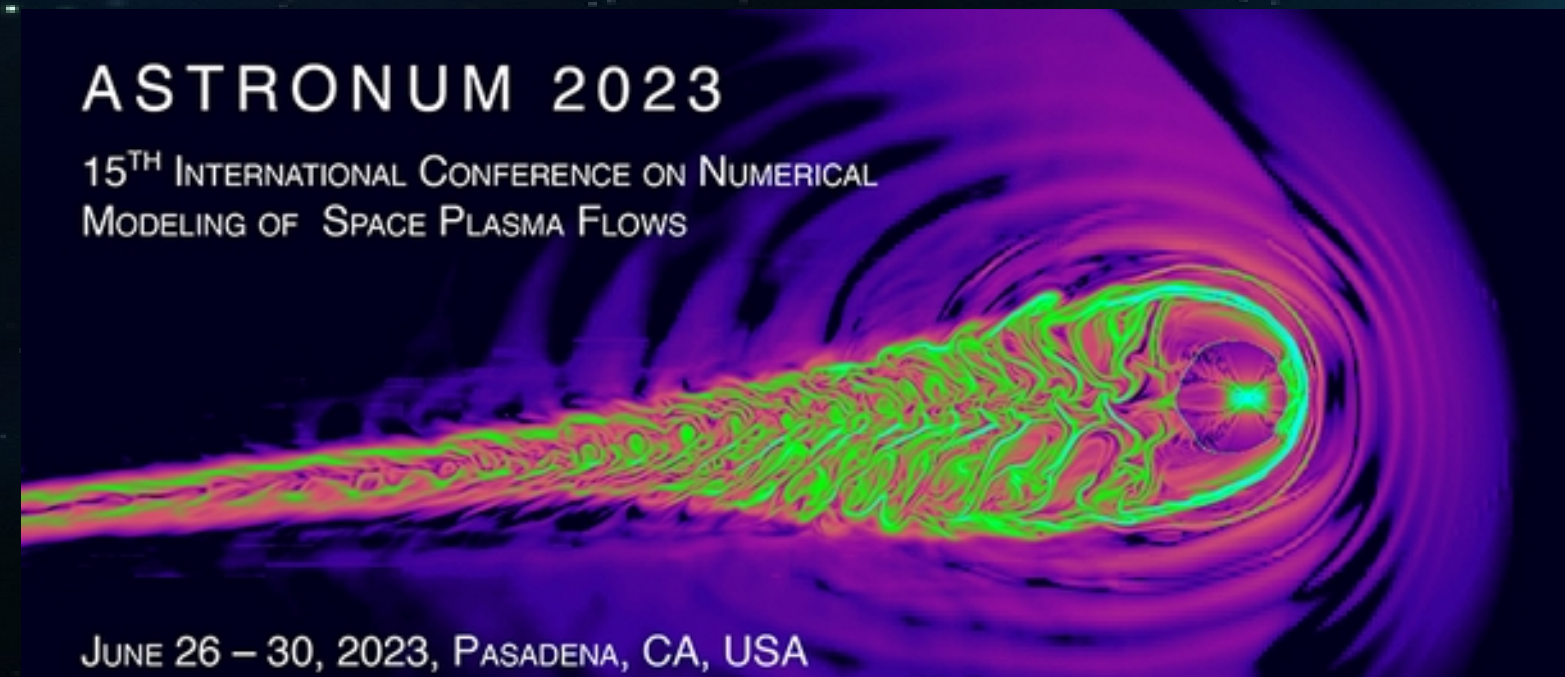
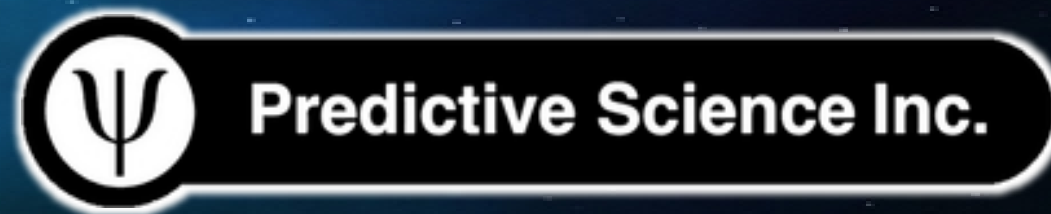
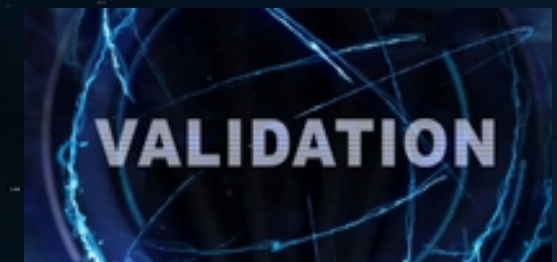


Evaluating a Practical Parabolic Time Step Limit for Unconditionally Stable Schemes in a Thermodynamic MHD Model

Ronald M. Caplan, Craig D. Johnston,
Lars K. S. Daldorff, and Jon A. Linker



- ① Unconditionally stable schemes
- ① The problem of large time steps
- ① Practical time-step limit
- ① MAS thermodynamic MHD model
- ① Test cases
- ① Validation
- ① Performance & scaling
- ① Summary & future development



Unconditionally Stable Schemes

- ⌘ Thermodynamic MHD models (like many others) have multiple time scales leading to vastly different explicit time-step stability requirements
- ⌘ In order to make simulations *tractable*, we need to exceed the most restrictive limits - here, we focus on the parabolic operators
- ⌘ Unconditionally stable time stepping schemes are guaranteed to be stable for any sized time step.
- ⌘ Implicit methods (using iterative matrix solvers)
- ⌘ Explicit methods (e.g. extended stability Runge Kutta)
- ⌘ **When exceeding explicit time-step limits, one must be careful about accuracy**



Unconditionally Stable Schemes

Implicit Backward-Euler (BE) + PCG

- ⊖ Backward Euler $\frac{u^{n+1} - u^n}{\Delta t} = F(u^{n+1})$
- ⊖ Yields a system of equations to solve
- ⊖ To avoid requiring nonlinear solvers, we linearize nonlinear terms (e.g. lagged diffusivity)
- ⊖ We use two non-communicating preconditioners:

PC1

Point-Jacobi
GPU friendly,
scalable, cheap,
not very effective

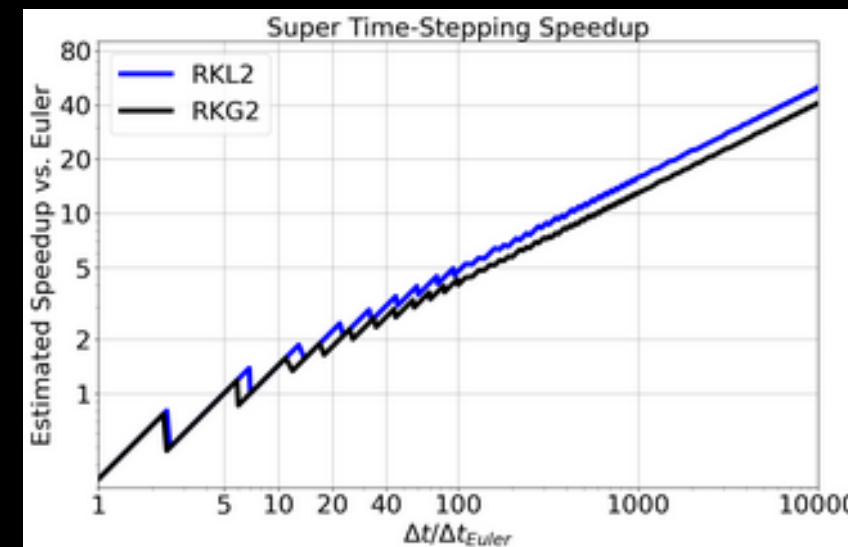
PC2

Non-overlapping domain
decomposition zero-fill
incomplete LU factorization
Less scalable, expensive,
much more effective

Explicit RKG2(3/2) Super Time-stepping

- ⊖ Extended Stability Runge-Kutta schemes: Explicit RK method with stages added for more stability, rather than for more accuracy
- ⊖ We use the 2nd-order Gegenbauer method (RKG2) [O'Sullivan (2019)] with an alpha of 3/2

"RKG allows for accurate modeling of solutions at early times and as such if early times are under investigation RKG is the optimal scheme. However, RKC, RKL, RKU, and RKG are all linearly stable and as such will all approach the correct solution asymptotically at long times." [Skaras et. al. (2021)]



TIP: Need to use odd # of stages in RKG2, otherwise amplification factor goes to 1 at highest mode!

Unconditionally Stable Schemes

Implicit BE + PCG



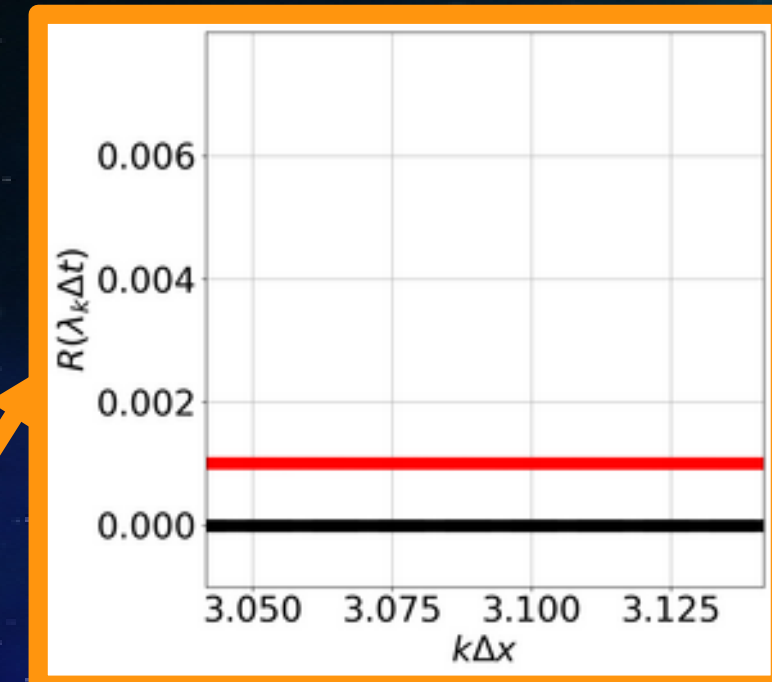
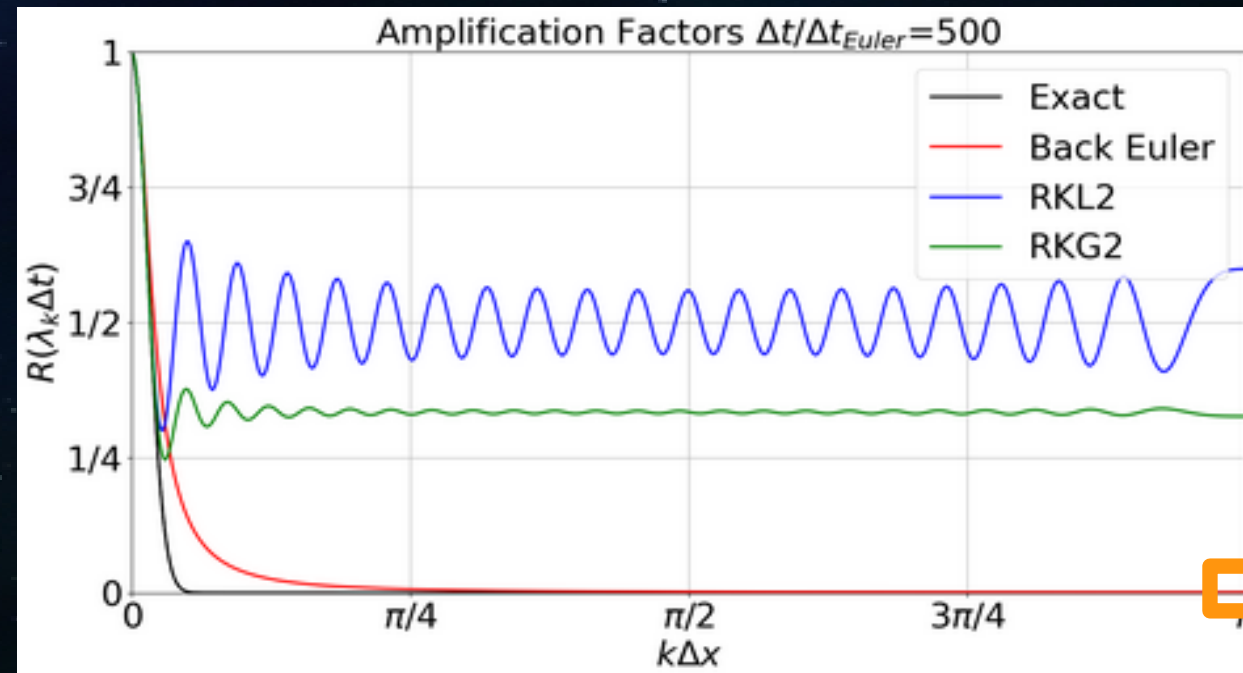
Explicit RKG2(3/2)

- ✓ Robust, proven method
- ✓ Can be very efficient
- ✓ **L-stable**
- ✗ Can be difficult to implement
- ✗ Requires good preconditioner to be efficient, (can be difficult to formulate and implement efficiently)
- ✗ Requires linear(ized) operator
- ✗ Global communication (dot products) hurts scaling
- ✗ Only 1st-order accurate

- ✓ Easy to implement
- ✓ Can include nonlinearities
- ✓ Vectorizable (GPU-friendly)
- ✓ No global synchronization points (better scaling)
- ✓ 2nd-order accurate
- ✗ Not as widely adopted & tested
- ✗ Can be slower than implicit methods
- ✗ **Only A-stable**

⊖ Example L-stable vs. A-stable amplification factors

⊖ A-stable method can have problems damping high wave modes over limited time scales

