Large-scale Finite Element Applications on High-Order Meshes

ELEMENT workshop, EPCC/virtual
October, 2020

Tzanio Kolev
Lawrence Livermore National Lab

LLNL-PRES-815918
This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

High-order finite elements are a good foundation for next-generation scalable multi-physics simulations

- **Large-scale parallel multi-physics simulations**
  - radiation diffusion
  - electromagnetic diffusion
  - compressible hydrodynamics
- **Finite elements naturally connect different physics**

$$H(grad) \xrightarrow{\nabla} H(curl) \xrightarrow{\nabla\times} H(div) \xrightarrow{\nabla\cdot} L_2$$

$H(grad)$	$H(curl)$	$H(div)$	L_2
"nodes"	"edges"	"faces"	"elems"
High-order kinematics	High-order MHD	High-order rad. diffusion	High-order thermodynamics

- **High-order finite elements on high-order meshes**
  - increased accuracy for smooth problems
  - sub-element modeling for problems with shocks
  - HPC utilization, FLOPs/bytes increase with the order
- **Need new (interesting!) R&D for full benefits**
  - meshing, discretizations, solvers, AMR, UQ, visualization, …

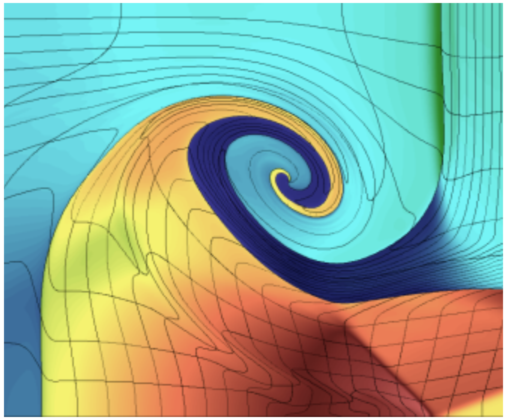

8th order Lagrangian simulation of shock triple-point interaction

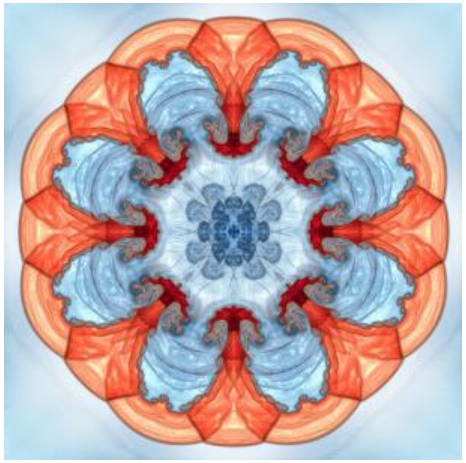

Inertial Confinement Fusion

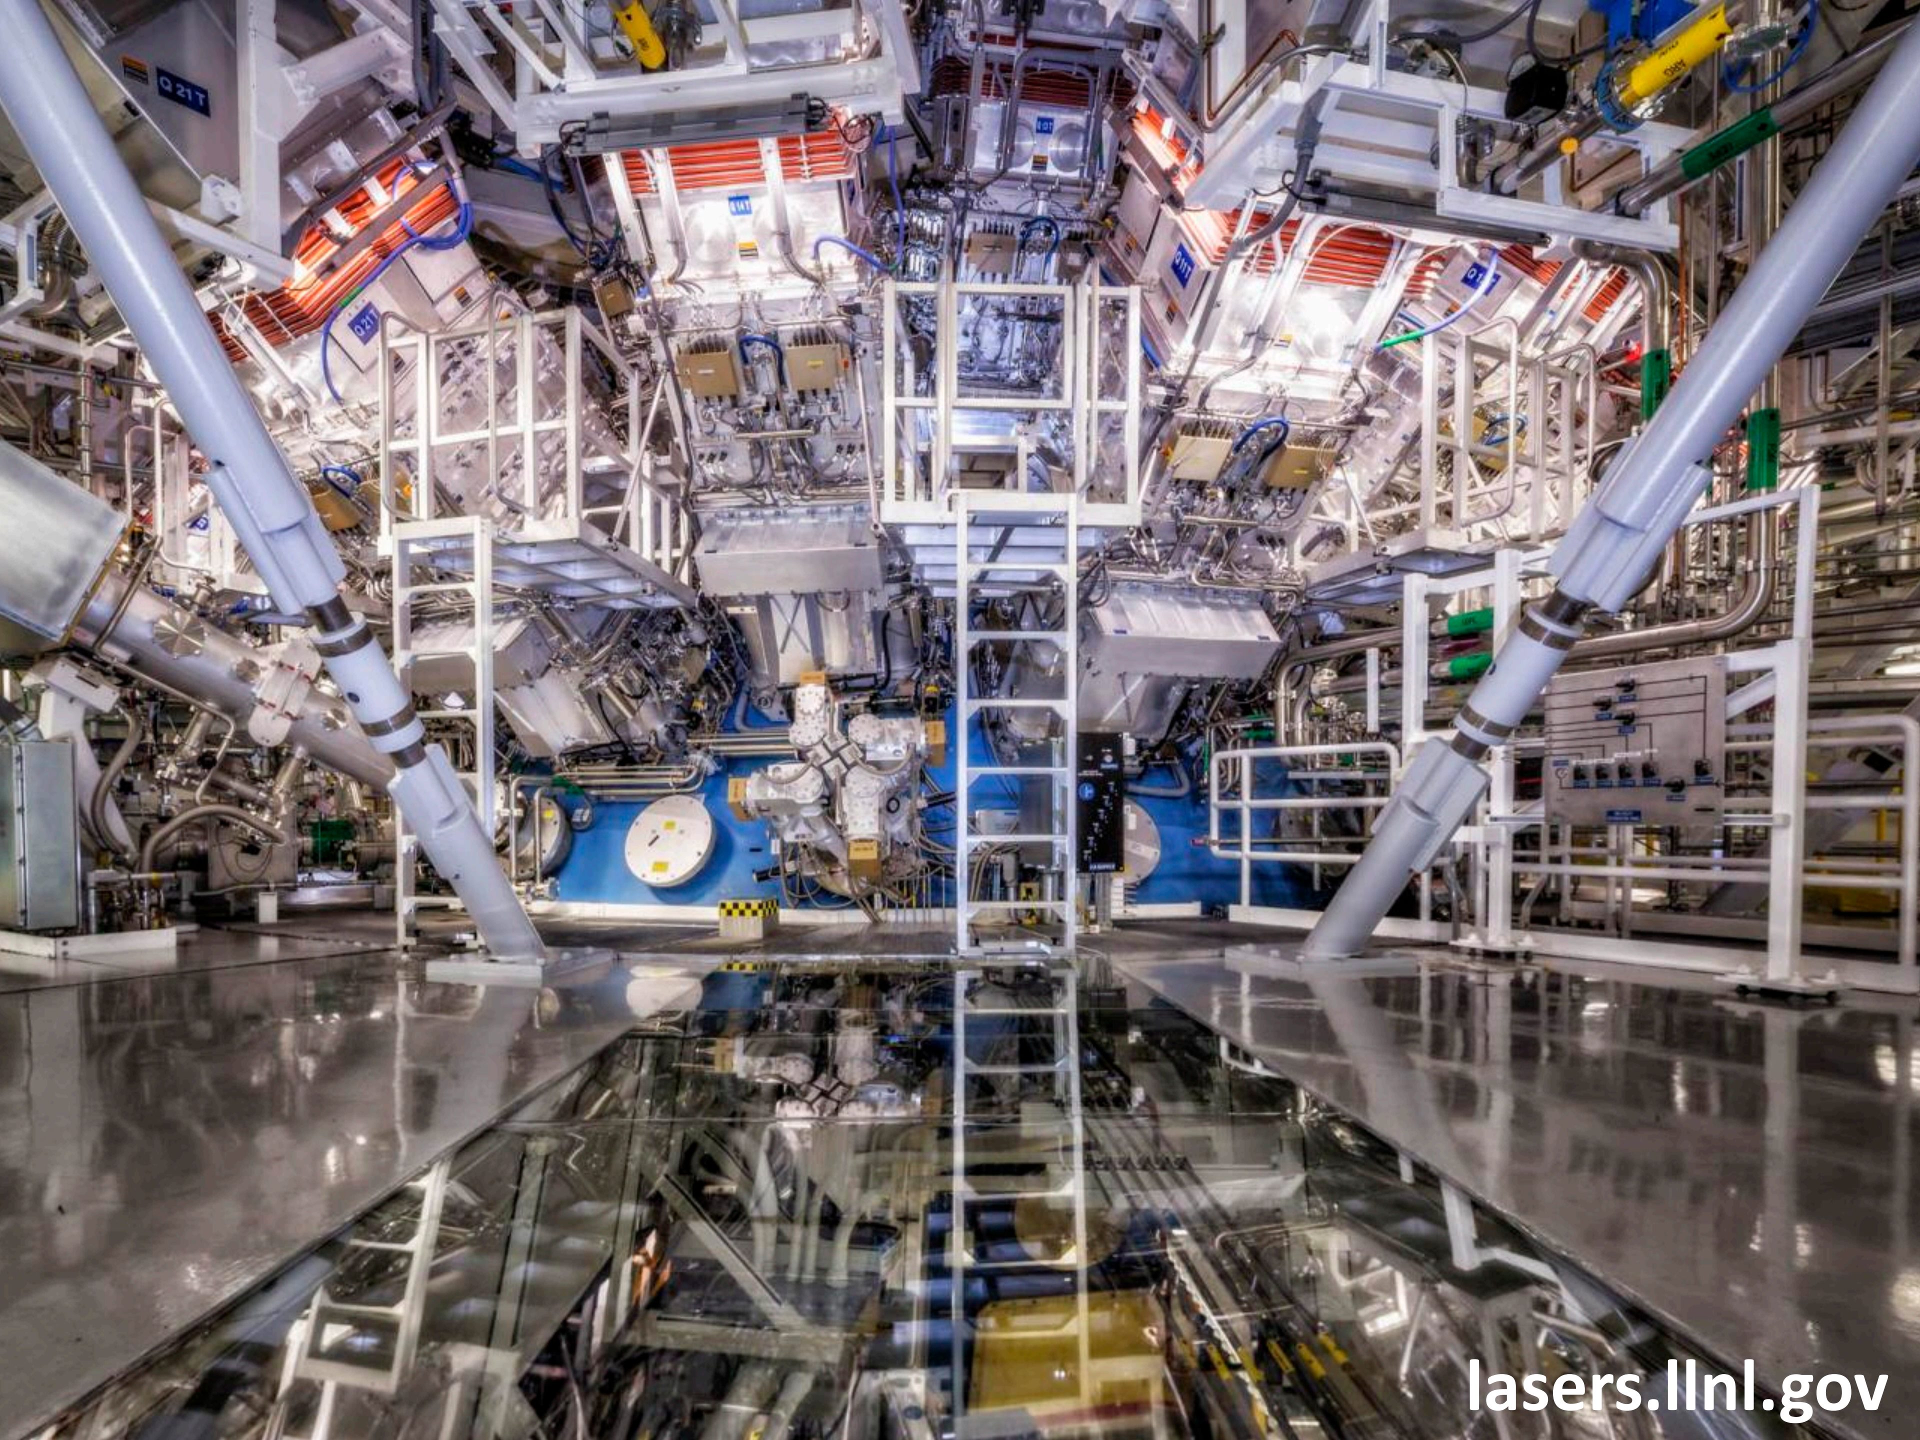
lasers.llnl.gov

We model shock hydrodynamics with high-order FEM in both Lagrange and Remap ALE phases

Lagrange phase
Physical time evolution
Based on physical motion

Remap phase
Pseudo-time evolution
Based on mesh motion

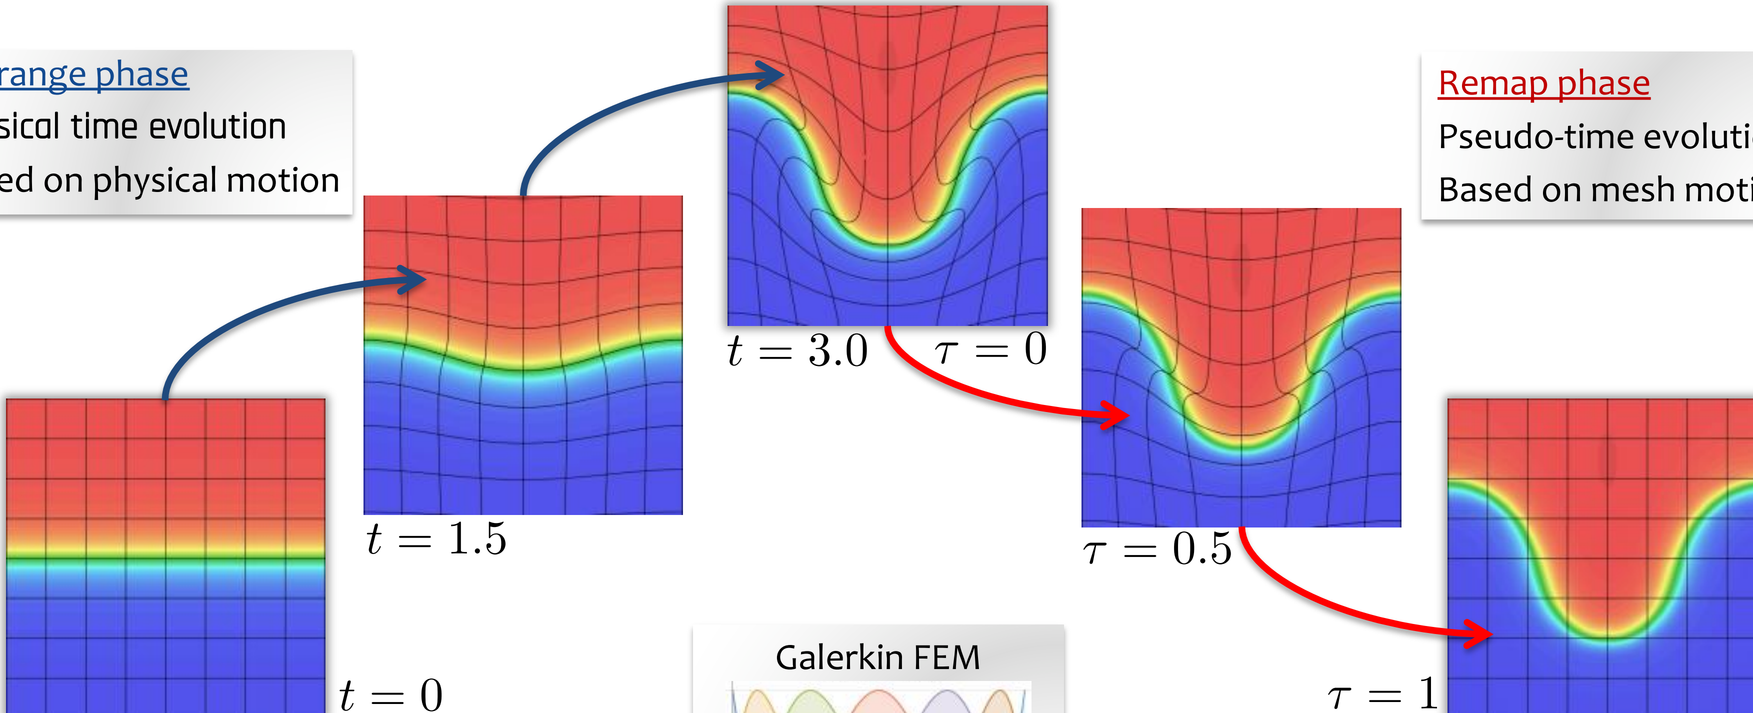

Galerkin FEM

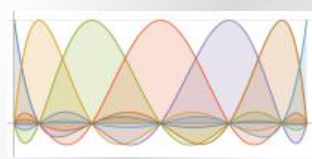

Gauss-Lobatto basis

DG

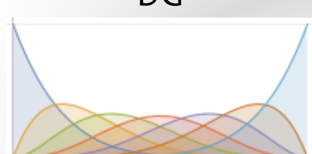

Bernstein basis

Lagrangian phase ($\vec{c} = \vec{0}$)

Momentum Conservation:	$\rho \frac{\mathrm{d}\vec{v}}{\mathrm{d}t} = \nabla \cdot \sigma$
Mass Conservation:	$\frac{\mathrm{d}\rho}{\mathrm{d}t} = -\rho \nabla \cdot \vec{v}$
Energy Conservation:	$\rho \frac{\mathrm{d}e}{\mathrm{d}t} = \sigma : \nabla \vec{v}$
Equation of Motion:	$\frac{\mathrm{d}\vec{x}}{\mathrm{d}t} = \vec{v}$

Advection phase ($\vec{c} = -\vec{v}_m$)

