

Foundations of Position-Based Rigid Body Dynamics

Tom Waterson *

April 12, 2026

*Epic Games, UK email:tom.watson@epicgames.com

Contents

1	Introduction	4
1.1	Motivation	4
1.2	Overview	4
	Notation	5
2	Rigid Body Configuration	7
2.1	Angular Velocity	8
2.1.1	The Path Ordered Exponential	8
2.1.2	Axis/Angle Representation and the Exponential Map	10
2.1.3	The Left Jacobian of $SO(3)$	10
2.2	Quaternion Representation	12
2.2.1	Quaternions from the Path Ordered Exponential	13
2.2.2	Relation to Axis/Angle Representation	14
3	Equations of Motion	15
3.1	Holonomic Constraints	18
3.2	Non-Holonomic Constraints	20
4	Solving the Constrained Equations of Motion	22
4.1	Solving the Unconstrained Equations of Motion	22
4.2	Finding a First Order Approximate Solution	23
4.3	Regularization	25
4.4	Interpretation: Connection to the Moore-Penrose Inverse and Least Squares Minimization	25
4.5	An Iterative Solution Using Gauss-Seidel	26
4.6	Comparison to the Jacobi Method	28
4.7	Convergence	28
4.7.1	Rate of convergence	29
4.8	A More Accurate Update Using Non-Linear Gauss-Seidel	32
4.9	Relaxation	32
4.9.1	Under-Relaxation	33
4.9.2	Over-Relaxation	33
4.9.3	Adaptive Relaxation	34
4.10	Warm Starting	34
4.11	Adding Non-Holonomic Constraints	34
4.12	Interpretation: Sequential Impulses	35
4.13	Interpretation: Position Based Dynamics	36
5	Stiff Constraint Examples	38
5.1	Point to Point Constraint	38
5.1.1	Comparison to an Impulse Based Approach	40
5.2	Symmetries and Conservation of Momentum	40
5.3	Inequality Constraints	42
5.4	Contact Constraint	43
5.5	Contact Velocity Constraint	44
5.6	Contact Friction Constraint	45

5.7	Angular Axis Constraint	46
6	Compliant Constraints	48
6.1	Generalized Implicit Non-Linear Forces	48
6.2	Connection to Vertex Block Descent	49
6.3	Connection to Lagrange Multiplier Method and PBD	49
6.4	Compliant Constraint Forces and XPBD	50
6.4.1	XPBD as a Regularization on PBD	52
6.5	Compliant Constraint Examples	53
6.5.1	Spring Damper	53
6.5.2	Volumetric Constraint	54
7	Augmented Lagrangian Method	56
7.1	Augmented Vertex Block Descent	57
7.1.1	AVBD, PBD and Geometric Stiffness	57
7.2	Convergence Comparison vs. PBD For Large Mass Ratios	58
7.3	Point to Point Constraint	59
7.3.1	Exact Solution	60
7.3.2	Iterative Solution	61
8	Convergence of Constraint Chains	63
8.1	Setup	63
8.2	Effect of Iteration Count	64
8.2.1	Relaxation	65
8.2.2	Jacobi Iterations	66
8.2.3	AVBD Iterations	67
8.3	Effect of Mass Ratios	68
8.3.1	Relaxation	69
8.3.2	AVBD Iterations	70
8.4	Results	70
9	Summary	71

1 Introduction

1.1 Motivation

Position-Based Dynamics (PBD) and its extensions are widely used in real-time physics simulation for games and interactive applications due to their robustness and computational efficiency. However, the relationship between these methods and the classical formulations of rigid body dynamics is often not presented in a unified or principled way.

This document develops a derivation of position-based rigid body dynamics from first principles using constrained Lagrangian mechanics. Starting from the configuration space of rigid bodies on $SE(3)$, we derive the equations of motion and linearize the system using a first order approximation and solve via Gauss-Seidel. We show how common position based methods, such as PBD and XPBD are manifestations of this approach. The goal is to provide a clear mathematical foundation for techniques that are widely used in practical physics engines while clarifying their connection to classical rigid body dynamics and constrained optimization. With this motivation, these notes differ from the majority of the literature on this subject, where the goal is to provide clear practical implementation details, and to show the efficacy of these methods. For some excellent notes on the practical approach to position based dynamics methods, see, for example [2, 3, 17, 15, 8, 10, 16].

These notes are an accumulation of notes that I have been making and refining over the years of implementing rigid body simulation in the games industry. The motivation was to provide a solid theoretical grounding for the techniques that are commonly used in industry. While the algorithms of PBD and its variants are well documented, their connection to classical mechanics is rarely presented in a unified or mathematically consistent way, particularly for the angular degrees of freedom, where the relationship between configuration variables, constraint Jacobians, and incremental updates is frequently left implicit. The goal here is to make that connection explicit and to illuminate the connections and shared foundations of position based methods.

1.2 Overview

The main substance of these notes is a derivation of position based dynamics for rigid bodies from first principles using constrained Lagrangian mechanics. In the first section, there is some background on rigid bodies and their configuration. In particular, I try to give a careful treatment of the representation of angular degrees of freedom.

The next section derives the equations of motion for a constrained set of rigid bodies. Here, we use some results from geometric mechanics that exploit the rotation group symmetry. Next is where I start to make the connection to real time simulation and I derive a first order approximation to the constrained equations of motion and show how this can be iteratively solved using Gauss-Seidel and discuss convergence properties of this algorithm.

From there I show how this approach connects to the algorithms of Position Based Dynamics. I cover PBD, XPBD, Vertex Block Descent (VBD) and Augmented Vertex Block Descent (AVBD), and also show the connection to a Sequential Impulse approach, to the Moore-Penrose pseudoinverse method, and to Tikhonov regularization.

In the final section, I analyze the convergence of a chain of connected rigid bodies with some of the different algorithms to show how they differ.


Notation

Symbol	Description
<i>Configuration and kinematics</i>	
$\mathbf{x} \in \mathbb{R}^3$	Position of the centre of mass
$\boldsymbol{\theta} \in \mathbb{R}^3$	Axis-angle rotation vector, $\boldsymbol{\theta} = \theta \hat{\boldsymbol{\theta}}$
$R(\boldsymbol{\theta}) \in SO(3)$	Rotation matrix, $R = e^{[\boldsymbol{\theta}]_{\times}}$
$\mathbf{X} = (\mathbf{x}, \boldsymbol{\theta})$	Generalised configuration of a single body
$\mathcal{V} = (\dot{\mathbf{x}}, \boldsymbol{\omega})$	Generalised velocity of a single body
$\boldsymbol{\omega} \in \mathbb{R}^3$	World-space angular velocity
$\boldsymbol{\omega}_b \in \mathbb{R}^3$	Body-space angular velocity, $\boldsymbol{\omega}_b = R^T \boldsymbol{\omega}$
q	Unit quaternion representing orientation
$\mathbf{p}(\mathbf{X})$	World-space position of a point on the body, $\mathbf{p} = \mathbf{x} + R(\boldsymbol{\theta})\mathbf{r}_0$
$\mathbf{r}_0 \in \mathbb{R}^3$	Fixed body-space offset of a point from the centre of mass
$\mathbf{r}(\theta)$	World-space offset, $\mathbf{r} = R(\boldsymbol{\theta})\mathbf{r}_0$
<i>Lie group structure</i>	
$[\cdot]_{\times}$	Map from \mathbb{R}^3 to 3×3 skew-symmetric matrices (hat map)
$SO(3)$	Special orthogonal group of 3×3 rotation matrices
$J_L(\boldsymbol{\theta})$	Left Jacobian of $SO(3)$, maps $\dot{\boldsymbol{\theta}}$ to $\boldsymbol{\omega}$: $\boldsymbol{\omega} = J_L \dot{\boldsymbol{\theta}}$
<i>Mass and inertia</i>	
m	Scalar mass of a body
$I(\theta)$	World-space inertia tensor, $I = R(\boldsymbol{\theta})I_0R(\boldsymbol{\theta})^T$
I_0	Body-space (diagonal) inertia tensor
I_b	Body-space inertia tensor
$M(\mathbf{X})$	Generalised mass matrix for a system of N bodies (block diagonal)
μ	Scalar effective mass for a single constraint
\mathcal{M}^{-1}	Inverse of the 6×6 effective mass matrix for a two-body system
<i>Constraints and forces</i>	
$\mathbf{c}(\mathbf{X})$	Vector of holonomic constraint functions
$\boldsymbol{\lambda}$	Vector of Lagrange multipliers (constraint forces)
$\mathcal{D}\mathbf{c}$	Left-trivialized constraint Jacobian, satisfies $\dot{\mathbf{c}} = \mathcal{D}\mathbf{c}\mathcal{V}$
$\dot{\mathcal{D}}\mathbf{c}$	Velocity-space constraint Jacobian $\partial\mathbf{c}/\partial\mathcal{V}$ for non-holonomic constraints
$\mathcal{D}^2\mathbf{c}$	Constraint Hessian using the left-trivialized derivative
$\Delta\mathbf{X}$	Incremental displacement applied to resolve constraints
$\mathbf{F} = (\mathbf{f}, \boldsymbol{\tau})$	Generalised force vector (force and torque)
$\boldsymbol{\tau}_w$	World-space torque

Symbol	Description
<i>Iterative solver</i>	
k	Iteration index
Δt	Timestep
ξ	SOR relaxation parameter ($\xi < 1$ under-relaxation, $\xi > 1$ over-relaxation)
G	Gauss-Seidel iteration matrix, $G = -(D + L)^{-1}L^T$
G_J	Jacobi iteration matrix, $G_J = -D^{-1}(L + L^T)$
G_ξ	SOR iteration matrix
$\rho(\cdot)$	Spectral radius of a matrix
D, L, L^T	Diagonal, strictly lower triangular, and strictly upper triangular parts of A
<i>Compliant constraints and regularization</i>	
A	Diagonal matrix of stiffness parameters in compliant constraints
B	Diagonal matrix of damping parameters
$\tilde{\alpha}_i$	XPBD compliance parameter, $\tilde{\alpha}_i = 1/(A_{ii}\Delta t^2)$
$\tilde{\beta}_i$	XPBD damping parameter, $\tilde{\beta}_i = B_{ii}\Delta t$
ρ	Diagonal matrix of penalty coefficients in the augmented Lagrangian method
ε	Tikhonov regularization parameter
<i>Convergence analysis</i>	
μ_k	Eigenvalues of the Jacobi iteration matrix, $\mu_k = \cos(k\pi/N)$
r_i	Diagonal dominance ratio for row i , $r_i = \sum_{j \neq i} A_{ij} / A_{ii} $
γ	Mass ratio m/M in the chain convergence analysis
α	Fraction of error to be removed (convergence target)

2 Rigid Body Configuration

Rigid bodies are an idealised type of object that does not deform. I.e., for any two points $\mathbf{p}_i, \mathbf{p}_j$ in the rigid body the distance $|\mathbf{p}_i - \mathbf{p}_j|$ remains constant.



They have no internal degrees of freedom, and only have six total degrees of freedom:

- Three to define the position \mathbf{x}
- Three to define the orientation $\boldsymbol{\theta}$

A rigid body rotates around its centre of mass and moves through space on the trajectory of its centre of mass, so these six coordinates are the coordinates of the rigid body centre of mass \mathbf{x} . The location of any point on the rigid body p at any time t can be deduced from the position of the centre of mass $\mathbf{x}(t)$, the orientation of the body $R(t)$ and the (fixed) offset of the point on the rigid body \mathbf{r}_0 :

$$\mathbf{p}(t) = \mathbf{x}(t) + R(t)\mathbf{r}_0. \quad (2.1)$$

The condition that distances between points on the rigid body remain fixed places conditions on the matrix R

$$\|\mathbf{p}_i - \mathbf{p}_j\|^2 = \|R(\mathbf{r}_i - \mathbf{r}_j)\|^2 = (R(\mathbf{r}_i - \mathbf{r}_j)) \cdot (R(\mathbf{r}_i - \mathbf{r}_j)) = (\mathbf{r}_i - \mathbf{r}_j)^T R^T R (\mathbf{r}_i - \mathbf{r}_j). \quad (2.2)$$

Therefore, we require $R^T R = \mathbb{1}$ so that distances remain the same, i.e. R is an orthogonal matrix, with $R^{-1} = R^T$.

Orthogonal matrices form a group: if two matrices R_i and R_j are orthogonal then their product will also be orthogonal

$$(R_i R_j)^T (R_i R_j) = R_j^T R_i^T R_i R_j = R_j^T \mathbb{1} R_j = \mathbb{1}. \quad (2.3)$$

This group is called the Orthogonal Group in three dimensions $O(3)$.

Since $R^T R = \mathbb{1}$, the determinant

$$\det(R^T R) = \det(R)^2 = 1, \quad (2.4)$$

so the determinant is either $+1$ or -1 . Matrices with $\det(R) = +1$ form a subgroup, since the product of two of matrices with determinant $+1$ is also a matrix with determinant $+1$. This is called the Special Orthogonal Group $SO(3)$ and represents rotations. Matrices with determinant -1 do not form a subgroup, as the product of two such matrices is a matrix with determinant $+1$. These represent reflections, and the product of two reflections is a rotation.

To justify that these two classes represent rotations and reflections respectively, we can look at two examples. First, a rotation about the z axis

$$R = \begin{pmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad \det(R) = \cos^2 \theta + \sin^2 \theta = 1. \quad (2.5)$$


And a reflection in the xy plane

$$R = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & -1 \end{pmatrix}, \quad \det(R) = -1. \quad (2.6)$$

2.1 Angular Velocity

The time derivative of the offset vector $\mathbf{r}(t) = R(t)\mathbf{r}_0$ is

$$\begin{aligned}\dot{\mathbf{r}}(t) &= \dot{R}(t)\mathbf{r}_0, \\ &= \dot{R}(t)R^T(t)\mathbf{r}.\end{aligned}\quad (2.7)$$



We know that the matrix $\dot{R}R^T$ must be anti-symmetric, since $RR^T = \mathbb{1}$ so

$$\frac{d}{dt}(RR^T) = \dot{R}R^T + R\dot{R}^T = \dot{R}R^T + (\dot{R}R^T)^T = 0. \quad (2.8)$$

Since $(\dot{R}R^T)$ is anti-symmetric, we can use it to define a vector $\boldsymbol{\omega}$ using the anti-symmetric matrix $[\boldsymbol{\omega}]_{\times}$

$$\dot{R}R^T = [\boldsymbol{\omega}]_{\times} = \begin{pmatrix} 0 & -\omega_z & \omega_y \\ \omega_z & 0 & -\omega_x \\ -\omega_y & \omega_x & 0 \end{pmatrix}. \quad (2.9)$$

Comparing $\dot{\mathbf{r}} = \dot{R}(t)R(t)\mathbf{r}$ to the standard result from mechanics

$$\dot{\mathbf{r}} = \boldsymbol{\omega} \times \mathbf{r}. \quad (2.10)$$

We identify $\boldsymbol{\omega}$ as the world space angular velocity.

Body space angular velocity $\boldsymbol{\omega}_b$ can be derived from the world space angular velocity using

$$R\boldsymbol{\omega}_b = \boldsymbol{\omega}, \quad (2.11)$$

so

$$\begin{aligned}[R\boldsymbol{\omega}_b]_{\times} &= [\boldsymbol{\omega}]_{\times}, \\ R[\boldsymbol{\omega}_b]_{\times}R^T &= \dot{R}R^T, \\ [\boldsymbol{\omega}_b]_{\times} &= R^T\dot{R},\end{aligned}\quad (2.12)$$

where we have used the fact that matrices A transform under a rotation R as $A' = RAR^T$.

2.1.1 The Path Ordered Exponential

Consider a body rotating with a constant angular velocity $\boldsymbol{\omega}$. For an infinitesimal time Δt a vector undergoes an infinitesimal rotation

$$\mathbf{r}' = \mathbf{r} + \boldsymbol{\omega} \times \mathbf{r}\Delta t = (\mathbb{1} + [\boldsymbol{\omega}]_{\times}\Delta t)\mathbf{r}. \quad (2.13)$$

We can build up a finite rotation as a product of infinitesimal rotations

$$\begin{aligned}R\mathbf{r} &= (\mathbb{1} + [\boldsymbol{\omega}\Delta t]_{\times})(\mathbb{1} + [\boldsymbol{\omega}\Delta t]_{\times}) \cdots (\mathbb{1} + [\boldsymbol{\omega}\Delta t]_{\times})\mathbf{r}, \\ &= \lim_{n \rightarrow \infty} \left(\mathbb{1} + \frac{1}{n}[\boldsymbol{\omega}\Delta T]_{\times} \right)^n,\end{aligned}\quad (2.14)$$

where $\Delta t = \Delta T/n$. This limit gives an exponential

$$R(\Delta T) = e^{[\boldsymbol{\omega}\Delta T]_{\times}} \quad (2.15)$$

This can also be seen in the case of constant angular velocity $\boldsymbol{\omega}(t) = \boldsymbol{\omega}_0$ by integrating the equation (2.9).

$$\ln(R) = \int \dot{R}R^T dt = \int [\boldsymbol{\omega}_0]_{\times} dt \quad (2.16)$$

It is tempting to want to generalize this for a general non-constant $\boldsymbol{\omega}(t)$ to

$$R(t) = e^{\int_0^t [\boldsymbol{\omega}(t')]_{\times} dt'} \quad (2.17)$$

but this is not exactly right, as different values of $[\boldsymbol{\omega}]_{\times}$ at different times do not commute with each other (since rotations do not commute). The answer is to use the *path ordered exponential*

$$\begin{aligned} R(t) &= P \exp \left(\int_0^t [\boldsymbol{\omega}(t')]_{\times} dt' \right), \\ &= 1 + \int_0^t [\boldsymbol{\omega}(t')]_{\times} dt' + \int_0^{t''} \int_{t'}^t [\boldsymbol{\omega}(t'')]_{\times} [\boldsymbol{\omega}(t')]_{\times} dt' dt'' + \dots \end{aligned} \quad (2.18)$$

for $0 < t' < t'' < t$. I.e. the omegas are ordered so that later times appear on the left.

This simplifies if we just want to take an incremental update going from some initial state $R(t)$ at time t to a new state at time $t + \Delta t$ for some small Δt . We can use

$$\begin{aligned} R(t + \Delta t) &= \exp \left(\int_t^{t+\Delta t} [\boldsymbol{\omega}(t')]_{\times} dt' \right) R(t), \\ &\approx \exp ([\boldsymbol{\omega}(t)]_{\times} \Delta t) R(t), \\ &= R(t) + [\boldsymbol{\omega}(t)\Delta t]_{\times} R(t) + \mathcal{O}(\Delta t^2), \end{aligned} \quad (2.19)$$


$$R(t + \Delta t) \approx (\mathbb{1} + [\boldsymbol{\omega}\Delta t]_{\times}) R(t) \quad (2.20)$$

where we assume that for a small enough time Δt the angular velocity can be considered to be constant. See David Tong's notes [23] for more detail on angular velocity, rotations and path ordered exponentials.

Note that, in practice, when updating a rotation matrix, it is important to retain its orthogonality, so in a rigid body simulation if we were updating rotation matrices in this way, we would have to orthonormalize the matrix after the update, which can be expensive, hence the general use of quaternions, which will be covered in section 2.2.

2.1.2 Axis/Angle Representation and the Exponential Map

The rotation matrix is a 3×3 matrix, so has 9 elements. But there are only three independent angular degrees of freedom. Later, when we come to look at rigid body constraints, we will need to be able to take the derivative of a function with respect to the angular degrees of freedom, so it is useful to find a representation that is a function of three independent variables. One such representation is the axis/angle representation.



Any orientation of a rigid body can be seen as a single rotation about a fixed axis \hat{n} by an angle θ . This statement is known as *Euler's rotation theorem*. Therefore, we can define a vector $\boldsymbol{\theta} = \theta \hat{\boldsymbol{\theta}}$, where the magnitude θ is the angle of rotation, and the normalized direction $\hat{\boldsymbol{\theta}} = \hat{n}$ is the axis of rotation, to encode a rotation.

This means that, using (2.15), we can write any rotation matrix as

$$R(\boldsymbol{\theta}) = e^{[\boldsymbol{\theta}]_{\times}}. \quad (2.21)$$

Technically, this exponential map is a map between elements R of the Lie group $SO(3)$ consisting of 3×3 orthogonal matrices with unit determinant and its corresponding Lie algebra elements $[\boldsymbol{\theta}]_{\times}$ in $\mathfrak{so}(3)$ consisting of 3×3 skew-symmetric matrices. See, for example [26] for more information on the theory of Lie groups.

This exponential can be expanded to give the very useful *Rodrigues Rotation Formula*

$$\begin{aligned} R(\boldsymbol{\theta}) &= e^{[\boldsymbol{\theta}]_{\times}}, \\ &= \mathbb{1} + [\boldsymbol{\theta}]_{\times} + \frac{1}{2!}[\boldsymbol{\theta}]_{\times}^2 + \frac{1}{3!}[\boldsymbol{\theta}]_{\times}^3 + \dots, \\ &= \mathbb{1} + [\boldsymbol{\theta}]_{\times} \left(1 - \frac{1}{3!}|\boldsymbol{\theta}|^2 + \dots\right) + [\boldsymbol{\theta}]_{\times}^2 \left(\frac{1}{2!} - \frac{1}{4!}|\boldsymbol{\theta}|^2 + \dots\right), \\ &= \mathbb{1} + \frac{\sin \theta}{\theta}[\boldsymbol{\theta}]_{\times} + \frac{1 - \cos \theta}{\theta^2}[\boldsymbol{\theta}]_{\times}^2 \end{aligned} \quad (2.22)$$

where we have used $[\boldsymbol{\theta}]_{\times}^3 = -|\boldsymbol{\theta}|^2[\boldsymbol{\theta}]_{\times}$.

2.1.3 The Left Jacobian of $SO(3)$

It's important to remember that $R(\boldsymbol{\theta})$ represents a single rotation, and we can't compose rotations by adding axis/angles as

$$e^{[\boldsymbol{\theta}_1 + \boldsymbol{\theta}_2]_{\times}} \neq e^{[\boldsymbol{\theta}_1]_{\times}} e^{[\boldsymbol{\theta}_2]_{\times}} \quad (2.23)$$

since $[\boldsymbol{\theta}_1]_{\times}$ and $[\boldsymbol{\theta}_2]_{\times}$ do not necessarily commute. They only commute if the rotations are about the same axis. What we can do is to ask what is $R(\boldsymbol{\theta} + \Delta\boldsymbol{\theta})$. The *Baker-Campbell-Hausdorff Formulas* can be used to relate small perturbations in $[\boldsymbol{\theta}]_{\times}$ to small perturbations in $R(\boldsymbol{\theta})$.

The BCH formula finds Z for $e^Z = e^X e^Y$

$$Z = X + Y + \frac{1}{2}[X, Y] + \frac{1}{12}[X, [X, Y]] + \dots \quad (2.24)$$

where $[A, B] = AB - BA$ is the commutator of A and B .

We want to use this to find a decomposition of the rotation for small angles, i.e. find a $J\Delta\theta$ such that $e^{[\theta+\Delta\theta]_\times} = e^{[J\Delta\theta]_\times} e^{[\theta]_\times}$. So $Z = [\theta + \Delta\theta]_\times$, $X = [J\Delta\theta]_\times$ and $Y = [\theta]_\times$

$$[\theta + \Delta\theta]_\times = [J\Delta\theta]_\times + \frac{1}{2}[[J\Delta\theta]_\times, \theta] + \frac{1}{12}[[\theta]_\times, [[\theta]_\times, [\Delta\theta]_\times]] + \dots, \quad (2.25)$$

where the dots indicate terms of higher order in $\Delta\theta$.

This series can be solved [5, 21] to give

$$e^{[\theta+\Delta\theta]_\times} \approx e^{[J_L(\theta)\Delta\theta]_\times} e^{[\theta]_\times}, \quad (2.26)$$

where $J_L(\theta)$ is given by

$$J_L(\theta) = \mathbb{1} + \frac{1 - \cos\theta}{\theta^2}[\theta]_\times + \frac{\theta - \sin\theta}{\theta^3}[\theta]_\times^2. \quad (2.27)$$

$J_L(\theta)$ is known as the *Left Jacobian of $SO(3)$* . Comparing to (2.20) we see that for small perturbations

$$\omega(t)\Delta t = \Delta\omega \approx J_L(\theta)\Delta\theta, \quad (2.28)$$

so J_L maps small changes in axis angle to small changes in angular velocity.

To see this more precisely, we can take another route. First, take the time derivative of the rotation matrix.

$$\begin{aligned} \dot{R}(t) &= \frac{d}{dt} e^{[\theta(t)]_\times}, \\ &= \frac{d}{dt} \lim_{N \rightarrow \infty} \left(\mathbb{1} + \frac{1}{N}[\theta(t)]_\times \right)^N, \\ &= \lim_{N \rightarrow \infty} \sum_{k=1}^N \left(\mathbb{1} + \frac{1}{N}[\theta(t)]_\times \right)^{N-k} \frac{1}{N}[\dot{\theta}(t)]_\times \left(\mathbb{1} + \frac{1}{N}[\theta(t)]_\times \right)^{k-1}. \end{aligned} \quad (2.29)$$

Note that $[\theta(t)]_\times$ and $[\dot{\theta}(t)]_\times$ do not necessarily commute. Now re-scale the interval into $[0, \dots, 1]$ using $\Delta s = \Delta k/N = 1/N$, since k are integers so Δk is 1. Taking the limit as $N \rightarrow \infty$ gives

$$\dot{R}(t) = \int_0^1 e^{(1-s)[\theta(t)]_\times} [\dot{\theta}(t)]_\times e^{s[\theta(t)]_\times} ds. \quad (2.30)$$

Now let $s' = 1 - s$

$$\begin{aligned} \dot{R}(t) &= - \int_1^0 e^{s'[\theta(t)]_\times} [\dot{\theta}(t)]_\times e^{(1-s')[\theta(t)]_\times} ds', \\ &= \int_0^1 e^{s[\theta(t)]_\times} [\dot{\theta}(t)]_\times e^{(1-s)[\theta(t)]_\times} ds, \\ &= \left(\int_0^1 e^{s[\theta(t)]_\times} [\dot{\theta}(t)]_\times e^{-s[\theta(t)]_\times} ds \right) R(t). \end{aligned} \quad (2.31)$$

Looking at the expression in the integral this is $R(s\boldsymbol{\theta})[\dot{\boldsymbol{\theta}}]_{\times}R(s\boldsymbol{\theta})^T$. We know that this is how a matrix R transforms under a rotation. A vector would transform as $R(s\boldsymbol{\theta})\dot{\boldsymbol{\theta}}$, so

$$\dot{R}(t)R^T(t) = [\omega(t)]_{\times} = \left[\left(\int_0^1 e^{s[\boldsymbol{\theta}(t)]_{\times}} ds \right) \dot{\boldsymbol{\theta}}(t) \right]_{\times}. \quad (2.32)$$

(Technically, the adjoint action of the group $SO(3)$ acting on its algebra $\mathfrak{so}(3)$ is $Ad_R[X]_{\times} = R[X]_{\times}R^T$. Elements of the Lie algebra can be mapped to vectors $\boldsymbol{x} \in \mathbb{R}^3$ with $Ad_R\boldsymbol{x} = R\boldsymbol{x}$).