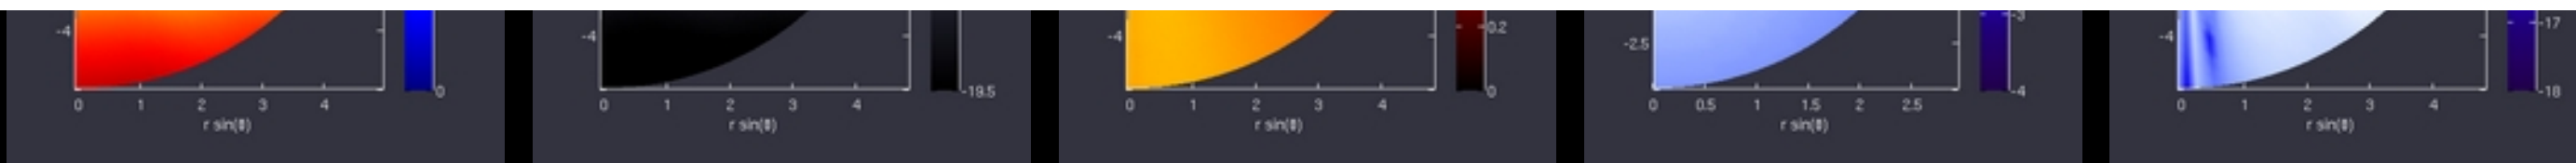
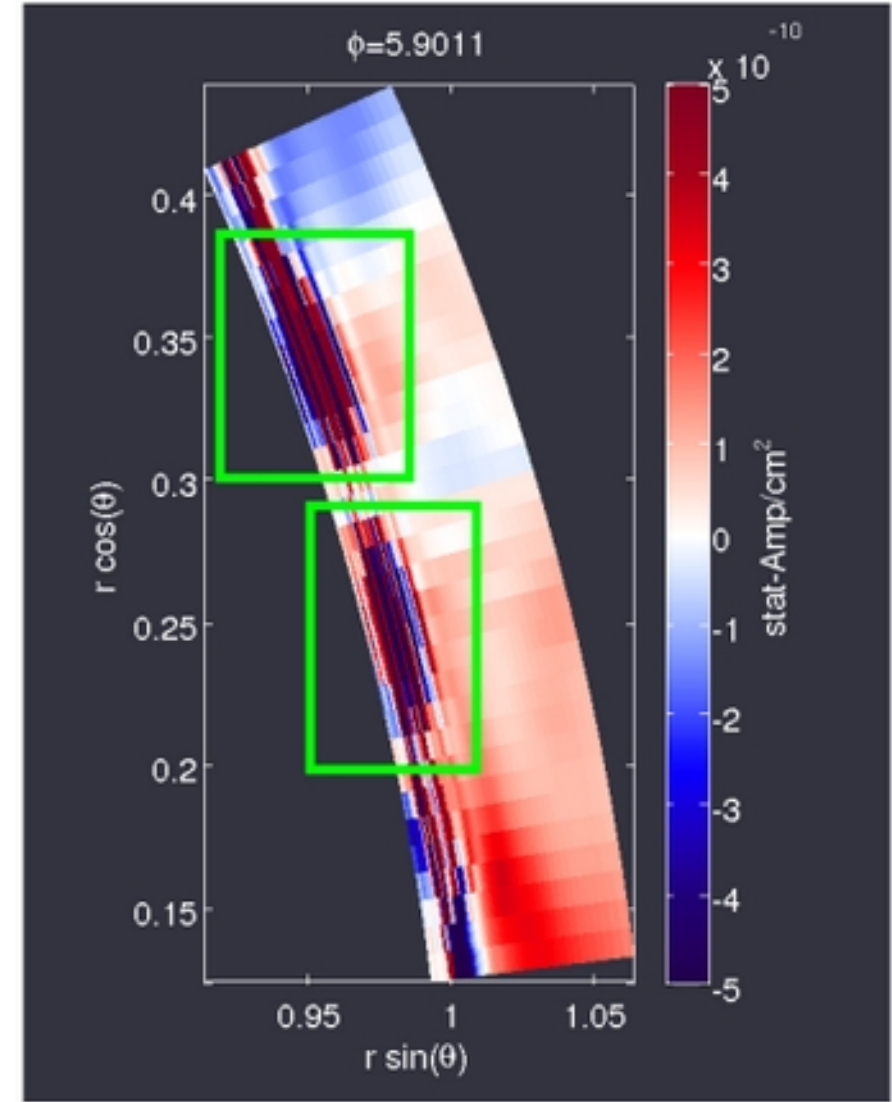
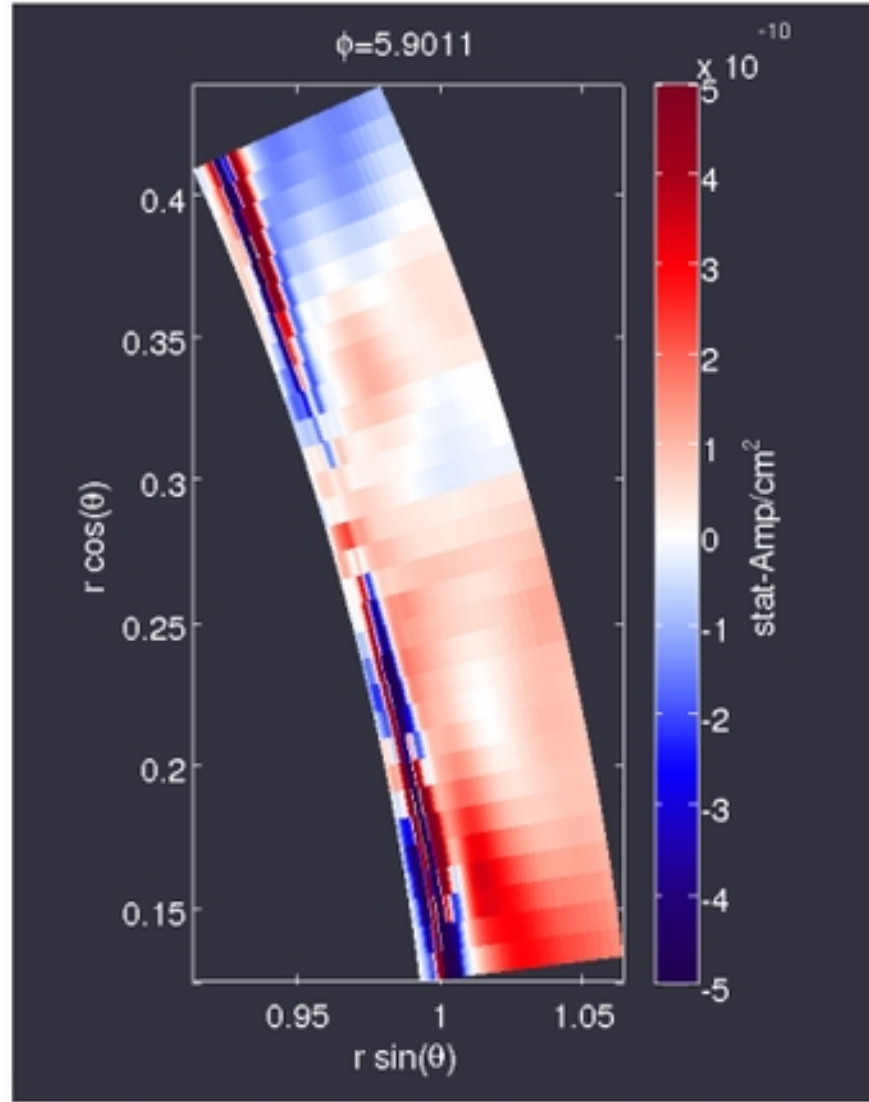
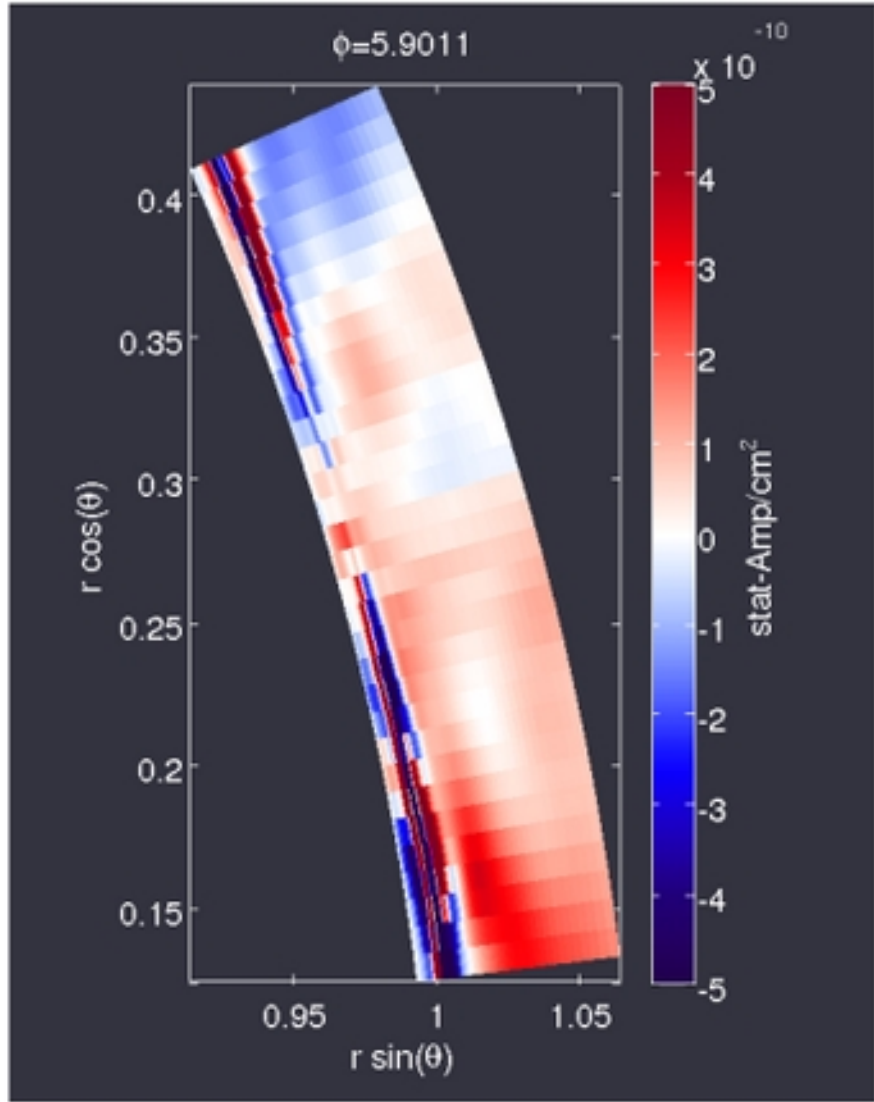


BE+PCG (PC2)

RKL2 (TC-only)

RKL2 (TC+Visc)

J_ϕ



The Problem with Large Time Steps

- Ⓧ Errors with very large time steps are a problem in L-stable methods, but worse in A-stable methods, as they often don't damp high wave modes efficiently
- Ⓧ Extended stability Runge-Kutta schemes fall into this latter category, and this issue can limit their applicability and robustness
- Ⓧ One option is to run the parabolic advance in a series of “outer” cycles (essentially reducing the time step for the operator) which lowers the errors and repeatedly damps high wave modes

“For general applications the universal approach is to try different numbers of steps and study any sensitivities.” [Dawes (2021)]

- Ⓧ **Is there a simple way to a-priory calculate the minimum number of outer cycles we need/want?**
- Ⓧ **How does adding these cycles affect performance?**



VS



A Practical Time Step Limit

Operator split
parabolic advance

$$\frac{\partial \mathbf{u}}{\partial t} = \mathbf{F}(\mathbf{u})$$

Discrete form
(1st expansion term)

$$\frac{u^{n+1} - u^n}{\Delta t} = F(u)$$

$$|u^{n+1} - u^n| = \Delta t |F(u)|$$

Max abs change in
u at grid cell k

$$|F_k(u)| \equiv \max(|F(u_i)|)$$

Bound the relative
change in u at the
location of maximum
absolute change

$$\frac{|u_k^{n+1} - u_k^n|}{|u_k|} < \alpha < 1$$

Practical time step limit: $\Delta t_p = \alpha \frac{|u_k|}{|F(u_k)|}$
 $\alpha = 0.95$

**Applied adaptively:
After each outer-cycle, recalculate!**

Viscosity: $\mathbf{F}_{\text{visc}}(\mathbf{v}) = \frac{1}{\rho} \nabla \cdot (\nu(\mathbf{x}) \rho \nabla \mathbf{v})$

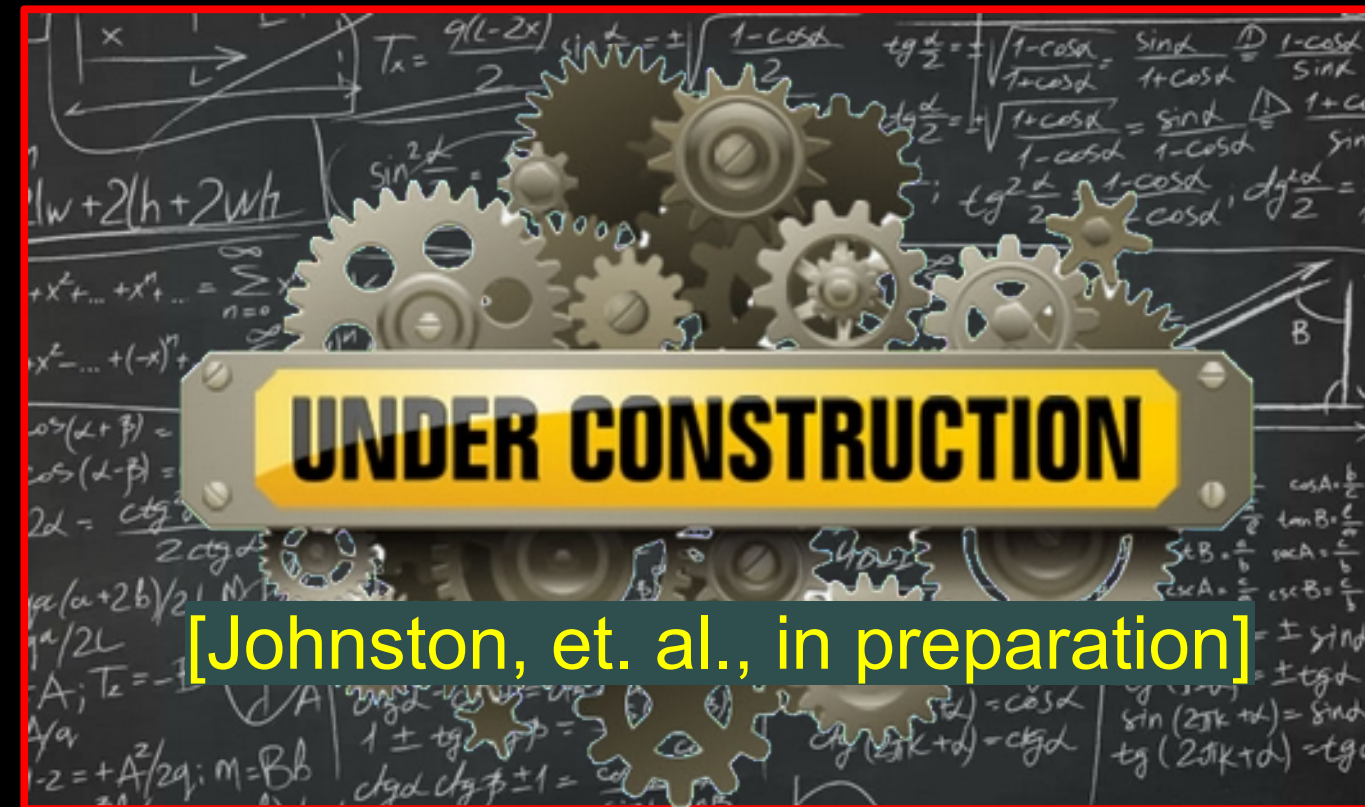
Thermal Conduction:

$$\mathbf{F}_{\text{tc}}(T) = \frac{(\gamma - 1) m_p}{2k} \frac{1}{\rho} \nabla \cdot (\kappa(T_0) \cdot \nabla T)$$

$$\kappa(T) = f_c(r) f_{\text{mod}}(T) \kappa_0 T^{5/2} \hat{\mathbf{b}} \hat{\mathbf{b}}$$

$$f_{\text{mod}}(T) = (1 + (T/T_{\text{cut}})^{-10})^{1/4}$$

Lagged diffusivity



MAS

MAGNETOHYDRODYNAMIC
ALGORITHM
OUTSIDE A SPHERE

Purpose:

General-purpose simulations of the corona and heliosphere for use with solar physics and space weather research

Model:

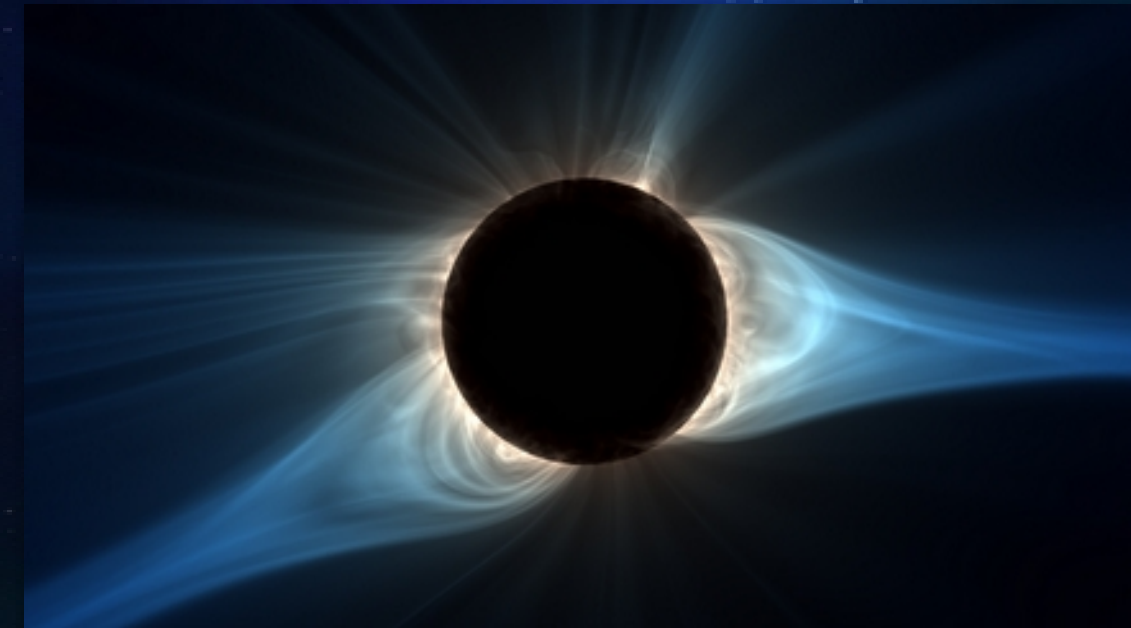
Spherical resistive thermodynamic MHD

Algorithms:

Implicit & explicit time-stepping with finite-difference stencils. Implicit steps use a sparse matrix preconditioned iterative solver

Code:

~70,000 lines of Fortran
parallelized with
MPI + OpenACC + StdPar



FREE WEBINAR

**Accelerating a Production
Solar MHD Code with
Fortran Standard Parallelism**

Ronald M. Caplan
Predictive Science Inc.

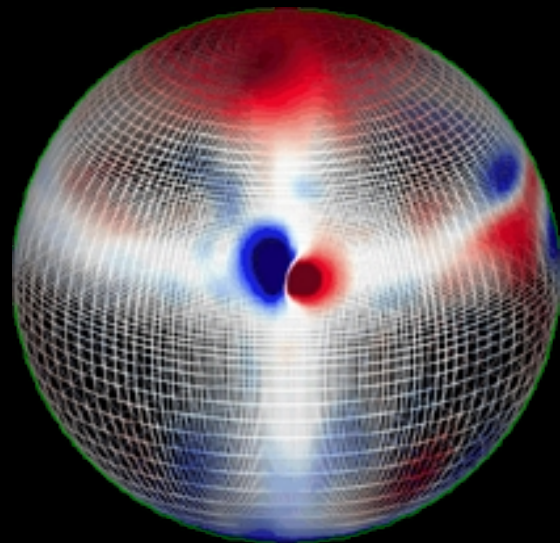
OpenACC

WEBINAR
DIGITAL EVENT
JULY 11, 2023
1 PM EDT/10 AM PDT

The MAS MHD Model

MAS

MAGNETOHYDRODYNAMIC ALGORITHM OUTSIDE & INSIDE



(r_i, θ_j, ϕ_k)

$$\frac{\partial \mathbf{A}}{\partial t} = \mathbf{v} \times (\nabla \times \mathbf{A}) - \frac{c^2 \eta}{4\pi} \nabla \times \nabla \times \mathbf{A}$$

RESISTIVITY

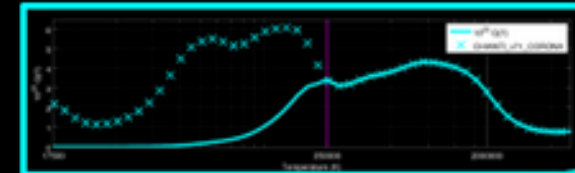
$$\frac{\partial \rho}{\partial t} = -\nabla \cdot (\rho \mathbf{v})$$

$$\frac{\partial T}{\partial t} = -\nabla \cdot (T \mathbf{v}) - (\gamma - 2) (T \nabla \cdot \mathbf{v}) + \frac{\gamma - 1}{2k} \frac{m_p}{\rho} \left[-\nabla \cdot (\mathbf{q}_1 + \mathbf{q}_2) - \frac{\rho^2}{m_p^2} Q(T) + H \right]$$

THERMAL CONDUCTION

$$\mathbf{q}_1 = -f(r) \beta_{\text{tot}}(T) \kappa_0 T^{5/2} \hat{\mathbf{b}} \hat{\mathbf{b}} \cdot \nabla T$$

$$\mathbf{q}_2 = (1 - f(r)) \frac{k}{(\gamma - 1)} \frac{\rho}{m_p} T \mathbf{v} \hat{\mathbf{b}} \hat{\mathbf{b}}$$



RADIATIVE COOLING

$$\frac{\rho^2}{m_p^2} Q(T) + H$$

CORONAL HEATING

$$H = H^* + \frac{\rho}{4\lambda_{\perp}} [|z_-| z_+^2 + |z_+| z_-^2]$$

$$\lambda_{\perp} = \lambda_0 \sqrt{\frac{B_w}{|\mathbf{B}|}} |z_{\pm}(r = R_{\odot})| = z_0$$

ALFVEN

$$\frac{\partial \epsilon_{\pm}}{\partial t} = -\nabla \cdot (\epsilon_{\pm} [\mathbf{v} \pm \mathbf{v}_A]) - \frac{\epsilon_{\pm}}{2} \nabla \cdot \mathbf{v}$$

$$\frac{\partial \mathbf{v}}{\partial t} = -\mathbf{v} \cdot \nabla \mathbf{v} + \frac{1}{\rho} \left[\frac{1}{c} \mathbf{J} \times \mathbf{B} - \nabla p - \nabla \left(\frac{\epsilon_+ + \epsilon_-}{2} \right) + \rho \mathbf{g} \right] + \frac{1}{\rho} \nabla \cdot (\nu \rho \nabla \mathbf{v}) + \nabla \cdot \left(S \rho \nabla \frac{\partial \mathbf{v}}{\partial t} \right)$$

VISCOSITY

SEMI-IMPLICIT OPERATOR

WAVE TURBULENCE

$$\frac{\partial z_{\pm}}{\partial t} = -(\mathbf{v} \pm \mathbf{v}_A) \cdot \nabla z_{\pm} - \frac{z_{\pm} |z_{\mp}|}{2\lambda_{\perp}} + \frac{z_{\pm}}{4} (\mathbf{v} \mp \mathbf{v}_A) \cdot \nabla (\ln \rho) + \frac{z_{\mp}}{2} (\mathbf{v} \mp \mathbf{v}_A) \cdot \nabla (\ln |\mathbf{v}_A|)$$

$\nabla \cdot \mathbf{B} = 0$	$p = 2kT\rho/m_p$	$\hat{\mathbf{b}} = \mathbf{B}/ \mathbf{B} $	$\beta_{\text{tot}}(T) = \begin{cases} (T/T_c)^{-1/2} & T < T_c \\ 1 & T \geq T_c \end{cases}$	$S = (\Delta t^2 k)^{-1} (C_p^2/(1-C_p)^2 - 1)$
$\mathbf{B} = \nabla \times \mathbf{A}$	$\mathbf{g} = -g_0 R_{\odot}^2 \hat{\mathbf{r}}/r^2$	$\mathbf{v}_A = \mathbf{B}/\sqrt{4\pi\rho}$	$T_c = 3.5 \times 10^6 \text{ K}$	$C_p = \Delta t k \cdot \nu$
$\mathbf{J} = \frac{c}{4\pi} \nabla \times \mathbf{B}$	$\gamma = 5/3$	$B_w = 6.09 \text{ G}$	$f(r) = 1 - 0.5 \tanh[(r - 10 R_{\odot})/R_{\odot}]$	$C_p^2 = 0.25 \Delta t^2 k^2 (\nu^2 + \mathbf{v}_A ^2)$
		$\nu^2 = \gamma p/\rho$		$\mathbf{P} = 4 (\Delta r^{-2} + (r \Delta \theta)^{-2} + (r \Delta \phi \sin \theta)^{-2})$

CORONA

HELIOSPHERE

EARTH



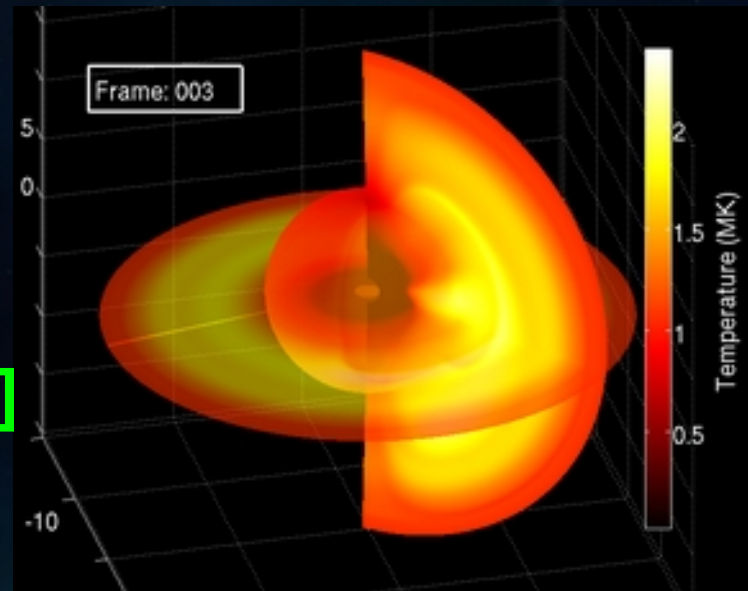
$r = 30 R_{\odot}$

Different components of the model are used depending on the domain and use-case

Test Cases

Test 1:

Modified low-res test case based on simulations used in [Reeves et. al. (2019)]



Resolution:
151x151x151
~ 3.4 million points

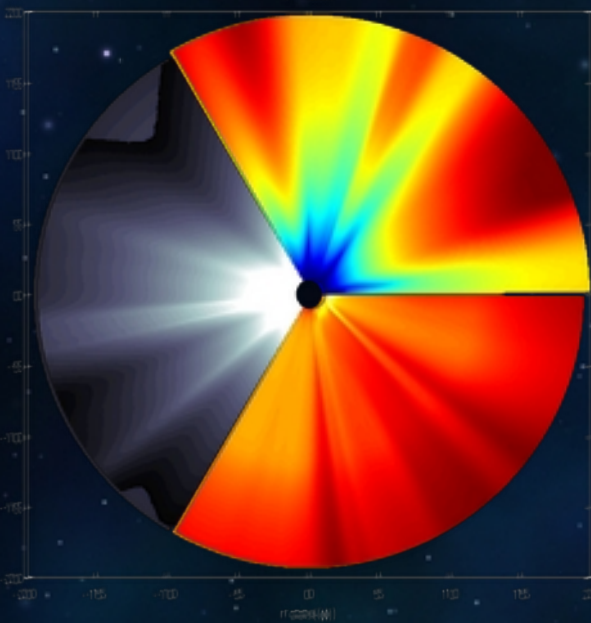
Thermodynamic
MHD relaxation

Test 2:

Modified test case used in Astronom 2016 paper [Caplan et. al. (2017)]



Resolution:
181x251x502
~ 22.8 million points



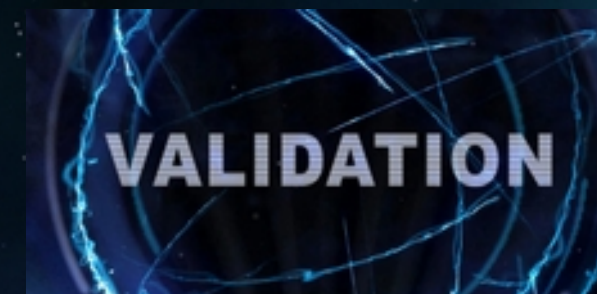
Thermodynamic
MHD relaxation

Validation runs:

Integrate relaxation for ~8 simulation-hours

Scaling runs (Test 2 only):

Integrate for 6 simulation-minutes starting with the ~8 simulation-hour relaxation (subtract restart loading time)



Questions to address:

- ❑ Can we get a solution as good (or better) as BE+PCG with RKG2 if we automatically outer-cycle at the practical time step limit?
- ❑ Can we get a “better” solution with BE+PCG if we outer-cycle it?
- ❑ How does outer-cycling affect performance?
- ❑ Is RKG2 competitive with BE+PCG?

