High-Frequency Nonlinear Model Predictive Control of a Manipulator

Sébastien Kleff^{1,2}, Avadesh Meduri¹, Rohan Budhiraja², Nicolas Mansard^{2,4}, Ludovic Righetti^{1,3}

Abstract—Model Predictive Control (MPC) promises to endow robots with enough reactivity to perform complex tasks in dynamic environments by frequently updating their motion plan based on measurements. Despite its appeal, it has seldom been deployed on real machines because of scaling constraints. This paper presents the first hardware implementation of closed-loop nonlinear MPC on a 7-DoF torque-controlled robot. Our controller leverages a state-of-the-art optimal control solver, namely Differential Dynamic Programming (DDP), in order to replan state and control trajectories at real-time rates (1kHz). In addition to this experimental proof of concept, an exhaustive performance analysis shows that our controller outperforms open-loop MPC on a rapid cyclic end-effector task. We also exhibit the importance of a sufficient preview horizon and full robot dynamics through comparisons with inverse dynamics and kinematic optimization.

I. INTRODUCTION

In order for robots to perform complex tasks such as collaborating with a human or running out in the real world, they must be able to anticipate and react. Optimal control offers a convenient framework to endow them with such capabilities as it characterizes state-dependent control policies from high-level objectives over a time horizon. Model predictive control (MPC) [1], [2] approximates those closed-loop policies online into locally optimal policies. Although well-known within the robotics community [1], [3], its hardware deployment is yet to be achieved when using high dimensional nonlinear dynamics, mainly due to scaling issues. In default of sufficient computation resources two main approaches can be found in the literature.

One approach is to separate planning and control into offline or low-frequency motion planning and online inverse dynamics (ID) [4]. Modern planners can generate complex and dynamic whole-body motions [5], [6] and ID controllers nowadays enjoy very fast rates thanks to efficient rigid-body dynamics algorithms [7], [8]. For instance the task function approach [9] which consists in solving IK/ID as the optimization of an operational space task metric can be solved in far less than 1ms [10]–[12]. However when these planners are not fast enough to allow online replanning, it becomes an issue in face of uncertainty such as external disturbances or modelling errors. Conversely, ID controllers [13], [14] are

Fig. 1: Pick-and-place-like task, with high dynamics and precision thanks to our state-feedback MPC, but keeping high compliance thanks to the low-level torque control.

unable to plan because the instantaneous linearization of the robot dynamics leads to a torque selection based solely on the next desired state. From an optimal control perspective ID can be regarded as the singular case of a zero preview horizon [15]. Yet anticipation at the control level is important for complex tasks, e.g. capturability-constrained balancing tasks in locomotion [16] or non-holonomic orientation control tasks through angular momentum [17], where current actions affect future states.

Another approach is to reduce the complexity of the MPC problem by reasoning over simplified models. For instance in [18] the manipulator dynamics are neglected and kinematic optimization problems are solved online. In aerial robotics, model simplifications enable to exploit structural advantages in the dynamics [19]. In legged locomotion, center-of-mass reductions, such as the linear inverted pendulum, can be used in a receding horizon controller to generate stable gaits thanks to its linearity and low dimension [20], [21]. But as simplified models are less expressive they artificially restrict the capabilities of the robot and thus limit the range of possible motions. In fact the importance of the full dynamics in complex motion generation was already acknowledged [22], [23].

Recent progress in trajectory optimization (TO) offers plausible solutions for preview control. The original continuous-time Optimal Control Problem (OCP) is transcribed into a finite dimensional nonlinear program (NLP) which can then be solved using parametric optimization techniques. In particular Differential Dynamic Programming (DDP) [24] has revealed its potential in whole-body motion

¹Tandon School of Engineering, New York University, Brooklyn, NY

²LAAS-CNRS, Université de Toulouse, CNRS, Toulouse, France

³Max Planck Institute for Intelligent Systems, Tübingen, Germany

⁴Artificial and Natural Intelligence Toulouse Institute, France

This work was in part supported by the European Union’s Horizon 2020 research and innovation program (grant agreement 780684 and European Research Councils grant 637935) and the National Science Foundation (grants 1825993, 1932187, 1925079 and 2026479).

