Torque Saturation in Bipedal Robotic Walking through Control Lyapunov Function Based Quadratic Programs

Kevin Galloway, Koushil Sreenath, Aaron D. Ames, and J. W. Grizzle

***Abstract*—This paper presents a novel method to address actuator saturation by directly incorporating user-defined input bounds in controller design. In particular, we consider the application of bipedal walking and show that our method (based on a quadratic programming (QP) implementation of a control Lyapunov Function (CLF)-based controller) enables gradual performance degradation while still continuing to walk under increasingly stringent input bounds. We draw on previous work by the authors which has demonstrated the effectiveness of CLF-based controllers for stabilizing periodic gaits for biped walkers [1]. The current work presents a framework which results in more effective handling of control saturations and provides a means for incorporating a whole family of user-defined constraints into the online computation of a CLF-based controller. The paper concludes with an experimental validation of the main results on the bipedal robot MABEL, demonstrating the usefulness of the QP-based CLF approach for real-time robotic control.**

I. Introduction

Biped locomotion presents an interesting control challenge, especially since the dynamic models are typically hybrid and underactuated. The method of Hybrid Zero Dynamics (HZD) [18], [19] has provided a rigorous and intuitive method for implementing periodic walking gaits in such robotic systems, by driving the system to a lower-dimensional zero dynamics manifold on which the walking gait exists as an exponentially stable periodic orbit. Typical experimental implementation of the HZD method has relied on input-output linearization with PD control to drive the system to the zero dynamics manifold [14], but recent work by the authors has demonstrated that control Lyapunov function (CLF)-based controllers can be used to effectively implement stable walking, both in simulation and in experimental contexts [1].

The work of K. Galloway is supported by DARPA Contract W91CRB-11-1-0002, K. Sreenath is supported by CMU departmental startup funds, A. D. Ames is supported by NSF grants CNS-0953823 and CNS-1136104, NHARP project 000512-0184-2009 and NASA contract NNX12AB58G, and J. Grizzle is supported in part by DARPA and by NSF grant ECCS-1231171.

A. D. Ames is with the Department of Mechanical Engineering, Texas A&M University, College Station, TX 77843, email: aames@tamu.edu.

J. W. Grizzle is with the Control Systems Laboratory, Electrical Engineering and Computer Science Department, University of Michigan, Ann Arbor, MI 48109, email: grizzle@umich.edu.

K. Galloway is with the Electrical and Computer Engineering Department, United States Naval Academy, Annapolis, MD 21401, email: kgallowa@usna.edu.

K. Sreenath is with the Departments of Mechanical Engineering and Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, email: koushils@cmu.edu.

A variant formulation known as an exponentially stabilizing control Lyapunov function (ES-CLF) provides a means for not only guaranteeing exponential stability of a system but also providing an explicit bound on the rate of convergence. In the case of hybrid systems (such as biped robots with impulsive foot-ground impact), a even stronger convergence property is required, and therefore we turn to rapidly exponentially stabilizing control Lyapunov functions (RES-CLF). This type of CLF, which will be reviewed in more detail in Section II, incorporates an additional tuning parameter which allows the user to directly control the rate of exponential convergence. The work in [1] established the key theoretical properties of CLF-based controllers in a hybrid context, and also presented a description of the successful experimental implementation of a CLF-based controller on the robotic testbed MABEL. However, it was also noted that the user-defined control saturations were active throughout a large portion of the walking experiment, and that these saturations had a significant impact on the actual performance of the CLF-based controller as compared to the predicted performance based on theoretical bounds. In this context the hard torque limits were "blindly" applied to the calculated CLF-based control torques, without explicit consideration of the potential effect on the controller performance.

The impact of actuator saturation in feedback systems is often detrimental to stability and performance, and it therefore has been the study of a large body of research. (See [2] for instance, which provides an extensive bibliography on the topic.) In the context of robotic biped locomotion, torque saturations can limit the ability to recover from disturbances and result in instability. Typically, torque saturation is considered during the design of walking gaits, where actuator limitations are included as inequality constraints for an offline gait-design optimization routine (see [9] for instance). However, while this approach can guarantee that the torques required on the periodic walking gait are within limits, it does not account for disturbances such as rough terrain or model uncertainties which demand higher torques during recovery phase. In other work, such as [4], torque saturations are incorporated into calculation of a feedback control designed to track a time-based reference trajectory, with tracking error traded off in order to keep torque controls within limits.

The main contribution of this paper is to provide a novel control design framework for application to bipedal robotics that enables gradual performance degradation while still continuing to walk under a range of stringent torque limits.

We achieve this through an alternative method of controller implementation based on quadratic programming, that not only preserves (as much as possible) the desirable performance characteristics promised by the CLF theory, but also respects the user-defined bounds on the inputs. While the consideration of saturation effects in feedback systems has been the subject of much study, to the authors' knowledge CLF-based controllers that respect saturations have not been addressed before, and this paper provides a constructive technique for doing so, while also presenting an experimental implementation on an embedded hard real-time system with a high control rate of 1 kHz. The use of QP in biped control can be found (for example) in [3] for realizing desired link accelerations, for maintaining balance after disturbances by modifying predefined reference trajectories [16], and for applying model predictive control approaches to biped control [20], [6]. The current work is unique in the fact that QP is used for implementing CLF-based controllers with saturation constraints.

The paper proceeds as follows. In Section II, we state the dynamics of the relevant model and review the results on CLF-based control of biped robots from [1]. Section III discusses the adverse effects of user-specified control input saturations on the CLF-based controller, providing the motivation for Section IV which introduces a new method for using quadratic programming to appropriately handle torque saturation constraints for the CLF-based controllers. Section V presents simulation and experimental results, and we conclude with a summary in Section VI.

II. Control Lyapunov Functions for Hybrid Systems Revisited

In this section we introduce the model for a biped robot and review the recent innovations introduced in [1] for using control Lyapunov functions to control such systems.