Tests:

SC1: 1 outer cycle (original)

SCA: Automatic adaptive
outer cycles
(practical time step)

BE+PCG	RKG2
SC1	SC1
SCA	SCA

Computational Environment

- Tests performed on both CPUs and GPUs
- On CPUs, the PC2 preconditioner is used for BE+PCG; on GPUs, PC1 is used
- On CPUs, the gfortran compiler is used with OpenMPI 4, on GPUs, the NV compiler is used with OpenMPI 3

In-house workstation with NVIDIA RTX 3090 Ti



CPU	Core i5-13600KF
GPU	RTX 3090 Ti
Peak DP FLOPs	0.625 TFLOP/s
Memory	24 GB
Memory Bandwidth	1008 GB/s

EPYC ROME 7742

SDSC Expanse 2xCPU Node

# CPUs x Model	(2x) EPYC 7742
# Total Cores	128 (we use 64)
Peak FLOP/s	7.0 TFLOP/s
Memory	256 GB
Total Memory Bandwidth	381.4 GB/s

A100 40GB

NCSA Delta 8xGPU Node

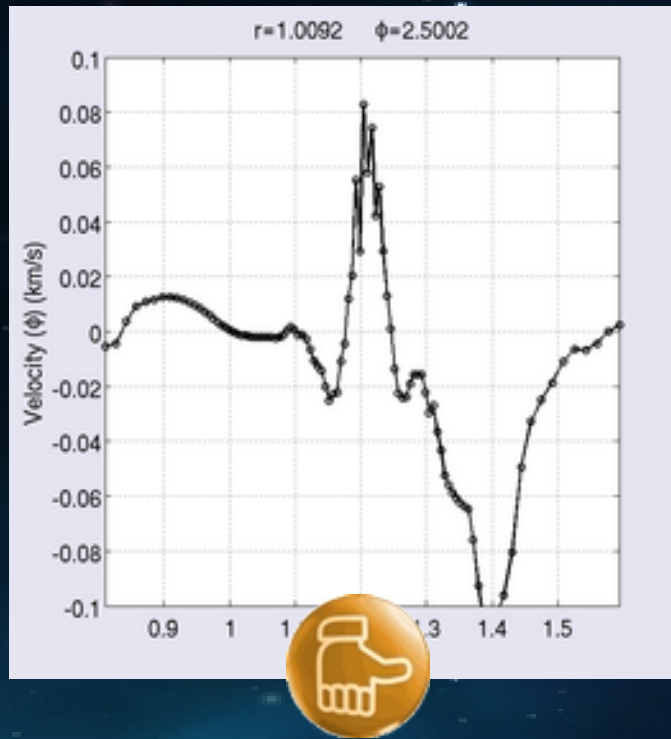
# CPUs x Model	(2x) EPYC 7742
# GPUs x Model	8x A100-40GB SXM4
Peak DP FLOP/s / GPU	9.8 TFLOP/s
Memory / GPU	40 GB
Memory Bandwidth/GPU	1555 GB/s

Allocations provided by:

