

Coupling of hierarchical multiphysics models and other mathematical issues in extreme scale computing

Peter Graf*, Ray Grout, Michael A. Sprague

Computational Science Center, National Renewable Energy Laboratory

Conceptions of extreme scale computing and their implications

Conceptions of extreme-scale simulation appear to come in two categories. The first is to use highly scalable codes based on direct numerical simulation (DNS) of a consistently derived mathematical model where relevant scales are captured; there are no subgrid or physics parameterizations. Here, grid refinement can lead to an exascale problem. DNS in fluid dynamics epitomizes this category. The challenge is to solve at sufficiently fine scale to resolve the relevant phenomena. The second category is to use multi-physics models that are an ad-hoc patchwork of single-physics models. Examples in this category include multi-scale/multi-physics battery simulations [8] and multi-turbine wind farm simulations that incorporate numerical weather prediction [2, 11] and turbine structural response [7].

We submit here that: (1) for the second category, basic issues such as convergence remain problematic to establish and, more fundamentally, there is an upper limit to the accuracy of this approach even with near-infinite computing capability; (2) problems of the first category tend to be “voracious” [1] computing problems, i.e., they are intractable at *any* scale; and (3) a practical approach to problems in the first category involves converting them to models of the second category, with its attendant problems. This discussion allows us to suggest the form extreme-scale computations must take.

First, in our experience, multi-physics models are typically based on physical insight used to formulate ad hoc representations of the physics necessary to represent the phenomena of interest. The potential for errors, inconsistencies and far-from-optimal convergence rates introduced through ad hoc coupling are unnerving. While convergence and other properties may be established for the individual single-physics models, this by no means guarantees convergence of the coupled systems. *Convergence in practice should not be used as de facto justification of methods that are not otherwise understood.* The mathematics of multi-scale multi-physics coupling is incomplete and should be further emphasized.

But ad hoc multi-physics models have a more fundamental problem. The scientific questions often motivate using ever-increasing computational power to presumably deliver ever-increasing accuracy, but this can be a false hope. The problem is that the single-physics submodels are designed to work over a certain range of scales before the bounds of the validity of the model are exceeded. Further refinement does not improve the solution but rather invalidates the basic assumptions made in formulating the model. Examples of this are numerous, such as turbulence models in LES calculations that require scale similarity between the subgrid and filtered fields and battery cell models that assume the existence of discrete electrodes and electrolytes within the domain over which the model is applied. *It is particularly troubling that many models have this implicit dependence on the exact level of space/time refinement.* For these models, the naive approach to scaling by refining grids or studying empirical convergence properties hits a road block that requires rethinking of our approach.

Unfortunately, the DNS approach has a significant problem as well. Because of the stiffness of the full multi-scale problem, the approach has a “voracious” aspect, where no amount of computational resources will suffice to capture scales of engineering relevance starting from first principles[†]. In practice, intractable DNS problems are made tractable by dividing the model into models of different scales, e.g., to break the stiffness. However, this results in separate models at different scales that must then be coupled. That is, the standard approach to the main difficulty encountered in the first category of extreme scale problems is to convert them to the second category, thus inheriting *its* problems.

The road ahead is not clear. On one hand, DNS models are problematic, even at exascale. On the other, multi-physics models would seem to have inherent limits to their scalability. We have no easy solution to offer. But this discussion leads to an approach roughly as follows: We need to accept the

* Corresponding author: peter.graf@nrel.gov [†] Actually, even if could take an unlimited number of time steps, fixed-precision calculations would suffer from roundoff errors that could still make the computation meaningless

inevitability of using models that must be coupled. We require careful analysis of the implications of the coupling on convergence rates and accuracy. We must design such simulations using a hierarchy of models with carefully quantified regimes of appropriateness to restore the “continuous refinement towards exact physical solution” property that the voracious computing formulations exhibit. This will enable us to combine models in a way that allows for predictable, extreme scaling.

Other issues deserving a “place at the extreme-scale table”

High-order methods: We believe spectral finite elements (SFEs) deserve more attention for extreme-scale computing. Here, SFEs are high-order finite elements with Lagrangian interpolant basis functions (with nodes at the Lobatto points) [9]. SFE methods are attractive for extreme-scale computing because, for a given level of accuracy, SFE models can be orders of magnitude smaller (in terms of degrees-of-freedom) compared to low-order methods. SFE methods are also attractive because they benefit greatly from tensor-product factorization in matrix-free evaluation of matrix-vector products. Further reduction in operations is attained when SFE weak forms are evaluated with nodal quadrature. SFEs can be orders of magnitude faster than low-order methods for a given level of accuracy. SFEs have shown great promise, for example, on peta-scale simulation [5] of incompressible flow, where elements are connected conformally. In DNS for turbulent flow, SFEs have shown an order-of-magnitude improvement in wall-clock time when compared to a low-order finite-volume solver [10]. However, looking forward to extreme-scale computing, there remain questions, foremost whether the data-structures required for SFEs will be suitable for efficient computation on future hardware.

