LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-JRNL-655765

A Subzone Reconstruction Algorithm for Efficient Staggered Compatible Remapping

D. P. Starinshak, J. M. Owen

June 16, 2014

Journal of Computational Physics

Disclaimer

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

A Subzone Reconstruction Algorithm for Efficient Staggered Compatible Remapping

D. P. Starinshak[a], J. Michael Owen[b]

[a]*Lawrence Livermore National Laboratory, AX Division, M/S L-38, P.O. Box 808, Livermore, CA 94550*
[b]*Lawrence Livermore National Laboratory, AX Division, M/S L-38, P.O. Box 808, Livermore, CA 94550*

Abstract

Staggered-grid Lagrangian hydrodynamics algorithms frequently make use of subzonal discretization of state variables for the purposes of improved numerical accuracy, generality to unstructured meshes, and exact conservation of mass, momentum, and energy. Coupling hydrodynamics to additional physics, such as heat conduction and radiation transport, further benefits from subzonal discretizations. For Arbitrary Lagrangian-Eulerian (ALE) methods using a geometric overlay, it is difficult to remap subzonal variables in an accurate and efficient manner due to the number of subzone-subzone intersections that must be computed. This becomes prohibitive in the case of 3D, unstructured, polyhedral meshes. A new procedure is outlined in this paper to avoid direct subzonal remapping. The new algorithm reconstructs the spatial profile of a subzonal variable using remapped zonal and nodal representations of the data. The reconstruction procedure is cast as an under-constrained optimization problem. Enforcing conservation at each zone and node on the remapped mesh provides the set of equality constraints; the objective function corresponds to a quadratic variation per subzone between the values to be reconstructed and a set of target reference values. Numerical results for various pure-remapping and hydrodynamics tests are provided. Ideas for extending the algorithm to staggered-grid radiation-hydrodynamics are discussed as well as ideas for generalizing the algorithm to include inequality constraints.

Email addresses: `starinshak1@llnl.gov` (D. P. Starinshak), `mikeowen@llnl.gov` (J. Michael Owen)

Keywords: Remapping, ALE methods

1. Background

In the following sections, we introduce the problem of remapping subzonal quantities in the context of staggered-mesh discretizations. We define the subzonal spatial discretization of interest to our method, highlight specific algorithms which employ this discretization, and describe the constraints and accuracy requirements of these algorithms.

Subzonal Discretization

Our main application in this paper is the use of a geometric overlay based remap with the spatially-staggered, Lagrangian-frame algorithm for hydrodynamics outlined in [2]. The Lagrangian hydrodynamic discretization of [2] derives a number of important properties–for instance, conservation, generality to unstructured grids, and improved accuracy–from the discretization of state variables on subzonal mesh elements. The main result and examples in this paper are the efficient reconstruction of such subzonal properties (in particular the mass) based on remapping the aggregate zonal and nodal properties. This approach is applicable so long as the primary physics discretization takes place on the zonal and/or nodal volumes, as is the case for the staggered Lagrangian hydrodynamic discretization we study here. These ideas can be generalized to other types of discretizations that employ subzonal quantities as well, such as flux-limited diffusion (FLD) radiation-hydrodynamics calculations on unstructured grids [20, 21, 22].

First, some notation concerning mesh subzones. Each zone is subdivided into subzonal elements, which we often refer to as *corners* and index by c. Each corner is a polyhedron in 3D or a quadrilateral in 2D, corresponding to a unique neighboring zone-node pair (z_c, n_c) on the Lagrangian mesh. The vertices of corner c consist of the centroid of zone z_c, the position of node n_c, and the centers of each of the faces of z_c that touch n_c. Figure 1 illustrates the corner discretization in 2D and 3D. (Similarly, a corner can be thought of as the intersection of zone element z_c on the primary mesh and node element n_c on the corresponding dual mesh.) Element volumes around zones V_z and nodes V_n are defined by sums of neighboring corner volumes V_c:

$$V_z = \sum_{c(z)} V_c \qquad \text{and} \qquad V_n = \sum_{c(n)} V_c. \tag{1}$$

Here $c(z)$ denotes the set of corners inside zone z, and $c(n)$ denotes the set of corner touching node n. The union of each corner element in $c(n)$ defines the mesh element associated with node n.

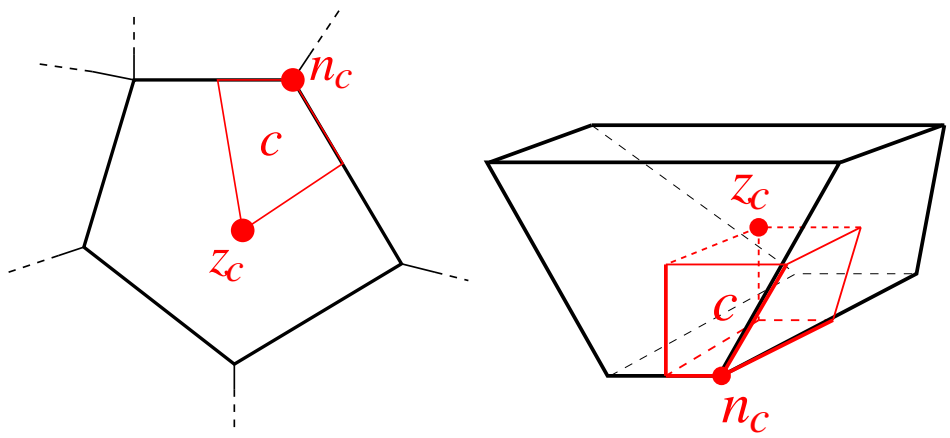

Figure 1: Corner subvolumes (c, in red) of a 2D (left) and 3D (right) zone.

Subzonal Densities and Compatible Hydrodynamics

We consider subzonal discretizations in the context of conservative hydrodynamics. We assume spatially-staggered variables: kinematic variables of velocity and kinetic energy at nodes; thermodynamic variables of internal energy at zone centers.

The compatible formalism for Lagrangian hydrodynamics is based on a consistent discretization of the fluid momentum and energy evolution equations [2]. Consistency is achieved through a subzonal discretization of the underlying fluid mass density. The zonal and nodal representations of density that define discrete energy and momenta respectively are coarsened representations of the subzonal density on their respective element volumes:

$$m_z = \sum_{c(z)} m_c \qquad \text{and} \qquad m_n = \sum_{c(n)} m_c, \tag{2}$$

where m_z, m_n, and m_c denote discrete masses centered at zone z, node n, and corner c respectively. The corresponding densities satisfy

$$\rho_\alpha = m_\alpha / V_\alpha, \quad \alpha = z, n, c. \tag{3}$$

Problem Statement

The existence of discrete state variables at zones, nodes, and subzones presents us with a considerable challenge in the context of Arbitrary Lagrangian-Eulerian (ALE) methods: in order to remap state variables from one mesh to another in a conservative fashion, accurate subzonal information must be maintained. The problem may be interpreted in terms of both flux-based and overlay-based remaps. Owen and Shashkov describe a flux-based strategy in [14] that computes consistent subzonal fluxes around nodes to preserve mass and momentum conservation. Subzonal fluxes are computed in a manner that reduces the local variation with respect to an optimal reference flux. However, the flux based algorithm of [14] implicitly assumes the topology of the mesh is not changing, only the geometry. For a geometric overlay-based remap we would like to relax this assumption and allow the target and donor mesh topologies to be unrelated. This paper builds on the ideas in [14] in the context of generic overlay-based remapping.

In an overlay remap, state variables are mapped from one spatial partition (donor mesh) to a second partition (target mesh) based on geometrical intersections between donor and target elements. For cell-centered discretizations, all state variables are collocated to zone centers; only zone-zone intersections need to be computed in the overlay. The computational complexity increases for staggered discretizations due to the fact that variables are centered at both zones and nodes. Zone-zone intersections are computed to remap quantities such as internal energy; node-node (i.e. dual mesh) intersections are computed to remap quantities such as momentum and kinetic energy.

Intersection operations can impose steep memory requirements and become computationally expensive, particularly if (i) mesh elements increase in complexity, as with unstructured, polygonal meshes (for instance, ReALE [10]); (ii) additional physics is coupled into the problem, thus increasing the number remapped variables; or (iii) the problem increases in dimensionality.

When state variables are discretized at mesh subzones, their information is lost during a geometric overlay if mesh intersections are limited to zones and nodes only. For compatible hydrodynamics, this means losing the very conservation property provided by the discretization in the first place. For staggered radiation-hydrodynamics, this means losing intrazonal temperature variation and reducing the accuracy of fluid-radiation energy coupling terms.

We recognize two general strategies for maintaining subzonal state information: (i) remap subzonal information directly from donor mesh to target

mesh or (ii) attempt to reconstruct subzonal data using information at neighboring zones and nodes. This paper describes a new algorithm for strategy (ii). We first summarize the state-of-the-art for strategy (i).

An accurate and efficient implementation of strategy (i) is described in [11] and successfully applied to staggered ReALE hydrodynamics calculations in [10]. The algorithm is efficient in the sense that geometric intersections are *only* computed between subzones. All variables are collocated to neighboring subzones to accomplish this. The collocation process, referred to as "gathering," computes subzonal values for all zonal and nodal quantities. After the remap, the gather stage is reversed in a process referred to as "scattering," and consistent zonal and nodal discretizations are recovered from their remapped state on subzones. The gather-remap-scatter process is designed to be conservative, reversible, and second-order accurate. The authors employ a local repair procedure based on [12] to ensure the remap is also monotonic.

Remapping subzones to subzones is an intuitive strategy, but it unfortunately does not solve the problem of computational complexity. The number of element intersections scales with the total number of elements in both donor and target meshes. For the simple case of quad zones in 2D, the number of subzone-subzone intersections is still twice the sum of both zone-zone and node-node intersections in pure element count. For arbitrary polyhedral zones in 3D, the number of subzones becomes much greater than the sum of zones and nodes together. Fortunately, the fact that the hydrodynamic equations are differenced on the nodal and zonal volumes suggests that by carefully remapping the nodal and zonal quantities and conservatively reconstructing the subzonal properties from this information should be adequate to maintain the overall accuracy of the hydrodynamic method.

We propose just such an alternative strategy here. Described in detail in the following section, this strategy involves reconstructing subzonal information based solely on zonal and nodal information. Our approach takes the form of a constrained optimization problem in which discrete conservation conditions serve as linear constraints on the subzonal field of interest. The reconstruction method generalizes to any subzonal quantity that is conserved by volume.

2. Reconstruction Algorithm

This section describes the algorithm for reconstructing subzonal information from discrete data at zones and nodes. In the case of staggered

compatible hydrodynamics, the quantity we reconstruct is subzonal density using remapped masses at zones and nodes. In general, the algorithm applies to any subzonal quantity that is conserved by volume.

The initial conditions of the algorithm are the post-remap zonal masses $\{m_z\}_{z=1}^{N_z}$, nodal masses $\{m_n\}_{n=1}^{N_n}$, as well as the subzonal volumes $\{V_c\}_{c=1}^{N_c}$ of the post-remap mesh. Subzonal mass densities ρ_c are unknown but assumed to be related to their local nodal and zonal masses.

We wish to solve for the set of corner densities satisfying the following conditions:

(i) The sum of sub-zone corner masses recovers the zonal mass.

(ii) The sum of nodal corner masses recovers the nodal mass.

(iii) Each corner density is as close as possible to a convenient reference state.

We interpret the statements above as a constrained optimization problem. Statements (i) and (ii) are linear constraints on the set of possible corner density values. Statement (iii) defines the optimality of the desired corner density field.

Using (1), (2), and (3), constraints (i) and (ii) are written in terms of local corner information:

$$(i): \quad m_z = \sum_{c(z)} m_c = \sum_{c(z)} V_c \rho_c = \rho_z \sum_{c(z)} V_c \quad \Longrightarrow \quad \sum_{c(z)} V_c(\rho_c - \rho_z) = 0. \tag{4}$$

and

$$(ii): \quad m_n = \sum_{c(n)} m_c = \sum_{c(n)} V_c \rho_c = \rho_n \sum_{c(n)} V_c \quad \Longrightarrow \quad \sum_{c(n)} V_c(\rho_c - \rho_n) = 0. \tag{5}$$

For statement (iii), we define a positive *relative variation* per corner element $(\rho_c/\mu_c - 1)^2$, where μ_c denotes some optimal reference state. (We will return to possible forms of μ_c shortly.) Summing over corners defines a global optimization functional

$$(iii): \quad \frac{1}{2} \sum_c \Big(\frac{\rho_c}{\mu_c} - 1 \Big)^2. \tag{6}$$

In summary, we cast (i)–(iii) as the following constrained optimization problem:

$$\begin{array}{ll} \text{Given:} & \rho_z,\ \rho_n,\ V_c \\ \text{Minimize:} & \frac{1}{2}\sum_c \Big(\frac{\rho_c-\mu_c}{\mu_c}\Big)^2 \\ \text{Subject to:} & \begin{cases} \sum_{c(z)} V_c(\rho_c-\rho_z)=0, & \forall z \\ \sum_{c(n)} V_c(\rho_c-\rho_n)=0, & \forall n \end{cases} \end{array} \tag{7}$$

We solve (7) by the technique of Lagrange multipliers. Define multipliers λ_z and λ_n corresponding to each zone constraint (4) and node constraint (5) respectively and define the global Lagrangian functional

$$L(\rho_c,\lambda_z,\lambda_n)=\frac{1}{2}\sum_c\Big(\frac{\rho_c-\mu_c}{\mu_c}\Big)^2+\sum_z\lambda_z\sum_{c(z)}V_c(\rho_c-\rho_z)+\sum_n\lambda_n\sum_{c(n)}V_c(\rho_c-\rho_n). \tag{8}$$

The gradient of L is

$$\begin{aligned} \frac{\partial L}{\partial \rho_c} &= \frac{\rho_c-\mu_c}{\mu_c^2}+V_c\lambda_{z_c}+V_c\lambda_{n_c} \\ \frac{\partial L}{\partial \lambda_z} &= \textstyle\sum_{c(z)} V_c(\rho_c-\rho_z) \\ \frac{\partial L}{\partial \lambda_n} &= \textstyle\sum_{c(n)} V_c(\rho_c-\rho_n) \end{aligned} \tag{9}$$

where, again, z_c is the unique zone containing corner c, and n_c is the unique node touching corner c. Setting $\partial L/\partial\rho_c=0$ gives an expression for ρ_c in terms of a perturbation on its reference value:

$$\rho_c=\mu_c\Big(1-\mu_cV_c(\lambda_{z_c}+\lambda_{n_c})\Big). \tag{10}$$

Setting $\partial L/\partial\lambda_n=0$ and substituting (10) gives an expression for λ_n in terms of values of λ_z at zones neighboring node n:

$$\lambda_n \;=\; -\frac{\sum_{c(n)}(\mu_cV_c)^2\lambda_{z_c}+\sum_{c(n)}(\rho_n-\mu_c)V_c}{\sum_{c(n)}(\mu_cV_c)^2} \;=\; -\frac{\sum_{c(n)}w_c\lambda_{z_c}+\Delta m_n}{W_n}. \tag{11}$$

For convenience, we define corner weights based on the square of the reference mass

$$w_c = (\mu_c V_c)^2 \tag{12}$$

as well as zone and node weights

$$W_z = \sum_{c(z)} w_c \qquad \text{and} \qquad W_n = \sum_{c(n)} w_c. \tag{13}$$

We also define zonal and nodal mass differences relative to the reference state:

$$\begin{aligned} \Delta m_z &= m_z - \textstyle\sum_{c(z)} \mu_c V_c = \sum_{c(z)} (\rho_z - \mu_c) V_c \\ \Delta m_n &= m_n - \textstyle\sum_{c(n)} \mu_c V_c = \sum_{c(n)} (\rho_n - \mu_c) V_c \end{aligned} \tag{14}$$

Finally, setting $\partial L / \partial \lambda_z = 0$ and substituting (10) and (11), gives a linear expression for each zonal Lagrange multiplier λ_z:

$$W_z \lambda_z \; - \; \sum_{n(z)} \frac{w_{c(z,n)}}{W_n} \sum_{c(n)} w_c \lambda_{n_c} \; = \; \sum_{n(z)} \frac{w_{c(z,n)}}{W_n} \Delta m_n \; - \; \Delta m_z, \tag{15}$$

where the notation $c(z, n)$ denotes the unique corner inside zone z and touching node n.

