THE UNIVERSITY of EDINBURGH

Edinburgh Research Explorer

Inverse Dynamics vs. Forward Dynamics in Direct Transcription Formulations for Trajectory Optimization

Citation for published version:
Ferrolho, H, Ivan, V, Merkt, W, Havoutis, I & Vijayakumar, S 2021, Inverse Dynamics vs. Forward Dynamics in Direct Transcription Formulations for Trajectory Optimization. in *2021 IEEE International Conference on Robotics and Automation (ICRA).* Institute of Electrical and Electronics Engineers (IEEE), pp. 12752-12758, 2021 IEEE International Conference on Robotics and Automation, Xi'an, China, 30/05/21. https://doi.org/10.1109/ICRA48506.2021.9561306

Digital Object Identifier (DOI):
10.1109/ICRA48506.2021.9561306

Link:
Link to publication record in Edinburgh Research Explorer

Document Version:
Peer reviewed version

Published In:
2021 IEEE International Conference on Robotics and Automation (ICRA)

General rights
Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy
The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

OPEN ACCESS

Download date: 25. Apr. 2024

Inverse Dynamics vs. Forward Dynamics in Direct Transcription Formulations for Trajectory Optimization

Henrique Ferrolho[1], Vladimir Ivan[1], Wolfgang Merkt[2], Ioannis Havoutis[2], Sethu Vijayakumar[1]

***Abstract*— Benchmarks of state-of-the-art rigid-body dynamics libraries report better performance solving the inverse dynamics problem than the forward alternative. Those benchmarks encouraged us to question whether that computational advantage would translate to direct transcription, where calculating rigid-body dynamics and their derivatives accounts for a significant share of computation time. In this work, we implement an optimization framework where both approaches for enforcing the system dynamics are available. We evaluate the performance of each approach for systems of varying complexity, for domains with rigid contacts. Our tests reveal that formulations using inverse dynamics converge faster, require less iterations, and are more robust to coarse problem discretization. These results indicate that inverse dynamics should be preferred to enforce the nonlinear system dynamics in simultaneous methods, such as direct transcription.**

I. Introduction

Direct transcription [1] is an effective approach to formulate and solve trajectory optimization problems. It works by converting the original trajectory optimization problem (which is *continuous* in time) into a numerical optimization problem that is *discrete* in time, and which in turn can be solved using an off-the-shelf nonlinear programming (NLP) solver. First, the trajectory is divided into segments and then, at the beginning of each segment, the system state and control inputs are explicitly discretized—these are the decision variables of the optimization problem. Due to this discretization approach, direct transcription falls under the class of *simultaneous* methods. Finally, a set of mathematical constraints is defined to enforce boundary and path constraints, e.g., initial and final conditions, or intermediate goals. In dynamic trajectory optimization, there exists a specific set of constraints dedicated to enforce the equations of motion of the system, the so-called *defect constraints*. This paper discusses different ways of defining these constraints, as well as their implications.

The dynamics defects are one of the most important constraints in optimization problems when planning highly dynamic motions for complex systems, such as legged robots. Satisfaction of these constraints ensures that the computed motion is reliable and physically consistent with the nonlinear dynamics of the system. The dynamics defect constraints are usually at the very core of optimal control formulations, and require computing rigid-body dynamics and their derivatives—which account for a significant portion of the optimization computation time. Therefore, it is of utmost importance to use an algorithm that allows to compute the dynamics of the system reliably, while achieving low computational time.

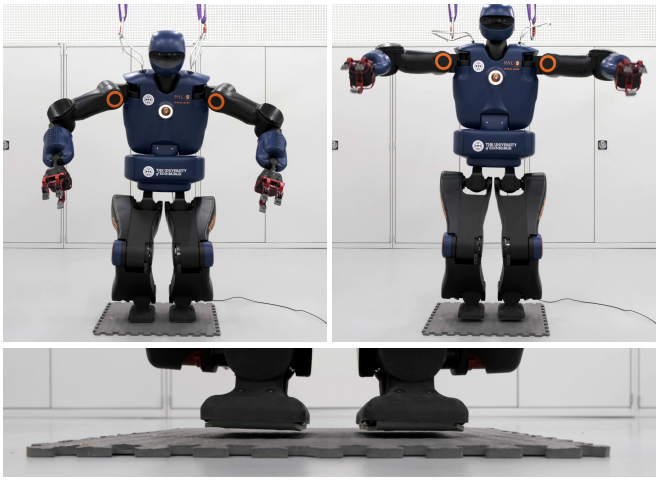

Fig. 1: Snapshots of the humanoid TALOS [2] jumping.

In the study of the dynamics of open-chain robots, the *forward dynamics* problem determines the joint accelerations resultant from a given set of joint forces and torques applied at a given state. On the other hand, the *inverse dynamics* problem determines the joint torques and forces required to meet some desired joint accelerations at a given state. In trajectory optimization, most direct formulations use forward dynamics to enforce dynamical consistency [3]. However, benchmarks have shown that most dynamics libraries solve the *inverse dynamics* problem (e.g., with the Recursive Newton-Euler Algorithm) faster than the *forward dynamics* problem (e.g., with the Articulated Body Algorithm) [4], [5]. For example, for the humanoid robot TALOS [2], the library Pinocchio [6] solves the inverse dynamics problem in just $4\,\mu s$, while the forward dynamics problem takes $10\,\mu s$. These differences in performance motivated us to question whether the computational advantage of inverse dynamics would translate to direct transcription—where the dynamics problem needs to be solved several times while computing the defect constraints. Moreover, there is biological evidence suggesting that inverse dynamics is employed by the nervous system to generate feedforward commands [7], while other studies support the existence of a forward model [8]—which increased our interest in this topic.

