



BRoken-Space Additive Schwarz (BRAS) for Mass Matrix Preconditioning

FEM@LLNL Seminar Series

Ben S. Southworth

**Joint with Oliver Krzysik, Golo A. Wimmer,
Tzanio Kolev, and Steven Walton**

Mass matrix inverses

Time integration

$$M\dot{\mathbf{u}} = \mathbf{f}(\mathbf{u}, t), \quad \dot{\mathbf{u}} = M^{-1}\mathbf{f}.$$

- Wave, Maxwell, transport, and compressible-flow discretizations
- High-order Lagrangian hydrodynamics
- Every Runge–Kutta stage requires one or more mass solves
- Mass inverse only globally coupled operation in otherwise element-local update.

Motivation: Mass matrices are local to assemble, but their inverse is global, and sparse approximations can be limiting factor in computational efficiency. Very little work has been done on improved mass preconditioning.

Block preconditioning

$$\mathcal{A} = \begin{bmatrix} M & G^T \\ G & -C \end{bmatrix}, \quad -S = C + GM^{-1}G^T.$$

Replacing M^{-1} by X produces

$$-\tilde{S} = C + GXG^T.$$

The mass approximation affects both the leading block and the approximate Schur complement.

Finite-element mass matrix assembly

For a conforming space $V_h = \text{span}\{\phi_i\}$,

$$M_{ij} = (\phi_j, \phi_i)_{L^2(\Omega)}.$$

Break continuity element by element:

$$V_h^{\text{br}} = \bigoplus_{K \in \mathcal{T}_h} V(K), \quad M_{\text{br}} = \bigoplus_{K \in \mathcal{T}_h} M_K.$$

The assembly map $P : \mathbb{R}^n \rightarrow \mathbb{R}^{n_{\text{br}}}$ copies global coefficients into their element-local occurrences. It also carries orientation signs in $H(\text{curl})$ and $H(\text{div})$.

Assembly:

$$\begin{aligned} M &= P^\top M_{\text{br}} P \\ &= \sum_{K \in \mathcal{T}_h} P_K^\top M_K P_K, \end{aligned}$$

where P_K denotes local interpolation to element K .

Local work

Each M_K^{-1} can be applied independently. The difficulty lies in reconciling local solutions across interfaces.

Overview of BRAS

1. **Construction.** Invert each element mass matrix exactly, then use a weighted average to return to the conforming space.
2. **Theory.** Assembly constraints determine the nonunit spectrum. Tensor-product de Rham complexes reduce to a one-dimensional interface problem.
3. **Results.** BRAS lowers condition numbers, iteration counts, and wallclock times on conforming, graded, and nonconforming meshes, for standalone and block solves.

$$M = P^{\top} M_{\text{br}} P \quad \Rightarrow \quad B_W = R_W M_{\text{br}}^{-1} R_W^{\top}$$

The same element kernels and assembly maps already present in a finite-element code provide the complete algorithm.

Outline

Weighted BRAS

Spectral theory

Numerical results and applications

BRAS in broken and conforming spaces

For $W = W^\top \succ 0$ on the broken coefficient space, define the weighted left inverse

$$A_W = P^\top W P, \quad R_W = A_W^{-1} P^\top W, \quad R_W P = I.$$

1. Lift residual

$$r \mapsto R_W^\top r$$

Broken copies receive a weighted share of the residual.

2. Solve locally

$$R_W^\top r \mapsto M_{\text{br}}^{-1} R_W^\top r$$

Every local element mass inverse is independent.

3. Average

$$z = R_W M_{\text{br}}^{-1} R_W^\top r$$

Weighted averaging restores conforming coefficients.

The weighted BRAS inverse $B_W = R_W M_{\text{br}}^{-1} R_W^\top$ is SPD and applies unchanged in H^1 , $H(\text{curl})$, and $H(\text{div})$.