The unknowns in (15) are the N_z zonal Lagrange multipliers λ_z. We write (15) as a linear system of the form

$$(\mathbf{W}_z + \mathbf{A})\lambda = \mathbf{b}.$$

Here λ is the vector of Lagrange multipliers $\{\lambda_z\}_{n=1}^{N_z}$, $\mathbf{b}$ is a length N_z vector of known quantities, $\mathbf{W}_z$ is an $N_z \times N_z$ diagonal matrix consisting of zone weights (13), and $\mathbf{A}$ is a sparse $N_z \times N_z$ matrix. The i-th entry of $\mathbf{b}$ corresponds to zone z_i:

$$b_{z_i} \; = \; \sum_{n(z_i)} \frac{w_{c(z_i,n)}}{W_n} \Delta m_n \; - \; \Delta m_{z_i}. \tag{16}$$

The (i, j)-th entry of $\mathbf{A}$ corresponds to zone pair (z_i, z_j):

$$A_{i,j} \; = \; - \sum_{n(z_i,z_j)} \frac{w_{c(z_i,n)} w_{c(z_j,n)}}{W_n}, \tag{17}$$

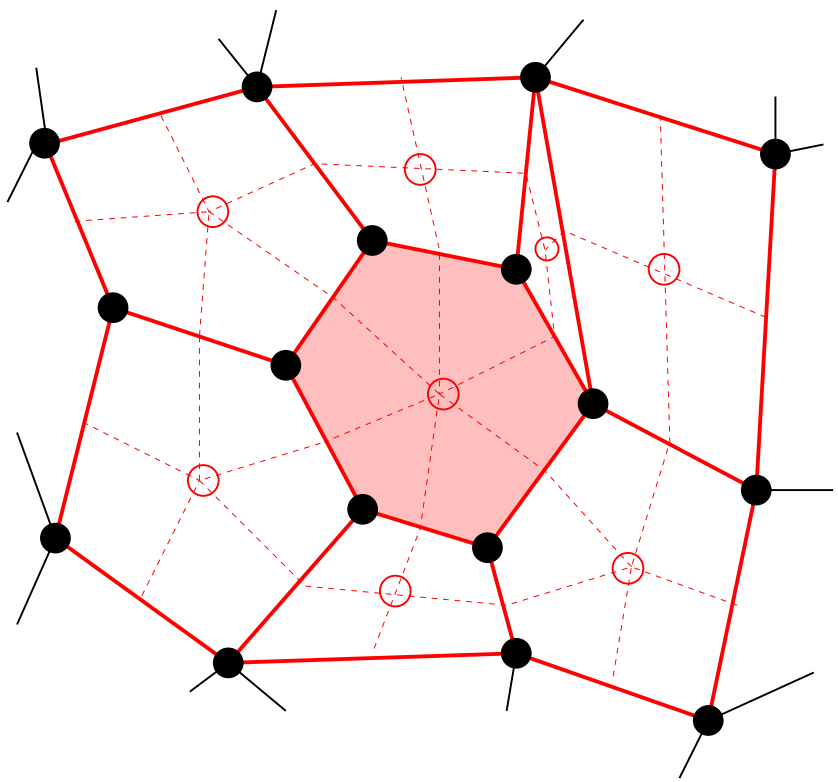

Figure 2: An example of the stencil of zones that share at least one node with with the shaded 2D polygonal zone.

where $n(z_i, z_j)$ is the set of nodes shared by both zones z_i and z_j. Matrix $\mathbf{A}$ is sparse in the sense that, if z_i and z_j do not share any nodes, $A_{i,j} = 0$. Alternatively, one can interpret the individual rows of matrix $\mathbf{A}$ using the topology of the mesh: the nonzero entries for row i correspond to the mesh stencil around z_i of all zones which share at least one node with z_i (see fig. 2).

It is instructive to write matrices $\mathbf{W}_z$ and $\mathbf{A}$ as products of geometrical sub-matrices based on the mesh connectivity. The sub-matrices are weighted

by corner volume for notational convenience. Define

$$\mathbf{\Psi}_{cz} = \begin{cases} V_c & , \text{ corner } c \text{ is contained in zone } z \\ 0 & , \text{ o.w.} \end{cases} \tag{18}$$

$$\mathbf{\Psi}_{cn} = \begin{cases} V_c & , \text{ corner } c \text{ touches node } n \\ 0 & , \text{ o.w.} \end{cases} \tag{19}$$

where $\mathbf{\Psi}_{cz}$ is size $N_c \times N_z$, and $\mathbf{\Psi}_{cn}$ is size $N_c \times N_n$. Similarly, let $\mathbf{W}_z = \text{diag}(\{W_z\}_z)$, $\mathbf{W}_n = \text{diag}(\{W_n\}_n)$, and $\mathbf{M}_c = \text{diag}(\{\mu_c\}_c)$ denote diagonal matrices of zone weights, node weights, and reference values respectively. Then,

$$\mathbf{W}_z = \mathbf{\Psi}_{cz}^T \mathbf{M}_c^2 \mathbf{\Psi}_{cz} \qquad \text{and} \qquad \mathbf{W}_n = \mathbf{\Psi}_{cn}^T \mathbf{M}_c^2 \mathbf{\Psi}_{cn} \tag{20}$$

by construction and similarly

$$\mathbf{A} \;=\; -\mathbf{\Psi}_{cz}^T \mathbf{M}_c^2 \mathbf{\Psi}_{cn} \mathbf{W}_n^{-1} \mathbf{\Psi}_{cn}^T \mathbf{M}_c^2 \mathbf{\Psi}_{cz}. \tag{21}$$

Finally, define

$$\mathbf{\Psi} \;=\; \mathbf{W}_z + \mathbf{A} \;=\; \mathbf{\Psi}_{cn}^T \mathbf{M}_c \Big(\mathbf{I}_c - \mathbf{M}_c \mathbf{\Psi}_{cz} \mathbf{W}_z^{-1} \mathbf{\Psi}_{cz}^T \mathbf{M}_c \Big) \mathbf{M}_c \mathbf{\Psi}_{cn}. \tag{22}$$

The matrix form of (15) is given by the system

$$\mathbf{\Psi} \lambda = \mathbf{b}. \tag{23}$$

The wrinkle involved with solving linear system (23) is that $\mathbf{\Psi}$ is a singular matrix. In general, the constrained minimization problem (7) is underdetermined: because $N_c \geq N_z + N_n$ in general, the combined constraints (4) and (5) are fewer than the total number of unknowns. Consequently, a unique set of constrained corner densities which minimized (6) does not exist. Singular value decomposition (SVD) of $\mathbf{\Psi}$ provides one well-known method of finding a best-fit solution to (23).

We begin by decomposing $\mathbf{\Psi}$ into components

$$\mathbf{\Psi} = \mathbf{U} \mathbf{w} \mathbf{V}^T, \tag{24}$$

such that $\mathbf{w}$ is a diagonal matrix containing the N_z singular values of $\mathbf{\Psi}$, and $\mathbf{U}$ and $\mathbf{V}$ are square unitary matrices. The pseudoinverse of $\mathbf{\Psi}$ is

$$\mathbf{\Psi}^+ = \mathbf{V} \mathbf{w}' \mathbf{U}^T, \tag{25}$$

where $\mathbf{w}'$ is the diagonal matrix $\mathbf{w}$ with all nonzero elements replaced by their reciprocal. The solution to (23) is then

$$\lambda = \mathbf{\Psi}^{+}\mathbf{b}. \tag{26}$$

Nodal Lagrange multipliers can then be computed from $\lambda = \{\lambda_z\}_{z=1}^{N_z}$ using (11), with final corner densities computed from (10).

Remark: Solving (23) does not necessarily require computing an SVD. In particular, the zero eigenvalue of $\mathbf{\Psi}$ is of multiplicity one (see appendix A). A preconditioned iterative solver will converge to the minimizing solution given an initial guess outside the null space of $\mathbf{\Psi}$. This is the solution approach utilized throughout this paper. We use the parallel preconditioner package Hypre to solve the linear system [6].

Remark: Penalty methods provide an alternative formulation to that of Lagrange multipliers for the purpose of constrained optimization. Additional functions may be incorporated into the full variation functional (6) to penalize unsatisfactory states, such as negative or non-monotonic states. Unfortunately, penalty methods do not satisfy equality constraints exactly, only in an iterative limit. This is undesirable for our purposes given our strict consistency conditions between zonal and nodal representations of mass to satisfy conservation.

Remark: Incorporating inequality constraints into the method of Lagrange multipliers is possible through the formalism of Karush-Kuhn-Tucker and complementary slackness [7, 8]. Inequality constraints may be used to enforce positivity in the resulting corner density solution or, stronger still, monotonicity with respect to neighboring zonal and nodal density values. See appendix C for a thorough discussion of these ideas as well as generalization of the above algorithm to incorporate inequality constraints.

3. Reference State

The reference state in (6) defines the optimal state for a given corner density. Intuitively, the reference state should depend on the density field local to each corner. For corner c, we make the following simplifying assumptions:

- μ_c is a function of local zone density ρ_{z_c} and local node density ρ_{n_c}.

- μ_c is bounded by the two local densities.

We consider six choices for μ_c in this article:

$$
\begin{aligned}
&\text{(i) Zone Density:} && \rho_{z_c} \\
&\text{(ii) Node Density:} && \rho_{n_c} \\
&\text{(iii) Arithmetic Average:} && \frac{\rho_{z_c} + \rho_{n_c}}{2} \\
&\text{(iv) Harmonic Average:} && \frac{2\rho_{z_c}\rho_{n_c}}{\rho_{z_c} + \rho_{n_c}} \\
&\text{(v) Local Maximum:} && \max(\rho_{z_c}, \rho_{n_c}) \\
&\text{(vi) Local Minimum:} && \min(\rho_{z_c}, \rho_{n_c})
\end{aligned}
\tag{27}
$$

Numerical tests comparing reference states for a variety of pure remapping problems are included in section 6.

Remark: The choice of reference state may be generalized beyond what is considered in this article. For instance, a different reference state may be chosen per-subzone based on local smoothness of the zonal and nodal density fields. In regions near large discontinuities, reference values that compute strictly positive densities are desirable; in smooth regions, the local zone value may be desirable. Note also that (i)–(iv) may be combined into two general forms of the arithmetic and harmonic average,

$$
\mu_c = \alpha_c \rho_{z_c} + (1 - \alpha_c)\rho_{n_c} \qquad \text{and} \qquad \mu_c = \frac{\rho_{z_c}\rho_{n_c}}{\alpha_c \rho_{n_c} + (1 - \alpha_c)\rho_{z_c}}, \tag{28}
$$

respectively, where $\alpha_c \in [0, 1]$ smoothly tracks the contributions of zone and node densities in the reference state per-subzone. Analysis of (28) is a topic of future work.

4. Repairing Subzonal Values

For certain choices of reference state, the algorithm may minimize the objective function (6) but compute subzonal values that are negative. Reconstructed values may also be non-monotonic, i.e. straying outside the bounds

set by their local zonal and nodal values. This is most likely to occur in regions with large discrepancies between zonal and nodal values, for instance when remapping a density field having steep discontinuities.

The reason for this is the extra degrees of freedom in the problem. The number of zonal and nodal constraints is less than the total number of subzones. Only the value of the reference state in each subzone may be used to drive the reconstructed values to ones that preserve positivity, monotonicity, etc.

We consider a simple local repair procedure to correct negative reconstructed values. Subzonal mass is redistributed within each zone in a manner that preserves zone mass exactly (4). The compromise involved in forcing positivity is a small error in node mass conservation (5).

The local repair process acts on zones containing at least one negative subzonal density. For zone z having remapped mass m_z, let $\mathcal{C}^-$ denote the set of subzones having negative reconstructed densities. Here $\mathcal{C}^-$ is a strict subset of $c(z)$ since the remap itself ensures $m_z > 0$. We proceed:

1. Sum the total negative subzonal mass within the zone:

$$M_- = \sum_{c \in \mathcal{C}^-} |\rho_c V_c|.$$

2. Replace all negative subzonal densities with a small positive constant value ϵ chosen to be smaller than the minimum density globally:

$$\rho_c \longleftarrow \epsilon \;\; \forall c \in \mathcal{C}^-.$$

   The amount of mass added to the zone by this replacement is

$$M_\epsilon = \sum_{c \in \mathcal{C}^-} \epsilon V_c.$$

   For all numerical tests, we use $\epsilon = 10^{-10}$.

3. Scale the remaining positive subzonal densities by the constant ratio

$$\rho_c \longleftarrow \left(\frac{m_z - M_\epsilon}{m_z + M_-} \right) \rho_c, \;\; \forall c \in c(z), \; c \notin \mathcal{C}^-.$$

   This is equivalent to subtracting a fraction of the total mass change $M_- + M_\epsilon$ in the zone from each positive subzone mass. The fraction

per subzone is the ratio of its initial mass to total positive zone mass $m_c/(m_z + M_-)$. This ensures all subzone masses are positive and constraint (4) is respected.

In the context of hydrodynamics, the ultimate consequence of redistributing subzonal mass within a zone is an error in momentum conservation. The size of this error scales with the total negative mass M_- and affects the discrete momentum values at each node around a corrected zone. Note that because some subzones gain mass and others lose it during the correction, the net affect on global momentum conservation has no definite sign. Furthermore, if velocity is constant, the correction will produce zero net momentum error.

Remark: Negative reconstructed values typically occur in the presence of steep discontinuities. They are rarely encountered in the context of hydrodynamic calculations because discontinuities—shocks, material fronts, etc.—tend to spread across more than a single zone due to numerical diffusion.

Remark: The small conservation error introduced by local repair is not significant for the types of problems we have considered. For the sake of completeness, we include two additional procedures for enforcing positivity that are fully conservative:

(i) Appendix B describes a global extension of the above local procedure. The global linear system in section 2 is modified by additional equality constraints imposed on subzones and solved a second time to remove negative values. The resulting reconstruction is positive and fully conservative.

(ii) Appendix C generalizes the constrained optimization problem to include inequality constraints. The method is fully conservative. Inequality constraints may also be chosen to enforce not only positivity but also monotonicity in the reconstructed subzonal field.

These algorithms are understandably more expensive than local repair. There is no free lunch for the staggered overlay problem: a full subzonal overlay can be conservative, positive, and monotonic [11] but is very expensive; conservative reconstruction is cheaper, but more work must be done to ensure positivity and monotonicity as well.

5. Subzonal Pressures

The subzonal reconstruction algorithm is well-optimized for use with the corner-based staggered compatible hydrodynamics algorithm of [2]. The discretization is sensitive to the total zonal and nodal masses but not to the individual subzonal values. As a result, the hydro algorithm performs well even if non-monotonicities creep into the reconstructed subzonal density field. The most important property to meet is that the subzonal corner masses sum to the correct remapped zonal and nodal values, which our algorithm ensures by construction.

The same cannot be said about subzonal densities used to compute subzonal pressures that serve as hourglass controls, as in [2]. Subzonal pressures computed from oscillatory subzonal mass densities can introduce numerical instabilities in the Lagrangian mesh. To avoid this complication, we use the anti-hourglassing filter of [3] instead of subzonal pressures to remove spurious grid motion.

Remark: Reconstruction does not forbid us from taking advantage of subzonal pressures to improve mesh motion. Similar to [14], we could use a separate and distinctly simpler definition of corner mass density solely for computing subzonal pressures. This separate subzonal density would not contribute to the momentum equation and therefore would not have to meet the same stringent conservation properties as the reconstructed values. The idea in [14] may be extended to an overlay remap by considering the change in mass inside a zone before and after an overlay:

$$m^r_{c,p} = \begin{cases} \dfrac{V^r_c}{V^r_z}\Delta m_z & , \quad \text{if zone mass increases} \\ \dfrac{m^r_z}{m_z} m_{c,p} & , \quad \text{if zone mass decreases} \end{cases}$$

where $m_{c,p}$ indicates the subzonal pressure corner mass, $\Delta m_z = m^r_z - m_z$ is the change in mass in zone z, and the superscript r indicates a post-remap quantity. This definition will reproduce the remapped zonal mass from summing $m^r_{c,p}$ in the zone, but not the remapped nodal masses. Such nodal reproduction is not necessary in the hourglass control and is therefore a perfectly acceptable compromise. The ability to use the subzonal corner pressure hourglass control will become important when we consider hydrodynamic examples on arbitrary polygonal grids. The kinematic filters of [3], which are specialized for quadrilaterals, are not appropriate for such calculations.

6. Numerical Results

The numerical results in this section result from applying the corner reconstruction algorithms described in this paper to data remapped using an existing geometrical overlay package named Overlink [5]. Overlink uses linear reconstruction with gradient limiting to achieve second-order accuracy and monotonicity in the overlay – arbitrary polyhedral intersections are computed using fast decomposition into tetrahedra. The full details of the algorithm are presented in [5]. For our purpose Overlink is a nice example of an existing geometric overlay remapping package intended for use with staggered data that we wish to generalize for use with subzonal corner data.

