CCMT

CENTER FOR COMPRESSIBLE MULTIPHASE TURBULENCE

Y1 Annual Report

1/27/2014-1/26/2015

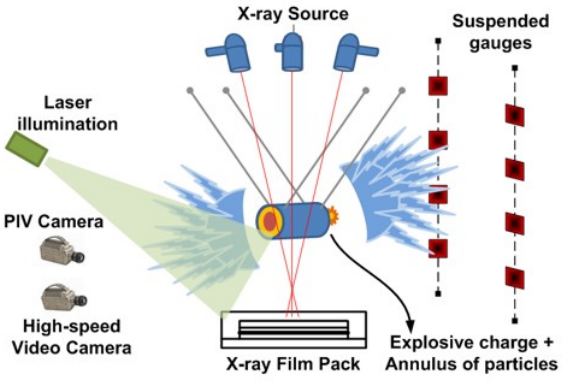

UF UNIVERSITY of FLORIDA

Center for Compressible Multiphase Turbulence

Table of Contents

1. Introduction

1.1 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 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:

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*

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.

1.2 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

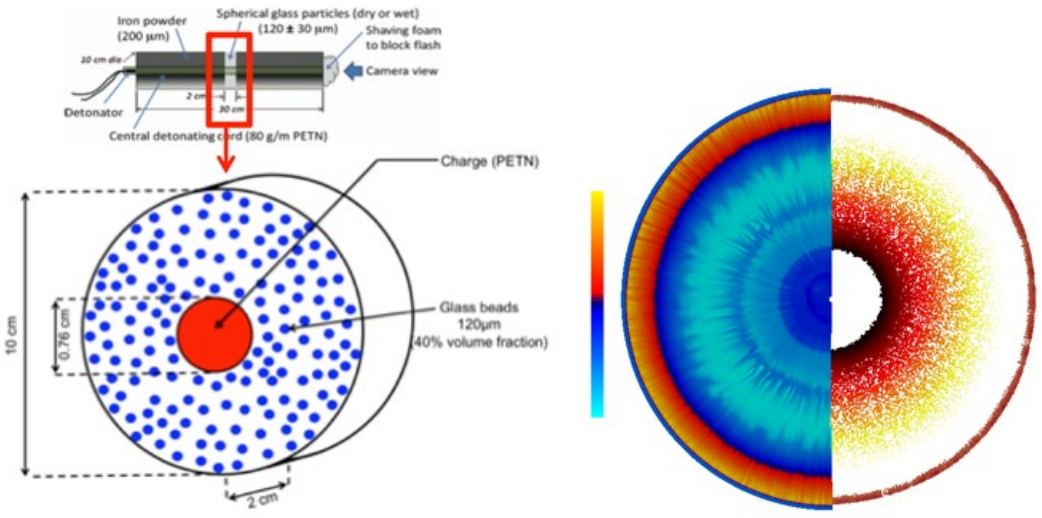

Figure 1.1. Experiment (left) and simulation (right). Left: Set up of Frost et al.'s experiment (top), and a zoom on the zone of interest (bottom), replicated in our simulation. Right: Left half: Gas Density. Right half: The red line is the shock front. The colored spheres are computational particles (darker tone is lower velocity). Not all particles are shown.

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.

1.3 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

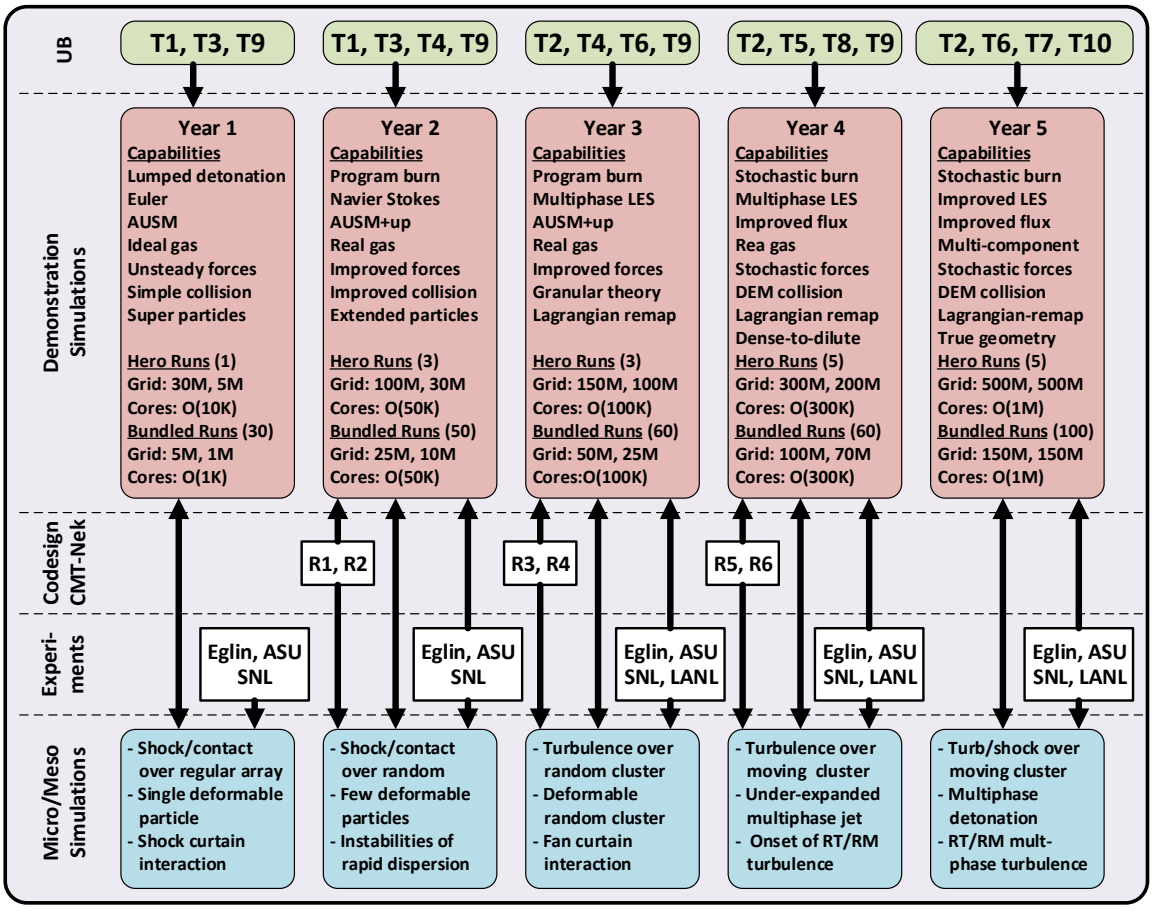

Figure 1.2. A roadmap showing the sequence of progressively more complex simulations to be performed over the five years of the center. Micro and mesoscale simulations will be performed along with macroscale simulations of the demonstration problem. Both *hero runs* and *bundled runs* will be performed. 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.

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, 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.

1.4 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.

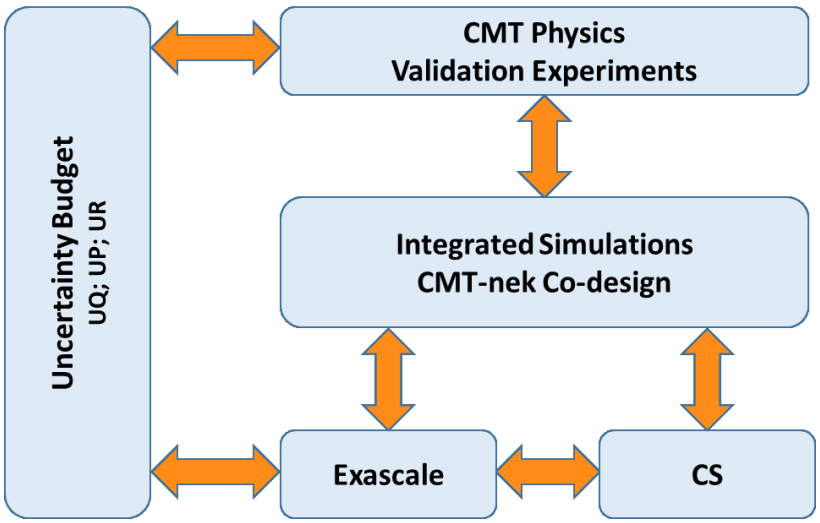

Figure 1.3. Integration – how the different pieces fit together. Here, UB drives the CMT physics, validation experiments, and exascale behavioral emulation/simulation. CS and exascale feeds into the CMT-nek codesign strategy.

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.

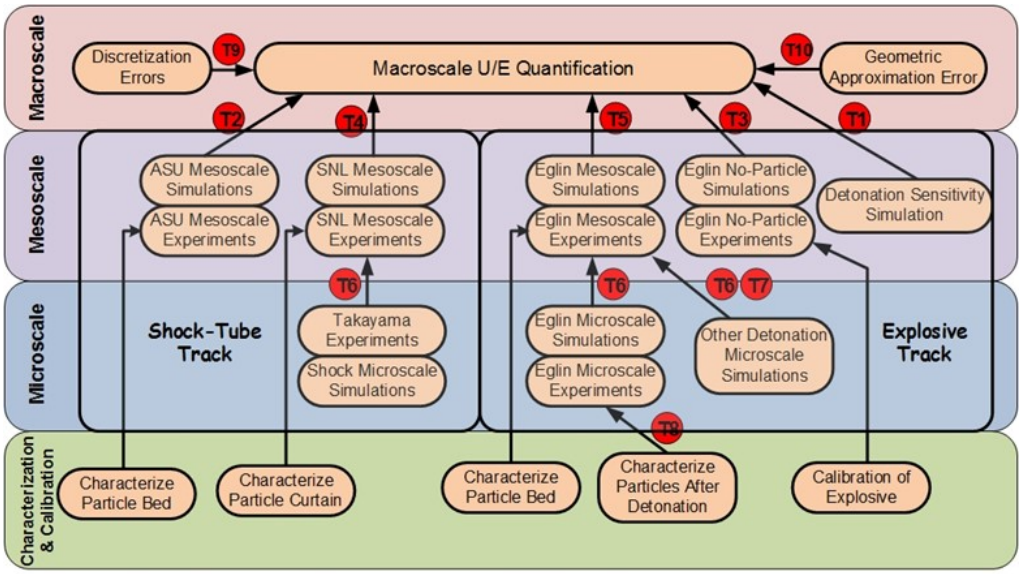

Figure 1.5. Uncertainty Quantification and Propagation Plan.

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.

Finally, Figure 1.6 shows the timeline for performing the different tasks. These tasks T1-T11 were previously described.

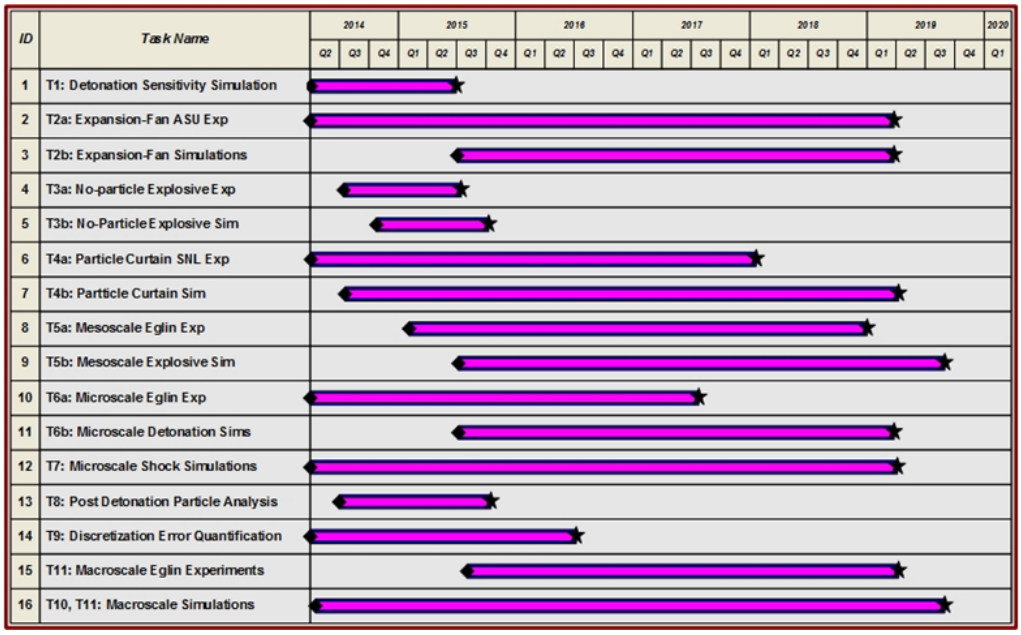

Figure 1.6. Uncertainty Quantification Task Timeline.

2. Macroscale Team

2.1 Overview

The goal of the Macroscale team is to solve a problem of Compressible Multiphase Turbulence (CMT) at an unprecedented level of physical detail and thereby to advance predictive simulation science. The chosen CMT problem for the initial phase of our research activities is an explosive dispersal of solid particles that consists of a PETN charge, which is surrounded by a layer of inert spherical glass particles of about 100μm in diameter (see Figure 2.1). The simulation performed by the Macroscale team will integrate the modeling developments and uncertainty reduction accomplished by the collaboration of the Microscale, Mesoscale, and Uncertainty Budget teams.

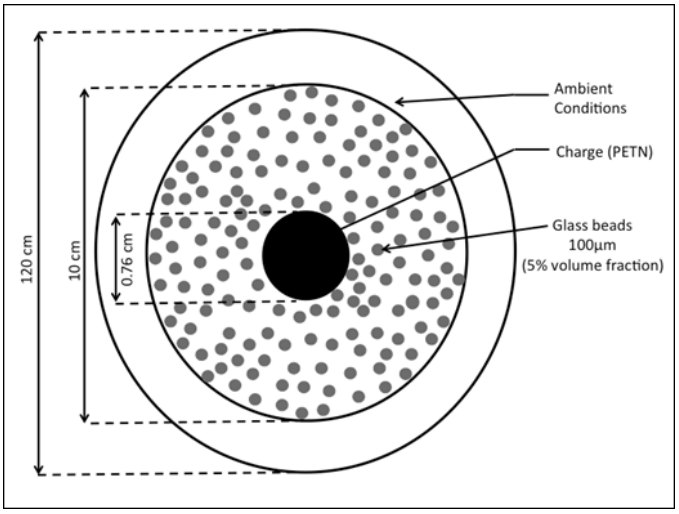

Figure 2.1. Schematic of the two-dimensional computational domain for the explosive dispersal of particles simulations. Not drawn to scale.

2.2 Code Improvement

The problem of explosive dispersal of particle is a very complex problem that requires a robust and versatile code to be able to consistently run simulations, while studying dependency of the problem on initial conditions, and achieving uncertainty reduction on numerous physics models. This problem is inherently multi-scale in space and in time, and, therefore, requires small time steps, a high grid resolution to properly track the fluid as well as a large number of computational particles to accurately capture the particle cloud behavior. Thus, most of the efforts of the Macroscale team during this first year of existence of CCMT have been geared toward preparing the code to operate in "production mode" on the DoE/NNSA supercomputers. Our code, *Rocflu,* developed under DOE ASAP, has been verified and validated against a wide range of complex flow problems. Also, it is well suited to study problems involving high temperature and pressure, such as an explosive problem.

Our first step was to verify that *Rocflu* could, indeed, handle an explosive dispersal of particle problem. Figure 2.2 shows that we have been successful in this task despite simplification in the physics included in the simulation. We were able to run a 3D simulation counting 30 million computational cells and 5 million computational particles for up to 200 microseconds of physical time. Next, we turned our focus on the limitation of the code that could become an issue when we perform large perform production runs, i.e., limitation in the number of processors that the code can handle, limitation in the computational grid size, etc. Numerous tests have been performed, including code profiling and scaling studies. The code has been successfully ported and run on

Vulcan, Lawrence Livermore National Laboratory's BG/Q supercomputer, and Mustang, Los Alamos National Laboratory's supercomputer. Our thorough evaluation of the code has allowed us to identify, and mitigate a memory and efficiency bottleneck in the initialization phase of our simulations as shown on Figure 2.3. Following this improvement in the initialization of *Rocflu*, and later on, in the communication tags in the MPI initialization routine, we have been able to successfully run a test with 16K processors on Mustang for the case presented on Figure 2.2. Currently, the size of the computational grid and the code scalability is the only constraint on the choice of the number of processors, *a priori*. More tests are in preparation to evaluate *Rocflu*'s limitations in terms of scaling.

On another front, we have significantly improved the processing of the simulation results by augmenting the code with the capability to generate, on-the-run, files directly readable by high performance computing parallel visualization software such as ParaView (developed by LANL, SNL and others) and VisIt (LLNL). With this improvement, we can track the progresses of our simulations as it is running and perform early analyses. This constitutes our first step towards co-processing for extreme scale production runs.

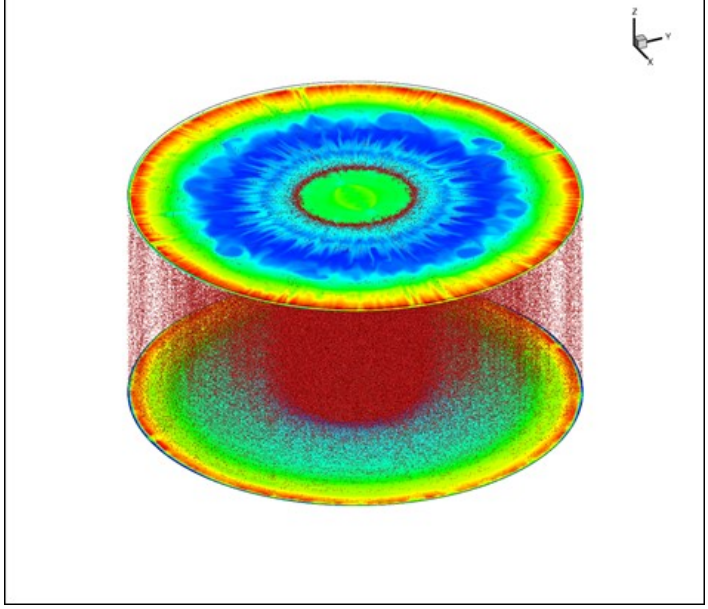

Figure 2.2. Density distribution at early stage of our simulation of an explosive dispersal of particle. The color scale goes from royal blue, low density, to red, high density. Computational particles are shown as red spheres.

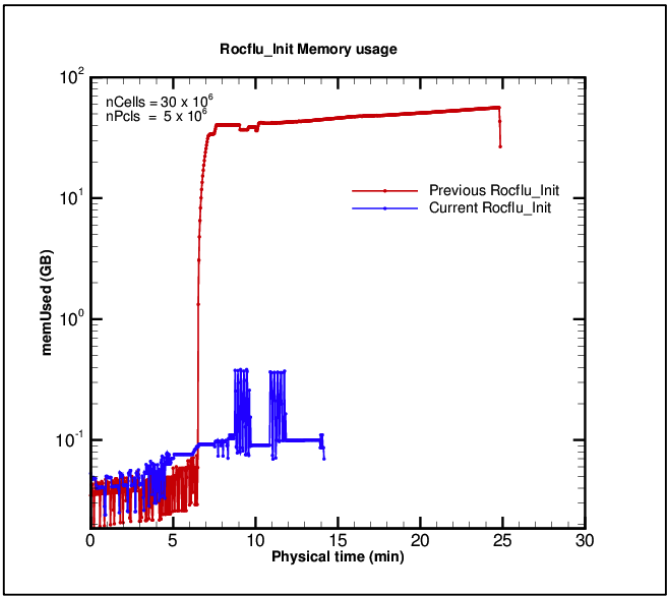

Figure 2.3. Memory usage and time required by the former Rocflu initialization (red) vs. new Rocflu initialization (blue).

As a final note, we have also made tremendous progress on our algorithm for shock and particle front extraction from the experimental videos. These algorithms will be used to process the results of the experiments to be run at Eglin Air Force Base. These extracted quantities serve as metric for assessing the quality of our simulations, and will help quantify the uncertainties due to numerical discretization and modeling.

