

PAPER • OPEN ACCESS

Advancing parabolic operators in thermodynamic MHD models II: Evaluating a Practical Time Step Limit for Unconditionally Stable Methods

To cite this article: Ronald M. Caplan *et al* 2024 *J. Phys.: Conf. Ser.* **2742** 012020

View the [article online](#) for updates and enhancements.

You may also like

- [Unconditionally secure digital signatures implemented in an eight-user quantum network](#)

Yoann Pelet, Ittoop Verghese Puthoor, Natarajan Venkatachalam *et al.*

- [Influence of the Complex Interface between Transport and Catalyst Layer on Water Electrolysis Performance](#)

Tien-Ching Ma, Andreas Hutzler, Boris Bensmann *et al.*

- [Polymer Electrolyte Water Electrolysis: Correlating Porous Transport Layer Structural Properties and Performance; Part I. Tomographic Analysis of Morphology and Topology](#)

Tobias Schuler, Ruben De Bruycker, Thomas J. Schmidt *et al.*



PRIME
PACIFIC RIM MEETING
ON ELECTROCHEMICAL
AND SOLID STATE SCIENCE

HONOLULU, HI
October 6-11, 2024

Joint International Meeting of

The Electrochemical Society of Japan (ECSJ)
The Korean Electrochemical Society (KECS)
The Electrochemical Society (ECS)

Early Registration Deadline:
September 3, 2024

**MAKE YOUR PLANS
NOW!**



Advancing parabolic operators in thermodynamic MHD models II: Evaluating a Practical Time Step Limit for Unconditionally Stable Methods

Ronald M. Caplan¹, Craig D. Johnston^{2,3}, Lars K. S. Daldorff^{4,3}, and Jon A. Linker¹

¹ Predictive Science Inc., 9990 Mesa Rim Road Suite 170, San Diego, CA 92121

² Department of Physics and Astronomy, George Mason University, Fairfax, VA 22030

³ Heliophysics Science Division, NASA Goddard Space Flight Center, Greenbelt, MD 20771

⁴ Catholic University of America, 620 Michigan Ave., N.E. Washington, DC 20064

E-mail: caplanr@predsci.com, craig.d.johnston@nasa.gov, lars.daldorff@nasa.gov, linkerj@predsci.com

Abstract. Unconditionally stable time stepping schemes are useful and often practically necessary for advancing parabolic operators in multi-scale systems. However, serious accuracy problems may emerge when taking time steps that far exceed the explicit stability limits. In our previous work, we compared the accuracy and performance of advancing parabolic operators in a thermodynamic MHD model using an implicit method and an explicit super time-stepping (STS) method. We found that while the STS method outperformed the implicit one with overall good results, it was not able to damp oscillatory behavior in the solution efficiently, hindering its practical use. In this follow-up work, we evaluate an easy-to-implement method for selecting a practical time step limit (PTL) for unconditionally stable schemes. This time step is used to ‘cycle’ the operator-split thermal conduction and viscosity parabolic operators. We test the new time step with both an implicit and STS scheme for accuracy, performance, and scaling. We find that, for our test cases here, the PTL dramatically improves the STS solution, matching or improving the solution of the original implicit scheme, while retaining most of its performance and scaling advantages. The PTL shows promise to allow more accurate use of unconditionally stable schemes for parabolic operators and reliable use of STS methods.

1. Introduction

Thermodynamic magnetohydrodynamic (TMHD) models, like many other models of complex physical systems, contain processes spanning widely varying time scales. This makes numerical integration difficult, as the time steps required to describe the fastest processes can be quite small. When using explicit time-stepping methods, the stability requirements for these fast processes can further restrict the time steps (often a great deal), making the simulations computationally intractable. To combat this problem, unconditionally stable time-stepping algorithms can be used, effectively eliminating the numerical time step restrictions. However, such methods are prone to miss-use, as integrating them with very large time steps can cause fast processes of the system to not be captured, and/or introduce large errors to the system.

In our previous work ([1], referred henceforth as ‘Paper I’), we compared two popular unconditionally stable time-stepping algorithms applied to the parabolic operators within the



Content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](#). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

Magnetospheric Algorithm outside a Sphere (MAS) TMHD model, which is used to study the solar corona and heliosphere (see references therein). These were the L-stable [2] implicit Backward Euler (BE) method solved with a preconditioned conjugate gradient (PCG) solver, and an explicit super time stepping (STS) scheme, specifically, the A-stable [2] second-order Legendre Extended Stability Runge-Kutta method (RKL2) [3]. It was found that the RKL2 method computationally outperformed BE+PCG (especially in scalability across many CPUs), but suffered from solution artifacts in some isolated regions of the model domain. These artifacts appeared to be the result of using too large a time step combined with the poor high-mode damping of the A-stable RKL2 method when applied to the artificial kinematic viscosity term used to damp small spurious oscillations arising from other parts of the model.

This issue raised concerns about the robust use of STS methods in general, and in MAS, led to limiting its production use to a small selection of run types. While the BE+PCG's L-stable properties made it more robust (avoiding the solution artifacts), we have observed cases where it too can suffer from weak damping of oscillations when given too large of a time-step. One way to mitigate this issue is to 'cycle' the parabolic operator, effectively reducing the time-step, and allowing the high modes to damp through successive application of the operator. However, it is difficult to know how many cycles are required a-priory [4], and overestimation can significantly reduce the computational performance.

In this work, we follow-up Paper I by evaluating a novel, easy-to-implement method for dynamically calculating a practical time step limit (PTL) for unconditionally stable time stepping schemes applied to parabolic PDEs. The PTL is implemented in MAS, where we use it to automatically set the minimum number of cycles needed for the operator-split parabolic operators over one full MAS time step (set by the flow CFL condition). To improve performance, the PTL is re-evaluated dynamically each cycle, reducing the total number of cycles needed over the full time step. As in Paper I, we test the PTL using an implicit scheme and an STS scheme. We use two real-world cases to compare solution quality, performance, and scaling.

