Tightly Coupled, Partitioned Time-Integration Methods*

Emil Constantinescu† Jed Brown Barry Smith

Argonne National Laboratory, Mathematics and Computer Science Division, Argonne, IL 60439
{emconsta,jedbrown,bsmith}@mcs.anl.gov

One of the most challenging aspects of fusion energy sciences (FES) simulations is the wide range of overlapping time scales. In this paper we discuss advanced time-integration algorithmic developments for integrated simulations of magnetic FES.

Adaptive, flexible, and robust time-integration algorithms with appropriate error monitors and controllers are needed in order to efficiently reach accurate and meaningful solutions. Due to the multiscale nature, optimal solutions can be obtained by using partitioned time integrators rather than a single strategy. Each partition is adapted to specific properties of a physics component or region, resulting in a series of implicit and explicit partitions such that stiffness and mesh refinement are directly addressed by the corresponding scheme. In this paradigm an essential component is the tight coupling among integrators and scalable (non)linear solvers that are tuned to run on large-scale machines. Therefore, the use of high-performance scientific computing libraries such as PETSc [3, 2, 4] is indispensable.

Context. The mathematical models of magnetic FES components are represented by multiscale PDEs. For instance, multiscale dynamics aspects arise in the edge plasma application from the rapid change of profiles across the pedestal, which plays a key role in determining overall fusion performance and component lifetimes. These PDEs are discretized in space and advanced in time by using the method of lines. The resulting ODEs need to be integrated in time. Time integration methods play a key role in FES simulations, and therefore efficient integration methods are a ubiquitous issue for all components.

Efficient integration of multiscale ODEs is challenging because the system contains both stiff and nonstiff components. Processes that evolve at a relatively slow rate (nonstiff) are generally handled more efficiently by low-cost explicit schemes, whereas processes that evolve relatively rapidly (stiff) are handled more efficiently by more expensive implicit schemes. Moreover, time-dependent problems can become stiff under mesh refinement. Explicit schemes have restrictive stability constraints, whereas implicit schemes have weak or no stability constraints on time steps; however, implicit schemes do have accuracy constraints that affect the nonlinear solver performance.

Time Integration. Fusion simulations have traditionally used classical Strang or Lie splittings [14] to address multiscale effects and preserve legacy codes. These are *loosely coupled* methods that integrate each component independently within a time step and work well with low-order methods. However, as the discretization methods have become more sophisticated, these strategies have shown significant drawbacks, including nonuniform or lack of convergence [12] and ineffective use of computational resources [13, 11]. Moreover, classical approaches lack a mechanism for error control across all scales.

Tightly Coupled Partitioned Time Integration. *Tightly coupled* methods require that all components be evaluated at every stage of the computation during one time step, which facilitates strict control and removal of splitting errors. Note that the fact that a method is explicit or implicit has nothing to do with the coupling. Partitioned methods can use different integrators for different components of the solution or system and can be seen as a generalization of fully implicit, explicit, and single-rate methods. Partitioned methods emerge as a robust and flexible strategy to tackle issues found in multiscale-multiphysics problems. Moreover, developing effective preconditioners is simplified because the implicit part is less complex. Therefore, tightly coupled partitioned methods can have the same positive properties as fully implicit schemes but can be significantly more efficient [10, 1, 9, 6].

Among partitioned methods, IMEX approaches are probably the best known for alleviating multiscale difficulties by allowing processes to be advanced simultaneously by either an explicit or implicit scheme [1, 6, 9]. Moreover, partitioned methods can accommodate a wide range between loosely coupled to tight temporal coupling. For instance, in gyrokinetic equations, partitioning the real and velocity space operators can facilitate the incorporation of the collision operator.

Development of time-stepping methods is guided by accuracy, stability, and performance, as well as other problem-dependent properties such as conservation. Efficient methods must also account for how well they utilize linear and nonlinear solvers [4] and the computational architectures [8, 6].

We now turn to the question of what is an efficient method. In the following example we illustrate how optimizing for different objectives can change the performance of time-stepping methods [7]. In Fig. 1 we show the performance of several partitioned methods used to integrate a partitioned nonstiff-advection–stiff-reaction problem [10]. We consider two second-order methods [ARK2], one optimized for accuracy, the other for stability, both having similar computational cost. We note a significant discrepancy in accuracy between the two methods. Based on this assessment one would choose [ARK2 – accuracy]; however, if stability were limiting the step size, then [ARK2 – stability] would likely be more efficient. Next we compare a third-order method [ARK3] with a method that has mixed orders. [EX rk4 - IM rk2] is a partitioned method whose cost is similar to that of [ARK3], but its two partitions are fourth order for the explicit scheme and second order for the implicit partition. Interesting in this case is that the mixed-order method, although it is overall second order, performs better than [ARK3] at larger steps. Such mixed-order methods can be used when accuracy in the stiff components is deemed less important, for instance, when high order accuracy of the Poisson-bracket $\mathbf{E} \times \mathbf{B}$ nonlinearity in the gyrokinetic equations is needed.