Three weight choices

Weight	Averaging	Result
$W = I$	Equal weight for every element copy	Original BRAS. Very cheap, but unaware of large differences in neighboring element masses.
$W = \text{diag}(M_{\text{br}})$	Weight by the local diagonal mass	Practical weighted BRAS. On a conforming mesh, $A_W = P^\top \text{diag}(M_{\text{br}})P = \text{diag}(M)$.
$W = M_{\text{br}}$	Full local mass weighting	Exact inverse, since $A_W = M$. This case identifies the ideal weighting but requires a global mass solve.

For $W = M_{\text{br}}$,

$$B_W = M^{-1} P^\top M_{\text{br}} M_{\text{br}}^{-1} M_{\text{br}} P M^{-1} = M^{-1} (P^\top M_{\text{br}} P) M^{-1} = M^{-1}.$$

$\Rightarrow$ Diagonal mass weighting is a cheap approximation to the exact weight.

Weighting on a strongly graded mesh

Consider three P_1 elements in one dimension with $h_0 + h_1 + h_2 = 1$, and

$$M_K = \frac{h_K}{6} \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix}, \quad \nu = \frac{h_{\max}}{h_{\min}}.$$

Identity weighting gives a bound proportional to ν . With $W_K = \text{diag}(M_K)$,


$$\frac{1}{2} W \preceq M_{\text{br}} \preceq \frac{3}{2} W.$$

Diagonal mass weighting removes the scale mismatch that causes identity-weighted BRAS to degrade under strong local grading.

h_1	$\kappa(B_I M)$	$\kappa(B_\Delta M)$
$3.33 \cdot 10^{-1}$	1.25	1.25
$3.33 \cdot 10^{-3}$	$7.55 \cdot 10^1$	1.00
$3.33 \cdot 10^{-5}$	$7.52 \cdot 10^3$	1.00
$3.33 \cdot 10^{-7}$	$7.52 \cdot 10^5$	1.00

Equal averaging lets a tiny element influence an interface as much as its much larger neighbor.

Nonconforming assembly with hanging nodes



Separate the assembly map into

$$\mathbb{R}^n \xrightarrow{E} \mathbb{R}^{n_{\text{msh}}} \xrightarrow{P_0} \mathbb{R}^{n_{\text{br}}}, \quad P = P_0 E.$$

- E interpolates constrained mesh DOFs from independent true DOFs
- P_0 copies mesh DOFs into element-local occurrences
- The mass identity remains $M = P^\top M_{\text{br}} P$
- For linear midpoint constraint, $u_h = \frac{1}{2}u_0 + \frac{1}{2}u_1$.

Weighted averaging on nonconforming meshes

For diagonal W , first assemble on the unconstrained mesh-DOF space:

$$D_W = P_0^\top W P_0.$$

D_W is diagonal, but the true-DOF averaging matrix is

$$A_W = P^\top W P = E^\top D_W E.$$

With unconstrained and hanging mesh DOFs ordered separately,

$$E = \begin{bmatrix} I \\ C \end{bmatrix}, \quad D_W = \begin{bmatrix} \Lambda_u & 0 \\ 0 & \Lambda_h \end{bmatrix}, \quad \boxed{A_W = \Lambda_u + C^\top \Lambda_h C}.$$

Only true DOFs that share a hanging-node interpolation stencil become coupled.

Apply A_W^{-1} by sparse factorization, component-wise Woodbury, or a diagonal approximation.

Hanging nodes change the averaging solve, not the element-local BRAS kernel.

Essentially matrix-free MFEM application

$$B_W = A_W^{-1} G_W A_W^{-1}, \quad G_W = P^\top W M_{\text{br}}^{-1} W P.$$

One application

1. Apply the right averaging solve A_W^{-1}
2. Prolong true DOFs to element DOFs with P
3. Apply $W M_{\text{br}}^{-1} W$ element by element
4. Assemble with $P^\top$
5. Apply the left averaging solve A_W^{-1}

Tensor-product elements

$$M_K = M_1 \otimes \cdots \otimes M_d,$$

$$M_K^{-1} = M_1^{-1} \otimes \cdots \otimes M_d^{-1}.$$

Sum factorization gives $\mathcal{O}(p^{d+1})$ element work and reuses one-dimensional reference inverses.

