CCMT

CENTER FOR COMPRESSIBLE MULTIPHASE TURBULENCE

Y5 Annual Report

1/27/2018-1/26/2019

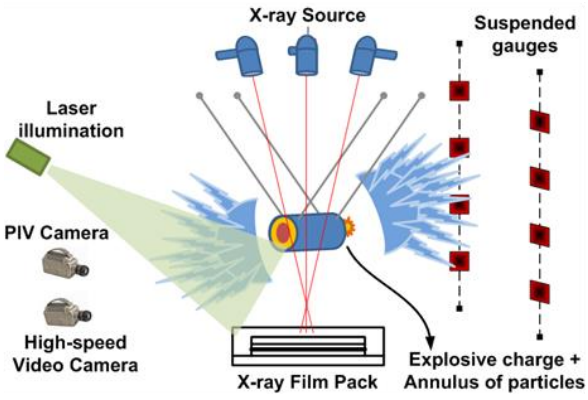

UF UNIVERSITY of FLORIDA

Table of Contents

1. Introduction 4
1.1 Years 1-5 Major Highlights 4
1.2 Background 8
1.3 Demonstration Problem 11
1.4 Simulation Roadmap 12
1.5 Integration 14
2. Macroscale Team 17
2.1 Overview 17
2.2 Demonstration Problem on CMT-nek 17
2.3 Simulations of Eglin microscale experiments 21
2.4 Summary and Future Work 23
3. Microscale Team 25
3.1 Goals and Motivation 25
3.2 Shock interaction with random distribution of particles 25
3.3 Developing a representative wake for force modeling 27
3.4 Shock interaction with a deformable particle 32
4. Experiments 35
4.1 ASU Experiments 35
4.1.1 Goals and Motivation 35
4.1.2 Introduction to ASU Experiments 35
4.1.3 New Equipment over 2017 35
4.1.4 PIV Experiments 35
4.1.5 Particle Image Velocimetry Triggered with Pressure Sensor Data 38
4.1.6 Horizontal Void Cracks 38
4.1.7 Pressure Sensor Data Examined with Differing Bed Heights 41
4.1.8 Main Shocktube Setup 43
4.2 Eglin AFB Experiments 43
4.2.1 Goals and Motivation 43
4.2.2 Microscale Experiments 43
4.2.3 Mesoscale Experiments 45
4.2.4 Macroscale Experiments 45

4.2.5 Summary 46
5. UB Team 47
5.1 Summary of Achievements 47
5.2 Overview 48
5.3 Validation, Uncertainty Quantification and Uncertainty Budget of Mesoscale Sandia Shock Tube Simulation 48
5.4 Validation, Uncertainty Quantification and Uncertainty Budget of Eglin experiments 51
5.5 Convergence study in 1D shock-particle interaction simulation 54
5.6 Reactive burn model parameter calibration in microscale 56
6. CMT-nek Code Development Team 60
6.1 Overview 60
6.2 Non-ideal equations of state and the 1-equation JWL model 60
6.3 Solution limiting and shock capturing 61
6.4 Hero runs 61
6.5 ASU simulations 64
6.6 Future work in 2019 64
7. CS Team 66
7.1 Overview 66
7.2 Load balancing CMT-nek 66
7.2.1 Experimental Results 66
7.3 Load balancing CMT-nek (version 2) 70
7.4 GPU implementation of CMT-nek 72
8. Exascale Team 75
8.1 Overview 75
8.2 BE-SST and architectural exploration using BE simulation 76
8.3 Trace-driven BE simulation of CMT-nek design space 77
8.3.1 Particle-workload distribution tool 78
8.3.2 Improvement to particle-workload distribution tool 79
8.3.3 Application of particle-workload distribution tool 81
8.4 FPGA-acceleration methods for rapid DSE space reduction & UQ 83
9. Deep Dives 86
9.1 Exascale Deep-dive 86
9.2 Multiphase Physics Deep-dive 88

9.3 Nek5000 Users/Developers Meeting 92
10. Publications 94
11. Conferences and Presentations 104
12. Workshops Held or Attended 113
13. Students and Staff Internships 115
13.1 Internships Completed 115
13.2 Internships Planned 116
13.3 Internships Not Yet Planned 116
13.4 Graduated Students 116
13.5 Students Who Will be Graduating 117
13.6 Placement of Staff 117
14. NNSA Laboratory Interactions 118

1. Introduction

1.1 Years 1-5 Major Highlights

The Center for Compressible Multiphase Turbulence (CCMT) has recently completed five years of the PSAAP II program. Below we give a brief summary of the highlights of the program. The highlights are discussed in more detail in the later sections of the annual report.

Research highlights

1. **Full-scale full-physics simulation of the demonstration problem.** Since day one, the center has focused on demonstrating its capabilities to run physically-meaningful Eulerian-Lagrangian simulations of the demonstration problem (explosive dispersal of a random-packed bed of metal particles following the detonation of an energetic material it surrounds) on the largest possible DOE unclassified supercomputers. Two of the major algorithmic co-design accomplishments were the development of a surrogate-based mixture equation of state for detonation products and air, and of a highly scalable (millions of MPI ranks) bin-based Lagrangian approach that enables highly efficient implementation of soft-sphere DEM particle collision model.
2. **Uncertainty budget.** The concept of uncertainty budget (UB) has been introduced and demonstrates as an effective tool for focused uncertainty reduction in both simulations and experiments. Uncertainty reduction exposes experimental weaknesses and reveals model errors, which are referred to as "useful failures" that can be improved upon in subsequent iterations. Our team has demonstrated the feasibility and the value of full physics-UQ of the demonstration problem, through systematic application of UB to four micro/mesoscale campaigns (Sandia Shock Tube, ASU Expansion Fan, Eglin Micro, Eglin Meso) and uncertainty propagation across scales to the macroscale. Continuous iterations between the UQ-team (experimental and numerical viewpoints) and Macro/Meso-scale simulation team have enabled the development of an iterative approach to uncertainty reduction in simulations & experiments.
3. **CMT-nek – a co-designed exascale-ready code.** The center has developed CMT-nek for higher-order-accurate simulations of compressible multiphase flows. CMT-nek is a discontinuous Galerkin spectral element code that is built upon the highly-scalable nek5000, which is a widely-used incompressible flow code. Both the compressible (CMT-nek) and incompressible (nek5000) tracks can now perform one, two and four-way coupled turbulent multiphase flow simulations. These enhanced capabilities are towards increasing the user base from current ~300 world-wide users to O(1000). Through careful co-design efforts between experts in multiphase flow physics, numerical methods and computer science, this code incorporates advanced dynamic load-balancing strategies, and algorithmic options for accuracy vs speed trade-off. With this code, multiphase flow simulations consisting of hundreds of millions of grid points and billions of Lagrangian particles have been run on $O(10^6)$ MPI ranks on the DoE high performance computers.
4. **Microscale simulations and development of PIEP.** The center has carried out fully-resolved three-dimensional simulations at the microscale to better understand fluid-

mediated particle-particle interactions, which has directly lead to the development of the pairwise interaction extended point-particle (PIEP) model. The PIEP framework has revolutionized Euler-Lagrange multiphase flow simulations in the following ways: (i) Allowed four to six orders of magnitude increase in resolution power through accurate sub-grid resolution. For example we can now perform multiphase simulations with an effective resolution of ten-of-trillions of grid points; (ii) PIEP-based Lagrangian models of sub-grid Reynolds stress and residual viscous stress models have been developed for closure of pseudo turbulence; (iii) PIEP-based particle-particle interaction force and torque maps have been used to create a rigorous framework for handling non-spherical particles such as cubes, ellipsoids, rods and plates.

5. **Forensic uncertainty and measurement processing uncertainty.** We proposed the concept of forensic uncertainty quantification using independent investigators to collect unbiased uncertainty information from both experiments and simulations. Also, when the quantity of interest (QoI) is being measured indirectly (not a direct measurement), there is a significant contribution to uncertainty through measurement processing by which the QoI is obtained from the indirect measurement. We showed that the measurement processing uncertainty is often much larger than measurement variability and developed ways to incorporate this in the overall uncertainty budget.

6. **Other UQ advances.** When the number of high-fidelity simulations are small due to expensive computation, we proposed multi-fidelity surrogate models to combine a large number of low-fidelity simulations with a small number of high-fidelity simulations. We showed that a similar level of accuracy could be achieved with much smaller number of high-fidelity simulations. Also, when a specific configuration (design point) is not possible to simulate or test, we proposed an extrapolation scheme using the method of multiple lines based on simulations at points along the line toward the design point. By choosing multiple lines toward the same design point, we obtained multiple predictions, by which extrapolation uncertainty can significantly be reduced.

7. **Complete transition from Rocflu to CMT-nek.** The center started with successful full-scale full-physics simulations of the demonstration problem, as well as the four micro/mesoscale campaigns with our in-house legacy hydrocode, Rocflu. During this period the center developed the fully compressible solver using the discontinuous Galerkin spectral element method, called CMT-nek. With the use of entropy viscosity and positivity-preserving solution limiting CMT-nek is able to capture strong shocks, while offering higher-order accuracy away from the shock for accurate representation of multiphase turbulence. In particular, the discrete element implementation of particle-particle collisions in CMT-nek has allowed simulations starting from very dense close-packed initial condition for the particle bed. Over the past two years the CMT-nek team and the Simulation-physics team have interacted to implement a complete transition to CMT-nek, with the following full range of micro/meso and macroscale simulations:

    a. Microscale simulations of an expansion fan over an arrays of spheres;
    b. Mesoscale simulations of ASU experiments on expansion fan interaction with a particle bed;
    c. Eglin microscale experiments of a few finite-sized particles driven by detonation;

d. Eglin mesoscale experiments of a particle bed driven by detonation;
e. Demonstration simulations of the Eglin blast pad explosive dispersal of a cylindrical bed of particles.

8. **Development of the proxy applications CMT-bone and CMT-bone-BE.** CMT-bone mimics the computational behavior of CMT-nek in terms of operation counts, memory access patterns for data and performance characteristics of hardware devices (memory, cache, floating point unit, etc.) and was validated using VERITAS tool developed at LLNL. While CMT-bone offers reduction in the code complexity by several factors over the application code CMT-nek, an even simpler proxy application with a much smaller footprint was needed for behavioral emulation. This led to the developed and extensive use of the proxy application CMT-bone-BE.

9. **Hybrid computing and energy tradeoff on CPU/GPU nodes.** The implementation was done using PGI CUDA-fortran on GPUs. Load balancing algorithms were developed to derive near optimal workload decomposition on CPU cores and GPU cores. Pareto optimal curves were derived which can help the user choose a system configuration based on his optimization goals. The system configuration includes choosing the optimal frequency for CPUs and GPUs.

10. **Dynamic load balancing for Eulerian-Lagrangian mixed algorithm.** Load balancing of Euler-Lagrange multiphase simulations on millions of MPI ranks for perfect scalability is a grand challenge problem. There are three competing factors: (i) particle-particle collision algorithm requires locality of neighboring particles, (ii) particle-fluid interpolation and projection algorithms require locality of particles and surrounding fluid information, and (iii) particles are not uniformly distributed over the entire fluid volume. In fact, in the demonstration problem all the particles are initially clustered within a small annular region and only over time disperse to occupy a large fluid volume. This demands a load balancing strategy that is dynamic and adaptive. We have developed algorithms that could automatically detect when to perform load balance and remap the particle and the mesh data structures. The remapping algorithms were optimized, resulting in overheads that are comparable to per time step cost of the multiphase flow solver. The dynamic load balancing strategy results in an order of magnitude reduction in the overall computational cost for a variety of problems involving dispersion of a bed of particles.

11. **Genetic algorithm based auto-tuning.** The core computational kernel of spectral element methodology involves dense matrix-matrix multiplication. Furthermore, in a three-dimensional simulation, there are several different variants of matrix-matrix multiplication that are used in elementary numerical operations such as x, y and z spatial derivatives, and interpolation to finer grid for de-aliasing. Thus, optimal implementation of each of these matrix multiplication operation is at the root of enhanced code performance. We have developed a genetic algorithm that rapidly searches and identifies the bed implementation of these matrix multiplication operations. Auto-tuning improved performance and energy requirements for a variety of architectures by upto 25%.

12. **Behavioral Emulation (BE) framework for co-design.** One of the key accomplishments of the center is the development of coarse-grained BE simulation methodology as a co-

design framework. The notable features of BE framework development include using symbolic regression for performance modeling and multi-fidelity surrogate model (with UQ team) for performance prediction. This effort of the center is highlighted by the development of the BE-SST tool, which is a parallel discrete-event BE simulator based on SST. BE-SST and its enhancements will be continually merged into Sandia's SST GitHub repository, making it open-source and accessible for everyone.

13. **Design space exploration using BE-SST.** The BE framework has been used to perform large-scale benchmarking and validation (on up to >100k MPI ranks) on Titan, Vulcan, and Quartz machines. The joint co-design efforts of the BE and CMT-nek teams include predictive simulations approaching a million MPI ranks using BE-SST and end-to-end application design space exploration case study to explore the performance vs accuracy tradeoff of different multiphase flow algorithmic options. One of the key innovations is the particle-workload prediction tool, which uses a novel trace-driven approach to predict dynamic workload of CMT-nek. While the architecture behavioral emulation object (arch-BEO) models the behavior of a specific supercomputer architecture and the application behavioral emulation object (app-BEO) models the behavior of the CMT-nek algorithms, the trace-driven approach is the key ingredient that models the details of how Lagrangian particles are distributed in a specific application.
14. **FPGA acceleration of BE simulation.** The center has developed and validated dataflow, pipelined approaches for FPGAs. We have demonstrated algorithmic and parametric design space pruning and exploration. This effort has achieved 10^6 times speedup over BE-SST.

Students and Staff

1. The number of students supported by core funding and cost share is 38 (Ph.D. 27; M.S. 10; B.S. 1). Of these, 25 students have graduated, and 13 will be graduating in the near future.
2. We have 2 students that have been awarded NSF Graduate Research Fellowships.
3. We had 1 student who won the MAE Best Dissertation Award (2017).
4. We have sent 19 students to the three NNSA laboratories as part of the student internship program. We will be sending 3 more students in the near future, raising the number of student internships to 22.
5. We have placed 8 Ph.D. students and 1 research staff member in postdoctoral or staff positions at the NNSA laboratories or in closely related defense agencies or companies. We fully anticipate placing 7 more students as postdocs at the NNSA laboratories in the coming year.
6. We have placed 3 Ph.D. students at universities as tenure-track faculty.
7. We have also placed 3 postdocs and 11 students in academia and industry.

Deep Dive Workshops

Over the past five years, the center has been very successful in organizing the following workshops and meetings to enhance collaboration and exchange of ideas with NNSA researchers and to disseminate the key accomplishments to the wider scientific community:

1. Deep dive on exascale & CS Issues; Feb. 3-4, 2015; University of Florida.

2. Deep dive on multiphase Physics – I; Oct. 13-14, 2016; Tampa, FL.
3. CMT-nek/nek5000 User & Developer Workshop; April 17-18, 2018; Tampa, FL.
4. Deep dive on multiphase Physics – II; Stanford, Florida co-lead; October22-23, 2018.
5. IUTAM Symposium on Dynamics and Stability of Fluid Interfaces; April 2-5, 2018; Gainesville, FL.
6. Workshop on Multiphase Flows; Nov. 15-16, 2018; Gainesville, FL.
7. Trilab Multiphase Workshop; proposed 3 day workshop; 2019; Drs. S. Balachandar and Duan Zhang.

Publications

1. 156 publications
2. 127 presentations

Video presentation at SC18

Educational Programs

1. Institute for Computational Engineering (ICE).
2. Course in Verification, Validation and Uncertainty Quantification; taught every third semester; Drs. N. Kim and R. Haftka.
3. Yearly specialized course in HPC for computational scientists (as part of the Computational Engineering Certificate); Dr. S .Ranka.
4. Graduate course on multiphase flows; Fall 2016 and 2018; Dr. Balachandar.
5. Exascale challenges and the NGEE work discussed in the reconfigurable computing course (EEL5721/4720) and digital Design course (EEL4712); Drs. H. Lam and G. Stitt.
6. CCMT center used as a motivational example in Introduction to Electrical and Computer Engineering (EEL3000); Drs. H. Lam and G. Stitt.
7. Parallel Computer Architecture (EEL6763); Dr. Ian Troxel.

1.2 Background

The University of Florida (UF) established a Center for Compressible Multiphase Turbulence (CCMT) on January 26, 2014 as part of the NNSA's Predictive Science Academic Alliance Program II (PSAAP-II) Single-Discipline Centers (SDC). The intellectual objectives of the Center are threefold: to radically advance the field of compressible multiphase turbulence (CMT) through rigorous first-principle multiscale modeling, to advance very large-scale predictive simulation science on present and near-future platforms, and to advance a co-design strategy that combines exascale emulation with a novel energy-constrained numerical approach. The Center is performing petascale, and working towards exascale, simulations of instabilities, turbulence and

Goals of CCMT

- *To radically advance the field of CMT*
- *To advance predictive simulation science on current and near-future computing platforms with uncertainty budget as backbone*
- *To advance a co-design strategy that combines exascale emulation, exascale algorithms, exascale CS*
- *To educate students and postdocs in exascale simulation science and place them at NNSA laboratories*

mixing in particulate-laden flows under conditions of extreme pressure and temperature to investigate fundamental problems of interest to national technological leadership. Towards this vision we are tackling the following challenges:

1) Target an important application that can only be enabled by exascale computing: We are solving a complex multiscale problem at an unprecedented level of physical detail and integration and thereby advance predictive simulation science. CMT poses a grand challenge to our understanding as it combines three complex physics: compressibility, multiphase flow and turbulence. CMT occurs often under extreme conditions of pressure and temperature, and as a result is not easily amenable to high-fidelity experiments and diagnostics. CMT presents a fascinating array of poorly-understood instability, transition, and turbulent processes manifest over a wide range of strongly interacting length and time scales. Current computational approaches involve models and closures that are developed from incomplete understanding, and as a result are largely empirical. Fully validated exascale simulation perhaps is the only path to fundamental breakthroughs that can lead us out of current empiricism.

2) Well-defined problem hierarchy leading to a demonstration problem: A multiscale approach from the microscale to the mesoscale and to the macroscale is being pursued for a systematic integrated investigation of the CMT physics. We have adopted a problem hierarchy that culminates at a signature demonstration problem of explosive dispersal of particles from a well-characterized initial condition, which fully exercises all the key complex processes of CMT. We pursue a coupling strategy where (i) fully resolved microscale simulations will lead to reduced order descriptions (interphase coupling models) to be employed at the mesoscale and (ii) partially resolved mesoscale simulations will lead to reduced order descriptions (multiphase large eddy simulation closures) to be employed at the macroscale. This will allow computational efficiency and high degree of parallelism at all levels of the hierarchy.

3) Simulation and experiment roadmaps for rigorous validation: We focus on integrated system-scale simulations of the demonstration problem from the outset using existing integrated code capabilities. Simultaneously, we also perform petascale simulations at the micro and mesoscales. Improvements to micro-to-meso and meso-to-macro coupling models will be systematically and periodically incorporated at the appropriate higher level. A layered systems engineering approach is used to organize and integrate physical subsystems with numerical, software and service components, to achieve progressively improved operational capability for system-scale simulations. We have developed a detailed simulation and experiment roadmap which allow rigorous step-by-step validation at each step of the problem hierarchy.

4) Develop novel uncertainty quantification (UQ) approaches for CMT: Detailed measurements from carefully chosen existing and planned experiments at the Air Force Research Laboratory Munitions Directorate (AFRL-RW), Sandia Multiphase Shock Tube facility and Los Alamos Center of Mixing under Extreme Conditions (CoMuEX) are used for rigorous quantification of uncertainties from the micro/mesoscales to the macroscale. We are engaged in vigorous uncertainty reduction through better characterization and instrumentation, rigorous calibration of the models, and improved numerical resolution. Simultaneous simulations and experiments at the micro, meso and macroscales of the problem hierarchy will allow us to both propagate up

uncertainty to higher scales, and to reduce uncertainty through iterative improvements at the lower scales. A particularly difficult aspect of CMT is that it is characterized by extreme events that are localized in space and time. A key innovation is the development of novel techniques for accurate characterization of probability tails in the uncertainty quantification of such rare but critical events.

5) Demonstrate integrated performance on current/near-future architectures: Modern many-core architectures (such as Intel MIC), that provide high raw gigaflops, have deep memory hierarchies and low overhead threading capabilities. We exploit these capabilities to optimally utilize both computational and energy resources. In particular, we will tackle load balance and performance challenges in terms of data and work decomposition for the CMT code framework. Different parallelization schemes will be considered for effectively implementing simulations at the microscale, mesoscale, and system-scale, especially for heterogeneous resources.

6) Develop methods for predicting performance on a variety of exascale architectures: While many exascale trends seem clear, there are far too many permutations in the design space to select one a priori. We leverage the unique Novo-G facility at the NSF-supported UF Center for High-Performance Reconfigurable Computing (CHREC) to emulate and evaluate a series of candidate exascale architectures. We are developing an unprecedented capability to behaviorally prototype in software and hardware a variety of promising (as defined by leading exascale initiatives) forms of next-generation exascale (i) device and node designs at the micro-level and (ii) communication and system architectures at the macro-level. We are conducting experiments with CMT-bone kernels, miniapps and skeleton-apps to evaluate promising architectures in terms of performance, energy, temperature, reliability, and scalability. Modeling, simulation, and estimation tools (e.g., those supported within the Sandia's Structural Simulation Toolkit (SST)) are being leveraged with our behavioral simulations and emulations.

7) Solutions for energy efficiency and thermal management: We are developing a framework for multi-element and multi-objective optimization that will simultaneously minimize energy and maximize performance. We exploit the data and task parallelisms within CMT application and its UQ implementation to develop innovative low complexity static and dynamic algorithms for scheduling, while considering important factors such as thermal constraints and leakage currents.

1.3 Demonstration Problem

We aim at solving a problem of Compressible Multiphase Turbulence (CMT) at an unprecedented level of physical detail and thereby advance predictive simulation science. The overarching demonstration problem consists of a cylindrical core of simple explosive pellet of about 10 grams will be surrounded by a cylindrical very-thin-walled glass jacket of larger diameter. The annular region between the pellet and the jacket will be filled with mono or polydisperse metal powder of spherical shape. The shape and amount of the explosive charge and the size distribution of the metal powder and its material (aluminum, steel, tungsten, etc.) are parameters that will be varied. The charge will be hung from a test fixture so that the effect of the ground and the surrounding structures will be eliminated during the initial phase of explosion and dispersion. The orientation of the test setup will be such that the resulting explosive dispersal of particles and the gas field can be highly accurately measured. The following features makes this problem a very good choice for demonstration: (i) the explosive dispersal exercises all the major CMT physics, (ii) the extreme conditions makes this a demanding test for predictive capability, (iii) this problem requires exascale for true predictive capability, and (iv) we have already performed similar experiments and validation-quality measurements. The explosive dispersal of solid particles problem displayed in Figure 1.1 and described by Frost *et al.* (Phys. Fluids, 24(9), 2012) was chosen for the initial phase of our research activities.

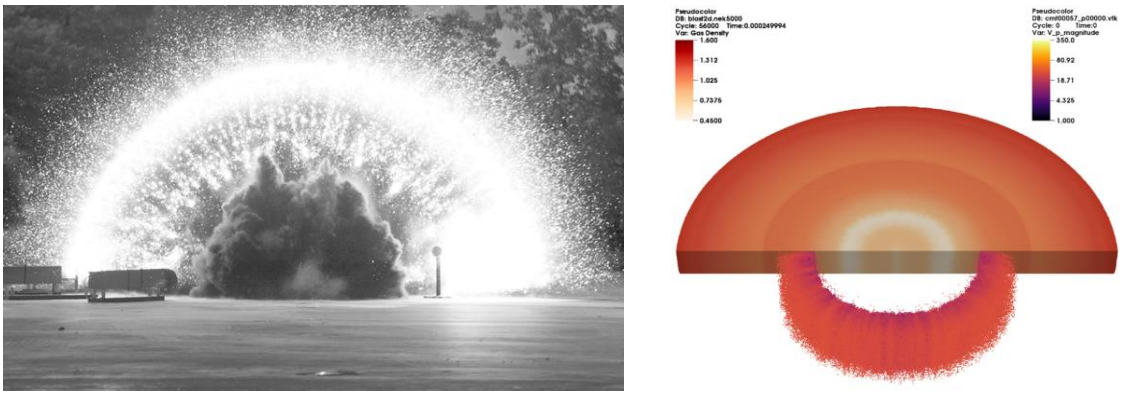

Figure 1.1. Left: Schematic of the three-dimensional computational domain for the explosive dispersal of particles simulations. Not drawn to scale. Right: Current three-dimensional hero run: 32768 ranks. Time shown, t=1.46ms.

1.4 Simulation Roadmap

The center is focused on integrated system-scale simulations of the demonstration problem from the outset using existing integrated-code capabilities. Figure 1.2 shows the roadmap of the proposed sequence of simulations. The following important considerations was used in constructing the roadmap: (i) Along with system-level simulations of the demonstration problem, we will perform increasingly more complex simulations at the micro and mesoscales. Based on these simulations, improvements will be made to micro-to-meso and meso-to-macro coupling models. (ii) To take maximum advantage of validation experiments, large numbers of simulations will be required for optimal calibration. We are using surrogate models to allow us to solve the multi-level optimization problem associated with selecting the physical constants that give the best match with the numerical model. (iii) Variations of the key control parameters (particle size,

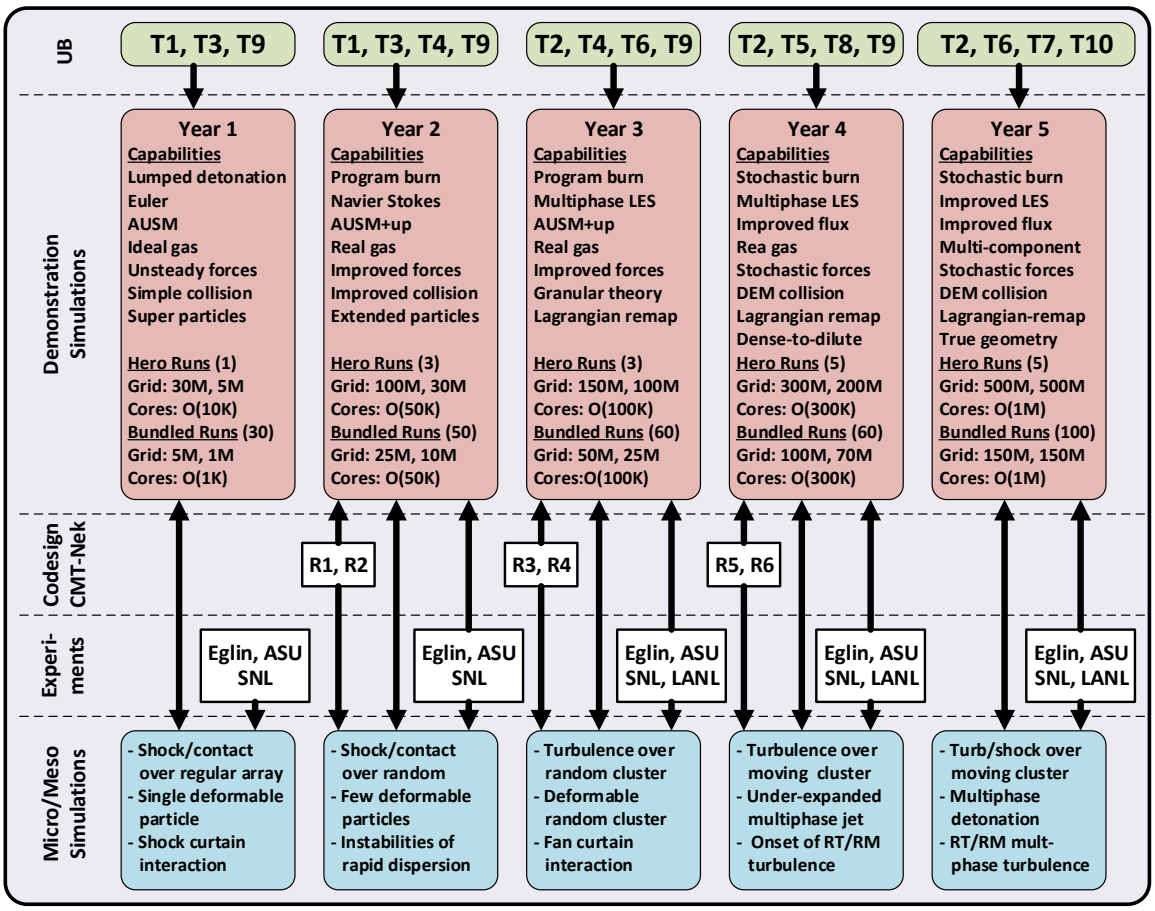

Figure 1.2. A roadmap showing the sequence of progressively more complex simulations to be performed over the five years of the center. Micro and mesoscale simulations will be performed along with macroscale simulations of the demonstration problem. Both *hero runs* and *bundled runs* will be performed. Bundled runs will be combinations of varying parameters and ensemble of realizations. This plan is based on future petascale platforms that are expected to be available through NNSA. Concurrent simulations will be accomplished using DAKOTA.

particle material, shock strength, etc.) will be guided by simulations that identify which combinations of parameters will elicit different modes of instability. (iv) Statistical variability will be explored through an ensemble of realizations under nominally identical conditions. (v) Simulations are currently being carried out concurrently as *bundled runs* using the DAKOTA toolkit. (vi) We anticipate increasingly larger petascale computational platforms to be available at the NNSA labs. (vii) We have and will continue to perform selective *hero runs* at super-high resolution to help quantify discretization errors to help assess the accuracy of the estimated uncertainties. (viii) UQ is being used to guide the selections of quantities to be measured with preference to those with low uncertainty, so as to avoid empty validation based on large error bars.

The Year-1 simulations of the demonstration problem employ simplified physics model: (i) a lumped detonation model, (ii) the single-phase AUSM+ flux scheme for the Euler gas equations with idea gas equations of state, (iii) the actual particles are approximated with computational super particles, (iv) gas-particle coupling is through point-particle models of quasi-steady and unsteady forces and heat transfer, and (v) particle-particle collisions are accounted using a simplified collision model. The corresponding hero and bundled runs represent our Year-1 starting point. The above roadmap shown in Figure 1.2 lays out year-by-year progression of more detailed simulations that incorporate additional physics through new and improved models. Furthermore, each year we plan to perform larger and larger hero runs as well as large array of bundles macroscale simulations for uncertainty quantification.

The simulation roadmap is driven from the top by Uncertainty Budget (UB). A detailed phenomenon identification and ranking analysis of the demonstration problem has identified 11 key sources of errors and uncertainties which are briefly listed below:

- T1: detonation process modeling
- T2: Multiphase turbulence modeling
- T3: Real gas thermodynamic and transport properties
- T4: Inter-particle collision modeling
- T5: Particle compaction modeling (during early stages of detonation/shock propagation)
- T6: Point particle modeling of gas-particle momentum (force) exchange
- T7: Point particle modeling of gas-particle thermal (heat-transfer) exchange
- T8: Particle deformation, sintering and break-up physics
- T9: Discretization (numerical) errors
- T10: Errors from geometric approximation (geometric differences in the details of experiments and simulations)
- T11: Experimental uncertainties and measurement errors

The key activity of UB effort will be to quantity the uncertainty in the zeroth and first order prediction metrics. The zeroth order prediction metrics of the demonstration problem are:

- The blast wave location as a function of time
- The average particle front and tail locations as a function of time
- The number of large-scale instabilities of the particulate front

The first order prediction metrics go beyond the zeroth order metrics and the details of the demonstration will be quantified with the following first order metrics:

- Time evolution of the pressure at selected points within 5% error
- Time evolution of the thermal load at selected points within 20% error
- Time evolution of average particle concentration within 15% error
- Evolution of particle front thickness due to instability and turbulent mixing within 10% error
- RMS turbulent velocity and pressure fluctuations at the particle front within 15% error,
- Time evolution of local particle size distribution within 15% error
- Multiphase turbulent spectra and correlation length scales within 20% error.

An important component of the yearly UB effort is to quantify contribution from the above 11 sources of errors and uncertainties to each of the prediction metrics. This quantification will allow us to focus on error/uncertainty reduction. Thus each year we will focus on uncertainty reduction and quantification through certain modeling and simulation activities. These are the UB drivers for the proposed roadmap and they are presented at the top row of Figure 1.2.

Figure 1.2 also presents the yearly releases of CMT-nek, the new code being co-designed through an integration of exascale higher-order algorithm with exascale emulation/ simulation. Also indicated are yearly coordination with the micro, meso and macroscale experiments to be performed at Eglin Air Force Base, Arizona State University (ASU), Sandia National Laboratory (SNL) multiphase shock tube facility and Los Alamos National Laboratory (LANL) Center of Mixing Under Extreme Conditions. The macroscale simulation road map will also be supported by the yearly progression of mico and mesoscale simulations, which is also indicated in Figure 1.2.

1.5 Integration

The Center recognizes the critical importance of tight integration for the success of the center. The center will be organized in terms of tasks and cross-cutting teams, rather than in terms of faculty and their research groups. The physics-based tasks are continuous and particulates phase modeling and simulation. In addition we have exascale (EX), computer sciences (CS) and uncertainty quantification (UQ) as the cross-cutting tasks that will interface and integrate the physics-based tasks. By ensuring faculty, research scientists, and postdocs contribute to multiple physics and/or cross-cutting tasks, we will achieve tight integration. This matrix organization, depicted in Figures 1.3 and 1.4, tears down discipline and departmental boundaries and allows close interaction. In addition, significant effort has gone into integrating the various disciplines.

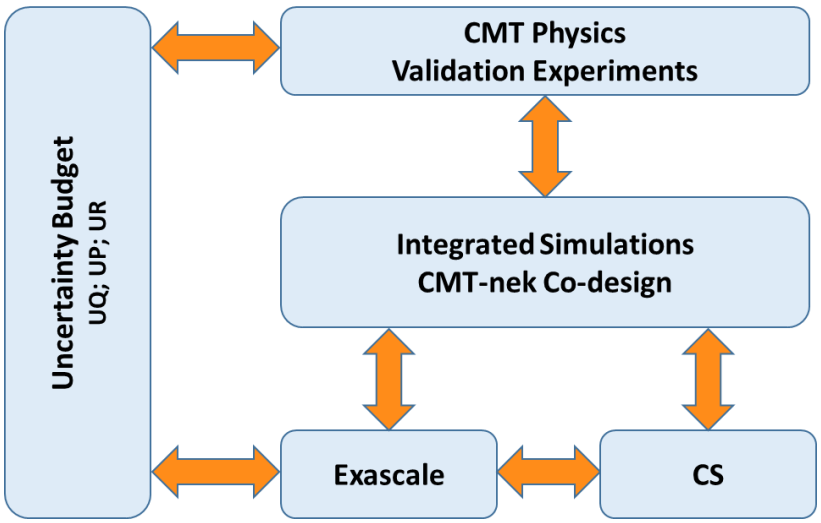

Figure 1.3. Integration – how the different pieces fit together. Here, UB drives the CMT physics, validation experiments, and exascale behavioral emulation/simulation. CS and exascale feeds into the CMT-nek codesign strategy.

Hour time slots	Exascale	CMT-nek	CS	Micro	Macro	UQ	Exp
Exascale	X	X	X			X	
CMT-nek	X	X	X	X	X		
CS	X	X	X				
Micro		X		X	X	X	
Macro		X		X	X	X	X
UQ	X			X	X	X	X

Figure 1.4. Management – tasks and teams. Teams include students, staff, and faculty. The Center is organized by physics-based tasks and cross-cutting teams, rather than by faculty and their research groups. All staff and large number of graduate students located on 2nd floor of PERC. All meetings held in PERC. Weekly interactions (black); Regular interactions (red).

The intellectual integration of the different simulation and experimental talks, across the three different scales (micro, meso and macro) is shown in Figure 1.5. Uncertainty quantification, propagation and reduction along the ten sources of errors/uncertainties (T1 to T10) forms the framework that connects and drives the different simulation and experimental activities of the center. The hierarchical flow of error/uncertainty information to the macroscale is shown.

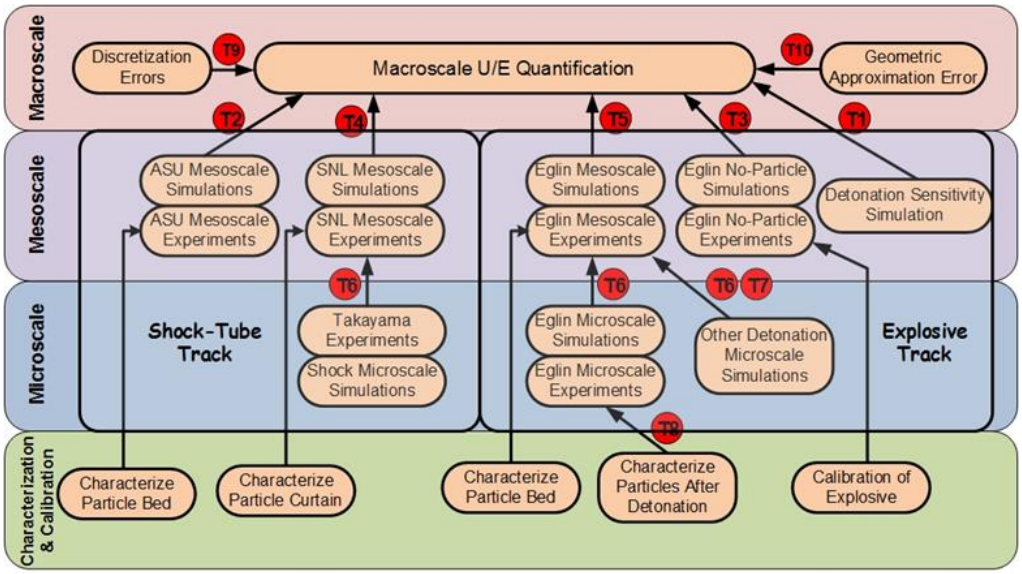

Figure 1.5. Uncertainty Quantification and Propagation Plan.

At the *microscale* the motion and thermal evolution of particles depends on the flow around them. In return, the particles modify the local flow by the formation of momentum and thermal wakes. Particle structures (chains and clusters) spontaneously form due to wake-wake, particle-wake and particle-particle interactions. At the *mesoscale*, due to inertial interaction with turbulence, particles preferentially accumulate. Also, flow instabilities can lead to large-scale structures in particle distribution. These nonuniformities have profound influence on their collective back influence on the flow. At the *macroscale* (or *system-scale*) the geometric details of the setup influence the coupling between the particles and expanding gas. Important aspects of the multiscale coupling strategy we are pursuing includes: (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 macroscale. With this strategy, the predictive capability at the system-scale can be thoroughly validated and uncertainty rigorously quantified as illustrated in Figure 1.5.

Note that the multiscale coupling strategy and the overall uncertainty quantification plan includes both a shock-tube track and an explosive track. We have been working with the Experimental Teams at the various locations and have discussed in detail the type of characterization, inputs, and output from the experiments for a meaningful UB approach.

2. Macroscale Team

2.1 Overview

The premier goal of the Macro/Mesoscale Team (MMST) is to run a problem of explosive dispersal of particles at an unprecedented level of details. Figure 2.1 provides an overview of the major phases of such a problem. Its second objective is to validate new state-of-the-art models developed at CCMT against experiments run by our collaborators. During year five, our efforts have focused on two targets: i) run production-scale simulations of the demonstration problem using CMT-nek, and ii) continue improving the simulations of Eglin Air Force Base experiments to better our validation.

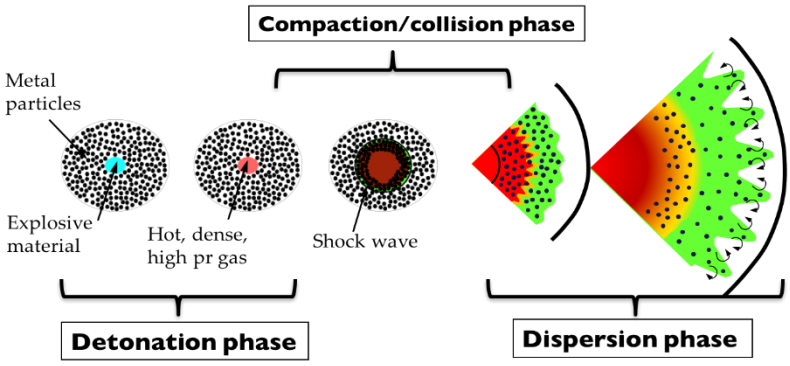

Figure 2.1. Sequence of events during our demonstration problem, an explosive dispersal of particles.

2.2 Demonstration Problem on CMT-nek

This year marked the full transition of the simulation of the demonstration problem from the original in-house code, Rocflu, to the center's developed code, CMT-nek. Therefore, the strategy for the three hero runs performed over the past year on CMT-nek has been to gradually ramp up the conditions and physics models. Eventually, the team will replicate true explosive charge conditions with full particle loading to fully capture the physical details and intricacies of the Eglin Blastpad experiment.

The first of these runs, labelled as CMT Hero Run 1, was a recreation of the first demonstration problem run performed on Rocflu but on a smaller domain. The radius of the charge was still taken to be 3.8 mm as done for the Frost simulations and the entire domain is a 30cm x 30cm x 1.5 mm box using 32.4 million degrees of freedom with periodic boundary conditions on all box edges. This run was performed in order to tune the artificial viscosity parameters in the code to handle the strong shocks that arise from the blast. To allow for this, the density and pressure inside of the charge was initialized at constant values of 23.5 $\frac{kg}{m^3}$ and 101.325 MPa respectively. The remaining gas was initialized at standard atmospheric conditions, leading to an initial pressure ratio of 100 for the gas. To reduce the number of additions to the code needed for this first run, both the charge and the ambient were governed by the ideal gas equations. One million computational particles of diameter 100 μm and density 33.33 $\frac{kg}{m^3}$ were placed outside of the charge up to an outer radius of 1cm at a volume fraction of 5%. The particle density was modified in order to better match the particle to charge mass ratio from the Frost experiments. This simulation was run out to 300 μs, at which point the blast wave reached the outer boundary of the domain, using LLNL's Vulcan

machine and 32768 ranks. Figure 2.1 shows the gas density and particle velocity contours at 250 μs showing the blast wave nearing the boundary as well as the expansion of the particle bed with what appears to be some small fingers at the outer edge.

