

On the Stability of Nonlinear Receding Horizon Control: A Geometric Perspective

Tyler Westenbroek^{*1}, Max Simchowitz^{*1}, Michael I. Jordan^{1,2}, S. Shankar Sastry¹

Abstract—

The widespread adoption of nonlinear Receding Horizon Control (RHC) strategies by industry has led to more than 30 years of intense research efforts to provide stability guarantees for these methods. However, current theoretical guarantees require that each (generally nonconvex) planning problem can be solved to (approximate) global optimality, which is an unrealistic requirement for the derivative-based local optimization methods generally used in practical implementations of RHC. This paper takes the first step towards understanding stability guarantees for nonlinear RHC when the inner planning problem is solved to first-order stationary points, but not necessarily global optima. Special attention is given to feedback linearizable systems, and a mixture of positive and negative results are provided. We establish that, under certain strong conditions, first-order solutions to RHC exponentially stabilize linearizable systems. Crucially, this guarantee requires that state costs applied to the planning problems are in a certain sense ‘compatible’ with the global geometry of the system, and a simple counter-example demonstrates the necessity of this condition. These results highlight the need to rethink the role of global geometry in the context of optimization-based control.

I. INTRODUCTION

The global stabilization of nonlinear systems is one of the most fundamental and challenging problems in control theory. In principle, the search for a stabilizing controller can be reduced to finding a control Lyapunov function [1], yet the search for such a function may be just as challenging as the search for a controller. Historically, the framing of these equivalent synthesis problems has been guided by two distinct perspectives: geometric control and optimal control.

Geometric control is a broad term that loosely refers to design methodologies which systematically exploit global system structure to achieve a control objective for a specific class of systems [2], [3], [4]. In the context of stabilization, this has led to constructive procedures for synthesizing stabilizing controllers or control Lyapunov functions for many important classes of systems (feedback linearizable, strict feedback, etc). Despite the broad impact of this collection of techniques, the traditional criticism of these approaches is that they often require significant human ingenuity and system-specific analysis to implement.

Conversely, optimal control offers a way to sidestep these challenges by formulating certain infinite-horizon problems which can be solved to synthesize optimal stabilizing

controllers automatically. Unfortunately, exact solutions to the infinite-horizon problem require solving the Hamilton-Jacobi-Bellman partial differential equation, whose exact solution is generally impractical in state spaces of large or even modest dimensionality [5]. The response to these limitations has been the rise of *receding horizon* or *model-predictive* control strategies [6], [7], [8], which attempt to stabilize the system by solving a sequence of more tractable open-loop optimal control problems along the system trajectory.

Due to the widespread adoption of these methods by practitioners, providing stability guarantees for nonlinear receding horizon methods has been an active area of research for theoreticians for more than 30 years. However, there is a distinct gap between the theoretical models of RHC studied in the literature and the practical implementations favored in industry. Specifically, while current theoretical results require that each optimal control problem posed in the RHC scheme can be solved to (near) global optimality, practical implementations of RHC, which usually employ local derivative-based search algorithms, can only be guaranteed to find approximate stationary points of each planning problem. Thus, since nonlinear RHC planning problems are generally nonconvex, providing stability guarantees for derivative-based nonlinear RHC implementations remains a largely open problem. Due to this gap, designing cost functions for derivative-based RHC schemes which reliably guide the local search algorithms towards stabilizing solutions remains a heuristic-driven and often time-consuming process.

In this paper we aim to link the geometric and optimal control perspectives, demonstrating how such a link provides new insights into the stability of derivative-based nonlinear RHC implementations. Specifically, by studying how the RHC cost functions interact with both the local and global geometry of the control system, we illustrate how the choice of cost function can lead either to provably stable behavior, or to failure modes where RHC control schemes get ‘stuck’ at undesirable stationary points. Our negative results are related through two counter-examples while our positive results are given in Theorem 1, which provides sufficient conditions which ensure that all (approximate) first-order stationary points of the RHC optimal control problems correspond to open-loop state trajectories which decay exponentially to the origin. We use this result to provide stability guarantees for nonlinear RHC when the implementation relies on derivative-based descent methods (Theorem 2), provided that the RHC planning horizon is of sufficient (though modest) length.

* Indicates equal contribution. ¹Department of Electrical Engineering and Computer Sciences, University of California at Berkeley, USA. ²Department of Statistics, University of California at Berkeley, USA.

Our sufficient conditions for exponential stability informally require (1) the state and input costs are strongly convex, (2) the Jacobian linearization of the dynamics along every trajectory is uniformly stabilizable, (3) the control cost is sufficiently small when compared to the state cost, (4) the (time-varying, affine) first-order expansion of the system dynamics along each trajectory satisfies a local ‘matching condition’, meaning that the affine term can be cancelled out by an appropriate choice of input, and (5) the nonlinearity in the input is sufficiently small. We note that assumption (1) is natural in practical implementations of RHC, and some (possibly weaker) stabilizability condition as in (2) is clearly necessary to ensure exponential stability. It is unclear whether (5) is necessary, and we leave its study to future work. We note that conditions (1) – (2) and (4) may be satisfied in some coordinate systems for the state but not others. Thus applying our sufficient conditions (or using them to design a ‘good’ RHC cost functional) may require finding an appropriate coordinate system for the system.

To shed light on conditions (3) and (4), and to establish their necessity, we examine a class of feedback linearizable systems which satisfy conditions (2), (4) and (5) in an appropriate choice of coordinates. For this class of systems conditions (1) and (4) require that the state costs are strongly convex *in the linearizing coordinates*. To demonstrate the necessity of this strong geometric condition, we examine a model for a flexible-joint manipulator which is full-state linearizable. A natural state cost is designed which is convex in the ‘original’ non-linearizing coordinates (where condition (4) is violated), but analysis reveals that the cost is non-convex in the linearizing coordinates (where condition (4) is satisfied). Thus, the chosen cost function is in some sense ‘incompatible’ with the geometry of the system as we are forced to pick a coordinate system in which either (1) or (4) are violated. Due to this mismatch, we are able to identify initial conditions from which derivative-based RHC schemes will fail to stabilize the system and get stuck at undesirable stationary points. To address condition (3) we also investigate a model for the simple inverted pendulum, which is in the class of linearizable systems discussed above. We demonstrate that even when the other four conditions are satisfied RHC may again fail to stabilize the system if the penalty on the input is too large.

Unlike prior works [7] which require global optimality for each RHC planning problem, the stability issues for derivative-based RHC implementations discussed above cannot be overcome by simply increasing the prediction horizon. Indeed, our counterexamples show that conditions (3) and (4) are necessary even when arbitrarily long planning horizons are used. The key difference here is that derivative-based planners can converge to overly-myopic sequences of control inputs, even with long planning horizons, due to the myopic nature of the optimization landscape (i.e., the presence of local minima).

In sum, the results of this paper indicate that the stability of derivative-based nonlinear RHC schemes may be fragile

unless the interaction between the geometry of the control system and cost functions are carefully considered. Fortunately, our positive results indicate that concepts from the geometric control literature may provide constructive techniques for designing RHC cost functionals which provably guide local search algorithms towards stabilizing solutions.

Further Background on RHC: For basic background on RHC, we refer the reader to any of a number of comprehensive reviews on RHC (see, e.g., [6], [9]). Simplifying considerably, previous theoretical nonlinear RHC formulations fall into either *constrained* approaches and *unconstrained* approaches. Constrained RHC formulations directly enforce stability by either constraining the terminal predictive state to lie at the origin [10], or using inequality constraints to force the system into a neighborhood containing the origin, and then stabilizing the system using a local controller [11]. The usual critique of these methods [12] is that the satisfaction of the relevant constraints may be overly demanding computationally in an online implementation. In contrast, unconstrained approaches implicitly enforce stability by either using an appropriate CLF as the terminal cost [12] or a sufficiently long prediction horizon [7]. As alluded to above, most of these stability guarantees require that a globally optimal solution can be found for each prediction problem. Several approaches provide stability guarantees using sub-optimal solutions, but generally require that an initial feasible solution is available [13], which may be restrictive in high-performance real-time scenarios, or require the availability of a CLF [8], [12], which implies that the stabilization problem has already been solved. Thus, in this paper we study unconstrained RHC formulations which use general terminal costs, and aim to provide stability guarantees which only require that a stationary point of each optimization problem can be found. We feel that this accurately reflects the spirit of optimization-based control—to stabilize the system with minimal system-specific knowledge—as well as the practical computational constraints facing practitioners.

