REPORT OF THE WORKSHOP ON FUTURE DIRECTIONS IN EXTREME SCALE COMPUTING FOR SCIENTIFIC GRAND CHALLENGES

Held 09-10 January 2020
Texas Advanced Computing Center
The University of Texas at Austin

Omar Ghattas, George Biros, Dan Stanzione,
Rick Stevens, John West

TACC
TEXAS ADVANCED COMPUTING CENTER

LEADERSHIP-CLASS
COMPUTING FACILITY

This material is based upon work supported by the National Science Foundation under Award No. 1925096. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Report of the Workshop on Future Directions in Extreme Scale Computing for Scientific Grand Challenges

NSF OAC-1925096
Texas Advanced Computing Center

Title	Report of the Workshop on Future Directions in Extreme Scale Computing for Scientific Grand Challenges
Document ID	LCCF-TEC.001
DCA Level	N/A
Type	Uncontrolled
Owner	John West Dan Stanzione

Change History

Date	Author	Version	Description of Change
Feb 2020	John West	.00a,b	Initial document
Mar 2020	Craig Stair, Dan Stanzione	.00c	Comments and cleanup
Mar 2020	George Biros, Omar Ghattas	.00d	Revisions for clarity in the problem setup
May 2020	John West	.01	Submit for CDR

Table of Contents

1 Executive Summary 4

2 Introduction and Background 5

3 Scientific Grand Challenges: A Summary of Inputs 6

3.1 Incorporation of Molecular-scale Processes 6

3.2 AI-Enhanced Science 8

3.3 Hypersonic Flight 8

3.4 Global Dynamics and the Great Earthquakes 9

3.5 Modeling Thermonuclear X-ray Bursts 9

3.6 Quantum Materials Engineering 10

3.7 Physics of Fundamental Particles 12

3.8 Turbulent Flows 12

4 Science Requirements for the LCCF 13

4.1 Increased Computing Power 14

4.2 Support the Full Lifecycle of a Computational Analysis 14

4.3 Use of Machine Learning Methods in Scientific Computing at Scale 14

4.4 Radical Application Change is Hard 14

5 Appendix A: Agenda 16

6 Appendix B: List of Science Attendees 18

7 Appendix C: Attendee Position Papers 19

1 Executive Summary

The NSF Leadership Class Computing Facility (LCCF) project held a workshop on January 9-10, 2020 at TACC on "Future Directions in Extreme Scale Computing for Scientific Grand Challenges". The goal of the workshop was to help determine the requirements of science and engineering applications on the cyberinfrastructure that must be developed to make significant progress toward resolving grand challenge problems.

Inputs from the workshop are being used to help identify the science requirements for the design of the computational capability of the LCCF. Although the grand challenges identified by each science domain are radically different, we found that they embody a limited set of common themes that will drive the design of the LCCF:

Numerous individual science tasks have demand for vastly greater computational scale and time; as such, a 10x baseline improvement in application performance should be considered a minimum for the Facility.

Much of the work of extracting science and engineering results from the workload happens outside the main simulation or analysis run, and is done in analysis of the produced results later, over a much longer time. Support for this "expanded" workflow at the appropriate scale is therefore critical for the facility, including throughout the data lifecycle.

The workload is evolving, with increasing incorporation of methods from artificial intelligence and machine learning in the science workloads, and increased emphasis on throughput at scale – this work is in addition to, rather than replacing simulation at scale.

Gradual evolution of code is possible, with proper incentives and sustained investments. Radical or rapid change to software in order to support new hardware is hard; the transition may outlive the hardware.

It is important to note that we are early in the requirements process, and that requirements gathering will continue over the course of the LCCF design period. The design drivers documented here will be updated in response to information gathered over the coming years. Additional workshops and community events will be held during the preliminary design period to address existing large science users not captured in this report, and will also seek to expand the requirements documented here related to the increased growing use of AI and ML in science and engineering research, as well as other disciplines that are emerging as significant users of NSF CI.

2 Introduction and Background

As deployed computational capability has increased, science and engineering research has adopted increasingly complex models in search of more accurate simulations. Models today are seeking to incorporate more of the wide range of physical processes that we observe in the real world: in the 1980s aerospace engineers modeled the flow of a laminar fluid over an airfoil in order to evaluate the potential of a new design. Today, engineers are using suites of codes coupled together that model turbulent aerodynamics, the structural response of the airframe, propulsion, acoustic propagation, and, in the case of hypersonic flight, the effects of mechanical ablation and multi-species chemical reactions that happen between the airframe and the atmosphere at high temperatures and pressures. Modeling this range of physical processes requires multiphysics models (often in multiple different application codes) simulating processes that happen on multiple time and space scales.

Increased computational capability has also expanded the ways in which we can use simulations to get a better understanding of the limitations of our science codes and our knowledge of the physical processes underlying them. Uncertainty quantification, optimal design, optimal control, and simulation-based decision-making all rely on principled exploration of the high-dimensional parameter spaces characterizing these applications.

Recent years have witnessed a dramatic increase in the availability of observed data—from satellites deployed at a global scale to monitor the planet's continents and oceans to clusters of microsensors providing data about engine combustion dynamics in real-time to electron microscopy images of nanoscale molecular structures. Computational models are increasingly assimilating this data into simulations to infer unknown or uncertain components of the models. In data-driven modeling, these components can represent initial/boundary conditions, sources, or parameters in a mechanistic model, often in the form of differential equations. In data-intensive modeling, the data can be used to infer reduced, statistical, or machine learning models for processes that are either too complex to simulate or for which mathematical models of sufficient fidelity do not yet exist.

Increased computer performance has come at the cost of increased hardware complexity and an accompanying increase in the difficulty of writing applications that take advantage of their potential power. Architectures and applications have traditionally evolved together, with each new generation of hardware requiring incremental adaptations of applications to exploit the new performance. The dramatic increase in complexity in recent years, however, has motivated the reconsideration, and even complete redesign, of the algorithms underlying the applications as well.

It is against this backdrop that the NSF Leadership Class Computing Facility (LCCF) project held a workshop on January 9-10, 2020 at TACC on "Future Directions in Extreme Scale Computing for Scientific Grand Challenges." The goal of the workshop was to help determine the requirements of science and engineering applications on the cyberinfrastructure that must be developed to make significant progress toward resolving grand challenge problems. In the 1980's the term "Grand Challenge" was given to a set of scientific problems whose solution would be enabled by HPC resources. Several Grand Challenge Problems were identified to motivate the accelerated funding of HPC activities in response to foreign competition.

In this report we use the term grand challenge in roughly the same way, but with the understanding that the high performance computer itself is only a part of the solution: they also require breakthroughs in algorithms, computational science, data management and visualization, software engineering, scientific workflows, and system architecture as well as a community of expertise built around the technological capabilities in these areas to ensure that the technologies (hardware and software) can be translated into practice.

The workshop was held under the auspices of the LCCF planning project at the Texas Advanced Computing Center (TACC) at The University of Texas at Austin. The NSF's design criterion for the LCCF is that the facility must provide at least 10x the capability of the Frontera system, deployed by TACC and #5 on the TOP500[1] list at the time of this writing. The LCCF is being designed following the NSF's Major Research and Experimental Facility Construction (MREFC) process for large-scale, long-life, research infrastructure. If awarded, the facility will operate for at least ten years.

Inputs from the workshop are being used to help identify the science drivers for the design of the computational capability of the LCCF, documented in the present report. Attendees (see Appendix B) were invited to submit a one-page position paper describing a grand challenge problem whose solution will have significant scientific or engineering impact and which requires next generation HPC resources. All submitted papers are included as an appendix in this report. Selected authors were also invited to present short talks at the workshop.

Attendees represented many of the disciplines that make the most sophisticated demands on cyberinfrastructure (CI), including: computational fluid dynamics, computational chemistry and materials, geosciences and geophysical flows, mechanical engineering, acoustics, seismology, quantum chromodynamics, chemically reacting flows, turbulence, biochemistry, astrophysics, plasma, and fusion physics. Rounding out the computational science focused attendees were those from enabling computational technologies, including computer science, applied mathematics, parallel algorithms, numerical optimization, artificial intelligence and machine learning, and software engineering.

3 Scientific Grand Challenges: A Summary of Inputs

In the text that follows we provide brief discussions, selected from papers discussed at the workshop, on representative scientific grand challenges. The discussions identify both the science challenges and the advances in CI and computational science needed to address them. We do not cover each submission in detail, but rather have selected a representative subset; the full text of all workshop submissions is included in the appendix.

3.1 Incorporation of Molecular-scale Processes

In many domains, computational tools of sufficient accuracy for design and decision-making require incorporation of highly-refined information about the molecular interactions which then drive macro-scale processes of interest. For example, the presence of atmospheric aerosols directly impacts the amount of radiation that reaches earth, cloud droplet formation, and optical properties of the atmosphere. These processes in turn affect weather, climate, food production, and human health, among other processes, that shape life on earth.

[1] www.top500.org

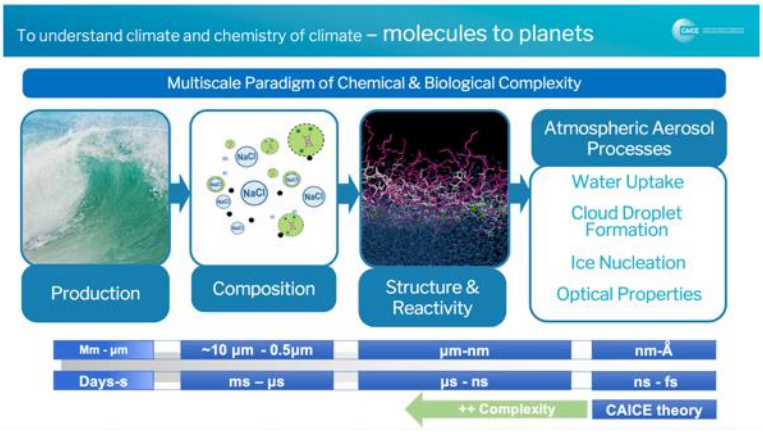

A central challenge in incorporating molecular-level processes into simulations with results of interest at the macro scale is difference among scales. For example, atmospheric aerosol simulations start with length dimensions ranging from nanometers to Angstroms and time dimensions from nanoseconds to femtoseconds, and range all the way up to meters and days for the macro changes that result from those molecular processes. In biological processes, understanding the molecular basis of disease to create better medicines starts at the molecular scale (Angstroms and femtoseconds) and proceeds to tissue scale, millimeters and seconds. Crossing scales creates significant challenges in communication, requires advanced frameworks in which to construct the simulation (and connect the components), and has to reconcile the coupling of codes and processes created for studying one phenomenon (molecular dynamics) into a workflow that applies the results to what is potentially a completely different phenomenon (tissue morphogenesis).

