Exascale Simulations of Canonical Turbulent Disperse Multiphase Flows

S. Balachandar (bala1s@ufl.edu), T.L. Jackson and P. Fischer
The Center for Compressible Multiphase Turbulence
University of Florida, Gainesville, FL

Our understanding of single-phase turbulence has immensely benefitted from direct and large eddy simulations. Although there remain many open questions, tera and petascale simulations have allowed us to explore the multifractal structure of turbulence and its statistical consequences. We believe exascale computations will allow us to gain similar unprecedented fundamental understanding of multiphase turbulence. The impact will be huge since multiphase turbulence is at the heart of many industrial, environmental and energy applications.

Disperse multiphase turbulence is significantly more complicated than single-phase turbulence in a number of ways. Fluctuations arise not only from classical non-linear turbulence cascade but also from the time-dependent random motion of the dispersed particles (pseudo-turbulence). On one hand, fluid turbulence affects the spatio-temporal distribution of particles to form patterns such as streaks, fingers, blobs and layers through processes such as preferential accumulation, turbophoresis and hindered settling. On the other hand, these patterns locally modify turbulence through mechanisms such as modified inertia, enhanced viscosity, stratification and wake-induced dissipation. Multiphase turbulence in the manageable limit when particles are much smaller than the range of turbulent scales has been widely studied through direct and large eddy simulations. Exascale computing powers offers a unique opportunity to gain similar understanding in the more complex regime of multiphase turbulence with finite-sized particles of diameter larger than the Kolmogorov scale.

We must pose address the following two questions: To advance multiphase turbulence do we need exascale computing (go beyond petascale)? and Is exascale computing sufficient to make significant advancement? Clearly, first-principle simulations (DNS) of real multiphase systems such as fluidized bed reactors, rain formation, and volcanic eruptions are far beyond even exascale computing. As the particle size relative to the turbulent eddies increases, the particle Reynolds number (based on diameter and relative velocity) increases quadratically and quickly reaches $O(10^3)$. Based on which, it can be estimated that with exascale computing we can perform direct numerical simulations of multiphase turbulence consisting of $O(10^5)$ finite-sized particles that strongly interact with the carrier fluid as well as among themselves through collisions. Although this is far short of most real systems, this allows us to study canonical problems such as particle-laden near-wall turbulence, shock propagation over a bed of particles and turbulent entrainment at a stratified particulate interface. In all these problems, it is important to consider at least tens of thousands of particles in order to accurately capture the mechanisms of turbulence-induced pattern formation and the back influence on multiphase turbulence.

Another very interesting feature of multiphase turbulence, that is well suited to exploit exascale computing power, is that it offers a well-defined multiscale approach that extends the above microscale DNS to mesoscale LES and onward to full system simulations. The particle size and mean inter-particle spacing introduce definite length scales, and based on the range of multiphase processes, we define microscale as the scale of particle clusters and patterns, mesoscale as the scale of multiphase flow instabilities and mixing, and macroscale as the scale of the full system. These continuum level scales are supported by atomistic scale processes (see figure 1).

The proposed strategy is to develop: (i) microscale-informed reduced-order descriptions (point-particle coupling models) to be employed at the mesoscale, and (ii) mesoscale-informed reduced-order descriptions (multiphase LES models) to be employed at the macroscales. With this strategy, the predictive capability at the system-scale can be thoroughly validated and uncertainty rigorously quantified.

We propose an approach that progresses from the fully resolved microscale, to the partially resolved mesoscale, and finally to the macroscale. At the microscale, every inter-particle interaction and turbulent flow features will be fully resolved. We will fully exploit the power of exascale to push the number of fully resolved particles to an unprecedented level. At the mesoscale, the individual particles will be comparable to the grid and thus flow scales of the order of the particle diameter will only be partially resolved and an adaptation of the smooth profile method will be used. However, flow features and turbulence of scale larger than the particles will be fully resolved. At the macroscale, particles are much smaller than the grid and will be represented as points and correspondingly for the continuous phase a multiphase LES representation will be pursued.

