paper, we present a robust predictive safety filter to provide a robustly safe control law. The approach extends the existing state-of-the-art by handling unknown perturbations and extending the safety results to time-varying systems with time-varying constraints. Furthermore, the proposed method

uses event-triggering to reduce the amount of online computation required. We discuss trade-offs in the design of the proposed control and implement it on a simple single integrator, a two-tank system, and a building system to demonstrate its efficacy.

The approach is dependent on the existence of an N -step robust, discrete-time barrier function, which for general time-varying systems is not straightforward to design. Existing methods can be used to design barrier functions for the time-invariant case, but a general method for time-varying systems is still an open problem and left for future work. Additionally, future work will extend the methodology to include adaptive model updates for reduced conservatism.

VII. APPENDIX

A. Barrier function check for the simple integrator

First, we define the following control to satisfy (6):

$$\hat{u}(x, k) = \begin{cases} 0, & \text{if } h(f(x, 0, k), k+1) \geq \mathcal{L}_h^x \mathcal{L}_d, \\ y(x, k) - \text{sign}(y(x, k)) \sqrt{\varepsilon - \mathcal{L}_h^x \mathcal{L}_d}, & \text{otherwise} \end{cases} \quad (23)$$

where $y(x, k) = \frac{1}{\Delta t}(-x + x_r(k+1))$. We leave it to the reader to see that (23) satisfies (6) for all $x \in \mathcal{C}(k)$, $\forall k \in \mathbb{N}$. Next we need to make sure this control satisfies the input constraints for which it can be shown that $|\hat{u}| \leq \frac{1}{\Delta t}(\mathcal{L}_h^k + \sqrt{\varepsilon} - \sqrt{\varepsilon - \mathcal{L}_h^x \mathcal{L}_d}) \leq 9.28$. Thus $\hat{u} \in \mathcal{U}$. Thus we have shown that h is a 1-step robust, discrete-time barrier function.

B. Barrier function check for the two-tank system

To show that h is a barrier function, we use a sampling-based approach. We first, define a feasible control to attempt to satisfy (6). This control consists of the following components: $y_1(x) = -\frac{1}{\Delta t}(x_1 - x_{r1}) + c_2\psi(x_1)$, $y_2(x) = -\frac{1}{\Delta t}(x_2 - x_{r2}) - c_2\psi(x_1) + c_2\psi(x_2)$. We can now define the proposed control noting that the condition (6) is only required when $x_k \in \mathcal{C}$. The condition (6) can be written as: $h(f(x, u)) = \varepsilon - \Delta t^2(c_1(1 - u_2)u_1 - y_1(x))^2 - \rho\Delta t^2(c_1u_1u_2 - y_2(x))^2 \geq \mathcal{L}_h^x \mathcal{L}_d(\mathcal{L}_f^x)^{N-1}$. Thus we can define an optimal point-wise control to satisfy the condition by: $\hat{u}(x) = \min_{u \in \mathcal{U}} (c_1(1 - u_2)u_1 - y_1(x))^2 + \rho(c_1u_1u_2 - y_2(x))^2$. The chosen control law to show that h is a barrier function is:

$$\bar{u}(x) = \begin{cases} 0, & \text{if } h(f(x_k, 0)) \geq \mathcal{L}_h^x \mathcal{L}_d(\mathcal{L}_f^x)^{N-1}, \\ \hat{u}(x), & \text{otherwise} \end{cases}$$

The control $\bar{u}$ was checked over a grid of 10,000 sampled states, equally spaced, containing $\mathcal{C}$ defined with endpoints $\{ [0.284, 0.284]^T, [0.284, 0.976]^T, [0.976, 0.284]^T, [0.976, 0.976]^T \}$. For the two scenarios presented in Section V-A, the sampling approach ensured that (6) held over the grid search, which suggested that h was an N -step robust, discrete-time barrier function.

