

Nonlinear Model Predictive Control of Robotic Systems with Control Lyapunov Functions

Ruben Grandia*, Andrew J. Taylor†, Andrew Singletary†, Marco Hutter*, and Aaron D. Ames†

*ETH Zürich, 8092 Zürich, Switzerland

†California Institute of Technology, Pasadena, CA 91125, USA

Abstract—The theoretical unification of Nonlinear Model Predictive Control (NMPC) with Control Lyapunov Functions (CLFs) provides a framework for achieving optimal control performance while ensuring stability guarantees. In this paper we present the first real-time realization of a unified NMPC and CLF controller on a robotic system with limited computational resources. These limitations motivate a set of approaches for efficiently incorporating CLF stability constraints into a general NMPC formulation. We evaluate the performance of the proposed methods compared to baseline CLF and NMPC controllers with a robotic Segway platform both in simulation and on hardware. The addition of a prediction horizon provides a performance advantage over CLF based controllers, which operate optimally point-wise in time. Moreover, the explicitly imposed stability constraints remove the need for difficult cost function and parameter tuning required by NMPC. Therefore the unified controller improves the performance of each isolated controller and simplifies the overall design process.

I. INTRODUCTION

Deploying autonomous and versatile robots into the real world comes with the challenge of ever increasing complexity in sensing, decision making, and actuation. One difficulty in designing controllers for such complex systems lies in the need to simultaneously meet a large set of design requirements. Achieving stable and safe behavior is often in conflict with performance objectives, and finding the right balance between these requirements can be a challenging task.

A disjoint, hierarchical approach is typically taken in this context. High-level trajectories are planned to satisfy performance objectives via computationally intensive nonlinear optimization, and local feedback controllers are separately designed to ensure stability. The goal of this paper is to directly integrate these two components by constructing controllers that simultaneously optimize performance along a horizon and satisfy local stability constraints in a computationally efficient manner. In particular, we seek to unify guarantee-based methods of nonlinear control with the optimization-based view in model predictive control. To this end, we combine the guarantees endowed by Control Lyapunov Functions (CLFs) [4, 44] with the optimal performance of Nonlinear Model Predictive Control (NMPC) [6, 29, 40].

Lyapunov methods are a powerful tool for certifying stability properties of nonlinear systems [21]. The use of Control Lyapunov Functions to synthesize stabilizing controllers for robotic platforms has become increasingly popular [11, 26, 36], often via quadratic programs (QPs) [2, 1]. Despite the

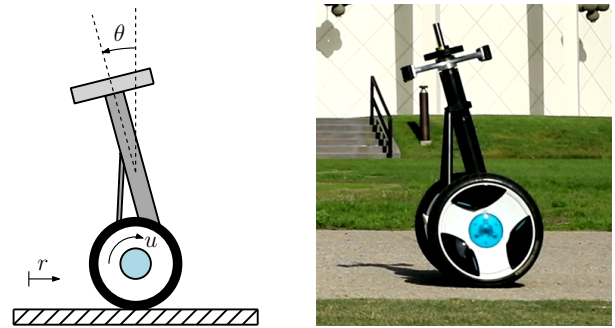


Fig. 1. Left: Segway model for simulation and control design. Right: Physical Segway system in outdoor experiment environment.

optimization-based formulation of these controllers, they often fail to achieve long-term optimal behavior. This deficiency arises due to the fact that the cost of these optimization problems fails to incorporate the future behavior of the system, but is instead point-wise optimal [10].

In contrast, Nonlinear Model Predictive Control (NMPC) emphasizes performance by solving a finite horizon optimal control problem online and applying the first element of the computed open-loop input trajectory to the system [40]. This optimization is repeatedly solved with each newly measured state to obtain the next control input trajectory. While this class of control laws can achieve strong performance in practice [12, 39, 13], and allows intuitive specification of the desired behaviour, additional assumptions must be met to certify closed-loop stability. In classical discrete-time NMPC, stability is guaranteed by an appropriately designed terminal penalty and terminal constraint [29, 7, 13].

The integration of Lyapunov methods with NMPC is not a new idea. Lyapunov methods have been used to construct stabilizing terminal conditions [19], or to analyze stability in the absence thereof [18]. Another approach incorporates the stability condition required by a CLF along the prediction horizon found with NMPC [38, 48]. As noted in [38], this approach has several desirable properties such as the absence of a terminal cost, stability for any horizon length, and recovery of the CLF-QP [2] or infinite horizon optimal controllers when considering the limiting behavior of the finite horizon. The idea of imposing stability constraints along the horizon has appeared in other forms such as contractive state constraints [8], and has been applied within the context of chemical process control [27, 47], economic cost functions [16], and switched nonlinear systems [30].

While this existing work has analyzed the stability and optimality properties obtained through the unification of CLFs and NMPC, there has been little attention to the practical and computational aspects of the resulting nonlinear optimization problem. Limited computational resources and fast system dynamics present a challenge to the deployment of these unified methods to modern robotic systems. Indeed, to the best of our knowledge, such a control scheme has not yet been applied to robotic systems experimentally.

To achieve the goal of experimental realizability, we develop a new methodology for combining CLFs with NMPC. We describe practical methods to efficiently solve the resulting nonlinear optimization problems and ultimately realize the proposed controllers in simulation and, for the first time, experimentally on a robotic platform; in this case, a Segway hardware platform seen in Fig. 1. While each proposed method provides theoretical stability guarantees, significant differences in computational efficiency and performance are observed. Furthermore, we find that the pairing of these control methodologies leads to improved performance over CLF methods and significantly reduced tuning of prediction horizon length and terminal conditions for NMPC methods.

Our paper is organized as follows. Section II provides a review of CLFs and the stability guarantees they yield, and reviews the NMPC problem and how it is solved in practice. In Section III we propose a set of methods for incorporating CLF stability constraints into the NMPC problem, and provide additional details on implementation. Lastly, in Section IV we provide results from both simulation and hardware that demonstrate the ability of this unified control approach to achieve stability and improve performance.

II. BACKGROUND

In this section we provide background information on Control Lyapunov Functions (CLFs) and Nonlinear Model Predictive Control (NMPC). This information supports the specific framework unifying CLFs and NMPC in Section III.

A. Control Lyapunov Functions

Consider a state space $\mathcal{X} \subset \mathbb{R}^n$ and a control input space $\mathcal{U} \subset \mathbb{R}^m$, where it is assumed \mathcal{X} is path-connected and $0 \in \mathcal{X}$. Consider the control-affine dynamic system given by:

$$\dot{x} = f(x) + g(x)u. \quad (1)$$

where $x \in \mathcal{X}$, $u \in \mathcal{U}$, and $f : \mathcal{X} \rightarrow \mathbb{R}^n$ and $g : \mathcal{X} \rightarrow \mathbb{R}^{n \times m}$ are Lipschitz continuous on \mathcal{X} . Further assume that $f(0) = 0$, or that the origin is an equilibrium point of the system. As in [21], we define a class \mathcal{K} function as a continuous function $\alpha : [0, a) \rightarrow \mathbb{R}_+$, with $a > 0$, $\alpha(0) = 0$ and α strictly monotonically increasing (denoted $\alpha \in \mathcal{K}$). If $a = \infty$ and $\lim_{r \rightarrow \infty} \alpha(r) = \infty$, then α is said to be a *class \mathcal{K}_∞ function* ($\alpha \in \mathcal{K}_\infty$). This type of function can be interpreted as a type of nonlinear gain function, noting the linear gain function $\alpha(r) = kr$ with $k > 0$ satisfies this definition. Given this definition, we define Control Lyapunov Functions (CLFs) as in [4, 25].

