

# Accelerating High-Order Mesh Optimization Using Finite Element Partial Assembly

## ICOSAHOM

14-18 August 2023



Ketan Mittal

J. S. Camier, V. Dobrev, P. Knupp, T. Kolev, V. Tomov



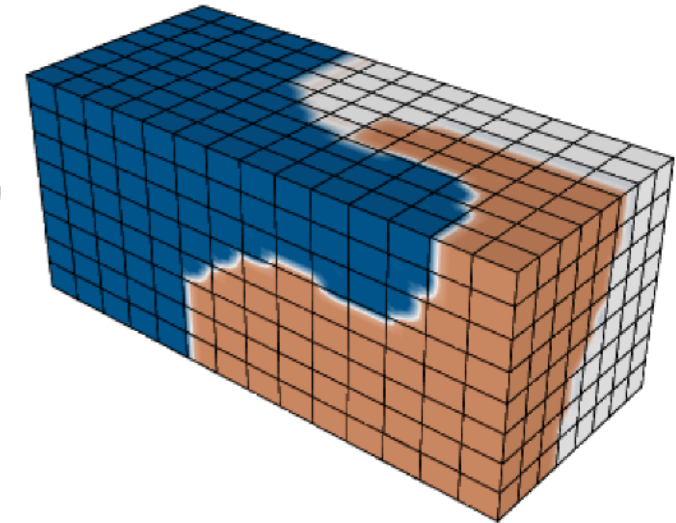
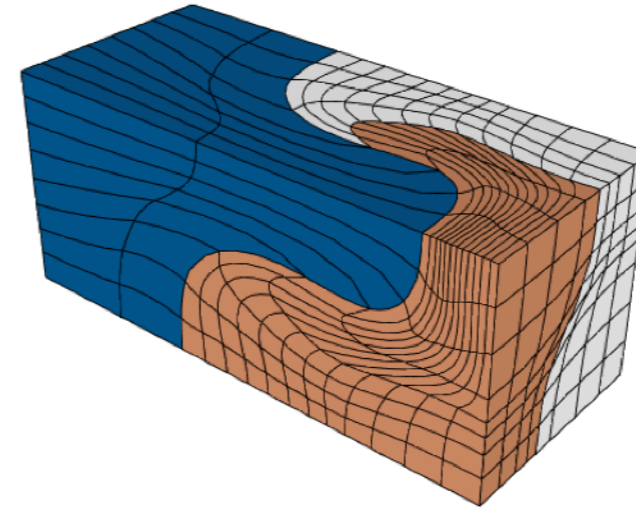
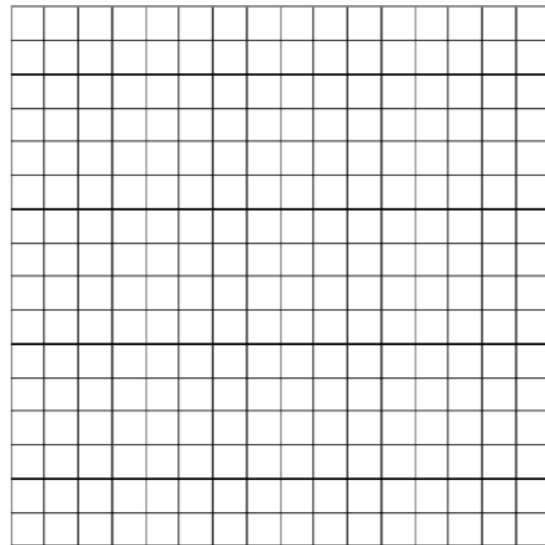
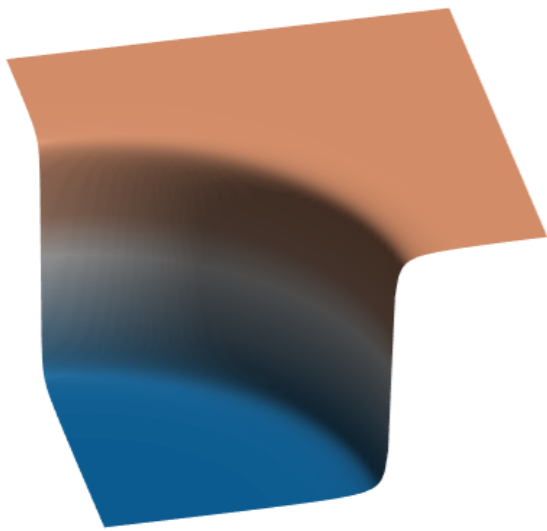
# Overview

---

- High-order meshes
- Target Matrix Optimization Paradigm (TMOP) for  $r$ -adaptivity
- Acceleration through Partial Assembly
- Acceleration through Metric Linearization
- Future work

# High-Order Mesh Optimization

- Why mesh optimization?



Outward propagating shock-wave

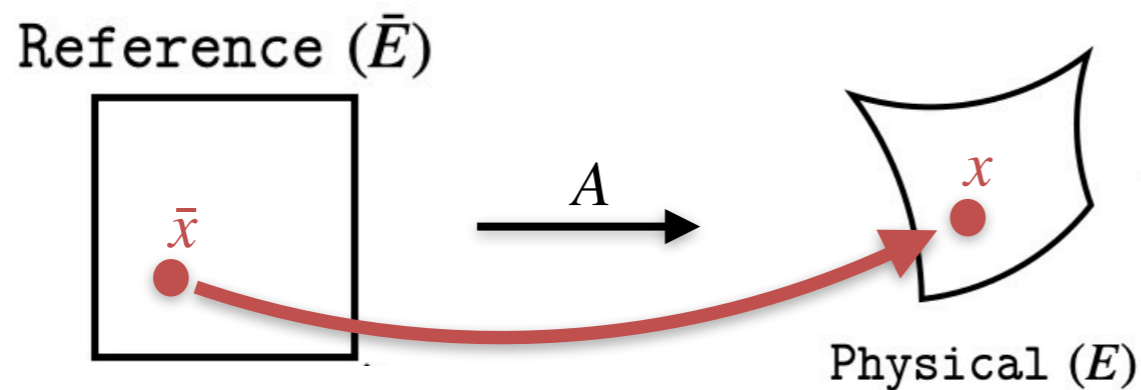
Multi-material Lagrangian Hydrodynamics

- Mesh optimization can help adapt the mesh to the solution and ultimately reduce error.
- Improve conditioning of the resulting system.

# High-Order Mesh Representation

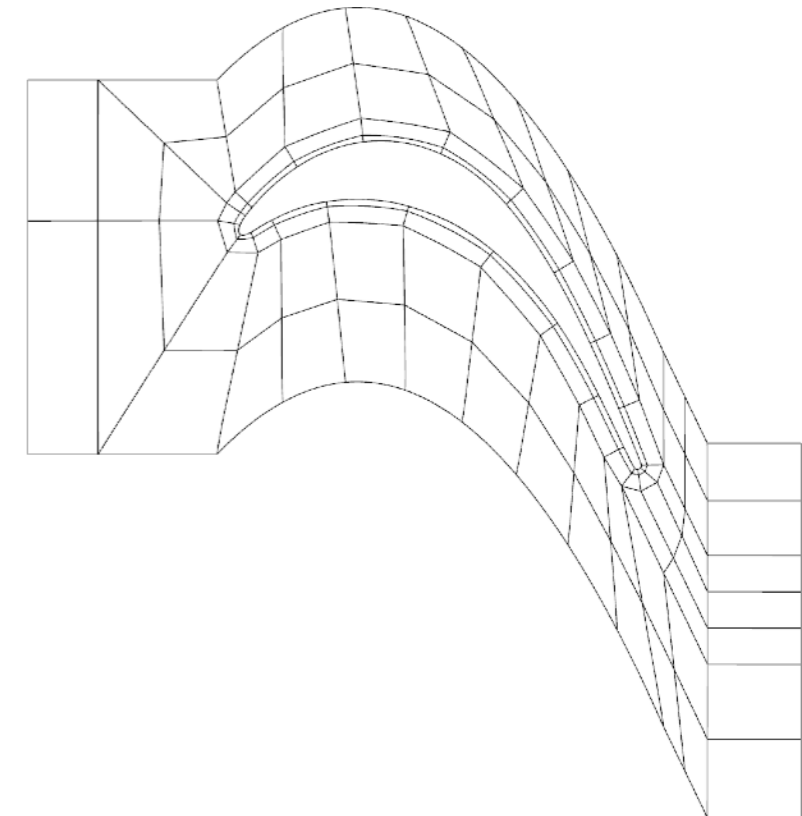
- Each element in the mesh is represented using scalar basis functions  $\{\bar{w}_i\}_{i=1}^{N_p}$  on the reference element  $\bar{E}$ .

$$x(\bar{x}) = \Phi_E(\bar{x}) \equiv \sum_{i=1}^{N_p} \mathbf{x}_{E,i} \bar{w}_i(\bar{x}), \quad \bar{x} \in \bar{E}, \quad x = x(\bar{x}) \in E$$



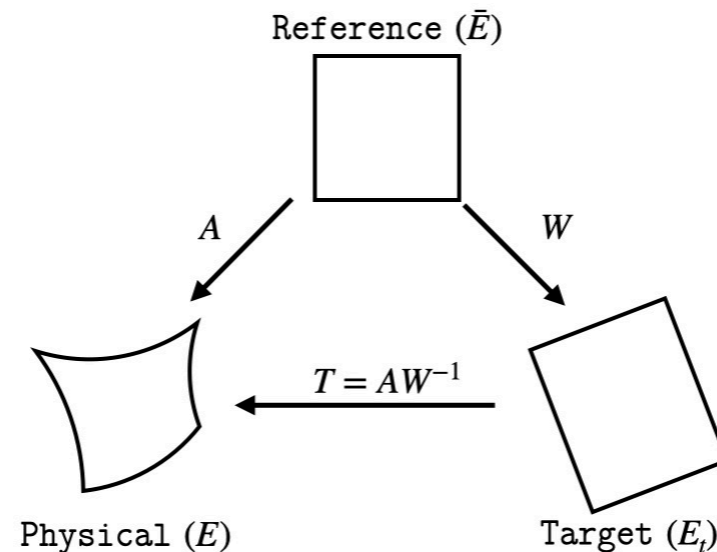
- The Jacobian of the transformation at each point represents the local deformation of the element with respect to the reference element:

$$A(\bar{x}) = \frac{\partial \Phi_E}{\partial \bar{x}} = \sum_{i=1}^{N_w} \mathbf{x}_{E,i} [\nabla \bar{w}_i(\bar{x})]^T$$



4th order mesh for a turbine blade

# Target-Matrix Optimization Paradigm (TMOP)



- Any Jacobian transformation can be represented using four geometric parameters:

$$W = \underbrace{\zeta}_{\text{[volume]}} \underbrace{R}_{\text{[rotation]}} \underbrace{Q}_{\text{[skewness]}} \underbrace{D}_{\text{[aspect-ratio]}}$$

- The transformation  $T$  from the active to target element can be defined using the Jacobian transformation  $A$  and  $W$ .