We would like to thank Maximilien Naveau, Johannes Pflöging and Julian Viereck for their advice in the implementation and experimental work.

generation [25] [17]. DDP is an iterative method based on the *shooting* approach [1], i.e it optimizes over controls while keeping states implicit, as opposed to *collocation* methods [26] which treat states explicitly as decision variables. While collocation methods can produce impressively realistic behaviors such as in [27] thanks to a proper globalization strategy, they are much slower than shooting methods which exploit the sparsity of the problem. Recent improvements based on multiple-shooting [28] such as FDDP [29] mitigate the poor globalization by allowing unfeasible initialization.

Hardware MPC implementations based on DDP-like algorithms can already be found in the literature [30], [31]. For instance in [32] the authors propose online replanning based on a DDP variant where an optimal open-loop torque is combined with a closed-loop PD controller, but sensory feedback is not directly used in the replanning loop. In [33] the authors use iLQR in a separate thread to replan with state feedback every 50ms while the robot is controlled at 5ms. In [34], the authors propose to accommodate a low replanning frequency (15 – 20Hz) with a control frequency of 400Hz through frequency shaping. While the proposed controller exhibits good performance it only solves kinematic TO and torques are retrieved through ID. Until the point where the MPC can be evaluated at the frequency of the motor driver inputs, an additional layer of control would be necessary and the MPC would act as a planner.

In this paper we present controller that computes torques at 1kHz over a time horizon based on the full dynamics model. We propose an original optimal control formulation of the pick-and-place task and use FDDP [29] to solve the OCP in an MPC fashion with two variants: an open-loop controller similar to [32] computing a feedforward torque combined with a joint impedance controller, and a closed-loop controller updating the feedforward torque based on sensory feedback. We show experimentally that the closed-loop controller outperforms the open-loop one and we compare it with ID and kinematic TO. To the best of our knowledge, it is the first experimental demonstration of nonlinear MPC at 1kHz with pure torque control on a manipulator without any additional stabilizing controller. The outcomes are of strong practical interest: the robot is able to versatily adapt to perturbation during a dynamic task, reach the target with arbitrary precision *and* offer a full compliance when e.g. an operator physically interacts with it. This paves the road to a new way of designing controllers of manipulator robots in applicative contexts. In Section II the OCP is formulated and solved through DDP-based MPC. In Section III the pick-and-place task is presented, which is then benchmarked on the real robot in Section IV.

II. DDP-BASED MPC

In this section we formulate the motion generation problem as an OCP solve it using FDDP. The two MPC controllers are presented.

A. OCP formulation

The problem is formulated as a continuous-time OCP

$$\min_{u(\cdot), x(\cdot)} \int_0^T L(x(t), u(t), t) dt + L_T(x(T)) \quad (1)$$

$$\text{s.t.} \quad \begin{cases} \dot{x}(t) = F(x(t), u(t)) & \forall t \in [0, T] \\ x(0) = x_0, x(t) \in \mathcal{X}, u(t) \in \mathcal{U} \end{cases} \quad (2)$$

where $x = (q, v)$ is the robot's state of positions and velocities, $u = \tau$ is the control torque, F the robot dynamics, L, L_T the running and terminal costs, $\mathcal{X}, \mathcal{U}$ are the sets of admissible states and controls.

B. DDP resolution

For more details about the original DDP algorithm we refer the reader to [25], [24]. The time horizon is divided into sub-intervals $[t_i, t_{i+1})$ of length $\delta > 0$ and the solution space of (1) is restricted to piecewise-constant open-loop controls

$$\begin{aligned} \min_{X, U} &= \sum_{i=0}^{N-1} l_i(x_i, u_i) + l_N(x_N) \\ \text{s.t.} &\begin{cases} x_{i+1} = f(x_i, u_i) & \forall i \in \{0, \dots, N-1\} \\ x_0 = x_0 \end{cases} \end{aligned} \quad (3)$$

where (X, U) is the state-control sequence, f is the Euler discretization of F and $l_i = \int_{t_i}^{t_{i+1}} L(x_i, u_i) dt$. DDP solves (3) through Bellman's recursion in 2 stages, given an initial guess (X^0, U^0) and LQ approximations of f, l