For multiscale simulations (whole cell, aerosol, etc.) we need improved coupled simulation modalities. Consider the interaction of particles at one scale with a phenomenon simulated on a mesh at a different scale – the simulation models have to be tightly coupled, with the mesh and the particles interacting with one another in real time, requiring two different kinds of simulations to interact by executing in separate but related processes. But the two simulation modalities may take wholly different approaches, e.g., different domain decomposition algorithms for performance and scaling. This is a challenge for computational science (common decomposition schemes and methods to pass information between scales) in addition to a science challenge among the related domains.

There are also significant challenges in computing at the molecular scale itself. Increasing complexity of the solution requires more particles, which drives a need for more processor nodes, and GPUs are playing an increasing role in this area. Common codes in this area such as NAMD2, GROMOS, and LAMMPS need fast inter-processor communications to simulate billions of atoms. Workflow and cyberinfrastructure frameworks are needed that will enable– at a single scale – researchers to enhance their ability to model and test combinations of different parameters and perturbations of a system of interest. Simulation and computing frameworks that allow real-time integration of data (e.g., from sensors or 4D microscopy data) are also needed in this community, as are better mechanisms for sharing difficult-to-gather measurements or novel methods with other researchers in the community.

3.2 AI-Enhanced Science

The intense commercial interest in machine learning and artificial intelligence have led to an interest in studying the application of these approaches to computational science. Traditional statistical and supervised machine learning approaches have a longer history in areas such as cosmology but, more recently, deep learning has been increasingly employed in a range of science application areas. For example, tokamak data can be used to train models to predict disruptions. In fluid dynamics a physical simulation and deep learning coupled framework called CFDNet has been designed to predict the primary physical properties of a fluid using a single convolutional neural network at its core; speedups between 2 and 7x have been observed without relaxing convergence constraints.

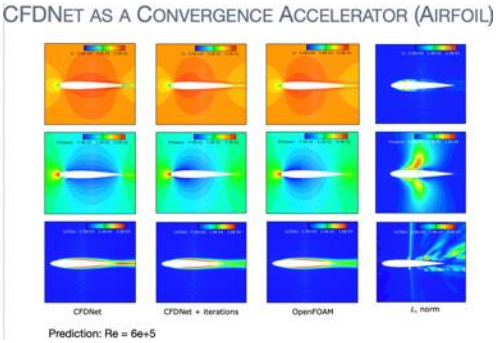

Experimental and observational data, as well as simulation data, are sources for data-derived models. Some efforts, such as neural tissue mapping from electron-microscopy, align well with mainstream machine learning successes in image classification. Others, such as deconstructing detector events and tokamak profile data, demand different kinds of neural networks. Deriving predictive reduced models from simulation data, such as turbulence and lattice QCD, requires quantitative validation against theory or against very expensive numerical simulations.

Effectively supporting AI and related methods in an HPC environment can present challenges, however. I/O systems need to be designed with the ability to stream in large amounts of data from disk or from outside data sources in the case of real-time observational data. The workflows combining physics-based models with these methods can also create a very complex simulation system, and the mass-market driven software in this area is not typically created with scaling to thousands of nodes in mind.

3.3 Hypersonic Flight

Hypersonic vehicles, characterized by sustained flight at speeds in excess of five times the speed of sound, are a critically needed technology that depends on the tight integration of chemistry, material science, structural dynamics, fluid mechanics, propulsion, and dynamics and control to be successful. Predicting their performance and robustness through modeling, simulation and analysis will require immense computational resources and innovative computational science to support the range of spatial scales (10^{-9} to 10^{1} m) and temporal scales (10^{-10} to 10^{2} s) that characterize the coupled interaction of the high-speed external flow field and the chemically active propulsion system with the vehicle's structural and material response and degradation along the vehicle's trajectory.

Effective modeling of hypersonic flight requires coupling codes from different scientific domains—density functional theory, direct simulation Monte Carlo, computational structural, thermal, fluid, and electricity and magnetism dynamics—to predict and model the evolution of the vehicle and its components and to estimate the uncertainty in the predictions.

An early and limited analysis of a generic hypersonic vehicle considered the detailed fluid dynamic response of a Mach 6 flow across a surface was completed using three Python coupled-codes for the fluid, thermal, and structural responses. No material degradation model was used

and the domain of interest was a very small fraction of the overall vehicle – both approximations which made the computation tractable but which rendered the result of limited usefulness. The simulation utilized half of Frontera for 48 hours to simulate 4 billion grid points (5 unknowns per point), and is the largest simulation of its type performed to date. Yet, in total even this simplified analysis only covered about 10^{-1} seconds in "real time"; at this rate, 36,000 days (about 100 years) would be required to simulate a 15-minute flight.[2]

3.4 Global Dynamics and the Great Earthquakes

Global tectonic plate motions in the Earth's crust are thought to be primarily driven by the negative buoyancy within subducted slabs, but there is little agreement in the scientific community on the strength and nature of resisting forces. Plate motions are likely related to how forces are transmitted through the slab, across the plate interface, and into the converging plate. Failure to reach consensus on the force balance of plate tectonics and mantle flow leaves fundamental questions unanswered, including causes for spatial variability in great earthquake occurrence. A great earthquake – one measuring 8 or higher on the Richter scale – happens roughly once per year and, when near areas with large populations, can cause devastating loss of life and economic activity. Despite the severity of the risk, the ambiguity surrounding their occurrence results in substantial uncertainty in the assessment of seismic hazards.

The magnitude 8.2 Chiapas earthquake of 2017 killed 98 people and affected an estimated 1.5 million people in Chiapas; 41,000 homes were damaged. Buildings in the city closest to the epicenter were reduced to rubble and 1.8 million people in Mexico City lost electricity.

Making progress will require a new generation of time-dependent plate-mantle models in which the details of plate boundaries are resolved in time and space. This will require global forward and inverse viscoelastic models to determine variations in mechanical coupling in subduction zones using plate motions, space-geodetic displacements, the record of large earthquakes, and the along strike variation in the depth of oceanic trenches. The models must have high resolution (e.g., locally 1 km or less) and fully incorporate the extreme nonlinearity in viscosity between the slab and mantle, hinge zone and rupturing interface between subducting and over-riding plates.

This will require iterative solvers for nonlinear saddle point problems with $\mathcal{O}(10^9)$ unknowns arising from discretization of partial differential equations in viscous fluid dynamics. While being much more complex, the resulting sparse systems share similarities with the HPCG benchmark: *i.e.*, they are memory-bound, involve collective communication (dot-products), and require multilevel preconditioners that utilize significant global communication. Since on future architectures, flop rates will likely be even faster compared to memory access, this is a problem where hardware may drive the development of numerical algorithms towards higher-order discretizations. Scientists will also use observations from plate motion, earthquakes and topography to deduce which mechanisms are able to explain our planet's evolution, and which

[2] At Mach 5, the level above which flight is considered to be hypersonic, flight speed is 1 mile per second; at Mach 9, 15 minutes is enough time to fly from D.C. to L.A.

are not. Such inverse problems require many repeated simulations, which is why strong parallel scaling of future algorithms and implementations is of particular importance.

3.5 Modeling Thermonuclear X-ray Bursts

Modeling X-ray bursts (XRBs) resulting from thermonuclear explosions allows scientists to probe the physics of dense nuclear matter, which feature heavily in the physics of neutron stars, the Big Bang, and the behavior of heavy atomic nuclei. Such modeling requires multiscale and multiphysics simulations, with spatial scales ranging from the size of the compact object (e.g., black holes, white dwarfs, etc.) down to the burning or dissipation scale. Temporal scales range from the hours of evolution that brought the star to the brink of explosion to the seconds-long explosion. Models must incorporate hydrodynamics, nuclear reactions, gravity, thermal conduction, magnetic fields, and radiation.

A problem of interest in studying XRBs is modeling the nucleosynthesis as a burning front propagates laterally through the accreted fuel layer on a neutron star. DOE leadership systems at $\mathcal{O}(10PF)$ enabled the first full hydrodynamics simulations in both the lateral and vertical directions, yielding insight into complex dynamics of the flame front with realistic microphysics. The simulations were challenging, and used approximations to both the length scales (using artificially high neutron star rotation rates to compress the lateral scales) and time scales (boosting the flame speed). New DOE systems at $\mathcal{O}(100PF)$ permit elimination of these approximations as the code runs $5-10x$ faster on GPUs than CPUs.

The increased power of next generation supercomputers will help to enable routine 3D simulations covering a larger portion of the neutron star surface, and include more realistic physics (such as coupled radiation), but improvements in algorithms (SDC coupling and higher order methods) are also needed. These simulations will be the first global simulations of burning fronts on neutron stars that capture the realistic physics and enable direct connections between simulation and observation.

3.6 Quantum Materials Engineering

Much of today's technology relies on the ability to control matter at the atomic scale. Solar photovoltaics, portable electronics, wireless communications, and electric vehicles all crucially depend on the unique properties of artificial materials manufactured with atomic precision. In this technological landscape, quantum-mechanical modeling of materials and extreme-scale computing have the potential to revolutionize the way we create disruptive new technology.

Over the last decade materials design from first principles has been making significant progress in validating electronic structure codes based on density functional theory (DFT), creating databases of DFT calculations for all known crystalline materials, introduction of data science techniques in materials science, and development of a software ecosystem for the automation and documentation of DFT computational workflows.

