RW convergence with NIMbnd

J King, E Howell, S Kruger (Tech-X)
D Barnes, C Ackay, J Finn (Tibbar Tech)

Resistive-wall with response matrix pre-CTTS

- Thin-wall resistive-wall implementation through auxiliary equation for normal B by Andi Becerra (outside of this talk)
- Uses GRIN to compute response matrix
- Two issues limited convergence:
  - First: related to how Andi ran cases (not covered here)
  - Second: GRIN uses methods that limit convergence to 2nd order
- Motivated development of NIMbnd by Tibbar Tech (Dan Barnes)
  - Use Nystrom collocation to discretize boundary
  - Internally interpolates to NIMROD GLL node locations on boundary
  - Uses special quadrature rules to resolve logarithmic singularities in equations
  - Convergence properties verified on two test problems independent of NIMROD
  - See https://ctts.pppl.gov/APS2018/Akcay.pdf

Major points

- Using growth rate to determine convergence rate works poorly
- Determining the convergence rate from spectral error works well
- Even with smooth profiles: x-direction error dominates convergence
- NIMbnd exhibits better convergence rate than GRIN and matches the convergence rate of the analytic cylindrical response matrix

Starting with cases based on those from AB's thesis

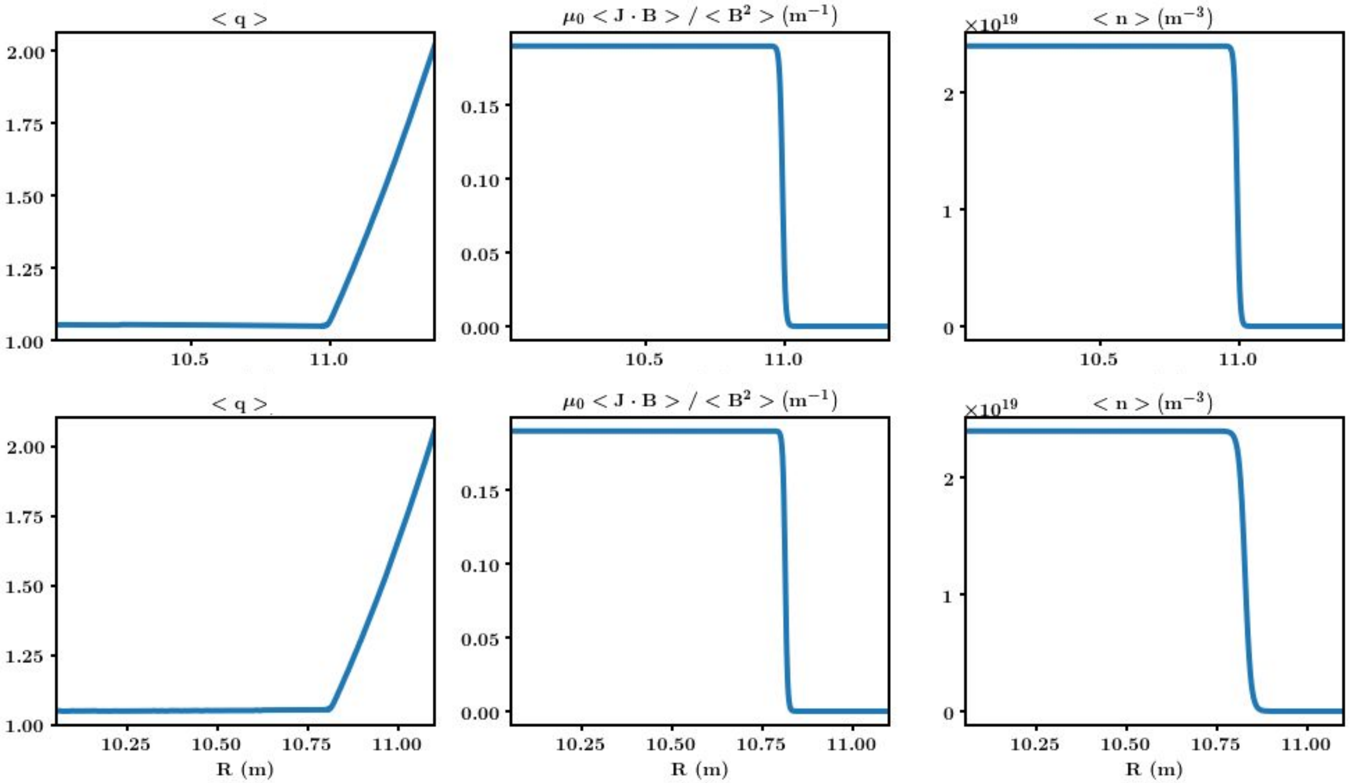

Cylinder (top)

The tanh width is larger

Large-aspect-ratio torus (bottom)