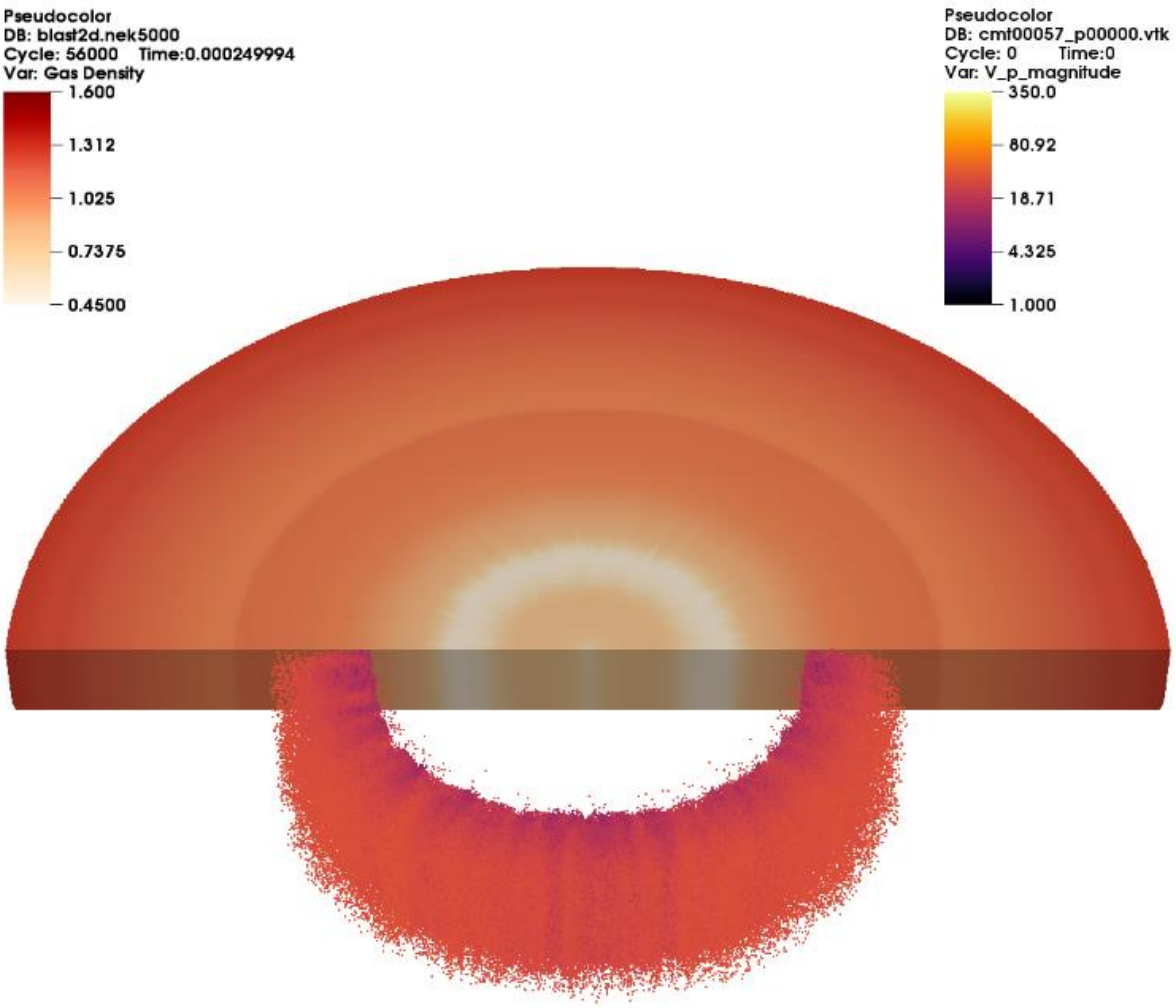

Figure 2.2.. Gas density and particle velocity magnitude for CMT Hero Run 1 at 250 μs.

The second run, CMT Hero Run 2, was designed to add CMT-nek's soft-sphere collision model to the physics of the demonstration problem. The same geometrical setup was used for the domain as CMT Hero Run 1 as well as the same initial conditions for the gas phase. The particle properties also remain the same, however now they were initialized at a uniform volume fraction of 60% and at an outer radius of 2 cm. The number of computational particles simulated was also increased from one to two million. With this initial setup, the compaction phase occurs almost instantaneously and the capability of the code to use the collision model and survive the

compaction phase was able to be tested. The case was run to 180 μs, again on Vulcan with 32768 ranks.

The third run, CMT Hero Run 3, was the first run to use the JWL equation of state mixture model which was imported from Rocflu into CMT-nek. The dimensions of the charge and particle bed remained unchanged from CMT Hero Run 2, but now the charge was taken to be Composition B to begin the transition to simulate the Eglin Blastpad cases. Then, the charge was initialized using a uniform density of 1712 kg/m^3 and pressure of 11.9 GPa, while the ambient was left at atmospheric conditions. The particles were still initialized at a volume fraction of 60%, but their density was now set to 8000 kg/m^3, matching the density of steel. The overall domain was then reduced to a box of 20 cm x 20cm x 1.5 mm, however the total number of degrees of freedom was left unchanged. This change in the grid parameters was due to the increased resolution required for CMT-nek to correctly handle the more extreme flow conditions imposed by the more realistic initial conditions. This case was again run on the Vulcan supercomputer but was the first demonstration problem simulation to cross the 100,000-rank threshold as it was run using 131,072 ranks. The final time of this simulation was 34 μs as this is when the blast wave from this case reached the outer boundary of the domain. Figures 2.2 and 2.3 show the particle bed at this final time while highlighting two key quantities. Figure 2.2 shows the particles colored by their local particle volume fraction. The inner edge of the particle bed is still clearly inside the compaction regime, at around 74% volume fraction, but the outer region of the bed shows some layering or flaking patterns as the outer region expands outwards and leaves the compacted region. This is reminiscent of the particle layering seen in both the experiments and simulations of the ASU expansion fan problem. Figure 2.3 shows the particles colored by the magnitude of the collisional forces acting on them. This shows that the collisional forces are highest in the compaction area, as expected, and also shows the layering effects at the outer region of the bed.

To conclude, Figures 2.4 shows the trajectory of the leading blast wave and the outer front of the particle bed for the three hero runs. The plots show the logical evolution of the trajectories for the increased physics that were introduced in each case. In terms of the blast wave trajectory, increasing the particle loading from CMT Hero Run 1 to CMT Hero Run 2 reduced the speed of the shock in some extent, but introducing a much stronger pressure ratio from realistic charge initial conditions and introducing the JWL equation of state led to, overall, the blast wave travelling much faster. The particle front plot demonstrates these same trends, with the outer front of the bed of particles replacing the blast wave.

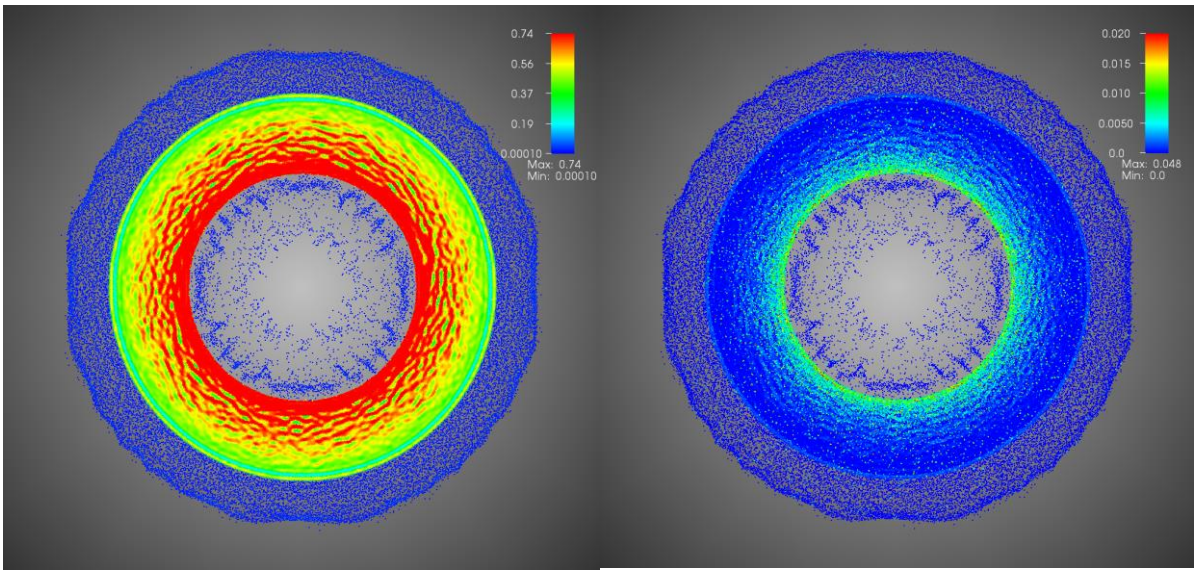

Figure 2.3. Individual particle locations colored by local particle volume fraction (2) and by imparted particle collisional force (3).

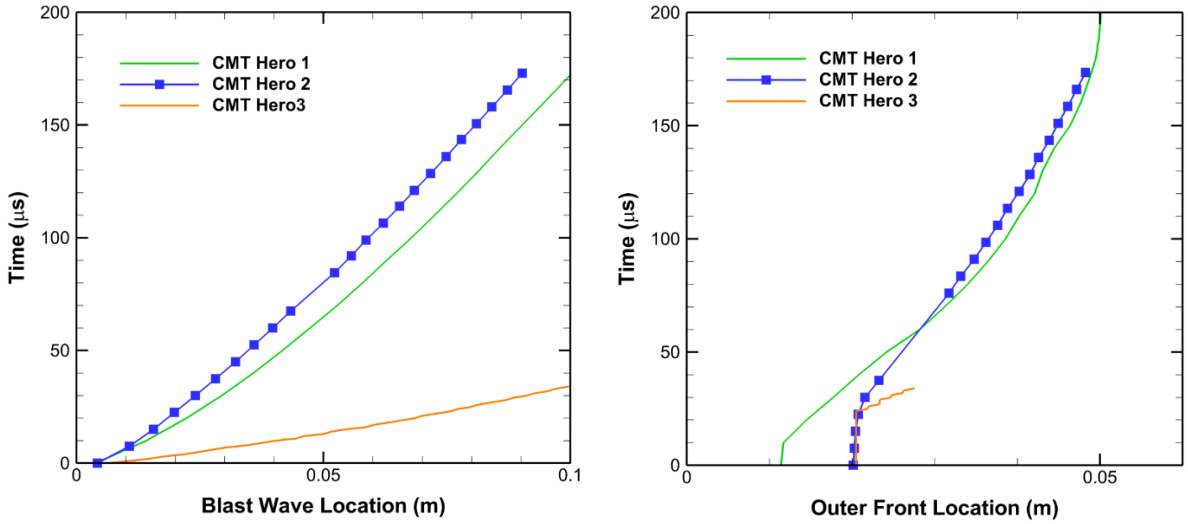

Figure 2.4. Plots of blast wave (left) and outer particle front location (right) as functions of simulation time.

2.3 Simulations of Eglin microscale experiments

The forensic UQ process has led our team to review the position of the explosive in the simulations of Eglin microscale experiments. It has been updated to match a second configuration of the experiment, in which the end of the explosive is located flush to the particle within the barrel. A grid resolution study for this problem has been performed by comparing the trajectories of the primary blast wave along the centerline in the simulations, as the cell spacing is gradually reduced and the initial condition is left unchanged. To ensure that maximum flow gradients and speeds would occur then, the initial conditions used for the grid resolution study contained more energy than those employed for simulations of the Eglin experiments. Illustrated in Figure 2.6, solution convergence for this problem was observed at 50 micrometers. The explosive nature along with the moderate geometric complexity of this problem make it suitable to serve as a test case for simulations employing detonation conditions in CMT-nek.

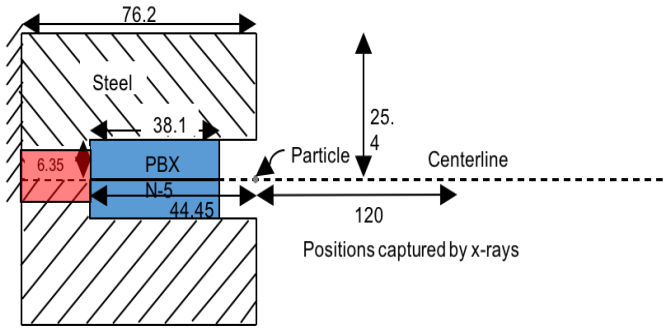

Figure 2.5. Eglin Air Force base microscale experiment configuration.

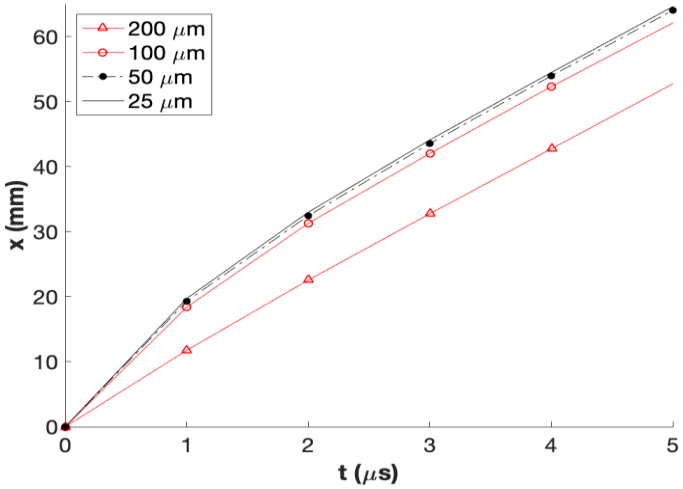

Figure 2.6. Early time shock trajectories of simulations performed for the grid resolution study. Only minor differences are observed between the simulations employing a cell spacing of 25 or 50 micrometers while other results remain in less agreement. A grid converged solution is observed at 50 micrometers.

Initial conditions for simulations of the Eglin microscale and mesoscale explosive experiments have been advanced to include a many of the uncertainties associated with the explosive and its detonation properties. The unreacted density, reaction zone thickness and heat of reaction are uncertainties of the explosive now considered in the simulations. The post-detonation profile of density for PBNX-5 used as initial condition in the simulations is showed in Figure 2.7. To quantify the effects of varying the explosive parameters on the post-detonation flow, the explosive products contact front and transverse shock position (TSP) are compared between simulations and experiments. The transverse shock is the portion of the primary shock that travels radially outward along the face of the barrel following the explosion.

Showed in Figure 2.8, simulation results of the TSP reveal high sensitivity to the explosive heat of reaction, a parameter that governs the detonation pressure of the initial condition. Following a study that included 45 unique explosive initial conditions, a most likely set of explosive parameters, used in the case of Figure 2.9 below, was obtained through uncertainty analysis based on shock wave position data from simulations and experiments.

The numerical modeling uncertainty of the explosive-specific parameters in the JWL equation of state has been preliminarily examined using a single explosive initial condition and Latin hypercube sampling of the JWL parameters. Forty samples of JWL parameters were obtained and used to simulate the post-detonation flow of explosive products. Results from these runs indicate a need for further investigation, which is currently underway.

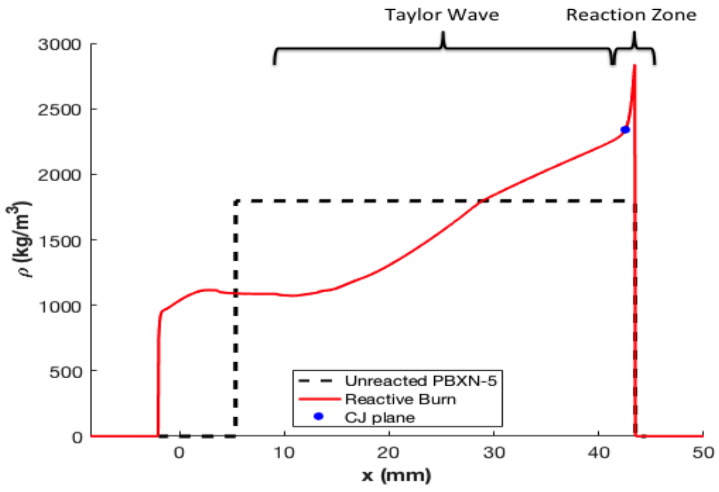

Figure 2.7. The post-detonation density profile of PBXN-5 is used as the explosive initial condition for simulations of the Eglin microscale experiments.

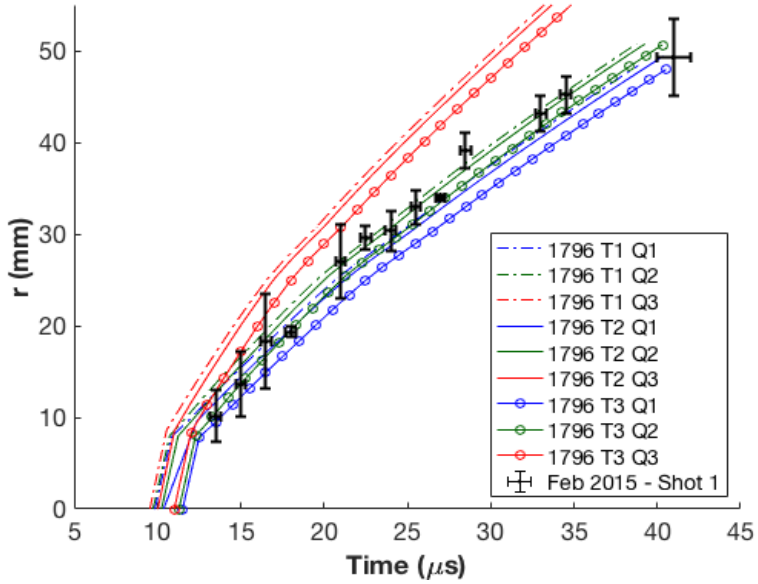

Figure 2.8. Experimental transverse shock position compared with simulation results. The initial conditions employed the mean experimental density for the explosive and vary the reaction zone thickness and heat of reaction, denoted by T[1-3] and Q[1-3], respectively, in the legend. Results indicate that the transverse shock is sensitive to the heat of reaction parameter and less sensitive to changes in reaction zone thickness.

2.4 Summary and Future Work

Throughout year five, the MMST has been rapidly implementing, testing and validating the simulations of the demonstration problem on its new production code, CMT-nek. Three massively parallel CMT Hero runs have been executed on Vulcan, demonstrating the new capability of the Center to produce physically realistic large-scale simulations of explosive dispersal of particles. The last missing piece in enabling direct comparisons of a CMT Hero run with the Eglin Blastpad experiment is the Reactive burn initial condition, which the MMST is already working on.

With regards to the simulations of Eglin Experiments, the integration of the MMST with the UQ team has proven its value once again. Uncertainties associated with the explosive and its detonation properties have been recognized as sources of significant variation in the transverse shock position. Our team has engaged in a meticulous study of these parameters as well as other explosive-specific parameters in the JWL equation of state to provide robust gas simulations that will not interfere in the analysis of the particles results.

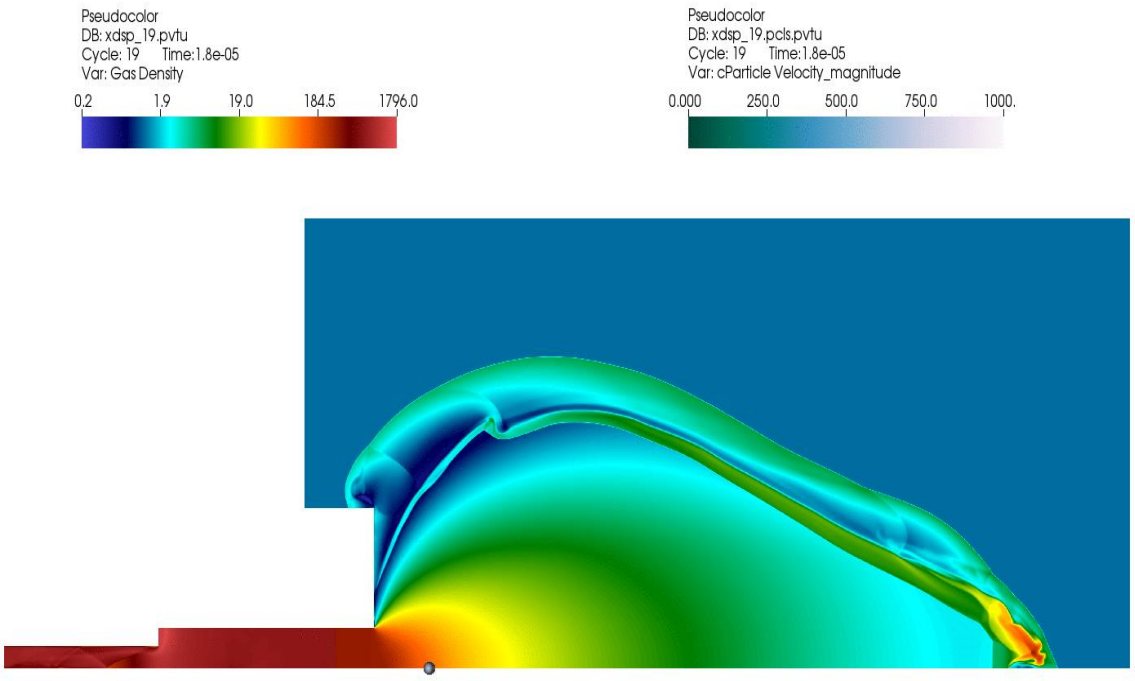

Figure 2.9. Simulation density contour at t_{sim} = 18 μs with multiple particle positions superimposed. This simulation employs the most likely parameters for the explosive initial condition obtained from the uncertainty analysis.

3. Microscale Team

3.1 Goals and Motivation

Goals of Microscale team include performing fully resolved simulations of shock and contact interacting with bed of particles. Purpose of these simulations is to understand the underlying complex physical mechanisms occurring during blast conditions and converting this understanding into models, which can be used to predict the particle motion and back effect of particles on the flow.

3.2 Shock interaction with random distribution of particles

One of the objectives of the microscale team is to perform fully resolved simulations of shock and contact interacting with bed of particles. The purpose of these simulations is to understand the underlying complex physical mechanisms occurring during blast conditions and converting this understanding into models, which can be used to predict the particle motion and back effect of particles on the flow. With these goals in mind, the microscale team has performed fully resolved inviscid simulations of incident shock interacting with randomly distributed particles. For these simulations we have neglected the effect of particle motion, since we are interested in the early time behavior of the force histories experienced by the particles.

Previously, we have studied in the depth the effect of the flow on the particles trough the force histories experienced by the particles. More recently we have been studying the effect of the particles on the flow. Random distribution of particles leads to fluctuation in velocity, pressure, and density fields. These fluctuations in turn influence the forces experienced by the particles.

To study the flow unsteadiness during shock particle interaction, we derive the volume averaged governing equations. They are given by

$$\frac{\partial(\varphi_g\langle\rho\rangle)}{\partial t}+\frac{\partial(\varphi_g\langle\rho\rangle\tilde{u}_i)}{\partial x_i}=0$$

$$\frac{\partial(\varphi_g\langle\rho\rangle\tilde{u}_i)}{\partial t}+\frac{\partial(\varphi_g\langle\rho\rangle\tilde{u}_i\tilde{u}_j)}{\partial x_j}=-\frac{\partial\left(\varphi_g\langle p\rangle\right)}{\partial x_i}+\frac{1}{V}\int_{S_d}pn_i dS-\frac{\partial(\varphi_g[\langle\rho u_i u_j\rangle-\langle\rho\rangle\tilde{u}_i\tilde{u}_j])}{\partial x_j}$$

$$\frac{\partial(\varphi_g\langle\rho\rangle\tilde{E})}{\partial t}+\frac{\partial(\varphi_g\langle\rho\rangle\tilde{u}_i\tilde{E})}{\partial x_i}$$
$$=-\frac{\partial\left(\varphi_g\langle p\rangle\tilde{u}_i\right)}{\partial x_i}+\frac{1}{V}\int_{S_d}pu_i n_i dS-\frac{\partial(\varphi_g[\langle pu_i\rangle-\langle\rho\rangle\tilde{u}_i])}{\partial x_i}-\frac{\partial(\varphi_g[\langle pu_i E\rangle-\langle\rho\rangle\tilde{u}_i\tilde{E}])}{\partial x_i}$$

The process of averaging results in Reynolds stress like term in the momentum equation. We refer to it as the pseudo turbulent Reynolds stress. Similar fluctuating terms appear in the energy equation as well. We also compute the RMS of velocity field and compare it against the mean

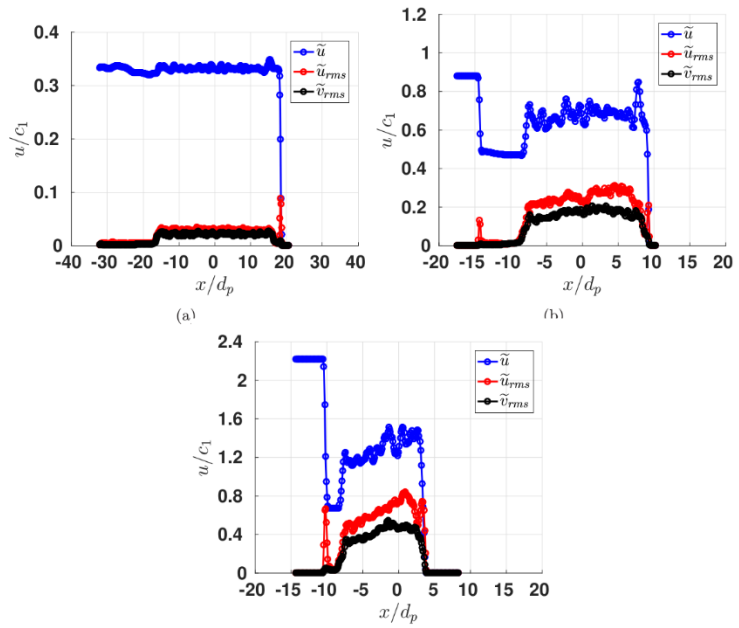

Figure 3.1.Plot of the non-dimensional Favre averaged streamwise velocity, $\tilde{u}$ (blue curve), RMS streamwise velocity, $\tilde{u}_{rms}$ (red curve), and RMS transverse velocity, $\tilde{v}_{rms}$ (black curve) for (a) $\varphi_1 = 2.5\%$ and $M_s = 1.22$, (b) $\varphi_1 = 20\%$ and $M_s = 1.66$, and (c) $\varphi_1 = 25\%$ and $M_s = 3$.

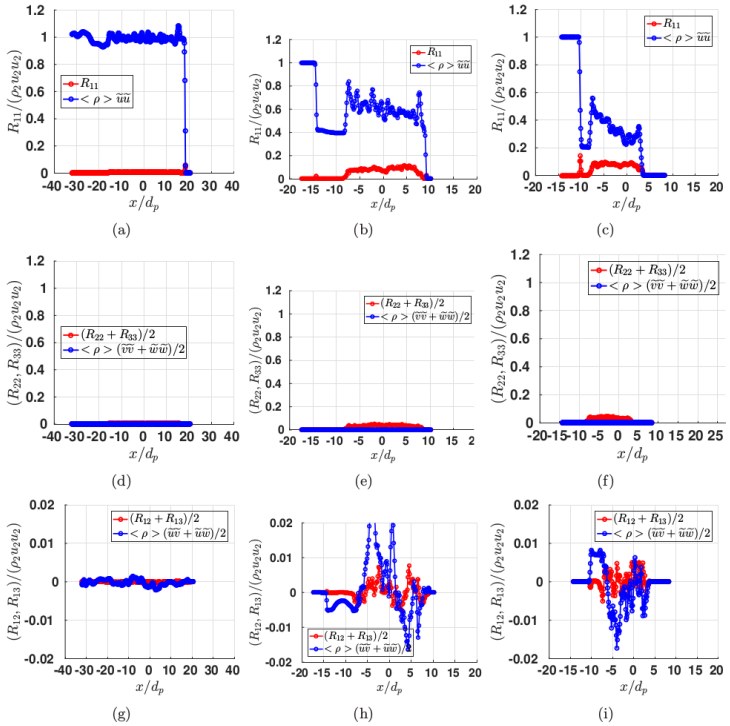

Figure 3.2. Plot of the non-dimensional pseudo turbulent Reynolds stress terms (a-c) R_{11} (d-f) $(R_{22} + R_{33})/2$ (g-i) $(R_{12} + R_{13})/2$; (a,d,g) $\varphi_1 = 2.5\%$ and $M_s = 1.22$, (b,e,h) $\varphi_1 = 20\%$ and $M_s = 1.66$, and (c,f,i) $\varphi_1 = 25\%$ and $M_s = 3$.

Favre averaged velocity. We can observe from the Fig. 3.1 that the RMS of streamwise velocity

and RMS of transverse velocity are of nearly the same magnitude, but they are less than 50% of the average streamwise velocity. We observe that the magnitude of the RMS velocity field increases as the particle volume fraction and incident shock Mach number increases.

We compare the magnitude of the pseudo turbulent Reynolds stress against the magnitude of the convective term in Fig. 3.2. Since, the magnitude of the RMS velocity field is quite small the magnitude of the pseudo turbulent Reynolds stress is also small for most of the cases considered.

3.3 Developing a representative wake for force modeling

One of the goals of the microscale team is to develop particle-scale informed models for use in meso- and macroscale simulations. The need for such models arises from the filtering/volume averaging necessary to perform such large simulations. Instead of acquiring the actual fluid properties at every location, meso- and macroscale simulations calculate filtered quantities. The filtered quantities are generally used as inputs for closure models to account for the average effect of neglected particle-scale details. The present work focuses on obtaining particle-scale fluid properties from filtered properties by utilizing the exact location of each particle, which is available in an Euler-Lagrange simulation. By utilizing direct numerical simulations (DNS), one can develop models to predict the interphase forces and Reynolds stress. These microscale quantities represent the terms of the macroscale fluid equations that must be modeled. The focus here is primarily on force modeling, while the Reynold stress will be considered in a future report.

The generalized Faxén's theorem allows for the expression of the particle force terms (quasi-steady, added mass, etc.) as functions of surface or volume averaged fluid properties. These fluid properties (velocity, vorticity, pressure, and stress) must be the *undisturbed* fluid properties. That is, the properties of the fluid without the presence of the particle of interest. For this reason, the particle of interest is hereby referred to as a *ghost particle* which should not disturb the flow. This name should not be confused by the computational ghost particles used to scale Euler-Lagrange simulations to multiple processors.

In a previous work, the pairwise-interaction extended point-particle (PIEP) model was introduced [1]. This model gives one the ability to account for fluid mediated particle-particle interactions that lie in the unresolved length-scale of Euler-Lagrange simulations. When constructing force maps for the PIEP model, the wake of an isolated particle was utilized along with the generalized Faxén's theorem. By placing a ghost particle in the wake of the isolated particle, the influence of the isolated particle on the ghost particle's force can be approximated with Faxén's theorem. The ghost particle is then moved throughout the isolated particles wake, and the forces are recomputed at various positions. By this means, a force map for the ghost particle is produced.

However, the force maps produced by an isolated wake are not accurate at high particle volume fractions (greater than 20%). This loss of accuracy can be explained by the following two hypotheses: (i) the wake of an isolated sphere is not relevant for a densely packed bed of particles. (ii) the Faxén's theorem approximation assumes that the ghost particle will not have any tertiary interactions with its neighboring particle. At higher volume fractions, one would expect these tertiary interactions to play a larger role.

This work focuses on hypothesis (i) concerning the wake of a particle at high volume fractions. Since the isolated particle's wake appears to be invalid at high volume fractions, a *representative wake* that is valid at these volume fractions must be found. The representative wake of a particle at a higher volume fraction will differ from that of a lower volume fraction. This is due to the increased presence of particles on all sides of the disturbing particle. To help determine the representative wake, DNS must be utilized.

The DNS was performed for a random array of monodispersed particles. The DNS data was produced by Akiki et al [2]. In that work, a random array of equally sized particles was distributed in a cubic domain with uniform probability. Periodic boundary conditions were implemented in the flow direction (y-direction) and a direction normal to the flow (x-direction). A no-stress boundary condition was implemented at the boundaries in the z-direction. The fluid was fully resolved around each particle by solving the incompressible Navier-Stokes equations with no-slip and no-penetration boundary conditions on each particle's surface. These no-slip and no-penetration conditions were satisfied using the immersed boundary method. To avoid the effects of the boundary conditions in the z-direction, only the middle 64% percent of the domain (in the z-directions) were considered for modeling purposes. This DNS was performed at various particle volume fractions and Reynolds numbers. For the range Reynolds numbers discussed in this report, the flow was steady. An example of this DNS can be found in Figure 3.3. With ample amounts of DNS, one can use the DNS to obtain a representative wake.

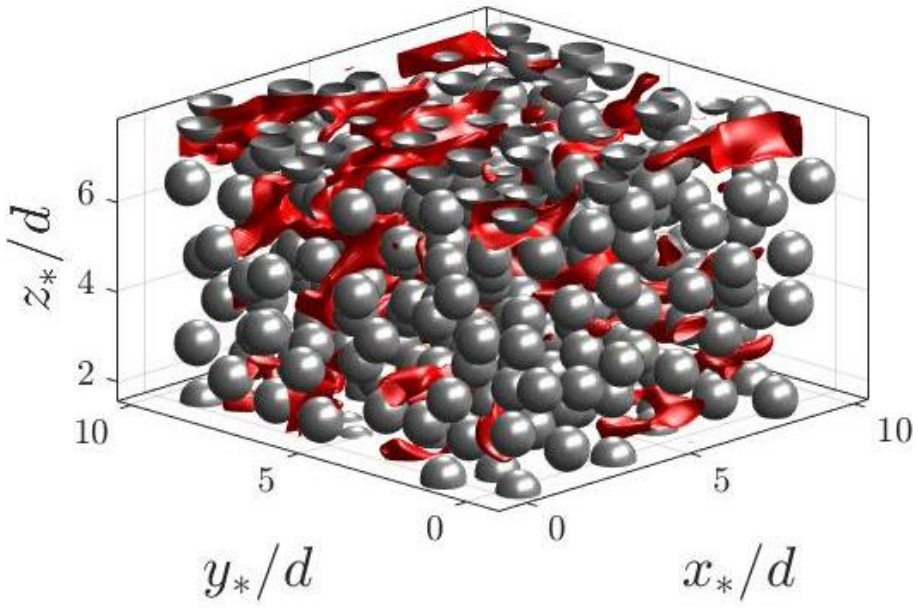

Figure 3.3: Isosurface plot of $|\mathbf{u}|/\langle|\mathbf{u}|\rangle = 1.6$ for Re = 16 and a particle volume fraction of 0.21

The simplest approach to find a representative wake is to take the average wake of each particle within the DNS. However, this approach is not useful for force modeling. The average wake accounts for the statistical disturbances of all neighboring particles. Therefore, the wake is no

longer *undisturbed* as the Faxén's theorem requires. It will statistically account for the presence of the ghost particle. For this reason, it is not useful for computing force maps. An example of a non-dimensional average wake can be seen at in Figure 3.4. As observed in the figure, the statistical disturbance of neighboring particles causes a "ringing" effect at higher volume fractions.

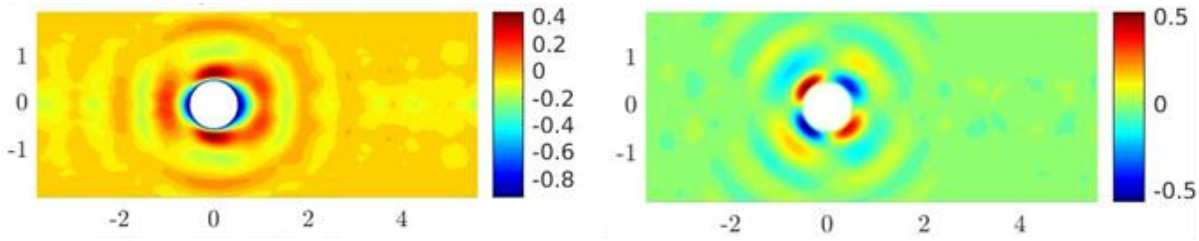

(a)Streamwise velocity *(b) Perpendicular velocity*

Figure 3.4: Non-dimensional average wake contours for a volume fraction of 45% and Re of 21 where the x-axis is the streamwise direction and the y-axis is the perpendicular direction.

Alternatively, an implicit approach can be utilized to determine the representative wake. This approach assumes that there is a linearly superposable wake (LSW), ***u****'(x)*, such that

$$\boldsymbol{u}(\boldsymbol{x}) = \boldsymbol{u}_0 + \sum_{k=1}^{N_p} \boldsymbol{u}'\left(\boldsymbol{x} - \boldsymbol{x}_{center\ ,k}\right) \ .$$

In this manner, ***u****'(x)*, can be superposed for each particle, and the velocity field, ***u****(x)*, of the DNS can be predicted. This assumption can be extended to the pressure of the fluid as well. Furthermore, the LSW can be determined via regression analysis, where a functional form of ***u****'(x)* is proposed, and the parameters within that functional form are optimized. An example of a LSW can be seen in Fig. 3.5. Note that the LSW is free from the "ringing" present in the average wakes in Figure 3.4.

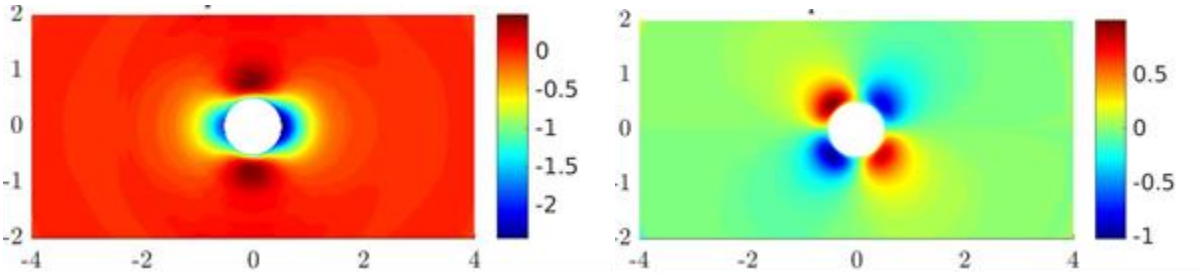

(a) *Streamwise velocity* *(b) Perpendicular velocity*

Figure 3.5: Linearly superposable wake contours for a volume fraction of 45% and Re of 21 where the x-axis is the streamwise direction and the y-axis is the perpendicular direction

Once regression is performed, the linearly superposable wakes were used to develop a force model utilizing the PIEP model framework. The resulting model was then applied to the DNS data to determine its accuracy. The resulting predictions from this new PIEP model were compare to the

PIEP model that uses a wake from an isolated particle [2]. Table 3.1 reports the R^2 values for the force in the stream-wise direction (drag) and force in the x-direction (lift). It includes the isolated wake's R^2 values along with the new model's R^2 for the sake of comparison. As expected, the LSW is more accurate than the isolated wake at high volume fractions. However, the models perform comparably at lower volume fractions.

		Isolated Wake		***LSW***	
φ	***Re***	*Drag*	*Lift*	*Drag*	*Lift*
10%	***40***	0.68	0.35	0.58	0.45
10%	***70***	0.65	0.29	0.61	0.35
10%	***173***	0.60	0.05	0.54	0.08
20%	***16***	0.41	0.29	0.71	0.60
20%	***89***	0.52	0.19	0.66	0.30
45%	***21***	0.12	0.02	0.56	0.43
45%	***115***	0.17	0.01	0.55	0.36

Table 3.1: R^2 values for PIEP model results for an isolated wake model and LSW model

In a previous report, the PIEP model was extended to higher volume fractions by using a data-driven approach to extract force and torque maps from stationary particle DNS data. These data-driven force maps were superposed onto the existing PIEP model to construct a *hybrid*, data-driven physics-informed PIEP model. These data-driven force maps are not separated into particles force terms (quasi-steady force, added mass force, pressure gradient force, etc.) like the LSW PIEP model. The particle force terms scale differently with particle motion and, therefore, the data-driven force maps cannot be appropriately scaled for dynamics simulations. However, the LSW PIEP model allows one to appropriately scale the force components for dynamic simulations. That being said, the data-driven force maps can be used to improve the LSW PIEP model for stationary cases. The R^2 values for the *hybrid* PIEP model with the LSW are reported in Table 3.2.

φ	***Re***	***Drag***	***Lift***
10%	***40***	0.69	0.76
10%	***70***	0.65	0.70
10%	***173***	0.55	0.65
20%	***16***	0.80	0.77
20%	***89***	0.70	0.65
45%	***21***	0.65	0.55
45%	***115***	0.66	0.54

Table 3.2: R^2 values for hybrid (with data-driven component) PIEP model with LSW

Likewise, the superposable wake can also be used to approximate the unfiltered flow field with the use of the wakes and the filtered flow properties. Let us consider each DNS to be the equivalent of one discretized volume inside a large-scale simulation. After filtering the DNS, the mean particle Reynolds number and volume fraction would be available. However, the sub-grid velocity and pressure fluctuations are not be available since they were filtered out. The superposable wake can be superposed at every particle location to reconstruct the unfiltered flow field from filtered properties. An example of this is seen in Figure 3.6 where a slice of DNS data is presented. This reconstruction process was tested for a range of Reynolds numbers and volume fractions. The corresponding coefficient of determination, R^2, for each Reynolds number and particle volume fraction is reported in Table 3.3. An R^2 value of 1 indicates the model is exact while a value of 0 indicates that the model fails to predict any of the variation in the fluid properties. Note that the DNS used in Figure 3.6 and Table 3.1 is not the same DNS used to develop the superposable wake.

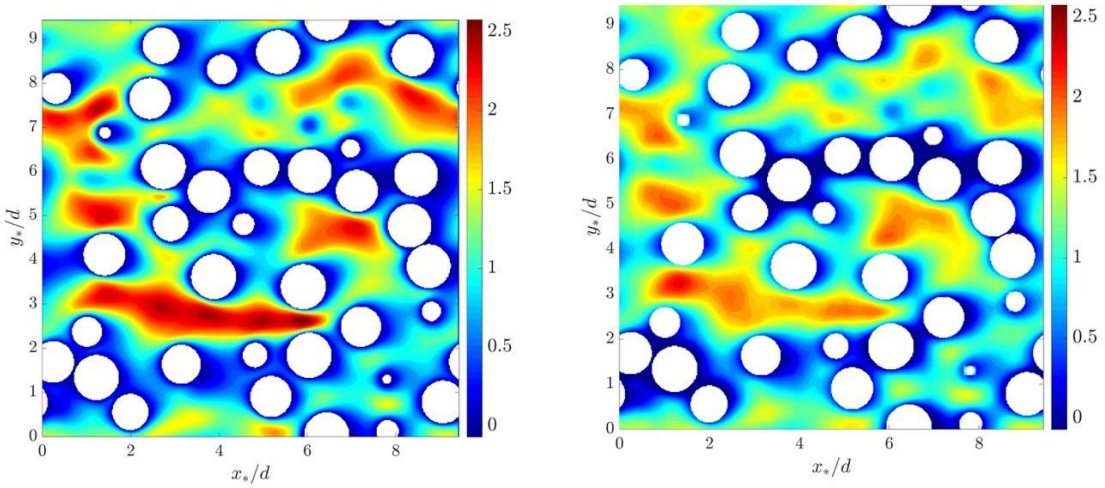

(a) *DNS* *(b) Superposable wake reconstruction*

Figure 3.6: Nondimensional streamwise velocity profiles at $z_*/d = 6.02$ for Re = 16 and a particle volume fraction of 0.21

φ	Re	$v_{streamwise}$	v_{perp}	p
10%	40	0.77	0.56	0.76
10%	70	0.70	0.48	0.74
10%	173	0.59	0.36	0.71
20%	16	0.83	0.71	0.58
20%	89	0.69	0.56	0.70
45%	3	0.71	0.61	0.66
45%	21	0.68	0.63	0.65
45%	115	0.56	0.54	0.60

Table 3.3: Resulting R^2 values for the superposable wakes for streamwise velocity, perpendicular velocity, and pressure

To summarize, even in the case of laminar flow, particle-scale variations in fluid properties will exist due to the presence of particles. When performing large-scale simulations, these particle-scale variations are filtered out to reduce the cost of the simulation. By using the superposable wake, these filtered variations can be reconstructed from the filtered properties. This superposable wake was computed from DNS data of flow around an array of particles via regression analysis. This superposable wake is useful for particle force modeling. The PIEP model is currently one of the best models for approximating interactions of particles in an array for particles, but the PIEP model loses accuracy at higher volume fractions. Prior to this reports findings, the PIEP model relied on the wake of an isolated particle to develop its force maps. This isolated particle's wake does not represent the wake of the particles at higher volume fractions. Once the superposable wake is implemented into the PIEP model's framework, the PIEP model can make more accurate predictions at higher volume fractions.