Most current efforts in these areas focus on properties that can be obtained from standard DFT calculations, namely calculations where electron-electron interactions are described by local or semi-local functionals, and electron-nucleus interactions are described by taking nuclei as

classical particles immobile in their equilibrium positions. Unfortunately, these calculations offer limited predictive insight into many functional properties that underpin energy and information technology, from light-matter interactions to temperature-dependent charge and spin transport.

One of the grand challenges of *ab initio* materials design is to overcome these limitations by systematically incorporating electron correlations, quantum fluctuations, and realistic temperature conditions in accurate post-DFT calculations of real materials.

As an example, the prediction of the electrical conductivity of materials, a key design parameter for electronic, thermoelectric, optoelectronic, photovoltaic, and plasmonic devices, requires (i) high-accuracy many-body GW band structures to obtain reliable carrier effective masses, and (ii) a fine sampling of all carrier-phonon scattering processes to obtain reliable carrier relaxation rates. The resulting calculations are already challenging for binary semiconductors like gallium nitride, with only a handful of atoms per crystalline unit cell. For more complex systems, like the recently discovered hybrid organic-inorganic perovskites which count almost fifty atoms per simulation cell, these calculations are almost prohibitive on the best available supercomputing architectures.

Screening a single compound (cubic perovskite $CsPbI_3$) at one temperature data point on MareNostrum 4 takes 96 hours on 3,840 cores. In a typical materials design project, it is common to screen several tens of different compounds, doping levels, and elastic strain conditions. As a result, at this time a systematic high-throughput screening of advanced functional properties such as the electrical conductivity is beyond reach.

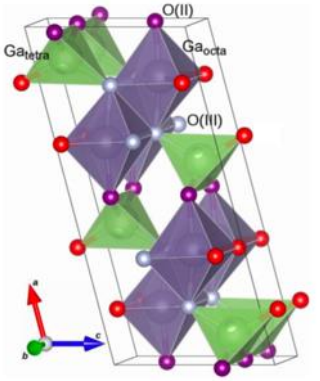

Figure 1 Ga2O3 calculation with EPW on 1,120 cores of Stampede2.

In order to be competitive on the front of advanced materials design with today's codes, and to deliver tangible impact in the form of new compounds, new functionalities, and new devices, the screening time for new compounds must be reduced. For example, improved scaling in the application (EPW) and a larger supercomputer would permit runs on 20-30k cores, reducing the total compute time to a few wall clock hours and enabling systematic materials screening over timescales on the order of days instead of weeks or months. Also, the runs tend to demand about 2x the memory that is typically configured per core, which forces the reservation of more cores than would otherwise be necessary: larger memory nodes are needed.

Looking to the future, local DFT functionals are gradually being replaced by non-local hybrid functionals, with an order-of-magnitude increase in computational cost. Simulating the electronic and optical properties of functional materials require GW quasi-particle methods on top of DFT, with an additional order-of-magnitude increase in computational cost over where we are today. Within the community itself, databases of the Materials Genome Initiative[3] contain almost exclusively DFT data. The calculations need to be expanded to include a broader array of materials properties such as optical absorption and luminescence, electrical and thermal transport, vibrational spectroscopy, photoelectron spectroscopy. All of these calculations are

[3] Materials Genome Initiative for Global Competitiveness. [Washington, D.C.]: Executive Office of the President, National Science and Technology Council, 2011.

much more computationally intensive than DFT and will increase the demand for capabilities to support materials development.

3.7 Physics of Fundamental Particles

An important scientific challenge in high energy or elementary particle physics is to find evidence of Beyond the Standard Model Physics. This might be done by producing new particles with higher energy collisions than previously studied, but it can also be done using very high precision measurements that test precise predictions of the Standard Model. The latter approach is being used to study the physics of the bottom quark by the LHCb experiment at the Large Hadron Collider at CERN, and the Belle II experiment at the KEK Laboratory. Lattice QCD calculations provide essential input to these experiments as they are the most precise way to make the required Standard Model predictions. As the experiments are ongoing and will increase their precision, it is necessary to increase the precision of the predictions of lattice QCD.

Lattice QCD involves a numerical treatment of the continuum field theory in which space time is treated as a discrete four-dimensional grid with lattice spacing a. This introduces systematic errors that are reduced when a is reduced. The bottom quark is quite heavy, and there are errors that depend on am_b, where m_b is the bottom quark mass. A grand challenge problem is to reduce the lattice spacing to 0.03 fm, so that the errors from the bottom quark are well controlled. We are currently doing calculations with a lattice spacing of 0.06 fm. We would use quark masses that are well tuned to the physical quark masses for the up, down, strange, and charm quarks. This approach largely eliminates the error from the chiral extrapolation that is needed when quarks heavier than in nature are used.

A lattice spacing of $a = 0.06fm$ results in a grid size of 963 × 192. Halving the lattice spacing (to get to $a = 0.03fm$) results in a grid size of 1923 × 384: 16 times larger than is currently used in production on Frontera. However, the computational cost would increase by more than a factor of 16 as the condition number of the sparse matrix which must be inverted increases, requiring more iterations, and the step size will also be decreased in the molecular dynamics-like evolution required to generate new gauge configurations. Research is currently underway into algorithmic advances that might reduce computational requirements such as multigrid solvers, and force gradient integrators.

Refining mass estimates will also require refining our understanding of the uncertainty of those estimates, resulting in a parameter space exploration. We have made a very rough estimate of the time required to create and analyze 1000 configurations for a variety of mesonic decays containing bottom and charm quarks. The analysis jobs are already in production on Frontera for the $a = 0.06fm$ ensemble, so one can estimate total time based on increasing the volume and an assumption about how many more iterations will be required with the smaller lattice spacing. Gottleib's group have done some test running for generation of a 1923 × 384 configuration on Frontera, but did not run long enough to equilibrate, so roughly know the performance of the solver, but don't know how many iterations will be required once equilibrated. A very rough estimate from these tests is 2.7 million node-days to generate 1000 configurations on Frontera – roughly a full year of dedicated system time on all 8,000 nodes.

3.8 Turbulent Flows

One of the most important challenges in fluid mechanics is to predict the effects of turbulence in a wide variety of flow situations. The inability to reliably predict the average behavior of many turbulent flows in systems such as aircraft, motor vehicles, ships, combustion systems, the

atmosphere and ocean, and many more, greatly hinders the ability to design or otherwise make decisions regarding such systems. As an example of the importance of turbulence, consider that approximately 28% of total US energy consumption is for the transportation of goods and people, and that the vast majority of that energy is dissipated in the turbulent flows involved.

A likely reason for the unreliability of current turbulence models (called Reynolds-Averaged Navier-Stokes or RANS models) is that they do not carry a sufficiently rich state-space of turbulence statistics. The most common RANS models track the mean velocity, the turbulent kinetic energy and the rate of dissipation of the energy (or similar). However, it is known that the dynamics of turbulence with these same characteristics can vary widely depending on details of the structure of the turbulence. The question then arises as to what is a sufficient set of statistical descriptors of that structure that would allow the determination of the dynamics. Promising states have been proposed, but their adequacy as a state space describing turbulence dynamics has not been thoroughly evaluated.

Arguably, a reason that such richer turbulence state spaces have not been explored is a lack of supporting data. The data that would be required includes values of potential statistical state quantities obtained from a wide variety of turbulent flows of varying complexity. The only viable source for such data is direct numerical simulations (DNS), which involve the solution of the Navier-Stokes equations on very large grids requiring extensive HPC resources. For example, a relatively recent DNS of a simple wall-bounded flow involved a simulation with 250 billion degrees of freedom and required 400 million core hours on Mira at Argonne. Many simulations of a wide range of more complex turbulent flows would be needed to determine an adequate turbulence state space and to inform a model for the evolution of such a state space.

The computational grand challenge, then, is to perform such DNS and the required analysis. In addition to a first-order computational challenge requiring extreme scale HPC, this represents a first-order data science challenge. Among the data science objectives would be identifying the necessary state dependencies, determining additional flow conditions for which simulations are needed, and learning the high-dimensional functional relations required to model the evolution of the turbulence state. Such a focused integrated computational and data analytic effort has the potential to revolutionize the modeling of turbulence, resulting in vastly improved abilities to design turbulent flow systems or otherwise make decisions regarding them.

4 Science Requirements for the LCCF

The grand challenges highlighted above and the many more included in the appendix (all of which informed this report) are unique to each scientific discipline and in some cases, even to the individual areas within a discipline. Despite the differences in the specific goals of each challenge, it is possible to extract a set of common computational demands from the scientific goals that will drive the design of the LCCF. The common requirements shown below are the first set extracted from our initial requirements meetings in November of 2019 (at SC19) and this workshop at TACC in early 2020. Requirements gathering will continue over the course of the LCCF design period and these drivers will be updated in response to information gathered over the coming years.

4.1 Increased Computing Power

Numerous individual science tasks have demand for vastly greater computational scale and time; as such, a 10x baseline improvement in application performance should be considered a minimum for the Facility.

A common theme among the science challenges is the need for increased raw computing power – this is a requirement across the board. In some cases, the science team is ready for a certain type of computing power, as documented in e.g. Zingale and Cheatham in the case of GPUs, and others are experimenting with the best use of GPUs as documented in Chang, Berzins, and Giustino. But in nearly all cases workshop participants indicated that a fundamental requirement for advancing the state of the knowledge in their field is increased computing power; as documented e.g. in Gottleib with a need for 2.7 million node days for a single parameter search, and Bodony with the requirement for 100 years of computing to simulate 15 minutes of hypersonic flight.

4.2 Support the Full Lifecycle of a Computational Analysis

Much of the work of extracting science and engineering results from the workload happens outside the main simulation or analysis run, and is done in analysis of the produced results later, over a much longer time. Support for this "expanded" workflow at the appropriate scale is therefore critical for the facility, including throughout the data lifecycle.

The CI community has long been aware that although the supercomputer is a necessary enabler of scientific discovery, it is not by itself sufficient. Thus, larger centers such as TACC offer a broad range of capabilities and expertise from data analysis and visualization to data publishing and code optimization.