Density not a FS quantity => smooths FS avg

New grid packing for toroidal case

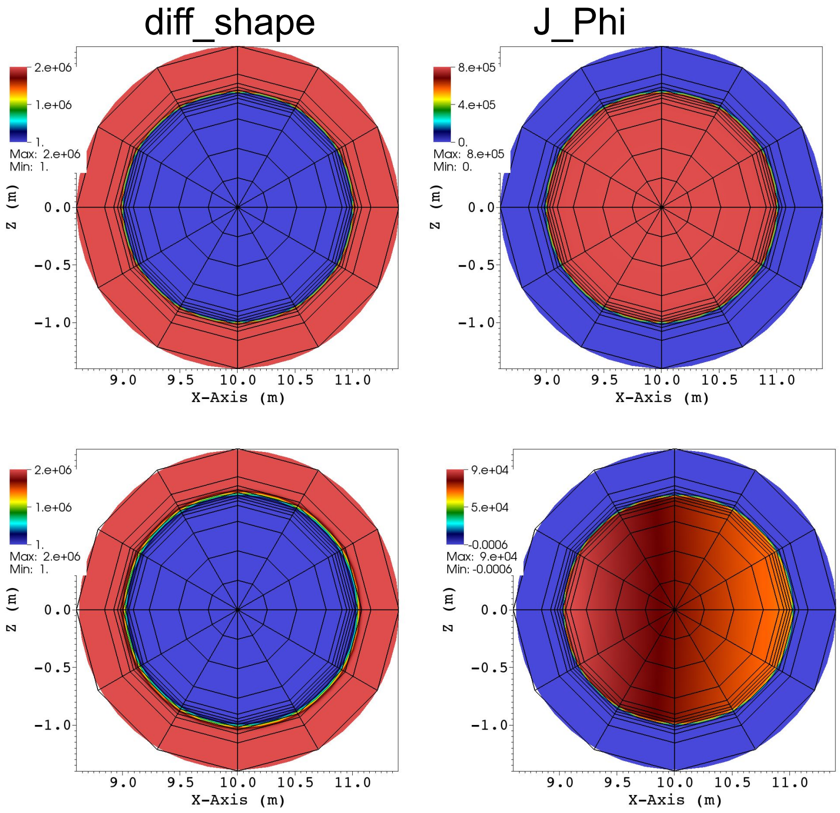

Cylindrical (top) packing slightly adjusted

Toroidal (bottom): recently implemented grid packing for geom='tor' used

Grid not flux aligned

Part I: cylindrical cases

- Use analytic response matrix
- Objectives:
  - Test if case is reasonable
  - Test if method of determining rate of convergence is correct

Convergence rates determined from growth rate are inconclusive

But the rate of convergence from least-squares fit to the growth rate makes no sense

Conclusion: this is a poor method of judging convergence

(5 highest-resolution points are used in all fits)

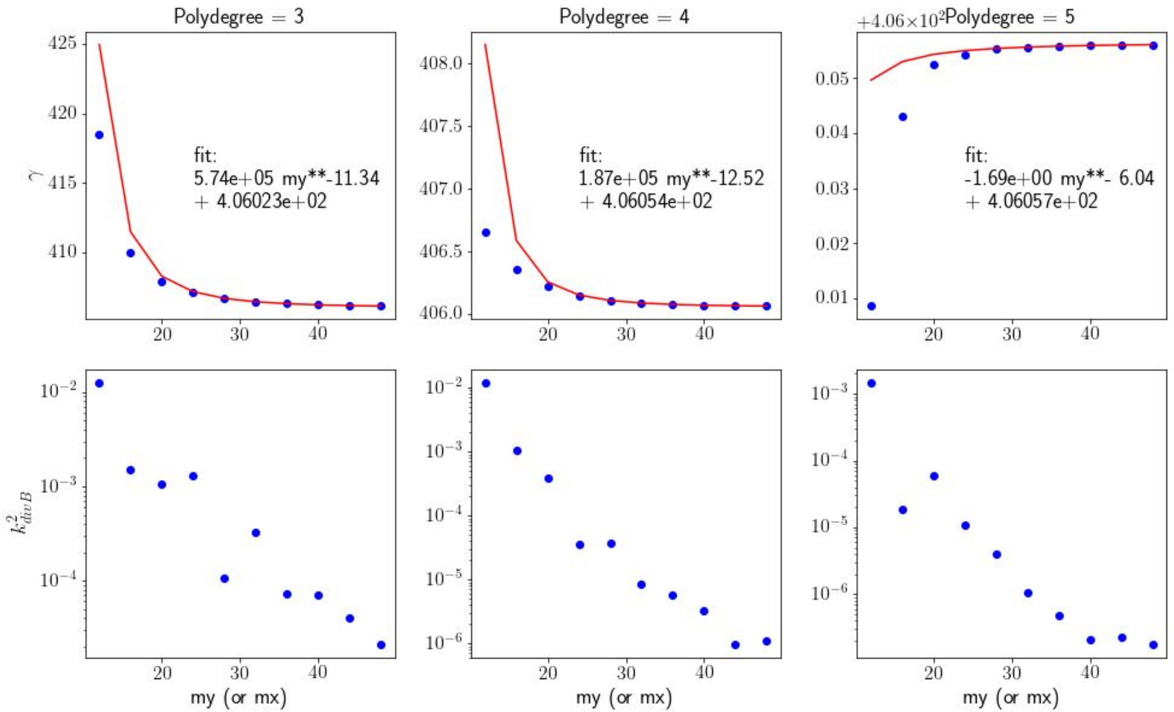

Spectral energy is a better metric for convergence rate

Within each element project to the Legendre basis:

$$a_{lm} = \frac{2l+1}{2}\frac{2m+1}{2}\int f(\mathbf{x})\, P_l(x)\, P_m(y)\, dxdy$$

Thus for poly_degree=p

$$f(\mathbf{x}) = \sum_{l,m=0}^{p} a_{lm} P_l(x)\, P_m(y)$$

Define total error as

$$error^2 = \left(\sum_{l=0}^{p} a_{lp}^2 + \sum_{m=0}^{p-1} a_{pm}^2\right) / \sum_{l,m=0}^{p} a_{lm}^2$$

