



# Improving an Iterative Physics Solver Using a Direct Method

Maciej Mizerski  
Technical Director, Roblox

**GAME DEVELOPERS CONFERENCE**  
MARCH 16–20, 2020 | #GDC20

# ROBLOX



# Physics Solver in Roblox

Computes the motion of constrained rigid bodies.





# Support for Complex Mechanisms



# Typical Physics: PGS Solver





# Roblox Physics: LDL-PGS Solver



Symbolic Phase  
Once per mechanism

Mechanism Structure

LDL Decomposition  
Program

Numeric Phase  
Every frame

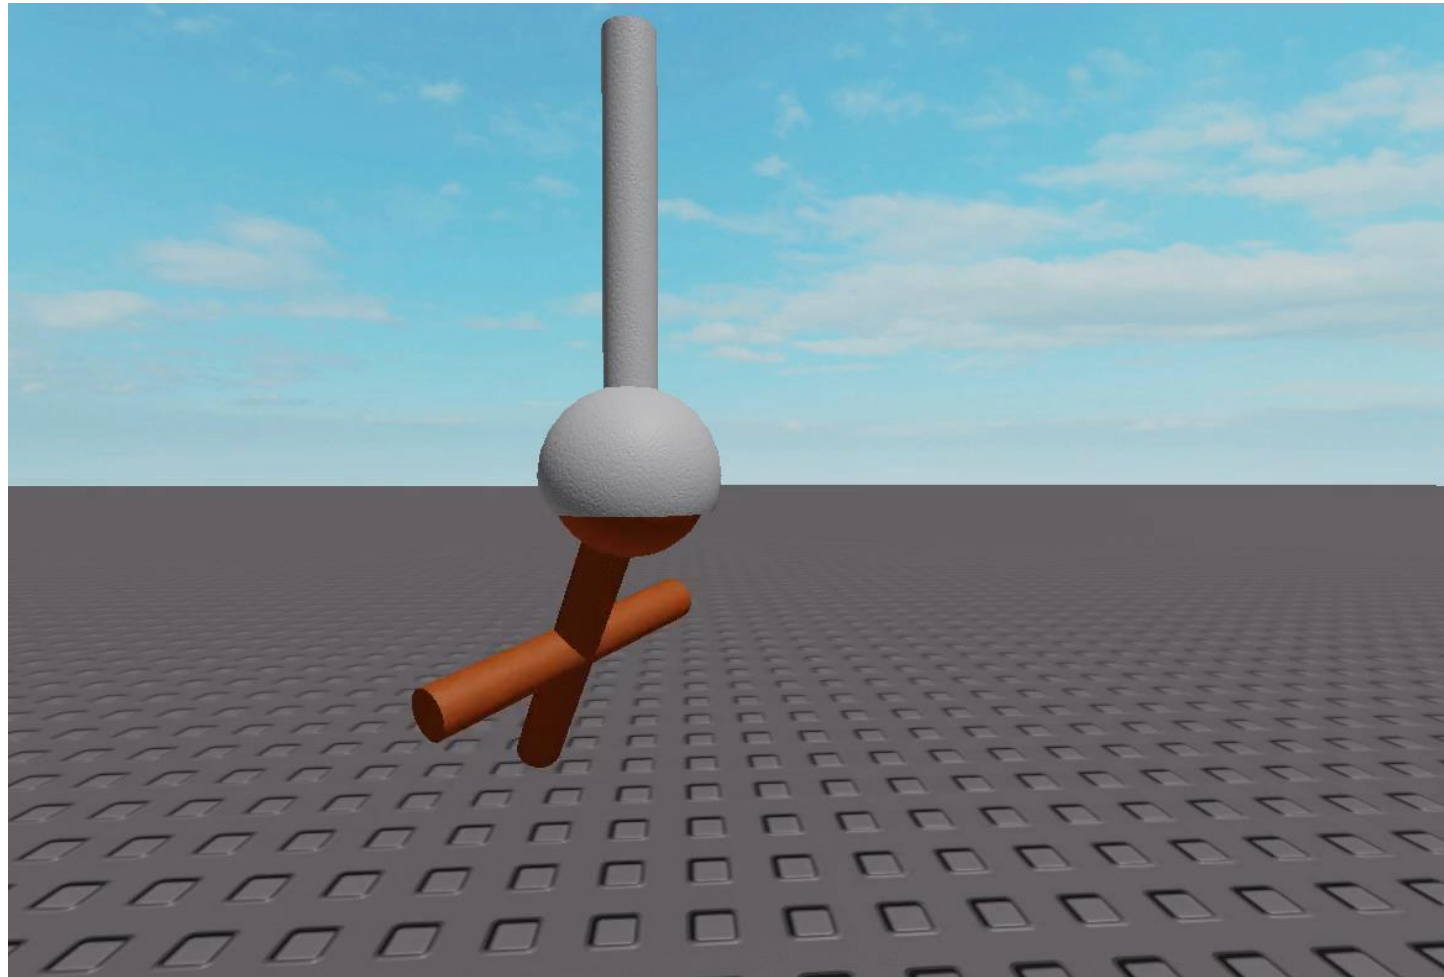
Mechanism State

N-1 iterations of  
PGS

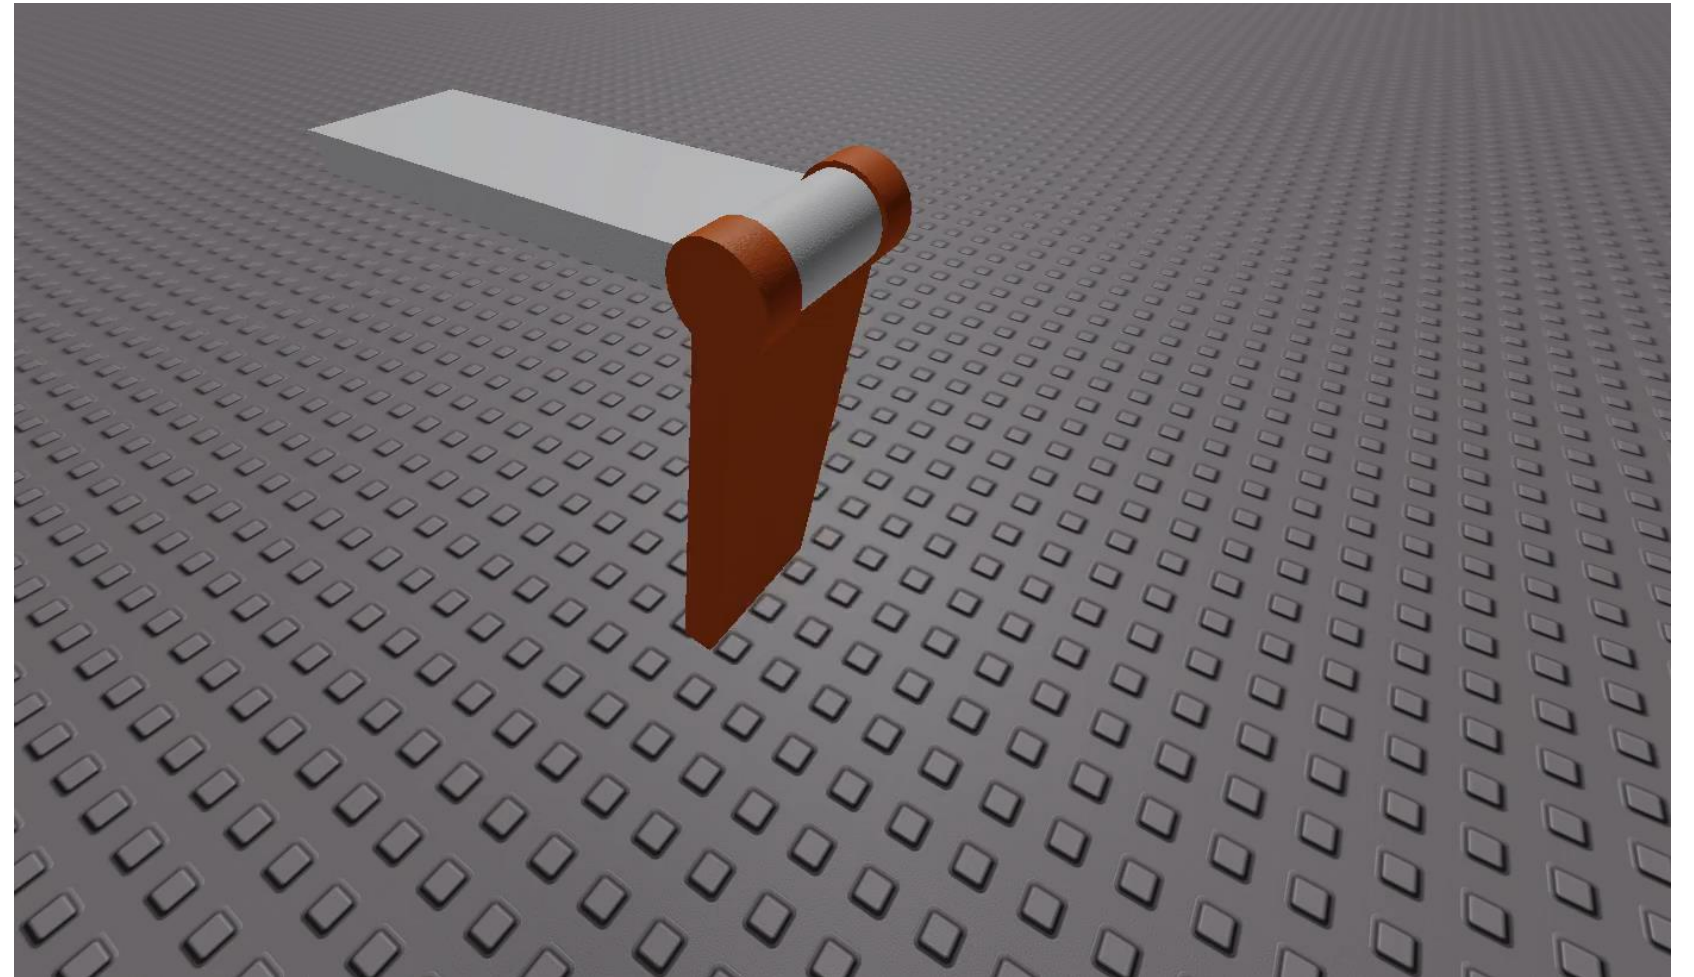
LDL Decomp

Constraint Impulses

# Constraint Examples



**Ball In Socket**



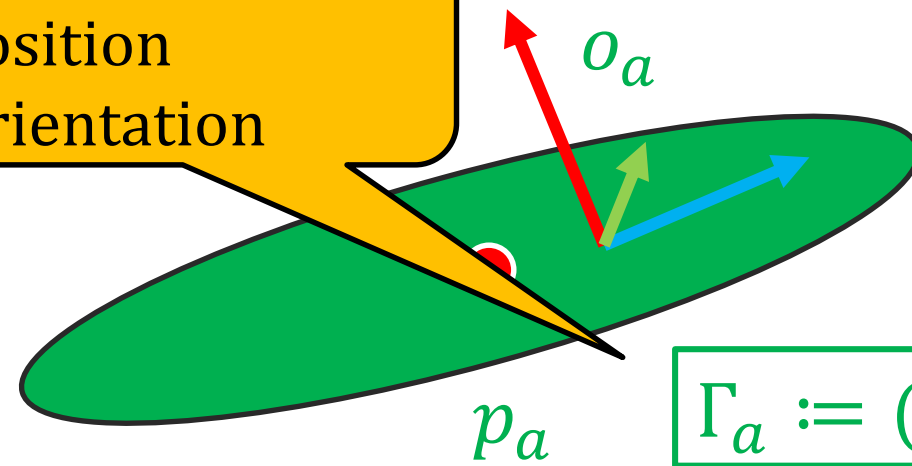
**Hinge**



# Constraint Examples

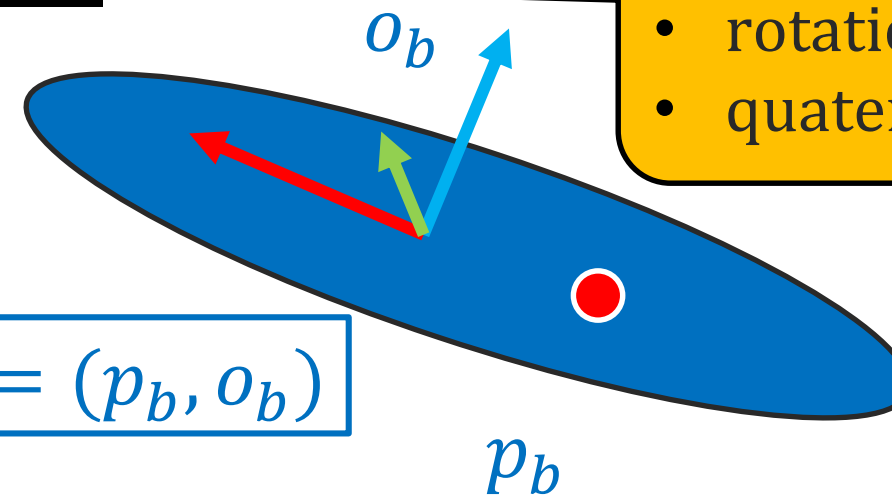
- Ball in Socket
- Hinge
- Rod (Distance Constraint)
- Prismatic
- Cylindrical
- Angular limit
- Positional limit
- Rope
- ...

Coordinates:  
1. Position  
2. Orientation



## Constraint

$$\Gamma_b := (p_b, o_b)$$



$o_a, o_b$ :  
• rotation matrices  
• quaternions

...

Differentiable function of coordinates:

$$c(\Gamma_a, \Gamma_b, \dots) \in \mathbb{R}^d$$

$\mathbb{R}^d = \text{Constraint Space}$   
 $d = \text{Degree}$

Bodies respect the constraint if:

$$c(\Gamma_a, \Gamma_b, \dots) = 0$$

Must be **regular** on the zero locus:

Jacobian has full rank  $\Leftrightarrow d = \text{number of DOF removed}$

# Ball in Socket

$x_a$  - pivot in local space of **a**  
 $x_b$  - pivot in local space of **b**

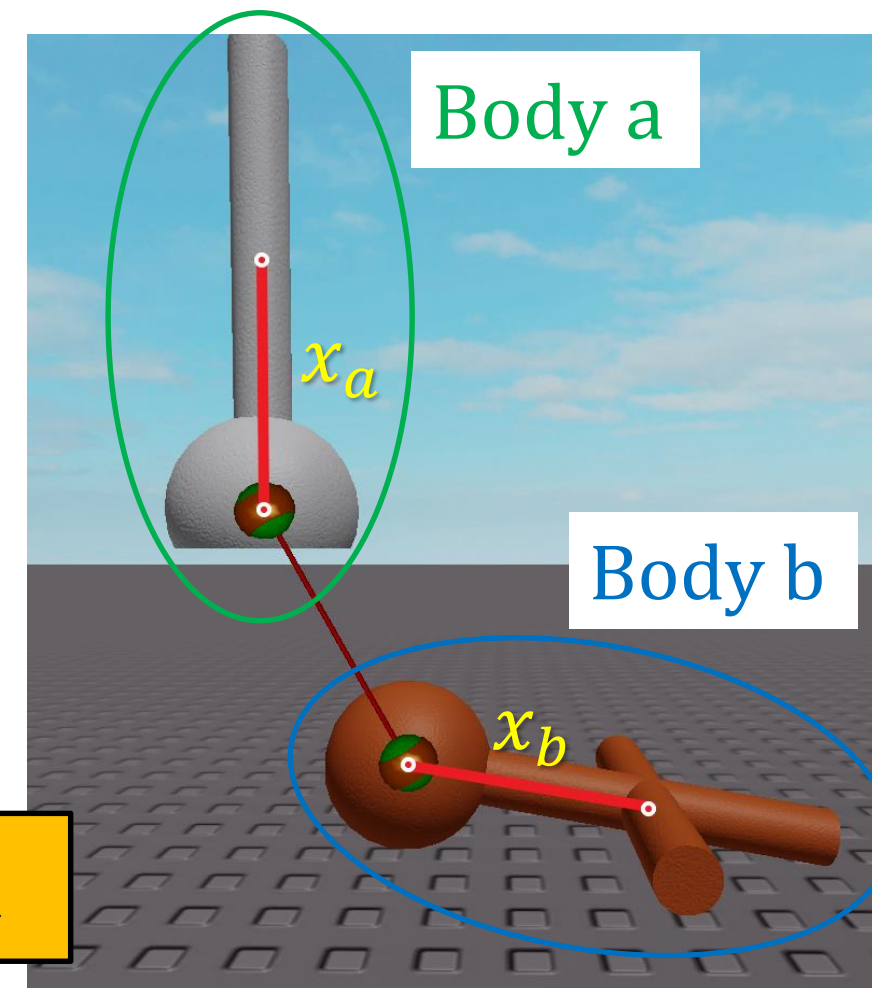
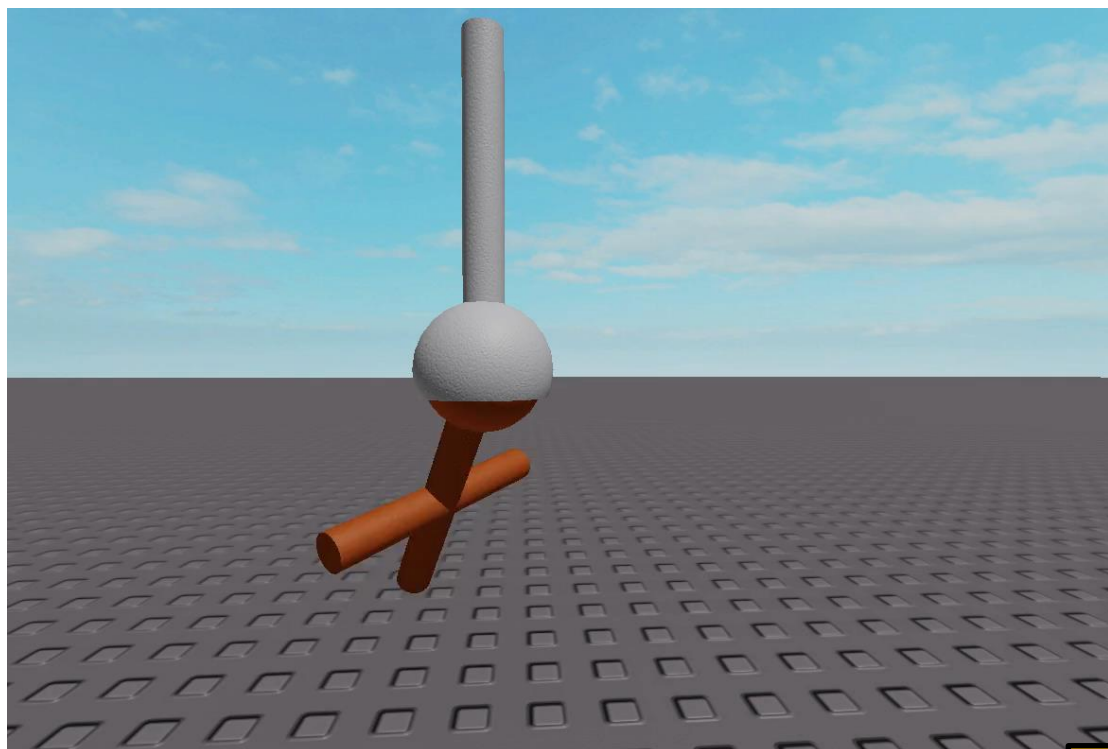
$x_a$  in world

$x_b$  in world

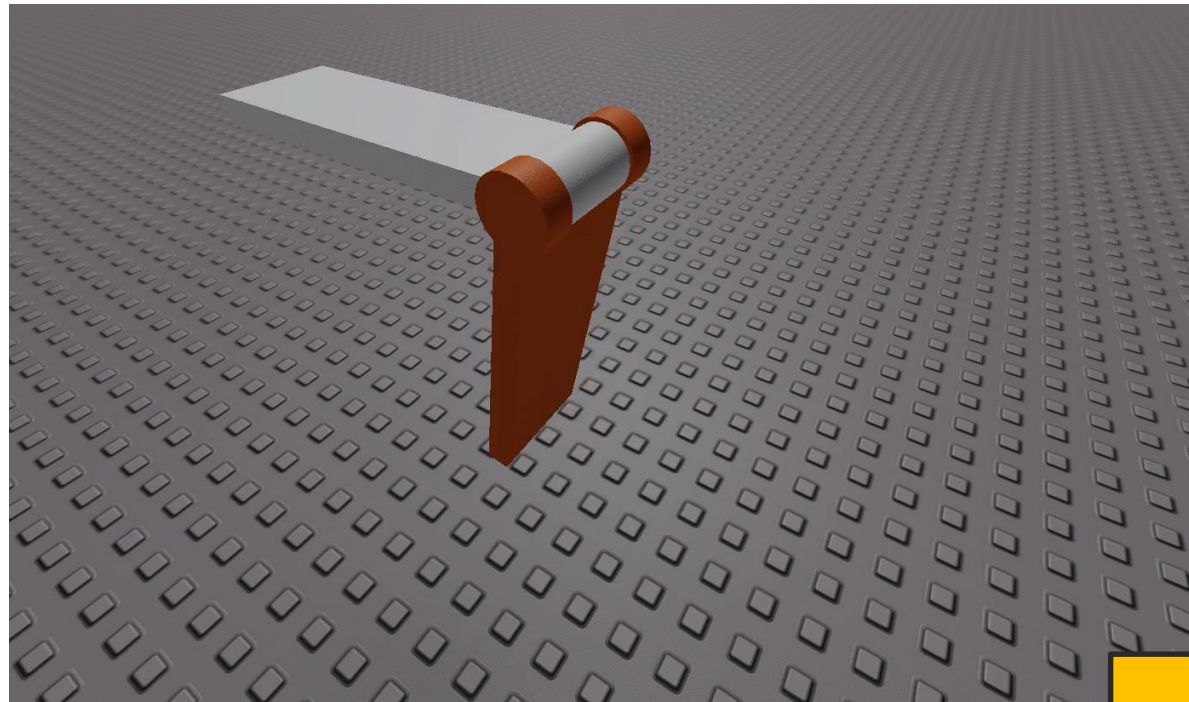
Constraint Function:

$$c(\Gamma_a, \Gamma_b) = (\underbrace{p_a + o_a x_a}_{\text{Same Location in World Space}}) - (\underbrace{p_b + o_b x_b}_{\text{Same Location in World Space}}) \in \mathbb{R}^3$$

Same Location in  
World Space







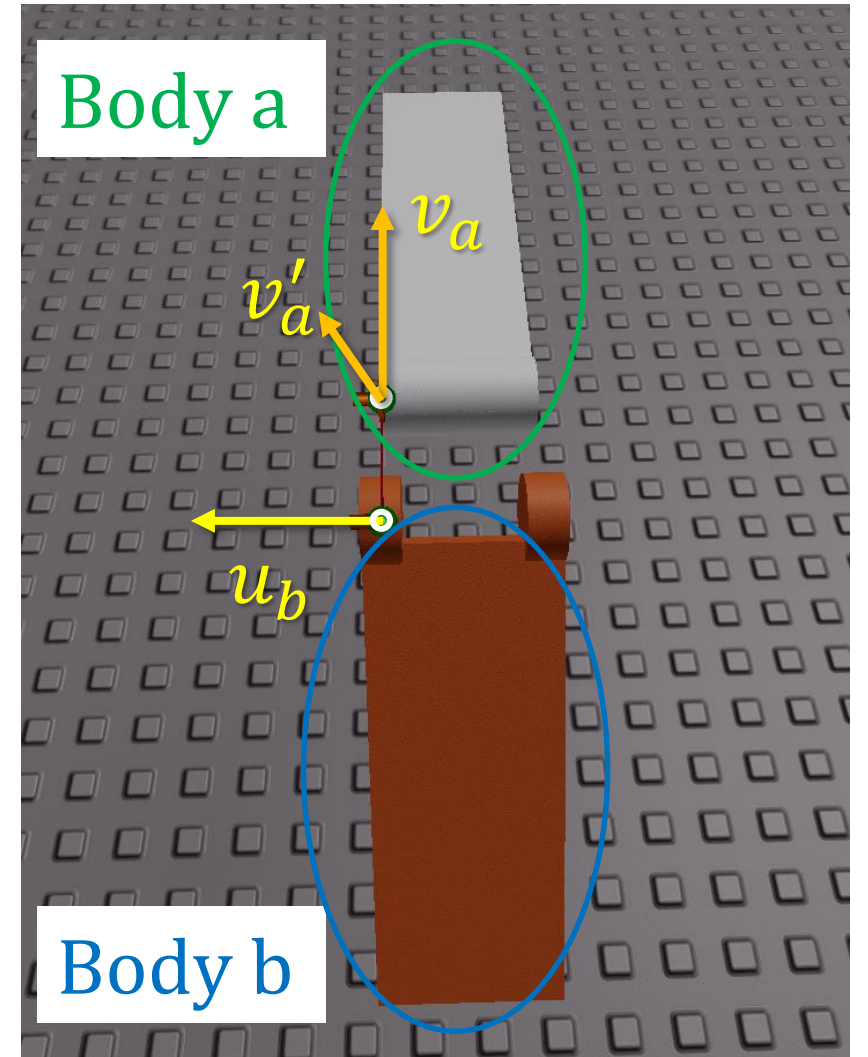
# Hinge

Ball in Socket

$u_b \perp v_a, v'_a$   
in world space

Body a

Body b



Constraint Function:

$$c(\Gamma_a, \Gamma_b) = \begin{bmatrix} (p_a + o_a x_a) - (p_b + o_b x_b) \\ (o_a v_a) \cdot (o_b u_b) \\ (o_a v'_a) \cdot (o_b u_b) \end{bmatrix} \in \mathbb{R}^5$$

# Velocity Space Constraint: Jacobian

Take the derivative of the constraint:

$$\frac{d}{dt} c(\Gamma_a, \Gamma_b, \dots) = \mathcal{J}_c \cdot \begin{bmatrix} \mathcal{V}_a \\ \mathcal{V}_b \\ \vdots \end{bmatrix}$$

Jacobian:  $d \times 6n$   
matrix

$$\mathcal{V}_\beta = \begin{bmatrix} v_\beta \\ \omega_\beta \end{bmatrix}$$

Linear Velocity

Angular Velocity

## Jacobian:

Body Space  
Velocities

Jacobian

Constraint Space  
Velocities

$$\mathcal{C}(\Gamma_a, \Gamma_b, \dots) = 0 \quad \Rightarrow \quad \mathcal{J}_c \begin{bmatrix} \mathcal{V}_a \\ \mathcal{V}_b \\ \vdots \end{bmatrix} = 0$$

