SAND2022-3781C

Smoothed Aggregation for Difficult Stretched Mesh and Coefficient Variation Problems

Ray Tuminaro
Center for Computing Research
Sandia National Laboratories

J. Hu @ Sandia, C. Siefert@ Sandia

Apr., 2022

4 smoothed aggregation adaptations

Talk centers on

• Alternative definition of matrix diagonal within prolongator smoother

• Alternative definition of filtered matrix within prolongator smoother

• Modification/post-process of interpolation operator

• ~~Modification/post-process of strength of connection algorithm~~

+

Numerical Results

Smoothed Aggregation (SA-AMG) review

Solve $A_1 u_1 = f_1$ via

Vcycle(A_l, f_l, u_l, l)

if $l \neq N_{level}$ {

$u_l \quad = \hat{S}_l(A_l, f_l, u_l)$

$f_{l+1} = P_l^T (f_l - A_l u_l)$

Vcycle$(A_{l+1}, f_{l+1}, 0, l+1)$

$u_l \quad = u_l + P_l u_{l+1}$

}

else $u_l = A_l^{-1} f_l$

$A_{l+1} = P_l^T A_l P_l$

$\hat{S}_l(\ )$ is relaxation or smoother

$$P_l = (I - \omega \bar{D}_l^{-1} \bar{A}_l) P_l^{(t)}$$

$\bar{A}_l$ is *filtered* A_l, D_l is diag $\overline{(A_l)}$

$P_l^{(t)}$ is piecewise constant tentative prolongator. Each column corresponds to 1 aggregate

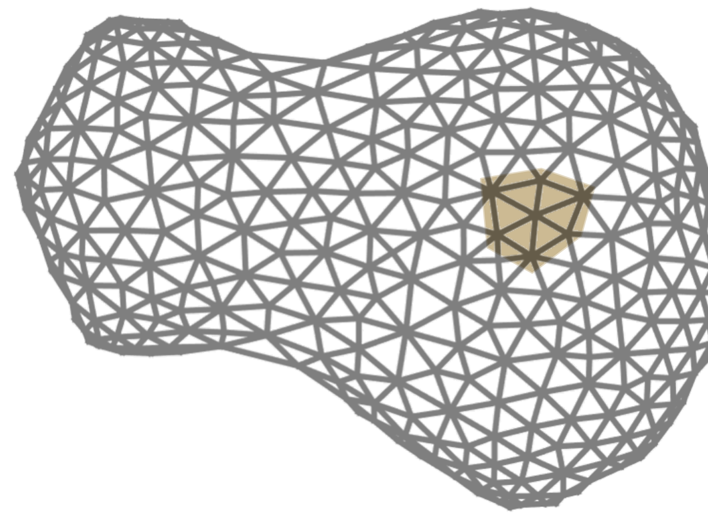

SA-AMG Details

$$P_l = (I - \omega \bar{D}_l^{-1} \bar{A}_l) P_l^{(t)}$$

$\bar{A}_l$ is *filtered* A_l, $\bar{D}_l$ is diag($\bar{A}_l$)

$$(\bar{A}_\ell)_{ij} = \begin{cases} (A_\ell)_{ij} & \text{if } \; i \neq j \text{ and } j \in \mathcal{S}_{\ell,i} \\ (A_\ell)_{ii} + \sum_{k \in \mathcal{W}_{\ell,i}} (A_\ell)_{ik} & \text{if } \; i = j \\ 0 & \text{otherwise} \end{cases}$$

Weak Connections

Strong Connections

Strong connections *usually* satisfy $|(A_\ell)_{ij}| \geq \theta \sqrt{(A_\ell)_{ii} (A_\ell)_{jj}}$

threshold

$$\omega = \frac{4}{3\, \rho(\bar{D}_l^{-1} \bar{A}_l)}$$

spectral radius

$P_l^{(t)}$ is piecewise constant tentative prolongator

Sandia National Laboratories

Four algorithm adaptations

- Alternative definition of $\overline{D}_l$
- Alternative definition of $\bar{A}_l$ using standard S_l
- Modification/post-process of P_l
- ~~Modification/post-process of S_l~~

