



HAL
open science

First Order Approximation of Model Predictive Control Solutions for High Frequency Feedback

Ewen Louis Dantec, Michel Taix, Nicolas Mansard

► **To cite this version:**

Ewen Louis Dantec, Michel Taix, Nicolas Mansard. First Order Approximation of Model Predictive Control Solutions for High Frequency Feedback. *IEEE Robotics and Automation Letters*, 2022, 7 (2), pp.4448 - 4455. 10.1109/LRA.2022.3149573 . hal-03419712v1

HAL Id: hal-03419712

<https://hal.science/hal-03419712v1>

Submitted on 8 Nov 2021 (v1), last revised 2 Feb 2022 (v2)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

First Order Approximation of Model Predictive Control Solutions for High Frequency Feedback

Ewen Dantec^{a,b,*}, Michel Taix^a, Nicolas Mansard^{a,b}

Abstract— The lack of computational power on mobile robots is a well-known challenge when it comes to implementing a real-time MPC scheme to perform complex motions. Currently the best solvers are barely able to reach 100Hz for computing the control of a whole-body legged model, while modern robots are expecting new torque references in less than 1ms. This problem is usually tackled by using a handcrafted low-level tracking control whose inputs are the low-frequency trajectory computed by the MPC. We show that a linear state feedback controller naturally arises from the optimal control formulation and can be used directly in the low-level control loop along with other sensitivities of relevant time-varying parameters of the problem. When the optimal control problem is solved by DDP, this linear controller can be computed for cheap as a by-product of the backward pass, and corresponds in part to the classical Riccati gains. A side effect of our proposition is to show that Riccati gains are valuable assets that must be used to achieve an efficient control and that they are not stiffer than the optimal control scheme itself. We propose a complete implementation of this idea on a full-scale humanoid robot and demonstrate its importance with real experiments on the robot Talos.

I. INTRODUCTION

Model-predictive control (MPC) amounts to frequently optimizing a prediction of the robot motion and executing the beginning of the resulting trajectory while waiting for a prediction update [1]. It has lately become popular in robotics to control complex multi-variable systems such as legged or aerial robots. MPC offers a viable method to solve large-scale optimal control problem (OCP) over a receding finite horizon, while efficiently dealing with constraints, non-linearity, model uncertainties and unforeseen disturbances [2], [3]. Historically, MPC was introduced to handle constrained, multi-variable industrial continuous processes which can be transcribed into linear models [4], [5], as found in the chemical or refining industry. Over the time, the ever increasing computational power of our computers has allowed the MPC paradigm to be applied to more complex systems such as 4-legged robots [6], collaborative manipulation involving dynamic obstacle [7] or quadrotor control [8].

MPC belongs to the class of trajectory optimization methods, where the trajectory planning problem is interpreted as an OCP. This approach allows to optimize simultaneously a task-related path and the associated control policy, while taking into account constraints and limitations [9], [10]. The tasks are classically described by a non-linear cost function




Fig. 1. The full scale humanoid Talos, used in the experiments, expects new torque references every 0.5ms. While such high frequencies is needed to unlock torque control, current MPC solvers cannot yet solve a whole-body problem at such a frequency. Another feedback must then be implemented, here to handle measurements of the robot state and target location.

which enforces constraints and defines the objectives hierarchy through cost weights. At each control cycle, starting from a state estimate of the system, an OCP is solved over a receding finite-time horizon and the first elements of the resulting trajectories are forwarded to a low-level execution process. The problem is then shifted from one sampling period before being solved again. Although this method works well for linear problems, a series of challenges arises when considering non-linear systems and complex tasks such as obstacle avoidance or contact planning. Since complexity varies with the cube of the system dimension, finding a global minimum in real time to a large-scale non-linear problem is very difficult to achieve [11]. For these reasons, the use of MPC in robotics has been, until recently, limited to systems with reduced dynamics [12], [13] or small scale [14], [15], [16].

Among the various numerical optimization schemes used to implement MPC, direct shooting [17] is a popular class of methods used in robotics. One of the most famous shooting algorithms is Differential Dynamic Programming (DDP) [18], [19], an efficient method featuring superlinear convergence rate with linear complexity in the horizon length, hence its success in robotics [20], [21], [22]. While it is originally formulated as a single shooting algorithm, it can be turned into a multiple shooting one [23], [17] at no extra cost [24], as we will recall later. This makes it easier to warm start the solver, either using the solution found at the previous control cycle [25], or deduced from an offline database of optimal solutions [26].

In order to tackle brutal disturbances and model uncertainties, it seems desirable to run the MPC at the highest

^a LAAS-CNRS, Université de Toulouse, CNRS, Toulouse, France

^b Artificial and Natural Intelligence Toulouse Institute, France

*corresponding author: edantec@laas.fr

possible frequency [27], [28], ideally around 1kHz. Direct MPC feedback at slower frequency is possible but drastically reduces the performances of the movements [29]. For large-scale systems such as legged robots, this may not be achievable despite some on-going exploratory works, e.g. contact-aware Linear Quadratic Regulator (LQR) controllers [30], parallelized LQR resolution [31] or multi-core GPU optimization solver [32]. Other works propose to use feedback MPC to provide an optimal feedback policy at a higher update rate than the re-computation of the whole problem. Those policies are either computed separately using another optimal control transcription [33] or based on an efficient Newton-type scheme [23]. In [34], a MPC scheme is deployed on a reduced model to rapidly generate a CoM trajectory and contact sequence, which are provided to a whole-body control block based on inverse dynamics. In [35], the feedback gains satisfy by design the contact constraints and dynamics constraints, and efficiently close the gap between the low-frequency high-level MPC and the low-level controller. However, both methods need to rely on a dedicated controller to produce a torque policy based on a state trajectory, bringing on an additional layer of complexity which should be fine-tuned.

It is known that the control solution of a constraint LQR is a continuous piece-wise affine function of the initial state [36]. Following this observation, we propose to deduce the optimal feedback policy from the OCP formulation and we show the connection between this policy and the DDP-induced Riccati gains. The important step is to demonstrate that Riccati gains can be interpreted as the sensitivities of the optimal solution with respect to the initial state of the OCP, and as such can be used to interpolate the optimal control between two DDP computations. Sensitivities have already been used to improve the warm-start in a receding horizon scheme [23] or for bilevel optimization [37]. Differentiating an MPC has also been proposed in the purpose of adding it as a layer of a neural network [38]. While we are only interested in getting the derivative of the initial control with respect to any external parameter, our proposed derivation leads to a more efficient backward pass to compute the corresponding gain, and elegantly connects it with the Riccati matrix. This is necessary as we are proposing to use it for high-frequency control. For small state changes, this is equivalent to computing a new solution of the MPC at a much higher frequency, reaching the expected computation of the low-level control. The same approach can be generalized to any parameter of the OCP, leading to a set of Riccati-like feedback gains which can be computed at low cost by extending the classical DDP backward pass. To illustrate our method, we performed a reactive reaching task involving 28 degrees of freedom, contacts and balance over a frequency rate of 100 Hz for the MPC, optimally interpolated at 2 kHz. We provide a demonstration of our whole-body MPC on the torque-controlled humanoid robot TALOS [39].

