Stable near-optimal control of nonlinear switched discrete-time systems: an optimistic planning-based approach

Mathieu Granzotto, Romain Postoyan, Lucian Buşoniu, Dragan Nešić, Jamal Daafouz

Stable near-optimal control of nonlinear switched discrete-time systems: an optimistic planning-based approach

Mathieu Granzotto, Romain Postoyan, Lucian Buşoniu, Dragan Nešić, and Jamal Daafouz

***Abstract*—Originating in the artificial intelligence literature, optimistic planning (OP) is an algorithm that generates near-optimal control inputs for generic nonlinear discrete-time systems whose input set is finite. This technique is therefore relevant for the near-optimal control of nonlinear switched systems for which the switching signal is the control, and no continuous input is present. However, OP exhibits several limitations, which prevent its desired application in a standard control engineering context, as it requires for instance that the stage cost takes values in $[0, 1]$, an unnatural prerequisite, and that the cost function be discounted. In this paper, we modify OP to overcome these limitations, and we call the new algorithm OP$_{\text{min}}$. We then analyze OP$_{\text{min}}$ under general stabilizability and detectability assumptions on the system and the stage cost. New near-optimality and performance guarantees for OP$_{\text{min}}$ are derived, which have major advantages compared to those originally given for OP. We also prove that a system whose inputs are generated by OP$_{\text{min}}$ in a receding-horizon fashion exhibits stability properties. As a result, OP$_{\text{min}}$ provides a new tool for the near-optimal, stable control of nonlinear switched discrete-time systems for generic cost functions.**

I. Introduction

Optimistic planning (OP) is an algorithm that computes near-optimal control inputs for generic nonlinear discrete-time systems and infinite-horizon discounted costs, provided the set of inputs is finite, see [16,23]. Given the current state, OP intelligently develops the tree of possible future states, which are enumerable, since the input set is finite. By prioritizing branches with better costs, which are optimistic approximations of the infinite-horizon cost, OP efficiently exploits the available computational power. It then returns an optimal sequence of inputs for a *finite*-horizon discounted cost, where the horizon depends on the given computational budget and the initial state. Guarantees on the mismatch between the obtained value function and the original infinite-horizon optimal cost are provided in [16] and are of the form $\frac{\gamma^{d(x)}}{1-\gamma}$, where $\gamma \in (0, 1)$ is the discount factor and $d(x)$ is the state-dependent horizon, which is related to the computation budget used by the algorithm. Hence, in general, for good near-optimality bounds, the discount factor γ has to be taken small.

OP is a priori well-suited for nonlinear switched discrete-time systems for which the control input is the switching signal [3], and no continuous input is present. This is appealing as the optimal control of switched systems remains an open problem, especially when dealing with nonlinear dynamics. Indeed, while the (near-)optimal control of switched linear discrete-time systems is addressed in, e.g., [1,7,30,36,38], the case of nonlinear switched systems is still unraveling and concentrates on continuous-time systems, see, e.g., [33,35,39]. Even so, in the mentioned works and references therein, algorithms are often presented for a particular class of systems, no explicit near-optimality bounds are given, tacitly assuming that the optimal solution is obtained, and the stability of the induced closed-loop is eluded, while stability is often essential in control applications. There is therefore a need for tools for the (near-)optimal stable control of general nonlinear switched systems: we propose a solution based on OP.

Unfortunately, it appears that we cannot apply OP "off-the-shelf" for standard control engineering problems as OP exhibits significant limitations. First, the stage cost has to take bounded values, e.g., in $[0, 1]$, which is not natural in control as this excludes quadratic stage costs for instance, and to constrain the stage cost to take values in $[0, 1]$ via a nonlinear transformation would change the sequence of optimal inputs. Second, the cost is discounted using factor $\gamma \in (0, 1)$, which has to be chosen small for good near-optimality guarantees as explained above, while γ needs to be close to 1 for the closed-loop system to be stable [9,12,25]. As a first contribution, we therefore modify OP to overcome these limitations, that is: the stage cost does not have to take values in a given bounded set and the cost is not discounted. We call this new algorithm OP$_{\text{min}}$. OP$_{\text{min}}$ returns a sequence of inputs, which minimizes a finite-horizon cost more efficiently than a brute-force approach in general, as explained in the paper. In addition, OP$_{\text{min}}$ is designed for minimizing costs instead of maximizing rewards as OP [16,23], which although apparently easy is in fact non-trivial when dealing with undiscounted and unbounded stage costs.

Mathieu Granzotto, Romain Postoyan and Jamal Daafouz are with the Université de Lorraine, CNRS, CRAN, F-54000 Nancy, France (e-mails: {name.surname}@univ-lorraine.fr).

Lucian Buşoniu is with the Department of Automation, Technical University of Cluj-Napoca, Memorandumului 28, 400114 Cluj-Napoca, Romania (e-mail: lucian.busoniu@aut.utcluj.ro).

Dragan Nešić is with the Department of Electrical and Electronic Engineering, University of Melbourne, Parkville, VIC 3010, Australia (e-mail: dnesic@unimelb.edu.au). His work was supported by the Australian Research Council under the Discovery Project DP170104099.

This work was supported by the France-Australia collaboration project IRP-ARS CNRS.

Moreover, our goal is also to ensure that a system controlled by $\mathrm{OP}_{\min}$ in a receding-horizon fashion exhibits stability guarantees, an aspect which is not addressed by OP works [3,16,23]. To do so and as a second contribution, we introduce a novel element to the algorithm, the so-called *stopping criterion*. Originally, at each call OP develops its tree until exhaustion of a computational budget or up to a given depth. In contrast, $\mathrm{OP}_{\min}$ develops a tree at each call until a stabilizing property is found among the calculated inputs. By incorporating this requirement in the algorithm itself via the stopping criterion, the stability of the induced closed-loop system is guaranteed under general stabilizability and detectability conditions, as discussed in more details below. Furthermore, the computational effort is adapted to the current plant state. This is important, as a fixed computational effort might be unfeasible or ill-suited in applications. Other recent works have considered introducing similar stopping criteria for different optimal control algorithms to ensure stability guarantees, see [24] in the context of nonlinear model predictive control and interior point solvers.

The stopping criterion allows to have a direct control on the near-optimality guarantees, that is, how the computed finite-horizon cost function compares to the original infinite-horizon optimal cost. We thus obtain a novel bound on the mismatch between the two costs, with the next desirable features: (i) it does not explode for $\gamma = 1$, contrary to the bound in [16]; (ii) it decreases as the state is close to a given attractor, while the bound $\frac{\gamma^{d(x)}}{1-\gamma}$ in [16] is a constant for a constant horizon. The latter implies that for some states $\mathrm{OP}_{\min}$ may stop with short horizons while ensuring good near-optimality properties, thus reducing computational costs. We also analyse how the endured cost functions along the system solutions when $\mathrm{OP}_{\min}$ is applied in a receding-horizon fashion compare to the original infinite-horizon optimal cost. We show that $\mathrm{OP}_{\min}$ does provide similar desirable performance properties when applied in closed-loop.

Concerning stability, we prove that a system whose inputs are generated by $\mathrm{OP}_{\min}$ in a receding-horizon fashion satisfies a semiglobal practical stability property, where the adjustable parameter is a decision vector used to tune the stopping criterion. We use a generic measuring function to define stability as in, e.g., [12,13,25], thus covering point and set stability in a unified manner. By strengthening the assumptions, we derive stronger stability properties, including a global exponential guarantee and we also prove that the stated stability properties are nominally robust [20]. These stability results differ from [12,13,25] where stability of systems whose inputs minimizes (discounted) finite-horizon costs is analysed. Indeed, the horizon of the cost in this paper is state-dependent, and not fixed, because of the way $\mathrm{OP}_{\min}$ operates. As a result, the stability analysis relies on a different Lyapunov function compared to [12], namely we exploit the infinite-horizon optimal cost function, which we believe is an interesting development in its own right. In addition, our analysis exploits the stopping criterion, while the latter is absent in [12,13,25].

We illustrate the results through the scenario where we are given a finite number of controllers, and we aim at optimally selecting one at any given time instant, while ensuring stability. Two examples are provided, for a cubic integrator and a flexible joint robotic arm, respectively.

Other tree-based algorithms have been considered in the literature for switched systems, albeit with different purposes. For instance, in [8], the stability of linear switched systems under arbitrarily switching is investigated, by generating a tree-like structure of possible future state transitions. The work in [21] considers a branch-and-bound approach for the discrete-time optimal control of switched linear systems and quadratic costs. On the other hand, (relaxed) dynamic programming approaches were considered in [29,31]. In particular, [29] approximates the infinite-horizon optimal control problem for linear switched systems, and [31] develops a value iteration approach exploiting homogeneity of the system and stage costs. The main difference between our present paper and these references is that we address nonlinear switched systems and generic costs.

This paper also conveys another important message. It illustrates how an optimal control algorithm from a different research field, namely artificial intelligence, can be adapted and tailored to solve an important control problem, here the near-optimal control of nonlinear switched discrete-time systems. It also demonstrates how control requirements, like stabilizability, detectability and stability, can be exploited to improve the original algorithm in terms of near-optimality guarantees and computational budget.

Compared to the preliminary version of this work in [11], the main novelty of this paper is the stopping criterion, while we were using a fixed computational budget in [11] as in [16]. This change reframes every theoretical result of the paper and provides major advantages in terms of computational cost as illustrated in examples (see Section V). In addition, (i) the algorithmic complexity of $\mathrm{OP}_{\min}$ is now investigated, (ii) a new case study on the optimal selection of feedback laws is presented; (iii) additional stability and performance results are provided.

The rest of the paper is organized as follows. Section II formally states the problem. $\mathrm{OP}_{\min}$ is presented in Section III, where its algorithmic complexity is analyzed. In Section IV, we analyse the near-optimality and stability properties of $\mathrm{OP}_{\min}$. In Section V, we apply $\mathrm{OP}_{\min}$ for the on-line optimal selection of feedback laws. The proofs are provided in Section VII, and some conclusions are drawn in Section VI. A contraction property of the finite optimal sequence is stated in the appendix.

Notation. Let $\mathbb{R} := (-\infty, \infty)$, $\mathbb{R}_{\geq 0} := [0, \infty)$, $\mathbb{Z}_{\geq 0} := \{0, 1, 2, \ldots\}$ and $\mathbb{Z}_{>0} := \{1, 2, \ldots\}$. We use (x, y) to denote $[x^\top, y^\top]^\top$, where $(x, y) \in \mathbb{R}^n \times \mathbb{R}^m$ and $n, m \in \mathbb{Z}_{>0}$. A function $\chi : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is of class $\mathcal{K}$ if it is continuous, zero at zero and strictly increasing, and it is of class $\mathcal{K}_\infty$ if it is of class $\mathcal{K}$ and unbounded. A continuous function $\beta : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is of class $\mathcal{KL}$ when $\beta(\cdot, t)$ is of class $\mathcal{K}$ for any $t \geq 0$ and $\beta(s, \cdot)$ is decreasing to 0 for any $s \geq 0$. The notation $\mathbb{I}$ stands for the identity map from $\mathbb{R}_{\geq 0}$ to $\mathbb{R}_{\geq 0}$. For any sequence $\boldsymbol{u} = [u_0, u_1, \ldots]$ of length $d \in \mathbb{Z}_{\geq 0} \cup \{\infty\}$ where $u_i \in \mathbb{R}^m$, $i \in \{0, \ldots, d\}$, and any $k \in \{0, \ldots, d\}$, we use $\boldsymbol{u}|_k$ to denote the first k elements of $\boldsymbol{u}$, i.e. $\boldsymbol{u}|_k = [u_0, \ldots, u_{k-1}]$ and $\boldsymbol{u}|_0 = \varnothing$ by convention. Let $f : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, we use $f^{(k)}$ for the composition of function f with itself k times, where $k \in \mathbb{Z}_{\geq 0}$, and $f^{(0)} = \mathbb{I}$. The Euclidean norm of a vector $x \in \mathbb{R}^n$ is denoted by $|x|$. The distance of a vector $x \in \mathbb{R}^n$ to non-empty set $\mathcal{A}$ is defined as $|x|_{\mathcal{A}} := \inf\{|z - x| : z \in \mathcal{A}\}$. The transpose of a matrix A is denoted by $A^\top$. We denote $\lambda_{\min}(P)$ and $\lambda_{\max}(P)$, the smallest and the largest eigenvalues of a symmetric real matrix P. Given a discrete-time dynamical system $x(k+1) = g(x(k), u(k))$ where $x(k) \in \mathbb{R}^n$ is the state variable at step $k \in \mathbb{Z}_{\geq 0}$ and $n \in \mathbb{Z}_{>0}$, we use the compact notation $x^+ = g(x, u)$ instead.

II. Problem Statement

Consider the system

$$x^+ = f_u(x), \tag{1}$$

with state $x \in \mathbb{R}^n$, control input $u \in \mathcal{U}$ where $\mathcal{U} := \{1, \ldots, M\}$ is a finite set of admissible inputs with $M \geq 2$, and $f_u : \mathbb{R}^n \to \mathbb{R}^n$ for every $u \in \mathcal{U}$. We use $\phi(k, x, \boldsymbol{u}|_k)$ to denote the solution to system (1) at time $k \in \mathbb{Z}_{\geq 0}$ with initial condition x and inputs sequence $\boldsymbol{u}|_k = [u_0, u_1, \ldots, u_{k-1}]$, with the convention $\phi(0, x, \cdot) = \phi(0, x, \varnothing) = x$.

We consider the infinite-horizon cost

$$J_\infty(x, \boldsymbol{u}) := \sum_{k=0}^{\infty} \ell_{u_k}(\phi(k, x, \boldsymbol{u}|_k)), \tag{2}$$

where $x \in \mathbb{R}^n$ is the initial state, $\boldsymbol{u}$ is an infinite sequence of admissible inputs, $\ell_u : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ is the stage cost given input $u \in \mathcal{U}$. Finding an infinite sequence of inputs which minimizes (2) given $x \in \mathbb{R}^n$ is very difficult in general, as the particular case of linear switched systems with quadratic stage cost already shows [22,38]. We therefore aim at generating sequences of inputs that *nearly* minimize (2) instead, in a sense made precise in the following. For this purpose, we revise optimistic planning (OP) as originally developed in [16]. We call this new algorithm OP_{min}. Furthermore, we aim at ensuring stability properties for the induced closed-loop system. We ensure the robustness of this stability property under additional regularity properties of f_u and ℓ_u, which will be made later, in Section IV-C.

The forthcoming analysis revolves around the general stabilizability and detectability assumptions on system (1) and stage cost ℓ stated next, as in, e.g., [12,13,25].

Standing Assumption 1 (SA1): There exist $\overline{\alpha}_V, \alpha_W \in \mathcal{K}_\infty$, continuous functions $W, \sigma : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$, and $\overline{\alpha}_W : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ continuous, non-decreasing and zero at zero, such that the following conditions hold.

(i) For any $x \in \mathbb{R}^n$, there exists an infinite sequence of admissible inputs $\boldsymbol{u}^*_\infty(x)$, called *optimal input sequence*, which minimizes (2), i.e. $V_\infty(x) := J_\infty(x, \boldsymbol{u}^*_\infty(x))$, and V_∞ is such that

$$V_\infty(x) \leq \overline{\alpha}_V(\sigma(x)). \tag{3}$$

(ii) For any $x \in \mathbb{R}^n$, $u \in \mathcal{U}$,

$$W(x) \leq \overline{\alpha}_W(\sigma(x)) \tag{4}$$

$$W(f_u(x)) - W(x) \leq -\alpha_W(\sigma(x)) + \ell_u(x). \tag{5}$$

□

Function $\sigma : \mathbb{R}^n \to \mathbb{R}_{\geq 0}$ in SA1 is a "measuring" function that we use to define stability, which depends on the problem. For instance, by defining $\sigma = |\cdot|, \sigma = |\cdot|^2$ or $\sigma : x \mapsto x^\top Q x$ with $Q = Q^\top > 0$, one would be studying the stability of the origin, and by taking $\sigma = |\cdot|_{\mathcal{A}}$, one would study stability of non-empty set $\mathcal{A} \subset \mathbb{R}^n$. General conditions to ensure the first part of item (i) can be found in [19]. The second part of item (i) is related to the stabilizability of system (1) with respect to stage cost ℓ_u in relation to σ. Indeed, it is shown in [13, Lemma 1] that a sufficient condition for (3) to hold is that the stage cost ℓ_u is uniformly globally exponentially controllable to zero with respect to σ for system (1), see [13, Definition 2]. On the other hand, item (ii) of SA1 is a detectability property of the stage cost ℓ_u with respect to σ, and is thus not related to item (i) of SA1. For example, when $\ell_u(x) \geq \sigma(x)$, one verifies item (ii) of SA1 with $W = 0$ and $\alpha_W = \mathbb{I}$. For more information on SA1 and "measuring" function σ, see, e.g., [12,13,25]. Note that we neither require ℓ_u to take values in $[0, 1]$ contrary to [16], nor that it is positive definite.

We are ready to present the algorithm.