Definition 1 (Control Lyapunov Functions). A continuously differentiable function $V : \mathcal{X} \rightarrow \mathbb{R}_+$ is a *Control Lyapunov Function (CLF)* for (1) on \mathcal{X} if there exists $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{K}_\infty$ such that for all $x \in \mathcal{X}$:

$$\alpha_1(\|x\|) \leq V(x) \leq \alpha_2(\|x\|) \quad (2)$$

$$\inf_{u \in \mathcal{U}} \dot{V}(x, u) \leq -\alpha_3(\|x\|), \quad (3)$$

This definition can be restated with $\alpha_1, \alpha_2, \alpha_3 \in \mathcal{K}$ with resulting stability guarantees holding locally. The existence of a CLF for (1) implies the existence of a continuous (except possibly at $x = 0$) state-feedback controller $k : \mathcal{X} \rightarrow \mathcal{U}$ that renders the origin globally asymptotically stable [4, 43]. It is possible to make k continuous at $x = 0$ if V satisfies the small-control property [44]. If the functions $\alpha_1, \alpha_2, \alpha_3$ take the form $\alpha_i(r) = c_i r^2$, $i = 1, 2, 3$, the resulting stability is global exponential stability, with the magnitude of the state upper bounded by a function exponentially decaying in time:

$$\|x(t)\| \leq M\|x(0)\|e^{-\gamma t} \quad (4)$$

with $M, \gamma > 0$. Similarly, the CLF can be upper bounded:

$$V(x(t)) \leq V(x(0))e^{-\gamma t} \quad (5)$$

This preceding bound will be useful for enforcing Lyapunov stability guarantees within the discrete time NMPC problem.

The CLF definition implies the existence of a point-wise set of stabilizing control inputs:

$$\mathcal{U}_{\text{clf}}(x) = \{u \in \mathcal{U} \mid \dot{V}(x, u) \leq -\alpha_3(\|x\|)\}. \quad (6)$$

Thus a CLF characterizes a stabilizing feedback controller as a controller $k : \mathcal{X} \rightarrow \mathcal{U}$ such that $k(x) \in \mathcal{U}_{\text{clf}}(x)$ for all $x \in \mathcal{X}$. Furthermore, upon selection of such a controller, the CLF is a certificate of stability for the closed loop system:

$$\dot{x} = f(x) + g(x)k(x). \quad (7)$$

Establishing that a given function V serves as a CLF for (1) is often done for robotic systems constructively by specifying a controller taking values in $\mathcal{U}_{\text{clf}}(x)$ for all $x \in \mathcal{X}$ [46, 22]. Note that for any $x \in \mathcal{X}$ the set $\mathcal{U}_{\text{clf}}(x)$ is described by an affine inequality in u due to the affine nature of the dynamics:

$$\dot{V}(x, u) = \frac{\partial V}{\partial x}(x) (f(x) + g(x)u). \quad (8)$$

Due to this, the CLF itself may then be used to synthesize a optimization-based controller with more desirable properties using quadratic programs [2, 1, 11]. Specifically, we obtain a feedback control law $k(x)$ that satisfies the inequality (3):

$$\begin{aligned} k(x) = \operatorname{argmin}_{u \in \mathcal{U}} & \quad \frac{1}{2}u^\top R u + q^\top u + p \\ \text{s.t.} & \quad \frac{\partial V}{\partial x}(x) (f(x) + g(x)u) \leq -\alpha_3(\|x\|) \end{aligned} \quad (9)$$

where R is positive definite and \mathcal{U} assumed to be a polytope. Feasibility of this optimization problem is guaranteed by the satisfaction of the constraint (3) and Lipschitz continuity of this controller has been studied in [32, 20].

This controller is point-wise optimal [10], and takes a greedy approach to specifying control inputs. This often leads to poor performance compared to even non-optimization based controllers as there is no consideration of the future behavior of the system when the input is chosen. In addition to challenges in achieving longer horizon optimality, these controllers face difficulty in implementation on robotic platforms. The stability guarantees endowed by these controllers assume a continuous-time implementation, which is not possible on many modern digital control systems. Instead, control inputs are chosen and held for a small interval of time in a zero-order-hold manner. Lyapunov stability of zero-order-hold systems has been studied utilizing an approximate discretization of the nonlinear dynamics [35, 34], or in the context of model predictive control [15, 33, 31, 14].

B. Nonlinear Model Predictive Control

Nonlinear Model Predictive Control (NMPC) offers an alternative to CLF-based methods for controlling nonlinear systems, and is inherently designed to resolve the challenges of longer horizon optimality at the expensive of higher online computational cost.

We consider a *direct NMPC* approach to transform the continuous optimal control problem into a finite dimensional nonlinear program (NLP) [6]. The continuous control signal $u(t)$ is parameterized over subintervals of the prediction horizon $[0, T]$ to obtain a finite dimensional decision problem. This creates a fixed grid of nodes $k \in \{0, \dots, N\}$ defining control times t_k separated by intervals of duration $\delta t = T/(N - 1)$. In this work, we consider a piecewise constant, or zero-order-hold, parameterization of the input. Denoting $x_k = x(t_k)$ and integrating the continuous dynamics in (1) over an interval leads to a discrete time representation of the dynamics:

$$x_{k+1} = f^d(x_k, u_k) = x_k + \int_{t_k}^{t_k + \delta t} f(x(\tau)) + g(x(\tau))u_k \, d\tau. \quad (10)$$

The integral in (10) is numerically approximated with an integration method of choice to achieve a desired approximation accuracy of the evolution of the continuous time system under the zero-order-hold commands.

The general NMPC problem presented below can be formulated by defining and evaluating a cost function and constraints on the grid of nodes. Here, we write the problem in parametric form, depending on the current measured state \hat{x} and additional parameters contained in p , and with a subset of the constraints implemented as soft-constraints with slack terms:

$$\min_{X, U, S} l_N(x_N, p) + \phi(s_N) + \sum_{k=0}^{N-1} l_k(x_k, u_k, p) + \phi(s_k) \quad (11a)$$

$$\text{s.t.} \quad x_0 - \hat{x} = 0, \quad (11b)$$

$$x_{k+1} - f_k^d(x_k, u_k) = 0, \quad k = 0, \dots, N-1, \quad (11c)$$

$$h_k(x_k, u_k, p) \leq s_k, \quad k = 0, \dots, N-1, \quad (11d)$$

$$h_N(x_N, p) \leq s_N, \quad (11e)$$

$$s_k \geq 0, \quad k = 0, \dots, N, \quad (11f)$$

where $X = [x_0^\top, \dots, x_N^\top]^\top$, $U = [u_0^\top, \dots, u_{N-1}^\top]^\top$, and $S = [s_0^\top, \dots, s_N^\top]^\top$ are the sequences of state, input, and slack variables respectively. The nonlinear cost and constraint functions l_k , h_k , l_N , and h_N , are allowed to vary depending on the node index k and are dependent on problem specific parameters p and the current measured state \hat{x} . The slack variables are penalized with $\phi(s) = z^\top s + \frac{1}{2}s^\top Zs$, an exact $\ell_1 - \ell_2$ penalty [42]. Collecting all decision variables into a vector, $w = [X^\top, U^\top, S^\top]^\top$, problem (11) can be framed as a general nonlinear program (NLP):

$$\min_w F(w, p) \quad \text{s.t.} \quad \begin{cases} G(w, p) = 0, \\ H(w, p) \leq 0. \end{cases} \quad (12)$$

C. Sequential Quadratic Programming (SQP)

Interior-Point Methods (IP) and Sequential Quadratic Programming (SQP) are two popular families of algorithms for solving general NLPs [37]. Additionally, the sparsity of (12) induced by the underlying structure of (11) can be exploited to obtain solutions in real-time and at a high sampling rate, which is necessary in many dynamic robotic applications. An overview of recent advances in sparsity exploiting algorithms and software tools is provided in [23].

