Aspect-Ratio Preserving Patch-Based TMOP with Linesweep Smoothing for ALE and ALE-AMR Meshes

A. Colaïtis[1,2], S. Guisset[3], J. Breil[3]

Abstract

We present a rezoning method for three-dimensional Arbitrary Lagrangian–Eulerian (ALE) and ALE with Adaptive Mesh Refinement (ALE-AMR) simulations. The method combines linesweep smoothing with an aspect-ratio-preserving Target-Matrix Optimization Paradigm (TMOP) formulation. It is designed to preserve anisotropic mesh structures produced by the Lagrangian phase while maintaining mesh validity and reducing shear, sliver formation, and cell bunching in block-structured, optionally nonconformal meshes. The approach combines a patch-local TMOP objective and a convexity-based disentanglement metric within an existing linesweep framework. The TMOP target matrix is built from local edge geometry and includes a prescribed aspect-ratio cap, which provides an explicit control parameter for mesh relaxation. The metric is evaluated on an extended nodal patch, so that the objective couples the motion of a node to the quality of neighboring elements affected by that motion. Combined with sequential updates, this ensures that accepted steps are tested against the updated quality of the local coupled patch, rather than only against the elements directly incident on the node. When invalid configurations are detected, the TMOP objective is replaced by a convexity objective that drives signed volumes back to positive values. Parallelism is recovered through a distance-1 coloring of the nodal graph, allowing independent node updates to be performed concurrently. Numerical experiments in two and three dimensions illustrate the roles of the linesweep, TMOP, and disentanglement components. The results show that the combined method reduces shear and controls the high-aspect-ratio tail through the prescribed cap, while retaining anisotropy inherited from the Lagrangian mesh. In the tested ALE-AMR configurations, including nonconformal AMR cases with large initial perturbations, the patch-based update also avoids the invalid intermediate meshes observed with bulk metric-optimization updates.

1. Introduction

Arbitrary Lagrangian-Eulerian (ALE) methods provide a way to choose the mesh motion independently of the material motion. Depending on the calculation, nodes may follow the material, remain fixed as in an Eulerian description, or move according to an intermediate prescription used to limit distortion, control resolution, or satisfy boundary constraints [1, 2, 3]. In indirect ALE hydrodynamics, this leads to a decomposition of a time step into a Lagrangian phase, a rezoning phase, and a conservative remapping phase [1, 4, 3]. The rezoner therefore plays a role beyond geometric smoothing. It determines how much of the mesh deformation produced during the Lagrangian phase is kept before remapping, and how much is relaxed to obtain a suitable mesh for the next hydrodynamic step.

This balance is especially important in compressible multi-material calculations. The Lagrangian mesh may contain useful information: compression of cells, alignment of mesh lines, or anisotropy of the local cell geometry may follow shocks, contact discontinuities, material interfaces, or flow-aligned structures. At the same time, this same motion can rapidly degrade the mesh in the presence of vorticity, shear, or strong compression. A practical rezoner must therefore reconcile two goals that can be opposed: preserve the

[1]Laboratory for Laser Energetics, 250 East River Rd, Rochester, NY 14623-1212
[2]Corresponding author: acol@lle.rochester.edu
[3]CEA/CESTA, 15 Avenue des Sablières CS 60001, 33116 Le Barp cedex, France

intrinsic Lagrangian information encoded in the evolving mesh, and prevent loss of validity and degradation of mesh quality. Beside mesh-quality optimization, a mesh rezoner can also act as a disentangler when necessary, and control sliver formation and cell bunching, since very small cells may drive the timestep down through CFL constraints [4, 5, 6, 7].

Several families of methods address part of this problem. Classical smoothing and mesh-motion algorithms can improve mesh regularity by moving nodes according to local averages, elliptic or mesh-motion equations, or line-based relaxations [8, 9, 10, 11]. Optimization-based methods instead minimize algebraic quality measures or objectives built from local Jacobian matrices, and can be used either for shape improvement or for untangling [12, 13, 14, 7, 15]. More recently, target-based formulations, and in particular the Target-Matrix Optimization Paradigm (TMOP), make it possible to compare the current element to a prescribed target element rather than to a single isotropic ideal metric [16, 17]. This distinction is useful here because a metric that is purely isotropic (in the mesh quality sense) will improve mesh quality at the detriment of local information contained in the mesh.

In an ALE-AMR setting, the same compromise is more constrained. While AMR is introduced to concentrate resolution near localized structures, coupling AMR with moving meshes requires the rezoning and remapping operations to remain consistent across refinement-level transitions [18, 19, 20, 21]. In block-structured nonconformal meshes, edge- and face-hanging nodes create additional geometric dependencies: moving a corner node may also move reconstructed hanging nodes, and therefore affect cells beyond a simple conformal nodal neighborhood. Moreover, the jump in cell sizes across a 2:1 level transition should not be interpreted by the rezoner as a mesh-quality defect. A robust ALE-AMR rezoner must therefore remain compatible with nonconformal connectivities and preserve physically meaningful anisotropy across level jumps.

1.1. Linesweep rezoning and its adaptation to nonconformal AMR

In previous work, we have presented an ALE-AMR strategy relying on a *weighted linesweep* rezoner [22, 21]. The method was designed to improve mesh quality while preserving a large fraction of the mesh's Lagrangian structure. It proceeds by defining, at each node, three logical mesh-line directions and computing a rezoned point from local one-dimensional line relaxations. In the undamped limit, the algorithm reduces to a transverse averaging that smooths aspect ratios while largely preserving the original spacing information along each direction. As damping increases, it progressively transitions toward an equal-space relaxation along mesh lines [10], which tends to isotropize the mesh. Because this operation can rezone large regions, the linesweep is activated only where a flow-induced line-rotation indicator exceeds a prescribed threshold, so that purely Lagrangian regions well aligned with the flow are minimally perturbed.

Extending this concept to block-structured nonconformal AMR meshes requires a robust definition of logical mesh lines across rotated blocks and junctions, a consistent treatment of hanging nodes, and safeguards for degenerate topological configurations at block corners and triple junctions. In [21], this was achieved by using node-local line directions, a local orientation-matching map across neighbors, and case-dependent fallbacks for locally reduced connectivities, while constraining boundary nodes along symmetry planes and lines. We refer to [21] for the complete formulation and the AMR-specific details.

1.2. Combination with metric optimization and shortcomings of the method

Linesweep rezoning is highly effective at reducing large-scale mesh distortion while preserving Lagrangian spacing, but it has known blind spots with respect to local shear, skewness, and cell bunching. In the previous approach, these blind spots were addressed by complementing the linesweep with a conventional nodal metric optimization stage driven by a cell-shearing metric evaluated on a nodal neighborhood [21]. Although this was conceptually attractive, extensive testing revealed two practical shortcomings that motivate the method proposed in this paper.

The first limitation is related to the way nodal updates are applied. In the prior scheme, the metric is evaluated on the current mesh, a descent direction is computed, and nodes are advanced with a line-search update. When this update is applied in bulk, all nodes move based on the same frozen configuration. As a result, even if each nodal step is acceptable with respect to its current neighbors, the simultaneous

motion of several neighboring nodes can invalidate elements once the full update is applied. Similarly, a local decrease predicted for fixed neighbors does not necessarily imply a decrease of the coupled objective after the geometry has changed. This is most problematic near small or nearly tangled elements, where admissible steps become very small and the algorithm can become both slow and fragile.

The second limitation concerns the optimization metric used previously; a normalized condition-number in which local edge directions are normalized before the standard condition number [12] is evaluated. This metric optimizes orthogonality and some measure of smoothness, while remaining compatible with the AMR framework because it is insensitive to local volume and edge-length ratios. However, this also makes the objective awkward for ALE rezoning. Thin but nearly orthogonal elements may not be strongly penalized, even though they lead to small volumes and restrictive timesteps. Conversely, using a non-normalized metric reintroduces sensitivity to volume differences around the node and tends to drive undesirable isotropization. Finally, in practice, the normalized condition number is not a sufficiently smooth barrier against tangling: near singular or inverted configurations, it does not provide a repulsive barrier away from invalidity, and interacts poorly with bulk updates and with small-step restrictions, which can trap the rezoner near singular configurations.

Taken together, these issues define the tradeoff addressed in this work. We want a metric that is sufficiently insensitive to relative volume differences to avoid erasing Lagrangian and AMR information, but that still penalizes dangerously small or highly distorted elements and provides a usable path away from invalid configurations. Combining several ad hoc metrics with thresholds can partially address this issue, for example by penalizing normalized condition number together with volume variations or by adding a separate global volume objective. However, such approaches quickly become parameter-heavy and difficult to make robust.

1.3. A new strategy: patch-based TMOP with dynamic targets and Gauss-Seidel updates

The approach proposed here modifies both the objective being optimized and the way nodal updates are applied.

We first replace the conventional metric-optimization stage [12, 13] by a TMOP objective [16, 17]. In TMOP, the quality of the current element is measured relative to a prescribed target element. The target represents the local shape that the rezoner is allowed to approach. This differs from standard metric optimization, where the objective is defined directly from the current element geometry, for example through a measure of shear or skewness. With TMOP, the target can include geometric information that should be retained during rezoning.

In the present method, this target is rebuilt from the local Lagrangian mesh geometry at each rezoning step and is chosen to preserve the anisotropy inherited from the Lagrangian phase. The optimizer therefore regularizes shear and skewness relative to this anisotropic local target, rather than driving the mesh toward an isotropic configuration. The target construction also includes a maximum aspect-ratio constraint. This provides a direct way to limit very thin elements that would otherwise reduce the timestep.

The linesweep and TMOP components are used for different parts of the rezoning problem. The linesweep is tied to the logical structure of the ALE-AMR mesh. It acts on extended mesh-line stencils and provides anisotropy-preserving smoothing of directional spacing. The TMOP formulation used here is local and unstructured. It operates on nodal tetrahedral neighborhoods and controls shear, skewness, and excessive distortion relative to prescribed local targets. The TMOP step alone is not sufficient in this setting: it regularizes the local shape relative to the target, but it does not provide the nonlocal smoothing needed to avoid irregular local minima. Adding this smoothing directly to the TMOP target would require reintroducing length- or volume-based information, and would therefore recover the isotropizing behavior that the method is intended to avoid. For this reason, the method combines the linesweep and TMOP steps rather than replacing one by the other.

We also change the update procedure. Instead of computing nodal corrections from a frozen mesh configuration and applying them in bulk, we use a patch-based objective on an extended nodal patch, referred to below as patch-1. The metric is evaluated on this enlarged patch so that the objective couples the motion of a node to the quality of neighboring elements affected by that motion. With a sequential

update, an accepted step is therefore tested against the updated quality of the local coupled patch, rather than only against the elements directly incident on the moved node. This avoids one of the main weaknesses of the previous bulk strategy, where neighboring nodes could each take acceptable steps with respect to the old mesh, but together produce an invalid or degraded configuration after the update.

A fully sequential update would be too expensive in parallel. To recover parallelism, we use a distance-1 coloring of the nodal adjacency graph, following standard parallel coloring and maximal-independent-set ideas [23, 24]. Nodes assigned the same color are chosen so that they do not directly interact through the patch-based objective or through hanging-node reconstruction. They can therefore be updated concurrently. This block Gauss-Seidel ordering is not strictly order-independent, but it provides a practical compromise between strictly sequential updates, which are robust but slow, and fully bulk updates, which are fast but fragile.

The same patch framework is also used for disentangling. When local inversions are detected, the TMOP quality objective is replaced by a convexity-based objective. The optimizer then drives the affected tetrahedra back toward positive signed volume before returning to the TMOP quality objective.