No Motion In  
Constraint Space



# Jacobian: Ball in Socket

$$c(\Gamma_a, \Gamma_b) = (p_a + o_a x_a) - (p_b + o_b x_b)$$

$$\frac{d}{dt}c = (v_a - o_a x_a \times \omega_a) - (v_b - o_b x_b \times \omega_b)$$

$$= \underbrace{\begin{bmatrix} Id & -(o_a x_a)^\times \\ & J_c^{(a)} \end{bmatrix}}_{J_c^{(a)}} - \underbrace{\begin{bmatrix} Id & (o_b x_b)^\times \\ & J_c^{(b)} \end{bmatrix}}_{J_c^{(b)}} \begin{bmatrix} v_a \\ \omega_a \\ v_b \\ \omega_b \end{bmatrix}$$

$$\underbrace{\hspace{10em}}_{J_c}$$

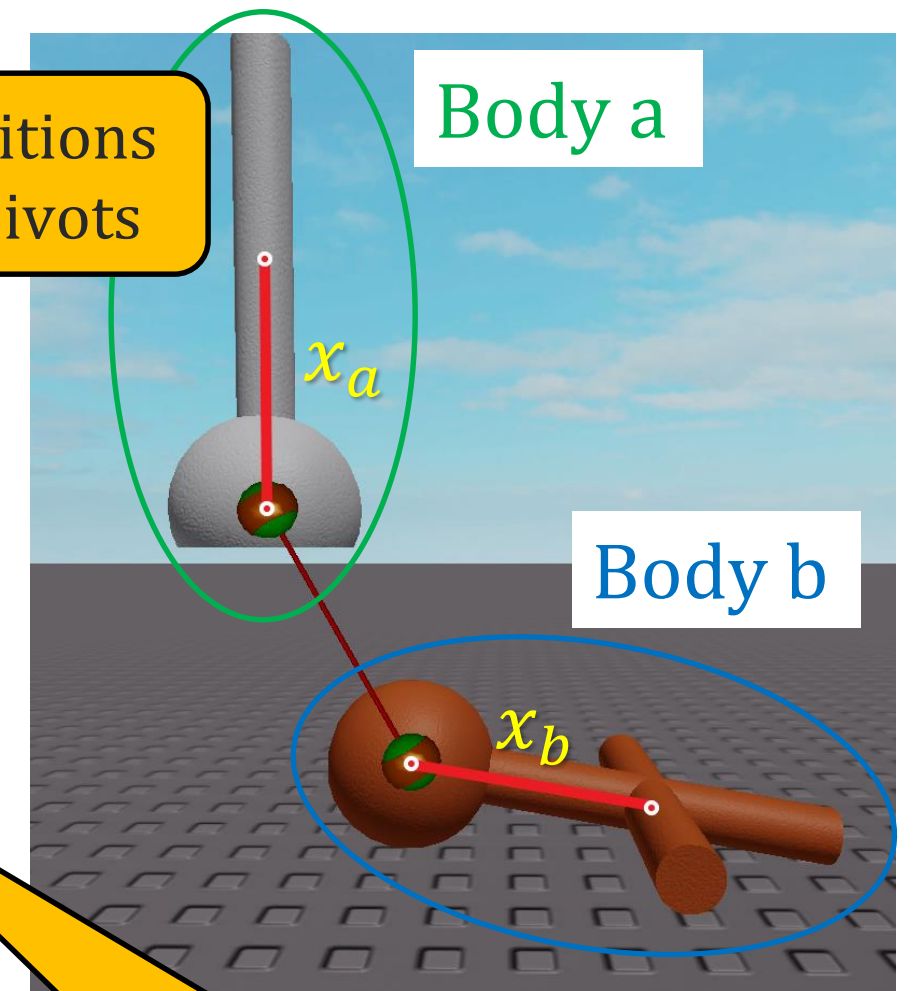
Relative positions  
of the two pivots

Body a

Body b

Relative velocity of  
the two pivots

$v^\times$ : cross-product  
matrix



# Mechanism

It's a collection of Constraints and Bodies:

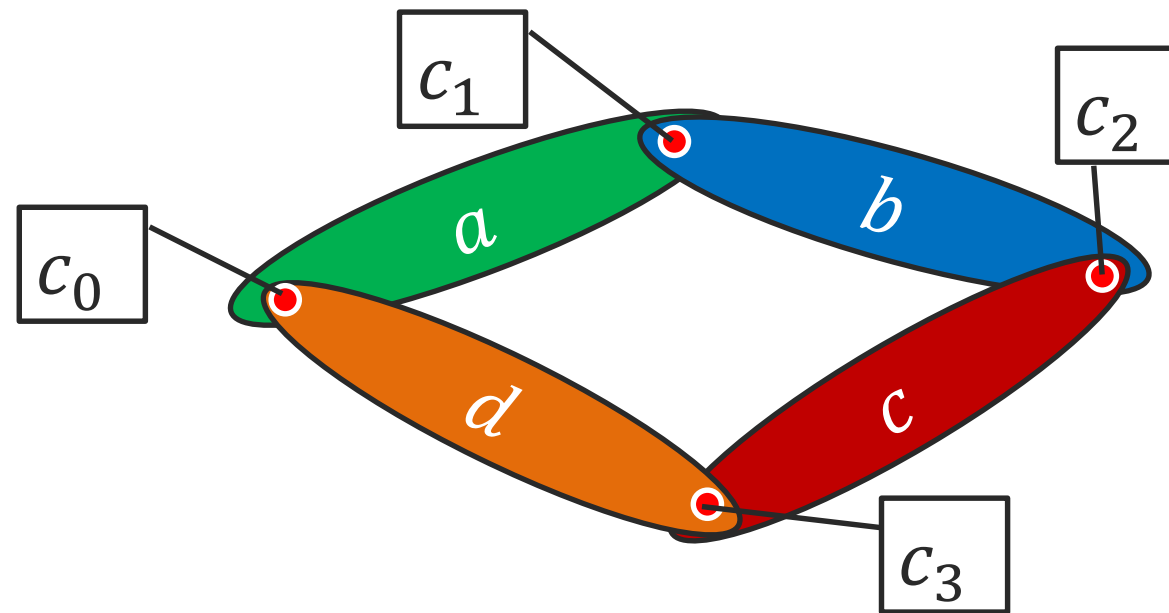
Global Constraint

$$\mathcal{C} = \begin{bmatrix} c_0 \\ c_1 \\ \vdots \end{bmatrix}, \quad \mathcal{B} = [a, b, \dots]$$

Each constraint:  
2 or 3 bodies

$$\frac{d}{dt} \mathcal{C}(\Gamma_a, \Gamma_b, \dots) = \begin{bmatrix} \mathcal{J}_0 \\ \mathcal{J}_1 \\ \vdots \end{bmatrix} \begin{bmatrix} \mathcal{V}_a \\ \mathcal{V}_b \\ \vdots \end{bmatrix} = \mathcal{J}_c \mathcal{V}_B$$

Each row: 12 or 18  
non-zero entries



$$\mathcal{J}_c = \begin{bmatrix} \underbrace{J_0^{(a)}}_a & 0 & 0 & \underbrace{J_0^{(d)}}_d \\ \underbrace{J_1^{(a)}}_a & \underbrace{J_1^{(b)}}_b & 0 & 0 \\ 0 & \underbrace{J_2^{(b)}}_b & \underbrace{J_2^{(c)}}_c & 0 \\ 0 & 0 & \underbrace{J_3^{(c)}}_c & \underbrace{J_3^{(d)}}_d \end{bmatrix} \begin{matrix} \left. \begin{matrix} \\ \\ \\ \end{matrix} \right\} c_0 \\ \left. \begin{matrix} \\ \\ \\ \end{matrix} \right\} c_1 \\ \left. \begin{matrix} \\ \\ \\ \end{matrix} \right\} c_2 \\ \left. \begin{matrix} \\ \\ \\ \end{matrix} \right\} c_3 \end{matrix}$$



# Euler Integration Step

$\mathcal{V}_0 :=$  Velocities at  $t = 0$

Velocities next time step:

$$\mathcal{V}_1 = \mathcal{V}_0 + \mathcal{W} \cdot \text{External Forces} \cdot dt + \mathcal{W} \mathcal{J}^t \Lambda$$

$$\mathcal{W} = \begin{bmatrix} m_0^{-1} Id & & & \\ & I_0^{-1} & & \\ & & \ddots & \\ & & & m_{n-1}^{-1} Id \\ & & & & I_{n-1}^{-1} \end{bmatrix}$$

