ISTITUTO ITALIANO
DI TECNOLOGIA
DYNAMIC LEGGED SYSTEMS

STANCE: Locomotion Adaptation over Soft Terrain

Shamel Fahmi, Michele Focchi, Andreea Radulescu, Geoff Fink,
Victor Barasuol, and Claudio Semini

Accepted on: October 03, 2019.
Published in: IEEE Transactions on Robotics (T-RO), 2019
To cite this paper:

S. Fahmi, M. Focchi, A. Radulescu, G. Fink, V. Barasuol, C. Semini, "STANCE: Locomotion Adaptation over Soft Terrain," *IEEE Trans. Robot.*, 2019.

For this and other publications from the Dynamic Legged Systems lab (DLS):
`https://dls.iit.it/dls-publications`

©2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

STANCE: Locomotion Adaptation over Soft Terrain

Shamel Fahmi, Michele Focchi, Andreea Radulescu, Geoff Fink, Victor Barasuol, and Claudio Semini

***Abstract*—Whole-Body Control (WBC) has emerged as an important framework in locomotion control for legged robots. However, most WBC frameworks fail to generalize beyond rigid terrains. Legged locomotion over soft terrain is difficult due to the presence of unmodeled contact dynamics that WBCs do not account for. This introduces uncertainty in locomotion and affects the stability and performance of the system. In this paper, we propose a novel soft terrain adaptation algorithm called STANCE: Soft Terrain Adaptation and Compliance Estimation. STANCE consists of a WBC that exploits the knowledge of the terrain to generate an optimal solution that is contact consistent and an online terrain compliance estimator that provides the WBC with terrain knowledge. We validated STANCE both in simulation and experiment on the Hydraulically actuated Quadruped (HyQ) robot, and we compared it against the state of the art WBC. We demonstrated the capabilities of STANCE with multiple terrains of different compliances, aggressive maneuvers, different forward velocities, and external disturbances. STANCE allowed HyQ to adapt online to terrains with different compliances (rigid and soft) without pre-tuning. HyQ was able to successfully deal with the transition between different terrains and showed the ability to differentiate between compliances under each foot.**

***Index Terms*—Whole-Body Control, Legged Robots, Compliance and Impedance Control, Optimization and Optimal Control**

I. Introduction

WHOLE-BODY CONTROL (WBC) frameworks have achieved remarkable results in legged locomotion control [1, 2, 3]. Their main feature is that they use optimization techniques to solve the locomotion control problem. WBC can achieve multiple tasks in an optimal fashion by exploiting the robot's full dynamics and reasoning about both the actuation constraints and the contact interaction. These tasks include balancing, interacting with the environment, and performing dynamic locomotion over a wide variety of terrains [3]. The tasks are executed at the robot's end effectors, but can also be utilized for contacts anywhere on the robot's body [4] or for a cooperative manipulation task between robots [5].

To date, most of the work done on WBC assumes that the ground is rigid (i.e., rigid contact consistent). However, if the robot traverses soft terrain (as shown in Fig. 1), the mismatch between the rigid assumption and the soft contact interaction can significantly affect the robot's performance and locomotion stability. This mismatch is due to the unmodeled contact dynamics between the robot and the terrain. In fact, under the rigid ground assumption, the controller can generate instantaneous changes to the Ground Reaction Forces (GRFs).

Manuscript received: April 26, 2019; Revised August 22, 2019; Accepted October 2, 2019. This paper was recommended for publication by Associate Editor X. XXXX and Editor X. XXXX upon evaluation of the reviewers' comments. (*Corresponding author: Shamel Fahmi.*)

All the authors are with the Dynamic Legged Systems Lab, Istituto Italiano di Tecnologia (IIT), Genova, Italy (email: firstname.lastname@iit.it).

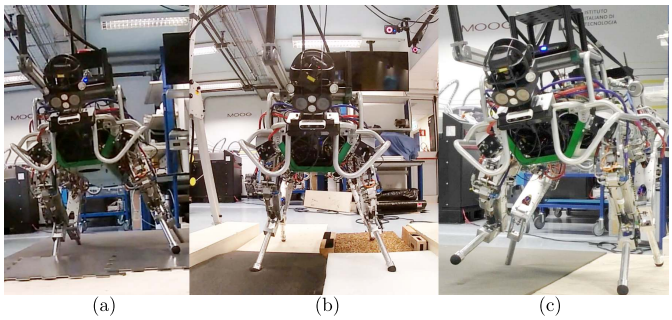

(a) (b) (c)

Fig. 1: HyQ traversing multiple terrains of different compliances.

This is equivalent to thinking that the terrain will respond with an infinite bandwidth.

In order to robustly traverse a wide variety of terrains of different compliances, the WBC must become *compliant contact consistent (c^3)*. Namely, the WBC should be terrain-aware. That said, a more general WBC approach should be developed that can adapt *online* to the changes in the terrain compliance.

A. Related Work - Soft Terrain Adaptation for Legged Robots

Locomotion over soft terrain can be tackled either from a control or a planning perspective. In the context of locomotion control, Henze *et al.* [6] presented the first experimental attempt using a WBC over soft terrain. Their WBC is based on the rigid ground assumption, but it allows for constraint relaxation. This allowed the humanoid robot TORO to adapt to a compliant surface. Their approach was further extended in [7] by dropping the rigid contact assumption and using an energy-tank approach. Despite balancing on compliant terrain, both approaches were only tested for one type of soft terrain when the robot was standing still.

Similarly, other works explicitly adapt to soft terrain by incorporating terrain knowledge (i.e., contact model) into their balancing controllers. For example, Azad *et al.* [8] proposed a momentum based controller for balancing on soft terrain by relying on a nonlinear soft contact model. Vasilopoulos *et al.* [9] proposed a similar hopping controller that models the terrain using a viscoplastic contact model. However, these approaches were only tested in simulation and for monopods.

In the context of locomotion planning, Grandia *et al.* [10] indirectly adapted to soft terrain by shaping the frequency of the cost function of their Model Predictive Control (MPC) formulation. By penalizing high frequencies, they generated optimal motion plans that respect the bandwidth limitations due to soft terrain. This approach was tested over three types of terrain compliances. However, it was not tested during transitions from one terrain to another. This approach showed

an improvement in the performance of the quadruped robot in simulation and experiment. However, the authors did not offer the possibility to change their tuning parameters online. Thus, they were not able to adapt the locomotion strategy based on the compliance of the terrain.

In contrast to the aforementioned work, other approaches relax the rigid ground assumption (hard contact constraint) but not for soft terrain adaptation purposes. For instance, Kim *et al.* [11] implemented an approach to handle sudden changes in the rigid contact interaction. This approach relaxed the hard contact assumption in their WBC formulation by penalizing the contact interaction in the cost function rather than incorporating it as a hard constraint. For computational purposes, Neunert *et al.* [12] and Doshi *et al.* [13] proposed relaxing the rigid ground assumption. Neunert *et al.* used a soft contact model in their nonlinear MPC formulation to provide smooth gradients of the contact dynamics to be more efficiently solved by their gradient based solver. The soft contact model did not have a physical meaning and the contact parameters were empirically chosen. Doshi *et al.* proposed a similar approach which incorporates a slack variable that expands the feasibility region of the hard constraint.

Despite the improvement in performance of the legged robots over soft terrain in the aforementioned works, none of them offered the possibility to adapt to the terrain *online*. Most of the aforementioned works lack a general approach that can deal with multiple terrain compliances or with transitions between them. Perhaps, one noticeable work (to date) in online soft terrain adaptation was proposed by Chang *et al.* [14]. In that work, an iterative soft terrain adaptation approach was proposed. The approach relies on a non-parametric contact model that is simultaneously updated alongside an optimization based hopping controller. The approach was capable of iteratively learning the terrain interaction and supplying that knowledge to the optimal controller. However, because the learning module was exploiting Gaussian process regression, which is computationally expensive, the approach did not reach real-time performance and was only tested in simulation, for one leg, under one experimental condition (one terrain).

B. Related Work - Contact Compliance Estimation in Robotics

For contact compliance estimation, we need to accurately model the contact dynamics and estimate the contact parameters online. In contact modeling, Alves *et al.* [15] presented a detailed overview of the types of parametric soft contact models used in the literature. In compliance estimation, Schindeler *et al.* [16] used a two stage polynomial identification approach to estimate the parameters of the Hunt and Crossley's (HC) contact model online. Differently, Azad *et al.* [17] used a Least Square (LS)-based estimation algorithm and compared multiple contact models (including the Kelvin-Voigt's (KV) and the HC models). Other approaches that are not based on soft contact models use force observers [18] or neural networks [19]. These aforementioned approaches in compliance estimation were designed for robotic manipulation tasks.

To date, the only work on compliance estimation in legged locomotion was the one by Bosworth *et al.* [20]. The authors presented two online (in-situ) approaches to estimate the ground properties (stiffness and friction). The results were promising and the approaches were validated on a quadruped robot while hopping over rigid and soft terrain. However, the estimated stiffness showed a trend, but was not accurate; the lab measurements of the terrain stiffness did not match the in-situ ones. Although the estimation algorithms could be implemented online, the robot had to stop to perform the estimation.

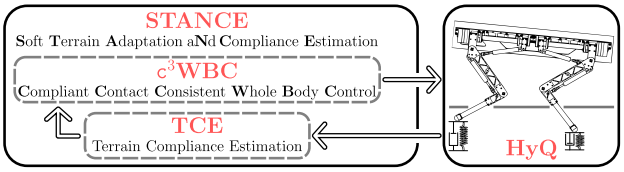

Fig. 2: An overview of the STANCE algorithm.

C. Proposed Approach and Contribution

In this work, we propose an online soft terrain adaptation algorithm: **S**oft **T**errain **A**daptation a**N**d **C**ompliance **E**stimation (STANCE). As shown in Fig. 2, STANCE consists of

- A Compliant Contact Consistent Whole-Body Control (c^3WBC) that is contact consistent to any type of terrain *given the terrain compliance*. This is done by extending the state-of-the-art WBC in [3], hereafter denoted as the Standard Whole-Body Control (sWBC). In particular, c^3WBC incorporates a soft contact model into the WBC formulation.
- A Terrain Compliance Estimator (TCE) which is an online learning algorithm that provides the c^3WBC with an estimate of the terrain compliance. It is based on the same contact model that is incorporated in the c^3WBC.

The main contribution of STANCE is that it can adapt to any type of terrain (stiff or soft) online without pre-tuning. This is done by closing the loop of the c^3WBC with the TCE. To our knowledge, this is the first implementation of such an approach in legged locomotion.

STANCE is meant to overcome the limitations of the aforementioned approaches in soft terrain adaptation for legged robots. Compared to previous works on WBC that tested their approach only during standing [6, 7], we test our STANCE approach during locomotion. Compared to other approaches [8, 9] that were tested on monopods in simulation, STANCE is implemented and tested in experiment on HyQ. Compared to previous work on soft terrain adaptation [10], STANCE can adapt to soft terrain *online* and was tested on multiple terrains with different compliances and with transitions between them. Compared to [14], our TCE is computationally inexpensive, which allows STANCE to run real-time in experiments and simulations. Compared to the previous work done on compliance estimation, we implemented our TCE on a legged robot which is, to the best of our knowledge, the first experimental validation of this approach. Differently from [20], our TCE approach could be implemented in parallel with any gait or task. We also achieved a more accurate estimation of the terrain compliance compared to [20].

As additional contributions, we discussed the benefits (and the limitations) of exploiting the knowledge of the terrain in WBC based on the experience gained during extensive experimental trials. To our knowledge, STANCE is the first work to present legged locomotion experiments crossing multiple terrains of different compliances.

