

Constrained Articulated Body Algorithms for Closed-Loop Mechanisms

Ajay Suresha Sathya  and Justin Carpentier , *Member, IEEE*

Abstract—Efficient rigid-body dynamics algorithms are instrumental in enabling high-frequency dynamics evaluation for resource-intensive applications (e.g., model-predictive control, large-scale simulation, and reinforcement learning), potentially on resource-constrained hardware. Existing recursive algorithms with low computational complexity are mostly restricted to kinematic trees with external contact constraints or are sensitive to singular cases (e.g., linearly dependent constraints and kinematic singularities), severely impacting their practical usage in existing simulators. This article introduces two original low-complexity recursive algorithms: the loop-constrained articulated body algorithm and proximal BBO (Brandl, Bae, and others), both based on a proximal dynamics formulation for forward simulation of closed-loop mechanisms. These algorithms are derived from first principles using nonserial dynamic programming, exhibit linear complexity in practical scenarios, and are numerically robust in the face of singular cases. They extend the existing constrained articulated body algorithm to handle internal loops and the pioneering BBO algorithm from the 1980s to singular cases. Both algorithms have been implemented by leveraging the open-source Pinocchio library, benchmarked in detail, and demonstrate state-of-the-art performance for various robot topologies, including over $6\times$ speed-ups compared to existing nonrecursive algorithms for high-degree-of-freedom systems with internal loops, such as recent humanoid robots.

Index Terms—Direct/inverse dynamics formulation, dynamics, humanoid robots, optimization and optimal control.

NOMENCLATURE

n Number of robot degrees of freedom.
 m Number of motion constraints (excluding spanning-tree joints).
 d Depth of the kinematic tree.
 $\mathcal{C}(\cdot)$ Returns cardinality of a set.
 \mathbf{q} Robot configuration.

Received 2 October 2025; accepted 18 November 2025. Date of publication 12 January 2026; date of current version 27 January 2026. This work was supported in part by the European Union’s Horizon Europe research and innovation programme through a Marie Skłodowska-Curie Postdoctoral Fellowship under Grant 101211945—ExTRAORDiNary, in part by the ARTIFACT project under Grant 101165695, in part by through the AGIMUS project under Grant 101070165, in part by the French government under the management of Agence Nationale de la Recherche through the project INEXACT under Grant ANR-22-CE33-0007-01, in part by the “PR[AI]RIE-PSAI” AI Cluster (ANR-23-IACL-0008), and in part by the Louis Vuitton ENS Chair on Artificial Intelligence. This article was recommended for publication by Associate Editor M. Posa and Editor P. M. Wensing upon evaluation of the reviewers’ comments. (Corresponding author: Ajay Suresha Sathya.)

The authors are with Inria, Département d’Informatique, École Normale Supérieure, PSL Research University, Paris 75013, France (e-mail: ajay.sathya@inria.fr; justin.carpentier@inria.fr).

Digital Object Identifier 10.1109/TRO.2026.3651683

\mathcal{Q} Robot configuration space.
 ν Generalized robot velocities.
 $\mathcal{T}_{\mathbf{q}}\mathcal{Q}$ Tangent space of \mathcal{Q} at \mathbf{q} .
 $\dot{\nu}$ Generalized robot accelerations.
 τ Generalized robot forces.
 $\mathcal{T}_{\mathbf{q}}^*\mathcal{Q}$ Dual tangent space of \mathcal{Q} at \mathbf{q} .
 \mathbf{v}_i i th link’s 6-D spatial velocity.
 \mathbb{M}^6 Motion vector space in spatial algebra.
 \mathbf{a}_i i th link’s 6-D 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 6-D spatial inertia tensor.
 \times Cross-product operator on spatial motion vectors.
 \times^* Cross-product operator on spatial force vectors.
 $\pi(i)$ i th link’s parent link.
 n_b Number of links in the mechanism.
 m_b Number of cut-joints in the mechanism.
 \mathcal{E} Set of all cut-joint indices.
 l_i^j Index of the j th link in the i th cut-joint.
 $\gamma(i)$ Set of i th link’s children link indices.
 \mathcal{S} Topologically ordered list of tree link indices.
 \mathcal{S}_r List \mathcal{S} reversed.
 $\text{SL}(i)$ Set of indices of supporting links of the i th loop.
 $\text{LS}(i)$ Set of indices of loops supported by the i th link.
 τ_i i th loop’s root link index.
 \mathcal{R}_i Set indexing loops rooted at the i th link.
 S_i Spans i th joint’s motion subspace.
 K_i^j Constraint matrix of the i th link for the j th loop.
 \mathbf{k}_i Desired constraint accelerations of the i th constraint.
 $\mathbf{a}_{b,i}$ i th link’s bias acceleration vector.
 M Joint-space inertia matrix (JSIM).
 $\dot{\nu}_{\text{tree}}$ Unconstrained spanning-tree generalized acceleration.
 \mathbf{a}_c Desired constraint accelerations.
 $\mathbf{J}_{\mathbf{f}_c}$ Geometric Jacobian of \mathbf{f}_c .
 λ Lagrange multipliers and constraint force magnitudes.
 μ Proximal operator / ALM parameter.
 Λ_μ Damped Delassus inverse matrix.
 M_μ Constraint augmented Inertia matrix.
 $H_{i,j}$ Inertia matrix coupling the i th and j th links in LCABA.
 \mathcal{N}_i Set of link indices neighboring the i th link in LCABA and $\text{LS}(i)$ in proxBBO.
 D_i Apparent constrained inertia felt at the i th joint.
 P_i Backward force propagation matrix at the i th joint.
 $\mathcal{S}^\mathcal{E}$ List of link indices in the LCABA elimination order.
 U_i $H_{i,i}S_i$.

m_c	Maximum number of neighbors for any link in LCABA.
$K_{i,j}$	Constraint matrix felt at the i th link due to the j th loop.
$L_{i,j}$	Constraint coupling matrix for the i th and j th cut-joint constraints.
\mathcal{U}_i	$LS(i) - \mathcal{R}_i$.
m_b	Maximum number loops supported by any link.

I. INTRODUCTION

EFFICIENT and reliable simulation of rigid-body dynamics is an important and extensively studied [1] problem in robotics. Efficient simulation is key for enabling computationally demanding downstream applications such as model-predictive control (MPC) [2] and reinforcement learning (RL) [3], and for generating large-scale synthetic datasets for modern data-hungry foundation models [4]. Faster simulation can unlock whole-body MPC for complex robots, allow for longer planning horizons that enhance optimality and safety, reduce the significant training and deployment time required for RL due to iterative reward and environment shaping, and even permit RL training on less expensive hardware. These simulation-driven techniques are at the forefront of research advancing reliable and real-time loco-manipulation with high-degree-of-freedom (DoF) robots, potentially in contact-rich scenarios.

Forward simulation algorithms in constrained scenarios, arising from bilateral [1], [5], unilateral, or frictional contact constraints [6], [7], [8], use a constrained dynamics algorithm (CDA) as a core inner solver, which solves the constrained forward dynamics problem. Constrained forward dynamics involves computing accelerations and constraint forces for a robot, given its configuration, velocity, applied forces, and desired constraint accelerations. To effectively simulate diverse loco-manipulation scenarios, the underlying CDA must efficiently handle a range of constraints, including those from contacts and kinematic loops, while being numerically robust to singular cases occurring at kinematic singularities or linearly dependent constraints.

Internal loops, i.e., kinematic loops of links that do not involve the ground link, constitute a particularly challenging class of constraints to simulate efficiently for existing CDAs [1], [5]. This constraint class is gaining importance as modern robot designs frequently incorporate kinematic loops to enhance mechanical properties. Even robots without inherent loops can form internal loops dynamically during operation, e.g., a robot hand grasping a cube in Fig. 1(a) or two humanoids in Fig. 1(b) jointly transporting a heavy object. Submechanisms such as gears and belt transmissions also result in internal loops [1, Sect. 9.6], which need to be accounted for to ensure simulation accuracy.

Despite their high relevance, only a few simulators, such as BULLET [9], MUJOCO [10], or SIMPLE [7], support internal loops, likely due to the challenge of simulating them efficiently. Even these select simulators use joint-space algorithms based on Featherstone's LTL algorithm [1], [11], [12]. However, the LTL algorithm does not fully exploit the available problem structure and is computationally expensive, with a complexity of $O(nd^2 + m^2d + md^2 + m^3)$, where n , d , and m are the




Fig. 1. Examples of closed-loops interactions or mechanisms classically encountered in robotics. (a) Several internal closed loops formed due to contact between the hand and the cube. (b) Humanoid platforms consisting of internal loops in addition to the loop formed due to collaborative manipulation.

robot DoFs, kinematic spanning tree depth, and the constraint dimensionality, respectively. SIMPLE employs an extension of the LTL algorithm, which we refer to as the proximal LTL (proxLTL) algorithm [5]. ProxLTL handles loops and leverages a proximal dynamics formulation to address singular cases exactly (up to solver tolerances) and efficiently. This avoids the need for approximate techniques such as Tikhonov regularization or expensive techniques such as the singular value decomposition (SVD). However, proxLTL retains the same computational complexity as the original LTL algorithm.

In contrast, for kinematic trees with only external loops (such as ground contact), there exist efficient recursive CDAs with lower computational cost, such as the PV algorithm [13], [14] with $O(n + m^2d + m^3)$ complexity and recent $O(n + m)$ complexity algorithms such as PV-soft, PV-early [15], proxPV, and constrained articulated body algorithm (constrainedABA) [16]. Efficient C++ implementations of these recursive algorithms are also available within the widely used open-source PINOCCHIO library [17]. There also exist pioneering recursive algorithms [18], [19] that can handle internal loops. These two works independently proposed what is practically the same algorithm, which we refer to as BBO (Brandl, Bae, and others) after the first authors of the two papers. The BBO algorithm is challenging to derive and implement and is vulnerable to singular cases. Likely due to these reasons, to the best of our knowledge, there is no efficient and user-friendly implementation of BBO available; it has not been benchmarked against the LTL algorithms and has seen limited use despite its potential advantages.

Addressing a key need for fast and reliable CDA that can handle internal loops, this article explores *recursive* algorithms that are robust to singular cases by leveraging the proximal dynamics formulation [5]. We introduce two efficient recursive CDAs: the loop-constrained articulated body algorithm (LCABA) and the proximal BBO (proxBBO) algorithm. LCABA and proxBBO extend the constrainedABA and proxPV algorithms [16] to handle internal loops, respectively. In addition, proxBBO generalizes the original BBO algorithm [18], [19] to be robust to singular cases using the proximal approach and is the recursive counterpart to the proxLTL algorithm [5]. These algorithms are derived by solving the quadratic program (QP) [20] associated with Gauss' principle of least constraint (GPLC) [21], [22],

[23] using dynamic programming (DP) [24] similarly to the approach in [15] and [16]. However, internal loops introduce a graph structure that requires nonserial DP [25, Ch. 10], [26], [27], leading to a variable elimination (also known as bucket elimination in probabilistic inference [28]) approach.

Article contributions: Our contributions are listed as follows.

- 1) *LCABA*: LCABA, a recursive efficient algorithm that can handle singular cases and internal loops with a best case complexity of $O(n + m)$, is derived by applying nonserial DP on the GPLC problem.
- 2) *ProxBBO*: The BBO algorithm [18], [19] is generalized to handle singular cases using a proximal dynamics formulation to obtain the proxBBO algorithm. The proxBBO algorithm is derived similarly to LCABA, but with a different variable elimination order, and also has a best case complexity of $O(n + m)$.
- 3) *Open-source and efficient implementations, detailed benchmarking, and analysis*: Both algorithms have been implemented in C++, leveraging the open-source library PINOCCHIO,¹ and are pipelined to be released to the public in the Pinocchio 4.0.0 release in January 2026. These algorithms are extensively benchmarked against the prevalent nonrecursive LTL algorithm for different robot topologies.

Article organization: The rest of this article is as follows. Section II reviews existing literature to situate our contributions, followed by Section III introducing the notation and necessary prior knowledge. The LCABA algorithm is derived, analyzed, and presented in an algorithmic form in Section IV, with an analogous treatment for the proxBBO algorithm in Section V. Section VI presents implementation details and benchmarking results, followed by a discussion in Section VII. Finally, Section VIII concludes this article and outlines future work.

II. RELATED WORK

CDAs can be broadly classified based on whether they solve forward, inverse, or hybrid dynamics [1]. In inverse dynamics, a robot's joint torques are computed given its configuration, velocity, acceleration, and constraint forces. In hybrid dynamics [1, Ch. 9], torques for a subset of the joints and accelerations for the remaining joints are provided as inputs, and the problem is to compute the accelerations and torques, respectively, of these joints. Due to our focus on simulation, this article focuses exclusively on *forward dynamics* CDAs, where the robot's joint accelerations and constraint forces are computed given its configuration, velocity, joint torques, and motion constraints. Forward dynamics CDAs can be further classified as recursive algorithms or joint-space algorithms (also called generalized/minimal/reduced coordinates). All the CDAs considered in this article leverage a spanning tree of the kinematic graph representing the mechanism, where links and joints correspond to nodes and edges, respectively. For completeness, we note an

alternate nonspanning tree approach, inspired by Baraff [29], that involves constructing a large and sparse system consisting of each link's Newton–Euler equations and joint constraint equations, which is then solved using a general-purpose sparse linear solver. This approach is typically not computationally competitive against spanning-tree-based algorithms [1].

CDAs can be further classified based on whether motion constraints are formulated implicitly or explicitly [1, eq. (3.11)]. The implicit formulation imposes constraints through additional equations that must be solved simultaneously with the equations of motion. The explicit constraint formulation, by computing a set of *independent* coordinates parameterizing the mechanism's constrained motion, projects the mechanism's equations of motion onto these independent coordinates and solves them.

In the rest of this section, we will review existing CDAs for unconstrained kinematic trees followed by general CDAs using explicit or implicit constraint formulations.

Unconstrained Kinematic Trees: Dynamics algorithms for even unconstrained kinematic trees can be joint-space-based or recursive. The joint-space approach computes the joint-space inertia matrix (JSIM) (typically using the efficient composite rigid body algorithm [30]). It then factorizes the JSIM efficiently using the LTL algorithm [11], which exploits branching-induced sparsity in the JSIM to reduce the factorization cost from $O(n^3)$ operations to $O(nd^2)$ operations. In contrast, the articulated body algorithm (ABA) [31], [32], [33], a recursive algorithm, has a linear computational complexity of $O(n)$ and scales better to high-DoF robot systems.

Implicit Constraint Approach: The implicit approach simultaneously solves the dynamics equations and constraint equations. External loops, arising commonly from robot–ground contacts, constitute a special class of implicit constraints that can be solved efficiently by cutting the loop at the ground link to preserve a tree structure. The joint-space LTL algorithm was extended in [12] to account for external loops. It exploits branching-induced sparsity to efficiently compute an expensive intermediate quantity known as the Delassus matrix [34], [35] (also named inverse operational space inertia matrix [36]). This results in an $O(nd^2 + m^2 d + md^2 + m^3)$ complexity algorithm. Carpentier et al. [5] introduced an extension of the LTL factorization in [11], that we call proxLTL, that leverages proximal algorithms [37] to cope with singular cases accurately. Among recursive algorithms for the implicit formulations of external loop constraints, the PV algorithm [13], [14] is a pioneering contribution with an $O(n + m^2 d + m^3)$ computational complexity. The PV algorithm was recently revisited in [15], providing a DP-based derivation for the PV algorithm by solving an equivalent discrete-time linear–quadratic regulator (LQR) problem [2]. Sathya et al. [15] further proposed two original algorithms, PV-soft and PV-early, by relaxing motion constraints with quadratic penalties and through early elimination of constraint forces, respectively. Both have a computational complexity of $O(n + m)$. However, PV-soft violates motion constraints, while PV-early is significantly challenging to implement, especially for singular cases and multi-DoF joints (e.g., when there are redundant constraints). Finally, Sathya

¹[Online]. Available: <https://github.com/stack-of-tasks/pinocchio> [17]

and Carpentier [16] applied the augmented Lagrangian method (ALM) [38], [39] on the proximal dynamics formulation [5] to derive an iterative algorithm, constrainedABA, which retains the optimal $O(n+m)$ complexity of PV-early, while handling singular cases and being significantly simpler to implement.

For both joint-space methods and recursive algorithms, handling internal loops is more challenging and represents a significant increase in computational cost. [1, Ch. 8] provides a detailed discussion on exploiting branching-induced sparsity for kinematic quantities associated with loop joints. Among joint-space methods, the proxLTL algorithm [5] extends the LTL algorithm to support internal loops and handles singular cases using proximal algorithms [40]. ProxLTL is also available as an efficient C++ implementation in the PINOCCHIO library [17]. Both LTL and proxLTL algorithms retain the $O(nd^2 + m^2d + md^2 + m^3)$ complexity for the internal loop case. Recursive algorithms for internal loops were pioneered by the BBO algorithm [18], [19], which independently rediscovered the PV algorithm and extended it to internal loops, resulting in a worst case $O(n + m^2d + m^3)$ complexity algorithm. Both these works also pioneered the early elimination of internal loop constraints, resulting in a best case computational complexity of $O(n+m)$ for local loops. The worst case complexity manifests only when all the loops are coupled with each other (e.g., external loops).

