Geometric adaptive smoothed aggregation multigrid for DG discretisations

Yulong Pan, Michael Lindsey, Per-Olof Persson

Department of Mathematics, University of California, Berkeley
Mathematics Department, Lawrence Berkeley National Laboratory

ICOSAHOM Montréal 2025

July 15, 2025

BERKELEY LAB
Bringing Science Solutions to the World

Yulong Pan, Michael Lindsey and Per-Olof Persson. Geometric adaptive smoothed aggregation multigrid for discontinuous Galerkin discretisations. Submitted to Journal of Computational Physics. https://arxiv.org/abs/2504.13373

Background

Discontinuous Galerkin (DG) methods

- Variant of Finite Element method allowing for discontinuities along element boundaries, with Finite Volume fluxes used for stabilisation
- Popular for discretisation of convection-diffusion equations in fluids and beyond
- High-order accurate, suitable for use on unstructured meshes in 2D/3D
- However, expensive to apply in practice, more work required to resolve large complex linear systems

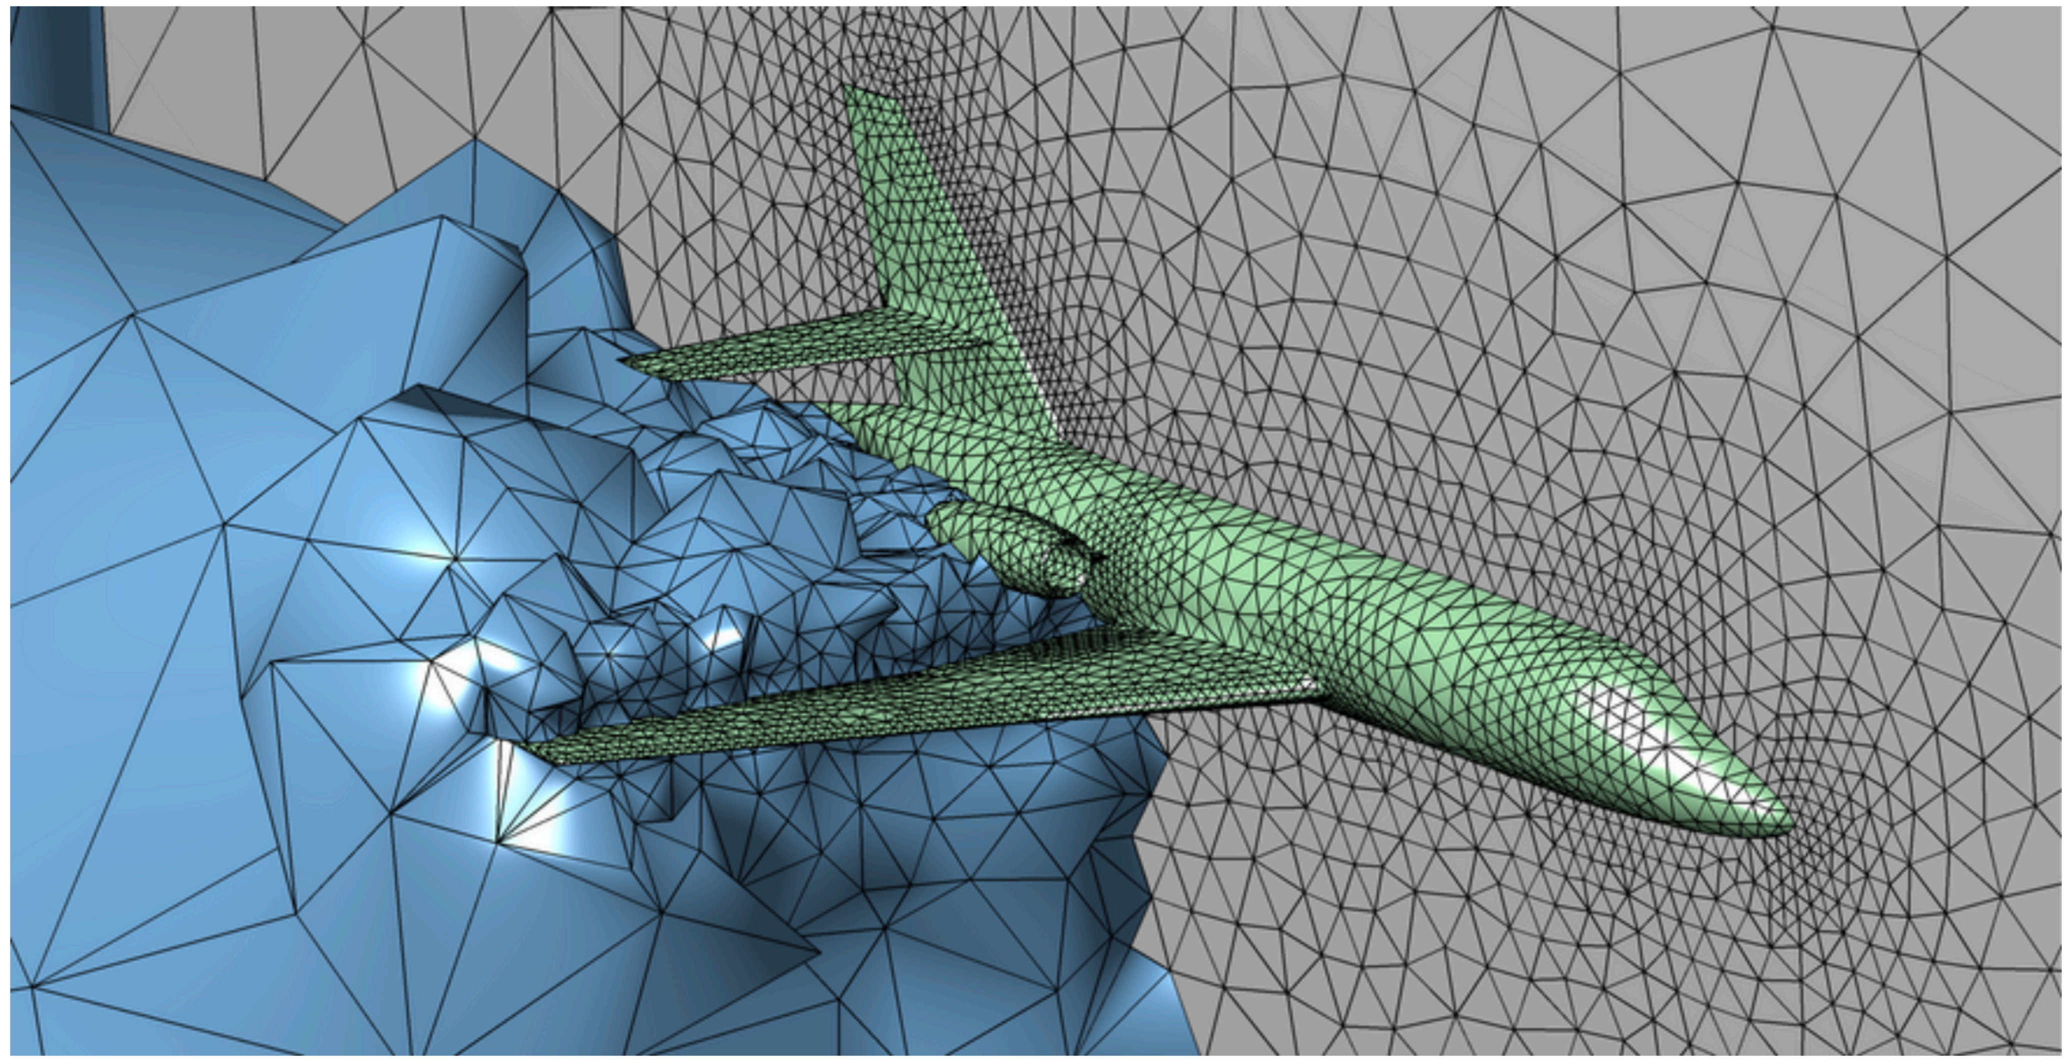

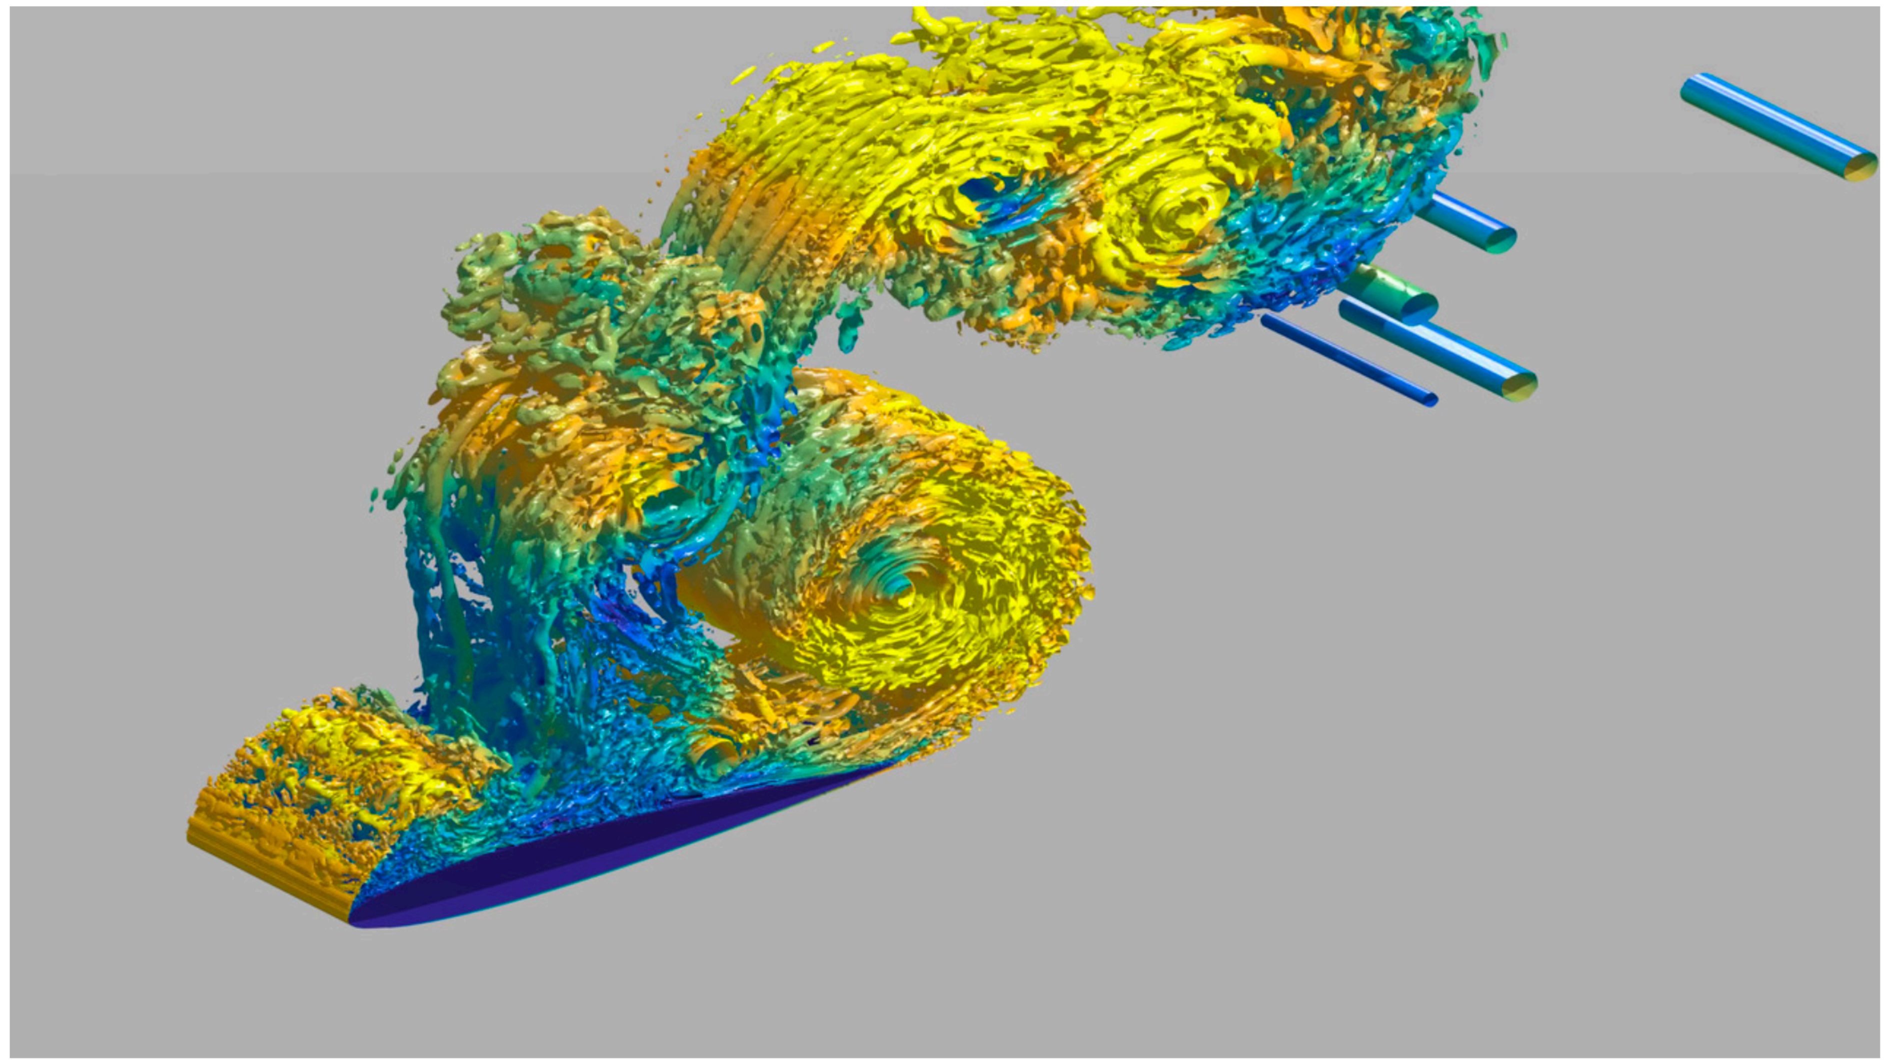

Linear solvers

- One advantage of DG - block sparsity patterns
- Some examples of popular solvers
    - General purpose: block Jacobi/block Gauss-Seidel
    - Hyperbolic: Incomplete LU factorisations
    - Elliptic: Multigrid
- However many techniques scale optimally OR cannot be applied general purpose for large classes of equations
- More work needed

Convection-diffusion equation

- Model equation for this talk:

$$\mathbf{v} \cdot \nabla u - \mu \Delta u = f$$

- After discretisation gives general linear system of form:

$$Au = f$$

- Challenges with this system:
    - Mesh of underlying domain is unstructured
    - Equation has both hyperbolic and elliptic character
    - Equations might be stiff