Sometimes these modifications have no positive or negative impact, but …

- ❖ sometimes quite helpful
- ❖ almost never detrimental*

Sandia National Laboratories

$\widetilde{D}_l$: an alternative to $\overline{D}_l$

Motivation

$$P_l = (I - \omega \overline{D}_l^{-1} \bar{A}_l) P_l^{(t)} \qquad \omega = \frac{4}{3\,\rho(\overline{D}_l^{-1} \bar{A}_l)}$$

Notice

- small $\omega \Longrightarrow \quad P_l \approx P_l^{(t)}$, which is suboptimal
- $\rho(\overline{D}_l^{-1} \bar{A}_l)$ may be *large* if just one entry of $\overline{D}_l$ is relatively *small*

- dropping lowers diagonal entry

$$\left(\bar{A}_\ell\right)_{ii} = \quad (A_\ell)_{ii} + \sum_{k \in \mathcal{W}_{\ell,i}} (A_\ell)_{ik}$$

when dropped $(A_\ell)_{ik} < 0$

$(\bar{A}_l)_{ii} < 0$ possible, e.g. 3D linear FE on cube with $1 \times 1 \times h_z$ elements for $h_z > \sqrt{7}/2$, dropping all but largest off-diagonal

$\widetilde{D}_l$: 1-norm approximation

Define

$$\gamma_{i,l} = \sum_j \left|(\bar{A}_l)_{ij}\right|, \quad s_{i,l} = \sum_j (\bar{A}_l)_{ij}$$

then

$$(\widetilde{D}_l)_{ii} = \begin{cases} \gamma_{i,l} & if \quad \gamma_{i,l} \geq 2\, s_{i,l} \\ 1 & if \quad \gamma_{i,l} \equiv 0 \\ 2\, s_{i,l} & otherwise \end{cases}$$

safeguards

Middle case corresponds to identically 0 row.

Last case enforces that

$$(P_l)_{ij} \geq {}^{1}\!/_{3}$$

when i is the center of an ideal j th aggregate, which would correspond to the aggregate's basis function peak.

Note: $\rho(\widetilde{D}_l^{-1}\bar{A}_l) \leq 1$

$\tilde{A}$: an alternative to $\bar{A}$ for diffusion operators*

Recall

$$\bar{A}_{ij} = \begin{cases} A_{ij} & \text{if } \; i \neq j \text{ and } j \in \mathcal{S}_i \\ A_{ii} + \sum_{k \in \mathcal{W}_i} A_{ik} & \text{if } \; i = j \\ 0 & \text{otherwise} \end{cases}$$

which might lead to negative or small diagonal entries … or more generally

$$\mathcal{D}(\bar{A}_i) \gg \mathcal{D}(A_i) \quad \text{where} \quad \mathcal{D}(A_i) = \frac{\sum_{j \neq i} |A_{ij}|}{|A_{ii}|}$$

Consider instead lumping that maintains row sums by modifying either

- retained positive off-diagonal entries
- diagonal entry
- retained negative entries

} *in this order of preference*

to enforce $\mathcal{D}(\bar{A}_i) \leq \tau \mathcal{D}(A_i)$ ← *user tolerance* (for τ)

* Dropping l subscript for remainder of presentation

Detailed lumping algorithm

$\bar{A}_i$ =Lump_AvoidSmallDiag($A_i, \mathcal{R}_i, \tau$)

Input:

- A_i — i^{th} row of matrix with entries to be dropped
- $\mathcal{R}_i$ — set of column indices in i^{th} row to be removed
- τ — tolerance indicating that $\mathcal{D}(\bar{A}_i)$ should not exceed $\tau\mathcal{D}(A_i)$

Output:

- $\bar{A}_i$ — matrix row where $\bar{A}_{ij} = 0$ for $j \in \mathcal{R}_i$ and $\bar{A}_i v = Av$ where v is a constant vector