6.1. Null Problem

We begin by considering the null problem of overlaying a specified density field onto the same mesh. This is an exact operation using Overlink, thereby allowing us to test the accuracy of different choices of reference density in reconstructing an underlying subzonal density field. We consider two analytic density distributions—linear and discontinuous—and test the reconstruction algorithm on various discretizations of these distributions.

In each test, we compute convergence of the reconstructed subzonal density values to their initial discrete values under L_1 norm. We compute the L_1 norm using the formula

$$L_1 = \frac{1}{M} \sum_{i=1}^{M} |a_i - b_i|, \tag{29}$$

where $\{a_i\}_i$ and $\{b_i\}_i$ are arbitrary sequences of length M. In each example, reconstruction is computed on $N \times N$ square meshes for $N \in \{8, 16, 32, 64, 128\}$.

Test 1: Linear Density Field

First is a linear density distribution. The continuous field is given by

$$\rho(\mathbf{x}) = \rho_0 + \nabla\rho \cdot \mathbf{x}. \tag{30}$$

The field is discretized on each subzone c by evaluating (30) at subzone center $\mathbf{x}_c$. Discrete zonal and nodal values are initialized by summing over subzones. In all subsequent tests, we use $\rho_0 = 10$ and $\nabla\rho = (1, 3)$.

Figure 3 contains convergence curves for each of the six reference states. Dashed lines of slope 1, 2, and 3 illustrate the convergence rates of the

various references. The results match expectations: the local minimum and maximum references are discontinuous functions of the local zonal and nodal values for each subzone; these references may only converge at first order. The zone density, node density, and harmonic average references are smooth functions and converge at second order. The harmonic average reference has smaller absolute error because the reference densities in each subzone are closer to the analytic distribution than the discrete zonal and nodal densities. The arithmetic average reference converges at third order. This is because the reconstruction only commits an error inside subzones on the boundary of the problem—a set having codimension 1. The reference density at all interior subzones is equal to the analytic distribution.

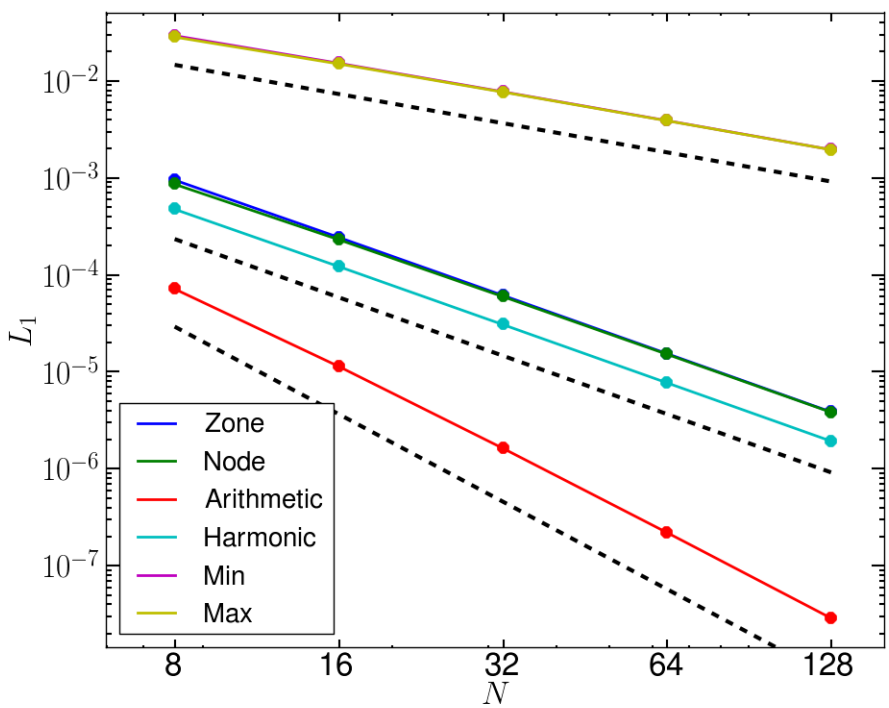

Figure 3: Null problem, linear field: L_1 convergence in subzonal density. Reference slopes 1, 2, and 3.

Test 2: Linear Density Field, Anti-Correlated Subzonal Field

We use the same continuous linear density field (30). Discrete zone densities ρ_z are initialized by evaluating (30) at each zone center $\mathbf{x}_z$. A subzonal density distribution is chosen to have the opposite variation of the analytic field:

$$\rho_{sub_z}(\mathbf{x}) = \rho_z - f\nabla\rho \cdot (\mathbf{x} - \mathbf{x}_z), \quad \mathbf{x} \in \{\text{zone}z\}. \tag{31}$$

The density at each subzone $c \in c(z)$ is (31) evaluated at subzone center $\mathbf{x}_c$. The scale factor f generalizes the subzonal density distribution. If $f = 0$, the subzonal density is flat; if $f = -1$, we recover the density distribution in Test 1; and if $f > 0$, the subzones are anti-correlated with the interzonal variation. Note that (4) is satisfied provided each zone position is equal to the average of its subzone positions. Figure 4 illustrates this density distribution on the zones, nodes, and subzones of an 8×8 square mesh. We use $f = 1$ to illustrate the subzonal field.

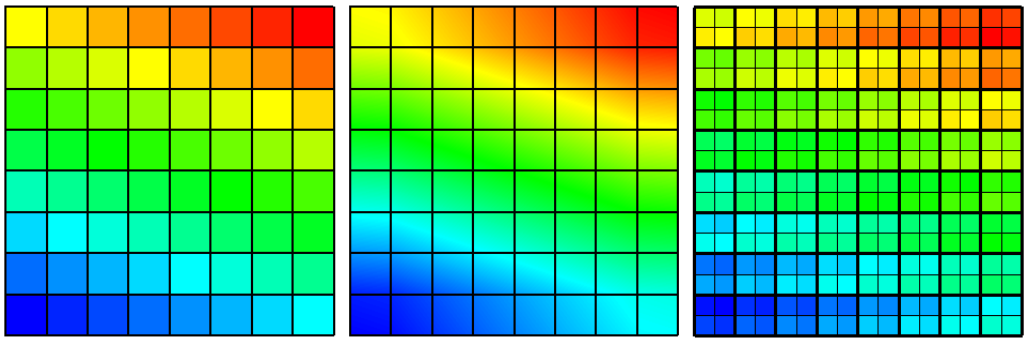

Figure 4: Initial discrete densities in zones (left), at nodes (center), and in subzones (right) for a linear field with anti- correlated subzonal variation.

Figure 5 contains convergence curves for two choices of scale factor f: a constant factor and a variable factor.

Left is the constant scale factor $f = 1$. Dashed black lines have reference slopes 1, 2, and 3. In this example, as N increases, the relative variation within each zone decreases; that is, the subzonal density field becomes smoother under mesh refinement. The relative errors and convergence rates of each reference match that in Test 1. This test is again advantageous for the arithmetic average reference since it is exact in all interior subzones. Unlike

Test 1, the zone density reference has lower absolute error than node density. This is because the discrete node densities are farther from the analytic profile at the boundary than the discrete zone densities when their values are initialized using (31).

The right plot contains a more challenging test for the reconstruction algorithm: a variable scale factor is chosen to increase with mesh resolution. The scale factor is as large as possible without having negative values in the initial discrete subzonal density field:

$$f_{max} = \frac{4\rho_0 N}{|\rho_x| + |\rho_y|}.$$

The resulting subzonal field is discontinuous across zones despite the fact that the coarsened zonal and nodal fields are smooth. The dashed black line has slope 1. All references converge at first order, as expected for the discontinuous nature of the subzonal field. The zone density reference has minimal absolute error by nearly an order of magnitude. This is again due to the zone values being closer to the analytic profile at boundary subzones and thus committing a smaller error there. The steeper subzonal slope only exacerbates this effect.

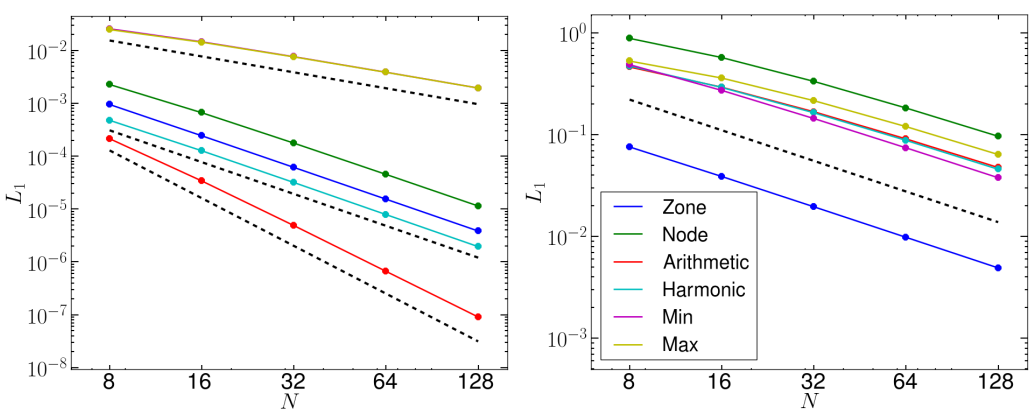

Figure 5: Null problem, linear field: L_1 convergence in subzonal density. Scale factor $f = 1$ with reference slopes 1, 2, and 3 (left). Maximum scale factor f_{max} with reference slope 1 (right).

Test 3: Discontinuous Density Field, Zero Subzonal Variation

We consider a discontinuous density profile. Let $\mathbf{r}$ point in the direction perpendicular to a linear discontinuity passing through point $\mathbf{c}$ and separating high and low density values ρ_H and ρ_L. The initial density profile is

$$\rho(\mathbf{x}) = \begin{cases} \rho_H & , \quad \mathbf{r} \cdot (\mathbf{x} - \mathbf{c}) < 0 \\ \rho_L & , \quad \mathbf{r} \cdot (\mathbf{x} - \mathbf{c}) \geq 0 \end{cases} . \tag{32}$$

We first initialize the discrete density profile so there is no subzonal variation. The discrete zone density ρ_z is computed by evaluating (32) at zone center $\mathbf{x}_z$. The discrete density in each subzone is equal to its associated zone density: $\rho_c = \rho_{z_c}$. For all subsequent examples, we set $\mathbf{r} = (4/3, 1)$, $\mathbf{c} = (1/2, 1/2)$, and $\rho_H = 1$. We vary ρ_L to test the robustness of each reference to large density jumps.

Figure 6 contains convergence curves for the L_1 error in subzonal density for two choices of ρ_L: 10^{-1} and 10^{-6}. The black dashed lines have reference slope of 1. The error curve for zone density is not provided because the reconstructed subzonal densities match their initial values to machine precision. This is to be expected for this particular test: the zone density reference automatically satisfies mass conservation constraints at all zones and nodes. Here, and in subsequent plots, solid markers indicate reconstructed densities that maintain positivity; 'X' markers indicate negative reconstructed values that have been repaired according to the nonconservative local procedure. It is worth mentioning that local repair produces an error in a very small set of subzones and is not a significant contribution to the global L_1 error.

The node density and local maximum references fail to preserve positivity for both density jumps and carry the largest absolute errors. The arithmetic average reference maintains positivity for smaller density jump but fails when the discontinuity is more extreme. The local minimum and harmonic average references are robust to computing negative states (even in an extreme case of $\rho_L = 10^{-12}$, not included.) The harmonic average outperforms the local minimum in all cases due to it being a smooth function. First order convergence is established for smaller density jumps. For the larger jump, the harmonic average and local minimum do not converge: errors are only committed at subzones along the discontinuity. The size of these errors and the number of subzones scale with the mesh resolution. Regardless, the errors are several orders of magnitude smaller than those committed by the other reference states.

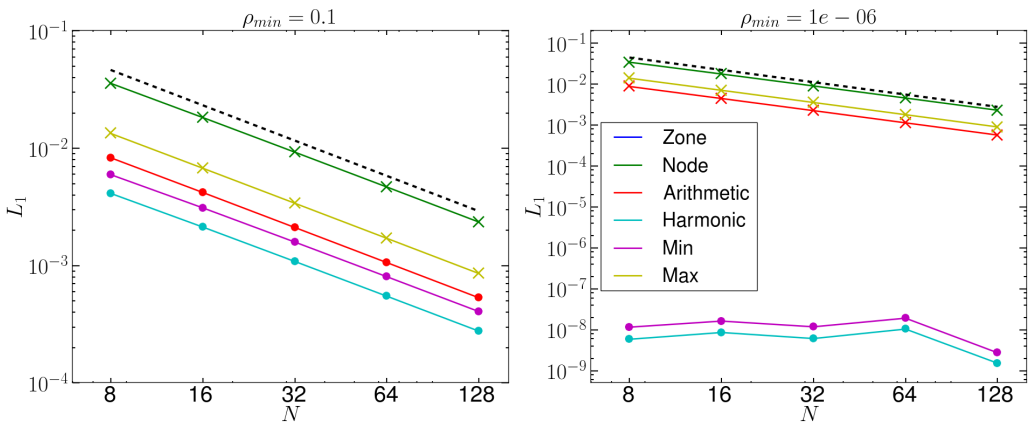

Figure 6: Null problem, discontinuous field: L_1 convergence in subzonal density for density jump of 10 (left) and 10^6 (right). No variation in the initial subzonal density. Positive reconstructions have solid markers, repaired negative reconstructions have X.

Test 4: Discontinuous Density Field

We also consider the discontinuous density profile (32) in this example. The discrete profile is initialized at the subzonal level: ρ_c is computed by evaluating (32) at subzone center $\mathbf{x}_c$. Discrete zonal and nodal densities are initialized by summing over local subzones. The overall effect is a transition region in the zonal and nodal density profiles: zone and node elements intersected by the discontinuity have an intermediate density value between ρ_L and ρ_R.

Figure 7 contains convergence curves for each reference state, again comparing for density jumps of 10 and 10^6. The density reference is no longer an exact solution, and we have included it here. Results are to be expected, with first-order convergence established for each reference state. Positivity of the reconstruction matches the results in Test 3: only the harmonic average and local minimum references reconstruct positive subzonal densities when the discontinuity is extreme.

6.2. Reconstructing a Linear Density Distribution

We next test how the reconstruction performs when remapping onto a target mesh with different topology and geometry. We begin with the linear density distribution (30), again with $\rho_0 = 10$ and $\nabla\rho = (1, 3)$, on a square

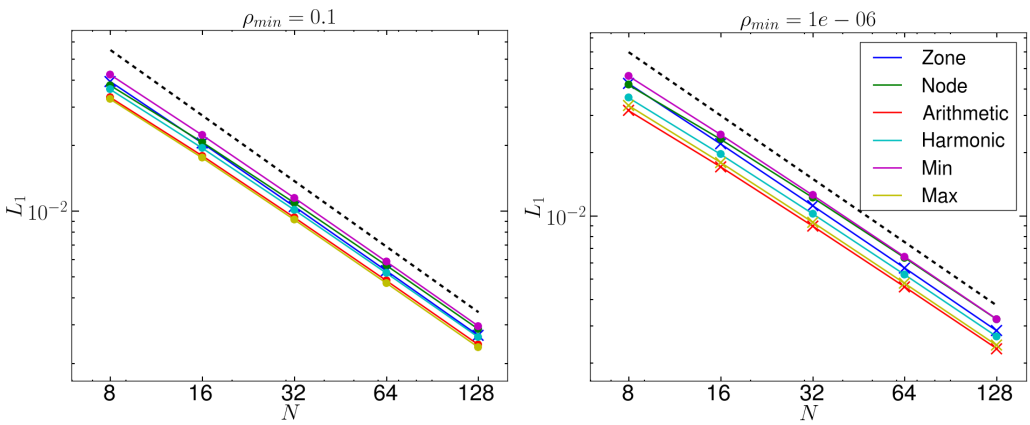

Figure 7: Null problem, discontinuous field: L_1 convergence in subzonal density for density jump of 10 (left) and 10^6 (right). Positive reconstructions have solid markers, repaired negative reconstructions have X.

$N \times N$ mesh. The subzonal variation follows the same distribution as the global density (i.e. equation (31) with $f = -1$).

We consider two target meshes: an unstructured mesh of arbitrary polygons and a structured mesh of high-aspect-ratio quads. The unstructured mesh consists of a random Voronoi tessellation having N^2 number of zones. The structured mesh consists of a $rN \times \lfloor \frac{N}{r} \rfloor$ lattice of rectangles, with $r > 1$. For all subsequent tests, we use $N \in \{8, 16, 32, 64, 128\}$ and $r = 4$. Figure 8 illustrates the zonal density on the initial mesh as well as its remapped values on the unstructured Voronoi mesh and high-aspect quad mesh.

