A Considered Approach to Multiphysics Problems at the Exascale: Coupled Until Proven Decoupled

David E. Keyes[1], Lois Curfman McInnes[2], Carol S. Woodward[3], William D. Gropp[4], Eric S. Myra[5]
with grateful acknowledgment to all participants of the 2011 ICiS Multiphysics Workshop

April 2013

Due to increasing demands that simulations capture all relevant influences on a system of interest, including those that belong in a feedback loop with that system but might otherwise be modeled exogenously, multiphysics simulations are becoming essential for predictive science and continue to be a primary motivation for exascale computing [1,2]. As stated in the 2010 ASCAC exascale report [1], "The great frontier of computational physics and engineering is in the challenge posed by high-fidelity simulations of real-world systems, that is, in truly transforming computational science into a fully predictive science. Real-world systems are typically characterized by multiple, interacting physical processes ('multiphysics'), interactions that occur on a wide range of both temporal and spatial scales."

However, the promise of coupled multiphysics simulations will be not realized effectively in extreme-scale computational environments in the principal way by which individual codes are coupled today, namely, through unmonitored divide-and-conquer operator splitting. Coupling individual simulations may introduce limitations on stability, accuracy, or robustness that are more severe than the limitations imposed by the individual components [3]. Furthermore, the data motion and data-structure conversions required to iterate between independent simulations for each component may be more costly in latency and electrical power than those of the individually tuned components. Thus, ''one plus one'' may cost significantly more than ''two'' and may be less amenable to scalable execution than expected.

Progress in verified and validated multiphysics simulations today and in extreme-scale multiphysics simulations in the future depends upon: (1) mathematical analyses that are applied dynamically to adapt the allowable degree of operator splitting and (2) hybrid mathematical-computational performance analyses to trade-off the cost of data motion against the predicted degradation of accuracy from economizing on data volume and synchronization induced by its exchange. This philosophy goes to the heart of engineering optimization, in which the work invested in each phase should be just enough (but no more) to guarantee that the overall performance is limited only by the most difficult phase to control. Our key proposals for extreme-scale multiphysics simulation are: (1) new computable metrics for determining the strength of coupling between operators in a multiphysics model must be developed; (2) algorithms that adapt accuracy controls, such as time step size and solution approach (explicit, implicit, etc.), of each physics component and the accuracy controls on the coupling between them, such as splitting method, must dynamically adapt to the complexities and costs of high-fidelity simulations; and (3) traditional flops-based cost metrics must be brought up to date by the new dominance of communication, including all types of memory-to-memory or node-to-node copying in terms of energy and time.

As examples, consider radiation hydrodynamics (RHD) and climate simulation. RHD models a system in which the evolution of radiation is coupled to the hydrodynamic evolution of a material medium, and it forms a vital area of study in astrophysics, high-energy-density physics, inertial-confinement fusion, and high-energy explosive phenomena. Coupling between radiation and matter is among the most significant challenges facing high-performance computation in RHD [4]. Coupling in leading-edge RHD problems varies not just quantitatively in space and time, but also qualitatively: diverse physical processes

---

[1] KAUST and Columbia University, kd2112@columbia.edu
[2] Argonne National Laboratory, mcinnes@mcs.anl.gov, corresponding author
[3] Lawrence Livermore National Laboratory, woodward6@llnl.gov
[4] University of Illinois at Urbana-Champaign, wgropp@illinois.edu
[5] University of Michigan, emyra@umich.edu

determine the nature of the coupling [5]. Hence, changes in dominant processes can result in changes in the dominant terms in the evolution equations as a calculation proceeds. Practical and historical considerations have almost always led to an operator-split approach, with the material component evolved explicitly and the radiation implicitly [5]—often through use of, what were originally, independent codes. Although this procedure is probably adequate in many situations, it cannot not be fully justified in general. For exascale RHD simulations, in particular, the long time scales inherent in many target problems, combined with potential splitting errors, are likely to lead to significant uncertainty in solution accuracy after many time steps. Similarly, climate simulations include numerous parameterizations for various sub-grid processes, including precipitation, cloud formation, radiation, and convective processes [6]. These processes have dependences on each other, and, as grid sizes are reduced, the order of computation of these processes can change results. Like RHD, for exascale problems where long time scales and highly complex models will be applied, significant uncertainty may arise.