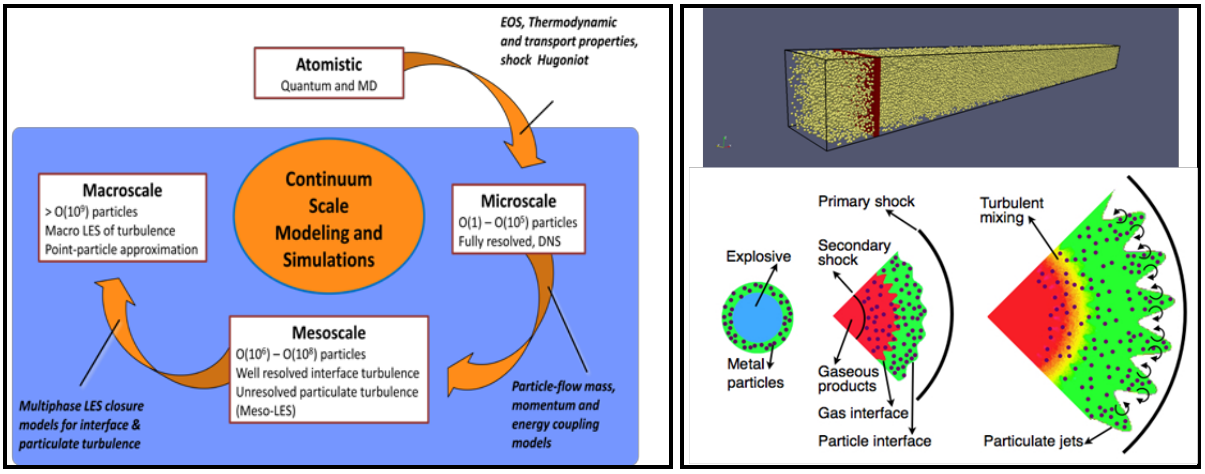

Figure 1: (a) Multiscale strategy for multiphase turbulence, (b) Microscale simulation of shock propagation over a bed of particles, (c) schematic of macroscale simulation of explosive dispersal (taken from UF PSAAP-II CCMT center presentation)

Cost and power constraints are driving exascale architecture design toward fine-grained many-core nodes featuring 10^3 to 10^4 cores. Exascale nodes are expected to have drastically reduced memory per core and to feature deep memory hierarchies with heterogeneous computational units. Intra- and internode data movement will be expensive, both in time and power. In response to these radical changes in exascale architecture, we have significantly refactored, and in some cases, completely re-planned the numerical approach to be adopted for multiphase flows, the coupling schema to be used for multiscale integration, and the approach for integration of UQ.

First, it is widely recognized that data movement will be expensive, and as a result performance can be significantly enhanced by exploiting locally structured discretizations with sizeable, curvilinear, blocks such that data can be accessed lexicographically without indirect addressing. The key to efficiency derives from use of the chain rule for differentiation, which separates a geometry independent part involving matrix operators that remain the same in all subdomains, from a geometry specific part which involves only pointwise multiplication. Second, it is important to recognize that many-core node architectures will result in increased internode latency, and thus places a premium on minimizing the number of initiated messages. In this light, we pursue a high-order discontinuous Galerkin method. This also has significantly reduced memory requirements for shadow nodes – the ratio of computation to number of shadow nodes in the memory has significant impact on the overall performance on GPUs.

We have initiated the development of a library that includes scalable generalized all-to-all communication and very lightweight near- and not-so-near neighbor exchanges. Finally, we pursue heterogeneous space-time discretization that performs computation only when and where it is needed. For increased computational efficiency, we will employ multi-element discretization with a hybrid spectral-WENO scheme to obtain the highest order accuracy in both regions dominated by turbulence and by discontinuities. At the meso and macroscales load balancing with billions of point-particles poses interesting challenges at the exascale.