The x-direction error is defined as

$$error_x^2 = \sum_{l=0}^{p} a_{lp}^2 / \sum_{l,m=0}^{p} a_{lm}^2$$

The y-direction error is defined as

$$error_y^2 = \sum_{}^{p} a_{pm}^2 / \sum_{}^{p} a_{lm}^2$$

Graphically for pd=3 by a_{lm}:

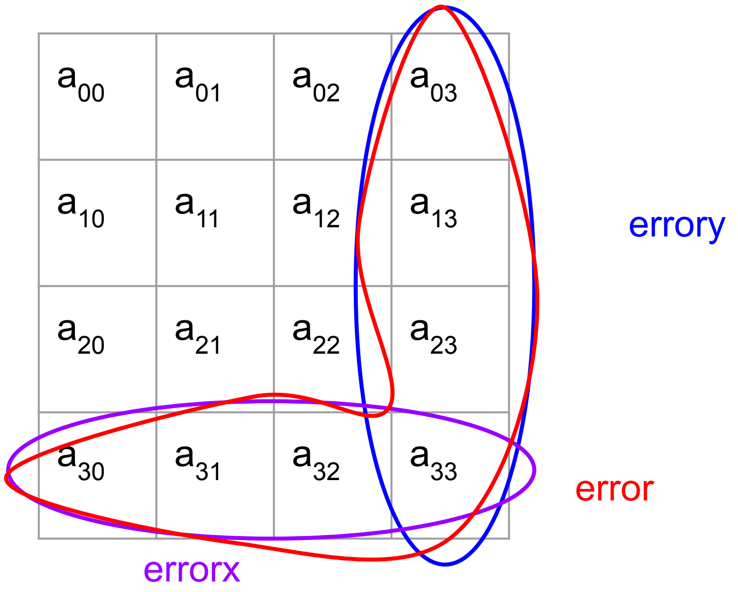

Spectral energy is a better metric for convergence rate

Overall convergence (top) is dominated by error in the x-direction (middle)

Expected convergence of my**pd is observed in the y-direction (bottom)

(tanh-gq ds_function is NOT used)

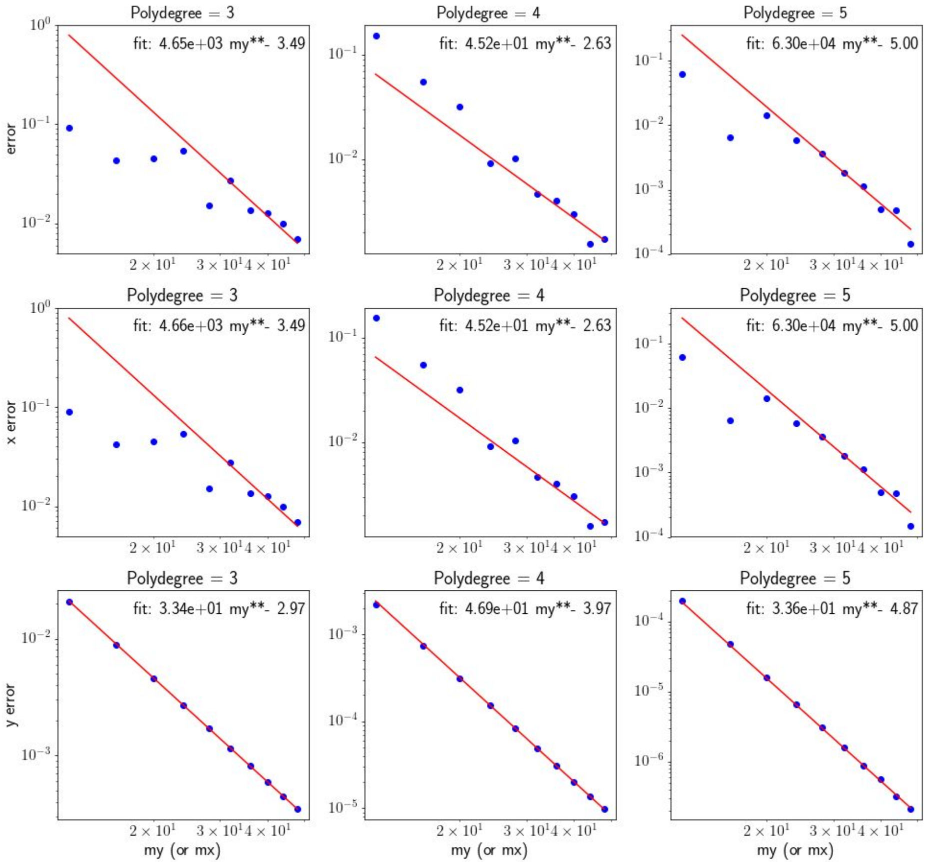

Part II: Large-aspect-ratio torus cases

- Objective: Compare convergence with GRIN and NIMbnd

GRIN: convergence by growth rate is mixed

(5 highest-resolution points are used in all fits)

