Sandia National Laboratories

SAND2021-5988C

Multilevel Methods for Maxwell's Equations

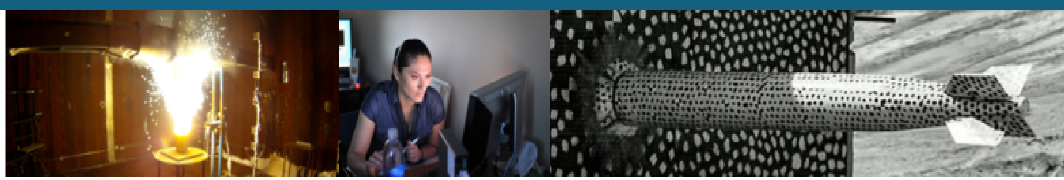

Chris Siefert, Ray Tuminaro, Christian Glusa and John Kaushagen
Sandia National Laboratories

SAND2021-XXXX C

Outline

- Introduction to the Eddy Current Maxwell's Equations
- Motivating Fourier Analysis
- A Closer Look at the Edge Hierarchy
- Results & Conclusions

Eddy Current Maxwell's Equations

- Start with Maxwell's Equations:

$$\nabla \cdot \mathbf{E} = \frac{\rho}{\varepsilon_0} \qquad \nabla \cdot \mathbf{B} = 0$$

$$-\nabla \times \mathbf{E} = \frac{\partial \mathbf{B}}{\partial t} \qquad \mathbf{H} = \mu_0^{-1}\mathbf{B} - \mathbf{M}$$

$$\mathbf{J} = \sigma \mathbf{E} \qquad \mathbf{M} = \left(\mu_0^{-1} - \mu^{-1}\right)\mathbf{B}$$

$$\nabla \times \mathbf{H} = \mu_0 \left(\mathbf{J} + \varepsilon_0 \frac{\partial \mathbf{E}}{\partial t}\right)$$

- Drop (nonlinear) magnetization, polarization, displacement currents, eliminate **B**:

$$\sigma \frac{\partial \mathbf{E}}{\partial t} + \nabla \times \mu^{-1} \nabla \times \mathbf{E} = 0$$

Eddy Current Maxwell's Equations

$$\sigma \frac{\partial \mathbf{E}}{\partial t} + \nabla \times \mu^{-1} \nabla \times \mathbf{E} = 0$$

- The resulting linear operator is singular if σ is zero.
- This has two important repercussions:
  - 3D nodal discretizations introduce unphysical phenomena (e.g. magnetic monopoles) → use edge elements.
  - Stock preconditioners don't scale (and provably so).
- Both of these facts have been known for a long time.

Basiscs of Edge Element Maxwell Multigrid

$E = \nabla\times A - \nabla\phi$

All Maxwell preconditioners have two sub-problems: edges and nodes:

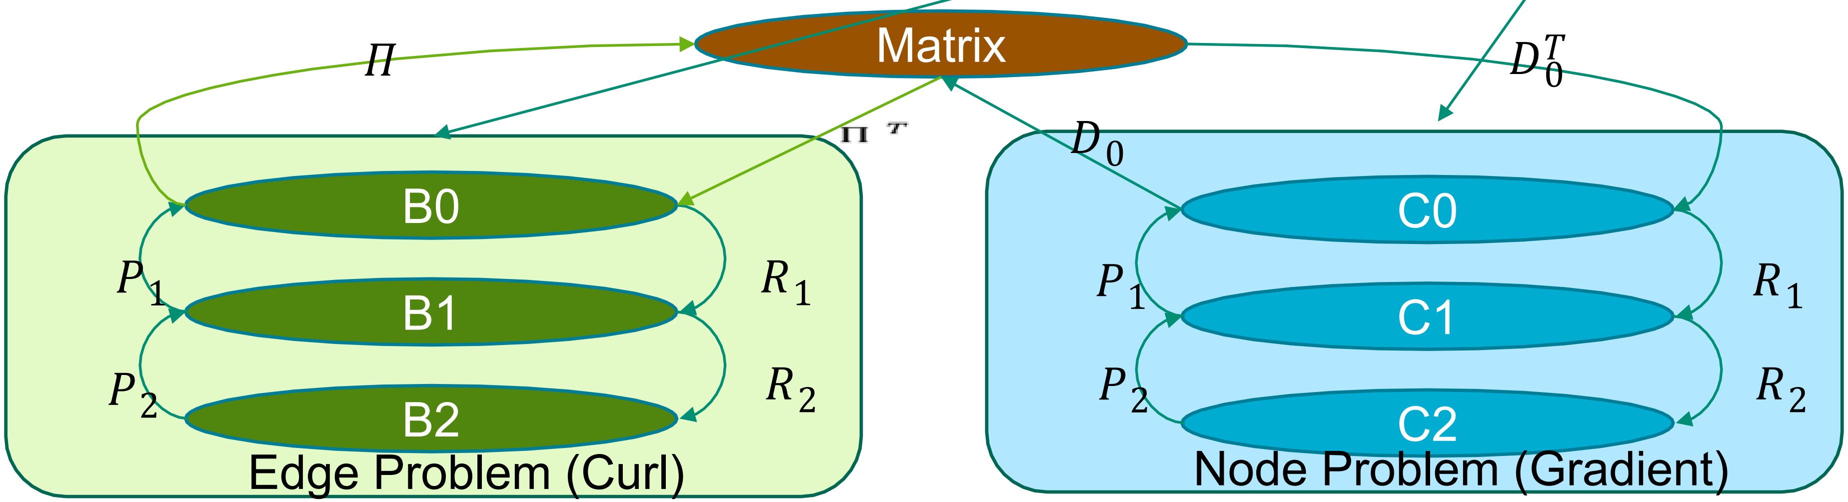

You can mathematically prove you need the node hierarchy (cf. Boonen et al, 2007).

- Most Maxwell preconditioners use the discrete gradient for the node hierarchy.
- The “special sauce” comes in how you handle the edge hierarchy.
- We’ll focus on auxiliary preconditioning style edge hierarchies, e.g. Bochev, et al. 2008 and Kolev, et al. 2008.

Why You Need the Nodal Hierarchy

- Tool: Turn to the Fourier Analysis of Boonen, et al., 2007 (shows damping factor):

Gauss Seidel (Edge Only)

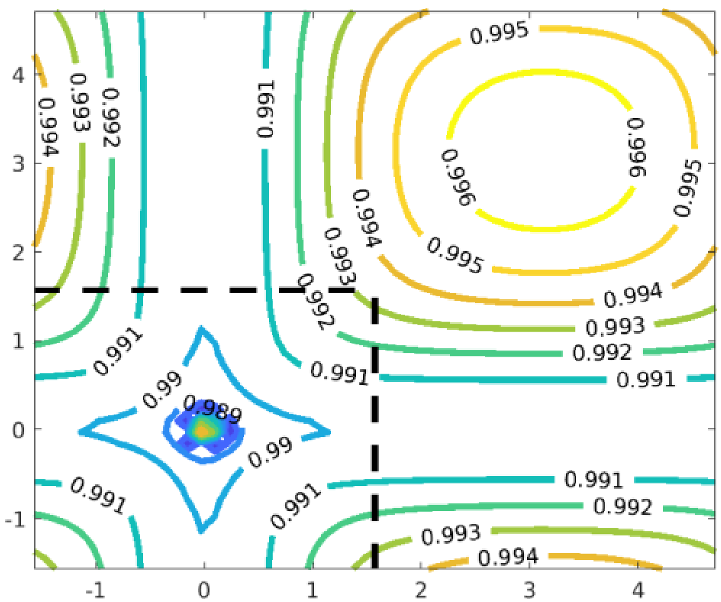

Hiptmair (Edge/Node)