A. Model

The dynamics for a biped robot (such as MABEL, the robot described in Section V) can be derived by the standard method of Lagrange and take the form

$$D(q)\ddot{q} + C(q,\dot{q})\dot{q} + G(q) = B(q)u, \tag{1}$$

where $q \in \mathscr{Q}$ is the robot configuration variable, u represents the motor control torques, and D, C and G are respectively the inertia matrix, Coriolis matrix, and gravity vector. In the case of MABEL the configuration vector q is 7-dimensional and is as described in [14] and depicted in Figure 7a. Reformulating the dynamics (1) as

$$\begin{bmatrix} \dot{q} \\ \ddot{q} \end{bmatrix} = f(q,\dot{q}) + g(q,\dot{q})u, \tag{2}$$

we also define output functions of the form

$$y(q) := H_0 q - y_d(\theta(q)), \tag{3}$$

where $\theta(q)$ is a strictly monotonic function of the configuration variable q, H_0 is an appropriately-sized matrix prescribing linear combinations of state variables to be controlled, and $y_d(\cdot)$ prescribes the desired evolution of these quantities. (See [14] for details.) The method of Hybrid Zero Dynamics (HZD) aims to drive these output functions (and their first derivatives) to zero, thereby imposing "virtual constraints" such that the system evolves on the lower-dimensional zero dynamics manifold, given by

$$Z = \{(q,\dot{q}) \in T\mathscr{Q} \mid y(q) = 0,\ L_f y(q,\dot{q}) = 0\}, \tag{4}$$

where L_f denotes the Lie derivative[10].

B. Input-output linearization

If $y(q)$ has vector relative degree 2, then the second derivative takes the form

$$\ddot{y} = L_f^2 y(q,\dot{q}) + L_g L_f y(q,\dot{q})u, \tag{5}$$

where the decoupling matrix $L_g L_f y(q,\dot{q})$ is invertible due to the vector relative degree assumption. Then defining

$$u^*(q,\dot{q}) := -(L_g L_f y(q,\dot{q}))^{-1} L_f^2 y(q,\dot{q}), \tag{6}$$

and applying a pre-control law of the form

$$u(q,\dot{q}) = u^*(q,\dot{q}) + \mu \tag{7}$$

or

$$u(q,\dot{q}) = u^*(q,\dot{q}) + (L_g L_f y(q,\dot{q}))^{-1}\mu \tag{8}$$

renders Z invariant (provided μ vanishes on Z). (Note that $u^*(q,\dot{q})$ is a feed-forward term representing the torque required to remain on Z.)

Under these assumptions, the dynamics (2) can be decomposed into zero dynamics states $z \in Z$ and transverse variables $\eta = \begin{bmatrix} y & \dot{y} \end{bmatrix}$. (See [19], [10] for details.) Under a pre-control law of the form (7) or (8), the closed-loop dynamics in terms of (η, z) take the form

$$\dot{\eta} = \bar{f}(\eta, z) + \bar{g}(\eta, z)\mu \tag{9}$$

$$\dot{z} = p(\eta, z). \tag{10}$$

For the work presented here, we will use the pre-control law (8) so that $\bar{f}(\eta,z) = F\eta$ and $\bar{g}(\eta,z) = G$, where

$$F = \begin{bmatrix} 0 & I \\ 0 & 0 \end{bmatrix}, \qquad G = \begin{bmatrix} 0 \\ I \end{bmatrix}. \tag{11}$$

The most common approach to controlling the transverse variables (i.e. driving η to zero) relies on input-output linearization with PD control, using (8) with

$$\mu = \begin{bmatrix} -\frac{1}{\varepsilon^2}K_P & -\frac{1}{\varepsilon}K_D \end{bmatrix}\eta, \tag{12}$$

where K_P and K_D are diagonal matrices chosen such that

$$A := \begin{bmatrix} 0 & I \\ -K_P & -K_D \end{bmatrix} \tag{13}$$

is Hurwitz.

C. CLF-based control

Recently, a new method based on control Lyapunov functions has been introduced in [1], which provides an alternative method for controlling the transverse variables. That method can be summarized as follows.

A function $V_\varepsilon(\eta)$ is a *rapidly exponentially stabilizing control Lyapunov function (RES-CLF)* for the system (9) if there exist positive constants $c_1, c_2, c_3 > 0$ such that for all $0 < \varepsilon < 1$ and all states (η, z) it holds that

$$c_1\|\eta\|^2 \le V_\varepsilon(\eta) \le \frac{c_2}{\varepsilon^2}\|\eta\|^2 \tag{14}$$

$$\inf_{\mu\in U}\left[L_{\bar f}V_\varepsilon(\eta,z) + L_{\bar g}V_\varepsilon(\eta,z)\mu + \frac{c_3}{\varepsilon}V_\varepsilon(\eta)\right] \le 0. \tag{15}$$

One way to generate a RES-CLF $V_\varepsilon(\eta)$ is to first solve the Lyapunov equation $A^TP+PA=-Q$ for P (where A is given by (13) and Q is any symmetric positive-definite matrix), and then define

$$V_\varepsilon(\eta) = \eta^T \begin{bmatrix}\frac{1}{\varepsilon}I & 0\\ 0 & I\end{bmatrix} P \begin{bmatrix}\frac{1}{\varepsilon}I & 0\\ 0 & I\end{bmatrix}\eta =: \eta^T P_\varepsilon \eta, \tag{16}$$

for which we have

$$\begin{aligned} L_{\bar f}V_\varepsilon(\eta,z) &= \eta^T(F^TP_\varepsilon + P_\varepsilon F)\eta,\\ L_{\bar g}V_\varepsilon(\eta,z) &= 2\eta^T P_\varepsilon G. \end{aligned} \tag{17}$$

Associated with a RES-CLF is the set of all μ for which (15) is satisfied,

$$K_\varepsilon(\eta,z) = \{\mu\in U : L_{\bar f}V_\varepsilon(\eta,z) + L_{\bar g}V_\varepsilon(\eta,z)\mu + \frac{c_3}{\varepsilon}V_\varepsilon(\eta) \le 0\},$$

and one can show that for any Lipschitz continuous feedback control law $\mu_\varepsilon(\eta,z)\in K_\varepsilon(\eta,z)$, it holds that

$$\|\eta(t)\| \le \frac{1}{\varepsilon}\sqrt{\frac{c_2}{c_1}}e^{-\frac{c_3}{2\varepsilon}t}\|\eta(0)\|, \tag{18}$$