---

Let $r_i \leftarrow \sum_{k\in\mathcal{R}_i} A_{ik}$
if $r_i > 0$ **then** $\bar{A}_{ii} \leftarrow A_{ii} + r_i$ // decreases $\mathcal{D}(\bar{A}_i)$
else {
 Let $\mathcal{K}_i^+ \leftarrow \{k \mid A_{ik} > 0 \wedge k \neq i \wedge k \notin \mathcal{R}_i\}$
 Let $\mathcal{K}_i^- \leftarrow \{k \mid A_{ik} < 0 \wedge k \neq i \wedge k \notin \mathcal{R}_i\}$
 Let $\kappa_i^+ \leftarrow \sum_{k\in\mathcal{K}_i^+} A_{ik}$; $\kappa_i^- \leftarrow \sum_{k\in\mathcal{K}_i^-} A_{ik}$;
 if $|r_i| \leq \kappa_i^+$ **then** $\bar{A}_{ij} \leftarrow A_{ij}(1+\delta_i)$ for $j \in \mathcal{K}_i^+$ where $\delta_i \leftarrow r_i/\kappa^+$
 else {
  $\bar{A}_{ij} \leftarrow 0$ for $j \in \mathcal{K}_i^+$ // zero out the $\mathcal{K}_i^+$ by distributing a
  $\hat{r}_i \leftarrow r_i + \kappa_i^+$ // portion of r_i $(= \kappa_i^+)$ to them
  if $\mathcal{K}_i^- == \emptyset$ **then** redistribute to $\mathcal{K}^+$ if possible or if not
    possible do not modify row i and return
  else {
   find largest positive $r_i^* < \min(d_{ii}, |\hat{r}_i|)$ such that $\mathcal{D}(\bar{A}_i) \leq \tau\mathcal{D}(A_i)$
   define $\bar{A}_i$ such that its only nonzero values are
    $\bar{A}_{ii} \leftarrow A_{ii} - r_i^*$
    $\bar{A}_{ij} \leftarrow A_{ij}(1+\delta_i)$ for $i \in \mathcal{K}_i^-$
   where $\delta_i \leftarrow (\hat{r}_i + r_i^*)/\kappa^-$
  }
 }
}

Sandia National Laboratories

$\tilde{P}$: post-process P

Force $0 \le \tilde{P}_{ij} \le 1$

via

$$\begin{cases} \widetilde{P}_i = \underset{\hat{P}_i}{\arg\min} \; ||\hat{P}_i - P_i||_2 \\ \quad\quad \text{subject to} \quad \hat{P}_{ij} = 0 \text{ if } P_{ij} = 0, \quad \hat{P}_{ij} \ge 0, \quad \hat{P}_{ij} \le 1, \quad \hat{P}_i v = P_i v \end{cases}$$

where $\tilde{P}_i$ is i^{th} row of $\tilde{P}$ & v is vector of all ones.

Solved by 2-step process when feasible solution exists

- shift all nonzero $\tilde{P}_{ij}$ in i^{th} row by δ_i
- enforce all bounds

Must determine δ_i to maintain $\tilde{P}_i v = P_i v$ when bounds enforced.

Done by assuming a minimal set of bound-violating entries, determining δ_i for this set. Add one entry to set if shift causes more bound-violating entries, repeat δ_i and so on.

$\widetilde{P}$: post-process P

Algorithm 3 Constrain_One_P_Row($\widetilde{P}_i$, lowerBound, upperBound)

---

optimal $\leftarrow$ false
$k \leftarrow 0$
$\ell_k \leftarrow \arg\min_{\ell_k} \ \widetilde{P}_i[\ell_k] > \mathsf{lowerBound}$
$u_k \leftarrow \arg\min_{u_k} \ \widetilde{P}_i[u_k] > \mathsf{upperBound}$
while (optimal == false) {

$\quad \delta_k \leftarrow \left(\sum_{j=\ell_k}^{u_k - 1} \widetilde{P}_i[j] - \sum_{j=1}^{nnz(\widetilde{P}_i)} \widetilde{P}_i[j] \right) / \left(u_k - \ell_k \right)$

$\quad$ // does $\widetilde{P}_i[\ell_k] - \delta_k$ still satisfy lowerBound & does $\widetilde{P}_i[u_k] - \delta_k$ still violate upperBound