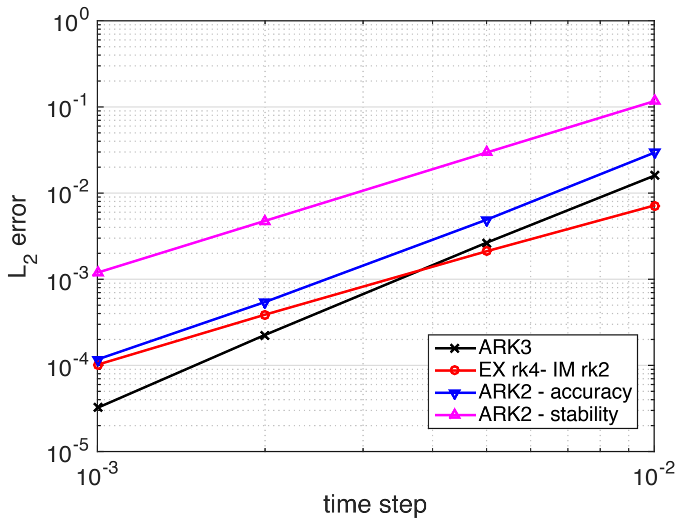

Figure 1: Performance of partitioned methods.

High-order methods are usually desirable, but they do not guarantee an accurate solution. Time step adaptivity can be used to assess and efficiently control the a posteriori solution accuracy. This creates a need for robust error estimation strategies [5] and global error control, which factor into the time-stepping method development as well. The accuracy shown in Fig. 1 is easy to compute for toy problems; however, in real complex simulations, one must rely on error estimation and control to gauge the true simulation accuracy. The error estimates can be severely misleading when using standard techniques (i.e., local truncation error) to estimate numerical errors [5], and therefore new strategies to efficiently estimate the errors need to be devised along with method development. One possibility is the new strategies based on general linear methods [5]. We also note that the time-integration strategies rely on effective scalable solvers.

Computing and Next Steps. Partitioned integrators provide a flexible and robust solution needed to address the challenges posed by multiphysics-multiscale problems in magnetic fusion energy simulations. These methods need to be optimized based on detailed criteria induced by accuracy and stability restrictions and accompanied by robust error estimators. Developing time-stepping algorithms in conjunction with high-performance scientific computing libraries such as PETSc [3, 2, 4] yields the most robust overall strategy because much of the partitioned schemes can be bundled up within the library, therefore reducing the complexity of the application codes. This allows efficient data movement, memory management, detailed control over convergence tolerances, and step rejection recourse. In this context a fair comparison of time to solution of a given accuracy among different integration strategies is missing in the literature. We believe that a simple FES test case can be used to highlight the advantages of partitioned methods.

References

[1] U. Ascher, S. Ruuth, and R. Spiteri. Implicit-explicit Runge-Kutta methods for time-dependent partial differential equations. *Applied Numerical Mathematics*, 25:151–167, 1997.

[2] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, K. Rupp, B. F. Smith, and H. Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 3.5, Argonne National Laboratory, 2014.

[3] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, K. Rupp, B. F. Smith, and H. Zhang. PETSc Web page. `http://www.mcs.anl.gov/petsc`, 2014.

[4] S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith. Efficient management of parallelism in object oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhäuser Press, 1997.

[5] E. Constantinescu. Estimating global errors in time stepping. *Submitted to SIAM Journal of Numerical Analysis (SINUM)*, 2015.

[6] E. Constantinescu and A. Sandu. Extrapolated implicit-explicit time stepping. *SIAM Journal on Scientific Computing*, 31(6):4452–4477, 2010.

[7] E. Constantinescu and A. Sandu. Optimal explicit strong-stability-preserving general linear methods. *SIAM Journal on Scientific Computing*, 32(5):3130–3150, 2010.

[8] E. Constantinescu and A. Sandu. Extrapolated multirate methods for differential equations with multiple time scales. *Journal of Scientific Computing*, 56(1):28–44, 2013.

[9] C. Kennedy and M. Carpenter. Additive Runge-Kutta schemes for convection-diffusion-reaction equations. *Appl. Numer. Math.*, 44(1-2):139–181, 2003.

[10] D. E. Keyes, L. C. McInnes, C. Woodward, W. Gropp, E. Myra, M. Pernice, J. Bell, J. Brown, A. Clo, J. Connors, E. Constantinescu, D. Estep, K. Evans, C. Farhat, A. Hakim, G. Hammond, G. Hansen, J. Hill, T. Isaac, X. Jiao, K. Jordan, D. Kaushik, E. Kaxiras, A. Koniges, K. Lee, A. Lott, Q. Lu, J. Magerlein, R. Maxwell, M. McCourt, M. Mehl, R. Pawlowski, A. P. Randles, D. Reynolds, B. Rivière, U. Rüde, T. Scheibe, J. Shadid, B. Sheehan, M. Shephard, A. Siegel, B. Smith, X. Tang, C. Wilson, and B. Wohlmuth. Multiphysics simulations: Challenges and opportunities. *International Journal of High Performance Computing Applications*, 27(1):4–83, Feb 2013.

[11] D. Knoll, L. Chacon, L. Margolin, and V. Mousseau. On balanced approximations for time integration of multiple time scale systems. *Journal of Computational Physics*, 185(2):583–611, 2003.

[12] J. Reisner, V. Mousseau, A. Wyszogrodzki, and D. Knoll. An implicitly balanced hurricane model with physics-based preconditioning. *Monthly weather review*, 133(4):1003–1022, 2005.

[13] D. Ropp, J. Shadid, and C. Ober. Studies of the accuracy of time integration methods for reaction–diffusion equations. *Journal of Computational Physics*, 194(2):544–574, 2004.

[14] G. Strang. On the construction and comparison of difference schemes. *SIAM Journal on Numerical Analysis*, 5(3):506–517, 1968.

---

[*] Topic of the whitepaper: D. Multiphysics and multiscale coupling.

[†] corresponding author