Expanding the boundaries of science with leadership-class computing systems strains all of the processes that enable scientists to prepare for simulation runs and analyze and share the results as documented in e.g. Bozdag, Zingale, Guistino, Chang, Aparna, Yeung, Moser, and others. As the CI and science community are learning from pre-exascale systems currently deployed, transferring and staging very large input data sets, identifying the right parameters to run jobs at a scale never before achieved, possibly modifying the code itself to support the new scale, checkpointing to preserve work, validating accuracy and proper code behavior (i.e., is the problem converging appropriately) during a run to avoid wasting time and computer resources, and then managing, analyzing, visualizing, storing, and potentially sharing the results all require careful reconsideration as the scale and complexity of computational simulations increase.

4.3 Use of Machine Learning Methods in Scientific Computing at Scale

The workload is evolving, with increasing incorporation of methods from artificial intelligence and machine learning in the science workloads, and increased emphasis on throughput at scale – this work is in addition to, rather than replacing simulation at scale.

As documented in e.g. Aparna, Chang, Bodony, and Bozdag, use of tools from machine learning and artificial intelligence is moving beyond exploration into early adoption. An important emerging use case for these methods in science simulation lies in augmenting existing (typically PDE-based) workflows with ML-based submodels to speed up the convergence or serve as

subgrid-scale model or closure or constitutive relation. Alternatively, ML models are becoming increasingly used as surrogates for (possibly) more accurate direct numerical simulations when it is prohibitive to execute them repeatedly (e.g., for parameter space exploration), or when the mathematical "physics-based" model is not sufficiently known. It is clear from workshop submissions that this is an emerging part of the workload and future systems must be able to provide appropriate hardware and software resources in addition to specialized expertise to integrate these disparate modes of computing. Specialized support will also be needed to modify the mass-market driven software in these fields to be relevant to the very large numbers of nodes that will be used in science simulations.

4.4 Radical Application Change is Hard

Gradual evolution of code is possible, with proper incentives and sustained investments. Radical or rapid change to software in order to support new hardware is hard; the transition may outlive the hardware.

Several of the workshop attendees with mature code bases were careful to emphasize that the code changes required to take advantage of the potential power of new architectures are not something that many science teams are well-suited to implementing, as documented in e.g., Gottleib, Park, Sundar, Yeung, Giustino, and Zingale, among others. To a first order approximation, workshop attendees prefer a model they are already familiar with; if that is not possible, then they emphasized that advanced code support is a necessary part of the LCCF to assist them in (i) making the needed changes and (ii) making them fast enough to be within the useful life of the machine. This quote from Zingale is illuminating:

> *"New architectures bring new programming challenges, and to date we've managed to keep pace with these developments. A key point is that funding needs to be provided to aid in preparing codes, not just for the science, especially for novel architectures."*

5 Appendix A: Agenda

Thursday January 9	
08:00 - 08:30	Breakfast
08:30 - 09:00	Welcome and charge - Dan Stanzione
09:00 - 9:15	Workshop/report structure - Omar Ghattas
09:15 - 09:30	Vignette #1: Mike Gurnis/Georg Stadler - Mantle Convection
09:30 - 09:45	Vignette #2: Steve Gottlieb - Lattice QCD
09:45 - 10:00	Vignette #3: Rich Vuduc HPC Performance
10:00 - 10:15	Vignette: #4: Ebru Bozdag - Global seismic inversion
10:15-10:45	Break
10:45 - 11:00	Vignette #5: Hari Sundar - Gravitational waves
11:00 - 12:00	Discussion Session #1: Algorithms of Today and Tomorrow (Moderators: George and Omar)
12:00 - 13:15	Working lunch: Application Survey and Discussion
13:15 - 13:30	Vignette #6: Rommie Amaro - Biochemistry
13:30 - 13:45	Vignette #17: Matthias Ihme - Turbulent combustion
13:45 - 14:00	Vignette #8: Martin Berzins - PDEs
14:00 - 14:15	Vignette #9: Michael Zingale - Astrophysics
14:15 - 14:30	Vignette #10: Feliciano Giustino - Materials
14:30 - 14:45	Vignette #11: C. S. Chang - Magnetic Fusion
14:45 - 15:15	Break
15:15 - 15:30	Vignette #12: Aparna Chandramowlishwaran
15:45 - 16:45	Discussion Session #2: HPC Performance (Moderator: Dan and Tommy) HPC Performance (Moderators: Dan and Tommy) Architecture (GPU or not GPU?), Libraries (PETSc, etc), Profiling/Debugging Tools
16:45 - 17:15	General Discussion

Friday January 10	
08:00 - 08:30	Breakfast
08:30 - 08:45	Vignette #13: Dan Bodony - Hypersonics
08:45 - 09:00	Vignette #14: George Park - Aircraft LES
09:00 - 09:15	Vignette #15: Eckart Meiburg - Environmental flows
09:15 - 09:30	Vignette #16: Tim Williams - AI for Science
09:30 - 09:45	Vignette #17: PK Yeung - Turbulence
09:45 - 10:00	Vignette #18: Patrick Heimbach - Ocean data assimilation
10:00 - 10:15	Break
10:15 - 11:30	Discussion session #3:The LCCF beyond the machine (Dan) What services should an LCCF provide beyond a large machine? Support Services, Data Lifecycle Management, Workforce/Diversity, Training, Outreach, Gateways, Software
11:30 - 12:00	Discussion Session #4: Life with the LCCF (Omar and George) What programs would make sense to support the applications community in the age of LCCF?
12:00 - 13:15	Working lunch: Discussion and workshop summary (Moderators: George, Dan, Omar)
Adjourn	

6 Appendix B: List of Science Attendees

Last	First	Email	Institution	Field
Berzins	Martin	mb@sci.utah.edu	U Utah	Computational math, HPC
Biros	George	biros@oden.utexas.edu	UT Austin	Mechanical engineering
Bodony	Daniel	bodony@illinois.edu	Illinois	CFD, aeroacoustics
Bozdag	Ebru	bozdag@mines.edu	Colorado School of Mines	Seismology
Chandramowlishwaran	Aparna	amowli@uci.edu	UC Irvine	HPC
Chang	Choong-Seock	cschang@pppl.gov	Princeton Plasma Physics Laboratory	Plasma physics
Cheatham	Thomas	tec3@utah.edu	U Utah	Computational biochemistry
Couch	Sean	couch@pa.msu.edu	Michigan State	Astrophysics
Ghattas	Omar	omar@oden.utexas.edu	UT Austin	Geosciences
Giustino	Feliciano	fgiustino@oden.utexas.edu	UT Austin	DFT / materials
Gligoric	Milos	gligoric@utexas.edu	UT Austin	Software Engineering
Gottlieb	Steve	sg@indiana.edu	Indiana U	QCD
Gurnis	Mike	gurnis@gps.caltech.edu	Caltech	Geophysics, geodynamics
Ihme	Matthias	mihme@stanford.edu	Stanford U	CFD, reacting flows
Maechling	Phil	maechlin@usc.edu	USC	Comp seismology
Meiburg	Eckart	meiburg@engineering.ucsb.edu	UCSB	CFD, geophysical flows
Moser	Bob	rmoser@oden.utexas.edu	UT Austin	CFD, turbulence
Park	George	gipark@seas.upenn.edu	U Penn	CFD
Stadler	Georg	stadler@cims.nyu.edu	NYU	Computational math, optimization
Sundar	Hari	hari@cs.utah.edu	U Utah	Computer science, HPC
Vuduc	Rich	richie@cc.gatech.edu	GA Tech	Computer science, HPC
Willams	Tim	zippy@anl.gov	ANL	Plasma fusion
Yeung	PK	pk.yeung@ae.gatech.edu	GA Tech	CFD, turbulence
Zingale	Michael	michael.zingale@stonybrook.edu	Stony Brook U	Computational astrophysics

7 Appendix C: Attendee Position Papers

Moving Beyond Present Notions of Exascale

Martin Berzins and Tamara Bidone, SCI Institute, University of Utah.

Given that GPU-based architectures will take us to Exascale somewhere in the 2021/2022 time frame and that they will do an amazing job of delivering hitherto unheard of computational power to bear on many challenging problems the need is to define example applications and to discuss what is needed next. One large class of applications, that will find it harder to make effective use of these architectures consists of problems that have a relatively low flops to data ratio. Many of these applications have roughly one flop per word. **Examples of such applications** are:

Weather codes with complex atmospheric physics and data assimilation.

Combustion code in which preserving the physical correctness of the solution is important.

Extreme physics codes with very severe shock waves and complex chemistry.

Manufacturing product design codes with very complex geometries are typical of many “real” applications.

More generally low-to medium order methods for partial differential equations codes.

Iterative linear solvers based upon methods such as conjugate gradients.

In all such applications t**he low flops to word ratio is a function of the underlying algorithm and not its coding.**

Such applications are limited on the GPU by the memory bus speed and also on some CPUs by unattainable vector peak rates. One exception is the ARM/Fujitsu chip which while being 7x slower than a summit gpu has the same memory bus speed. The question for the future is whether or not it is possible to design and build computer architectures through co-design that are even more applications friendly.

An example Exascale problem with one flop per word from combustion

One particular application being readied for Exascale on the DOE Aurora system using the Uintah Arches code is the Boiler Simulation Facility (BSF), is a 13 × 3 × 3 m3 tangentially fired boiler firing approximately 5000 pounds of coal per hour. The facility represents the combustion and radiation physics of a full size boiler, while being small enough to take high-quality, well-characterized measurements. In order to resolve a significant fraction of the energy spectrum millimeter of sub-millimeter grid spacing is required. At a spacing of 0.1mm and assuming that say there are 100K "Summit-type" GPUs in an Exascale machine this would mean roughly 1000 128 cubed patches per GPU.