The paper is organized as follows. In Sec. 2, we describe the parabolic operators in the MAS model and present a brief derivation of the PTL. Then, in Sec. 3, we describe the test setup, including the unconditionally stable methods, the test cases, and the computational environment. Solutions computed with and without using the PTL are evaluated in Sec. 4, including performance and scaling results. We conclude with a discussion in Sec. 5.

2. A practical time step limit for integrating parabolic operators

Often, parts of a physical model that have smaller time scales than others are operator-split, such that each operator is integrated over the overall time step using the partial solution of the operator before it. The order and time steps of the splitting can be done in several ways, each with a cost of adding a splitting error to the system [5]. Each split operator is advanced over the time step independently by integrating

$$\frac{\partial \mathbf{u}}{\partial t} = \mathbf{F}(\mathbf{u}), \quad (1)$$

where $\mathbf{F}(\mathbf{u}) \equiv \mathbf{F}(\mathbf{x}, \mathbf{u}, \dots)$ is the currently split operator.

The MAS model¹ described in Paper I (see references within) uses an overall time step set by a flow CFL condition, and operator splits several terms that would require very small time steps. These include the parabolic operators of artificial kinematic viscosity and Spitzer thermal conduction [6] given by

$$\mathbf{F}_{\text{visc}}(\mathbf{v}) = \frac{1}{\rho} \nabla \cdot (\nu(\mathbf{x}) \rho \nabla \mathbf{v}), \text{ and } \mathbf{F}_{\text{tc}}(T) = \frac{(\gamma - 1) m_p}{2 k \rho} \nabla \cdot \left(f_c(r) f_m(T_0) \kappa_0 T_0^{5/2} \hat{\mathbf{b}} \hat{\mathbf{b}} \cdot \nabla T \right), \quad (2)$$

¹ See predsci.com/mas for additional details and references

where \mathbf{v} , T , and ρ are the plasma flow velocity, temperature, and density respectively, $\nu(\mathbf{x})$ is the coefficient of viscosity, $\hat{\mathbf{b}} = |\mathbf{B}|/\mathbf{B}$ is the normalized direction of the magnetic field, $\gamma = 5/3$ is the adiabatic index, m_p is the proton mass, k is Boltzman's constant, κ_0 is the Spitzer coefficient for thermal conduction, $f_m(T)$ is used to increase the parallel thermal conductivity at low temperatures which, in conjunction with an inverse modification to radiative cooling, broadens the transition region [7], $f_c(r)$ is a profile that limits the radial extent that collisional thermal conduction is active, and T_0 is the previous step's temperature used to keep the operator linear through lagged diffusivity [8].

If the numerical stability time step limit of the operator-split parabolic term is smaller than the overall step, it can be cycled at its stable limit (at extreme computational cost), or an unconditionally stable scheme can be utilized. As described in Sec. 1, the latter can lead to problems when the overall time step is very large compared to the explicit limit. In such cases, the operator can be cycled (with many less cycles than an explicit scheme) to reduce solution artifacts. However, the number of cycles is typically set by experimentation (see the discussion in Ref. [4]) which can be time consuming to determine, and is therefore not practical for large MHD models.

A novel method to compute a practical time-step limit (PTL) for a parabolic operator is introduced in Ref [9]. Given a parabolic equation of the form of Eq. 1, the goal is to maintain the sign of the difference in adjacent cells when integrating from time t^n to $t^{n+1} = t^n + \Delta t$:

$$\Delta \mathbf{u}_{\vec{i}}^{n+1} \Delta \mathbf{u}_{\vec{i}}^n \geq 0, \quad (3)$$

where \vec{i} is the grid point in question and $\Delta \mathbf{u}$ is the difference between the solution at \vec{i} and its cell neighbor. In general, this condition can be considered in all directions and dimensions. For simplicity, we show one direction, where we define $\Delta \mathbf{u}_i^n = \mathbf{u}_i^n - \mathbf{u}_{i-1}^n$. Using the first-order Taylor expansion of the solution \mathbf{u}^{n+1} in Δt and dropping higher order terms, Eq. 3 becomes

$$(\Delta \mathbf{u}_i^n)^2 + \Delta t \Delta \mathbf{u}_i^n \Delta \mathbf{F}_i^n \geq 0,$$

where $\Delta \mathbf{F}_i^n = \mathbf{F}_i(\mathbf{u}^n) - \mathbf{F}_{i-1}(\mathbf{u}^n)$. The only case where this condition is not unconditionally satisfied is when $\Delta \mathbf{u}_i^n \neq 0$, $\Delta \mathbf{F}_i^n \neq 0$, and $\Delta \mathbf{u}_i^n \Delta \mathbf{F}_i^n < 0$, in which case

$$\Delta t \leq -\frac{\Delta \mathbf{u}_i^n}{\Delta \mathbf{F}_i^n}. \quad (4)$$

The condition can be checked in all directions surrounding grid cell \vec{i} that meet the criteria over every grid cell, and the minimum time step calculated. However, Ref. [9] found (in their cases) that the minimum Δt always occurred at the location of the maximum absolute value of the components of \mathbf{F} (denoted as F_{\max}). This allowed them to avoid needing to compute the conditions over the whole grid and sidestep numerical sensitivity issues near values very close to zero. Here, we make the assumption that their result extends to our cases and only compute the time step conditions at the grid location of F_{\max} , which we denote as \vec{k} . The PTL can then be defined as

$$\Delta t_{\text{PTL}} \leq \min \left[-\frac{u_{\vec{k}}^n - u_{\vec{k}+\vec{i}}^n}{F_{\vec{k}}(u^n) - F_{\vec{k}+\vec{i}}(u^n)} \right], \quad (5)$$

where the minimum is taken for all \vec{i} one cell away in the 6 orthogonal grid directions.