3.4 Shock interaction with a deformable particle

Physical behavior in shock-particle interaction has been studied extensively with the approximation of non-deformable particle. However, the deformation effect becomes important when the shock pressure becomes higher than the yield stress of the particle material. Such situations appear in high shock Mach number flows in air as well as shock–particle interactions in a liquid medium. To understand the interaction of a shock with a deformable particle, fully-resolved multi-component inviscid simulations were performed. Analysis of the simulation was primarily focused on baroclinic vorticity production and its influence on unsteady drag coefficient.

Numerical method

The dynamics of an aluminum spherical particle exposed by a shock at Ms = 1.6 in Nitromethane is investigated utilizing a high–resolution axisymmetric solver for the Euler equations that allows for multi–material interface and shock propagation in both the particle and surrounding medium. Viscous effect is neglected in the present study for high Reynolds number flow in early stage of shock-particle interaction.

$$\frac{\partial(\rho_k\phi_k)}{\partial t}+\nabla\cdot(\rho_k\phi_k\mathbf{u})=0,$$

$$\frac{\partial(\rho\mathbf{u})}{\partial t}+\nabla\cdot(\rho\mathbf{u}\mathbf{u})+\nabla P=0,$$

$$\frac{\partial E}{\partial t}+\nabla\cdot(\mathbf{u}(E+P))=0,$$

$$\frac{\partial\phi_1}{\partial t}+\mathbf{u}\cdot\nabla\phi_1=0,$$

$$E=\rho\left(e+\frac{1}{2}\mathbf{u}^2\right),$$

$$P = (\gamma - 1)\rho e - \gamma P^{\infty}.$$

Here ρ_k is the k_{th} component density, **u** is the velocity, P is the pressure, E is the total energy, φ_k is the volume fractions of materials, e is the specific internal energy, γ and P^{∞} are the parameter for the stiffened gas equation of state.

The interface compression technique is applied to the field of density and volume fraction to reduce the smearing effect at the material interfaces.

$$\frac{\partial \phi_1}{\partial \tau} = \mathbf{n} \cdot \nabla \left(\varepsilon \left| \nabla \phi_1 \right| - \phi_1 \left(1 - \phi_1 \right) \right),$$

$$\frac{\partial \rho_k \phi_k}{\partial \tau} = H\mathbf{n} \cdot \nabla \left(\varepsilon \mathbf{n} \cdot \nabla \rho_k \phi_k - \left(1 - 2\phi_k \right) \nabla \rho_k \phi_k \right),$$

where H is the Heaviside function, **n** is the normal to the interface, τ is the pseudo–time for the iterative relaxation, and ε is a grid dependent parameter. Computational condition is shown in Figure.3.7.

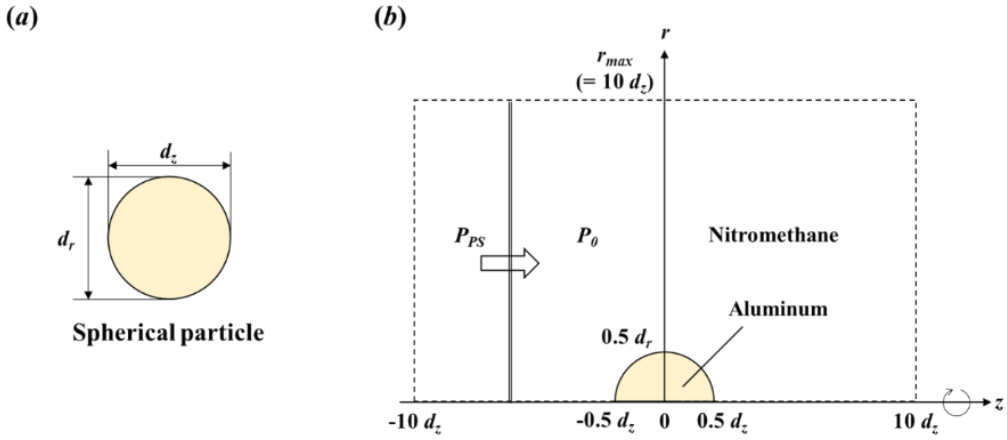

Figure 3.7. Computational condition.

Mechanism of quasi-steady drag by baroclinic vorticity

Figure 3.8 shows the time–evolution of baroclinic vorticity and pressure for the case of a Ms = 1.6 shock passing through a particle. As the shock propagates over the particle, the baroclinic vorticity is generated due to the large misalignment of density gradient and pressure gradient in the flow field (not shown). At t = 5, the baroclinic vorticity generated at the material interface starts to move to the downstream owing to the large velocity near the top of the particle. Simultaneously,

sharpened edge of the particle is formed owing to the rotational flow arising from the moving vortex. It is clearly observed that the vortex is trapped by the particle edge after t = 6. This trapping of vortex leads to quasi-steady drag even after the passage of the shock.

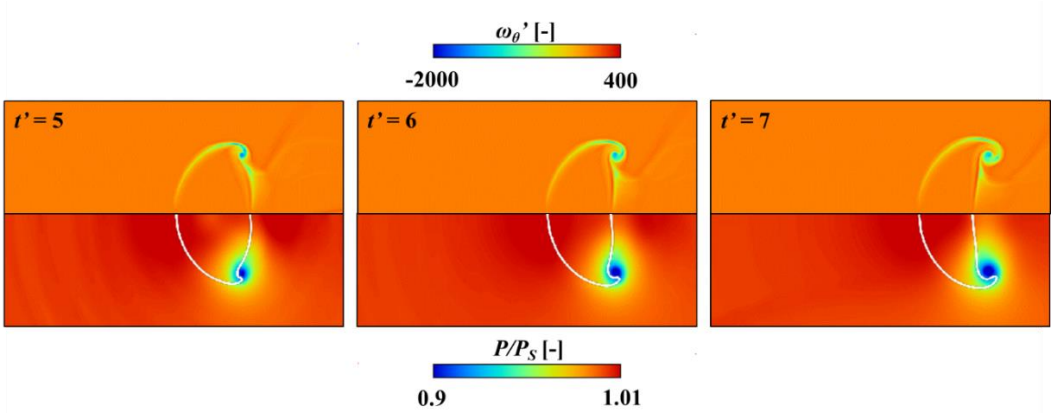

Figure 3.8. Baroclinic vorticity (upper) and pressure (lower) plots.

Figure 3.9 shows the unsteady drag coefficient for the case of a Ms = 1.6 shock passing through a particle. The drag coefficient increases and reaches its maximum value when the incident shock interacts with the particle. After the passage of the shock, drag coefficient rapidly decreases. Then, the drag coefficient reaches its second peak at the instant of vortex trapping. Finally, the drag coefficient gradually decreases to zero. These results indicate that the compressible flow pattern dominates the early stages of the unsteady drag coefficient in the shock–particle interaction and the baroclinic vorticity production causes the quasi steady drag in later time. These results suggest the importance of baroclinic vorticity production to the unsteady drag coefficient in shock–particle interaction after the passage of shock.

Current models do not consider the viscous effect as well as the three-dimensional effect. To better understand the vorticity production in shock-particle interaction, we plan to carry our three–dimensional viscous simulations in future work.

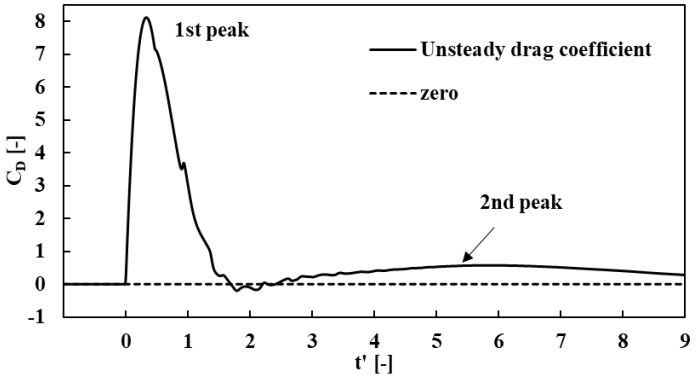

Figure 3.9. Unsteady drag coefficient.

4. Experiments

4.1 ASU Experiments

4.1.1 Goals and Motivation

Turbulent multiphase flows are highly complex. The simulation team at the University of Florida is working to simulate an exploding cylindrical charge in its entirety. This multifaced problem has been broken down into several smaller segments, including the decompression of a densely-packed particle bed.

The experimental team at Arizona State University is working to gather experimental data on the decompression of a densely-packed bed to help validate the early stage codes being developed at the University of Florida.

A particle bed composed of small (< 1mm diameter) glass beads is placed in the bottom of a vertical shock tube. A diaphragm is placed above the particle bed and everything above the diaphragm is vacuumed down to a low pressure (p_1 < 20kPa). Everything below the diaphragm remains at atmospheric pressure (p_4 = 101.3kPa). Once the diaphragm is ruptured, the pressure differential causes material (air and particles) to rapidly accelerate. The time frame of interest during each realization is on the order of milliseconds.

4.1.2 Introduction to ASU Experiments

The Laboratory for Energetic Flow and Turbulence at Arizona State University is using a vertical shock tube to examine the expansion fan behind a shock wave. The expansion fan travels into the high pressure region, which contains a particle bed of small glass beads. The shock wave will travel upward into the low pressure region. High-speed video, pressure sensor data, and PIV data is recorded.

4.1.3 New Equipment over 2017

The new pressure sensors and data acquisition device arrived at ASU's experimental facility. With this new setup, the team is able to obtain pressure data with an appropriate temporal resolution required to see the dynamic pressure changes, including shocks and reflected shocks. Additionally, with the new computer, we are able to synchronize the pressure sensor data and the video data. As of April 2017, all the timing cards, data acquisition cards, and data acquisition devices are connected and working properly.

4.1.4 PIV Experiments

In order to explore the gas dynamics in the shock tube, we setup particle image velocimetry (PIV) experiments. Our PIV experiments use an illuminated planar cross-section (approximately 2mm thick) of the flow with a Litron Nano L 200-15 pulsed Nd:YAG laser. The laser is synchronized with a high-speed camera so the flow is sufficiently illuminated when images are collected. As the laser illuminates both glass particles (from the bed) and tracer particles (passively following the gas flow), we use 10 micron (D_{50} = 10 μm) silver-coated hollow glass spheres to serve as passive gas tracer particles. These tracer particles are sufficiently smaller than the O(100 micron)

glass particles, so we are able to differentiate between the gas tracer particles and glass particles by using a size threshold in the PIV analysis.

A possible change to future PIV experiments could include using fluorescent particles with wavelength-filtered cameras to better distinguish between the particles. An additional method we are considering would be using smoke to trace the gas flow, similar to preliminary experiments that were performed in Fall 2016. By using smoke, we could perform techniques based on laser induced fluorometry (LIF) to evaluate the evolution of the gas concentration as it compresses and decompresses along the shock tube.

By using these optical techniques, we have several main objectives to accomplish. Using a vertical light sheet that passes through the center of the tube allows for 2-dimensional 2-component velocity measurements of both the gas and dilute particle bed. It is critical not to perform these measurements in a dense particle bed, as neither the laser light sheet nor camera will be able to penetrate through the measurement region of interest. Thus, these measurements are performed above the bed. A schematic of this setup is shown in Figure 4.1. With resulting PIV analysis of glass particle and gas tracer particle velocity fields, we will be able to quantify the transition to turbulence and relationship between gas and particle motions. Similarly, in this setup, if we use smoke or other gas tracers, we can perform LIF-based measurements of vertical motions in the shock tube.

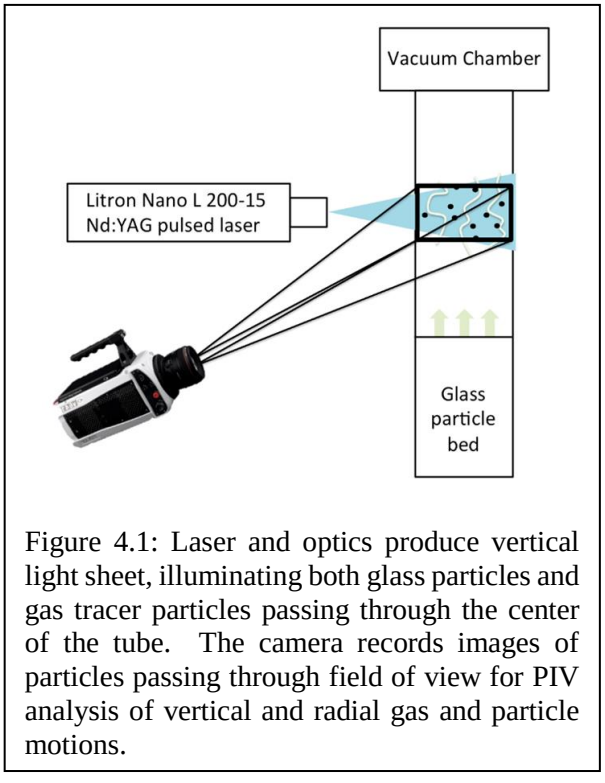

Figure 4.1: Laser and optics produce vertical light sheet, illuminating both glass particles and gas tracer particles passing through the center of the tube. The camera records images of particles passing through field of view for PIV analysis of vertical and radial gas and particle motions.

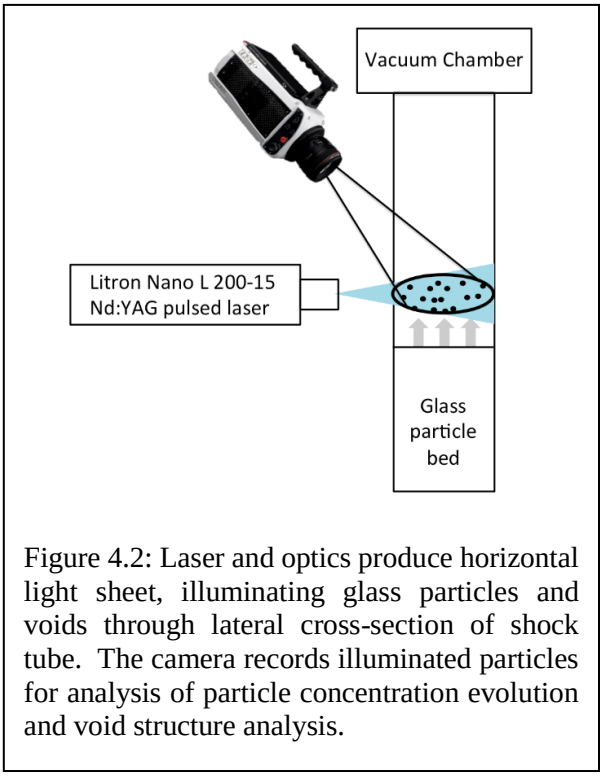

Figure 4.2: Laser and optics produce horizontal light sheet, illuminating glass particles and voids through lateral cross-section of shock tube. The camera records illuminated particles for analysis of particle concentration evolution and void structure analysis.

Alternately, we can orient the optics such that the laser produces a horizontal light sheet through the shock tube, as we aim to measure the flux of glass particles as a function of time and distance from the initial bed height. As prior optical measurements performed by Heather Zunino have focused on the front of the cylindrical tube, this will allow us to see the entire cross-section of the flow and center of the tube, as long as the glass particle bed is sufficiently dilute to allow for penetration of the laser light sheet and optical access of the camera. These measurements will be advantageous in that voids in the interior of the tube may be visible and quantifiable. A schematic of this setup is shown in Figure 4.2. Note that this method is not meant to incorporate PIV analysis techniques, but it is possible we will be able to detect lateral motions through these experiments.

With the new timing system purchased from National Instruments, we will be able to synchronize the optical measurements with the pressure sensors. When the diaphragm bursts, the significant pressure drop above the diaphragm will be recorded via the PXIe-4492 card, and this signal will trigger the PCI-6602 card to initiate optical measurements by sending synchronized laser and camera pulses. The pressure sensors were not previously linked to optical measurements, so the new system provides an advantage in providing the relative timing between the diaphragm burst and onset of motion in the glass particle bed and in the surrounding gas.

In the initial PIV experiments Images were collected with the Phantom v641 high-speed camera. We synchronize the camera with the laser output by using a Digital Delay / Pulse Generator developed by Berkeley Nucleonics Corporation (BNC). The pulse generator has 8 available outputs, 4 of which control the laser flash lamps and Q-switches for each laser head, and 2 of which control the camera (one pulse per frame within a pair of images). At first, the PIV data collection operated continuously during experiments but wasn't yet synchronized with the diaphragm burst or inception of the shock or rarefaction waves. The BNC signal generator has an available port to receive a trigger from an external source, and so we were able to incorporate a trigger generated by the pressure signal to be able to relate the timing of PIV images to the physical processes occurring within the shock tube. The initial images are shown below in Figure 4.3.

In exploring the timing of image collection in PIV, we found we are limited to a sampling frequency f_S = 14.1 Hz for maximum laser power, but we may be able to increase this frequency by diminishing the intensity levels of the laser heads. Within an image pair, we are able to use separation times (Δt) as low as approximately 5 microseconds. We are expecting to use values of Δt around 50 μs to 100 μs according to our understanding of the speed of the waves from previous pressure measurements in Heather Zunino's experiments.

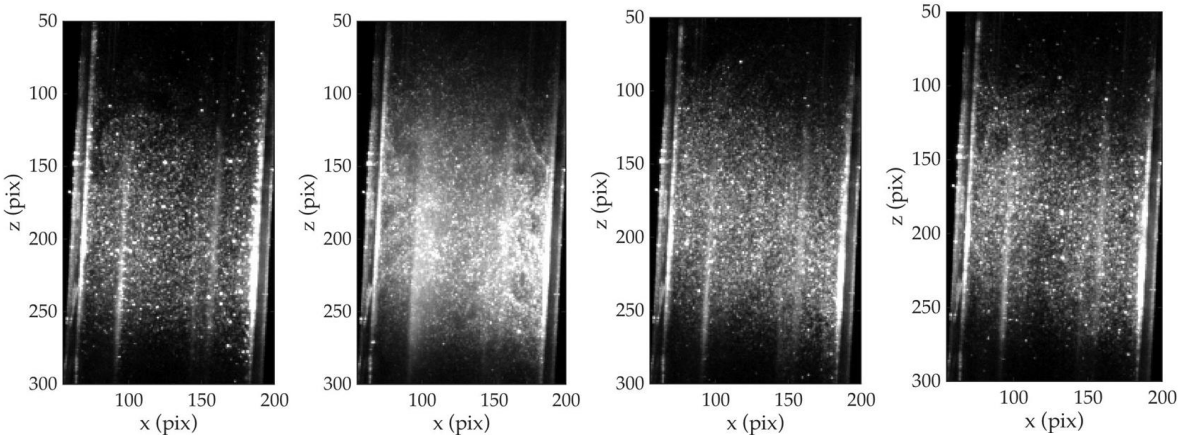

Figure 4.3: Sample images of 10 μm illuminated tracer particles in response to a pressure ratio P4/P1 = 5. The first image on the left is the first image with visible particles after the diaphragm burst. Each subsequent image from left to right is separated by a lag of 40 ms. This is an air-only experiment, with tracer particles initially at rest at the base of the shock tube (~60 cm below the diaphragm); for future experiments with particles, the tracers will be seeded at the top of the bed. Note also that these images do not show image pairs with correlated velocity fields; future experiments will be collected in pairs to deduce gas motions within images separate by Δt.

While preparing the PIV experimental setup, we also worked to improve the shock tube itself, replacing parts of the nylon joints to ensure vertical alignment of the tubes and adequate sealing between sections. We also mounted a snorkel exhaust fume hood near the pump to collect stray particles that escape during the depressurization, as the silver-coated particles can be quite dangerous if inhaled.

4.1.5 Particle Image Velocimetry Triggered with Pressure Sensor Data

The team at ASU was able to trigger a particle image velocimetry (PIV) experiment with the pressure sensors, meaning we can measure the gas flow in the expansion region relative to the shock location in the shocktube. A brief set of PIV frames, along an image of the difference between the two frames is shown below in Figure 4.4.

4.1.6 Horizontal Void Cracks

As the dense particle bed expands, it has several interesting flow features. When the rarefaction wave impinges on the particle bed interface, the bed begins to swell and then is broken up by "horizontal void cracks." These cracks then expand and break down further into smaller cells. By using some simple image processing techniques, the horizontal void cracks that appear in experiments with varying initial pressure ratios, particle sizes, and initial bed are plotted Figures

4.5 – 4.7 below. The dark streaks are the cracks near the edge of the glass cylinder as they grow and move in time.

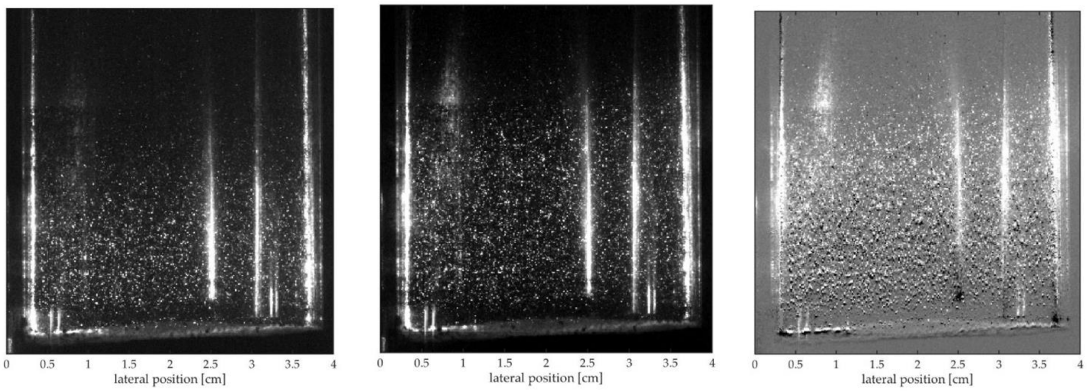

Figure 4.4: Shown are two PIV frames (left and middle) and the difference between the two frames (right), in which the upward motion is represented as traveling between dark pixels to bright pixels. The timing between the two PIV frames is 80 microseconds. The delay between the shock triggering the system to record and the first frame is 300 microseconds. The vertical field of view is approximately 3.5cm of gas and 0.5cm of the particle bed.

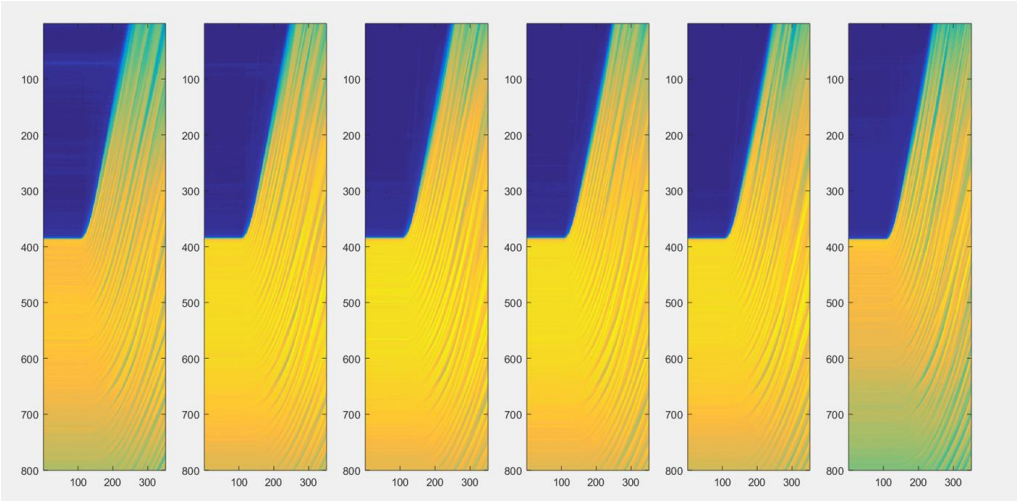

Figure 4.5: Bed composed of 212-297 micron particles. The independent variable in these plots is time (in seconds x 10^{-4}) and the dependent variable is the axial coordinate along the length of shocktube (in pixels). Each of the six figures above is created by averaging the intensity across 25 pixels in a single image, plotted as a single column, and then the process is continued for nearly 350 frames. The first slice above contains information from the left wall of the shocktube to the first 25 pixels, the second slice is from 25 to 50 pixels, and the pattern continues until the right wall of the shocktube.

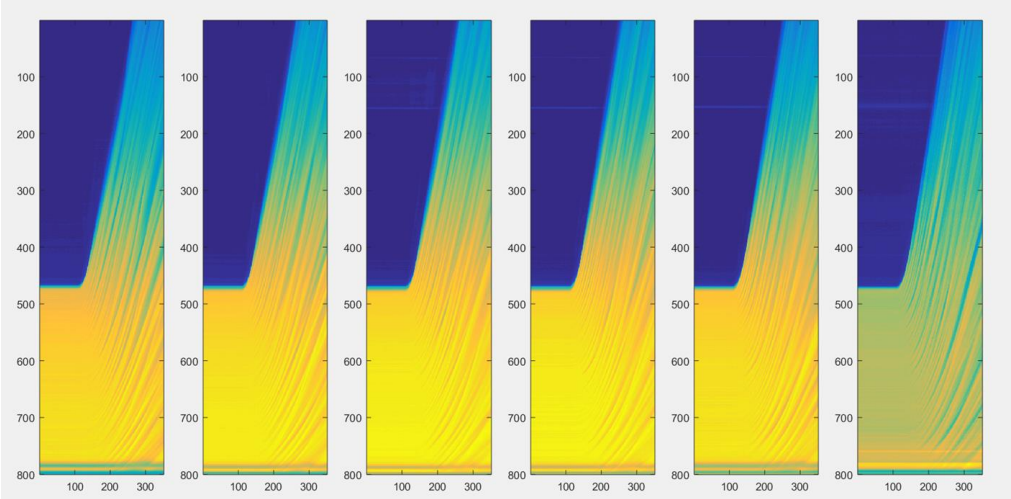

Figure 4.6: Bed composed of 150-212micron particles. This figure is created the same way as Figure 4.4.

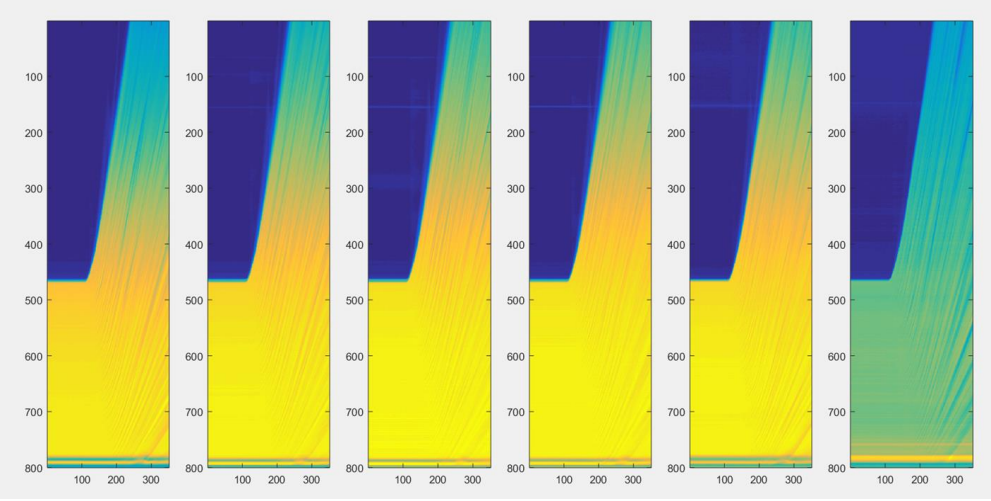

Figure 4.7: Bed composed of 44 - 90micron particles. This figure is created the same way as Figures 4.5 and 4.6.

4.1.7 Pressure Sensor Data Examined with Differing Bed Heights

The pressure sensor data from 5 realizations of same experiment were ensemble averaged. Contained in each experimental data set, is the data from any redundant pressure sensors located at the same axial location along the shocktube. This yields four pressure traces. There are two above the diaphragm (1 pressure sensor at the very top and three 32cm above the diaphragm), showing the shock, and two below (2 pairs of pressure sensors), showing the rarefaction wave. See Figures 4.8-4.10.

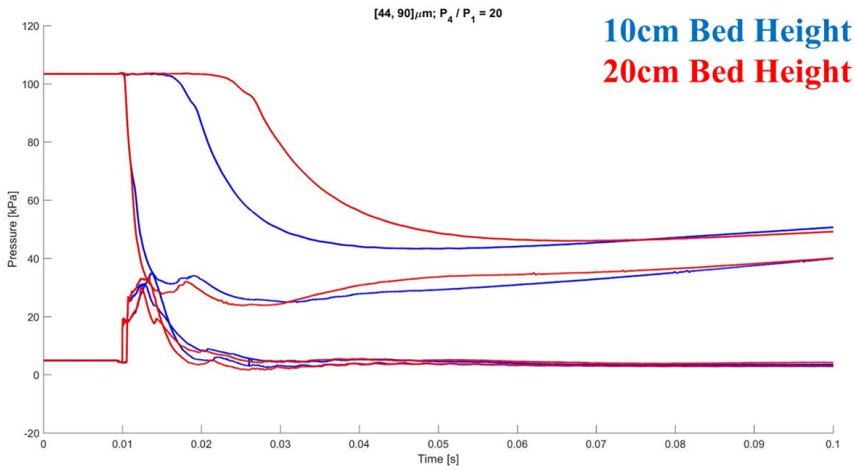

Figure 4.8: Two experimental data sets with the same particle size range [44, 90] microns, but differing initial bed heights. Notice that the pressure traces located from above particle bed, seem unaffected. The rarefaction wave travels more slowly though the taller particle bed.

.

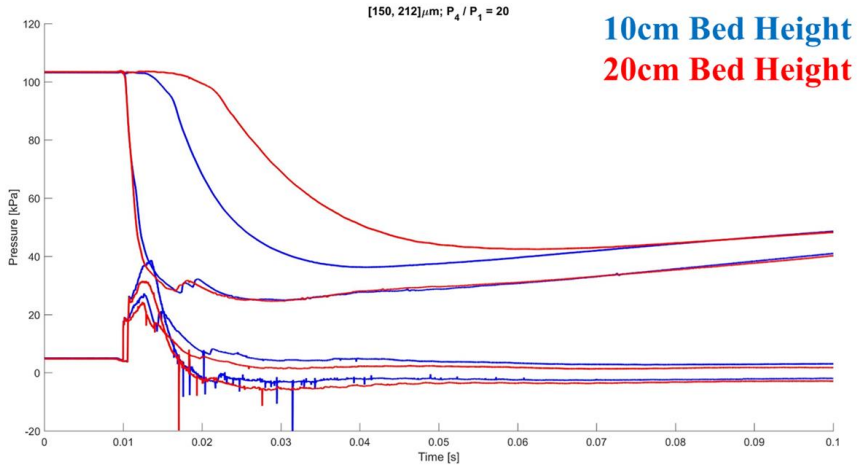

Figure 4.9: This plot is the constructed in the same way as Figure 4.8, however the particles in these experiments range between [150, 212] microns. The same effect is seen, although it is not as dramatic.

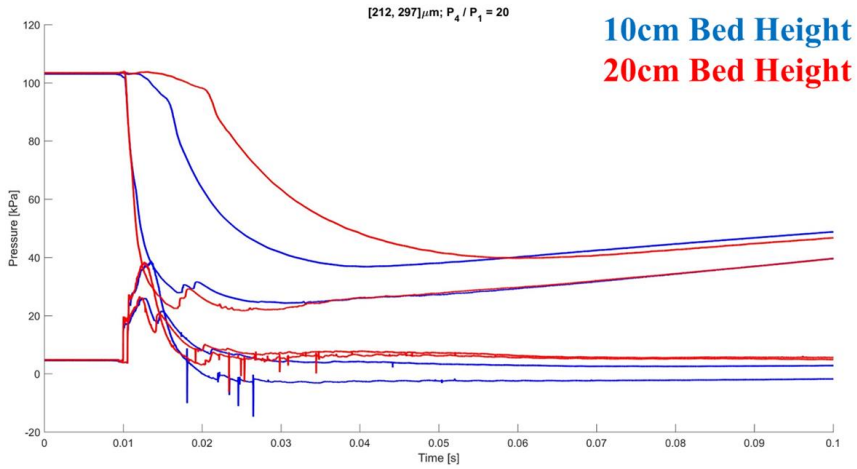

Figure 4.10: This plot is the constructed in the same way as Figures 4.8 and 4.9, however the particles in these experiments range between [212, 297] microns. Again, the same effect is seen, although it is even less dramatic.

4.1.8 Main Shocktube Setup

Additionally, new glass segments were ordered and installed on the shocktube. The new glass segments have minimal optical defects and have a much more uniform thickness and length than the previously installed glass segments. The new segments have been carefully measured and documented and that data has been sent to the team at UF.

With the new segments, the shocktube has been realigned so that the axis along the length (z – axis) is nearly perfectly perpendicular to the ground. There was an issue with some of the glass segments and connections not being perfectly straight, but that issue has been corrected and documented.

4.2 Eglin AFB Experiments

4.2.1 Goals and Motivation

The primary goal of the experiments conducted at Eglin Air Force Base is to provide validation quality data at the micro, meso, and macroscales. AFRL/RW has completed initial experiments at the micro- and meso-scales as described in this section, and the data have been given to the UF-CCMT modeling team for uncertainty analysis. The experiments include:

a. Microscale experiments with a no-particle, detonation-only, planar geometry in support of T3, uncertainties due to thermodynamic (EOS) and transport properties;
b. Mesoscale experiments with a planar particle bed in support of T5, compaction modeling uncertainty;
c. Microscale experiments with a planar particle cluster in support of T6/T7, uncertainty in finite Re, Ma and volume fraction dependent drag and heat transfer; and
d. Macroscale experiments with an annular particle bed in support of T0, validation of the demonstration problem.

4.2.2 Microscale Experiments

Twelve small-scale explosive experiments were performed at Eglin AFB in February 2015. These experiments are considered microscale in that a small number of particles are of interest. The test data include shock arrival times, high-speed video, x-ray images, and witness panel results. The twelve experiments are equally divided into gas-valve tests (compressed helium, tests 1-6) and explosively-driven test (tests 7-12). Table salt was substituted for tungsten particles as noted in the table in an attempt to visualize the gas flow patterns with the fine particles acting as tracers.

The microscale experiments were conducted at the Advanced Warhead Experimental Facility (AWEF) at Eglin AFB, FL. All tests utilize the same basic experimental setup. The pressure array is a 3x3 grid of pressure probes, slightly offset from the test centerline with the middle probe removed due to the high likelihood of fragment impact. Four x-ray heads were used to increase the ability to track the particles' location and velocity within the fireball.

Compressed helium provides an optically transparent shock wave and was used in the gas-valve experiments refine the settings on the diagnostic equipment, particularly the cameras. During the second day of testing, the compressed helium driver was replaced by an explosive charge. For

these experiments the explosive charge consisted of three stacked N5 pellets (each 0.5” in length and diameter) initiated by an RP-83 detonator.

Data from the experimental diagnostics are detailed below. For test 8, one of the delay generators malfunctioned, resulting in a loss of Simacon camera data. The delay generator was replaced in shot 9, but an improper arming sequence resulted in a loss of all data.

The pressure probes from the compressed helium tests show a relatively weak shock (≤1 psi) and a sharp pressure rise. Alternately, the pressure probes from the explosive tests show a much stronger shock (8-23 psi) with a complex pressure rise, exhibited in a non-noise signal, followed by an oscillation, then a sharp rise in signal.

High speed images were recorded for the Phantom 6.11, Phantom Miro M310 and SIMACON. The Phantom 6.11 was used to capture the detonation products down range. The SIMACON camera was used to capture the blast wave at close range, but has a faster framing rate (but limited number of images) than the Phantom Miro M310. The Phantom Miro M310 was used to capture the blast wave at close range.

Four x-ray heads were mounted above the particle driver. The x-ray heads were triggered on a timing delay of 10-20 microseconds; the particle velocity is of primary interest. The large number and small size of the salt particles in test 10 and 11 precluded accurate velocity measurements.

Thin aluminum witness panels were used to determine the particles’ far-field impact locations in tests 7, 8, 9, and 12. The witness panels were placed 66” from the particle driver, and the center of each panel was determined using a bore sighting laser. No particles were located or recovered.

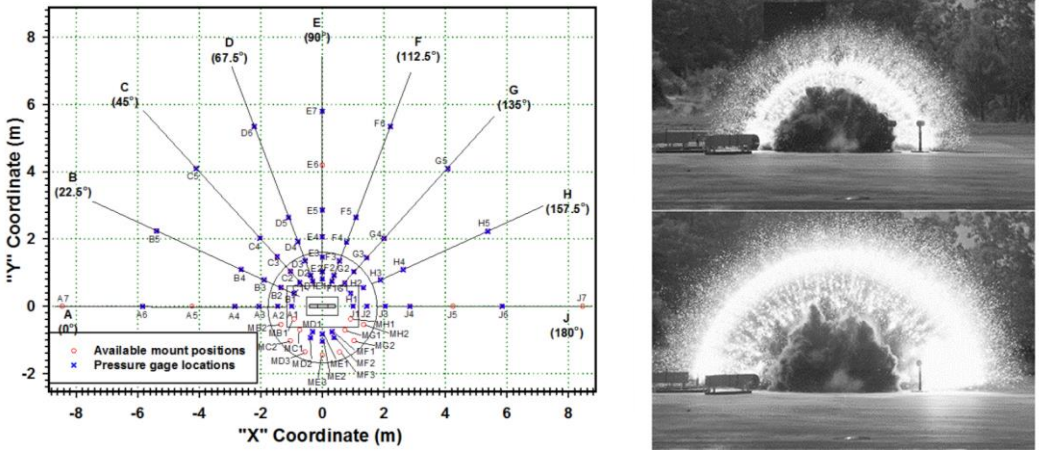

Figure 4.11. Left: Locations of 54 pressure transducers. Right: Two time sequence of typical blast pad experiment.

4.2.3 Mesoscale Experiments

Twenty-two mesoscale explosive experiments were performed at Eglin AFB in October-November 2015. The diagnostics and setup are the same as the microscale experiments. The 22 experiments are divided into gas-valve tests (tests 1-12) and explosively-driven tests (tests 13-22). The first tests were completed with a coarse steel powder. Tungsten powder was used for the remaining tests, where the best configuration of Tungsten required some experimentation, as seen in tests 8, 9, and 10. It was determined that the Tungsten powder was ejected most consistently when pressed in alcohol between sheets of tissue; tests 10-22 used this method.

4.2.4 Macroscale Experiments

The macroscale experiments were completed this past summer. Kyle Hughes (UF) was present at Eglin AFB during the testing to direct the testing so that satisfactory uncertainty quantification could be performed in future work. All six shots were successfully completed with accompanying

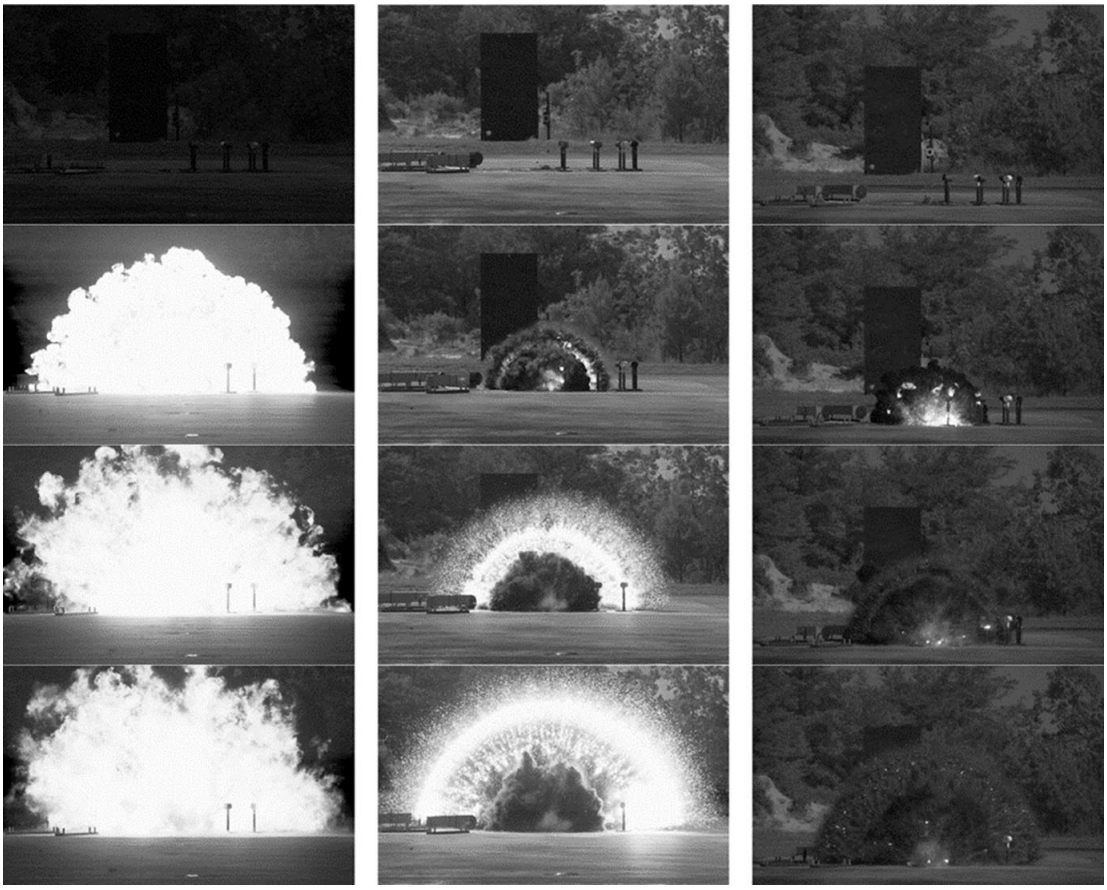

Figure 4.12. High-speed imagery of the three test configurations viewed down the long axis of the cylindrical charge. Time stamps (from top to bottom) are 0, 1.67, 2.42, and 3.67 ms for a) bare Composition B, b) Composition B with tungsten liner, and c) Composition B with steel liner (note magnifications are not equal for all test conditions).

video and pressure traces collected. Figure 4.11 shows the locations of the pressure probes for each of the experiments, as well as time sequence of a typical experiment. This data is currently being analyzed. Figure 4.12 shows the time history of the event.

4.2.5 Summary

The micro and mesoscale experiments performed at Eglin AFB, FL, present the UF CCMT staff with a first-look at diagnostic and measurement techniques. The ultimate objective is to provide the UF-CCMT with high quality data for validation and verification of their turbulent multiphase models.