Contributions and novelty. The method combines elements that have been studied separately, including line-based smoothing, optimization-based rezoning, untangling metrics, and TMOP-style quality metrics. The contribution of this work is their integration into a patch-based workflow suitable for ALE-AMR and block-structured meshes. In particular, we introduce: (i) a patch-1, corner-based TMOP formulation with dynamic per-tetrahedron targets that preserve anisotropy already present in the Lagrangian mesh while allowing an explicit aspect-ratio cap; (ii) a two-mode optimization strategy that uses the aspect-ratio-preserving TMOP objective when the local nodal patch is valid, and switches to a convexity-driven disentangling objective when local inversions are detected; and (iii) a parallel execution strategy based on conservative distance-1 node coloring and block Gauss-Seidel updates, designed to remain valid in the presence of hanging-node reconstruction across AMR level transitions.

The paper is organized as follows. In Sec. 2, we present the objective metrics, namely the TMOP and convexity objectives, and detail the patch-based framework and block Gauss-Seidel update. In Sec. 3, we present numerical test cases. Finally, in Sec. 4, we present our conclusions.

2. Rezoning algorithm

2.1. Mesh topology and geometric notation

We first introduce the mesh-topology notation used to define the nodal neighborhoods and tetrahedral stencils on which the metrics are evaluated. The numerical quantities associated with the TMOP objective, the convexity metric, and the line-search procedure are introduced later, when each part of the algorithm is described.

- $\mathcal{P}$: set of mesh nodes.
- $\mathcal{C}$: set of mesh cells.
- $\mathcal{P}(c)$: set of nodes associated with cell c, including hanging nodes introduced by non-conformal refinement.
- $\mathcal{C}(p)$: set of cells incident on node p.
- $\mathcal{P}_0(p)$: set of corner nodes incident on node p.
- $\mathcal{S}_p$: nodal-star tetrahedral stencil associated with node p; this is a set of oriented corner-based tetrahedra, each having p as its apex.
- $\mathcal{Z}_p$: extended nodal patch, referred to as patch-1 in what follows and used for patch-based optimization; $\mathcal{Z}_p$ contains $\mathcal{S}_p$ together with the tetrahedra whose apex node belongs to $\mathcal{P}_0(p)$, i.e. whose geometry depends on $\mathbf{x}_p$.
- $\tau = (\mathbf{x}_p, \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$: an oriented tetrahedron used in the metric evaluation.
- $\alpha(\tau)$: signed Jacobian determinant of tetrahedron τ, used as a signed volume surrogate.

2.2. Cell structure and corner nodes

In the present implementation, the mesh is a block-structured hexahedral mesh with optional non-conformal AMR. The block and refinement topology is managed with `p4est` [25]: the domain is represented as a forest of octrees, where each tree corresponds to a logically Cartesian hexahedral root block, and where inter-tree connectivity defines the global block topology. Local refinement splits a hexahedral cell into eight children per refinement level, and the deepest elements of the forest define the active computational cells.

The method therefore targets structured or block-structured hexahedral meshes, including both conformal and non-conformal AMR configurations. Independantly, each mesh block is logically Cartesian, but their union can be performed arbitrarily as long as face connectivity is well defined, such that the global mesh is in general not logically Cartesian.

In the non-conformal AMR case, hanging nodes appear on edges and faces at 2:1 refinement transitions. Because of these non-conformal neighbor relations, an active cell may possess a variety of local topological configurations. An isolated unrefined cell has the standard hexahedral topology, with six quadrilateral faces and eight corner nodes. When neighboring cells are refined, however, additional hanging nodes appear on the corresponding faces and edges of the coarse cell. In general, the cells may have faces with four to eight nodes (due to edge-hanging refinements on neighbors), and in the extreme case may be surrounded by up to twenty-four quadrilateral faces when the current cell is not refined but all neighboring cells are.

In each cell, we differentiate between *corner nodes*, i.e. the eight original vertices of the root (unrefined) hexahedral cell, and the other nodes of the cell are topologically hanging nodes. This distinction is used throughout the paper to define some nodal neighborhoods, patches, and geometric constructions that rely on the invariant corner-based structure of the mesh while remaining compatible with non-conformal refinement. An illustration is given in Fig. 1.

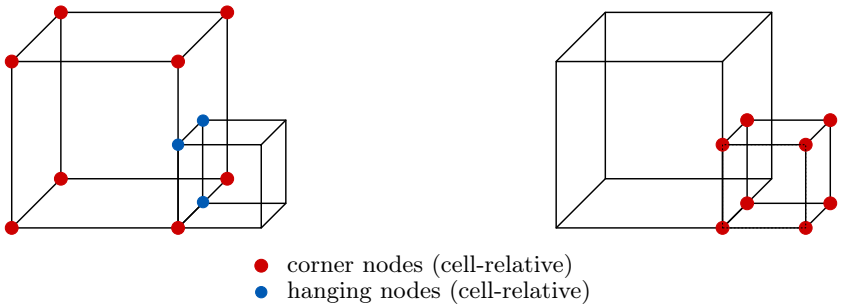

Figure 1: Illustration of cell-relative corner and hanging nodes in a non-conformal AMR interface. A coarse and a fine hexahedral cell share only a quarter of a face due to the level difference, inducing edge- and face-hanging nodes on the interface. The corner nodes of the cells are highlighted in red for the coarse cell (right) and the fine cell (left). The blue nodes are hanging in the coarse cell.

2.3. Definition of the stencil for nodal metric computations

The mesh cells have a polyhedral structure that varies depending on local and neighborhing topology. As in [12, 21], we consider quality metrics that are evaluated using the Jacobian properties of tetrahedra incident on the nodes. One such tetrahedra is illustrated in Fig. 2. The approach in [21] uses a node-star tetrahedra stencil denoted $\mathcal{S}_p$, and in what follows, we extend it to a patch-1 stencil denoted $\mathcal{Z}_p$.

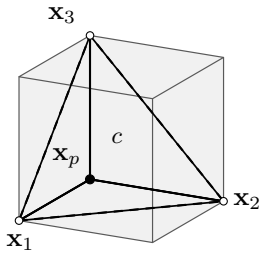

Figure 2: Illustration of a tetrahedron τ used for nodal metric evaluation. In a hexahedral cell c (light grey), and for a node p located at $\mathbf{x}_p$, we define a tetrahedron $\tau = (\mathbf{x}_p, \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$ by gathering the three corner nodes of c that are edge-connected to p. Such a tetrahedron is built in each cell incident on p, yielding the nodal-star stencil $\mathcal{S}_p$ used for the metric evaluation.

2.3.1. Node-star stencil $\mathcal{S}_p$

We start by defining the usual node-star stencil used in nodal-based metric evaluation, adapted to the non-conformal mesh. Let $\mathbf{x}_p \in \mathbb{R}^3$ denote the coordinate of node p, and the *star* of p be the set of corner nodes incident on p, denoted $\mathcal{P}_0(p)$. For each node p, let $\mathcal{S}_p$ be the set of (oriented) tetrahedra in the star of p constructed from the cells $\mathcal{C}(p)$. Each tetrahedron $\tau \in \mathcal{S}_p$ contains p as its apex, and is constructed from the corner nodes incident on p in that cell, as illustrated in Fig. 3(left). All the nodes involved in the tetrahedra set $\mathcal{S}_p$ are in the node set $\mathcal{P}_0(p)$.

Here we underline that the metric optimization is only applied - through $\mathcal{S}_p$ - to the corner nodes of the mesh. These define the independent geometric degrees of freedom of the AMR topology. As such, the hanging nodes are not optimized independently; instead, their positions are reconstructed from their parent entities (in this paper, by barycentric interpolation). This design preserves the AMR topology and avoids ill-conditioned optimization problems. The method retains its AMR specificity through the metric formulation itself, namely; volume-insensitive, target-based measures that preserve aspect ratios across 2:1 level transitions, and through the complementary action of the line-sweep rezoner.

2.3.2. Patch-1 stencil $\mathcal{Z}_p$

We now define the extended node stencil (patch-1 stencil) $\mathcal{Z}_p$ as the corner-based patch of tetrahedra around node p, constructed as:

$$\mathcal{Z}_p = \mathcal{S}_p \cup \left[\bigcup_{p' \in \mathcal{P}_0(\mathrm{p})} \tau \in \mathcal{S}_{p'} \mid p \in \tau \right],$$

that is, $\mathcal{Z}_p$ is the union of all tetras with an apex in $\mathcal{P}_0(p)$, plus the tetras in $\mathcal{S}_p$. Said differently, these are all the tetras that have p as one of their vertex. An illustration in 2D is given in Fig. 3(right). For a Cartesian 3D mesh and away from boundary conditions, $\mathcal{S}_p$ contains 8 tetrahedra (one for each cell incident on the node), and $\mathcal{Z}_p$ contains 32 tetrahedra; 8 from $\mathcal{S}_p$ and 3 more for each cell incident on p, i.e. tetras in those cell that have p as one of their node.

Note that because $\mathcal{Z}_p$ only includes tetra defined from $\mathcal{C}(p)$ by construction, the distance-1 coloring algorithm introduced in Sec. 2.7 remains compatible with a parallelization scheme where the ghost mesh consists of a single layer of nodal neighbors.

2.4. Definition of the optimization metric

In what follows we will denote the nodes of the tetrahedron $\tau_j(\mathbf{x}_p) = (\mathbf{x}_p, \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$, where the index j signifies that the vertices $\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3$ depend on which $\tau_j \in \mathcal{Z}_p$ is selected. The ordering of vertices in τ is defined such that their signed volume surrogate $\alpha(\tau) > 0$ when the mesh is known to be "valid", where the volume surrogate reads:

$$\alpha(\tau_j) = \det\big[\, \mathbf{x}_1 - \mathbf{x}_p,\ \mathbf{x}_2 - \mathbf{x}_p,\ \mathbf{x}_3 - \mathbf{x}_p \,\big],$$

noting that the true signed volume is $\alpha/6$. When the mesh is regridded through AMR, one has to reconstruct this ordering, which can be done from topological relations.

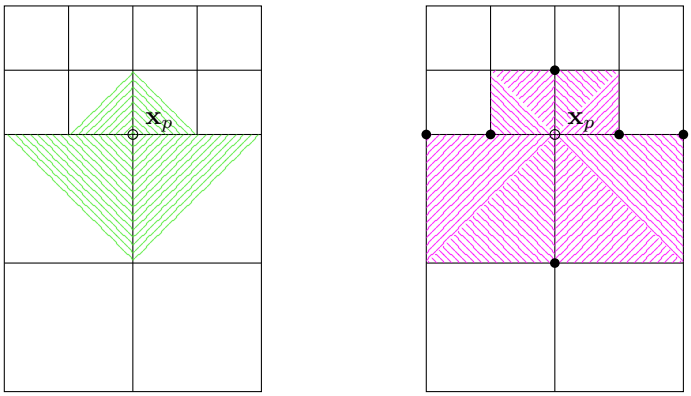

Figure 3: 2D schematic illustration of the stencil used for a node p indicated by a circle. The left panel shows the usual $\mathcal{S}_p$ stencil, where each green triangle (tetra in 3D) is formed by gathering the 2 (3 in 3D) nodes in each cell that are incident on p. The right panel shows the additional magenta triangles added to obtain $\mathcal{Z}_p$, i.e. the triangles only incident onto p by 1 node, such that the union of the left and right panel triangles form $\mathcal{Z}_p$. The black circle indicates the positions of nodes p' associated to the magenta triangles with respect to the stencils $\mathcal{S}_{p'}$, i.e. the nodes p' for which the tetra metrics ϕ are evaluated. These triangles also illustrate how only corner nodes are considered for cells containing topologically-hanging nodes.

2.4.1. TMOP formulation

We now define the TMOP quantities used in the local objective. The general TMOP framework follows [13, 17], and the specific choice of target matrix used here is adapted to the ALE rezoning objective, namely preserving the local Lagrangian aspect ratios while allowing an explicit aspect-ratio cap.