The paper is structured as follows. Sec. II introduces the basic formulation of optimal control for whole body motion in contact. Sec. III proves that Riccati gains can

be interpreted as sensitivity of the problem with respect to the initial state, and gives a generic method to extract other sensitivities. Sec. IV explains the implementation of these gains in the low-level control of our scheme. Sec. V presents the experimental results on the humanoid robot TALOS.

II. WHOLE BODY OPTIMAL CONTROL WITH CONTACT

A. Contact-Constrained Rigid body dynamics

Similarly to [40] and [21], we consider a rigid body system with n_j joints and n_p rigid contacts with the environment and model the dynamics of this constrained multi-body system through the following equation:

$$\begin{bmatrix} \mathbf{M} & \mathbf{J}_c^\top \\ \mathbf{J}_c & \mathbf{0} \end{bmatrix} \begin{bmatrix} \ddot{\mathbf{q}} \\ -\boldsymbol{\lambda} \end{bmatrix} = \begin{bmatrix} \mathbf{S}^\top \boldsymbol{\tau} - \mathbf{b} \\ -\mathbf{J}_c \dot{\mathbf{q}} \end{bmatrix} \quad (1)$$

where \mathbf{M} is the joint-space inertia matrix, \mathbf{b} the generalized non-linear forces, $\mathbf{q} \in SE(3) \times \mathbb{R}^{n_j}$ the configuration vector which accounts for the n_j actuated joints positions and the free-flyer joint, $\dot{\mathbf{q}}$ the velocity vector laying in the tangent space of $SE(3) \times \mathbb{R}^{n_j}$, $\ddot{\mathbf{q}}$ the acceleration vector, \mathbf{S} the actuation matrix typically selecting the actuated joints, $\boldsymbol{\tau}$ the joint torques, $\boldsymbol{\lambda} = [\boldsymbol{\lambda}_1 \cdots \boldsymbol{\lambda}_{n_p}]^\top$ and $\mathbf{J}_c = [\mathbf{J}_1 \cdots \mathbf{J}_{n_p}]^\top$ the concatenation of vectors of contact forces and concatenation of contact Jacobian matrices. In this context, forces $\boldsymbol{\lambda}$ abstractly represents either 3D forces for punctual contacts or spatial 6D forces for planar contacts, expressed in their respective contact frame. They have to respect the contact model described by the cone K_p :

$$\forall p = 0..n_p, \boldsymbol{\lambda}_p \in K_p. \quad (2)$$

We then classically consider the state of the robot to be $\mathbf{x} = (\mathbf{q}, \dot{\mathbf{q}}) \in \mathbb{R}^{n_x}$, $\dot{\mathbf{x}} = (\dot{\mathbf{q}}, \ddot{\mathbf{q}}) \in \mathbb{R}^{n_{dx}}$, and the control to be $\mathbf{u} = \boldsymbol{\tau} \in \mathbb{R}^{n_u}$. In [21], we showed that with this formulation, forces become consequences of the state $(\mathbf{q}, \dot{\mathbf{q}})$ and the control \mathbf{u} , and thus our dynamics becomes independent from these variables. Consequently, (1) leads to a force-free partial derivative equation which we numerically integrate into a next-state equation:

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t) \quad (3)$$

B. Optimal control for legged robots

Based on (3), we can formulate a continuous OCP then transcript it into a nonlinear program in order to be solved. We use a multiple-shooting formulation, as it combines the efficiency of shooting formulations [41] with numerical stability [17]. Our discretized OCP then takes the following form:

$$\begin{aligned} \min_{\underline{\mathbf{x}}, \underline{\mathbf{u}}} & \sum_{t=0}^{T-1} \ell(\mathbf{x}_t, \mathbf{u}_t, t) + \ell_T(\mathbf{x}_T) \\ \text{s.t.} & \quad \mathbf{x}_0 = \mathbf{f}_0 \\ & \quad \forall t = 0..T, \quad \mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t) \end{aligned} \quad (4)$$

where the decision variables are the state $\underline{\mathbf{x}} = (\mathbf{x}_t)_{t=0..T}$ and control $\underline{\mathbf{u}} = (\mathbf{u}_t)_{t=0..T-1}$ trajectories, T is the number of knots, ℓ and ℓ_T are the running and terminal cost functions, used to define the tasks the robot has to perform (goal

tracking, center of mass tracking, etc.) and to regularize state and control trajectories, and \mathbf{f}_0 is the initial state. An OCP for locomotion often contains admissibility constraints on state and control; however, as recent progresses in handling hard constraints for large OCP [20], [42], [43] still suffer from practical limitations in the perspective of real robot experiments, we focus here on constraint-free OCP (without loss of generality for our theoretical developments [38], but targeting our experimental setup). Thus, we formulate these constraints as penalties of various forms in our cost function, and describe them in Section IV.

C. Formulating the LQR

Our LQR formulation is obtained by linearizing (4) around an initial trajectory $(\underline{\mathbf{x}}, \underline{\mathbf{u}})$:

$$\begin{aligned} \min_{\Delta \mathbf{x}, \Delta \mathbf{u}} & \left(\sum_{t=0}^{T-1} \frac{1}{2} [\Delta \mathbf{x}_t^\top, \Delta \mathbf{u}_t^\top] \begin{bmatrix} L_{xx} & L_{xu} \\ L_{ux} & L_{uu} \end{bmatrix} \begin{bmatrix} \Delta \mathbf{x}_t \\ \Delta \mathbf{u}_t \end{bmatrix} \right. \\ & \left. + [\ell_x \quad \ell_u] \begin{bmatrix} \Delta \mathbf{x}_T \\ \Delta \mathbf{u}_T \end{bmatrix} + \frac{1}{2} \Delta \mathbf{x}_T^\top L_{xx} \Delta \mathbf{x}_T + \ell_x \Delta \mathbf{x}_T \right) \quad (5) \\ \text{s.t.} \quad & \Delta \mathbf{x}_0 = \mathbf{f}_0 \\ & \forall t = \{0, \dots, T-1\}, \\ & \Delta \mathbf{x}_{t+1} = \mathbf{F}_x \Delta \mathbf{x}_t + \mathbf{F}_u \Delta \mathbf{u}_t + \mathbf{f}_{t+1} \end{aligned}$$

For the sake of readability, we will denote $A_b = \frac{\partial A}{\partial b}$, except when not suitable (e.g. $\frac{\partial \Delta u}{\partial x}$). In (5), ℓ_x, ℓ_u and L_{xx}, L_{xu}, L_{uu} are the gradients and Hessians of the cost function (indices t have been dropped to simplify derivative notations). Similarly, F_x and F_u are the Jacobians of the dynamics computed at each (x_t, u_t) . We write \mathbf{f}_t the drift in dynamics, which represents the change in state when the control is 0. Note that \mathbf{f}_t is typically omitted in DDP, but is introduced here to handle our multiple shooting transcription.