D'Alembert's Principle

**Solver:** finds  $\Lambda$  such that:

$$\mathcal{J} \mathcal{V}_1 = 0$$

Lagrangian Multipliers  
aka Impulses

# Linearized Constraint Equation

Valid only for  
equality constraints

$$\mathcal{K}\Lambda = \mathcal{R}$$

$$\mathcal{K} := J\mathcal{W}J^t$$

- Symmetric
- Positive semidefinite

$$\mathcal{R} := -J(\dots)$$

# PGS Solver

Collisions and  
Friction

Inequality  
Constraints

Equality  
Constraints



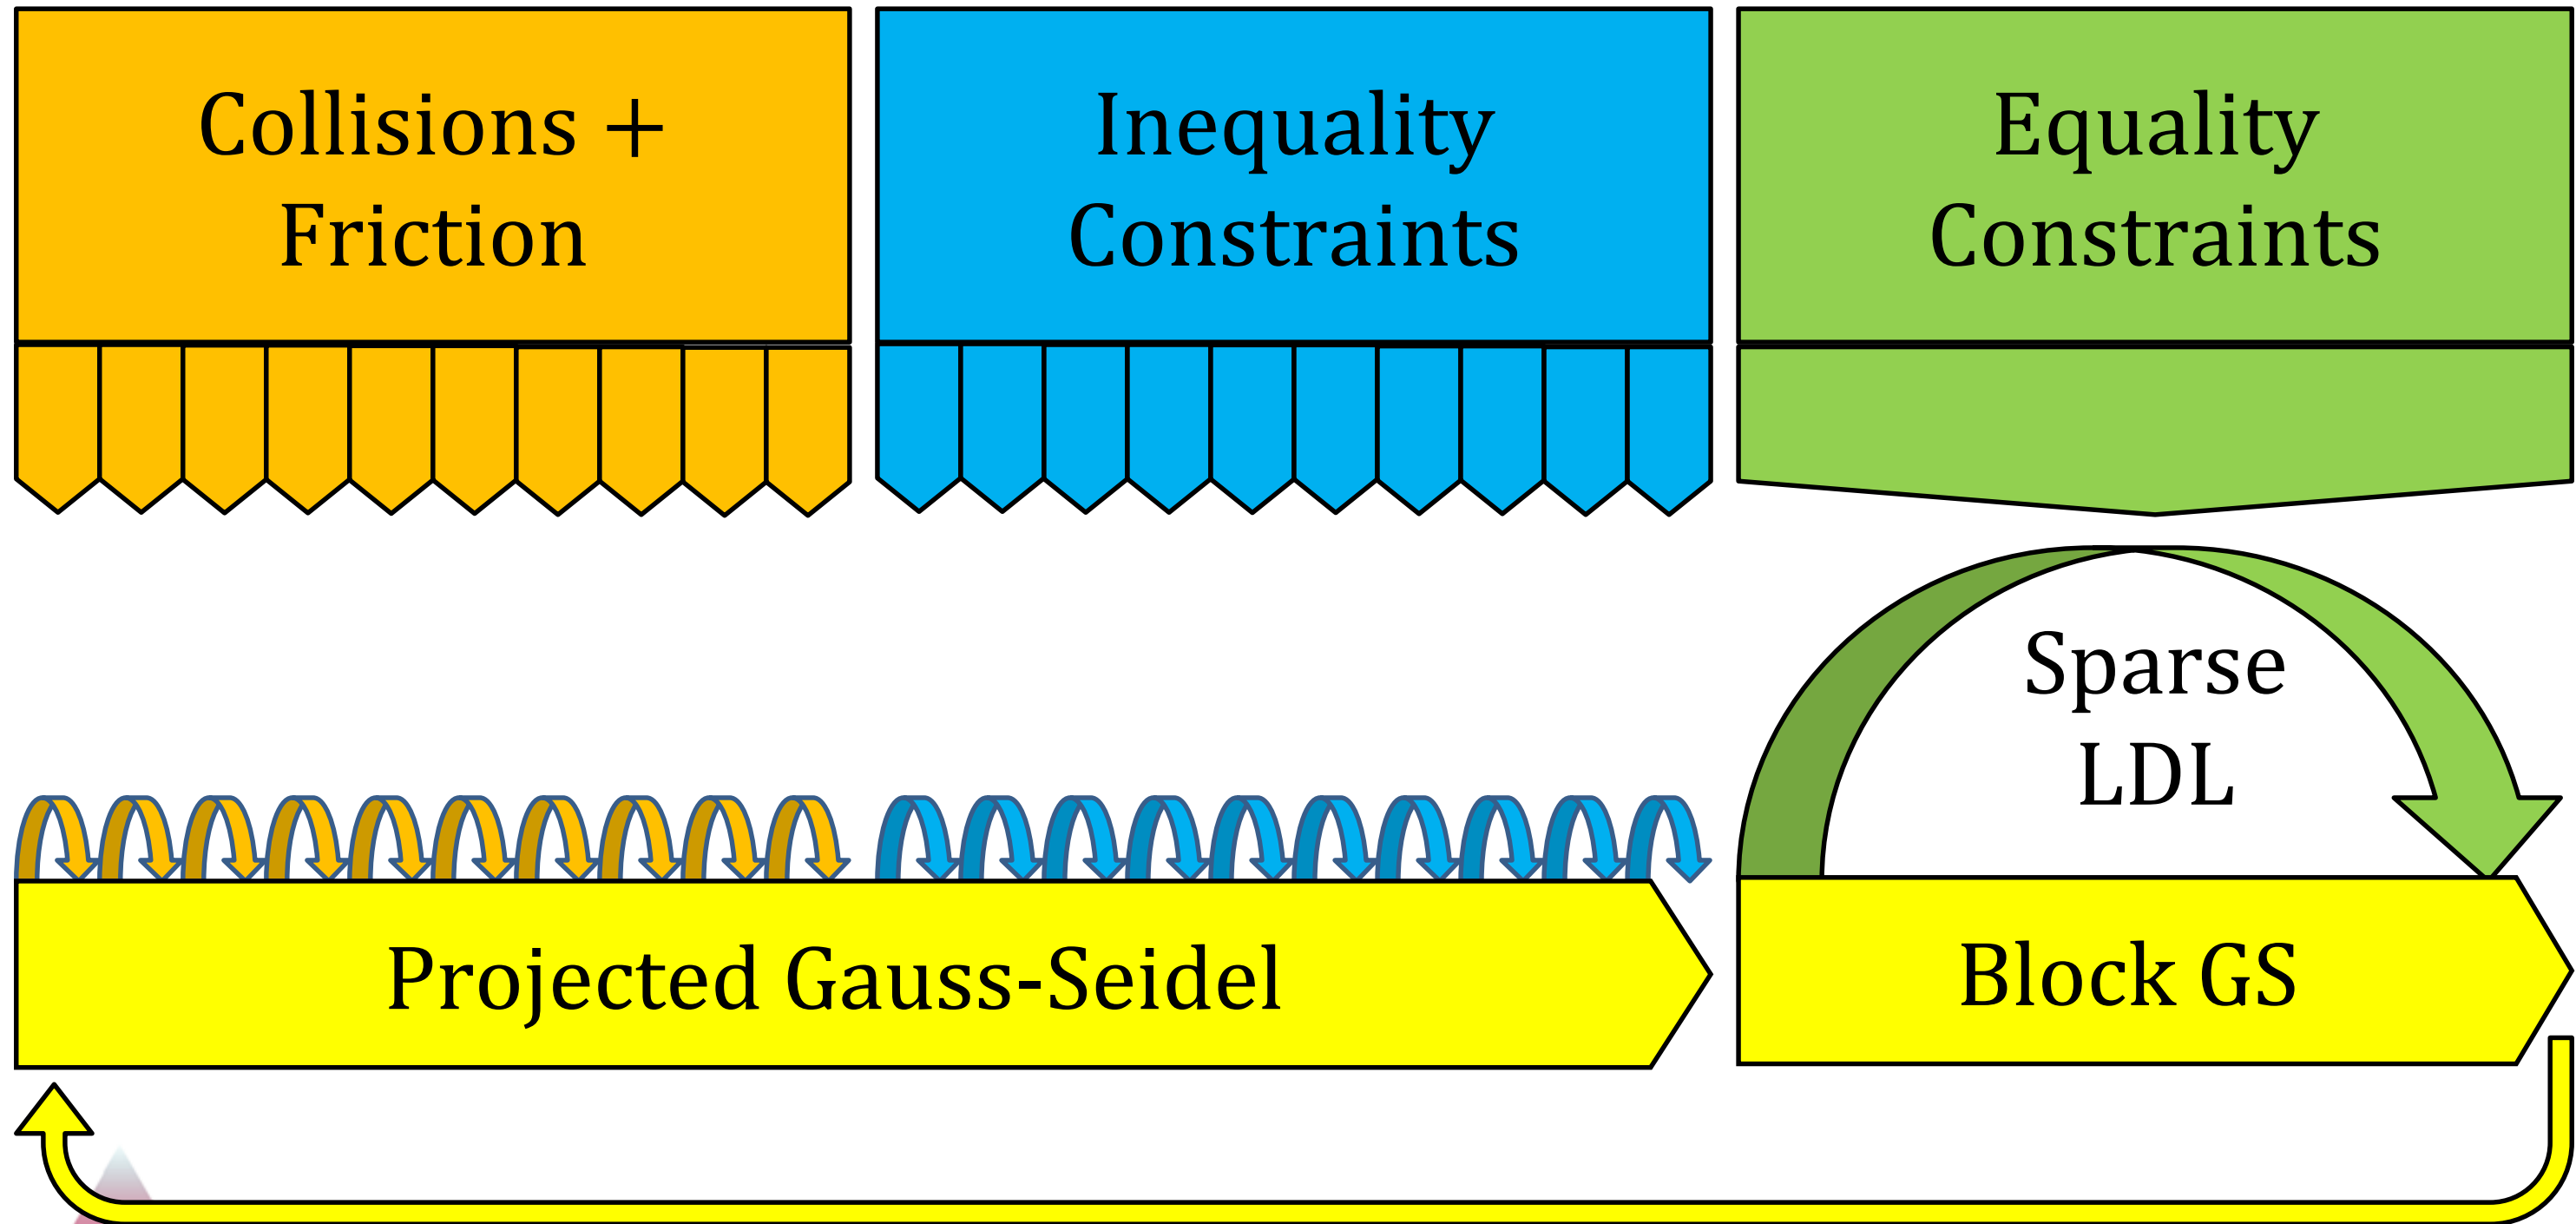
Projected Gauss-Seidel

Gauss-Seidel

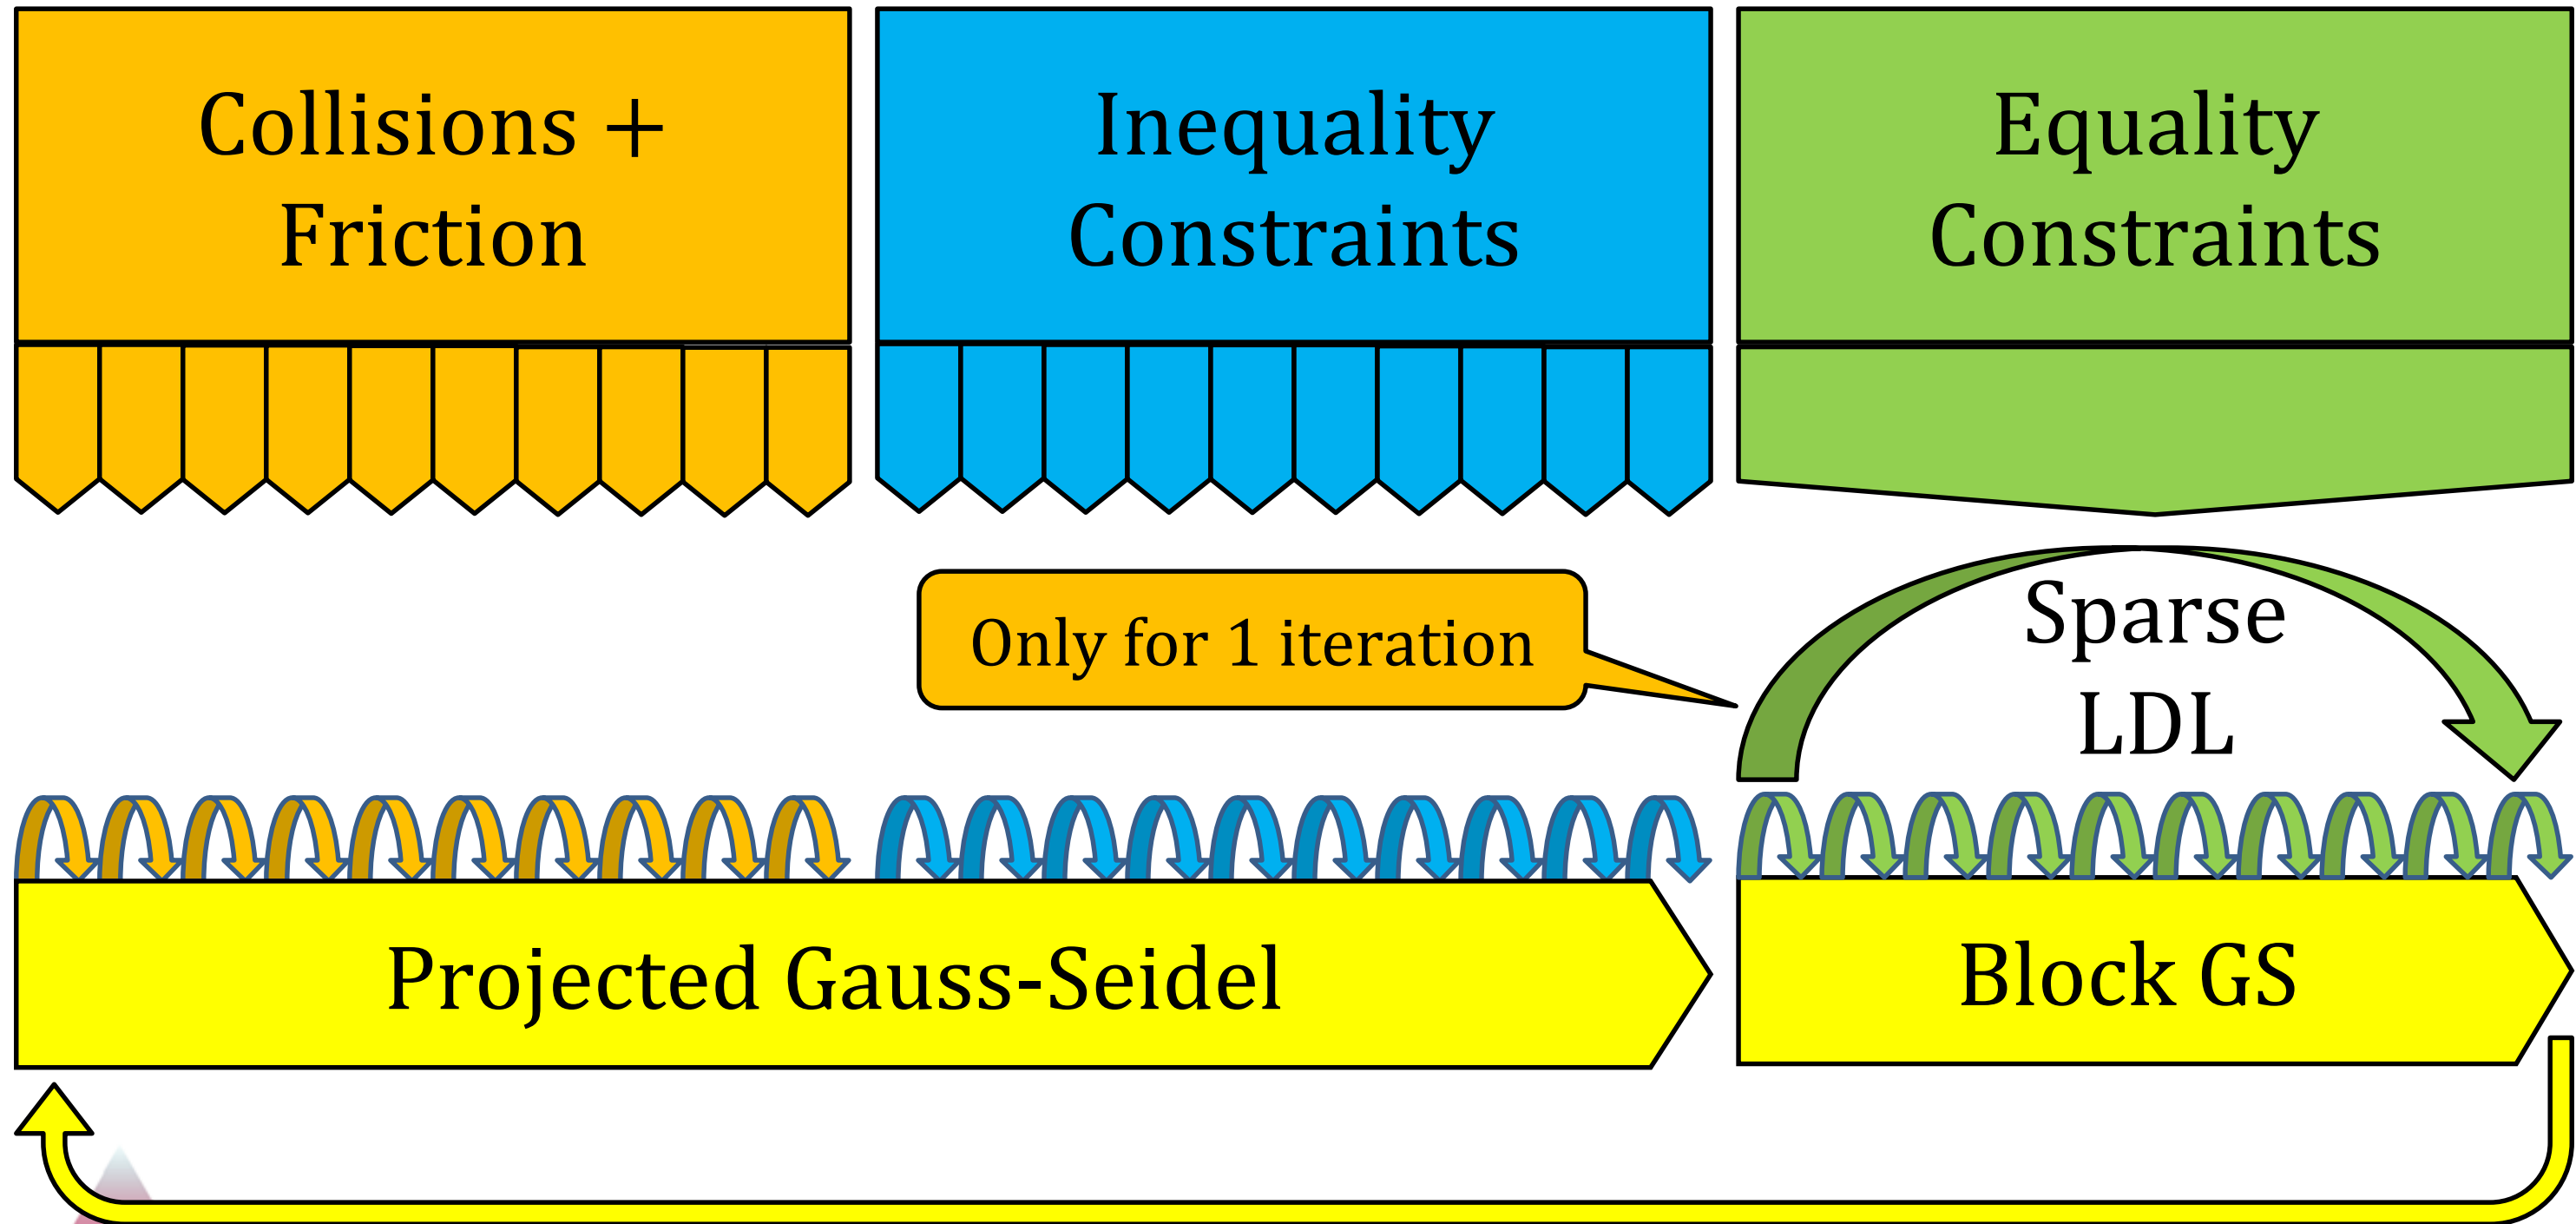




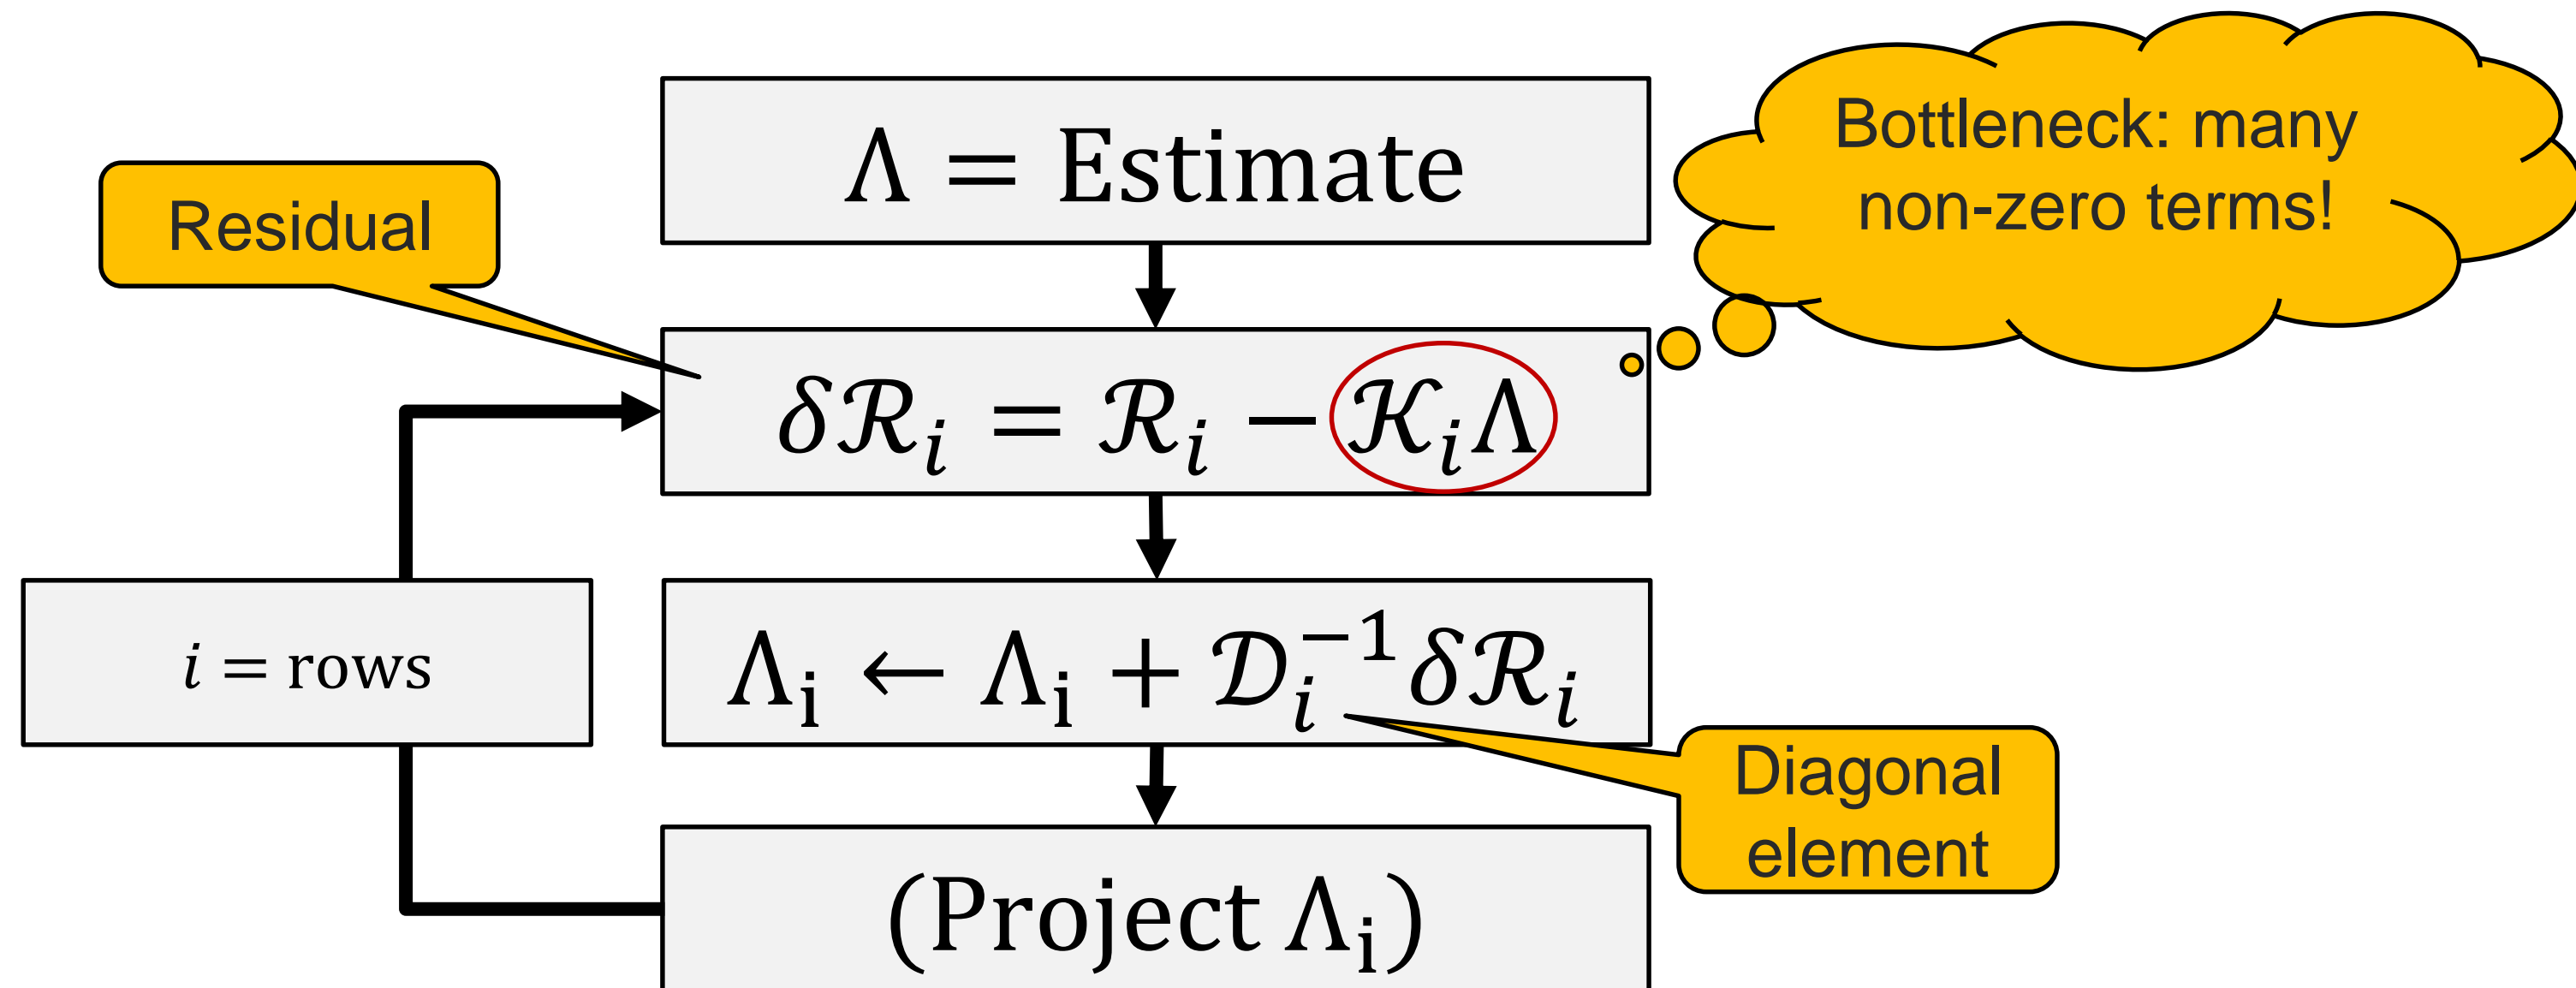
# LDL-PGS Solver



# LDL-PGS Solver



# Solving $\mathcal{K}\Lambda = \mathcal{R}$ : (Projected) Gauss-Seidel



# Faster Calculation of $\mathcal{K}_i \Lambda$

Constraint Equation:

$$\mathcal{K} \Lambda = \mathcal{R}$$

$$(\mathcal{J} \mathcal{W} \mathcal{J}^t) \Lambda = \mathcal{R}$$

Associativity of matrix multiplication

$$\mathcal{J}(\underbrace{\mathcal{W} \mathcal{J}^t \Lambda}_{\delta \mathcal{V}}) = \mathcal{R}$$

Velocity Changes

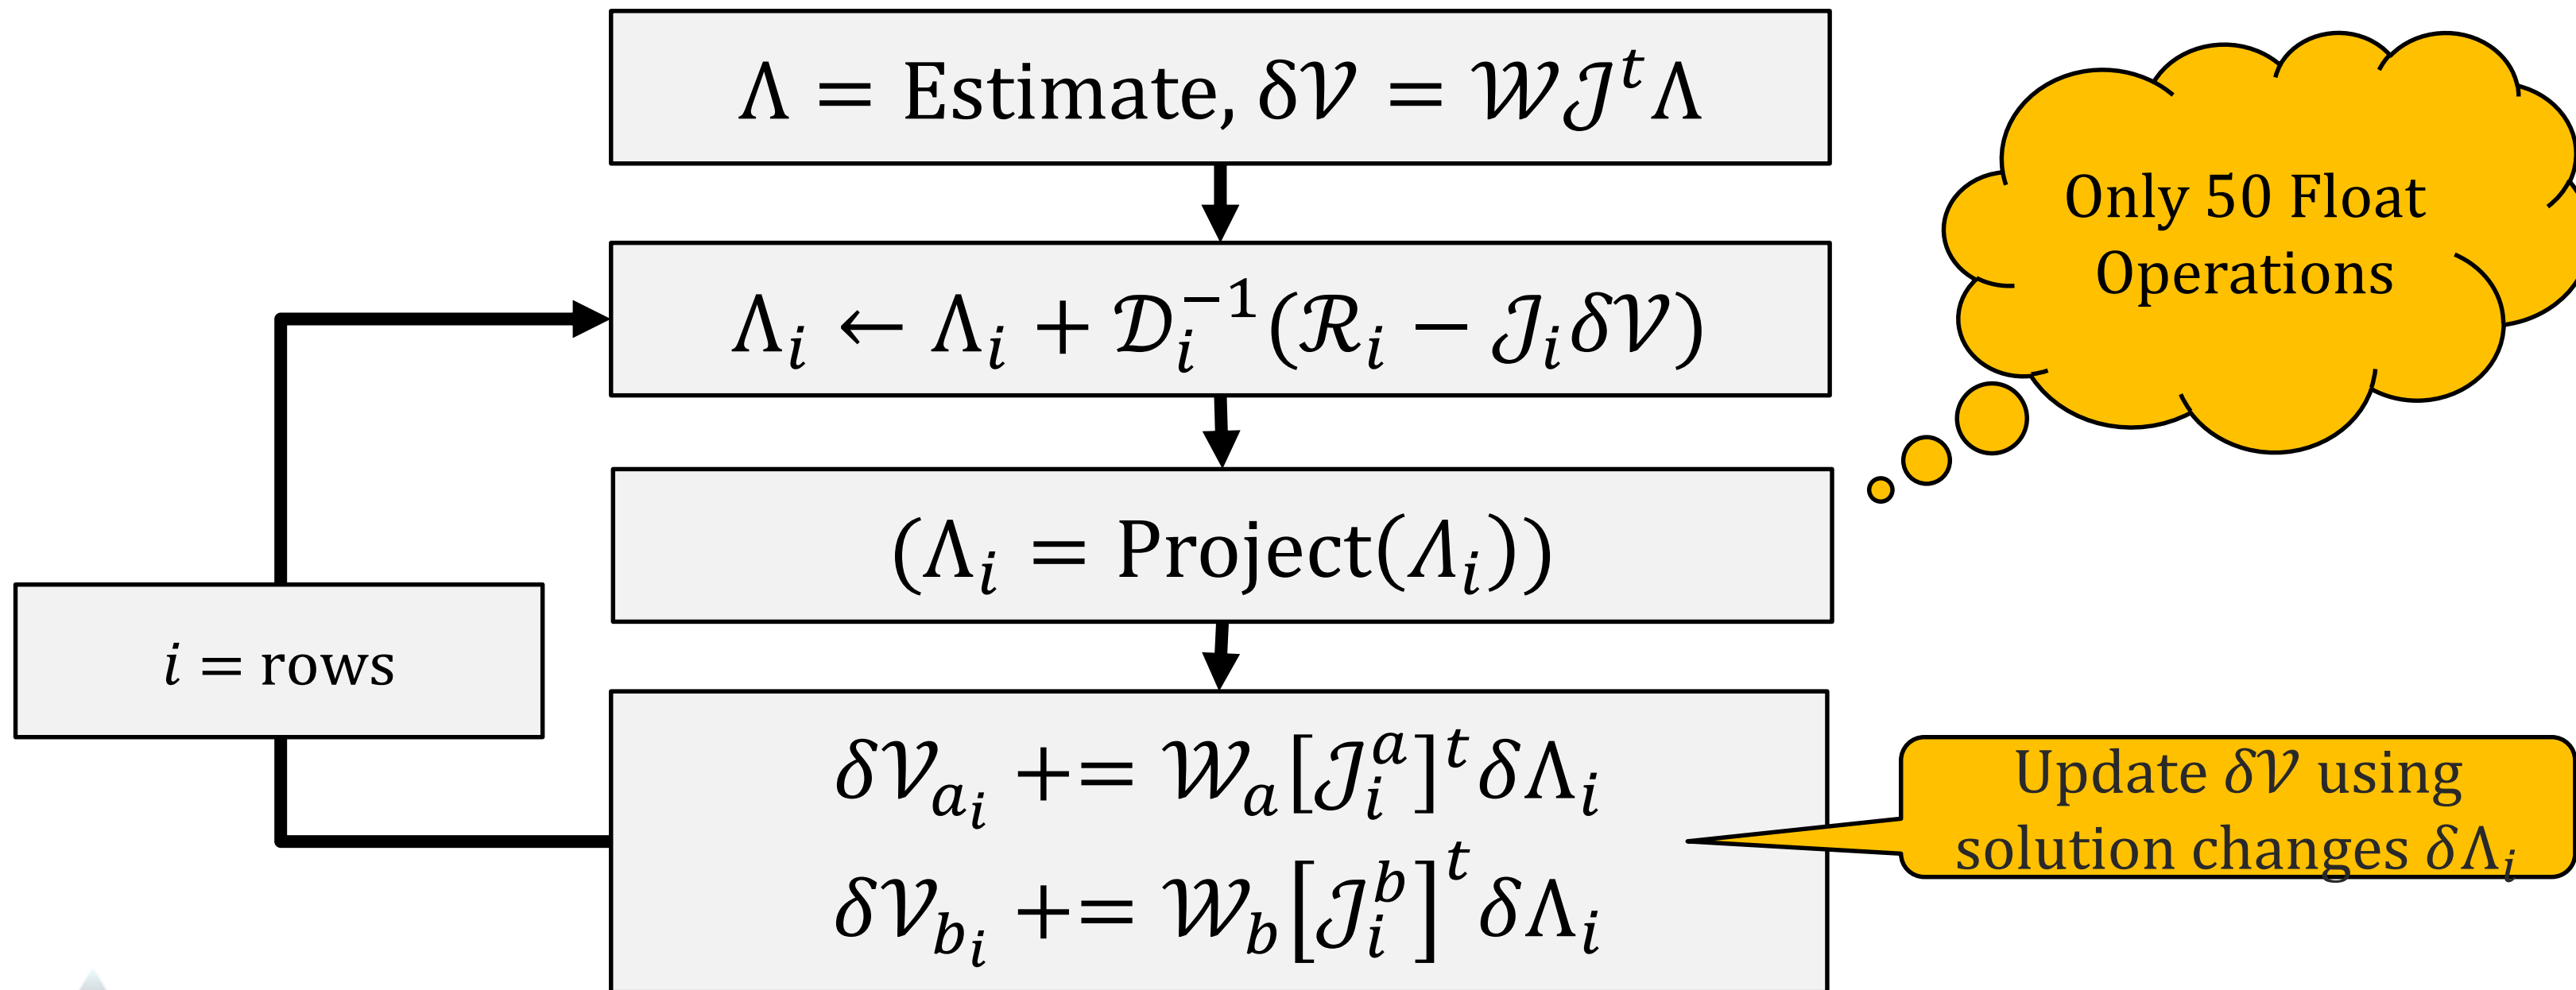
In Gauss-Seidel replace:

$$\mathcal{K}_i \Lambda \leftrightarrow \mathcal{J}_i \delta \mathcal{V}$$

Only 12 (or 18) Terms!



# Impulse Solver (Kaczmarz Method)



# Block Structure

Row partition:

$$\pi = \{\pi_0, \pi_1, \dots\}$$

$$\pi_i = \{\pi_{i,0}, \pi_{i,1}, \dots\}$$

List of rows for each  $i$

$$\mathcal{J}_\pi = \begin{bmatrix} \mathcal{J}_{\pi_0} \\ \mathcal{J}_{\pi_1} \\ \vdots \end{bmatrix}, \quad \mathcal{J}_{\pi_i} = \begin{bmatrix} \mathcal{J}_{\pi_{i,0}} \\ \mathcal{J}_{\pi_{i,1}} \\ \vdots \end{bmatrix}$$

# Partitioned Constraint Matrix

Symmetric Matrices

Transposed

$$\mathcal{K} = \begin{bmatrix} \mathcal{J}_{\pi_0} \mathcal{W} \mathcal{J}_{\pi_0}^t & \mathcal{J}_{\pi_0} \mathcal{W} \mathcal{J}_{\pi_1}^t & \cdots \\ \mathcal{J}_{\pi_1} \mathcal{W} \mathcal{J}_{\pi_0}^t & \mathcal{J}_{\pi_1} \mathcal{W} \mathcal{J}_{\pi_1}^t & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix}$$

$\mathcal{N}$  for “Node” Matrix

$$= \begin{bmatrix} \mathcal{N}_0 & \mathcal{E}_{01} & \cdots \\ \mathcal{E}_{10} & \mathcal{N}_1 & \cdots \\ \vdots & \vdots & \ddots \end{bmatrix}$$

$\mathcal{E}$  for “Edge” Matrix

# Block Impulse Solver

$$\Lambda = \text{Estimate}, \delta\mathcal{V} = \mathcal{W} \mathcal{J}^t \Lambda$$

Inverse of a square matrix

$$\Lambda_{\pi_i} \leftarrow \Lambda_{\pi_i} + \mathcal{N}_i^{-1} (\mathcal{R}_{\pi_i} - \mathcal{J}_{\pi_i} \delta\mathcal{V})$$

Block solver doesn't support projection.

$i$  = elements of partition

$$\delta\mathcal{V} += \mathcal{W} \mathcal{J}_{\pi_i}^t \delta\Lambda_{\pi_i}$$



We need to solve:

$$\delta\Lambda_{\pi_i} = \mathcal{N}_i^{-1} \delta r$$

Small dimensions  $\Rightarrow$  easily invert  $\mathcal{N}_i$

Ex: natural partition based on constraints:

- a ball-in-socket (dim = 3x3)
- a hinge (dim = 5x5)
- a cylindrical (dim = 4x4)
- ...

$\mathcal{N}_i$  are invertible because constraints are **regular**

Idea: group all equality constraints together

$\pi_0$  : rows of all equality constraints

$\pi_1, \pi_2, \dots$  : other individual rows

H for holonomic  $\mathcal{H} := \mathcal{N}_0 = \mathcal{J}_{\pi_0} \mathcal{W} \mathcal{J}_{\pi_0}^t$

constraint matrix of equality constraints.

# How to evaluate?

$$\mathcal{H}^{-1}(\dots)$$

Potential issues with  $\mathcal{H}$ :

1. Not invertible
2. Large dimensions
3. Large dense submatrices

Solutions:

1. Regularize: add compliance
2. Use a sparse method
3. Increase sparsity: split bodies

# Compliance and Regularization

$\mathcal{H}$  is symmetric positive semidefinite.

It may be not be invertible, but

$$\tilde{\mathcal{H}} := \mathcal{H} + \begin{bmatrix} \epsilon_0 & & \\ & \epsilon_1 & \\ & & \ddots \end{bmatrix}$$

$\epsilon_i$  are not unit free,  
scale dependent

Is invertible for any  $\epsilon_0, \epsilon_1, \dots > 0$

Use a scaled diagonal:

$$\epsilon_i = \epsilon \cdot \mathcal{H}_{i,i} \text{ for small } \epsilon > 0$$

Adds compliance to constraints!



# Compliance and Regularization

Traditional PGS engines also use compliance to stabilize solutions

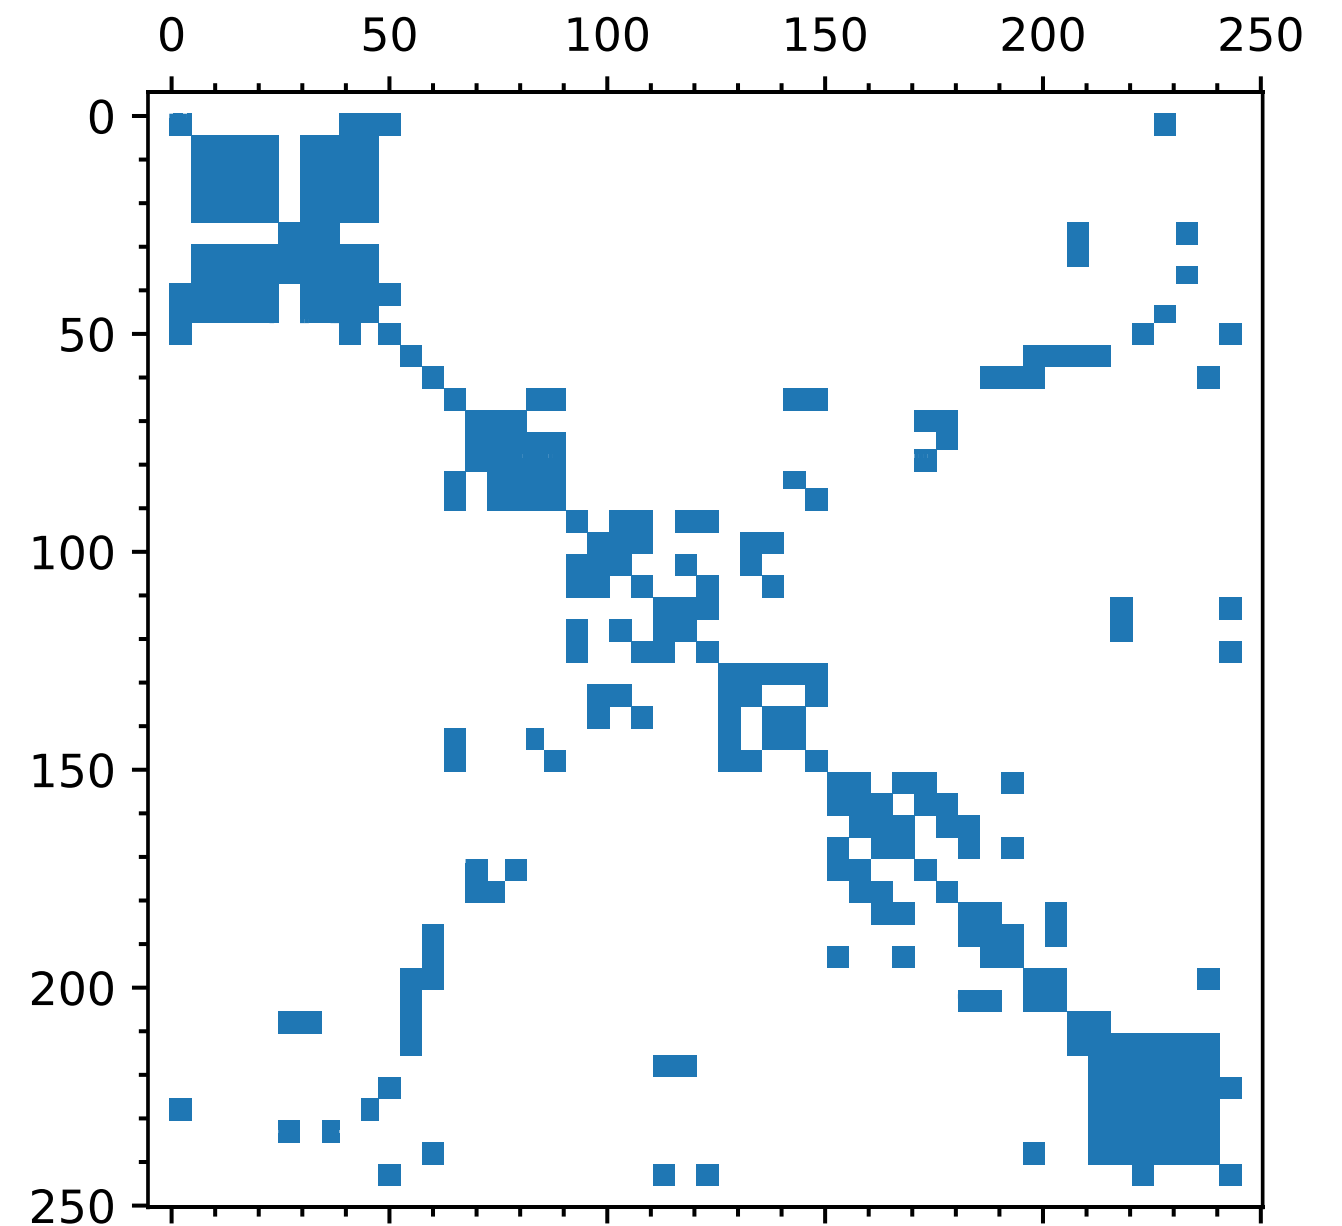
PGS for  $\tilde{\mathcal{H}}\Lambda = \mathcal{R}$



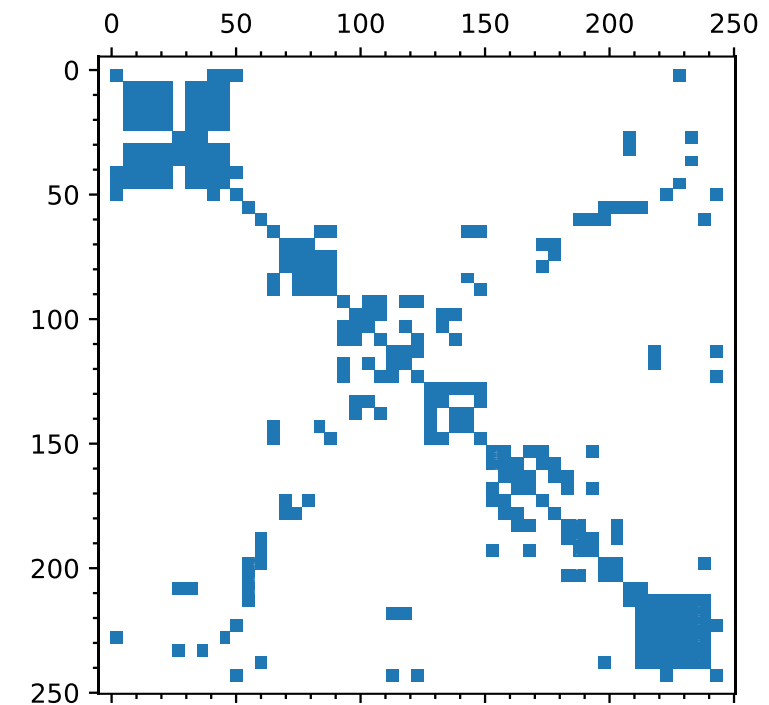
PGS for  $\mathcal{H}\Lambda = \mathcal{R}$   
with  $\Lambda_i \leftarrow (1 + \epsilon_i)^{-1} \Lambda_i$   
for each row

CFM in Bullet/ODE  
(Constraint Force  
Mixing)

# Sparse Methods



# Sparse Methods



$\#Constraints = 50$

$Dim(\mathcal{H}) = 246$

$Density(\mathcal{H}) = 12\%$

Sparse Cholesky LDL:

$Flops = 90k$

$Performance (Core i9) = 70\mu s$

# Cholesky LDL Decomposition

Any Symmetric Positive Definite matrix can be decomposed as

$$\mathcal{A} = \mathcal{L}\mathcal{D}\mathcal{L}^t$$

Where

$\mathcal{L}$  – lower triangular with 1s on diagonal

$\mathcal{D}$  – diagonal

This is useful because both  $\mathcal{L}^{-1}$  and  $\mathcal{D}^{-1}$  can be efficiently evaluated.