REFERENCES

- [1] A. Ames, X. Xu, J. Grizzle, and P. Tabuada, "Control barrier function based quadratic programs for safety critical systems," *IEEE Transactions on Automatic Control*, vol. 62, no. 8, pp. 3861–3876, 2017.
- [2] T. X. Nghiem, J. Drgoňa, C. Jones, Z. Nagy, R. Schwan, B. Dey, A. Chakrabarty, S. Di Cairano, J. A. Paulson, A. Carron, M. N. Zeilinger, W. Shaw Cortez, and D. L. Vrabie, "Physics-informed machine learning for modeling and control of dynamical systems," in *American Control Conference*, 2023, pp. 3735–3750.
- [3] L. Brunke, M. Greeff, A. W. Hall, Z. Yuan, S. Zhou, J. Panerati, and A. P. Schoellig, "Safe learning in robotics: From learning-based control to safe reinforcement learning," *Annual Review of Control, Robotics, and Autonomous Systems*, vol. 5, no. Volume 5, 2022, pp. 411–444, 2022.
- [4] V. Freire and M. M. Nicotra, "Systematic Design of Discrete-Time Control Barrier Functions Using Maximal Output Admissible Sets," *IEEE Control Systems Letters*, vol. 7, pp. 1891–1896, 2023.
- [5] P. Akella and A. D. Ames, "A Barrier-Based Scenario Approach to Verifying Safety-Critical Systems," *IEEE Robotics and Automation Letters*, vol. 7, no. 4, pp. 11 062–11 069, 2022.
- [6] W. Shaw Cortez and D. V. Dimarogonas, "Safe-by-design control for Euler–Lagrange systems," *Automatica*, vol. 146, p. 110620, 2022.
- [7] A. Clark, "Verification and Synthesis of Control Barrier Functions," in *2021 60th IEEE Conference on Decision and Control (CDC)*, Dec. 2021, pp. 6105–6112.
- [8] A. Robey, H. Hu, L. Lindemann, H. Zhang, D. V. Dimarogonas, S. Tu, and N. Matni, "Learning Control Barrier Functions from Expert Demonstrations," in *IEEE Conference on Decision and Control*, 2020, pp. 3717–3724.
- [9] K. Long, C. Qian, J. Cortés, and N. Atanasov, "Learning barrier functions with memory for robust safe navigation," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 4931–4938, 2021.
- [10] A. Wiltz, X. Tan, and D. V. Dimarogonas, "Construction of Control Barrier Functions Using Predictions with Finite Horizon," in *IEEE Conference on Decision and Control*, 2023, pp. 2743–2749.
- [11] P. Roque, W. Shaw Cortez, L. Lindemann, and D. V. Dimarogonas, "Corridor MPC: Towards Optimal and Safe Trajectory Tracking," in *American Control Conference*, 2022, pp. 2025–2032.
- [12] K. P. Wabersich and M. N. Zeilinger, "Predictive Control Barrier Functions: Enhanced Safety Mechanisms for Learning-Based Control," *IEEE Transactions on Automatic Control*, vol. 68, no. 5, pp. 2638–2651, 2023.
- [13] J. Zeng, Z. Li, and K. Sreenath, "Enhancing Feasibility and Safety of Nonlinear Model Predictive Control with Discrete-Time Control Barrier Functions," in *IEEE Conference on Decision and Control*, 2021, pp. 6137–6144.
- [14] L. Brunke, S. Zhou, and A. P. Schoellig, "Robust Predictive Output-Feedback Safety Filter for Uncertain Nonlinear Control Systems," in *IEEE Conference on Decision and Control*, 2022, pp. 3051–3058.
- [15] F. P. Bejarano, L. Brunke, and A. P. Schoellig, "Multi-Step Model Predictive Safety Filters: Reducing Chattering by Increasing the Prediction Horizon," in *IEEE Conference on Decision and Control*, 2023, pp. 4723–4730.
- [16] A. Leeman, J. Köhler, S. Bennani, and M. Zeilinger, "Predictive safety filter using system level synthesis," in *Annual Learning for Dynamics and Control Conference*. PMLR, 2023, pp. 1180–1192.
- [17] J. F. Fisac, A. K. Akametalu, M. N. Zeilinger, S. Kaynama, J. Gillula, and C. J. Tomlin, "A General Safety Framework for Learning-Based Control in Uncertain Robotic Systems," *IEEE Transactions on Automatic Control*, vol. 64, no. 7, pp. 2737–2752, 2019.
- [18] S. Li and O. Bastani, "Robust Model Predictive Shielding for Safe Reinforcement Learning with Stochastic Dynamics," in *IEEE International Conference on Robotics and Automation*, 2020, pp. 7166–7172.
- [19] K. P. Wabersich, L. Hewing, A. Carron, and M. N. Zeilinger, "Probabilistic Model Predictive Safety Certification for Learning-Based Control," *IEEE Transactions on Automatic Control*, vol. 67, no. 1, pp. 176–188, 2022.
- [20] H.-N. Nguyen, "Improved Prediction Dynamics for Robust MPC," *IEEE Transactions on Automatic Control*, vol. 68, no. 9, pp. 5445–5460, 2023.
- [21] J. Berberich, J. Köhler, M. A. Müller, and F. Allgöwer, "Linear Tracking MPC for Nonlinear Systems—Part I: The Model-Based Case," *IEEE Transactions on Automatic Control*, vol. 67, no. 9, pp. 4390–4405, 2022.