5. UB Team

5.1 Summary of Achievements

a. Validation and uncertainty quantification (UQ) of the mesoscale shock tube simulation
    a. Wrapping up validation and UQ of the mesoscale shocktube simulation
    b. Found that UQ was very useful for identifying modeling errors in the simulation under development
    c. Found that experimental data is often obtained by elaborate measurement processing, which is one of major uncertainty sources. The paper that emphasizes the significance of measurement processing uncertainty has been accepted in the Journal of VVUQ
    d. Achieved the grid convergence of the hydrodynamic force in a Euler-Lagrange simulation for a single particle by using a finite-size particle model
    e. Quantified the discretization error for quasi-steady, pressure gradient and added mass force for a single particle
    f. Investigated the effect of the finite size of the particle on the convergence of the point particle models for the shock-particle interaction problem with many particles and moderate volume fraction
    g. Noise characterization and reduction in the prediction metric of the macroscale cylindrical detonation model
b. UQ of Eglin experiments
    a. Lessons learned and uncertainty quantification results obtained from investigation of past microscale experiments by an independent investigator, labeled forensic uncertainty quantification, has been submitted for a review in the Journal of VVUQ
    b. Mesoscale gas gun experimental design were significantly improved to meet simulation capabilities in initial testing. However, equipment failure prevented completion of the tests
    c. Successful firing of the macroscale tests were performed. Uncertainty quantification personnel were on-site to ensure gathering of all necessary uncertainty quantification data and documentation of the testing process
c. Convergence of point particle models in Euler-Lagrange Simulations
    a. Faxén form of coupling is utilized to quantify the spatial discretization error.
    b. A non-monotonic behavior in the rate of change in the Downstream Front Position (DFP) for the standard point particle method as the mesh size reduces. DFP does not converge.
    c. A monotonic decrease in the rate of change in the Downstream Front Position (DFP) for the finite-size particle as the mesh size reduce. No evidence for a traditional convergence.
d. Reactive burn model calibration
    a. Prior information of the model parameters is obtained from a variety of sources including an expert's opinion, literature, and direct measurement.
    b. For calibration, three different methods are used: nonlinear least square method, naïve Bayesian method, and Bayesian calibration
    c. All three method give similar estimate on the explosive density and heat release Q, while there is significant difference on the estimate of reactive zone thickness.

5.2 Overview

The primary objective of the UB team is to estimate the error in compressible multiphase turbulence (CMT) models to assess the prediction capability of the hero simulation based on the model. Often model error estimation is suffered by large uncertainty. Thus, uncertainty reduction (UR) was emphasized and UQ was utilized to prepare UR. For systematic UR, the effects of the uncertainty sources on the total uncertainty were categorized and UR priority was made for reducible epistemic uncertainties.

For the Sandia shocktube simulation, about 70% of the initial uncertainty has been reduced. It was found that UQ was very useful for identifying invalidities of the simulation under development and experiment measurement process could be a large uncertainty source.

Validation and UQ of simulation and experiment coupling efforts have been being made based on the UQ and validation framework that was developed for the shock tube simulation. The Eglin macroscale experiments were carried out in 2017, UQ of the experiments are being carried out.

The key models need to be individually improved through subscales, micro-, meso- and macro-scales. Remaining error and uncertainty in the improved models will be propagated through hero simulations to quantify the influences of the models on the hero predictions. The secondary objective is to support other disciplines for quantifying and reducing the uncertainties of their applications and model improvements.

In 2018, calibration on the reactive burn model was performed using Eglin microscale simulation and experiment.

5.3 Validation, Uncertainty Quantification and Uncertainty Budget of Mesoscale Sandia Shock Tube Simulation

The UB team has finished the Sandia campaign of four key campaigns. The error in the particle force model was estimated and the uncertainty in the error were quantified. Figure 5.1 shows comparisons between experimental measurements and one-, two- and three-dimensional shock tube simulation predictions. The three-dimensional simulation prediction captures the measured behavior of the particle curtain. The uncertainty is no longer a hindrance to observe the error since the uncertainty in the both simulation and experiment sides is small. It is believed that the error in the particle force model dominates the prediction error, the good agreement indicates that the error in the particle force model is small. However, one- and two-dimensional predictions have larger errors than the three-dimensional prediction. The error was interpreted as geometrical approximation error due to particle movement restriction. For example, a particle in the one-dimensional simulation can move only one-direction. Two-dimensional simulation has less restriction than the one-dimensional simulation but still it does not allow the degree of freedom as the three-dimensional simulation does.
The on-going study is to quantify the error and uncertainty in the hero simulation prediction due to the particle force model error and uncertainty. Note that the hero simulation embeds the same particle force model with the three-dimensional shock tube simulation. They were propagated through the first CMT-nek based hero simulation “CMT-nek Hero 1” shown in Fig. 5.2(a). Figure

5.2(b) shows a comparison between predictions with (red curve) and without propagation (green band). The green band denotes the 95% confidence interval of the propagated uncertainty due to the uncertainty in the particle force model. The bias of the band from the red curve denotes the influence of the error due to the particle force model estimated from the Sandia campaign. The UB team will apply errors and uncertainties of other CMT models as they were estimated from ASU and Eglin campaigns.

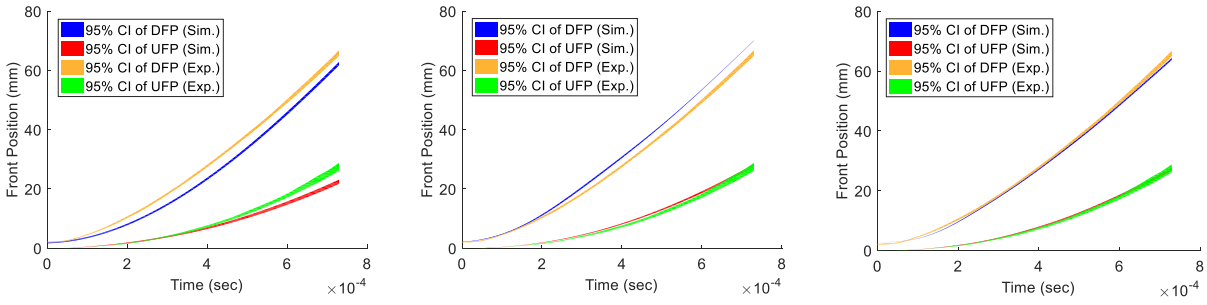

(a) One-dimensional prediction and experimental measurement

(b) Two-dimensional prediction and experimental measurement

(c) Three-dimensional prediction and experimental measurement

Figure 5.1: Similarity between experimental measurement and predictions of one-, two- and three-dimensional simulations

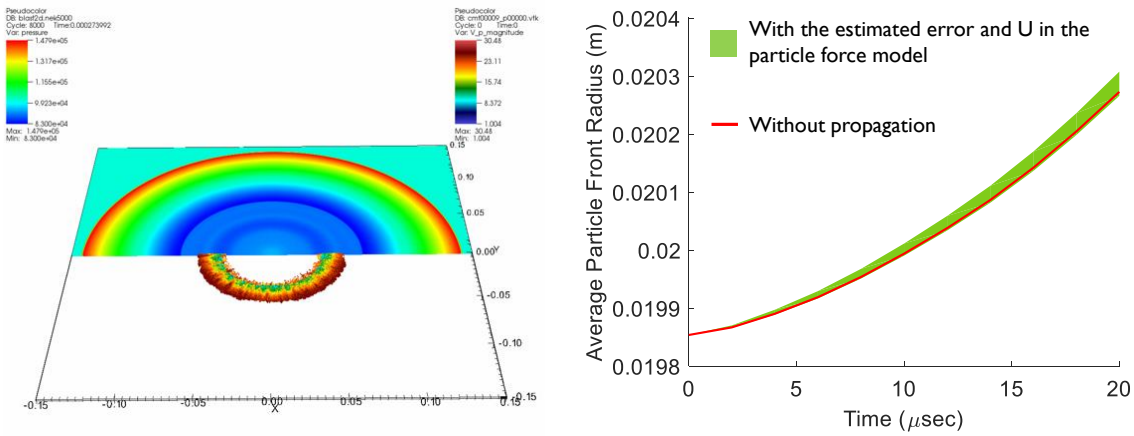

(a) CMT-nek Hero 1 simulation: pressure and particle movement predictions

(b) The error and uncertainty in the average particle front radius prediction due to the error and uncertainty in the particle force model

Figure 5.2: Model error estimate with uncertainty

Since the particle force model is composed of five force component models (quasi-static, added mass, unsteady viscous, pressure gradient and inter particle force), the individual force models should be assessed to pinpoint which force component model is most responsible. The on-going

study is to run the three-dimensional simulation with particle force perturbations for the five force components. This is a what-if type analysis that what if a particle force component is different from the original one. For example, the simulation run was made by multiplying 1.05 on the quasi-static force value calculated using the particle force model. Then one more run was made by perturbing all the forces. The reason of the runs is two folds: 1) to observe the difference between the perturbed and unperturbed results and identify the contribution of the force models on the disagreement, and 2) to see if the all force perturbation can be captured by a linear combination of the individual force perturbations.

Figure. 5.3 shows comparisons between the perturbed simulation runs and the initial unperturbed simulation run in terms of DFP and UFP. For DFP, the most of difference between the all force perturbation and the initial simulation can be explained by the perturbation of the quasi-static force perturbation and the pressure gradient perturbation. For UFP, the influence of quasi-static perturbation is dominant.

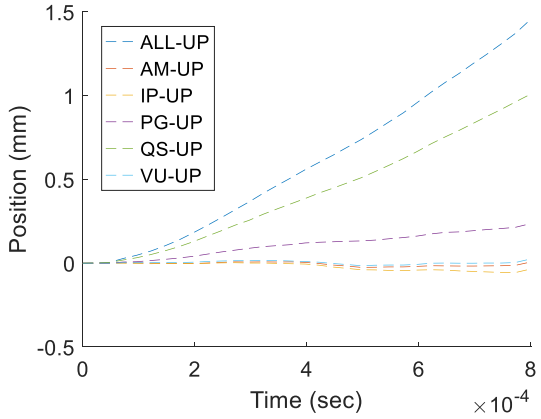

(a) Differences between the perturbed and unperturbed simulation runs in terms of DFP

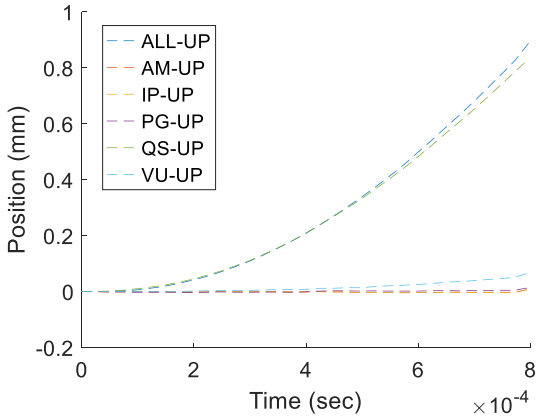

(b) Differences between the perturbed and unperturbed simulation runs in terms of DFP

Figure 5.3: Difference between the perturbed and unperturbed simulation runs (The legends of the graphs indicate which curve is the difference between the prediction with perturbing which force components and the unperturbed simulation. ALL: prediction with all force perturbation, AM: prediction with added mass force perturbation, IP: prediction with inter-particle force perturbation, PG: prediction with pressure gradient force perturbation, QS: prediction with quasi-static force perturbation, VU: prediction with viscous unsteady force perturbation, and UP: unperturbed simulation prediction)

The observation indicates that the error in the pressure gradient force model can explain the disagreement shown in Fig. 5.1(c). This is because perturbing pressure gradient force made the speed of DFP faster while it did not make change for UFP. With the fact that the three-dimensional simulation slightly underestimated the speed of DFP but it does not the speed of UFP, such disagreement is likely to be reduced by improving the pressure gradient model. We are discussing with the physics team on this matter.

5.4 Validation, Uncertainty Quantification and Uncertainty Budget of Eglin experiments

Kyle Hughes is working on UQ study of Eglin microscale experiments. The Eglin experimental personnel, led by Chi Mai, has involved the uncertainty budget teams closely in the design of the experiments based on lessons learned from forensic uncertainty quantification of previous experiments. In the designing of new experiments there are several driving questions:

1. What were the major sources of uncertainty from previous experiments? Can we mitigate these effects in future experiments?
2. What were the important diagnostics in the past experiments? What has provided the most benefit in terms of validation metrics?
3. What configurations will provide the most benefit for comparison to simulations?

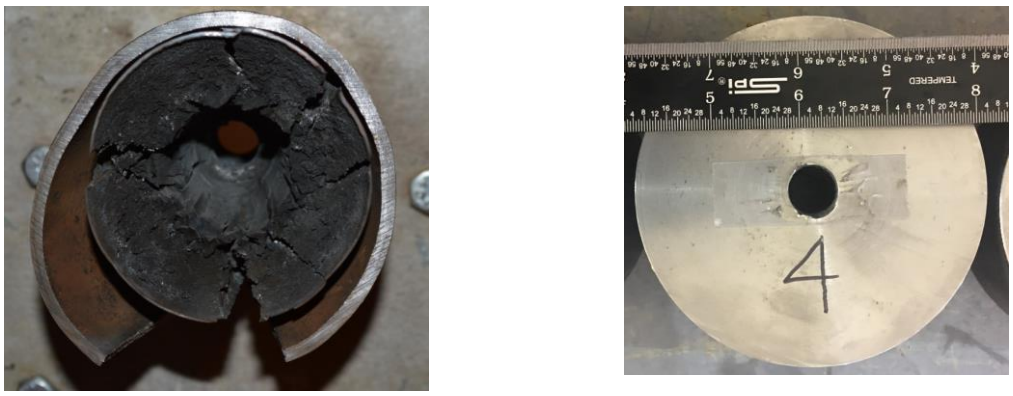

a) b)

Figure 5.4: Post-shot pictures of the casings used in previous explosive experiments. a) Casing is largely destroyed in previous microscale experiments. b) Casing from proton radiography experiments show relatively mild deformation.

While the full details of the new experiments are not presented here, a few highlights are presented. First, one of the largest sources of uncertainty of the previous experiments is the large casing fragmentation and deformation present in previous experiments. The new casings will be doubled in size, similar to the Eglin mesoscale experiments, to prevent the casing from significantly fracturing (see Figure 5.4). The prevention of fragmentation will allow easier processing of x-ray imaging, one of the most important diagnostics. Secondly, knowledge of the gas behavior was only available from the high-speed contact line images in previous experiments. The pressure probes were too far away to easily include in simulations. The pressure probes will be moved significantly closer to allow their inclusion in the simulation domain for the next set of experiments. Third, the modeling of the explosive is still a significant challenge. Several of the shots planned for the future experiments will be "bare charge," with no particles in the test, to allow validation of the explosive modeling. Multiple shots of the single particle test case will be repeated to build further evidence for the drag model on a lone particle.

Uncertainty quantification for Eglin microscale experiments. A promising new approach to performing uncertainty quantification (UQ) is being developed for past experiments that involves the adoption of the perspective of a crime scene investigator. Like a crime scene investigator, the UQ personnel collects and documents "evidence" about the experiment to reconstruct an accurate picture of the experiment. In addition, after-the-fact characterization of the measured inputs may be performed in the laboratory to confirm details, similar to a forensics laboratory. In addition, it is postulated that there is a distinct advantage in have a third party beside the simulationist and the experimentalist perform the investigation. The relative independence of the investigator allows for a critical evaluation of discrepancies between experiments and simulations. A journal article is currently being drafted, to be submitted to AIAA Journal, exploring the concept of forensic UQ with application to the microscale experiments.

Uncertainty quantification for Eglin mesoscale experiments.

Kyle Hughes (UF) assisted Mike Jenkins (AFRL) as they attempted to complete the gas gun experiments. However, the mesoscale experiments were unable to be completed this summer due to equipment issues. The laser synchronizer was faulty and had to be sent to the manufacturer for repairs. The tests will be resumed once the PIV setup is returned to operational status and time is available in the testing chamber.

One of the challenges for the current setup is the fragmentation of the sabot. Large pieces of sabot can be seen following immediately after the particles and is a significant source of uncertainty. The velocity of the sabot is being reduced to eliminate the fragmentation while still maintaining large particle velocities.

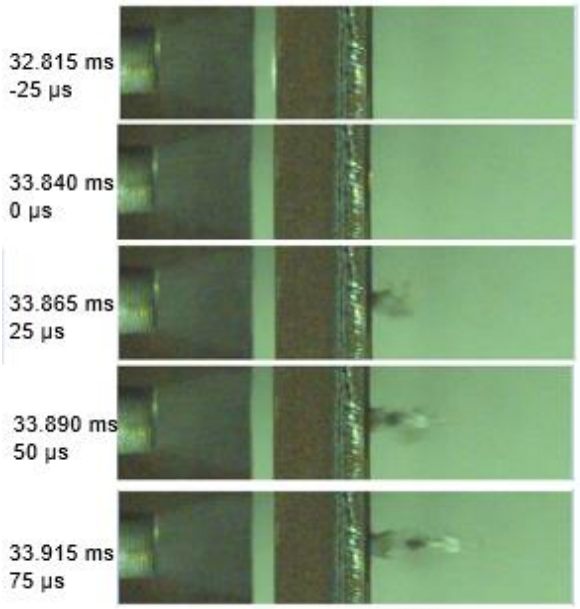

Figure 5.5: Test firing of a sabot carrying 1.0 grams of tungsten powder. The sabot impacts the stripper plate at the 0 µs mark. The dark cloud is the ejection of the tungsten powder and then the white sabot fragments can be seen quickly overtaking the powder.

Uncertainty quantification for Eglin macroscale experiments

Particles were reported by the manufacturer as sieved between 75-125 μm. The particle size distribution of the steel particles was quantified from a series of 25 scanning electron microscope (SEM) images. Images are segmented with a global threshold. The equivalent diameter, d_e, of the particles is obtained from the following relation:

$$d_e = C\sqrt{\frac{4A}{\pi}}$$

Where A is the number of enclosed pixels for each particle and C is the calibration constant obtained from the SEM image. Figure 5.6 contains an example image and the accompanying segmentation results (with the boundaries of the particles shown in red). Two distributions of particles were apparent during the post-processing. The small particle sizes (1-20 μm) and the large particle sizes (60-240 μm). The small particle sizes appear to be an artifact of the manufacturing process. The large particles have many small nodules on their surface. It is postulated that subsequent handling after sieving constantly breaks the small nodules from the surface of the large particles and gives rise to their presence in the sample. Particle size results are presented as histograms in Figure 5.7 and summary statistics in Table 5.1. As shown in Figure 5.7, both small and large particle sizes appear to differ significantly from normality. The Kolmogorov-Smirnov test performed with 95% confidence rejected the null hypothesis of normality for both small and large particles. Fitting of the distributions with 16 different distributions showed the log normal distribution with a minimum negative log likelihood.

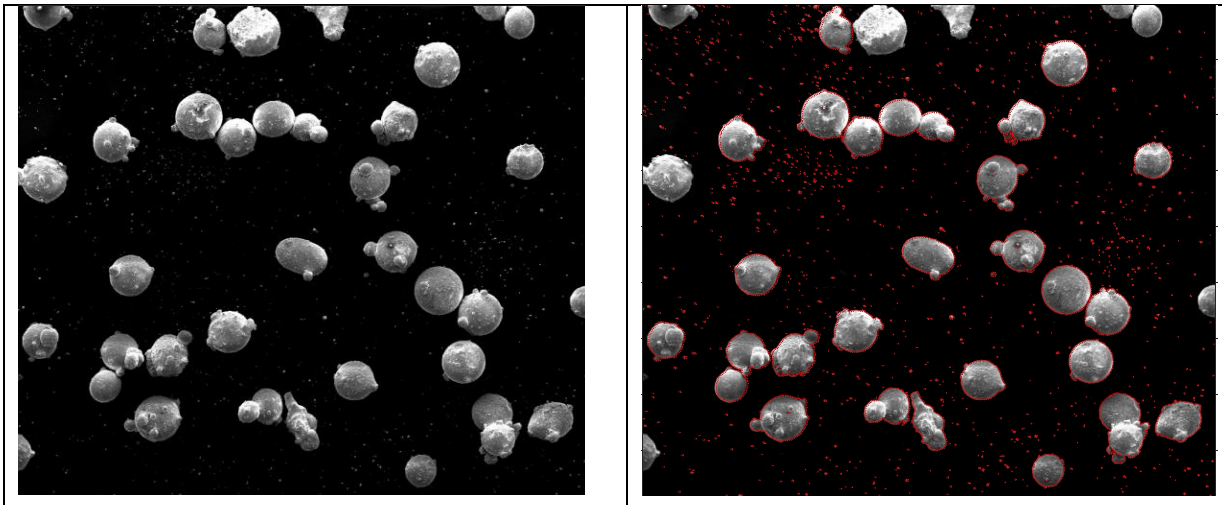

a) **b)**

Figure 5. 6: Sample of SEM images used to obtain particle the size distribution. a) SEM image showing both large and small particles present. b) Red traces indicate the selected boundaries of the particles. Particles partially on the boundary are rejected.

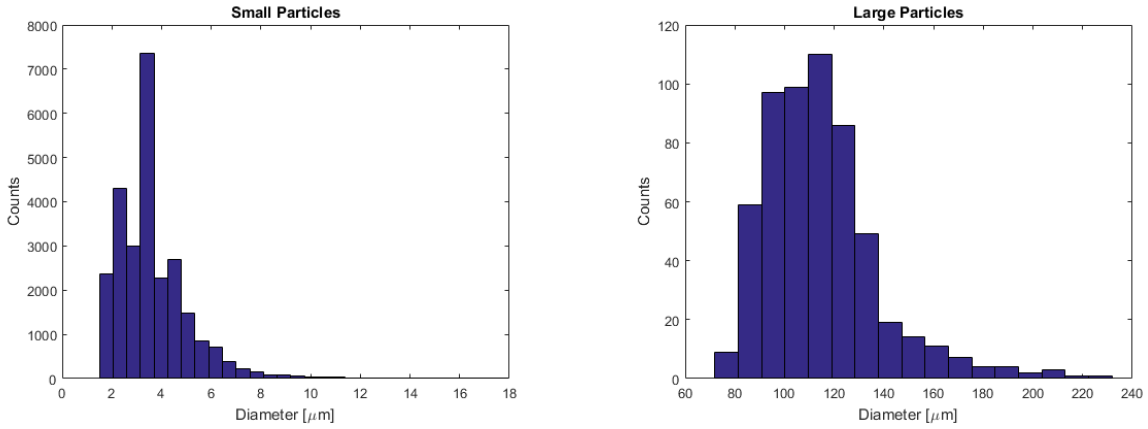

a) **b)**

Figure 5.7: Particle size distributions. a) Small particle sizes. b) Large particle sizes.

Table 5.1: Summary statistics of the two particle size distributions. (The hypothesis of following a normal distribution is rejected for both small and large particles)

	n	μ [μm]	σ [μm]	Kolmogorov-Smirnov (95%)	*p*-value
Small Particles	26228	3.6	1.5	Reject	0
Large Particles	575	115	23	Reject	0.00063

5.5 Convergence study in 1D shock-particle interaction simulation

Investigation of Convergence of Point Particle Models in Euler-Lagrange Simulations of Shock-Particle Interaction

The purpose of this study is to investigate the convergence of Euler-Lagrange (E-L) simulation under the mesh refinement. The finding of this study is meant to use for the validation of multiphase shock tube simulation.

Point particle methods are extensively used in simulating E-L multiphase dispersed flow. However, numerical convergence and accuracy of these methods under mesh refinement is still an open question. The standard approach of approximating the fluid-particle coupling at the particle center fails to converge as the Eulerian grid is reduced below particle size. For the example of shock-particle interaction, the particle does not feel any force due to the presence the shock unless the shock reaches very close to the center of the particle. This causes a very sharp change in the particle force over a very short period when the shock crosses over the particle. The change gets

sharper as we increase the mesh size resolution. Since we deal with a two-way or four-way coupled simulation, the particle momentum transfer to the Eulerian fluid is significant where the fluid disturbance by the particle is in order of the particle size. Hence, projecting the particle momentum onto the few Eulerian girds where they are located very close to the center of the particle causes the fluid to experience similar to the particle; a very sharp change in momentum at the very early stage when the shock passed over the particle that it depends on the mesh resolution. We need an approach that takes the finite size of the particle into account. Here we use the generalized Faxén form to compute the force on a particle and examine the limit of the finite-size particle to compare the results against standard point particle method. This is done by apportioning the different force components on the particle to fluid cells based on the fraction of particle volume in the cell (e.g., pressure gradient and added mass force) or the fraction of particle surface area in the cell (e.g., quasi-steady force). The Faxèn form also allows for the apportioning of the coupling force back to the fluid in a spatially distributed manner. The contribution to the surface or volume average from each grid cell can be used to apply the force back to the fluid in all the grid cells covered by the particle. The convergence of this approach will be examined and compared to the usual approach of feeding the force back to a few grid cells that surround the particle center.

As we illustrated in previous reports, the convergence is achieved for a well-formulated force model and back coupling for the finite-size particle for the case study of a single stationary and moving particle. We also showed on the previous reports that in the one-way coupled context, even though the time history of standard point particle force experienced by a particle subjected to a planar shock continues to change with increasing mesh resolution, the overall impulse on the particle reasonably converges. As a result, the long-term motion of the particle after the passage of the shock can be well predicted even though the early time motion is grid dependent. However, this will not be true in the case of four-way coupling as the shock propagates through a cloud of particles. The very strong force at early time, whose magnitude depends on the mesh resolution, can lead to inter-particle collisions, and thereby affect long term evolution. Thus, four-way coupling has the potential to adversely influence long term convergence as well.

Now we illustrate and compare the particle motion under the mesh refinement for both methods for the case of many moving particles with moderate volume fraction. We chose the downstream particle front position (DFP) or the average position of the most extreme particles at downstream as the quantity of interest. In the present work, we refine the mesh while keeping the number of the computational particles per cell constant. The particle diameter chose to be 115 μm for all cases with L = 2 mm as an initial particle cloud thickness. The mesh size and the number of computational particles varying from dx = 400 μm and N = 500 particles to dx = 6.25 μm and N = 32,000 particles. The simulation performed under the moderate regime of 21% initial particle volume fraction (PVF).

Figure 5.8 shows the DFP for the standard point particle and finite size particle methods. The DFP is defined by the average location of the most extreme right 5 to 15 percent of the particles at the downstream particle front where. It is computed for the time frame of t = 350 μsec after the shock hits the first particle. The system is chaotic, too much averaging leads to degradation of the information available to describe the system behavior, while too little averaging leads to statistics that are noisy and difficult to interpret. The horizontal axis in figure 5.8(a) and (b) indicate the mesh size normalized by the particle diameter and the vertical axis show DFP normalized by the

initial particle cloud thickness. Our justification here is based on the left three points of figure 5.8(a) and (b); the smallest mesh sizes that we examined. According to figure 5.8(a) and 5.8(b), the variation in DFP is reduced for both methods as the particle statistics increases. Based on figure 5.8(a) the rate of change of DFP for the standard point particle method is not monotonically decreases while figure 8(b) shows a monotonic decrease in the rate of change in DFP for the finite size particle method as the mesh size reduces. However, there is not a clear sign of the convergence in neither of the two methods. Another observation we obtained by comparing figure 5.8(a) vs. (b) is that DFP possesses a different value for each method.

Even though the finite size particle method converges for a single particle, we cannot conclude a convergence for the case of four-way coupled simulation. Taking the finite size of the particle into account by itself does not guarantee the convergence of the four-way coupled E-L simulation. Other physics and numerical challenges are also tied to changing the mesh size. So, further investigation is needed.

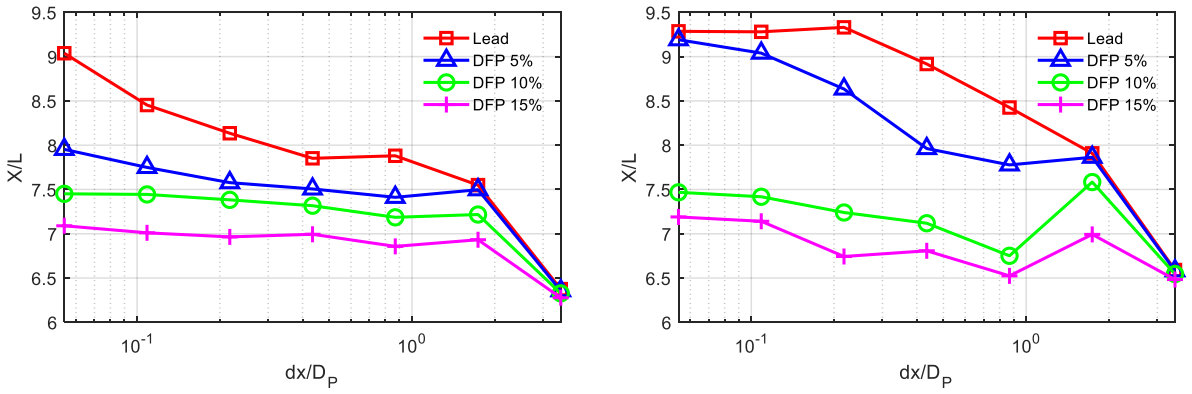

(a) Standard point particle

(b) Finite size particle

Figure 5.8: Mesh refinement for standard point particle vs. finite size particle

5.6 Reactive burn model parameter calibration in microscale

Chanyoung Park has finished calibration of reactive burn model parameters. There are three main parameters in the model: explosive density, reactive zone (RZ) thickness, and heat release. The quantity of interest was transverse shock location. Figure 5.9 shows the location of shock.

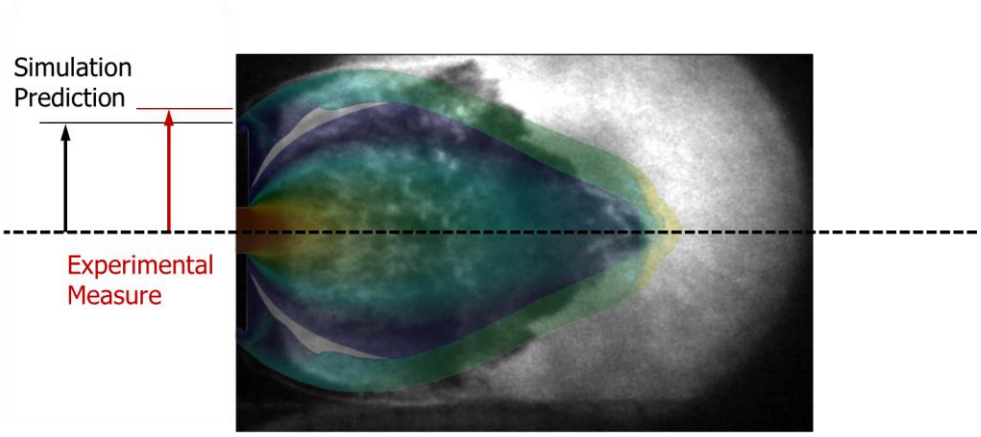

Figure 5.9: Measurement of transverse shock location

Prior information of the three parameters are determined from pycnometer measurement, literature, an expert's opinion, or combination of more than two sources. Table 5.2 summarizes the uncertain parameters used in the simulation.

Table 5.2: Summary of uncertain parameters

Parameter	Quantity	Method
U1: explosive density	Mean = 1795 kg/m^3 std = 2.9 kg/m^3 Bi-modal	Derived
U2: RZ thickness	[0.365, 1.4]mm	Expert's opinion
U3: Heat release Q	[11.6, 14.2] MJ/kg	Literature and opinion
U4: Particle diameter	Mean 2.0156 mm std = 0.0073mm Weibull	Direct measurement (Micrometer) 52 samples
U5: Particle density	Mean = 15540 kg/m^3 std = 250 kg/m^3 Normal	Direct measurement (Pycnometer) 12 samples
U6: Initial radial position	[0, 0.254] mm	
U7 – U11: JWL	A, B, w, R_1, R_2	Literature and opinion

Figure 5.10 shows the uncertainty propagation of the reactive burn model parameters. The red line is the prediction mean while the grey area shows the uncertainty of prediction. Measurement from experiment is plotted on top of the simulation result along with its error bar. Figure 5.11 is the result of the global sensitivity analysis. Heat release Q is the dominant source of uncertainty.

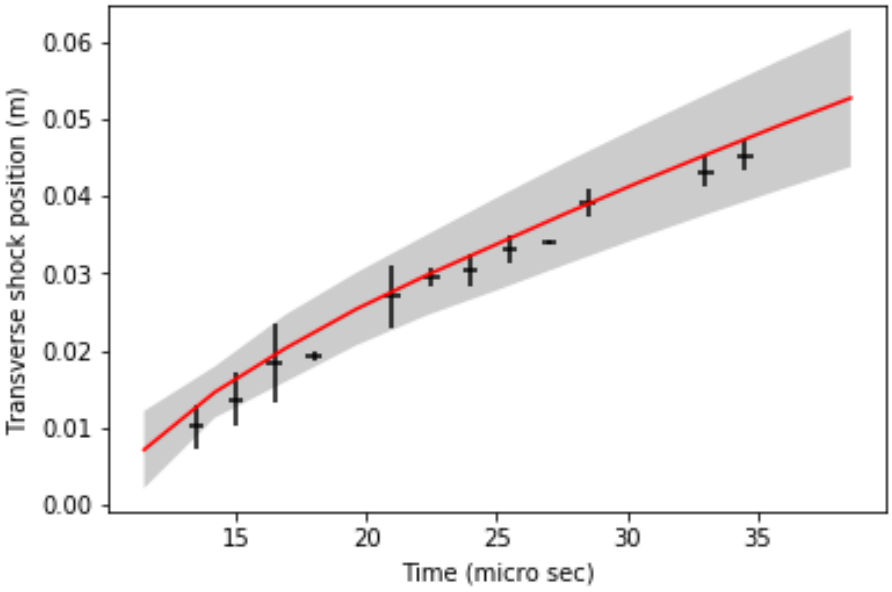

Figure 5.10: Uncertainty propagation on transverse shock position

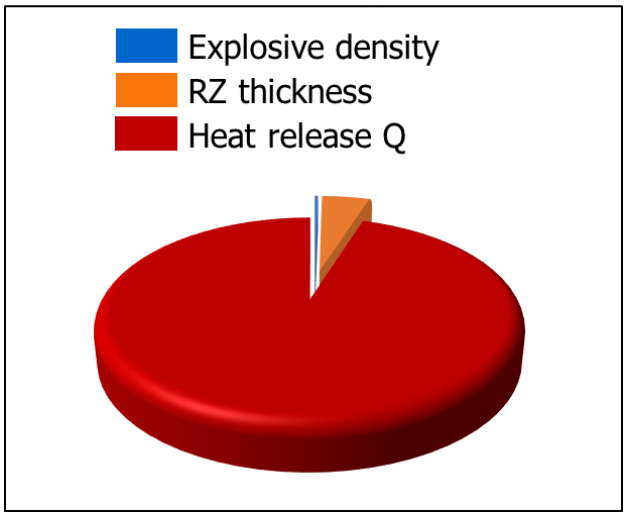

Figure 5.11: Global sensitivity analysis of reactive burn model parameters

On the calibration of the model parameters, three different methods are used: nonlinear least square method, naïve Bayesian, and Bayesian calibration. The difference between naïve Bayesian and Bayesian calibration method is the existence of discrepancy function. All three methods give similar estimate for explosive density and heat release Q, while there is significant difference in RZ thickness estimation from Bayesian calibration. Table 5.3 summarizes the result.

Table 5.3: Summary of parameter calibration

	LSQ	Naïve Bayesian	Bayesian
Exp. Density	1787 kg/m^3	1798 kg/m^3	1796 kg/m^3
RZ thickness	1.7 mm	1.7 mm	0.88 mm
Q heat	12.9 MJ/kg	12.9 MJ/kg	12.5 MJ/kg

Based on the calibration result, the transverse shock location and the particle position are compared. Figure 5.12 displays the result.

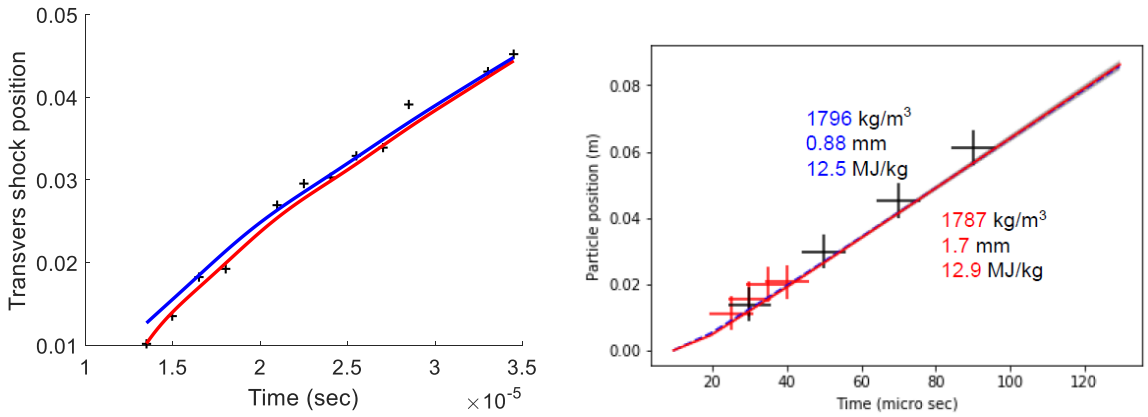

(a) Transverse shock location (b) Particle position (m)

Figure 5.12: Result of reactive burn model calibration

6. CMT-nek Code Development Team

6.1 Overview

CCMT completed its transition to CMT-nek for mesoscale and macroscale flow regimes in 2018, and the challenges of applying CMT-nek to a sequence of hero runs representing the demonstration problem brought about a pivot in shock capturing strategy and a migration towards the state of the art in stable nodal discontinuous Galerkin methods. The major advancements of 2018 related to CMT-nek were the acceptance of the final microscale simulations in a *Physical Review Fluids* paper (Marjanovic *et al*), the robustness added by positivity-preserving limiters to CMT-nek, as well as our first runs with non-ideal state equations (as introduced above by the macroscale team).

One of the challenges in performing simulations was load-balancing the particles. When the year began, particles were stored in memory with their surrounding grid points. This poses an issue when particles only occupy a small portion of the domain since only a select few processors are allocated to solving the particle phase. One key enabling technology for large-scale multiphase simulations with CMT-nek has been the development of a highly parallel particle algorithm. This algorithm, called the binned ghost particle (BGP) algorithm, has allowed for simulation of larger problem sizes than previous approaches even under extreme load-imbalances. By design, the BGP algorithm scales independently of the underlying imbalances. The BGP algorithm has also been shown to scale up to 100,000 MPI ranks totaling over 300 million particles. Furthermore, the scaling trends suggest that this trend will continue past even 1 million MPI ranks.

Hero runs at conditions representing an explosion of Composition B into the atmosphere showed CMT-nek's ability to tackle the demonstration problem in production simulations, and statistical analysis continues apace to extract new insights into blast wave propagation through compacted particle beds. We also look forward to exciting new research directions in artificial viscosity and the entropy-stable flux functions in the discontinuous Galerkin spectral element method (DGSEM).

6.2 Non-ideal equations of state and the 1-equation JWL model

The lumped parameters used to represent the detonation products as initial conditions to the demonstration problem are at thermodynamic conditions that are not well-represented by an ideal gas. The equation of state of Jones, Wilkins & Lee (1968) (JWL) has been used to represent detonation products in early simulations of the demonstration problem and other experimental configurations in CCMT. This EOS as well as the 1-equation mixture model used by the Macroscale Team to "blend" JWL with calorically perfect air at atmospheric conditions were both implemented in CMT-nek and validated against RocFlu.

However, the 1-equation mixture model does not have an algebraic expression for entropy as a function of other state variables. Therefore, evaluating the entropy residual by numerically integrating the 1-equation mixture model's expression for entropy was deemed prohibitively expensive for the entropy viscosity method (EVM). We recognized the need for improvements to our shock capturing methodology by changing our artificial viscosity method and augmenting it with safeguards against unphysical solution values. The major results and all non-ideal-gas hero runs were first-order in 2018, applying the viscosity and mass diffusivity computed from the

element-wise maximum wave speed everywhere in the domain. To reduce its values lower than were needed to calibrate EVM in ideal gas simulations, a limiting procedure now automatically rescales the solution in each element to avoid negative density and energy.

6.3 Solution limiting and shock capturing

Rescaling the solution variables within each element to correct local violations of physical solution bounds is called "solution limiting" in the literature and is a widespread practice in the discontinuous Galerkin community. CMT-nek now uses limiters to preserve the positivity of density (Zhang & Shu (2010) *J. Comp. Phys.* **229**:8918-8934) and internal energy at each Runge-Kutta stage. The internal energy limiter applies Jensen's inequality between each GLL node and the internal energy computed from the averages of the conserved variables in a given element. This strategy was adapted from the entropy-bounding solution limiter of Lv & Ihme (2015) *J. Comp. Phys.* **295**:715-739, whose first author, Prof. Yu Lv of Mississippi State University, briefed CCMT and assisted the CMT-nek development team in a visit invited in June.

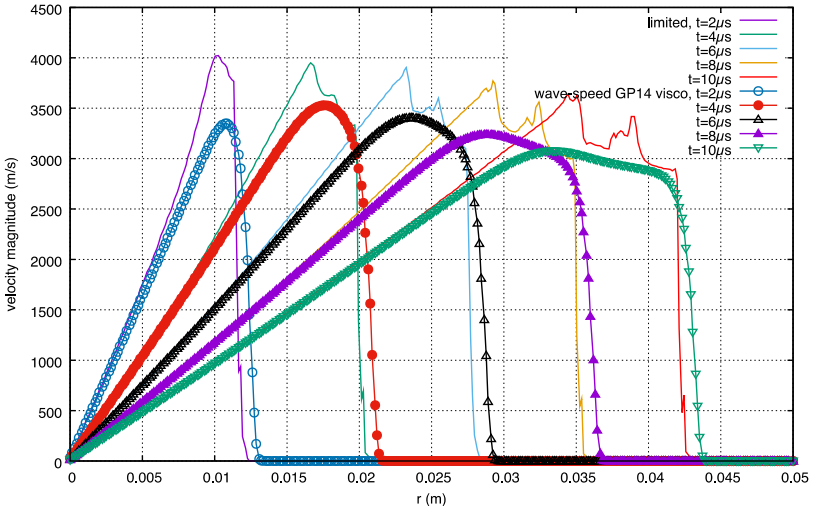

Figure 6.1. CMT-nek velocity magnitude profiles from a cylindrical 2D blast wave into the atmosphere initialized to the lumped static conditions of a 3.8mm charge of Composition B. A simulation using wave-speed-based viscosity in the Guermond-Popov (2014) stress tensor (symbols) is compared to a simulation using positivity-preserving limiters in density and energy (solid curves) with no explicit artificial viscosity at all. Polynomial order is 3, and mesh spacing is 0.11875mm.

Results in Figure 6.1 show an example blast wave from a 3.8mm-diameter charge of Composition B using the 1-equation JWL mixture model EOS. Both first-order wave speed viscosity and a solution computed using the solution limiter with no viscosity recover the same wave speeds. The limiter preserves a sharper shock, but it does not guard against overshoots. Work will continue in 2019 on deriving a new artificial viscosity method able to localize artificial viscosity more robustly than EVM while suppressing overshoots that the limiter allows.

6.4 Hero runs

CCMT's evolution towards production use of CMT-nek proceeded through three "hero runs" in 2018. This is discussed in the Macroscale group section.

Hero 1 began at a low pressure ratio and dilute volume fraction, which was then increased to a dense, precomputed packing for Hero 2. Hero 3 changed both the equation of state (EOS) and increased the pressure ratio to 140,000, representing Composition B detonation products in the 1-equation JWL mixture model.