Despite its low computational complexity, the BBO algorithm has not been adopted in existing simulators, perhaps due to its implementation complexity, sensitivity to singular cases, absence of open-source implementations, and lack of benchmarking. Considering the demonstrated computational speed-ups provided by the recursive algorithms [15], [16] for external loops, there is an unexplored opportunity to exploit recursive algorithms for internal loops, provided that they can be made robust to singular cases. This is addressed in this article.

Explicit Constraint Approach: The explicit constraint formulation directly parameterizes the effective DoFs of the constrained system using independent coordinates, which are generally obtained by computing the null space of the implicit constraint formulation. This scales cubically, with the worst case manifesting for mechanisms with a large loop consisting of $\propto n$ links. Both external and internal loops are approached identically in the explicit constraint approach and are, therefore, similar in computational cost and implementation difficulty. In the joint-space approach, the JSIM and the joint torques are projected onto the independent coordinates and solved, which typically corresponds to $O(n^2n_m + nn_m^2 + n_m^3)$ operations, where n_m is the mobility of the constrained system. *Local* loops, e.g., due to a four-bar linkage, permit an efficient recursive approach through linear constraint embedding (LCE) [41], [42], [43]. LCE aggregates the links constituting a loop and defines aggregate-level generalizations of rigid-body quantities such as spatial inertia, motion, and force vectors. This aggregation transforms a kinematic graph into a tree of aggregated links, to which the articulated-body algorithm is straightforwardly adapted, resulting in a recursive algorithm. When all loops are local, LCE-ABA has a best case computational complexity of $O(n+m)$. It is particularly well-suited to handle local loops resulting from submechanisms such as gears, where the explicit

constraint formulation is readily available and the loops are local. This algorithm has recently been implemented and open-sourced in [43]. The explicit approach has limited generality, as it becomes inefficient for large loops or coupled loops, which can occur frequently, e.g., for external loops when a robot makes multipoint contact with the ground. It is, moreover, sensitive to singular cases that can occur when loop-closure constraints become linearly dependent.

III. PRELIMINARIES

This section introduces the notation used in this article, the connectivity graph representation of a mechanism, and GPLC. We also review the two equivalent QP solver approaches, the ALM and the dual proximal point algorithm (PPA), which will be used to derive LCABA and proxBBO, respectively. This is followed by a brief introduction to nonserial DP and joint-space CDAs.

A. Notations

Lowercase symbols (x), bold-faced lowercase symbols (\mathbf{x}), and uppercase symbols (X) represent scalars, vectors, and matrices, respectively. The $\mathcal{C}(\cdot)$ operator returns the cardinality of a given set. The operator $:=$ defines the left-side symbol with the right-side expression. The operators \leftarrow , \leftarrow^+ , and \leftarrow^- assigns, increments, and decrements, respectively, to the left-side variable with the right-side expression. The sets of symmetric positive-definite and symmetric positive-semidefinite matrices of size $m \times m$ are denoted as \mathbb{S}_{++}^m and \mathbb{S}_+^m , respectively. Any variable x overset with a bar and indexed with a set \mathcal{Y} as \bar{x}_y concatenates all $x_i \forall i \in \mathcal{Y}$. Any list \mathcal{X} is reversed and denoted by \mathcal{X}_r .

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, where \mathcal{Q} and $\mathcal{T}_{\mathbf{q}}\mathcal{Q}$ are the robot's configuration space and \mathcal{Q} 's tangent space at \mathbf{q} , respectively. Let $\boldsymbol{\tau} \in \mathcal{T}_{\mathbf{q}}^*\mathcal{Q} \simeq \mathbb{R}^n$ be the generalized forces exerted on the robot. If the generalized coordinates \mathbf{q} span the Euclidean space \mathbb{R}^n , as it does for robot manipulator arms, \mathbf{q} , $\boldsymbol{\nu}$, and $\dot{\boldsymbol{\nu}}$ are simply the joint positions, velocities $\dot{\mathbf{q}}$, and accelerations $\ddot{\mathbf{q}}$, respectively. When \mathcal{Q} is a nontrivial manifold, such as for floating-base robots, $\boldsymbol{\nu}$ and $\dot{\boldsymbol{\nu}}$ are not the time derivatives of \mathbf{q} , but are related to $\dot{\mathbf{q}}$ and $\ddot{\mathbf{q}}$ through the tangent space of the local parameterization of \mathcal{Q} [44]. We use Featherstone's spatial algebra [1] for rigid body quantities. The 6-D spatial velocity and acceleration of a rigid body indexed by i are $\mathbf{v}_i \in \mathbb{M}^6$ and $\mathbf{a}_i \in \mathbb{M}^6$, respectively, where \mathbb{M}^6 is the 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 of the i th link is $H_i \in \mathbb{I}^{6 \times 6} \simeq \mathbb{S}_{++}^6$ and maps \mathbb{M}^6 to \mathbb{F}^6 . \times and \times^* are the cross-product operators on the spatial motion and force vectors. Refer [1] for more details on the spatial algebra.

B. Kinematic Graph

Consider a mechanism with n_b links modeled via a connectivity graph, as shown in Fig. 2(a). The mechanism's links and joints are denoted by nodes and edges, respectively. The ground




Fig. 2. (a) Kinematic graph for an illustrative mechanism. (b) Spanning tree for the graph in Fig. 2(a).

is the zeroth link, also called the root link. In floating-base mechanisms, such as legged robots, the floating-base link is connected to the root through a “free-flying joint” that permits unrestricted relative motion. A tree is a special type of graph that is acyclic, i.e., there is a unique path between any two nodes. A subgraph of a graph consists of a subset of the original graph’s links and joints. A spanning tree is a subgraph that is a tree containing all nodes of the original graph. The original graph’s edges absent in the spanning tree are termed cut edges or cut joints. For a given joint indexed i in the spanning tree connecting two links, the link closer to the root and the other link are termed the joint’s parent and child links, respectively, and are numbered $\pi(i)$ and i , respectively. All the nonroot links and joints are numbered topologically from 1 to n_b , such that $\pi(i) < i$, and the cut-joints are numbered from $n_b + 1$ to $n_b + m_b$. Let the lists $\mathcal{S} = \{1, 2, \dots, n_b\}$ and $\mathcal{E} = \{n_b + 1, n_b + 2, \dots, n_b + m_b\}$ index the nonroot links/joints and cut-joints respectively. Let l_i^1 and l_i^2 index the two links connected by cut-joint i . Let the set $\gamma(i) = \{j \in \mathcal{S} \mid \pi(j) = i\}$ consist of the i th link’s children links. For the illustrative graph in Fig. 2(a), Fig. 2(b) shows a spanning tree that is appropriately numbered, where spanning tree joints are shown as directed edges from parents to children links, and cut-joints are shown as dashed edges.

A spanning tree defines m_b *fundamental loops*, where loop i is formed by adding cut-joint i to the spanning tree. From now on, we will refer to each fundamental loop simply as a loop. Let $\text{SL}(i)$ (supporting links) be the set of indices of links constituting loop i . Similarly, let the set $\text{LS}(i)$ (loops supported) contain indices of loops that link i supports. A loop i ’s root, defined as $\tau_i = \min(\text{SL}(i))$, is the link with the smallest index in the loop. Let the set \mathcal{R}_i denote the set of loops for which link i is the loop root. Two loops i and j are considered to be coupled if they share at least one joint.

C. Gauss’ Principle of Least Constraint

We now recall GPLC [21], [23], an optimization-based mechanics formulation, which states that a constrained rigid body’s acceleration under forces is the minimizer of the following strongly convex QP [13], [15]:

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

$$\text{subject to} \quad \mathbf{a}_i = \mathbf{a}_{\pi(i)} + S_i \dot{\boldsymbol{\nu}}_i + \mathbf{a}_{b,i}, \quad i \in \mathcal{S} \quad (1b)$$

$$K_i^1 \mathbf{a}_{l_i^1} + K_i^2 \mathbf{a}_{l_i^2} = \mathbf{k}_i, \quad i \in \mathcal{E} \quad (1c)$$

$$\mathbf{a}_0 = -\mathbf{a}_{\text{grav}} \quad (1d)$$

where \mathbf{f}_i is the resultant spatial force on link i due to external forces and bias forces ($-\mathbf{v}_i \times^* H_i \mathbf{v}_i$), $\boldsymbol{\nu}_i \in \mathbb{R}^{n_i}$. The vectors $\dot{\boldsymbol{\nu}}_i \in \mathbb{R}^{n_i}$ and $\boldsymbol{\tau}_i \in \mathbb{R}^{n_i}$ are the i th joint’s generalized velocities, generalized accelerations and the joint torques respectively, where n_i is the i th joint’s DoFs. The acceleration recurrence equation in (1b) explicitly formulates the spanning tree joint constraints. The column vectors of the matrix $S_i \in \mathbb{R}^{6 \times n_i}$ span the i th joint’s motion subspace. The i th link’s bias acceleration is $\mathbf{a}_{b,i} := \dot{S}_i \boldsymbol{\nu}_i$, which, during the common case of the joint axis being fixed relative the parent link, is

$$\mathbf{a}_{b,i} = \mathbf{v}_i \times S_i \boldsymbol{\nu}_i.$$

The cut-joint motion constraints are implicitly formulated in (1c). Here, $K_i^1, K_i^2 \in \mathbb{R}^{m_i \times 6}$ are the constraint matrices on the links indexed l_i^1 and l_i^2 , respectively, whose relative motion is constrained by the i th cut-joint, and $\mathbf{k}_i \in \mathbb{R}^{m_i}$ is desired constraint accelerations, where $m_i = 6 - n_i$ is the constraint dimension. Each row vector of K_i^1 (or K_i^2) is an element in \mathbb{F}^6 . A uniform acceleration field of minus acceleration due to gravity is applied by fixing the root node acceleration in (1d). This strategy [33] spares us from adding each link’s weight to \mathbf{f}_i , improving computational efficiency.

Joint-space GPLC formulation results in the following QP problem:

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

$$\text{subject to} \quad J_{f_c}(\mathbf{q}) \dot{\boldsymbol{\nu}} = \mathbf{a}_c(\mathbf{q}, \boldsymbol{\nu}) \quad (2b)$$

where $M(\mathbf{q}) \in \mathbb{S}_{++}^n$, $\dot{\boldsymbol{\nu}}_{\text{free}}, J_{f_c} \in \mathbb{R}^{m \times n}$, and $\mathbf{a}_c \in \mathbb{R}^m$ are the JSIM, unconstrained spanning-tree joint accelerations, constraint Jacobian, and desired constraint accelerations, respectively. The terms J_{f_c} and \mathbf{a}_c are typically computed from (1c) by substituting \mathbf{a}_i with the corresponding kinematic Jacobian equations

$$\mathbf{a}_i = J_i \dot{\boldsymbol{\nu}} + \dot{J}_i \boldsymbol{\nu}. \quad (3)$$

D. QP Solver Approaches

To solve a strongly convex QP of the form

$$\underset{\mathbf{x}}{\text{minimize}} \quad \frac{1}{2} \mathbf{x}^T Q \mathbf{x} + \mathbf{q}^T \mathbf{x} \quad (4a)$$

$$\text{subject to} \quad A \mathbf{x} = \mathbf{b} \quad (4b)$$

where $Q \in \mathbb{S}_{++}^n$, we review two equivalent approaches: the ALM [20], [38], [39] and the dual PPA [37], [40], which underpin LCABA and proxBBO, respectively. These approaches do not assume constraint linear independence (A need not have full row-rank) and are efficient by leveraging the Cholesky decomposition [45]. Alternate strategies for redundant constraints such as Tikhonov regularization or truncated SVD [45] either undesirably bias the optimal \mathbf{x} toward the origin or incur high computational costs.

The QP's Lagrangian function [46] is defined as

$$\mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}_x) := \frac{1}{2} \mathbf{x}^T Q \mathbf{x} + \mathbf{q}^T \mathbf{x} + \boldsymbol{\lambda}_x^T (A \mathbf{x} - \mathbf{b}) \quad (5)$$

where $\boldsymbol{\lambda}_x$ are the QP's dual variables.

The ALM augments the Lagrangian function with a quadratic penalty on constraint violation, defining the augmented Lagrangian function (ALF)

$$\mathcal{L}^A(\mathbf{x}, \boldsymbol{\lambda}_x) := \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}_x) + \frac{\mu}{2} \|A \mathbf{x} - \mathbf{b}\|^2 \quad (6)$$

and alternately minimizes and maximizes the ALF w.r.t the primal and dual variables in an iterative fashion

$$\mathbf{x}^{k+1} = \left(Q + \frac{\mu}{2} A^T A \right)^{-1} (-\mathbf{q} - A^T \boldsymbol{\lambda}_x^k + \mu A^T \mathbf{b}) \quad (7a)$$

$$\boldsymbol{\lambda}_x^{k+1} = \boldsymbol{\lambda}_x^k + \mu (A \mathbf{x}^{k+1} - \mathbf{b}) \quad (7b)$$

until a termination criterion is met.

Dual PPA: The QP's dual function is

$$g(\boldsymbol{\lambda}_x) := \min_{\mathbf{x}} \mathcal{L}(\mathbf{x}, \boldsymbol{\lambda}_x) \quad (8)$$

where minimizing \mathcal{L} w.r.t \mathbf{x} is always possible since Q is positive definite, yielding

$$g(\boldsymbol{\lambda}_x) = -\frac{1}{2} \boldsymbol{\lambda}_x^T (A Q^{-1} A^T) \boldsymbol{\lambda}_x - (A Q^{-1} \mathbf{q} + \mathbf{b})^T \boldsymbol{\lambda}_x. \quad (9)$$

Note that the dual Hessian $A Q^{-1} A^T \in \mathbb{S}_+^m$ makes the dual function concave, but not necessarily strongly concave since A need not have full row-rank. The QP can be solved by maximizing the dual function [note that optimal \mathbf{x}^* can be recovered from optimal $\boldsymbol{\lambda}^*$ using (7a)].

The proximal operator for a convex function $f: \mathbb{R}^n \rightarrow \mathbb{R}$ is defined as

$$\text{prox}_{\mu, f}(\mathbf{x}^k) := \underset{\mathbf{x}}{\text{argmin}} \left\{ f(\mathbf{x}) + \frac{1}{2\mu} \|\mathbf{x} - \mathbf{x}^k\|^2 \right\} \quad (10)$$

where $\mu \in (0, \infty)$. The PPA minimizes the function f by performing fixed-point iterations using the proximal operator

$$\mathbf{x}^{k+1} := \text{prox}_{\mu, f}(\mathbf{x}^k) \quad (11)$$

until a termination criterion is met. The dual function $g(\boldsymbol{\lambda}_x)$ needs to be maximized to solve the optimization problem. This is achieved by minimizing $-g(\boldsymbol{\lambda}_x)$ using the PPA

$$\boldsymbol{\lambda}_x^{k+1} := \text{prox}_{\mu, -g}(\boldsymbol{\lambda}_x^k).$$

The shifted regularization term added in (10) makes each inner problem solved by the proximal operator strongly convex, even if the dual function is not strongly concave. For convex problems, PPA enjoys theoretical convergence guarantees [40, Sect. 4.1] for any choice of $\mu > 0$ under very mild conditions, assuming that the minimizer exists. However, the specific choice of μ , which can be interpreted as a step-size parameter in PPA, can significantly affect the convergence rate of both the ALM and the PPA methods. Proximal algorithms [40] have been found to be particularly effective in robotics [47], [48]. They are easy to implement, support warm-starting, can handle infeasible initial guesses and even infeasible problems (by returning the closest feasible solution [49]), and typically require a few iterations

(each of which is efficient) to converge to desired levels of accuracy for robot dynamics problems [5], [7], [16].

E. Joint-Space Algorithms

The joint-space GPLC problem formulation in (2) can also be solved using dual PPA or the ALM, resulting in the proxLTL [5] and proxLTLs algorithms, respectively. These two algorithms are the joint-space counterparts of proxBBO and LCABA, respectively, and will be reviewed in this section.

ProxLTL Algorithm: The first iteration of dual PPA is equivalent to solving the following Karush-Kuhn-Tucker (KKT) system [5]:

$$\begin{bmatrix} -\frac{1}{\mu} I_{m \times m} & J_{\mathbf{f}_c} \\ J_{\mathbf{f}_c}^T & M \end{bmatrix} \begin{bmatrix} \boldsymbol{\lambda}^1 \\ \dot{\mathbf{v}}^1 \end{bmatrix} = \begin{bmatrix} \mathbf{a}_c - \frac{1}{\mu} \boldsymbol{\lambda}^0 \\ M \dot{\mathbf{v}}_{\text{free}} \end{bmatrix} \quad (12)$$

where both M and $J_{\mathbf{f}_c}$ have the spanning-tree-induced sparsity pattern [1, Sect. 8.9], [11]. This sparsity was exploited in [5] to efficiently factorize the KKT system above using the UDU^T Cholesky decomposition (reverse of the typical LDL^T ordering). This approach turns out to compute the dual function as an intermediate quantity and the upper left-block of the U matrix corresponds to the Cholesky factor of the damped Delassus matrix ($\Lambda_\mu := J_{\mathbf{f}_c}^T M^{-1} J_{\mathbf{f}_c} + \frac{1}{\mu} I_{m \times m}$). The Cholesky factor is reused to efficiently compute the subsequent dual PPA iterations

$$\Lambda_\mu \boldsymbol{\lambda}^{k+1} = J_{\mathbf{f}_c} \dot{\mathbf{v}}_{\text{free}} - \mathbf{a}_c + \frac{1}{\mu} \boldsymbol{\lambda}^k. \quad (13)$$

ProxLTLs Algorithm: Applying ALM on the joint-space GPLC problem yields the following updated equations:

$$\dot{\mathbf{v}}^{k+1} = M_\mu^{-1} (\dot{\mathbf{v}}_{\text{free}} + J_{\mathbf{f}_c}^T [\mu \mathbf{a}_c - \boldsymbol{\lambda}^k]) \quad (14a)$$

$$\boldsymbol{\lambda}^{k+1} = \boldsymbol{\lambda}^k + \mu (J_{\mathbf{f}_c} \dot{\mathbf{v}} - \mathbf{a}_c) \quad (14b)$$

where

$$M_\mu = M + J_{\mathbf{f}_c}^T (\mu I_{m \times m}), \quad J_{\mathbf{f}_c} \in \mathbb{S}_{++}^n. \quad (15)$$

The proxLTLs was mentioned in [5], but not implemented, while it has been first considered in [16] for the restricted case of kinematic trees with external loops. In this restricted case, M_μ 's sparsity pattern is identical to that of M . However, this property no longer holds for mechanisms with internal loops, where M_μ contains a dense block corresponding to all the joints comprising a loop. To see why, consider a row of the $J_{\mathbf{f}_c}$ matrix corresponding to a cut-joint. All the columns of this row corresponding to the joints in the loop will be nonzero in general, and a μ -weighted outer product of this row with itself yields a dense block in M_μ . This loss of sparsity significantly complicates the implementation of proxLTLs and also reduces its computational efficiency; therefore, this algorithm is not considered in this article.