An example Exascale problem from cell biology again with one flop per word

Understanding single cell behavior is essential for research in biological sciences. The ability of cells to sense mechanical stimuli, or cell mechanosensing, regulates a plethora of physiological functions, including migration, wound healing and tissue morphogenesis. The response of cells to mechanical forces occurs by converting mechanical stimuli into biochemical pathways, corresponding to changes in functional relations between cytoskeletal and adhesion proteins. Mechanistic understanding of cell mechanosensing via cytoskeleton and adhesions remodeling remains elusive due to a lack of high resolution information from experiments. Computational modeling offers a powerful solution to characterize cell mechanosensing because it allows to simulate in space and time the reorganization of cytoskeletal and adhesion proteins into assemblies that sense, sustain, generate and transmit forces and deformations. Such models involve complex particle systems. For the study of semiflexible actin filaments, a Brownian Dynamics based algorithm have been developed. The models sowed that: the relative concentrations of the various protein components determine the morphology of the cytoskeletal network in different environments; stress transmission and sustainability in contractile actomyosin bundles, such as stress fibers and lamellipodium actin arc. We developed a computational model of contractile ring assembly using particle-based phenomenological models, which revealed that: interactions between filamentous proteins and molecular motors generate contractile rings for successful cytokinesis, depending on the initial distribution of the molecular components; properties of the cytoplasmic reticulum govern the displacement of attachment proteins at the cell cortex. Such particle models involve billions of particles to simulate a single cell and the extension to multiple cells would require corresponding larger particle sizes.

Modeling, Simulation and Analysis of Hypersonic Vehicles: A Grand Challenge Problem for Next Generation HPC Resources

Daniel J. Bodony (bodony@illinois.edu)
Blue Waters Associate Professor
Department of Aerospace Engineering
University of Illinois at Urbana-Champaign

Hypersonic vehicles, characterized by sustained flight at speeds in excess of five times the speed of sound, are a critically needed technology that depend the tight integration of chemistry, material science, structural dynamics, fluid mechanics, propulsion, and dynamics and control to be successful. Predicting their performance and robustness through modeling, simulation and analysis (MS&A) will require immense computational resources and innovative computational science to support the range of spatial scales (10^{-9} to 10^{1} m) and temporal scales (10^{-10} to 10^{2} s) that characterize the coupled interaction of the high-speed external flow field and the chemically active propulsion system with the vehicle's structural and material response and degradation along the vehicle's trajectory. Successful MS&A will require coupling codes from different scientific domains—density functional theory, direct simulation Monte Carlo, computational structural, thermal, fluid, and electricity & magnetism dynamics—to predict and model the evolution of the vehicle and its components and to estimate the uncertainty in the predictions.

An early and limited analysis[1] of a deflected and compliant control surface of a generic hypersonic vehicle is shown in the figure below. The simulation considered the detailed fluid dynamic response of a Mach 6 flow impinging on a thermally and structurally-compliant control surface that was deflected into the flow and was performed using three Python coupled-codes for the fluid, thermal, and structural responses. No material degradation model was used and the domain of interest was a very small fraction of the overall vehicle. The simulation utilized the entire Frontera machine for 48 hours, generated approximately 10^{-1} seconds worth of "real time" data, and represents the largest simulation of its type. Expanding simulations of this type to meet the needs of hypersonic vehicles is a grand challenge for next generation HPC resources.

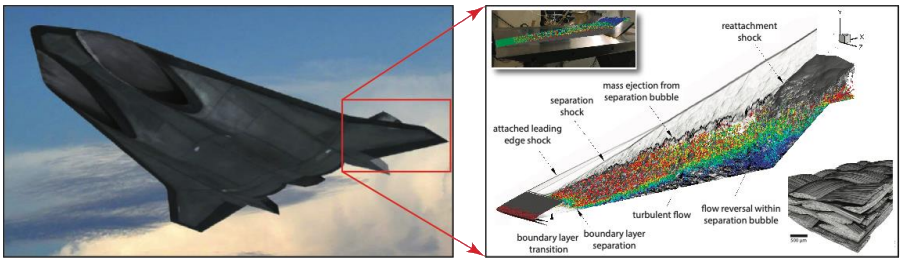

Left figure: generic hypersonic vehicle. Right figure: direct numerical simulation of Mach 6 flow impinging on a compliant and deflected control surface with corresponding experimental validation (upper inset). Lower inset shows carbon weave of control surface.

[1] Work supported by the AFOSR (Dr. Ivett Leyva, PM). Experiments performed in NASA Langley's 20 inch Mach 6 tunnel by Dr. Stuart Laurence, University of Maryland.

Global Full-Waveform Inversion of the Earth's Interior

Ebru Bozdag

Department of Geophysics, Colorado School of Mines

The foundational pillars of seismological research are theory, computation, and data. During an era of unprecedented growth of data united with the theory under the capabilities of high-performance computing (HPC) offer tremendous opportunities for seismologists to refine our understanding of the inner dynamics of our planet. Imaging Earth's interior based on global full-waveform inversion has been a long-standing challenge in seismology, mainly due to its intense computational requirements. The adjoint-state method, an iterative full-waveform inversion technique, efficiently incorporates the full nonlinearity of 3D seismic wave propagation in iterative inversions (e.g., Tromp et al. 2005; Fichtner et al. 2006; Plessix 2009; Virieux & Operto 2009). "Adjoint tomography" is a powerful technique for the following reasons: (1) the full nonlinearity of 3D seismic wave propagation is taken into account; (2) 3D background models are used to compute Fréchet derivatives and the models are iteratively updated; (3) as a result of (1)–(2), the amount of usable data steadily increases from iteration to iteration, thus enabling the extraction of more information from seismograms, in the ideal case the use of entire three-component seismograms with every wiggle; and (4) the crust and mantle are inverted jointly avoiding the commonly used crustal corrections which can potentially bias, specifically our inference of the upper mantle anisotropy.

After the construction of the first seismic tomographic models (Aki et al. 1977, Dziewonski et al. 1977; Sengupta & Toksoz 1977), the first global mantle models came out in 1984 (e.g., Woodhouse & Dziewonski 1984; Nataf et al. 1984). Meanwhile, Lailly (1983) and Tarantola (1984) formulated and adapted the theory of adjoint methods for seismology to capture the full physics of seismic wave propagation in imaging. Mainly due to computational challenges, it took until 2009 to see the first applications of adjoint methods in regional- and continental-scale earthquake seismology (e.g., Tape at al. 2009; Fichtner et al. 2009, Zhu et al. 2012). The first-generation global adjoint models of the Earth's mantle, GLAD-M15 (Bozdag et al. 2016), then its successor GLAD-M25 (Lei et al. in revision) constructed based on a larger data set of 1420 earthquakes came out more than 30 years after the original theory, using Oak Ridge National Laboratory's "Titan" and "Summit" systems. GLAD-M15 and GALD-M25 are both elastic models with transverse isotropy confined to the upper mantle. Next-generation models aim to capture better physics in the inverse problem with adequate parameterization (i.e., anelasticity, general anisotropy, source parameters, internal topographic variations, etc.). Eventually, the goal is to take advantage of the entire database of ever-increasing broad-band (from land and ocean-bottom seismometers) and emerging datasets (from MERMAIDs, fiber optics, etc.) of both earthquakes and ambient noise in inversions to improve global data coverage. Analysis and incorporation of all available data in inversions bring out the necessity of properly coupling big data with HPC systems as well as developing and maintaining computational tools for discovery and reproducibility in science. At the scale of globe, going down to 1 Hz in global simulations is the ultimate goal to perform whole-Earth inversions including the core which is the subject of exascale computing together with assimilating all available data. Exascale resources will also allow for proper resolution and uncertainty quantifications, which are the major bottlenecks in large-scale tomographic problems to assess the robustness of the constructed images.

Position paper

Author: C.S. Chang, Princeton Plasma Physics Laboratory, Princeton University
Scientific Area: magnetic fusion
Projects: SciDAC High-fidelity Boundary Plasma Simulation (HBPS),
ECP High-fidelity Whole-device Modeling (ECP WDMApp)
ECP Co-design Center for Particle-based Applications (CoPA)
Scientific application code name: XGC

Design, construction, and operation of ITER and future fusion demo reactors are extremely expensive and time consuming. First-principles based predictive simulation can save the time and the money from the design phase to the operation phase, possibly accelerating the fusion reactor commercialization by a significant factor.

Performance of magnetic fusion plasma is determined by multi-scale self-organization of multi-physics phenomena. Large free energy from the hot plasma induces not only the machine-size scale dynamical evolution and instabilities (~10m), but also the microscale turbulence that spans all the way to the electron gyro radius (~10μm). Time scale spans over a wider range (100s – 10^{-8}s). Since they interact together to produce the properties of the burning plasma, they ultimately need to be simulated together self-consistently using first-principles governing equations. In the edge plasma close to the enclosing wall, the problem becomes even more demanding due to the interaction with material wall that produces neutral particles with atomic interaction with the plasma, and the complicated geometry conditions. Number of charge species involved can be up to 100 in ITER.

To fit the first-principles magnetic fusion simulations into today's ~10-100 PF leadership class computers, large-scale instabilities are simulated independently from microscale turbulence, the electron gyro radius scale turbulence is simulated only in a thin local annulus, and the number of charge species can only be allowed to be a few. Also, the full 6D (3X+3V) simulation of the whole-volume tokamak simulation is not possible in today's leadership class computers. The dimension is reduced to 5D (3X+2V) by limiting its validity to the slower physics than ion gyration motion.

Due to the large-scale scientific data produced by the simulations (~100PF per day from a 100PF computer), data needs to be analyzed, visualized and reduced in situ and on-the-fly before they are moved to the filesystem. On-memory, in-situ application of machine learning tools is essential to achieve this goal.

We have found that the high-fidelity, 5D kinetic PIC code XGC performs equally well on both GPU (Summit) and CPU (Cori/Theta) accelerator architectures. The GPU usability is limited not only by the data-transfer latency, but also by its memory size. Due to the large-scale scientific data it produces, a fast I/O is a necessity. On-node NVRAMs are found to be highly helpful for this purpose. Node-to-node and node-to-filesystem MPI communication speed will be an important factor. We are looking forward to the Slingshot interconnect architecture on Perlmutter, Aurora, and Frontier. Some robust fault resistant and tolerant methods, via both software and hardware, will be greatly helpful as we go forward to exascale computing. We also find that a portability-enabling ecosystem, such as Kokkos, is extremely helpful. By leveraging Kokkos, a single-version XGC is performance-portable on both Summit GPUs and Cori/Theta GPUs.

