

Spectral Thresholds for State Redistribution and Defect Subspaces in Cut-Cell Finite-Volume Operators

Justo E. Karell*
Stevens Institute of Technology

September 8, 2026

Mathematics Subject Classification (2020): 65M12, 65M08, 76M12

Abstract

In finite-volume methods on embedded-boundary meshes, arbitrarily small cut cells can arise near complex boundaries. Let $\alpha \in (0, 1]$ denote the cut-cell fraction relative to a regular cell. Because the finite-volume flux balance is divided by cell volume, a time step chosen for the regular grid can produce $O(\alpha^{-1})$ coefficients in the cut-cell update, creating the small-cell problem. For a one-dimensional periodic upwind discretization, we show that this instability is carried by a discrete eigenmode that localizes at the cut cell. We then blend the unstabilized update with a two-cell merge-left redistribution operator, which replaces the provisional values in the cut-cell neighborhood by their volume-weighted average, and derive the local spectral stability threshold associated with the localized mode. If L denotes the Courant–Friedrichs–Lewy (CFL) number based on the cut-cell width, then for fixed $L > 2$ the threshold redistribution fraction approaches $(L - 2)/(L - 1)$ as the relative cut-cell width tends to zero. We also prove that the redistribution correction becomes asymptotically aligned with the unstable mode and that analogous localization occurs for a broader class of cut-cell operators. In a second-order finite-volume benchmark, the model-derived partial redistribution preserves the observed convergence rate while reducing smooth-solution error. For the sine-wave test, the model-derived choice reduces the interpolated grid-cell count required to reach a one-period discrete L^1 error of 10^{-3} by 13.3%–27.1% relative to full redistribution. The corresponding estimates are 282 versus 325 cells for $\alpha = 0.10$ and 248 versus 340 cells for $\alpha = 0.15$. In separate 100-period higher-order tests, every $s = s_0$ and $s = 1$ run completes the integration, with measured mass drift below 8×10^{-14} .

1 Introduction

Embedded-boundary and cut-cell methods represent complex physical boundaries while retaining a simple background mesh. Their main explicit-time-stepping difficulty is geometric. If a boundary leaves only a fraction $\alpha \in (0, 1]$ of a background control volume, then a conservative finite-volume update divides the net numerical flux by this reduced volume. A time step that satisfies the background Courant–Friedrichs–Lewy (CFL) condition can therefore produce $O(\alpha^{-1})$ coefficients in the cut-cell row of the update matrix.

Several methods remove this small-cell restriction. Flux redistribution [5], h-box methods [4], explicit-implicit updates [10, 11], and domain-of-dependence stabilization [7, 13, 9, 12] modify the

*Department of Computer Science, Stevens Institute of Technology, Hoboken, NJ, USA. justokarell@gmail.com.

local update or its effective domain of dependence. State redistribution (SRD) instead forms a provisional finite-volume update, merges small cells with neighboring cells into larger neighborhoods, and redistributes the merged states while preserving conservation [2]. Weighted SRD reduces dissipation through geometry-dependent weights [8]; Berger and Giuliani [3] further construct a weighted SRD scheme with monotonicity, total-variation-diminishing, and GKS-stability results in the configurations covered by their analysis. Update-Magnitude SRD (UM-SRD) blends redistribution according to the magnitude of the provisional update [14]. The scalar calibration studied here is distinct from these weight constructions: the redistribution operator is held fixed and only its overall blend with the base update is varied.

A closely related stability analysis is Berger’s study of cut-cell merging [1]. For one-dimensional linear advection with first-order upwinding, that work uses GKS stability theory to quantify the admissible time step after cell merging and shows that the usual cut-cell CFL restriction can be overly conservative. Berger and Giuliani [3], by contrast, construct geometry-dependent weighted SRD coefficients and prove stability properties of that weighted algorithm. The question here is different from both: the base update and the redistribution operator are held fixed, only the scalar blend between them is varied, and the blend is selected by following the localized eigenvalue of the full discrete operator to the unit circle. The resulting quantity is therefore a spectral calibration of redistribution strength, not a time-step bound or a new geometric weighting formula.

Let

$$\overline{\mathbb{D}} := \{z \in \mathbb{C} : |z| \leq 1\}$$

denote the closed unit disk. We consider the following calibration problem: given a finite-volume operator with a localized eigenvalue outside $\overline{\mathbb{D}}$ and a specified redistribution operator, determine the scalar blend at which that eigenvalue re-enters $\overline{\mathbb{D}}$. We analyze this problem for positive-speed scalar linear advection on a one-dimensional periodic mesh of N cells with one cut cell, discretized by first-order upwind with merge-left SRD as the redistribution operator. Let N denote the total number of cells and let λ denote the background CFL number. Define the local cut-cell CFL number by

$$L = \lambda_\alpha = \frac{\lambda}{\alpha}.$$

The unstabilized operator has a localized eigenvalue near $1 - L$, so $L > 2$ places that eigenvalue outside $\overline{\mathbb{D}}$. For a square matrix M , write $\sigma(M)$ for its spectrum and

$$\rho(M) := \max_{\mu \in \sigma(M)} |\mu|.$$

Throughout the paper, *spectral stability* means $\rho(A_s) \leq 1$. Because A_s is generally nonnormal, power boundedness and finite transient amplification remain separate questions. Let A denote the base update matrix and A_{SRD} the fully redistributed update matrix. We blend them as

$$A_s = (1 - s)A + sA_{\text{SRD}}, \quad 0 \leq s \leq 1,$$

where s is the redistribution fraction: $s = 0$ gives the base update and $s = 1$ gives full redistribution. We then derive the full N -cell characteristic equation. For fixed $L > 2$ and $N \geq 4$, the localized eigenvalue crosses the value -1 at

$$s_{-1}^{(N)}(\alpha, L) = \frac{L - 2}{L - 1} - \frac{L(L - 2)}{2(L - 1)^2} \alpha^2 + O(\alpha^3).$$

All other nonconstant eigenvalues are inside the unit disk near this crossing for sufficiently small α . Consequently,

$$s_0(L) = \frac{L - 2}{L - 1}$$

is asymptotically sufficient and differs from the exact local spectral stability threshold by $O(\alpha^2)$. The contribution is a derivation of the redistribution fraction from the localized unstable branch of the discrete spectrum.

The unstable eigenvector converges to the cut-cell coordinate vector as $\alpha \rightarrow 0$, so the unstable mode is concentrated at the geometric defect. Merge-left SRD changes the state in an asymptotically aligned direction. Consequently, varying s directly changes the amplification factor of the localized branch used to define the spectral threshold. Eigenvalue spectra of embedded-boundary discretizations have previously been used as a stability diagnostic [6]; here the characteristic equation identifies a localized amplification mode and its spectral crossing analytically.

Two asymptotic regimes are used and are kept separate throughout. In the *fixed-background-CFL regime*, λ is fixed while $\alpha \rightarrow 0$, so $L = \lambda/\alpha \rightarrow \infty$. This is the regime used for the singular-eigenvalue and defect-subspace analyses. In the *fixed-local-CFL regime*, L is fixed while $\alpha \rightarrow 0$, so $\lambda = \alpha L \rightarrow 0$. This is the regime in which the spectral-threshold expansion is proved. Later higher-order experiments use finite α and fixed background CFL; there $s_0(L)$ is tested as a model-derived rule, not invoked as a theorem for the higher-order scheme.

At fixed regular-grid CFL λ , the local CFL is $L = \lambda/\alpha$, so the leading prescription becomes

$$s_0(\lambda/\alpha) = \frac{\lambda - 2\alpha}{\lambda - \alpha} = 1 - \frac{\alpha}{\lambda - \alpha} \rightarrow 1 \quad (\alpha \rightarrow 0).$$

Thus the difference $1 - s_0$ from full redistribution is largest for moderately restrictive cut cells. For arbitrarily small cut cells at a fixed regular-grid time step, s_0 approaches 1.

We next identify which parts of the model calculation extend beyond one cut cell. After re-ordering the unknowns into defect and regular coordinates, a class of cut-cell operators contains a singular finite-dimensional block. We prove that its limiting defect matrix controls the unbounded eigenvalues and that the corresponding invariant subspace converges to the defect-coordinate subspace. This separates the structural localization result from the model-specific formula for the spectral threshold.

The main contributions are:

1. an explicit characteristic-polynomial analysis of the one-cut-cell periodic upwind operator, including the singular eigenvalue and eigenvector localization;
2. a definition of merge-left SRD on the same full operator and a proof that its correction becomes asymptotically aligned with the localized unstable direction in the volume-weighted inner product;
3. a full- N derivation of the local spectral stability crossing and the leading-order prescription $s_0(L) = (L - 2)/(L - 1)$ for $L > 2$;
4. a singular-block theorem relating a finite-dimensional defect matrix to the unbounded spectrum and localized invariant subspace, with multi-cut-cell, multidimensional, and finite-volume consequences; and
5. numerical evidence from a separate second-order benchmark that the model-derived partial redistribution preserves the observed convergence order, lowers smooth-solution error, and reaches a prescribed smooth-solution error with a smaller interpolated total cell count than full redistribution.

The threshold theorem applies to the periodic first-order upwind model with merge-left SRD. Extension to nonlinear conservation laws, systems, higher-order reconstructions, adjacent cut cells,

or multidimensional SRD neighborhoods requires deriving the corresponding defect block and tracking the action of the chosen redistribution on the defect-associated spectrum.

The paper follows the argument in that order. Section 2 defines the model problem and derives the localized unstable mode. Section 3 defines merge-left SRD, proves its alignment with the mode, and derives the local spectral stability threshold for the full periodic operator. Section 4 develops the abstract defect-subspace framework. Sections 5–8 give multi-cut-cell, multidimensional, and finite-volume consequences. Section 9 then visualizes the spectral mechanism and tests the model-derived partial choice in a second-order benchmark.