F. Nonserial Dynamic Programming

Nonserial DP [25, Ch. 10], [26], [27] is a straightforward generalization of the traditional serial DP [24] to nonchain graphs. Suppose that the function to be optimized is given by

$$f(a, b, c, d, e) = p(a, d) + q(b, c) + r(b, e) + s(d, e)$$




Fig. 3. (a) Visualizing the structure of the objective function minimized by nonserial DP. (b) Graph structure after optimizing over nodes a and b .

whose structure can be visualized in Fig. 3(a), with the variables depicted as nodes and the functions depicted as edges. A function depending on multiple variables is depicted by a clique involving the corresponding nodes and vice versa. DP successively eliminates variables by optimizing over them to compute the optimal “cost-to-go” functions that depend on the remaining variables until all variables are eliminated.

Let $\mathcal{N}(g)$ be the set of all the nodes neighboring g in the graph. Let $\hat{f}(g, \mathcal{N}(g))$ be the sum of functions depending on g , and $\tilde{f}(\mathcal{N}(g))$ be the sum of functions depending only on variables in $\mathcal{N}(g)$. Optimizing over the variable g gives the function

$$\hat{f}^*(\mathcal{N}(g)) = \min_g \left\{ \hat{f}(g, \mathcal{N}(g)) \right\} \quad (16)$$

which, in general, depends on all elements in $\mathcal{N}(g)$ and is thereby represented graphically as a clique comprising the nodes in $\mathcal{N}(g)$. The graph’s connectivity is updated by adding edges between any unconnected nodes in $\mathcal{N}(g)$. Then, the function $\tilde{f}(\mathcal{N}(g))$ is updated with the terms obtained from optimizing over g as follows:

$$\tilde{f}(\mathcal{N}(g)) \stackrel{\pm}{\leftarrow} \hat{f}^*(\mathcal{N}(g)). \quad (17)$$

This step is repeated until all the variables are eliminated.

As an illustrative example, suppose that DP optimizes over variables in the order a, b, c, d, e for the case depicted in Fig. 3(a). After eliminating a , we get

$$\tilde{f}(d) \leftarrow \min_a p(a, d) \quad (18)$$

where the connectivity structure of the graph does not change since a has only one neighbor. At the next step, minimizing over b results in the following update:

$$\tilde{f}(c, e) \leftarrow \min_b \{r(b, c) + s(b, e)\} \quad (19)$$

resulting in a new edge being added between the nodes c and e . This new graph connectivity is plotted in Fig. 3(b). Similarly, the function is sequentially optimized over c, d , and e to get a constant function, which is the optimum.

DP is generally intractable due to the curse of dimensionality. This is exacerbated in nonserial DP, where the cost of representing and optimizing the cost-to-go function grows exponentially with the number of neighbors of the variable being eliminated. However, just as LQR is a tractable special case for serial DP, solving GPLC over graphs also turns out to be a tractable problem because the cost at each stage can be efficiently parameterized as a quadratic form. Finally, it is worth noting that the traditional serial DP is a special case of the nonserial DP algorithm, where each leaf node of the graph has only one

neighbor, due to which eliminating that variable does not modify the graph structure.

IV. LOOP-CONSTRAINEDABA

This section derives the LCABA algorithm, presents it in an algorithmic form, and analyzes its computational complexity.

A. LCABA Derivation

The derivation applies the ALM to the GPLC problem from (1) and solves each inner primal problem [see (7a)] using DP. We will follow the notation and style of [15] for this DP-based derivation.

The ALF associated with GPLC (1) is given by

$$\begin{aligned} \mathcal{L}^A(\dot{\nu}, \mathbf{a}, \boldsymbol{\lambda}) = & \sum_{i \in \mathcal{S}} \left\{ \frac{1}{2} \mathbf{a}_i^T H_i \mathbf{a}_i - \mathbf{f}_i^T \mathbf{a}_i - \boldsymbol{\tau}_i^T \dot{\nu}_i \right\} + \\ & \sum_{i \in \mathcal{E}} \left\{ \boldsymbol{\lambda}_i^T \left(K_i^1 \mathbf{a}_{l_i^1} + K_i^2 \mathbf{a}_{l_i^2} - \mathbf{k}_i \right) + \right. \\ & \left. \frac{\mu}{2} \| K_i^1 \mathbf{a}_{l_i^1} + K_i^2 \mathbf{a}_{l_i^2} - \mathbf{k}_i \|^2 \right\} \end{aligned} \quad (20)$$

where the spanning tree joint constraints [see (1b)] are excluded, as they are eliminated by substitution, similar to single-shooting transcription [2] in optimal control.

The ALM’s inner primal problem is defined by

$$\dot{\nu}^{k+1}, \mathbf{a}^{k+1} = \underset{\dot{\nu}, \mathbf{a}}{\operatorname{argmin}} \mathcal{L}^A(\dot{\nu}, \mathbf{a}, \boldsymbol{\lambda}^k). \quad (21)$$

Rearranging the constraint terms for the i th cut-joint in \mathcal{L}^A yields the quadratic form:

$$\begin{aligned} & \frac{1}{2} \begin{bmatrix} \mathbf{a}_{l_i^2} \\ \mathbf{a}_{l_i^1} \end{bmatrix}^T \begin{bmatrix} \mu K_i^{2T} K_i^2 & \mu K_i^{2T} K_i^1 \\ \mu K_i^{1T} K_i^2 & \mu K_i^{1T} K_i^1 \end{bmatrix} \begin{bmatrix} \mathbf{a}_{l_i^2} \\ \mathbf{a}_{l_i^1} \end{bmatrix} - \\ & \begin{bmatrix} K_i^{2T} (\mu \mathbf{k}_i - \boldsymbol{\lambda}_i^k) \\ K_i^{1T} (\mu \mathbf{k}_i - \boldsymbol{\lambda}_i^k) \end{bmatrix}^T \begin{bmatrix} \mathbf{a}_{l_i^2} \\ \mathbf{a}_{l_i^1} \end{bmatrix}. \end{aligned} \quad (22)$$

The off-diagonal blocks reveal an inertial coupling term $\mu K_i^{1T} K_i^2$ between links connected by a cut-joint. To account for this in DP, we hypothesize the following quadratic form for the optimal cost-to-go function for the intermediate DP subproblem at a spanning tree leaf link indexed i :

$$\begin{aligned} V_i^*(\mathbf{a}_i, \bar{\mathbf{a}}_{\mathcal{N}_i}) := & \frac{1}{2} \begin{bmatrix} \mathbf{a}_i \\ \bar{\mathbf{a}}_{\mathcal{N}_i} \end{bmatrix}^T \begin{bmatrix} H_{i,i} & \bar{H}_{i,\mathcal{N}_i}^T \\ \bar{H}_{i,\mathcal{N}_i} & \bar{H}_{\mathcal{N}_i,\mathcal{N}_i} \end{bmatrix} \begin{bmatrix} \mathbf{a}_i \\ \bar{\mathbf{a}}_{\mathcal{N}_i} \end{bmatrix} \\ & - \begin{bmatrix} \mathbf{f}_i \\ \bar{\mathbf{f}}_{\mathcal{N}_i} \end{bmatrix}^T \begin{bmatrix} \mathbf{a}_i \\ \bar{\mathbf{a}}_{\mathcal{N}_i} \end{bmatrix} \end{aligned} \quad (23)$$

where \mathcal{N}_i denotes the set of neighboring link indices for the i th link. Note that initially, before any eliminations, \mathcal{N}_i consists of the cut-joint neighbors of the i th link, represented by the red dashed edges in Fig. 2(b). The terms $\bar{\mathbf{a}}_{\mathcal{N}_i}$, $\bar{H}_{i,\mathcal{N}_i}$, $\bar{H}_{\mathcal{N}_i,\mathcal{N}_i}$, and $\bar{\mathbf{f}}_{\mathcal{N}_i}$ concatenate the accelerations, coupled inertias, and forces

of the links indexed in \mathcal{N}_i . For all $j, k \in \mathcal{N}_i$

$$\bar{\mathbf{a}}_{\mathcal{N}_i} = \left[\dots \quad \mathbf{a}_j^T \quad \dots \right]^T, \quad \bar{\mathbf{f}}_{\mathcal{N}_i} = \left[\dots \quad \mathbf{f}_j^T \quad \dots \right]^T$$

$$\bar{H}_{i,\mathcal{N}_i} = \begin{bmatrix} \vdots \\ H_{j,i} \\ \vdots \end{bmatrix}, \quad \bar{H}_{\mathcal{N}_i,\mathcal{N}_i} = \begin{bmatrix} H_{k,k} & \dots & H_{k,j} & \dots \\ \vdots & \ddots & \vdots & \ddots \\ H_{k,j}^T & \dots & H_{j,j} & \dots \\ \vdots & \ddots & \vdots & \ddots \end{bmatrix}. \quad (24)$$

The quadratic form coefficients are initialized by iterating over the cut-joint constraints. For a cut-joint j , the updates are

$$\begin{bmatrix} H_{l_j^1, l_j^1} \\ H_{l_j^2, l_j^2} \\ H_{l_j^1, l_j^2} \end{bmatrix} \stackrel{\pm}{\leftarrow} \mu \begin{bmatrix} K_j^{1T} K_j^1 \\ K_j^{2T} K_j^2 \\ K_j^{1T} K_j^2 \end{bmatrix}, \quad \begin{bmatrix} \mathbf{f}_{l_j^1} \\ \mathbf{f}_{l_j^2} \end{bmatrix} \stackrel{\pm}{\leftarrow} \begin{bmatrix} K_j^{1T}(\mu \mathbf{k}_j - \boldsymbol{\lambda}_j^k) \\ K_j^{2T}(\mu \mathbf{k}_j - \boldsymbol{\lambda}_j^k) \end{bmatrix}. \quad (25)$$

The DP elimination process begins by selecting a leaf link indexed i and eliminating its acceleration \mathbf{a}_i and joint acceleration $\dot{\boldsymbol{\nu}}_i$. In addition to its cut-joint neighbors \mathcal{N}_i , link i is connected to its parent link indexed $\pi(i)$ via the acceleration recurrence relation. Thus, upon eliminating the i th link, the resulting cost-to-go function depends on links in $\{\pi(i)\} \cup \mathcal{N}_i$, by the recurrence relation

$$V^*(\mathbf{a}_{\pi(i)}, \bar{\mathbf{a}}_{\mathcal{N}_i}) \leftarrow \frac{1}{2} \mathbf{a}_{\pi(i)}^T H_{\pi(i), \pi(i)} \mathbf{a}_{\pi(i)} - \mathbf{f}_{\pi(i)}^T \mathbf{a}_{\pi(i)} + \bar{\mathbf{a}}_{\mathcal{N}_i}^T \bar{H}_{\pi(i), \mathcal{N}_i} \mathbf{a}_{\pi(i)} + \min_{\dot{\boldsymbol{\nu}}_i, \mathbf{a}_i} \left\{ V_i^*(\mathbf{a}_i, \bar{\mathbf{a}}_{\mathcal{N}_i}) - \boldsymbol{\tau}_i^T \dot{\boldsymbol{\nu}}_i \right\}. \quad (26)$$

Substituting \mathbf{a}_i using the acceleration recurrence relation (1b) transforms the minimization above to

$$\min_{\dot{\boldsymbol{\nu}}_i} \left\{ V_i^*(\mathbf{a}_{\pi(i)} + S_i \dot{\boldsymbol{\nu}}_i + \mathbf{a}_{b,i}, \bar{\mathbf{a}}_{\mathcal{N}_i}) - \boldsymbol{\tau}_i^T \dot{\boldsymbol{\nu}}_i \right\}. \quad (27)$$

Expanding V_i^* using (23) and collecting terms with $\dot{\boldsymbol{\nu}}_i$ yields the unconstrained QP

$$\begin{aligned} \text{minimize}_{\dot{\boldsymbol{\nu}}_i} \quad & \frac{1}{2} \dot{\boldsymbol{\nu}}_i^T D_i \dot{\boldsymbol{\nu}}_i - \left[\boldsymbol{\tau}_i + S_i^T \left(\mathbf{f}_i - \bar{H}_{i,\mathcal{N}_i}^T \bar{\mathbf{a}}_{\mathcal{N}_i} \right. \right. \\ & \left. \left. - H_{i,i} (\mathbf{a}_{\pi(i)} + \mathbf{a}_{b,i}) \right) \right]^T \dot{\boldsymbol{\nu}}_i \end{aligned} \quad (28)$$

where $D_i := S_i^T H_{i,i} S_i$ is the i th link's projected inertia on the joint subspace, with $D_i \in \mathbb{S}_{++}^{n_i}$ [19]. Solving this QP yields the optimal joint accelerations

$$\dot{\boldsymbol{\nu}}_i^* = D_i^{-1} \left[\boldsymbol{\tau}_i + S_i^T \left(\mathbf{f}_i - \bar{H}_{i,\mathcal{N}_i}^T \bar{\mathbf{a}}_{\mathcal{N}_i} - H_{i,i} (\mathbf{a}_{\pi(i)} + \mathbf{a}_{b,i}) \right) \right]. \quad (29)$$

Substituting $\dot{\boldsymbol{\nu}}_i^*$ into (26) results in the updates

$$H_{\pi(i), \pi(i)} \stackrel{\pm}{\leftarrow} P_i H_{i,i} \quad (30a)$$

$$\mathbf{f}_{\pi(i)} \stackrel{\pm}{\leftarrow} P_i (\mathbf{f}_i - H_{i,i} \mathbf{a}_{b,i}) - H_{i,i} S_i D_i^{-1} \boldsymbol{\tau}_i \quad (30b)$$

$$\bar{H}_{\pi(i), \mathcal{N}_i} \stackrel{\pm}{\leftarrow} \bar{H}_{i,\mathcal{N}_i} P_i^T \quad (30c)$$

$$\bar{H}_{\mathcal{N}_i, \mathcal{N}_i} \stackrel{\pm}{\leftarrow} \bar{H}_{i,\mathcal{N}_i} S_i D_i^{-1} S_i^T \bar{H}_{i,\mathcal{N}_i}^T \quad (30d)$$

$$\bar{\mathbf{f}}_{\mathcal{N}_i} \stackrel{\pm}{\leftarrow} \bar{H}_{i,\mathcal{N}_i} \left[P_i^T \mathbf{a}_{b,i} + S_i D_i^{-1} (\boldsymbol{\tau}_i + S_i^T \mathbf{f}_i) \right] \quad (30e)$$

where

$$P_i = I_{6 \times 6} - H_{i,i} S_i D_i^{-1} S_i^T.$$

After eliminating the i th link, it is removed from the neighbor lists \mathcal{N}_j for all $j \in \mathcal{N}_i$

$$\mathcal{N}_j \leftarrow \mathcal{N}_j - \{i\}, \quad \text{for } j \in \mathcal{N}_i. \quad (31)$$

Eliminating link i introduces mutual coupling between all links in $\mathcal{N}_i \cup \{\pi(i)\}$ [see (30c) and (30d)], requiring neighbor list updates

$$\mathcal{N}_j \leftarrow \mathcal{N}_j \cup ((\mathcal{N}_i \cup \{\pi(i)\}) - \{j\}), \quad \text{for } j \in \mathcal{N}_i \cup \{\pi(i)\}. \quad (32)$$

Undefined $H_{j,k}$ terms are initialized to $0_{6 \times 6}$ before executing (30c) and (30d).

The process repeats for the next selected leaf link in the spanning tree until all are eliminated. Note that a link k becomes a leaf once all its children $\gamma(k)$ are eliminated. The k th link's neighbors \mathcal{N}_k at this DP stage will have been recursively updated using (31) and (32) whenever a neighbor or child is eliminated. Furthermore, V_k^* retains the quadratic form hypothesized in (23), updated via (30). Therefore, the quadratic form hypothesis in (23) can be shown to hold inductively throughout the DP process.