Useful to compute  $\mathcal{A}^{-1}v$ :

$$\begin{aligned}\mathcal{A}^{-1}v &= (\mathcal{L}^{-t}\mathcal{D}^{-1}\mathcal{L}^{-1})v \\ &= \mathcal{L}^{-t}(\mathcal{D}^{-1}(\mathcal{L}^{-1}v))\end{aligned}$$

Where

$$\mathcal{L}^{-1}(\cdot), \mathcal{L}^{-t}(\cdot)$$

Computed by back-substitution and

$$\mathcal{D}^{-1}(\cdot)$$

Are scalar products.

# Block LDL

Suppose  $\mathcal{A}$  is SPD and has a block structure:

$$\mathcal{A} = \begin{bmatrix} \begin{bmatrix} * & * \end{bmatrix} & \begin{bmatrix} * & * & * \end{bmatrix} & \begin{bmatrix} * \end{bmatrix} \\ \begin{bmatrix} * & * \end{bmatrix} & \begin{bmatrix} * & * & * \end{bmatrix} & \begin{bmatrix} * \end{bmatrix} \\ \begin{bmatrix} * & * \end{bmatrix} & \begin{bmatrix} * & * & * \end{bmatrix} & \begin{bmatrix} * \end{bmatrix} \\ \begin{bmatrix} * & * \end{bmatrix} & \begin{bmatrix} * & * & * \end{bmatrix} & \begin{bmatrix} * \end{bmatrix} \end{bmatrix}$$

# Block LDL

It can be decomposed as

$$\mathcal{A} = \mathcal{L}\mathcal{D}\mathcal{L}^t$$

$$\mathcal{D} = \begin{bmatrix} \begin{bmatrix} * & * \\ * & * \end{bmatrix} & & \\ & \begin{bmatrix} * & * & * \\ * & * & * \\ * & * & * \end{bmatrix} & \\ & & [*] \end{bmatrix}, \quad \mathcal{L} = \begin{bmatrix} \begin{bmatrix} 1 & & \\ & 1 & \\ & & 1 \end{bmatrix} & & \\ & \begin{bmatrix} * & * \\ * & * \\ * & * \end{bmatrix} & & \\ & & \begin{bmatrix} 1 & & \\ & 1 & \\ & & 1 \end{bmatrix} & \\ & & & \begin{bmatrix} * & * & * \\ & 1 & \end{bmatrix} \end{bmatrix}$$

$\mathcal{D}$  and  $\mathcal{L}$  inherit the block structure from  $\mathcal{A}$ .

Block LDL vs LDL?

**Performance: block operations faster than scalar!**

# LDL: Algorithms

There are at least 2 algorithms:

- Gaussian Elimination ← we'll use this one
- Doolittle Algorithm

# Schur Complement

Multiply first row by  $\varepsilon \mathcal{N}^{-1}$  on the left

$$\begin{bmatrix} \mathcal{N} & \varepsilon^t \\ \varepsilon & \mathcal{S} \end{bmatrix}$$

Multiply first column by  $\mathcal{N}^{-1} \varepsilon^t$  on the right

Subtract from second row

$$\begin{bmatrix} \mathcal{N} & \varepsilon^t \\ \textcircled{0} & \mathcal{S} - \varepsilon \mathcal{N}^{-1} \varepsilon^t \end{bmatrix}$$

Subtract from second column

Eliminated

$$\begin{bmatrix} \mathcal{N} & \textcircled{0} \\ 0 & \mathcal{S} - \varepsilon \mathcal{N}^{-1} \varepsilon^t \end{bmatrix}$$

Eliminated

# Schur Complement

$$\begin{bmatrix} \mathcal{N} & \varepsilon^t \\ \varepsilon & \mathcal{S} \end{bmatrix}$$

$$= \underbrace{\begin{bmatrix} Id & \\ \varepsilon \mathcal{N}^{-1} & Id \end{bmatrix}}_{\text{Lower Triangular}} \underbrace{\begin{bmatrix} \mathcal{N} & 0 \\ 0 & \mathcal{S} - \varepsilon \mathcal{N}^{-1} \varepsilon^t \end{bmatrix}}_{\text{Block Diagonal}} \underbrace{\begin{bmatrix} Id & \mathcal{N}^{-1} \varepsilon^t \\ & Id \end{bmatrix}}_{\text{Upper Triangular}}$$

Block Gaussian Elimination = Recursive Schur complements



# Block Gaussian Elimination

$$\begin{array}{c}
 \boxed{\varepsilon_{*,0}} \quad \left[ \begin{array}{c|ccc}
 \mathcal{N}_0 & \varepsilon_{10}^t & \varepsilon_{20}^t & \dots \\
 \hline
 \varepsilon_{10} & \mathcal{N}_1 & \varepsilon_{21}^t & \dots \\
 \varepsilon_{20} & \varepsilon_{21} & \mathcal{N}_2 & \dots \\
 \vdots & \vdots & \vdots & \ddots
 \end{array} \right] \quad \begin{array}{l} \boxed{\varepsilon_{*,0}^t} \\ \boxed{\mathcal{S}_0} \end{array}
 \end{array}$$

$$= \begin{bmatrix} \mathcal{N}_0 & \varepsilon_{*,0}^t \\ \varepsilon_{*,0} & \mathcal{S}_0 \end{bmatrix}$$

$$= \begin{bmatrix} Id & \\ \varepsilon_{*,0} \mathcal{N}_0^{-1} & Id \end{bmatrix} \begin{bmatrix} \mathcal{N}_0 & 0 \\ 0 & \mathcal{S}_0 - \varepsilon_{*,0} \mathcal{N}_0^{-1} \varepsilon_{*,0}^t \end{bmatrix} \begin{bmatrix} Id & \mathcal{N}_0^{-1} \varepsilon_{*,0} \\ & Id \end{bmatrix}$$

# Block Gaussian Elimination

Block  
Diagonals

Already  
Eliminated

$$\begin{bmatrix} \ddots & 0 & 0 \\ 0 & \mathcal{N}_i & \epsilon_{*,i}^t \\ 0 & \epsilon_{*,i} & \mathcal{S}_i \end{bmatrix}$$

$\mathcal{S}_{i-1}$

$\equiv$

$$\underbrace{\begin{bmatrix} Id & & \\ & Id & \\ \mathcal{L}_{*,i} & & Id \end{bmatrix}} \begin{bmatrix} \ddots & & \\ & \mathcal{N}_i & 0 \\ & 0 & \mathcal{S}_i - \epsilon_{*,i} \mathcal{L}_{*,i}^t \end{bmatrix} \underbrace{\begin{bmatrix} Id & & \\ & Id & \mathcal{L}_{*,i}^t \\ & & Id \end{bmatrix}}$$

$$[\mathcal{L}_{*,i}] \leftarrow \text{Elementary Matrices} \rightarrow [\mathcal{L}_{*,i}]^t$$

Schur Complement:

1. Invert  $\mathcal{N}_i$

2. Compute:

$$\mathcal{L}_{*,i} = \epsilon_{*,i} \mathcal{N}_i^{-1}$$

3. Reduce:

$$\mathcal{S}_i \leftarrow \mathcal{S}_i - \epsilon_{*,i} \mathcal{L}_{*,i}^t$$

# Block LDL

$$\mathcal{D} = \begin{bmatrix} \mathcal{N}_0 & & \\ & \ddots & \\ & & \mathcal{N}_{n-1} \end{bmatrix}$$

$$\mathcal{L} = [\mathcal{L}_{*,0}][\mathcal{L}_{*,1}] \cdots [\mathcal{L}_{*,n-2}] = \begin{bmatrix} 1 & & & & \\ \uparrow & 1 & & & \\ & \uparrow & \ddots & & \\ \mathcal{L}_{*,0} & \mathcal{L}_{*,1} & & \ddots & \\ \downarrow & \downarrow & & & \mathcal{L}_{*,n-2} & 1 \end{bmatrix}$$

$$\mathcal{H} = \mathcal{L} \mathcal{D} \mathcal{L}^t$$

# Inverse Operator

$$\mathcal{H}^{-1}v = \mathcal{L}^{-t}(\mathcal{D}^{-1}(\mathcal{L}^{-1}v))$$

$$\begin{aligned}\mathcal{L}^{-1} &= [-\mathcal{L}_{*,n-2}] \cdots [-\mathcal{L}_{*,1}] [-\mathcal{L}_{*,0}] \\ \mathcal{L}^{-t} &= [-\mathcal{L}_{*,0}]^t [-\mathcal{L}_{*,1}]^t \cdots [-\mathcal{L}_{*,n-2}]^t\end{aligned}$$

Don't evaluate  
these products.  
Use as is.

$$\mathcal{D}^{-1} = \begin{bmatrix} \mathcal{N}_0^{-1} & & \\ & \ddots & \\ & & \mathcal{N}_{n-1}^{-1} \end{bmatrix}$$

Computed  
during  
elimination

# Sparse Block LDL

Constraints

$$\mathcal{H} = \begin{bmatrix} \mathcal{N}_0 & \varepsilon_{10}^t & \varepsilon_{20}^t & \dots & 0 \\ \varepsilon_{10} & \mathcal{N}_1 & 0 & \dots & \varepsilon_{n-1,1}^t \\ \varepsilon_{20} & 0 & \mathcal{N}_2 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & \varepsilon_{n-1,1} & 0 & \dots & \mathcal{N}_{n-1} \end{bmatrix}$$

Where some  $\varepsilon_{ij}$ 's are 0.

Can we reduce the operation count in the Gaussian Elimination?

Yes absolutely! Here is how:

- Sparse matrices  $\Leftrightarrow$  Graphs
- Gaussian elimination  $\Leftrightarrow$  Process on graphs
- Pack the sparse matrix data in memory



## Body Graph

- Nodes: Rigid Bodies
- Edges: Constraints

## Constraint Graph

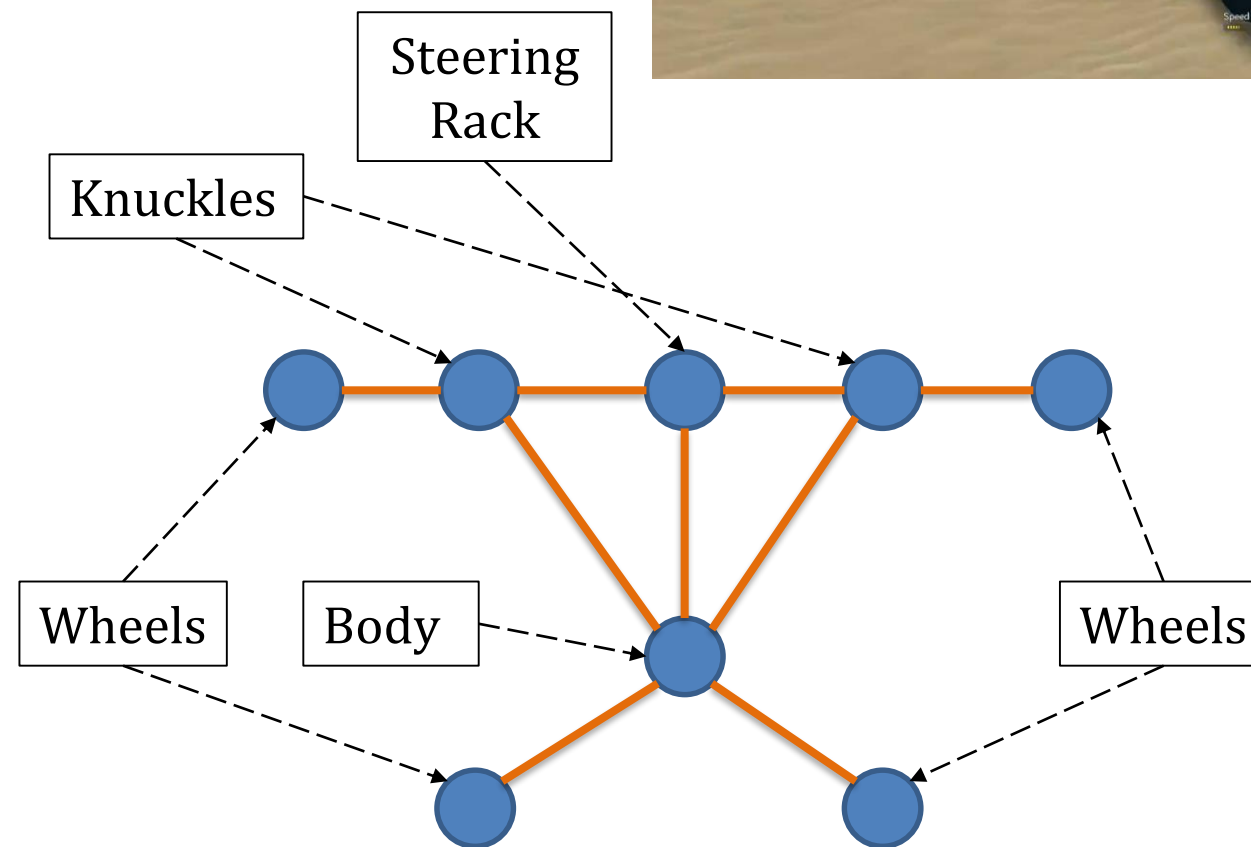
It is the **Edge Graph** of the Body Graph:

- **Nodes: Constraints**
- **Edges: Common bodies between constraints**

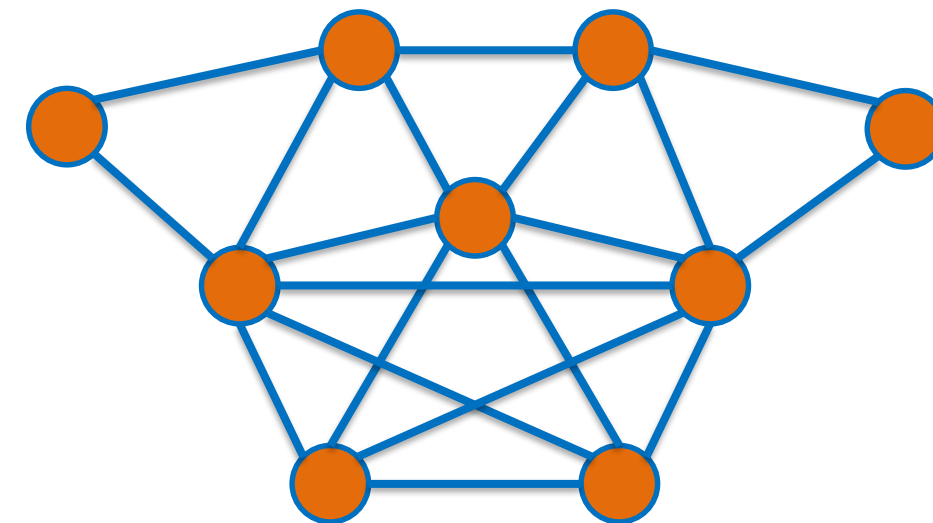
Constraint Graph = Graph of the Constraint Matrix

# Constraint Graph

Constraint Matrix	Constraint Graph
Diagonal block $\mathcal{N}_i$	Node $\mathbf{n}_i$
Off-diagonal block $\mathcal{E}_{ji}$	Edge $\mathbf{e}_{ji}$ between $\mathbf{n}_i$ and $\mathbf{n}_j$
Gaussian Elimination	Graph Elimination

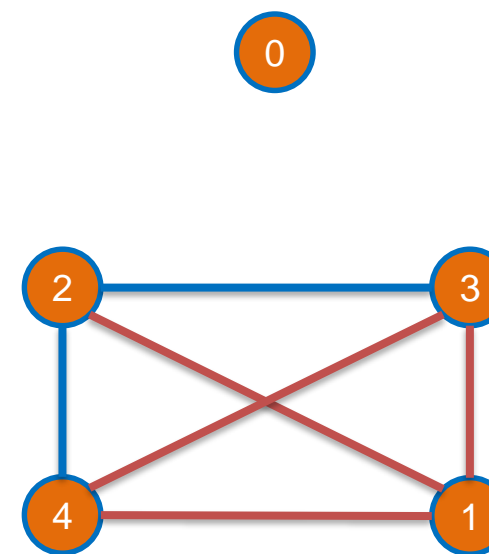
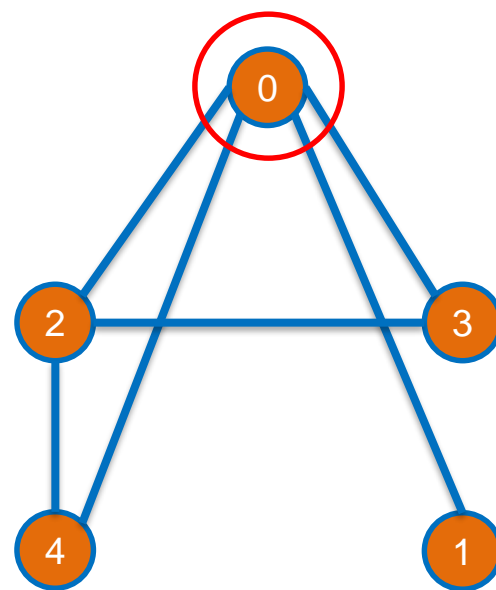


Body Graph

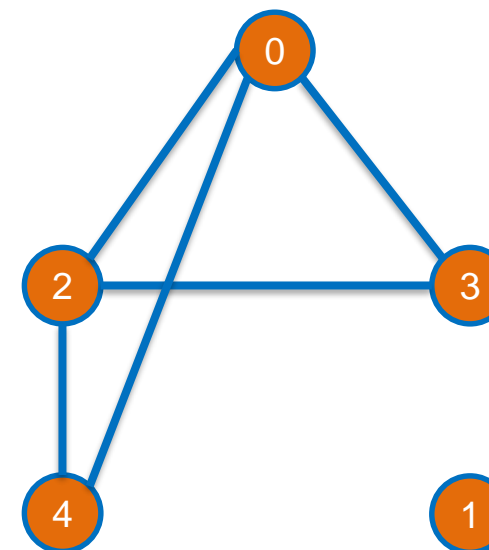
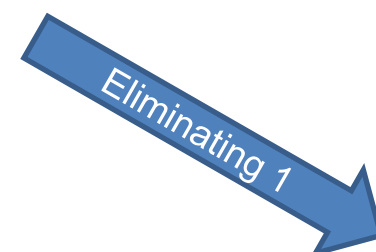


Constraint Graph

# Gaussian Elimination on Graphs



4 new edges



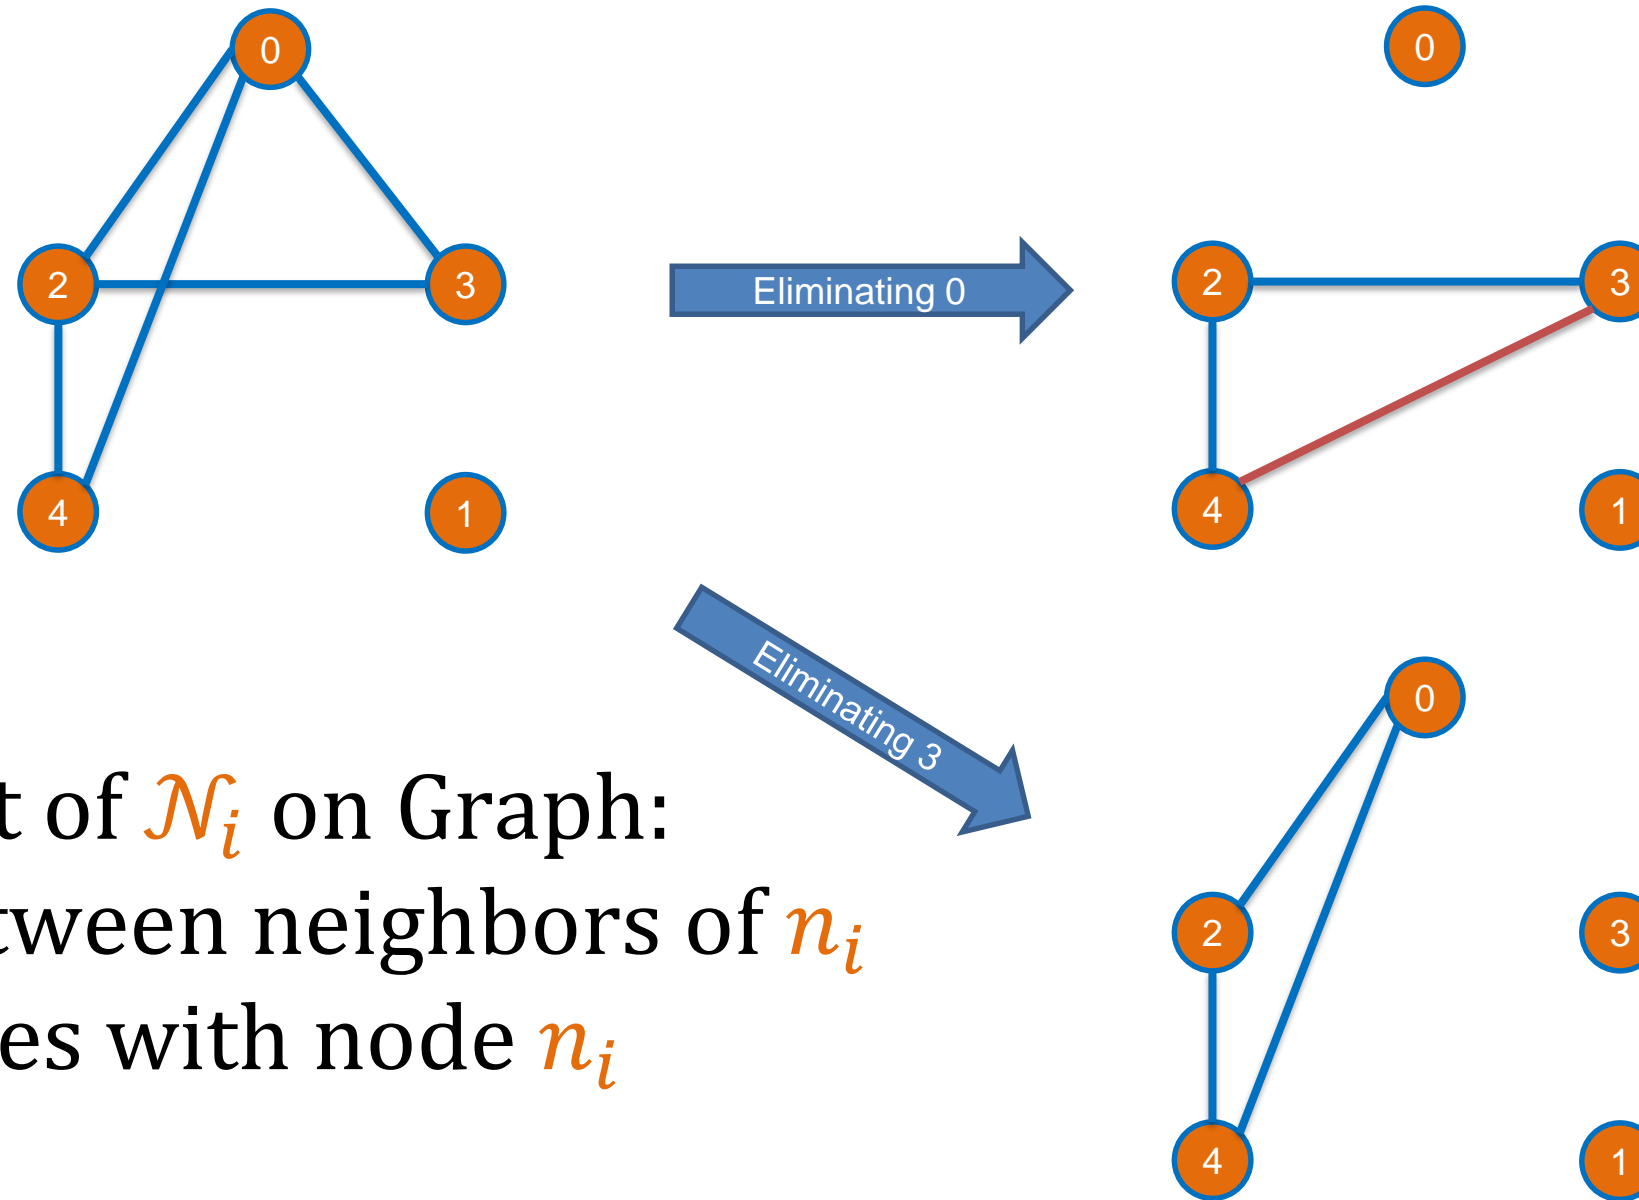
no new edges

$n_i$  is called the **Pivot**.

