

M-ABD: Scalable, Efficient, and Robust Multi-Affine-Body Dynamics

ZHIYONG HE, University of Utah, USA
DEWEN GUO, University of Utah, USA
MINGHAO GUO, MIT, USA
YILI ZHAO, USC, USA
WOJCIECH MATUSIK, MIT, USA
HAO SU, UCSD, USA
CHENFANFU JIANG, UCLA, USA
PETER YICHEN CHEN, UBC, Canada
YIN YANG, University of Utah, USA




Fig. 1. **A huge pulley system.** This paper proposes an affine body dynamics (ABD) based framework that integrates multiple rigid/near-rigid links via kinematic joints. Such a multi-ABD system is designed based on a co-rotated formulation of a single affine body. This strategy makes an ABD system matrix pre-factorized even with the implicit integration. We then generalize all the commonly-seen joint constraints to the affine coordinate using either a linear or a nonlinear formulation. When joint constraints are exactly enforced with KKT, we can project any complex multibody system into the dual space with minimized freedoms. As most articulated systems are well-structured, we exploit topological connections among joints to quickly solve complex systems. The teaser figure reports a representative example of a large-scale multibody system consisting of 1,076,748 links. Our method robustly simulates this system at $h = 10^{-2}$ s using only one iteration for each time step. The simulation takes 904 ms only with a single CPU thread.

Simulating large-scale articulated assemblies poses a significant challenge due to the numerical stiffness and geometric complexity of jointed structures. Conventional rigid body solvers struggle with the high nonlinearity induced by rotation parameterization. This difficulty becomes more pronounced for multiple two-way-coupled bodies. This paper introduces a novel framework that leverages the linear kinematic mapping of Affine Body Dynamics (ABD). As ABD targets near-rigid objects, the constitutive variations of different materials become negligible, which justifies a co-rotational approach to

isolate geometric nonlinearities of the system. This insight enables the use of constant system matrices that can be pre-factorized throughout the simulation, even with fully implicit integration schemes. To manage the high DOF counts of large-scale systems, we map primal body coordinates onto a compact dual space defined by minimal joint degrees of freedom. By solving the resulting KKT systems, our method ensures exact constraint enforcement and physically accurate motion propagation. We provide a suite of specialized solvers tailored for diverse joint topologies, including chains, trees, closed loops, and irregular networks. Experimental results show that our approach achieves interactive rates for systems with hundreds of thousands of bodies on a single CPU core, while maintaining excellent stability at large time steps.

CCS Concepts: • **Computing methodologies** → **Physical simulation.**

Additional Key Words and Phrases: Rigid body dynamics, Multibody dynamics, Reduced model, Kinematic constraints

ACM Reference Format:

Zhiyong He, Dewen Guo, Minghao Guo, Yili Zhao, Wojciech Matusik, Hao Su, Chenfanfu Jiang, Peter Yichen Chen, and Yin Yang. 2026. M-ABD: Scalable,

Authors' Contact Information: Zhiyong He, University of Utah, Salt Lake City, USA, zhiyong.he@utah.edu; Dewen Guo, University of Utah, Salt Lake City, USA, guodw.sh@gmail.com; Minghao Guo, MIT, Cambridge, USA, guomh2014@gmail.com; Yili Zhao, USC, Los Angeles, USA, yilizhao.cs@gmail.com; Wojciech Matusik, MIT, Cambridge, USA, wojciech@mit.edu; Hao Su, UCSD, San Diego, USA, academic@haosu.ai; Chenfanfu Jiang, UCLA, Los Angeles, USA, chenfanfu.jiang@gmail.com; Peter Yichen Chen, UBC, Vancouver, Canada, pyc@csail.mit.edu; Yin Yang, University of Utah, Salt Lake City, USA, yin.yang@utah.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License.
© 2026 Copyright held by the owner/author(s).
ACM 1557-7368/2026/7-ART99
<https://doi.org/10.1145/3811276>

Efficient, and Robust Multi-Affine-Body Dynamics. *ACM Trans. Graph.* 45, 4, Article 99 (July 2026), 19 pages. <https://doi.org/10.1145/3811276>

1 Introduction

Rigid bodies and articulated multibody systems are central to computer graphics, robotics, fabrication, and control. Multibody dynamics is formulated based on the representation of individual rigid bodies and the mathematical modeling of their joint constraints. In the traditional rigid body dynamics (RBD) framework [Baraff 1997; Bender et al. 2014; Erleben et al. 2005], a rigid object is defined by a minimal set of six degrees of freedom (DOFs), representing its global translation and rotation. This formulation assumes that the body’s local geometry remains undeformed under any external forces or joint interactions. The material is therefore considered infinitely stiff. A drawback of RBD is that the relation between the spatial coordinates, i.e., x, y, z coordinates, and RBD coordinates, i.e., a certain parameterization of the rotation and translation vector, is nonlinear. This nonlinearity may be of less importance for a single body. However, in the context of a multibody system, the time-varying constraint Jacobian leads to a time-varying system matrix if the joint constraints need to be exactly enforced, e.g., via KKT (Karush-Kuhn-Tucker) systems [Boyd and Vandenberghe 2004]. Existing methods often rely on explicit integration with highly conservative time steps or employ penalty methods to approximate constraints, while fully implicit RBD-multibody simulation is less common. As the system scales up, these methods fail to enforce joint constraints, leading to numerical instability and constraint drift.

Affine Body Dynamics [Lan et al. 2022] serves as a more versatile alternative. By expanding the kinematic space from 6 DOFs to 12 DOFs, ABD replaces nonlinear rotations with a linear mapping. This linearity ensures that the Jacobian remains constant w.r.t. the body’s spatial coordinates, whereas the nonlinear rigidity is enabled with material stiffness. The resultant material nonlinearity requires the re-assembly and factorization of the system matrix at every iteration. As the kinematic space is expanded to 12 DOFs, ABD incurs higher computational costs than the RBD model.

We propose a scalable multibody simulation framework based on a dedicated numerical treatment of ABD. Our approach leverages a key observation: high material stiffness suppresses *constitutive nonlinearity*, and the key challenge lies in the *geometric nonlinearity*. Fortunately, it is known that difficulty can be resolved effectively through a co-rotational formulation [Müller and Gross 2004]. By extracting the body’s rotation, we land a constant system matrix that can be pre-factorized even using an implicit time integration scheme. This strategy serves as the core building block of our framework.

We provide a detailed analysis of commonly used joint constraints under the affine coordinate, and detail their possible linear and/or nonlinear formulations. For multibody systems consisting of a large number of stiff components, rather than resorting to penalty-based constraint approximation, we exactly enforce joint constraints via the KKT form to ensure accuracy and stability of the simulation. Since each body’s primal system is pre-factorized, projecting the global primal problem onto the dual space becomes highly efficient. By adopting a semi-linearization for each joint with minimal DOFs, we reach a dual problem whose dimension is significantly

smaller than that of the original primal version. We propose efficient algorithms for different mechanical structures, including a block-tridiagonal solver, an extension of the Featherstone algorithm [Featherstone 2008] for ABD-based multibody systems, a low-dimensional inverse approximation when kinematic loops exist, and a bidirectional Gauss-Seidel optimizer for dense joint networks.

We conduct extensive experiments to validate this new framework. Our method outperforms existing rigid body simulators and algorithms for both accuracy and efficiency. The proposed approach demonstrates excellent robustness, maintaining stability at a time step of 0.01 *sec* even for systems containing one million bodies. Because the joint constraints are strictly enforced via the KKT form, our simulator accurately propagates motion across the entire system. Even on a single-threaded CPU, we enable interactive simulation of highly complex multibody systems. Code and data for this paper are at <https://minsugly.github.io/mabd/>.

2 Related Work

Articulated rigid body simulation has a long history in graphics and robotics, with a central challenge being how to enforce joint constraints and contacts robustly at scale. Classical rigid body contact can be formulated as a Linear Complementarity Problem (LCP) dating back to early work on unilateral constraints and non-penetration [Baraff 1989; Moreau 1985], and many subsequent methods focus on improving robustness and efficiency of constraint solves and impact handling [Smith et al. 2012; Stewart 2000]. For articulated systems in particular, prior work explored impulse-based formulations for contact and joints [Mirtich and Canny 1994; Weinstein et al. 2006], adaptive or reduced representations for speed [Redon et al. 2005], and scalable treatments of highly constrained motion [Sueda et al. 2011]. Recent improvements include stabilization and geometric-acceleration ideas for constrained dynamics [Andrews et al. 2017; Tournier et al. 2015], as well as structured factorizations and Schur-complement techniques that exploit articulated sparsity [Enzenhöfer et al. 2019; Peiret et al. 2019]. These approaches highlight the importance of exploiting kinematic structure when solving large constrained systems, particularly for chains, trees, and lightly loopy mechanisms.

A canonical example of exploiting kinematic structure is the articulated-body algorithm (ABA) introduced by Featherstone [2008]. ABA achieves linear-time complexity for tree-structured articulated systems by recursively propagating articulated inertia and bias forces along the kinematic hierarchy. Subsequent work extends this idea to implicit integration, loop handling, and constraint-consistent formulations, and ABA-style recursion remains a foundational building block in robotics and multibody dynamics. Our solver draws conceptual inspiration from this line of work, but adapts the recursion and condensation strategy to affine generalized coordinates and KKT-based constraint projection.

Practical engines and toolkits such as MuJoCo, and other widely-used multibody platforms emphasize numerical robustness, efficient sparse solves, and engineered contact handling for application-driven workloads [Todorov et al. 2012]. These implementations illustrate important trade-offs between representation choice, solver

complexity, and runtime performance that guide design choices in both graphics and robotics-grade simulators.

Beyond rigid-link mechanisms, many graphics phenomena with long, filamentary structures (e.g., cables, hair, vegetation) are often modeled using continuum-inspired reduced formulations. Work on discrete elastic rods, strands, and vegetation emphasizes specialized reduced models and continuum mechanics to achieve efficiency for very long or highly elastic structures [Bergou et al. 2008]. These models focus on elastic deformation rather than exact articulated joints. Our articulated ABD formulation complements this literature by targeting scenarios in which rigid-link connectivity, exact joint enforcement, and large-scale constraint propagation dominate the computational cost.

Efficient articulated solvers in maximal or hybrid coordinates have also been developed with near-linear scaling in the number of bodies for many practical scenes [Wang et al. 2019]. Related efforts in robotics further explore structure-aware and parallel variants of articulated-body algorithms to improve scalability and numerical robustness in deep kinematic trees. Differentiable articulated simulators broaden these formulations to learning, system identification, and control, where stability and solver structure are critical [Geilinger et al. 2020; Macklin et al. 2019; Werling et al. 2021]. Complementary to these lines, quaternion-based constrained rigid body formulations provide an alternative, implementation-friendly parameterization and constraint handling strategy [Maloisel et al. 2025]. Many production simulators adopt semi-implicit stepping with stabilized velocity constraints (e.g., Baumgarte-type stabilization), which can exhibit drift and requires careful tuning; implicit position-level formulations instead solve constraints to a prescribed tolerance [Baumgarte 1972; Erez et al. 2015].

Beyond rigid-only pipelines, coupling rigid and deformable bodies introduces heterogeneous degrees of freedom and typically demands special-case treatments [Bai and Liu 2014; Jain and Liu 2011; Kaufman et al. 2008; Shinar et al. 2008]. Unified constraint views such as position-based dynamics can couple rigid and deformable components with customized iterative solvers, but do not generally provide strong global guarantees or exact constraint satisfaction [Deul et al. 2016; Müller et al. 2007, 2020]. In contrast, optimization-based time integration has recently enabled stronger robustness properties for contact and constraints. Incremental Potential Contact (IPC) formulates contact (and friction) via barrier energies inside a variational integrator, yielding intersection-free behavior under appropriate numerics [Li et al. 2020], and has been extended to codimensional objects [Li et al. 2021]. For rigid bodies, intersection-free formulations directly in $SE(3)$ have also been studied [Ferguson et al. 2021], offering a complementary perspective on robust contact handling through geometric integration.

Geometric and variational integrators on Lie groups preserve structure such as momentum, symplectic form, and orthogonality, and provide desirable long-term behavior for articulated mechanisms [Hairer et al. 2006; Marsden and West 2001]. While attractive theoretically, these methods are typically formulated in minimal-coordinate $SE(3)$ settings and introduce configuration-dependent Jacobians for constraint enforcement. By contrast, our affine co-rotational approach trades a modest increase in degrees of freedom

for constant kinematic mappings that interact more favorably with KKT-based condensation.

Our work is most closely related to affine-coordinate rigid modeling and optimization-based multibody frameworks. Affine Body Dynamics (ABD) represents each rigid body with a 12-DoF affine generalized coordinate and enforces rigidity via an orthogonality energy, while benefiting from Euclidean trajectories that simplify continuous collision detection and contact handling in collision-heavy scenes [Lan et al. 2022]. Compared to minimal-coordinate formulations on $SE(3)$, the resulting linear kinematic mappings and constant Jacobians can significantly simplify constraint assembly and projection. Chen et al. propose a unified Newton-barrier multibody framework that combines optimization-based time integration, IPC-style contact handling, and a general treatment of articulation constraints, including an efficient change-of-variables approach for linear equality constraints [Chen et al. 2022]. Augmented Vertex Block Descent (AVBD) [Giles et al. 2025] extends VBD [Chen et al. 2024] to handle stiff constraints via augmented penalty energies and GPU parallelism. While effective for deformable body simulations with mixed soft and stiff elements, AVBD approximates constraints with penalties rather than enforcing them exactly. Our method targets articulated rigid body systems where exact joint enforcement is essential.

Building on these foundations, our focus is on articulated rigid body simulation with affine generalized coordinates while designing a solver that is simple to author, preserves key geometric invariants under large rotations, and exploits problem structure so that the dominant computational cost depends on constraint connectivity rather than the total state dimension. In particular, we draw inspiration from articulated-body algorithms and Schur-complement-based elimination, but tailor them to the affine setting so that per-body systems can be pre-factorized and reused, enabling efficient dual-space solves for chains, trees, and more general joint graphs.