III. OP_{min}

A. Main idea

The algorithm evaluates *finite*-horizon costs given any initial state $x \in \mathbb{R}^n$

$$J_d(x, \boldsymbol{u}_d) := \sum_{k=0}^{d} \ell_{u_k}(\phi(k, x, \boldsymbol{u}_d|_k)), \tag{6}$$

where $d \in \mathbb{Z}_{>0}$ is a horizon, and $\boldsymbol{u}_d = [u_0, u_1, ..., u_d] \in \mathcal{U}^{d+1}$. OP_{min} searches for optimal input sequences which minimizes exactly cost (6), given state x and a state-dependent finite horizon $d(x) \in \mathbb{Z}_{>0}$, that is

$$V_{d(x)}(x) := \min_{\boldsymbol{u}_{d(x)}} J_{d(x)}(x, \boldsymbol{u}_{d(x)}). \tag{7}$$

We denote by $\boldsymbol{u}^*_{d(x)}(x)$ a corresponding optimal input sequence of length $d(x)$, which may be non-unique. Hence, $V_{d(x)}(x) = J_{d(x)}(x, \boldsymbol{u}^*_{d(x)}(x))$. The horizon $d(x)$ in (7) is selected by the algorithm itself. In particular, OP_{min} iteratively increases the horizon d in (6) up to a horizon $d(x)$, which depends on the initial state x. Horizon $d(x)$ corresponds to the first d-horizon optimal cost that verifies the next criterion, i.e. $d(x)$ is such that

$$\sigma(\phi(d(x), x, \boldsymbol{u}^*_{d(x)}(x)|_{d(x)})) \leq c_{\text{stop}}(\varepsilon, x), \tag{8}$$

where

- $\phi(d(x), x, \boldsymbol{u}^*_{d(x)}(x)|_{d(x)})$ is the state, with some abuse of terminology[1], reached by applying the finite-horizon sequence $\boldsymbol{u}^*_{d(x)}(x)$;
- $c_{\text{stop}}(\varepsilon, x) \geq 0$ is a stopping function, which we design and which may depend on state vector x and a vector of tuneable parameters $\varepsilon \in \mathbb{R}^{n_\varepsilon}$, with $n_\varepsilon > 0$. The design of c_{stop} is explained in the sequel.

As we will show, by controlling the "size" of the last state $\phi(d(x), x, \boldsymbol{u}^*_{d(x)}(x)|_{d(x)})$ through function c_{stop} and parameter

[1] Strictly speaking $\phi(d(x), x, \boldsymbol{u}^*_{d(x)}(x)|_{d(x)})$ is the value of the solution to (1) initialized at $x \in \mathbb{R}^n$ at step $d(x)$ with inputs $\boldsymbol{u}^*_{d(x)}(x)$.

ε, we control directly, for each $x \in \mathbb{R}^n$, the mismatch between $V_{d(x)}(x)$ and the optimal value function associated to cost (2) at x, i.e. $V_\infty(x)$ defined in item (i) of SA1. For $d(x)$ to be finite for any $x \in \mathbb{R}^n$, we also rely on the following assumption on c_{stop}, which is made without loss of generality as we design c_{stop}.

Standing Assumption 2 (SA2): For any $\varepsilon \in \mathbb{R}^{n_\varepsilon}$ and any $x \in \mathbb{R}^n$ with $\sigma(x) > 0$, $c_{\text{stop}}(\varepsilon, x) > 0$. $\square$

SA2 formalizes which stopping functions c_{stop} guarantee that OP_{min} terminates. Possible candidate functions are, e.g., $c_{\text{stop}}(\varepsilon, x) = \alpha(\sigma(x))$ for some $\alpha \in \mathcal{K}$ in which case there is no parameter ε, $c_{\text{stop}}(\varepsilon, x) = |\varepsilon|\alpha(\sigma(x))$, $c_{\text{stop}}(\varepsilon, x) = |\varepsilon|$, with $\varepsilon \in \mathbb{R}\backslash\{0\}$, or combinations like $c_{\text{stop}}(\varepsilon, x) = \max\{|\varepsilon_1|\sigma(x), |\varepsilon_2|\}$ for $\varepsilon = (\varepsilon_1, \varepsilon_2) \in (\mathbb{R}\backslash\{0\})^2$ and $x \in \mathbb{R}^n$. We stress that c_{stop} is not required to be positive definite, i.e., we accept $c_{\text{stop}}(\varepsilon, x) > 0$ for $\sigma(x) = 0$. By shaping the terminating function c_{stop}, we can tighten (or relax) near-optimality properties as shown in the sequel.

Remark 1: Model predictive control often similarly solves a finite-horizon problem with terminal constraints, see, e.g., [32] for results where the horizon is varying as in (7). Here, not only the horizon is state-dependent, but the terminal set constraint itself is also state-dependent. $\square$

Altogether, given any $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$, the cost function that OP_{min} explicitly calculates is denoted by

$$V_\varepsilon(x) := J_{d(x)}(x, \boldsymbol{u}^*_\varepsilon(x)). \tag{9}$$

We use the notation V_ε in (9), instead of $V_{d(x)}$, to emphasize that the returned cost function is parameterized by ε. Likewise, we denote by $\boldsymbol{u}^*_\varepsilon(x) := \boldsymbol{u}^*_{d(x)}(x)$ a sequence of inputs that has cost $V_\varepsilon(x)$ and verifies (8). Problem (9) is implementable when $d(x)$ is finite, as the input set $\mathcal{U}$ is finite. In this case, a brute-force approach can solve it by developing all possible sequences. However, this is computationally intensive, in particular when $d(x)$ is large, as the computational cost grows exponentially with the horizon. OP_{min} instead intelligently explores the possible sequences to solve (9) with potentially larger horizons with the same computation budget compared to a brute-force approach in general [16], as the computational cost for a given horizon grows with smaller exponential base, see Section III-C.

The next statement ensures that, given any $x \in \mathbb{R}^n$, $d(x)$ as defined in (8) is finite.

Proposition 1: For all $x \in \mathbb{R}^n$ and any $\varepsilon \in \mathbb{R}^{n_\varepsilon}$, $d(x)$ in (8) is finite and $V_\varepsilon(x)$ in (9) is thus well-defined. $\square$

B. Algorithm description

OP_{min} explores the possible choices of inputs *optimistically* until the stopping condition in (8) is verified, and is inspired by [16]. The computational resources utilized for this purpose are denoted as a budget $B \in \mathbb{Z}_{>0}$, which corresponds to $B+1$ leaf expansions, and which adapts to the state x. We denote by $\mathcal{T}$ the exploration tree from initial state $x \in \mathbb{R}^n$, constructed from admissible input sequences and their respective cost (6). A leaf is a node of $\mathcal{T}$ with no children, and the set of all leaves of $\mathcal{T}$ is denoted $\mathcal{L}(\mathcal{T})$. At iteration $i \in \mathbb{Z}_{\geq 0}$, a leaf $L_i \in \mathcal{L}(\mathcal{T})$ is fully expanded. That is, for every $u \in \mathcal{U}$, we add a child to L_i labeled by the resulting state $f_u(L_i)$, which are new leaves of $\mathcal{T}$; after this, L_i is no longer a leaf, but becomes an inner node. We denote with a slight abuse of notation $\boldsymbol{u}(L_i)$ the input sequence from the root x to the state of leaf L_i. We also denote by $J(L_i) := J_{d(i)}(x, \boldsymbol{u}(L_i))$ the cost (6) given by the sequence that takes x to the state of leaf L_i, with $d(i) := \text{depth}(L_i) - 1$, where depth$(\cdot)$ is the number of edges (or inputs) from the root to L_i. The algorithm expands the leaf with minimal associated cost $J(L)$ among all non-expanded leaves $L \in \mathcal{T}$, and we denote such leaf by L_i. The algorithm terminates when an optimal sequence candidate is found with $\sigma(L_i) \leq c_{\text{stop}}(\varepsilon, x)$, that is, $\boldsymbol{u}(L_i)$ verifies $\sigma(\phi(d, x, \boldsymbol{u}^*_d(x))) \leq c_{\text{stop}}(\varepsilon, x)$, see (8). This sequence exists according to Proposition 1. The algorithm is formalized in Algorithm 1.

Algorithm 1 Algorithm for OP_{min}

Input: $c_{\text{stop}}(\varepsilon, \cdot)$, state x
Output: depth explored $d(x)$, sequence $\boldsymbol{u}^*_\varepsilon(x)$, cost $V_\varepsilon(x)$
Initialisation:
1: $d, i \leftarrow -1$
2: tree $\mathcal{T} \leftarrow \{x, \varnothing, 0\}$ {the empty sequence and cost 0}
Optimistic exploration:
3: **while** true **do**
4: $i = i + 1$
5: Find optimistic leaf $L_i \in \arg\min_{L \in \mathcal{L}(\mathcal{T})} J(L)$
Add to $\mathcal{T}$ the children of L_i:
6: for each child c of L_i, $\mathcal{T} \leftarrow \mathcal{T} \cup \{c, \boldsymbol{u}(c), J(c)\}$
7: **if** $d < \text{depth}(L_i) - 1$ **then** {Sequence $\boldsymbol{u}^*_{d+1}(x)$ found}
8: $S \leftarrow L_i$
9: $d \leftarrow \text{depth}(L_i) - 1$
10: **if** $\sigma(L_i) \leq c_{\text{stop}}(\varepsilon, x)$ **then** break {Leaf S selected}
11: **end if**
12: **end while**
13: $B \leftarrow i$
14: **return** $d(x) \leftarrow d$, $\boldsymbol{u}^*_\varepsilon(x) \leftarrow \boldsymbol{u}(S)$ and $V_\varepsilon(x) \leftarrow J(S)$

Steps in lines 5-6 of Algorithm 1 are the optimistic exploration. Any sequence of inputs from descendants (children, children of children and so on) of a node N will have costs J greater than N, as $\ell_u(x) \geq 0$ for any $x \in \mathbb{R}^n$ and $u \in \mathcal{U}$. The optimistic choice then guarantees that $J(L_{i+1}) \geq J(L_i)$ for any iteration $i \in \mathbb{Z}_{\geq 0}$. This implies that the first leaf to be expanded at a depth $d' + 1$ will be a suitable candidate for $V_{d(x)}(x)$, which is in turn tested for the terminating constraint (8). The output cost is $V_\varepsilon(x)$, corresponding to the first finite-horizon input sequence that verifies (8), calculated with a varying budget B, which depends on x. Note that the expansion of the tree is independent from the "leaf selection" step, and is fully determined by the optimistic selection of leaves. The terminating condition in line 10 is guaranteed to be eventually verified when Proposition 1 holds, as formalized in the next proposition.

Proposition 2: Let $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$, then Algorithm 1 terminates with $\boldsymbol{u}(S) = \boldsymbol{u}^*_\varepsilon(x)$ and $J(S) = V_\varepsilon(x)$. $\square$

Remark 2: Compared to the version of OP_{min} presented in our preliminary work [11], the computational budget is not fixed and adapts to $d(x)$. This has the potential advantage of using less computations, see the example in Section V-A. $\square$

In the following subsection, we study the benefits of optimistic exploration compared to a brute-force approach.

C. Algorithmic complexity

OP_{min} solves (9) by generating an exploration tree $\mathcal{T}$. As Algorithm 1 in general only expands certain leaves and not others, we save computational power compared to a brute-force approach. This is shown in the original OP [16], and we extend that analysis here for OP_{min}. The computational cost of OP_{min} is related to the set of leaves it may expand, which, given $x \in \mathbb{R}^n$, is quantified by

$$\mathcal{T}^*(x) := \{\boldsymbol{u}_d : d \in \mathbb{Z}_{\geq 0}, V_\infty(x) \geq J_d(x, \boldsymbol{u}_d)\}. \tag{10}$$

We call $\mathcal{T}^*(x)$ the near-optimal tree at $x \in \mathbb{R}^n$, which is composed of all input sequences that have a cost smaller than optimal cost $V_\infty(x)$. Note that due optimistic exploration, OP_{min} expands a leaf from exploration tree $\mathcal{T}$ with smallest cost: no leaf with cost larger than $V_\infty(x)$ may be expanded. Hence, OP_{min} only considers sequences that belong to $\mathcal{T}^*(x)$. We have the next result.

Proposition 3: Let $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$, then Algorithm 1 only expands leaves with sequences in $\mathcal{T}^*(x)$. □

We have then that $\mathcal{T}^*(x) \subset \mathcal{T}$. We characterize the algorithmic complexity of Algorithm 1 by the *branching factor* of tree $\mathcal{T}^*(x)$, defined as follows.

Definition 1: For any $x \in \mathbb{R}^n$, the branching factor of tree $\mathcal{T}^*(x)$ is the smallest value $\flat(x) \in (1, M]$ for which there exists a constant $C(x) > 0$ so that $|\mathcal{T}_d^*(x)| \leq C(x) \cdot (\flat(x))^d$, for all $d \in \mathbb{Z}_{\geq 0}$, where $|\mathcal{T}_d^*(x)|$ denotes all nodes of $\mathcal{T}^*(x)$ at depth d. □

The branching factor $\flat(x)$ takes values between 1 and M for any $x \in \mathbb{R}^n$, where M is the number of inputs. For lower computational costs, it is desirable to have $\flat(x)$ close to 1. The branching factor $\flat(x)$ depends on the problem and the state x. For example, if one were to consider M identical controllers, OP_{min} would be forced to explore all branches, independently of system and stage cost, as no controller is better than the others. In this case, $\flat(x) = M$ for any $x \in \mathbb{R}^n$. However, the branching factor may be much smaller in applications, see Section V. It is hard to estimate the branching factor a priori, but instead it can be examined a posteriori, as done in the examples provided in Section V. Overall, in order to find a finite optimal sequence for horizon $d(x)$, OP_{min} requires computational budget $B \leq C(x)(\flat(x))^{d(x)}$, which is exponential (and the price to pay for a general algorithm), but with lower cost than brute-force search when $\flat(x) < M$.

On the other hand, the choice of stopping criterion has an important impact on the computational cost of OP_{min}. Indeed, horizon $d(x)$ is typically small when c_{stop} is large, hence requiring less exploration than $d(x)$ large (and c_{stop} small). However, in that case, we may not have good optimality or stability guarantees, as it will be seen in the following section. One then has to consider the appropriate balance between computational budget and system performance. Particular applications may have suitable heuristics for the choice of stopping criterion or refined node expansion strategies. In general, this question must be studied on a case by case basis, and is thus out-of-scope of this paper as we present a general theory and do not concentrate on specific examples.

IV. Main results

In this section, we first analyze the near-optimality properties of OP_{min}. We then provide conditions under which system (1), whose inputs are generated in a receding-horizon fashion by OP_{min}, exhibits stability properties. Robustness of this stability property is ensured afterwards under mild regularity properties. Finally, we analyse the cost along solutions to system (1) controlled by OP_{min}, thereby providing performance guarantees of the closed-loop system.

A. Relationship between V_ε and V_∞

Algorithm 1 is able to calculate $V_\varepsilon(x)$ exactly for any given $x \in \mathbb{R}^n$. However it is not obvious how $V_\varepsilon(x)$ relates to $V_\infty(x)$, which is the original optimal value function we aim for. Since ℓ_u is not constrained to take values in a given compact set, and we do not consider discounted costs, the tools used in [16] to analyze this relationship are no longer applicable. We overcome this issue by exploiting SA1 and the stopping criterion in the next theorem.

Theorem 1: For any $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$,

$$V_\varepsilon(x) \leq V_\infty(x) \leq V_\varepsilon(x) + v_\varepsilon(x), \tag{11}$$

where $v_\varepsilon(x) := \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, x))$ with $\overline{\alpha}_V$ from SA1. □

The lower-bound in (11) trivially holds from the optimality of $V_\varepsilon(x) = V_{d(x)}(x)$ as $d(x) < \infty$, and the fact that $\ell_u(x) \geq 0$ for any $x \in \mathbb{R}^n$ and $u \in \mathcal{U}$. The upper-bound, on the other hand, implies that the infinite-horizon cost is at most $v_\varepsilon(x)$ away from the finite-horizon $V_\varepsilon(x)$. The error term $v_\varepsilon(x)$ is small when so is $c_{\text{stop}}(\varepsilon, x)$ as $\overline{\alpha}_V \in \mathcal{K}_\infty$, which again we can tune. We can therefore make $V_\varepsilon(x)$ as close as desired to $V_\infty(x)$ by adjusting c_{stop}; the price to pay will be more computations.

Remark 3: Compared to the term given in [16], which we recall is $\frac{\gamma^{d(x)}}{1-\gamma}$ for a discount factor $\gamma \in (0, 1)$, v_ε in (11) is finite in the absence of a discount factor. Moreover, we can directly tune $v_\varepsilon(x)$ via $c_{\text{stop}}(\varepsilon, x)$, which is not the case in [16]. By exploiting stabilizability and detectability properties in SA1, we have obtained an error bound that forfeits the assumption $\ell_u \in [0, 1]$, accepts the undiscounted case $\gamma = 1$, depends on the selected stopping condition c_{stop}, and is not necessarily uniform in x. □

B. Stability

We now consider the scenario where system (1) is controlled in a receding-horizon fashion by OP_{min} as defined by Algorithm 1. That is, at each time instant $k \in \mathbb{Z}_{\geq 0}$, the first element of the optimal sequence $\boldsymbol{u}_\varepsilon^*(x_k)$ is calculated by OP_{min}, and then applied to system (1). This leads to the closed-loop system

$$x^+ \in f_{\mathcal{U}_\varepsilon^*(x)}(x) =: F_\varepsilon^*(x), \tag{12}$$

where $f_{\mathcal{U}_\varepsilon^*(x)}(x)$ is the set $\{f_u(x) : u \in \mathcal{U}_\varepsilon^*(x)\}$ and $\mathcal{U}_\varepsilon^*(x) := \big\{u_0 : \exists u_1, \ldots, u_{d(x)} \in \mathcal{U} \text{ such that } V_\varepsilon(x) = J_{d(x)}(x, [u_0, \ldots, u_{d(x)}])\big\}$ is the set of the first input of $d(x)$-horizon optimal input sequences at x, with $d(x)$ as defined in (8). We denote by $\phi(k, x)$ a solution to (12) at time $k \in \mathbb{Z}_{\geq 0}$ with initial condition $x \in \mathbb{R}^n$, with some abuse of notation.

We assume next that c_{stop} can be made small as desirable by taking $|\varepsilon|$ sufficiently small. As we are free to design c_{stop} as wanted, this is without loss of generality.