[1]School of Informatics, University of Edinburgh, United Kingdom.
[2]Oxford Robotics Institute, University of Oxford, United Kingdom.
This research is supported by EPSRC UK RAI Hub for Offshore Robotics for Certification of Assets (ORCA, EP/R026173/1), EU H2020 project Memory of Motion (MEMMO, 780684), and EPSRC as part of the Centre for Doctoral Training in Robotics and Autonomous Systems at Heriot-Watt University and The University of Edinburgh (EP/L016834/1).
Email address: henrique.ferrolho@ed.ac.uk

In this work, we present a trajectory optimization framework for domains with rigid contacts, using a direct transcription approach. Particularly, our formulation allows to define dynamics defect constraints employing either forward dynamics or inverse dynamics. We defined a set of evaluation tasks across different classes of robot platforms, including fixed- and floating-base systems, with point and surface contacts. Our results showed that inverse dynamics leads to significant improvements in computational performance when compared to forward dynamics—supporting our initial hypothesis.

II. Related Work

RigidBodyDynamics.jl (RBD.jl) [9], RBDL [10], Pinocchio [6], and RobCoGen [11] are all state-of-the-art software implementations of key rigid-body dynamics algorithms. Recently, Neuman *et al.* [5] benchmarked these libraries and revealed interesting trends. One such trend is that implementations of inverse dynamics algorithms have faster runtimes than forward dynamics.[1] Koolen and Deits [4] also compared RBD.jl with RBDL, and their results showed that solving inverse dynamics was at least two times faster than solving forward dynamics for the humanoid robot Atlas. Both of these studies only consider computation time of rigid-body dynamics; they do not provide insight into how these algorithms perform when used in trajectory optimization.

Lee *et al.* [12] have proposed Newton and quasi-Newton algorithms to optimize motions for serial-chain and closed-chain mechanisms using inverse dynamics. However, they used relatively simple mechanisms for which analytic derivatives can be obtained. In this work, we are interested in dynamic motions of complex mechanisms in domains with contact, for which the derivation of analytic derivatives is an error-prone process, involving significant effort.

In the same spirit, Erez and Todorov [13] generated a running gait for a humanoid based on inverse dynamics under external contacts. This method allowed them to formulate an *unconstrained* optimization where all contact states can be considered equally, contact timings and locations are optimized, and reaction forces are computed using a smooth and invertible contact model [14] with convex optimization. However, their approach requires "helper forces", as well as tuning of contact smoothness and of the penalty parameters on the helper forces to achieve reasonable-looking behavior. In contrast, our approach does not require helper forces or any tuning whatsoever; we consider contact forces as decision variables and model contacts rigidly. Another difference is that we formulate a *constrained* optimization problem and enforce the nonlinear system dynamics with hard constraints, which results in high-fidelity motions. This is especially important for deployment on real hardware, where dynamic consistency and realism are imperative. The main focus of our paper is not the contact problem, and we assume contact locations and contact times are known *a priori*.

[1] RobCoGen is an exception to this observation as it implements a hybrid dynamics solver which has a higher computational cost, and is significantly different from the implementations used by the other libraries.

Finally, to the best of our knowledge, there is no current work directly comparing inverse dynamics against forward dynamics in the context of direct methods. Posa *et al.* [15] also identified that a formal comparison is important, but missing so far. They argued that one of the reasons for this was that the field had not yet agreed upon a set of canonical and hard problems. In this paper, we tackle this issue, and compare the two approaches on robots of different complexity on a set of dynamic tasks.

The main contributions of this work are:

1) A direct transcription formulation that uses *inverse dynamics* to enforce physical consistency, for constrained trajectory optimization in domains with rigid contacts.
2) Evaluation of the performance of direct transcription formulations using either forward or inverse dynamics, for different classes of robot platforms: a fixed-base manipulator, a quadruped, and a humanoid.
3) Comparison of performance for different linear solvers, and across strategies to handle the barrier parameter of the interior point optimization algorithm.

We validated our trajectories in full-physics simulation and with hardware experiments. We also open-sourced a version of our framework for fixed-base robots, TORA.jl [16].

III. Trajectory Optimization

A. Robot Model Formulation

We formulate the model of a legged robot as a free-floating base B to which limbs are attached. The motion of the system can be described with respect to (w.r.t.) a fixed inertial frame I. We represent the position of the free-floating base w.r.t. the inertial frame, and expressed in the inertial frame, as ${}_I\boldsymbol{r}_{IB} \in \mathbb{R}^3$; and the orientation of the base as $\boldsymbol{\psi}_{IB} \in \overline{\mathbb{R}}^3$, using modified Rodrigues parameters (MRP) [17], [18]. The joint angles describing the configuration of the limbs of the robot (legs or arms) are stacked in a vector $\boldsymbol{q}_j \in \mathbb{R}^{n_j}$, where n_j is the number of actuated joints. The generalized coordinates vector $\boldsymbol{q}$ and the generalized velocities vector $\boldsymbol{v}$ of this floating-base system may therefore be written as

$$\boldsymbol{q} = \begin{bmatrix} {}_I\boldsymbol{r}_{IB} \\ \boldsymbol{\psi}_{IB} \\ \boldsymbol{q}_j \end{bmatrix} \in \mathbb{R}^3 \times \overline{\mathbb{R}}^3 \times \mathbb{R}^{n_j}, \quad \boldsymbol{v} = \begin{bmatrix} \boldsymbol{\nu}_B \\ \dot{\boldsymbol{q}}_j \end{bmatrix} \in \mathbb{R}^{n_v}, \quad (1)$$

where the twist $\boldsymbol{\nu}_B = \begin{bmatrix} {}_I\boldsymbol{v}_B & {}_B\boldsymbol{\omega}_{IB} \end{bmatrix}^\top \in \mathbb{R}^6$ encodes the linear and angular velocities of the base B w.r.t. the inertial frame expressed in the I and B frames, and $n_v = 6 + n_j$.