2 One-Dimensional One-Cut-Cell Model

2.1 Model problem and update operator

Consider the scalar advection equation

$$u_t + au_x = 0, \quad a > 0, \quad (1)$$

on a periodic one-dimensional domain discretized into N cells indexed periodically by $i = 0, \dots, N-1$. The periodic setting isolates the small-volume singularity from physical boundary-condition effects; it is a model operator for the cut-cell mechanism rather than a geometrically literal one-dimensional embedded boundary. Let h denote the width of a regular background cell, let Δt be the time step, and define the background CFL number

$$\lambda = \frac{a\Delta t}{h}, \quad 0 < \lambda \leq 1.$$

All cells have length h except one cut cell c , whose length is αh with $\alpha \in (0, 1)$. Let U_i^n denote the cell average in cell i at time t^n , and let $U^n = (U_0^n, \dots, U_{N-1}^n)^\top$. Following the notation of Berger and Giuliani for SRD, we denote the *provisional* base-scheme update by $\hat{U}$. The positive-speed first-order upwind step is

$$\hat{U}_i = U_i^n - \frac{a\Delta t}{|C_i|} (U_i^n - U_{i-1}^n), \quad |C_i| = \begin{cases} \alpha h, & i = c, \\ h, & i \neq c. \end{cases}$$

Thus $\hat{U} = AU^n$, where the provisional update matrix A is defined by

$$(AU)_i = (1 - \lambda_i)U_i + \lambda_i U_{i-1}, \quad \lambda_i = \begin{cases} \lambda_\alpha := \lambda/\alpha, & i = c, \\ \lambda, & i \neq c, \end{cases} \quad (2)$$

with periodic indexing. Let $e_i \in \mathbb{R}^N$ denote the i th canonical coordinate vector. The only singular row as $\alpha \rightarrow 0$ is the cut-cell row. If A_{bg} denotes the uniform periodic upwind matrix at CFL λ , then

$$A = A_{\text{bg}} + (\lambda_\alpha - \lambda)e_c(e_{c-1} - e_c)^\top. \quad (3)$$

Equation (2) is the algebraic form of the small-cell problem. The same flux difference that is divided by h on a regular cell is divided by αh on the cut cell. Thus the local CFL number is $\lambda_\alpha = \lambda/\alpha$. The analysis below determines how this single singular row changes the spectrum.

2.2 Characteristic polynomial and the localized root

By periodic translation we may set $c = 0$. Let $Av = \mu v$. On every regular row $i = 1, \dots, N-1$,

$$(\mu - 1 + \lambda)v_i = \lambda v_{i-1}.$$

For $\mu \neq 1 - \lambda$, define

$$q = \frac{\lambda}{\mu - 1 + \lambda}. \quad (4)$$

The excluded value $\mu = 1 - \lambda$ is not an eigenvalue. Indeed, the regular-row equations then imply

$$v_0 = \dots = v_{N-2} = 0,$$

and the cut-cell row subsequently gives $v_{N-1} = 0$. Thus every eigenvalue is represented by the q -parametrization below.

Then

$$v_i = q^i v_0, \quad i = 1, \dots, N-1, \quad (5)$$

and

$$\mu(q) = 1 - \lambda + \frac{\lambda}{q}. \quad (6)$$

Substituting $v_{N-1} = q^{N-1}v_0$ into the cut-cell row gives

$$P_\alpha(q) := q^N - (1 - \alpha)q - \alpha = 0. \quad (7)$$

Conversely, every nonzero root of (7) gives an eigenvalue through (6). Since $P_\alpha(0) = -\alpha$, all roots are nonzero for $\alpha > 0$.

The localized root lies near the zero of the linear part of P_α . Set

$$\tilde{q} = -\frac{\alpha}{1 - \alpha}.$$

Writing $q = \tilde{q} + \eta$ and using $(1 - \alpha)\tilde{q} + \alpha = 0$ gives

$$\eta = \Phi_\alpha(\eta) := \frac{(\tilde{q} + \eta)^N}{1 - \alpha}.$$

On the disk

$$|\eta| \leq r_\alpha := \frac{2|\tilde{q}|^N}{1 - \alpha},$$

we have $r_\alpha = O(\alpha^N) = o(|\tilde{q}|)$. Therefore, for sufficiently small α , Φ_α maps the disk into itself and

$$\sup_{|\eta| \leq r_\alpha} |\Phi'_\alpha(\eta)| = O(\alpha^{N-1}) < 1.$$

The contraction mapping theorem gives a unique root in this disk:

$$q_u = -\frac{\alpha}{1 - \alpha} + O(\alpha^N). \quad (8)$$

Theorem 2.1 (Unique unbounded eigenvalue). *Fix $N \geq 3$ and $\lambda \in (0, 1]$. As $\alpha \rightarrow 0$, the matrix A has a unique unbounded eigenvalue*

$$\mu_u = 1 - \frac{\lambda}{\alpha} + O(\alpha^{N-2}), \quad (9)$$

while all other eigenvalues remain bounded.

Proof. The root (8) gives

$$\frac{1}{q_u} = -\frac{1-\alpha}{\alpha} + O(\alpha^{N-2}),$$

and substitution into (6) yields (9).

It remains to show that no second root approaches zero. Choose a fixed $C > 1$. For sufficiently small α , the circle $|q| = C\alpha$ contains $\tilde{q}$. Write

$$P_\alpha(q) = g_\alpha(q) + q^N, \quad g_\alpha(q) = -(1-\alpha)q - \alpha.$$

On $|q| = C\alpha$,

$$|g_\alpha(q)| \geq \alpha((1-\alpha)C - 1), \quad |q^N| = C^N \alpha^N.$$

For sufficiently small α , $|q^N| < |g_\alpha(q)|$ on the circle. Rouché's theorem therefore gives exactly one root inside $|q| < C\alpha$.

For the remaining roots,

$$P_\alpha(q) \longrightarrow q(q^{N-1} - 1)$$

coefficientwise. The $N - 1$ nonzero limiting roots are simple and lie on the unit circle, so the remaining roots of P_α stay bounded away from zero. Equation (6) then shows that their eigenvalues remain bounded. $\square$

Corollary 2.2 (Limiting spectral stability boundary). *Fix $N \geq 3$ and let $L = \lambda/\alpha$ remain in a compact subset of $(0, \infty)$ as $\alpha \rightarrow 0$, so that $\lambda = \alpha L$. The localized branch satisfies*

$$\mu_u = 1 - L + O(\alpha^{N-1}). \quad (10)$$

The constant branch is exactly $\mu = 1$, and every other branch satisfies

$$\mu_k = 1 + \alpha L(q_k^{-1} - 1) + O(\alpha^2), \quad q_k^{N-1} = 1, \quad q_k \neq 1. \quad (11)$$

Consequently, on every compact subset of $(0, 2)$, the full base operator is spectrally stable for all sufficiently small α , while on every compact subset of $(2, \infty)$ it is spectrally unstable for all sufficiently small α . Thus $L = 2$ is the limiting spectral stability boundary. No finite- α spectral stability claim is made at $L = 2$ itself.

Proof. The root expansion (8) is independent of λ . Substituting $\lambda = \alpha L$ into (6) gives

$$\mu_u = 1 - \alpha L + \alpha L \left[-\frac{1-\alpha}{\alpha} + O(\alpha^{N-2}) \right] = 1 - L + O(\alpha^{N-1}),$$

which proves (10). The root $q = 1$ of (7) is exact and gives the conservation eigenvalue $\mu = 1$. The other roots converge to the nontrivial $(N - 1)$ st roots of unity. For such a root $q_k = e^{i\varphi_k}$, analytic perturbation and (6) give (11), hence

$$|\mu_k|^2 = 1 + 2\alpha L(\cos \varphi_k - 1) + O(\alpha^2) < 1$$

for sufficiently small α . If $L \leq 2 - \varepsilon$, then $|1 - L| \leq 1 - \varepsilon$ and (10) places the localized branch strictly inside the unit disk. If $L \geq 2 + \varepsilon$, then $|1 - L| \geq 1 + \varepsilon$ and the localized branch lies outside it. This proves the stated full-operator result away from the boundary $L = 2$. $\square$

The value 2 should not be interpreted as a universal cut-cell CFL condition. It is the limiting boundary for the particular periodic first-order upwind operator (2). A finite- α version of the localized crossing can be written directly in terms of the localized root $q_u(\alpha)$ of (7). Since that root is independent of L , setting $\mu_u = -1$ in (6) gives

$$L_*(\alpha, N) = \frac{2q_u(\alpha)}{\alpha[q_u(\alpha) - 1]} = 2 + O(\alpha^{N-1}). \quad (12)$$

Thus the finite- α displacement of the localized crossing from $L = 2$ is extremely small when N is large. This formula identifies the crossing of the localized branch; the full spectral statement at finite α still requires the remaining branches to lie inside the unit disk, as in the proof of Corollary 2.2.

Theorem 2.3 (Eigenvector localization). *Normalize the eigenvector associated with μ_u by $(v_u)_c = 1$. Let $d_+(c, i) \in \{0, \dots, N-1\}$ be the number of forward upwind-recurrence steps from c to i around the periodic mesh. Then*

$$|(v_u)_i| = O\left(\alpha^{d_+(c, i)}\right), \quad \|v_u - e_c\|_2 = O(\alpha). \quad (13)$$

Proof. For the localized root, $q_u = O(\alpha)$. Equation (5) therefore gives

$$(v_u)_i = q_u^{d_+(c, i)}$$

under the normalization $(v_u)_c = 1$. The component at the nearest forward regular cell is $O(\alpha)$, and subsequent components are higher powers of α . Summing their squares gives (13). $\square$

Theorem 2.3 shows that the unstable eigenvector is concentrated at the cut cell: each forward recurrence step introduces another factor $O(\alpha)$. This localization explains why a redistribution supported on the cut-cell neighborhood can alter the defect-associated eigenvalue efficiently.

3 State Redistribution and the Spectral Threshold

3.1 Merge-left SRD and the correction direction

We first define the redistribution operator used throughout the analysis. Following Berger–Giuliani notation, $\widehat{U}$ denotes a provisional base-scheme state and $\widehat{Q}$ a merged neighborhood state. In the operator analysis we suppress the time superscript: for any input vector $U \in \mathbb{R}^N$, let $\widehat{U} = AU$. Merge-left SRD uses the two-cell merging neighborhood

$$\mathcal{N}_c = \{c-1, c\}$$

and forms the volume-weighted neighborhood state

$$\widehat{Q}_c = \frac{\widehat{U}_{c-1} + \alpha \widehat{U}_c}{1 + \alpha}. \quad (14)$$

In the present two-cell model, the fully redistributed one-step operator A_{SRD} is defined by

$$(A_{\text{SRD}}U)_i = \begin{cases} \widehat{Q}_c, & i = c-1 \text{ or } i = c, \\ \widehat{U}_i, & \text{otherwise.} \end{cases} \quad (15)$$




Figure 1: One-dimensional cut-cell geometry and merge-left state redistribution. The diagram is schematic and is not drawn to scale. The cut cell c has physical width αh . After the base finite-volume step, the provisional cell averages $\hat{U}_{c-1}$ and $\hat{U}_c$ are merged over the neighborhood $\mathcal{N}_c = \{c-1, c\}$ to form the volume-weighted state $\hat{Q}_c$, which replaces both provisional values.

Because (14) is weighted by the physical cell lengths, the local mass $h\hat{U}_{c-1} + \alpha h\hat{U}_c$ is preserved.

Figure 1 summarizes the geometry and the merge-left postprocessing used in the analysis. It is a definition schematic rather than a numerical result.