If we define

$$J_L(\boldsymbol{\theta}) = \int_0^1 e^{s[\boldsymbol{\theta}]_{\times}} ds \quad (2.33)$$

then

$$\boldsymbol{\omega} = J_L(\boldsymbol{\theta})\dot{\boldsymbol{\theta}}. \quad (2.34)$$

It can be verified that this is the same J_L as we got from the BCH expansion by using the Rodrigues rotation formula and performing the integral

$$\begin{aligned} J_L(\boldsymbol{\theta}) &= \int_0^1 e^{s[\boldsymbol{\theta}]_{\times}} ds, \\ &= \int_0^1 \left(\mathbb{1} + \frac{\sin s\theta}{\theta} [\boldsymbol{\theta}]_{\times} + \frac{1 - \cos s\theta}{\theta^2} [\boldsymbol{\theta}]_{\times}^2 \right) ds, \\ &= \mathbb{1} + \frac{1 - \cos \theta}{\theta^2} [\boldsymbol{\theta}]_{\times} + \frac{\theta - \sin \theta}{\theta^3} [\boldsymbol{\theta}]_{\times}^2. \end{aligned} \quad (2.35)$$

In summary:

J_L is known as the left Jacobian of $SO(3)$ and maps the derivative of the axis angle vector to the world space angular velocity.

$$\boldsymbol{\omega} = J_L(\boldsymbol{\theta})\dot{\boldsymbol{\theta}}$$

It can be used to find how infinitesimal changes in the axis/angle $\boldsymbol{\theta}$ update the rotation $R = e^{[\boldsymbol{\theta}]_{\times}}$

$$e^{[\boldsymbol{\theta}+\Delta\boldsymbol{\theta}]_{\times}} \approx e^{[J_L(\boldsymbol{\theta})\Delta\boldsymbol{\theta}]_{\times}} e^{[\boldsymbol{\theta}]_{\times}} = e^{[\Delta\boldsymbol{\omega}]_{\times}} R(\boldsymbol{\theta}). \quad (2.36)$$

This is very useful for rigid body simulations where we are taking small perturbations in the configuration space, and taking derivatives of constraints with respect to the constraint rotation.

2.2 Quaternion Representation

This section is mostly tangential to the rest of the notes, which, on the whole, use the axis/angle and exponential map representation of rotations. The exception is in integrating rigid body rotations, where maintaining orthogonality is important, so we use the quaternion representation as it is more efficient to normalize.

In this section, I will be using and assuming some familiarity with geometric algebra, the natural and most intuitive setting for quaternions. A good reference for this in the context of classical physics is [13]. For readers familiar with Hamilton's formulation of quaternions, the end result will be familiar and you can safely replace I with the unit complex number i .

When using quaternions, we define a point on a rigid body by its centre of mass $\mathbf{x}(t)$ and a quaternion $q(t)$ representing the orientation, so that a point on the body is

$$\mathbf{p}(t) = \mathbf{x}(t) + q(t)\mathbf{r}_0\bar{q}(t), \quad (2.37)$$

where \bar{q} is the quaternion conjugate. The condition that the distance between two points on the body remains fixed means that, using the geometric product

$$\|\mathbf{p}_i - \mathbf{p}_j\|^2 = q(\mathbf{r}_i - \mathbf{r}_j)\bar{q}q(\mathbf{r}_i - \mathbf{r}_j)\bar{q} \quad (2.38)$$

So we require that $q\bar{q} = \bar{q}q = 1$ for this to be invariant.

The time derivative of a vector $\mathbf{r}(t) = q(t)\mathbf{r}_0\bar{q}(t)$ is

$$\begin{aligned} \dot{\mathbf{r}}(t) &= \dot{q}(t)\mathbf{r}_0\bar{q}(t) + q(t)\mathbf{r}_0\dot{\bar{q}}(t), \\ &= \dot{q}(t)\bar{q}(t)\mathbf{r}(t) + \mathbf{r}(t)q(t)\dot{\bar{q}}(t), \\ &= (\dot{q}(t)\bar{q}(t))\mathbf{r}(t) - \mathbf{r}(t)(\dot{q}(t)\bar{q}(t)). \end{aligned} \quad (2.39)$$

Now define the bivector $\mathbf{\Omega} = 2\dot{q}(t)\bar{q}$. This is then

$$\begin{aligned} \dot{\mathbf{r}}(t) &= \frac{1}{2}(\mathbf{\Omega}\mathbf{r} - \mathbf{r}\mathbf{\Omega}), \\ &= \mathbf{\Omega} \cdot \mathbf{r}. \end{aligned} \quad (2.40)$$

$\mathbf{\Omega}$ is the dual of the angular velocity vector $\boldsymbol{\omega}$

$$\mathbf{\Omega} = I\boldsymbol{\omega}, \quad (2.41)$$

where I is the unit pseudoscalar, with the property $I^2 = -1$.

2.2.1 Quaternions from the Path Ordered Exponential

Just as we found a differential equation for the rotation matrix $\dot{R} = [\boldsymbol{\omega}]_{\times}R$, we can find the equivalent differential equation for the quaternion q . Starting from our definition of the angular velocity bivector $\mathbf{\Omega} = 2\dot{q}\bar{q}$ we have

$$\dot{q}(t) = \frac{1}{2}\mathbf{\Omega}(t)\bar{q}(t). \quad (2.42)$$

This equation is identical in form to the one we solved for rotation matrices. If the body is rotating with a constant angular velocity bivector $\mathbf{\Omega}$, we can integrate this directly to find the quaternion at time t :

$$q(t) = \exp\left(\frac{1}{2}\mathbf{\Omega}t\right). \quad (2.43)$$

However, for a general non-constant angular velocity $\mathbf{\Omega}(t)$, the bivectors at different times do not commute. Just as in the rotation matrix case, we must use the *path ordered exponential* to build up the finite rotation from infinitesimal steps:

$$q(t) = P \exp\left(\frac{1}{2}\int_0^t \mathbf{\Omega}(t')dt'\right) \quad (2.44)$$

In practice, for physics simulations and games, we usually want to perform an incremental update over a small timestep Δt . We can assume the angular velocity is roughly constant over this small

window, giving the discrete update step:

$$\begin{aligned}
q(t + \Delta t) &= \exp\left(\frac{1}{2} \int_t^{t+\Delta t} \boldsymbol{\Omega}(t') dt'\right) q(t), \\
&\approx \exp\left(\frac{1}{2} \boldsymbol{\Omega}(t) \Delta t\right) q(t), \\
&\approx \exp\left(\frac{1}{2} I \boldsymbol{\omega}(t) \Delta t\right) q(t).
\end{aligned} \tag{2.45}$$

This quaternion update is generally favoured in numerical integration as normalizing the quaternion to avoid numerical drift is computationally easier than orthonormalizing a 3×3 matrix.

2.2.2 Relation to Axis/Angle Representation

To connect this back to our axis/angle representation $\boldsymbol{\theta} = \theta \hat{\mathbf{n}}$, we need to evaluate the exponential map of the bivector. Let \hat{B} be the unit bivector representing the plane of rotation, which is the dual to the rotation axis $\hat{\mathbf{n}}$ such that $\hat{B} = I \hat{\mathbf{n}}$. If a body rotates by a total angle θ in this plane, the integral of the angular velocity over that time is $\int \boldsymbol{\Omega} dt = \theta \hat{B}$. Substituting this into our exponential solution, the quaternion representing a rotation by $\boldsymbol{\theta}$ is:

$$q(\boldsymbol{\theta}) = \exp\left(\frac{1}{2} \theta \hat{B}\right). \tag{2.46}$$

Because the square of a unit bivector in 3D Euclidean space is negative ($\hat{B}^2 = -1$), the Taylor series expansion of this exponential gives Euler's formula for complex numbers:

$$\begin{aligned}
\exp\left(\frac{1}{2} \theta \hat{B}\right) &= 1 + \frac{1}{2} \theta \hat{B} + \frac{1}{2!} \left(\frac{1}{2} \theta \hat{B}\right)^2 + \frac{1}{3!} \left(\frac{1}{2} \theta \hat{B}\right)^3 + \dots \\
&= \left(1 - \frac{1}{2!} \left(\frac{\theta}{2}\right)^2 + \dots\right) + \hat{B} \left(\frac{\theta}{2} - \frac{1}{3!} \left(\frac{\theta}{2}\right)^3 + \dots\right) \\
&= \cos\left(\frac{\theta}{2}\right) + \hat{B} \sin\left(\frac{\theta}{2}\right).
\end{aligned} \tag{2.47}$$

This is the link between the axis/angle representation and the quaternion. A unit quaternion q is an even-grade multivector composed of a scalar part and a bivector part:

$$q = \cos\left(\frac{\theta}{2}\right) + I \hat{\mathbf{n}} \sin\left(\frac{\theta}{2}\right). \tag{2.48}$$

In standard (Hamilton) quaternion notation, the bivector $I \hat{\mathbf{n}}$ is mapped directly to the imaginary vector components i, j, k . Notice the appearance of the half-angle $\theta/2$. This occurs because the vector transformation law $\mathbf{r} = q \mathbf{r}_0 \bar{q}$ applies the quaternion twice (once on the left, once on the right). To rotate a vector by a full angle θ , each side of the sandwich product must contribute half of the rotation.

3 Equations of Motion

The goal is to find the equations of motion for a rigid body. To do this we follow a Lagrangian derivation. The Lagrangian for a system is the difference between the kinetic and potential energy. For a single rigid body this is

$$L = \frac{1}{2}m|\dot{\mathbf{x}}|^2 + \frac{1}{2}\boldsymbol{\omega} \cdot I(\boldsymbol{\theta})\boldsymbol{\omega} - V(\mathbf{x}, \boldsymbol{\theta}), \quad (3.1)$$

for mass m , inertia $I(\boldsymbol{\theta})$ and potential energy $V(\mathbf{x}, \boldsymbol{\theta})$. For standard variational mechanics, we want to make this a function of \mathbf{X} and it's time derivative $\dot{\mathbf{X}}$. In this notation

The Lagrangian for a single rigid body is

$$L(\mathbf{X}, \dot{\mathbf{X}}) = \frac{1}{2}\dot{\mathbf{X}} \cdot M(\mathbf{X})\dot{\mathbf{X}} - V(\mathbf{X}) \quad (3.2)$$

where

$$M(\mathbf{X}) = \begin{pmatrix} m\mathbb{1} & 0 \\ 0 & J_L(\boldsymbol{\theta})^T I(\boldsymbol{\theta}) J_L(\boldsymbol{\theta}) \end{pmatrix}, \quad \mathbf{X} = \begin{pmatrix} \mathbf{x} \\ \boldsymbol{\theta} \end{pmatrix}, \quad \dot{\mathbf{X}} = \begin{pmatrix} \mathbf{v} \\ J_L(\boldsymbol{\theta})^{-1}\boldsymbol{\omega} \end{pmatrix}. \quad (3.3)$$

Notice that we have factored in the left Jacobian $J_L(\boldsymbol{\theta})$ into the mass matrix $M(\mathbf{X})$, so that we get an equation in terms of $\dot{\mathbf{X}}$, and we have used it to map $\dot{\boldsymbol{\theta}}$ to the angular velocity $\boldsymbol{\omega}$.

In standard Lagrangian mechanics we vary the action according to

$$\begin{aligned} \delta S &= \int dt \delta L(\dot{\mathbf{X}}, \mathbf{X}), \\ &= \int dt \left(\frac{\partial L}{\partial \dot{\mathbf{X}}} \delta \dot{\mathbf{X}} + \frac{\partial L}{\partial \mathbf{X}} \delta \mathbf{X} \right), \\ &= \int dt \left(-\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{\mathbf{X}}} \right) + \frac{\partial L}{\partial \mathbf{X}} \right) \delta \mathbf{X} + \left[\frac{\partial L}{\partial \dot{\mathbf{X}}} \delta \mathbf{X} \right]_{\text{boundary}}. \end{aligned} \quad (3.4)$$

Where the last step uses integration by parts. The term $\delta \mathbf{X}$ is taken to be zero on the boundary, so we are left with the Euler-Lagrange equations of motion

$$-\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{\mathbf{X}}} \right) + \frac{\partial L}{\partial \mathbf{X}} = 0, \quad (3.5)$$

which, for our rigid body Lagrangian gives

$$\frac{d}{dt} \left(M(\mathbf{X})\dot{\mathbf{X}} \right) = -\frac{\partial V}{\partial \mathbf{X}} + \frac{1}{2}\dot{\mathbf{X}} \cdot \frac{\partial M}{\partial \mathbf{X}} \dot{\mathbf{X}}. \quad (3.6)$$

For the linear term this gives Newton's second law for the linear motion of the centre of mass

$$m\ddot{\mathbf{x}} = -\frac{\partial V}{\partial \mathbf{x}} = \mathbf{f}. \quad (3.7)$$

for force \mathbf{f} .

The angular degrees of freedom give

$$\frac{d}{dt} \left(I(\boldsymbol{\theta})\dot{\boldsymbol{\theta}} \right) = -\frac{\partial V}{\partial \boldsymbol{\theta}} + \frac{1}{2}\dot{\boldsymbol{\theta}} \cdot \frac{\partial I}{\partial \boldsymbol{\theta}} \dot{\boldsymbol{\theta}} \quad (3.8)$$

This approach can be used (see, for example [6]) to recover Euler's equations

$$I(\boldsymbol{\theta})\dot{\boldsymbol{\omega}} + \boldsymbol{\omega} \times (I(\boldsymbol{\theta})\boldsymbol{\omega}) = \boldsymbol{\tau} \quad (3.9)$$

for a torque $\boldsymbol{\tau}$, but the algebra is complicated due to the third order tensor term $\frac{\partial I}{\partial \boldsymbol{\theta}}$.

Instead, we can improve on the standard Lagrangian approach when dealing with actions that are symmetric under some group action by making a variation that preserves the group symmetry. In general, this approach leads to the *Euler-Poincaré equations of motion*.

Start by considering the variation of a rotation matrix R . Since $R^T R = 1$ we have

$$\delta R^T R + R^T \delta R = 0, \quad (3.10)$$

so

$$(R^T \delta R)^T + (R^T \delta R) = 0 \quad (3.11)$$

meaning that $R^T \delta R$ must be antisymmetric. Say

$$R^T \delta R = [\eta]_{\times}. \quad (3.12)$$

Now we use the body space angular velocity from (2.12) $\boldsymbol{\omega}_b = R^T \dot{R}$ So

$$\begin{aligned} [\delta \boldsymbol{\omega}_b]_{\times} &= \delta R^T \dot{R} + R^T \delta \dot{R}, \\ &= (R[\eta]_{\times})^T \dot{R} + R^T \frac{d}{dt}(R[\eta]_{\times}), \\ &= -[\eta]_{\times} [\boldsymbol{\omega}_b]_{\times} + [\boldsymbol{\omega}_b]_{\times} [\eta]_{\times} + [\dot{\eta}]_{\times}, \\ &= [\boldsymbol{\omega}_b \times \eta]_{\times} + [\dot{\eta}]_{\times}. \end{aligned} \quad (3.13)$$

So the variation of $\boldsymbol{\omega}_b$ is

$$\delta \boldsymbol{\omega}_b = \dot{\eta} + \boldsymbol{\omega}_b \times \eta. \quad (3.14)$$

We now write the Lagrangian in terms of $\boldsymbol{\omega}_b$

$$L(\boldsymbol{\omega}_b) = \frac{1}{2} \boldsymbol{\omega}_b \cdot I_b \boldsymbol{\omega}_b \quad (3.15)$$

and take the variation of the action

$$\begin{aligned} \delta S &= \int dt \delta L(\boldsymbol{\omega}_b), \\ &= \int dt \frac{\partial L}{\partial \boldsymbol{\omega}_b} \delta \boldsymbol{\omega}_b, \\ &= \int dt (I_b \boldsymbol{\omega}_b) \cdot (\dot{\eta} + \boldsymbol{\omega}_b \times \eta), \\ &= \int dt \left(-\frac{d}{dt} (I_b \boldsymbol{\omega}_b) + (I_b \boldsymbol{\omega}_b) \times \boldsymbol{\omega}_b \right) \cdot \eta + [(I_b \boldsymbol{\omega}_b) \cdot \eta]_{\text{boundary}}, \end{aligned} \quad (3.16)$$

which then gives the Euler equations of motion in body space

$$I_b \dot{\boldsymbol{\omega}}_b + \boldsymbol{\omega}_b \times (I_b \boldsymbol{\omega}_b) = 0. \quad (3.17)$$

This can be transformed into world space by multiplying on the left by R

$$I(\boldsymbol{\theta})\dot{\boldsymbol{\omega}} + \boldsymbol{\omega} \times (I(\boldsymbol{\theta})\boldsymbol{\omega}) = \boldsymbol{\tau}, \quad (3.18)$$

where we have added back the torque term $\boldsymbol{\tau}$ for completeness.

Now, since we can express the inertia in terms of a body space diagonal matrix I_0 and the orientation of the body $R(\boldsymbol{\theta})$ as

$$I(\boldsymbol{\theta}) = R(\boldsymbol{\theta})I_0R(\boldsymbol{\theta})^T \quad (3.19)$$

then the time derivative of the inertia is

$$\begin{aligned} \dot{I}(\boldsymbol{\theta}) &= \frac{d}{dt}R(\boldsymbol{\theta})I_0, \\ &= \dot{R}(\boldsymbol{\theta})I_0R(\boldsymbol{\theta}) + R(\boldsymbol{\theta})I_0\dot{R}(\boldsymbol{\theta})^T, \\ &= \dot{R}(\boldsymbol{\theta})R(\boldsymbol{\theta})^T I(\boldsymbol{\theta}) + I(\boldsymbol{\theta})R(\boldsymbol{\theta})\dot{R}(\boldsymbol{\theta})^T. \\ &= \dot{R}(\boldsymbol{\theta})R(\boldsymbol{\theta})^T I(\boldsymbol{\theta}) - I(\boldsymbol{\theta}) \left(\dot{R}(\boldsymbol{\theta})R(\boldsymbol{\theta})^T \right), \\ &= [\boldsymbol{\omega}]_{\times} I(\boldsymbol{\theta}) - I(\boldsymbol{\theta}) [\boldsymbol{\omega}]_{\times}, \end{aligned} \quad (3.20)$$

where we have used $\frac{d}{dt}(RR^T) = \dot{R}R^T + R\dot{R}^T = 0$.

Using this, we can write Euler's equations of motion as

$$\frac{d}{dt}(I(\boldsymbol{\theta})\boldsymbol{\omega}) = \boldsymbol{\tau} \quad (3.21)$$

and the combined equations of motion as

$$\frac{d}{dt}(M(\mathbf{X})\boldsymbol{\nu}) = \mathbf{F} \quad (3.22)$$

where

$$\boldsymbol{\nu} = \begin{pmatrix} \dot{\mathbf{x}} \\ \boldsymbol{\omega} \end{pmatrix} \quad (3.23)$$

is the velocity vector and

$$\mathbf{F} = \begin{pmatrix} \mathbf{f} \\ \boldsymbol{\tau} \end{pmatrix} \quad (3.24)$$

is the generalized force vector. I.e. Force is equal to the derivative of the momentum $M(\mathbf{X})\boldsymbol{\nu}$.

This can be trivially generalized to a system of N bodies:

Going back to the Euler-Poincaré formulation it is simpler to see where the factor of J_L^{-T} multiplying the generalized external torque τ_θ and the constraint torque $\left(\frac{\partial \mathbf{c}}{\partial \theta}\right)^T \boldsymbol{\lambda}$ comes from.

With the addition of the constraints and a potential term $V(\theta)$, the reduced Lagrangian becomes

$$L(\boldsymbol{\omega}_b, \theta, \boldsymbol{\lambda}) = \frac{1}{2} \boldsymbol{\omega}_b \cdot I_b \boldsymbol{\omega}_b - V(\theta) + \boldsymbol{\lambda} \cdot \mathbf{c}(\theta). \quad (3.32)$$

When we vary the Lagrangian, in addition to varying with respect to $\boldsymbol{\omega}_b$ we have to vary with respect to $\boldsymbol{\lambda}$ and θ . The variation with respect to $\boldsymbol{\lambda}$ just gives leads to the constraint equations $\mathbf{c}(\theta) = 0$. For the variation with respect to θ we need to compute $\delta\theta$ in terms of the variation η . The left Jacobian comes from perturbing the rotation

$$\begin{aligned} R(\theta + \delta\theta) &= e^{[\theta + \delta\theta]_\times}, \\ &\approx e^{[J_L(\theta)\delta\theta]_\times} e^{[\theta]_\times}, \\ &\approx R(\theta) + [J_L(\theta)\delta\theta]_\times R(\theta). \end{aligned} \quad (3.33)$$

So

$$\delta R R^T = [J_L \delta\theta]. \quad (3.34)$$

But from (3.12) we had $R^T \delta R = [\eta]$ so

$$[J_L \delta\theta]_\times = R[\eta]R^T = [R\eta]_\times \quad (3.35)$$

which gives

$$\delta\theta = J_L^{-1} R\eta. \quad (3.36)$$

Using this

$$\frac{\partial L}{\partial \theta} \delta\theta = \left(R^T J_L^{-T} \left(-\frac{\partial V}{\partial \theta} + \left(\frac{\partial \mathbf{c}}{\partial \theta} \right)^T \boldsymbol{\lambda} \right) \right) \cdot \eta \quad (3.37)$$

Plugging this into the variation of the action, we get

$$\begin{aligned} \delta S &= \int dt \delta L(\boldsymbol{\omega}_b), \\ &= \int dt \frac{\partial L}{\partial \boldsymbol{\omega}_b} \delta \boldsymbol{\omega}_b, \\ &= \int dt (I_b \boldsymbol{\omega}_b) \cdot (\dot{\eta} + \boldsymbol{\omega}_b \times \eta), \\ &= \int dt \left(-\frac{d}{dt} (I_b \boldsymbol{\omega}_b) + (I_b \boldsymbol{\omega}_b) \times \boldsymbol{\omega}_b + R^T J_L^{-T} \left(-\frac{\partial V}{\partial \theta} + \left(\frac{\partial \mathbf{c}}{\partial \theta} \right)^T \boldsymbol{\lambda} \right) \right) \cdot \eta + \mathbf{c} \cdot \delta \boldsymbol{\lambda}, \end{aligned} \quad (3.38)$$

The term in the brackets is zero, so, multiplying through by R to convert to world space, we get the equations of motion

$$I(\theta) \dot{\boldsymbol{\omega}} + \boldsymbol{\omega} \times (I(\theta) \boldsymbol{\omega}) = J_L^{-T} \left(\boldsymbol{\tau}_\theta + \left(\frac{\partial \mathbf{c}}{\partial \theta} \right)^T \boldsymbol{\lambda} \right). \quad (3.39)$$

Rewriting using (3.20) gives

$$\frac{d}{dt} (I(\theta) \boldsymbol{\omega}) = \boldsymbol{\tau}_w + J_L^{-T} \left(\frac{\partial \mathbf{c}}{\partial \theta} \right)^T \boldsymbol{\lambda}, \quad (3.40)$$

We can simplify the notation:

Define the *left-trivialized differential* of the constraint function as

$$\mathcal{D}c = \frac{\partial c}{\partial \boldsymbol{\theta}} J_L^{-1}. \quad (3.41)$$

in this way

$$\dot{c} = \frac{\partial c}{\partial \boldsymbol{\theta}} \dot{\boldsymbol{\theta}} = \frac{\partial c}{\partial \boldsymbol{\theta}} J_L^{-1} \boldsymbol{\omega} = (\mathcal{D}c) \boldsymbol{\omega}. \quad (3.42)$$

The left trivialized differential transforms the derivative from a rate of change with respect to the axis/angle coordinates $\boldsymbol{\theta}$ to a rate of change with respect to the angular velocity $\boldsymbol{\omega}$.

Now looking at linear equation of motion. This is

$$\frac{d}{dt}(m\dot{\mathbf{x}}) = \mathbf{f} + \left(\frac{\partial \mathbf{c}}{\partial \mathbf{x}}\right)^T \boldsymbol{\lambda}. \quad (3.43)$$

We extend the definition of $\mathcal{D}c$ to include positional derivatives and apply to the vector of constraint functions \mathbf{c}

$$\mathcal{D}\mathbf{c} = \begin{pmatrix} \frac{\partial c_i}{\partial \mathbf{x}_i} & \frac{\partial c_i}{\partial \boldsymbol{\theta}_i} J_L(\boldsymbol{\theta}_i)^{-1} & \cdots & \frac{\partial c_i}{\partial \mathbf{x}_N} & \frac{\partial c_i}{\partial \boldsymbol{\theta}_N} J_L(\boldsymbol{\theta}_N)^{-1} \\ \vdots & & \ddots & & \\ \frac{\partial c_M}{\partial \mathbf{x}_i} & \frac{\partial c_M}{\partial \boldsymbol{\theta}_i} J_L(\boldsymbol{\theta}_i)^{-1} & \cdots & \frac{\partial c_M}{\partial \mathbf{x}_N} & \frac{\partial c_M}{\partial \boldsymbol{\theta}_N} J_L(\boldsymbol{\theta}_N)^{-1} \end{pmatrix} \dots \quad (3.44)$$

Using this notation we can summarise all of the equations of motion with constraints as

$$\frac{d}{dt}(M(\mathbf{X})\mathcal{V}) = \mathbf{F} + (\mathcal{D}\mathbf{c})^T \boldsymbol{\lambda}, \quad (3.45)$$

where now

$$M = \begin{pmatrix} m_i & & & & \\ & I(\boldsymbol{\theta}_i) & & & \\ & & \ddots & & \\ & & & m_N & \\ & & & & I(\boldsymbol{\theta}_N) \end{pmatrix}. \quad (3.46)$$

So, if we use the left trivialized derivative $\mathcal{D}\mathbf{c}$ then the angular equations are in the space of angular velocity, rather than the axis/angle space.

3.2 Non-Holonomic Constraints

Non-Holonomic Constraints are constraints that are a function of velocities $\dot{\mathbf{X}}$ as well as configuration \mathbf{X} . These are handled differently than holonomic constraints. As described in [7] and [12], if you try to add the non-holonomic constraints to the Lagrangian and take the variation, as we did with holonomic constraints, you get the incorrect answer. The reason is that, as the Lagrange multipliers represent physical constraint forces, substituting the constraints into the Lagrangian before differentiation effectively assumes that the constraint forces can be derived from a potential, which is not necessarily true for non-holonomic systems.

Instead, we can use the Lagrange-d'Alembert Principle. This states that the non-conservative constraint forces do no virtual work. This modifies our Lagrangian variation by adding an extra term to enforce this condition

$$\delta \int L(\dot{\mathbf{X}}, \mathbf{X}) dt + \int Q \cdot \delta \mathbf{X} = 0 \quad (3.47)$$

where Q are the non-conservative constraint forces.

If we have a non-holonomic constraint of the form $c(\dot{\mathbf{X}}, \mathbf{X}) = \mathbf{a}(\mathbf{X}) \cdot \dot{\mathbf{X}} = 0$ this means that we are preventing displacement in the direction $\mathbf{a}(\mathbf{X})$, so the constraint force Q must be proportional to that direction $Q_i = \lambda_i \mathbf{a}_i(\mathbf{X})$.