Multigrid (MG) methods

- Originally designed for elliptic equations on structured domains
- Main idea:
  - Recursively solve linear algebra problem on nested hierarchy
  - Successive levels resolve distinct *high* frequencies
  - Achieve asymptotically optimal linear O(N) runtime

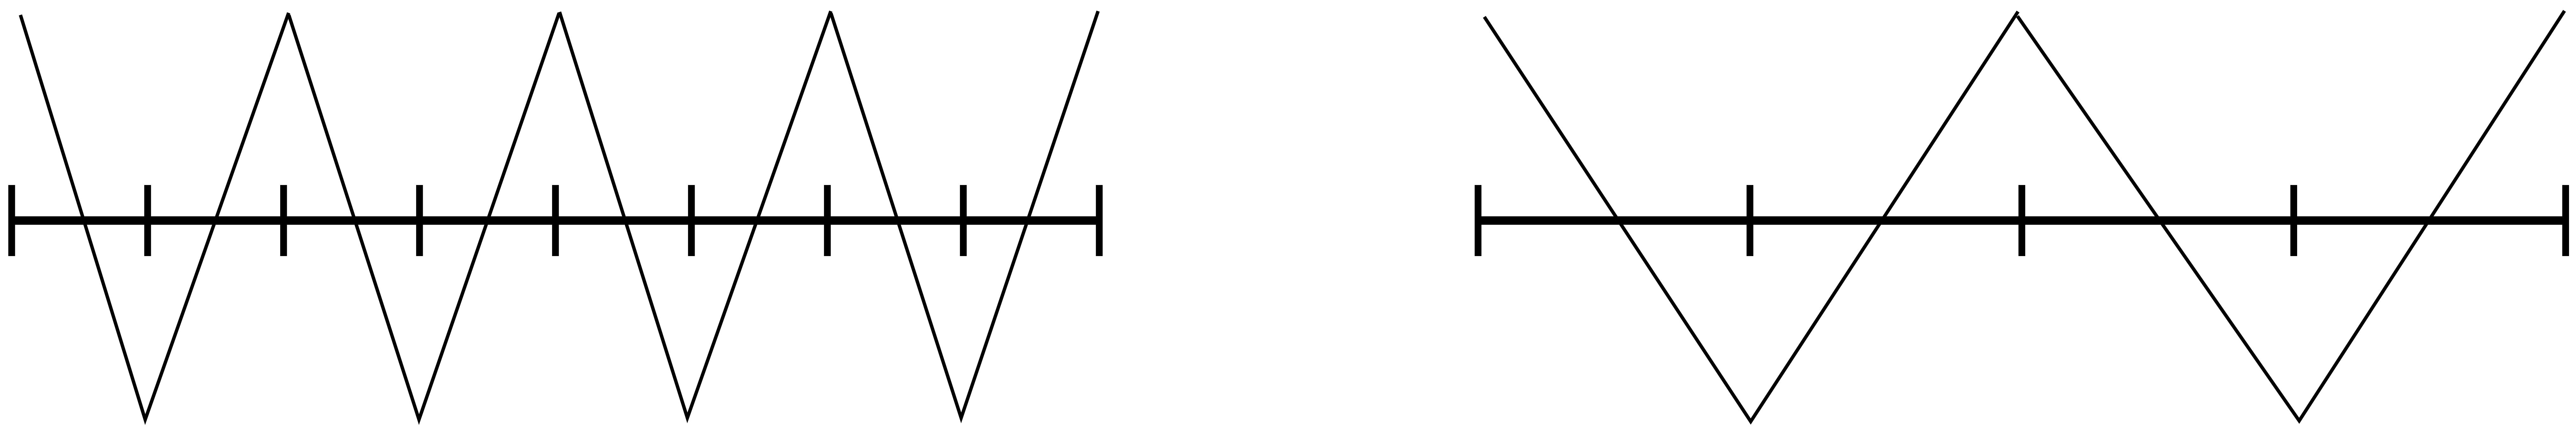

Multigrid (MG) methods

- Ingredients in general MG method:
  - Restriction R_k^{k+1}
    - Transfer of residual from higher to lower level
  - Smoother S_k
    - Relax high frequency modes at each level, e.g. Jacobi, Gauss-Seidel
  - Interpolation T_{k+1}^k
    - Transfer of solution from lower to higher level

Geometric multigrid (GMG)

- Hierarchy is formed using mesh discretisation, generally nested mesh constructions are used

- Direct polynomial injection generally used for interpolation operator, restriction taken to be its adjoint

$$R_k^{k+1} = \left(T_{k+1}^{k}\right)^T$$

- BUT, hard to form on unstructured meshes, some approaches include element agglomeration[1], mesh decimation, non-nested triangulations



[1]Yulong Pan, Per-Olof Persson. Agglomeration-based geometric multigrid solvers for Compact Discontinuous Galerkin discretisations on unstructured meshes. *J. Comp. Phys.*, Vol 454, 110906, April 2022. https://arxiv.org/abs/2012.08024

Algebraic multigrid (AMG)

- Introduced as an alternative to geometric multigrid methods
- Multigrid hierarchy is constructed blind to the underlying mesh, instead hierarchy is formed using only entries of the matrix
- Help ease the problem of constructing unstructured mesh hierarchies
- Two main variants: Classical Ruge-Stuben and Smoothed Aggregation AMG

- However, reliance on the matrix means its entries must be readily accessible
- Can be difficult to make fully matrix free

Smoothed aggregation (SA)

- Introduced in Vanek et al., *Algebraic multigrid by smoothed aggregation for second and fourth order*. (1996)
- Fundamentally a method of constructing interpolation/restriction operators for AMG
- Assume given matrix A, two partitions B_1, B_2 of its degrees of freedom defining next level in MG hierarchy
  - Also assume given a set of vectors $V = \{v_1, \ldots, v_j\}$ representing the column space of the next level in AMG hierarchy

Smoothed aggregation (SA)

- Assume given matrix A, two partitions B_1, B_2 of its degrees of freedom
    - Also given a set of vectors $V = \{v_1, \ldots, v_j\}$ representing the column space of the next level in AMG hierarchy

$$
\begin{matrix} B_1 \left\{ \begin{matrix} \\ \\ \\ \end{matrix} \right. \\ B_2 \left\{ \begin{matrix} \\ \\ \\ \end{matrix} \right. \end{matrix}
\begin{pmatrix}
v_{11} & & v_{1j} \\
\vdots & \cdots & \vdots \\
v_{k1} & & v_{kj} \\
v_{(k+1)1} & & v_{(k+1)j} \\
\vdots & \cdots & \vdots \\
v_{n1} & & v_{nj}
\end{pmatrix}
\longrightarrow
\begin{pmatrix}
v_{11} & & v_{1j} & & & \\
\vdots & \cdots & \vdots & & & \\
v_{k1} & & v_{kj} & & & \\
& & & v_{(k+1)1} & & v_{(k+1)j} \\
& & & \vdots & \cdots & \vdots \\
& & & v_{n1} & & v_{nj}
\end{pmatrix}
$$

$$V \qquad\qquad\qquad\qquad P_1^0$$

- Coarse modes V are split according to partitions in block diagonal matrix

Smoothed aggregation (SA)