The correction made by full SRD is

$$D(U) := A_{\text{SRD}}U - AU.$$

From (14),

$$D_{c-1} = \frac{\alpha(\hat{U}_c - \hat{U}_{c-1})}{1 + \alpha}, \quad D_c = -\frac{\hat{U}_c - \hat{U}_{c-1}}{1 + \alpha},$$

so

$$D(U) = \Delta(U)(\alpha e_{c-1} - e_c), \quad \Delta(U) = \frac{\hat{U}_c - \hat{U}_{c-1}}{1 + \alpha}. \quad (16)$$

For the upwind operator, define the two neighboring jumps

$$\delta_{c-1} = U_{c-1} - U_{c-2}, \quad \delta_c = U_c - U_{c-1}.$$

Then

$$\Delta(U) = \frac{\lambda \delta_{c-1} + (1 - \lambda_\alpha) \delta_c}{1 + \alpha}. \quad (17)$$

We use the volume-weighted inner product and its induced vector norm:

$$\langle x, y \rangle_\alpha = \alpha x_c \bar{y}_c + \sum_{i \neq c} x_i \bar{y}_i, \quad \|x\|_\alpha^2 = \langle x, x \rangle_\alpha. \quad (18)$$

For nonzero real vectors x and y , define their volume-weighted alignment by

$$\mathcal{A}_\alpha(x, y) := \frac{|\langle x, y \rangle_\alpha|}{\|x\|_\alpha \|y\|_\alpha} \in [0, 1]. \quad (19)$$

For sufficiently small α , the localized root q_u is real: the polynomial P_α has real coefficients, and uniqueness of the root in the contraction disk forces it to equal its complex conjugate. We therefore take $v_u \in \mathbb{R}^N$.

Theorem 3.1 (Alignment of merge-left SRD with the unstable mode). *Let v_u be normalized by $(v_u)_c = 1$. For every real state $U \in \mathbb{R}^N$ for which $\Delta(U) \neq 0$,*

$$1 - \mathcal{A}_\alpha(D(U), v_u) = O(\alpha) \quad (\alpha \rightarrow 0). \quad (20)$$

Proof. By (16),

$$D_c = -\Delta, \quad D_{c-1} = \alpha\Delta.$$

With $c = 0$, Theorem 2.3 gives $(v_u)_{c-1} = q_u^{N-1} = O(\alpha^{N-1})$ and $\|v_u - e_c\|_2 = O(\alpha)$. Therefore

$$\begin{aligned} \langle D, v_u \rangle_\alpha &= \alpha\Delta (v_u)_{c-1} - \alpha\Delta (v_u)_c \\ &= -\alpha\Delta + O(\alpha^N \Delta), \\ \|D\|_\alpha^2 &= \alpha^2 |\Delta|^2 + \alpha |\Delta|^2 = \alpha(1 + \alpha) |\Delta|^2, \\ \|v_u\|_\alpha^2 &= \alpha + O(\alpha^2). \end{aligned}$$

Hence

$$\frac{\langle D, v_u \rangle_\alpha}{\|D\|_\alpha \|v_u\|_\alpha} = -\frac{\Delta}{|\Delta|} + O(\alpha),$$

and taking the absolute value gives (20). $\square$

The theorem determines only the asymptotic direction of the SRD correction: the correction and the localized eigenvector approach the same one-dimensional subspace in the volume-weighted geometry. Eigenvalue motion and spectral stability are determined next from the blended operator itself.

3.2 Local spectral stability threshold

Define the partial-redistribution operator

$$A_s = (1 - s)A + sA_{\text{SRD}}, \quad 0 \leq s \leq 1. \quad (21)$$

Thus $s = 0$ is the base upwind update and $s = 1$ is full merge-left SRD.

We now use the fixed-local-CFL regime. Fix

$$L := \lambda_\alpha > 0, \quad \lambda = \alpha L,$$

and let $\alpha \rightarrow 0$ with L fixed. Set $c = 0$.

For an eigenvector $A_s v = \mu v$, first consider the exceptional value $\mu = 1 - \alpha L$, for which the recurrence parametrization below is undefined. The unchanged rows then imply

$$v_0 = \cdots = v_{N-3} = 0.$$

The two modified rows admit a nonzero solution only if

$$s(L - 1 + 2L\alpha) = L(1 + \alpha).$$

If $L - 1 + 2L\alpha \leq 0$, this equation has no solution with $0 \leq s \leq 1$. If $L - 1 + 2L\alpha > 0$, then for sufficiently small α we have $\alpha L < 1$, and hence

$$\frac{L(1 + \alpha)}{L - 1 + 2L\alpha} > 1.$$

Thus $\mu = 1 - \alpha L$ is not an eigenvalue for $0 \leq s \leq 1$ once α is sufficiently small.

For every remaining eigenvalue, define

$$q = \frac{\alpha L}{\mu - 1 + \alpha L}.$$

The rows $i = 1, \dots, N - 2$ are unchanged by redistribution and satisfy

$$v_i = q^i v_0, \quad \mu(q) = 1 - \alpha L + \frac{\alpha L}{q}. \quad (22)$$

Moreover, $v_0 \neq 0$. If $v_0 = 0$, the recurrence gives

$$v_1 = \dots = v_{N-2} = 0,$$

while the $c = 0$ equation gives

$$\left[(1 - s)L + \frac{s}{1 + \alpha} \right] v_{N-1} = 0.$$

The coefficient is strictly positive for $L > 0$ and $0 \leq s \leq 1$, so $v_{N-1} = 0$, contradicting $v \neq 0$. We may therefore normalize $v_0 = 1$ and write $w = v_{N-1}$. Applying the base operator to the eigenvector, define the provisional vector $v^b := Av$. The two provisional components entering the merge are

$$\widehat{v}_{N-1} = \alpha L q^{N-2} + (1 - \alpha L)w, \quad \widehat{v}_0 = Lw + 1 - L. \quad (23)$$

The corresponding merged neighborhood value is

$$\widehat{Q}_0 = \frac{\widehat{v}_{N-1} + \alpha \widehat{v}_0}{1 + \alpha}.$$

The two modified eigenvalue equations are therefore

$$\mu w = (1 - s)\widehat{v}_{N-1} + s\widehat{Q}_0, \quad \mu = (1 - s)\widehat{v}_0 + s\widehat{Q}_0.$$

Eliminating w gives

$$\begin{aligned} G_N(q; \alpha, L, s) = & sq \left[L\alpha^2(q - 1) + L\alpha(2q^{N-1} - q^{N-2} - 1) + L(q^{N-1} - 1) \right. \\ & \left. + \alpha(1 - q) + (1 - q^{N-1}) \right] \\ & - L(1 + \alpha)[\alpha(q - 1) + q^N - q]. \end{aligned} \quad (24)$$

For $\alpha > 0$, the N roots of $G_N = 0$ correspond to the N eigenvalues of A_s through (22). At $\alpha = 0$,

$$G_N(q; 0, L, s) = q(1 - q^{N-1})[L - (L - 1)s]. \quad (25)$$

Thus one branch emerges from $q = 0$ and the other $N - 1$ branches emerge from the $(N - 1)$ st roots of unity.

Theorem 3.2 (Local spectral stability threshold). *Fix $N \geq 4$ and $L > 2$. Define*

$$s_0(L) = \frac{L - 2}{L - 1}, \quad q_-(\alpha, L) = \frac{\alpha L}{\alpha L - 2}, \quad r_- = q_-^{N-2}.$$

For all sufficiently small $\alpha > 0$, the localized eigenvalue branch of A_s equals -1 at a unique value of s in a neighborhood of $s_0(L)$, namely

$$s_{-1}^{(N)}(\alpha, L) = \frac{(1 + \alpha) [L^2 \alpha r_- - L^2 \alpha + 2L\alpha + 2L - 4]}{L^2 \alpha^2 r_- - L^2 \alpha^2 + L^2 \alpha r_- - L^2 \alpha + L\alpha r_- + 2L\alpha^2 + 3L\alpha + 2L - 2\alpha - 2}. \quad (26)$$

As $\alpha \rightarrow 0$ with L fixed,

$$s_{-1}^{(N)}(\alpha, L) = \frac{L - 2}{L - 1} - \frac{L(L - 2)}{2(L - 1)^2} \alpha^2 + O(\alpha^3). \quad (27)$$

Moreover, there exist $\alpha_0 > 0$ and a neighborhood $\mathcal{I}$ of $s_0(L)$ such that for $0 < \alpha < \alpha_0$ and $s \in \mathcal{I}$,

$$\rho(A_s) \leq 1 \iff s \geq s_{-1}^{(N)}(\alpha, L). \quad (28)$$

Proof. Setting $\mu = -1$ in (22) gives

$$q = q_-(\alpha, L) = \frac{\alpha L}{\alpha L - 2} = O(\alpha).$$

Substitution into (24) yields an equation affine in s ; solving it gives (26). Because $N \geq 4$, $r_- = O(\alpha^{N-2}) = O(\alpha^2)$, and expansion of the rational expression gives (27). In particular, the coefficient of α vanishes.

It remains to show that this crossing is the local spectral stability boundary of the full periodic matrix. At $s = s_0(L)$, equation (25) has one simple root at $q = 0$ and $N - 1$ simple roots

$$q_k = e^{2\pi i k / (N-1)}, \quad k = 0, \dots, N-2.$$

The root $q_0 = 1$ persists exactly for every α and s because $A_s \mathbf{1} = \mathbf{1}$, where $\mathbf{1} = (1, \dots, 1)^\top$; it is the conservation eigenvalue $\mu = 1$.

For each $q_k \neq 1$, the implicit-function theorem gives a nearby root $q_k(\alpha, s) = q_k + O(\alpha)$, uniformly for s in a sufficiently small neighborhood of s_0 . Using (22),

$$\mu_k = 1 + \alpha L(q_k^{-1} - 1) + O(\alpha^2),$$

and therefore, with $q_k = e^{i\varphi_k}$,

$$|\mu_k|^2 = 1 + 2\alpha L(\cos \varphi_k - 1) + O(\alpha^2) < 1$$

for sufficiently small α . Thus every nonconstant regular branch is strictly inside the unit disk near the crossing.

For the branch issuing from $q = 0$, set $q = \alpha z$ in (24). The leading equation is

$$z = -\frac{L}{L - (L-1)s},$$

and (22) gives

$$\mu_u(\alpha, s) = -(L-1)(1-s) + O(\alpha), \quad \partial_s \mu_u = L-1 + O(\alpha) > 0$$

near s_0 . The exact value (26) belongs to this branch because $q_- = O(\alpha)$, and there $\mu_u = -1$. Hence the localized eigenvalue enters the unit disk as s increases through $s_{-1}^{(N)}$. The regular nonconstant branches are already strictly inside, and the constant branch remains at 1. This proves (28). $\square$

Corollary 3.3 (Leading-order spectral threshold). *For fixed $L > 2$ and $N \geq 4$,*

$$s_0(L) - s_{-1}^{(N)}(\alpha, L) = \frac{L(L-2)}{2(L-1)^2} \alpha^2 + O(\alpha^3) > 0$$

for sufficiently small α , where

$$s_0(L) = \frac{L-2}{L-1}. \tag{29}$$

Hence $s = s_0(L)$ is asymptotically sufficient for the full periodic model and exceeds the exact local spectral stability threshold only by $O(\alpha^2)$. For fixed $0 < L < 2$, Corollary 2.2 shows that the unstabilized operator $s = 0$ is already spectrally stable for sufficiently small α .

