

Lawrence Berkeley National Laboratory

LBL Publications

Title

Adaptive clipping-and-redistribution algorithms for bounded and conservative high-order interpolations applied to discontinuous and reactive flows

Permalink

<https://escholarship.org/uc/item/2zq1s63s>

Journal

International Journal for Numerical Methods in Fluids, 95(5)

ISSN

0271-2091

Authors

Overton-Katz, Nathaniel
Gao, Xinfeng
Johansen, Hans
et al.

Publication Date

2023-05-01

DOI

10.1002/flid.5165

Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

Adaptive Clipping-and-Redistribution Algorithms for Bounded and Conservative High-Order Interpolations Applied to Discontinuous and Reactive Flows

Nathaniel Overton-Katz¹ | Xinfeng Gao¹ | Hans Johansen² | Stephen M. Guzik¹

¹Computational Fluid Dynamics and Propulsion Laboratory, Colorado State University, Fort Collins, CO, 80523, USA

²Applied Numerical Algorithms Group, Lawrence Berkeley National Laboratory, Berkeley, CA, USA

Correspondence

*Nathaniel Overton-Katz, Email: noverton@colostate.edu

Summary

A new adaptive clipping-and-redistribution method is presented which provides bounds-preservation for multi-dimensional interpolation in the context of high-order finite-volume discretizations with adaptive mesh refinement (AMR). The underlying finite-volume method (FVM) for the computational fluid dynamics applications is fourth-order accurate for smooth solutions and utilizes AMR for computational efficiency in solving multiscale problems involving turbulence and combustion. High-order interpolation between different AMR levels is required. However, this operation often leads to numerical issues because combustion species must have physical bounds preserved. The present study overcomes two major challenges in the development of the high-order interpolation method. First, the method needs to be bound-preserving near extrema or discontinuities to prevent the emergence of unphysical oscillations while maintaining fourth-order accuracy in smooth flows. Second, the method needs to satisfy the conservation requirement in multiple dimensions, particularly in the context of curvilinear coordinate transformations. Additionally, the method is designed to be localized and computationally inexpensive. The new interpolation scheme is demonstrated by solving reacting flows, which are extremely sensitive to unphysical overshoots in conserved quantities. The test problems are shock-induced H_2-O_2 combustion and a C_3H_8 -air flame in a practical bluff-body combustor. Results show the method prevents new extrema near discontinuities while maintaining high-order accuracy in smooth regions. In particular, the method is extremely beneficial for combustion with stiff chemistry. With the proposed new method, even if flame fronts cross AMR interfaces or new grids are created in the vicinity of the flame, solution stability is retained.

KEYWORDS:

Adaptive Clipping-and-Redistribution, Adaptive Mesh Refinement, High-Order Finite-Volume Method, Computational Combustion

NOTATION

The notation is generally explained as it is introduced. Bold type is reserved for real physical vectors, such as solution variables of a system of the governing equations. However, a bold symbol can also represent vectors in integer space. Matrices are named with capital letters, such as \mathbf{N} , and the entries are named with subscripts, $N_{i,j}$, where i and j are the indices for the row and the column, respectively. The vector symbol \vec{x} is used for vectors (or column matrices) depending on space dimensions. Some symbols are listed and defined here for convenience.

D	The spatial dimension
\mathcal{V}_i	The volume of cell i
$\vec{\mathbf{F}}$	flux dyad, e.g., $[\mathbf{F}_1, \dots, \mathbf{F}_D]$
\mathbf{F}	flux vector, e.g., $[F_1, \dots, F_N]$
\mathbf{U}	solution vector, e.g., $[U_1, \dots, U_N]$
F_j	the j^{th} component of vector \mathbf{F}
i	grid indices, e.g., (i, j, k) in 3D
e^d	unit vector in direction d
$\langle \cdot \rangle$	the cell-averaged or face-averaged quantity
\vec{x}	physical space, e.g., (x, y, z)
x_d	the d^{th} component of \vec{x}
$\vec{\xi}$	computational space, e.g., (ξ, η, ζ)
ξ_d	the d^{th} component of $\vec{\xi}$
h	grid spacing
\mathcal{C}	grid coarsening operator
\mathcal{C}^{-1}	grid refinement operator
ℓ	grid refinement level
n_{ref}	grid refinement factor

1 | INTRODUCTION

Numerical strategies for finite-volume methods (FVMs) such as high-order schemes and adaptive mesh refinement (AMR) on structured grids have proven themselves valuable for improving computational efficiency¹. However, these methods are not without challenges, especially when solving combustion physics. The purpose of AMR is to better resolve regions of interest in a solution in order to reduce the discretization error. AMR requires interpolation in *both space and time* at the boundaries of grid refinements, to enable local time stepping at each resolution. When the underlying finite-volume method is high-order accurate for smooth flows, as in the present study, oscillations can appear near extrema or discontinuities due to the use of a *high-order conservative interpolation scheme*. These overshoots are particularly troublesome when they become unphysical, and ideally AMR would *tag* (mark locations for refinement) discontinuous features before they form. When solving compressible reactive flows, shock waves and flame fronts are features of interest. A high-order interpolant near or crossing these features often brings about oscillations, resulting in undesirable characteristics such as negative mass fractions, or such that the sum of all species' mass fractions are greater than unity in some localized region. Eventually, these unphysical phenomena can destroy the solution.

Two types of interpolation are required for the finite-volume AMR strategy discussed herein. Interpolation is used to fill ghost cells surrounding a patch of fine grid cells, and interpolation is used to define new fine grids as the solution evolves. For both cases, standard practice is to adjust the fine grids so that discontinuities always remain well away from refinement boundaries, and the interpolations are of smooth solution profiles. Having a discontinuous feature encroach on a grid refinement boundary is avoided. However, for reasons of computational cost, this is not always practical. In many simulation strategies, regions are refined based on criteria other than following a discontinuity. For example, if reproducing an experiment, one may add fixed levels of refinement that match a physical observation window to focus computational effort, and minimize error, in that window. Strategies such as embedded-large-eddy-simulation (embedded-LES) and embedded-direct-numerical-simulation (embedded-DNS), to be discussed shortly, have similar requirements. For some problems, such as combustion described by highly sensitive reaction mechanisms, tiny overshoots and undershoots in bounded species mass fractions may be enough to corrupt the solution,

in regions well away from the identified flame discontinuity. Although it is not the preferred way to apply AMR, there are strong motivations for robustly and accurately interpolating in the vicinity of discontinuities. The goal of this work is to create high-order accurate AMR methods that are more robust to tagging and refinement variability, especially when AMR interfaces are near regions with strong gradients.

To demonstrate the problem specifically, an example of high-order interpolation is shown in Fig. 1 where significant erroneous features appear due to mesh refinement in patches highlighted by violet bubbles. The figure shows the distribution of CH_2O species mass fraction from combustion of a C_3H_8 -air mixture. Further details for this case are explained in Section 5.2, but here we mainly point out the oscillations produced by the high-order interpolation. These seemingly small noises, resulting from mesh refinement, have been shown to grow and interact with the flow dynamics, eventually leading to numerical instability. This is observed to be especially problematic for chemically reacting flows where small differences in mass fractions are able to push the chemical reaction to unphysical states that lead to solution failure due to the stiff nature of the reactions. Therefore, for combustion with stiff chemistry, this necessitates a high-order interpolation approach that is not only bound-preserving but also maintains conservation while providing a high-order accurate solution in smooth regions. By contrast, for calorically perfect gases, even at high Mach numbers, it is the authors' experience that bounded interpolation is not necessary as long as discontinuities are contained in fine grids.

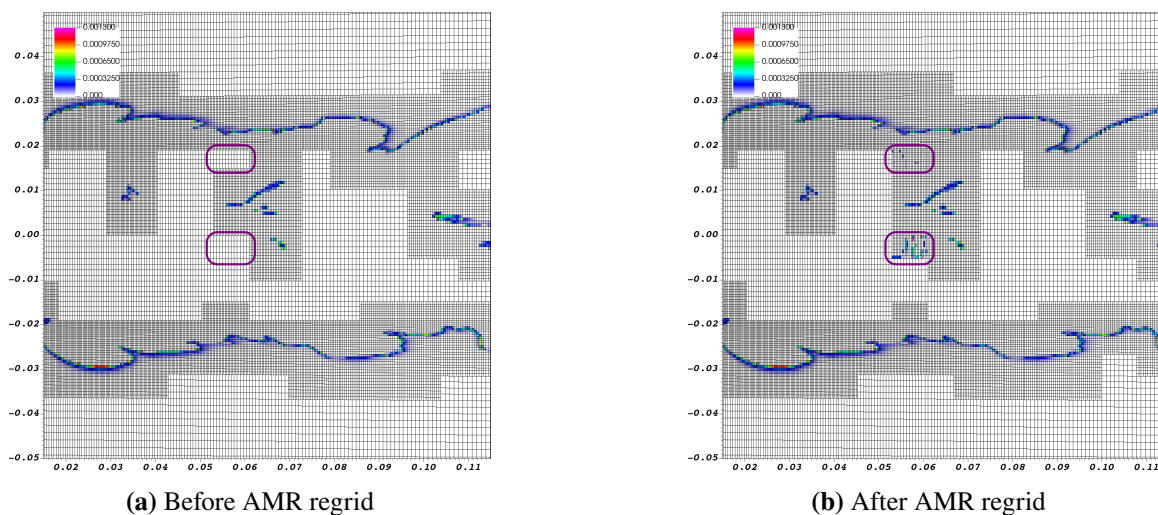


FIGURE 1 Flame front in the bluff body combustion case, shown by mass fractions of species CH_2O . The grid is adaptively refined to capture the flame front, and regrided between the left and right images. The addition of new fine grid cells introduces overshoots which significantly disrupt the flow. The regions with overshoot are indicated by violet bubbles.

Another example is controlling oscillations that may appear from high-order interpolation when studying turbulent flow applications utilizing embedded-LES and embedded-DNS methods². In these approaches, only a small subset of the domain has sufficient resolution for a detailed query of the physics of interest using LES or DNS modeling of turbulence. Elsewhere, the domain can be considered as approximating the influence of realistic geometries and boundary conditions using far more affordable methods. With additional control of error, e.g., via data assimilation, it is a way forward when full LES or DNS is otherwise completely infeasible due to the computational expense. Since embedded-DNS and embedded-LES cases can only afford limited amounts of mesh refinement, discontinuous flow features often must span AMR interfaces. This is illustrated by Fig. 2, where only a small window of a flame front is resolved using DNS. In this example, the discontinuous flame front must cross multiple different levels of grid refinement. To reduce computational time, there is also interest in allowing the flame to only develop on the coarse level for an amount of time before refining and using it as an initial condition for the embedded-DNS region. For these applications, it is important to have an algorithm which is able to mitigate the oscillations present in high-order interpolation to prevent unwanted numerical instabilities due to abrupt AMR-driven changes in grid resolution.

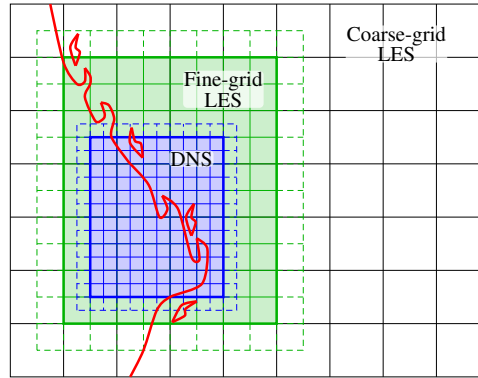


FIGURE 2 Illustration of embedded-DNS. Due to computational limitations, only a small window of DNS spanning a flame front is affordable. In order to develop the large scale flow features of a flame, multiple levels of grid refinement must be used. This results in discontinuous features spanning different refinement levels, which can destabilize the solution.

In summary, we have devised a new high-order bound-preserving interpolation scheme, which improves the robustness of AMR for high-order FVMs when modeling combusting flows. The new high-order adaptive clipping-and-redistribution (HO-ACR) method has been implemented, verified, and validated with our in-house CFD code Chord^{3,1,4,5,6,7,8,9,10}. Chord employs fourth-order finite-volume discretizations in space and fourth-order Runge-Kutta methods (both fully explicit and additive implicit-explicit) for time integration. AMR in both space and time is key to enable computational efficiency for multiscale problems. Mapped structured grids (e.g., transforming a complex physical geometric domain to a Cartesian computational domain) are used to accommodate the geometric complexity and utilize AMR. Chord is capable of solving a wide range of compressible fluid dynamics with combustion and turbulence including Reynolds-averaged Navier-Stokes and large eddy simulation for turbulence. Chord is built upon the Chombo framework¹¹ and implemented using flat MPI for parallel computing on high performance computers with 10^5 or more cores.

1.1 | Existing Methods to Cope with Unphysical Numerical Solutions

Traditional CFD approaches for managing interpolation overshoots often utilize a measure to detect discontinuities and locally revert to low-order methods in their presence¹². There is a wealth of well-developed methods on this topic in the form of flux limiters¹³ and the family of ENO schemes¹⁴.

Common methods for bounds preservation in CFD schemes exist in the form of limiters. The PPM limiter¹³ is well known for its application in high-order schemes, and relative simplicity. This method operates by constructing high-order approximations at cell faces, while ensuring these constructions do not generate new extrema and are monotonic. Furthermore, solution smoothness is accounted for in order to prevent excessive limiting and reduction in the order of accuracy in regions of smooth extrema. Although well-developed, in general limiters are only designed for face reconstructions and most often applied in one dimension at a time. Approaches for limiters in multiple dimensions do exist¹⁵, but are not easily modified to achieve strict bound-preserving cell-based interpolation.

In a similar spirit to limiters, the flux corrected transport (FCT) methods utilize low-order solutions to provide bounds for high-ordered flux reconstructions. Unlike limiters, FCT methods have been extended to operate in multiple dimensions. The FCT-like maximum preserving principle (MMP) in the work by Anderson et al.¹⁶ enforces the flux correction by instead adjusting Discontinuous Galerkin elements. This adjustment requires satisfying flux constructions, while redistributing and preserving element mass. Despite the approach having many valuable properties, the aims are subtly different from ours and can not be directly translated into a FVM scheme.

One approach that is multi-dimensional, conservative, and solves the spatial interpolation problem we face is the CENO method described by Ivan et al.¹⁷. This approach formulates a smoothness measure to identify regions which are non-smooth. In non-smooth regions a low-order interpolation is used, while a high-order interpolation is employed elsewhere. Although it is an effective approach, it requires a tuned coefficient for smoothness detection. Additionally, large stencils of data can be required which necessitate expensive data communication, but further developments of this approach are able to reduce stencil sizes by more frequently moving smaller amounts of information¹⁸. Implementing this interpolation approach other than within

the CENO scheme would be non-trivial since it requires careful consideration of limiter choice and significant changes to communication patterns.

Another approach for the multi-dimensional, conservative interpolation problem is the clipping-and-redistribution algorithm described by Hilditch et al.¹⁹. After a high-order interpolation, any solution outside physically imposed bounds is clipped off. The sum of clipped quantities is tracked and redistributed elsewhere inside the domain to maintain conservation. This approach is developed to prevent physically constrained overshoots based upon predetermined bounds. Additionally, the redistribution happens globally and thus requires solving an implicit system which is potentially expensive. Nevertheless, this work is the inspiration for the algorithm that we present. We seek to adaptively enforce bounds near discontinuities and reduce the cost of redistribution. In addition, the present algorithm operates on mapped quantities since the underlying CFD infrastructure uses mapped structured grids to accommodate complex geometries.

None of the above approaches fully match our required criteria for refining complex combusting flows. Although the interpolation method in the CENO scheme is the nearest to our goals, it is not straightforward to incorporate into our existing code. In the present study, we develop the high-order adaptive clipping-and-redistribution (HO-ACR) method and successfully demonstrate the design properties of the interpolation: conservative, solution-dependent bounds-preservation, efficiency, interpolation without compromising accuracy in smooth regions. Additionally, the HO-ACR algorithm is designed to operate as a correction to a high-order interpolation. This makes it non-invasive and relatively simple to implement independent of the scheme it is used in.

The rest of paper is organized as follows. Section 2 describes the foundational numerical methodology including that of the finite-volume method, AMR, and unbounded interpolations used to fill ghost cells and populate newly refined grid levels. Section 3 presents the details of the newly developed HO-ACR algorithm, including the solution procedures and simple verification tests. In Section 4, the HO-ACR algorithm is implemented in Chord and verified by advection physics. Results and discussions are provided in Section 5, which both validate the methodology and demonstrate the effectiveness for complex flows. Finally, conclusions are drawn in Section 6 and future follow-up work is discussed.

2 | BASE METHODOLOGY

Although not the focus of this paper, the high-order FVM on mapped structured grids is briefly described for convenience since the new HO-ACR method is developed, implemented, and applied in this test bed. A more complete description is available for Cartesian grids in McCorquodale and Colella²⁰ and for mapped grids in Guzik et al.^{1,21}.