The optimum of (5) is characterized by the gradients of the associated Lagrangian vanishing. This condition can be formulated as a linear equality called Karush-Kuhn-Tucker (KKT) condition, which binds the primal variables $\Delta \mathbf{z} = (\Delta \mathbf{x}, \Delta \mathbf{u})$ to the dual Lagrange multipliers $\boldsymbol{\mu} = (\boldsymbol{\mu}_0, \dots, \boldsymbol{\mu}_T)$ associated with the dynamics constraints:

$$\begin{bmatrix} \mathcal{L}_{zz} & \mathbf{h}_z^\top \\ \mathbf{h}_z & 0 \end{bmatrix} \begin{bmatrix} \Delta \mathbf{z} \\ \boldsymbol{\mu} \end{bmatrix} = \begin{bmatrix} -\mathcal{L}_z \\ \mathbf{g} \end{bmatrix} \quad (6)$$

where the notations adopted here will become clear when each block will be properly defined in the next section. While we will later exploit the KKT equation to formulate our contribution, it is in general not easy to efficiently invert the KKT matrix to solve the OCP.

D. Solving the OCP using the KKT conditions

DDP solves (5) by using Bellman's principle [44], which naturally exploits the sparsity of the Markovian nature of the dynamics constraints. Eventually, each descent step can be produced by an iterative scheme which pre-computes the descent direction in a backward pass (going from N to 0), and then finds the optimal step by line search along this direction in a forward pass (going from 0 to N) [45]. In this

forward pass, the optimal state and control $(\Delta \mathbf{x}_t, \Delta \mathbf{u}_t)$ and optimal dual variable $\boldsymbol{\mu}_t \forall t = 0..T-1$ can be expressed as:

$$\begin{aligned} \Delta \mathbf{u}_t &= \mathbf{K}_t \Delta \mathbf{x}_t + \mathbf{k}_t \\ \boldsymbol{\mu}_{t+1} &= \mathbf{V}_{t+1} (\mathbf{F}_x \Delta \mathbf{x}_t + \mathbf{F}_u \Delta \mathbf{u}_t) + \mathbf{v}_{t+1} \\ \mathbf{x}_{t+1} &= f(\mathbf{x}_t, \mathbf{u}_t) \end{aligned} \quad (7)$$

Here the gains and value function are computed during the backward pass of the DDP: starting from $\mathbf{V}_T = L_{xx}$ and $\mathbf{v}_T = \ell_x + L_{xx} \mathbf{f}_T$, we have, $\forall t = 0..T-1$:

$$\begin{aligned} \mathbf{K}_t &= -\mathbf{Q}_{uu}^{-1} \mathbf{Q}_{ux} \\ \mathbf{k}_t &= -\mathbf{Q}_{uu}^{-1} \mathbf{q}_u \\ \mathbf{V}_t &= \mathbf{Q}_{xx} - \mathbf{Q}_{xu} \mathbf{K}_t \\ \mathbf{v}_t &= \mathbf{q}_x - \mathbf{Q}_{xu} \mathbf{k}_t + \mathbf{V}_{t+1} \mathbf{f}_t \end{aligned} \quad (8)$$

The backward pass also computes the derivative of the quality function:

$$\begin{aligned} \mathbf{q}_x &= \ell_x + \mathbf{F}_x^\top \mathbf{v}_{t+1} \\ \mathbf{q}_u &= \ell_u + \mathbf{F}_u^\top \mathbf{v}_{t+1} \\ \mathbf{Q}_{xx} &= L_{xx} + \mathbf{F}_x^\top \mathbf{V}_{t+1} \mathbf{F}_x \\ \mathbf{Q}_{xu} &= L_{xu} + \mathbf{F}_x^\top \mathbf{V}_{t+1} \mathbf{F}_u \\ \mathbf{Q}_{uu} &= L_{uu} + \mathbf{F}_u^\top \mathbf{V}_{t+1} \mathbf{F}_u \end{aligned} \quad (9)$$

III. SENSITIVITY ANALYSIS OF THE OPTIMAL SOLUTION

Suppose that our OCP depends on a parameter $\boldsymbol{\theta}$ which varies with time. If the variation of $\boldsymbol{\theta}$ is non-negligible over the computation duration of the OCP, then the optimal solution is no longer correct since it has been computed over an outdated value of $\boldsymbol{\theta}$. Determining the sensitivity of the OCP with respect to $\boldsymbol{\theta}$ will allow us to approximate a control correction proportional to the variation of the parameter. Such approach, where we differentiate optimization problems, has already been studied for contact simulation [46] and neural network applied to QP [47].

A. Defining sensitivity as partial derivative of the OCP

Let us view the OCP described by (4) as a classical non-linear problem parametrized by a given $\boldsymbol{\theta} \in \mathbb{R}^{n_\theta}$. We define the optimal state and control trajectory $\underline{\mathbf{x}}^* = (\mathbf{x}_t^*)_{t=0..T}$ and $\underline{\mathbf{u}}^* = (\mathbf{u}_t^*)_{t=0..T-1}$. For the sake of simplicity we denote by $\mathbf{z}^* = (\underline{\mathbf{x}}^*, \underline{\mathbf{u}}^*) \in \mathbb{R}^{n_z}$ with $n_z = T n_u + (T+1) n_x$ solution of the following discretized problem:

$$\begin{aligned} \mathcal{P}(\boldsymbol{\theta}) = \mathbf{z}^* &= \min_{\mathbf{z}} \ell(\mathbf{z}, \boldsymbol{\theta}) \\ \text{s.t.} \quad & h(\mathbf{z}, \boldsymbol{\theta}) = 0 \end{aligned} \quad (10)$$

with $h : \mathbb{R}^{n_z} \times \mathbb{R}^{n_\theta} \rightarrow \mathbb{R}^{(T+1) \times n_x}$ the $T+1$ discretized dynamics constraints given by $\mathbf{x}_{t+1} = f(\mathbf{x}_t, \mathbf{u}_t) \forall t = 0..T-1$ and $\mathbf{x}_0 = \mathbf{f}_0$. We assume ℓ and h are continuously twice differentiable, and that the Mangasarian-Fromovitz constraint qualification holds for our nonlinear problem, in order to be able to apply the KKT optimality conditions. From here the Lagrangian of the problem writes $\mathcal{L}(\mathbf{z}, \boldsymbol{\mu}, \boldsymbol{\theta}) = \ell(\mathbf{z}, \boldsymbol{\theta}) + \boldsymbol{\mu}^\top h(\mathbf{z}, \boldsymbol{\theta})$ and the KKT optimality

conditions state that if z^* is solution of the problem (10) for θ , then:

$$\begin{aligned} \mathcal{L}_z(z^*, \mu, \theta) &= \ell_z(z^*, \theta) + \mu^T h_z(z^*, \theta) = 0 \\ h(z^*, \theta) &= 0 \end{aligned} \quad (11)$$