- Coarse modes V are split according to partitions in block diagonal matrix

$$
\begin{matrix} B_1 \left\{ \begin{matrix} \\ \\ \\ \end{matrix} \right. \\ \\ B_2 \left\{ \begin{matrix} \\ \\ \\ \end{matrix} \right. \end{matrix}
\begin{pmatrix}
v_{11} & & v_{1j} \\
\vdots & \cdots & \vdots \\
v_{k1} & & v_{kj} \\
v_{(k+1)1} & & v_{(k+1)j} \\
\vdots & \cdots & \vdots \\
v_{n1} & & v_{nj}
\end{pmatrix}
\longrightarrow
\begin{pmatrix}
v_{11} & & v_{1j} & & & \\
\vdots & \cdots & \vdots & & & \\
v_{k1} & & v_{kj} & & & \\
& & & v_{(k+1)1} & & v_{(k+1)j} \\
& & & \vdots & \cdots & \vdots \\
& & & v_{n1} & & v_{nj}
\end{pmatrix}
$$

$$V \qquad\qquad\qquad\qquad P_1^0$$

- To finish defining prolongation operator, smoothing is applied (e.g. Jacobi)

$$T_1^0 = S_0 P_1^0$$

- Smooth out high frequency modes introduced by partitioning

GMG vs AMG

Geometric MG

- Mesh hierarchies
- Difficult to generalise on unstructured meshes
- Matrix A not explicitly necessary

Algebraic MG

- Algebraic hierarchies inferred from A
- Entirely blind to underlying mesh
- Matrix A needed explicitly

- Is there a way to combine the two to obtain fast method suitable for unstructured meshes?

Further challenges

- MG difficult to generalise to both hyperbolic and elliptic equations
    - Some previous work include L-AIR solvers (Southworth et al. 2017)
- MG methods are also sensitive to choice of DG numerical flux
    - In particular for Laplacian, different constructions required for Interior Penalty (IP) vs Local DG (LDG) fluxes (Fortunato et al. 2019)
- Can also struggle on stiff systems

*Apologies for the myriad abbreviations, we know there are a lot of them in this field unfortunately

Geometric adaptive SA multigrid for DG

Adaptive smoothed aggregation (αSA)

- Introduced by Brezina et al. *Adaptive smoothed aggregation (αSA) multigrid*. (2005)
- In original SA algorithm, coarse modes V are given a priori based off expected low frequency modes of Laplacian (e.g. constant mode)
    - Doesn't necessarily generalise to other problems
- Idea of αSA: adaptively find coarse modes by applying MG smoother on some random vectors
- Why? The leftover modes should span the column space of the next level of the hierarchy!

Adaptive smoothed aggregation (αSA)

1. Assume given matrix A, and partitions of its degrees of freedom $B_1, \ldots B_m$
2. Define smoother S
   e.g. Jacobi, $S = I - \omega D^{-1}A$, $\omega = 2/3$, $D =$ diagonal of A
3. ***Form random matrix with random Gaussian entries*** $(b_1 \ldots b_r)$
4. ***Apply S to each random vector*** p ***times to obtain coarse space vectors*** V
5. Partition V according to $B_1, \ldots B_m$
6. ***Apply SVD to each diagonal block to filter which modes to keep***
7. Form prolongation operator as before

Adaptive smoothed aggregation (αSA)

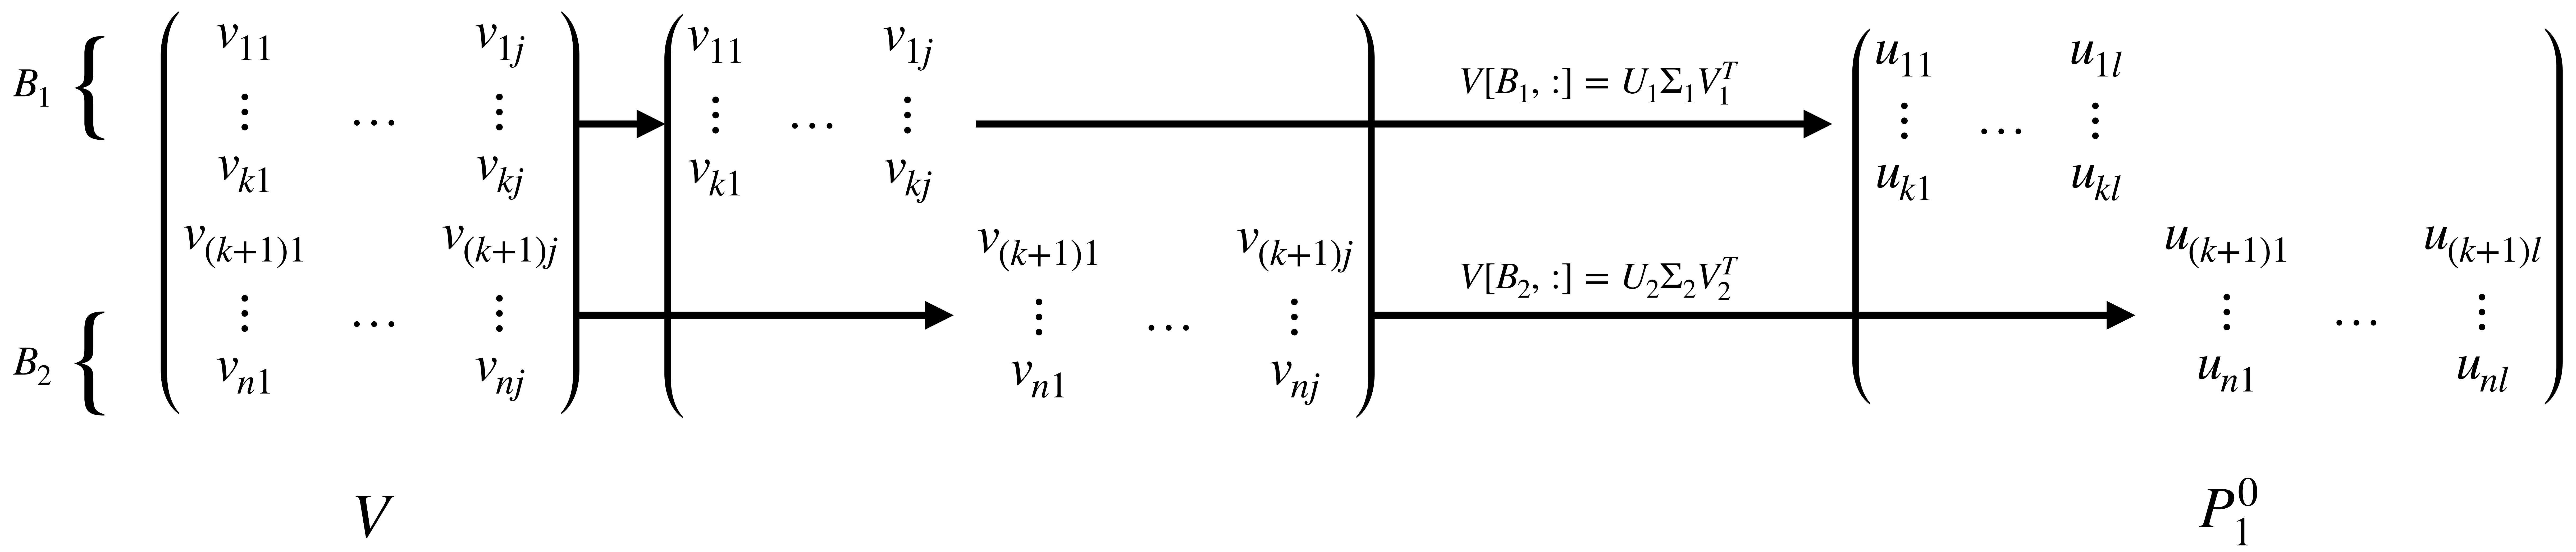

- Additional SVD step to filter out relevant modes in each block
- Difference in αSA vs SA lies in how V is formed
    - Pre-determined in SA, adaptive found via smoothing random vectors in αSA
- Partitions $B_1, \ldots B_m$ are formed using standard AMG techniques, which we will not delve into in this talk