3 Single-object Stiff Dynamics

To make our exposition self-contained, we start with a brief review of RBD [Baraff 1997] and ABD [Lan et al. 2022] for a single rigid or near-rigid object. Classic RBD defines the position of any material point as $\mathbf{x}_i = \mathbf{R}\bar{\mathbf{x}}_i + \mathbf{t}$ for a given rotation matrix $\mathbf{R}$ and translation $\mathbf{t}$. Here, $\mathbf{x}_i$ and $\bar{\mathbf{x}}_i$ denote the displaced and rest-shape positions of point i . In other words, the rigidity of the object is *kinematically* prescribed.

3.1 ABD Kinematics

ABD allows the object to deform under a uniform linear transformation $\mathbf{A}(t)$ and a translation $\mathbf{t}(t)$ such that:

$$\mathbf{x}_i = \mathbf{A}(t)\bar{\mathbf{x}}_i + \mathbf{t}(t). \quad (1)$$

The rigidity of the object is approached with an elasticity potential of:

$$E_A = \int_{\Omega} \Psi d\Omega. \quad (2)$$

Ψ is the energy density function, and many known hyperelastic models can be used as long as they are rotation-invariant. For instance, the vanilla ABD [Lan et al. 2022] used a simple polynomial

energy of:

$$\Psi = k_A \|\mathbf{A}\mathbf{A}^\top - \mathbf{I}_3\|_F^2, \quad (3)$$

where $\|\cdot\|_F^2$ is the Frobenius norm; $\mathbf{I}_3$ is a 3×3 identity matrix; and k_A is the affine stiffness of the body.

The relaxation from $\mathbf{R} \in SO(3)$ to $\mathbf{A}$ offers both advantages and disadvantages compared with RBD. On the positive side, the affine DOFs can be used to store elastic energies, which naturally integrates with energy-based collision and contact frameworks such as Incremental Potential Contact (IPC) [Li et al. 2020]. Suppose the body is discretized by a tetrahedral mesh with N nodes. The nodal positions $\mathbf{x}$ a.k.a. the *spatial coordinates* are linearly related to the affine DOFs such that:

$$\mathbf{x}_i = \mathbf{J}_i \mathbf{q} = [\bar{\mathbf{x}}_i^\top \mathbf{I}_3, \mathbf{I}_3] \mathbf{q} \Leftrightarrow \mathbf{x} = \underbrace{[\mathbf{J}_1^\top, \dots, \mathbf{J}_N^\top]^\top}_{\mathbf{J}} \mathbf{q} \quad (4)$$

where $\otimes$ denotes the Kronecker product, and the generalized affine coordinate $\mathbf{q} = [\text{vec}^\top(\mathbf{A}), \mathbf{t}^\top]^\top \in \mathbb{R}^{12}$ is constructed by stacking the three column vectors of $\mathbf{A}$ and the translation vector $\mathbf{t}$. $\mathbf{J}$ is the coordinate Jacobian and is constant when the rest-shape model is given.

At the other end of the spectrum, ABD needs six more DOFs than RBD, which doubles the size of the system matrix. At each time step, we need to assemble and factorize a 12×12 system matrix. This is notably slower than an RBD system. Therefore, it is often believed that the multi-ABD framework is less scalable than the multi-RBD framework, especially when there exists a large number of interconnected stiff objects/links.

3.2 ABD Dynamics

With the N -node discretization introduced above, we can formulate its equilibrium at each time step as a variational optimization of:

$$\mathbf{x}^{n+1} = \arg \min_{\mathbf{x}} E(\mathbf{x}), \quad E(\mathbf{x}) = E_I(\mathbf{x}) + \Psi(\mathbf{x}), \quad (5)$$

where the subscript $n+1$ indicates the time step index. E_I is the inertia potential defined as:

$$E_I(\mathbf{x}) = \frac{1}{2h^2} (\mathbf{x} - \hat{\mathbf{x}})^\top \mathbf{M} (\mathbf{x} - \hat{\mathbf{x}}).$$

h is the time step size. $\mathbf{M}$ is the $3N \times 3N$ nodal mass matrix. $\hat{\mathbf{x}}$ is a known vector depending on the kinematic state of the system from the previous time step. For instance if the implicit Euler is used, $\hat{\mathbf{x}} = \mathbf{x}^n + h\dot{\mathbf{x}}^n + h^2\mathbf{M}^{-1}\mathbf{f}_{ext}$ with $\mathbf{f}_{ext}$ being the external force.

We use Newton's method to solve the optimization of Eq. (5), leading to a linear system of:

$$\mathbf{H}\delta\mathbf{x} = -\mathbf{g}, \quad (6)$$

where

$$\mathbf{H} = \frac{\partial^2 E(\mathbf{x})}{\partial^2 \mathbf{x}} = \frac{1}{h^2} \mathbf{M} + \mathbf{K}(\mathbf{x}), \quad \mathbf{g} = \frac{\partial E(\mathbf{x})}{\partial \mathbf{x}}$$

are the Hessian and gradient of the total variational energy. $\mathbf{K}(\mathbf{x}) = \partial^2 \Psi / \partial \mathbf{x}^2$ is a.k.a. the tangent stiffness matrix.

We can then project Eq. (6) into the column space of $\mathbf{J}$:

$$\mathbf{H}_A \delta \mathbf{q} = -\mathbf{J}^\top \mathbf{g}, \quad \mathbf{H}_A = \frac{1}{h^2} \mathbf{M}_A + \mathbf{K}_A, \quad (7)$$

where $\mathbf{M}_A = \mathbf{J}^\top \mathbf{M} \mathbf{J}$ is the generalized mass matrix. Unlike its counterpart in RBD, $\mathbf{M}_A$ is constant and is pre-computed. The generalized

stiffness matrix $\mathbf{K}_A = \mathbf{J}^\top \mathbf{K}(\mathbf{x}) \mathbf{J}$ represents the main computation hurdle. Since *any* rotation-invariant energy is nonlinear w.r.t. $\mathbf{x}$ (and thus w.r.t. $\mathbf{q}$), we need to re-assemble $\mathbf{K}$ at each Newton iteration, and project the resulting $\mathbf{K}$ to the affine subspace.

3.3 A Co-rotated Formula

Eq. (7) solves a 12×12 linear system at each Newton step. Its assembly and solve require a larger system than the 6×6 RBD counterpart, and the system matrix must be rebuilt at every iteration due to the nonlinear material model. We show that it is possible to completely resolve the limitation and have ABD reclaim a performance lead over RBD via a co-rotated formulation.

Both RBD and ABD are nonlinear. Nevertheless, their sources of the nonlinearity are drastically different. Specifically, the nonlinearity of RBD lies in the map between the rigid coordinate and nodal positions. In contrast, the nonlinearity of ABD stems from the geometric nonlinearity of the stiff material i.e., the requirement of being rotation-invariant. In a multibody system, we focus on highly stiff objects, meaning the actual deformation of an affine body is small or nearly vanished. The small deformation prior plays a critical role: it suggests the strain-stress relations among different energy models are indistinguishable in the *material space*, which is nearly identical to the one at the rest shape. Mathematically, it means:

$$\begin{aligned} \mathbf{K}(\mathbf{x})\delta\mathbf{x} &\approx \text{diag}_N(\mathbf{R})\tilde{\mathbf{K}}(\text{diag}_N(\mathbf{R}^\top)\delta\mathbf{x}) \\ &= \underbrace{(\text{diag}_N(\mathbf{R})\tilde{\mathbf{K}}\text{diag}_N(\mathbf{R}^\top))}_{\tilde{\mathbf{K}}} \delta\mathbf{x}, \end{aligned} \quad (8)$$

where $\tilde{\mathbf{K}}$ is the stiffness matrix evaluated at the rest shape, and $\text{diag}_N(\mathbf{R}) \in \mathbb{R}^{3N \times 3N}$ is a block-diagonal matrix of N copies of $\mathbf{R}(\mathbf{q}) \in \mathbb{R}^{3 \times 3}$. $\mathbf{R}(\mathbf{q})$ is the rotation matrix relating the object's rest-shape orientation and current orientation. It can be conveniently computed by applying the polar decomposition based on the current affine coordinate $\mathbf{q}$:

$$\mathbf{R}(\mathbf{q}) = \mathbf{A}(\mathbf{A}^\top \mathbf{A})^{-\frac{1}{2}}, \quad \mathbf{A}(\mathbf{q}) = \text{vec}^{-1}([\mathbf{I}_9, \mathbf{0}_{9 \times 3}]\mathbf{q}). \quad (9)$$

$\tilde{\mathbf{K}}$ is a high-quality approximation of the ground-truth $\mathbf{K}(\mathbf{x})$. The approximation error in Eq. (8) vanishes as $\mathbf{A} \rightarrow \mathbf{R}$. Under the close-to-rigid assumption, *this error is virtually a numeric zero in practice*.

For any rotation $\mathbf{R}$, it is easy to see $\mathbf{R}\mathbf{x}_i = (\mathbf{R}\mathbf{A})\bar{\mathbf{x}}_i + \mathbf{R}\mathbf{t}$ by left-multiplying $\mathbf{R}$ at both sides of Eq. (1). This leads to:

$$\mathbf{x} = \mathbf{J}\mathbf{q} \Rightarrow \text{diag}_N(\mathbf{R})\mathbf{x} = \text{diag}_N(\mathbf{R})\mathbf{J}\mathbf{q} = \mathbf{J}\text{diag}_4(\mathbf{R})\mathbf{q},$$

which holds for any $\mathbf{q}$, and we therefore have:

$$\text{diag}_N(\mathbf{R})\mathbf{J} = \mathbf{J}\text{diag}_4(\mathbf{R}). \quad (10)$$

Eq. (10) reveals an important feature of ABD: the spatial coordinate $\mathbf{x}$ and generalized coordinate $\mathbf{q}$ co-rotate i.e., with $\text{diag}_N(\mathbf{R})$ and $\text{diag}_4(\mathbf{R})$ under the same global rotation.

Bearing this property in mind, the reduced stiffness matrix can be simplified as:

$$\begin{aligned} K_A &= J^\top \tilde{K} J = J^\top \text{diag}_N(R) \tilde{K} \text{diag}_N(R^\top) J \\ &= \text{diag}_4(R) \underbrace{J^\top \tilde{K} J}_{\tilde{K}_A} \text{diag}_4(R^\top), \end{aligned} \quad (11)$$

where the second equality substitutes $\tilde{K}$ with Eq. (8), and the last equality applies the co-rotation property (Eq. (10)) to pass $\text{diag}_N(R)$ and $\text{diag}_N(R^\top)$ through $J^\top$ and J respectively. $\tilde{K}_A$ is the rest-shape generalized stiffness matrix, and it can be pre-computed.

Similarly, the generalized mass is also rotation-invariant. This follows from Eq. (10) together with the fact that a lumped mass matrix commutes with $\text{diag}_N(R)$, which is the standard choice for quasi-rigid bodies:

$$\begin{aligned} M_A &= J^\top M J = \text{diag}_4(R) J^\top \text{diag}_N(R^\top) M \text{diag}_N(R) J \text{diag}_4(R^\top) \\ &= \text{diag}_4(R) M_A \text{diag}_4(R^\top). \end{aligned} \quad (12)$$

Eq. (7) can be re-written as:

$$\text{diag}_4(R) \left(\frac{1}{h^2} M_A + \tilde{K}_A \right) \text{diag}_4(R^\top) \delta q = -J^\top g. \quad (13)$$

The r.h.s. of Eq. (13) consists of the inertia term and the elastic force. The former is straightforward; the latter integrates the elasticity gradient over each element e on the mesh to the affine coordinate such that:

$$\begin{aligned} -J^\top g &= \sum_e J_e^\top \frac{\partial \int_{\Omega_e} \Psi d\Omega_e}{\partial \mathbf{x}_e} = \sum_e V_e J_e^\top \left(\frac{\partial \Psi}{\partial F_e} : \frac{\partial F_e}{\partial \mathbf{x}_e} \right) \\ &= \left(\sum_e V_e J_e \right)^\top \left(\frac{\partial \Psi}{\partial A} : \frac{\partial F_e}{\partial \mathbf{x}_e} \right) = \bar{J}^\top f = f_A, \end{aligned} \quad (14)$$

where $F_e \in \mathbb{R}^{3 \times 3}$ is the deformation gradient of the element. $\partial F_e / \partial \mathbf{x}_e$ is a constant and sparse 3rd tensor for the linear tetrahedron (e.g., see [Kim and Eberle 2020]). For an ABD mesh, all elements share the same affine deformation, so $F_e = A$ and, assuming uniform material properties, $\partial \Psi / \partial F_e = \partial \Psi / \partial A$ is constant across all the elements. As a result, the double contraction can be moved out of the summation sign. $\bar{J} \in \mathbb{R}^{12 \times 12}$ is volume-weighted Jacobian, which can be pre-computed. It converts the aggregated spatial force $f \in \mathbb{R}^{12}$ to the affine coordinate f_A . Therefore, the complexity of computing the r.h.s. of Eq. (13) is independent of the mesh resolution.

Because of the small deformation prior, we opt for the linear material model for Ψ . The affine energy used in [Lan et al. 2022] i.e., Eq. (3) can also be used, but it fails to incorporate both Lamé parameters (μ, λ) . The simplification of linear elasticity gives:

$$\frac{\partial \Psi}{\partial A} = \mu(A + A^\top - 2I_3) + \lambda \text{tr}(A - I_3) I_3. \quad (15)$$

Linear elasticity offers the first-order approximation of most rotation-invariant hyperelastic models around the rest shape. Because each Newton iteration also linearizes the equation of motion, *the error induced by material simplification is hidden by the error induced by the Newton linearization*. As a result, the convergence of Newton's method is never impacted by material simplification.

3.4 Single-object ABD

The final system for a single-body ABD now becomes:

$$\bar{H}_A \text{diag}_4(R^\top) \delta q = \text{diag}_4(R^\top) f_A, \quad \bar{H}_A = \frac{1}{h^2} M_A + \tilde{K}_A. \quad (16)$$

Eq. (16) features a constant $\bar{H}_A$, which is pre-factorized. In other words, we avoid any matrix factorization, even with an implicit time integration. Note that this co-rotated formulation solves the same variational problem as the original ABD (Eq. (5)); it is a solver strategy rather than a physical approximation. We measure convergence against the vanilla ABD residual (i.e., the strong-form stress equilibrium), so upon convergence the solution is identical to that of standard Newton-based ABD. Because deformations of stiff bodies are inherently small, the method exhibits near-Newton convergence: in practice, the residual reaches 10^{-12} – 10^{-13} , matching double-precision machine epsilon.