MFEM already provides the element restriction, prolongation, and AMR constraint operators.

Outline

Weighted BRAS

Spectral theory

Numerical results and applications

A general weighted-BRAS bound

Assume the local weight approximates the broken mass matrix:

$$\alpha W \preceq M_{\text{br}} \preceq \beta W, \quad 0 < \alpha \leq \beta.$$

Then the weighted-BRAS-preconditioned spectrum satisfies

$$1 \leq \lambda(B_W M) \leq \boxed{\mathcal{K}(\alpha, \beta) = \frac{(\alpha + \beta)^2}{4\alpha\beta}}.$$

In comparison, the assembled weight $A_W = P^\top W P$ satisfies

$$\tilde{\kappa}(M, A_W) \leq \frac{\beta}{\alpha}.$$

- Corresponding Chebyshev factors satisfy $\rho_{\text{cheb}}(\mathcal{K}) = \rho_{\text{cheb}}(\beta/\alpha)^2$.
- Generic theory predicts about half as many CG iterations for BRAS as Jacobi.
- Theory can be tight (e.g., GLL basis), but in other cases is very loose (Bernstein).

Assembly constraints control the spectrum

Let $\mathcal{C} \in \mathbb{R}^{q \times n_{\text{br}}}$ encode continuity:

$$\ker(\mathcal{C}) = \text{range}(P), \quad q = n_{\text{br}} - n.$$

Scale by W :

$$\tilde{M}_{\text{br}} = W^{-1/2} M_{\text{br}} W^{-1/2}, \quad \tilde{\mathcal{C}} = \mathcal{C} W^{-1/2},$$

and define the $q \times q$ constraint-space matrices

$$G_- = \tilde{\mathcal{C}} \tilde{M}_{\text{br}}^{-1} \tilde{\mathcal{C}}^\top, \quad G_0 = \tilde{\mathcal{C}} \tilde{\mathcal{C}}^\top, \quad G_+ = \tilde{\mathcal{C}} \tilde{M}_{\text{br}} \tilde{\mathcal{C}}^\top.$$

Then

$$\text{spec}_{\neq 1}(B_W M) = \text{spec}_{\neq 1}(G_0^{-1} G_+ G_0^{-1} G_-).$$

A full problem on n true DOFs reduces to a problem on the q interface constraints.

BRAS is exact on large subspace

For a suitable $Z_C \in \mathbb{R}^{n \times q}$,

$$B_W = M^{-1} + Z_C G^{-1} Z_C^T.$$

Therefore,

$$B_W - M^{-1} \succeq 0, \quad \text{rank}(B_W - M^{-1}) \leq q.$$

Unit eigenvalues

At least $n - q$ eigenvalues of $B_W M$ equal one when $q \leq n$.

Nonunit eigenvalues

Every departure from exact inverse is from enforcing interelement continuity.

- In 1D H^1 , q is number of interior interfaces, independent of p .
- Higher order adds interior element modes on which BRAS is already exact.

Tensor-product de Rham complexes

Use a closed one-dimensional H^1 space V_c and an open discontinuous space V_o . In three dimensions,

$$V^0 : V_c \otimes V_c \otimes V_c \subset H^1,$$

$$V^1 : (V_o \otimes V_c \otimes V_c)\mathbf{e}_x \oplus \text{cyclic terms} \subset H(\text{curl}),$$

$$V^2 : (V_c \otimes V_o \otimes V_o)\mathbf{e}_x \oplus \text{cyclic terms} \subset H(\text{div}),$$

$$V^3 : V_o \otimes V_o \otimes V_o \subset L^2.$$

For $W = \text{diag}(M_{\text{br}})$, assembly, mass matrices, and BRAS all tensorize. Open factors are exact because their DOFs are not shared between elements.