Geometric αSA for DG

- Main idea: Form mesh partitions $B_1, \ldots B_m$ using mesh information
- To form each partition block B_i, agglomerate elements together using graph partitioning algorithms (e.g. METIS)
- Advantage: Removes explicit dependence on matrix A, αSA construction allows for generalisation to unstructured meshes
    - Combine advantages of GMG and smoothed aggregation AMG

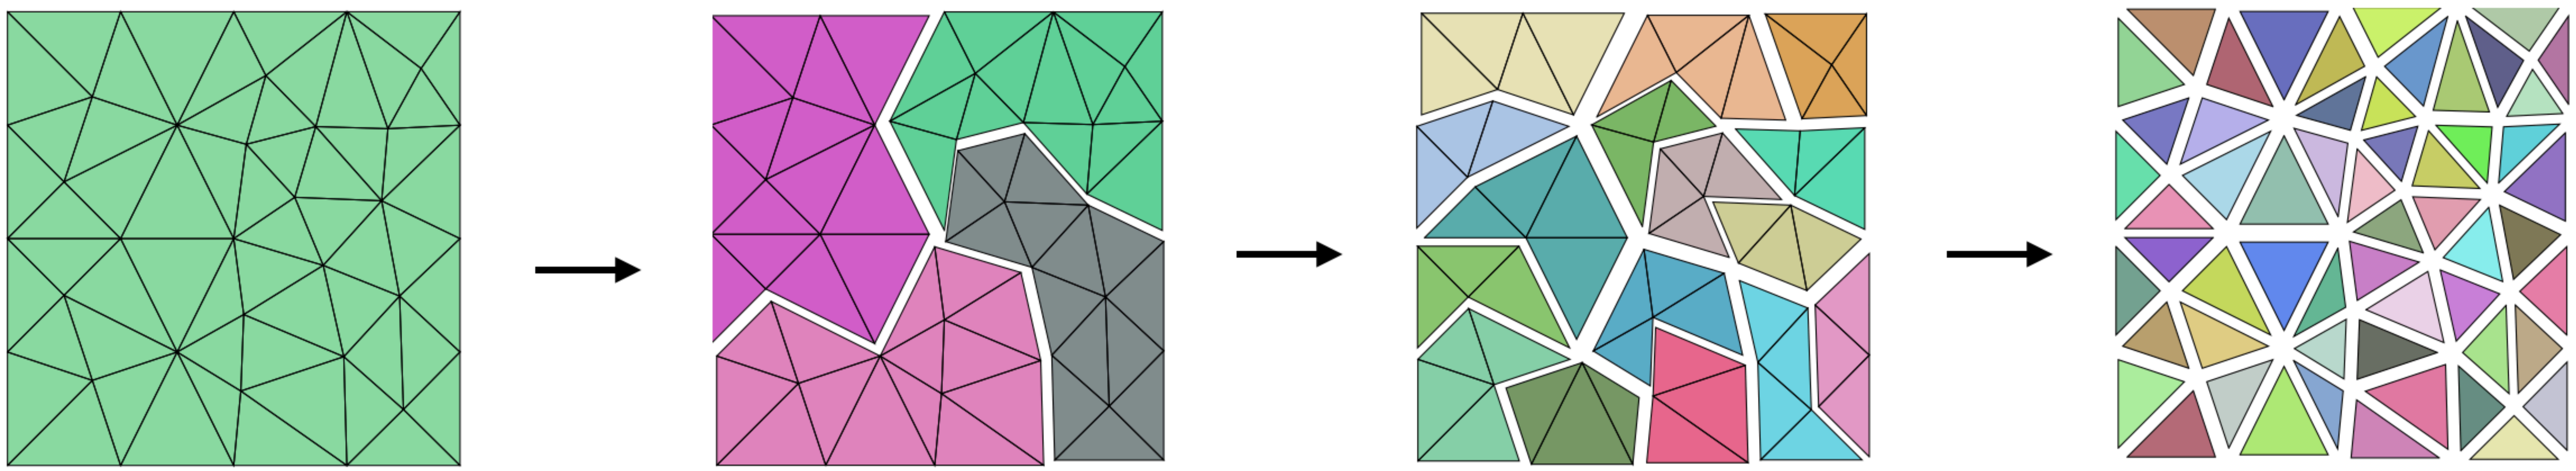

h- vs h*- multigrid

- Two types of geometric coarsening possible in DG
    - Inter-element h-multigrid, by combining elements via agglomeration
    - Intra-element h*-multigrid, by coarsening modes within each element (akin to standard p-multigrid)