With the system matrix pre-factorized, the polar decomposition that extracts the rotation component from the affine coordinate q i.e., Eq. (9) becomes a major computational hurdle. We want to remind that simulating ABD using Eq. (16) is already highly efficient and much faster than implicit RBD. Nevertheless, we can further push the performance by skipping Eq. (9) and replacing it with a simple normalization operation. As the object is highly stiff, $A \approx R$ is already close to a rotation matrix, when Ra rotates a vector a , instead of computing R out of $A(q)$, we only enforce $\|Aa\| = \|a\|$ to make sure A preserves the length of a such that:

$$Ra \approx \frac{\|a\|}{\|Aa\|} Aa. \quad (17)$$

By skipping the polar decomposition with Eq. (17), the ABD simulation procedure only involves lightweight level 1 and 2 BLAS operations, making implicit ABD outperform both implicit and explicit RBD. The pseudo code for one Newton iteration of a single object ABD is summarized in Alg. 1.

ALGORITHM 1: One Newton-step ABD integration

Input: $M_A, \tilde{K}_A, q = [q_1^\top, q_2^\top, q_3^\top, q_4^\top]^\top, \bar{J}, f, h$
Output: δq

- 1 $H_A \leftarrow \frac{1}{h^2} M_A + \tilde{K}_A$
- 2 $f_A = [f_{A,1}^\top, f_{A,2}^\top, f_{A,3}^\top, f_{A,4}^\top]^\top \leftarrow \bar{J}^\top f$
- 3 $A \leftarrow [q_1, q_2, q_3]$
- 4 **for** $k = 1$ **to** 4 **do**
- 5 $l_k^2 \leftarrow f_{A,k}^\top \cdot f_{A,k}$
- 6 $f_{A,k} \leftarrow A^\top f_{A,k}$
- 7 $f_{A,k} \leftarrow \sqrt{\frac{l_k^2}{f_{A,k}^\top \cdot f_{A,k}}} \cdot f_{A,k}$
- 8 **end**
- 9 Solve $\delta p = [\delta p_1^\top, \delta p_2^\top, \delta p_3^\top, \delta p_4^\top]^\top$ via $H_A \delta p = f_A$
- 10 **for** $k = 1$ **to** 4 **do**
- 11 $l_k^2 \leftarrow \delta p_k^\top \cdot \delta p_k$
- 12 $\delta p_k \leftarrow A \delta p_k$
- 13 $\delta q_k \leftarrow \sqrt{\frac{l_k^2}{\delta p_k^\top \cdot \delta p_k}} \cdot \delta p_k$
- 14 **end**




Fig. 2. **Performance comparison ABD vs. RBD.** We simulate a rolling cylinder for 10K steps and compare the total simulation using several variations of RBD and ABD. As visualized on the right, the vanilla implicit ABD is the least efficient (161 ms) as it needs to assemble and solve a 12×12 matrix for each iteration. It is much slower than implicit RBD (44 ms) and explicit RBD (32 ms). With our co-rotated formulation, the system matrix for implicit ABD is pre-factorized, which significantly reduces the computation time from 161 ms to 34 ms. By skipping polar decomposition, the performance of (implicit) ABD reaches 27 ms, which is over 20% faster than explicit RBD. The experiment is tested on an i7 CPU with single thread.

A simple benchmark test is reported in Fig. 2, where we simulate a rigid rolling cylinder for 10K step (with $h = 0.01$ sec) with ABD and RBD. The vanilla implicit ABD i.e., the original algorithm proposed in [Lan et al. 2022] uses 161 ms, while implicit and explicit RBD need 44 ms and 32 ms, respectively. Our co-rotated ABD (with polar decomposition) takes 34 ms, which is nearly as fast as explicit RBD already. Skipping polar decomposition further shortens the total computation time to 27 ms.

4 Multi-object ABD

The co-rotated formulation eliminates the major disadvantage of ABD. Based on this, we show how to incorporate different kinematic joints to assemble a multi-ABD system while maintaining computational efficiency. Without loss of generality, we consider basic cases where two affine bodies are constrained with a bilateral joint.

4.1 Control points and control tetrahedron of ABD




Fig. 3. **Control points and tetrahedron.** We can map the generalized coordinate of an affine body to any non-degenerate tetrahedron or the control tetrahedron. Its four corners are named control points of the affine body. Concatenating their positions gives the CP coordinate.

We re-parameterize the affine coordinate $\mathbf{q}$ as a collection of four *control points* (CPs), which forms a bijective map between $\mathbf{q}$ and CPs' spatial coordinates.

Let $\tilde{\mathbf{y}}_1, \tilde{\mathbf{y}}_2, \tilde{\mathbf{y}}_3, \tilde{\mathbf{y}}_4$ be the spatial coordinates of four nodes of a tetrahedron corresponding to the rest-shape state of an affine body, and we name this tetrahedron the control tetrahedron (CT). When an affine deformation $\mathbf{q} = [\text{vec}^T(\mathbf{A}), \mathbf{t}^T]^T$ occurs, the deformed positions of CPs are $\mathbf{y}_1, \mathbf{y}_2, \mathbf{y}_3$, and $\mathbf{y}_4$. They are concatenated as $\mathbf{y} = [\mathbf{y}_1^T, \mathbf{y}_2^T, \mathbf{y}_3^T, \mathbf{y}_4^T]^T \in \mathbb{R}^{12}$, which serves as another set of generalized coordinate or the *CP coordinate* of the affine body. It is easy to verify that:

$$\mathbf{y} = [\mathbf{Y}^T \otimes \mathbf{I}_3, \mathbf{1}_{4 \times 1} \otimes \mathbf{I}_3] \mathbf{q} = \mathbf{T} \mathbf{q}, \quad \mathbf{Y} = [\tilde{\mathbf{y}}_1, \tilde{\mathbf{y}}_2, \tilde{\mathbf{y}}_3, \tilde{\mathbf{y}}_4]. \quad (18)$$

As long as the CT is non-degenerate, we have $\mathbf{q} = \mathbf{T}^{-1} \mathbf{y}$, meaning the Jacobian matrix $\frac{\partial \mathbf{q}}{\partial \mathbf{y}} = \mathbf{T}^{-1}$ is constant.

The choice of CT is flexible. As shown in Fig. 4, the affine body can be fully enclosed by the CT, or partially overlap with it, or is completely separate from the CT. We then define joint constraints via the CP coordinate (instead of $\mathbf{q}$). Since any spatial point associated with an affine body is a linear function of $\mathbf{q}$ (Eq. (18)), each joint simply selects the points relevant to its constraint definition. When a body participates in multiple joints, each joint may choose its own set of control points. They are all linear functions of the same $\mathbf{q}$, so no special treatment is required.

4.2 Ball joint

A ball joint eliminates relative translational motion between two stiff bodies α and β , while permitting their rotations. Let $\mathbf{y}_k^\alpha$ and $\mathbf{y}_k^\beta$ be their CPs with $k = 1, 2, 3, 4$. We enforce the ball constraint by setting one of the CPs of each body, say $\mathbf{y}_1^\alpha$ and $\mathbf{y}_1^\beta$, to be at the prescribed location of the ball joint. It becomes a positional constraint on the CP coordinate:

$$\mathbf{S}_B^\alpha \mathbf{y}^\alpha - \mathbf{S}_B^\beta \mathbf{y}^\beta = \mathbf{0} \Leftrightarrow \mathbf{S}_B^\alpha \mathbf{T}^\alpha \mathbf{q}^\alpha - \mathbf{S}_B^\beta \mathbf{T}^\beta \mathbf{q}^\beta = \mathbf{0}, \quad (19)$$

where we use the superscript to denote the body index in the system. $\mathbf{S}_B \in \mathbb{R}^{3 \times 12}$ is a selection matrix picking the pertaining DOFs out of the CP coordinate $\mathbf{y}$. A ball joint is a 3-DOF constraint. As a result, two ball-joint-connected affine bodies possess $12 + 12 - 3 = 21$ DOFs in total.

4.3 Hinge joint

A hinge joint or a revolute joint permits one rotational DOF i.e., the rotation about a specified axis, while all translational motion and all other rotations are constrained.


The relaxation to affine DOFs makes the hinge joint formulation handy. After mapping $\mathbf{q}$ to $\mathbf{y}$, one can enforce a hinge joint by requiring two edges on the CTs of two stiff bodies to align with the prescribed rotation axis, with their ends overlapped. In other words, the hinge joint becomes the superposition of two ball joints, and the corresponding affine hinge is a 6-DOF constraint (as shown in Fig. 4). This strategy was used in previous ABD-based frameworks such as [Chen et al. 2022] as it is particularly implementation-friendly.

Alternatively, we propose a more compact encoding of a hinge joint, keeping its dimensionality the same as its RBD counterpart. As illustrated in Fig. 4, we compute a hinge rotation $\mathbf{R}_H$ via the Rodrigues formula, which aligns the prescribed rotation axis upright along the local $\mathbf{y}$ axis:

$$\mathbf{R}_H = \mathbf{I}_3 + [\mathbf{v}]_\times + \frac{[\mathbf{v}]_\times^2}{1 + a_y}, \quad (20)$$

where $\mathbf{a} = [a_x, a_y, a_z]^T$ is the unit rotation axis of the hinge. $\mathbf{v} = \mathbf{a} \times \mathbf{e}_y$, where $\mathbf{e}_y = [0, 1, 0]^T$.

$\mathbf{R}_H$ allows us to convert $\mathbf{y}$ to the local frame: $\tilde{\mathbf{y}} = \text{diag}_4(\mathbf{R}_H) \mathbf{y}$. With the local CP coordinate, the hinge constraint includes a ball constraint, fixing one CP from each body on the rotation axis. The other two constraints apply to the other CP from each body, which is also on the axis, but only at their local x and z coordinates. In other words,






Fig. 4. **Rotational joints.** The affine coordinate linearizes rotational constraints between affine bodies. By pre-parameterizing $\mathbf{q}$ to the CP coordinate $\mathbf{y}$, a ball joint becomes a point-point constraint; a hinge joint can be modeled as a 6-DOF edge-edge constraint; and a universal joint is the superposition of two hinge joints (with a virtual intermediate body). If we transfer the interconnected bodies into the local frame, the hinge joint can be enforced with 5 DOFs, and the universal joint only consumes 4 DOFs i.e., with the minimum DOFs needed.

it restricts the movement of the CP to be only along the rotation axis (i.e., see the inset). Unlike in RBD-based multibody formulation, the distance between two CPs is not explicitly preserved. Instead, the distance change will be strongly penalized by the material stiffness. Therefore, the hinge constraint can be written as a 5-dimensional system:

$$S_H^\alpha \text{diag}_4(R_H)T^\alpha \mathbf{q}^\alpha - S_H^\beta \text{diag}_4(R_H)T^\beta \mathbf{q}^\beta = \mathbf{0}, \quad (21)$$

where $S_H \in \mathbb{R}^{5 \times 12}$ is the corresponding selection matrix.

4.4 Universal joint


A universal joint constrains all translational motion and one rotational motion, while permitting rotation about two orthogonal axes. A convenient implementation of the universal joint may be treating it as the composition of two perpendicular hinge joints. As shown in Fig. 4, we can insert a virtual affine body whose CT has two orthogonal edges aligned with the two rotation axes of the universal joint. The universal constraint can then be enforced by adding two hinges at those two edges. This strategy needs 12 more primal DOFs for the added stiff body and consumes 12 DOFs. The total DOFs of the resulting system are $12 + 12 + 12 - 12 = 24$.

A compact formulation is also possible for the universal joint, which only needs a minimum set of 4 DOFs. Let $\mathbf{a}_1 \perp \mathbf{a}_2$ be a pair of orthogonal axes of the joint. We can build a rotation matrix $\mathbf{R}_U$ via:

$$\mathbf{R}_U = [\mathbf{a}_1 \times \mathbf{a}_2, \mathbf{a}_1, \mathbf{a}_2]^\top, \quad (22)$$

such that $\mathbf{R}_U \mathbf{a}_1$ and $\mathbf{R}_U \mathbf{a}_2$ align with the local y axis and z axis respectively. Within the local frame, the universal joint can be realized as the composition of a ball joint and another equality constraint to make sure $\mathbf{a}_2$ can only move in the plane normal to $\mathbf{a}_1$, whereas the length of $\mathbf{a}_2$ is maintained via the material stiffness.

Take the setup in Fig. 4 as an example. Since $\mathbf{R}_U$ is built to rotate $\mathbf{a}_1$ and $\mathbf{a}_2$ to be parallel to local y and z axes, the CP-CP constraint of $\tilde{\mathbf{y}}_2^\alpha = \tilde{\mathbf{y}}_2^\beta$ removes the relative translational movement between α and β . The additional equality constraint keeps the local y coordinates of the first and the second CPs of β identical i.e., $\tilde{y}_1^\beta = \tilde{y}_2^\beta$. This constraint restricts $\tilde{\mathbf{y}}_2^\beta$ to always reside in the local $\tilde{x}\tilde{z}$ plane and hence to be orthogonal to $\mathbf{a}_1$ (see the inset).



The formulation of a universal joint is then:

$$S_U^\alpha \text{diag}_4(\mathbf{R}_U)T^\alpha \mathbf{q}^\alpha - S_U^\beta \text{diag}_4(\mathbf{R}_U)T^\beta \mathbf{q}^\beta = \mathbf{0}. \quad (23)$$

Here, $S_U \in \mathbb{R}^{4 \times 12}$ only picks four relevant local CP coordinates.