II. Robot model

Consider a legged robot with n Degrees of Freedom (DoFs) and c feet. The total dimension of the feet operational space n_f can be separated into stance ($n_{\rm st} = 3c_{\rm st}$) and swing feet ($n_{\rm sw} = 3c_{\rm sw}$) where $c_{\rm st}$ and $c_{\rm sw}$ are the number of stance and swing legs respectively. Assuming that all external forces are exerted on the stance feet, the robot dynamics is written as

$$\underbrace{\begin{bmatrix} M_{\rm com} & 0_{3\times3} & 0_{3\times n} \\ 0_{3\times3} & M_\theta & M_{\theta j} \\ 0_{n\times3} & M_{\theta j}^T & M_j \end{bmatrix}}_{M(q)} \underbrace{\begin{bmatrix} \ddot{x}_{\rm com} \\ \dot{\omega}_b \\ \ddot{q}_j \end{bmatrix}}_{\ddot{q}} + \underbrace{\begin{bmatrix} h_{\rm com} \\ h_\theta \\ h_j \end{bmatrix}}_{h(q,\dot{q})} = \begin{bmatrix} 0_{3\times1} \\ 0_{3\times1} \\ \tau_j \end{bmatrix} + \underbrace{\begin{bmatrix} J_{{\rm st},{\rm com}}^T \\ J_{{\rm st},\theta}^T \\ J_{{\rm st},j}^T \end{bmatrix}}_{J_{\rm st}(q)^T} F_{\rm grf} \tag{1}$$

where $q \in SE(3) \times \mathbb{R}^n$ denotes the generalized robot states consisting of the Center of Mass (CoM) position $x_{\rm com} \in \mathbb{R}^3$, the base orientation $R_b \in SO(3)$, and the joint positions $q_j \in \mathbb{R}^n$. The vector $\dot{q} = [\dot{x}_{\rm com}^T \ \omega_b^T \ \dot{q}_j^T]^T \in \mathbb{R}^{6+n}$ denotes the generalized velocities consisting of the velocity of the CoM $\dot{x}_{\rm com} \in \mathbb{R}^3$, the angular velocity of the base $\omega_b \in \mathbb{R}^3$, and the joint velocities $\dot{q}_j \in \mathbb{R}^n$. The vector $\ddot{q} = [\ddot{x}_{\rm com}^T \ \dot{\omega}_b^T \ \ddot{q}_j^T]^T \in \mathbb{R}^{6+n}$ denotes the corresponding generalized accelerations. All Cartesian vectors are expressed in the world frame Ψ_W unless mentioned otherwise. $M \in \mathbb{R}^{(6+n)\times(6+n)}$ is the inertia matrix. $h \in \mathbb{R}^{6+n}$ is the force vector that accounts for Coriolis, centrifugal, and gravitational forces. $\tau_j \in \mathbb{R}^n$ are the actuated joint torques, $F_{\rm grf} \in \mathbb{R}^{n_{\rm st}}$ is the vector of GRFs (contact forces). The Jacobian matrix $J \in \mathbb{R}^{n_{\rm f}\times(6+n)}$ is separated into swing Jacobian $J_{\rm sw} \in \mathbb{R}^{n_{\rm sw}\times(6+n)}$ and stance Jacobian $J_{\rm st} \in \mathbb{R}^{n_{\rm st}\times(6+n)}$ which can be further expanded into $J_{{\rm st},{\rm com}} \in \mathbb{R}^{n_{\rm st}\times3}$, $J_{{\rm st},\theta} \in \mathbb{R}^{n_{\rm st}\times3}$, and $J_{{\rm st},j} \in \mathbb{R}^{n_{\rm st}\times n}$. The feet velocities $v \in \mathbb{R}^{n_{\rm f}}$ are separated into stance $v_{\rm st} \in \mathbb{R}^{n_{\rm st}}$ and swing $v_{\rm sw} \in \mathbb{R}^{n_{\rm sw}}$ feet velocities. Similarly, the feet accelerations $\dot{v} \in \mathbb{R}^{n_{\rm f}}$ are separated into stance $\dot{v}_{\rm st} \in \mathbb{R}^{n_{\rm st}}$ and swing $\dot{v}_{\rm sw} \in \mathbb{R}^{n_{\rm sw}}$ feet accelerations. The feet forces $F = [F_{\rm st}^T \ F_{\rm sw}^T]^T \in \mathbb{R}^{n_{\rm f}}$ are also separated into stance $F_{\rm st} \in \mathbb{R}^{n_{\rm st}}$ and swing $F_{\rm sw} \in \mathbb{R}^{n_{\rm sw}}$ feet forces. We split the robot dynamics (1) into an unactuated floating base part (the first 6 rows) and an actuated part (the remaining n rows) as

$$M_u(q)\ddot{q} + h_u(q,\dot{q}) = J_{{\rm st},u}(q)^T F_{\rm grf} \tag{2a}$$
$$M_a(q)\ddot{q} + h_j(q,\dot{q}) = \tau_j + J_{{\rm st},j}(q)^T F_{\rm grf} \tag{2b}$$

where $M_u \in \mathbb{R}^{6\times6+n}$ and $M_a \in \mathbb{R}^{n\times6+n}$ are sub matrices of M, $h_u = [h_{\rm com}^T \ h_\theta^T]^T \in \mathbb{R}^6$ and $h_j \in \mathbb{R}^n$ are sub vectors of h, and $J_{{\rm st},u} = [J_{{\rm st},{\rm com}}^T \ J_{{\rm st},\theta}^T]^T$. Finally, we define the gravito-inertial wrench as $W_{\rm com} = M_u(q)\ddot{q} + h_u(q,\dot{q}) \in \mathbb{R}^6$.

III. Standard Whole-Body Controller (sWBC)

This section summarizes the sWBC as detailed in [3]. Besides the WBC, our locomotion framework includes a locomotion planner, state estimator and a low-level torque controller as shown in Fig. 3. Given high-level user inputs, the planner generates the desired trajectories for the CoM, trunk orientation and swing legs, and provides them to the WBC. The state estimation provides the WBC with the estimated states of the robot.

The objective of the sWBC is to ensure the execution of the trajectories provided by the planner while keeping the robot balanced and reasoning about the robot's dynamics, actuation limits and the contact constraints [3]. We denote the execution of the trajectories provided by the planner as *control tasks*. These control tasks alongside the aforementioned constraints define the WBC problem. The control problem is casted as a Whole-Body Optimization (WBOpt) problem via a Quadratic Program (QP) which solves for the optimal generalized accelerations and contact forces at each iteration of the control loop [21]. The optimal solution of the WBC is then mapped into joint torques that are sent to the low-level torque controller.

A. Control Tasks

We categorize the sWBC *control tasks* into: 1) a *trunk task* that tracks the desired trajectories of the CoM position and trunk orientation, and 2) a *swing task* that tracks the swing feet trajectories [3]. Similar to a PD+ controller [22], both tasks are achieved by a Cartesian-based impedance controller with a feed-forward term. The feedforward terms are added in order to improve the tracking performance of the tasks when following the trajectories from the planner [6, 23]. The tracking of the trunk task is obtained by the desired wrench at the CoM $W_{\rm com,d} \in \mathbb{R}^6$. This is generated by a Cartesian impedance at the CoM, a gravity compensation term, and a feed-forward term. Similarly, the tracking of the swing task can be obtained by the virtual force $F_{\rm sw,d} \in \mathbb{R}^{n_{\rm sw}}$. This is generated by a Cartesian impedance at the swing foot and a feed-forward term. As in [3], we can also write the swing task at the acceleration level by defining the desired swing feet velocities $\dot{v}_{\rm sw,d} \in \mathbb{R}^{n_{\rm sw}}$ as

$$\dot{v}_{\rm sw,d} = \dot{v}_{\rm sw,ff} + {\rm K}_{\rm sw}\Delta x_{\rm sw} + {\rm D}_{\rm sw}\Delta v_{\rm sw} \tag{3}$$

where ${\rm K}_{\rm sw}, {\rm D}_{\rm sw} \in \mathbb{R}^{n_{\rm sw}\times n_{\rm sw}}$ are positive definite PD gains, $\Delta x_{\rm sw} = x_{\rm sw,d} - x_{\rm sw} \in \mathbb{R}^{n_{\rm sw}}$ and $\Delta v_{\rm sw} = v_{\rm sw,d} - v_{\rm sw} \in \mathbb{R}^{n_{\rm sw}}$ are tracking errors of the swing foot position and velocity, respectively, and $\dot{v}_{\rm sw,ff}$ is a feed-forward term.

B. Whole-Body Optimization

To accomplish the sWBC objective (the control tasks in Section III-A and constraints), we formulate the WBOpt problem presented in Formulation 1 and detailed in [3].

1) Decision Variables: As shown in Formulation 1, we choose the generalized accelerations $\ddot{q}$ and the contact forces $F_{\rm grf}$ as the decision variables $u = [\ddot{q}^T \ F_{\rm grf}^T]^T \in \mathbb{R}^{6+n+n_{\rm st}}$. Later in this subsection, we will augment the vector of decision variables with a slack term $\eta \in \mathbb{R}^{n_{\rm sw}}$.

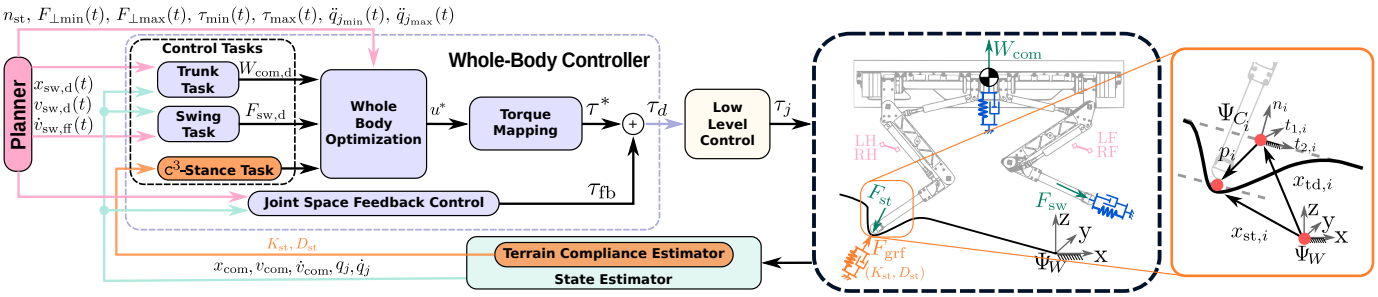

Fig. 3: Overview of the WBC in our locomotion framework. The dashed black box presents the definition of the HyQ's legs (Left-Front (LF), Right-Front (RF), Left-Hind (LH) and Right-Hind (RH)) and the generated wrenches. The solid orange box presents the terminologies used for the soft contact model. Ψ_{W} is the world frame and Ψ_{C_i} is the local contact frame for the leg i fixed at the touch down position.

2) Cost: The cost function (5) consists of two terms. The first term ensures the tracking of the trunk task by minimizing the two-norm of the tracking error between the actual W_{com} and desired $W_{\mathrm{com,d}}$ CoM wrenches. The second term in (5) regularizes the solution and penalizes the slack variable.

3) Physical Consistency: The equality constraint (6) enforces the physical consistency between F_{grf} and $\ddot{q}$ by ensuring that the contact wrenches due to F_{grf} will sum up to W_{com}. This is done by imposing the unactuated dynamics (2a) as an equality constraint.

4) Stance Task: To remain *contact consistent*, we incorporate the *stance task* that enforces the stance legs to remain in contact with the terrain. Since the sWBC is assuming a rigid terrain, the stance feet are forced to remain stationary in the world frame, i.e., $v_{\mathrm{st}} = \dot{v}_{\mathrm{st}} = 0$ (see [3]). As a result, we incorporate the rigid contact model in the sWBC formulation as an equality constraint at the acceleration level (7) in order to have a direct dependency on the decision variables. In detail, since $v_{\mathrm{st}} = J_{\mathrm{st}}\dot{q} = 0$, differentiating once with respect to time yields $\dot{v}_{\mathrm{st}} = J_{\mathrm{st}}\ddot{q} + \dot{J}_{\mathrm{st}}\dot{q} = 0$.