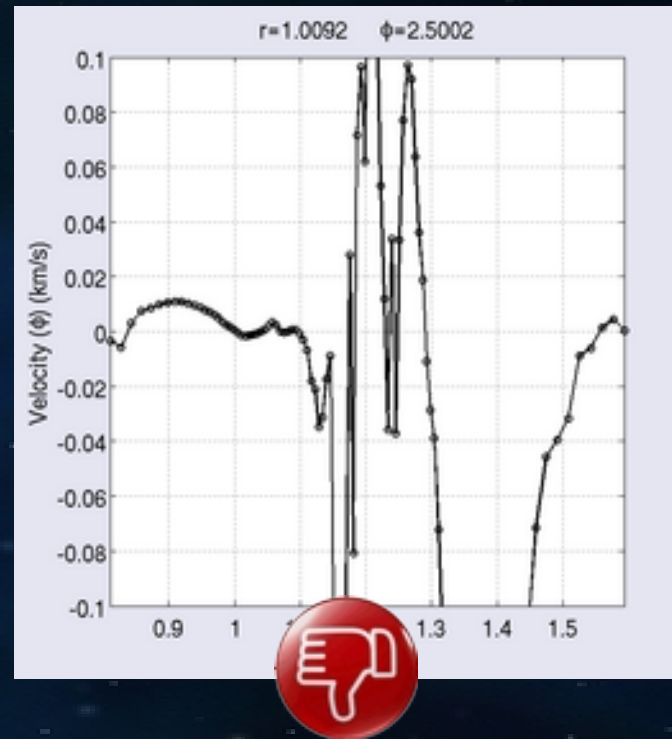


Validation Results Test 1

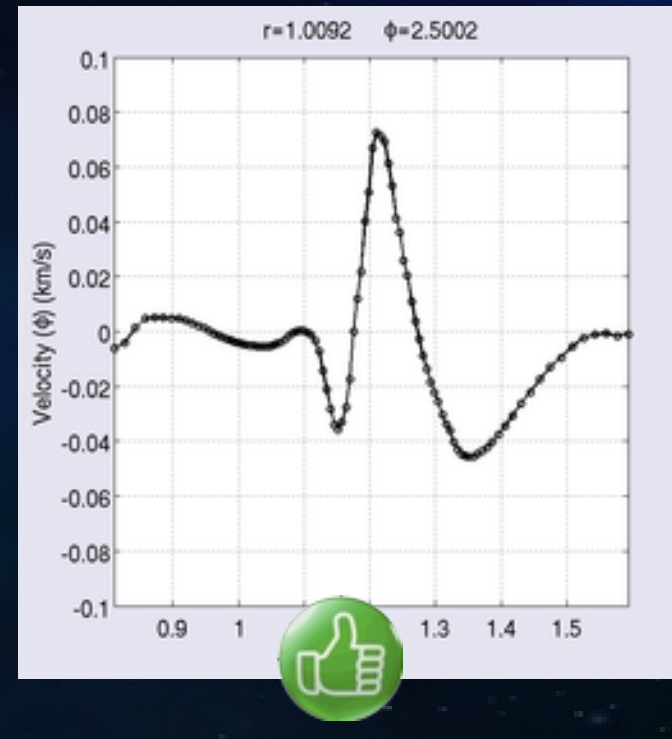
BE+PCG SC1



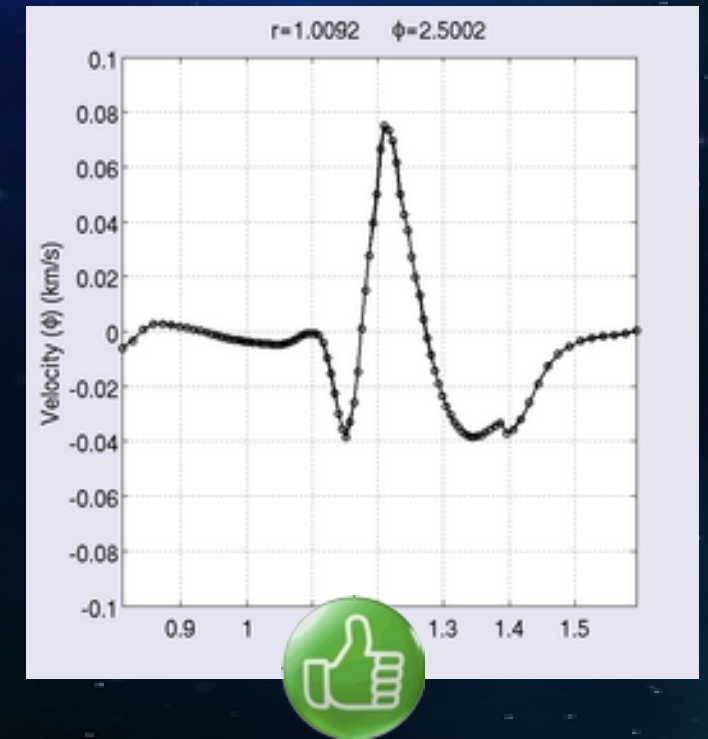
RKG2 SC1



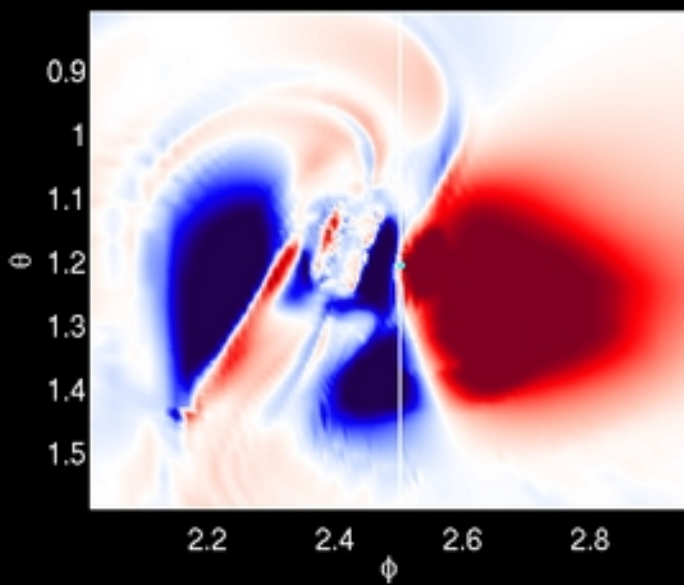
BE+PCG SCA



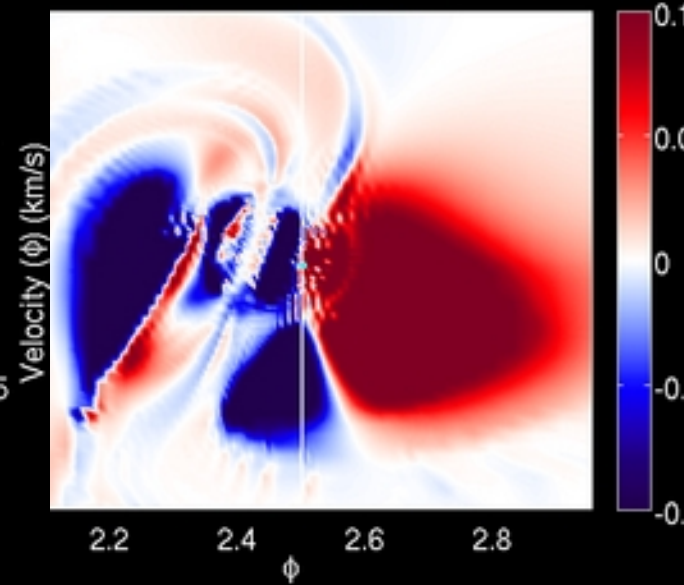
RKG2 SCA



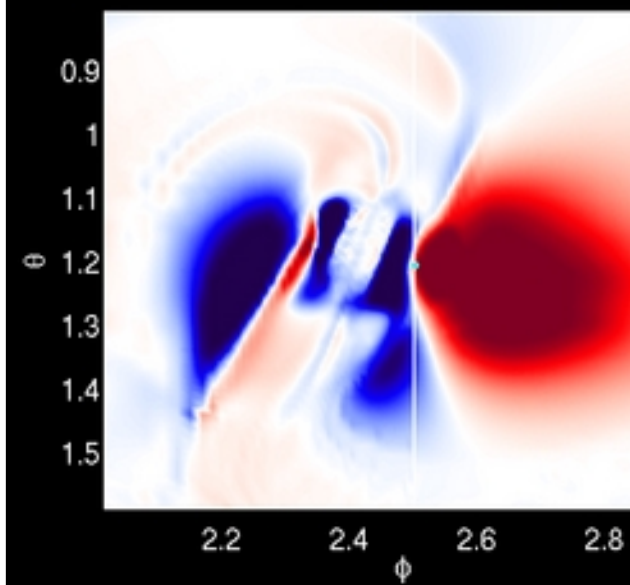
$r=1.0002$



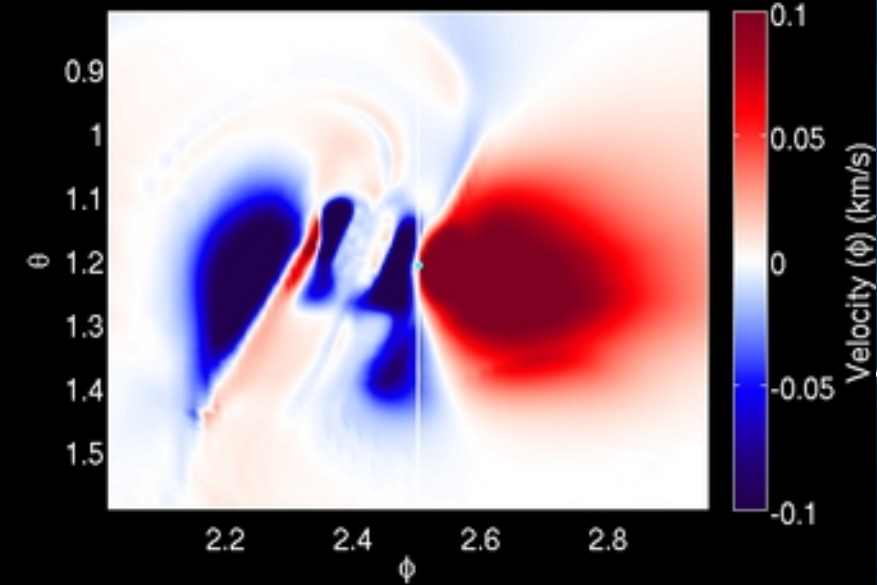
$r=1.0002$



$r=1.0002$

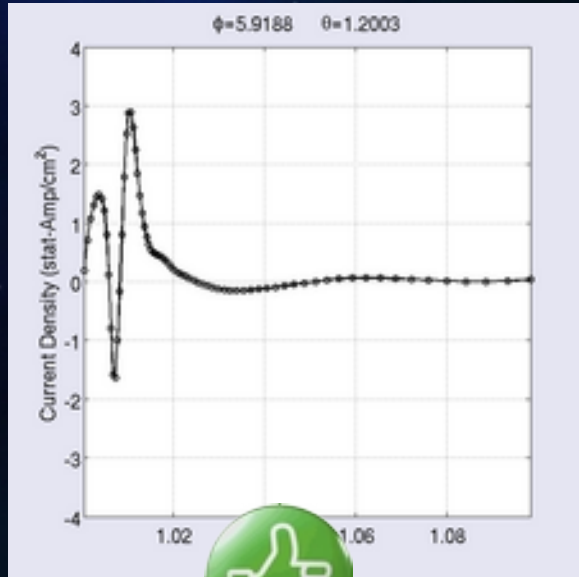


$r=1.0002$

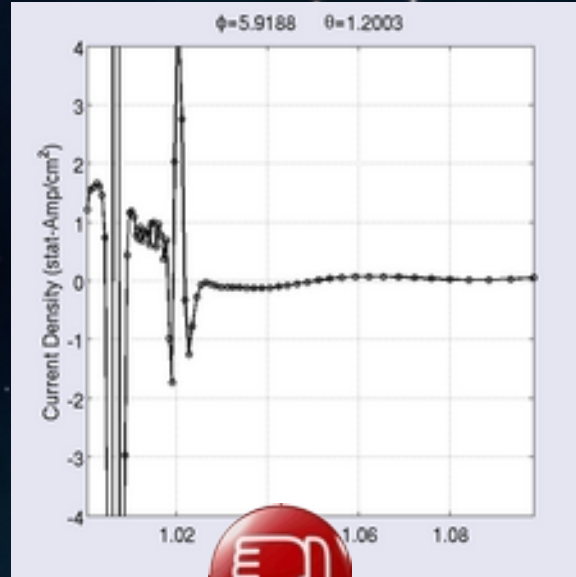


Validation Results Test 2

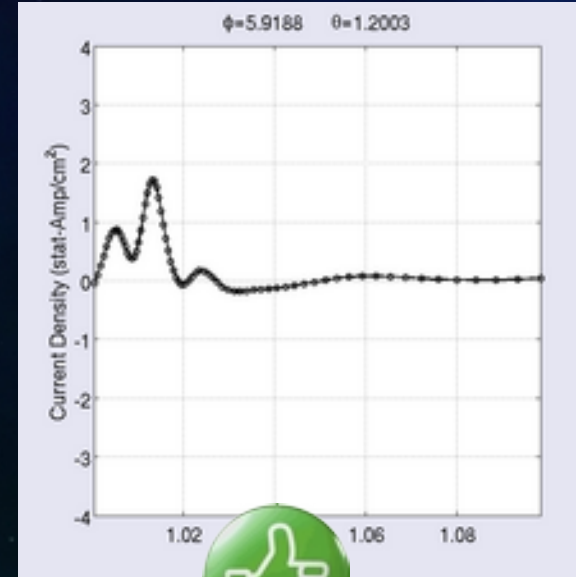
BE+PCG SC1



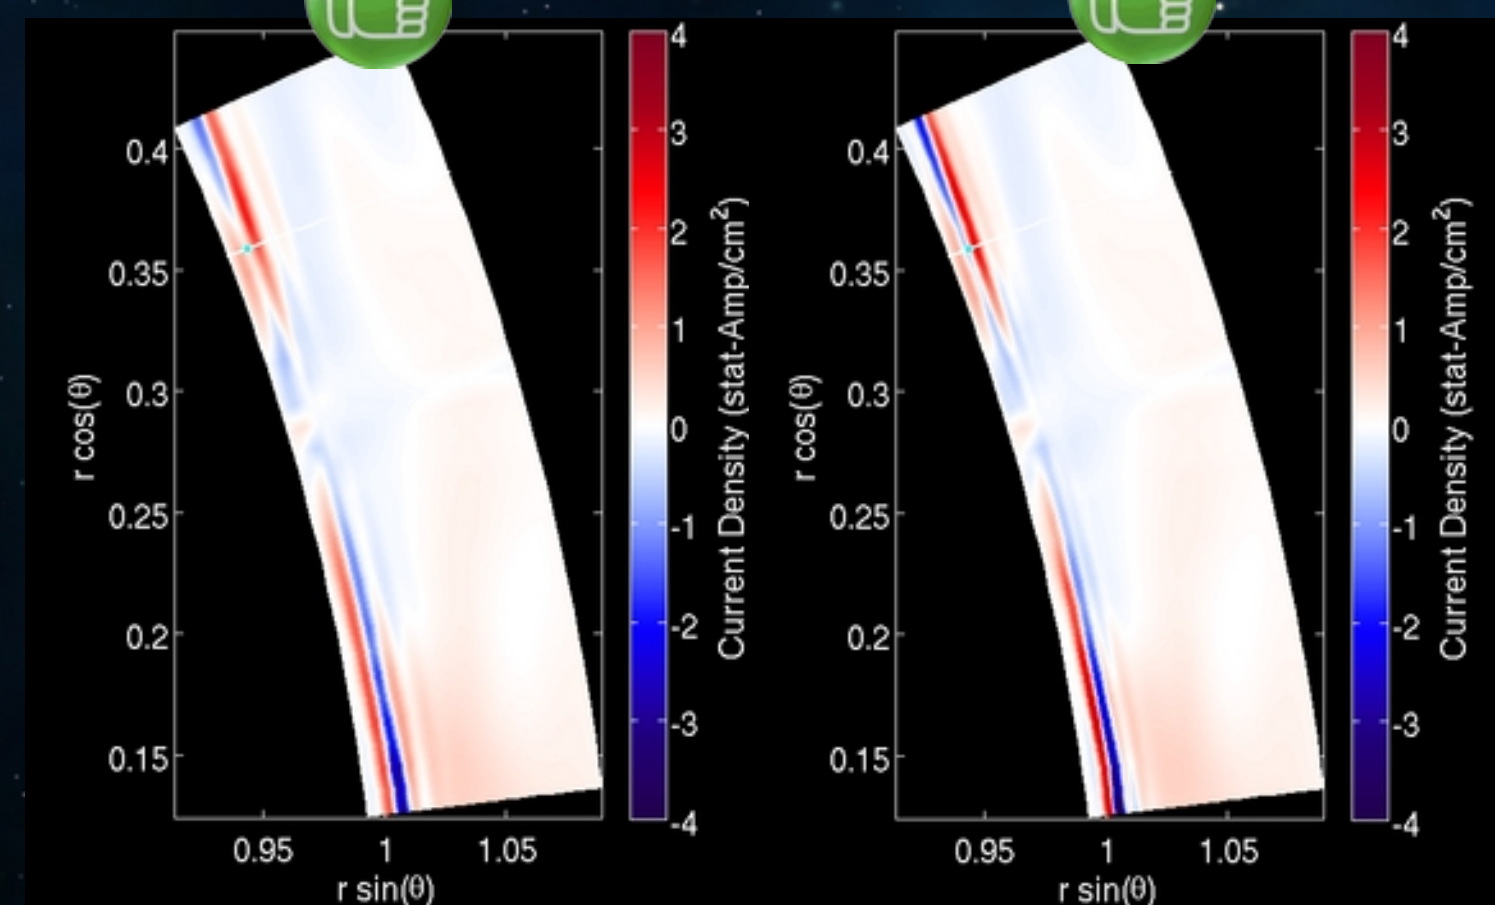
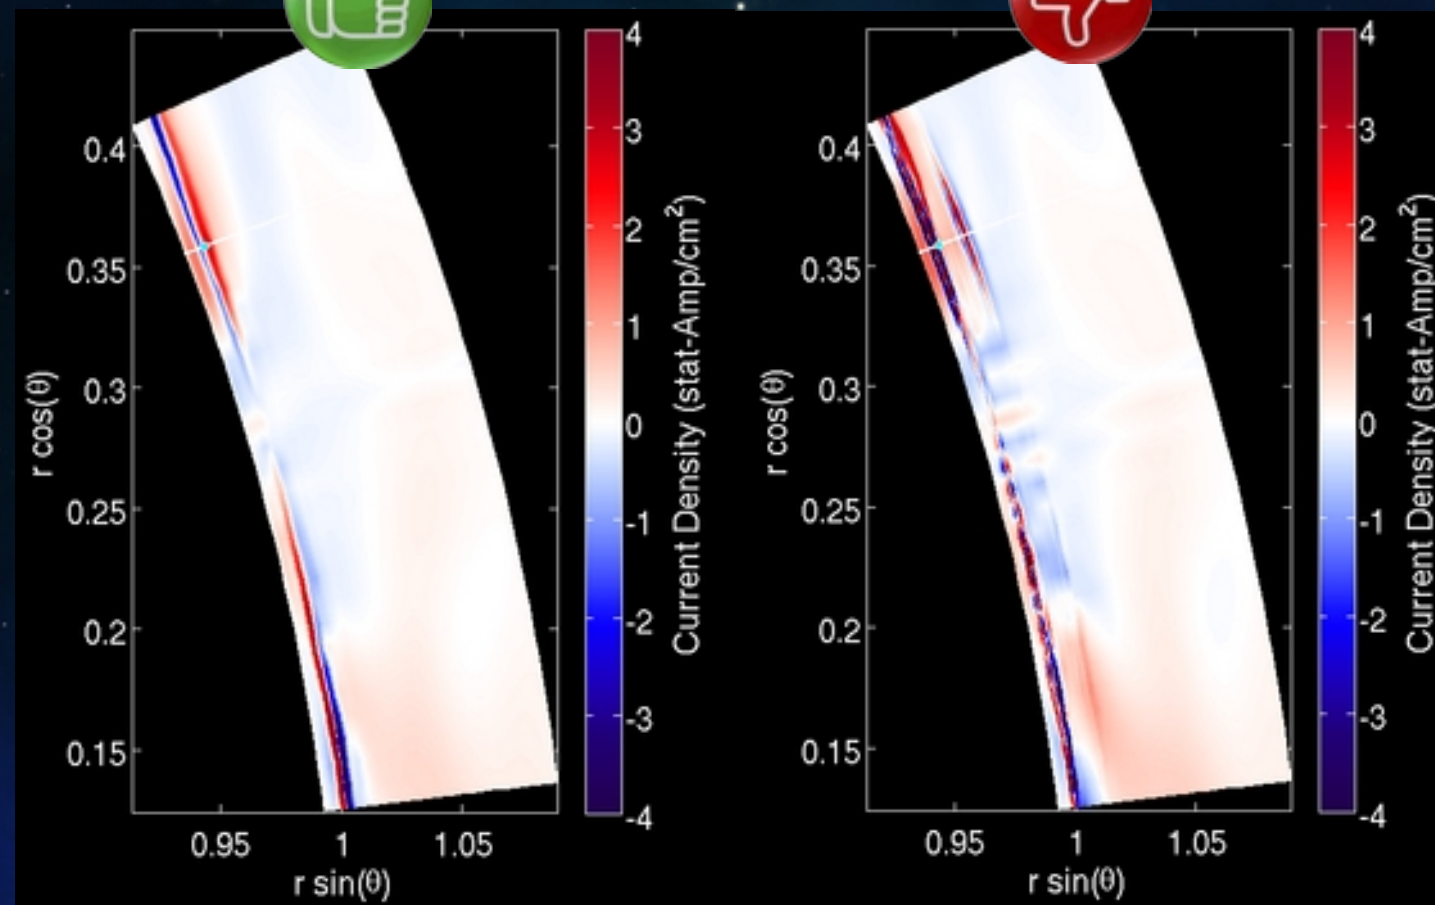
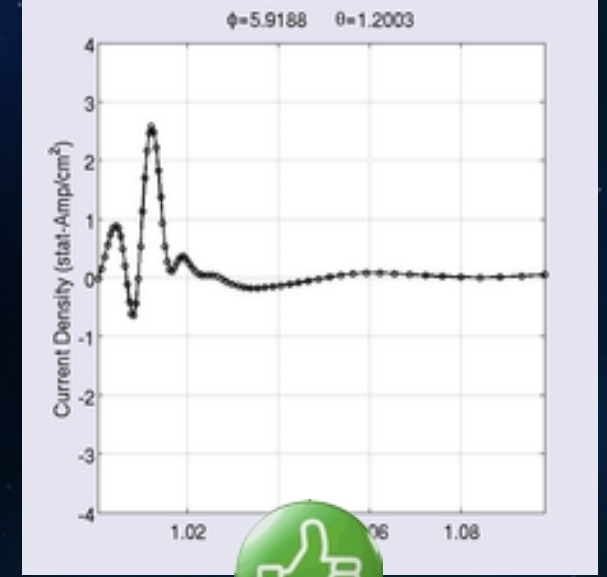
RKG2 SC1