$\Rightarrow$ Only the closed factors contribute nonunit eigenvalues.

Tensor products preserve clusters

Suppose each closed one-dimensional factor satisfies

$$\tilde{\kappa}(M_C, M_{\text{bras},c}) \leq \kappa_1 d.$$

For the degree- r space $V^r(d)$,

$$\tilde{\kappa}(M, M_{\text{bras}}) \leq \kappa_1 d^{d-r}.$$

The exponent counts closed tensor factors:

- H^1 : $d - r = d$
- $H(\text{curl})$ in 3D: $d - r = 2$
- $H(\text{div})$: $d - r = 1$
- L^2 : $d - r = 0$, hence exact


$\implies$ tensor de Rham complex reduces to one closed 1d mass problem.

If the 1D spectrum lies in $\{1\} \cup [\alpha, \beta]$, then


$$\text{spec}(M_{\text{bras}}^{-1} M) \subset \{1\} \cup \bigcup_{\ell=1}^{d-r} [\alpha^\ell, \beta^\ell].$$

There are at most $d - r + 1$ clusters.

Polynomial-degree dependence in 1d





-□- GaussLobatto: diag -□- ClosedGL: diag
 -○- GaussLobatto: BRAS -○- ClosedGL: BRAS



-□- Positive: diag -□- ClosedUniform: diag
 -○- Positive: BRAS -○- ClosedUniform: BRAS

Uniform 1D mesh with eight elements, × shows theoretical predictions.

Tensor-product verification



Verification of theory on 3D Cartesian product mesh with direction-dependent grading.
Black $\times$ show theory, open markers show numerical computation.

Nonunit cluster predicts CG convergence

For 1d closed space, a reflection-invariant reference basis, and $n_{1d} \geq 2$,

$$\text{spec}(M_{\text{bras}}^{-1}M) = \underbrace{\{1\}^{n_{1d}(p-1)+2}}_{\text{unit cluster}} \cup \underbrace{\{\mu_{\mathcal{G}}(x) : x \in \text{spec}(X_h)\}}_{\text{at most } n_{1d}-1 \text{ nonunit eigenvalues}},$$

for $\text{spec}(X_h) \subset [-1, 1]$. Define endpoints of nonunit cluster by

$$\lambda_- = \min_{x \in \text{spec}(X_h)} \mu_{\mathcal{G}}(x), \quad \lambda_+ = \max_{x \in \text{spec}(X_h)} \mu_{\mathcal{G}}(x).$$



Then

$$\underbrace{\tilde{\kappa}_{\text{c}}(M, M_{\text{bras}})}_{\text{full spectral condition number}} = \lambda_+,$$

$$\underbrace{\tilde{\kappa}(\mathcal{G})}_{\text{condition number of the nonunit cluster}} = \frac{\lambda_+}{\lambda_-}.$$

$\tilde{\kappa}(\mathcal{G})$ can decay with p , where nonunit eigenvalues cluster more tightly.

Eigenvalue clustering predicts CG behavior



Uniform mesh, $n_{1d} = 64$, random right-hand side, relative tolerance 10^{-5} . Global condition number may grow modestly with p , but **we observe super convergence in CG w.r.t. p .**

Outline

Weighted BRAS

Spectral theory

Numerical results and applications

Target classes of problems

Standalone solves

$$Mu = f$$

Triangles, tetrahedra, quadrilaterals, and hexahedra in H^1 , $H(\text{curl})$, and $H(\text{div})$. Random right-hand sides and relative tolerances 10^{-8} to 10^{-10} .

Nonuniform meshes

Strong conforming grading, static hanging-node meshes, and dynamic AMR. Tests compare identity and diagonal-mass weights across nodal and Bernstein bases.