Once the PTL is found, the number of cycles to integrate the operator at the PTL over the larger overall time-step can be set. However, since the operator is parabolic and smooths the solution with each application, the PTL can be dynamically re-evaluated and applied after each cycle. This can increase the PTL size after each step, leading to much fewer overall cycles, increasing performance. A detailed formulation and analysis of the PTL is discussed in Ref [9]. Here, we test applying the PTL to a full production code (MAS) on real-world problems.

3. Test setup

3.1. Unconditionally stable schemes

In Paper I, we described our implementation of the first-order accurate, L-stable BE+PCG scheme, which has been successfully used in the MAS code for many years. The implementation uses two choices of preconditioners (PC1 and PC2). PC1 (point-Jacobi) is inexpensive to formulate and apply, but limited in its effectiveness to reduce solver iterations, while PC2 (non-overlapping ILU0 factorization) is more expensive to formulate and apply, but also more effective at reducing iterations. PC1 is easily vectorized, making it simple to implement efficiently on vector-optimized hardware (e.g. GPUs), while the standard implementation of PC2 is not vectorizable, requiring the use of alternative algorithms and libraries for use on GPUs [10, 11]. The current version of MAS only supports PC1 on GPUs, while PC2 is used when running on CPUs (an important factor in comparing MAS's CPU and GPU performance). The BE+PCG has two important drawbacks - the requirement of global communications for the dot products which hinders scaling, and the necessity of using a linear operator.

The other scheme described and tested in Paper I was the explicit second-order extended stability Runge-Kutta Legendre super time stepping (RKL2) scheme [12, 3]. RKL2 is easy to implement, second-order accurate, can handle a non-linear operator, has no global synchronization points, and is vectorizable yielding efficient implementation on GPUs. However, it is only A-stable and does not damp high modes well for large time steps.

The second-order Gegenbauer method (RKG2) recently introduced in Ref. [13] has similar performance to the RKL2 scheme, but with a better damping amplification factor for high modes. In Fig. 1, we show the estimated speedups of RKL2 and RKG2, and compare the amplification factors of the BE, RKL2, and RKG2 with $\alpha = 3/2$ (see Ref. [14]) schemes to the exact solution of a simple 1D heat equation problem. We see that the RKG2 scheme's amplification factor

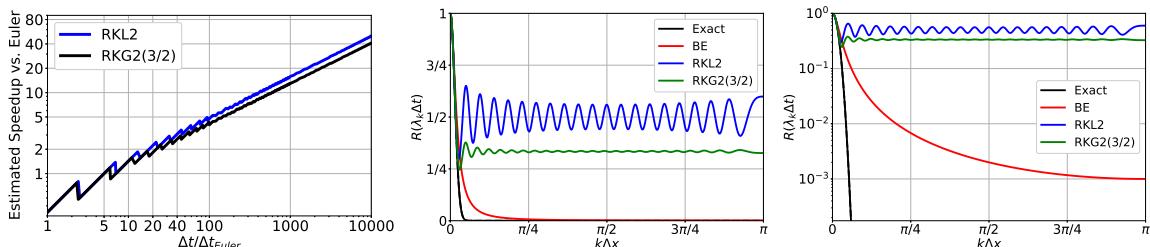


Figure 1: Left: Estimated speedup of the RKL2 and RKG2(3/2) schemes compared to explicit Euler (approximating one Euler step to have the same computational cost as one STS iteration). Middle: Amplification factors for the BE, RKL2, and RKG2 schemes when applied to a 1D uniform grid heat equation discretized with a second-order central finite difference at a time-step 500 times larger than the explicit Euler limit. Right: Same as middle, but on a log scale to show the difference in the BE scheme's damping rate compared to the exact solution.

is qualitatively similar to the RKL2, but saturates to a lower value and is less variable across modes. The RKG2 is also shown to be similar in performance to RKL2. We note that even the BE scheme's damping of modes diverges from the exact solution greatly, but overall it is much better than that of the STS schemes.

Since the inability to efficiently damp high modes is a key factor in the solution artifacts we observe in Paper I, we choose to use RKG2 over RKL2 in this work. We also note that it is critically important to ensure that the number of STS iterations of the RKG2 method be odd, otherwise the amplification factor approaches one at the highest modes. Details of the RKG2 scheme implementation are shown in [Appendix A](#).

3.2. Test cases

To test the effects and performance of the PTL, we utilize two production MAS simulation runs. The first test (Test 1) is a lower resolution version of a thermodynamic MHD relaxation based on simulations used in Ref. [15]. This test has 3.4 million highly non-uniform grid points and uses a constructed lower radial magnetic field boundary combining a global dipole with a localized bi-pole representing an active region. The second test (Test 2) is a slightly modified (small changes in the viscosity values) version of the thermodynamic MHD relaxation used in Paper I with a size of 22.8 million grid points. This test uses real-world observed surface radial magnetic fields for the lower boundary.

For comparing solutions with and without using the PTL, we run each test's relaxation for a total of 8 hours of simulated physical time. For performance and scaling results, we only use Test 2, where we start the simulation with the final output of an 8-hour relaxation, and continue the run for another 6 simulated physical minutes. Since these performance runs are much shorter simulations than those used in production, we subtract off the loading time of the relaxation solution initial condition from the wall clock time.