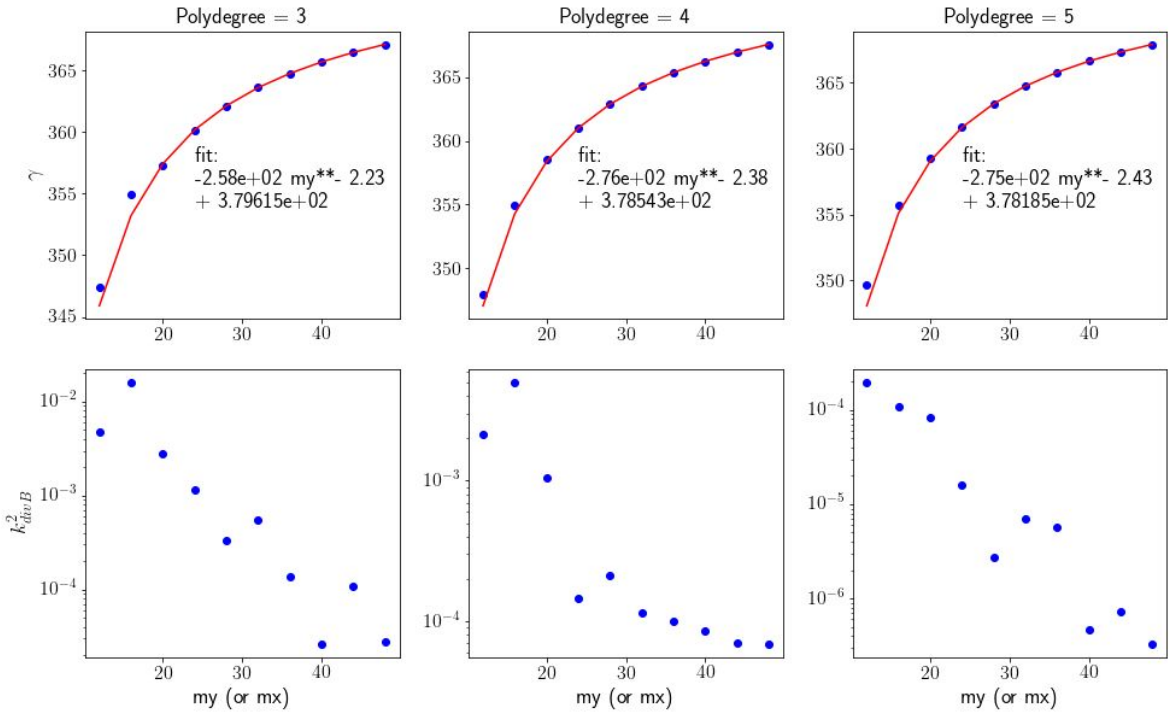

GRIN: spectral energy y-dir convergence rate is low

Convergence in the y-direction (associated with the resistive-wall boundary) is lower than expected

Total error is dominated by the x-direction

Reminder: x/y directions and radial/poloidal coordinates are slightly mixed for toroidal cases without grid flux alignment

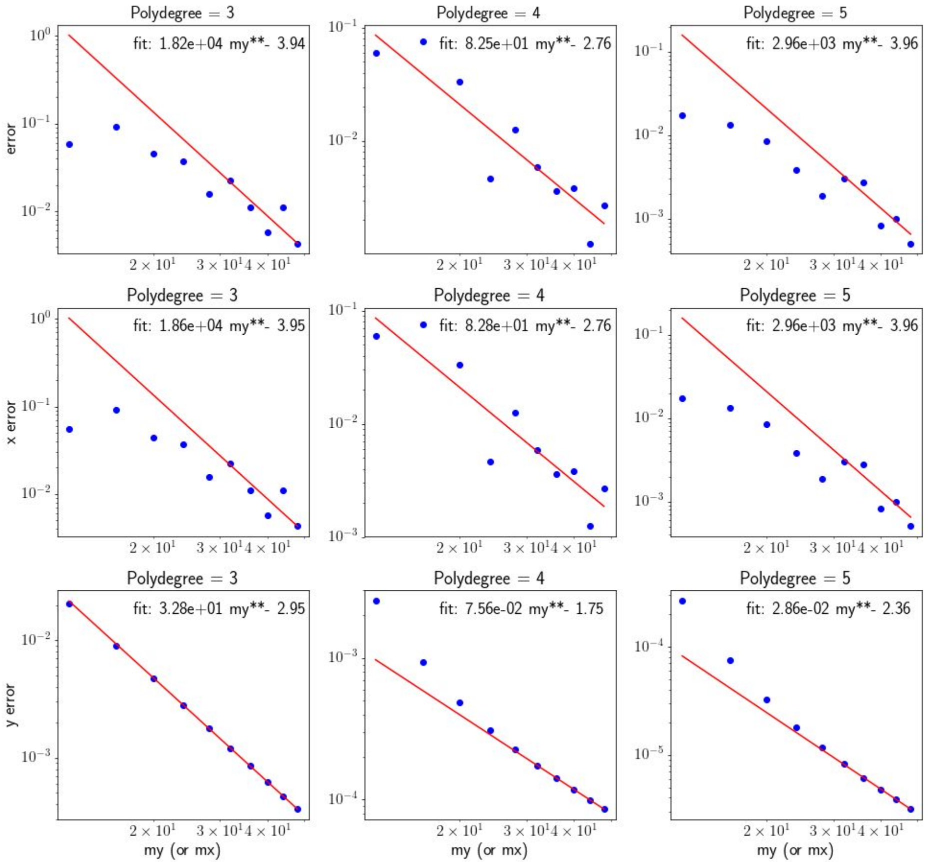

NIMbnd: convergence by growth rate is mixed

Comparable result to GRIN

(5 highest-resolution points are used in all fits)