2.3 Physics of Explosive Dispersal of Particles: Initial Conditions Effects

The problem of homogeneous explosions into ambient particle clouds involves many complex intertwined physical phenomena, including multiple occurrences of Rayleigh-Taylor (RT) and Richtmyer-Meshkov (RM) instabilities. The Macroscale team was interested in numerically investigating the effects of initial perturbations on the gas and the particle phase during the early moments of an explosive dispersal of particles for initially low volume fraction of particles.

In one case, only the density of the high-energy products is disturbed with the addition of a single mode perturbation in the azimuthal direction while maintaining the particle volume fraction unchanged. In contrast, in the second case, the particle volume fraction (surrounding the high-energy material) is disturbed with the addition of the same single wavelength perturbation as in the previous case, and no change is made in the high-energy material.

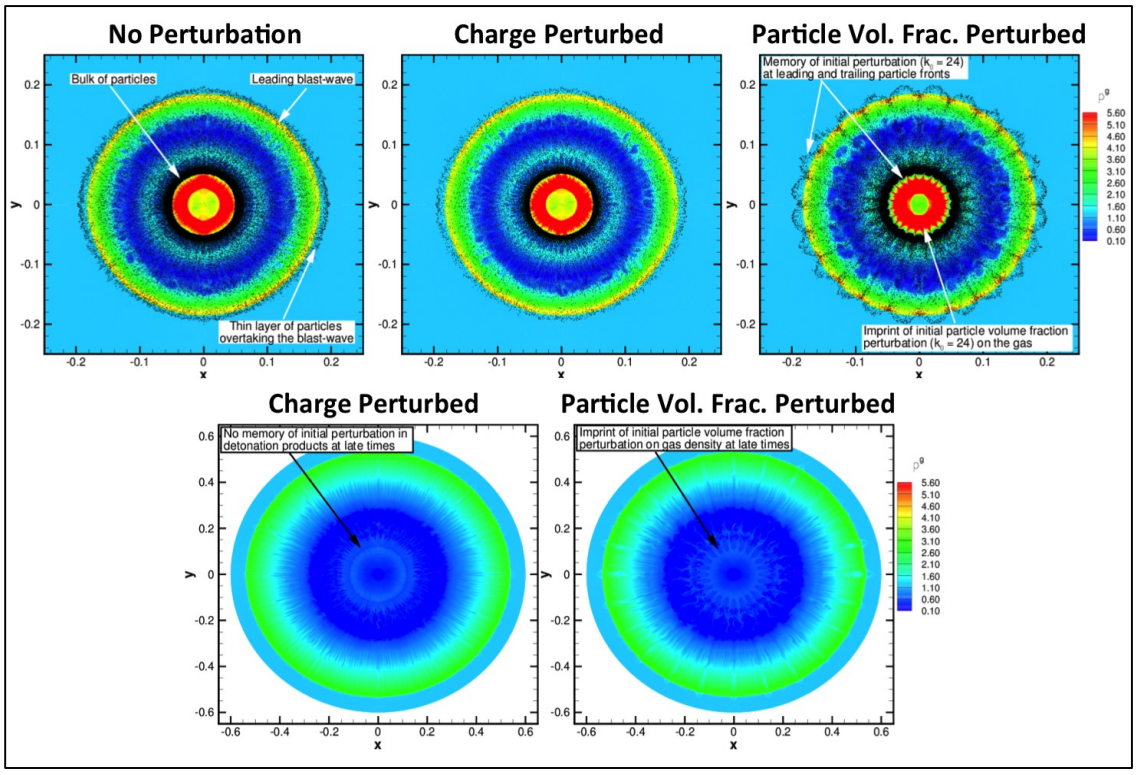

Figure 2.4. Summary of the effects of an initial single wavelength perturbation in the azimuthal direction in the charge vs. in the particle volume fraction.

As summarized by the set of images composing Figure 2.4, our simulations suggest that a single wavelength perturbation in the rapidly expanding detonation products does not have a significant impact on the early evolution of neither the gas phase nor the particle phase. In particular, the Rayleigh-Taylor instability occurring in the gas phase shows no notable difference with the unperturbed case, i.e. the instability displays features characteristic of a low-amplitude broadband initial perturbation. Equivalently, the particle volume fraction does not display any distinctive structure or spatial organization. In contrast, the equivalent single wavelength perturbation in the initial particle volume fraction distribution lingers for the duration of our simulations. Indeed, the innermost front of the particle bed volume fraction shows typical instability fingering at a wavelength equal to the wavelength of the initial perturbation in the particle volume fraction. Similarly, the hydrodynamic instabilities developing in the gas phase exhibits the same dominant wavelength.

2.4 Collaborative Effort

The Macroscale team has been working in close collaboration with the Uncertainty Budget (UB) team. The first and most important step in this collaboration was about education. The Macroscale team has been educating the UB team members on the physics of Compressible Multiphase Turbulence, and in particular on explosive dispersal of particles, on shock tube experiments, and on numerical simulations of these phenomena. In turn, The UB team has been educating the Macroscale team on uncertainty quantification and budget techniques, and methods. This effort continues on a weekly basis, and is crucial to our progresses in uncertainty reduction.

Next, the Macroscale and UB teams have coupled and successfully tested *Rocflu* and the uncertainty quantification scripting library developed by Sandia National Laboratory, DAKOTA. Uncertainty quantification studies can now be performed Mustang (LANL) and Vulcan (LLNL). Such study will take place in year two.

Also, the Macroscale team has been in close collaboration with the Microscale team, in particular regarding the improvements made to *Rocflu*. Current development of our code is now in complete collaboration with the Microscale team, so that each improvement benefits to all *Rocflu* users.

Finally, the Macroscale team is interacting on a regular basis with the experimental team at Eglin Air Force base, and with the experimental team at Arizona State University. Discussions on future experiments and their diagnostics are the main topic of discussion, along with early experimental test results.

2.5 Plan for Year Two

Regarding improvement to *Rocflu,* the emphasis in year two will be two-fold. First, the Macro- and Micro- teams will be integrating full co-processing capabilities to Rocflu with the introduction of Catalyst, following a suggestion of Greg Weir (Sandia) of the AST team. This new capability will allow for increased productivity for “hero” runs as well as bundled UQ runs. Second, a list of new physics model has been identified, and will be added to *Rocflu.* It includes a more appropriate

equation of state, new volume fraction computation, a new flux scheme among other improvements.

Regarding the physics of explosive dispersal of particles, the Macroscale team's exploratory study has paved the way for more detailed investigations. First, one needs to investigate the physical ground of the particles riding on the blast wave as seen on figure 4. Then, a quantitative study of the impact of the designed initial conditions should be conducted. A thorough account of the mixing and turbulence properties associated with carefully designed single and multimode initial perturbations shall be studied.

Finally, the collaboration with the Microscale team will strengthen in year two as progress will be made on integrating the co-processing tool. Similarly, the collaboration with the UB team will continue as *Rocflu* can now operate in production mode and new experimental results will be available.

3. Microscale Team

3.1 Overview

The purpose of the microscale simulations is to develop point-particle (PP) and other related models that are needed for the mesoscale and macroscale simulations of explosive dispersal of particles. The focus is limited to non-reactive particles. These particle models are required to compute (a) instantaneous aerodynamic force on the particle and (b) instantaneous net heat transfer between the particle and the surrounding. The commonly employed standard models for drag and heat transfer have been derived under simplifying assumptions. The objective here is to develop extended models in finite Mach, Reynolds, and volume fraction regimes – to be used in meso and macroscale simulations. A strategy for a sequence of microscale simulations has been devised that will allow systematic development of the hybrid surrogate models that are applicable at conditions representative of the explosive dispersal application. The ongoing microscale simulations seek to examine particle force dependence on: (a) Mach number, (b) Reynolds number, and (c) volume fraction (different particle arrangements such as cubic, FCC, BCC and random).

3.2 Year 1 Research

For the microscale simulations, two compressible flow solvers are being employed. *Rocflu* is a parallel finite volume solver (developed under the ASAP program) for unstructured, body-fitted meshes. It can solve inviscid or viscous compressible flows and is being used for microscale simulations for non-deforming particles to study shock-particle and contact-particle interactions including effects of Reynolds number, Mach number, volume fraction. *RocSDT* is a compressible code based on level set methodology on Cartesian grids and can handle shock-particle interactions, multi-material interfaces, material deformation, general equations of state, and chemistry. It complements *Rocflu* in two ways: (a) it allows particle deformation, and (b) it allows more flexibility for arbitrary/random particle clusters.

Grid generation and determining the adequacy of grid resolution for the microscale simulations is an important task. As a first step, comprehensive 2-D and 3-D mesh generation capabilities were developed for use by all CCMT members. Microscale simulations at the CCMT require an assortment of 2-D and 3-D grids as part of our effort to model the forces experienced by particles in high speed flows. Meshing capabilities for handling arbitrary particle shapes, sizes, locations, and numbers were developed. GRIDPRO is being used to generate particle surface meshes which is then used as input to TetGen to generate 3-D unstructured grids. A filter has been written to convert TetGen output to Rocflu-compatible format. Furthermore, a series of simulations was performed to study the effect of surface mesh resolution and global mesh resolution. While generating the grid, one can either control the global element size, which in turn will generate a conforming surface mesh on the particle, or control global and surface element size. Based on these simulations, two grid generation criteria for particle simulations have been developed: (a) a global mesh criterion (specifying a uniform mesh volume for each element or a maximum volume constraint as a function of particle diameter, and (b) a particle surface mesh criterion (specifying the area of each element on particle surface as a function of particle diameter).

These criteria are used to generate the single- and multiple-particle grids employed in the ongoing 3D microscale simulations.

On the particle force modeling front, a methodology has been developed by CCMT personnel to extract an inviscid force kernel on particles in uniform flow by employing an impulsive particle force (delta acceleration) to a single spherical particle – very good agreement with analytical and published results was demonstrated in the subcritical Mach number range (M=0 to 0.6). The goal of the current microscale simulations is to extend this methodology to higher Mach number regimes. To this end, single particle simulations at Mach numbers greater than 0.6 have been conducted. A generalized shock-particle model developed by CCMT personnel was also

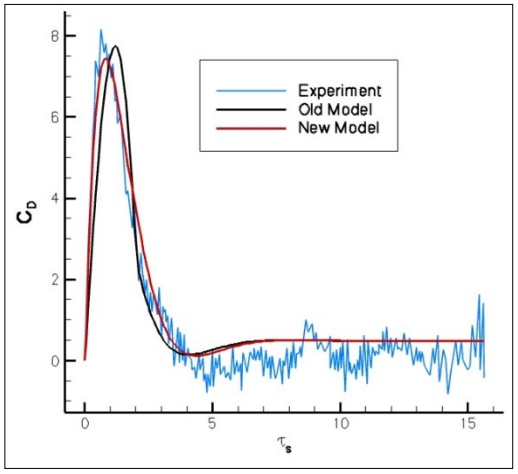

Figure 3.1. Shock-particle interaction modeling.

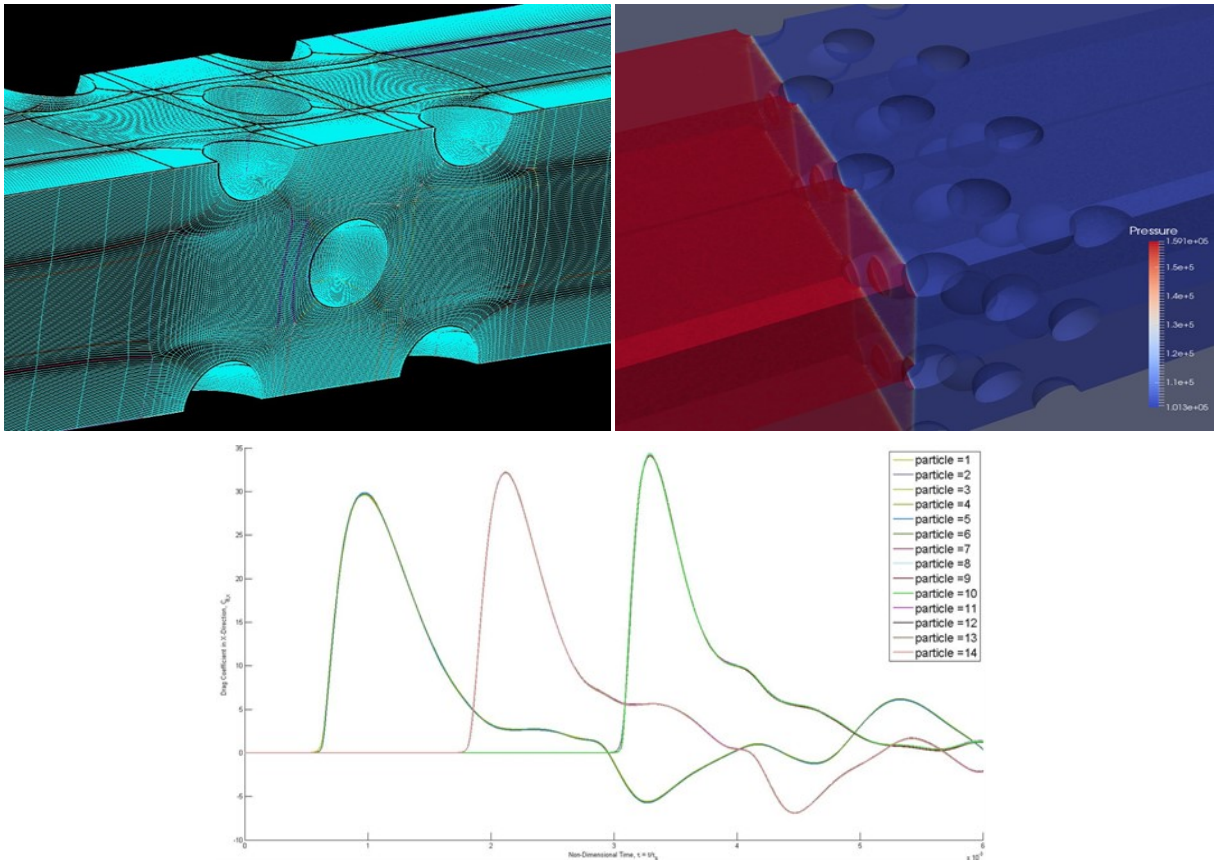

Figure 3.2. 3-D multi-particle simulations using FCC particle arrangement. A single FCC unit cell simulation contains 14 particles (volume fraction 10%). Top left: grid. Top right: Pressure distribution as a shock passes over the particles. Bottom: drag history on the 3 planes of particles.

implemented to test the results of these microscale simulations. Figure 3.1 shows a plot of force experienced by an isolated and fixed particle, when a shock wave impinges on it. Experimental results are from Sun et al. Models labelled "Old" and "New" have been proposed by CCMT personnel. The new model (which currently does not include viscous unsteady effects) shows a good agreement with the experimental data for early times. Ongoing work seeks to include viscous unsteady effects and to develop a model which can predict force on moving particles. Furthermore, an investigation of a contact surface-particle interaction has also been initiated.

Another focus of the microscale group is the investigation of the effect of finite volume fraction on particle force histories by performing simulations of clusters of particles interacting with shocks and contact surfaces. This requires large 3-D grids and necessitates large scale parallel computations. To this end, the microscale team members interacted with the computing resources teams at LLNL and LANL to successfully compile the Rocflu code on Vulcan, Cab, and Mustang clusters. Simulations using Rocflu are currently being performed at the LLNL machines. As an example, sample results of a 3-D multi-particle simulation using an FCC particle arrangement are shown in Figure 3.2. A single FCC unit cell simulation containing 14 particles (corresponding to volume fraction of 10%) is subjected to a normal shock. A view of the 3-D mesh is shown on the top left of the figure and the pressure distribution as the shock passes over the particles is shown on the top right. The bottom plot in Figure 3.2 shows the drag histories on the 3 planes of particles.

For the 3-D shock-particle investigations, it was recognized that the grid resolution study for single-particle simulations should be formalized to produce numerical error estimates as a function of flow field conditions. To get an estimate of the peak integrated force on an infinite resolution grid we are using a Richardson extrapolation procedure. This process is executed by making a series of three or four simulations using the same conditions on grids that have finer and finer elements in the domain. An extrapolation fit can be obtained by using the results from three of the four simulations. The fourth simulation can be used to estimate the predictive error of the model (the Richardson extrapolation fit of the data). Alternatively, a non-linear regression can be performed using all of the simulation results to try and best estimate the parameters of the Richardson extrapolation. Performing this procedure will give us an idea of how far away the simulations are from the grid independent solution. This information can also be used to determine an estimate of the uncertainty in a simulation that does not use a fine enough grid for grid independence as would be the case where the computational cost of adding the required number of elements would be prohibitively expensive in terms of time and processors used. Additionally, a similar approach will be adopted for multiple-particle simulations.

3.3 Plan for Year Two

A plan for a suite of simulations using Rocflu for 3-D multi-particle simulations has been developed. These simulations will fill a design space of Mach number, volume fraction, and particle packing type as depicted in Figure 3.3. The volume fraction will be varied from 10 percent to 40 percent to encapsulate the range between dilute and dense particle packing. A series of three-dimensional meshes has been generated to investigate the design space shown in the diagrams. Sample meshes for FCC, BCC and random particle distribution are shown in Figure 3.4.

Simulations for a combination of four Mach numbers and four volume fractions are currently being conducted for an FCC arrangement of particles. Following this, BCC and random particle arrangements will be investigated. Future plans include investigation of sequences of fully-resolved microscale simulations consisting of an array of particles subjected to more realistic time-dependent flows that progressively better approximate the actual problem of explosive dispersal.

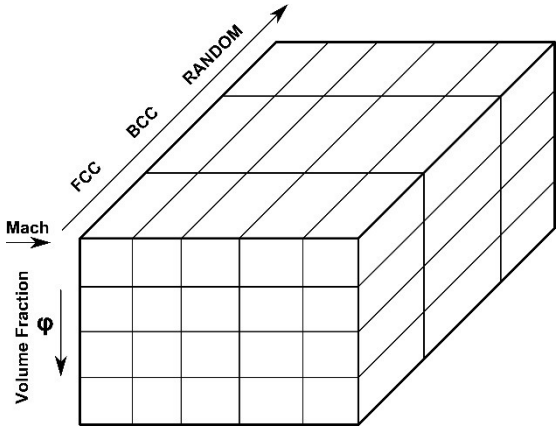

Figure 3.3. Design space of Mach number, volume fraction, and particle packing type (FCC, BCC, random).

