

Constrained Articulated Body Dynamics Algorithms

Ajay Suresha Sathya^{1,2} and Justin Carpentier¹

Abstract—Rigid-body dynamics algorithms have played an essential role in robotics development. By finely exploiting the underlying robot structure, they allow the computation of the robot kinematics, dynamics, and related physical quantities with low complexity, enabling their integration into chipsets with limited resources or their evaluation at very high frequency for demanding applications (e.g., model predictive control, large-scale simulation, reinforcement learning, etc.). While most of these algorithms operate on constraint-free settings, only a few have been proposed so far to adequately account for constrained dynamical systems while depicting low algorithmic complexity. In this article, we introduce a series of new algorithms with reduced (and lowest) complexity for the forward simulation of constrained dynamical systems. Notably, we revisit the so-called articulated body algorithm (ABA) and the Popov-Vereshchagin algorithm (PV) in the light of proximal-point optimization and introduce two new algorithms, called constrainedABA and proxPV. These two new algorithms depict linear complexities while being robust to singular cases (e.g., redundant constraints, singular constraints, etc.). We establish the connection with existing literature formulations, especially the relaxed formulation at the heart of the MuJoCo and Drake simulators. We also propose an efficient and new algorithm to compute the damped Delassus inverse matrix with the lowest known computational complexity. All these algorithms have been implemented inside the open-source framework Pinocchio and depict, on a wide range of robotic systems ranging from robot manipulators to complex humanoid robots, state-of-the-art performances compared to alternative solutions of the literature.

I. INTRODUCTION

Efficient rigid body dynamics algorithms [1] have played an essential role in robotics development. By finely exploiting the underlying robot structure to compute the robot kinematics, dynamics, and related physical quantities with low computational complexity, these algorithms enable dynamics evaluation in chip sets with limited resources and at high frequencies for demanding applications (e.g., computed torque control, model predictive control, large-scale simulation, reinforcement learning, etc.).

Most simulators [2]–[6] currently use such low-complexity algorithms only in constraint-free settings. Though a few low-complexity algorithms have been proposed for constrained dynamical systems [7]–[10], they suffer from being fairly complex to derive and implement. Also, they cannot adequately deal with singular cases (e.g., redundant constraints, singular constraints, etc.). Perhaps due to these issues, the usage of efficient constrained dynamics algorithms in simulators is

currently low. Addressing these issues, we present a series of new constrained dynamics algorithms based on proximal algorithms in this paper that are simple and effectively handle singular cases. Our proximal perspective also presents a unified framework to interpret several existing disparate algorithms as special cases of the presented algorithms.

A. Related work

The first $O(n)$ complexity unconstrained forward dynamics algorithm, where n is the degrees of freedom (DoF) of the robot, was proposed by Vereshchagin in [11], but was unknown in west for decades. Their approach solved the linear quadratic regulator (LQR) problem associated with the Gauss’ principle of least constraint (GPLC) [12], which is an optimization-based formulation of classical mechanics, using the dynamic programming (DP) principle [13]. Nearly a decade later, a practically identical algorithm [14] called the articulated body algorithm (ABA) was proposed, which was derived using a different perspective of cleverly propagating the solution of Newton-Euler equations through a kinematic tree. The ABA algorithm was further refined in [15] with insights such as improvements in computational efficiency through the use of local frames. Vereshchagin’s LQR-ABA connection was also independently discovered by Western researchers in [16], where the similarities between the ABA algorithm and Kalman filter were recognized to derive ABA using filtering theory literature. However, the approach in [16] is fairly complex, and [11] (see [10] for an expository derivation) remains a significantly clearer connection between LQR and ABA.

These $O(n)$ complexity algorithms are efficient compared to computing the joint-space inertia matrix (JSIM) [17] and factorizing it, which has a complexity of $O(n^3)$. However, branching in kinematic trees induces sparsity in the JSIM, which can be exploited by the Cholesky decomposition (LTL algorithm) developed in [18] with $O(nd^2)$ complexity, where d is the depth of the kinematic tree. The LTL algorithm is faster than the $O(n)$ algorithms for small robots ($n < 7$) and can be competitive even for robots as large as quadrupeds ($n = 18$), due to branching in the kinematic tree.

The unconstrained forward dynamics algorithm in [11] was quickly extended to account for motion constraints in the Popov-Vereshchagin algorithm (PV algorithm) [7], [19] with $O(n + m^2d + m^3)$ complexity, where m is the dimensionality of motion constraints, by encoding the constraints as equality constraints in the associated LQR problem. However, the PV algorithm remained sparsely known and used in the robotics community. Addressing this issue, a recent work [10] provides an accessible and expository derivation of the PV

¹ Inria - Département d’Informatique de l’École normale supérieure, PSL Research University. {ajay.sathya, justin.carpentier}@inria.fr
²MECO Research Team, Department of Mechanical Engineering, KU Leuven and Flanders Make@KU Leuven, Belgium {ajay.sathya}@kuleuven.be

algorithm along with two original extensions, namely, PV-soft and PV-early algorithms, each with only $O(n+m)$ complexity. The PV-soft algorithm relaxes the motion constraints with a quadratic penalty, allowing it to obtain low complexity. However, it does not solve the constraints exactly. On the other hand, the PV-early algorithm exactly solves the motion constraints and yet achieves its low complexity by aggressively eliminating the Lagrange multipliers associated with motion constraints as early as possible. This prevents the accumulation of Lagrange multipliers during the DP recursions, that occurs in the original PV algorithm contributing to the $O(m^2 + m^3)$ complexity terms. The PV-early algorithm uses the singular value decomposition (SVD) [20] of a symmetric matrix at every joint, whose rank is equal to the respective joint's DoF, for early elimination of the Lagrange multipliers. While an efficient analytical formula exists for computing SVD of symmetric rank one (SR-1) matrices due to single DoF joints, the general applicability of the PV-early algorithm is impacted by the computational expense of SVD for multi-DoF joints. Moreover, the PV algorithm is sensitive to singular/redundant constraints, and the popular approach of Tikhonov regularization employed in [10] does not adequately solve constrained dynamics as it can bias solutions towards the origin and adversely affect constraint satisfaction.

The most widely used constrained dynamics algorithm for kinematic trees [21] solves the problem in joint space, resulting in a relatively higher computational complexity of $O(nd^2 + m^2d + md^2 + m^3)$ despite exploiting branching-induced sparsity. Recent work [22] generalized this approach to robustly handle singular/redundant constraints using proximal algorithms [23]. A similar proximal perspective can be applied to the low-complexity constrained dynamics algorithms to derive efficient and robust algorithms, forming the basis of this paper.

The main bottleneck in the joint-space constraint dynamics algorithms [21], [22] is the computation and factorization of the Delassus matrix [24], [25], also known as the inverse operational space inertia (OSIM) matrix [26], which itself has a computational complexity of $O(nd^2 + m^2d + m^3)$ [21]. The Delassus matrix computation is crucial to many tasks, such as contact simulation, differentiating constrained dynamics, and operational space control, just to name a few. Therefore, unsurprisingly, several dedicated low-complexity algorithms have been proposed for computing it. One of the first such algorithms was the Kreuz-Jain-Rodriguez (KJR) algorithm [27], which used a three-sweep computation structure similar to ABA to achieve $O(n+m^2d+m^3)$. A further improvement on this propagation-based algorithm was proposed in EFPA [28] using extended force propagators (EFP) [29] to achieve a lower complexity of $O(n+md+m^3)$. Recent work [10] also showed that the PV algorithm computes the Delassus matrix as an intermediate quantity, which yielded another efficient algorithm, PV-OSIM, which requires only two computation sweeps, unlike EFPA or KJR. Despite its higher computational complexity of $O(n+m^2d+m^3)$, PV-OSIM was found to be more efficient [10] than the EFPA for most realistic robots. Finally, the PV-OSIM's computational complexity was further optimized in [30] by using the compositionality of

EFPs to the lowest known complexity of $O(n+m^3)$. All the above algorithms have the m^3 complexity term due to the factorization of the Delassus matrix, which is generally dense. Specifically for robots with branching at the base (e.g., quadrupeds, humanoids), it was shown [10] that using matrix inversion lemma (MIL) [31] could accelerate the inverse Delassus matrix computation by up to 30%.

Though we restrict our scope in this paper to kinematic trees with motion constraints relative to the inertial ground frame, we mention constrained dynamics algorithms for closed loops for the sake of completeness. Most robot simulators do not cope with closed-loop mechanisms properly, and often, when they do, they resort to using inefficient methods like using general-purpose linear solvers to solve the contact KKT matrix. Recent work in [22] extended Featherstone's LTL algorithm [18], [21] to account for loop closure constraints with the same computational complexity of $O(n^2d+nd^2+md^2+m^3)$ as the original LTL algorithm and provided an implementation within the open source robot dynamics simulator library PINOCCHIO [6]. However, more efficient algorithms were proposed in the past, which can be an improvement over [22]. In the late 1980s, the PV algorithm was independently rediscovered and extended to a more general class of mechanisms, including those with internal closed loops in [9] and [8]. Both these algorithms are practically identical with the worst-case complexity of $O(n+m^2d+m^3)$ and employ loop-cutting with zero mass phantom links to be able to use a propagation-based algorithm similar in spirit to ABA. However, these algorithms are quite complex and are not yet supported in robotics simulators.

B. Contributions

Our paper contains five main contributions, three of which correspond to algorithmic contributions:

- (i) **The constrainedABA algorithm:** We introduce a $O(n+m)$ complexity-constrained dynamics algorithm for trees called constrainedABA (cABA). Its derivation and implementation are significantly simpler than the existing $O(n+m)$ algorithms like the PV-early while being more efficient and more generally applicable to singular cases.
- (ii) **The proxPV algorithm:** Revisiting the PV algorithm in the light of proximal-point optimization, we introduce proxPV that makes the PV algorithm robust to singular cases with minimal computational overhead.
- (iii) **The cABA-OSIM algorithm:** We introduce the most efficient known algorithm with $O(n+m^2)$ complexity to compute the damped Delassus inverse matrix, called cABA-OSIM. In contrast, all the existing algorithms have an additional $O(m^3)$ complexity term in the general case due to the Delassus matrix factorization.
- (iv) **Analysis and detailed benchmarking:** Our proximal point perspective presents a unifying framework highlighting connections between seemingly disparate solvers like constrainedABA, PV-early, and the soft Gauss' principle formulation at the heart of the MUJOCO simulator and PV-soft algorithms. We present a detailed benchmarking of the different algorithms regarding the number

of operations and computational speed for various robot models.

- (v) **Open-source software implementation:** The constrainedABA algorithm proposed here is available to the community in the open-source software toolbox PINOCCHIO¹ [6].

C. Paper organization

In the rest of this section, we introduce notation used in the paper. In the next section, we review and introduce proximal formulation of constrained dynamics, as well as the proxLTL, proxLTLs and PV algorithms. Section.III subsequently derives the constrained ABA algorithm, which is followed by an algorithmic presentation in Section IV. Subsequently, the fourth proximal algorithm proxPV is introduced in Section V, which is followed by the damped Delassus inverse algorithm in Section III. The connections between the proposed algorithms and existing algorithms are analyzed in Section VII before performing detailed benchmarking of the proposed algorithms in Section VIII. The paper is finally concluded with final discussions and conclusions in Section IX and Section X.

D. Notation

Lower-case symbols (x), bold-faced lower-case symbols (\mathbf{x}), and upper-case symbols (X) are scalars, vectors, and matrices, respectively. Let \mathbb{S}_+^k and \mathbb{S}_{++}^k denote the space of symmetric positive semidefinite and positive definite matrices of size $k \times k$.

Let $\mathbf{q} \in \mathcal{Q}$, $\boldsymbol{\nu} \in \mathcal{T}_{\mathbf{q}}\mathcal{Q} \simeq \mathbb{R}^n$ and $\dot{\boldsymbol{\nu}}$ be the robot generalized configuration, generalized velocity, and generalized accelerations respectively. \mathcal{Q} , $\mathcal{T}_{\mathbf{q}}\mathcal{Q}$ and n are the robot's configuration space, \mathcal{Q} 's tangent space at \mathbf{q} and degrees of freedom (DoF) respectively. Let $\boldsymbol{\tau} \in \mathcal{T}_{\mathbf{q}}^*\mathcal{Q} \simeq \mathbb{R}^n$ be the generalized forces exerted on the robot. Let $\mathcal{S} = \{1, 2, \dots, n_b\}$ be an ordered list of a kinematic tree's link indices such that $i < j$ if the i^{th} link is an ancestor of the j^{th} link. By convention, we assume the 0th link to be a global inertial frame. For floating base robots, 1st link is the floating base. The acceleration of the i^{th} joint is denoted as $\dot{\boldsymbol{\nu}}_i$. Occasionally, we refer to the accelerations of a set of joints $\{\dot{\boldsymbol{\nu}}_i, i \in \gamma\}$, which we will conveniently represent as $\dot{\boldsymbol{\nu}}_\gamma$. The difference between $\dot{\boldsymbol{\nu}}_i$ and $\dot{\boldsymbol{\nu}}_\gamma$ will be clear from the context.

We will use Featherstone's spatial algebra [1] to refer to rigid body quantities. Let ${}^i X_j \in SE(3)$ (special Euclidean group in 3 dimensions) be the spatial transform of the j^{th} link in the i^{th} link's frame expressed as a 6×6 spatial transformation matrix. The 6D spatial velocity and acceleration of a rigid body i is $\mathbf{v}_i \in \mathbb{M}^6$ and $\mathbf{a}_i \in \mathbb{M}^6$ respectively, where \mathbb{M}^6 is a spatial motion vector space. The spatial forces acting on the i^{th} body is $\mathbf{f}_i \in \mathbb{F}^6$, where \mathbb{F}^6 is the spatial force vector space that is dual to \mathbb{M}^6 . The spatial inertia matrix $H_i \in \mathbb{I}^{6 \times 6}$ (represented as a \mathbb{S}_{++}^6 matrix) maps \mathbb{M}^6 to \mathbb{F}^6 . The cross-product operators on motion and force vectors are denoted as \times and \times^* respectively. Refer [1] for more details on the spatial algebra. Table I gives a comprehensive list of the symbols used all over the paper.

Symbol	Meaning
d	Depth of the kinematic tree.
\mathbf{q}	Robot configuration.
\mathcal{Q}	Robot configuration space.
$\boldsymbol{\nu}$	Generalized robot velocities.
$\mathcal{T}_{\mathbf{q}}\mathcal{Q}$	Tangent space of \mathcal{Q} at \mathbf{q} .
$\dot{\boldsymbol{\nu}}$	Generalized robot accelerations.
$\boldsymbol{\tau}$	Generalized robot forces.
$\mathcal{T}_{\mathbf{q}}^*\mathcal{Q}$	Dual tangent space of \mathcal{Q} at \mathbf{q} .
\mathcal{S}	Topologically ordered list of tree link indices.
\mathbf{v}_i	i^{th} link's 6D spatial velocity.
\mathbb{M}^6	Motion vector space in spatial algebra.
\mathbf{a}_i	i^{th} link's 6D spatial acceleration.
\mathbf{f}_i	6D spatial forces acting on the i^{th} link.
\mathbb{F}^6	Force vector space, dual of \mathbb{M}^6 .
H_i	i^{th} link's 6D spatial inertia tensor.
\times	Cross-product operator on spatial motion vectors.
\times^*	Cross-product operator on spatial force vectors.
\mathbf{f}_c	Motion constraint function in generalized coordinate space.
M	Robot inertia matrix in generalized coordinates (JSIM).
$\dot{\boldsymbol{\nu}}_{\text{free}}$	Generalized accelerations when unconstrained.
\mathbf{v}_c^*	Desired constraint velocities.
\mathbf{a}_c^*	Desired constraint accelerations.
$\mathbf{J}_{\mathbf{f}_c}$	Geometric Jacobian of \mathbf{f}_c .
\mathbf{c}	Generalized forces due to gravity, centripetal and Coriolis effects.
$\boldsymbol{\gamma}_{\mathbf{f}_c}$	Constraint accelerations due to Coriolis effects.
$\boldsymbol{\lambda}$	Lagrange multipliers and constraint force magnitudes.
\mathcal{L}	Lagrangian of GPLC.
Λ	OSIM or inverse Delassus matrix.
μ	Proximal operator parameter.
\mathbf{g}	Dual function of GPLC.
Λ_μ	Damped Delassus inverse matrix.
M_μ	Constraint augmented Inertia matrix.
\mathcal{S}_i	Spans i^{th} joint's motion subspace.
$\mathbf{a}_{b,i}$	i^{th} link's bias acceleration vector.
K_i	i^{th} link's constraint matrix.
\mathbf{k}_i	i^{th} link's desired constraint accelerations.
\mathbf{a}_{grav}	Spatial acceleration-due-to-gravity vector.
$\pi(i)$	i^{th} link's parent link.
$\gamma(i)$	Set of i^{th} link's children links.
\tilde{H}_i	i^{th} link's constraint augmented inertia tensor Eq. (32).
$\tilde{\mathbf{f}}_i$	i^{th} link's constraint augmented force vector Eq. (32).
$\text{desc}(i)$	Set of all descendants of the i^{th} link.
H_i^A	i^{th} link's constrained articulated body inertia.
\mathbf{f}_i^A	i^{th} link's resultant force due to all constraints and forces on descendant links.
D_i	Apparent constrained inertia felt at the i^{th} joint.
P_i	Backward force propagation matrix at the i^{th} joint.
$\text{ES}(i)$	Set of motion-constrained links among i^{th} link's descendants.
\mathcal{S}_r	List \mathcal{S} reversed.
\mathcal{E}	Set of all links that are motion-constrained (end-effectors).
$\mathcal{C}(\cdot)$	Returns cardinality of a set.
$\mathcal{S}_{\mathcal{E}}$	Sublist of \mathcal{S} , whose links support an end-effector.
$\mathcal{S}_{\mathcal{E}r}$	List $\mathcal{S}_{\mathcal{E}}$ reversed.
$\Delta \mathbf{f}_i$	Force update vector in constrainedABA reduced sweep.
K_i^A	Constraint matrix felt at i^{th} link for all $\text{ES}(i)$ constraints.
L_i^A	Intermediate term to recursively accumulate Λ^{-1} .
R_E^{-1}	Diagonal constraint weighting matrix in soft Gauss' principle.
${}^j P_i$	EFP - backward propagates i^{th} link's forces to j^{th} link.
$\text{cca}(i, j)$	i^{th} and j^{th} links' closest common ancestor.
${}^j \Omega_i$	Spatial inverse inertia of the i^{th} link considering a sub-tree rooted at the j^{th} link.
\mathcal{N}	Set of branching links supporting strictly more of constraints than any child link.
${}^j K_i$	Constraint-space EFP.
${}^i L_{e_j, e_k}$	Delassus matrix block for e_j and e_k constraints for a sub-tree rooted at i^{th} link.