Therefore, these non-holonomic constraints can be added to the system as

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{\mathbf{X}}} \right) - \frac{\partial L}{\partial \mathbf{X}} = \mathbf{a} \cdot \boldsymbol{\lambda} = \left(\frac{\partial \mathbf{c}}{\partial \dot{\mathbf{X}}} \right)^T \boldsymbol{\lambda}. \quad (3.48)$$

I.e. exactly the same form as the holonomic constraints, but with the constraint function differentiated with respect to the velocity $\dot{\mathbf{X}}$.

Note that, as with the previous section, we have to be careful about distinguishing between differentiation with respect to the axis/angle variable $\dot{\boldsymbol{\theta}}$ and the angular velocity $\boldsymbol{\omega}$. The constraint torque needs to be multiplied by J_L^{-T} in order to bring it to world space

$$\boldsymbol{\tau}_w = J_L^{-T} \left(\frac{\partial \mathbf{c}}{\partial \dot{\boldsymbol{\theta}}} \right)^T \boldsymbol{\lambda} \quad (3.49)$$

But also

$$\frac{\partial \mathbf{c}}{\partial \dot{\boldsymbol{\theta}}} = \frac{\partial \mathbf{c}}{\partial \boldsymbol{\omega}} \frac{\partial \boldsymbol{\omega}}{\partial \dot{\boldsymbol{\theta}}} = \frac{\partial \mathbf{c}}{\partial \boldsymbol{\omega}} J^L \quad (3.50)$$

since $\boldsymbol{\omega} = J_L \dot{\boldsymbol{\theta}}$. So the constraint torque in world space is

$$\left(\frac{\partial \mathbf{c}}{\partial \boldsymbol{\omega}} \right)^T \boldsymbol{\lambda}. \quad (3.51)$$

As before with the holonomic constraints, we extend the definition of $\dot{\mathcal{D}}$ to include linear derivatives.

The left trivialized velocity derivative is defined as

$$\dot{\mathcal{D}} \mathbf{c} = \begin{pmatrix} \frac{\partial c_i}{\partial \dot{\mathbf{x}}_i} & \frac{\partial c_i}{\partial \boldsymbol{\omega}_i} & \cdots & \frac{\partial c_i}{\partial \dot{\mathbf{x}}_N} & \frac{\partial c_i}{\partial \boldsymbol{\omega}_N} \\ \vdots & & \ddots & & \\ \frac{\partial c_M}{\partial \dot{\mathbf{x}}_i} & \frac{\partial c_M}{\partial \boldsymbol{\omega}_i} & \cdots & \frac{\partial c_M}{\partial \dot{\mathbf{x}}_N} & \frac{\partial c_M}{\partial \boldsymbol{\omega}_N} \end{pmatrix}. \quad (3.52)$$

With this, the constrained equations of motion for both holonomic and non-holonomic constraints are

$$\frac{d}{dt} (M(\mathbf{X}) \mathcal{V}) = \mathbf{F} + (\mathcal{D} \mathbf{c}_h)^T \boldsymbol{\lambda}_h + (\dot{\mathcal{D}} \mathbf{c}_{nh})^T \boldsymbol{\lambda}_{nh}, \quad (3.53)$$

4 Solving the Constrained Equations of Motion

In this section we show how to solve the system of equations of motion as a first order approximation and show the relationship between this solution and

1. A Pseudo-Inverse approach to solving the constraint equations
2. A Sequential Impulse approach
3. Traditional Position Based Dynamics

4.1 Solving the Unconstrained Equations of Motion

We solve the constrained equations of motion incrementally as a first order approximation to the solution to the unconstrained (zero constraint forces) equations, so first we will briefly cover integrating the unconstrained system.

Finding the incremental time advancement of the unconstrained equations of motion to find the provisional state \mathbf{X}^{t+1} , $\dot{\mathbf{X}}^{t+1}$ can be done with a standard numerical integration scheme. The standard in real time simulation is a semi-implicit (symplectic) Euler integration scheme due to it's stability. In this simple scheme, first the velocity is updated, then the new velocity it used to update the position.

For linear degrees of freedom the equation of motion is

$$m\ddot{\mathbf{x}} = \mathbf{f} \quad (4.1)$$

and the discrete integration step for symplectic Euler is

$$\begin{aligned} \dot{\mathbf{x}}^{t+1} &= \dot{\mathbf{x}}^t + \frac{1}{m} \mathbf{f} \Delta t, \\ \mathbf{x}^{t+1} &= \mathbf{x}^t + \dot{\mathbf{x}}^{t+1} \Delta t. \end{aligned} \quad (4.2)$$

For angular degrees of freedom the equation of motion is

$$\frac{d}{dt}(I\boldsymbol{\omega}) = \boldsymbol{\tau} \quad (4.3)$$

and the discrete integration step for symplectic Euler is

$$\begin{aligned} \boldsymbol{\omega}^{t+1} &= \boldsymbol{\omega}^t + I^{-1} (\boldsymbol{\tau} - \boldsymbol{\omega}^t \times (I\boldsymbol{\omega}^t)) \Delta t, \\ \mathbf{q}^{t+1} &= e^{\frac{1}{2} I \boldsymbol{\omega}^{t+1} \Delta t} \mathbf{q}^t \end{aligned} \quad (4.4)$$

where the inertia $I(t)$ is evaluated at the current configuration at t and we are using the quaternion representation for orientation. We could equally choose to use the rotation matrix representation where

$$R^{t+1} = e^{[\boldsymbol{\omega}^{t+1} \Delta t]_{\times}} R^t \approx (\mathbb{1} + [\boldsymbol{\omega}^{t+1} \Delta t]_{\times}) R^t \quad (4.5)$$

but prefer the quaternion approach for reasons discussed in section 2.2.

Note that, for simplicity, we have shown the gyroscopic term $\boldsymbol{\omega} \times I\boldsymbol{\omega}$ evaluated at time t , i.e. it is treated explicitly. In practice, if the angular velocity is high and the inertia term is non-spherical, then this can cause instability, so an implicit integration scheme (evaluating at $t+1$) or sub-stepping over smaller time steps may be required. However, the remainder of this section only assumes that a solution at $t + 1$ can be found that is first order accurate in Δt , and the subsequent derivations not affected by which scheme is used.

4.2 Finding a First Order Approximate Solution

Now, given the solution to the unconstrained equations of motion, we assume that the solution to the constrained equations at $t + 1$ is given by a small perturbation $\Delta \mathbf{X}$ away from the solution to the unconstrained equations at $t + 1$. I.e. for $\boldsymbol{\lambda} = 0$ we solve the unconstrained equations to find $\mathbf{X}^{t+1}, \dot{\mathbf{X}}^{t+1}$ using a standard rigid body integration scheme, then we apply a perturbation $\Delta \mathbf{X}$ where

$$\Delta \mathbf{X} = \begin{pmatrix} \Delta \mathbf{x}_i \\ \Delta \boldsymbol{\omega}_i \\ \vdots \\ \Delta \mathbf{x}_N \\ \Delta \boldsymbol{\omega}_N \end{pmatrix} \quad (4.6)$$

and we define

$$\Delta \boldsymbol{\omega}_i = J_L(\boldsymbol{\theta}_i) \Delta \boldsymbol{\theta}_i. \quad (4.7)$$

This is interpreted as an angular displacement. I.e. small change in angular velocity integrated over time Δt .

With this definition, the updated configuration is

$$\begin{aligned} \mathbf{X}_i^{t+1} &\rightarrow \mathbf{X}_i^{t+1} + \Delta \mathbf{X}_i, \\ \dot{\mathbf{X}}_i^{t+1} &\rightarrow \dot{\mathbf{X}}_i^{t+1} + J_L(\mathbf{X}_i)^{-1} \Delta \mathbf{X}_i / \Delta t. \end{aligned} \quad (4.8)$$

We solve the constrained equations to first order in $\Delta \mathbf{X}$ and $\boldsymbol{\lambda}$ to find an approximate solution to the constrained equations at $t + 1$. As long as the time step is small, and so the displacement due to applying the constraint forces over Δt is small, then this remains a valid approximation.

Start with the linear degrees of freedom. We have

$$\begin{aligned} \frac{1}{\Delta t} (m(\dot{\mathbf{x}}^{t+1} \Delta t + \Delta \mathbf{x}) - m\dot{\mathbf{x}}^{t+1} \Delta t) &= \mathbf{f}(\mathbf{x}^{t+1} + \Delta \mathbf{x}) \Delta t + (\mathcal{D}\mathbf{c})^T \boldsymbol{\lambda} \Delta t \\ &\Rightarrow M \Delta \mathbf{x} = (\mathcal{D}\mathbf{c})^T \boldsymbol{\lambda} \Delta t^2, \end{aligned} \quad (4.9)$$

where

$$M = m\mathbb{1} - \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \Delta t^2 \quad (4.10)$$

as we have assumed that the unconstrained equation is satisfied at $(\mathbf{x}^{t+1}, \dot{\mathbf{x}}^{t+1})$. Now expand the constraint equation, as this must be satisfied at $\mathbf{x}^{t+1} + \Delta \mathbf{x}$

$$\mathbf{c}(\mathbf{x}^{t+1} + \Delta \mathbf{x}) = \mathbf{c}(\mathbf{x}) + (\mathcal{D}\mathbf{c}) \Delta \mathbf{x} + \mathcal{O}(\Delta \mathbf{x})^2 = 0. \quad (4.11)$$

Taking these together we find an equation for $\boldsymbol{\lambda}$

$$(\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T) \boldsymbol{\lambda} \Delta t^2 = -\mathbf{c}(\mathbf{x}^{t+1}). \quad (4.12)$$

We solve this for $\boldsymbol{\lambda}$ and then obtain $\Delta \mathbf{x}$ from

$$\Delta \mathbf{x} = M^{-1} (\mathcal{D}\mathbf{c})^T \boldsymbol{\lambda} \Delta t^2. \quad (4.13)$$

Now we do the same for the angular equations of motion. Start with

$$\frac{d}{dt} (I(\boldsymbol{\theta}) \boldsymbol{\omega}) = \boldsymbol{\tau}(\boldsymbol{\theta}) + (\mathcal{D}\mathbf{c})^T \boldsymbol{\lambda}. \quad (4.14)$$

Expanding, using the assumption that $\boldsymbol{\theta}^{t+1}, \dot{\boldsymbol{\theta}}^{t+1}$ solve the unconstrained equation, we get

$$\frac{1}{\Delta t} (I(\boldsymbol{\theta}^{t+1} + \Delta\boldsymbol{\theta})(\boldsymbol{\omega}^{t+1}\Delta t + \Delta\boldsymbol{\omega}) - I(\boldsymbol{\theta}^{t+1})\boldsymbol{\omega}^{t+1}\Delta t) = \boldsymbol{\tau}_w(\boldsymbol{\theta}^{t+1} + \Delta\boldsymbol{\theta})\Delta t + (\mathcal{D}\mathbf{c})^T \boldsymbol{\lambda}\Delta t$$

\Rightarrow

$$\begin{aligned} & (\mathbb{1} + [\Delta\boldsymbol{\omega}]_{\times})I(\boldsymbol{\theta}^{t+1})(\mathbb{1} - [\Delta\boldsymbol{\omega}]_{\times})(\boldsymbol{\omega}^{t+1}\Delta t + \Delta\boldsymbol{\omega}) - I(\boldsymbol{\theta}^{t+1})\boldsymbol{\omega}^{t+1}\Delta t \\ & = \boldsymbol{\tau}_w(\boldsymbol{\theta}^{t+1})\Delta t^2 + (\mathcal{D}\boldsymbol{\tau})^T \Delta\boldsymbol{\omega}\Delta t^2 + (\mathcal{D}\mathbf{c})^T \boldsymbol{\lambda}\Delta t^2. \end{aligned} \quad (4.15)$$

which simplifies to

$$\tilde{I}\Delta\boldsymbol{\omega} = (\mathcal{D}\mathbf{c})^T \boldsymbol{\lambda}\Delta t^2, \quad (4.16)$$

Where

$$\tilde{I} = I + I[\boldsymbol{\omega}^{t+1}\Delta t]_{\times} - [I\boldsymbol{\omega}^{t+1}\Delta t]_{\times} - (\mathcal{D}\boldsymbol{\tau})^T \Delta t^2 \quad (4.17)$$

is the modified inertia matrix. Note that for spherical inertia or for small $\boldsymbol{\omega}^{t+1}\Delta t$ this reduces to the standard world space inertia. This modification is a result of the precession term in the equations of motion and means, for instance, a spinning body configured as a gyroscope will be more difficult to rotate when applying the $\Delta\boldsymbol{\theta}$ due to the constraint correction.

In practice, this term can usually be ignored, as the time step Δt is taken to be small and so $|\boldsymbol{\omega}\Delta t| \ll 1$ and usually inertias in rigid body simulation are reasonably uniform $I_{ii}/I_{jj} \sim 1$. These assumptions break down for fast spinning objects with highly non-spherical inertia, where gyroscopic effects are important. In this case, although the motion can be correctly modelled if we include the $\boldsymbol{\omega} \times I\boldsymbol{\omega}$ term in the integration, constraint interactions will be incorrect - e.g. trying to tip a spinning gyroscope off of its rotation axis with a contact constraint, the resistance should scale like $(\boldsymbol{\omega}\Delta t)^2$ but if we ignore this term it will be just as easy to push as if it wasn't spinning.

A perhaps more important argument for ignoring this term is that its addition makes the mass matrix non-symmetric, which means that the matrix $\mathcal{D}\mathbf{c} I^{-1}(\mathcal{D}\mathbf{c})^T$ is not necessarily symmetric. This means that the system is no longer guaranteed to converge, as we explore in the next section. For these reasons, we will omit this correction term in our analysis from now on, and assume $\tilde{I} = I$. It is easy to add it back though by just replacing I with \tilde{I} , but the convergence analysis will not be valid in this case.

Back to the constraint function, we again expand to first order

$$\mathbf{c}(\boldsymbol{\theta}^{t+1} + \Delta\boldsymbol{\theta}) = \mathbf{c}(\boldsymbol{\theta}^{t+1}) + \mathcal{D}\mathbf{c} \Delta\boldsymbol{\omega} + \mathcal{O}(\Delta t^2) = 0, \quad (4.18)$$

and solving gives

$$\begin{aligned} & \left(\mathcal{D}\mathbf{c} I^{-1} (\mathcal{D}\mathbf{c})^T \right) \boldsymbol{\lambda}\Delta t^2 = -\mathbf{c}(\boldsymbol{\theta}^{t+1}), \\ & \Delta\boldsymbol{\omega} = I^{-1} (\mathcal{D}\mathbf{c})^T \boldsymbol{\lambda}\Delta t^2. \end{aligned} \quad (4.19)$$

This brings the angular and linear equations into the same form. We can then combine all degrees of freedom together into a single linear system.

First solve the following linear system to find $\boldsymbol{\lambda}$

$$\left(\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T\right) \boldsymbol{\lambda} \Delta t^2 = -c(\mathbf{X}^{t+1}) \quad (4.20)$$

and then use this to solve for $\Delta\mathbf{X}$

$$\Delta\mathbf{X} = M^{-1} (\mathcal{D}\mathbf{c})^T \boldsymbol{\lambda} \Delta t^2, \quad (4.21)$$

This equation for $\boldsymbol{\lambda}$ which is then used to solve for $\Delta\mathbf{X}$ is a linear system that can be solved using a number of numerical methods.

4.3 Regularization

A useful extension to the method of Lagrange multipliers is a method of regularization by applying a quadratic term in the Lagrange multiplier to the Lagrangian. The idea is to add a term $\frac{1}{2}\epsilon|\boldsymbol{\lambda}|^2$ to the Lagrangian with ϵ being a small constant that ensures the Lagrange multipliers (and hence the constraint forces) remain small

$$L(\mathbf{X}, \dot{\mathbf{X}}, \boldsymbol{\lambda}) \rightarrow \frac{1}{2} \dot{\mathbf{X}} \cdot M(\mathbf{X}) \dot{\mathbf{X}} - V(\mathbf{X}) + \boldsymbol{\lambda} \cdot \mathbf{c}(\mathbf{X}) + \frac{1}{2} \epsilon |\boldsymbol{\lambda}|^2. \quad (4.22)$$

This new term alters the constraint equation

$$\mathbf{c}(\mathbf{X}^{t+1} + \Delta\mathbf{X}) = \mathbf{c}(\mathbf{X}^{t+1}) + \mathcal{D}\mathbf{c} \Delta\mathbf{X} + \epsilon\boldsymbol{\lambda} = 0. \quad (4.23)$$

The result of regularization is that equation to solve for $\boldsymbol{\lambda}$ now becomes

$$\left(\epsilon\mathbb{1} + \mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T\right) \boldsymbol{\lambda} \Delta t^2 = -c(\mathbf{X}^{t+1}). \quad (4.24)$$

for some small ϵ

This addition term ensures that if $\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T$ becomes singular, for instance if there are degenerate constraints, then the inversion of the matrix for $\boldsymbol{\lambda}$ does not also become singular. This is particularly important if the matrix solver uses a direct solve for all or part of the matrix.

This type of regularization is generally called *Tikhonov Regularization* [22], and is equivalent to compliance in XPBD, as we will see later.

4.4 Interpretation: Connection to the Moore-Penrose Inverse and Least Squares Minimization

The Moore-Penrose inverse [18] (also called the pseudoinverse) of a matrix A^+ is a generalization of the matrix inverse for non-square matrices. For a full row rank matrix A (so that (AA^T) is invertible) this is defined as

$$A^+ = A^T(AA^T)^{-1}. \quad (4.25)$$

The Lagrange multiplier method gave us a solution for $\Delta \mathbf{X}$ of

$$\Delta \mathbf{X} = -M^{-1} (\mathcal{D}\mathbf{c})^T \left(\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right)^{-1} \mathbf{c}(\mathbf{X}^{t+1}). \quad (4.26)$$

If M were equal to the identity then this would reduce to

$$\Delta \mathbf{X} = -(\mathcal{D}\mathbf{c})^T \left(\mathcal{D}\mathbf{c} (\mathcal{D}\mathbf{c})^T \right)^{-1} \mathbf{c}(\mathbf{X}^{t+1}). \quad (4.27)$$

This is exactly the Moore-Penrose inverse solution to the constraint equation at first order in $\Delta \mathbf{X}$

$$\mathcal{D}\mathbf{c} \Delta \mathbf{X} = -\mathbf{c}(\mathbf{X}^{t+1}). \quad (4.28)$$

So the method of Lagrange multipliers applied to the first order approximation of the constrained equations of motion reduces to a mass weighted version of the pseudoinverse of the constraint Jacobian.

The linearization of the constrained equations of motion is equivalent to solving the constraint equations subject to a least squares minimization of $\Delta \mathbf{X}$ with a distance measure using “metric” M .

$$\min_{\Delta \mathbf{X}} \frac{1}{2} \Delta \mathbf{X}^T \cdot M \Delta \mathbf{X} \quad (4.29)$$

subject to the constraint

$$\mathcal{D}\mathbf{c} \Delta \mathbf{X} = -\mathbf{c}(\mathbf{X}^{t+1}). \quad (4.30)$$

We are finding a solution that satisfies the constraint equations while minimizing the energy term $\frac{1}{2} \Delta \mathbf{X} \cdot M \Delta \mathbf{X}$.

This is an extension of the Moore-Penrose Pseudoinverse method $A^+ = A^T (AA^T)^{-1}$ to a space with an inner product \langle, \rangle_M .

Tikhonov regularization can be used together with the pseudoinverse [4] for ill posed problems where $(AA^T)^{-1}$ does not exist. In this case for small ϵ the inverse is corrected by

$$A^+ = A^T (\epsilon \mathbb{1} + AA^T)^{-1} \quad (4.31)$$

and, applying this to the solution of the constraint equations with a metric M we recover our previous result

$$\Delta \mathbf{X} = -M^{-1} (\mathcal{D}\mathbf{c})^T \left(\epsilon \mathbb{1} + \mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right)^{-1} \mathbf{c}(\mathbf{X}^{t+1}) \quad (4.32)$$

4.5 An Iterative Solution Using Gauss-Seidel

We now turn to the method for solving the system of equations iteratively. We need to solve λ from

$$\left(\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right) \lambda \Delta t^2 = -\mathbf{c}(\mathbf{X}^{t+1}), \quad (4.33)$$

which can then be used to compute $\Delta \mathbf{X}$ using

$$\Delta \mathbf{X} = M^{-1} (\mathcal{D}\mathbf{c})^T \lambda \Delta t^2. \quad (4.34)$$

The equation is a linear system $A\boldsymbol{\lambda} = -\mathbf{c}$, where $A = (\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T) \Delta t^2$ is a square matrix. Writing out the rows explicitly, this is

$$\begin{aligned} A_{11}\lambda_1 + A_{12}\lambda_2 + \cdots + A_{1m}\lambda_m &= -c_1, \\ A_{21}\lambda_1 + A_{22}\lambda_2 + \cdots + A_{2m}\lambda_m &= -c_2, \\ &\vdots \\ A_{m1}\lambda_1 + A_{m2}\lambda_2 + \cdots + A_{mm}\lambda_m &= -c_m \end{aligned}$$

which can be solved using the Gauss-Seidel method. See, for example, [24].

We start with an initial guess of $\boldsymbol{\lambda}^{(0)} = 0$, then solve each equation i in turn for λ_i using the latest value from previous solutions, and then iterate to a solution:

$$\begin{aligned} \lambda_1^{(1)} &= \frac{1}{A_{11}} \left(-c_1 - A_{12}\lambda_2^{(0)} - \cdots - A_{1m}\lambda_m^{(0)} \right), \\ \lambda_2^{(1)} &= \frac{1}{A_{22}} \left(-c_2 - A_{21}\lambda_1^{(1)} - A_{23}\lambda_3^{(0)} \cdots - A_{2m}\lambda_m^{(0)} \right), \\ &\vdots \\ \lambda_i^{k+1} &= \frac{1}{A_{ii}} \left(-c_i - \sum_{j<i} A_{ij}\lambda_j^{k+1} - \sum_{j>i} A_{ij}\lambda_j^k \right), \\ &= \lambda_i^k - \frac{1}{A_{ii}} (\mathbf{c} + A\boldsymbol{\lambda}^*)_i, \end{aligned} \quad (4.35)$$

where $\boldsymbol{\lambda}^*$ is the latest value of all of the λ s.

For our system this is

$$\lambda_i^{k+1} \Delta t^2 = \lambda_i^k \Delta t^2 - \left(\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right)_{ii}^{-1} \left(\mathbf{c}(\mathbf{X}^{t+1}) + \left(\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right) \boldsymbol{\lambda}^* \Delta t^2 \right)_i. \quad (4.36)$$

Using (4.34) we can update the displacements incrementally per iteration

$$\Delta \mathbf{X}'_j = \Delta \mathbf{X}_j + M_i^{-1} (\mathcal{D}\mathbf{c})_{ij} (\boldsymbol{\lambda}^{k+1} - \boldsymbol{\lambda}^k)_i \Delta t^2 \quad (4.37)$$

and use this to rewrite the update for λ_i as

$$\lambda_i^{k+1} \Delta t^2 = \lambda_i^k \Delta t^2 - \left(\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right)_{ii}^{-1} \left(\mathbf{c}(\mathbf{X}^{t+1}) + \mathcal{D}\mathbf{c} \Delta \mathbf{X} \right)_i. \quad (4.38)$$

Summary:

The equations of motion can be solved using a Gauss-Seidel update. We start with an initial guess for a solution for $\boldsymbol{\lambda}$, typically $\boldsymbol{\lambda} = 0$, and set all of the displacements $\Delta \mathbf{X} = 0$. Then for $k = 0$ to $N - 1$ iterations we solve one row at a time for each constraint's Lagrange multiplier using the current value of the displacement

$$\lambda_i^{k+1} \Delta t^2 = \lambda_i^k \Delta t^2 - \left(\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right)_{ii}^{-1} \left(\mathbf{c}(\mathbf{X}^{t+1}) + \mathcal{D}\mathbf{c} \Delta \mathbf{X} \right)_i. \quad (4.39)$$

and then update the current displacement

$$\Delta \mathbf{X}'_j = \Delta \mathbf{X}_j + M_i^{-1} (\mathcal{D}\mathbf{c})_{ij} (\boldsymbol{\lambda}^{k+1} - \boldsymbol{\lambda}^k)_i \Delta t^2. \quad (4.40)$$

In the end we are left with a set of displacements that resolve the constraints and be use to update the configuration of the system to give a first order solution to the constrained equations of motion at time $t + 1$

$$\begin{aligned} \dot{\mathbf{x}}_i^{t+1} &\rightarrow \mathbf{x}_i^{t+1} + \Delta \mathbf{x}_i / \Delta t, \\ \boldsymbol{\omega}_i^{t+1} &\rightarrow \boldsymbol{\omega}_i^{t+1} + \Delta \boldsymbol{\omega}_i / \Delta t, \\ \mathbf{x}^{t+1} &\rightarrow \mathbf{x}_i + \Delta \mathbf{x}_i, \\ q(\boldsymbol{\theta}^{t+1}) &\rightarrow e^{\frac{1}{2} I \Delta \boldsymbol{\omega}_i} q(\boldsymbol{\theta}^{t+1}) \end{aligned} \quad (4.41)$$

4.6 Comparison to the Jacobi Method

The Jacobi method is similar to the Gauss-Seidel method except that instead of feeding the result from the previous row into the solution for the next row, each row is solved independently using the solution from the previous iteration. In our system this means instead of the Gauss-Seidel update

$$\lambda_i^{k+1} = \frac{1}{A_{ii}} \left(-c_i - \sum_{j < i} A_{ij} \lambda_j^{k+1} - \sum_{j > i} A_{ij} \lambda_j^k \right), \quad (4.42)$$

we would have a Jacobi update of

$$\lambda_i^{k+1} = \frac{1}{A_{ii}} \left(-c_i - \sum_{j \neq i} A_{ij} \lambda_j^k \right). \quad (4.43)$$

The Jacobi update is known to converge slower than Gauss-Seidel, which we will demonstrate with a simple example in section 8, but as solving each λ is independent, the update can be parallelized much more easily than Gauss-Seidel, making it popular for GPU based simulations.

4.7 Convergence

It's useful to look at the convergence criteria for the Gauss-Seidel update of the linear system (4.20). If we take the unmodified inertia, then the matrix $A = \mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T$ is symmetric positive definite (SPD), since M is SPD. I.e. for any \mathbf{x} , $\mathbf{x} \cdot A\mathbf{x} \geq 0$. A is strictly positive definite if $(\mathcal{D}\mathbf{c})^T \mathbf{x} = 0$ implies that $\mathbf{x} = 0$. I.e. $\mathcal{D}\mathbf{c}$ has full row rank, meaning that all of the constraints are linearly independent. The Gauss-Seidel update will converge for any SPD matrix A , so our system will always converge if all of the constraints are linearly independent.

To show this, we construct an energy function

$$E(\boldsymbol{\lambda}) = \frac{1}{2} \boldsymbol{\lambda}^T A \boldsymbol{\lambda} + \boldsymbol{\lambda}^T \mathbf{c}. \quad (4.44)$$

This is a convex function for SPD matrix A with a unique minimum at $A\boldsymbol{\lambda}^* = -\mathbf{c}$. Taking a single λ_i and fixing the energy for all of the other λ_j s the energy as a function of λ_i is

$$E(\lambda_i) = \frac{1}{2}A_{ii}|\lambda_i|^2 + \lambda_i \left(\sum_{j \neq i} A_{ij}\lambda_j + c_i \right), \quad (4.45)$$

and minimizing with respect to λ_i gives

$$\frac{\partial E(\lambda_i)}{\partial \lambda_i} = A_{ii}\lambda_i + \sum_{i \neq j} A_{ij}\lambda_j + c_i = \mathbf{r}_i = 0 \quad (4.46)$$

for a residual \mathbf{r}_i . The value of λ_i^{k+1} that sets this residual to zero is

$$\begin{aligned} \lambda_i^{k+1} &= -A_{ii}^{-1} \left(\sum_{i \neq j} A_{ij}\lambda_j^k + c_i \right), \\ &= \lambda_i^k - A_{ii}^{-1} (A\boldsymbol{\lambda} + \mathbf{c})_i, \end{aligned} \quad (4.47)$$

i.e. exactly the iterative update we defined in (4.39).