Finally we suppose that for each $\mu \in \mathbb{R}^{(T+1) \cdot n_x}$ satisfying the KKT conditions and each $\delta z \neq 0 \in \mathbb{R}^{n_z}$ so that $h_z(z^*, \theta)\delta z = 0$, we have the following inequality:

$$\delta z^T \mathcal{L}_{zz}(z^*, \mu, \theta) \delta z > 0 \quad (12)$$

Under these assumptions, Shapiro [48] showed that the function $\mathcal{P}(\cdot)$ is directionally differentiable and that for each $\theta \in \mathbb{R}^{n_\theta}$, $d \in \mathbb{R}^{n_\theta}$, there exists λ verifying the KKT conditions such that $\mathcal{P}'(\theta; d)$, derivative along the direction d , is the unique solution of the following quadratic program:

$$\begin{aligned} \min_{\delta z} \quad & \frac{1}{2} \delta z^T \mathcal{L}_{zz}(z, \mu, \theta) \delta z + \delta z^T \mathcal{L}_{z\theta}(z, \mu, \theta) d \\ \text{s.t.} \quad & h_z(z, \theta) \delta z = 0 \end{aligned} \quad (13)$$

Since in our case, the Lagrangian multiplier μ is unique due to the upper diagonal form of the constraints matrix, it turns out we can easily compute the directional derivative of our problem along the direction d :

$$\begin{bmatrix} \delta z \\ \delta \mu \end{bmatrix} = \begin{bmatrix} \mathcal{L}_{zz} & h_z^T \\ h_z & 0 \end{bmatrix}^{-1} \begin{bmatrix} -\mathcal{L}_{z\theta} d \\ 0 \end{bmatrix} \quad (14)$$

with $\delta \mu$ a Lagrangian multiplier increment we are not interested in.

One can recognize in (14) the inverse of the KKT matrix already met in (6), yet not multiplied here by $(-\mathcal{L}_z, g)$ but by $(-\mathcal{L}_{z\theta}, 0)$. As explained in Sec. II-D, DDP can be seen as an efficient way of computing the inverse of the KKT matrix; we then propose to extend the DDP to compute (14). This is equivalent to replacing the terms (ℓ_x, ℓ_u, f_t) in (8) and (9) with $(\mathcal{L}_{x\theta}, \mathcal{L}_{u\theta}, 0)$, then solving the backward pass as usual. Note here that only the quantities (q_x, q_u, v_t, k_t) are affected by this change:

$$\begin{aligned} k_\theta &= -Q_{uu}^{-1} Q_{u\theta} \\ v_\theta &= \mathcal{L}_{x\theta} \quad \text{for } t = T \\ v_\theta &= Q_{x\theta} - Q_{xu} k_\theta \\ Q_{x\theta} &= \mathcal{L}_{x\theta} + F_x^T v'_\theta \\ Q_{u\theta} &= \mathcal{L}_{u\theta} + F_u^T v'_\theta \end{aligned} \quad (15)$$

As before, t is dropped in index when it conflicts with another notation to improve readability. The value derivative v_θ at time $t+1$ is denoted v'_θ . Equation (7) then becomes:

$$\frac{\partial \Delta u_t}{\partial \theta} = K_t \frac{\partial \Delta x_t}{\partial \theta} + k_\theta \quad (16)$$

This gives us $\frac{\partial \Delta u_0}{\partial \theta}$, sensitivity of the first control of the linearized problem with respect to the parameter θ . Since the OCP (4) has the same derivatives as the LQR (5) at each

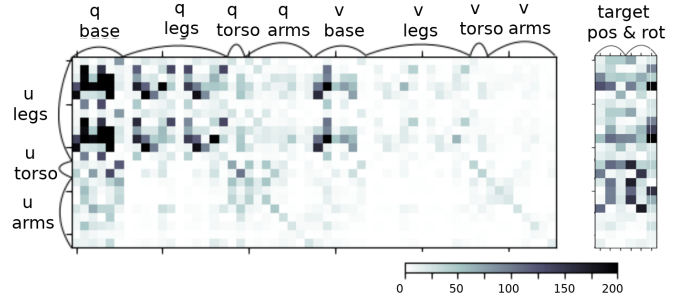


Fig. 2. Left: matrix of the absolute Riccati gain K_0 for a whole-body OCP involving end-effector tracking and biped stabilization. Right: matrix of the absolute placement (position and rotation) gains for a whole-body OCP involving tracking a reference placement with the left hand while stabilizing balance.

knot, the sensitivity of both problems are equal and we can write

$$\frac{\partial u_0^*}{\partial \theta} = \frac{\partial \Delta u_0^*}{\partial \theta} \quad (17)$$

B. Interpreting Riccati gains as sensitivity to initial state

A feedback policy can be obtained from (16) when θ is measured at a higher frequency than the DDP solves (4). In the experimental part of this paper, we demonstrate a visual feedback policy where θ is the position of an end-effector target. Another noticeable case of interest arises when θ is the initial robot state x_0 over which the DDP is computed.