TABLE I: List of symbols and notations used in the article.

¹<https://github.com/stack-of-tasks/pinocchio>

II. PROXIMAL REFORMULATION OF CONSTRAINED DYNAMICS

A. Constrained dynamics in generalized coordinates

Constrained dynamics. Suppose that the robot motion is constrained by the following holonomic (sometimes also called rheonomic [1, Sec.3.4] due to time-dependence) constraint,

$$\mathbf{f}_c(\mathbf{q}, t) = 0, \quad (1)$$

where $\mathbf{f}_c : \mathcal{Q} \times \mathbb{R} \mapsto \mathbb{R}^m$ is a smooth function, differentiating which with respect to time gives

$$J_{\mathbf{f}_c}(\mathbf{q}, t)\boldsymbol{\nu} = \mathbf{v}_c^*(\mathbf{q}, t), \quad (2)$$

where $J_{\mathbf{f}_c} \in \mathbb{R}^{m \times n}$ is the geometric Jacobian of the constraints, and $\mathbf{v}_c^* = -\frac{\partial \mathbf{f}_c}{\partial t} \in \mathbb{R}^m$. We note that even non-holonomic constraints can be expressed in the form of Eq. (2) [32, Section 7.1]. We will omit explicitly writing the time-dependence from now on for notational brevity. We differentiate the constraint Eq. (2) with respect to time one more time and impose them in Gauss' principle of least constraint [12], [33], [34], according to which a robot at state $(\mathbf{q}, \boldsymbol{\nu})$ undergoes accelerations that minimize the following equality-constrained strongly convex quadratic program (QP):

$$\underset{\boldsymbol{\nu}}{\text{minimize}} \quad \frac{1}{2} \|\dot{\boldsymbol{\nu}} - \dot{\boldsymbol{\nu}}_{\text{free}}(\mathbf{q}, \dot{\boldsymbol{\nu}}, \boldsymbol{\tau})\|_{M(\mathbf{q})}^2 \quad (3a)$$

$$\text{subject to} \quad J_{\mathbf{f}_c}(\mathbf{q})\dot{\boldsymbol{\nu}} + \dot{J}_{\mathbf{f}_c}(\mathbf{q}, \boldsymbol{\nu})\boldsymbol{\nu} = \mathbf{a}_c^*, \quad (3b)$$

where $M(\mathbf{q}) \in \mathbb{S}_{++}^n$ is the joint-space inertia matrix (JSIM) and $\dot{\boldsymbol{\nu}}_{\text{free}}(\mathbf{q}, \dot{\boldsymbol{\nu}}, \boldsymbol{\tau})$ is the joint-space acceleration in the absence of constraints

$$\dot{\boldsymbol{\nu}}_{\text{free}}(\mathbf{q}, \dot{\boldsymbol{\nu}}, \boldsymbol{\tau}) := M^{-1}(\mathbf{q})(\boldsymbol{\tau} - \mathbf{c}(\mathbf{q}, \boldsymbol{\nu})), \quad (4)$$

where $\mathbf{c}(\mathbf{q}, \boldsymbol{\nu})$ is the generalized force vector due to gravity, Coriolis and centripetal effects. $\mathbf{a}_c^* = \dot{\mathbf{v}}_c^* \in \mathbb{R}^m$ is the desired constraint acceleration. In practice the constraints accumulate numerical integration errors during simulation and can be stabilized by including a constraint error (residual of Eq. (1)) negative feedback term in \mathbf{a}_c^* , with methods like Baumgarte's stabilization [35]. Let

$$\boldsymbol{\gamma}_{\mathbf{f}_c}(\mathbf{q}, \boldsymbol{\nu}) := \dot{J}_{\mathbf{f}_c}(\mathbf{q}, \boldsymbol{\nu})\boldsymbol{\nu}, \quad (5)$$

be the bias accelerations. The variable dependencies will be dropped for brevity whenever obvious from the context.

Constrained dynamics Lagrangian. The solution to the QP above is the primal-dual saddle point of the Lagrangian [36]

$$(\dot{\boldsymbol{\nu}}^*, \boldsymbol{\lambda}^*) \in \underset{\boldsymbol{\lambda}}{\text{arg max}} \underset{\dot{\boldsymbol{\nu}}}{\text{min}} \mathcal{L}(\dot{\boldsymbol{\nu}}, \boldsymbol{\lambda}), \quad (6)$$

where

$$\mathcal{L}(\dot{\boldsymbol{\nu}}, \boldsymbol{\lambda}) := \frac{1}{2} \|\dot{\boldsymbol{\nu}} - \dot{\boldsymbol{\nu}}_{\text{free}}\|_M^2 + \boldsymbol{\lambda}^T (J_{\mathbf{f}_c} \dot{\boldsymbol{\nu}} + \boldsymbol{\gamma}_{\mathbf{f}_c} - \mathbf{a}_c^*), \quad (7)$$

and $\boldsymbol{\lambda} \in \mathbb{R}^m$ are the Lagrange multipliers (also called dual variables) of the QP in Eq. (3).

Since Eq. (3) is an equality-constrained QP, its primal-dual saddle point is immediately given by solving the Karush-Kuhn-Tucker (KKT) [36] system implied by the first-order necessary optimality conditions [37],

$$\begin{bmatrix} M & J_{\mathbf{f}_c}^T \\ J_{\mathbf{f}_c} & 0_{m \times m} \end{bmatrix} \begin{bmatrix} \dot{\boldsymbol{\nu}} \\ \boldsymbol{\lambda} \end{bmatrix} = \begin{bmatrix} M \dot{\boldsymbol{\nu}}_{\text{free}} \\ -\boldsymbol{\gamma}_{\mathbf{f}_c} + \mathbf{a}_c^* \end{bmatrix}. \quad (8)$$

Solving the KKT system. Because M is positive definite, the KKT system in Eq. (8) is most often solved by eliminating $\dot{\boldsymbol{\nu}}$ using

$$\dot{\boldsymbol{\nu}}^* = \dot{\boldsymbol{\nu}}_{\text{free}} - M^{-1} J_{\mathbf{f}_c}^T \boldsymbol{\lambda}, \quad (9)$$

back-substituting which in Eq. (6), gives the dual function

$$g(\boldsymbol{\lambda}) = -\frac{1}{2} \boldsymbol{\lambda}^T \Lambda^{-1} \boldsymbol{\lambda} + (J_{\mathbf{f}_c} \dot{\boldsymbol{\nu}}_{\text{free}} + \boldsymbol{\gamma}_{\mathbf{f}_c} - \mathbf{a}_c^*)^T \boldsymbol{\lambda}, \quad (10)$$

where $\Lambda^{-1}(\mathbf{q}) := J_{\mathbf{f}_c} M^{-1} J_{\mathbf{f}_c}^T$ is the so-called Delassus matrix [24], [25], also known as the inverse operational space inertial matrix (inverse OSIM) [26].

Remark 1. $\Lambda^{-1} \in \mathbb{S}_+^m$ and the dual function above is concave. Furthermore, if $J_{\mathbf{f}_c}$ is full row-rank, $\Lambda^{-1} \in \mathbb{S}_{++}^m$, the dual function $g(\boldsymbol{\lambda})$ is strongly concave, Λ^{-1} is invertible and the OSIM matrix Λ exists.

The dual function is maximized to obtain the optimal Lagrange multipliers by solving the linear equation

$$\Lambda^{-1} \boldsymbol{\lambda}^* = J_{\mathbf{f}_c} \dot{\boldsymbol{\nu}}_{\text{free}} + \boldsymbol{\gamma}_{\mathbf{f}_c} - \mathbf{a}_c^*. \quad (11)$$

The frequently used approach to solve the equation above is computing M using the composite rigid body algorithm (CRBA) [17] and its Cholesky factorization using Featherstone's LTL algorithm [18], which exploits branching-induced sparsity. Then, the LTL algorithm is exploited to efficiently compute Λ^{-1} as well [21], which is finally factorized using dense Cholesky factorization. However, as mentioned in Section I-A, this approach is computationally expensive with a total computational complexity of $O(nd^2 + m^2d + md^2 + m^3)$, where d is the depth of the kinematic tree, n the number of joints and m is the number of constraints.

Nonuniqueness of solutions. In practice, $J_{\mathbf{f}_c}$ often does not have full rank or loses it due to redundant constraints or kinematic singularities, making the Λ^{-1} singular. This prevents the use of efficient algorithms like Cholesky factorization. Typically, Tikhonov regularization (adding a positive definite diagonal matrix to Λ^{-1}) is employed to address this issue. But Tikhonov regularization can be interpreted as a weighted ℓ_2 norm squared penalty on constraint force magnitudes ($\|\boldsymbol{\lambda}\|_R^2$). This penalizes high constraint forces thereby biasing $\boldsymbol{\lambda}$ towards the origin, preventing constraint satisfaction and can pose numerical issues if the regularization parameter is too small. An alternate approach that does not introduce such bias is the truncated singular value decomposition (SVD) [20]. However, this is computationally more expensive than Cholesky factorization and is therefore not commonly used.

B. Proximal point algorithm

An exact and efficient alternative to Tikhonov regularization or SVD that we will leverage is the proximal point algorithm (PPA) [23], [38], which is effective for robotics problems [22], [39], [40], and most often requiring few iterations (each of which is efficient) to converge for robot dynamics problems [22].

We first introduce the proximal operator of a convex function, $f : \mathbb{R}^n \rightarrow \mathbb{R}$

$$\text{prox}_{\mu, f}(\mathbf{x}^k) := \mathbf{arg min}_{\mathbf{x}} \left\{ f(\mathbf{x}) + \frac{1}{2\mu} \|\mathbf{x} - \mathbf{x}^k\|^2 \right\}, \quad (12)$$

where $\mu \in (0, \infty)$ is the proximal operator parameter, which can be interpreted as a step-size [23]. Please note that we slightly abuse the notation and use ‘:=’ instead of ‘∈’ above because the optimization problem above has a unique solution [23, Section 1.1], hence **arg min** returns a singleton solution set. We will continue to use ‘:=’ or ‘=’ in the rest of the paper when the arg optimization problems have a unique solution. The PPA minimizes $f(\mathbf{x})$ via fixed-point iterations on the proximal operator until a termination criterion is met

$$\mathbf{x}^{k+1} = \text{prox}_{\mu, f}(\mathbf{x}^k). \quad (13)$$

C. Constrained dynamics using the two QP approaches

We now describe two QP approaches to solve Eq. (3), that are mathematically equivalent but differ computationally.

1) *Dual proximal point method (proxLTL)*: Applying PPA to optimize the dual function in Eq. (10) gives

$$\boldsymbol{\lambda}^{k+1} = \text{prox}_{\mu, -g}(\boldsymbol{\lambda}^k) = \mathbf{arg min}_{\boldsymbol{\lambda}} -g(\boldsymbol{\lambda}) + \frac{1}{2\mu} \|\boldsymbol{\lambda} - \boldsymbol{\lambda}^k\|^2, \quad (14a)$$

$$= \Lambda_{\mu} \left(J_{\mathbf{f}_c} \dot{\mathbf{v}}_{\text{free}} + \boldsymbol{\gamma}_{\mathbf{f}_c} - \mathbf{a}_c^* + \frac{1}{\mu} \boldsymbol{\lambda}^k \right), \quad (14b)$$

where $\Lambda_{\mu}^{-1} := \Lambda^{-1} + \frac{1}{\mu} I$ is the damped Delassus matrix. $\Lambda_{\mu}^{-1} \in \mathbb{S}_{++}^m$ and can be factorized efficiently using Cholesky decomposition. This factorization is computed once and reused, making each fixed point iteration fast. This algorithm, which we will call proxLTL, was proposed in [22] and was accompanied by an efficient C++ implementation in the PINOCCHIO library [6]. ProxLTL converges to the exact solution that satisfies motion constraints unlike the Tikhonov regularization approach.

Remark 2. *All the algorithmic developments in a paper can be trivially extended to the case $\Lambda_{\mu}^{-1} := \Lambda^{-1} + R_E^{-1}$, where R_E is a positive definite diagonal matrix. Such a choice can permit variable weights on different constraints.*

Assuming that the constraints are correctly modeled for physical systems, the primal problem in Eq. (3) is always feasible (meaning that the constraints are satisfiable). However, numerically, the primal problem in Eq. (3) can be infeasible, especially when there are redundant constraints and Baumgarte terms [35] are added to \mathbf{a}_c^* for constraint stabilization. Then, the dual problem is unbounded above, and the PPA iterations do not converge. However, in this situation, it has been shown

[41]–[43] that the primal residual converges to a desirable value that is optimal in the least squares sense during the PPA iterations. Therefore, the termination criteria for the convergence of the PPA should monitor the least-squares error given by

$$\|\Lambda^{-T}(\Lambda^{-1} \boldsymbol{\lambda} - [J_{\mathbf{f}_c} \dot{\mathbf{v}}_{\text{free}} + \boldsymbol{\gamma}_{\mathbf{f}_c} - \mathbf{a}_c^*])\|_{\infty}. \quad (15)$$

Advantageously, the PPA algorithm converges to the optimal value for feasible problems and to the least squares solution for numerically infeasible problems automatically without any modification to the algorithm.

2) *Augmented Lagrangian method (proxLTLs)*: An alternative to proxLTL that can solve the QP in Eq. (3) exactly is the augmented Lagrangian method [44], [45] (ALM), where the augmented Lagrangian function is

$$\mathcal{L}^A(\dot{\mathbf{v}}, \boldsymbol{\lambda}) := \mathcal{L}(\dot{\mathbf{v}}, \boldsymbol{\lambda}) + \frac{\mu}{2} \|J_{\mathbf{f}_c} \dot{\mathbf{v}} + \boldsymbol{\gamma}_{\mathbf{f}_c} - \mathbf{a}_c^*\|^2. \quad (16)$$

ALM iterations alternately optimize \mathcal{L}^A over its primal and dual variables

$$\dot{\mathbf{v}}^{k+1} = M_{\mu}^{-1} \left\{ M \dot{\mathbf{v}}_{\text{free}} - J_{\mathbf{f}_c}^T \left(\boldsymbol{\lambda}^k + \mu (\boldsymbol{\gamma}_{\mathbf{f}_c} - \mathbf{a}_c^*) \right) \right\}, \quad (17a)$$

$$\boldsymbol{\lambda}^{k+1} = \boldsymbol{\lambda}^k + \mu \left(J_{\mathbf{f}_c} \dot{\mathbf{v}}^{k+1} + \boldsymbol{\gamma}_{\mathbf{f}_c} - \mathbf{a}_c^* \right), \quad (17b)$$

where $M_{\mu} := M + \mu J_{\mathbf{f}_c}^T J_{\mathbf{f}_c}$ is the augmented JSIM, with the influence of the constraints from the quadratic term in the augmented Lagrangian function. This formulation, which we will call proxLTLs (proxLTL-soft), is interesting as M_{μ} can be efficiently computed with a small modification to the CRBA (see Section II-E) and efficiently factorized using the LTL algorithm, leading to an algorithm that has a lower computational complexity of $O(nd^2 + md)$, which is significantly faster than the $O(nd^2 + m^2d + md^2 + m^3)$ complexity of the proxLTL algorithm. This difference arises because proxLTLs does not compute or factorize the Λ^{-1} matrix. While the idea of proxLTLs was briefly mentioned in [22], it was neither implemented nor explored further, which we will do in this paper.