2.1 | High-Order Finite-Volume Methods

Consider hyperbolic PDE's of the form

$$\frac{\partial}{\partial t} \mathbf{U} + \vec{\nabla} \cdot \vec{\mathbf{F}}(\mathbf{U}) = 0. \quad (1)$$

FVMs approximate the integral form of the conservation laws with many contiguous control volumes \mathcal{V}_i in physical space, \vec{x} ,

$$\frac{\partial}{\partial t} \int_{\vec{x}(\mathcal{V}_i)} \mathbf{U} d\vec{x} + \int_{\vec{x}(\mathcal{V}_i)} \vec{\nabla}_x \cdot \vec{\mathbf{F}} d\vec{x} = 0. \quad (2)$$

The underlying FVM code employs mapped grids to accommodate geometric complexities and the Cartesian grid in the computational domain also facilitates the use of AMR. The Cartesian grid, with coordinates $\vec{\xi}$, is transformed into physical space, with coordinates \vec{x} by the forward mapping $\vec{x} = X(\vec{\xi})$, as illustrated in Fig. 3. Using the mapping function, geometries of moderate complexity can be solved while maintaining the computational advantages of Cartesian grids including tractability of AMR.

For mapped grids, the integral equations are transformed into computational space, $\vec{\xi}$, by using grid metric terms such that

$$\frac{\partial}{\partial t} \int_{\mathcal{V}_i} J \mathbf{U} d\vec{\xi} + \int_{\mathcal{V}_i} \vec{\nabla}_\xi \cdot (N^I \vec{\mathbf{F}}) d\vec{\xi} = 0, \quad (3)$$




FIGURE 3 A representation of the transformation between computational and physical space.

where the transformation matrix, $\mathbf{N}^T = J \vec{\nabla}_x \vec{\xi}$, describes the grid metrics and the metric Jacobian $J \equiv \det(\vec{\nabla}_x \vec{\xi})$. Cell-averaged quantities are defined as

$$\langle \mathbf{U} \rangle = \frac{1}{V_i} \int_{V_i} \mathbf{U} dV_i, \quad (4)$$

and together with the divergence theorem yield a discretized form

$$\frac{d}{dt} \langle J\mathbf{U} \rangle_i + \frac{1}{h} \sum_{d=1}^D \left(\langle \mathbf{N}_d^T \vec{\mathbf{F}} \rangle_{i+\frac{1}{2}e^d} - \langle \mathbf{N}_d^T \vec{\mathbf{F}} \rangle_{i-\frac{1}{2}e^d} \right) = 0, \quad (5)$$

where the subscript d denotes the d^{th} row of \mathbf{N}^T . The term $\langle \mathbf{N}_d^T \vec{\mathbf{F}} \rangle$ is devised such that free-stream preservation is maintained¹. This formulation is exact until numerical methods are used to approximate fluxes and time derivatives.

2.2 | Adaptive Mesh Refinement with Mapped Grids

Adaptive mesh refinement is effective in solving multi-scale problems. When discontinuities appear in a solution, high-order methods must locally revert to low-order methods around discontinuities. The decrease in order of the scheme is necessary and adversely impacts the solution accuracy. To remedy this, locally refining the grid near the discontinuities with AMR is the best option available. Traditionally, discontinuities are well-embedded inside the fine grid. Layers of buffer cells ensure that interpolations to fill ghost cells and the fine patches are kept away from any discontinuities. However, there are situations in combusting flows and embedded-DNS where it is not practical to fully refine every discontinuity. Building upon this base methodology for AMR, we develop the HO-ACR method in Section 3 to make AMR more robust for non-ideal circumstances. While the details of AMR in Chord can be found in McCorquodale and Colella²⁰ and Guzik et al.¹, relevant portions of the AMR algorithm are described here to facilitate understanding of the proposed HO-ACR method.

Structured grids are defined by integer points $\mathbf{i} \in \mathbb{Z}^D$ which mark cell centers, where D is the number of spatial dimensions. The problem domain to operate in is a subset of the integer lattice $\Gamma \subset \mathbb{Z}^D$. For AMR, a hierarchy of grids $\Gamma^0 \dots \Gamma^{\ell_{\max}}$ exist, each with a different cell size and nested such that $\Gamma^\ell = C_{n_{\text{ref}}}(\Gamma^{\ell+1})$. Fig. 4 illustrates a 3-level AMR grid. The integer n_{ref}^ℓ is the refinement ratio between level ℓ and $\ell + 1$ so that the Cartesian mesh spacing $h^\ell = n_{\text{ref}}^\ell h^{\ell+1}$. The operator C is a coarsening of the grid, and the inverse C^{-1} is the refinement. Adaptive mesh refinement calculations are performed on a hierarchy of nested meshes $\Omega^\ell \subset \Gamma^\ell$, with $\Omega^\ell \supset C_{n_{\text{ref}}}(\Omega^{\ell+1})$. The grid levels are considered as overlapping rather than embedded. At level ℓ all cells inside Ω^ℓ are labeled valid and all cells outside Ω^ℓ are labeled invalid. Valid ghost cells are used to communicate data on a single level split over multiple computational units. Invalid ghost cells are those used to transfer information at the interface between a coarse and fine level. It is assumed that there are a sufficient number of cells on level ℓ separating the level $\ell + 1$ from the level $\ell - 1$ such that interpolations to fill invalid ghost cells on finer levels can be independently performed.

In addition to the refinement in space, AMR can make use of sub-cycling for refinement in time, as detailed by McCorquodale and Colella²⁰. Briefly, sub-cycling updates each grid level using a different time step size to reduce computational costs.




FIGURE 4 A three-level grid with $n_{\text{ref}} = 2$, with each level shown by a different color. Nesting sufficient for one cell to separate level $\ell + 1$ from level $\ell - 1$ exists. One layer of invalid ghost cells are shown on the computational domain using the dashed lines for each level.

Solutions on level $\ell + 1$ use a time step sized as $\Delta t_{\ell+1} = \Delta t_{\ell} / n_{\text{ref}}^{\ell}$. Levels $\ell + 1$ and ℓ are synchronized at the beginning of each time step on ℓ , and operate independently between these time synchronization points. During this synchronization, the solution values on level ℓ are updated to match the cell-averages from $\ell + 1$, and fluxes at coarse-fine interfaces are corrected to ensure conservation. Flux correction replaces the fluxes on the coarse grid with those from the fine grid at the interface between the two levels. At intermediate steps on the fine level $\ell + 1$, ghost cells are interpolated from the coarse level ℓ in space and time.

2.3 | Conservative Interpolation Schemes for AMR

With an AMR grid, each level is intended to advance independently with data transfer between levels occurring at only refinement boundaries. For the FVM, the coarse cell-average is defined exactly by the average of fine cell-averages. From coarse to fine, data is interpolated. There are two AMR operations where interpolation is used directly. When regridding, a new level $\ell + 1$ is created, and its valid region is interpolated from existing coarse data on level ℓ . This interpolation must be conservative to maintain the strict conservation property of finite-volume schemes. At AMR interfaces, a number of operations to communicate between levels ℓ and $\ell + 1$ rely on having invalid ghost cells of level $\ell + 1$ interpolated from level ℓ . An example grid with one layer of invalid ghost cells is show in Fig. 4. This interpolation needs not be conservative because invalid ghost cells are only intermediate states used for flux reconstruction, although enforcing conservation may improve the quality of the solution¹. These two interpolations between levels have slightly different implementations, but both follow this procedure:

1. For each coarse cell \mathbf{i}^{ℓ} , determine an interpolation stencil $\mathcal{I}(\mathbf{i}^{\ell})$. The stencil must have sufficient degrees of freedom for the desired order of accuracy.
2. Construct a scheme-order multi-dimensional polynomial using the data in $\mathcal{I}(\mathbf{i}^{\ell})$ while maintaining conservation of $\langle \mathbf{U} \rangle_{\mathbf{i}^{\ell}}$.
3. For each desired fine cell $\mathbf{i}^{\ell+1} \in C^{-1}(\mathbf{i}^{\ell})$, evaluate the cell-averaged value using the prior reconstruction.

The interpolation procedure for both newly refined cells and invalid ghost cells as described by McCorquodale and Colella²⁰ is a form of k-exact reconstruction²². This process is prone to problematic oscillations when done at high-order, independent of whether a conservation constraint is enforced. Due to this similarity, bound-preservation can be enforced using the same method for both newly refined cells and invalid ghost cells.

2.3.1 | Multi-Dimensional Polynomial Interpolation

When interpolating an AMR solution, the idea is to construct a piecewise interpolant, $\mathbf{U}_i(\vec{\xi})$, for a coarse cell, \mathbf{i} . To achieve a specified order of accuracy, a multi-dimensional Taylor expansion about the center of cell \mathbf{i} is used. Multi-dimensional notation

is defined using \mathbf{q} as an integer vector of size D . Vector powers of \mathbf{q} of vector quantities $\vec{\xi}$ are expressed as

$$(\vec{\xi} - \vec{\xi}_i)^{\mathbf{q}} = \prod_{d=1}^D (\xi_d - \xi_{i,d})^{q_d}. \quad (6)$$

A multi-dimensional form of the Taylor series up to order Q , based on cell centers, $\vec{\xi}_i$, is expressed as

$$\mathbf{U}_i(\vec{\xi}) = \sum_{\|\mathbf{q}\|_1 < Q} \frac{1}{\mathbf{q}!} \mathbf{U}_i^{(\mathbf{q})}(\vec{\xi}_i) (\vec{\xi} - \vec{\xi}_i)^{\mathbf{q}} + O(h^Q), \quad (7)$$

where $\mathbf{U}_i^{(\mathbf{q})}(\vec{\xi}_i)$ is the \mathbf{q}^{th} derivative in multiple dimensions (e.g., $\frac{\partial^{(q_1+q_2+q_3)} \mathbf{U}}{\partial \xi^{q_1} \partial \eta^{q_2} \partial \zeta^{q_3}}$), and h is the cell spacing.

Choosing $c_{\mathbf{q}} = \frac{1}{\mathbf{q}!} \mathbf{U}_i^{(\mathbf{q})}(\vec{\xi}_i)$ as the unknowns to be solved, the Taylor polynomial in Eq. (7) working with cell-averaged quantities, defined in Eq. (4), can be written compactly as

$$\langle \mathbf{U}_i(\vec{\xi}) \rangle_i = \frac{1}{V_i} \int_{V_i} \sum_{\|\mathbf{q}\|_1 < Q} c_{\mathbf{q}} (\vec{\xi} - \vec{\xi}_i)^{\mathbf{q}} + O(h^Q) d\vec{\xi} \quad (8)$$

$$= \sum_{\|\mathbf{q}\|_1 < Q} c_{\mathbf{q}} \langle (\vec{\xi} - \vec{\xi}_i)^{\mathbf{q}} \rangle_i + O(h^Q). \quad (9)$$

For a unique construction, $K = \binom{D+Q-1}{D}$ terms (binomial coefficients) are needed. In general, there is no obvious unique choice for how to select a stencil to satisfy exactly this many constraints.

2.3.2 | Conservative Least Squares Interpolation on Mapped Grids

Since there is no obvious choice of the stencil, it is preferred to use a stencil $\mathcal{I}(i)$ of size $N > K$ so long as $i \in \mathcal{I}(i)$. This allows a stencil choice that appears more regular and has flexibility near boundaries²⁰. In practice we use interpolation stencils as shown by McCorquodale et al.²⁰. For interior solutions, this stencil uses a radius of one cell from the interpolation center with an additional cell in each face-normal direction, resulting in 13 cells in 2D and 33 in 3D. A similar stencil is used near boundaries, but shifted towards the interior to maintain the total number of cells. (See²⁰ for specific stencil construction details.)

Since the stencil size is overdetermined, the least squares method must be used to create the interpolating polynomial of best fit. This still retains the desired interpolation order of accuracy and has the added benefit of allowing extra constraints²⁰.

For interpolation between AMR levels, preserving scalar conservation is required. To fulfill this, a constrained least squares problem of the form

$$\min_{\mathbf{B}\mathbf{x} = \mathbf{d}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2, \quad (10)$$

is formulated, where the unknowns \mathbf{x} are a vector of the coefficients $c_{\mathbf{q}}$. The conservation constraint $\mathbf{B}\mathbf{x} = \mathbf{d}$ is enforced using the equation

$$\sum_{i^{\ell+1} \in C^{-1}(i^{\ell})} \sum_{\|\mathbf{q}\|_1 < Q} c_{\mathbf{q}} \langle J \hat{\xi}^{\mathbf{q}} \rangle_{i^{\ell+1}} = \langle J\mathbf{U} \rangle_{i^{\ell}}, \quad (11)$$

while the objectives $\mathbf{A}\mathbf{x} - \mathbf{b}$ to minimize are described by

$$\sum_{\|\mathbf{q}\|_1 < Q} c_{\mathbf{q}} \langle J \hat{\xi}^{\mathbf{q}} \rangle_{j^{\ell}} = \langle J\mathbf{U} \rangle_{j^{\ell}} \quad j^{\ell} \in \mathcal{I}(i^{\ell}). \quad (12)$$

Once the interpolating function has been constructed by solving coefficients $c_{\mathbf{q}}$, fine cells are evaluated as

$$\langle J\mathbf{U} \rangle_{i^{\ell+1}} = \sum_{\|\mathbf{q}\|_1 < Q} c_{\mathbf{q}} \langle J \hat{\xi}^{\mathbf{q}} \rangle_{i^{\ell+1}}. \quad (13)$$

Due to the conservation constraint, it can be shown that

$$\langle J\mathbf{U} \rangle_{i^{\ell}} = \sum_{i^{\ell+1} \in C^{-1}(i^{\ell})} \langle J\mathbf{U} \rangle_{i^{\ell+1}}, \quad (14)$$

using the refinement operator C^{-1} . The values of $\langle J \hat{\xi}^{\mathbf{q}} \rangle$ are computed using a fourth-order product formula and the interpolation step is shown to be freestream-preserving¹. When filling invalid ghost cells (see Fig. 4), interpolation of conservative values $\langle \mathbf{U} \rangle$ is sufficient, which is solved using constant values of J in space. For regridding operations where valid cells are filled, conservation of mapped quantities $\langle J\mathbf{U} \rangle$ is required.

This method of least squares interpolation is said to be k -exact²², that is polynomial solutions up to order Q are interpolated *exactly* using the approach. Any least squares remainder is of the desired order for smooth solutions. Using this approach, fourth-order accurate methods with AMR have been demonstrated^{20,1}.

3 | ADAPTIVE CLIPPING-AND-REDISTRIBUTION ALGORITHM

The adaptive clipping-and-redistribution, or ACR, method is designed as a post-processing operation independent of the high-order interpolation. This simplifies implementation since it requires no fundamental changes to the existing high-order interpolation algorithm, and can be adaptively enabled. The method maintains conservation, avoids significant data exchange, and does not require predetermined bounds. The procedure for ACR is outlined by the following steps:

1. Interpolate the fine data from the coarse data using a high-order conservative scheme as described in Section 2.3.
2. Establish minimum and maximum bounds for interpolation values of each fine cell. These bounds are created using local coarse cell information.
3. Clip each fine cell to fall within the established bounds and track the solution quantity change.
4. Redistribute the clipped excess to local regions that can accept it while remaining inside their bounds.

A few main features deserve emphasis and clarification. The primary objective of the ACR method is to prevent new extrema while maintaining solution conservation. Notably, this does not strictly enforce monotonicity of the solution. It is important to properly establish the local clipping bounds since their selection directly impacts the resulting solution. Furthermore, the extension of the ACR to high-order is desirable for maximum accuracy. At the end of the section, we provide verification of the new method.

3.1 | Establishing Local Clipping Bounds

Local clipping bounds provide a feasible range that the interpolation of each fine cell must lie within to prevent generation of new extrema. For each fine cell i on AMR level $\ell + 1$, an interval $\beta_{i^{\ell+1}}$ containing the valid interpolated solution is determined. Bounds are created from the minimum and maximum values of a neighborhood $\mathcal{I}(C(i^{\ell+1}))$ of coarse data on level ℓ about the fine cell $i^{\ell+1}$. In this work, the neighborhood $\mathcal{I}(C(i^{\ell+1}))$ is chosen to be the same as the neighborhood from the interpolation stencils. However, it need not be identical. Clipping bounds are determined as

$$\beta_{i^{\ell+1}} = \left[\min (\langle \mathbf{U} \rangle_{j^\ell}, j^\ell \in \mathcal{I}(C(i^{\ell+1}))), \max (\langle \mathbf{U} \rangle_{j^\ell}, j^\ell \in \mathcal{I}(C(i^{\ell+1}))) \right]. \quad (15)$$

The ACR method will ensure the interpolated fine level respects the bounds from the existing coarse level. As a result, the ACR method does not require explicit bounds to maintain physically realizable solutions, so long as the coarse solution itself is physically realizable.

The only requirement on local neighborhood choice $\mathcal{I}(C(i^{\ell+1}))$ is that it must contain the coarse cell $C(i^{\ell+1})$; however other choices do impact the behavior of the ACR method. For example, the smallest clipping neighborhood may be set to include only the coarse cell $C(i^{\ell+1})$. This removes any flexibility in the interpolation, and reverts to a piecewise constant interpolation. In the other extreme, the clipping neighborhood is expanded to the entire domain, and the ACR method will only restrict the creation of new global extrema. Our compromise is that clipping stencils should have the same locality as interpolation stencils; smaller neighborhoods are overly restrictive, while larger neighborhoods are not restrictive enough and are computationally impractical.