Remark On Proofs: The proofs of Lemmas 1 and 3, Proposition 1, and Theorem 2 can be found in [14].

II. PRELIMINARIES

This paper studies control systems of the form

$$\dot{x}(t) = F(x(t), u(t)), \quad (1)$$

where $x \in \mathbb{R}^n$ is the state and $u \in \mathbb{R}^m$ the input, and $\dot{x}(t) = \frac{d}{dt}x(t)$ denotes time derivatives. We make the following assumptions about the vector field $F: \mathbb{R}^n \times \mathbb{R}^m \rightarrow \mathbb{R}^n$:

Assumption 1: *The origin is an equilibrium point of (1), namely, $F(0, 0) = 0$.*

Assumption 2: *The vector field F is continuously differentiable. Furthermore, there exist constants $L_F > 0$ such that for each $x_1, x_2 \in \mathbb{R}^n$ and $u_1, u_2 \in \mathbb{R}^m$ we have:*

$$\|F(x_1, u_1) - F(x_2, u_2)\|_2 \leq L_F (\|x_1 - x_2\|_2 + \|u_1 - u_2\|_2).$$

Taken together, these standard assumptions support the global existence and uniqueness of solutions to (1) on compact intervals of time [15, Proposition 5.6.5].

The primary object of study in this paper will be finite horizon cost functionals $J_T(\cdot; x_0): \mathcal{U}_T \rightarrow \mathbb{R}$ of the form

$$J_T(\tilde{u}; x_0) = \int_0^T Q(\tilde{x}(\tau)) + R(\tilde{u}(\tau))d\tau + V(\tilde{x}(T))$$

$$\text{s.t. } \dot{\tilde{x}}(t) = F(\tilde{x}(t), \tilde{u}(t)), \quad \tilde{x}(0) = x_0, \quad (2)$$

where $T > 0$ is a finite prediction horizon, $x_0 \in \mathbb{R}^n$ is the initial condition for (1), and the space of admissible inputs is given by $\mathcal{U}_T := \mathcal{L}^2([0, T], \mathbb{R}^m) \cap \mathcal{L}^\infty([0, T], \mathbb{R}^m)$. Here, $Q: \mathbb{R}^n \rightarrow \mathbb{R}$ is the running cost applied to the state, $R: \mathbb{R}^m \rightarrow \mathbb{R}$ is the running cost applied to the input, and $V: \mathbb{R}^n \rightarrow \mathbb{R}$ a penalty for the terminal state. For now we assume that each of these maps is continuously differentiable.

A. Linearizations and (Approximate) Stationary Points

Next, we briefly review a few basic facts from the calculus of variations which are essential for understanding out results. We endow $\mathcal{L}^2([0, T], \mathbb{R}^m)$ with the usual inner product and norm, denoted $\langle \cdot, \cdot \rangle: \mathcal{L}^2([0, T], \mathbb{R}^m) \times \mathcal{L}^2([0, T], \mathbb{R}^m) \rightarrow \mathbb{R}$ and $\| \cdot \|_2: \mathcal{L}^2([0, T], \mathbb{R}^m) \rightarrow \mathbb{R}$. Under Assumptions 2 and 3, directional (Fréchet) derivatives of $J_T(\cdot, x_0)$ are guaranteed to exist [15, Theorem 5.6.8] as there is a well-defined gradient at each point in the optimization space. We denote the directional Fréchet derivative of $J_T(\cdot, x_0)$ at the point $\tilde{u} \in \mathcal{U}_T$ in the direction $\delta u \in \mathcal{U}_T$ by $DJ_T(\tilde{u}; x_0; \delta u)$. The gradient $\nabla J_T(\tilde{u}; x_0) \in \mathcal{L}^2([0, T], \mathbb{R}^m)$ is the unique object satisfying, for each $\delta u \in \mathcal{U}_T$,

$$DJ_T(\tilde{u}; x_0; \delta u) = \int_0^T \langle \nabla J_T(\tilde{u}; x_0)(t), \delta u(t) \rangle dt, \quad (3)$$

or more compactly, $DJ_T(\tilde{u}; x_0; \delta u) = \langle \nabla J_T(\tilde{u}; x_0), \delta u \rangle$. The following notions from the optimization literature are crucial for understanding our technical results:

Definition 1: We say that an input \tilde{u} is a **first-order stationary point** (FOS) if $\nabla J_T(\tilde{u}; x_0) = 0$. We say that \tilde{u} is an ϵ -FOS if $\|\nabla J_T(\tilde{u}; x_0)\|_2 \leq \epsilon$.

In practice, derivative-based descent algorithms take an infinite number of iterations to converge to exact stationary points, thus our analysis will primarily focus on the approximate stationary points of $J_T(\cdot, x_0)$, as these can be reached in a finite number of iterations.

Finally, we discuss how to calculate the gradient $\nabla J_T(\tilde{u}; x_0)$ using first-order expansions of the system dynamics and cost functions. Let $(\tilde{x}(\cdot), \tilde{u}(\cdot))$ denote the state-input pair of (1) defined on the interval $[0, T]$ such that $\tilde{x}(0) = x_0$ and define for each $t \in [0, T]$

$$\tilde{A}(t) = \frac{\partial}{\partial x} F(\tilde{x}(t), \tilde{u}(t)), \quad \tilde{B}(t) = \frac{\partial}{\partial u} F(\tilde{x}(t), \tilde{u}(t)).$$

Definition 2: Let $(\tilde{x}(\cdot), \tilde{u}(\cdot))$, $\tilde{A}(\cdot)$ and $\tilde{B}(\cdot)$ be defined as above. We refer to the time-varying linear system $(\tilde{A}(\cdot), \tilde{B}(\cdot))$ as the **Jacobian linearization** of the vector field F along the trajectory $(\tilde{x}(\cdot), \tilde{u}(\cdot))$.

The Jacobian linearization $(\tilde{A}(\cdot), \tilde{B}(\cdot))$ can be used to construct a first-order approximations to trajectories near $(\tilde{x}(\cdot), \tilde{u}(\cdot))$ as follows. Let $\delta u \in \mathcal{U}_T$ be a small admissible perturbation to the input and let $(\hat{x}(\cdot), \hat{u}(\cdot))$ satisfy (1) and $\hat{u} = \tilde{u} + \delta u$ and $\hat{x}(0) = \tilde{x}(0) = x_0$. Then a first-order approximation to $\hat{x}(\cdot)$ is given by

$$\hat{x}(\cdot) \approx \bar{x}(\cdot) := \tilde{x}(\cdot) + \delta x(\cdot) \quad (4)$$

where $\delta x: [0, T] \rightarrow \mathbb{R}^n$ solves $\dot{\delta x}(t) = \tilde{A}(t)\delta x(t) + \tilde{B}(t)\delta u(t)$ with $\delta x(0) = 0$, and the approximation in (4) suppresses higher-order terms involving δu . Regarding gradients as row vectors, by [15, Theorem 5.6.3] we have

$$\nabla J_T(\tilde{u}; x_0)(t) = p(t)\tilde{B}(t) + \nabla R(\tilde{u}(t)) \quad (5)$$

where the *co-state* $p: [0, T] \rightarrow \mathbb{R}^{1 \times n}$ satisfies

$$-\dot{p}(t) = p(t)\tilde{A}(t) + \nabla Q(\tilde{x}(t)) \quad p(t) = \nabla V(\tilde{x}(T)). \quad (6)$$

Thus, equations (5) and (6) reveal that the gradient of the objective can be efficiently computed using a ‘backwards pass’ along the nominal trajectory $(\tilde{x}(\cdot), \tilde{u}(\cdot))$ and the linearizations of the vector field and costs along this curve.