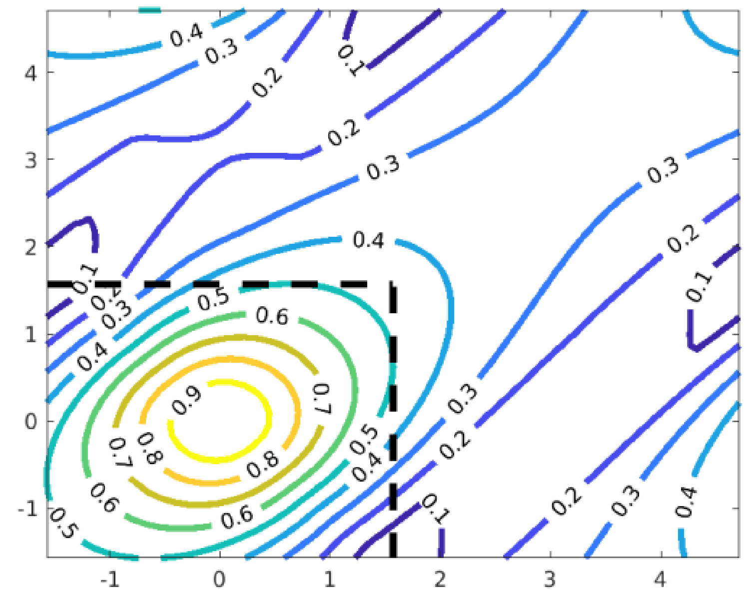

- Lower left corner should be damped by coarse grid; everything else should be damped by the smoother (but GS doesn't!).

Caveat:
2D analysis,
for σ=0.01

Outline

- Introduction to the Eddy Current Maxwell's Equations
- Motivating Fourier Analysis
- A Closer Look at the Edge Hierarchy
- Results & Conclusions

An Analyst's Hypothesis

- Analyst observation: Iteration counts increase logarithmically with CFL.
- Is this behavior to be expected? Theory (e.g. Kolev, et al., 2008) says it *shouldn't* be.
- Fourier analysis, but this time as a function of diffusive CFL: $\frac{\Delta t}{\sigma\mu\,\Delta x^2}$

