

Moment-based scale-bridging algorithms for multiphysics kinetic simulations at the exascale. D. A. Knoll, L. Chacón (LANL), C. T. Kelley (NCSU), K. Smith (MIT)

The DoE will invest significant resources to bring exascale computing on line. Exascale computing will likely employ hierarchical, heterogeneous computing architectures, featuring multiple levels of parallelism, each delivered by different architectural solutions. It is paramount for DOE that this hardware is utilized to its fullest extent. Achieving high utilization levels will require the development of a new generation of computational solver algorithms. The most successful of these new algorithms will simultaneously 1) embrace emerging architecture characteristics so as to optimally utilize the compute capability and do so in a resilient fashion, and 2) ensure that a demonstrable advancement in predictive simulation capability is achieved in key application areas of importance to DoE and the nation. *This position paper argues for further development of moment-based scale-bridging algorithms which not only adapt, but in fact embrace, the types of heterogeneous architectures that will underpin upcoming exascale computers.*

The thesis of this position paper is that, for an algorithm to exploit hierarchical exascale computing successfully, it must become hierarchical itself. A hierarchical algorithm is one that can exploit different levels of physical description by mapping them to corresponding levels in the computing hierarchy. In doing so, each level of parallelism deals with the physical description level best suited to its architectural specifications. Kinetic transport systems (Boltzmann-like models) feature a natural hierarchical formulation based on moments. Moments are rigorously obtained by weighed integration of the distribution function over velocity, energy, and/or other parameters describing phase space. Thus, moments are functions of configuration space only. Governing equations for the moments can be obtained by successive integration of the kinetic transport equation. Moment equations, however, require closures for well-posedness. It is important to understand that moment equations involve no physical approximations when closed self-consistently with the kinetic model.

The target applications for this evolving algorithmic approach are multiphysics kinetic systems. Multiphysics kinetic systems are systems that 1) consider a variety of physical models that are interacting in a strong nonlinear fashion, and 2) at least one of the required physical components demands a kinetic treatment for physical fidelity, with the solution of a corresponding transport equation in a multiple dimensional phase space (e.g., three spatial dimensions plus additional dimensions describing velocity or energy space). The potential application space modeled by multiphysics kinetic systems is very large, including a variety of plasmas (space weather, magnetic fusion), high-energy-density physics and ICF applications, engineering combustion systems, climate and weather systems, multiphysics nuclear reactor simulation, materials science and many others. In all these problems, it is the need for a phase-space representation of kinetic physics components which dominates the computational effort (>80-90%). *In the vast majority of such systems, the kinetic components are currently approximated in some fashion to make system-scale simulations tractable. Such approximations are frequently the weak link in increased predictive simulation.* This position paper argues for a novel algorithmic approach that builds upon an evolving foundation to move past such approximations (to achieve physics fidelity) while simultaneously embracing emerging exascale computer architectures. This is a classic “win-win” situation, which will ensure the advancement of predictive simulation capability while effectively utilizing exascale computer architectures. Our research team has

already demonstrated a proof of principle of such a “win-win”.

We use the phrase “high-order” (HO) to describe the kinetic transport system, and “low order” (LO) to describe the consistent moment system. Moment-based scale-bridging algorithms utilize a hierarchical kinetic formulation that features LO and HO descriptions collocated everywhere in the domain. Both descriptions are coupled via closures and fields. The original purpose of considering the LO description [1-3] was as a means for algorithmic acceleration via time-scale segregation (akin to a multigrid approach), which in turn enabled one to address stiff time scales at the LO component level (which is far more amenable to physics-based preconditioning [4-10] than the HO component). The foundations of algorithmic moment-based scale-bridging ideas reside in statistical mechanics, kinetic theory, and transport theory. These ideas have matured in steady-state neutron transport for nuclear reactors [11-14]. The extension from neutron transport to kinetic plasma simulation is natural with the inclusion Maxwell’s equations. Early related ideas can be found in [2-3], and our recent advances are presented in [15-19]. We have also made significant recent advances in applying these ideas to thermal radiation transport [20-22].

Solver-wise, HO-LO consistency and the strict preservation of conservation laws (which determine long-term accuracy) demand tight nonlinear coupling. This, in turn, will provide some important benefits such as robustness, numerical stability, and superior accuracy. However, tight coupling requires practical, scalable nonlinear solution strategies, which overcome key obstacles such as onerous memory requirements. Significant progress towards the practicality of fully nonlinear approaches for kinetic multiphysics systems has recently been made. We currently have two related, yet somewhat distinct approaches. One is based on a standard Picard iteration between the HO and LO problems [1,11,13,17], and could be easily adapted to use Anderson acceleration techniques [23]. A different approach is to solve the HO problem directly with JFNK through the use of kinetic enslavement (or nonlinear elimination) [12,15]. Here one can use the LO problem as the preconditioner and achieve significant acceleration [12,18].

With the push towards exascale, the hierarchical nature of the HO-LO approach possesses another important advantage: an efficient mapping of the algorithm to heterogeneous computer architectures. In particular, LO (moment) equations are low-dimensional, and hence memory friendly. They are best suited for the highest level of parallelism. The HO (kinetic) description is high-dimensional and very compute intensive, and hence very suitable to lower (less reliable) levels of parallelism. Some key advantages of HO-LO formulations for exascale computing are: **Minimization of data movement**: the HO component can be completely contained at the appropriate level of the architecture (the hardware accelerator). Particles “live” in the accelerator (or isolated on nodes), and are never communicated elsewhere. Only (low dimensional) moments and fields are communicated up and down the architecture hierarchy. **Resiliency**: The redundant nature of the HO-LO approach provides a first firewall against soft and hard hardware failures. Compute-intensive HO descriptions, which will live in the less reliable architectures, may be best served with a Lagrangian description based on particles, which is naturally resilient against soft failures [24]. **Mixed precision**: HO computations can be done in single precision to minimize memory storage and maximize computational efficiency, while LO evolution can be done in double precision [16]. **Asynchrony**: a nonlinearly enslaved HO-LO description allows one to exploit it at a very coarse algorithmic level because all HO descriptions can be enslaved to the same LO formulation. Thus, several HO models representing different coupled physics do not interact with one another, and therefore can be integrated concurrently.

