Scale-Based Parallelization for Turbulent Flow Simulations Enabled by Extreme-Scale Computing

Jeremy A. Templeton, Sandia National Laboratories

Emerging computer architectures present an opportunity to develop new algorithms to simulate turbulent flows that will enhance our ability to gain physical understanding at a drastically reduced computational cost. Realizing this opportunity requires exploiting what has heretofore been the greatest challenge associated with turbulent simulation: the multiscale nature of turbulent flows. Methods taking advantage of the statistical similarity of the multiple scales to distribute computational work have the potential to transform turbulence simulations from using a brute force, tightly-coupled parallelization strategy to one based on sampling and data analysis.

Background. Because turbulence possesses a large bandwidth of energy containing scales, direct numerical simulations (DNS) are extremely costly. Large-eddy simulation (LES) low-pass filters the Navier-Stokes equations so larger energy containing motions are captured while the effect of small-scale eddies is represented via a sub-grid scale (SGS) model, offering potentially significant cost savings[1]. However, for many scientifically and technologically relevant flow configurations, SGS models contain significant model form errors, limiting the utility of the resulting solutions.

Scale-Based Parallelization. Most existing LES/DNS parallelization approaches use spatial decomposition to split a mesh across available processors and update all processors' parts in lock-step. Communication is required after each update and during solution of linear systems spanning the entire mesh. Requiring a high degree of communication and concurrency creates computational issues at extreme scale, as well as incurring significant inefficiency by resolving turbulent fluctuations based on their physical length and time scales rather than their information content. An alternative is to take advantage of the mathematical structure of LES to develop a combined solution method/SGS model enabled by extreme scale computing. The strategy is to parallelize in scale, rather than space or time, using statistical methods which exploit the mathematical connection between LES at larger scales and LES/DNS at smaller scales. As shown in Fig. 1, multiple scales, from small through meso to large, can be assembled both hierarchically, to provide SGS models to larger scale problems, and laterally, to sample a flow configuration arising in a larger scale problem. Data analysis methods would synthesize the result of each simulation into a set of SGS models designed for use at the next largest scale, as indicated by the arrows pointing upwards in Fig. 1. For the approach to work, the smaller scale simulations would need to be conditioned on the larger scale state so they could provide the correct sub-grid fluctuations. Capturing these physics can be augmented by coordinating the small scale samples to explore both the physical parameter space as well as the stochastic parameter space, as suggested in Fig. 1 by the horizontal arrows. Exascale simulation using scale-based parallelization would involve learning what are the important physics in a given flow and adaptively constructing an appropriate SGS model for them.

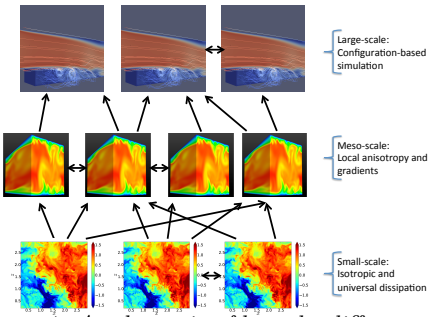

Figure 1. A schematic of how the different simulations at each scale can interact. Algorithms will be needed to exploit the multiscale physics as increasingly small scale behavior becomes increasingly universal.

[1]Deck, Renard, & Laraufie, *Phys. Fluids* 2014

Benefits for Exascale. Scale-based parallelization has the potential to address some significant difficulties anticipated with extreme-scale computing: fault tolerance and data movement. Benefits would be realized by using sampling approaches requiring large numbers of moderately sized computations. Because the sampling structure can be partially laterally aligned across scales, no single computation must successfully complete in order to obtain useful statistics. Further, individual simulations can be restarted on a different set of processors without causing the global simulation to fail, mitigating the impact of hardware failures. Simulations would need only communicate with each to perform data analysis for SGS model construction. As these models require coarse-grained aggregate information only, significant amounts of data movement will not be necessary. With these advantages, a scale-based parallelization is a good fit for exascale architectures.

Benefits for Turbulence Simulations. By decomposing a turbulence simulation into a set of targeted sample simulations, scale-based parallelization could: 1) reduce computational cost, and 2) develop accurate, data-driven SGS models. Compute time would be allocated based on information content, rather than physical extent, of the small scales. In addition, there would be no need to tie SGS model forms and model constants to their present values, which display significant variability from flow to flow. Instead, the samples would be used to formulate and calibrate models on the fly through data analysis of the sample simulations. Results from a proof of concept study of similar methods applied to molecular dynamics simulations of heat transfer[2], shown in Fig. 2, suggest that non-parametric (Gaussian process) constitutive models can be inferred from small-scale samples with accuracy that is both high and quantified, with a cost savings of potentially 500X.

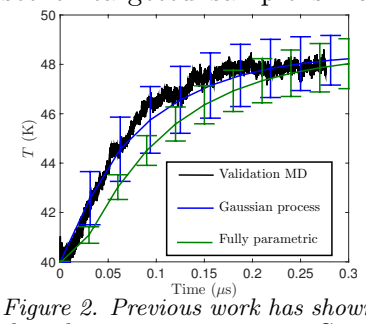

Figure 2. Previous work has shown data-driven non-parametric Gaussian processes can better predict data with uncertainty bounds commensurate with the model error than parametric approaches.

Research Needs. Realizing scale-based parallelization requires development of new technologies in applied mathematics, computer science, and turbulence theory. New importance sampling methods will be needed which can estimate uncertainty in an existing SGS model and evaluate where new samples will add the most information. Similarly, data analysis methods will be needed to synthesize the sample data in SGS models for larger scales, including potentially non-parametric models. If non-parametric SGS models are used, novel linear and non-linear solvers will be needed which can efficiently solve matrix equations lacking much of the current structure provided by existing approaches (*e.g.* matrices arising from diffusion equations when eddy viscosities are used). Providing an infrastructure for scale-based parallelization will require scheduling and prioritization algorithms to optimally allocate hardware resources. Distributed schedulers would be advantageous over a single server to minimize communication and increase robustness. Finally, turbulence modeling theories rigorously connecting small and large scales are required to precisely specify the overall mathematical problem which must be solved. Development is needed of non-parametric models that can satisfy physically-motivated constraints (*e.g.* diffusivity must be positive) such that equations using them as SGS models are solvable, and we must also learn how to incorporate these types of models into turbulent simulation software. In addition, methods to instantiate turbulence simulations conditioned on a larger scale state will be required to enable conditional sampling.

Summary. Scale-based decomposition through coordinated sampling and data analysis offers an alternative to existing approaches for parallel decomposition of turbulent flow simulations at the exascale. Not only can it leverage anticipated computing resources by reducing data movement and increasing fault tolerance, it also offers the possibility of greatly improving the accuracy and efficiency with which we can compute turbulent flows important to many DOE applications.

[2]Salloum & Templeton, *IJUQ* 2014.