Considering a second order approximation, the only derivatives of the problem which depend on x_0 are $\mathcal{L}_{x,0}$ and $\mathcal{L}_{u,0}$, so by recurrence with (15), it is straightforward to deduce that $v_\theta = 0 \forall t = 1..T$ and that $k_\theta = -Q_{uu}^{-1} \mathcal{L}_{u\theta} \forall t = 0..T$. The derivative of the first optimal control with respect to its initial state then simply writes:

$$\frac{\partial \Delta u_0}{\partial x_0} = K_0 \frac{\partial \Delta x_0}{\partial x_0} + k_\theta = K_0 + k_\theta \quad (18)$$

For problems where \mathcal{L}_{ux} is null at $t = 0$, (18) becomes $\frac{\partial \Delta u_0}{\partial x_0} = K_0$. As a consequence, the Riccati gains K_0 can be used as a feedback term to approximate the optimal control between two DDP computations, given that the current state of the robot is provided at a sufficiently high frequency. The resulting feedback policy can be seen as a first order Taylor development of the MPC: suppose that the MPC has been solved at a given observed state \hat{x} , denoted by $u = mpc(\hat{x})$. The linear feedback at a new observation \hat{x}' will then approximate the MPC solution:

$$mpc(\hat{x}') \approx mpc(\hat{x}) + K_0(\hat{x}' - \hat{x}) \quad (19)$$

It has been considered that the Riccati gains might be too stiff to be actually used for feedback [49]. Yet we see here that they are nothing more than a linear interpolation of the MPC feedback, and, for a sufficiently high frequency, they lead to a fair numerical approximation of the MPC, hence are not stiffer than the MPC itself.

The structure of a typical Riccati gain is presented in Fig. 2. It is interesting to note that the Riccati matrix presents a strong diagonal in position: the most contributing correction of a given joint torque strongly depends on the

position error of this joint. It can also be noted that the feedback correction associated to the base position of the robot is very high, and affects mainly the leg controls. This is expected since the legs are the main drivers of the robot base position. Finally, one can notice that the velocity feedback term of the Riccati matrix are small compared to the position feedback terms, which is reassuring since velocity estimates are typically more noisy, and LQR controllers tend to create brutal velocity feedbacks [30].

IV. MPC IMPLEMENTATION WITH RICCATI FEEDBACK

A. An OCP for reaching, balancing and switching contacts

The OCP formulated for the experiments is composed of five different costs:

- a state regularization cost:
 $\ell_1(\mathbf{x}) = (\mathbf{x} - \mathbf{x}_0)^T \mathbf{R}_x (\mathbf{x} - \mathbf{x}_0)$ with \mathbf{R}_x a positive definite weight matrix and \mathbf{x}_0 the default state;
- a control regularization cost:
 $\ell_2(\mathbf{u}) = (\mathbf{u} - \mathbf{u}_0)^T \mathbf{R}_u (\mathbf{u} - \mathbf{u}_0)$ with \mathbf{R}_u a positive definite weight matrix and \mathbf{u}_0 the gravity-compensating torque in default state;
- a goal-tracking cost:
 $\ell_3(\mathbf{x}) = a(\mathbf{p}(\mathbf{x}) - \mathbf{p}_d)$ with $a : \mathbb{R}^3 \rightarrow \mathbb{R}$ a quadratic activation function, $\mathbf{p}(\mathbf{x})$ current end-effector position and \mathbf{p}_d desired end-effector position;
- a CoM tracking cost:
 $\ell_4(\mathbf{x}) = \|\mathbf{c}(\mathbf{x}) - \mathbf{c}_d\|$ with $\mathbf{c}(\mathbf{x})$ current CoM and \mathbf{c}_d desired CoM;
- a kinematic limit cost:
 $\ell_5(\mathbf{x}) = \|\max(\mathbf{x} - \mathbf{x}_u, \mathbf{0}) + \min(\mathbf{x} - \mathbf{x}_l, \mathbf{0})\|$ with \mathbf{x}_u the upper bound and \mathbf{x}_l the lower bound of the joints positions and velocities.

This OCP is solved until convergence before starting the motion. Then, a single DDP iteration is performed, using the previous solution as an initial warm-start, before iterating with the latest sensor measurements. Two quantities are estimated from sensors: the robot state \mathbf{x}_0 (from joint encoders and base IMU), and the tracking target position \mathbf{p}_d (from motion capture camera). This formulation is sufficient to ensure stable reaching motions with both feet in contact and to reject external disturbances in real time.

B. Low-level control

The low-level torque control is composed of a proportional-derivative feedback on the joint torque measurement, plus a feedforward term which compensates for the intrinsic dynamics of the joints which are not considered in the whole-body model. Unexpected dynamics such as motor inertia, high frictions or inner flexibility of the harmonic-drive are thus handled by the low-level control. This way every joint behaves as an ideal joint from the point of view of the high-level control MPC.

C. Riccati interpolation of the MPC

Our DDP scheme produces the optimal control along with the sensitivities associated with the initial state and desired position of the end-effector. The first sensitivity $\mathbf{K}_0 = \frac{\partial \mathbf{u}_0}{\partial \mathbf{x}_0}$ has been introduced in Sec. III-B and is directly obtained

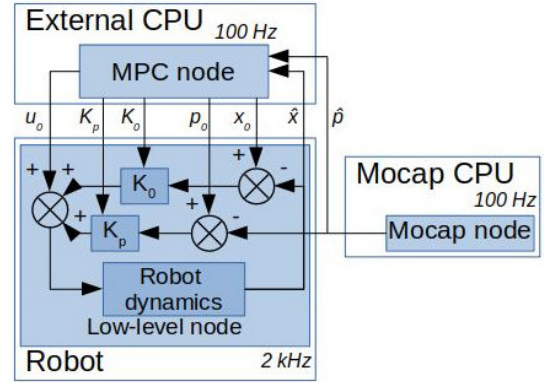


Fig. 3. Diagram of the ROS implementation of the Riccati feedback MPC.

during the backward pass of the DDP. The second sensitivity $\mathbf{K}_p = \frac{\partial \mathbf{u}_0}{\partial \mathbf{p}_d}$ is obtained by noting that the only cost which depends on the target position is the goal-tracking cost $\ell_3(\mathbf{x})$. The only non-zero derivative of the Lagrangian gradient then writes:

$$\mathcal{L}_{\mathbf{x}\mathbf{p},t} = -\mathbf{J}_{ee}^T \mathbf{A} \quad (20)$$

with $\mathbf{J}_{ee}(\mathbf{x})$ Jacobian of the end-effector when the state is \mathbf{x} and \mathbf{A} Hessian of the activation function. Re-injecting in (15), we obtained the sensitivity \mathbf{K}_p whose structure is presented Fig. 2. The optimal torque sent to the low-level control finally writes:

$$\mathbf{u}^* = \mathbf{u}_0 + \mathbf{K}_0(\hat{\mathbf{x}} - \mathbf{x}_0) + \mathbf{K}_p(\hat{\mathbf{p}} - \mathbf{p}_0) \quad (21)$$

with \mathbf{u}_0 , \mathbf{K}_0 , \mathbf{K}_p the feedforward, state feedback and target feedback terms updated at about 60 Hz, $\hat{\mathbf{x}}$ the measured state updated at 2 kHz, $\hat{\mathbf{p}}$ the current target position updated at 100Hz, and \mathbf{x}_0 , \mathbf{p}_0 the initial state and target position at the time of the last DDP computation.