Advocating an examination of coupling strength before pursuing any particular approach for solving a multiphysics problem, we propose ‘‘coupled until proven decoupled’’ as a perspective worthy of 21st-century simulation purposes and resources. The situations outlined above mandate that at first approach to a coupled multiphysics problem, one should conduct a full study of the coupling regimes of importance and the strength of coupling within these regimes prior to algorithm selection. Once a solid understanding of the strength of coupling is acquired, efficient algorithms can be targeted. To develop this understanding, rigorous estimates derivable for two-physics linear problems, as in [7], where computable bounds on functionals of interest are related to the norms of off-diagonal blocks, must be generalized and developed into computable metrics for strength of coupling.

At the extreme scale, a multifaceted and adaptive solution approach across spatial and temporal domains will be necessary. Using computable error bounds and strength of coupling metrics, adaptive algorithms must be developed to respect strongly coupled regimes yet exploit well-separated regimes for efficiency. These algorithms should leverage work in IMEX as well as implicit, semi-implicit, standard splitting, and explicit temporal discretizations. In addition, spatial discretizations, such as mortar methods and adaptive meshing techniques, must be extended for multiphysics problems. The main challenges are exploiting separation of scales where possible while still maintaining stability and consistency of the entire solution approach. Of necessity at the exascale, these adaptive methods must be developed in tandem with new architectures so that data structures at the heart of algorithm implementations can be as efficient as possible. The move toward more tightly coupled multiphysics approaches looms as a prime opportunity for adaptation to the exascale, with its memory bandwidth stringencies [8,9,10].

A recent report [4], which incorporates perspectives of applied mathematicians, computer scientists, and domain scientists, advocates that we must fundamentally rethink approaches to multiphysics modeling, algorithms, and solvers with attention to issues of data motion, data structure conversions, and overall software design. *We need sustained investment in research on the mathematical analysis necessary to ensure that splitting and coupling schemes are accurate, stable, robust, consistent, and implemented correctly.* Key research topics include: development of metrics for strength of coupling, formal analysis of splitting methods, development of multiphysics model problems on which splitting and adaptive methods can be tested, techniques for dynamic adaptation of coupling approaches as simulations progress, methodologies and software for coupled codes, approximations that are both compact in memory and mapped to efficient data structures that minimize transformations needed between them, reduction of data transfers, and methods for relaxation of synchrony. Programming paradigms and mathematics both need to be revisited, with attention to less-synchronous algorithms employing work stealing, so that different physics components can complement each other in cycle and resource scavenging without interference.

Research of this rich interplay between applied mathematics and practical computer science constraints is essential to exploit exascale resources to achieve the vision of predictive multiphysics science.

References:

[1] R. Rosner et al., *The opportunities and challenges of exascale computing*. Advanced Scientific Computing Advisory Committee (ASCAC) Subcommittee on Exascale Computing, Office of Science, U.S. Department of Energy, 2010.

[2] D. Brown, P. Messina, et al., *Scientific grand challenges: Crosscutting technologies for computing at the exascale*. Office of Science, U.S. Department of Energy, 2010.

[3] D. Estep, V. Ginting, D. Ropp, J. N. Shadid, and S. Tavener, An a posteriori–a priori analysis of multiscale operator splitting, *SIAM Journal on Numerical Analysis* 46(3): 1116–1146, 2008.

[4] D. E. Keyes, L. C. McInnes, C. Woodward, W. D. 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. Riviere, 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: 4–83, 2013.

[5] J. Castor, *Radiation Hydrodynamics*, Cambridge University Press, Cambridge, 2004.

[6] A. C. M. Beljaars, Numerical schemes for parametrizations, in *Seminar on Numerical Methods in Atmospheric Models*, vol II: 1-42, 1991.

[7] J. P. Whiteley, K. Gillow, S. J. Tavener, and A C. Walter, Error bounds on block Gauss–Seidel solutions of coupled multiphysics problems, *International Journal of Numerical Methods in Engineering* 88(12): 1219–1237, 2011.

[8] D. Brown et al., *Applied mathematics at the U.S. Department of Energy: Past, present, and a view to the future*. Office of Science, U.S. Department of Energy, 2008.

[9] J. Ang, K. Evans, A. Geist, M. Heroux, P. Hovland, O. Marques, L.C. McInnes, E. Ng, S. Wild, *Report on the workshop on extreme-scale solvers: Transitions to future architectures*, Office of Advanced Scientific Computing Research, U.S. Department of Energy, 2012.

[10] J. Dongarra, P. Beckman, et al., The International Exascale Software Project roadmap. *International Journal of High Performance Computing Applications* 25: 3–60, 2011.