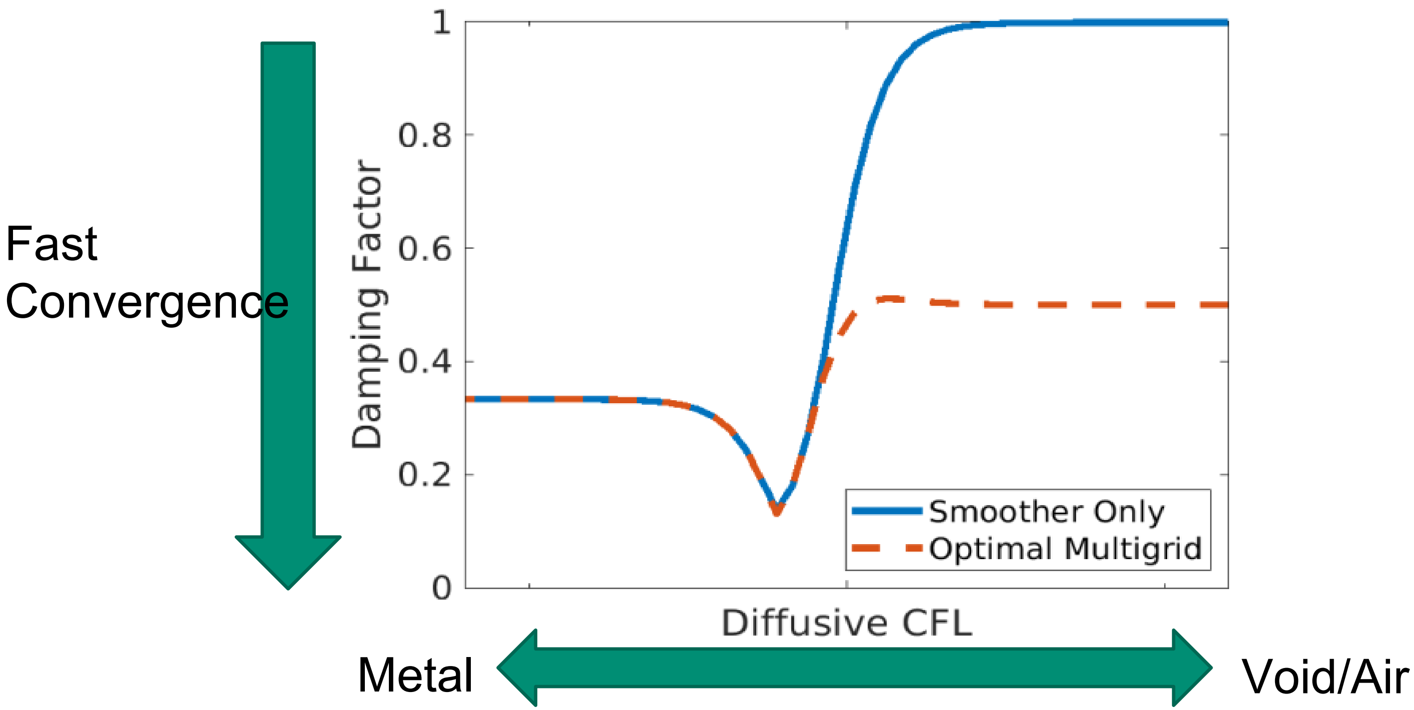

- Takeaway: We should *not* expect to see rising iteration counts, outside of the transition to the "high CFL" regime.

Outline

- Introduction to the Eddy Current Maxwell’s Equations
- Motivating Fourier Analysis
- A Closer Look at the Edge Hierarchy
- Results & Conclusions

Edge Element Stencil

- Curl-Curl Operator:

$$
\begin{bmatrix}
-\frac{\partial^2}{\partial y^2} - \frac{\partial^2}{\partial z^2} & \frac{\partial}{\partial x \partial y} & \frac{\partial}{\partial x \partial y} \\
\frac{\partial}{\partial x \partial y} & -\frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial z^2} & \frac{\partial}{\partial y \partial z} \\
\frac{\partial}{\partial x \partial z} & \frac{\partial}{\partial y \partial z} & -\frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial y^2}
\end{bmatrix}
$$

- Operator is anisotropic! *xx*-component does not contain *x*-derivative!
- We can see this in the stencil (2D example, following Boonen et al., 2007):

$$
\frac{1}{h^2}
\begin{bmatrix}
\circ & -1 & \circ \\
-1 & & 1 \\
\bullet & 2 & \bullet \\
1 & & -1 \\
\circ & -1 & \circ
\end{bmatrix}
$$

Auxiliary Preconditioning Π Operator

- Both Bochev et al. and Kolev et al., use an operator, Π, which moves from a single vector basis on edges to a vector of scalar basis functions on node.

- Π can be written as follows:

$$\Pi = \left[\frac{diag(D_0 x_n)|D_0|}{2}, \frac{diag(D_0 y_n)|D_0|}{2}, \frac{diag(D_0 z_n)|D_0|}{2}\right]$$

Where D_0 is the discrete gradient operator. This gives us entries of O(h).

- Π preserves the stencil anisotropy property when applied. Leading Taylor series terms for $\Pi^T A\, \Pi$ are:

$$h^3 \begin{bmatrix} -\frac{\partial^2}{\partial y^2} - \frac{\partial^2}{\partial z^2} & \frac{\partial}{\partial x \partial y} & \frac{\partial}{\partial x \partial y} \\ \frac{\partial}{\partial x \partial y} & -\frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial z^2} & \frac{\partial}{\partial y \partial z} \\ \frac{\partial}{\partial x \partial z} & \frac{\partial}{\partial y \partial z} & -\frac{\partial^2}{\partial x^2} - \frac{\partial^2}{\partial y^2} \end{bmatrix}$$

- Not surprising: Barring *h*, this is what is supposed to be discretized.