Block systems

$$\begin{bmatrix} M & G^T \\ G & -C \end{bmatrix}$$

Mixed Poisson and a mixed anisotropic biharmonic (fourth-order). BRAS appears in the mass block and inside an AMG-preconditioned approximate Schur complement.

Standalone mass solves on simplicial meshes

Mass-matrix solve time ratio $t_{\text{diag}}/t_{\text{bras}}$.



Problem setup


- About 10^6 conforming DOFs
- $p = 1, \dots, 4$
- CG tolerance 10^{-10}
- Firedrake implementation
- Setup time excluded
- Original $W = I$ BRAS

At $p = 4$, representative iteration counts were


	Jacobi	BRAS
2D H^1 Bernstein	180	23
3D $H(\text{div})$	215	30

⇒ BRAS was faster in every tested standalone case despite its more expensive application.


Conforming grading: Sod shock tube



Initial density and grading

Final density at $t = 0.2$

Laghos PA velocity mass matrix on 128×8 extrusion of Sod problem. H^1 , $\rho = 4$, PCG tolerance 10^{-8} , cell-width jump $J \in \{1, 4, 8, 16, 32, 64\}$.



Nonconforming AMR in two dimensions

CG iteration counts for base ref 2, 4 refinement levels, and nonconforming limit of 10. Each entry is **diagonally weighted BRAS** (**identity-weighted BRAS**) / **Jacobi**; 5000* denotes the iteration limit.

 H^1
 $H(\text{div})$

Basis	$p = 3$	$p = 6$
GaussLobatto	10(38)/23	9(40)/21
ClosedGL	11(38)/27	10(41)/26
ClosedUniform	12(39)/31	19(52)/67
Positive	30(38)/324	47(56)/5000*


Basis	$p = 3$	$p = 6$
GaussLobatto	8(8)/14	7(9)/12
ClosedGL	9(9)/26	19(19)/90
ClosedUniform	20(19)/69	45(33)/563

- For Bernstein H^1 at $p = 6$, weighted BRAS reduces iteration count by more than 100 $\times$.
- Identity weighting can be better in $H(\text{div})$, especially at higher order.

Nonconforming AMR in three dimensions

CG iteration counts for base ref 0, 4 refinement levels, and nonconforming limit of 10. Each entry is diagonally weighted BRAS (identity-weighted BRAS) / Jacobi; 5000* denotes iteration limit.

Space	Basis	$p = 3$	$p = 5$
H^1	GaussLobatto	13(87)/35	12(94)/32
H^1	ClosedGL	14(84)/44	13(91)/42
H^1	ClosedUniform	14(85)/52	22(107)/104
H^1	Positive	51(53)/1663	88(92)/5000*
$H(\text{curl})$	GaussLobatto	10(24)/24	9(27)/22
$H(\text{curl})$	ClosedGL	11(26)/43	16(36)/89
$H(\text{curl})$	ClosedUniform	22(40)/118	54(83)/634
$H(\text{div})$	GaussLobatto	8(10)/14	7(9)/13
$H(\text{div})$	ClosedGL	12(10)/40	24(22)/173
$H(\text{div})$	ClosedUniform	29(23)/241	90(66)/3446




The preferred weight depends on the space and basis.

For ClosedUniform $H(\text{div})$ at $p = 5$, identity weighting requires 66 iterations, compared with 90 for diagonal weighting and 3446 for Jacobi.


Laghos: repeated mass solves on evolving meshes

Laghos discretizes velocity in $V_h \subset H^1(\Omega)^d$ and internal energy in a discontinuous thermodynamic space. The semidiscrete equations take the form

$$M_v \dot{\mathbf{v}} = -\mathbf{F}\mathbf{1}, \quad M_e \dot{\mathbf{e}} = \mathbf{F}^\top \mathbf{v}.$$



Sedov blast




Triple-point problem


Kinematic mesh evolves every step. With AMR, element sizes and hanging-node constraints also change. Laghos with AMR uses full weighted, nonconforming BRAS.