Microscale simulations of multi-material shock particle interaction problems using *RocSDT* have also been conducted to include the effects of particle deformation, initial shape and volume fraction (with random packing) in particle force models. The methodology in RocSDT allows freedom in choice of particle shape, size, and number in simulation as well as the ability to understand the transient particle deformation dependence on various parameters including: (a) particle material, (b) medium material, (c) multiple particles, (d) incoming shock pressure and speed, (e) medium to particle impedance ratio, (f) particle shape and orientation to shock, etc. Studies conducted thus far include the effect of particle deformation as well as presence of multiple particles on particle drag histories during shock-particle interactions. The first part of this study investigated the transient drag coefficient experienced by an aluminum spherical particle under various shock-loading conditions. Under these conditions, the aluminum particle was essentially non-deforming. It was found that the peak drag coefficient decreased with increasing incident shock Mach number. In addition to this, linear arrays of five aluminum particles with varying inter-particle spacing were simulated (see Figure 3.5). Post-shock pressure multiplication was found to occur as the incident shock wave passed over a particle, which would cause higher peak drag coefficients for the subsequent particle. The pressure multiplication was found to saturate after the fourth particle. The inter-particle spacing was varied to determine its effect on the peak drag

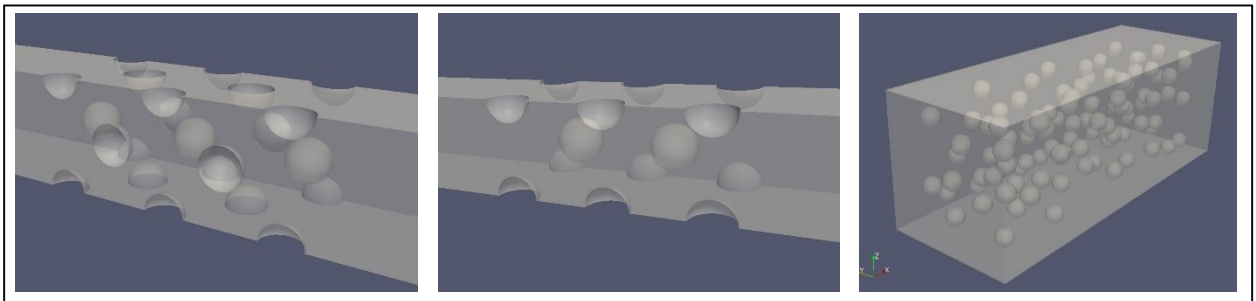

Figure 3.4. Sample grids for 3D multi-particle simulations using: (a) FCC arrangement, (b) BCC arrangement, (c) random particle distribution.

coefficient on the fifth particle. The drag coefficient of the fifth particle is normalized with the drag coefficient of the first particle. It was found that there was an inter-particle spacing where the magnitude of pressure multiplication reaches a maximum. The effect of inter-particle spacing is shown in Figure 3.5(b). The horizontal axis represents the non-dimensional ratio of particle diameter to inter-particle spacing, while the vertical axis represents the percent increase of peak drag coefficient felt by the fifth particle with respect to a single particle. The peak drag coefficient of a single particle was found to be 7.57.

Ongoing and future work using *RocSDT* will – through incorporation of Mie-Gruneisen equation of state and stronger incident shock conditions – allow the aluminum particle to deform. The initial particle shape will also allowed be vary between a sphere and an ellipsoid. The effect of initial particle shape and deformation on the transient drag coefficient curve will be investigated. With the aforementioned changes in mind, the study will include investigation arrays of particles.

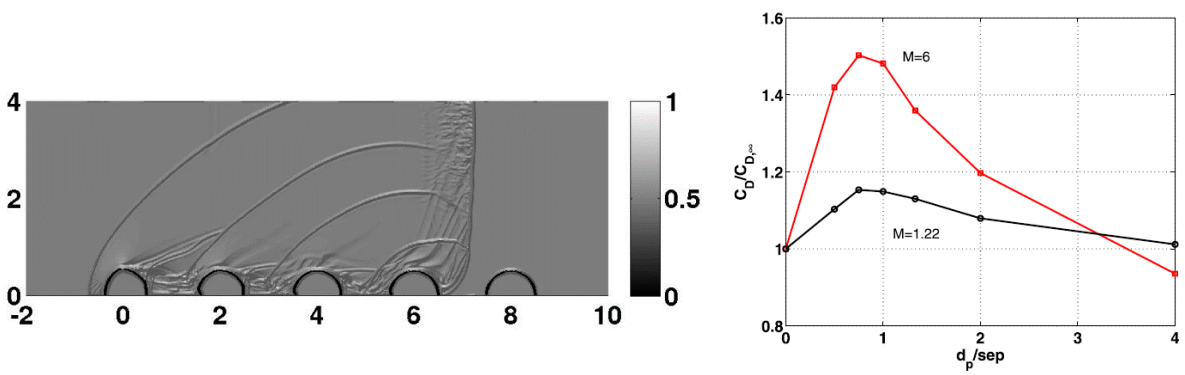

Figure 3.5. Five-particle shock interaction simulation: (a) contour plot of numerical Schlieren for a representative simulation, and (b) effect of inter-particle spacing on peak drag coefficient.

4. Experiments

4.1 ASU Experiments

4.1.1 Goals and Motivation

Multi-phase studies involving compressible flow are very complicated. The gas phase may move separately from solid particles. The particles may affect the gas flow and the particle structures may also be highly influenced by the gas phase. In addition, it's possible for the particles to generate turbulence.

The computational model required to describe this type of flow must be highly complex. There is a great need, especially in the early stages of the model development, for a simple experiment that can be used for early validation of the computational codes developed by the PSAAP center. This experiment will involve simpler physics than the spherical charge capstone experiment.

The experimental team at Arizona State University proposed a simple shock tube experiment with a particle bed at the base. The shock tube will be initially separated into two regions: a lower region at atmospheric pressure (~101kPa) where the particle bed is contained, and an upper region close to vacuum (~1kPa). When the two sections merge, a shock wave is formed and the particle bed will expand upward into the lower pressure region.

The idea of designing a simple experiment like this one, is that we can examine the real physics of the flow. Dr. Kirsten Chojnicki's Ph.D. research at Arizona State University involved looking at shock velocity, mixture velocity, and the particle drag coefficient, in a similar shock tube experiment to the one proposed by ASU. In addition to looking at these quantities, we'd like to examine the gas and particle velocity fields, and the particle-gas interfacial instabilities.

By carefully designing the experiment and doing many tests, we believe we can reduce the scatter seen in Dr. Chojnicki's data. We can do this by more precisely controlling particle properties, including the size distribution. Usually, when labeling the diameter of the diameter of the particles, the expected variation is as high as plus or minus the mean. We can decrease the size distribution by suspending the particles in a fluid and using fractional separation—or fractionation.

We are also interested in incorporating Schlieren imaging into our experiment. Schlieren experiments use a point source of light to illuminate fluid flow. As the pressure changes throughout the flow, the density changes—meaning the index of refraction also changes. In Schlieren experiments these inhomogeneities in density can be explored by examining the differences in dark regions and bright regions appearing on the camera images.

4.1.2 Description of Experiment

The basic profile of the shock tube will be a six foot tall glass tube with a six inch by six inch square footprint. A particle bed will be placed at the bottom of the shock tube. The particle bed will be composed of small glass spheres. The glass spheres will have diameters between 90 and

150 microns. A mylar diaphragm will be used to separate the high pressure (101kPa) from the low pressure region (1kPa).

The experiment was designed in this way to maximize optical access to the flow. The wide windows will allow for particle image velocimetry measurements and Schlieren measurements. Using optical techniques, we can measure the contact line velocity, the gas velocity, and the particle volume concentration. We can also investigate the particle bed interface and how it evolves in time. The main parameters that will be used to vary the experiment will be the initial pressure ratio (typically referred to as p4/p1) and the particle diameter.

The diaphragm separating the low pressure region from the high pressure region will be put well above the particle bed. Please see Figure 4.1. This will ensure the diaphragm does not interfere with the expansion of the particle bed. The sidewalls of the shock tube will have many PCB Piezotronics Model 113B22 pressure sensors. They have a sensitivity of +/- 10% and output 0.145 mV/kPa. Using these pressure sensors, we will be able to track the shock wave as it travels upward into the test section. We will also be able to examine the pressure fluctuations along the length of the shock tube, including the particle bed, during the experiment.

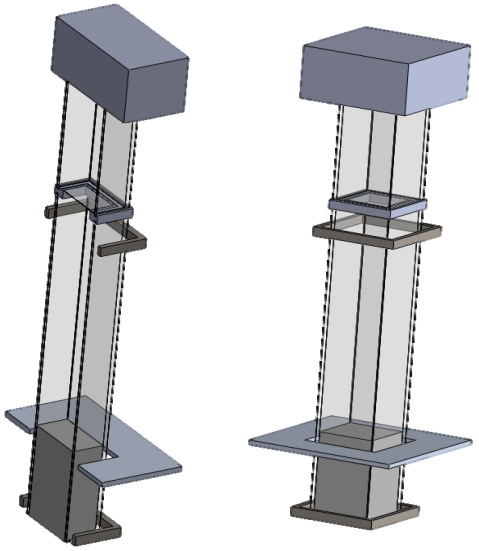

Figure 4.1. Basic design of the 6' tall vertical shock tube. It has a large footprint of 6" x 6". A vacuum chamber is located at the top and the particle bed will be at the base. The diaphragm separating the low pressure region from the high pressure region is located about one-third of the way down from the top of the shock tube. It is well out of the way of any of the particle bed motion.

4.1.3 Accomplishments during FY 2014

Heather Zunino has completed several experiments on a shock tube (from Dr. Amanda Clarke's lab at ASU) similar to the one she will build to verify theoretical models for particle behavior in a suddenly expanded environment. From these experiments, she has obtained voltage measurements and high-speed video data for $\frac{p4}{p1}$ ratios of ~10-20.

In Figure 4.2, three pressure traces are shown. They are labeled according to the distance between the associated pressure sensor and the diaphragm. In Figure 4.3, the pressure traces from two different experiments are compared. This data shows several interesting events:

- There is a small bump in the trace for the pressure sensor located 33cm above the diaphragm (blue), which is associated with the diaphragm burst.

- The shockwave is unsupported—note the pressure decreases rapidly after the initial jump. In this experiment, there was a very small air gap between the particle bed and the diaphragm. With larger air gaps, pressure behind the shock wave does not decrease as rapidly.
- The pressure traces located at 33cm and 65cm from the diaphragm show obvious additional sharp jumps after the initial shock front passes. These are caused by a reflected shockwave traveling downward from the top of the shock tube. Notice the pressure sensor located at 97cm has two apparent peaks. The second peak is the reflected shock as it begins to descend down the tube. The reflected shock decreases in energy as it travels past the pressure sensors located at 65cm (orange) and then 33cm (blue).

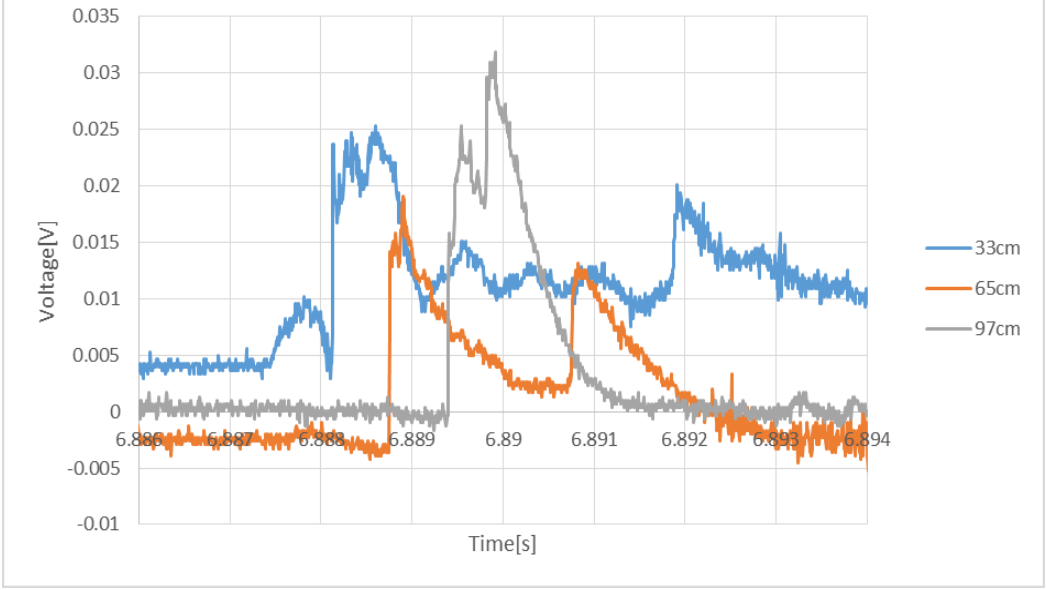

Figure 4.2. Pressure data [V] from shock tube experiment.

After doing experiments on the shock tube in Dr. Clarke's lab, there have been numerous design alterations to Heather's experiment. The current diaphragm design in Dr. Clarke's shock tube has been problematic and difficult to control. It doesn't burst the same way every experiment and sometimes it fails altogether. For this reason, Heather has decided to alter her proposed design. The diaphragm component of Heather's experimental setup will be a 3D-printed square connection between the driving section and the driven section. There will be a small groove on the top of the 3D-printed connection. This groove will be lined with a thin layer of porcelain. A thin layer of conductive material will be permanently attached to the top of the porcelain. Thin mylar will then be laid over the entire component.

During the experiment, an electric current will run through the conductive material and evenly burn through the mylar diaphragm along the edges. From the experiments performed on

Dr. Clarke's shock tube, it was seen in the high-speed videos that the diaphragm is quickly blown upward into the vacuum tank and seems to minimally effect the flow later on.

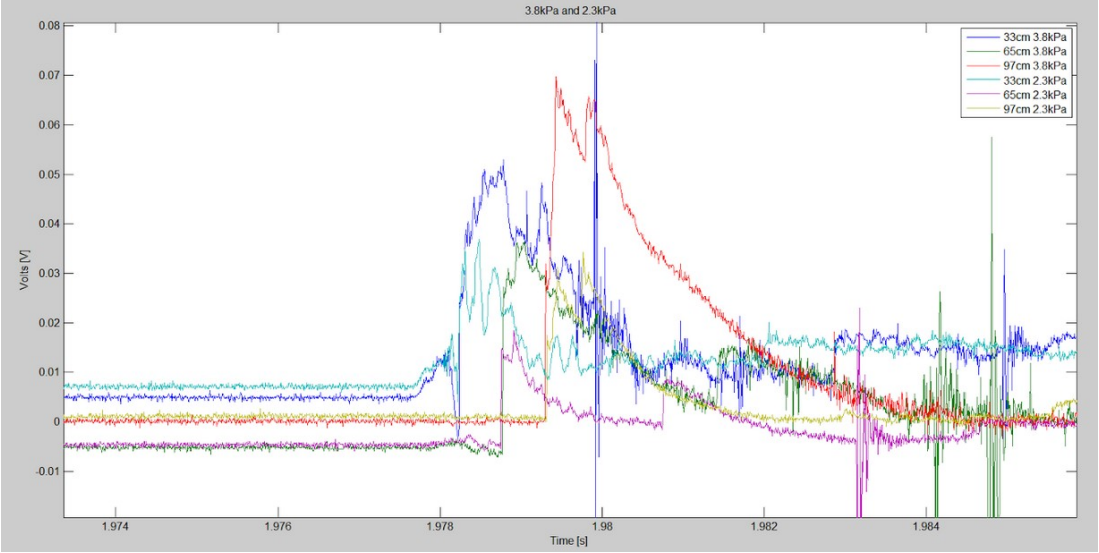

Figure 4.3. Pressure data [V] from two separate shock tube experiments at with P1 pressures at 3.8kPa and 2.3kPa. Notice the shock strength is higher for the lower P1 pressure, but the shock speeds remain similar.

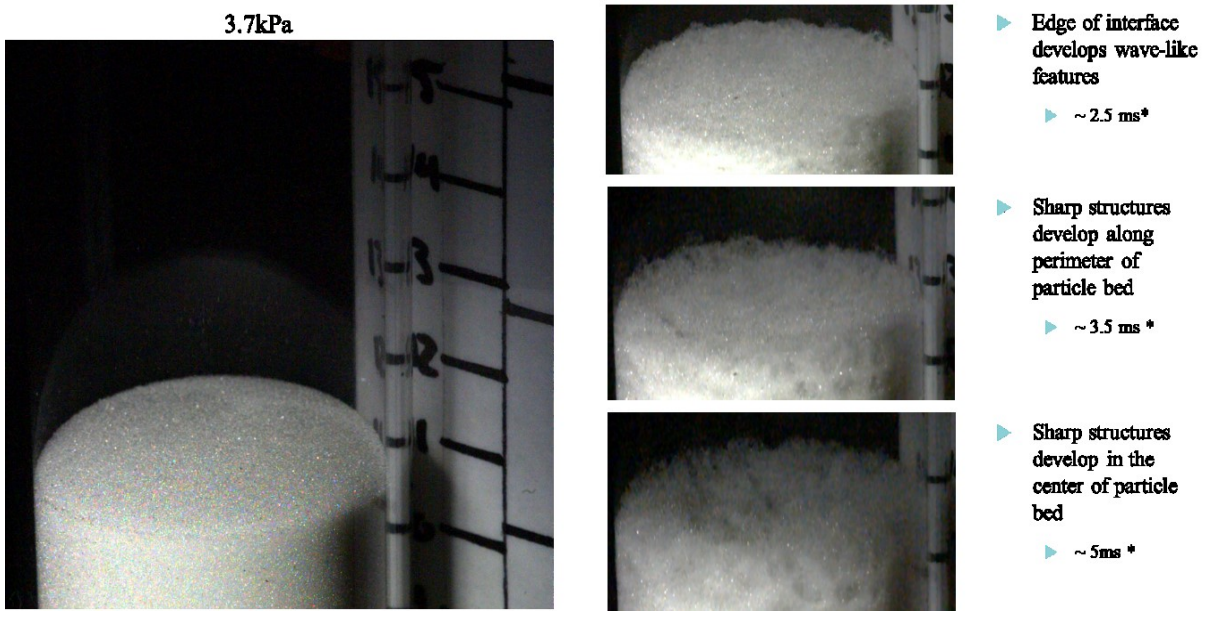

Figure 4.4. Initially, the particle bed interface is very flat. Several stages of the interface evolution are seen. *All times recorded are relative to the moment the particle bed begins to move.

Heather has also taken high-speed videos of the particle bed expanding. She has found that the particle bed interface appears relatively flat during the first stage of expansion (<2ms after the particle bed starts to move). Around 2.5 milliseconds after the first movement in the particle bed is detected, the edges of the particle bed interface appear to have a sinusoidal wavelike structure—see Figure 4.4. Approximately one millisecond later, the wavelike structures transform into more sharp, peak-like structures along the edges of the particle bed. After another 1.5 milliseconds, these sharp structures have spread into the middle of the particle bed.

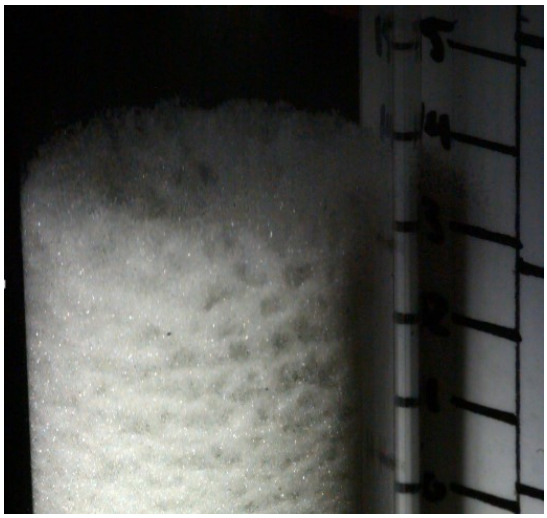

Figure 4.5. Semi-homogeneous air cells appear in the expanding particle bed. These cell-like structures begin to appear around 3ms after the particle bed begins to move and continue to form throughout the experiment.

While the interface is evolving, the body of the particle bed is also changing. A pressure drop has been measured by examining the first signs of particle movement in the particle bed. The pressure drop appears to be traveling down into the particle bed at approximately 33m/s when the initial pressure ratio is $\frac{p4}{p1} = 27$. After this pressure drop travels through the particle bed, well-defined air cells appear, semi-homogeneously. These cells are particle-vacant regions—see Figure 4.5.

It is possible these cell-like structures appear because the glass beads are over 2,000 x's heavier than the air mixed in with the particle bed—so, the air accelerates much more quickly. The air will travel the path of least resistance – through small vacancies in the particle bed, and then expand further – pushing the glass beads.

The cells were a point of particular interest, and further experiments were performed to explore this phenomena. It turns out that the cell distribution and size homogeneity may have a particle diameter distribution size dependence—please see Figures 4.6 and 4.7.

4.1.4 Plan for Year Two

During the next fiscal year, Heather will continue to do experiments on Dr. Clarke's shock tube at ASU. She will examine the shock strength as a function of the gap between the diaphragm and the particle bed. She will also continue analyze the high-speed video data.