i.e. the rate of exponential convergence to the zero dynamics manifold can be directly controlled with the constant ε through $\frac{c_3}{\varepsilon}$. There are various methods for finding a feedback control law $\mu_\varepsilon(\eta,z)\in K_\varepsilon(\eta,z)$; in practical applications, it is often important to select the control law of minimum norm. If we let $c_3 = \frac{\lambda_{\min}(Q)}{\lambda_{\max}(P)}$ and define

$$\begin{aligned} \psi_{0,\varepsilon}(\eta,z) &= L_{\bar f}V_\varepsilon(\eta,z) + \frac{c_3}{\varepsilon}V_\varepsilon(\eta,z)\\ \psi_{1,\varepsilon}(\eta,z) &= L_{\bar g}V_\varepsilon(\eta,z)^T, \end{aligned} \tag{19}$$

then this pointwise min-norm control law [7] can be explicitly formulated as

$$\mu_\varepsilon(\eta,z) = \left\{ \begin{array}{ll} -\frac{\psi_{0,\varepsilon}(\eta,z)\psi_{1,\varepsilon}(\eta,z)}{\psi_{1,\varepsilon}(\eta,z)^T\psi_{1,\varepsilon}(\eta,z)} & \text{if} \quad \psi_{0,\varepsilon}(\eta,z) > 0\\ 0 & \text{if} \quad \psi_{0,\varepsilon}(\eta,z) \le 0, \end{array} \right\} \tag{20}$$

wherein we can take $\mu = \mu_\varepsilon$ in (8).

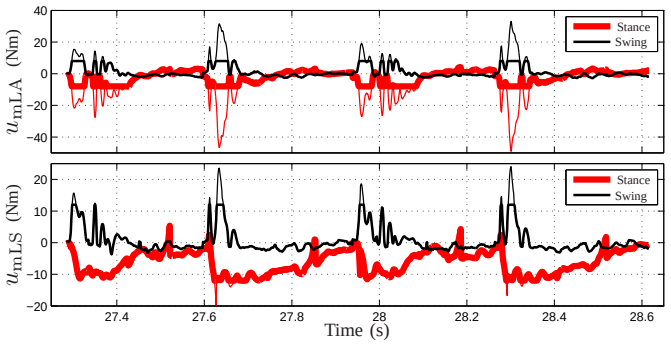

Fig. 1: Motor torques (from the MABEL experiment described in [1]) for the stance and swing legs for 4 consecutive steps of walking with the minimum-norm CLF-based controller given in (20). The thicker plots indicate the experimental (saturated) torques, while the thinner plots are the raw (unsaturated) torques computed by the CLF-based controller. For the leg angle motor (top graph), the raw (desired) control torque is at times more than 400% of the (actually implemented) saturated value. Moreover, this occurs over a significant duration of the step. Note that the symbols $u_{\mathrm{mLA}}, u_{\mathrm{mLS}}$ indicate the motor torques at the *leg angle* and *leg shape* coordinates respectively, which are linear combinations of the thigh and knee angles [13].

III. Adverse effects of torque saturation on the CLF-based controller

The approach described in Section II was successfully implemented on the robotic testbed MABEL, producing a stable walking gait. (See [1] for a description of the experiment and a reference to the online video.) However, analysis of the experimental data reveals that the user-imposed saturations on the control torque inputs were active throughout much of the experiment (see Figure 1) and significantly affected the implementation of the CLF-based control method. Though necessary to prevent unsafe or damaging motions, these saturation constraints were not applied in a manner that appropriately preserved the qualities of the CLF-based controller, and therefore the nominal bounds given by (15) and (18) were frequently violated.

Limits for control inputs are typically imposed by the user to ensure that motor torque specifications are not exceeded. When the calculated ideal control effort frequently exceeds the prescribed bounds and must therefore be truncated, the controller performance is degraded and theoretical performance measures may be violated, as in the experiment described above. More importantly, when a control input is saturated, the system runs in open-loop and is no longer able to respond to increasing errors in tracking, often leading to eventual failure.

Designing controllers which respect such bounds is important, and therefore a variety of approaches have been developed, such as quasi-linear control [5], which offers one solution for a special class of systems. The main objective of the current work is to present a method for implementing CLF-based controllers for a general class of nonlinear systems

in a manner which respects the user-specified input bounds.

IV. Formulating the CLF Min-Norm controller as a Convex Optimization

To design such a controller, we proceed by recognizing that the pointwise min-norm controller in (20) can be equivalently expressed as a convex optimization problem formulated as

$$\begin{aligned} \min_{\mu} \quad & \mu^T \mu \\ \text{s.t.} \quad & \psi_{0,\varepsilon}(\eta, z) + \psi_{1,\varepsilon}(\eta, z)\ \mu \leq 0. \end{aligned} \tag{21}$$

The inequality constraint above enforces the bound on the time-derivative of the CLF given by (15), which can be equivalently expressed as $\dot{V}_\varepsilon(\eta) \leq -c_3/\varepsilon\ V_\varepsilon(\eta)$. The solution of this convex optimization problem is then exactly the controller specified in (20).

Once we have expressed the pointwise min-norm controller as a convex optimization problem, we can introduce bounds on the control input in the form of additional constraints for the convex optimization problem. However, for these additional constraints to be satisfied, we first need to relax the bound on the time-derivative of the CLF. We do this by requiring $\dot{V}_\varepsilon(\eta) \leq -c_3/\varepsilon\ V_\varepsilon(\eta) + d_1$, where d_1 is typically a small positive quantity. The new optimization problem is formulated as

$$\begin{aligned} \min_{\mu, d_1} \quad & \mu^T \mu + p_1\ d_1^2 \\ \text{s.t.} \quad & \psi_{0,\varepsilon}(\eta, z) + \psi_{1,\varepsilon}(\eta, z)\ \mu \leq d_1, \\ & (L_g L_f y(q, \dot{q}))^{-1}\ \mu \geq (u_{min} - u^*), \\ & (L_g L_f y(q, \dot{q}))^{-1}\ \mu \leq (u_{max} - u^*). \end{aligned} \tag{22}$$

where p_1 is a large positive number that represents the penalty of relaxing the inequality constraints and u^* is defined by (6).