Dynamic AMR: Sedov blast


Quarter domain $[0, 1]^2$, RK4 with $p \in \{2, 3, 4\}$, dynamic refinement and derefinement. Iterations averaged over accepted time steps. *Full assembly used for all matrices.*



Initial mesh



Final density and mesh



Mass-inverse bounds transfer to Schur complements

Consider

$$-S_C = C + GM^{-1}G^T, \quad -\tilde{S}_C = C + GXG^T, \quad X \succ 0.$$

For $C = 0$, set $\mathbf{g} = G^T \mathbf{q}$. Then

$$\frac{\mathbf{q}^T GXG^T \mathbf{q}}{\mathbf{q}^T GM^{-1}G^T \mathbf{q}} = \frac{\mathbf{g}^T X \mathbf{g}}{\mathbf{g}^T M^{-1} \mathbf{g}}.$$

Therefore,

$$\boxed{\tilde{\kappa}(S_0, \tilde{S}_0) \leq \tilde{\kappa}(M, X^{-1})}.$$

More sharply, the Schur complement samples the mass approximation only on $\text{range}(G^T)$. A positive C can further moderate the error.

Improving the mass inverse improves the approximate Schur complement, but only the gradient or divergence range matters.

Block preconditioning: mixed Poisson


Mixed problem $\sigma - \nabla u = 0$ and $\nabla \cdot \sigma = -f$ gives

$$\begin{bmatrix} M & B^\top \\ B & 0 \end{bmatrix} \begin{bmatrix} \sigma \\ u \end{bmatrix} = \begin{bmatrix} 0 \\ -f \end{bmatrix}$$

for $\text{BDM}_p/\text{DG}_{p-1}$ mixed pair.
MINRES uses unweighted BRAS or Jacobi in the flux mass block and one AMG V-cycle for

$$-\tilde{S}_X = BXB^\top.$$

Tris/tets, $p = 1, 2$, random source, tolerance 10^{-10} . At $p = 1$, “hybrid” uses BRAS only for flux block.




$\Rightarrow$ At $p = 2$, BRAS improves both blocks and gives best time, with $1.3\text{--}2\times$ speedup.

Block preconditioning: mixed anisotropic biharmonic

Two-dimensional mixed anisotropic biharmonic system in $H^1 \times H^1$

$$\begin{bmatrix} M & C^\top \\ -C & M \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} = \begin{bmatrix} 0 \\ \mathbf{f} \end{bmatrix}.$$

FGMRES(30) uses a lower triangular preconditioner with $\tilde{S}_X = M + CXC^\top$, for $X \in \{\text{diag}(M)^{-1}, B_W\}$. Inner mass and Schur solves use CG with 1 LS-AMG-DD for the Schur complement.



Conclusions

Algorithm

- Exact element mass inverses in the broken space
- Weighted averaging back to true DOFs
- Diagonal mass weights handle strong local grading
- Hanging-node constraints enter only through assembly

Results

- Assembly constraints determine the nonunit spectrum
- Tensor-product de Rham theory predicts strong clustering
- BRAS is particularly effective at high order and for difficult bases
- Faster standalone and block solves in the reported tests

$$B_W = R_W M_{\text{br}}^{-1} R_W^{\top}$$

Local finite-element ingredients produce sparse approximation to global inverse.

Papers

1. O. A. Krzysik, B. S. Southworth, and G. A. Wimmer, **Broken-space Additive Schwarz Mass Inverse Approximations and (Block) Preconditioning.**
2. O. A. Krzysik, B. S. Southworth, and G. A. Wimmer, **Constraint-space Spectral Theory for Broken-space Additive Schwarz Mass Preconditioning.**
3. O. A. Krzysik, B. S. Southworth, T. Kolev, and S. T. Walton, **Mass Matrix Preconditioning on Nonuniform and Nonconforming Meshes with Broken-space Additive Schwarz.**

Questions?