Figure 9 gives L_1 convergence results for the reconstructed subzone densities relative to their analytic values from (30). For the Voronoi mesh (left), all reference values demonstrate first-order convergence. Arithmetic and harmonic averages are identical and outperform the other references on an arbitrary polygonal grid. In the high-aspect quad case (right), all smooth references converge at second order and non-smooth references (local minimum and maximum) at first order. The regularity of the high-aspect quad zones allows for second-order convergence in the reconstructed subzonal densities; second-order convergence on a polygonal mesh can be achieved provided greater regularity and mesh quality–for instance, a centroidal Voronoi tessellation composed of regular hexagons.

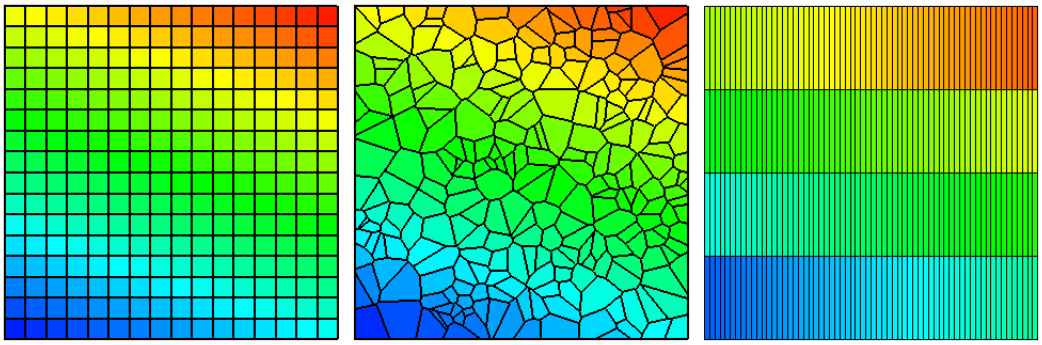

Figure 8: Initial linear zonal density (left) and remapped zonal density on unstructured polygons (center) and high- aspect-ratio quads (right).

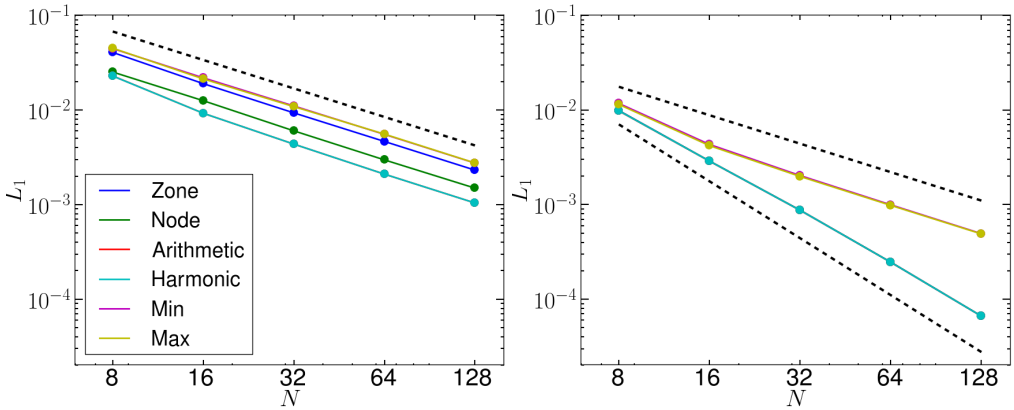

Figure 9: Reconstruct linear density field: L_1 convergence in subzonal density. Unstructured polygonal target mesh with reference slope 1 (left) and high-aspect-ratio target mesh with reference slopes 1 and 2 (right).

6.3. Reconstructing a Density Discontinuity

We repeat the above test for a density discontinuity of size 10. We initialize the density fields in the same manner as Test 4 for the null problem tests. Here $\rho_L = 10^{-1}$.

We again map onto an unstructured random Voronoi tessellation and a structured grid of high-aspect quads. Figure 10 illustrates this in zone density. Figure 11 gives convergence results for each reference state. Again,

solid markers indicate positive reconstructed densities, and 'X' markers indicate locally repaired negative densities. Because of the modest density jump and relative grid regularity, all references produce positive reconstructions for the high-aspect quad case. Only the zone density and local maximum references compute negative subzonal densities for the random Voronoi case. This is due to significant grid irregularity: the existence of small edges in the Voronoi mesh leads to disparate subzonal volumes. Negative reconstructed values tend to occur in or near zones having small edges. Improving the quality of the grid, for instance by deleting or expanding small edges in the mesh, reduces negative reconstructions.

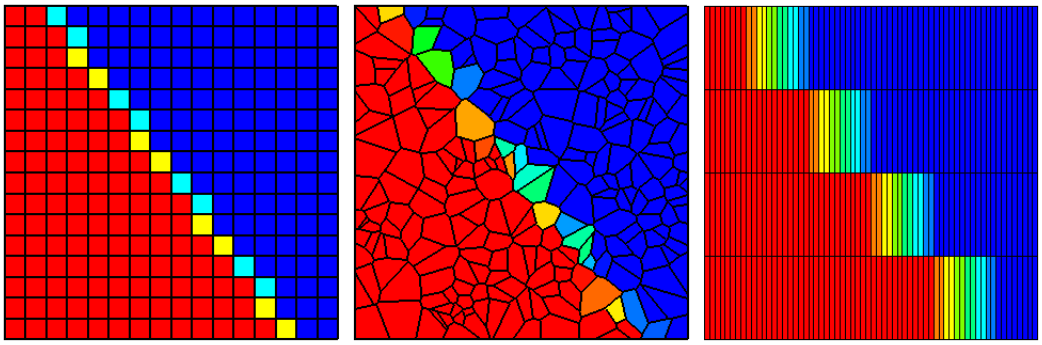

Figure 10: Initial discontinuous zonal density (left) and remapped zonal density on unstructured polygons (center) and high-aspect-ratio quads (right).

We also measure how well monotonicity is preserved in the reconstructed subzonal density field. For each test, we compute the ratio

$$r_{mono} = \min_c \rho_c \Big/ \max_c \rho_c. \tag{33}$$

If the reconstruction is monotone, r_{mono} will be 0.1 for this problem. Ratio values are provided in table 1 and table 2 for the unstructured Voronoi and high-aspect quad target meshes respectively. Empty entries correspond to cases of negative reconstructed densities. In all instances, the reconstructed values are non-monotonic, with more extreme ratio values occurring for the unstructured Voronoi target mesh. It is worth reiterating that the monotonicity being measured here is in the *subzonal* densities: the directly remapped zonal and nodal densities maintain monotonicity exactly. Although the subzonal densities do not respect the monotonicity limits of their coarsened nodal

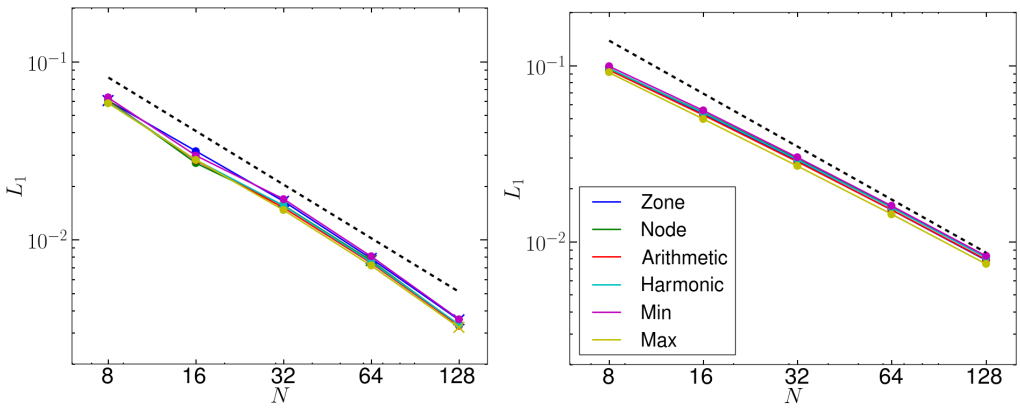

Figure 11: Reconstruct discontinuous density field: L_1 convergence in subzonal density. Unstructured polygonal target mesh (left) and high-aspect-ratio target mesh (right), each with reference slopes of 1. Positive reconstructions have solid markers, repaired negative reconstructions have X.

and zonal values, they do rigorously reproduce the remapped zonal and nodal mass densities due to the constraints imposed. This is adequate in the context of the staggered hydrodynamic method towards which reconstruction is applied; it is the compromise we accept in order to allow the significant computational savings compared with directly remapping subzonal quantities.

From the data collected, the zone density, node density, and local maximum reference states appear unsatisfactory due to both the size of monotonicity errors and their tendency to compute negative values. The local minimum reference is most robust to computing negative states; however, it is a discontinuous function of local data and has poor convergence properties for smooth fields as a result. The arithmetic and harmonic averages have similar convergence properties, but the latter is more robust to computing negative states. The remaining tests use the harmonic average reference for this reason. Additional reference comparisons are included to aid discussion, but the harmonic average is our choice moving forward.

6.4. Cyclic Remap Tests

This is a remapping test based on the work of [9] and used in [14] to test flux-based ALE remap strategies. The initial mesh configuration consists of a lattice of quads. Mesh nodes are iteratively moved according to the following

N	Zone	Node	Arith. Avg.	Harm. Avg.	Local Min	Local Max
8		0.0266	0.0436	0.0566	0.0667	0.0455
16	0.0027	0.0716	0.0583	0.0671	0.0560	0.0453
32		0.0331	0.0198	0.0281	0.0295	0.0176
64			0.0354	0.0350	0.0425	0.0222
128		0.0129	0.0260	0.0359	0.0342	

Table 1: Reconstruct discontinuous density field: monotonicity metric convergence for unstructured polygonal target mesh. (Monotone reconstruction has metric value 0.1.) Missing entries indicate cases that failed due to negative reconstructed densities.

N	Zone	Node	Arith. Avg.	Harm. Avg.	Local Min	Local Max
8	0.0547	0.0149	0.0646	0.0665	0.0782	0.0565
16	0.0530	0.0141	0.0651	0.0663	0.0782	0.0572
32	0.0524	0.0142	0.0654	0.0663	0.0782	0.0575
64	0.0522	0.0143	0.0655	0.0664	0.0783	0.0577
128	0.0522	0.0143	0.0656	0.0664	0.0783	0.0578

Table 2: Reconstruct discontinuous density field: monotonicity metric convergence for high-aspect quad target mesh. (Monotone reconstruction has metric value 0.1.)

periodic displacement:

$$
\begin{aligned}
x_k(x_0, y_0, \tau) &= x_0 + F(\tau)\sin(2\pi\tau x_0)\sin(2\pi\tau y_0) \\
y_k(x_0, y_0, \tau) &= y_0 + F(\tau)\sin(2\pi\tau x_0)\sin(2\pi\tau y_0) \\
F(\tau) &= \begin{cases} 0.2\tau & , \quad \tau \in [0,\ 0.5] \\ 0.2(1-\tau) & , \quad \tau \in [0.5,\ 1] \end{cases} \\
\tau &= k/k_{max}
\end{aligned}
\tag{34}
$$

Here k is the remap iteration, (x_0, y_0) is the initial position of each node, (x_k, y_k) is the k-th position, and τ is the pseudo-time period of the displacement. Mesh nodes return to their original positions after k_{max} total iterations.

Initial background density and velocity fields are iteratively overlaid onto the new displaced mesh. We consider two initial fields: one smooth and one discontinuous.

Smooth Fields

We consider smooth profiles for density and velocity

$$\begin{aligned} \rho(x,y) &= \rho_{back} + \rho_{amp}\sin(2\pi x)\sin(2\pi y) \\ u(x,y) &= \frac{x}{\sqrt{x^2+y^2}} V_{amp}\sin(2\pi x)\sin(2\pi y) \\ v(x,y) &= \frac{y}{\sqrt{x^2+y^2}} V_{amp}\sin(2\pi x)\sin(2\pi y) \end{aligned} \tag{35}$$

Figure 12 illustrates these profiles for $\rho_{back} = \rho_{amp} = 1$ and velocity amplitude $V_{amp} = 1$.

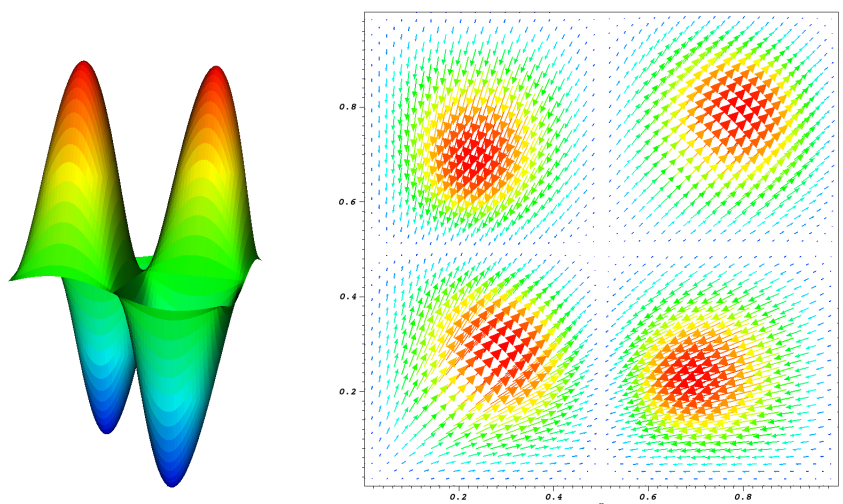

Figure 12: Surface plot of the initial smooth density field (left) and vector plot of the initial smooth velocity field (right).

Differences in accuracy among the various reference states is negligible due to the smoothness of the underlying profiles. Figure 13 reports L_1 convergence results for the remapped zonal density, nodal density, and radial velocity for the harmonic average reference state only. Results correspond to an $N \times N$ square mesh with $N \in \{32, 64, 128, 256, 512\}$. The overlay remap

package is formally second-order accurate. Our reported convergence rates converge the order of accuracy of the overlay: 2.04, 1.97, and 2.12 for zonal density, nodal density, and radial velocity respectively.

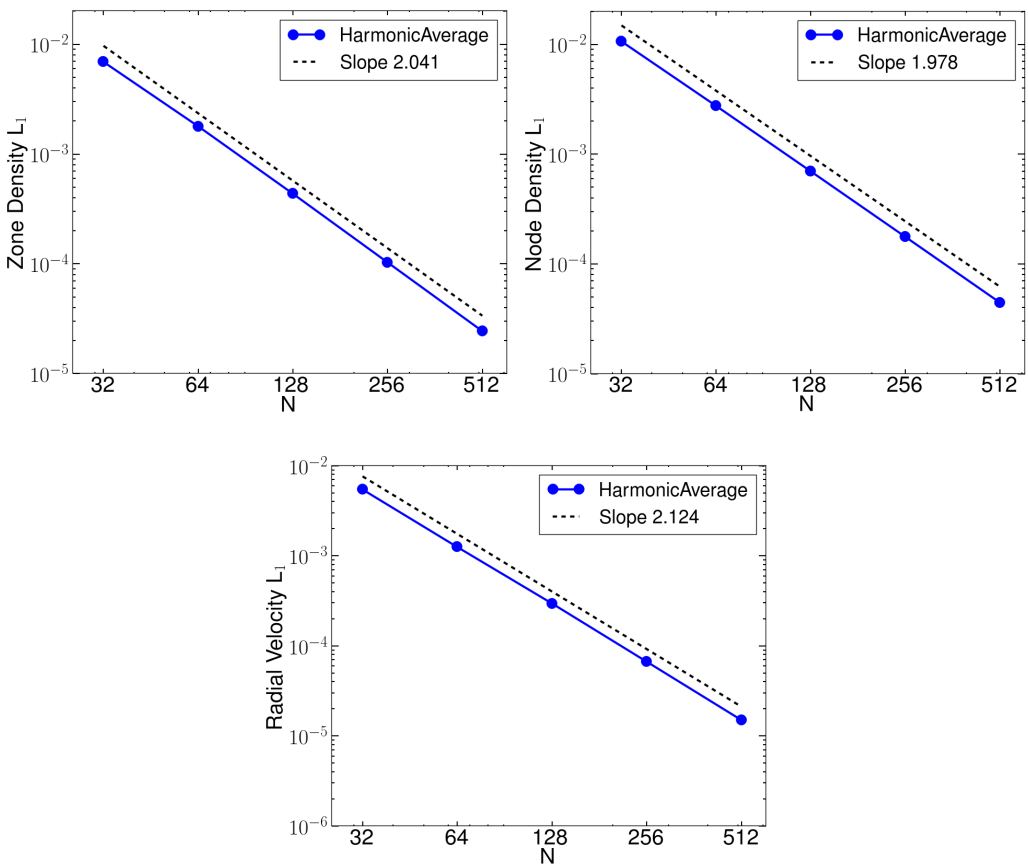

Figure 13: Cyclic remap, smooth fields: L_1 convergence in remapped zonal density, nodal density, and radial velocity with linear fit.