3.2 | High-Order Extension

In regions where data is smooth and in state variables with no strict physical bounds (such as velocity), high-order interpolation is expected to perform well without assistance from the ACR method. Reducing the interpolation order in these smooth regions can compromise physics that are correctly captured by the high-order scheme¹³. In such cases, interpolation of smooth data is expected to accurately generate subtle new extrema. This is illustrated in Fig. 5 where high-order interpolation correctly captures a Gaussian profile, and enforcing bounds via the ACR method results in a poor quality result. For this purpose the ACR

method is extended to high-order, hereafter called the HO-ACR method. The HO-ACR method is adaptive in that the operation is not engaged in smooth regions, while being engaged elsewhere. Regions are determined as smooth if all second derivative constructions share the same sign. The second derivatives are computed for each coarse cell i^ℓ , and compared to each of the directly adjacent neighbors. If the entire set share the same sign, the region is determined as smooth and the HO-ACR for fine cells $i^{\ell+1} \in C^{-1}(i^\ell)$ is disabled. It is at the discretion of the researcher whether high-order detection is used for any of the state variables. HO-ACR is used for all results presented herein except for the bluff-body combustor in Section 5.2.




FIGURE 5 An example of the high-order interpolation correctly creating new extrema of a smooth Gaussian profile.

Computing the curvature using a 1st-order accurate method is sufficient. Specifically, curvature is evaluated using a 2nd-order centered scheme as given by

$$\frac{\partial^2 \mathbf{U}_i}{\partial \xi_d^2} = \frac{\langle \mathbf{U} \rangle_{i-e^d} - 2\langle \mathbf{U} \rangle_i + \langle \mathbf{U} \rangle_{i+e^d}}{h^2} + O(h^2), \quad (16)$$

where h is the cell spacing. For locations near the physical boundaries, 1st-order one-sided formulations such as

$$\frac{\partial^2 \mathbf{U}_i}{\partial \xi_d^2} = \frac{\langle \mathbf{U} \rangle_i - 2\langle \mathbf{U} \rangle_{i+e^d} + \langle \mathbf{U} \rangle_{i+2e^d}}{h^2} + O(h), \quad (17)$$

are used as required.

3.3 | Local Linear Redistribution

Once clipping is completed, properly redistributing the clipped quantities is necessary to maintain conservation. One of our design goals for the redistribution scheme is to be computationally inexpensive by avoiding data exchange. A redistribution region is chosen such that the redistribution can be solved explicitly without additional data exchange by choosing non-overlapping redistribution regions \mathcal{R} . The natural choice for such regions is the coarse cells where data is being interpolated from, such that $\mathcal{R} = C(i^{\ell+1})$.

Many CFD algorithms use a mapped grid technique to accommodate moderately complex geometries, meaning that a general curvilinear coordinate transformation is used to map the physical domain to the computational domain. Operations on the mapped quantities $\langle J\mathbf{U} \rangle$ require additional care to enforce bounds on the physical solution $\langle \mathbf{U} \rangle$. Bounds can only be enforced on conservative quantities $\langle \mathbf{U} \rangle$ because gradients in $\langle J\mathbf{U} \rangle$ may appear solely as the result of grid mapping. However, regridding requires conservation of the mapped variables as in Eq. (11). The mapped HO-ACR algorithm must first solve for a 2nd-order conservative variable

$$\overline{\langle \mathbf{U} \rangle}_i = \frac{\langle J\mathbf{U} \rangle_i}{\langle J \rangle_i} + O(h^2). \quad (18)$$

On Cartesian grids, the same HO-ACR algorithm can be applied directly to conservative quantities $\langle \mathbf{U} \rangle$ by setting J to one, and Eq. (18) then becomes exact.

The method for localized linear redistribution begins by determining the overshoot from each fine cell $\delta \langle \mathbf{U} \rangle_i$ with

$$\delta \overline{\langle \mathbf{U} \rangle}_i = \begin{cases} \overline{\langle \mathbf{U} \rangle}_i - \beta_i^{\max} & \overline{\langle \mathbf{U} \rangle}_i > \beta_i^{\max} \\ \overline{\langle \mathbf{U} \rangle}_i - \beta_i^{\min} & \overline{\langle \mathbf{U} \rangle}_i < \beta_i^{\min} \\ 0 & \text{otherwise} \end{cases}, \quad (19)$$

where β_i^{\max} and β_i^{\min} are respective maximum and minimum of the previously determined clipping bounds β_i (e.g., an interval). The mapped quantities that are clipped off are computed by

$$\delta\langle\overline{JU}\rangle_i = \delta\langle\overline{U}\rangle_i\langle J\rangle_i + O(h^2), \quad (20)$$

which again is a 2nd-order operation. Using the clipping bounds, the extra mass to be redistributed in region \mathcal{R} is given by

$$E_{\mathcal{R}} = \sum_{j \in \mathcal{R}} \delta\langle\overline{JU}\rangle_j. \quad (21)$$

To determine where mass may be pushed to, each fine cell has a maximum allowable difference calculated as

$$\alpha_i = \begin{cases} \beta_i^{\max} - \langle\overline{U}\rangle_i & E_{\mathcal{R}} > 0 \\ \langle\overline{U}\rangle_i - \beta_i^{\min} & E_{\mathcal{R}} < 0 \\ 0 & \text{otherwise} \end{cases}. \quad (22)$$

Using this, each fine cell can be updated using the clipping-and-redistribution as

$$\langle\widetilde{JU}\rangle_i = \langle JU\rangle_i - \delta\langle\overline{JU}\rangle_i + \frac{\alpha_i E_{\mathcal{R}}}{\sum_{j \in \mathcal{R}} \alpha_j}. \quad (23)$$

In effect, this drops the order of accuracy of the interpolation to 1st-order in regions which are fully clipped. This mapping process only needs to be 2nd-order, since the HO-ACR method effectively reduces the interpolation order in non-smooth regions where clipping is engaged. Note that the smoothness check for the high-order extension must operate on $\langle U \rangle$.

3.3.1 | Conservation Proof

The ACR method can be shown to maintain conservation, by demonstrating

$$\sum_{i \in \mathcal{R}} \langle\widetilde{JU}\rangle_i = \sum_{i \in \mathcal{R}} \langle JU\rangle_i. \quad (24)$$

Assuming the prior equation holds, and substituting it into Eq. (23), conservation can be shown by

$$\begin{aligned} \sum_{i \in \mathcal{R}} \delta\langle\overline{JU}\rangle_i &= \sum_{i \in \mathcal{R}} \frac{\alpha_i E_{\mathcal{R}}}{\sum_{j \in \mathcal{R}} \alpha_j} \\ &= E_{\mathcal{R}}, \end{aligned}$$

which is true by the definition of Eq. (21), and therefore

$$\sum_{i \in \mathcal{R}} \langle\widetilde{JU}\rangle_i = \sum_{i \in \mathcal{R}} \langle JU\rangle_i$$

so long as

$$\sum_{i \in \mathcal{R}} \beta_i^{\max} \geq \sum_{i \in \mathcal{R}} \langle\overline{U}\rangle_i \geq \sum_{i \in \mathcal{R}} \beta_i^{\min}.$$

In general $\sum_{i \in \mathcal{R}} \langle\widetilde{U}\rangle_i \neq \sum_{i \in \mathcal{R}} \langle U \rangle_i$ except for the case of Cartesian grids in physical space when J is one.

3.3.2 | 1D Verification

The property of bound-preservation of the new method is verified using a range of 1D profiles from smooth to non-smooth, including 6 tests — (a) a step function, (b) a smooth step function, (c) a piecewise linear function, (d) a Gaussian profile, (e) a piecewise cubic function, and (f) a discrete delta function. Fig. 6 shows the spatial profiles of solution variable $\langle U \rangle$ and compares the performance of the HO-ACR method for the 6 cases on a Cartesian grid (by setting $J = 1$ uniformly).

The conservative interpolation is 4th-order accurate using a centered 5-point centered stencil. For all cases, the HO-ACR method successfully clips off the overshoots (e.g., cases a, b, c, e, f), while preserving the smooth extrema (e.g., case d). Additionally, all 1D cases maintain conservation to machine precision.

Despite the improvement in interpolation near sharp features, there is still room for improvement in the solution. In cases (c) and (e), the high-order interpolation at one coarse-cell away from the discontinuities does not correctly capture the coarse polynomial solution. This is due to the interpolation stencil spanning the discontinuities, and as a result oscillations are present in the high-order conservative interpolation. While the HO-ACR method is not able to correct the high-frequency oscillations,

it still does no worse than the existing high-order interpolation. It is important to note, the objective of the HO-ACR method is to prevent new extrema near discontinuous features and the method does not enforce monotonic interpolations of non-smooth coarse data. The emphasis of the HO-ACR method is to maintain solutions that are physically bounded, but leave the quality of these solutions up to the base interpolation scheme.

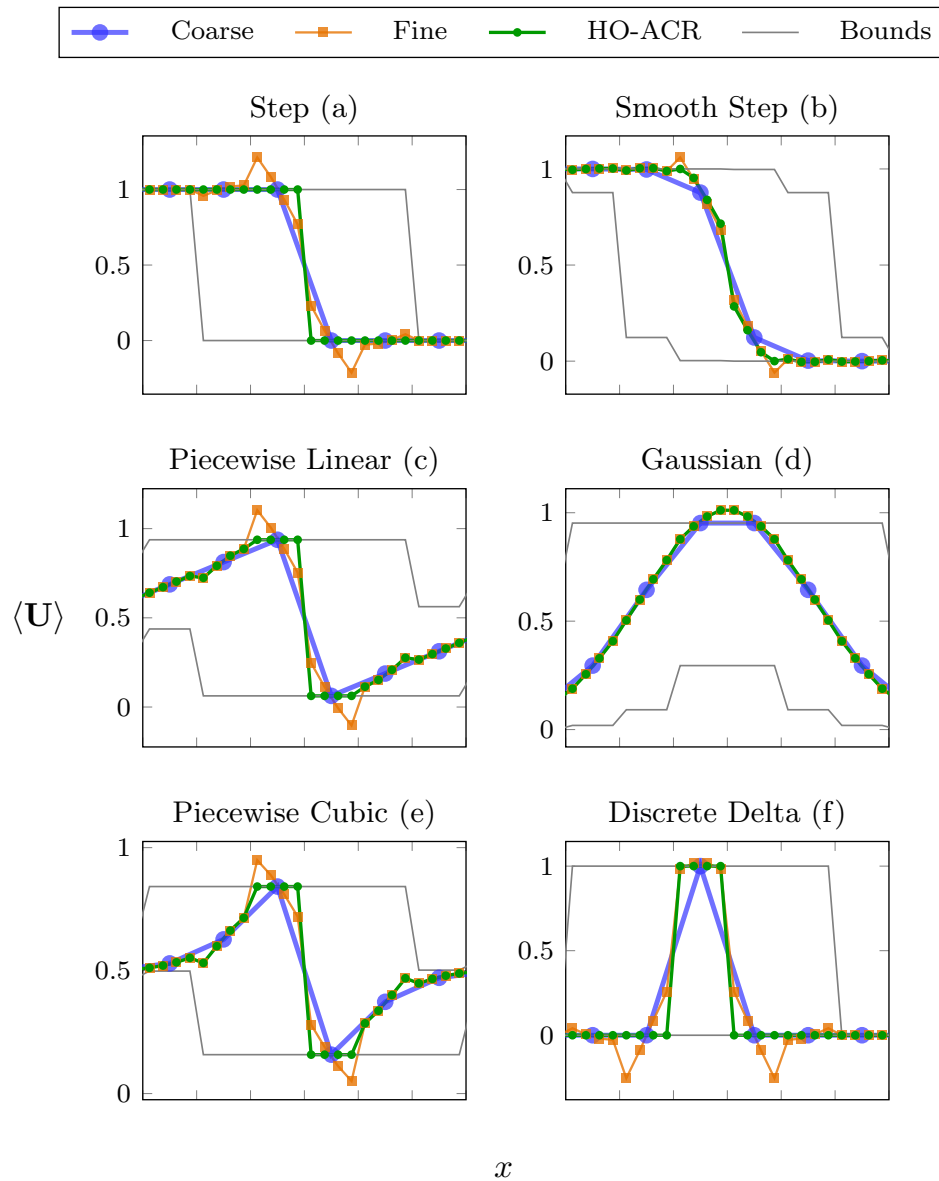


FIGURE 6 1D interpolation showing the HO-ACR method on Cartesian grids. The interpolation is 4th-order and conservative using a 5-point centered stencil. The x axis shows the coarse cells, and y axis the 4th-order cell-averaged values of the solution. Blue lines indicate coarse solution values, orange lines the coarse to fine interpolation, and green lines the interpolation with addition of the HO-ACR method. Gray lines show the bounding envelope used by the HO-ACR method for non-smooth regions.

Finally, we demonstrate the accuracy for some cases. For the following tests, we start with exact solution cell averages on progressively finer grids, and apply the base interpolation and HO-ACR algorithms using a factor of four refinement. The resulting convergence rates for L_1 and L_∞ norms over the $O(N)$ domain are plotted in Fig. 7.

The two smooth cases, the Gaussian from Fig. 6(d) and smooth step from Fig. 6(b), behave as expected in Fig. 7(a) and indicate 4th-order accuracy as anticipated. Note the results are nearly identical for the base and HO-ACR scheme in smooth

regions with sufficient resolution. The smooth step on coarser grids is clearly under resolved, resulting in reduced convergence rates. However, the HO-ACR reduces the error slightly at these coarse resolutions.

Fig. 7(b) shows the errors for the piecewise linear and piecewise cubic cases, corresponding to the profiles in Fig. 6(c) and (e) respectively. In both of these cases interpolation errors next to the discontinuities dominate, while interpolation is exact more than two coarse cells away (note the outermost cells in Fig. 6 (c) and (e)). Because of the discontinuities, the base method has an $O(1)$ overshoot, so that the L_∞ norm is constant and does not converge with refinement. Meanwhile, the HO-ACR method shows local 1st-order convergence. In both cases the error of the L_1 norms are one-order smaller than the L_∞ norms because a constant number of cells interpolate the discontinuity, and the L_1 norms scale these by the cell size. Thus, while the base interpolation is 1st-order, the HO-ACR method increases this to 2nd-order.

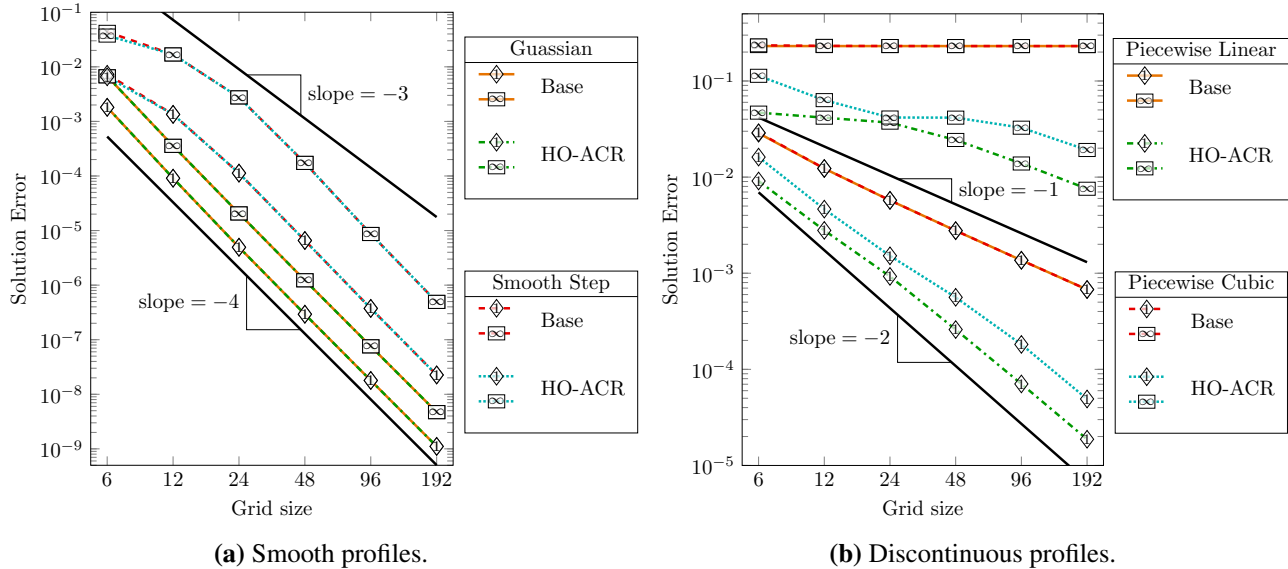


FIGURE 7 Convergence rates of one dimensional profiles. In all plots, the L_1 norms are represented by the diamond marks containing a 1, and the L_∞ norms are represented by the squares marks containing a ∞ .