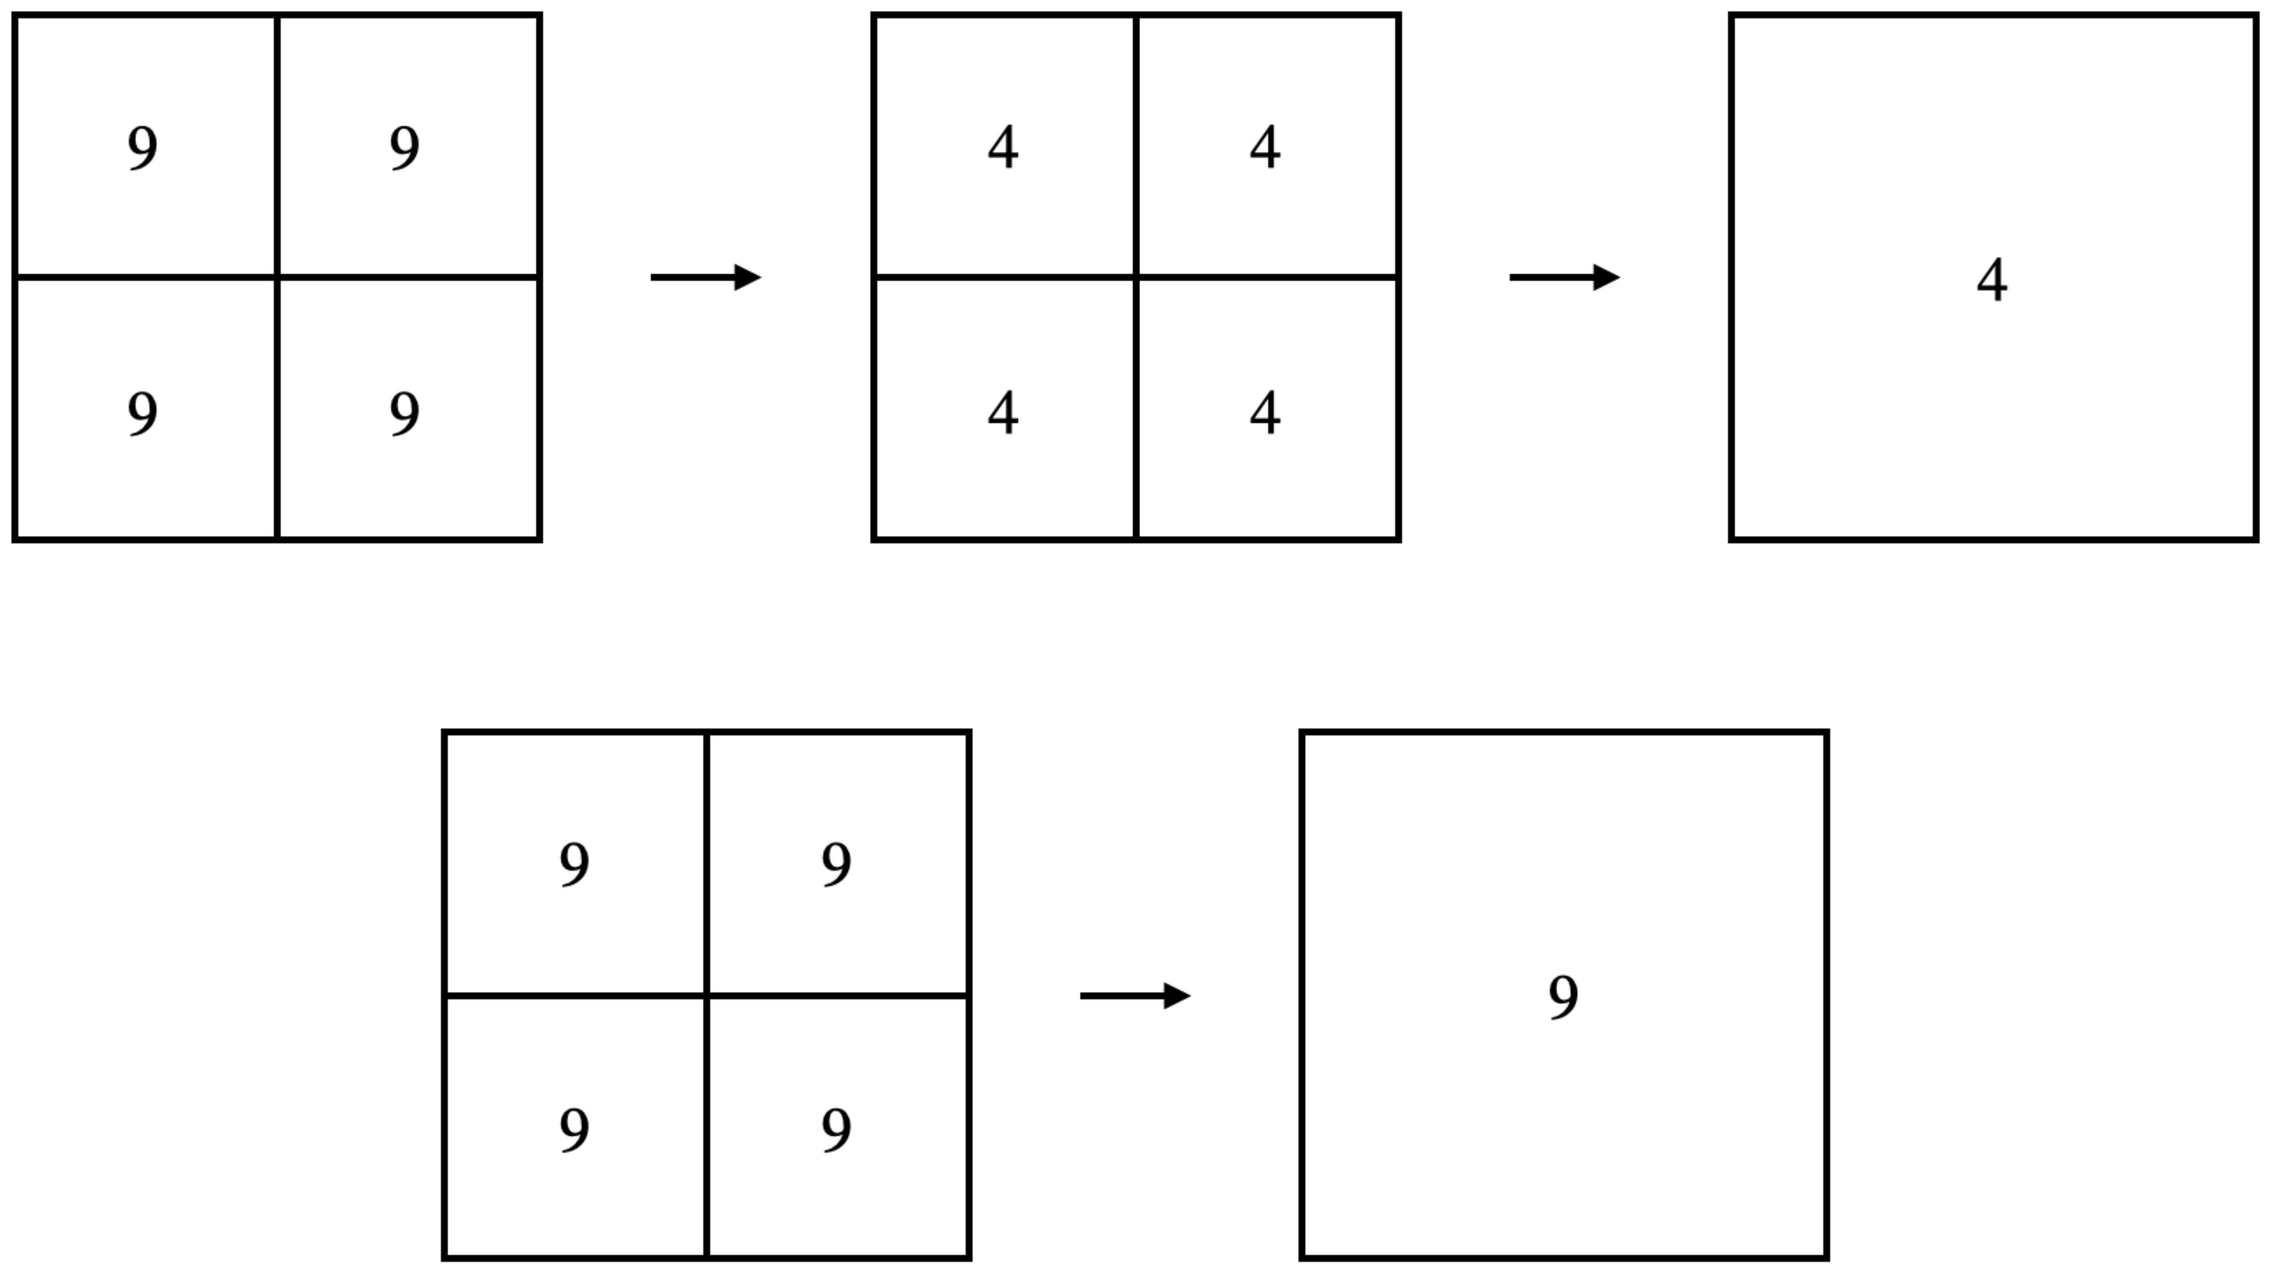

Geometric αSA for DG

- Overall algorithm (assume given matrix A^k and mesh), at each level k:

1. Form mesh partitions for next level via agglomeration
2. Form smoother S_k, and apply to random Gaussian vectors b_j^k
3. Partition random vectors and form prolongation T_{k+1}^k using αSA procedure
4. Form restriction operator $R_k^{k+1} = \left(T_{k+1}^k\right)^T$
5. Form operator at next level $A^{k+1} = R_k^{k+1} A^k T_{k+1}^k$
6. Go to next level $k + 1$

Numerical examples

Poisson's equation (various numerical fluxes)