- [22] A. Dhar and S. Bhasin, "Indirect Adaptive MPC for Discrete-Time LTI Systems With Parametric Uncertainties," *IEEE Transactions on Automatic Control*, vol. 66, no. 11, pp. 5498–5505, 2021.
- [23] S. Brüggemann and R. R. Bitmead, "Forward-looking persistent excitation in model predictive control," *Automatica*, vol. 136, p. 110033, 2022.
- [24] A. Agrawal and K. Sreenath, "Discrete Control Barrier Functions for Safety-Critical Control of Discrete Systems with Application to Bipedal Robot Navigation," in *Proceedings of Robotics: Science and Systems*, 2017, pp. 1–10.
- [25] Y. Xiong, D.-H. Zhai, M. Tavakoli, and Y. Xia, "Discrete-Time Control Barrier Function: High-Order Case and Adaptive Case," *IEEE Transactions on Cybernetics*, pp. 1–9, 2022.
- [26] R. Takano, H. Oyama, and M. Yamakita, "Application of Robust Control Barrier Function with Stochastic Disturbance Model for Discrete Time Systems," *IFAC-PapersOnLine*, vol. 51, no. 31, pp. 46–51, 2018.
- [27] R. K. Cosner, P. Culbertson, A. J. Taylor, and A. D. Ames, "Robust Safety under Stochastic Uncertainty with Discrete-Time Control Barrier Functions," no. arXiv:2302.07469, 2023.
- [28] M. Ahmadi, A. Singletary, J. W. Burdick, and A. D. Ames, "Safe Policy Synthesis in Multi-Agent POMDPs via Discrete-Time Barrier Functions," in *IEEE Conference on Decision and Control*, 2019, pp. 4797–4803.
- [29] J. Zeng, B. Zhang, and K. Sreenath, "Safety-Critical Model Predictive Control with Discrete-Time Control Barrier Function," in *American Control Conference*. New York: Ieee, 2021, pp. 3882–3889.
- [30] S. Liu, J. Zeng, K. Sreenath, and C. A. Belta, "Iterative Convex Optimization for Model Predictive Control with Discrete-Time High-Order Control Barrier Functions," in *American Control Conference*, 2023, pp. 3368–3375.
- [31] A. Thirugnanam, J. Zeng, and K. Sreenath, "Safety-Critical Control and Planning for Obstacle Avoidance between Polytopes with Control Barrier Functions," in *International Conference on Robotics and Automation*, 2022, pp. 286–292.
- [32] M. H. Cohen, C. Belta, and R. Tron, "Robust Control Barrier Functions for Nonlinear Control Systems with Uncertainty: A Duality-based Approach," in *IEEE Conference on Decision and Control*, 2022, pp. 174–179.
- [33] X. Tan and D. V. Dimarogonas, "Compatibility checking of multiple control barrier functions for input constrained systems," in *IEEE Conference on Decision and Control*, 2022, pp. 939–944.
- [34] J. Drgoňa, J. Arroyo, I. Cupeiro Figueroa, D. Blum, K. Arendt, D. Kim, E. P. Ollé, J. Oravec, M. Wetter, D. L. Vrabie, and L. Helsen, "All you need to know about model predictive control for buildings," *Annual Reviews in Control*, vol. 50, pp. 190–232, 2020.
- [35] H. K. Khalil, *Nonlinear Systems*. Upper Saddle River, N.J. : Prentice Hall, c2002., 2002.
- [36] P. Pauli, A. Koch, J. Berberich, P. Kohler, and F. Allgöwer, "Training Robust Neural Networks Using Lipschitz Bounds," *IEEE Control Systems Letters*, vol. 6, pp. 121–126, 2022.
- [37] M. Fazlyab, A. Robey, H. Hassani, M. Morari, and G. J. Pappas, "Efficient and accurate estimation of lipschitz constants for deep neural networks," in *Proceedings of the International Conference on Neural Information Processing Systems*. Red Hook, NY, USA: Curran Associates Inc., 2019, no. 1025, pp. 11 427–11 438.
- [38] F. Blanchini and S. Miani, *Set-Theoretic Methods in Control*, ser. Systems & Control : Foundations & Applications. Birkhäuser, 2015.
- [39] J. Drgoňa, A. Tuor, and D. Vrabie, "Learning Constrained Parametric Differentiable Predictive Control Policies With Guarantees," *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, pp. 1–12, 2024.
- [40] J. Drgoňa, K. Kiš, A. Tuor, D. Vrabie, and M. Klaučo, "Differentiable predictive control: Deep learning alternative to explicit model predictive control for unknown nonlinear systems," *Journal of Process Control*, vol. 116, pp. 80–92, Aug. 2022.
- [41] J. Drgono, A. Tuor, J. Koch, M. Shapiro, and D. Vrabie, "NeuroMANCER: Neural Modules with Adaptive Nonlinear Constraints and Efficient Regularizations," 2023. [Online]. Available: <https://github.com/pnnl/neuromancer>
- [42] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "CasADi – A software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, 2019.
- [43] J. Drgoňa, M. Kvasnica, M. Klaučo, and M. Fikar, "Explicit stochastic mpc approach to building temperature control," in *IEEE Conference on Decision and Control*, 2013, pp. 6440–6445.