Momentum Conservation:	$\frac{\mathrm{d}(\rho\vec{v})}{\mathrm{d}\tau} = \vec{v}_m \cdot \nabla(\rho\vec{v})$
Mass Conservation:	$\frac{\mathrm{d}\rho}{\mathrm{d}\tau} = \vec{v}_m \cdot \nabla\rho$
Energy Conservation:	$\frac{\mathrm{d}(\rho e)}{\mathrm{d}\tau} = \vec{v}_m \cdot \nabla(\rho e)$
Mesh velocity:	$\vec{v}_m = \frac{\mathrm{d}\vec{x}}{\mathrm{d}\tau}$

* ***"High-order multi-material ALE hydrodynamics"***, SISC 2018

High-order finite elements on high-order meshes
lead to more robust and reliable Lagrangian simulations

	High-order	Low-order SGH
Robustness	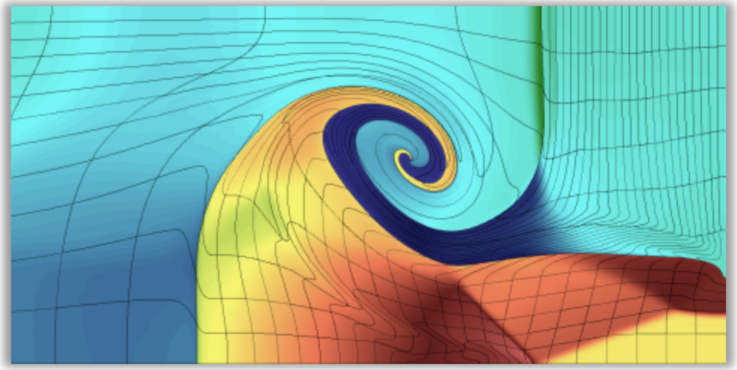	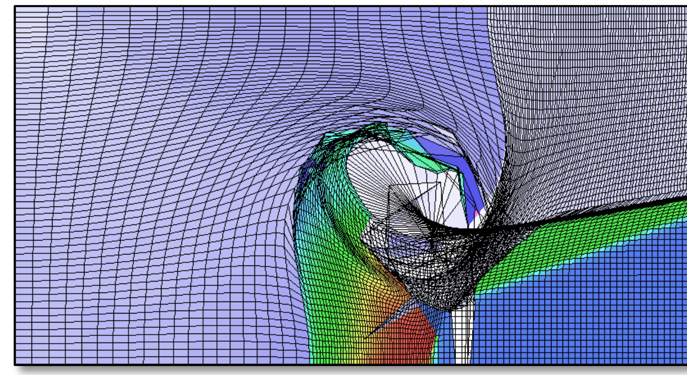
Symmetry preservation	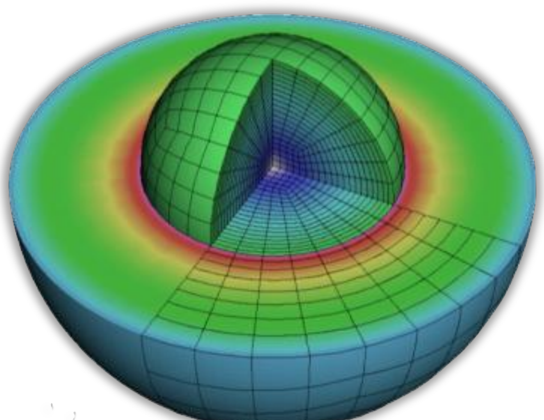	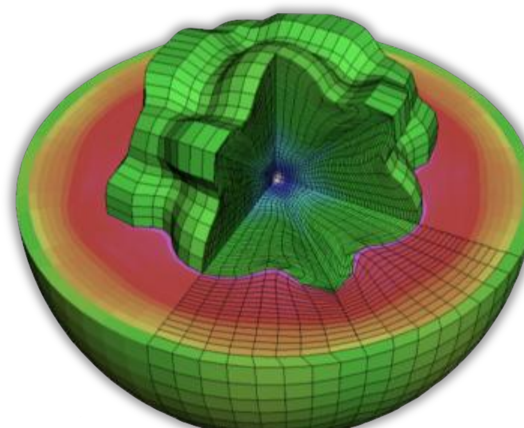

* BLAST project, ***www.llnl.gov/casc/blast***

Moving meshes: high-order mesh optimization and interpolation between meshes

We target *high-order curved elements* + *unstructured meshes* + *moving meshes*

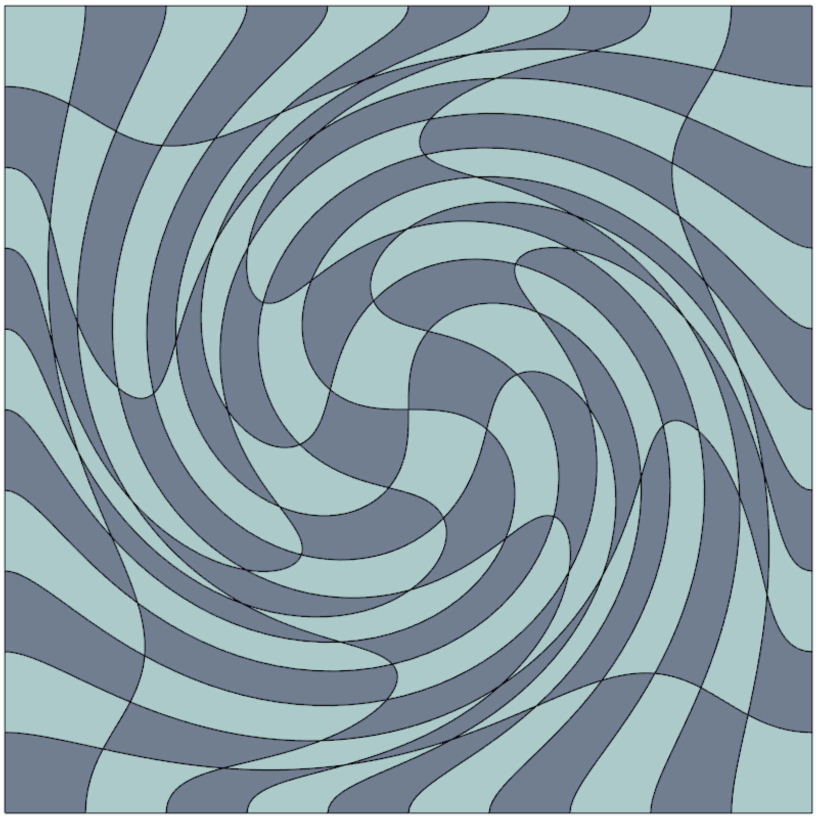

High-order mesh relaxation in MFEM (neo-Hookean evolution)

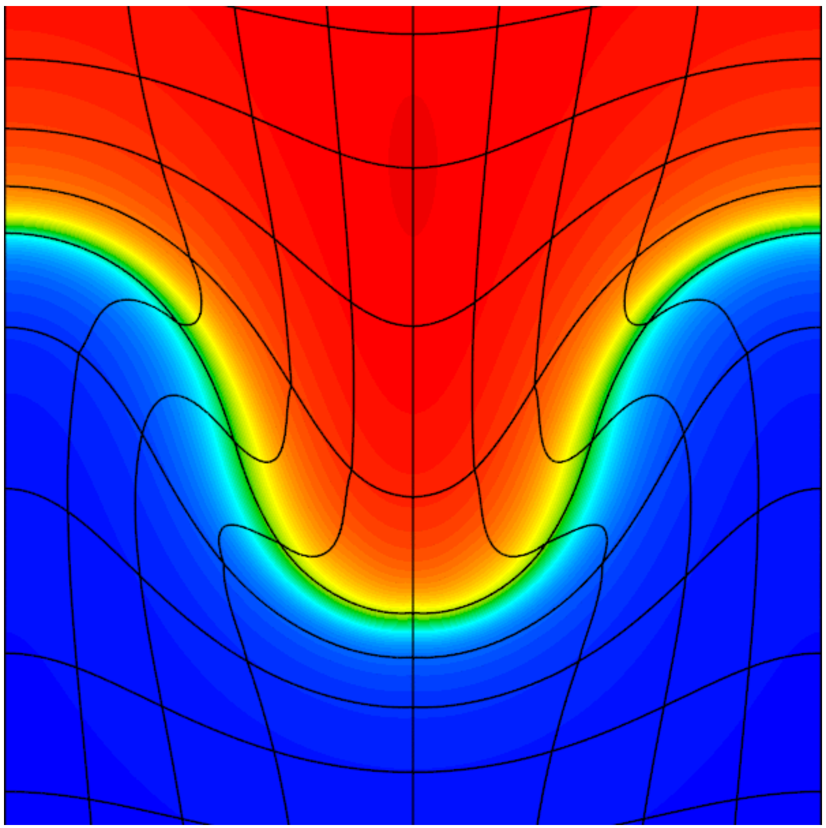

Advection-based interpolation (DG pseudo-time remap in BLAST)