Coarsening the Edge Problem

- By construction of Π, this anisotropy is still there, even if we rotate the mesh.

- Anisotropy is well known to cause issues for multigrid especially when the anisotropy is not mesh aligned. You need to semicoarsen in this case:

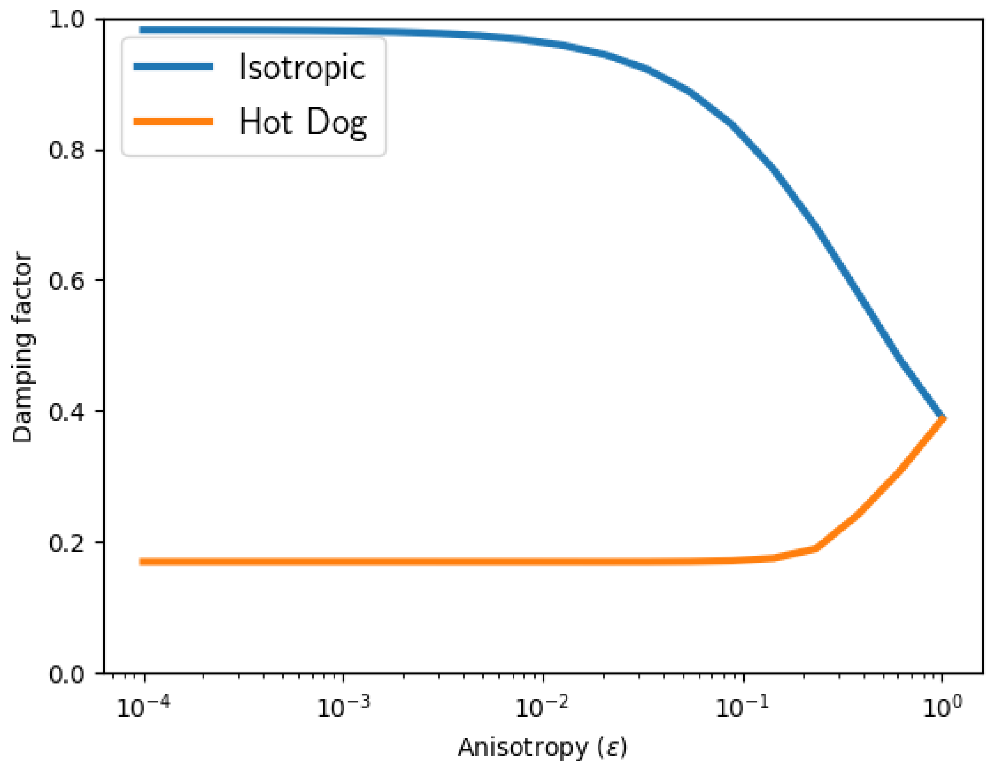

Plot generated using Hannah Rittich's LFA-lab software *github.com/hrittich/lfa-lab*

Coarsening the Edge Problem

- What makes matters worse here is that each component has different anisotropy!

$$\begin{bmatrix} -\frac{\partial^2}{\partial y^2}-\frac{\partial^2}{\partial z^2} & \frac{\partial}{\partial x \partial y} & \frac{\partial}{\partial x \partial y} \\ \frac{\partial}{\partial x \partial y} & -\frac{\partial^2}{\partial x^2}-\frac{\partial^2}{\partial z^2} & \frac{\partial}{\partial y \partial z} \\ \frac{\partial}{\partial x \partial z} & \frac{\partial}{\partial y \partial z} & -\frac{\partial^2}{\partial x^2}-\frac{\partial^2}{\partial y^2} \end{bmatrix}$$

- Result: Standard “nodal” coarsening will not work well with point smoothers!

- Each component (x,y,z) needs to be coarsened separately…. But cannot be allowed to mix!

- Approach to take: Composite prolongator based on block-diagonalization of matrix.

- Some Classical AMG codes do this by default, but SA-AMG codes generally

And What About Boundary Conditions?