3) *Both approaches are equivalent*: ProxLTL and proxLTLs are mathematically equivalent because dual PPA and the ALM are equivalent [42], [46]. Theoretically, the sequence of $\boldsymbol{\lambda}^k$ they compute are equal. However, they differ numerically and computationally, with the difference arising due to the order in which the primal and dual variables are optimized. We now show the equivalence of proxLTL and proxLTLs by recovering the ALM iterations in Eq. (17a) and Eq. (17b) from the PPA iterations in Eq. (14).

ProxLTL in Eq. (14) maximizes the dual function $g(\cdot)$, that is obtained after minimizing the Lagrangian primal variables in the saddle point problem (see Eq. (6)). However, since each PPA iteration solves a feasible sub-problem that is strongly convex and strongly concave in primal and dual variables, re-

spectively with a unique primal-dual solution, this elimination ordering can be reversed [47]

$$\text{prox}_{\mu, -g}(\boldsymbol{\lambda}^k) = \mathbf{arg} \min_{\boldsymbol{\lambda}} \left[-\min_{\dot{\boldsymbol{\nu}}} \mathcal{L}(\dot{\boldsymbol{\nu}}, \boldsymbol{\lambda}) + \frac{1}{2\mu} \|\boldsymbol{\lambda} - \boldsymbol{\lambda}^k\|^2 \right], \quad (18a)$$

$$= \mathbf{arg} \max_{\boldsymbol{\lambda}} \left[\min_{\dot{\boldsymbol{\nu}}} \mathcal{L}(\dot{\boldsymbol{\nu}}, \boldsymbol{\lambda}) - \frac{1}{2\mu} \|\boldsymbol{\lambda} - \boldsymbol{\lambda}^k\|^2 \right], \quad (18b)$$

$$= \mathbf{arg} \min_{\boldsymbol{\lambda}} \max_{\dot{\boldsymbol{\nu}}} \left[\mathcal{L}(\dot{\boldsymbol{\nu}}, \boldsymbol{\lambda}) - \frac{1}{2\mu} \|\boldsymbol{\lambda} - \boldsymbol{\lambda}^k\|^2 \right], \quad (18c)$$

where we define $\mathbf{arg} \min_{\boldsymbol{\lambda}} \max_{\dot{\boldsymbol{\nu}}} (\cdot) := \left\{ \boldsymbol{\lambda}^{k+1} : (\boldsymbol{\lambda}^{k+1}, \dot{\boldsymbol{\nu}}^{k+1}) \in \mathbf{arg} \min_{\boldsymbol{\lambda}} \max_{\dot{\boldsymbol{\nu}}} (\cdot) \right\}$. The maximizer over $\boldsymbol{\lambda}$ is obtained by solving for gradient stationarity

$$\boldsymbol{\lambda}^{k+1} = \boldsymbol{\lambda}^k + \mu (J_{\mathbf{f}_c} \dot{\boldsymbol{\nu}} + \boldsymbol{\gamma}_{\mathbf{f}_c} - \mathbf{a}_c^*), \quad (19)$$

substituting which back in Eq. (18c) in the place of $\boldsymbol{\lambda}$ and optimizing for the primal variables $\dot{\boldsymbol{\nu}}$ gives

$$\dot{\boldsymbol{\nu}}^{k+1} = M_{\mu}^{-1} \left\{ M \dot{\boldsymbol{\nu}}_{\text{free}} - J_{\mathbf{f}_c}^T \left(\boldsymbol{\lambda}^k + \mu (\boldsymbol{\gamma}_{\mathbf{f}_c} - \mathbf{a}_c^*) \right) \right\}, \quad (20)$$

which can be substituted back in Eq. (19) to compute the numerical value of $\boldsymbol{\lambda}^{k+1}$. Comparing the equations Eq. (19) and Eq. (20) with Eq. (17b) and Eq. (17a) verifies that proxLTL iterations and the ALM iterations are identical.

D. Constrained dynamics in maximal coordinates

Both proxLTL and proxLTLs algorithms have lower complexity counterparts that can be derived by applying dynamic programming (DP) on Gauss' principle formulated in the so-called maximal coordinates. Consider a link i constrained as

$$\mathbf{f}_c^{i,m}({}^0X_i, t) = \mathbf{0}, \quad (21)$$

where $\mathbf{f}_c^{i,m} : SE(3) \times \mathbb{R} \mapsto \mathbb{R}^{m_i}$ is a smooth function, differentiating which with respect to time gives

$$K_i({}^0X_i, t) \mathbf{v}_i = \mathbf{k}_i^v({}^0X_i, t), \quad (22)$$

where $K_i \in \mathbb{R}^{m_i \times 6}$ is the geometric Jacobian of $\mathbf{f}_c^{i,m}$ w.r.t 0X_i and $\mathbf{k}_i^v = -\frac{\partial \mathbf{f}_c^{i,m}}{\partial t} \in \mathbb{R}^{m_i}$. Differentiating the constraint at the velocity-level one more time to get the constraint at the acceleration-level and imposing it within the maximal coordinate GPLC formulation [10] gives

$$\mathbf{minimize}_{\dot{\boldsymbol{\nu}}, \mathbf{a}} \sum_{i=1}^{n_b} \left\{ \frac{1}{2} \mathbf{a}_i^T H_i \mathbf{a}_i - \mathbf{f}_i^T \mathbf{a}_i \right\} \quad (23a)$$

$$\mathbf{subject\ to} \quad \mathbf{a}_i = \mathbf{a}_{\pi(i)} + S_i \dot{\boldsymbol{\nu}}_i + \mathbf{a}_{b,i}, \quad i = 1, 2, \dots, n_b, \quad (23b)$$

$$K_i \mathbf{a}_i = \mathbf{k}_i, \quad i = 1, 2, \dots, n_b, \quad (23c)$$

$$\mathbf{a}_0 = -\mathbf{a}_{\text{grav}}, \quad (23d)$$

where \mathbf{f}_i is the resultant spatial force acting on the i^{th} link including the bias forces $(-\mathbf{v}_i \times {}^*H_i \mathbf{v}_i)$ and $\mathbf{k}_i = -\dot{K}_i \mathbf{v}_i + \dot{\mathbf{k}}_i^v \in \mathbb{R}^{m_i}$. All the spatial quantities are expressed in the inertial frame in our subsequent derivations for simplicity of notation. Since the GPLC is expressed in the maximal coordinates here, unlike the problem in Eq. (3), the joint constraints are also included in Eq. (23b). $\pi(i)$ is the parent

link of the i^{th} link in the kinematic tree, $\boldsymbol{\nu}_i \in \mathbb{R}^{n_i}$ is the i^{th} joint's generalized velocities, $\dot{\boldsymbol{\nu}}_i \in \mathbb{R}^{n_i}$ is the i^{th} joint's generalized accelerations. S_i is the i^{th} joint's motion subspace matrix of size $6 \times n_i$, with n_i being the i^{th} joint's DoF. Each column vector of S_i is an element in \mathbb{M}^6 .

$$\mathbf{a}_{b,i} := \mathbf{v}_i \times S_i \boldsymbol{\nu}_i, \quad (24)$$

is the i^{th} link's bias acceleration. The motion constraints are encoded in Eq. (23c), where K_i , expressed as a matrix in $\mathbb{R}^{m_i \times 6}$ and $\mathbf{k}_i \in \mathbb{R}^{m_i}$, are the constraint matrix and the desired constraint accelerations. Each row vector of K_i is an element in \mathbb{F}^6 . A uniform acceleration field of minus acceleration-due-to-gravity is added to all the links by fixing the root node acceleration in Eq. (23d). This approach [15] spares us from having to add the weight of each link to \mathbf{f}_i thereby providing some computational speed up.

We now connect the constraints in Eq. (23c) to the joint-space formulation constraints in Eq. (3b). Connection in the other direction is not as interesting, since the constraints are more naturally imposed on individual links. Suppose that $J_i \in \mathbb{R}^{6 \times n_i}$ is the geometric Jacobian of the i^{th} link, mapping the generalized velocity $\boldsymbol{\nu}$ to the joint spatial velocity \mathbf{v}_i . Then,

$$J_{f_c}(\mathbf{q}, t) := \begin{bmatrix} K_1 J_1 \\ \vdots \\ K_{n_b} J_{n_b} \end{bmatrix}; \dot{J}_{f_c}(\mathbf{q}, \boldsymbol{\nu}, t) := \begin{bmatrix} \dot{K}_1 J_1 + K_1 \dot{J}_1 \\ \vdots \\ \dot{K}_{n_b} J_{n_b} + K_{n_b} \dot{J}_{n_b} \end{bmatrix} \\ \mathbf{a}_c^*(\mathbf{q}, \boldsymbol{\nu}, t) := [\dot{\mathbf{k}}_1^{vT} \quad \dots \quad \dot{\mathbf{k}}_{n_b}^{vT}]^T \quad (25)$$

E. Accelerating M_{μ} computation.

We now discuss the connection between the link inertias H_i , the JSIM M and the augmented JSIM M_{μ} . From [48, Eq. 8.57], we have

$$M = \begin{bmatrix} J_1 \\ \vdots \\ J_n \end{bmatrix}^T \begin{bmatrix} H_1 & & \\ & \ddots & \\ & & H_n \end{bmatrix} \begin{bmatrix} J_1 \\ \vdots \\ J_n \end{bmatrix}. \quad (26)$$

Considering the definition of M_{μ} and Eq. (25), we get the equation

$$M_{\mu} = \begin{bmatrix} J_1 \\ \vdots \\ J_n \end{bmatrix}^T \begin{bmatrix} H_1 + \mu K_1^T K_1 & & \\ & \ddots & \\ & & H_n + \mu K_n^T K_n \end{bmatrix} \begin{bmatrix} J_1 \\ \vdots \\ J_n \end{bmatrix} \quad (27)$$

The equation above implies that the inertia of every constrained link can be updated as above and the CRBA can be called to compute M_{μ} directly. This provides a minor speed-up compared to first separately computing M using CRBA and later adding the $\mu J_{f_c}^T J_{f_c}$ terms. In the rest of this paper, we will use this strategy to speed up M_{μ} computation in the proxLTLs algorithm.

F. The PV algorithm

We now briefly review the PV algorithm [7], [19], an efficient constrained dynamics algorithm with its origin in the 1970s, see [10] for an expository derivation. It is a

low-complexity counterpart to Featherstone's LTL algorithm, with a $O(n + m^2d + m^3)$ complexity. Similarly to the LTL approach, it first eliminates the primal variables from the Lagrangian of the optimization problem in Eq. (23), but by using DP, to obtain the dual function $g(\lambda)$. The dual function is optimized to compute the optimal dual values, which is then used to compute the resulting accelerations \mathbf{a}_i during a roll-out. It has a three-sweep structure analogous to ABA. However, the PV algorithm assumes the invertibility of Λ^{-1} , and Tikhonov regularization was used in [10] to make the algorithm applicable to singular cases. However, this biases the constraint forces to be closer to the origin, so the motion constraints will not be satisfied accurately.

III. CONSTRAINED ABA DERIVATION

In this section, we will derive the constrained ABA algorithm. Further discussed later, constrained ABA is the low-complexity, recursive analog of the proxLTLs algorithm (see Section II-C2), that is derived from the maximal coordinate GPLC formulation (see Eq. (23)). Therefore, similarly to proxLTLs, we will solve the QP in Eq. (23) using the ALM. This derivation closely follows the DP approach from the PV algorithm derivation in [10], which applied DP on the Lagrangian, resulting in a three-sweep algorithm. In contrast, we will apply DP on the augmented Lagrangian, which gives a similar three-sweep algorithm, and propose two efficient reduced sweeps associated with the subsequent ALM iterations.

All the link acceleration terms \mathbf{a}_i s can be expressed as a function of the joint accelerations $\dot{\nu}_i$ s and the base link acceleration \mathbf{a}_0 , and these constraints will be eliminated through substitution using Eq. (23b) and Eq. (23d). Therefore, we include only the motion constraints Eq. (23c) as the joint constraints in the Lagrangian of Eq. (23) defined as

$$\mathcal{L}_m(\dot{\nu}, \lambda) := \sum_{i=1}^{n_b} \left\{ \frac{1}{2} \mathbf{a}_i^T H_i \mathbf{a}_i - \mathbf{f}_i^T \mathbf{a}_i + \lambda_i^T (K_i \mathbf{a}_i - \mathbf{k}_i) \right\}, \quad (28)$$

and the augmented Lagrangian is

$$\mathcal{L}_m^A(\dot{\nu}, \lambda) := \mathcal{L}_m(\dot{\nu}, \lambda) + \frac{\mu}{2} \sum_{i=1}^{n_b} \|K_i \mathbf{a}_i - \mathbf{k}_i\|^2. \quad (29)$$

Similarly to Section II-C2, applying ALM to Eq. (23) provides the update equations

$$\dot{\nu}^{k+1} = \min_{\dot{\nu}} \mathcal{L}_m^A(\dot{\nu}, \lambda^k), \quad (30a)$$

$$\lambda_i^{k+1} = \lambda_i^k + \mu \{K_i \mathbf{a}_i^{k+1} - \mathbf{k}_i\}, \quad \forall i \in \mathcal{S}, \quad (30b)$$

where λ_i are the dual variables corresponding to the i^{th} link's motion constraints. We first solve Eq. (30a) leveraging the dynamic programming principle.

A. DP-based derivation

Eq. (30a) is an unconstrained QP over $\dot{\nu}$. As a first step to solving it, we collect all the \mathcal{L}_m^A terms (see Eq. (28)

and Eq. (29)) that are quadratic and linear in \mathbf{a}_i to obtain

$$\mathcal{L}_m^A(\dot{\nu}, \lambda^k) = \sum_{i=1}^{n_b} \left\{ \frac{1}{2} \mathbf{a}_i^T (H_i + \mu K_i^T K_i) \mathbf{a}_i - \left(\mathbf{f}_i + K_i^T (\mu \mathbf{k}_i - \lambda_i^k) \right)^T \mathbf{a}_i + \text{const} \right\}, \quad (31)$$

where we remind readers that all the \mathbf{a}_i s are functions of $\dot{\nu}$, and 'const' consists of all the remaining terms in \mathcal{L}_m^A that do not depend on \mathbf{a}_i . Considering the algebraic similarity of equation above with Eq. (23), let us 'update' the H_i and \mathbf{f}_i terms

$$\tilde{H}_i := H_i + \mu K_i^T K_i, \quad \tilde{\mathbf{f}}_i := \mathbf{f}_i + K_i^T (\mu \mathbf{k}_i - \lambda_i^k), \quad (32)$$

to get the unconstrained QP problem

$$\begin{aligned} \text{minimize}_{\dot{\nu}, \mathbf{a}} \quad & \sum_{i=1}^{n_b} \left\{ \frac{1}{2} \mathbf{a}_i^T \tilde{H}_i \mathbf{a}_i - \tilde{\mathbf{f}}_i^T \mathbf{a}_i \right\}, \\ \text{where} \quad & \mathbf{a}_i = \mathbf{a}_{\pi(i)} + S_i \dot{\nu}_i + \mathbf{a}_{b,i}, \quad i = 1, 2, \dots, n_b, \\ \text{and} \quad & \mathbf{a}_0 = -\mathbf{a}_{\text{grav}}. \end{aligned} \quad (33)$$

The problem above is algebraically identical to an unconstrained forward dynamics problem and can be readily solved using the ABA. For completeness, we now derive the ABA algorithm using Vereshchagin's DP approach [11], which is algebraically identical to the Riccati recursion derived for the standard discrete-time linear quadratic regulator (LQR) problem [10], [49]. Using the terminology from LQR literature, let us define the 'cost-to-go' function

$$V_i(\mathbf{a}_i, \dot{\nu}_{\text{desc}(i)}) := \sum_{j=\{i\} \cup \text{Desc}(i)} \left\{ \frac{1}{2} \mathbf{a}_j^T \tilde{H}_j \mathbf{a}_j - \tilde{\mathbf{f}}_j^T \mathbf{a}_j \right\}, \quad (34)$$

as the GPLC objective considering the i^{th} link and all its descendants. Then the optimal cost-to-go function is

$$V_i^*(\mathbf{a}_i) := \min_{\dot{\nu}_{\text{desc}(i)}} V_i(\mathbf{a}_i, \dot{\nu}). \quad (35)$$

With LQR being one of the few problems where DP is computationally tractable, we know that $V_i^*(\mathbf{a}_i)$ is exactly parametrized as quadratic form (see [49]), which we will hypothesize

$$V_i^*(\mathbf{a}_i) = \frac{1}{2} \mathbf{a}_i^T H_i^A \mathbf{a}_i - \mathbf{f}_i^{AT} \mathbf{a}_i. \quad (36)$$