Discontinuous Fields

We consider two discontinuous regions of the mesh: a high-density cylindrical shell, having $\rho_{shell} = 1$ and unit, outward-pointing velocity $\mathbf{u}_{shell} = \hat{\mathbf{x}}_r$ and a low-density background having $\rho_{background} = 1/J$ and zero velocity. Here $\mathbf{x}_r = \mathbf{x} - (0.5, 0.5)$ defines a radial position and $J > 1$ defines a jump in

density. The shell has inner and outer radii $r_{inner} = 0.25$ and $r_{outer} = 0.45$, centered on the unit square. Figure 14 gives the initial zonal density and nodal velocity magnitude.

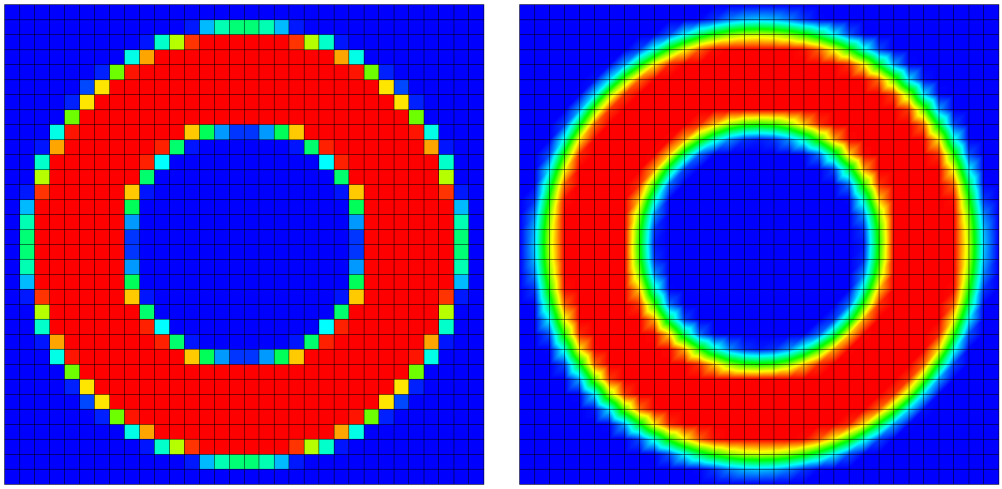

Figure 14: Initial zonal density (left) and velocity magnitude (right) for the discontinuous cyclic remap test.

Figure 15 gives the nodal velocity field at the time of maximum distortion ($\tau = 1/2$) and at the final time ($\tau = 1$) in which the mesh returns to its initial configuration. Velocity is an important quantity to consider when judging the accuracy of a remap for a staggered discretization. Velocity itself is not a conserved quantity; it is remapped instead through the node-centered momentum density $\mathbf{p}_n = \rho_n \mathbf{u}_n$. The post-remap velocity is recovered by $\widetilde{\mathbf{u}}_n = \widetilde{\mathbf{p}}_n / \widetilde{\rho}_n$. If the post-remap subzonal densities do not satisfy (5) precisely, as they are constrained to do in the reconstruction procedure, then linear momentum will not be conserved.

Figure 16 gives the pseudo-time history of total momentum for density jumps of $J = 10$ and $J = 1000$. We compare the ability of the various reference states to respect an initial state having zero total momentum. Discontinuous jumps in total momentum correspond to negative reconstructed subzonal densities. By applying local repair to negative states, we violate nodal mass conservation, thus creating a small error in momentum. When the density discontinuity is small, most reference states (with the obvious ex-

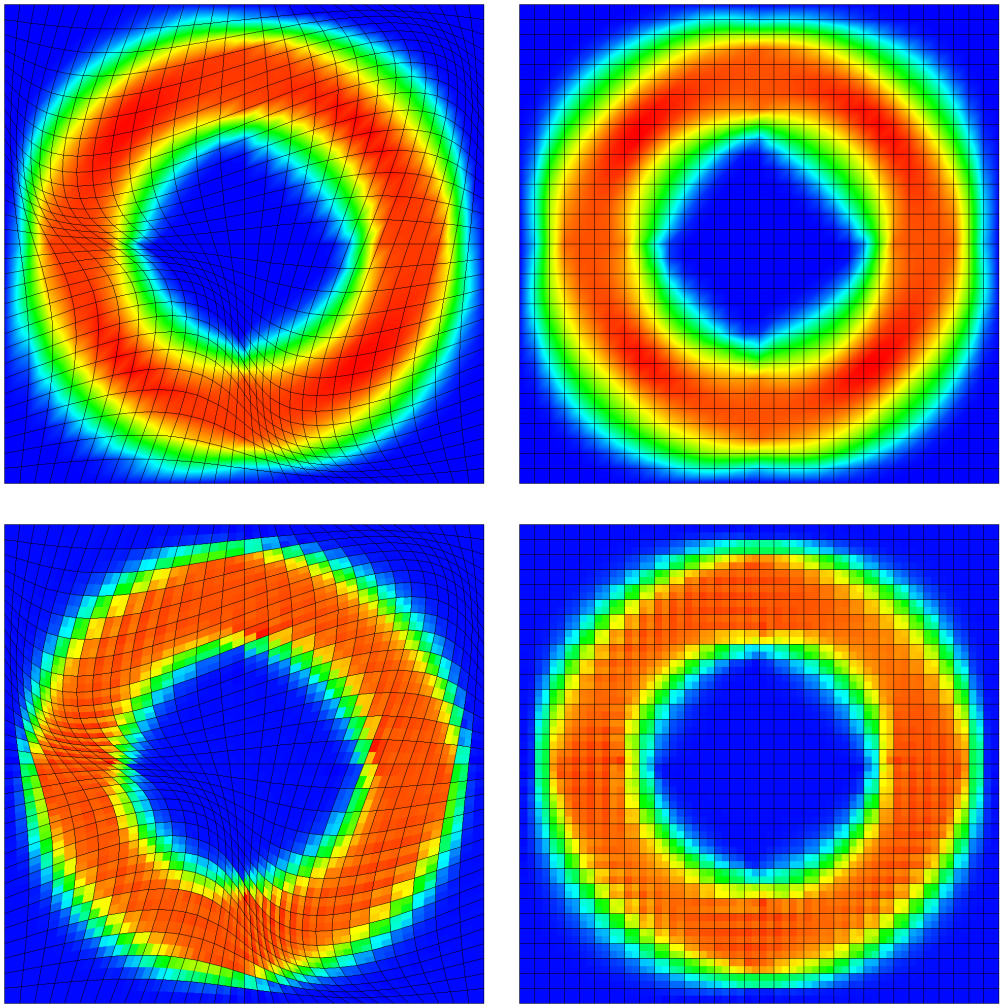

Figure 15: Cyclic remap, discontinuous fields: velocity magnitude (top row) and subzonal density (bottom row). Fields at maximum mesh distortion ($\tau = 1/2$, left column) and after distortion is reversed ($\tau = 1$, right column). Subzonal density reconstruction with harmonic average reference was applied to make these plots.

ception of nodal density) compute only positive reconstructed values and respect momentum conservation to floating point precision. When the density discontinuity is increased, more reference states produce negative subzonal

densities (with the exception of the harmonic average and local minimum). For $J = 1000$, the nodal density reference ultimately fails: repeated conservation errors eventually produce an ill-conditioned linear system, and the linear solver's preconditioner diverges. It should also be pointed out that even though the zone density reference computes a negative subzonal density state at maximum displacement, the total conservation error in momentum is only 10^{-6} when local repair is applied. Nonetheless, the harmonic average and local minimum succeed in avoiding negative reconstructed densities, do not consequently require repair, and maintain conservation to machine accuracy in all cases. Clearly these two references are preferable for robustness and positivity, though as noted in the previous examples the harmonic average reference is more accurate than the local minimum and therefore should be preferred overall.

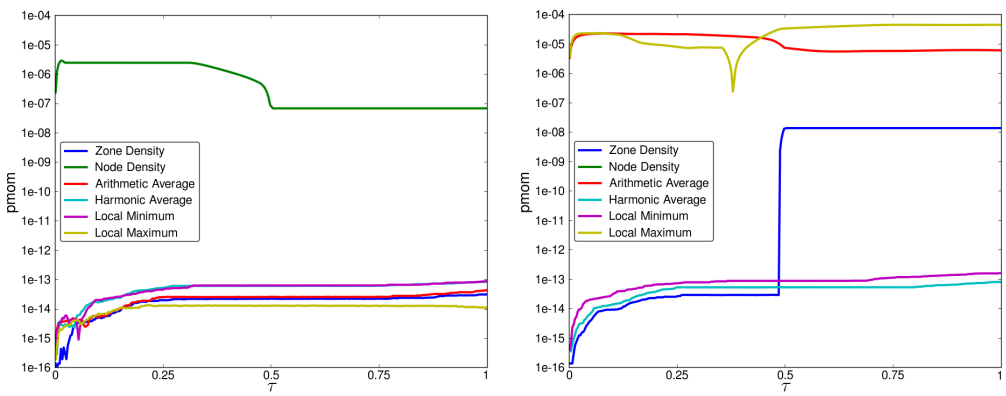

Figure 16: Cyclic remap, discontinuous fields: total momentum as a function of pseudo-time τ for initial density jumps of $J = 10$ (left) and $J = 10^3$ (right).

Figure 17 compares the time history of the ratio of minimum to maximum subzonal density for the two initial density discontinuities and for each reference state. This quantity will be constant in pseudo-time for a monotone remap and reconstruction. We have not imposed a monotonicity constraint on the reconstructed densities, so the monotonicity violations do not come as a surprise. (See appendix C for a discussion on monotonicity constraints.) Note that a similar ratio for zonal and nodal density *does* remain constant. This reassures us that the geometric overlay itself is monotone; all violations

seen in the reconstructed subzonal densities result from the fact the new algorithm does not have monotonicity constraints. If desired, the addition of a repair procedure [12] or additional inequality constraints (appendix C) could be used to mitigate monotonicity violations in the reconstructed values.

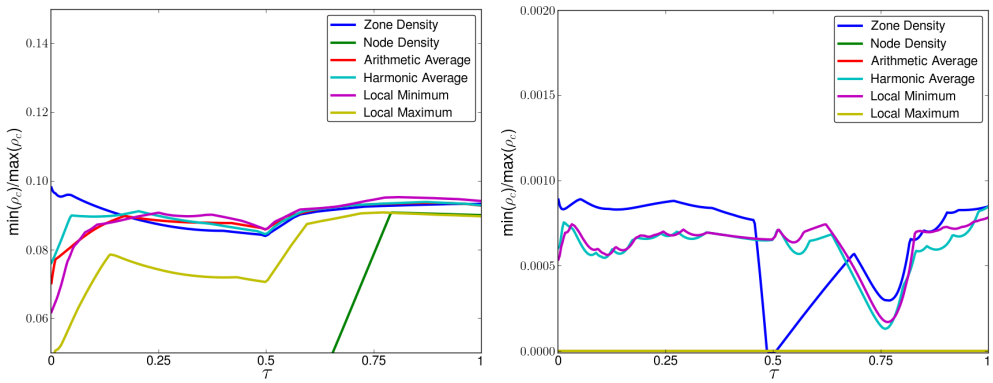

Figure 17: Cyclic remap, discontinuous fields: monotonicity ratio as a function of pseudo-time τ for initial density jumps $J = 10$ (left) and $J = 10^3$ (right). (Monotone reconstruction has constant value $1/J$.)

Figure 18 gives L_1 convergence results for zonal density, nodal density, and radial velocity. The existence of discontinuities in density and velocity imply that the maximum possible convergence rate is one. We report rates of 0.66, 0.54, and 0.55 for zonal density, nodal density, and radial velocity respectively. This is consistent with results in [14] for a comparable flux-based ALE procedure.

6.5. Sedov-Taylor Blast Wave

We now consider a hydrodynamic test involving a Lagrangian update followed by a overlay remap. We analyze the Sedov-Taylor blast wave [15, 16, 17]; this models the shock expansion of an initially pressureless, constant-density gas due to a high-energy point source. Exact similarity solutions exist in planar (1D), cylindrical (2D), and spherical (3D) geometries. We initialize the blast wave with a unit energy density spike in the center of a zero-pressure ideal gas having $\gamma = 1.4$ and unit initial density. This yields a peak shock density of 6. Each simulation is advanced until the analytic solution predicts the shock position to be at a radial distance of 0.4.

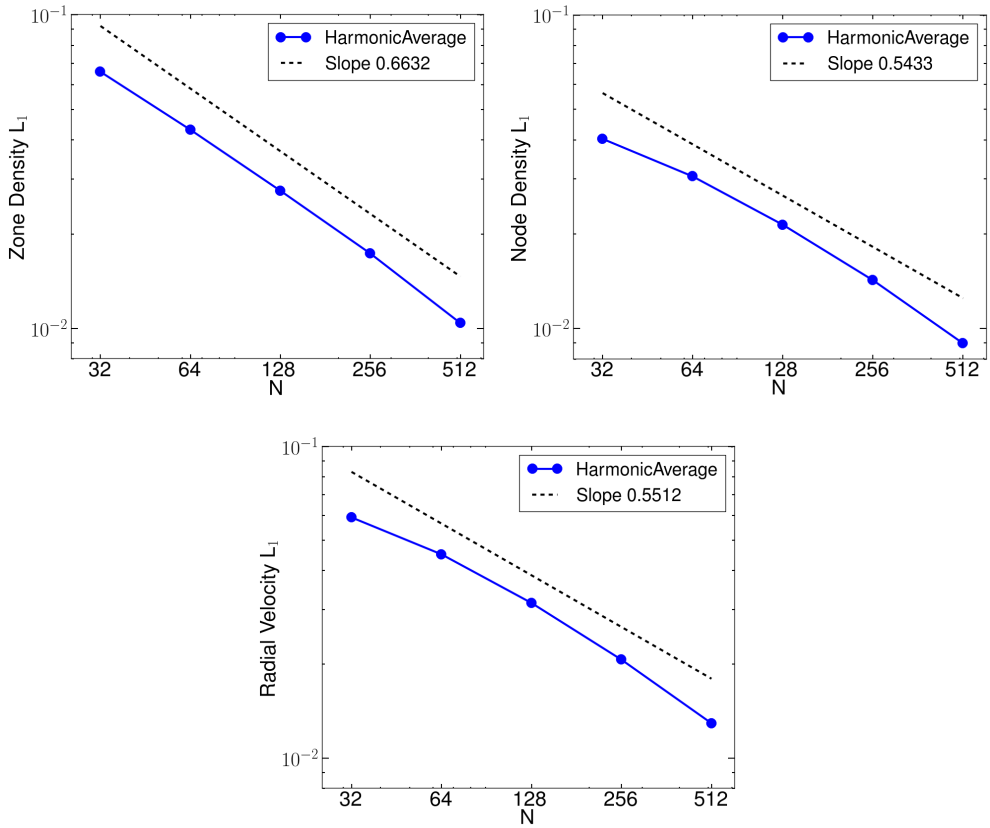

Figure 18: Cyclic remap, discontinuous fields: L_1 convergence in remapped zonal density, nodal density, and radial velocity with linear fit.

We compute numerical solutions on two separate grids: a regular $N \times N$ square grid and a radially-symmetric unstructured grid of arbitrary polygons. Solution variables are updated using the staggered, compatible Lagrangian hydro discretization [2]. Mesh nodes move due to the Lagrangian-frame fluid velocity. At the end of each hydro step, we perform an ALE-based remap: all node positions are relaxed using one iteration of the Winslow-Crowley-based elliptic mesh smoothing algorithm described in [1]. Updated solution variables are then overlaid onto the relaxed mesh to complete the simulation cycle.

The test setup is intended to verify the convergence properties of the new

reconstruction algorithm by applying it at every cycle of a simplified ALE calculation. The ALE strategy itself—relaxing every mesh node mesh at every cycle and overlaying the solution—is not intended to be more numerically-accurate. In fact, a purely-Lagrangian solution maintains a steeper shock profile, achieves a higher peak shock density, and keeps mesh resolution closer to the shock front in each Sedov example. By comparison, our remap step regularizes the mesh at the shock front and introduces additional numerical diffusion during the overlay. Our goal is to establish numerical convergence to the analytic solution and compare results to [14], in which identical Sedov tests were performed using a conservative flux-based ALE remap.