3.3. Computational environment

Since Paper I, the MAS code has been updated to run on NVIDIA GPUs through the use of OpenACC [16] and Fortran's 'do concurrent' standard parallelism constructs [17]. We therefore test the methods on both CPU and GPU systems. For the solution comparison run of Test 1, we utilize an in-house workstation with an NVIDIA RTX 3090Ti GPU, while for Test 2 we use 32 nodes of the Expanse system at the San Diego Supercomputer Center (SDSC). For testing the parallel scaling performance on CPUs, we use between one and 32 nodes of Expanse, while on GPUs, we use between one and eight GPUs on a single node at the Delta system at the National Center for Supercomputing Applications (NCSA). The details of the hardware and software configurations used are shown in [Appendix B](#).

We note that the MAS code is highly memory-bandwidth bound, making the maximum memory bandwidth the best indication of expected relative performance between systems, given the same algorithm. However, as discussed in Sec. 3.1, MAS uses a more effective preconditioner when run on CPUs than on GPUs, making the relative performance between CPU and GPU a practical time-to-solution comparison, and not an apples to apples hardware comparison.

4. Results

The key questions we want to address are:

- (i) Can using the PTL with the RKG2 scheme obtain a solution similar (or better) to the original BE+PCG scheme?
- (ii) Does using the PTL with the BE+PCG scheme improve *its* solution?
- (iii) How does the PTL affect performance for both BE+PCG and RKG2 schemes?
- (iv) Does using the PTL allow RKG2 to be competitive with the original BE+PCG scheme in both solution quality and performance?

4.1. Solution comparisons

To address the first two questions, we run the test cases with both BE+PCG and RKG2, each with and without using the PTL and compare the solutions. Like in Paper I, the overall global solution looks qualitatively the same using all methods (not shown here). Our focus here is on the locations and quantities within the solutions that exhibit oscillatory behavior, such as the zoomed views in Paper I's Fig. 7. For Test 1, this issue can be seen in the ϕ -component of the velocity field near the solar surface. A zoomed portion of the solution, along with 1D cuts through the view are shown in Fig. 2.

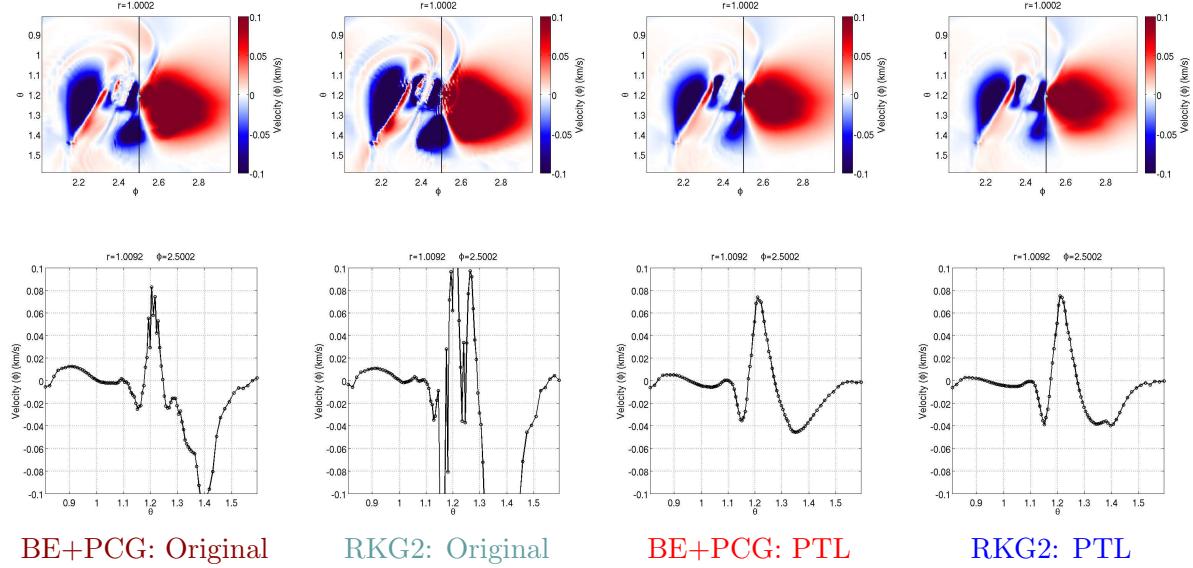


Figure 2: Comparison run results for Test 1. A small section of the solution output of v_ϕ is shown in the $\theta - \phi$ plane (top) with the black line indicating the cut that is plotted in 1D (bottom). The PTL is shown to improve the quality of the solution when using either RKG2 and BE+PCG. The scheme used for each result is shown at the bottom of each column, where the text color corresponds to the performance result figures in the next section.

In this case, we notice that even the original BE+PCG scheme does not fully damp out the run's oscillations, implying that either the viscosity coefficient of the problem is not high enough, or that the time step is so large, not even the BE+PCG scheme can apply the specified viscosity correctly. The latter seems to be the case since the PTL run with BE+PCG eliminates the oscillations, answering question (ii) in the affirmative.

For the RKG2 scheme, we see that the original implementation does a poor job applying the viscosity to damp oscillations and/or is adding its own oscillations due to the large time step (see Ref. [9]). In contrast, when using the PTL with RKG2, the oscillations have been damped, and the resulting solution is qualitatively similar to the BE+PCG solution with PTL, with only a small 'notch' near $\theta = 1.4$ being a noticeable difference. The quality improvement of the solution when using the PTL is also seen clearly in the 2D cut planes, where the PTL has drastically cleaned up the original RKG2 solution. In this case, we have shown that the answer to question (i) is 'yes', as the PTL RKG2 is not only 'as good' as the original BE+PCG, but, in this case, is substantially better.