SQP approaches offer a distinct advantage in that successive problem instances may be warm-started with solutions from preceding instances. This serves to further decrease computation time as it is often only feasible to take a single SQP step per control iteration [24]. As NMPC computes optimal control inputs over a horizon, successive instances of (11) are similar and portions of the preceding optimal control sequence can be used to warm-start the following iteration, enabling convergence across multiple iterations of the problem, rather than iterating until convergence on one instance of the problem [9].

SQP based methods apply Newton-type iterations to KKT optimality conditions for (12), assuming some regularity conditions on the constraints [28]. The Lagrangian of the NLP in (12) is defined as:

$$\mathcal{L}(w, \lambda, \mu, p) = F(w, p) + \lambda^\top G(w, p) + \mu^\top H(w, p), \quad (13)$$

with Lagrange multipliers λ , and $\mu \geq 0$ corresponding to equality and inequality constraints, respectively. The Newton iterations can be equivalently computed by solving the following potentially non-convex QP [37]:

$$\min_{\delta w} \nabla_w F(w_i, p)^\top \delta w + \frac{1}{2} \delta w^\top B_i \delta w \quad (14a)$$

$$\text{s.t.} \quad G(w_i, p) + \nabla_w G(w_i, p)^\top \delta w = 0, \quad (14b)$$

$$H(w_i, p) + \nabla_w H(w_i, p)^\top \delta w \leq 0, \quad (14c)$$

where the decision variables, $\delta w = w - w_i$, define the update step relative to the current iteration w_i , and the Hessian $B_i = \nabla_w^2 \mathcal{L}(w_i, \lambda_i, \mu_i, p)$. Computing the solution to (14) provides a decision variable update, δw_i , and updated Lagrange multipliers λ_i^{QP} and μ_i^{QP} . These iterations are ran until the variables w_i , λ_i , and μ_i converge. This iterative approach is summarized in Algorithm 1.

Algorithm 1 Sequential Quadratic Programming (SQP)

Given $p, w_0, \lambda_0, \mu_0, F, G, H$
Initialize $(i, w_i, \lambda_i, \mu_i) \leftarrow (0, w_0, \lambda_0, \mu_0)$
while NotConverged(w_i, λ_i, μ_i) **do**
 compute $\nabla_w F(w_i, p), B_i, H(w_i, p), \nabla_w H(w_i, p),$
 $G(w_i, p), \nabla_w G(w_i, p).$

 $(\delta w_i, \lambda_i^{QP}, \mu_i^{QP}) \leftarrow \text{Solve (14)}$
 $w_{i+1} \leftarrow w_i + \delta w_i$
 $\lambda_{i+1} \leftarrow \lambda_i^{QP}$
 $\mu_{i+1} \leftarrow \mu_i^{QP}$
 $i \leftarrow i + 1$
end while

Convergence of the SQP algorithm leads to a state and input sequence, X^* and U^* , respectively. The first element of U^* , denoted as u_0 , can be applied to the system, after which the SQP algorithm is run again to determine a new control input sequence. The application of SQP as a subroutine within a NMPC feedback controller is provided in Algorithm 2.

III. UNIFYING CLFs WITH NMPC

In this section we explore different ways of integrating the stability based CLF-QP and performance driven NMPC controllers discussed in Section II. These different methods will be evaluated experimentally in Section IV.

The NMPC framework presented in Algorithm 2 can be interpreted as a closed-loop feedback controller, $k_{\text{NMPC}} : \mathcal{X} \rightarrow \mathcal{U}$. As described in Section II-A, for each state $x \in \mathcal{X}$, a CLF defines a point-wise set of stabilizing control inputs $\mathcal{U}_{\text{clf}}(x)$ given in (6). To inherit the stability guarantees provided by the CLF, we need to restrict the NMPC controller to these stabilizing inputs, such that $k_{\text{NMPC}}(x) \in \mathcal{U}_{\text{clf}}(x)$ for all $x \in \mathcal{X}$.

As noted in [38], if only the first input in the input sequence is applied before the input sequence is recomputed, the restriction of the NMPC controller to stabilizing inputs can be achieved by directly imposing the CLF condition only on the first input, subject to the current measured state \hat{x} :

$$h_{\text{CLF}}(\hat{x}, u_0) = \frac{\partial V}{\partial x}(\hat{x})(f(\hat{x}) + g(\hat{x})u_0) \leq -\alpha_3(\|\hat{x}\|). \quad (15)$$

As in the case of the controller (9), for a given state, this constraint is affine in the decision variable u_0 . Due to this, the SQP subroutine (14) contains the CLF constraint without approximation, and will therefore, in the same way as the CLF-QP, compute a stabilizing control input after solving just one QP. In the context of dynamic robotic platforms, this provides an advantage over general NMPC in that it does not require the computational cost of multiple Newton iterations to converge to a potentially stabilizing control input.

Beyond the constraint on the first input, any further modifications to the NMPC problem are done explicitly to increase performance or accommodate the discrete time implementation of control inputs. In the following, we will discuss addi-

Algorithm 2 SQP - NMPC

Given w^0, λ^0, μ^0
Initialize $(j, w^j, \lambda^j, \mu^j) \leftarrow (0, w^0, \lambda^0, \mu^0)$
while ControllerIsRunning() **do**
 $\hat{x} \leftarrow \text{StateEstimation}()$
 $y_{\text{ref}} \leftarrow \text{Commands}()$
 $p \leftarrow (\hat{x}, y_{\text{ref}})$
 $(w^{j+1}, \lambda^{j+1}, \mu^{j+1}) \leftarrow \text{Solve SQP}(p, w^j, \lambda^j, \mu^j)$
 $U^* \leftarrow \text{ExtractInputSequence}(w^{j+1})$
 $u_0 \leftarrow \text{ExtractFirstInput}(U^*)$
 $f(x) + g(x)u \leftarrow \text{ApplyInput}(u_0)$
 $j \leftarrow j + 1$
end while

tional constraints that serve to increase performance beyond that of the controller given by (9) while achieving stability.

A. Extended Horizon Constraints

While the preceding CLF constraint (15) enforces the selection of a stabilizing control input, the resulting theoretical stability properties rely on the controller being applied continuously in time, as noted in Section II-A. As this is not possible in practice, it is desirable to incorporate additional stability constraints that enforce stable behavior when control is implemented in a zero-order-hold fashion. NMPC is an advantageous framework for these types of constraints as future states to reflect desired stability properties.