From Bellman's recursion [13], the optimal cost-to-go function of a link and its children are related as

$$V_i^*(\mathbf{a}_i) = \min_{\dot{\nu}_{\gamma(i)}} \left\{ \frac{1}{2} \mathbf{a}_i^T \tilde{H}_i \mathbf{a}_i - \tilde{\mathbf{f}}_i^T \mathbf{a}_i + \sum_{j \in \gamma(i)} V_j^*(\mathbf{a}_j) \right\} \quad (37)$$

$$= \min_{\dot{\nu}_{\gamma(i)}} \left\{ \frac{1}{2} \mathbf{a}_i^T \tilde{H}_i \mathbf{a}_i - \tilde{\mathbf{f}}_i^T \mathbf{a}_i + \sum_{j \in \gamma(i)} V_j^*(\mathbf{a}_i + S_j \dot{\nu}_j + \mathbf{a}_{b,j}) \right\}. \quad (38)$$

Substituting the assumed quadratic form from Eq. (36) for all the $V_j^*(\cdot)$ terms above yields

$$V_i^*(\mathbf{a}_i) = \min_{\dot{\mathbf{v}}_{\gamma(i)}} \left\{ \frac{1}{2} \mathbf{a}_i^T \tilde{H}_i \mathbf{a}_i - \tilde{\mathbf{f}}_i^T \mathbf{a}_i + \sum_{j \in \gamma(i)} \left[(\mathbf{a}_i + S_j \dot{\mathbf{v}}_j + \mathbf{a}_{b,j})^T H_j^A (\mathbf{a}_i + S_j \dot{\mathbf{v}}_j + \mathbf{a}_{b,j}) - \mathbf{f}_j^{AT} (\mathbf{a}_i + S_j \dot{\mathbf{v}}_j + \mathbf{a}_{b,j}) \right] \right\}. \quad (39)$$

Collecting all the terms involving $\dot{\mathbf{v}}_j$ gives a positive definite quadratic function in $\dot{\mathbf{v}}_j$, which is minimized by solving for gradient stationarity condition

$$\dot{\mathbf{v}}_j^* = D_j^{-1} S_j^T \{ \mathbf{f}_j^A - H_j^A (\mathbf{a}_i + \mathbf{a}_{b,j}) \}, \quad (40)$$

where $D_j := S_j^T H_j^A S_j$ and $D_i \in \mathbb{S}_{++}^{n_i}$. Substituting $\dot{\mathbf{v}}_j^*$'s expression from above back in Eq. (39) results in a quadratic expression in \mathbf{a}_i that is identical to the hypothesized quadratic form in Eq. (36), with quadratic form's coefficients given by the recursive update equations

$$H_i^A = \tilde{H}_i + \sum_{j \in \gamma(i)} P_j H_j^A, \quad (41a)$$

$$\mathbf{f}_i^A = \tilde{\mathbf{f}}_i + \sum_{j \in \gamma(i)} P_j \{ \mathbf{f}_j^A - H_j^A \mathbf{a}_{b,j} \}, \quad (41b)$$

where $P_j := I_{6 \times 6} - H_j^A S_j D_j^{-1} S_j^T$ is the projection matrix [14] that propagates inertia and force quantities from a child link to its parent link after subtracting the component that causes the child joint's motion. For all leaf-nodes j in the kinematic tree, we set $H_j^A = \tilde{H}_j$ and $\mathbf{f}_j^A = \tilde{\mathbf{f}}_j$, following which we compute these terms for all the non-leaf links tree recursively using Eq. (41). Thus, we can inductively show that the cost-to-go function hypothesis in Eq. (36) is valid for all links.

After the recursion in Eq. (41) reaches the root node 0 of the tree, the value of $-\mathbf{a}_{\text{grav}}$ is substituted in place of \mathbf{a}_0 , and in a forward-sweep from root to leaves (also called roll-out in DP), the resulting joint accelerations and link accelerations are computed using Eq. (40) and Eq. (23b) respectively.

Unsurprisingly, the recursive equations in Eq. (41) are similar to the ABA [14], as constraint-related terms were absorbed into inertia and force terms in Eq. (32). Despite the algebraic similarity, the H_i^A and \mathbf{f}_i^A terms in Eq. (41) may not be interpreted as the physical articulated body inertia and the resultant spatial forces because of the mathematical modification in Eq. (32) due to the augmented Lagrangian terms of the constraints.

Details on \mathbf{f}_i : As explained in [10]

$$\mathbf{f}_i = T_i \boldsymbol{\tau}_i - \mathbf{v}_i \times^* H_i \mathbf{v}_i + \mathbf{f}_i^{\text{ext}} - \sum_{j \in \gamma(i)} T_j \boldsymbol{\tau}_j,$$

consisting of forces due to joint torques, reaction forces from children joint torques, bias forces and external forces. T_i , expressed as an $n_i \times 6$ matrix, is the i^{th} joint's active force subspace matrix, such that $T_i \boldsymbol{\tau}_i$ is the force exerted on the i^{th} link. Each column of T_i is an element of \mathbb{F}^6 . The active force

subspace matrix and the motion subspace matrices satisfy the constraint [1, Eq. 35]

$$S_i^T T_i = I_{n_i}. \quad (42)$$

We use \mathbf{f}_i in our derivation and algorithm presentation for notational simplicity. However, please note that T_i does not need to be explicitly computed because the algorithm eventually requires only the inner product of S_i and \mathbf{f}_i^A . in Eq. (40), where T_i gets cancelled due to Eq. (42). The resultant force on the i^{th} link due to the torque of its children joints $\boldsymbol{\tau}_j$ is the the sum of reaction force (that is part of \mathbf{f}_i and the backward propagation of $\boldsymbol{\tau}_j$'s contribution to \mathbf{f}_j (using the backpropagation equation in Eq. (41)) is given by

$$\begin{aligned} -T_j \boldsymbol{\tau}_j + P_j T_j \boldsymbol{\tau}_j &= -T_j \boldsymbol{\tau}_j + (I_6 - H_j^A S_j (D_j)^{-1} S_j^T) T_j \boldsymbol{\tau}_j, \\ &= -H_j^A S_j D_j^{-1} \boldsymbol{\tau}_j, \end{aligned}$$

thereby not requiring T_j computation as well.

IV. CONSTRAINED ABA ALGORITHM

We now present constrainedABA derived in the previous section in an algorithmic form. Introducing some notation, let $ES(i)$ be the set of indices of the end-effector supported by the i^{th} link, which implies that these end-effectors are descendants of the i^{th} link that are in \mathcal{E} . Let $\mathcal{S}_{\mathcal{E}}$ be the sublist of \mathcal{S} , consisting of only those links $i \in \mathcal{S}$ that support constraints, meaning that the cardinality of $ES(i)$ is greater than 0, denoted by $\mathcal{C}(ES(i)) > 0$. $\mathcal{S}_{\mathcal{E}^r}$ is a reversed version of $\mathcal{S}_{\mathcal{E}}$.

We first present the three-sweep algorithm, which is computed only during the first ALM iteration. For subsequent ALM iterations, we present a fast reduced two-sweep algorithm, that re-uses many of the terms computed in the three-sweep algorithm. Finally, we present the constraintABA algorithm itself and perform a complexity analysis. Please note that there are some additional computation optimizations to avoid unnecessary computations. Performing a second forward sweep to compute accelerations for links not supporting constraints ($\mathcal{S} - \mathcal{S}_{\mathcal{E}}$) is not useful during ALM iterations since they do not have any constraints that would require the ALM's Lagrange multiplier updates. For the same reason, no backward sweep is required for these links from the second ALM iteration onwards, as there are no updated constraint forces that need to be transmitted back. Therefore, to avoid unnecessary computation, the reduced sweeps are not performed for these links until the ALM iterations have converged.

A. Three-sweep constrainedABA algorithm

Algorithm 1 lists the three sweep part of constrainedABA. It is identical to the original articulated-body algorithm [1], [14] except for the line 5, where the inertia matrices and the spatial force terms are updated to account for the influence of the constraint. It requires modifying existing ABA implementations by adding only a few lines of code to obtain the presented algorithm, thereby facilitating its implementation in existing simulators and rigid-body dynamics libraries such as RBDL [50], Pinocchio [6] or Drake [51].

Algorithm 1 ConstrainedABA three-sweep algorithm

Require: \mathbf{q} , $\boldsymbol{\nu}$, $\boldsymbol{\tau}$, S_i s, K_i s, \mathbf{k}_i s, μ , $\boldsymbol{\lambda}^k$, robot model

First forward sweep

- 1: **for** i in \mathcal{S} **do**
- 2: $\mathbf{v}_i = \mathbf{v}_{\pi(i)} + S_i \boldsymbol{\nu}_i$
- 3: $\mathbf{a}_{b,i} = \mathbf{v}_i \times S_i \boldsymbol{\nu}_i$
- 4: $\mathbf{f}_i^A \leftarrow \mathbf{f}_i + T_i \boldsymbol{\tau}_i - \mathbf{v}_i \times * H_i \mathbf{v}_i + \mathbf{f}_i^{\text{ext}}; \mathbf{f}_{\pi(i)}^A \leftarrow \mathbf{f}_{\pi(i)}^A - T_i \boldsymbol{\tau}_i$
- 5: $H_i^A \leftarrow H_i + \mu K_i^T K_i; \mathbf{f}_i^A \leftarrow \mathbf{f}_i^A + K_i^T (\mu \mathbf{k}_i - \boldsymbol{\lambda}_i^k)$

Backward sweep

- 6: **for** i in \mathcal{S}_r **do**
- 7: $D_i = S_i^T H_i^A S_i; P_i = (I_{6 \times 6} - H_i^A S_i (D_i)^{-1} S_i^T)$
- 8: $\mathbf{f}_{\pi(i)}^A \leftarrow \mathbf{f}_{\pi(i)}^A + P_i (\mathbf{f}_i^A - H_i^A \mathbf{a}_{b,i})$
- 9: $H_{\pi(i)}^A \leftarrow H_{\pi(i)}^A + P_i H_i^A$

Second forward sweep (roll-out)

- 10: **for** i in $\mathcal{S}_{\mathcal{E}}$ **do**
 - 11: $\dot{\boldsymbol{\nu}}_i^{(k+1)*} = (D_i)^{-1} S_i^T \left\{ \mathbf{f}_i^A - H_i^A (\mathbf{a}_{\pi(i)}^{k+1} + \mathbf{a}_{b,i}) \right\}$
 - 12: $\mathbf{a}_i^{k+1} = \mathbf{a}_{\pi(i)}^{k+1} + S_i \dot{\boldsymbol{\nu}}_i^{(k+1)*} + \mathbf{a}_{b,i}$
-

B. Reduced two-sweep algorithm

For the subsequent ALM iterations, we now derive a fast reduced two-sweep algorithm that computes fast sweeps by re-using as many terms computed in Algorithm 1 as possible. The main observation is that only the force update in Eq. (32) changes between the ALM iterations, and we compute only the modification to the force and acceleration terms due to this force update in the two-sweep algorithm. This force term also changes only for those links that support constraints $\mathcal{S}_{\mathcal{E}}$.

Even in Eq. (32), only the $\boldsymbol{\lambda}_i$ term changes between ALM iterations. These dual variables are updated using Eq. (30b). Substituting Eq. (30b) in Eq. (32), we see that the spatial force terms change by $\mu K_i^T (\mathbf{k}_i - K_i \mathbf{a}_i^{k+1})$. We introduce an additional term $\Delta \mathbf{f}_i$ to keep track of these force updates. The reduced sweep algorithm is presented in Algorithm 2.

Despite the reduced sweep algorithm appearing to be distinct from the three-sweep traditional ABA algorithm, we note that it is strictly a subset of the three-sweep ABA algorithm. This facilitates implementing the reduced sweeps since much the source code of an existing ABA implementation can be conveniently reused.

Algorithm 2 ConstrainedABA reduced two-sweep algorithm

Require: $\Delta \mathbf{f}_i$ s, μ , K_i s, \mathbf{k}_i s, \mathbf{a}_i s, P_i s, D_i s, S_i s, \mathbf{f}_i^A s, H_i^A s, robot model

Backward sweep

- 1: **for** i in $\mathcal{S}_{\mathcal{E}_r}$ **do**
- 2: $\Delta \mathbf{f}_i \leftarrow \Delta \mathbf{f}_i + \mu K_i^T (\mathbf{k}_i - K_i \mathbf{a}_i^k)$
- 3: $\Delta \mathbf{f}_{\pi(i)} \leftarrow \Delta \mathbf{f}_{\pi(i)} + P_i \Delta \mathbf{f}_i$

Second forward sweep (roll-out)

- 4: **for** i in $\mathcal{S}_{\mathcal{E}}$ **do**
 - 5: $\dot{\boldsymbol{\nu}}_i^{(k+1)*} = D_i^{-1} S_i^T \left\{ \mathbf{f}_i^A + \Delta \mathbf{f}_i - H_i^A (\mathbf{a}_{\pi(i)}^{k+1} + \mathbf{a}_{b,i}) \right\}$
 - 6: $\mathbf{a}_i^{k+1} = \mathbf{a}_{\pi(i)}^{k+1} + S_i \dot{\boldsymbol{\nu}}_i^{(k+1)*} + \mathbf{a}_{b,i}$
 - 7: $\Delta \mathbf{f}_i = \mathbf{0}_6$
-

C. ConstrainedABA algorithm

Using the three-sweep and reduced two-sweep algorithms, we present the entire constrainedABA in Algorithm 3. The algorithm takes as input the numerical tolerance $\epsilon \in \mathbb{R}$, $\epsilon > 0$, and the maximum number of iterations to determine the convergence of the ALM method.

Algorithm 3 ConstrainedABA

Require: \mathbf{q} , $\boldsymbol{\nu}$, $\boldsymbol{\tau}$, μ , max_iter, K_i s, \mathbf{k}_i s, ϵ , $\dot{\boldsymbol{\nu}}^0$, $\boldsymbol{\lambda}^0$ robot model

- 1: Execute the three-sweep algorithm in Algorithm 1.
 - 2: **for** i in \mathcal{S} **do**
 - 3: $\Delta \mathbf{f}_i \leftarrow \mathbf{0}_6$
 - 4: **for** k in range(1, max_iter) **do**
 - 5: **for** $i \in \mathcal{E}$ **do**
 - 6: $\boldsymbol{\lambda}^k = \boldsymbol{\lambda}^{k-1} + \mu \{ K_i \mathbf{a}_i^k - \mathbf{k}_i \}$
 - 7: **if** $\min(\|\dot{\boldsymbol{\nu}}^k - \dot{\boldsymbol{\nu}}^{k-1}\|_{\infty}, \|K_i \mathbf{a}_i^k - \mathbf{k}_i\|_{\infty}) < \epsilon$ **then**
 - 8: break
 - 9: Execute reduced-two sweep algorithm in Algorithm 2
 - 10: **for** i in $\mathcal{S} - \mathcal{S}_{\mathcal{E}}$ **do**
 - 11: $\dot{\boldsymbol{\nu}}_i^{(k+1)*} = (D_i)^{-1} S_i^T \left\{ \mathbf{f}_i^A - H_i^A (\mathbf{a}_{\pi(i)}^{k+1} + \mathbf{a}_{b,i}) \right\}$
 - 12: $\mathbf{a}_i^{k+1} = \mathbf{a}_{\pi(i)}^{k+1} + S_i \dot{\boldsymbol{\nu}}_i^{(k+1)*} + \mathbf{a}_{b,i}$
-

D. Complexity analysis

We now analyze the computational complexity of Algorithm 1, where each line has a fixed computational cost except for the line 5, which is proportional to number of rows in K_i , corresponding to the number of constraints acting on the i^{th} link m_i . The line 5 is executed at all constrained links. Therefore, this line requires $O(m)$ number of operations in total. All the other lines are executed at most n number of times (once per joint) and therefore require $O(n)$ number of operations. Combining these costs, we get $O(m+n)$ as the computational complexity for the three-sweep algorithm.

Very similar analysis can be made for the reduced sweep Algorithm 2, where the line 2 requires $O(m)$ operations while the rest of the lines have a fixed cost and require $O(n)$ operations, bringing the total computational complexity of the two-sweep algorithm to $O(m+n)$ as well.

Within constrainedABA (see Algorithm 3), the three-sweep algorithm is computed once, and the two-sweep algorithm is computed at most max_iter–1 number of times, which is a fixed number supplied by the user. Empirically, we find that constrainedABA requires only a few iterations (nearly always under 5), as we shall see in the benchmarking (see Sec. VIII) on a diverse set of mechanisms. In practice, this brings the total computational complexity of constrainedABA to $O(m+n)$.

V. PROXIMAL PV ALGORITHM

As mentioned in Section II-F, the PV algorithm in literature [10] requires linear independence of the motion constraints and in practice employs Tikhonov regularization to handle the singular cases. However, this prevents the motion constraints

from being satisfied even when feasible. This drawback of the PV algorithm can be straightforwardly resolved by employing the PPA (see Section II-C1) to obtain a proximal algorithm version of the PV algorithm, that we shall call proxPV. We only discuss the modifications to the original PV algorithm that will provide us with the proxPV algorithm. For a detailed derivation of the PV algorithm, the readers are referred to [10].