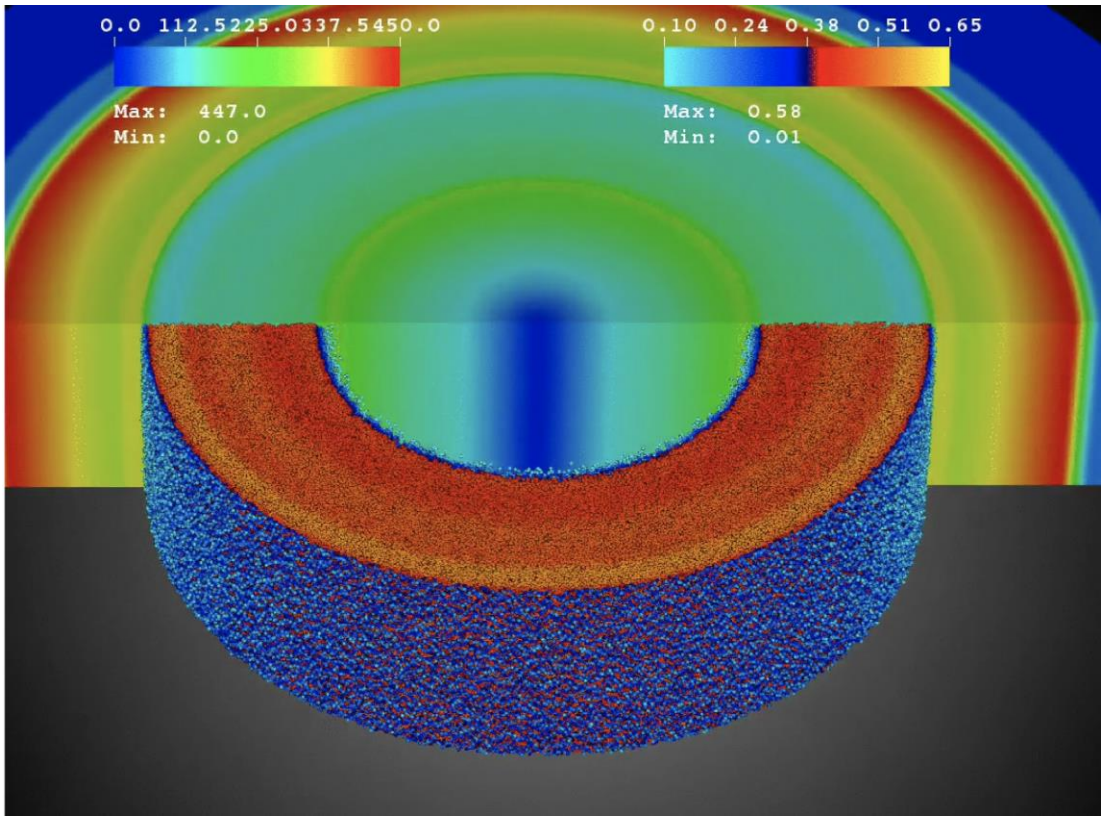

Figure 6.2. Visualization of CMT-Hero 2 simulation. The particles are shown in the bottom half colored by their local volume fraction. In the top half, the fluid velocity is shown.

	CMT-Hero 1 ✓	CMT-Hero 2 ✓	CMT-Hero 3 ✓	CMT-Hero 4
Case	Frost	Frost	Blastpad	Blastpad
Simulation Time	0.3 ms	0.3 ms	2.5 ms	2.5 ms
Degrees of Freedom	32.4 M	32.4 M	120 M	240 M
MPI ranks	32 768	32 768	65 536	131 072
No. computational particles	1 M	2 M	5 M	10 M
Initial Particle volume fraction	5 %	60 %	60 %	60 %
Gas-Particle Coupling	2-way	4-way	4-way	4-way
Reactive burn	No	No	No	Yes
EOS	Ideal Gas	Ideal Gas	JWL	JWL

6.5 ASU simulations

Another simulation that has been performed is the mesoscale simulation of an expansion fan propagating into a bed of particles at conditions corresponding to the ASU multiphase shock tube

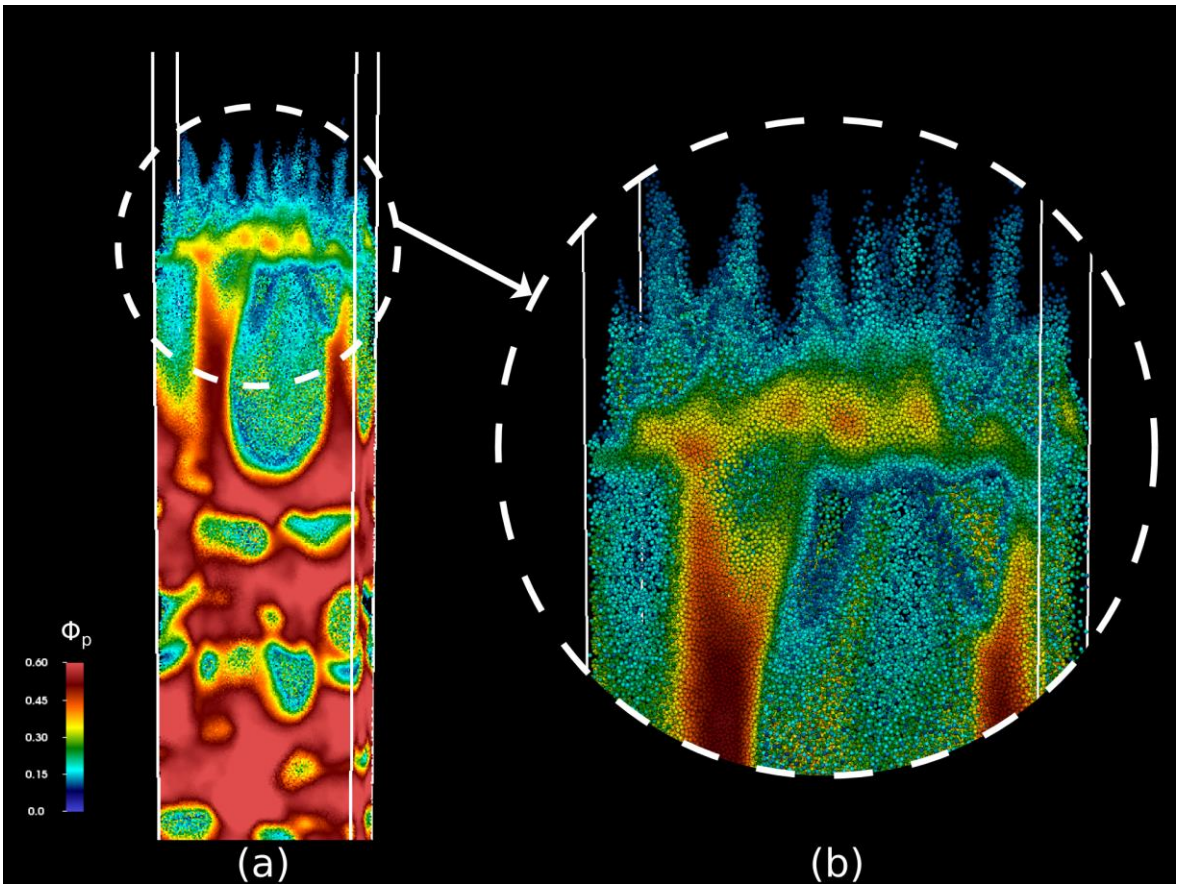

Figure 6.3. Simulation of ASU experiment. The picture depicts the particle volume fraction of each individual particle. In (a) the front of the expanding particles is shown. In (b) a zoomed view showing each individual particle is given.

experiment. A single instant in time is shown in Figure 6.3, with particles colored by their local volume fraction. In frame (a), a small portion of an expanding particle bed is shown at 20 ms after the diaphragm of the shock tube has burst. Void regions as well as jetting at the particle front are apparent. In frame (b), the individual spherical particles at the uppermost portion of the expanding bed are shown. The data from the simulations are currently being analyzed.

6.6 Future work in 2019

CCMT is changing CMT-nek to follow a novel and transformative approach to DGSEM that has evolved rapidly in the last five years. Research at Sandia National Laboratory and NASA (Carpenter *et al* (2014) *SIAM J. Sci. Comput.***36**(5):B835-B867) has demonstrated a novel way of exploiting the summation-by-parts (SBP) property of the finite difference matrix used in spectral elements to derive energy-preserving and even entropy-conserving "two-point-forms" of the

volume integrals of DGSEM. These forms are easily modified to err on the side of dissipation, and the resulting guarantee that global entropy will at worst increase is termed "entropy stability." The substantial rewrite required to get CMT-nek to follow the two-point approach has already shown more stable and accurate results in preserving an under-resolved homentropic vortex than traditional dealiasing using energy-preserving flux functions (Figure 6.4. See also Gassner *et al* (2016) *J. Comp. Phys* **327**:39-66). Work continues on rewriting the artificial viscous fluxes to be consistent with this form and incorporating entropy-stable fluxes for ideal gases (Ismail & Roe (2009) *J. Comp. Phys.* **228**:5410-5436). An obvious research direction open to the center involves deriving entropy-stable fluxes for non-ideal equations of state. Such fluxes will empower DGSEM to extend the benefits and multiscale resolution of high-order methods to energetic materials and challenging compressible flows. Nevertheless, simpler energy-preserving forms should retain this property and economically stabilize DGSEM for any equation of state.

Replacement of EVM with the Navier-Stokes regularization popularized by Cook & Cabot (2005) *J. Comp. Phys.* **203**:379-385 is another simplifying step that will remove the temporary restriction

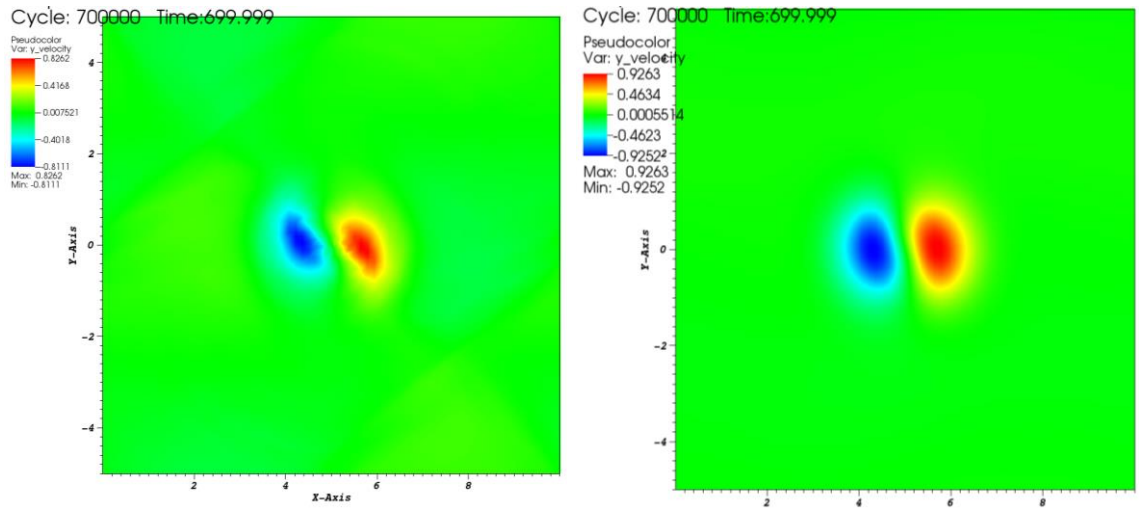

Figure 6.4. Pseudocolor plot of *y*-velocity in the advection of a steady compressible homentropic vortex through a periodic domain 70 times in a CMT-nek simulation at mesh spacing $h=10/4$ and polynomial order of 24. Traditional dealiasing (left) is seen to suffer more corruption than the two-point form (right) using the energy-conserving fluxes of Kennedy & Gruber (2008) *J. Comp. Phys.* **227**:1676-1700 in the flux differencing form derived by Gassner *et al* (2016) *J. Comp. Phys* **327**:39-66. Fully entropy-stable fluxes (currently in testing) should force the vortex to decay at very long time without any significant asymmetry.

to overly diffused production runs. Limiters will compliment any artificial viscosity scheme. Future hero runs will also need improvement in boundary conditions, and time-evolving "reactive-burn" initial conditions that will also receive development support in 2019.

7. CS Team

7.1 Overview

The research effort of our CS team has been towards optimizing CMT-nek in terms of performance, power and energy consumption. Towards that end, we completed the experiments on dynamic load balancing and submitted our load balancing paper. We have started implementing a GPU version of CMT-nek. We present below some results we obtained during these implementations.

7.2 Load balancing CMT-nek

7.2.1 Experimental Results

Problem description

To test the load-balancing algorithm, we used a test case that has been devised to mimic some of the key features of particle-laden, explosively driven flows that the load-balancing algorithm proposes to overcome. The test deals with expansion fans in one dimension which are simple compressible flows. The problem domain is a rectangular prism that extends from 0 to 0.0802 in the y and z directions and from −2.208 to 6.0 in the x direction. Note that the units in this case are non-dimensional. The particles are assigned between −1.0 and −0.5 in x direction, where the difference between the left (x = −1.0) and right (x = −0.5) boundaries determines the initial volume fraction of particles. The left boundary is often adjusted to obtain a different initial volume fraction.

Experiments on Quartz

Figure 7.1 shows the overhead of three load balancing strategies, namely, centralized, distributed and hybrid load-balancing algorithms on Quartz. It is a weak scaling with 4 elements per MPI rank and about 343 particles on an average per element. The variable *lelt* was set to 16. Each spectral element consists of 5 × 5 × 5 grid points. The overhead includes time taken for each of the following steps: 1) remapping elements to processors; 2) packing, sending, and unpacking received elements and particles; and 3) reinitialization of data structures that are used in computation. The horizontal axis represents the number of MPI ranks while the vertical axis represents the time in seconds taken to load balance the application.

The overhead incurred by a load-balancing step increases with the number of MPI ranks. Ideally, the distributed algorithm should take less time than the centralized algorithm with increasing MPI ranks since there is no processor P0 bottleneck in it. However, on Quartz the centralized algorithm is faster due to a higher ratio of communication-time to computation-time on the system and the distributed algorithm is rich in communication especially in MPI_ALLGATHERV. The hybrid algorithm, eliminates calls to MPI_ALLGATHERV, as well as, the part in the centralized algorithm where all processors send their element loads to P0. As we can see from Figure 7.1, the hybrid algorithm was the fastest. The actual overhead for 65, 520 MPI ranks for centralized, distributed and hybrid was 0.33, 0.57 and 0.31 seconds, respectively. Compared to the time per time step which was 0.17 seconds, the overhead expressed as number of time steps was 1.94 for the centralized, 3.35 for the distributed, and 1.82 for the hybrid algorithm. This makes dynamic load balancing practical for a large class of simulations. For these experiments, the total number

of time steps was 100, and load-balancing took place every 10 steps. Thus, we found that the overhead for load balancing is low and scales very well with the number of processors.

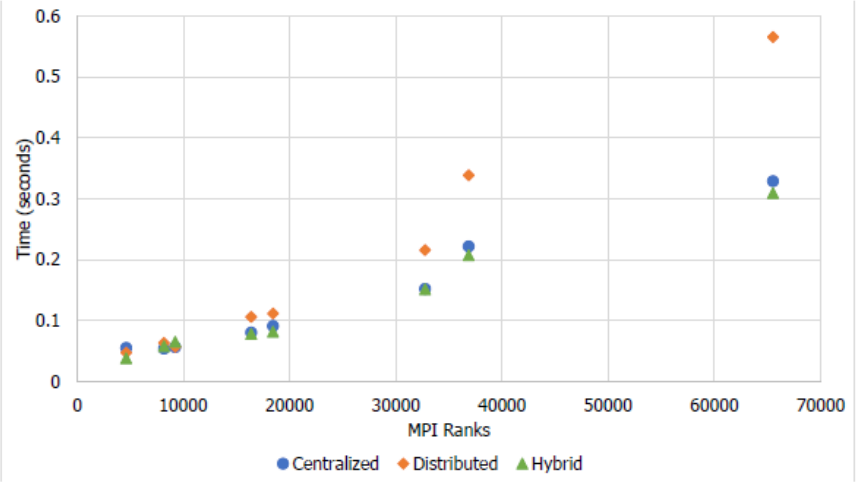

Figure 7.1 On Quartz, total overhead for a load balancing step for centralized, distributed and hybrid algorithms. It is a weak scaling with 4 elements per MPI rank, 5 × 5 × 5 grid points per element, and about 343 particles per element. The actual overhead expressed as number of time steps for 65, 520 MPI ranks was 1.94 for the centralized, 3.35 for the distributed, and 1.82 for the hybrid algorithm.

The load-balanced and non-load-balanced (original) codes were run on 67, 206 MPI ranks on Quartz, that is 1, 867 nodes with 36 cores per node. The grid size per element was 5× 5× 5 and the total number of elements was 900, 000. The variable *lelt* was set to 120 elements. The total number of particles was 1.125×10^9, obtained as 1250 particles per element on an average. Initially, the percent of elements that have particles is 6.1% of the total number of elements.

Figure 7.2 compares a trace of the CPU time taken per simulation time step for load-balanced versus the original. Adaptive hybrid load balancing was used in this example. The average time per time step for the original and the load balanced versions were 9.92 and 0.995 seconds, respectively. Thus, we gained an overall speedup of 9.97 using load balancing algorithm. During the duration of simulation, apart from the compulsory load balancing that happens before simulation time step 1, CMT-nek begins load balanced at 4, 077 time steps and after that, it will automatically load balance by itself (the small blue dot there represents the time taken by next step after load balance). Original version did not finish in 2.2 hours.

Experiments on Vulcan

We now evaluate the load-balancing algorithms on Vulcan. Figure 7.3 shows the total overhead for a load-balancing step using the centralized, distributed and hybrid load-balancing algorithms.

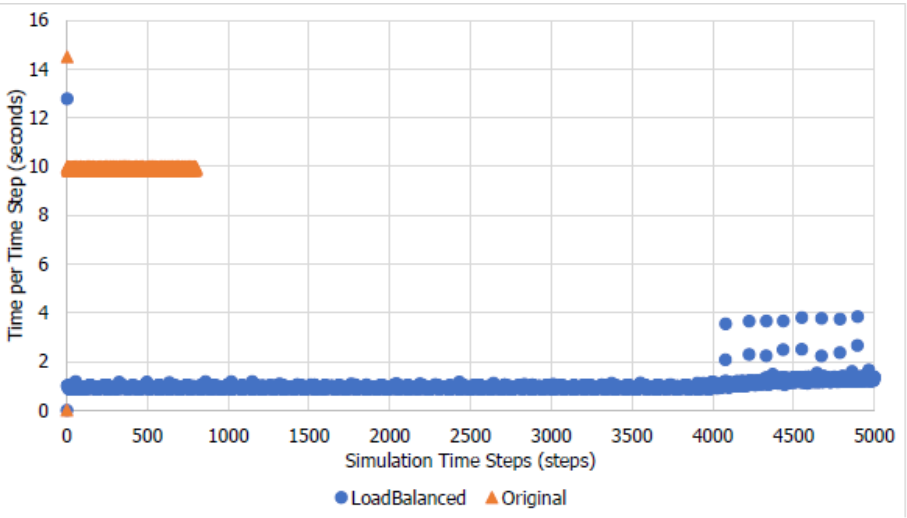

Figure 7.2 Performance comparison between load-balanced and original versions of CMT-nek on Quartz. They were run on 67, 206 MPI ranks, that is 1, 867 nodes with 36 cores per node. Adaptive hybrid load balancing was used. The average time per time step taken by the original version and the load balanced version were 9.92 and 0.995 seconds, respectively, giving us an overall speed-up factor of 9.97. Original version did not finish in 2.2 hours.

It is a weak scaling study so problem size increases proportionally to the number of MPI rank that is 2 elements per MPI rank and 343 particles per element on average. As we can see from Figure 7.3, the load-balancing overhead increases with an increasing number of total MPI ranks.

Especially, distributed algorithm was faster than the centralized and hybrid algorithms. That is because of a lower ratio of communication-time to computation-time on this platform. The actual overhead for 393, 216 MPI ranks for centralized, distributed and hybrid algorithm was 1.00, 0.77 and 0.84 seconds respectively. Compared to the time per time step which was 0.33 seconds, the overhead expressed as the number of time steps was 3.03 time steps for the centralized, 2.33 for the distributed, and 2.55 for the hybrid algorithm. The variable *lelt*, which is the maximum number of elements on an MPI rank, was set to 8 for these overhead runs. Of these experiments, the total number of time step was 100 and load balancing took place every 10 steps. Again, we can see from these results that the overhead for load balancing is low and scales very well with the number of processors.

The load-balanced and original codes were run on 65, 536 MPI ranks, that is 16, 384 nodes with 4 cores per node. The grid size per element was 5× 5× 5 and the total number of elements was 900, 000. The total number of particles was 1.125×10^9 obtained as 1250 particles per element on average. Initially, the percent of elements that have particles was 6.1% of the total number of elements. The variable lelt was set to 140 for the load-balanced version. Figure 7.4 shows the differences in performance of load-balanced versus the original CMT-nek on Vulcan. The original version didn't finish in 5 hours. The time per time step for the original and the load-balanced versions was 20.00 and 2.52 seconds, respectively, giving us an overall speed-up of 7.9. Load balancing happened before simulation time step 1. There was no need to load balance

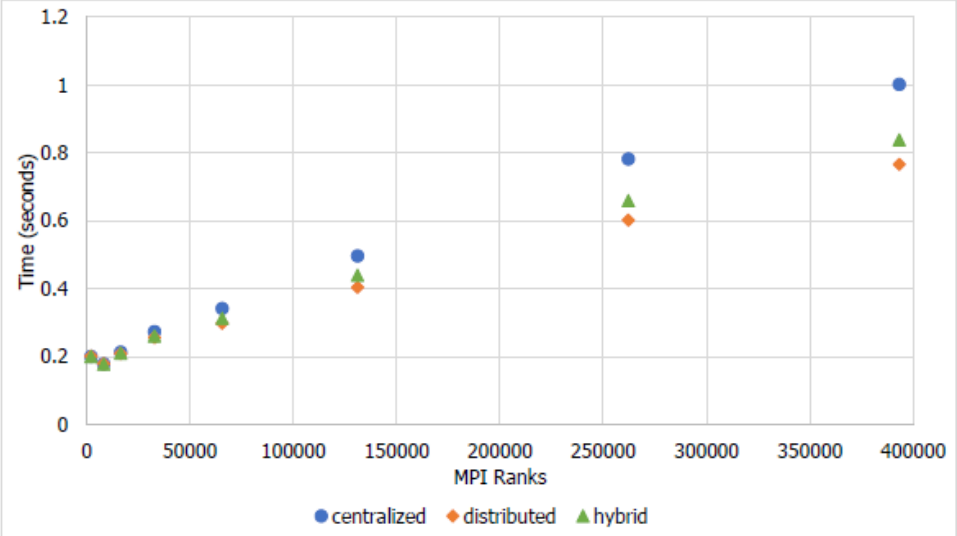

Figure 7.3 On Vulcan, total overhead for a load-balancing step for centralized, distributed and hybrid algorithms. It is a weak scaling with 2 elements per MPI rank, 5 × 5 × 5 grid points per element, and about 343 particles per element. The actual overhead expressed as the number of time steps for 393, 216 MPI ranks was 3.03 for the centralized, 2.33 for the distributed, and

after that since the time per time step didn't increase over the threshold set to trigger load balancing.

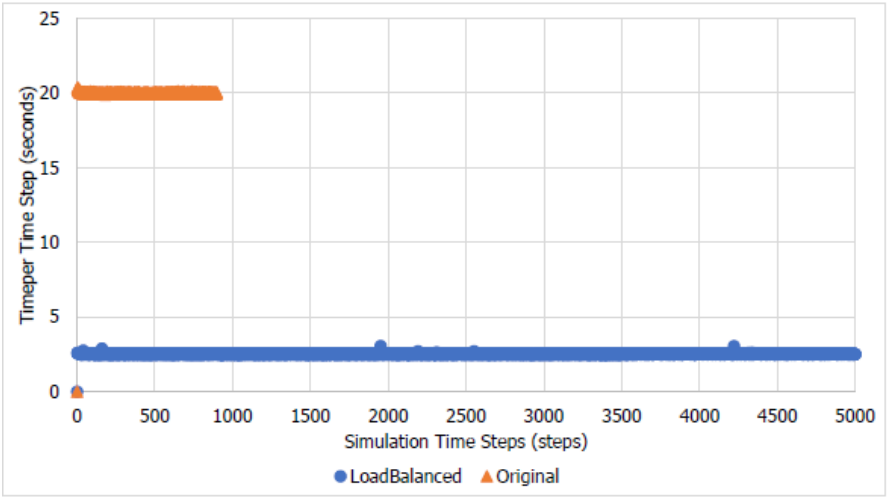

Figure 7.4 Performance comparison between load-balanced and original versions of CMT-nek on Vulcan. They were run on 65, 536 MPI ranks, that is 16, 384 nodes with 4 cores per node. Adaptive distributed load balancing was used. The average time per time step for the original and the load balanced versions was 20.00 and 2.52 seconds, respectively, giving us an overall speed-up of 7.9. Original code did not finish in 5 hours.

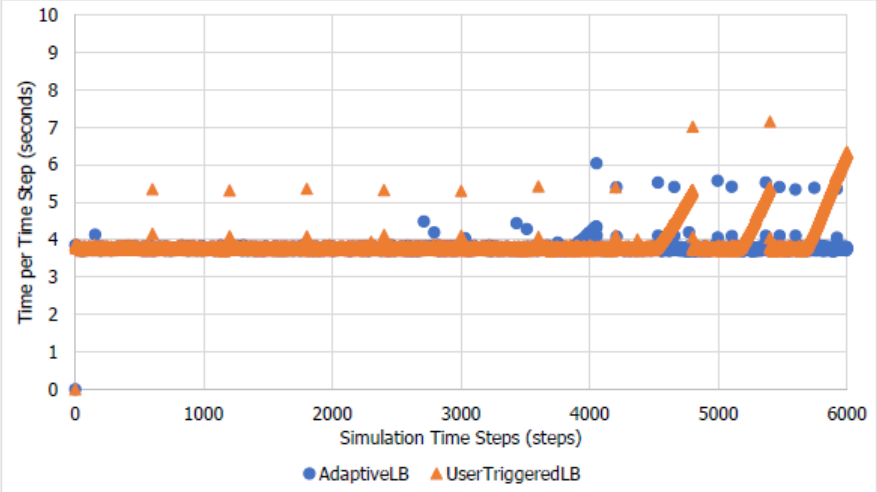

Figure 7.5 Performance comparison between adaptive load balanced and user-triggered load balanced versions of CMT-nek on Vulcan. They were run on 32,768 MPI ranks, that is 8,192 nodes with 4 cores per node. Distributed load balancing algorithm was used. The time per time step for the user-triggered and the adaptive load-balanced versions was 4.17 and 3.78 seconds, respectively, for the last 2,000 steps, giving us an overall improvement of 9.4%.

Figure 7.5 shows a comparison between the adaptive load-balancing and user-triggered load-balancing algorithms. For the user-triggered load-balancing algorithm, the k = 500, thus load balance is triggered every 500 time steps. As we can see from the figure, there is no performance degradation in the first 4, 000 time steps, making any load balancing redundant during this time. However, right after step 4, 000 performance degrades sharply, requiring frequent load balancing. The user-triggered load-balancing algorithm is insensitive to these performance variations and continues to load balance every 500 time steps. The average time per time step from step 4, 000 to step 6, 000 taken by the adaptive and user-triggered load balancing versions was 3.78 and 4.17 seconds, respectively. Thus, adaptive load-balancing algorithm gained an overall improvement of 9.4% compared to the user specified triggered load-balancing algorithm, and further the load balancing happens automatically without requiring any intervention by the user.

7.3 Load balancing CMT-nek (version 2)

We updated our load balancing code to be compatible to the major changes made to the Nek5000 code. This was more effort than anticipated because of the extent to which the Nek5000 code base was changed. Most of the changes in Nek5000 are changes in file names, rearrangement of code in the existing files, changes in subroutine names, compilation structure, third party software handling. CMT-nek also had its own fair share of changes in the meantime. The one relevant to load balancing was the addition of 4-way coupling methods in particle-particle interactions. A significant chunk of the effort was verifying the accuracy of the results after the merge.

We also implemented a second version of our load balancing algorithm. In the first version as described in the previous section, the load in our load balancing algorithm was based on the ratio of number of particles to the number of elements, that is based on how many elements and particles

were present in the current processor. However, this only works for the particle-element problem. We wanted to generalize it to work for a variety of problems. So, we implemented a time based load balancing algorithm. We first measure the computation time spent for each processor for the fluid part and particle part. Then based on these measured times, we predict the load on each element, that is the computational time for each element plus the computational time for the particles in this element. After doing this, we used this time measurement to find the percent imbalance between the various processors for that time step and input that information to the load balancing algorithms.

We applied the time-based load balancing algorithm on a cylindrical test case. In this test, a bed of particles is driven by a high-pressure gas. The high-pressure gas is placed at the center of the domain in a circular zone of radius 3.8mm. The domain extends from -5cm to 5cm in both x- and y-directions and 1.5mm in the z-direction. The ring of particles is placed around the charge in a 2cm radius bed. The bed is packed with an initial volume fraction of 60% and contains around 2 million particles. The pressure ratio between the high-pressure center and the ambient is 100. The particle-particle collisions are tracked using a soft-sphere collision model making this a 4-way coupled simulation.

Figure 7.6 shows a comparison between the new time-based load balancing and the original code with no load balancing. In this figure, the x-axis represents the number of time steps, and the y-axis represents the time per time step. The average time per time step for the load balanced one is about 1.12 second, while that for the original one is about 1.38 seconds. So, the load balanced algorithm gives about 1.24 times speedup in the execution time.

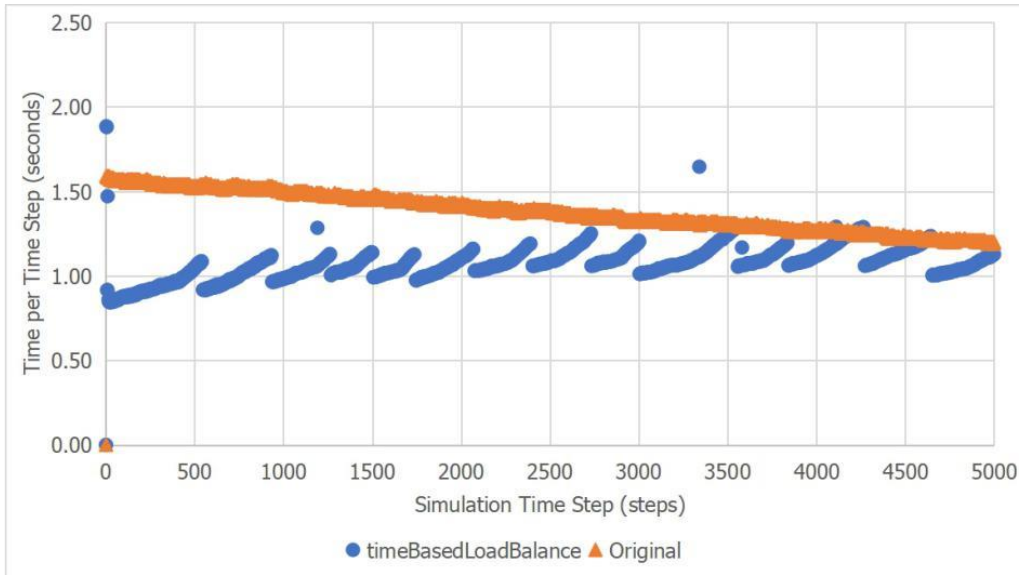

Figure 7.6: Comparison of load balancing CMT-nek using the time-based algorithm, versus a non-load balanced code. There are 134,832 elements each containing 4x4x4 grid points and the total number of particles is 2,091,923. The results are on 32,768 MPI ranks, that is 8192 nodes with 16 cores per node on IBM BGQ platform, Vulcan.

Figure 7.7 shows a comparison between the algorithm-based load balancing presented in the previous section and the original code. In this figure, the x-axis represents the number of time steps, and the y-axis represents the time per time step. The average time per time step for the load balanced one is about 0.72 second, while that for the original one is about 1.38 seconds. So, the load balanced algorithm gives about 1.9 times speedup in the execution time.

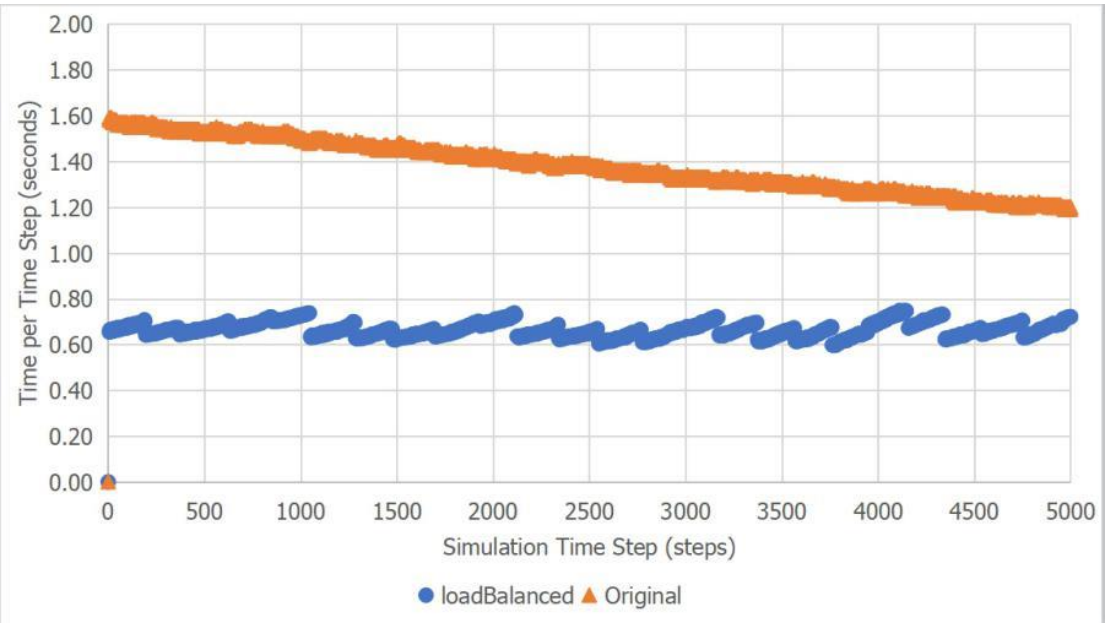

Figure 7.7 Comparison of load balancing CMT-nek using an algorithm-based load balancing, versus a non-load balanced code. There are 134,832 elements each containing 4x4x4 grid points and the total number of particles is 2,091,923. The results are on 32,768 MPI ranks, that is 8192 nodes with 16 cores per node on IBM BGQ platform, Vulcan.

7.4 GPU implementation of CMT-nek

The co-design team has been working to implement a version of a GPU-ized CMT-nek where spectral elements will be processed by both GPU and CPU simultaneously. We just finished porting the fluid portion of CMT-nek on GPU and validating the results against those on a CPU only implementation.

The CPU-GPU interaction model is as follows. All computationally intensive work is performed on GPU while the host CPU core aides in inter process communication involving the GPU and other CPU cores on the current node as well as across nodes. All the remaining CPU cores are used in computation also.

The high level architecture of this implementation is as follows. A spectral element is assigned to a streaming multiprocessor (SM) to minimize array data being read from the global memory. The number of threads dispatched equals the number of grid points (lx *ly *lz) for cases when calculations are required for each grid point. The number of threads dispatched equals the (number

of faces * lx * lz) when calculations are required for each face. Lx, ly and lz are the number of grid points along x, y and z direction, respectively, in the Cartesian coordinates.

We next describe some of the challenges in porting CMT-nek. From a hardware perspective, memory capacity was the primary constraint. 48KB of cache is available per SM. On the other hand, to store a fluid property on 512 grid points of one element (say, with lx=ly=lz=8), requires an array of size 4KB. Thus, only 12 such arrays can be stored, whereas there are about 50 such property arrays to be stored. To solve this problem we may reduce the number of grid points (which is possible by increasing the number of elements so that accuracy is not compromised). Or we may also split the kernels so that each kernel requires only a few arrays.

Another hardware constraint is that there are only 64K registers per SM. For the example where there are 512 grid points per element, each thread can be allotted only 128 registers. Additional data will have to be stored in the global memory, thereby reducing the potential speedup. This problem was solved by splitting the kernels and by changing the order of calculations to reuse registers.

An important challenge from coding perspective is that several Nek5000 subroutines for spectral element methods must be ported to GPU to achieve more speedup. These are currently executed on CPU and the results are copied to GPU. Also, all MPI calls require the GPU to stop processing and instead transfer data to CPU for inter process communication.

We also used CuBLAS matrix multiplication functions and these were used to replace native CMT-nek matrix multiplications. We also used new smaller testcases to test the code. Currently, we are at a point in the validation exercise where we are testing medium to large size problems.

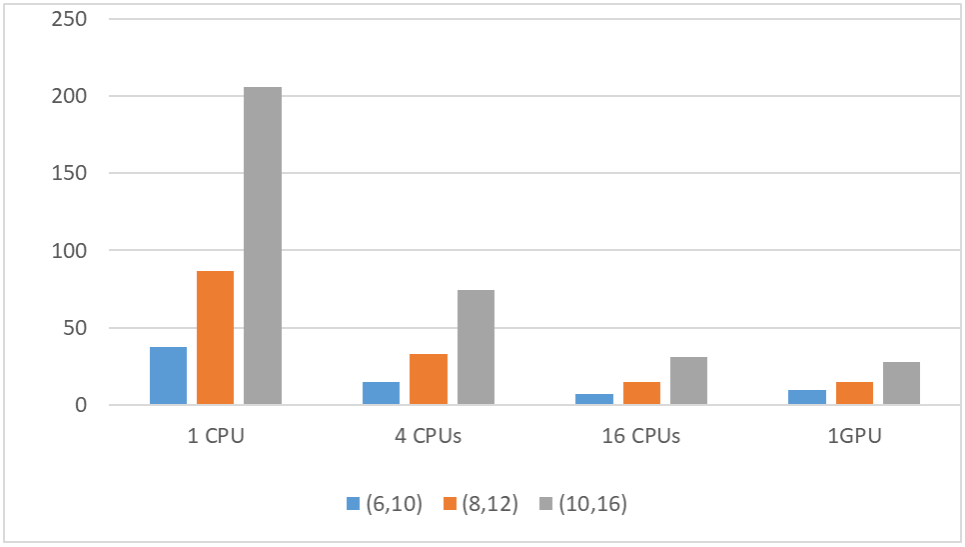

Figure 7.8 Comparison of GPU time to run CMT-nek versus the times needed by different number of CPUs. Each problem has 125 elements, for different grid sizes. GPU gives better advantage when the problem size is large.

Figure 7.8 shows some preliminary results in terms of the time taken by a GPU on a single node versus all the CPUs available on that node. Each of the three problems considered here has a total of 125 elements. The three problems differ in their sizes represented by a tuple (a, b), where a is the number of initial grid points and b is the number of grid points on a finer mesh obtained using spectral interpolation on the original grid. GPU is more advantageous to use when the problem size is large, as is true for the second and third problem sizes and in these cases the GPU outperforms the 16 CPUs available on the same node.

Titan@ORNL was used to validate the implementation. Except for surface routines all other CMT-nek subroutines have been verified to produce correct results. We are currently working to validate the two surface routines – namely, fluxes_full_field and surface_integral_full. Some of the functions that were verified to function properly are those used to update mesh and grid points, to compute primitive variables such as velocity and thermodynamic state from conserved variables, Jacobian computations, and gradient computations.

The main challenges faced while validating the results are floating point round off errors that make the GPU results slightly different than those of the CPU but the differences seem to propagate and increase with the time steps. For these tests we have to device a completely different validation workflow, which would include visualization and a statistical comparison of the data being generated.

Our main goal for the current year is to validate the GPU based CMT-nek code for fluids, using medium to large sized testcases.

8. Exascale Team

8.1 Overview

The Exascale Behavioral Emulation (BE) research team focuses on the use of BE methods to enable rapid design-space exploration (DSE) of design strategies, optimizations, and parallelization for extreme-scale systems up to Exascale. In Behavioral Emulation, low-level details are abstracted away to improve scalability, while retaining sufficient accuracy for design-space exploration. Such exploration supports the project by providing insight into how to optimize CMT-nek for potential existing, near-future, or notional architectures. A concept diagram illustrating the BE co-design process is shown in Figure 8.1.

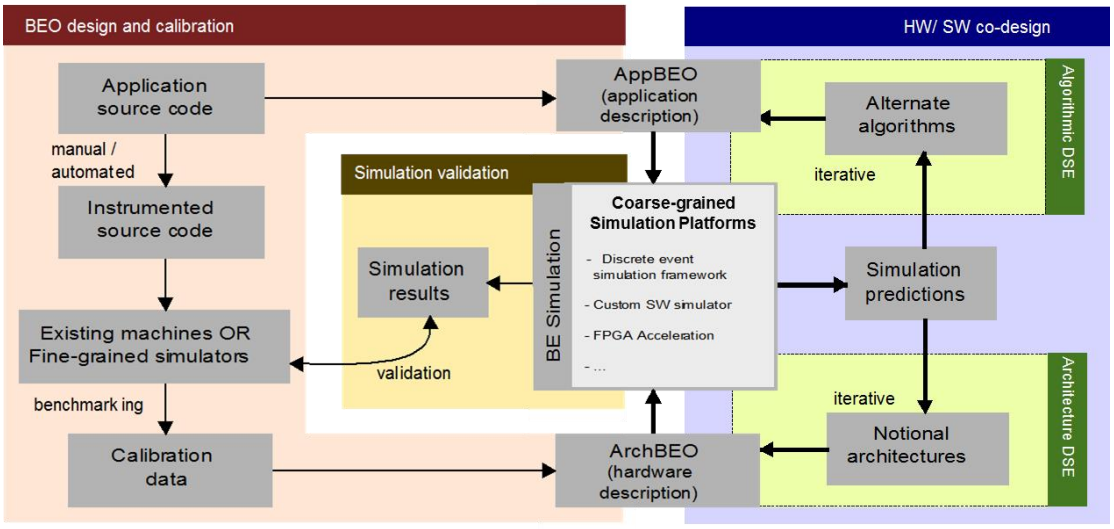

Figure 8.1. Co-design using Behavioral

In this section of the annual report, we will report our progress for this past year and the ongoing plans for the following tasks:

8.2 BE-SST simulator and architectural exploration using BE simulation

8.3 Trace-driven BE simulation of CMT-nek design space

8.4 FPGA-acceleration methods for rapid DSE space reduction & UQ

In Section 8.2, we first give an update on the enhancement of the BE-SST simulator and tools for the report period. This is followed by an example of how BE-SST is used for coarse-grained simulation of CMT-bone-BE. The example will demonstrate DSE on existing architectures, for validation, and for performance prediction on notional architectures. In addition, we will describe how BE is combined with Multi-Fidelity Surrogates (MFS) for low-cost model validation. In Section 8.3, the use of trace-driven BE simulation is illustrated with a case study: design space exploration of CMT-nek. Finally, In Section 8.4, we will describe our progress on using FPGAs to accelerate behavioral emulation for rapid DSE.

8.2 BE-SST and architectural exploration using BE simulation

After establishing the BE-SST simulator as our primary simulation platform, the main objective this past year was to achieve better scalability while also improving accuracy and performance. Also in this reporting period, we added support for symbolic regression for the creation of performance models, in addition to the existing interpolation API. Finally, we are adding support for fat-tree network topology for communication modeling. BE-SST, with these extensions, was then used to perform validation of our simulations on existing architecture, followed by performance predictions (in terms of execution time) on notional architectures.