The exact expression (26) is useful for verifying the finite- α crossing of the model operator, but it is intentionally not used as the higher-order prescription. It depends on the specific first-order, one-cut-cell periodic operator and on N , whereas the higher-order benchmark changes the reconstruction, time integrator, and number of cut cells. Using $s_0(L)$ avoids assigning model-specific finite- α precision to that transfer. For reference, at $N = 150$ and $L = 4$, $s_0 - s_{-1}^{(150)}$ is approximately 0.0050 for $\alpha = 0.10$ and 0.0122 for $\alpha = 0.15$.

The two statements meet at the limiting boundary $L = 2$, but the theorem does not assert that $s = 0$ is spectrally stable there for every finite α . For $L > 2$, $s_0(L)$ increases monotonically from 0 and approaches 1 as $L \rightarrow \infty$. Thus the model spectral threshold can be substantially below the full-redistribution value $s = 1$ for moderate L , while $s_0(L) \rightarrow 1$ as $L \rightarrow \infty$.

Remark 3.4 (Spectral stability versus power boundedness). *Theorem 3.2 establishes a local spectral stability threshold, meaning $\rho(A_s) \leq 1$. The matrices A_s are generally nonnormal, so this condition alone does not imply power boundedness, i.e., $\sup_{n \geq 0} \|A_s^n\| < \infty$ in a fixed matrix norm, and it does not exclude finite transient amplification. The theorem is also local in s and asymptotic in α at fixed L ; it does not prove that every $s \geq s_{-1}^{(N)}(\alpha, L)$ is spectrally stable for arbitrary finite α , nor does it transfer the same threshold to other fluxes, reconstructions, nonlinear systems, or multidimensional redistribution rules. Section 9 therefore reports finite-time amplification separately from the spectral crossing.*




Figure 2: Localized spectral branch and eigenvector for $N = 150$, $\alpha = 0.10$, and $L = 4$. *Left:* $\mu_u(s)$ with $\mu_u(s_{-1}^{(150)}) = -1$, where $s_{-1}^{(150)} \approx 0.6617$ and $s_0 = 2/3 \approx 0.6667$. *Right:* $|(v_u)_{c+j}|/|(v_u)_c|$ at $s = 0$, $s = s_{-1}^{(150)}$, and $s = 1$, where j is a signed periodic offset and $c + j$ is interpreted modulo N .

Figure 2 displays the two quantities entering Theorem 3.2: the branch $\mu_u(s)$ and its normalized eigenvector. The mode is already localized at $s = 0$. Increasing s moves $\mu_u(s)$ toward the unit circle, with $\mu_u(s_{-1}^{(N)}) = -1$. The distinction between $s_{-1}^{(N)}$ and s_0 is the finite- α correction quantified by (27).

4 Singular-Block Framework

The preceding calculation has one singular row. We now isolate the part of the argument that does not depend on the one-dimensional upwind stencil. The section has two steps. First, Theorem 4.2

identifies the unbounded eigenvalue cluster from a finite-dimensional defect matrix. Second, Theorem 4.3 shows that the associated invariant subspace concentrates on the defect coordinates. The later cut-cell results apply these two statements with concrete finite-volume blocks.

Let $A_\alpha \in \mathbb{C}^{N \times N}$ be a family of matrices. After reordering the unknowns, suppose the first m coordinates represent geometric defects and the remaining $N - m$ coordinates are regular. Write

$$A_\alpha = \begin{pmatrix} -\alpha^{-p}\Gamma + R_\alpha & B_\alpha \\ C_\alpha & Q_\alpha \end{pmatrix}, \quad p > 0. \quad (30)$$

The matrix $\Gamma \in \mathbb{C}^{m \times m}$ is called the *defect matrix*. The coordinate defect subspace is

$$\mathcal{C} = \mathbb{C}^m \oplus \{0\}.$$

Throughout this section, I_k denotes the $k \times k$ identity matrix, I denotes an identity matrix whose dimension is clear from context, and $\|\cdot\|$ denotes any fixed subordinate matrix norm.

Assumption 4.1 (Singular-block structure). *As $\alpha \rightarrow 0$,*

$$\|\alpha^p R_\alpha\| \rightarrow 0, \quad \|B_\alpha\| = O(\alpha^{-p}), \quad \|C_\alpha\| = O(1), \quad \|Q_\alpha\| = O(1),$$

and Γ is invertible. The dimensions m and N are fixed.

The upper-right block may have the same asymptotic order as the defect block. This is needed for conservative cut-cell methods because a singular cut-cell row contains both a large diagonal coefficient and large coefficients multiplying its regular neighbors. The lower blocks remain bounded because regular-cell updates are not divided by a vanishing volume.

Theorem 4.2 (Singular-block spectrum). *Under Assumption 4.1, A_α has exactly m unbounded eigenvalues, counted with algebraic multiplicity. Let $\gamma_1, \dots, \gamma_m$ denote the eigenvalues of Γ , repeated according to algebraic multiplicity. The unbounded eigenvalues of A_α can be labeled $\mu_1(\alpha), \dots, \mu_m(\alpha)$ so that*

$$\alpha^p \mu_k(\alpha) \longrightarrow -\gamma_k, \quad k = 1, \dots, m. \quad (31)$$

All remaining eigenvalues are bounded. If, in addition, Γ is diagonalizable and $R_\alpha = O(1)$, then the unbounded eigenvalues may be matched with the eigenvalues γ_k of Γ so that

$$\mu_k(\alpha) = -\frac{\gamma_k}{\alpha^p} + O(1). \quad (32)$$

Proof. For $|\mu| > \|Q_\alpha\| + 1$, Schur complementation gives

$$\det(\mu I - A_\alpha) = \det(\mu I - Q_\alpha) \det \mathcal{S}_\alpha(\mu),$$

where

$$\mathcal{S}_\alpha(\mu) = \mu I_m + \alpha^{-p}\Gamma - R_\alpha - B_\alpha(\mu I - Q_\alpha)^{-1}C_\alpha.$$

Set $\nu = \alpha^p \mu$. Uniformly for ν in compact subsets of $\mathbb{C} \setminus \{0\}$,

$$\alpha^p \mathcal{S}_\alpha(\nu/\alpha^p) = \nu I_m + \Gamma + o(1).$$

Indeed, $\alpha^p R_\alpha \rightarrow 0$, and

$$\alpha^p B_\alpha(\nu \alpha^{-p} I - Q_\alpha)^{-1} C_\alpha = O(\alpha^p).$$

Choose disjoint contours enclosing the distinct eigenvalue clusters of $-\Gamma$ and excluding 0. Uniform convergence of the determinant on each contour and Rouché's theorem preserve the algebraic

number of roots inside each contour. Hence there are exactly m such eigenvalues, counted with algebraic multiplicity, and they can be labeled so that $\alpha^p \mu_k(\alpha) \rightarrow -\gamma_k$.

Suppose a further eigenvalue were unbounded. Since $\|A_\alpha\| = O(\alpha^{-p})$, after passing to a subsequence its scaled value $\nu_\alpha = \alpha^p \mu_\alpha$ has a finite limit. A nonzero limit would belong to $-\sigma(\Gamma)$ and is already counted. If $\nu_\alpha \rightarrow 0$ while $|\mu_\alpha| \rightarrow \infty$, then

$$\alpha^p \mathcal{S}_\alpha(\mu_\alpha) \rightarrow \Gamma,$$

which is invertible, a contradiction. Thus the remaining spectrum is bounded.

For the final refinement, assume that Γ is diagonalizable and $R_\alpha = O(1)$. On the unbounded spectral scale, the Schur equation can be written

$$\det(\mu I_m + \alpha^{-p} \Gamma + H_\alpha(\mu)) = 0, \quad H_\alpha(\mu) = -R_\alpha - B_\alpha(\mu I - Q_\alpha)^{-1} C_\alpha,$$

and $\|H_\alpha(\mu)\| \leq C$ uniformly whenever $|\mu| \asymp \alpha^{-p}$. Write $\Gamma = V \Lambda V^{-1}$ with Λ diagonal and let $\kappa(V) = \|V\| \|V^{-1}\|$. For $M_0(\mu) = \mu I_m + \alpha^{-p} \Gamma$,

$$\|M_0(\mu)^{-1}\| \leq \frac{\kappa(V)}{\text{dist}(\mu, -\alpha^{-p} \sigma(\Gamma))}.$$

Here $\text{dist}(z, S) := \inf_{w \in S} |z - w|$ denotes the distance from a complex number z to a set $S \subset \mathbb{C}$. Hence, if the distance in the denominator exceeds $\kappa(V)C$, then $\|M_0(\mu)^{-1} H_\alpha(\mu)\| < 1$ and the Neumann-series criterion shows that $M_0(\mu) + H_\alpha(\mu)$ is invertible. Every unbounded eigenvalue must therefore lie within an $O(1)$ neighborhood of $-\alpha^{-p} \sigma(\Gamma)$. Combined with the algebraic multiplicity count already established for each cluster, this yields (32). $\square$

Theorem 4.3 (Defect-subspace convergence). *Under Assumption 4.1, let $\mathcal{U}_\alpha$ be the invariant subspace associated with the m unbounded eigenvalues. Then*

$$\|P_{\mathcal{U}_\alpha} - P_{\mathcal{V}}\| = O(\alpha^p), \tag{33}$$

where $P_{\mathcal{V}}$ denotes the Euclidean orthogonal projector onto a subspace $\mathcal{V}$.

Proof. Scale the matrix:

$$\tilde{A}_\alpha := \alpha^p A_\alpha = \begin{pmatrix} -\Gamma + \alpha^p R_\alpha & \alpha^p B_\alpha \\ \alpha^p C_\alpha & \alpha^p Q_\alpha \end{pmatrix}.$$

By Theorem 4.2, the nonzero spectral cluster converges to $-\sigma(\Gamma)$ and remains separated from the cluster at 0. Choose a fixed positively oriented contour $\mathfrak{C}$ enclosing $-\sigma(\Gamma)$ and excluding 0. The Riesz projector

$$\Pi_\alpha = \frac{1}{2\pi i} \int_{\mathfrak{C}} (zI - \tilde{A}_\alpha)^{-1} dz$$

has range $\mathcal{U}_\alpha$ and rank m .

