

Algorithms for Exascale Computational Mesoscience

Dmitry Karpeev, Juan de Pablo, Olle Heinonen

Background. Mesoscopic phenomena refer to dynamic states of matter composed of persistent excitations residing between the atomistic and engineering (macroscopic) scales.

Spatial and temporal averaging at this scale are sufficiently strong so that the state of the material can be accurately described by continuum *phases*. The behavior of these phases (e.g., magnetization or polarization fields) is governed by PDEs, the so-called *phasefield* models. The averaging is too weak, however, to overcome *fluctuations*, which result in spontaneous rare events (e.g., magnetic domain flipping) and are modeled as noise of appropriate kind. Fluctuations cause a violation of local charge neutrality resulting in long-range electrostatic interactions, leading to locally correlated, but globally heterogeneous evolving structures – *domains* and *interfaces* (e.g., magnetic domains and domain walls). Mesoscopic theories have been very useful in describing complex phenomena and predicting their macroscopic consequences in systems as diverse as high-temperature superconductors, liquid crystals, ferroelectric materials and complex fluids. Computational mesoscience is increasingly important, but for it to be viable at exascale its inherent algorithmic challenges must be addressed.

The presence of (i) strong heterogeneities, and (ii) strong fluctuations sets computational mesoscience apart from the engineering PDE models, which are largely based on smoothly varying constitutive relations (see, e.g., [1]). Material heterogeneity produces many local minima in the energy landscape, and fluctuations replace simple relaxation by the relatively rare transitions from one metastable state to the next. This introduces significant *serialization* in the time-stepping or minimization algorithms, which spend most of the time integrating fluctuations about a local minimum. Furthermore, in order for the implicit solvers to be effective at maximizing the stable timestep size of the dissipative dynamics driving mesoscopic, optimal multilevel preconditioners (see, e.g., [2]). In the presence of strong heterogeneities the construction of effective coarse-grid problems becomes expensive and can be highly nontrivial.

How should these algorithmic challenges be addressed in the context of future exascale architectures? These architectures are expected to require increasingly finely-grained parallelism, algorithmic fault-tolerance, minimization of communication, all in the context of a shrinking relative memory size and bandwidth. What are the exascale-optimal multilevel preconditioner constructions for *mesoscale* models with their strong and evolving heterogeneities? How should long-range interactions be handled? What are the best methods to capture the fluctuation-induced rare events? How can the inherently sequential time integration take advantage of the parallelism of the modern architectures and the exascale hardware? All of these questions have to be addressed in concert in the context of specific requirements of computational mesoscience.

Position. It is our position that an effective approach to addressing the above challenges can be based on (i) the *statistical* nature of mesoscale models, (ii) their *homogenization* properties, together with (iii) modern *statistical learning* methods.

Ensemble simulation. Because the structure of material heterogeneities and fluctuations can only be known probabilistically, only statistical properties of the system can be reliably observed and computed. This suggests probing the system using *ensembles* of simulations to calculate the *statistics* of metastable states and transitions between them, increasing available parallelism. The

essential idea is that progress towards the next transition is accelerated if the current metastable state is sampled *in parallel* by many realization of the current quasistationary distribution. This is the approach that was described in [3] and analyzed mathematically in [4]. A similar idea was proposed earlier in the context of Monte-Carlo relaxation [5], where realizations of multiple nearby ensembles were used in concert to accelerate convergence. First a set of N *independent* realizations of the quasistationary distribution about the current metastable state is prepared. The simulation can then proceed in parallel until the next transition, sampling the transition events at an N -fold speedup and affording N -fold parallelism. This offers a potentially critical advantage on modern architectures where the improvement in the single thread performance may not be as significant as the increase in hardware thread parallelism.

Homegenization. The construction and application of the coarse-grained operators -- the coarse-grid problem and the far-field component of long-range interactions -- can remain a substantial bottleneck in mesoscale simulations. The heterogeneity in mesoscale systems, however, frequently exhibits a long-range ordering -- its statistics are stationary¹ with respect to macroscopic spatial translations. Then, as observed in [6], grid coarsening results in the *homogenization* of stochasticity, similar and related to charge *screening*: the coarse-grained operators are essentially deterministic, insensitive to ensemble realization exchange. This suggests that while the system remains in the same quasistationary state, the coarse-grained operators can be frozen and need to be rebuilt only after the transition. Furthermore, they can be inexpensively reused among the different realizations of the same or nearby ensembles, and even applied to multiple coarse-grained states at once, potentially alleviating the memory bandwidth limitations.

Statistical learning. Long-time integration of mesoscale systems can be rather expensive for an exhaustive probing of their parameter spaces. In the context of uncertainty quantification this is sometimes addressed by replacing the simulation with a much simpler *surrogate model*, which approximates the original behavior learned from a small number of sample simulations. Success of this approach requires factoring out (e.g., via polynomial chaos expansion) significant correlations in the heterostructure, leaving a small number of stochastic parameters to sample. While the strong disorder of mesoscale systems precludes this, we propose to apply this method to the much smoother coarse-grained problem. In addition to accelerating coarse-grained operator construction, surrogate models can improve the resilience of exascale mesoscopic simulations, simplifying their checkpointing and the construction of good initial metastable states. An important role in the construction of surrogate interpolants may be played by numerical *adjoint* calculation, which requires an efficient handling of simulation checkpointing, relating to another important topic of modern computational mesoscience -- *Big Data* management and analysis.

Software. Finally, we feel the research ideas outlined above must be implemented as part of existing solver libraries, such as PETSc [2], and tested on concrete problems. This may require a considerable redesign of the algorithms and datastructures, strongly coupling them to statistical samplers, enabling shared solves across ensembles of simulations and incorporating adjoint calculations in a more fundamental way.

¹ Or *locally* stationary, in the sense that they partition into a small number of approximately stationary subsets (similar to decision trees).

References.

- [1] D. Keyes, et al., *DOE Workshop Report on Multiphysics Simulations: "Multiphysics Simulations: Challenges and Opportunitie"s*, Argonne National Laboratory, [Technical Memorandum ANL/MCS-TM-321](#), 2012.
- [2] Portable Extensible Toolkit for Scientific computing (PETSc), <http://www.mcs.anl.gov/petsc>
- [3] A. Voter, *Parallel replica method for dynamics of infrequent events*, Phys. Rev. B, 57(22):985-988, 1998.
- [4] C. Le Bris, T. Levievre, M. Luskin, D. Perez, *A mathematical formalization of the parallel replica dynamics*, *Monte Carlo Methods and App.* 18(2):119-146
- [5] R. H. Swendsen and J. Wang, *Replica Monte-Carlo simulation of spin glasses*, Phys. Rev. Lett., 57: 2607:2609
- [6] G. C. Papanicolaou, S. R. S. Varadhan, *Diffusion with random coefficients*, Essays in Statistics and Probability, edited by G. Kallianpur, P.R. Krishnaiah and J.K. Ghosh, North Holland, 1982, pp. 253-262.