For $\tau_j = (\mathbf{x}_p, \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3) \in \mathcal{Z}_p$, we define the current element Jacobian:

$$J_j = \big[\, \mathbf{x}_1 - \mathbf{x}_p,\ \mathbf{x}_2 - \mathbf{x}_p,\ \mathbf{x}_3 - \mathbf{x}_p \,\big], \qquad \det J_j = \alpha(\tau_j).$$

This Jacobian maps the local tetrahedral coordinate frame to the current physical tetrahedron. The sign of $\det J_j$ gives the orientation of the tetrahedron (positive means that it is convex), while its magnitude is proportional to its volume.

We define the raw target edge lengths:

$$a^0_{i,j} = \|\mathbf{x}_{i,j} - \mathbf{x}_p\|_2, \qquad i = 1, 2, 3,$$

where the superscript 0 denotes the uncapped value. We then introduce a parameter $r_{\max} \geq 1$ that represents the maximum target aspect ratio we wish to preserve. When it is active, the smallest target lengths are raised so that the ratio between the largest and smallest target lengths does not exceed $r_{\max}$:

$$a^0_{\max,j} = \max_i a^0_{i,j}, \qquad a_{i,j} = \max\left(a^0_{i,j}, \frac{a^0_{\max,j}}{r_{\max}}\right).$$

If no cap is used, we set $a_{i,j} = a^0_{i,j}$. These three (optionally capped) lengths define the local Lagrangian aspect ratios that we wish to preserve in the TMOP target.

We then define the unit-determinant diagonal target:

$$T_j = \operatorname{diag}\left(\frac{a_{1,j}}{g_j}, \frac{a_{2,j}}{g_j}, \frac{a_{3,j}}{g_j}\right), \qquad T_j^{-1} = \operatorname{diag}\left(\frac{g_j}{a_{1,j}}, \frac{g_j}{a_{2,j}}, \frac{g_j}{a_{3,j}}\right),$$

where the scalar $g_j = (a_{1,j}a_{2,j}a_{3,j})^{1/3}$ is the geometric average of the three target edge lengths. Normalizing by g_j removes the overall size of the local tetrahedron from the target, so that the target matrix retains only relative edge-length information. By construction, $\det T_j = 1$. This is important because the target encodes

local aspect ratios without encoding the local volume scale. The TMOP objective therefore regularizes the element relative to the anisotropic target, without directly penalizing the volume differences naturally present in Lagrangian and AMR configurations.

The aspect-ratio cap acts only on the target geometry. It prevents the ideal element from becoming arbitrarily thin, while preserving the dominant anisotropy and without enforcing isotropy. In practice, we have found this to be an effective control parameter for limiting extreme aspect ratios that naturally arise in Lagrangian calculations, particularly in three dimensions. In such cases, strongly elongated cells can become geometrically ill-conditioned, with cell centroids and face barycenters drifting arbitrarily close to, or even beyond, opposing faces. This effect is especially pronounced in spherical geometries due to curvature effects. While the numerical scheme itself may formally tolerate these configurations, they tend to make physics models and geometric remapping fragile.

Finally, following the standard TMOP construction, we define the target-normalized Jacobian as:

$$M_j = J_j T_j^{-1}.$$

This matrix represents the current tetrahedron after the prescribed target anisotropy has been factored out. The TMOP nodal objective metric is then applied to M_j.

2.4.2. TMOP nodal objective

Among the possible TMOP metrics acting on the target-normalized Jacobian M_j, we adopt the Inverse Mean Ratio (IMR) metric [13]. It is defined as:

$$\phi_j^{\mathrm{TMOP}}(\mathbf{x}_p) = \begin{cases} \dfrac{\|M_j\|_F^2}{3\,\det(J_j)^{2/3}}, & \det(J_j) > 0, \\ \phi_{\max}, & \det(J_j) \le 0 \quad \text{(large barrier)}, \end{cases}$$

where $\phi_{\max}$ is a large number, e.g. 10^{300}, and $\|\cdot\|_F$ denotes the Frobenius norm $\|M_j\|_F^2 = \mathrm{tr}(M_j^T M_j)$.

The choice of IMR metric is motivated by its intrinsic properties. For det-1 target matrices, the IMR functional attains its minimum when M_j is orthogonal, i.e. when $M_j^T M_j = I$ [13, 17]. In this case, the physical Jacobian can be written as $J_j = R T_j$ with R an orthogonal matrix, meaning that the element matches the prescribed anisotropic target up to a rigid rotation. Geometrically, this corresponds to a right-angle nodestar configuration once the target anisotropy has been factored out. Deviations from orthogonality in M_j-such as shear or skew are therefore directly penalized by the IMR metric, while the relative edge-length ratios encoded in T_j are preserved. In contrast, alternative metrics such as the condition number admit broad flat minima and do not uniquely penalize skewed but well-conditioned configurations [26, 13], while purely norm-based metrics tend to couple shape and size and may promote isotropization or excessive element shrinking [14].

2.4.3. TMOP objective aggregation

The TMOP contribution of each tetrahedron $\tau_j \in \mathcal{Z}_p$ is measured by the scalar value $\phi_j^{\mathrm{TMOP}}(\mathbf{x}_p)$ defined above, which quantifies shape distortion relative to the prescribed anisotropic target through the IMR metric applied to the target-normalized Jacobian M_j. These per-tetra contributions are aggregated into a single nodal objective using an RMS norm:

$$\Phi_p(\mathbf{x}_p) = \Big(\frac{1}{|\mathcal{Z}_p|} \sum_{\tau_j \in \mathcal{Z}_p} \big(\phi_j^{\mathrm{TMOP}}(\mathbf{x}_p)\big)^2 \Big)^{1/2},$$

such that the objective represents the magnitude of the elements distortion in the distance-1 patch surrounding node p.

Here we emphasize the following design philosophy of the algorithm. Because Φ_p operates on the extended patch $\mathcal{Z}_p$, tetrahedra in the immediate neighborhood of p contribute directly to the local objective. In valid configurations, the IMR metric increases smoothly as $\det J_j = \alpha(\tau_j)$ approaches zero from above, and therefore provides a barrier-like resistance to loss of orientation before an inversion occurs. If an invalid

configuration is encountered, the metric is replaced by the large value $\phi_{\max}$ and the algorithm switches to the convexity-driven disentangling objective. Combined with the sequential line-search update, this makes accepted moves sensitive to the validity constraints of both the current node and its immediate neighbors.

2.5. Definition of the disentanglement metric

The TMOP optimization framework described above is dependent on $\det(J_i) \equiv \alpha(\tau_j) > 0$ for all elements in $\mathcal{Z}_p$. When that is not the case, our method no longer describe a smooth path away from invalidity. In that case, we switch to a disentanglement (convexity) metric optimization. In general, the use of a signed Jacobian determinant as a validity indicator is standard in optimization-based mesh untangling. For polygonal elements, a negative signed area or volume, equivalently a negative Jacobian determinant, indicates an inverted element, and several untangling methods are based on increasing the minimum signed area or volume in a local submesh [7, 27, 6]. In the present work, we use the same basic validity criterion, but transfer it to the corner-based tetrahedral stencils introduced above. Here, the role of the convexity-driven objective is to restore positive signed volumes before returning to TMOP quality optimization, by driving invalid tetrahedra back above a local positive signed-volume threshold defined from $\mathcal{S}_p$.

We start by defining the local threshold $\beta(p)$ used in disentangling mode. If all tetrahedra in $\mathcal{Z}_p$ have positive signed volume, no disentanglement is required and $\beta(p)$ is not used. Otherwise, $\beta(p)$ is computed from the positive volumes in $\mathcal{S}_p$ as:

$$\beta(p) = \begin{cases} \eta \,\text{median}\,\{\,\alpha(\tau) :\ \tau \in \mathcal{S}_p,\ \alpha(\tau) > 0\,\}\,, & \text{if at least two } \alpha(\tau) > 0, \\ \eta\,\alpha(\tau^+), & \text{if exactly one tetrahedron } \tau^+ \text{ satisfies } \alpha(\tau^+) > 0, \\ 0, & \text{if } \alpha(\tau) \le 0 \ \ \forall \tau \in \mathcal{S}_p. \end{cases}$$

where $\eta \in [0, 1]$ is a parameter, here set to $\eta = 0.01$, and where *median* denotes the statistical median of the finite set of positive signed volumes. Given $\beta(p) \ge 0$, we define a convexity penalty for each tetrahedron in the patch $\mathcal{Z}_p$ using a hinge function:

$$\phi_j^{\text{dis}}(\mathbf{x}_p) = \max\{0,\ \beta(p) - \alpha(\tau_j)\}\,, \qquad \tau_j \in \mathcal{Z}_p. \tag{1}$$

This penalty is zero when the signed volume surrogate $\alpha(\tau_j)$ is above the local convexity threshold, and grows linearly when the tetrahedron is inverted or insufficiently convex. In the limiting case $\beta(p) = 0$, which occurs when no positive tetrahedral volume is available in the nodal star, the penalty simply drives negative signed volumes back toward zero and then positive values. This formulation is the classical hinge penalty associated with the constraint $\alpha(\tau_j) \ge \beta(p)$.

The nodal disentanglement metric is then aggregated using an RMS norm:

$$\Psi_p(\mathbf{x}_p; \beta) = \left(\frac{1}{|\mathcal{Z}_p|} \sum_{\tau_j \in \mathcal{Z}_p} \left(\phi_j^{\text{dis}}(\mathbf{x}_p)\right)^2 \right)^{1/2}. \tag{2}$$

As written here, the disentanglement metric aims to restore validity, i.e. enforce $\alpha(\tau_j) > 0$, and to drive the inverted or nearly degenerate tetrahedra toward a level of convexity that is comparable with the surrounding valid elements. The threshold $\beta(p)$ defines a local target volume scale, expressed as a fraction of the typical positive volumes in the nodal neighborhood of p. More precisely, $\beta(p)$ is defined with respect to elements in the nodal star $\mathcal{S}_p$, while the metric function is optimized over the patch-1 stencil $\mathcal{Z}_p$. In this sense, the disentangler aims to recover a minimally robust configuration rather than a marginally valid one, which avoids repeated transitions between valid and invalid states and reduces sensitivity to small perturbations during the following TMOP phase. This convexity/volume-driven untangling metric is conceptually similar to the class of methods presented in [7].

The use of the median in $\beta(p)$ is an implementation choice; in anisotropic Lagrangian or AMR configurations, the signed volumes in $\mathcal{S}_p$ may have strong variation around the node. An arithmetic average can then

be dominated by a single large tetrahedron, producing an overly large convexity threshold, or be strongly affected by a small number of nearly degenerate tetrahedra. The median is a more robust local estimate of the typical positive signed volume. In our tests, this choice was more stable than using the average, especially near nonuniform AMR transitions and highly anisotropic cells.

2.6. Nodal function minimization

We have defined the per-node objectives:

$$F_p(\mathbf{x}_p) = \begin{cases} \Psi_p(\mathbf{x}_p; \beta(p)), & \text{if any } \alpha(\tau) \le 0 \text{ in } \mathcal{Z}_p \quad \text{(disentangling mode)}, \\ \Phi_p(\mathbf{x}_p), & \text{otherwise} \quad \text{(TMOP mode)}. \end{cases}$$

In this section, we describe how these objectives are optimized for a single node. The block Gauss-Seidel update is given later in Sec. 2.7. Here, this is achieved with a gradient-descent step bounded by a local maximum radius. An activation function is also used to control step length and activation of the optimization.

2.6.1. Optimization radius

In order to accommodate both mesh-quality optimization and disentanglement, we distinguish two step-length scales associated with node p.