Additionally, the materials to build the 6' vertical shock tube will be acquired. As the budget is limited, it is important to finalize the designs before acquiring materials. To help with this, the computational team and the experimental team at ASU have begun to meet monthly via Skype. This has helped align the goals of the computational team with the experimental team. Several additions to the shock tube design have been made since the meetings began, including

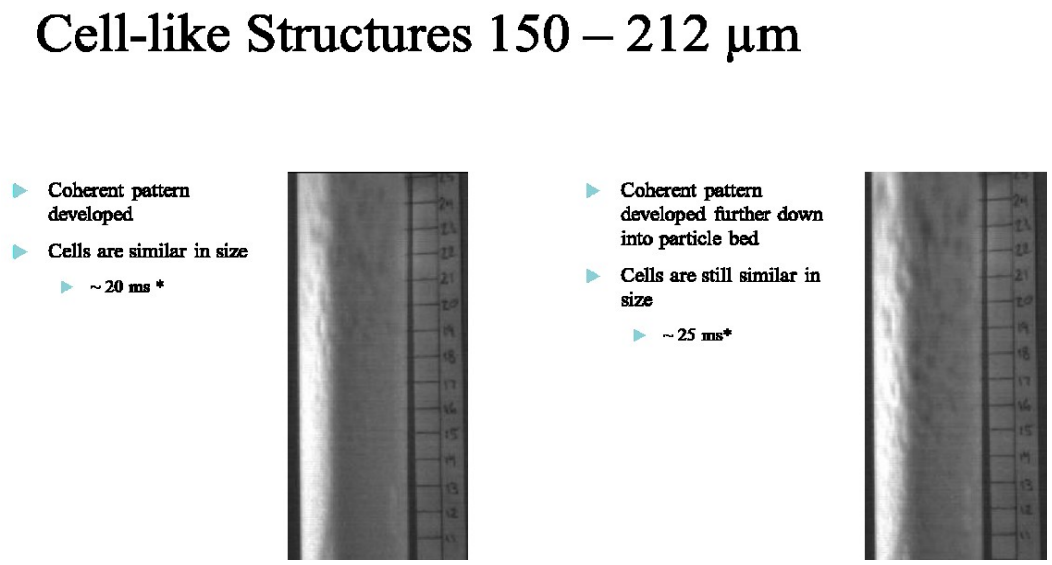

Figure 4.6. Air cells, similar in size, propagate into the depth of the particle bed. They appear to be distributed evenly. The particle diameters in this experiment have a Gaussian distribution between 150-212 microns. *All times listed are relative to the first apparent motion in the particle bed.

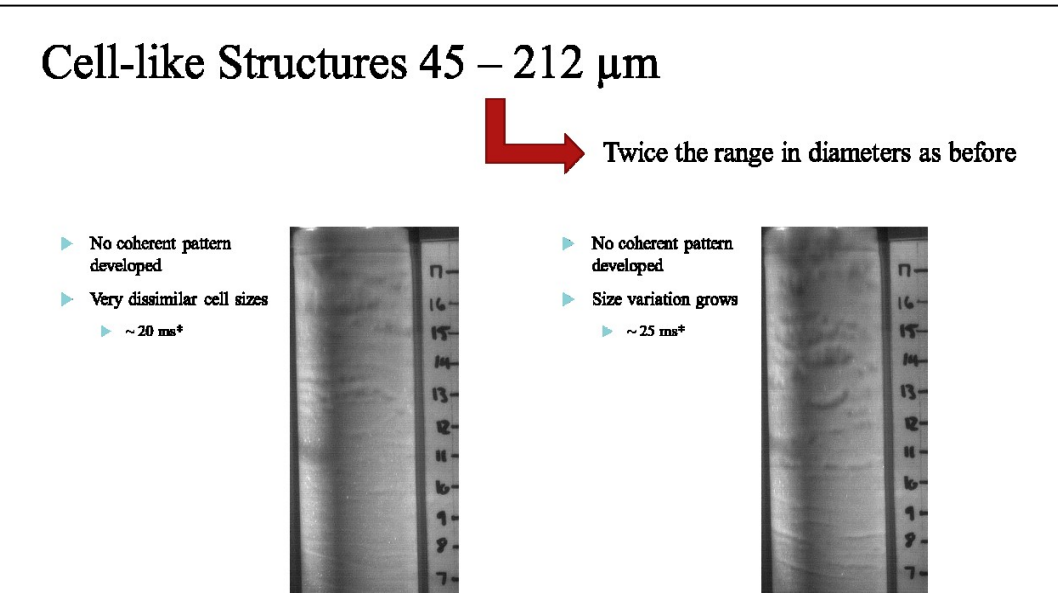

Figure 4.7. When the particle diameter distribution is nearly doubled in size and poorly mixed, the air cells no longer appear similar in size or evenly distributed laterally or horizontally. *All times listed are relative to the first apparent motion in the particle bed.

additional pressure sensors to test the variations in pressure at a single location along the length of the shock tube.

These meetings will continue for the duration of the project. To obtain a better understanding of the mechanisms at work, Heather will complete some trade studies to find out which high-speed video camera would be best for this project. The early stage gas-particle activity at the surface of the particle bed is very important. It evolves at such a short time scale—so, it is important to have a fast frame rate. However, there are significant budgetary restraints. Heather will test several cameras from Vision Research to make sure the camera the ASU experimental team purchases will provide the highest quality images at the fastest frame rate and will remain within the budget.

4.2 Eglin AFB Experiments

4.2.1 Goals and Motivation

The primary goals of the experiments being conducted at Eglin Force Base is to provide validation quality data at the micro, meso, and macroscales. Experiments are currently being carried out at the micro and mesoscales, while experiments at the macroscale will start in Year 3. Figures 4.8-4.11 highlight the overall Y1+ experimental plans and Y1 accomplishments.

- Year 1+: Microscale experiments
  - Millimeter-sized particles; single/few particles; planar geometry
  - Controlled experiments where
    - a few well-characterized finite-sized metal particles are placed outside a well-characterized explosive in a precise manner
    - particles embedded inside the explosive interact with the detonation wave and the post-detonation flow
    - complex particle arrays (*e.g.*, stacked particles or spaced particles) are embedded in a frangible, inert matrix material that is impedance-matched to the explosive
- Year 2+: Mesoscale experiments
  - 10-100 μm particles; $>10^3$ particles; planar geometry
- Year 2+: Macroscale experiments
  - 10-100 μm particles; $>10^3$ particles; cylindrical geometry

Figure 4.8. Experimental plans for Eglin AFB covering micro, meso, and macroscale problems.

Objectives	Parameters/Diagnostics
Accurate extraction of particle position, velocity and acceleration in the near & intermediate fields	• Position vs time / X-ray images & high speed video – Velocity – derivative – Acceleration – double derivative
Quantify the deformation of the particle	• Soft catch • 3-D scan of deformed particles
Extraction of shock waves and flow field in the region of the particles in the near & intermediate fields	• Light transmission / high speed video with strong back-lighting • Fireball temperature / Fourier Transform Infrared (FTIR) video • Shock waves / piezoelectric pressure transducers
Uncertainty quantification	• Repeat selected experiments

Figure 4.9. Objectives for microscale experiments for Y1.

- Well characterized explosive (precision explosive charges)
  - Composition PBXN-5 explosive
    - High performance explosive
    - Pressed 0.5” OD x 0.5” L pellets for good density control
    - L/D=~3 charge (stacked pellets with interface control)
      - Sufficient length for steady-state detonation (> DDT length)
      - Minimal explosive charge for better near-field diagnostics
  - 2” OD mild steel case. Heavy radial confinement ensures:
    - Fixed boundary conditions
    - Near-planar detonation waves
  - RP-83 Exploding bridge wire (EBW) detonator
  - Well-characterized metal particles (spheres)
  - Tungsten alloy (ρ=17g/cc, diameter = 2mm)

Figure 4.10. Y1 Experiment Design Goals & Approaches.

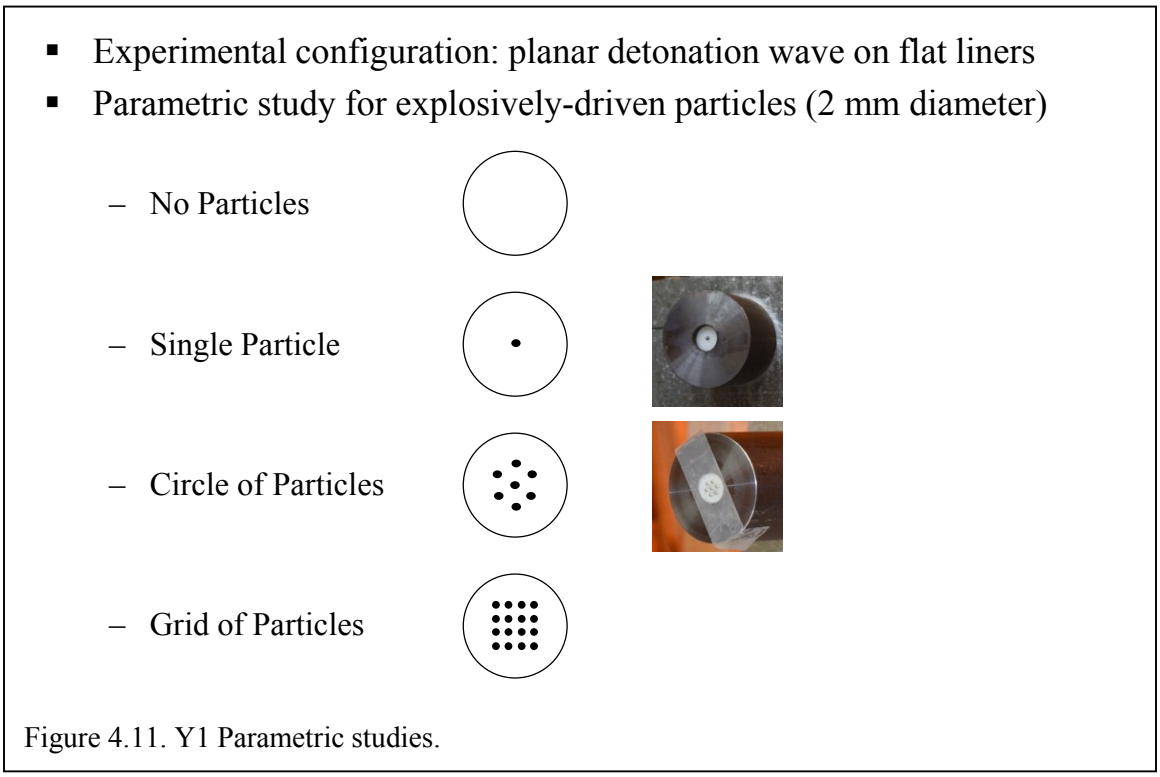

Figure 4.11. Y1 Parametric studies.

4.2.2 Experimental Plan

A detailed analysis of all errors and uncertainties in the macroscale simulation of the demonstration problem has revealed various sources of errors and uncertainties. The Center will use a combination of micro and mesoscale experiments and simulations to quantify these errors and uncertainties and propagate them to the macroscale. These include both shock-tube track and explosive track micro and mesoscale experiments and simulations.

T3: Uncertainties due to thermodynamic (EOS) and transport properties

Purpose: Our simulations will make several assumptions: (a) use program burn model for detonation, (b) use JWL equation of state for detonation products, (c) use ideal gas law for ambient air, (d) assume the stainless steel encasing of the explosive to be non-deformable, (e) assume perfect axisymmetry. In this task we evaluate errors involved in these assumptions (or idealizations).

Strategy: We will perform no-particle detonation-only experiments (or with very few particles, so that detonation gas flow will remain unaffected). Compare against gas-only simulations. Differences will highlight gas-only modeling errors.

Desired experimental features:

- Geometric axisymmetric
- Very strong stainless steel encasing, which undergoes minimal deformation
- Well known explosive whose energy release and EOS of products are known

Desired experimental measurements for comparison:

- Pressure time histories at a number of well-characterized spatial locations
- High speed images of the expanding detonation products into shock-heated air; information 3-D structure with multiple cameras will be useful
- Potential quantitative measurements of this expanding gas front

Parametric Variation:

- None. If the explosive to be used changes, then this calibration check must be repeated

Statistical repetition:

- 2 or 3 repetitions will be desirable
- We could repeat the no-particle experiment with a thicker stainless steel casing

T5: Compaction modeling uncertainty

Purpose: Our meso and macroscale simulations will employ a compaction model to describe the very early stages of the intense shock wave propagating over a dense bed of particle. The form of this compaction model and the model parameters to be used are not well understood and therefore will involve significant uncertainty. Here by comparing the mesoscale simulation results with those from mesoscale experimental measurement, we plan to isolate and quantify uncertainties of this model.

Strategy: We will perform mesoscale experiments with a bed of particles. Corresponding simulations will also be performed. By comparing the results we will evaluate uncertainties in the compaction model.

Desired experimental features:

- Geometric axisymmetric
- Very strong stainless steel encasing, which undergoes minimal deformation
- Well known explosive whose energy release and EOS of products are known
- Rigorous characterization of particles – size distribution, deviation from sphericity, amount of particles in the bed

Desired experimental measurements for comparison:

- Pressure time histories at a number of well-characterized spatial locations
- High speed short and long time images of the leading and trailing particle fronts
- Amplitude and growth of instabilities at the leading and trailing particle fronts
- High speed images of the expanding detonation products
- Potential quantitative (PIV) measurements of expanding gas front
- Soft collect the particles to examine potential deformation

Parametric Variation:

- Vary particle size (eg: 100 micron, 50 micron, 200 micron) while keeping total mass of the bed the same – this will be the first set of experiments
- Vary the mass of the bed, while keeping particle size the same
- Start with a uniform (flat) bed, but add patterns in subsequent experiments

Statistical repetition:

- 3 repetitions for a few of the experiments will be desirable

T6/T7: Uncertainty in finite Re, Ma and volume fraction dependent drag and heat transfer

Purpose: Our meso and macroscale simulations will treat particles as just points – the flow details around each and every particle will not be resolved. Therefore we will employ drag and heat transfer models that will depend on velocity and temperature of the particle and the surrounding gas. The question is how good are our models at the Reynolds numbers, Mach numbers and finite volume fractions typical of explosive dispersal. We want to evaluate errors and uncertainties associated with our force and heat transfer models.

Strategy: We will perform microscale experiments with a cluster of few larger size particles, whose trajectory will be tracked more precisely (which is not possible at the mesoscale). The particle position vs time will be used to back out acceleration, from which the force acting on the particle can be calculated. This force vs time will be used to validate the force model.

Desired experimental features:

- Very strong stainless steel encasing, which undergoes minimal deformation
- Well known explosive whose energy release and EOS of products are known
- Precise initial location of all the particles

Desired experimental measurements for comparison:

- Pressure time histories at a number of well-characterized spatial locations
- Short time x-ray images of particle position vs time
- Long time trajectories of individual particles (which particle went where)
- High speed images of the expanding detonation products
- Soft catch the particles to examine potential deformation

Parametric Variation:

- Vary particle arrangement (one, two, three and more circle of particles) – first set of experiments
- Vary how particles are held in place – investigate its sensitivity
- Vary particle size and particle material

Statistical repetition:

- 3 repetitions for a few of the experiments will be desirable

T0: Validation of the demonstration problem

Purpose: The purpose of all the other meso and microscale activities is to quantify errors and uncertainties and propagate them to the macroscale simulation of the demonstration problem. The question then is Can the macroscale simulation + propagated uncertainty accurately predict the results of actual macroscale experiments of the demonstration problem + associated experimental uncertainties.

Strategy: We will perform macroscale experiments of the demonstration problem with an annular bed of particles. Macroscale simulations will also be performed and validated against the experiments. Large errors/uncertainties will drive our uncertainty reduction effort.

Desired experimental features:

- Very light annular casing – one that will break away isotropically with minimal resistance
- Well known explosive whose energy release and EOS of products are known
- Precise details of the bed of particles
- Similar to Forst experiments, but under controlled conditions and with more instrumentation

Desired experimental measurements for comparison:

- Pressure time histories at a number of well-characterized spatial locations
- Short time x-ray images of particle front position vs time
- Long time high speed images of particles front and its instability
- High speed images of the expanding detonation products

Parametric Variation:

- Vary particle bed thickness
- Vary particle size and particle material
- Consider Oreo arrangement
- Introduce controlled initial perturbation

Statistical repetition:

- 3 repetitions for a few of the experiments will be desirable

5. UB Team

5.1 Overview

The main objective of the UB team is to reduce the uncertainty in both simulations and experiments to a level that would allow meaningful validation of the simulation code. A second objective is to help the exascale team reduce the uncertainty in their predictions to permit confidence in emulation of future exascale architectures. This requires an extensive uncertainty quantification effort, which poses two challenges. First, the UB team has to work collaboratively on identifying important uncertainties with the simulation, experiment, and exascale teams. Second, the computational requirements for good uncertainty quantification is orders of magnitude more expensive than simulation alone. To address the first challenge, we selected the mesoscale shock tube configuration to establish working relationship and procedures to collaborate with the physics and experimentalist teams, mostly because experimental results for that scale are already available. Similarly, we chose a problem of predicting matrix multiplication execution times for collaboration with the exascale team. To address the computational challenge we have been evaluating and developing several surrogate based methods.

5.2 Validation of Mesoscale Shock Tube Simulation

As a preliminary study for subscale and overall validations, we are engaging to validate the mesoscale simulation, the shock tube with a particle curtain. Figure 5.1 (a) shows the configuration of the experiment where a shock is created from high pressure gas in the driver section into the driven section by bursting a diaphragm between two sections. The shock interacts with a particle curtain in the test section as is illustrated in Figure 5.1 (b). We plan to validate the capability of predicting the motion of the particles due to the shock by comparison with experimental observations. The prediction metrics (PM) to be used for validation are the upstream and downstream edge locations of the particle curtain.

An important outcome of validation is quantification of the model uncertainty. As a first step in uncertainty quantification (UQ), we identified and classified uncertainty sources and errors,

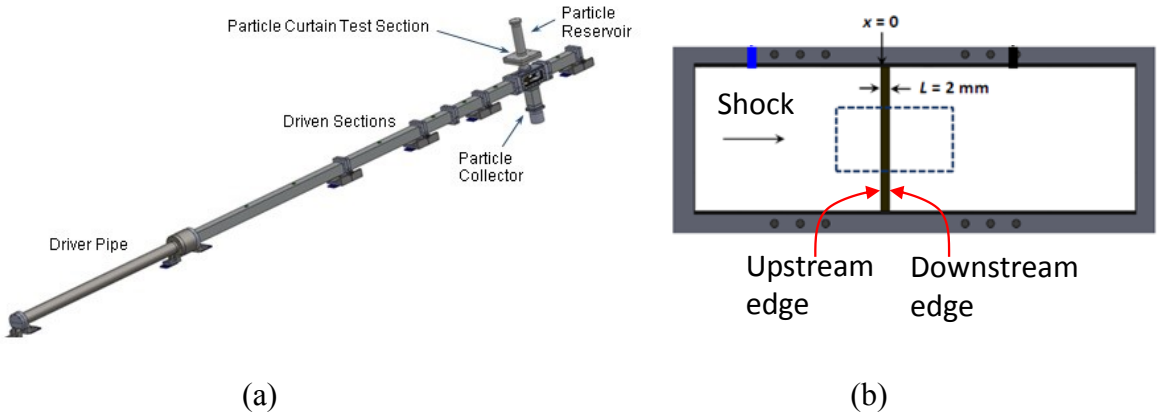

Figure 5.1. Schematic figure of the shock tube (a) and particle curtain in test section (b).

as shown in Figure 5.2. For validation, the measured PM and the calculated PM are compared and the discrepancy is due to uncertainties, including physical model, numerical model and numerical solution errors and measurement uncertainties in input and PMs and numerical solution error. By quantifying the other uncertainties and comparing it to the discrepancy between simulation and experiment, we also get an idea of the possible range of magnitudes of the physical model errors.

The model uncertainty of the 1-D simulation code (Rocflu lite) has been thus quantified. Figure 5.3(a) shows calculated PMs with the model uncertainty of 95% confidence in terms of time. The width of the band represents the model uncertainty in the simulation and the band says that the true edge locations will be within the bands with 95% probability. Figure 5.3(b) shows the discrepancy between calculated PMs without UQ and measured PMs from four repeated experiments. The variations between symbols from measured PMs represent the uncertainty in experiments.

