Preconditioning over Multiple Fourier Components in NIMROD and NIMSTELL

Carl Sovinec

Department of Engineering Physics
University of Wisconsin-Madison

Center for Tokamak Transient Simulation Meeting
October 16, 2022

Supported by US DOE grants DE-SC0018642 and DE-SC0018001

Center for Tokamak Transient Simulation

Motivation

- Nonlinear asymmetry leads to implicit coupling among toroidal Fourier harmonics in NIMROD.
- Toroidal shaping of the mesh and equilibria couple harmonics in NIMSTELL.
- Preconditioning over multiple Fourier components is important for efficiency.

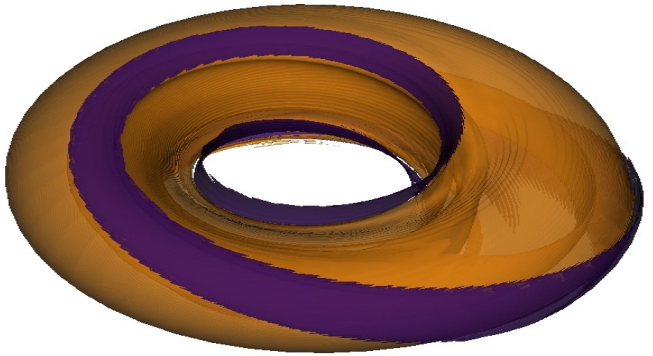

Isosurfaces of $J_{\parallel}/B$ from a nonlinear VDE computation with NIMROD.

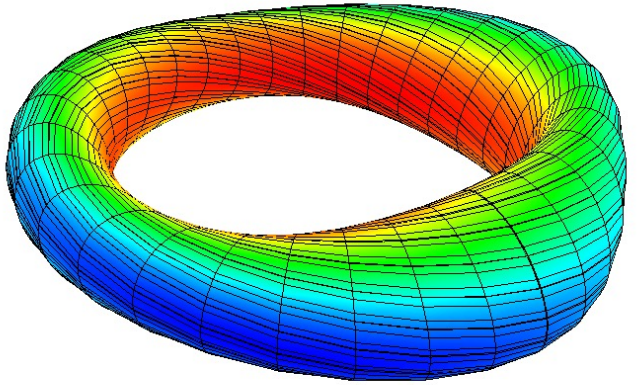

Example $|\boldsymbol{B}|$ of 3D vacuum field that NIMSTELL solved on a twisting mesh.

Many approaches to improve toroidal preconditioning have been tested in NIMROD.

1. Partial Gauss-Seidel-like updates with limited coupling within process layers helped in some cases.
2. Flexible GMRES (FGMRES) alternating between block Fourier diagonal (bFd) and block toroidal-plane diagonal is not better than bFd alone.
3. Kyle Bunkers tried Fourier multigrid and found that coarse-representation solves did not help the full solve.
4. 1D solves over toroidal angle showed little synergistic benefit with bFd with additive, multiplicative, or FGMRES approaches.
5. Fully coupled static condensation to reduce the system size reduces iteration count by ~ 50%.

Direct solves over multiple Fourier components and the full poloidal plane has been implemented in NIMSTELL.

- Modifications to the ***V***-matrix routine produces systems for bands of Fourier components.
  - Numerical integration is over all toroidal mesh points for FFTs.
  - Bands may overlap.
  - Systems use real algebra for separated real/imaginary components.