We first define a *characteristic* geometric length scale d_p^{char} based on the nodal star:

$$d_p^{\text{char}} = \min_{\tau \in \mathcal{S}_p} \text{dist}[\mathbf{x}_p, \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}],$$

where $\text{dist}[\mathbf{x}_p, \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}]$ is the impact-parameter distance between the position $\mathbf{x}_p$ and the plane of the opposite face $\{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}$ of tetrahedron τ. This quantity provides the local geometric scale used to bound metric optimization steps.

When at least one tetrahedron in the nodal star is inverted, we additionally define a disentanglement length scale d_p^{dis} driven by the inverted configurations:

$$\tau^{\min} = \arg \min_{\tau \in \mathcal{S}_p \,:\, \alpha(\tau) < 0} \alpha(\tau), \qquad d_p^{\text{dis}} = \text{dist}[\mathbf{x}_p, \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3\}_{\tau^{\min}}],$$

which provides a characteristic length scale associated with the most severely inverted tetrahedron (in the signed-volume sense). In practice, the inversion of a node typically produces several inverted tetrahedra in the nodal star, such that the descent direction is generally not aligned with the normal direction of a single face. As a result, this step length is not sufficient to disentangle a node in a single update.

Moreover, d_p^{dis} can become arbitrarily small as $\mathbf{x}_p$ approaches the plane of an inverted tetrahedron. A similar degeneracy can arise in cells with very high aspect ratio, where the impact-parameter distance to a face may be vanishingly small despite finite edge lengths. To prevent the optimization radius from collapsing in such configurations, we introduce an additional characteristic edge-based length scale:

$$e_p^{\text{char}} = \text{median}\{\, \|\mathbf{x}_p - \mathbf{x}_{p'}\|_2 \;:\; p' \in \mathcal{P}_0(p) \,\},$$

defined as the median edge length connecting node p to its neighboring vertices in the nodal star.

The maximum admissible step length for node p is then defined as:

$$\ell_p = \begin{cases} \gamma_{\text{opt}}(p)\, \max\!\big(10^{-2} e_p^{\text{char}},\, d_p^{\text{char}}\big), & \text{if } \alpha(\tau) > 0; \forall \tau \in \mathcal{S}_p, \\ \gamma_{\text{dis}}(p)\, \max\!\big(10^{-2} e_p^{\text{char}},\, d_p^{\text{dis}}\big), & \text{otherwise.} \end{cases}$$

The coefficients $\gamma_{\text{opt}}(p)$ and $\gamma_{\text{dis}}(p)$ are dimensionless activation factors. In the present implementation, the TMOP activation is binary:

$$\gamma_{\text{opt}}(p) = \begin{cases} \gamma_{\text{opt}}^0, & \Phi_p(\mathbf{x}_p) > \Phi_p^{\text{thr}}, \\ 0, & \Phi_p(\mathbf{x}_p) \le \Phi_p^{\text{thr}}, \end{cases} \qquad \gamma_{\text{opt}}^0 = 50,$$

where $\Phi_p^{\text{thr}} \geq 1$ is the prescribed activation threshold. In disentangling mode, invalid nodes are assigned

$$\gamma_{\text{dis}}(p) = \gamma_{\text{dis}}^0, \qquad \gamma_{\text{dis}}^0 = 1.1.$$

In disentangling mode, this formulation ensures that an inverted configuration ($\alpha(\tau) < 0$) admits a step length sufficient to move $\mathbf{x}_p$ across the plane of the opposite face, thereby allowing the tetrahedron to recover a positive signed volume. This geometric interpretation is illustrated in Fig. 4. When several tetrahedra are inverted simultaneously, the descent direction is generally not aligned with a single face normal, and validity is recovered through successive line-search iterations.

In TMOP mode, the step length remains scaled by the smallest local geometric scale, preventing excessive motion relative to neighboring elements, which is desirable for swept-face remapping algorithms that proceed in a neighbor-to-neighbor fashion. The relatively large value $\gamma_{\text{opt}}^0 = 50$ gives active TMOP nodes enough freedom during the line search and was found to reduce the sensitivity of the result to the order in which color classes are processed for block Gauss-Seidel updates (see Sec. 2.7).

Note that ℓ_p does not, by itself, prevent a node from transitioning from a valid to an invalid configuration. Enforcing such a constraint at this stage would be overly restrictive and could lead to vanishingly small step lengths in regions of fine resolution, effectively preventing mesh improvement. Instead, robustness and validity are ensured by the patch-based objectives and the sequential update strategy.

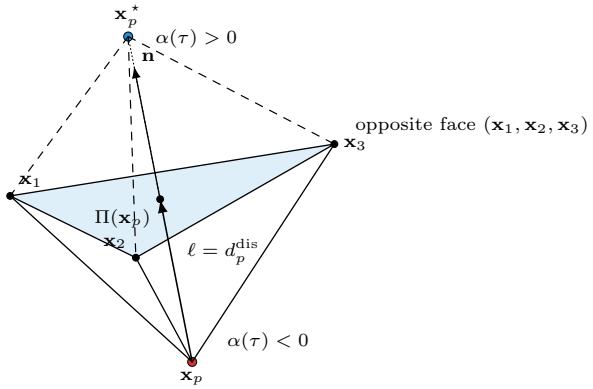

Figure 4: Illustration of how the choice of ℓ_p must be sufficient to allow for flipping an inverted nodal tetrahedron back to a convex state. The opposite face $(\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$ defines a plane $\Pi(\mathbf{x}_p)$ with unit normal $\mathbf{n}$. The tetrahedron $\tau(\mathbf{x}_p) = (\mathbf{x}_p, \mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3)$ is inverted when $\mathbf{x}_p$ lies on the wrong side of this plane ($\alpha(\tau) < 0$). Choosing a step-length scale ℓ at least equal to the orthogonal distance from $\mathbf{x}_p$ to the plane (here $\ell = d_p^{\text{dis}}$) guarantees that $\mathbf{x}_p$ can be moved across the plane, yielding a valid configuration $\mathbf{x}_p^\star$ with $\alpha(\tau) > 0$.

2.6.2. Gradient descent and line search

The objective gradient is approximated in a local orthonormal basis $\{\mathbf{e}_1(p), \mathbf{e}_2(p), \mathbf{e}_3(p)\}$ using central finite differences:

$$\partial_{e_i} F_p(\mathbf{x}_p) \;\approx\; \frac{F_p(\mathbf{x}_p + h\,\mathbf{e}_i(p)) - F_p(\mathbf{x}_p - h\,\mathbf{e}_i(p))}{2h}, \qquad i = 1, 2, 3,$$

with a local step size $h = \max(10^{-8} e_p^{\text{char}}, \min(10^{-2} e_p^{\text{char}}, \theta \ell_p))$ where $\theta = \frac{1}{4}$, and R_p is the rotation matrix:

$$R_p = \big[\, \mathbf{e}_1(p)\ \mathbf{e}_2(p)\ \mathbf{e}_3(p) \,\big] \in \mathbb{R}^{3\times 3},$$

whose columns are the basis vectors expressed in Cartesian coordinates.

The basis $\mathbf{e}_i$ is chosen according to the global mesh geometry (Cartesian, cylindrical, or spherical) and modified as needed to conform with local symmetry constraints (e.g. boundary conditions). More details on that procedure, as well as the handling of free boundaries, are given in [21]. The gradient

expressed in the local frame $\nabla_e F_p = (\partial_{e_1} F_p, \partial_{e_2} F_p, \partial_{e_3} F_p)$ is then mapped back to Cartesian coordinates as $\nabla F_p(\mathbf{x}_p) = R_p \, \nabla_e F_p$, which defines a descent direction in Cartesian space as $\mathbf{s}_p = -\frac{\nabla F_p}{\|\nabla F_p\|_2}$.

The nodal position is updated using a backtracking line search along $\mathbf{s}_p$, with trial steps bounded by the mode-dependent maximum length ℓ_p defined previously. We initialize the trial step with:

$$t_0 = \ell_p, \tag{3}$$

and test a decreasing sequence of candidate steps:

$$t_m = t_0 \, B^m, \qquad m = 0, 1, 2, \ldots, \tag{4}$$

where m is the backtracking index and $B \in (0, 1)$ is the reduction factor. We select the first t_m satisfying the Armijo condition [28]:

$$F_p(\mathbf{x}_p + t_m \, \mathbf{s}_p) \leq F_p(\mathbf{x}_p) - A \, t_m \, \|\nabla F_p(\mathbf{x}_p)\|_2, \tag{5}$$

where A is a constant. If no such t_m is found, the step yielding the largest decrease of F_p is selected. In this work, the Armijo parameters are set to $A = 0.2$ and $B = 0.7$.

2.6.3. Coupling disentanglement with local smoothing

In severely distorted configurations, the feasible region of the disentanglement metric may become extremely small, particularly when multiple neighboring nodes are invalid. In such cases, a purely local disentanglement update can stall, even though a nearby valid configuration exists. To address this issue, the disentanglement iterations are periodically complemented by a mesh rezoning optimization step applied only to invalid nodes and to nodes adjacent to invalid ones. Here, this is done when the disentangler fails to disentangle a node for more than 20 consecutive steps. The optimization entails a weighted line-sweep pass and a TMOP pass. This hybrid strategy has been found to provide a robust escape mechanism in highly tangled regimes.

2.7. Distance-1 node coloring for block Gauss-Seidel updates

Up to this point, we have described an optimization framework that improves a nodal quality metric while guaranteeing that neither the node itself nor its immediate neighbors can be tangled. This property follows from the use of barrier-based objective functions together with a patch-1 stencil, and it relies on the assumption that a node p is moved sequentially with respect to all nodes involved in its local patch $\mathcal{Z}_p$. In other words, p must not be rezoned simultaneously with any node whose position directly enters the evaluation of F_p.

In a large-scale parallel environment, rezoning nodes strictly one by one is clearly impractical. Instead, we rely on a node coloring strategy that allows groups of nodes to be rezoned concurrently, provided that these nodes are sufficiently separated so that their optimization problems are independent. In the present context, two nodes p and p' are considered adjacent (i.e. coupled with respect to the optimization) if the position of p' appears in the patch $\mathcal{Z}_p$ used to evaluate F_p.

To avoid missing any coupling between nodes updated at the same time, we build the adjacency graph from cell connectivity. Two nodes p and p' are declared adjacent if p' belongs to the set of cell nodes associated with any cell incident on p, i.e. if p' is a node of a cell in $\mathcal{C}(p)$. This definition extends beyond the patch $\mathcal{Z}_p$ used in the evaluation of F_p. This extension is required in the nonconformal framework because some nodes are topologically hanging and are not optimized independently. Instead, we recall that their position is reconstructed from their parent nodes, which are themselves cell nodes of the incident coarse cell. As a consequence, moving a corner node may indirectly affect the position of hanging nodes that do not explicitly belong to $\mathcal{Z}_p$, but are nevertheless coupled through this parent–child relationship.

The resulting cell-based distance-1 nodal adjacency therefore provides a conservative over-approximation of the true optimization coupling. It guarantees that no two nodes updated simultaneously can interact directly through the patch-based objective or indirectly through hanging-node reconstruction. As a result, all nodes of a given color class can be rezoned in parallel without violating the assumptions underlying the patch-based optimization framework.

Let $\mathcal{P}$ denote the set of mesh nodes and $\mathcal{C}$ the set of mesh cells. For a cell $c \in \mathcal{C}$, let $\mathcal{P}(c) \subset \mathcal{P}$ be the set of nodes associated with c, including hanging nodes introduced by non-conformal refinement. The nodal adjacency graph is defined as $G = (V, E)$ with $V = \mathcal{P}$ and

$$(p, q) \in E \quad \Longleftrightarrow \quad \exists\, c \in \mathcal{C} \text{ such that } p \neq q, \;\; p \in \mathcal{P}(c) \text{ and } q \in \mathcal{P}(c),$$

where adjacency reflects potential coupling in the optimization process.

A distance-1 coloring of this graph therefore guarantees that all nodes belonging to the same color class can be rezoned in parallel without violating the assumptions underlying the patch-based optimization framework.