Using the same 6 cases as shown in Fig. 6, we demonstrate the mapped HO-ACR performance on mapped grids. The setting of Fig. 8 and Fig. 9 is the same as that of Fig. 6 — the sequence of cases (a-f), the line color and symbols, and the interpolation scheme. Cases in Fig. 8 use a mapping function of $x = a \tan(b\xi)$ while in Fig. 9 the mapping function is $x = \frac{1}{2}\xi^2$. The mapped HO-ACR performs exceptionally well for both cases. The clipped profiles are bound-preserving and smooth extrema-preserving, while maintaining conservation to machine precision. As in the prior example, the cases (c) and (e) do not correctly match the coarse polynomial solution due to the high-order interpolation. Although there is room for improvement in regions away from the discontinuity, the solution is still improved as a whole over the unbounded high-order interpolation. Notably these small regions of error do not violate any existing coarse bounds of the solution.

Algorithm 1 presents the mapped HO-ACR for $\langle JU \rangle$ of AMR level $\ell + 1$. The same algorithm can be applied directly to conservative quantities $\langle U \rangle$ on a Cartesian grid by setting J to be one.

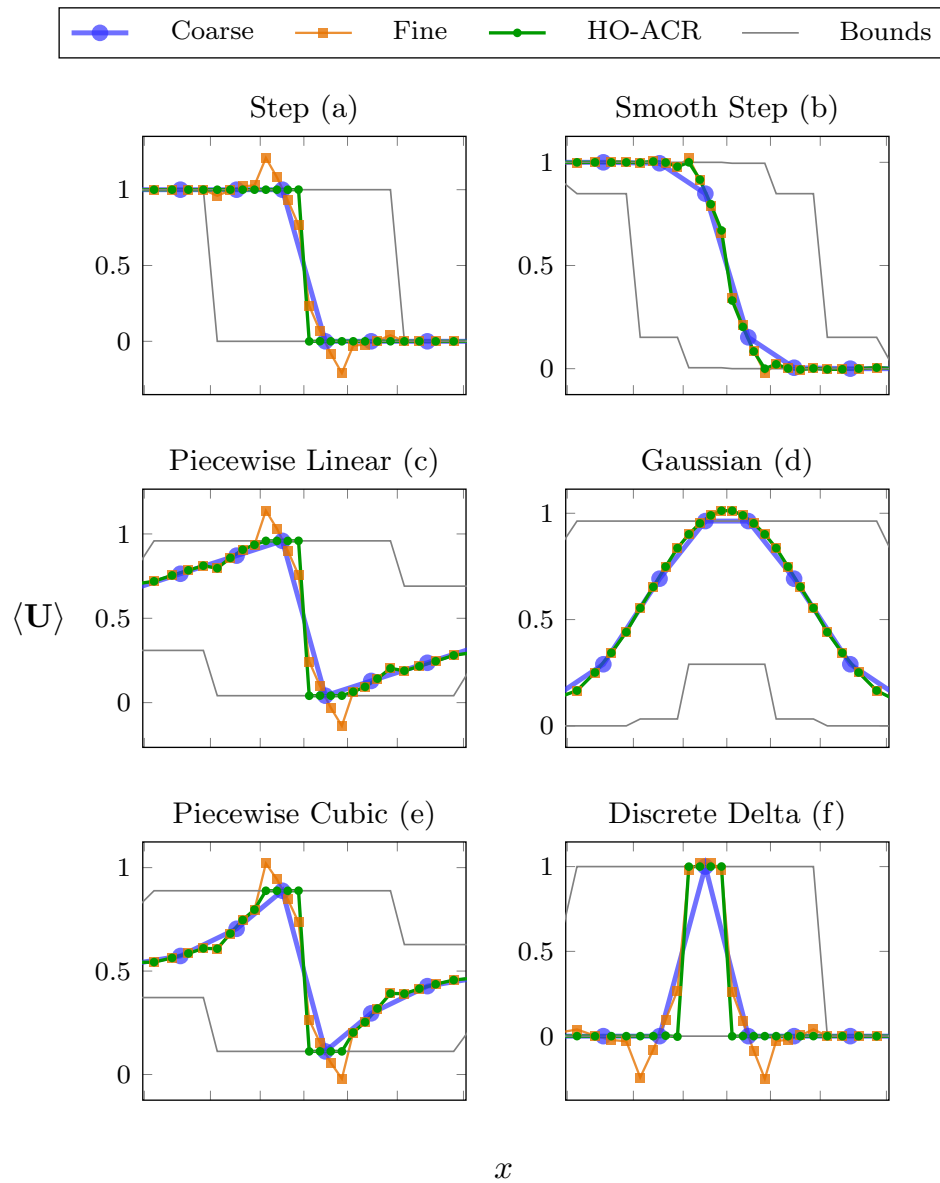


FIGURE 8 1D interpolation using the HO-ACR method on a mapped grid defined by $x = a \tan(b\xi)$, where the scaling factors are $a = (2 \cot(\frac{\pi}{8}))^{-1}$ and $b = \frac{3\pi}{4}$ for the domain centered about zero. The interpolation is 4th-order and conservative using a 5 point centered stencil. The x axis shows the coarse cells, and y axis the 4th-order cell-averaged values of the solution. The blue lines shows the values of coarse solutions, the orange lines the coarse to fine interpolation, and green lines the interpolation with addition of the HO-ACR method. Gray lines show the bounding envelope used by the HO-ACR method for non-smooth regions.

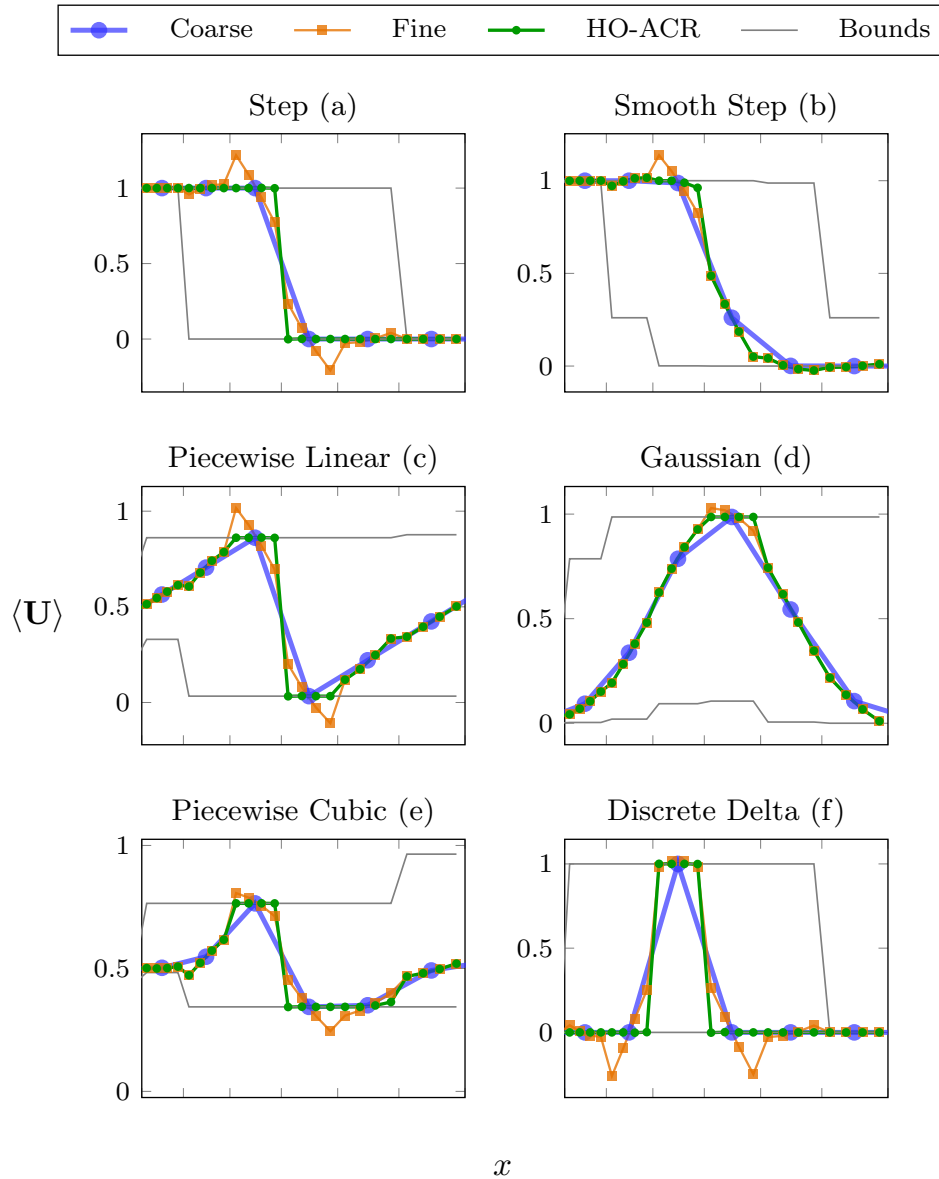


FIGURE 9 1D interpolation using the HO-ACR method on a mapped grid defined by $x = \frac{1}{2}\xi^2$ for positive ξ . The interpolation is 4th-order and conservative using a 5-point centered stencil. The x axis shows the coarse cells, and y axis the 4th-order cell-averaged values of the solution. The blue lines shows the values of coarse solutions, the orange lines the coarse to fine interpolation, and green lines the interpolation with addition of the HO-ACR method. Gray lines show the bounding envelope used by the HO-ACR method for non-smooth regions.

Algorithm 1 Mapped HO-ACR for $\langle \mathbf{JU} \rangle$ of level $\ell + 1$

for $i^\ell \in C(\Omega^{\ell+1})$ **do** ▷ Determine the clipping bounds
 for $i^{\ell+1} \in C^{-1}(i^\ell)$ **do**
 $\beta_{i^{\ell+1}}^{\min} = \min(\langle \mathbf{U} \rangle_{j^\ell}, \mathbf{j}^\ell \in \mathcal{I}(i^\ell))$
 $\beta_{i^{\ell+1}}^{\max} = \max(\langle \mathbf{U} \rangle_{j^\ell}, \mathbf{j}^\ell \in \mathcal{I}(i^\ell))$
 end for
end for
for $i^\ell \in C(\Omega^{\ell+1})$ **do** ▷ Clip-and-redistribute
 if $\text{sign}(\frac{\partial^2}{\partial \xi_d^2} \mathbf{U}_{i^\ell}) \neq \text{sign}(\frac{\partial^2}{\partial \xi_d^2} \mathbf{U}_{i^\ell \pm e^b})$ *for all* $b, d \in \mathbf{D}$ **then**
 for $i^{\ell+1} \in \mathcal{R}(i^{\ell+1})$ **do**
 $\overline{\langle \mathbf{U} \rangle}_{i^{\ell+1}} = \frac{\langle \mathbf{JU} \rangle_{i^{\ell+1}}}{\langle \mathbf{J} \rangle_{i^{\ell+1}}}$
 $\delta \overline{\langle \mathbf{U} \rangle}_{i^{\ell+1}} = \begin{cases} \overline{\langle \mathbf{U} \rangle}_{i^{\ell+1}} - \beta_{i^{\ell+1}}^{\max} & \overline{\langle \mathbf{U} \rangle}_{i^{\ell+1}} > \beta_{i^{\ell+1}}^{\max} \\ \overline{\langle \mathbf{U} \rangle}_{i^{\ell+1}} - \beta_{i^{\ell+1}}^{\min} & \overline{\langle \mathbf{U} \rangle}_{i^{\ell+1}} < \beta_{i^{\ell+1}}^{\min} \\ 0 & \text{otherwise} \end{cases}$
 $\delta \overline{\langle \mathbf{JU} \rangle}_{i^{\ell+1}} = \delta \overline{\langle \mathbf{U} \rangle}_{i^{\ell+1}} \langle \mathbf{J} \rangle_{i^{\ell+1}}$
 end for
 $E_{\mathcal{R}} = \sum_{j^{\ell+1} \in \mathcal{R}} \delta \overline{\langle \mathbf{JU} \rangle}_{j^{\ell+1}}$
 for $i^{\ell+1} \in \mathcal{R}(i^{\ell+1})$ **do**
 $\alpha_{i^{\ell+1}} = \begin{cases} \beta_{i^{\ell+1}}^{\max} - \overline{\langle \mathbf{U} \rangle}_{i^{\ell+1}} & E_{\mathcal{R}} > 0 \\ \overline{\langle \mathbf{U} \rangle}_{i^{\ell+1}} - \beta_{i^{\ell+1}}^{\min} & E_{\mathcal{R}} < 0 \\ 0 & \text{otherwise} \end{cases}$
 $\widetilde{\langle \mathbf{JU} \rangle}_{i^{\ell+1}} = \langle \mathbf{JU} \rangle_{i^{\ell+1}} - \delta \overline{\langle \mathbf{JU} \rangle}_{i^{\ell+1}} + \frac{\alpha_{i^{\ell+1}} E_{\mathcal{R}}}{\sum_{j^{\ell+1} \in \mathcal{R}} \alpha_{j^{\ell+1}}}$
 end for
 end if
end for

4 | VERIFICATION

The new HO-ACR algorithm is verified by a set of test cases involving the advection of smooth and sharp profiles on mapped grids. The features to be verified are: (1) ensuring high-order accuracy for smooth flow problems; (2) bound-preservation for non-smooth ones; and (3) maintaining conservation. Specifically the algorithm is examined for interpolation operations involving both the invalid ghost cells and the regridding process. The HO-ACR method is applied for invalid ghost cells and newly refined cells without fundamentally changing the base interpolation algorithm. The verification is first performed with the advection of a sharp profile in 1D, then advection of both a smooth profile and non-smooth profile in 2D. A more complicated case is explored with the solid body rotation of smooth and non-smooth profiles in 2D. All three features are demonstrated and assessed in each case.

4.1 | Advection of a 1D Sharp Profile

To explore the effects of the HO-ACR method in Chord, a one dimensional scalar advection test using an Euler equation solver is explored. The initial condition top-hat profile is defined by

$$\phi = \begin{cases} 1 & \text{if } 0.6875 > x > 0.3125 \\ 0 & \text{otherwise} \end{cases}$$

on a one dimensional periodic Cartesian grid of length $L = 2$. This initial profile aligns the discontinuity with cell boundaries on the coarse grid with spacing $\Delta x = 1/16$, which preserves the sharp discontinuity. The profile is advanced to $t = 0.5$ with speed $u = 1$. A time step size of $\Delta t = 5e-3$ is used on the coarse level, and finer levels are sub-cycled in time.

To investigate the HO-ACR method, three versions of “extreme” AMR are compared using Chord with and without the HO-ACR method. Each of these chosen “extreme” cases violates standard practice by intentionally refining in the vicinity of discontinuous features, and anticipates the emergence of new oscillations in the solution without the HO-ACR method. In each case, a single level of AMR with a refinement factor of 4 is used. Results are shown in figures 10, 11, and 12. As references, solutions run entirely on the fine and coarse grids are shown. Due to the nature of the 4th-order finite volume method, small overshoots and undershoots appear around discontinuous features as they evolve in time, even on the coarse grid. Ideally the results in these test cases with “extreme” AMR should never be worse than the coarse solution, and good results should match the fine solution.

The first case (10) examines the effects of delaying regridding, by initializing on the coarse grid and first advancing for one coarse time step. Beginning at the second time step, the solution is either refined using the base interpolation algorithm of Chord or using the HO-ACR method. The solutions after the second time step are shown in Fig. 10(a). A significant overshoot is observed in the base algorithm; the HO-ACR method effectively reduces these overshoots. We quantify the improvement in Table 1 by measuring the fine solution’s overshoot relative to the analytic solution. The base algorithm has overshoots that are an order of magnitude larger than the reference fine solution, while the HO-ACR modification has slightly smaller errors than the fine solution. Although this represents a significant improvement over the base algorithm, the HO-ACR method does not fully remove the small oscillations near the discontinuity, which is reflected in the larger L_2 norm than the fine solution.

To see the long-term impact, the solutions are advanced to 100 time steps, with results in Fig. 10(b). The overshoots from refining using the base algorithm are still present, while the solution using the HO-ACR solution closely matches the reference fine case, with extrema quantified in Table 2. The base algorithm has extrema that are nearly an order of magnitude larger than the fine solution, while the HO-ACR has extrema that almost match the fine solution. Comparing the difference in extrema demonstrates the improvement of HO-ACR, which is nearly two orders of magnitude smaller. To summarize, improper refinement of the initial condition by tagging one time step too late has significant and lasting consequences in the evolution of a solution. However, with the use of the HO-ACR method we are able to mitigate the effects of imperfect initial tagging. In more complex cases where discontinuous features appear spontaneously, this indicates that the HO-ACR method may be more accurate and robust with new flow features.

The second case examines the effects of ghost cell interpolation by specifying an initial profile with discontinuities on the coarse grid, shown in Fig. 11. Directly ahead of the profile is a region of AMR fixed in space for $x > 0.75$. Fig. 11(a) shows the solution as the leading edge enters the refinement region. The base algorithm produces a leading negative undershoot on the fine region, as well as an overshoot in the coarse region. The HO-ACR method prevents both of these new extrema, with minor oscillations from the 4th-order evolution scheme. In Fig. 11(b), the solution has been advanced until the profile has mostly




FIGURE 10 1D advection with refinement spanning the entire domain at the beginning of the 2nd time step, with tick marks representing the finest grid spacing. The orange lines are the solution using the base algorithm of Chord, and green lines show the solution with addition of the HO-ACR method. The blue lines show the coarse solution values with no AMR, and black lines show the solution purely on the fine grid. The plot window follows the profile in time and the labeled x -axis locations indicate the analytic of solution discontinuities.