Time-parallelization via exploiting short memory and its use in UQ studies: We can parallelize in time when forcing in systems causes rapid loss of dependence on initial conditions. For example, in a building energy simulation, the energy use in July depends on the weather in July, not the energy use in January. Thus one can divide a year of simulation into independent short periods. We start each period from rough initial guesses *before* the true start of each target period so that initial conditions are “forgotten” by the true start of the target period.

This method is inherently useful in uncertainty quantification as long as the uncertainty in the forcing is greater than the uncertainty introduced by the time-parallelization. To verify this, we can perform time-series analysis within the time periods by examining the autocorrelation of individual system variables. This results in estimates of the length of the system memory, which can be used in an online manner to tune the length of “warm up” periods. The mathematical challenge here is the interplay between the strength and uncertainty of the external forcing (e.g., weather) and the system’s memory of its initial conditions (e.g., current energy use) in the general setting of parallel uncertainty quantification of time-dependent simulations with external stochastic forcing. Also of interest is combining this “memory decay” approach with iterative time-parallelization methods [4].

Surrogate-based optimization for discrete and mixed integer problems: In surrogate-based optimization, search steps using approximations alternate with corrections by occasional higher fidelity calculations. This general scheme can be applied in a wide variety of settings. Convergence has been proved for the continuous case (e.g., [3]). However, there has been insufficient emphasis on the case of discrete design variables. For example, we have used such an approach in the context of electronic structure optimization [6]. Here, we compute the full eigenbasis for a given atomic configuration, then select a subset of the eigenvectors and project the Schroedinger equation on this basis. The reduced basis is used to solve (much more quickly) several additional candidate atomic configurations. When the configuration gets sufficiently far from the one used to create the reduced basis, we again solve the full system and repeat the process. Of particular interest is the tradeoff (w.r.t. computational resources, convergence, accuracy, etc.) between sampling (i.e. to build better models) and search (i.e. to find the optimum), especially in the context of nonconvex global problems, where the additional tradeoff between local and global search must be considered.

References

- [1] Ashby et al. The opportunities and challenges of exascale computing. Technical report, AS-CAC Subcommittee on Exascale Computing, 2010. Summary Report of the Advanced Scientific Computing Advisory Committee (ASCAC) Subcommittee.
- [2] M. J. Churchfield, S. Lee, P. J. Moriarty, L. A. Martinez, S. Leonardi, G. Vijayakumar, and J. G. Brasseur. A large-eddy simulation of wind-plant aerodynamics. In *50th AIAA Aerospace Sciences Meeting*, number AIAA-2012-537, 2012.
- [3] M. S. Eldred and D. M. Dunlavy. Formulations for surrogate-based optimization with data fit, multifidelity, and reduced-order models. In *Proceedings of 11th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference AIAA*, 2006.
- [4] M. Emmett and M. L. Minion. Toward an efficient parallel in time method for partial differential equations. *Comm. App. Math. and Comp. Sci.*, 7(1):105–132, 2012.
- [5] P. Fischer, J. Lottes, D. Pointer, and A. Siegel. Petascale algorithms for reactor hydrodynamics. *Journal of Physics: Conference Series*, 125:012076, 2008.
- [6] P. A. Graf and W. B. Jones. A projection based multiscale optimization method for eigenvalue problems. *J. of Global Optimization*, 39(2):235–245, October 2007.
- [7] J. M. Jonkman. The new modularization framework for the FAST wind turbine CAE tool. In *Proceedings of the 32nd ASME Wind Energy Symposium, 51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, 2012.
- [8] G. H. Kim, K. Smith, K. J. Lee, S. Santhanagopalan, and A. Pesaran. Multi-domain modeling of lithium-ion batteries encompassing multi-physics in varied length scales. *Journal of The Electrochemical Society*, 158:A955–A969, 2011.
- [9] A. T. Patera. A spectral element method for fluid dynamics: Laminar flow in a channel expansion. *Journal of Computational Physics*, 54:468–488, 1984.
- [10] M. A. Sprague, M. J. Churchfield, A. Purkayastha, P. J. Moriarty, and S. Lee. A comparison of nek5000 and openfoam for dns of turbulent channel flow, 2010. Presented at the Nek5000 Users meeting; http://www.mcs.anl.gov/~fischer/nek5000/sprague_nek5000_dec2010.pdf.
- [11] M. A. Sprague, P. J. Moriarty, M. J. Churchfield, K. Gruchalla, S. Lee, J.K. Lundquist, J. Michalakes, and A. Purkayastha. Computational modeling of wind-plant aerodynamics. In *Proceedings of SciDAC 2011*, 2011.