* ***"Monotonicity in high-order curvilinear finite element arbitrary Lagrangian–Eulerian remap"***, IJNMF 2015

Algorithms for high-order mesh optimization

- High-order mesh positions are discretized via position vector and a FE basis:

$$\boldsymbol{x} = (\boldsymbol{x}_1 \ldots \boldsymbol{x}_N)^T, \quad x_q(\bar{x}_q) = \sum_{i=1}^{N} \boldsymbol{x}_i \bar{w}_i(\bar{x}_q)$$

- $\{\overline{w}_i\}_1^{N_E}$ spans Q_k for quadrilateral / hexahedral elements.
- $\{\overline{w}_i\}_1^{N_E}$ spans P_k for triangular / tetrahedral elements.
- Reference -> physical Jacobian is given by the basis functions' gradients:

$$A_q(x) = \frac{\partial x_q}{\partial \bar{x}_q} = \sum_{i=1}^{N} \boldsymbol{x}_i [\nabla \bar{w}_i(\bar{x}_q)]^T$$

- To optimize the curved mesh, we move its nodes by changing $\boldsymbol{x}$.
- Topology is preserved.

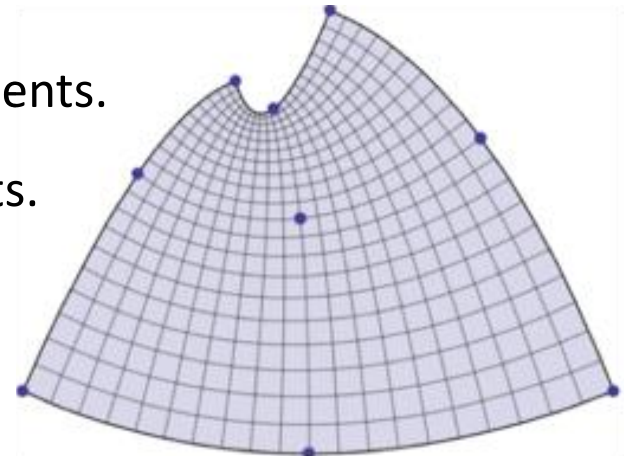

Example of a single Q_2 element

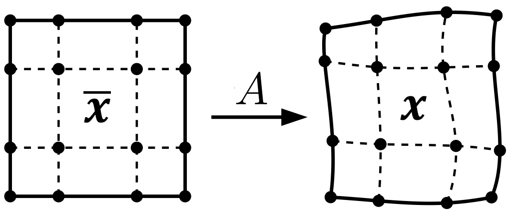

3rd order transformation

* Many papers by: ***Shephard***, ***Shontz***, ***Roca***, ***Geuzaine***, ***Johnen***, ***Persson***, ***Panozzo***

Target-Matrix Optimization Paradigm (TMOP)

- *Target-Matrix Optimization Paradigm (TMOP)*
  - Extended P. Knupp's theory to high-order meshes.
- *Application-specific target elements, W*
  - Allow tailoring to different apps. Examples: ideal, ideal + specified size.
- *Point-based mesh quality metric $\mu(T)$*
  - Can measure shape, size and alignment independently. computed on quadrature point level. Examples:

$$\mu_2^{shape} = \frac{|T|^2}{2\det(T)} - 1 \quad \mu_{55}^{size} = 0.5\left(\det(T) - 1\right)^2$$

- *Global quality functional and minimization*
  - Hessian-based methods need $\partial^2\mu/\partial T^2$.

$$\frac{\partial F(\boldsymbol{x})}{\partial \boldsymbol{x}} = 0\,, \quad \text{where} \quad F(x) = \sum_K \int_{K_t} \mu(T(x))$$

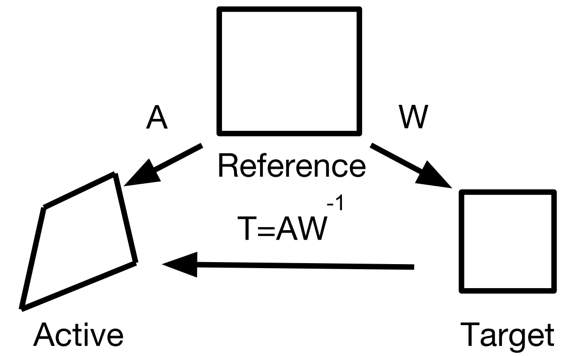

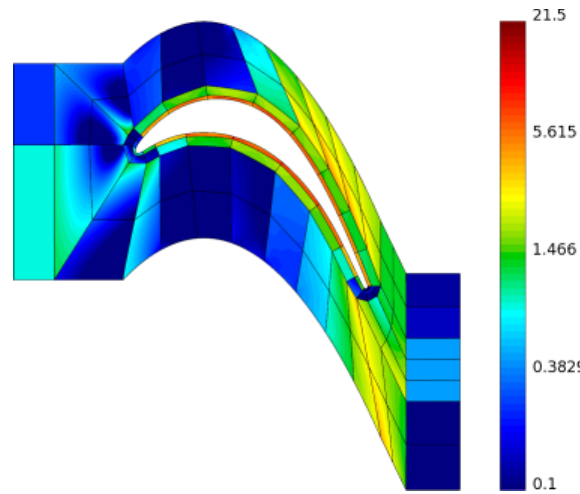

Target construction for solution-based adaptivity

- Jacobian decomposition: $W =$ [volume] [skew] [orientation] [aspect ratio].
- Size/volume and aspect-ratio adaptivity to improve mesh quality around high-gradient regions of the solution (η)

$$W = \underbrace{\begin{bmatrix} \sqrt{s} & 0 \\ 0 & \sqrt{s} \end{bmatrix}}_{\text{volume}} \underbrace{\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}}_{\text{skew}} \underbrace{\begin{bmatrix} 1 & \cos\phi \\ 0 & \sin\phi \end{bmatrix}}_{\text{orientation}} \underbrace{\begin{bmatrix} \frac{1}{\sqrt{r}} & 0 \\ 0 & \sqrt{r} \end{bmatrix}}_{\text{aspect ratio}}$$

$$s \propto ||\nabla \eta||$$
$$r \propto \nabla_x \eta / \nabla_y \eta$$

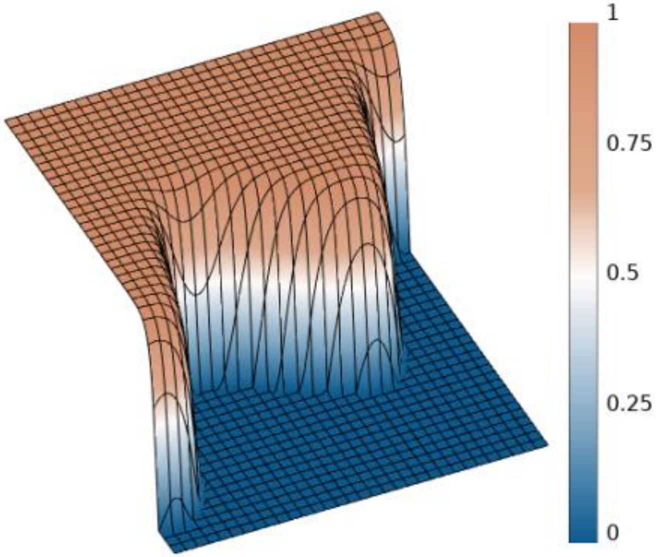

Material indicator on the original mesh

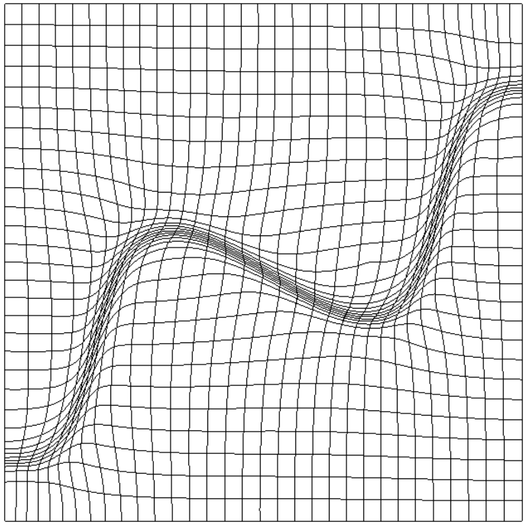

$$\mu_7(T) = |T - T^{-t}|^2$$

Mesh optimized for shape + size

* ***"Simulation-Driven Optimization of High-Order Meshes in ALE Hydrodynamics"***, 2020