$\quad$ if $\delta_k < min(\widetilde{P}_i[\ell_k] - \mathsf{lowerBound}, \ \widetilde{P}_i[u_k] - \mathsf{upperBound}))$, optimal $\leftarrow$ true
$\quad$ else {
$\quad\quad$ if $\widetilde{P}_i[\ell_k] - \mathsf{lowerBound} \leq \widetilde{P}_i[u_k] - \mathsf{upperBound}$, $\ell_k \leftarrow \ell_k + 1$
$\quad\quad$ else $u_k \leftarrow u_k + 1$
$\quad$ }
$\quad k \leftarrow k + 1$
}
for $j = 1 : \ell_k - 1$, $\quad \widetilde{P}_i[j] \leftarrow$ lowerBound
for $j = \ell_k : u_k - 1$, $\quad \widetilde{P}_i[j] \leftarrow \widetilde{P}_i[j] - \delta_k$
for $j = u_k : nnz(\widetilde{P}_i)$, $\quad \widetilde{P}_i[j] \leftarrow$ upperBound

Results: Refined Boundary Mesh

Poisson

BCs

u(top) = 0

Neumann elsewhere

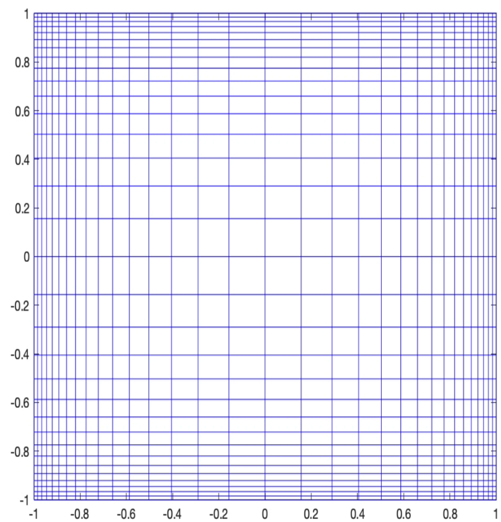

Algorithm Choice			$\theta = .05$ AMG complex.	its.	AMG Setup	Solve
			1.33	19	.21	.16
		$\tilde{P}_\ell$	1.33	20	.23	.18
	$\tilde{A}_\ell$		1.33	15	.27	.14
	$\tilde{A}_\ell$	$\tilde{P}_\ell$	1.33	14	.28	.12
$\tilde{D}_\ell$			1.33	14	.22	.12
$\tilde{D}_\ell$		$\tilde{P}_\ell$	1.33	14	.26	.12
$\tilde{D}_\ell$	$\tilde{A}_\ell$		1.33	14	.26	.12
$\tilde{D}_\ell$	$\tilde{A}_\ell$	$\tilde{P}_\ell$	1.33	14	.29	.12

2D Linear Elasticity

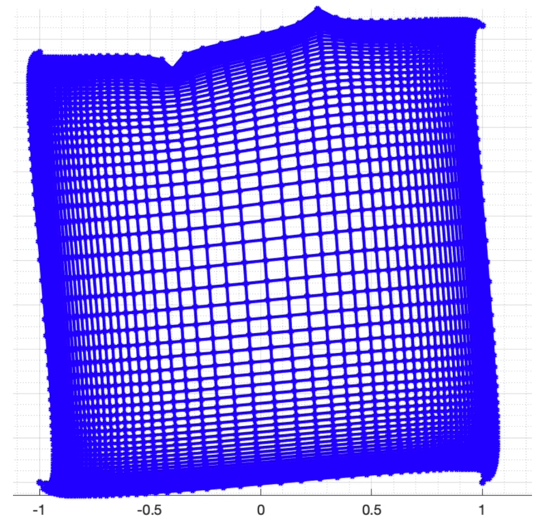

BCs

u(top corners) Dirichlet x displacements

u(bottom corners) Dirichlet x, y displacements

Neumann elsewhere

Algorithm Choice		$\theta = .05$ AMG complex.	its.	AMG Setup	Solve
		1.31	60	.84	2.18
	$\tilde{P}_\ell$	1.31	50	.89	1.66
$\tilde{D}_\ell$		1.31	59	.87	2.11
$\tilde{D}_\ell$	$\tilde{P}_\ell$	1.31	48	.90	1.56

linear FE on 256 x 256 quad mesh, CG residual reduced by 10^{-8}, $\tau = 1.1$ for $\tilde{A}_l$, 4 MG levels

distance Laplacian used for strength, p=2 Cheby relaxation, rotations not used for elasticity