For fixed-base manipulators, the generalized vectors of coordinates and velocities can be simplified to $\boldsymbol{q} = \boldsymbol{q}_j \in \mathbb{R}^{n_j}$ and $\boldsymbol{v} = \dot{\boldsymbol{q}}_j \in \mathbb{R}^{n_j}$, due to the absence of a free-floating base.

B. Problem Formulation

We tackle the motion planning problem using trajectory optimization; more specifically, using a *direct transcription* approach. The original problem is *continuous* in time, so we start by converting it into a numerical optimization problem that is *discrete* in time. For that, we divide the trajectory into N equally spaced segments, $t_I = t_1 < \cdots < t_M = t_F$,

where t_I and t_F are the start and final instants, respectively. This division results in $M = N + 1$ discrete *mesh points*, for each of which we explicitly discretize the states of the system, as well as the control inputs. Let $x_k \equiv x(t_k)$ and $u_k \equiv u(t_k)$ be the values of the state and control variables at the k-th mesh point. We treat $x_k \triangleq \{\boldsymbol{q}_k, \boldsymbol{v}_k\}$ and $u_k \triangleq \{\boldsymbol{\tau}_k, \boldsymbol{\lambda}_k\}$ as a set of NLP variables, and formulate the trajectory optimization problem as

$$\text{find } \boldsymbol{\xi} \quad \text{s.t. } x_{k+1} = f(x_k, u_k), \;\; x_k \in \mathcal{X}, \;\; u_k \in \mathcal{U}, \quad (2)$$

where $\boldsymbol{\xi}$ is the vector of decision variables, $x_{k+1} = f(x_k, u_k)$ is the state transition function incorporating the nonlinear system dynamics, and $\mathcal{X}$ and $\mathcal{U}$ are sets of feasible states and control inputs enforced by a set of equality and inequality constraints. The decision variables vector $\boldsymbol{\xi}$ results from aggregating the generalized coordinates, generalized velocities, joint torques, and contact forces at every[2] mesh point, i.e.,

$$\boldsymbol{\xi} \triangleq \{\boldsymbol{q}_1, \boldsymbol{v}_1, \boldsymbol{\tau}_1, \boldsymbol{\lambda}_1, \cdots, \boldsymbol{q}_N, \boldsymbol{v}_N, \boldsymbol{\tau}_N, \boldsymbol{\lambda}_N, \boldsymbol{q}_M, \boldsymbol{v}_M\}. \quad (3)$$

Similarly to Winkler *et al.* [19] and differently to Erez and Todorov [13], we transcribe the problem by only making use of hard constraints; and satisfaction of those constraints is a necessary requirement for the computed motions to be physically feasible and to complete the task successfully. This design decision is motivated by the fact that considering a cost function requires expert knowledge to carefully tune the weighting parameters that control the trade-off between different objective terms. Optimizing an objective function also requires additional iterations and computation time. Nonetheless, for the sake of completion, one of the experiments we present later in this paper does include and discuss the minimization of a cost function.

For tasks where the robot makes or breaks contacts with the environment, we assume contact locations and contact timings are known *a priori*. This assumption allows us to enforce zero contact forces for mesh points where the robot is not in contact with the environment, and therefore our formulation does not require any actual complementarity constraints. On the other hand, such assumption depends on pre-determined contact sequences specified either by a human or by a contact planner (such as [19], [20], [21]).

C. Problem Constraints

1) Bounds on the decision variables: We constrain the joint positions, velocities, and torques to be within their corresponding lower and upper bounds.

2) Initial and final joint velocities: We enforce the initial and final velocities of every joint to be zero: $\boldsymbol{v}_1 = \boldsymbol{v}_M = \boldsymbol{0}$.

3) End-effector pose: We enforce end-effector poses with $f^{\text{fk}}(\boldsymbol{q}_k, i) = \boldsymbol{p}_i$, where $f^{\text{fk}}(\cdot)$ is the forward kinematics function, i refers to the i-th end-effector of the robot, and $\boldsymbol{p}_i \in SE(3)$ is the desired pose.

4) Contact forces: For mesh points where the robot is *not* in contact with the environment, we enforce the contact forces at the respective contact points to be zero: $\boldsymbol{\lambda}_k = \boldsymbol{0}$.

[2]The control inputs at the final state need not be discretized.

5) Friction constraints: We model friction at the contacts with linearized friction cones, in the same way as [22].

6) System dynamics: We enforce nonlinear whole-body dynamics, $\dot{x} = f(x, u)$, with *defect* constraints. The approach used to define these constraints is the main subject of this paper, and the next section explains this in detail.