Application to solution adaptation in multi-material ALE

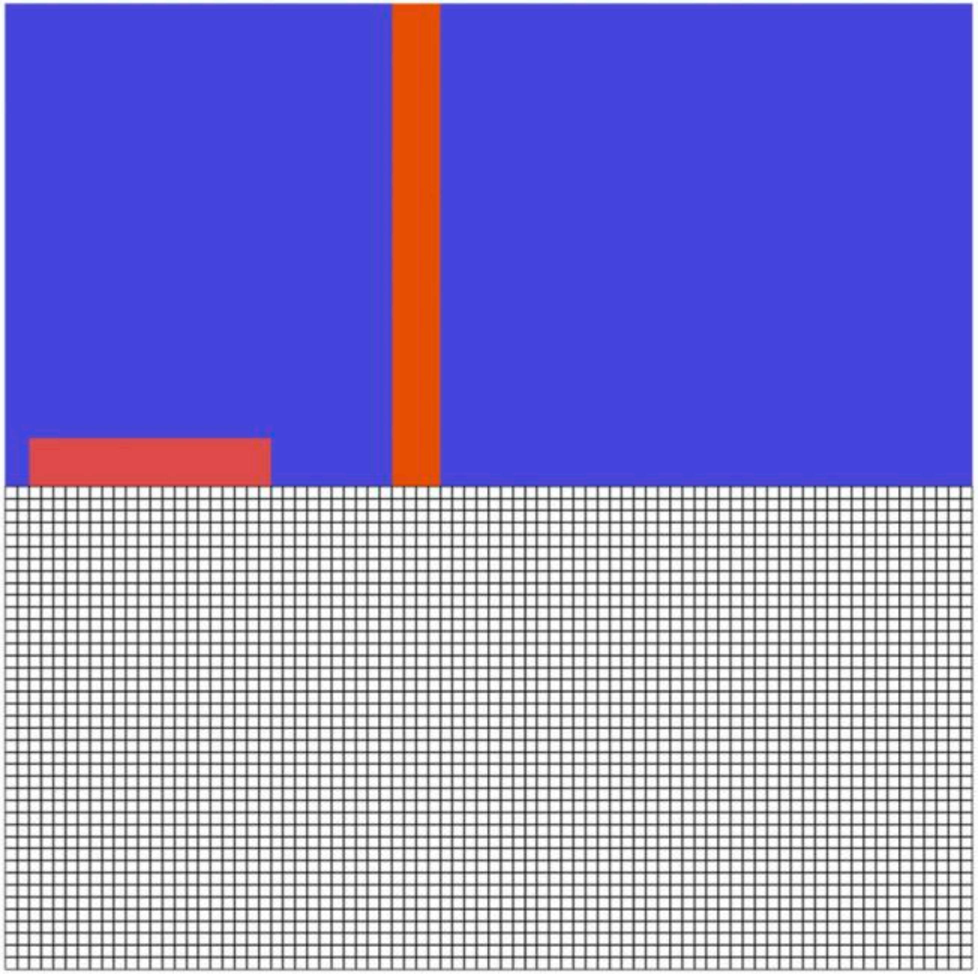

High-velocity gas impact. Mesh adapted to (high-order) material indicators.

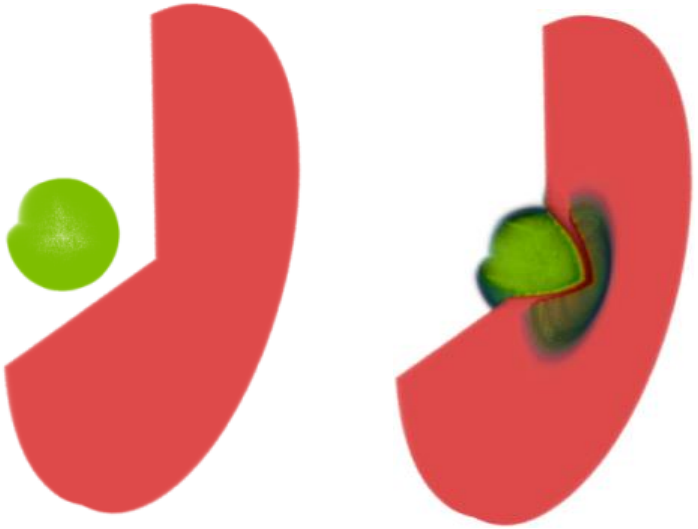

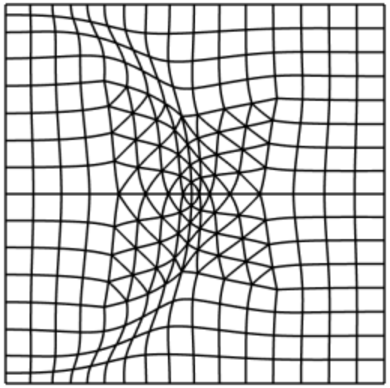

TMOP extends to higher space dimensions, different element types

Algorithms for unstructured non-conforming AMR

Adaptive mesh refinement on library level

- Conforming local refinement on simplex meshes
- *Non-conforming refinement for quad/hex meshes*
- h-refinement with fixed p (p-refinement in a branch)

General approach

- any high-order finite element space, H1, H(curl), H(div), … on any high-order curved mesh
- 2D and 3D
- arbitrary order hanging nodes
- anisotropic refinement
- derifenement
- serial and parallel, including parallel load balancing
- independent of the physics (easy to incorporate in applications)

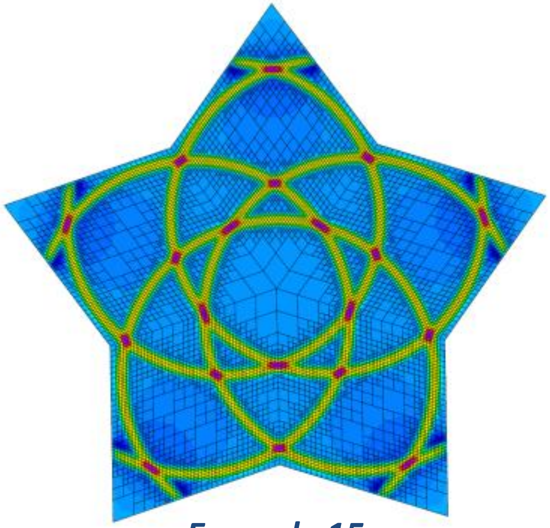

Example 15

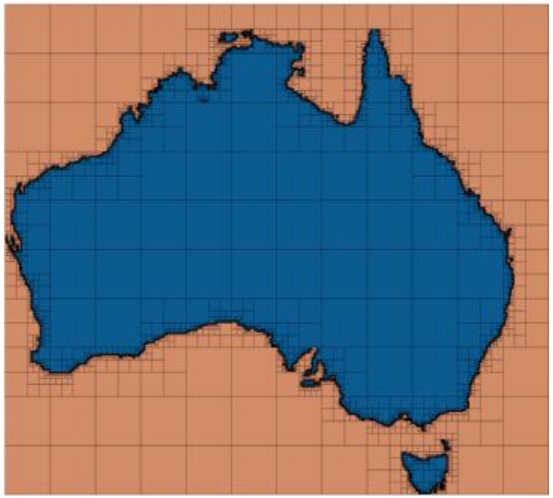

Shaper miniapp

Nonconforming variational restriction

High-order elements

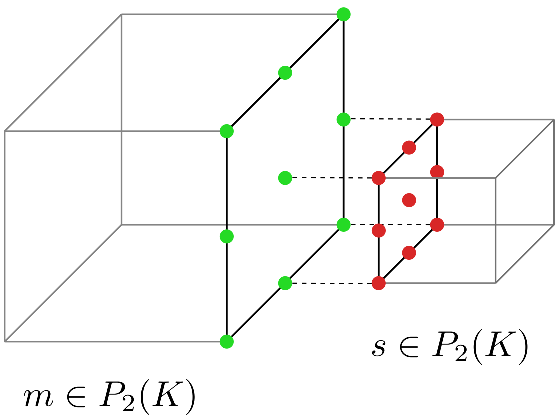

Indirect constraints

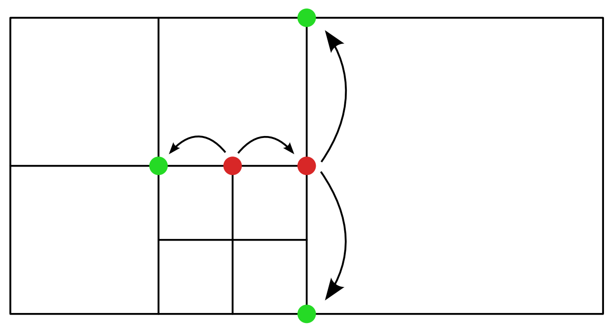

more complicated in 3D…

H(curl) elements

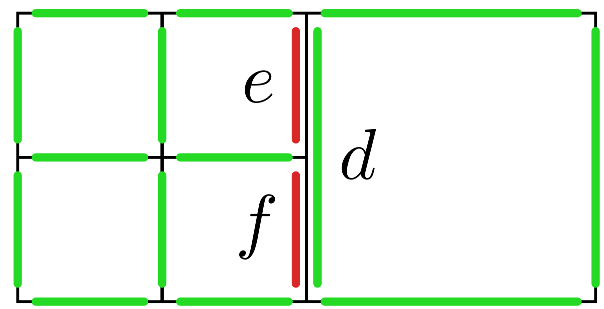

constraint: e = f = d/2

Global interpolation

$$y = Px, \quad P = \begin{bmatrix} I \\ W \end{bmatrix}$$

x *– conforming dofs*

y *– nonconforming dofs*
(unconstrained + constrained)