The PV algorithm computes the dual function $\mathbf{g}(\boldsymbol{\lambda})$ (both Λ^{-1} and the affine terms) at the end of its backward sweep. Similarly to proxLTL, the weighted regularization $\frac{1}{\mu}I$ is added to Λ^{-1} to get Λ_μ^{-1} , which is always invertible and can be efficiently factorized using the Cholesky factorization. Then the dual function is maximized using the PPA until convergence and the resulting robot accelerations are computed in the third sweep. Effectively, only three lines are changed in the PV algorithm to get the proxPV algorithm.

Though we do not include a full derivation, Algorithm 4 presents the entire proxPV algorithm to make this paper self-contained. Every term in Algorithm 4 also has a physical interpretation, and we refer readers to [10] for a detailed discussion. Algorithm 4 requires additional notations defined as follows: for $\forall j \in \gamma(i)$, let

$$\boldsymbol{\lambda}_i^A = [\boldsymbol{\lambda}_i, \dots, \boldsymbol{\lambda}_j^{AT}, \dots]^T, K_i^A = \begin{bmatrix} K_i \\ \vdots \\ K_j^A \\ \vdots \end{bmatrix},$$

and:

$$L_i^A = \begin{bmatrix} L_i & & & \\ & \ddots & & \\ & & L_j^A & \\ & & & \ddots \end{bmatrix}.$$

These three quantities collect all the terms related to motion constraints on the subtree rooted at the i^{th} link to make the recursive formulae in the backward sweep work.

Complexity analysis: ProxPV's complexity analysis is very similar to that performed in [10]. The primary bottlenecks of the Algorithm are line 10, which requires $O(m_i^2)$ operations at each link and the line 13 where the dual Hessian is once factorized requiring $O(m^3)$ operations, further proximal iterations re-use this factorization and only require $O(m^2)$ operations. Similarly to the previous section, considering that practical examples require only a few proximal iterations, the practical computational complexity of the proxPV algorithm is $O(n + m^2d + m^3)$.

VI. EFFICIENT ALGORITHMS TO COMPUTE Λ_μ

In this section, we introduce a new algorithm, named cABA-OSIM, to compute the damped Delassus matrix inverse Λ_μ with the lowest computational complexity of $O(n + m^2)$ compared to the next best algorithm of $O(n + m^3)$ complexity from literature [30]. Λ_μ^{-1} is the Hessian of the dual function of the proximal formulation of constrained dynamics (see Eq. (10)),

Algorithm 4 Proximal PV solver

Require: \mathbf{q} , $\boldsymbol{\nu}$, $\boldsymbol{\tau}$, K_i s, \mathbf{k}_i s, ϵ , max_iter, λ_0^0 , robot model

First forward sweep

- 1: **for** i in \mathcal{S} **do**
- 2: $\mathbf{v}_i = \mathbf{v}_{\pi(i)} + S_i \boldsymbol{\nu}_i$
- 3: $\mathbf{a}_{b,i} = \mathbf{v}_i \times S_i \boldsymbol{\nu}_i$
- 4: $\mathbf{f}_i^A \leftarrow \mathbf{f}_i^A + T_i \boldsymbol{\tau}_i - \mathbf{v}_i \times^* H_i \mathbf{v}_i + \mathbf{f}_i^{\text{ext}}$; $K_i^A \leftarrow K_i$;
 $\mathbf{l}_i \leftarrow -\mathbf{k}_i$; $L_i^A \leftarrow \mathbf{0}_{m_i \times m_i}$ $H_i^A \leftarrow H_i$;
 $\mathbf{f}_{\pi(i)}^A \leftarrow \mathbf{f}_{\pi(i)}^A - T_i \boldsymbol{\tau}_i$

Backward sweep

- 5: **for** i in \mathcal{S}_r **do**
- 6: $D_i = S_i^T H_i^A S_i$; $P_i = (I_{6 \times 6} - H_i^A S_i (D_i)^{-1} S_i^T)$
- 7: $\mathbf{f}_{\pi(i)}^A \leftarrow \mathbf{f}_{\pi(i)}^A + P_i (\mathbf{f}_i^A - H_i^A \mathbf{a}_{b,i})$
- 8: $H_{\pi(i)}^A \leftarrow H_{\pi(i)}^A + P_i H_i^A$
- 9: $K_{\pi(i)}^A \leftarrow \begin{bmatrix} K_{\pi(i)}^A \\ K_i^A P_i^T \end{bmatrix}$
- 10: $L_{\pi(i)}^A \leftarrow \begin{bmatrix} L_{\pi(i)}^A & \\ & L_i^A + K_i^A S_i (D_i)^{-1} S_i^T K_i^{AT} \end{bmatrix}$
- 11: $\mathbf{l}_{\pi(i)} \leftarrow \begin{bmatrix} \mathbf{l}_{\pi(i)} \\ \mathbf{l}_i + K_i^A \{ \mathbf{a}_{b,i} + S_i D_i^{-1} S_i^T (\mathbf{f}_i^A - H_i^A \mathbf{a}_{b,i}) \} \end{bmatrix}$
- 12: **for** k in range(1, max_iter) **do**
- 13: $\boldsymbol{\lambda}_0^k = (L_0^A + \frac{1}{\mu} I)^{-1} (\frac{\boldsymbol{\lambda}_0^{k-1}}{\mu} + K_0 \mathbf{a}_0 + \mathbf{l}_0)$
- 14: **if** $\|L_0^{AT} (L_0^A \boldsymbol{\lambda}_0^k - K_0 \mathbf{a}_0 - \mathbf{l}_0)\|_\infty < \epsilon$ **then**
- 15: **break**
- 16: $\boldsymbol{\lambda}_0^{A*} = \boldsymbol{\lambda}_0^k$

Second forward sweep (roll-out)

- 17: **for** i in \mathcal{S} **do**
 - 18: $\boldsymbol{\nu}_i^* = D_i^{-1} S_i^T \{ \mathbf{f}_i^A - H_i^A (\mathbf{a}_{\pi(i)} + \mathbf{a}_{b,i}) - K_i^{AT} \boldsymbol{\lambda}_i^{A*} \}$
 - 19: $\mathbf{a}_i = \mathbf{a}_{\pi(i)} + S_i \boldsymbol{\nu}_i^* + \mathbf{a}_{b,i}$
-

and is an important quantity used in constrained dynamics simulation [10], [22] and operational space control [26].

The cABA-OSIM algorithm exploits a linear algebra identity known as the matrix inversion lemma [31] (also called Morrison-Sherman-Woodbury formula), which states that

$$(A + UCV)^{-1} = A^{-1} - A^{-1}U(C^{-1} + VA^{-1}U)^{-1}VA^{-1}, \quad (43)$$

where $A \in \mathbb{R}^{n \times n}$, $U \in \mathbb{R}^{n \times m}$, $C \in \mathbb{R}^{m \times m}$, $V \in \mathbb{R}^{m \times n}$, and A and C matrices are assumed to be invertible. The identity is useful for computing the inverse of the left-hand-side in Eq. (43), when it is easy to compute the inverses of A and C .

We recall that

$$\Lambda_\mu^{-1} := \frac{1}{\mu} I_{m \times m} + J_{\mathbf{f}_c} M^{-1} J_{\mathbf{f}_c}^T,$$

to which applying the matrix inversion lemma gives,

$$\Lambda_\mu = \mu I_{m \times m} - \mu J_{\mathbf{f}_c} (M_\mu)^{-1} J_{\mathbf{f}_c}^T, \quad (44)$$

where we recall that

$$M_\mu := M + \mu J_{\mathbf{f}_c}^T J_{\mathbf{f}_c}.$$

The second term in the right-hand side of Eq. (44) is algebraically identical to the Delassus matrix expression, which is of the form $\Lambda^{-1} := J_{\mathbf{f}_c} (M)^{-1} J_{\mathbf{f}_c}^T$. Therefore, we can augment

the inertia matrix and leverage existing efficient algorithms for computing the Delassus matrix [21], [28], [30] followed by scalar multiplications and subtraction from a diagonal matrix to compute the damped Delassus matrix inverse. We recall from Section II-E that the augmented inertia matrix in the maximal coordinates takes the form of updating all the constrained links' inertias $H_i \leftarrow H_i + \mu K_i^T K_i$. This implies that one of the two efficient maximal coordinates-based Delassus matrix algorithms, namely EFPA [28] or PV-OSIMr [30] can be used within cABA-OSIM.

Between the two options, we select the PV-OSIMr algorithm due to its lower computational complexity of $O(n + m^2)$ compared to $O(m + md + m^2)$ of EFPA. Since the inertia augmentation and the evaluation of Eq. (44) require at most $O(m^2)$ number of additional operations, the total complexity of cABA-OSIM is $O(n + m^2)$. This contrasts with all the existing (that we are aware of) Λ_μ computation algorithms that require an additional $O(m^3)$ number of operations to factorize the Delassus matrix. We present the cABA-OSIM algorithm in the coming subsection, where we also include the computations within PV-OSIMr for completeness.

A. cABA Delassus inverse algorithm

Algorithm 5 presents the cABA-OSIM algorithm. We now review the terms from [30] needed to parse this algorithm and follow that with a brief explanation. Let $\text{anc}(i)$ and $\text{desc}(i)$ be the set of all link indices that are ancestors and descendants of the i^{th} link in the tree respectively. Let $\text{ES}(i) = \mathcal{E} \cap (\text{desc}(i) \cup \{i\})$ be the set of end-effectors supported by the i^{th} link. Let \mathcal{N} be the set of all 'branching' links that support more strictly more constraints than their children links, $\mathcal{N} = \{0\} \cup \mathcal{E} \cup \{i \in \mathcal{S} | \text{ES}(j) \subset \text{ES}(i), \forall j \in \gamma(i)\}$.

Let ${}^j P_i$ denote the extended force propagator (EFP), first discussed in [52], which directly propagates spatial forces acting on the i^{th} link backward in the kinematic tree to an ancestor link j . ${}^i P_i = I_{6 \times 6}$. EFP's transpose ${}^j P_i^T$ directly propagates the spatial accelerations of the j^{th} link to the descendant link i (see [10], [52] for a derivation). Let ${}^j \Omega_i$ be the spatial inverse inertia matrix, which maps spatial forces acting on the i^{th} link to the i^{th} link's spatial acceleration considering the motion of all the joints connecting the j^{th} and the i^{th} links. Let $\text{cca}(i, j) = \mathbf{max}\{(\text{anc}(i) \cup \{i\}) \cup (\text{anc}(j) \cup \{j\})\}$ be the closest common ancestor link of the i^{th} and the j^{th} links.

The only additions compared to the PV-OSIMr algorithm [30] are in line 4, where the inertia and constraint forces are updated and in line 23, where the Delassus matrix is computed using the matrix inversion lemma. Algorithm 5 presentation is simplified relative to PV-OSIMr in [30] by absorbing multiple "end-effectors" on the same link i within a single constraint matrix K_i and by omitting a minor optimization, whereby the EFP of a constrained link, that supports no other end-effector, is initialized with K_i instead of $I_{6 \times 6}$. For more details, readers are referred to [30].

Line 11 recursively computes the extended force propagator (EFP) from a link $i \in \mathcal{N}$ to its closest ancestor in \mathcal{N} , denoted as $\mathcal{N}_{\text{anc}(i)} = \mathbf{max}\{\mathcal{N} \cap \text{anc}(i)\}$. Similarly, line 12 recursively computes the spatial inverse inertia apparent at a link $i \in \mathcal{N}$

considering the joints connecting the i^{th} link and $\mathcal{N}_{\text{anc}(i)}^{\text{th}}$ link. Then, in a limited outward sweep over $i \in \mathcal{N}$, the spatial inverse inertia considering the motion of all ancestor joints ${}^0 \Omega_i$ is computed in line 15. Finally, the constraint space EFP ${}^j K_i^T$, which maps λ_i to the spatial forces acting on the j^{th} link is computed for all end-effector links $i \in \mathcal{E}$ and for all links $j \in \mathcal{N} \cap \text{anc}(i)$ in line 18. Note that ${}^i K_i = K_i$, and ${}^{\pi(i)} P_i = I_6, \forall i \in \mathcal{E}$.

This gives all the quantities required for assembling the Delassus matrix for M_μ , given by

$$L_0^A = \begin{bmatrix} {}^0 L_{1,1}^A & \cdots & {}^0 L_{1,m_b}^A \\ \vdots & \ddots & \vdots \\ {}^0 L_{m_b,1}^A & \cdots & {}^0 L_{m_b,m_b}^A \end{bmatrix}$$

whose blocks are computed in line 22, for $i, j \in \mathcal{E}$, where $m_b = \mathcal{C}(\mathcal{E})$. That L_0^A is the Delassus matrix, was shown in [10]. Each block ${}^0 L_{i,j}^A$ encodes the coupling between the i^{th} and j^{th} links. The right-hand side term in line 22 can be interpreted as follows, it first maps λ_j to spatial forces on the $\text{cca}(i, j)^{\text{th}}$ link, then computes the spatial acceleration of the $\text{cca}(i, j)^{\text{th}}$ link resulting from these forces which actually causes the the constraints on the i^{th} and j^{th} link to get coupled, before finally transmitting this acceleration to the constraint accelerations of the i^{th} link. Finally, the damped Delassus inverse matrix, the output of the cABA-OSIM algorithm, is computed in line 23 using the matrix inversion lemma.

A limitation of this approach is that this optimal complexity algorithm can only compute the damped Delassus inverse and not the undamped Delassus inverse matrix. However, this is not a major limitation because undamped Delassus matrix might not even exist for redundant constraint formulations. Even when the undamped inverse exists, it is a common practice [53] to use the damped inverse due to the risk of kinematic singularities.

B. Accelerating cABA for floating-base robots (cABA-OSIMf)

We now propose a minor structure-exploiting acceleration for cABA-OSIM for certain robot structures, namely floating-base robots with branching at the base. This structure is found commonly, e.g., in quadrupeds and humanoids. For these mechanisms, the Delassus matrix (as well as its inverse) is dense, without any structural zeros because all the constraints acting on different branches get coupled due to the floating-base's motion. However, L_1^A , where 1 is the floating base's index, would have a block-diagonal structure, with each block corresponding to constraints from a branch. It is computationally more efficient to exploit this block-diagonal sparsity and not explicitly compute the dense L_0^A matrix. This idea is implemented by the cABA-OSIM-fast (cABA-OSIMf) that we introduce in the rest of this subsection.

Using the notation from the PV algorithm, L_1^A is related to L_0^A by the equation

$$L_0^A = L_1^A + K_1^A (H_1^A)^{-1} K_1^{AT}, \quad (45)$$

where accounting for the coupling introduced between the constraints on different branches due to the spatial inverse inertia of the floating base link destroys block-diagonal sparsity

Algorithm 5 The cABA-OSIM algorithm

Require: \mathbf{q} , K_i s, robot model, μ
First forward sweep

- 1: **for** i in \mathcal{S} **do**
- 2: $H_i^A \leftarrow H_i$;
- 3: **if** $i \in \mathcal{E}$ **then**
- 4: $H_i^A \leftarrow H_i^A + \mu K_i^T K_i$; ${}^i K_i \leftarrow \mu K_i$

Backward sweep

- 5: **for** i in \mathcal{S}_r **do**
 - 6: **if** $i \in \mathcal{N}$ **then**
 - 7: $d(i) \leftarrow i$; ${}^i P_i = I_{6 \times 6}$; ${}^i \Omega_i \leftarrow 0_{6 \times 6}$
 - 8: $D_i = S_i^T H_i^A S_i$; $P_i = (I_{6 \times 6} - H_i^A S_i (D_i)^{-1} S_i^T)$
 - 9: $H_{\pi(i)}^A \leftarrow H_{\pi(i)}^A + P_i H_i^A$;
 - 10: **if** $ES(i) \neq \emptyset$ **then**
 - 11: $\pi^{(i)} P_{d(i)} = P_i {}^i P_{d(i)}$
 - 12: $\pi^{(i)} \Omega_{d(i)} = {}^i \Omega_{d(i)} + {}^i P_{d(i)}^T S_i (D_i^{-1}) S_i^T {}^i P_{d(i)}$
 - 13: $d(\pi(i)) \leftarrow d(i)$
 - 14: **for** i in $\mathcal{N} - \{0\}$ **if** $\mathcal{N}_{\text{anc}(i)} \neq 0$ **do**
 - 15: ${}^0 \Omega_i = \mathcal{N}_{\text{anc}(i)} \Omega_i + \mathcal{N}_{\text{anc}(i)} P_i^T {}^0 \Omega_{\mathcal{N}_{\text{anc}(i)}} \mathcal{N}_{\text{anc}(i)} P_i$
 - 16: **for** $e_i \in \mathcal{E}$ **do**
 - 17: **for** i in $(\mathcal{N}_r - \{0\}) \cap (\{\text{anc}(e_i) \cup \{e_i\}\})$ **if** $\mathcal{N}_{\text{anc}(i)} \neq 0$ **do**
 - 18: $\mathcal{N}_{\text{anc}(i)} K_{e_i} = {}^i K_{e_i} \mathcal{N}_{\text{anc}(i)} P_i^T$
 - 19: **for** $i, j \in \mathcal{E}$ **if** $i \leq j$ **do**
 - 20: ${}^0 L_{i,j}^A = \text{cca}(i,j) K_i ({}^0 \Omega_{\text{cca}(i,j)}) \text{cca}(i,j) K_j^T$
 - 21: **if** $i \neq j$ **then**
 - 22: ${}^0 L_{j,i}^A = {}^0 L_{j,i}^{AT}$
 - 23: $\Lambda_\mu = \mu I_{m \times m} - L_0^A$
-

pattern of L_1^A . The key idea behind cABA-OSIMf is to not explicitly evaluate the right-hand side of the equation above but to directly multiply this right-hand side with vectors.