The formulation in (22) deals with the non-ideal context of saturated control inputs and therefore cannot ensure the same type of stability claims as those provided by Theorem 2 of [1], since relaxations in the bound on $\dot{V}_\varepsilon$ result in a loss of the RES-CLF quality for V_ε. However, given a prescribed convergence bound and a set of saturation constraints, the control described by (22) is guaranteed to perform at least as well as any other controller in the sense that it will keep V_ε in the smallest possible level set. In this sense, the CLF-based controller (22) can "match" the performance of any other controller in regards to bounding the growth of the RES-CLF V_ε. We also note that, though (22) as formulated does not guarantee Lipschitz continuity of the resultant controller, the work in [12] provides sufficient conditions to ensure Lipschitz continuity for these types of problems. Moreover, the recent result in [17] does not assume Lipschitz continuity.

Remark 1: We note that (22) can also be formulated with "soft" bounds on the control inputs, such that the control input u in (8) satisfies $u_{min} - d_2 \leq u \leq u_{max} + d_3$, with d_2, d_3 typically small positive quantities. This alternative formulation provides the control designer with parameters to trade off violation of the bound on the time-derivative of the CLF with that of the saturation bound on the control input. However, in most practical cases the bounds on the inputs appear as hard bounds which cannot be relaxed, and the current work will focus only on this case.

Remark 2: Note that in (22) we have depicted u_{min} and u_{max} as constants. However, since the convex optimization problem is to be solved at every instant in time, these values can be specified as functions of time or system state, leading to *dynamic torque saturation*. This approach will be explicitly demonstrated in the next section, where one case of our simulation trials involves input bounds which vary as a function of the system state.

Remark 3: In Section II-B we presented an input-output linearizing controller based on PD control, given by (8) with (12). As formulated, the controller has no built-in means for dealing with saturation constraints, but we note that this controller can also be formulated as a convex optimization problem analogous to (22), as

$$\begin{aligned} \min_{u, d_1, d_2, d_3, d_4} \quad & u^T u + \sum_{i=1}^{4} p_i\ d_i^2 \\ \text{s.t.} \quad & L_g L_f h u = -L_f^2 h - \frac{\alpha_0}{\varepsilon} h - \frac{\alpha_1}{\varepsilon^2} L_f h \\ & \qquad + [d_1, d_2, \ldots, d_4]^T, \\ & u \geq u_{min}, \\ & u \leq u_{max}. \end{aligned} \tag{23}$$

However, unlike the CLF-based controller in (22), this formulation does not provide a clear correspondence between the relaxations d_i and performance of the controller. This highlights one of the main advantages of using the QP implementation of the CLF-based controller over the IO controller (either the original implementation or the QP version). Under active saturation constraints, the CLF-based controller relaxes the bound on $\dot{V}_\varepsilon$ just enough to balance the conflicting requirements between performance and saturation constraints. In contrast, the original IO controller ((8) with (12)) "blindly" saturates controls, and the QP version (23) relaxes an equality constraint in a manner that does not clearly correlate to controller performance.

V. Simulation and Experimental Results

In this section we present both numerical simulation and experimental results to validate the performance of the control methods described in Section IV. Since experimental testing on MABEL was the ultimate goal, the numerical studies were conducted first on a simple model of MABEL, followed by simulations on a complex model of MABEL developed in [13], which closely replicates the experimental setup. This latter model includes a compliant ground model as well as a model that allows for stretch in the cables between the transmission pulleys. MABEL is a 5-link bipedal robot with point feet and series-compliant actuation for improved agility and energy efficiency. The experimental setup has been described previously in [14] and is illustrated in Figure 7. For the simulations and experiments described here, the four output functions in (3) were defined by the absolute pitch angle of the torso, the leg angle (*LA*) for the swing leg, and the appropriately scaled leg-shape motor position (*mLS*) for the swing and stance legs. The four control inputs are the leg-angle motor torque ($u_{\mathrm{mLA_{st}}}, u_{\mathrm{mLA_{sw}}}$) and leg-shape motor torque ($u_{\mathrm{mLS_{st}}}, u_{\mathrm{mLS_{sw}}}$) for the stance and swing legs respectively.

A. Numerical simulation

1) CLF-QP controller under various torque bounds: The numerical simulation results presented here employ the CLF-based controller with hard input constraints, as in (22). We consider four separate cases with different control bounds, given by

$$A \quad : \quad \left\{ \begin{bmatrix} -8 \\ -12 \\ -8 \\ -12 \end{bmatrix} \leq \begin{bmatrix} u_{\mathrm{mLA_{st}}} \\ u_{\mathrm{mLS_{st}}} \\ u_{\mathrm{mLA_{sw}}} \\ u_{\mathrm{mLS_{sw}}} \end{bmatrix} \leq \begin{bmatrix} 8 \\ 12 \\ 8 \\ 12 \end{bmatrix} \right.$$

$$B \quad : \quad \left\{ \begin{bmatrix} -5 \\ -8 \\ -2 \\ -2 \end{bmatrix} \leq \begin{bmatrix} u_{\mathrm{mLA_{st}}} \\ u_{\mathrm{mLS_{st}}} \\ u_{\mathrm{mLA_{sw}}} \\ u_{\mathrm{mLS_{sw}}} \end{bmatrix} \leq \begin{bmatrix} 4 \\ 4 \\ 4 \\ 4 \end{bmatrix} \right.$$

$$C \quad : \quad \left\{ \begin{bmatrix} -4 \\ -8 \\ -2 \\ -2 \end{bmatrix} \leq \begin{bmatrix} u_{\mathrm{mLA_{st}}} \\ u_{\mathrm{mLS_{st}}} \\ u_{\mathrm{mLA_{sw}}} \\ u_{\mathrm{mLS_{sw}}} \end{bmatrix} \leq \begin{bmatrix} 1 \\ 1 \\ 1 \\ 1 \end{bmatrix} \right.$$

$$D \quad : \quad \left\{ u^*(\theta) + \begin{bmatrix} -4 \\ -7 \\ -1 \\ -1 \end{bmatrix} \leq \begin{bmatrix} u_{\mathrm{mLA_{st}}} \\ u_{\mathrm{mLS_{st}}} \\ u_{\mathrm{mLA_{sw}}} \\ u_{\mathrm{mLS_{sw}}} \end{bmatrix} \leq u^*(\theta) + \begin{bmatrix} 4 \\ 7 \\ 1 \\ 1 \end{bmatrix} \right.$$

where $u^*(\theta)$ is the nominal value of (6) along the periodic orbit, regressed as a 5th-order Bezier polynomial of $\theta(q)$. Note that in case D, the bounds are specified dynamically as a function of the state of the robot, resulting in dynamic torque saturation.