In particular, consider the bound in time on the Lyapunov function established with exponential stability in (5). While this bound is continuous in time, comparison principles [21] may be used to formulate an analogous discrete time bound at the k th node:

$$h_{\text{LLS}}(x_k, \hat{x}) = V(x_k) - V(\hat{x})e^{-\gamma k \cdot \delta t} \leq 0. \quad (16)$$

The constraints given by h_{CLF} and h_{LLS} can be combined in varying ways along the length of the prediction horizon. The basic approach, denoted *CLF-0* and presented in (17a) to (17f), implements the h_{CLF} constraint only at the initial node. This is extended in the *CLF-All* approach, where the h_{CLF} constraint is enforced at each node in the horizon in (17g). These constraints are no longer affine since at nodes $k \geq 1$ both state x_k and input u_k are decision variables, and are non-linearly coupled in (15).

In the approach denoted *LLS-N*, we enforce the h_{CLF} constraint at the first node and enforce the h_{LLS} constraint at only the final node in the prediction horizon in (17h). A similar bound, which relies on evaluating V at the final node after the system has been simulated under a control law given by Sontag's universal formula [43], is enforced in [38]. Our bound differs in that it is controller independent and only relies on the bound from exponential stability. Lastly, in the approach denoted *LLS-All*, the level set constraint h_{LLS} is applied at each node in (17h), forming an exponentially contracting

funnel along the horizon. Additionally, for all formulations, the inputs are bounded $\underline{u} \leq u_k \leq \bar{u}$, enforcing $u_k \in \mathcal{U}$.

CLF-NMPC:

$$\min_{X,U,S} \phi(s_N) + \sum_{k=0}^{N-1} \frac{1}{2} u_k^\top u_k + \phi(s_k) \quad (17a)$$

$$\text{s.t} \quad x_0 - \hat{x} = 0, \quad (17b)$$

$$x_{k+1} - f_k^d(x_k, u_k) = 0, \quad k = 0, \dots, N-1, \quad (17c)$$

$$\underline{u} \leq u_k \leq \bar{u}, \quad k = 0, \dots, N-1, \quad (17d)$$

$$s_k \geq 0, \quad k = 0, \dots, N, \quad (17e)$$

$$\text{CLF-0:} \quad h_{CLF}(\hat{x}, u_0) \leq s_0, \quad (17f)$$

Additionally, for

$$\text{CLF-ALL:} \quad h_{CLF}(x_k, u_k) \leq s_k, \quad k = 1, \dots, N-1, \quad (17g)$$

$$\text{LLS-N:} \quad h_{LLS}(x_N, \hat{x}) \leq s_N, \quad (17h)$$

$$\text{LLS-All:} \quad h_{LLS}(x_k, \hat{x}) \leq s_k, \quad k = 1, \dots, N. \quad (17i)$$

B. Quadratic Approximation Strategy

When applying the SQP algorithm presented in Section II-B to the nonlinear formulations presented in (17), the quadratic subproblem (14) is repeatedly solved. As we seek to deploy NMPC on dynamic robotic platforms, it is critical that these optimization problems are well conditioned and do not provide difficulty to numerical solvers. In particular, when B_i in (14a) is positive semi-definite (p.s.d), the resulting QP is convex and can be efficiently solved [23, 45].

To ensure this, an approximate (p.s.d) Hessian can be used instead of the full Hessian of the Lagrangian. For (17), the objective function has a least-squares form, i.e. $F(w, p) = \frac{1}{2} \|R(w, p)\|^2$, in which case the Gauss-Newton approximation,

$$B_i \approx \nabla_w R(w_i, p)^\top \nabla_w R(w_i, p), \quad (18)$$

proves effective in practice [5, 17]. This neglects the curvature of $R(w, p)$, as well as the contribution by the curvature of the constraints. We use this strategy for the *CLF-0* and *CLF-All* formulations.

For the *LLS-N* and *LLS-All* formulations, we retain the contribution of the LLS constraints in the Hessian approximation. The second term in (16) is independent of the decision variables, and the properties of this constraint are thus directly determined by the structure of the CLF V . In particular, if the CLF used is convex (as is the case for many constructive techniques for producing CLFs [46, 22]), then the approximation

$$B_i \approx \nabla_w R(w_i, p)^\top \nabla_w R(w_i, p) + \sum_k \mu_{k,LLS} \nabla_w^2 h_{k,LLS}, \quad (19)$$

remains positive definite, and, as we will show in Section IV-C, improves convergence compared to a Gauss-Newton approximation.

C. Baseline Comparisons

To understand how unifying these two control methodologies impacts performance and stability, it is necessary to compare against baseline controllers given by both (9) and (11). In particular, elements should be shared between controllers to limit the impact of tuning on performance. To this end, we begin by synthesizing a CLF using feedback-linearization based constructive techniques discussed in [46], which enables the consideration of underactuated systems.

Consider an output $y : \mathcal{X} \rightarrow \mathbb{R}^k$ with relative degree 2 [41] and $k \leq m$, a time-varying reference trajectory $y : \mathbb{R}_+ \rightarrow \mathbb{R}^k$, and define the tracking error $e : \mathcal{X} \times \mathbb{R}_+ \rightarrow \mathbb{R}^k$:

$$e(x, t) = y(x) - y_d(t). \quad (20)$$

A feedback-linearizing controller exists, $k_{\text{tbl}} : \mathcal{X} \rightarrow \mathcal{U}$ that yields linear closed-loop error dynamics given by:

$$\dot{\eta}(x, t) = A\eta(x, t), \quad \eta = \begin{bmatrix} e(x, t) \\ \dot{e}(x, t) \end{bmatrix} \quad (21)$$

where the eigenvalues of $A \in \mathbb{R}^{2k \times 2k}$ have negative real part. For any positive definite $Q \in \mathbb{R}^{2k \times 2k}$, the Continuous Time Lyapunov Equation (CTLE):

$$A^\top P + PA = -Q \quad (22)$$

has a positive definite solution $P \in \mathbb{R}^{2k \times 2k}$. This enables the synthesis of a quadratic (and convex) Lyapunov function $V : \mathcal{X} \times \mathbb{R}_+ \rightarrow \mathbb{R}$ given by $V(x, t) = \eta(x, t)^\top P \eta(x, t)$ with time derivative $\dot{V}(x, t) = -\eta(x, t)^\top Q \eta(x, t)$, which is negative definite. Furthermore, the existence of the feedback linearizing controller k_{tbl} implies that V is a CLF, as:

$$\inf_{u \in \mathcal{U}} \dot{V}(x, t, u) \leq -\lambda_{\min}(Q) \|\eta(x, t)\|_2^2. \quad (23)$$

This CLF can be used in the following *CLF-QP* controller that achieves exponential stability with $\gamma = \frac{\lambda_{\min}(Q)}{\lambda_{\max}(P)}$. The constraint is slacked for numerical conditioning with $z, Z \geq 0$.

CLF-QP:

$$\min_{u \in \mathcal{U}, s \in \mathbb{R}_+} \frac{1}{2} u^\top u + zs + \frac{1}{2} Zs^2 \quad (24)$$

$$\text{s.t.} \quad \dot{V}(x, t, u) \leq -\lambda_{\min}(Q) \|\eta(x, t)\|_2^2 + s$$

To synthesize a baseline *NMPC-β* controller, elements from the construction of the CLF can be utilized. In particular, Q can be used as a running cost on the state, and the terminal value of the CLF can be penalized as in [19]:

NMPC-β:

$$\min_{X,U} \beta V(x_N, t_N) + \sum_{k=0}^{N-1} \eta(x_k, t_k)^\top Q \eta(x_k, t_k) + \frac{1}{2} u_k^\top u_k \quad (25a)$$

$$\text{s.t} \quad x_0 - \hat{x} = 0, \quad (25b)$$

$$x_{k+1} - f_k^d(x_k, u_k) = 0, \quad k = 0, \dots, N-1, \quad (25c)$$

$$\underline{u} \leq u_k \leq \bar{u}, \quad k = 0, \dots, N-1, \quad (25d)$$

That is, the CLF is used as a terminal cost and scaled up with the parameter β . As noted in [40] if β is selected large

enough, stability can be achieved without the need to specify a terminal state constraint. The baseline NMPC problem has constraints on the initial condition and dynamic evolution given by (11b) and (11c), but does not include any CLF-based constraints. The inputs are also constrained such that $u \in \mathcal{U}$.