4.5 Prismatic joint

When two bodies are connected by a prismatic joint, only relative translation along a prescribed axis is allowed. Similar to the hinge joint, we can find a rotation $\mathbf{R}_P$ such that the translational direction $\mathbf{a} = [a_x, a_y, a_z]^\top$ is parallel to the local y axis:

$$\mathbf{R}_P = \mathbf{I}_3 + [\mathbf{v}]_\times + \frac{[\mathbf{v}]_\times^2}{1 + a_y}. \quad (24)$$

A prismatic joint in the framework of ABD constrains a different set of local CT coordinates than a hinge joint. As shown in Fig. 5, we initialize CTs so that two edges align with the translational direction $\mathbf{a}$. After $\mathbf{R}_P$ is applied, we enforce the local $\tilde{x}$ and $\tilde{z}$ coordinates from $\tilde{\mathbf{y}}_2^\alpha$, $\tilde{\mathbf{y}}_1^\beta$, and $\tilde{\mathbf{y}}_2^\beta$ be the same. This leads to four constraint equations in $\tilde{\mathbf{y}}$ coordinate:

$\tilde{x}_2^\alpha - \tilde{x}_1^\beta = 0$, $\tilde{z}_2^\alpha - \tilde{z}_1^\beta = 0$, $\tilde{x}_1^\alpha - \tilde{x}_2^\beta = 0$, and $\tilde{z}_1^\alpha - \tilde{z}_2^\beta = 0$. The last constraint equation requires either $\tilde{\mathbf{y}}_4^\beta$ or $\tilde{\mathbf{y}}_3^\beta$ to stay with the current plane perpendicular to $\mathbf{a}$. To this end, we can simply fix $\tilde{x}$ or $\tilde{z}$ coordinate of any of them. In the figure, we set $\tilde{x}_3^\beta = 0$. Therefore, the constraint formulation of the prismatic joint is:

$$S_P^\alpha \text{diag}_4(\mathbf{R}_P)T^\alpha \mathbf{q}^\alpha - S_P^\beta \text{diag}_4(\mathbf{R}_P)T^\beta \mathbf{q}^\beta = \mathbf{0}. \quad (25)$$

The selection matrices $S_P \in \mathbb{R}^{5 \times 12}$ have different structures on α and β .

4.6 Linear or nonlinear

In the classic RBD-based multibody system, the constraint equations for all types of kinematic joints in RBD are nonlinear. We say a constraint type is linear or nonlinear, meaning its gradient w.r.t. the generalized coordinate used for simulation i.e., $\partial \mathbf{C}(\mathbf{q}) / \partial \mathbf{q}$ is constant or varying based on the current $\mathbf{q}$.




Fig. 5. **Prismatic joint.** A prismatic joint filters out-of-axis movements. The last constraint keeps β non-rotatable.

An ABD-based multibody system on the other hand, offers more flexibility. Because the CP coordinate $\mathbf{y}$ and the affine coordinate $\mathbf{q}$ are linearly related via Eq. (18). The ball joint is a linear constraint. Similarly, a 6-DOF hinge constraint and a 12-DOF universal constraint are also linear since they are defined on the (global) CP coordinate $\mathbf{y}$. This has been considered a major advantage of ABD. When solving the multibody system in a primal way, the constant constraint Jacobian eases the implementation. For instance, when approximating joint constraints with quadratic penalty energies [Chen et al. 2022], the penalty Hessian is constant. If we keep the rank of the hinge joint and universal joint minimal, the equality constraints are enforced at the local CP coordinate $\tilde{\mathbf{y}}$, which relates the global CP coordinate $\mathbf{y}$ via a varying rotation ($\mathbf{R}_H$, $\mathbf{R}_U$). That said, the 5-DOF hinge joint and the 4-DOF universal joint are nonlinear.

In our implementation, we opt for the nonlinear version of the hinge and universal joints as of Eqs. (21) and (23) for more compact representations. This choice is closely related to how we would like to solve the corresponding multibody system. As to be detailed in Sec. 5, we design a primal-dual formulation that solves the system in the dual space, exploiting the re-factorized primal matrix of an affine body. A compact joint representation yields a smaller-sized dual matrix.

Normally, relative motions between two joint-linked bodies cannot be arbitrarily large. We combine the strain-limiting [Provot et al. 1995] and explicit penalty to tackle unilateral limits. Specifically, when we observe an out-of-range joint DOF θ , we clamp it to its nearest valid setup $\theta \leftarrow \hat{\theta}$, and keep $\theta = \hat{\theta}$ for another Newton iteration while applying a dual-space penalty force as $k(\theta - \hat{\theta})$ at r.h.s. of Eq. (26). In other words, we try to avoid any implicit mechanism for inequality joint constraints so that we can fully exploit the pre-factorized per-body Hessian.

5 Dual-space KKT

We now consider an ABD-based multibody system, consisting of M bodies or links and K joints. The corresponding KKT system at each Newton step is:

$$\begin{bmatrix} \tilde{\mathbf{H}} & \nabla^T \tilde{\mathbf{C}} \\ \nabla \tilde{\mathbf{C}} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \delta \tilde{\mathbf{q}} \\ \delta \tilde{\boldsymbol{\lambda}} \end{bmatrix} = \begin{bmatrix} \tilde{\mathbf{f}}_A \\ \mathbf{0} \end{bmatrix}, \quad (26)$$

where $\tilde{\mathbf{H}} = \text{diag}_M(\mathbf{H}_A^j)$ is a block-diagonal matrix, whose 12×12 diagonal blocks are the reduced Hessian $\mathbf{H}_A^j$ i.e., Eq. (7). The superscript $(\cdot)^j$ is used to denote the body or joint index, and $(\cdot)$ suggests the variable is global i.e., it stacks the local variables from all the bodies or joints. For instance, $\delta \tilde{\mathbf{q}} \in \mathbb{R}^{12M}$ is the global primal unknown, and $\tilde{\mathbf{f}}_A \in \mathbb{R}^{12M}$ is the global primal r.h.s. They concatenate the regular affine coordinate $\mathbf{q}^j$ and the affine force $\mathbf{f}_A^j$ of each body. $\nabla \tilde{\mathbf{C}}$ is the generalized constraint gradient such that $\nabla \tilde{\mathbf{C}} = \partial \tilde{\mathbf{C}} / \partial \tilde{\mathbf{q}}$. $\delta \tilde{\boldsymbol{\lambda}}$ is the dual DOFs or the Lagrange multipliers. Let C_k be the number of constraint equations of each joint. The dimensionality of $\delta \tilde{\boldsymbol{\lambda}}$ is $\sum_{k=1}^K C_k$. In general, the r.h.s. of the constraint row should be $-\tilde{\mathbf{C}}(\mathbf{q}_n)$, the negative constraint residual from the previous iteration. Our method often converges in a single iteration, making this residual negligibly small; we therefore set it to $\mathbf{0}$ in Eq. (26).

When multiple iterations are needed, this term can be included at no additional cost.

Instead of using a minimal-coordinate formulation to solve Eq. (26), we eliminate the primal variable and solve the system in the dual space. The reason is twofold. First, $\sum_{k=1}^K C_k$ is often a much smaller number than $12M$, making the dual version of Eq. (26) a much smaller problem. Second, the dual-space solve better exploits the co-rotated single-body ABD solver of Eq. (16). Expanding the first line of Eq. (26) yields:

$$\tilde{\mathbf{H}} \delta \tilde{\mathbf{q}} + \nabla^T \tilde{\mathbf{C}} \delta \tilde{\boldsymbol{\lambda}} = \tilde{\mathbf{f}}_A \Rightarrow \delta \tilde{\mathbf{q}} = \tilde{\mathbf{H}}^{-1} (\tilde{\mathbf{f}}_A - \nabla^T \tilde{\mathbf{C}} \delta \tilde{\boldsymbol{\lambda}}). \quad (27)$$

Substituting Eq. (27) back to the second line of Eq. (26) leads to:

$$(\nabla \tilde{\mathbf{C}} \tilde{\mathbf{H}}^{-1} \nabla^T \tilde{\mathbf{C}}) \delta \tilde{\boldsymbol{\lambda}} = \nabla \tilde{\mathbf{C}} \tilde{\mathbf{H}}^{-1} \tilde{\mathbf{f}}_A. \quad (28)$$

On the r.h.s. of Eq. (28), $\tilde{\mathbf{H}}^{-1} \tilde{\mathbf{f}}_A$ represents the constraint-free single ABD solve for all the bodies. It can be efficiently handled with Eq. (16) i.e., Alg. 1. The result is then projected to the tangent space of the constraint. The dual matrix $\nabla \tilde{\mathbf{C}} \tilde{\mathbf{H}}^{-1} \nabla^T \tilde{\mathbf{C}}$ is block-sparse. $\nabla \tilde{\mathbf{C}}$ is also a sparse matrix as each joint only influences its incident affine bodies, and $\tilde{\mathbf{H}}^j$ is pre-factorized. Therefore, $\tilde{\mathbf{H}}^{-1} \nabla^T \tilde{\mathbf{C}}$ can be efficiently evaluated at each body. If all the joints are linear, meaning $\nabla \tilde{\mathbf{C}}$ is constant, computing $\tilde{\mathbf{H}}^{-1} \nabla^T \tilde{\mathbf{C}}$ using Eq. (16) still needs to rotate $\nabla \mathbf{C}^k$ into the local frame attached at each affine body based on $\mathbf{q}^j$. Therefore, we care less about the linearity of a joint constraint but prefer having its rank minimized.




Fig. 6. **Block-sparse dual matrix.** The constraint gradient matrix $\nabla \tilde{\mathbf{C}}$ is block-sparse, and the dual matrix is also block-sparse. An off-diagonal block is non-zero iff two joints are on the same affine body.

5.1 Constraint gradient

Except for the ball constraint whose gradient is constant, the hinge, universal, and prismatic constraints can be uniformly written as:

$$\mathbf{C}^k (\mathbf{R}_{Joint}(\mathbf{q}^\alpha, \mathbf{q}^\beta), \mathbf{q}^\alpha, \mathbf{q}^\beta) = 0, \quad (29)$$

where $\mathbf{R}_{Joint}$ refers to the necessary rotation i.e. $\mathbf{R}_H$, $\mathbf{R}_U$, or $\mathbf{R}_P$ in Eqs. (21), (23) and (25), that converts $\mathbf{y}$ to $\tilde{\mathbf{y}}$. Here, we only consider two affine bodies α and β connected by the k -th joint. The corresponding constraint gradient is:

$$\nabla \mathbf{C}^k = \frac{\partial \mathbf{C}^k}{\partial \mathbf{R}_{Joint}} : \frac{\partial \mathbf{R}_{Joint}}{\partial \mathbf{q}^\alpha} + \frac{\partial \mathbf{C}^k}{\partial \mathbf{R}_{Joint}} : \frac{\partial \mathbf{R}_{Joint}}{\partial \mathbf{q}^\beta} + \frac{\partial \mathbf{C}^k}{\partial \mathbf{q}^\alpha} + \frac{\partial \mathbf{C}^k}{\partial \mathbf{q}^\beta}. \quad (30)$$

Without loss of generality, we focus on the 5-DOF hinge constraint defined as Eq. (21). In Eq. (30), $\frac{\partial \mathbf{C}^k}{\partial \mathbf{q}^\alpha}$ and $\frac{\partial \mathbf{C}^k}{\partial \mathbf{q}^\beta}$ are easy, which treat $\mathbf{R}_{Joint}$ as constant:

$$\frac{\partial \mathbf{C}^k}{\partial \mathbf{q}^\alpha} = \mathbf{S}_H^\alpha \text{diag}_4(\mathbf{R}_H) \mathbf{T}^\alpha, \quad \frac{\partial \mathbf{C}^k}{\partial \mathbf{q}^\beta} = -\mathbf{S}_H^\beta \text{diag}_4(\mathbf{R}_H) \mathbf{T}^\beta. \quad (31)$$

Similarly, the constraint derivatives w.r.t. the rotation when treating $\mathbf{q}^\alpha$ and $\mathbf{q}^\beta$ as constants are:

$$\frac{\partial C^k}{\partial \text{vec}(\mathbf{R}_H)} = S_H^\alpha \left(\text{vec}^{-1}(\mathbf{T}^\alpha \mathbf{q}^\alpha)^\top \otimes \mathbf{I}_3 \right) - S_H^\beta \left(\text{vec}^{-1}(\mathbf{T}^\beta \mathbf{q}^\beta)^\top \otimes \mathbf{I}_3 \right), \quad (32)$$

where vec^{-1} reshapes $\mathbf{T}\mathbf{q}$ to a 3×4 matrix.

We also notice that $\mathbf{R}_H = \Delta \mathbf{R}_H^\alpha \mathbf{R}^{\alpha^\top}$ or $\mathbf{R}_H = \Delta \mathbf{R}_H^\beta \mathbf{R}^{\beta^\top}$, where $\Delta \mathbf{R}_H^\alpha$ (or $\Delta \mathbf{R}_H^\beta$) is the constant rotation that aligns a designated CT edge of body α (or β) with the hinge axis, and $\mathbf{R}^\alpha$ is the current rotation of body α . That said, $\mathbf{R}_H$ can always be considered as a body-wise rotation converting the regular CP coordinate $\mathbf{y}$ to the local affine frame followed by $\Delta \mathbf{R}_H^\alpha$ (or $\Delta \mathbf{R}_H^\beta$) to align the CT edge with the hinge axis. Both $\Delta \mathbf{R}_H^\alpha$ and $\Delta \mathbf{R}_H^\beta$ are constant since the CTs are set. The rotation gradient (in the vectorized form) is then:

$$\frac{\partial \text{vec}(\Delta \mathbf{R}_H \mathbf{R}^\top)}{\partial \mathbf{q}} = (\mathbf{I}_3 \otimes \Delta \mathbf{R}_H) \frac{\partial \text{vec}(\mathbf{R}^\top)}{\partial \mathbf{q}}. \quad (33)$$

Here, we further ignore the superscripts of α or β for better clarity.

An implicit (and accurate) constraint gradient evaluates the gradient of the rotation $\mathbf{R}$, which is cumbersome and often numerically unstable for an RBD-based multibody system. Fortunately, in ABD-based models, we always have $\mathbf{R} \approx \mathbf{A}$ and $\mathbf{q} = [\text{vec}^\top(\mathbf{A}), \mathbf{t}^\top]^\top$ becomes a linear function of $\mathbf{A}$. Therefore, Eq. (33) can be efficiently estimated as:

$$\frac{\partial \text{vec}(\Delta \mathbf{R}_H \mathbf{R}^\top)}{\partial \mathbf{q}} \approx (\mathbf{I}_3 \otimes \Delta \mathbf{R}_H) \frac{\partial \text{vec}(\mathbf{A}^\top)}{\partial \mathbf{q}} = (\mathbf{I}_3 \otimes \Delta \mathbf{R}_H) \tilde{\mathbf{I}}_{9 \times 12}. \quad (34)$$

$\tilde{\mathbf{I}}_{9 \times 12}$ is a constant permutation matrix, which re-indexes entries in $\text{vec}(\mathbf{A}^\top)$ to $\mathbf{q}$ with zero-padding at translation DOFs.