We compute the coloring using a distributed-memory variant of the Jones-Plassmann / Luby-style scheme [23, 24]: each node p is assigned a deterministic pseudo-random priority $\pi(p)$ (here obtained by hashing the global node identifier with a SplitMix64 hash [29]), and the algorithm repeatedly extracts a maximal independent set of currently uncolored nodes by selecting *local maxima* of π among uncolored neighbors. This selection step is synchronized across MPI partitions by exchanging the current color array on ghost nodes. Once a set of winners is determined, all winners are colored simultaneously using a greedy *first-fit* rule (smallest available color not used by already-colored neighbors). The procedure is iterated until all internal nodes are colored. This approach is closely related to classic parallel coloring heuristics based on maximal independent set selection and has good practical scalability in distributed settings. An example in 2D is given in Fig. 5.

To improve robustness and scalability in a distributed-memory setting, we apply the selection in a staged manner that prioritizes nodes with fully local neighborhoods before processing nodes near partition interfaces. This strategy reduces sensitivity to remote information while preserving the correctness of the distance-1 coloring.

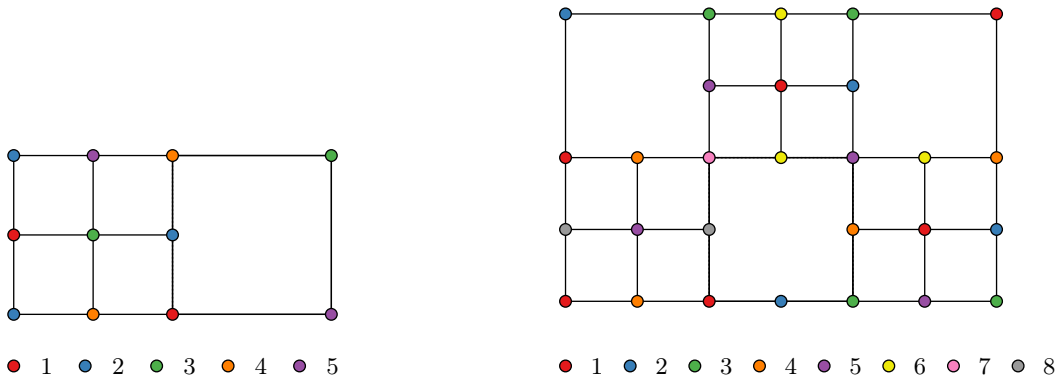

Figure 5: 2D illustration of distance-1 coloring under *cell-induced adjacency* with a 2:1 refinement constraint. (left) simple case requiring 5 colors and (right) more complex case requiring 8 colors.

3. Test cases

This section describes the mesh-only test suite used to assess (i) aspect-ratio control, (ii) global quality improvement in the presence of bunching and shear-like defects, and (iii) convergence speed and robustness of the disentangler.

When both linesweep and metric optimization are used, one rezoning iteration consists of one linesweep pass followed by one metric optimization pass (be it the normed-CN baseline evoked at the beginning or the patch-based TMOP pass). The linesweep pass first updates the mesh-line spacing and provides the nonlocal smoothing component. The metric optimization pass is then applied to the updated mesh, using the block Gauss-Seidel/coloring procedure described in Sec. 2.7 for TMOP, and the baseline Jacobi-like bulk update for the normed-CN. This second pass aims to reduce local shear and, for the TMOP, excessive aspect ratios relative to the dynamic targets. Hanging-node coordinates are reconstructed after each pass.

In [21], weights and damping parameters are introduced to control the specific behavior of the rezoner. Here, however, we only aim to showcase limitations of the linesweep and of the metric optimization that are fundamental to the method and that appear in all regimes, including at maximum rezoning. For this reason, all tests presented in this section are performed with the linesweep operating in equal-space mode, i.e. with the damping chosen so that the line relaxation reduces to the equal-space update along each logical mesh line [10, 22, 21]. In particular, using the weighted variant of the linesweep would only exacerbate the limitations highlighted here, as it would further slow down the relaxation of strongly bunched or distorted meshes. The metric optimization is also activated on all eligible nodes, rather than only through criteria-based activation functions such as e.g. in [21]. As such, the configurations considered below represent a conservative setting in which the observed limitations are not caused by under-relaxation, but instead reflect fundamental behaviors of the underlying methods. We recall that the normed-CN baseline refers to the normalized condition-number metric used in [21] to complement the linesweep, and which shortcomings were highlighted in the introduction. In this metric, the local edge directions are normalized before the condition number is evaluated, so that the objective mainly measures the loss of orthogonality between incident edges while remaining insensitive to local volume and edge-length ratios.

All test cases are conducted in a $[0, 1]$ square/cubic domain with wall-like (symmetry) boundary conditions. In addition to the geometric deformations prescribed in the tests and given below, we superimpose a *macro-zone* nonconformal refinement pattern: the unit domain is split into 2^d macro blocks (2×2 in 2D, $2\times2\times2$ in 3D), and alternating macro blocks are refined by one level (0-1 pattern). This introduces coarse-fine interfaces and hanging nodes while keeping the refinement pattern simple.

To quantify mesh quality during the test, we monitor simple geometric proxies for aspect ratio and shear, evaluated locally at each corner node of the tetrahedral decomposition. The aspect ratios are defined from the node star tetrahedra $\tau \in \mathcal{S}_p$ using the same definition as in Sec. 2.4.1. To characterize shear, we measure the maximum deviation from orthogonality between any pair of incident edges in the node-star tetrahedra at the node:

$$\mathrm{shear}(p) = \max_{\tau\in\mathcal{S}_p} \max_{i<j} \left| \frac{(\mathbf{x}_i - \mathbf{x}_p)\cdot(\mathbf{x}_j - \mathbf{x}_p)}{\|\mathbf{x}_i - \mathbf{x}_p\|_2 \, \|\mathbf{x}_j - \mathbf{x}_p\|_2} \right|.$$

This quantity corresponds to the maximum absolute cosine of the angles between edge pairs and therefore ranges from 0 (orthogonal edges) to 1 (collinear edges), providing a direct proxy for local shear and sliver-like configurations.

Both metrics are evaluated over all tetrahedra incident on all interior nodes. We report their maximum, median, and 95th percentile values, which together provide a description of worst-case behavior, typical mesh quality, and the upper tail of the distribution. Notably, the 95th percentile is the value below which 95% of the sampled metric values lie; equivalently, only the worst 5% of the samples have a larger value. Since larger aspect-ratio and shear values correspond to poorer local quality, this statistic describes the high-distortion tail of the mesh while being less sensitive than the maximum to a single isolated outlier.

3.1. 2D bunching layer; high aspect-ratio and cell bunching

This test targets cell bunching and aspect-ratio control. It initializes thin, high-aspect-ratio and bunched-up elements in the vicinity of a tight cell layer that connects to the boundaries through a moderately sheared region.

We consider the unit square $\Omega = [0, 1]^2$ with reference coordinates $(\xi, \eta) \in [0, 1]^2$. The initial mesh is a Cartesian grid in (ξ, η), mapped to a deformed configuration (x, y) that preserves the rectangular boundary. We define the coordinate mapping:

$$\begin{aligned} w(\xi,\eta) &= \big(4\xi(1-\xi)\big)^2 \big(4\eta(1-\eta)\big)^2, \\ s(\eta) &= \tanh\!\left(\frac{\eta-\eta_0}{\delta}\right), \\ x(\xi,\eta) &= \xi, \\ y(\xi,\eta) &= \eta + B\, w(\xi,\eta)\, s(\eta), \\ \eta_0 = 0.50, \quad \delta = 0.05, \quad B &= 0.35, \end{aligned}$$

which induces an interior compression/expansion band while keeping $x = \xi$ and $y = \eta$ on all boundaries. We compare; linesweep only, linesweep + normed-CN optimizer, patch-based TMOP only, and linesweep + patch-based TMOP with two different values of the aspect-ratio cap. Resulting meshes after 200 optimization iterations are shown in Fig. 6.

The initial mesh is shown in Fig. 6a. This mesh is valid and, in a purely Lagrangian setting, would not necessarily be considered of poor quality since it is smooth. However, the strong bunching of the cells creates very high–aspect-ratio elements with small volumes, which would significantly reduce the timestep through CFL-like constraints. Cells are also sheared in the region transitioning from the bunched cells to the left/right boundaries.

Figure 6b shows that the linesweep algorithm is not able to significantly de-bunch the mesh on its own, despite operating at maximum strength in equal-space mode. This is because, in a smooth configuration, each node is locally close to its equal-space position, such that the relaxation proceeds extremely slowly. We also note a significant break in the mesh alignment around the 2:1 junction, which is due to this particular checkerboard refinement configuration and is also supposed to be handled by the complement optimization metric (normed-CN or TMOP). The normed-CN, introduced in [21] primarily for shear control, is likewise unable to significantly reduce the bunching, as shown in Fig. 6c, since this metric is largely insensitive to volume ratios. The sheared region does appear less distorted, with mesh lines rotating to accommodate it; however, small local deformations develop in the same region. These artifacts are indicative of convergence toward a local minimum of the non-linear metric being optimized. The results with TMOP-only are shown in Fig. 6d. As mentioned earlier, optimising with our patch-based TMOP but without smoothing the aspect ratios does not yield a good quality mesh, and significant perturbations are present. The mesh quality improves significantly when combining it with the linesweep, as shown in Figs. 6e and 6f. Imposing an aspect-ratio cap of $r_{\max} = 100$ yields a mesh with reduced overall shear, but little de-bunching, as shown in Fig. 6e. In contrast, enforcing a tighter aspect-ratio cap of $r_{\max} = 10$ highlights the effectiveness of the method, as seen in Fig. 6f. In this configuration, the mesh is both significantly de-bunched and less sheared. By prescribing an explicit target aspect ratio, one also prevents over-relaxation of the high-resolution region, which is essential to preserve the information carried by the Lagrangian mesh. In all cases, the transition between size-1 and size-2 AMR elements remains well behaved, reflecting the specific treatments introduced to ensure consistency across refinement levels. We also underline that during the optimization process, the linesweep with normed-CN method (that does not use a patch-based metric nor a block Gauss-Seidel descent) produced 5 intermediate invalid meshes.

The evolution of the aspect ratio and shear metrics as a function of rezoning iteration is shown in Fig. 7. The initial maximum shear in the mesh is ~ 0.65 and the initial maximum aspect ratio is ~ 330. The linesweep combined with TMOP using a high aspect-ratio cap ($r_{\max} = 100$) behaves similarly to the pure linesweep in terms of aspect-ratio control, which is expected given the loose constraint. However, it significantly reduces shear while the linesweep alone is largely ineffective in this respect. Reducing the cap to $r_{\max} = 10$ results in a much faster decrease of the aspect ratio for the p-95, which reaches below 20 after 200 iterations. Although the final shear levels are comparable between the TMOP runs with $r_{\max} = 10$ and $r_{\max} = 100$, a transient spike in the maximum shear is observed for the $r_{\max} = 10$ case. This behavior is explained by the high aspect ratio in the initial configuration: enforcing a strict aspect-ratio target induces an intermediate state in which maximum shear is temporarily amplified before being relaxed. In practice, such large initial shear would not arise, as ALE rezoning would be applied continuously between each Lagrangian phase. The median and p-95 shear values remain nearly identical between the two TMOP configurations. The linesweep combined with the normed-CN yields the second lowest median shear but the high maximum shear. It also produces the large maximum aspect ratios, while its median and p-95 values are comparable to those obtained with TMOP at $r_{\max} = 100$. Finally, the TMOP-only case gives the poorest overall result; it reaches a low median shear, but retains large maximum and p-95 values for both aspect ratio and shear.