Once all links are eliminated, the optimal joint accelerations are calculated in reverse of the elimination order using (29), solving the ALM's inner primal problem. Subsequently, dual variables are straightforwardly updated [see (7b)] using the constraint violations

$$\boldsymbol{\lambda}_i^{k+1} = \boldsymbol{\lambda}_i^k + \mu \left(K_i^1 \mathbf{a}_{l_i^1}^{k+1} + K_i^2 \mathbf{a}_{l_i^2}^{k+1} - \mathbf{k}_i \right). \quad (33)$$

When $\pi(i) \in \mathcal{N}_i$, a special case occurs. Equation (30c) evaluates the off-diagonal matrix blocks $\bar{H}_{\pi(i), \mathcal{N}_i}$, with $\pi(i) \in \mathcal{N}_i$. Since the Hessian is symmetric, the symmetric counterpart of the update in (30c) must be added to $H_{\pi(i), \pi(i)}$

$$H_{\pi(i), \pi(i)} \leftarrow H_{\pi(i), \pi(i)} + (H_{i,\pi(i)} P_i^T)^T. \quad (34)$$

Efficiency of subsequent primal iterations: Subsequent ALM iterations' inner problems only need to recompute force and acceleration terms affected by updated Lagrange multipliers. Since inertia terms H are constant, a reduced recursion is devised, which avoids expensive matrix operations in (25) and (30) and only performs operations needed for force and acceleration updates. This is akin to reusing a factorized linear system and will be detailed later in this section.

Elimination ordering: A spanning tree may have multiple leaf links, allowing various elimination orders, with differing computational costs. The cost of eliminating a link scales quadratically with the link's neighbor count ($|\mathcal{N}_i|$) and introduces coupling among numerous links. Finding an optimal elimination order for variable elimination is generally NP-hard [50]. However, effective greedy heuristics exist, such as minimum degree [51], nested dissection [52], and the Cuthill-McKee algorithm [53]. LCABA adopts the minimum degree heuristic: eliminating the




Fig. 4. (a)–(e) Graphical illustration of LCABA’s elimination steps.

leaf link with the fewest neighbors, breaking ties randomly. Let m_c denote the maximum neighbor count encountered.

Constraints with respect to ground: The constraint model in (1) also supports external contact constraints (e.g., quadruped foot contacts) without internal loops. Assume that l_i^1 th link is the ground link without loss of generality. Since ground acceleration $\mathbf{a}_0 = -\mathbf{a}_{\text{grav}}$ is constant and known, it is directly substituted. As \mathbf{a}_0 is not a decision variable, only $H_{l_i^1, l_i^2}$ and $\mathbf{f}_{l_i^2}$ need to be updated in (25). To account for $H_{l_i^1, l_i^2}$ coupling, the force update needs to be modified as follows:

$$\mathbf{f}_{l_i^2} \stackrel{\pm}{\leftarrow} K_i^{2T} (\mu \mathbf{k}_i - \mu K_i^1 \mathbf{a}_0 - \boldsymbol{\lambda}_i^k). \quad (35)$$

Since inertial coupling is absorbed into the force update, the ground link and l_i^2 th link are not neighbors, substantially simplifying the DP. Remaining steps proceed as derived. If all loops are external, LCABA reduces to constrainedABA [16]. Thus, LCABA generalizes constrainedABA to support internal loops, while retaining constrainedABA’s efficiency for external loops.

Remark 1: Although the constraint model (1) constrains relative motion between two links, the presented algorithm readily supports multilink constraints of the form

$$\sum_j \{K_i^j \mathbf{a}_{l_i^j}\} = \mathbf{k}_i \quad (36)$$

with trivial modifications to the quadratic form updates in (25) to compute coupling among all the links involved in this multilink constraint.

However, efficiency drops if many links are involved, as cost scales quadratically with neighbor count. Thus, LCABA is not recommended for global constraints involving $\propto n$ links (e.g., center-of-mass (CoM) constraints).

B. LCABA Illustration

Fig. 4 illustrates LCABA elimination for an 11-link mechanism with two cut-joints (links 6&8 and 8&11). Solid edges depict spanning-tree constraints; dashed red edges depict cut-joint coupling. Link 6 is eliminated first (minimum degree

heuristic), having the fewest neighbors among the leaf-links 6, 8, and 11 (though not strictly since the eighth link has the same number of neighbors). This elimination couples links 5 and 8. Subsequently, leaves 5 and 4 are eliminated, coupling links 3 and 8 [see Fig. 4(b)]. Next, leaves 9–11 are eliminated [see Fig. 4(c)]. Since links 3 and 8 were already coupled, eliminating link 9 only modifies this existing coupling. Links 8 and 7 are then eliminated [see Fig. 4(d)]. Since the remaining links have no neighbors, LCABA proceeds identically to ABA. If link 8 were eliminated first, it would couple links 6 and 11 [see Fig. 4(e)]. This increases cost, as each leaf link would then have two neighbors (note that neighbors are links with inertial coupling which may not include the parent). This demonstrates the benefit of the minimum degree heuristic.

C. LCABA Algorithm

This section presents LCABA in an algorithmic form. Let $\mathcal{S}^{\mathcal{E}}$ be the order in which spanning-tree links are eliminated. The first ALM iteration is presented in Algorithm 1, followed by the more efficient reduced-sweep for subsequent ALM iterations in Algorithm 2, before presenting the whole LCABA in Algorithm 3. The terms in parentheses, such as $(U_i D_i^{-1})$, are stored in variables to avoid their re-computation. Since blocks in coupling matrices satisfy $H_{i,j} = H_{j,i}^T$, only $H_{i,j}$, for $i < j$, is computed and stored. However, this aspect is avoided in the algorithm for clarity. The reduced sweeps depicted in Algorithm 2 compute only the delta changes in forces and accelerations due to constraint force updates in an ALM iteration.

1) *Forward Kinematic Sweep:* The link velocities and the bias accelerations are computed as

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

Link inertias and the link wrenches due to external forces and bias forces are initialized as

$$H_{i,i} \leftarrow H_i, \quad \mathbf{f}_i \leftarrow \mathbf{f}_i^{\text{ext}} - \mathbf{v}_i \times^* H_i \mathbf{v}_i.$$

2) *Resultant Torques, Apparent Inertia, Inverse Inertia, and Projection Matrix:* They are computed as

$$\mathbf{u}_i = \boldsymbol{\tau}_i + S_i^T \mathbf{f}_i, \quad U_i = H_{i,i} S_i, \quad D_i = S_i^T U_i, \quad D_i^{-1};$$

$$P_i = I_{6 \times 6} - U_i (D_i^{-1} S_i^T).$$

3) *Backward Inertia and Forces:* The propagated inertias and forces to the parent link are updated as

$$H_{\pi(i),i}^a = H_{i,i} - (U_i D_i^{-1}) U_i^T; \quad H_{\pi(i),\pi(i)} \stackrel{\pm}{\leftarrow} H_{i,i}^a$$

$$\mathbf{f}_{\pi(i)} \stackrel{\pm}{\leftarrow} \mathbf{f}_i - H_{i,i}^a \mathbf{a}_{b,i} - (U_i D_i^{-1}) \mathbf{u}_i.$$

4) *Second Forward Sweep:* The resulting joint and link accelerations at the $(k+1)$ th iteration are computed as

$$\boldsymbol{\nu}_i^{(k+1)} = D_i^{-1} \mathbf{u}_i - (U_i D_i^{-1})^T (\mathbf{a}_{\pi(i)} + \mathbf{a}_{b,i})$$

$$\mathbf{a}_i^{k+1} = \mathbf{a}_{\pi(i)}^{k+1} + S_i \boldsymbol{\nu}_i^{(k+1)} + \mathbf{a}_{b,i}.$$

D. LCABA Complexity Analysis

The worst case computational complexity of LCABA is now analyzed, starting with the three-sweep Algorithm 1. The

Algorithm 1: LCABA Three-Sweep Algorithm.

Require: $\mathbf{q}, \nu, \tau, S_i s, H_i s, K_i s, \mathbf{k}_i s, \mathbf{f}_i^{\text{ext}}, \mathcal{S}^\mathcal{E}, \mu, \lambda^0, \mathcal{E}$

- 1: **for** i in \mathcal{S} \triangleright First forward sweep
- 2: FK and initialize inertias and forces (see Section IV-C1).
- 3: $\mathcal{N}_i \leftarrow \{\}$
- 4: **for** j in \mathcal{E} \triangleright Process constraints: update inertia, forces, and neighbors
- 5: $H_{l_j^1, l_j^1} \stackrel{\pm}{\leftarrow} \mu K_j^1 T K_j^1, H_{l_j^2, l_j^2} \stackrel{\pm}{\leftarrow} \mu K_j^2 T K_j^2;$
- 6: **if** $H_{l_j^1, l_j^2}$ is undefined \triangleright Create new edges
- 7: $H_{l_j^1, l_j^2} \leftarrow 0_{6 \times 6};$
- 8: $\mathcal{N}_{l_j^1} \leftarrow \mathcal{N}_{l_j^1} \cup \{l_j^2\}; \mathcal{N}_{l_j^2} \leftarrow \mathcal{N}_{l_j^2} \cup \{l_j^1\}$
- 9: $H_{l_j^1, l_j^2} \stackrel{\pm}{\leftarrow} \mu K_j^1 T K_j^2, \mathbf{f}_{l_j^1} \stackrel{\pm}{\leftarrow} K_j^1 T (\mu \mathbf{k}_j - \lambda_j^k),$
- 10: $\mathbf{f}_{l_j^2} \stackrel{\pm}{\leftarrow} K_j^2 T (\mu \mathbf{k}_j - \lambda_j^k)$
- 11: **for** i in $\mathcal{S}^\mathcal{E}$ \triangleright Backward sweep
- 12: Compute resultant torques, apparent inertia, inverse inertia, and projection matrix (see Section IV-C2).
- 13: **if** $\mathcal{C}(\mathcal{N}_i) > 0$ \triangleright Update terms after i th link's elimination
- 14: **for** j in $\mathcal{N}_i \cup \{\pi(i)\}$ \triangleright Update connections
- 15: $\mathcal{N}_j \leftarrow \mathcal{N}_j - \{i\};$ \triangleright Remove the i th link from neighbors
- 16: **for** k in $(\mathcal{N}_i \cup \{\pi(i)\}) - \{j\}$
- 17: $H_{k,j} \leftarrow 0_{6 \times 6};$ **if** $H_{k,j}$ undefined
- 18: $\mathcal{N}_j \leftarrow \mathcal{N}_j \cup \{k\}$ \triangleright Update \mathcal{N}_j
- 19: Compute (30c) to (30e) \triangleright Update coupling inertias and forces.
- 20: **if** $\pi(i) \in \mathcal{N}_i$ \triangleright Resymmetrize if parent is a neighbor
- 21: $H_{\pi(i), \pi(i)} \stackrel{\pm}{\leftarrow} (H_{i, \pi(i)} P_i^T)^T;$
- 22: **if** $\pi(i) > 0$
- 23: Backward inertia and forces propagation (see Section IV-C3)
- 24: **for** i in $\mathcal{S}_r^\mathcal{E}$ \triangleright Second forward sweep (Rollout)
- 25: $\mathbf{u}_i \leftarrow (\bar{H}_{i, \mathcal{N}_i} S_i)^T \bar{\mathbf{a}}_{\mathcal{N}_i};$
- 26: Compute joint and link accelerations (see Section IV-C4)

first forward sweep from line 1 to line 10 requires $O(n + m)$ operations. In the second forward sweep, line 25 requires $O(|\mathcal{N}_i|)$ operations per joint, while the next line requires a fixed number of operations, bringing the second forward sweep's complexity to $O(n + m_c n)$. Note that m_c the maximum neighbor count among all links, given by $\max_{i \in \mathcal{S}} \mathcal{C}(\mathcal{N}_i)$. Similarly to ABA and constrainedABA, the backward sweep is the most computationally expensive part, with line 19 requiring $O(\mathcal{C}(\mathcal{N}_i)^2)$ operations at each joint to update the inertial coupling by computing (30d). This brings the total computational complexity of the three-sweep algorithm to $O(n + m + m_c^2 n)$. Note that loops are local and not coupled in many practical cases, e.g., the four-bar submechanisms of the Digit robot in Fig. 1(b). Even when there is coupling among loops, typically only a few loops ($m_c \sim 3$) participate in such coupling, bringing the

Algorithm 2: LCABA Reduced Two-Sweep Algorithm.

Require: $\Delta \mathbf{f}_i s, K_i s, (\bar{H}_{\mathcal{N}_i} S_i) s, U_i s, D_i s, S_i s, \mathcal{S}^\mathcal{E}$

- 1: **for** i in $\mathcal{S}^\mathcal{E}$ \triangleright Reduced backward sweep
- 2: $\Delta \mathbf{u}_i \leftarrow S_i^T \Delta \mathbf{f}_i$ \triangleright Resultant torque changes
- 3: **if** $\mathcal{C}(\mathcal{N}_i) > 0$ \triangleright Update forces on neighbors due to $\Delta \mathbf{u}_i$
- 4: $\Delta \bar{\mathbf{f}}_{\mathcal{N}_i} \stackrel{\pm}{\leftarrow} (\bar{H}_{i, \mathcal{N}_i} S_i)(D_i^{-1} \Delta \mathbf{u}_i)$
- 5: **if** $\pi(i) > 0$ \triangleright Update forces on parent due to $\Delta \mathbf{u}_i$
- 6: $\Delta \mathbf{f}_{\pi(i)} \stackrel{\pm}{\leftarrow} \Delta \mathbf{f}_i - U_i (D_i^{-1} \Delta \mathbf{u}_i)$
- 7: $\Delta \mathbf{a}_0 \leftarrow 0_{6 \times 1}$
- 8: **for** i in $\mathcal{S}_r^\mathcal{E}$ \triangleright Reduced forward sweep
- 9: **if** $\mathcal{C}(\mathcal{N}_i) > 0$ \triangleright Update $\Delta \mathbf{u}_i$ due to neighbor accelerations
- 10: $\Delta \mathbf{u}_i \leftarrow (\bar{H}_{i, \mathcal{N}_i} S_i)^T \Delta \bar{\mathbf{a}}_{\mathcal{N}_i}$
- 11: $\Delta \dot{\nu}_i^{(k+1)} = D_i^{-1} \Delta \mathbf{u}_i - (U_i D_i^{-1})^T \Delta \mathbf{a}_{\pi(i)};$
- 12: $\Delta \mathbf{a}_i = \Delta \mathbf{a}_{\pi(i)} + S_i \Delta \dot{\nu}_i^{(k+1)}$

Algorithm 3: LCABA.

Require: $\mathbf{q}, \nu, \tau, S_i s, H_i s, K_i s, \mathbf{k}_i s, \mathbf{f}_i^{\text{ext}}, \mathcal{S}^\mathcal{E}, \mu, \lambda^0, \mathcal{E}, \epsilon, \max_iter$

- 1: Execute the three-sweep algorithm in Algorithm 1.
- 2: **for** k in range(1, max_iter)
- 3: **for** i in \mathcal{S}
- 4: $\Delta \mathbf{f}_i \leftarrow \mathbf{0}_6;$
- 5: **for** i in \mathcal{E} \triangleright ALM multiplier update
- 6: $\Delta \mathbf{k}_i \leftarrow K_i^1 \mathbf{a}_{l_i^1} + K_i^2 \mathbf{a}_{l_i^2} - \mathbf{k}_i$ \triangleright Compute constraint residual;
- 7: $\Delta \mathbf{f}_{l_i^2} \leftarrow \mu K_i^2 T \Delta \mathbf{k}_i; \Delta \mathbf{f}_{l_i^1} \leftarrow \mu K_i^1 T \Delta \mathbf{k}_i;$
- 8: $\lambda_i^{k+1} \stackrel{\pm}{\leftarrow} \mu \Delta \mathbf{k}_i;$
- 9: **if** $\min(\|\dot{\nu}^k - \dot{\nu}^{k-1}\|_\infty, \|\Delta \bar{\mathbf{k}}\|_\infty) < \epsilon$
- 10: **break** \triangleright Terminate if converged
- 11: Execute reduced-two sweep algorithm in Algorithm 2
- 12: $\dot{\nu}^{k+1} \leftarrow \dot{\nu}^k + \Delta \dot{\nu};$

effective complexity of LCABA to the best case complexity of $O(n + m)$.

V. PROXBBO

This section presents the proxBBO algorithm, which generalizes the state-of-the-art recursive algorithm BBO [18], [19] to the proximal dynamics formulation, enabling the handling of singular constraints.

A. ProxBBO Derivation

ProxBBO is derived using the dual PPA discussed in Section III-D, similarly to the LCABA derivation. ProxBBO uses DP to compute the following dual PPA iteration:

$$\lambda^{k+1} := \underset{\lambda}{\operatorname{argmin}} \left\{ - \left(\min_{\dot{\nu}} \mathcal{L}(\dot{\nu}, \mathbf{a}, \lambda) \right) + \frac{1}{2\mu} \|\lambda - \lambda^k\|^2 \right\} \quad (37)$$

where

$$\begin{aligned} \mathcal{L}(\dot{\boldsymbol{\nu}}, \mathbf{a}, \boldsymbol{\lambda}) = & \sum_{i \in \mathcal{S}} \left\{ \frac{1}{2} \mathbf{a}_i^T H_i \mathbf{a}_i - \mathbf{f}_i^T \mathbf{a}_i - \boldsymbol{\tau}_i^T \dot{\boldsymbol{\nu}}_i \right\} \\ & + \sum_{i \in \mathcal{E}} \left\{ \boldsymbol{\lambda}_i^T \left(K_i^{-1} \mathbf{a}_{i_1} + K_i^2 \mathbf{a}_{i_2} - \mathbf{k}_i \right) \right\}. \end{aligned} \quad (38)$$

Note that the proximal regularization term in (37) does not depend on $\dot{\boldsymbol{\nu}}$, so the proximal term can be pushed inside the inner minimization problem to get

$$\boldsymbol{\lambda}^{k+1} := \underset{\boldsymbol{\lambda}}{\operatorname{argmax}} \left\{ \underset{\dot{\boldsymbol{\nu}}}{\min} \left(\mathcal{L}(\dot{\boldsymbol{\nu}}, \mathbf{a}, \boldsymbol{\lambda}) - \frac{1}{2\mu} \|\boldsymbol{\lambda} - \boldsymbol{\lambda}^k\|^2 \right) \right\}. \quad (39)$$

This max–min problem will be solved using DP.