SPE10 oil/gas subsurface flow

Darcy Flow

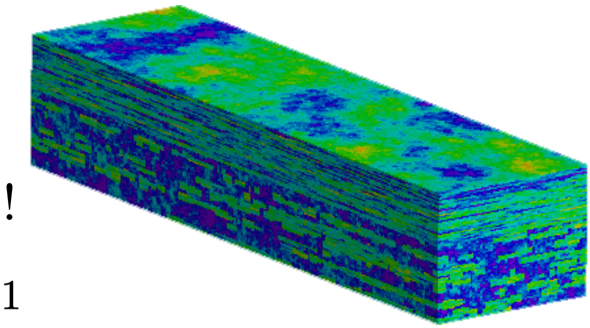

porosity range of 10^{12}

All results use $\widetilde{D}_l$ …. SA-AMG did not converge without $\widetilde{D}_l$!

Algorithm Choice	$\theta = .02$ AMG complex.	$\theta = .02$ its.	$\theta = .05$ AMG complex.	$\theta = .05$ its.	$\theta = .1$ AMG complex.	$\theta = .1$ its.
$\widetilde{D}_l$	1.85	24	2.00	27	2.60	31
$\widetilde{D}_l$ $\tilde{P}_\ell$	1.85	26	2.00	28	2.60	31
$\tilde{A}_\ell$ $\widetilde{D}_l$	1.85	24	2.00	19	2.60	17
$\tilde{A}_\ell$ $\widetilde{D}_l$ $\tilde{P}_\ell$	1.85	22	2.00	19	2.60	16

Algorithm Choice	$\theta = .02$ AMG Setup	$\theta = .02$ Solve	$\theta = .05$ AMG Setup	$\theta = .05$ Solve	$\theta = .1$ AMG Setup	$\theta = .1$ Solve
$\widetilde{D}_l$	2.49	1.86	2.64	2.25	3.15	3.62
$\widetilde{D}_l$ $\tilde{P}_\ell$	2.74	2.31	2.87	2.31	3.46	3.67
$\tilde{A}_\ell$ $\widetilde{D}_l$	2.93	1.86	3.13	1.63	3.69	2.27
$\tilde{A}_\ell$ $\widetilde{D}_l$ $\tilde{P}_\ell$	3.18	1.69	3.31	1.60	3.99	2.13

finest mesh 660K dofs, CG residual reduced by 10^{-8} , $\tau = 1.1$ for $\tilde{A}_l$, 5 MG levels, standard strength measure, p=2 Cheby relaxation