Modifications to Algorithm 5: The cABA-OSIM shown in Algorithm 5 is easily modified to compute L_1^A instead of L_0^A . Physically, L_1^A is the Delassus matrix assuming that link 1 (the floating-base) is fixed. Therefore, we need to modify Algorithm 5 to not add the inverse inertia due to the 1st joint. To do this, the backward sweep is run over only $\mathcal{S}_r - \{1\}$ joints. Similarly $\mathcal{N} \leftarrow \mathcal{N} - \{1\}$ and $\mathcal{N}_r \leftarrow \mathcal{N}_r - \{1\}$. The lines 15, 20, 22 and 23 are modified to

$${}^1 \Omega_i = \mathcal{N}_{\text{anc}(i)} \Omega_i + \mathcal{N}_{\text{anc}(i)} P_i^T {}^1 \Omega_{\mathcal{N}_{\text{anc}(i)}} \mathcal{N}_{\text{anc}(i)} P_i, \quad (46a)$$

$${}^1 L_{i,j}^A = \text{cca}(i,j) K_i^A ({}^1 \Omega_{\text{cca}(i,j)}) \text{cca}(i,j) K_j^{AT}, \quad (46b)$$

$${}^1 L_{j,i}^A = {}^1 L_{j,i}^{AT}, \quad (46c)$$

$$\Lambda_\mu = \mu I_{m \times m} - L_1^A - K_1^A (H_1^A)^{-1} K_1^{AT}. \quad (46d)$$

However, the right-hand-side in Eq. (46d) is not to be explicitly evaluated. Matrix-vector products with the damped Delassus matrix Λ_μ are evaluated by directly multiplying the right-hand-side with the vector.

The best-case scenario for cABA-OSIMf is when constraints on the system are evenly distributed among the branches originating from the floating-base. Let b be the number of branches at the base. Then the constraints per branch is $\frac{m}{b}$. The cost of computing the block diagonal term in L_1^A using PV-OSIMr for each branch is $O(\frac{m^2}{b^2})$ plus an additional

cost term proportional to the number of links in the branch. Accumulating this cost for all the branches requires $O(\frac{m^2}{b} + n)$ operations for computing L_1^A . Finally, evaluating the right-hand side of Eq. (45) has a cost of $O(m)$ bringing the total computational complexity of cABA-OSIMf in this best-case to $O(\frac{m^2}{b} + n + m)$.

VII. CONNECTIONS TO MUJOCo AND THE PV ALGORITHMS

In this section, we discuss the connections between the algorithms presented in this paper and the ones implemented in MUJOCo [54] corresponding to PV-soft and the PV-early algorithms developed in [10].

A. Connection to MUJOCo and PV-soft

MUJOCo [5], [54], a popular simulator in robotics and RL communities, with a wide user base, solves a soft version of the Gauss' principle. In MUJOCo, the hard motion constraints are relaxed with a quadratic penalty. For a simulation problem with equality constraints, the optimization problem solved by MUJOCo is of the form

$$\underset{\dot{\mathbf{v}}}{\text{minimize}} \quad \frac{1}{2} \|\dot{\mathbf{v}} - \dot{\mathbf{v}}_{\text{free}}\|_{M(\mathbf{q})}^2 + \|\mathbf{J}_{\mathbf{f}_c}(\mathbf{q})\dot{\mathbf{v}} + \gamma_{\mathbf{f}_c} - \mathbf{a}_c^*\|_{R_E^{-1}}^2, \quad (47)$$

where $R_E \in \mathcal{S}_{++}^m$ is a diagonal matrix chosen by the user, and which can be interpreted as a compliance term.

Comparing the optimization problem in Eq. (47) with the proxLTLs formulation in Eq. (16), we see that for $R_E^{-1} = \mu I_{m \times m}$ and $\lambda^k = \mathbf{0}_m$, one iteration of proxLTLs is identical to the soft-Gauss problem solved in MUJOCo. Please note from Remark 2 that proxLTLs is trivially extendable to support a positive definite diagonal matrix R_E^{-1} instead of the weighted identity matrix $\mu I_{m \times m}$ we have chosen as the compliance term. Therefore, MUJOCo's algorithm for kinematic trees can be considered to be a special case of proxLTLs. A practical consequence of this insight is that codebases using soft Gauss principle such as MUJOCo can be extended readily to obtain the proxLTLs algorithm that converges to a solution that satisfies motion constraints. Computationally, even our first iteration of the proxLTLs differs from MUJOCo's implementation by using the approach mentioned in Section II-E to speed up M_μ computation.

Analogously in the maximal coordinates, the problem solved by the PV-soft algorithm is minimizing

$$\underset{\dot{\mathbf{v}}}{\text{minimize}} \quad \sum_{i=1}^{n_b} \left\{ \frac{1}{2} \mathbf{a}_i^T H_i \mathbf{a}_i - \mathbf{f}_i^T \mathbf{a}_i + \frac{\mu}{2} \|K_i \mathbf{a}_i - \mathbf{k}_i\|^2 \right\},$$

which can be seen to be a special case of constraintABA's objective function (see Eq. (29)), when $\lambda_k = \mathbf{0}_{m_k}$. So, PV-soft can be interpreted as the first ALM iteration of the constraintABA.

B. The proximal PV-early solver

ConstrainedABA is closely related to the $O(n + m)$ complexity PV-early algorithm proposed in [10]. ConstrainedABA can be derived by applying the PV-early algorithm to solve the proximal function in Eq. (14), but this time in maximal coordinates. This involves adding proximal ℓ_2 regularization to the maximal coordinate Lagrangian in Eq. (28) to get

$$\text{prox}_{\mu, \mathcal{L}_m}(\dot{\mathbf{v}}, \boldsymbol{\lambda}^k) := \arg \min_{\boldsymbol{\lambda}} \max_{\dot{\mathbf{v}}} \sum_{i=1}^{n_b} \left\{ \frac{1}{2} \mathbf{a}_i^T H_i \mathbf{a}_i - \mathbf{f}_i^T \mathbf{a}_i + \boldsymbol{\lambda}_i^T (K_i \mathbf{a}_i - \mathbf{k}_i) - \frac{1}{2\mu} \|\boldsymbol{\lambda}_i - \boldsymbol{\lambda}_i^k\|^2 \right\}, \quad (48)$$

which makes the dual problem strongly convex. In the PV-early algorithm, the dual variables are aggressively eliminated as early as possible. Applying this idea to the problem above in Eq. (48), all the dual variables can be immediately eliminated since the problem is now strongly convex in the dual variables. Imposing gradient stationarity KKT condition w.r.t the dual variables of (48) gives

$$-\frac{1}{\mu} \boldsymbol{\lambda}_i^* + \frac{1}{\mu} \boldsymbol{\lambda}_i^k + (K_i \mathbf{a}_i - \mathbf{k}_i) = 0. \quad (49)$$

Re-arranging the terms gives the expression for optimal $\boldsymbol{\lambda}_i$

$$\boldsymbol{\lambda}_i^* = \boldsymbol{\lambda}_i^k + \mu(K_i \mathbf{a}_i - \mathbf{k}_i). \quad (50)$$

Finally, substituting them back into the problem in Eq. (48) gives the purely primal objective function

$$\min_{\dot{\mathbf{v}}} \sum_{i=1}^{n_b} \left\{ \frac{1}{2} \mathbf{a}_i^T H_i \mathbf{a}_i - \mathbf{f}_i^T \mathbf{a}_i + \frac{\mu}{2} \|K_i \mathbf{a}_i - \mathbf{k}_i\|^2 + \boldsymbol{\lambda}_i^{kT} (K_i \mathbf{a}_i - \mathbf{k}_i) \right\}.$$

The objective function above is identical to the maximal coordinate ALM formulation (see Eq. (29)) used to derive the constrainedABA algorithm. Further computations by the PV-early solver would mirror the three-sweep algorithm of constrainedABA. Therefore, evaluating the proximal operator using the PV-early solver is an alternative derivation of the constrainedABA. Compared to this approach, our ALM-based derivation in Section III is more accessible.

Previous work [10] had presented PV-soft and PV-early as two separate algorithms. However, from a proximal perspective, we have shown that PV-soft is identical to one iteration of constraintABA and that constraintABA can be derived by applying the PV-early algorithm on the proximal problem in Eq. (48). This reveals an interesting connection between the PV-soft and PV-early algorithms, namely that PV-soft can be derived from the PV-early algorithm using the proximal formulation.

VIII. COMPUTATIONAL BENCHMARKING

We now present the benchmarking results of this paper's algorithms on diverse robots, namely Kuka Iiwa (7 DoF chain), Solo [55]/Unittree Go1 (18 DoF tree), Talos (50 DoF tree) [56], and Atlas with two shadow-hands attached to each wrist (84 DoF tree). We also study the computational scaling

on a chain mechanism and a balanced binary tree mechanism of varying numbers of links. This studies the impact of different robot topologies on the presented algorithms.

Constraints on a hand, fingertip, or feet in the benchmarks are represented as H_{m_i} , T_{m_i} or F_{m_i} respectively where m_i is the constraint dimension. $m_i = 3$ for connect-type constraints where the constrained frame's translation is constraint, but its rotation is unconstrained. This can model a quadruped's foot contact constraints. $m_i = 6$ for weld-type constraints, where the constrained link's motion is restricted in all six directions. Notations such as $F_{m_i}^2$ would indicate that F_{m_i} constraint is imposed on each of the two robot feet. T_3^i would imply that connect-type constraints were imposed on i different fingertips of the Shadow Hand. We also consider the typical case of modeling a humanoid's foot contact with 4 connect-type constraint, which is denoted by F_3^4 (note that this symbol for a quadruped would instead imply a connect-type constraint on each of its four feet).

We first discuss our implementation of the algorithms in Section VIII-A, followed by benchmarking the constrained dynamics algorithms and the damped Delassus inverse algorithms in Section VIII-B and, Section VIII-C respectively.

A. Efficient C++ implementation

We first prototyped the algorithms in a MATLAB environment using CASADI [57] followed by a C++ implementation within the PINOCCHIO library [6]. ConstrainedABA and proxPV can be found in the `pv.hxx` file², while both proxLTL and proxLTLs can be found in the `constrained-dynamics.hxx` file³. The cABA-OSIM algorithm can be found in the `PvDelassusMatrix` function in the `delassus.hxx` file⁴. With CASADI's SX expressions, we can conveniently count the number of atomic operations such as additions, divisions, multiplications, sin/cos computations required by each algorithm. As explained in [10], this allows comparing the algorithms in the classical manner based on the number of operations, which is both implementation and computer architecture agnostic. However, on modern CPU architectures, vectorization, and memory access can significantly influence computation timings, therefore we also compare based on the computation timings of a C++ implementation of the algorithms within the PINOCCHIO library. However, this timing comparison is admittedly influenced by code implementation details, compilers as well as on CPU architecture. The LTL implementation [22] in the PINOCCHIO library is particularly efficient due to its usage of CPU vectorization. Similarly, vectorizing constrainedABA and proxPV in a general manner is not straightforward and is deemed to be outside this paper's scope.

All timings were benchmarked on a 13th Gen Intel® Core™ i9-13950HX laptop CPU running Ubuntu 22.04LTS operating system. We present both results

²<https://github.com/stack-of-tasks/pinocchio/blob/d091266a94b5e2357862548ecfe0836e230a2c75/include/pinocchio/algorithm/pv.hxx>

³<https://github.com/stack-of-tasks/pinocchio/blob/912aa3dba3942c24639ce4c6be4c5ed1b4eca15a/include/pinocchio/algorithm/constrained-dynamics.hxx>

⁴<https://github.com/stack-of-tasks/pinocchio/blob/d091266a94b5e2357862548ecfe0836e230a2c75/include/pinocchio/algorithm/delassus.hxx>

with Intel’s Turbo Boost both turned off to minimize CPU frequency variability and with Turbo Boost turned on to indicate the speeds achievable on a modern laptop in operational scenarios, when Turbo Boost is typically active. The code was compiled using Clang-14 compiler with the usual optimized compilation flags `-O3 -march=native`. All the listed computation timings are averaged over 100000 samples.

B. Constrained dynamics

We now benchmark constrainedABA, proxPV, proxLTLs, and proxLTL algorithms in terms of the number of operations and the computation timings of their C++ implementations.

1) *Number of operations*: Fig. 1 lists the number of operations, with their internal breakdown for proximal iterations, for the different algorithms on different robots. H^A refers to the cost of computing the articulated body inertia matrices in constrainedABA and proxPV algorithms. This operation is more expensive in constrainedABA due to the inertia matrix and constraint forces updates in Eq. (32). M_μ^{-1} and M^{-1} blocks for proxLTLs and proxLTL denote the cost of computing and factorizing M_μ and M blocks, respectively. Computing M_μ^{-1} is more expensive than M^{-1} due to the JSIM augmentation step in Eq. (27). The Λ block for proxPV and proxLTL denotes the cost of factorizing the Λ^{-1} matrix. Finally, the proximal block for all the algorithms denotes the number of operations required by three proximal iterations.

The proximal iterations for constrainedABA and proxLTLs are more expensive than for proxPV and proxLTL respectively, as the latter two algorithms only solve the smaller dual problem during proximal iterations. ConstrainedABA’s proximal iterations are more expensive than that of proxLTLs because the former performs a reduced backward and forward sweep over all the joints supporting constraints, while proxLTLs computes this using the constraint Jacobian and Cholesky decomposition of M_μ . Proximal iterations of proxPV and proxLTL should cost the same, but proxPV’s proximal iterations are slightly more expensive because of an implementation detail. We have employed an early elimination at the base strategy [10], which speeds up factorization of Λ^{-1} (as seen in the reduced size of the Λ block for proxPV), but increases the cost of solving for the dual variables.

For the Go1 robot, the operation count of all the algorithms is similar, with constrainedABA requiring approximately 20% fewer operations than proxLTL. For larger robots as well as an increased number of constraints, the lower complexity algorithms clearly demonstrate their superiority.

2) *Computation timings*: Fig. 2 lists the computational timings benchmarking of the C++ implementations of the algorithms with the TurboBoost disabled. We find constrainedABA to be faster than the other algorithms for all the considered robots and constraint scenarios, except for one case involving the Kuka iiwa robot arm with seven DoFs. Both constrainedABA and proxPV algorithms become relatively more efficient than the proxLTL algorithm for bigger robots due to their linear complexity w.r.t n . ConstrainedABA and proxLTLs were found to scale gracefully with an increasing

number of constraints compared to the proxLTL or the proxPV algorithms due to their linear complexity w.r.t m . However, as expected from the previous section, the computational cost of each proximal iteration is greater for constrainedABA, which performs a reduced backward and forward sweep, compared to proxPV or proxLTL, which only solves the dual problem iteratively. Since TurboBoost is enabled in practice while simulating robotic systems, Table II includes the mean and standard deviations for the computation timings with TurboBoost enabled, for reference. In Table II, we include a redundant constraint case when there are 4 3D constraints at the convex hull of the robots’s feet in order to capture the contact wrench cone [58]. It also considers end-effector constrained serial chains composed of 3 DoF joints, consisting of 7 (21 DoF total) or 14 (42 DoF) links.

Computational scaling for chains. For an end-effector constrained chain (with 6D weld-type constraint), we now study the computational scaling of the algorithms w.r.t increasing chain sizes. The timings results are plotted for one proximal iteration and three proximal iterations in Fig. 3a and, Fig. 3b respectively. As expected, we observe a linear increase in computational cost for constrainedABA and proxPV and a superlinear increase for the proxLTL algorithm as the number of links increases. ConstrainedABA is over 4x faster than proxLTL for a chain of 100 links if we allow only one proximal iteration. However, when three proximal iterations are allowed, constrainedABA becomes nearly as expensive as the proxPV algorithm. As a testament to the efficient vectorized implementation of LTL in PINOCCHIO, the cost of proxLTL and proxLTLs increased only by 4x for 100 links compared to 50 links in Fig. 3a. Alternatively, if we used the dense Cholesky decomposition in the EIGEN library in place of PINOCCHIO’s LTL, the cost for 100 link chain was found to be 8x higher than the cost for 50 link chain. This is expected due to the cubic complexity of LTL for chains.