# TMOP Mesh Quality Metrics

- Quality metric  $\mu(T)$  is a measure of the deviation between the active and target Jacobian transformation.
- Different metrics depend on different geometric parameters.
  - Shape metric - depends on Skew ( $Q$ ) and Aspect-ratio ( $D$ ).  $\mu_2(T) = 0.5 \frac{|T|^2}{\det(T)} - 1$
  - Size metric - depends on  $\zeta$ .  $\mu_{77}(T) = 0.5 \left( \det(T) - \frac{1}{\det(T)} \right)^2$
  - Other kinds include Alignment, Shape + Size, Shape + Alignment, etc.
  - We typically deploy Shape + Size metrics but seldom also use Alignment metrics.
- $\mu(T)$  typically non-linear and defined such that its minima is  $T = I$  with  $\mu(I) = 0$ .

# Node Movement with TMOP

- Using the quality metric and the Jacobian transformation  $T$ , the TMOP objective function is defined as:

$$F(\mathbf{x}) = \sum_{E \in \mathcal{M}} F_E(\mathbf{x}_E) = \sum_{E(\mathbf{x}_E)} \int_{E_t} \mu(T(\mathbf{x})) d\mathbf{x}_t$$

where  $\mathbf{x}$  represents mesh coordinates. The element-by-element integral is computed as:

$$\sum_{E \in \mathcal{M}} \int_{E_t} \mu(T(\mathbf{x}_t)) d\mathbf{x}_t = \frac{1}{N_E} \sum_{E \in \mathcal{M}} \sum_{\mathbf{x}_q \in E_t} w_q \det(W(\bar{\mathbf{x}}_q)) \mu(T(\mathbf{x}_q))$$

- In practice, we can use multiple metrics with different spatial weights.

# Node Movement with TMOP

- $r$ -adaptivity -  $F(\mathbf{x})$  is minimized using a technique such as the Newton's method to optimize the mesh.

- Node movement direction: 
$$\Delta \mathbf{x}_k = - \underbrace{[\partial^2 F(\mathbf{x}_k)]^{-1}}_{\mathcal{H}(\mathbf{x}_k)} \underbrace{\partial F(\mathbf{x}_k)}_{\mathcal{J}(\mathbf{x}_k)}$$

- Solution using MINRES with preconditioning:  $\mathcal{H}(\mathbf{x}_k)\Delta \mathbf{x}_k = \mathcal{J}(\mathbf{x}_k) \leftrightarrow Ay = b$ .

- Newton update with line-search:  $\mathbf{x}_{k+1} = \mathbf{x}_k - \alpha \Delta \mathbf{x}_k$ .  $\alpha$  is backtracked starting from 1.0 until:

- $F(\mathbf{x}_{k+1}) \leq 1.2F(\mathbf{x}_k)$

- $|\mathcal{J}(\mathbf{x}_{k+1})| \leq 1.2|\mathcal{J}(\mathbf{x}_k)|$

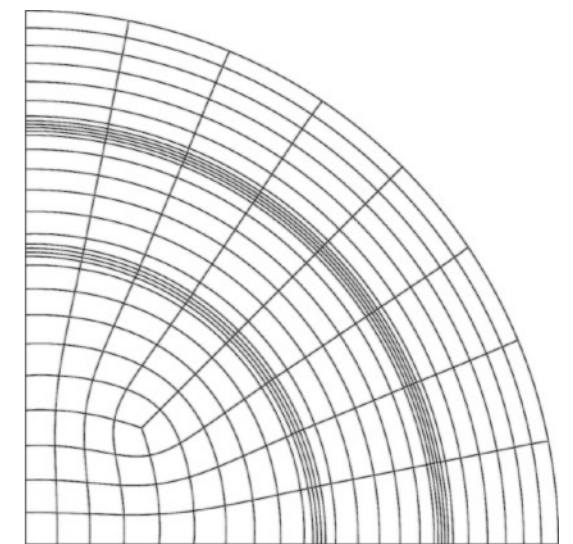
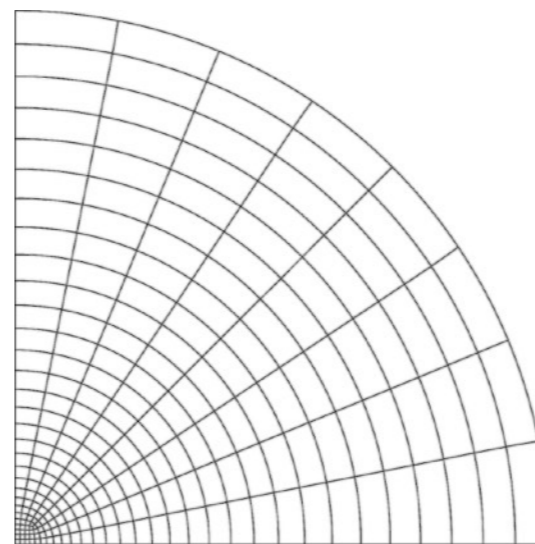
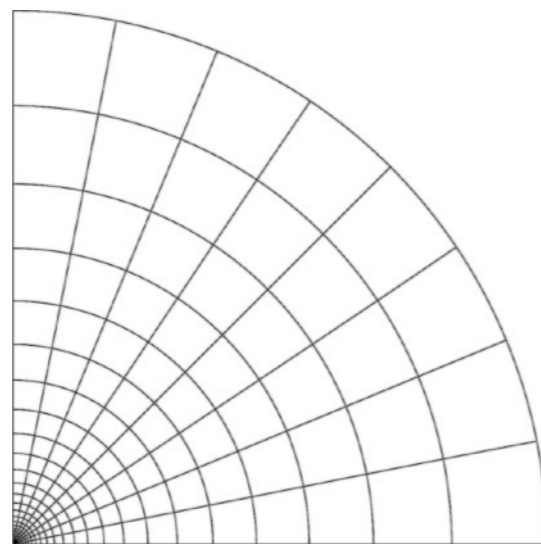
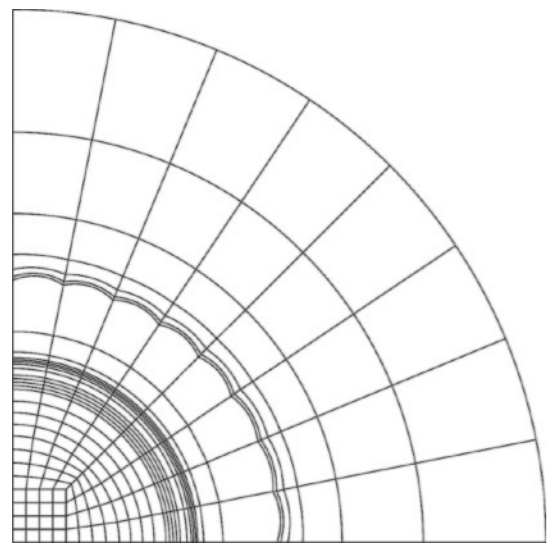
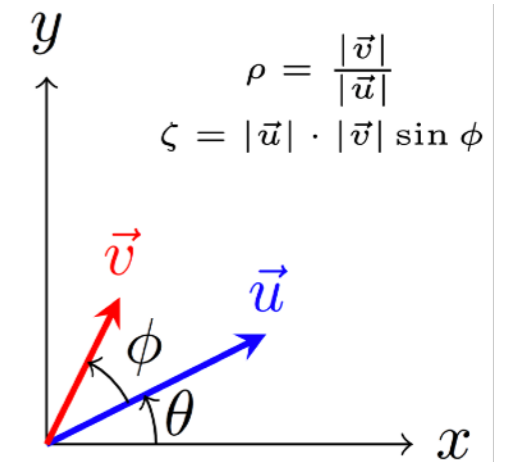
- $\min(\det(A(\mathbf{x}_{k+1}))) > \epsilon_{\det}$ .



# Geometric $r$ -adaptivity

- TMOP for  $r$ -adaptivity:

$$W = \sqrt{\frac{\zeta}{\sin \phi}} \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} 1 & \cos \phi \\ 0 & \sin \phi \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{\rho}} & 0 \\ 0 & \sqrt{\rho} \end{bmatrix}$$



Original 2nd order mesh

Ideal shape + shape-metric.

Ideal shape, equal size +  
shape-metric.

Ideal shape, spatially varying  
size + shape+size -metric.

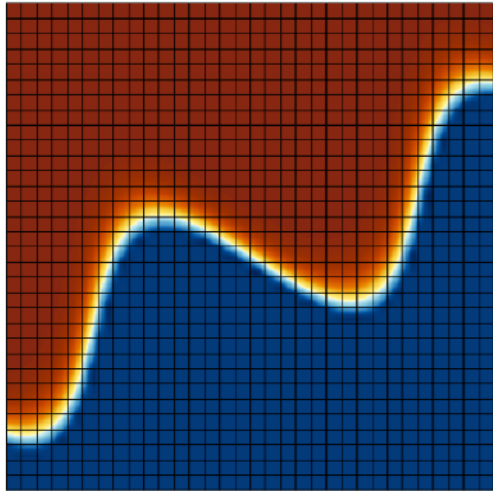
$$\phi = \frac{\pi}{2}, \rho = 1, \mu_{Sh}(T)$$

$$\zeta = \frac{\mathcal{V}}{N_E}, \phi = \frac{\pi}{2}, \rho = 1, \mu_{ShSz}(T)$$

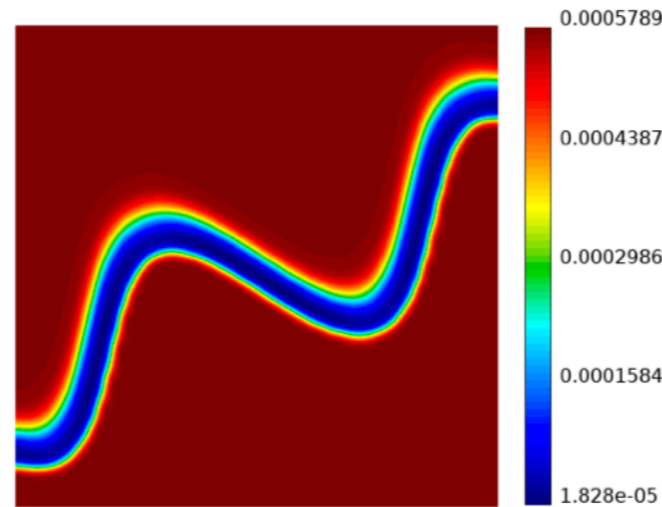
$$\zeta(\mathbf{x}), \phi = \frac{\pi}{2}, \rho = 1, \mu_{ShSz}(T)$$

"The target-matrix optimization paradigm for high-order meshes." *SIAM Journal on Scientific Computing* (2019).