Sandia National Laboratories

Wind Turbine

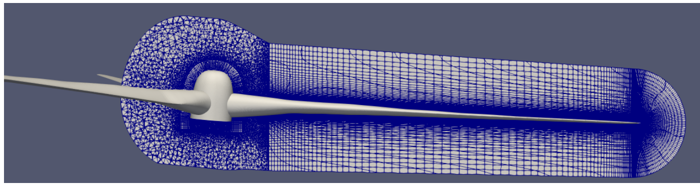

uncoupled formulation

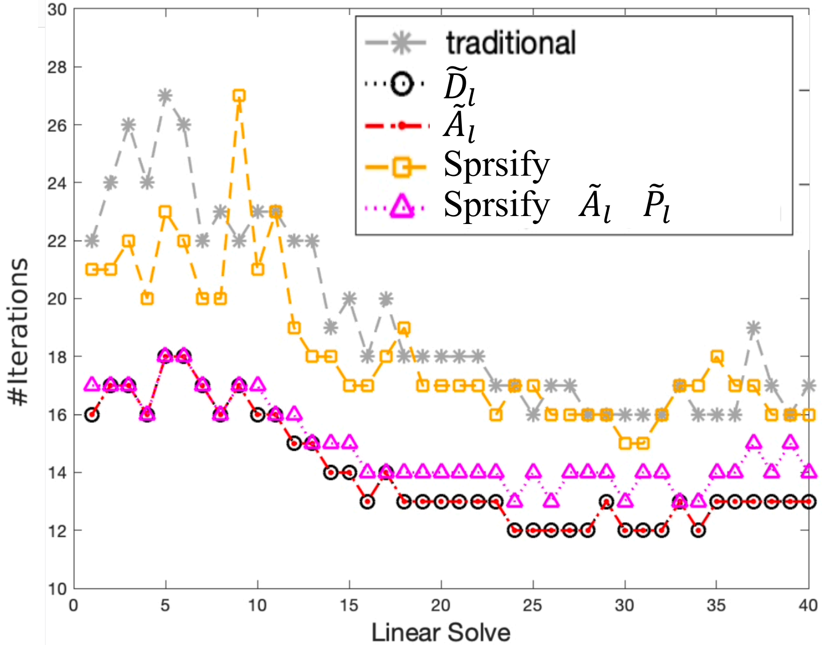

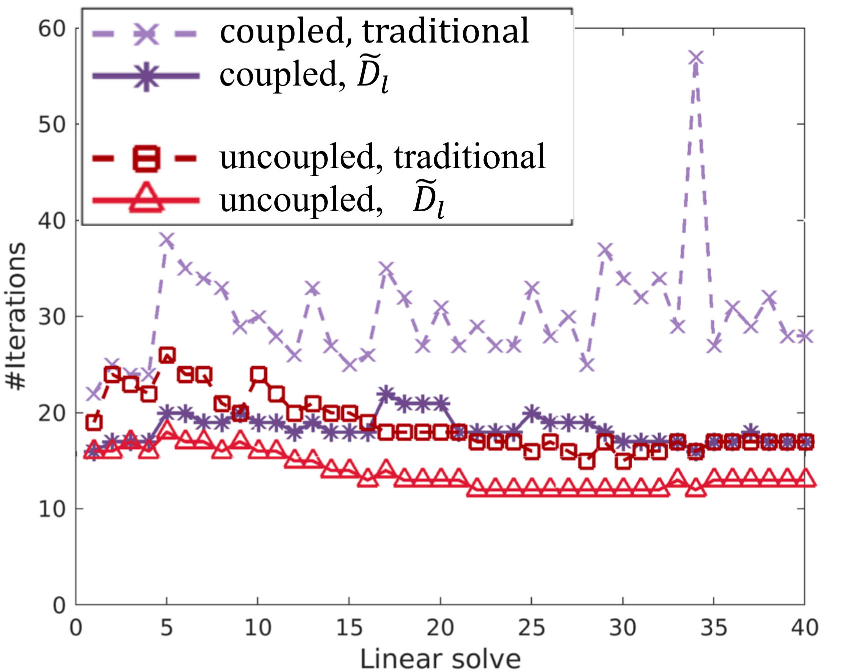

uncoupled formulation

Algorithm Choice	AMG Setup	Solve	AMG Complexity
traditional	3.15	7.78	1.66
$\tilde{D}_\ell$	3.52	4.41	1.65
$\tilde{A}_\ell$ Sprsfy $\tilde{P}_\ell$	5.47	3.16	1.61

finest mesh 23M dofs, GMRES residual reduced by 10^{-5} , $\tau = 1.1$ for $\tilde{A}_l$, 5 MG levels, standard strength measure, p=2 Cheby relaxation, $\theta = .02$