- For Dirichlet boundaries on the original edge element system, you get this:

$$\boldsymbol{n} \times \boldsymbol{E} = 0$$

- This constraints **E** only in directions tangent to the interface.

- Treating a Dirichlet node in the edge auxiliary preconditioner as a Dirichlet point in all 3 directions is not right.

- A roller-like condition (Dirichlet only in tangent directions) is probably what we want.

- Yang and Arnold (2019) have recent work in this area that we're still trying to parse.

Outline

- Introduction to the Eddy Current Maxwell’s Equations
- Motivating Fourier Analysis
- A Closer Look at the Edge Hierarchy
- Results & Conclusions

Test Problems

- Problem 1: Uniform mesh w/ 40 or 80 elements per dim and, $\sigma = 1$.

- Problem 2: 90 degree wedge of radial mesh (with hole near the origin) and $\mu = \mu_0$ and σ = 1e-2 with a narrow shell of σ = 1e6 in the middle of the domain.

- Test: Multigrid-preconditioned CG w/ tolerance of 1e-10.
  - 1 SGS sweep pre-and post. Multigrid for node hierarchy is fixed.
  - Compare block diagonal to nodal coarsening w/ and w/o dropping.

- Software: Trilinos/MueLu github.com/trilinos/Trilinos

Computational Results: Uniform Mesh

	NX=40, nodrop	NX=40,drop	NX=80,nodrop	NX=80,drop
Nodal, 2 Level	12	12	16	16
Nodal, Multilevel	11	11	15	15
BlockDiag, 2 Level	12	13	16	12
BlockDiag, Multilevel	11	12	15	13

On the small problem, all methods are similar.

Computational Results: Uniform Mesh

	NX=40, nodrop	NX=40,drop	NX=80,nodrop	NX=80,drop
Nodal, 2 Level	12	12	16	16
Nodal, Multilevel	11	11	15	15
BlockDiag, 2 Level	12	13	16	12
BlockDiag, Multilevel	11	12	15	13

On the small problem, all methods are similar.

On larger, problems, the block diagonal version does better

Computational Results: Uniform Mesh

	NX=40, nodrop	NX=40,drop	NX=80,nodrop	NX=80,drop
Nodal, 2 Level	12	12	16	16
Nodal, Multilevel	11	11	15	15
BlockDiag, 2 Level	12	13	16	12
BlockDiag, Multilevel	11	12	15	13

On the small problem, all methods are similar.

On larger, problems, the block diagonal version does better

If you add dropping, the iteration counts on the larger mesh don't increase by 25% → Improved scalability!

Computational Results: Radial Mesh

	Ref=1, nodrop	Ref=1,drop	Ref=2,nodrop	Ref=2,drop
Nodal, 2 Level	28	28	39	38
Nodal, Multilevel	26	26	36	35
BlockDiag, 2 Level	28	28	39	37
BlockDiag, Multilevel	26	24	36	31

- Modest improvement in iteration count w/ dropping enabled.
- Does not seem to scale optimally even though it scales better.
- We’re that hoping improving boundary conditions will make an impact here.

Conclusions

- Revisited auxiliary preconditioners for eddy current Maxwells' equations.

- Analyst's observation led to identification of problems in previous approaches.

- Edge hierarchy dofs were not being coarsened correctly.
    - Need the "block diagonal approach"

- Edge hierarchy boundary conditions are problematic (still work in progress).

- Modest improvements to performance so far.

Kolev and Vassilevski. Parallel Auxiliary Space AMG for H(Curl) Problems. J. Comp. Math, 2008.
Bochev, Hu, Siefert and Tuminaro. An Algebraic Multigrid Approach Based on a Compatible Gauge Reformulation of Maxwell's Equations. SIAM J. Sci. Comput. 2008.
Boonen, van Lent and Vandewalle, Local Fourier Analysis of Multigrid for the Curl-Curl Equation, SIAM J. Sci. Compt., 2008.