	Absolute Overshoot			Overshoot Relative to Fine	
	Base	HO-ACR	Fine	Base	HO-ACR
Max	1.621e-1	3.950e-2	4.259e-2	1.195e-1	—
Min	-1.621e-1	-3.950e-2	-4.259e-2	-1.195e-1	—
L_2	8.757e-4	1.690e-4	1.511e-4	5.969e-4	—
L_1	1.189e-2	7.144e-4	8.014e-4	1.047e-3	—

TABLE 1 Interpolation overshoots for 1D advection with AMR added at the beginning of the 2nd time step, measured at 2 time steps. Absolute overshoot is defined by solution values exceeding the bounds $[0, 1]$. The overshoot relative to fine data are calculated using solution values exceeding the fine solution peaks at the matching time. Since the HO-ACR method does not exceed the fine solution peaks, there are no relative errors to measure.

entered the fine region. The extrema persist in the base algorithm while the HO-ACR method keeps them suppressed. Extrema at the end time are quantified in Table 3 relative to the analytic and fine solutions. The overshoots of the HO-ACR method are only marginally larger than the fine solution, while the base algorithm roughly doubles the overshoot. Surprisingly, the minimum values from the HO-ACR are even smaller than those in the fine solution. Because of the smoother solution from the HO-ACR method, the L_2 norm of its extrema is smaller than the fine solution's, and the L_1 norm nearly matches, while the base algorithm case is significantly worse. The differences in extrema between the base algorithm and the HO-ACR become more pronounced when compared relative to the fine solution. Relative to the extrema allowed by the fine solution, the HO-ACR method reduces the overshoot in L_2 and L_1 norms by over two orders of magnitude from the base algorithm. To summarize, this test demonstrates better robustness around sharp features and AMR interfaces when using HO-ACR. In realistic cases, this shows that discontinuous features can interact with AMR interfaces without causing excessive overshoots and undershoots.

	Absolute Overshoot			Overshoot Relative to Fine	
	Base	HO-ACR	Fine	Base	HO-ACR
Max	9.623e-2	1.635e-2	1.374e-2	8.249e-2	2.611e-3
Min	-9.623e-2	-1.634e-2	-1.375e-2	-8.248e-2	-2.597e-3
L_2	2.793e-4	5.335e-5	6.872e-5	2.211e-4	4.744e-6
L_1	5.503e-3	1.230e-3	1.187e-3	3.842e-3	3.794e-5

TABLE 2 Interpolation overshoots for 1D advection with AMR added at the beginning of the 2nd time step, measured at 100 time steps. Absolute overshoot is defined by solution values exceeding the bounds $[0, 1]$. Overshoot relative to the fine solution is calculated using values exceeding the minimum and maximum present in the fine solution at the same time.

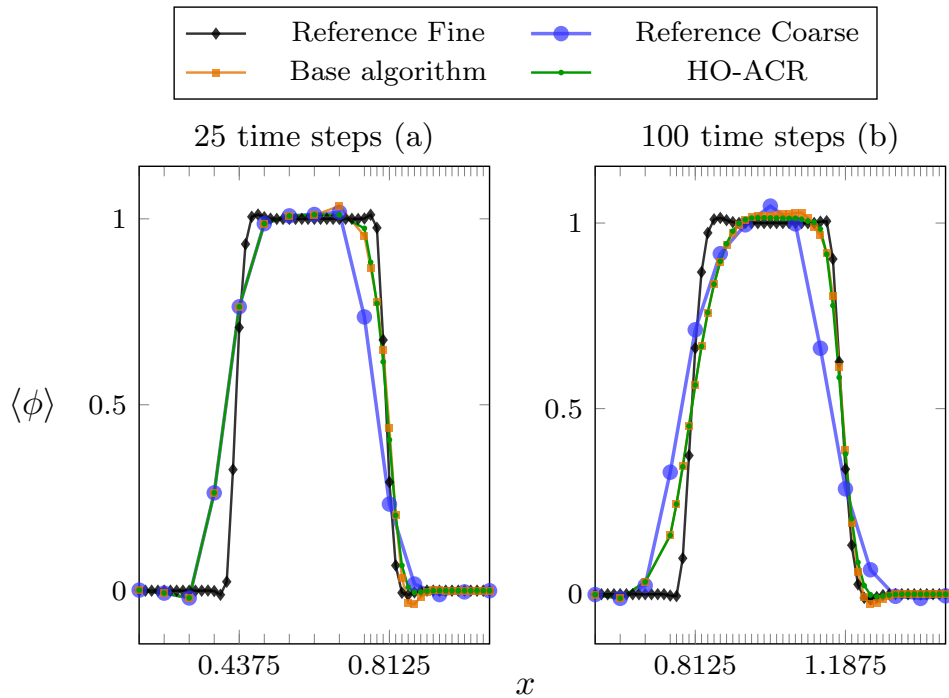


FIGURE 11 1D advection with AMR at a fixed location in space, with refinement shown by the tick marks. The orange lines are the solution using the base algorithm of Chord, and green lines show the solution with the HO-ACR method. The blue lines show the values of the coarse solution with no AMR, and the black lines the solution purely on the fine grid. The plot window is adjusted to follow the profile in time, and the labeled points of the x -axis are the analytic location of the discontinuities.

The third case, shown in Fig. 12, defines the initial profile with immediate refinement based on the solution gradients. The mesh is regrided at the beginning of each time step over the entire run. However, the refinement is applied minimally so that only one coarse cell on each side of the discontinuity is refined. This causes insufficient spacing between the discontinuity and the interpolation stencil for filling invalid ghost cells. After the time $t = 0.125$, Fig. 12(a) shows the base algorithm solution has new extrema emerge in both the coarse and fine regions, and begins to move away from the reference fine solution. The HO-ACR method prevents these new extrema and is able to closely match the reference fine solution with less than half the number of cells. This becomes more pronounced as the solutions advance to the final time $t = 0.5$ (Fig. 12(b)). At the end time Table 4 again shows the extrema of the base algorithm exceeding the analytic and fine solutions. Comparing extrema of the fine and HO-ACR solutions shows good agreement. Although we are hesitant to recommend this as a general practice, this test shows that using the HO-ACR method enables smaller regions of refinement (and reduced computational cost) with minimal loss of solution quality.

	Absolute Overshoot			Overshoot Relative to Fine	
	Base	HO-ACR	Fine	Base	HO-ACR
Max	2.820e-2	1.411e-2	1.374e-2	1.446e-2	3.713e-4
Min	-2.742e-2	-1.113e-2	-1.375e-2	-1.368e-2	—
L_2	7.006e-5	2.011e-5	6.872e-5	6.332e-5	9.153e-7
L_1	3.338e-3	1.725e-3	1.187e-3	1.073e-3	7.237e-6

TABLE 3 Interpolation overshoots for 1D advection with AMR at a fixed location in space, measured at 100 time steps. Absolute overshoot is defined by solution values exceeding the bounds $[0, 1]$. Overshoot relative to the fine solution is calculated using values exceeding the minimum and maximum present in the fine solution at the same time. Since the HO-ACR method has a smaller minimum value than the fine solution, no relative minimum value is recorded.

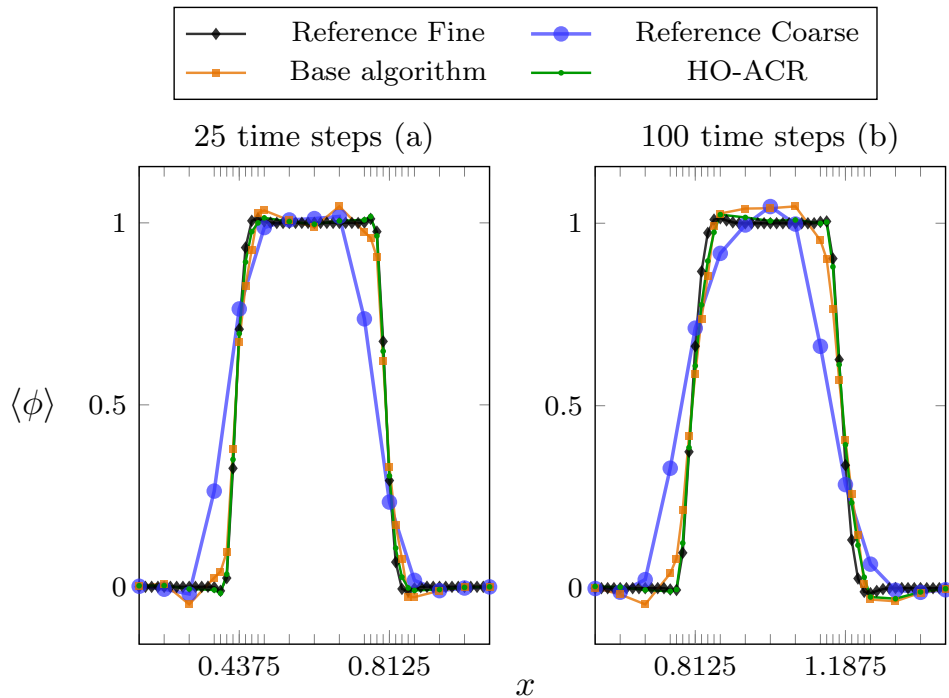


FIGURE 12 1D advection with AMR using the minimal tagging of two coarse cells spanning the solution discontinuity, with refinement shown by the tick marks. The orange lines are the solution using the base algorithm of Chord, and green lines show the solution using the HO-ACR method. The blue lines show the values of the coarse solution with no AMR, and the black lines indicate the solution purely on the fine grid. The plot window is adjusted to follow the profile in time, and the labeled points of the x-axis indicate the analytic solution discontinuities.

4.2 | Advection of a Smooth Profile

In order to ensure the HO-ACR correctly allows for smooth interpolation on two dimensional mapped grids, a Gaussian profile in density is specified by

$$\rho = \rho_0 + s(r) \Delta \rho e^{-(100r^2)}, \quad (25)$$

where

$$s(r) = \begin{cases} 0 & : |4r| \geq 1 \\ \cos^6\left(\frac{\pi}{8}r\right) & : |4r| < 1 \end{cases}, \quad (26)$$

$\rho_0 = 1.4$, $\Delta \rho = 0.14$, and r specifies the distance of a point from initial profile center at $(0.25, 0.375)$ in the periodic domain, $[0, 1]^D$ where D is 2. Constant pressure and velocity are set as 1 and $(1.0, 0.5)$, respectively. An initial profile is specified on a

	Absolute Overshoot			Overshoot Relative to Fine	
	Base	HO-ACR	Fine	Base	HO-ACR
Max	4.619e-2	2.358e-2	1.374e-02	3.245e-2	9.839e-3
Min	-4.435e-2	-2.874e-2	-1.375e-02	-3.060e-2	-1.499e-2
L_2	3.829e-4	1.969e-4	6.872e-5	2.443e-4	7.677e-5
L_1	3.087e-3	1.945e-3	1.187e-3	1.795e-3	3.893e-4

TABLE 4 Interpolation overshoots for the 1D advection, after 100 time steps, with minimal refinement of two coarse cells across the solution discontinuity. Absolute overshoot is defined by solution values exceeding the bounds $[0, 1]$. Overshoot relative to the fine solution is calculated using values exceeding the minimum and maximum present in the fine solution at the same time.

coarse grid and advected in time. The grid mapping is defined as

$$x_d = \xi_d + \prod_{p=1}^D \sin(2\pi\xi_p) \quad d = 1, 2. \quad (27)$$

To test invalid ghost cell interpolation, a fixed region of AMR (also called fixed AMR throughout the paper) is added as shown in Fig. 13(a). The initial profile does not initially fall inside the refined region, but is advected into it over time. Testing regridding is done by initially starting with the coarse resolution only, then refining the solution profile based on the gradient of ρ with a level of AMR after a specified time of 0.25 time units, as shown in Fig. 13(b).

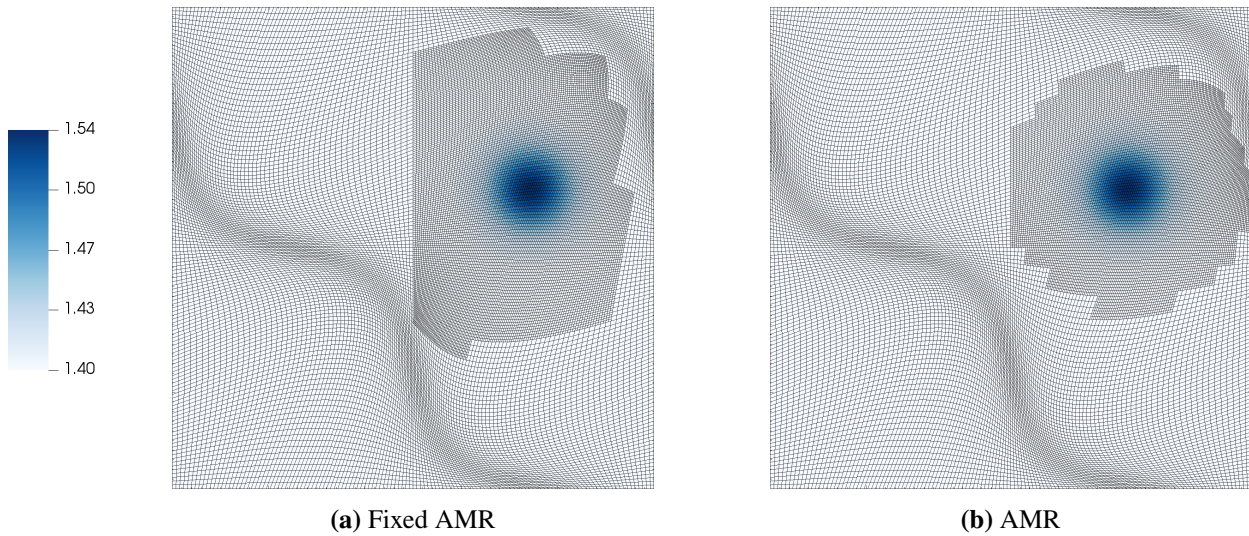


FIGURE 13 Warped meshes and solutions of density at a point in time for testing invalid ghost cell and regrid interpolation.

For a smooth profile, the HO-ACR scheme is designed to allow for new interpolated extrema to maintain the 4th-order accuracy of the underlying finite-volume scheme. The profile is advected for a total of 0.5 time units, after which solution errors are calculated by comparing the numerical and exact solutions. Tables 5 and 6 show the measured solution errors in L_1 , L_2 , L_∞ for grids sizes 64^2 , 128^2 , 256^2 and 512^2 , and the convergence rates in L_1 , L_2 , L_∞ between each two consecutive grids, respectively. Clearly, the measures verify the high-order accuracy for both the fixed AMR and regridded AMR cases using a smooth profile.

Conservation is measured by computing a relative difference in the volume weighted sum of the solution. Ideally, these differences should be within machine precision of zero between the beginning and the end of a run. Using the fixed AMR case, Chord with the HO-ACR method is measured to have a relative difference in conservative values of $2.104e-14$, whereas base Chord has value of $2.326e-14$ on the 256^2 case. This small difference shows that the HO-ACR method is conservative. For the regridded AMR case, the relative conservation difference for the 256^2 mesh is measured as $3.416e-13$ and $2.777e-13$ using Chord with and without the HO-ACR method, respectively.

Coarse Cells	L_1	Rate(L_1)	L_2	Rate(L_2)	L_∞	Rate(L_∞)
64^2	3.971e-05	—	2.586e-04	—	4.227e-03	—
128^2	2.675e-06	3.89	1.734e-05	3.90	2.919e-04	3.86
256^2	1.708e-07	3.97	1.108e-06	3.97	1.850e-05	3.98
512^2	1.080e-08	3.98	6.980e-08	3.99	1.166e-06	3.99

TABLE 5 Convergence rates for advection of a Gaussian profile using a fixed region of AMR.

Coarse Cells	L_1	Rate	L_2	Rate	L_∞	Rate
64^2	3.878e-05	—	2.409e-04	—	4.101e-03	—
128^2	2.607e-06	3.90	1.631e-05	3.88	2.714e-04	3.92
256^2	1.660e-07	3.97	1.040e-06	3.97	1.719e-05	3.98
512^2	1.042e-08	3.99	6.530e-08	3.99	1.081e-06	3.99

TABLE 6 Convergence rates for advection of a Gaussian profile using AMR.

4.3 | Advection of a Sharp Profile

In this test, using a similar setup to the Gaussian profile, a sharp profile is advected in a periodic domain allowing for verification that the HO-ACR method is bound-preserving. The domain is of size $[0, 1]^D$ where D is 2. The sharp profile is the “balls and jack’s” inspired by Anderson et al.¹⁶. This profile is composed of three features prescribed in density $\rho_0 + \Delta\rho$ where $\rho_0 = 1.4$ and $\Delta\rho = 0.14$, as shown in Fig. 14. Relative to a profile center of $(0.25, 0.375)$, a smaller circular shell is located at distance $(0.1, -0.1)$ from the center with an outer radius of 0.07 and an inner radius of 0.04. A larger circular shell is located at distance $(0.1, 0.1)$ from the profile center with an outer radius of 0.10 and an inner radius of 0.07. Finally, the jack base is placed at the distance $(-0.11, -0.11)$ from the profile center, with one arm specified by $[-0.04, -0.095] \times [-0.25, -0.125]$ the other by $[-0.095, -0.04] \times [-0.125, -0.25]$ and both are rotated counter-clockwise by 45° .