Assumption 1: There exists $\theta : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, with $\theta(\cdot, s) \in \mathcal{K}$ and $\theta(s, \cdot)$ non-decreasing for any $s > 0$, such that $c_{\text{stop}}(\varepsilon, x) \leq \theta(|\varepsilon|, \sigma(x))$ for any $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$. $\square$

Example of functions c_{stop} which satisfy Assumption 1 are $c_{\text{stop}}(\varepsilon, x) = |\varepsilon|\sigma(x)$, $c_{\text{stop}}(\varepsilon, x) = \max\{|\varepsilon_1|\alpha(\sigma(x)), |\varepsilon_2|\}$ for $\varepsilon = (\varepsilon_1, \varepsilon_2)$, $\alpha \in \mathcal{K}$ and $x \in \mathbb{R}^n$ to give a few. The next theorem provides Lyapunov properties for system (12) that we use to derive the main stability result afterwards.

Theorem 2: Let $Y := V_\infty + W$, the following holds.

(i) For any $x \in \mathbb{R}^n$,

$$\underline{\alpha}_Y(\sigma(x)) \leq Y(x) \leq \overline{\alpha}_Y(\sigma(x)), \tag{13}$$

where $\underline{\alpha}_Y := \alpha_W, \overline{\alpha}_Y := \overline{\alpha}_V + \overline{\alpha}_W$, with $\alpha_W, \overline{\alpha}_V, \overline{\alpha}_W$ from SA1.

(ii) For any $x \in \mathbb{R}^n$, $\varepsilon \in \mathbb{R}^{n_\varepsilon}$, $v \in F_\varepsilon^*(x)$,

$$Y(v) - Y(x) \leq -\alpha_Y(\sigma(x)) + \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, x)) \tag{14}$$

where $\alpha_Y = \alpha_W$, with α_W and $\overline{\alpha}_V$ from SA1, and c_{stop} comes from (8). $\square$

Item (i) states that Y is positive definite and radially unbounded with respect to the set $\{x : \sigma(x) = 0\}$. Item (ii) of Theorem 2 shows that Y strictly decreases along the solutions to (12) up to a perturbative term $\overline{\alpha}_V(c_{\text{stop}}(\varepsilon, x))$, which can be made as small as desired by selecting $|\varepsilon|$ close to 0 as $\overline{\alpha}_V(c_{\text{stop}}(\varepsilon, x)) \leq \overline{\alpha}_V(\theta(|\varepsilon|, \sigma(x)))$, per Assumption 1.

Remark 4: Similar Lyapunov constructions are employed in [12,13]. The difference here is that the horizon in cost (9) is not fixed as in [12,13] and depends on the state. We circumvent this problem in Theorem 2 by using the *infinite-horizon* optimal value function V_∞ in the definition of the Lyapunov function Y (and not the finite-horizon optimal value function as in [12,13]), which we believe is an interesting result on its own. $\square$

The next theorem provides stability guarantees for system (12).

Theorem 3: Consider system (12) and suppose c_{stop} verifies Assumption 1. There exists $\beta \in \mathcal{KL}$ such that, for any $\delta, \Delta > 0$, there exists $\varepsilon^* > 0$ such that for any $x \in \{z \in \mathbb{R}^n : \sigma(z) \leq \Delta\}$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$ with $|\varepsilon| < \varepsilon^*$, any solution $\phi(\cdot, x)$ to system (12) satisfies, for all $k \in \mathbb{Z}_{\geq 0}$

$$\sigma(\phi(k, x)) \leq \max\{\beta(\sigma(x), k), \delta\}. \tag{15}$$

$\square$

Theorem 3 provides a semiglobal practical stability property for the set $\{z : \sigma(z) = 0\}$. This implies that solutions to (12), with initial state x such that $\sigma(x) \leq \Delta$, where Δ is any given (arbitrarily large) strictly positive constant, will converge to the set $\{z : \sigma(z) \leq \delta\}$, where δ is any given (arbitrarily small) strictly positive constant, by taking ε^* sufficiently close to 0, thereby making c_{stop} sufficiently small. An explicit formula for ε^* is given in the proof of Theorem 3 in Section VII, which is nevertheless subject to some conservatism. The result should rather be appreciated qualitatively, in the sense that (15) holds for small enough ε^*.

By strengthening SA1, we can provide stronger properties under a particular class of stopping criterion, namely $c_{\text{stop}}(\varepsilon, x) \leq |\varepsilon|\sigma(x)$ for any $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$. The next result ensures a semiglobal asymptotic stability property.

Corollary 1: Suppose the following holds.

(i) There exist $L, \bar{a}_V, a_W > 0$, such that SA1 holds with $\overline{\alpha}_V(s) \leq \bar{a}_V s$, $\overline{\alpha}_W(s) \leq \bar{a}_W s$ and $\alpha_W(s) \geq a_W s$ for any $s \in [0, L]$.

(ii) For any $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$, $c_{\text{stop}}(\varepsilon, x) \leq |\varepsilon|\sigma(x)$.

Let $\varepsilon^* > 0$ and $\Delta > L$ be such that

$$\varepsilon^* < \min\left\{1, \frac{a_W}{\bar{a}_V}, \frac{\overline{\alpha}_V^{-1}(\frac{1}{2}\alpha_W(L))}{\Delta}\right\}. \tag{16}$$

Then, there exists $\beta \in \mathcal{KL}$, such that, for any $x \in \{z \in \mathbb{R}^n : \sigma(z) \leq \Delta\}$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$ such that $|\varepsilon| < \varepsilon^*$, any solution $\phi(\cdot, x)$ to system (12) satisfies

$$\sigma(\phi(k, x)) \leq \beta(\sigma(x), k) \tag{17}$$

for all $k \in \mathbb{Z}_{\geq 0}$. $\square$

The stability property in (17) corresponds to (15) with $\delta = 0$, thus ensuring a semiglobal asymptotic stability property. Inequality (16) can always be verified by taking ε^* small, since the right-hand side is strictly positive. When the sublinear properties in item (i) of Corollary 1 are valid for $L = \infty$, we have the next stronger result.

Corollary 2: Suppose the following holds.

(i) There exist $\bar{a}_V, a_W > 0$, such that SA1 holds with $\overline{\alpha}_V \leq \bar{a}_V \cdot \mathbb{I}$, $\overline{\alpha}_W \leq \bar{a}_W \cdot \mathbb{I}$ and $\alpha_W \geq a_W \cdot \mathbb{I}$.

(ii) For any $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$, $c_{\text{stop}}(\varepsilon, x) \leq |\varepsilon|\sigma(x)$.

Let $\varepsilon^* > 0$ be such that

$$\varepsilon^* < \frac{a_W}{\bar{a}_V}. \tag{18}$$

Then, for any $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$ such that $|\varepsilon| \leq \varepsilon^*$, any solution $\phi(\cdot, x)$ to system (12) satisfies

$$\sigma(\phi(k, x)) \leq \frac{\bar{a}_V + \bar{a}_W}{a_W}\left(1 - \frac{a_W - |\varepsilon|\bar{a}_V}{\bar{a}_V + \bar{a}_W}\right)^k \sigma(x) \tag{19}$$

for all $k \in \mathbb{Z}_{\geq 0}$. $\square$

Corollary 2 ensures a uniform global exponential stability property of set $\{x : \sigma(x) = 0\}$ for system (12). Indeed, in (19), decay rate $1 - \frac{a_W - |\varepsilon|\bar{a}_V}{\bar{a}_V + \bar{a}_W} \in (0, 1)$ as $|\varepsilon| \leq \varepsilon^* < \frac{a_W}{\bar{a}_V}$ in view of (18), hence $\left(1 - \frac{a_W - |\varepsilon|\bar{a}_V}{\bar{a}_V + \bar{a}_W}\right)^k \to 0$ as $k \to \infty$. Furthermore, the estimated decay rate can be tuned via ε from 1 to $1 - \frac{a_W}{\bar{a}_V + \bar{a}_W}$ as $|\varepsilon|$ decreases to zero.

Remark 5: Items (i)-(ii) of Corollary 2 are sufficient conditions for global exponential stability. If only global asymptotic stability is required, the stopping criterion can be selected as $c_{\text{stop}}(\varepsilon, x) \leq \overline{\alpha}_V^{-1}(\frac{1}{2}\alpha_W(\sigma(x)))$ for all $x \in \mathbb{R}^n$, where $\overline{\alpha}_V$ and α_W come from SA1. Furthermore, if only global practical stability properties are required, the stopping criterion can be selected as $c_{\text{stop}}(\varepsilon, x) \leq |\varepsilon|$. $\square$

C. Nominal robustness

To ensure that the stability properties ensured in Section IV-B are robust to so-called ρ-perturbations, as defined in, e.g., [20], we can rely on two conditions according to [20, Theorem 2.8]. First, the set-valued mapping F_ε^* in (12) needs to be such that $F_\varepsilon^*(x)$ is nonempty and compact for any $x \in \mathbb{R}^n$.

This is the case since $\mathcal{U}$ and $d(x)$ are finite. Compactness of $F^*_\varepsilon(x)$ proceeds from the compactness of $\mathcal{U}^*_\varepsilon(x)$, $\mathcal{U}^*_\varepsilon(x)$ being a closed non-empty subset of finite set $\mathcal{U}$, given that f_u is continuous which we assume in the upcoming lemma. Second, the Lyapunov function used to prove stability has to be continuous. In our case, the Lyapunov function constructed in Section IV-B is $Y = V_\infty + W$. Since W is continuous by SA1, we need V_∞ to be continuous. The next proposition ensures this is the case under extra conditions on f_u, ℓ_u and σ. The result follows from [25, Theorem 3] with $\gamma = 1$ and $\mathcal{U} = \{0, \dots, M\}$, and its proof is therefore omitted.

Lemma 1: Suppose the following holds.

(i) f_u and ℓ_u are continuous for all $u \in \mathcal{U}$.

(ii) For every $M \geq 0$, set $\{x : \sigma(x) \leq M\}$ is compact.

Then V_∞ is continuous on $\mathbb{R}^n$. □

Item (ii) of Lemma 1 means that σ is radially unbounded, which is the case when $\sigma(x) \geq \alpha_\sigma(|x|_\mathcal{A})$ for any $x \in \mathbb{R}^n$, for a non-empty compact set $\mathcal{A}$ and $\alpha_\sigma \in \mathcal{K}_\infty$.

D. Performance guarantees

In Section IV-A, we have provided relationships between the finite-horizon cost V_ε and the infinite-horizon cost V_∞. This is an important feature of OP_{min}, but this does not directly provide us with information on the actual value of the cost function (2) *along solutions* to (12). Indeed, we do not implement the whole sequence $\boldsymbol{u}^*_\varepsilon(x)$ given by OP_{min} at x in (12), instead we proceed in a receding-horizon fashion. Therefore, we analyse a different cost called running cost [14] defined as

$$\mathcal{V}^{\text{run}}_\varepsilon(x) := \Bigg\{ \sum_{k=0}^{\infty} \ell_{\mathcal{U}^*_\varepsilon(\phi(k,x))}(\phi(k,x)) : \phi(\cdot, x) \text{ is a solution to (12)} \Bigg\}, \tag{20}$$

where $\ell_{\mathcal{U}^*_\varepsilon(\phi(k,x))}(\phi(k,x))$ is the actual stage cost incurred at time-step k. It has to be noted that $\mathcal{V}^{\text{run}}_\varepsilon(x)$ is a set, since solutions of (12) are not necessarily unique. Each element $V^{\text{run}}_\varepsilon(x) \in \mathcal{V}^{\text{run}}_\varepsilon(x)$ corresponds then to the cost of a solution of (12). Clearly, $V^{\text{run}}_\varepsilon(x)$ is not necessarily finite, as the stage costs may not decrease to 0 in view of Theorem 3. Indeed, only practical convergence is ensured in Theorem 3 in general. We thus first look at the average cost defined as

$$\mathcal{V}^{\text{avg}}_\varepsilon(x) := \Bigg\{ \lim_{N\to\infty} \frac{1}{N} \sum_{k=0}^{N} \ell_{\mathcal{U}^*_\varepsilon(\phi(k,x))}(\phi(k,x)) : \phi(\cdot, x) \text{ is a solution to (12)} \Bigg\}. \tag{21}$$

As before, $\mathcal{V}^{\text{avg}}_\varepsilon(x)$ is a set of possible averages, where $V^{\text{avg}}_\varepsilon(x) \in \mathcal{V}^{\text{run}}_\varepsilon(x)$ is the average of a possible solution of (12). We provide the next guarantee on each element of $\mathcal{V}^{\text{avg}}_\varepsilon(x)$.

Theorem 4: Consider system (12), and suppose Assumption 1 and Theorem 3 hold with tuple $(\varepsilon^*, \delta, \Delta)$. For any $\varepsilon \in \mathbb{R}^{n_\varepsilon}$ such that $|\varepsilon| < \varepsilon^*$, any $x \in \{z \in \mathbb{R}^n : \sigma(z) \leq \Delta\}$, and $V^{\text{avg}}_\varepsilon(x) \in \mathcal{V}^{\text{avg}}_\varepsilon(x)$, it follows that

$$0 \leq V^{\text{avg}}_\varepsilon(x) \leq \overline{\alpha}_V(\theta(|\varepsilon|, \delta)), \tag{22}$$

where $\overline{\alpha}_V$ and θ comes from SA1 and Assumption 1, respectively. □

Theorem 4 shows that, if δ is small, so is the average running cost, and the latter can be made as close to 0 as desired by taking $|\varepsilon|$ small, as $\theta(\cdot, \delta) \in \mathcal{K}$ according to Assumption 1 and $\overline{\alpha}_V \in \mathcal{K}_\infty$. Note that the average cost associated to the *infinite*-horizon cost (2) is zero as $V_\infty(x) < \infty$ according to item (i) of SA1. Hence, the mismatch between the latter and the elements of $\mathcal{V}^{\text{avg}}_\varepsilon(x)$ can be made as small as desired. Furthermore, the upper-bound in (22) is uniform with respect to x. Theorem 3 plays a vital role in Theorem 4, as it guarantees that the state converges to the attractor $\{z \in \mathbb{R}^n : \sigma(z) \leq \delta\}$ from any initial condition $x \in \{z \in \mathbb{R}^n : \sigma(z) \leq \Delta\}$.

The average cost provides information about the performance along solutions to (12) in the long run, typically once these have converged to attractor $\{x \in \mathbb{R}^n : \sigma(x) \leq \delta\}$. To quantify performance in the transient, i.e. before the solution has entered and stays forever in the attractor, we propose to consider what we call the "cost-to-attractor" function defined as

$$\mathcal{V}^{\text{cta}}_\varepsilon(x) := \Bigg\{ \sum_{k=0}^{N(x)} \ell_{\mathcal{U}^*_\varepsilon(\phi(k,x))}(\phi(k,x)) : \phi(\cdot, x) \text{ is a solution to (12)} \Bigg\}, \tag{23}$$

where $N(x) > 0$ is an integer such that for every $n > N(x)$, $\sigma(\phi(n,x)) \leq \delta$. That is, in contrast to $V^{\text{run}}_\varepsilon(x)$ in (20), where the series goes up to infinity, here we truncate the series earlier at $N(x)$, when the state has reached once and for all the attractor $\{z \in \mathbb{R}^n : \sigma(z) \leq \delta\}$. Given the semiglobal practical stability property of Theorem 3, $N(x)$ is well-defined for any $x \in \{z \in \mathbb{R}^n : \sigma(z) \leq \Delta\}$ and any $\Delta > 0$, provided we select ε^* sufficiently small. We give the following property for $\mathcal{V}^{\text{cta}}_\varepsilon(x)$.

Theorem 5: Consider system (12), and suppose Assumption 1 and Theorem 3 hold with tuple $(\varepsilon^*, \delta, \Delta)$. For any $\varepsilon \in \mathbb{R}^{n_\varepsilon}$ such that $|\varepsilon| < \varepsilon^*$, any $x \in \{z \in \mathbb{R}^n : \sigma(z) \in (\delta, \Delta]\}$, and $V^{\text{cta}}_\varepsilon(x) \in \mathcal{V}^{\text{cta}}_\varepsilon(x)$, it follows that

$$0 \leq V^{\text{cta}}_\varepsilon(x) \leq V_\varepsilon(x) + \sum_{k=0}^{N(x)} \overline{\alpha}_V\big(\theta(|\varepsilon|, \max\{\beta(\sigma(x), k)\}, \delta\}\big), \tag{24}$$

where $\overline{\alpha}_V$, β and θ come from SA1, Theorem 3 and Assumption 1, respectively. □

Theorem 5 implies that the *cost-to-attractor* $\{x \in \mathbb{R}^n : \sigma(x) \leq \delta\}$ at x is upper-bounded by V_ε and an error term which can be controlled by ε. In contrast to the average cost in Theorem 4, we observe in Theorem 5 the role of the decay rate of β on the cost-to-attractor: the faster $\beta(x, \cdot)$ decays, the smaller $N(x)$ and the smaller the error term in (24).

On the other hand, when the set $\{x \in \mathbb{R}^n : \sigma(x) = 0\}$ is globally exponentially stable as in Corollary 2, the elements of $\mathcal{V}^{\text{run}}_\varepsilon(x)$ in (20) are finite and satisfy the next property.