Simulations of a representative walking step with the controller (22) were run for each of Cases A-D; the corresponding RES-CLF V_ε and its time derivative are presented in Figure 2, and the resulting input torques and tracking errors are illustrated in Figure 3. The saturation effects are most visible in the plots in the first and third rows of the figure; as expected, more restrictive torque limits result in increased tracking error. However, we observe that the degradation in performance is gradual and walking stability is still maintained for all cases (A-D) of input saturation.

To illustrate the effect of saturation on the walking limit cycle, we also carry out simulations on the complex model of MABEL. We use the controller given by (22) in closed-loop and analyze the phase portrait of the torso angle, subject to several different saturation values. Figure 4 illustrates the torso phase portrait for 15 steps of walking, and we observe that stricter saturations result in (gradual) deterioration in tracking, as evidenced by deviations of the limit cycle from the nominal orbit. The saturation values used here differ from those used in the simulations described in the first part of this section (since the complex model differs significantly from the simple model and the required torques for walking are different), but the approach is analogous, with bounds becoming increasingly restrictive proceeding from Case I to Case IV.

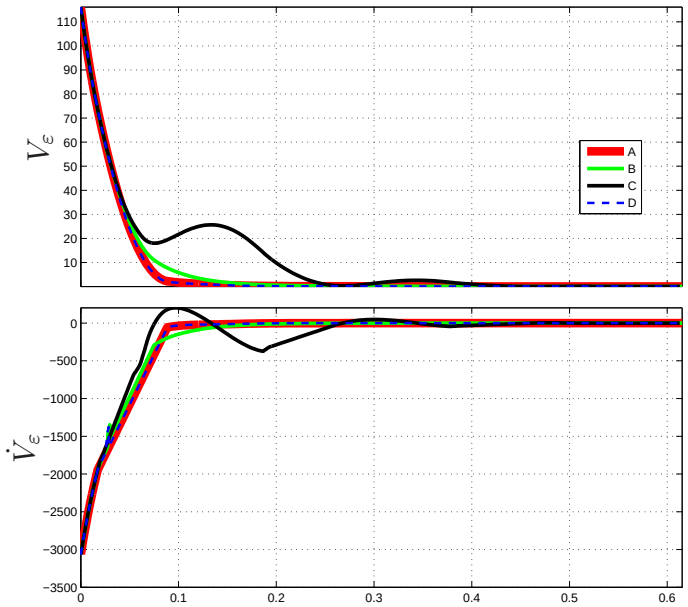

Fig. 2: The RES-CLF V_ε and its derivative for the numerical simulations described in Section V-A. The figures depict the results for four different torque bounds, and we note that walking stability is maintained in all cases.

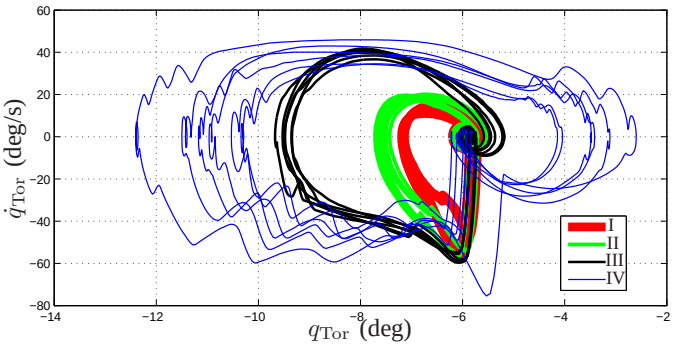

Fig. 4: Phase portrait of the torso angle for four different cases of input bounds: Case I corresponds to $-8 \leq u_{\mathrm{mLA}} \leq 8$, $-12 \leq u_{\mathrm{mLS}} \leq 12$, Case II corresponds to $-6 \leq u_{\mathrm{mLA}} \leq 4$, $-6.5 \leq u_{\mathrm{mLS}} \leq 4$, Case III corresponds to $-4 \leq u_{\mathrm{mLA}} \leq 2$, $-6 \leq u_{\mathrm{mLA}} \leq 3$, and Case IV corresponds to a dynamic torque bound that is a function of the state of the robot. Observe that stricter saturations result in (gradual) deterioration in tracking, as evidenced by deviations of the limit cycle from the nominal orbit.

2) Comparison of CLF-based controller with IO-linearizing PD control: Having demonstrated that the CLF-QP controller is capable of functioning, albeit at degraded performance, under various levels of torque bounds, we will now attempt to compare the four controllers presented in this paper. In this section the controllers are termed as (a) *IO controller* referring to the input-output linearizing controller, (8) with (12); (b) *CLF controller* referring to the CLF-based min-norm controller, (20); (c) *CLF-QP controller* referring to the CLF-based min-norm controller posed as a quadratic program with additional input constraints, (22); and (d) *IO-QP controller*