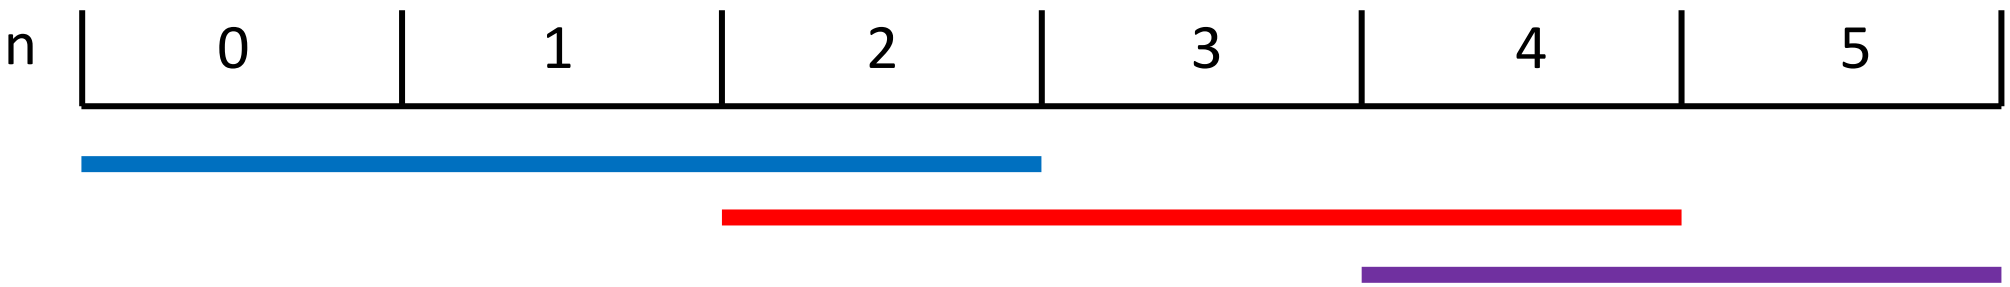

Example set of matrices with *stride* = 2 and *overlap* = 1.

A parameter scan uses a cylindrical interchange mode in a twisted mesh.

- The *m*=4, *n*=1 mode is resonant at $r = 0.371$, where $D_s = 0.443$.
- Twisting the mesh one turn puts equilibrium $\boldsymbol{B}_{\text{pol}}$ into the n=1 component.

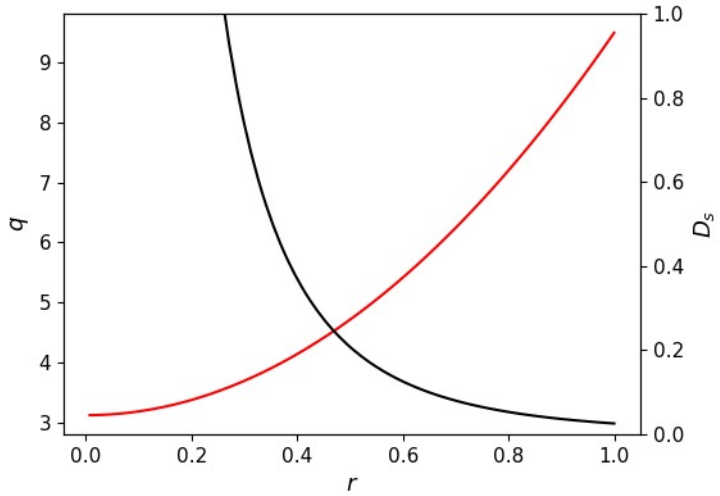

Profiles of safety factor (red) and Suydam parameter (black).

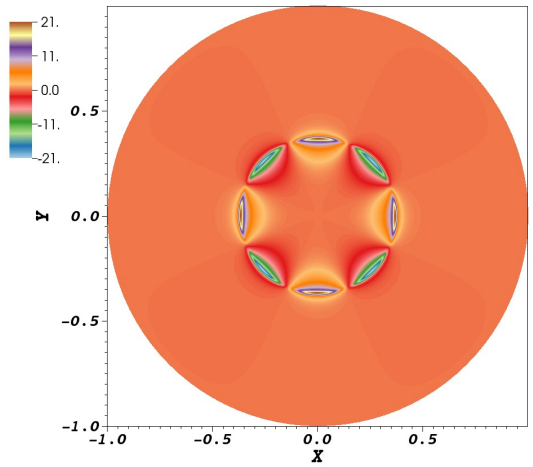

NIMSTELL linear result on A_z eigenfunction component with all dissipation parameters set to 10^{-7}.