B. Convex Time-Varying Approximations to $J_T(\cdot, x_0)$

Our primary goal throughout the paper is study the properties of (approximate) stationary points of $J_T(\cdot; x_0)$, and our primary analytical tool will be a family of convex approximations constructed using the Jacobian linearization around particular trajectories of the system. To begin constructing these approximations, observe that the evolution of the estimate $\bar{x}(\cdot)$ in (4) is given by

$$\begin{aligned} \hat{\tilde{x}}(t) &= \tilde{x}(t) + \delta \tilde{x}(t) \\ &= F(\tilde{x}(t), \tilde{u}(t)) + \tilde{A}(t)\delta x(t) + \tilde{B}(t)\delta u(t) \\ &= \tilde{A}(t)\bar{x}(t) + \tilde{B}(t)\bar{u}(t) + \tilde{d}(t), \quad \text{where} \quad (7) \\ \tilde{d}(t) &:= F(\tilde{x}(t), \tilde{u}(t)) - \tilde{A}(t)\tilde{x}(t) - \tilde{B}(t)\tilde{u}(t). \quad (8) \end{aligned}$$

We use these time-varying dynamics to approximate $J_T(\cdot; x_0)$ near the point \tilde{u} with the cost functional $J_T^{\text{jac}}(\cdot; x_0, \tilde{u}): \mathcal{U}_T \rightarrow \mathbb{R}$ defined as follows:

$$J_T^{\text{jac}}(\bar{u}; x_0, \tilde{u}) = \int_0^T Q(\bar{x}(\tau)) + R(\bar{u}(\tau))d\tau + V(\bar{x}(T))$$

$$\text{s.t. } \dot{\bar{x}}(t) = \tilde{A}(t)\bar{x}(t) + \tilde{B}(t)\bar{u}(t) + \tilde{d}(t), \quad \bar{x}(0) = x_0. \quad (9)$$

The following result, which follows from a direct comparison of the formulas for the gradients of $J_T(\cdot; x_0)$ and $J_T^{\text{jac}}(\cdot; x_0, \tilde{u})$ at the point \tilde{u} , motivates this construction:

Lemma 1: For any input $\tilde{u}(\cdot) \in \mathcal{U}_T$ we have $J_T(\tilde{u}; x_0) = J_T^{\text{jac}}(\tilde{u}; x_0, \tilde{u})$ and $\nabla J_T(\tilde{u}; x_0) = \nabla J_T^{\text{jac}}(\tilde{u}; x_0, \tilde{u})$.

III. SUFFICIENT CONDITIONS FOR EXPONENTIALLY DECAYING FIRST-ORDER STATIONARY POINTS

We begin our analysis by providing sufficient conditions which ensure that *all* (approximate) stationary points of $J_T(\cdot; x_0)$ decay exponentially to the origin at a rate that is independent of $T \geq 0$ and $x_0 \in \mathbb{R}^n$. This is a strong condition which will enable us to provide global exponential

stability guarantees for RHC schemes which use derivative-based iterative optimization schemes in Section IV-A. After stating the main result and its proof, we draw the connection to feedback linearization and investigate examples which highlight the necessity of some of our stronger assumptions.

A. Sufficient Conditions for Exponentially Decaying (Approximate) First-Order Stationary Points

We begin by introducing the Assumptions required for the proof of Theorem 1. We emphasize some of these assumptions are highly dependent on the particular set of coordinates chosen for the state, and may be satisfied in certain coordinate systems but not others. Thus, applying our sufficient conditions requires finding a coordinate system in which the following conditions hold. Later we relate these conditions to the coordinate systems that arise when performing feedback linearization, which succinctly capture the underlying geometry of the control system.

Our first assumption is strong convexity and smoothness of the running and terminal cost functions:

Assumption 3: We assume that $Q(\cdot), R(\cdot), V(\cdot)$ are twice-continuously differentiable functions, with $Q(0) = R(0) = V(0) = 0$, whose Hessians satisfy the pointwise bounds $\alpha_Q I \preceq \nabla^2 Q \preceq \beta_Q I$, $\alpha_R I \preceq \nabla^2 R \preceq \beta_R I$, and $\alpha_V I \preceq \nabla^2 V \preceq \beta_V I$ for constants $0 < \alpha_Q \leq \beta_Q$, $0 < \alpha_R \leq \beta_R$, and $0 \leq \alpha_V \leq \beta_V$.

Note that when Assumption 3 is satisfied the optimization in (9) is strongly convex. Thus, by Lemma 1, \tilde{u} is a stationary point of $J_T(\cdot, x_0)$ if and only if it is the global minimizer of $J_T^{\text{jac}}(\cdot; x_0, \tilde{u})$. Due to the convexity of the approximation, it is much easier to study the properties of stationary points of $J_T(\cdot, x_0)$ using $J_T^{\text{jac}}(\cdot; x_0, \tilde{u})$ rather than the original functional. This observation is a key insight in our proof technique. The following result extends the above discussion to approximate stationary points of $J_T(\cdot, x_0)$:

Lemma 2: (Approximate FOS) Suppose Assumption 3 holds, and that \tilde{u} is an ϵ -FOS of $J_T(\cdot; x_0)$. Then,

$$J_T(\tilde{u}; x_0) \leq \min_u J_T^{\text{jac}}(u; x_0, \tilde{u}) + \frac{\epsilon^2}{2\alpha_R}.$$

Proof: Assumption 3 implies that $J_T^{\text{jac}}(\tilde{u}; x_0, \tilde{u}) - \alpha_R \|\tilde{u}\|^2$ is convex, and thus the Polyak-Łojasiewicz inequality holds: $J_T^{\text{jac}}(\tilde{u}; x_0, \tilde{u}) \leq \min_u J_T^{\text{jac}}(u; x_0, \tilde{u}) + \|\nabla J_T^{\text{jac}}(\tilde{u}; x_0, \tilde{u})\|_2^2 / 2\alpha_R$. Since \tilde{u} is an ϵ -FOS of $J_T(\cdot; x_0)$ and $\nabla J_T(\tilde{u}; x_0) = \nabla J_T^{\text{jac}}(\tilde{u}; x_0, \tilde{u})$, it follows that $J_T^{\text{jac}}(\tilde{u}; x_0, \tilde{u}) \leq \min_u J_T^{\text{jac}}(u; x_0, \tilde{u}) + \epsilon^2 / 2\alpha_R$. Using $J_T(\tilde{u}; x_0) = J_T^{\text{jac}}(\tilde{u}; x_0, \tilde{u})$ concludes. ■

Next we place restrictions on the local structure of the control system along each of its trajectories:

Assumption 4: Along each system trajectory $(\tilde{x}(\cdot), \tilde{u}(\cdot))$ the drift term in (8) satisfies $\tilde{d}(t) \in \text{range}(\tilde{B}(t))$.

Assumption 5: There exists $\gamma > 0$ such that for each time horizon $T \geq 0$, $x_0 \in \mathbb{R}^n$ and system trajectory $(\tilde{x}(\cdot), \tilde{u}(\cdot))$ of length T with $\tilde{x}(0) = x_0$ we have

$$\inf_{\tilde{u}(\cdot)} \int_0^T \|\hat{x}(t)\|_2^2 + \|\hat{u}(t)\|_2^2 dt + \|\hat{x}(T)\|_2^2 \leq \gamma \|x_0\|_2^2, \quad (10)$$

where $\hat{x}(t) = \tilde{A}(t)\hat{x}(t) + \tilde{B}(t)\hat{u}(t)$ and $\hat{x}(s) = x_0$ and $(\tilde{A}(\cdot), \tilde{B}(\cdot))$ is the Jacobian linearization along $(\tilde{x}(\cdot), \tilde{u}(\cdot))$.