3.2. 3D swirl: strong shear at moderate aspect ratio

We now consider a three-dimensional rotated core with a thin sheared layer. Compared to the 2D bunching test, this configuration exhibits more moderate aspect ratios and weaker volumetric bunching, but introduces a genuinely three-dimensional shear with non-trivial morphology. As such, this test is not

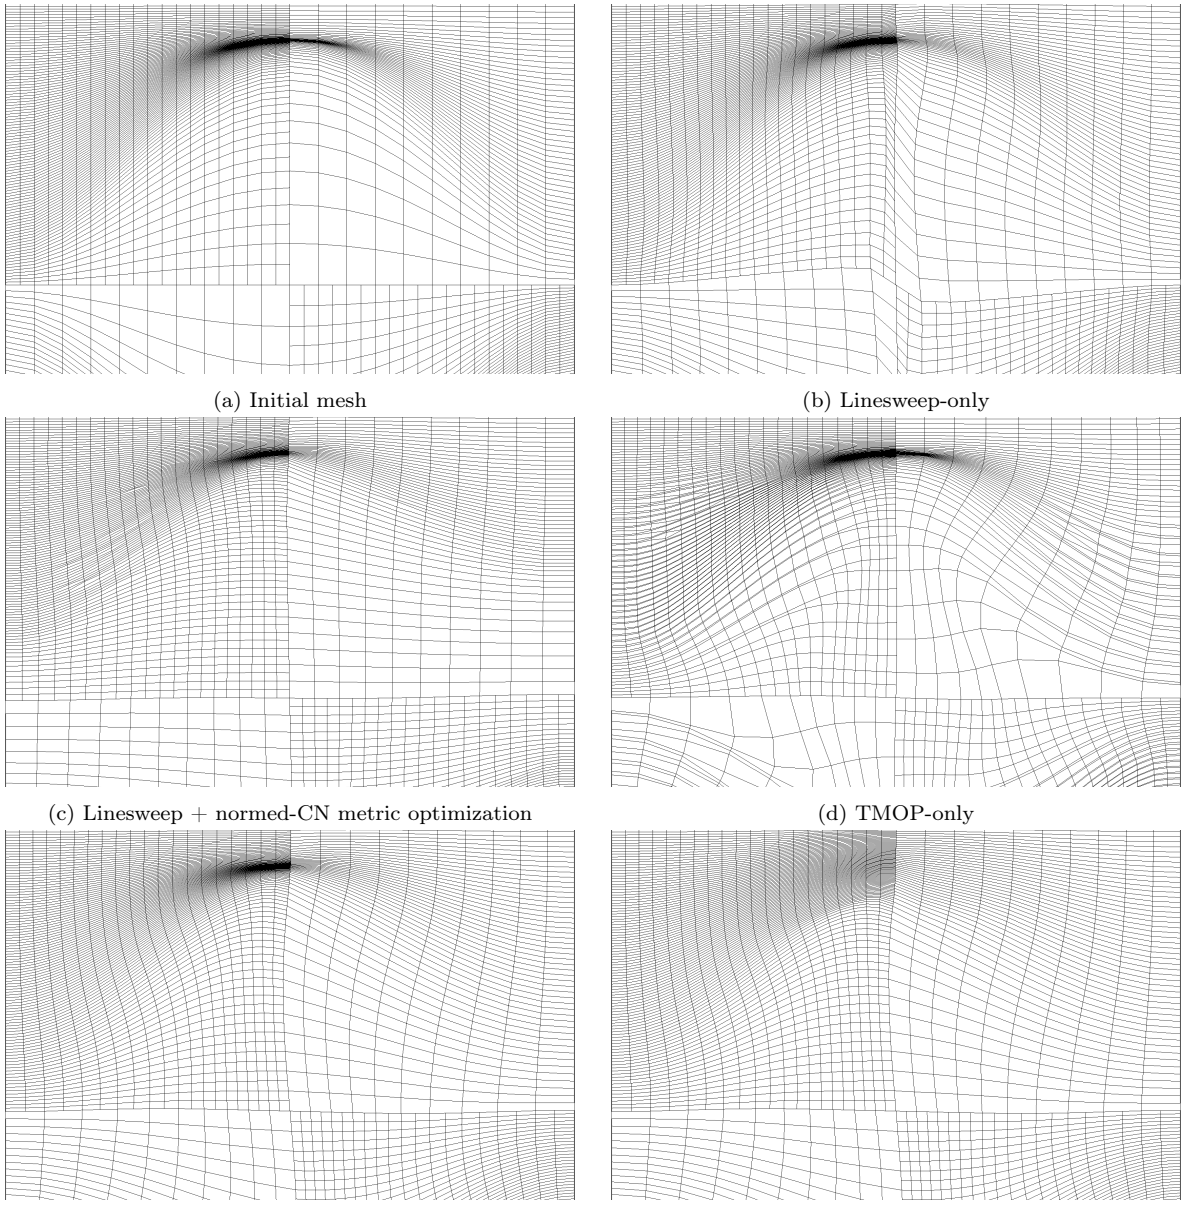

(a) Initial mesh

(b) Linesweep-only

(c) Linesweep + normed-CN metric optimization

(d) TMOP-only

(e) Linesweep + patch-based TMOP (AR cap 100)

(f) Linesweep + patch-based TMOP (AR cap 10)

Figure 6: Meshes in the 2D bunching layer test. Panel (a) shows the initial mesh. All other panels show the result of 200 optimization steps.

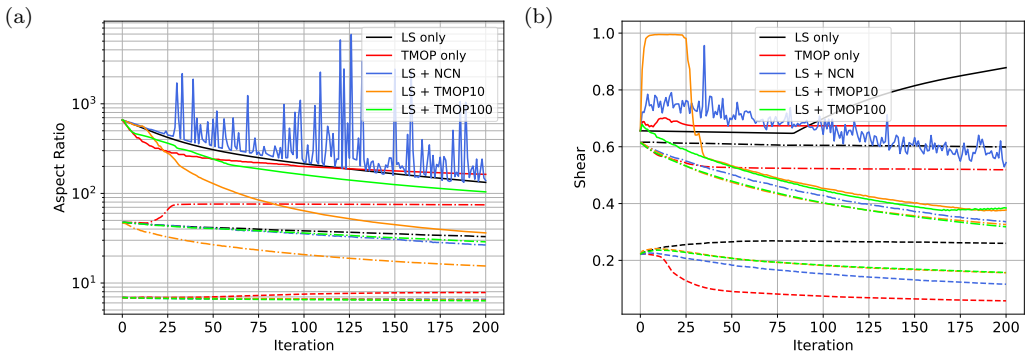

Figure 7: Statistical mesh metrics as a function of rezoning iteration number for the 2D bunching layer test, showing (a) aspect-ratio r as defined from the node star in Sec. 2.4.1, and (b) shear, defined as the maximum $|\cos\theta|$ angle in the node star tetrahedra. Each panel shows the median value (dashed lines), the maximum value (solid lines) and the 95-th percentile value (dot dashed line). Data is overlaid from the linesweep only case (black), TMOP-only case (red), linesweep with normed-CN (blue), linesweep with TMOP with $r_{\max} = 10$ (orange) and $r_{\max} = 100$ (green).

primarily intended to demonstrate aspect-ratio control. Instead, its main purpose is to assess the robustness of the rezoning algorithms in the presence of strong 3D shear, and in particular their ability to preserve mesh validity throughout the optimization.

The deformed mesh is constructed through a smooth transformation of the reference configuration,

$$\mathbf{x}(\mathbf{X}) = \mathbf{c} + (1 - w)\,\mathbf{r} + w\,\mathcal{R}_{\mathbf{a}}\big(\theta(\rho)\big)\,\mathbf{h}(\mathbf{r}), \tag{6}$$

where $\mathbf{X} = (\xi, \eta, \zeta)^T \in [0, 1]^3$ are reference coordinates, $\mathbf{c} = (1/2, 1/2, 1/2)$ is the domain center, $\mathbf{r} = \mathbf{X} - \mathbf{c}$, $\rho = \|\mathbf{r}\|_2$, $\mathcal{R}_{\mathbf{a}}(\theta)$ is a rotation matrix around axis $\mathbf{a}$ and with angle θ, and $\mathbf{h}(\mathbf{r})$ defines the radially compressed configuration that is rotated.

By construction, the mapping reduces to the identity on $\partial\Omega$, while the interior undergoes a combined rotation and radial compression. The envelope function is written as a product of one-dimensional boundary functions:

$$w(\xi, \eta, \zeta) = \omega(\xi)\,\omega(\eta)\,\omega(\zeta), \qquad \omega(u) = \big(4u(1 - u)\big)^2. \tag{7}$$

The function ω is equal to zero at $u = 0$ and $u = 1$, and reaches one at $u = 1/2$. Thus, w smoothly blends the deformation to zero on all six boundary faces and leaves the boundary fixed.

The rotation angle is prescribed through a smooth mask:

$$\theta(\rho) = \theta_0\, s(\rho), \qquad s(\rho) = \frac{1}{2}\left(1 - \tanh\left(\frac{\rho - R}{\delta}\right)\right), \tag{8}$$

with $R = 0.30$, $\delta = 0.05$, and $\theta_0 = 88°$. The rotation axis in $\mathcal{R}_{\mathbf{a}}(\theta)$ is set to $\mathbf{a} = (0, 0, 1)^T$.

To generate a thin shell of compressed elements in the shear transition region, we define:

$$\mathbf{h}(\mathbf{r}) = \mathbf{r} + \varepsilon_r\, \operatorname{sech}^2\left(\frac{\rho - R}{\delta}\right) \frac{\mathbf{r}}{\rho + \varepsilon}, \tag{9}$$

with $\varepsilon_r = -0.15$ and $\varepsilon = 10^{-12}$. Here, $\operatorname{sech}(x) = 1/\cosh(x)$, so the compression is localized near $\rho = R$. This produces a thin layer aligned with the shear transition and yields a smooth 3D swirl with localized strong shear and mild radial bunching.

We compare results between the pure linesweep and linesweep combined with normed-CN metric optimization, TMOP with $r_{\max} = 3$ and TMOP with $r_{\max} = 10$. The rezoning algorithms are run for 200 iterations. In this regime, the final meshes produced by the normed-CN and by the TMOP-based approaches are visually and quantitatively similar (see Fig. 8), in part because the linesweep is operated in equal-space mode such that little explicit aspect-ratio control is required. However, a key distinction emerges during

the optimization process: the normed-CN approach transiently generates a large number of invalid meshes (typically $\mathcal{O}(30)$ in this test), whereas the TMOP-based approach remains valid throughout. This difference highlights the importance of formulating a consistent optimization workflow when dealing with strongly sheared three-dimensional configurations. Again, because aspect ratios are more tame in this test, final mesh results for $r_{\max} = 3$ (not shown) look similar to those with $r_{\max} = 10$, but we do observe a difference in convergence speed between the two, as made explicit later. The inability for the linesweep-only to handle shear is illustrated in Fig. 8b. The normed-CN optimization, specifically formulated for shear, is highly efficient here and produces a mesh with the rotation almost undone, as seen in Fig. 8c. The TMOP case has not completely undone the rotation, and some aspect ratio information remains contained in the mesh, as seen in Fig. 8d. The mesh remains of high quality in that case. The slower relaxation of the rotated structure is expected, because the TMOP target is constructed to preserve part of the local aspect-ratio information instead of driving the mesh directly toward an isotropic configuration.