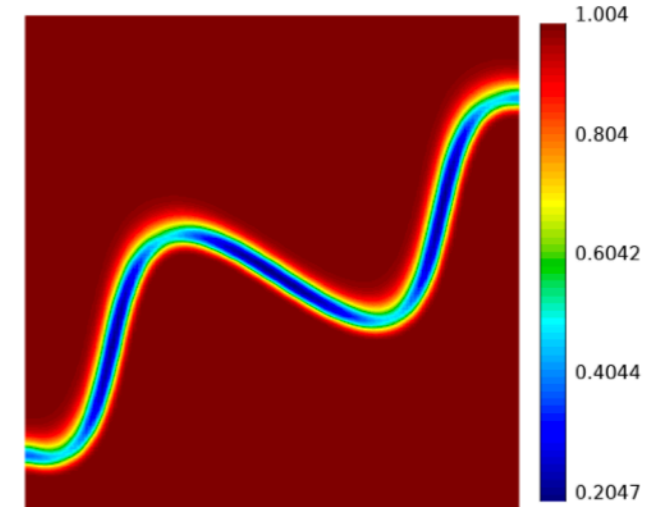
# Simulation-driven Adaptivity



Simulation data material indicator ( $\eta$ )



Size -  $\zeta \propto 1/|\nabla\eta|$



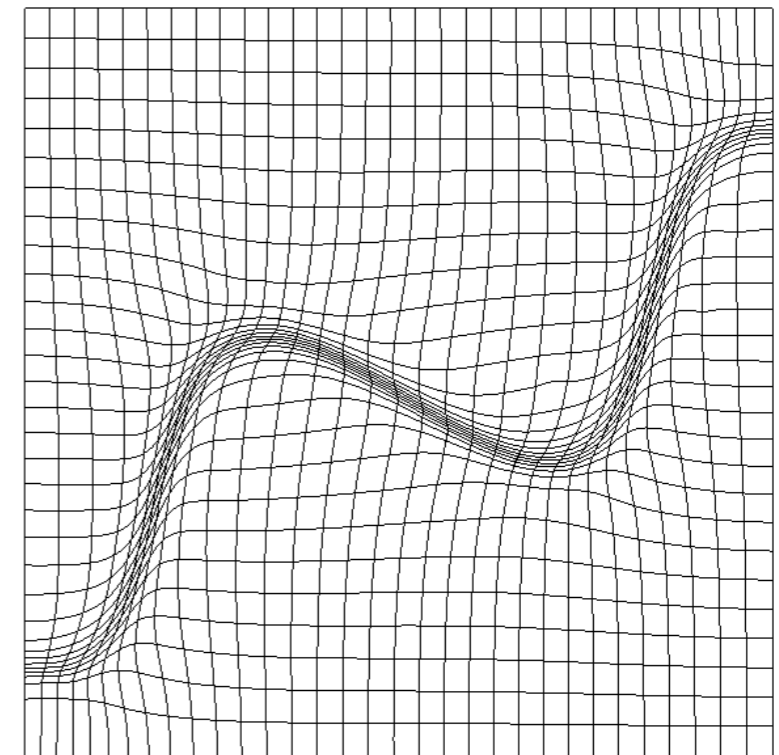
Aspect-Ratio -  $\rho \propto |\eta_x/\eta_y|$

$$W = \sqrt{\frac{\zeta}{\sin \phi}} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & \cos \phi \\ 0 & \sin \phi \end{bmatrix} \begin{bmatrix} \frac{1}{\sqrt{\rho}} & 0 \\ 0 & \sqrt{\rho} \end{bmatrix},$$

- $\phi = \frac{\pi}{2}$  for an ideal square.
- Use a Shape + Size polyconvex metric,  $\mu_{80} = (1 - \gamma)\mu_2 + \gamma\mu_{77}$ .

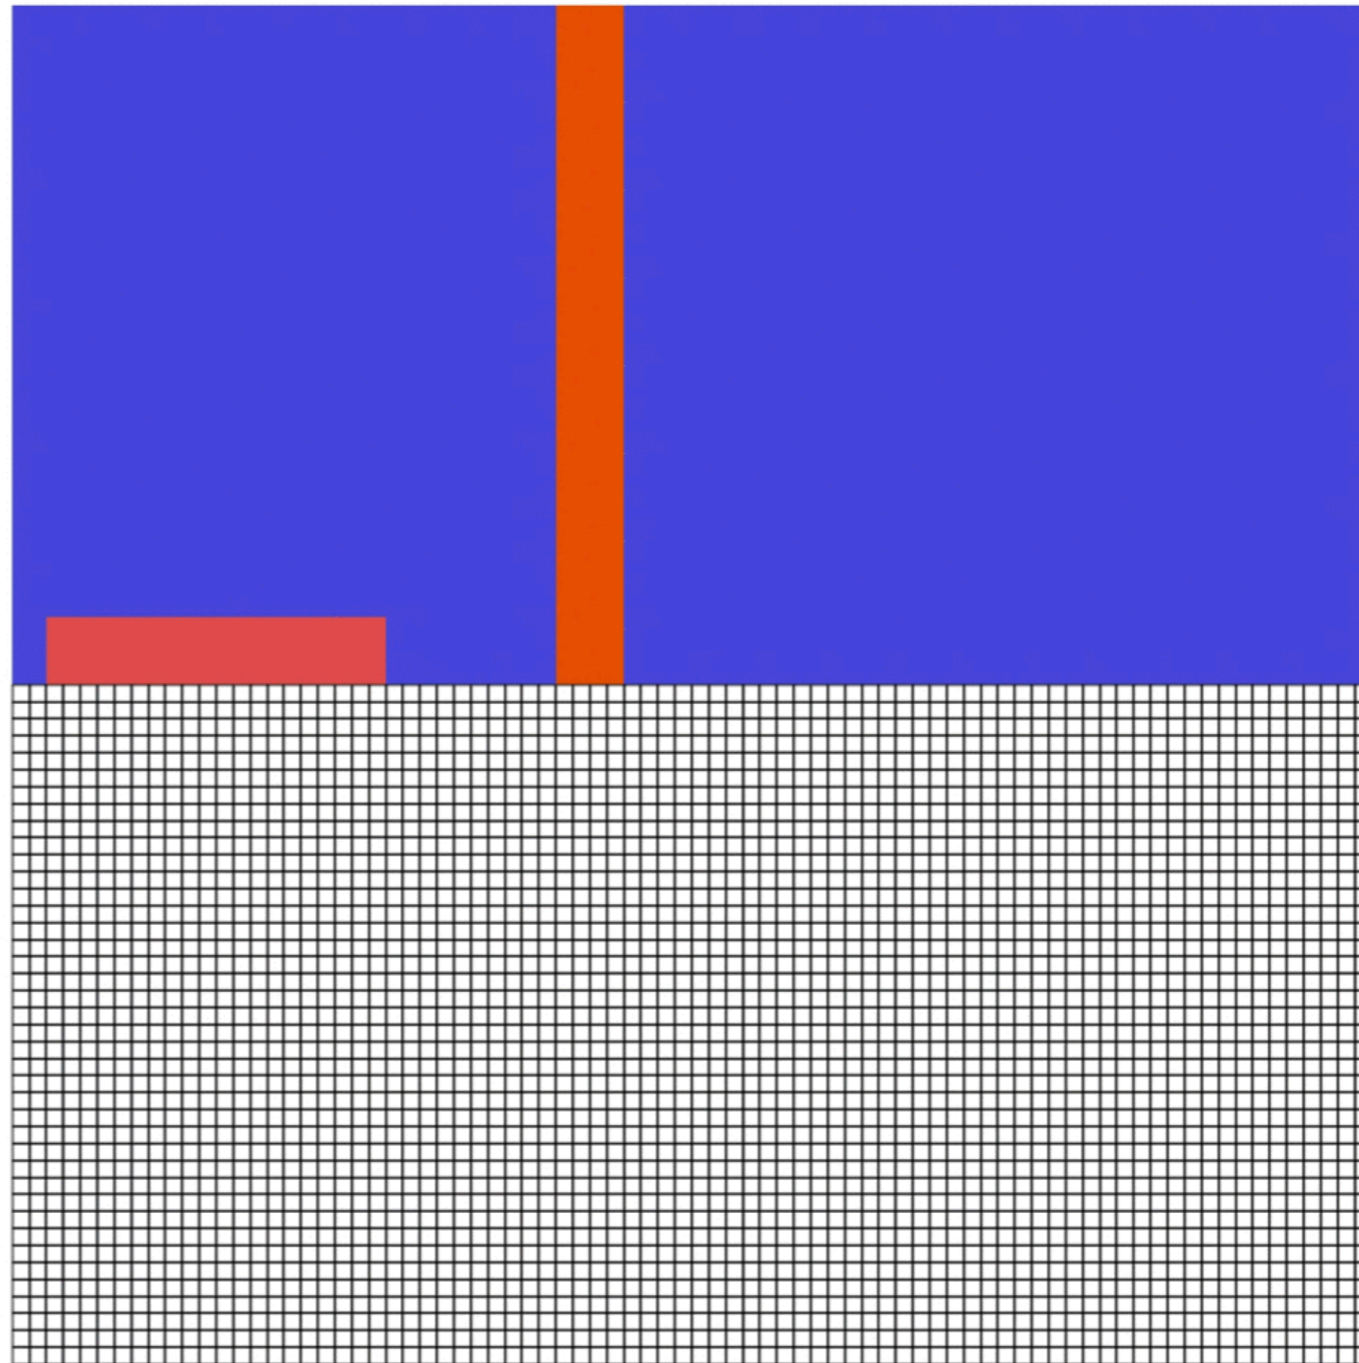
$$\mu_2(T) = 0.5 \frac{|T|^2}{\det(T)} - 1 \quad \mu_{77}(T) = \frac{1}{2} \left( \tau - \frac{1}{\tau} \right)^2$$

- Note:  $\eta$  must be remapped between and after Newton iterations.



"Simulation-driven optimization of high-order meshes in ALE hydrodynamics." Computers & Fluids (2020).

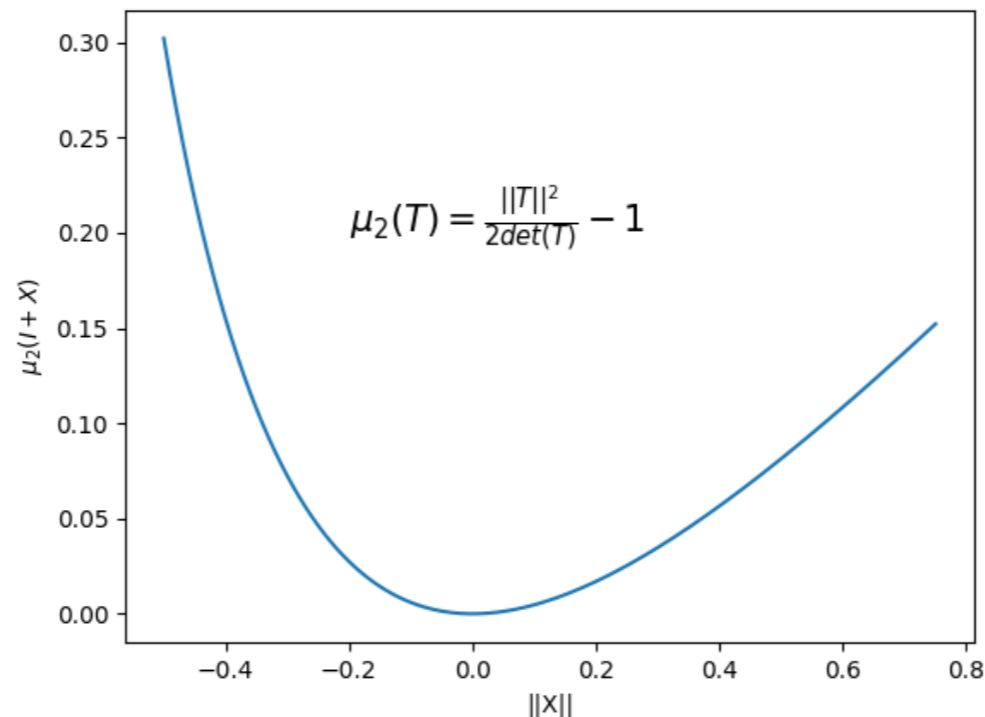
# Simulation-driven Adaptivity



Simulation-driven adaptivity with TMOP for multi-material ALE.