D. ROS architecture

The architecture is illustrated in Fig. 3, and is composed of two parallel processes running on independent CPUs: one for the MPC, the other for the low-level torque control. In addition, a motion capture system (mocap) measures at 100Hz the position of the target to be reached and sends it to the other two nodes.

In order to synchronise the different frequencies of our processes, a ROS publisher-subscriber architecture is used to build the communication framework between all running nodes. Every ROS node uses the latest available data on the topics it subscribes to, and publishes the output once computation is done.

V. EXPERIMENTAL RESULTS

A. Disturbances experiment

B. Experimental setup

Our experimental setup aims at proving two hypothesis:

- 1) The Riccati gains introduced in Sec. (III-B) play the role of a state-feedback term and contribute to improve the overall behaviour of the MPC by approximating




Fig. 4. Whole-body tracking experiment without Riccati gains. After 4 seconds, the desired torque oscillations become too large, and the securities of the robot are triggered, shutting down all motors.

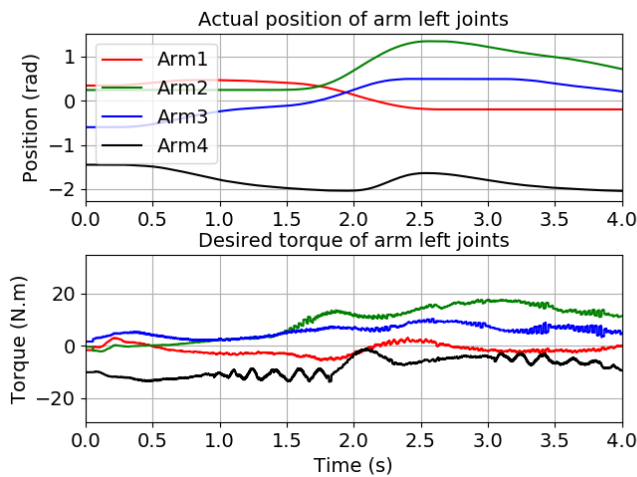


Fig. 5. Whole-body tracking experiment with Riccati gains. Although some oscillations remain, the overall tracking behaviour is good.

the optimal control policy at the low-level control frequency;

- 2) Other sensitivities can be used as well inside the low-level control in order to adapt the solution to a fast variation of one of the parameters of the problem, or even to replace the MPC node in case it stops working.

In order to prove these hypotheses, three different experimental protocols have been set up on the TALOS robot [39], in which the MPC node is running on a powerful external computer (AMD Ryzen 5950X, 32 cores and 4.9GHz with 64 Go of RAM) whereas the low-level control node is running on the robot internal computer. We used Pinocchio to enable fast computations of costs, dynamics and their derivatives [50], [51]. Two of the experiments are based on the same whole-body tracking task, which requires the robot to reach a moving target with the end-effector, here its left hand, while standing on the ground and keeping its balance. We used 22 joints in our model (wrists and neck are kept fixed for practical reasons), plus the free flyer state which is




Fig. 6. Disturbance experiment: the MPC is shut down and external disturbances are manually applied.

estimated by an observer provided by the manufacturer. The tracking weights have been selected to be high as compared to others gains, in order to make the problem dynamically challenging to solve in real time.

Our companion video, also available at <https://link.infini.fr/riccatifedbackantec>, illustrates these experiments as well as other dynamic motions made possible by the use of Riccati interpolation.

C. Riccati feedback experiment

To validate hypothesis 1), the whole-body tracking task was performed with and without the use of Riccati gains in the low-level control loop. For this experiment, the target p_d is not estimated from sensor but set to an arbitrary sinusoidal trajectory. Fig. 4 shows that for this particular task, the MPC alone produces an unstable control trajectory which will eventually trigger the robot inner securities on torque and velocity. On the contrary, Fig. 5 illustrates the good control trajectory computed by the MPC node and interpolated between each control cycle using the Riccati gains. In this example the oscillatory behaviour remains limited and the task is successfully performed.

Our second experimental protocol aims at showing that the Riccati gains alone produce a decent feedback control able to stabilize the robot under external disturbances. The experiment consists in launching the MPC with no goal-tracking cost, then shutting it down and applying external disturbances to the robot. The policy sent to the low-level control then becomes $u^* = u_0 + K_0(\hat{x} - x_0)$, with \hat{x} being the only varying quantity. This policy is an approximation of the optimal stabilizing policy around x_0 . Here (u_0, K_0, x_0) are drawn from the last DDP computation of the high-level control. As can be seen in Fig. 6, the Riccati gains efficiently reject the perturbations and keep the robot in a default stabilizing position. While our goal is not to promote a purely linear feedback, this experiment illustrates how large the stable domain of our linear policy can be.

D. Placement feedback experiment

Our third experimental protocol aims at implementing feedback policy on the target position, as described in Sec. IV-C. This time the mocap system is used to provide

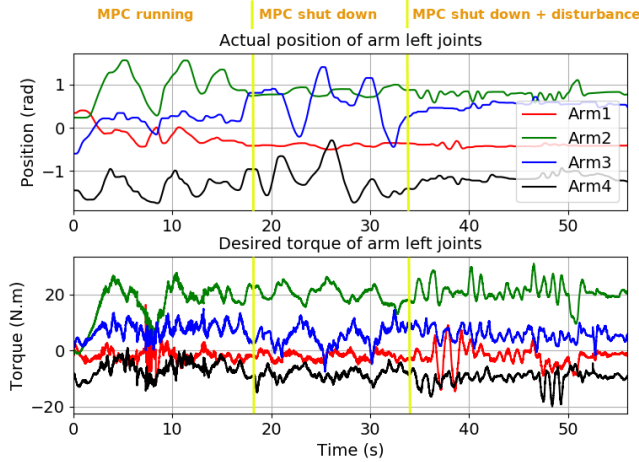


Fig. 7. Placement feedback experiment state and control plot: MPC is shut down at $t = 18s$, and external disturbances start to be applied at $t = 33s$.

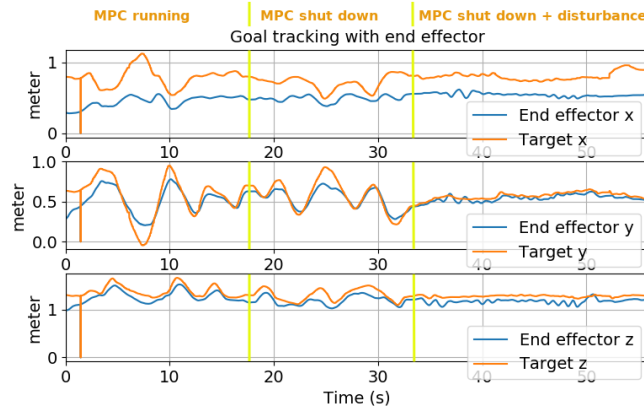


Fig. 8. Placement feedback experiment tracking plot.

the desired placement of the end-effector at a frequency of approximately 100 Hz. Given that this frequency is similar to the MPC node frequency, the target feedback is small compared to the feedforward term u_0^* . To better observe the feedback effect, the MPC has been intentionally shut down so that the drift in desired end-effector placement becomes significant over time.