Write

$$zI - \tilde{A}_\alpha = \begin{pmatrix} X_\alpha(z) & -\alpha^p B_\alpha \\ -\alpha^p C_\alpha & Y_\alpha(z) \end{pmatrix},$$

where

$$X_\alpha(z) = zI + \Gamma - \alpha^p R_\alpha, \quad Y_\alpha(z) = zI - \alpha^p Q_\alpha.$$

Since $\mathfrak{C}$ stays a positive distance from 0, $Y_\alpha(z)^{-1} = O(1)$ uniformly on $\mathfrak{C}$. The Schur complement

$$\mathcal{S}_\alpha(z) = X_\alpha(z) - \alpha^{2p} B_\alpha Y_\alpha(z)^{-1} C_\alpha$$

satisfies

$$\mathcal{S}_\alpha(z) = zI + \Gamma + o(1)$$

uniformly on $\mathfrak{C}$, because $\alpha^{2p}B_\alpha Y_\alpha(z)^{-1}C_\alpha = O(\alpha^p)$. Hence $\mathcal{S}_\alpha(z)^{-1} = O(1)$ uniformly on the contour.

The block inverse formula now gives

$$(zI - \tilde{A}_\alpha)_{21}^{-1} = Y_\alpha(z)^{-1}\alpha^p C_\alpha \mathcal{S}_\alpha(z)^{-1} = O(\alpha^p),$$

and

$$(zI - \tilde{A}_\alpha)_{22}^{-1} = Y_\alpha(z)^{-1} + Y_\alpha(z)^{-1}\alpha^p C_\alpha \mathcal{S}_\alpha(z)^{-1}\alpha^p B_\alpha Y_\alpha(z)^{-1}.$$

The second term is $O(\alpha^p)$. The first term, $Y_\alpha(z)^{-1}$, is analytic inside $\mathfrak{C}$ because the spectrum of $\alpha^p Q_\alpha$ remains near 0, outside the contour. Its contour integral is therefore zero. It follows that

$$(\Pi_\alpha)_{21} = O(\alpha^p), \quad (\Pi_\alpha)_{22} = O(\alpha^p).$$

Similarly,

$$(\Pi_\alpha)_{11} = I_m + o(1),$$

because the contour encloses all m eigenvalues of $-\Gamma$.

For small α , $(\Pi_\alpha)_{11}$ is invertible. The m vectors $\Pi_\alpha(x, 0)^\top$ therefore span $\text{range}(\Pi_\alpha)$ and show that

$$\mathcal{U}_\alpha = \{(x, K_\alpha x) : x \in \mathbb{C}^m\}, \quad K_\alpha = (\Pi_\alpha)_{21}(\Pi_\alpha)_{11}^{-1} = O(\alpha^p).$$

The standard orthogonal-projector formula for the graph of K_α then gives $\|P_{\mathcal{U}_\alpha} - P_{\mathcal{C}}\| = O(\|K_\alpha\|) = O(\alpha^p)$. $\square$

Remark 4.4. *The projector theorem concerns the whole invariant subspace and therefore does not require the defect eigenvalues to be simple. This distinction matters when two cut cells have the same leading defect coefficient: individual eigenvectors may not have a unique limiting labeling, while the m -dimensional unbounded spectral subspace still converges to the defect-coordinate subspace.*

Theorems 4.2 and 4.3 isolate the general mechanism: a singular defect block creates a finite cluster of unbounded eigenvalues, and the corresponding invariant subspace concentrates on the coordinates where the singular scaling occurs. Computationally, this identifies a low-dimensional object to analyze: when the number of geometric defects $m \ll N$, the singular spectral behavior is carried by an m -dimensional defect-associated subspace rather than requiring calibration against the full N -dimensional state space. The next sections recover the concrete cut-cell consequences.

5 Multiple Cut Cells

Let

$$J = \{j_1, \dots, j_m\}$$

be non-adjacent cut cells. Assume their volume fractions shrink at comparable rates. Write the volume fraction of cut cell j_k as $\alpha_k = \theta_k \alpha$, where α is a common small parameter, θ_k is a fixed positive scale factor, and the scale factors satisfy fixed bounds $0 < \theta_- \leq \theta_k \leq \theta_+ < \infty$:

$$\alpha_k = \theta_k \alpha, \quad 0 < \theta_- \leq \theta_k \leq \theta_+ < \infty, \quad (34)$$

with fixed m , N , and background CFL $\lambda \in (0, 1)$. Reorder the unknowns so the cut-cell coordinates come first.

For first-order positive-speed upwind, non-adjacency means that the upwind neighbor of a cut cell is regular. The cut-cut block is therefore diagonal:

$$(A_\alpha)_{JJ} = I_m - \frac{1}{\alpha}\Gamma, \quad \Gamma = \text{diag}\left(\frac{\lambda}{\theta_1}, \dots, \frac{\lambda}{\theta_m}\right).$$

The cut-to-regular block is $O(\alpha^{-1})$, while all regular rows are bounded. Thus Assumption 4.1 holds with $p = 1$.

Corollary 5.1 (Separated multi-cut-cell defect subspace). *Under (34), the periodic first-order upwind operator has exactly m unbounded eigenvalues. After matching with multiplicity,*

$$\mu_k(\alpha) = -\frac{\lambda}{\theta_k \alpha} + O(1), \quad k = 1, \dots, m. \quad (35)$$

If $\mathcal{U}_\alpha$ is the invariant subspace associated with these eigenvalues and

$$\mathcal{C} = \text{span}\{e_{j_1}, \dots, e_{j_m}\},$$

then

$$\|P_{\mathcal{U}_\alpha} - P_{\mathcal{C}}\| = O(\alpha). \quad (36)$$

Proof. The defect matrix is diagonal and hence diagonalizable, while the remainder in the cut-cut block is $I_m = O(1)$. Apply Theorems 4.2 and 4.3. $\square$

Equation (36) is the invariant-subspace version of the one-cell localization theorem. When two cut cells have the same leading coefficient, individual eigenvectors in the unbounded spectral cluster need not select unique coordinate vectors. The invariant statement is that the entire unbounded spectral subspace converges to the span of the cut-cell coordinates; this subspace remains well defined when leading defect eigenvalues are repeated.

Adjacent cut cells can introduce singular couplings inside the defect block and therefore require a different Γ . Allowing m or N to grow with α^{-1} is also outside the fixed-dimensional theorem. Likewise, widely separated volume scales require either several small parameters or a rescaling different from (34).

6 One Cut Cell in d Dimensions

Consider a periodic first-order upwind discretization in d dimensions. Let λ_r be the directional background CFL number in coordinate direction r , and define the summed background CFL number

$$\lambda_\Sigma := \sum_{r=1}^d \lambda_r.$$

The corresponding local cut-cell CFL number is $\lambda_{\alpha,\Sigma} := \lambda_\Sigma/\alpha$. For one cut cell c of volume fraction α , the cut-cell row is

$$(A_\alpha U)_c = \left(1 - \frac{\lambda_\Sigma}{\alpha}\right) U_c + \sum_{r=1}^d \frac{\lambda_r}{\alpha} U_{c-\hat{r}}, \quad (37)$$

where $\hat{r}$ denotes one grid-index step in coordinate direction r , so $c - \hat{r}$ is the upwind neighbor of the cut cell in that direction. Regular rows remain $O(1)$.

After moving the cut-cell coordinate first, Assumption 4.1 holds with $m = 1$, $p = 1$, scalar defect matrix $\Gamma = \lambda_\Sigma$, and bounded remainder in the defect diagonal.

Corollary 6.1 (Multidimensional localized defect mode). *For fixed grid size and fixed directional CFL numbers, the operator has one unbounded eigenvalue*

$$\mu_\alpha = -\lambda_{\alpha,\Sigma} + O(1) = -\frac{\lambda_\Sigma}{\alpha} + O(1), \quad (38)$$

while all remaining eigenvalues are bounded. If the associated eigenvector is normalized by $(v_\alpha)_c = 1$, then

$$\|v_\alpha - e_c\|_2 = O(\alpha). \quad (39)$$

Proof. Apply Theorems 4.2 and 4.3 with $\Gamma = \lambda_\Sigma$. Since the scalar defect matrix is diagonalizable and the defect remainder is bounded, the $O(1)$ eigenvalue refinement applies. $\square$

The multidimensional corollary is a leading-order spectral result. The geometry of a real embedded boundary affects face areas, neighboring cut cells, reconstruction, and the redistribution neighborhoods. Those features enter through the actual assembled defect matrix and cannot be replaced by the scalar L without a separate derivation.

7 Finite-Volume Conditions for the Defect Structure

The abstract block assumptions are useful only if they can be checked from a discretization. We therefore state sufficient conditions at the assembled operator level and then derive a local formula for a one-dimensional bounded-flux scheme. We state sufficient conditions rather than an equivalence because cancellations among face contributions can invalidate conclusions based only on face-level scaling estimates.

Let J_c and J_r denote the cut-cell and regular-cell index sets, respectively. For index sets X and Y , $(A_\alpha)_{XY}$ denotes the block of A_α with rows in X and columns in Y . Assume $\alpha_k = \theta_k \alpha$ with the fixed bounds in (34).

Theorem 7.1 (Sufficient assembled-operator conditions). *Suppose that, after ordering the cut-cell coordinates first, there exists $p > 0$ and an invertible matrix Γ such that*

$$\alpha^p (A_\alpha)_{J_c J_c} \longrightarrow -\Gamma, \quad (40)$$

and

$$\|(A_\alpha)_{J_c J_r}\| = O(\alpha^{-p}), \quad \|(A_\alpha)_{J_r J_c}\| + \|(A_\alpha)_{J_r J_r}\| = O(1). \quad (41)$$

Then A_α satisfies Assumption 4.1.

Proof. Set

$$R_\alpha = (A_\alpha)_{J_c J_c} + \alpha^{-p} \Gamma, \quad B_\alpha = (A_\alpha)_{J_c J_r}, \quad C_\alpha = (A_\alpha)_{J_r J_c}, \quad Q_\alpha = (A_\alpha)_{J_r J_r}.$$

Equation (40) is exactly $\alpha^p R_\alpha \rightarrow 0$, and (41) supplies the remaining bounds. $\square$

Let C_c denote the cut-cell control volume, $|C_c|$ its volume, ∂C_c its boundary, f a face on that boundary, $|f|$ the face measure, and $F_f(U)$ the numerical flux through face f with the chosen outward orientation. A conservative finite-volume update on C_c has the form

$$U_c^{n+1} = U_c^n - \frac{\Delta t}{|C_c|} \sum_{f \in \partial C_c} |f| F_f(U). \quad (42)$$

The singular scaling can therefore come from the small volume $|C_c|$, from reconstruction coefficients inside F_f , from shrinking or nonshrinking face measures, or from combinations of these effects. The assembled conditions above keep all of those factors visible instead of assuming that the volume fraction alone determines the power of α .

8 One-Dimensional Bounded-Flux Schemes

We now connect the block framework to the local flux coefficients in the simplest setting where an explicit formula is available. Consider a scalar one-dimensional conservation law with one cut cell of width αh . Linearize the two numerical fluxes adjacent to the cut cell. Let $\beta_{c\pm 1/2,j}(\alpha)$ denote the resulting coefficient multiplying cell value U_j , where j ranges over the finite flux stencil:

$$F_{c+1/2}(U) = \sum_j \beta_{c+1/2,j}(\alpha) U_j, \quad F_{c-1/2}(U) = \sum_j \beta_{c-1/2,j}(\alpha) U_j. \quad (43)$$

Only finitely many coefficients are nonzero.