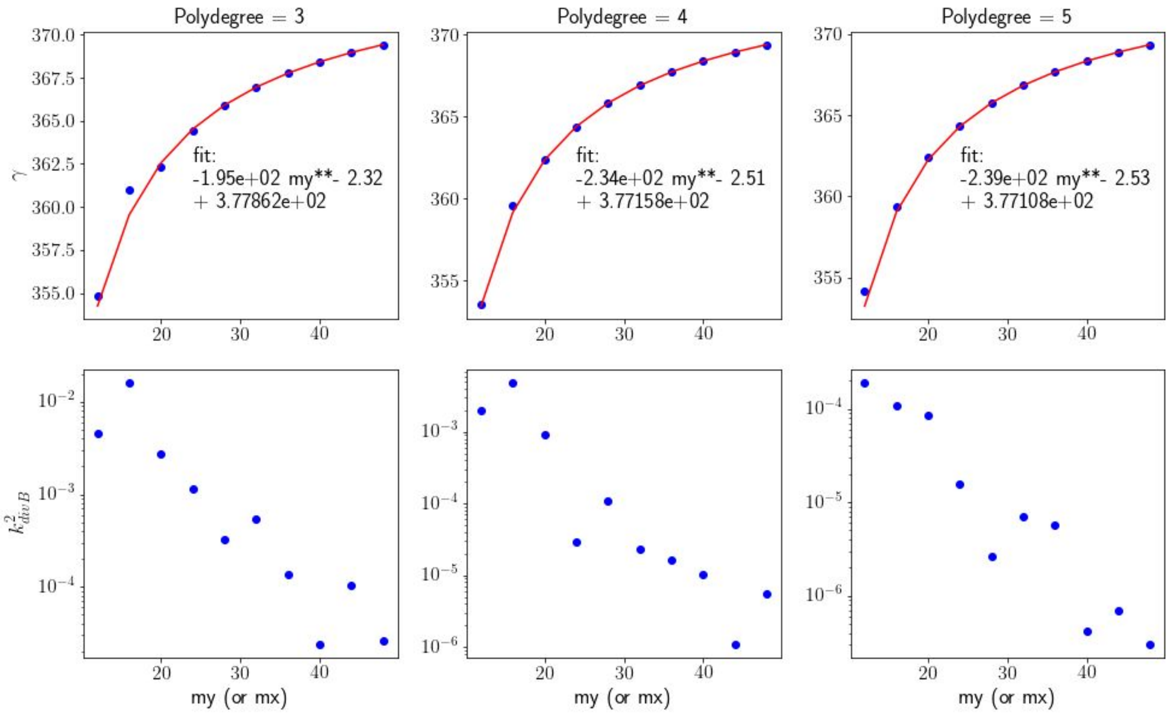

NIMbnd: spectral energy y-dir convergence rate is good

Convergence in the y-direction (associated with the resistive-wall boundary) is at rate polydegree

Total error is dominated by the x-direction

Reminder: x/y directions and radial/poloidal coordinates are slightly mixed for toroidal cases without grid flux alignment

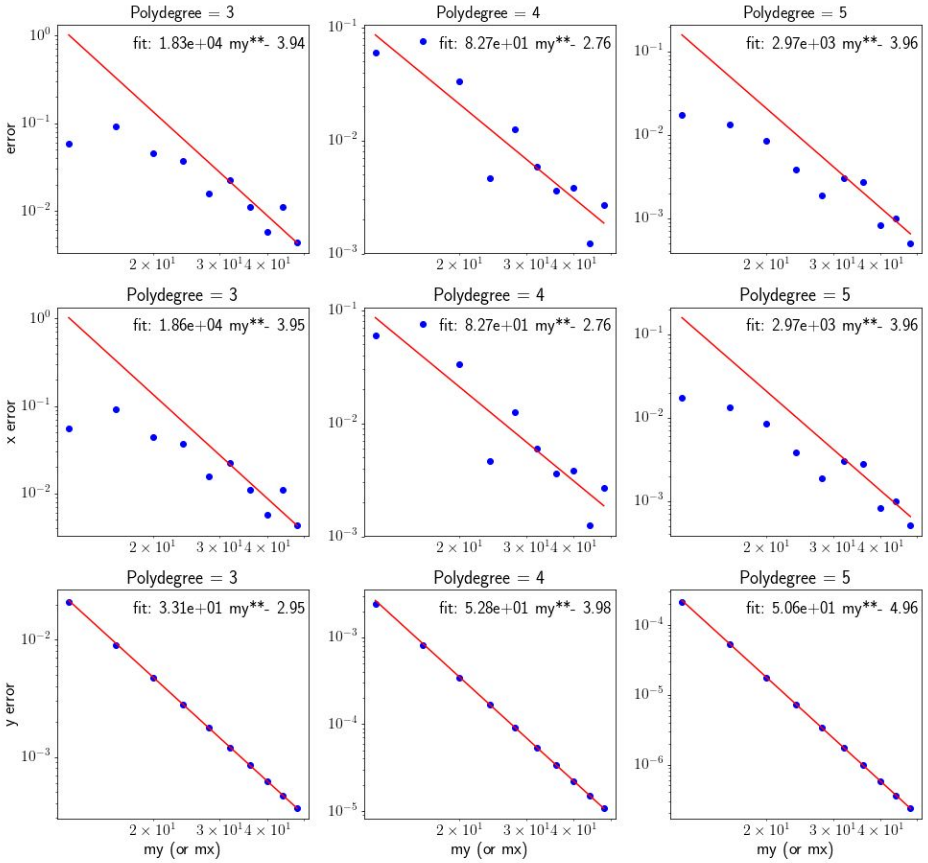

Outstanding issues with NIMbnd

- Implementation mature and easy to use
- Use of derivative matrix produces the same results (not shown)
- Open question: use transpose or x/y flip operation on NIMbnd matrix
  - Results here are with x/y flip
  - Cihan uses transpose
  - Tests (not shown) show comparable results on large-aspect-ratio torus
- Need better test case: working on diverted, realistic-aspect-ratio case with shaping (next slides)