Computational scaling for balanced binary trees. Finally, we benchmark the various algorithms on a balanced binary tree mechanism, a structure that most favours the LTL-based algorithms since the kinematic tree depth d is minimal for a balanced binary tree with n links if we permit a branching factor of 2. Each branch in the tree has three links before it splits into two descendant branches. We constrain all the leaf links with a 3D connect-type constraint. Therefore, a balanced binary tree with n links has $(n \bmod 6 + 1)3$ constraints and the results are plotted in Fig. 3c and Fig. 3d for one and three proximal iterations respectively. Therefore, we can also see the algorithms’ scaling w.r.t constraints in this test case. As expected, we observe constrainedABA to scale linearly. However, surprisingly, proxLTLs is also particularly efficient. This is because its complexity for the favorable binary tree example is $O(n \log_2^2(n) + md)$ linearly, and it benefits from the efficient vectorized implementation of LTL [22] in the PINOCCHIO library. Both proxPV and proxLTL scale poorly due to their cubic complexity w.r.t. to the number of constraints.




Fig. 1: Benchmarking the number of operations of the constrained dynamics algorithms with three proximal iterations.




Fig. 2: Computational benchmarking of the C++ implementation of the constrained FD algorithms with TurboBoost disabled. The ‘prox’ component in the bar graphs corresponds to the computational cost of three proximal iterations for each algorithm.

3) *Convergence results:* Empirically, the proximal method was found to converge quite reliably in about three iterations. This should not be surprising since we are solving small equality-constrained strongly convex QPs. Fig. 4 plots a representative convergence whisker plot for 1000 randomly generated examples for a Talos robot with a challenging redundant constraint formulation. μ was chosen to be 10^6 as it provided a reasonable trade-off between fast convergence and numerical issues. Interestingly, we find the proxLTLs and constrainedABA to converge to smaller errors compared to proxLTL and proxPV algorithms. This is because, in the former algorithms, the primal error is fed back in the ALM

formulation, effectively providing iterative refinement for free, while in the latter algorithms, the dual problem is solved till convergence before forward-sweep to get the primal variables.

We now investigate the convergence of constrainedABA to a high accuracy of 10^{-10} for the various robots and constraint scenarios considered in Table III. Similarly to the previous paragraph, $\mu = 10^6$ for all the robots and is kept constant during the iterations. ConstrainedABA was implemented on 100000 examples with randomly sampled positions, velocities and torques. We select the the constraint scenario for which constrainedABA performed worst for a given robot, and for this scenario, the percentage of cases that converged to

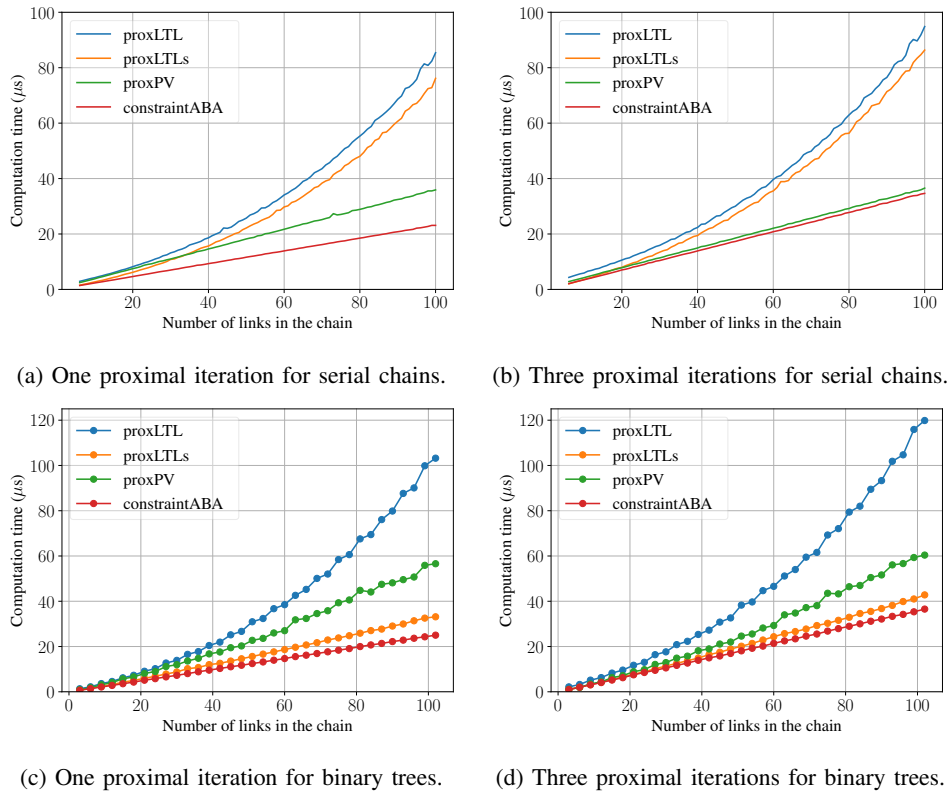


Fig. 3: Benchmarking the computational scaling of the different algorithms for serial chains and binary trees.

TABLE II: Computational timings of the algorithms in μs with Turbo Boost enabled and is averaged over 100000 samples with the standard deviation displayed in the parentheses. Each algorithm is allowed three proximal iterations.

System	cABA	PV	LTLs	LTL
Iiwa- H_3	1.07(0.40)	1.03(0.23)	1.01 (0.35)	1.58(0.51)
Iiwa- H_6	1.10(0.45)	1.43(0.51)	1.03 (0.427)	2.10(0.34)
Solo- F_3^2	2.00 (0.60)	2.25(0.69)	2.23(0.76)	3.18(0.74)
Solo- F_3^4	2.45(0.54)	3.05(0.73)	2.35 (0.73)	4.15(0.5)
Talos- F_6^2	5.43 (1.08)	6.83(0.60)	7.32(0.54)	10.6(0.79)
Talos- $F_6^2 H_6^2$	6.69 (0.65)	(1.0)	7.77(0.6)	15.1(0.73)
Talos- F_3^8	6.16 (1.2)	8.76(1.1)	7.47(1.2)	14.3(1.2)
Atlas SR- F_6^2	8.36 (1.34)	10.1(1.3)	13.8(1.0)	18.1(1.4)
Atlas SR- $F_6^2 T_3^5$	10.8 (1.4)	14.9(0.7)	14.8(1.5)	26.8(1.7)
Atlas SR- $F_6^2 T_3^{10}$	12.5 (1.2)	21.6(1.2)	15.2(1.2)	37.5(1.3)
Chain 7 SP- H_6	1.91 (0.74)	2.42(0.97)	2.21(0.83)	3.53(0.90)
Chain 14 SP- H_6	3.80 (1.0)	4.54(1.0)	6.21(1.1)	7.67(1.4)

high accuracy within a given number of iterations is plotted in Fig. 5. While a large majority of the problems (all for the humanoid platforms) considered were solved to a high-accuracy within 5 iterations. About 0.5% of cases for the Iiwa and Solo required more than 10 iterations. Notably, Iiwa's and Solo's convergence improves significantly for higher and lower values of μ respectively. This highlights the limitations of keeping μ fixed and calls for μ -scheduling similarly to [39]. However, this aspect must be left to the outer solver (e.g. [59]) simulating physics using the this paper's algorithms as the inner solver.




Fig. 4: Benchmarking convergence of the different proximal dynamics algorithms on a Talos robot with 21 constraints ($F_6^2 H_6$ and 3D constraint on elbow). This leads to a redundant constraint formulation that requires proximal methods. A whisker plot of convergence of the constraint residual is plotted for 10 proximal iterations $\mu = 10^6$ for 1000 randomly generated examples.

C. Benchmarking cABA-OSIM

We now benchmark the damped Delassus inverse algorithms, cABA-OSIM and cABA-OSIMf (an accelerated version of cABA-OSIM tailored for floating-base robots with branching at the base) introduced in Section VI. We compare our cABA-OSIM and cABA-OSIMf with the existing state-of-

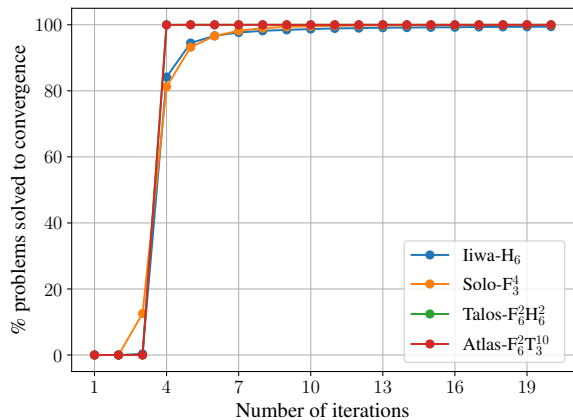


Fig. 5: Benchmarking the percentage of problems solved to a high accuracy of 10^{-10} within the plotted number of iterations. $\mu = 10^6$ for all the robot setups. 10000 samples were randomly generated.

the-art OSIM algorithms PV-OSIMr [30] and PV-OSIMf [10], along with the standard higher-complexity LTL-OSIM [21] algorithm. Similarly to the constrained dynamics algorithms, we first compare the algorithms by the number of operations before moving on to the computation timings of their C++ implementations in the PINOCCHIO library.

Number of operations Fig. 6 lists the number of operations required by different OSIM algorithms for the Unitree Go1 robot, Talos robot and the Atlas robot with a Shadow Hand attached to each wrist in Fig. 6a, Fig. 6b, Fig. 6c respectively. In the plots, H^A refers to the cost of computing the articulated body inertias, which is slightly more expensive for the cABA algorithms compared to the PV algorithms due to additional inertia updates performed in line 4 of the Algorithm 5. The LTL-OSIM algorithm is in general the most expensive algorithm and scales poorly to higher dimensional robots as well. Both the PV-OSIMr and the LTL-OSIM algorithms scale poorly to the increasing number of constraints compared to cABA-OSIM, due to the cubic complexity they incur during the Delassus matrix factorization.

Since cABA-OSIMf and PV-OSIMf are both tailored for floating-base robots with branching at the base, they outperform the other algorithms in benchmark cases, where the robots conform to this structure. However, for a higher number of constraints, the lower complexity cABA-OSIMf algorithm scales better because it does not require the factorization of L_1^A , unlike the PV-OSIMf algorithm.

Computation timings Table III lists the computation timings for a subset of the algorithms from the previous section, that we implemented in C++. Similarly to the constrained dynamics benchmarking, we enabled `TURBOBOOST` since this is more reflective of the actual usage of these algorithms. However, there are two main differences compared to the previous benchmarks. Firstly, we also include the cost for computing the damped Delassus inverse using LTL in the

TABLE III: Computational timings in μs of the damped Delassus inverse algorithms with Turbo Boost enabled are displayed below. The timings are averaged over 100000 samples and the standard deviations are displayed in the parentheses. * LTL-chol computes the Cholesky decomposition of Λ_μ^{-1} unlike the other algorithms which actually compute Λ_μ .

System	cABA-f	cABA	LTL-inv	LTL-chol*
Iiwa - H ₃	0.74(0.7)	0.74(0.76)	0.66 (0.7)	0.58(0.7)
Iiwa - H ₆	0.82 (0.8)	0.81 (0.8)	0.84(0.6)	0.67(0.8)
Solo - F ₃ ²	1.63 (0.8)	1.75(0.8)	1.64 (0.9)	1.44(0.8)
Solo - F ₃ ⁴	2.00 (1.0)	2.38(1.0)	2.64(1.1)	1.95(1.1)
Talos - F ₆ ²	4.28 (1.6)	4.62(1.8)	6.74(1.8)	6.01(1.66)
Talos - F ₆ ² H ₆ ²	5.51 (1.5)	5.86(1.4)	11.3(1.3)	8.79(1.6)
Atlas SR - F ₆ ²	6.70 (1.4)	6.82 (1.4)	13.1(2.9)	12.6(2.6)
Atlas SR - F ₆ ² T ₃ ⁵	8.63 (1.7)	9.21(1.83)	21.4(2.5)	17.8(2.1)
Atlas SR - F ₆ ² T ₃ ¹⁰	12.5 (2.0)	13.4(2.1)	37.2(5.1)	26.3(2.5)

last column, to make the comparison fairer with our cABA-OSIM algorithms, which compute the inverse and not just the Cholesky decomposition of Λ_μ^{-1} . Computing the inverse is interesting because it makes computing analytical derivatives faster because of vectorization effects [60]. Secondly, the cABA-OSIM and cABA-OSIMf algorithms are implemented in a global frame instead of a local frame compared to the previous sections. Using global frame negatively affects cABA-OSIM and cABA-OSIMf more than the LTL methods because of the loss of sparsity in the computation of D_i and P_i (note that S_i typically has only one non-zero element for a local frame implementation, while it is dense for a global frame implementation). This choice was made to facilitate our future work on using the Delassus inverse matrix to compute analytical derivatives of constrained dynamics, which is faster with global frames.

As expected cABA-OSIM and cABA-OSIMf significantly outperform the LTL methods for larger robots. However, for smaller robots including an 18 DoF quadruped, the LTL method remains competitive.

IX. DISCUSSION

The proximal formulation of the equality-constrained dynamics problem results in four different algorithms, namely proxLTL, proxLTLs, proxPV, and constrainedABA depending on the usage of maximal or minimal coordinates during derivation and the order of elimination of the primal and dual variables, as shown in the overview figure Fig. 7. These algorithms generalize several existing algorithms. For instance, MUJOCO’s equality-constrained dynamics algorithm and PV-soft [10] can be considered to be the first iteration of proxLTLs and constrainedABA respectively. Our analysis also revealed that PV-soft can be recovered by applying PV-early on the proximal formulation, thereby connecting two existing algorithms that were considered to be separate.

Though applying PV-early’s idea of aggressively eliminating dual variables to the proximal dynamics formulation immediately gives us the first iteration of constrainedABA, constrainedABA generalizes this idea by introducing proximal iterations. Moreover, PV-early’s superior performance in [10]




Fig. 6: Benchmarking the number of operations of the damped Delassus inverse/factorization algorithms.

	λ before \dot{v}	\dot{v} before λ
maximal	<p>constrainedABA $O(n + m)$</p>	<p>proxPV $O(n + m^2d + m^3)$</p>
minimal	<p>proxLTLs $O(nd^2 + md)$</p>	<p>proxLTL $O(nd^2 + m^2d + md^2 + m^3)$</p>

Fig. 7: Overview of the proximal dynamics algorithms.

relied on an efficient formula to compute the SVD of a symmetric rank-1 (SR1) 6×6 matrix arising due to one DoF joints and would become inefficient for more general higher DoF joints. However, constrainedABA avoids this limitation by eliminating all dual variables, thereby not requiring any SVD computation, and is therefore readily applicable to robots with higher DoF joints without compromising performance. However, we must note that we expect the vanilla PV-early algorithm (without proximal regularization) to marginally outperform constrainedABA for robots that predominantly possess single DoF joints and no singular constraints, as it does not require an iterative procedure. However, constrainedABA is significantly easier to implement than the vanilla PV-early algorithm, remains efficient for multi-DoF joints, and is valid even in singular cases.

ConstrainedABA, the main contribution of this paper, is the first algorithm with a $O(n + m)$ complexity that can reliably deal with singular constraints without biasing solutions towards the origin or resorting to expensive SVD computation. We also derived cABA-OSIM, a $O(n + m^2)$ complexity algorithm, to compute the damped Delassus inverse algorithm, also known as the OSIM. All existing OSIM algorithms have an extra m^3 term in the general case due to the factorization of the Delassus matrix. Since the damped Delassus inverse matrix is an $m \times m$ matrix and its computation requires considering all the joints, the asymptotic complexity achieved by us is optimal.

In addition to deriving these algorithms, we implemented them in C++, benchmarked them for various robots, and studied their scaling for chains and binary trees. The comparison between the algorithms was made based on the number of

operations and the computation timings of an efficient C++ implementation. The lower-complexity algorithms, constrainedABA, and proxPV were significantly more efficient than proxLTLs and proxLTL regarding the number of operations for larger robots like humanoids. However, the performance difference between the low-complexity and high-complexity algorithms was reduced for the C++ implementation. This is because the high-complexity algorithms enjoy CPU vectorization benefits from an efficient implementation in [22]. Devising a similarly efficient vectorized implementation for the low-complexity algorithms is far from straightforward as there is an inherent sequentiality in the recursive algorithm, so this aspect is not considered within the scope of this article. Nevertheless, constrainedABA emerged to be the fastest of the four algorithms for quadrupeds and humanoids even for the C++ implementation, being over 2x faster than proxLTL, the previously existing state-of-the-art C++ implementation. However, the proxLTLs algorithm turned out to be surprisingly competitive with constrainedABA by requiring only 20% more time for a heavily constrained Talos robot.

Our optimal complexity algorithm, cABA-OSIM, was also benchmarked with existing algorithms in terms of the number of operations required and computation timing of a C++ implementation. While it was found to require over 3x fewer operations compared to the LTL-OSIM methods, this difference was reduced for the C++ timings because of the efficient vectorized implementation of LTL-OSIM in the PINOCCHIO library. Despite this reduction, the cABA-OSIM remained over 2x faster than the LTL method for a large robot-like Talos.