5) Friction and Normal Contact Force: The inequality constraint (11) enforces the friction constraints by ensuring that the contact forces lie inside the friction cones. This is done by limiting the tangential component of the GRFs $F_{\mathrm{grf},\parallel}$. The inequality constraint (12) enforces constraints on the normal component of the GRFs $F_{\mathrm{grf},\perp}$. This includes the unilaterality constraints which encodes that the legs can only push on the ground by setting an "almost-zero" lower bound $F_{\min}$ to $F_{\mathrm{grf},\perp}$. They also allow a smooth loading/unloading of the legs, and set a varying upper bound $F_{\max}$ to $F_{\mathrm{grf},\perp}$. For the detailed implementation of the inequality constraints (11) and (12), refer to [23].

6) Swing Task: We implement the tracking of the swing task (in Section III-A) at the acceleration level (3) rather than the force level since we can express the swing feet velocities $\dot{v}_{\mathrm{sw}}$ as a function of $\ddot{q}$ which is a decision variable, i.e., $\dot{v}_{\mathrm{sw}}(q) = J_{\mathrm{sw}}\ddot{q} + \dot{J}_{\mathrm{sw}}\dot{q}$. This task could be encoded as an equality constraint $\dot{v}_{\mathrm{sw}} = \dot{v}_{\mathrm{sw,d}}$. Yet, it is important to relax this hard constraint when the joint kinematic limits are reached (see [3]). Hence, the swing task is encoded in (13) by an inequality constraint that bounds the solution around the original hard constraint and a slack term η that is penalized for its non-zero values in the cost function (5) and is constrained to remain non-negative in (13).

Formulation 1 Whole-Body Optimization: sWBC Vs. c³WBC

(Trunk Task)	$\min_{u} \ \|W_{\mathrm{com}} - W_{\mathrm{com,d}}\|^2_Q + \|\boldsymbol{u}\|^2_{\boldsymbol{R}}$	(5)
(Decision Variables)	$u = [\ddot{q}^T \ F_{\mathrm{grf}}^T \ \eta^T \ \boldsymbol{\epsilon}^T]^T$	
	s.t.:	
(Physical Consistency)	$M_u\ddot{q} + h_u = J_{\mathrm{st},u}^T F_{\mathrm{grf}}$	(6)
~~(Stance Task)~~	~~$\dot{v}_{\mathrm{st}} = J_{\mathrm{st}}\ddot{q} + \dot{J}_{\mathrm{st}}\dot{q} = 0$~~	(7)
(c³-Stance Task)	$\boldsymbol{F_{\mathrm{grf}} = K_{\mathrm{st}}\epsilon + D_{\mathrm{st}}\dot{\epsilon}}$	(8)
	$\boldsymbol{\dot{v}_{\mathrm{st}} = J_{\mathrm{st}}\ddot{q} + \dot{J}_{\mathrm{st}}\dot{q} = -\ddot{\epsilon}}$	(9)
	$\boldsymbol{\epsilon \geq 0}$	(10)
(Friction)	$\|F_{\mathrm{grf},\parallel}\| \leq \mu\|F_{\mathrm{grf},\perp}\|$	(11)
(Normal Contact Force)	$F_{\min} \leq F_{\mathrm{grf},\perp} \leq F_{\max}$	(12)
(Swing Task)	$-\eta \leq \dot{v}_{\mathrm{sw}} - \dot{v}_{\mathrm{sw,d}} \leq \eta, \ \eta \geq 0$	(13)
(Torque Limits)	$\tau_{\min} \leq \tau_j \leq \tau_{\max}$	(14)
(Joint Limits)	$\ddot{q}_{j_{\min}} \leq \ddot{q}_j \leq \ddot{q}_{j_{\max}}$	(15)

7) Torque and Joint Limits: The torque and joint limits are enforced in the inequality constraints (14) and (15), respectively.

8) Torque Mapping: The WBOpt (5)-(7), (11)-(15) generates optimal joint accelerations $\ddot{q}_j^*$ and contact forces F_{grf}^*, that are mapped into optimal joint torques τ^* and sent to the low-level controller using the actuated dynamics (2b) as

$$\tau^* = M_a\ddot{q}^* + h_j - J_{\mathrm{st},j}^T F_{\mathrm{grf}}^* \tag{4}$$

C. Feedback Control

The computation of the optimal torques τ^* relies on the inverse dynamics in (4) which might be prone to model inaccuracies [24]. In order to tackle this issue, the desired torques τ_d sent to the lower level control could combine the optimal torques τ^* in (4) with a feedback controller τ_{fb} as shown in Fig. 3. The feedback controller improves the tracking performance if the dynamic model of the robot becomes less accurate [24]. The feedback controller is a proportional-derivative (PD) joint space impedance controller [23].

Remark 1: Throughout this work and similar to [3] and [21], we found it sufficient to use only the inverse-dynamics term

(the optimal torques τ^*) and not the joint feedback part. This is due to the fact that we can identify the parameters of our dynamic model with sufficient accuracy as detailed in [25]. That said, we carried out the simulation and experiment without any need of the feedback loop.

IV. C^3 Whole Body Controller

Over soft terrain, the feet positions are non-stationary and are allowed to deform the terrain. Thus, the rigid contact assumption of the stance task (7) in the sWBC does not hold anymore and should be dropped. To be c^3, the interaction between the stance feet and the soft terrain must be governed not just by the robot dynamics but also by the soft contact dynamics. That said, the c^{3}WBC extends the sWBC by: 1) modeling the soft contact dynamics and incorporating it as a *stance task* similar to the control tasks in Section III-A, and 2) encoding the stance task in the WBOpt as a function of the decision variables. The differences between the sWBC and the c^{3}WBC are highlighted in **boldface** in Formulation 1.

Remark 2: The term *contact consistent* is a well-established term in the literature that was initially introduced in [26]. It implies formulating the control structure to account for the contact with the environment. The term c^3 is an extension of the term contact consistent. Hence, c^3 implies formulating the control structure to account for the *compliant* contact with the environment.

A. c^3-Stance Task

We model the soft contact dynamics with a simple explicit model (the KV model). This consists of 3D linear springs and dampers normal and tangential to the contact point [12]. The normal direction of this impedance emulates the normal terrain deformation while the tangential ones emulate the shear deformation. Although several models that accurately emulate contact dynamics are available [27, 28, 15], we implemented the KV model for several reasons. First, since the model is linear in the parameters, it fits our QP formulation. Second, estimating the parameters of the KV model is computationally efficient. As a result, using this model, we can run a learning algorithm online which would be challenging if a model similar to [14] is used. For a legged robot with point-like feet, for each stance leg i, we formulate the contact model in the world frame as

$$F_{\mathrm{grf,i}} = k_{\mathrm{st},i} p_i + d_{\mathrm{st},i} \dot{p}_i \tag{16}$$

where $k_{\mathrm{st},i} \in \mathbb{R}^{3\times3}$, $d_{\mathrm{st},i} \in \mathbb{R}^{3\times3}$, $F_{\mathrm{grf,i}} \in \mathbb{R}^3$, $p_i \in \mathbb{R}^3$, and $\dot{p}_i \in \mathbb{R}^3$ are the terrain stiffness, the terrain damping, the GRFs, the penetration and the penetration rate of the i-th stance leg, all expressed in the world frame, respectively (see Fig. 3). We define p_i and $\dot{p}_i$ as

$$p_i = x_{\mathrm{td},i} - x_{\mathrm{st},i}, \qquad \dot{p}_i = 0 - v_{\mathrm{st},i} \tag{17}$$

where $x_{\mathrm{td},i} \in \mathbb{R}^3$ denotes the position of the contact point of foot i at the touchdown in the world frame. By appending all of the stance feet, the contact model can be re-written in a compact form as

$$F_{\mathrm{grf}} = K_{\mathrm{st}} p + D_{\mathrm{st}} \dot{p} = K_{\mathrm{st}}(x_{\mathrm{td}} - x_{\mathrm{st}}) - D_{\mathrm{st}} v_{st} \tag{18}$$

where $K_{\mathrm{st}} \in R^{n_{\mathrm{st}}\times n_{\mathrm{st}}}$ and $D_{\mathrm{st}} \in R^{n_{\mathrm{st}}\times n_{\mathrm{st}}}$ are the block-diagonal stiffness and damping matrices of the terrain of all the stance feet, respectively, and $x_{\mathrm{td}} \in \mathbb{R}^{n_{st}}$ are the touchdown positions of all the stance feet.

Similar to Section III-A, we deal with the contact model (18) as another WBC task (alongside the trunk and swing (3) tasks). We can think of (18) as a desired stance task that keeps the WBC c^3. This stance task is achieved by a Cartesian impedance at the stance foot which is represented by the impedance of the terrain (K_{st} and D_{st}). This similarity makes us encode the contact model in the WBOpt as a stance constraint similar to what we did for the swing task in Section III-B. Hereafter, we refer to this stance task as the c^3-stance task (see Fig. 3).

B. Whole-Body Optimization Revisited

The c^3-stance task is included in the WBOpt by writing the soft contact model (18) as a function of the decision variables. Ideally, we can directly reformulate (18) as a function of F_{grf} and $\dot{v}_{\mathrm{st}}$. Indeed, $\dot{v}_{\mathrm{st}}$ can be expressed as a function of the joint accelerations $\ddot{q}$ which is a decision variable (as explained in [3]). By numerically integrating $\dot{v}_{\mathrm{st}}$ (once to obtain v_{st} and twice to obtain x_{st}), we can associate F_{grf} with $\dot{v}_{\mathrm{st}}$. This approach requires the knowledge of x_{td} to compute p which might be prone to estimation errors and it requires a reset of the integrator at every touchdown.

We choose a more convenient approach which is to add the *desired* foot penetration ϵ as an extra decision variable in the WBOpt formulation. The difference between p and ϵ is that p is the *actual* penetration due to the interaction with the soft contact while ϵ is the *desired* penetration in the world frame generated from the optimization problem. Both variables imply the same physical phenomenon (the soft contact deformation). That said, we can rewrite F_{grf} in (18) as a function of ϵ and $\dot{\epsilon}$ (by numerically differentiating ϵ) without the previous knowledge of x_{td}, which is advantageous.

To do so, ϵ is appended to the vector of decision variables u and regularized in (5). Then, we incorporate (18) directly as a function of F_{grf} and ϵ as in the equality constraint (8). We numerically differentiate ϵ to obtain $\dot{\epsilon}_k = \frac{\epsilon_k - \epsilon_{k-1}}{\Delta t}$. To maintain physical consistency, we need to enforce an additional constraint between the desired penetration ϵ and the contact acceleration ($\ddot{\epsilon} = -\dot{v}_{\mathrm{st}}$). This is encoded as an equality constraint as shown in (9). To do so, we numerically differentiate ϵ twice to obtain $\ddot{\epsilon}_k = \frac{\epsilon_k - 2\epsilon_{k-1} + \epsilon_{k-2}}{\Delta t^2}$. We also ensure the consistency of the physical contact model throughout the optimization problem by ensuring that the penetration is always positive in (10).

We also consider the loading and unloading phase, explained in [3] and [23], to be terrain-aware. We tune the loading and unloading phase period $T_{l/u}$ for each leg to follow the settling time of a second order system response that is a function of the terrain compliance and the robot's mass [29]. Hence, $T_{l/u}$ is

$$T_{l/u} = 4.6/\sqrt{\frac{k_{\mathrm{st,i}}}{m_e}} \tag{19}$$

where m_e is the equivalent mass felt at the robot's feet (i.e., the weight of the robot m_R spread across its stance feet $m_e = m_R/n_{\rm st}$) and the constant in the numerator represents a 1% steady-state error.

Finally, the WBOpt (5), (6), (8)-(15) generates optimal joint accelerations $\ddot{q}_j^*$ and contact forces $F_{\rm grf}^*$, that are mapped into optimal joint torques τ^* and sent to the low-level controller using the actuated part of the robot's dynamics as shown in (4). Note that similar to the sWBC, we found it sufficient to use only the inverse-dynamics term (the optimal torques τ^*) and not the joint feedback part.

As explained above, adding ϵ as a decision variable involved adding two constraints in the optimization which increases the problem size and the computation time. Yet, we are still able to run the c^{3}WBC in real-time. The advantage of our approach is that the knowledge of the touchdown position $x_{\rm td}$ is not required. We only need the previous two time instances of the penetration ϵ_{k-1} and ϵ_{k-2} that we already computed in the previous control loops.