- 1) *Backward pass* : a quadratic model of the Hamiltonian around the current guess is propagated starting from the terminal node and the corresponding LQRs are solved recursively for $i = N, \dots, 0$
- 2) *Forward pass* : the resulting LQR policy is simulated and the current guess is updated

The algorithm returns optimal state sequence and control policy (X^*, U^*) with feedforward control k^* and feedback gains K^*

$$u_i^*(x) = k_i^* + K_i^*(x - x_i^*) \quad \forall i \in \{0, \dots, N-1\} \quad (4)$$

The specificity of the FDDP solver used is to accept unfeasible initialization to enhance classical DDP globalization strategy. This is done by using a multiple-shooting transcription which forces the discretized state trajectory to match the shooting nodes: the equality constraint in (3) is changed into $x_{i+1} = f(x_i, u_i) - \bar{f}_{i+1}$ where $\bar{f}_{i+1}$ represent the gap between the current integrated state and the next shooting node x_{i+1} . Detailed explanations of the FDDP solver can be found in [29]. Also note that the inequality constraints are not explicitly taken into account in the solver but implicitly using cost penalization. It enables us to enjoy the efficiency of shooting methods, while enforcing hard-constraints using collocation [35] or recent constrained-DDP solvers [36] would be computationally more demanding.

Fig. 2: Open-loop MPC: replan based on predictions. Planning and control are separated. This control scheme is representative of MPC controllers found in the literature such as in [32].

C. MPC control schemes

At each MPC cycle, problem (3) is solved with a horizon $T_h = N_h \delta$ and its solution is used as an initial guess for the next cycle, i.e. $(X^0, U^0) \leftarrow (X^*, U^*)$ where x_0^* is replaced by some initial state x_0 . The open-loop and closed-loop MPC differ in the way x_0 is selected.

1) *Open-loop MPC*: The plan is updated based on the MPC predictions (see Figure 2), i.e. $x_0 \leftarrow x_1^*$. The control policy yields

$$u_{OL}(\hat{x}) = u_0^*(x_1^*) + PD(x_1^*, \hat{x}) \quad (5)$$

$$= k_0^* + K_0^*(x_1^* - x_0^*) - K_P(\hat{q} - q_1^*) - K_D(\hat{v} - v_1^*)$$

where $\hat{x} = (\hat{q}, \hat{v})$ is the measured/estimated state and $K_P, K_D > 0$ are joint PD gains. Since planning and control are synchronous, $x_0 \equiv x_0^*$ is reset to x_1^* at each cycle so the feedback term vanishes and (5) reduces to

$$u_{OL}(\hat{x}) = k_0^* - K_P(\hat{q} - q_1^*) - K_D(\hat{v} - v_1^*) \quad (6)$$

2) *Closed-loop MPC*: The plan is updated based on direct state feedback (see Figure 3), i.e. $x_0 \leftarrow \hat{x}$. The control policy yields

$$u_{CL}(\hat{x}) = u_0^*(\hat{x}) \quad (7)$$

$$= k_0^* + K_0^*(\hat{x} - x_0^*)$$

Since the $x_0 \equiv x_0^*$ is reset to $\hat{x}$ at each cycle, the control law (7) only includes a feedforward term

$$u_{CL}(\hat{x}) = k_0^* \quad (8)$$

The feedback gain K_0^* compensates small deviations around x_0^* and can be used to interpolate the control trajectory over $[t_0, t_1]$ when the planning rate is lower than the control rate. But as seen from equation (8) the feedback term is not used when planning and control are synchronous.

III. PICK-AND-PLACE TASK FORMULATION

We propose here an original formulation of a cyclic end-effector task. First we introduce an acyclic reaching task.

Fig. 3: Closed-loop MPC: direct feedback from sensors. This control scheme is representative of an "ideal" MPC.