2D Blast Wave on a Regular Grid

As an example of the necessity for a subzonal remap treatment that respects both the zonal and nodal remapped masses, we first compare the results of a flattened reconstruction to the new reconstruction algorithm with harmonic average reference state. Nodal density before and after a remap are inconsistent for the flattened reconstruction, and linear momentum is consequently not conserved. The effect is most apparent in pressure. Figure 19 compares the computed final pressure profile on a 128×128 square grid using the two reconstruction strategies. In the flattened reconstruction example, the shock front lags the analytic location, and the post-shock state is incorrect. The new reconstruction algorithm conserves linear momentum to round-off and correctly captures the analytic shock front.

Figures 20 and 21 compare the zone density contours and radial profiles for the flattened and reconstructed calculations respectively. The solution using reconstruction captures the shock front more accurately and maintains a steeper peak density. Numerical diffusion at the shock front is to be expected and is due to our relaxing the mesh and overlaying the solution at every cycle. Our reasons for doing so are discussed above.

Figure 22 contains the reconstructed subzonal density and radial profile for this problem. The reconstructed density matches the same solution profile as fig. 21. However, monotonicity errors are apparent throughout the solution. Distribution of corner density within a zone is not constant; in some places, it is even misaligned with the interzonal density variation, with overshoots and undershoots visible across zone faces. The lack of strict monotonicity manifests as increased scatter in the radial plot. The Lagrange hydro is not sensitive to this effect since the coarsened zonal and nodal representations of density are each monotone. We have overlaid their values in

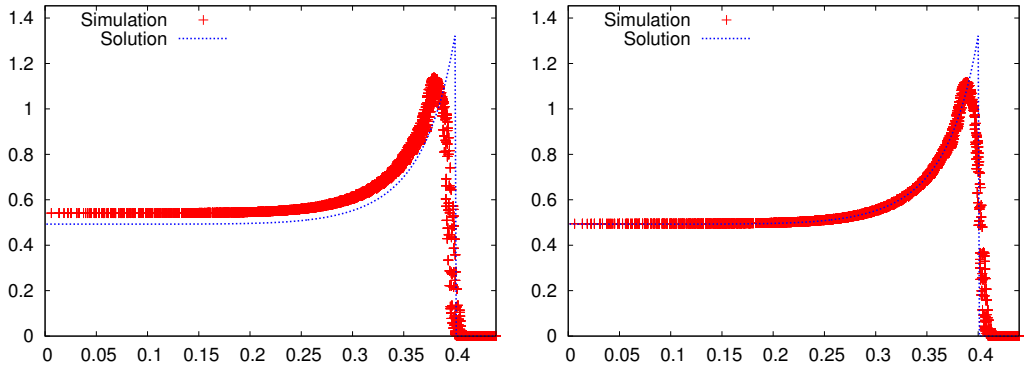

Figure 19: Sedov problem on 128×128 square mesh. Final pressure profile for flattened reconstruction of the subzonal densities (left) and a variation-diminishing reconstruction with harmonic average reference (right).

fig. 22 to better illustrate this.

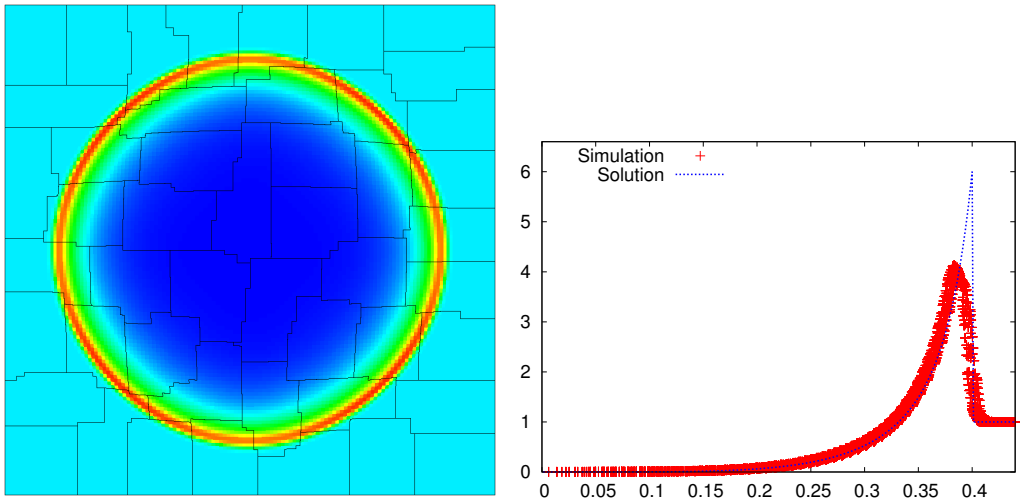

Figure 20: Sedov problem on 128×128 square mesh. Zonal density using flattened reconstruction.

Figure 23 compares convergence rates for flattened subzonal densities and reconstructed values using harmonic average reference. L_1 error is evaluated between the computed density, radial velocity, and pressure and their analytic

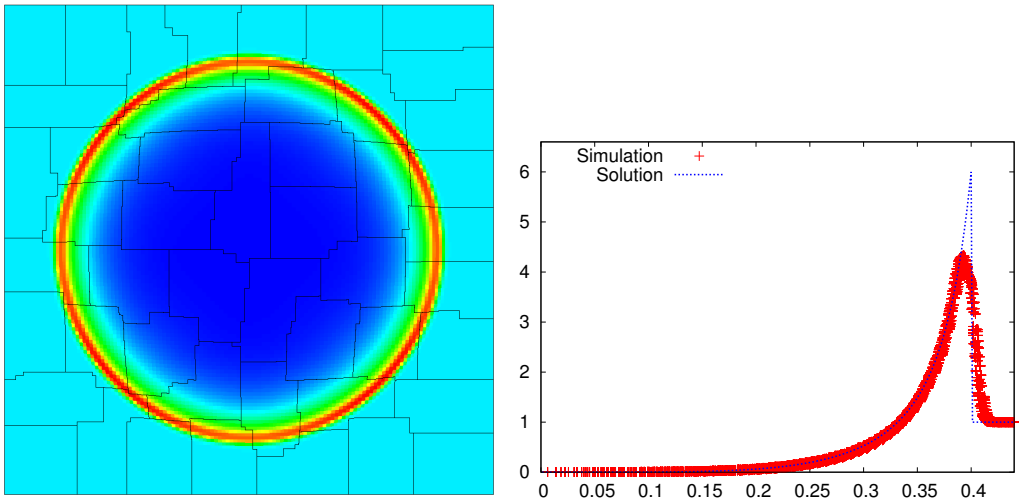

Figure 21: Sedov problem on 128×128 square mesh. Zonal density using variation-diminishing reconstruction with harmonic average reference.

values at the same spatial locations. Results were computed on an $N \times N$ square mesh for $N \in \{32, 64, 128, 256, 512\}$. Conservation errors cause the flattened reconstruction to diverge under mesh refinement, as expected for a shock-dominated test problem such as this. The presence of a strong, symmetric shock front and a relatively-stagnant post-shock state makes the Sedov blast wave a notoriously difficult problem to demonstrate optimal first-order convergence in most hydro schemes. We report converge rates of 0.72, 0.78, and 0.83 in density, radial velocity, and pressure respectively. Results are consistent with those obtained in [14] using the same Lagrangian hydro solver.

2D Blast Wave on an Unstructured Polygonal Grid

We repeat the above calculation on a grid composed of arbitrary polygons. We employ the same Voronoi gridding strategy as [14]: the initial mesh consists of a Voronoi tessellation constructed from regularly-spaced rings of generating points. For N_r radial rings, we have a radial grid spacing of $1/N_r$ between rings. Generating points are distributed on each ring so that the circumferential spacing matches the radial grid spacing as best as possible. We consider meshes corresponding to $N_r \in \{16, 32, 64\}$.

Figures 24 and 25 compare the zone density contours and radial pro-

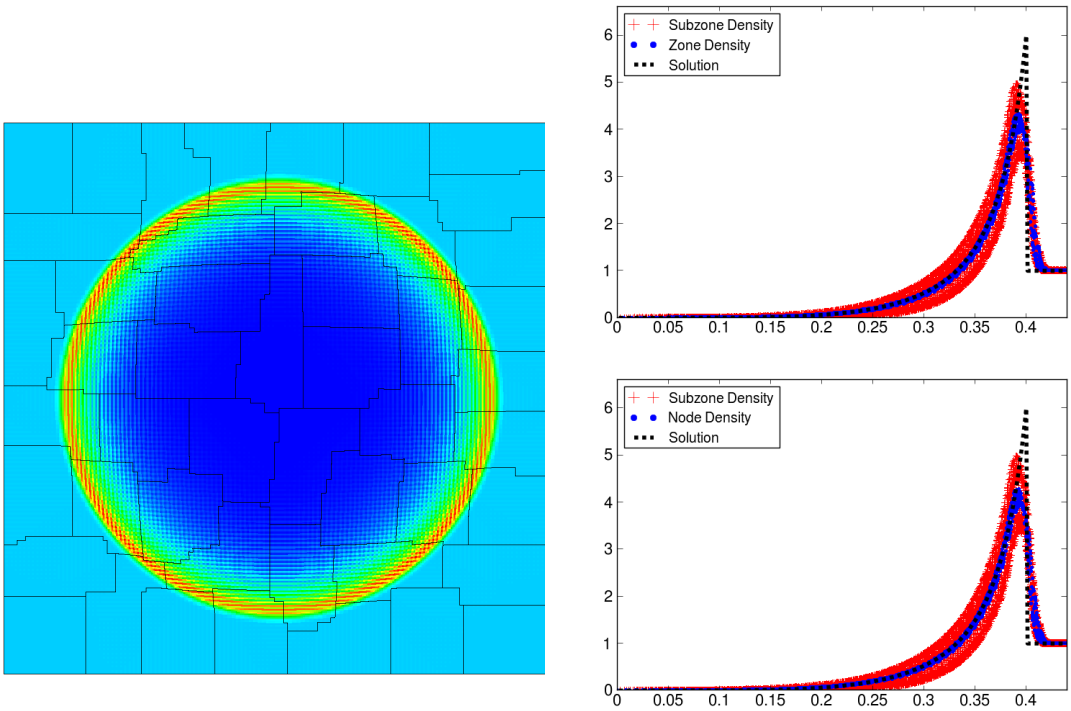

Figure 22: Sedov problem on 128×128 square mesh. Reconstructed subzonal density with harmonic average reference. Corresponding nodal and zonal densities are overlaid.

files for the flattened and reconstructed calculations respectively. Again, the reconstructed solution attains a steeper peak density and more-accurately captures the shock front. The effect is more exaggerated on an unstructured polygonal grid compared to the structured quad results. The reconstructed subzonal densities are illustrated in fig. 26. As before, the coarsened density fields are overlaid onto the radial plot to illustrate that zonal and nodal density remain monotone.

Figure 27 compares convergence rates for the flattened and variation-diminishing reconstruction. As with the structured mesh calculation, the flattened solution diverges under mesh refinement due to conservation errors. We report convergence rates of 0.58, 0.82, and 0.64 for density, radial velocity, and pressure respectively for the reconstructed solution. Results are again comparable to those in [14]

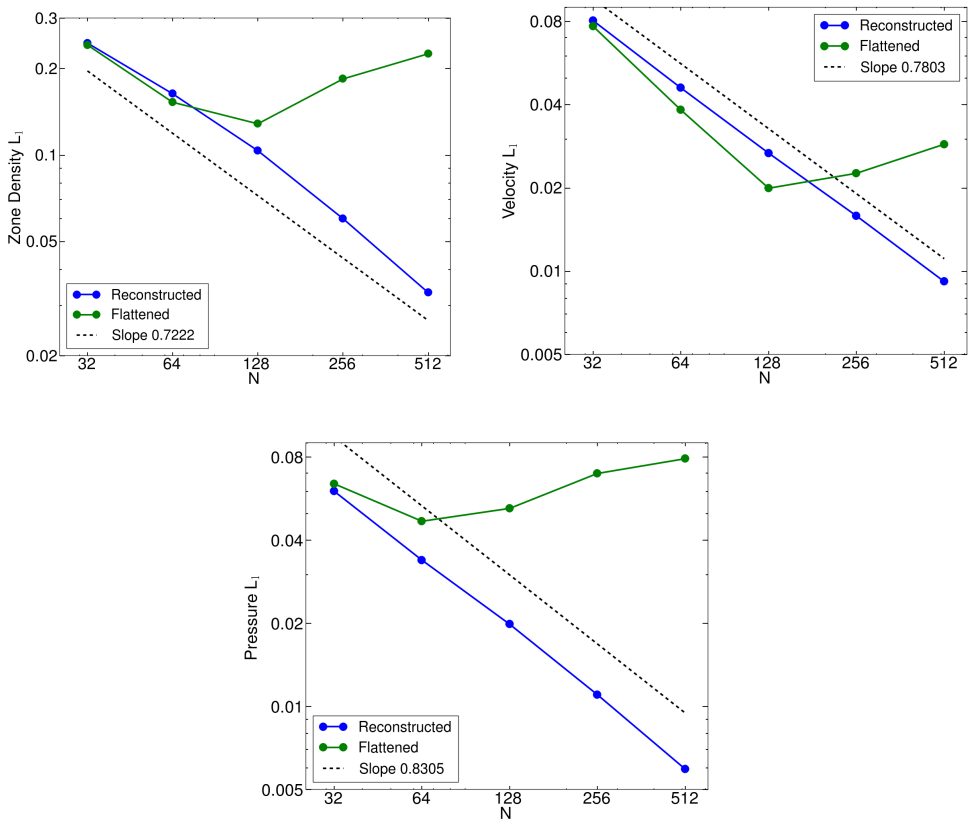

Figure 23: Sedov problem on a structured quad mesh. Plotted are L_1 convergence in zone density, radial velocity, and pressure comparing a flattened reconstruction to a variation-diminishing reconstruction with harmonic average reference.

7. Future Work: Generalization to Subzonal Temperatures

The mathematical framework outlined in section 2 may be used to reconstruct any subzonal variable, provided it can be expressed as a conserved quantity per unit volume. In the case of spatially-staggered diffusion radiation transport, it is often advantageous to have a subzonal description of the material temperature to improve the accuracy of radiation-hydrodynamic coupling terms (i.e. energy exchange through thermal emission and absorption) [18, 19]. The goal of subzonal reconstruction in this instance is the preservation of subzonal temperature variation through an overlay remap

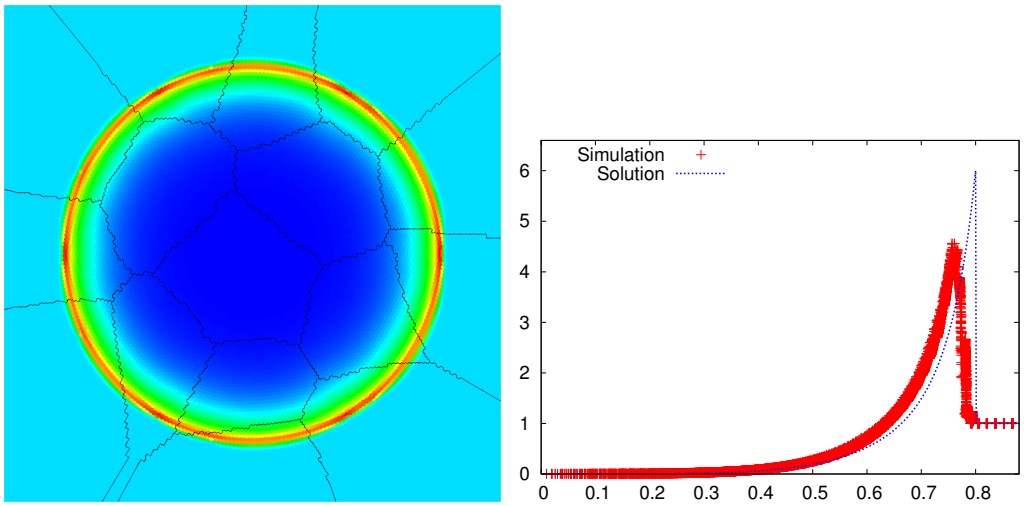

Figure 24: Sedov problem on radial Voronoi mesh with $N_r = 64$. Zonal density using flattened reconstruction.