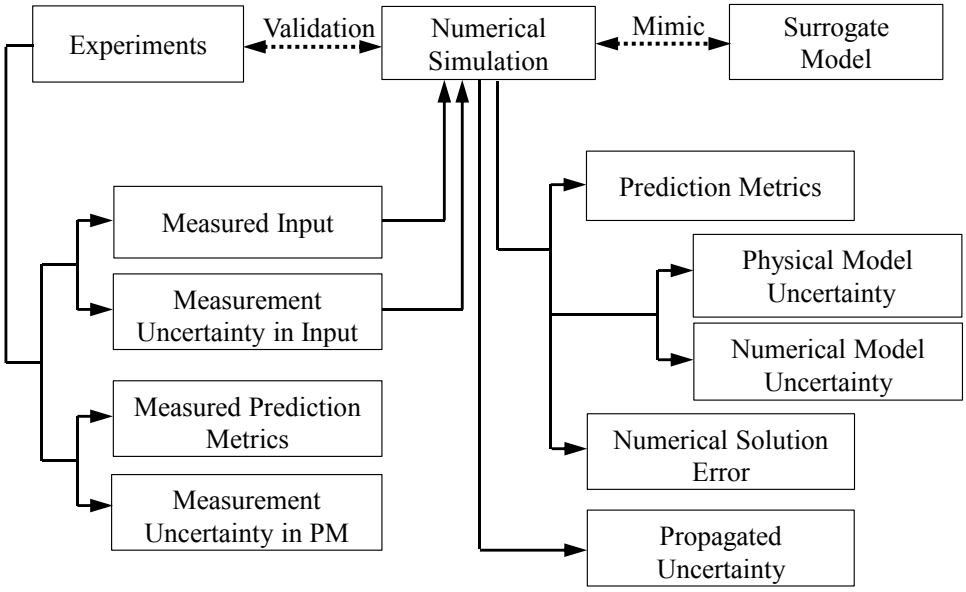

Figure 5.2. Sources of uncertainty and errors in validation process.

The width of the band represents the uncertainty in the validation process. To reduce it we have launched an uncertainty reduction (UR) effort. In this validation, we found that the uncertainty in the particle diameter of the experiment affects the model uncertainty most, so that reducing the uncertainty in the diameter is an efficient way to narrow down the uncertainty.

From the shock tube experiment, local volume fraction variation and curtain thickness variation are observed whereas the 1-D simulation uses a simplified curtain model of a constant thickness and a homogeneous volume fraction. A 3-D simulation is being prepared to quantify the effects of these simplifications.

A graduate student, Giselle Fernandez, studies 1-D simulation with help of Bertrand Rollin who is in charge of macro/mesoscale simulation, Christopher Neal a graduate student working for micro-scale simulation and Angela Diggs in Eglin Air Force Base. We have also interacted with Justin Wagner in Sandia National Labs to identify and understand uncertainty sources in shock tube experiments and we identified uncertainty sources.

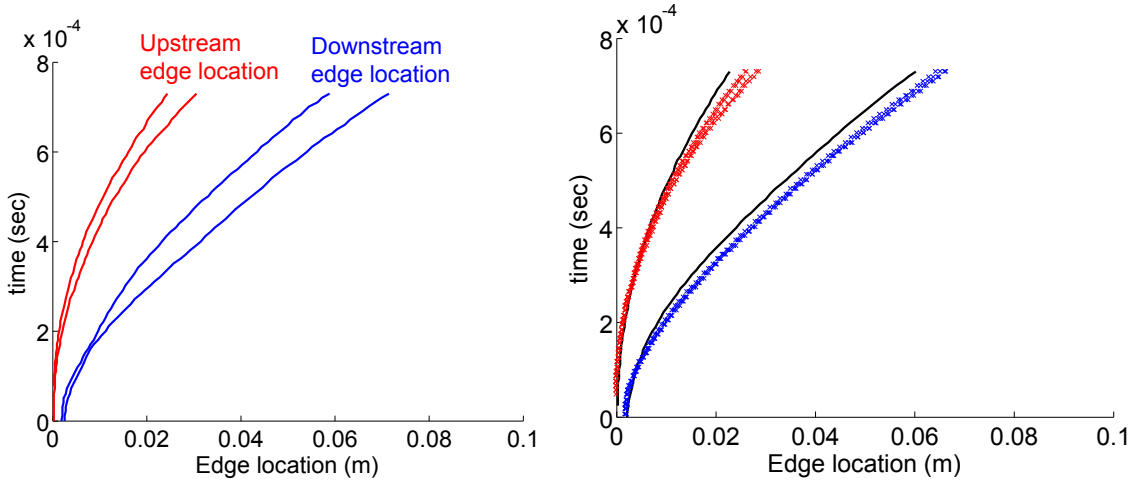

(a) (b)

Figure 5.3. Predicted upstream and downstream particle curtain edge locations (PMs). (a) Calculated PMs with model uncertainty of 95% confidence. (b) A comparison between a calculated PM and measured PMs from four repeated experiments.

5.3 Multifidelity Surrogate Model for Uncertainty Quantification (UQ)

Since UQ requires computational resources more than one order of magnitude higher than simulation, we plan to use a surrogate model that is a cheap algebraic function mimicking an expensive simulation response. Fitting a surrogate requires data obtained from simulation runs but even obtaining sufficient data to fit a surrogate with reasonable accuracy may not be affordable for high fidelity simulations. Multi-fidelity surrogate (MFS) compensates for insufficient data from expensive high fidelity simulations. In particular, we expect to use 1-D simulations as low-fidelity source and 3-D simulations as high-fidelity source.

Figure 5.4 shows a demonstration example of building a MFS using high and low fidelity data sets and a developed MFS building tool. The left figure shows high and low fidelity data sets from the corresponding functions which represent high and low fidelity simulation responses for a single input variable over [0,1]. The right figure shows that the fitted MFS (red line) captures the high fidelity function by combining two data sets and the estimated prediction uncertainty (blue shade) captures the actual discrepancy between the fit and the true high fidelity function.

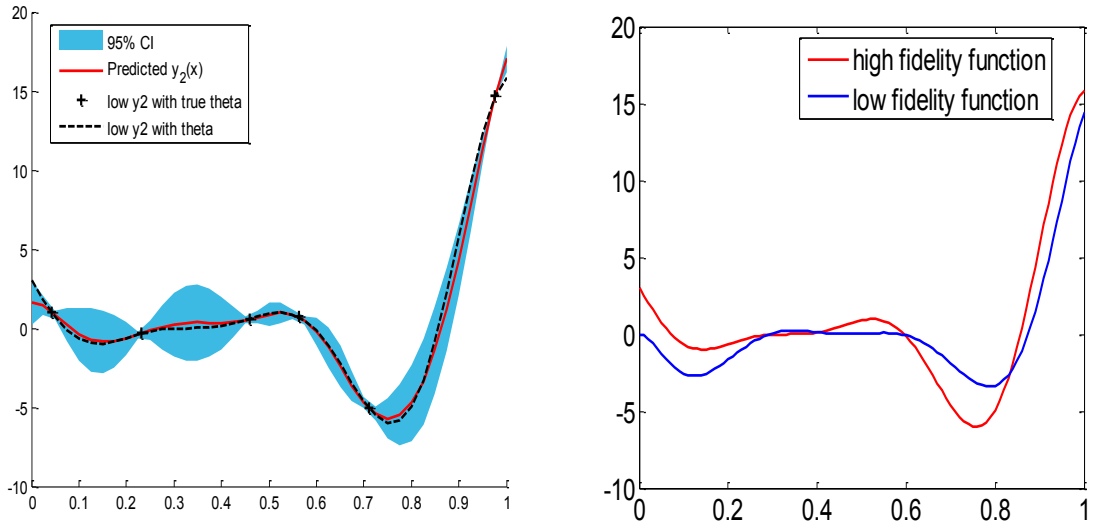

Figure 5.4. Fitted multi fidelity surrogates (red solid line) based on 6 samples (black crosses) from the high fidelity function and 36 samples from the low fidelity function and estimated prediction error (blue shades).

There are well developed alternative MFS formulations, and we are studying them to have better understanding of their features. For example, we observed that an MFS with a model discrepancy function generally performs well but it is bad when the number of high fidelity data points is small. A calibration based MFS outperforms the MFS with a model discrepancy function. We also carried out study for developing best sampling strategy for MFSs.

5.4 Extrapolation

We expect to need extrapolation in the validation work, both on the physics and exascale sides. We are developing methodology for extrapolation to an inaccessible point and estimating extrapolation uncertainty at that point. The key is using multiple 1-D surrogates for multi-dimensional functions.

The methodology was implemented and examined with two benchmark problems: 1) drag coefficient prediction for sphere particles and 2) computer performance prediction for matrix multiplication.

For the first problem, extrapolation of a drag coefficient for high Reynolds number and supersonic Mach number is made. Since experiments at high Reynolds number and supersonic Mach number are very expensive, data are not available and extrapolation is favored. Figure 5.5 (a) shows the extrapolation point with a hollow circle and inaccessible domain on log(Re) and Mach number space. Figure 5.5 (b), (c) and (d) show extrapolations on three lines and each line was fitted with 6 data points. Each line uses a normalized parameter α on the line and α=1 is the extrapolation point.

We found that extrapolation on Mach number line makes better extrapolation than the other two lines and the estimated uncertainty shows the same phenomenon that the methodology works for the problem. Bayesian combination of the three lines automatically weighted to most accurate line most heavily to yield very accurate estimate at the extrapolation point.

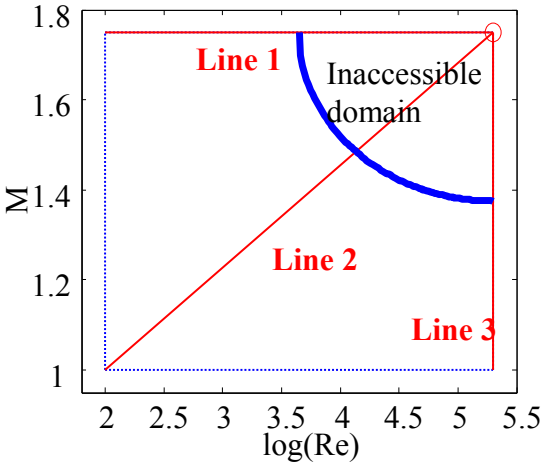

(a) The extrapolation point in the inaccessible domain with three lines

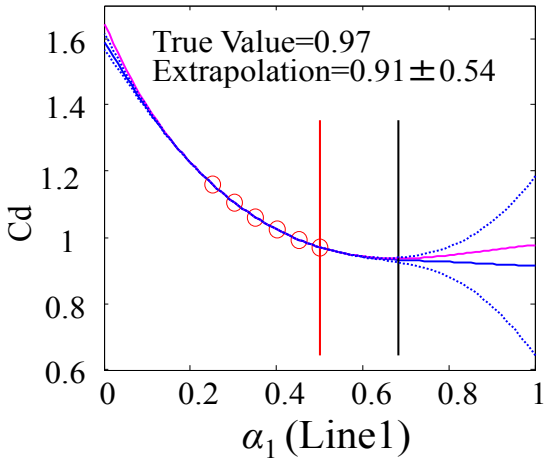

(b) Extrapolation with Line 1 for variable Reynold number and a fixed Mach number

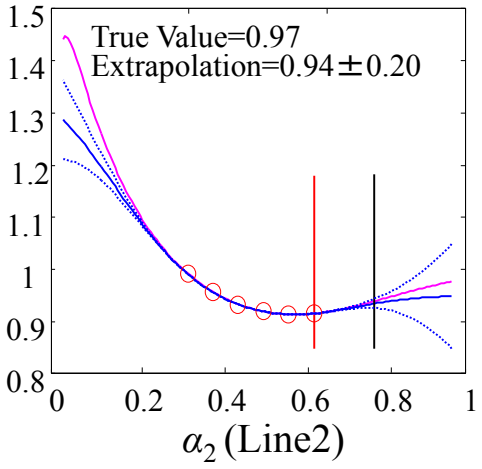

(c) Extrapolation with Line 2 on a hyper diagonal

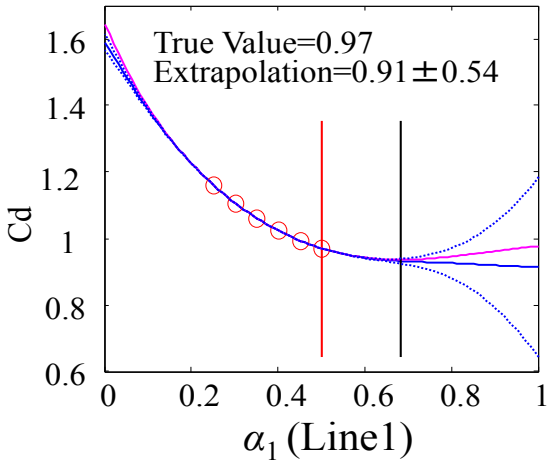

(d) Extrapolation with Line 3 for variable Mach number and a fixed Reynold number

Figure 5.5. Drag coefficient extrapolation for M=1.8 and log(Re)=5.5.

Extrapolation could be adopted to estimate computational cost of numerical tasks. Since execution time of matrix multiplication is a basic indicator of computer performance, extrapolation of the matrix multiplication time for a new computer can help to predict its capability. For this problem, the data is noisy due to competing demands on the processor, and so we are examining extrapolation methods that deal well with noisy data. Currently ridge regression is the tool we are examining for this purpose.

We have also realized that some of our models are weaker for certain combinations of particle density and diameter, usually at the boundary of the expected range. Covering the parameter range domain with a grid of simulations permits us to extrapolate from inside the domain to the boundary. When the extrapolated results differ substantially from the simulations, it raises a red flag that motivates us to check whether the simulation is doing a good job. In some cases, refining the mesh corrected the problem. In others we are looking to model improvements under way to see if they will bring the extrapolated results in line with the simulations.

5.5 The Microscale drag force model development support

For the microscale, we have developed a physical algebraic surrogate model for calculating kernel functions used to calculate force on a spherical particle for a given Mach number.

Getting a kernel function for a given Mach number requires expensive DNS. As Fig. 5.6 shows with the crosses, the kernel functions are obtained at certain time steps, the kernel functions has in the past been fitted with an exponential sum with more than 60 coefficients for interpolation. Our goal was to construct a simple parametric surrogate for kernel function for time and Mach number.

Here we constructed a physical-algebraic surrogate using knowledge about physics of kernel functions and the known solution for M=0. From Figure 5.6, we elicit the information that the kernels can be reproduced by stretching the kernel for M=0 in the time direction for different Mach numbers. Since the kernel function for M=0 is analytically defined, a surrogate for kernel functions can be derived based on the kernel function of M=0 with time stretching functions.

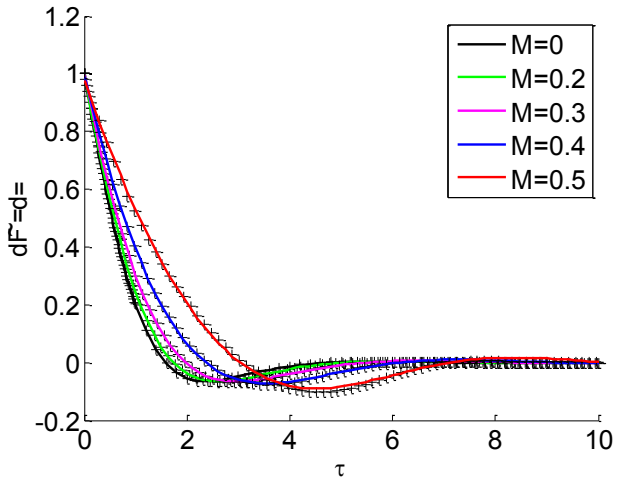

Figure 5.6. Kernel function data (crosses) and fitted surrogates (colored solid lines) for different Mach numbers.

The developed physical-algebraic surrogate significantly reduces the number of coefficients 60 to 6 with acceptable accuracy reduction. A

comparison between the fitted curves and data are shown in Figure 5.6.

However, extrapolation of a kernel function for Mach number (i.e. M=1) for given kernels for M={0, 0.1, 0.2, 0.3, 0.4, 0.5} increases error rapidly with extrapolation distance. Further research will be carried out for suppressing the extrapolation error.

6. CMT-nek Code Development Team

6.1 Overview

In 2014, development began on CMT-nek from a production release of nek5000 and the governing equations of compressible multiphase flow given by Powers (2004) Phys. Fluids 16, 2975. After familiarization with existing scalar transport method and code in nek5000 and a study of literature about discontinuous Galerkin (DG) methods in fluid mechanics, the CMT-nek development group applied the strong form of the weighted residual theorem to the governing equations. CMT-nek inherits element topology, approximation polynomials and high-order spatial discretization from nek5000 and hooks into nek5000 as a branch of execution and source code (see Figure 6.1), yet solves a fundamentally different class of problems that exceed the capabilities of nek5000. The discontinuous Galerkin method replaces the constraint of solution continuity at the boundaries of neighboring elements with numerical fluxes, which are informed by the physics of compressible flow and analogous to their role in finite volume methods. The numerical flux functions are taken from the implementation of the improved Advection Upstream Splitting Method (AUSM+) of Liou (1996) J. Comp. Phys. 129 in *RocFlu*. *RocFlu*'s validated methods and established code base have provided CMT-nek with boundary conditions for inviscid compressible flow and ideal-gas state equations.

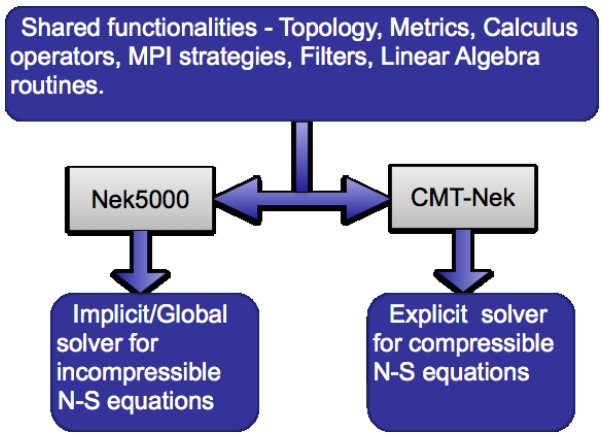

Figure 6.1. Schematic that shows the association of CMT-nek and nek5000 with respect to shared functionalities and their fundamental difference, i.e., the governing equations that are solved.

6.2 Year 1 Research

In the first fiscal year, CMT-nek has achieved several important milestones (see Table 6.1). By leveraging the existing functionalities of nek5000 and translating established code base of RocFlu, CMT-nek development group was able to obtain an explicit solver for Euler equations with limited multiphase coupling in the form of a nozzling term in momentum equation in four months (April-August 2014). From August 2014 to November 2014, CMT-nek was symbolically debugged and tested with the problems listed in Table 6.2. These test problems constitute verification of the method and its implementation for inviscid flows of a single phase without shocks in two and three dimensions.

Subsonic/Supersonic flow in a 2-D domain test simply ensured that the initialized uniform flow could be maintained over an arbitrary number of iterations, serving as a sanity check. One-dimensional rarefaction in two- and three-dimensional domains is a time evolving flow furnishing

canonical solutions demonstrating the physical coupling among the governing equations. While the solution is not smooth, it does not have jump discontinuities; nonmonotonicity introduced by the polynomial discretization merely reduces the rate of convergence. This test case verified the AUSM+ implementation and subsonic outflow boundary condition upon comparison of the CMT-nek solution with the numerical solution obtained from the Riemann solver of Toro (1997) Riemann Solvers and Numerical Methods for Fluid Dynamics (see Figure 6.2 and 6.3).