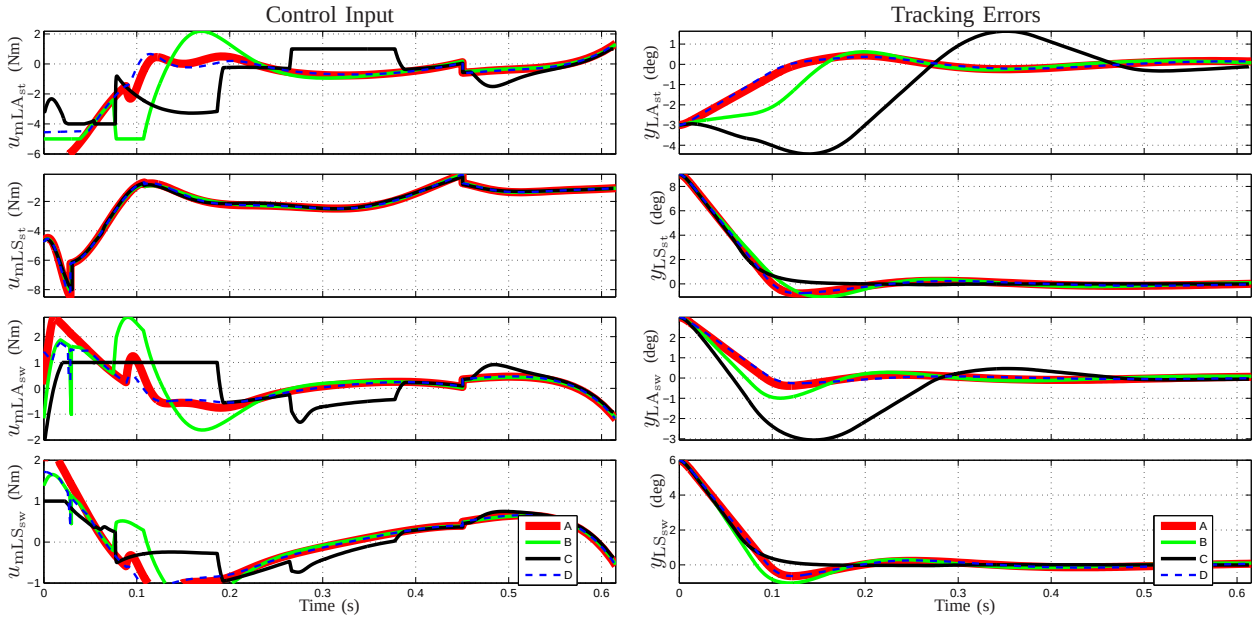

(a) (b)

Fig. 3: (a) Motor torques for the stance (top two figures) and swing legs (bottom two figures), and (b) Corresponding errors in tracking the output (3), based on the numerical simulations described in Section V-A. Each figure depicts the results for four different cases of input bounds. Walking stability is maintained in each case, but we note that the stringent torque bounds in Case C result in control inputs that are only piece-wise continuous. For obtaining Lipschitz continuous control inputs, see additional required conditions in [12].

referring to the input-output controller posed as a quadratic program along with the additional input bound constraints, (23). For each of these controllers, one step of walking is simulated with an initial error and with the restrictive input saturation constraints of case C.

It must be noted that direct comparison of the performance of the CLF controller and IO controllers is difficult and somewhat anecdotal because of the heavy dependance on parameter tuning. We note that for the CLF-QP controller, performance depends on

- selection of the RES-CLF V_ε,
- the relaxation penalty p_1,
- and the parameter ε which dictates the bound on $\dot{V}_\varepsilon$,

while the IO controller is dependent on the selection of ε and the parameters K_P and K_D. For this comparison, we use the same ε for all controllers, however the relaxation penalty for the CLF controller and the PD gains for the IO controller are selected separately. A study of best procedures for tuning the CLF controller and for comparison of controller performance is not the subject of the current work, but presents an interesting field of study for future research.

The controllers are compared in Table I, and graphical results of numerical simulations are presented in Figures 5-6. In the particular simulations at hand, Table I illustrates that under the same conditions, the CLF-QP controller spends the least amount of time having one or more actuators in saturation and also results in the most energy efficient gait, as computed by the specific cost of mechanical transport [15]. However, as noted previously, comparison of controller performance is somewhat anecdotal due to the reliance on paramter tuning and thus the results in Table I should be viewed accordingly. The comparison does suggest that the (non-QP) CLF controller performs the worst under input saturations since the controller has no awareness of saturation constraints, and thus even when the actuators are not in saturation the controller does not act aggressively to reduce the large errors that have built up. Figure 5 illustrates the RES-CLF V_ε and its time-derivative for all the controllers. For the two controllers which do not incorporate knowlege of the input saturations (i.e. the CLF and IO controller) the V_ε grows considerably, although the IO controller is able to quickly decrease the errors once the calculated control torques are within saturation limits. Figure 6b illustrates the tracking errors for the controllers. Note that the CLF controller is unable to control the growing errors and results in instability under these stringent torque bounds.

Controller	% Time in Saturation	C_{mt}
IO (8), (12)	68%	0.021
CLF (20)	91%	0.092
CLF-QP (22)	**23%**	**0.008**
IO-QP (23)	65%	0.016

TABLE I: Comparison between the different types of controllers presented in this paper when under hard input saturation. The second column represents the percentage of time for which one or more actuators are in saturation, and the third column presents the specific cost of mechanical transport.

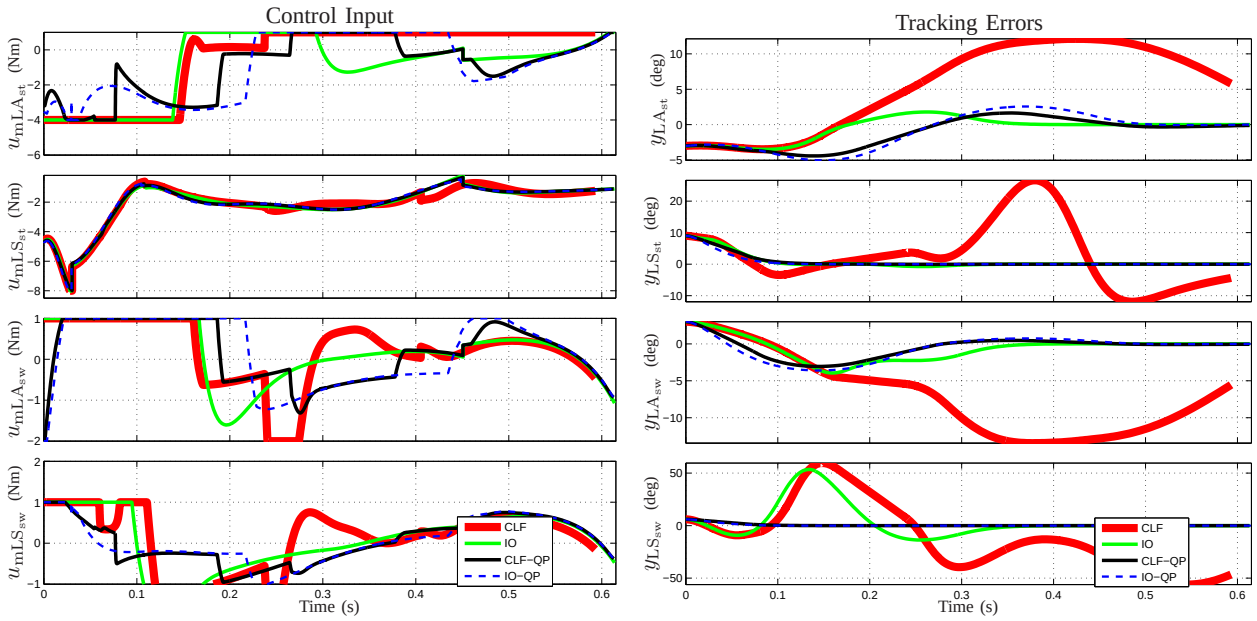

(a) (b)

Fig. 6: (a) Motor torques for the stance (top two figures) and swing legs (bottom two figures), and (b) Corresponding errors in tracking the output (3), based on the numerical simulations described in Section V-A. Each figure depicts the results for the four controllers presented in this paper with hard input saturation. Only the CLF controller leads to instability, while the IO, IO-QP and CLF-QP controllers stabilize to the periodic walking gait.