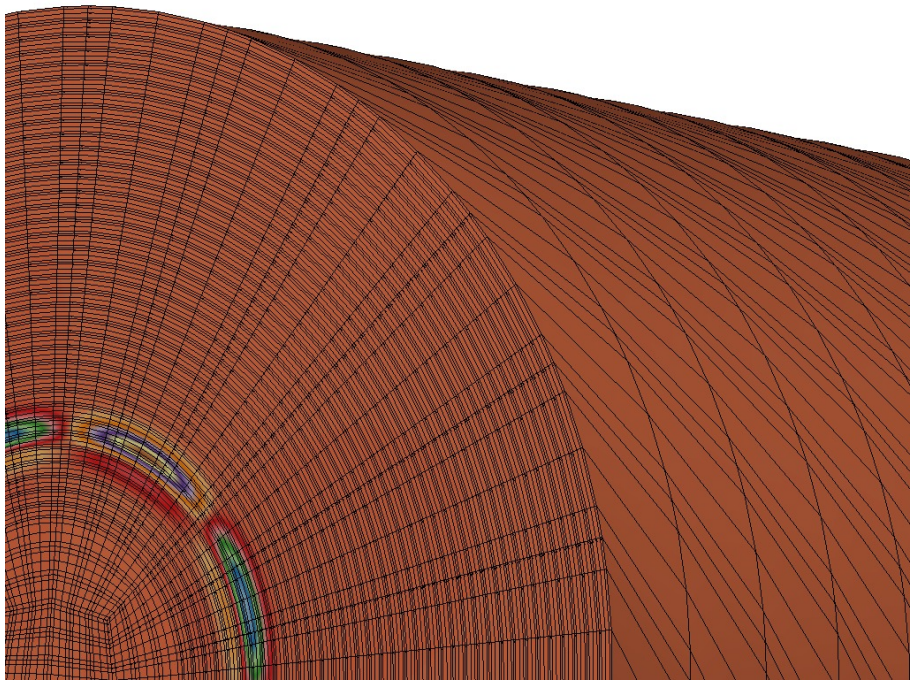

Twisting does not match the mode helicity.

Thirteen computations with a small mesh compare performance when varying stride and overlap.

- Elements are bicubic. Fourier representation is $0 \leq n \leq 5$.
- Outer part of the mesh is a ring of 18x24 elements; 540 total elements.
- The twisted mesh leads to strong ± 1 coupling among Fourier components in the semi-implicit operator for $\Delta \boldsymbol{V}$.
- Two computations use block Fourier diagonal (“bFd”) preconditioning for the *V*-advance.
- Ten have varied stride (“st”) and overlap (“ov”) with real-algebra solves for the *V*-advance.
- The ”st6 dg” uses our “diagonal” preconditioner applied to the full matrix.
- Computations were run on 5 cores of a shared-memory workstation without layer parallelism.
- Timings include the advances for ***A***, *n*, and *T* that use bFd with SuperLU_DIST (small contributions).

Iteration is weakly dependent on stride; it jumps to low numbers with an overlap of 2.

- Computations are run 10 steps, and average iteration count for ΔV is shown.
- Computations with overlap needed GMRES, which affects timings. Others used CG.