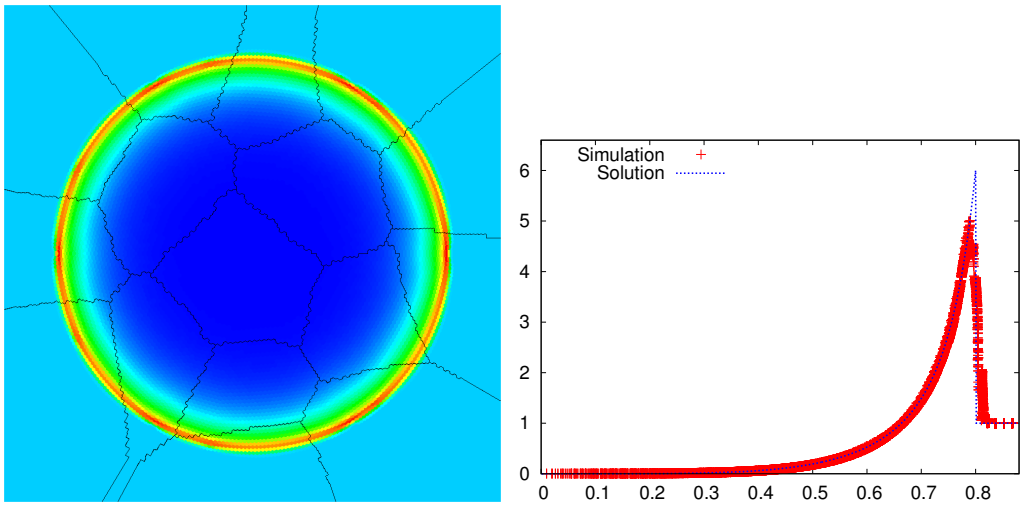

Figure 25: Sedov problem on radial Voronoi mesh with $N_r = 64$. Zonal density using variation-diminishing reconstruction with harmonic average reference.

while maintaining thermal energy conservation.

We define a thermal energy at each subzone $\epsilon_c = V_c C_{V_c} T_c$ using the sub-

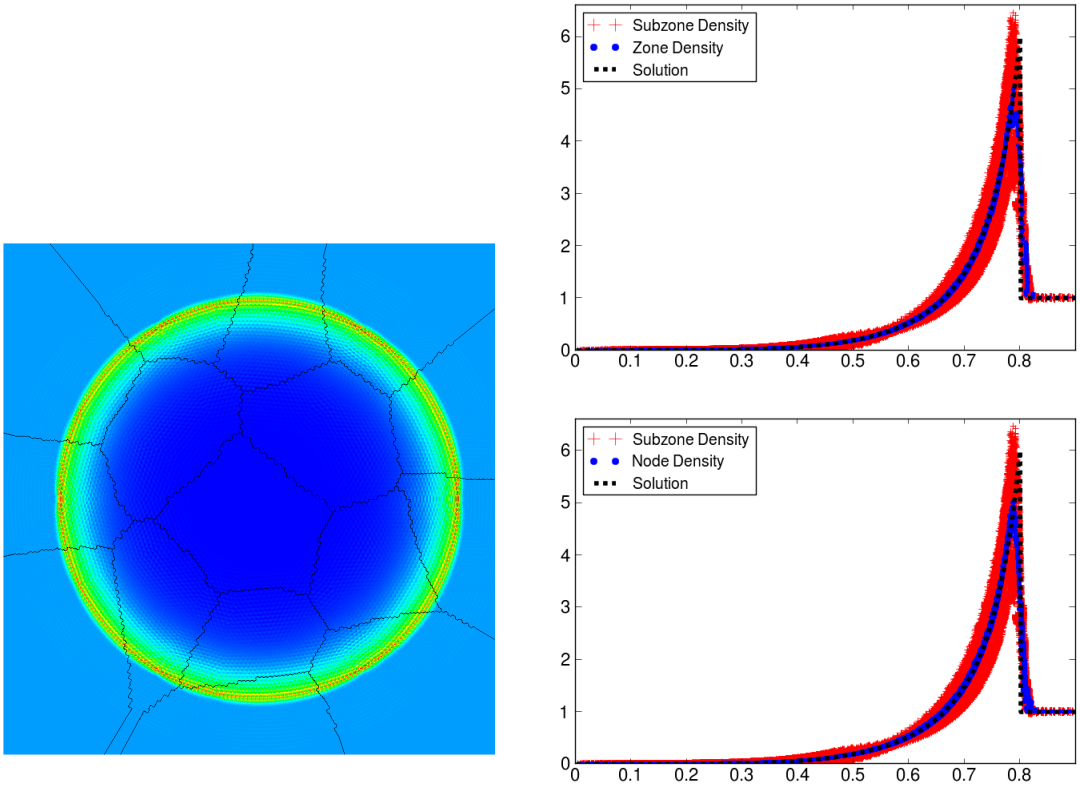

Figure 26: Sedov problem on radial Voronoi mesh with $N_r = 64$. Reconstructed subzonal density with harmonic average reference. Corresponding nodal and zonal densities are overlaid.

zonal temperature T_c and specific heat capacity per unit volume C_{V_c}. Zonal and nodal representations of ϵ are defined by summing around local subzones, then passed through the overlay remap. The reconstruction algorithm that solves (23) proceeds as before, with mass m replaced by thermal energy ϵ. A reference temperature τ_c in each subzone defines the cost function (6) in the constrained optimization.

Subzonal temperature reconstruction is future work due to its strict monotonicity requirements. The staggered compatible hydro algorithm [2] used in this paper is not sensitive to non-monotonicities in the subzonal density field. The same cannot be said for subzonal temperatures in staggered radiation-

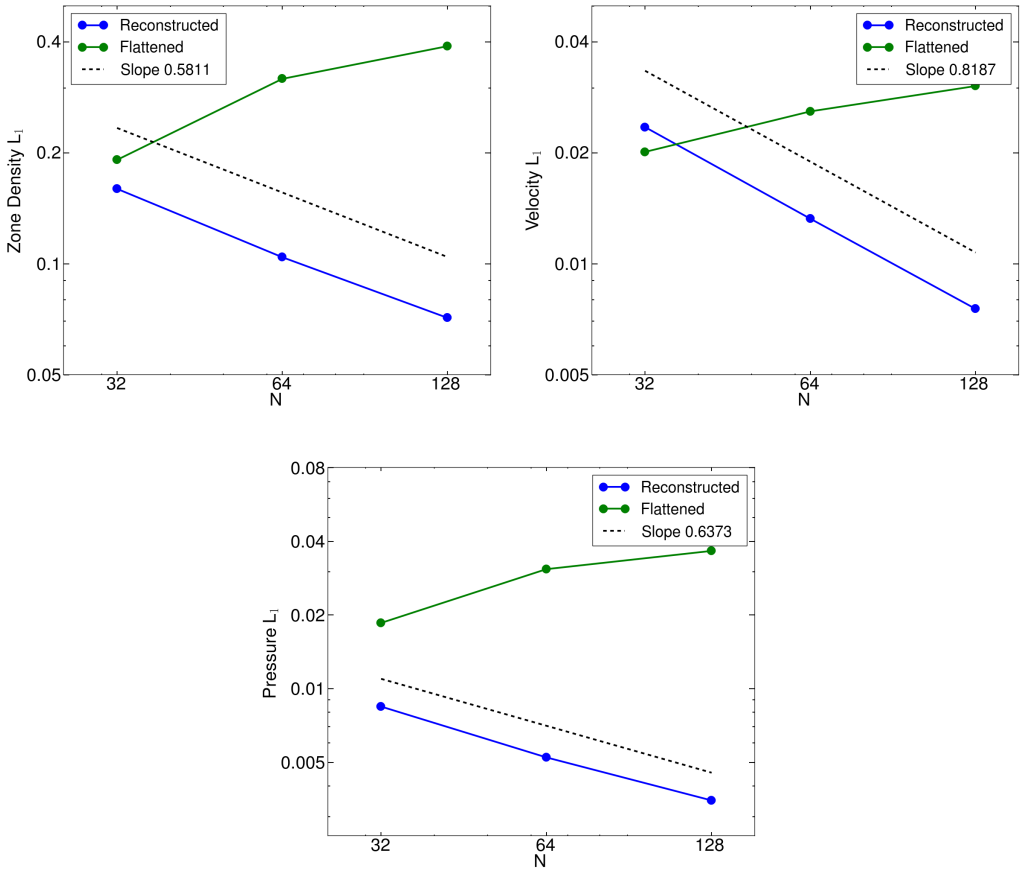

Figure 27: Sedov problem on a radial Voronoi mesh. Plotted are L_1 convergence in zone density, radial velocity, and pressure comparing a flattened reconstruction to a variation-diminishing reconstruction with harmonic average reference.

hydrodynamics; oscillations in the subzonal temperature field can trigger spurious emission/absorption sources and possibly disrupt states of thermodynamic equilibrium.

Future experiments involving subzonal temperature reconstruction will include corrective measures to return monotonicity to the reconstructed values, as in [12], or direct enforcement of inequality constraints during the constrained optimization procedures, as outlined in appendix C.

8. Conclusions

We have introduced a new framework for reconstructing subzonal quantities within a staggered discretization following an overlay-based remap. Examples focus on reconstructing subzonal density to enforce compatibility between zonal and nodal representations of mass. The driving principle of this derivation is maintaining conservation of linear momentum in staggered Lagrangian hydrodynamics. An extension of the framework for reconstructing subzonal temperatures is discussed, with specific application to FLD radiation transport on unstructured grids.

Numerical results are very favorable. In pure remapping tests, we demonstrate the capability of the algorithm to reconstruct smooth and discontinuous fields on both structured and unstructured meshes. In hydrodynamics tests, we demonstrate improved shock capturing due to exact momentum conservation. Convergence in L_1 with respect to mesh resolution is established and meets expectations.

It is possible for the reconstruction algorithm as posed to compute negative subzonal values, which are invalid for mass densities. Test cases with arbitrarily large density discontinuities were shown to produce such negative densities for certain choices of the reference density. The harmonic average outperformed all other references in pure remapping tests due to its robustness to negative reconstructions and its established second-order convergence for smooth fields.

A local procedure has been presented for repairing negative reconstructed states if and when they occur. Local repair is not strictly conservative at nodes; however, the errors are small, carry no definite sign, and do not impact momentum conservation if velocity is constant. We further note that hydrodynamics tests are less likely to require repair than pure remapping problems. The steep discontinuities that trigger negative reconstructions in pure remapping tests are less present in hydrodynamics problems due to numerical diffusion.

Future work will involve extending the algorithm to reconstructing subzonal temperatures and applying it to a series of FLD radiation-hydrodynamics problems. Doing so will require incorporating inequality constraints into the constrained optimization procedure to better control monotonicity during reconstruction. The compatible hydro algorithm [2] used in this article is not sensitive to oscillations in the subzonal mass density field; staggered-grid radiation-hydrodynamics will certainly be sensitive to oscillations in subzonal

temperature. Formulation of a monotonicity-preserving optimization process as a fixed-point quadratic programming problem is presented in appendix C and will be pursued in future installments of this work.

[1] P. Anninos, Kull ALE: II. Grid Motion on Unstructured Arbitrary Polyhedral Meshes, Technical Report: UCRL-ID-147297-PT-2, Lawrence Livermore National Laboratory, 2002.

[2] E. Caramana, D. Burton, M. Shashkov, P. Whalen, The Construction of Compatible Hydrodynamics Algorithms Utilizing Conservation of Total Energy, *Journal of Computational Physics*, 146, pp. 227–262, 1998.

[3] D. Flanagan and Belytschko T. Uniform Strain Hexahedron and Quadrilateral with Orthogonal Hourglass Control, *International Journal on Numerical Methods in Engineering*, 17(5), pp. 679–706, 1981.

[4] J. Gondzio, Multiple Centrality Corrections in a Primal-Dual Method for Linear Programming, *Computational Optimization and Applications*, 6, pp. 137–156, 1996.

[5] J. Grandy, Conservative Remapping and Region Overlays by Intersecting Arbitrary Polyhedra, *Journal of Computational Physics*, 148, pp. 433–466, 1999.

[6] hypre. High performance preconditioners. http://www.llnl.gov/CASC/linearsolvers/

[7] W. Karush, Minima of Functions of Several Variables with Inequalities as Side Constraints, MS Dissertation, Department of Mathematics, University of Chicago, 1939.

[8] H. Kuhn and A. Tucker, Nonlinear Programming, Proceedings of 2nd Berkeley Symposium on Mathematical Statistics and Probability, pp. 481–492, 1951.

[9] R. Liska, M. Shashkov, P. Vachal, B. Wendroff, Optimization-Based Synchronized Flux-Corrected Conservative Interpolation (Remapping) of Mass and Momentum for Arbitrary Lagrangian-Eulerian Methods, *Journal of Computational Physics*, 229, pp. 1467–1497, 2010.

[10] R. Loubere, P.-H. Maire, M. Shashkov, J. Breil, S. Galera, ReALE: A Reconnection-Based Arbitrary Lagrangian-Eulerian Method, *Journal of Computational Physics*, 229, pp. 4724–4761, 2010.

[11] R. Loubere and M. Shashkov, A Subcell Remapping Method on Staggered Polygonal Grids for Arbitrary Lagrangian-Eulerian Methods, *Journal of Computational Physics*, 209, pp. 105–138, 2005.

[12] L. Margolin and M. Shashkov, Remapping Recovery and Repair on a Staggered Grid, *Comput. Methods Appl. Engrg.*, 193, pp. 4139–4155, 2004.

[13] S. Mehrotra, On the Implementation of a Primal-Dual Interior Point Method, *SIAM Journal on Optimization*, 2, pp. 576–601, 1992.

[14] J. M. Owen, M. Shashkov, Arbitrary Lagrangian Eulerian Remap Treatments Consistent with Staggered Compatible Total Energy Conserving Lagrangian Methods, *Journal of Computational Physics*, 273, pp. 520–547, 2014.

[15] L. Sedov, Propagation of Strong Shock Waves, *Journal of Applied Mathematics and Mechanics*, v. 10, pp. 241–250, 1946.

[16] L. Sedov, Similarity and Dimensional Methods in Mechanics, Academic Press, New York, 1959.

[17] J Lockwood-Taylor, An Exact Solution of the Spherical Blast Wave Problem, *Philosophical Magazine*, v. 46, pp. 317–320, 1955.

[18] M. Adams, Subcell Balance Methods for Radiative Transfer on Arbitrary Grids, *Transport Theory and Statistical Physics*, 26, pp. 385–431, 1997.

[19] T. Eaton, M. Adams, A New Corner-Balance/Linear-Discontinuous Method for Transport in Slab Geometry, *Transactions of the American Nuclear Society*, 70, 158, 1994.

[20] T. Kaiser, Evolution of a Zone-Centered Field with a Node-Centered Diffusion Solver on an Unstructured Mesh, American Physical Society, 44th Annual Meeting of the Division of Plasma, abstract #BP1.109, 2002.

[21] T. Palmer, Discretization of the Diffusion Equation on Unstructured Polygonal Meshes in Two Dimensions, *Annals of Nuclear Energy*, 28, pp. 1851, 2001.

[22] T. Bailey, M. Adams, B. Yang, M. Zika, A Piecewise Linear Finite Element Discretization of the Diffusion Equation for Arbitrary Polyhedral Grids, Mathematics and Computation, Supercomputing, Reactor Physics and Nuclear and Biological Applications Conference (ANS), 2005.

Prepared by LLNL under Contract DE-AC52-07NA27344.

A. Analysis of the Singular Matrix

Consider the linear system (23). The matrix $\mathbf{\Psi}$ may be written

$$\mathbf{\Psi} = \mathbf{\Psi}_{cn}^T \Big(\mathbf{I} - \mathbf{\Psi}_{cz} (\mathbf{\Psi}_{cz}^T \mathbf{\Psi}_{cz})^{-1} \mathbf{\Psi}_{cz}^T \Big) \mathbf{\Psi}_{cn}. \tag{A.1}$$

Without loss of generality, we have replaced the diagonal matrix of positive reference values $\mathbf{M}_c$ with the identity matrix. We wish to explore the null space of $\mathbf{\Psi}$ through its sub-matrices. Let

$$\mathbf{P} = \mathbf{\Psi}_{cz} (\mathbf{\Psi}_{cz}^T \mathbf{\Psi}_{cz})^{-1} \mathbf{\Psi}_{cz}^T. \tag{A.2}$$

It is easy to check that the matrix $\mathbf{P}$ is a projector (i.e. $\mathbf{P}^2 = \mathbf{P}$). The form of (A.2) implies that $\mathbf{P}$ is an orthogonal projector onto the range of matrix $\mathbf{\Psi}_{cz}$. In addition,

$$\mathbf{Q} = \mathbf{I} - \mathbf{P} \tag{A.3}$$

is a projector orthogonal to $\mathbf{P}$: the range of $\mathbf{Q}$ is the null space of $\mathbf{P}$ and vice versa.