In the language of nonlinear control theory, Assumption 4 is known as a *matching condition* [4, Chapter 9.4]. The assumption implies that the drift term $\tilde{d}(t)$ can be ‘cancelled out’ by choosing the input $\tilde{u}(t) = \tilde{B}^\dagger \tilde{d}(t)$. Meanwhile, the parameter $\gamma > 0$ in Assumption 5 measures the difficulty (in terms of a simple \mathcal{L}^2 cost) of stabilizing the Jacobian Linearizations along system trajectories.

Roughly speaking, our first three conditions ensure that the state costs are in a certain sense ‘compatible’ with the local (first-order) geometry of the control system, meaning that at each point in the optimization space they guide local search algorithms to find an input \tilde{u} which drives the corresponding predictive trajectory \tilde{x} towards the origin. Indeed, by Lemma 2, when Assumptions 3, 4 and 5 all hold, at each point $\tilde{u} \in \mathcal{U}_T$ the functional $J_T(\cdot; x_0)$ has the same local structure (up to first-order approximations) as an optimal control problem with convex costs and stabilizable time-varying dynamics, namely, $J^{\text{jac}}(\cdot; x_0, \tilde{u})$. For this local convex approximation it is much clearer to see how the state costs yield state trajectories which decay to the origin. Our second counter-example investigates a situation where there does not exist a coordinate system in which both Assumption 3 and Assumption 4 can be satisfied simultaneously and local search algorithms can produce predictive trajectories which get ‘stuck’ at undesirable equilibria.

Our last technical condition, which is made Assuming 4 already holds, effectively bounds how costly it is to ‘cancel out’ the affine drift term $\tilde{d}(t)$ along each trajectory:

Assumption 6: There exists $L_x, L_u > 0$ such that for each system trajectory $(\tilde{x}(\cdot), \tilde{u}(\cdot))$ defined on $[0, T]$ we have $\|\tilde{B}^\dagger(t)\tilde{d}(t)\| \leq L_x \|\tilde{x}(t)\| + L_u \|\tilde{u}(t)\|$ for each $t \in [0, T]$. Moreover, these constants satisfy

$$L_u^2 \leq \frac{\alpha_R}{8\beta_R} \quad \text{and} \quad L_x^2 \leq \frac{\alpha_Q}{8\beta_R} \quad (11)$$

In particular, Assumption 6 states that the cost of rejecting $\tilde{d}(t)$ can only grow linearly with $\tilde{x}(t)$ and $\tilde{u}(t)$. The constant L_x can be made arbitrarily large by re-scaling the relative magnitudes of the state and input costs (so that $\alpha_Q \gg \beta_R$). However, since $\alpha_R \leq \beta_R$, the condition implies that L_u can be at most $\frac{1}{\sqrt{8}}$, which effectively limits how nonlinear the control system is with respect to the input.¹ As our first counter example demonstrates, when Assumption 6 is violated local search algorithm may again get ‘stuck’ at undesirable stationary points. The intuition for this failure mode is that even when Assumptions 4 and 5 are satisfied if $\tilde{d}(t)$ grows too quickly it may appear ‘too costly’ (from the perspective of optimization algorithms which only have access to first-order information) to reject $\tilde{d}(t)$ and drive the system to the origin.

Under these assumptions we obtain our main result:

¹We remark that the factors of $\frac{1}{8}$ in (11) can be replaced by any constant in the interval $(0, 1)$ and the proof of Theorem 1 will go through with minor modifications. However fixing a specific constant simplifies the statement of the main result.

Theorem 1: Suppose Assumptions 1 to 6 hold. Then for each $T \geq 0$, $x_0 \in \mathbb{R}^n$ and every $\tilde{u} \in \mathcal{U}_T$ which is an ϵ -FOS of $J_T(\cdot; x_0)$ the following hold. If $\alpha_V > 0$, then $\forall s \in [0, T]$,

$$\|\tilde{x}(s)\|^2 \leq C_0 \cdot (C_1 e^{-\frac{s}{c_1}} \cdot \|x_0\|^2 + C_2 \epsilon^2),$$

where $C_0 = 6L_F + \frac{2L_F\alpha_Q}{\alpha_R} + \frac{2\alpha_Q}{\alpha_V}$, $C_1 = 4\gamma \max\{\beta_V, \beta_R, \beta_Q\}$ and $C_2 = \frac{1}{2\alpha_R} (1 + 8\beta_R \max\{\frac{L_x^2}{\alpha_Q}, \frac{L_u^2}{\alpha_R}\})$. More generally, for any $\alpha_V \geq 0$, it holds that for all $s \in [0, T]$ and $\delta > 0$,

$$\|\tilde{x}(s)\|^2 \leq C_0^\delta \cdot (C_1^\delta e^{-\frac{s}{c_1}} \cdot \|x_0\|^2 + C_2 \epsilon^2)$$

where $C_0^\delta := 6L_F + \frac{2L_F\alpha_Q}{\alpha_R} + \min\{\frac{2}{\delta}, \frac{2\alpha_Q}{\alpha_R}\}$, $C_1^\delta := e^{\frac{\delta}{c_1}} C_1$.

Note that when $\epsilon = 0$ taking the square root of both sides of either bound in the statement of the theorem demonstrates that the stationary point is exponentially decaying. We emphasize that the rate of decay is uniform across all stationary points corresponding to different initial conditions $x_0 \in \mathbb{R}^n$ and prediction horizons $T > 0$. This uniformity is essential for our RHC stability results in Section IV-A. There we stipulate how long the prediction horizon $T > 0$ and how small the optimality parameter $\epsilon > 0$ must be each time a planning problem is solved to ensure stability.

B. Proof of Theorem 1

Let $(\tilde{x}(\cdot), \tilde{u}(\cdot))$ be as in the statement of the theorem, and let $\tilde{\mathcal{V}}(s) := J_{T-s}(\tilde{u}_{[s,T]}; \tilde{x}(s))$ for $s \in [0, T]$. Then the following bounds hold (under the assumptions of Theorem 1):

Lemma 3: If $\alpha_V > 0$ then for each $0 \leq s' \leq s \leq T$ we have $\|\tilde{x}(s)\|^2 \leq \frac{1}{\alpha_Q} C_0 \cdot \tilde{\mathcal{V}}(s')$. Alternatively, if $\alpha_V \geq 0$ then for each $0 \leq \delta \leq T$, $s' \leq T - \delta$ and $s' \leq s \leq T$ we have $\|\tilde{x}(s)\|^2 \leq \frac{1}{\alpha_Q} C_0^\delta \cdot \tilde{\mathcal{V}}(s')$.

Lemma 4: If $\tilde{u}(t)$ is an ϵ -FOS of $J_T(\cdot, x_0)$, then for each $s \in [0, T]$ we have $\tilde{\mathcal{V}}(s) \leq \alpha_Q (C_1 \|\tilde{x}(s)\|^2 + \frac{C_2}{2} \epsilon^2)$.

The proof of Lemma 4 can be found in the Appendix, while the proof of Lemma 3 can be found in [14]. However, we note that this under bound can be produced by under bounding the cost-to-go for $J_T(\cdot, x_0)$, and that similar under bounds have appeared in the literature [12].