Realistic-aspect-ratio case -- work in progress

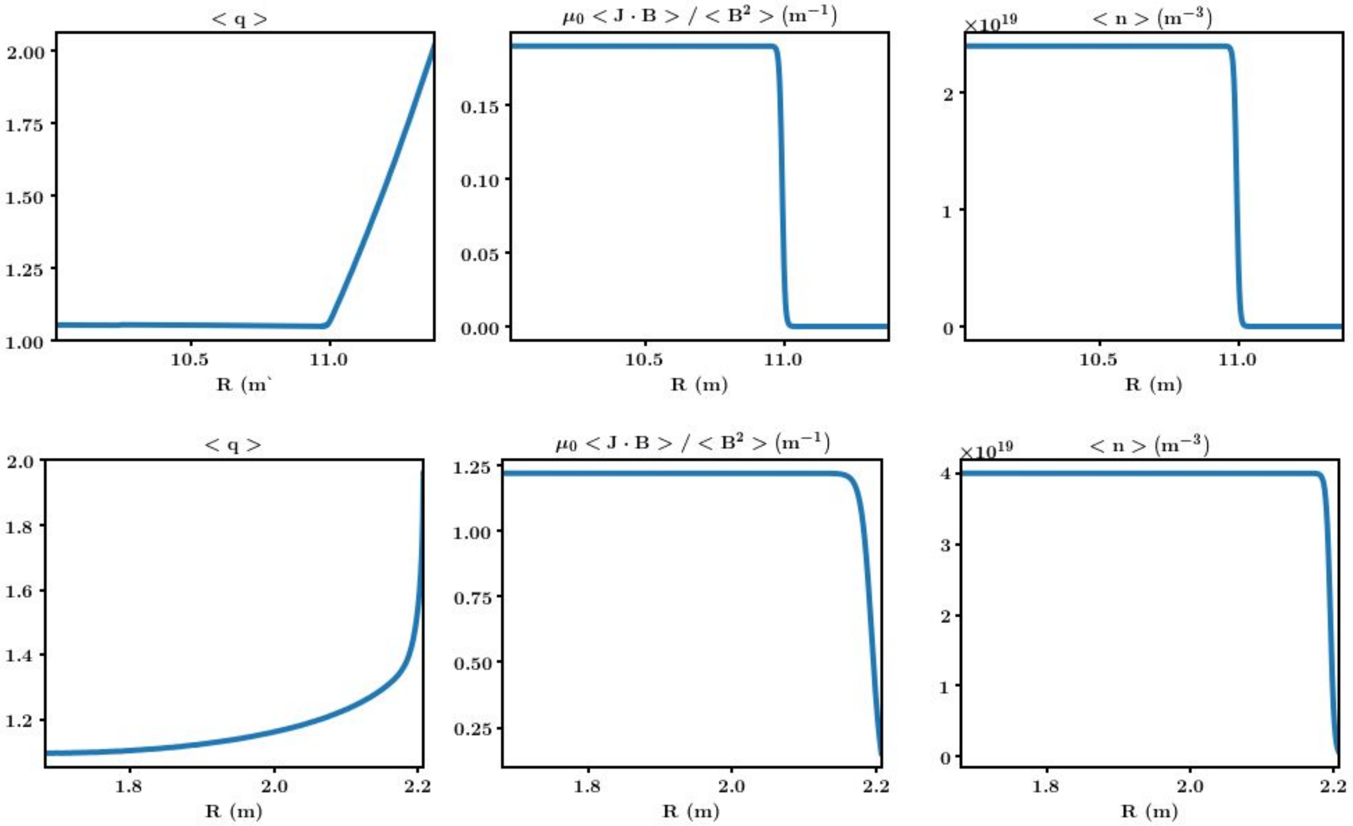

Cylinder (top)

Realistic aspect-ratio torus (bottom)