From the Lagrangian's structure, we anticipate linear–quadratic terms depending on \mathbf{a}_i and the Lagrange multipliers associated with the cut-joints of loops supported by the i th link. Therefore, we hypothesize the optimal cost-to-go *Lagrangian* for the DP step at a spanning-tree leaf-link to have the following quadratic form similarly to [15]:

$$\begin{aligned} V_i^{\mathcal{L}^*}(\mathbf{a}_i, \bar{\boldsymbol{\lambda}}_{\mathcal{N}_i}) := & \frac{1}{2} \begin{bmatrix} \mathbf{a}_i \\ \bar{\boldsymbol{\lambda}}_{\mathcal{N}_i} \end{bmatrix}^T \begin{bmatrix} H_{i,i} & \bar{K}_{i,\mathcal{N}_i}^T \\ \bar{K}_{i,\mathcal{N}_i} & -\bar{L}_{\mathcal{N}_i,\mathcal{N}_i} \end{bmatrix} \begin{bmatrix} \mathbf{a}_i \\ \bar{\boldsymbol{\lambda}}_{\mathcal{N}_i} \end{bmatrix} \\ & - \begin{bmatrix} \mathbf{f}_i \\ \bar{\mathbf{k}}_{\mathcal{N}_i} \end{bmatrix}^T \begin{bmatrix} \mathbf{a}_i \\ \bar{\boldsymbol{\lambda}}_{\mathcal{N}_i} \end{bmatrix} \end{aligned} \quad (40)$$

where $\bar{\boldsymbol{\lambda}}_{\mathcal{N}_i}$, $\bar{K}_{i,\mathcal{N}_i}$, $\bar{L}_{\mathcal{N}_i,\mathcal{N}_i}$, and $\bar{\mathbf{k}}_{\mathcal{N}_i}$ aggregate the dual variables, constraint matrices, dual Hessian terms, and desired constraint accelerations for all the loop constraints supported by the link i such that $\mathcal{N}_i = \text{LS}(i)$.

For every $j, k \in \mathcal{N}_i$, these terms are defined as follows:

$$\begin{aligned} \bar{\boldsymbol{\lambda}}_{\mathcal{N}_i} = & \left[\dots \quad \boldsymbol{\lambda}_j^T \quad \dots \right]^T, \quad \bar{\mathbf{k}}_{\mathcal{N}_i} = \left[\dots \quad \mathbf{k}_j^T \quad \dots \right]^T \\ \bar{K}_{i,\mathcal{N}_i} = & \begin{bmatrix} \vdots \\ K_{j,i} \\ \vdots \end{bmatrix}, \quad \bar{L}_{\mathcal{N}_i,\mathcal{N}_i} = \begin{bmatrix} L_{k,k} & \dots & L_{k,j} & \dots \\ \vdots & \ddots & \vdots & \ddots \\ L_{k,j}^T & \dots & L_{j,j} & \dots \\ \vdots & \ddots & \vdots & \ddots \end{bmatrix}. \end{aligned} \quad (41)$$

At the start of the DP elimination, the aforementioned terms are initialized to zero and updated by iterating over the cut-joints $j \in \mathcal{E}$ as follows:

$$K_{j,l_j^1} \leftarrow K_j^1, \quad K_{j,l_j^2} \leftarrow K_j^2 \quad (42a)$$

$$L_{j,j} \leftarrow \frac{1}{\mu} I_{m_j, m_j} \quad (42b)$$

$$\mathbf{k}_j \leftarrow \frac{1}{\mu} \boldsymbol{\lambda}_j^k. \quad (42c)$$

For each joint $i \in \mathcal{S}$

$$H_{i,i} \leftarrow H_i.$$

Similarly to LCABA, the DP recurrence relation for proxBBO is given by

$$\begin{aligned} V^{\mathcal{L}^*}(\mathbf{a}_{\pi(i)}, \bar{\boldsymbol{\lambda}}_{\mathcal{N}_i}) = & \frac{1}{2} \mathbf{a}_{\pi(i)}^T H_{\pi(i), \pi(i)} \mathbf{a}_{\pi(i)} - \mathbf{f}_{\pi(i)}^T \mathbf{a}_{\pi(i)} + \\ & \bar{\boldsymbol{\lambda}}_{\mathcal{N}_{\pi(i)}}^T \bar{K}_{\pi(i), \mathcal{N}_{\pi(i)}} \mathbf{a}_{\pi(i)} + \underset{\dot{\boldsymbol{\nu}}_i, \mathbf{a}_i}{\min} \left\{ V_i^{\mathcal{L}^*}(\mathbf{a}_i, \bar{\boldsymbol{\lambda}}_{\mathcal{N}_i}) - \boldsymbol{\tau}_i^T \dot{\boldsymbol{\nu}}_i \right\}. \end{aligned} \quad (43)$$

To solve the aforementioned minimization problem, link i 's acceleration \mathbf{a}_i is again eliminated via substitution using the acceleration recurrence relation to get the following simplified unconstrained QP, similar to (28):

$$\begin{aligned} \underset{\dot{\boldsymbol{\nu}}_i}{\text{minimize}} \quad & \frac{1}{2} \dot{\boldsymbol{\nu}}_i^T D_i \dot{\boldsymbol{\nu}}_i - \left[\boldsymbol{\tau}_i + S_i^T \left(\mathbf{f}_i - \bar{K}_{i,\mathcal{N}_i}^T \bar{\boldsymbol{\lambda}}_{\mathcal{N}_i} - \right. \right. \\ & \left. \left. H_{i,i} (\mathbf{a}_{\pi(i)} + \mathbf{a}_{b,i}) \right) \right]^T \dot{\boldsymbol{\nu}}_i \end{aligned} \quad (44)$$

the optimization of which yields

$$\dot{\boldsymbol{\nu}}_i^* = D_i^{-1} \left[\boldsymbol{\tau}_i + S_i^T \left(\mathbf{f}_i - \bar{K}_{i,\mathcal{N}_i}^T \bar{\boldsymbol{\lambda}}_{\mathcal{N}_i} - H_{i,i} (\mathbf{a}_{\pi(i)} + \mathbf{a}_{b,i}) \right) \right]. \quad (45)$$

Substituting the optimal $\dot{\boldsymbol{\nu}}_i^*$ expression back into (43) gives a quadratic form for the function $V^*(\mathbf{a}_{\pi(i)}, \bar{\boldsymbol{\lambda}}_{\mathcal{N}_i})$ with the following quadratic form coefficient updates:

$$H_{\pi(i), \pi(i)} \stackrel{\pm}{\leftarrow} P_i H_{i,i} \quad (46a)$$

$$\mathbf{f}_{\pi(i)} \stackrel{\pm}{\leftarrow} P_i (\mathbf{f}_i - H_{i,i} \mathbf{a}_{b,i}) - H_{i,i} S_i D_i^{-1} \boldsymbol{\tau}_i \quad (46b)$$

$$\bar{K}_{\pi(i), \mathcal{N}_i} \stackrel{\pm}{\leftarrow} \bar{K}_{i,\mathcal{N}_i} P_i^T \quad (46c)$$

$$\bar{L}_{\mathcal{N}_i, \mathcal{N}_i} \stackrel{\pm}{\leftarrow} \bar{K}_{i,\mathcal{N}_i} S_i D_i^{-1} S_i^T \bar{K}_{i,\mathcal{N}_i}^T \quad (46d)$$

$$\bar{\mathbf{k}}_{\mathcal{N}_i} \stackrel{\pm}{\leftarrow} \bar{K}_{i,\mathcal{N}_i} \left[P_i^T \mathbf{a}_{b,i} + S_i D_i^{-1} (\boldsymbol{\tau}_i + S_i^T \mathbf{f}_i) \right]. \quad (46e)$$

The inertia and force recursions in (46a) and (46b) are identical to the ABA and LCABA equations in (30). However, the $H_{i,i}$ and \mathbf{f}_i terms computed by proxBBO differ from the corresponding terms in LCABA because proxBBO's Lagrangian [see (38)] does not contain the quadratic penalty terms of the ALF used in LCABA. The i th link's constraint matrix $\bar{K}_{i,\mathcal{N}_i}$ is backpropagated to the parent link $\pi(i)$ in (46c). The set $\text{LS}(\pi(i))$ for each link can be recursively computed using the following update rule:

$$\mathcal{N}_{\pi(i)} \leftarrow \mathcal{N}_{\pi(i)} \cup \mathcal{N}_i. \quad (47)$$

Early elimination: The spanning-tree leaf-links can all be eliminated sequentially using the recursive formulas above before solving for the optimal Lagrange multipliers. However, this would introduce expensive coupling between all the loops, eventually leading to $O(n + m^2n + m^3)$ operations. To counter this, the authors of [18] and [19] propose eliminating loop constraints as soon as all links supporting the corresponding loop are eliminated. We adopt this approach for proxBBO as well.

The last link to be eliminated among the links supporting loop j is the loop's root link $i = \tau_j$ from its definition in Section III. The set of loops supported by link i , \mathcal{N}_i , can be partitioned

into two sets: 1) the set of loops for which the link i is a root denoted as \mathcal{R}_i and 2) the remaining neighbor loops $\mathcal{U}_i := \mathcal{N}_i - \mathcal{R}_i$. The optimal cost-to-go Lagrangian function from (40) is also expanded based on this partition

$$V_i^{\mathcal{L}^*}(\mathbf{a}_i, \bar{\boldsymbol{\lambda}}_{\mathcal{U}}, \bar{\boldsymbol{\lambda}}_{\mathcal{R}}) := \frac{1}{2} \begin{bmatrix} \mathbf{a}_i \\ \bar{\boldsymbol{\lambda}}_{\mathcal{U}_i} \\ \bar{\boldsymbol{\lambda}}_{\mathcal{R}_i} \end{bmatrix}^T \begin{bmatrix} H_{i,i} & \bar{K}_{i,\mathcal{U}_i}^T & \bar{K}_{i,\mathcal{R}_i}^T \\ \bar{K}_{i,\mathcal{U}_i} & -\bar{L}_{\mathcal{U}_i,\mathcal{U}_i} & -\bar{L}_{\mathcal{U}_i,\mathcal{R}_i} \\ \bar{K}_{i,\mathcal{R}_i} & -\bar{L}_{\mathcal{U}_i,\mathcal{R}_i}^T & -\bar{L}_{\mathcal{R}_i,\mathcal{R}_i} \end{bmatrix} \begin{bmatrix} \mathbf{a}_i \\ \bar{\boldsymbol{\lambda}}_{\mathcal{U}_i} \\ \bar{\boldsymbol{\lambda}}_{\mathcal{R}_i} \end{bmatrix} - \begin{bmatrix} \mathbf{f}_i \\ \bar{\mathbf{k}}_{\mathcal{U}_i} \\ \bar{\mathbf{k}}_{\mathcal{R}_i} \end{bmatrix}^T \begin{bmatrix} \mathbf{a}_i \\ \bar{\boldsymbol{\lambda}}_{\mathcal{U}_i} \\ \bar{\boldsymbol{\lambda}}_{\mathcal{R}_i} \end{bmatrix}. \quad (48)$$

The Lagrange multipliers associated with the loops in \mathcal{R}_i are now eliminated, leading to the following updates to the DP cost function and the i th link's neighbor set:

$$V_i^{\mathcal{L}^*}(\mathbf{a}_i, \bar{\boldsymbol{\lambda}}_{\mathcal{U}}) \leftarrow \max_{\bar{\boldsymbol{\lambda}}_{\mathcal{R}}} V_i^{\mathcal{L}^*}(\mathbf{a}_i, \bar{\boldsymbol{\lambda}}_{\mathcal{U}}, \bar{\boldsymbol{\lambda}}_{\mathcal{R}}) \quad (49)$$

$$\mathcal{N}_i \leftarrow \mathcal{U}_i. \quad (50)$$

The optimizer $\bar{\boldsymbol{\lambda}}_{\mathcal{R}_i}^*$ of (49) is given by the necessary first-order optimality conditions of the corresponding QP problem

$$\bar{\boldsymbol{\lambda}}_{\mathcal{R}_i}^* = \bar{L}_{\mathcal{R}_i,\mathcal{R}_i}^{-1} (\bar{K}_{i,\mathcal{R}_i}^T \mathbf{a}_i - \bar{L}_{\mathcal{U}_i,\mathcal{R}_i}^T \bar{\boldsymbol{\lambda}}_{\mathcal{U}_i} - \bar{\mathbf{k}}_{\mathcal{R}_i}) \quad (51)$$

where $\bar{L}_{\mathcal{R}_i,\mathcal{R}_i}$ is invertible because it is initialized as a positive-definite diagonal matrix due to the proximal regularization [see (42)], which is followed by adding symmetric positive-semidefinite matrices to its diagonal blocks in (46d). Substituting optimal Lagrange multipliers back into the original DP cost function in (48) gives the following recursive formulae for the coefficients of $V_i^{\mathcal{L}^*}(\mathbf{a}_i, \bar{\boldsymbol{\lambda}}_{\mathcal{U}})$:

$$\begin{bmatrix} H_{i,i} & \bar{K}_{i,\mathcal{U}_i}^T \\ \bar{K}_{i,\mathcal{U}_i} & -\bar{L}_{\mathcal{U}_i,\mathcal{U}_i} \end{bmatrix} \leftarrow \begin{bmatrix} \bar{K}_{i,\mathcal{R}_i}^T \\ -\bar{L}_{\mathcal{U}_i,\mathcal{R}_i} \end{bmatrix} \bar{L}_{\mathcal{R}_i,\mathcal{R}_i}^{-1} \begin{bmatrix} \bar{K}_{i,\mathcal{R}_i}^T \\ -\bar{L}_{\mathcal{U}_i,\mathcal{R}_i} \end{bmatrix}^T \quad (52a)$$

$$\begin{bmatrix} \mathbf{f}_i \\ \bar{\mathbf{k}}_{\mathcal{U}_i} \end{bmatrix} \leftarrow \begin{bmatrix} \bar{K}_{i,\mathcal{R}_i}^T \\ -\bar{L}_{\mathcal{U}_i,\mathcal{R}_i} \end{bmatrix} \bar{L}_{\mathcal{R}_i,\mathcal{R}_i}^{-1} \bar{\mathbf{k}}_{\mathcal{R}_i}. \quad (52b)$$

This way, all links and loop constraints are eliminated from the leaf links down to the root link. The second forward sweep then computes the numerical values of the joint accelerations and the optimal Lagrange multipliers using (45) and (51), respectively. Note that similarly to the LCABA algorithm, the subsequent proximal iterations are efficient because only relatively inexpensive acceleration and force computations in (42c), (45), (46b), (46e), (51), and (52b) need to be evaluated.

B. ProxBBO Illustrative Example

ProxBBO is illustrated in Fig. 5 on the same mechanism used in the LCABA illustration in Section IV-B. Compared to LCABA, ProxBBO introduces additional nodes for the Lagrange multipliers of each cut-joint constraint, depicted as red nodes in Fig. 5(a). Similarly to LCABA, elimination of each link

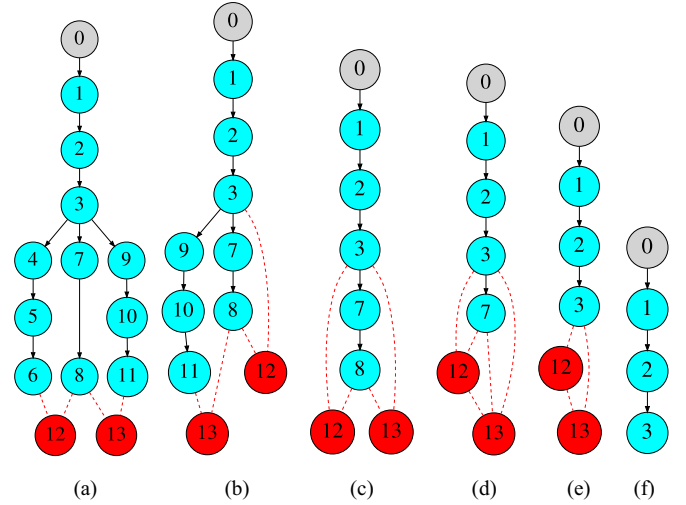


Fig. 5. (a)–(f) Graphical illustration of the ProxBBO's elimination steps. Compared to LCABA, additional red nodes are introduced to represent cut-joint Lagrange multipliers.

or constraint introduces coupling between all the neighbors (including a link's parent) of the eliminated link or constraint. Eliminating links 4–6 results in a graph where constraint 12 is coupled with the link 3, as seen in Fig. 5(b). Then, eliminating links 9–11 couples constraint 13 and link 3, as seen in Fig. 5(c). Eliminating link 8 introduces coupling between constraints 12 and 13, as seen in Fig. 5(d), since link 8 supports both the loops. Upon eliminating the next leaf link, link 7, we arrive at link 3 in Fig. 5(e), which is the root link of both loops 12 and 13. Both constraints are then eliminated to get the graph in Fig. 5(f), after which the elimination steps proceed identically to the ABA algorithm.