By the Fundamental Theorem of Calculus,

$$\begin{aligned} -\frac{d}{ds} \tilde{\mathcal{V}}(s) &= Q(\tilde{x}(s)) + R(\tilde{u}(s)) \geq \alpha_Q \|\tilde{x}(s)\|^2 \\ &\geq \frac{1}{c_1} \tilde{\mathcal{V}}(s) - \frac{\alpha_Q C_2 \epsilon^2}{2c_1}, \end{aligned}$$

where the last line uses Lemma 4. Integrating the bound and again invoking Lemma 4,

$$\begin{aligned} \tilde{\mathcal{V}}(s) &\leq \exp(-\frac{s}{c_1}) \tilde{\mathcal{V}}(0) + \frac{\alpha_Q C_2 \epsilon^2}{2c_1} \int_{t=0}^s \exp(-\frac{t}{c_1}) dt \quad (12) \\ &\leq \exp(-\frac{s}{c_1}) \tilde{\mathcal{V}}(0) + \frac{\alpha_Q}{2} C_2 \epsilon^2 \\ &\leq \alpha_Q \cdot (C_1 e^{-\frac{s}{c_1}} \cdot \|\tilde{x}(0)\|^2 + C_2 \epsilon^2). \end{aligned}$$

Finally, in the case where $\alpha_V > 0$, Lemma 3 lets us convert the above bound to one on $\|\tilde{x}(s)\|^2$, replacing α_Q with C_0 , as desired. In the case where $\alpha_V = 0$, for each $\delta \in [0, T]$

application of Lemma 3 yields the desired result for each $s \in [0, T - \delta]$. For each $s \in [T - \delta, T]$ Lemma 3 yields

$$\begin{aligned} \|\tilde{x}(s)\|_2^2 &\leq C_0^\delta \cdot (C_1 e^{-\frac{T-\delta}{c_1}} \cdot \|\tilde{x}(0)\|^2 + C_2 \epsilon^2) \\ &\leq C_0^\delta \cdot (e^{\frac{\delta}{c_1}} C_1 e^{-\frac{s}{c_1}} \cdot \|\tilde{x}(0)\|^2 + C_2 \epsilon^2). \end{aligned}$$

C. Connection to Feedback Linearization

We begin by applying our sufficient conditions to feedback linearizable systems, perhaps the most widely studied and well-characterized class of systems in the nonlinear geometric control literature [4, Chapter 9]. Roughly speaking, a system is feedback linearizable if it can be transformed into a linear system using state feedback and a coordinate transformation. Formally, we say that (1) is feedback linearizable if it is both control-affine, namely, of the form

$$F(x, u) = f(x) + g(x)u,$$

where $f: \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $g: \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$, and if there exists a change of coordinates $\xi = \Phi(x)$, where $\Phi: \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a diffeomorphism, such that in the new coordinates the dynamics of the system are of the form

$$\dot{\xi} = \hat{A}\xi + \hat{B}[\hat{f}(\xi) + \hat{g}(\xi)u] := \hat{F}(\xi, u), \quad (13)$$

where $\hat{A} \in \mathbb{R}^{n \times n}$ and $\hat{B} \in \mathbb{R}^{n \times m}$ define a controllable pair (\hat{A}, \hat{B}) , $\hat{f}: \mathbb{R}^n \rightarrow \mathbb{R}^n$ and $\hat{g}: \mathbb{R}^n \rightarrow \mathbb{R}^{n \times m}$ is such that $\hat{g}(\xi)$ is invertible for each $\xi \in \mathbb{R}^n$. We will let $\hat{g}_i(\xi)$ denote the i -th column of $\hat{g}(\xi)$. We emphasize that this *global* transformation is distinct from the local Jacobian linearizations employed earlier. In this case the application of the feedback rule $u(\xi, v) = \hat{g}^{-1}(\xi)[- \hat{f}(\xi) + v]$, where $v \in \mathbb{R}^m$ is a new ‘virtual’ input, results in $\dot{\xi} = \hat{A}\xi + \hat{B}v$. In essence, feedback linearization reveals a linear structure underlying the global geometry of the system. Clearly \hat{F} satisfies Assumption 4, and the following proposition provides sufficient conditions for Assumptions 5 and 6 to hold in the new coordinates:

Proposition 1: Suppose that (1) is feedback linearizable, and let \hat{f} , \hat{g} , \hat{A} and \hat{B} be as defined above. Assume that *i)* there exists $\hat{L} > 0$ such that $\|\frac{d}{d\xi} \hat{f}(\xi)\| < \hat{L}$ for each $\xi \in \mathbb{R}^n$ and *ii)* $\hat{g}(\cdot)$ is constant on \mathbb{R}^n . Then there exists $\gamma > 0$ such that along each trajectory $(\tilde{\xi}(\cdot), \tilde{u}(\cdot))$ of \hat{F} the associated Jacobian linearization $(\tilde{A}(\cdot), \tilde{B}(\cdot))$ is γ -stabilizable. Furthermore the drift term $\tilde{d}(t) = \hat{f}(\xi(t)) - \hat{A}(t)\tilde{\xi}(t)$ satisfies $\|\tilde{B}^\dagger(t)\tilde{d}(t)\|_2 \leq 2\hat{L}\|\tilde{\xi}(t)\|_2$.

Remark 1: Suppose that the representations of the state running and terminal costs, $\hat{Q} := Q \circ \Phi^{-1}$ and $\hat{V} := V \circ \Phi^{-1}$, are convex in the linearizing coordinates and satisfy pointwise bounds as in Assumption 3. Further assume that the assumptions made of \hat{F} in Proposition 1 hold. Then the conclusions of Theorem 1 can be applied to the representation of $J_T(\cdot, \xi_0)$ in the linearizing coordinates by rescaling \hat{Q} and R appropriately.

Thus, the global linearizing coordinates provide a useful tool for verifying the sufficient conditions in Theorem 1. Moreover, as we illustrate with our counter-examples, they

also provide insight into what goes wrong in cases where the state costs do not lead to stabilizing behavior.

While we illustrate this point further with our examples, let us briefly remark on the necessity of the conditions in Proposition 1 (for our analysis). First, note that along a given solution $(\tilde{\xi}(\cdot), \tilde{u}(\cdot))$ we have $\tilde{A}(t) = \hat{A} + \hat{B}[\frac{d}{d\xi}f(\tilde{\xi}(t)) + \sum_{i=1}^m \frac{d}{d\xi}g_i(\tilde{\xi}(t))\tilde{u}_i(t)]$ and $\tilde{B}(t) = \hat{B}\hat{g}(\tilde{\xi}(t))$ and $\tilde{d}(t) = \tilde{A}(t)\tilde{\xi}(t) - \frac{d}{d\xi}f(\tilde{\xi}(t))\tilde{\xi}$. When condition *ii*) is violated, even when $\frac{d}{d\xi}\hat{g}(\xi)$ can be bounded globally, the linear growth of $\tilde{A}(t)$ with respect to $\tilde{u}(t)$ may make the pair $(\tilde{A}(\cdot), \tilde{B}(\cdot))$ more difficult to stabilize (in the sense of Assumption 5) for large values of the input. Moreover, in this case $\tilde{d}(t)$ will have quadratic cross-terms in $\tilde{\xi}(t)$ and $\tilde{u}(t)$, which may violate the growth conditions in Assumption 6. Similar issues arise when $\frac{d}{d\xi}\hat{f}(\xi)$ is not bounded globally. This occurs, for example, in Lagrangian mechanical systems wherein the Coriolis terms display quadratic growth in the generalized velocities of the system. The core challenge in each of these cases, from the perspective of our analysis, is that without making additional structural Assumptions beyond those in Proposition 1 it is difficult to rule out cases where the time-varying approximation to the dynamics along some trajectory of the system is arbitrarily difficult to stabilize.