The approximate error of $\mathbf{R} \approx \mathbf{A}$ can also be easily eliminated by knowing the fact that $\frac{\partial \text{vec}(\Delta \mathbf{R}_H \mathbf{R}^\top)}{\partial \mathbf{q}}$ is the stacking of 12 (i.e., the dimensionality of $\mathbf{q}$) vectorized skew-symmetric matrices. Let $\mathbf{s}_\ell \in \mathbb{R}^9$ be the ℓ -th row of $\frac{\partial \text{vec}(\Delta \mathbf{R}_H \mathbf{R}^\top)}{\partial \mathbf{q}}$, we “skew-symmetrize” $\mathbf{s}_\ell$ by setting each of its elements as:

$$\begin{aligned} s_{\ell,1} &\leftarrow 0, & s_{\ell,5} &\leftarrow 0, & s_{\ell,9} &\leftarrow 0, & (s_{\ell,2}, s_{\ell,4}) &\leftarrow \pm \frac{1}{2}(|s_{\ell,2}| + |s_{\ell,4}|), \\ (s_{\ell,3}, s_{\ell,7}) &\leftarrow \pm \frac{1}{2}(|s_{\ell,3}| + |s_{\ell,7}|), & (s_{\ell,6}, s_{\ell,8}) &\leftarrow \pm \frac{1}{2}(|s_{\ell,6}| + |s_{\ell,8}|). \end{aligned} \quad (35)$$

The gradients of universal and prismatic joints can be obtained in a similar manner.

From the above derivation, it should be noted that in an ABD-based multibody system, $\tilde{\mathbf{C}}$ can be well approximated as a quadratic function of $\tilde{\mathbf{q}}$ since $\mathbf{R}^j$ is relaxed to $\mathbf{A}^j$. As a result, $\nabla \tilde{\mathbf{C}}$ becomes a linear function of $\tilde{\mathbf{q}}$ i.e., see Eq. (32), which lumps multiple skew-symmetric matrices of $\frac{\partial \mathbf{R}_{Joint}}{\partial \mathbf{q}^j}$ into the constraint’s tangent space.

In our implementation, we only keep three entries for each rotation gradient, and the corresponding computation is lightweight. Because $\tilde{\mathbf{H}}$ is pre-factorized, building the dual matrix is also fast. At this point, we can use any linear solver like Cholesky decomposition or sparse $LDL^\top$ factorization [Lacoursière 2007] to solve the dual matrix and convert the solution back to the primal variable via Eq. (27). Since the dual matrix is typically one-third or one-fourth the size of the primal system, we observe a significant speedup. Meanwhile, our framework enforces all the joints’ constraints via

KKT. Those constraints are well satisfied in most cases just using one Newton iteration.

Articulated systems are often man-made and well-structured, which endows the dual matrices with specialized patterns. Next, we show how to leverage some common structures to optimize the computation. Our goal is to achieve efficient multibody simulation, even on a single CPU thread.

5.2 Case I: Joint chain

Many multibody systems consist of a single kinematic chain (e.g., see Fig. 6), such as robotic arms. On such articulated structure, we have $M - K = 1$, and each joint couples exactly two bodies. Assume the k -th joint is C_k -DOF, and it couples bodies j and $j + 1$. On the kinematic chain, we can re-index joints and bodies such that $k = j$, and the corresponding constraint block becomes:

$$\nabla C_j^j \delta \mathbf{q}^j + \nabla C_{j+1}^j \delta \mathbf{q}^{j+1} = \mathbf{0}, \quad (36)$$

which leads to two dense blocks in the dual matrix:

$$\begin{aligned} \mathbf{D}^j &= \nabla C_j^j (\mathbf{H}_A^j)^{-1} \nabla^\top C_j^j + \nabla C_{j+1}^j (\mathbf{H}_A^{j+1})^{-1} \nabla^\top C_{j+1}^j \in \mathbb{R}^{C_j \times C_j}, \\ \mathbf{B}^j &= \nabla C_{j+1}^j (\mathbf{H}_A^{j+1})^{-1} \nabla^\top C_{j+1}^j \in \mathbb{R}^{C_j \times C_{j+1}}, \end{aligned} \quad (37)$$

and the r.h.s. dense block is:

$$\mathbf{b}^j = \nabla C_j^j (\mathbf{H}_A^j)^{-1} \mathbf{f}_A^j + \nabla C_{j+1}^j (\mathbf{H}_A^{j+1})^{-1} \mathbf{f}_A^{j+1} \in \mathbb{R}^{C_j}. \quad (38)$$

In other words, the dual matrix $\nabla \tilde{\mathbf{C}} \tilde{\mathbf{H}}^{-1} \nabla^\top \tilde{\mathbf{C}}$ becomes *block-tridiagonal* such that:

$$\begin{bmatrix} \mathbf{D}^1 & \mathbf{B}^1 & & \\ \mathbf{B}^{1^\top} & \mathbf{D}^2 & \ddots & \\ & \ddots & \ddots & \mathbf{B}^{K-1} \\ & & \mathbf{B}^{K-1^\top} & \mathbf{D}^K \end{bmatrix} \begin{bmatrix} \delta \lambda^1 \\ \delta \lambda^2 \\ \vdots \\ \delta \lambda^K \end{bmatrix} = \begin{bmatrix} \mathbf{b}^1 \\ \mathbf{b}^2 \\ \vdots \\ \mathbf{b}^K \end{bmatrix}. \quad (39)$$

It is known that such a block-tridiagonal system can be solved with $O(K)$ block-wise eliminations using a block Thomas algorithm [Press 2007]. $\{\delta \lambda_k\}_{k=1}^K$ are known. Specifically, for each body j , we have

$$\mathbf{H}_A^j \delta \mathbf{q}^j + \nabla^\top C_j^{j-1} \delta \lambda^{j-1} + \nabla^\top C_j^j \delta \lambda^j = \mathbf{f}_A^j. \quad (40)$$

Therefore, each body’s update boils down to the explicit form of:

$$\delta \mathbf{q}^j = \mathbf{H}_A^{j-1} \left(\mathbf{f}_A^j - \nabla^\top C_j^{j-1} \delta \lambda^{j-1} - \nabla^\top C_j^j \delta \lambda^j \right). \quad (41)$$

5.3 Case II: Joint tree (ABD-ABA)

A generalization of a chain of joints is a tree of joints — the whole articulated system forms a kinematic hierarchy such that each body has exactly one parent, and the system has exactly one root body. Such a structure allows us to solve the kinematic propagation level by level. This is known as the Articulated Body Algorithm (ABA) or Featherstone algorithm [Featherstone 2008]. We show that this strategy can be conveniently (and more efficiently) transplanted to an ABD-based multibody system, and we name the resulting algorithm ABD-ABA. The pseudo code is listed in Alg. 2. Similar to the classic ABA, the core idea is to treat each link as a pre-computed quadratic model in ABD increments and to propagate subtree contributions through a structured condensation, so that the dominant

ALGORITHM 2: ABD Articulated Body Algorithm (ABD-ABA)

Input: $\{H_A^j, q^j, \dot{q}^j, M_A^j, W_{ext}^j, S^j\}_{j=1}^M, h$
Output: $\{\delta q^j\}_{j=1}^M$

```

1 for  $j = 1$  to  $M$  do
2    $A^j \leftarrow A(q^j)$ 
3    $G^j \leftarrow G(A^j)$ 
4    $f_{A,ext}^j \leftarrow G^{j\top} W_{ext}^j$ 
5    $f_A^j \leftarrow f_{A,ext}^j + \frac{1}{h} M_A^j \dot{q}^j$ 
6    $\hat{H}_A^j \leftarrow H_A^j, \hat{f}_A^j \leftarrow f_A^j$ 
7 end
8 for  $link\ j$  (leaf-to-root) do
9    $p \leftarrow \text{parent}(j)$ 
10   $S_{abd}^j \leftarrow E(A^j) S^j$ 
11   $U^j \leftarrow \hat{H}_A^j S_{abd}^j, D^j \leftarrow S_{abd}^{j\top} U^j$ 
12  solve  $\alpha^j$  via  $D^j \alpha^j = S_{abd}^{j\top} \hat{f}_A^j$ 
13   $\Delta H_A^j \leftarrow \hat{H}_A^j - U^j (D^j)^{-1} U^{j\top}, \Delta f_A^j \leftarrow \hat{f}_A^j - U^j \alpha^j$ 
14   $R_{rel}^j \leftarrow A^{p\top} A^j$ 
15   $T^j \leftarrow \text{diag}_4(R_{rel}^{j\top})$ 
16   $\hat{H}_A^p \leftarrow \hat{H}_A^p + T^{j\top} \Delta H_A^j T^j, \hat{f}_A^p \leftarrow \hat{f}_A^p + T^{j\top} \Delta f_A^j$ 
17 end
18 for  $link\ j$  (root-to-leaf) do
19    $p \leftarrow \text{parent}(j)$ 
20    $r^j \leftarrow -\nabla C_p^j \delta q^p$ 
21   solve  $(\delta q^j, \delta \lambda^j)$  via Eq. (57)
22 end

```

cost depends on the kinematic connectivity rather than the full state dimension.

For the link j , we denote its velocity as $\dot{q}^j = [\dot{q}_1^{j\top}, \dot{q}_2^{j\top}, \dot{q}_3^{j\top}, \dot{t}^{j\top}]^\top$ and build a linear map $G(A^j) \in \mathbb{R}^{6 \times 12}$ that converts ABD velocity to the spatial twist $V^j = [\omega^{j\top}, v^{j\top}]^\top$:

$$G(A^j) = \begin{bmatrix} \frac{1}{2} [q_1^j]_\times & \frac{1}{2} [q_2^j]_\times & \frac{1}{2} [q_3^j]_\times & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & \mathbf{0}_{3 \times 3} & I_3 \end{bmatrix}, \quad V^j = G \dot{q}^j. \quad (42)$$

Equivalently, it yields:

$$v^j = \dot{t}^j, \quad \omega^j = \frac{1}{2} (q_1^j \times \dot{q}_1^j + q_2^j \times \dot{q}_2^j + q_3^j \times \dot{q}_3^j). \quad (43)$$

Given an external spatial wrench $W_{ext}^j = [\tau_{ext}^{j\top}, f_{ext}^{j\top}]^\top$, the corresponding affine force in ABD coordinates can be derived following the principle of virtual work i.e., $f_{A,ext}^{j\top} \dot{q}^j = W_{ext}^{j\top} V^j$ such that:

$$f_A^j = f_{A,ext}^j + \frac{1}{h} M_A^j \dot{q}^j, \quad f_{A,ext}^j = G(A^j)^\top W_{ext}^j. \quad (44)$$

To see why we do not need an explicit gyroscopic term, we start with the equation of motion of RBD:

$$W_{dyn}^j = I^j \ddot{V}^j + V^j \times^* (I^j V^j), \quad (45)$$

where $I^j \in \mathbb{R}^{6 \times 6}$ is the spatial inertia of the rigid body. $\times^*$ is the force cross operator. We apply time differentiation at both side of $V^j = G(A^j) \dot{q}^j$, which gives $\ddot{V}^j = G \ddot{q}^j + \dot{G} \dot{q}^j$. We can then project

Eq. (45) to the affine coordinate such that:

$$f_{A,dyn}^j = \underbrace{G^\top I^j G}_{M_A^j} \ddot{q}^j + \left(G^\top I^j \dot{G} \dot{q}^j + G^\top (V^j \times^* (I^j V^j)) \right). \quad (46)$$

On the other hand, the kinetic energy in ABD coordinates is $T^j = \frac{1}{2} \dot{q}^{j\top} M_A^j \dot{q}^j$. The Euler-Lagrange equation can be written as:

$$f_{A,dyn}^j = M_A^j \ddot{q}^j + \left(\dot{M}_A^j \dot{q}^j - \frac{\partial T^j}{\partial \dot{q}^j} \right). \quad (47)$$

From Eqs. (46) and (47), we can find the following correspondence:

$$G^\top I^j \dot{G} \dot{q}^j + G^\top (V^j \times^* (I^j V^j)) = \dot{M}_A^j \dot{q}^j - \frac{\partial T^j}{\partial \dot{q}^j}. \quad (48)$$

In ABD, M_A^j is constant and T^j does not depend on q^j , so $\dot{M}_A^j = 0$ and $\frac{\partial T^j}{\partial \dot{q}^j} = 0$. Therefore, we have:

$$G^\top I^j \dot{G} \dot{q}^j + G^\top (V^j \times^* (I^j V^j)) = 0. \quad (49)$$

That said, the projected spatial bias wrench exactly cancels the contribution from $\dot{G}$. Consequently, the gyroscopic term is not needed in ABD equations of motion.

During the upward condensation, each local matrix and force is accumulated sequentially such that i.e., $\hat{H}_A^j \leftarrow H_A^j$, and $\hat{f}_A^j \leftarrow f_A^j$. We first perform the leaf-to-root recursion within ABD coordinates. In standard Featherstone ABA, each joint k has a subspace matrix $S^k \in \mathbb{R}^{6 \times m_k}$, which maps joint velocity to the relative spatial twist by $V_{rel} = S^k \dot{\theta}^k$. Here, $m_k = 6 - C_k$ gives the joint's DOF count. To express the same admissible motion in with ABD coordinates, we use the rigid-motion embedding of $E(A^j) \in \mathbb{R}^{12 \times 6}$ such that:

$$E(A^j) = \begin{bmatrix} -[q_1^j]_\times & \mathbf{0}_{3 \times 3} \\ -[q_2^j]_\times & \mathbf{0}_{3 \times 3} \\ -[q_3^j]_\times & \mathbf{0}_{3 \times 3} \\ \mathbf{0}_{3 \times 3} & I_3 \end{bmatrix}, \quad \dot{q}^j = E(A^j) V^j. \quad (50)$$

This embedding is consistent with the twist map: for rigid motions, $V = G(A) \dot{q}$ and $G(A) E(A) = I_6$.