Mathematical formulation of DGSEM for compressible multiphase Navier-Stokes equations using the approximation polynomials and higher order spatial discretization of nek5000	April 2014
Overarching framework of CMT-nek and its association with nek5000	May 2014
Explicit solver for Euler equations – isentropic flows, with limited multiphase coupling in the form of throttling term in the momentum equation.	August 2014
Verification and validation of CMT-nek using several physical flow cases	November 2014
Establish spectral convergence of CMT-nek	December 2014
Detailed implementation plan for incorporating viscous terms in CMT-nek	January 2014

Table 6.1. List of milestones reached in the first fiscal year, January 27th 2014 to January 27th 2015.

The previous test cases, although simulated on a 2-D and 3-D grid, were 1-D flows. The test of subsonic flow over a cylinder and Prandtl-Meyer expansion fan ensure that CMT-Nek was verified for a truly 2D flow field. Figure 6.4 shows the Mach lines as Mach 2 flow turns over a circular wall. This solution is verified using the analytical solution for Prandtl-Meyer expansion fan. Figure 6.5 (a) and (b) shows the pseudocolor plot of local Mach number and static pressure around the cylinder. This solution has been compared to the Rayleigh-Janzen low-Mach-number asymptotic expansion due to Imai (1941) Proc. Physico-Math. Soc. Japan 23. (see Figure 6.6). Figure 6.7 shows the time evolution of coefficient of drag for the cylinder, which agrees well with the theory (Cd is zero for steady flow). It should be noted that the uniform flow over a cylinder case was run on LANL machine Mustang using 480 MPI ranks.

Another important test case was the isentropic vortex advected in a periodic domain. We have adopted the formulation given in Yee, Vinukor, Djomehri (2000) J. Comp. Phys. 162. This flow is smooth and uses periodic boundary conditions, which makes it an ideal test case for establishing the convergence rate of CMT-nek (explicit compressible inviscid solver). Several cases were simulated using different advection velocity, domain size and the shape of the elements. Figure 6.8 shows the shape of the vortex that is advected in the domain. Figure 6.9 shows the convergence rate plot. Different lines in the plot correspond to either different advection velocity or change in the shape of the elements. Through these simulations we have established that CMT-nek indeed has spectral convergence.

The current version of CMT-nek has the potential to be deployed to perform microscale simulations for CCMT. As an important step in that direction, CMT-nek group stated simulation of subsonic flow over a sphere. This case fulfills two objectives – to verify CMT-nek for 3-D flow and to establish feasibility of performing microscale simulations for CCMT. These simulations will be performed in two stages; stage 1 will attain a steady state solution and in stage 2 the inviscid unsteady force kernel will be extracted by imposing the sphere with an impulsive acceleration profile. The extracted kernel will be compared with the exact solution given by Parmar et.al. (2008) Phil. Trans. R. Soc. A. 366 to verify the convergence rate of CMT-nek. Towards the end of the first fiscal year we began with these simulations. Figure 6.10 shows pseudocolor plot of static pressure and velocity magnitude on the surface of the sphere. These runs are performed on Mustang at LANL using 128 MPI ranks. More simulations are currently underway and we envision the exercise to be completed in the first quarter of second fiscal year.

Uniform subsonic/Supersonic flow in a 2D domain	At various Mach numbers, one element to multiple elements in the domain and different in-element resolution.
Rarefaction – Hot and pressurized can of gas emptying into ambient	Multiple elements in the domain and different in element resolution.
Prandtl-Meyer Expansion fan	Inflow Mach = 2, turning angle = 20^{o}. Multiple elements and different in element resolution.
Subsonic flow over a cylinder	Various inflow Mach numbers, Multiple elements in the domain and different in-element resolution.
Isentropic vortex in a 2D periodic domain	Different advection velocity, Multiple elements in the domain and different in-element resolution.
Subsonic flow over a sphere	Various inflow Mach numbers, Multiple elements in the domain and different in-element resolution.

Table 6.2. List of test problems performed to verify the explicit inviscid flow solver in CMT-nek.

6.3 Second Year Plans

The current version of CMT-nek provides a strong foundation to build new capabilities that are paramount for CCMT. CMT-nek development team has made detailed plans of CMT-nek code development for the second fiscal year. Table 6.3 shows the list of capabilities that would be added in the second fiscal year. Also, shown in the table are several test cases that would be used to verify CMT-nek.

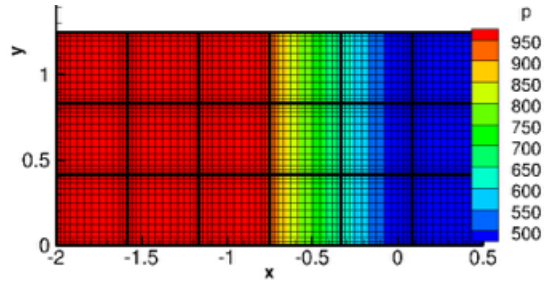

Figure 6.2. Pseudo-color plot of pressure on top of the mesh that was used to simulate rarefaction case. Note that the rarefaction is moving from right to left. The left vertical boundary is a wall. The gas in the domain was initially at high pressure and high temperature, while the ambient outside is at a lower pressure and temperature. The right vertical boundary is an outflow through which the hot and pressurized gas is being released.

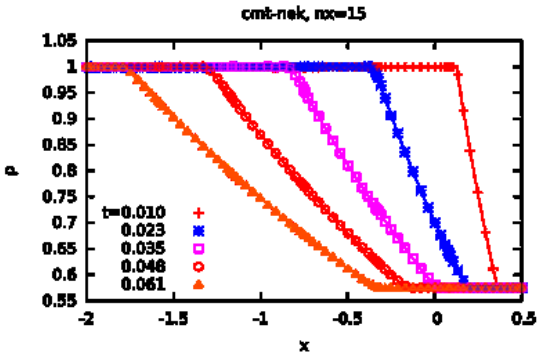

Figure 6.3. Evolution of density in the domain at different time. Symbols represent the solution obtained from CMT-nek and the lines represent the exact solution (obtained from the numerical solution provided by Riemann solver, E.F. Toro 1997).

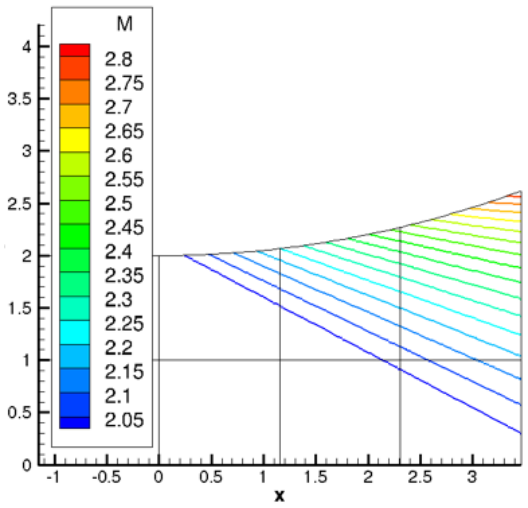

Figure 6.4. Prandtl-Meyer expansion fan. Mach 2 flow turning over a circular arc by 20°. Also shown in the figure are the elements that make up the domain.

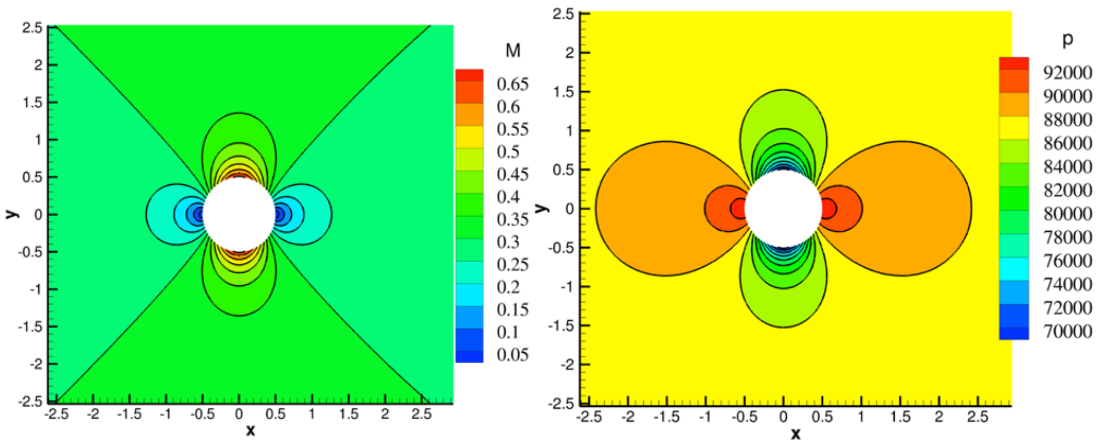

Figure 6.5. Pseudocolor plots of (a) Mach number and (b) static pressure near the cylinder. Note that free stream Mach number is 0.3.

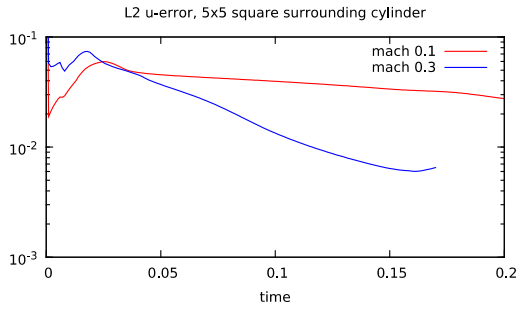

Figure 6.6. Evolution of L2 error obtained by comparing CMT-nek solution with the Rayleigh-Janzen low Mach number asymptotic expansion.

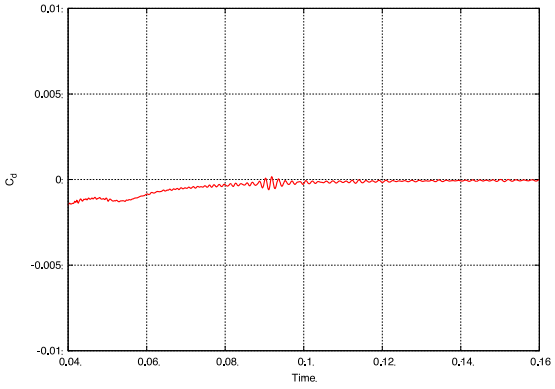

Figure 6.7. Coefficient of drag for M=0.3 flow over a cylinder.

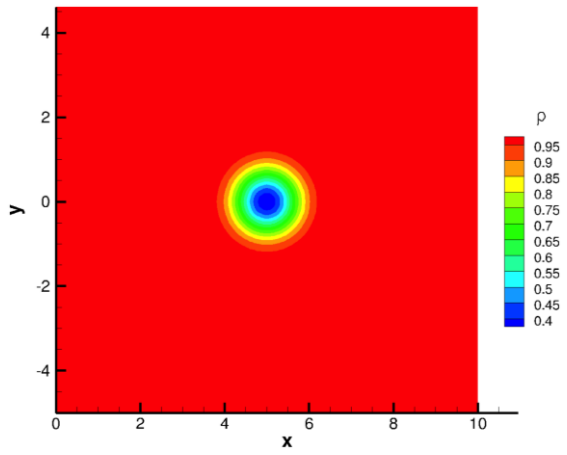

Figure 6.8. Pseudocolor of density in an isentropic vortex.

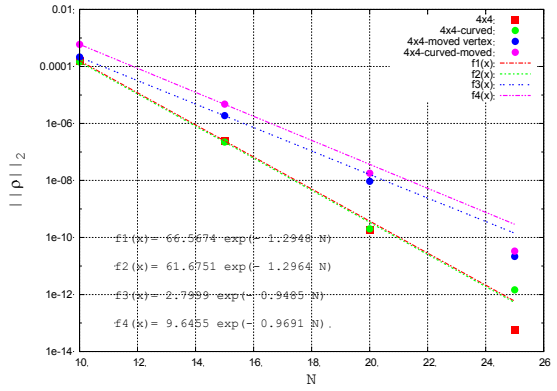

Figure 6.9. Convergence rate for various cases using deformed elements.

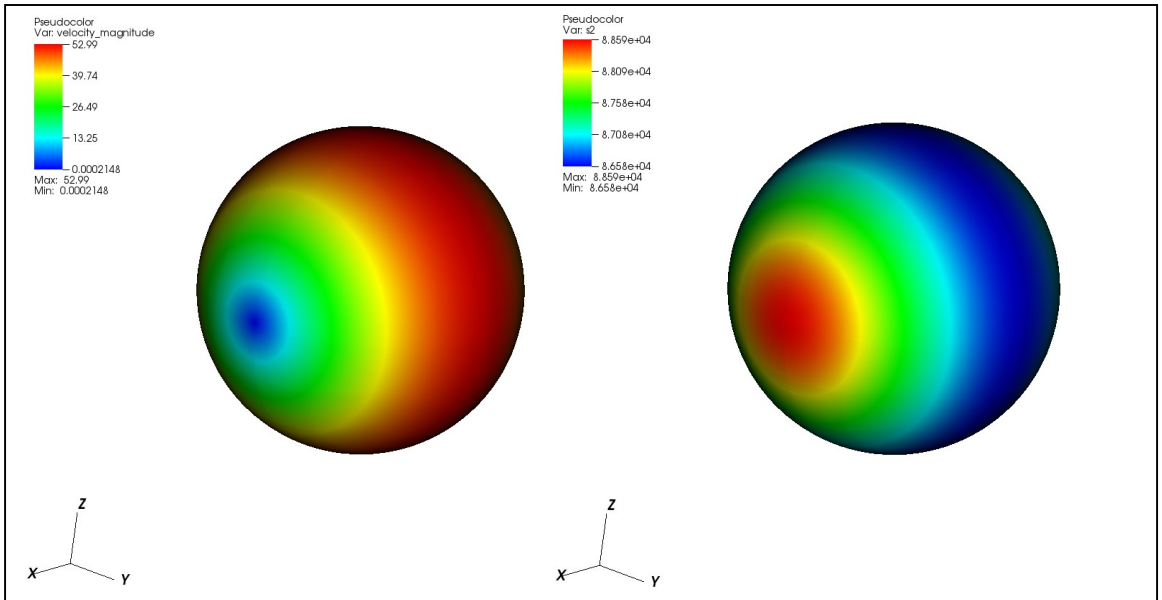

Figure 6.10. Pseudocolor plots of (a) Mach number and (b) static pressure near the cylinder. Note that free stream Mach number is 0.3.

Capability	Test cases	Time frame
Euler equations subsonic	• Inviscid unsteady force kernel for subsonic flow over a sphere • Simulate flow over array of spheres (microscale simulations for CCMT)	April 2015
Compressible Navier-Stokes	• Exact solutions, canonical asymptotics (vortex and BL) • Viscous flow over a sphere, low Mach, low Re • Compressible isotropic turbulence	May 2015
Lagrangain point particles	• Inertial particles in compressible isotropic turbulence, 1-way coupling	August 2015
Shock capturing	• SOD problem • Sedov problem • Shock propagating over array of spheres (microscale simulations for CCMT)	September 2015
Dense to Dilute volume fractions	• Moving shock over curtain of particles and cylindrical explosive dispersal of particles (Macroscale simulation for CCMT)	November 2015

Table 6.3. List of capabilities that would be added to CMT-nek in the second fiscal year. Also shown are the projected completion dates and list of test cases to be used to verify the code.

7. CS Team

7.1 Overview

The primary goal of the research of our CS team has been to improve performance, power consumption and thermal characteristics of the CMT-nek application, and in the process find new strategies and optimization ideas that can be applied to other applications as well. During the last year we concentrated on the derivative computation kernel of CMT-nek. We implemented a genetic algorithm based autotuning method, which gives us highly efficient versions of the derivative computation kernel of CMT-nek for a given platform. While autotuning is a standard method, most prior work does exhaustive search of the design space. Our work introduces a smart search mechanism in this area.

7.2 Autotuning

Over the past year, we focused our efforts on improving the performance and power consumption characteristics of the spectral element solver which is the most compute intensive kernel of Nek5000 and is being reused in CMT-nek. The fundamental algorithm of this kernel is small matrix-matrix multiply that is used to numerically compute functional derivatives. We have used a number of optimization techniques on CPU and GPU to achieve improved results compared to the original Nek5000. On CPU, the main optimization technique used was auto tuning and we tested our results on IBM Blue Gene/Q, 64-core AMD Opteron 6378 and AMD-Fusion. Figure 7.1 shows the improvement in performance and energy obtained on IBM BG/Q for a function matrix of size 10x10x10. In the figure, 'u' is a function that is differentiated along (r,s,t) axes. Three different basic versions of matrix-matrix multiply were autotuned. These versions are namely, 5loop-fused, 4loop and 4loop-fused. The 4loop version is the regular matrix-matrix multiply algorithm used to multiply a 3-D functional matrix with a 2-D operator matrix. The 4loop-fused version is a variant of the regular algorithm where two loops are fused. The 5loop version is another variant that includes the loop over all spectral elements in the optimization, and the 5loop-fused if the figure is a version of the 5loop algorithm with two or more fused loops. The performance improvement is 40, 30 and 11 percent respectively, with an aggregate improvement of 30 percent with respect to the original implementation of the Nek5000 kernel. The memory access pattern of duds prevents loop fusion, and hence the duds group of bars does not contain a 4loop-fused version.

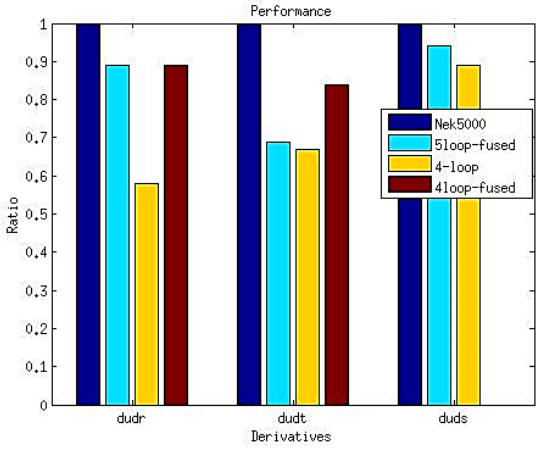

Figure 7.1. Performance graph for partial derivatives dudr, dudt and duds, normalized with respect to performance of original Nek5000 kernel.

On GPU we employed a number of optimization techniques based on the data access pattern of this application.

Figure 7.2 presents the energy requirements for a subset consisting of high performance matrix multiplication variants for dudt 4loop-fused version, and Figure 7.3 presents the total energy consumption for the different version of codes.

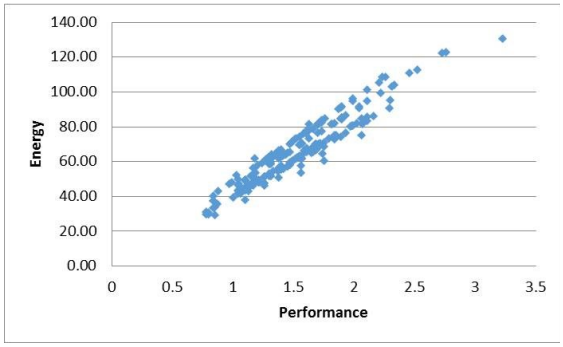

Figure 7.2. Energy of the multiplication variants for dudt 4loop-fusedversion measured on IBM BG/Q. Energy was measured in Joules and performance in seconds. The graph shows that energy and performance are very closely related. We therefore studied the top 32 best performing variants for each version for energy analysis.

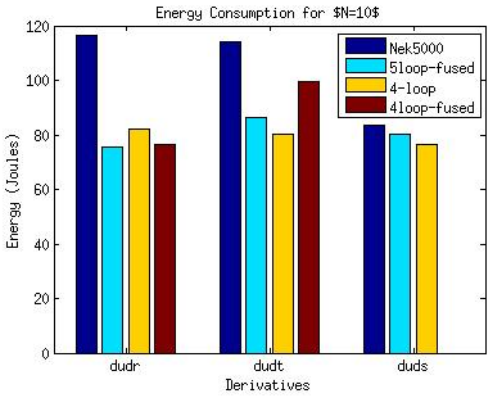

Figure 7.3. Total energy consumption of the top 32 fastest variant for each version of the basic codes namely 5loop-fused, 4loop and 4loop-fused and how they compare with Nek5000.