D. Counterexamples

Relative Weighting of State and Input Costs: We first consider the simple inverted pendulum in Figure 1. The states are $(x_1, x_2) = (\theta, \dot{\theta})$, where θ is the angle of the arm from vertical. The dynamics are governed by

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \end{bmatrix} = \begin{bmatrix} x_2 \\ k \sin(x_2) + u \end{bmatrix},$$

where $k = g\ell$ with $g > 0$ the gravitational constant and $\ell > 0$ the length of the arm. The cost to be minimized is

$$J_T(\cdot, x_0) = \int_0^T \|\tilde{x}(t)\|_Q^2 + r\|\tilde{u}(t)\|_2^2 dt + \|\tilde{x}(T)\|_{Q_V}^2,$$

where the scalar parameter $r > 0$ is used to control the relative weighting of the input and state costs and

$$Q_V = \begin{bmatrix} 1/4 & 1/(5\sqrt{2}) \\ 1/(5\sqrt{2}) & 1 \end{bmatrix} \quad Q = \begin{bmatrix} 1 & -1/4 \\ -1/4 & 1 \end{bmatrix}.$$

Note that the dynamics are linearizable, as they are already in the form (13). Moreover, applying Proposition 1, one can show that Assumptions 3 through 6 are satisfied with parameters $\alpha_Q = \beta_Q = \frac{3}{4}$, $\alpha_R = \beta_R = r$, $L_x = k$ and $L_u = 0$. Considering $k = 10$, we find that if $r < \frac{3}{4k^2} = \frac{3}{400}$ the sufficient conditions for exponential stability of Theorem 1 are satisfied. However, if r is not small enough then there may exist undesirable first order stationary points of the cost functional. Specifically, consider the initial condition $x_0 = (\frac{3\pi}{4}, 0)^T$ and the control signal $\tilde{u}(\cdot) \equiv -k \sin(\frac{3\pi}{4}) = -\frac{10}{\sqrt{2}}$, which generates the trajectory $\tilde{x}(\cdot) \equiv x_0$. The costate along this arc is $p(\cdot) \equiv (\frac{3\pi}{16}, \frac{3\pi}{20\sqrt{2}})$ and the gradients of the objective is given by $\nabla J_T(\tilde{u}; x_0)(t) = p(t) + ru(t)$. Thus, we see that if we choose $r = \frac{3\pi}{200}$ then we will have




Fig. 1. (a) Schematic for the simple inverted pendulum (b) schematic for the inverted pendulum with a flexible joint.

$\nabla J_T(\tilde{u}; x_0)(\cdot) \equiv 0$, which demonstrates that \tilde{u} is a stationary point of the cost function. Thus, while RHC stability results which rely on global optimality of each planning problem predict stabilizing behavior for sufficiently large $T > 0$ [7], this example demonstrates that algorithms which only find first order-stationary points may be ‘too myopic’ to guarantee stability unless the input cost is small enough.

Structure of Local Drift Term: Next, consider the flexible link manipulator depicted in Figure 1. The state is $(x_1, x_2, x_3, x_4)^T = (\theta_1, \dot{\theta}_1, \theta_2, \dot{\theta}_2)^T$, where θ_1 is the angle of the arm from vertical and θ_2 is the difference between the angle of the arm and the internal angle of the motor. The dynamics are

$$\begin{bmatrix} \dot{x}_1 \\ \dot{x}_2 \\ \dot{x}_3 \\ \dot{x}_4 \end{bmatrix} = \begin{bmatrix} x_2 \\ -K_1 \sin(x_1) + K_2 x_3 \\ x_4 \\ -K_1 \sin(x_1) - (K_2 + K_3)x_3 - u \end{bmatrix},$$

where $K_1 = g\ell$, $K_2 = \frac{k}{M}$ and $K_3 = \frac{k}{I}$, where g is the gravitational constant, ℓ is the length of the arm, $k = 10$ is the spring coefficient, M the mass of the arm and I the internal inertia of the motor. For concreteness, we will assume that these physical parameters are such that $K_1 = 20$ and $K_2 = K_3 = 1$. We apply cost functionals of the form

$$J_T(\cdot; x_0) = \int_0^T \|\tilde{x}(t)\|_Q + r\|\tilde{u}(t)\|_2^2 dt + \|\tilde{x}(T)\|_{Q_V},$$

$$Q_f := \begin{bmatrix} 1 & \frac{1}{4} & 0 & 0 \\ \frac{1}{4} & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \quad Q_V := \begin{bmatrix} 5 & -1 & -\frac{1}{4} & 0 \\ -1 & 5 & 0 & 0 \\ -\frac{1}{4} & 0 & 5 & 0 \\ 0 & 0 & 0 & 5 \end{bmatrix}.$$

While the running and terminal costs satisfy Assumption 3 in the x coordinates, Assumption 4 is violated for this parameterization of the control system. Consider the initial condition $x_0 = (\pi, 0, 0, 0)^T$. Note that the input $\tilde{u}(\cdot) \equiv 0$ generates the trajectory $\tilde{x}(\cdot) \equiv x_0$. The reader may verify that the costate along this trajectory is $p(t) \equiv (\pi, \frac{1}{4}\pi, 0, 0)$ and the gradient at this point in the optimization space is defined by $J_T(\tilde{u}; x_0)(\cdot) \equiv 0$. Note that this is true for every choice of prediction horizon $T > 0$ and choice of the scaling parameter $r > 0$. Thus, regardless of the prediction horizon, algorithms which find first-order stationary points may get stuck at this undesirable equilibrium.

We can also see how the proposed cost function fails to guide local search algorithms to stabilizing solutions by

studying its structure in a set of linearizing coordinates. Indeed, consider coordinates defined by $\xi = \Phi(x)$ where

$$\Phi(x) = (x_1, x_2, -k_1 \sin(x_1) + k_2 x_3, -k_1 \cos(x_1)x_2 + k_2 x_4).$$

These coordinates are obtained by input-output linearizing the dynamics with the output $y = x_1$ (see [4, Chapter 9]), and in the new coordinates the system can be shown to be of the form (13). Thus, in the new coordinates Assumption 4 is satisfied, however, Assumption 3 is not satisfied. Indeed, due to the nonlinearities in Φ , the reader may verify that the maps $z \rightarrow \|\Phi^{-1}(z)\|_Q^2$ and $z \rightarrow \|\Phi^{-1}(z)\|_{Q_V}$ are not convex. Thus, in these coordinates, the local structure of the state costs attracts trajectories towards the undesirable equilibrium point we identified above.

IV. FIRST-ORDER STABILITY GUARANTEES FOR RECEDING HORIZON CONTROL

In *receding horizon control* (RHC) or *model predictive control*, a planner solves $\inf_{\tilde{u}(\cdot)} J_T(\tilde{u}, x(t))$, where $x(t)$ is the current state of the real world system. The planner then applies the resulting open loop predictive control until a new state measurement is received and the process can be repeated. As discussed above, most formal stability guarantees require that an (approximate) globally optimal solution is found for each (generally nonconvex) planning problem. Applying Theorem 1, we provide the first stability guarantees for a formal model of nonlinear RHC which only requires planning to approximate stationary points.

A. First-Order Receding Horizon Control

Many practical implementations of RHC use a technique known as *warm starting*, where the predictive control returned during each optimization phase is used to construct the ‘initial guess’ for the subsequent planning problem. This approach has proven highly effective for systems which require rapid re-planning to maintain stability [16].

To model this approach, we define the *first-order receding horizon control* strategy, denoted FO-RHC, as follows. First a prediction horizon $T > 0$ and a replanning interval $\delta \in (0, T]$ are chosen and a sequence of replanning times $t_k = k\delta$ for $k \in \mathbb{N}$ are defined. Next, the process takes in an initial condition of the physical system $x_0 \in \mathbb{R}^n$ and a warm-start control $\bar{u}_0 \in \mathcal{U}_T$ specified by the user. We let $(\mathbf{x}(\cdot, x_0, \bar{u}_0), \mathbf{u}(\cdot, x_0, \bar{u}_0))$ denote the resulting trajectory produced by the control scheme described below.