Theorem 6: Consider system (12) and suppose the conditions of Corollary 2 hold. For any ε such that $|\varepsilon| < \varepsilon^*$, $x \in \mathbb{R}^n$, and $V^{\text{run}}_\varepsilon(x) \in \mathcal{V}^{\text{run}}_\varepsilon(x)$, it follows that

$$V_\infty(x) \leq V^{\text{run}}_\varepsilon(x) \leq V_\infty(x) + w_\varepsilon \sigma(x), \tag{25}$$

with $w_\varepsilon := \dfrac{\bar{a}_V}{a_W}(\bar{a}_V + \bar{a}_W)^2 \dfrac{|\varepsilon|}{a_W - \bar{a}_V|\varepsilon|}$, where constants come from Corollary 2. □

The inequality $V_\infty(x) \leq V_\varepsilon^{\text{run}}(x)$ of Theorem 6 directly follows from the optimality of V_∞. The inequality $V_\varepsilon^{\text{run}}(x) \leq V_\infty(x) + w_\varepsilon \sigma(x)$ provides a relationship between the running cost $V_\varepsilon^{\text{run}}(x)$ and the infinite-horizon cost at state x, $V_\infty(x)$. The latter inequality in (25) confirms the intuition coming from Theorem 1 that a smaller stopping criterion leads to tighter near-optimality guarantees. That is, when $|\varepsilon| \to 0$, $w_\varepsilon \to 0$ and $V_\varepsilon^{\text{run}}(x) \to V_\infty(x)$ for any $x \in \mathbb{R}^n$, provided that Corollary 2 holds. In contrast with Theorem 1, stability of system (12) plays a role in Theorem 6. Indeed, the term $\frac{1}{a_W - \bar{a}_V|\varepsilon|}$ in the expression of w_ε shows that the running cost is large when $|\varepsilon|$ is close to $\frac{a_W}{\bar{a}_V}$, hence when stability is not guaranteed the running cost might be unbounded.

Remark 6: The running cost for the original OP was considered in [5], and it was found to perform at worst like the finite sequence, i.e. $V_{\gamma,\bar{d}}^{\text{run}}(x) \leq V_{\gamma,\infty}(x) + \frac{\gamma^{\bar{d}}}{1-\gamma}$, where OP calculates at x an optimal input sequence $\boldsymbol{u}^*_{\gamma,\bar{d}}(x)$ for finite-horizon *discounted* cost $J_{\gamma,\bar{d}}(x, \boldsymbol{u}_{\bar{d}}) = \sum_{k=0}^{\bar{d}} \gamma^k \ell_{u_k}(\phi(k, x, \boldsymbol{u}_{\bar{d}}|_k))$ for some horizon $\bar{d} \in \mathbb{Z}_{>0}$ and discount factor $\gamma \in (0,1)$. Compared to the bound derived for OP, the bound in Theorem 6 has similar benefits as Theorem 1, namely: we are not limited to $\ell_u \in [0,1]$, it is finite for undiscounted costs ($\gamma = 1$), and when $\sigma(x)$ is small follows $w_\varepsilon \cdot \sigma(x)$ small. Moreover, the mismatch is smaller for faster decays, i.e. for smaller $1 - \frac{a_W - |\epsilon|\bar{a}_V}{\bar{a}_V + \bar{a}_W}$. □

Remark 7: Inequality (25) can be written as a relationship of the finite-horizon costs in view of Theorem 6. In particular, we have $V_\varepsilon(x) \leq V_\infty(x) \leq V_\varepsilon^{\text{run}}(x) \leq V_\varepsilon(x) + w_\varepsilon \cdot \sigma(x)$, for any $x \in \mathbb{R}^n$. Hence, $V_\varepsilon(x)$, which is returned by the algorithm at the initial time, can be used to upper and lower bound $V_\varepsilon^{\text{run}}(x)$ from the first call of OP$_{\text{min}}$ at initial state x. □

Remark 8: Inequality (25) can be written as a relative relationship of the true optimal cost [14] when $\sigma(x) > 0$, in view of Theorem 2. In particular, we have $\frac{V_\varepsilon^{\text{run}}(x) - V_\infty(x)}{V_\infty(x) + W(x)} \leq \frac{w_\varepsilon}{a_W}$ under the conditions of Corollary 2, which can be made as small as desired by tunning ε. □

V. Applications

A relevant application of the results of Sections III and IV is when we are given a finite number of feedback laws and we would like to optimally switch among them to minimize a cost function given by (2), while ensuring the stability of the closed-loop. We first discuss the general case in Section V-A and illustrate it on a cubic integrator. We then provide results tailored to the uniting control case [26]–[28], which we apply to a flexible robot arm. These examples illustrate the reduced computational effort and the desired near-optimal properties of OP$_{\text{min}}$, compared respectively to our prior work [11] and a to uniting control approach.

A. General case

We have the plant model

$$\begin{aligned} x^+ &= f(x, \kappa) \\ \kappa &= g(u, x), \end{aligned} \tag{26}$$

where $x \in \mathbb{R}^n$ is the state and $\kappa \in \mathbb{R}^{n_\kappa}$, $n_\kappa \in \mathbb{Z}_{>0}$, is the input generated by the feedback law. The latter is given by $\kappa = g(u, x)$, where $u \in \{1, \ldots, M\}$ is the index of the controller and M is the number of feedback laws. In that way, $x^+ = f(x, g(u, x))$, which is the same form as (1). The objective is to select u to minimize cost (2) while ensuring stability. We can directly apply the results of Section IV for this purpose provided SA1 and SA2 are satisfied. An example is provided below, where we compare the computational budget of OP$_{\text{min}}$ with the one utilised by its preliminary version in [11].

We consider the cubic integrator from [13, Example 1], $x_1^+ = x_1 + u$, $x_2^+ = x_2 + u^3$, i.e. where $(x_1, x_2) := x \in \mathbb{R}^2$ and $u \in \mathbb{R}$. It was verified in [13] that an open-loop sequence of inputs drives the system to $x = 0$ in a finite number of steps. This open-loop sequence can be expressed as three feedback laws $g(1, x) = -x_1$, $g(2, x) = x_2^{\frac{1}{3}}$ and $g(3, x) = \left(-\frac{1}{2} + \sqrt{\frac{7}{12}}\right) x_2^{\frac{1}{3}}$, which are successively applied. We propose here to switch between these gains to minimize cost (2), with $\ell_u(x) = |x_1|^3 + |x_2| + |g(u, x)|^3$ for any $x \in \mathbb{R}^2$ and $u \in \{1, 2, 3\}$. Note that we cannot design a local linear quadratic regulator for this system, due to the lack of stabilizability of the linearized model at the origin. We therefore consider the switched system

$$\begin{aligned} x_1^+ &= x_1 + g(u, x) \\ x_2^+ &= x_2 + g(u, x)^3 \end{aligned} \tag{27}$$

with $u \in \{1, 2, 3\}$.

To apply OP$_{\text{min}}$, we need to ensure that the required assumptions hold. The first part of item (i) of SA1 holds for the same reasons as in [13]. By taking $\sigma(x) = |x_1|^3 + |x_2|$ for any $x \in \mathbb{R}^2$, SA1 holds $\alpha_W = \mathbb{I}$, $W = \overline{\alpha}_W = 0$ and $\overline{\alpha}_V = 14\mathbb{I}$, as in [13]. With $c_{\text{stop}}(\varepsilon, x) = |\varepsilon|\sigma(x)$ with $\varepsilon \in \mathbb{R}$, we verify the conditions of Corollary 2 with $a_W = 1$, $\bar{a}_V = 14$ and $\bar{a}_W = 0$, and, by taking any $|\varepsilon| < \epsilon^* = \frac{a_W}{\bar{a}_V} = \frac{1}{14} \approx 0.07$, we ensure the global exponential stability of the origin. Consequently, Theorem 6 also holds. In particular, by taking $\varepsilon = \frac{a_W}{2\bar{a}_V} = \frac{1}{28} \approx 0.035$, we derive that $\sigma(\phi(k, x)) \leq 14 \left(\frac{27}{28}\right)^k \sigma(x)$ and that $V_\infty(x) - V_\varepsilon^{\text{run}}(x) \leq w_\varepsilon \sigma(x)$ holds with $w_\varepsilon = \bar{a}_V^2 = 196$. For such ε and initial state $x = (-1, 1.5)$ we observe in simulations, see Figure 1, that both x_1 and x_2 converge to zero, as ensured by Corollary 2. The bound on ε given above is subject to some conservatism. In fact, OP$_{\text{min}}$ finds the best input sequence observed with ε as high as $\frac{1}{9} \approx 0.111$, and convergence to the origin for ε as high as $\frac{10}{12} \approx 0.833$.

In Table I, we provide the estimates of $V_\varepsilon^{\text{run}}(x)$ for $x = (10, 15)$ for different choices of ε, and the associated computational budget utilized by OP$_{\text{min}}$, see Section III-B for the definition of the budget. In particular, we denote by B_{max} the maximum utilized budget and by B_{avg} the average budget, i.e. the mean of budgets across time steps of the simulation. Moreover, we provide an estimate $\hat{\flat}(x)$ of branching factor $\flat(x)$ defined in Section III-C. Here, we estimate $\hat{\flat}(x)$ for exploration tree $\mathcal{T}$ at x as $\frac{1}{d(x)} \sum_{k=1}^{d(x)} \frac{|\mathcal{T}_k|}{|\mathcal{T}_{k-1}|}$, i.e. the average growth rate of exploration tree $\mathcal{T}$. The calculated running cost becomes smaller when we decrease parameter ε, which is consistent with Theorem 6. We also see how the computational budget B

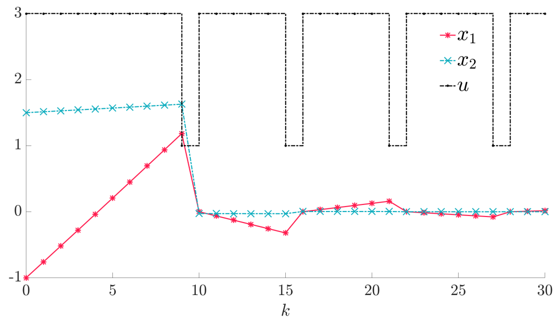

Fig. 1: State and input evolution for $\mathrm{OP}_{\min}$ with $\varepsilon = 0.035$ and $x = (-1, 1.5)$.

adapts to fulfill the stopping criterion. This is a clear advantage over the preliminary version of $\mathrm{OP}_{\min}$ in [11], where budget was fixed in such way to guarantee that the horizon explored was large enough to guarantee stability. A budget of $B \geq \frac{3^{73}-1}{2}$ was required in [11, Corollary 2] to ensure stability in this problem, which is clearly unfeasible. Here, as we do not have to estimate the budget a priori, we do not have to assume the worst case branching factor, $M = 3$, or conservative horizon, $d(x) \geq 72$. Instead, the algorithm now leverages the true branching factor $\flat(x) \approx 1.4$ of the near optimal tree $\mathcal{T}^*(x)$, as indicated by the empirical observation that $\hat{\flat}(x) \to 1.4$ when $\varepsilon \to 0$ in Table I, and adapts the required horizon on the fly, hence the significant reduction in computational budget. On the other hand, $\mathrm{OP}_{\min}$ as in [11] has the advantage that the budget is fixed in advance, which may be suitable for some real-time implementations, given that good estimates of the branching factor and required horizon are available a priori.

ε	$V_\varepsilon^{\text{run}}(x)$	$B_{\max}$	B_{avg}	$\hat{\flat}(x)$
5	∞	2	2	—
0.910	57770	460	159.4	1.977
0.830	22697	460	137.7	1.977
0.590	13757	790	222.2	1.948
0.145	13757	4762	2289.5	1.664
0.110	12609	4762	1986.6	1.664
0.070	12609	4762	1986.6	1.664
0.035	12609	9095	4503.9	1.527
0.001	12609	14695	11265	1.392

TABLE I: Estimated running cost of $\mathrm{OP}_{\min}$, associated computational budget utilized by $\mathrm{OP}_{\min}$, and average branching factor along solutions for various values of ε and initial state $x = (10, 15)$. Symbol "∞" signifies that the cost is infinite (i.e. the state does not converge to zero). Symbol "—" implies that the estimate branching factor is undefined, since $d(x) = 0$.

B. A uniting control approach

A particular instance of Section V-A is when $M = 2$ and one controller is locally optimal for cost (2), and corresponds to $u = 1$, and the other has global stability properties and has index $u = 2$, like in uniting control [26]–[28]. Then $\mathrm{OP}_{\min}$ can be used to optimally switch between these controllers as explained next.

1) Main result : We consider system (26) where $x \in \mathbb{R}^{n_x}$ is the state, $\kappa \in \mathbb{R}^{n_\kappa}$ is the feedback law output, which is parameterized by $u \in \{1, 2\}$, $n_\kappa \in \mathbb{Z}_{>0}$. Vector field f is assumed to be continuously differentiable. We focus on quadratic infinite-horizon costs of the form

$$J_\infty(x, \boldsymbol{u}) = \sum_{k=0}^{\infty} \phi_k^\top Q \phi_k + \kappa_k^\top R \kappa_k, \tag{28}$$

where ϕ_k and κ_k are respectively, with some abuse of notation, the state and input the given feedback law at time-step $k \in \mathbb{Z}_{\geq 0}$ and sequence of controller choices $\boldsymbol{u} := [u_0, u_1, \ldots]$, and $Q \in \mathbb{R}^{n_x \times n_x}$, $R \in \mathbb{R}^{n_\kappa \times n_\kappa}$ are symmetric and positive definite matrices. We aim to minimize (28) over the choices of u. We assume that the global controller satisfies the next properties.

Assumption 2: There exists $P \in \mathbb{R}^{n_x \times n_x}$ symmetric, positive definite matrix and $a, b > 0$ such that, for $U : x \mapsto x^\top P x$, the following hold for any $x \in \mathbb{R}^n$.

(i) $U(f(x, g(2, x))) - U(x) \leq -aU(x)$.
(ii) $|g(2, x)| \leq b|x|$. □

Item (i) of Assumption 2 implies that U is an exponential Lyapunov function for system $x^+ = f(x, g(2, x))$. Item (ii) of Assumption 2 means that the norm of the feedback law is upper-bounded by a linear term in $|x|$, which is the case when $g(2, x)$ is linear for instance. Design techniques to verify Assumption 2 can be found in, e.g., [2,6,10,17], for given classes of systems.

To design the local optimal feedback law, we rely on the next assumption.

Assumption 3: Let $A := \frac{\partial f}{\partial x}|_{(0,0)}$ and $B := \frac{\partial f}{\partial u}|_{(0,0)}$. The pair (A, B) is stabilizable. □

In view of Assumption 3, we can design the optimal controller for the linearized model of (26) at the origin, i.e. $x^+ = Ax + Bu$ and cost $\tilde{\ell}(x, \kappa) = x^\top Q x + \kappa^\top R \kappa$, where Q, R come from (28). This local controller is given by $g(u_{\text{local}}, x) = -Kx$ for any $x \in \mathbb{R}^m$ with $K := (B^\top P_{\text{local}} B + R)^{-1} B^\top P_{\text{local}} A$. Matrix P_{local} is the unique solution of the discrete Riccati equation $P_{\text{local}} = A^\top P_{\text{local}} A - A^\top P_{\text{local}} B (R + B^\top P_{\text{local}} B)^{-1} B^\top P_{\text{local}} A + Q$.

Remark 9: The conditions of Assumption 2 can be relaxed. Possible extensions include the cases where the stage cost in (28) is only quadratic in a neighborhood of the origin, or where the global controller only ensures stability properties with respect to a neighborhood of the origin. These extensions are left for future work not to blur the main message of this paper. □

We choose the stopping criterion as $c_{\text{stop}}(\varepsilon, x) = |\varepsilon| x^\top P x$, with P from Assumption 2, for any $x \in \mathbb{R}^{n_x}$ and some $\varepsilon \in \mathbb{R} \backslash \{0\}$.

The next statement guarantees that the standing assumptions stated in Sections II and III are verified.

Proposition 4: Consider system (26) where Assumption 2 is verified. The following hold.

(i) SA1 is verified with $\sigma(x) = x^\top P x$ for any $x \in \mathbb{R}^n$, $\alpha_W := \frac{\lambda_{\min}(Q)}{\lambda_{\max}(P)}$, $\overline{\alpha}_V := \frac{\nu_1}{1-e^{-\nu_2}} \mathbb{I}$ for $\nu_1 := \frac{\lambda_{\max}(Q) + b\lambda_{\max}(R)}{\lambda_{\min}(P)}$, $\nu_2 := \ln(1-a)^{-1}$ and $W = \overline{\alpha}_W := 0$.
(ii) SA2 and Assumption 1 are verified with $\theta(\varepsilon, \sigma(x)) = \varepsilon \sigma(x)$, for any $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R} \backslash \{0\}$. □

As a result, we can tune ε according to Corollary 2 to endow the corresponding system (26) with global exponential stability and performance guarantees as formalized next.

Proposition 5: Consider system (26) where Assumption 2 is verified. Let $\varepsilon \in (0, \frac{a_W}{\bar{a}_V})$, where $\bar{a}_V = \frac{\nu_1}{1-e^{-\nu_2}}$, $a_W = \frac{\lambda_{\min}(Q)}{\lambda_{\max}(P)}$ with Q from (28), P from Assumption 2 and ν_2, ν_1 from Proposition 4. Let $\sigma(x) = x^\top P x$ for any $x \in \mathbb{R}^{n_x}$. The following hold for any $x \in \mathbb{R}^n$.

(i) For any $k \in \mathbb{Z}_{\geq 0}$, $\sigma(\phi(k,x)) \leq \frac{\bar{a}_V}{a_W}\left(1 + \varepsilon - \frac{a_W}{\bar{a}_V}\right)^k \sigma(x)$.

(ii) $|V_\varepsilon^{\text{run}}(x) - V_\infty(x)| \leq \frac{\bar{a}_V^3}{a_W} \frac{\varepsilon}{a_W - \bar{a}_V \varepsilon}$. □

Proposition 5 is the application of Corollary 2 and Theorem 6, its proof is therefore omitted.