A. Static pose reaching task

The task is to bring the end-effector position $p(x)$ to a desired end-effector position $\bar{p}$. The running cost is split into a goal tracking term, regularization terms and a barrier term

$$l(x, u) = l^1(x) + l^2(x) + l^3(u) + l^4(x) \quad (9)$$

$$= w\|p(x) - \bar{p}\|^2 + x^T Q x + u^T R u + B(x)$$

And the terminal cost is similarly defined as

$$l_N(x) = l_N^1(x) + l_N^2(x) + l_N^3(x) \quad (10)$$

$$= w_N\|p(x) - \bar{p}\|^2 + x^T Q_N x + B_N(x)$$

where the weights $w, w_N > 0$ penalize the deviation from the end-effector goal, the weight matrices $Q, R, Q_N \succ 0$ regularize states and controls and $B(\cdot), B_N(\cdot)$ are weighted quadratic barriers enforcing the state limits $(\underline{x}, \bar{x})$

$$B(x) = \begin{cases} 0 & \text{if } \underline{x} \leq x \leq \bar{x} \\ (x - \underline{x})^T B(x - \underline{x}) & \text{if } x \leq \underline{x} \\ (\bar{x} - x)^T B(\bar{x} - x) & \text{if } x \geq \bar{x} \end{cases} \quad (11)$$

Note that the OCP for this acyclic task needs to be defined only once as the weights are fixed. Typically a high terminal cost is set on the end-effector goal ($w_N \gg w$).

B. Pick-and-place task

The pick-and-place motion consists in alternatively reaching desired end-effector positions p_a and p_b at predefined times characterized by the cycle duration T_C . We propose to use a time-varying cost function similar to (9),(10), the main difference being that w becomes an increasing sequence of weights over each half-cycle and $\bar{p}$ is modified in order to cyclically penalize targets p_a and p_b . We also add to l^1 a term penalizing the end-effector velocity in order to ensure that the robot comes to a rest each time p_a or p_b is reached. Also, as shown in Figures 2 and 3 the OCP is solved at the control rate ($1kHz$) while the problem formulation (3) only allows an update of the cost function weights at the OCP sampling rate δ . In order to mitigate discontinuities in the solution, the OCP can be updated at each control cycle t

$$l^1(x, t) = w(t)\|p(x) - p^{ab}(t)\|^2 + v(t)\|\dot{p}(x)\|^2 \quad (12)$$

Fig. 4: Weights profile for pick-and-place task.

where $w(t), v(t), p^{ab}(t)$ are designed to encode the cyclic nature of the pick-and-place task, i.e. the alternation of goals p_a and p_b and a zero end-effector velocity at switching times

$$w(t) = w^a(t) + w^b(t) \quad (13)$$

$$p^{ab}(t) = \frac{1}{w(t)}(w^a(t)p^a + w^b(t)p^b) \quad (14)$$

$$v(t) = \epsilon w(t) \quad (15)$$

where $w^a(t), w^b(t) > 0$ penalize the distances to p_a, p_b and $\epsilon > 0$ is a scaling factor. Without going into algorithmic details we explain here how these weights are calculated. The general idea behind our cost design is to gradually enforce a soft terminal constraint on the end-effector position every $\frac{T_C}{2}$. We use a periodic function $h(t)$ that is monotonically increasing over each half-cycle

$$h(t) = \int_t^{t+\delta} e^{\alpha(\tau-\beta)} d\tau \quad \forall t \in [k\frac{T_C}{2}, (k+1)\frac{T_C}{2}] \quad (16)$$

where the parameters $\alpha, \beta > 0$ control the slope and sharpness of the exponential. Then $w^a(t), w^b(t)$ are defined from $h(t)$ such that they are anti-cyclic (see Figure 4).

IV. EXPERIMENTS

In this section we benchmark the two controllers of Section II-C on a 7-DoF torque-controlled KUKA LWR iiwa R820 14 (see Figure 1) for the pick-and-place task described in Section III. We also analyze the performances of the closed-loop MPC as it tends toward an ID controller and a kinematic TO.

A. Setup

The FDDP solver used in our experiments is available in the Crocoddyl library [29] which benefits from automatic differentiation and analytical rigid-body dynamics from Pinocchio [8]. Our MPC controllers were implemented in C++ and the real-time data flow was handled using the Dynamic Graph library [37] along with the Fast Robot Interface (FRI) extension of the KUKA Sunrise Workbench (v1.3). We used an Intel(R) Xeon(R) W-2145 CPU @ 3.70GHz CPU running on RT-Preempt. The KUKA internal controller applies gravity compensation by default while the policies (6), (8) already take into account the full dynamics of the robot. In order to avoid compensating twice for gravity we

(a) End-effector trajectories (desired in thin dashed line).