IV. System Dynamics

The equations of motion for a floating-base robot that interacts with its environment can be written as

$$\boldsymbol{M}(\boldsymbol{q})\dot{\boldsymbol{v}} + \boldsymbol{h}(\boldsymbol{q}, \boldsymbol{v}) = \boldsymbol{S}^\top \boldsymbol{\tau} + \boldsymbol{J}_s^\top(\boldsymbol{q})\boldsymbol{\lambda}, \quad (4)$$

where $\boldsymbol{M}(\boldsymbol{q}) \in \mathbb{R}^{n_v \times n_v}$ is the mass matrix, and $\boldsymbol{h}(\boldsymbol{q}, \boldsymbol{v}) \in \mathbb{R}^{n_v}$ is the vector of Coriolis, centrifugal, and gravity terms. On the right-hand side of the equation, $\boldsymbol{\tau} \in \mathbb{R}^{n_\tau}$ is the vector of joint torques commanded to the system, and the selection matrix $\boldsymbol{S} = [\boldsymbol{0}_{n_\tau \times (n_v - n_\tau)} \quad \mathbb{I}_{n_\tau \times n_\tau}]$ selects which degrees of freedom (DoF) are actuated. We consider that all limb joints are actuated, thus $n_\tau = n_j$. The vector $\boldsymbol{\lambda} \in \mathbb{R}^{n_s}$ denotes the forces and torques experienced at the contact points, with n_s being the total dimensionality of all contact wrenches. The support Jacobian $\boldsymbol{J}_s \in \mathbb{R}^{n_s \times n_v}$ maps the contact wrenches $\boldsymbol{\lambda}$ to joint-space torques, and it is obtained by stacking the Jacobians which relate generalized velocities to limb end-effector motion as $\boldsymbol{J}_s = [\boldsymbol{J}_{C_1}^\top \quad \cdots \quad \boldsymbol{J}_{C_{n_c}}^\top]^\top$, with n_c being the number of limbs in contact. For fixed-base robots that are not subject to contact forces, we can simplify the equations of motion to $\boldsymbol{M}(\boldsymbol{q})\dot{\boldsymbol{v}} + \boldsymbol{h}(\boldsymbol{q}, \boldsymbol{v}) = \boldsymbol{\tau}$.

In order to enforce the equations of motion of nonlinear systems, we define a set of equality constraints within our framework, the so-called *defect constraints*. Usually, these constraints are defined using a forward dynamics algorithm, but in this paper we argue that using inverse dynamics can be more computationally advantageous.

The standard problem of forward dynamics computes the joint accelerations resultant from commanding torques and applying forces to the robot at a given state, i.e.,

$$\dot{\boldsymbol{v}}_k^* = f^{\text{fd}}(\boldsymbol{q}_k, \boldsymbol{v}_k, \boldsymbol{\tau}_k, \boldsymbol{\lambda}_k), \quad (5)$$

where $f^{\text{fd}}(\cdot)$ is the function that solves forward dynamics. The asterisk $(\cdot)^*$ denotes intermediately computed values, whereas terms without an asterisk are NLP variables. Using the semi-implicit Euler method as the integration scheme and $h = (t_F - t_I)/N$ as the integration time step, we can compute the state of the robot after h seconds. First, we integrate $\dot{\boldsymbol{v}}_k^*$ to compute the next generalized velocities $\boldsymbol{v}_{k+1}^* = \boldsymbol{v}_k + h\,\dot{\boldsymbol{v}}_k^*$. Then, we can compute the time derivative of the generalized coordinates, $\dot{\boldsymbol{q}}_{k+1}^*$, from those velocities, $\boldsymbol{v}_{k+1}^*$. In turn, we integrate that time derivative to compute the next generalized coordinates, $\boldsymbol{q}_{k+1}^* = \boldsymbol{q}_k + h\,\dot{\boldsymbol{q}}_{k+1}^*$. After these calculations, we end up with two different values for the state of the system at mesh point $k+1$: one from the discretized NLP variables, and another computed as a result of the controls applied to the system at mesh point k. To enforce dynamical consistency, we define the defect constraints as

$$\boldsymbol{q}_{k+1}^* - \boldsymbol{q}_{k+1} = \boldsymbol{0} \quad \text{and} \quad \boldsymbol{v}_{k+1}^* - \boldsymbol{v}_{k+1} = \boldsymbol{0}. \quad (6)$$

However, there is an alternative way to enforce dynamical consistency: with inverse dynamics. In contrast to (5), inverse dynamics computes the joint torques and forces required to meet desired joint accelerations at a given state, i.e.,

$$\boldsymbol{\tau}_k^* = f^{\mathrm{id}}(\boldsymbol{q}_k, \boldsymbol{v}_k, \dot{\boldsymbol{v}}_k^*, \boldsymbol{\lambda}_k), \tag{7}$$

where $f^{\mathrm{id}}(\cdot)$ is the function that solves the inverse dynamics problem, and the desired joint accelerations can be calculated implicitly with $\dot{\boldsymbol{v}}_k^* = (\boldsymbol{v}_{k+1} - \boldsymbol{v}_k)/h$. Similarly to the forward dynamics case, we compute $\dot{\boldsymbol{q}}_{k+1}^*$ from $\boldsymbol{v}_{k+1}$, and integrate it to compute the next generalized coordinates $\boldsymbol{q}_{k+1}^*$. And finally, we define the dynamics defect constraints as

$$\boldsymbol{q}_{k+1}^* - \boldsymbol{q}_{k+1} = \boldsymbol{0} \quad \text{and} \quad \boldsymbol{\tau}_k^* - \boldsymbol{\tau}_k = \boldsymbol{0}. \tag{8}$$

Notice that the main difference between equations (6) and (8) is that forward dynamics enforces consistency of the generalized velocities whereas inverse dynamics enforces consistency of joint torques commanded to the system.