Moving to Test 2, in Fig. 3 we show zoomed-in portions of the ϕ -component of the current density for each solution, along with 1D cuts through the view. We see similar results as in Test 1. The original BE+PCG scheme does not have much oscillatory behavior in this case, but when using the PTL, the solution changes and becomes much smoother overall. The original RKG2 scheme once again suffers from oscillatory behavior, but is cleaned up substantially when applying the PTL. In this case, the PTL RKG2 solution is even closer qualitatively to the PTL BE+PCG than in Test 1.

We can therefore conclude (at least for these cases) that the answer to both questions (i) and (ii) is 'yes', as the PTL improves the solution of both schemes, and allows the RKG2 scheme to yield a very close solution to BE+PCG, with the PTL RKG2 solution being a significant improvement over the original BR+PCG scheme. However, to know whether it is practical to

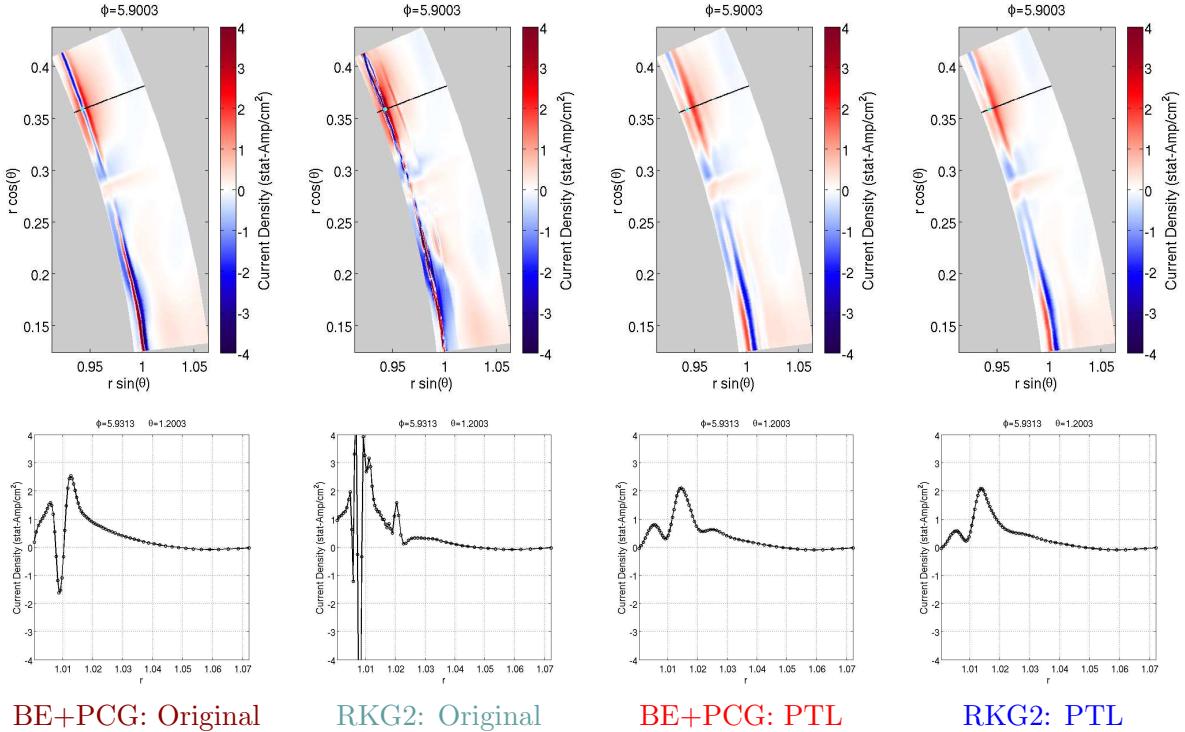


Figure 3: Comparison run results for Test 2. A small section of the solution output of j_ϕ is shown in the $r - \theta$ plane (top) with the black line indicating the cut that is plotted in 1D (bottom). The PTL is shown to improve the quality of the solution when using both RKG2 and BE+PCG. Notation is the same as Figure 2.

use the PTL for either scheme, these solution results need to be viewed in conjunction with each scheme's computational performance.

4.2. Performance

We start by looking at the performance results for the solution comparison runs shown in the previous section, which we show in Fig. 4. In both tests, we see that the original RKG2 scheme is faster than the original BE+PCG scheme as expected. Using the PTL on BE+PCG makes the code slower by up to a factor of two, with the slowdown for viscosity being a factor of around four. While this method did improve the solution over the original scheme, the improvement may not be worth the performance hit. Looking at the PTL RKG2 runs, we see that, while they are slower than the original RKG2 runs, they are comparable to the run time of the original BE+PCG scheme. Since the solution for the PTL RKG2 scheme is an improvement to the original BE+PCG solution, and comparable to the PTL BE+PCG scheme's solution, the small increase in run time is acceptable, and makes the RKG2 competitive with the original BE+PCG scheme, answering question (iv) in the affirmative for this case.

In Paper I, the performance advantages of the STS scheme over BE+PCG were even more pronounced when looking at parallel scaling across multiple CPUs. Here, we test the scaling with and without the PTL to see if the analysis above remains valid, and if the RKG2's advantage increases. The timing results for the two parabolic operators (thermal conduction and viscosity) portion of the runs are shown in Fig. 5. On the CPUs, we see that the BE+PCG exhibits expected scaling, and for the thermal conduction operator on large numbers of nodes, exhibits 'super-scaling'. This super scaling is likely due to the very large size and speed of the EPYC