The experiment is composed of three different parts: first, the MPC node is started and the OCP is solved at approximately 60 Hz. Then, at $t = 18s$, we shut down the MPC but keep the target moving to illustrate the effect of the placement feedback gains. From this moment the control is only due to state and target feedbacks. Finally, at $t = 33s$, the target is kept fixed and external disturbances are applied on the robot. Using the sensitivities as an approximation for the optimal solution is relevant as long as the current OCP, defined by the current state and target position, does not vary too much from the OCP computed just before shutting down the MPC. The architecture of this scheme is the same as in Fig. 3, without the MPC node to update the terms u_0, p_0, x_0 .

The resulting plots of this experiment are presented in Fig 7 and Fig. 8. As expected, switching off the MPC and relying only on the Riccati gains force the robot state to remain in the vicinity of the initial state used by the MPC node to compute the last OCP solution. As a consequence,




Fig. 9. Placement feedback experiment time plot.

the joints 1 and 2 are barely moving after switching off the MPC, because their inertial load is higher than the two other arm joints. Nevertheless, placement gains provide a coarse approximation of the optimal control and are sufficient to move the end effector toward the desired target and reject external perturbations, as observed in Fig. 8.

Fig. 9 provides the computation cost of the DDP and the extra loop computing K_p . As expected, K_p is much cheaper to compute as it always corresponds to back-propagating only 6 columns, as opposed to the classical backward pass which works with the n_x columns of K .

VI. CONCLUSION

This paper highlights the idea that the classical feedback gains obtained during the Riccati recursion of the DDP can be used as sensitivities with respect to the initial state of the OCP. These gains allow to interpolate the optimal control at 2kHz in order to produce dynamical yet stable motions which would otherwise fail because of the MPC computation time delay. This has been illustrated on a whole-body tracking movement on the full scale humanoid robot Talos. Additionally, other sensitivities can be computed and used at 2kHz in order to adapt the control to a rapidly changing parameter, e.g. the desired end-effector position. This is especially useful to deduce at high frequency a good approximation of the optimal policy when small perturbations are observed by fast sensors. This work paves the way for the implementation of a real-time whole-body torque MPC with dynamic contacts, able to perform challenging tasks like walking or interacting with the environment, at arbitrary high feedback frequency. Theoretical extensions to hand constraints and contact-invariant formulation and experimental extensions to dynamic locomotion will be considered in the future.

REFERENCES

- [1] J. Rawlings, D. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation and Design*. Nob Hill Publishing, 2017.
- [2] T. A. Johansen, "Chapter 1 introduction to nonlinear model predictive control and moving horizon estimation," 2011.
- [3] M. Huba, S. Skogestad, M. Fikar, M. Hovd, T. A. Johansen, and B. Rohal-Ilkiv, "Selected topics on constrained and nonlinear controls," *STU Bratislava – NTNU Trondheim*, 2011.