The main subject of this paper revolves around the two formulations explained above to enforce the nonlinear system dynamics: *forward dynamics* vs. *inverse dynamics*. We developed our framework with both options in mind, and we are able to easily toggle between one approach and the other, which was particularly useful for our experiments.

V. Experiments and Results

This section is organized into four subsections:

A. Compares the computation time and number of solver iterations required to find locally-optimal solutions;
B. Evaluates the robustness of each approach as problem discretization gets more coarse (larger time steps);
C. Analyzes the performance of each formulation for the minimization of a cost function; and finally,
D. Shows hardware validation of the planned motions.

All evaluations were carried out in a single-threaded process on an Intel i7-6700K CPU with clock frequency fixed at $4.0\,\mathrm{GHz}$, and $32\,\mathrm{GB}$ $2133\,\mathrm{MHz}$ memory. The framework we propose has been implemented in Julia [23], using the rigid-body dynamics library RBD.jl [9], and the optimization library Knitro [24]. To solve the formulated NLP problems, we used the interior-point method of Waltz *et al.* [25].

A. Evaluation of Convergence

In order to evaluate and compare forward dynamics against inverse dynamics in the context of direct transcription, we used our framework to specify tasks in the form of numerical optimization problems for different types of robots: a manipulator, a quadruped, and a humanoid. Those robots were selected as they allows us to evaluate the formulations for distinct features: fixed- and floating-base systems, single-point and surface contacts, and low and high dimensionality. For each task on each robot, we solved the optimization problem twice: first defining the defect constraints with forward dynamics, and then with inverse dynamics. The *only* changing factor was the toggling between forward and inverse dynamics for the definition of the defect constraints; every other aspect of the formulation was kept unchanged.

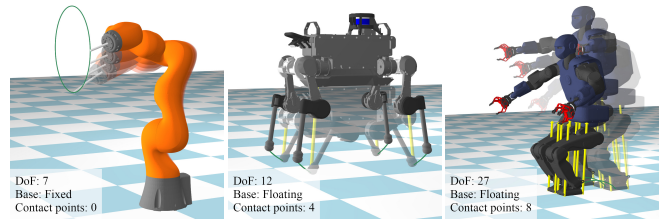

(a) Manipulator (b) Quadruped (c) Humanoid

Fig. 2: Left to right: KUKA iiwa tracing a circular path, ANYmal jumping in-place, and TALOS jumping forward.

The performance of general NLP solvers is greatly affected by the linear solver used for solving the linear systems of equations of the problem. For this reason, we tested different state-of-the-art linear solvers exhaustively. For interior-point methods, another important factor that affects performance is the update strategy of the barrier parameter. Therefore, for all of our evaluations, we tested the different strategies available within the Knitro [24] library exhaustively.

In the remainder of this subsection, we present the task specifications for each robot and indicate all the parameters for reproducibility. Then, we present the results we obtained for each task, which evaluated the solver's performance in terms of computation time and number of iterations taken by the solver until a locally-optimal solution was found.

1) Manipulator: We evaluated the different formulations using a fixed-base robot arm with seven DoF (shown in Figure 2a). We specified the end-effector to trace a circular path given by $[0.5, 0.2\cos\theta, 0.8 + 0.2\sin\theta] \forall \theta \in [0, 2\pi]$. The total duration was set to $2.0\,\mathrm{s}$ and the trajectory was discretized at $150\,\mathrm{Hz}$, resulting in a total of 301 mesh points.

2) Quadruped: The quadruped robot we used is shown in Figure 2b. This system is more complex than the manipulator due to its floating-base, more DoF (three motors per leg), and because it needs to handle contact forces. We defined a jumping task by enforcing the contact forces to be zero for a short period of time. The trajectory was discretized at $100\,\mathrm{Hz}$, the total duration of the motion was $2.0\,\mathrm{s}$, and the interval specified for the flight-phase was $[1.0, 1.2]$ s. We did not constrain feet positions during the flight-phase, which allowed the solver to converge to a solution where the feet swing most naturally according to the system dynamics.

3) Humanoid: Finally, we considered the humanoid robot shown in Figure 2c. This robot is more complex than the quadruped because it has 27 DoF[3] (seven per arm, six per leg, and one at the torso), and its feet cannot be simplified to single-point contacts. We also defined a jumping task for this robot. The motion duration was $1.2\,\mathrm{s}$, discretized at $125\,\mathrm{Hz}$, and the interval for the flight-phase was $[0.5, 0.8]$ s.

For all the tasks, the initial guess was a fixed standing configuration and zero velocities, torques, and contact forces.

The results of the experiments on these robots are shown in Table I, where smaller numbers indicate better performance. The rows of the table are grouped according to robot, dynamics, and linear solver. Each row shows the time taken

[3]The real robot has more DoF: grippers, neck, and one more DoF at the torso. For simplicity, we assumed those joints were fixed to zero.

TABLE I: Computation time[4] (in seconds) and number of iterations (within parenthesis) for each robot. The best computation time for each dynamics and each robot is highlighted in bold.