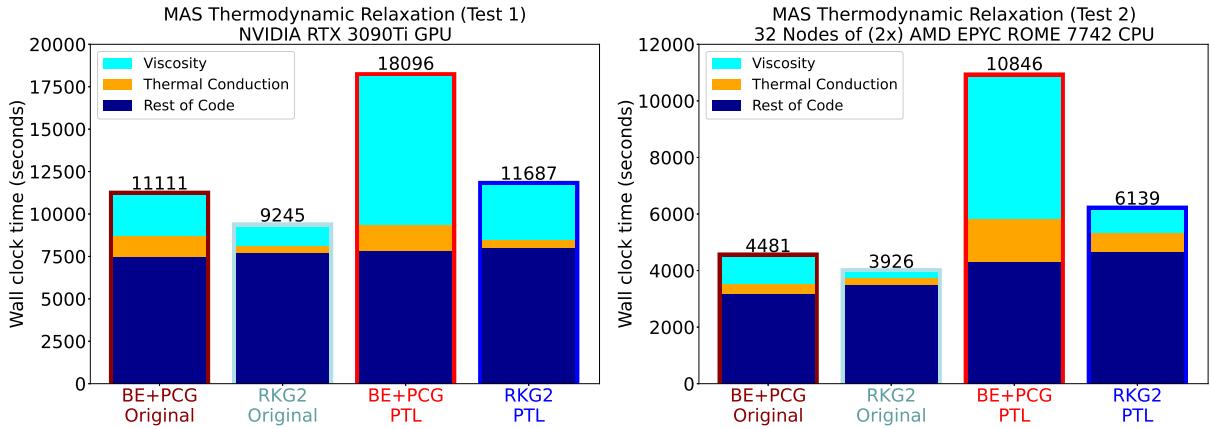


Figure 4: Timing results for the solution comparison runs in Sec. 4.1 for Test 1 (left) and Test 2 (right).

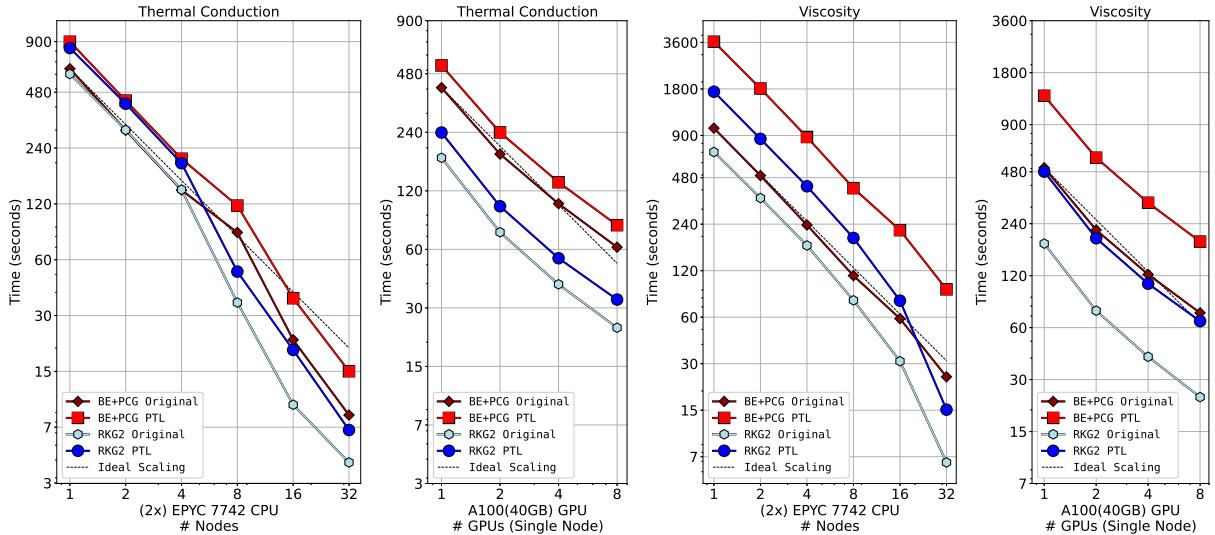


Figure 5: Test 2 scaling results for the thermal conduction (left two plots) and viscosity (right two plots) operators on CPUs (left) and GPUs (right).

CPU's cache, such that the local portion of the grid (which decreases with increasing number of CPU cores) is fitting into the fast cache. The RKG2 exhibits stronger super scaling for both the thermal conduction and viscosity operators, and is overall faster than the BE+PCG. When using the PTL, each method runs slower than their original implementations as expected, up to a factor of four. However, while the RKG2 with PTL runs slower than the original non-PTL BE+PCG method on small numbers of CPU nodes, when run on many CPU nodes, the super scaling causes it to be faster. Therefore, we see that the PTL RKG2 scheme can outperform the original BE+PCG scheme, while producing a better solution as shown in Sec. 4.1.

On the GPUs, all methods do not scale perfectly, likely caused by GPU-GPU communication overheads and the problem size not being large enough to fully saturate the GPUs. In this case, the RKG2 is always faster than the BE+PCG. This is expected since the GPU runs only use the less effective PC1 preconditioner for BE+PCG. For thermal conduction, the PTL RKG2 method outperforms the original BE+PCG method by a factor of 2, while for viscosity the two run times are almost the same. Given the superior solution quality of the PTL RKG2 over the original

BE+PCG method, these results strongly favor the use of PTL with RKG2 over BE+PCG.

In all the cases shown here, we see that using PTL with BE+PCG slows down the operator considerably, by up to a factor of 4. Therefore, although the PTL with BE+PCG yields a better solution than without PTL, the drop in performance may be unacceptable. This implies that the PTL has made STS methods competitive with BE+PCG, answering question (iv) in the affirmative.