# Computational Cost of High-Order Mesh Optimization

- Nonlinear mesh quality metric  $\rightarrow$  problem is not quadratic so need more than 1 Newton iteration.
  - Each Newton iteration requires assembly of matrix  $\mathcal{H}$ .
  - Each Newton iteration has  $O(10-100)$  MINRES iterations where we do matrix-vector products using  $\mathcal{H}$ .



# Computational Cost of High-Order Mesh Optimization

- Minimizing the TMOP objective function entails solving  $\mathcal{H}(\mathbf{x}_k)\Delta\mathbf{x}_k = \mathcal{J}(\mathbf{x}_k)$ .

$$F(\mathbf{x}) = \sum_{E \in \mathcal{M}} F_E(\mathbf{x}_E) = \sum_{E(\mathbf{x}_E)} \int_{\bar{E}} \mu(T(\mathbf{x})) d\bar{\mathbf{x}}$$

- Assume  $W = I$ :

$$\frac{\partial F(\mathbf{x})}{\partial x_{a,i}} = \int_{\bar{E}} \frac{\partial \mu}{\partial T(\mathbf{x})} : \frac{\partial T(\mathbf{x})}{\partial x_{a,i}} d\bar{\mathbf{x}} = \int_{\bar{E}} \frac{\partial \mu}{\partial T} : \left( \frac{\partial A}{\partial x_{a,i}} \right) d\bar{\mathbf{x}}$$

$$a = 1, \dots, d, \quad i = 1, \dots, N_x$$

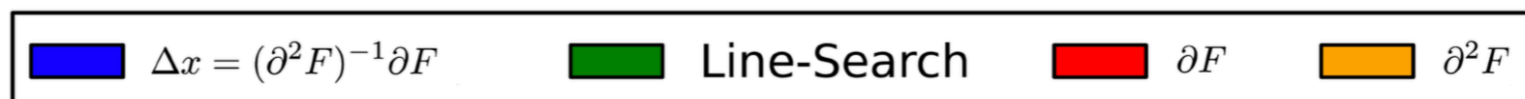
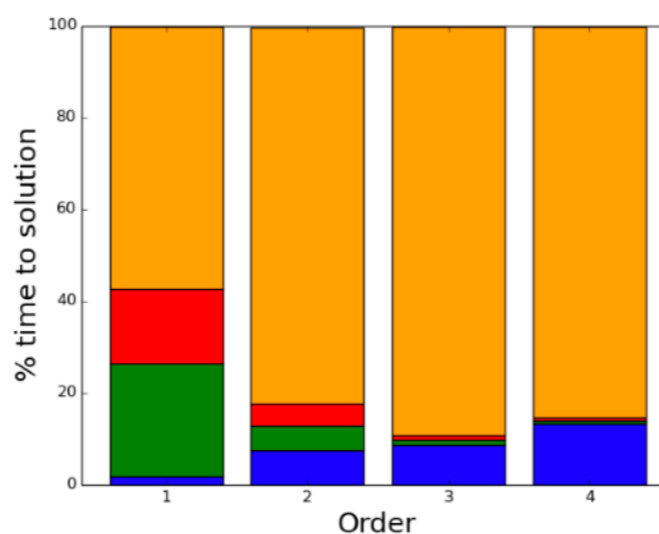
$$\frac{\partial^2 F(\mathbf{x})}{\partial x_{b,j} \partial x_{a,i}} = \int_{\bar{E}} \frac{\partial}{\partial x_{b,j}} \left[ \frac{\partial \mu}{\partial T} : \left( \frac{\partial A}{\partial x_{a,i}} \right) \right] d\bar{\mathbf{x}}$$

$$a, b = 1, \dots, d, \quad i, j = 1, \dots, N_x$$

# Computational Cost of High-Order Mesh Optimization

- Simple implementation (Full assembly):
  - Construct and store the global matrix  $\mathcal{H}$  for the entire mesh at each Newton iteration.
  - Easy to setup but computationally expensive (prohibitive as  $p$  increases):

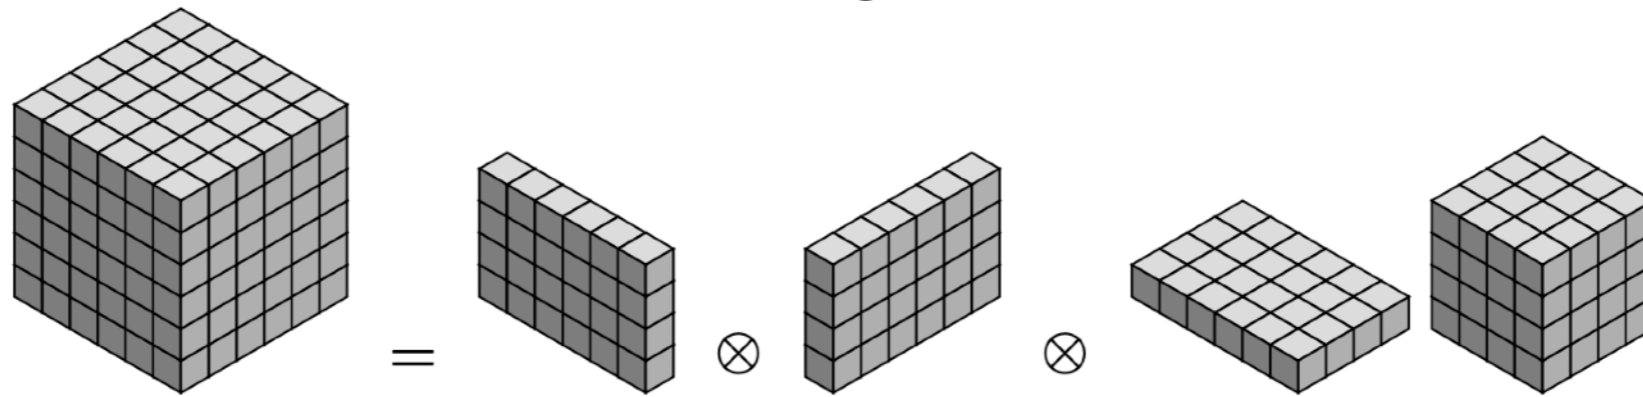
Method	Storage	Assembly	Evaluation
Traditional full assembly + matvec	$\mathcal{O}(p^{2d})$	$\mathcal{O}(p^{3d})$	$\mathcal{O}(p^{2d})$



# Partial Assembly

- Acceleration for tensor-product elements (quads/hexes) using partial-assembly and matrix-free evaluation.

- Construct  $nD$  Operators as a Kronecker product ( $\otimes$ ) of 1D operators.



$$B_{3D} = (B_{1D} \otimes B_{1D} \otimes B_{1D})u$$

- Store only locally assembled 1D matrices.
- At each iteration, store only quadrature point data and use locally stored 1D operators to perform element-by-element matrix-vector products.

Method	Storage	Assembly	Evaluation
Traditional full assembly + matvec	$\mathcal{O}(p^{2d})$	$\mathcal{O}(p^{3d})$	$\mathcal{O}(p^{2d})$
Partial assembly + matrix-free action	$\mathcal{O}(p^d)$	$\mathcal{O}(p^d)$	$\mathcal{O}(p^{d+1})$

- Well suited for GPUs due to low storage complexity and matrix-vector products.

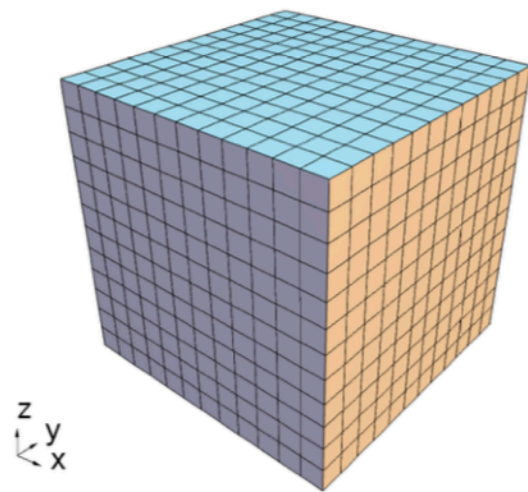
# Kershaw Benchmark

- Easy-to-setup benchmark for timing high-order mesh optimization.
- Two parameters  $(\epsilon_y, \epsilon_z) \in (0,1]^2$  control the element deformation.
- In our tests, we use a  $24 \times 24 \times 24$  mesh, 9 quadrature points in each direction in an element, and a shape metric  $\mu_{303} = \frac{|T|^2}{3\tau^{2/3}} - 1$ .