Remark 10: We do not show that we recover the properties of the local optimal controller in a neighborhood of the origin. This is left for future work, however we illustrate next that in simulations this is indeed the case. □

2) Example: We consider the flexible joint robotic arm model from [34, Section 4], discretized by an Euler scheme with $T > 0$, that gives

$$x^+ = x + T(A_p x + B_p \kappa - E_p \phi(x)), \tag{29}$$

where $x \in \mathbb{R}^4$, $\kappa \in \mathbb{R}^1$,

$$A_p = \begin{bmatrix} 0 & 1 & 0 & 0 \\ -48.6 & -1.25 & 48.6 & 0 \\ 0 & 0 & 0 & 1 \\ 19.5 & 0 & -16.17 & 0 \end{bmatrix}, \tag{30}$$

$B_p = (0, 21.6, 0, 0)$, $E_p = (0, 0, 0, 3.33)$. System (29) has a nonlinearity of type Lur'e, due term $\phi(x) = x_3 + \sin(x_3)$. We fix $T = 0.1$. The infinite-horizon cost is given by

$$J_\infty(x, \boldsymbol{u}) := \sum_{k=0}^{\infty} \phi_k^\top Q \phi_k + R\kappa_k^2, \tag{31}$$

where $Q = I_{4\times4}$, with $I_{4\times4}$ is the identity matrix of dimension 4, and $R = 1$. The jump map of system (29) is continuously differentiable at the origin, and the linearized system at the origin is controllable. Hence, Assumption 3 is verified and we design a local controller $g(1,x) = -K_{\text{local}} x$ that optimizes (31) for linear system of (29) around the origin ($u_{\text{local}} = 1$), as in Subsection V-B.1. On the other hand, we design the global controller $g(2,x) = -K_{\text{global}} x$ with $K_{\text{global}} = [3.6\,,\ 0.9\,,\ -1.5\,,\ \ 0.3]$, which verifies LMI conditions found in [10, Theorem 2]. In fact, we slightly modify the LMI[2] in [10, Theorem 2] to guarantee $U(f(x, -K_{\text{global}} x)) - U(x) \leq -a'(x^\top Q x + x^\top K_{\text{global}} R K_{\text{global}} x)$ for some $a' > 0$. This is done to ensure a less conservative estimate than Proposition 5 and conclude that cost (28) for $u = 2$ is given (and upper-bounded) by

$$\sum_{k=0}^{\infty} x_k^\top Q x_k + x_k \top K_{\text{global}} R K_{\text{global}} x \leq \frac{1}{a'} x^\top P x =: \bar{a}_V \sigma(x),$$

where

$$P = 10^4 \cdot \begin{bmatrix} 5.31 & 0.35 & -2.62 & 0.96 \\ 0.35 & 0.03 & -0.20 & 0.05 \\ -2.62 & -0.20 & 2.97 & -0.24 \\ 0.96 & 0.05 & -0.24 & 0.27 \end{bmatrix}. \tag{32}$$

[2]By adding $a(Q + K_{\text{global}}^\top R K_{\text{global}})$, where $a > 0$ is a decision variable, to the block $-P$ of the LMI from [10, Theorem 2]

We calculate $a_W = \dfrac{\lambda_{\min}(Q)}{\lambda_{\max}(P)} = 1.40 \cdot 10^{-5}$ and $\bar{a}_V = \dfrac{1}{a'} = 0.0504$ and SA1 holds with $W = \bar{a}_W = 0$, similarly as in Proposition 5. Take $c_{\text{stop}}(\varepsilon, x) = |\varepsilon| x^\top P x$ with $\varepsilon \in \mathbb{R}\backslash\{0\}$ for any $x \in \mathbb{R}^n$, hence SA2 holds. Moreover Assumption 1 holds with $\theta(|\varepsilon|, \sigma(x)) = c_{\text{stop}}(\varepsilon, x)$. Therefore Corollary 2 and Theorem 6 follow by taking $\varepsilon \in (0, \frac{a_W}{\bar{a}_V})$, where $\frac{a_W}{\bar{a}_V} = 2.77 \cdot 10^{-4}$. Hence, for inputs of system (29) given by

$$\kappa = g(u, x) = \begin{cases} -[0.6, 0.6, -0.7, -0.2]x & \text{when } u = 1 \\ -[3.6, 0.9, -1.5,\ \ 0.3]x & \text{when } u = 2, \end{cases}$$

we can utilize $\text{OP}_{\min}$ to unite both controllers which, given an appropriate choice of ε, will calculate $u_k \in \{1, 2\}$ that preserves the global stability ensured by feedback law $x \mapsto g(2, x)$ while having the option to utilize the locally optimal controller $g(1, x_k)$. We choose $\varepsilon = 10^{-9}$. In Figure 2, we simulate the closed-loop system with initial state $x_0 = (10, -10, -10, -10)$, with input κ either given by $\text{OP}_{\min}$, the local controller or the global controller. We observe that: 1) the local controller indeed only locally stabilizes the origin, as the state fails to converge to the origin; 2) $\text{OP}_{\min}$ prioritizes the local controller, but opts for the global stabilizing one when necessary, see Figure 3.

In Table II, we compare the running cost of $\text{OP}_{\min}$ versus the running costs obtained with the local and global controllers, respectively. We observe that $\text{OP}_{\min}$ outperforms both the local and global controller. We also compare $\text{OP}_{\min}$ to a uniting controller in Table II. The uniting controller is implemented by employing the global controller when $\sigma(x) \geq \xi$ and the local one when $\sigma(x) < \xi$, with $\xi = 10^6$ and no hysteresis. The threshold is a priori large. However, since $\sigma(x) = x^\top P x$ with $\lambda_{\max}(P) \approx 7 \cdot 10^7$, the switch to the local controller only happens when the state is close to 0, as desired. Interestingly, $\text{OP}_{\min}$ also outperforms the uniting controller. We point to Figure 3 for a simple explanation: in spite of the large state, $\text{OP}_{\min}$ first selects $u = 1$, only briefly selecting $u = 2$. $\text{OP}_{\min}$ is free to chose the most (optimally) favourable controller. On the other hand, the uniting controller *has to* employ the global controller for large states, which is sub-optimal. Similar results have been obtained for various values of ξ from 10^7 to 10^2. Moreover, we have observed an average branching factor $\hat{\mathfrak{b}}(x)$ of 1.032 across the 4 given initial states, which is significantly smaller than the worst case $M = 2$ (and close to the best possible value, 1, on which $\mathcal{T}^*$ contains a single path).

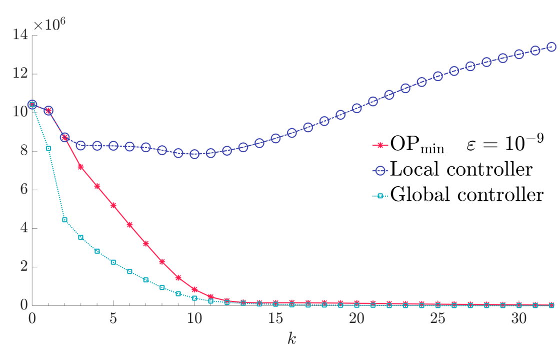

Fig. 2: $\sigma(\phi(\cdot, x))$ for inputs given by $\text{OP}_{\min}$, $g(1, x)$ and $g(2, x)$ for $x_0 = -10 \cdot (-1, 1, 1, 1)$.

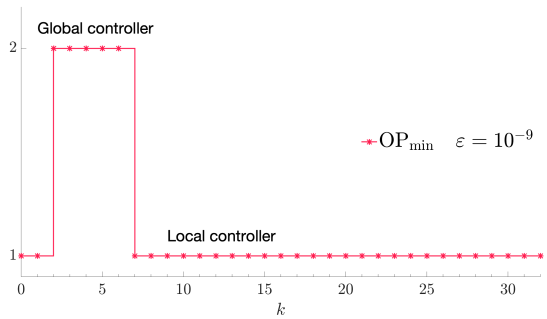

Fig. 3: Input selection of $\mathrm{OP}_{\min}$ for $x_0 = -10 \cdot (-1, 1, 1, 1)$.

		Controller			
		$\mathrm{OP}_{\min}$	Local	Global	Uniting
Initial States	$-10 \cdot (-1, 1, 1, 1)$	**70453**	∞	83365	83295
	$10 \cdot (1, 1, 1, 1)$	**15714**	∞	16046	15998
	$(1, 1, 1, 1)$	**49.75**	**49.75**	203.77	**49.75**
	$-(1, 2, 3, 4)$	**682.48**	**682.48**	791.3	751.79

TABLE II: Estimated running cost for the different controllers and various initial conditions. The symbol "∞" implies that the state is not converging towards the origin, hence the cost explodes. The minimum of each line is given in bold.

VI. Conclusion

We have presented and analysed a planning-based approach for the near-optimal, stable control of general nonlinear switched discrete-time systems where the control input is the switching signal. It will be interesting in future work to study the potential of $\mathrm{OP}_{\min}$ for stochastic problems inspired by [4]. Another path of interest would be to further study the stopping criterion and conservatism in the bound of ε^*. For example, this could be done by studying the particular case of linear switched systems with quadratic stage costs, in which we could then compare with the related literature, e.g., [1,30,37]. A more in-depth analysis of the optimally uniting control problem would also be a relevant research direction, see Remark 10.

VII. Proofs

A. Proof of Proposition 1

Let $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$. In view of Lemma 2 in the appendix, we distinguish two cases. When $\sigma(x) = 0$, it follows that $\sigma(\phi(d, x, \boldsymbol{u}_d^*(x))) \le \alpha_W^{-1} \circ \left(\mathbb{I} - \alpha_Y \circ \overline{\alpha}_Y^{-1}\right)^d \circ \overline{\alpha}_Y(0) = 0$ for any $d \in \mathbb{Z}_{>0}$ according to item (i) of Lemma 2. Hence, $\sigma(\phi(d, x, \boldsymbol{u}_d^*(x))) \le c_{\text{stop}}(\varepsilon, x)$ for any $d \in \mathbb{Z}_{>0}$ since $c_{\text{stop}}(\varepsilon, x) \ge 0$. When $\sigma(x) > 0$, $\sigma(\phi(d, x, \boldsymbol{u}_d^*(x))) \le \alpha_W^{-1} \circ \left(\mathbb{I} - \alpha_Y \circ \overline{\alpha}_Y^{-1}\right)^d \circ \overline{\alpha}_Y(\sigma(x))$ and this upper-bound can be made arbitrarily close to 0 by increasing d, according to item (iii) of Lemma 2. Hence, there exists a finite d sufficiently large such that $\sigma(\phi(d, x, \boldsymbol{u}_d^*(x))) \le \alpha_W^{-1}(\left(\mathbb{I} - \alpha_Y \circ \overline{\alpha}_Y^{-1}\right)^d (\overline{\alpha}_Y(\sigma(x))) \le c_{\text{stop}}(\varepsilon, x)$ as $c_{\text{stop}}(\varepsilon, x) > 0$ for $\sigma(x) > 0$, and we take $\boldsymbol{u}_\varepsilon^*(x) = \boldsymbol{u}_d^*(x)$. Thus, for any $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$ there is a finite $d \in \mathbb{Z}_{>0}$ such that $\sigma(\phi(d, x, \boldsymbol{u}_\varepsilon^*(x))) \le c_{\text{stop}}(\varepsilon, x)$. As a consequence, $d(x)$ in (8) is finite and so is V_ε in (9).

B. Proof of Proposition 2

Let $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$. First, we show that S exactly calculates cost $V_{d'}^* := J_{d'}(x, \boldsymbol{u}_{d'}^*(x))$ for some $d' \in \mathbb{Z}_{>0}$. The optimal property of output S to Algorithm 1 is fully determined in the particular iteration in which it is updated. Hence, let $\mathcal{T}_i$ be the tree to be expanded at iteration $i \in \mathbb{Z}_{\ge 0}$, in which S is updated. We show now that the selected leaf S with cost $J(S)$, where $J(S)$ is the cost associated to leaf S, attains the optimum of horizon $d' := \text{depth}(S) - 1$, that is $J(S) = V_{d'}^*$. Since $V_{d'}^* \le J(S)$ by the optimality of $\boldsymbol{u}_{d'}^*(x)$, it suffices to prove $V_{d'}^* \ge J(S)$. For this purpose, we proceed by contradiction, and we assume that $V_{d'}^* < J(S)$. It follows from the fact that the input set $\mathcal{U}$ is finite that a sequence that attains the optimum $V_{d'}^*$ exists, i.e. there is a node $N \ne S$, decedent of root x and possibly not in $\mathcal{T}_i$, with cost $J(N) = V_{d'}^*$. Since $\ell_u(x) \ge 0$ for any $x \in \mathbb{R}^n$ and $u \in \{1, \dots, M\}$, any ancestor (parents, parents of parents and so on) of N will have a lower cost than $J(N)$. Hence, let L_i' be the ancestor of N such that $L_i' \in \mathcal{L}(\mathcal{T}_i)$, thus $J(L_i') \le J(N)$. Then, we have $J(L_i') \le J(N) = V_{d'}^* < J(S)$, that is $J(L_i') < J(S)$. However, S is the optimistically chosen leaf $S = L_i$, and $J(S) = J(L_i) \le J(L)$ for any leaf $L \in \mathcal{L}(\mathcal{T}_i)$, hence for leaf L_i', it follows that $J(L_i') < J(S) \le J(L_i')$, which is impossible. We have attained a contradiction, therefore $V_{d'}^* < J(S)$ is false, which implies $J(S) \le V_{d'}^*$ and since $V_{d'}^* \le J(S)$, we conclude $V_{d'}^* = J(S)$. $J(S) \le V_{d'}^*$ and since $V_{d'}^* \le J(S)$, we conclude $V_{d'}(x) = J(S)$. Thus, at every update of S, a new optimal sequence is found with increased horizon $d' \leftarrow d' + 1$. By Proposition 1, $d(x)$ is well-defined and there exists a sequence $\boldsymbol{u}_{d(x)}^*(x) =: \boldsymbol{u}_\varepsilon^*(x)$ such that $\sigma(\phi(d(x), x, \boldsymbol{u}_\varepsilon^*(x))) \le c_{\text{stop}}(\varepsilon, x)$. In other words, $\sigma(L_i) \le c_{\text{stop}}(\varepsilon, x)$ is bound to be verified in a finite number of expansion, and $d(x)$ as defined in (8) holds with $d(x) = \text{depth}(L_i) - 1$ holds. Therefore, Algorithm 1 is guaranteed to terminate, with outputs $d(x)$ and $S = \{\phi(d(x), x, \boldsymbol{u}_\varepsilon^*(x)), \boldsymbol{u}_\varepsilon^*(x), V_\varepsilon(x)\}$ fully determined.

C. Proof of Proposition 3

Let $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$. First, note that the root in the exploration tree $\mathcal{T}$ of Algorithm 1 is always expanded, and corresponds to the empty sequence in $\mathcal{T}^*(x)$. Moreover, for any exploration tree $\mathcal{T}$ of Algorithm 1, there exists a leaf $L' \in \mathcal{L}(\mathcal{T})$ such that $J(L') \le V_\infty(x)$. Indeed, a truncated subsequence of $\boldsymbol{u}^*(x)$ will do. At iteration i, due to optimistic exploration, Algorithm 1 selects leaf L_i such that $J(L_i) \le J(L)$ for all $L \in \mathcal{L}(\mathcal{T})$, hence $J(L_i) \le J(L') \le V_\infty(x)$. Thus, by definition of (10), $\boldsymbol{u}(L_i)$ is in $\mathcal{T}^*(x)$ and the proof is complete.

D. Proof of Theorem 1

Let $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$, $d(x) \in \mathbb{Z}_{>0}$ as in (8), optimal sequence $[u_0^*, u_1^*, \dots, u_{d(x)}^*] := \boldsymbol{u}_\varepsilon^*(x)$, cost $V_\varepsilon(x)$ defined in (9) are well-formed by Proposition 1. Since V_ε is a finite-horizon optimal cost, $V_\varepsilon(x) \le V_\infty(x)$. On the other hand, consider the infinite-horizon sequence $\boldsymbol{u} = [u_0^*, u_1^*, \dots u_{d(x)-1}^*, \boldsymbol{u}_\infty^*(\phi(d(x), x, \boldsymbol{u}_\varepsilon^*(x)|_{d(x)}))]$ which exists by item (i) of SA1. It follows from the optimality of $V_\infty(x)$ that $V_\infty(x) \le J_\infty(x, \boldsymbol{u})$, and from the definition of $\boldsymbol{u}$ that $J_\infty(x, \boldsymbol{u}) = V_\varepsilon(x) + V_\infty(\phi(d(x), x, \boldsymbol{u}_\varepsilon^*(x)|_{d(x)}))$, which is finite. By invoking item (i) of SA1, we derive $V_\infty(x) \le V_\varepsilon(x) + \overline{\alpha}_V(\sigma(\phi(d(x), x, \boldsymbol{u}^*_\varepsilon(x)|_{d(x)})))$, and, by the definition of $d(x)$ in (8), $V_\infty(x) \le V_\varepsilon(x) + \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, x))$.

E. Proof of Theorem 2

Let $\varepsilon \in \mathbb{R}^{n_\varepsilon}$, $x \in \mathbb{R}^n$ and $v \in F^*_\varepsilon(x)$, which is well-defined in view of Proposition 1. There exists $[u_0^*, u_1^*, \ldots, u^*_{d(x)}] = \boldsymbol{u}^*_\varepsilon(x)$ such that $v = f_{u_0^*}(x)$ and $\boldsymbol{u}^*_\varepsilon(x)$ is an optimal input sequence for system (1) and cost (6) with horizon $d(x)$, which also verifies (8). Hence $V_\varepsilon(x) = J_{d(x)}(x, \boldsymbol{u}^*_\varepsilon(x))$.