The above analysis was for the two parabolic operators in isolation. However, since these operators are only part of the overall run time (around 30%, see Paper I) of the MAS code, we also look at the effects on the total wall clock time, which we show in Fig. 6. In the CPU case,

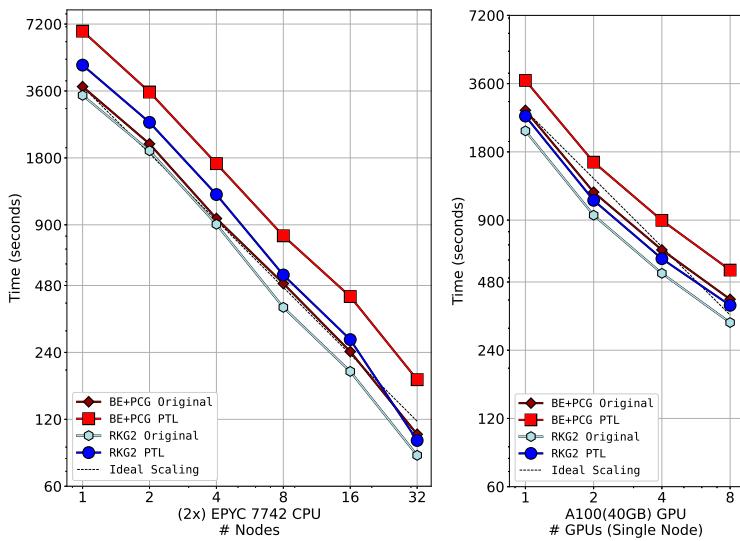


Figure 6: Test 2 scaling results for total Wall clock time (with startup time removed) on CPUs (left) and GPUs (right).

the overall code scales ideally, with a small amount of super-scaling when run on 32 CPU nodes. Given the amount of super scaling shown in Fig. 5, this shows that different sections of the code have varying levels of scaling efficiency.

We once again see that using the RKG2 scheme yields the best performance overall. The RKG2 scheme with PTL is slower than the original BE+PCG scheme for small numbers of CPU nodes, but has similar performance for more than 8 CPU nodes, and surpasses its performance for 32 CPU nodes. Using PTL with BE+PCG is once again much slower. On GPUs, the RKG2 with PTL closely matches the performance of the original BE+PCG scheme on all numbers of GPUs. Given the improvements of the solution using the PTL with RKG2 over the original RKG2 and BE+PCG, the new scheme is overall advantageous, where the performance loss on low numbers of CPU nodes is likely acceptable.

5. Discussion

In this paper, we have followed up our analysis in Paper I comparing unconditionally stable explicit STS schemes to implicit schemes for parabolic operators in a thermodynamic MHD model by applying an easy-to-use practical time step limit (PTL). This limit is designed to ensure a consistent solution structure from one time step to the next and is implemented as inner steps at a dynamically recalculated time step. Using the PTL was shown to improve the solution for both the RKG2 STS scheme and for the implicit BE+PCG scheme. Using the PTL reduces performance of the operators, in some cases by a large factor. However, it was found

that the RKG2's original performance and scaling advantage over the BE+PCG scheme allow it to be competitive against the BE+PCG scheme even when using the PTL, all while exhibiting a much better solution than the original RKG2. The results presented imply the following answers to the questions in Sec. 4:

- (i) Can using the PTL with the RKG2 scheme obtain a solution similar (or better) to the original BE+PCG scheme?
Yes, as shown in Figs. 2 and 3.
- (ii) Does using the PTL with the BE+PCG scheme improve *its* solution?
Yes, as shown in Figs. 2 and 3.
- (iii) How does the PTL affect performance for both BE+PCG and RKG2 schemes?
In both cases, the performance decreases, however due to the original performance advantage of the RKG2 scheme, the PTL applied to RKG2 exhibits similar or better performance than the original BE+PCG.
- (iv) Does using the PTL allow RKG2 to be competitive with the original BE+PCG scheme in both solution quality and performance?
Yes, as shown in Figs. 2,3,4 and 6

We note that these answers are implied by the results obtained here, and are not guaranteed to be true in all cases. We have tested the PTL in other models, such as our flux evolution code HipFT², and our magnetic map smoothing tool called Diffuse [18]. In both these models, the PTL always avoided solution artifacts that were shown to occur when using very large time steps, with very little decline in performance. Work is proceeding on the theoretical formulation and analysis of the PTL, including careful testing, especially in nonlinear cases [9]. While the performance when applied to real-world simulations will be model and scheme dependent, the results shown here are encouraging that using the PTL can be used as a simple high performance way to avoid unwanted behavior of parabolic operators when using large time steps. The PTL also may be a great way to ensure proper behavior of STS methods, making them a more robust and practical option for the integration of parabolic operators in multi physics models such as thermodynamic MHD.

Appendix A.

The RKG2 scheme with $\alpha = 3/2$ [13, 14] is given by

$$\begin{aligned}
 u_0 &= u^n \\
 y_0 &= F(u_0) \\
 u_1 &= u_0 + \tilde{\mu}_1 \Delta t y_0 \\
 \text{do } k &= 2 : s \\
 u_k &= \mu_k u_{k-1} + \nu_k u_{k-2} \\
 &\quad + (1 - \mu_k - \nu_k) u_0 \\
 &\quad + \tilde{\mu}_k \Delta t F(u_{k-1}) + \gamma_k \Delta t y_0 \\
 \text{enddo} \\
 u^{n+1} &= u_s,
 \end{aligned}
 \quad
 \begin{aligned}
 \text{do } k &= 3 : s \\
 b_k &= \frac{4(k-1)(k+4)}{3k(k+1)(k+2)(k+3)} \\
 \mu_k &= \left(2 + \frac{1}{k}\right) \frac{b_k}{b_{k-1}} \\
 \nu_k &= -\left(\frac{1}{k} + 1\right) \frac{b_k}{b_{k-2}} \\
 \tilde{\mu}_k &= \mu_k w \\
 \gamma_k &= \left(\frac{k(k+1)}{2} b_{k-1} - 1\right) \tilde{\mu}_k,
 \end{aligned}
 \quad
 \text{enddo}$$