(b) 2-norm of the tracking error $\hat{x} - x_1^*$ and average (thin dashed) Fig. 5: The closed-loop controller generates smooth profiles that are well tracked (in green) while the open-loop controller generates abrupt switches that are difficult to track (in red). High PD gains render the robot stiff and the controller close to instability.

subtract the gravity torque $g(x_1^*)$ from our policies, i.e. the actual torque commands sent to the robot are

$$u_{OL}(\hat{x}) = k_0^* + PD(\hat{x}, x_1^*) - g(x_1^*) \quad (17)$$

$$u_{CL}(\hat{x}) = k_0^* - g(x_1^*) \quad (18)$$

The actuator rotor inertia was estimated to $0.1kg.m^2$ and added to the diagonal of the inertia matrix.

B. Open-loop vs closed-loop MPC

We compare here the MPC policies (17) and (18). The regularization weights Q were set to 10^{-2} for the joint positions and 10^{-1} for the joint velocities, the control regularization R to 10^{-2} and the state limit weight B to 50. The running cost weights are defined according to equation (16) with $T_C = 2s$, $\alpha = 40$, $\beta = 0.85$ and $\epsilon = 0.02$. The MPC parameters were set to $\delta = 30ms$ and $N_h = 30$ (i.e. we optimize over a $0.9s$ horizon). Figure 5 shows the end-effector trajectories and tracking errors between MPC predictions and measurements for both controllers. We can see that the open-loop controller generates aggressive desired trajectories requiring unreasonably high PD gains to be tracked. Furthermore it is less optimal for the task as the total running cost was 31934 (against 19978 for the closed-loop controller). In order to improve the tracking performance, the cost function weights can be reduced and the MPC horizon increased for the

Fig. 6: Reducing the cost weights ($\alpha = 20$, $\beta = 0.8$) open-loop MPC (red) results in a smaller tracking error (average norm 0.7 vs 1.6 in Figure 5) but a loss in the end-effector precision, as seen by comparison with the original closed-loop MPC trajectory (green) duplicated from Figure 5.

open-loop controller but this results in a poor goal tracking precision and the MPC solving time exceeding $1ms$. As a result we had to modify the task and trade-off a significant part of the task precision against compliance in order to improve the tracking, as shown in Figure 6. In conclusion the trade-off between MPC tracking and task fulfillment appears challenging to arbitrate with the open-loop controller since it leads either to a high gains control loop close to instability, or to a poor task performance with a longer solving time. This experimental data shows that the open-loop control scheme is not suited for online use as the predictions do not match the actual capabilities of the robot. This confirms that the state feedback is crucial to generate online smooth and realistic trajectories. The attached video¹ shows qualitatively the response of each controller during a real human-robot interaction. The stiff open-loop MPC compensates poorly an external push whereas the closed-loop controller exhibits a more compliant, yet effective, response.

C. Comparison with ID

When the horizon of an OCP collapses, the trajectories consists in a single point optimized under the instantaneous dynamics constraints and it is equivalent to ID, as shown in Appendix. As a consequence we expect to observe a reduced anticipation ability of the closed-loop controller as it tends toward an ID controller. In order to exhibit this phenomena we look at the joint velocity constraint saturation in the pick-and-place task when $T_h \rightarrow 0$. For the same task as in the first experiment, Figure 7 shows joint velocities obtained for $T_h = 900ms$ and $T_h = 360ms$ as well as the velocity constraints saturation and Figure 8 shows the average saturation over all joints as a function of T_h . As expected when $T_h \rightarrow 0$, the velocity constraint saturation increases which confirms that short horizons are less able to anticipate constraints. This resulted in a higher running cost (14182 against 10048) so the shorter horizon is less optimal for the task. This experimental data shows that a sufficiently

Fig. 7: The velocity constraint saturation peaks more significantly at each half-cycle when the horizon is smaller: short-sighted controllers are less able to anticipate constraints.

Fig. 8: Average constraint saturation over all joints as a function of the time horizon. Shorter horizon controllers hit the constraint more often.

long horizon in control schemes is important to avoid myopic behaviors.