As all the bodies are tree-structured, we also have $M - K = 1$ and can re-index joints and bodies such that $k = j$. The ABD joint subspace simply becomes:

$$S_{abd}^j = E(A^j) S^j \in \mathbb{R}^{12 \times m_j}. \quad (51)$$

Given S^j , we eliminate the joint DOFs of link j via a condensation analogous to articulated inertia recursion. This is achieved by assembling $U^j = \hat{H}_A^j S_{abd}^j$, $D^j = S_{abd}^{j\top} U^j \in \mathbb{R}^{m_j \times m_j}$, and solving:

$$D^j \alpha^j = S_{abd}^{j\top} \hat{f}_A^j. \quad (52)$$

The condensed contributions of subtree j (with joint coordinates eliminated) are:

$$\Delta H_A^j = \hat{H}_A^j - U^j (D^j)^{-1} U^{j\top}, \quad \Delta f_A^j = \hat{f}_A^j - U^j \alpha^j. \quad (53)$$

Here, ΔH_A^j and Δf_A^j encode how the subtree rooted at j reacts to a parent-frame perturbation once the joint DOFs at (p, j) are internally resolved, where p is the index of the parent of body j .

Because our ABD increment stacks four 3-vectors (three columns of A plus t), and the relative rotation is $R_{rel} = A^{p\top} A^j$, we use the

block-diagonal transform $X^j = \text{diag}_4(R_{rel}^\top)$ to rotate the condensed contribution into the parent coordinates, and accumulate:

$$\hat{H}_A^p += X^{j^\top} \Delta H_A^j X^j, \quad \hat{f}_A^p += X^{j^\top} \Delta f_A^j. \quad (54)$$

After the upward pass, each link j stores an articulated (condensed) pair of $(\hat{H}_A^j, \hat{f}_A^j)$ that summarizes the effect of the entire subtree rooted at j . We then perform a root-to-leaf traversal to compute the per-link ABD increment δq^j while enforcing the joint constraints at each parent-child attachment.

Let $C^j(q^p, q^j) = 0$ be the joint constraint between the parent p and child j . Linearizing at the current configuration yields:

$$\nabla C_p^j \delta q^p + \nabla C_j^j \delta q^j = 0. \quad (55)$$

Given that δq^p is already known in the downward pass, we rewrite it as a linear equality in the child increment:

$$\nabla C_j^j \delta q^j = r^j, \quad r^j = -\nabla C_p^j \delta q^p, \quad (56)$$

where ∇C_p^j and ∇C_j^j are in the same coordinate frame as δq^j .

The increment δq^j is obtained by solving the local KKT of:

$$\begin{bmatrix} \hat{H}_A^j & \nabla C_j^{j^\top} \\ \nabla C_j^j & 0 \end{bmatrix} \begin{bmatrix} \delta q^j \\ \delta \lambda^j \end{bmatrix} = \begin{bmatrix} \hat{f}_A^j \\ r^j \end{bmatrix}. \quad (57)$$

Similar to its global counterpart i.e., Eq. (26), we eliminate the primal unknown δq^j , and solve for the multiplier first. As a result, the global nonlinear solve is decomposed into a series of lightweight joint-size local solves.

5.4 Case III: Joint loop

A joint loop is made by a chain of bodies whose first and last bodies are also constrained by a joint. In other words, it rolls back to a joint chain if we remove any affine body from the system.

Let M be the body we temporarily remove, and it is constrained by joints K and 1. We partition the global KKT of Eq. (26) into two blocks of DOFs. The first block contains all the primal DOFs from body $j = 1$ to $j = M - 1$, as well as all the dual DOFs from joints $k = 2$ to $k = K - 1$. The body M and pertaining joints K and 1 are in the other block such that:

$$\begin{bmatrix} \mathcal{A} & C^\top \\ C & \mathcal{D} \end{bmatrix} \begin{bmatrix} w_{\mathcal{A}} \\ w_{\mathcal{D}} \end{bmatrix} = \begin{bmatrix} b_{\mathcal{A}} \\ b_{\mathcal{D}} \end{bmatrix}, \quad (58)$$

where $\mathcal{A}$ represents the system of the joint chain, which can be converted to Eq. (39) and efficiently solved with block Thomas. Therefore, we solve Eq. (58) by computing the Schur complement $\mathcal{S} = \mathcal{D} - C\mathcal{A}^{-1}C^\top$, which is a low-rank system, and substituting the $w_{\mathcal{D}}$ back to the first line of Eq (58) to retrieve unknowns in the first block.

This idea can be further generalized for loops on a joint tree. In this case, $\mathcal{A}$ represents an ABD-ABA system. As long as the total number of loops is limited, we can always use Schur complement to solve the sub-system corresponding to DOFs at the “loop breakers”.

5.5 Case IV: Joint graph

The most generic situation is when the joint topology of the multi-body system forms a graph. In this case, we propose a multi-directional

Table 1. **Experiment statistics.** For each experiment, we report the number of links (**# Link**), the total number of constraint equations (**# Cons.**), the time step size h , the constraint residual tolerance $\|r\|$, the Young’s modulus E , and the per-step simulation time (**Sim.**) in milliseconds. At each time step, Newton iterations proceed until the constraint residual drops below $\|r\|$. Both **# Iter.** and **Sim.** are averaged over all time steps. In most scenes, convergence is reached in a single iteration.

Scene	# Link	# Cons.	h	$\ r\ $	# Iter.	E	Sim.
Joint net (Fig. 14: 100 × 100)	30K	120K	10 ms	1E-7	1	1E9	84 ms
A pulley system (Fig. 15)	1.5K	4.5K	10 ms	1E-7	1	1E8	2 ms
A huge pulley system (Fig. 1)	1M	3M	10 ms	1E-5	1	1E8	904 ms
Willow (Fig. 16: top)	21K	63K	10 ms	1E-6	1	1E8	18 ms
Pear (Fig. 16: bottom)	29K	87K	10 ms	1E-6	1	1E8	23 ms
Net cloak (Fig. 17)	12K	48K	10 ms	1E-7	1	1E9	33 ms
Armadillo (Fig. 18)	2.7K	10.8K	10 ms	1E-6	7.3	1E6	116 ms
Ragdolls (Fig. 19)	1.5K	6K	5 ms	1E-6	1	1E6	17 ms
Falling joints (Fig. 20)	1.4K	2.9K	10 ms	1E-6	1	1E9	54 ms
Protein (Fig. 22)	14K	56K	–	1E-7	1	1E6	14 ms

block Gauss-Seidel procedure if a direct matrix factorization is prohibitive in a single-thread simulation environment. We first decompose the joint graph into a set of overlapping chains. Each chain defines a block of dual variables that can be solved exactly using the block Thomas algorithm described in Sec. 5.2. One sweep of the block Gauss-Seidel method iterates over all chains sequentially: for each chain, we fix the dual variables of all other chains at their current values and solve the chain’s block tridiagonal system. The procedure repeats until the overall dual residual converges. This differs from a standard Gauss-Seidel iteration, which updates one dual variable (i.e., one joint) at a time. By treating an entire chain as a single block, our method propagates constraint corrections across many joints simultaneously, leading to faster convergence especially for elongated or branching structures.

6 Experimental Results

We implemented the M-ABD framework on a desktop equipped with an AMD 9950X3D CPU, using Eigen [Guennebaud et al. 2010] and MKL [Wang et al. 2014] for linear algebra operations. We specifically target a single-thread CPU implementation, as multibody simulations in tasks like large-scale data synthesis for robotic training are typically parallelized across multiple independent instances. For complex systems with massive joint networks, the dual solver is amenable to multi-core CPU or GPU parallelization, e.g., via Gauss-Seidel with graph-coloring. Nevertheless, this is not the scope of this paper. The statistics of the experimental setup and timing information is reported in Tab. 1. Please refer to the supplementary video for additional animation results.

6.1 Single-body comparison

We first quantitatively analyze the accuracy and reliability of the co-rotated ABD model for single stiff object simulation. We monitor the linear momentum, angular momentum, and total energy to validate that our method matches expected physical results under several standard benchmarks. Throughout this subsection, the “implicit RBD baseline” uses implicit Euler integration with the same time step as our method, while the “reference” curves are generated by




Fig. 7. **Spinning box.** We compare ABD (solid curves) against an implicit RBD baseline (dashed curves) for a cube with initial $\mathbf{p}_0 = [100, 0, 0]^\top \text{ kg m/s}$ and $\mathbf{L}_0 = [0, 100, 0]^\top \text{ kg m}^2/\text{s}$. Co-rotated ABD formulation closely matches the behavior of an implicit RBD.

a fourth-order Runge-Kutta (RK4) integrator at $h = 10^{-4} \text{ s}$ unless otherwise noted.

In Fig. 7, we compare our co-rotated ABD with a standard implicit RBD baseline using a single cube moving on a frictionless surface. The size of the cube is 0.1 m ($\rho = 10^3 \text{ kg/m}^3, E = 10^9 \text{ Pa}, \nu = 0.3$) with initial momenta $\mathbf{p}_0 = [100, 0, 0]^\top$ and $\mathbf{L}_0 = [0, 100, 0]^\top$. To initialize the motion, we compute the target spatial twist $\mathbf{V}_0 = [\omega_0^\top, v_0^\top]^\top$ and map it to ABD generalized velocities via Eq. (42). Similarly, external spatial wrenches $\mathbf{W}_{ext} = [\tau_{ext}^\top, f_{ext}^\top]^\top$ are converted to ABD generalized forces $f_{A,ext}$ using Eq. (44). Following the principle of virtual work, the equivalent spatial wrench can be obtained as:

$$\mathbf{W}_{ext} = \mathbf{E}(\mathbf{A})^\top f_{A,ext}. \quad (59)$$

The variations of the momentum are plotted in the figure. We observe that linear momentum is preserved for all step sizes, while angular momentum shows dissipation that increases with h . Our ABD method matches the RBD baseline, and the conservation improves as h decreases. This confirms that affine linearization does not violate the D'Alembert principle. The error mainly comes from the numerical damping of the time integrator.

In another benchmark shown in Fig. 8, we use a T-handle model to evaluate how well the co-rotated ABD formulation captures the coupling among rotational DOFs. In a zero-gravity environment, we initialize the body with a nonzero angular velocity aligned with its intermediate principal axis. According to the intermediate-axis theorem [Goldstein et al. 1950], the T-handle should undergo periodic flipping. We set the initial angular speed to $\omega_0 = 3 \text{ rad/s}$ and introduce a subtle asymmetry in the mass distribution. The motion is simulated using various time step sizes and compared against the reference. Similar to Fig. 9, the co-rotation ABD formula closely matches implicit RBD simulations. Yet, our method is over 30% faster.

In Fig. 9, we test a heavy top with a fixed pivot to see if our ABD method captures classic gyroscopic effects like precession and nutation. The motion is driven by gravity, with the top initialized with a 5° tilt and an angular speed of 10 rad/s . ABD accurately




Fig. 8. **T-handle.** Our co-rotated ABD integrator accurately reproduces the intermediate-axis instability per the intermediate-axis theorem. Left: We plot the body-frame angular velocity along the intermediate principal axis, which undergoes periodic sign flips. Right: We plot the energy loss over time, which increases during the rapid flipping events and remains nearly constant during the quasi-steady spinning phases between flips.




Fig. 9. **Heavy top.** We simulate a standard heavy top with a fixed pivot under gravity. The model is initialized with an angular speed of 10 rad/s and a 5° tilt. We plot the precession velocity and the nutation angle over time using our method and implicit RBD under different time step sizes. Our method captures the expected coupling between spin and nutation, and the slight inertia asymmetry induces small oscillations in the nutation angle.

reproduces the expected motion and converges to the reference as h decreases. Notably, ABD shows less sensitivity to the time step size than the implicit RBD baseline, which suffers from more significant waveform distortion at coarser steps.

In Fig. 10, we simulate a physical pendulum starting from a horizontal position. The pendulum swings freely under gravity, and we compare our ABD method and an implicit RBD baseline against the analytic solution based on elliptic integrals of:

$$\theta(t) = \frac{\pi}{2} - 2 \arcsin(\kappa \operatorname{sn}(K(\kappa) - \omega_{\text{lin}} t, \kappa)), \quad (60)$$

where $K(\kappa)$ is the complete elliptic integral of the first kind and $\operatorname{sn}(\cdot, \kappa)$ is the Jacobi elliptic sine. Our method tracks the analytic reference more closely than the implicit RBD baseline and provides more stable results for both the motion and joint forces.




Fig. 10. **Physical pendulum.** We simulate a physical pendulum with a fixed pivot, released from a horizontal configuration with zero initial velocity under gravity. Top: We plot the pendulum angle over time, where the reference curve is from the elliptic-integral (Eq. 60). Bottom: We plot the magnitude of the joint force over time, which varies periodically and peaks near the turning points of the swing. As the time step decreases, both the angle trajectory and the joint-force waveform produced by ABD better match the reference. Larger time steps lead to accumulated phase drift.

6.2 Multi-body comparison

We now test our method on various multibody systems, focusing on constraint satisfaction and energy loss. Our primary competitors include MuJoCo [Todorov et al. 2012], Bullet [Coumans 2015], and PhysX [NVIDIA Corporation 2025] which are popular off-the-shelf articulated body simulation solutions. Another important peer technique is quaternion-based multibody system or (VQ) [Maloisel et al. 2025]. VQ uses quaternion to parameterize the rotational motion of a rigid body, and it solves the full articulated dynamics with Newton’s method. VQ also exactly enforces the joint constraints as our method. Nevertheless, this method is less scalable and becomes prohibitive as the total number of bodies increases.

Fig. 11 shows a 163-link chain with a heavy end mass released from a coiled configuration. The mass ratio between the end mass and a single link is 352:1. As the chain drops and straightens, it creates a massive impulsive load. We compare our results with two MuJoCo setups using either equality constraints or articulated joints. The time step size is $h = 10^{-3}$ s for all tests. Under sharp external impulses, MuJoCo’s equality constraints show visible joint separation as seen in the zoomed view, while its articulated model becomes unstable and fails with a NaN error. In contrast, our method remains stable using only one iteration with all joint constraints exactly satisfied. The affine DOFs facilitate the conversion of impulse into elastic potential energy and eventually into kinetic energy. As a result, our method captures the natural vertical oscillations after the chain reaches full tension. We note that our method remains stable and produces the same result when $h = 10^{-2}$ s. In the example, PhysX and Bullet diverge. AVBD [Giles et al. 2025] avoids divergence with 30 iterations, but the chain exhibits jittering that cannot be eliminated even with increased damping. This is because using an augmented Lagrangian approach to enforce constraints




Fig. 11. **Chain with a heavy end mass.** We compare our method against two MuJoCo configurations (equality constraint and articulated joint) and AVBD [Giles et al. 2025]. When $h = 10^{-3}$ s, MuJoCo’s equality constraints exhibit joint drift and its articulated model fails to converge. Bullet and PhysX both crash in this experiment. AVBD (30 iterations) avoids divergence but exhibits jittering after the chain settles. Our method remains stable and naturally captures the elastic vibrations after the chain straightens even for $h = 10^{-2}$ s.

can introduce artificial stiffness when the dual variables (Lagrange multipliers) are updated alongside penalty terms.