Figures 7.4 and 7.5 show the results of our optimization strategy as applied to the derivative computation algorithm on GPU. THE GPU optimization strategy revolves around the fact that the derivative operator matrix is a common operator and can be brought just once into device memory and can also be stored in registers. We implemented two variants of derivative computation – "Combined" and "Separate". The "Combined" implementation brings a chunk of the function matrix to device memory and computes partial results for *dudr*, *duds* and *dudt* from that chunk, then brings the next chunk of the function matrix and adds to the partially computed results for *dudr*, *duds*, and *dudt*, and this process continues for the entire function matrix. The "Separate" implementation computes the whole of *dudr* first, then *duds* and finally *dudt*. In doing so, the same chunk of the function matrix is brought in the device memory multiple times, as opposed to only once in the "Combined" case.

7.3 Autotuning with Genetic Algorithms

Our initial efforts on autotuning focused on exhaustive search of over the design parameter space. For the next part we conducted research on how to cut down on the search time and have developed a novel genetic algorithm based driver for an autotuned flow identifying efficient algorithms for different matrix sizes. The table in Figure 7.6 highlights the quality of the results we got using our driver for the matrix size 10x10x10 on AMD Opteron.

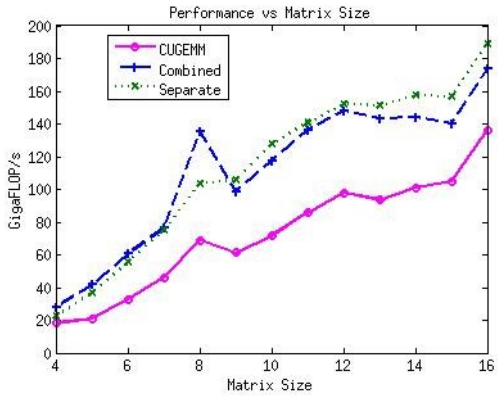

Figure 7.4. Results of our GPU optimizations for the matrix multiplication kernels. CUGEMM refers to the work by Jhurani et al. which speeds up small matrix-matrix multiplications significantly over the existing BLAS libraries that are optimal for large matrix sizes. 'Combined' represents our strategy in which we compute dudr, duds and dudt in one single kernel, whereas 'Separate' represents our strategy of computing the three derivatives separately. We achieve over 180 GFLOPS for matrix size 16x16x16 on Tesla K20c GPU.

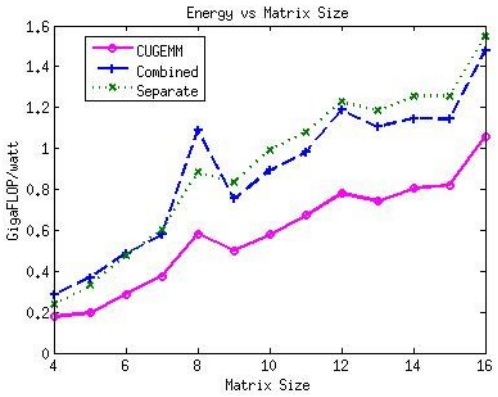

Figure 7.5. Performance per watt graph for each matrix size. The performance per watt is over 1.5 on GPU. From our results, GPUs are better than CPU cores in terms of performance per watt, by a factor of 2.5 to 4.

Derivative	Minimum Time	Nek5000 Time	CMT-nek Time	Number of variants analyzed
$\frac{\partial u}{\partial r}$	0.39 s	0.62 s	0.41 s	444 variants
$\frac{\partial u}{\partial s}$	0.18 s	0.31 s	0.20 s	393 variants
$\frac{\partial u}{\partial t}$	0.45 s	0.88 s	0.45 s	446 variants

Figure 7.6. A comparison among various runtimes, for an input matrix size of 10x10x10, is shown. The second column titled "Minimum Time" shows the runtime by the best algorithm for each of the derivatives listed in the first column. We obtained the minimum runtime after running more than 30,000 matrix multiplication variants for each derivative type. The third column titled "Nek5000 Time" shows the runtime by the current Nek50000 implementation. The fourth column titled "CMT-nek Time" shows the runtime of the best variant reported by our genetic algorithm based driver. The fifth column shows the number of variants analyzed by our driver to report these variants. This figure shows how close we got to the minimum runtime after analyzing only about 1% of the variants in the exhaustive search case.

On Teller, which is an advanced AMD-Fusion platform at Sandia National Laboratory, we measured energy for the matrix multiplication variants, and found that the variants that consumed less energy are also the ones that have faster runtimes. Figures 7.7 and 7.8 show the energy consumption and runtimes of the energy optimal variants for various matrix sizes as determined using autotuning driven by our genetic algorithm.

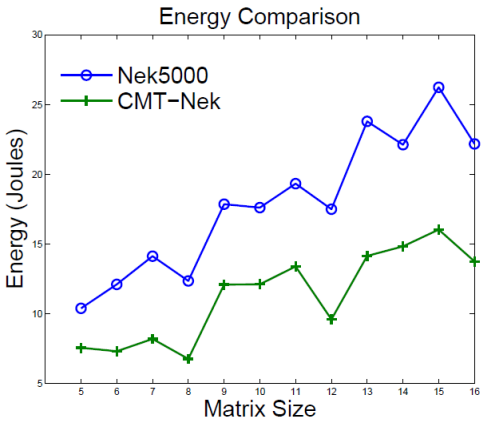

Figure 7.7. Energy consumption in Joules by the optimal variants found using our genetic algorithms for different matrix sizes. The matrix sizes range from 5x5x5 through 16x16x16. Energy was measured using PowerInsight. There was up to 45% improvement in energy consumption compared to the original Nek5000 version.

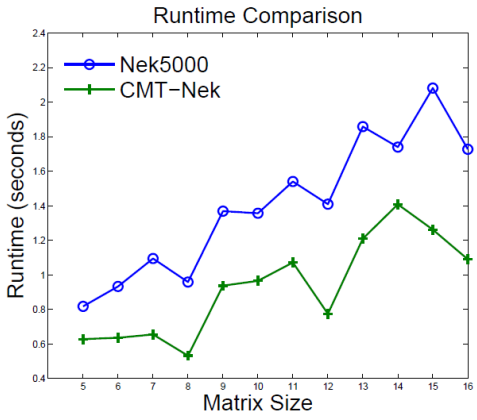

Figure 7.8. Runtime comparison in seconds by the energy optimal variants in Figure 7.7. Up to 45% improvement in runtime was observed, compared to the original NEK50000 version of these algorithms.

7.4 Second Year Plans

Below is a list of tasks that the CS team expects to complete in the coming year.

1) Continuing our work on performance and energy optimization, we will apply our autotuning technique on spectral interpolation of Nek5000. This code also uses matrix-matrix multiplication but in this case the matrices are rectangular and presents additional optimization parameters.
2) We will run a complete CMT-nek example with our tuned library of matrix multiplication routines for derivative computation and spectral interpolation and observe the overall improvement in performance and energy over Nek5000.
3) We will work on power modeling and a method to predict component power on multicore processors.

4) We will also analyze thermal characteristics of applications. We will first explore the existing temperature measurement facilities available at the National Laboratories as well as at the UF high performance computing system – Hipergator, and if needed set up our own temperature measurement apparatus to conduct research on how temperature on a chip can be controlled using different workloads on the cores. We will leverage our preliminary work on thermal aware allocation of workload for data parallel processors.

8. Exascale Team

8.1 Overview

A key goal for the Integrated Code Development and Performance effort is to identify the communication and computationally intensive portions of the CMT-nek code, understand its thermal and energy issues, and develop methods to perform design-space exploration (DSE) to optimize and auto-tune the code on existing computing systems. In a similar manner, the goal for the Exascale behavioral emulation effort is to develop methods and simulation/emulation tools to perform DSE for the identified communication and computationally intensive portions of CMT-nek. In this case, however, the objective of the DSE is to provide insights and guidelines to the code development team for algorithmic exploration and design choices for the development of CMT-Nek on notional, future-generation computing systems, up to Exascale.

8.2 Overview of Behavioral Emulation

Traditional simulation and emulation approaches face significant challenges when used for analyzing design tradeoffs of Exascale systems; key among them is optimizing the delicate balance among simulation speed, model accuracy, and design scalability. To overcome the limitations of analytical modeling (fast but inaccurate), cycle-accurate software simulation (accurate but prohibitively slow) or cycle-accurate hardware emulation (fast and accurate but too complex), we propose coarse-grained simulation and emulation that enables rapid virtual prototyping of extreme-scale systems and follows a hierarchical multi-scale approach called ***behavioral emulation (BE)***. In our approach, as shown in Figure 8.1, an important way to manage Exascale complexity is by abstracting the high-level behavior of a system as a function of its components and their interactions. The fundamental primitive of behavioral emulation is a ***Behavioral Emulation Object (BEO)***, an abstract model that stands in as a surrogate for and mimics the behavior of the component under study. Another important way to manage Exascale complexity and accuracy is

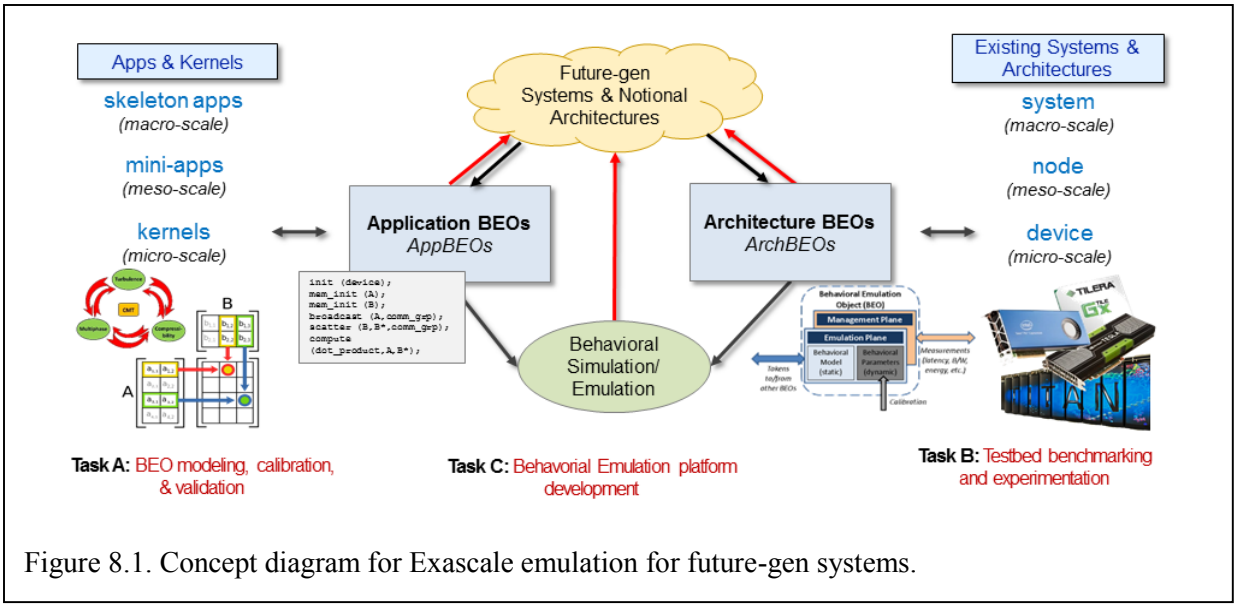

Figure 8.1. Concept diagram for Exascale emulation for future-gen systems.

to use a ***multi-scale approach***, dividing the simulation or emulation into three levels of abstraction: micro-, meso-, and macro-scale.

As shown in Figure 8.1, the project is divided into three major tasks. Task A is Behavioral Emulation Object (BEO) modeling, calibration, and validation of application kernels (App BEOs) and systems (architectural BEOs). Task B is testbed experimentation and benchmarking used to calibrate and validate the developed models (BEOs). Task C is development of the behavioral emulation platforms. Currently, we are exploring the use and development of three platform types: discrete-event simulator, an SMP simulator on many-core computer, and hardware emulator on a reconfigurable computer. The three major tasks have been divided in six major research thrusts. For this first-year annual report, progress for the first four research thrusts (T1 – T4) is summarized as follows.

8.2.1 T1: Behavioral Characterization

The aim of behavioral emulation is to explore and evaluate the performance of Exascale applications and architectures by mimicking their behavior. A key task is to abstract away low-level functional details of applications and architectures and modeling them at multiple levels. Shown in Figure 8.2 is an illustration of the behavioral emulation workflow. In Step 1, we explore different modeling strategies for architectural components (such as processors, accelerators, networks, etc.), which are modeled and calibrated, both in terms of computation and communication. Performance models are developed based on sampling (running kernel code on target platforms) and using these samples for interpolation. Currently, we are using the Kriging method for multi-dimensional interpolation.

These models are then calibrated and validated using inputs that are not in the initial sample set. Currently, we have developed, calibrated, and validated models using the kernels shown in Figure 8.2 on two multi-core platforms that are available in our lab (Tile-Gx36 and Xeon Phi).

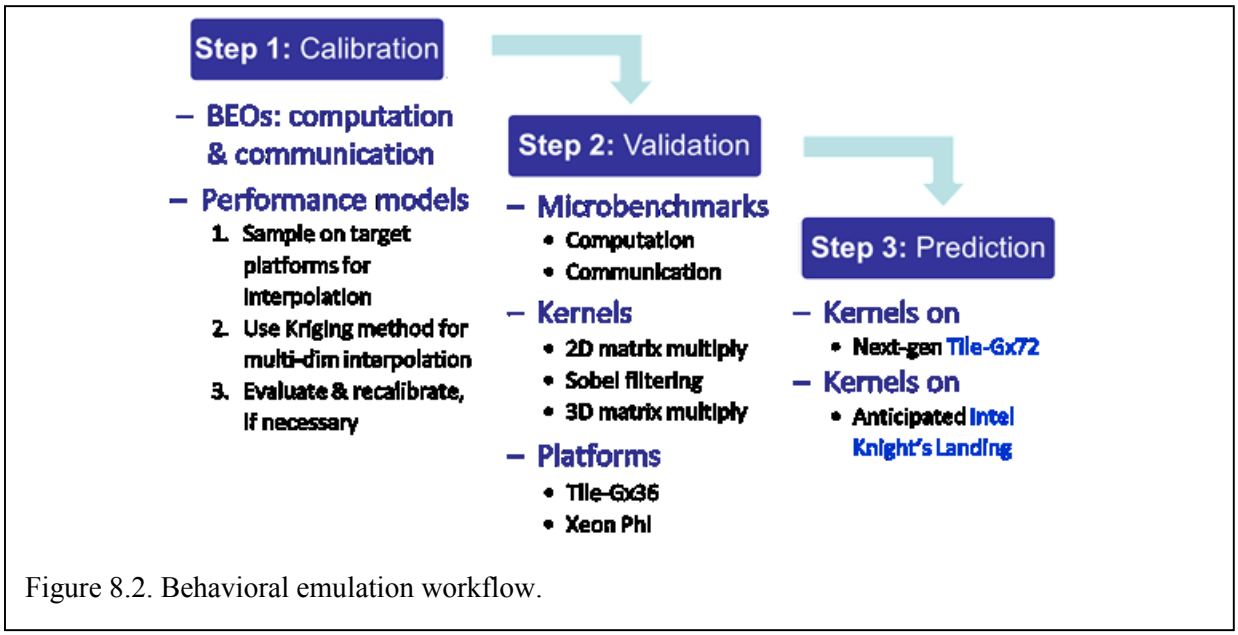

Figure 8.2. Behavioral emulation workflow.

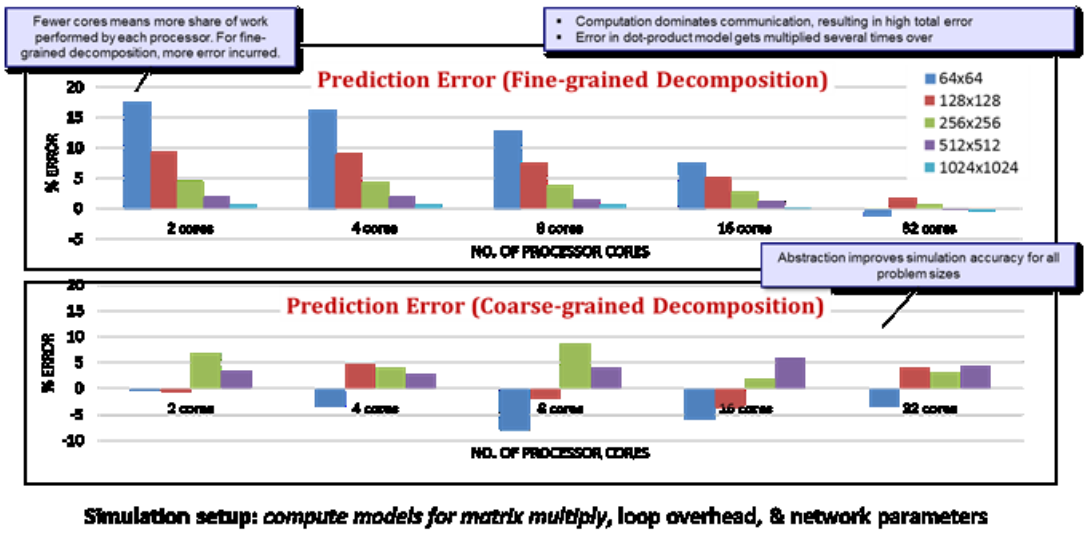

Figure 8.3. Validation of different problem sizes of 2D matrix multiply on mesh device (Tile-Gx36).

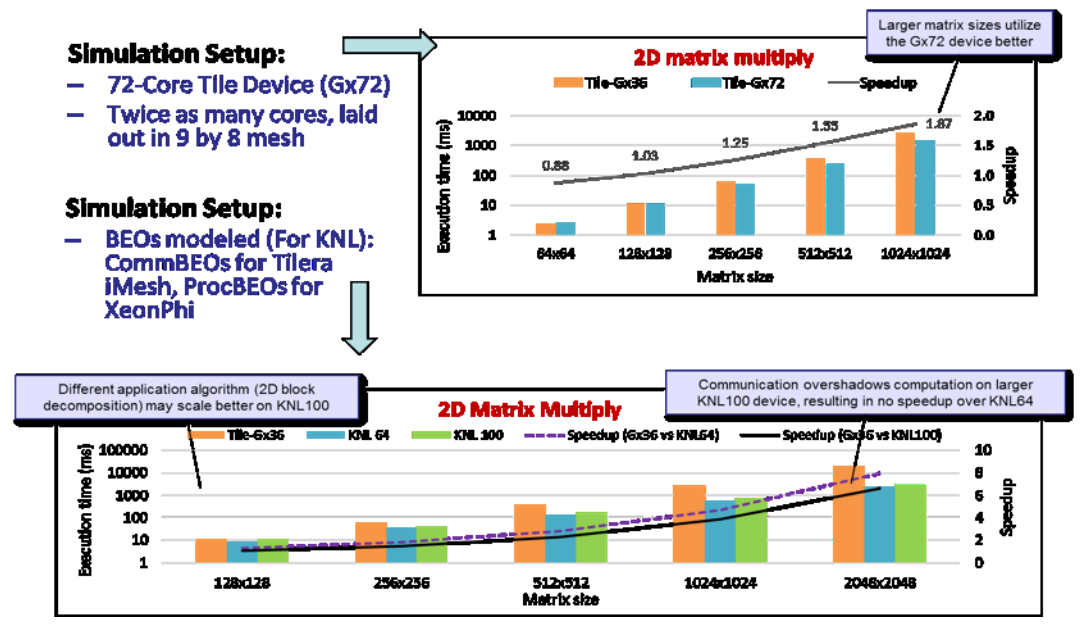

Figure 8.4. Prediction for Tile-Gx72 & Intel KNL.

Example results of model validation for parallel 2D matrix multiply on a Tile-Gx36 platform are shown in Figure 8.3.

The results in Figure 8.3 show that a higher abstraction of compute details improves simulation accuracy at a one-time cost of training effort. Also the accuracy of simulation is a function of domain, number of samples, and other Kriging parameters.