² Soon to be released at <https://www.github.com/predsci/hipft>

where u^n is the current solution, u^{n+1} is the solution after integrating the full time step Δt , F is the parabolic operator, u_k is the solution at STS iteration k , and

$$w = \frac{6}{(s+4)(s-1)}, \quad b_0 = 1, \quad b_1 = \frac{1}{3}, \quad b_2 = \frac{1}{15}, \quad \mu_1 = 1, \\ \tilde{\mu}_1 = w, \quad \mu_2 = \frac{1}{2}, \quad \nu_2 = -\frac{1}{10}, \quad \tilde{\mu}_2 = \mu_2 w, \quad \gamma_2 = 0.$$

The number of required iterations s is given by

$$s = \left\lceil \frac{1}{2} \sqrt{25 + 24 \frac{\Delta t}{\Delta t_{\text{Euler}}}} - \frac{3}{2} \right\rceil,$$

where Δt_{Euler} is the explicit Euler stability time step. We note that it is critically important to ensure that the number of iterations be odd (by adding an iteration if needed), otherwise the amplification factor approaches one at the highest modes.

Appendix B.

The details of the hardware and software specifications used in this paper are shown in Table B1.

SDSC Expanse CPU Node		NCSA Delta GPU(8x) Node	Local GPU Workstation
CPU	(2x) AMD EPYC Rome 7742	(2x) AMD EPYC Milan 7763	Intel Core-i5 Raptor Lake 13600KF
GPU	N/A	(8x) NVIDIA A100-40GB-SXM4	NVIDIA RTX 3090 Ti
CPU Cores	128	128	14
DRAM	128 GB	320 GB (GPUs)	24 GB (GPU)
Peak Memory Bandwidth	410 GB/s	12,440 GB/s (GPUs)	1,008 GB/s (GPUs)
Max Flops	9 TFlops	78 TFlops (GPUs)	0.625 TFlops (GPUs)
Network	HDR InfiniBand	HPE SlingShot	N/A
MPI Library	OpenMPI 4.1.3	OpenMPI 3.1.5	OpenMPI 3.1.5
Compiler	GCC 10.2.0	NV HPC SDK 23.5	NV HPC SDK 23.5
OS	Rocky 8.7	Red Hat Enterprise 8.8	Mint 21.2

Table B1: System hardware and software configuration of SDSC Expanse, NCSA Delta, and our local workstation used for the results shown in this paper.

Acknowledgments

Work at Predictive Science Inc. was supported by the NASA LWS Strategic Capabilities Program (grant 80NSSC22K0893), the NSF PREEVENTS program (grant ICER1854790), and the NSF/NASA SWQU program (grants AGS 2028154 and 80NSSC20K1582). CDJ and LKSD acknowledge support from the NASA GSFC Heliophysics Internal Scientist Funding Model program (PI Jim Klimchuk). CDJ also acknowledges support from the NASA LWS Focused Science Topic programs: NNH21ZDA001N-LWS “The Origin of the Photospheric Magnetic Field: Mapping Currents in the Chromosphere and Corona” (PI Pete Schuck) and NNH17ZDA001N-LWS “Investigating Magnetic Flux Emergence with Modeling and Observations to Understand the Onset of Major Solar Eruptions” (PI Mark Linton). Computational resources at the San Diego Supercomputer Center and the National Center for Supercomputing Applications were provided by the NSF-supported ACCESS program (allocation MCA03S014). We thank Puget Systems for providing the RTX 3090 Ti GPU for our local workstation.

References

- [1] Caplan R M, Mikić Z, Linker J A and Lionello R 2017 *Journal of Physics: Conference Series* **837** 012016
- [2] Bui T 1979 *Journal of the ACM (JACM)* **26** 483–493
- [3] Meyer C D, Balsara D S and Aslam T D 2014 *Journal of Computational Physics* **257** 594–626
- [4] Dawes A 2021 *Computers and Fluids* **214** 104762 ISSN 0045-7930
- [5] Simpson M J and Landman K A 2008 *Mathematics and Computers in Simulation* **77** 9–21
- [6] Priest E 2014 *Magnetohydrodynamics of the Sun* (Cambridge University Press)
- [7] Linker J A, Lionello R, Mikić Z and Amari T 2001 *Journal of Geophysical Research: Space Physics (1978–2012)* **106** 25165–25175
- [8] Galligani E 2012 *Advances in Difference Equations* **2012** 1–21
- [9] Johnston C D, Daldoff L K S and Caplan R M 2024 *In Preparation*
- [10] Dufrechou E and Ezzatti P 2018 *2018 26th Euromicro International Conference on Parallel, Distributed and Network-based Processing (PDP)* pp 196–203
- [11] Ma W and Cai X C 2021 *The International Journal of High Performance Computing Applications* **35** 121–135
- [12] Meyer C D, Balsara D S and Aslam T D 2012 *Monthly Notices of the Royal Astronomical Society* **422** 2102–2115
- [13] O'Sullivan S 2019 *Journal of Computational Physics* **388** 209–223
- [14] Skaras T, Saxton T, Meyer C and Aslam T D 2021 *Journal of Computational Physics* **425** 109879
- [15] Reeves K K, Török T, Mikić Z, Linker J and Murphy N A 2019 *The Astrophysical Journal* **887** 103
- [16] Caplan R M, Linker J A, Mikić Z, Downs C, Török T and Titov V S 2019 *Journal of Physics: Conference Series* **1225** 012012
- [17] Caplan R M, Stulajter M M and Linker J A 2023 *2023 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)* pp 582–590
- [18] Stulajter M M, Caplan R M and Linker J A 2022 *Accelerator Programming Using Directives* ed Bhalachandra S, Daley C and Melesse Vergara V (Cham: Springer International Publishing) pp 3–21 ISBN 978-3-030-97759-7