D. Comparison with kinematic TO

Taking into account the full dynamics in the preview horizon is important in order to generate dynamically consistent motions. This can be verified by reducing the control regularization: when $R \rightarrow 0$, the optimizer should neglect the dynamics and return kinematically optimized trajectories that are dynamically not relevant. Indeed, for a fixed-based manipulator, any desired accelerations can be realized through ID at the cost of potentially very high torques. Figure 9 shows the end-effector positions, joint velocities and torque profiles

¹<https://peertube.laas.fr/videos/watch/76fcf68d-8509-4257-a3a4-c1d0eca7a8ba>

Fig. 9: When the control regularization term decreases, the trajectories are less relevant: the velocity tends toward a square signal and the torque exhibits sharp peaks

obtained for $R = 10^{-2}$ and $R = 10^{-4}$. As expected when the control is less penalized, the MPC generates abrupt torque and velocity switches that are not consistent with the robot mechanical capabilities.

V. CONCLUSION

We introduced the first torque MPC controller running at $1kHz$ on a real robot. We showed experimentally that this controller is able to achieve with good performance dynamic cyclic tasks while ensuring a compliant response to external disturbances thanks to its high-frequency state feedback. We also exhibited the importance of the preview horizon and full dynamics in the performance and thereby its superiority against kinematic optimization and ID which is merely a singular case of MPC with a collapsed horizon.

As future work we would like to improve the efficiency of our controller in order to increase the preview horizon and achieve more dynamic tasks, to incorporate hard constraints, and to extend this approach to more complex scenarios involving contacts with the environment, with the ultimate objective of applying it to multi-contact locomotion.

APPENDIX

We show here that ID is equivalent to an OCP with a collapsed horizon. Task-Space Inverse Dynamics (TSID) computes joint torques achieving a desired task-space impedance by solving a QP

$$\begin{aligned} \min_{\tau} \quad & \alpha_1 \|J(q)a + PD_1(p(q), \dot{p}(q, v))\|^2 \\ & + \alpha_2 \|a + PD_2(q, v)\|^2 + \alpha_3 \|\tau\|^2 \\ \text{s.t.} \quad & a = M^{-1}(\tau - b) \end{aligned} \quad (19)$$

where q, v, a are the joint positions, velocities and accelerations, $p(q)$ is the end-effector position, $J = \frac{\partial p}{\partial q}(q)$ is the Jacobian, PD_1, PD_2 are task-space and joint-space impedances, $b = b(q, v)$ summarizes Coriolis, gravity and centrifugal effects. Consider the OCP (3) with $N = 1$

$$\begin{aligned} \min_{u_0, x_0, x_1} \quad & l_0(x_0, u_0) + l_1(x_1) \\ \text{s.t.} \quad & \begin{cases} x_1 = f(x_0, u_0) \\ x_0 = x_0 \end{cases} \end{aligned} \quad (20)$$

Selecting a cost function encoding a task-space objective with joint-space regularization, e.g.

$$\begin{aligned} l_0(x, u) &= \|p(q)\|^2 + \|Jv\|^2 + \|q - \bar{q}\|^2 + \|v\|^2 + \|u\|^2 \\ l_1(x) &= \|p(q)\|^2 + \|Jv\|^2 + \|q - \bar{q}\|^2 + \|v\|^2 \end{aligned} \quad (21)$$

where $\bar{q}$ is a reference posture. Since DDP uses the shooting approach, the state is not an explicit decision variable so the singular-horizon OCP (20) is equivalent to

$$\begin{aligned} \min_{u_0} \quad & \|u_0\|^2 + \|p(q_1)\|^2 + \|Jv_1\|^2 + \|q_1 - \bar{q}\|^2 + \|v_1\|^2 \\ \text{s.t.} \quad & a_1 = M^{-1}(q_0)(\tau - b(q_0, v_0)) \end{aligned} \quad (22)$$