V. Terrain Compliance Estimation

The purpose of the TCE is to estimate *online* the terrain parameters (namely $K_{\rm st}$ and $D_{\rm st}$) based on the states of the robot. It is a stand-alone algorithm that is decoupled from the c^{3}WBC. The TCE uses the contact model (18). Based on that, the current measurement of the *contact states* (contact status α, GRFs $F_{\rm grf}$, the penetration p, and the penetration rate $\dot{p}$) of each leg i at every time step are required. Given the contact states, we use supervised learning to learn the terrain parameters. As shown in Fig. 4, the TCE consists of two main modules: contact state estimation (Section V-A) and supervised learning (Section V-B). The contact state estimation module estimates the contact states and provides it to the supervised learning module that collects these data and computes the estimates of the terrain parameters.

A. Contact State Estimation

The contact states are estimated solely from the current states of the robot by the state estimator. The GRFs are estimated from the torques and the joint states, and the penetration and its rate are estimated from the floating base (trunk) states and the joint states.

1) GRFs Estimation: To estimate the GRF, we use actuated part of the dynamics in (2b) as

$$F_{{\rm grf},i} = \alpha_i J_{j,i}^{-T}(M_{a,i}\ddot{q}_i + h_{j,i} - \tau_{j,i}) \tag{20}$$

where $F_{{\rm grf},i}$, $J_{j,i}$, $M_{a,i}$, $\ddot{q}_i$, $h_{j,i}$, and $\tau_{j,i}$ correspond to $F_{\rm grf}$, $J_{{\rm st},j}$, M_a, $\ddot{q}$, h_j, and τ_j for the i-th leg, respectively. Additionally, α_i is the contact status variable that detects if there is a contact in the i-th leg or not. The contact is detected when the GRF exceed a certain threshold $F_{\rm min}$. Hence, α_i computed along the normal direction of the i-th leg n_i as:

$$\alpha_i = \begin{cases} 1, & \text{if } n_i^T(J_{j,i}^{-T}(M_{a,i}\ddot{q}_i + h_{j,i} - \tau_{j,i}) \geq F_{\rm min} \\ 0, & \text{otherwise} \end{cases} \tag{21}$$

2) Penetration Estimation: As shown in (17), we estimate the penetration and its rate using the stance feet positions $x_{{\rm st},i}$ and velocities $v_{{\rm st},i}$, and the touchdown position $x_{{\rm td},i}$ all in the *world* frame. To estimate the feet states in the world frame, we use the forward kinematics and the base state in the world frame. Thus, the penetration and its rate are written as

$$p_i = x_{{\rm td},i} - x_{{\rm st},i} = x_{{\rm td},i} - x_b - R_B^W x_{{\rm st},i}^B \tag{22}$$

$$\dot{p}_i = -v_{{\rm st},i} = -v_b - R_B^W v_{{\rm st},i}^B - (\omega_b \times R_B^W) x_{{\rm st},i}^B \tag{23}$$

where $x_b \in \mathbb{R}^3$ and $v_b \in \mathbb{R}^3$ are the base position and velocity in the world frame, respectively. The terms $x_{{\rm st},i}^B$ and $v_{{\rm st},i}^B$ are the stance feet position and velocity of the i-th leg in the base frame, respectively. The terms $R_B^W \in SO(3)$ and ω_b are the rotation matrix mapping vectors from the base frame to the world frame and the base angular velocity, respectively. The touch down positions are obtained using a height map.

3) Contact States Mapping: Since, the KV model consists of 3D linear springs and dampers, normal and tangential to the contact point, this makes the stiffness and damping matrices diagonal with respect to the contact frame. However, if expressed in the world frame, the stiffness and damping matrices become dense. Thus, if we formulate the KV model in the contact frame rather than the world frame, we estimate less number of elements per matrix per leg: three elements instead of nine. Henceforth, the KV model in the TCE should be formulated with respect to the contact frame rather than the world frame to reduce the computational complexity. To do so, the GRFs (20), the penetration (22) and its rate (23) of the i-th leg are transformed from the world frame Ψ_W to the contact frame Ψ_{C_i} as (see Fig. 3)

$$F_{{\rm grf},i}^C = R_W^{C_i} F_{{\rm grf},i} \tag{24}$$

$$p_i^C = R_W^{C_i} p_i \tag{25}$$

$$\dot{p}_i^C = R_W^{C_i} \dot{p}_i \tag{26}$$

where the superscript $\bullet^C$ refers to the contact frame and $R_W^{C_i}$ is the rotation matrix mapping from the world Ψ_W to the contact Ψ_{C_i} frames for the i-th leg. Note that the transformation (26) is linear since the contact frame is *fixed* with respect to the world frame at the touch down position (i.e., $\dot{R}_W^{C_i} = 0$).

B. Supervised Learning

Considering the contact model in the contact frame and using the estimated contact states (24)-(26), we learn the terrain parameters online via supervised learning. In particular, we use weighted linear least squared regression. The algorithm is treated as a batch algorithm with m-examples such that, at every time instant k, we gather samples from the previous m time instances and compute the terrain parameters [30].

For the k-th time instant, of the i-th leg in the d-th direction ($d \in \{n_i, t_{1,i}, t_{2,i}\}$, see Fig. 3), the terms $F_{{\rm grf},i}^{C,d}(k)$, $p_i^{C,d}(k)$, and $\dot{p}_i^{C,d}(k)$ are estimated as shown in Section V-A where $\bullet_i^d(k)$ refers to the k-th time instance of the i-th leg

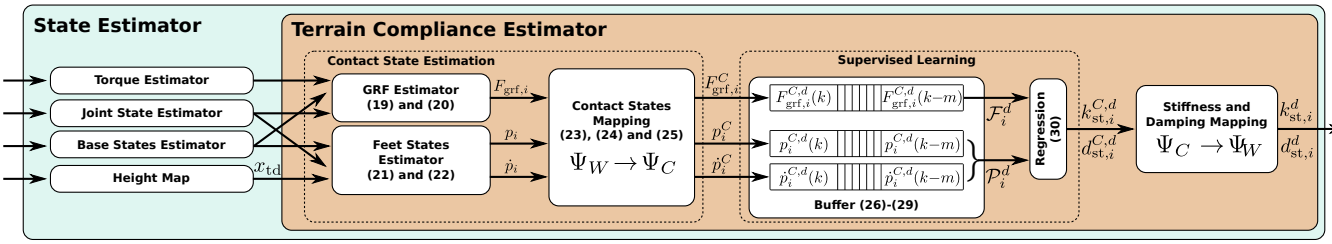

Fig. 4: Overview of the TCE's architecture inside the state estimator.

in the d-th direction. That said, we construct the following objects (buffers) with size m

$$\mathcal{F}_i^d = \begin{bmatrix} F^{C,d}_{\text{grf},i}(k) & \cdots & F^{C,d}_{\text{grf},i}(k-m) \end{bmatrix}^T \quad (27)$$

$$\mathrm{P}_i^d = \begin{bmatrix} p^{C,d}_i(k) & \cdots & p^{C,d}_i(k-m) \end{bmatrix}^T \quad (28)$$

$$\dot{\mathrm{P}}_i^d = \begin{bmatrix} \dot{p}^{C,d}_i(k) & \cdots & \dot{p}^{C,d}_i(k-m) \end{bmatrix}^T \quad (29)$$

$$\mathcal{P}_i^d = \begin{bmatrix} \mathrm{P}_i^d & \dot{\mathrm{P}}_i^d \end{bmatrix} \quad (30)$$

where $\mathcal{F}_i^d \in \mathbb{R}^m$ is the GRFs buffer and $\mathcal{P}_i^d \in \mathbb{R}^{m\times 2}$ is the penetration, and penetration rate buffer. Given $\mathcal{F}_i^d$ and $\mathcal{P}_i^d$ as inputs and outputs of the learning algorithm respectively, we estimate the terrain impedance parameters as $I_i^d = \begin{bmatrix} k^{C,d}_{\text{st},i} & d^{C,d}_{\text{st},i} \end{bmatrix}^T \in \mathbb{R}^2$ using the analytical solution

$$I_i^d = (\mathcal{P}_i^{dT} W \mathcal{P}_i^d)^{-1} \mathcal{P}_i^{dT} W \mathcal{F}_i^d \quad (31)$$

where $k^{C,d}_{\text{st},i} \in \mathbb{R}$ and $d^{C,d}_{\text{st},i} \in \mathbb{R}$ are the terrain stiffness and damping parameters expressed in the contact frame. The matrix $W \in \mathbb{R}^{m\times m}$ is a weighting matrix used to penalize the error on most recent sample compared to the less recent ones and thus, giving more importance to the most recent samples.

All of the legs in the learning algorithm are decoupled. We found it advantageous to treat each leg separately because the robot can be standing on a different terrain at each foot.

C. Implementation Details

Algorithm 1 sketches the entire TCE process. To initialize the buffers, we acquire samples when the robot is at full stance and return the first estimate of I_i^d once the buffers are full. After initialization, we acquire samples and update the buffers only when the leg is at stance.

The buffers are continuously updated in a sliding window fashion. When a leg finishes the swing phase and is at a new touch down, it continues to use the previous samples from the previous stance phase. This is advantageous since it gives a smooth transition between terrains, but it adds a delay.

Remark 3: Since the c^{3}WBC formulation is based in the world frame, it is essential to map the estimated stiffness and damping matrices back to the world frame before providing them to the c^{3}WBC (see Fig. 4).

Remark 4: The TCE can be used with any arbitrary terrain geometry given the terrain normal and thus $R_W^{C_i}$. The terrain normal n_i at the contact point i can be provided by a height map that is generated via an RGBD sensor.

Algorithm 1 Terrain Compliance Estimation

```
initialize the buffers (F_i^d and P_i^d) and I_i^d
for each iteration k do
    for each leg i do
        if leg is in contact (α_i == 1) then                    (21)
            for each direction d do
                estimate F^d_{grf,i}(k)                         (20)
                estimate p_i^d(k)                               (22)
                estimate ṗ_i^d(k)                               (23)
                transform F^d_{grf,i}(k) into F^{C,d}_{grf,i}(k) (24)
                transform p_i^d(k) into p_i^{C,d}(k)            (25)
                transform ṗ_i^d(k) into ṗ_i^{C,d}(k)            (26)
                update buffers F_i^d and P_i^d                  (27)-(30)
                solve for I_i^d                                 (31)
            end for
            map the estimated parameters to Ψ_W
        end if
    end for
end for
```

VI. Experimental Setup

A. State Estimation

We implemented our approach on HyQ [31] which is equipped with a variety of sensors. Each leg contains two load-cells, one torque sensor, and three high-resolution optical encoders. A tactical-grade Inertial Measurement Unit (IMU) (KVH 1775) is mounted on its trunk. Of particular importance to this experiment is the Vicon motion capture system (MCS). It is a multi-camera infrared system capable of measuring the pose of an object with high accuracy. During experiments, an accurate and non-drifting estimate of the position of the feet in the world frame is required to calculate the real penetration for the TCE. Typically, HyQ works independently of external sensors (e.g., MCS or GPS), however, soft terrain presents problems for state estimators [6]. This was re-affirmed in experiment.

The current state estimator [32] relies upon fusion of IMU and leg odometry data at a high frequency and uses lower frequency feedback from cameras or lidars to correct the drift. The leg odometry makes the assumption that the ground is rigid. On soft terrain, the estimator has difficulties in determining when a foot is in contact with the ground (i.e., is the foot in the air, or compressing the surface?). These errors cause the leg odometry signal to drift jeopardizing the estimation. Although, incorporating vision information could

be a possibility to correct for the drift in the estimation, improving state estimation on soft terrain is an ongoing area of research and is out of the scope of this paper.

Despite the drifting problem, we used the current state estimator [32] in our WBC because the planner in Fig. 3 has a re-planning feature that makes our WBC robust against a drifting state estimator [33]. However, the TCE still requires an accurate and non drifting estimate of the feet position in the world frame. Therefore, to validate the TCE, we used an external MCS that completely eliminates the drift problem. The MCS measures the pose of a special marker array placed on the head of the robot. Then the position of the feet in the world frame was calculated online by using the MCS measurement and the forward kinematics of the robot.