- [4] M. L. Darby, "Industrial mpc of continuous processes," In: *Baillieul J., Samad T. (eds) Encyclopedia of Systems and Control*, 2013.
- [5] E. Fernandez-Camacho and C. Bordons-Alba, *Model Predictive Control in the Process Industry*. Springer, London, 1995.
- [6] F. Farshidian, E. Jelavic, A. Satapathy, M. Gifithaler, and J. Buchli, "Real-time motion planning of legged robots: A model predictive control approach," *IEEE-RAS 17th Humanoids*, 2017.
- [7] M. Krämer, C. Rösmann, F. Hoffmann, and T. Bertram, "Model predictive control of a collaborative manipulator considering dynamic obstacles," *Optim Control Appl Meth.* 2020;41:1211–1232., 2020.
- [8] K. Alexis, C. Papachristos, G. Nikolakopoulos, and A. Tzes, "Model predictive quadrotor indoor position control," *19th Mediterranean Conference on Control Automation (MED)*, p. 1247–1252, 2011.
- [9] J. Rawlings, E. Meadows, and K. Muske, "Nonlinear model predictive control: A tutorial and survey," *IFAC Advanced Control of Chemical Processes*, 1994.
- [10] R. Findeisen and F. Allgöwer, "An introduction to nonlinear MPC," *Benelux Meeting on Systems and Control*, 2020.
- [11] S. J. Qin and T. A. Badgwell, "An overview of nonlinear model predictive control applications," *Nonlinear Predictive Control, Springer-Verlag*, pp 369—392, 2000.
- [12] S. Kajita, F. Kanehiro, K. Kaneko, K. Fujiwara, K. Harada, K. Yokoi, and H. Hirukawa, "Biped walking pattern generation by using preview control of zero-moment point," *IEEE ICRA*, 2003.
- [13] D. Kim, J. D. Carlo, B. Katz, G. Bledt, and S. Kim, "Highly dynamic quadruped locomotion via whole-body impulse control and model predictive control," *arXiv preprint:1909.06586*, 2019.
- [14] M. Neunert, C. De Crousaz, F. Furrer, M. Kamel, F. Farshidian, R. Siegwart, and J. Buchli, "Fast nonlinear model predictive control for unified trajectory optimization and tracking," in *IEEE ICRA*, 2016.
- [15] M. Geisert and N. Mansard, "Trajectory generation for quadrotor based systems using numerical optimal control," in *IEEE ICRA*, 2016.
- [16] B. Houska and M. Diehl, "Robustness and stability optimization of power generating kite systems in a periodic pumping mode," in *IEEE International Conference on Control Applications*, 2010.
- [17] H. G. Bock and K.-J. Plitt, "A multiple shooting algorithm for direct solution of optimal control problems," *IFAC Proceedings Volumes*, vol. 17, no. 2, pp. 1603–1608, 1984.
- [18] D. Q. Mayne, "Differential dynamic programming—a unified approach to the optimization of dynamic systems," *Control and Dynamics Systems*, vol. 10, pp. 179–254, 1973.
- [19] Y. Tassa, T. Erez, and W. D. Smart, "Receding horizon differential dynamic programming," *Advances in Neural Information Processing Systems*, vol. 20, p. 1465–1472, 2008.
- [20] T. A. Howell, B. E. Jackson, and Z. Manchester, "Altro: A fast solver for constrained trajectory optimization," *IEEE/RSJ IROS*, 2019.
- [21] R. Budhiraja, J. Carpentier, C. Mastalli, and N. Mansard, "Differential dynamic programming for multi-phase rigid contact dynamics," *IEEE-RAS Humanoids*, 2018.
- [22] B. Jackson, Z. Manchester, and F. Farshidian, "Software tools for real-time optimal control," *Workshop at RSS*, 2021.
- [23] M. Diehl, H. G. Bock, and J. P. Schlöder, "A real-time iteration scheme for nonlinear optimization in optimal feedback control," *SIAM Journal on control and optimization*, 2005.
- [24] C. Mastalli, R. Budhiraja, W. Merkt, G. Saurel, B. Hammoud, M. Naveau, J. Carpentier, L. Righetti, S. Vijayakumar, and N. Mansard, "Crocodyl: An efficient and versatile framework for multi-contact optimal control," *IEEE ICRA*, 2020.
- [25] M. Diehl, R. Findeisen, S. Schwarzkopf, I. Uslu, F. Allgöwer, H. Bock, E. Gilles, and J. Schlöder, "An efficient algorithm for nonlinear model predictive control of large-scale systems. part i: Description of the method." *Automatisierungstechnik*. 50. 557-567, 2002.
- [26] N. Mansard, A. D. Prete, M. Geisert, S. Tonneau, and O. Stasse, "Using a memory of motion to efficiently warm-start a nonlinear predictive controller," *IEEE ICRA*, 2018.
- [27] B. Henze, A. Werner, M. A. Roa, G. Garofalo, J. Engelsberger, and C. Ott, "Control applications of toro—a torque controlled humanoid robot," in *IEEE-RAS Int. Conf. Humanoid Robots*, 2014.
- [28] A. Herzog, L. Righetti, F. Grimmering, P. Pastor, and S. Schaal, "Balancing experiments on a torque-controlled humanoid with hierarchical inverse dynamics," in *International Conference on Intelligent Robots and Systems*, 2014.
- [29] E. Dantec, R. Budhiraja, A. Roig, T. Lembono, G. Saurel, O. Stasse, P. Fernbach, S. Tonneau, S. Vijayakumar, S. Calinon, M. Taix, and N. Mansard, "Whole body model predictive control with a memory of motion: Experiments on a torque-controlled talos," *IEEE ICRA*, 2021.
- [30] S. Mason, N. Rotella, S. Schaal, and L. Righetti, "Balancing and walking using full dynamics lqr control with contact constraints," *IEEE-RAS Humanoids*, 2016.
- [31] F. Laine and C. Tomlin, "Parallelizing lqr computation through endpoint-explicit riccati recursion," *IEEE Conference on Decision & Control (CDC)*, 2019.
- [32] B. Chrétien, A. Escande, and A. Kheddar, "Gpu robot motion planning using semi-infinite nonlinear programming," *IEEE Transactions on Parallel and Distributed Systems*, 2016.
- [33] M. Posa, S. Kuindersma, and R. Tedrake, "Optimization and stabilization of trajectories for constrained dynamical systems," *IEEE International Conference on Robotics and Automation (ICRA)*, 2016.
- [34] C. Mastalli, I. Havoutis, M. Focchi, D. Caldwell, and C. Semini, "Motion planning for quadrupedal locomotion: Coupled planning, terrain mapping and whole-body control," *IEEE Transactions on Robotics*, 2020.
- [35] R. Grandia, F. Farshidian, R. Ranftl, and M. Hutter, "Feedback mpc for torque-controlled legged robots," *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019.
- [36] A. Bemporad, M. Morari, V. Dua, and E. N. Pistikopoulos, "The explicit linear quadratic regulator for constrained systems," *Automatica* 38(1):3-20, 2002.
- [37] W. Chen, Z. Shao, and L. Biegler, "A bilevel nlp sensitivity-based decomposition for dynamic optimization with moving finite elements," *AICHE Journal*, vol. 60, 2014.
- [38] B. Amos, I. Rodriguez, J. Sacks, B. Boots, and Z. Kolter, "Differentiable mpc for end-to-end planning and control," *NeurIPS*, 2018.
- [39] O. Stasse, T. Flayols, R. Budhiraja, K. Giraud-Esclasse, J. Carpentier, J. Mirabel, A. D. Prete, P. Souères, N. Mansard, F. Lamiroux, J.-P. Laumond, L. Marchionni, H. Tome, and F. Ferro, "Talos: A new humanoid research platform targeted for industrial applications," *IEEE-RAS 17th Humanoids*, 2017.
- [40] R. Featherstone, *Rigid Body Dynamics Algorithms*. Springer, 2014.
- [41] Y. Tassa, T. Erez, and E. Todorov, "Synthesis and stabilization of complex behaviors through online trajectory optimization," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2012.
- [42] S. Kazdadi, J. Carpentier, and J. Ponce, "Equality constrained differential dynamic programming," *IEEE ICRA*, 2021.
- [43] R. Verschueren, G. Frison, D. Kouzoupis, N. van Duijkeren, A. Zanelli, B. Novoselnik, J. Frey, T. Albin, R. Quirynen, and M. Diehl, "Acados: a modular open-source framework for fast embedded optimal control," 2019.
- [44] R. Bellman, "The theory of dynamics programming," *the Rand Corporation*, 1954.
- [45] Y. Tassa, "Theory and implementation of biomimetic motor controllers," Ph.D. dissertation, Hebrew University of Jerusalem, 2011.
- [46] Q. L. Lidec, I. Kalevatykh, I. Laptev, C. Schmid, and J. Carpentier, "Differentiable simulation for physical system identification," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 3413-3420, 2021.
- [47] B. Amos and J. Z. Kolter, "Optnet: Differentiable optimization as a layer in neural networks," *International Conference on Machine Learning, PMLR*, 2017.
- [48] A. Shapiro, "Sensitivity analysis of nonlinear programs and differentiability properties of metric projections," *SIAM Journal on Control and Optimization* 26, 1988.
- [49] S. Mason, N. Rotella, S. Schaal, and L. Righetti, "Balancing and walking using full dynamics lqr control with contact constraints," in *IEEE-RAS Humanoids*, 2016.
- [50] J. Carpentier and N. Mansard, "Analytical derivatives of rigid body dynamics algorithms," in *Robotics: Science and Systems*, 2018.
- [51] J. Carpentier, G. Saurel, G. Buondonno, J. Mirabel, F. Lamiroux, O. Stasse, and N. Mansard, "The pinocchio c++ library : A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives," *IEEE/SICE Int. Symp. on System Integration*, 2019.