Plots end at separatrix

Realistic-aspect-ratio case -- work in progress

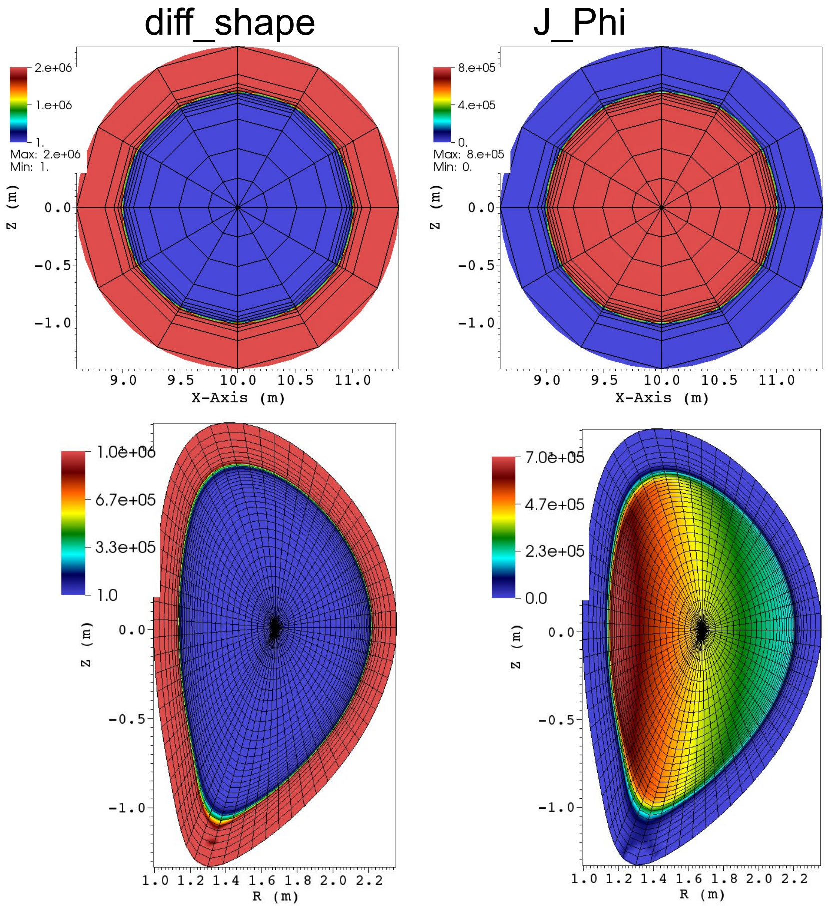

Cylindrical (top)

Realistic aspect ratio, Toroidal, Diverted (bottom)

Aspect ratio = 2.3

NIMbnd produces a lower kdivb error than GRIN

- Neither cases shows “smooth” convergence in growth rate
- Challenges relative to large AR cases:
  - Grad-Shafranov solve error
  - Convergence associated with x-point singularity
  - Shaped boundary/plasma profiles
  - Spectral broadening in the poloidal Fourier modes
- Expect that we need to run preprocessor more carefully for better results

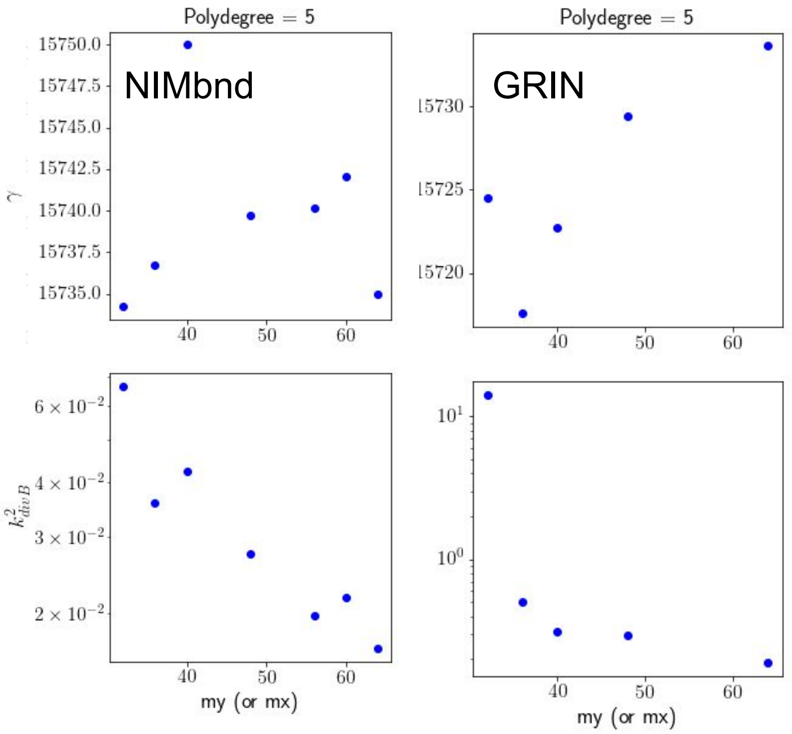

NIMbnd spectral error is lower than GRIN

- Neither cases shows “smooth” convergence in growth rate
- Challenges relative to large AR cases:
  - Grad-Shafranov solve error
  - Convergence associated with x-point singularity
  - Shaped boundary/plasma profiles
  - Spectral broadening in the poloidal Fourier modes
- Expect that we need to run preprocessor more carefully for better results