Fig. 12 shows a 10×10 network of ball joints consisting of 280 links. We use a relatively aggressive time step of $h = 1/30$ s to test robustness among various multibody frameworks. In this test, our method uses only one iteration per step, while PhysX uses 10, and MuJoCo (equality constraint)/Bullet use 30. Our method exactly enforces all joint constraints, preserving the net’s geometry under gravity. On the other hand, all three other methods show visible gaps at the joints, indicating inaccurate constraint enforcement under these settings. VQ [Maloisel et al. 2025] solves the global system implicitly using Newton’s method. It also yields high-quality results. Nevertheless, VQ is considerably slower for two reasons. First, it solves a full primal-dual KKT system involving both body and constraint DOFs, whereas our approach projects onto a compact dual system (Eq. (28)) scaled strictly to the number of constraint DOFs. Second, VQ requires a costly reassembly and refactorization of the system matrix at every Newton iteration. In contrast, our method pre-factorizes the co-rotated $\bar{H}$, reducing the dual system bottleneck to inexpensive back-substitutions. While both methods employ a direct solver, ours minimizes the per-iteration cost by avoiding redundant global factorizations. In this example, our method uses less than 1 ms to finish one time step, while VQ needs 27 ms. VQ quickly becomes prohibitive for large-scale instances such as the one shown in Figs. 1 or 14.

In Fig. 13, we put a cylinder collider under the ball-joint net to test contact-involved simulations. We use the same time step and iteration budgets as in Fig. 12. As the simulation becomes more complex, both MuJoCo and Bullet fail to converge, and the links become scattered. PhysX does not diverge, we can still see the joint gaps from the result. Our method remains stable with a single iteration — the ball-joint net drapes on the collider naturally. We




Fig. 12. **Ball-joint net I.** A 10×10 net (280 links) is pinned along one boundary and sags under gravity with $h = 1/30$ s. Our method uses a single iteration per step and produces an expected result. MuJoCo (30 iterations, equality constraint), Bullet (30 iterations), and PhysX (10 iterations) all fail to enforce the constraints and produce noticeable artifacts. VQ solves the global system using Newton’s method. As a result, it also produces high-quality result with all the constraints exactly satisfied. Nevertheless, it is much slower than our method. Our method uses less than 1 ms for one step while VQ needs nearly 30 ms.




Fig. 13. **Ball-joint net II.** With an external collider, our method still performs with a single iteration per time step ($h = 1/30$ s. All other three methods fail to produce correct results.)

further test the scalability of our method. As shown in Fig. 14, we increase the size of the ball-joint network to 20×20 , 50×50 , and 100×100 , respectively. We set $h = 10^{-2}$ s and use a single iteration per step for all cases. Despite the increased complexity, our method remains stable and produces coherent draping without any instabilities. All the other competitors fail to complete this example.

6.3 Complex multi-ABD systems

Our next set of experiments focuses on evaluating the proposed multi-ABD algorithm using the dual KKT solve under complex articulations and a large number of links.

Fig. 14. **Ball-joint net III.** We increase the sizes of the net to 20×20 , 50×50 , and 100×100 and simulate each case with a fixed time step size of $h = 10^{-2}$ s using one iteration. Our method remains stable while all the existing methods fail to handle such large-scale simulations.

6.3.1 Joint chain. As discussed in Sec. 5.2, our framework features a block-tridiagonal structure if an articulated system forms a chain of joints. This enables a linear-time solve only with one CPU thread. To validate this advantage, we showcase simulations of two pulley systems connected by ball joints as shown in Figs. 15 and 1. We set the time step size as $h = 10^{-2}$ s. Similar to Fig. 14, MuJoCo, PhysX, and Bullet all fail in these examples (even with a highly conservative setting with $h = 10^{-4}$ s).

6.3.2 Joint graph. As the connectivity among joints becomes more complex, we need to solve the dual matrix of Eq. (28). If the system size is moderate, a direct solver could be used. Otherwise, we can use a multi-directional GS procedure to solve along multiple joint chains. Fig. 15 shows a lifting-and-release cycle of a pulley system




Fig. 15. **A pulley system.** We simulate a lifting-and-release cycle of a pulley system with a chain with 1.5K joints. Our method stably handles the rapid tensioning as the heavy anchor rises and captures the inertial oscillations of the pulleys during the subsequent release. The time step size is $h = 0.01$ s




Fig. 16. **Tree as a multibody system.** We simulate a willow tree (21K links, top) and a pear tree (29K links, bottom) under winds. Our ABD-ABA framework handles these branched structures with a single iteration per step, requiring approximately 20 ms of computation time per frame on a single CPU thread. The resulting animation is high-quality and realistic.

consisting of 1.5K joints. The major portion of DOFs is on the long chain linked to a heavy anchor. During loading, the chain tightens and the anchor rises as tension moves through the pulleys. Once the drive is released, the anchor falls and the chain slides back. A brief back-and-forth motion occurs during unloading, caused by the redistribution of tension and the pulleys' inertia. Fig. 1 shows a stress test of a similar pulley setup, scaled up to 1,076,748 links. Our method remains stable with $h = 10^{-2}$ s and only needs one iteration for each step. Meanwhile, all the joint constraints are




Fig. 17. **A multi-ABD cloak.** A “cloak” is made of 11.7K links connected with joints. It interacts with a moving avatar at $h = 10^{-2}$ s. Our method captures the rich, garment-like dynamics of the structure while maintaining stable constraint enforcement and contact handling using only one iteration per step. MuJoCo, Bullet, and PhysX all fail in this example.

exactly enforced. Running at a single CPU thread, our method only takes 904 ms on average to simulate one step.

6.3.3 Joint hierarchy. Our ABD-ABA framework generalizes joint chain systems to multi-level topologies, allowing it to handle large-scale, complex joint hierarchies (Sec. 5.3). Fig. 16 showcases this with a willow tree (21K links) and a pear tree (29K links). To build these models, we extract skeletal graphs from the tree geometries (as highlighted in the figure) and organize the links by branching depth for efficient traversal. Each joint includes penalty energies to model real-world branch compliance. A prescribed wind field applies aerodynamic loads, which are accumulated as external wrenches and mapped to generalized forces in ABD coordinates. Our solver advances the state at $h = 10^{-2}$ s using only a single iteration per step.




Fig. 18. **Armadillo**. We simulate the interaction between a multi-ABD net and a Neo-Hookean Armadillo. Our method handles the strong coupling between articulated constraints and volumetric elasticity. With CP coordinate, ABD-based system can be seamlessly coupled with FEM models.




Fig. 19. **Ragdolls**. We drop 27 ragdolls onto a stiff (with Young's modulus being 10^6) joint net. Our method gives stable results under high-speed collisions. In this experiment, we have time step size $h = 1/200$ s, and our simulation uses one iteration for each step.

The simulation takes approximately 20 ms per step for these models, which provides an alternative approach to real-time simulation of complex plants.

A representative example is shown in Fig. 17, where a cloak-like jointed net interacts with a moving avatar. The cloak consists of 11,750 links that slide frictionlessly over the body. As in previous examples, we set $h = 10^{-2}$ s and use only one iteration per step. Despite being a collection of rigid links and joints, the entire structure behaves like a flexible garment, producing rich dynamic effects as it follows the body's movement. Notably, none of the other simulators i.e., MuJoCo, Bullet, or PhysX manage to run this example, whereas our method remains stable throughout the sequence. Fig. 19 gives another example, where 27 ragdolls fall onto a corner-pinned ball-joint net (1,160 links). Each ragdoll is an 11-link articulated body with joint penalties for bending and twisting. We use a reduced Young's modulus of 10^6 to allow small deflections of the network. The experiment involves dense impacts and changing contact topologies between the ragdolls and the net, as well as among the ragdolls themselves. We use a smaller time step size of $h = 1/200$ s to facilitate the collision detection of fast-movement ragdolls. Our method uses 33 ms to process one step.

6.4 More versatile results

ABD formulation shares the same underlying framework as FEM (finite element method). The control tetrahedron of an affine body essentially makes it a single-element model. Therefore, ABD-based multibody systems naturally integrate with deformable body simulations. Fig. 18 demonstrates the coupling between a multi-ABD net and a Neo-Hookean Armadillo. In this example, we drop the Armadillo onto the net and then twist the boundary anchors to induce shear and torsion. This example highlights the stability of our solver under strong coupling between articulated constraints and volumetric elasticity.

Fig. 20 shows snapshots of another collision-heavy scene. In this experiment, we drop 720 jointed pairs into a bowl container, including 180 ball joints, 180 universal joints, 180 hinge joints, and 180 prismatic joints. All the constraints are exactly satisfied during the simulation. Collisions among articulated bodies and with the collider are handled by an implicit penalty method: contact potentials are added to the primal energy and resolved together with joint constraints within the same Newton iteration. Our framework is also compatible with barrier-based methods such as IPC [Li et al. 2020].

A motivation for our ABD-based multibody solver is to provide an efficient and robust algorithm for Embodied AI training and testing. In many large-scale or real-time applications, each simulation scene is allocated with limited computational resources. Our framework is designed to maintain physical stability and constraint integrity even under these constraints, using minimal solver iterations and large time steps. This robustness allows for the reliable evaluation of agents across diverse, contact-rich environments without the numerical failures that often plague traditional simulators in resource-constrained scenarios. To this end, Fig. 21 showcases this efficiency in a robotics context. We integrate our ABD framework with inverse kinematics (IK) and motion planning to drive a Franka Panda manipulator through pick-and-place tasks. Our solver handles joint actuation and contact interactions between the gripper and objects using a single iteration per step. The results show that the arm reliably executes grasping and transport while objects remain stable during stacking and settling.




Fig. 20. **Falling joints.** We drop 720 articulated bodies connected by different joint types, including 180 ball joints, 180 universal joints, 180 hinge joints, and 180 prismatic joints. Our solver robustly maintains all joint configurations under dense stacking and repeated impacts.




Fig. 21. **Multi-body ABD for Embodied AI.** We integrate our ABD solver into a standard robotics pipeline to perform pick-and-place tasks. Even when allocated limited computational resources, specifically a single iteration per step, our method maintains the joint integrity and contact stability required for training and testing Embodied AI agents. The solver robustly handles the transition from planned motion to contact-rich interaction as objects are grasped, transported, and stacked.

Fig. 22 demonstrates a cross-disciplinary application of our framework in structural biology. A polypeptide backbone can be modeled as an articulated chain because its flexibility is primarily determined by a subset of rotatable bonds. We represent the backbone as a sequence of rigid units connected by rotatable dihedrals ϕ and ψ . In this example, the original protein dynamics sequence is a 40 ns trajectory of the SARS-CoV-2 Spike glycoprotein (PDB: 6VXX), used to generate temporally sparse observations; for biological context on SARS-CoV-2 recognition by full-length human ACE2, see Walls et al. [2020]; Yan et al. [2020]. To reconstruct the motion of a viral protein undergoing a large-scale conformational change, we

solve for IK keyframes from sparse observations and roll out multi-ABD dynamics to recover a continuous trajectory. This procedure produces a smooth animation that fills the gaps between temporally sparse markers. The recovered motion maintains the folded topology and structural coherence without any discontinuities or numerical fragmentation. This result illustrates that the proposed ABD-based multibody framework is a versatile tool for simulating complex articulated systems across different scientific domains, including the study of large-scale structural transitions in biological macromolecules.




Fig. 22. **Protein unfolding.** A polypeptide backbone is modeled as an articulated chain of rigid links connected by ball joints parameterized by the dihedral angles ϕ and ψ , which yields a rigid-link abstraction of backbone motion. Given a sparse time sequence of marker observations $\{\mathbf{y}_t\}$ (e.g., C_α positions), we solve an IK problem to recover generalized coordinates $\mathbf{q}_t$ at the observed frames, and we then roll out the ABD dynamics between these IK keyframes to obtain a temporally continuous backbone motion. Left: rendered initial protein structure with a zoom-in illustrating the locally rigid backbone geometry used to define each link. Middle and right: backbone representations at the initial and final frames of the reconstructed sequence. A region that undergoes a large conformational change between the initial and final states is highlighted.

7 Conclusion & Limitations

This paper presents an efficient framework for simulating large-scale articulated systems. By separating geometric nonlinearity with a co-rotational formulation, we turn the ABD problem into a constant-matrix system, which can be pre-factorized even with implicit integration. Our analysis of joint constraints in the affine coordinate space provides both linear and nonlinear constraint models, which allows the simulation to capture complex mechanical interactions accurately. To handle different mechanical structures, we developed a collection of specialized solvers, including a block-tridiagonal solver, a generalized Featherstone algorithm, a low-dimensional approximation for kinematic loops, and a bidirectional Gauss-Seidel optimizer for dense joint networks. By projecting the problem into a minimal-DOF dual space, it is possible to solve the KKT on a single-threaded CPU at a scale not previously achieved. We believe our method offers a novel and reliable solution for many demanding tasks in embodied AI, mechanical design, and interactive graphics.

For performance, we adopt a co-rotated formulation with a pre-factorized system matrix and project onto a compact dual space to minimize the problem size. This is complemented by topology-aware solvers—block Thomas, ABD-ABA, and block Gauss-Seidel—that significantly reduce per-iteration overhead. For robustness, we employ KKT-based constraint enforcement and implicit integration, which prevents joint drift and maintains stability even at large time steps. Serving both goals, our linear coordinate mapping for ABD streamlines the constraint Jacobian for faster computation while eliminating the singularities of rotational parameterizations.

Our method also has limitations. While all the equality constraints induced by joints are exactly and robustly processed with KKT, inequality constraints remain a challenge. If two bodies are about to collide with each other, it is possible to employ collision barriers [Li et al. 2020] to resolve the collision. However, doing so puts nonlinear

off-diagonal entries in the primal matrix, and our pre-factorized per-body solve is less helpful. In order to fully exploit the co-rotated ABD model, we plan to generalize the proposed primal-dual algorithm for inequality constraints. This strategy will need to pivot active primal DOFs, and the key problem is how to quickly identify the right set of DOFs to be activated (i.e., converting to equality constraints). Our method solves the dual matrix using Newton’s method on the CPU. It is also possible to use GPU-based solvers to further accelerate the computation. This is particularly useful for large-scale molecular simulations such as the protein example in Fig. 22.