Extreme Scale Computing: Continuing challenges in modeling the structure, dynamics and interactions of biomolecules across biologically relevant time and size scales

Author: Thomas Cheatham, University of Utah

For more than half a century, chemists and physicists have labored to atomistically model the structure, dynamics and interactions of biomolecules through molecular mechanics and various simulation approaches [1]. Such methods are widely applied to understand biomolecular function, to model biomolecular folding, elucidate conformational distributions and molecular interactions, and for computer aided drug and biomolecule design. Whereas simulations in the 90's pushed nanoseconds of dynamics, in recent years state of the art simulations on modest sized biomolecules push millisecond timescales. Yet, there is a desire to model larger biomolecular assemblies and to reach biologically relevant timescales approaching seconds or longer, which implies that the biomolecular simulation community could saturate any computational resource available (as they have consistently done for at least the last 30 years where a significant portion of the academically available computational resources have been used for biomolecular simulation). Great strides have been made in modeling larger assemblies such as viruses, organelles, ribosomes, and nucleosomes [2-6], yet even with very large computational resources simulations of such assemblies barely achieve a nanosecond of conformational sampling. Yet, for rather small model systems of small proteins or nucleic acid tetranucleotides or tetraloops, complete sampling often requires 100's of microseconds (which was only achieved in the past few years with supercomputers or specialized hardware) [7-10], and the challenge of complete sampling only gets harder (or more time consuming) as the molecular assemblies get bigger. Promising developments include the use of ensemble methods, better free energy methods, and further improvements in the underlying potentials and force fields. As the computational methods and representations improve, we continue to see growth in the long-tail of biomolecular simulators (often from experimental labs) who are perfectly happy with simulations on the nano- to micro-second timescale to provide insight into their favorite biomolecular system or design problem while a few pioneering groups with access to the most extreme resources will keep pushing the size and time scales while continuing to improve the methods. So, in summary, expect growth in demand for all types of computational resources from the long tail to extreme scale by our biomolecular simulation community (and p.s. our community loves GPU's as many of the community codes are well to highly optimized on GPU's, specifically NVIDIA GPU's, and it is unclear how well we will adapt to newer generation GPU's from alternative vendors or FPGA's).

1. Levitt, M., *The birth of computational structural biology.* Nat Struct Biol, 2001. **8**(5): p. 392-3.
2. Hadden, J.A. and J.R. Perilla, *All-atom virus simulations.* Curr Opin Virol, 2018. **31**: p. 82-91.
3. Lee, C.T. and R.E. Amaro, *Exascale Computing: A New Dawn for Computational Biology.* Comput Sci Eng, 2018. **20**(5): p. 18-25.
4. Jung, J., W. Nishima, M. Daniels, G. Bascom, C. Kobayashi, A. Adedoyin, M. Wall, A. Lappala, D. Phillips, W. Fischer, C.S. Tung, T. Schlick, Y. Sugita, and K.Y. Sanbonmatsu, *Scaling molecular dynamics beyond 100,000 processor cores for large-scale biophysical simulations.* J Comput Chem, 2019. **40**(21): p. 1919-1930.
5. Sanbonmatsu, K.Y., *Large-scale simulations of nucleoprotein complexes: ribosomes, nucleosomes, chromatin, chromosomes and CRISPR.* Curr Opin Struct Biol, 2019. **55**: p. 104-113.
6. Feig, M. and Y. Sugita, *Whole-Cell Models and Simulations in Molecular Detail.* Annu Rev Cell Dev Biol, 2019. **35**: p. 191-211.
7. Dror, R.O., R.M. Dirks, J.P. Grossman, H. Xu, and D.E. Shaw, *Biomolecular simulation: a computational microscope for molecular biology.* Annu Rev Biophys, 2012. **41**: p. 429-52.
8. Bergonzo, C., N.M. Henriksen, D.R. Roe, J.M. Swails, A.E. Roitberg, and T.E. Cheatham, 3rd, *Multidimensional Replica Exchange Molecular Dynamics Yields a Converged Ensemble of an RNA Tetranucleotide.* J. Chem. Theory Comput., 2014. **10**(1): p. 492-499.
9. Galindo-Murillo, R., D.R. Roe, and T.E. Cheatham, 3rd, *On the absence of intrahelical DNA dynamics on the μs to ms timescale.* Nature Commun., 2014. **5**: p. 5152.
10. Galindo-Murillo, R., D.R. Roe, and T.E. Cheatham, 3rd, *Convergence and reproducibility in molecular dynamics simulations of the DNA duplex d(GCACGAACGAACGAACGC).* Biochim Biophys Acta, 2015. **1850**(5): p. 1041-58.

Grand Challenge Computational Problem: Massive Stellar Explosions

Core-collapse supernovae (CCSNe) are among the most extreme laboratories for nuclear physics in the universe. Stellar core collapse and the violent explosions that follow give birth to neutron stars and black holes, and in the process synthesize most of the elements heavier than helium throughout the universe. The behavior of matter at supranuclear densities is crucial to the CCSN mechanism, as are strong and weak interactions. Beyond Standard Model behavior of neutrinos may also impact the CCSN mechanism. Despite the key role CCSNe play in many aspects of astrophysics, and decades of research effort, we still do not fully understand the details of these explosions and their observable outcomes. This is due in large part to the limitations of past and current computational platforms. Leadership-class computing hardware and software has the potential to dramatically accelerate progress on this important astrophysical problem.

CCSNe have historically been a grand challenge problem in computational astrophysics since the field emerged 60 years ago, always pushing the envelope of what was possible computationally. This is due to the combination of complex physical processes that are crucial to the CCSN mechanism and the vast range of scales that must be captured in order to simulate them. Simulating the collapse of a massive stellar core and its subsequent explosion requires modeling magnetohydrodynamics, general relativistic gravity, the equation of state of matter at supranuclear densities, nuclear reaction kinetics, and multidimensional neutrino radiation transport. This last aspect of the problem, neutrino transport, is the most challenging and the principle driver of the computational expense of CCSN simulations. In addition to physical complexity, the CCSN mechanism is fundamentally multidimensional and multi-scale. The neutrino-heated medium behind the CCSN shock is convective and highly turbulent, with Reynolds numbers on the order of 10^{17}, while simultaneous large scale asymmetric motion in the problem can fundamentally drive important dynamics. The range of scales of the full problem is vast, going from the radius of the star itself, about 10^8 km, down to the width of the proto-neutron star atmosphere where the neutrinos decouple from the matter, scales of 100s of m. This necessitates the use of some form of non-uniform gridding, such as adaptive mesh refinement (AMR). Only with petascale platforms have high-fidelity 3D simulations even been possible, but the problem is far from solved. Vastly more, and more sophisticated, 3D simulations are required, demanding beyond-petascale hardware and software.

Executing even a single high-fidelity 3D CCSN simulation is challenging, but fully addressing the CCSN problem requires many such simulations. CCSN arise from an enormous variety in initial conditions: stars of masses ranging from $\sim$7 to $\sim$100 solar masses, with wide variation in initial compositions, spin rates, and mass loss histories. In addition, the all-too-common presence of a binary companion can alter the evolution of CCSN progenitors in important ways, expanding the parameter space of initial conditions to include the characteristics of the secondary star and the properties of the binary orbit. All this means that the parameter space of CCSN initial conditions is of high dimensionality. We have only barely begun to sample even one dimension in this space.

A Grand Challenge Problem for next generation extreme-scale computing platforms then is to simulate a sufficiently large ensemble of CCSNe to adequately sample the parameter space of initial conditions with accuracy enough to enable predictions for the impact of CCSNe on the evolution of the Universe. On current petascale platforms such as *Theta* at ALCF, a single high-fidelity 3D CCSN simulations requires about 200k node-hours to 0.5 s of evolution following core bounce. In order to simulate enough time to fully capture the growth of the explosion energy and bulk of the nucleosynthesis, this time scale must be increased by roughly a factor of 10. Conservatively estimating that next-generation platforms will realize a per-node performance increase of 5x, this brings the cost of a single, long-time scale CCSN simulation to 400k node-hours. Each simulation should easily scale to 1000 nodes, meaning about 17 wall-clock days per simulation. Sampling the full parameter space of CCSN initial conditions will not be possible in 3D even on the next generation platforms, so we choose to limit the parameter dimensions to just three: progenitor mass, composition, and initial spin rate. Sampling this space with 10x3x3 simulations (90 total) would be

Extreme-scale computing to enable quantum materials engineering

TACC Workshop on Future Directions in Extreme Scale Computing for Scientific Grand Challenges

Feliciano Giustino
Oden Institute for Computational Engineering and Sciences and Department of Physics
University of Texas at Austin

Much of today's technology relies on the ability to control matter at the atomic scale. Solar photovoltaics, portable electronics, wireless communications, electric vehicles, the internet, all crucially depend on the unique properties of artificial materials manufactured with atomic precision. In this technological landscape, quantum-mechanical modeling of materials and extreme-scale computing have the potential to revolutionize the way we create disruptive new technology.

Over the last decade materials design from first principles has been making tremendous progress along four directions: (i) Validation and verification of electronic structure codes based on density functional theory (DFT). (ii) Creation of databases of DFT calculations for all known crystalline materials. (iii) Introduction of data science techniques, such as machine learning, in materials science. (iv) Development of a software ecosystem for the automation and documentation of DFT computational workflows.

Most current efforts in (i)-(iv) focus on properties that can be obtained from standard DFT calculations, namely calculations where electron-electron interactions are described by local or semilocal functionals, and electron-nucleus interactions are described by taking nuclei as classical particles immobile in their equilibrium positions. Unfortunately these calculations offer limited predictive insight into many functional properties that underpin energy and information technology, from light-matter interactions to temperature-dependent charge and spin transport. In fact, predictive first-principles calculations of these materials properties require sub-100 meV accuracy, and at this energy scale many-body electron-electron and electron-phonon effects cannot be ignored.

One of the grand challenges of *ab initio* materials design is to overcome these limitations by systematically incorporating electron correlations, quantum fluctuations, and realistic temperature conditions in accurate post-DFT calculations of real materials.