From items (i) and (ii) of SA1, we have $Y(x) = V_\infty(x) + W(x) \le \overline{\alpha}_V(\sigma(x)) + \overline{\alpha}_W(\sigma(x)) =: \overline{\alpha}_Y(\sigma(x))$. On the other hand, we have from item (ii) of SA1 that $\alpha_W(\sigma(x)) \le W(x) + \ell_{u_0^*}(x)$ since $W(f_{u_0^*}(x)) \ge 0$. This implies that $\alpha_W(\sigma(x)) \le W(x) + V_\varepsilon(x) \le W(x) + V_\infty(x) = Y(x)$. Hence item (i) of Theorem 2 holds with $\underline{\alpha}_Y = \alpha_W$.

Consider the sequence $\hat{\boldsymbol{u}} := [u_1^*, u_2^*, \ldots, u^*_{d(x)-1}, \bar{\boldsymbol{u}}]$ where $\bar{\boldsymbol{u}} := \boldsymbol{u}^*_\infty(\phi(d(x), x, \boldsymbol{u}^*_\varepsilon(x)|_{d(x)}))$, $\boldsymbol{u}^*_\varepsilon(x)|_{d(x)} = [u_0^*, \ldots, u^*_{d(x)-1}]$ and ϕ denotes the solution of system (1). The sequence $\hat{\boldsymbol{u}}$ consists of the first $d(x)$ elements of $\boldsymbol{u}^*_\varepsilon(x)$ after u_0^*, followed by an optimal input sequence of infinite length at state $\phi(d(x), x, \boldsymbol{u}^*_\varepsilon(x)|_{d(x)})$, which exists according to item (i) of SA1. Sequence $\bar{\boldsymbol{u}}$ minimizes $J_\infty(\phi(d(x), x, \boldsymbol{u}^*_\varepsilon(x)|_{d(x)}), \bar{\boldsymbol{u}})$ by virtue of item (i) of SA1. From the definition of cost J_d in (6) and $V_\infty(v)$ in view of item (i) of SA1,

$$\begin{aligned} V_\infty(v) &\le J_\infty(v, \hat{\boldsymbol{u}}) \\ &= J_{d(x)-1}(v, \hat{\boldsymbol{u}}|_{d(x)-1}) \\ &\quad + J_\infty(\phi(d(x)-1, v, \hat{\boldsymbol{u}}|_{d(x)-1}), \bar{\boldsymbol{u}}). \end{aligned} \tag{33}$$

From Bellman optimality principle, we have $V_\varepsilon(x) = V_{d(x)}(x) = \ell_{u_0^*}(x) + V_{d(x)-1}(v) = \ell_{u_0^*}(x) + J_{d(x)-1}(v, \hat{\boldsymbol{u}}|_{d(x)-1})$, hence

$$J_{d(x)-1}(v, \hat{\boldsymbol{u}}|_{d(x)-1}) = V_\varepsilon(x) - \ell_{u_0^*}(x). \tag{34}$$

Moreover, by item (i) of SA1,

$$\begin{aligned} &J_\infty(\phi(d(x)-1, v, \hat{\boldsymbol{u}}|_{d(x)-1}), \bar{\boldsymbol{u}}) \\ &\quad \le \overline{\alpha}_V(\sigma(\phi(d(x)-1, v, \hat{\boldsymbol{u}}|_{d(x)-1}))). \end{aligned} \tag{35}$$

Consequently, in view of (33), (34) and (35),

$$\begin{aligned} V_\infty(v) &\le V_\varepsilon(x) - \ell_{u_0^*}(x) \\ &\quad + \overline{\alpha}_V(\sigma(\phi(d(x)-1, v, \hat{\boldsymbol{u}}|_{d(x)-1}))). \end{aligned} \tag{36}$$

Since $\phi(d(x)-1, v, \hat{\boldsymbol{u}}|_{d(x)-1}) = \phi(d(x), x, \boldsymbol{u}^*_\varepsilon(x)|_{d(x)})$ and (8) holds, $\sigma(\phi(d(x)-1, v, \hat{\boldsymbol{u}}|_{d(x)-1})) = \sigma(\phi(d(x), x, \boldsymbol{u}^*_\varepsilon(x)|_{d(x)})) \le c_{\text{stop}}(\varepsilon, x)$. Therefore,

$$V_\infty(v) \le V_\varepsilon(x) - \ell_{u_0^*}(x) + \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, x)). \tag{37}$$

By Theorem 1, $V_\varepsilon(x) \le V_\infty(x)$, thus

$$V_\infty(v) \le V_\infty(x) - \ell_{u_0^*}(x) + \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, x)). \tag{38}$$

By invoking item (ii) of SA1, we derive $V_\infty(v) + W(v) \le V_\infty(x) + W(x) - \alpha_W(\sigma(x)) + \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, x))$, and since $Y = V_\infty + W$, the proof is completed with $\alpha_Y := \alpha_W$.

F. Proof of Theorem 3

Let $\Delta, \delta > 0$. We select $\varepsilon^* > 0$ such that

$$\theta(\varepsilon^*, \underline{\alpha}_Y^{-1}(\widetilde{\Delta})) < \overline{\alpha}_V^{-1}(\frac{1}{2}\widetilde{\alpha}_Y(\widetilde{\delta})), \tag{39}$$

where $\widetilde{\alpha}_Y := \alpha_W \circ \overline{\alpha}_Y^{-1}$, $\widetilde{\Delta} := \overline{\alpha}_Y(\Delta)$, $\widetilde{\delta} := \left(\mathbb{I} - \frac{\widetilde{\alpha}_Y}{2}\right)^{-1} \circ \underline{\alpha}_Y(\delta)$ and θ comes from Assumption 1. Note that $\left(\mathbb{I} - \frac{\widetilde{\alpha}_Y}{2}\right)^{-1}$ is indeed of class $\mathcal{K}_\infty$ as we assume without loss of generality that[3] $\mathbb{I} - \widetilde{\alpha}_Y \in \mathcal{K}_\infty$, hence $\mathbb{I} - \widetilde{\alpha}_Y + \frac{\widetilde{\alpha}_Y}{2} \in \mathcal{K}_\infty$ and so is its inverse. Inequality (41) can always be verified by taking ε^* sufficiently small since $\theta(\cdot, \underline{\alpha}_Y^{-1}(\widetilde{\Delta})) \in \mathcal{K}$, and $\overline{\alpha}_V^{-1}(\frac{1}{2}\widetilde{\alpha}_Y(\widetilde{\delta})) > 0$. It follows from $\theta(\cdot, s) \in \mathcal{K}$ for any $s > 0$ and $\theta(s, \cdot)$ is non-decreasing for any $s \ge 0$, that $\theta(|\varepsilon|, \underline{\alpha}_Y^{-1}(s)) \le \theta(\varepsilon^*, \underline{\alpha}_Y^{-1}(\widetilde{\Delta}))$ for any $s \in [0, \widetilde{\Delta}]$ and $|\varepsilon| < \varepsilon^*$. Furthermore, from Assumption 1 and item (i) of Theorem 2, we derive $c_{\text{stop}}(\varepsilon, x) \le \theta(|\varepsilon|, \underline{\alpha}_Y^{-1}(Y(x)))$. Thus, in view of (39),

$$c_{\text{stop}}(\varepsilon, x) \le \overline{\alpha}_V^{-1}\left(\frac{1}{2}\widetilde{\alpha}_Y(\widetilde{\delta})\right) \tag{40}$$

for any x such that $Y(x) \le \widetilde{\Delta}$. On the other hand, we have $\overline{\alpha}_V^{-1}(\frac{1}{2}\widetilde{\alpha}_Y(\widetilde{\delta})) \le \overline{\alpha}_V^{-1}(\frac{1}{2}\widetilde{\alpha}_Y(s))$ for any $s \in [\widetilde{\delta}, \infty)$. Hence, for any $x \in \mathbb{R}^n$ such that $Y(x) \in [\widetilde{\delta}, \widetilde{\Delta}]$ and $|\varepsilon| < \varepsilon^*$,

$$\overline{\alpha}_V(c_{\text{stop}}(\varepsilon, x)) \le \frac{\widetilde{\alpha}_Y(\widetilde{\delta})}{2} \le \frac{\widetilde{\alpha}_Y(Y(x))}{2}. \tag{41}$$

Let $x \in \mathbb{R}^n$ with $\sigma(x) \le \Delta$ and $v \in F^*_\varepsilon(x)$. In view of (40) and items (i) and (ii) of Theorem 2,

$$Y(v) - Y(x) \le -\widetilde{\alpha}_Y(Y(x)) + \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, x)). \tag{42}$$

Since $\sigma(x) \le \Delta$, $Y(x) \le \overline{\alpha}_Y(\sigma(x)) \le \overline{\alpha}_Y(\Delta) = \widetilde{\Delta}$. Consider $Y(x) \in [0, \widetilde{\delta})$. Since $c_{\text{stop}}(\varepsilon, x) \le \overline{\alpha}_V^{-1}(\frac{1}{2}\widetilde{\alpha}_Y(\widetilde{\delta}))$ holds for $Y(x) \le \widetilde{\Delta}$, it holds here. Furthermore, since $\mathbb{I} - \widetilde{\alpha}_Y \in \mathcal{K}_\infty$ holds without loss of generality, and in view of (42),

$$\begin{aligned} Y(v) &\le Y(x) - \widetilde{\alpha}_Y(Y(x)) + \overline{\alpha}_V(\theta(\varepsilon^*, \sigma(x))) \\ &\le (\mathbb{I} - \widetilde{\alpha}_Y)(\widetilde{\delta}) + \frac{1}{2}\widetilde{\alpha}_Y(\widetilde{\delta}). \end{aligned} \tag{43}$$

Given the definition of $\widetilde{\delta}$,

$$Y(v) \le \left(\mathbb{I} - \frac{\widetilde{\alpha}_Y}{2}\right)(\widetilde{\delta}) = \underline{\alpha}_Y(\delta). \tag{44}$$

When[4] $Y(x) \ge \widetilde{\delta}$, we derive from (41) that $-\widetilde{\alpha}_Y(Y(x)) + \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, x)) \le -\frac{1}{2}\widetilde{\alpha}_Y(Y(x))$. Thus, from (42),

$$Y(v) - Y(x) \le -\tfrac{1}{2}\widetilde{\alpha}_Y(Y(x)). \tag{45}$$

In view of (44) and (45), it follows for any $k \in \mathbb{Z}_{\ge 0}$ that

$$Y(\phi(k+1, x)) \le \max\left\{(\mathbb{I} - \tfrac{1}{2}\widetilde{\alpha}_Y)(Y(x)), \underline{\alpha}_Y(\delta)\right\}, \tag{46}$$

[3] If that is not the case we can always find $\alpha' \in \mathcal{K}_\infty$ such that $\mathbb{I} - \widetilde{\alpha}_Y \le \mathbb{I} - \alpha'$. Indeed, as $\alpha_Y = \alpha_W = \underline{\alpha}_Y \le \overline{\alpha}_Y$ holds from (i), which in turn implies $\alpha_Y \circ \overline{\alpha}_Y^{-1} \le \mathbb{I}$, hence $s - \alpha_Y \circ \overline{\alpha}_Y(s) \ge 0$ for all $s \ge 0$ and equality holds if and only if $s = 0$. Therefore there exists $\alpha' \in \mathcal{K}_\infty$ such that $\mathbb{I} - \alpha' \in \mathcal{K}_\infty$ and $\mathbb{I} - \alpha_Y \circ \overline{\alpha}_Y^{-1} \le \mathbb{I} - \alpha'$, by Lemma B.1 [18]. A similar property is derived in [15].

[4] It might be of interest to assume $\underline{\alpha}_Y(\delta) \le \widetilde{\Delta}$ as to the set $\{x : Y(x) \ge \widetilde{\delta}\}$ be non-empty, but it is not necessary.

where $\phi(k,x)$ is a solution starting at x for system (12). Furthermore, when $Y(x) \le \underline{\alpha}_Y(\delta)$, $Y(v) \le \underline{\alpha}_Y(\delta)$ follows. Indeed, if $Y(x) \in [\widetilde{\delta}, \widetilde{\Delta}]$, $Y(v) \le Y(x) \le \underline{\alpha}_Y(\delta)$ from (45), and if $Y(x) \in [0, \widetilde{\delta})$, we deduce $Y(v) \le \underline{\alpha}_Y(\delta)$ from (44). Hence the set $\{z \in \mathbb{R}^n : Y(z) \le \underline{\alpha}_Y(\delta)\}$ is forward invariant for system (12). By iterating (46), we obtain

$$Y(\phi(k,x)) \le \max\left\{\widetilde{\beta}(Y(x),k), \underline{\alpha}_Y(\delta)\right\}, \tag{47}$$

where $\widetilde{\beta}(s,k) = \left(\mathbb{I} - \frac{1}{2}\widetilde{\alpha}_Y\right)^{(k)}(s)$ for any $s \ge 0$, with $\widetilde{\beta} \in \mathcal{KL}$ as $\lim_{k\to\infty}\left(\mathbb{I} - \frac{1}{2}\widetilde{\alpha}_Y\right)^{(k)}(s) = 0$ for any $s \ge 0$ per the proof of item (ii) of Lemma 2 from the Appendix, since $\left(\mathbb{I} - \frac{1}{2}\widetilde{\alpha}_Y\right)(s) < s$ for $s > 0$ and $\left(\mathbb{I} - \frac{1}{2}\widetilde{\alpha}_Y\right)(0) = 0$. Finally, invoking $\underline{\alpha}_Y(\sigma(x)) \le Y(x) \le \overline{\alpha}_Y(\sigma(x))$, we deduce

$$\sigma(\phi(k,x)) \le \max\left\{\underline{\alpha}_Y^{-1}\left(\widetilde{\beta}(\overline{\alpha}_Y(\sigma(x)),k)\right), \delta\right\}. \tag{48}$$

Thus (15) holds with $\beta(s,k) = \underline{\alpha}_Y^{-1}\left(\widetilde{\beta}(\overline{\alpha}_Y(s),k)\right)$ for any $s \ge 0$ and $k \in \mathbb{Z}_{\ge 0}$.

G. Proof of Corollary 1

Let $\Delta > 0$, $x \in \mathbb{R}^n$ be such that $\sigma(x) \le \Delta$. We select ε^* as in (16) and let $\varepsilon \in \mathbb{R}^{n_\varepsilon}$ such that $|\varepsilon| \le \varepsilon^*$ and $v \in F_\varepsilon^*(x)$. Note that ε^* in (16) is well defined since the right-hand side is strictly positive. We will follow the same arguments as proof of Theorem 3, however applying the sublinear bounds of Corollary 1. From item (ii) of Theorem 2, $Y(v) - Y(x) \le -\alpha_W(\sigma(x)) + \overline{\alpha}_V(c_{\text{stop}}(\varepsilon,x))$. We use the following strategy. First, we show that $Y(v) - Y(x) \le -\frac{\mu}{\bar{a}_V + \bar{a}_W}Y(x)$ holds for some $\mu > 0$ when $\sigma(x) \in [0, L]$ since $\varepsilon^* < \frac{a_W}{\bar{a}_V}$. Then, we show that $Y(v) - Y(x) \le -\frac{1}{2}\widetilde{\alpha}_Y(Y(x))$ holds for $\widetilde{\alpha}_Y = \alpha_W \circ \overline{\alpha}_Y^{-1}$ when $\sigma(x) \in (L, \Delta]$, given $\varepsilon^* < \frac{\overline{\alpha}_V^{-1}(\frac{1}{2}\alpha_W(L))}{\Delta}$. To conclude, we combine the two inequalities and we defer to the proof of Theorem 3.

Let, for the moment, x be such that $\sigma(x) \le L$. From item (i) of Corollary 1, we have that $-\alpha_W(\sigma(x)) \le -a_W\sigma(x)$ holds as $\sigma(x) \le L$, and similarly that $\overline{\alpha}_V(c_{\text{stop}}(\varepsilon,x)) \le \bar{a}_V|\varepsilon|\sigma(x)$, since $|\varepsilon|L < \varepsilon^* L < L$ follows from item (ii) of Corollary 1 and (16). It follows then that $Y(v) - Y(x) \le -\alpha_W(\sigma(x)) + \overline{\alpha}_V(c_{\text{stop}}(\varepsilon,x)) \le (-a_W + \bar{a}_V|\varepsilon|)\sigma(x)$ holds. Since $|\varepsilon| \le \varepsilon^* < \frac{a_W}{\bar{a}_V}$, we derive $-a_W + \bar{a}_V|\varepsilon| < 0$, hence, there exists $\mu > 0$ such that $-a_W + \bar{a}_V|\varepsilon| < -\mu$. We derive $Y(v) - Y(x) \le -\mu\sigma(x)$. On the other hand, we have $Y(x) \le \overline{\alpha}_Y(\sigma(x)) \le (\bar{a}_V + \bar{a}_W)\sigma(x)$ in view of item (i) of Theorem 2, hence, $-(\bar{a}_V + \bar{a}_W)\sigma(x) \le -Y(x)$. Since $\mu > 0$, we derive

$$Y(v) - Y(x) \le -\frac{\mu}{\bar{a}_V + \bar{a}_W}Y(x). \tag{49}$$

When $\sigma(x) \in (L, \Delta]$, it follows that $Y(v) - Y(x) \le -\alpha_W(\sigma(x)) + \overline{\alpha}_V(c_{\text{stop}}(\varepsilon,x)) \le -\alpha_W(\sigma(x)) + \overline{\alpha}_V(\varepsilon^*\Delta)$ from item (ii) of Corollary 1. As $\varepsilon^* < \frac{\overline{\alpha}_V^{-1}(\frac{1}{2}\alpha_W(L))}{\Delta}$, $\overline{\alpha}_V(\varepsilon^*\Delta) < \frac{1}{2}\alpha_W(L)$ holds. Since $\sigma(x) > L$ and $\alpha_W \in \mathcal{K}_\infty$, we have that $\frac{1}{2}\alpha_W(L) < \frac{1}{2}\alpha_W(\sigma(x))$, hence $\overline{\alpha}_V(\varepsilon^*\Delta) < \frac{1}{2}\alpha_W(\sigma(x))$ and $Y(v) - Y(x) < -\alpha_W(\sigma(x)) + \frac{1}{2}\alpha_W(\sigma(x)) = -\frac{1}{2}\alpha_W(\sigma(x))$. Then, in view of item (i) of Theorem 2, we have $Y(x) \le \overline{\alpha}_Y(\sigma(x))$ that implies $\alpha_W \circ \overline{\alpha}_Y^{-1}(Y(x)) \le \alpha_W(\sigma(x))$, and conclude

$$Y(v) - Y(x) \le -\frac{1}{2}\widetilde{\alpha}_Y(Y(x)), \tag{50}$$

where $\widetilde{\alpha}_Y = \alpha_W \circ \overline{\alpha}_Y^{-1}$.