After the models have been validated (with an acceptable margin of error), these models can be extended based on information about future apps and notional future architectures (e.g., roadmap and NDA information from FastForward 2 vendors). Using these models, we can perform design-space exploration and predict application performance on anticipated and notation systems. Figure 8.4 shows two such predictions: (a) extending the model of an existing 36-core Tile device to the next-generation 72-core Tile device, (b) extending the model of an existing Intel Xeon Phi device to notional 64-core and 100-core Knights Landing (KNL) devices connected via an anticipated mesh interconnect.

The progress for Research Thrust T1 (Behavioral Characterization) in 2014 is summarized as follows:

- Explored, designed and developed BEO models for the TileGx-36 many-core processor. The device presents SIMD cores and on-chip 2D-mesh communication network which is useful to evaluate the robustness of our approach. We scaled the TileGx-36 model to study the performance of TileGx-72, the largest many-core device available from Tilera.
- Using the validated TileGx-36 models, we designed and calibrated BEOs using extensive benchmarking for notional many-core processors devices with on-chip 2D mesh network. Each core (ProcBEO) in the notional device represents an Intel Xeon Phi core, the x86-based AMD Opteron core, IBM Power 7 core, or the ARM A9 core.
    - Note: the 64-bit Xeon processing core in Xeon Phi (Knights Corner) is representative and important for studies in high-end computing, but its interconnect (bidirectional ring) is not, whereas the interconnect in Tile processors (i.e., 2D mesh) is more representative for studies in high-end computing and likely what Knights Landing will feature, and thus our mix.
- With input from the CCMT physics team, we focused on modeling AppBEOs for key application kernel closely related to CMT-nek (e.g., spectral element solver). Simulations of these AppBEOs running on various ProcBEOs are important in developing CMT-nek code targeted towards notional Exascale systems.
    - Note: Experimentation with these application kernels involves the determination of optimal parameter selection and problem decomposition given a specific system configuration. This data is of critical importance to the physics group in developing CMT-nek code targeted towards notional exascale systems.
- We have successfully developed BE prototypes in VisualSim and C++. Our C++ based SMP simulator supports easy representation of application BEOs, automated training of architecture BEO models, and fast simulation.

Plans for 2015:

- Publish in a conference the behavioral characterization work; investigate the effect of coarse-grained vs fine-grained decomposition of computation kernels on the accuracy of simulation predictions.
- With input from FastForward2 vendors, develop architecture BEOs representative of anticipated Exascale node architectures.

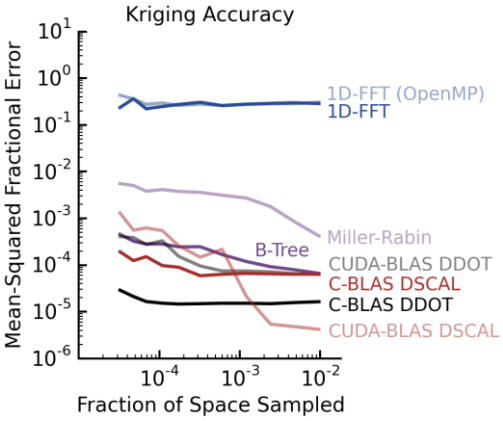

Figure 8.5. Kriging accuracy for single-input benchmarks.

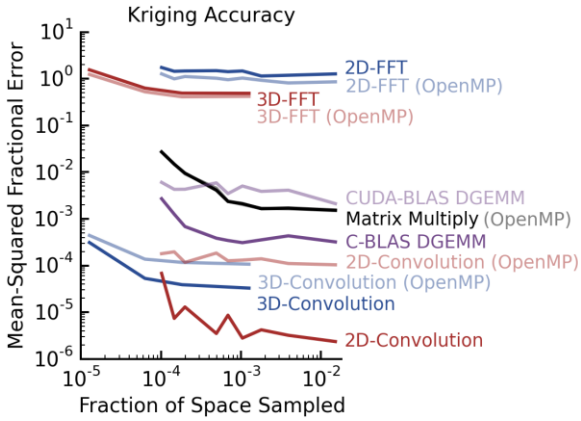

Figure 8.6. Kriging accuracy for multi-input benchmarks.

- With input from the CCMT Physics and CS teams, develop application BEOs for key computation and communication kernels in CMT-nek; Provide feedback to Physics and CS teams on varying key application parameters (effecting workload distribution) and their effect on application performance on various architectures.
- With input from the CCMT UB team, determine an effective experimental procedure which can be used to calibrate the simulator with measured results.
- With input from the CCMT UB team, determine an approach for characterizing uncertainty in the simulator based on measurement and simulation uncertainties.

8.2.2 T2: Parameter Estimation

Behavioral emulation necessitates that simulators do not perform cycle-accurate (or otherwise complex) operations, but the simulator must still know the time required for a specific application to run on the simulated architecture. The goal of the Parameter Estimation research thrust is to estimate values of specific non-integral numerical parameters (e.g., computation time) before or during simulation without having access to real generators (e.g., the target CPU) of those parameters. Our approach is to take a number of representative samples (of the parameter of interest) on the target platform or simulator, within the expected domain. Using these samples, we can interpolate the time required for any particular operation at simulation. The current focus of these parameters is computation time, but communication time, power, and other information will be necessary in future iterations of our simulation platforms. The current method of interpolation is Kriging, and the present research focuses within this task are to determine optimal Kriging parameters, and good methods of sampling within the domain.

The progress for this research thrust in 2014 is summarized as follows. We have created a substantial set of benchmark software for use in surrogate-modeling experiments. Using these experiments, we have determined the efficacy of Kriging, as a computation-time surrogate, by evaluating its capability to predict the time required for each benchmark under a variety of conditions. Figures 8.5 and 8.6 show sample results illustrating the accuracy of Kriging for a

variety of computational applications, demonstrating that Kriging is a sufficient computation-time surrogate for a large class of applications.

The plans for the Parameter Estimation research thrust in 2015 is summarized as follows:

- Publish in a conference the results of the Kriging work.
- Begin investigation of methods that may supersede Kriging for the classes of benchmark for which Kriging is not well suited.
- Work tightly with the CCMT UB team to find methodological improvements for current sampling and interpolation techniques.

8.2.3 T3: Synchronization & Congestion Modeling

Behavioral emulation relies on BEOs and their interactions to model a system under study. In order to model and simulate extreme-scale systems up to and including Exascale, it is vital to study various synchronization schemes and their simulation speed vs. scalability tradeoffs. Based upon this analysis, a fast and scalable technique can be formulated to maintain causality and inherent parallelism. Of equal importance is the ability to model with reasonable accuracy the communication fabric behavior while balancing simulation speed and scalability. Adapting synchronization and congestion-modeling techniques to support simulation experiments with millions of behavioral objects is essential for fast and scalable emulation of Exascale systems. After developing and experimenting with the basic modeling methods for communication in Research Thrust T1 (Behavioral Characterization), we have begun the exploration of methods for synchronization and congestion modeling.

Progress in 2014 thus far includes:

- BEO models feature vendor-specific communication (Tilera iLib) and OpenMP at micro-scale, and MPI at meso- and macroscales; researched various network congestion modeling techniques to balance the accuracy, speed, and scalability in our multiscale simulation approach.
- Verified BE simulations of on-chip iMesh network (without a strict timing model) for various application case studies in VisualSim and SMP simulator prototypes.
- Used Xeon Phi ProcBEOs and Tile-Gx36 mesh-network CommBEOs to predict the performance of applications of several notional many-core devices, including those representative of the anticipated Intel Knights Landing device (results will be included in the planned conference manuscript on behavioral characterization).

Plans for 2015:

- Explore synchronization techniques that can ignore causality or are optimistic; Research various network congestion modeling techniques to balance the accuracy, speed, and scalability in our multi-scale simulation approach.
- Adopt and integrate with an open-source, scalable, well supported simulation backend optimized for the purpose of behavioral emulation; Continue experimentation with our in-house SMP simulator.

- Identify communication patterns used in CMT-nek, understand the communication algorithms, and develop abstract representation for AppBEOs.

8.2.4 T4: Reconfigurable Architectures

It is our belief that it is necessary to go beyond the computing power of conventional parallel software simulator efforts to study the performance and scale of Exascale systems. The goal of this research thrust is to investigate, develop, and optimize an infrastructure to provide FPGA-centric hardware-supported behavioral emulation of extreme-scale systems up to and including Exascale. We have focused primarily on the challenge of efficiently mapping the developed BEOs from previously discussed thrusts onto the University of Florida's Novo-G reconfigurable supercomputer. Over the past year we have developed an initial single-FPGA prototype (NGEEv1) that is functionally identical to the current SMP BE software simulator and capable of offering significant performance improvement based on initial testing (for three cases), as shown in Figure 8.7. The top table shows the results for an existing 6X6 (36-node) Tile processor. The second table shows the results for the next-generation 72-node Tile processor (prediction error is unknown since the next-gen Tile device does not exist yet). The third table shows results for an anticipated notational Intel Knights Landing processor. The left side of each table shows the simulated time by our hardware emulator and the corresponding prediction error, both consistent with the software SMP simulator. The right side of each table shows simulation time (i.e., simulation execution time) for the FPGA emulator compared with the execution time of the SMP simulator. As you can see, at least for these three data points, FPGA-centric hardware-supported behavioral emulation show great promise, capable of offering on the order of 100x performance improvement.

We have begun initial planning for a redesign of the initial prototype (NGEEv2) that takes into account planned directions for the core BE modeling methodology over the next year and beyond. These efforts require the exploration of various techniques for BE mapping to a reconfigurable-computing infrastructure including, but not limited to, creating a fabric of BEOs with intra- and inter-FPGA BEO communication strategies, device-level optimization on BEO architectures to maximize BEO density, investigation of resource sharing for performance and scalability tradeoffs at various granularities, exploration of system-level communication, and control and coordination strategies, while exploiting a multidimensional interconnect fabric between hundreds of FPGAs. Below is a summary of T4 progress over the past year and proposed areas of focus in the coming year.

Progress in 2014:

- Continued development of functional NGEEv1 prototype.
  - Current results are based on un-optimized designs (i.e., max-resource implementation); current core density of 90 for Stratix IV, 256 for Stratix V with each core contains one each of appBEO, procBEO, & commBEO; appBEO scripts stored in on-chip block RAMs as memory initialization files (MIFs); Proc interpolation resources replaced with MIF pre-processing; Explicit emulation of target communication fabric without

congestion modeling; Separate management plane fabric collecting management tokens for postmortem analysis (e.g., simulation visualization).

- Conducted performance comparison with software-based SMP simulator for actual Tile-GX36, notional Tile-GX72 devices (micro-scale), and notional Knights Landing mesh architecture executing matrix multiply and image filtering mini-apps running on single Stratix-IV FPGA; Observed ~100x performance improvement for case studies selected.
- Performed first stages of multi-FPGA development for NGEEv1; Developed new multi-FPGA addressing scheme; Performed scalability prediction experiment on Novo-G# upgrade (i.e., Stratix-V devices with direct FPGA interconnect in 3D-torus topology).
- Initial planning for next NGEE design (NGEEv2).

Plans for 2015:

- Extend SMP simulator performance comparison with NGEEv1 to new set of system architectures (e.g., anticipated Intel Xeon Phi KNL) and new set of CMT-nek centric application case studies.
- Scale NGEEv1 prototype to multiple FPGAs; Updated scalability prediction experiments of NGEEv1 on Novo-G# incorporating results from multi-FPGA experiments.
- NGEEv2 design and initial prototype.

	Tile 6x6	Simulated Time (Consistent with SMP results)	Prediction Error	FPGA Simulation Time	SMP Simulation Time	Speedup
	2D MM 1024x1024 across 36 cores	2.82×10^9 ns	-0.35%	35.7us	4.82ms	~135x
	Sobel 800x600 image across 36 cores	9.27×10^7 ns	-2.61%	54.2us	8.08ms	~149x
Next-gen	Tile 9x8	Simulated Time (Consistent with SMP results)	Prediction Error	FPGA Simulation Time	SMP Simulation Time	Speedup
	2D MM 1024x1024 across 72 cores	1.66×10^9 ns	To be determined	81.1us	10.24ms	~126x
	Sobel 800x600 image across 72 cores	2.58×10^7 ns	To be determined	102.0us	17.94ms	~176x
Anticipated	KNL 9x8	Simulated Time (Consistent with SMP results)	Prediction Error	FPGA Simulation Time	SMP Simulation Time	Speedup
	2D MM 1024x1024 across 72 cores	5.87×10^8 ns	To be determined	81.1us	10.24ms	~126x
	Sobel 800x600 image across 72 cores	1.37×10^7 ns	To be determined	102.0us	17.94ms	~176x

Figure 8.7. Performance comparison between SMP behavioral simulator and Novo-G behavioral emulator.

9. 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	
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.

10. Publications

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. 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.
3. 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.
4. 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.
5. 2014 67th Annual Meeting of APS Division of Fluid Dynamics.
    a. 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)*.*
    b. 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.
    c. 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.
    d. 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.
    e. Jackson, T., Sridharan, P., Zhang, J., & Balachandar, S. (2014). Shock propagation over a deformable particle. Bulletin of the American Physical Society, 59.
6. 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.
7. Matsumura, Y. and Jackson, T.L. (2014). Numerical simulation of fluid flow through random packs of polydisperse cylinders. Physics of Fluids, Vol. 26, 123302.
8. 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.

9. 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.
10. 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.
11. Mankbadi, M. R., & Balachandar, S. (2014). Multiphase effects on spherical Rayleigh-Taylor interfacial instability. Physics of Fluids (1994-present), 26(2), 023301.
12. 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.
13. Sridharan, P., Jackson, T.L., Zhang, J. and Balachandar, S. (2015). Shock interaction with 1-D array of particles in air. Journal of Applied Physics, Vol. 117, 075902.
14. 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, accepted.
15. Annamalai, S., Rollin, B., Ouellet, F., Neal, C., Jackson, T.L. and Balachandar, S. (2015). Effects of initial perturbations in the early moments of an explosive dispersal of particles. Proceedings of the IWPCTM 2014, ASME Journal of Fluids Engineering, submitted.
16. 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.
17. Park, C., Haftka, R.T. and Kim, N. (2015). The Effect of Ignoring Dependence between Failure Modes on Evaluating System Reliability. Submitted to Structural Multi-disciplinary Optimization.
18. Park, C., Choi, J., and Haftka, R.T. (2015). Teaching a Verification and Validation Course using Simulations and Experiments with Paper Helicopter, submitted to The Journal of Verification, Validation and Uncertainty Quantification (2015).
19. Banerjee, T. and Ranka, S. (2015). Performance and Energy Benchmarking of Spectral Element Solvers. To be submitted to SuperComputing 2015.
20. Banerjee, T. and Ranka, S. (2015). Power modeling and predicting method on multi-core processors. In preparation.
21. 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, (accepted).
22. Annamalai, S. and Balachandar, S. (2015). Generalized Faxen's theorem in time-domain for unsteady viscous compressible flows over a sphere, to be submitted to Phys. Rev. Lett.
23. Schwarzkopf, J.D., Balachandar, S. and Buttler, W.T. (2015). Compressible Multiphase Flows, in Handbook of Multiphase Flows, Ed. Stathis Michalides, CRC Press (to appear).
24. McGrath, T.P., St Clair, J.G. and Balachandar, S. (2015). A compressible multiphase model for dispersed particle flows with application from dense to dilute regimes, submitted Phys. Fluids.

25. Annamalai, S. and Balachandar, S. (2015). Acoustic levitation force on a finite-sized rigid particle, droplet or bubble in a viscous compressible medium. Submitted to Phys. Fluids.
26. Cook, C.R., Balachandar, S. and Chung, J.N. (2015). A modified Characteristic-Based Split (CBS) projection method for compressible flows of real fluids, to be submitted to J. Comput. Phys.

11. Conferences and Presentations

1. Jackson, T.L. 67th Annual Meeting, Division of Fluid Dynamics, American Physical Society, San Francisco, CA, Nov 2014. ``Shock propagation over a deformable particle".
2. S. Annamalai, C. Neal, F. Ouellet, B. Rollin, T.J. Jackson, & S. Balachandar, *Numerical Simulation of Explosive Dispersal of Particles in Cylindrical Geometry,* IWPCTM 2014, San Francisco, California, USA.
3. B. Rollin, S. Annamalai, C. Neal, T.J. Jackson, & S. Balachandar, *Numerical Study of Explosive Dispersal of Particles,* DFD2014, Bulletin of the American Physical Society, Vol. 59, Number 20 (November 2014).
4. Chanyoung Park, Joo-Ho Choi, Raphael T. Haftka, "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*
5. Chanyoung Park, University of Florida, "Dakota – An Overview", CCMT Seminar, 3:00 P.M., Tuesday, April 8, 2014, 221 MAE-A.
6. Behavioral Modeling & Emulation of Extreme-scale Systems, Workshop on Modeling & Simulation of Systems & Applications, Seattle, WA, August 13-14, 2014.
7. Invited Talk. Prof. Nam-Ho Kim. "Prioritizing Efforts to Reduce Uncertainty in System Models based on Uncertainty Budget", *Seoul National University, Seoul, Korea, July 7, 2014.*
8. A.M.D. Jost_ and J. Zhang, "Numerical Study of Intermittent Laminar Bubble Bursting and Vortex Shedding on an NACA 643-618 Airfoil", AIAA Meeting, 2015.

12. Workshops Held or Attended

1. Dr. Chanyoung Park. Attended "A Short Course on Uncertainty Quantification", Stanford CA, June 2-4 2014.
2. Presented "Dakota - Tutorial", CCMT workshop, 4:00 P.M., Thursday, February 19, 2015, Feb 19 2015.
3. Salishan Conference, April 2014.
4. 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.
5. Deep Dive Workshop. Held at the University of Florida on Feb 3-4, 2015.
6. "Good Software Engineering Practices and Beyond" Workshop - Internal workshop - organized by Bertrand Rollin - Macroscale team, held Feb 19, 2015.

7. 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.

13. Students and Staff Internships

13.1 Internships Completed

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.
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

13.2 Internships Planned

1. 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.
2. Chris Hajas, M.S. Student, US, ECE, Dr. Herman Lam, core. May 18-August 18, 2015 at LLNL with Dr. Maya Gokhale.
3. Christopher Neal, Ph.D. Student, US, MAE, Dr. S. Balachandar, core. June 14-August 20, 2015 at LLNL with Dr. Kambiz Salari.
4. Carlo Pascoe, Ph.D. Student, US, ECE, Dr. Herman Lam, core. Will intern summer 2015 at LLNL with Dr. Maya Gokhale.
5. Giselle Fernandez, Ph.D. Student, Argentina, MAE, Drs. Haftka and Kim, core. Internship planned to Sandia for Fall 2015.

13.3 Internships Not Yet Planned

1. Subramanian Annamalai, Ph.D. Student, India, MAE, Dr. S. Balachandar, cost share (Internship not required)
2. Angela Diggs, Ph.D. Student, US, MAE, Dr. S. Balachandar (other funding, internship not required)
3. Fred Ouellet, Ph.D. Student, US, MAE, Dr. S. Balachandar, core
4. Kasim Alli, MS/Ph.D. Student, US, ECE, Dr. Greg Stitt, core

5. Ryan Blanchard, BS, US, ECE, Dr. Herman Lam, core
6. Hugh Miles, BS, US, ECE, Dr. Greg Stitt, core
7. Dylan Rudolph, Ph.D. Student, US, ECE, Dr. Greg Stitt, part core and part leveraged
8. Prashanth Sridharan, Ph.D., US, MAE, Dr. Balachandar, core
9. David Zwick, Ph.D., US, MAE, Dr. Balachandar, core
10. Dr. Tania Banerjee, Center Staff, US, core (one week required)
11. Goran Marjanovic, Ph.D., MAE, Dr. Balachandar, core
12. Yash Metha, Ph.D., MAE, Dr. Balachandar, cost share (internship not required)
13. Yiming Zhang, Ph.D., China, MAE, Drs. Haftka and Kim, cost share (internship not required)

14. NNSA Laboratory Interactions

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
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