W *– constrained dofs interpolation*

$$dim(x) \leq dim(y)$$

Variational restriction

$$P^T A P x = P^T b \qquad y = Px$$

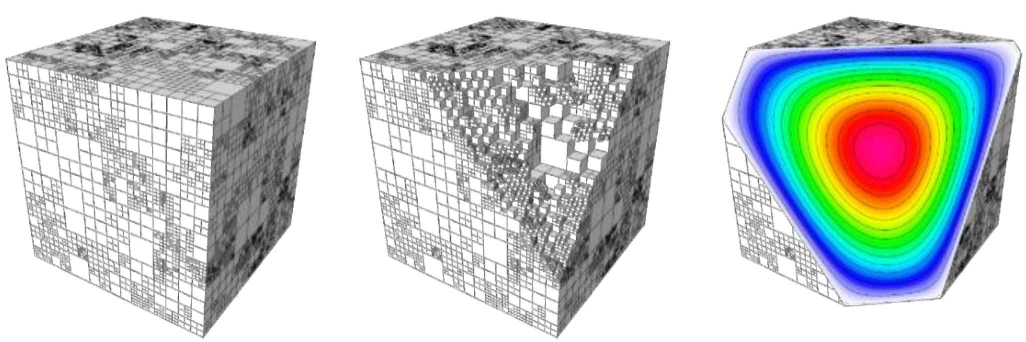

* ***"Non-Conforming Mesh Refinement for High-Order Finite Elements"***, SISC 2019

Lagrangian dynamic AMR on Sedov blast

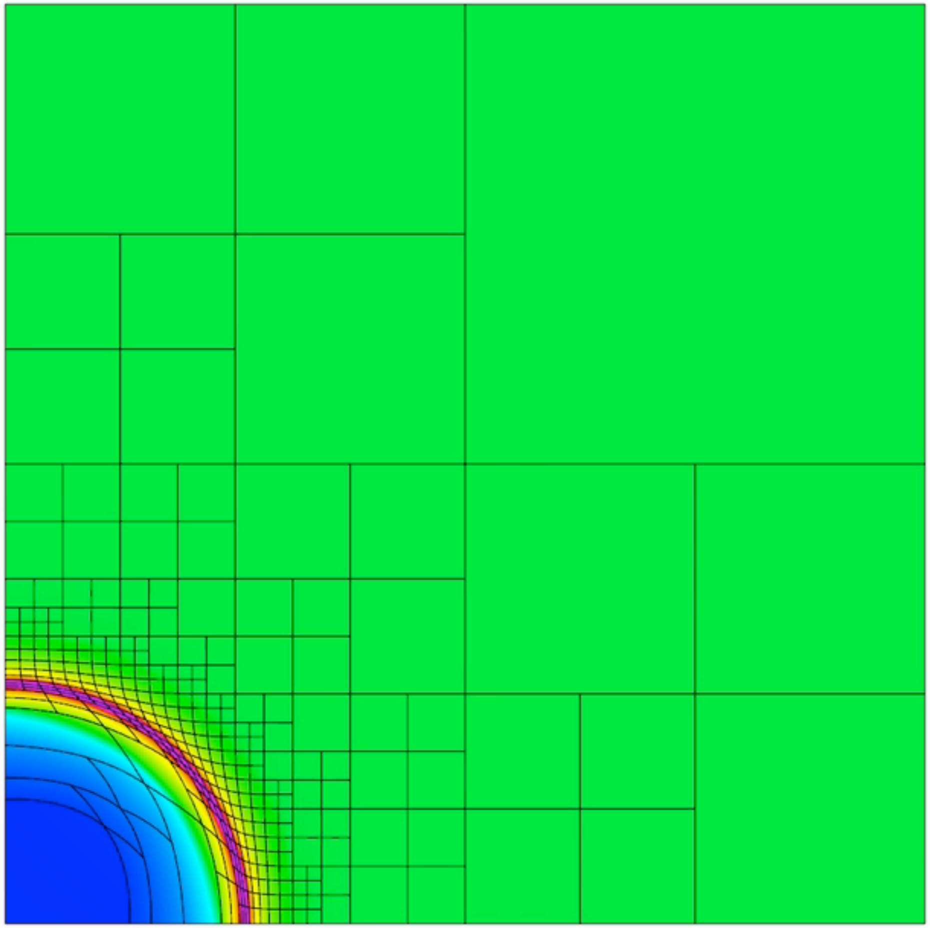

Adaptive, viscosity-based refinement and derefinement, 2nd order Lagrangian Sedov

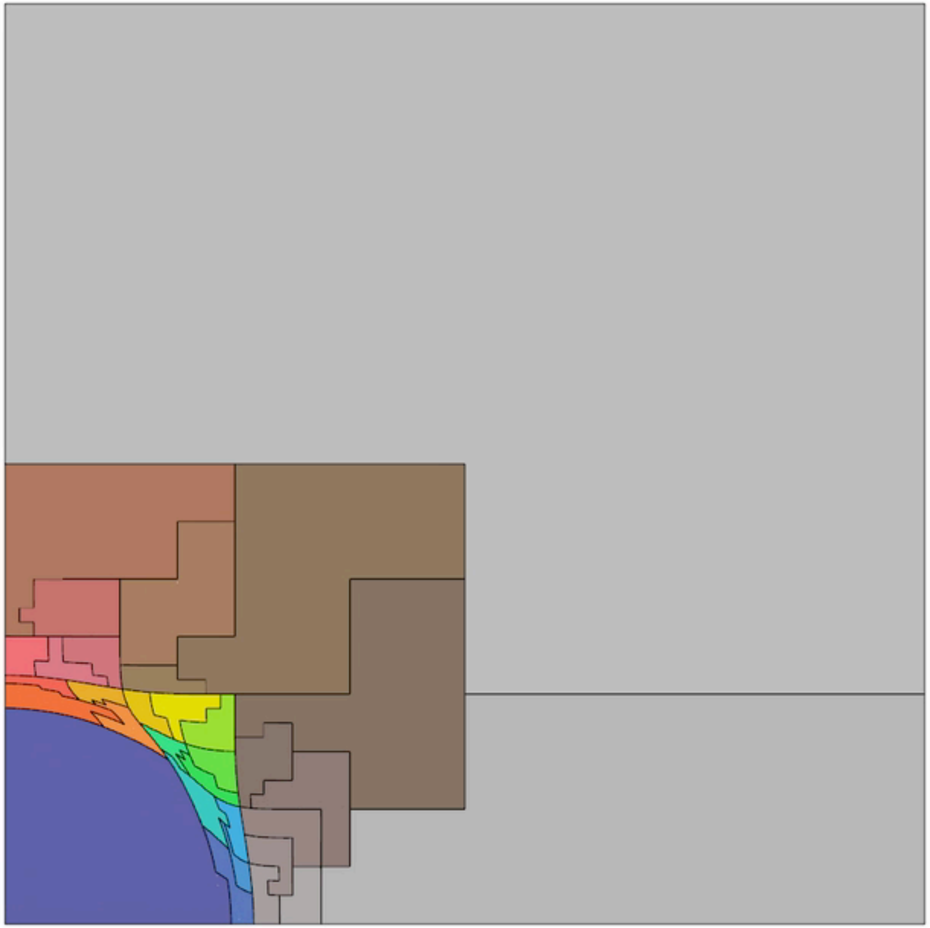

Parallel load balancing based on space-filling (Hilbert) curve partitioning, 16 cores