Schur Complement of  $n_i$ :

1. Add edges between neighbors of  $n_i$
2. Eliminate edges with node  $n_i$

# Gaussian Elimination on Graphs



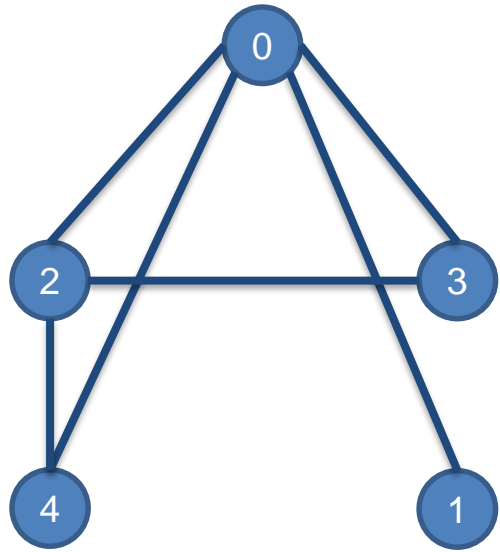
1 new edge

no new edges

Schur Complement of  $\mathcal{N}_i$  on Graph:

1. Add edges between neighbors of  $n_i$
2. Eliminate edges with node  $n_i$

# Gaussian Elimination on Graphs



A **perfect elimination order** for this graph is:

[1, 3, 0, 2, 4]

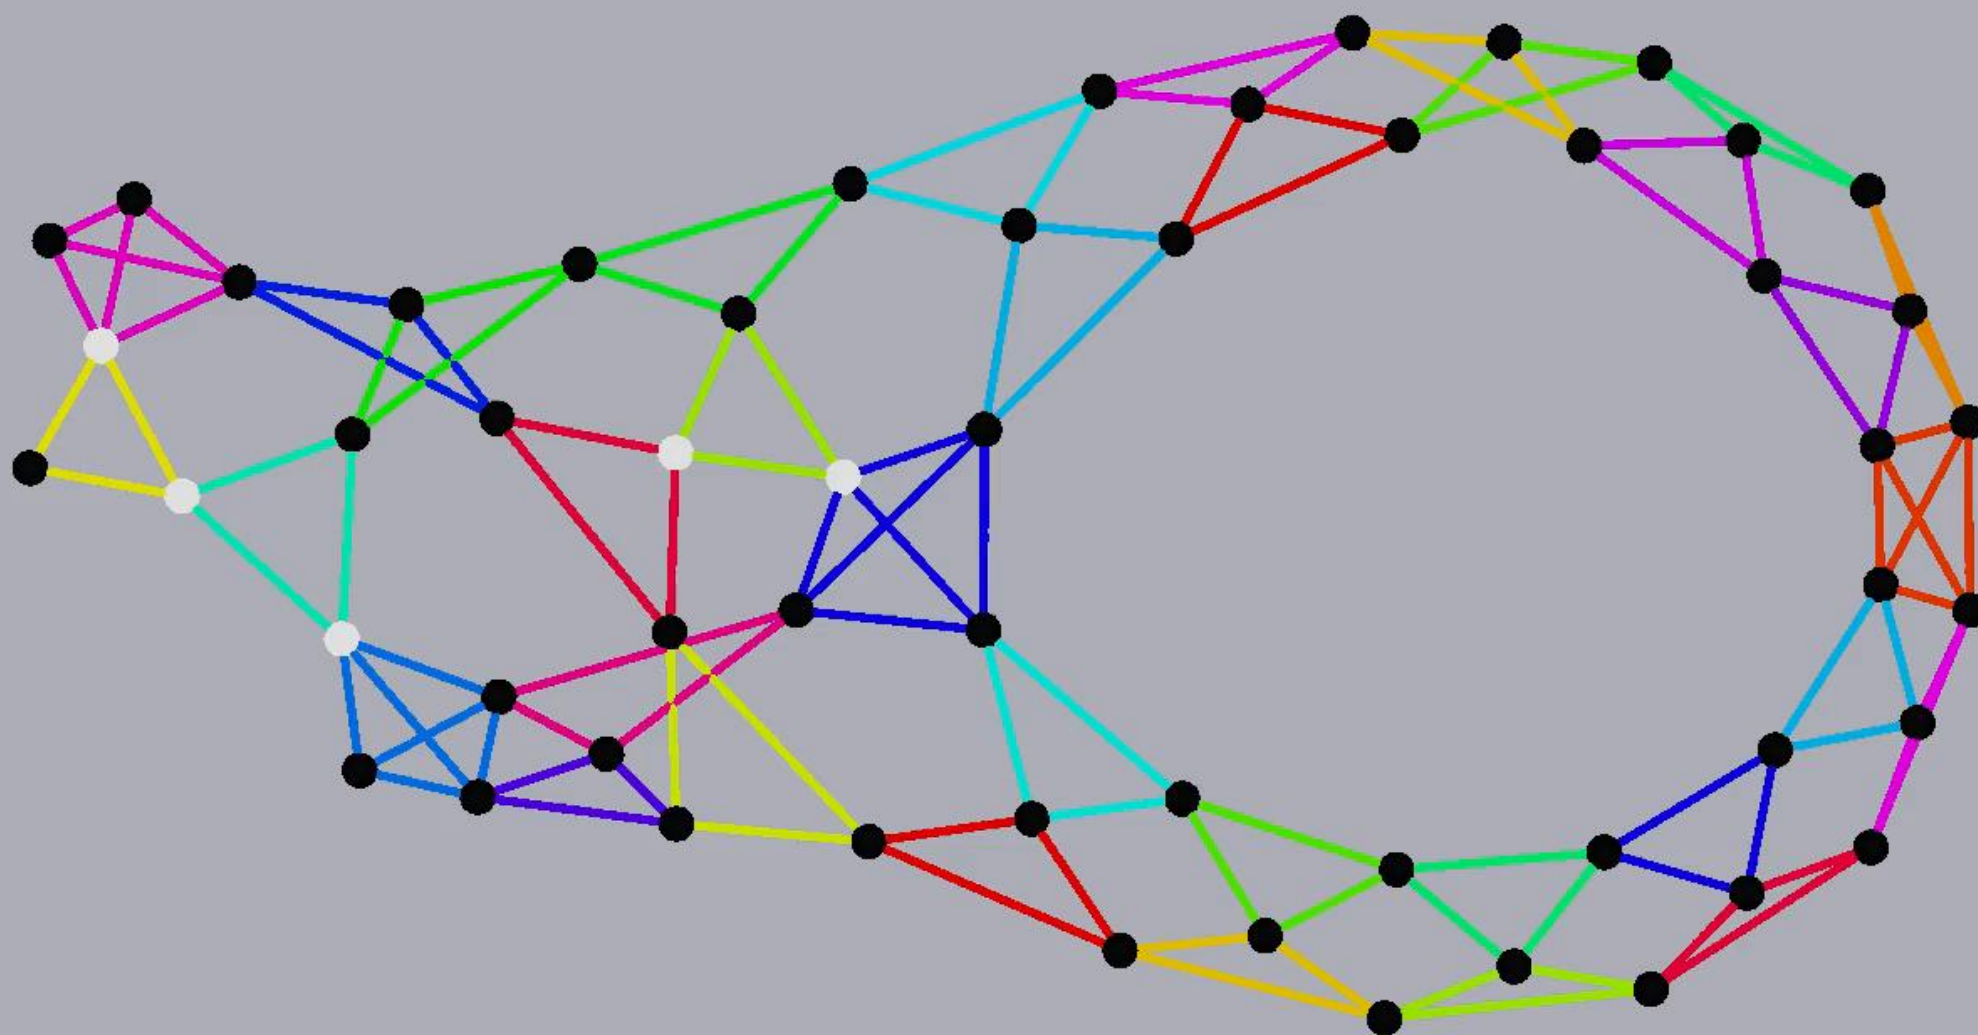
Does not always exist!

Elimination Game: minimize the number of new edges

Finding such ordering is **NP-Complete**.

There are good **heuristics**!





# Gaussian Elimination on Graphs

Heuristics for Graph Elimination:

- **Minimum Degree Algorithm (MDA)**
  - Fast but generates mediocre ordering
- **Minimum Edge Creation Algorithm (MECA)**
  - More expensive but better ordering
  - We only need to compute this once!

# Gaussian Elimination on Graphs

## **Minimum Edge Creation Algorithm (MECA):**

At each step, eliminate the pivot that creates a minimum number of new edges...

# Gaussian Elimination on Graphs

Graph Elimination gives us:

1. An ordered sequence of nodes (**pivot sequence**):

$$Pivots = [n_0, n_1, \dots]$$

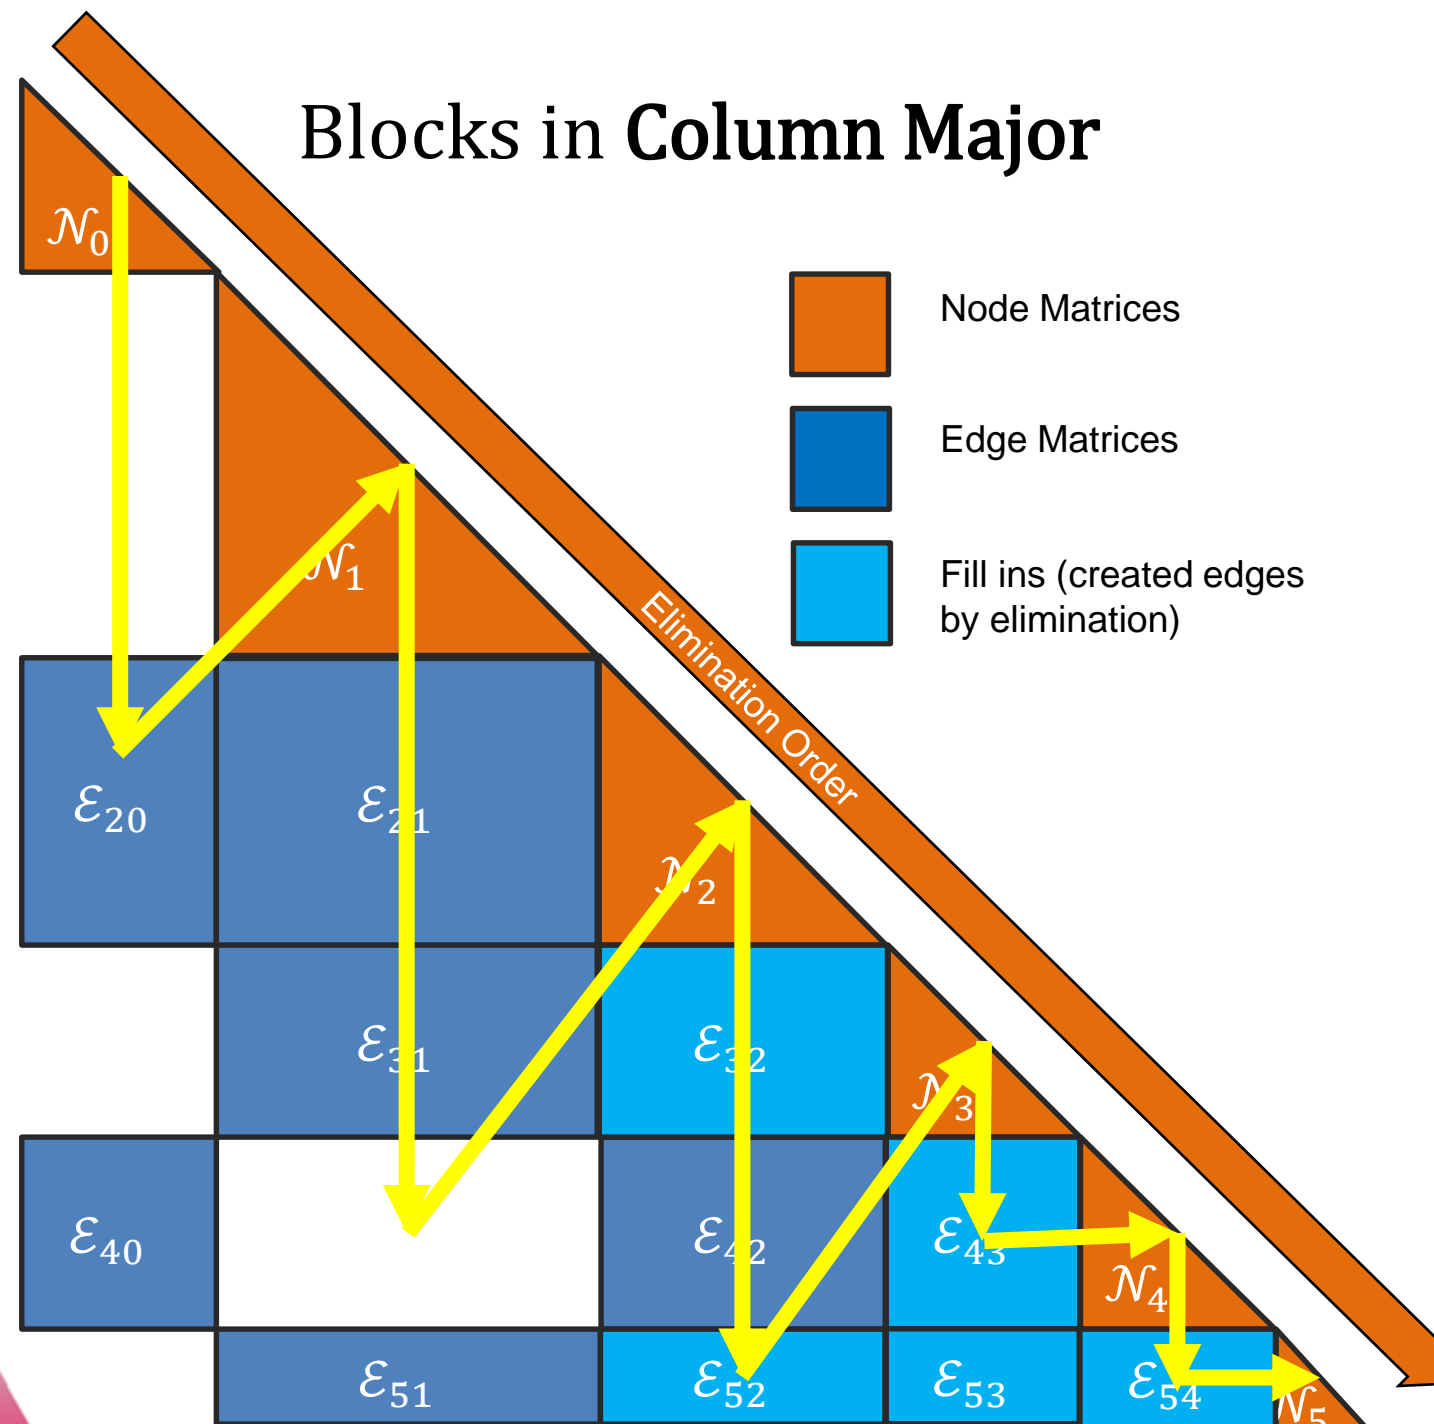
2. For each pivot  $n_i$  a sequence of eliminated edges (**edge sequences**):

$$Elim(n_i) = [e_{j_0,i}, e_{j_1,i}, \dots]$$

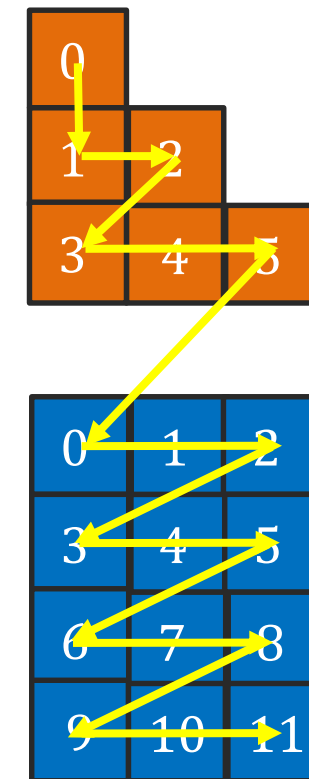
These sequences should be sorted in pivot order:

$$n_{j_0} < n_{j_1} < \dots$$

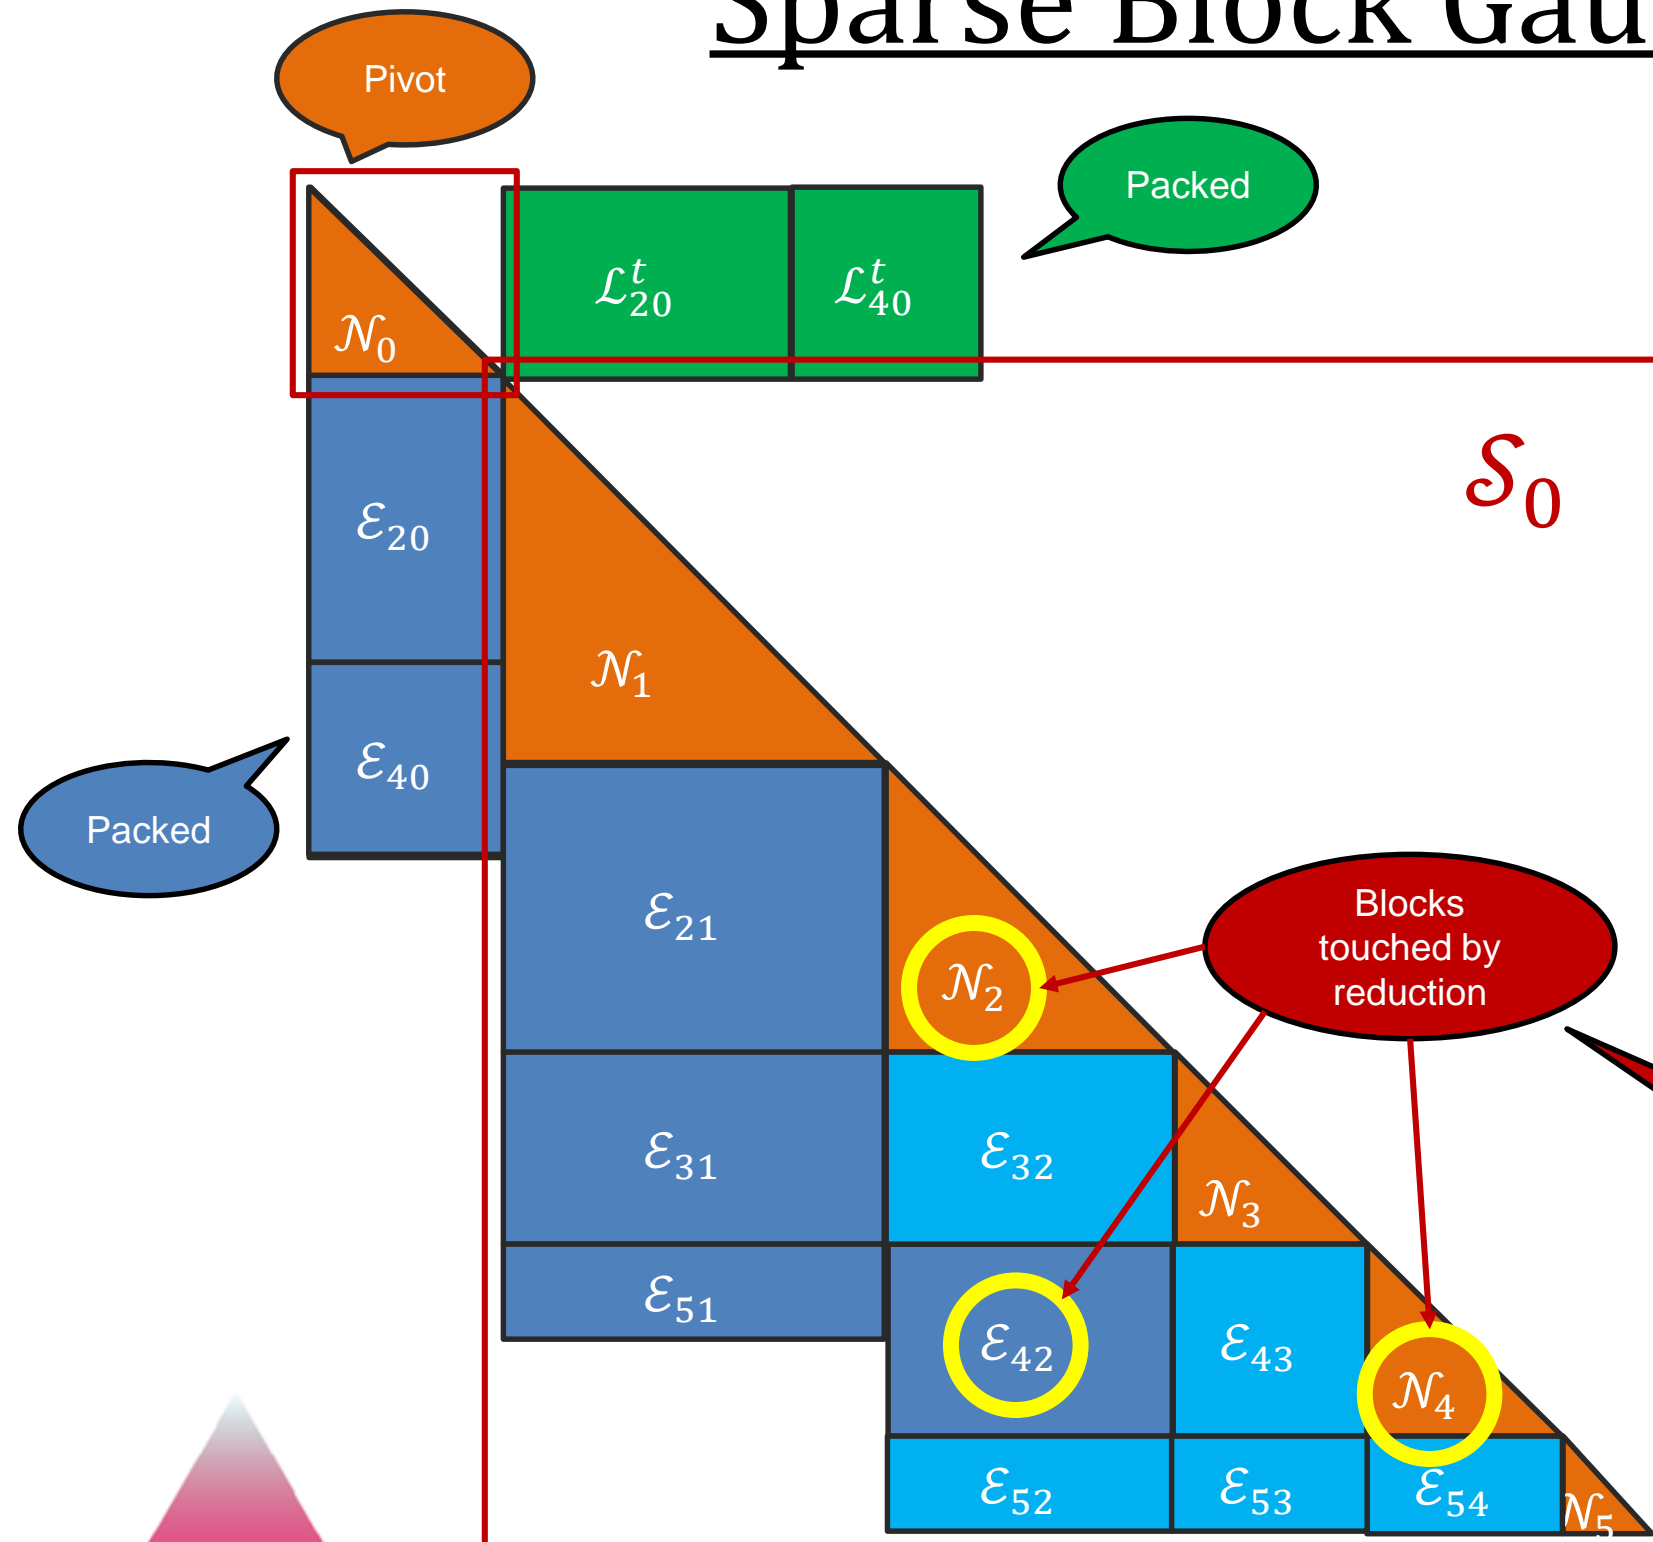
# Sparse Block Matrix Memory Layout



Elements of blocks in **Row Major**



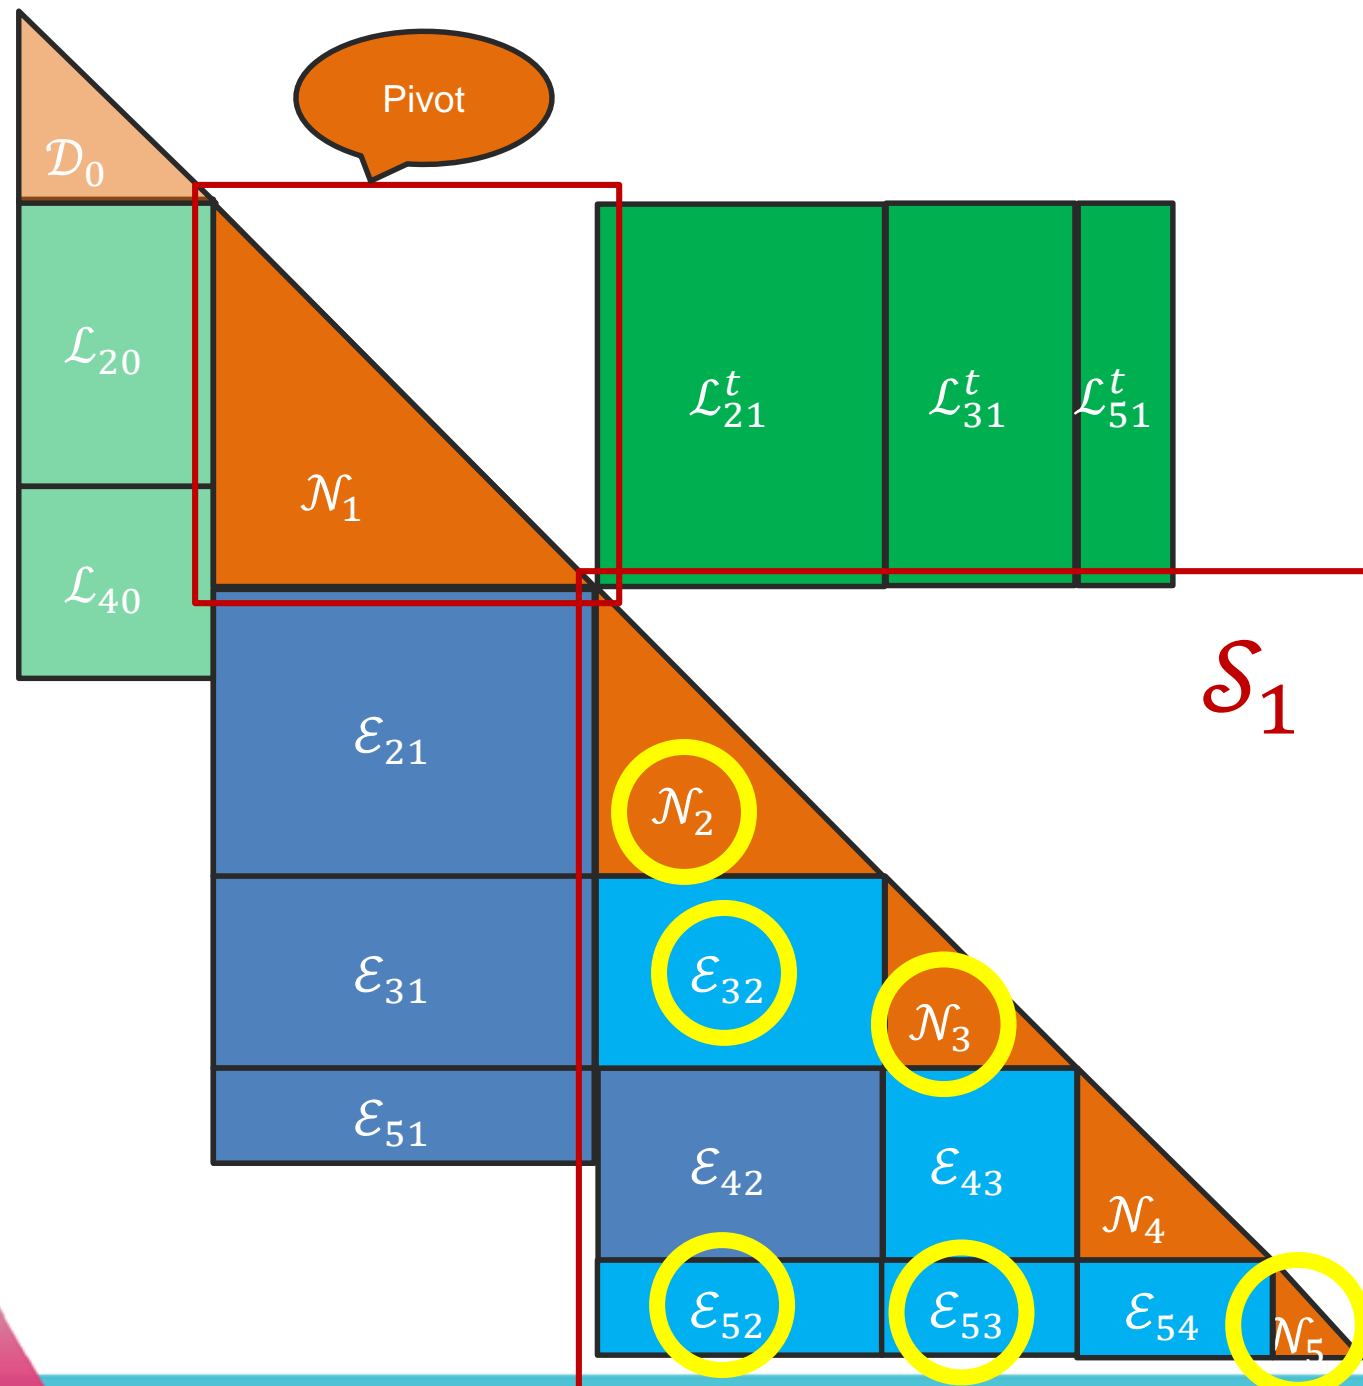
# Sparse Block Gaussian Elimination



1. LDL decompose the pivot  $\mathcal{N}_0$  in-place,  $\mathcal{D}_0 := LDL(\mathcal{N}_0)$
2. Compute  $\mathcal{L}_{*,0} = \mathcal{E}_{*,0} \mathcal{N}_0^{-1}$  in a temp buffer
3. Reduce  $\mathcal{S}_0 \leftarrow \mathcal{S}_0 - \mathcal{E}_{*,0} \mathcal{L}_{*,a}^t$
4. Replace  $\mathcal{E}_{*,0} \leftarrow \mathcal{L}_{*,0}$