Theorem 8.1 (Defect coefficient for bounded-flux schemes). *Assume*

$$\beta_{c\pm 1/2,j}(\alpha) = O(1)$$

uniformly as $\alpha \rightarrow 0$, all regular rows of the assembled update are $O(1)$, and

$$g_\alpha := \frac{\Delta t}{h} (\beta_{c+1/2,c}(\alpha) - \beta_{c-1/2,c}(\alpha)) = \Gamma_0 + O(\alpha), \quad \Gamma_0 \neq 0. \quad (44)$$

Then the assembled operator satisfies Assumption 4.1 with $m = 1$, $p = 1$, scalar defect coefficient Γ_0 , bounded R_α , $B_\alpha = O(\alpha^{-1})$, and bounded lower blocks.

Proof. The cut-cell update is

$$U_c^{n+1} = U_c^n - \frac{\Delta t}{\alpha h} (F_{c+1/2}(U) - F_{c-1/2}(U)).$$

The coefficient multiplying U_c is therefore

$$1 - \frac{g_\alpha}{\alpha} = -\frac{\Gamma_0}{\alpha} + O(1).$$

Every other coefficient in the cut-cell row is $O(\alpha^{-1})$ because the linearized flux coefficients are bounded. By assumption, the regular rows remain bounded. These are precisely the block estimates of Assumption 4.1. $\square$

By (44), Γ_0 is the leading coefficient of U_c in the dimensionless flux difference $(\Delta t/h)(F_{c+1/2} - F_{c-1/2})$ as $\alpha \rightarrow 0$. For positive-speed first-order upwind,

$$\beta_{c+1/2,c} = a, \quad \beta_{c-1/2,c} = 0,$$

so $\Gamma_0 = \lambda$. The abstract theorem then recovers the leading unbounded eigenvalue $-\lambda/\alpha + O(1)$, consistent with Theorem 2.1.

The value 2 and the spectral-threshold formula in Section 3.2 are specific to the complete characteristic equation of the first-order upwind model with merge-left SRD. The defect coefficient Γ_0 alone determines the leading singular spectral scale, not a scheme-independent spectral stability threshold.

If reconstruction introduces another inverse power of the cut-cell scale, then the defect row can be more singular. For example, a reconstruction coefficient of size $O(\alpha^{-1})$ combined with division by the cut-cell volume produces $O(\alpha^{-2})$ entries. In that case the framework is applied with $p = 2$ after computing the nonzero scaled limit $\alpha^2(A_\alpha)_{J_c J_c}$. This observation concerns the spectral scaling of the assembled operator; it does not by itself imply a spectral stability threshold for a higher-order reconstruction.

9 Numerical Experiments

The numerical study has three components. We first evaluate the spectrum of the exact periodic upwind operator with merge-left SRD and then compute its finite-time induced amplification near the spectral threshold. We finally test the model-derived redistribution fraction in a separate second-order piecewise-linear benchmark, measuring convergence, fixed-error cell count, layout sensitivity, and 100-period behavior. The higher-order tests are empirical and are not covered by Theorem 3.2.

9.1 Full-matrix spectral crossing

Figure 2, placed after Theorem 3.2, shows the full spectral mechanism for $N = 150$, $\alpha = 0.10$, and $L = 4$. The localized eigenvalue begins near $1 - L = -3$, crosses $\mu = -1$ at

$$s_{-1}^{(150)} \approx 0.661654,$$

and then lies inside the unit disk. The leading prescription $s_0 = (L - 2)/(L - 1) = 2/3$ is slightly larger, consistent with the negative $O(\alpha^2)$ correction in (27). The associated eigenvectors remain localized near the cut cell throughout the crossing.

Figure 3 plots the spectrum in the complex μ -plane. The orange curve traces the localized branch $\mu_u(s)$. The gray markers are the remaining individual eigenvalues of the full matrix at $s = s_{-1}^{(N)}$; together they form the inner spectral loop. At $s = 0$, $\mu_u(s)$ lies outside $|\mu| = 1$; at $s = s_{-1}^{(N)}$, $\mu_u = -1$; and for larger s the branch lies inside the unit disk. The black point at $\mu = 1$ is the conservation eigenvalue.

To visualize the asymptotic correction itself, Figure 4 plots the exact crossing (26) of the full periodic operator for $L \in \{3, 4, 6\}$ and $N = 150$. The left panel shows convergence toward the leading value $s_0(L)$. The right panel scales the difference by α^2 ; as $\alpha \rightarrow 0$ the curves approach

$$\frac{L(L - 2)}{2(L - 1)^2},$$

the coefficient predicted by (27). This figure is a visualization of the exact formula, not independent evidence for it.

9.2 Finite-time amplification near the spectral threshold

Because A_s is nonnormal, transient growth can persist after the spectral crossing. With the cut cell ordered first, set

$$W_\alpha := \text{diag}(\alpha, 1, \dots, 1).$$

For a matrix M , let

$$\|M\|_\alpha := \sup_{x \neq 0} \frac{\|Mx\|_\alpha}{\|x\|_\alpha} = \|W_\alpha^{1/2} M W_\alpha^{-1/2}\|_2. \quad (45)$$

Figure 5 plots $\|A_s^n\|_\alpha$ through $n = 200$ for $N = 80$ and fixed $L = 4$.

For $\alpha = 0.10$, taking s only 0.01 below the exact crossing gives $\rho(A_s) = 1.0301$ and $\|A_s^{200}\|_\alpha \approx 582$; the eigenvalue outside $\mathbb{D}$ produces geometric modal growth. At the exact crossing the maximum over the first 200 steps is approximately 2.01. At the slightly larger model prescription $s_0 = 2/3$, the maximum is approximately 2.01 and $\|A_s^{200}\|_\alpha$ has returned to 1. For $\alpha = 0.15$, the corresponding maxima at the exact crossing and at s_0 are approximately 2.34 and 2.26. These finite-time computations show that crossing the unit circle eliminates exponential growth of the localized eigenmode while finite nonnormal amplification remains possible; power boundedness is a separate question.




Figure 3: Localized spectral crossing for $N = 150$, $\alpha = 0.10$, and $L = 4$. The unit circle $|\mu| = 1$ is the spectral-stability boundary. The orange curve traces the localized branch $\mu_u(s)$ for $0 \leq s \leq 1$. The gray markers denote the remaining individual eigenvalues of the full matrix $A_{s_{-1}^{(N)}}$ evaluated at the exact crossing $s_{-1}^{(150)} = 0.661654\dots$; collectively they form the inner spectral loop and remain strictly inside the unit circle. The black point at $\mu = 1$ is the conservation eigenvalue, and the green marker identifies the localized crossing $\mu_u(s_{-1}^{(150)}) = -1$. The asymptotic threshold $s_0 = 2/3$ lies slightly above this finite- α crossing.

9.3 Higher-order accuracy benchmark

The spectral-threshold theorem is proved only for the first-order periodic upwind model with merge-left SRD. We therefore test the model-derived partial choice in a separate second-order piecewise-linear benchmark and measure its effect on accuracy for

$$u_t + u_x = 0$$

on the unit periodic domain. For unit advection speed, let $T := 1$ denote one advection period.

The mesh contains m separated cut cells of width αh and $N - m$ regular cells of width h on a unit periodic domain. Thus

$$(N - m)h + m\alpha h = 1, \quad h = \frac{1}{N - m + m\alpha}.$$

For the evenly spaced layouts used in the convergence tests, the zero-based cut-cell indices are

$$c_j = \left\lfloor 5 + \frac{j}{m-1}(N-11) \right\rfloor, \quad j = 0, \dots, m-1,$$

for $m > 1$ (and $c_0 = 5$ for $m = 1$). Let x_i denote the center of cell i and $x_{i+1/2}$ its right face. For two real numbers r_1 and r_2 , define

$$\text{minmod}(r_1, r_2) = \begin{cases} \text{sign}(r_1) \min\{|r_1|, |r_2|\}, & r_1 r_2 > 0, \\ 0, & r_1 r_2 \leq 0. \end{cases}$$




Figure 4: Finite- α correction to the local spectral stability threshold for $N = 150$. *Left*: exact crossing $s_{-1}^{(N)}(\alpha, L)$ from (26); dashed horizontal lines are $s_0(L) = (L - 2)/(L - 1)$. *Right*: the scaled difference $(s_0 - s_{-1}^{(N)})/\alpha^2$ approaches $L(L - 2)/(2(L - 1)^2)$ as $\alpha \rightarrow 0$, consistent with (27).

The piecewise-linear slope and left state at the right face are

$$\ell_i = \minmod\left(\frac{U_i - U_{i-1}}{x_i - x_{i-1}}, \frac{U_{i+1} - U_i}{x_{i+1} - x_i}\right), \quad U_{i+1/2}^- = U_i + \ell_i(x_{i+1/2} - x_i).$$

For positive unit advection speed, the numerical flux is $F_{i+1/2} = U_{i+1/2}^-$. If h_i is the width of cell i , define the finite-volume spatial operator

$$(\mathcal{L}_h(U))_i = -\frac{F_{i+1/2}(U) - F_{i-1/2}(U)}{h_i}.$$

To make the time discretization and the placement of redistribution explicit, let V denote an arbitrary stage state and define its provisional forward Euler update by

$$\hat{U}(V) = V + \Delta t \mathcal{L}_h(V).$$

Let $\mathcal{R}$ denote conservative merge-left redistribution: for every cut cell c , $\mathcal{R}$ replaces $\hat{U}_{c-1}$ and $\hat{U}_c$ by their volume-weighted average and leaves all other components unchanged. The cut cells are separated so these two-cell neighborhoods do not overlap. For a prescribed scalar $s \in [0, 1]$, define one redistributed forward Euler stage by

$$\mathcal{E}_s(V) = (1 - s)\hat{U}(V) + s\mathcal{R}(\hat{U}(V)). \quad (46)$$

The second-order strong-stability-preserving Runge-Kutta method (SSPRK2) used in the experiments is then

$$U^{(1)} = \mathcal{E}_s(U^n), \quad U^{n+1} = \frac{1}{2}U^n + \frac{1}{2}\mathcal{E}_s(U^{(1)}). \quad (47)$$

Thus redistribution is applied after each forward Euler stage, exactly as in the numerical implementation.

We use the nominal background CFL $\lambda_0 = 0.4$ and define

$$L_0 := \frac{\lambda_0}{\alpha}, \quad s = s_0(L_0) = \frac{L_0 - 2}{L_0 - 1}.$$




Figure 5: Finite-time induced amplification near the spectral crossing for $N = 80$ and $L = 4$. *Left:* $\alpha = 0.10$. *Right:* $\alpha = 0.15$. For $s = s_{-1}^{(N)} - 0.01$, $\rho(A_s) > 1$ and $\|A_s^n\|_\alpha$ grows throughout the displayed interval. For $s \geq s_{-1}^{(N)}$, the displayed curves remain below 2.35 for $0 \leq n \leq 200$, showing finite transient amplification despite $\rho(A_s) \leq 1$.