BE+PCG SCA



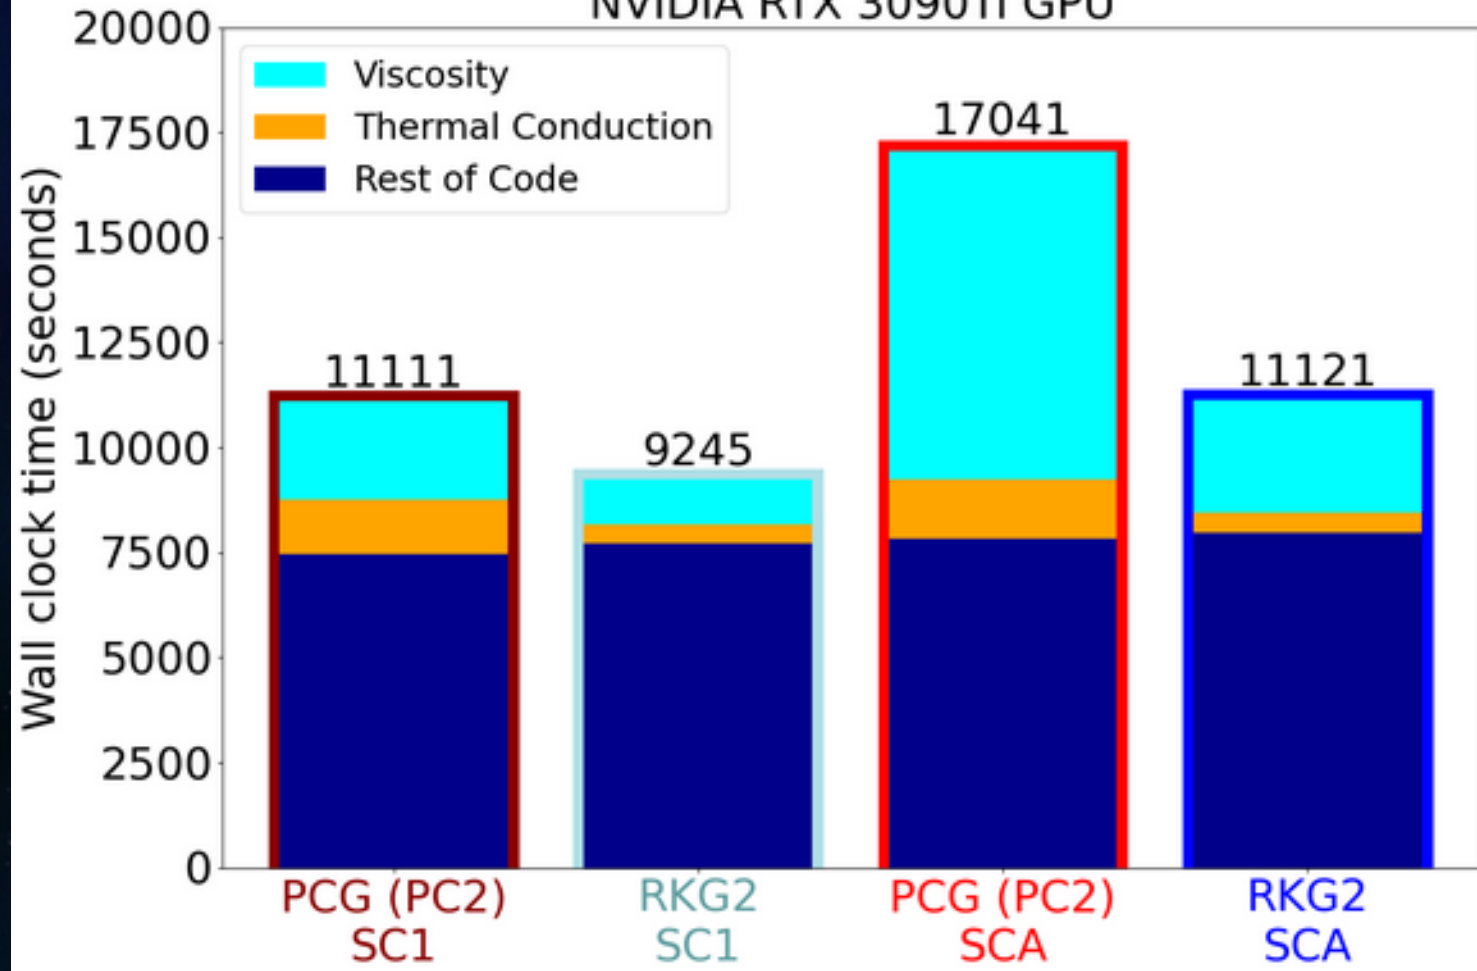
RKG2 SCA



Validation Timing Results

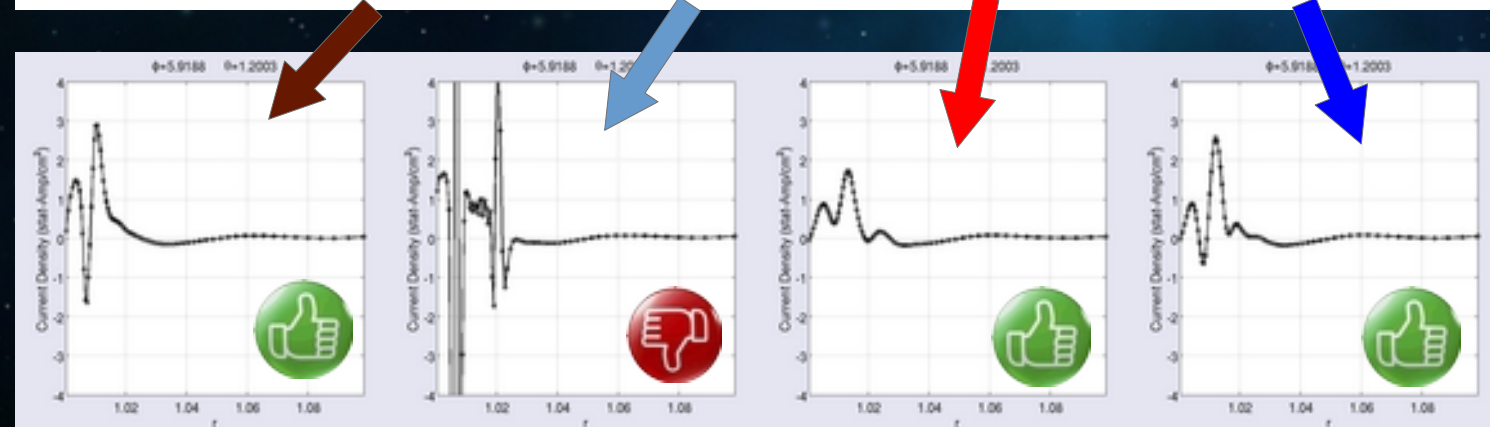
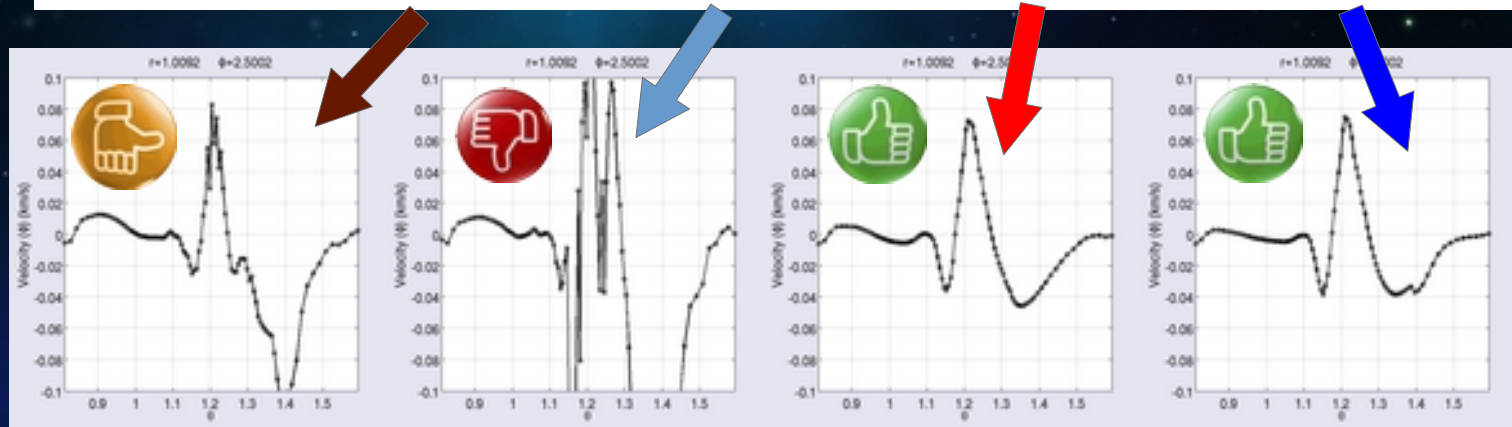
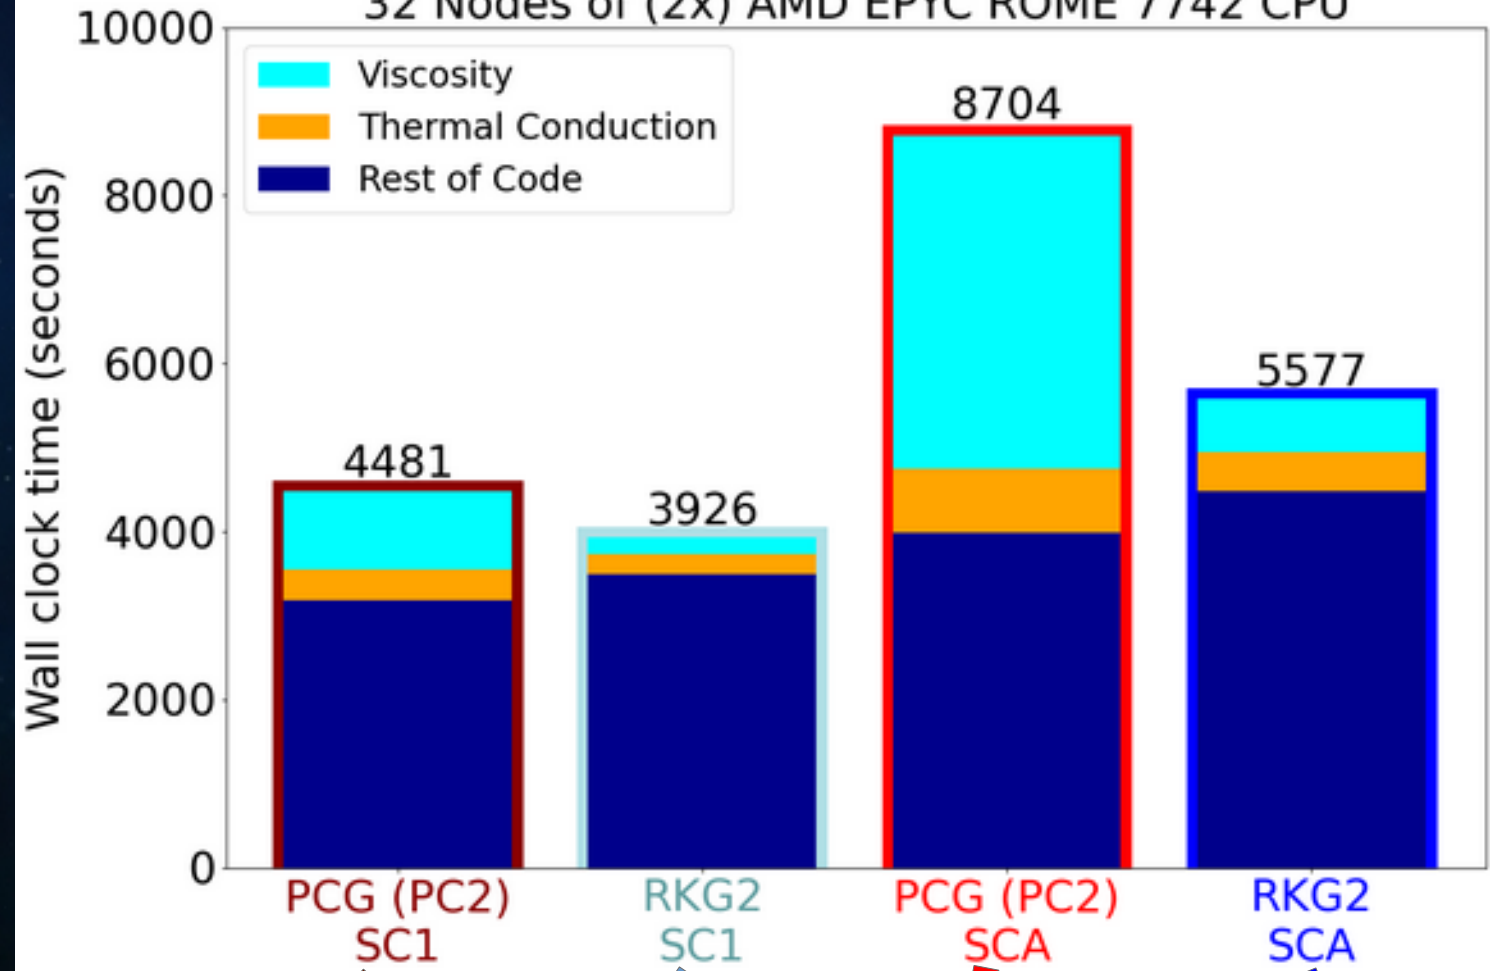
Test 1:

MAS Thermodynamic Relaxation (Test 1)
NVIDIA RTX 3090Ti GPU



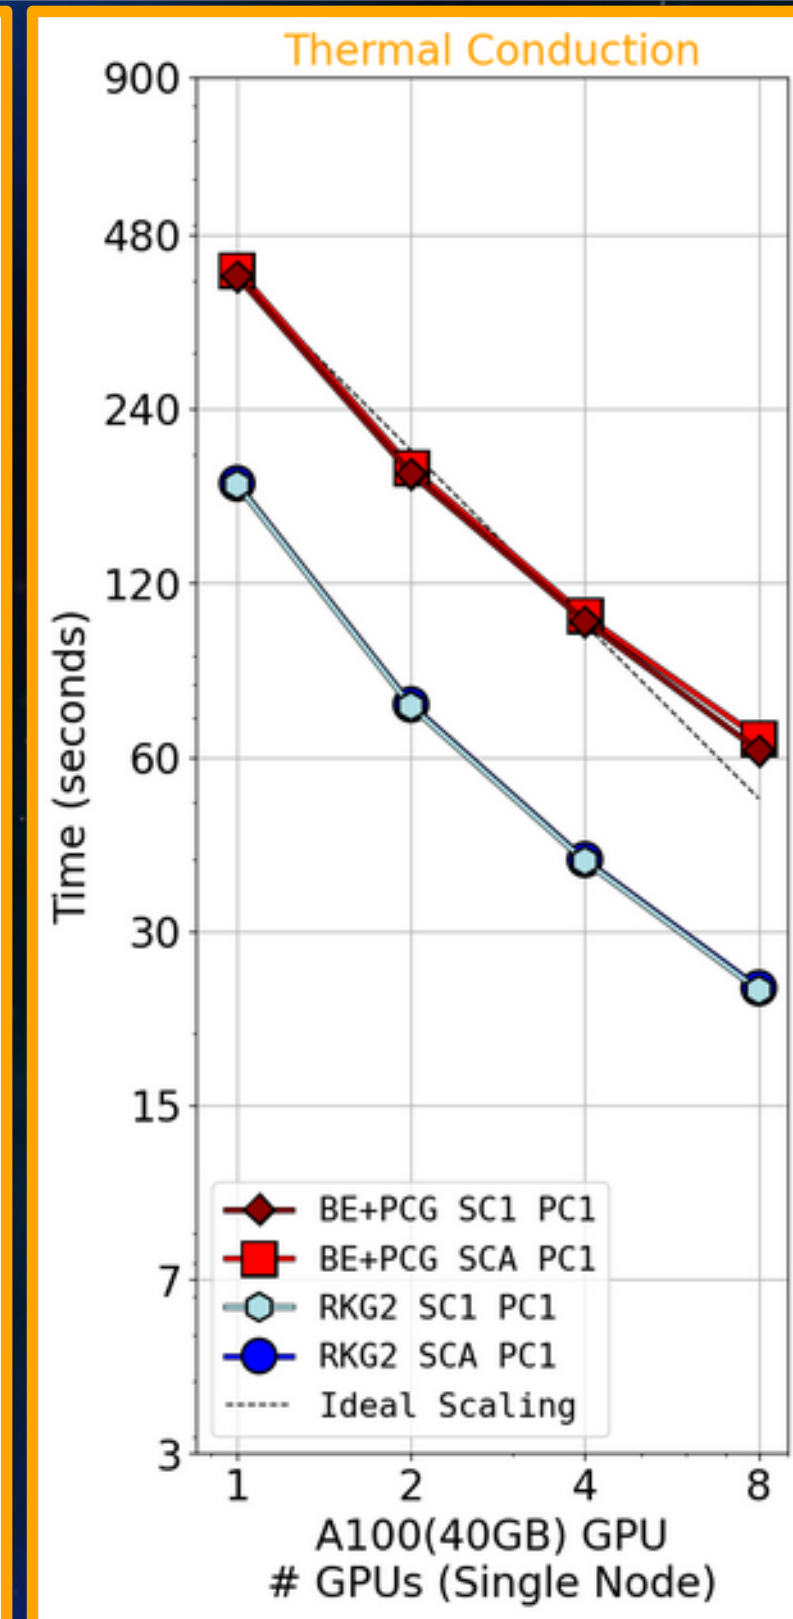
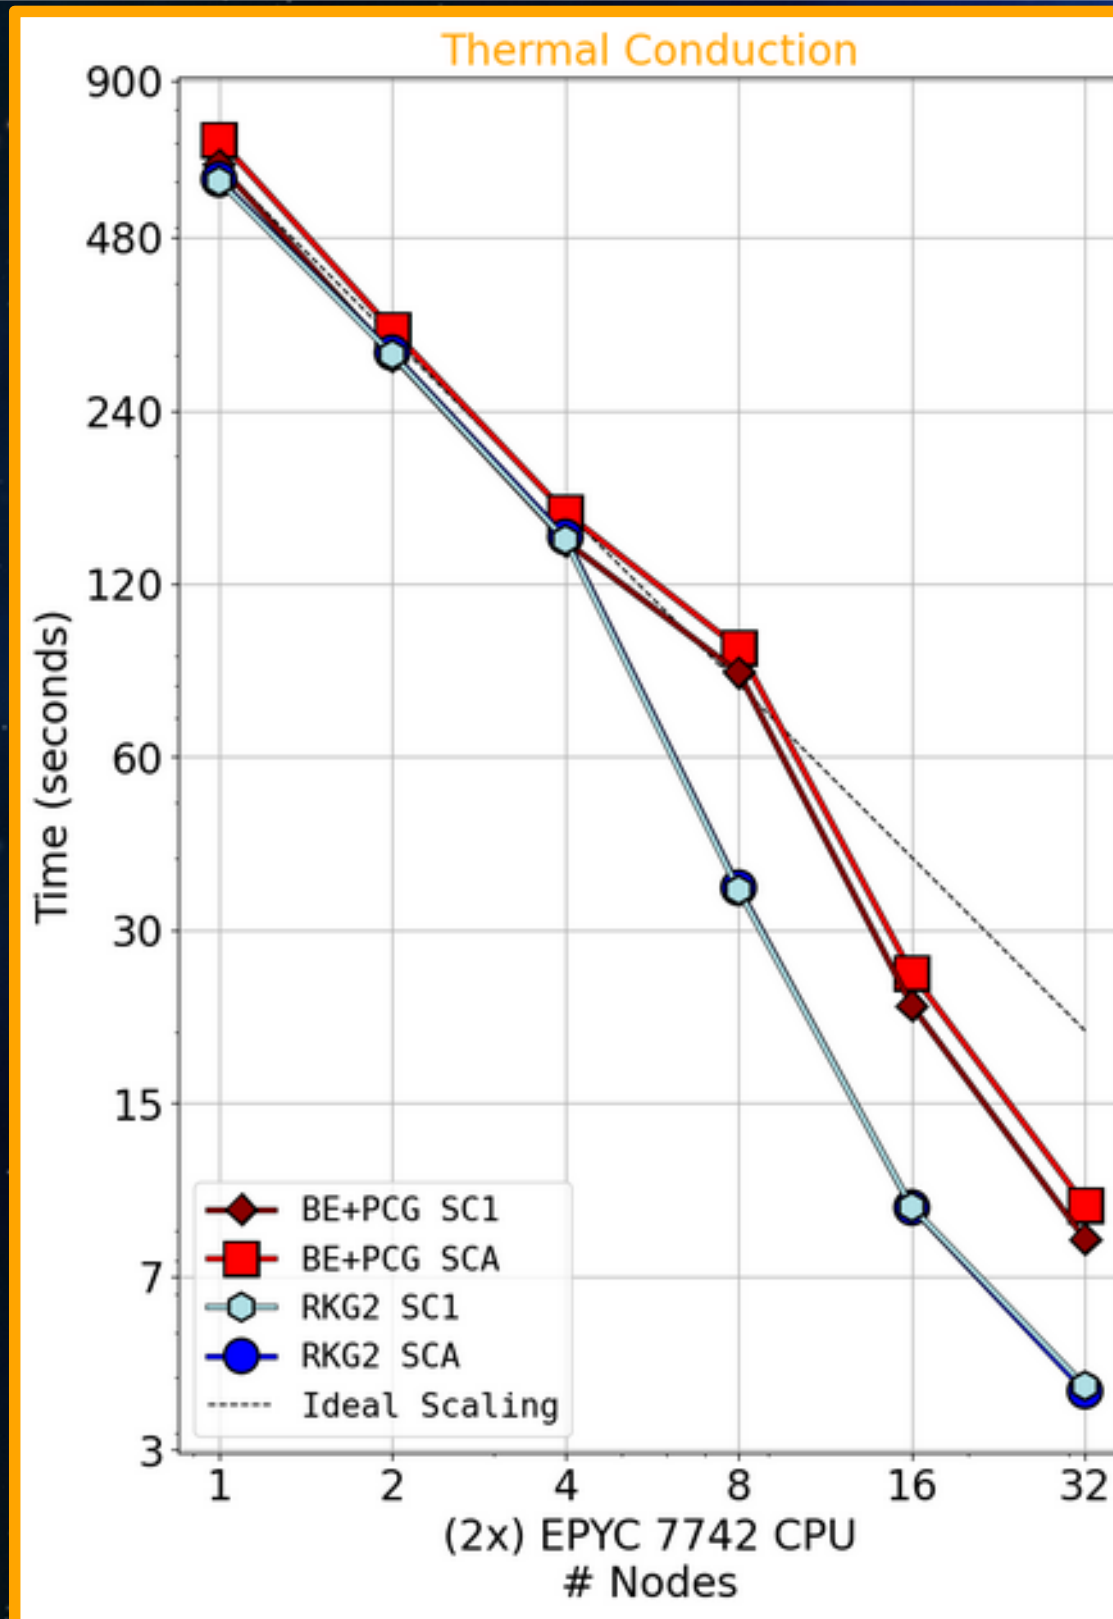
Test 2:

MAS Thermodynamic Relaxation (Test 2)
32 Nodes of (2x) AMD EPYC ROME 7742 CPU



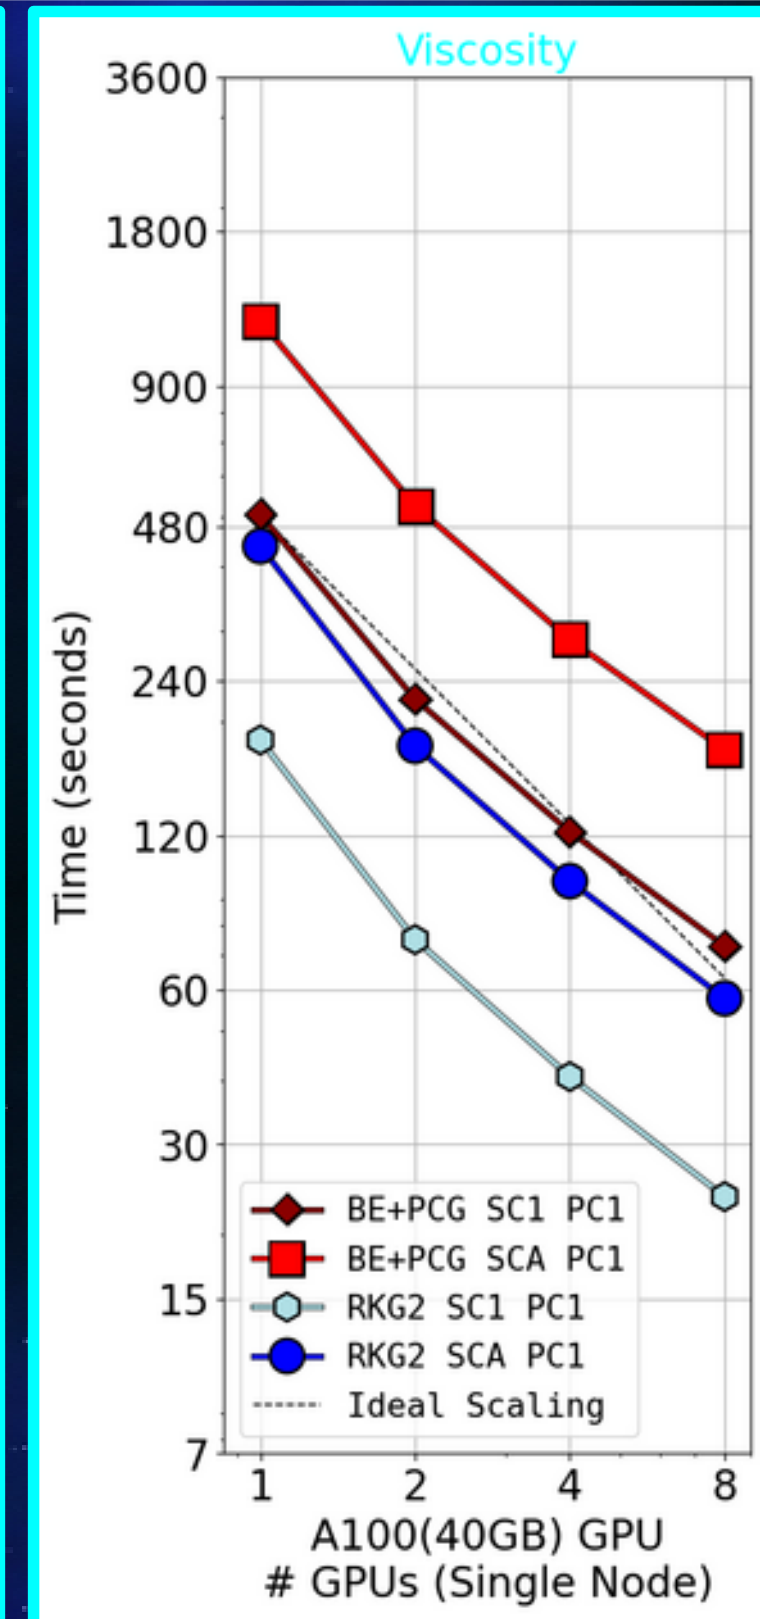
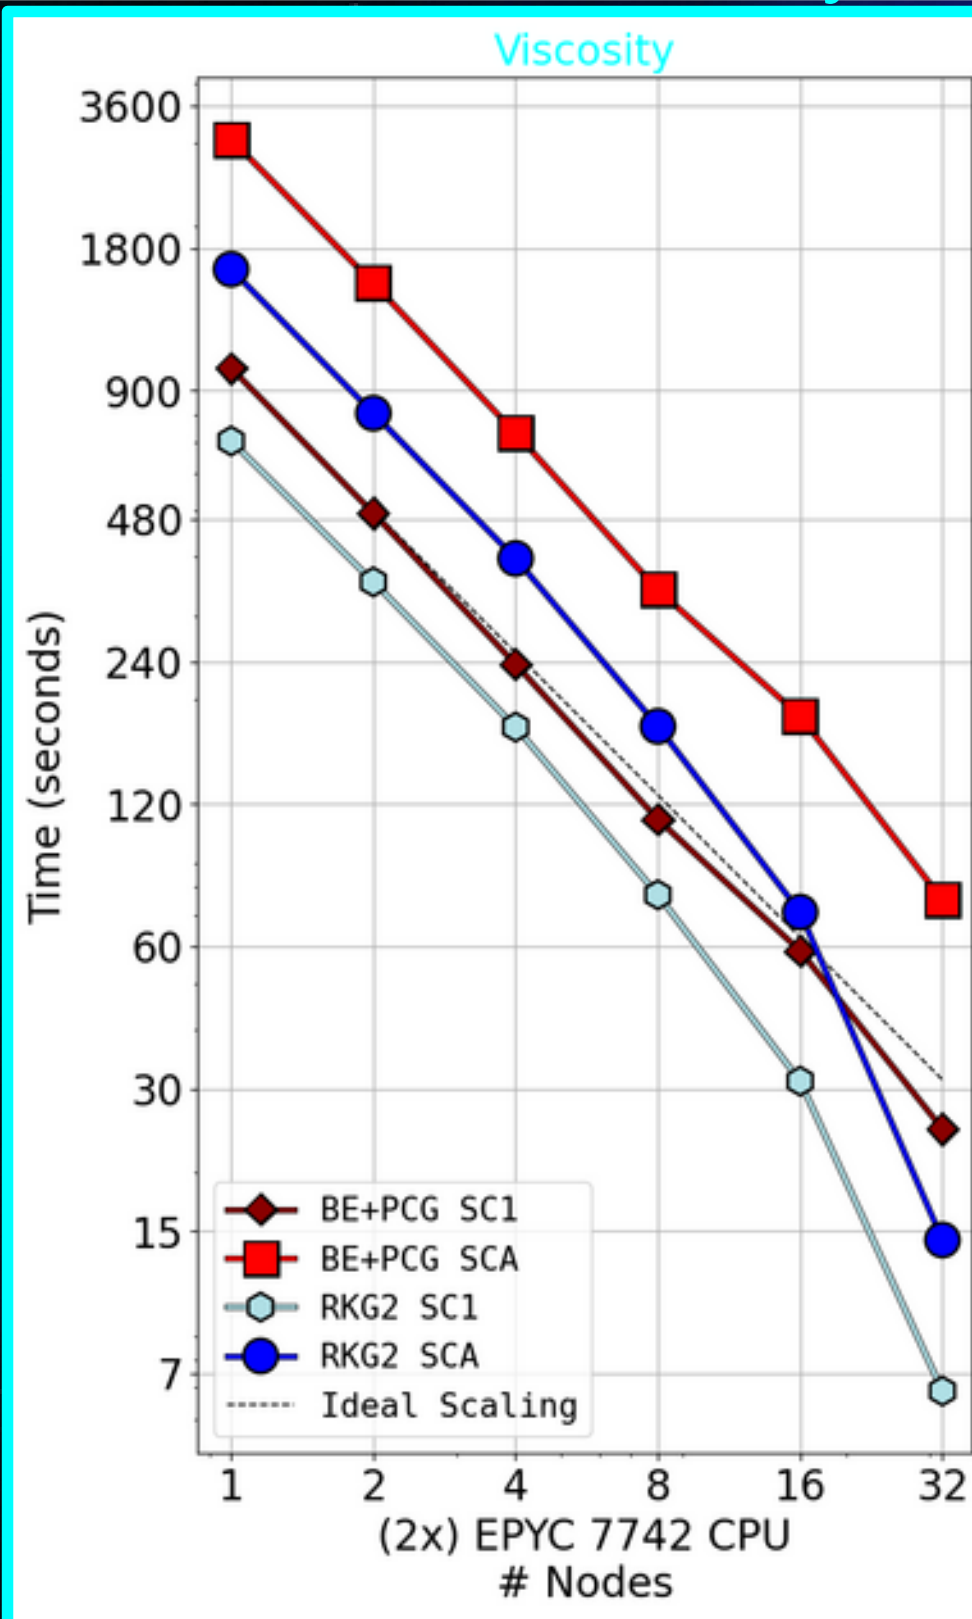
Scaling Results Test 2 Thermal Conduction

- Very few outer cycles were needed in this test for TC
- Therefore, SCA has similar run times as SC1
- Scaling is better with RKG2 than BE+PCG (PC2) on CPUs
- Run time for RKG2 is much faster than BE+PCG (PC1) on GPUs
- Scaling of both schemes similar on multi-GPU single server runs