The Euler 2^{nd} -order integration with step $\delta \ll 1$ leads to

$$\begin{aligned} p(q_1) &\simeq p(q_0) + Jv_0\delta + Ja_1\delta^2 \\ Jv_1 &\simeq Jv_0\delta + Ja_1\delta \end{aligned} \quad (23)$$

and using (23) in (22), it can be shown that the collapsed OCP can be written under the form

$$\begin{aligned} \min_{\tau} \quad & \beta_1 \|Ja + K_1p(q) + K_2Jv\|^2 \\ & + \beta_2 \|a + K_3(q - \bar{q}) + K_4v\|^2 \\ \text{s.t.} \quad & a = M^{-1}(\tau - b) \end{aligned} \quad (24)$$

where $(K_1, K_2), (K_3, K_4), \beta_1, \beta_2$ are task-space PD gains, joint-space PD gains, and weights all depending on δ (and the OCP cost weights if any). This problem is equivalent to (19), which concludes the proof that TSID is merely the particular case of an OCP with a singular horizon.

REFERENCES

- [1] M. Diehl, H. G. Bock, H. Diedam, and P.-B. Wieber, "Fast Direct Multiple Shooting Algorithms for Optimal Robot Control," in *Fast Motions in Biomechanics and Robotics*. Springer Berlin Heidelberg, 2005, pp. 65–93.
- [2] D. Q. Mayne, "Model predictive control: Recent developments and future promise," *Automatica*, vol. 50, no. 12, pp. 2967–2986, 2014.
- [3] V. Duchaine, S. Bouchard, and C. Gosselin, "Computationally efficient predictive robot control," *IEEE/ASME Transactions on Mechatronics*, vol. 12, pp. 570 – 578, 2007.
- [4] O. Khatib, "A Unified Approach for Motion and Force Control of Robot Manipulators: The Operational Space Formulation," *IEEE Journal on Robotics and Automation*, vol. 3, no. 1, pp. 43–53, 1987.
- [5] B. Ponton, A. Herzog, A. Del Prete, S. Schaal, and L. Righetti, "On Time Optimization of Centroidal Momentum Dynamics," in *IEEE International Conference on Robotics and Automation*, 2018.
- [6] J. Carpentier and N. Mansard, "Multicontact locomotion of legged robots," *IEEE Transactions on Robotics*, vol. 34, no. 6, pp. 1441–1460, 2018.
- [7] R. Featherstone, *Rigid Body Dynamics Algorithms*. Berlin, Heidelberg: Springer-Verlag, 2008.
- [8] 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," in *IEEE/SICE International Symposium on System Integration*, 2019.
- [9] L. Saab, O. E. Ramos, F. Keith, N. Mansard, P. Soueres, and J. Y. Fourquet, "Dynamic whole-body motion generation under rigid contacts and other unilateral constraints," *IEEE Transactions on Robotics*, vol. 29, no. 2, pp. 346–362, 2013.
- [10] A. Del Prete, N. Mansard, O. E. Ramos, O. Stasse, and F. Nori, "Implementing Torque Control with High-Ratio Gear Boxes and Without Joint-Torque Sensors," *International Journal of Humanoid Robotics*, vol. 13, no. 1, p. 1550044, 2016.
- [11] A. Herzog, N. Rotella, S. Mason, F. Grimmering, S. Schaal, and L. Righetti, "Momentum control with hierarchical inverse dynamics on a torque-controlled humanoid," *Autonomous Robots*, vol. 40, no. 3, pp. 473–491, 2016.
- [12] P. M. Wensing, L. R. Palmer, and D. E. Orin, "Efficient recursive dynamics algorithms for operational-space control with application to legged locomotion," *Autonomous Robots*, vol. 38, no. 4, p. 363–381, 2015.
- [13] L. Righetti, J. Buchli, M. Mistry, M. Kalakrishnan, and S. Schaal, "Optimal distribution of contact forces with inverse-dynamics control," *International Journal of Robotics Research*, vol. 32, no. 3, pp. 280–298, 2013.
- [14] M. Mistry and L. Righetti, "Operational Space Control of Constrained and Underactuated Systems," in *Robotics: Science and Systems VII*. The MIT Press, 2012.
- [15] P. Geoffroy, N. Mansard, M. Raison, S. Achiche, and E. Todorov, "From Inverse Kinematics to Optimal Control," *Advances in Robot Kinematics*, pp. 409–418, 2014.
- [16] P. B. Wieber, "Viability and predictive control for safe locomotion," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2008.