At each t_k for $k \in \mathbb{N}$ a warm-start routine generates an initial guess $\tilde{u}_k(\cdot) = \tilde{u}_k(\cdot; x_0, \bar{u}_0) \in \mathcal{U}_T$ for the problem $J_T(\cdot; \mathbf{x}(\cdot, x_0, \bar{u}_0))$; a simple choice for such a routine is presented momentarily. The local search method then optimizes the problem using the chosen initial guess, and produces the predictive control $\tilde{u}_k(\cdot) = \tilde{u}_k(\cdot; x_0, \bar{u}_0)$. Note that both of these quantities depend on both the initial condition of the system and the initial warm-start control specified by the user. The predictive control is constructed via $\tilde{u}_k(\cdot) = u_T^{\text{plan}}(\cdot, \mathbf{x}(t_k, x_0, \bar{u}_0), \tilde{u}_k) \in \mathcal{U}_T$, where the map u_T^{plan} is used to model how the chosen search algorithm selects a predictive

control given for a given initial condition and warm-start input. Finally, the actual control $\mathbf{u}(t, x_0, \bar{u}_0) = \tilde{u}_k(t - t_k)$ is applied on the interval $[t_k, t_{k+1})$, and the process repeats.

Assumption 7: We assume that, for any $\hat{x}_0 \in \mathbb{R}^n$, $\bar{u} \in \mathcal{U}_T$, the planned solution $\tilde{u} = u_T^{\text{plan}}(\cdot, \hat{x}_0, \bar{u}) \in \mathcal{U}_T$ satisfies the following two conditions with parameter $\epsilon_0 > 0$:

- 1) $J_T(\tilde{u}; \hat{x}_0) \leq J_T(\bar{u}; \hat{x}_0)$; and,
- 2) \tilde{u} is an $\epsilon_0 J_T(\tilde{u}; \hat{x}_0)^{1/2}$ -FOS of $J_T(\cdot; \hat{x}_0)$.

The rationale for the first condition is that many popular trajectory optimization methods are *descent methods*, and therefore only decrease the value of the functional J_T . The second condition is reasonable because such methods converge to approximate first-order stationary points, even for nonconvex landscapes [17]. The normalization by $J_T(\hat{x}_0, \hat{u})^{1/2}$ affords geometric stability in Theorem 2 by ensuring the optimization terminates close enough to a stationary point for each planning problem as the system trajectory approaches the origin.

It remains to specify how the warm-starts \bar{u}_k are produced for $k \geq 1$. We propose selecting \bar{u}_k with δ -delay, continuing until time T , and then applying zero input:

$$\bar{u}_{k+1}(t, x_0, \bar{u}_0) = \begin{cases} \tilde{u}_k(t + \delta, x_0, \bar{u}_0) & t \in [0, T - \delta] \\ 0 & t \in (T - \delta, T]. \end{cases}$$

While more sophisticated warm-starts may be adopted in practice, the above is preferable for the present analysis because (a) it does not require further system knowledge, and (b) is amenable to transparent stability guarantees.

B. Sufficient Conditions for Exponential Stability of FO-RHC

Finally, we apply Theorem 1 and its assumptions to provide sufficient conditions for the stability of FO-RHC. In order to obtain exponential convergence, we will require that, given a desired replanning interval $\delta > 0$, the prediction horizon $T > 0$ is sufficiently large and the optimality parameter $\epsilon_0 > 0$ in Assumption 7 is sufficiently small:

Theorem 2: Let the assumptions in Theorem 1 hold. Further assume that the search algorithm chosen for FO-RHC satisfies the conditions in Assumption 7. Then for any prediction horizon $T > 0$, replanning interval $\delta \in (0, T]$ and optimality parameter $\epsilon_0 < \sqrt{2\alpha_Q \mathcal{C}_2}$, and each initial condition for the physical system $x_0 \in \mathbb{R}^n$ and initial warm-start decision variable $\bar{u}_0 \in \mathcal{U}_T$ the system trajectory $(\mathbf{x}(\cdot, x_0, \bar{u}_0), \mathbf{u}(\cdot, x_0, \bar{u}_0))$ generated by the corresponding FO-RHC scheme satisfies

$$\|\mathbf{x}(t_k, x_0, \bar{u}_0)\|_2 \leq \sqrt{M(\delta, T, \epsilon_0)} e^{\eta(\delta, T, \epsilon_0)t_k} \|x_0\|_2,$$

for each $k \in \mathbb{N}$ where we define

$$\begin{aligned} M(\delta, \epsilon_0) &:= \mathcal{C}_0^\delta \mathcal{C}_1 \left(1 - \frac{\alpha_Q}{2} \mathcal{C}_2 \epsilon_0^2\right)^{-1} \\ \eta(\delta, T, \epsilon_0) &:= \frac{1}{2\delta} \ln \left(e^{-\delta/\mathcal{C}_1} + \mathcal{T}(\delta, T) + \mathcal{E}(\delta, \epsilon_0)\right) \\ \mathcal{E}(\delta, \epsilon_0) &:= \frac{1}{2} \mathcal{C}_0^\delta \mathcal{C}_2 e^{2L_F \delta} ((\delta + 1)\alpha_Q + \alpha_V) \epsilon_0^2 \\ \mathcal{T}(\delta, T) &:= \mathcal{C}_0^\delta \mathcal{C}_1 (\delta\alpha_Q + \alpha_V) e^{-\frac{T}{\mathcal{C}_1} + \delta(\frac{1}{\mathcal{C}_1} + 2L_F)}. \end{aligned}$$

To interpret the above constants, first note that for a fixed replanning interval $\delta > 0$ we have $\lim_{\epsilon_0 \rightarrow 0} M(\delta, \epsilon_0) = \bar{C}_0^\delta \bar{C}_1$ and $\lim_{\epsilon_0 \rightarrow 0} \eta(\delta, T, \epsilon_0) = \frac{1}{2\bar{C}_1}$. Thus, in the limiting case FO-RHC recovers the exponential rate of convergence predicted by Theorem 1. Next, note that $\eta(\delta, T, \epsilon_0)$ will only be negative if $e^{-\frac{\delta}{\bar{C}_1}} + \mathcal{T}(\delta, T, \epsilon_0) + \mathcal{E}(\delta, \epsilon_0) < 1$. Thus, for our estimate on the rate of convergence to be exponentially decaying we require that T is (at least) as big as $T \geq \ln(\bar{C}_0^\delta \bar{C}_1(\delta\alpha_Q + \alpha_V)) + \delta(1 + \bar{C}_1 2L_F)$.

V. FUTURE DIRECTIONS

There are many important directions for future work. First, it should be determined whether the strong assumptions required for Theorem 1 can be relaxed. Concretely, in the context of control affine systems of the form $F(x, u) = f(x) + g(x)u$, it remains to be determined under what conditions practical RHC methods can stabilize the system when the drift term $\tilde{d}(t)$ does not grow linearly in $\tilde{x}(t)$ and $\tilde{u}(t)$ (which can happen when the growth of f is super-linear or when g is not constant). The primary shortcoming of our proof technique is that we completely ‘cancel out’ $\tilde{d}(t)$ when constructing sub-optimal controls in the proof of Theorem 1, and in some cases this cancellation may be too ‘costly’ to obtain a useful upper bound. However, the disturbance term may actually be useful at certain points if it helps drive the system towards the origin, making complete cancellation unnecessary. It is interesting to note that feedback linearizing controllers, which can also cancel out ‘useful’ nonlinearities, often receive a similar criticism. Thus, we suspect that in general further structural assumptions on the system dynamics could lead to positive results.