- Test problem: $-\Delta u = f$ in 3D, $\Omega = [-1,1]^3$, $p = 1$, mesh refinement

nRef	IP	LDG	IP(pCG)	LDG(pCG)
1	7	7	6	6
2	7	8	6	7
3	7	9	6	7

#Iterations

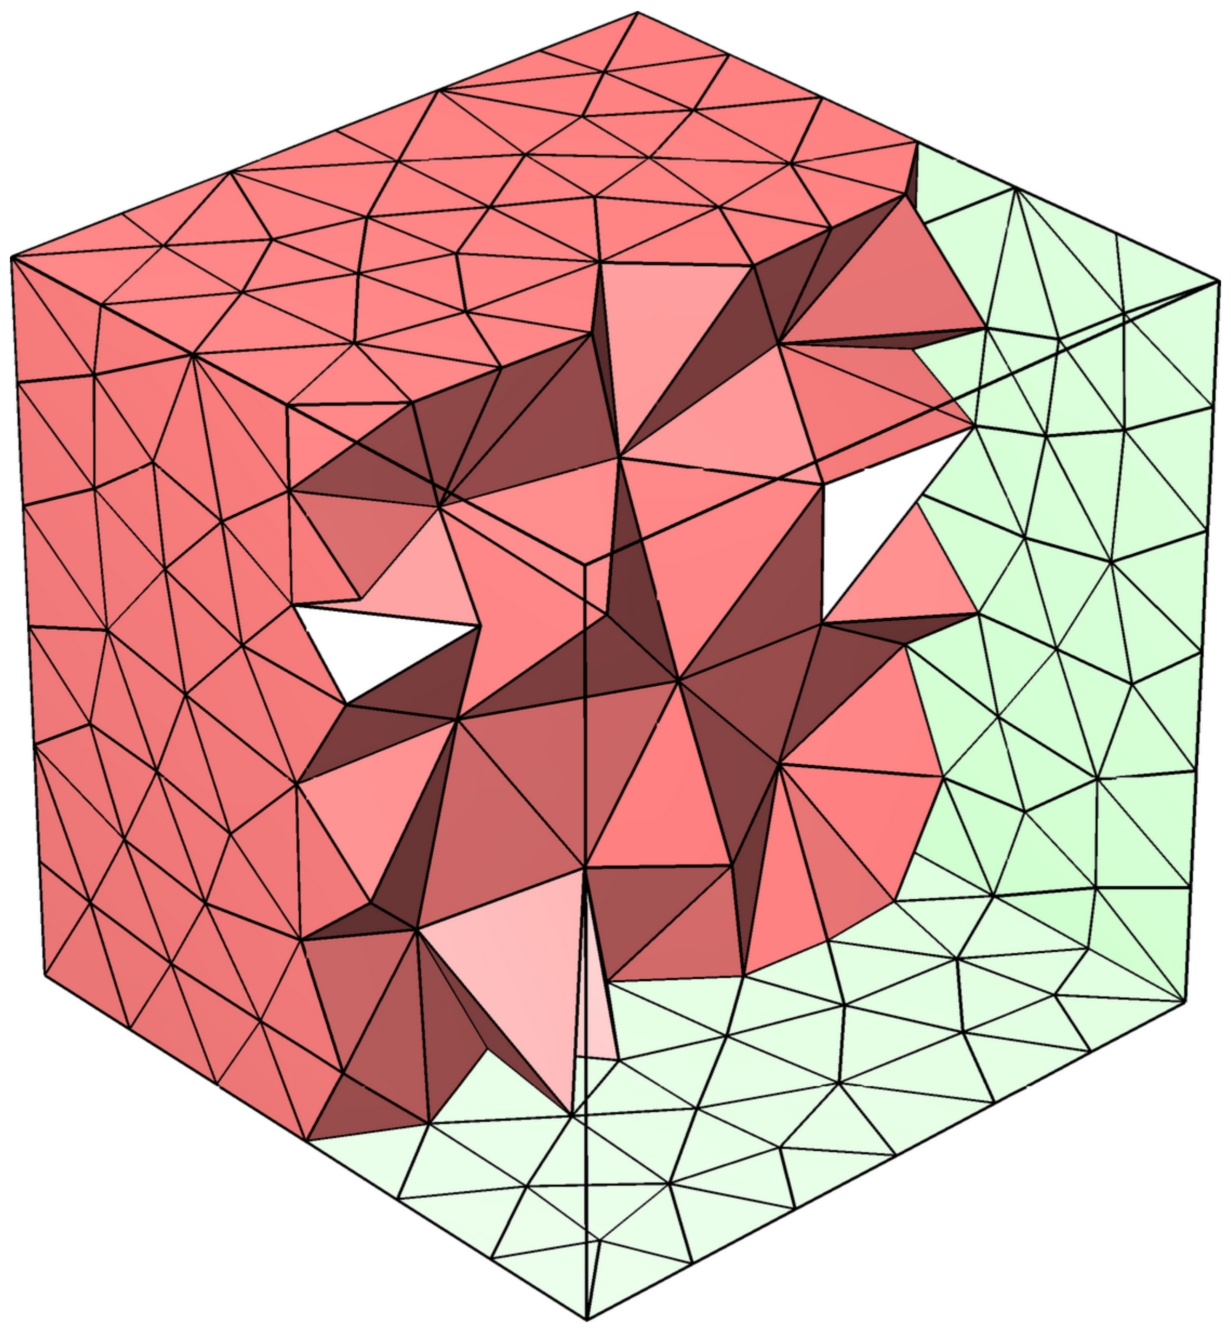

k	dof	nnz
0	32,728	1,835,008
1	5,206	1,706,338
2	882	310,538
3	100	9,422

MG hierarchy at nRef=2

Poisson's equation (h- vs h*- multigrid)

- Test problem: $-\Delta u = f$ in 3D, $\Omega = [-1,1]^3$, fixed mesh, variable p

p	h*-	h-	h*-(pCG)	h-(pCG)
1	8	8	7	6
2	8	9	6	7
3	10	9	7	7

#Iterations

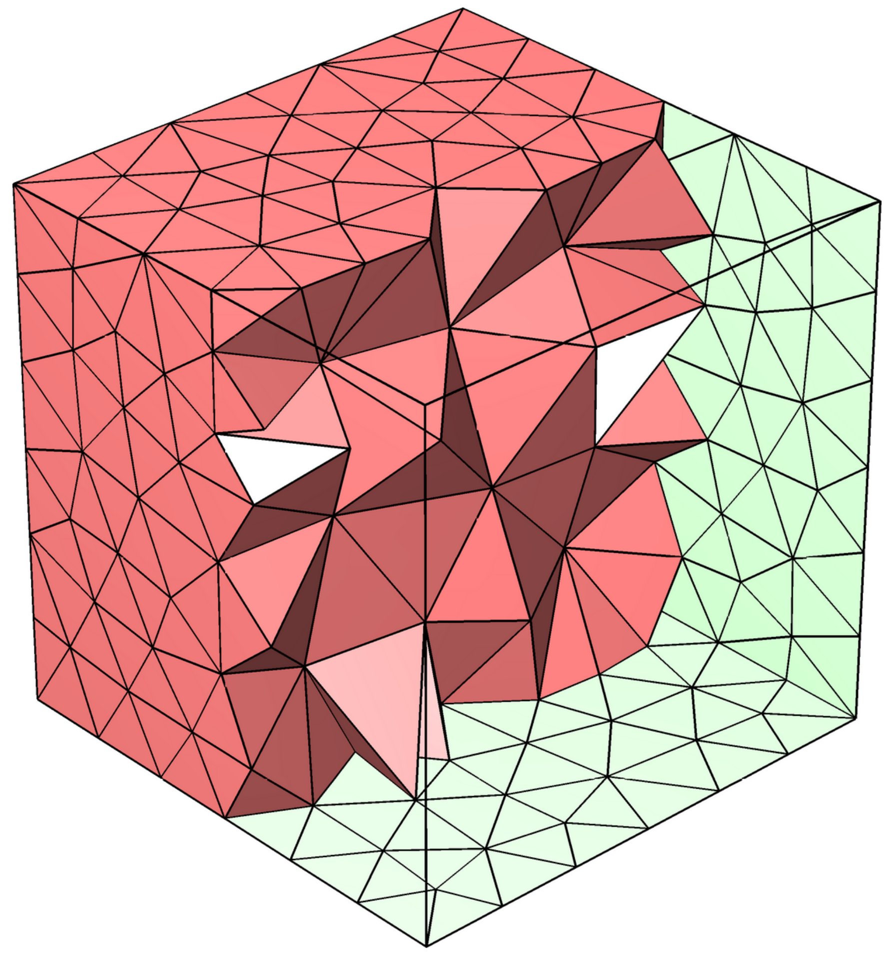

k	dof	nnz
0	110,592	20,901,888
1	56,029	5,411,417
2	9,574	5,387,940
3	3,524	640,246