FIGURE 14 Initial density profile of the sharp “balls and jack’s”.

Beginning with the initial conditions as shown in Fig. 14 and a fixed region of AMR, the profile is advected forward in time for 0.492843 time units. Fig. 15 shows the fixed region of mesh refinement and compares the solution profiles obtained by Chord on the left and the inclusion of HO-ACR algorithm on the right. To emphasize, the HO-ACR method establishes the bounds from the solution on the coarse grid without AMR, serving as the basis to determine overshoots. Despite only bounding invalid ghost cells, a significant reduction in regions that experience overshoots beyond the coarse regions is seen when using the HO-ACR

method. Although the improvement may not be immediately present based on the number of points producing new extrema, the magnitude of these overshoots is greatly reduced as shows by the tabulated data in Fig. 15. Some new extrema are to be expected, as the high-order finite-volume algorithm includes other unbounded operations.

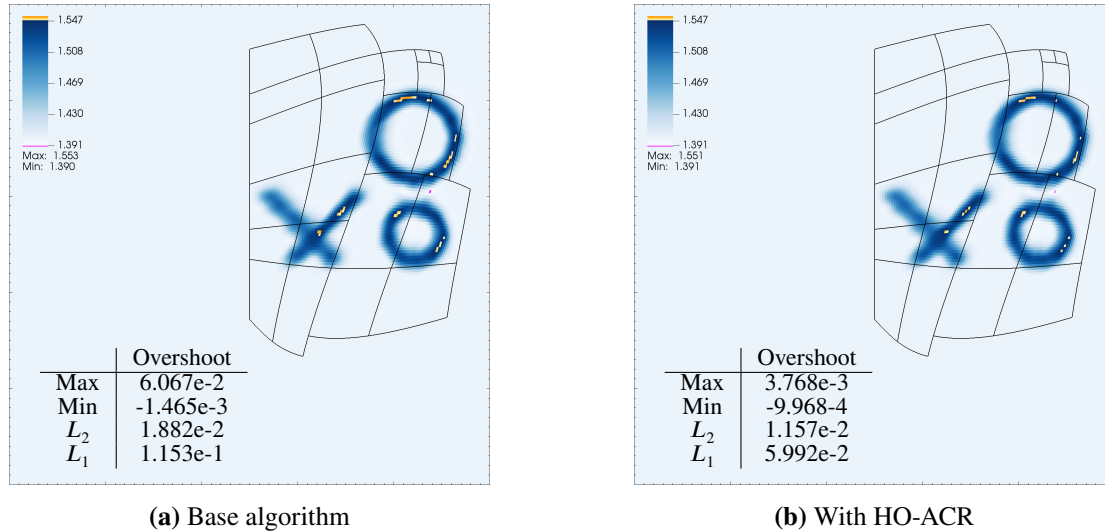


FIGURE 15 AMR at a fixed location in space, testing invalid ghost cell interpolation.

To test the HO-ACR method for the regridding process, the initial sharp profile in Fig. 14 is advected forward in time for 0.492843 time units. A new region of refinement which tracks the profile is then added as shown in Fig. 16, where the refinement spans the periodic boundary. The figure compares the solution profiles obtained by Chord without (a) and with (b) the HO-ACR algorithm. Clearly, the HO-ACR algorithm prevents many of the new extrema when regridding sharp solutions as in Fig. 16. Notice from the HO-ACR solution profile that a small number of values appear to violate the bounds. However, this is due to the fact that some 4th-order operations have happened between the final application of the HO-ACR and the output of the solution. These operations will naturally produce new extrema beyond prior solution bounds. One example is the conversion from one set of cell-averaged quantities to another²³, such as done for computing primitives from conservative values or conservative values from the mapped values.

Table 7 shows the solution growth of extrema, their L_1 , and L_2 -norms. For the fixed AMR test, the HO-ACR method is shown to have reduced the magnitude of the new maximum by a factor of 16, minimum by a factor 1.5, and only modest reduction in the L_1 and L_2 -norms. More impressively, the adaptive test reduces the magnitude of the new maximum by a factor of 8, and the minimum by a factor of 25. Decreases of over an order of magnitude are seen by the L_1 and L_2 -norms. As in the prior section, conservation is measured and the HO-ACR method is found to have negligible impact on the total solution conservation. The relative conservation difference of the 256^2 fixed AMR case is measured as $3.318e-15$ and $5.214e-15$ using Chord with and without the HO-ACR method, respectively. Similarly, the regridded case measures relative conservation of $6.099e-13$ and $1.054e-13$ using Chord with and without the HO-ACR method.

	Fixed		Adaptive	
	Base	HO-ACR	Base	HO-ACR
Max	6.067e-2	3.768e-3	3.910e-3	5.010e-4
Min	-1.465e-3	-9.968e-4	-4.094e-3	-1.62e-4
L_2	1.882e-2	1.157e-2	9.567e-3	5.514e-4
L_1	1.153e-1	5.992e-2	3.906e-2	9.031e-4

TABLE 7 Interpolation overshoots for the advected “ball and jacks” case.

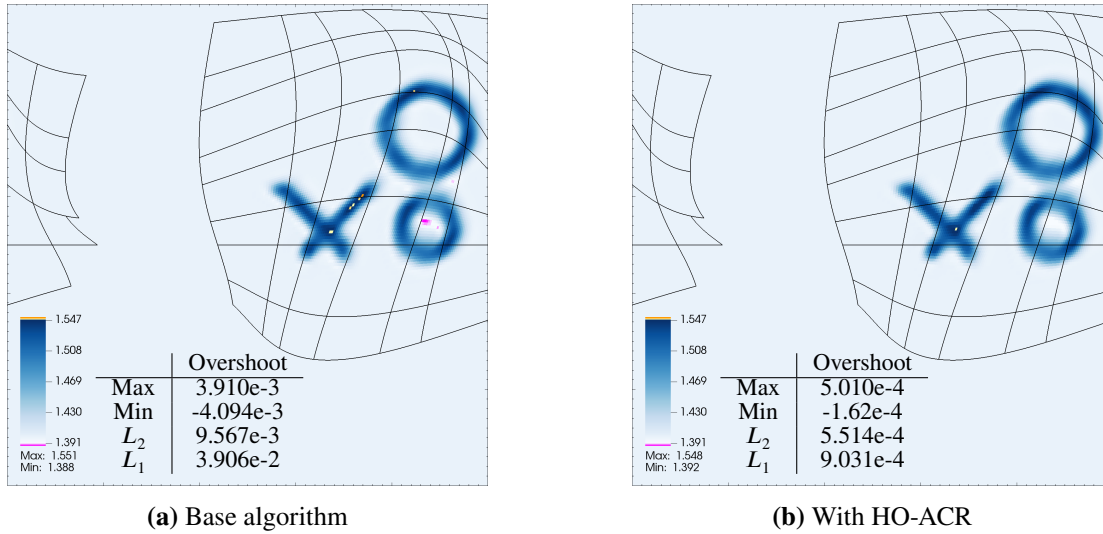


FIGURE 16 AMR added at this point in time, testing regrid interpolation.

4.4 | Solid Body Rotation

Another standard test problem for the HO-ACR method uses 2D solid body rotation to advect a passive scalar quantity with different localized initial profiles. The velocity field is defined by the profile $u = \omega(y, -x)$, where the rotational speed is $\omega = 20\pi$ in the clockwise direction. The initial profiles (Anderson et al. ²⁴) are specified by

$$\phi = \phi_0 + \Delta\phi(\phi_{\text{cone}}(r_{\text{cone}}) + \phi_{\text{bump}}(r_{\text{bump}}) + \phi_{\text{cyl}}(r_{\text{cyl}})) \quad (28)$$

where $\phi_0 = 0$, $\Delta\phi = 1$, and $r_{\text{cone}}, r_{\text{bump}}, r_{\text{cyl}}$ specify the distances from the initial profile center as $(\cos(\pi/6), \sin(\pi/6))$, $(\cos(3\pi/2), \sin(3\pi/2))$, and $(\cos(5\pi/6), \sin(5\pi/6))$ respectively. The cone, bump, and notched cylinder profiles are given as

$$\phi_{\text{cone}}(r) = \begin{cases} 0 & \text{if } r > 1 \\ 1 - r & \text{otherwise} \end{cases} \quad (29)$$

$$\phi_{\text{bump}}(r) = \begin{cases} 0 & \text{if } r > 1 \\ \cos(\pi r) & \text{otherwise} \end{cases} \quad (30)$$

$$\phi_{\text{cyl}}(r) = \begin{cases} 0 & \text{if } r > 1 \\ 0 & \text{if } (1.3r_j > y > 0.7r_j) \text{ and } (x > 1.4r_j) \\ 1 & \text{otherwise} \end{cases} \quad (31)$$

and shown as evaluated on the initial mesh in Fig. 17(a). The structured grid is defined by an annulus centered at the origin with outer radius $\sqrt{8}$ and inner radius $0.1\sqrt{8}$. The base mesh uses 64 cells in the radial direction and 128 in the angular. The inner and outer boundary conditions are specified by Dirichlet conditions using the analytic solution values. The solution is advanced in time for one rotation, or 0.1 units in time, after which the profiles should exactly match their initial conditions.

This test is motivated by tagging regions other than where the solution discontinuities are, and forcing the discontinuities to pass through the refined regions. For this imperfect application of AMR, the primary objective for the scalar solution is that it is stable and produces no new extrema as it is regrided and advects across AMR boundaries. This case is examined using three patches of counter rotating AMR at a refinement factor of 2. These patches are initially centered about $\pi/2$, $7\pi/6$, and $11\pi/6$ with widths of $\pi/8$, respectively marked A, B, and C in Fig. 17(a). The AMR patches are rotated in the opposite direction of the scalar advection by an angular velocity of $-5\omega/6$, and regrided every 10 time steps. Initial profiles are defined only in coarse regions with this choice of AMR, whereas the final profiles primarily fall inside refined regions, indicated with black boxes in Fig. 17. Results from using the base algorithm of Chord and with the addition of the HO-ACR method are examined and shown in Fig. 17(c) and Fig. 17(d), respectively. A reference using a fine grid for the entire domain is shown in Fig. 17(b). In each of these cases values outside the analytic solution bounds of $[0, 1]$ are highlighted, with overshoots in the yellow-orange color, and undershoots in magenta.

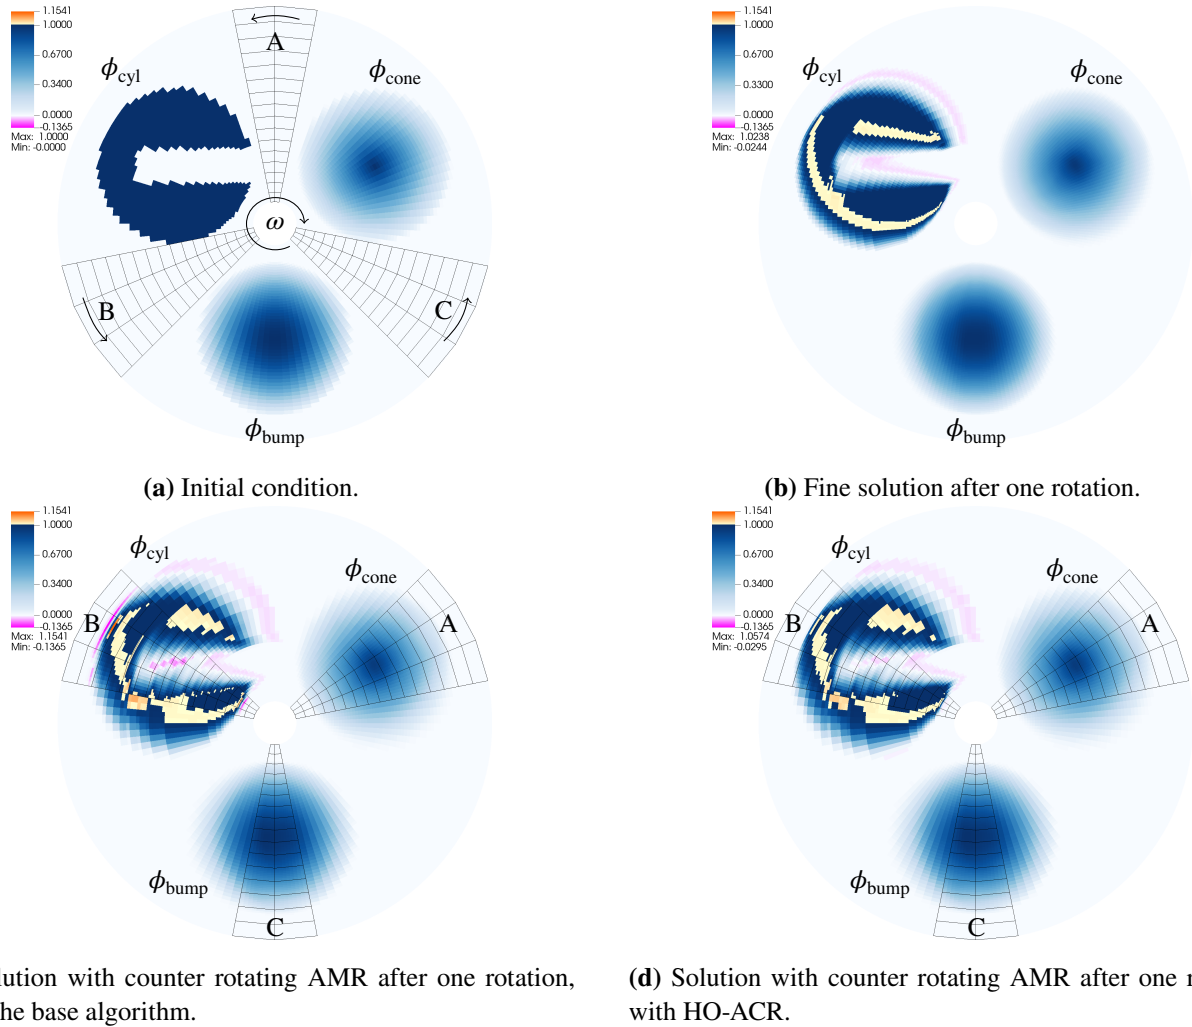
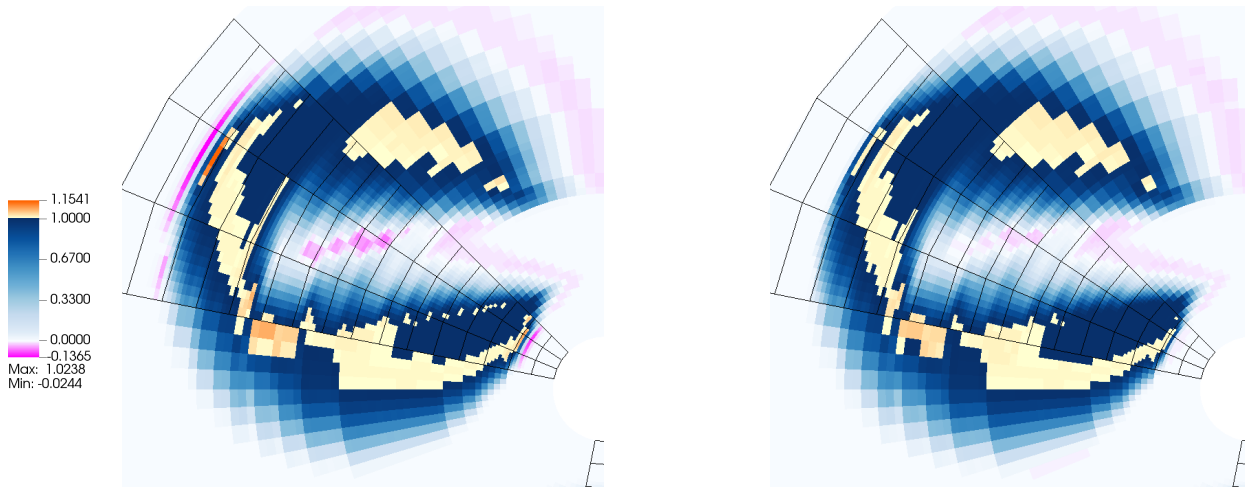


FIGURE 17 Solution of a passive scalar after advecting for one clockwise rotation. Negative species mass fractions are shown in magenta, and positive overshoots beyond the coarse solution are shown in yellow-orange. Black lines are a coarsened representation of the AMR region. Subfigures (b)–(d) show the results for a fine-only mesh, the base AMR algorithm with moving refinement domains, and the HO-ACR method with the same meshes, respectively. Each solution profile is labeled as ϕ_{cone} , ϕ_{bump} , and ϕ_{cyl} matching the definitions given in Eq. (31). The refinement regions labeled as A, B, and C, are regridded to make $5/6^{\text{th}}$ of a rotation counter-clockwise.