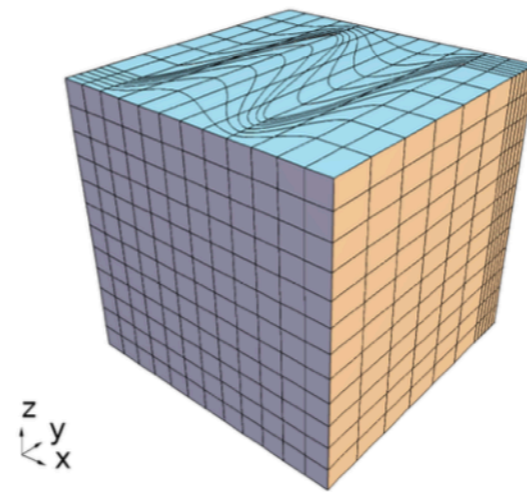
```

1 double right(const double eps, const double x) // 1D transformation at right boundary
2 {
3     return (x <= 0.5) ? (2-eps)*x : 1+eps*(x-1);
4 }
5 double left(const double eps, const double x) // 1D transformation at left boundary
6 {
7     return 1-right(eps,1-x);
8 }
9 double step(const double a, const double b, double x)
10 {
11     if (x <= 0) { return a; }
12     if (x >= 1) { return b; }
13     return a + (b-a)*(x*x*x*(x*(6*x-15)+10)); // Smooth transition from a to b
14 }
15 void kershaw(const double epsy, const double epsz,
16             const double x, const double y, const double z,
17             double &X, double &Y, double &Z) // (x,y,z) -> (X,Y,Z) Kershaw transform
18 {
19     X = x;
20     int layer = x*6.0;
21     double lambda = (x-layer/6.0)*6;
22     switch (layer)
23     {
24     case 0:
25         Y = left(epsy, y);
26         Z = left(epsz, z);
27         break;
28     case 4:
29         Y = step(left(epsy, y), right(epsy, y), lambda);
30         Z = step(left(epsz, z), right(epsz, z), lambda);
31         break;
32     case 2:
33         Y = step(right(epsy, y), left(epsy, y), lambda/2);
34         Z = step(right(epsz, z), left(epsz, z), lambda/2);
35         break;
36     case 3:
37         Y = step(right(epsy, y), left(epsy, y), (1+lambda)/2);
38         Z = step(right(epsz, z), left(epsz, z), (1+lambda)/2);
39         break;
40     default:
41         Y = right(epsy, y);
42         Z = right(epsz, z);
43         break;
44     }
45 }

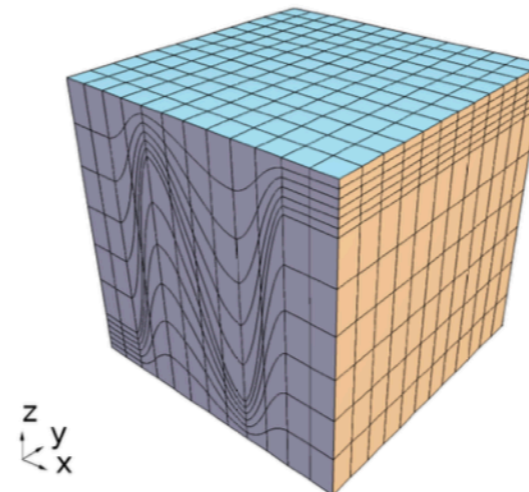
```



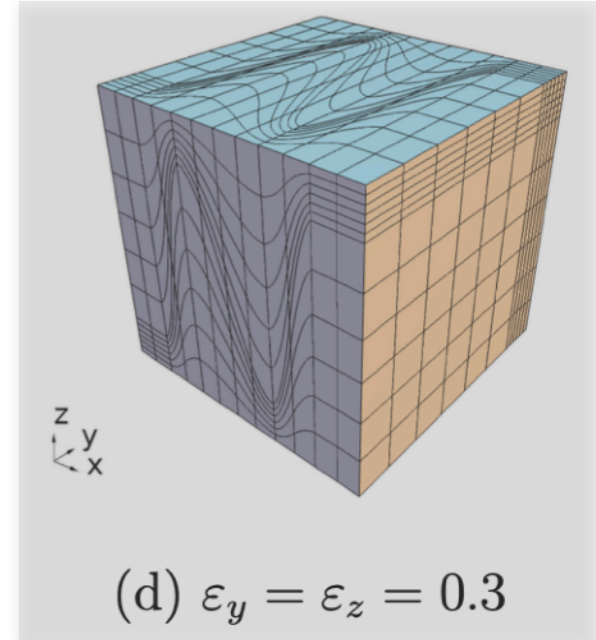
(a)  $\epsilon_y = \epsilon_z = 1$



(b)  $\epsilon_y = 0.3, \epsilon_z = 1$



(c)  $\epsilon_y = 1, \epsilon_z = 0.3$

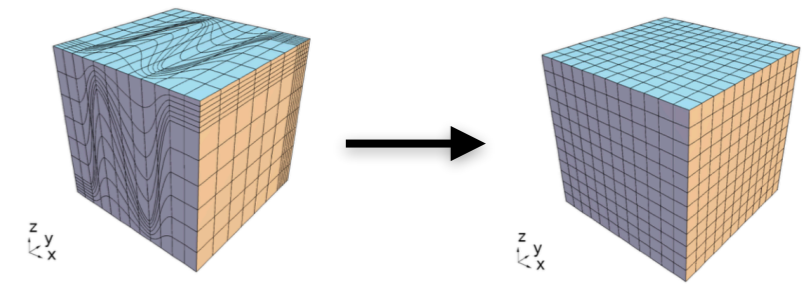


(d)  $\epsilon_y = \epsilon_z = 0.3$

Camier et al. *Accelerating high-order mesh optimization using finite element partial assembly on GPUs*. Journal of Computational Physics (2023).



# Kershaw Benchmark: Timing Results

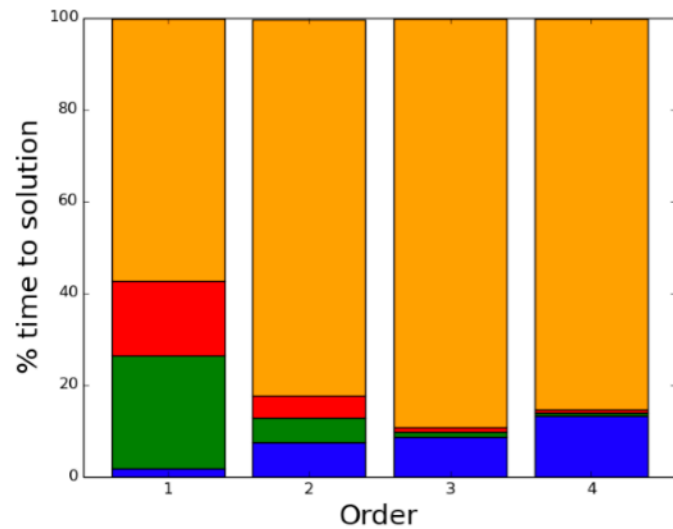
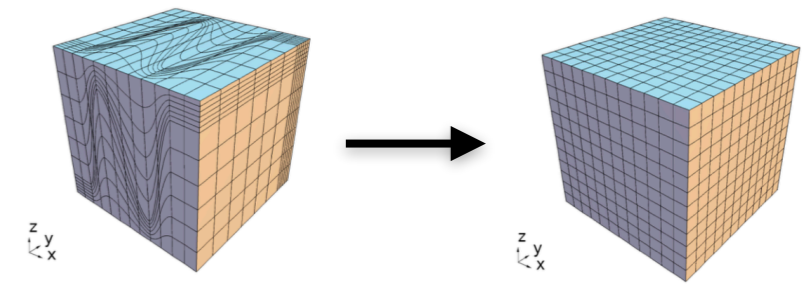


- Timing comparison on Lassen, a Livermore Computing supercomputer, for full- and partial-assembly on CPU vs partial-assembly on GPU.
  - CPU - 36 cores with 44 CPUs per core.
  - GPU - 1 core with 1 GPU and 4 CPUs per core.

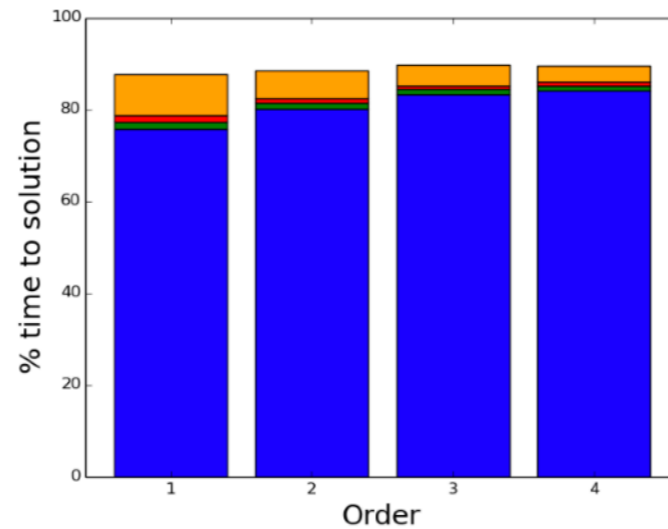
	Time to solution (sec)			
	$p = 1$	$p = 2$	$p = 3$	$p = 4$
CPU <sup>FA</sup>	2.9	31.1	489.6	2868.8
CPU <sup>PA</sup>	18.0	41.0	128.5	298.0
GPU <sup>PA</sup>	0.4	0.9	3.9	8.5
Speedup (GPU <sup>PA</sup> vs CPU <sup>PA</sup> )				
	<b>42×</b>	<b>43×</b>	<b>32×</b>	<b>35×</b>

$\mathcal{O}(30\times)$  speed-up on GPUs versus CPUs  
PA beneficial on CPUs for higher  $p$ .