Using $e_i = (A\boldsymbol{\lambda} + \mathbf{c})_i$ we compute the change in energy

$$\begin{aligned} E(\lambda_i^{k+1}) - E(\lambda_i^k) &= \frac{1}{2}A_{ii}|\lambda_i^k - A_{ii}^{-1}e_i|^2 + (\lambda_i^k - A_{ii}^{-1}e_i) \left(\sum_{i \neq j} A_{ij}\lambda_j^k + c_i \right) \\ &\quad - \frac{1}{2}A_{ii}|\lambda_i^k|^2 + \lambda_i^k \left(\sum_{j \neq i} A_{ij}\lambda_j^k + c_i \right), \\ &= -\lambda_i^k e_i + \frac{1}{2}A_{ii}^{-1}|e_i|^2 - A_{ii}^{-1}e_i \left(\sum_{j \neq i} A_{ij}\lambda_j^k + c_i \right), \\ &= -\frac{1}{2}A_{ii}^{-1}|e_i|^2. \end{aligned} \quad (4.48)$$

Since A is SPD then $A_{ii} > 0$ and so the energy decreases each iteration. Combining this with the fact that the energy is bounded below by the minimum $A\boldsymbol{\lambda}^* = -\mathbf{c}$ means that the solution converges to $\boldsymbol{\lambda}^*$.

4.7.1 Rate of convergence

To look at the rate of convergence we write the update as

$$\lambda_i^{k+1} = A_{ii}^{-1} \left(-c_i - \sum_{j < i} A_{ij}\lambda_j^{k+1} - \sum_{j > i} A_{ij}\lambda_j^k \right) \quad (4.49)$$

as, with Gauss-Seidel, the result from solving a previous row of λ_j^{k+1} with $j < i$ is fed into the update for λ_i^{k+1} . So this is

$$A_{ii}\lambda_i^{k+1} + \sum_{j < i} A_{ij}\lambda_j^{k+1} = -\sum_{j > i} A_{ij}\lambda_j^k - c_i. \quad (4.50)$$

Writing the matrix A in terms of its diagonal D , its lower triangular matrix L and upper U , where $U = L^T$ due to A being symmetric, we have

$$(L + D)\boldsymbol{\lambda}^{k+1} = -L^T\boldsymbol{\lambda}^k - \mathbf{c} \quad (4.51)$$

so

$$\boldsymbol{\lambda}^{k+1} = -(L + D)^{-1}L^T\boldsymbol{\lambda}^k - (L + D)^{-1}\mathbf{c}. \quad (4.52)$$

Letting

$$\boldsymbol{\epsilon}^k = \boldsymbol{\lambda}^k - \boldsymbol{\lambda}^* \quad (4.53)$$

be the error at iteration k from the solution $\boldsymbol{\lambda}^*$ defined by $A\boldsymbol{\lambda}^* = -\mathbf{c}$, then

$$\boldsymbol{\epsilon}^{k+1} = -(D + L)^{-1}L^T\boldsymbol{\epsilon}^k. \quad (4.54)$$

So, for an initial error $\boldsymbol{\epsilon}_0$ the error goes like

$$\boldsymbol{\epsilon}^k = G^k\boldsymbol{\epsilon}_0 \quad (4.55)$$

for

$$G = -(D + L)^{-1}L^T. \quad (4.56)$$

This means that the convergence of the system depends on the spectral radius $\rho(G)$. Express $\boldsymbol{\epsilon}$ in terms of the basis of eigenvalues of G

$$\boldsymbol{\epsilon} = \sum_i c_i v_i \quad (4.57)$$

for constants c_i and eigenvalues v_i . From this, if α_i are the eigenvalues of G then

$$\boldsymbol{\epsilon}^k = \sum_i \alpha_i^k c_i v_i. \quad (4.58)$$

So the error becomes dominated by the largest eigenvalue α_i . This needs to be less than one for the error to shrink, hence the convergence criteria is

$$\rho(G) < 1 \quad (4.59)$$

and how fast the system will converge depends on how small the spectral radius of G is.

In practice, it is often more practical to use the diagonal dominance of the matrix A in order to give an intuition on the rate of convergence of the system. Strict diagonal dominance is defined as

$$|A_{ii}| > \sum_{j \neq i} |A_{ij}|. \quad (4.60)$$

The claim is that for a strictly diagonally dominant matrix A the Gauss-Seidel update will converge, and the more diagonally dominant, the faster the rate of convergence. To see this consider the update for the first row

$$\epsilon_1^{k+1} = -A_{11}^{-1} \sum_{j>1} A_{1j} \epsilon_j^k. \quad (4.61)$$

Taking absolute values

$$\begin{aligned} |\epsilon_1^{k+1}| &\leq |A_{11}|^{-1} \sum_{j>1} |A_{1j}| |\epsilon_j^k|, \\ &\leq r_1 \|\boldsymbol{\epsilon}^k\|_{\max}, \end{aligned} \quad (4.62)$$

where

$$r_1 = \frac{\sum_{j>1} A_{1j}}{A_{11}}, \quad (4.63)$$

which for a diagonally dominant matrix is less than 1, and $\|\epsilon^k\|_{\max}$ is the maximum error from the previous iteration. So

$$|\epsilon_1^{k+1}| < \|\epsilon^k\|_{\max}. \quad (4.64)$$

The error for the second row is

$$\epsilon_2^{k+1} = -A_{22}^{-1} \left(A_{21}\epsilon_1^{k+1} + \sum_{j>2} A_{1j}\epsilon_j^k \right). \quad (4.65)$$

Taking absolutes

$$\begin{aligned} |\epsilon_2^{k+1}| &\leq |A_{22}|^{-1} \left(A_{21}|\epsilon_1^{k+1}| + \sum_{j>2} |A_{2j}|\epsilon_j^k \right), \\ &\leq \left(\frac{A_{21}}{|A_{22}|} + \frac{\sum_{j>2} A_{2j}}{|A_{22}|} \right) \|\epsilon^k\|_{\max}, \end{aligned} \quad (4.66)$$

where we have used the previous result for ϵ_1^{k+1} .

Continuing, row by row we get

$$|\epsilon_i^{k+1}| \leq \frac{\sum_{i \neq j} A_{ij}}{A_{ii}} \|\epsilon^k\|_{\max} = r_i \|\epsilon^k\|_{\max} \quad (4.67)$$

and overall

$$\|\epsilon^{k+1}\|_{\max} \leq r \|\epsilon^k\|_{\max} \quad (4.68)$$

where

$$r = \max_i \frac{\sum_{i \neq j} A_{ij}}{A_{ii}}. \quad (4.69)$$

So, if A is strictly diagonally dominant then $r < 1$ and each iteration the error is reduced, hence the system converges.

Summary:

The Gauss-Seidel algorithm for finding the solution to $A\mathbf{x} = \mathbf{b}$ is guaranteed to converge for any symmetric positive definite matrix A . For the constrained equations of motion $A = \mathcal{D}\mathbf{c} M^{-1}(\mathcal{D}\mathbf{c})^T$ is symmetric since mass matrix is symmetric and positive definite if the constraints are linearly independent and so $\mathcal{D}\mathbf{c}$ has full row rank

The rate of convergence depends on the spectral radius $\rho(G)$ of the Gauss-Seidel iteration matrix $G = -(D+L)^{-1}L^T$, where we split $A = (D+L+L^T)$ into its diagonal, lower triangular and upper triangular parts. The per iteration error goes like

$$|\epsilon^k| \sim \rho(G)^k \|\epsilon_0\|. \quad (4.70)$$

Diagonal dominance is determined by the ratio

$$r = \max_i \frac{\sum_{i \neq j} A_{ij}}{A_{ii}}. \quad (4.71)$$

The more diagonally dominant the matrix is (the smaller the value of r) the faster the rate of convergence.

4.8 A More Accurate Update Using Non-Linear Gauss-Seidel

We can improve on the result (4.39) by noticing that the term $\mathbf{c}(\mathbf{X}^{t+1}) + \mathcal{D}\mathbf{c} \Delta\mathbf{X}$ is the first order expansion of $\mathbf{c}(\mathbf{X}^{t+1} + \Delta\mathbf{X})$, i.e. the constraint function evaluated at the most recent value of the displacement $\Delta\mathbf{X}$.

In order to obtain a more accurate update, we can replace the linearized term $\mathbf{c}(\mathbf{X}^{t+1}) + \mathcal{D}\mathbf{c} \Delta\mathbf{X}$ with the full constraint evaluation at the latest version of $\Delta\mathbf{X}$.

$$\lambda_i^{k+1} \Delta t^2 = \lambda_i^k \Delta t^2 - \left(\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right)_{ii}^{-1} \mathbf{c}(\mathbf{X}^{t+1} + \Delta\mathbf{X})_i, \quad (4.72)$$

This turns the update from a linear solution to (4.33) into a non-linear iterative algorithm.

This is now straying from the standard Gauss-Seidel method, but it is demonstrated to converge better when applied in [15]. This non-linear update is particularly important when solving rigid body systems with rotation when the time step is large, as errors due to using a linear approximation to rotations can be large, leading to inaccuracies in the solution and instability.

Another thing that is often done to improve the convergence of the result is to recompute the constraint derivatives in the “resistance” term $\left(\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right)_{ii}^{-1}$ for each iteration. However, this is often expensive, and having an accurate measure of the error is the more important factor in improving stability and the robustness of the result in the case of large angular velocities and/or time steps.

Given that this update strays from the linearized Gauss-Seidel update, we can’t use the same convergence arguments that we did in the previous section. However, if the constraints are smooth then the constraint function is well approximated by its first order approximation near the solution and the nonlinear Gauss-Seidel update converges asymptotically to the linear Gauss-Seidel update and the local convergence analysis holds, and if the constraint row is diagonally dominant then we are still reducing the per iteration error in the non-linear regime.

What we gain by this method though is that we are no longer accumulating linearization error, and so the frame to frame stability in cases of large angular velocity is greatly improved. This is the practical reason for the better convergence observed in [15].

4.9 Relaxation

A well known modification to the Gauss-Seidel algorithm is to introduce a relaxation parameter ξ . This can be used to control how much of the error per constraint is fixed up per iteration, and thus can be used to control the convergence of the system.

To implement relaxation we modify the λ update (4.72) with the addition of a relaxation factor ξ

$$\lambda_i^{k+1} \Delta t^2 = \lambda_i^k \Delta t^2 - \xi \left(\mathcal{D} \mathbf{c} M^{-1} (\mathcal{D} \mathbf{c})^T \right)_{ii}^{-1} \mathbf{c}(\mathbf{X}^{t+1} + \Delta \mathbf{X}^k)_i, \quad (4.73)$$

The relaxation factor scales the per frame correction to λ .

This relaxation factor can help with the convergence of the system. Under-relaxation with $\xi < 1$ can allow for systems that are otherwise divergent to converge, and over-relaxation $\xi > 1$ can help speed up the convergence of an already convergent system.

Previously, in (4.48) we saw the change in energy per iteration $\Delta E = -\frac{1}{2} A_{ii}^{-1} |e_i|^2$. With the addition of the relaxation parameter, this becomes

$$\Delta E = -\xi \left(1 - \frac{\xi}{2} \right) A_{ii}^{-1} |e_i|^2. \quad (4.74)$$

So the energy decreases, and the system converges, if

$$0 < \xi < 2 \quad (4.75)$$

With the additional relaxation factor, the iteration matrix goes to

$$G_\xi = (D + \xi L)^{-1} ((1 - \xi)D - \xi L^T), \quad (4.76)$$

which reduces to the regular Gauss-Seidel iteration matrix for $\xi = 1$.

4.9.1 Under-Relaxation

If $\xi < 1$ this is called under-relaxation. With relaxation the update of the per iteration error (4.68) goes to

$$|\epsilon^{k+1}| \leq (1 - \xi(1 - r_i)) \|\epsilon^k\|_{\max} \quad (4.77)$$

so converges for

$$1 - \xi(1 - r_i) < 1. \quad (4.78)$$

This means that having $\xi < 1$ worsens the convergence rate but at the trade off of greater stability. We are reducing the per constraint per iteration step size, fixing less of the error each frame. This slows convergence but can also reduce potential oscillation between competing constraint targets or errors due to highly non-linear constraints.

4.9.2 Over-Relaxation

If $1 < \xi < 2$ this is called successive over-relaxation, or SOR. We are effectively extrapolating the error, which can speed up convergence, but may cause divergence depending on how large the spectral radius of the iteration matrix is.

The optimal relaxation factor ξ^* for a given system is related to the spectral radius of the Jacobi iteration matrix $\rho(G_J)$ by [25].

$$\xi^* \sim \frac{2}{1 + \sqrt{1 - \rho(G_J)}}. \quad (4.79)$$

where, from (4.43) it can be seen that the Jacobi iteration matrix is

$$G_J = -D^{-1}(L + L^T). \quad (4.80)$$

Under special conditions where the matrix $A = \mathcal{D}\mathbf{c} M^{-1}(\mathcal{D}\mathbf{c})^T$ is consistently ordered, e.g. has a block tri-diagonal structure (as it would have for a simple chain of constrained bodies) the spectral radius of the Jacobi matrix is related to the spectral radius of the Gauss-Seidel matrix by $\rho(G_{GS}) = \rho(G_J)^2$.

If the spectral radius of the iteration matrix is low (which will be true if A is diagonally dominant) then an optimal value of ξ can be found in the range $1 < \xi < 2$ that can help to speed up convergence.

4.9.3 Adaptive Relaxation

A good balance is modify the relaxation per iteration based on the measured convergence of the system. If $|\epsilon^{k+1}| < |\epsilon^k|$ then the system is converging and we can increase ξ . If $|\epsilon^{k+1}| \geq |\epsilon^k|$ the system is diverging and we can try to recover by decreasing ξ .

4.10 Warm Starting

The iteration loop presented in Algorithm 1 sets the Lagrange multipliers to zero at the start of each update. This is a safe choice as the Lagrange multipliers correspond to the constraint forces and the goal is to find a minimum set of constraint forces to simultaneously solve all of the constraints. However, for stable systems where the constraints are not changing from frame to frame, such as solving a chain of connected bodies, the result from the previous frame is often a good guess for what the result will be for the subsequent frame, so we can speed up convergence by starting the next frame's solve using the resulting Lagrange multipliers from the previous frame. For stable systems we are effectively allowing for convergence of the system over multiple frames, which can mean greater stability and reduce the number of iterations required per frame. However, for highly dynamic systems where constraints are changing significantly between frames this can produce instability as we can end up injecting forces into the system that do not solve the new set of constraints well.

In practice, if the set of constraints

4.11 Adding Non-Holonomic Constraints

We saw in 3.2 that the equations of motion including non-holonomic constraints follow the holonomic constraint formulation but now with the constraint force

$$(\dot{\mathcal{D}}\mathbf{c})^T \boldsymbol{\lambda}. \quad (4.81)$$

The derivation of the linearized update directly follows the holonomic constraint derivation, leading to a per iteration update of

$$\begin{aligned} \lambda_i^{k+1} \Delta t^2 &= \lambda_i^k \Delta t^2 - \xi \left(\dot{\mathcal{D}}\mathbf{c} M^{-1} \left(\dot{\mathcal{D}}\mathbf{c} \right)^T \right)_{ii}^{-1} \mathbf{c}(\dot{\mathbf{X}} + \Delta \mathbf{X}^k / \Delta t)_i, \\ \Delta \mathbf{X}^{k+1} &= \Delta \mathbf{X}^k + M^{-1}(\dot{\mathcal{D}}\mathbf{c})^T (\lambda^{k+1} - \lambda^k) \Delta t^2. \end{aligned} \quad (4.82)$$

So with this update we can see that velocity constraints and position constraints can be part of the same solver loop, since they are both incrementally updating the displacement $\Delta \mathbf{X}$. Alternatively,

we can treat the velocity solve as being separate from the position solve and solve the position constraints first, then solve the velocity constraints, with the $\Delta\mathbf{X}$'s from the velocity solve only updating the velocities and not the positions.

Algorithm 1 Solver Update

- 1: **Input:** Predicted state \mathbf{X}^{t+1} , mass matrix M , timestep Δt
- 2: **Initialize:** $\boldsymbol{\lambda}^0 = 0$ (or a fraction of $\boldsymbol{\lambda}$ for warm starting), $\Delta\mathbf{X}^0 = 0$
- 3: **for** each constraint i **do**
- 4: Compute the constraint derivative and effective mass

$$\mu_i = \left(\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right)_{ii}^{-1}$$

- 5: **end for**
- 6: **for** $k = 0$ to $N - 1$ **do**
- 7: **for** each constraint i **do**
- 8: Compute the Lagrange multiplier update using the latest value of the constraint error

$$\lambda_i^{k+1} \Delta t^2 = \lambda_i^k \Delta t^2 - \xi \mu_i c(\mathbf{X}^{t+1} + \Delta\mathbf{X}^k)$$

- 9: Update the incremental displacement

$$\Delta\mathbf{X}_j^{k+1} = \Delta\mathbf{X}_j^k + M_j^{-1} (\mathcal{D}\mathbf{c})_{ij} (\lambda_i^{k+1} - \lambda_i^k) \Delta t^2$$

- 10: **end for**
- 11: **end for**
- 12: **State Update:**

$$\begin{aligned} \mathbf{x}^{t+1} &\rightarrow \mathbf{x}^{t+1} + \Delta\mathbf{x} & \dot{\mathbf{x}}^{t+1} &\rightarrow \dot{\mathbf{x}}^{t+1} + \Delta\mathbf{x}/\Delta t \\ q^{t+1} &\rightarrow e^{\frac{1}{2}I\Delta\omega} q^{t+1} & \boldsymbol{\omega}^{t+1} &\rightarrow \boldsymbol{\omega}^t + \Delta\boldsymbol{\omega}/\Delta t \end{aligned}$$

4.12 Interpretation: Sequential Impulses

We rewrite (4.72) by introducing the quantity

$$\mathbf{j} = (\mathcal{D}\mathbf{c})^T \boldsymbol{\lambda} \Delta t, \tag{4.83}$$

which has dimensions $[MLT^{-1}]$, i.e. of an impulse. Using the position update $\Delta\mathbf{X} = M^{-1} (\mathcal{D}\mathbf{c})^T \Delta t^2 \boldsymbol{\lambda}$, we obtain

$$\mathbf{j} = M \frac{\Delta\mathbf{X}}{\Delta t}, \tag{4.84}$$

so that \mathbf{j} represents the generalized impulse satisfying $\mathbf{j} = M\Delta\dot{\mathbf{X}}$.

Define

$$\mu_i = \left(\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right)_{ii}^{-1}, \tag{4.85}$$

The scalar μ_i is the effective mass, measuring the resistance of the system to motion in the direction that resolves the i -th constraint.

With these substitutions the Gauss–Seidel update becomes

$$\mathbf{j}_i^{k+1} = \mathbf{j}_i^k - \frac{\mu_i}{\Delta t} \left((\mathcal{D}\mathbf{c})^T \mathbf{c}(\mathbf{X}^{t+1} + M^{-1}\mathbf{j}^k \Delta t) \right)_i. \quad (4.86)$$

To relate this to a force-based formulation, consider the quadratic penalty potential

$$V(\mathbf{X}) = \frac{1}{2}k\|\mathbf{c}(\mathbf{X})\|^2. \quad (4.87)$$

Its associated force is

$$\mathbf{F} = -k(\mathcal{D}\mathbf{c})^T \mathbf{c}. \quad (4.88)$$

Substituting into the previous expression yields

$$\mathbf{j}_i^{k+1} = \mathbf{j}_i^k + \left(\frac{\mu_i}{k\Delta t^2} \right) \mathbf{F}^k \Delta t. \quad (4.89)$$

The factor Δt^2 arises from converting a position-level penalty into an impulse-level update. As a dimensional check, if \mathbf{c} represents a distance constraint so that $[\mathbf{c}] = L$, then $[k] = MT^{-2}$ and $[\mu_i] = M$, making $(\mu_i/(k\Delta t^2))$ dimensionless as required.

If we identify $\mathbf{j} = (\mathcal{D}\mathbf{c})^T \boldsymbol{\lambda} \Delta t$ as the impulse produced by the constraint, then the Gauss–Seidel update can be interpreted as incrementally applying impulses based on the per iteration configuration, and is equivalent to the sequential impulse method [8].

4.13 Interpretation: Position Based Dynamics

The equation (4.73) for the Lagrange multipliers λ_i together with the subsequent $\Delta \mathbf{X}$ update (4.34) can be directly related to the original position based dynamics update [17]. In this paper, they dealt with only point particles, so no angular degrees of freedom, and they defined an iterative update of the positions as

$$\delta \mathbf{p}_i = -s w_i \nabla_{\mathbf{p}_i} C(\mathbf{p}_i, \dots, \mathbf{p}_n), \quad (4.90)$$

with

$$s = \frac{C(\mathbf{p}_i, \dots, \mathbf{p}_n)}{\sum_i w_i |\nabla_{\mathbf{p}_i} C(\mathbf{p}_i, \dots, \mathbf{p}_n)|^2} \quad (4.91)$$

and $w_i = 1/m_i$.

In our notation, this is just

$$s = \left(\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right)^{-1} \mathbf{c}(\mathbf{x}^{t+1} + \Delta \mathbf{x}^k) \quad (4.92)$$

and

$$\begin{aligned} \lambda_i^{k+1} &= \lambda_i^k - s, \\ \Delta \mathbf{x}^{k+1} &= \Delta \mathbf{x}^k - M^{-1} (\mathcal{D}\mathbf{c})^T (\lambda_i^{k+1} - \lambda_i^k). \end{aligned} \quad (4.93)$$

In the PBD paper, they introduce a constraint compliance by multiplying the corrections at each iteration by a constant k . In our formulation, we can see that this is equivalent to adding a relaxation $k = \xi$ to the update. So, it is obvious that the compliance is just slowing the convergence of the hard constraint, making it effectively a soft constraint, but with the softness dependent on the number of iterations. The authors address this in their paper using a modified stiffness $k' = 1 - (1 - k)^n$, where n is the number of iterations.

The Gauss-Seidel method applied to solve the linearized constrained equations of motion reproduces exactly the PBD algorithm.


5 Stiff Constraint Examples

5.1 Point to Point Constraint

Define a constraint between two bodies i and j such that the distance between a point \mathbf{p}_i on body i and a point \mathbf{p}_j on body j is zero.

Let \mathbf{r}_0 be the offset of the point \mathbf{p} from the centre of mass \mathbf{x} of body so that

$$\mathbf{p}(\mathbf{X}) = \mathbf{x} + R(\boldsymbol{\theta})\mathbf{r}_0 = \mathbf{x} + \mathbf{r}(\boldsymbol{\theta}), \quad (5.1)$$



The constraint function sets the distance between the two points to be zero

$$c(\mathbf{X}) = \|\mathbf{p}_i - \mathbf{p}_j\| = \|\delta\mathbf{p}\|. \quad (5.2)$$

We can find the constraint derivative by using $\dot{c} = \mathcal{D}c \mathcal{V}$ with

$$\mathcal{V} = \begin{pmatrix} \mathbf{x}_i \\ \boldsymbol{\omega}_i \\ \mathbf{x}_j \\ \boldsymbol{\omega}_j \end{pmatrix} \quad (5.3)$$

So computing

$$\begin{aligned} \dot{c}(\mathbf{X}) &= (\dot{\mathbf{p}}_i - \dot{\mathbf{p}}_j) \cdot \hat{\delta\mathbf{p}}, \\ &= (\dot{\mathbf{x}}_i + \boldsymbol{\omega}_i \times \mathbf{r}_i - \dot{\mathbf{x}}_j - \boldsymbol{\omega}_j \times \mathbf{r}_j) \cdot \hat{\delta\mathbf{p}} \end{aligned} \quad (5.4)$$

where $\hat{\delta\mathbf{p}}$ is the unit vector in the direction $\delta\mathbf{p}$

$$\hat{\delta\mathbf{p}} = \frac{\delta\mathbf{p}}{\|\delta\mathbf{p}\|} \quad (5.5)$$

gives

$$\mathcal{D}c = \begin{pmatrix} \hat{\delta\mathbf{p}}^T & (\mathbf{r}_i \times \hat{\delta\mathbf{p}})^T & -\hat{\delta\mathbf{p}}^T & -(\mathbf{r}_j \times \hat{\delta\mathbf{p}})^T \end{pmatrix}. \quad (5.6)$$

Define the effective mass μ for this constraint by

$$\begin{aligned} \mu^{-1} &= \mathcal{D}c M^{-1} (\mathcal{D}c)^T, \\ &= \begin{pmatrix} \hat{\delta\mathbf{p}}^T & (\mathbf{r}_i \times \hat{\delta\mathbf{p}})^T & -\hat{\delta\mathbf{p}}^T & -(\mathbf{r}_j \times \hat{\delta\mathbf{p}})^T \end{pmatrix} \\ &\quad \begin{pmatrix} m_i \mathbb{1} & & & \\ & I_i & & \\ & & m_j \mathbb{1} & \\ & & & I_j \end{pmatrix}^{-1} \begin{pmatrix} \hat{\delta\mathbf{p}} \\ \mathbf{r}_i \times \hat{\delta\mathbf{p}} \\ -\hat{\delta\mathbf{p}} \\ -\mathbf{r}_j \times \hat{\delta\mathbf{p}} \end{pmatrix}, \\ &= \left(\frac{1}{m_i} + \frac{1}{m_j} \right) \mathbb{1} + (\mathbf{r}_i \times \hat{\delta\mathbf{p}}) \cdot I_i (\mathbf{r}_i \times \hat{\delta\mathbf{p}}) + (\mathbf{r}_j \times \hat{\delta\mathbf{p}}) \cdot I_j (\mathbf{r}_j \times \hat{\delta\mathbf{p}}), \\ &= \hat{\delta\mathbf{p}} \cdot \left(\left(\frac{1}{m_i} + \frac{1}{m_j} \right) \mathbb{1} - [\mathbf{r}_i]_{\times} I_i^{-1} [\mathbf{r}_i]_{\times} - [\mathbf{r}_j]_{\times} I_j^{-1} [\mathbf{r}_j]_{\times} \right) \hat{\delta\mathbf{p}}. \end{aligned} \quad (5.8)$$

This is the projection of the inverse of the effective mass matrix \mathcal{M}^{-1} for the two body system along the direction $\hat{\delta\mathbf{p}}$.