The reference fine solution, in Fig. 17(b), shows some small new extrema near the discontinuous features of the notched cylinder due to the 4th-order flux calculations in Chord. Even with imperfect AMR, the cone and bump profiles do not show any overshoot. This is expected since the scheme’s numerical dissipation smooths continuous profiles over time, and AMR interpolation is designed for smooth solutions. With this in mind, the continuous cone and bump regions of the solutions show no distinguishable differences with or without the HO-ACR method in Fig. 17(c), which indicate it is not changing the algorithm substantially in smooth regions. However, the more extreme case of the notched cylinder using the counter rotating AMR with Chord in Fig. 17(c), shows significantly larger overshoots and undershoots. These errors are the most pronounced at the edges of the notched cylinder, especially near the inner and outer boundaries, parallel to the velocity. These larger extrema are quantified in Table 8, labeled as the “base” columns, where the peak errors are seen to be an order of magnitude larger than what is observed in the fine solution. In contrast, the HO-ACR method shown in Fig. 17(d) greatly reduces the extrema in the solution and closely resembles the fine solution. For better visual comparison between the two cases, Fig. 18 shows solutions in a window containing only the notched cylinder. The extrema are measured in Table 8, (columns labeled as “HO-ACR”) where we see that the peak value errors are an order of magnitude smaller than the base case, and comparable to the fine case. The same trend is apparent in



(a) Solution with counter rotating AMR after one rotation, using the base algorithm of Chord. (b) Solution with counter rotating AMR after one rotation, with HO-ACR.

FIGURE 18 Close-up of the notched cylinder results after advecting for one rotation, as shown in Fig. 17. Negative species mass fractions are shown in magenta, and positive overshoots beyond the coarse solution are shown in yellow-orange. Black lines are a coarsened representation of the AMR region.

the L_1 and L_2 measure of the errors, with the HO-ACR overshoot an order of magnitude smaller than the base case and closer to the fine solution.

	Absolute Overshoot			Overshoot Relative to Fine	
	Base	HO-ACR	Fine	Base	HO-ACR
Max	1.541e-1	5.741e-2	2.381e-2	1.303e-1	3.362e-2
Min	-1.365e-1	-2.954e-2	-2.435e-2	1.121e-1	-5.197e-3
L_2	5.243e-5	2.577e-5	1.485e-5	1.607e-5	2.216e-6
L_1	1.258e-3	8.091e-4	6.144e-4	2.806e-4	1.794e-5

TABLE 8 Measured interpolation overshoots for the solid body rotation case. The data for the absolute overshoot are calculated from solution values exceeding the analytic bounds $[0, 1]$. The overshoots relative to the fine solution are calculated at the same time.

5 | RESULTS AND DISCUSSION

The verified HO-ACR algorithm is applied to solve two types of cases. In the first case, a H_2 bubble is advected through a strong shock in 2D. This presents numerical challenges, including properly tracking of the thin H_2 - O_2 flame fronts and the strong discontinuities of shock waves. The second case is the C_3H_8 -air combustion in a bluff body combustor in which the chemistry is extremely stiff. In both these cases, the HO-ACR method is applied to invalid ghost cells and newly refined cells with no fundamental changes to the base algorithm. The investigation has demonstrated that the HO-ACR method has played a critical role in ensuring the stability and accuracy in the process of obtaining solutions. Immediately what follows is the description of each problem. Then, results are presented and discussed.

5.1 | H₂-O₂ Shock Induced Combustion

The configuration is shown in Fig. 19 and follows the setup used in a previous study⁷. A Mach 2 steady planar shock is located 0.75 cm to the right of the origin and parallel to the y -axis. A hydrogen bubble is located just upstream of the shock. The upper and lower boundaries of the y -direction are periodic, while x -min boundary condition is Dirichlet and the x -max is extrapolated. The initial velocities $U_u = 1.24 \times 10^5$ cm/s and $U_s = 4.34$ cm/s are the velocities upstream and downstream of the shock. The hydrogen mass fraction is initialized to

$$c_{\text{H}_2} = \frac{1}{2} \left[1 + \tanh \left(\frac{r_c - r}{C_2} \right) \right], \quad r = \sqrt{(x - x_0)^2 + (y - y_0)^2}, \quad (32)$$

with r_c the radius of the bubble and the coordinate (x_0, y_0) the center of the bubble. The coefficient C_2 determines the sharpness of the interface between the H₂ bubble and the surrounding air. The case parameters are $C_2 = 3 \times 10^{-3}$ cm⁻¹, $r_c = 0.28$ cm, $(x_0, y_0) = (0.4, 0.75)$ cm, and $\lambda = 1.5$ cm. The mass fractions for the surrounding air are simplified to $c_{\text{N}_2} = 0.767$ and $c_{\text{O}_2} = 0.233$. The H₂-O₂ combustion is a 9-species, 19-reaction mechanism described by Billet et al.²⁵.

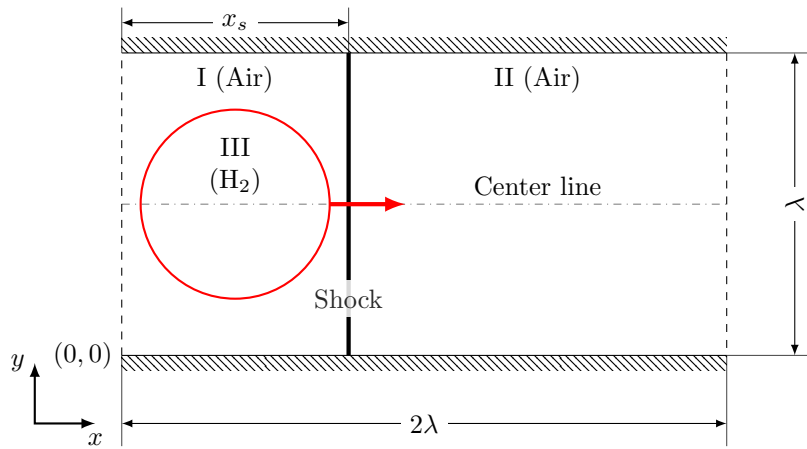


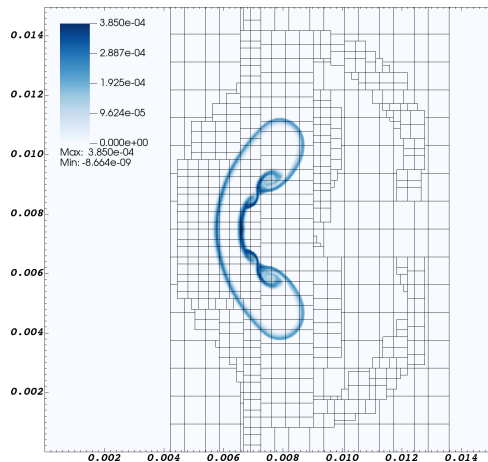
FIGURE 19 Shock bubble case configuration.

A traditional approach for this case would use AMR to completely refine around the discontinuities. Instead, an imperfect AMR tagging scenario that is representative of what might occur in an embedded-DNS simulation is considered. Assume only a small window of refinement can be afforded in the solution process, and due to this limitation significant features will cross AMR interfaces. Additionally, refinement is only added once the flow is developed. Ideally, the HO-ACR algorithm will handle this imperfect tagging; that is improve accuracy over a solution that uses no AMR whatsoever while maintaining solution stability.

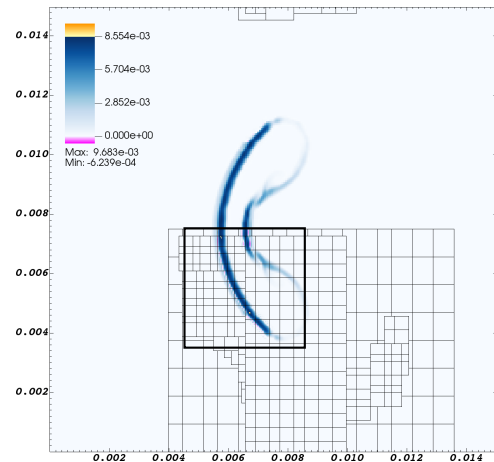
In total, three simulations are performed. Two cases are run for comparison: one with base Chord (without the HO-ACR method), and the other with the HO-ACR method. The base mesh has 512×256 cells. Both cases begin with only the base mesh until $t = 5.3 \mu\text{s}$. Then two AMR levels with a refinement ratio of 2 are added based on the gradients of density and pressure for the lower half of the domain, with a small buffer wrapping around the periodic boundary. For comparison, the third case is run by using AMR on the entire domain from the initial solution as traditional applied, which is considered as the “ideal” case. The goal here is to demonstrate that the HO-ACR method enables us to obtain a quality solution even if imposing AMR in arbitrary space-time and crossing strong discontinuities. This strategy may be performed for embedded-DNS or embedded-LES where only a portion of the space-time mesh is appropriately refined due to computational expense.

Fig. 20 shows the solution of H₂O at $t = 5.3 \mu\text{s}$. In the figure, there are 4 subfigures: (a) with the “ideal” case where AMR has been used for the entire run; (b) showing an initially coarse solution from Chord with refinement added; (c) and (d) comparing the closeups of solutions between Chord without and with the HO-ACR algorithm for the domain regridded on the bottom half, as indicated by the black rectangles in the prior subfigures. After the initial refinement as seen in the subfigure (c), negative species mass fractions, as indicated by magenta, are large, and positive overshoots beyond the coarse solution are shown in yellow-orange. While mass fractions of species must be always positive to be physically meaningful, Chord does have a small tolerance for negative mass fractions to cope with the numerical instability in the nonlinear solvers employed during the solution process¹⁰.

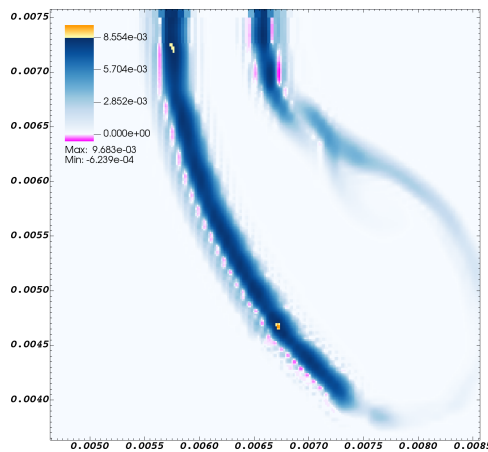
In the subfigure (d), the magnitude of the negative mass fractions is reduced by an order of magnitude, which is sufficient to have kept the overall algorithm stable and allow the solution to develop properly. Note that for intermediate chemistry-dependent operations, Chord first renormalizes the mass fractions to avoid nonphysical quantities. The correction first bounds each species mass fraction between 0 and 1 and then renormalizes so that the sum of all species mass fractions equals one²⁶. At the end of each time step, a similar procedure is applied to the conserved solution state except that the bounds are expanded by a small tolerance. We find the correction process handles small nonphysical perturbations from sources such as time integration well, but large errors from grid refinement are too drastic to correct, motivating the need for the HO-ACR method. Overall, reducing nonphysical overshoots at the source (in this case, from interpolation), vastly mitigates the impact of renormalization on the behavior of the dynamical system. Although not the focus of this work, balancing chemistry is an important discussion, and we refer readers to works by Christopher et al.¹⁰ and Owen et al.²⁷ for details of the approach Chord uses.




(a) Full AMR the entire run time.



(b) Solution regridded on half the domain at current time, using Chord.



(c) Close-up view of the boxed region in (b) for the solution using Chord.

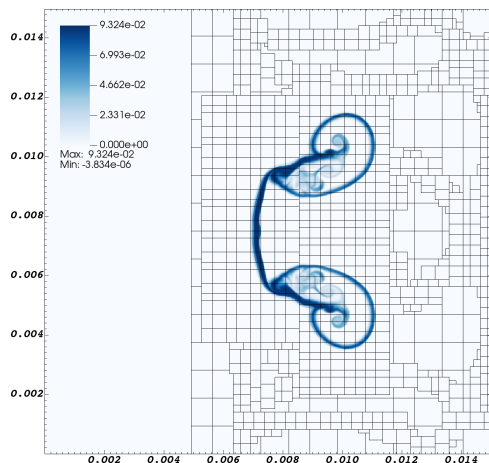


(d) Close-up view of the boxed region in (b) for the solution using Chord with HO-ACR.

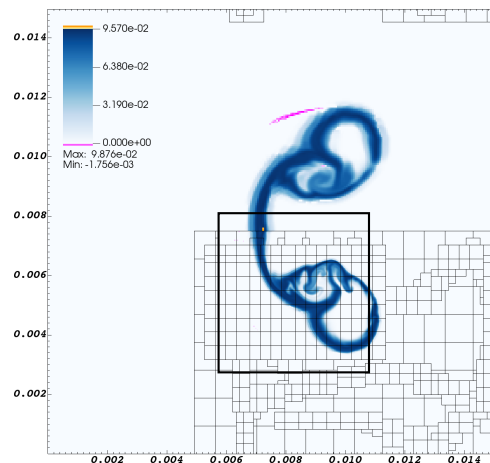
FIGURE 20 Solution of H_2O at $t = 5.3 \mu\text{s}$. In subfigures (b)-(d) negative species mass fractions are shown in magenta, and positive overshoots beyond the coarse solution are shown in yellow-orange.

Fig. 21 shows the H_2O after continuing to evolve the solution to $t = 10 \mu\text{s}$. The layout of this figure is the same as Fig. 21. Although not matching the “ideal” solution, the partially added refinement is seen to begin capturing fine scale features not visible in the coarse solution while using significantly less computational resources. This provides an ability to assess the validity of embedded-DNS or LES, although we refrain from making judgment here. Instead, we focus on the quality of the interpolations.


The magnitude of the solution difference with and without the HO-ACR method grows smaller in comparison to that of the earlier time. There is a significant overshoot along the AMR interface that is effectively prevented by the HO-ACR algorithm. Despite what may appear to be minor differences, this is of significant impact for chemically reacting flows with stiff reactions where minor errors can push the solution to unphysical states that are unsolvable. Running further in time, this overshoot would eventually destabilize the solution if the HO-ACR method is not applied.



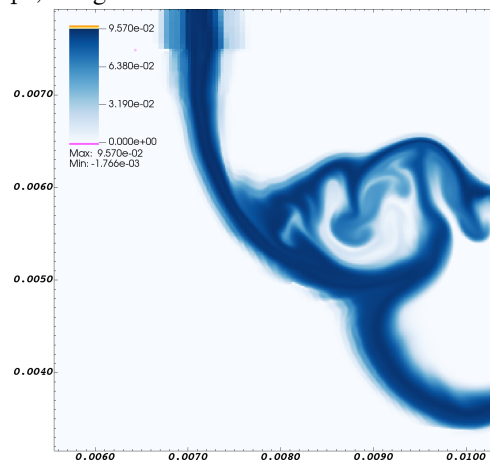
(a) With full AMR the entire run time.



(b) Current solution after AMR added on half the domain at $t = 5.3 \mu\text{s}$, using Chord.



(c) Close-up view of the boxed region in (b) for the solution using Chord.



(d) Close-up view of the boxed region in (b) for the solution using Chord with HO-ACR.

FIGURE 21 Solution H_2O at $t = 10 \mu\text{s}$. In subfigures (b)-(d) negative species mass fractions are shown in magenta, and positive overshoots beyond the coarse solution of the HO-ACR case (d) are shown in yellow-orange.

5.2 | C_3H_8 -Air Combustion in a Bluff Body Combustor

The premixed combustion of C_3H_8 -air is simulated in a bluff body combustor^{28,29}. As shown in Fig. 22, the triangular bluff body lies in a rectangular channel. At the inlet, a gaseous mixture of 4.01% C_3H_8 , 22.36% O_2 , and 73.64% N_2 by mass fraction flows into the domain at a velocity of 15.7 m/s. The mixture has a temperature of 310 K and a pressure of 101,325 Pa. The flow has a bulk Mach number of 0.053 and bulk Reynolds number of 50,000. The reaction mechanism for C_3H_8 -air is a stiff system including 25-species and 66-reactions³⁰. To ensure stability, only the ACR method is explored for this application and the high-order extension is disabled. We revert to the low order (strictly bound-preserving) ACR method for this case, because

the reacting species are highly sensitive to any overshoots allowed by the HO-ACR method. The initial conditions in the domain




FIGURE 22 The geometry of the bluff body combustor.

are set to the same values as the inlet, with the exception of a small region surrounding the bluff body. Near the bluff body, the gas mixture is initialized to 12% CO₂, 6.54% H₂O, 5.14% O₂, and 73.62% N₂, with a temperature of 1300 K. This hot spot provides the ignition of the combustion.




FIGURE 23 Initial grid boxes in physical space for the bluff body combustor.

A base mesh of 11 008 cells are used, with two levels of refinement at a ratio of 2. The 2nd level is fixed behind the bluff body and spans the experimental observation window. The 3rd level dynamically tracks the flame and boundary layers. The bluff body case uses the conforming mapped multiblock technique to fit a structured mesh to the geometry. The initial grid boxes are shown in Fig. 23. Species OH and CH₂O are good flame markers. Therefore, during the run time, cells defining the 3rd level are tagged based on $c_{OH} \times c_{CH_2O} > 2 \times 10^{-9}$.