Acknowledgments

We thank the anonymous reviewers for their constructive feedback. Yin Yang is partially supported by NSF 2301040. Chenfanfu Jiang is supported by NSF 2153851, TRI, Sony, and Nvidia. We would also like to acknowledge the kind support from XtalPi and Foundry BioSciences.

References

- Sheldon Andrews, Marek Teichmann, and Paul G. Kry. 2017. Geometric Stiffness for Real-time Constrained Multibody Dynamics. *Computer Graphics Forum* 36, 2 (2017), 235–246. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13122 doi:10.1111/cgf.13122
- Yunfei Bai and C. Karen Liu. 2014. Coupling cloth and rigid bodies for dexterous manipulation. In *Proceedings of the 7th International Conference on Motion in Games (Playa Vista, California) (MIG ’14)*. Association for Computing Machinery, New York, NY, USA, 139–145. doi:10.1145/2668064.2668066
- D. Baraff. 1989. Analytical methods for dynamic simulation of non-penetrating rigid bodies. *SIGGRAPH Comput. Graph.* 23, 3 (July 1989), 223–232. doi:10.1145/74334.74356
- David Baraff. 1997. An introduction to physically based modeling: rigid body simulation I—unconstrained rigid body dynamics. *SIGGRAPH course notes* 82 (1997).
- J. Baumgarte. 1972. Stabilization of constraints and integrals of motion in dynamical systems. *Computer Methods in Applied Mechanics and Engineering* 1, 1 (1972), 1–16. doi:10.1016/0045-7825(72)90018-7
- Jan Bender, Kenny Erleben, and Jeff Trinkle. 2014. Interactive Simulation of Rigid Body Dynamics in Computer Graphics. *Computer Graphics Forum* 33, 1 (2014), 246–270. arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.12272 doi:10.1111/cgf.12272

- Miklós Bergou, Max Wardetzky, Stephen Robinson, Basile Audoly, and Eitan Grinspun. 2008. Discrete elastic rods. In *ACM SIGGRAPH 2008 Papers* (Los Angeles, California) (SIGGRAPH '08). Association for Computing Machinery, New York, NY, USA, Article 63, 12 pages. doi:10.1145/1399504.1360662
- Stephen Boyd and Lieven Vandenbergh. 2004. *Convex optimization*. Cambridge university press.
- Anka He Chen, Ziheng Liu, Yin Yang, and Cem Yuksel. 2024. Vertex Block Descent. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2024)* 43, 4, Article 116 (07 2024), 16 pages. doi:10.1145/3658179
- Yunuo Chen, Minchen Li, Lei Lan, Hao Su, Yin Yang, and Chenfanfu Jiang. 2022. A unified newton barrier method for multibody dynamics. *ACM Trans. Graph.* 41, 4, Article 66 (July 2022), 14 pages. doi:10.1145/3528223.3530076
- Erwin Coumans. 2015. Bullet Physics Simulation. In *ACM SIGGRAPH 2015 Courses*. ACM. doi:10.1145/2776880.2792704
- Crispin Deul, Patrick Charrier, and Jan Bender. 2016. Position-based rigid-body dynamics. *Comput. Animat. Virtual Worlds* 27, 2 (March 2016), 103–112. doi:10.1002/cav.1614
- Andreas Enzenhöfer, Nicolas Lefebvre, and Sheldon Andrews. 2019. Efficient block pivoting for multibody simulations with contact. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games* (Montreal, Quebec, Canada) (I3D '19). Association for Computing Machinery, New York, NY, USA, Article 2, 9 pages. doi:10.1145/3306131.3317019
- Tom Erez, Yuval Tassa, and Emanuel Todorov. 2015. Simulation tools for model-based robotics: Comparison of Bullet, Havok, MuJoCo, ODE and PhysX. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*. 4397–4404. doi:10.1109/ICRA.2015.7139807
- Kenny Erleben, Jon Sporring, Knud Henriksen, and Henrik Dohlmann. 2005. *Physics-based animation*. Vol. 79. Charles River Media Hingham.
- Roy Featherstone. 2008. *Rigid body dynamics algorithms*. Springer.
- Zachary Ferguson, Minchen Li, Teseo Schneider, Francisca Gil-Ureta, Timothy Langlois, Chenfanfu Jiang, Denis Zorin, Danny M. Kaufman, and Daniele Panozzo. 2021. Intersection-free rigid body dynamics. *ACM Trans. Graph.* 40, 4, Article 183 (July 2021), 16 pages. doi:10.1145/3450626.3459802
- Moritz Geilinger, David Hahn, Jonas Zehnder, Moritz Bäcker, Bernhard Thomaszewski, and Stelian Coros. 2020. ADD: analytically differentiable dynamics for multi-body systems with frictional contact. *ACM Trans. Graph.* 39, 6, Article 190 (Nov. 2020), 15 pages. doi:10.1145/3414685.3417766
- Chris Giles, Elie Diaz, and Cem Yuksel. 2025. Augmented Vertex Block Descent. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2025)* 44, 4 (08 2025), 12 pages. doi:10.1145/3731195
- Herbert Goldstein, Charles P Poole, and John Safko. 1950. *Classical mechanics*. Vol. 2. Addison-wesley Reading, MA.
- Gaël Guennebaud, Benoît Jacob, et al. 2010. Eigen: A C++ Template Library for Linear Algebra. <https://eigen.tuxfamily.org>.
- Ernst Hairer, Christian Lubich, and Gerhard Wanner. 2006. *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations* (2 ed.). Springer. doi:10.1007/3-540-30666-8
- Sumit Jain and C. Karen Liu. 2011. Controlling physics-based characters using soft contacts. In *Proceedings of the 2011 SIGGRAPH Asia Conference* (Hong Kong, China) (SA '11). Association for Computing Machinery, New York, NY, USA, Article 163, 10 pages. doi:10.1145/2024156.2024197
- Danny M. Kaufman, Shinjiro Sueda, Doug L. James, and Dinesh K. Pai. 2008. Staggered projections for frictional contact in multibody systems. In *ACM SIGGRAPH Asia 2008 Papers* (Singapore) (SIGGRAPH Asia '08). Association for Computing Machinery, New York, NY, USA, Article 164, 11 pages. doi:10.1145/1457515.1409117
- Theodore Kim and David Eberle. 2020. Dynamic deformables: implementation and production practicalities. In *Acm siggraph 2020 courses*. 1–182.
- Claude Lacoursière. 2007. *Ghosts and Machines: Regularized Variational Methods for Interactive Simulations of Multibodies with Dry Frictional Contacts*. Ph. D. Dissertation. Umeå University. <https://api.semanticscholar.org/CorpusID:115550875>
- Lei Lan, Danny M. Kaufman, Minchen Li, Chenfanfu Jiang, and Yin Yang. 2022. Affine body dynamics: fast, stable and intersection-free simulation of stiff materials. *ACM Trans. Graph.* 41, 4, Article 67 (July 2022), 14 pages. doi:10.1145/3528223.3530064
- Minchen Li, Zachary Ferguson, Teseo Schneider, Timothy Langlois, Denis Zorin, Daniele Panozzo, Chenfanfu Jiang, and Danny M. Kaufman. 2020. Incremental potential contact: intersection-and inversion-free, large-deformation dynamics. *ACM Trans. Graph.* 39, 4, Article 49 (Aug. 2020), 20 pages. doi:10.1145/3386569.3392425
- Minchen Li, Danny M. Kaufman, and Chenfanfu Jiang. 2021. Codimensional incremental potential contact. *ACM Trans. Graph.* 40, 4, Article 170 (July 2021), 24 pages. doi:10.1145/3450626.3459767
- Miles Macklin, Kenny Erleben, Matthias Müller, Nuttapong Chentanez, Stefan Jeschke, and Viktor Makoviychuk. 2019. Non-smooth Newton Methods for Deformable Multi-body Dynamics. *ACM Trans. Graph.* 38, 5, Article 140 (Oct. 2019), 20 pages. doi:10.1145/3338695
- Guirec Maloisel, Ruben Grandia, Christian Schumacher, Espen Knoop, and Moritz Bäcker. 2025. A Versatile Quaternion-Based Constrained Rigid Body Dynamics. *ACM Trans. Graph.* 44, 4, Article 156 (July 2025), 17 pages. doi:10.1145/3730872
- J. E. Marsden and M. West. 2001. Discrete mechanics and variational integrators. *Acta Numerica* 10 (2001), 357–514. doi:10.1017/S096249290100006X
- Brian Mirtich and John F. Canny. 1994. *Impulse-based Dynamic Simulation*. Technical Report UCB/CSD-94-815. <http://www2.eecs.berkeley.edu/Pubs/TechRpts/1994/5494.html>
- J. J. Moreau. 1985. Standard Inelastic Shocks and the Dynamics of Unilateral Constraints. In *Unilateral Problems in Structural Analysis*, Gianpietro Del Piero and Franco Maceri (Eds.). Springer Vienna, Vienna, 173–221.
- Matthias Müller and Markus H Gross. 2004. Interactive Virtual Materials.. In *Graphics interface*, Vol. 2004. 239–246.
- Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. 2007. Position based dynamics. *J. Vis. Comun. Image Represent.* 18, 2 (April 2007), 109–118. doi:10.1016/j.jvcir.2007.01.005
- Matthias Müller, Miles Macklin, Nuttapong Chentanez, Stefan Jeschke, and Tae-Yong Kim. 2020. Detailed rigid body simulation with extended position based dynamics. In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (Virtual Event, Canada) (SCA '20). Eurographics Association, Goslar, DEU, Article 10, 12 pages. doi:10.1111/cgf.14105
- NVIDIA Corporation. 2025. NVIDIA PhysX SDK. <https://github.com/NVIDIA-Omniverse/PhysX>. Accessed: 2026-01-04.
- Albert Peiret, Sheldon Andrews, József Kövecses, Paul G. Kry, and Marek Teichmann. 2019. Schur Complement-based Substructuring of Stiff Multibody Systems with Contact. *ACM Trans. Graph.* 38, 5, Article 150 (Oct. 2019), 17 pages. doi:10.1145/3355621
- William H Press. 2007. *Numerical recipes 3rd edition: The art of scientific computing*. Cambridge university press.
- Xavier Provot et al. 1995. Deformation constraints in a mass-spring model to describe rigid cloth behaviour. In *Graphics interface*. Canadian Information Processing Society, 147–147.
- Stephane Redon, Nico Galoppo, and Ming C. Lin. 2005. Adaptive dynamics of articulated bodies. In *ACM SIGGRAPH 2005 Papers* (Los Angeles, California) (SIGGRAPH '05). Association for Computing Machinery, New York, NY, USA, 936–945. doi:10.1145/1186822.1073294
- Tamar Shinar, Craig Schroeder, and Ronald Fedkiw. 2008. Two-way coupling of rigid and deformable bodies. In *Proceedings of the 2008 ACM SIGGRAPH/Eurographics Symposium on Computer Animation* (Dublin, Ireland) (SCA '08). Eurographics Association, Goslar, DEU, 95–103.
- Breannan Smith, Danny M. Kaufman, Etienne Vouga, Rasmus Tamstorf, and Eitan Grinspun. 2012. Reflections on simultaneous impact. *ACM Trans. Graph.* 31, 4, Article 106 (July 2012), 12 pages. doi:10.1145/2185520.2185602
- David E. Stewart. 2000. Rigid-Body Dynamics with Friction and Impact. *SIAM Rev.* 42, 1 (2000), 3–39. arXiv:https://doi.org/10.1137/S0036144599360110 doi:10.1137/S0036144599360110
- Shinjiro Sueda, Garrett L. Jones, David I. W. Levin, and Dinesh K. Pai. 2011. Large-scale dynamic simulation of highly constrained strands. *ACM Trans. Graph.* 30, 4, Article 39 (July 2011), 10 pages. doi:10.1145/2010324.1964934
- Emanuel Todorov, Tom Erez, and Yuval Tassa. 2012. MuJoCo: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 5026–5033. doi:10.1109/IROS.2012.6386109
- Maxime Tournier, Matthieu Nesme, Benjamin Gilles, and François Faure. 2015. Stable constrained dynamics. *ACM Trans. Graph.* 34, 4, Article 132 (July 2015), 10 pages. doi:10.1145/2766969
- Alexandra C Walls, Young-Jun Park, M Alejandra Tortorici, Abigail Wall, Andrew T McGuire, and David Veisler. 2020. Structure, function, and antigenicity of the SARS-CoV-2 spike glycoprotein. *Cell* 181, 2 (2020), 281–292.
- Endong Wang, Qing Zhang, Bo Shen, Guangyong Zhang, Xiaowei Lu, Qing Wu, and Yajuan Wang. 2014. Intel math kernel library. In *High-Performance Computing on the Intel® Xeon Phi™: How to Fully Exploit MIC Architectures*. Springer, 167–188.
- Ying Wang, Nicholas J. Weidner, Margaret A. Baxter, Yura Hwang, Danny M. Kaufman, and Shinjiro Sueda. 2019. RedMax: efficient & flexible approach for articulated dynamics. *ACM Trans. Graph.* 38, 4, Article 104 (July 2019), 10 pages. doi:10.1145/3306346.3322952
- R. Weinstein, J. Teran, and R. Fedkiw. 2006. Dynamic simulation of articulated rigid bodies with contact and collision. *IEEE Transactions on Visualization and Computer Graphics* 12, 3 (2006), 365–374. doi:10.1109/TVCG.2006.48
- Keenon Werling, Dalton Omens, Jeongseok Lee, Ioannis Exarchos, and C. Karen Liu. 2021. Fast and Feature-Complete Differentiable Physics for Articulated Rigid Bodies with Contact. *CoRR abs/2103.16021* (2021). arXiv:2103.16021 <https://arxiv.org/abs/2103.16021>
- Renhong Yan, Yuan Yuan Zhang, Yaning Li, Lu Xia, Yingying Guo, and Qiang Zhou. 2020. Structural basis for the recognition of SARS-CoV-2 by full-length human ACE2. *Science* 367, 6485 (2020), 1444–1448.