Define the inverse of the effective mass as \mathcal{M}^{-1} and its projection in the constraint direction as μ^{-1}

$$\mathcal{M}^{-1} = \left(\frac{1}{m_i} + \frac{1}{m_j} \right) \mathbb{1} - [\mathbf{r}_i]_{\times} I_i^{-1} [\mathbf{r}_i]_{\times} - [\mathbf{r}_j]_{\times} I_j^{-1} [\mathbf{r}_j]_{\times}, \quad (5.9)$$

$$\mu^{-1} = \hat{\delta\mathbf{p}} \cdot \mathcal{M}^{-1} \hat{\delta\mathbf{p}}. \quad (5.10)$$

The effective mass measures the resistance to movement of the two body system

Using this, the update is

$$\lambda^{k+1} = \lambda^k - \xi\mu \|\delta\mathbf{p}(\mathbf{X}^{t+1} + \Delta\mathbf{X})\|, \quad (5.11)$$

with

$$\begin{aligned} \Delta\mathbf{X}' &= \Delta\mathbf{X} + M^{-1} (\mathcal{D}\mathbf{c})^T (\lambda^{k+1} - \lambda^k), \\ &= \begin{pmatrix} \frac{1}{m_i} \mathbb{1} & & & \\ & I_i^{-1} & & \\ & & \frac{1}{m_j} \mathbb{1} & \\ & & & I_j^{-1} \end{pmatrix} \begin{pmatrix} \hat{\delta\mathbf{p}} \\ \mathbf{r}_i \times \hat{\delta\mathbf{p}} \\ -\hat{\delta\mathbf{p}} \\ -\mathbf{r}_j \times \hat{\delta\mathbf{p}} \end{pmatrix} (\lambda^{k+1} - \lambda^k), \end{aligned} \quad (5.12)$$

Overall:

Constraint update for a point to point constraint:

Delta:

$$\Delta = -\xi\mu \delta\mathbf{p}(\mathbf{X}^{t+1} + \Delta\mathbf{X})$$

Lagrange multiplier update:

$$\lambda^{k+1} = \lambda^k - \Delta$$

Configuration update:

$$\begin{aligned} \Delta\mathbf{x}'_i &= \Delta\mathbf{x}_i + \frac{1}{m_i} \Delta, \\ \Delta\boldsymbol{\omega}'_i &= \Delta\boldsymbol{\omega}_i + I_i^{-1} \mathbf{r}_i \times \Delta, \\ \Delta\mathbf{x}'_j &= \Delta\mathbf{x}_j - \frac{1}{m_j} \Delta, \\ \Delta\boldsymbol{\omega}'_j &= \Delta\boldsymbol{\omega}_j - I_j^{-1} \mathbf{r}_j \times \Delta. \end{aligned} \quad (5.13)$$

5.1.1 Comparison to an Impulse Based Approach

We can see that the constraint update is essentially an incremental impulse update

$$\begin{aligned}
(m_i \Delta \mathbf{x}_i / \Delta t)' &= (m_i \Delta \mathbf{x}_i / \Delta t) + \Delta, \\
(I_i \Delta \boldsymbol{\omega}_i / \Delta t)' &= (I_i \Delta \boldsymbol{\omega}_i / \Delta t) + \mathbf{r}_i \times \Delta, \\
(m_j \Delta \mathbf{x}_j / \Delta t)' &= (m_j \Delta \mathbf{x}_j / \Delta t) - \Delta, \\
(I_j \Delta \boldsymbol{\omega}_j / \Delta t)' &= (I_j \Delta \boldsymbol{\omega}_j / \Delta t) - \mathbf{r}_j \times \Delta.
\end{aligned} \tag{5.14}$$

The only difference being that in an incremental impulse update you would typically be computing the relative velocity between the two points on i and j , $\delta \dot{\mathbf{p}}$, rather than computing the change in position coordinates \mathbf{x} and $\boldsymbol{\theta}$ then using that to find the exact error $\delta \mathbf{p}(\mathbf{X}^{t+1} + \Delta \mathbf{X})$ and dividing by the time step. A velocity/impulse based approach results in a linearization of this error, which can cause problems when the angular velocities are high.

5.2 Symmetries and Conservation of Momentum

The solution to the point to point constraint (5.13) preserves both linear and angular momentum. The change in linear momentum of the system is

$$\begin{aligned}
\delta \mathbf{P} &= m_i \dot{\mathbf{x}}_i^{t+1} + m_j \dot{\mathbf{x}}_j^{t+1} - m_i \dot{\mathbf{x}}_i^t - m_j \dot{\mathbf{x}}_j^t, \\
&= m_i \Delta \mathbf{x}_i / \Delta t + m_j \Delta \mathbf{x}_j / \Delta t, \\
&= m_i \frac{1}{m_i} \mathbf{d} - m_j \frac{1}{m_j} \mathbf{d}, \\
&= 0.
\end{aligned} \tag{5.15}$$

and the change in angular momentum about the origin is

$$\begin{aligned}
\delta \mathbf{L} &= I_i \boldsymbol{\omega}_i^{t+1} + \mathbf{x}_i^{t+1} \times (m_i \dot{\mathbf{x}}_i^{t+1}) + I_j \boldsymbol{\omega}_j^{t+1} + \mathbf{x}_j^{t+1} \times (m_j \dot{\mathbf{x}}_j^{t+1}), \\
&= \mathbf{x}_i \times (m_i \Delta \mathbf{x}_i / \Delta t) + I_i \Delta \boldsymbol{\omega}_i / \Delta t - \mathbf{x}_j \times (m_j \Delta \mathbf{x}_j / \Delta t) + I_j \Delta \boldsymbol{\omega}_j / \Delta t, \\
&= (\mathbf{x}_i + \mathbf{r}_i - \mathbf{x}_j - \mathbf{r}_j) \times \mathbf{d}, \\
&= \delta \mathbf{p} \times \mathbf{d}, \\
&= 0.
\end{aligned} \tag{5.16}$$

So, in order for the angular momentum to be preserved, \mathbf{d} needs to be parallel to the vector between the two constrained points $\delta \mathbf{p}$, or the two constrained points need to be coincident, i.e. $\delta \mathbf{p} = 0$.

With our choice of constraint function $c(\mathbf{X}) = \|\delta \mathbf{p}(\mathbf{X})\|$, \mathbf{d} is proportional to $\delta \mathbf{p}$ and so angular velocity is preserved. However, we could have chosen a different constraint function. For example, some engines [10] use a vector constraint $c(\mathbf{X}) = \delta \mathbf{p}(\mathbf{X})$. This also works, but there is a difference in the per iteration update. Instead of getting a scalar resistance μ_i the resistance is the full effective mass matrix \mathcal{M} . This means that \mathbf{d} is now proportional to $\mathcal{M} \delta \mathbf{p}$, but, as \mathcal{M} is a full 3×3 matrix with off diagonal elements, this is no longer necessarily parallel to $\delta \mathbf{p}$. Therefore, if the per iteration update does not leave $\delta \mathbf{p} = 0$ (say, there is some relaxation) then angular momentum will not be conserved, and is only conserved when the system converges.

The form of constraint required to ensure that the update conserves momentum can be understood using Noether's theorem, which relates symmetries of the action to conserved quantities [12].

Consider shifting all position coordinates of by a small constant vector $\boldsymbol{\epsilon}$: $\mathbf{x}_i \rightarrow \mathbf{x}_i + \boldsymbol{\epsilon}$. Then

$$\begin{aligned}
S &\rightarrow \int dt \left[L(\dot{\mathbf{X}}, \mathbf{X}) + \sum_i \frac{\partial L}{\partial \mathbf{x}_i} \cdot \boldsymbol{\epsilon} + \boldsymbol{\lambda} \cdot \left(c(\mathbf{X}) + \sum_i J_i \boldsymbol{\epsilon} \right) \right], \\
&= S + \int dt \sum_i \left[-\frac{\partial V}{\partial \mathbf{x}_i} + J_i^T \boldsymbol{\lambda} \right] \cdot \boldsymbol{\epsilon}, \\
&= S + \int dt \sum_i \left[\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{\mathbf{x}}_i} \right) \right] \cdot \boldsymbol{\epsilon}, \\
&= S + \int dt \frac{d}{dt} \left[\sum_i (m_i \dot{\mathbf{x}}_i) \right] \cdot \boldsymbol{\epsilon}. \tag{5.17}
\end{aligned}$$

So, if the action is symmetric under $\mathbf{x} \rightarrow \mathbf{x} + \boldsymbol{\epsilon}$, then that implies that $\sum_i (m_i \dot{\mathbf{x}}_i)$ is a constant, i.e. the total linear momentum of the system is conserved. This means we need constraints to be symmetric under translations in order to preserve linear momentum. E.g. a two body constraint should be a function of $(\mathbf{x}_i - \mathbf{x}_j)$.

The conservation of angular momentum follows in the same way from rotational symmetry of the action. Consider an infinitesimal rotation of the system $R(\boldsymbol{\epsilon})$, so that

$$\begin{aligned}
\boldsymbol{\theta}_i &\rightarrow \boldsymbol{\theta}_i + \boldsymbol{\epsilon}, \\
\mathbf{x}_i &\rightarrow \mathbf{x}_i + \boldsymbol{\epsilon} \times \mathbf{x}_i. \tag{5.18}
\end{aligned}$$

To see the quantity that needs to be conserved in order to preserve this symmetry use

$$\begin{aligned}
\delta \left(L(\dot{\mathbf{X}}, \mathbf{X}) + \boldsymbol{\lambda} \cdot c(\mathbf{X}) \right) &= \sum_i \left[\frac{\partial L}{\partial \dot{\mathbf{X}}_i} \cdot \delta \dot{\mathbf{X}}_i + \left(J_i^T \boldsymbol{\lambda}_i + \frac{\partial L}{\partial \mathbf{X}_i} \right) \cdot \delta \mathbf{X}_i \right], \\
&= \sum_i \left[\mathbf{P}_i \cdot \delta \dot{\mathbf{X}}_i + \dot{\mathbf{P}}_i \cdot \delta \mathbf{X}_i \right], \tag{5.19}
\end{aligned}$$

where \mathbf{P}_i is the momentum $\frac{\partial L}{\partial \dot{\mathbf{X}}_i} = M_i \dot{\mathbf{X}}_i$. So

$$\frac{d}{dt} \sum_i \left(\dot{\mathbf{P}}_i \cdot \delta \mathbf{X}_i \right) = 0. \tag{5.20}$$

For the case of an infinitesimal rotation, this gives

$$\frac{d}{dt} \sum_i \left[\mathbf{x}_i \times (m_i \dot{\mathbf{x}}_i) + I_i \dot{\boldsymbol{\theta}}_i \right] \cdot \boldsymbol{\epsilon} = 0. \tag{5.21}$$

Therefore, the total angular momentum of the system is conserved if the action is symmetric under an infinitesimal rotation. Since we can compose a finite rotation from a product of infinitesimal rotations, angular momentum is conserved if the system is invariant under arbitrary rotations.

For example, a scalar like $\|\delta \mathbf{p}\|$ is invariant under rotations, since, if we apply a rotation R to the system

$$\begin{aligned}
\|\delta \mathbf{p}\| &= [\delta \mathbf{p} \cdot \delta \mathbf{p}]^{1/2}, \\
&\rightarrow [(R\delta \mathbf{p}) \cdot (R\delta \mathbf{p})]^{1/2}, \\
&= [\delta \mathbf{p} \cdot (R^T R \delta \mathbf{p})]^{1/2}, \\
&= \|\delta \mathbf{p}\| \tag{5.22}
\end{aligned}$$

since $R^T R = \mathbb{1}$.

But also, a constraint $\mathbf{c} = \mathbf{p}_i - \mathbf{p}_j$ is fine

$$\boldsymbol{\lambda} \cdot (\mathbf{p}_i - \mathbf{p}_j) \rightarrow (R\boldsymbol{\lambda}) \cdot R(\mathbf{p}_i - \mathbf{p}_j) = \boldsymbol{\lambda} \cdot (\mathbf{p}_i - \mathbf{p}_j), \quad (5.23)$$

since the Lagrange multipliers are also transformed under rotations of the coordinate system.

A constraint like

$$\mathbf{c} = \mathbf{p}_i - \mathbf{p}_j - \mathbf{L} \quad (5.24)$$

that might try to enforce that the two points maintain a distance L apart, however, would not be rotationally invariant and so would not conserve angular momentum.

So, in order to preserve linear momentum, constraints must be invariant under translations and, in order to preserve angular momentum, the constraint term must be invariant under rotation.

It can also be shown that energy conservation is a result of time translation invariance, which will be true unless the constraint has a direct time dependence, e.g. a damping term.

Conserved quantities come from symmetries of the Lagrangian. These symmetries must be respected by the constraint term $\boldsymbol{\lambda} \cdot \mathbf{c}(\mathbf{X})$

1. Conservation of linear momentum comes from translational symmetry $\mathbf{X} \rightarrow \mathbf{X} + \mathbf{D}$ of the coordinates
2. Conservation of angular momentum comes from rotational symmetry $\mathbf{X} \rightarrow R\mathbf{X}$ of the coordinates

5.3 Inequality Constraints

A standard extension to Gauss-Seidel is Projected Gauss-Seidel where, as well as solving inequality constraints $\mathbf{c}(\mathbf{X}) = 0$, we can extend the system to also solve inequality constraints $\mathbf{c}(\mathbf{X}) \geq 0$. This extension comes from treating the system as a Linear Complementarity Problem (LCP). Starting from the linearized constraint

$$\mathbf{c}(\mathbf{X}^{t+1} + \Delta\mathbf{X}) \approx \mathbf{c}(\mathbf{X}^{t+1}) + \mathcal{D}\mathbf{c}\Delta\mathbf{X}, \quad (5.25)$$

the inequality $\mathbf{c}(\mathbf{X}^{t+1} + \Delta\mathbf{X}) \geq 0$ leads to the convex quadratic program

$$\min_{\Delta\mathbf{X}} \frac{1}{2} \Delta\mathbf{X} \cdot M \Delta\mathbf{X} \quad \text{s.t.} \quad \mathbf{c}(\mathbf{X}^{t+1}) + \mathcal{D}\mathbf{c}\Delta\mathbf{X} \geq 0. \quad (5.26)$$

Eliminating $\Delta\mathbf{X}$ gives the LCP:

$$0 \leq \boldsymbol{\lambda} \perp \mathbf{c}(\mathbf{X}^{t+1}) + \left(\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right) \boldsymbol{\lambda} \geq 0, \quad (5.27)$$

which is solved using a Projected Gauss-Seidel update of

$$\begin{aligned} \lambda_i^{k+1} &= \lambda_i^k - \left(\mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right)_{ii} \mathbf{c}(\mathbf{X}^{t+1} + \Delta\mathbf{X}^k)_i, \\ \lambda_i^{k+1} &= \begin{cases} \lambda_i^* & \text{if } \lambda_i^* \geq 0 \\ 0 & \text{otherwise} \end{cases} \end{aligned} \quad (5.28)$$

See, for example, [20] where they derive the projected Gauss-Seidel update as a solution to the LCP. Alternatively, [8] gives an impulse centric approach, where it is obvious that clamping the impulse to to be one way in the constraint direction gives the correct solution.

5.4 Contact Constraint

A contact constraint is an example of an inequality constraint. The constraint is that for two points \mathbf{p}_i and \mathbf{p}_j on bodies i and j the distance between the points in the direction of a contact normal direction $\hat{\mathbf{n}}$ is positive.

$$c(\mathbf{X}) = (\mathbf{p}(\mathbf{X}_i) - \mathbf{p}(\mathbf{X}_j)) \cdot \hat{\mathbf{n}} \geq 0. \quad (5.29)$$

where, as in 5.1

$$\begin{aligned} \mathbf{p}(\mathbf{X}) &= \mathbf{x} + R(\boldsymbol{\theta})\mathbf{r}_0, \\ &= \mathbf{x} + \mathbf{r}(\boldsymbol{\theta}) \end{aligned} \quad (5.30)$$

for an offset \mathbf{r} from a centre of mass \mathbf{x} .

The constraint derivative is

$$\mathcal{D}\mathbf{c} = (\hat{\mathbf{n}}^T \quad (\mathbf{r}_i \times \hat{\mathbf{n}})^T \quad -\hat{\mathbf{n}}^T \quad -(\mathbf{r}_j \times \hat{\mathbf{n}})^T) \quad (5.31)$$

leading to an inverse effective mass of

$$\begin{aligned} \mu^{-1} &= \mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T, \\ &= \hat{\mathbf{n}} \cdot \mathcal{M}^{-1} \hat{\mathbf{n}} \end{aligned} \quad (5.32)$$

where \mathcal{M}^{-1} is the inverse of the effective mass matrix (5.10). I.e., we are projecting the inverse of the effective mass matrix in the direction $\hat{\mathbf{n}}$ to measure the resistance to movement in that direction.

With this, the update becomes

Contact normal constraint update

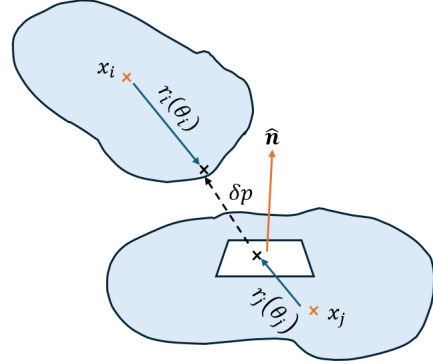
Lagrange multiplier update with clamping:

$$\begin{aligned} \lambda_i^* &= \lambda_i^k - \xi \mu \delta\mathbf{p}(\mathbf{X}^{t+1} + \Delta\mathbf{X}^k) \cdot \hat{\mathbf{n}}, \\ \lambda_i^{k+1} &= \begin{cases} \lambda_i^* & \text{if } \lambda_i^* \geq 0 \\ 0 & \text{otherwise} \end{cases} \end{aligned} \quad (5.33)$$

Delta:

$$\Delta = (\lambda_i^{k+1} - \lambda_i^k) \hat{\mathbf{n}}. \quad (5.34)$$

Configuration update:



$$\begin{aligned}
\Delta \mathbf{x}'_i &= \Delta \mathbf{x}_i + \frac{1}{m_i} \Delta, \\
\Delta \boldsymbol{\omega}'_i &= \Delta \boldsymbol{\omega}_i + I_i^{-1} \mathbf{r}_i \times \Delta, \\
\Delta \mathbf{x}'_j &= \Delta \mathbf{x}_j - \frac{1}{m_j} \Delta, \\
\Delta \boldsymbol{\omega}'_j &= \Delta \boldsymbol{\omega}_j - I_j^{-1} \mathbf{r}_j \times \Delta.
\end{aligned} \tag{5.35}$$

5.5 Contact Velocity Constraint

For collisions with restitution e the final relative velocity along the contact normal should be equal to $-e$ times the incoming relative velocity.

$$\delta \dot{\mathbf{p}}(\dot{\mathbf{X}}^{t+1}) \cdot \hat{\mathbf{n}} = -e \delta \dot{\mathbf{p}}(\dot{\mathbf{X}}^t) \cdot \hat{\mathbf{n}} \tag{5.36}$$

with

$$\dot{\mathbf{p}} = \dot{\mathbf{x}} + \boldsymbol{\omega} \times \mathbf{r}. \tag{5.37}$$

We can make this into a velocity constraint

$$c(\dot{\mathbf{X}}) = -(1+e) \delta \dot{\mathbf{p}} \cdot \hat{\mathbf{n}}. \tag{5.38}$$

This is an example of a non-holonomic constraint, covered in section 3.2. To see that this constraint is the correct one, just consider linear velocities along \mathbf{n} for now. The update would be

$$\begin{aligned}
\dot{\mathbf{x}}_i &\rightarrow \dot{\mathbf{x}}_i - \frac{1}{m_i} \mu (1+e) (\dot{\mathbf{x}}_i - \dot{\mathbf{x}}_j), \\
\dot{\mathbf{x}}_j &\rightarrow \dot{\mathbf{x}}_j + \frac{1}{m_j} \mu (1+e) (\dot{\mathbf{x}}_i - \dot{\mathbf{x}}_j), \\
\Rightarrow (\dot{\mathbf{x}}_i - \dot{\mathbf{x}}_j) &\rightarrow (\dot{\mathbf{x}}_i - \dot{\mathbf{x}}_j) + \mu^{-1} \mu (1+e) (\dot{\mathbf{x}}_i - \dot{\mathbf{x}}_j), \\
&= -e (\dot{\mathbf{x}}_i - \dot{\mathbf{x}}_j).
\end{aligned} \tag{5.39}$$

This is an example of a non-holonomic constraint covered in section 3.2. The constraint derivative is

$$\dot{\mathcal{D}}c = (\hat{\mathbf{n}}^T \quad (\mathbf{r}_i \times \hat{\mathbf{n}})^T \quad -\hat{\mathbf{n}}^T \quad -(\mathbf{r}_j \times \hat{\mathbf{n}})^T), \tag{5.40}$$

which is identical to $\mathcal{D}c$ we derived from the contact normal constraint. So, the effective mass is also the same

$$\mu^{-1} = \hat{\mathbf{n}} \cdot \mathcal{M}^{-1} \hat{\mathbf{n}}. \tag{5.41}$$

With this, the update is

Contact normal velocity constraint with restitution
Lagrange multiplier update with clamping:

$$\begin{aligned}
\lambda_i^* &= \lambda_i^k + \xi \mu (1+e) \delta \dot{\mathbf{p}}(\dot{\mathbf{X}}^{t+1} + \Delta \dot{\mathbf{X}}^k / \Delta t) \cdot \hat{\mathbf{n}}, \\
\lambda_i^{k+1} &= \begin{cases} \lambda_i^* & \text{if } \lambda_i^* \geq 0 \\ 0 & \text{otherwise} \end{cases}
\end{aligned} \tag{5.42}$$

Delta:

$$\Delta = (\lambda_i^{k+1} - \lambda_i^k) \hat{\mathbf{n}}. \quad (5.43)$$

Configuration update:

$$\begin{aligned} \Delta \mathbf{x}'_i &= \Delta \mathbf{x}_i + \frac{1}{m_i} \Delta, \\ \Delta \boldsymbol{\omega}'_i &= \Delta \boldsymbol{\omega}_i + I_i^{-1} \mathbf{r}_i \times \Delta, \\ \Delta \mathbf{x}'_j &= \Delta \mathbf{x}_j - \frac{1}{m_j} \Delta, \\ \Delta \boldsymbol{\omega}'_j &= \Delta \boldsymbol{\omega}_j - I_j^{-1} \mathbf{r}_j \times \Delta. \end{aligned} \quad (5.44)$$

This is almost the same as (5.35). The difference is the displacement term

$$\begin{aligned} \text{Position :} \quad \Delta &= -\xi \mu \delta \mathbf{p}(\mathbf{X}^{t+1} + \Delta \mathbf{X}^k) \cdot \hat{\mathbf{n}} \hat{\mathbf{n}}, \\ \text{Velocity :} \quad \Delta &= \xi \mu (1 + e) \delta \dot{\mathbf{p}}(\mathbf{X}^{t+1} + \Delta \mathbf{X}^k / \Delta t) \cdot \hat{\mathbf{n}} \hat{\mathbf{n}} \Delta t \end{aligned} \quad (5.45)$$

and

$$\delta \mathbf{p}(\mathbf{X}^{t+1} + \Delta \mathbf{X}^k) = \delta \mathbf{p}(\mathbf{X}^t) + \delta \dot{\mathbf{p}}(\mathbf{X}^{t+1} + \Delta \mathbf{X}^k / \Delta t) \Delta t. \quad (5.46)$$

So, to first order, the velocity constraint can be formulated as a position constraint with just a different target relative position.

5.6 Contact Friction Constraint

Friction acts to oppose relative tangential velocity at the contact point. The dry friction model divides the force into two modes: static friction for when two surfaces are not moving relative to each other, and dynamic friction when they are in relative motion in the contact plane. The model is that if the magnitude of the force is less than the static friction limit $\nu_s \|\mathbf{N}\|$ for some contact normal force \mathbf{N} and static friction coefficient ν_s then the force prevents the points moving relative to each other in the contact plane. If, however, the magnitude of the force required to keep the points from moving relative to each other in the contact plane exceeds that limit then the bodies will start moving and a dynamic friction force of $\nu_d \|\mathbf{N}\| \hat{\mathbf{s}}$, where $\hat{\mathbf{s}}$ is the direction of the relative surface velocity and ν_d is a dynamic friction coefficient. The model is valid for $\nu_d \leq \nu_s$.

Let the relative velocity at the contact points be

$$\delta \dot{\mathbf{p}} = \dot{\mathbf{p}}_i - \dot{\mathbf{p}}_j = (\dot{\mathbf{x}}_i + \boldsymbol{\omega}_i \times \mathbf{r}_i) - (\dot{\mathbf{x}}_j + \boldsymbol{\omega}_j \times \mathbf{r}_j). \quad (5.47)$$

We decompose this into normal and tangential components:

$$\delta \dot{\mathbf{p}}_{\perp} = \delta \dot{\mathbf{p}} - (\delta \dot{\mathbf{p}} \cdot \hat{\mathbf{n}}) \hat{\mathbf{n}}. \quad (5.48)$$

Where

$$\hat{\mathbf{t}} = \frac{\delta \dot{\mathbf{p}}_{\perp}}{\|\delta \dot{\mathbf{p}}_{\perp}\|}. \quad (5.49)$$

We impose the constraint

$$c(\dot{\mathbf{X}}) = \delta \dot{\mathbf{p}} \cdot \hat{\mathbf{t}}. \quad (5.50)$$

The derivative with respect to velocities is

$$\dot{\mathcal{D}}\mathbf{c} = (\hat{\mathbf{t}}^T \quad (\mathbf{r}_i \times \hat{\mathbf{t}})^T \quad -\hat{\mathbf{t}}^T \quad -(\mathbf{r}_j \times \hat{\mathbf{t}})^T). \quad (5.51)$$

Giving an effective mass of

$$\mu_t^{-1} = \hat{\mathbf{t}} \cdot \mathcal{M}^{-1} \hat{\mathbf{t}}, \quad (5.52)$$

where \mathcal{M}^{-1} is the inverse effective mass matrix used in the normal contact constraint. The constraint update is


Contact friction constraint update

$$\begin{aligned} \boldsymbol{\lambda}^* &= \boldsymbol{\lambda}^k - \xi \mu_t \delta \dot{\mathbf{p}}(\mathbf{X}^{t+1} + \Delta \mathbf{X}^k / \Delta t) \cdot \hat{\mathbf{t}} \hat{\mathbf{t}}, \\ \boldsymbol{\lambda}^{k+1} &= \begin{cases} \boldsymbol{\lambda}^* & \text{if } \|\boldsymbol{\lambda}^*\| \leq \nu_s \|\mathbf{N}\| \\ \nu_d \|\mathbf{N}\| \hat{\boldsymbol{\lambda}}^* & \text{otherwise} \end{cases} \end{aligned} \quad (5.53)$$

with the incremental position and orientation updates as in the update for the contact normal constraint (5.35).

5.7 Angular Axis Constraint

For this constraint example we will look at a purely angular constraint. The goal is to limit the rotation of one body to be around an axis defined in the space of another body. Imagine a wheel constrained to rotate around an axis in the space of a car body, for instance.



There are a few different ways of defining this constraint, but we choose

$$c(\boldsymbol{\theta}) = \sin^{-1} \|\hat{\mathbf{n}}_i \times \hat{\mathbf{n}}_j\|, \quad (5.54)$$

where $\hat{\mathbf{n}}$ are axes on the respective bodies that we want to keep aligned.