Because each corner c corresponds to a unique zone z and a unique node n, each row of $\mathbf{\Psi}_{cz}$ and $\mathbf{\Psi}_{cn}$ has exactly one nonzero element, implying the two matrices are full-rank. Therefore, the null space of

$$\mathbf{\Psi} = \mathbf{\Psi}_{cn}^T \mathbf{Q} \mathbf{\Psi}_{cn}$$

is equal in size to the null space of $\mathbf{Q}$. It remains to find the null space of $\mathbf{Q}$.

Let $\mathbf{x} \in \text{null}(\mathbf{\Psi})$, then there exists some $\mathbf{w} \in \text{range}(\mathbf{\Psi}_{cn})$ such that $\mathbf{w} = \mathbf{\Psi}_{cn}\mathbf{x}$ and $\mathbf{Q}\mathbf{w} = 0$. We note

(*i*) $\mathbf{Q}$ and $\mathbf{P}$ are orthogonal: $\mathbf{w} \in \text{null}(\mathbf{Q}) \implies \mathbf{w} \in \text{range}(\mathbf{P})$.

(*ii*) $\mathbf{P}$ projects onto the range of $\mathbf{\Psi}_{cz}$.

Using (*i*) and (*ii*), there must exist $\mathbf{y} \in \text{range}(\mathbf{\Psi}_{cz})$ such that $\mathbf{w} = \mathbf{\Psi}_{cz}\mathbf{y}$. Finally, for $\mathbf{x}$ to be in the null space of $\mathbf{\Psi}$, there must exist $\mathbf{y}$ such that

$$\mathbf{\Psi}_{cz}\mathbf{y} = \mathbf{\Psi}_{cn}\mathbf{x}. \tag{A.4}$$

We wish to show that the space spanned by $\mathbf{x}$ satisfying (A.4) is of dimension 1. We propose this space consists of the set of constant vectors. Geometrically, this is plausible: the action of $\mathbf{\Psi}_{cn}$ takes node-centered data

and projects it to its corners (weighted by corner volume). Similarly for $\mathbf{\Psi}_{cz}$. Each corner has a unique zone and a unique node. If node-centered data $\mathbf{x}$ and zone-centered data $\mathbf{y}$ satisfy (A.4), the data must be constant.

Proof: Suppose not. Then there exists some non-constant node-centered data $\mathbf{x}$ and zone- centered data $\mathbf{y}$ satisfying (A.4). The i-th elements of $(\mathbf{\Psi}_{cn}\mathbf{x})$ and $(\mathbf{\Psi}_{cz}\mathbf{y})$ correspond to corner c_i:

$$(\mathbf{\Psi}_{cn}\mathbf{x})_i = V_{c_i} x_{n_{c_i}} \qquad \text{and} \qquad (\mathbf{\Psi}_{cz}\mathbf{x})_i = V_{c_i} y_{z_{c_i}}.$$

For (A.4) to hold, $x_{n_c} = y_{z_c}$ at every corner c. In other words, the values y_z for zones around node n must be equal to x_n, and the values x_n for nodes around zone z must be equal to y_z. For this to hold, both x_n and y_z must take on the same value at every node and zone respectively. Therefore, only constant data may satisfy $\mathbf{x} \in \text{null}(\mathbf{\Psi})$ and (A.4).

B. Global Repair of Negative Reconstructed Densities

We generalize the optimization problem in section 2 to include additional constraints at subzones. More precisely, we augment (7) with equality constraints at a subset of subzones:

$$\begin{aligned}
&\text{Given:} && \rho_z,\ \rho_n,\ V_c \\
&\text{Minimize:} && \frac{1}{2}\sum_c \left(\frac{\rho_c - \mu_c}{\mu_c}\right)^2 \\
&\text{Subject to:} && \begin{cases} \displaystyle\sum_{c(z)} V_c(\rho_c - \rho_z), & \forall z \\ \displaystyle\sum_{c(n)} V_c(\rho_c - \rho_n), & \forall n \\ \rho_c = \epsilon_c \quad \forall c \in \mathcal{C} \end{cases}
\end{aligned} \tag{B.1}$$

Where ϵ_c is the constrained value of ρ_c, and $\mathcal{C}$ may be any subset of subzones. The system remains under-constrained provided $N_z + N_n + |\mathcal{C}| < N_c$.

We interpret the solution to (B.1) in terms of a repair procedure to correct unsatisfactory values in the solution to (7) (i.e. negatives, non-monotonicities). Let $\rho_c^{(0)}$ denote the solution to the original reconstruction problem under some

arbitrary reference state. We solve (B.1) to obtain the correction $\rho_c^{(1)}$, where the reference state is given by the previous solution $\rho_c^{(0)}$. (The advantage of choosing the previous solution as the reference state ensures that the reference automatically satisfies conservation.)

The set $\mathcal{C}$ consists of a collection of subzones sharing some undesirable property we wish to repair. For positivity constraints, one may consider the subset

$$\mathcal{C} = \mathcal{C}_-^{(0)} = \{c \ : \ \rho_c^{(0)} \leq 0\}. \tag{B.2}$$

For a stricter constraint such as monotonicity, one may consider

$$\mathcal{C} = \mathcal{C}_{mono}^{(0)} = \left\{c \ : \ \rho_c^{(0)} \notin \left[\min(\rho_{z_c}, \rho_{n_c}) \ , \ \max(\rho_{z_c}, \rho_{n_c})\right]\right\}. \tag{B.3}$$

For the purposes of this paper, we limit our focus to (B.2) and consider monotonicity constraints to be a source of future work.

Finally, we have freedom to choose the set of constrained subzonal values $\{\epsilon_c\}_{\mathcal{C}}$. For positivity constraints, we select a local positive value for each subzone. One option is to use

$$\epsilon_c = \alpha \min(\rho_{z_c}, \rho_{n_c}), \tag{B.4}$$

where $0 < \alpha \leq 1$. Note that as α approaches 1, the positivity constraint becomes stricter.

The method of Lagrange multipliers solves (B.1). The Lagrangian is

$$\begin{aligned} L(\rho_c, \lambda_z, \lambda_n, \lambda_c) \ = \ \frac{1}{2}\sum_c \left(\frac{\rho_c - \mu_c}{\mu_c}\right)^2 \ &+ \ \sum_z \lambda_z \sum_{c(z)} V_c(\rho_c - \rho_z) \\ &+ \ \sum_n \lambda_n \sum_{c(n)} V_c(\rho_c - \rho_n) \ , \\ &+ \ \sum_c \lambda_c \chi_c V_c(\rho_c - \epsilon_c) \end{aligned} \tag{B.5}$$

where λ_c denotes the Lagrange multipliers for the subzonal constraints and χ_c is an indicator function that is 1 for $c \in \mathcal{C}$ and 0 otherwise. Formulas (11) and (15) for the nodal and zonal Lagrange multipliers hold, with w_c, Δm_z,

and Δm_n replaced by

$$\begin{aligned} w_c &= (\mu_c V_c)^2 (1-\chi_c) \\ \Delta m_z &= m_z - \sum_{c(z)} \Big((1-\chi_c)\mu_c + \chi_c \epsilon_c \Big) V_c \\ \Delta m_n &= m_n - \sum_{c(n)} \Big((1-\chi_c)\mu_c + \chi_c \epsilon_c \Big) V_c \end{aligned} \quad . \tag{B.6}$$

The linear system (23) is modified

$$\mathbf{\Psi} = \mathbf{\Psi}_{cn}^T \mathbf{M}_c (\mathbf{I}_c - \mathbf{X}_c) \Big(\mathbf{I}_c - \mathbf{M}_c \mathbf{\Psi}_{cz} \mathbf{W}_z^{-1} \mathbf{\Psi}_{cz}^T \mathbf{M}_c \Big) (\mathbf{I}_c - \mathbf{X}_c) \mathbf{M}_c \mathbf{\Psi}_{cn}, \tag{B.7}$$

where $\mathbf{X}_c = \mathrm{diag}(\chi_c)$. The updated subzonal densities are

$$\rho_c = (1-\chi_c)\mu_c + \chi_c \epsilon_c - V_c \mu_c^2 (1-\chi_c) \left(\lambda_{z_c} + \lambda_{n_c} \right). \tag{B.8}$$

Remark: For problems with steep discontinuities, it is possible that a single global repair step is insufficient to remove all negative subzonal states. When this occurs, an iterative repair may be performed. Given state $\rho_c^{(i)}$, we solve (B.1) for new subzonal densities $\rho_c = \rho_c^{(i+1)}$ using reference $\mu_c = \rho_c^{(i)}$ and subzone subset

$$\mathcal{C} = \mathcal{C}_-^{(i)} = \{ c \ : \ \rho_c^{(j)} \leq 0, \ \forall j \leq i \}. \tag{B.9}$$

We define $\mathcal{C}_-^{(i)}$ such that, once a subzone is constrained, it is constrained for all iterations to ensure convergence of the iterative repair.

C. Optimization with Equality and Inequality Constraints: Quadratic Programming Formulation

The optimization problem (7) may be expressed as a quadratic programming (QP) problem:

$$\begin{aligned} &\text{minimize:} \quad f(\mathbf{x}) = \frac{1}{2}\mathbf{x}^T Q \mathbf{x} + \mathbf{c}^T \mathbf{x} \\ &\text{subject to:} \quad A\mathbf{x} = \mathbf{b} \quad \text{and} \quad C\mathbf{x} \geq \mathbf{d} \end{aligned} \tag{C.1}$$

where Q is a symmetric, positive-definite matrix, $\mathbf{c}$ is a constant vector, and A and C are full-rank matrices comprising the equality and inequality constraints respectively on the target variable $\mathbf{x}$.

We follow the Karush-Kuhn-Tucker formalism [7, 8] for optimization with mixed constraints. We define Lagrange multipliers λ and γ corresponding to the equality and inequality constraints respectively. An extra variable $\mathbf{s}$ enforces the inequality constraints through complementary slackness.

We solve for fixed points of the Lagrangian

$$L(\mathbf{x};\lambda,\gamma,\mathbf{s}) = \frac{1}{2}\mathbf{x}^T Q\mathbf{x} + \mathbf{c}^T\mathbf{x} + \lambda^T(A\mathbf{x}-\mathbf{b}) + \gamma^T(C\mathbf{x}-\mathbf{s}-\mathbf{d}). \tag{C.2}$$

satisfying complementarity condition

$$\gamma_i s_i = 0, \quad \gamma_i \geq 0, \quad s_i \geq 0 \tag{C.3}$$

for each inequality constraint i. Solving for the fixed points of (C.2) corresponds to solving the following coupled systems:

$$\mathbf{c} + Q\mathbf{x} - A^T\lambda - C^T\gamma = 0 \tag{C.4a}$$

$$A\mathbf{x} - \mathbf{b} = 0 \tag{C.4b}$$

$$C\mathbf{x} - \mathbf{s} - \mathbf{d} = 0 \tag{C.4c}$$

$$S\Gamma\mathbb{1}_e = 0 \tag{C.4d}$$

$$\mathbf{s},\gamma \geq 0 \tag{C.4e}$$

where $S = \mathrm{diag}(\mathbf{s})$, $\Gamma = \mathrm{diag}(\gamma)$, and $\mathbb{1}_e$ denotes a vector of ones.

For the problem of subzonal reconstruction, we have convenient forms for the coefficients of (C.4). The target variable is the relative variation with respect to reference μ per subzone:

$$x_c = \frac{\rho_c - \mu_c}{\mu_c}, \quad c = 1, ..., N_c \tag{C.5}$$

The objective function simplifies to

$$Q = I_c \quad \text{and} \quad \mathbf{c} = 0.$$

Zonal (4) and nodal (5) equality constraints are written compactly

$$\lambda = \begin{bmatrix} \lambda_z \\ \lambda_n \end{bmatrix}, \qquad A = \begin{bmatrix} \mathbf{\Psi}_{cz}^T\mathbf{M}_c \\ \mathbf{\Psi}_{cn}^T\mathbf{M}_c \end{bmatrix}, \qquad \mathbf{b} = \begin{bmatrix} \mathbf{m}_z - \mathbf{\Psi}_{cz}^T\mu \\ \mathbf{m}_n - \mathbf{\Psi}_{cn}^T\mu \end{bmatrix}, \tag{C.6}$$

where λ is length $N_z + N_n$.

The inequality constraints may be used to enforce bounds on the reconstructed variables. Enforcing one-sided bounds for $\{\rho_c\}_c$ of the form $\rho_c \geq \rho_c^{min}$ corresponds to N_c inequality constraints of the form

$$C = I_c, \qquad \mathbf{d} = \mathbf{M}_c^{-1}(\rho^{min} - \mu). \tag{C.7}$$

Likewise, two-sided bounds of the form $\rho_c^{min} \leq \rho_c \leq \rho_c^{max}$ corresponds to $2N_c$ inequality constraints of the form

$$\gamma = \begin{bmatrix} \gamma^{min} \\ \gamma^{max} \end{bmatrix}, \qquad C = \begin{bmatrix} I_c \\ -I_c \end{bmatrix}, \qquad \mathbf{d} = \begin{bmatrix} \mathbf{M}_c^{-1}(\rho^{min} - \mu) \\ -\mathbf{M}_c^{-1}(\rho^{max} - \mu) \end{bmatrix}. \tag{C.8}$$

Enforcing (C.7) is useful for eliminating negative subzonal reconstruction values. Enforcing the stricter constraints in (C.8) may be used to preserve monotonicity locally.

An iterative interior-point method may be used to solve (C.1). Given initial state $\mathbf{v}^{(i)}$ for $\mathbf{v} \in \{\mathbf{x}, \lambda, \gamma, \mathbf{s}\}$ satisfying $\gamma_i^{(i)} > 0$ and $s_i^{(i)} > 0$, we solve for step lengths $\Delta\mathbf{v}$ such that $\mathbf{v} + \Delta\mathbf{v}$ solves (C.4a)– (C.4c) and satisfies (C.4e):

$$\begin{bmatrix} Q & -A^T & -C^T & 0 \\ A & 0 & 0 & 0 \\ C & 0 & 0 & -I \\ 0 & 0 & S & \Gamma \end{bmatrix} \begin{bmatrix} \Delta\mathbf{x} \\ \Delta\lambda \\ \Delta\gamma \\ \Delta\mathbf{s} \end{bmatrix} = - \begin{bmatrix} r_{var} \\ r_{eq} \\ r_{ineq} \\ r_{slack} \end{bmatrix}. \tag{C.9}$$

Using current data $\mathbf{v}^{(i)}$, the residuals are computed

$$r_{var} = Q\mathbf{x}^{(i)} + \mathbf{c} - A^T\lambda^{(i)} - C^T\gamma^{(i)} \tag{C.10a}$$

$$r_{eq} = A\mathbf{x}^{(i)} - \mathbf{b} \tag{C.10b}$$

$$r_{ineq} = C\mathbf{x}^{(i)} - \mathbf{s}^{(i)} - \mathbf{d} \tag{C.10c}$$

$$r_{slack} = S^{(i)}\Gamma^{(i)}\mathbb{1}_e \tag{C.10d}$$

corresponding to fixed points of the Lagrangian in the solution variable, equality multipliers, inequality multipliers, and slackness condition respectively. Step lengths are limited by positive factor $\alpha^{(i)}$ such that complementarity holds:

$$0 < \alpha^{(i)} \leq \max_{\alpha} \left\{ \gamma^{(i)} + \alpha\Delta\gamma^{(i)} > 0 \ , \ \mathbf{s}^{(i)} + \alpha\Delta\mathbf{s}^{(i)} > 0 \right\} \tag{C.11}$$

The state update is then $\mathbf{v}^{(i+1)} = \mathbf{v}^{(i)} + \alpha^{(i)}\Delta\mathbf{v}^{(i)}$. Iteration continues until the residual vector falls below a given tolerance.

Remark: The system (C.9) must be solved each iteration with new coefficients and residuals. This is the most computationally intensive part of the fixed-point method. Efficient versions, such as [4, 13], incorporate multiple sub-steps by applying higher-order corrections to the minimization path. The coefficient matrix remains fixed during the entire step; it is typically factored and applied to a different residual for each sub-step. Significant savings in numbers of iterations to convergence have been demonstrated.

Remark: System (C.9) may be factored to produce a single linear system of size N_z for the zonal equality Lagrange multiplier step lengths. The factored system consists of an $N_z \times N_z$ sparse matrix identical in structure to (23) described in section 2. The step lengths of the remaining state variables are obtained through simple linear transformations.