		Linear Solver	Barrier strategy (bar_murule) adaptiv	dampmpc	quality	Average time per iteration (s)
KUKA iiwa	Fwd D	MA27	**0.40** (5)	0.49 (6)	0.43 (5)	0.08 ± 0.002
		MA57	0.41 (5)	0.48 (6)	0.42 (5)	0.08 ± 0.001
		MA97	0.46 (5)	0.56 (6)	0.50 (5)	0.09 ± 0.002
	Inv D	MA27	**0.20** (4)	0.26 (5)	0.23 (4)	0.05 ± 0.002
		MA57	0.22 (4)	0.28 (5)	0.24 (4)	0.05 ± 0.001
		MA97	0.27 (4)	0.33 (5)	0.29 (4)	0.07 ± 0.002
ANYmal B	Fwd D	MA27	2.26 (7)	2.80 (9)	2.34 (7)	0.31 ± 0.021
		MA57	2.80 (10)	2.49 (9)	32.13 (99)	0.29 ± 0.020
		MA97	**2.21** (7)	2.76 (9)	2.26 (7)	0.31 ± 0.009
	Inv D	MA27	2.99 (13)	2.60 (11)	2.18 (9)	0.23 ± 0.005
		MA57	2.90 (13)	2.54 (11)	**2.10** (9)	0.22 ± 0.004
		MA97	3.21 (13)	2.81 (11)	2.37 (9)	0.25 ± 0.007
TALOS	Fwd D	MA27	—	—	14.82 (15)	0.94 ± 0.115
		MA57	—	13.47 (19)	44.23 (66)	0.68 ± 0.020
		MA97	13.42 (18)	**11.48** (15)	12.92 (16)	0.76 ± 0.026
	Inv D	MA27	8.38 (15)	8.02 (13)	38.59 (70)	0.57 ± 0.035
		MA57	7.36 (15)	**6.59** (13)	36.88 (73)	0.50 ± 0.014
		MA97	7.43 (15)	6.61 (13)	41.84 (82)	0.50 ± 0.012

to solve the optimization problem for each barrier update strategy, as well as the total number of iterations (within parenthesis). The last column shows the time spent on each iteration, averaged over all of the update strategies.

In general, we can see that the computation time depends mostly on the complexity of the system, regardless of linear solver or barrier update strategy; i.e., solving the manipulator task was faster than solving the quadruped task, which in turn was faster than the humanoid task. More importantly, for each robot and given the same choice of linear solver and update strategy, the computation time of inverse dynamics was better than forward dynamics. We can also see that the number of iterations required to solve the problem did not change significantly (apart from a few exceptions). This indicates that the difficulty of the problem itself did not change with the different dynamics defects; it just took longer to solve using forward dynamics—as supported by the information in the last column of the table.

B. Robustness to Coarser Problem Discretizations

In the next experiment, we compare the ability of each formulation to handle trajectories discretized using fewer mesh points. We defined the same quadruped jumping task repeatedly, but transcribed it with different resolutions. First, we divided the trajectory into equally spaced segments with a time step of $h = 0.01$; we solved the optimization problem and took the resulting trajectory as our baseline. Then, we incrementally changed h, making the problem more coarse each time, and compared the obtained trajectories against the baseline. The problems were initialized with a nominal configuration repeated for each point, and zero velocities, torques and contact forces. The results of this experiment are shown in Figure 3 and Table II.

[4]This table (and future tables) show the minimum time value observed over 10 trials. Reporting the minimum time is more reliable than the median or the mean, since all measured noise is positive, as explained in [26].

TABLE II: Computation time and number of iterations required by different problem discretizations, for the quadruped jump.

Frequency	Forward Dyn. Time (s)	Forward Dyn. Iter.	Inverse Dyn. Time (s)	Inverse Dyn. Iter.
100 Hz	3.289	10	3.295	13
90 Hz	3.916	14	2.774	13
80 Hz	5.227	21	1.821	9
70 Hz	4.062	19	1.344	8
60 Hz	1.518	9	1.193	8
50 Hz	2.226	16	0.983	8
40 Hz	1.099	9	0.785	8
30 Hz	0.731	8	0.534	7
20 Hz	0.433	7	0.354	7

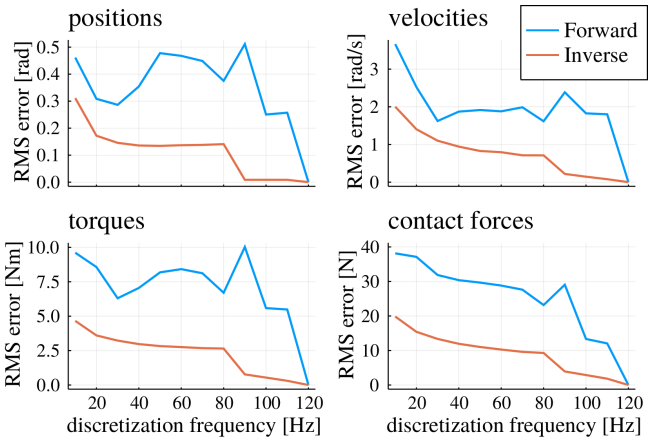

Fig. 3: Root-mean-square error (RMSE) of joint positions, velocities, torques, and contact forces of each formulation for different discretizations, using a baseline of 120 Hz.

The plots in Figure 3 show that the solutions deviate more from the baseline as the number of mesh points used to discretize the problem decreases (in the x-axis, from right to left). But more importantly, the plots reveal that the rate at which deviation occurs is significantly different depending on the formulation. We can see that the root-mean-square error (RMSE) of the formulation using inverse dynamics is significantly lower than that of the forward.