The reason for choosing the sine of the angle being zero as the constraint, rather than, say looking at the cosine with $c(\boldsymbol{\theta}) = \hat{\mathbf{n}}_i \cdot \hat{\mathbf{n}}_j - 1$ is that around the point where the angle between the two is small sine is changing rapidly, whereas cosine is fairly flat, so the derivative around this point is more reliable with the sine formulation.

Taking the derivative with respect to $\boldsymbol{\theta}_i$ and $\boldsymbol{\theta}_j$ gives

$$\mathcal{D}\mathbf{c} = (\hat{\mathbf{w}}^T \quad -\hat{\mathbf{w}}^T) \quad (5.55)$$

where

$$\hat{\mathbf{w}} = \frac{\hat{\mathbf{n}}_i \times \hat{\mathbf{n}}_j}{\|\hat{\mathbf{n}}_i \times \hat{\mathbf{n}}_j\|} \quad (5.56)$$

is the unit rotation axis required to align the axes.

Using this, the inverse effective inertia is

$$\mu^{-1} = \left[\hat{\mathbf{w}} \cdot (I_i^{-1} + I_j^{-1}) \hat{\mathbf{w}} \right]^{-1}, \quad (5.57)$$

i.e. the projection of the effective inertia in the direction $\hat{\mathbf{w}}$.

With this, the update is

Axis alignment constraint

Lagrange multiplier update:

$$\lambda^{k+1} = \lambda_i^k - \mu \left(\sin \|\hat{\mathbf{n}}_i(\boldsymbol{\theta}_i^{t+1} + \Delta\boldsymbol{\theta}_i) \times \hat{\mathbf{n}}_j(\boldsymbol{\theta}_j^{t+1} + \Delta\boldsymbol{\theta}_j)\| \right) \quad (5.58)$$

Delta:

$$\Delta = (\lambda^{k+1} - \lambda^k) \hat{\mathbf{w}}. \quad (5.59)$$

Configuration update:

$$\begin{aligned} \Delta\boldsymbol{\omega}'_i &= \Delta\boldsymbol{\omega}_i + I_i^{-1}\Delta, \\ \Delta\boldsymbol{\omega}'_j &= \Delta\boldsymbol{\omega}_j - I_j^{-1}\Delta. \end{aligned} \quad (5.60)$$

6 Compliant Constraints

6.1 Generalized Implicit Non-Linear Forces

So far, we have assumed that we can, by some other method / integration scheme, find the solutions to the equations of motion when there are no constraints. But if we have non-linear forces and torques and we want to solve for them implicitly (i.e. taking the force term at $F(\mathbf{X}^{t+1})$ when solving for \mathbf{X}^{t+1} is not necessarily easy. But in general updating with implicit forces and torques is a good thing for stability, so we want to find a solution. We proceed using a similar first order expansion that we made when solving the constrained equation, except that this time we assume that we have a solution $\mathbf{X}^{t+1}, \dot{\mathbf{X}}^{t+1}$ that satisfies the equations of motion in the absence of forces and torques and then add a small displacement $\Delta\mathbf{X}$ to the solution and solve the full equations of motion up to first order in $\Delta\mathbf{X}$.

The equations of motion are

$$\frac{d}{dt}(M(\mathbf{X})\mathcal{V}) = \mathbf{F}(\mathbf{X}), \quad (6.1)$$

Taking $\mathbf{X}^{t+1} \rightarrow \mathbf{X}^{t+1} + \Delta\mathbf{X}$ we get

$$M(\mathbf{X}^{t+1} + \Delta\mathbf{X})(\dot{\mathbf{X}}^{t+1}\Delta t + \Delta\dot{\mathbf{X}}) - M(\mathbf{X}^{t+1})\dot{\mathbf{X}}^{t+1}\Delta t = \mathbf{F}(\mathbf{X}^{t+1})\Delta t^2 + (\mathcal{D}\mathbf{F})^T \Delta\mathbf{X}\Delta t^2, \quad (6.2)$$

which gives a linear equation for $\Delta\mathbf{X}$

Linearized system for an increment $\Delta\mathbf{X}$ required to add to the solution \mathbf{X}^{t+1} of the free equations of motion that solves to first order the equations of motion with generalized forces:

$$(M + K)\Delta\mathbf{X} = \mathbf{F}(\mathbf{X}^{t+1})\Delta t^2 \quad (6.3)$$

with

$$K = -(\mathcal{D}\mathbf{F})^T \Delta t^2 \quad (6.4)$$

Since K is not necessarily block diagonal, we need to use an iterative approach to solve this system for $\Delta\mathbf{X}$. We again use the Gauss-Seidel algorithm. Following a similar method to section 4.5 results in an iterative update

Gauss-Seidel iteration for solving the force equation (6.3)

$$\Delta\mathbf{X}_i^{k+1} = \Delta\mathbf{X}_i^k + \xi [(M + K)_{ii}]^{-1} \left[\mathbf{F}_i(\mathbf{X}^{t+1} + \Delta\mathbf{X}^k)\Delta t^2 - M_i\Delta\mathbf{X}_i^k \right] \quad (6.5)$$

This is of the same form as the iterative update for the constraint forces, and so can be added as an additional term in the constraint solve update in the same iteration loop. The result is that, as we update the displacement $\Delta\mathbf{X}$ each iteration, we correct \mathbf{X} with a change that corresponds to the force to be applied.

6.2 Connection to Vertex Block Descent

Notice in this approach that, instead of solving iteratively per constraint, we are solving per body. \mathbf{F}_i is the sum over all of the forces on body i . This is the same methodology as applied in Vertex Block Descent (VBD) methods [9].

In the vertex block descent paper they define a local variational energy equivalent to

$$G(\mathbf{X}) = \frac{m}{2} \|\Delta \mathbf{x}\|^2 + E(\mathbf{X}) \quad (6.6)$$

which they minimize, leading to a linear system of

$$H\Delta \mathbf{x} = \mathbf{f}, \quad (6.7)$$

with

$$\begin{aligned} \mathbf{f} &= -m\Delta \mathbf{x} - \frac{\partial E}{\partial \mathbf{x}}, \\ H &= m\mathbb{1} + \frac{\partial^2 E}{\partial \mathbf{x}^2} \end{aligned} \quad (6.8)$$

We saw in 4.4 that this kind of minimization is equivalent to solving the constrained equations of motion as a first order perturbation on the unconstrained equations. Now we are finding a first order perturbation that solves the full equations of motion as a perturbation on the free equations of motion. So, our generalized force method is the extension of VDB to rigid bodies, identifying

$$\begin{aligned} \frac{\partial E}{\partial \mathbf{x}} &\sim \mathbf{F}, \\ \frac{\partial^2 E}{\partial \mathbf{x}^2} &\sim \mathcal{D}\mathbf{F}, \\ M(\mathbf{X}) &\sim m \end{aligned} \quad (6.9)$$

The method presented for solving for general implicit non-linear forces is a rigid body version of the method of Vertex Block Descent.

6.3 Connection to Lagrange Multiplier Method and PBD

We can reformulate the system to be solved by introducing an intermediate variable $\boldsymbol{\lambda}$ as

$$\Delta \mathbf{X} = M^{-1}(M + K)^T \boldsymbol{\lambda}. \quad (6.10)$$

Using this, the equation (6.3) becomes

$$((M + K)M^{-1}(M + K)^T) \boldsymbol{\lambda} = \mathbf{F}(\mathbf{X}^{t+1})\Delta t^2. \quad (6.11)$$

From this we get a Gauss-Seidel iterative update of

$$\boldsymbol{\lambda}_i^{k+1} = \boldsymbol{\lambda}_i^k + \xi \left((M + K)M^{-1}(M + K)^T \right)_{ii}^{-1} \left(\mathbf{F}(\mathbf{X}^{t+1})\Delta t^2 - (M + K)M^{-1}(M + K)^T \boldsymbol{\lambda}^k \right)_i \quad (6.12)$$

with

$$\Delta \mathbf{X}_i^{k+1} = \Delta \mathbf{X}_i^k + M_{ii}^{-1}(M + K)_i^T (\boldsymbol{\lambda}_i^{k+1} - \boldsymbol{\lambda}_i^k). \quad (6.13)$$

Rearranging the λ^{k+1} update we get

$$\lambda_i^{k+1} = \lambda_i^k + \xi \left((M + K)M^{-1}(M + K)^T \right)_{ii}^{-1} \left(\mathbf{F}(\mathbf{X}^{t+1} + \Delta\mathbf{X}^k)\Delta t^2 - M\Delta\mathbf{X}^k \right)_i \quad (6.14)$$

Compare this to (4.73). We see that they correspond if we interpret

$$\begin{aligned} \mathbf{c}(\mathbf{X} + \Delta\mathbf{X}) &= M\Delta\mathbf{X} - \mathbf{F}(\mathbf{X}^{t+1} + \Delta\mathbf{X})\Delta t^2, \\ \mathcal{D}\mathbf{c} &= M - \mathcal{D}\mathbf{F} \Delta t^2 \\ &= M + K. \end{aligned} \quad (6.15)$$

We can transform the per body solve from the generalized implicit force system (6.3) into a per constraint solve by identifying the first order equations for the forces as constraints in a Lagrange multiplier system

$$\mathbf{c}(\mathbf{X}) = M\Delta\mathbf{X} - \mathbf{F}(\mathbf{X}^{t+1} + \Delta\mathbf{X})\Delta t^2 = 0. \quad (6.16)$$

So, enforcing a constraint that the integrated force evaluated at the current configuration must produce a displacement of $M\Delta\mathbf{X}$.

With this, we are transforming the problem back to iterating over “constraints” with multipliers λ rather than iterating over bodies, where each constraint enforces an individual force equation rather than summing over all of the forces. While this does affect the algorithm and the convergence of the solution, we are in the end solving the same system and so both should converge to the same answer.

6.4 Compliant Constraint Forces and XPBD

The method of Lagrange multipliers is good for modelling stiff constraints where we want the constraint equation $\mathbf{c}(\mathbf{X}) = 0$ to be satisfied exactly. Often though we want to model systems with some compliance in the constraints, as this can help with stability and also be used to model things with spring-like behaviour. For these types of systems, we add a potential term rather than using a constraint.

$$V = \mathbf{c}(\mathbf{X}) \cdot A\mathbf{c}(\mathbf{X}), \quad (6.17)$$

where $\mathbf{c}(\mathbf{X})$ are the constraint functions and A is a diagonal matrix of stiffness parameters. In addition we want to model damping. This is not naturally handled by the standard Lagrangian method, as that assumes conservative forces. The usual modification is to add a Rayleigh dissipation function [12]

$$\mathcal{F} = \dot{\mathbf{X}} \cdot D\dot{\mathbf{X}} \quad (6.18)$$

and then the equations of motion are modified to

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{\mathbf{X}}} \right) - \frac{\partial L}{\partial \mathbf{X}} - \frac{\partial \mathcal{F}}{\partial \dot{\mathbf{X}}}. \quad (6.19)$$

In our case we want to apply damping to the constraint function, so we want a term

$$\dot{\mathbf{c}}(\mathbf{X}) \cdot B \dot{\mathbf{c}}(\mathbf{X}) \quad (6.20)$$

for some diagonal matrix of damping constants B . Using $\dot{\mathbf{c}}(\mathbf{X}) = \mathcal{D}\mathbf{c}\dot{\mathbf{X}}$ we find

$$D = (\mathcal{D}\mathbf{c})^T B \mathcal{D}\mathbf{c}. \quad (6.21)$$

With these additional terms the equations of motion become

$$\frac{d}{dt} (M(\mathbf{X})\mathcal{V}) = -(\mathcal{D}\mathbf{c})^T \mathbf{A}\mathbf{c}(\mathbf{X}) - (\mathcal{D}\mathbf{c})^T B \mathcal{D}\mathbf{c}\dot{\mathbf{X}}. \quad (6.22)$$

so the force is

$$\mathbf{F} = -(\mathcal{D}\mathbf{c})^T \left(\mathbf{A}\mathbf{c}(\mathbf{X}) + B \mathcal{D}\mathbf{c}\dot{\mathbf{X}} \right). \quad (6.23)$$

The equation (6.5) tells us how to iteratively update the system given this force. First we need to find K we do this by looking at $\mathbf{F}(\mathbf{X}^{t+1} + \Delta\mathbf{X})$

$$\begin{aligned} \mathbf{F}(\mathbf{X}^{t+1} + \Delta\mathbf{X}) &= -(\mathcal{D}\mathbf{c})^T \Big|_{\mathbf{X}^{t+1} + \Delta\mathbf{X}} \left(\mathbf{A}\mathbf{c}(\mathbf{X}^{t+1} + \Delta\mathbf{X}) + B \mathcal{D}\mathbf{c}|_{\mathbf{X}^{t+1} + \Delta\mathbf{X}} (\dot{\mathbf{X}} + \Delta\mathbf{X}/\Delta t) \right), \\ &= F(\mathbf{X}^{t+1}) - (\mathcal{D}\mathbf{c})^T (A + B/\Delta t) \mathcal{D}\mathbf{c}\Delta\mathbf{X} \\ &\quad - (\mathcal{D}^2\mathbf{c})^T \left(\mathbf{A}\mathbf{c}(\mathbf{X}^{t+1}) + B \mathcal{D}\mathbf{c}\dot{\mathbf{X}} \right) - (\mathcal{D}\mathbf{c})^T B (\mathcal{D}^2\mathbf{c}\Delta\mathbf{X}) \dot{\mathbf{X}} \\ &= F(\mathbf{X}^{t+1}) - (\mathcal{D}\mathbf{c})^T (A + B/\Delta t) \mathcal{D}\mathbf{c}\Delta\mathbf{X} \\ &\quad - \left(\mathbf{A}\mathbf{c}(\mathbf{X}^{t+1}) + 2B \mathcal{D}\mathbf{c}\dot{\mathbf{X}} \right)^T \mathcal{D}^2\mathbf{c}\Delta\mathbf{X}. \end{aligned} \quad (6.24)$$

Where $\mathcal{D}^2\mathbf{c}$ is the constraint Hessian using the left trivialized derivative (3.44). Therefore, comparing to (6.3)

$$K = (\mathcal{D}\mathbf{c})^T (A\Delta t^2 + B\Delta t) \mathcal{D}\mathbf{c} + \left(A\Delta t^2 \mathbf{c}(\mathbf{X}^{t+1}) + 2B\Delta t \mathcal{D}\mathbf{c}\dot{\mathbf{X}} \Delta t \right)^T \mathcal{D}^2\mathbf{c}. \quad (6.25)$$

The constraint Hessian term is neglected in XPBD, equivalent to Gauss–Newton approximation, so we will also assume it to be small and ignore it for the purposes of comparison to XPBD. This leaves

$$K = (\mathcal{D}\mathbf{c})^T (A\Delta t^2 + B\Delta t) \mathcal{D}\mathbf{c}. \quad (6.26)$$

Using this

$$\left(M + (\mathcal{D}\mathbf{c})^T (A\Delta t^2 + B\Delta t) \mathcal{D}\mathbf{c} \right) \Delta\mathbf{X} = -(\mathcal{D}\mathbf{c})^T \left(\mathbf{A}\mathbf{c}(\mathbf{X}) + B \mathcal{D}\mathbf{c}\dot{\mathbf{X}} \right). \quad (6.27)$$

Again, we add a variable $\boldsymbol{\lambda}$ with $\Delta\mathbf{X} = M^{-1} (\mathcal{D}\mathbf{c})^T \boldsymbol{\lambda}$. This gives

$$\left(M + (\mathcal{D}\mathbf{c})^T (A\Delta t^2 + B\Delta t) \mathcal{D}\mathbf{c} \right) M^{-1} (\mathcal{D}\mathbf{c})^T \boldsymbol{\lambda} = -(\mathcal{D}\mathbf{c})^T \left(\mathbf{A}\mathbf{c}(\mathbf{X}) + B \mathcal{D}\mathbf{c}\dot{\mathbf{X}} \right). \quad (6.28)$$

Simplifying gives

$$\left(\mathbb{1} + (A\Delta t^2 + B\Delta t) \mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right) \boldsymbol{\lambda} = -\mathbf{A}\mathbf{c}(\mathbf{X}) - B \mathcal{D}\mathbf{c}\dot{\mathbf{X}}. \quad (6.29)$$

The Gauss-Seidel iterative update for $\boldsymbol{\lambda}$ is then

$$\boldsymbol{\lambda}_i^{k+1} = \boldsymbol{\lambda}_i^k - \xi \left(\mathbb{1} + (A\Delta t^2 + B\Delta t) \mathcal{D}\mathbf{c} M^{-1} (\mathcal{D}\mathbf{c})^T \right)_{ii}^{-1} \left(A\Delta t^2 \mathbf{c}(\mathbf{X}^{t+1} + \Delta\mathbf{X}^k) \right) \quad (6.30)$$

$$+ B \Delta t \mathcal{D}\mathbf{c}(\dot{\mathbf{X}}^{t+1} \Delta t + \Delta\mathbf{X}^k) + \boldsymbol{\lambda}_i^k \Big|_i. \quad (6.31)$$

To make the connection with XPBD we take

$$A_{ii}\Delta t^2 = \frac{1}{\tilde{\alpha}_i}, \quad (6.32)$$

$$B_{ii}\Delta t = \tilde{\beta}_i, \quad (6.33)$$

$$\gamma_i = \frac{\tilde{\alpha}_i\tilde{\beta}_i}{\Delta t}. \quad (6.34)$$

Using these, the update becomes

$$\boldsymbol{\lambda}_i^{k+1} = \boldsymbol{\lambda}_i^k - \xi \left(\alpha_i + (1 + \gamma_i) \mathcal{D}\mathbf{c}M^{-1}(\mathcal{D}\mathbf{c})^T \right)_{ii}^{-1} \left(\mathbf{c}(\mathbf{X}^{t+1} + \Delta\mathbf{X}^k) \right. \quad (6.35)$$

$$\left. + \gamma_i \mathcal{D}\mathbf{c}(\dot{\mathbf{X}}^{t+1}\Delta t + \Delta\mathbf{X}^k) + \tilde{\alpha}_i \boldsymbol{\lambda}_i^k \right)_i \quad (6.36)$$

which is exactly the update presented in the XPBD paper [15].

Summary:

In XPBD, instead of enforcing constraints through a Lagrange multiplier approach, constraint penalty terms are added to the potential energy

$$V = \mathbf{c}(\mathbf{X}) \cdot A\mathbf{c}(\mathbf{X}) + \dot{\mathbf{c}}(\mathbf{X}) \cdot B \dot{\mathbf{c}}(\mathbf{X}), \quad (6.37)$$

These penalty terms add a spring and damping term for each constraint function allowing for soft constraint resolution. Previously, in PBD, softness was added to a Lagrange multiplier system system using solver relaxation, which is iteration dependent. This was the big benefit of XPBD, and allowed for more consistent and stable systems.

The forces $\frac{\partial V}{\partial \mathbf{X}}$ are naturally solved using the per body method of (6.5). However, in their paper they treat the forces as constraints, as in (6.16) making the connection back to PBD and allowing for per constraint rather than per body solver iterations.

6.4.1 XPBD as a Regularization on PBD

Note that the presence of the additional $\tilde{\alpha}_i$ term in the denominator acts similarly to the regularization parameter we explored in section 4.3. For large stiffness values $\tilde{\alpha} \rightarrow 0$ and we get back to something like the Lagrange multiplier result. For simplicity, let the damping be zero and just focus on the spring like term. Let $k_i\Delta t^2 \rightarrow \infty$ so that $\tilde{\alpha}_i \rightarrow \epsilon$ for $\epsilon \rightarrow 0$. Then the XPBD update becomes

$$\boldsymbol{\lambda}_i^{k+1} = \boldsymbol{\lambda}_i^k - \xi \left(\epsilon + \mathcal{D}\mathbf{c}M^{-1}(\mathcal{D}\mathbf{c})^T \right)_{ii}^{-1} \left(\mathbf{c}(\mathbf{X}^{t+1} + \Delta\mathbf{X}) + \epsilon \boldsymbol{\lambda}_i^k \right)_i. \quad (6.38)$$

So for infinitely stiff spring coefficients k_i the update is exactly the same as the PBD update, but for non-infinite springs we get a relaxation term in the denominator, like we did when explicitly adding regularization in section 4.3, which helps with convergence of ill-posed problems (e.g. over constrained systems).

For large but not infinite stiffness, XPBD behaves like regularized PBD update, making it more stable.


6.5 Compliant Constraint Examples

6.5.1 Spring Damper

Consider a system of two rigid bodies connected by a spring at an offset \mathbf{r}_i and \mathbf{r}_j from the respective centres of mass of the two bodies. The linear force is given by

$$\mathbf{f} = -k\delta\mathbf{p}(\mathbf{X}) - \nu\delta\dot{\mathbf{p}}(\mathbf{X}), \quad (6.39)$$

where k is the spring coefficient, ν is the damping coefficient, and $\delta\mathbf{p}(\mathbf{X})$ is the vector from the connection point on body i to body j .



$$\mathbf{p} = \mathbf{x} + \mathbf{r}(\boldsymbol{\theta}). \quad (6.40)$$

The torque is $\boldsymbol{\tau}_w = \mathbf{r} \times \mathbf{f}$, so the overall generalized force is

$$\mathbf{F}(\mathbf{X}) = \begin{pmatrix} \mathbf{f} \\ \mathbf{r}_i \times \mathbf{f} \\ -\mathbf{f} \\ -\mathbf{r}_j \times \mathbf{f} \end{pmatrix} = \begin{pmatrix} \mathbb{1} \\ [\mathbf{r}_i]_{\times} \\ -\mathbb{1} \\ -[\mathbf{r}_j]_{\times} \end{pmatrix} \mathbf{f}. \quad (6.41)$$

This force can be derived from a constraint

$$\mathbf{c}(\mathbf{X}) = \delta\mathbf{p}(\mathbf{X}) \quad (6.42)$$

and using the XPBD method of adding constraints as a spring-like potential

$$V = k\|\delta\mathbf{p}\|^2 + \nu\|\delta\dot{\mathbf{p}}\|^2 \quad (6.43)$$

with

$$\mathcal{D}\mathbf{c} = (\mathbb{1} \quad -[\mathbf{r}_i]_{\times} \quad -\mathbb{1} \quad [\mathbf{r}_j]_{\times}) \quad (6.44)$$

and

$$\begin{aligned} \mathbf{F}(\mathbf{X}) &= -(\mathcal{D}\mathbf{c})^T (k\mathbf{c}(\mathbf{X}) + \nu\dot{\mathbf{c}}(\mathbf{X})), \\ &= \begin{pmatrix} \mathbb{1} \\ [\mathbf{r}_i]_{\times} \\ -\mathbb{1} \\ -[\mathbf{r}_j]_{\times} \end{pmatrix} (-k\delta\mathbf{p} - \nu\delta\dot{\mathbf{p}}). \end{aligned} \quad (6.45)$$

We use the constraint derivative to compute the inverse effective mass

$$\begin{aligned} \mathcal{M}^{-1} &= (\mathbb{1} \quad -[\mathbf{r}_i]_{\times} \quad \mathbb{1} \quad [\mathbf{r}_j]_{\times}) \begin{pmatrix} \frac{1}{m_i}\mathbb{1} & & & \\ & I_i & & \\ & & \frac{1}{m_j}\mathbb{1} & \\ & & & I_j \end{pmatrix}^{-1} \begin{pmatrix} \mathbb{1} \\ [\mathbf{r}_i]_{\times} \\ -\mathbb{1} \\ -[\mathbf{r}_j]_{\times} \end{pmatrix}, \\ &= \left(\frac{1}{m_i} + \frac{1}{m_j} \right) \mathbb{1} - [\mathbf{r}_i]_{\times} I_i^{-1} [\mathbf{r}_i]_{\times} - [\mathbf{r}_j]_{\times} I_j^{-1} [\mathbf{r}_j]_{\times}. \end{aligned} \quad (6.46)$$

Comparing to (6.31) we see find the per iteration update:

Spring damper constraint update.

Lagrange multiplier update:

$$\boldsymbol{\lambda}^{k+1} = \boldsymbol{\lambda}^k - \xi (\mathbb{1} + \alpha \mathcal{M}^{-1})^{-1} \left(k \Delta t^2 \delta \mathbf{p}(\mathbf{X}^{t+1} + \Delta \mathbf{X}) + \nu \Delta t \delta \dot{\mathbf{p}}(\mathbf{X}^{t+1} + \Delta \mathbf{X}) \Delta t + \boldsymbol{\lambda}^k \right), \quad (6.47)$$

where $\alpha = k \Delta t^2 + \nu \Delta t$

Delta:

$$\Delta = \boldsymbol{\lambda}^{k+1} - \boldsymbol{\lambda}^k \quad (6.48)$$

Configuration update:

$$\begin{aligned} \Delta \mathbf{x}'_i &= \Delta \mathbf{x}_i + \frac{1}{m_i} \Delta, \\ \Delta \boldsymbol{\omega}'_i &= \Delta \boldsymbol{\omega}_i + I_i^{-1} \mathbf{r}_i \times \Delta, \\ \Delta \mathbf{x}'_j &= \Delta \mathbf{x}_j - \frac{1}{m_j} \Delta, \\ \Delta \boldsymbol{\omega}'_j &= \Delta \boldsymbol{\omega}_j - I_j^{-1} \mathbf{r}_j \times \Delta, \end{aligned} \quad (6.49)$$

6.5.2 Volumetric Constraint

Up until this point we have been developing constraints between two rigid bodies, but it's worth showing how the method can be used for constraints between multiple bodies. As an example, consider a constraint between four particles making up the nodes of a tetrahedron. We introduce a constraint on the volume V of the tetrahedron


$$c(\mathbf{X}) = 6(V(\mathbf{X}) - V_0). \quad (6.50)$$

If \mathbf{x}_i are the positions of the particles then the volume of the tetrahedron is

$$V(\mathbf{X}) = \frac{1}{6} (\mathbf{x}_i - \mathbf{x}_0) \cdot ((\mathbf{x}_j - \mathbf{x}_0) \times (\mathbf{x}_3 - \mathbf{x}_0)). \quad (6.51)$$

As we are dealing with point particles, there are no angular degrees of freedom and the constraint derivative is

$$\begin{aligned} \mathcal{D}c &= \left(\frac{\partial c}{\partial \mathbf{x}_0} \quad \frac{\partial c}{\partial \mathbf{x}_i} \quad \frac{\partial c}{\partial \mathbf{x}_j} \quad \frac{\partial c}{\partial \mathbf{x}_3} \right), \\ &= \left(-[(\mathbf{d}_i \times \mathbf{d}_j) + (\mathbf{d}_j \times \mathbf{d}_3) + (\mathbf{d}_3 \times \mathbf{d}_i)]^T \quad (\mathbf{d}_j \times \mathbf{d}_3)^T \quad (\mathbf{d}_3 \times \mathbf{d}_i)^T \quad (\mathbf{d}_i \times \mathbf{d}_j)^T \right), \\ &= 2 \left(\mathbf{A}_0^T \quad \mathbf{A}_i^T \quad \mathbf{A}_j^T \quad \mathbf{A}_3^T \right), \end{aligned} \quad (6.52)$$



where $\mathbf{d}_i = (\mathbf{x}_i - \mathbf{x}_0)$ and \mathbf{A}_i is the directed area (area times normal) of the triangle opposite the point i , e.g. $\mathbf{A}_i = \frac{1}{2}\mathbf{d}_j \times \mathbf{d}_3$.