The performance prediction on existing architecture was previously performed through BE-SST simulations of CMT-bone-BE of Vulcan and Titan. The average simulation accuracy was approximately 94%, thus satisfactorily validating our models to be used for notional prediction. We extended this work by simulating hypothetical Vulcan and Titan systems with increased core counts (dubbed Vulcan+ and Titan+ respectively) as our first step towards notional prediction.

One of the highlights the past year was the architectural exploration. We showcased the plug-and-play capability of BE-SST by switching between different compute nodes and interconnect topology. For this exploration, Quartz compute nodes, being an up-to-date architecture, were used. Our experiments paired Quartz compute nodes with with 3D and 5D Torus interconnect topology; i.e., the performance prediction of CMT-bone-BE on Quartz with 5D torus (let us call Machine 1) was compared against Quartz with 3D torus (let us call Machine 2). We also compared the performance of Vulcan against Machine 1 and Titan against Machine 2. In the latter comparison, we are trying to analyze the performance of the compute nodes while keeping the interconnect topology constant.

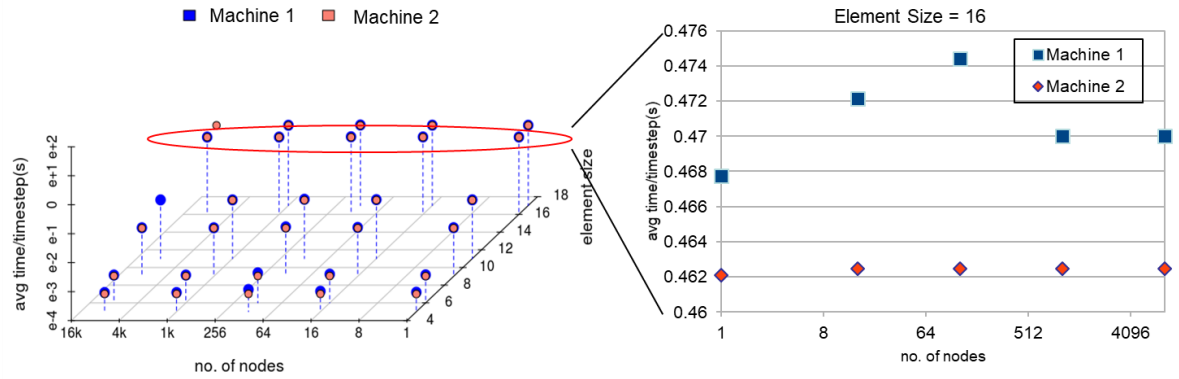

Figure 8.2. Machine 1 vs Machine 2 for node architecture comparison

The results for Machine 1 vs. Machine 2 are shown in Figure 8.2. As explained earlier, the node architecture is fixed as Quartz compute node; whereas the interconnect topology is varied between 5D- and 3D-Torus (for Machine 1 and Machine 2, respectively). As expected, we observed very similar performance for Machine 1 and Machine 2, since CMT-bone-BE is compute dominant, with minimum communication cost. The result on right of Figure 8.2 shows that Machine 2 (with 3D-Torus topology) is slightly faster than Machine 1 (with 5D-Torus) due to low latency on Titan's 3D-Torus network.

The result for the node architecture comparison is shown in Figure 8.3. In this case, the interconnect topology is fixed, first to 5D-Torus, and the performance of the node architectures of Quartz and Vulcan are compared. Then, the interconnect topology is fixed to 3D-Torus and the performance of the node architectures of Quartz and Titan are again compared. Figure 8.3 indicates that Machine 1 and Machine 2 with Quartz compute node is faster than Vulcan compute node by almost 20 times and Titan compute node by almost 5 times. This is an expected result, as Quartz's Intel Xeon CPUs are newer with faster compute nodes compared the Vulcan's Power7 and Titan's AMD Opteron. These results showcase the capability of BE-SST's plug-and-play feature. This feature will be used going forward to study different notional architectures in our exploration of abstract machines.

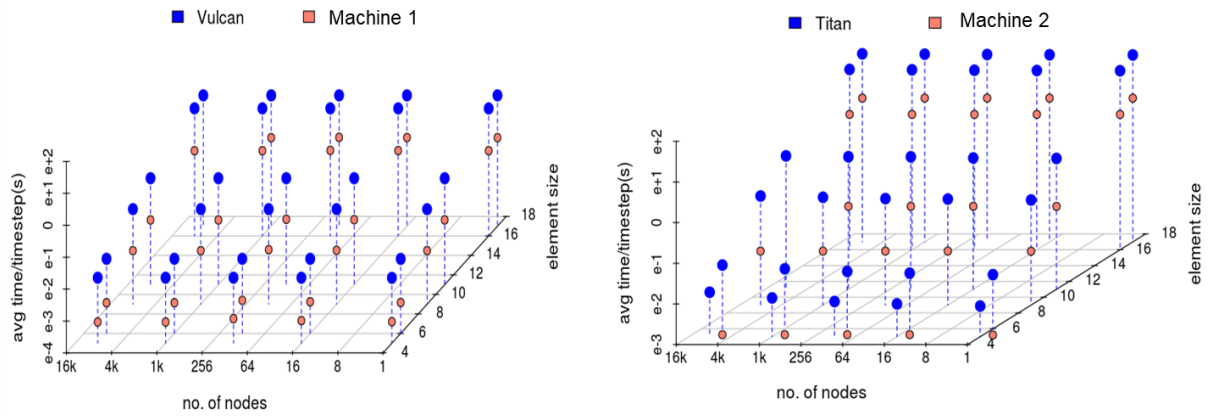

Figure 8.3. Vulcan vs Machine 1 (left) and Titan vs Machine 2 (right) for interconnect topology comparison

Finally, we have started to include abstract networking topologies into the BE-SST framework. Our goal is to expand the breadth of our design-space exploration research by allowing new network topologies, both existing and theoretical, to be simulated in BE-SST. We can combine this feature with our existing application and architecture BEOs to take more defined steps toward notional machine simulation. We have started adding non-Cartesian network topologies into BE-SST. Currently, we are working on implementing fat-tree topology as this is becoming more common on newer HPC systems.

Going forward, the plan is to merge BE-SST into SST's GitHub repository to make it open-source.

8.3 Trace-driven BE simulation of CMT-nek design space

In general, BE simulations can be classified into two categories. In the first category, the workload on each of the processors in the system is known prior to execution. The workload consists of both the computation and communication workload on each processor and in turn on the entire system. The workload is generally presented in terms of the problem size of the application of interest, which would be given as the input parameters to the simulation. In some cases, where the workload is not fixed but we have information on how it varies throughout the execution, we use an approximation function to predict the workload to perform the simulation.

However, there is a second category in which the workload of an application varies *dynamically* and the variation is dependent on the specific problem. Additionally, this variation is not quantifiable using the input parameters. For this case, we have introduced a novel *trace-driven simulation* method for Behavioral Emulation, as described in this section.

8.3.1 Particle-workload distribution tool

To perform a trace-driven simulation, a trace from the actual execution is used to calculate the workload at each stage of execution. However, collecting traces for every run (even for the same type of problem) to perform BE simulations defeats the main purpose of performing simulation. The main purpose of the simulator is to predict and study the behavior of an application or a specific use-case of application on a specific machine or architecture without actually executing it on the platform.

To support efficient trace-driven simulation, we have developed a particle-workload distribution tool to predict the workload of CMT-nek on a different system for a given specific problem trace. The key advantage for this tool is that we can use a single trace to calculate the workload on any system.

To demonstrate and validate our tool, the ASU Shock-tube experiment was used as a case study. BE-SST was used to simulate the CMT-nek particle-solver kernel, which is the most expensive kernel in the CMT-nek application. As mentioned, in case of trace-driven simulation, the workload of this kernel is dependent on the number of particles each processor owns, which is dynamic. Figure 8.4 below shows an example of how the particles may distribute across the processors. We randomly plotted two ranks while running the Shock-tube demonstration problem using 2048 cores on Vulcan. As you can see, there is a distinction on workload (which is dependent on number of particles) among different processors, and there is a variation within a processor during execution.

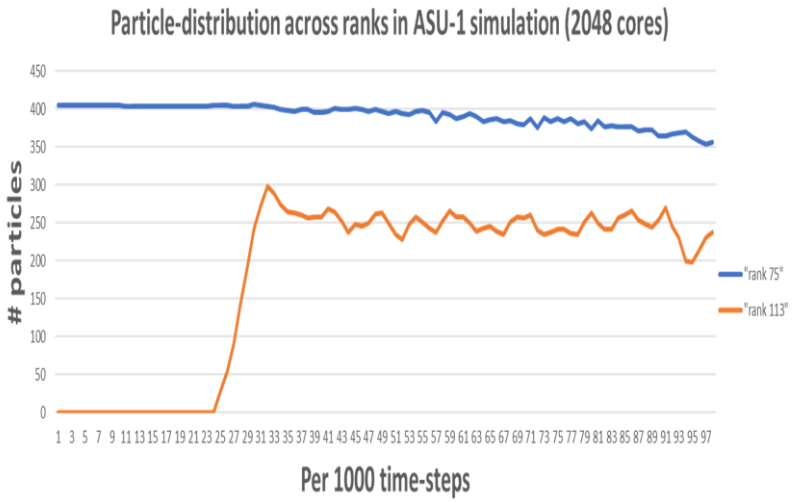

Figure 8.4: Particle-workload distribution across different ranks of CMT-nek application on 2048 cores Vulcan machine performing ASU-1(shock-tube) simulation.

In such cases, where there is a huge variation of workload among processors, we need a trace to perform accurate BE-simulation in order to predict the actual performance of the system. We performed a strong-scaling study where, for a given problem size, we have seen how the workload varies for different machine configurations. Figure 8.5 summarizes the strong-scaling results for two different mapping algorithms and its effect on computation cost, communication cost, and resource utilization.

As shown in Figure 8.5(a), as the number of cores increases, the peak workload on each processor decreases. Figure 8.5 shows that there is a clear advantage in using a load-balancing mapping algorithm which takes into consideration the effect of particles while mapping elements to processors. Such a load-balancing algorithm will have less peak-workload as compared to the default mapping algorithm (genmap) which does not consider the particle workload (Figure 8.5(a). Similarly, you can see that by using a load-balancing algorithm, the percentage of the processors with zero workload (zero particles) can be reduced, from **85%** to **15%** (Figure 8.5(b)). In the strong-scaling study, one of the costs that we expect to increase is the communication cost. Figure 8.5(c) shows the increase in the communication cost for different machine configurations. The load-balancing algorithm results in more communication, as the particles are now more evenly distributed; however, the increased communication cost is insignificant relative to the saved computation cost.

The main take-away point is that using the particle-workload distribution tool, we have *predicted* how the workload varies on different machine configurations by just a *single trace*, without actually running the application on the platform using these configurations. As a result, using this tool in conjunction with BE-SST, we can find the optimal machine configuration for a specific problem by predicting the computation cost, communication cost and resource utilization. Currently we have an implementation of the particle-workload distribution tool; and are working on integrating the output of the tool into BE-SST in order to calculate the actual performance used in determining the optimal configuration.

8.3.2 Improvement to particle-workload distribution tool

In the second quarter of the reporting period, we made key changes to the particle-workload distribution tool. One such change was to incorporate a synthetic trace generator. In the earlier version of the tool, although a single trace was sufficient to predict the workload on any machine configuration, if we changed the number of particles and/or elements, we would need to collect a new trace in order to calculate workload predictions. To eliminate this limitation, we introduce a synthetic trace generator, which takes an existing trace as an input along with scaling ratios of particle and element and generates a new synthetic trace based on new scaling ratios. One key assumption that we make is even though we increase/decrease the number of particles/elements, as the problem specifications are the same, the particles move in a similar fashion. Hence, for a specific demonstration problem, we can use a relatively smaller problem set to collect a trace and use it to predict the workload of bigger simulations, as long as we are simulating the same problem. This approach significantly reduces the time taken to collect the trace.

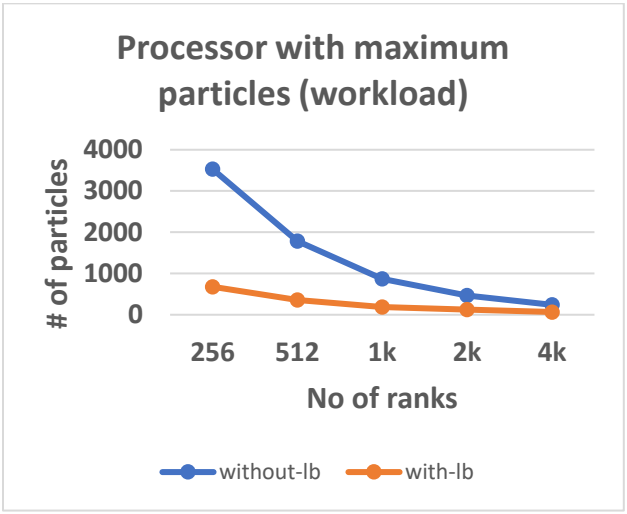

(a)

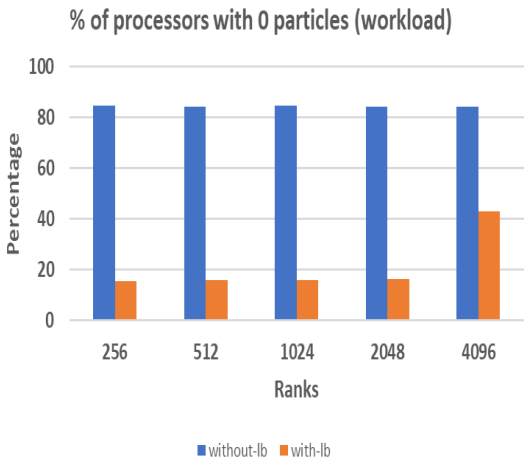

(b)

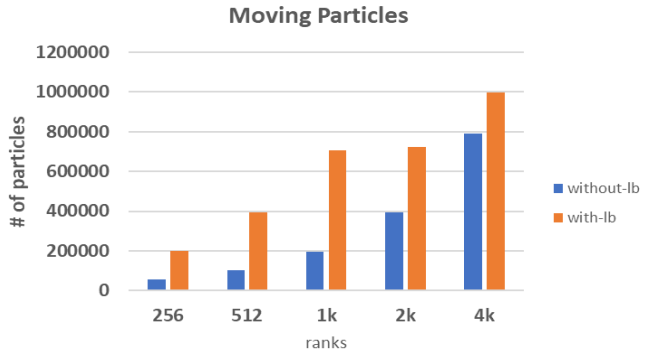

(c)

Figure 8.5: Strong-scaling study: (a) Peak Particle-workload (b) Percentage of processors with 0 workload and (c) number of particles being communicated across different configurations for ASU-1 (shock-tube) problem simulated using CMT-nek.

More recently, we updated the particle workload distribution tool based on the recent changes in the CMT-nek algorithm. CMT-nek now includes two-way and four-way coupling, which calculates the particle effect on the surrounding fluid grid points, and also calculates the collision forces among particles. Two-way coupling generates ghost particles if the particles in question have influence on the neighboring fluid gridpoints that are outside the current processor. These ghost particles need to be sent to the corresponding processors whose fluid gridpoints are influenced. Hence, the computation and communication cost not only depend on the number of actual particles, but also on the ghost particles residing on each processor. The size of the bin depends on the particle's *Zone of Influence*, which in turn determines the number of ghost particles. In case of four-way coupling, each particle checks with all the other particles within the

surrounding bins and calculates the collision forces. In both cases mentioned above, the size of the bin affects the total computation and communication cost. Using particle-workload prediction tool, we are aiming to predict the computation and communication workload associated with the ghost particles for different bin sizes using the application trace.

8.3.3 Application of particle-workload distribution tool

In collaboration with the CMT-nek code development team, we have identified a set of key questions to be answered using the particle-workload :

1. What is the optimal configuration to run for a given problem? (or) How the resource utilization varies across different processor configurations?
2. How would the problem scale for a greater number of elements?
3. What is effect of projection filter on application performance?

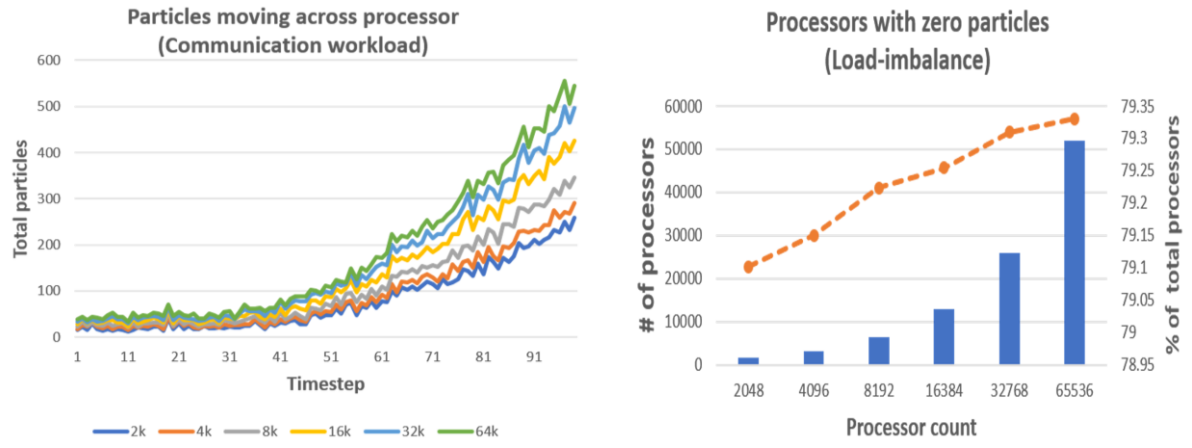

Figure 8.6: ASU- Cagnolis experiment. (a) Communication workload of the experiment. (b). Load imbalance of the experiment.

To answer these questions, the simulation problem of the ASU-1 Cagnolis experiment is used. The simulation contains 4.2 million particles and 74,304 elements, with an element grid-size as 5. The original trace has been collected at a frequency of 10 time-steps. We tried to answer the first question, by predicting the computation and communication workload for the given problem size on different processor counts. Figure 8.6 shows the strong scaling study where the plot shows how many particles are moving across each processor for different processor configurations. This is indicative of the communication workload of the processor and across the system. As you can see, more particles are being moved from the processor as the simulation progresses, which is expected as the particles start moving in the simulation. Also, the number of particles being moved across processors increases with the processor count. As a result, while increasing the processor count reduces the per-processor particle count, inter-processor communication increases. Using the particle-workload distribution tool, we can quantify the computation and communication workload in terms of number of particles, which when fed into the BE-simulator, predicts the total performance using the derived BE performance models

In order to answer the second question, we doubled the number of elements along both X and Z directions and predicted how that would affect the particle workload across processors. In the original trace, there is a total of 74,304 elements in a 16x774x6 box configuration. We created a synthetic trace for 32x774x12 and 64x774x24 configurations, thereby increasing the total elements by 4x and 16x respectively. Based on the results, we observed there is no significant change in the particle workload distribution with increase in element count. These predictions were validated with actual runs, where we increased the elements by 4x and 16x in the configurations mentioned above. The validation results are similar to the predictions. Figure 8.7 shows the load-imbalance factor (# of processors with 0 particles) for different element configurations with corresponding validation results. As you can observe, the predictions and validation match, and we observe a slight decrease in the load imbalance with more elements.

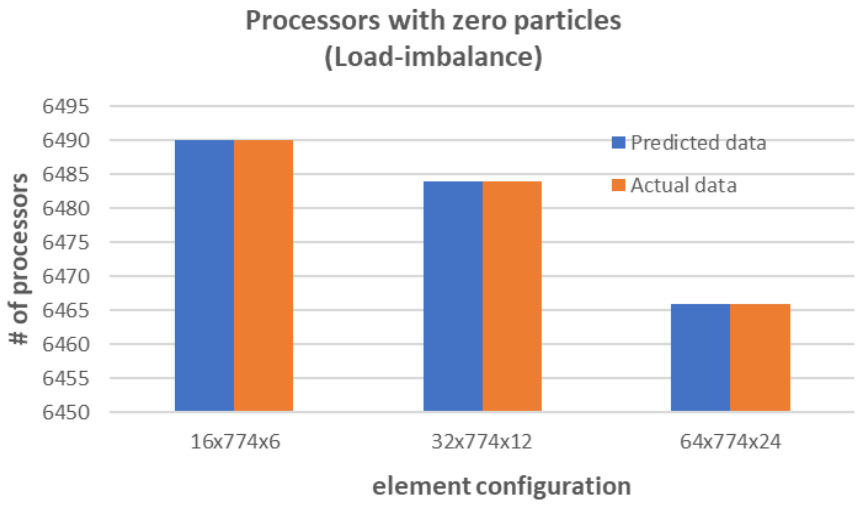

Figure 8.7: Comparison of Predicted and Actual particle workload for different element configurations

Finally, the third question considers the effect of the filter-size on the application performance. Filter-size determines the zone of influence, which is used to calculate the number of ghost particles present in a process. Most of the expensive subroutines in particle-solver phase are performed on ghost particles as well as the actual particles. Hence, the number of ghost particles present in a processor also affects the computation time. Similar to particle communication, ghost particles are also communicated across processors at each time-step. Depending on the number of ghost particles being moved across processors, corresponding communication time could be significant. Figure 8.8 shows the results from actual runs of the effect of the filter-size on (a) number of ghost particles and (b) total execution time. We are currently working on predicting the performance using the workload-distribution tool for varying filter-sizes.

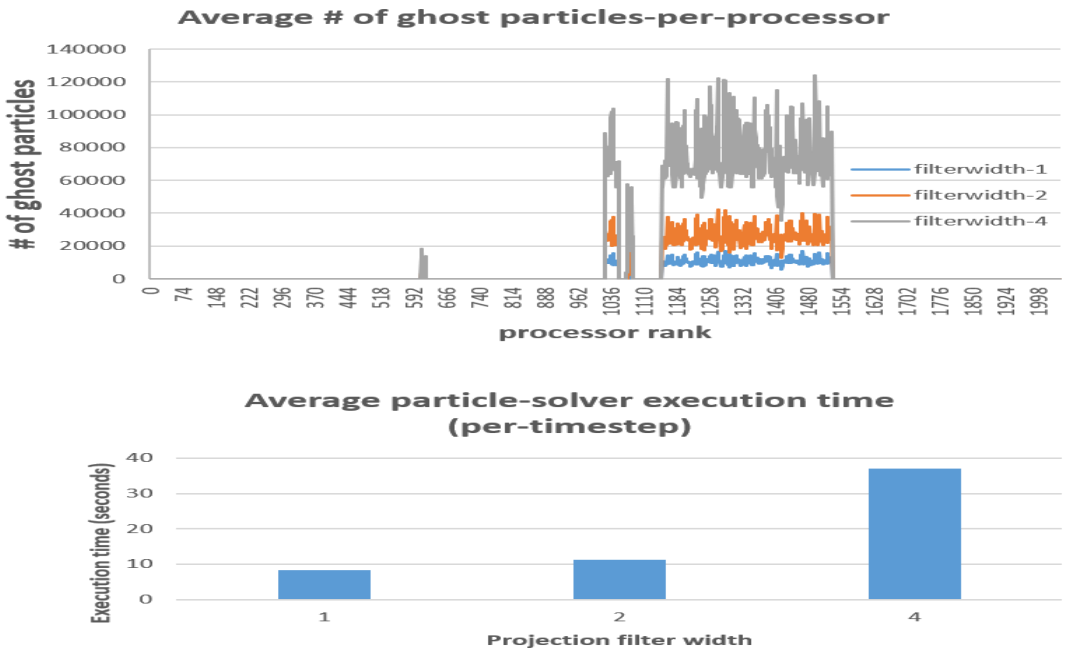

Figure 8.8: The effect of filter size on (a) the number of ghost particles per processor, and (b) average timestep length.

8.4 FPGA-acceleration methods for rapid DSE space reduction & UQ

The BE approach is based on distributed discrete-event simulation (DES) between BE objects (BEOs). BEOs mimic the abstract behavior of applications or architectures using high-level surrogate models to predict each event's impact on the targeted simulation objectives (e.g., performance, reliability, power). Although this approach enables faster simulation than traditional cycle-accurate approaches, Exascale simulation with existing tools could take minutes, hours, or even days to complete a single simulation. These lengthy simulations place very practical limits on DSE and Uncertainty Quantification (UQ) efforts that often require thousands, or even millions of independent simulations. To address this issue, we are exploring an FPGA-accelerated approach based on pipelined DES that focuses not necessarily on improved performance for a single simulation, but instead on increased simulation throughput. By focusing on throughput, we unlock the potential for huge performance gains when the problem under study calls for numerous independent simulations (e.g., DSE, Monte Carlo simulation).

We have previously reported performance for FPGA-pipelined simulations of CMT-bone-BE configurations up to 2,147,483,648 threads, which showed simulation/event throughput that is many orders-of-magnitude faster than the BE-SST software simulator. For this reporting period, our FPGA acceleration efforts have mainly focused on creation/integration of a tool that auto generates performance-model pipelines based on provided symbolic-regression equations and the

development of an end-to-end DSE case study meant to showcase the utility and appropriate use cases for our FPGA and BE-SST simulators. Using the CMT-nek particle-solver kernel discussed in section B, we consider a huge design space containing hundreds of thousands of parametric and algorithmic combinations (Table 8.1). In this case study, we use FPGA simulation to explore and reduce the entire design space, before using BE-SST to more closely analyze candidates of interest we determine from the FPGA data.

Parameter	Range
Element-size	5-25 (21)
Element-count	4-512 (8)
particles/gridpoint (α)	0.1-10 (19)
# of ranks	16-1M (17)
Algorithmic options:	
Time Integration	Rk3,bdf
Interpolation	Barycentric, Reduced Barycentric, Trilinear interpolation

Table 8.1: Parametric and Algorithmic Design Space of CMT-nek particle-solver kernel constituting of a total 273904 design candidates.

Table 8.1 shows a total of 273,904 design candidates. Although BE-SST is highly scalable and performs simulations of larger systems within a few minutes, use of BE-SST on all design space candidates mentioned above would be practically infeasible. Analysis of BE-SST throughput parallelized on 64 cores of HiPerGator (a UF HPC system) indicates that exploration of the entire design space (1 simulation per 273904 candidates) would take 5.5 years to complete. Performing 100 Monte-Carlo simulations per candidate increases this to 550 years. The FPGA simulator can explore the entire design space with 100 Monte-Carlo simulations per candidate in only 61 minutes (Table 8.2 shows the throughputs of BE-FPGA and BE-SST along with the relative speedup).

Ranks	No of events	FPGA Avg. Simulation time (microseconds)	BE-SST* Avg. Simulation time (seconds)	FPGA speedup
1024	24,576	3.53	8.234	2.33×10^6
2048	49,152	7.06	6.5	9.21×10^5
4096	98,304	14.12	17.3	1.22×10^6
8192	196,608	28.24	31	1.10×10^6
16k	393,216	56.48	120	2.12×10^6
32k	786,432	112.96	145	1.28×10^6
64k	1,572,864	225.92	405	1.79×10^6
128K	3,145,728	451.84	531	1.17×10^6
256K	6,291,456	903.68	1569	1.74×10^6
512K	12,582,912	1807.36	7416	4.10×10^6
1M	25,165,824	3614.72	X	X

Table 8.2: Throughput and relative speedup of BE-FPGA and BE-SST.

Going forward, we are developing a second DSE case study with additional parametric, algorithmic, and architectural design-space parameters and meant to provide useful feedback to the CMT-nek development team. Our study aims to provide a detailed cost benefit analysis for candidate algorithm implementations before code is developed and integrated into CMT-nek. The case study will span across a configurable number of FPGAs in Novo-G# (a large-scale reconfigurable computer), which will allow us to test the scalability of this approach for more than a single FPGA.

Also in this reporting period, we have begun work on attempting to accelerate the kernels of the CMT-nek application itself, so as to reduce computation run time on real experiments. In order to achieve this, we seek to extract parallelism inherent in the application behavior and accelerate it with custom circuitry on reconfigurable FPGA hardware. Our current target system for this experimentation is a new Intel technology, which combines Skylake Xeon CPU cores and Arria 10 FPGA fabric on a single package, with the goal of significantly decreasing communication overhead (a common bottleneck for FPGA applications) via shared memory and on-package communication. Using the FPGA to exploit parallelism and the tightly-coupled CPU for inherently serial behavior, we hope to be able to show performance speedup over a traditional CPU, as well as better performance-per-watt over a GPU accelerator due to the low-power nature of FPGA technology. In addition to speeding up CMT-nek, we wish to show a high level of scalability for the Xeon + FPGA system for theoretical workloads with larger communication requirements based on the combination of lower cost on-package communication and energy-efficient parallel computation.

9. Deep Dives

9.1 Exascale Deep-dive

The University of Florida held a Deep Dive Workshop on Feb 3-4, 2015. The Agenda is presented below, and the presentation slides can be found on the CCMT webpage at:

https://www.eng.ufl.edu/ccmt/events/workshops/

Agenda:

Deep Dive
University of Florida
February 3-4, 2015

Current Attendee List:

Bob Voigt	NNSA HQ	rvoigt@krellinst.org
Matt Bement	LANL	bement@lanl.gov
David Daniel	LANL	ddd@lanl.gov
Dave Nystrom	LANL	wdn@lanl.gov
Maya Gokhale	LLNL	maya@llnl.gov
Martin Schulz	LLNL	schulzm@llnl.gov
Jim Ang	SNL	jaang@sandia.gov
Arun Rodrigues	SNL	afrodri@sandia.gov
Jeremy Wilke	SNL	jjwilke@sandia.gov
S. Balachandar “Bala”	University of Florida	bala1s@ufl.edu
Alan George	University of Florida	george@hcs.ufl.edu
Rafi Haftka	University of Florida	haftka@ufl.edu
Herman Lam	University of Florida	hlam@ufl.edu
Sanjay Ranka	University of Florida	ranka@cise.ufl.edu
Greg Stitt	University of Florida	gstitt@ece.ufl.edu
Tom Jackson	University of Florida	tlj@ufl.edu
Tania Banerjee	University of Florida	tmishra@cise.ufl.edu

University of Florida Students:

Dylan Rudolph	rudolph@hcs.ufl.edu
Nalini Kumar	nkumar@hcs.ufl.edu
Carlo Pascoe	carlo.pascoe@gmail.com
Kasim AlliKasim	kasimalli490@yahoo.com
Chris Hajas	chrishajas@ufl.edu
Mohammed Gadou	mgadou@ufl.edu
Michael Retherford	

UF Deep dive agenda:

Tuesday, February 3, 2015

8:20 Van pickup at Hilton

8:30 – 9:00 Breakfast

9:00 – 9:30 **Welcome and Deep-Dive Overview (3 Sessions)**

1. Behavioral emulation (BE): modeling & simulation/emulation methods
2. CS issues (performance, energy, and thermal)
3. Use of reconfigurable computing to accelerate behavioral emulation

* Each of the three deep-dive sessions is designed to be ***interactive*****:** a combination of short presentations by UF and Tri-lab researchers, intermixed with discussion, demonstrations, etc.

9:30 – 11:30 **Session 1: Behavioral Emulation: Modeling & Simulation/Emulation Methods**

- UF topics:
  - Behavioral characterization
  - Parameter estimation
- Tri-lab topics:
  - Overview of FastForward 2 and DesignForward 2 (Jim Ang, SNL)
  - Multi-scale architectural simulation with the Structural Simulation Toolkit (Arun Rodrigues, SNL)

11:30 – 12:30 Lunch

12:30 – 2:00 **Session 1 (continued): Behavioral Emulation: Beyond Device Level**

- UF topics:
  - Synchronization for speed
  - Congestion modeling
  - Behavioral characterization & modeling beyond device level
- Tri-lab topics:
  - Using discrete event simulation for programming model exploration at extreme-scale (Jeremy Wilke, SNL)
  - ASC next-generation code projects (David Daniel, LANL)

2:00 – 5:00 **Session 2: CS Issues (Performance, Energy, and Thermal)**

- UF topics:
  - Performance and autotuning for hybrid architectures
  - Energy and thermal optimization
  - Dynamic load balancing
- Tri-lab topics:
  - Performance, energy, and thermal benchmarking (Jim Ang, SNL)
  - Why power is a performance issue: utilizing overprovisioned systems (Martin Schulz, LLNL)

* There will be an afternoon coffee break in this time slot

6:30 Dinner (University Hilton)

Wednesday February 4, 2015

8:20	Van pickup
8:30 – 9:00	Breakfast

9:00 – 11:00	**Session 3: Use of Reconfigurable Computing to Accelerate Behavioral Emulation**
	• UF topics:
	  ○ Efficient mapping of behavioral emulation objects (BEOs) onto a system of FPGAs
	  ○ Demo of current single FPGA prototype
	  ○ Transitioning to multiple FPGAs
	  ○ Challenges associated with maximizing emulation speed while maintaining scalability/usability
	• Tri-lab topic:
	  • FPGA-based emulation of processing near memory (Maya Gokhale, LLNL)

11:00 – 12:00	Open discussion and planning for action items
12:00	Box lunch; transportation to airport as needed.

9.2 Multiphase Physics Deep-dive

The University of Florida held a Multiphase Physics Deep Dive Workshop on October 6-7, 2016. The Agenda is presented below, and the presentation slides can be found on the CCMT webpage at:

https://www.eng.ufl.edu/ccmt/events/workshops/

Agenda:

Multiphase Physics Deep-dive Workshop

October 6-7, 2016 Attendee List

Georges Akki	gakiki@ufl.edu	University of Florida
Subramanian Annamalai	subbu.ase@ufl.edu	University of Florida
Marco Arienti	marient@sandia.gov	Sandia
S. Balachandar	bala1s@ufl.edu	University of Florida
Ankur Bordoloi	ankur@lanl.gov	LANL
Alexander Brown	albrown@sandia.gov	Sandia
Jesse Capecelatro	jcaps@umich.edu	University of Michigan
Seungwhan Chung	schung58@illinois.edu	University of Illinois
Paul Crittenden	pcritte@ufl.edu	University of Florida
William Dai	dai@lanl.gov	LANL
Angela Diggs	angela.diggs.1@us.af.mil	Eglin AFB
Timothy Dunn	dunn13@llnl.gov	LLNL
Brad Durant	neoncrash@ufl.edu	University of Florida
John Eaton	eatonj@stanford.edu	Stanford University

Mahdi Esmaily	mesmaily@stanford.edu	Stanford University
Giselle Fernandez	giselleferndandez@ufl.edu	University of Florida
Marianne Francois	mmfran@lanl.gov	LANL
Joshua Garno	jgarno@ufl.edu	University of Florida
Jason Hackl	jason.hackl@ufl.edu	University of Florida
Alan Harrison	alanh@lanl.gov	LANL
Jeremy Horwitz	horwitz1@stanford.edu	Stanford University
Kyle Hughes	kylethughes89@ufl.edu	University of Florida
Thomas Jackson	tlj@ufl.edu	University of Florida
Rahul Koneru	rahul.koneru@ufl.edu	University of Florida
Allen Kuhl	kuhl2@llnl.gov	LLNL
Ali Mani	alimani@stanford.edu	Stanford University
Yash Mehta	ymehta@ufl.edu	University of Florida
Chandler Moore	wcm0015@tigermail.auburn.edu	University of Florida
Brandon Morgan	morgan65@llnl.gov	LLNL
Balu Nadiga	balu@lanl.gov	LANL
Fady Najjar	najjar2@llnl.gov	LLNL
Brandon Osborne	bosborne3@ufl.edu	University of Florida
Fred Ouellet	f.ouellet@ufl.edu	University of Florida
John Parra-Alvarez	jcparraa@gmail.com	University of Utah
Katherine Prestridge	kpp@lanl.gov	LANL
Bertrand Rollin	bertrandrollin@gmail.com	Embry-Riddle
Kevin Ruggirello	kruggir@sandia.gov	Sandia
Kambiz Salari	salari1@llnl.gov	LLNL
Shane Schumacher	scschum@sandia.gov	Sandia
Philip Smith	philip.smith@utah.edu	University of Utah
Sean Smith	sean.t.smith@utah.edu	University of Utah
Prashanth Sridharan	shan1130@ufl.edu	University of Florida
Jeff St. Clair	jeff.stclair@navy.mil	University of Florida
Cameron Stewart	csstewart10@ufl.edu	University of Florida
Jeremy Thornock	jthornock@gmail.com	University of Utah
Gretar Tryggvason	gtryggva@nd.edu	University of Notre Dame
Markus Uhlmann	markus.uhlmann@kit.edu	Karlsuhe Institute of Tech.
Laura Villafane	lvillafa@stanford.edu	Stanford University
Robert Voigt	rvoigt@krellinst.org	ASC AST
Seng Keat Yeoh	syeoh2@illinois.edu	XPACC University of Illinois
Duan Zhang	dzhang@lanl.gov	LANL
Ju Zhang	jzhang@fit.edu	Florida Institute of Tech.

Agenda Multiphase Physics Deep-Dive, October 6-7, 2016
St. Petersburg Marriott Clearwater,

12600 Roosevelt Blvd, North St. Petersburg, FL 33716, Phone: 727-572-7800

Thursday Oct 6th, 2016

7:30 – 8:30 Breakfast (Provided)

8:30 – 8:45 Welcome (S. Balachandar)

8:45 – 10:15 Overviews

Gretar Tryggvason Challenges and opportunities in fully resolved simulations of multi fluid flows

Ali Mani Overview of computational modeling at Stanford PSAAP: particle-laden flows subject to radiative heating

Kambiz Salari Research activities for energetic dispersal of particles

10:15 – 10:30 Coffee

10:30 – 12:00 Overviews

Alex Brown Multiphase Methods for Modeling Fire Environments

Phil Smith Multi-phase flow modeling at Utah PSAAP – predictivity in application

S. Balachandar Overview of multiphase flow computational strategy at UF PSAAP

12:00 – 12:15 Further Discussion

12:15 – 1:30 Lunch (Provided)

1:30 – 3:00 Macroscale

Duan Zhang Equations and Closures for Deformation and Flow of Continuous and Disperse Materials

Allen Khul 3-Phase Model of Explosion Fields

Marco Arienti Multiphase Flow Simulation Strategies at the CRF

3:00 – 3:15 Coffee

3:15 – 5:15 Microscale

Jeremy Horwitz Point-particle modeling for two-way-coupled problems: Challenges, verification, and physics-based improvements

Georges Akiki Extended point particle model

Jesse Capecelatro Recent insights on turbulence modeling of strongly-coupled particle-laden flows

Tom Jackson Microscale simulations of shock particle interaction

6:30 – 9:00 Dinner (Provided – all attendees)

UF UNIVERSITY of FLORIDA Center for Compressible Multiphase Turbulence

Friday Oct 7th, 2016

7:30 – 8:30 Breakfast (Provided)

8:30 – 10:30 Modeling & numerical methods

Sean Smith Particle dynamics: coal-specific modeling

A. Subramaniam Microscale modeling based on Generalized Faxen theorem

Alan Harrison Modeling of Ejecta Particles in the FLAG Continuum Mechanics Code

Markus Uhlmann Large scale microscale simulations and modeling opportunities

10:30 – 10:45 Coffee

10:45 – 12:15 Experiments and simulations

Ankur Bordoloi Experimental measurements of drag on shocked particles

Laura Villafañe Including real experimental effects in validation of numerical models for confined particle-laden flows

Fady Najjar Meso-scale Simulations of Shock-Particle Interactions

12:15 – 1:30 Lunch (Provided)

1:30 – 3:00 Meso/macroscale

Mahdi Esmaily A systematic study of turbophoresis by four-way-coupled simulation of Stokesian particles in channel flow

Balu Nadiga Bayesian Analysis of Inter-Phase Momentum Transfer in the Dispersed Eulerian Formulation of Multiphase Flow

John Parra-Álvarez Eulerian Models and Polydispersity Treatment for Dilute Gas-Particle Flows

3:00-4:00 Discussion and Closing remarks

9.3 Nek5000 Users/Developers Meeting

6th Nek5000 Users/Developers Meeting
Renaissance Tampa International Plaza Hotel
April 17-18, 2018

UF UNIVERSITY of FLORIDA
Center for Compressible Multiphase Turbulence

Agenda for 6th Nek5000 Users/Developers Meeting

Tuesday April 17, 2018

1	9:00	Balachandar
2	9:30	Zwick
3	9:50	Zhai
4	10:10	Mavriplis
	10:30	BREAK
5	11:00	Schlatter
6	11:30	Peplinski
7	11:50	Offermans
	12:10	LUNCH (Provided)
8	13:30	Merzari
9	14:00	Tomboulides
10	14:30	Shaver
11	14:45	Kaneko
	15:00	BREAK
12	15:30	Hackl
13	15:50	Peet
14	16:10	Min
15	16:30	Ratnayaka
	16:45	DISCUSSION
	17:15	Break for Dinner (Provided)

Wednesday April 18, 2018

16	9:00	Fischer
17	9:20	Mittal
18	9:50	Lu
19	10:10	Dutta
	10:30	BREAK
20	11:00	Obabko

21	11:20	Carisik
22	11:40	Yuan
23	12:00	Bello Maldonado
	12:15	LUNCH (Provided)
24	13:30	Kerkemeier: Discussion & Future Plans
	15:00	Meeting adjourns

TITLES:

1.	Siva Balachandar	UFL	Nek5000 + multiphase
2.	David Zwick	UFL	Scalable discrete element method
3.	Keke Zhai	UFL	Dynamic load balancing for CMT-nek
4.	Catherine Mavriplis	Ottowa	Adaptive DG
5.	Philipp Schlatter	KTH	ExaFLOW: PGAS-based gather-scatter
6.	Adam Peplinski	KTH	Nonconforming: stability and implementation
7.	Nicolas Offermans	KTH	Adjoint error estimators
8.	Elia Merzari	ANL	Nek5000 in reactor problems
9.	Ananias Tomboulides	AUTH/ANL	Two phase and RANS models
10.	Dillon Shaver	ANL	Boiling using two-fluid two-phase
11.	Kento Kaneko	UIUC	Heat transfer simulations; Plan 5
12.	Jason Hackl	UFL	DG-based shock capturing
13.	Yulia Peet	ASU	FSI and validation
14.	Misun Min	ANL	Drift diffusion and implicit models
15.	Thilina Ratnayaka	UIUC	Nek5000 & libCEED
16.	Paul Fischer	UIUC/ANL	Nek5000 / HPC / GPUs
17.	Ketan Mittal	UIUC	Mesh smoothing & Neknek extensions
18.	Li Lu	UIUC	Nonlinear regularization for the SEM
19.	Som Dutta	UIUC	Fast semi-implicit particle tracking
20.	Aleksandr Obabko	ANL	Million-rank simulations
21.	Lane Carisik	Kairos	Nek5000 in industry
22.	Haomin Yuan	ANL	Flow-induced vibration
23.	Pedro Bello-Maldonado	UIUC	Low-Order and FAS preconditioners
24.	Stefan Kerkemeier	self	Where we are heading

10. Publications

2014

1. Annamalai, S., Neal, C., Ouellet, F., Rollin, B., Jackson, T.L. & Balachandar, S. (2014). Numerical Simulation of Explosive Dispersal of Particles in Cylindrical Geometry. Slides available online on the IWPCTM 2014 website - https://iwpctm.llnl.gov/index.html.
2. Annamalai, S., Parmar, M., Mehta, Y., & Balachandar, S. (2014). Modeling of hydrodynamic forces on a finite-sized spherical particle due to a planar shock wave. Bulletin of the American Physical Society, 59.
3. Akiki, G., & Balachandar, S. (2014). Immersed Boundary Methods on Non-Uniform Grids for Simulation of a Fully Resolved Bed of Particles in a Near-Wall Turbulent Flow. Bulletin of the American Physical Society, 59.
4. Rollin, B., Annamalai, S., Neal, C., Jackson, T., & Balachandar, S. (2014). Numerical Study of Explosive Dispersal of Particles. Bulletin of the American Physical Society, 59.
5. Jackson, T., Sridharan, P., Zhang, J., & Balachandar, S. (2014). Shock propagation over a deformable particle. Bulletin of the American Physical Society, 59.
6. Annamalai, S., Parmar, M. K., Ling, Y., & Balachandar, S. (2014). Nonlinear Rayleigh–Taylor Instability of a Cylindrical Interface in Explosion Flows. Journal of Fluids Engineering, 136(6), 060910.
7. Annamalai, S., Balachandar, S., & Parmar, M. K. (2014). Mean force on a finite-sized spherical particle due to an acoustic field in a viscous compressible medium. Physical Review E, 89(5), 053008.
8. Mankbadi, M. R., & Balachandar, S. (2014). Multiphase effects on spherical Rayleigh-Taylor interfacial instability. Physics of Fluids (1994-present), 26(2), 023301.
9. Kumar, N., Pascoe, C., Rudolph, D., Lam, H., George, A., and Stitt, G. (2014). Multi-scale, Multi-objective, Behavioral Modeling & Emulation of Extreme-scale Systems. Workshop on Modeling & Simulation of Systems & Applications, Seattle, WA, August 13-14, 2014.
10. Zunino, H., Adrian, R.J., Ding, L., and Prestridge, K. (2014). New in-situ, non-intrusive calibration. 67th Annual Meeting of the APS Division of Fluid Dynamics, San Francisco, CA, November 23-25, 2014.
11. Adrian, R.J., Wu, X., Moin, P., Baltzer, J.R. (2014). Osborne Reynolds pipe flow: direct numerical simulation from laminar to fully-developed turbulence. 67th Annual Meeting of the APS Division of Fluid Dynamics, San Francisco, CA, November 23-25, 2014.
12. Adrian, R. co-authored several presentations. These include: *Triple Pulse Particle Image Velocimeter/Accelerometer Measurements of Flow-Structure Interaction* (S. Gogineni), *Effect of Small Roughness Elements on Thermal Statistics of Turbulent Boundary Layer at Moderate Reynolds Number* (A. Doosttalab), *Multi-Scale Coherent Structure Interactions in Rayleigh-Benard Convection* (P. Sakievich), and *New in-situ, non-intrusive calibration* (H.A. Zunino). *Optimization and Application of Surface Segmentation Technique for Tomographic PIV* (L. Ding). 2014 67th Annual Meeting of APS Division of Fluid Dynamics. Bulletin of the American Physical Society, 59.
13. Chen Q., R. J. Adrian, Q. Zhong, D. Li, X. Wang (2014). "Experimental study on the role of spanwise vorticity and vortex filaments in the outer region of open-channel flow", J. Hydraulic Res., 1-14.