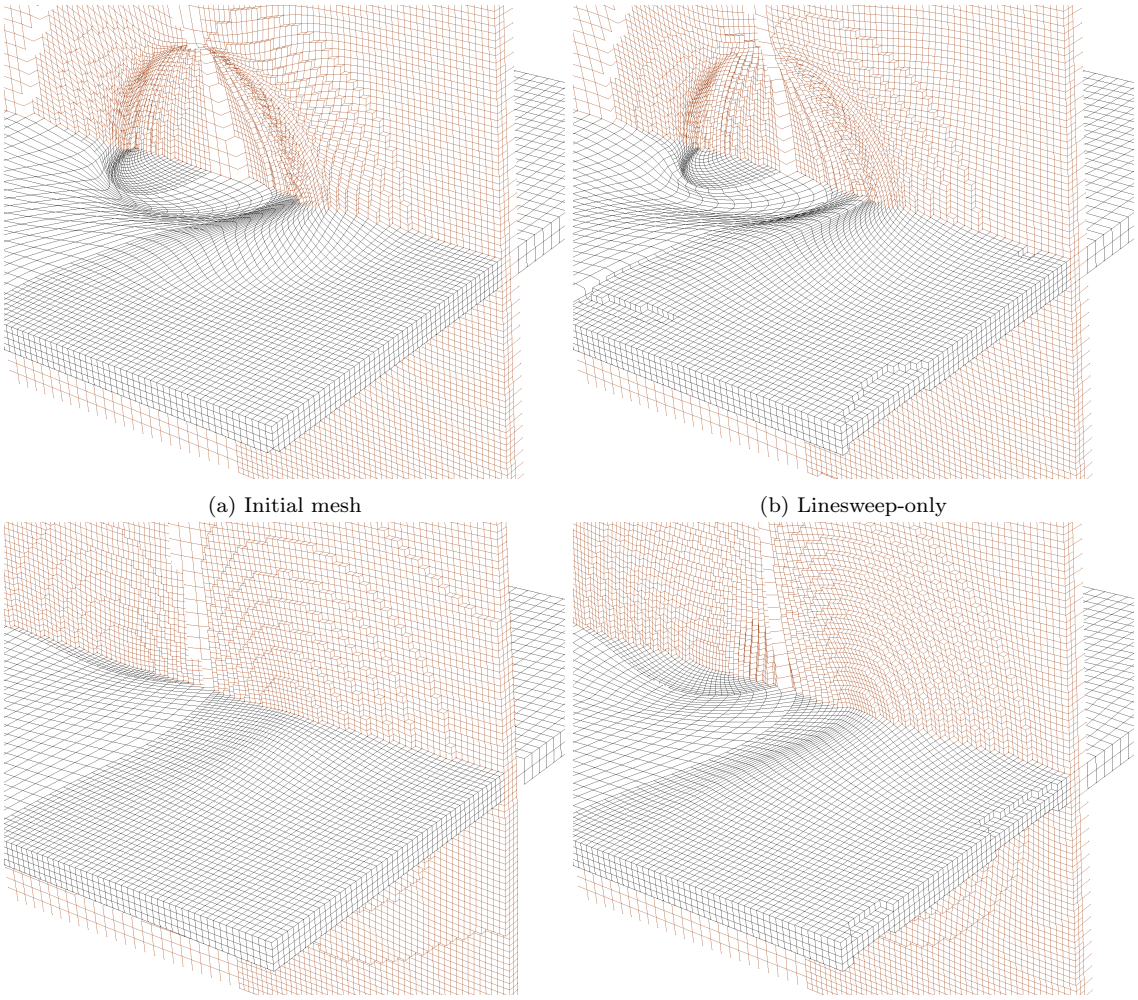

(a) Initial mesh

(b) Linesweep-only

(c) Linesweep + normed-CN metric optimization

(d) Linesweep + patch-based TMOP (AR cap 10)

Figure 8: Meshes sections in the 3D swirl test, showing mesh cells around $z = 0.5$ in black $y = 0.5$ in brown. Panel (a) shows the initial mesh. All other panels show the result of 200 optimization steps.

We now compare the shear and aspect-ratio statistics. The initial maximum shear in the mesh is ~ 1, corresponding to nearly coplanar tetrahedra in the nodal stars, and the initial maximum aspect ratio is

~ 15. Results are shown in Fig. 9. We first observe that the linesweep alone is able to reduce aspect-ratio variations, consistently with the fact that it operates here in equal-space mode, but is unable to reduce shear effectively: the maximum shear remains essentially unchanged and the p-95 decreases only very slowly. By contrast, both TMOP cases and the normed-CN case rapidly reduce the p-95 and maximum shear, with TMOP being more efficient at reducing the maximum value at early times. The TMOP cases and normed-CN also reduce the p-95 and maximum aspect ratios. Varying the value of $r_{\max}$ is shown to directly affect the rate at which aspect ratios equalize, which is a key property of the TMOP approach presented here: it provides a straightforward tunable mechanism to control aspect-ratio relaxation in the mesh, and thus to better preserve the structure of the underlying Lagrangian mesh.

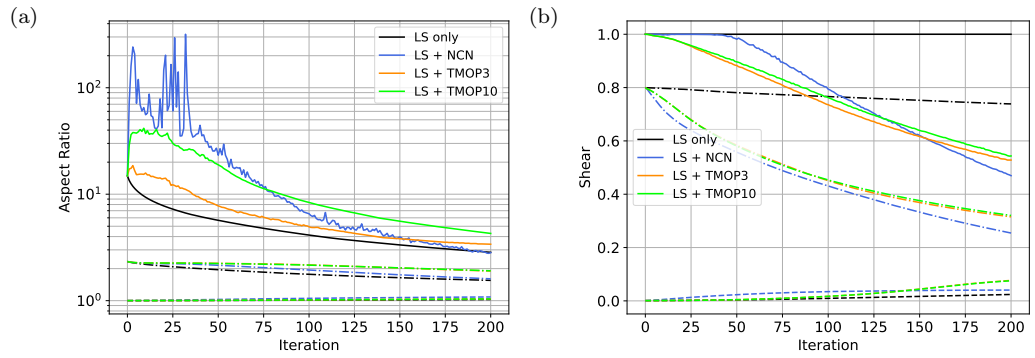

Figure 9: Statistical mesh metrics as a function of rezoning iteration number for the 3D swirl test, showing (a) aspect-ratio r as defined from the node star in Sec. 2.4.1, and (b) shear, defined as the maximum $|\cos\theta|$ angle in the node star tetrahedra. Each panel shows the median value (dashed lines), the maximum value (solid lines) and the 95-th percentile value (dash-dotted line). Data is overlaid from the linesweep only case (black), linesweep with normed-CN (blue), linesweep with TMOP with $r_{\max} = 3$ (orange) and $r_{\max} = 10$ (green).

3.3. Single-node inversion recovery

To isolate disentangler convergence, we start from a valid Cartesian mesh in $\Omega = [0, 1]^3$ and select an interior node p with position $\mathbf{x}_p$ at the mesh center $[0.5, 0.5, 0.5]$. We displace that node outside of the initial domain along $+\mathbf{e}_z$,

$$\mathbf{x}_p \leftarrow \mathbf{x}_p + 10\,\mathbf{e}_z,$$

which creates a small, localized set of inverted tetrahedra in the nodal stars of p and of its immediate neighbor along $+\mathbf{e}_z$ (typically eight inverted tetrahedra for the standard hexahedral corner-based decompositions considered here). We then run (i) the linesweep alone and (ii) the disentangler alone (convexity mode), and report the number of iterations required to restore validity,

$$\min_{\tau \in \mathcal{S}_p} \alpha(\tau) > 0,$$

and to reach the target convexity threshold (i.e. $\alpha(\tau) \geq \beta(p)$ for all $\tau \in \mathcal{Z}_p$). This test is intentionally localized and provides a direct measure of disentangler convergence speed in terms of iteration numbers. In this test, both the linesweep and the convexity metric use the same disentangler/smoothing coupling as described in Sec. 2.6.3 (the linesweep disentangler combines linesweep with normed-CN neighboring relaxation while the convexity disentangler combines the convexity metric with TMOP neighborhing relaxation).

In this particular test, we find that the linesweep+normed-CN approach takes 133 iterations in equal-space mode to recover a valid mesh, while the patch-based convexity metric+TMOP converges in 40 iterations.

3.4. Randomly perturbed nonconformal AMR cube

To assess robustness in a larger setting and in presence of nonconformal interfaces, we initialize a 2:1 balanced AMR mesh on the unit domain using the macro-zone 0-1 refinement pattern described above. We apply a random perturbation to interior (non-boundary) nodes. For each such node p we compute a local length scale δ_p as the minimum distance between $\mathbf{x}_p$ and any other vertex sharing a face with p. We then draw an isotropic unit direction $\mathbf{u}_p$ (uniform on a sphere) and set

$$\mathbf{x}_p \leftarrow \mathbf{x}_p + \sigma\, \delta_p\, \mathbf{u}_p$$

where σ is the prescribed perturbation amplitude. This perturbation produces localized, scale-consistent distortions of the mesh. By setting σ to a large value (here 15), a large portion of the mesh nodes are initially invalid. The initial mesh contains 45200 nodes, and the initial number of invalid mesh nodes is 39006. We start the test by disentangling the mesh and follow it by mesh optimization.

We find that the convexity-based approach achieves full mesh disentanglement in 133 iterations, while the linesweep-based approach requires 280 iterations. The resulting disentangled meshes are shown in Fig. 10(b,d); both retain significant deformations at this stage. Meshes after 10 and 50 subsequent optimization steps are shown in Fig. 10(c) and (e) for the linesweep- and convexity-based approaches, respectively. The linesweep+normed-CN approach converges more rapidly toward a regular mesh in this case, as the TMOP-based optimization attempts to preserve the aspect ratios induced by the initial random perturbation of the mesh.

4. Conclusions

We have presented a rezoning method for three-dimensional ALE and ALE-AMR meshes that combines line-based smoothing with a patch-based, aspect-ratio-preserving TMOP formulation. The method targets a difficulty specific to ALE rezoning: the mesh should not be relaxed indiscriminately toward an isotropic target, because local anisotropy may contain useful Lagrangian information. At the same time, the mesh must remain valid and avoid excessive shear, near-degenerate elements, and cell bunching. These requirements are harder to satisfy in ALE-AMR, where nonconformal nodes and 2:1 level transitions introduce geometric features that should not automatically be treated as defects.

The method separates these roles between two components. The linesweep acts on extended mesh-line stencils and provides the long-range smoothing needed to regulate spacing and avoid irregular local minima. The TMOP component acts locally, on corner-based extended tetrahedral stencils, and controls shear and skewness relative to a dynamic anisotropic target. This target is constructed from the local Lagrangian edge geometry and normalized so that it preserves aspect-ratio information without directly penalizing local volume differences. The optimization can therefore reduce local distortion without removing anisotropy inherited from the Lagrangian phase or from AMR level transitions. The aspect-ratio cap introduced in the target gives an explicit control on the amount of relaxation: it limits near-degenerate sliver elements without forcing the target to become isotropic.

The update procedure is also part of the method. The TMOP and convexity objectives are evaluated on an extended patch-1 stencil, so that the optimization step accounts for the elements incident on the node being moved and for neighboring elements affected by that motion. In valid configurations, the TMOP objective provides a barrier-like response as the mesh approaches loss of validity. When invalid configurations are detected, the algorithm switches to a convexity objective that restores positive signed volumes before returning to TMOP quality optimization. The nodal updates are applied in a Gauss-Seidel-like manner, with each accepted move based on the current mesh rather than on a frozen bulk state. Parallelism is obtained through distance-1 coloring of the nodal graph, which allows non-interacting nodes to be updated concurrently while remaining compatible with hanging-node reconstruction across AMR interfaces.