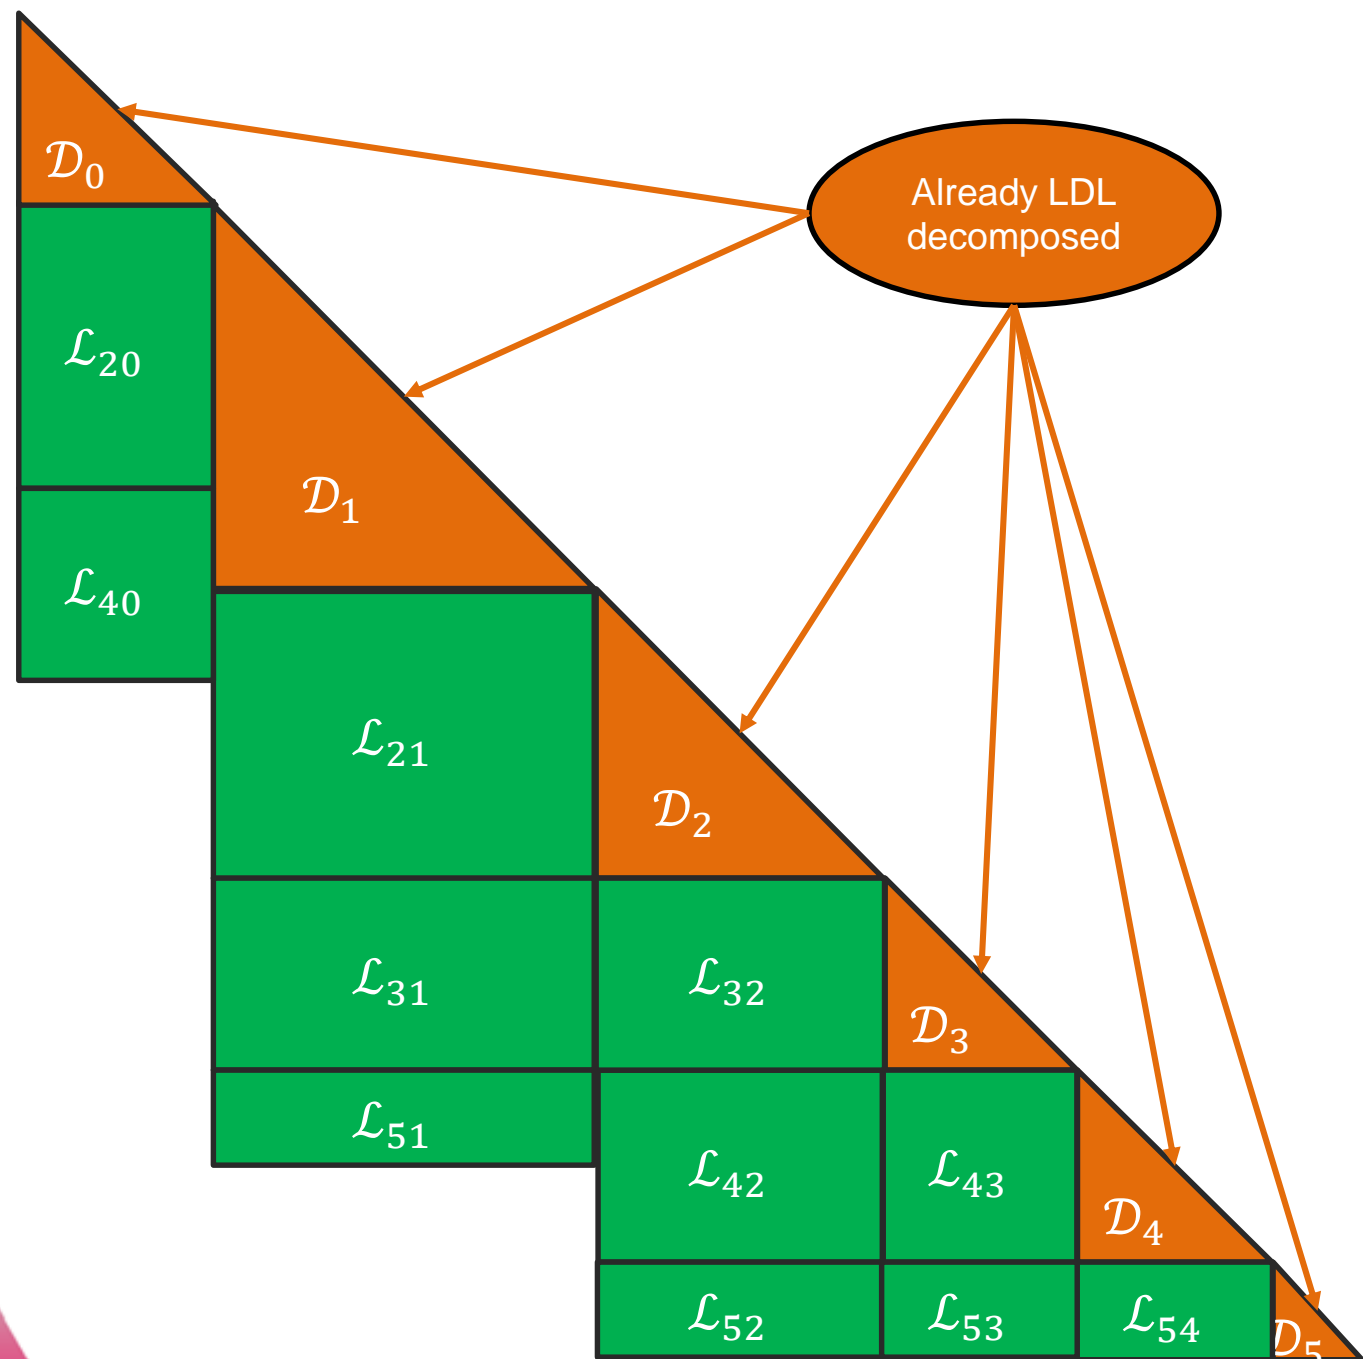


# Sparse Block Gaussian Elimination

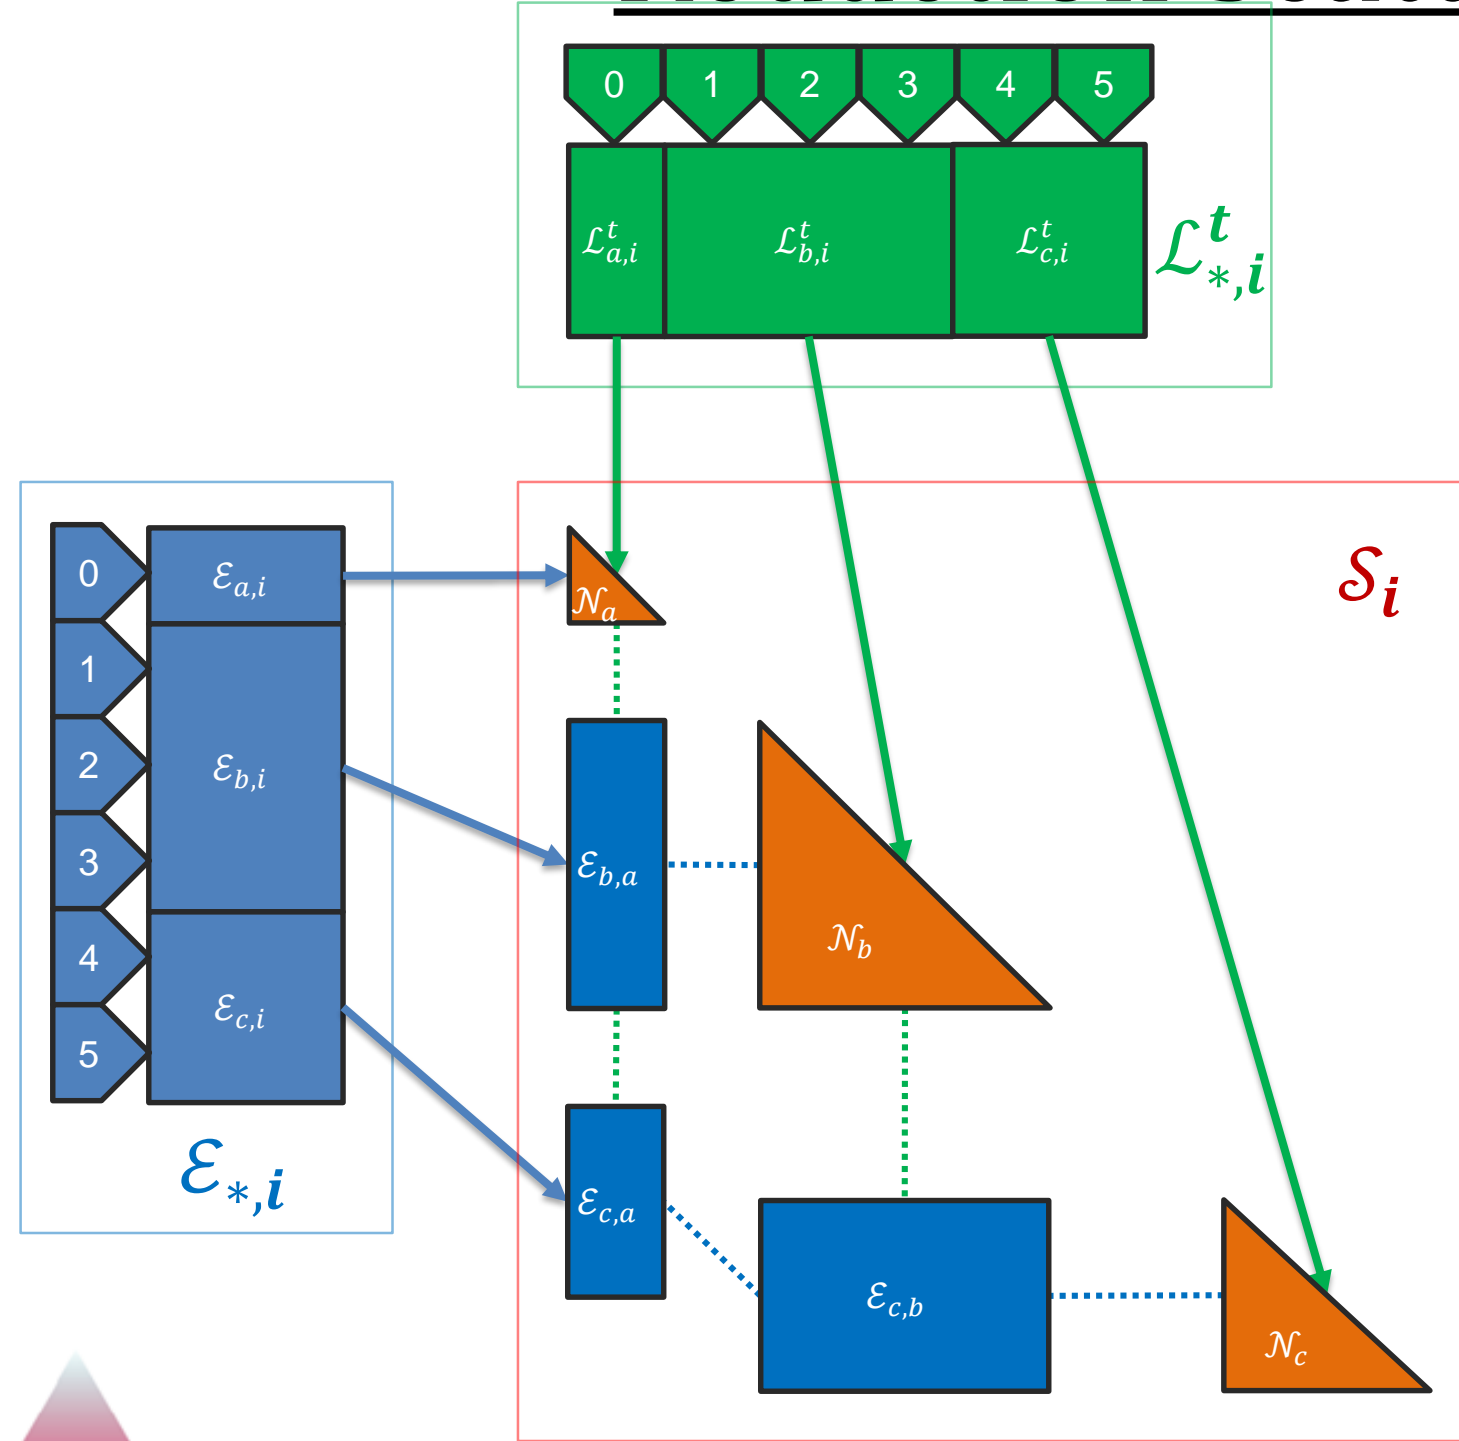


1. LDL decompose the pivot  $\mathcal{N}_1$  in-place,  $\mathcal{D}_1 := LDL(\mathcal{N}_1)$
2. Compute  $\mathcal{L}_{*,1} = \epsilon_{*,1} \mathcal{N}_1^{-1}$  in a temp buffer
3. Reduce  $\mathcal{S}_1 \leftarrow \mathcal{S}_1 - \epsilon_{*,1} \mathcal{L}_{*,1}^t$
4. Replace  $\epsilon_{*,1} \leftarrow \mathcal{L}_{*,1}$

# Sparse Block LDL Decomposition



# Reduction Scattering Indexation



k	Row of $E$	Column of $L^t$	$RSI[k]$
0	0	0	$\&\mathcal{N}_a[0,0]$
1	1	0	$\&\mathcal{E}_{b,a}[0,0]$
2	1	1	$\&\mathcal{N}_b[0,0]$
3	2	0	$\&\mathcal{E}_{b,a}[1,0]$
4	2	1	$\&\mathcal{N}_b[1,0]$
5	2	2	$\&\mathcal{N}_b[1,1]$
6	3	0	$\&\mathcal{E}_{b,a}[2,0]$
7	3	1	$\&\mathcal{N}_b[2,0]$
8	3	2	$\&\mathcal{N}_b[2,1]$
9	3	3	$\&\mathcal{N}_b[2,2]$
10	4	0	$\&\mathcal{E}_{c,a}[0,0]$
$\vdots$	$\vdots$	$\vdots$	$\vdots$
14	5	5	$\&\mathcal{N}_c[1,1]$

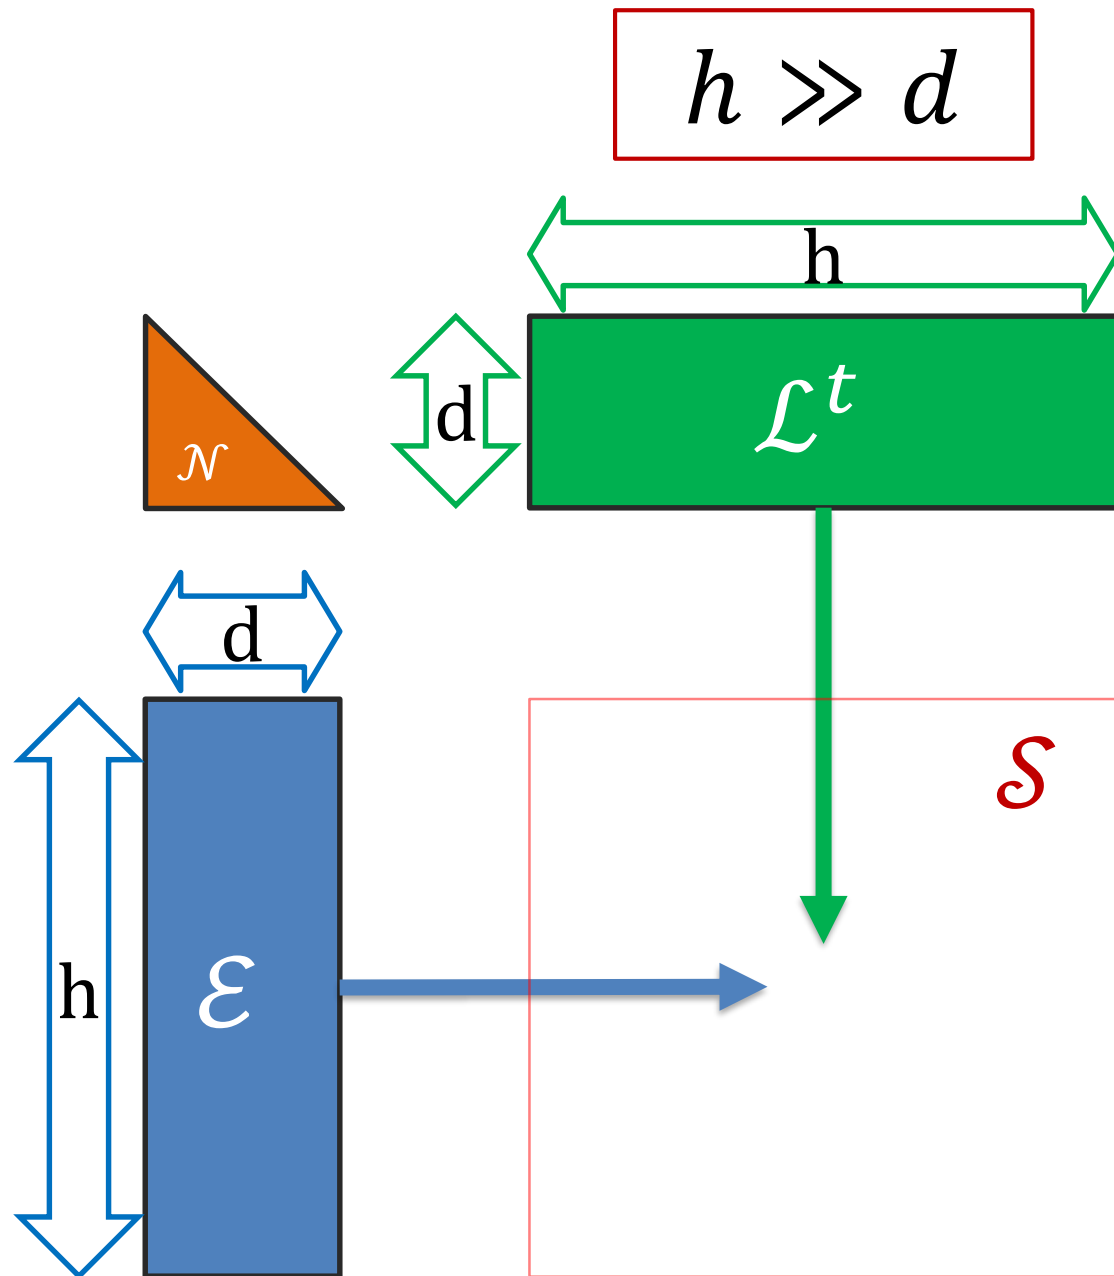
Size of  $RSI$  :  

$$\frac{h(h+1)}{2}$$

Here:  
 $h = 5,$   
 $size = 15$

$RSI$  better use indices  $\Rightarrow$  less memory!  
 Precompute during symbolic phase.

# Performance of Block LDL



Operation count:

1.  $LDL(\mathcal{N}) : \frac{d^3}{6}$

2.  $\mathcal{L} = \varepsilon \mathcal{N}^{-1} : hd^2$

3.  $\mathcal{S} = \mathcal{S} - \varepsilon \mathcal{L}^t : d \frac{h(h+1)}{2}$

Memory Access:

1.  $LDL(\mathcal{N}) : \frac{d(d+1)}{2}$

2.  $\mathcal{L} = \varepsilon \mathcal{N}^{-1} : 2hd$

3.  $\mathcal{S} = \mathcal{S} - \varepsilon \mathcal{L}^t : 2hd + \frac{h(h+1)}{2}$

Modern processors:

- Good at floating point operations
- Bad at memory access

$$\frac{\text{Op Count}}{\text{Mem Access}} \sim d$$

$\Rightarrow$  Fastest for large  $d$ , as long as  $d \ll h$

Dominant

# Reduction Step: 70% CPU Time

Edge Matrix  
Width

Schur  
Complement

Edge Matrix

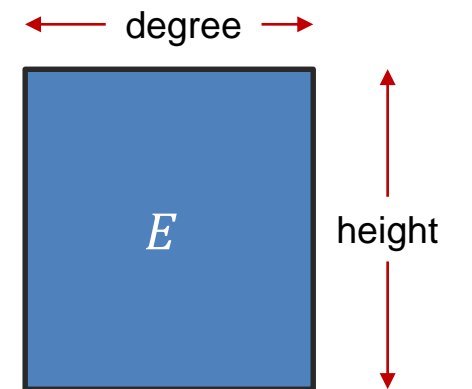
L Matrix

Reduction  
Scattering  
Indexation

Edge Matrix  
Height

```
template<uint degree>
void reduce(float* S, const float* E, const float* L, const uint* rsi, uint height)
{
    for (uint i = 0; i < height; i++)
    {
        for (uint j = 0; j <= i; j++)
        {
            S[ *rsi ] -= dotProduct<degree>(E + i*Dim, L + j*Dim);
            rsi++;
        }
    }
}
```

Loop Unroll and  
Inline!



# Implementation

```
template<uint deg> void eliminatePivot (float* S, float* LTemp, float* N, float* E,
    const uint* rsi, uint height)
{
    // Dense LDL decomposition of N
    ldlDecompose<deg>(N);

    // Compute  $L = E * N^{-1}$ 
    computeL<deg>(LTemp, N, E, height);

    // Reduce  $S -= E * L^t$ 
    reduce<deg>(S, E, LTemp, rsi, height);

    // Don't need the Edge Matrix anymore
    // Overwrite with L Matrix
    copy(E, LTemp, height * deg);
}
```

~5% of CPU Time  
No Unrolling Necessary

~25% of CPU  
Time... Unroll!

~70% of CPU  
Time

# Implementation

Use switch statement:

```
for (const Pivot& pivot : pivots)
{
    switch (pivot.degree)
    {
        case 0: break;
        case 1: eliminatePivot<1>(...); break;
        case 2: eliminatePivot<2>(...); break;
        case 3: eliminatePivot<3>(...); break;
        case 4: eliminatePivot<4>(...); break;
        case 5: eliminatePivot<5>(...); break;
        case 6: eliminatePivot<6>(...); break;
        default: eliminatePivot(..., pivot.dimension); break; // Not templated for Degree > 6
    }
}
```



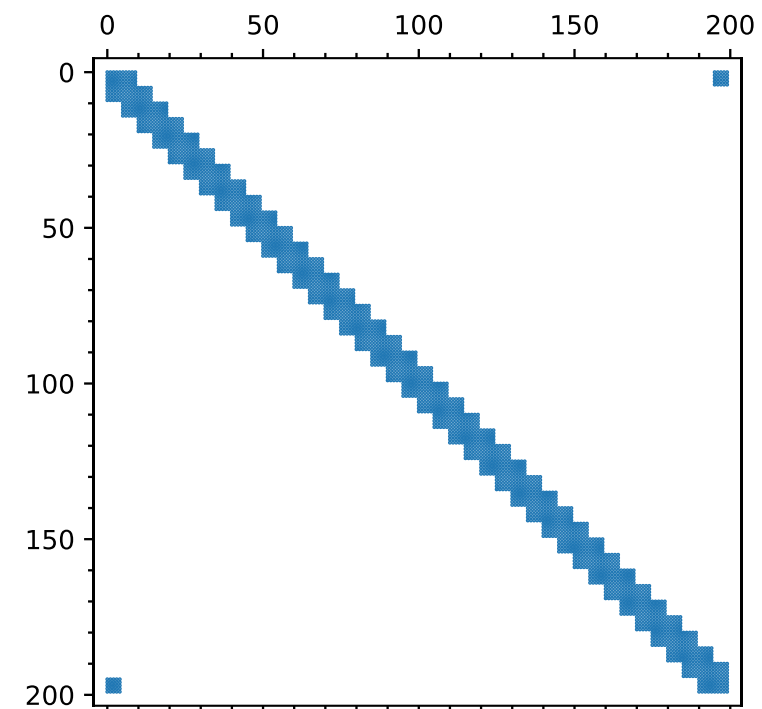
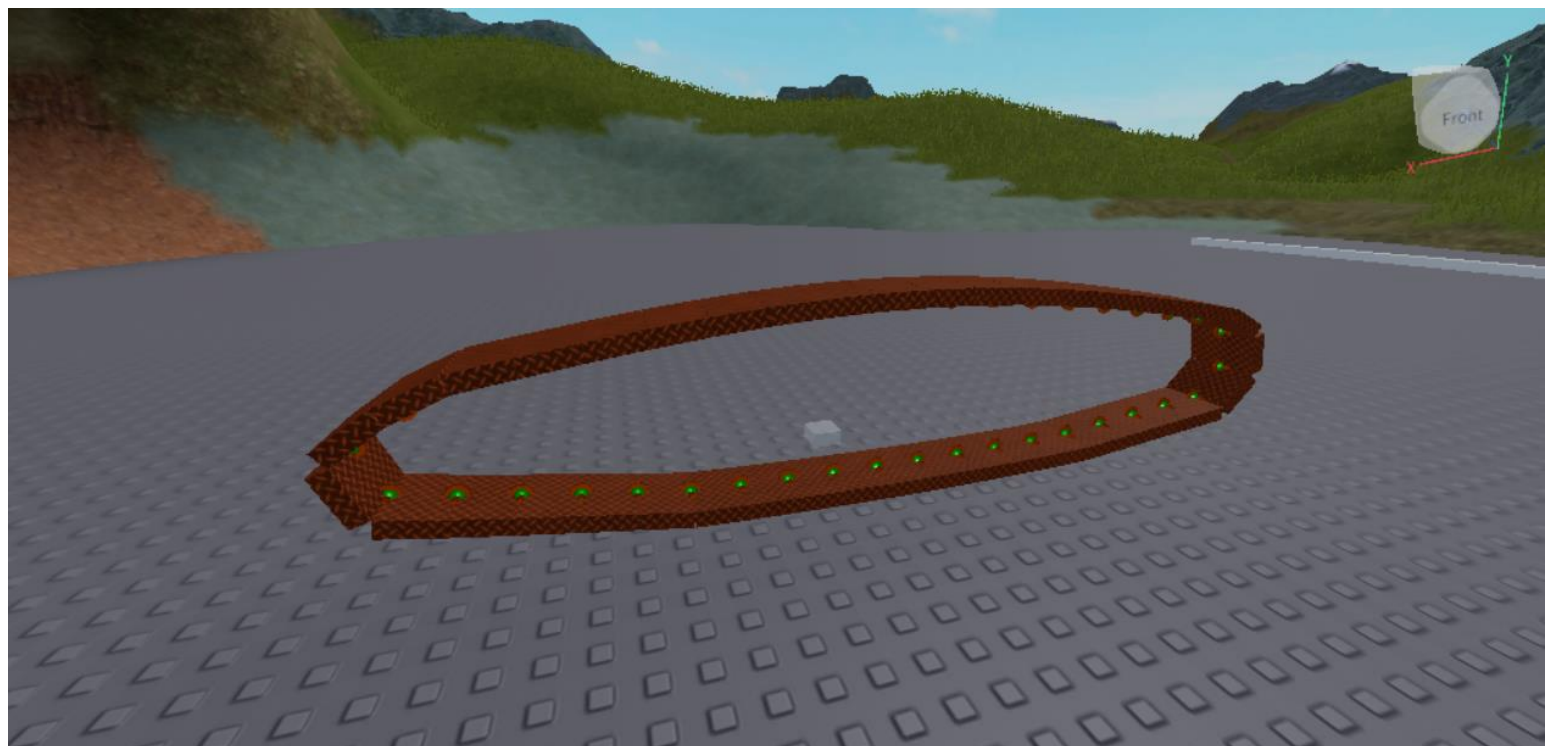
# Tracked Vehicle