B. Terrain Compliance Estimator TCE Settings

In this work we used a sliding window of $m = 250$ samples (or 1 s for a control loop running at 250 Hz). Despite the general formulation, in this paper we estimate the terrain parameters only for the direction normal to the terrain, and assume that the tangential directions are the same. We carried out the simulation and experiment on a horizontal plane. Thus, the rotation matrix $R_W^{C_i}$ is identity. Furthermore, we did not estimate the damping parameter due to the inherent noise in the feet velocity signals that would jeopardize the estimation. The damping term $D_{\text{st}} v_{st}$ in (18) is less dominant in computing the GRFs compared to the stiffness term. This is because the feet velocities in the world frame v_{st} are usually orders of magnitude smaller than the penetration during stance, and the damping parameter D_{st} is usually orders of magnitude smaller than the stiffness parameter as shown in [28].

C. Tuning of the Low Level Control

During experiments, we found that the low level torque loop creates system instabilities when interacting with soft environments. In particular, when we used the same set of (high) torque gains in the low level control loop tuned for rigid terrain, we noticed joint instabilities when walking over soft terrain. This is because interacting with soft terrain reduces the stability margins of the system. Thus, keeping a high bandwidth in the inner torque loop given the reduced stability margin will cause system instability. In our previous work [34], we experimentally validated that increasing the torque gain of the inner loop can indeed cause system instabilities. In fact, this is a well know issue in haptics [35]. As a result, reducing the bandwidth by decreasing the torque gains in the inner torque loop was necessary to address these instabilities.

Our control design is a nested architecture consisting of the WBC and the low level torque control in which, both control loops contribute to the system stability [34, 36]. Over soft terrain, the dynamics of the environment also plays a role and must be considered in analyzing the stability of the system. That said, there is a nontrivial relationship between soft terrain and the stability of a nested control loop architecture, and a formal and thorough analysis is an ongoing work.

TABLE I
Mean Absolute Tracking Error (MAE) [N] of the GRFs in Simulation using sWBC, c^{3}WBC and STANCE over Multiple Terrains.

Terrain	**sWBC**	**c^{3}WBC**	**STANCE**
Soft	7.7261	7.4419	**6.3547**
Moderate	8.0594	**7.4585**	7.9889
Rigid	**4.889**	6.6523	5.128

VII. Results

In this section, we evaluate the proposed approach on HyQ in simulation and experiment. We compare *three* approaches: the sWBC which is the baseline, the c^{3}WBC which is our proposed WBC without the TCE, and STANCE which incorporates both the c^{3}WBC and TCE. We show the extent of improvement given by the c^{3}WBC controller with respect to the sWBC as well as the importance of the TCE during locomotion over multiple terrains with different compliances. We set the same parameters and gains throughout the entire simulations and experiments, unless mentioned otherwise. The results are shown in the accompanying video[1].

A. Simulations

To render soft terrain in simulation, we used the Open Dynamics Engine (ODE) physics engine [37]. We used ODE because it is easily integrable with our framework, and it is numerically fast and stable for stiff and soft contacts [38]. Moreover, ODE can render soft contacts that emulates physical parameters (using the SI units N/m and Ns/m for springs and dampers, respectively) unlike other engines that uses non-physical ones [39]. ODE's implicit solver uses linear springs and dampers for their soft constraints which fits perfectly with our contact model (18). In this way, we have a controlled simulation environment where we can emulate any terrain compliance by manipulating its stiffness K_t and damping D_t parameters similar to our contact model. Throughout the simulation, we use four types of terrains with the following parameters: soft T_1 ($K_t = 3500$ N/m), moderate T_2 ($K_t = 8000$ N/m), stiff T_3 ($K_t = 10000$ N/m), and rigid T_4 ($K_t = 2 \times 10^6$ N/m) all with the same damping ($D_t = 400$ Ns/m).

1) Locomotion over Multiple Terrains: We evaluate the three approaches with the robot walking at 0.05 m/s over the terrains: T_1 (soft), T_2 (moderate), and T_4 (rigid). We provided the c^{3}WBC with the terrain parameters of the moderate terrain T_2 for all the three simulations. We do that in order to test the performance of c^{3}WBC if given the real terrain parameters (in case of T_2) or inaccurate parameters (in case of T_1 and T_4). In this simulation, we compare the actual values of $F_{\text{grf},\perp}$ against the optimal values $F^*_{\text{grf},\perp}$(solution of the WBOpt) as well as the actual penetration p against the desired penetration ϵ of the LF leg. We have omitted the other three feet for space as all four legs have the same performance. The results are shown in Fig. 5. The Mean Absolute Tracking Error (MAE) of the GRFs in these simulations are presented in Table I. The MAE of the GRFs is defined as: $\text{MAE} = \frac{1}{T}\int_0^T |F_{\text{grf}} - F^*_{\text{grf}}| dt$.

[1]Link: `https://youtu.be/0BI4581DFjY`

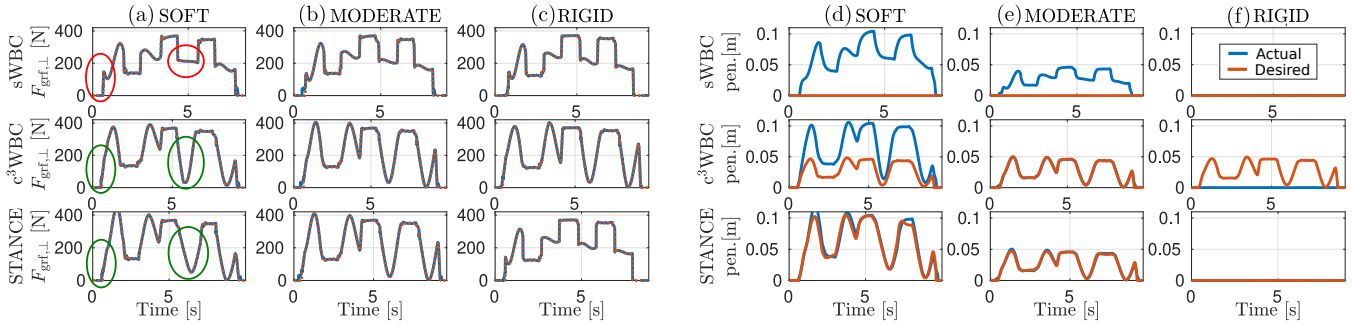

Fig. 5: Simulation. Comparison of sWBC, c^{3}WBC, and STANCE over three type of terrains: soft T_1 ($K_t = 3500$ N/m), moderate T_2 ($K_t = 8000$ N/m), and rigid T_4 ($K_t = 2 \times 10^6$ N/m) all with the same damping ($D_t = 400$ Ns/m). (a)-(c): The actual and desired contact forces in the normal direction of the LF leg for one gait cycle. (d)-(f): The actual p and desired ϵ penetration in the normal direction of the LF leg for one gait cycle. The red and green ellipses highlight the performance of the three approaches in adapting to soft terrain.

Fig. 5a captures the effect of the three approaches on the GRFs over soft terrain. We can see that a WBC based on a rigid contact assumption (sWBC) assumes that it can achieve an infinite bandwidth from the terrain and thus supplying an instantaneous change in the GRFs as highlighted by the red ellipses in Fig. 5a. On the other hand, STANCE and c^{3}WBC were both capable of attenuating this effect as highlighted by the green ellipses. For the reasons explained earlier in this paper and in [10], instantaneous changes in the GRFs are undesirable over soft terrain. This resulted in an improvement in the tracking of the GRFs in STANCE and c^{3}WBC compared to sWBC as shown in Table I. Moreover, by comparing c^{3}WBC and STANCE over soft terrain T_1, we can see that the shape of the GRFs did not differ. However, the tracking of the GRFs in STANCE is better than the c^{3}WBC. This shows that suppling the c^{3}WBC with the incorrect values of the terrain parameters deteriorates the GRFs tracking performance. Fig. 5b shows the GRFs on a moderate terrain. Since the c^{3}WBC is provided with the exact terrain parameters of T_2, we can perceive the c^{3}WBC as STANCE with a perfect TCE on moderate terrain. As a result, Table I shows that in this set of simulations, c^{3}WBC outperformed STANCE in the GRFs tracking. This shows that a more accurate TCE can result in a better GRFs tracking. Additionally, Fig. 5c shows the GRFs on rigid terrain. We can see that the sWBC resulted in a typical (desired) shape of the GRFs for a crawl motion in rigid terrain [33]. STANCE showed a shape of the GRFs similar to the sWBC which is expected since the TCE provided STANCE with parameters similar to the rigid terrain. However, for c^{3}WBC, the GRFs shape did not change compared to the other three terrains. As shown in Table I, the best tracking to the GRFs was by the sWBC, which was expected since the sWBC was designed for rigid terrain. However, sWBC was only slightly better than STANCE due to small estimation errors from the TCE.

Fig. 5a-c show the superiority of STANCE compared to sWBC and c^{3}WBC. STANCE adapted to the three terrains by estimating their parameters and supplying them to the WBC. This resulted in changing the shape of the GRFs accordingly that improved the tracking of the GRFs. Unlike STANCE, the sWBC and the c^{3}WBC both are contact consistent for only *one* type of terrain which resulted in a deterioration of the

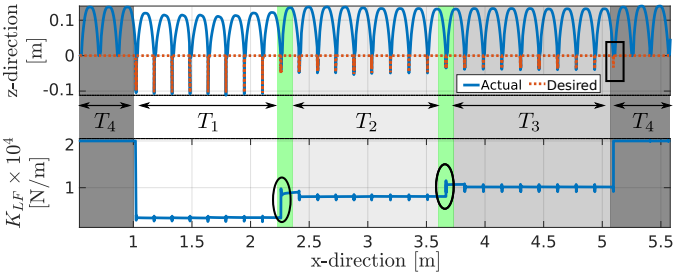

Fig. 6: Simulation. Traversing multiple terrains of different compliances (T_4, T_1, T_2, T_3, T_4). Top: Tracking of the desired terrain penetration of the LF leg in the xz-plane. Bottom: Estimated terrain stiffness of the LF leg. For readability purposes we only plot estimated values less than 2×10^4. The green shaded areas highlight the overlap between terrains that results in higher estimated stiffness (black ellipses).

GRFs tracking over the other types of terrains. The advantages of STANCE compared to sWBC and c^{3}WBC are also shown in Fig. 5d-f. Since the sWBC is always assuming a rigid contact, the penetration ϵ was always zero throughout the three terrains. Similarly, since the c^{3}WBC alone is aware only of one type of terrain, it is always assuming the same contact model, in which the desired penetration ϵ was similar throughout the three terrains. STANCE, however, was capable of predicting the penetration correctly for all the three terrains.

In general, even if the contact model is for soft contacts, STANCE was capable of correctly predicting the penetration of the robot even in rigid terrain (zero penetration). This resulted in STANCE adapting to rigid, soft and moderate terrains by means of adapting the GRFs and correctly predicting the penetration.

2) Longitudinal Transition Between Multiple Terrains: We show the adaptation of STANCE when walking and transitioning between multiple terrains. We test the accuracy of the TCE and the effect of closing the loop of the c^{3}WBC with the TCE on the feet trajectories and terrain penetration. In this simulation, HyQ is traversing five different terrains, starting and ending with a rigid terrain: T_4, T_1, T_2, T_2, T_4. The results are presented in Fig. 6. The top plot presents the actual foot position against the desired penetration ϵ of the LF leg in the xz-plane of the world frame. The origin of the z-direction (normal direction) is the uncompressed terrain height. Thus, trajectories below zero represent the penetration of the

TABLE II
Mean μ [N/m], Standard Deviation σ [N/m], and Percentage Error of the Estimated Terrain Stiffness of the LF Leg in Simulation.