As an example, the prediction of the electrical conductivity of materials, a key design parameter for electronic, thermoelectric, optoelectronic, photovoltaic, and plasmonic devices, requires (i) high-accuracy many-body GW band structures to obtain reliable carrier effective masses, and (ii) a fine sampling of all carrier-phonon scattering processes to obtain reliable carrier relaxation rates. The resulting calculations are already challenging for binary semiconductors like gallium nitride, with only a handful of atoms per crystalline unit cell. For more complex systems, like the recently discovered hybrid organic-inorganic perovskites which count almost fifty atoms per simulation cell, these calculations are almost prohibitive on the best available supercomputing architectures. Typical HPC resources required for these calculations are as follows. The numbers refer to the cubic perovskite CsPbI3 with 5 atoms per simulation cell, on MareNostrum 4 from BSC (Intel Xeon Platinum 8160 CPU @2.10GHz, 48 cores per node):

- GW calculations for quasiparticle bands and effective masses: 48h wall-time on 3840 cores;
- EPW calculations of carrier mobility via solution of the *ab initio* Boltzmann transport equation for one temperature: 48h wall-time on 3840 cores.

These figures refer to a single temperature datapoint for one compound. We have been able to test Quantum ESPRESSO/EPW on up to 15K cores, corresponding to a theoretical peak performance of 1.2 petaflops (325 nodes), but most calculations are executed using fewer nodes in order to minimize job queuing time.

In a typical materials design project it is common to screen several tens of different compounds, doping levels, and elastic strain conditions. As a result, at this time a systematic high-throughput screening of advanced functional properties such as the electrical conductivity is beyond reach. Therefore, the scope and impact of materials design projects are largely shaped/constrained by the availability of computational resources.

In order to be competitive on the front of advanced materials design, and to deliver tangible impact in the form of new compounds, new functionalities, new devices, and related intellectual property rights, it would be useful if the design of LCCF at TACC could take the following considerations into account:

- If we could reduce the total compute time for calculations of the type above to a few wall-time hours, then we could perform systematic materials screening over timescales of the order of days instead of weeks/months. This speedup would give us a strong competitive advantage.
  This would require us to run routinely on 20-30,000 cores, which will only be possible if we will be able to improve the parallel scaling of our code EPW and if the LCCF will be large enough to schedule jobs of this size with a short queuing time.
- There is an increasing interest in using hybrid CPU/GPU architectures, and the Quantum ESPRESSO suite has an experimental GPU release. We have plans to experiment with GPUs with our flagship software EPW, but the general feeling is that the workhorse of most of our calculations will be large CPU clusters due to memory requirements. Indeed it would be useful to have high-memory nodes for memory-hungry tasks. The general philosophy in large open-source computational materials science projects is to prioritize code simplicity, transparency, and readability over economical memory management. As an example, we recently introduced some point-to-point communications in EPW to save on memory, but we are reverting back to global communications since our choice of using ScaLAPACK and distributed memory was not met with much enthusiasm as it adds unnecessary complexity.
- We will need to keep improving the parallel scaling of our codes. While my group was in the UK we had a positive experience with a EU consortium called POP Center of Excellence in Computing Applications. They analysed the scalability and load balancing of EPW using Score-P and Scalasca, and helped us improving the scaling tenfold. This kind of software engineering support would be highly desirable to continuously update the codes and make the most of new architectures.
- Perhaps slightly unrelated to the LCCF project, it would be extremely useful if TACC could provide facilities to run computational summer schools. For example we are trying to organize a summer school on materials design and one of the challenges is to find a location at the University of Texas at Austin that can host 70-80 Linux desktops in a classroom setting.

Lattice QCD for Flavor Physics

An important scientific challenge in high energy or elementary particle physics is to find evidence of Beyond the Standard Model Physics. This might be done by producing new particles with higher energy collisions than previously studied, but it can also be done using very high precision measurements that test precise predictions of the Standard Model. The latter approach is being used to study the physics of the bottom quark by the LHCb experiment at the Large Hadron Collider at CERN, and the Belle II experiment at the KEK Laboratory. High precision measurements are also planned at BES III in Beijing to study the lighter charm quark. Lattice QCD calculations provide essential input to these experiments as they are the most precise way to make the required Standard Model predictions. Currently, there are several experimental measurements that are a couple of standard deviations from the theoretical predictions. As the experiments are ongoing and will increase their precision, it is necessary to increase the precision of the predictions of lattice QCD.

Lattice QCD involves a numerical treatment of the continuum field theory in which space time is treated as a discrete four-dimensional grid with lattice spacing a. This introduces systematic errors that are reduced when a is reduced. The bottom quark is quite heavy, and there are errors that depend on am_b, where m_b is the bottom quark mass. A grand challenge problem is to reduce the lattice spacing to 0.03 fm, so that the errors from the bottom quark are well controlled. We are currently doing calculations with a lattice spacing of 0.06 fm. We would use quark masses that are well tuned to the physical quark masses for the up, down, strange, and charm quarks. This approach largely eliminates the error from the chiral extrapolation that is needed when quarks heavier than in Nature are used.

With $a = 0.06$ fm we are using a grid of size $96^3 \times 192$. Thus, when we halve the lattice spacing we would use a grid of size $192^3 \times 384$. This problem has a volume 16 times as large as what we are currently using in production jobs on Frontera. The computational cost would increase by more than a factor of 16 as the condition number of the sparse matrix we need to invert will increase requiring more iterations. We also expect to have to decrease the step size in the molecular dynamics-like evolution required to generate new gauge configurations. We are looking at algorithmic advances that might reduce computational requirements such as multigrid solvers, and force gradient integrators. As part of the Exascale Computing Project, we are about to test a multigrid solver for staggered quarks on Summit.

We have made a very rough estimate of the time required to create and analyze 1000 configurations for a variety of mesonic decays containing bottom and charm quarks. The analysis jobs are already in production on Frontera for the $a = 0.06$ fm ensemble, so we can estimate that time based on increasing the volume and an assumption about how many more iterations we will require with the smaller lattice spacing. Currently, we are using 16.5 hours on 128 nodes to analzye a configuration. We assume with the larger volume, we would use 16 times as many nodes, and that the time would be about 50% longer because of increased iterations. (This does not anticipate decreased performance with a larger node count.) So, each configuration would take 24 hours on 1920 nodes. To analyze 1000 configurations would take 2 million node-days on Frontera. We have done some test running for configuration generation of a $192^3 \times 384$ configuration on Frontera, but did not run long enough to equilibrate, so roughly know the performance of the solver, but don't know how many iterations will be required once equilibrated. A very, very rough estimate is 2.7 million node-days to generate 1000 configurations on Frontera.

Bridging short and long time scales in global plate tectonics: A grand challenge computational problem

Michael Gurnis (gurnis@gps.caltech.edu) and Georg Stadler (stadler@cims.nyu.edu)
Seismological Laboratory, Caltech, CA and Courant Institute, New York University, NY

Global plate motions are primarily driven by the negative buoyancy within subducted slabs, but there is little consensus on the strength and nature of resisting forces. Plate motions are likely related to how forces are transmitted through the slab, across the plate interface, and into the converging plate. Failure to reach consensus on the force balance of plate tectonics and mantle flow leaves fundamental questions unanswered, including causes for spatial variability in great earthquake occurrence.

Plate subduction causes great earthquakes, but the ambiguity surrounding their occurrence results in substantial uncertainty in the assessment of seismic hazards. Some subduction zones, like Chile and the Aleutians, experience great earthquakes, while others do not. For a long time, this was thought to be simply related to plate age and convergence rate, but this correlation failed to account for recent great earthquakes. We have shown that global dynamic models can reproduce the overall state-of-stress in subduction zones and thus a connection between tectonic forces and the occurrence of great earthquakes remains quite viable, although there is no widely accepted global model.

End member approaches have been used. Global flow models with 1 km resolution have allowed us to solve the steady-state, long time-scale motion of plates as a Stokes flow problem (Fig. 1). An emerging class of models use 2D cross-sections with visco-elastic constitutive relations and lead to earthquake-like ruptures and reveal intriguing earthquake-cycling phenomena, but are not self-consistent with plate motions. These approaches can be brought closer together in a new generation of computations that bridge the time- and spatial-scales while assimilating the vast geophysical constraints. We could achieve a much deeper understanding on the forces driving and resisting plate motions, the link between great earthquakes and plate motions, and putative connections between great earthquakes.

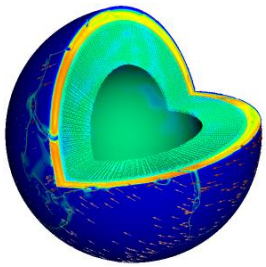

Figure 1: Forward model with $\sim$1 km resolution. Shown are effective viscosity and velocity vectors.

Making progress will require a new generation of time-dependent plate-mantle models in which the details of plate boundaries are resolved in time and space. This will require global forward and inverse, visco-elastic models to determine variations in mechanical coupling in subduction zones using plate motions, space-geodetic displacements, the record of large earthquakes, and the along strike variation in the depth of oceanic trenches. The models must have high resolution (e.g., locally 1 km or less) and fully incorporate the extreme non-linearity in viscosity between the slab and mantle, hinge zone and rupturing interface between subducting and over-riding plates.

This will require iterative solvers for nonlinear saddle point problems with $\mathcal{O}(10^9)$ unknowns arising from discretization of partial differential equations in viscous fluid dynamics. While being much more complex, the resulting sparse systems share similarities with the HPCG benchmark, i.e., they are memory-bound, involve collective communication (dot-products), and require multi-level preconditioners that utilize significant global communication. Since on future architectures, flops will likely be even faster compared to memory access, this is a problem where hardware may influence the development of numerical algorithms, e.g., towards higher-order discretizations or for a preference for computations on-the-fly instead of precomputing quantities. We will also use observations from plate motion, earthquakes and topography to deduce which mechanisms are able to explain our planet's evolution, and which are not. Such inverse problems require many repeated simulations, which is why strong parallel scaling of future algorithms and implementations is of particular importance.