Using this, the effective mass is

$$\mu = \frac{1}{4} \left(\sum_{i=0}^3 \frac{1}{m_i} \|\mathbf{A}_i\|^2 \right)^{-1}. \quad (6.53)$$

Using this in an XPBD update we get

$$\lambda^{k+1} = \lambda^k - \xi(\tilde{\alpha} + (1 + \gamma)\mu)^{-1} \left(6(V(\mathbf{x}^{t+1} + \Delta\mathbf{x}^k) - V_0) + 2\gamma \sum_{i=0}^3 \mathbf{A}_i \cdot (\dot{\mathbf{x}}^{t+1} + \Delta\mathbf{x}^k/\Delta t) + \tilde{\alpha}\lambda^k \right) \quad (6.54)$$

with

$$\Delta\mathbf{x}' = \Delta\mathbf{x} + \frac{1}{m_i} 2\mathbf{A}_i(\lambda^{k+1} - \lambda^k). \quad (6.55)$$

So, each body is pushed in the direction of the normal of the triangle opposite to it.

7 Augmented Lagrangian Method

The augmented Lagrangian method [14, 19] combines the spirit of the compliant constraint method and the method of Lagrange multipliers.

We saw in (4.29) that solving the constrained Lagrangian to first order in $\Delta \mathbf{X}$ is equivalent to solving

$$\min_{\Delta \mathbf{X}} \frac{1}{2} \Delta \mathbf{X} \cdot M \Delta \mathbf{X} + \boldsymbol{\lambda} \cdot (\mathbf{c}(\mathbf{X}^{t+1}) + \mathcal{D}\mathbf{c}\Delta \mathbf{X}) \Delta t^2 \quad (7.1)$$

The augmented Lagrangian method adds an additional penalty term to the Lagrangian $\frac{1}{2} \mathbf{c}(\mathbf{X}) \cdot \rho \mathbf{c}(\mathbf{X})$ for a diagonal matrix of penalty coefficients ρ . For our linearized system this becomes

$$\min_{\Delta \mathbf{X}} \frac{1}{2} \Delta \mathbf{X} \cdot M \Delta \mathbf{X} + \boldsymbol{\lambda} \cdot (\mathbf{c}(\mathbf{X}^{t+1}) + \mathcal{D}\mathbf{c}\Delta \mathbf{X}) \Delta t^2 \quad (7.2)$$

$$+ \frac{1}{2} (\mathbf{c}(\mathbf{X}^{t+1}) + \mathcal{D}\mathbf{c}\Delta \mathbf{X} + \Delta \mathbf{X} \cdot \mathcal{D}^2 \mathbf{c} \Delta \mathbf{X}) \cdot \rho (\mathbf{c}(\mathbf{X}^{t+1}) + \mathcal{D}\mathbf{c}\Delta \mathbf{X} + \Delta \mathbf{X} \cdot \mathcal{D}^2 \mathbf{c} \Delta \mathbf{X}) \Delta t^2. \quad (7.3)$$

Where $\mathcal{D}^2 \mathbf{c}$ is the constraint Hessian using the left trivialized derivative (3.44). Note that I am including terms of order $\Delta \mathbf{X}^2$ and $\Delta \mathbf{X} \cdot \boldsymbol{\lambda}$ in the function as we want the result of the minimization to be linear in $\Delta \mathbf{X}$ and $\boldsymbol{\lambda}$,

The solution to this is

The Augmented Lagrangian system

$$H \Delta \mathbf{X} = - (\mathcal{D}\mathbf{c})^T (\boldsymbol{\lambda} + \rho \mathbf{c}(\mathbf{X}^{t+1})) \Delta t^2 \quad (7.4)$$

where

$$H = M + (\mathcal{D}\mathbf{c})^T \rho \Delta t^2 \mathcal{D}\mathbf{c} + \mathbf{c}(\mathbf{X}^{t+1}) \cdot \rho \Delta t^2 \mathcal{D}^2 \mathbf{c}. \quad (7.5)$$

The augmented Lagrangian method solves this iteratively by first fixing $\boldsymbol{\lambda}^k$ and solving for $\Delta \mathbf{X}$ and then updating the Lagrange multiplier

$$\boldsymbol{\lambda}^{k+1} = \boldsymbol{\lambda}^k + \rho \mathbf{c}(\mathbf{X}^{t+1} + \Delta \mathbf{X}) \quad (7.6)$$

This is then repeated for $k = 1$ to N iterations.

Notice that if the solution converges then $\boldsymbol{\lambda}^{k+1} \sim \boldsymbol{\lambda}^k$ and so $\mathbf{c}(\mathbf{X}^\infty) \sim 0$ and the constraint equations are satisfied.

Using a Gauss-Seidel to solve for the displacements we have

$$\Delta \mathbf{X}_i^{k+1} = \Delta \mathbf{X}_i^k - \xi H_{ii}^{-1} \left((\mathcal{D}\mathbf{c})^T (\boldsymbol{\lambda}^k + \rho \mathbf{c}(\mathbf{X}^{t+1})) \Delta t^2 + H \Delta \mathbf{X}^k \right)_i. \quad (7.7)$$

This can be rewritten in the same way as section 4.8 as a non-linear Gauss-Seidel update:

Non-linear iteration update for the augmented Lagrangian method

$$\Delta \mathbf{X}_i^{k+1} = \Delta \mathbf{X}_i^k - \xi H_{ii}^{-1} \left((\mathcal{D}\mathbf{c})^T (\boldsymbol{\lambda}^k + \rho \mathbf{c}(\mathbf{X}^{t+1} + \Delta \mathbf{X})) \Delta t^2 + M \Delta \mathbf{X}^k \right)_i. \quad (7.8)$$

I.e. we compute the constraint exactly at the current state $\mathbf{X}^{t+1} + \Delta\mathbf{X}$. Compare this to (6.5) and we see that this is just the force update we computed earlier with

$$\mathbf{F}(\mathbf{X}) = -(\mathcal{D}\mathbf{c})^T (\boldsymbol{\lambda}^k + \rho \mathbf{c}(\mathbf{X})). \quad (7.9)$$

After calculating the displacement update at $\boldsymbol{\lambda}^k$, we then update the Lagrange multipliers:

$$\boldsymbol{\lambda}_j^{k+1} = \boldsymbol{\lambda}_j^k + (\rho \mathbf{c}(\mathbf{X}^{t+1} + \Delta\mathbf{X}))_j. \quad (7.10)$$

7.1 Augmented Vertex Block Descent

The method presented here is essentially the update from the Augmented Vertex Block Descent (AVBD) method [11], but extended to rigid bodies (AVBD just presents an update for particles without angular degrees of freedom). However, in the algorithm presented in the AVBD paper, some additional features are presented to improve performance and stability. One such feature they present is the use of an approximation for the Hessian constraint Hessian term $\mathbf{c}(\mathbf{X}^{t+1}) \cdot \rho \Delta t^2 \partial^2 \mathcal{D}^2 \mathbf{c}$. They use a diagonal approximation, which is more efficient but also, importantly, it ensures that H remains symmetric positive definite, and the system remains convergent.

7.1.1 AVBD, PBD and Geometric Stiffness

By doing some rearranging, we can make an interesting comparison to PBD. Let D be the Hessian term

$$D = \mathbf{c}(\mathbf{X}^{t+1}) \cdot \rho \Delta t^2 \partial^2 \mathcal{D} \mathbf{c}. \quad (7.11)$$

Then the equation to be solved is

$$H' \Delta \mathbf{X} = -(\mathcal{D}\mathbf{c})^T (\boldsymbol{\lambda} + \rho \mathbf{c}(\mathbf{X}^{t+1})) \Delta t^2, \quad (7.12)$$

where

$$\begin{aligned} H' &= (M' + (\mathcal{D}\mathbf{c})^T \rho \Delta t^2 \mathcal{D}\mathbf{c}), \\ M' &= M + D. \end{aligned} \quad (7.13)$$

The inverse of H' can be found using the Woodbury matrix identity

$$(A + UCV)^{-1} = A^{-1} - A^{-1}U(C^{-1} + VA^{-1}U)^{-1}VA^{-1}. \quad (7.14)$$

Using this

$$H'^{-1} = M'^{-1} - M'^{-1}(\mathcal{D}\mathbf{c})^T ((\rho \Delta t^2)^{-1} + \mathcal{D}\mathbf{c} M'^{-1}(\mathcal{D}\mathbf{c})^T)^{-1} (\mathcal{D}\mathbf{c})^T M'^{-1} \quad (7.15)$$

and

$$M' \Delta \mathbf{X} = -(\mathcal{D}\mathbf{c})^T ((\rho \Delta t^2)^{-1} + \mathcal{D}\mathbf{c} M'^{-1}(\mathcal{D}\mathbf{c})^T)^{-1} (\rho \Delta t^2)^{-1} (\boldsymbol{\lambda} + \rho \mathbf{c}(\mathbf{X}^{t+1})) \Delta t^2. \quad (7.16)$$

Multiplying through by $\mathcal{D}\mathbf{c} M'^{-1}$ gives

$$\mathcal{D}\mathbf{c} \Delta \mathbf{X} = -\mathcal{D}\mathbf{c} M'^{-1}(\mathcal{D}\mathbf{c})^T ((\rho \Delta t^2)^{-1} + \mathcal{D}\mathbf{c} M'^{-1}(\mathcal{D}\mathbf{c})^T)^{-1} (\rho \Delta t^2)^{-1} (\boldsymbol{\lambda} + \rho \mathbf{c}(\mathbf{X}^{t+1})) \Delta t^2. \quad (7.17)$$

$$(\mathbb{1} + (\rho \Delta t^2)^{-1} (\mathcal{D}\mathbf{c} M'^{-1}(\mathcal{D}\mathbf{c})^T)^{-1}) \mathcal{D}\mathbf{c} \Delta \mathbf{X} = -\mathbf{c}(\mathbf{X}^{t+1}) - (\rho \Delta t^2)^{-1} \boldsymbol{\lambda} \Delta t^2. \quad (7.18)$$

Then, using $\mathbf{c} = \mathcal{D}\mathbf{c} \Delta\mathbf{X}$

$$\mathcal{D}\mathbf{c} M'^{-1}(\mathcal{D}\mathbf{c})^T \lambda \Delta t^2 = \mathbf{c}(\mathbf{X}^{t+1}). \quad (7.19)$$

So, the equation for the Lagrange multipliers is the same as in the PBD case, but with

$$M \rightarrow M + D \quad (7.20)$$

(Note, we get $+\mathbf{c}$ due to the usual sign convention we're following for this section for λ where $\mathbf{c} > 0$ gives $\lambda > 0$).

With AVBD, the D term is approximated by its diagonal. Since M is also block diagonal then it is simple to find the inverse M'^{-1} and the matrix $\mathcal{D}\mathbf{c} M'^{-1}(\mathcal{D}\mathbf{c})^T$ is symmetric and the system is guaranteed to converge.

AVBD does two things:

1. It converts the problem to a new matrix inverse problem with different convergence properties, where the Gauss-Seidel update is per body instead of per constraint
2. It adds in a diagonal Hessian term that can improve accuracy and convergence.

However, the addition of the Hessian term can be applied to any of the methods by making an adjustment to the mass matrix

$$M \rightarrow M + \mathbf{c}(\mathbf{X}^{t+1}) \cdot \rho \Delta t^2 \partial^2 \mathcal{D}\mathbf{c}. \quad (7.21)$$

where AVBD takes the diagonal of the Hessian term to make the calculation simpler and to guarantee convergence.

This Hessian term here is also called geometric stiffness and is discussed in more detail in [1].

7.2 Convergence Comparison vs. PBD For Large Mass Ratios

It's interesting to compare the matrix we're inverting with the augmented Lagrangian method and the force based method of section 6.1 to the matrix inverse from PBD.

$$\begin{aligned} \text{PBD: } A &= \mathcal{D}\mathbf{c} M^{-1}(\mathcal{D}\mathbf{c})^T, \\ \text{AVBD: } A &= M + (\mathcal{D}\mathbf{c})^T \rho \Delta t^2 \mathcal{D}\mathbf{c} + \mathbf{c}(\mathbf{X}^{t+1}) \cdot \rho \Delta t^2 \mathcal{D}^2 \mathbf{c}, \\ \text{Force: } A &= M - (\mathcal{D}\mathbf{F})^T \Delta t^2 \end{aligned} \quad (7.22)$$

In PBD, the mass matrix is sandwiched between the constraint derivative terms, so mass terms can appear in the off-diagonals with terms like $(\mathcal{D}\mathbf{c})_i M_i^{-1} (\mathcal{D}\mathbf{c})_j$. The diagonal will contain the full constraint effective mass inverse $A_{ii} \sim \mu_{ij}^{-1} = 1/m_i + 1/m_j$, whereas off diagonals can contain a single body inverse mass $\sim 1/m_j$

This means that, when looking at the per row diagonal dominance

$$r_i = \frac{\sum_{j \neq i} |A_{ij}|}{|A_{ii}|} \quad (7.23)$$

with PBD we can get terms that depend on the the mass ratio of the bodies $r_i \sim \frac{m_i}{m_i+m_j}$. If the mass of the shared body m_i is large then this ~ 1 and the system will struggle to converge.

For the augmented Lagrangian method and for the force based method, the mass term appears on its own and is block diagonal. This means that we will only ever get a factor of m in the denominator $|A_{ii}|$ and a large mass here will actually improve convergence.

The benefit of the augmented Lagrangian method is that it effectively combines the Lagrange multiplier method of PBD with the constraint penalty method of XPBD, and in doing so can support soft constraints as well as hard constraints. In XPBD, hard constraints require the stiffness term in the constraint penalty function to go to infinity, but with the Augmented Lagrangian method, as it also includes a Lagrange multiplier term, that is not the case.

In addition, unlike PBD, the convergence of the augmented Lagrangian method is not highly sensitive to the mass ratios of the bodies.

7.3 Point to Point Constraint

For comparison, we can implement the point to point stiff constraint example in section 5.1 using the Augmented Lagrangian method.

The constraint is

$$c(\mathbf{X}) = \|\delta\mathbf{p}\| \quad (7.24)$$

for

$$\delta\mathbf{p} = \mathbf{p}_i - \mathbf{p}_j \quad (7.25)$$

where

$$\mathbf{p} = \mathbf{x} + \mathbf{r}(\boldsymbol{\theta}). \quad (7.26)$$

From this we get a constraint derivative

$$\mathcal{D}c = \begin{pmatrix} \hat{\delta\mathbf{p}}^T & (\mathbf{r}_i \times \hat{\delta\mathbf{p}})^T & -\hat{\delta\mathbf{p}}^T & -(\mathbf{r}_j \times \hat{\delta\mathbf{p}})^T \end{pmatrix}. \quad (7.27)$$

and the Hessian is

$$\mathcal{D}^2c = \frac{1}{|\delta\mathbf{p}|} \begin{pmatrix} (\mathbb{1} - P) & -(\mathbb{1} - P)[\mathbf{r}_i]_{\times} & -(\mathbb{1} - P) & (\mathbb{1} - P)[\mathbf{r}_j]_{\times} \\ [\mathbf{r}_i]_{\times}(\mathbb{1} - P) & -[\mathbf{r}_i]_{\times}(\mathbb{1} - P)[\mathbf{r}_i]_{\times} & -[\mathbf{r}_i]_{\times}(\mathbb{1} - P) & [\mathbf{r}_i]_{\times}(\mathbb{1} - P)[\mathbf{r}_j]_{\times} \\ -(\mathbb{1} - P) & (\mathbb{1} - P)[\mathbf{r}_i]_{\times} & (\mathbb{1} - P) & -(\mathbb{1} - P)[\mathbf{r}_j]_{\times} \\ -[\mathbf{r}_j]_{\times}(\mathbb{1} - P) & [\mathbf{r}_j]_{\times}(\mathbb{1} - P)[\mathbf{r}_i]_{\times} & [\mathbf{r}_j]_{\times}(\mathbb{1} - P) & -[\mathbf{r}_j]_{\times}(\mathbb{1} - P)[\mathbf{r}_j]_{\times} \end{pmatrix} \quad (7.28)$$

with P being the projection matrix

$$P = \hat{\delta\mathbf{p}} \hat{\delta\mathbf{p}}^T. \quad (7.29)$$

We can also compute $(\mathcal{D}c)^T \mathcal{D}c$

$$(\mathcal{D}c)^T \mathcal{D}c = \begin{pmatrix} P & -P[\mathbf{r}_i]_{\times} & -P & P[\mathbf{r}_j]_{\times} \\ [\mathbf{r}_i]_{\times}P & -[\mathbf{r}_i]_{\times}P[\mathbf{r}_i]_{\times} & -[\mathbf{r}_i]_{\times}P & [\mathbf{r}_i]_{\times}P[\mathbf{r}_j]_{\times} \\ -P & P[\mathbf{r}_i]_{\times} & P & -P[\mathbf{r}_j]_{\times} \\ -[\mathbf{r}_j]_{\times}P & [\mathbf{r}_j]_{\times}P[\mathbf{r}_i]_{\times} & [\mathbf{r}_j]_{\times}P & -[\mathbf{r}_j]_{\times}P[\mathbf{r}_j]_{\times} \end{pmatrix}. \quad (7.30)$$

Plugging these two into

$$H = M + \rho\Delta t^2(\mathcal{D}\mathbf{c})^T\mathcal{D}\mathbf{c} + \rho\Delta t^2c(\mathbf{X})\mathcal{D}^2\mathbf{c}(\mathbf{X}) \quad (7.31)$$

gives

$$H = \left(M + \rho\Delta t^2 \begin{pmatrix} \mathbb{1} & -[\mathbf{r}_i]_{\times} & -\mathbb{1} & [\mathbf{r}_j]_{\times} \\ [\mathbf{r}_i]_{\times} & -[\mathbf{r}_i]_{\times}[\mathbf{r}_i]_{\times} & -[\mathbf{r}_i]_{\times} & [\mathbf{r}_i]_{\times}[\mathbf{r}_j]_{\times} \\ -\mathbb{1} & [\mathbf{r}_i]_{\times} & \mathbb{1} & -[\mathbf{r}_j]_{\times} \\ -[\mathbf{r}_j]_{\times} & [\mathbf{r}_i]_{\times}[\mathbf{r}_i]_{\times} & [\mathbf{r}_j]_{\times} & -[\mathbf{r}_j]_{\times}[\mathbf{r}_j]_{\times} \end{pmatrix} \right), \quad (7.32)$$

since $c = |\delta\mathbf{p}|$. The projection operators cancel. This means that we can write H in a simple form letting

$$Q = (\mathbb{1} \quad -[\mathbf{r}_i]_{\times} \quad -\mathbb{1} \quad [\mathbf{r}_j]_{\times}). \quad (7.33)$$

With this, the H matrix for this particular system can be written as

$$H = M + Q^T(\rho\Delta t^2)Q. \quad (7.34)$$

In the standard Gauss-Seidel algorithm, we solve the equation

$$H\Delta\mathbf{X} = (\mathcal{D}\mathbf{c})^T(\boldsymbol{\lambda} + \rho\mathbf{c}(\mathbf{X}^{t+1}))\Delta t^2 \quad (7.35)$$

row by row or, in the rigid body case, body by body (i.e. 6 rows at a time). Before we do that, it is interesting to look at what the exact solution is for the full inverse of the 12×12 matrix H in this case of a single constraint between two bodies.

7.3.1 Exact Solution

The inverse of this can be found using the Woodbury matrix identity

$$(A + UCV)^{-1} = A^{-1} - A^{-1}U(C^{-1} + VA^{-1}U)^{-1}VA^{-1}. \quad (7.36)$$

Using this we get

$$H^{-1} = M^{-1} - M^{-1}Q^T \left(\frac{1}{\rho\Delta t^2}\mathbb{1} + QM^{-1}Q^T \right)^{-1} QM^{-1}. \quad (7.37)$$

Notice that $\left(\frac{1}{\rho\Delta t^2}\mathbb{1} + QM^{-1}Q^T \right)^{-1}$ is a 3×3 matrix inversion

$$\left(\frac{1}{\rho\Delta t^2}\mathbb{1} + QM^{-1}Q^T \right)^{-1} = \rho\Delta t^2 (\mathbb{1} + \rho\Delta t^2\mu^{-1})^{-1}, \quad (7.38)$$

where μ is the reduced mass matrix

$$\mu = \left(\frac{1}{m_i} + \frac{1}{m_j} \right) \mathbb{1} - [\mathbf{r}_i]I_i^{-1}[\mathbf{r}_i] - [\mathbf{r}_j]I_j^{-1}[\mathbf{r}_j]. \quad (7.39)$$

Using this, the formula for $\Delta\mathbf{X}$ is (assuming $\boldsymbol{\lambda}^0 = 0$)

$$M\Delta\mathbf{X}^{k+1} = -Q^T (\mathbb{1} + \rho\Delta t^2\mu^{-1})^{-1} \rho\Delta t^2 (\mathbf{p}(\mathbf{X}_i^{t+1}) - \mathbf{p}(\mathbf{X}_j^{t+1})). \quad (7.40)$$

So, if

$$\mathbf{f} = -(\mathbb{1} + \rho\Delta t^2\mu^{-1})^{-1} \rho\Delta t^2 (\mathbf{p}(\mathbf{X}_i^{t+1}) - \mathbf{p}(\mathbf{X}_j^{t+1})) \quad (7.41)$$

then

$$\begin{aligned}
\Delta \mathbf{x}_i &= \frac{1}{m_i} \mathbf{f}, \\
\Delta \boldsymbol{\omega}_i &= I_i^{-1} \mathbf{r}_i \times \mathbf{f}, \\
\Delta \mathbf{x}_j &= -\frac{1}{m_j} \mathbf{f}, \\
\Delta \boldsymbol{\omega}_j &= -I_j^{-1} \mathbf{r}_j \times \mathbf{f}
\end{aligned} \tag{7.42}$$

This is the result for a soft constraint. Notice that it is of the same form as the result for the point constraint except that now μ^{-1} is the inverse of the full effective mass matrix.

7.3.2 Iterative Solution

Solving per body instead of for the full matrix, we can again use the Woodbury matrix identity, except this time for

$$H_i = M_i + \rho \Delta t^2 Q_i^T Q_i \tag{7.43}$$

with

$$M_i = \begin{pmatrix} m_i \mathbb{1} & 0 \\ 0 & I_i \end{pmatrix} \quad Q_i = (\mathbb{1} \quad -[\mathbf{r}_i]_{\times}) \tag{7.44}$$

The inverse is

$$H_i^{-1} = M_i^{-1} - M_i^{-1} Q_i^T \left(\frac{1}{\rho \Delta t^2} \mathbb{1} + Q_i M_i^{-1} Q_i^T \right)^{-1} Q_i M_i^{-1}. \tag{7.45}$$

Note that $\left(\frac{1}{\rho \Delta t^2} \mathbb{1} + Q_i M_i^{-1} Q_i^T \right)^{-1}$ is

$$\left(\frac{1}{\rho \Delta t^2} \mathbb{1} + Q_i M_i^{-1} Q_i^T \right)^{-1} = \rho \Delta t^2 (\mathbb{1} + \rho \Delta t^2 \mu_i^{-1})^{-1}, \tag{7.46}$$

where μ_i is the mass matrix

$$\mu_i = \frac{1}{m_i} \mathbb{1} - [\mathbf{r}_i] I_i^{-1} [\mathbf{r}_i]. \tag{7.47}$$

Using this, the formula for $\Delta \mathbf{X}_i$ is

$$M_i \Delta \mathbf{X}_i^{k+1} = Q_i^T (1 + \rho \Delta t^2 \mu_i^{-1})^{-1} \left(\rho \Delta t^2 Q_i \Delta \mathbf{X}_i^k - (\lambda^k \hat{\delta} \mathbf{p}^k + \rho \delta \mathbf{p}^k) \Delta t^2 \right). \tag{7.48}$$

So, if

$$\mathbf{f} = -(1 + \rho \Delta t^2 \mu_i^{-1})^{-1} \left(\rho \Delta t^2 Q_i \Delta \mathbf{X}_i^k - (\lambda^k \hat{\delta} \mathbf{p}^k + \rho \delta \mathbf{p}^k) \Delta t^2 \right) \tag{7.49}$$

then

$$\begin{aligned}
\Delta \mathbf{x}_i &= \frac{1}{m_i} \mathbf{f}, \\
\Delta \boldsymbol{\omega}_i &= I_i^{-1} \mathbf{r}_i \times \mathbf{f},
\end{aligned} \tag{7.50}$$

The formula for $\Delta \mathbf{X}_j$ will be the same but with \mathbf{f} replaced by $-\mathbf{f}$. There is a difference from the exact solution though. On the first iteration, if $\boldsymbol{\lambda}^0 = 0$, when solving for \mathbf{X}_i^{k+1} , \mathbf{f} is

$$\mathbf{f}_i^1 = -(\mathbb{1} + \rho \Delta t^2 \mu_i^{-1})^{-1} \rho \Delta t^2 \left(\mathbf{p}(\mathbf{X}_i^{t+1}) - \mathbf{p}(\mathbf{X}_j^{t+1}) \right), \tag{7.51}$$

i.e. the constraint function is evaluated in the unperturbed state. However, as we are using Gauss-Seidel, when we go so solve for $\Delta \mathbf{X}_j^{k+1}$ we use the latest value of $\Delta \mathbf{X}_i$, so

$$\mathbf{f}_j^1 = \left(\mathbb{1} + \rho \Delta t^2 \mu_j^{-1} \right)^{-1} \rho \Delta t^2 \left(\mathbf{p}(\mathbf{X}_i^{t+1} + \Delta \mathbf{X}_i^{k+1}) - \mathbf{p}(\mathbf{X}_j^{t+1}) \right), \quad (7.52)$$

i.e. the constraint is evaluated in the new configuration for body i of $\mathbf{X}_i^{t+1} + \Delta \mathbf{X}_i^{k+1}$, and the mass term is μ_j and not μ_i . We do not apply an equal and opposite impulse due to the constraint in a single iteration, and so momentum is not conserved per iteration. However, over the course of multiple iterations the system should converge to the correct solution.

To illustrate this consider, for simplicity, the case of particles instead of bodies, where we just have linear degrees of freedom

$$\begin{aligned} \mathbf{f}_i^1 &= -\frac{\rho \Delta t^2}{m_i + \rho \Delta t^2} \delta \mathbf{x}^{t+1}, \\ \Delta \mathbf{x}_i^1 &= -\frac{\rho \Delta t^2 / m_i}{1 + \rho \Delta t^2 / m_i} \delta \mathbf{x}^{t+1}, \\ \mathbf{f}_j^1 &= \frac{\rho \Delta t^2}{m_j + \rho \Delta t^2} \left(\delta \mathbf{x}^{t+1} - \frac{\rho \Delta t^2 / m_i}{1 + \rho \Delta t^2 / m_i} \delta \mathbf{x}^{t+1} \right), \\ &= \frac{\rho \Delta t^2}{(m_j + \rho \Delta t^2)(1 + \rho \Delta t^2 / m_i)} \delta \mathbf{x}^{t+1}, \\ \Delta \mathbf{x}_j^1 &= \frac{\rho \Delta t^2 / m_j}{(1 + \rho \Delta t^2 (1/m_i + 1/m_j) + (\rho \Delta t^2)^2 / (m_i m_j))} \delta \mathbf{x}^{t+1}, \end{aligned} \quad (7.53)$$

Even after one iteration, $\Delta \mathbf{x}^1$ is approaching the correct solution of

$$\Delta \mathbf{x}_j^1 = \frac{\rho \Delta t^2 / m_j}{(1 + \rho \Delta t^2 (1/m_i + 1/m_j))} \delta \mathbf{x}^{t+1} \quad (7.54)$$

with an error in the denominator of order $(\rho \Delta t^2)^2$.