Lawrence Livermore National Laboratory

BLAST

hr-adaptivity for nonconforming high-order meshes

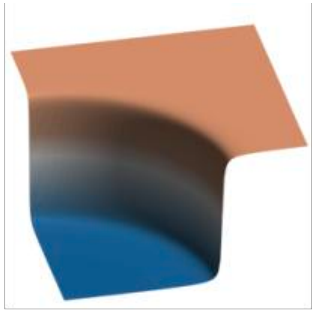

solution

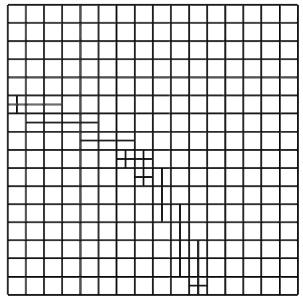

h-adaptivity

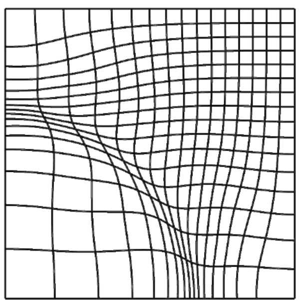

r-adaptivity

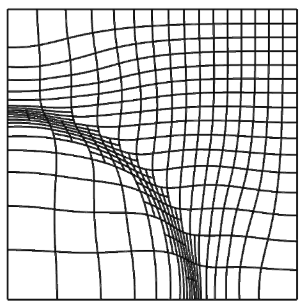

hr-adaptivity

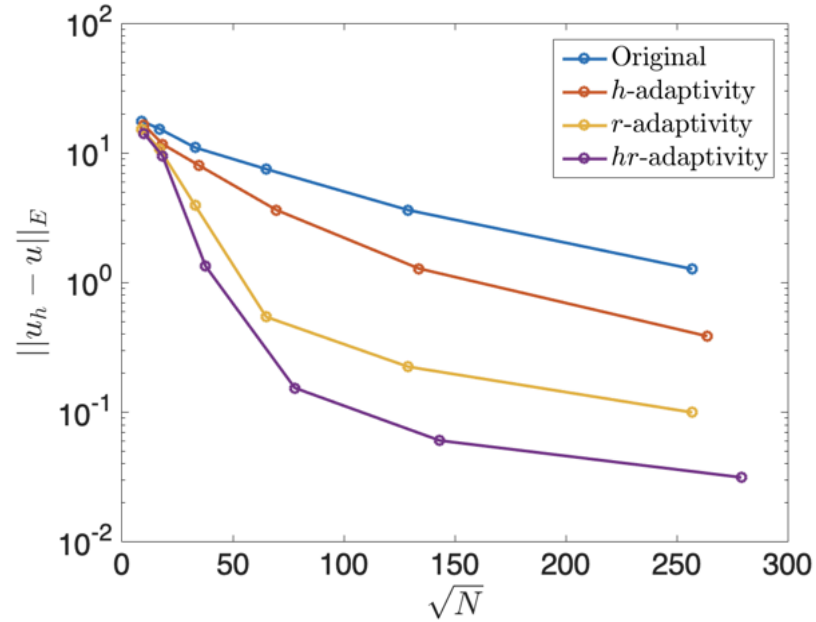

hr-adaptivity requires many fewer degrees than r- or h-adaptivity for the same error

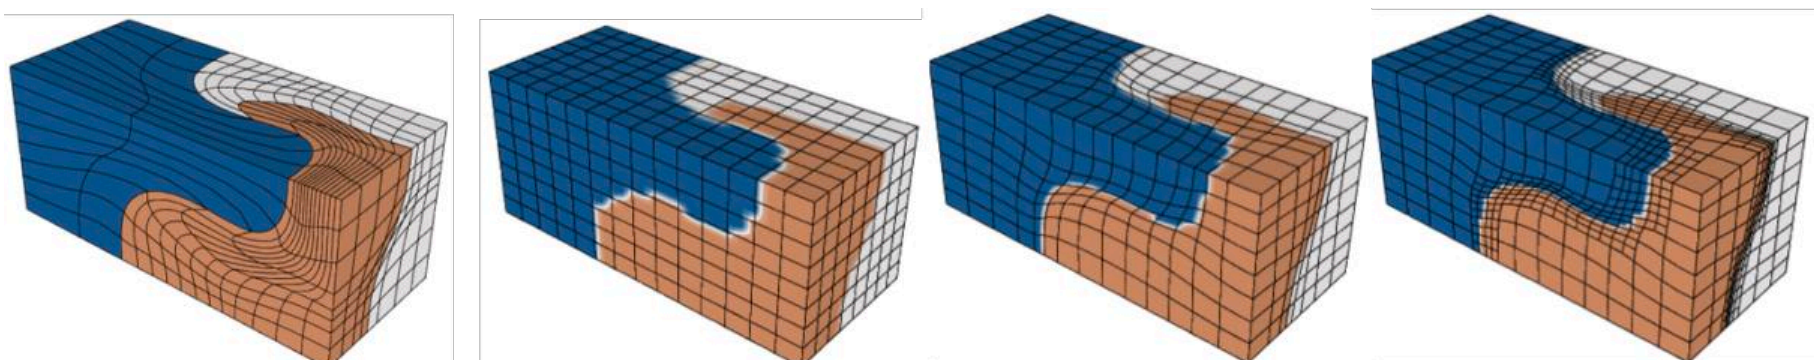

hr-adaptivity leads to 20x reduction in error relative to a Cartesian mesh and 8x relative to r-adaptivity mesh

Lawrence Livermore National Laboratory * ***"hr-adaptivity for nonconforming high-order meshes with the target matrix optimization paradigm"***, 2020 ethos

Robust + efficient algorithms for high-order applications

CEED: *High-performance algorithms on advanced architectures*

EXASCALE COMPUTING PROJECT

CEED EXASCALE DISCRETIZATIONS

ceed.exascaleproject.org

ETHOS: *High-order mesh optimization*

ethos

www.llnl.gov/casc/ethos

MFEM: *Finite element methods research + software*

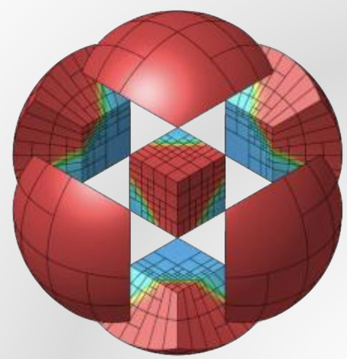

mfem.org

BLAST: *High-order ALE shock multi-physics code*

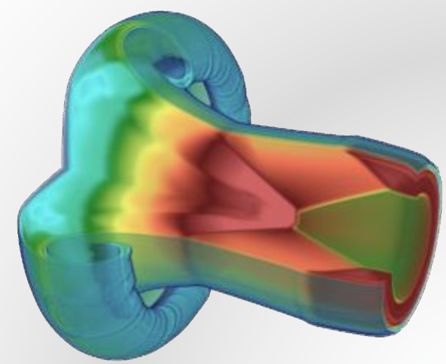

www.llnl.gov/casc/blast

GPU port of high-order mesh optimization

- The variational formulation allows to use *partial assembly* extensively.

$$F(x) := \sum_{E \in \mathcal{M}} \int_{E_t} \mu(T(x_t))dx_t$$

$$A = \frac{\partial^2 F}{\partial x^2}$$

$$\boxed{A = P^T G^T B^T D B G P}$$

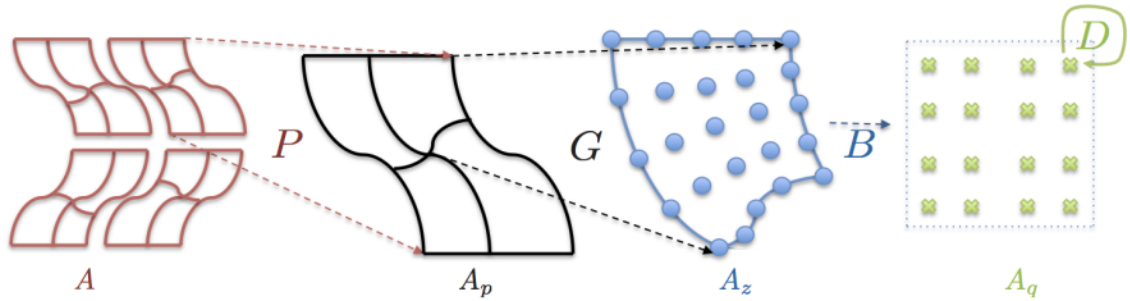

- Naturally separates FE-based (P, G, B) and TMOP-based operations (D).
- TMOP-based (D) PA / GPU kernels: $T(x),\ \mu(T),\ \frac{\partial \mu(T)}{\partial T},\ \frac{\partial^2 \mu(T)}{\partial T^2}$
- FE-based (P, G, B) PA / GPU kernels:
  - Integral of F, and the nonlinear form $\partial F / \partial x$.
  - Local action of the Hessian $\partial^2 F / \partial x^2$, and its diagonal.
- Initial speedups relative to full assembly:
  - 3D, 4th order mesh, 448 elements: 6.5x (1 CPU PA), 102x (1 GPU PA)
  - full-scale runs on Sierra (16K V100 GPUs)