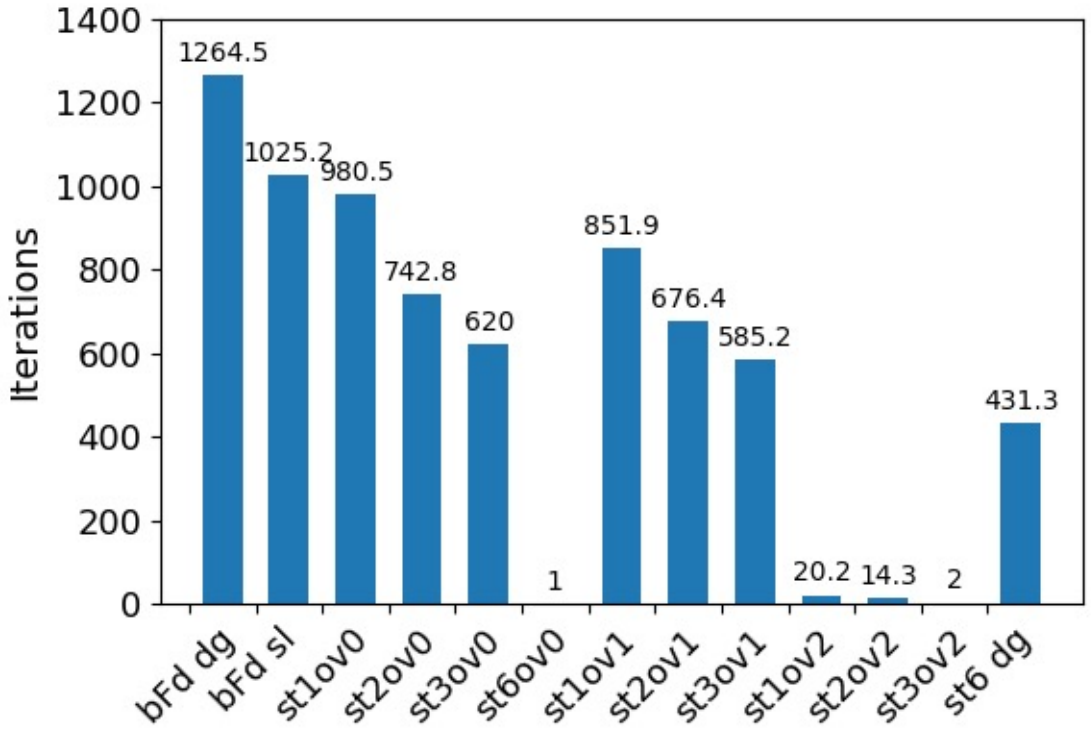

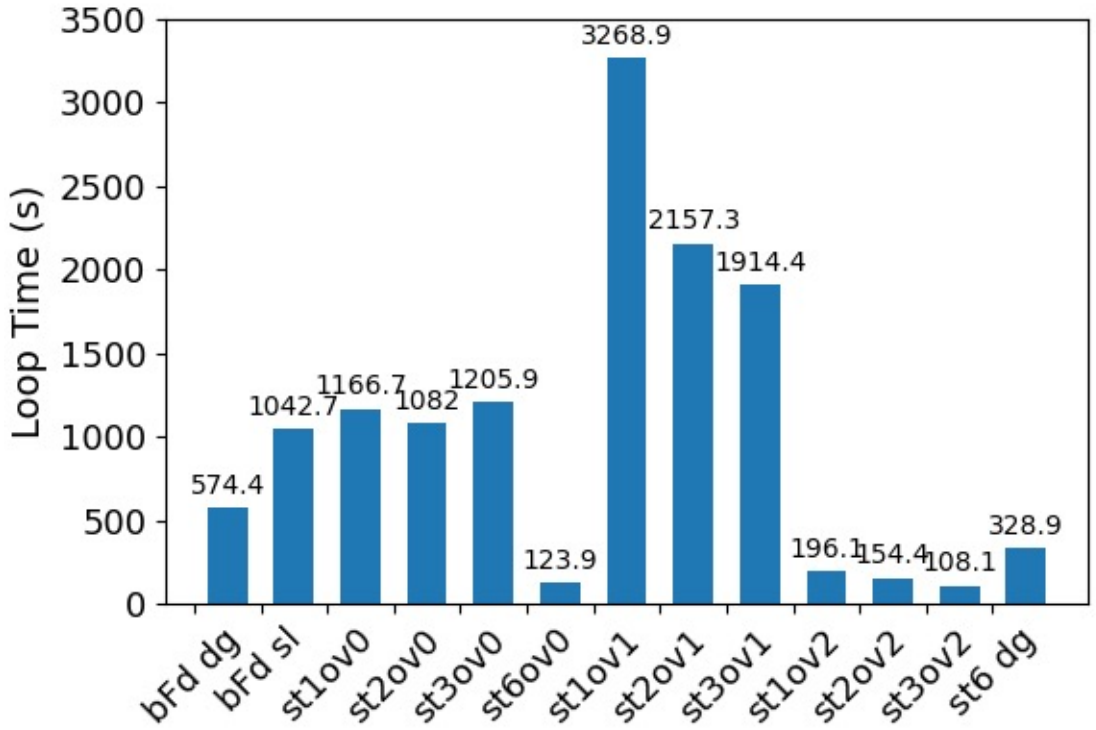

Comparing matrix and solve times shows a tradeoff.