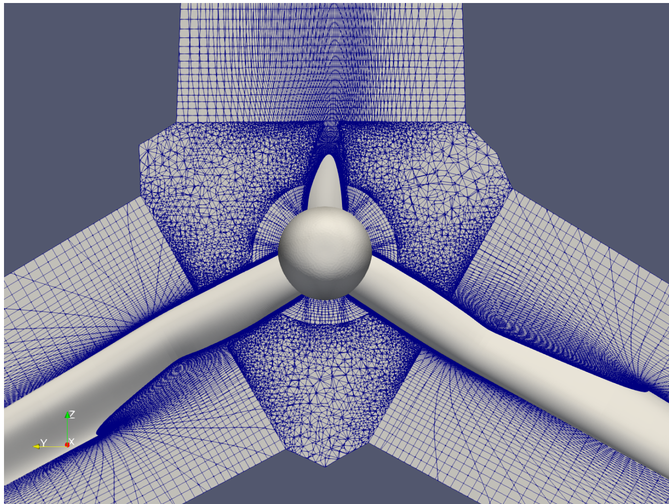

Wind Turbine Robustness Study

coupled formulation

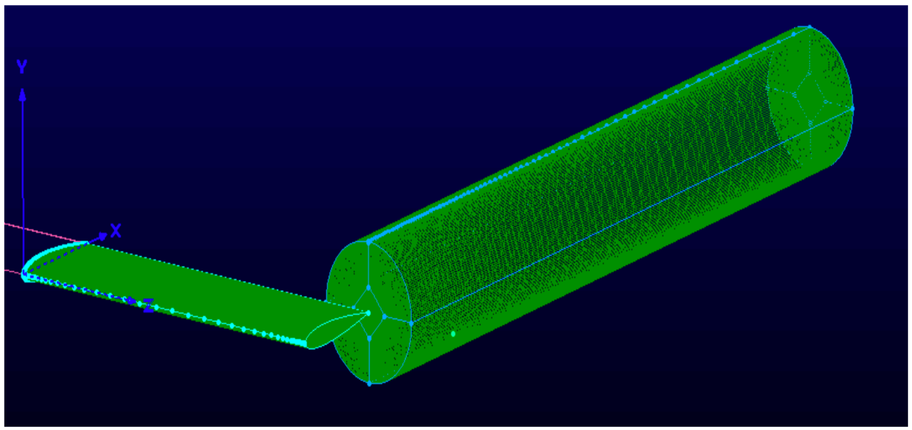

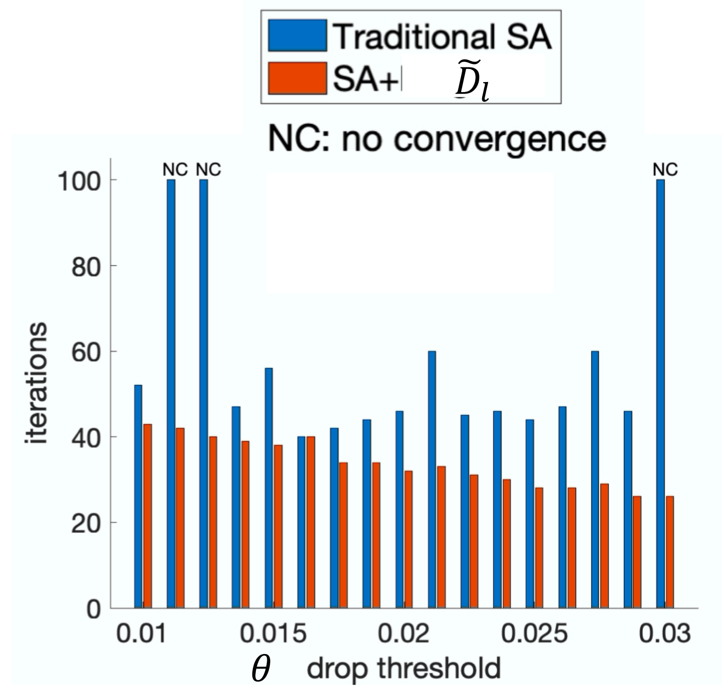

$\widetilde{D}_l$ needed to avoid erratic failures that occur for traditional SA !

Sandia National Laboratories

Concluding Remarks

- 3 SA-AMG adaptations presented
  - Alternative definition of $\overline{D}_l$
  - Alternative definition of $\bar{A}_l$ using standard $\mathcal{S}_l$
  - Modification/post-process of P_l

- Sometimes help improve robustness or convergence of SA-AMG, rarely hurt
  - Poisson & linear elasticity on stretched mesh
  - SPE10 benchmark
  - wind calculation

Support provided by the US Dept. of Energy, Office of Science,
Advanced Scientific Computing Research (ASCR)
Applied Mathematics program

Sandia National Laboratories