14. Matsumura, Y. and Jackson, T.L. (2014). Numerical simulation of fluid flow through random packs of cylinders using immersed boundary method. Physics of Fluids, Vol. 26, 043602.
15. Matsumura, Y. and Jackson, T.L. (2014). Numerical simulation of fluid flow through random packs of polydisperse cylinders. Physics of Fluids, Vol. 26, 123302.
16. Anderson, M.J., Jackson, T.L., Wasistho, B., and Buckmaster, J. (2014). A physics-based hot-spot model for pore collapse in HMX. 15th International Detonation Symposium, San Francisco, CA, July 13-18, pp. 951-961.
17. Anderson, M.J., Jackson, T.L., Wasistho, B., and Buckmaster, J. (2014). A physics-based hot-spot model for pore collapse in HMX. 46th JANNAF Combustion Subcommittee Meeting, Albuquerque, NM, December 8-11, 2014.
18. Chen Q., R. J. Adrian, Q. Zhong, D. Li, X. Wang (2014). "Experimental study on the role of spanwise vorticity and vortex filaments in the outer region of open-channel flow", J. Hydraulic Res., 1-14.
19. Hengxing Tan and Sanjay Ranka (2014). Thermal-aware Scheduling for Data Parallel Workloads on Multi-Core Processors, Proceedings of ISCC 2014.

2015

20. Sridharan, P., Jackson, T.L., Zhang, J. and Balachandar, S. (2015). Shock interaction with one-dimensional array of particles in air. Journal of Applied Physics, Vol. 117, 075902.
21. Thakur, S., Neal, C., Mehta, Y., Sridharan, P., Jackson, T.L. and Balachandar, S. (2015). Microscale Simulations of Shock Interaction with Large Assembly of Particles for Developing Point-Particle Models. SHOCK15 Meeting, American Physical Society.
22. Zhang, J., Jackson, T.L., Sridharan, P. and Balachandar, S. (2015). Towards a mass and volume conserving interface reinitialization scheme for a diffuse interface methodology (for shock-particle interaction). AIP Conf. Proc. 1793, 150005 (4 pages).
23. G. Akiki, T.L. Jackson, S. Balachandar (2015). Mean and Fluctuating Force Distribution in a Random Array of Spheres. Bulletin of the American Physical Society, Vol. 60.
24. T.L. Jackson, P. Sridharan, J. Zhang, S. Balachandar (2015). Numerical Simulation of Shock Interaction with Deformable Particles Using a Constrained Interface Reinitialization Scheme. Bulletin of the American Physical Society, Vol. 60.
25. Diggs, A., Balachandar, S. (2015). Modeling and Simulation Challenges in Eulerian-Lagrangian Computations of Shock-Driven Multiphase Flows, Bulletin of the American Physical Society, Vol. 60.
26. McGrath, T., St. Clair, J. and Balachandar, S. (2015). An extended pressure equilibrium model for multiphase flows – application to shock-induced particle dispersion. APS Shock Compression of Condensed Matter.
27. Zhang, J., Jackson, T.L. and Balachandar, S. (2015). Numerical simulation of shock/detonation deformable particle interaction with constrained interface reinitialization. APS Shock Compression of Condensed Matter.
28. Annamalai, A. and Balachandar, S. (2015). Mean force on a finite-sized rigid particle, droplet, or bubble in a viscous compressible medium. Physics of Fluids, Vol. 27(10), 103304.

29. Annamalai, S., Balachandar, S. and Mehta, Y. (2015). Analytic expressions for first order correction to inviscid unsteady forces due to surrounding particles in a multiphase flow. APS DFD abstracts.
30. Matsumura, Y., Jenne, D. and Jackson, T.L. (2015). Numerical simulation of fluid flow through random packs of ellipses. Physics of Fluids, Vol. 27, 023301.
31. Jackson, T.L., Buckmaster, J., Zhang, J. and Anderson, M. (2015). Pore collapse in an energetic material from the micro-scale to the macro-scale. Combustion Theory and Modeling, Vol. 19(3), pp. 347-381.
32. Gillman, A., Amadio, G., Matous, K. and Jackson, T.L. (2015). Third-order thermos-mechanical properties for packs of Platonic solids using statistical micromechanics. Proceedings of the Royal Society of London A, Vol. 471, 20150060.
33. Sakievich, P.J., Peet Y.T., Adrian, R.J. (2015). Large-scale, coherent structures in wide-aspect ratio, turbulent, Rayleigh-Benard convection. Ninth International Symposium on Turbulence and Shear Flow Phenomena. July 2015.
34. Wu, X., Moin. P, Adrian, R.J., Baltzer, J.R. (2015). Osborne Reynolds pipe flow: Direct simulation from laminar through gradual transition to fully developed turbulence. Proc Natl Acad Sci U.S.A. 2015 Jun 30; 112(26):7920-4. DOI: 10.1073/pnas.1509451112. Epub 2015 Jun 15.
35. Chojnicki K. N., A. B. Clarke, J.C. Phillips, R. J. Adrian (2015). The evolution of volcanic plume morphologies from short-lived eruptions. Geology, v. 43 no. 8 p. 707-710. July 10, 2015, DOI:10.1130/G36642.1.
36. J. Zhang, T.L. Jackson (2015). Detonation initiation with thermal deposition due to pore collapse in energetic materials - Towards the coupling between micro- and macroscales. Bulletin of the American Physical Society, Vol. 60.
37. Zhang, J. and Jackson, T.L. (2015). Detonation Initiation with Thermal Deposition due to Pore Collapse in Energetic Materials - Towards the Coupling between Micro- and Macroscale. AIAA Paper No. 2015-4097, 51st AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit, 27-29 July 2015, Orlando, FL.
38. Amadio, G. and Jackson, T.L. (2015). A new packing code for creating microstructures of propellants and explosives. AIAA Paper No. 2015-4098, 51st AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit, 27-29 July 2015, Orlando, FL.
39. Banerjee, T., Ranka, S., (2015). A Genetic Algorithm Based Autotuning Approach for Performance and Energy Optimization, The 6th International Green and Sustainable Computing Conference, 2015.
40. N. Kumar, M. Shringarpure, T. Banerjee, J. Hackl, S. Balachandar, H. Lam, A. George, S. Ranka, (2015). CMT-bone: A mini-app for Compressible Multiphase Turbulence Simulation Software. Workshop on Representative Applications, co-located with IEEE Cluster 2015, Chicago, IL, USA, Sep 8-11 2015.
41. Dylan Rudolph and Greg Stitt (2015). An Interpolation-Based Approach to Multi-Parameter Performance Modeling for Heterogeneous Systems, IEEE International Conference on Application-specific Systems, Architectures and Processors (ASAP), July 2015.
42. N. Kumar, C. Hajas, A. George, H. Lam, G. Stitt (2015). Multi-scale, “Multi-objective Behavioral Emulation of Future-gen Applications and Systems”, 2015 Salishan Conference on High-Speed Computing, April 27-30, 2015, Gleneden Beach, Oregon.

43. N. Kumar, A. George, H. Lam, G. Stitt, S. Hammond (2015). "Understanding Performance and Reliability Trade-offs for Extreme-scale Systems using Behavioral Emulation", Workshop on Modeling & Simulation of Systems and Applications (ModSim 2015), August 12-14, 2015, Seattle, Washington.
44. N. Kumar, M. Shringarpure, T. Banerjee, J. Hackl, S. Balachandar, H. Lam, A. George, S. Ranka (2015). CMT-bone: A mini-app for Compressible Multiphase Turbulence Simulation Software. Workshop on Representative Applications, co-located with IEEE Cluster 2015, Chicago, IL, USA, Sep 8-11 2015.

2016

45. Annamalai, S., Rollin, B., Ouellet, F., Neal, C., Jackson, T.L. and Balachandar, S. (2016). Effects of initial perturbations in the early moments of an explosive dispersal of particles. ASME Journal of Fluids Engineering, Vol. 138, 070903 (9 pages).
46. Sridharan, P., Jackson, T.L., Zhang, J., Balachandar, S., and S. Thakur (2016). Shock interaction with deformable particles using a constrained interface reinitialization scheme. Journal of Applied Physics, Vol. 119, 064904, 18 pages.
47. Akiki, G., and Balachandar, S. (2016). Immersed boundary method with non-uniform distribution of Lagrangian markers for a non-uniform Eulerian mesh. Journal of Computational Physics. Vol. 307, pp. 34-59. DOI 10.1016/j.jcp.2015.11.019.
48. Mehta, Y., Jackson, T. L., Zhang, J., and Balachandar, S., (2016). "Numerical investigation of shock interaction with one-dimensional transverse array of particles in air". Journal of Applied Physics 119(10), p. 104901.
49. Mehta, Y., Neal, C., Jackson, T. L., Balachandar, S., and Thakur, S., (2016). "Shock interaction with three-dimensional face centered cubic array of particles". Phys. Rev. Fluids, Vol. 1, 054202 (27 pages).
50. Akiki, G., Jackson, T.L., and Balachandar, S. (2016). Force variation within arrays of mono-disperse spherical particles. Physical Review Fluids, Vol. 1, 044202 (29 pages).
51. Ling, Y., Balachandar, S. and Parmar, M. (2016). Interphase heat transfer and energy coupling in turbulent dispersed multiphase flows. Physics of Fluids, Vol. 28(3), 033304.
52. McGrath, T., St. Clair, J. and Balachandar, S. (2016). A compressible two-phase model for dispersed particle flows with application from dense to dilute regimes. Journal of Applied Physics, Vol. 119(17), 174903.
53. Diggs, A. and Balachandar, S. (2016). Evaluation of methods for calculating volume fraction in Eulerian-Lagrangian multiphase flow simulations. Journal of Computational Physics, Vol.313, pp. 775-798.
54. Schwarzkopf, J.D., Balachandar, S. and Butler, W.T. (2016). Compressible multiphase flow. Multiphase Flow Handbook, Vol. 455.
55. Koneru, R., Rollin, B., Ouellet, F., Annamalai, S. and Balachandar, S. (2016). Simulations of shock wave interaction with a particle cloud. APS DFD abstracts.
56. Neal, C., Mehta, Y., Salari, K., Jackson, T.L. and Balachandar, S. (2016). Shock interaction with random spherical particle beds. APS DFD abstracts.
57. Akiki, G., Jackson, T.L., Balachandar, S. (2016). Pairwise interaction extended point particle (PIEP) model for a random array of spheres. APS DFD abstracts.
58. Mehta, Y., Neal, C., Jackson, T.L. and Balachandar, S. (2016). Shock particle interaction – fully resolved simulations and modeling. APS DFD abstracts.

59. Ouellet, F., Park C., Rollin, B. and Balachandar, S. (2016). A multi-fidelity surrogate model for handling real gas equations of state. APS DFD abstracts.
60. Akiki, G., Jackson, T.L. and Balachandar, S. (2016). Quantifying and modeling the force variation within random arrays of spheres. ASME 2016 IMECE abstracts.
61. Zhang, J. and Jackson, T.L. (2016). Direct detonation initiation with thermal deposition due to pore collapse in energetic materials – Towards the coupling between micro- and macroscale. Combustion Theory and Modelling, 1218053.
62. C. Park, & N. H. Kim, (2016). Safety envelope for load tolerance of structural element design based on multi-stage testing. *Advances in Mechanical Engineering*, *8*(9).
63. C. Park, R. T. Haftka, & N. H. Kim, (2016). Remarks on multi-fidelity surrogates. *Structural and Multidisciplinary Optimization*, 1-22.
64. Y. Zhang, C. Park, N. H. Kim., R. T. Haftka (2016). "Function Prediction at One Inaccessible Point using Converging Lines", Journal of Mechanical Design, accepted for publication in Jan 2017.
65. Doosttalab A., G. Araya, J. Newman, R. J. Adrian, K. Jansen, L. Castillo (2016). "Effect of small roughness elements on thermal statistics of a turbulent boundary layer at moderate Reynolds number" J. Fluid Mech. 787, 84-115.
66. Doosttalab A., S. Dharmarathne, M. Tutkun M., R. J. Adrian, L. Castillo (2016). "Analysis of velocity structures in a transitionally rough turbulent boundary layer" Chapter 5 in A. Pollard et al. (eds.), Whither Turbulence and Big Data in the 21st Century, DOI 10.1007/978-3-319-41217-7_5.
67. Ding L, R J Adrian (2016). N-pulse Particle Image Velocimetry-accelerometry for Unsteady Flow-structure Interaction" Measurement Science and Technology 28(1).
68. Sakievich, P., Y. Peet and R. J. Adrian (2016). "Large-Scale Thermal Motions of Turbulent Rayleigh-Bénard Convection in a Wide Aspect-Ratio Cylindrical Domain", Int'l. J. Heat and Fluid Flow, http://dx.doi.org/10.1016/j.ijheatfluidflow.2016.04.011.
69. Chaudhury, R, V. Atlasman, G. Pathangey, N. Pracht, R. J. Adrian and D. H. Frakes (2016). "A high performance pulsatile pump for aortic flow experiments in 3-dimensional models", Cardiovascular Engineering & Technology 7(2), DOI:10.1007/s13239-016-0260-3.
70. Chaudhury, R, R. J. Adrian and D. H. Frakes (2016). "Prediction of flow and velocity waveforms in in vitro cardiovascular flow experiments", IEEE Trans. on Biomed. Engr., Vol. 63.
71. T. Banerjee, J. Hackl, M. Shringarpure, T. Islam, S. Balachandar, T. Jackson and S. Ranka (2016). "CMT-bone – A Proxy Application for Compressible Multiphase Turbulent Flows", HiPC.
72. M. Gadou, T. Banerjee and S. Ranka (2016). "Multiobjective Optimization of CMT-bone on Hybrid Processors", IGSC.
73. N. Kumar, C. Pascoe, C. Hajas, H. Lam, G. Stitt, and A. George (2016). "Behavioral Emulation for Scalable Design-Space Exploration of Algorithms and Architectures", 2016 Workshop on Exascale Multi/Many Core Computing Systems (E-MuCoCoS), June 23, 2016, Frankfurt, Germany.
74. C. Pascoe, N. Kumar, K. Alli, H. Lam, G. Stitt, and A. George (2016). "FPGA-Pipelined Discrete-Event Simulations for Accelerated Behavioral Emulation of Extreme-Scale Systems", Workshop on Modeling & Simulation of Systems and Applications (ModSim 2016), August 10-11, 2016, Seattle, Washington.

2017

75. Annamalai, S., Balachandar, S. Sridharan, P. and Jackson, T.L. (2017). Pressure evolution equation for the particulate phase in inhomogeneous compressible disperse multiphase flows. Physical Review Fluids, Vol. 2, 024301 (23 pages).
76. Akiki, G., Jackson, T.L., and Balachandar, S. (2017). Pairwise interaction extended point-particle model for a random array of monodisperse spheres. Journal of Fluid Mechanics, Vol. 813, pp. 882-928.
77. P. Crittenden & S. Balachandar (2017). The impact of the form of the Euler equations for radial flow in cylindrical and spherical coordinates on numerical conservation and accuracy, submitted to Shockwaves (2017).
78. F. Ouellet, C. Park, B. Rollin, and S. Balachandar, S. (2017). Analysis of a Multi-Fidelity Surrogate for Handling Real Gas Equations of State, Proceedings of the 20th Biennial APS Conference on Shock Compression of Condensed Matter, St-Louis, MO.
79. Moore, W.C., Akiki, G., Balachandar, S. (2017). A Hybrid Point-Particle Force Model That Combines Physical and Data-Driven Approaches, to be submitted.
80. Akiki, G., Moore, W. C., Balachandar, S. (2017). "Pairwise-interaction extended point-particle model for particle-laden flows," Journal of Computational Physics, Vol. 351, pp. 329-357.
81. Hughes, K., Diggs, A., Littrell, D., Balachandar, S., Haftka, R., Kim, N., Park, C. (2017). Uncertainty quantification of experiments on a small number of explosively driven particles. 55th AIAA Aerospace Sciences Meeting, 1463.
82. Nili, S., Park, C., Haftka, R., Balachandar, S. and Kim, N. (2017). Sensitivity analysis of force models for a four-way coupled Eulerian-Lagrangian dispersed multiphase flow. 23rd AIAA Computational Fluid Dynamics Conference, 3800.
83. Thakur, S., Neal, C., Mehta, Y., Sridharan, Jackson, T.L., and Balachandar, S. (2017). Microscale simulations of shock interaction with large assemble of particles for developing point-particle models. AIP Conference Proceedings 1793(1), 150007.
84. Diggs, A. and Balachandar, S. (2017). Modeling and simulation challenges in Eulerian-Lagrangian computations of multiphase flows, AIP Conference Proceedings, 1793(1), 150008.
85. Zhang, J., Jackson, T.L. and Balachandar, S. (2017). Towards a mass and volume conserving interface reinitialization scheme for a diffuse interface methodology (for shock particle interaction), AIP Conference Proceedings, 1793(1), 150005.
86. Annamalai, S. and Balachandar, S. (2017). Faxen form of time-domain force on a sphere in unsteady spatially varying viscous compressible flows. Journal of Fluid Mechanics. Vol.816, pp. 381-411.
87. McGrath, T., St. Clair, J. and Balachandar, S. (2017). Modeling compressible multiphase flows with dispersed particles in both dense and dilute regimes, Shock Waves, 1-12.
88. Zwick, D., Sakhaee, E., Balachandar, S. and Entezari, A. (2017). Accurate signal reconstruction for higher order Lagrangian-Eulerian back-coupling in multiphase turbulence, Fluid Dynamic Research, Vol. 49(5), 055507.
89. Aref, H. and Balachandar, S. (2017). A first course in computational fluid dynamics, Cambridge University Press.

90. Nili, S., Park, C., Haftka, R., Kim, N. and Balachandar, S. (2017). Effect of Finite Particle Size on Convergence of Point Particle Models in Euler-Lagrange Multiphase Dispersed Flow, Bulletin of the American Physical Society 62.
91. Zwick, D., Hackl, J. and Balachandar, S. (2017). Scalable Methods for Eulerian-Lagrangian Simulation Applied to Compressible Multiphase Flows, Bulletin of the American Physical Society 62.
92. Garno, J., Ouellet, F., Koneru, R., Balachandar, S. and Rollin, B. (2017). Predictive Capability of the Compressible MRG Equation for an Explosively Driven Particle with Validation, Bulletin of the American Physical Society 62.
93. Marjanovic, G., Hackl, J., Annamalai, S., Jackson, T.L. and Balachandar, S. (2017). Fully resolved simulations of expansion waves propagating into particle beds, Bulletin of the American Physical Society 62.
94. Ouellet, F., Park, C., Koneru, R., Balachandar, S. and Rollin, B. (2017). A Multi-Fidelity Surrogate Model for the Equation of State for Mixtures of Real Gases, Bulletin of the American Physical Society 62.
95. Koneru, R., Rollin, B., Ouellet, F., Park, C. and Balachandar, S. (2017). Euler-Lagrange Simulations of Shock Wave-Particle Cloud Interaction, Bulletin of the American Physical Society 62.
96. Moore, C., Akiki, G. and Balachandar, S. (2017). A Hybrid Physics-Based Data-Driven Approach for Point-Particle Force Modeling, Bulletin of the American Physical Society 62.
97. Liu, K. and Balachandar, S. (2017). Pairwise Interaction Extended Point-Particle (PIEP) model for multiphase jets and sedimenting particles, Bulletin of the American Physical Society 62.
98. Mehta, Y., Salari, K., Jackson, T.L., Balachandar, S. and Thakur, S. (2017). Strong Shock Propagating Over a Random Bed of Spherical Particles, Bulletin of the American Physical Society 62.
99. Durant, B., Hackl, J. and Balachandar, S. (2017). Long-time stability effects of quadrature and artificial viscosity on nodal discontinuous Galerkin methods for gas dynamics, Bulletin of the American Physical Society 62.
100. Hackl, J., Shringarpure, M., Fischer, P. and Balachandar, S. (2017). Shock capturing in discontinuous Galerkin spectral elements via the entropy viscosity method, Bulletin of the American Physical Society 62.
101. Osborne, B., Jackson, T.L. and Balachandar, S. (2017). Density Discontinuity Interaction with a Structured Array of Particles, Bulletin of the American Physical Society 62.
102. Zhang, J. and Jackson, T.L. (2017). Direct detonation initiation with thermal deposition due to pore collapse in energetic materials – Towards the coupling between micro- and macroscale. Combustion Theory and Modelling, Vol. 21(2), pp. 248-273; 1218053.
103. Jackson, T.L., and Zhang, J. (2017). Density-based kinetics for mesoscale simulations of detonation initiation in energetic materials. Combustion Theory Modelling, Vol. 21(4), pp. 749-769; 1296975.
104. Zhang, J., Jackson, T.L., and Jost, A.M.D. (2017). Effects of Air Chemistry and Stiffened EOS of Air in Numerical Simulations of Bubble Collapse in Water. Physical Review Fluids, Vol. 2, 053603.

105. Jackson, T.L., Antoine, A.M.D., Zhang, J., Sridharan, P. and Amadio, G. (2017). Multi-dimensional mesoscale simulations of detonation initiation in energetic materials with density-based kinetics. Combustion Theory Modelling, 1401121 (25 pages).
106. Jackson, T.L. and Zhang, J. (2017). Mesoscale simulations of energetic materials using density-based kinetics. SIAM Sixteenth International Conference on Numerical Combustion, April 3-5, Orlando, Florida.
107. B.T. Bojko, M.L. Gross, and T.L. Jackson (2017). Investigating coupled micro and mesoscale combustion models in 3D numerical simulations of heterogeneous propellants. JANNAF Paper, presented at the 2017 48th Combustion JANNAF Joint Meeting, Newport News, VA, December 4-7.
108. M. Gadou, T. Banerjee and S. Ranka, "Multiobjective evaluation and optimization of CMT-bone on multiple CPU/GPU systems", in preparation (2017).
109. Ouellet, F., Annamalai, S. and Rollin, B., (2017). January. Effect of a bimodal initial particle volume fraction perturbation in an explosive dispersal of particles. In AIP Conference Proceedings (Vol. 1793, No. 1, p. 150011). AIP Publishing.
110. Fernández-Godino, M. G., Park, C., Kim, N. H., & Haftka, R. T. (2016). Review of multi-fidelity models. *arXiv preprint arXiv:1609.07196*.
111. H. Zunino, R.J. Adrian, A.B. Clarke, B.A. Johnson (2017). "Exploring the Early Structure of a Rapidly Decompressed Particle Bed," American Physical Society Division of Fluid Dynamics Meeting, Denver, CO, November 2017.
112. B.A. Johnson, H. Zunino, R.J. Adrian, A.B. Clarke (2017). "Gas and particle motions in a rapidly decompressed flow," American Physical Society Division of Fluid Dynamics Meeting, Denver, CO, November 2017.
113. R.J. Adrian, X. Wu, P. Moin (2017). "Turbulent spots and scalar flashes in pipe transition," American Physical Society Division of Fluid Dynamics Meeting, Denver, CO, November 2017.
114. C. Park, R. T. Haftka, and N. H. Kim (2017). Low-Fidelity Scale Factor Improves Bayesian Multi-Fidelity Prediction by Reducing Bumpiness of Discrepancy Function, submitted to *Structural Multidisciplinary Optimization*.
115. N. Qui, C. Park, Y. Gao, J. Fang, G. Sun, and N. H. Kim (2017). Sensitivity-Based Parameter Calibration and Model Validation under Model Error, *ASME Journal of Mechanical Design*, 140(1), 011403.
116. Bae, S., Kim, N. H., Park, C. and Kim, Z. (2017). Confidence Interval of Bayesian Network and Global Sensitivity Analysis. *AIAA Journal*, 55(11), pp. 3916-3924.
117. Zhang, Y., Kim, N.-H., Park, C., and Haftka, R. T. (2017). "Multi-Fidelity Surrogate Based on Single Linear Regression," arXiv preprint arXiv:1705.02956.
118. Tania Banerjee, Jason Hackl, Mrugesh Shringarpure, Tanzima Islam, S. Balachandar, Thomas Jackson, Sanjay Ranka (2017). "A New Proxy Application for Compressible Multiphase Turbulent Flows", Elsevier Sustainable Computing: Informatics and Systems, Volume 16, 2017, Pages 11-24.
119. Mohamed Gadou, Tania Banerjee, Meena Arunachalam, Sanjay Ranka (2017). "Multiobjective evaluation and optimization of CMT-bone on multiple CPU/GPU systems", Journal of Sustainable Computing: Informatics and Systems, 2017.
120. Y. Zhang, A. Neelakantan, N. Kumar, C. Park, R. Haftka, N. H. Kim, and H. Lam (2017). "Multi-fidelity Surrogate Modeling for Application/Architecture Co-design", Performance

Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS17), Denver, CO, Nov. 13, 2017.

121. C. Pascoe, S. Chenna, G. Stitt, H. Lam (2017). "A FPGA-Pipelined Approach for Accelerated Discrete-Event Simulation of HPC Systems", Heterogeneous High-performance Reconfigurable Computing (H2RC), Denver, CO, Nov. 13, 2017.
122. Zhang Y., Neelakantan A. et al. Multi-fidelity Surrogate Modeling for Application/Architecture Co-design. In: Jarvis S., Wright S., Hammond S. (eds) High Performance Computing Systems. Performance Modeling, Benchmarking, and Simulation. PMBS 2017. Lecture Notes in Computer Science, vol 10724. Springer, Cham.

2018

123. Mehta, Y., Neal, C., Salari, K., Jackson, T., Balachandar, S., & Thakur, S. (2018). "Propagation of a strong shock over a random bed of spherical particles," Journal of Fluid Mechanics, 839, 157-197. doi:10.1017/jfm.2017.909.
124. Ling, S. and Balachandar, S. (2018). Simulation and scaling analysis of a spherical particle-laden blast wave, Shock Waves, 1-14.
125. Cook, C.R., Balachandar, S., Chung, J.N., and Vu-Quoc, L. (2018). A generalized characteristic-based split projection method for Navier-Stokes with real fluids. International Journal of Heat and Mass Transfer, Vol. 124, pp. 1045-1058.
126. Fernandez-Godino, M.G., Haftka, R., Balachandar, S. Gogu, C. and Bartoli, N. (2018). Noise Filtering and Uncertainty Quantification in Surrogate based Optimization, AIAA Non-Deterministic Approaches Conference, 2176.
127. Hughes, K., Diggs, Park, C., A., Littrell, D., Haftka, R., Kim, N., and Balachandar, S. (2018). Simulation-Driven Experiments of Macroscale Explosive Dispersal of Particles, AIAA Aerospace Sciences Meeting, 1545.
128. Annamalai, S., Balachandar, S.(2018). Towards Combined Deterministic and Statistical Approaches to Modeling Dispersed Multiphase Flows, Droplets and Sprays, 7-42.
129. F. Ouellet, C. Park, B. Rollin, R.T. Haftka and S. Balachandar (2018). "A Multi-Fidelity Surrogate Model for Handling Gas Mixture Equations of State", *Journal of Computational Physics*, in preparation for submission.
130. Johnson, B.A., Cowen, E.A. (2018). "Turbulent boundary layers absent mean shear" Journal of Fluid Mechanics 835, pp. 217 - 251.
131. A. Ramaswamy, N. Kumar, H. Lam, and G. Stitt (2018). "Scalable Behavioral Emulation of Extreme-Scale Systems Using Structural Simulation Toolkit", submitted to the International Parallel and Distributed Processing Symposium (IPDPS), Vancouver, CA. May 21-25, 2018.
132. Keke Zhai, Tania Banerjee, David Zwick, Jason Hackl, Sanjay Ranka (2018). "Dynamic Load Balancing for Compressible Multiphase Turbulence", acccepted to ICS 2018
133. Mohamed Gadou, Sankeerth Reddy, Tania Banerjee, Sanjay Ranka (2018). "Multi-objective Optimization on DVFS based Hybrid Systems", submitted to ICS 2018
134. Fernandez-Godino, M.G., Haftka, R.T., Balachandar, S., Gogu, C., Bartoli, N. and Dubreuil, S., "Noise Filtering and Uncertainty Quantification in Surrogate based Optimization", *2018 AIAA Non-Deterministic Approaches Conference*, p. 2176 (2018).

135. Yiming Zhang, Nam H. Kim, Chanyoung Park and Raphael T. Haftka, Multi-fidelity surrogate based on single linear regression, AIAA Journal, Vol. 56, No. 12, pp. 4944-4952, 2018. doi: 10.2514/1.J057299.
136. Chanyoung Park, Raphael T. Haftka, and Nam H. Kim, Low-fidelity scale factor improves Bayesian multi-fidelity prediction by reducing bumpiness of discrepancy function, Structural and Multidisciplinary Optimization, Vol. 58, No. 2, pp. 399-414, 2018. doi: 10.1007/s00158-018-2031-2
137. Marjanovic, G. M., Hackl, J. F., Shringarpure, M., Annamalai, S., Jackson, T.L. & Balachandar, S. (2018) Inviscid simulations of expansion waves propagating into face-centered cubic structured particle beds at low volume fractions. *Phys. Rev. Fluids* **3**, 094301.
138. Mohamed Gadou, Tania Banerjee, Meena Arunachalam, Sanjay Ranka, "Multiobjective evaluation and optimization of CMT-bone on Intel Knights Landing", accepted in IGSC, 2018.
139. Ajay Ramaswamy, Nalini Kumar, Aravind Neelakantan, Herman Lam, and Greg Stitt. 2018. Scalable Behavioral Emulation of Extreme-Scale Systems Using Structural Simulation Toolkit. In Proceedings of the 47th International Conference on Parallel Processing (ICPP 2018). ACM, New York, NY, USA, Article 17, 11 pages
140. Yiming Zhang, Aravind Neelakantan, Chanyoung Park, Nam H. Kim, Herman Lam, and Raphael T. Haftka. "Adaptive Sampling with Varying Sampling Cost for Design Space Exploration". AIAA Journal 0 0:0, 1-12 , https://doi.org/10.2514/1.J057470

2019

141. Zhang, J. and Jackson, T.L. (2019). Effect of microstructure on the detonation initiation in energetic materials. Shock Waves, Vol. 29(2), pp. 327-338. doi: 10.1007/s00193-017-0796-7.
142. Balachandar, S., Liu, K., and Lakhote, M. (2019). Self-induced velocity correction for improved drag estimation in Euler-Lagrange point-particle simulations. Journal of Computational Physics, Vol. 376, pp. 160-185.
143. Mehta, Y., Salari, K., Jackson, T.L. and Balachandar, S. (2019). Effect of Mach number and volume fraction in air-shock interacting with a bed of randomly distributed spherical particles. Physical Review Fluids, Vol. 4(1), 014303.
144. Bojko, B.T., Gross, M., and Jackson, T.L. (2019). Investigating dimensional effects on predicting burning rates of heterogeneous solid propellants. AIAA Scitech 2019 Forum, paper 1240.
145. Brown, K.C., Sankaran, V., and Jackson, T.L. (2019). Mesoscale modeling of solid propellant burn rates. AIAA Scitech 2019 Forum, paper 1236.
146. Fernandez-Godino, M. Giselle, S. Balachandar, and Raphael T. Haftka. "On the Use of Symmetries in Building Surrogate Models", Journal of Mechanical Design, 141(6), p.061402. (2019)
147. M.G. Fernandez-Godino, F. Ouellet, R.T. Haftka, S. Balachandar, "Early Time Evolution of Circumferential Perturbation of Initial Particle Volume Fraction in Explosive Cylindrical Multiphase Dispersion", Journal of Fluids Engineering, *under review* (2019).

148. F. Ouellet, C. Park, B. Rollin, R.T. Haftka and S. Balachandar, "A Kriging Surrogate Model for Computing Gas Mixture Equations of State", Journal of Fluids Engineering, accepted (2019).
149. Moore, W. C., Balachandar, S. (2019). "Lagrangian Modeling of Pseudo-Turbulence in Multiphase Flow Using Superposable Wakes," In progress.
150. Moore, W. C., Akiki, G., Balachandar, S. (2019). "A Hybrid Point-Particle Force Model That Combines Physical and Data-Driven Approaches, "Journal of Computational Physics, Accepted.
151. Chanyoung Park, Justin Matthew, Nam H. Kim, and Raphael T. Haftka, Epistemic uncertainty stemming from measurement processing - A case study of multiphase shock tube experiments, Journal of Verification, Validation and Uncertainty Quantification, accepted, 2019.
152. M. Giselle Fernández-Godino, Chanyoung Park, Nam H. Kim and Raphael T. Haftka, Review of multi-fidelity surrogates; When are they Worthwhile?, AIAA Journal, accepted, 2019.
153. Yiming Zhang, Aravind Neelakantan, Chanyoung Park, Nam H. Kim, Herman Lam and Raphael T. Haftka, Adaptive sampling with varying sampling cost for design space exploration, AIAA Journal, in press, 2019.
154. Hackl, J. F., Shringarpure, M. S., Koneru, R. B., Delchini, M.-O. G. & Balachandar, S. (2019). A shock capturing discontinuous Galerkin spectral element method for curved geometry using entropy viscosity. *Computers & Fluids*, under review.
155. Mohamed Gadou, Sankeerth Reddy, Tania Banerjee, Sanjay Ranka, "Multi-objective Optimization on DVFS based Hybrid Systems", under preparation for submission to iWAPT, 2019.
156. Keke Zhai, Tania Banerjee, David Zwick, Jason Hackl, Rahule Koneru, Sanjay Ranka, "Dynamic Load Balancing for A Mesh-based Scientific Application", submitted to Journal of Supercomputing, 2019.

11. Conferences and Presentations

2014

1. S. Balachandar (2014). "Center for Compressible Multiphase Turbulence – Overview", PSAAP-II Kickoff Meeting, Albuquerque, NM, December 9-11, 2013.
2. Sanjay Ranka, Herman Lam (2014). CCMT – Extreme Scale CS Research, PSAAP-II Kickoff Meeting, Albuquerque, NM, December 9-11, 2013.
3. S. Balachandar (2014). "Center for Compressible Multiphase Turbulence – Overview", ASC PI Meeting, Kansas City, MO, February 26, 2014.
4. S. Balachandar (2014). "Center for Compressible Multiphase Turbulence – Overview and CMT Physics", CCMT Road trip to Sandia National Laboratories, Albuquerque, NM, March 24, 2014.
5. R.T. Haftka (2014). "Center for Compressible Multiphase Turbulence – V&V and Uncertainty Quantification Research and Innovation", CCMT Road trip to Sandia National Laboratories, Albuquerque, NM, March 24, 2014.

6. Herman Lam (2014). "Center for Compressible Multiphase Turbulence – Exascale Emulation", CCMT Road trip to Sandia National Laboratories, Albuquerque, NM, March 24, 2014.
7. S. Balachandar (2014). "Center for Compressible Multiphase Turbulence – Overview and CMT Physics", CCMT Road trip to Los Alamos National Laboratory, Los Alamos, NM, March 25, 2014.
8. R.T. Haftka (2014). "Center for Compressible Multiphase Turbulence – V&V and Uncertainty Quantification Research and Innovation", CCMT Road trip to Los Alamos National Laboratory, Los Alamos, NM, March 25, 2014.
9. Herman Lam (2014). "Center for Compressible Multiphase Turbulence – Exascale Emulation", CCMT Road trip to Los Alamos National Laboratory, Los Alamos, NM, March 25, 2014.
10. S. Balachandar (2014). "Fundamental Advances in Compressible Multiphase Flows – Potential Relevance to Liquid Atomization and Spray Systems", Keynote Lecture, Institute for Liquid Atomization and Spray Systems, Portland, OR, May 18-21, 2014.
11. S. Balachandar (2014). "The Need for Fundamental Advances in Compressible Multiphase Flows – Shock-particle Interaction to Explosive Dispersal", Department of Aerospace Engineering, Iowa State University, Ames, Iowa, March 6, 2014.
12. S. Balachandar (2014). "Shock-Particle Interaction to Explosive Particle Dispersal – What Fundamental Advances in Compressible Multiphase Flows are Needed", University of Stuttgart, Stuttgart, Germany, May 7, 2014.
13. Subramanian Annamalai (2014). University of Florida, "Rocflu – An Overview", CCMT Seminar, 2:00 P.M., Wednesday, February 5, 2014, 210 MAE-B.
14. Jackson, T.L. (2014). 67th Annual Meeting, Division of Fluid Dynamics, American Physical Society, San Francisco, CA, Nov 2014. ``Shock propagation over a deformable particle".
15. S. Annamalai, C. Neal, F. Ouellet, B. Rollin, T.J. Jackson, & S. Balachandar (2014). Numerical Simulation of Explosive Dispersal of Particles in Cylindrical Geometry, IWPCTM 2014, San Francisco, California, USA.
16. Heather Zunino (2014). H. Zunino, R.J. Adrian, L. Ding, K. Prestridge, "New in-situ, non-intrusive calibration", 67th Annual Meeting of the APS Division of Fluid Dynamics, San Francisco, CA, November 23-25, 2014.
17. Ronald Adrian: R.J. Adrian, X. Wu , P. Moin , J.R. Baltzer. (2014). Osborne Reynolds pipe flow: direct numerical simulation from laminar to fully-developed turbulence, 67th Annual Meeting of the APS Division of Fluid Dynamics, San Francisco, CA, November 23-25, 2014.
18. B. Rollin, S. Annamalai, C. Neal, T.J. Jackson, & S. Balachandar (2014). Numerical Study of Explosive Dispersal of Particles*;* DFD2014, Bulletin of the American Physical Society, Vol. 59, Number 20 (November 2014).
19. Chanyoung Park, Joo-Ho Choi, Raphael T. Haftka (2014). Teaching a Verification and Validation Course using Simulations and Experiments with Paper Helicopter, *ASME 2014 Verification and Validation Symposium, Las Vegas, NV, May 7-9, 2014*
20. Chanyoung Park (2014). University of Florida, "Dakota – An Overview", CCMT Seminar, 3:00 P.M., Tuesday, April 8, 2014, 221 MAE-A.

21. H. Lam (2014). Behavioral Modeling & Emulation of Extreme-scale Systems, Workshop on Modeling & Simulation of Systems & Applications, Seattle, WA, August 13-14, 2014.
22. Invited Talk. Prof. Nam-Ho Kim (2014). Prioritizing Efforts to Reduce Uncertainty in System Models based on Uncertainty Budget, *Seoul National University, Seoul, Korea, July 7, 2014.*

2015

23. A.M.D. Jost and J. Zhang (2015). Numerical Study of Intermittent Laminar Bubble Bursting and Vortex Shedding on an NACA 643-618 Airfoil, AIAA Meeting, 2015.
24. B. Rollin, S. Annamalai, F. Ouellet (2015). A Study of Interfacial Instability in Explosive Dispersal of Particles, 19th Biennial APS Conference on Shock Compression of Condensed Matter, Tampa, FL (2015).
25. B. Rollin, F. Ouellet, S. Annamalai, & S. Balachandar (2015). Numerical Study of Explosive Dispersal of Particles, DFD2015, Bulletin of the American Physical Society, Vol. 60, Number 21 (November 2015).
26. B. Rollin (2015). *Center for Compressible Multiphase Turbulence,* Stewardship Science Academic Programs (SSAP) Symposium, Santa Fe, NM.
27. B. Rollin (2015). Progress on Simulation of Explosive Dispersal of Particles, LANL, Los Alamos, NM.
28. B. Rollin (2015). Toward Predictive Complex Flow Simulations, invited seminar, Embry-Riddle Aeronautical University, Daytona Beach, FL.
29. Zunino, H., Adrian, R.J., Clarke, A.B. (2015). Experimental studies of gas-particle mixtures under sudden expansion. 68th Annual Meeting of the APS Division of Fluid Dynamics, Boston, MA, November 22-24, 2015.
30. Zunino, H., Adrian, R.J., Ding, L. (2015). Non-intrusive calibration technique. Particle Image Velocimetry 2015 meeting, Santa Barbara, CA, September 14-16, 2015.
31. Chanyoung Park, R. T. Haftka, and N. H. Kim (2015). Experience with Several Multi-fidelity Surrogate Frameworks, 11th World Congresses of Structural and Multidisciplinary Optimization, Sydney, Australia, Jun, 2015.
32. Yiming Zhang, N. H. Kim, C. Park, R. T. Haftka (2015). One-dimensional Function Extrapolation Using Surrogates, 11th World Congresses of Structural and Multidisciplinary Optimization, Sydney, Australia, Jun, 2015.
33. Yiming Zhang, N. H. Kim, R. T. Haftka and C. Park (2015). Function Extrapolation at One Inaccessible Point Using Converging Lines, ASME 2015 International Design Engineering Technical Conference & Computers and Information in Engineering Conference, Boston, MA, USA, August 2-5, 2015.
34. Banerjee, T. (2015). A Genetic Algorithm Based Autotuning Approach for Performance and Energy Optimization, 6th International Green and Sustainable Computing Conference, Las Vegas, NV, Dec 2015.
35. Banerjee, T. (2015). Optimizing Nek5000 kernels for performance and energy, on May 27, 2015 at LLNL.
36. N. Kumar, A. George, H. Lam, G. Stitt, S. Hammond (2015). Understanding Performance and Reliability Trade-offs for Extreme-scale Systems using Behavioral Emulation, 2015 Workshop on Modeling & Simulation of Systems and Applications (ModSim 2015), Seattle, Washington.