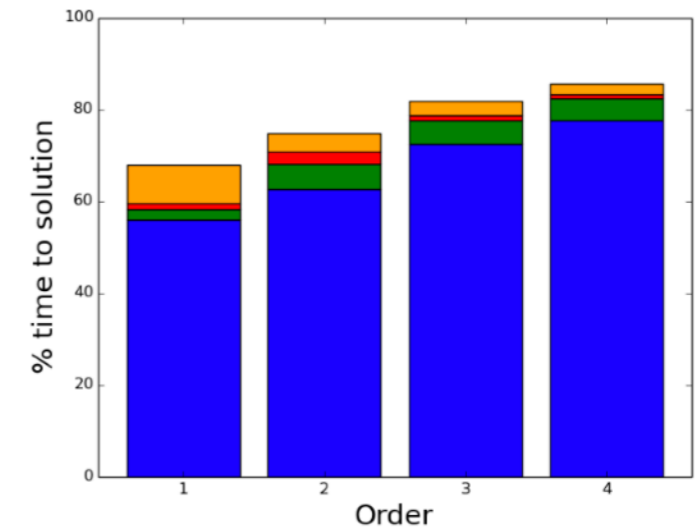
# Kershaw Benchmark: Timing Results



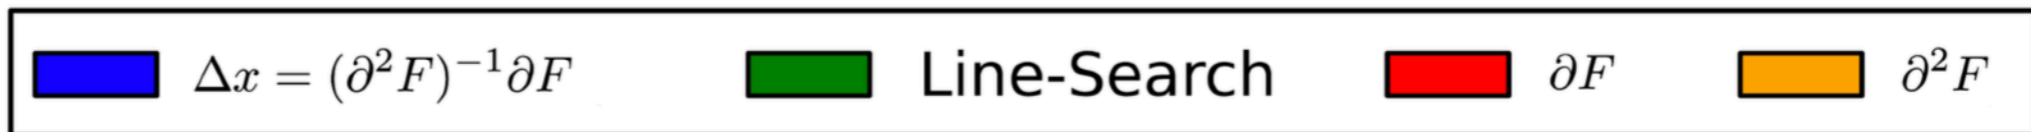
(a) CPU<sup>FA</sup>



(b) CPU<sup>PA</sup>



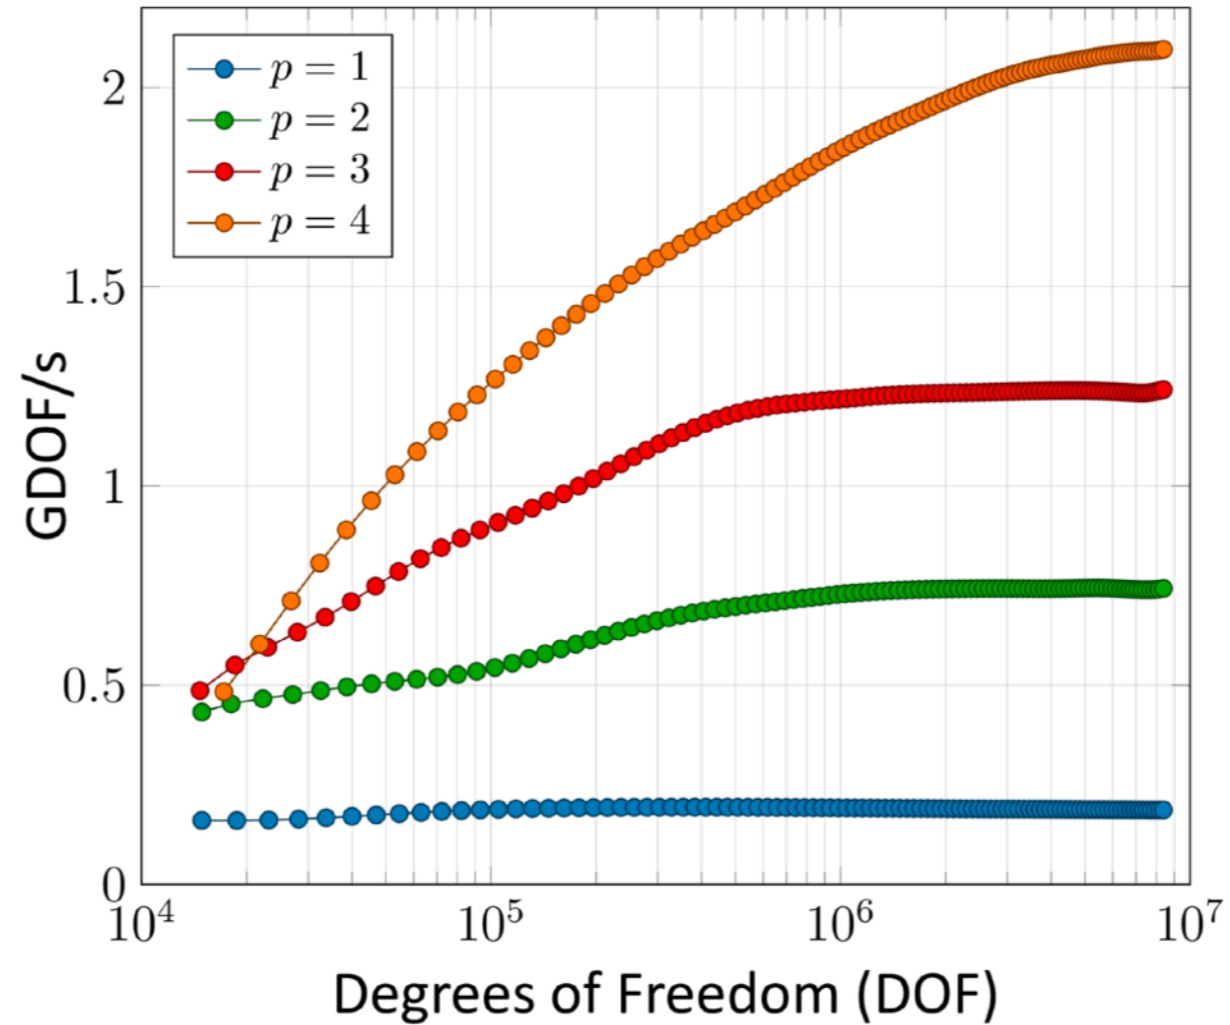
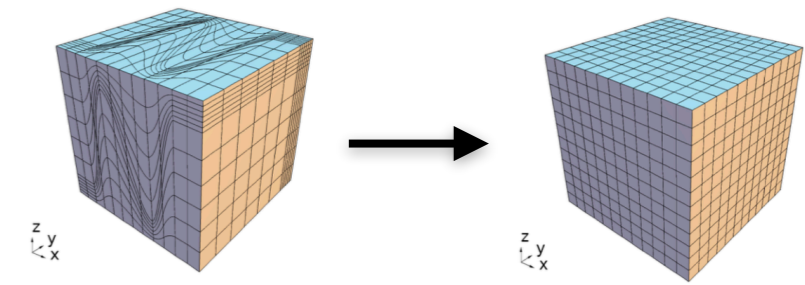
(c) GPU<sup>PA</sup>



Partial-assembly virtually eliminates the assembly cost associated with  $\mathcal{H}$ .

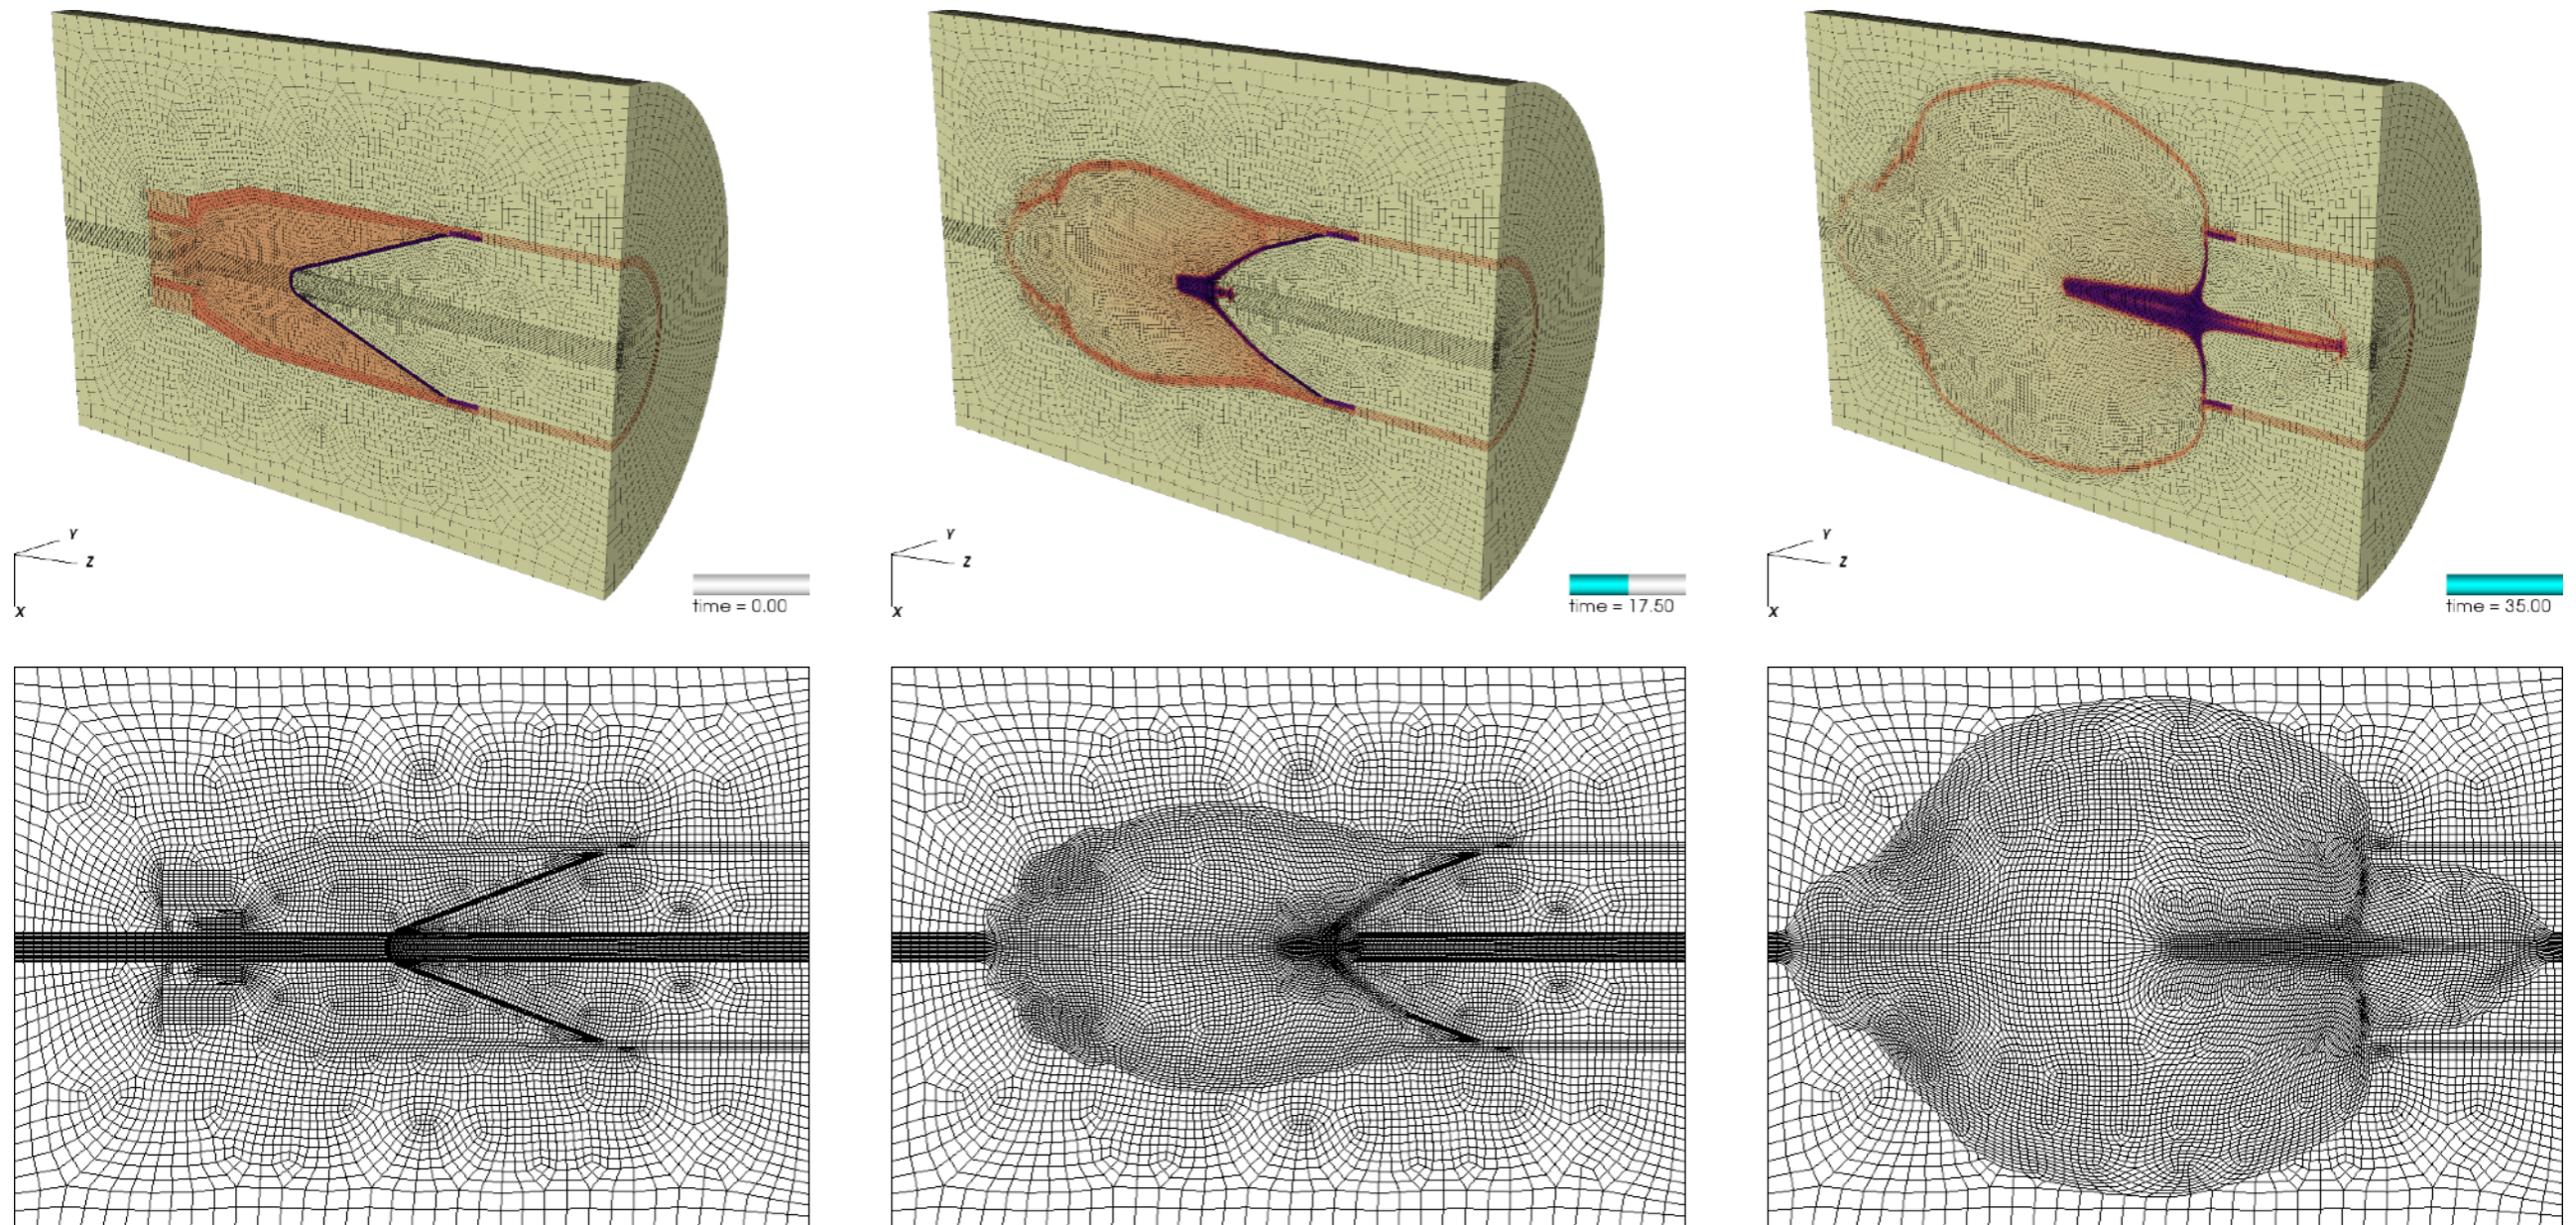
GPUs provide further acceleration.