C. ProxBBO Algorithm

This section presents ProxBBO in an algorithmic form. Similarly to LCABA in Section IV-C, the three-sweep algorithm corresponding to the first proximal iteration is presented in Algorithm 4, followed by the reduced two-sweep algorithm for the subsequent iterations in Algorithm 5 and the entire ProxBBO algorithm in Algorithm 6.

D. ProxBBO Complexity Analysis

Let the maximum number of loops supported by a link be

$$m_b = \max_{i \in \mathcal{S}} \mathcal{C}(\mathcal{N}_i) \quad (53)$$

in the ProxBBO algorithm. The three-sweep algorithm from Algorithm 4 corresponding to the first proximal iteration dominates the computational cost of the algorithm. Reasoning similarly to the LCABA analysis, the first forward sweep can be shown to require $O(n + m)$ operations. The second forward sweep is more expensive than LCABA due to line 23 requiring $O(\mathcal{C}(\mathcal{N}_i)^2)$ operations per joint, bringing the second forward sweep's cost to $O(n + m_b^2 n)$ operations. The backward sweep is the most expensive part, where backpropagating the constraint

Algorithm 4: proxBBO Three-Sweep Algorithm.

Require: $\mathbf{q}, \nu, \tau, S_i s, H_i s, K_i s, \mathbf{k}_i s, \mathbf{f}_i^{\text{ext}}, \mu, \lambda^0, \mathcal{E}$

- 1: **for** i in \mathcal{S} \triangleright First forward sweep
- 2: FK and initialize inertias and forces (see Section IV-C1).
- 3: $\mathcal{N}_i \leftarrow \{\}$
- 4: **for** j in \mathcal{E} \triangleright Initialize constraint terms and connections
- 5: $l_j^1 K_j \leftarrow K_j^1; l_j^2 K_j \leftarrow K_j^2;$
- 6: $L_{j,j} \leftarrow \frac{1}{\mu} I_{m_j, m_j}; \mathbf{k}_j \leftarrow \frac{1}{\mu} \lambda_j^0;$
- 7: $\mathcal{N}_{l_j^1} \leftarrow \mathcal{N}_{l_j^1} \cup \{j\}; \mathcal{N}_{l_j^2} \leftarrow \mathcal{N}_{l_j^2} \cup \{j\};$
- 8: **for** i in \mathcal{S}_r \triangleright Backward sweep
- 9: **if** $\mathcal{C}(\mathcal{R}_i) > 0$ \triangleright Early eliminate constraints rooted at i
- 10: $\mathcal{U}_i \leftarrow \mathcal{N}_i - \mathcal{R}_i; \mathcal{N}_i \leftarrow \mathcal{U}_i;$ \triangleright Update neighbor set
- 11: Update $H_{i,i}, \bar{K}_{i,\mathcal{U}_i}, \bar{L}_{\mathcal{U}_i, \mathcal{U}_i}, \mathbf{f}_i,$ and $\bar{\mathbf{k}}_{\mathcal{U}_i}$ after constraint elimination using (52);
- 12: Compute resultant torques, apparent inertia (its inverse) and projection matrix (see Section IV-C2).
- 13: **if** $\mathcal{C}(\mathcal{N}_i) > 0$ \triangleright Propagate remaining constraints to parent
- 14: $\mathcal{R}_{\pi(i)} \leftarrow \mathcal{U}_i \cup (\mathcal{N}_i \cap \mathcal{N}_{\pi(i)}); \mathcal{N}_{\pi(i)} \leftarrow \mathcal{N}_i;$
- 15: Update $\bar{K}_{\pi(i), \mathcal{N}_i}, \bar{L}_{\mathcal{N}_i, \mathcal{N}_i}$ and $\mathbf{k}_{\mathcal{N}_i}$ using (46c)–(46e);
- 16: **if** $\pi(i) > 0$
- 17: Backpropagate inertia and force to parent link (see Section IV-C3).
- 18: **for** i in \mathcal{S} \triangleright Second forward sweep (roll-out)
- 19: **if** $\mathcal{C}(\mathcal{N}_i) > 0$ \triangleright Add constraint forces to joint torques
- 20: $\mathbf{u}_i \leftarrow (\bar{K}_{i, \mathcal{N}_i} S_i)^T \bar{\lambda}_{\mathcal{N}_i}^1$
- 21: Compute resulting joint and link accelerations (see Section IV-C4).
- 22: **if** $\mathcal{C}(\mathcal{R}_i) > 0$ \triangleright Compute eliminated constraint forces
- 23: $\bar{\lambda}_{\mathcal{R}_i}^1 = \bar{L}_{\mathcal{R}_i, \mathcal{R}_i}^{-1} (\bar{K}_{i, \mathcal{R}_i} \mathbf{a}_i^1 - \bar{L}_{\mathcal{U}_i, \mathcal{R}_i}^T \bar{\lambda}_{\mathcal{U}_i}^1 - \bar{\mathbf{k}}_{\mathcal{R}_i})$

matrices, constraint couplings, and the constraint accelerations in line 15 incur $O(\mathcal{C}(\mathcal{N}_i))$, $O(\mathcal{C}(\mathcal{N}_i)^2)$, and $O(\mathcal{C}(\mathcal{N}_i))$ operations, respectively, at each joint, bringing their total worst case cost to $O(m_b^2 n)$ operations. ProxBBO also factorizes the constraint coupling matrices in line 11 during early elimination, incurring a cubic cost in the number of eliminated constraints, upper-bounded by $O(m_b^3)$ at each joint. If all the loops are coupled, this cost can even reach $O(m_b^3)$ operations. This brings the backward sweep cost to $O(n + m_b^3 n)$. The reduced two-sweep algorithm in Algorithm 5 reuses the factorization from the three-sweep algorithm and incurs a lower cost of $O(n + m_b^2 n)$ operations. Therefore, the total worst case complexity of the proxBBO three-sweep algorithm, being dominated by the backward sweep computations, is $O(n + m_b^3 n)$. While the worst case computational complexity of LCABA required pathological cases to manifest, the proxBBO algorithm’s worst case complexity is more likely to be encountered in the common case

Algorithm 5: proxBBO Reduced Two-Sweep Algorithm.

Require: $\Delta \mathbf{f}_i s, \Delta \mathbf{u}, K_i s, (\bar{K}_{i, \mathcal{N}_i} S_i) s, U_i s, D_i s, S_i s$

- 1: **for** i in \mathcal{S}_r \triangleright Reduced backward sweep
- 2: **if** $\mathcal{C}(\mathcal{R}_i) > 0$ \triangleright Updates from early eliminated constraints
- 3: $\begin{bmatrix} \Delta \mathbf{f}_i \\ \Delta \bar{\mathbf{k}}_{\mathcal{U}_i} \end{bmatrix} \leftarrow \begin{bmatrix} \bar{K}_{i, \mathcal{R}_i}^T \\ -\bar{L}_{\mathcal{U}_i, \mathcal{R}_i} \end{bmatrix} \bar{L}_{\mathcal{R}_i, \mathcal{R}_i}^{-1} \Delta \bar{\mathbf{k}}_{\mathcal{R}_i};$
- 4: $\Delta \mathbf{u}_i \leftarrow S_i^T \Delta \mathbf{f}_i$ \triangleright Compute change in joint torque;
- 5: **if** $\mathcal{C}(\mathcal{N}_i) > 0$ \triangleright Update desired constraint accelerations
- 6: $\Delta \bar{\mathbf{k}}_{\mathcal{N}_i} \leftarrow (\bar{K}_{i, \mathcal{N}_i} S_i) (D_i^{-1} \Delta \mathbf{u}_i);$
- 7: **if** $\pi(i) > 0$ \triangleright Propagate force changes to parent
- 8: $\Delta \mathbf{f}_{\pi(i)} \leftarrow \Delta \mathbf{f}_i - U_i (D_i^{-1} \Delta \mathbf{u}_i);$
- 9: $\Delta \mathbf{a}_0 \leftarrow \mathbf{0}_6;$
- 10: **for** i in \mathcal{S} \triangleright Reduced forward sweep
- 11: **if** $\mathcal{C}(\mathcal{N}_i) > 0$ \triangleright Add constraint forces to joint torques
- 12: $\Delta \mathbf{u}_i \leftarrow -(\bar{K}_{i, \mathcal{N}_i} S_i)^T \Delta \bar{\lambda}_{\mathcal{N}_i};$
- 13: $\Delta \dot{\nu}_i^{(k+1)} = D_i^{-1} \Delta \mathbf{u}_i - (U_i D_i^{-1})^T \Delta \mathbf{a}_{\pi(i)};$
- 14: $\Delta \mathbf{a}_i = \Delta \mathbf{a}_{\pi(i)} + S_i \Delta \dot{\nu}_i^{(k+1)};$
- 15: **if** $\mathcal{C}(\mathcal{R}_i) > 0$ \triangleright Compute eliminated constraint forces
- 16: $\Delta \bar{\lambda}_{\mathcal{R}_i}^{k+1} = \bar{L}_{\mathcal{R}_i, \mathcal{R}_i}^{-1} (\bar{K}_{i, \mathcal{R}_i} \Delta \mathbf{a}_i^{k+1} - \bar{L}_{\mathcal{U}_i, \mathcal{R}_i}^T \Delta \bar{\lambda}_{\mathcal{U}_i}^{k+1} - \Delta \bar{\mathbf{k}}_{\mathcal{R}_i});$

Algorithm 6: proxBBO.

Require: $\mathbf{q}, \nu, \tau, S_i s, H_i s, K_i s, \mathbf{k}_i s, \mathbf{f}_i^{\text{ext}}, \mu, \lambda^0, \mathcal{E}, \epsilon, \text{max_iter}$

- 1: Execute the three-sweep algorithm in Algorithm 4.
- 2: **for** k in range(1, max_iter)
- 3: **for** i in \mathcal{S}
- 4: $\Delta \mathbf{f}_i \leftarrow \mathbf{0}_6$
- 5: **for** i in \mathcal{E} \triangleright Update right-hand side of dual PPA problem
- 6: $\Delta \mathbf{k}_i^k \leftarrow \frac{1}{\mu} \Delta \lambda_i^k;$
- 7: **if** $\min(\|\Delta \dot{\nu}\|_\infty, \|\Delta \bar{\mathbf{k}}_\mathcal{E}\|_\infty) < \epsilon$ \triangleright Check convergence
- 8: break
- 9: Execute reduced two-sweep algorithm in Algorithm 5;
- 10: $\dot{\nu}^{k+1} \leftarrow \dot{\nu}^k + \Delta \dot{\nu};$
- 11: $\bar{\lambda}_\mathcal{E}^{k+1} \leftarrow \bar{\lambda}_\mathcal{E}^k + \Delta \bar{\lambda}_\mathcal{E};$

when the loops are external and coupled, e.g., ground contact for a quadruped.

VI. EXPERIMENTS

This section discusses the C++ implementation of LCABA and proxBBO, presents the computational benchmarking of the algorithms on various robot setups, and investigates the scaling of the algorithms for different topologies.

A. Implementation

The recursive algorithms LCABA² and proxBBO³ are implemented in C++ using the efficient open-source dynamics library PINOCCHIO [17] and are staged for release in `Pinocchio 4.0.0` version in January 2026. These algorithms are computationally benchmarked with the joint-space algorithm proxLTL, whose state-of-the-art version [5] is implemented in PINOCCHIO. All these three algorithms, being implemented in C++ and identically leveraging PINOCCHIO's efficient rigid-body dynamics functions, contribute to benchmarking fairness. The proxLTL implementation in Pinocchio is particularly mature with implementation improvements since [5] and leverages vectorization, which empirically gives it quadratic scaling as opposed to the theoretically expected cubic scaling [16], making the comparison between the recursive and the joint-space algorithms particularly fair toward the joint-space algorithms. All timings were benchmarked on a laptop running Ubuntu 22.04 LTS with an Intel Core Ultra 7 165H CPU, and the code was compiled using the Clang 21.1.2 compiler.

To simulate the constrained system, our implementation supports the semi-implicit Euler method

$$\boldsymbol{\nu}_{k+1} = \boldsymbol{\nu}_k + \Delta t \dot{\boldsymbol{\nu}}_k \quad (54)$$

$$\mathbf{q}_{k+1} = \mathbf{q}_k \oplus \Delta t \boldsymbol{\nu}_{k+1} \quad (55)$$

where \oplus denotes integration on manifold, supporting configuration spaces, including Lie groups due to floating base. Numerical integration can accumulate constraint drift over time. To stabilize the constraints against this drift, we support Baumgarte stabilization technique [54], where a proportional–derivative feedback term on the position and velocity terms of the constraint is added to the desired constraint acceleration term \mathbf{k}_i in (1c).

We note that there is diverse and rich literature on integration schemes for constrained systems [55], [56]. Users can readily integrate our algorithms within any of the explicit integration schemes. However, implicit or half-explicit DAE schemes [56, Sect. VII] typically rely on Newton iterations, whose Jacobian may not match the problem structure assumed in (1), except for special cases such as [57]. Therefore, the presented algorithms require nontrivial extensions to be integrated within such schemes. Investigating this is left for future work.

B. Benchmarking on Robot Setups

We benchmark the presented algorithms on five robot scenarios consisting of internal closed loops. We start with the 16-DoF Allegro hand (AH)⁴ holding a cube with its four fingertips. Connect-type 3-D constraints, which allow relative rotation but not relative translation, are imposed at the contact point between the fingertips and the cube. The next example involves two AHs collaboratively holding a cube. Subsequently, the two AHs

²`include/pinocchio/algorithm/loop-constrained-aba.hxx`

³`include/pinocchio/algorithm/proximal-bbo.hxx`

⁴Accessed: 19 August 2024. [Online]. Available: https://github.com/Gepetto/example-robot-data/tree/8d899847c8e7531a3d723b9647a79748056b0414/robots/allegro_hand_description/urdf

TABLE I

COMPUTATIONAL TIMINGS OF THE PROPOSED ALGORITHMS LCABA AND PROXBBO COMPARED WITH PROXLTL [5] UP TO TWO SIGNIFICANT DIGITS IN MICROSECONDS IS AVERAGED OVER 100 000 SAMPLES

System	LCABA	BBO	LTL	ABA
AH-cube- $T_{3D}^4(1)\{22\}$ [12]	4.1	5.1	4.4	2.0
AH-cube- $T_{3D}^4(3)$	5.6	7.3	5.9	–
AH ² -cube- $T_{3D}^8(1)\{38\}$ [24]	7.6	15	13	4.7
AH ² -cube- $T_{3D}^8(3)$	9.7	21	16	–
Hum-AH ² -cube- $T_{3D}^8(1)\{73\}$ [24]	13	26	25	9.0
Hum-AH ² -cube- $T_{3D}^8(3)$	15	33	28	–
Hum-AH ² -cube- $F_{6D}^2 T_{3D}^8(1)\{73\}$ [36]	13	31	32	9.0
Hum-AH ² -cube- $F_{6D}^2 T_{3D}^8(3)$	16	41	37	–
Digit- $F_{6D}^2 C_{6D}^6(1)\{44\}$ [48]	9.1	12	25	5.3
Digit- $F_{6D}^2 C_{6D}^6(3)\{44\}$ [48]	12	17	30	–
Digit-cube- $F_{6D}^2 C_{6D}^6 T_{3D}^2(1)\{50\}$ [54]	11	14	31	5.6
Digit-cube- $F_{6D}^2 C_{6D}^6 T_{3D}^2(3)$	14	20	38	–
Digit ² -cube- $F_{6D}^4 C_{6D}^{12} T_{3D}^4(1)\{94\}$ [108]	22	37	160	11
Digit ² -cube- $F_{6D}^4 C_{6D}^{12} T_{3D}^4(3)$	29	50	180	–

Timings of ABA [1] is provided for reference. The number of proximal/ALM iterations executed are indicated in the parentheses, system DoF within curly braces, and constraint dimension within the box brackets. Note that 6-D constraints used for each loop in Digit's leg lead to a redundant constraint formulation.

holding the cube are attached to a humanoid⁵ robot's wrists to investigate the scaling of the algorithms. Then, 6-D `weld`-type constraints are imposed between the humanoid robot's feet and the ground, testing the algorithms on a combination of internal and external loops. Then, we consider the Digit humanoid robot, which has three closed loops on each leg, standing with `weld`-type constraints on the feet. This is followed by a Digit robot standing and holding a box with its wrists. Finally, we consider the example of two Digit platforms collaboratively holding a box with their wrists while standing.