IV. SIMULATION & EXPERIMENTAL RESULTS

In this section we provide details on the implementation of the methods established in Section III on a Segway platform, and discuss simulation and experimental results.

A. Segway System & Implementation

Dynamics simulations provide an environment for assessing attainable levels of performance of the various approaches. The simulated dynamics model reflects a modified Ninebot E+ Segway platform, seen in Fig. 1. We consider a planar representation of the Segway, with state $x = [r \ \theta \ \dot{r} \ \dot{\theta}]^\top \in \mathbb{R}^4$ where r is the horizontal position and θ is the pitch angle. The input to the Segway, $u = [i] \in [-20, 20]$, is current to the Segway motors. The equations of motion are derived via the Newton-Euler equations for an asymmetric, two-wheeled inverted pendulum with torque input. The asymmetry of the system leads to an unforced equilibrium at $x_e = [0, \theta_e, 0, 0]$ with $\theta_e = 0.138$. In the NMPC controllers, a forward Euler time discretization is used with $\delta t = 0.01$ s.

State estimation on the physical Segway is done with wheel encoders and IMU data from a VectorNav VN-100. All computations are performed on board on an ARM Cortex-A57 (quad-core) @ 2GHz CPU running the ERIKA3 RTOS. For each NMPC formulation, all functions, gradients, and Hessians in (14) are found using the CasADi auto-differentiation framework [3]. This leads to a QP with a fixed sparsity pattern, which we solve with the sparsity-exploiting solver OSQP [45]. We solve a single QP per control iteration, unless otherwise stated.

B. Simulation Results

To compare the behavior of the different control approaches, we considered a stabilization task. In particular, we simulated the system under each controller with a fixed initial condition and an objective of stabilizing to the unforced equilibrium point x_e . The performance of each controller was quantified by the average input norm over a 2 second time horizon, which provides an assessment of the total input used. These averages appear in Table I.

TABLE I. Average input norm along the simulation horizon against prediction horizon length (N) for different controller formulations defined in Section III. Absences indicate failure to stabilize the system.

N	1	10	20	30	40	50
<i>CLF-QP</i>	1.085	1.085	1.085	1.085	1.085	1.085
<i>CLF-0</i>	1.085	1.085	1.085	1.085	1.085	1.085
<i>CLF-All</i>	1.085	1.072	0.952	0.849	0.794	0.769
<i>LLS-N</i>	1.083	0.957	0.889	0.842	0.808	0.784
<i>LLS-All</i>	1.083	0.956	0.887	0.839	0.805	0.782
<i>NMPC-0.1</i>	-	-	3.232	2.435	2.036	1.783
<i>NMPC-1</i>	-	3.026	2.019	1.732	1.574	1.471
<i>NMPC-10</i>	0.828	0.607	0.704	0.823	0.926	1.006

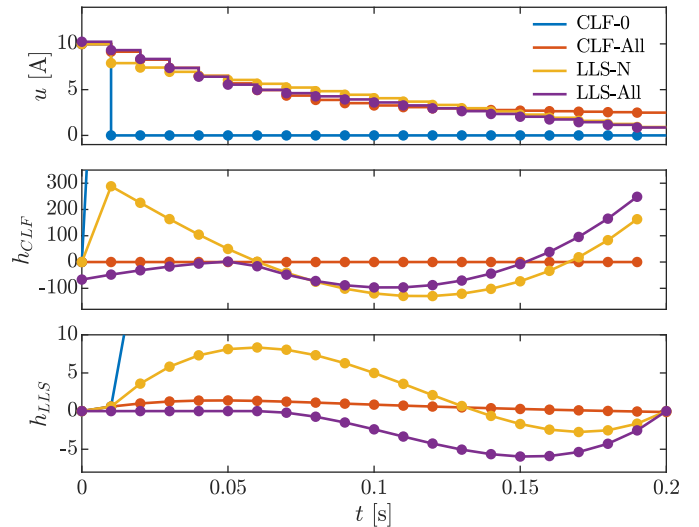


Fig. 2. Initial *CLF-NMPC* solutions stabilizing to x_e for $N = 20$ and $\hat{x} = [0, \frac{\pi}{8}, 0, 0]^\top$. **Top:** The optimal input sequence determined by each controller, with the *CLF-0* controller dropping to zero beyond the first node. **Middle:** Evaluation of the bound on \dot{V} in (15). The controllers with the h_{CLF} constraint meet this bound at all required points, while the two *LLS* controllers violated it at various points along the trajectory. **Bottom:** Evaluation of the level set bound in (16). This bound is satisfied at all necessary points by the *LLS* controllers, but is violated by the *CLF-All* controller due to the zero-order-hold implementation.

We see that the *CLF-QP* and *CLF-0* controllers select identical inputs over the entire trajectory and have the largest average input of any *CLF*-based controllers. This indicates that the addition of a prediction horizon to the baseline *CLF-QP* controller only improves its performance if the stability h_{CLF} constraint is applied further along the horizon. We also see that with no horizon ($N = 1$), all of the *CLF*-based controllers recover similar performance to the baseline *CLF-QP*. If we consider the controllers that impose constraints further along the horizon we see that their average input consistently decreases with horizon length, with the *CLF-All* controller marginally outperforming the *LLS* controllers at the longest horizon lengths.

The performance of the baseline *NMPC- β* controllers heavily depended on the weighting of the terminal cost. At shorter horizons the *NMPC-0.1* and *NMPC-1* controllers failed to stabilize the system, and saw improved performance as the horizon grew longer. In contrast, the *NMPC-10* controller demonstrated the best performance of all controllers at shorter horizons, but saw worsening performance as the horizon grew longer. This illustrates the issues that arise when both stability and performance are achieved through the cost, rather than decoupled through constraints as in the *CLF*-based formulations.