3 Components:

- Main body
- Left track
- Right track

# Track



$\#Constraints = 40$

$Dim(\mathcal{H}) = 200$

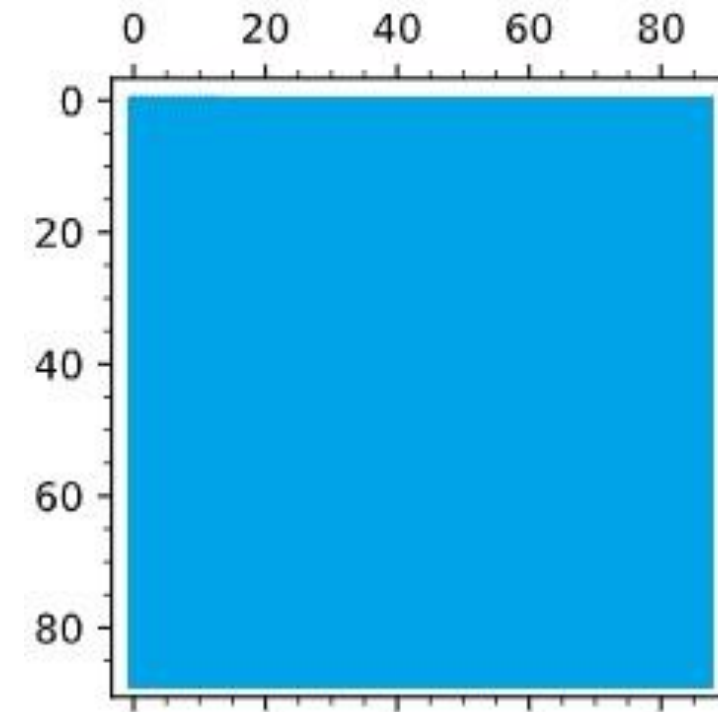
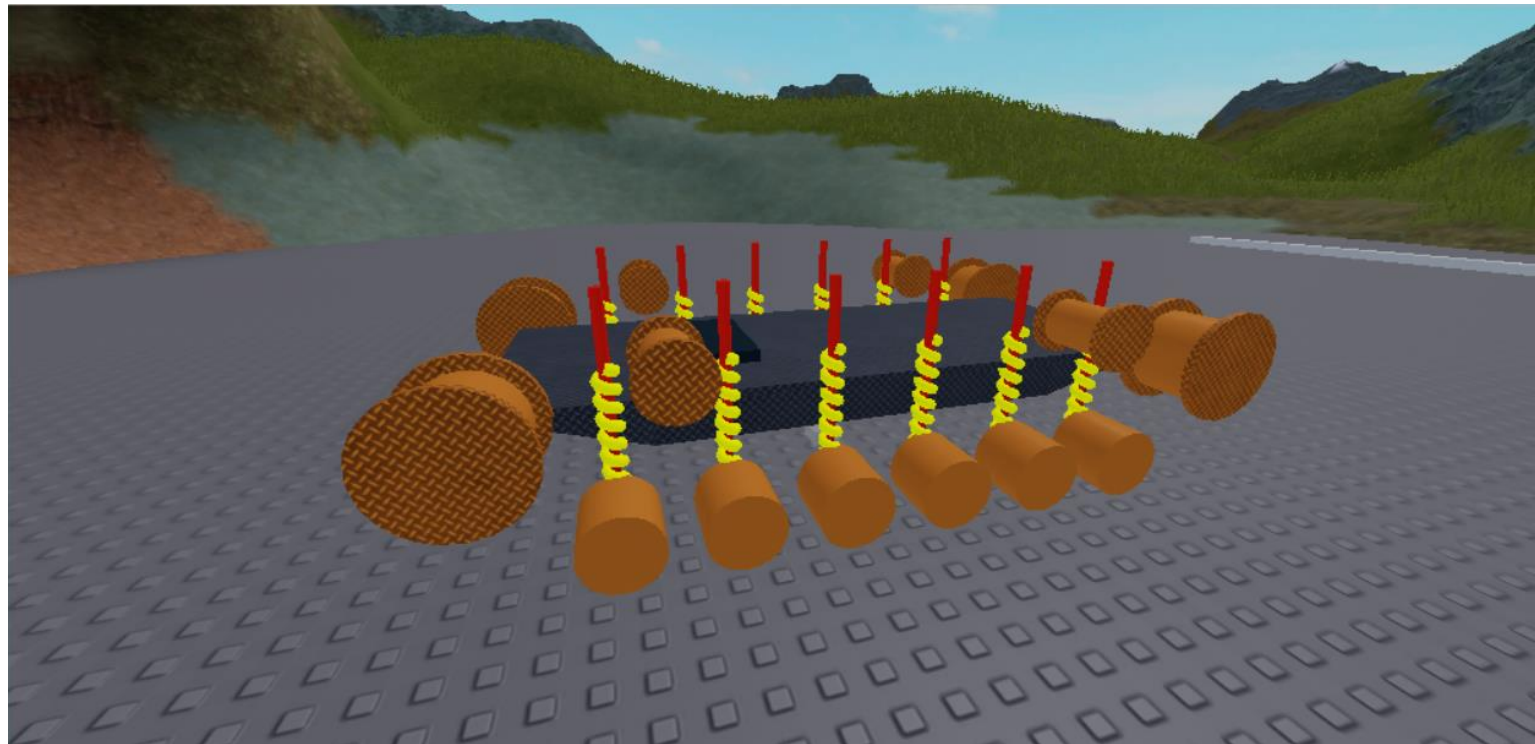
$Density(\mathcal{H}) = 7.5\%$

Sparse Cholesky LDL:

$Flops = 43k$

$Performance = 35\mu s$

# Multi-Wheeled Vehicle



$\#Constraints = 20 (= n)$

$Dim(\mathcal{H}) = 88$

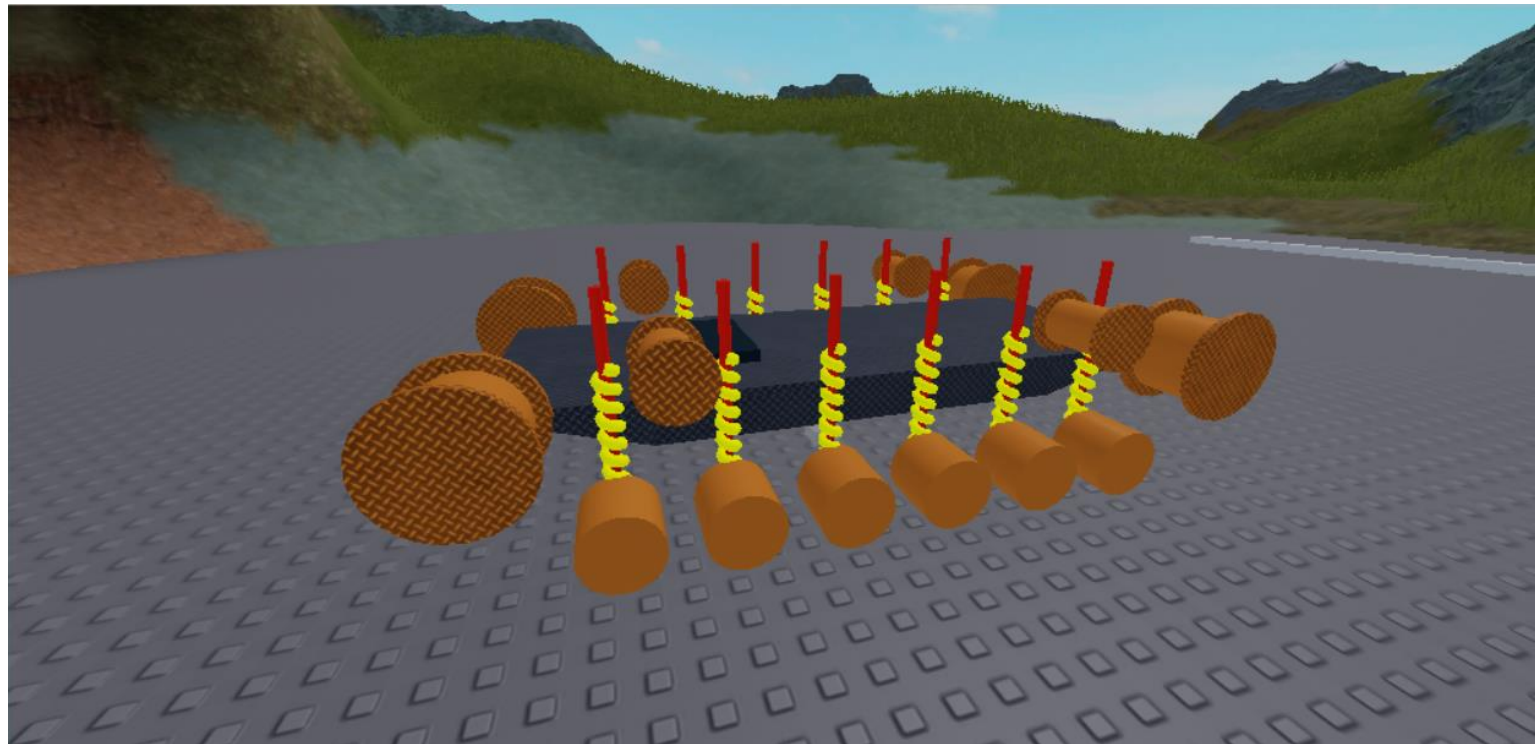
$Density(\mathcal{H}) = 100\%$

$Flops = 250k = O(n^3)$

There is a solution: **Body Shattering**



# Multi-Wheeled Vehicle



$\#Constraints = 20 (= n)$

$Dim(\mathcal{H}) = 88$

$Density(\mathcal{H}) = 100\%$

$Flops = 250k = O(n^3)$



**Body Shattering**

$\#Constraints = 29$

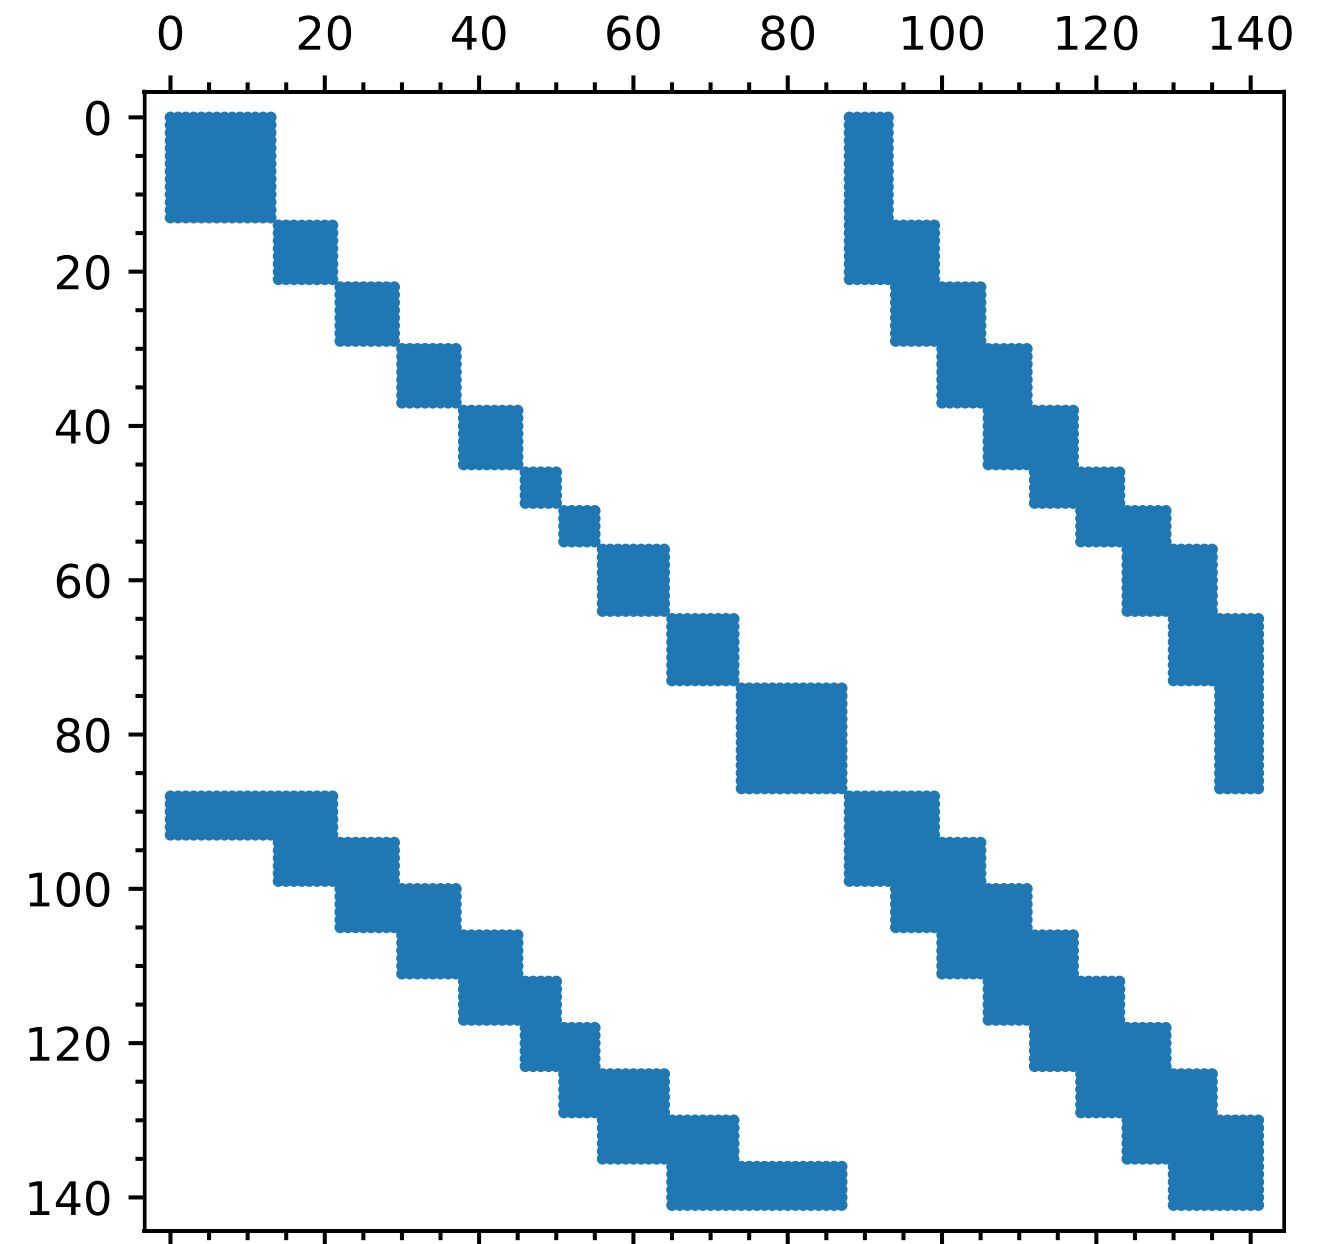
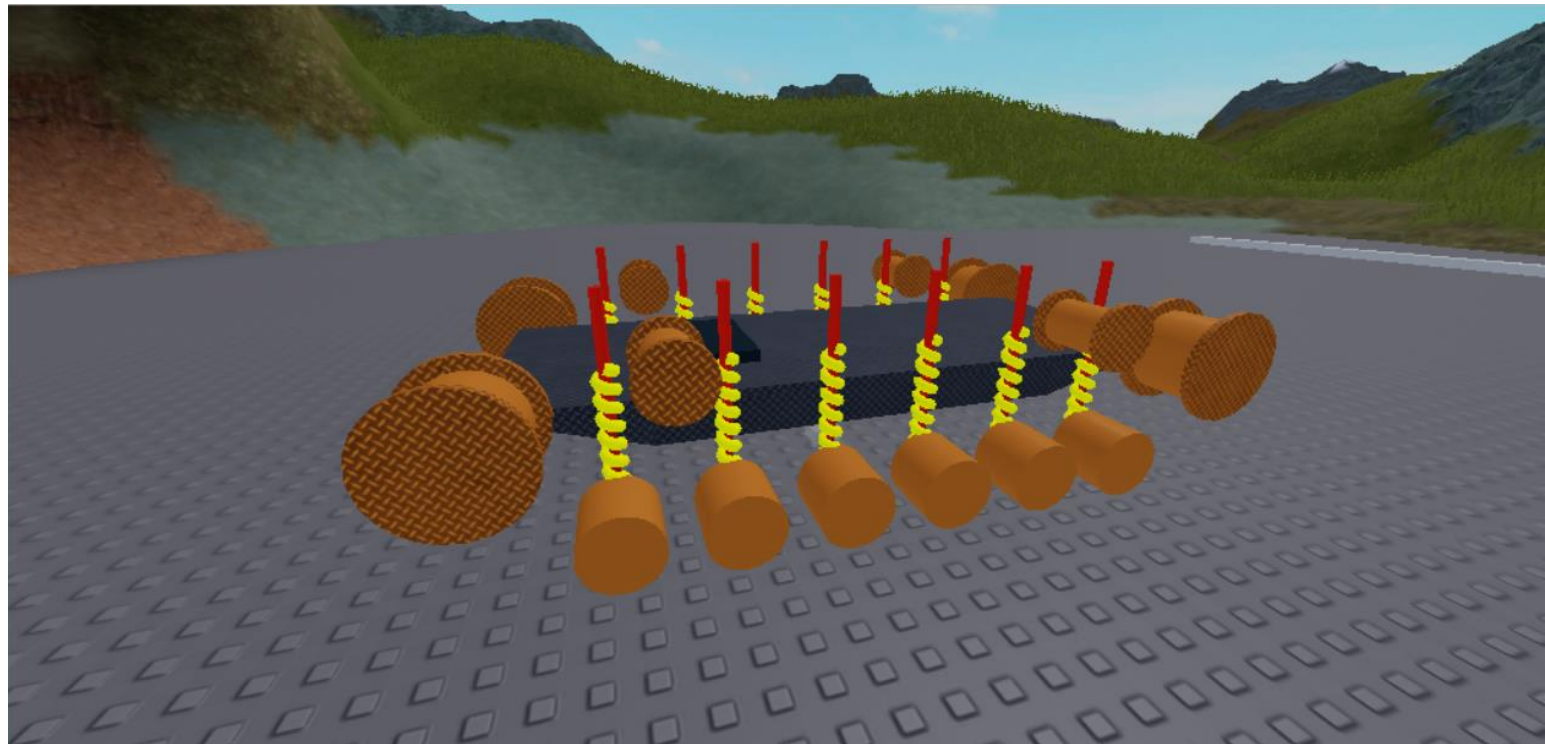
$Dim(\mathcal{H}) = 142$

$Density(\mathcal{H}) = 17.5\%$

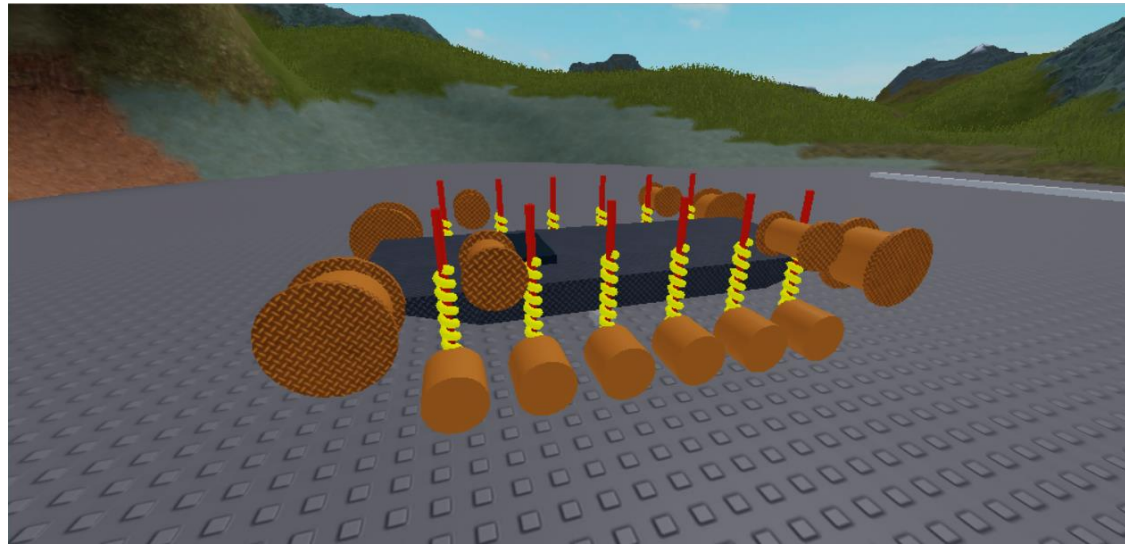
$Flops = 35k = O(n)$

$Performance = 34 \mu s$

# Shattered Multi-Wheeled Vehicle



# Dense Submatrices and Body Shattering



Large dense submatrix



Body with many constraints

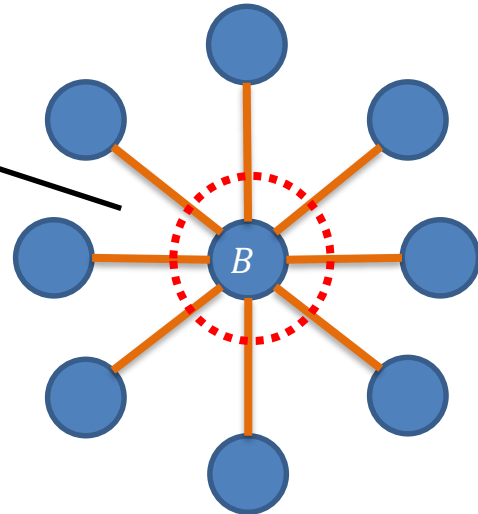
## Body Shattering:

- Find bodies with many constraints (total degree  $> 20$ )
- Split into smaller equal shards
- Join using rigid joints
- Distribute constraints over the shards