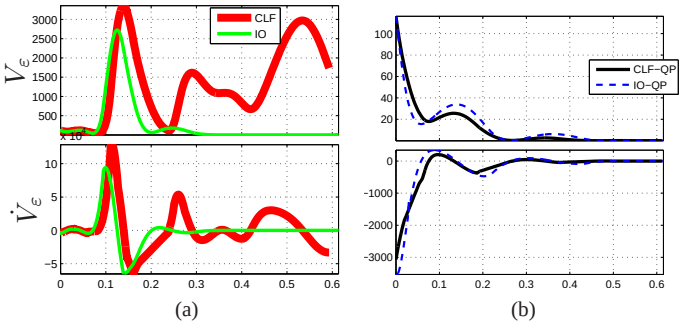

(a) (b)

Fig. 5: The RES-CLF V_ε and its derivative for the numerical simulations described in Section V-A. The figures depict the results for the four controllers presented in the paper while under stringent torque limits.

B. *Experimental results*

Motivated by the favorable numerical simulation results, we proceed to test the controller experimentally on MABEL. Experimental implementation of the CLF controller at real-time speeds is a challenging task, since it requires computation of the system dynamics (2), the Lie derivatives of the output (3), and the CLF controller terms (19), as well as the solving of a convex optimization problem. In order to meet hard real-time constraints of 1 kHz, these computations must be completed in less than 1 ms. By employing the custom-code generation method CVXGEN [11] for solving constrained quadratic programs, we are able to solve the optimization problem in a few hundred microseconds and meet the 1 kHz update requirement, making experimental implementation feasible.

In this experiment, we implemented the CLF controller described in (22), with the CLF-bound penalty set at $p_1 = 50$ and with torque bounds u_{min}, u_{max} chosen such that $-8 \leq u_{\mathrm{mLA}} \leq 8$, $-12 \leq u_{\mathrm{mLS}} \leq 12$. This experiment resulted in 70 steps of walking for MABEL and is portrayed in the video in [8]. (A photo sequence depicting one representative step is also shown in Figure 8.) Figure 9 illustrates the resultant control torques; we observe that the user-specified control bounds are respected, as evidenced by the flattened control signals at the boundary areas. Note that the green squares on the plot depict the time instances at which control bounds are not met, which occur at moments in which the convex optimization algorithm is not able to converge within the specified time constraints. These occurrences are isolated and have no affect on the experimental system since a motor is not able to respond to them. Figure 10 illustrates the Lyapunov function and its time derivative for this experiment. The fact that the Lypanuov function V_ε increases at some points where the calculated $\dot{V}_\varepsilon$ is negative is most likely due to model uncertainty.

VI. Conclusion

We have presented a novel method that explicitly addresses input saturation in the feedback control design for achieving walking in bipedal robots. The resulting controller enables

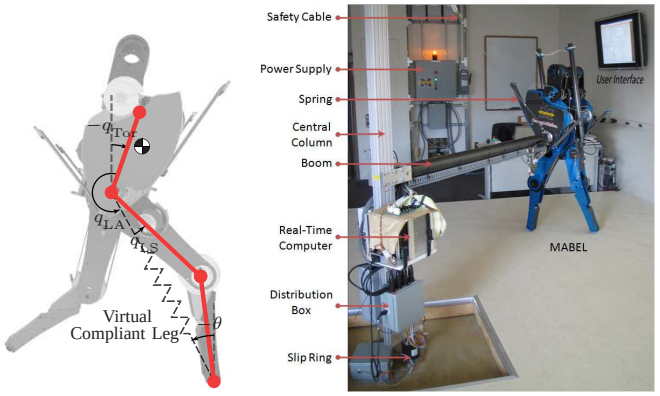

(a) Coordinates (b) MABEL experimental setup

Fig. 7: Experimental setup for bipedal robot MABEL and associated coordinates. (From [14].)

Fig. 8: A photo sequence depicting one representative step from the second experiment described in Section V-B.

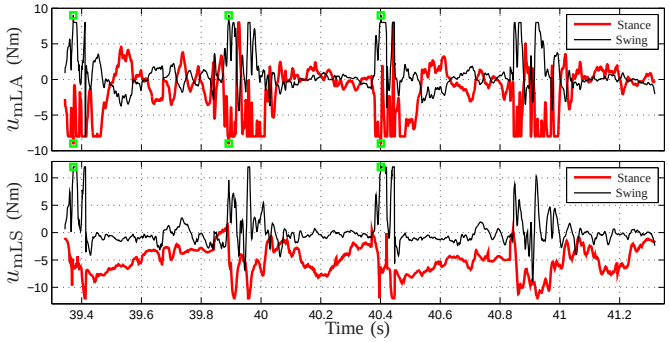

Fig. 9: Motor torques for the stance and swing legs for 4 consecutive steps of walking with the CLF controller with convex optimization and strict torque limits. The convex optimization was tasked to enforce the magnitude of the LA and LS motor torques to be within 8 Nm and 12 Nm respectively. The green square markers on the plots indicate isolated time instances at which the user-specified torque bound was exceeded by the convex optimization. This occurs when the convex optimization fails to converge within the maximum number of allowed iterations, a limit required to ensure the hard real-time constraints are met for experimental implementation.