Table II shows the computation time (in seconds) and the number of iterations required to solve the quadruped task using different discretizations. We can see that the time required to solve the problem using inverse dynamics follows a clear pattern: it decreases as the problem gets more coarse; and the same goes for the number of iterations. In contrast, a pattern does not seem to exist for forward dynamics.

The results shown in Figure 3 and Table II provide strong evidence that defining the defect constraints with inverse dynamics is the approach more robust to different problem discretizations, both in terms of deviation from realistic solutions and in terms of computation performance.

C. Optimization with an Objective Function

Thus far we have analyzed the trajectory optimization performance for feasibility problems. However, in optimization, it is common to define a cost function to be minimized (or a value function to be maximized). In this next experiment, we evaluate the performance of our formulation when a cost function is considered. We minimize the

(a) Initial configuration (b) Take-off (c) Full-flight phase (d) Landing (e) Final configuration

Fig. 4: Snapshots of ANYmal [27] performing a $0.5\,\mathrm{m}$-long jump. The length of the black tape on the ground is $0.5\,\mathrm{m}$.

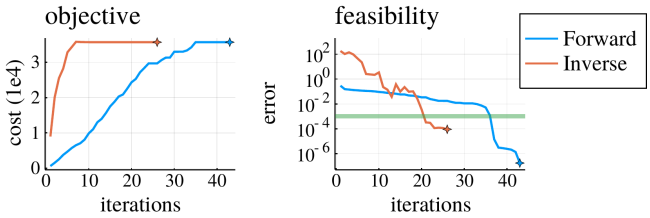

Fig. 5: Evolution of the cost and the feasibility error during convergence. The faint-green line at $y = 10^{-3}$ denotes the tolerance under which we consider a problem to be feasible.

actuator torques and ground-reaction contact forces with $\min_{\boldsymbol{\xi}} \sum_{k=1}^{M-1} \frac{\boldsymbol{\tau}_k^\top \boldsymbol{\tau}_k + \boldsymbol{\lambda}_k^\top \boldsymbol{\lambda}_k}{M-1}$. We tested this cost function on the quadruped jump task, with the MA57 linear solver and `adaptive` barrier parameter update strategy.

Both formulations converged to very similar solutions: the RMSE between the two trajectories was 0.038. The final objective value was $3.567\,801 \times 10^4$ and $3.567\,804 \times 10^4$ for forward and inverse dynamics, respectively. Despite converging to similar solutions, the formulation employing inverse dynamics finished in $6.208\,\mathrm{s}$, showing better performance than the formulation using forward dynamics which took $14.570\,\mathrm{s}$. The time in seconds corresponds to the minimum value measured over a total of 10 samples.

Figure 5 shows the evolution of the cost and feasibility error throughout the optimization. The star-shaped marker denotes the point at which the local minimum of the problem was found. In the left plot, we can see that inverse dynamics reached values close to the optimal cost much earlier than forward dynamics. In the right plot, we can see that inverse dynamics required less iterations than forward dynamics to cross the faint-green line, which marks the point at which the error becomes acceptable to be considered feasible. Inverse dynamics converged in 26 iterations, and forward dynamics converged in 43 iterations. Inadvertently, one advantage of the forward formulation was that its final feasibility error was smaller than that of inverse dynamics.

D. Hardware Validation

We conducted real-world experiments with ANYmal [27] and TALOS [2] to validate the trajectories computed with our framework. The motion planning is performed offline and then the trajectories are sent to the controller for playback. To execute the whole-body motions, we commanded each joint with feedforward torque and feedback on joint position and velocity. For the quadruped, we updated the references for each joint's position, velocity, and torque at $400\,\mathrm{Hz}$. The decentralized motor controller at every joint closes the loop compensating for friction effects. On the humanoid, we updated the references at $2\,\mathrm{kHz}$, and a centralized controller compensates for the motor dynamics and friction.

Figure 1 and Figure 4 contain snapshots of the jumps realized with the humanoid and with the quadruped, respectively. These experiments can be seen in our supplementary video: `https://youtu.be/pV4s7hzUgjc`. Jumping motion is challenging to execute in real hardware because it includes a severely underactuated phase when the robot is fully off the ground. Nonetheless, our controller is able to execute our planned trajectories reliably, attesting the dynamical consistency of our formulation.

VI. Discussion

The results of this work indicate that direct transcription implementations relying on forward dynamics to define the defect constraints can be reformulated with inverse dynamics to see an increase in performance, for both feasibility or minimization problems, and without sacrificing the feasibility of the solutions to the optimization problem. An additional reason to prefer inverse dynamics is robustness to coarser discretizations, both in terms of computation efficiency and faithfulness of solutions with respect to finer discretizations.

When minimizing a cost function, the locally-optimal solutions computed with either formulation are essentially the same. However, when an objective function is not considered, the formulations may diverge to different solutions. Experimentally, we have observed that the solutions computed with inverse dynamics are easier to perform in real hardware. The reasons behind this divergence are not yet clear to us, and this is something we plan to investigate in future work.

Erez and Todorov [13] observed a striking feature in their results: an emergent coordination between legs and opposite arms during a running gait. In this work, for the humanoid jumping task, we also observed such emerging behavior: the resulting motions swing the arms upwards to build-up energy before the take-off instant. Both in [13] and our work, these features originated without any explicit modeling—reaffirming the power of dynamic trajectory optimization.