To more clearly understand the possible behaviors of the various *CLF-NMPC* controllers, we visualize the solutions to each controller obtained at the initial condition, i.e. $\hat{x} = x(0)$ in Fig. 2. As the *CLF-0* controller is only required to satisfy the stability constraint at the initial node, its input quickly drops to zero and the h_{CLF} constraint is violated in the next step in the horizon. In contrast, the *CLF-All* controller meets the

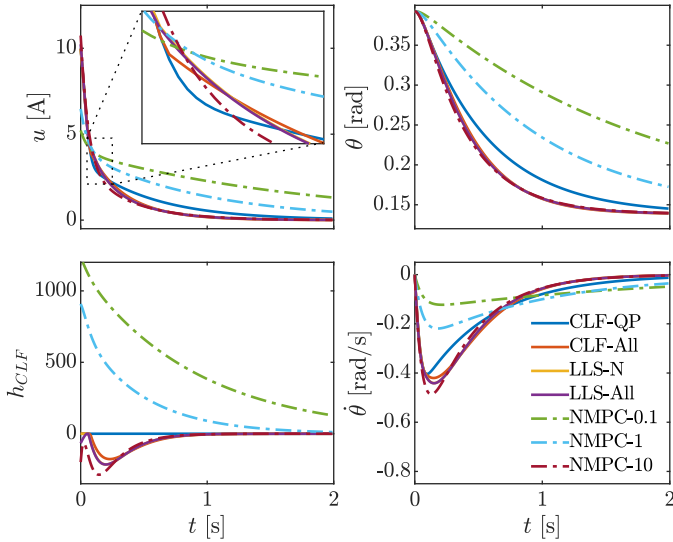


Fig. 3. Simulation results stabilizing to x_e for $x(0) = [0, \frac{\pi}{8}, 0, 0]^T$. The evaluation of the CLF constraint (15) along the simulation horizon is shown in the bottom-left plot. The system input, pitch angle, and pitch rate are shown in the top-left, top-right, and bottom-right respectively. The lower weighted *NMPC*- β controllers take significantly more time to converge to the equilibrium point.

h_{CLF} constraint along the entirety of the horizon as required. Despite meeting this constraint on \dot{V} at each node, the *CLF-All* controller slightly fails to meet the implied level-set bounds along the horizon. This is due to the fact that the constraint is checked only at the beginning of each interval and is not required to hold over the interval which the control input is held over.

The *LLS* controllers both satisfy the level constraints at the required points, with the *LLS-N* controller violating the level set bounds earlier in the horizon. Despite meeting these level set bounds and the required points, both controllers violate the associated bound on \dot{V} , with the *LLS-All* controller satisfying the h_{LLS} constraint loosely early in the trajectory before satisfying it tightly at the end of the trajectory.

The evolution of the system under these controllers with a horizon length $N = 30$ is captured in Fig. 3. We see that all *CLF*-based controllers satisfy the stability constraint (15) during the entire simulation. The most significant difference in the behavior of these controllers arises in the magnitude of the input taken early in the trajectory. The best performing controllers applying higher inputs earlier in the trajectory to stabilize the system quickly and avoid accruing input over more of the trajectory.

We additionally perform a simulation in the opposite direction, driving the Segway from its unforced equilibrium x_e to a desired state, $x_d = [0.0, \frac{\pi}{8}, 0.0, 0.0]^T$. The resulting trajectories are seen in Fig. 4. In this scenario we see that the *CLF-QP* controller demonstrates significantly different behavior from the *CLF-NMPC* controllers. In particular, the *CLF-QP* controller takes an initial input to drive the system towards the desired state, and then allows gravity to carry it to this point. This results in a large overshoot past the

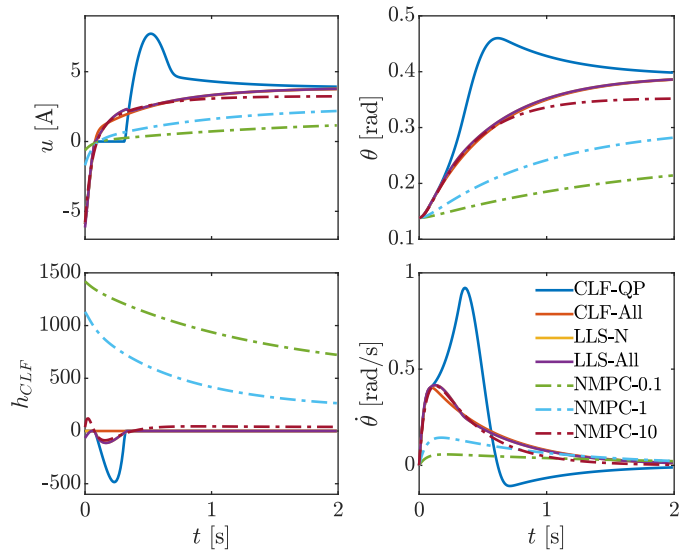


Fig. 4. Simulation results stabilizing to $x_d = [0, \frac{\pi}{8}, 0, 0]^T$ from x^* . The evaluation of the CLF constraint (15) along the simulation horizon is shown in the bottom-left plot. The system input, pitch angle, and pitch rate are shown in the top-left, top-right, and bottom-right respectively. The *CLF-QP* controller overshoots the desired state, while the *CLF-NMPC* controllers slowly approach it. The *NMPC*- β controllers do not converge due to the trade-off between input and state error in the cost function.

desired state, where as the *CLF-NMPC* controllers approach the equilibrium point more slowly.

The *NMPC*- β controllers fail to converge to the desired state. This arises due to the fact that the cost on the input is not centered about the input that makes the desired state an equilibrium point, i.e. $|u - u_{ref}|_2^2$. Because both the input and the state error appear in the cost, larger error is accepted in the state to minimize the input. This highlights the flexibility in choosing the cost function in the *CLF-NMPC* formulations. General cost functions that do not necessarily obtain their minimum at the goal can be used, which opens up the possibility of using economic cost functions [16].

C. Numerical Results

To understand the feasibility of deploying these control approaches on the computationally limited Segway platform, we investigated the convergence rate of each formulation on a single instance of the *NMPC* problem. We also considered variants of the *LLS* formulations using the Gauss-Newton approximation in (18) (denoted *LLS-N*^{GN} and *LLS-All*^{GN}) and compare them to using the modified Hessian in (19).

We execute Algorithm 1 for each formulation until the constraints are sufficiently satisfied, $\|c(w)\|_1 \leq 10^{-6}$, and no further progress is made in cost, $|F(w_i) - F(w_{i-1})| \leq 10^{-6}$. The optimization is fully cold-started such that all decision variables and Lagrange multipliers are initialized to zero. The problem is set up with $\hat{x} = x_d = [0, \frac{\pi}{8}, 0, 0]^T$ and a horizon length of $N = 30$. The step size, constraint satisfaction, and first order optimality are plotted in Fig. 5.

The *CLF-0* and *NMPC*- β formulations converge rapidly as they have a quadratic cost function and affine constraints. The only nonlinearity in these problems therefore arises from the

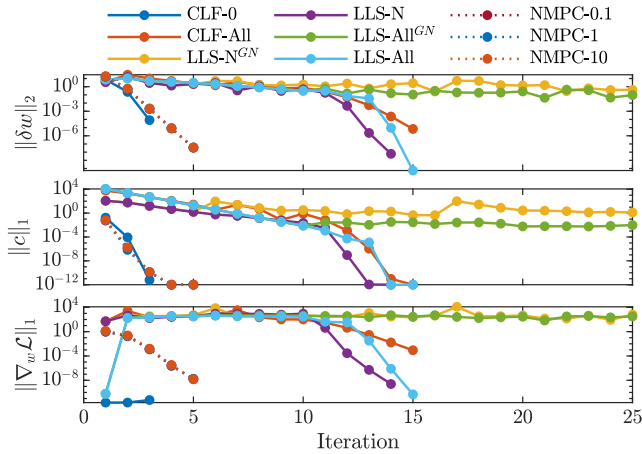


Fig. 5. Convergence of Algorithm 1 applied to a cold-started NMPC problem with $\hat{x} = x_d = [0, \frac{\pi}{8}, 0, 0]^\top$. **Top:** Step size, $\|\delta w\|_2$, **Middle:** Constraint violation, $\|c(w)\|_1$, **Bottom:** First order optimality condition, $\|\nabla_w \mathcal{L}\|_1$. The *CLF-0* and *NMPC- β* methods converge quickly, while the *CLF-NMPC* methods take longer. The *LLS* formulations without the modified Hessian in (19) stop progressing beyond a certain point.

system dynamics, making the QP subproblem a good model for the full problem. The *CLF-All*, *LLS-N*, and *LLS-All* have nonlinear constraints along the horizon, and therefore take significantly longer to converge. Nonetheless, the convergence rate accelerates over the last few iterations, indicating that these formulations can still perform well when starting close to the solution. Finally, the *LLS-N^{GN}*, and *LLS-All^{GN}* formulations initially decrease the constraint violation, but fail to make further progress after a certain point.