REFERENCES:

- [1] Gol'din, V. Ya., USSR Comput. Math. Math. Phys., 4:136, 1967.
- [2] R.J. Mason, Implicit moment particle simulation of plasmas, J. Comput. Phys., 41(2): 233–244, 1981.
- [3] J. U. Brackbill and D. W. Forslund. An implicit method for electromagnetic plasma simulation in two dimensions., J. Comput. Phys., 46:271, 1982.
- [4] L. Chacón. An optimal, parallel, fully implicit Newton-Krylov solver for three-dimensional visco-resistive magnetohydrodynamics. Phys. Plasmas, 15:056103, 2008.
- [5] L. Chacón and D. A. Knoll. A 2D high- β Hall MHD implicit nonlinear solver. J. Comput. Phys., 188(2):573–592, 2003.
- [6] L. Chacón, D. A. Knoll, and J. M. Finn. Implicit, nonlinear reduced resistive MHD nonlinear solver. J. Comput. Phys., 178(1):15–36, 2002.
- [7] D A Knoll and D E Keyes. Jacobian-Free Newton Krylov methods: A survey of approaches and applications. J. Comp. Phys., 193:357–397, 2004.
- [8] V.A. Mousseau D A Knoll and W.J. Rider. Physics-based preconditioning and the Newton-Krylov method for non-equilibrium radiation diffusion. J. Comp. Phys., 160:743-765, 2000.
- [9] J.M. Reisner, V.A. Mousseau, A.A. Wyszogrodzki, and D A Knoll. An implicitly balanced hurricane model with physics-based preconditioning. Mon. Wea. Rev., 133:1003-1022, 2005.
- [10] H. Park, R. R. Norgaliev, R.C. Martineau and D A Knoll. On physics-based preconditioning and the Navier-Stokes equations. J. Comp. Phys., 228:9131-9146, 2009.
- [11] K.S. Smith and J.D. Rhodes III. Full-core, 2-D LWR core calculations with CASMO-4E. PHYSOR 2002, Seoul, Korea, 2002.
- [12] D. A. Knoll, H. Park, and K. Smith. Application of the Jacobian-free Newton-Krylov method to nonlinear acceleration of transport source iteration in slab geometry. Nucl. Sci. Eng., 167:122–132, 2011.
- [13] H. Park, D.A. Knoll, and C. K. Newman. Nonlinear acceleration of transport criticality problems. Nucl. Sci. Eng., 172:52–65, 2012.
- [14] Jeff Willert, C T Kelley, D A Knoll, and H K Park. Hybrid Deterministic/Monte Carlo Neutronics, SIAM J. Sci. Comput., 2013 (accepted)
- [15] G. Chen, L. Chacón, and D. C. Barnes. An energy- and charge-conserving, implicit, electrostatic particle-in-cell algorithm. J. Comput. Phys., 230:7018–7036, 2011.
- [16] G. Chen, L. Chacón, and D. C. Barnes. An efficient mixed-precision, hybrid CPU-GPU implementation of a fully implicit particle-in-cell algorithm. J. Comput. Phys., 231(16):5374–5388, 2012.
- [17] W. Taitano, D. A. Knoll, L. Chacón, and G. Chen. Development of a consistent and stable fully implicit moment method for Vlasov-Ampere particle-in-cell (PIC) system. SIAM J. Sci. Comput., 2013. (accepted).
- [18] G. Chen and L. Chacón. Moment preconditioning for a nonlinear kinetic particle-in-cell solver. J. Comput. Phys., 2013. In preparation.
- [19] J. Payne, D.A. Knoll, A. McPherson, W. Taitano, L. Chacon, G. Chen, S. Pakin, Computational Co-Design for a Multiscale Plasma Application: A Process and Initial Results, LA-UR-13-22909 (submitted to SC13)
- [20] H. Park, D. A. Knoll, R. M. Rauenzahn, J. D. Densmore, and A. B. Wollaber. A consistent, moment-based, multiscale-solution approach for thermal radiative transfer problems. Transport. Theo. Stat. Phys., 41:284-303, 2012.
- [21] H. Park, D. A. Knoll, R. M. Rauenzahn, C. K. Newman, J. D. Densmore, and A. B. Wollaber. An efficient and time accurate, moment-based scale-bridging algorithm for thermal

radiative transfer problems. *SIAM J. Sci. Comput.*, 2013. (accepted).

[22] H. Park, J. D. Densmore, A. B. Wollaber, D. A. Knoll, and R. M. Rauenzahn. Monte Carlo solution methods in a moment-based scale-bridging algorithm for thermal radiative transfer problems: comparison with Fleck and Cummings. *Int. Conf. on Math. and Comp. Meth. Applied to Nuclear Science and Engineering.*, Sun Valley, ID, May 5-9, 2013.

[23] Donald G. Anderson. Iterative Procedures for Nonlinear Integral Equations. *Journal of the ACM*, 12(4):547–560, October 1965

[24] S. Mukherjee. *Architecture design for soft errors*. Morgan Kaufmann, 2008.