# Kershaw Benchmark: Throuput



Throughput over 1 Newton iteration on 4 NVIDIA V100s

# GPU Acceleration for MultiMaterial ALE with Solution-Driven Adaptivity

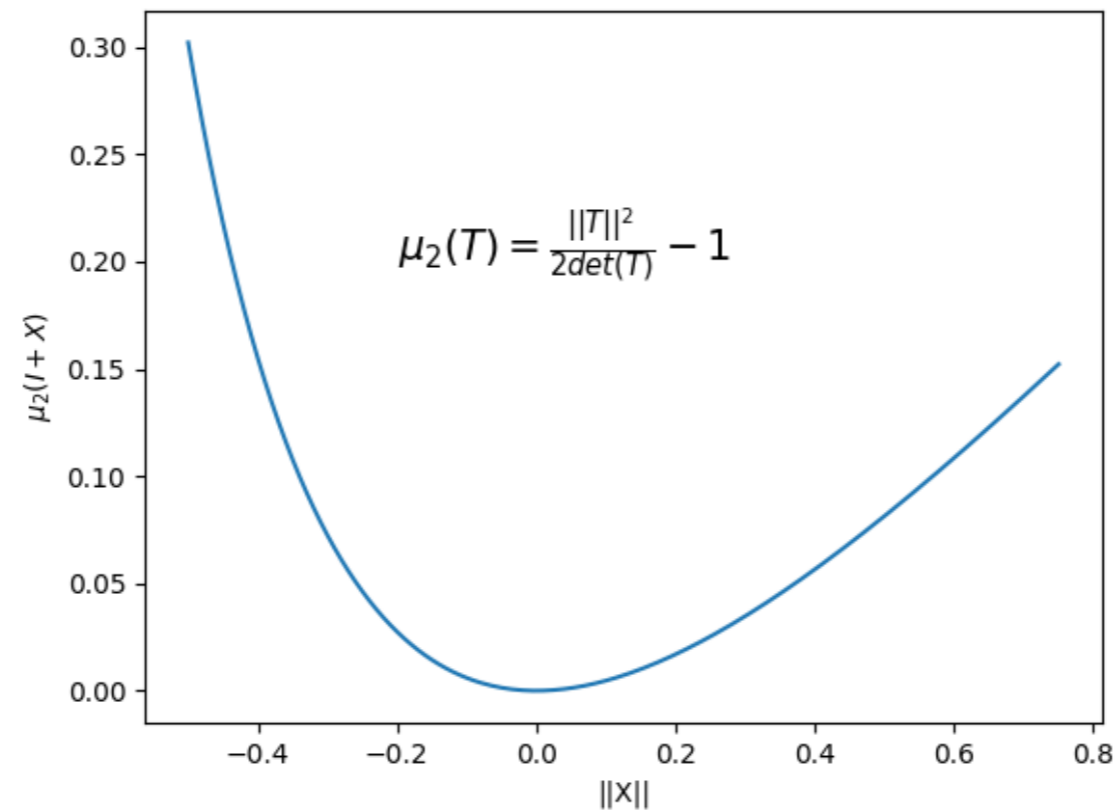


Density (top) and 2nd order mesh (bottom) for the ALE shaped charge GPU simulation.

20x speed-up for TMOP step in the solver.

# Metric Linearization

- TMOP problem requires multiple Newton iterations due to the non-linearity from the mesh quality metric.



- We can linearize the problem using Taylor expansion when:
  - The target matrix  $W$  is constant throughout the domain. [Dependence of Hessian only on  $A$ ]
  - Deviation of current jacobian  $A$  is *small* with respect to the target Jacobian  $W$ .

# Metric Linearization

- Linearize the metric around the minima  $I$  using  $T = I + X$ :

$$\mu(T) = \mu(I + X) = \mu(I) + X : \left. \frac{\partial \mu}{\partial T} \right|_{T=I} + \frac{1}{2} X : \left. \frac{\partial^2 \mu}{\partial T^2} \right|_{T=I} : X + \dots$$

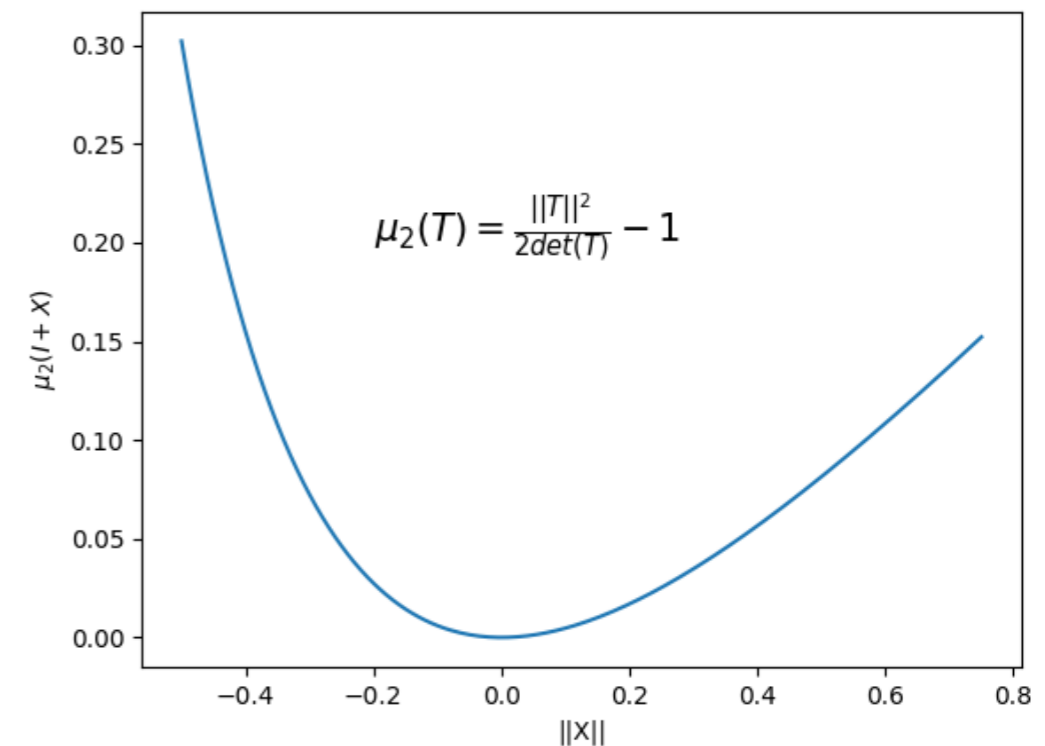
$$\mu(T) = \frac{1}{2} X : \left. \frac{\partial^2 \mu}{\partial T^2} \right|_{T=I} : X$$

$$\frac{\partial \mu}{\partial T} = X : \left. \frac{\partial^2 \mu}{\partial T^2} \right|_{T=I}$$