Table I lists the benchmarking results. The robot name is listed on the left, with the superscript indicating the number of robots. The constraints are listed as T, F, or H, depending on whether the constraint is imposed on the fingertip, feet, or hand. The three closed loops in each of Digit's legs are modeled by cutting a link involved in the loop and introducing a fixed joint, denoted by C . The subscript on the constraint indicates the type of constraint, and the superscript indicates the number of such imposed constraints. For each example, the timings are listed for one proximal/ALM iteration and three iterations to indicate how the cost might scale for a higher number of iterations. The computation timings are averaged over 100 000 samples and are reported in microseconds. The benchmarking was done in Ubuntu's terminal mode to avoid interference with the background processes affecting benchmarking results. Intel's `Turbo Boost` was left turned ON since we did not observe appreciable differences in computation timings between different runs. For reference, the computation timings of the vanilla ABA algorithm for unconstrained dynamics are also reported in the last column. The ALM parameter was set to an unrealistically low value of

⁵[Online]. Available: https://github.com/stack-of-tasks/pinocchio/blob/25714c7d738b08e98201871757811525db74f2aa/models/simple_humanoid.urdf

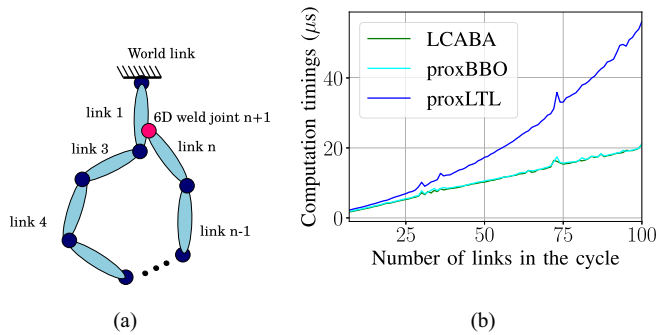


Fig. 6. Computational scaling of the different CDAs for cyclic mechanisms. (a) Cycle with n links. (b) Computation timings in microseconds for the first prox/ALM iteration for the cycle mechanism.

$\mu = 10^1$ to ensure that all the algorithms will execute three iterations without converging to obtain the timings for three iterations. For a single AH grasping a cube, both proxLTL and LCABA are competitive. Note that the AH’s topology, with its extensive branching and short depth, best suits the proxLTL algorithm. The resulting JSIM matrix enjoys a favorable sparsity structure that is even block-diagonal. The proxBBO algorithm is more expensive than the other algorithms for this contact-rich AH example because all the loops are coupled through the free-floating cube leading to its worst case (more on the worst case in Section VI-C) cubic complexity. LCABA emerges as the most efficient algorithm and remains so for the rest of the cases with larger robots due to its lower computational complexity. For the Digit robot, both the recursive algorithms scale better than the higher complexity proxLTL algorithm, with LCABA being the fastest among the recursive algorithms. ProxBBO’s improved performance on the Digit platform is due to limited coupling among constraints due to the closed loops in the legs being local. For the case of two Digits holding a box together, LCABA is even over $6\times$ faster than the proxLTL algorithm.

C. Scaling Results

This section studies how the algorithms scale for different robot topologies, starting with a single loop with a varying number of links, followed by a chain of loops, a worst case mechanism topology where every loop is coupled with every other loop, and finally, a topology that particularly favors LCABA over proxBBO.

Single Loop: A schematic diagram of a cycle of n links is shown in Fig. 6(a), with the last link connected to the first link through a 6D weld joint (shown in red) and Fig. 6(b) shows the computation timings in μs for the first prox/ALM iteration. The proxLTL algorithm scales superlinearly with the number of links, as expected. Note that its cost is empirically observed to be quadratic and not the theoretically expected cubic cost due to the efficient implementation that leverages vectorization. Both the recursive algorithms display similar scaling since the cost of propagating a single constraint through a loop is similar for both the algorithms.

Chain of Loops: Next, we benchmark the algorithms on a chain of loops, as shown in Fig. 7(a), with each loop consisting of

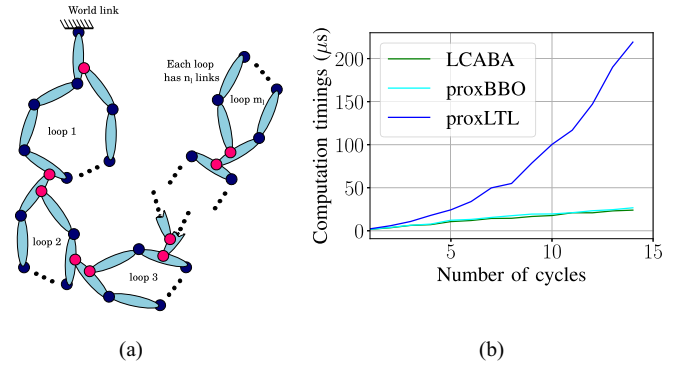


Fig. 7. Computational scaling of the different CDAs for cyclic mechanisms. (a) Chain of m_i loops with each loop consisting of n_i links. (b) Computation timings in microseconds for the first prox/ALM iteration for the cycle mechanism.

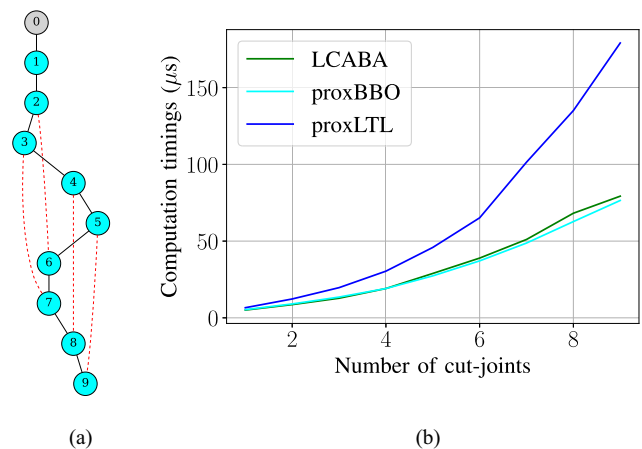


Fig. 8. Computational scaling of the different CDAs for cyclic mechanisms. (a) Simple worst case mechanism topology. (b) Computation timings in microseconds for the first prox/ALM iteration for the cycle mechanism.

seven links. The computation timings plotted in Fig. 7(b) follow a similar trend as the single-loop case for the same reasons, with proxBBO being slightly faster than LCABA and proxLTL being the slowest.

Worst Case Mechanism Topology: The algorithms are next benchmarked on a mechanism topology shown in Fig. 8(a), where all the loops are coupled. Each loop formed contains at least one joint from every other loop. The computation timings in Fig. 8(b) show that all the algorithms scale superlinearly, and the results indicate that the speed-up provided by the recursive algorithms is not as high in the worst case scenario, where they are subject to a similar cubic computational complexity, which is fortunately not encountered in practice.

Favorable Mechanism Topology: We now consider a topology where there is branching arising from a single link, and the branch tips are connected with each other with a 6-D constraint, as seen in Fig. 9(a). LCABA’s minimum degree heuristic ensures that the neighbor count for a leaf link being eliminated does not exceed one. The computation timings shown in Fig. 9(b) demonstrate that LCABA’s scaling is much better than proxBBO, which, in turn, scales better than proxLTL for this mechanism topology.

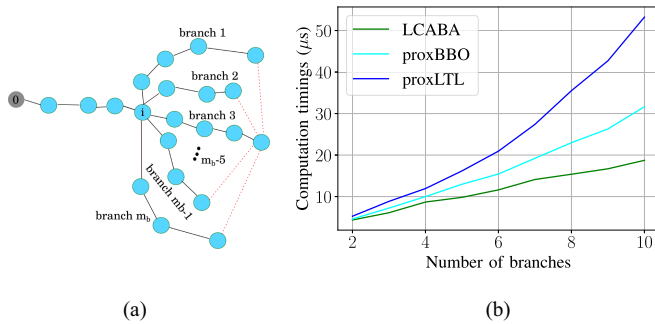


Fig. 9. Computational scaling for a mechanism topology that favors LCABA over proxBBO. (a) Favorable mechanism topology with branching from a single link i . (b) Computation timings in microseconds for the first prox/ALM iteration as the number of branches m_b increases.




Fig. 10. Benchmarking convergence of the different algorithms on the two AHs grasping a cube (second case in Table I with 24 constraints T_{3D}^8). This leads to a redundant constraint formulation that requires proximal methods. The constraint residual's inf norm is shown for 11 proximal/ALM iterations for $\mu = 10^5$ for 10 000 randomly generated examples. The box depicts mean and standard deviation, while line depicts the maximum and minimum values.

D. Convergence of the Algorithms

This section investigates the convergence of the constraint residuals (the ℓ_∞ norm of the residuals to be precise) and the numerical stability of the presented algorithms. Benchmarking is done for the second case in Table I with two AHs grasping a cube with their fingertips, which leads to a singular case due to redundant constraint formulation. The results, plotted in Fig. 10 for $\mu = 10^5$, depict mean and standard deviation of the constraint residual with the boxes and maximum and minimum values of the constraint residual with the lines. These statistics, generated over 10 000 randomly sampled robot positions and control inputs, indicate rapid convergence for all three algorithms within a few iterations. Faster convergence was observed for higher values of μ . However, the ALM behind LCABA is known to be numerically sensitive to large quadratic penalty parameters μ . Therefore, the numerical sensitivity of the different algorithms was investigated by comparing the solution ($\dot{\nu}$) of different algorithms against the relatively numerically stable solution of proxLTL with $\mu = 10^3$ and the results are plotted in Fig. 11 over different values of μ . As expected, the ALM-based LCABA algorithm's solution deviated from the reference solution with increasing μ values, while both proxBBO and proxLTL demonstrated numerical stability up to $\mu = 10^{11}$.

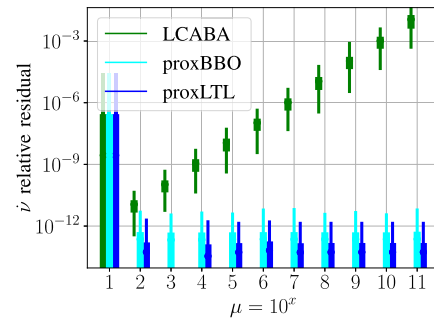


Fig. 11. Relative residual ($\frac{\|\dot{\nu} - \dot{\nu}_{\text{ref}}\|}{\|\dot{\nu}_{\text{ref}}\|}$) is plotted for the algorithms as μ increases, using the relatively numerically stable proxLTL algorithm with $\mu = 10^3$ as the reference over 10 000 randomly generated samples. The box depicts mean and standard deviation, while line depicts the maximum and minimum values.

Values of μ between 10^5 and 10^7 appear to provide a good balance between fast convergence and numerical stability for all the algorithms. We observed similar behavior across other setups and tasks, where these values lead to convergence within a tolerance of 10^{-6} typically within three iterations. Note that an effective value of μ needs to be $> 10^3$ times the largest eigenvalue of the JSIM, since the relative magnitude of the penalty parameter and the objective function of the QP determines the convergence of the ALM. This eigenvalue can be estimated using power iteration, which can be computed efficiently in a low-complexity manner using recursive Newton-Euler Algorithm (RNEA)-based sweeps. However, we do not pursue this direction of automatic selection of the μ parameter because it is fairly easy to tune. If users of the presented algorithms find themselves with a mechanism whose JSIM spectrum is significantly different enough for the recommended range of μ to not work, they can easily tune μ by going high enough that the convergence is fast without encountering numerical issues.

VII. DISCUSSIONS

This section critically discusses the presented algorithms, LCABA and proxBBO, their connections to existing literature, and potential directions for extensions. We first discuss how they generalize Riccati recursion to graphs. We then discuss the impact of the choice of spanning tree on the computational efficiency of the algorithms. We then highlight the connections between proxBBO and LCABA and how they can be combined to form a unified algorithm before connecting the algorithms to factor graphs and probabilistic inference. Finally, we discuss the choice of implicit versus explicit constraint formulations and its implications for the presented algorithms.

A. Generalizing Riccati Recursion to Graphs

The ABA [31], [32] and the PV algorithms [13] are known [15], [58] to generalize the celebrated Riccati recursion to tree-structured unconstrained and constrained equivalent LQR problems respectively. Such tree-structured Riccati recursions, developed in parallel in control and optimization literature [59], [60], are useful for solving stochastic optimal control problems

with scenario trees [61]. The presented recursive algorithms eliminate links from leaves to root for a spanning tree using the joint acceleration recurrence relation similarly to dynamics equation elimination via substitution in LQR solvers, effectively making them a generalization of Riccati recursion beyond tree structure to general graphs with loops. A straightforward occurrence of loops in an LQR problem is in periodic optimal control problems, where the periodicity constraint imposes the initial and terminal states to be equal. Efficient numerical algorithms for such periodic optimal control have been well studied (see [62], [63], and references therein). However, they do not appear to have been generalized to general graph structures.

Applying the presented algorithms to control problems for combining scenario trees and periodicity constraints, as well as studying the convergence and stabilization properties of such controllers, is a promising avenue for future work. Permitting graph structure in optimal control further enables interesting applications such as periodicity constraints at different frequencies for different subsets of states or for enforcing synchronization constraints in multiagent systems at different time instants.

B. Spanning Tree Selection

It is important to choose the spanning tree that is assumed as an input to the presented algorithms appropriately, especially since it can significantly impact their computational efficiency. In practical scenarios, the spanning-tree choice is often straightforward. External contact constraints, e.g., finger-cube contact constraints from Fig. 1(a), are modeled as cut joints, and the robot joints are included in the spanning tree to prioritize the mechanism's kinematic consistency. Even for robots with kinematic loops like the Digit robot, the actuated joints and the floating-base joint are typically chosen to get a spanning tree, and the closed-loop constraints and the submechanism constraints are modeled as cut joints. This also often yields a favorable spanning tree for the presented algorithms.

It may be desirable to algorithmically automate the optimal spanning tree selection by solving a secondary discrete-optimization problem for a given mechanism and algorithm. However, it is well known that finding an optimal elimination order, even without the restriction of conforming to a spanning-tree ordering, is an NP-complete problem [50]. Developing an effective algorithm for this problem is nontrivial and is unlikely to provide significant speed-ups for many existing robot topologies over manual spanning-tree selection. Therefore, this aspect is decidedly considered out of the scope of this article so as not to overload it.

Relaxing spanning-tree elimination order: The spanning-tree-based elimination ordering can be relaxed to fully leverage heuristics from numerical linear algebra such as minimum degree or minimum fill-in. However, such an approach does not exploit “term-level” sparsity inherent in joint acceleration recurrence relations. It is likely to be less efficient than the presented algorithms for most practical robots. It also results in a significantly more complex algorithm; for example, eliminating a joint and the corresponding link in the middle of a chain of one DoF joints results in constructing a new fictitious two DoF joint between the link's parent and child link after tedious

calculations, and the new constraint, moreover, does not, in general, have term-level sparsity. Overall, this approach is sensitive to singularities necessitating expensive pivoting methods and corresponds to general-purpose sparse linear solvers, which are known to be less efficient than specialized algorithms [1].

C. Relation to Factor Graphs and Probabilistic Inference

The GPLC problem for a mechanism with loops can also be interpreted as a probabilistic inference problem over a Bayesian network [64]. Consider the link accelerations, joint accelerations, and joint torques as random variables, with a Bayesian prior on each link's acceleration given by a Gaussian distribution whose covariance is the link's inverse inertia matrix and mean is the link's acceleration in the absence of joint constraints. Each joint imposes a deterministic constraint between the accelerations of different links. Being an equivalent representation of the GPLC problem, solving for the maximum likelihood solution of this network, conditioned upon the applied joint torques, provides the solution to the constrained dynamics problem. The factor graph perspective is not merely theoretically interesting, since nonserial DP and the variable elimination perspective used to derive this article's algorithms are the same ideas underpinning inference algorithms [28] over factor graphs. This implies that the dynamics algorithms from this article are suitable for probabilistic inference over factor graphs whose priors and constraints conform to the GPLC problem structure.

D. Connections Between proxBBO and LCABA

LCABA can be derived from proxBBO's proximal formulation by eliminating Lagrange multipliers before eliminating any primal variable. Starting with the proximal formulation in (39), swap the order in which λ and $\dot{\nu}$ are eliminated to get

$$\dot{\nu}^{(k+1)} = \underset{\dot{\nu}}{\operatorname{argmin}} \max_{\lambda} \left\{ \mathcal{L}(\dot{\nu}, \lambda) - \frac{1}{2\mu} \|\lambda - \lambda^{(k)}\|^2 \right\} \quad (56)$$

where solving the inner minimization problem yields LCABA's ALM formulation from (20) for the outer minimization problem. The equivalence between ALM and dual PPA is well known in optimization [65] and has also been observed between constrainedABA and proxPV in the context of CDAs for kinematic trees and external loops [16].

Therefore, the main difference between LCABA and proxBBO is the elimination ordering of the variables from the proximal formulation. ProxBBO can be more efficient for constraints involving a high number of links compared to the LCABA algorithm, making it suitable for constraints like CoM constraints. This suggests that a unified algorithm that can switch to proxBBO or LCABA based on the type of constraints and the connectivity graph structure can offer some speed-up. Like spanning-tree selection, this second-order enhancement requires discrete optimization and is left for future work. Similarly to constrainedABA [16], LCABA also generalizes the compliant constraint model in MUJoCo [8], with the first ALM iteration corresponding to solving MUJoCo's soft-Gauss principle (if dual variable initial guess $\lambda^0 = \mathbf{0}$) and the subsequent iterations converging to the rigid constraint model.

E. Explicit Versus Implicit Constraint Formulations

This article focused exclusively on the implicit constraint formulation. Explicit constraint formulations are readily specified only for a limited set of constraints like four-bar linkages with single-DoF joints and gear submechanisms. In other cases, they need to be derived from the implicit formulation, leading to an additional cost that can get particularly expensive for large loops. Another common strategy to obtain an explicit constraint formulation is to assign a subset of joints as *independent*, which is, however, prone to kinematic singularities. Compared to these approaches, the implicit formulation is more readily specified, can handle a wider variety of constraints efficiently, and is less prone to singularities.