For the next iteration λ gets updated

$$\lambda^1 = \rho \|\delta \mathbf{p}(\mathbf{x}^{t+1} + \Delta \mathbf{x}^k)\| \quad (7.55)$$

then subsequent updates bring $\Delta \mathbf{x}_i$ and $\Delta \mathbf{x}_j$ closer to the exact solution.

The iterative update converges quickly to the exact solution, but per iteration, since forces are applied asymmetrically, momentum is not necessarily conserved.

where the diagonal elements are the inverse effective mass between adjacent bodies

$$\mu_{i\ i+1}^{-1} = \hat{\boldsymbol{p}}_{i\ i+1} \cdot \left(\left(\frac{1}{m_i} + \frac{1}{m_{i+1}} \right) \mathbb{1} - [\mathbf{a}_i]_{\times} I_i^{-1} [\mathbf{a}_i]_{\times} - [\mathbf{b}_{i+1}]_{\times} I_{i+1}^{-1} [\mathbf{b}_{i+1}]_{\times} \right) \hat{\boldsymbol{p}}_{i\ i+1} \quad (8.5)$$

and the off diagonal elements represent coupling between adjacent constraints

$$\alpha_{i\ i+2} = -\hat{\boldsymbol{p}}_{i\ i+1} \cdot \left(\frac{1}{m_{i+1}} \mathbb{1} - [\mathbf{a}_{i+1}]_{\times} I_{i+1}^{-1} [\mathbf{b}_{i+1}]_{\times} \right) \hat{\boldsymbol{p}}_{i+1\ i+2}. \quad (8.6)$$

In general, the coupling will be strongest when $\hat{\boldsymbol{p}}_{i\ i+1} = \hat{\boldsymbol{p}}_{i+1\ i+2}$, e.g. the bodies are all lined up in a straight line. In this case, the correction to solve one constraint will directly feed into the error for the adjacent constraint the matrix reduces to

$$\begin{pmatrix} \left(\frac{1}{m_1} + \frac{1}{m_2} \right) & -\frac{1}{m_2} & & & \\ -\frac{1}{m_2} & \left(\frac{1}{m_2} + \frac{1}{m_3} \right) & & & \\ & & \ddots & & \\ & & & \left(\frac{1}{m_{N-1}} + \frac{1}{m_N} \right) & \end{pmatrix}. \quad (8.7)$$

Notice that only the first and last rows are strictly diagonally dominant. For the other rows

$$\frac{\sum_{i \neq j} |A_{ij}|}{|A_{ii}|} = \frac{\frac{1}{m_i} + \frac{1}{m_{i+1}}}{\frac{1}{m_i} + \frac{1}{m_{i+1}}} = 1. \quad (8.8)$$

The system is still guaranteed to converge, since the matrix is SPD, but the rate of convergence may be slow.

8.2 Effect of Iteration Count

To examine the rate of convergence, we find the spectral radius of the Gauss-Seidel iteration matrix $G = -(D + L)^{-1} L^T$. First, assume that all masses are equal, say $m_i = 1$, then the matrix is

$$\begin{pmatrix} 2 & -1 & & & \\ -1 & 2 & -1 & & \\ & -1 & \ddots & & \\ & & & \ddots & \\ & & & & 2 \end{pmatrix}. \quad (8.9)$$

In this case, the iteration matrix is

$$G = \begin{pmatrix} 0 & \frac{1}{2} & & & \\ 0 & \frac{1}{4} & \frac{1}{2} & & \\ 0 & \frac{1}{8} & \frac{1}{4} & \frac{1}{2} & \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \quad (8.10)$$

and the spectral radius is

$$\rho(G) = \cos^2 \left(\frac{\pi}{N} \right). \quad (8.11)$$

So for large N the spectral radius goes to

$$\rho(G) \rightarrow 1 - \frac{\pi^2}{N^2} \text{ as } N \rightarrow \infty. \quad (8.12)$$

So, as you would expect, as the number of links in the chain goes to infinity, the spectral radius goes to 1 and so the rate of convergence goes to zero.

In general, the error from one iteration to the next will decrease like

$$\begin{aligned} \|\epsilon^k\| &\sim |\rho(G)|^k \|\epsilon_0\|, \\ &\sim \left| \cos\left(\frac{\pi}{N}\right) \right|^{2k} \|\epsilon_0\|, \\ &\sim \exp\left(-\frac{k\pi^2}{N^2}\right) \|\epsilon_0\| \text{ for large } N \end{aligned} \quad (8.13)$$

So, to reduce the error to $1 - \alpha$ times the original error takes

$$k \sim \frac{N^2}{\pi^2} \ln\left(\frac{1}{1 - \alpha}\right) \quad (8.14)$$

iterations. See figure 2 for an example plot showing the iterations required to remove different fractions of the original error in a chain of length 10.

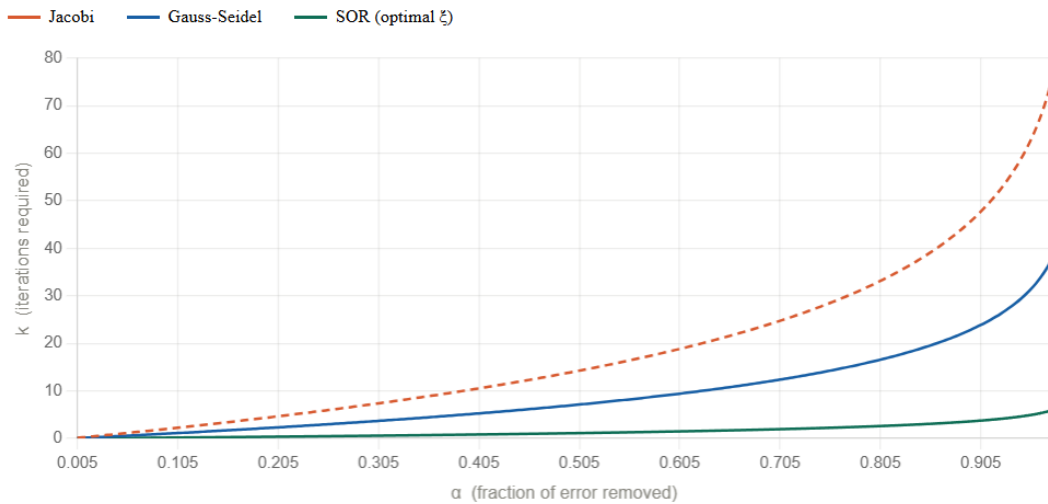


Figure 2: Comparing Gauss-Seidel vs. Jacobi vs. Gauss-Seidel with optimal SOR methods on the number of iterations required to reduce a fraction of the error in a chain of 10 bodies. The optimal relaxation in the SOR method is $\xi^* \sim 1.5$

8.2.1 Relaxation

The addition of a relaxation factor ξ modifies the iteration matrix by

$$G_\xi = (D + \xi L)^{-1}((1 - \xi)D - \xi L^T). \quad (8.15)$$

To find the eigenvalues λ of G_ξ we use two standard results. First, for the equal-mass chain the eigenvalues of the Jacobi iteration matrix G_J are

$$\mu_k = \cos\left(\frac{k\pi}{1 + N}\right), \quad k = 1, \dots, N. \quad (8.16)$$

Second, since A is tridiagonal (consistently ordered), the eigenvalues λ of G_ξ are related to the Jacobi eigenvalues μ_k by Young's relation [25]

$$(\lambda + \xi - 1)^2 = \xi^2 \mu_k^2 \lambda. \quad (8.17)$$

Rearranging, this is a quadratic equation in λ :

$$\lambda^2 - (\xi^2 \mu_k^2 - 2(\xi - 1))\lambda + (\xi - 1)^2 = 0 \quad (8.18)$$

with discriminant

$$\Delta = \xi^4 \mu_k^4 - 4\xi^2 \mu_k^2 (\xi - 1). \quad (8.19)$$

For $\xi = 1$ (standard Gauss-Seidel), $\Delta = \mu_k^4 > 0$ and both roots are real. As $\xi \rightarrow 2$, $\Delta \rightarrow 16\mu_k^2(\mu_k^2 - 1)$, which is negative since $|\mu_k| < 1$, giving two complex conjugate roots with magnitude $|\lambda| = \xi - 1 \rightarrow 1$. So $|\lambda|$ is minimized by the value of ξ for which the discriminant is zero, since that is the boundary between real and complex roots. Setting $\Delta = 0$ for the dominant Jacobi eigenvalue $\mu_1 = \cos(\frac{\pi}{N})$ gives

$$\xi^{*2} \mu_1^2 = 4(\xi^* - 1). \quad (8.20)$$

Using $\sqrt{1 - \mu_1^2} = \sin(\frac{\pi}{N})$, this gives

$$\xi^* = \frac{2}{1 + \sin(\frac{\pi}{N})}. \quad (8.21)$$

When $\Delta = 0$ the quadratic has a repeated root

$$\lambda^* = \xi^* - 1 = \frac{1 - \sin(\frac{\pi}{N})}{1 + \sin(\frac{\pi}{N})}, \quad (8.22)$$

In the limit of large N this gives

$$\xi^* \sim 2\left(1 - \frac{\pi}{N}\right), \quad \rho(G_{\xi^*}) \sim 1 - \frac{2\pi}{N}. \quad (8.23)$$

This is a substantial improvement over the standard Gauss-Seidel result $\rho(G) \sim 1 - \frac{\pi^2}{N^2}$, as the iteration count to remove a fraction α of the error is now

$$k \sim \frac{N}{2\pi} \ln\left(\frac{1}{1 - \alpha}\right), \quad (8.24)$$

scaling with N instead of N^2 .

The practical difficulty with SOR in general scenarios is that we do not always have a fixed constraint topology, so we cannot pre-compute an optimal ξ .

8.2.2 Jacobi Iterations

It is instructive to briefly look at the Jacobi method for solving the linear system. This is the same as the Gauss-Seidel method in that we solve row by row, but the difference being that we do not propagate the result from the solution of the previous row to the solution of the next row. In this method, the iteration matrix is

$$G_J = -D^{-1}(L + L^T). \quad (8.25)$$

For the case of a straight chain with equal masses, this is

$$G_J = \begin{pmatrix} 0 & \frac{1}{2} & & & \\ \frac{1}{2} & 0 & \frac{1}{2} & & \\ & \frac{1}{2} & 0 & \frac{1}{2} & \\ & & \ddots & \ddots & \ddots \end{pmatrix} \quad (8.26)$$

and the spectral radius is

$$\rho(G_J) = \cos\left(\frac{\pi}{N}\right). \quad (8.27)$$

As A is a tridiagonal matrix, Young's theorem holds and

$$\rho(G) = \rho(G_J)^2. \quad (8.28)$$

and, for large N , to reduce the error to $1 - \alpha$ times the original error takes

$$k \sim \frac{2N^2}{\pi^2} \ln\left(\frac{1}{1 - \alpha}\right), \quad (8.29)$$

which is exactly twice the number of iterations required by the Gauss-Seidel algorithm. This is why Gauss-Seidel is generally preferred.

8.2.3 AVBD Iterations

In the AVBD method, instead of inverting $\mathcal{D}\mathbf{c} M^{-1}(\mathcal{D}\mathbf{c})^T$, we are inverting H , where

$$H = M + (\mathcal{D}\mathbf{c})^T \rho \Delta t^2 \mathcal{D}\mathbf{c} + \mathbf{c}(\mathbf{X}^{t+1}) \cdot \rho \Delta t^2 \frac{\partial^2 \mathbf{c}}{\partial \mathbf{X}^2}. \quad (8.30)$$

For the straight chain of equal masses this is

$$H = \begin{pmatrix} (m + \rho \Delta t^2) \mathbb{1} & -\rho \Delta t^2 \mathbb{1} & & & \\ -\rho \Delta t^2 \mathbb{1} & (m + 2\rho \Delta t^2) \mathbb{1} & -\rho \Delta t^2 \mathbb{1} & & \\ & -\rho \Delta t^2 \mathbb{1} & (m + 2\rho \Delta t^2) \mathbb{1} & -\rho \Delta t^2 \mathbb{1} & \\ & & & \ddots & \\ & & & & (m + \rho \Delta t^2) \mathbb{1} \end{pmatrix} \quad (8.31)$$

This is a block tridiagonal matrix with additional terms on the boundaries (so the first and last diagonal elements are different from the interior). The boundary differences do not affect the spectral radius for large N so we can use standard result for the spectral radius of the iteration matrix of

$$\rho(G) \sim \left(\frac{2\rho \Delta t^2}{m + 2\rho \Delta t^2}\right)^2 \cos^2\left(\frac{\pi}{1 + N}\right). \quad (8.32)$$

This is the same as for the PBD but with the additional scaling factor

$$\beta = \left(\frac{2\rho \Delta t^2}{m + 2\rho \Delta t^2}\right). \quad (8.33)$$

And now reducing the error to $1 - \alpha$ times the original error requires roughly

$$k \sim \left(\frac{1}{1 - 2\frac{N^2}{\pi^2} \ln \beta}\right) \frac{N^2}{\pi^2} \ln\left(\frac{1}{1 - \alpha}\right) \quad (8.34)$$

In the case of a large stiffness value $\rho \rightarrow \infty$ so $\beta \rightarrow 1$ and this reduces to the PBD result. For smaller stiffness values $\beta < 1$ and the number of required iterations is reduced. This can also be seen by looking at the diagonal dominance of the matrix H

$$r_i = \frac{\sum_{i \neq j} |A_{ij}|}{|A_{ii}|} = \frac{2\rho\Delta t^2}{m + 2\rho\Delta t^2} \quad (8.35)$$

When $m \gg \rho\Delta t^2$, r_i becomes small and $k \rightarrow 0$. This makes sense, as the system is weak springs trying to pull together infinite masses, so the solution is no movement and no iterations are required. What this does show is that in general AVBD, like PBD, requires the number of iterations to scale like N^2 in order to solve the chain to the same error.

8.3 Effect of Mass Ratios

Consider the matrix (8.7). So far, we have assumed all masses are equal. If we now make the end masses m_1 and m_N different, with $m_i = m$, for $i = 2, \dots, N-1$ and $m_1 = m_N = M$, with $M > m$. This changes the diagonal dominance of the first and last rows

$$r_{\text{ends}} = \frac{\frac{1}{m}}{\frac{1}{m} + \frac{1}{M}} \sim 1, \text{ for } M \gg m. \quad (8.36)$$

This means that all rows now have r either 1 or close to 1. This has a big impact on convergence. If just one end had a light mass then there is at least one row with a smaller r , which eventually allows the system to converge as the error correction from that row can propagate up the chain. With both ends having a heavy mass, the error correction on the ends will only solve a fraction of the error, this then has to propagate back up the chain to the other end, where again a fraction of the error will be solved. This then repeats.

To analyze the convergence of this we start by looking at the Jacobi iteration matrix $G_J = -D^{-1}(L + L^T)$. If we find the spectral radius of this we can then infer the spectral radius of the Gauss-Siedel iteration matrix using the Young's theorem relation $\rho(G) = \rho(G_J)^2$, which holds for triangular matrices, like we have for the chain example.

Explicitly, setting $m_i = 1$ for simplicity, the matrix $\mathcal{Dc} M^{-1}(\mathcal{Dc})^T$ is

$$A = \mathcal{Dc} M^{-1}(\mathcal{Dc})^T = \begin{pmatrix} 1 + \gamma & -1 & & & \\ -1 & 2 & -1 & & \\ & -1 & \ddots & & \\ & & & \ddots & \\ & & & & 1 + \gamma \end{pmatrix}, \quad (8.37)$$

where $\gamma = \frac{m}{M}$ is the mass ratio.

If the μ is an eigenvalue of G_J , then

$$A\mathbf{x} = (1 - \mu)D\mathbf{x}. \quad (8.38)$$

This is a generalized eigenvector problem. For $\gamma \rightarrow 0$ the matrix A has an eigenvector $\mathbf{v} = (1 \ 1 \ \dots \ 1)^T$ with eigenvalue zero. We can use this using a Rayleigh quotient to find an approximate eigenvalue λ closest to zero, which will give the dominant eigenvalue of G_J $\mu = 1 - \lambda$ which

is the eigenvalue closest to 1.

$$\begin{aligned}\lambda &\sim \frac{\mathbf{v}^T A \mathbf{v}}{\mathbf{v}^T D \mathbf{v}}, \\ &\sim \frac{2\gamma}{2\gamma + 2 + 2(N-2)}, \\ &\sim \frac{\gamma}{N} \text{ for large } N\end{aligned}\tag{8.39}$$

so

$$\mu \sim 1 - \frac{\gamma}{N}\tag{8.40}$$

and the spectral radius for the Gauss-Seidel iteration matrix is

$$\begin{aligned}\rho(G) &\sim \left(1 - \frac{\gamma}{N}\right)^2, \\ &\sim 1 - \frac{2\gamma}{N}.\end{aligned}\tag{8.41}$$

Using this, the number of iterations required to reduce the error by a factor of $1 - \alpha$ is for small γ

$$k \sim \frac{N}{2\gamma} \ln \left(\frac{1}{1 - \alpha} \right).\tag{8.42}$$

So, say you have a chain of length $N = 10$, for a mass ratio of 1 we have the previous result $k \sim 10 \ln(1/(1 - \alpha))$, and with a mass ratio of 100 this will be $k \sim 500 \ln(1/(1 - \alpha))$, so a mass ratio of 100 means having to increase the number of iterations by $\sim 50\times$ to solve for the same correction to the initial error.

8.3.1 Relaxation

Following the same procedure from the previous section 8.2.1 we can find an optimal relaxation value in this case of large masses at either end

$$\begin{aligned}\xi^* &= \frac{2}{1 + \sqrt{1 - \rho(G)}}, \\ &\sim \frac{2}{1 + \sqrt{2\gamma/N}}.\end{aligned}\tag{8.43}$$

With this, the optimal spectral radius is

$$\rho(G_{\xi}^*) \sim 1 - 2\sqrt{\frac{2\gamma}{N}}\tag{8.44}$$

and the number of iterations goes like

$$k \sim \frac{1}{2} \sqrt{\frac{N}{2\gamma}} \ln \left(\frac{1}{1 - \alpha} \right).\tag{8.45}$$

This is an improvement over the non-relaxed result, going from a N/γ scaling to a $\sqrt{N/\gamma}$ scaling. For example, with a chain length of 10 and a mass ratio of 100 the iterations now go like $k \sim 11 \ln(1/(1 - \alpha))$ compared to $k \sim 500 \ln(1/(1 - \alpha))$ for no relaxation.

8.3.2 AVBD Iterations

For AVBD, setting the two end masses to M makes the matrix H become

$$H = \begin{pmatrix} (M + \rho\Delta t^2)\mathbb{1} & -\rho\Delta t^2\mathbb{1} & & & \\ -\rho\Delta t^2\mathbb{1} & (m + 2\rho\Delta t^2)\mathbb{1} & -\rho\Delta t^2\mathbb{1} & & \\ & -\rho\Delta t^2\mathbb{1} & (m + 2\rho\Delta t^2)\mathbb{1} & -\rho\Delta t^2\mathbb{1} & \\ & & & \ddots & \\ & & & & (M + \rho\Delta t^2)\mathbb{1} \end{pmatrix} \quad (8.46)$$

The effect of this is to improve the diagonal dominance for the first and last constraint rows

$$\begin{aligned} r_{\text{ends}} &= \frac{\rho\Delta t^2}{M + \rho\Delta t^2}, \\ &\sim \frac{\rho\Delta t^2}{M} \text{ for large } M \end{aligned} \quad (8.47)$$

Proceeding using the Rayleigh quotient method used in the PBD case, we find

$$\rho(G) \sim \left(\frac{\rho\Delta t^2 N}{M + \rho\Delta t^2 N} \right)^2 \quad (8.48)$$

and

$$k \sim \frac{\rho\Delta t^2 N}{2M} \ln \left(\frac{1}{1 - \alpha} \right), \quad (8.49)$$

where the assumption that N is large but $\rho\Delta t^2/M$ is not as large. In this range the number of iterations scales with N instead of N^2 and the convergence is improved.

8.4 Results

These results are for the simplified chain, where all of the bodies are in a straight line and the angular terms vanish.

Both AVBD and PBD require the iteration count to scale with the square of the number of bodies in the chain in order to correct the same error, but AVBD can soften the constraint resolution through the ρ parameter, which will allow stretching but improve convergence. Using SOR can greatly improve the convergence of Gauss-Seidel, but finding the optimal relaxation parameter requires prior knowledge and analysis of the constraint matrix.

When large masses are added to either end of the chain the convergence is greatly affected. For large mass ratio $1/\gamma$, Gauss Siedel iterations scale by N/γ , meaning the mass ratio is directly proportional to the required number of iterations. AVBD does better in this scenario in the regime $\rho\Delta t^2/M \sim 1$ and the iterations scale by N .

9 Summary

We first went through a derivation of position based dynamics and showed the following

- Position Based Dynamics can be derived as a first order approximate solution to the constrained rigid body equations of motion
- This is equivalent to solving the constraint equations at first order subject to a least squares minimization over the metric defined by the mass matrix M

The convergence section proved the following

- For a symmetric mass matrix, the Gauss-Seidel solution for the PBD rigid body system is guaranteed to converge.
- The rate of convergence depends on the spectral radius of the Gauss-Seidel update matrix, and can be estimated by the diagonal dominance of the matrix $\mathcal{Dc} M^{-1}(\mathcal{Dc})^T$

Next, we went through finding a first order solution to the general equations of motion for non-linear forces and showed:

- Solving the linearized equations of motion for forces instead of constraints is equivalent to the Vertex Block Descent method where the iteration loop is over bodies instead of constraints.
- XPBD swaps constraints with Lagrange multipliers for compliant constraints implemented through potential penalty functions
- This naturally fits into the VBD framework, but XPBD transforms the problem into a constraint based update through a change of variables
- The XPBD update is equivalent to the PBD least squares minimization but with the addition on a regularization term proportional to the inverse of the spring stiffness.

Next, we went through the Augmented Lagrangian method and showed:

- The Augmented Lagrangian method mixes both the constraint penalty method and the Lagrange multiplier method, supporting both stiff and compliant constraints
- AVBD is a flavour of this that approximates the Hessian term by taking only the diagonal. This ensures convergence.
- The Hessian term is known as geometric stiffness, and can be included in all position based methods as a modification to the mass matrix.

Finally, the convergence section went through the convergence of a simple chain of bodies and showed

- For chains with uniform mass both PBD and AVBD require the number of iterations to scale with the number of bodies squared for convergence.
- Systematic Over Relaxation can greatly improve convergence. An optimal SOR parameter can reduce the convergence scaling to N instead of N^2 .
- The body based iteration in VBD, AVBD and force based update methods means that the convergence of these methods is not adversely affected by high mass ratios, unlike constraint based methods like PBD.

References

- [1] Sheldon Andrews, Marek Teichmann, and Paul G. Kry. Geometric stiffness for real-time constrained multibody dynamics. *Computer Graphics Forum*, 36(2):235–246, 2017.
- [2] David Baraff. An introduction to physically based modeling: Rigid body simulation. *SIGGRAPH Course Notes*, 1997.
- [3] David Baraff. Rigid body simulation ii: Nonpenetration constraints. *SIGGRAPH Course Notes*, 1999.
- [4] Joao Barata and Mahir Hussein. The moore-penrose pseudoinverse. a tutorial review of the theory. *Brazilian Journal of Physics*, 42, 10 2011.
- [5] Timothy D. Barfoot. *State Estimation for Robotics*. Cambridge University Press, Cambridge, UK, 2017.
- [6] Dennis S. Bernstein, Ankit Goel, and Omran Kouba. Deriving euler’s equation for rigid-body rotation via lagrangian dynamics with generalized coordinates. *Mathematics*, 11(12):2727, June 2023.
- [7] Anthony M. Bloch. *Nonholonomic Mechanics and Control*, volume 24 of *Interdisciplinary Applied Mathematics*. Springer, New York, 2003.
- [8] Erin Catto. Iterative dynamics with temporal coherence. In *Game Developers Conference (GDC)*, 2005.
- [9] Anka He Chen, Ziheng Liu, Yin Yang, and Cem Yuksel. Vertex block descent. *ACM Transactions on Graphics*, 43(4):1–16, 2024.
- [10] Kenny Erleben. *Game Physics Engine Development*. CRC Press, Boca Raton, FL, 2 edition, 2010.
- [11] Chris Giles, Elie Diaz, and Cem Yuksel. Augmented vertex block descent. *ACM Transactions on Graphics*, 44(4):1–12, July 2025.
- [12] Herbert Goldstein, Charles P. Poole, and John L. Safko. *Classical Mechanics*. Addison-Wesley, San Francisco, 3 edition, 2002.
- [13] David Hestenes. *New Foundations for Classical Mechanics*, volume 15 of *Fundamental Theories of Physics*. Reidel / Kluwer Academic Publishers, Dordrecht, Netherlands, 1986. Second revised edition published 1999, ISBN 978-0792355144.
- [14] Magnus R. Hestenes. Multiplier and gradient methods. *Journal of Optimization Theory and Applications*, 4(5):303–320, 1969.
- [15] Miles Macklin, Matthias Müller, and Nuttapon Chentanez. Xpbd: Position-based simulation of compliant constrained dynamics. *ACM SIGGRAPH Motion in Games*, 2016.
- [16] Ian Millington. *Game Physics Engine Development*. Morgan Kaufmann, 2007.
- [17] Matthias Müller, Bruno Heidelberger, Marcus Hennix, and John Ratcliff. Position Based Dynamics. In Cesar Mendoza and Isabel Navazo, editors, *Vriphys: 3rd Workshop in Virtual Reality, Interactions, and Physical Simulation*. The Eurographics Association, 2006.

- [18] Roger Penrose. A generalized inverse for matrices. *Proceedings of the Cambridge Philosophical Society*, 51:406–413, 01 1955.
- [19] M. J. D. Powell. A method for nonlinear constraints in minimization problems. *Optimization*, pages 283–298, 1969.
- [20] Morten Silcowitz, Sarah Niebe, and Kenny Erleben. Projected gauss–seidel subspace minimization method for interactive rigid body dynamics. In *Proceedings of the International Conference on Computer Graphics Theory and Applications (GRAPP)*. SciTePress, 2010.
- [21] Joan Solà. Quaternion kinematics for the error-state kalman filter. Technical report iri-tr-16-02, Institut de Robòtica i Informàtica Industrial (CSIC–UPC), 2016. Also available as arXiv:1711.02508 [cs.RO].
- [22] A Tikhonov and V Arsenin. *Solution of Ill-Posed Problem*. 01 1977.
- [23] David Tong. Rigid body dynamics. <https://www.damtp.cam.ac.uk/user/tong/dynamics/three.pdf>, 2019. Lecture notes, Department of Applied Mathematics and Theoretical Physics, University of Cambridge.
- [24] Lloyd N. Trefethen and David Bau III. *Numerical Linear Algebra*. SIAM, Philadelphia, PA, USA, 1997.
- [25] Richard S. Varga. *Matrix Iterative Analysis*. Springer, New York, 2 edition, 2000.
- [26] Anthony Zee. *Group Theory in a Nutshell for Physicists*. In a Nutshell. Princeton University Press, 2016.