- [17] R. Budhiraja, J. Carpentier, C. Mastalli, and N. Mansard, "Differential Dynamic Programming for Multi-Phase Rigid Contact Dynamics," in *IEEE-RAS 18th International Conference on Humanoid Robots*, 2018.
- [18] M. Krämer, C. Rösmann, F. Hoffmann, and T. Bertram, "Model predictive control of a collaborative manipulator considering dynamic obstacles," *Optimal Control Applications and Methods*, vol. 41, no. 4, pp. 1211–1232, 2020.
- [19] A. Boeuf, J. Cortés, R. Alami, and T. Siméon, "Planning agile motions for quadrotors in constrained environments," in *IEEE International Conference on Intelligent Robots and Systems*, 2014.
- [20] 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," in *IEEE International Conference on Robotics and Automation*, 2003.
- [21] A. Herdt, H. Diedam, P.-B. Wieber, D. Dimitrov, K. Mombaur, and M. Diehl, "Online walking motion generation with automatic foot step placement," *Advanced Robotics*, vol. 24, pp. 719–737, 04 2010.
- [22] D. E. Orin, A. Goswami, and S. H. Lee, "Centroidal dynamics of a humanoid robot," *Autonomous Robots*, vol. 35, no. 2-3, pp. 161–176, 2013.
- [23] P.-B. Wieber, "Holonomy and nonholonomy in the dynamics of articulated motion," in *Fast motions in biomechanics and robotics*. Springer Berlin Heidelberg, 2006, pp. 411–425.
- [24] D. Mayne, "A second-order gradient method for determining optimal trajectories of non-linear discrete-time systems," *International Journal of Control*, vol. 3, no. 1, pp. 85–95, 1966.
- [25] Y. Tassa, T. Erez, and E. Todorov, "Synthesis and stabilization of complex behaviors through online trajectory optimization," *IEEE International Conference on Intelligent Robots and Systems*, 2012.
- [26] M. Kelly, "An introduction to trajectory optimization: How to do your own direct collocation," *SIAM Review*, vol. 59, no. 4, pp. 849–904, 2017.
- [27] I. Mordatch, E. Todorov, and Z. Popović, "Discovery of complex behaviors through contact-invariant optimization," *ACM Transactions on Graphics*, vol. 31, no. 4, 2012.
- [28] M. Diehl, H. Bock, J. P. Schlöder, R. Findeisen, Z. Nagy, and F. Allgöwer, "Real-time optimization and nonlinear model predictive control of processes governed by differential-algebraic equations," *Journal of Process Control*, vol. 12, no. 4, pp. 577–585, 2002.
- [29] 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," in *International Conference on Robotics and Automation*, 2020.
- [30] M. Gifftthaler, M. Neunert, M. Stauble, J. Buchli, and M. Diehl, "A Family of Iterative Gauss-Newton Shooting Methods for Nonlinear Optimal Control," in *IEEE International Conference on Intelligent Robots and Systems*, 2018.
- [31] 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," 2016.
- [32] M. Neunert, F. Farshidian, A. W. Winkler, and J. Buchli, "Trajectory Optimization Through Contacts and Automatic Gait Discovery for Quadrupeds," *IEEE Robotics and Automation Letters*, vol. 2, no. 3, pp. 1502–1509, 2017.
- [33] J. Koenemann, A. Del Prete, Y. Tassa, E. Todorov, O. Stasse, M. Bennewitz, and N. Mansard, "Whole-body model-predictive control applied to the hrp-2 humanoid," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2015.
- [34] R. Grandia, F. Farshidian, R. Ranftl, and M. Hutter, "Feedback mpc for torque-controlled legged robots," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019.
- [35] M. Posa, S. Kuindersma, and R. Tedrake, "Optimization and stabilization of trajectories for constrained dynamical systems," *IEEE International Conference on Robotics and Automation*.
- [36] T. Howell, B. Jackson, and Z. Manchester, "Altro: A fast solver for constrained trajectory optimization," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2019.
- [37] N. Mansard, O. Stasse, P. Evrard, and A. Kheddar, "A versatile generalized inverted kinematics implementation for collaborative working humanoid robots: The stack of tasks," in *International Conference on Advanced Robotics*, 2009.