There exists Jain's [41] LCE approach that exploits explicit constraint formulation to propose efficient recursive algorithms. LCE, while particularly suited for local loops, can get expensive for larger loops (like the dual-arm manipulation constraint or the feet-ground contact constraint for legged robots) because its cost increases cubically with the number of joints supporting the largest loop. While not as optimized for local loops as LCE, the presented algorithms can handle a wider variety of constraints efficiently. Investigating LCE's speedup compared to the presented algorithms for mechanisms with only local loops is interesting. Still, due to the complexity of implementing LCE within PINOCCHIO for fair comparison, this is considered outside the scope of this article. Moreover, the LCE approach is compatible with the GPLC derivation [43] and can be embedded in the derivation of proxBBO or LCABA to obtain a hybrid algorithm that uses implicit formulation for larger loops and explicit formulation for local loops. However, it is nontrivial to combine these approaches, and the computational benefit of such a hybrid algorithm compared to LCABA/proxBBO is unclear. Answering these questions must be left for future work so as not to overload this article.

VIII. CONCLUSION

This article culminates the development of low-complexity CDAs, proxBBO and LCABA, in the context of the proximal dynamics formulation [5], by efficiently extending proxPV and constrainedABA [16], respectively, to a wide class of mechanisms with internal closed loops. The presented algorithms leverage proximal/ALM iterations, enabling them to account for singular cases due to redundant constraints and singular configurations in a straightforward manner. It also contributes to revisiting and reviving recursive low-complexity CDAs, which had fallen out of favor in modern simulators compared to joint-space algorithms, by demonstrating compelling computational speed-ups compared to the state-of-the-art joint-space algorithm proxLTL.

LCABA matches proxLTL's performance for lower dimensional robots while providing over $6\times$ speed-ups for higher dimensional robots like humanoids with several internal closed loops. LCABA also typically outperforms proxBBO because it does not require additional factorizations to eliminate constraints and benefits from additional flexibility in choosing elimination orderings. It is also simpler to implement. However, proxBBO

may be preferable for constraints that involve a large number of bodies, such as CoM or momentum constraints, and is numerically more robust to high penalty parameters.

Limitations of the presented algorithms were also identified, such as LCABA being numerically sensitive to high penalty parameters μ . The presented algorithms also likely require non-trivial modifications to support implicit integration schemes, as well as some half-explicit schemes. Moreover, the proposed recursive algorithms assume that a spanning tree is provided as input. They are also limited to a spanning tree elimination order and implicit constraint formulations. Finally, the proposed CDAs are currently limited to equality constraints, while many practical applications involve motion constraints arising from contacts, which are inequality constraints or frictional contact constraints.

Future work: Future work will involve addressing these limitations. Most importantly, we will next explore strategies to support inequality constraints and frictional contact constraints leveraging the proposed algorithms. A natural avenue for such an extension is to use the alternating direction method of multipliers (ADMM)-style approach to handle frictional contact proposed in [7]. The inner subproblem solved at each ADMM iteration in this approach directly matches the equality-constrained proximal dynamics problem solved in this article. Another promising avenue for this extension would be to explore interior-point-based approaches [66], which also require solving a sequence of equality-constrained problems. However, the contribution of the barrier term to the KKT conditions does not directly match the proximal dynamics problem and, therefore, will require additional, but minor, modifications to the presented algorithms.

The assumption of a spanning tree can be relaxed by developing a discrete-optimization framework to compute the optimal spanning tree or even to relax the spanning-tree elimination ordering constraint when beneficial. The proposed algorithms can also be extended to support explicit constraint formulations by exploring a combination of the proposed algorithms with the LCE [41] approaches. Whether implicit integration schemes, with their energy dissipative properties, are beneficial in the context of robot control needs to be explored. If found to be relevant, we can also explore extending our algorithms to support such schemes. However, this extension is expected to be nontrivial. Finally, differentiability of the proximal operator also makes it suitable for gradient-based optimization methods.

Through this future research agenda, this article's algorithms have the potential to serve as the algorithmic foundation for speeding up contact-rich simulation and computation-intensive control applications such as MPC. Beyond mechanics, the presented algorithms can also find applications in control and estimation by effectively generalizing Riccati recursion to general graphs and through the factor graph connection.

ACKNOWLEDGMENT

Views and opinions expressed are those of the authors 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*. Berlin, Germany: Springer, 2014.
- [2] J. B. Rawlings, D. Q. Mayne, and M. Diehl, *Model Predictive Control: Theory, Computation, and Design*, vol. 2. Madison, WI, USA: Nob Hill Publishing, 2017.
- [3] R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction*, 2nd ed. Cambridge, MA, USA: The MIT Press, 2018.
- [4] R. Firoozi et al., "Foundation models in robotics: Applications, challenges, and the future," *Int. J. Robot. Res.*, vol. 44, no. 5, pp. 701–739, 2025.
- [5] J. Carpentier, R. Budhiraja, and N. Mansard, "Proximal and sparse resolution of constrained dynamic equations," in *Proc. Robot., Sci. Syst. Conf.*, 2021.
- [6] 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, Apr. 2018.
- [7] J. Carpentier, Q. L. Lidec, and L. Montaut, "From compliant to rigid contact simulation: A unified and efficient approach," in *Proc. Robot., Sci. Syst. Conf.*, Delft, The Netherlands, Jul. 2024.
- [8] E. Todorov, "Convex and analytically-invertible dynamics with contacts and constraints: Theory and implementation in MuJoCo," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2014, pp. 6054–6061.
- [9] E. Coumans and Y. Bai, "PyBullet, a Python module for physics simulation for games, robotics and machine learning," 2016–2021. [Online]. Available: <http://pybullet.org>
- [10] E. Todorov, T. Erez, and Y. Tassa, "MuJoCo: A physics engine for model-based control," in *Proc. IEEE/RSJ Int. Conf. Int. Robots. Syst.*, 2012, pp. 5026–5033.
- [11] 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.
- [12] R. Featherstone, "Exploiting sparsity in operational-space dynamics," *Int. J. Robot. Res.*, vol. 29, no. 10, pp. 1353–1368, 2010.
- [13] A. F. Vereshchagin, "Modeling and control of motion of manipulatory robots," *Sov. J. Comput. Syst. Sci.*, vol. 27, no. 5, pp. 29–38, 1989.
- [14] J. P. Popov, A. F. Vereshchagin, and S. L. Zenkevič, *Manipulacionnyye roboty: Dinamika i algoritmy*. Moscow, Russia: Nauka, 1978.
- [15] 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 Trans. Robot.*, vol. 40, pp. 729–749, 2024.
- [16] A. S. Sathya and J. Carpentier, "Constrained articulated body dynamics algorithms," *IEEE Trans. Robot.*, vol. 41, pp. 430–449, 2025.
- [17] J. Carpentier et al., "The pinocchio C++ library: A fast and flexible implementation of rigid body dynamics algorithms and their analytical derivatives," in *Proc. IEEE/SICE Int. Symp. Syst. Integr.*, 2019, pp. 614–619.
- [18] M. Otter, H. Brandl, and R. Johanni, "An algorithm for the simulation of multibody systems with kinematic loops," in *Proc. 7th World Congr. Theory Mach. Mech.*, 1987, pp. 407–411.
- [19] D.-S. Bae and E. J. Haug, "A recursive formulation for constrained mechanical system dynamics: Part II. Closed loop systems," *J. Struct. Mech.*, vol. 15, no. 4, pp. 481–506, 1987.
- [20] J. Nocedal and S. Wright, *Numerical Optimization*. Berlin, Germany: Springer, 2006.
- [21] C. F. Gauß, "Über ein neues allgemeines grundgesetz der mechanik," *J. für die reine und angewandte Mathematik*, vol. 4, pp. 232–235, 1829.
- [22] H. Bruyninckx and O. Khatib, "Gauss' principle and the dynamics of redundant and constrained manipulators," in *Proc. IEEE Int. Conf. Robot. Autom.*, 2000, pp. 2563–2568.
- [23] F. E. Udwardia and R. E. Kalaba, *Analytical Dynamics: A New Approach*. Cambridge, U.K.: Cambridge Univ. Press, 1996.
- [24] R. Bellman, "Dynamic programming," *Science*, vol. 153, no. 3731, pp. 34–37, 1966.
- [25] R. Aris, *Discrete Dynamic Programming: An Introduction to the Optimization of Staged Processes*. London, U.K.: Blaisdell, 1964.
- [26] U. Bertele and F. Brioschi, "Contribution to nonserial dynamic programming," *J. Math. Anal. Appl.*, vol. 28, no. 2, pp. 313–325, 1969.
- [27] U. Bertele and F. Brioschi, *Nonserial Dynamic Programming*. Orlando, FL, USA: Academic, 1972.
- [28] R. Dechter, "Bucket elimination: A unifying framework for reasoning," *Artif. Intell.*, vol. 113, nos. 1/2, pp. 41–85, 1999.
- [29] D. Baraff, "Linear-time dynamics using lagrange multipliers," in *Proc. 23rd Annu. Conf. Comput. Graph. Interact. Techn.*, 1996, pp. 137–146.
- [30] M. W. Walker and D. E. Orin, "Efficient dynamic computer simulation of robotic mechanisms," *J. Dyn. Syst., Meas., Control*, vol. 104, no. 3, pp. 205–211, 1982, doi: [10.1115/1.3139699](https://doi.org/10.1115/1.3139699).
- [31] A. Vereshchagin, "Computer simulation of the dynamics of complicated mechanisms of robot-manipulators," *Eng. Cybern.*, vol. 12, pp. 65–70, 1974.
- [32] R. Featherstone, "The calculation of robot dynamics using articulated-body inertias," *Int. J. Robot. Res.*, vol. 2, no. 1, pp. 13–30, 1983.
- [33] 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 Proc. Vol.*, vol. 19, no. 14, pp. 95–100, 1986.
- [34] B. Brogliato, *Nonsmooth Mechanics*, vol. 3. Berlin, Germany: Springer, 1999.
- [35] C. Duriez, F. Dubois, A. Kheddar, and C. Andriot, "Realistic haptic rendering of interacting deformable objects in virtual environments," *IEEE Trans. Vis. Comput. Graph.*, vol. 12, no. 1, pp. 36–47, Jan./Feb. 2006.
- [36] O. Khatib, "A unified approach for motion and force control of robot manipulators: The operational-space formulation," *IEEE J. Robot. Autom.*, vol. 3, no. 1, pp. 43–53, Feb. 1987.
- [37] R. T. Rockafellar, "Monotone operators and the proximal point algorithm," *SIAM J. Control Optim.*, vol. 14, no. 5, pp. 877–898, 1976.
- [38] M. J. Powell, "A method for nonlinear constraints in minimization problems," in *Optimization*. New York, NY, USA: Academic, 1969, pp. 283–298.
- [39] M. R. Hestenes, "Multiplier and gradient methods," *J. Optim. Theory Appl.*, vol. 4, no. 5, pp. 303–320, 1969.
- [40] N. Parikh and S. Boyd, "Proximal algorithms," *Found. Trends Optim.*, vol. 1, no. 3, pp. 127–239, 2014.
- [41] A. Jain, "Recursive algorithms using local constraint embedding for multibody system dynamics," in *Proc. Int. Des. Eng. Tech. Conf. Comput. Inf. Eng. Conf.*, 2009, vol. 49019, pp. 139–147.
- [42] A. Jain, "Multibody graph transformations and analysis: Part II: Closed-chain constraint embedding," *Nonlinear Dyn.*, vol. 67, pp. 2153–2170, 2012.
- [43] M. Chignoli, N. Adrian, S. Kim, and P. M. Wensing, "Recursive rigid-body dynamics algorithms for systems with kinematic loops," 2023, [arXiv:2311.13732v1](https://arxiv.org/abs/2311.13732v1).
- [44] R. M. Murray, Z. Li, and S. S. Sastry, *A Mathematical Introduction to Robotic Manipulation*. Boca Raton, FL, USA: CRC Press, 2017.
- [45] G. H. Golub and C. F. Van Loan, *Matrix Computations*. Baltimore, MD, USA: Johns Hopkins Univ. Press, 2013.
- [46] S. P. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [47] A. Bambade, S. El-Kazdadi, A. Taylor, and J. Carpentier, "Prox-QP: Yet another quadratic programming solver for robotics and beyond," in *Proc. Robot., Sci. Syst. Conf.*, 2022.
- [48] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, "OSQP: An operator splitting solver for quadratic programs," *Math. Program. Comput.*, vol. 12, no. 4, pp. 637–672, 2020.
- [49] A. Chiche and J. C. Gilbert, "How the augmented Lagrangian algorithm can deal with an infeasible convex quadratic optimization problem," *J. Convex Anal.*, vol. 23, no. 2, pp. 425–459, 2016.
- [50] M. Yannakakis, "Computing the minimum fill-in is NP-complete," *SIAM J. Algebr. Discrete Methods*, vol. 2, no. 1, pp. 77–79, 1981.
- [51] H. M. Markowitz, "The elimination form of the inverse and its application to linear programming," *Manage. Sci.*, vol. 3, no. 3, pp. 255–269, 1957.
- [52] A. George and J. W. Liu, *Computer Solution of Large Sparse Positive Definite*. Englewood Cliffs, NJ, USA: Prentice-Hall, 1981.
- [53] E. Cuthill and J. McKee, "Reducing the bandwidth of sparse symmetric matrices," in *Proc. 24th Nat. Conf.*, 1969, pp. 157–172.
- [54] J. Baumgarte, "Stabilization of constraints and integrals of motion in dynamical systems," *Comput. Methods Appl. Mech. Eng.*, vol. 1, no. 1, pp. 1–16, 1972.
- [55] M. Haddouni, V. Acary, S. Garreau, J.-D. Beley, and B. Brogliato, "Comparison of several formulations and integration methods for the resolution of DAEs formulations in event-driven simulation of nonsmooth frictionless multibody dynamics," *Multibody Syst. Dyn.*, vol. 41, no. 3, pp. 201–231, 2017.
- [56] G. Wanner and E. Hairer, *Solving Ordinary Differential Equations II*, vol. 375. Berlin, Germany: Springer, 1996.
- [57] V. Brasey and E. Hairer, "Symmetrized half-explicit methods for constrained mechanical systems," *Appl. Numer. Math.*, vol. 13, no. 1–3, pp. 23–31, 1993.
- [58] G. Rodriguez, "Kalman filtering, smoothing, and recursive robot arm forward and inverse dynamics," *IEEE J. Robot. Autom.*, vol. 3, no. 6, pp. 624–639, Dec. 1987.

- [59] M. C. Steinbach, "Tree-sparse convex programs," *Math. Methods Oper. Res.*, vol. 56, no. 3, pp. 347–376, 2003.
- [60] G. Frison, D. Kouzoupis, M. Diehl, and J. B. Jørgensen, "A high-performance Riccati based solver for tree-structured quadratic programs," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 14399–14405, 2017.
- [61] R. Marti, S. Lucia, D. Sarabia, R. Paulen, S. Engell, and C. de Prada, "Improving scenario decomposition algorithms for robust nonlinear model predictive control," *Comput. Chem. Eng.*, vol. 79, pp. 30–45, 2015.
- [62] J. J. Hench and A. J. Laub, "Numerical solution of the discrete-time periodic Riccati equation," *IEEE Trans. Autom. Control*, vol. 39, no. 6, pp. 1197–1210, Jun. 1994.
- [63] Y. Yang, "An efficient LQR design for discrete-time linear periodic system based on a novel lifting method," *Automatica*, vol. 87, pp. 383–388, 2018.
- [64] M. Xie and F. Dellaert, "A unified method for solving inverse, forward, and hybrid manipulator dynamics using factor graphs," 2019, *arXiv:1911.10065*.
- [65] R. T. Rockafellar, "A dual approach to solving nonlinear programming problems by unconstrained optimization," *Math. Program.*, vol. 5, no. 1, pp. 354–373, 1973.
- [66] V. Acary, P. Armand, H. Minh Nguyen, and M. Shpakovych, "Second-order cone programming for frictional contact mechanics using interior point algorithm," *Optim. Methods Softw.*, vol. 39, pp. 634–663, 2024.



Ajay Suresha Sathya received the bachelor's degree in mechanical engineering from NITK Surathkal, Mangaluru, India, in 2016, and the master's and Ph.D. degrees in mechanical engineering from KU Leuven, Leuven, Belgium, in 2018 and 2023, respectively.

He has been a Postdoctoral Researcher with Inria and École Normale Supérieure, Paris, France, in the Willow research team since December 2023. His research interests include algorithmic foundations of robot dynamics and control for developing agile and safe robots.



Justin Carpentier (Member, IEEE) graduated in 2014 with a double master in applied mathematics and electrical engineering from École Normale Supérieure Paris-Saclay, Paris, France, in 2014, and the Ph.D. degree in robotics from the University of Toulouse, Toulouse, France, in 2017.

He did his Ph.D. research with the Gepetto team, Laboratory for Analysis and Architecture of Systems, French National Center for Scientific Research, Toulouse, working on the computational foundations of legged locomotion. He is currently a Researcher with Inria and École Normale Supérieure, Paris, where he has been heading the Willow research team since 2023. He is also the leading developer and manager of widely used open-source robotics software, including Pinocchio, ProxSuite, HPP-FCL, and Aligator. His research interests include the interface of optimization, machine learning, computer vision, simulation, and control for robotics, with applications ranging from agile locomotion to dexterous manipulation.

Dr. Carpentier received an ERC Starting Grant focusing on laying the algorithmic and computational foundations of Artificial Motion Intelligence in 2024.