Parallel AMR scaling to ~400K MPI tasks

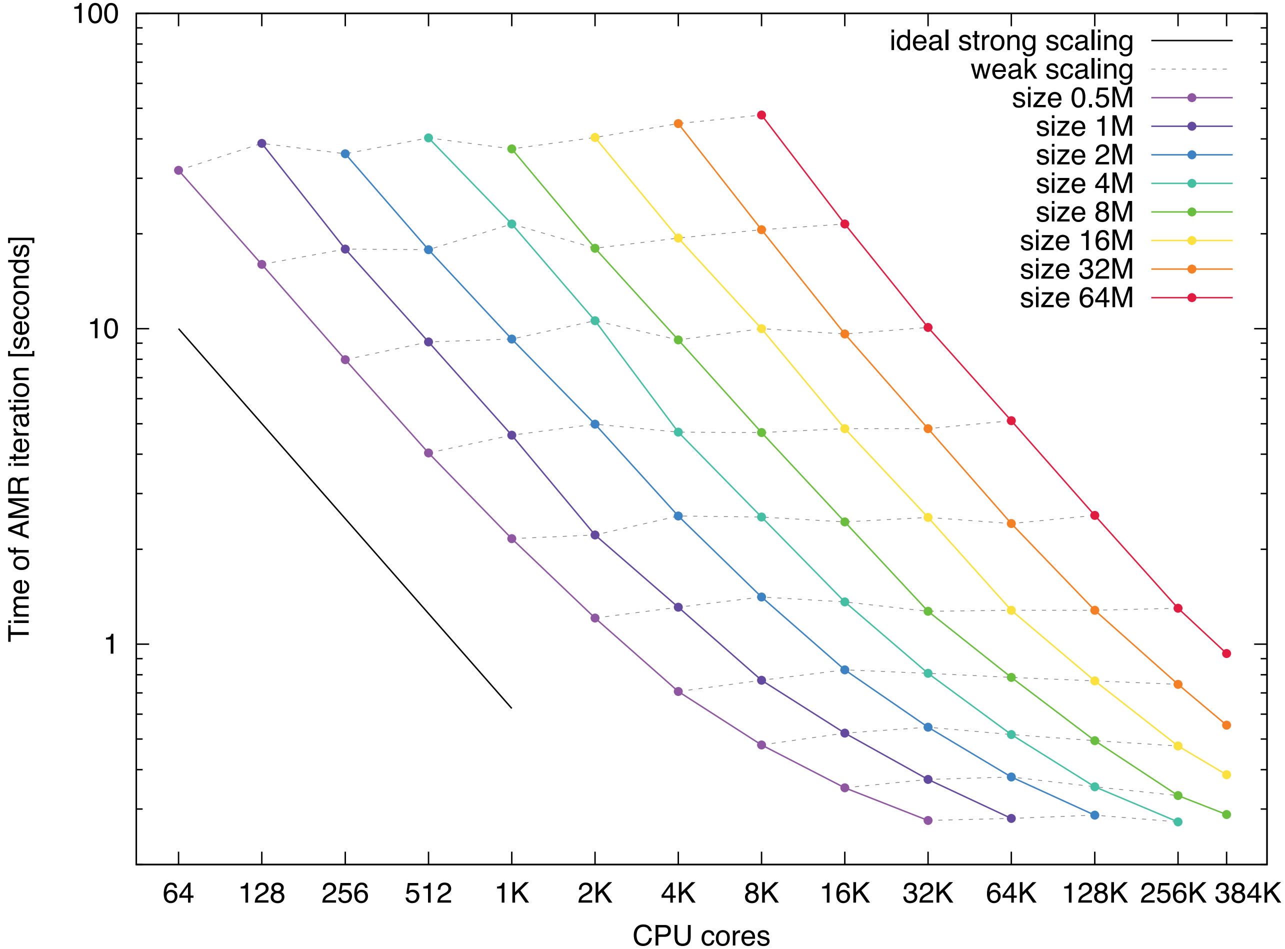

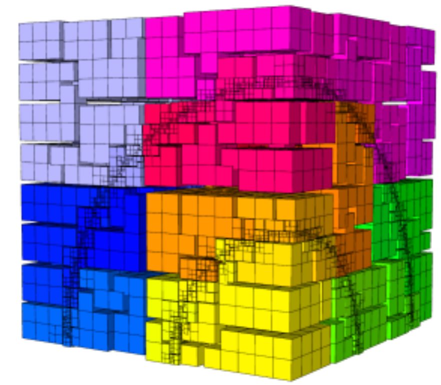

Parallel decomposition (2048 domains shown)

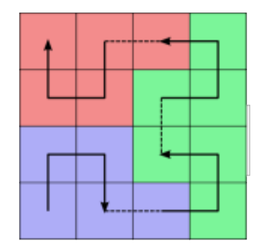

Parallel partitioning via Hilbert curve

- weak+strong scaling up to ~400K MPI tasks on BG/Q
- **measure AMR only components**: interpolation matrix, assembly, marking, refinement & rebalancing (no linear solves, no "physics")

Field & Mesh Specification

ceed.exascaleproject.org/fms

- ***FMS is a new lightweight API/specification describing:***
  1. mesh topology
  2. finite element fields defined on the mesh

- Both mesh and fields can be general, high-order.
- Mesh nodes are just one of the fields.

- All mesh entities of all dimensions are represented: *vertices, edges, faces, elements:*
  - mesh entities described by downward adjacencies.
  - fields described by dofs associated with interior of entities

- Visualization: ***next version of VisIt will support FMS!***
  - Native support in MFEM, CEED partners
  - Binary and ASCII I/O formats via Conduit

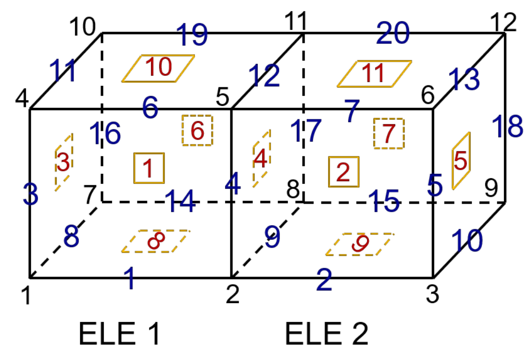

Element → Face → Edge → Vertex

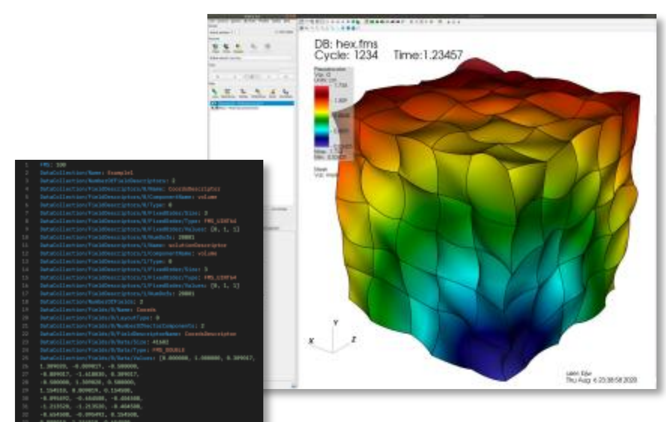

FMS version 0.2 to be released soon – let us know what you think!

Current and future work

- *High-order finite elements on high-order meshes show promise for HPC multi-physics simulations*
- Some ongoing research:
  - GPU-oriented algorithms and performance optimization on modern architectures
  - Matrix-free scalable preconditioners
  - Combination of TMOP and AMR (*hr*-adaptivity)
  - Approximate preservation of discrete surfaces
  - Improved nonlinear solvers
- Papers and additional details:

  `people.llnl.gov/kolev1`
- Open-source finite element software:

mfem.org

CEED

EXASCALE DISCRETIZATIONS

Q4 Rayleigh-Taylor single-material ALE on 256 processors

CASC

Lawrence Livermore National Laboratory

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.
LLNL-PRES-815918