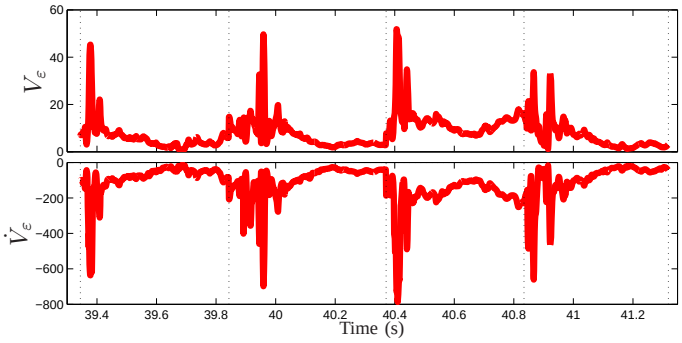

Fig. 10: The RES-CLF V_ε and its time-derivative for 4 consecutive steps of walking with the CLF controller with convex optimization and hard constraints on saturations. Note that the time derivative of V_ε is computed from the experimental data on the best model of the system we have. There are instance in this plot when $\dot{V}_\varepsilon$ is negative while V_ε is increasing, which is most likely due to model uncertainty.

gradual performance degradation while still continuing to walk under a range of stringent torque limits. We accomplish this through an alternative method for implementing the point-wise min-norm CLF-based controller described in (20) in a manner that more appropriately handles input saturations. Numerical simulation as well as experimental implementation has demonstrated that these control methods can be very useful in practice, even in systems which require a high real-time control update rate. This method has great potential for effectively dealing with saturations in a variety of contexts, such as power-limited systems which could progressively lower user-defined torque saturations as the battery charge decreases, thereby prolonging the last bit of battery charge while allowing system performance to gracefully degrade. In addition to dynamic torque saturation, we also note that this approach provides a method for incorporating a whole family of user-defined constraints into the online calculation of controller effort for the types of systems described here. Future work will consider the effects of varying ε throughout the gait, which may result in an improved trade-off between convergence rate and saturation response over the course of the step.

References

[1] Aaron D. Ames, Kevin Galloway, J. W. Grizzle, and Koushil Sreenath. Rapidly Exponentially Stabilizing Control Lyapunov Functions and Hybrid Zero Dynamics. *To appear in IEEE Trans. Automatic Control*, 2014.

[2] Dennis Bernstein and Anthony Michel. A chronological bibliography on saturating actuators. *International Journal of Robust and Nonlinear Control*, 5:375–380, 1995.

[3] David Braun and Michael Goldfarb. A control approach for actuated dynamic walking in biped robots. *IEEE Transactions on Robotics*, 25:1292–1303, 2009.

[4] C. Chevallereau, A. Formal'sky, and B. Perrin. Low energy cost reference trajectories for a biped robot. *Proc. IEEE Intl. Conf. on Robotics and Automation*, pages 1398–1404, 1998.

[5] ShiNung Ching, Yongsoon Eun, Cevat Gokcek, Pierre T. Kabamba, and Semyon M. Meerkov. *Quasilinear Control: Performance Analysis and Design of Feedback Systems with Nonlinear Sensors and Actuators*. Cambridge University Press, 2011.

[6] Dimitar Dimitrov, Alexander Sherikov, and Pierre-Brice Wieber. A sparse model predictive control formulation for walking motion generation. *Proc. IEEE Intl. Conf. on Intelligent Robots and Systems*, pages 2292–2299, 2011.

[7] R. A. Freeman and P. V. Kokotović. *Robust Nonlinear Control Design*. Birkhäuser, 1996.

[8] K. Galloway, K. Sreenath, A. Ames, and J. Grizzle. Walking with CLF Control + Dynamic Torque Saturation (2013). Youtube Video. [Online]. Available: http://youtu.be/rc1FSXpfrrM.

[9] Michael Hardt, Kenneth Kreutz-Delgado, and J. William Helton. Optimal biped walking with a complete dynamical model. *Proc. 38th IEEE Conf. Decision and Control*, pages 2999–3004, 1999.

[10] A. Isidori. *Nonlinear Control Systems: An Introduction*. Springer-Verlag, Berlin, Germany, 2nd edition, 1989.

[11] Jacob Mattingley and Stephen Boyd. CVXGEN: a code generator for embedded convex optimization. *Optimization and Engineering*, 13(1):1–27, March 2012.

[12] Benjamin Morris, Matthew J. Powell, and Aaron D. Ames. Sufficient conditions for the Lipschitz continuity of QP-based multi-objective control of humanoid robots. *Proc. 52nd IEEE Conf. Decision and Control*, pages 2920–2926, 2013.

[13] Hae-Won Park, Koushil Sreenath, Jonathan Hurst, and Jessy W. Grizzle. Identification of a bipedal robot with a compliant drivetrain: Parameter estimation for control design. *IEEE Control Systems Magazine (CSM)*, 31(2):63–88, April 2011.

[14] K. Sreenath, H.W. Park, I. Poulakakis, and J. W. Grizzle. A compliant hybrid zero dynamics controller for stable, efficient and fast bipedal walking on MABEL. *IJRR*, 30(9):1170–1193, 2011.

[15] Koushil Sreenath, Hae-Won Park, Ioannis Poulakakis, and Jessy W. Grizzle. Compliant hybrid zero dynamics controller for achieving stable, efficient and fast bipedal walking on MABEL. *International Journal of Robotics Research*, 30(9):1170–1193, August 2011.

[16] T. Sugihara and Y. Nakamura. Whole-body cooperative balancing of humanoid robot using COG jacobian. *Proc. IEEE/RSJ Int. Conf. Intell. Robots and Systems*, 3:2575–2580, 2002.

[17] A. R. Teel, R. Goebel, B. Morris, A. D. Ames, and J. W. Grizzle. A stabilization result with application to bipedal locomotion. *Proc. 52nd IEEE Conf. Decision and Control*, pages 2030–2035, 2013.

[18] E. Westervelt, J.W. Grizzle, and D.E. Koditschek. Hybrid zero dynamics of planar biped walkers. *IEEE Transactions on Automatic Control*, 48(1):42–56, January 2003.

[19] E. R. Westervelt, J. W. Grizzle, C. Chevallereau, J. H. Choi, and B. Morris. *Feedback Control of Dynamic Bipedal Robot Locomotion*. CRC Press, Boca Raton, 2007.

[20] P-B. Wieber. Trajectory free linear model predictive control for stable walking in the presence of strong perturbations. *Proc. 6th IEEE-RAS Intl. Conf. on Humanoid Robots*, pages 137–142, 2006.