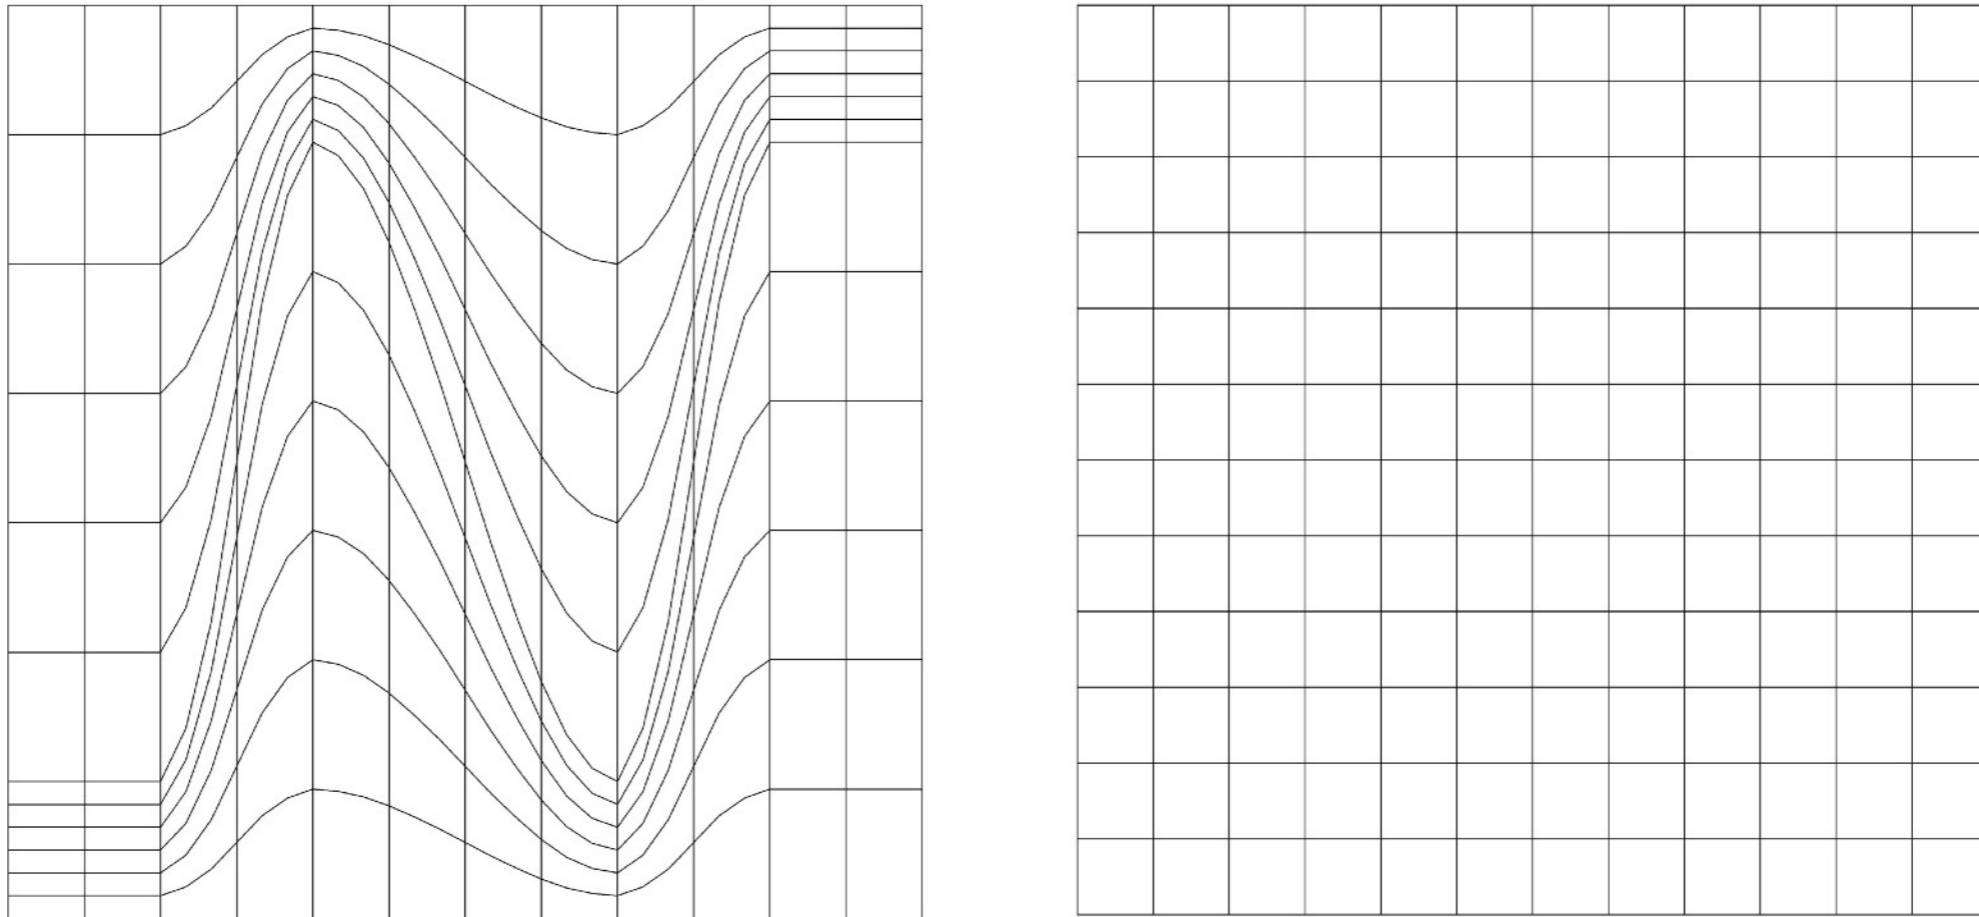
$$\frac{\partial^2 \mu}{\partial T^2} = \left. \frac{\partial^2 \mu}{\partial T^2} \right|_{T=I}$$

Compute  $\left. \frac{\partial^2 \mu}{\partial T^2} \right|_{T=I}$  once and re-use at all quadrature points.

Problem is quadratic so 1 Newton iteration is sufficient.



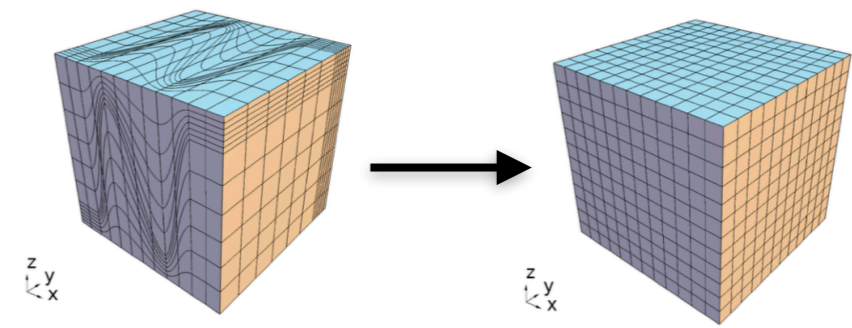
# Metric Linearization - Results



Kershaw transformed 2nd order mesh optimized using ideal shape + shape metric.

- Linearized problem requires 1 Newton iteration with 82 MINRES iterations.
- Non-linear form requires 8 Newton iterations with a total of 1457 MINRES iterations.

# Metric Linearization + GPU acceleration - Kershaw



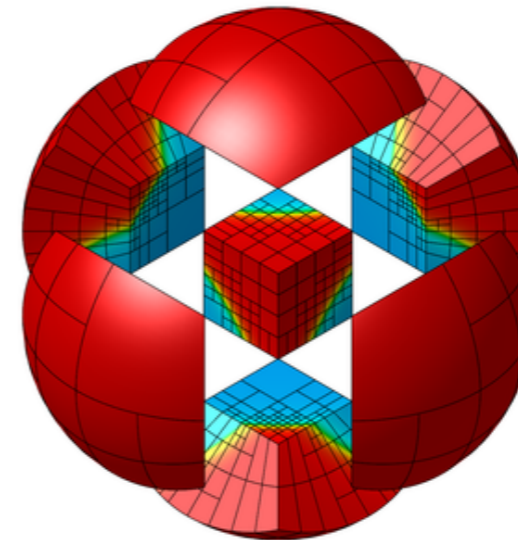
- Further  $\mathcal{O}(10\times)$  speed-up on GPUs + Partial assembly + Linearization in comparison to GPUs + Partial assembly.
- $\mathcal{O}(300\times)$  speed-up in comparison to CPUs + Partial assembly.

	Time to solution (sec)			
	$p = 1$	$p = 2$	$p = 3$	$p = 4$
CPU <sup>PA</sup>	18.0	41.0	128.5	298.0
GPU <sup>PA</sup>	0.4	0.9	3.9	8.5
GPU <sup>PA+Linearized</sup>	0.05	0.16	0.32	0.8
	Speedup (GPU <sup>PA+Linearized</sup> vs GPU <sup>PA</sup> )			
	<b>8×</b>	<b>5×</b>	<b>10×</b>	<b>10×</b>
	Speedup (GPU <sup>PA+Linearized</sup> vs CPU <sup>PA</sup> )			
	<b>360×</b>	<b>256×</b>	<b>400×</b>	<b>372×</b>

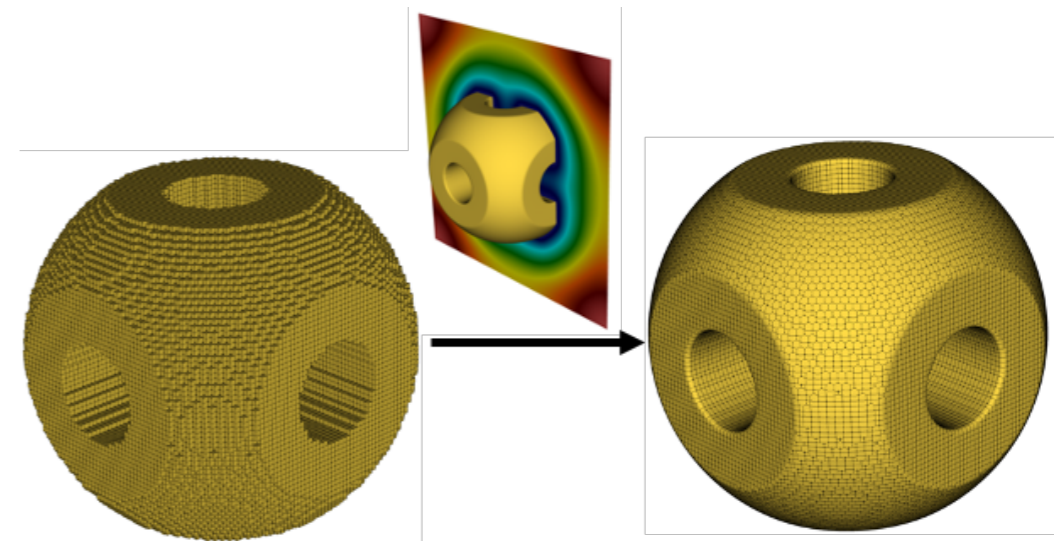


# Summary & Future Work

- High-order mesh optimization using TMOP.
- $\mathcal{O}(30x)$  speed up for mesh optimization using partial-assembly.
- Another  $\mathcal{O}(10x)$  gain from metric linearization.
- Functionality based on open-source high-order FEM library, MFEM.
  - Learn more about it at the virtual MFEM Community workshop in October:  
[www.mfem.org/workshop](http://www.mfem.org/workshop).
- Partial assembly and matrix-free action for other TMOP-based functionalities in future.



[mfem.org](http://mfem.org)



*Barrera et al. High-Order Mesh Morphing for Boundary and Interface Fitting to Implicit Geometries. Computer Aided Design (2023).*



# CASC

Center for Applied  
Scientific Computing



#### **Disclaimer**

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.