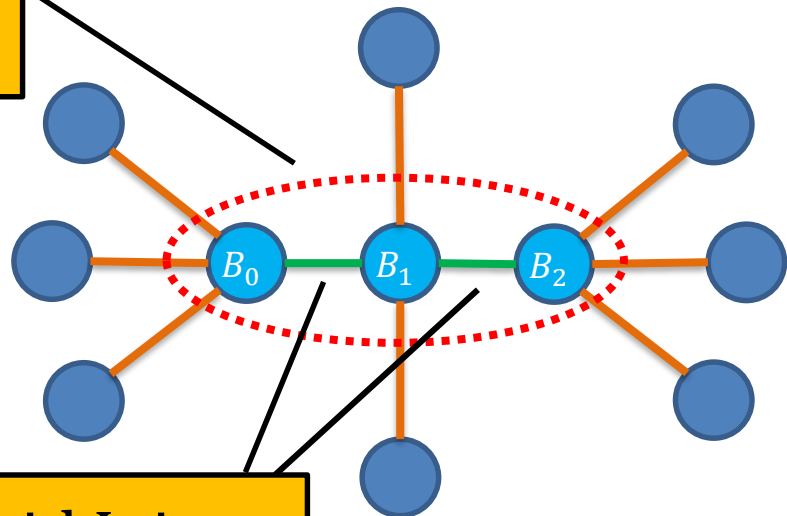
# Body Shattering

## Body Graph

B has many constraints



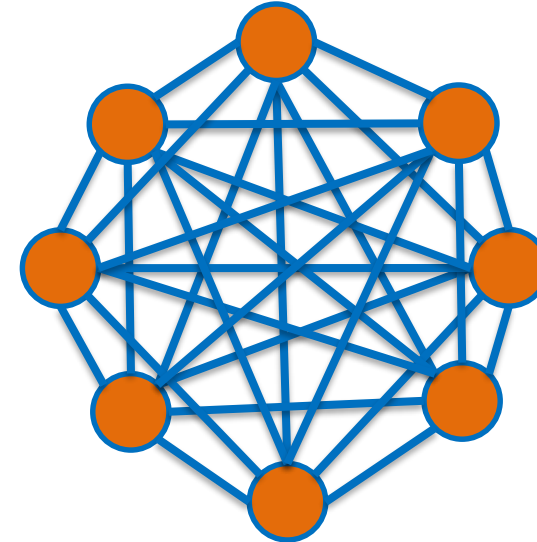
Shatter  $B$  into  $B_0, B_1, B_2$



Rigid Joints

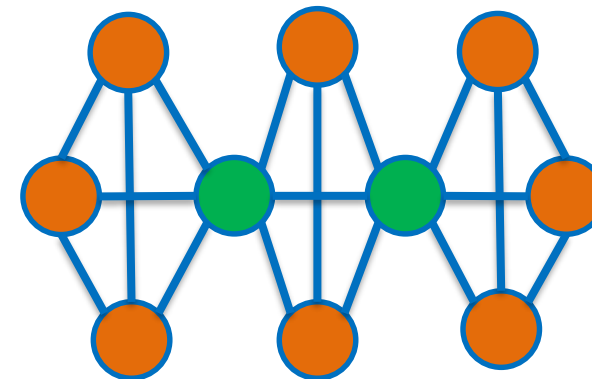
## Constraint Graph

Edge Graph



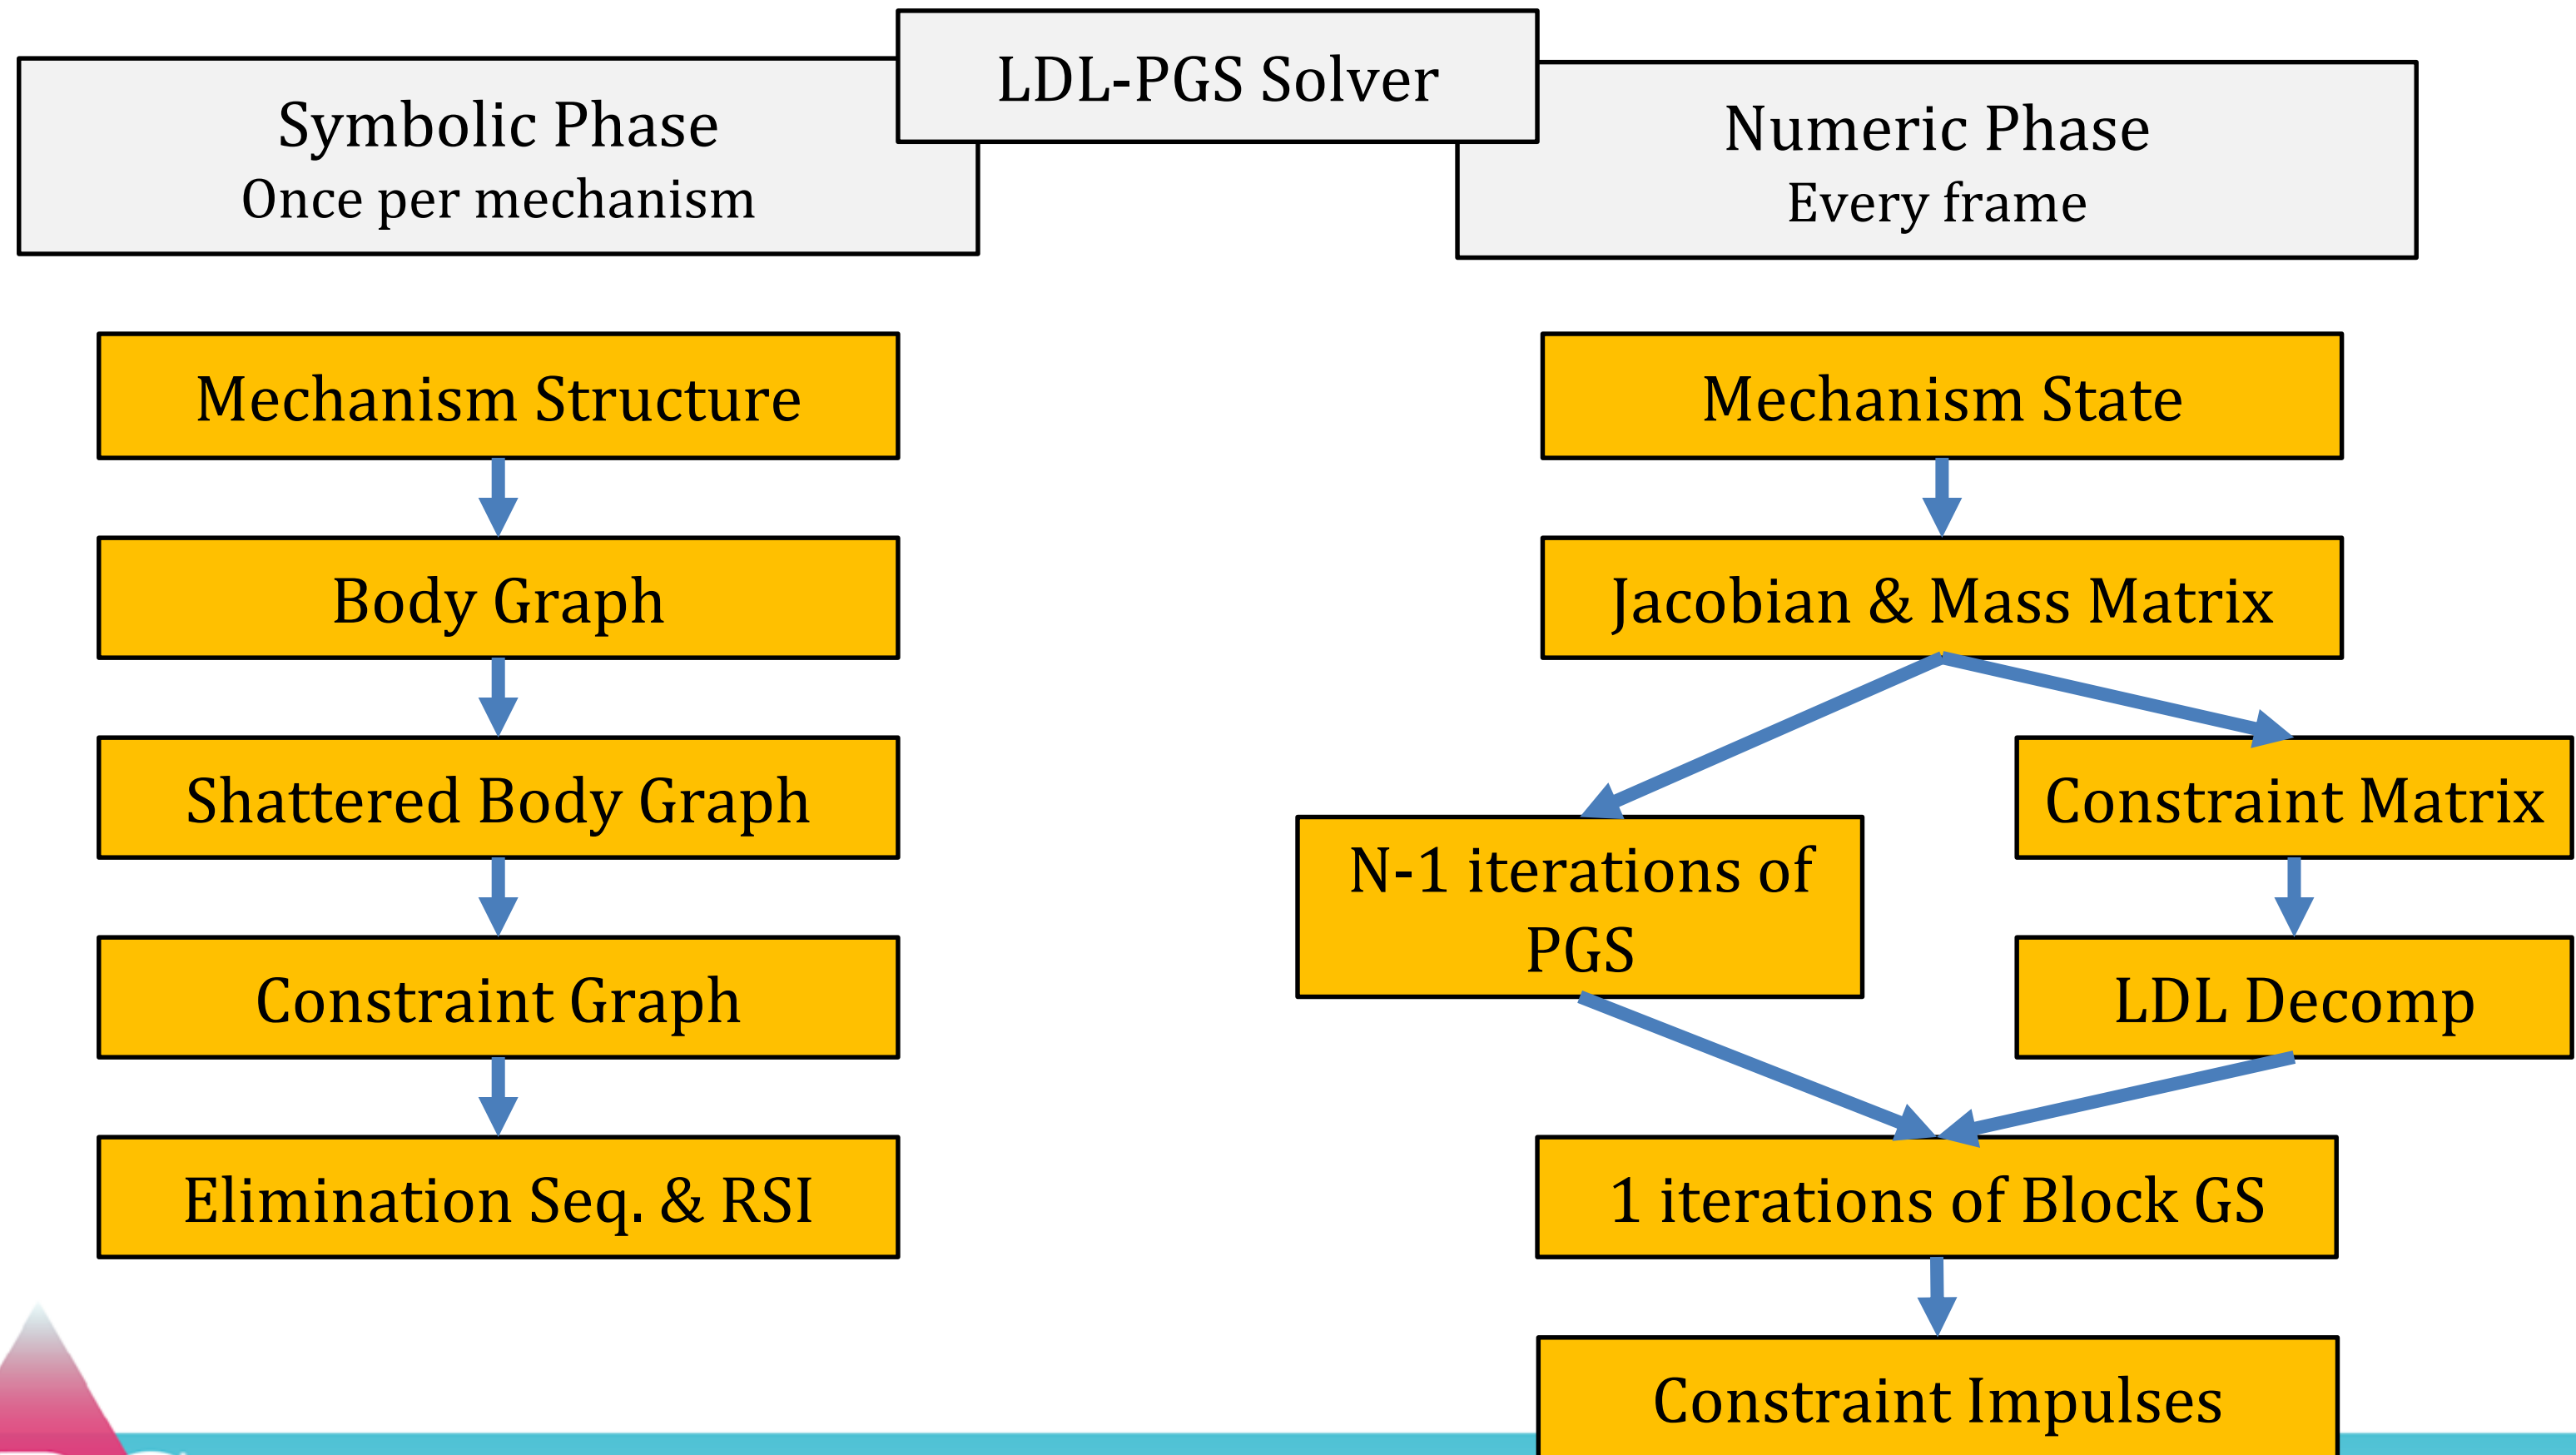
100% dense!

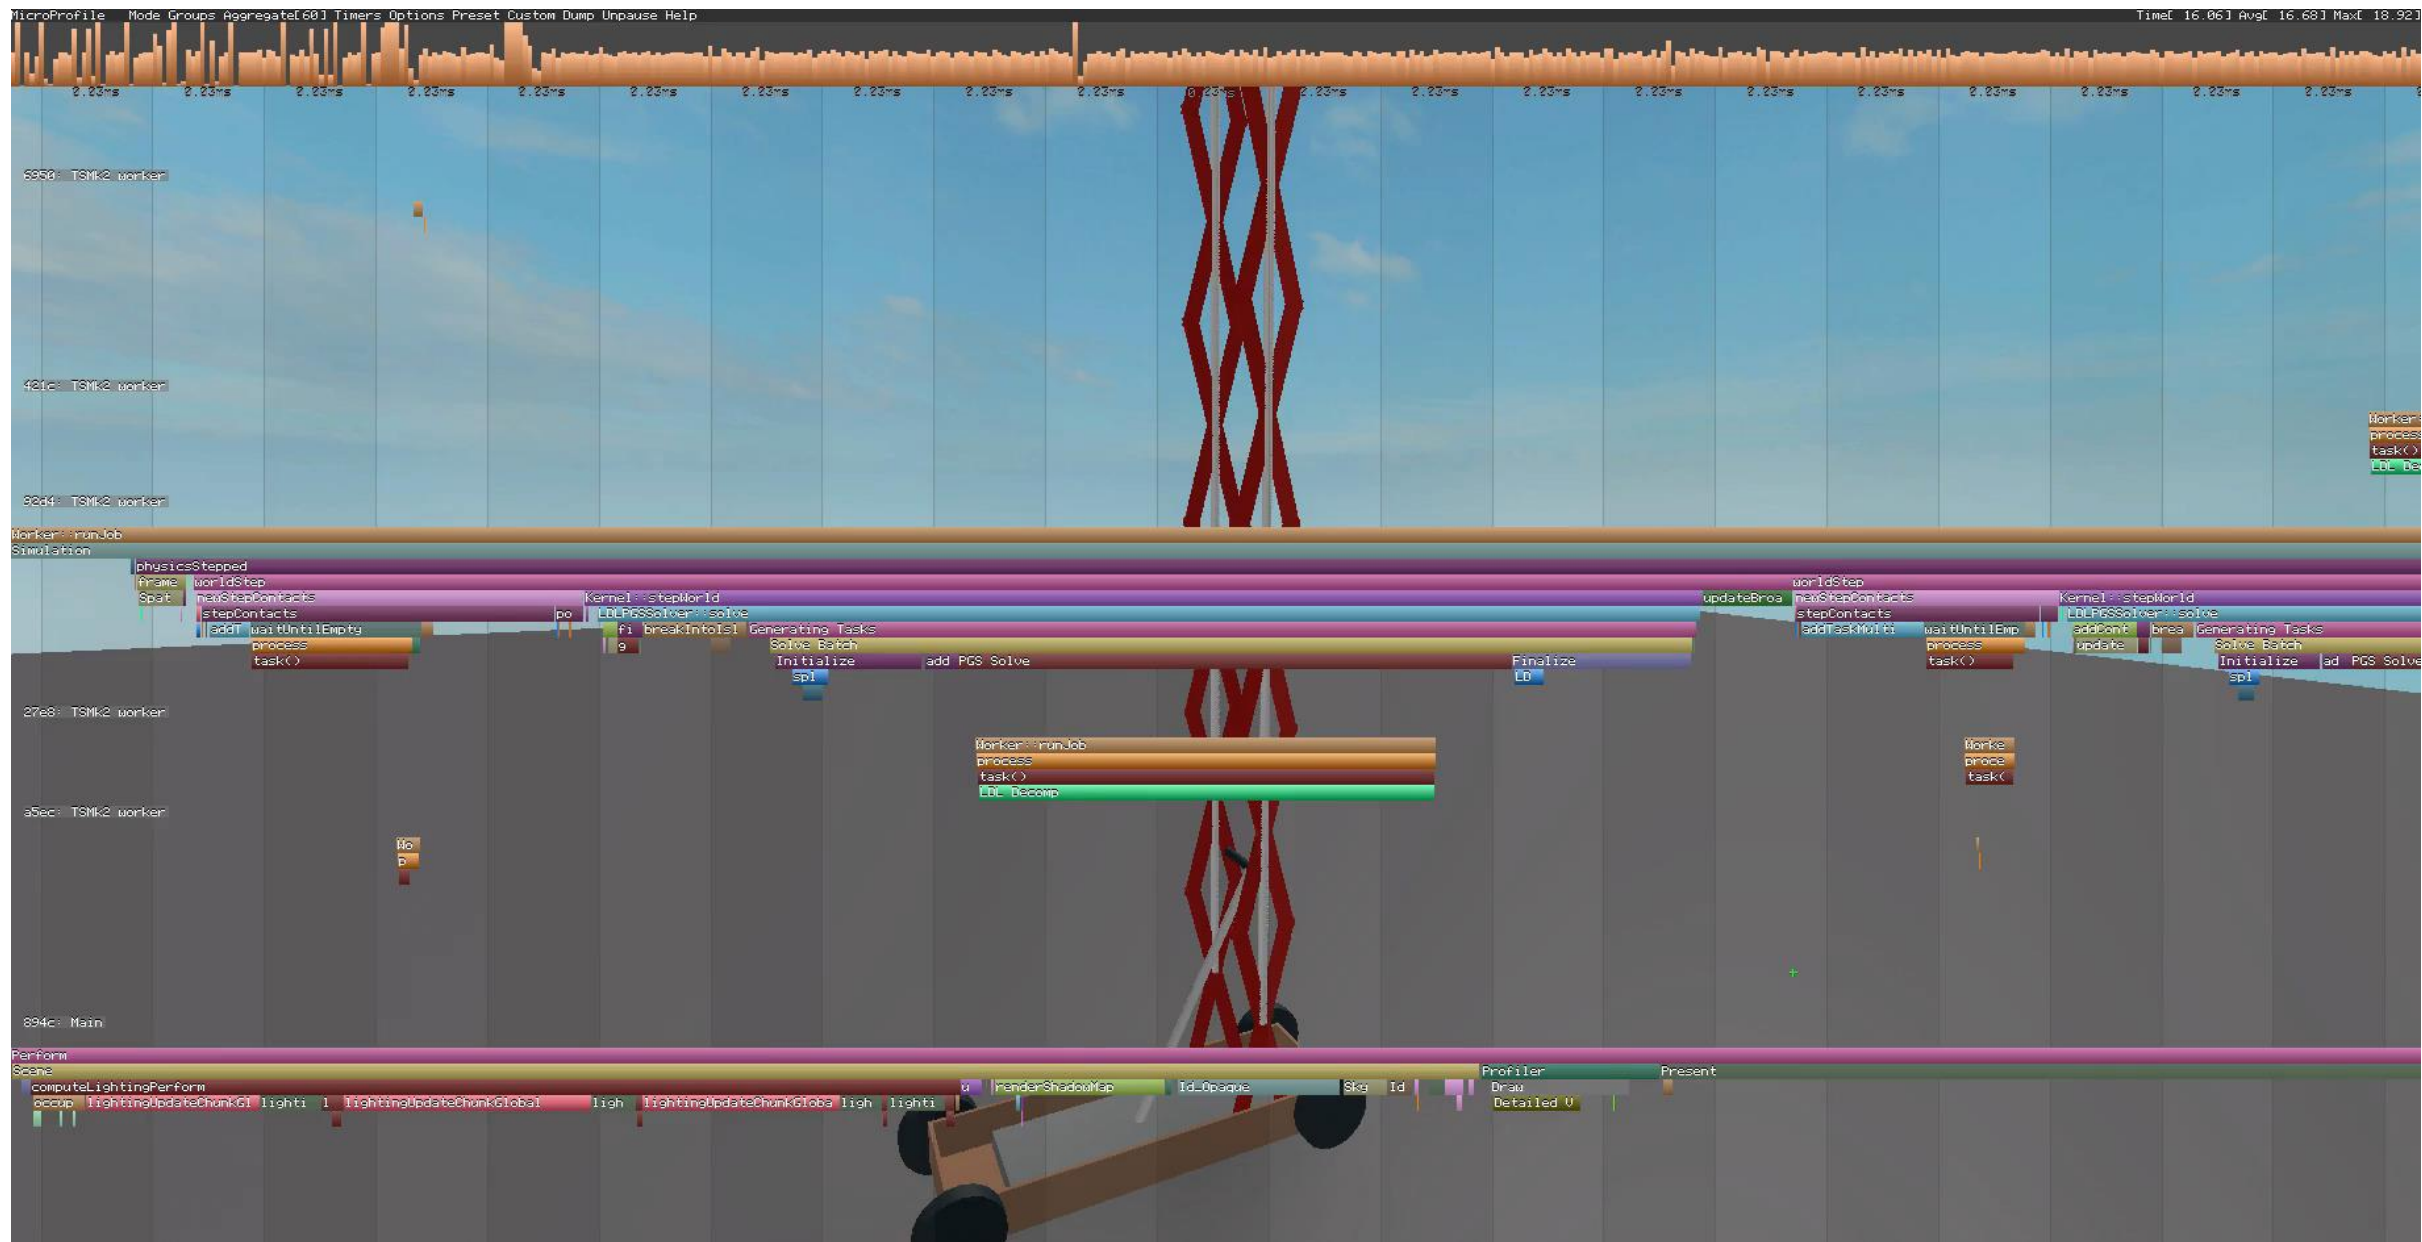
Edge Graph

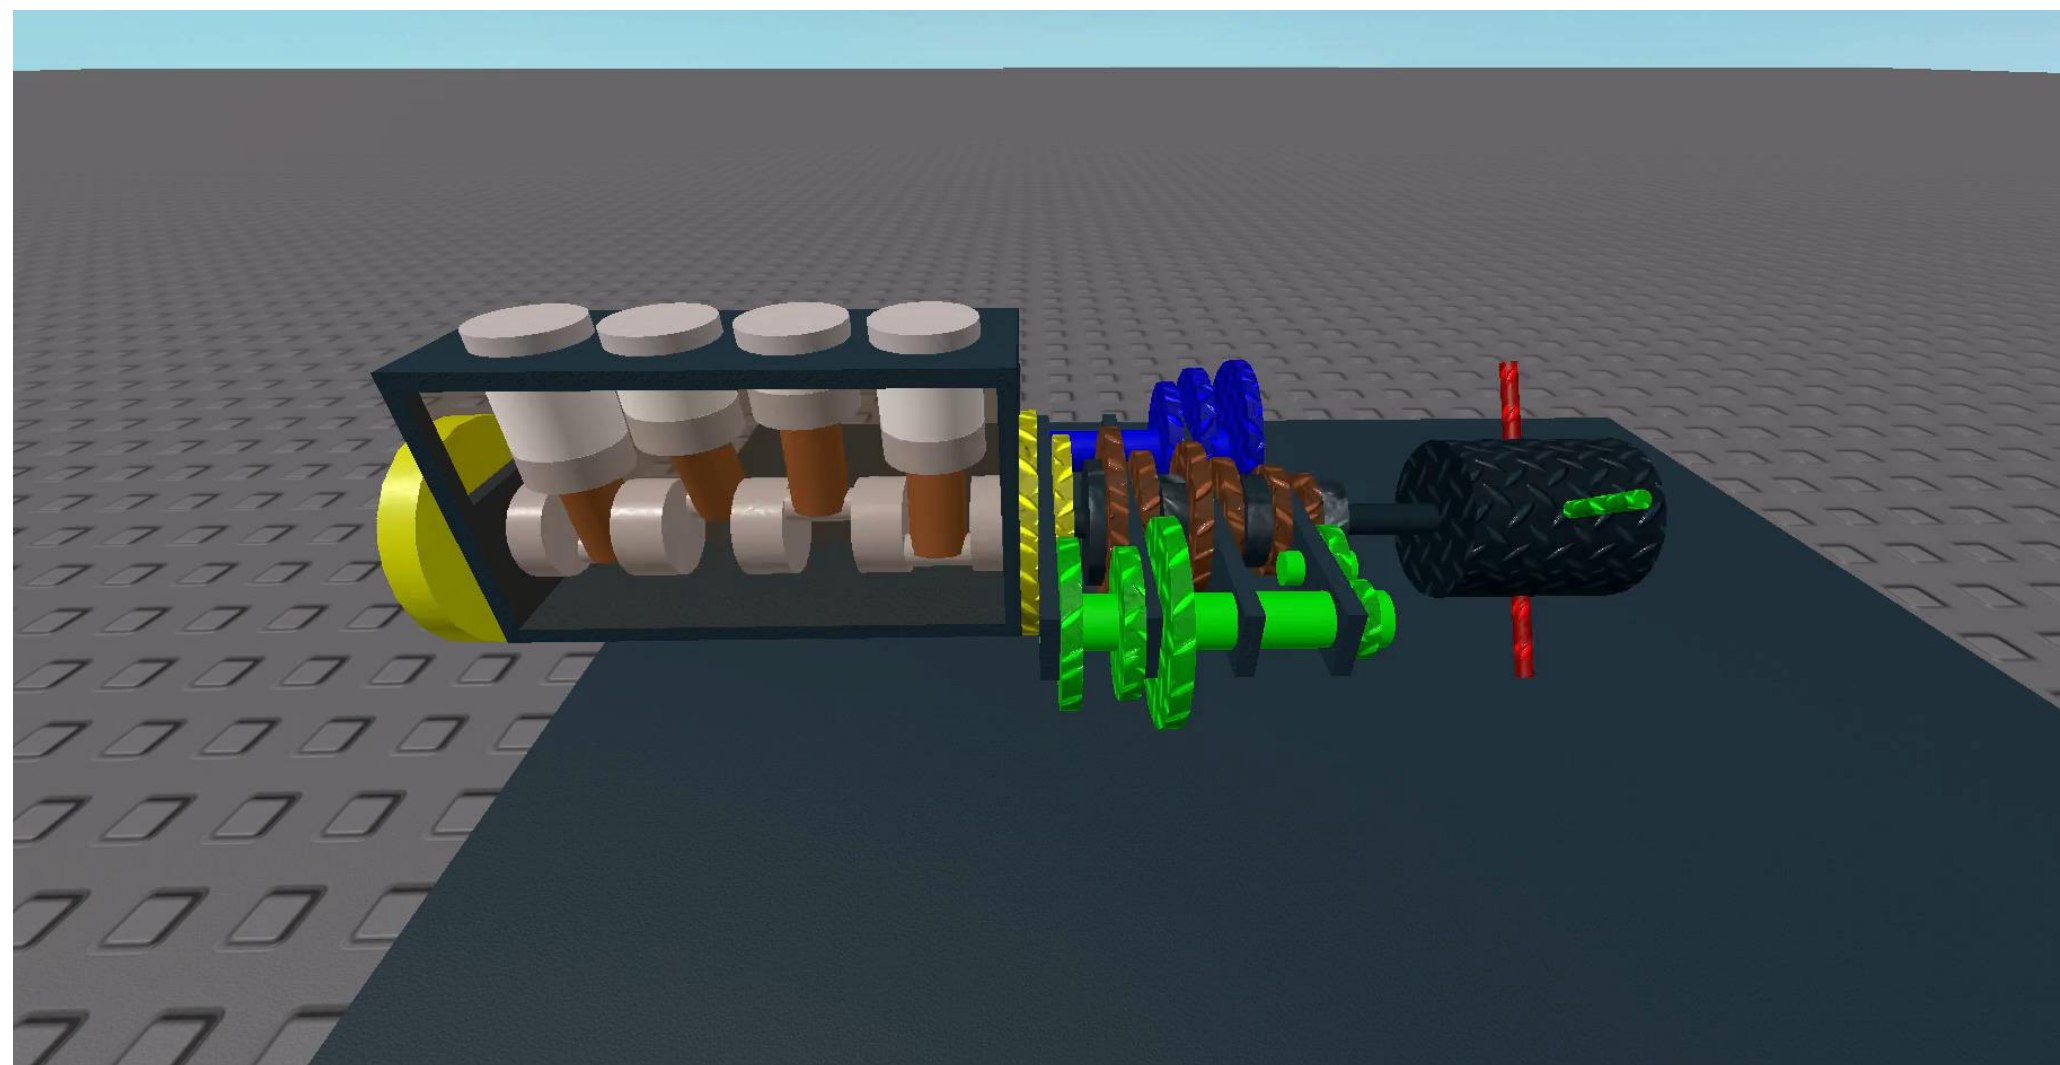


More constraints but sparser!









# References

- Tim Davis, “Direct Methods for Sparse Linear Systems”
- Tim Davis’s Lectures on YouTube
- Kenny Erleben, “Physics Based Animation”
- Bullet Physics Forum

Many thanks to Roblox and the Simulation Team for caring about awesome physics!