Terrain	Actual Stiffness	Mean $\mu \pm$ STD σ	% Error
T_1	3500	3530 $\pm$ 200	0.9%
T_2	8000	8110 $\pm$ 400	1.4%
T_3	10000	10110 $\pm$ 400	1.1%
T_4	2000000	2240000 $\pm$ 740000	12%

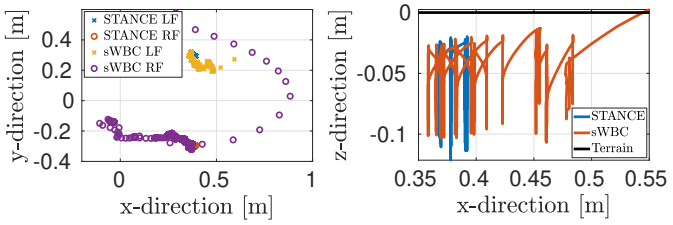

Fig. 7: Simulation. Comparing sWBC and STANCE under aggressive trunk maneuvers. Left: Top view of the front feet (RF and LF) positions. Right: Side view of the LF position.

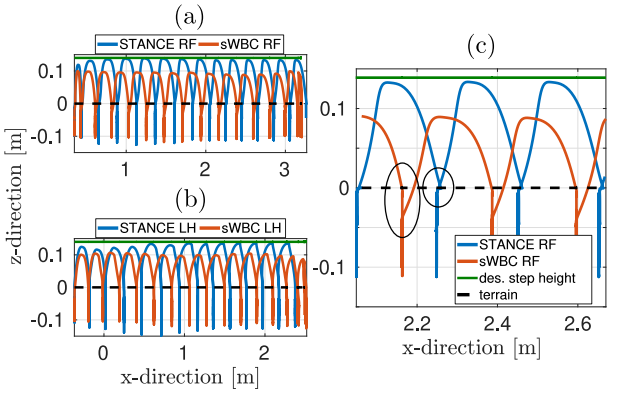

Fig. 8: Simulation. Speed test. Increasing the desired forward velocity from 0.05 to 0.3 m/s. Left: Side view of the RF (a) and LH (b) positions. Right: (c) Closeup section of the top left plot. The green lines are the desired step height. The black dashed lines are the terrain height.

LF leg. The bottom plot shows the history of the estimated terrain stiffness of the TCE of the LF leg. Table II reports the mean, standard deviation, and percentage error[2] of the estimated terrain stiffness of the LF leg against the ground truth value set in ODE. The table shows that the TCE had an estimation accuracy below 2% for the soft terrains T_1, T_2 and T_3. However, the estimation accuracy of the rigid terrain was lower than that of the soft terrains. This is expected since on a rigid terrain, the penetrations are (almost) zero. Thus, a small inaccurate penetration estimation due to any model errors could result in a lower estimation accuracy. Apart from the rigid case, the standard deviation is always below 6% of the ground truth value. Fig. 6 shows that STANCE is always c^3, the actual foot position is always consistent with the desired penetration during stance. We can see that, when HyQ is standing over rigid terrain, both the actual foot position and desired penetration are zero. As HyQ walks, over the soft terrains, the penetration is highest in the softest terrain and smallest in the stiffest terrain.

In the simulation environment, we overlapped the terrains to prevent the feet from getting stuck between them. This overlap created a transition (highlighted in green in the figure) which resulted in a stiffer terrain. The overlap was captured by the TCE and resulted in a slight increase in the estimated parameters as highlighted by the two ellipses in the lower plot. We also noticed a lag in estimation, due to a filtering effect, since the TCE is using the most recent m-samples. As highlighted by the black box in Fig. 6, HyQ was on rigid terrain (actual penetration is zero) while STANCE still perceived it as being on T_3 (desired penetration is non-zero).

3) Aggressive trunk maneuvers: We tested sWBC and STANCE under aggressive trunk maneuvers by commanding desired sinusoidal trajectories at the robot's height (0.05 m amplitude and 1.8 Hz frequency) and at roll orientation (0.5 rad amplitude and 1.5 Hz frequency) over the soft terrain T_1. The results are shown in Fig. 7. The left plot shows a top view of the actual front feet (LF and RF) positions in the world frame. The right plot shows a side view of the actual LF foot position in the world frame. We notice that the feet of HyQ are always in contact with the terrain in STANCE which is expected since STANCE is c^3. Unlike STANCE, the feet did not remain in contact with the terrain in the sWBC. This is clearly seen in Fig. 7 where HyQ lost contact multiple times. This resulted in the robot falling over in the sWBC case as shown in the video.

4) Speed Test: We carried out a simulation where HyQ walks over soft terrain T_1, starting with a forward velocity of 0.05 m/s until it reaches 0.3 m/s with an acceleration of 0.005 m/s^2. In this simulation, we compare STANCE against the sWBC. Fig. 8a and Fig. 8b show the actual trajectories of the RF and LH legs in the world frame, respectively. Fig. 8c shows a closeup section of the RF leg's trajectory. The simulation shows that STANCE was c^3 over the entire simulation while the sWBC was not.

In particular, STANCE was able to remain in contact with the terrain that allowed HyQ to start the swing phase directly from the terrain height. Unlike STANCE, the sWBC is not terrain aware and did not remain c^3 which resulted in starting the swing trajectory while still being inside the deformed terrain. This is highlighted by the two ellipses in the right plot. Additionally, the compliance contact consistency property of STANCE enabled the robot to maintain the desired step clearance (i.e., achieving the desired step height of 0.14 cm) compared to sWBC. Most importantly, as shown in the accompanying video, the sWBC failed to complete the simulation and could not achieve the final desired forward velocity; It fell at a speed of 0.21 m/s. Note that both approaches could reach higher velocities with a more dynamic gait (trot). However, this simulation is not focusing on analyzing the maximum speed that the two approaches can reach but rather the differences between these approaches at a higher crawl speeds.

5) Power Test: In this test, we compare the power consumption using STANCE and sWBC on HyQ during walking over the soft terrain T_1 at different forward velocities (0.05 m/s, 0.15 m/s and 0.25 m/s). Fig. 9 presents the energy

[2] The percentage error is defined as: % Error $= |\frac{\text{Estimate}-\text{Actual}}{\text{Actual}}| \times 100$

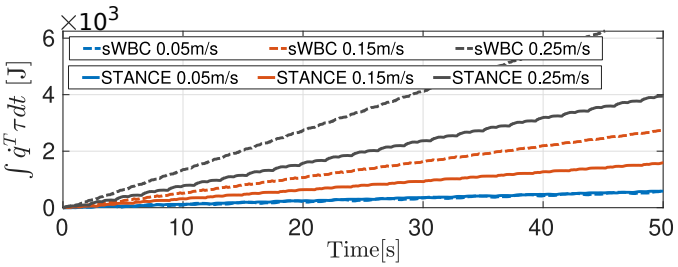

Fig. 9: Simulation. Power consumption comparison between sWBC and STANCE with different forward velocities (0.05 m/s, 0.15 m/s and 0.25 m/s).

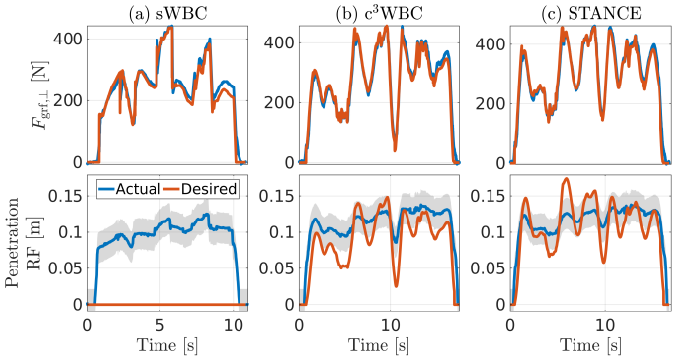

Fig. 10: Experiment. Comparing sWBC, c^{3}WBC and STANCE over a soft foam block ($K_t = 2400$ N/m). Top: Tracking of the GRFs of the RF leg. Bottom: Tracking of the foot penetration. The gray shaded areas represent the uncertainty of the measurements.

plots of STANCE and sWBC. The plot shows that STANCE requires less power than the sWBC because it knows how the terrain will deform. STANCE exploits the terrain interaction to achieve the motion. The difference in consumed energy is negligible at 0.05 m/s but becomes significant at higher speeds.

B. Experiment

We validated the simulation presented in Section VII-A on the real platform. We analyzed sWBC, c^{3}WBC (with fixed terrain parameters) and STANCE as well as the performance of the TCE module itself.

A foam block of 160 cm × 120 cm × 20 cm was selected as a soft terrain for these experiments. To obtain a ground truth of the foam stiffness, we carried out indentation tests on a 50 cm^3 sample of the foam with a stress-strain machine that covers the range of penetration of interest for our robot (below 0.15 cm). The indentation test showed a softening behavior of the foam with an average stiffness of 2400 N/m. The MAE of the GRFs of the upcoming experiments are shown in Table III.

1) Locomotion over Soft Terrain: In this experiment, HyQ is walking over the foam with a forward velocity of 0.07 m/s using the three approaches. The results are presented in Fig. 10 that shows the actual and desired $F_{\text{grf},\perp}$ and penetration of the RF leg. The shaded gray area in the lower plots of Fig. 10 represents the uncertainty in the estimation of the foot position (see Section VI-A). In these experiments, all three approaches performed well; none of them failed. However, the shape of GRFs were different within the three approaches. As in Section VII-A1, since sWBC is rigid contact consistent, the

TABLE III
Mean Absolute Tracking Error (MAE) [N] of the GRFs using sWBC, c^{3}WBC and STANCE under Different Sets of Experiments.

Description	**sWBC**	**c^{3}WBC**	**STANCE**
Soft Terrain (Sec. VII-B1)	73.9042	68.5581	**61.8207**
Longitudinal Trans. (Sec. VII-B2)	70.5276	64.2636	**60.6285**
Lateral Trans. (Sec. VII-B3)	73.0766	-	**53.0107**

TABLE IV
Mean μ [N/m], Standard Deviation σ [N/m], and Percentage Error of the Estimated Terrain Stiffness of the Four Legs in Experiment over Soft Terrain (2400 N/m).

Leg	Mean $\mu \pm$ STD σ	% Error
LF	2186 ± 166	9%
RF	2731 ± 173	14%
LH	2368 ± 317	1%
RF	2078 ± 331	13%

desired GRFs were designed for rigid contacts. Unlike sWBC, STANCE is c^3, which was capable of changing the shape of the GRFs. This is highlighted in Table III in which STANCE outperformed sWBC in the tracking of the GRFs.

In simulation, when we provided the c^{3}WBC with the true value of the stiffness, the MAE of the GRFs was better. However, in this experiment, providing the value obtained from the indentation tests to the c^{3}WBC resulted in a worse GRFs MAE. This outperformance of STANCE compared to the c^{3}WBC in this experiment could be because of the TCE. To clarify, the actual terrain compliances are not constant, but since the TCE is online, it is able to capture these changes in the terrain compliances as well as model errors. As shown in the accompanying video, STANCE had a smoother transition during crawling compared to sWBC. We found the robot transitioning from swing to stance more aggressively in sWBC than STANCE. Such smooth behavior was also noticed in [10].

Table IV shows the mean, standard deviation, and percentage error of the estimated terrain stiffness of all the four legs against the ground truth value (2400 N/m) obtained from the indentation tests. The table shows that the accuracy of the TCE in simulation is better than in experiments. This is expected since in simulation, the TCE has a perfect knowledge of the feet penetration. However, the accuracy of our TCE is better compared to [20] in which the percentage error exceeded 50% (the actual stiffness was more than double that of the estimated one in [20]).

2) Longitudinal Transition Between Multiple Terrains: Similar to Section VII-A2, we compare the three approaches while transitioning between the foam block and a rigid pallet. We added a pad between between the two terrains to avoid the foot getting stuck (see Fig. 1a). Fig. 11a-c show the actual position and the desired penetration of the RF leg in the xz-plane for the three approaches. Fig. 11d shows the estimated terrain stiffness of the TCE for all four feet.