We have found that $Y(v) - Y(x)$ decreases for all $Y(x) \in (0, \overline{\alpha}_Y(\Delta)]$. In particular, by $-\frac{\mu}{\bar{a}_V + \bar{a}_W}Y(x)$ for $\sigma(x) \in [0, L]$ and by $-\frac{1}{2}\widetilde{\alpha}_Y(Y(x))$ elsewhere, that is $Y(v) - Y(x) \le -\min\left\{\frac{\mu}{\bar{a}_V + \bar{a}_W}\mathbb{I}, \frac{1}{2}\widetilde{\alpha}_Y\right\}(Y(x))$. The desired result is then derived by following the final steps of Theorem 3 to construct β.

H. Sketch of proof of Corollary 2

Let $x \in \mathbb{R}^n$. We select ε^* as in (18) and let $\varepsilon \in \mathbb{R}^{n_\varepsilon}$ such that $|\varepsilon| \le \varepsilon^*$ and $v \in F_\varepsilon^*(x)$. In particular, we have shown that $Y(v) - Y(x) \le -\frac{\mu}{\bar{a}_V + \bar{a}_W}Y(x)$ holds for $\sigma(x) \in [0, L)$ in Corollary 1. In this proof, we derive from (18) that $Y(v) - Y(x) \le -\left(\frac{a_W - |\varepsilon|\bar{a}_V}{\bar{a}_V + \bar{a}_W}\right)Y(x)$ holds for any $\sigma(x) \ge 0$. Note that we do not require $\varepsilon^* < 1$ since $\overline{\alpha}_V(|\varepsilon|\sigma(x)) \le \bar{a}_V|\varepsilon|\sigma(x)$ is guaranteed to hold for any $x \in \mathbb{R}^n$ and $\varepsilon \in \mathbb{R}^{n_\varepsilon}$. We now proceed with the same argument as the proof of Corollary 2 in [25]. Let $x \in \mathbb{R}^n$ and denote $\phi(k,x)$ be a corresponding solution to (12) at time $k \in \mathbb{Z}_{\ge 0}$, it holds that $Y(\phi(k,x)) \le \left(1 - \frac{a_W - |\varepsilon|\bar{a}_V}{\bar{a}_V + \bar{a}_W}\right)^k Y(x)$. Since $Y(x) \ge \alpha_W(\sigma(x)) \ge a_W\sigma(x)$ and $Y(x) \le \overline{\alpha}_Y(\sigma(x)) \le (\bar{a}_V + \bar{a}_W)\sigma(x)$ holds from item (i) of Theorem 2 and item (i) of Corollary 2 for any $x \in \mathbb{R}^n$, it follows from $Y(\phi(k,x)) \le \left(1 - \frac{a_W - |\varepsilon|\bar{a}_V}{\bar{a}_V + \bar{a}_W}\right)^k Y(x)$ that $a_W\sigma(\phi(k,x)) \le \left(1 - \frac{a_W - |\varepsilon|\bar{a}_V}{\bar{a}_V + \bar{a}_W}\right)^k (\bar{a}_V + \bar{a}_W)\sigma(x)$ hence $\sigma(\phi(k,x)) \le \frac{\bar{a}_V + \bar{a}_W}{a_W}\sigma(x)\left(1 - \frac{a_W - |\varepsilon|\bar{a}_V}{\bar{a}_V + \bar{a}_W}\right)^k$ and the proof is concluded.

I. Proof of Theorem 4

Let $\Delta, \delta > 0$, $x \in \mathbb{R}^n$ such that $\sigma(x) \le \Delta$. We select ε^* as in Theorem 3. Let $\varepsilon \in \mathbb{R}^{n_\varepsilon}$ such that $|\varepsilon| < \varepsilon^*$, $\phi(k+1,x) \in F_\varepsilon^*(\phi(k,x))$ for any $k \in \mathbb{Z}_{\ge 0}$ where ϕ is a solution to (12) initialized at x. For the sake of convenience, we denote $\ell(x,u) := \ell_u(x)$ for any $u \in \mathcal{U}$. Consider

$$V_\varepsilon^{\text{avg}}(x) = \lim_{N\to\infty}\frac{1}{N}\sum_{k=0}^{N}\ell(\phi(k,x), u_k^r), \tag{51}$$

where $u_k^r \in \mathcal{U}_\varepsilon^*(\phi(k,x))$ such that $\phi(k+1,x) = f_{u_k^r}(\phi(k,x))$. Note that indeed $V_\varepsilon^{\text{avg}}(x) \in \mathcal{V}_\varepsilon^{\text{avg}}(x)$. The lower-bound $0 \le V_\varepsilon^{\text{avg}}(x)$ in (22) follows immediately from $\ell(x,u) \ge 0$ for any $x \in \mathbb{R}^n$ and $u \in \mathcal{U}$. On the other hand, we derive from (37) that, for any $k \in \mathbb{Z}_{\ge 0}$,

$$\begin{aligned} &V_\infty(\phi(k+1,x)) - V_\varepsilon(\phi(k,x)) \\ &\quad\le -\ell(\phi(k,x), u_k^r) + \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, \phi(k,x))). \end{aligned} \tag{52}$$

Hence

$$\begin{aligned} \ell(\phi(k,x), u_k^r) \le\ & V_\varepsilon(\phi(k,x)) - V_\infty(\phi(k+1,x)) \\ &+ \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, \phi(k,x))), \end{aligned} \tag{53}$$

from which we deduce, for any $N \geq 0$,

$$\begin{aligned}
&\sum_{k=0}^{N} \ell(\phi(k,x), u_k^r) \\
&\quad \leq V_\varepsilon(\phi(0,x)) - V_\infty(\phi(1,x)) + \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, \phi(0,x))) \\
&\qquad + V_\varepsilon(\phi(1,x)) - V_\infty(\phi(2,x)) + \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, \phi(1,x))) \\
&\qquad + \dots \\
&\qquad + V_\varepsilon(\phi(N(x),x)) - V_\infty(\phi(N+1,x)) \\
&\qquad + \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, \phi(N(x),x))) \\
&\quad \leq V_\varepsilon(\phi(0,x)) + \sum_{k=0}^{N} \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, \phi(k,x))), \qquad (54)
\end{aligned}$$

since $V_\varepsilon(\phi(k,x)) - V_\infty(\phi(k,x)) \leq 0$ for all $k \in \mathbb{Z}_{\geq 0}$ according to Theorem 1 and $V_\infty(\phi(N+1,x)) \geq 0$. According to Assumption 1, $c_{\text{stop}}(\varepsilon, \phi(k,x)) \leq \theta(|\varepsilon|, \sigma(\phi(k,x)))$, and since Theorem 3 holds, $\sigma(\phi(k,x)) \leq \max\{\beta(k, \sigma(x)), \delta\}$ as $\sigma(x) \leq \Delta$. Hence, by direct substitution in (51),

$$\begin{aligned}
V_\varepsilon^{\text{avg}}(x) \leq \lim_{N\to\infty} \frac{1}{N} \Bigg(& V_\varepsilon(\phi(0,x)) \\
& + \sum_{k=0}^{N} \overline{\alpha}_V(\theta(\varepsilon, \max\{\beta(k, \sigma(x))), \delta\}) \Bigg).
\end{aligned}$$

We break the sum in two parts. Let $H(x) \in \mathbb{Z}$ be such that $\beta(\sigma(x), k) > \delta$ for $k \in \{0, \dots, H(x)\}$ and $\beta(\sigma(x), k) \leq \delta$ for $k \in \{H(x)+1, \dots\}$ when $\sigma(x) > \delta$, otherwise, if $\sigma(x) \leq \delta$, we define it $H(x) = -1$. Integer $H(x)$ exists and is finite since $\beta \in \mathcal{KL}$. It follows that $\beta(\sigma(x), k) \leq \delta$ for $k \in \{H(x)+1, \dots\}$. Hence

$$\begin{aligned}
V_\varepsilon^{\text{avg}}(x) \leq \lim_{N\to\infty} \frac{1}{N} \Bigg(& V_\varepsilon(x) \\
& + \sum_{k=0}^{\min\{H(x),N\}} \overline{\alpha}_V\Big(\theta(\varepsilon, \beta(k, \sigma(x)))\Big) \\
& + \sum_{k=H(x)+1}^{N} \overline{\alpha}_V(\theta(\varepsilon, \delta)) \Bigg), \qquad (55)
\end{aligned}$$

where $\sum_{k=0}^{-1} = 0$ by convention. It follows that $\frac{1}{N}\left(V_\varepsilon(x) + \sum_{k=0}^{\min\{H(x),N\}} \overline{\alpha}_V(\theta(\varepsilon, \beta(k, \sigma(x))))\right) \to 0$ as $N \to \infty$, which implies

$$\begin{aligned}
V_\varepsilon^{\text{avg}}(x) &\leq \lim_{N\to\infty} \frac{1}{N} \sum_{k=H(x)+1}^{N} \overline{\alpha}_V(\theta(|\varepsilon|, \delta)) \\
&\leq \lim_{N\to\infty} \frac{1}{N} \sum_{k=0}^{N} \overline{\alpha}_V(\theta(|\varepsilon|, \delta)). \qquad (56)
\end{aligned}$$

Hence $V_\varepsilon^{\text{avg}}(x) \leq \overline{\alpha}_V(\theta(|\varepsilon|, \delta)) \lim_{N\to\infty} \frac{1}{N} \sum_{k=0}^{N} 1$ and Theorem 4 holds.

J. Proof of Theorem 5

Let $\Delta, \delta > 0$, $x \in \mathbb{R}^n$ such that $\sigma(x) \leq \Delta$. We select ε^* as in Theorem 3. Let $\varepsilon \in \mathbb{R}^{n_\varepsilon}$ such that $|\varepsilon| < \varepsilon^*$, $\phi(k+1,x) \in F_\varepsilon^*(\phi(k,x))$ for any $k \in \mathbb{Z}_{\geq 0}$ where ϕ is a solution to (12) initialized at x, and $N(x)$ is such that for any $n > N(x)$, $\phi(n,x) \leq \delta$ which exists since $\sigma(\phi(k,x)) \leq \max\{\beta(\sigma(x), k), \delta\}$ for any $k \in \mathbb{Z}_{\geq 0}$ and $\beta \in \mathcal{KL}$ according to (15). For the sake of convenience, we denote $\ell(x,u) = \ell_u(x)$ as in the proof of Theorem 4. Consider

$$V_\varepsilon^{\text{cta}}(x) := \sum_{k=0}^{N(x)} \ell(\phi(k,x), u_k^r), \qquad (57)$$

where $u_k^r \in \mathcal{U}_\varepsilon^*(\phi(k,x))$ such that $\phi(k+1,x) = f_{u_k^r}(\phi(k,x))$. From (54) in the proof of Theorem 4, we have that

$$\sum_{k=0}^{N} \ell(\phi(k,x), u_k^r) \leq V_\varepsilon(\phi(0,x)) + \sum_{k=0}^{N} \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, \phi(k,x))), \qquad (58)$$

for any $N \geq 0$. Theorem 5 holds by taking $N = N(x)$, and by invoking Assumption 1 and Theorem 3, that is, $\overline{\alpha}_V(c_{\text{stop}}(\varepsilon, \phi(k,x))) \leq \overline{\alpha}_V(\theta(|\varepsilon|, \max\{\beta(\sigma(x), k), \delta\}))$.

K. Sketch of Proof of Theorem 6

Let $x \in \mathbb{R}^n$, $\varepsilon \in \mathbb{R}^{n_\varepsilon}$ such that $|\varepsilon| < \varepsilon^*$ where ε^* is selected as in Corollary 2, $\phi(k+1,x) \in F_\varepsilon^*(\phi(k,x))$ for any $k \in \mathbb{Z}_{\geq 0}$ where ϕ is a solution to (12) initialized at x. The proof follows by following the steps of the proof of Theorem 4, in particular inequality (54), however summed with $N \to \infty$. That is,

$$V_\varepsilon^{\text{run}}(x) \leq V_\varepsilon(\phi(0,x)) + \sum_{k=0}^{\infty} \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, \phi(k,x))). \qquad (59)$$

All that remains is to compute a bound on $\sum_{k=0}^{\infty} \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, \phi(k,x)))$, which is possible by recalling that $\sigma(\phi(k,x)) \leq \frac{\bar{a}_V + \bar{a}_W}{a_W} \sigma(x) \left(1 - \frac{a_W - |\varepsilon| \bar{a}_V}{\bar{a}_V + \bar{a}_W}\right)^k$ holds from Corollary 2 and $\overline{\alpha}_V(c_{\text{stop}}(\varepsilon, \phi(k,x))) \leq \bar{a}_V |\varepsilon| \sigma(\phi(k,x))$ as the conditions of Corollary 2 are assumed to hold. Specifically, $\sum_{k=0}^{\infty} \overline{\alpha}_V(c_{\text{stop}}(\varepsilon, \phi(k,x))) \leq |\varepsilon| \bar{a}_V \frac{(\bar{a}_V + \bar{a}_W)}{a_W} \sigma(x) \sum_{k=0}^{\infty} \left(1 - \frac{a_W - |\varepsilon| \bar{a}_V}{\bar{a}_V + \bar{a}_W}\right)^k$, which provides (25) as $\sum_{k=0}^{\infty} \left(1 - \frac{a_W - |\varepsilon| \bar{a}_V}{\bar{a}_V + \bar{a}_W}\right)^k = \frac{\bar{a}_V + \bar{a}_W}{a_W - \bar{a}_V |\varepsilon|}$. The lower bound $V_\infty(x) \leq V_\varepsilon^{\text{run}}(x)$ follows from the optimality of $V_\infty(x)$. Since (59) holds for an arbitrary solution of (12), $\phi(k+1,x) = f_{u_k^r}(\phi(k,x))$ for any $k \in \mathbb{Z}_{\geq 0}$, the resulting bound holds for any $V_\varepsilon^{\text{run}}(x) \in \mathcal{V}_\varepsilon^{\text{run}}(x)$.

L. Proof of Proposition 4

Let $x \in \mathbb{R}^n$. From item (i) of Assumption 2, there exists $P \in \mathbb{R}^n \times \mathbb{R}^n$ symmetric, positive definite matrix and $a > 0$ such that, for solution $\phi_{\text{global}}(k,x)$ to system (26) with feedback law $g(2,x)$ initialized at x, the following holds for $k \in \mathbb{Z}_{\geq 0}$.

$$\begin{aligned}
&\phi_{\text{global}}(k+1,x)^\top P \phi_{\text{global}}(k+1,x) - \phi_{\text{global}}(k,x)^\top P \phi_{\text{global}}(k,x) \\
&\qquad \leq -a \phi_{\text{global}}(k,x)^\top P \phi_{\text{global}}(k,x).
\end{aligned}$$

Hence, $\phi_{\text{global}}(k+1,x)^\top P \phi_{\text{global}}(k+1,x) \leq (1-a)\phi_{\text{global}}(k,x)^\top P \phi_{\text{global}}(k,x)$. By iteration and recalling that $\phi_{\text{global}}(0,x) = x$, we derive

$$\sigma(\phi_{\text{global}}(k+1,x)) \leq (1-a)^k \sigma(x), \qquad (60)$$

where $\sigma(x) = x^\top P x$. We show next that $\ell_2(\phi_{\text{global}}(k,x)) \leq \nu_1 \sigma(x) e^{-\nu_2 k}$ for some $\nu_1, \nu_2 > 0$ and $\ell_2(x) := x^\top Q x +$

$g(2,x)^\top Rg(2,x)$. Since Q is positive definite, $x^\top Qx \le \lambda_{\max}(Q)|x|^2$. Furthermore $|x|^2 \le \frac{1}{\lambda_{\min}(P)}x^\top Px$, hence $x^\top Qx \le \frac{\lambda_{\max}(Q)}{\lambda_{\min}(P)}x^\top Px$. Similarly for R and invoking item (ii) of Assumption 2, we obtain that, $g(2,x)^\top Rg(2,x) \le \lambda_{\max}(R)|g(2,x)|^2 \le \lambda_{\max}(R)b|x|^2 \le \frac{b\lambda_{\max}(R)}{\lambda_{\min}(P)}x^\top Px$. It follows from (60) that

$$\ell_2(\phi_{\text{global}}(k,x)) \le \nu_1\sigma(x)e^{-\nu_2 k}, \tag{61}$$

where $\nu_1 = \frac{\lambda_{\max}(Q)+b\lambda_{\max}(R)}{\lambda_{\min}(P)}$ and $\nu_2 = \ln(1-a)^{-1}$. By invoking [13, Lemma 1], we derive that the second part of item (i) of SA1 holds with $\overline{\alpha}_V := \frac{\nu_1}{1-e^{-\nu_2}}\mathbb{I}$. In particular, $V_\infty(x) \le \overline{\alpha}_V(\sigma(x)) = \frac{\nu_1}{1-e^{-\nu_2}}x^\top Px$, which is finite for all $x \in \mathbb{R}^n$, thus [19, Theorem 2] is verified and the first part of item (i) of SA1 holds, hence item (i) of Proposition 4 is verified. On the other hand, since $\lambda_{\min}(Q)|x|^2 \le x^\top Qx$ holds, $\frac{\lambda_{\min}(Q)}{\lambda_{\max}(P)}\sigma(x) = \frac{\lambda_{\min}(Q)}{\lambda_{\max}(P)}x^\top Px \le x^\top Qx \le \ell_u(x)$ for any $u \in \{1,\ldots,M\}$, hence item (ii) of SA1 is verified with $\alpha_W := \frac{\lambda_{\min}(Q)}{\lambda_{\max}(P)}\mathbb{I}$ and $W = \overline{\alpha}_W = 0$. We have proved that item (ii) of Proposition 4 holds. Item (iii) of Proposition 4 follows immediately since $c_{\text{stop}}(\varepsilon,x) = \theta(|\varepsilon,\sigma(x)) := \varepsilon\sigma(x)$ by our choice of c_{stop}, hence Assumption 1 holds. Furthermore, $\varepsilon\sigma(x) > 0$ when $\sigma(x) > 0$ follows from $\varepsilon \in \mathbb{R}_{>0}$, hence SA2 holds.