The numerical tests illustrate the roles of the different components. In the 2D bunching-layer test, pure linesweep and linesweep combined with a normalized condition-number metric did not efficiently relax strongly bunched high-aspect-ratio cells. Pure TMOP reduced some aspect-ratio and shear measures, but did not recover a smooth mesh by itself. The combination of linesweep and TMOP gave better overall

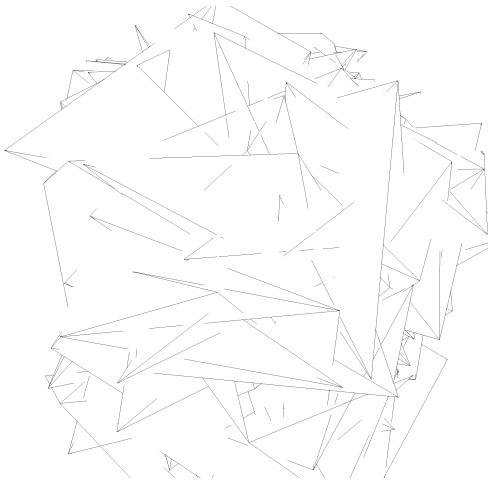

(a) Initial mesh

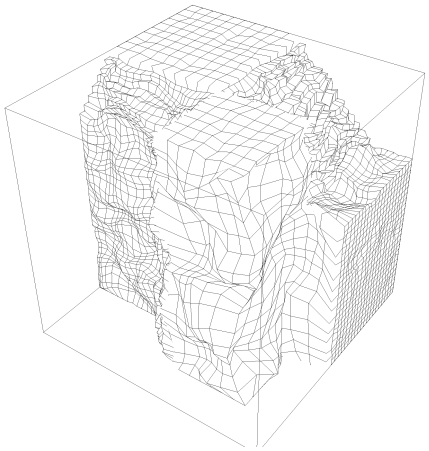

(b) Disentangled mesh with linesweep-based approach (280 iterations)

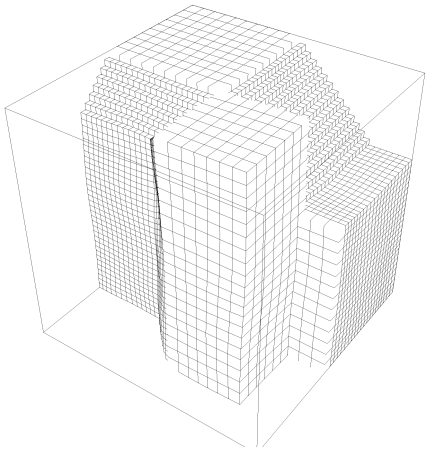

(c) Disentangled mesh regularized by 10 iterations with linesweep-based approach

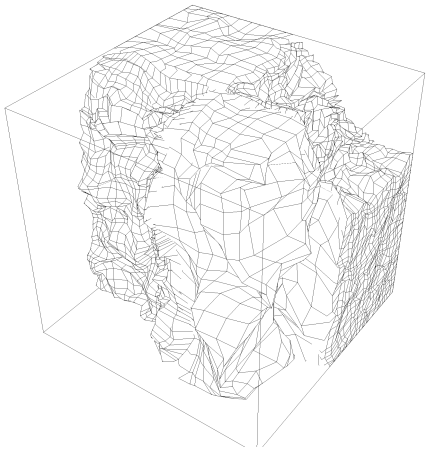

(d) Disentangled mesh with convexity-based approach (133 iterations)

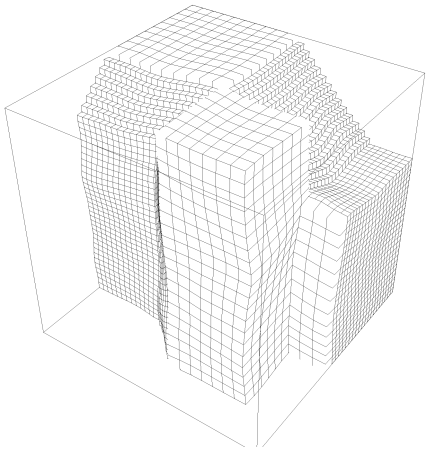

(e) Disentangled mesh regularized by 50 iterations with convexity-based approach

Figure 10: Meshes sections in the 3D randomly perturbed nonconformal cube case.

mesh quality. Changing the TMOP aspect-ratio cap produced a faster decrease of the high-percentile aspect-ratio tail, showing that the cap acts as an effective control parameter. In the 3D swirl test, line-based smoothing alone did not remove the genuinely three-dimensional shear, whereas both linesweep+TMOP and linesweep+normed-CN reduced shear. The important difference was robustness during the optimization: the bulk normed-CN approach produced transient invalid meshes, while the patch-based TMOP update remained valid throughout the test. The inversion-recovery and randomly perturbed AMR tests further showed that the convexity-TMOP strategy can recover from localized and distributed invalid configurations, including in the presence of hanging nodes and block junctions.

These results indicate that the proposed method provides a practical framework for mesh-quality control in ALE and ALE-AMR calculations. TMOP does not replace linesweep smoothing; rather, the two components address different defects. Linesweep smoothing controls nonlocal spacing and mesh-line regularity, while target-based optimization controls local shear, slivers, and excessive aspect ratios without discarding Lagrangian anisotropy. The method is modular, compatible with nonconformal AMR topology, and provides a direct way to tune mesh relaxation through the prescribed aspect-ratio cap.

The present study was limited to mesh-only tests, chosen to isolate the behavior of the rezoner. A next step is to evaluate the method in fully coupled ALE-AMR hydrodynamic calculations, including conservative remapping, material-interface motion, and repeated rezoning over many time steps. Such tests will also make it possible to quantify the effect of the aspect-ratio cap on timestep control, remap accuracy, and preservation of physically relevant Lagrangian structures. In addition, although the distance-1 coloring provides a parallel update structure, dedicated strong- and weak-scaling studies will be needed to assess the cost of the patch-based optimization in large production calculations.

5. Acknowledgment

This work is supported by the U.S. Department of Energy (DOE), Office of Science, Fusion Energy Sciences, under Award No. DE-SC0024863: IFE-COLoR and AWD-00007026: INFUSE.

References

[1] Cyril W. Hirt, Anthony A. Amsden, and James L. Cook. An arbitrary lagrangian-eulerian computing method for all flow speeds. *Journal of Computational Physics*, 14(3):227–253, 1974.

[2] David J. Benson. Computational methods in Lagrangian and Eulerian hydrocodes. *Comput. Methods Appl. Mech. Eng.*, 99:235–394, 1992.

[3] Andrew J. Barlow, Mikhail J. Shashkov, Pierre-Henri Maire, Robert N. Rieben, and William J. Rider. Arbitrary lagrangian–eulerian methods for modeling high-speed compressible multimaterial flows. *Journal of Computational Physics*, 322:603–665, 2016.

[4] Patrick M. Knupp, Len G. Margolin, and Mikhail J. Shashkov. Reference Jacobian optimization-based rezone strategies for Arbitrary Lagrangian-Eulerian methods. *J. Comput. Phys.*, 176:93–128, 2002.

[5] Markus Berndt, Milan Kucharik, and Mikhail J. Shashkov. Using the feasible set method for rezoning in ALE. *Procedia Computer Science*, 1(1):1885–1892, 2010. ICCS 2010.

[6] Pavel Váchal, Rao V. Garimella, and Mikhail J. Shashkov. Untangling of 2d meshes in ALE simulations. *Journal of Computational Physics*, 196(2):627–644, 2004.

[7] Lori A. Freitag and Paul Plassmann. Local optimization-based simplicial mesh untangling and improvement. *International Journal for Numerical Methods in Engineering*, 49(1–2):109–125, 2000.

[8] Alan M Winslow. Numerical solution of the quasilinear poisson equation in a nonuniform triangle mesh. *Journal of Computational Physics*, 1(2):149–172, 1966.

[9] Hyung Taek Ahn, Larisa Branets, and Graham F. Carey. Moving boundary simulations with dynamic mesh smoothing. *International Journal for Numerical Methods in Fluids*, 64(8):887–907, 2010.

[10] Jin Yao and Douglas Stillman. An equal-space algorithm for block-mesh improvement. *Procedia Engineering*, 163:199–211, 2016. 25th International Meshing Roundtable.

[11] Pavel Váchal and Pierre-Henri Maire. Discretizations for weighted condition number smoothing on general unstructured meshes. *Computers and Fluids*, 46(1):479–485, 2011. 10th ICFD Conference Series on Numerical Methods for Fluid Dynamics (ICFD 2010).

[12] Patrick M. Knupp. Achieving finite element mesh quality via optimization of the jacobian matrix norm and associated quantities. part ii—a framework for volume mesh optimization and the condition number of the jacobian matrix. *Int. J. Numer. Methods Eng.*, 48(8):1165–1185, 2000.

[13] Patrick M. Knupp. Algebraic mesh quality metrics. *SIAM Journal on Scientific Computing*, 23(1):193–218, 2001.

[14] Patrick M. Knupp. A matrix norm approach to mesh quality metrics. *SIAM Journal on Scientific Computing*, 24(3):1071–1092, 2003.

[15] Jibum Kim, Thap Panitanarak, and Suzanne M. Shontz. A multiobjective mesh optimization framework for mesh quality improvement and mesh untangling. *International Journal for Numerical Methods in Engineering*, 94(1):20–42, 2013.

[16] Patrick M. Knupp. Introducing the target-matrix paradigm for mesh optimization via node-movement. *Engineering with Computers*, 28(4):419–429, 2012.

[17] Veselin A. Dobrev, Patrick Knupp, Tzanio V. Kolev, Ketan Mittal, and Vladimir Z. Tomov. The target-matrix optimization paradigm for high-order meshes. *SIAM Journal on Scientific Computing*, 41(1):B50–B68, 2019.

[18] R.W. Anderson, N.S. Elliott, and R.B. Pember. An arbitrary lagrangian–eulerian method with adaptive mesh refinement for the solution of the euler equations. *J. Comp. Phys.*, 199(2):598–617, 2004.

[19] Jacob I. Waltz and Jozsef Bakosi. A coupled ALE–AMR method for shock hydrodynamics. *Comput. & Fluids*, 167:359–371, 2018.

[20] Arnaud Colaïtis, Sébastien Guisset, and Jérôme Breil. A cell-centered amr-ale framework for 3d multi-material hydrodynamics. part i: Lagrangian and indirect euler amr algorithms. *Journal of Computational Physics*, 552:114701, 2026.

[21] Arnaud Colaïtis, Sébastien Guisset, and Jérôme Breil. A cell-centered amr-ale framework for 3d multi-material hydrodynamics. part ii: linesweep ale rezoning for nonconformal block-structured amr meshes. *Journal of Computational Physics*, 552:114702, 2026.

[22] Jérôme Breil, Guillaume Damour, Sébastien Guisset, and Arnaud Colaïtis. 3d mesh regularization within an ALE code using a weighted line sweeping method. *Computers & Fluids*, 292:106591, 2025.

[23] Mark T. Jones and Paul E. Plassmann. A parallel graph coloring heuristic. *SIAM Journal on Scientific Computing*, 14(3):654–669, 1993.

[24] Michael Luby. A simple parallel algorithm for the maximal independent set problem. *SIAM Journal on Computing*, 15(4):1036–1055, 1986.

[25] Carsten Burstedde, Lucas C. Wilcox, and Omar Ghattas. `p4est`: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees. *SIAM Journal on Scientific Computing*, 33(3):1103–1133, 2011.

[26] David A. Field. Qualitative measures for initial meshes. *International Journal for Numerical Methods in Engineering*, 47(4):887–906, 2000.

[27] Patrick M. Knupp. Hexahedral mesh untangling & algebraic mesh quality metrics. In Steven J. Owen, editor, *Proceedings of the 9th International Meshing Roundtable, IMR 2000, New Orleans, Louisiana, USA, October 2-5, 2000*, pages 173–183, 2000.

[28] Larry Armijo. Minimization of functions having lipschitz continuous first partial derivatives. *Pacific Journal of Mathematics*, 16(1):1–3, 1966.

[29] Sebastiano Vigna. Splitmix64 reference implementation, 2015. Public-domain reference implementation, see prng.di.unimi.it/splitmix64.c.