D. Experimental Results

In addition to simulation, we also demonstrate the ability of this unified control approach on the physical Segway in Fig. 1. In particular, we define a desired angular trajectory in time:

$$\theta_d(x, t) = \theta_e - K_v(\dot{r} - \dot{r}_d(t)), \quad (26)$$

where $\dot{r}_d(t)$ is a commanded velocity and $K_v = 0.3$ is a velocity feedback gain, leading to error dynamics given by $e(x, t) = \theta - \theta_d(x, t)$. These dynamics are used to synthesize a quadratic CLF as per Section III. This CLF was used to formulate a *CLF-QP* controller given by (24), a *NMPC-1* and *NMPC-10* controller with cost given by (25a), and *CLF-All*, *LLS-N*, and *LLS-All* controllers. Each controller was used to track the desired trajectory and then stabilize the system. The results of these experiments can be seen in Fig. 6 and the accompanying video¹. We see that all controllers except the *NMPC-1* controller are able to stabilize the system. The *CLF-QP* displays aggressive behavior compared to the *NMPC* controllers as it does not incorporate a prediction horizon. The average input and control frequency along the experimental horizon is seen in Table II. We see that at a horizon of $N = 30$ the *NMPC-10* controller has the best performance, with the *NMPC-1* controller omitted due to failure to stabilize. Of the *CLF-NMPC* controllers the *CLF-All* controller has the best performance. We note that although the baseline *NMPC-10* controller outperforms the proposed methods, this required tuning of the cost function and matches the behavior seen in

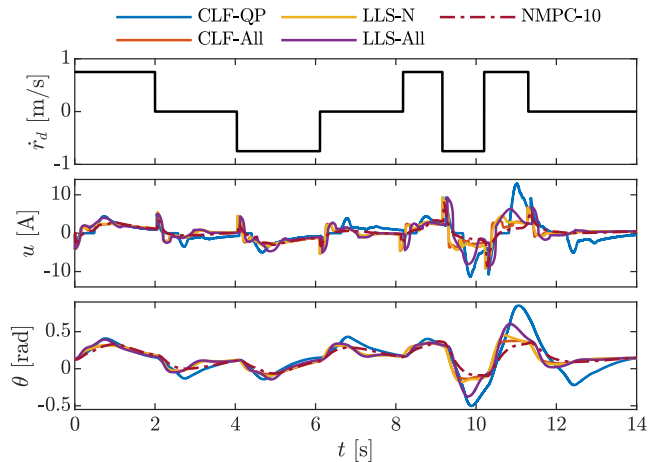


Fig. 6. Experimental results from trajectory tracking. **Top:** Desired velocity profile, **Middle:** Input profile, **Bottom:** Pitch angle profile. The *CLF-QP* controller displays more aggressive behavior due to no prediction horizon, while the *NMPC-1* controller fails to stabilize the system.

simulation at this horizon length. We see that the *CLF-NMPC* methods have a higher computational cost than the *NMPC-10* controller. This follows as the NLP has additional constraints related to stability that must be met. In that sense, *LLS-N* is an appealing approach among the *CLF-NMPC* methods, as it imposes only two stability constraints.

TABLE II. Average input norm (\bar{u}) and computation time (t_{CPU}) in ms along the experiment horizon with prediction horizon $N = 30$ for the different controller formulations defined in Section III.

	CLF-QP	CLF-All	LLS-N	LLS-All	NMPC-10
\bar{u}	2.081	1.594	1.666	1.898	1.152
t_{CPU}	1.25	5.56	4.17	6.13	3.11

V. CONCLUSION

In conclusion, we have presented a novel set of approaches for unifying CLFs and NMPC on robotic platforms with limited computational resources. The use of a SQP algorithm with modified Hessian was proposed to efficiently solve the resulting nonlinear optimization problem. The different unified formulations were analyzed in simulation, for the first time demonstrated on hardware, and were shown to improve performance beyond baseline CLF and NMPC methods. In particular for forced equilibria, the *CLF-NMPC* method converges without modifications while the cost-driven baseline NMPC does not. Furthermore, the unified methods all achieved stability where as the stability baseline NMPC methods was sensitive to cost function parameters. As system complexity increases, such manual tuning becomes increasingly difficult. In this space, we see an opportunity for the presented *CLF-NMPC* methods, where stability is explicitly embedded and requires no further tuning.

ACKNOWLEDGMENTS

R.Grandia and M. Hutter are supported via the European Union's Horizon 2020 research and innovation programme under grant agreement No 780883. A. Taylor, A. Singletary, and A. Ames are supported via DARPA awards HR00111890035 and NNN12AA01C, and NSF awards 1923239 and 1924526.