Appendix

We show that, for any $d \in \mathbb{Z}_{>0}$ and $x \in \mathbb{R}^n$, the following properties hold for any finite-horizon optimal sequence $\boldsymbol{u}_d^*(x)$.

Lemma 2: Let $x \in \mathbb{R}^n$. For any $d \in \mathbb{Z}_{>0}$ and $\boldsymbol{u}_d^*(x)$, the following hold.

(i) $\sigma(\phi(d,x,\boldsymbol{u}_d^*(x)|_d)) \le \alpha_W^{-1} \circ \left(\mathbb{I} - \alpha_W \circ \overline{\alpha}_Y^{-1}\right)^{(d)} \circ \overline{\alpha}_Y(\sigma(x))$, with $\overline{\alpha}_Y = \overline{\alpha}_V + \overline{\alpha}_W$ and $\alpha_W, \overline{\alpha}_V, \overline{\alpha}_W \in \mathcal{K}_\infty$ comes from SA1.

(ii) Function $\mathbb{I} - \alpha_W \circ \overline{\alpha}_Y^{-1}$ contracts to zero, that is, for any $s > 0$, $\left(\mathbb{I} - \alpha_W \circ \overline{\alpha}_Y^{-1}\right)^{(d)}(s) < \left(\mathbb{I} - \alpha_W \circ \overline{\alpha}_Y^{-1}\right)^{(d-1)}(s)$ and for any $s \ge 0$, $\lim_{d\to\infty}\left(\mathbb{I} - \alpha_W \circ \overline{\alpha}_Y^{-1}\right)^{(d)}(s) = 0$.

Proof. Let $x \in \mathbb{R}^n$ and $d \in \mathbb{Z}_{>0}$. A d-horizon optimal sequence $\boldsymbol{u}_d^*(x)$ such that $V_d(x) = J_d(x,\boldsymbol{u}_d^*(x))$ exists for any $d \in \mathbb{Z}_{>0}$ in view of (6) as the input set $\mathcal{U}$ is finite. Let $Y_d := V_d + W$, where W comes from item (ii) of SA1. We have that: (a) $\underline{\alpha}_Y(\sigma(x)) \le Y_d(x) \le \overline{\alpha}_Y(\sigma(x))$ holds with $\underline{\alpha}_Y = \alpha_W$ and $\overline{\alpha}_Y = \overline{\alpha}_V + \overline{\alpha}_W$ as [12, Theorem 1] or [13, Theorem 1] applies; (b) with $Y_0(\phi(d,x,\boldsymbol{u}_d^*(x))) \le \left(\mathbb{I} - \alpha_Y \circ \overline{\alpha}_Y^{-1}\right)^{(d)}(Y_d(x))$ with $\alpha_Y := \alpha_W$ according to [12, (32)], and $\mathbb{I} - \alpha_Y \circ \overline{\alpha}_Y^{-1} \in \mathcal{K}_\infty$ (see footnote 3 in page 11). By applying (b) in (a), we obtain $\sigma(\phi(d,x,\boldsymbol{u}_d^*(x))) \le \underline{\alpha}_Y^{-1} \circ \left(\mathbb{I} - \alpha_Y \circ \overline{\alpha}_Y^{-1}\right)^d(\overline{\alpha}_Y(\sigma(x)))$, and item (i) of Lemma 2 holds. In view of $\left(\mathbb{I} - \alpha_Y \circ \overline{\alpha}_Y^{-1}\right), \alpha_Y \circ \overline{\alpha}_Y^{-1} \in \mathcal{K}_\infty$, it follows that $s - \alpha_Y \circ \overline{\alpha}_Y^{-1}(s) < s$ for any $s > 0$. Hence, by composing $\left(\mathbb{I} - \alpha_Y \circ \overline{\alpha}_Y^{-1}\right)$ to both sides $d-1$ times, we conclude $\left(\mathbb{I} - \alpha_Y \circ \overline{\alpha}_Y^{-1}\right)^{(d)}(s) < \left(\mathbb{I} - \alpha_Y \circ \overline{\alpha}_Y^{-1}\right)^{(d-1)}(s)$ when $s > 0$, i.e. strictly decreasing in d for $s > 0$, and 0 when $s = 0$. Hence, $\mathbb{I} - \alpha_Y \circ \overline{\alpha}_Y^{-1}$ is contractive to zero and item (ii) of Lemma 2 holds. ∎

References

[1] D. Antunes and W. P. M. H. Heemels. Linear quadratic regulation of switched systems using informed policies. *IEEE Transactions on Automatic Control*, 62(6):2675–2688, 2017.

[2] G. I. Bara and M. Boutayeb. A new sufficient condition for the static output feedback stabilization of linear discrete-time systems. In *IEEE Conference on Decision and Control*, pages 4723–4728, San Diego, USA, 2006.

[3] L. Buşoniu, J. Daafouz, M. C. Bragagnolo, and I.-C. Morărescu. Planning for optimal control and performance certification in nonlinear systems with controlled or uncontrolled switches. *Automatica*, 78:297 – 308, 2017.

[4] L. Buşoniu and R. Munos. Optimistic planning for markov decision processes. In *Artificial Intelligence and Statistics*, pages 182–189, 2012.

[5] L. Buşoniu, R. Postoyan, and J. Daafouz. Near-optimal strategies for nonlinear and uncertain networked control systems. *IEEE Transactions on Automatic Control*, 61(8):2124–2139, 2016.

[6] X. Chang and G. Yang. New results on output feedback h_∞ control for linear discrete-time systems. *IEEE Transactions on Automatic Control*, 59(5):1355–1359, 2014.

[7] G. S. Deaecto and J. C. Geromel. Stability and performance of discrete-time switched linear systems. *Systems & Control Letters*, 118:1 – 7, 2018.

[8] F. Dercole and F. Della Rossa. Tree-based algorithms for the stability of discrete-time switched linear systems under arbitrary and constrained switching. *IEEE Transactions on Automatic Control*, 64(9):3823–3830, 2018.

[9] V. Gaitsgory, L. Grüne, M. Höger, C. M. Kellett, and S. R. Weller. Stabilization of strictly dissipative discrete time systems with discounted optimal control. *Automatica*, 93:311 – 320, 2018.

[10] C. A. C. Gonzaga, M. Jungers, and J. Daafouz. Stability analysis of discrete-time Lur'e systems. *Automatica*, 48(9):2277 – 2283, 2012.

[11] M. Granzotto, R. Postoyan, L. Buşoniu, D. Nešić, and J. Daafouz. Optimistic planning for the near-optimal control of nonlinear switched discrete-time systems with stability guarantees. In *IEEE Conference on Decision and Control*, Nice, France, 2019.

[12] M. Granzotto, R. Postoyan, L. Buşoniu, D. Nešić, and J. Daafouz. Finite-horizon discounted optimal control: stability and performance. In *IEEE Transactions on Automatic Control*, 2020.

[13] G. Grimm, M. J. Messina, S. E. Tuna, and A. R. Teel. Model predictive control: for want of a local control Lyapunov function, all is not lost. *IEEE Transactions on Automatic Control*, 50(5):546–558, 2005.

[14] L. Grüne and A. Rantzer. On the infinite horizon performance of receding horizon controllers. *IEEE Transactions on Automatic Control*, 53(9):2100–2111, 2008.

[15] L. Grüne, E. D. Sontag, and F. R. Wirth. Asymptotic stability equals exponential stability, and ISS equals finite energy gain — if you twist your eyes. *Systems & Control Letters*, 38(2):127 – 134, 1999.

[16] J.-F. Hren and R. Munos. Optimistic planning of deterministic systems. In *European Workshop on Reinforcement Learning*, pages 151–164, Villeneuve d'Ascq, France, 2008.

[17] S. Ibrir. Static output feedback and guaranteed cost control of a class of discrete-time nonlinear systems with partial state measurements. *Nonlinear Analysis: Theory, Methods & Applications*, 68(7):1784 – 1792, 2008.

[18] Z.-P. Jiang and Y. Wang. Input-to-state stability for discrete-time nonlinear systems. *Automatica*, 37(6):857–869, 2001.

[19] S. Keerthi and E. Gilbert. An existence theorem for discrete-time infinite-horizon optimal control problems. *IEEE Transactions on Automatic Control*, 30(9):907–909, 1985.

[20] C. M. Kellett and A. R. Teel. On the robustness of $\mathcal{KL}$-stability for difference inclusions: Smooth discrete-time Lyapunov functions. *SIAM Journal on Control and Optimization*, 44(3):777–800, 2005.

[21] B. Lincoln and B. Bernhardsson. LQR optimization of linear system switching. *IEEE Transactions on Automatic Control*, 47(10):1701–1705, 2002.

[22] B. Lincoln and A. Rantzer. Relaxing dynamic programming. *IEEE Transactions on Automatic Control*, 51(8):1249–1260, Aug 2006.

[23] R. Munos. From bandits to Monte-Carlo tree search: The optimistic principle applied to optimization and planning. *Foundations and Trends® in Machine Learning*, 7(1):1–129, 2014.

[24] A. Pavlov, I. Shames, and C. Manzie. Early termination of NMPC interior point solvers: Relating the duality gap to stability. In *European Control Conference*, pages 805–810, 2019.

[25] R. Postoyan, L. Buşoniu, D. Nešić, and J. Daafouz. Stability analysis of discrete-time infinite-horizon optimal control with discounted cost. *IEEE Transactions on Automatic Control*, 62(6):2736–2749, 2017.

[26] C. Prieur. Uniting local and global controllers with robustness to vanishing noise. *Mathematics of Control, Signals and Systems*, 14(2):143–172, 2001.

[27] C. Prieur and L. Praly. Uniting local and global controllers. In *IEEE Conference on Decision and Control*, pages 1214–1219, Phoenix, USA, 1999.

[28] C. Prieur and A. R. Teel. Uniting local and global output feedback controllers. *IEEE Transactions on Automatic Control*, 56(7):1636–1649, 2011.

[29] A. Rantzer. Relaxed dynamic programming in switching systems. *IEEE Proceedings - Control Theory and Applications*, 153(5):567–574, 2006.

[30] P. Riedinger. A switched LQ regulator design in continuous time. *IEEE Transactions on Automatic Control*, 59(5):1322–1328, 2014.

[31] M. Rinehart, M. Dahleh, and I. Kolmanovsky. Value iteration for (switched) homogeneous systems. *IEEE Transactions on Automatic Control*, 54(6):1290–1294, 2009.

[32] P. O. M. Scokaert and D. Q. Mayne. Min-max feedback model predictive control for constrained linear systems. *IEEE Transactions on Automatic Control*, 43(8):1136–1142, 1998.

[33] X. Xu and P. J. Antsaklis. Results and perspectives on computational methods for optimal control of switched systems. In *Hybrid Systems: Computation and Control*, pages 540–555, Berlin, Heidelberg, 2003.

[34] F. Zhang, H. L. Trentelman, G. Feng, and J.M.A. Scherpen. Absolute stabilization of Lur'e systems via dynamic output feedback. *European Journal of Control*, 44:15 – 26, 2018.

[35] H. Zhang, C. Qin, and Y. Luo. Neural-network-based constrained optimal control scheme for discrete-time switched nonlinear system using dual heuristic programming. *IEEE Transactions on Automation Science and Engineering*, 11(3):839–849, 2014.

[36] W. Zhang, J. Hu, and A. Abate. On the value functions of the discrete-time switched LQR problem. *IEEE Transactions on Automatic Control*, 54(11):2669–2674, 2009.

[37] W. Zhang, J. Hu, and A. Abate. A study of the discrete-time switched LQR problem. *ECE TechnicalReports*, (Paper 384), 2009.

[38] W. Zhang, J. Hu, and A. Abate. Infinite-horizon switched LQR problems in discrete time: A suboptimal algorithm with performance analysis. *IEEE Transactions on Automatic Control*, 57(7):1815–1821, 2012.

[39] F. Zhu and P. J. Antsaklis. Optimal control of hybrid switched systems: A brief survey. *Discrete Event Dynamic Systems*, 25(3):345–364, 2015.

Mathieu Granzotto received his engineering degree in Control and Automation from UFSC (Brazil) in 2016. In 2019, he received his Ph.D. in Control Theory from Université de Lorraine (France), where he is a Temporary Research and Teaching Attaché at the "Centre de Recherche en Automatique de Nancy" (France). His research interests include non-linear systems, optimal control and dynamic programming methods.

Romain Postoyan received the "Ingénieur" degree in Electrical and Control Engineering from ENSEEIHT (France) in 2005. He obtained the M.Sc. by Research in Control Theory & Application from Coventry University (United Kingdom) in 2006 and the Ph.D. in Control Theory from Université Paris-Sud (France) in 2009. In 2010, he was a research assistant at the University of Melbourne (Australia). Since 2011, he is a CNRS researcher at the "Centre de Recherche en Automatique de Nancy" (France). He obtained the "Habilitation à Diriger des Recherches" from Université de Lorraine in 2019. He served/serves as an Associate Editor for the journals: Automatica, IEEE Control Systems Letters and IMA Journal of Mathematical Control and Information.

Lucian Buşoniu received his Ph.D. degree cum laude from the Delft University of Technology, the Netherlands, in 2009. He is a professor with the Department of Automation at the Technical University of Cluj-Napoca, where he leads the group on Robotics and Nonlinear Control. He has previously held research positions in the Netherlands and in France. He serves on the editorial board of the Elsevier journal Engineering Applications of Artificial Intelligence. His research interests include nonlinear optimal control using artificial intelligence and reinforcement learning techniques, robotics, and multiagent systems. His publications include among others several influential review articles and a book on reinforcement learning.

Dragan Nešić is a Professor at the Department of Electrical and Electronic Engineering (DEEE) at The University of Melbourne, Australia. He currently serves as Associate Dean Research at the Melbourne School of Engineering. He received his Bachelor of Engineering (BE) degree in Mechanical Engineering from The University of Belgrade, Serbia (1990), and his Ph.D. degree from Systems Engineering, RSISE, Australian National University, Canberra, Australia (1997). His research interests include networked control systems, reset systems, extremum seeking control, hybrid control systems, event-triggered control, security and privacy in cyber-physical systems, and so on. He published more than 400 peer reviewed journal/conference papers in top outlets in his field. He is a Fellow of the Institute of Electrical and Electronic Engineers (IEEE, 2006) and Fellow of the International Federation for Automatic Control (IFAC, 2019). He was awarded Doctorate Honoris Causa by the University of Lorraine, France (2019) and Humboldt Research Award (2020) by the Alexander von Humboldt Foundation. He was also awarded a Future Fellowship (2010-2014) and Australian Professorial Fellowship (2004-2009) by the Australian Research Council (ARC), as well as a Humboldt Research Fellowship (2003-2004). He was invited to serve as a Distinguished Lecturer of the Control Systems Society (CSS), IEEE (2008-2012). He was a co-recipient of the George S. Axelby Outstanding Paper Award for the Best Paper in IEEE Transactions on Automatic Control (2018). He served as an Associate Editor for the journals Automatica (2003-2015), IEEE Transactions on Automatic Control (2004-2008), Systems and Control Letters (2001-2010), European Journal of Control (2007-2011) and as a General Co-Chair of IEEE Conference on Decision and Control (CDC), Melbourne (2017). He currently serves as an Associate Editor for the IEEE Transactions on Control of Network Systems (since 2016). He also served as a Member of the Board of Governors, Control Systems Society (CSS), IEEE (2011-2016).

Jamal Daafouz is a Full Professor at University de Lorraine (France) and researcher at CRAN - CNRS. In 1994, he received the Dipl.Ing. degree from INSA Toulouse. He prepared his Ph.D. at LAAS-CNRS Toulouse and he received the Ph.D. in Automatic Control from INSA Toulouse, in 1997. He also received the "Habilitation à Diriger des Recherches" from INPL (University de Lorraine), Nancy, in 2005.

His research interests include analysis, observation and control of uncertain systems, switched systems, hybrid systems, delay and networked systems with a particular interest for convex based optimisation methods.

In 2010, Jamal Daafouz was appointed as a junior member of the Institut Universitaire de France (IUF). He served as an associate editor of the following journals: Automatica, IEEE Transactions on Automatic Control, European Journal of Control and Non linear Analysis and Hybrid Systems. He is senior editor of the journal IEEE Control Systems Letters.