h*-MG hierarchy at p=2

k	dof	nnz
0	110,592	20,901,888
1	18,500	16,494,040
2	3,104	2,342,854

h-MG hierarchy at p=2

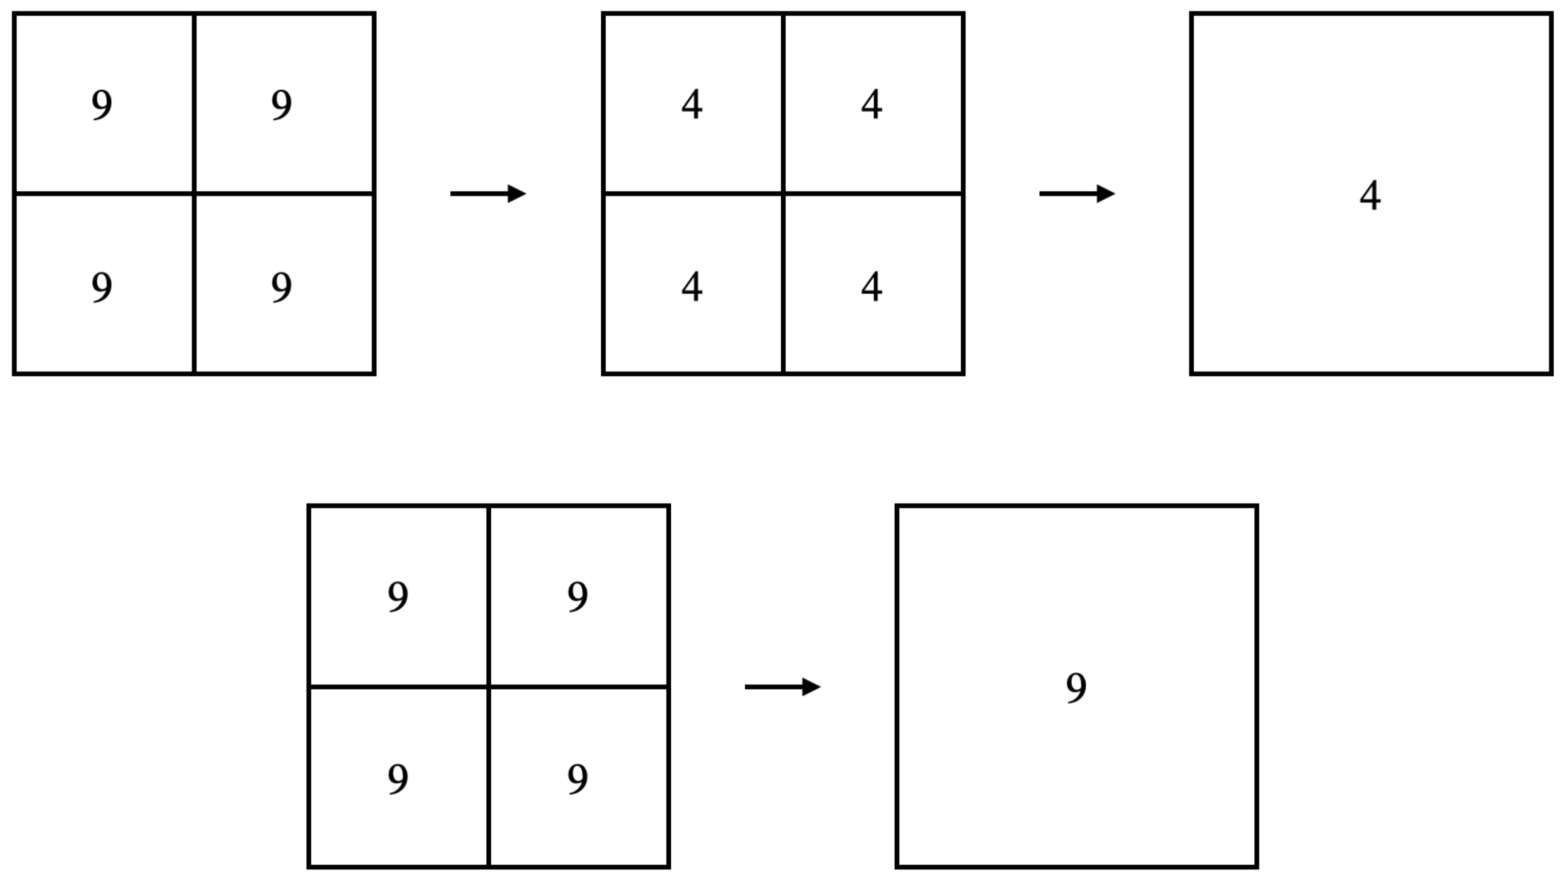

Convection-diffusion

- Test problem: $\mathbf{v} \cdot \nabla u - \mu \Delta u = f$ in 3D, $\Omega = [-1,1]^3$, $p = 1$, mesh refinement

#Iterations - Pe = 0

nRef	nDof	Our method	block Jacobi	AMG (classical)	AMG (SA)
0	**4,188**	9	184	38	36
1	**29,180**	10	430	64	58
2	**214,260**	11	1462	121	91
3	**1,660,556**	13	3838	309	185

#Iterations - Pe = 1000

nRef	nDof	Our method	block Jacobi	AMG (classical)	AMG (SA)
0	**4,188**	10	46	-	-
1	**29,180**	12	83	-	2473
2	**214,260**	12	279	719	507
3	**1,660,556**	17	589	367	325

#Iterations - Pe = 100

nRef	nDof	Our method	block Jacobi	AMG (classical)	AMG (SA)
0	**4,188**	9	76	61	66
1	**29,180**	10	181	57	74
2	**214,260**	12	451	75	100
3	**1,660,556**	16	892	173	172

#Iterations - Pe = ∞

nRef	nDof	Our method	block Jacobi	AMG (classical)	AMG (SA)
0	**4,188**	9	37	-	-
1	**29,180**	13	63	-	-
2	**214,260**	17	118	-	-
3	**1,660,556**	25	416	-	-

Comparison of pGMRES iterations

Conclusion

Conclusion

- Novel geometric multigrid method suitable for DG discretisations
- Applicable on unstructured meshes
- Can be applied uniformly for a variety of numerical fluxes without each requiring different treatments
- Combines aspects of GMG/AMG solvers to reduce reliance on explicit structure of matrix operators
- Excellent solver performance observed in practice for convection-diffusion problems

Thank you!

Yulong Pan, Michael Lindsey and Per-Olof Persson. Geometric adaptive smoothed aggregation multigrid for discontinuous Galerkin discretisations. Submitted to Journal of Computational Physics. https://arxiv.org/abs/2504.13373