37. N. Kumar, C. Hajas, A. George, H. Lam, G. Stitt (2015). Multi-scale, Multi-objective Behavioral Emulation of Future-gen Applications and Systems, 2015 Salishan Conference on High-Speed Computing, Gleneden Beach, Oregon, April 27-30, 2015.
38. Ranka,S. (2015). "Scalable Network Simulations", Random access talk, 2015 Salishan Conference on High-Speed Computing, Gleneden Beach, Oregon.
39. Zhang, J. and Jackson, T.L. (2015). Detonation Initiation with Thermal Deposition due to Pore Collapse in Energetic Materials - Towards the Coupling between Micro- and Macroscale. 51st AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit, 27-29 July 2015, Orlando, FL.
40. Amadio, G. and Jackson, T.L. (2015). A new packing code for creating microstructures of propellants and explosives. 51st AIAA/ASME/SAE/ASEE Joint Propulsion Conference and Exhibit, 27-29 July 2015, Orlando, FL.
41. G. Akiki, T Jackson, S Balachandar (2015). Mean and Fluctuating Force Distribution in a Random Array of Spheres. 68th Annual Meeting, Division of Fluid Dynamics, APS, Boston, MA, Nov. 2015.
42. TL Jackson, P Sridharan, J Zhang, S Balachandar (2015). Numerical Simulation of Shock Interaction with Deformable Particles Using a Constrained Interface Reinitialization Scheme. 68th Annual Meeting, Division of Fluid Dynamics, APS, Boston, MA, Nov. 2015.
43. J. Zhang, T.L. Jackson (2015). Detonation initiation with thermal deposition due to pore collapse in energetic materials - Towards the coupling between micro- and macroscales. 68th Annual Meeting, Division of Fluid Dynamics, APS, Boston, MA, Nov. 2015.
44. Dylan Rudolph and Greg Stitt (2015). An Interpolation-Based Approach to Multi-Parameter Performance Modeling for Heterogeneous Systems, IEEE International Conference on Application-specific Systems, Architectures and Processors (ASAP), July 2015.

2016

45. Yiming Zhang, N. H. Kim, Chanyoung Park, and R. T. Haftka (2016). Function Extrapolation of Noisy Data using Converging Lines, AIAA Modeling and Simulation Technologies Conference, San Diego, CA, USA, 4-8 January 2016.
46. Chanyoung Park, Giselle Fernández-Godino, R. T. Haftka, and N. H. Kim (2016). Validation, Uncertainty Quantification and Uncertainty Reduction for a Shock Tube Simulation, 18th AIAA Non-Deterministic Approaches Conference, San Diego, CA, USA, 4-8 January 2016.
47. B. Rollin, F. Ouellet, R. Koneru, and. S. Annamalai (2016). "Eulerian-Lagrangian Simulation of an Explosive Dispersal of Particles", DFD 2016, Bulletin of the American Physical Society, Vol. 61, No. 20.
48. R. Koneru, B. Rollin, F. Ouellet, S. Annamalai, S. Balachandar (2016). "Simulations of a shock wave interaction with a particle cloud", DFD 2016, Bulletin of the American Physical Society, Vol. 61, No. 20.
49. S. Annamalai and S. Balachandar (2016). "Generalized Faxen's theorem: Evaluating first-order (hydrodynamic drag) and second-order (acoustic radiation) forces on finite-sized rigid particles, bubbles and droplets in arbitrary complex flows", DFD 2016, Bulletin of the American Physical Society, Vol. 61, No. 20.
50. F. Ouellet, C. Park and B. Rollin, S. Balachandar (2016). "A Multi-Fidelity Surrogate

Model for Handling Real Gas Equations of State", DFD 2016, Bulletin of the American Physical Society, Vol. 61, No. 20.

51. Akiki, G., Jackson, T.L., and Balachandar, S. (2016). Quantifying and modeling the force variation within random arrays of spheres. Proceedings of the ASME 2016 International Mechanical Engineering Congress & Exposition, Nov. 11-17, 2016, Phoenix, AZ.
52. Jackson, T.L. and Zhang, J. (2016). Detonation Initiation of Energetic Materials Using Density-based Kinetics -Towards the Coupling between Micro and Mesoscale. Proceedings of the ASME 2016 International Mechanical Engineering Congress & Exposition, Nov. 11-17, 2016, Phoenix, AZ.
53. Zhang, J. and Jackson, T.L. (2016). Numerical simulation of blast-wave-particle and contact interaction induced by a detonation in condensed matter. AIAA Paper No. 2016-5105, 52st AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit, 25-27 July 2016, Salt Lake City, Utah.
54. M. G. Fernández-Godino, C. Park, N. H. Kim, and R. T. Haftka (2016). "Review of Multifidelity Surrogate Models". ECCOMAS Congress 2016, Uncertainty Quantification in CFD and Fluid Structure Interaction, Crete Island, Greece, June 5-10, 2016.
55. C. Park, M. G. Fernández-Godino, N.H. Kim, and R.T. Haftka (2016). "Validation, Uncertainty Quantification and Uncertainty Reduction for a Shock Tube Simulation". 18th AIAA Non-Deterministic Approaches Conference, San Diego, CA, USA. January 4-8, 2016.
56. M. G. Fernández-Godino, A. Diggs, C. Park, N. H. Kim, and R. T. Haftka (2016). "Anomaly Detection using Groups of Simulations". 18th AIAA Non-Deterministic Approaches Conference, San Diego, CA, USA. January 4-8, 2016.
57. Chanyoung Park, R T. Haftka, and N. H. Kim (2016). "Investigation of the Effectiveness of Multi-fidelity Surrogates on Extrapolation", ASME 2016 IDETC/CIE conference, Charlotte, NC, August 21-24, 2016.
58. Yiming Zhang, N. H. Kim, Chanyoung Park, and R. T. Haftka (2016). Function Extrapolation of Noisy Data using Converging Lines, AIAA Modeling and Simulation Technologies Conference, San Diego, CA, USA, 4-8 January 2016.
59. T. Banerjee, J. F. Hackl, M. Shringarpure, T. Islam, S. Balachandar, T. Jackson and S. Ranka (2016). "CMT-bone – A proxy Application for Compressible Multiphase Turbulent Flows", 23rd IEEE International Conference on High Performance Computing, Data and Analytics (HiPC). December 19-22, 2016, Hyderabad, India.
60. B.A. Johnson & E.A. Cowen (2016). Turbulent Boundary Layers and Sediment Suspension Absent Mean Flow-Induced Shear: An Experimental Study; American Geophysical Union, San Francisco, CA, December 2016.
61. B.A. Johnson & E.A. Cowen (2016). "Turbulent Boundary Layers and Sediment Suspension Absent Mean Flow-Induced Shear: An Experimental Study," American Physical Society Division of Fluid Dynamics Meeting, Portland, OR, November 2016.
62. H.A. Zunino , R.J. Adrian, A.B. Clarke (2016). American Physical Society Division of Fluid Dynamics Meeting, Portland, OR, November 2016.
63. Banerjee, T. (2016). "Optimizing CMT-nek for performance and power consumption", on 6/30/2016 at Sandia National Laboratories, Livermore, CA.
64. Banerjee, T. (2016). "Multiobjective Optimization of CMT-bone on Hybrid Processors", on 11/8/2016 at Hangzhou, China.

65. "Multiobjective Algorithms for Hybrid Multicore Processors", Keynote Speech by Sanjay Ranka on 11/9/2016 at Hangzhou, China.
66. G. Akiki, T.L. Jackson, S. Balachandar (2016). "Quantifying and modeling the force variation within random arrays of spheres" Society of Engineering Science 53rd Annual Technical Meeting, College Park, MD, October 2016.
67. G. Akiki, T.L. Jackson, S. Balachandar (2016). "Pairwise Interaction Extended Point Particle (PIEP) Model for a Random Array of Spheres" American Physical Society Division of Fluid Dynamics Meeting, Portland, OR, November 2016.

2017

68. Chanyoung Park, Raphael T. Haftka and Nam Ho Kim (2017). "Simple Alternative to Bayesian Multi-Fidelity Surrogate Framework", 55th AIAA Aerospace Sciences Meeting, AIAA SciTech Forum, (AIAA 2017-0135).
69. Kyle Hughes, Angela Diggs, Don Littrell, Sivaramakrisnan Balachandar, Raphael T. Haftka, Nam Ho Kim, Chanyoung Park, and Myles DelCambre (2017). "Uncertainty Quantification of Experiments on a Small Number of Explosively-Driven Particles", 55th AIAA Aerospace Sciences Meeting, AIAA SciTech Forum, (AIAA 2017-1463).
70. M. G. Fernández-Godino, F. Ouellet, S. Balachandar, and R. T. Haftka (2017). "Multi-fidelity surrogate-based optimization as a tool to study the physics in explosive dispersal of particles". 12th World Congress of Structural and Multidisciplinary Optimisation (WCSMO12), Braunschweig, Germany, June 5-9, 2017.
71. M. G. Fernández-Godino, F. Ouellet, S. Balachandar, and R. T. Haftka (2017). "Noise quantification in the study of instabilities during explosive dispersal of solid particles". V&V ASME Conference, Las Vegas, Nevada, May 3-5, 2017.
72. M. G. Fernández-Godino, R. T. Haftka, S. Balachandar, C. Gogu, S. Dubreuil, and N. Bartoli (2017). "Noise Filtering and Uncertainty Quantification in Surrogate based Optimization". 20th AIAA Non-Deterministic Approaches Conference, Kissimmee, FL, USA. January 8-12, 2018.
73. Moore, W.C., Akiki, G., Balachandar, S. (2017). "A Hybrid Physics-Based Data-Driven Approach for Point-Particle Force Modeling," APS Division of Fluid Dynamics 70th Annual Meeting (2017).
74. Moore, W.C., Balachandar, S. (2017). "Machine Learning Approach for the Development of Point-Particle Force Models," AIAA Region II Student Conference, (2017).
75. Mehta, Y., Neal, C., Jackson, T. L., Balachandar, S., and Thakur, S (2017). "Strong shock propagating over a random bed of spherical particles". 70th Annual Meeting of APS DFD, Denver, CO (2017).
76. Mehta, Y. (2017). "Results on fully resolved simulations and modeling of shock particle interaction," Computation weekly seminar, LLNL, CA (2017).
77. Mehta, Y. (2017). "Microscale Simulations," AST Review of PSAAP II center CCMT, UF, FL (2017).
78. J. Garno, F. Ouellet, R. Koneru, S. Balachandar, and B. Rollin (2017). "Predictive Capability of the Compressible MRG Equation for an Explosively Driven Particle with Validation", DFD2017, Bulletin of the American Physical Society, Vol. 62, No. 14, (2017).

79. R. Koneru, B. Rollin, F. Ouellet, C. Park, and S. Balachandar (2017). "Euler-Lagrange Simulations of Shock Wave-Particle Cloud Interaction", DFD2017, Bulletin of the American Physical Society, Vol. 62, No. 14, (2017).
80. F. Ouellet, C. Park, R. Koneru, S. Balachandar, and B. Rollin (2017). "A Multi-Fidelity Surrogate Model for the Equation of State for Mixtures of Real Gases", DFD2017, Bulletin of the American Physical Society, Vol. 62, No. 14, (2017).
81. F. Ouellet, C. Park, B. Rollin, and S. Balachandar, S. (2017). "Analysis of a Multi-Fidelity Surrogate for Handling Real Gas Equations of State", 20th Biennial APS Conference on Shock Compression of Condensed Matter, St-Louis, MO (2017).
82. B. Rollin, F. Ouellet, R. Koneru, J. Garno, and B. Durant (2017). "Effects of Initial Particle Distribution on an Energetic Dispersal of Particles", DFD2017, Bulletin of the American Physical Society, Vol. 62, No. 14, (2017).
83. B. Rollin, R. Koneru, and F. Ouellet (2017). "A Twist on the Richtmyer-Meshkov Instability", DFD2017, Bulletin of the American Physical Society, Vol. 62, No. 14, (2017).
84. T.L. Jackson (2017). SIAM Sixteenth International Conference on Numerical Combustion, Orlando, FL, April 3-5, 2017.
85. T.L. Jackson (2018). "Three-dimensional Mesoscale Simulations of Detonation Initiation in Energetic Materials with Density-based Kinetics'", 20th APS Shock Compression in Condensed Matter (SCCM), St. Louis, MO, July 9-14, 2017.
86. T.L. Jackson (2017). "Multidimensional Mesoscale Simulations of Detonation Initiation in Energetic Materials", Invited presentation at Los Alamos National Laboratory, Los Alamos, NM, August, 2017.
87. T.L. Jackson (2017). "Mesoscale Simulation of Solid Propellant Combustion (Rocfire), 48th JANNAF Combustion Meeting, Newport News, VA, December 4-7, 2017.
88. Chanyoung Park, Raphael T. Haftka and Nam Ho Kim (2017). "Simple Alternative to Bayesian Multi-Fidelity Surrogate Framework", 55th AIAA Aerospace Sciences Meeting, AIAA SciTech Forum, (AIAA 2017-0135)
89. Bae, S. J., Kim, N. H., Park, C., & Kim, Z. (2017). "Confidence Interval of Bayesian Network and Global Sensitivity Analysis", 55th AIAA Aerospace Sciences Meeting, AIAA SciTech Forum, (AIAA 2017-0595).
90. Zhang, Y**.,** Neelaantan, A., Kumar,N., Park, C., Haftka, R., Kim, N., Lam, H. (2017). "Multi-fidelity Surrogates of Abstract Application and Architecture Models for Predicting Application Performance", Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems (PMBS), International Workshop on IEEE.
91. Nili S, Park C, Haftka RT, Balachandar S, Kim NH. (2017). Sensitivity Analysis of Force Models for a Four-Way Coupled Eulerian-Lagrangian Dispersed Multiphase Flow. In23rd AIAA Computational Fluid Dynamics Conference 2017 (p. 3800).
92. Nili S, Park C, Haftka RT, Kim NH, Balachandar S. (2017). Effect of Finite Particle Size on Convergence of Point Particle Models in Euler-Lagrange Multiphase Dispersed Flow. Bulletin of the American Physical Society. 2017 Nov 20;62.
93. Kyle Hughes, Chanyoung Park, Raphael Haftka Nam-Ho Kim (2017). "Forensic Uncertainty Quantification of Explosive Dispersal of Particles," APS Shock Compression of Condensed Matter 2017, St. Louis, MI (BAPS.2017.SHOCK.B8.2).
94. Kyle Hughes, Angela Diggs, Don Littrell, S. Balachandar, Raphael Haftka, Nam-Ho Kim, Chanyoung Park, and Myles DelCambre (2017). "Uncertainty Quantification of

Experiments on a Small Number of Explosively-Driven Particles", 55th AIAA Aerospace Sciences Meeting, AIAA SciTech Forum, (AIAA 2017-1463).

2018

95. Rahul Koneru and Fred Ouellet, "Shock-particle interaction and explosive dispersal of particles", Workshop on Physics and Modeling of Dispersed Multiphase Flows, 22-23 October 2018, Stanford University, California (2018)
96. B. Rollin, R. Koneru, F. Ouellet, J. Garno, "Simulations of a Shock-Driven Instability Developing from a Curtain of Particles", DFD2018, Bulletin of the American Physical Society (2018).
97. Liu, K. and Balachandar, S. "Applications of pairwise interaction extended point-particle model: mid-field spray control". DFD2018, Bulletin of the American Physical Society (2018).
98. Balachandar, S., Moore, W., Akiki, G., and Liu, K. "A novel Euler-Lagrange method that incorporates fully resolved physics using pairwise interaction extended point-particle (PIEP) model". DFD2018, Bulletin of the American Physical Society (2018).
99. Salinas, J., Shringarpure, M., Cantero, M., and Balachandar, S. "Ambient fluid entrainment and basal drag in turbidity currents". DFD2018, Bulletin of the American Physical Society (2018).
100. Ling, Y., Balachandar, S., and Jian, X. "Asymptotic scaling laws for spherical and cylindrical finite-source blast waves". DFD2018, Bulletin of the American Physical Society (2018).
101. Hughes, K., Prestridge, K., Kim, N.H., Haftka, R., and Balachandar, S. "Proton radiography of explosively dispersed metal particles while under vacuum". DFD2018, Bulletin of the American Physical Society (2018).
102. Mehta, Y., Jackson, T.L., and Balachandar, S. "Fully resolved simulations of air-shock interacting with randomly distributed spherical particles". DFD2018, Bulletin of the American Physical Society (2018).
103. Subramaniam, S. and Balachandar, S. "Extending current mathematical formulations of multiphase flow to regions of strong inhomogeneity". DFD2018, Bulletin of the American Physical Society (2018).
104. Zhou, K. and Balachandar, S. "Investigation of direct forcing immersed boundary method". DFD2018, Bulletin of the American Physical Society (2018).
105. Fujisawa, K., Sridharan, P., Jackson, T.L, Zhang, J. and Balachandar, S. "Vortex formation in shock interaction with a deformable ellipsoidal particle", DFD2018, Bulletin of the American Physical Society (2018).
106. J. Garno, F. Ouellet, R. Koneru, T.L. Jackson, S. Balachandar, and B. Rollin, "Assessing the Drag Law for an Explosively Driven Particle using Experiments and Uncertainty Quantification", DFD2018, Bulletin of the American Physical Society (2018).
107. F. Ouellet, R. Koneru, J. Garno, B. Rollin, and S. Balachandar, "Energetic Dispersal of a Carefully Perturbed Bed of Particles", DFD2018, Bulletin of the American Physical Society (2018).
108. R. Koneru, B. Rollin, C. Park, F. Ouellet, and S. Balachandar, "Euler-Lagrange Simulations of Shock-Particle Cloud Interaction", DFD2018, Bulletin of the American Physical Society (2018).

109. B. Rollin, R. Koneru, and F. Ouellet, “Numerical Investigation of a Shock-Driven Perturbed Dense Curtain of Solid Particles”, IWPCTM16, Marseille, France (2018).
110. F. Ouellet, R. Koneru, B. Rollin, and S. Balachandar “Effect of Initial Particle Configurations on the Behavior of a Particle Cloud Following a Detonation Wave”, IWPCTM16**,** Marseille, France (2018).
111. Fernandez-Godino, M.G., Haftka, R.T., Balachandar, S., Gogu, C., Bartoli, N. and Dubreuil, S., “Noise Filtering and Uncertainty Quantification in Surrogate based Optimization”, *2018 AIAA Non-Deterministic Approaches Conference* (2018).
112. Moore, W. C., Balachandar, S., "Analysis of particle wakes for PIEP modeling in Euler-Lagrange simulations," APS Division of Fluid Dynamics 71th Annual Meeting, Atlanta, GA, 2018.
113. Moore, W. C., Balachandar, S., "Towards Euler-Lagrange simulations with fully resolved physics," Workshop on Physics and Modeling of Dispersed Mulitphase Flows, Stanford, CA, 2018.
114. Kyle T. Hughes, Chanyoug Park, Nam H. Kim, Raphael T. Haftka, S. Balachandar, Angela, Diggs, and Don Littrell, Simulation-driven experiments of macroscale explosive dispersal of particles, AIAA Science and Technology Forum and Exposition, January 8-12, 2018, Kissimmee, Florida.
115. Justin T. Mathew, Chanyoung Park, Nam H. Kim, Raphael T. Haftka, Epistemic uncertainty stemming from measurement processing – A case study of multiphase shock tube experiments, AIAA Science and Technology Forum and Exposition, January 8-12, 2018, Kissimmee, Florida.
116. Chanyoung Park, Nam H. Kim, and Raphael T. Haftka, Including □ in multi-fidelity surrogate prediction can make discrepancy extrapolation accurate by reducing bumpiness, AIAA Science and Technology Forum and Exposition, January 8-12, 2018, Kissimmee, Florida.
117. Chanyoung Park, Samaun Nili, Justin T. Mathew, Nam H. Kim, and Raphael T. Haftka, Uncertainty investigation for shock tube simulation error quantification, AIAA Science and Technology Forum and Exposition, January 8-12, 2018, Kissimmee, Florida.
118. Yiming Zhang, Nam H. Kim, Chanyoung Park, and Raphael T. Haftka, Effect of varying test cost on design of experiments, AIAA Science and Technology Forum and Exposition, January 8-12, 2018, Kissimmee, Florida.
119. Chanyoung Park, Nam H. Kim, Raphael T. Haftka, Least bumpiness calibration with extrapolative bias correction, Proceedings of the ASME 2018 IDETC/CIE Conference, August 26-29, 2018, Quebec City, Canada
120. Yiming Zhang, Nam H. Kim, and Raphael T. Haftka, A strategy for adaptive sampling when sampling cost is variable in design space, Asian Congress of Structural and Multidisciplinary Optimization, May 21 - 24, 2018, Dalian, China
121. Nili, Samaun, Chanyoung Park, Raphael Haftka, Nam Kim, and S. Balachandar, Convergence of Point Particle Models in Euler-Lagrange Simulations of Shock-Particle Interaction, *Bulletin of the American Physical Society* (2018).
122. David Zwick, Highly resolved Euler-Lagrange multiphase flow simulation applied to dispersed two-phase flow, Sandia National Laboratories, Albuquerque, NM, October 30, 2018.

123. Jason Hackl, CMT-nek: a shock-capturing discontinuous Galerkin spectral element method for compressible multiphase flow, Naval Surface Warfare Center Indian Head Explosive Ordnance Disposal Technical Division, Indian Head, MD, October 12, 2018.
124. David Zwick, A novel strategy for load-balanced exascale simulations of 4-way coupled Euler-Lagrange simulations, Workshop on Physics and Modeling of Dispersed Multiphase Flows, Stanford University, October 22, 2018.
125. David Zwick, Highly resolved Euler-Lagrange multiphase flow simulation applied to dispersed two-phase flow, Los Alamos National Laboratory, Los Alamos, NM, September 20, 2018.
126. David Zwick, A Scalable Discrete Element Method in Nek5000 and CMT-nek, 6th Nek5000 Users/Developers Meeting, Tampa, FL, April 17, 2018.
127. Jason Hackl, CMT-nek: a shock capturing discontinuous G , 6th Nek5000 Users/Developers Meeting, Tampa, FL, April 17, 2018.

12. Workshops Held or Attended

2014

1. Dr. Chanyoung Park. Attended "A Short Course on Uncertainty Quantification", Stanford CA, June 2-4 2014.
2. Dr. Tania Banerjee. Salishan Conference, April 2014.
3. 2014 International Workshops on the Physics of Compressible Turbulent Mixing, hosted by Lawrence Livermore National Laboratory, 31 August 2014 - 5 September 2014, San Francisco, California, USA.
4. Professor Nam-Ho Kim. Attended, "A Short Course on Uncertainty Quantification", Stanford CA, June 2-4 2014.
5. Heather Zunino presented tomographic PIV at a PIV workshop to a group of faculty, post-docs, and Ph.D. students at Instituto Superior Técnico (IST) in Lisbon, Portugal. She was invited to stay for this workshop and present after meeting with several researchers from IST at the Laser Symposium this summer.

2015

6. Presented "Dakota - Tutorial", CCMT workshop, February 19, 2015.
7. S. Balachandar and T.L. Jackson. Co-organized a mini-symposium titled, Turbulence and Mixing in Shock-Driven Multiphase Flows. APS (APS) Topical Group on Shock Compression of Condensed Matter (SCCM), Tampa, Florida, June 2015.
8. Deep Dive Workshop. Held at the University of Florida on Feb 3-4, 2015.
9. "Good Software Engineering Practices and Beyond" Workshop - Internal workshop - organized by Bertrand Rollin - Macroscale team, held Feb 19, 2015.
10. Heather Zunino presented tomographic PIV at a PIV workshop to a group of faculty, post-docs, and Ph.D. students at Instituto Superior Técnico (IST) in Lisbon, Portugal. She was invited to stay for this workshop and present after meeting with several researchers from IST at the Laser Symposium summer 2015.
11. Dr. Tania Banerjee Attended the 6th International Green and Sustainable Computing Conference, 2015.

2016

12. Dr. Sanjay Kumar and Ph.D. students Kasim Alli and Carlo Pascoe (2016). Attended the WEST workshop. Dr. Kumar gave a talk.
13. Jason Hackl, David Zwick, Goran Marjanovic and Bradford Durant presented at the 5th Nek5000 Users/Developers Meeting on August 10-12 at Massachusetts Institute of Technology. CMT-nek followed the official version of Nek5000 in migrating to http://www.github.com/Nek5000.
14. Deep Dive Workshop on Multiphase Physics. Held in St. Petersburg, Florida on October 6-7, 2016.
15. Rahul Koneru and Fred Ouellet attended "Scaling your science on Mira" workshop, May 24-25, Argonne National Laboratory.

2017

16. A Boot Camp on CMT-nek, November 29, 2017 at CCMT, University of Florida, Gainesville. Organized by B.Rollin and J. Hackl.
17. T.L. Jackson, F. Najjar, and H. Najm (2017). Session Organizer for Focus Topic titled, Uncertainty quantification (UQ) in compressible high-speed flows, APS Shock Compression of Condensed Matter (SCCM), St. Louis, Missouri, July 9-14, 2017.

2018

18. 6th Nek5000 User/Developer meeting, 17-18 April 2018, Tampa, Florida
19. Workshop on Physics and Modeling of Dispersed Multiphase Flows, 22-23 October 2018, Stanford University, California

13. Students and Staff Internships

13.1 Internships Completed

2014

1. Heather Zunino, Ph.D. Student (ASU), US, Dr. R. Adrian. Ms. Zunino completed her 10-consecutive week stay at Los Alamos National Laboratory, under the mentorship of Dr. Kathy Prestridge this summer (May-August 2014). Dr. Prestridge is the Extreme Fluids Team leader in the Physics-23 group at LANSCE. Project: Vertical Shock Tube (Calibration and Tomographic PIV), Horizontal Shock Tube (Particle Tracking Program).
2. Kevin Cheng, MS Student, Florida. Lawrence Livermore National Laboratory. Mentor: Dr. Maya Gokhale, Dr. Scott Lloyd. Project: An Emulation Framework for Tracing near Memory Computation. US, Dr. Alan George, ECE, MS (graduated Fall 2014), core.
3. Dr. Chanyoung Park, Postdoc, CCMT. Visited Sandia National Laboratories, Albuquerque NM, on March 24-28 2014.

2015

4. Dr. Jason Hackl, Postdoc, CCMT. Visited LLNL from February 23-27, 2015. LLNL. Sam Schofield, Robert Nourgaliev, Rob Rieben, Tzanio Kolev, Fady Najjar, David Dawson. CMT-nek.
5. Dr. Bertrand Rollin, Staff Scientist, CCMT. March 16-20, 2015, LANL.
6. Nalini Kumar, Ph.D. Student, India, ECE, Dr. Alan George, part cost share and part leveraged. (Internship not required). March-May, 2015. LLNL. Dr. James Ang.
7. Chris Hajas, M.S. Student, US, ECE, Dr. Herman Lam, core. May 18-August 18, 2015 at LLNL with Dr. Maya Gokhale.
8. Christopher Neal, Ph.D. Student, US, MAE, Dr. S. Balachandar, core. June 14-August 20, 2015 at LLNL with Dr. Kambiz Salari. Chris' work was highlighted in the Stewardship Science Academic Programs (SSAP) Annual, which highlights the exceptional research activities and people supported by the SSAP (PSAAP is under the SSAP).
9. Carlo Pascoe, Ph.D. Student, US, ECE, Dr. Herman Lam, core. Will intern summer 2015 at LLNL with Dr. Maya Gokhale.
10. Giselle Fernandez, Ph.D. Student, Argentina, MAE, Drs. Haftka and Kim, core. Visited Sandia National Laboratories, Albuquerque NM, Oct 12-Dec 25, 2015.
11. Dr. Tania Banerjee, PhD., one week internship in May, 2015 to LLNL.
12. Justin Thomas Mathew, MS Student (2015), Drs. Haftka and Kim, core. Visited Los Alamos National Laboratory. T6-Theoretical Biology and Biophysics group, Dr. Nick Hengartner, studying and developing extensions of epidemiological mathematical models of infectious disease.

2016

13. Dr. Chanyoung Park, Postdoc, CCMT. Feb 22-26, 2016 at LLNL with Dr. Samuel P. Schofield.
14. Dr. Jason Hackl (Postdoc, CCMT) visited Sandia National Labs, Albuquerque, New Mexico from January 16-20 to present CMT-nek to Greg Weirs and other personnel.

15. David Zwick, Ph.D., US, MAE, Dr. Balachandar, core. Internship to Sandia National Lab, May-August, 2016, Drs. John Pott, Kevin Ruggirello.
16. Goran Marjanovic. Ph.D., US, MAE, Dr. Balachandar, core. Internship to Sandia National Lab, Aug-Nov, 2016, Drs. Paul Crozier, Stefan Domino.
17. Georges Akiki, Ph.D. Student, US, MAE, Dr. Balachandar, core. Internship to LANL. May-August, 2016, Dr. Marianne Francois.
18. Dr. Tania Banerjee, PhD., one week internship in June, 2016 to Sandia.

2017

19. Maria Giselle Fernandez-Godino, PhD Candidate, CCMT. May 06- July 26, 2017 at University of Toulouse, Toulouse France with Dr. Christian Gogu.
20. Kyle T. Hughes, Ph.D. Student, MAE, Drs, R. Haftka and N. Kim, core. August 14-December 8, 2017 at LANL with Dr. Katherine Prestridge.
21. Paul Crittenden, Ph.D. Student, US, MAE, Dr. Balachandar, core. Internship to LLN, Spring, 2017, Drs. Kambiz Salari and Sam Schofield.
22. Mohamed Gadou, Ph.D. Candidate, CISE, Dr. Ranka, core. Internship to LANL; Summer, Dr. Galen Shipman.
23. Trokon Johnson, Ph.D. Student, US, ECE, Drs. H. Lam and G. Stitt. Internship to LANL, Summer, 2017, Drs. Cristina Garcia- Cardona, Brendt Wohlberg, Erik West.
24. Yash Mehta, Ph.D. Student, MAE, Dr. Balachandar. Internship to LLNL; Summer, 2017, Dr. Kambiz Salari.

2018

25. Frederick Ouellet, PhD Candidate, CCMT. ISTI Co-Design Summer School 2018, Los Alamos National Laboratory. Supervisor: Cristoph Junghans
26. Prashanth Sridharan, Ph.D., US, MAE, Dr. Balachandar, core, Internship to LANL Summer 2018.
27. Brad Durant, Ph.D., MAE, Dr. Balachandar, core. Internship to Los Alamos National Lab, Los Alamos, Summer 2018, Dr. Joseph Schmidt in the Theoretical Division (NTA-XTD).

13.2 Internships Planned

1. Josh Garno, Ph.D. Student, US, MAE, Dr. S. Balachandar, core; summer 2019, LANL

13.3 Internships Not Yet Planned

2. Ryan Blanchard, Ph.D., US, ECE, Dr. Herman Lam, core
3. Chandler Moore, Ph.D. Student, US,MAE, Dr. Balachandar, core

13.4 Graduated Students

1. Kevin Cheng, MS Student (2014). Dr. Alan George, ECE.
2. Hugh Miles, BS Student (2015). US, ECE, Dr. Greg Stitt.
3. Chris Hajas, M.S. Student (2015). US, ECE, Dr. Herman Lam.
4. Angela Diggs, Ph.D. (2015). US, MAE, Dr. S. Balachandar (other funding, internship not required). Currently employed at Eglin AFB and working with the Center.

5. Subramanian Annamalai, Ph.D. (2015). MAE, Dr. S. Balachandar.
6. Dylan Rudolph, MS, Student, US, ECE, Dr. Greg Stitt, part core and part leveraged.
7. Kasim Alli, M.S. Student, US, ECE, Dr. Greg Stitt.
8. Parth Shah, M.S. Student (2016). India, ECE, Dr. Herman Lam.
9. Georges Akiki, Ph.D. (2016). MAE, Dr. S. Balachandar.
10. Justin Thomas Mathew, M.S. (2017). MAE, Drs. Haftka and Kim.
11. Sankeerth Mogili, M.S. (2017). Dr, Ranka, CISE, UF.
12. Ajay Ramaswamy, M.S. (2017). ECE, Dr. Herman Lam.
13. Nalini Kumar, Ph.D. (August, 2017). Dr. H. Lam, ECE. Intel, Santa Clara CA.
14. Chris Neal,M.S.(2017). MAE, Dr. Balachandar.
15. Brandon Osborne, M.S. (2018). US, MAE, Dr. Balachandar.
16. Yiming Zhang, Ph.D. (2018). MAE, Drs. Haftka and Kim, GE Global Research, NY.
17. Mohamed Gadou, Ph.D. (2018). CISE, Dr. Ranka, Bloomberg, NY.
18. Yash Mehta, Ph.D. (2018). MAE, Dr. Balachandar, Postdoc, LANL.
19. Kyle Hughes, Ph.D. (2018).MAE, Drs. Haftka and Kim, Postdoc, LANL.
20. Giselle Fernandez, Ph.D. (2018). MAE, Drs. Haftka and Kim, Postdoc, LANL.
21. Goran Marjanovic, Ph.D. (2018). MAE, Dr. Balachandar, Raytheon.
22. Paul Crittenden, Ph.D. (2018). MAE, Dr. Balachandar, Adjunct Professor, Mathematics, UNF.
23. Prashanth Sridharan, Ph.D. (2018). MAE, Drs. Balachandar and Jackson.
24. Cameron Stewart, Ph.D. (2018). MAE, Dr. Balachandar (leveraged funding), Naval Surface Warfare Center, Indian Head, MD.
25. David Zwick, Ph.D. (2019). MAE, Dr. Balachandar, Staff Scientist, Sandia.

13.5 Students Who Will be Graduating

26. Carlo Pascoe, Ph.D.
27. Heather Zunino, Ph.D.
28. Rahul Koneru, Ph.D.
29. Samaun Nili, Ph.D.
30. Fred Ouellet, Ph.D.
31. Keke Zhai, Ph.D.
32. Brad Durant, Ph.D.
33. Josh Garno, Ph.D.
34. Sai Chenna, Ph.D.
35. Trokon Johnson, Ph.D.
36. Chandler Moore, Ph.D.
37. Aravind Neelakantan, Ph.D.
38. Ryan Blanchard, Ph.D.

13.6 Placement of Staff

1. Dr. Bertrand Rollin, Staff Scientist, CCMT. Faculty position at Embry-Riddle, Fall 2015.
2. Dr. Mrugesh Shringarpure, Postdoc, CCMT. Researcher, Exxon Mobil, Spring 2016.
3. Dr. Subramanian Annamalai, PhD (2015). Postdoc through March 2017, currently employed at Optym.

4. Dr. Georges Akiki, PhD (2016), Dr. S. Balachandar, MAE; Postdoc thru March 2017, Currently Postdoctoral Associate, LANL.
5. Dr. Chanyoung Park. Staff Scientist through December 2018. Currently Caterpillar, Illinois.

14. NNSA Laboratory Interactions

2014

1. Rob Cunningham, LANL. Setting up "Florida" file sharing group on Mustang
2. Blaise Barney, LLNL. Setting up account on Vulcan and Surface.
3. Greg Weir, SNL. Introduction to Catalyst.
4. Nathan Fabian, SNL. Introduction to Catalyst.
5. Don Frederick, LLNL. Issue with submitting a run on Vulcan.
6. John Gyllenhaal, LLNL. Help building and running Rocflu on Vulcan.
7. Jan Nunes, LLNL. Account request on Edge
8. Discussions with Donald Frederick of Lawrence Livermore National Laboratory related to MPI issues on the Vulcan computing cluster – discussion was relevant to parallel operation of the Rocflu code.
9. Telecon with Paraview Catalyst developers and users (Greg Weirs, Nathan Fabian, Kenneth Moreland at Sandia National Laboratory) at about deploying the Catalyst library into Rocflu for in-situ visualization.
10. Worked with Greg Lee (Livermore Computing Center)–who is a software debugging/ troubleshooting expert at LLNL–to get Rocflu to run at scale on the Vulcan computing cluster at LLNL.
11. Interactions with Rich Cook (Livermore Computing) who is the visualization expert at LLNL–we are working with him to get in-situ visualization capabilities using Catalyst integrated into Rocflu.
12. Interacted with David DeBonis at Sandia National Laboratories to get setup with using PowerInsight for power and energy measurements.
13. Interacted with Robert Cunningham and Amanda Bonnie at Los Alamos National Laboratory for temperature measurements using the data collection tool Lightweight Distributed Metric Service.
14. Interacted with Justin A. Too and Daniel J. Quinlan at the Lawrence Livermore National Laboratory on compilation and installation of ROSE.
15. Interacted with Blaise Barney at the Lawrence Livermore National Laboratory on various system issues, including access to CVS.
16. Road trip to SNL (Albuquerque) and LANL, March 24 – 26, 2014 (Herman, Bala, and Rafi)
17. SNL (Albuquerque: Jim Ang, Arun Rodrigues, Scott Hemmert, Simon Hammond - from SST team at SNL, Albuquerque, NM
18. SNL (Livermore): Jeremiah Wilke from SST Macro team at SNL, Livermore, CA

2015

19. LLNL: Maya Gokhale regarding proposal to analyze memory access demands of CMT-Nek kernels and evaluate the potential benefits of utilizing stacked memories with embedded logic functions to increase performance
20. Dr. Steve Beresh (SNL), visit CCMT and gave talk, Thursday April 23, 2015
21. The members of the microscale group (in particular Chris Neal and Yash Mehta, another graduate student) have worked closely with the staff at LLNL's Livermore Computing center to resolve an issue that Rocflu was having when it was scaled up and executed on the BG/Q machine Vulcan. We have also been in touch with Rich Cook at LLNL's Livermore Computing to continue our work with Paraview's Catalyst library integration into Rocflu. During the fall and spring, Chris Neal has been in touch with Dr. Kambiz Salari to keep him up-to-date on the progress of Chris's research. Chris and Yash were aided by Blaise Barney (LLNL) in renewing their computing accounts at LLNL.
22. Drs. Jason Hackl and Mrugesh Shringarpure and David Zwick and Goran Marjanovic. Visited Argonne National Lab for the 2015 Nek User/Developer meeting to commit CMT-nek to the nek5000 repository and work intensively with Prof. Paul Fischer.
23. Barry Rountree, LLNL. Tania worked closely with Barry's group to come up to speed on using RAPL to measure and bound power on Intel platform
24. Tanzima Islam, LLNL. Tania is working with Tanzima on validating CMT-bone proxy app
25. David DeBonis, SNL. Tania interacted with David to come up to speed on using PowerInsight for power and energy measurements.
26. Patrick Lavin, a summer intern with Barry Rountree, worked on the performance issue and used vectorization effectively on the derivative optimization code getting about 40% improvement using vectorization alone. Tania will follow up with verifying if it is possible to get further improvement with CHiLL based autotuning approach.
27. S. Balachandar and H. Lam attended and presented at the 2015 ASC PI meeting at Monterey, CA.
28. S. Balachandar visited LANL and interacted with Drs. Kathy Prestridge, Robert Gore, and John Schwarkopf in November 2015.

2016

29. S. Balachandar and S. Ranka attended and presented at the 2016 ASC PI meeting at Las Vegas, NV.
30. Tania Banerjee collaboratied with Jeremy Wilke, Gwen Voskuilen and Arun Rodriguez on MLM.
31. Tania Banerjee collaborated with Tanzima Islam to validate proxy app CMT-bone.
32. Justin Mattews and Chanyoung Park collaborated with Justin Wagner (Sandia).
33. Kyle Hughes collaborated with Kathy Prestridge (LANL) in sharing p-rad experiments of rapidly dispersed pre-fragmented particles.
34. Jeff St. Clair and Balachandar are interacting with Fady Najjar (LLNL) in performing ALE3D simulations of intense shock propagation over a close-packed bed of deformable aluminum particles.
35. Balachandar interacted with John Schwarzkef (LANL) in completing a review article on compressible multiphase flow for the second evolution of multiphase flow handbook.
36. Balachandar and Jackson organized and hosted the Deepdive workshop on multiphase flows. Approximately 15 research scientists from the three NNSA laboratories attended.

37. Yash Metha interacted with Kambiz Salari (LLNL) on three-dimensional simulations of a strong shock propagating through a random bed of particles. This work has been written up and is in the final stages of submission as an archival journal paper.
38. Tom Jackson, Fady Najjar (LLNL), and Habib Najm (Sandia) are organizing a Focus Topic session on UQ in high speed flows for the upcoming APS SCCM meeting in July.
39. The UQ team has regular teleconference calls with Greg Weirs (Sandia).

2017

40. T.L. Jackson, F. Najjar, and H. Najm (2017). Session Organizer for Focus Topic titled, Uncertainty quantification (UQ) in compressible high-speed flows, APS Shock Compression of Condensed Matter (SCCM), St. Louis, Missouri, July 9-14, 2017.
41. T.L. Jackson, visited with Dr. Mark Short, LANL, August 2017.
42. The UQ team has regular teleconference calls with Greg Weirs (Sandia).
43. S. Balachandar and T.L. Jackson interacting with Dr. F.M. Najjar on fully resolved inviscid simulations of flow over a particle close to a wall.
44. S. Balachandar attended the 2017 ASCI PI meeting and presented at Monterey, CA
45. Jeff St Clair and Balachandar are using the Navy code Dismas along with the Sandia solid mechanics code Peridynamics to solve the flow structure interaction problem at detonation conditions.
46. Yash Meta is continuing his interaction with Dr. Kambiz Salari, in both processing his shock propagation over a random array of particles and process the data from shock propagation over a moving array of particles.
47. From Sandia Kevin Ruggirello and Shane Schumacher attended David Zwick's PhD proposal in September 2017 via teleconference.
48. Dr. Jason Hackl visited Sandia National Laboratory January 16-19, 2017, and presented. His interactions with Drs. Bill Rider, Travis Fisher, Greg Weirs and John Shadid were a wealth of education on testing, assessing convergence, and crucial issues in entropy stability of high-order schemes.
49. Dr. Jason Hackl attended the 5th International Workshop on High-Order CFD Methods, AIAA SciTech, Kissimmee, FL, January 6-7, 2018, and met with Dr. Travis Fisher (SNL) to discuss artificial viscosity and entropy splitting.
50. Chanyoung Park: Met Vicente Romero from Sandia and discussed the idea of cost based adoptive sampling and presented the Sandia shocktube validation and UQ to him at the SciTech 2018 conference.
51. Kyle Hughes: Met Justin Wagner (Sandia) and Daniel Guildenbecher (Sandia) and shared his experience on the P-rad experiment and the Eglin tests at the SciTech 2018 conference.
52. Justin Mathew: Met Justin Wagner (Sandia) and received feedback about our UQ study for the Sadia shock-tube experiments at the SciTech 2018 conference.
53. Herman Lam and students meet with Maya Gokhale at SuperComputing 2017.