From Fig. 11a-b we see that both sWBC and c^{3}WBC did not adapt to terrain changes. Since both controllers are designed for a specific constant terrain, the desired penetration did not change from soft to rigid. In the sWBC, there is no tracking of the penetration, and in the c^{3}WBC, the tracking of the

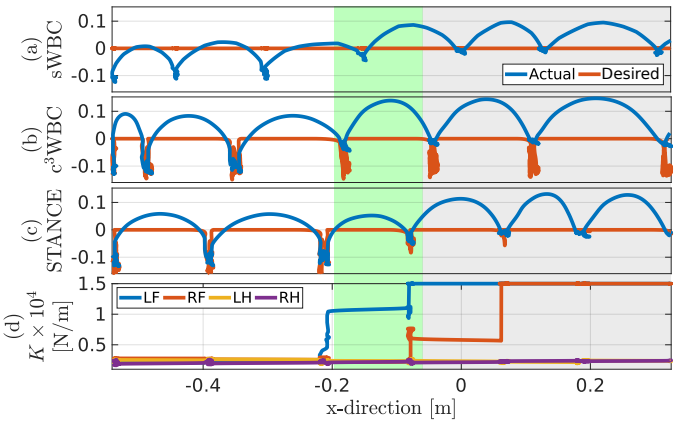

Fig. 11: Experiment. Longitudinal transition from soft to rigid terrain. The first three plots show the tracking of the desired foot penetration of RF leg using the three approaches (the sWBC, c^{3}WBC with fixed terrain stiffness and STANCE. The fourth plot shows the stiffness estimated by the TCE for the four legs.

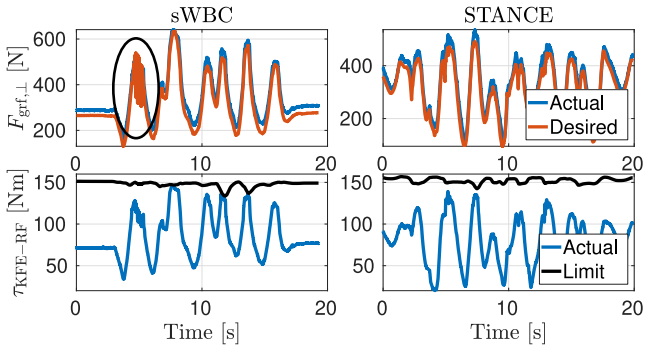

Fig. 12: Experiment. The sWBC and STANCE under disturbances over soft terrain. Top: The actual and desired $F_{\mathrm{grf},\perp}$ in sWBC and STANCE, respectively. Bottom: The actual torque and torque limits of the Knee Flexion-Extension (KFE) joint of the RF leg in sWBC and STANCE, respectively.

penetration is good only when the leg is on the foam where the stiffness is consistent to the one used in the controller. On the other hand, as shown in Fig. 11c-d, STANCE changes its parameters when facing a different terrain; it was capable of adapting its desired penetration to the type of terrain. In fact, the desired penetration was non-zero on soft terrain and was almost zero on rigid terrain. This again resulted in STANCE achieving the best GRFs tracking as shown in Table III.

Fig. 11d shows the importance of having a TCE for each leg. The estimated terrain parameters are different between the legs where the hind legs are on the foam while the rigid ones transitioning from foam to rigid. The figure also shows that the LF leg walked over the rigid terrain before the RF and that the TCE captures the intermediate stiffness estimation due to the rubber pad (see video).

3) Lateral Transition Between Multiple Terrains: Unlike the previous experiment, we set the foam and the pallet laterally as shown in Fig. 1c and in the accompanying video. This is a more challenging scenario for stability reasons. In particular, the robot must extend its leg further in the soft terrain maintain the trunk's balance. Consequently, since the width of HyQ's torso is smaller than its length, the Zero Moment Point (ZMP) is more likely to get out of the support polygon. The GRFs MAE in Table III show that STANCE can outperform sWBC during both longitudinal and lateral transitions.

4) External Disturbances over Soft Terrain: In this experiment, we test the sWBC and STANCE when the user applies a disturbance on HyQ. The results are shown in Fig. 12. The top plots show the actual and desired $F_{\mathrm{grf},\perp}$ in sWBC and STANCE, respectively. The bottom plots show the actual torque and torque limits of the Knee Flexion-Extension (KFE) joint of the RF leg in sWBC and STANCE, respectively. In the accompanying video, we can qualitatively see that with STANCE, the feet of HyQ keep moving to remain c^3 with the terrain. On the other hand, the sWBC kept its feet stationary. This behavior was also reported by [6].

Most importantly, we noticed that HyQ reaches the torque limits in the sWBC as shown in Fig. 12. However, in STANCE, since the robot was constantly moving its feet, hence re-distributing its forces, the torque limits were not reached. This behavior was also reflected on the GRFs in which, the GRFs were resonating in the sWBC as highlighted by the ellipse in Fig. 12.

TABLE V
Mean μ [N/m] and Standard Deviation σ [N/m] of the Estimated Terrain Stiffness of the Four Legs in Experiments (see Fig. 1b).

Leg	Mean $\mu \pm$ STD σ
LF	448400 $\pm$ 165100
RF	55200 $\pm$ 48400
LH	2645000 $\pm$ 336000
RH	1393000 $\pm$ 442000

5) TCE's Performance over Multiple Terrains: We analyze the performance of the TCE on HyQ over multiple terrains with various softnesses. The softness of the four used terrains are shown in Fig. 1b. The estimated stiffness (mean and standard deviation) under each leg is shown in Table V. As shown in the table, the robot can differentiate between the types of terrain. Although we did not measure the true stiffness value of these terrains, we can observe their softness in the video and Fig. 1b and compare it to the values in Table V.

C. Computational Analysis

STANCE is running online which means that we can estimate the terrain compliance (using the TCE) continuously while walking, and run the entire framework without breaking real-time requirements. We validated the first argument by showing that indeed the TCE can continuously estimate the terrain compliance. Hereafter, we validate the second argument by analyzing the computational complexity of STANCE and compare it against the sWBC. Since our WBC framework is running at 250 Hz, it is essential that the computation does not exceed the 4 ms time frame. Hence, we conducted a simulation in which we calculated the time taken to process the entire framework without the lower level control (ie., the state estimator, the planner and the WBC) that is running on a different real-time thread at 1 kHz. We compared the computation time on an Intel Core i7 quad core CPU in the case of STANCE (the c^{3}WBC and the TCE) and the sWBC. We used the same parameters and gains as in Section VII-A1.

The results show that the average processing time taken was 0.68 ms and 0.74 ms for the sWBC and STANCE respectively. In both cases, the maximum computation time was always below 2 ms.

VIII. Conclusions

We presented a soft terrain adaptation algorithm called STANCE: **S**oft **T**errain **A**daptation a**N**d **C**ompliance **E**stimation. STANCE can adapt online to any type of terrain compliance (stiff or rigid). STANCE consists of two main modules: a compliant contact consistent whole-body controller (c^3WBC) and a terrain compliance estimator (TCE). The c^3WBC extends our previously implemented WBC (sWBC) [3], such that it is contact consistent to any type of compliant terrain given the terrain parameters. The TCE estimates online the terrain compliance and closes the loop with the c^3WBC. Unlike previous works on WBC, STANCE does not assume that the ground is rigid. Stance is computationally lightweight and it overcomes the limitations of the previous state of the art approaches. As a result, STANCE can efficiently traverse multiple terrains with different compliances. We validated STANCE on our quadruped robot HyQ over multiple terrains of different stiffness in simulation and experiment. This, to the best of the authors' knowledge, is the first experimental validation on a legged robot of closing the loop with a terrain estimator.

Incorporating the terrain knowledge makes STANCE c^3. This allows STANCE to generate smooth GRFs that are physically consistent with the terrain, and continuously adapt the robot's feet to remain in contact with the terrain. As a result, the tracking error of the GRFs and the power consumption were reduced, and the impact during contact interaction was attenuated. Furthermore, STANCE is more robust in challenging scenarios. As demonstrated, STANCE made it possible to perform aggressive maneuvers and walk at high walking speeds over soft terrain compared to the state of the art sWBC. In the standard case, the contact is lost because the motion of the terrain is not taken into account. On the other hand, there are minor differences in performance between STANCE and the sWBC for less dynamic motions.

STANCE can efficiently transition between multiple terrains with different compliances, and each leg was able to independently sense and adapt to the change in terrain compliance. We also tested the capability of the TCE in discriminating between different terrains. The insights gained in simulation have been confirmed in experiment.

In future works, we plan to implement an algorithm to improve the TCE. In particular, we plan on using onboard sensors, such as a camera, instead of relying on the external measurements from an MCS. We also plan to explore other non-linear contact models in the TCE and the c^3WBC.

Acknowledgements

We thank Gustavo Medrano-Cerda, Roy Featherstone, and all the DLS lab members for the help provided during this work.

References

[1] F. Farshidian, E. Jelavi, A. W. Winkler, and J. Buchli, "Robust whole-body motion control of legged robots," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, Vancouver, Canada, Sep. 2017, pp. 4589–4596, DOI:10.1109/IROS.2017.8206328.

[2] C. Dario Bellicoso, F. Jenelten, P. Fankhauser, C. Gehring, J. Hwangbo, and M. Hutter, "Dynamic locomotion and whole-body control for quadrupedal robots," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, Vancouver, Canada, Sep. 2017, pp. 3359–3365, DOI:10.1109/IROS.2017.8206174.

[3] S. Fahmi, C. Mastalli, M. Focchi, and C. Semini, "Passive whole-body control for quadruped robots: Experimental validation over challenging terrain," *IEEE Robot. Automat. Lett. (RA-L)*, vol. 4, no. 3, pp. 2553–2560, Jul. 2019, DOI:10.1109/LRA.2019.2908502.

[4] B. Henze, A. Dietrich, M. A. Roa, and C. Ott, "Multi-contact balancing of humanoid robots in confined spaces: Utilizing knee contacts," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, Vancouver, Canada, Sep. 2017, pp. 697–704, DOI:10.1109/IROS.2017.8202227.

[5] K. Bouyarmane, K. Chappellet, J. Vaillant, and A. Kheddar, "Quadratic programming for multirobot and task-space force control," *IEEE Trans. Robot. (T-RO)*, vol. 35, no. 1, pp. 64–77, Feb. 2019, DOI:10.1109/TRO.2018.2876782.

[6] B. Henze, M. A. Roa, and C. Ott, "Passivity-based whole-body balancing for torque-controlled humanoid robots in multi-contact scenarios," *Int. J. Robot. Res. (IJRR)*, vol. 35, no. 12, pp. 1522–1543, Jul. 2016, DOI:10.1177/0278364916653815.

[7] B. Henze, R. Balachandran, M. A. Roa-Garzn, C. Ott, and A. Albu-Schffer, "Passivity analysis and control of humanoid robots on movable ground," *IEEE Robot. Automat. Lett. (RA-L)*, vol. 3, no. 4, pp. 3457–3464, Oct. 2018, DOI:10.1109/LRA.2018.2853266.

[8] M. Azad and M. N. Mistry, "Balance control strategy for legged robots with compliant contacts," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, Seattle, USA, May 2015, pp. 4391–4396, DOI:10.1109/ICRA.2015.7139806.

[9] V. Vasilopoulos, I. S. Paraskevas, and E. G. Papadopoulos, "Monopod hopping on compliant terrains," *Robot. Auton. Syst.*, vol. 102, pp. 13–26, Apr. 2018, DOI:10.1016/j.robot.2018.01.004.

[10] R. Grandia, F. Farshidian, A. Dosovitskiy, R. Ranftl, and M. Hutter, "Frequency-aware model predictive control," *IEEE Robot. Automat. Lett. (RA-L)*, vol. 4, no. 2, pp. 1517–1524, Apr. 2019, DOI:10.1109/LRA.2019.2895882.

[11] D. Kim, S. Jorgensen, J. Lee, J. Ahn, J. Luo, and L. Sentis, "Dynamic locomotion for passive-ankle biped robots and humanoids using whole-body locomotion control," *arXiv preprint*, pp. 1–18, 2019. [Online]. Available: https://arxiv.org/abs/1901.08100

[12] M. Neunert, M. Stuble, M. Giftthaler, C. D. Bellicoso, J. Carius, C. Gehring, M. Hutter, and J. Buchli, "Whole-body nonlinear model predictive control through contacts for quadrupeds," *IEEE Robot. Automat. Lett. (RA-L)*, vol. 3, no. 3, pp. 1458–1465, Jul. 2018, DOI:10.1109/LRA.2018.2800124.