¹<https://youtu.be/weNv-FIRKiE>

REFERENCES

- [1] A. D. Ames and M. Powell. Towards the unification of locomotion and manipulation through control Lyapunov functions and quadratic programs. In *Control of Cyber-Physical Systems*, pages 219–240. Springer, 2013.
- [2] A. D. Ames, K. Galloway, K. Sreenath, and J. W. Grizzle. Rapidly exponentially stabilizing control Lyapunov functions and hybrid zero dynamics. *IEEE Transactions on Automatic Control*, 59(4):876–891, 2014.
- [3] J. A. E. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl. CasADi – A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1):1–36, 2019.
- [4] Z. Artstein. Stabilization with relaxed controls. *Nonlinear Analysis: Theory, Methods & Applications*, 7(11):1163–1173, 1983.
- [5] H. G. Bock. Recent advances in parameter identification techniques for ode. In *Numerical treatment of inverse problems in differential and integral equations*, pages 95–121. Springer, 1983.
- [6] H.G. Bock and K.J. Plitt. A Multiple Shooting Algorithm for Direct Solution of Optimal Control Problems. *IFAC Proceedings Volumes*, 17(2):1603 – 1608, 1984.
- [7] H. Chen and F. Allgöwer. A Quasi-Infinite Horizon Nonlinear Model Predictive Control Scheme with Guaranteed Stability. *Automatica*, 34(10):1205 – 1217, 1998.
- [8] S. L. de Oliveira Kothare and M. Morari. Contractive model predictive control for constrained nonlinear systems. *IEEE Transactions on Automatic Control*, 45(6):1053–1071, June 2000.
- [9] M. Diehl, H.G. Bock, J. P. Schlöder, R. Findeisen, Z. Nagy, and F. Allgöwer. Real-time optimization and nonlinear model predictive control of processes governed by differential-algebraic equations. *Journal of Process Control*, 12(4):577 – 585, 2002.
- [10] R. A. Freeman and P. V. Kokotović. Inverse optimality in robust stabilization. *SIAM journal on control and optimization*, 34(4):1365–1391, 1996.
- [11] K. Galloway, K. Sreenath, A. D. Ames, and J. W. Grizzle. Torque saturation in bipedal robotic walking through control Lyapunov function-based quadratic programs. *IEEE Access*, 3:323–332, 2015.
- [12] C. E. Garcia, D. M. Prett, and M. Morari. Model predictive control: theory and practice—a survey. *Automatica*, 25(3):335–348, 1989.
- [13] L. Grüne and J. Pannek. Nonlinear model predictive control. In *Nonlinear Model Predictive Control*, pages 45–69. Springer, 2017.
- [14] L. Grüne, D. Nešić, and J. Pannek. *Model Predictive Control for Nonlinear Sampled-Data Systems*, pages 105–113. Springer Berlin Heidelberg, 2007.
- [15] É. Gyurkovics and A. M. Elaiw. Stabilization of sampled-data nonlinear systems by receding horizon control via discrete-time approximations. *Automatica*, 40(12):2017 – 2028, 2004.
- [16] M. Heidarinejad, J. Liu, and P. D. Christofides. Economic model predictive control of nonlinear process systems using Lyapunov techniques. *AIChE Journal*, 58(3):855–870, 2012.
- [17] B. Houska, H. J. Ferreau, and M. Diehl. An auto-generated real-time iteration algorithm for nonlinear MPC in the microsecond range. *Automatica*, 47(10):2279–2285, 2011.
- [18] A. Jadbabaie and J. Hauser. On the stability of receding horizon control with a general terminal cost. *IEEE Transactions on Automatic Control*, 50(5):674–678, May 2005.
- [19] A. Jadbabaie, J. Yu, and J. Hauser. Unconstrained receding-horizon control of nonlinear systems. *IEEE Transactions on Automatic Control*, 46(5):776–783, May 2001.
- [20] M. Jankovic. Robust control barrier functions for constrained stabilization of nonlinear systems. *Automatica*, 96:359–367, 2018.
- [21] H.K. Khalil. *Nonlinear Systems - 3rd Edition*. PH, Upper Saddle River, NJ, 2002.
- [22] S. Kolathaya and S. Veer. PD based Robust Quadratic Programs for Robotic Systems. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6834–6841. IEEE, 2019.
- [23] D. Kouzoupis, G. Frison, A. Zanelli, and M. Diehl. Recent Advances in Quadratic Programming Algorithms for Nonlinear Model Predictive Control. *Vietnam Journal of Mathematics*, 46(4):863–882, Dec 2018.
- [24] W. C. Li and L. T. Biegler. A Multistep, Newton-Type Control Strategy for Constrained, Nonlinear Processes. In *1989 ACC*, pages 1526–1527, 1989.
- [25] Y. Lin and E. D. Sontag. A universal formula for stabilization with bounded controls. *Sys. & Control Letters*, 16(6):393–397, 1991.
- [26] W. Ma, S. Kolathaya, E. R. Ambrose, C. M. Hubicki, and A. D. Ames. Bipedal robotic running with DURUS-2D: Bridging the gap between theory and experiment. In *Proceedings of the 20th International Conference on Hybrid Systems: Computation and Control*, pages 265–274. ACM, 2017.
- [27] M. Mahmood and P. Mhaskar. Constrained control Lyapunov function based model predictive control design. *International Journal of Robust and Nonlinear Control*, 24(2):374–388, 2014.
- [28] O.L. Mangasarian and S. Fromovitz. The Fritz John necessary optimality conditions in the presence of equality and inequality constraints. *Journal of Mathematical Analysis and Applications*, 17(1):37 – 47, 1967.
- [29] D.Q. Mayne, J.B. Rawlings, C.V. Rao, and P.O.M. Scokaert. Constrained model predictive control: Stability and optimality. *Automatica*, 36(6):789 – 814, 2000.
- [30] P. Mhaskar, N. H. El-Farra, and P. D. Christofides. Predictive control of switched nonlinear systems with scheduled mode transitions. *IEEE Transactions on Automatic Control*, 50(11):1670–1680, 2005.

- [31] P. Mhaskar, N. H. El-Farra, and P. D. Christofides. Stabilization of nonlinear systems with state and control constraints using Lyapunov-based predictive control. *Systems & Control Letters*, 55(8):650 – 659, 2006.
- [32] B. J. Morris, M. J. Powell, and A. D. Ames. Continuity and smoothness properties of nonlinear optimization-based feedback controllers. In *2015 54th IEEE Conference on Decision and Control (CDC)*, pages 151–158. IEEE, 2015.
- [33] D. Nešić and L. Grüne. A receding horizon control approach to sampled-data implementation of continuous-time controllers. *Systems & Control Letters*, 55(8):660–672, 2006.
- [34] D. Nesic and A. R. Teel. A framework for stabilization of nonlinear sampled-data systems based on their approximate discrete-time models. *IEEE Transactions on Automatic Control*, 49(7):1103–1122, 2004.
- [35] D. Nešić, A.R. Teel, and E.D. Sontag. Formulas relating KL stability estimates of discrete-time and sampled-data nonlinear systems. *Systems & Control Letters*, 38(1):49 – 60, 1999.
- [36] Q. Nguyen and K. Sreenath. Optimal Robust Control for Bipedal Robots through Control Lyapunov Function based Quadratic Programs. In *Robotics: Science and Systems*, 2015.
- [37] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, New York, NY, USA, second edition, 2006.
- [38] J. A. Primbs, V. Nevistic, and J. C. Doyle. A receding horizon generalization of pointwise min-norm controllers. *IEEE Transactions on Automatic Control*, 45(5):898–909, 2000.
- [39] S. J. Qin and T. A. Badgwell. A survey of industrial model predictive control technology. *Control engineering practice*, 11(7):733–764, 2003.
- [40] J. B. Rawlings, D. Q. Mayne, and M. M. Diehl. *Model Predictive Control: Theory, Computation, and Design*. Nob Hill, 2nd edition edition, 2017.
- [41] S. S. Sastry. *Nonlinear Systems: Analysis, Stability and Control*. Springer, NY, 1999.
- [42] P. O. M. Scokaert and J. B. Rawlings. Feasibility issues in linear model predictive control. *AIChE Journal*, 45(8):1649–1659, 1999.
- [43] E. D. Sontag. Smooth stabilization implies coprime factorization. *IEEE transactions on automatic control*, 34(4):435–443, 1989.
- [44] E. D. Sontag. A ‘universal’ construction of Artstein’s theorem on nonlinear stabilization. *Systems & control letters*, 13(2):117–123, 1989.
- [45] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd. OSQP: An operator splitting solver for quadratic programs. In *2018 UKACC 12th International Conference on Control (CONTROL)*, pages 339–339. IEEE, 2018.
- [46] A. J. Taylor, V. D. Dorobantu, H. M. Le, Y. Yue, and A. D. Ames. Episodic Learning with Control Lyapunov Functions for Uncertain Robotic Systems. *arXiv preprint arXiv:1903.01577*, 2019.
- [47] Z. Wu, F. Albalawi, Z. Zhang, J. Zhang, H. Durand, and P. D. Christofides. Control Lyapunov-Barrier Function-Based Model Predictive Control of Nonlinear Systems. In *2018 Annual American Control Conference (ACC)*, pages 5920–5926, 2018.
- [48] J. Yu, A. Jadbabaie, J. Primbs, and Y. Huang. Comparison of nonlinear control design techniques on a model of the Caltech ducted fan. *Automatica*, 37(12):1971 – 1978, 2001.