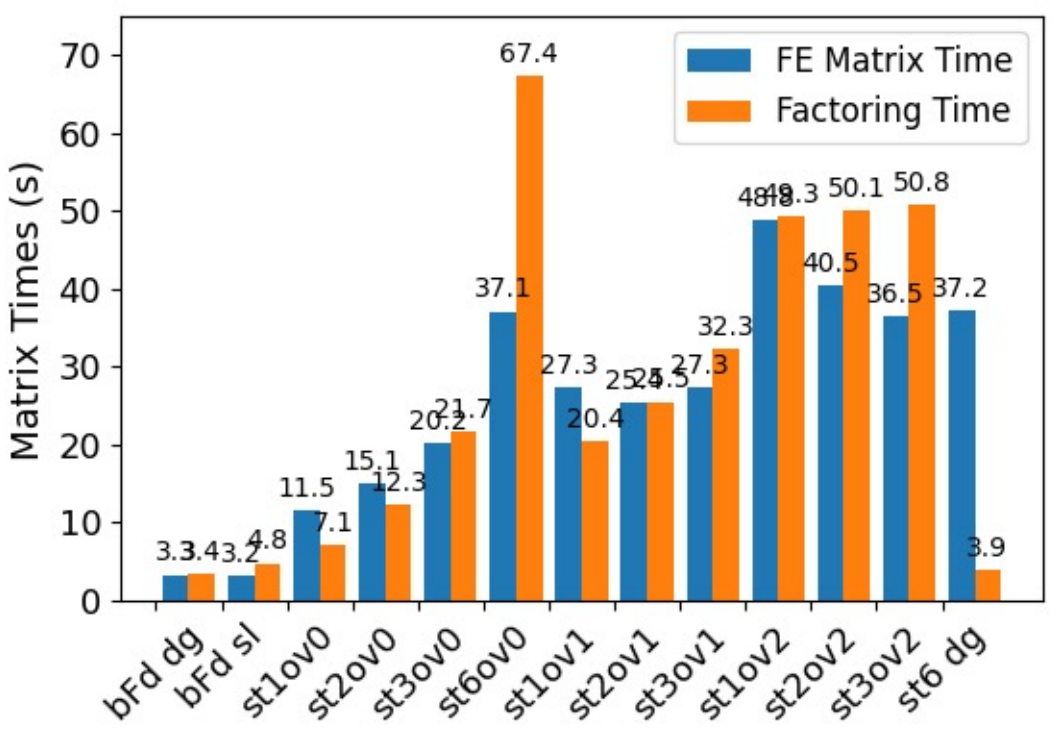

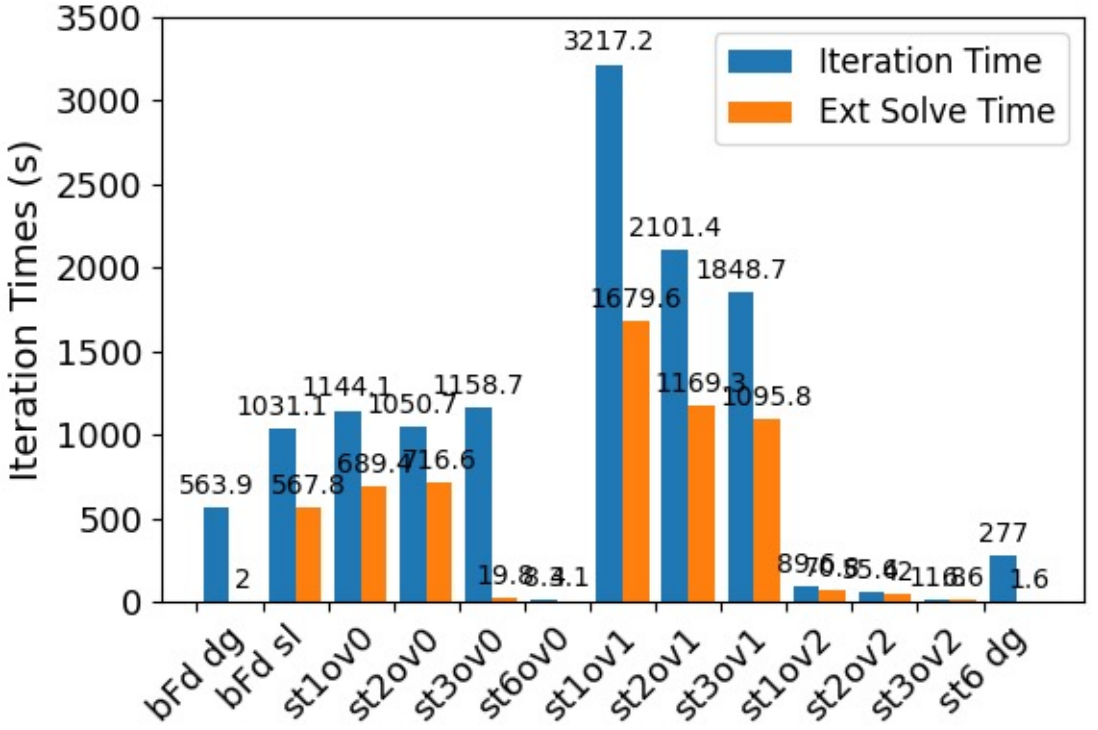

- The reasonable performance of bFd-diagonal probably stems from the small problem size.

Is a complex matrix for $n \geq 0$ possible?

- The working implementation uses real algebra with the system split into real and imaginary degrees of freedom.
- Our complex Fourier expansion uses:

$$A(\zeta) = A_0 + \sum_{n=1}^{N} \left[A_n e^{in\zeta} + cc. \right] = A_0 + \sum_{n=1}^{N} \left[2A_{n_r} \cos(n\zeta) - 2A_{n_i} \sin(n\zeta) \right]$$

- A complex matrix for $n \geq 0$ would use ½ of the memory.
- There are also fewer operations.
  - A test shows SLU_DIST factoring time is reduced by 2.

A complex matrix implies a restriction.

- Real-algebra Fourier basis functions are $\cos(n\zeta)$ and $-\sin(n\zeta)$.
- Complex-algebra Fourier basis functions are $\exp(in\zeta)$.
- If the system can be represented by a complex matrix,

$$\begin{pmatrix} a & b \\ c & d \end{pmatrix}\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{pmatrix} r_1 \\ r_2 \end{pmatrix} \rightarrow \begin{pmatrix} a_r + ia_i & b_r + ib_i \\ c_r + ic_i & d_r + id_i \end{pmatrix}\begin{pmatrix} x_{1_r} + ix_{1_i} \\ x_{2_r} + ix_{2_i} \end{pmatrix} = \begin{pmatrix} r_{1_r} + ir_{1_i} \\ r_{2_r} + ir_{2_i} \end{pmatrix}$$

then the equivalent real matrix is restricted:

$$\begin{pmatrix} a_r & -a_i & b_r & -b_i \\ a_i & a_r & b_i & b_r \\ c_r & -c_i & d_r & -d_i \\ c_i & c_r & d_i & d_r \end{pmatrix}\begin{pmatrix} x_{1_r} \\ x_{1_i} \\ x_{2_r} \\ x_{2_i} \end{pmatrix} = \begin{pmatrix} r_{1_r} \\ r_{1_i} \\ r_{2_r} \\ r_{2_i} \end{pmatrix}$$

A complex implementation indicates limitations.