[13] N. Doshi, K. Jayaram, B. Goldberg, Z. Manchester, R. Wood, and S. Kuindersma, "Contact-implicit optimization of locomotion trajectories for a quadrupedal microrobot," in *Proc. Robot.: Sci. and Syst. (RSS)*, Pittsburgh, USA, 2018, pp. 1–10.

[14] A. H. Chang, C. M. Hubicki, J. J. Aguilar, D. I. Goldman, A. D. Ames, and P. A. Vela, "Learning to jump in granular media: Unifying optimal control synthesis with gaussian process-based regression," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, Singapore, Singapore, May 2017, pp. 2154–2160, DOI:10.1109/ICRA.2017.7989248.

[15] J. Alves, N. Peixinho, M. T. da Silva, P. Flores, and H. M. Lankarani, "A comparative study of the viscoelastic constitutive models for frictionless contact interfaces in solids,"

Mech. Mach. Theory, vol. 85, pp. 172–188, Mar. 2015, DOI:10.1016/j.mechmachtheory.2014.11.020.

[16] R. Schindeler and K. Hashtrudi-Zaad, "Online identification of environment huntcrossley models using polynomial linearization," *IEEE Trans. Robot. (T-RO)*, vol. 34, no. 2, pp. 447–458, Apr. 2018, DOI:10.1109/TRO.2017.2776318.

[17] M. Azad, V. Ortenzi, H. Lin, E. Rueckert, and M. Mistry, "Model estimation and control of compliant contact normal force," in *Proc. IEEE/RAS Int. Conf. Humanoid Robot. (Humanoids)*, Cancun, Mexico, Nov. 2016, pp. 442–447, DOI:10.1109/HUMANOIDS.2016.7803313.

[18] F. Coutinho and R. Corteso, "Online stiffness estimation for robotic tasks with force observers," *Control Eng. Pract.*, vol. 24, pp. 92–105, Mar. 2014, DOI:10.1016/j.conengprac.2013.11.002.

[19] F. Coutinho and R. Corteso, "A neural-based approach for stiffness estimation in robotic tasks," in *Proc. Int. Conf. Adv. Robot. (ICAR)*, Montevideo, Uruguay, Nov. 2013, pp. 1–7, DOI:10.1109/ICAR.2013.6766499.

[20] W. Bosworth, J. Whitney, Sangbae Kim, and N. Hogan, "Robot locomotion on hard and soft ground: Measuring stability and ground properties in-situ," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, Stockholm, Sweden, May 2016, pp. 3582–3589, DOI:10.1109/ICRA.2016.7487541.

[21] A. Herzog, N. Rotella, S. Mason, F. Grimminger, S. Schaal, and L. Righetti, "Momentum control with hierarchical inverse dynamics on a torque-controlled humanoid," *Auton. Robot.*, vol. 40, no. 3, pp. 473–491, Mar. 2016, DOI:10.1007/s10514-015-9476-6.

[22] R. Ortega, J. A. L. Perez, P. J. Nicklasson, and H. J. Sira-Ramirez, *Passivity-based control of Euler-Lagrange systems: mechanical, electrical and electromechanical applications*, 1st ed., ser. Commun. Control Eng. Springer Science & Business Media, 1998, DOI:10.1007/978-1-4471-3603-3.

[23] M. Focchi, A. del Prete, I. Havoutis, R. Featherstone, D. G. Caldwell, and C. Semini, "High-slope terrain locomotion for torque-controlled quadruped robots," *Auton. Robot.*, vol. 41, no. 1, pp. 259–272, Jan. 2017, DOI:10.1007/s10514-016-9573-1.

[24] L. Righetti, J. Buchli, M. Mistry, M. Kalakrishnan, and S. Schaal, "Optimal distribution of contact forces with inverse-dynamics control," *Int. J. Robot. Res. (IJRR)*, vol. 32, no. 3, pp. 280–298, Jan. 2013, DOI:10.1177/0278364912469821.

[25] G. Tournois, M. Focchi, A. Del Prete, R. Orsolino, D. G. Caldwell, and C. Semini, "Online payload identification for quadruped robots," in *Proc. IEEE/RSJ Int. Conf. Intell. Robot. Syst. (IROS)*, Vancouver, Canada, Sep. 2017, pp. 4889–4896, DOI:10.1109/IROS.2017.8206367.

[26] Jaeheung Park and O. Khatib, "Contact consistent control framework for humanoid robots," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, Orlando, USA, May 2006, pp. 1963–1969, DOI:10.1109/ROBOT.2006.1641993.

[27] M. Azad, R. Featherstone *et al.*, "Modeling the contact between a rolling sphere and a compliant ground plane," in *Proc. Australas. Conf. Robot. Automat. (ACRA)*, Brisbane, Australia, Dec. 2010, pp. 100–107.

[28] L. Ding, H. Gao, Z. Deng, J. Song, Y. Liu, G. Liu, and K. Iagnemma, "Footterrain interaction mechanics for legged robots: Modeling and experimental validation," *Int. J. Robot. Res. (IJRR)*, vol. 32, no. 13, pp. 1585–1606, Oct. 2013, DOI:10.1177/0278364913498122.

[29] G. F. Franklin, J. D. Powell, and A. Emami-Naeini, *Feedback Control of Dynamic Systems*, 7th ed. Upper Saddle River, USA: Prentice Hall Press, 2014.

[30] F. Stulp and O. Sigaud, "Many regression algorithms, one unified model: A review," *Neural Networks*, vol. 69, pp. 60–79, Sep. 2015, DOI:10.1016/j.neunet.2015.05.005.

[31] C. Semini, N. G. Tsagarakis, E. Guglielmino, M. Focchi, F. Cannella, and D. G. Caldwell, "Design of HyQ a hydraulically and electrically actuated quadruped robot," *Proc. Inst. Mech. Eng. Part I J. Syst. Control Eng.*, vol. 225, no. 6, pp. 831–849, 2011, DOI:10.1177/0959651811402275.

[32] S. Nobili, M. Camurri, V. Barasuol, M. Focchi, D. Caldwell, C. Semini, and M. Fallon, "Heterogeneous sensor fusion for accurate state estimation of dynamic legged robots," Cambridge, USA, Jul. 2017, pp. 1–9, DOI:10.15607/RSS.2017.XIII.007.

[33] M. Focchi, R. Orsolino, M. Camurri, V. Barasuol, C. Mastalli, D. G. Caldwell, and C. Semini, *Heuristic Planning for Rough Terrain Locomotion in Presence of External Disturbances and Variable Perception Quality*. Springer, 2018, vol. 132.

[34] M. Focchi, G. A. Medrano-Cerda, T. Boaventura, M. Frigerio, C. Semini, J. Buchli, and D. . Caldwell, "Robot impedance control and passivity analysis with inner torque and velocity feedback loops," *Control Theory Technol.*, vol. 14, no. 2, pp. 97–112, May 2016, DOI:10.1007/s11768-016-5015-z.

[35] T. Hulin, A. Albu-Schffer, and G. Hirzinger, "Passivity and stability boundaries for haptic systems with time delay," *IEEE Trans. Control Syst. Technol.*, vol. 22, no. 4, pp. 1297–1309, Jul. 2014, DOI:10.1109/TCST.2013.2283372.

[36] M. Mosadeghzad, G. A. Medrano-Cerda, N. Tsagarakis, and D. G. Caldwell, "Impedance control with inner pi torque loop: Disturbance attenuation and impedance emulation," in *Proc. Int. Conf. Robot. Biomimetics (ROBIO)*, Shenzhen, China, Dec. 2013, pp. 1497–1502.

[37] R. Smith *et al.* (2005) Open dynamics engine. [Online]. Available: http://www.ode.org/

[38] E. Catto, "Soft constraints - reinventing the spring," in *Proc. Game Developers Conf. (CDC)*, 2011. [Online]. Available: http://box2d.org/files/GDC2011/GDC2011_Catto_Erin_Soft_Constraints.pdf

[39] T. Erez, Y. Tassa, and E. Todorov, "Simulation tools for model-based robotics: Comparison of Bullet, Havok, MuJoCo, ODE and PhysX," in *Proc. IEEE Int. Conf. Robot. Automat. (ICRA)*, Seattle, USA, May 2015, pp. 4397–4404, DOI:10.1109/ICRA.2015.7139807.

Shamel Fahmi (S'19) was born in Cairo, Egypt. He received the B.Sc. degree in mechatronics from the German University in Cairo, Egypt, in 2015, and the M.Sc. degree in systems and control from the University of Twente, the Netherlands, in 2017. In December 2017, he joined the Dynamic Legged Systems (DLS) lab at Istituto Italiano di Tecnologia (IIT) for the Ph.D. degree. His research interests include robotics, controls, optimization, and learning for dynamical systems.

Michele Focchi received the B.Sc. degree and the M.Sc. degree in Control System Engineering from Politecnico di Milano. After gaining some R&D experience in the industry, in 2009, he joined Istituto Italiano di Tecnologia (IIT) where he developed a micro-turbine for which he obtained an international patent. He received the Ph.D. degree in robotics, getting involved in the Hydraulically-actuated Quadruped (HyQ) robot project, in 2013. He is currently a Researcher at the Dynamic Legged Systems (DLS) lab at IIT. His research interests focus on pushing the performances of quadruped robots in traversing unstructured environments, by using optimization-based planning strategies to perform dynamic planning.

Andreea Radulescu received the B.Sc. degree in Engineering in Automatic Control and Applied Informatics from the Polytechnic University of Bucharest, Romania. She received her M.Sc. degree and Ph.D degree in Intelligent Robotics from the University of Edinburgh, Scotland, in 2011 and 2016, respectively. From 2016 to 2019 she was a postdoc researcher in the Dynamic Legged Systems (DLS) lab at Istituto Italiano di Tecnologia (IIT). Her research interests include optimal control, planning, machine learning and using variable impedance actuators for systems in domains with contacts. She is currently a robotics research engineer at Dyson Technology Ltd.

Geoff Fink (S'08-M'18) received the B.Sc. degree in computer engineering from the University of Alberta, Canada, in 2007, the M.Sc. degree in computer and electrical engineering from the University of Guadalajara, Mexico, in 2011, and the Ph.D. degree in control systems from the University of Alberta, Canada, in 2018. In 2018 he joined the the Dynamic Legged Systems (DLS) lab at Istituto Italiano di Tecnologia (IIT) as a postdoc researcher. His research interests include robotics, sensing, perception, state estimation and SLAM.

Victor Barasuol was born in São Miguel do Oeste/SC, Brazil. He received the Diploma in electrical engineering from Universidade do Estado de Santa Catarina (UDESC) in 2006. He received the M.Sc. degree in electrical engineering and the Ph.D degree in automation and systems engineering from Universidade Federal de Santa Catarina (UFSC) in 2008 and 2013, respectively. He is currently a researcher at the Dynamic Legged Systems (DLS) lab at Istituto Italiano di Tecnologia (IIT). He is an expert in motion generation and control for quadruped robots with emphasis in reactive actions using proprioceptive and exteroceptive sensor feedback.

Claudio Semini (S'07-M'10) received the M.Sc. degree in electrical engineering and information technology from ETH, Zurich, Switzerland, in 2005, and the Ph.D. degree from Istituto Italiano di Tecnologia (IIT), Genoa, Italy, in 2010. From 2004 to 2006, he visited the Hirose Laboratory at Tokyo Tech, and worked at Toshibas R&D Center, Japan. From 2007 to 2010, during his doctorate, he developed the hydraulic quadruped robot HyQ and worked on its control. From 2010 to 2012, he was a postdoc at the same department. He is currently a tenure-track researcher and head of the Dynamic Legged Systems (DLS) Laboratory at IIT. His research interests include the development and control of versatile legged robots for real-world environments. He is a co-founder and co-chair of the IEEE RAS Technical Committee of Robot Mechanisms and Design.