Scaling Results Test 2 Viscosity

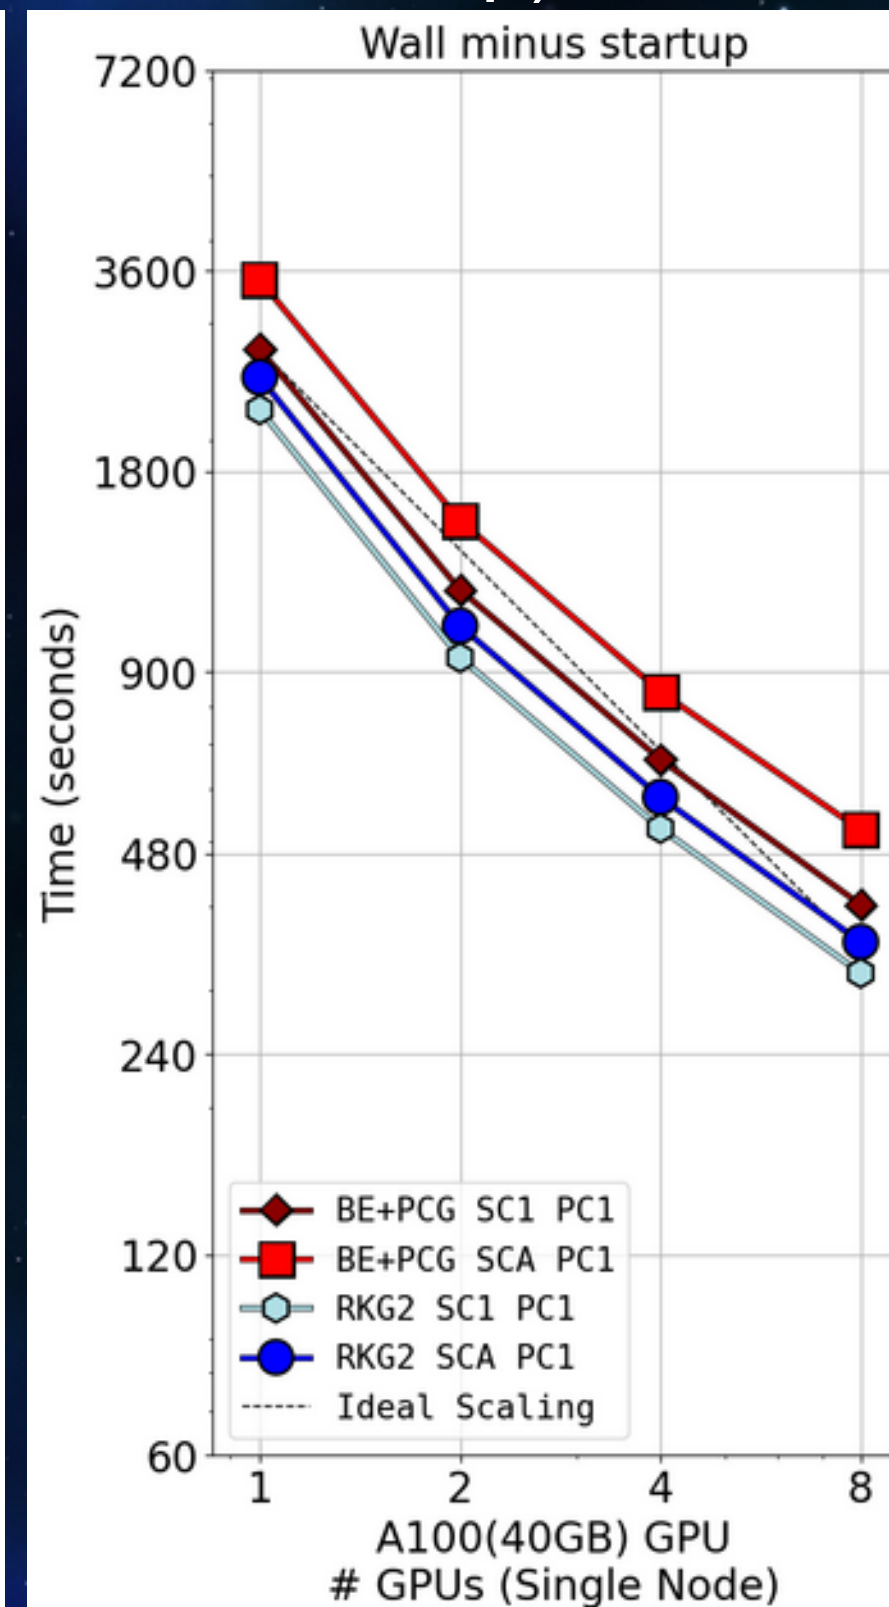
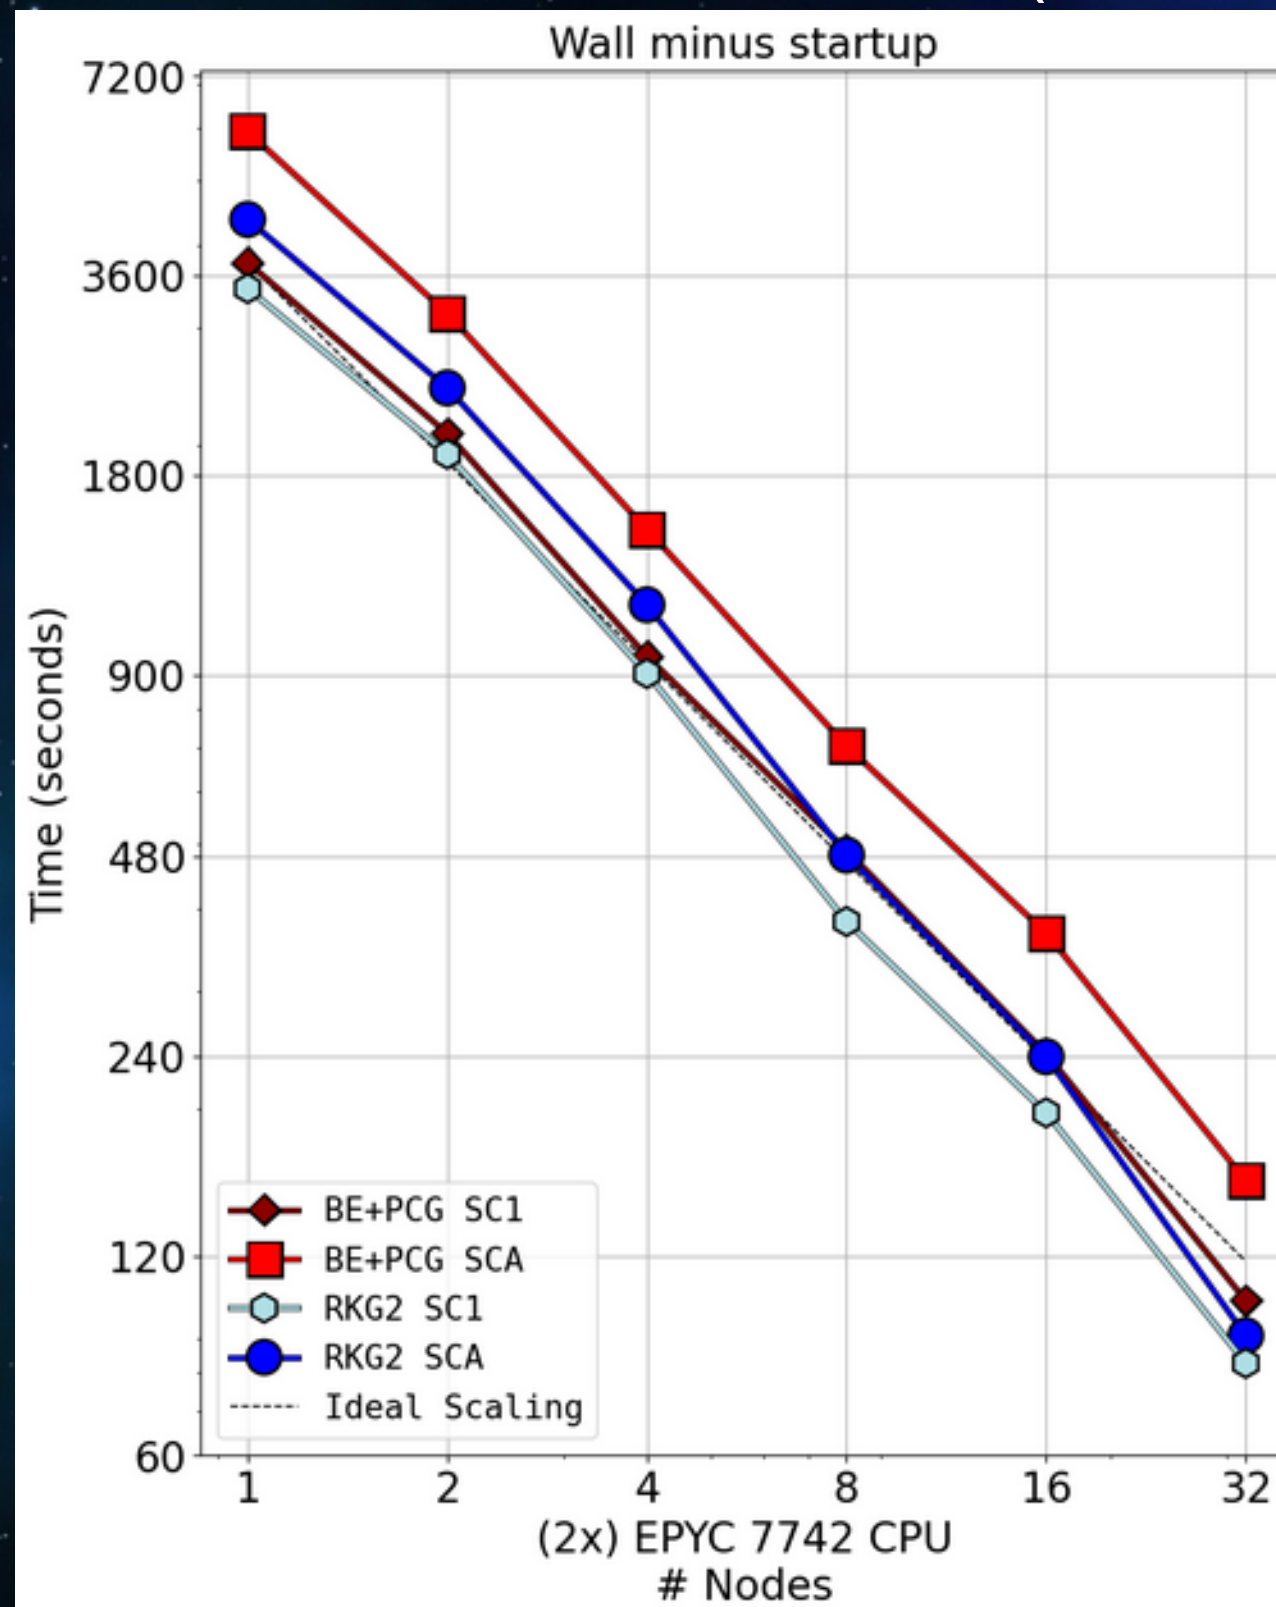
- ⌘ The test problem required ~10 outer cycles for viscosity
- ⌘ Therefore, auto-cycle has slower run time than single cycle
- ⌘ Scaling is better with RKG2 than PCG (PC2) on CPUs
- ⌘ Run time for RKG2 is faster than PCG (PC1) on GPUs with similar scaling
- ⌘ RKG2 auto-cycle faster than PCG single cycle:
 - With PCG (PC1), always the case
 - With PCG (PC2), it's faster with maximum CPUs due to better scaling and PC2 effectiveness decline



Scaling Results Test 2 **Total Wall Clock** (minus restart startup)

⊖ RKG2 auto-cycle has comparable or better performance and scaling than BE+PCG single-cycle

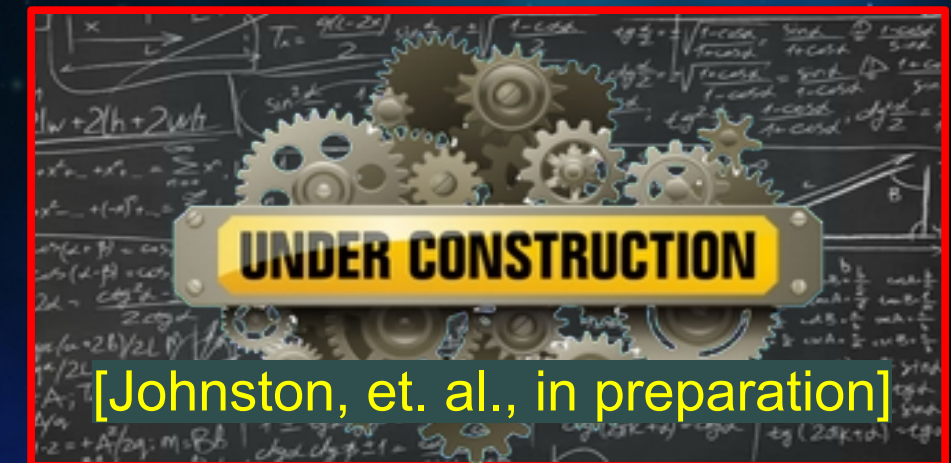
⊖ BE+PCG auto-cycle significantly slower than single-cycle



Results

- Can we get a solution as good (or better) as BE+PCG with RKG2 if we automatically outer-cycle at the practical time step limit?
 - **Yes (in our cases)**
- Can we get a “better” solution with BE+PCG if we outer-cycle it?
 - **Yes (in our cases)**
- How does outer-cycling affect performance?
 - **For BE+PCG, significant decrease in performance (in our cases)**
 - **For RKG2, small decrease in performance (in our cases)**
- Is RKG2 competitive with BE+PCG?
 - **Yes (in our cases)**

How general are these results???



Summary

- Unconditionally stable methods are often necessary when conditionally stable schemes' time-step limits for operators (e.g. parabolic) are too restrictive for practical simulations
- Using “too large” of a time step can create large errors including qualitative changes in the operator’s effect and/or inability to damp high wave modes (oscillations)
- We have tested an easy-to-calculate practical time step limit for such schemes that can be applied as an operator-split outer cycle
- The new limit can help fix known solution issues in extended stability Runge-Kutta (super time stepping) methods that are due to their poor amplification factors at high wave modes
- Testing the method with RKG2 implemented in the MAS MHD code recovered the solution accuracy of the BE+PCG method with negligible effect on performance & scaling
- Work is proceeding on the theoretical formulation/modification of the practical time step, including careful testing, especially in nonlinear cases [Johnston et al, in preparation]

$$\Delta t_p = \alpha \frac{|u_k|}{|F(u_k)|}$$

$$k : |F(u_k)| = \max |F(u)|$$