For a one-period run, let

$$n_{\text{step}} := \left\lceil \frac{T}{\lambda_0 h} \right\rceil, \quad \Delta t := \frac{T}{n_{\text{step}}}, \quad \lambda_{\text{act}} := \frac{\Delta t}{h} \leq \lambda_0,$$

Full redistribution uses $s = 1$. The partial choice uses the nominal value L_0 , not the slightly smaller

$$L_{\text{act}} := \frac{\lambda_{\text{act}}}{\alpha}.$$

This keeps the prescribed redistribution fraction fixed across grid refinement rather than allowing the final-step adjustment in Δt to change s . Over the convergence grids $N \in \{100, 200, 400, 800\}$, L_{act} ranges from 2.6638 to 2.6663 when $\alpha = 0.15$ ($L_0 = 2.6667$), and from 3.9912 to 3.9990 when $\alpha = 0.10$ ($L_0 = 4$). The difference is below 0.9% of L_0 in every reported convergence run. This remains an empirical transfer of the model-derived rule to the higher-order discretization.

In this benchmark, evaluating $s_0(L_0)$ adds only the scalar formula (29); no eigenvalue computation, matrix assembly, or iterative calibration is performed during time stepping. The redistribution stage is the same merge-left operation used for $s = 1$, followed by the scalar blend in (46).

For the smooth sine-wave test we use

$$u(x, 0) = \frac{1}{2} + \frac{1}{2} \sin(2\pi x).$$

The Gaussian test uses

$$u(x, 0) = \exp \left[-\frac{1}{2} \left(\frac{d_{\mathbb{T}}(x, 0.25)}{0.06} \right)^2 \right], \quad d_{\mathbb{T}}(x, x_0) = \min_{k \in \mathbb{Z}} |x - x_0 + k|.$$

Reference cell averages are evaluated with 16-point Gauss–Legendre quadrature on each cell. Because the exact periodic solution returns to its initial profile at every integer multiple of T , the same quadrature rule provides the reference averages at $t = kT$. Let $\bar{u}_i(t)$ denote this reference cell

average in cell i at time t , and define

$$E_{1,h}(t) := \sum_{i=0}^{N-1} h_i |U_i(t) - \bar{u}_i(t)|, \quad E_{\infty,h}(t) := \max_i |U_i(t) - \bar{u}_i(t)|. \quad (48)$$

When comparing redistribution fractions, $E_{1,h}^{(s)}(t)$ and $E_{\infty,h}^{(s)}(t)$ denote the corresponding errors obtained with blend s . We also record the discrete mass drift

$$E_{M,h}(t) := \left| \sum_{i=0}^{N-1} h_i U_i(t) - \sum_{i=0}^{N-1} h_i U_i(0) \right|. \quad (49)$$

The one-period convergence tests report these quantities at $t = T$. Figure 6 shows convergence of $E_{1,h}(T)$ with ten separated cut cells and $N \in \{100, 200, 400, 800\}$. Least-squares slopes of $\log E_{1,h}(T)$ versus $\log h$ are used for the reported fitted orders. Both partial and full redistribution retain approximately second-order convergence. For $\alpha = 0.15$ ($L \approx 2.67$, $s_0 = 0.4$), fitted orders are 1.88 for partial redistribution and 1.94 for full redistribution. For $\alpha = 0.10$ ($L = 4$, $s_0 = 2/3$), the fitted orders are 1.92 and 1.94.

The effect is therefore primarily a change in the error constant. At $\alpha = 0.15$, the sine-wave $E_{1,h}(T)$ is reduced relative to full SRD by approximately 48.3%, 47.8%, 45.1%, and 41.0% as N increases from 100 to 800. At $\alpha = 0.10$, the reductions are approximately 25.4%, 25.1%, 24.8%, and 21.5%. Gaussian tests show the same trend with smaller gains. The quantitative accuracy claims below are restricted to the sine and localized Gaussian profiles.




Figure 6: Grid convergence of $E_{1,h}(T)$ for the smooth periodic sine-wave benchmark using the model-derived partial redistribution $s = s_0$ and full redistribution $s = 1$. Both choices retain approximately second-order convergence; fitted orders are reported in the text. The principal effect of partial redistribution is a lower error constant.

Table 1 reports the raw sine and localized-Gaussian values of $E_{1,h}(T)$ used for the convergence fits and fixed-error interpolation.

α	profile	N	$E_{1,h}^{(s_0)}(T)$	$E_{1,h}^{(1)}(T)$
0.15	sine	100	6.186342×10^{-3}	1.195548×10^{-2}
0.15	sine	200	1.512853×10^{-3}	2.899806×10^{-3}
0.15	sine	400	3.948102×10^{-4}	7.194984×10^{-4}
0.15	sine	800	1.065206×10^{-4}	1.805239×10^{-4}
0.15	Gaussian	100	2.859174×10^{-2}	3.890978×10^{-2}
0.15	Gaussian	200	8.543866×10^{-3}	1.167405×10^{-2}
0.15	Gaussian	400	2.606660×10^{-3}	3.116492×10^{-3}
0.15	Gaussian	800	7.118494×10^{-4}	8.034344×10^{-4}
0.10	sine	100	8.225149×10^{-3}	1.102404×10^{-2}
0.10	sine	200	1.990182×10^{-3}	2.658857×10^{-3}
0.10	sine	400	4.960524×10^{-4}	6.595976×10^{-4}
0.10	sine	800	1.303314×10^{-4}	1.661054×10^{-4}
0.10	Gaussian	100	3.260590×10^{-2}	3.754897×10^{-2}
0.10	Gaussian	200	9.581360×10^{-3}	1.114065×10^{-2}
0.10	Gaussian	400	2.767331×10^{-3}	3.019190×10^{-3}
0.10	Gaussian	800	7.398722×10^{-4}	7.847185×10^{-4}

Table 1: Measured one-period values of $E_{1,h}(T)$ for the sine and Gaussian tests. The values are direct simulation outputs; no interpolation is used in this table.

9.4 Interpolated total cell count at fixed $E_{1,h}(T)$

For the sine-wave data, the accuracy difference can also be expressed through the interpolated total cell count needed to reach a prescribed error. Let N_{est} denote the interpolated total cell count. For the sine-wave benchmark, we estimate N_{est} at $E_{1,h}(T) = 10^{-3}$ by linear interpolation in $\log N$ versus $\log E_{1,h}(T)$ between the measured $N = 200$ and $N = 400$ runs. The target error is bracketed by those two measured cell counts for both redistribution choices and both values of α , so no extrapolation is used.

For $\alpha = 0.15$, $N_{\text{est}} = 248$ for partial redistribution and $N_{\text{est}} = 340$ for full redistribution, a reduction of 27.1%. For $\alpha = 0.10$, the corresponding values are $N_{\text{est}} = 282$ and $N_{\text{est}} = 325$, a reduction of 13.3%. These are interpolated cell-count estimates, not wall-clock timing measurements.

9.5 Dependence on cut-cell count and layout

The preceding test uses ten approximately evenly separated cut cells. We therefore repeat the one-period sine and Gaussian tests on $N = 400$ cells while varying the number of non-adjacent cut cells over $m \in \{1, 2, 5, 10, 20\}$. For each m we use one evenly spaced layout and ten pseudorandom non-adjacent layouts. The pseudorandom layouts are generated with NumPy’s PCG64 generator using seed 20260808: indices are sampled without replacement from $\{5, \dots, N-6\}$ and a draw is rejected if any two selected cut cells are adjacent. The file `partial_srd_layout_robustness_seeded_raw.csv` records the cut indices for every accepted layout. All other numerical choices are unchanged. For each layout define the relative one-period error reduction

$$\mathcal{G}_1 := 100 \frac{E_{1,h}^{(1)}(T) - E_{1,h}^{(s_0)}(T)}{E_{1,h}^{(1)}(T)}. \quad (50)$$




Figure 7: Interpolated total cell count N_{est} at $E_{1,h}(T) = 10^{-3}$ in the sine-wave benchmark. Each estimate is obtained by log-log interpolation between the measured $N = 200$ and $N = 400$ runs; no extrapolation is used. The interpolation gives $N_{\text{est}} = 248$ versus 340 at $\alpha = 0.15$ and $N_{\text{est}} = 282$ versus 325 at $\alpha = 0.10$ for $s = s_0$ versus $s = 1$, respectively.

Thus $\mathcal{G}_1 > 0$ is equivalent to $E_{1,h}^{(s_0)}(T) < E_{1,h}^{(1)}(T)$.

Every $s = s_0$ and $s = 1$ run reaches $T = 1$, and $E_{M,h}(T) < 1.7 \times 10^{-15}$ in all seeded-layout runs. The accuracy difference grows with the number of redistributed neighborhoods. For a Gaussian profile with $\alpha = 0.15$, the mean $\mathcal{G}_1$ is 0.7%, 1.6%, 7.0%, 16.8%, and 36.2% as m increases from 1 to 20. For the smooth sine wave, the corresponding means are -0.7%, 0.9%, 21.0%, 44.6%, and 52.1%. At $\alpha = 0.10$ the same qualitative trend holds with smaller gains because $s_0 = 2/3$ is closer to full redistribution. Once $m \geq 5$, $\mathcal{G}_1 > 0$ for every tested layout for both profiles.

Practical guidance. For $m = 1$ or 2, the measured gain is small and can change sign. Within this benchmark, there is therefore little accuracy-based reason to replace full redistribution when only one or two isolated cut cells are active. The case for s_0 becomes stronger when several separated redistribution neighborhoods are present and the accumulated dissipation from full redistribution materially affects the error.

9.6 100-period higher-order test

We extend the sine and Gaussian benchmark to 100 advection periods with $N = 200$ and ten evenly spaced cut cells. The same nominal background CFL $\lambda_0 = 0.4$ is used. Partial redistribution uses $s_0 = 2/3$ for $\alpha = 0.10$ and $s_0 = 0.4$ for $\alpha = 0.15$; full redistribution uses $s = 1$.

All eight $s = s_0$ and $s = 1$ runs complete 100 periods, and $E_{M,h}(t) < 8 \times 10^{-14}$ at every recorded integer period. At $\alpha = 0.10$, $E_{1,h}(100T)$ is 0.148 for $s = s_0$ and 0.178 for $s = 1$ in the sine case, and 0.159 and 0.168, respectively, in the Gaussian case. At $\alpha = 0.15$, the corresponding




Figure 8: $\mathcal{G}_1$ from (50) versus the number m of separated cut cells. Each point is the mean over one evenly spaced and ten pseudorandom layouts; shaded bands show the observed range. *Left:* $\alpha = 0.15$. *Right:* $\alpha = 0.10$. Seed 20260808 and every accepted cut-cell index set are recorded with the raw data.

values are 0.123 and 0.187 for sine, and 0.151 and 0.171 for Gaussian. These are finite-duration numerical observations and are interpreted only as empirical boundedness over the tested interval. For the control runs, failure is declared when a non-finite component occurs or when $\|U^n\|_\infty > 10^8$. Under this criterion, the same higher-order scheme with $s = 0$ fails before $0.06T$ for both $\alpha = 0.10$ profiles at $N = 200$, whereas both $\alpha = 0.15$ control runs complete 100 periods. Table 2 collects the corresponding partial- and full-redistribution measurements.