APPENDIX

In view of Lemma 2, to obtain a bound on $\tilde{\mathcal{V}}(s)$ it suffices to bound $\mathcal{V}^{\text{jac},*}(s) := \inf_{\tilde{u}_{[s,T]}} J_{T-s}^{\text{jac}}(\cdot, \tilde{x}(s), \tilde{u}_{[s,T]})$. Moreover, we can bound $\mathcal{V}^{\text{jac},*}(s)$ by bounding $J_{T-s}^{\text{jac}}(\tilde{u}_{[s,T]}; \tilde{x}(t), \tilde{u}_{[s,T]})$ for any (possibly suboptimal) control $\tilde{u}_{[s,T]}$; for simplicity, let us drop the $[s, T]$ -subscript going forward. We select $\tilde{u}(t) = \tilde{u}_1(t) + \tilde{u}_2(t)$, where $\tilde{u}_1(t)$ satisfies $\tilde{B}(t)\tilde{u}_1(t) = -\tilde{d}(t)$, and where \tilde{u}_2 witnesses γ -stabilizability at time s as in Assumption 5.

With this choice of $\tilde{u}(t)$ the dynamics of $\tilde{x}(t)$ in J_{T-s}^{jac} are $\frac{d}{dt}\tilde{x}(t) = \tilde{A}(t)\tilde{x}(t) + \tilde{B}(t)\tilde{u}(t) + \tilde{d}(t) = \tilde{A}(t)\tilde{x}(t) + \tilde{B}\tilde{u}_2(t)$, and, writing out J_{T-s}^{jac} explicitly, we obtain

$$\mathcal{V}^{\text{jac},*}(s) \leq \int_{t=s}^T (Q(\tilde{x}(t)) + R(\tilde{u}(t)))dt + V(\tilde{x}(T)). \quad (14)$$

By the elementary bound $\|\tilde{u}(t)\|^2 \leq 2\|\tilde{u}_1(t)\|^2 + 2\|\tilde{u}_2(t)\|^2$, the following holds for constant $c = \max\{\beta_V, 2\beta_R, \beta_Q\}$,

$$\mathcal{V}^{\text{jac},*}(s) \leq 2\beta_R \int_{t=s}^T \|\tilde{u}_1(t)\|^2 dt \quad (15)$$

$$+ c \left(\int_{t=s}^T (\|\tilde{x}(t)\|^2 + \|\tilde{u}(t)\|^2) dt + \|\tilde{x}(T)\|^2 \right). \quad (16)$$

To bound (16), we observe that $\frac{d}{dt}\tilde{x}(t) = \tilde{A}(t)\tilde{x}(t) + \tilde{B}(t)\tilde{u}(t) + \tilde{d}(t) = \tilde{A}(t)\tilde{x}(t) + \tilde{B}(t)\tilde{u}_2(t)$, which corresponds

to the $\hat{x}(t)$ dynamics in the definition of γ -stabilizability; thus, (16) is at most $c \cdot \gamma \|\tilde{x}(s)\|^2$.

To bound (15), we use (6) to bound $\|\tilde{u}_1(t)\|^2 \leq 2L_x^2 \|\tilde{x}(t)\|^2 + 2L_u^2 \|\tilde{u}(t)\|^2 \leq 2c'(Q(\tilde{x}(t)) + R(\tilde{u}(t)))/\beta_R$, where $c' = \beta_R \max\{\frac{L_x^2}{\alpha_Q}, \frac{L_u^2}{\alpha_R}\}$. Hence, in view of Lemma 2, and the principle of optimality:

$$2\beta_R \int_{t=s}^T \|\tilde{u}_1(t)\|^2 dt \leq 4c'\beta_R \int_{t=s}^T (Q(\tilde{x}(t)) + R(\tilde{u}(t)))dt \\ \leq 4c'\beta_R \tilde{\mathcal{V}}(s) \leq 4c'\beta_R (\mathcal{V}^{\text{jac},*}(s) + \frac{\epsilon^2}{2\alpha_R})$$

Putting the bounds together and rearranging:

$$(1 - 4\beta_R c') \mathcal{V}^{\text{jac},*}(s) \leq c \cdot \gamma \|\tilde{x}(s)\|^2 + \frac{2c'\beta_R}{\alpha_R} \epsilon^2$$

Under Assumption 6, we have $4\beta_R c' \leq 1/2$, so that

$$\mathcal{V}^{\text{jac},*}(s) \leq 2c \cdot \gamma \|\tilde{x}(s)\|^2 + 4c'\beta_R \epsilon^2 / \alpha_R. \quad (17)$$

We recognize $2c\gamma \leq \bar{C}_1$, and $\frac{4c'\beta_R}{\alpha_R} = \bar{C}_2 - \frac{1}{2\alpha_R}$, and invoke Lemma 2 to obtain the desired bound. ■

REFERENCES

- [1] E. D. Sontag, “A ‘universal’ construction of artstein’s theorem on nonlinear stabilization,” *Systems and Control Letters*, vol. 13, no. 2, pp. 117 – 123, 1989.
- [2] A. Isidori and C. I. Byrnes, “Output regulation of nonlinear systems,” *IEEE transactions on Automatic Control*, vol. 35, no. 2, pp. 131–140, 1990.
- [3] P. V. Kokotovic, “The joy of feedback: nonlinear and adaptive,” *IEEE Control Systems Magazine*, vol. 12, no. 3, pp. 7–17, 1992.
- [4] S. Sastry, *Nonlinear systems: analysis, stability, and control*. Springer Science & Business Media, 1999, vol. 10.
- [5] R. Bellman, “Dynamic programming,” *Science*, vol. 153, no. 3731, pp. 34–37, 1966.
- [6] D. Q. Mayne, J. B. Rawlings, C. V. Rao, and P. O. Scokaert, “Constrained model predictive control: Stability and optimality,” *Automatica*, vol. 36, no. 6, pp. 789–814, 2000.
- [7] A. Jadbabaie and J. Hauser, “On the stability of unconstrained receding horizon control with a general terminal cost,” in *Proceedings of the 40th IEEE Conference on Decision and Control (Cat. No. 01CH37228)*, vol. 5. IEEE, 2001, pp. 4826–4831.
- [8] J. A. Primbs, V. Nevistić, and J. C. Doyle, “Nonlinear optimal control: A control lyapunov function and receding horizon perspective,” *Asian Journal of Control*, vol. 1, no. 1, pp. 14–24, 1999.
- [9] M. Morari and J. H. Lee, “Model predictive control: past, present and future,” *Computers & Chemical Engineering*, vol. 23, no. 4-5, pp. 667–682, 1999.
- [10] D. Q. Mayne and H. Michalska, “Receding horizon control of nonlinear systems,” in *Proceedings of the 27th IEEE Conference on Decision and Control*. IEEE, 1988, pp. 464–465.
- [11] H. Michalska and D. Q. Mayne, “Robust receding horizon control of constrained nonlinear systems,” *IEEE transactions on automatic control*, vol. 38, no. 11, pp. 1623–1633, 1993.
- [12] A. Jadbabaie, J. Yu, and J. Hauser, “Unconstrained receding-horizon control of nonlinear systems,” *IEEE Transactions on Automatic Control*, vol. 46, no. 5, pp. 776–783, 2001.
- [13] P. O. Scokaert, D. Q. Mayne, and J. B. Rawlings, “Suboptimal model predictive control (feasibility implies stability),” *IEEE Transactions on Automatic Control*, vol. 44, no. 3, pp. 648–654, 1999.
- [14] T. Westenbroek, M. Simchowiz, M. I. Jordan, and S. S. Sastry, “On the stability of nonlinear receding horizon control: a geometric perspective.” [Online]. Available: https://people.eecs.berkeley.edu/~msimchow/mpc_local.pdf
- [15] E. Polak, *Optimization: algorithms and consistent approximations*. Springer, 2012, vol. 124.
- [16] H. G. Bock, M. Diehl, D. Leineweber, and J. Schlöder, “Efficient direct multiple shooting in nonlinear model predictive control,” *Scientific Computing in Chemical Engineering II*, vol. 2, pp. 218–227, 1999.
- [17] D. P. Bertsekas, “Nonlinear programming,” *Journal of the Operational Research Society*, vol. 48, no. 3, pp. 334–334, 1997.