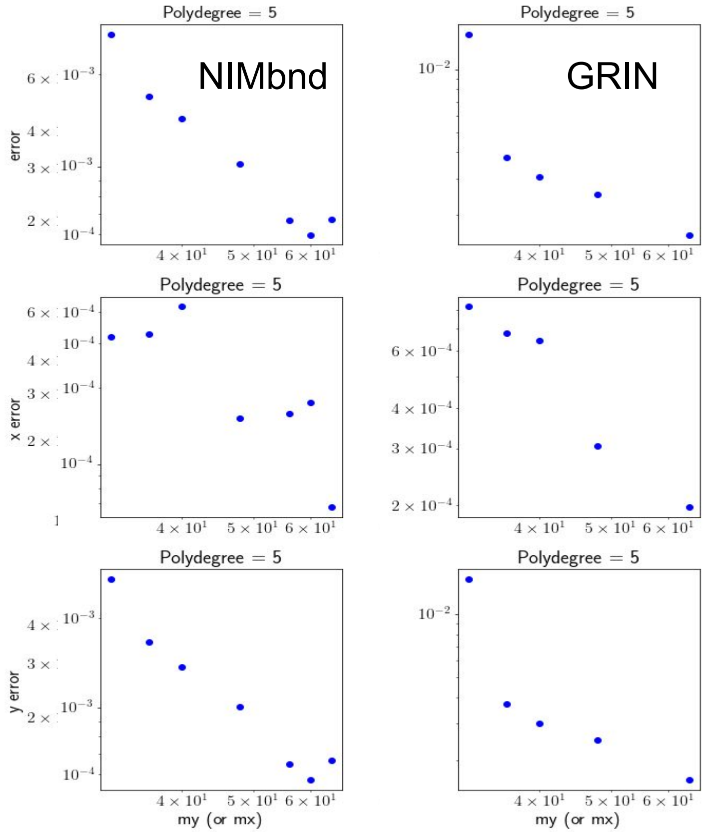

Needs for NTM / locked mode disruption simulation

- Resistive wall (this talk; through Tibbar Tech / Tech-X collaboration)
- Neoclassical closures (Howell https://ctts.pppl.gov/APS2018/Howell.pdf)
- Improvements to Grad-Shafranov solver to control drives (work by Howell presented in https://ctts.pppl.gov/Sherwood2018/King.pdf)
  - E.g. increase current to drive TM or modify profiles to stabilize
- Field error interaction and nonlinear simulation challenges (Howell today)

Using GPUs in NIMROD

J King, E Howell, S Kruger (Tech-X)
C Sovinec, B Cornille (UW-Madison)
S Taheri (UWash) C Kim (SLS2)
A Spencer (USU)

Topics:

- We have developed a plan for placing the NIMROD FE integration on GPUs
- Test toy code performance with NIMROD-like data structures and OpenACC

Challenge: GPU FE integration with Fourier decomp

- Fourier decomposition presently requires communication during FE integration for matrix-vector-dot product and RHS integration
- Anticipate this synchronization to be costly on GPUs
- GPU plan: use a different decomposition with all Fourier modes on each GPU (use (b) not (a) for data layout)

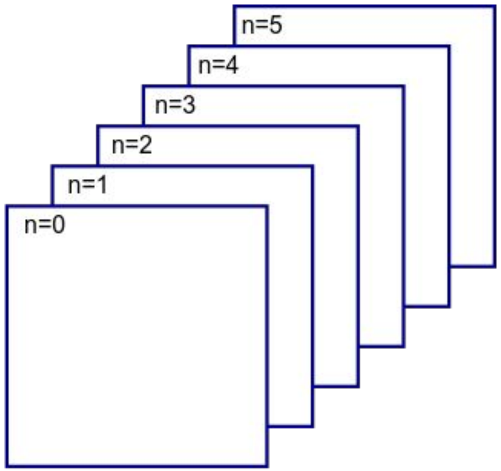

(a) toroidal (φ) decomposition into Fourier modes of RZ-FE planes where complex coefficients correspond to each mode #

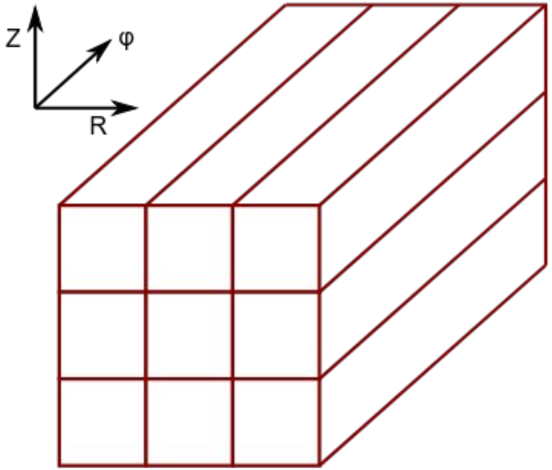

(b) decomposition into RZ-blocks containing either all Fourier modes or a real-space representation in φ

Developed test code to for toy implementation with OpenACC

```
DO iloop=1,50
  !$acc kernels present(ef,be,ve,alpha,integrand) async(1)
  !$acc loop gang independent
  DO ibl=1,nbl
    !$acc loop worker independent
    DO im=1,nmodes
      !$acc loop collapse(2) vector
      DO ie=1,ne
        DO ig=1,ng
          ef(ig,ie,im,1,ibl) = -1.0*(ve(ig,ie,im,2,ibl)*be(ig,ie,im,3,ibl)-ve(ig,ie,im,3,ibl)*be(ig,ie,im,2,ibl))
          ef(ig,ie,im,2,ibl) = -1.0*(ve(ig,ie,im,3,ibl)*be(ig,ie,im,1,ibl)-ve(ig,ie,im,1,ibl)*be(ig,ie,im,3,ibl))
          ef(ig,ie,im,3,ibl) = -1.0*(ve(ig,ie,im,1,ibl)*be(ig,ie,im,2,ibl)-ve(ig,ie,im,2,ibl)*be(ig,ie,im,1,ibl))
        ENDDO
      ENDDO
    ENDDO
```

Toy code allows quick experimentation with data structure and layout

- Start with implementation in global scope (1)
  - All code in single file permits compiler optimization
  - Speedup observed on GPU (2)
- Implementation with full data structures is work in progress (3-5)

	(1) baretest	(2) baretest-acc	(3) base	(4) reorder	(5) acc
CPU cores	3	GPU	3	3	GPU
Par Language	OpenACC	OpenACC	OpenMP	OpenMP	OpenACC
time (s)	0.34	0.06	0.31	0.29	1.33
ordering	ng/ne	ng/ne	nqty	ng/ne	ng/ne

GPU summary

- Early plans and progress for GPU usage in place
- Early results on old hardware (NVIDIA K20c) need to use newer machines
- Unfortunately, we were not accepted for a NESAP award
    - Plan: keep politely asking NERSC for development resources