- The complex operation is effective for the ***V***-matrix equation with non-twisting meshes.
  - The Krylov method takes only 1 or 2 steps when including all Fourier components in the matrix.
- It does not converge with a twisting mesh.
  - Printing some components of the *real-algebra* version shows that real and imaginary parts of diagonals are not the same.
  - This violates the restriction on the previous page!

A simple example shows the limitation.

- Consider diagonal matrix elements that result from a diffusion operation with diffusivity $D(\zeta) = \begin{cases} D_0, \frac{3\pi}{4} \leq \zeta \leq \frac{5\pi}{4} \\ 0, \text{otherwise} \end{cases}$.
- "Real" phase, $\cos\zeta$: $D_0 \int_{\frac{3\pi}{4}}^{\frac{5\pi}{4}} d\zeta \left(\frac{d}{d\zeta}\cos\zeta\right)^2 = D_0\left(\frac{\pi}{4} - \frac{1}{2}\right)$
- "Imaginary" phase, $-\sin\zeta$: $D_0 \int_{\frac{3\pi}{4}}^{\frac{5\pi}{4}} d\zeta \left(\frac{d}{d\zeta}\sin\zeta\right)^2 = D_0\left(\frac{\pi}{4} + \frac{1}{2}\right)$
- Complex representation, $\exp i\zeta$: $D_0 \int_{\frac{3\pi}{4}}^{\frac{5\pi}{4}} d\zeta \left(\frac{d}{d\zeta}e^{i\zeta}\right)\left(\frac{d}{d\zeta}e^{-i\zeta}\right) = D_0\frac{\pi}{2}$
- This example indicates that we must separate the phases in real matrices.

MPI communication for decomposition over ζ is now complete.

- Matrix construction subdivides NIMROD's grid blocks.
  - Integration over all ζ is needed.
  - The decomposition is the same as used for projecting vectors.
  - Asynchronous communication is used to gather contributions.
- Preconditioning operations need to share Fourier components.
  - Matrix-band harmonics may differ from layer decomposition.
  - Asynchronous communication is used to exchange harmonics.

Static condensation over 3D can be used in other ways.

- Static condensation is element-based matrix partitioning,

$$\left(\underline{\boldsymbol{M}}_{11} - \underline{\boldsymbol{M}}_{12}\underline{\boldsymbol{M}}_{22}^{-1}\underline{\boldsymbol{M}}_{21}\right)\underline{x}_1 = \underline{b}_1 - \underline{\boldsymbol{M}}_{12}\underline{\boldsymbol{M}}_{22}^{-1}\underline{b}_2 \ ,$$

  where $\underline{x}_2 = \underline{\boldsymbol{M}}_{22}^{-1}\left(\underline{b}_2 - \underline{\boldsymbol{M}}_{21}\underline{x}_1\right)$ has the degrees of freedom from element interiors.
- In the banded systems discussed above, the static condensation is over each band.
- Forming $\underline{\boldsymbol{M}}_{12}$, $\underline{\boldsymbol{M}}_{21}$, and $\underline{\boldsymbol{M}}_{22}$, involves element-local FE computations, which is straightforward for layer parallelism.
- Will the block-Fourier diagonals of the Schur complement (top equation) make a scalable preconditioner in combination with the 3D static condensation?

Conclusions and Discussion

- Direct solves over multiple Fourier components have been implemented.
  - Full matrix computation and factoring will not be tractable for large problems.
- Overlapping 2 Fourier components shows good performance in the test case.
  - This may be related to the strong ± 1 coupling.
  - CG with additive overlap should be possible.
- Parallel domain decomposition has just been implemented.
- 3D static condensation with bFd is worth trying.
- The developments will be ported to NIMROD.