X. CONCLUSION

We have applied the proximal point algorithm from optimization literature to equality-constrained robot dynamics problems for kinematic trees. The proximal formulation enables handling challenging singular cases efficiently and solves motion constraints exactly up to machine precision often within 5 iterations. Focusing on the algorithmic aspect of proximal constrained dynamics algorithms, our extensive treatment derived four different algorithms, proxLTL, proxLTLs, proxPV, and constrainedABA depending on the choice of maximal/minimal coordinate formulation and the primal/dual elimination order. Two of these (constrainedABA

and proxPV) are original contributions of this article, while the third (proxLTLs) is implemented and benchmarked for the first time, though previously mentioned in [22].

ConstrainedABA, this article’s main contribution, is the first linear complexity algorithm that deals with singular cases without resorting to expensive SVD computation, that we are aware of, and perhaps the simplest. The proximal formulation and the matrix inversion lemma also led to a beautiful result, where any Delassus matrix algorithm could be used to directly compute the damped Delassus inverse matrix, avoiding the expensive matrix factorization step. By using PV-OSIMr algorithm, we obtained cABA-OSIM algorithm with an optimal complexity of $O(n + m^2)$. The algorithms are implemented in C++ as a part of the Pinocchio library and are extensively benchmarked. ConstrainedABA emerged as the fastest out of the four algorithms for larger robots like quadrupeds and humanoids. The higher complexity proxLTLs surprisingly remained competitive even for larger robots, when heavily constrained, due to linear complexity in constraint dimension and an efficient vectorized C++ implementation.

The proximal formulation generalizes and establishes connections between existing algorithms like MUJoCo’s solver, PV-soft, and PV-early. It is a powerful formulation that enables efficient trading-off between MUJoCo-style compliance and an expensive SVD-style rigid contact during singular cases (and also in general) through fast proximal iterations. ConstrainedABA and proxLTLs, in particular, are fairly straightforward to implement, by introducing only a few new lines of additional code compared to Featherstone’s original ABA and LTL algorithms. This facilitates their implementation in existing simulators like MUJoCo, DRAKE, or RAISIM, to name a few of the most popular within the robotics community.

The favorable properties of the proximal point algorithm such as fast proximal iterations, differentiability of the proximal operator, and the cABA-OSIM set us up for several interesting future directions. We expect constrainedABA and cABA-OSIM, in particular, to serve as the computational backbone for our planned future work on extending proximal methods to efficiently deal with inequality constraints [39], [61], frictional contacts [4], [62]–[65], and analytical derivatives [6], [43], [60], [66].

ACKNOWLEDGEMENTS

The authors warmly thank Antoine Bambade for helpful discussions on the proximal-point method and augmented Lagrangian algorithms.

This work was supported by the Belgium government through the FWO project GOD1119N of the Research Foundation - Flanders (FWO - Flanders), by the French government under the management of Agence Nationale de la Recherche through the project INEXACT (ANR-22-CE33-0007-01) and as part of the “Investissements d’avenir” program, reference ANR-19-P3IA-0001 (PRAIRIE 3IA Institute), by the European Union through the AGIMUS project (GA no.101070165) and the Louis Vuitton ENS Chair on Artificial Intelligence. Views and opinions expressed are those of the author(s) only and do not necessarily reflect those of the European Union or

the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

REFERENCES

- [1] R. Featherstone, *Rigid body dynamics algorithms*. Springer, 2014.
- [2] E. Coumans and Y. Bai, “Pybullet, a python module for physics simulation for games, robotics and machine learning,” <http://pybullet.org>, 2016–2021.
- [3] J. Lee, M. X. Grey, S. Ha, T. Kunz, S. Jain, Y. Ye, S. S. Srinivasa, M. Stilman, and C. K. Liu, “DART: Dynamic animation and robotics toolkit,” *The Journal of Open Source Software*, vol. 3, no. 22, p. 500, Feb 2018. [Online]. Available: <https://doi.org/10.21105/joss.00500>
- [4] J. Hwangbo, J. Lee, and M. Hutter, “Per-contact iteration method for solving contact dynamics,” *IEEE Robot. Autom. Lett.*, vol. 3, no. 2, pp. 895–902, 2018. [Online]. Available: www.raisim.com
- [5] E. Todorov, T. Erez, and Y. Tassa, “Mujoco: A physics engine for model-based control,” in *Proc. IEEE/RSJ Int. Conf. Int. Robots. Syst.* IEEE, 2012, pp. 5026–5033.
- [6] J. Carpentier, G. Saurer, 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 *2019 IEEE/SICE International Symposium on System Integration (SII)*. IEEE, 2019, pp. 614–619.
- [7] A. F. Vereshchagin, “Modeling and control of motion of manipulative robots,” *Soviet Journal of Computer and Systems Sciences*, vol. 27, no. 5, pp. 29–38, 1989.
- [8] D.-S. Bae and E. J. Haug, “A recursive formulation for constrained mechanical system dynamics: Part ii. closed loop systems,” *Journal of Structural Mechanics*, vol. 15, no. 4, pp. 481–506, 1987.
- [9] M. Otter, H. Brandl, and R. Johanni, “An algorithm for the simulation of multibody systems with kinematic loops,” in *Proceedings of the 7th World Congress on Theory of Machines and Mechanisms, IFToMM, Sevilla, Spain, 1987*.
- [10] A. S. Sathya, H. Bruyninckx, W. Decré, and G. Pipeleers, “Efficient constrained dynamics algorithms based on an equivalent lqr formulation using gauss’ principle of least constraint,” *IEEE Transactions on Robotics*, vol. 40, pp. 729–749, 2024.
- [11] A. Vereshchagin, “Computer simulation of the dynamics of complicated mechanisms of robot-manipulators,” *Eng. Cybernet.*, vol. 12, pp. 65–70, 1974.
- [12] C. F. Gauß, “Über ein neues allgemeines grundgesetz der mechanik.” 1829.
- [13] R. Bellman, “Dynamic programming,” *Science*, vol. 153, no. 3731, pp. 34–37, 1966.
- [14] R. Featherstone, “The calculation of robot dynamics using articulated-body inertias,” *Int. J. Robot. Res.*, vol. 2, no. 1, pp. 13–30, 1983.
- [15] H. Brandl, R. Johanni, and M. Otter, “A very efficient algorithm for the simulation of robots and similar multibody systems without inversion of the mass matrix,” *IFAC Proceedings Volumes*, vol. 19, no. 14, pp. 95–100, 1986.
- [16] G. Rodriguez, “Kalman filtering, smoothing, and recursive robot arm forward and inverse dynamics,” *IEEE Journal on Robotics and Automation*, vol. 3, no. 6, pp. 624–639, 1987.
- [17] M. W. Walker and D. E. Orin, “Efficient dynamic computer simulation of robotic mechanisms,” 1982.
- [18] R. Featherstone, “Efficient factorization of the joint-space inertia matrix for branched kinematic trees,” *Int. J. Robot. Res.*, vol. 24, no. 6, pp. 487–500, 2005.
- [19] J. P. Popov, A. F. Vereshchagin, and S. L. Zenkevič, *Manipulacionnyje roboty: Dinamika i algoritmy*. Nauka, 1978.
- [20] G. H. Golub and C. F. Van Loan, *Matrix computations*. JHU press, 2013.
- [21] R. Featherstone, “Exploiting sparsity in operational-space dynamics,” *Int. J. Robot. Res.*, vol. 29, no. 10, pp. 1353–1368, 2010.
- [22] J. Carpentier, R. Budhiraja, and N. Mansard, “Proximal and sparse resolution of constrained dynamic equations,” in *Proc. Robot., Sci. Syst.*, 2021.
- [23] N. Parikh, S. Boyd *et al.*, “Proximal algorithms,” *Foundations and trends® in Optimization*, vol. 1, no. 3, pp. 127–239, 2014.
- [24] É. Delassus, “Mémoire sur la théorie des liaisons finies unilatérales,” in *Annales scientifiques de l’École normale supérieure*, vol. 34, 1917, pp. 95–179.
- [25] B. Brogliato, *Nonsmooth mechanics*. Springer, 1999, vol. 3.

- [26] 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.
- [27] K. Kreutz-Delgado, A. Jain, and G. Rodriguez, "Recursive formulation of operational-space control," *Int. J. Robot. Res.*, vol. 11, no. 4, pp. 320–328, 1992.
- [28] P. Wensing, R. Featherstone, and D. E. Orin, "A reduced-order recursive algorithm for the computation of the operational-space inertia matrix," in *Proc. IEEE Int. Conf. Robot. Autom.* IEEE, 2012, pp. 4911–4917.
- [29] K.-S. Chang and O. Khatib, "Efficient recursive algorithm for the operational space inertia matrix of branching mechanisms," *Advanced Robotics*, vol. 14, no. 8, pp. 703–715, 2001.
- [30] A. S. Sathya, W. Decré, and J. Swevers, "Pv-osimr: A lowest order complexity algorithm for computing the delassus matrix," *IEEE Robot. Autom. Lett.*, vol. 9, no. 11, pp. 10224–10231, 2024.
- [31] J. Sherman and W. J. Morrison, "Adjustment of an inverse matrix corresponding to a change in one element of a given matrix," *The Annals of Mathematical Statistics*, vol. 21, no. 1, pp. 124–127, 1950.
- [32] R. M. Murray, Z. Li, and S. S. Sastry, *A mathematical introduction to robotic manipulation*. CRC press, 2017.
- [33] F. E. Udewadia and R. E. Kalaba, *Analytical dynamics : a new approach*. Cambridge: Cambridge University press, 1996.
- [34] H. Bruyninckx and O. Khatib, "Gauss' principle and the dynamics of redundant and constrained manipulators," in *Proc. IEEE Int. Conf. Robot. Autom.*, vol. 3. IEEE, 2000, pp. 2563–2568.
- [35] J. Baumgarte, "Stabilization of constraints and integrals of motion in dynamical systems," *Computer methods in applied mechanics and engineering*, vol. 1, no. 1, pp. 1–16, 1972.
- [36] S. P. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [37] J. Nocedal and S. Wright, *Numerical optimization*. Springer Science & Business Media, 2006.
- [38] R. T. Rockafellar, "Monotone operators and the proximal point algorithm," *SIAM journal on control and optimization*, vol. 14, no. 5, pp. 877–898, 1976.
- [39] A. Bambade, S. El-Kazdadi, A. Taylor, and J. Carpentier, "Prox-qp: Yet another quadratic programming solver for robotics and beyond," in *RSS 2022-Robotics: Science and Systems*, 2022.
- [40] W. Jallet, A. Bambade, N. Mansard, and J. Carpentier, "Constrained differential dynamic programming: A primal-dual augmented lagrangian approach," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 13 371–13 378.
- [41] O. Güler, "On the convergence of the proximal point algorithm for convex minimization," *SIAM journal on control and optimization*, vol. 29, no. 2, pp. 403–419, 1991.
- [42] A. Chiche and J. C. Gilbert, "How the augmented lagrangian algorithm can deal with an infeasible convex quadratic optimization problem," *Journal of Convex Analysis*, vol. 23, no. 2, 2016.
- [43] A. Bambade, F. Schramm, A. Taylor, and J. Carpentier, "QPLayer: efficient differentiation of convex quadratic optimization," in *2024 International Conference on Learning Representations (ICLR)*, 2023.
- [44] M. R. Hestenes, "Multiplier and gradient methods," *Journal of optimization theory and applications*, vol. 4, no. 5, pp. 303–320, 1969.
- [45] M. J. Powell, "A method for nonlinear constraints in minimization problems," *Optimization*, pp. 283–298, 1969.
- [46] R. T. Rockafellar, "A dual approach to solving nonlinear programming problems by unconstrained optimization," *Mathematical programming*, vol. 5, no. 1, pp. 354–373, 1973.
- [47] R. Rockafellar, "Minimax theorems and conjugate saddle-functions," *Mathematica Scandinavica*, vol. 14, no. 2, pp. 151–173, 1964.
- [48] K. M. Lynch and F. C. Park, *Modern robotics*. Cambridge University Press, 2017.
- [49] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model predictive control: theory, computation, and design*. Nob Hill Publishing Madison, 2017, vol. 2.
- [50] M. L. Felis, "Rbdl: an efficient rigid-body dynamics library using recursive algorithms," *Autonomous Robots*, vol. 41, no. 2, pp. 495–511, 2017.
- [51] R. Tedrake and the Drake Development Team, "Drake: Model-based design and verification for robotics," 2019. [Online]. Available: <https://drake.mit.edu>
- [52] K. W. Lilly, *Efficient dynamic simulation of multiple chain robotic systems*. The Ohio State University, 1989.
- [53] S. R. Buss and J.-S. Kim, "Selectively damped least squares for inverse kinematics," *Journal of Graphics tools*, vol. 10, no. 3, pp. 37–49, 2005.
- [54] E. Todorov, "Convex and analytically-invertible dynamics with contacts and constraints: Theory and implementation in mujoco," in *Proc. IEEE Int. Conf. Robot. Autom.* IEEE, 2014, pp. 6054–6061.
- [55] F. Grimmering, A. Meduri, M. Khadiv, J. Viereck, M. Wüthrich, M. Naveau, V. Berenz, S. Heim, F. Widmaier, T. Flayols *et al.*, "An open torque-controlled modular robot architecture for legged locomotion research," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3650–3657, 2020.
- [56] O. Stasse, T. Flayols, R. Budhiraja, K. Giraud-Esclasse, J. Carpentier, J. Mirabel, A. Del Prete, P. Souères, N. Mansard, F. Lamiraux *et al.*, "Talos: A new humanoid research platform targeted for industrial applications," in *2017 IEEE-RAS 17th International Conference on Humanoid Robotics (Humanoids)*. IEEE, 2017, pp. 689–695.
- [57] J. A. Andersson, J. Gillis, G. Horn, J. B. Rawlings, and M. Diehl, "Casadi: a software framework for nonlinear optimization and optimal control," *Mathematical Programming Computation*, vol. 11, no. 1, pp. 1–36, 2019.
- [58] S. Caron, Q.-C. Pham, and Y. Nakamura, "Stability of surface contacts for humanoid robots: Closed-form formulae of the contact wrench cone for rectangular support areas," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 5107–5112.
- [59] J. Carpentier, L. Montaut, and Q. L. Lidec, "From compliant to rigid contact simulation: a unified and efficient approach," *arXiv preprint arXiv:2405.17020*, 2024.
- [60] J. Carpentier and N. Mansard, "Analytical derivatives of rigid body dynamics algorithms," in *Proc. Robot. Sci. Syst.*, 2018.
- [61] W. Jallet, A. Bambade, E. Arlaud, S. El-Kazdadi, N. Mansard, and J. Carpentier, "PROXDDP: Proximal Constrained Trajectory Optimization," Dec. 2023, working paper or preprint. [Online]. Available: <https://inria.hal.science/hal-04332348>
- [62] P. C. Horak and J. C. Trinkle, "On the similarities and differences among contact models in robot simulation," *IEEE Robotics and Automation Letters*, vol. 4, no. 2, pp. 493–499, 2019.
- [63] A. M. Castro, F. N. Permenter, and X. Han, "An unconstrained convex formulation of compliant contact," *IEEE Transactions on Robotics*, vol. 39, no. 2, pp. 1301–1320, 2022.
- [64] Q. L. Lidec, W. Jallet, L. Montaut, I. Laptev, C. Schmid, and J. Carpentier, "Contact Models in Robotics: a Comparative Analysis," *arXiv preprint arXiv:2304.06372*, 2023.
- [65] Q. Le Lidec and J. Carpentier, "Reconciling RaiSim with the Maximum Dissipation Principle," Feb. 2024, working paper or preprint. [Online]. Available: <https://hal.science/hal-04438175>
- [66] K. Werling, D. Omens, J. Lee, I. Exarchos, and C. K. Liu, "Fast and Feature-Complete Differentiable Physics Engine for Articulated Rigid Bodies with Contact Constraints," in *Proceedings of Robotics: Science and Systems*, Virtual, July 2021.



Ajay Sathya is a postdoctoral researcher at Inria and École Normale Supérieure, in the Willow research team since December 2023. He obtained his Bachelors degree from NITK Surathkal, India in 2016 and Masters and PhD degrees from KU Leuven, Belgium in 2018 and 2023 respectively. His research interests include developing efficient dynamics simulators and high-performance robot controllers.



Justin Carpentier is a researcher at Inria and École Normale Supérieure, heading the Willow research team since 2023. He graduated from École Normale Supérieure Paris-Saclay in 2014 and received a Ph.D. in Robotics in 2017 from the University of Toulouse. He did his Ph.D. in the Gepetto team at LAAS-CNRS in Toulouse, working on the computational foundations of legged locomotion. In 2024, he received an ERC Starting Grant focusing on laying the algorithmic and computational foundations of Artificial Motion Intelligence. His research interests

lie at the interface of optimization, machine learning, computer vision, simulation, and control for robotics, with applications ranging from agile locomotion to dexterous manipulation. He is also the main developer and manager of widely used open-source robotics software, among them Pinocchio, ProxSuite, HPP-FCL, and Aligator.