Fig. 24 shows the CH₂O contour in the region immediately downstream of the trailing edge of the bluff body. On the left, the solution obtained from base Chord after regriding displays small artifacts arising from interpolation overshoots as highlighted by the violet bubbles. On the right, the presence of these noises in the solution is diminished greatly by instead applying the ACR method. Similarly, Fig. 25 compares the OH mass distributions between the base and the ACR algorithms after regriding for the same physical location and solution time. While the differences in OH are barely visible in Fig. 25, after 80 time steps, these grow into quite large perturbations as shown in Fig. 26. The effectiveness of using the ACR to suppress the interpolation overshoots is clearly seen in Fig. 26, as highlighted by gray bubbles. Through the numerical experiments, the ACR algorithm is found to be necessary to ensure the overall simulation stability.

With the ACR method, a full simulation of the three-dimensional version of the bluff body case was performed. The z-direction is defined with width of 76.2 mm, and enforced with periodic boundaries. The base mesh is approximately 330 000 cells. One additional level with a refinement factor of 2 is added. Fig. 27 shows the instantaneous isosurfaces of the vorticity, temperature, mass fractions of H₂O and OH, respectively, for the 3D C₃H₈-air flame at a solution time of $t = 246$ ms. Without using AMR, it is much more expensive to capture the flow and flame dynamics near the bluff body as shown in Fig. 27. Furthermore, without using the ACR method, this 3D combustion simulation has numerical difficulty where the flame crosses AMR interfaces. One could refine the flame everywhere in the domain to avoid numerical issues from AMR, but this would require significantly more computing resources for little benefit – high fidelity is only required in the experimental observation window.

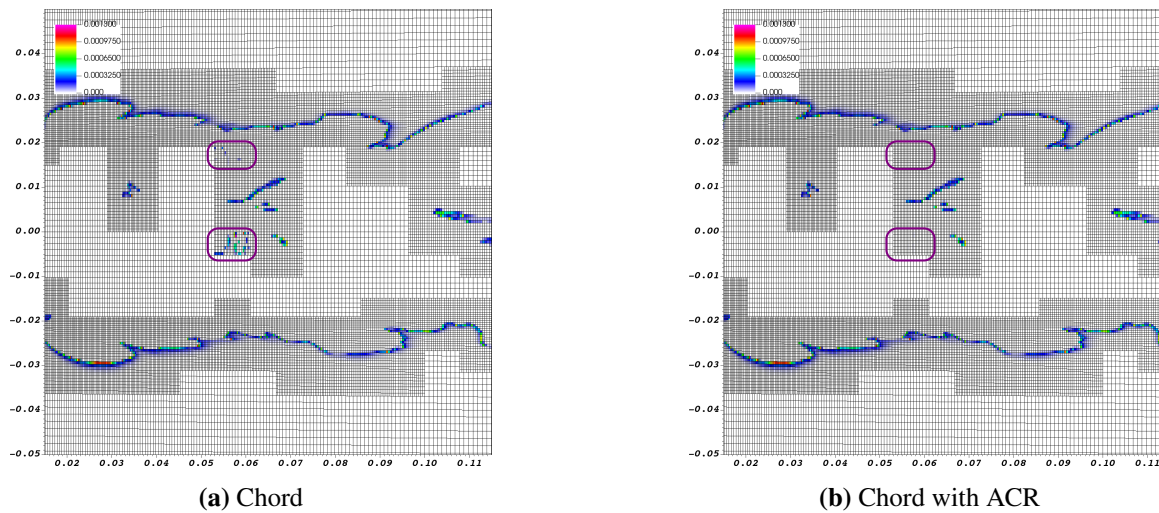


FIGURE 24 Solution of CH_2O after regridding. Small artifacts from interpolation overshoot are present in the base algorithm, but suppressed using the ACR method in regions highlighted by the violet bubbles.

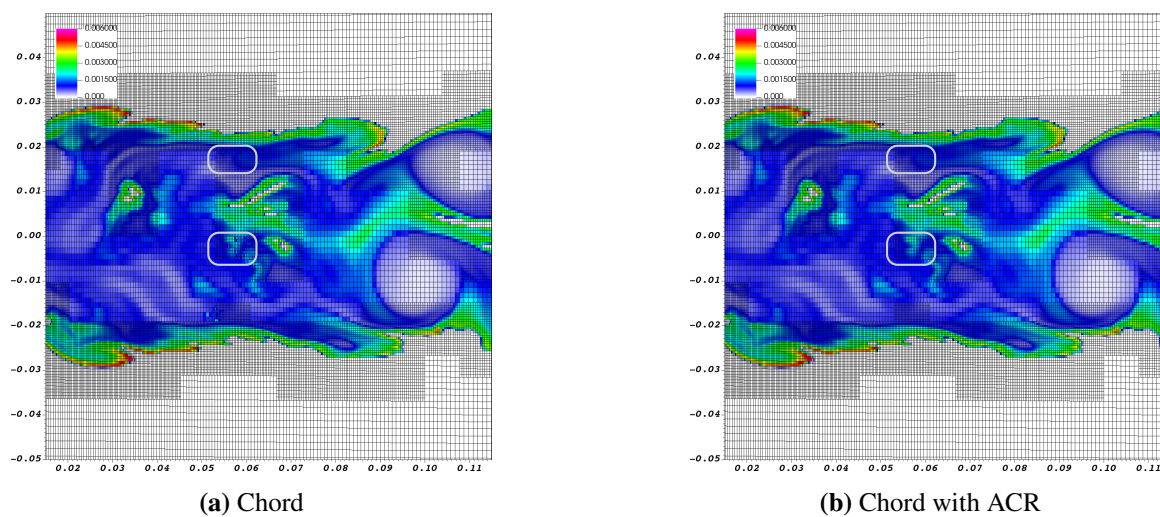


FIGURE 25 Solution of OH after regridding. Small artifacts from interpolation overshoot are present in the base algorithm, but suppressed using the ACR method in regions highlighted by the light gray bubbles.

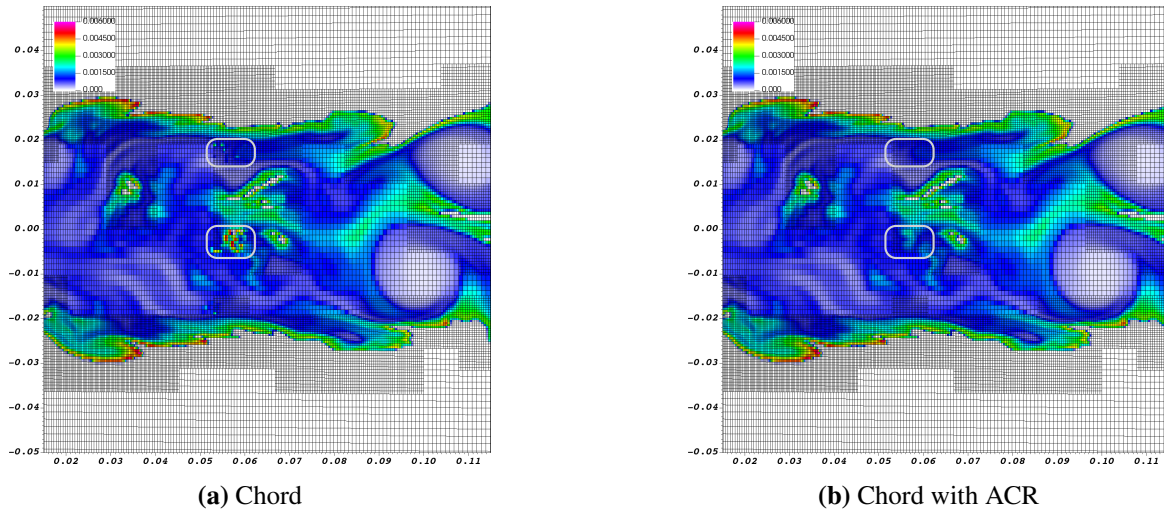


FIGURE 26 Solution of OH run 80 time steps after regridding. The originally small artifacts have developed into significant erroneous features in the flow. After this point the chemistry of the base algorithm is no longer numerically stable, while solution using the ACR method encounters no such issues.




FIGURE 27 Vorticity, temperature, $c_{\text{H}_2\text{O}}$, and c_{OH} of the 3D Bluff Body C_3H_8 -air case at solution time $t = 246$ ms.

6 | CONCLUSIONS AND FUTURE WORK

In the present study, the HO-ACR algorithm has been implemented, verified, and validated in a fourth-order finite-volume CFD software solving complex fluid dynamics problems with the presence of chemical reactions, shock waves, and realistic geometries. For such problems, AMR is essential for capturing multiscale flow features with computational efficiency. However, the interpolation operation required by the AMR process often induces unwanted overshoots near discontinuities which destabilizes the overall simulation, particularly for flows with combustion. This is one of the main motivations for the development of the HO-ACR algorithm. The HO-ACR method has been verified to ensure high-order accuracy for smooth flow problems, preserve bounds for non-smooth ones, and maintain conservation on mapped grids. On challenging problems such as the shock-induced H_2-O_2 combustion, the HO-ACR method ensures better bounds-preservation in solutions even if a fine level is arbitrarily imposed across or near discontinuities. Furthermore, the application of the method to a C_3H_8 -air premixed flame in the bluff body combustor has demonstrated that the engagement of HO-ACR is imperative to ensure the solution stability. Note that the chemical kinetics used in this combustion problem, involving 24-species and 66-reactions, is very stiff and the bluff body geometry increases the stiffness through the nonlinear interactions between the fluid and flame dynamics.

Future work will be focused on extending applications of the HO-ACR method to the computational geometries represented by general mapped multi-block domains. This would require having some form of bound-preserving property in the interpolation operation across blocks, where the grid metrics may also be discontinuous. Interpolation between blocks requires a somewhat different approach for bounds preservation. Interpolation of ghost cells between different meshes can not reasonably have a conservation constraint, but with careful consideration may use clipping and redistribution for maintaining bounds across different mappings.

7 | ACKNOWLEDGMENTS

We would like to thank the Department of Defense High Performance Computing Modernization Program for their compute time and support. This research was supported by Department of Defense United States Air Force (DOD-USAF-AirForce) under the award number FA9550-18-1-0057 and by National Science Foundation under the award number 1723191. Johansen is supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics Program under contract No. DE-AC02-05CH11231.

References

1. Guzik SM, Gao X, Owen LD, McCorquodale P, Colella P. A Freestream-Preserving Fourth-Order Finite-Volume Method in Mapped Coordinates with Adaptive-Mesh Refinement. *Comput. Fluids* 2015; 123: 202–217.
2. Holgate J, Skillen A, Craft T, Revell A. A Review of Embedded Large Eddy Simulation for Internal Flows. *Arch. Comput. Method. E.* 2018: 1-18.
3. Gao X, Guzik SMJ, Colella P. A Fourth-Order Boundary Treatment for Viscous Fluxes on Cartesian Grid Finite-Volume Methods. AIAA 2014-1277, 52nd AIAA Aerospace Sciences Meeting; : 2014.
4. Guzik SM, Gao X, Olschanowsky C. A high-performance finite-volume algorithm for solving partial differential equations governing compressible viscous flows on structured grids. *Comput. Math Appl.* 2016; 72: 2098–2118.
5. Gao X, Wang Y, Overton N, May I. Estimation of Flame Speed Model Parameter Using Ensemble Kalman Filter Algorithm. In; ; 2015. WSSCI 2015-134IE-0013.
6. Gao X, Owen LD, Guzik SMJ. A Parallel Adaptive Numerical Method with Generalized Curvilinear Coordinate Transformation for Compressible Navier-Stokes Equations. *Int. J. Numer. Meth. Fluids* 2016; 82: 664–688.
7. Owen L, Guzik S, Gao X. A Fourth-Order Finite-Volume Algorithm for Compressible Flow with Chemical Reactions on Mapped Grids. 2017(AIAA 2017-4498). doi: 10.2514/6.2017-4498

8. Owen LD, Guzik SM, Gao X. A high-order adaptive algorithm for multispecies gaseous flows on mapped domains. *Comput. Fluids* 2018; 170: 249–260.
9. Walters S, Guzik S, Gao X. Evaluation of the Stretched-Vortex Subgrid-Scale Model for Large-Eddy Simulation with a Fourth-Order Finite Volume Algorithm. In: No. AIAA 2019-1886. 2019 AIAA SciTech Forum. ; 2019. <https://doi.org/10.2514/6.2019-1886>.
10. Christopher J, Guzik SM, Gao X. High-order implicit-explicit additive Runge–Kutta schemes for numerical combustion with adaptive mesh refinement. *Int. J. Numer. Meth. Fl.* 2022; 94(8): 1082-1110. doi: 10.1002/fld.5084
11. Adams M, Colella P, Graves DT, et al. Chombo Software Package for AMR Applications—Design Document. Tech. Rep. LBNL-6616E, Lawrence Berkeley National Laboratory; ; 2015.
12. Van Leer B. Towards the ultimate conservative difference scheme. IV. A new approach to numerical convection. *J. Comput. Phys.* 1977; 23(3): 276-299.
13. Colella P, Sekora M. A limiter for PPM that preserves accuracy at smooth extrema. *J. Comput. Phys.* 2008; 227(15): 7069–7076.
14. Jiang G, Shu CW. Efficient implementation of weighted ENO schemes. *J. Comput. Phys.* 1996; 126: 202–228.
15. Michalak C, Ollivier-Gooch C. Accuracy Preserving Limiter for the High-Order Accurate Solution of the Euler Equations. *J. Comput. Phys.* 2009; 228(23): 8693–8711. doi: 10.1016/j.jcp.2009.08.021
16. Anderson R, Dobrev V, Kolev T, et al. High-order local maximum principle preserving (MPP) discontinuous Galerkin finite element method for the transport equation. *J. Comput. Phys.* 2017; 334: 102 - 124. doi: 10.1016/j.jcp.2016.12.031
17. Ivan L, De Sterck H, Susanto A, Groth C. High-order central ENO finite-volume scheme for hyperbolic conservation laws on three-dimensional cubed-sphere grids. *J. Comput. Phys.* 2015; 282: 157 - 182. doi: 10.1016/j.jcp.2014.11.002
18. Freret L, Ivan L, Sterck H, Groth CP. High-Order Finite-Volume Method with Block-Based AMR for Magnetohydrodynamics Flows. *J. Sci. Comput.* 2019; 79(1): 176–208. doi: 10.1007/s10915-018-0844-1
19. Hilditch J, Colella P. A projection method for low Mach number fast chemistry reacting flow. Paper AIAA 97-0263, AIAA; ; 1997.
20. McCorquodale P, Colella P. A high-order finite-volume method for conservation laws on locally refined grids. *Comm. App. Math. Comput. Sci.* 2011; 6(1): 1–25.
21. Guzik SM, McCorquodale P, Colella P. A Freestream-Preserving High-Order Finite-Volume Method for Mapped Grids with Adaptive-Mesh Refinement. In: No. AIAA 2012-0574. AIAA. ; 2012. <https://doi.org/10.2514/6.2012-574>.
22. Barth T, Frederickson P. High-order solution of the Euler equations on unstructured grids using quadratic reconstruction. Paper AIAA-90-0013, AIAA; ; 1990.
23. Guzik SMJ, Gao X, Weisgraber T, Alder B, Colella P. An Adaptive Mesh Refinement Strategy with Conservative Space-Time Coupling for the Lattice-Boltzmann Method. In: No. AIAA 2013-0866. AIAA. ; 2013. <https://doi.org/10.2514/6.2013-866>.
24. Anderson RW, Dobrev VA, Kolev TV, Rieben RN. Monotonicity in high-order curvilinear finite element arbitrary Lagrangian–Eulerian remap. *Int. J. Numer. Meth. Fl.* 2015; 77(5): 249-273. doi: 10.1002/fld.3965
25. Billet G, Giovangigli V, Gassowski dG. Impact of volume viscosity on a shock–hydrogen-bubble interaction. *Combust. Theory Model.* 2008; 12(2): 221-248.
26. Owen LD. *A Fourth-Order Solution-Adaptive Finite-Volume Algorithm for Compressible Reacting Flows on Mapped Domains*. Ph.D. dissertation. Colorado State University, ; 2019.
27. Owen LD, Gao X, Guzik SM. Techniques for improving monotonicity in a fourth-order finite-volume algorithm solving shocks and detonations. *J. Comput. Phys.* 2020; 415.

28. Paxton B, Fugger C, Rankin B, Caswell A. Development and characterization of an experimental arrangement for studying bluff-body-stabilized turbulent premixed propane-air flames. In: No. AIAA 2019-0118. AIAA Scitech 2019 Forum. ; 2019. <https://doi.org/10.2514/6.2019-0118>
29. Comer A, Ihme M, Li C, et al. Proceedings of the Third Model Validation for Propulsion (MVP 3) Workshop.; 2019. AIAA SciTech Forum.
30. Zettervall N, Nordin-Bates K, Nilsson EJK, Fureby C. Large Eddy Simulation of a premixed bluff body stabilized flame using global and skeletal reaction mechanisms. *Combust. Flame* 2017; 179: 1–22.

How to cite this article: Nathaniel Overton-Katz, Xinfeng Gao, Hans Johansen, and Stephen M. Guzik (2022), Adaptive Clipping-and-Redistribution Algorithms for Bounded and Conservative High-Order Interpolations Applied to Discontinuous and Reactive Flows, *Int. J. Numer. Meth. Fl.*, .