α	profile	s	$E_{1,h}(100T)$	$\max_{0 \leq k \leq 100} E_{M,h}(kT)$
0.10	sine	s_0	0.148	4.85×10^{-14}
0.10	sine	1	0.178	7.74×10^{-14}
0.10	Gaussian	s_0	0.159	1.47×10^{-14}
0.10	Gaussian	1	0.168	2.35×10^{-14}
0.15	sine	s_0	0.123	1.11×10^{-14}
0.15	sine	1	0.187	2.53×10^{-14}
0.15	Gaussian	s_0	0.151	3.00×10^{-15}
0.15	Gaussian	1	0.171	7.61×10^{-15}

Table 2: One-hundred-period higher-order measurements for $N = 200$ and ten evenly spaced cut cells. Every listed run reaches $100T$. The last column is the maximum recorded mass drift over the integer-period samples.

The higher-order experiments lie outside the scope of Theorem 3.2. They quantify the accuracy difference between $s = s_0$ and $s = 1$ and record finite-time boundedness over the reported 100-period tests.

Scope of the numerical evidence

While Theorem 3.2 applies to the first-order periodic upwind operator with merge-left SRD in the fixed-local-CFL asymptotic regime, the singular-block framework in Section 4 is not tied to that specific stencil. The higher-order experiments test whether the model-derived value s_0 transfers to a standard second-order test bed using finite α , multiple separated cut cells, minmod reconstruction,




Figure 9: Measured $E_{1,h}(kT)$, $k = 0, \dots, 100$, for the higher-order benchmark with $N = 200$ and ten evenly spaced cut cells. Both methods complete each of the four combinations of profile and α , with $E_{M,h} < 8 \times 10^{-14}$ at the recorded integer periods. Left column: sine; right column: Gaussian. Top row: $\alpha = 0.10$; bottom row: $\alpha = 0.15$.

SSPRK2, and stage-wise redistribution. The observed transfer is empirical: the experiments do not establish a stability theorem for the higher-order scheme or a universal SRD threshold, and the results are not claimed for PPM, WENO, implicit integrators, adjacent cut cells, nonlinear systems, or multidimensional weighted-SRD neighborhoods. Those cases require constructing the corresponding defect-associated operator and analyzing how its spectrum moves under the chosen redistribution.

Reproducibility summary

Table 3 consolidates the numerical choices used in the higher-order experiments.

The accompanying raw-data files are `partial_srd_grid_convergence.csv`, `partial_srd_layout_robustness_verified_threshold_check.csv`, `verified_transient_time_history_summary.csv`, `verified_transient_time_long_time_N200_alpha_010_history_notation.csv`, and `long_time_N200_alpha_015_history_notation.csv`. The seeded layout file records every accepted cut-cell index set.

quantity	value or definition
domain and speed	$x \in [0, 1)$ periodic, $a = 1$, $T = 1$
mesh scale	$h = (N - m + m\alpha)^{-1}$
cut-cell width	αh
reconstruction	piecewise linear with minmod limiter
time integrator	SSPRK2, redistribution after each Euler stage
nominal background CFL	$\lambda_0 = 0.4$
partial blend	$s = s_0(L_0)$, $L_0 = \lambda_0/\alpha$
reference averages	16-point Gauss–Legendre quadrature per cell
Gaussian width	0.06 in $d_{\mathbb{T}}(x, 0.25)$
layout seed	NumPy PCG64, seed 20260808
control failure criterion	non-finite state or $\ U^n\ _{\infty} > 10^8$

Table 3: Numerical settings used in the higher-order benchmark. The convergence section gives the deterministic cut-cell placement formula, and the seeded-layout data file records every pseudo-random accepted index set.

10 Conclusion

The periodic upwind model with merge-left SRD yields a quantitative spectral calibration rule. For sufficiently small α , the base update has a localized eigenvalue outside $\mathbb{D}$ whose eigenvector is concentrated at the cut cell. The merge-left SRD correction approaches the same direction in the volume-weighted geometry. Tracking the localized branch $\mu_u(s)$ therefore determines the spectral redistribution threshold.

For the full N -cell blended operator

$$A_s = (1 - s)A + sA_{\text{SRD}},$$

fixed $L > 2$, and $N \geq 4$, the localized eigenvalue re-enters the closed unit disk through $\mu = -1$ at

$$s_{-1}^{(N)}(\alpha, L) = \frac{L-2}{L-1} - \frac{L(L-2)}{2(L-1)^2} \alpha^2 + O(\alpha^3).$$

Thus

$$s_0(L) = \frac{L-2}{L-1}$$

is asymptotically sufficient and differs from the exact local spectral stability threshold by $O(\alpha^2)$ in the regime of the theorem. The quantity $s_0(L)$ is the leading-order redistribution fraction associated with the return of μ_u to the unit disk. For moderate fixed $L > 2$, $s_0(L)$ can be substantially below 1, while $s_0(L) \rightarrow 1$ as $L \rightarrow \infty$.

The theorem is spectral; uniform power boundedness is a separate property. Direct finite-time calculations in the volume-weighted induced norm show transient amplification of approximately 2.01 for $\alpha = 0.10$ and 2.34 for $\alpha = 0.15$ at the exact $L = 4$ crossing, even though no eigenvalue lies outside the unit disk. Thus the spectral threshold eliminates exponential growth of the localized mode, while nonnormal transient growth can remain.

The broader singular-block analysis identifies what is structural and what is model-specific. A finite-dimensional defect matrix controls the unbounded spectrum of a class of singularly perturbed finite-volume operators, and the associated invariant subspace converges to the defect-coordinate subspace. The existence and localization of the defect-associated modes therefore generalize more readily than the scalar threshold formula. For a different flux, reconstruction, nonlinear system, or multidimensional redistribution rule, the corresponding defect block and the action of redistribution on its spectrum must be derived again.

The higher-order benchmark quantifies the accuracy difference between the two redistribution fractions. In the second-order minmod benchmark, partial and full redistribution retain essentially the same smooth-solution convergence order, but partial redistribution has a lower error constant. When the sine-wave convergence data are compared at fixed $E_{1,h}(T)$, log-log interpolation shows that partial redistribution reaches $E_{1,h}(T) = 10^{-3}$ with an estimated 27.1% fewer grid cells for $\alpha = 0.15$ and 13.3% fewer for $\alpha = 0.10$ than full redistribution.

In separate 100-period tests with $N = 200$, all eight $s = s_0$ and $s = 1$ runs for the sine and Gaussian profiles complete the integration, and $E_{M,h}(t) < 8 \times 10^{-14}$ at every recorded integer period. These observations provide finite-duration evidence of bounded evolution for the model-derived partial choice in this benchmark. Uniform stability remains outside the numerical claim.

The present result is a model theorem. A next extension is to replace the scalar one-cut-cell mode by the defect-associated spectral cluster of a multidimensional weighted-SRD neighborhood. If $\sigma_{\text{def}}(A_s)$ denotes the eigenvalue cluster continuously associated with the local defect modes, a local model could determine

$$s_* = \inf \left\{ s : \max_{\mu \in \sigma_{\text{def}}(A_s)} |\mu| \leq 1 \right\}.$$

This would turn the spectral mechanism developed here into an algorithm for calibrating redistribution from the local defect spectrum rather than prescribing full redistribution independently of it.

References

- [1] M. J. Berger. A note on the stability of cut cells and cell merging. *Applied Numerical Mathematics*, 96:180–186, 2015. doi:10.1016/j.apnum.2015.05.003.
- [2] M. Berger and A. Giuliani. A state redistribution algorithm for finite volume schemes on cut cell meshes. *Journal of Computational Physics*, 428:109820, 2021.
- [3] M. Berger and A. Giuliani. A new provably stable weighted state redistribution algorithm. *SIAM Journal on Scientific Computing*, 46(5):A2848–A2873, 2024. doi:10.1137/23M1597484.
- [4] M. J. Berger, C. Helzel, and R. J. LeVeque. H-box methods for the approximation of hyperbolic conservation laws on irregular grids. *SIAM Journal on Numerical Analysis*, 41(3):893–918, 2003.
- [5] P. Colella, D. T. Graves, B. J. Keen, and D. Modiano. A Cartesian grid embedded boundary method for hyperbolic conservation laws. *Journal of Computational Physics*, 211(1):347–366, 2006.
- [6] D. Devendran, D. T. Graves, H. Johansen, and T. Ligocki. A fourth-order Cartesian grid embedded boundary method for Poisson’s equation. *Communications in Applied Mathematics and Computational Science*, 12(1):51–79, 2017.
- [7] C. Engwer, S. May, A. Nüßing, and F. Streitbürger. A stabilized DG cut cell method for discretizing the linear transport equation. *SIAM Journal on Scientific Computing*, 42(6):A3677–A3703, 2020.
- [8] A. Giuliani, A. S. Almgren, J. B. Bell, M. J. Berger, M. T. Henry de Frahan, and D. Rangarajan. A weighted state redistribution algorithm for embedded boundary grids. *Journal of Computational Physics*, 464:111305, 2022. doi:10.1016/j.jcp.2022.111305.
- [9] A. Giuliani. A two-dimensional stabilized discontinuous Galerkin method on curvilinear embedded boundary grids. *SIAM Journal on Scientific Computing*, 44(1):A389–A415, 2022. doi:10.1137/21M1396277.
- [10] S. May and M. J. Berger. An explicit implicit scheme for cut cells in embedded boundary meshes. *Journal of Scientific Computing*, 71(3):919–943, 2017.
- [11] S. May and F. Laakmann. Accuracy analysis for explicit-implicit finite volume schemes on cut cell meshes. *Communications on Applied Mathematics and Computation*, 6(4):2239–2264, 2024.
- [12] L. Petri, G. Birke, C. Engwer, and H. Ranocha. The domain-of-dependence stabilization for cut-cell meshes is fully discretely stable. *SMAI Journal of Computational Mathematics*, 12:187–218, 2026. arXiv:2508.05372.
- [13] S. May and F. Streitbürger. DoD stabilization for non-linear hyperbolic conservation laws on cut cell meshes in one dimension. *Applied Mathematics and Computation*, 419:126854, 2022.
- [14] J. E. Karell. Update-magnitude state redistribution (UM-SRD): A shut-off extension of weighted SRD for cut-cell methods. Submitted to *Communications in Applied Mathematics and Computational Science*, 2026. Paper ID: 260523-Karell.