In recent work [28], we took into account uncertainty and robustness to disturbances using direct transcription. Considering uncertainty usually incurs additional computational cost due to more complex problem formulations. With the findings from this paper, we plan to redefine the dynamics defect constraints in that work with inverse dynamics, improving the performance of our robustness framework and making it more competitive.

REFERENCES

[1] J. T. Betts, *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*, 2nd ed. SIAM, 2010.

[2] O. Stasse, T. Flayols *et al.*, "TALOS: A new humanoid research platform targeted for industrial applications," in *IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*, 2017.

[3] D. Pardo, L. Möller *et al.*, "Evaluating Direct Transcription and Nonlinear Optimization Methods for Robot Motion Planning," *IEEE Robotics and Automation Letters*, vol. 1, no. 2, pp. 946–953, 2016.

[4] T. Koolen and R. Deits, "Julia for robotics: simulation and real-time control in a high-level programming language," in *International Conference on Robotics and Automation (ICRA)*, 2019.

[5] S. M. Neuman, T. Koolen *et al.*, "Benchmarking and Workload Analysis of Robot Dynamics Algorithms," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2019.

[6] J. Carpentier, G. Saurel *et al.*, "The Pinocchio C++ library – A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives," in *IEEE International Symposium on System Integrations (SII)*, 2019.

[7] N. Schweighofer, M. A. Arbib, and M. Kawato, "Role of the cerebellum in reaching movements in humans. i. distributed inverse dynamics control," *European Journal of Neuroscience*, 1998.

[8] B. Mehta and S. Schaal, "Forward Models in Visuomotor Control," *Journal of Neurophysiology*, vol. 88, no. 2, pp. 942–953, 2002.

[9] T. Koolen and contributors, "Rigidbodydynamics.jl," 2016. [Online]. Available: https://github.com/JuliaRobotics/RigidBodyDynamics.jl

[10] M. L. Felis, "RBDL: an efficient rigid-body dynamics library using recursive algorithms," *Autonomous Robots*, pp. 1–17, 2016.

[11] M. Frigerio, J. Buchli *et al.*, "RobCoGen: a code generator for efficient kinematics and dynamics of articulated robots, based on Domain Specific Languages," *Journal of Software Engineering for Robotics (JOSER)*, vol. 7, no. 1, pp. 36–54, 2016.

[12] Sung-Hee Lee, Junggon Kim *et al.*, "Newton-type algorithms for dynamics-based robot movement optimization," *IEEE Transactions on Robotics (T-RO)*, vol. 21, no. 4, pp. 657–667, 2005.

[13] T. Erez and E. Todorov, "Trajectory optimization for domains with contacts using inverse dynamics," in *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2012, pp. 4914–4919.

[14] E. Todorov, "A convex, smooth and invertible contact model for trajectory optimization," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 1071–1076.

[15] M. Posa, C. Cantu, and R. Tedrake, "A direct method for trajectory optimization of rigid bodies through contact," *The International Journal of Robotics Research (IJRR)*, vol. 33, no. 1, pp. 69–81, 2014.

[16] H. Ferrolho and contributors, "TORA.jl," 2020. [Online]. Available: https://github.com/JuliaRobotics/TORA.jl

[17] P. G. Gormley, "Stereographic Projection and the Linear Fractional Group of Transformations of Quaternions," *Proceedings of the Royal Irish Academy. Mathematical and Physical Sciences*, vol. 51, 1945.

[18] G. Terzakis, M. Lourakis, and D. Ait-Boudaoud, "Modified Rodrigues Parameters: An Efficient Representation of Orientation in 3D Vision and Graphics," *Journal of Mathematical Imaging and Vision*, 2018.

[19] A. W. Winkler, D. C. Bellicoso *et al.*, "Gait and Trajectory Optimization for Legged Systems through Phase-based End-Effector Parameterization," *IEEE Robotics and Automation Letters*, 2018.

[20] S. Tonneau, D. Song *et al.*, "SL1M: Sparse L1-norm Minimization for contact planning on uneven terrain," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2019.

[21] T. Stouraitis, I. Chatzinikolaidis *et al.*, "Online Hybrid Motion Planning for Dyadic Collaborative Manipulation via Bilevel Optimization," *IEEE Transactions on Robotics (T-RO)*, vol. 36, no. 5, 2020.

[22] S. Caron, Q. Pham, and Y. Nakamura, "Leveraging Cone Double Description for Multi-contact Stability of Humanoids with Applications to Statics and Dynamics," in *Robotics: Science and System*, 2015.

[23] J. Bezanson, A. Edelman *et al.*, "Julia: A Fresh Approach to Numerical Computing," *SIAM Review*, vol. 59, 2017.

[24] R. H. Byrd, J. Nocedal, and R. A. Waltz, *Knitro: An Integrated Package for Nonlinear Optimization*. Springer US, 2006.

[25] R. A. Waltz, J. L. Morales *et al.*, "An interior algorithm for nonlinear optimization that combines line search and trust region steps," *Mathematical Programming*, vol. 107, no. 3, pp. 391–408, 2006.

[26] J. Chen, J. Revels, and A. Edelman, "Robust benchmarking in noisy environments," in *Proceedings of the 20th IEEE High Performance Extreme Computing Conference*, Waltham, USA, 2016.

[27] M. Hutter, C. Gehring *et al.*, "ANYmal - toward legged robots for harsh environments," *Advanced Robotics*, 2017.

[28] H. Ferrolho, W. Merkt *et al.*, "Optimizing Dynamic Trajectories for Robustness to Disturbances Using Polytopic Projections," in *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2020.