High-Throughput Ensemble Calculations for Multiphysics and Turbulent Reacting Flows

Matthias Ihme
Department of Mechanical Engineering, Stanford University, Stanford CA 94305

Motivation: Reliable simulations of turbulent combustion can have a significant impact on improving existing combustion systems, on the acceleration of design-cycles, and on identifying novel combustion concepts for enabling the utilization of renewable fuel sources and the reductions of unwanted pollutant emissions. Specific applications include gas-turbines for aviation and power generation, multi-mode internal combustion engines, as well as chemical and process engineering and other propulsion systems for space access and defense. However, the reliable prediction of these combustion processes requires the consideration of complex physical processes, involving the elemental reactions among a large set of chemical species, localized heat release, turbulence, multi-phase flow, complex geometries, and heat-transfer by flame-wall coupling and radiation.

While existing engineering design-evaluations largely rely on low-order models, RANS-simulations, or empirical methods, scale-resolving simulations (such as large-eddy simulation or detailed numerical simulations) have shown to provide improved predictions of turbulent combustion, scalar transport by turbulent mixing, and pollution formation. However, because of the computational complexity, these high-fidelity simulations are largely performed in the form of "hero" calculations that push the state-of-the-art on spatial resolution, describing geometric complexity, or incorporating increasingly more complex physical processes. As such, these simulations are limited to few single-points calculations. By taking advantage of emerging exa-scale computing infrastructure, unprecedented opportunities arise to achieve a paradigm-shift by enabling *high-throughput ensemble calculations* in which multiple synchronous or asynchronous calculations are performed to enable engineering system optimization, sensitivity analysis, and geo-parametric design exploration.

Grant Challenge: A grant challenge that is facilitated through next-generation HPC exascale computing facility is *high-throughput ensemble calculations* of multiphysics and turbulent reacting flows. Enabling such calculation can impact the following aspects:

- ***Optimization and Parameter Exploration:*** Ensemble calculations employing high-fidelity and scale-resolving multiphysics modeling tools will support engineering design evaluation, integrated system optimization, and the rapid exploration of system operating envelopes.
- ***Extreme-Event Analysis:*** Apart from supporting engineering design-evaluations, high-throughput ensemble calculations would enable to push boundaries in our scientific ability to examine rare events that include extinction, engine-knock, onset of thermoacoustic instabilities or ignition by non-equilibrium energy deposition. Ensemble simulations would enable sampling low-frequency and high-impact events and obtain robust statistical information of light-tailed distributions and for importance sampling

- ***Data Assimilation:*** With relevance to the development of improved computational models, ensemble calculations can play an important role in data-assimilation, by combining incomplete and sparse experimental data with erroneous models for state-estimation, parameter estimation, and model evaluations. Since statistical assimilation methods rely on large ensembles to obtain robust estimators of prior-error covariance matrices, high-throughput ensemble calculations would open opportunities to expose these analysis techniques to turbulent multiphysics and combustion applications. With the advent of high-speed simultaneous and multidimensional measurements, opportunities arise in utilizing these ensemble techniques for state-estimation, accelerating numerical simulations, parameter estimations, and model evaluation.

Exploring Cohesive Sediment Dynamics in Turbulent Flows via Extreme Scale Computing

Eckart Meiburg
Department of Mechanical Engineering
University of California at Santa Barbara
Santa Barbara, CA 93106

Cohesive sediment is ubiquitous in ecologically sensitive environments such as rivers, lakes, estuaries, fisheries and benthic habitats. Reliable predictions of contaminant and nutrient transport in such settings require accurate models of cohesive sediment dynamics, which we currently lack. For cohesive sediment, which commonly refers to particles below approximately 63μm in size, interparticle cohesive and adhesive forces due to electric charges frequently dominate over the hydrodynamic and gravitational forces known to govern noncohesive sediment. These interparticle forces can trigger a process known as flocculation, which results in the formation of aggregates much larger than the individual grains. At the same time, turbulent stresses can act to break up these flocs generated by the cohesive forces, so that a delicate balance emerges between coalescence and break-up. Consequently, the dynamics of cohesive sediment in turbulent flows is significantly more complex than that of its noncohesive counterpart, with important implications for particle/floc size distributions and their effective settling rates. Cohesive forces furthermore strongly affect the erodibility of sediment deposits on the seafloor, which in turn influences fluvial and oceanic sediment transport processes.

We aim to explore the dynamics of cohesive sediment in turbulent environments via a series of increasingly complex computational investigations, from one-way coupled simulations to grain-resolving direct numerical simulations that are fully `four-way' coupled. These simulation results will serve as basis for formulating scaling laws that capture the dynamics of cohesive sediment in turbulence, which are suitable for implementation into existing larger-scale sediment transport models.

One-way coupled simulations are currently already possible, cf. fig. 1. We also have developed the computational code for fully four-way coupled direct numerical simulations based on an Immersed Boundary Method. However, carrying out such simulations for fully developed turbulent flows will push the limits of the largest computational resources, as can easily be seen from the following scaling argument: At the smallest scale, each primary particle will have to be resolved by about 20 grid cells. In order to study the flocculation process of these primary particles in turbulence, the Kolmogorov scale needs to be O(20) particle diameters. And to get a reasonable energy spectrum of the turbulent flow, the computational domain needs to extend over O(20) Kolmogorov scales. In this way, a full three-dimensional simulation will necessitate $O(10{,}000^3) = 10^{12}$ grid points, which renders the dynamics of cohesive sediment in turbulent flows a Scientific Grand Challenge that will require Extreme Scale Computing.

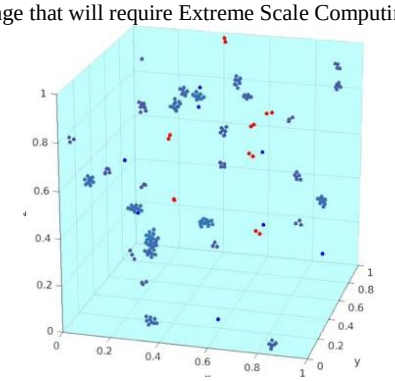

Fig. 1: One-way coupled simulations of cohesive sediment in turbulent flow.

A Computational Grand Challenge in Turbulence Simulation and Modeling

Robert Moser
Oden Institute, University of Texas

One of the most important challenges in fluid mechanics is to predict the effects of turbulence in a wide variety of flow situations. The inability to reliably predict the average behavior of many turbulent flows in systems such as aircraft, motor vehicles, ships, combustion systems, the atmosphere and ocean, and many more, greatly hinders the ability to design or otherwise make decisions regarding such systems. As an example of the importance of turbulence, consider that approximately 28% of total US energy consumption is for the transportation of goods and people,[†] and the vast majority of that energy is dissipated in the turbulent flows involved. The fact that practical models for the effects of turbulence in such flows are so unreliable can thus be seen as a major problem.

A likely reason for the unreliability of current turbulence models (called Reynolds-Averaged Navier-Stokes or RANS models) is that they do not carry a sufficiently rich state-space of turbulence statistics. The most common RANS models track the mean velocity, the turbulent kinetic energy and the rate of dissipation of the energy (or similar). However, it is known that the dynamics of turbulence with these same characteristics can vary widely depending on details of the structure of the turbulence. The question then arises as to what is a sufficient set of statistical descriptors of that structure that would allow the determination of the dynamics. A promising richer statistical state has been proposed, the structure tensors of [1], however, their adequacy as a state space describing turbulence dynamics has not been thoroughly evaluated. Further, there may in fact be other or different statistical descriptors that would be appropriate, but there has been little effort to discover them.

Arguably, a reason that such richer turbulence state spaces have not been explored is a lack of supporting data. The data that would be required includes values of potential statistical state quantities obtained from a wide variety of turbulent flows of varying complexity. The only viable source for such data is direct numerical simulations (DNS), which involve the solution of the Navier-Stokes equations on very large grids requiring extensive HPC resources. For example, a relatively recent DNS of a simple wall bounded flow involved a simulation with 250 billion degrees of freedom and required 400 million core hours on Mira at ALCF[2]. Many simulations of a wide range of more complex turbulent flows would be needed to determine an adequate turbulence state space and to inform a model for the evolution of such a state space. Many 10's of simulations may be needed for the basic case of turbulence in simple fluids, with many more needed when additional physics such as scalar transport, combustion, buoyancy, or multi-phase are included.

The computational grand challenge, then, is to perform such DNS and the required data-analysis. In addition to a first-order computational challenge requiring extreme scale HPC, this represents a first-order data science challenge. Among the data science objectives would be identifying the necessary state dependencies, determining additional flow conditions for which simulations are needed, and learning the high-dimensional functional relations required to model the evolution of the turbulence state. Such a focused integrated computational and data analytic effort has the potential to revolutionize the modeling of turbulence, resulting in vastly improved abilities to design turbulent flow systems or otherwise make decisions regarding them. Such an effort would be impossible absent extreme-scale HPC computational resources.

[1] Stavros C Kassinos, William C Reynolds, and Mike M Rogers. One-point turbulence structure tensors. *Journal of Fluid Mechanics*, 428:213–248, 2001.

[2] Myoungkyu Lee and Robert D. Moser. Direct numerical simulation of turbulent channel flow up to $Re_\tau \approx 5200$. *Journal of Fluid Mechanics*, 774:395–415, 7 2015. ISSN 1469-7645.

[†] http://needtoknow.nas.edu/energy/energy-use/transportation/

AIRCRAFT LES: A GRAND-CHALLENGE PROBLEM IN COMPUTATIONAL AERODYNAMICS

George Ilhwan Park[1]
[1]Department of Mechanical Engineering and Applied Mechanics, University of Pennsylvania, Philadelphia, PA

Kinetic energy of turbulence is generated at large scales controlled by boundary conditions, but it is destroyed at the smallest scales of motion known as the Kolmogorov scale. The ratio of these two length scales increases rapidly with Reynolds number, and turbulence at sufficiently high Reynolds numbers becomes a multiscale continuum phenomenon by itself. Solid walls present in nearly all engineering applications add another dimension in this scale landscape, as they limit the sizes of the near-wall energy-containing eddies to their distance to the wall. Consequently, the scale separation pronounced in the bulk outer region progressively depletes toward the wall. This has significant ramifications on the cost of computer simulation of external aerodynamics of wings and aircraft geometries. The facts that the chord Reynolds number is very high ($Re_c = 10\text{M} \sim 40\text{M}$), that the boundary layer (BL) grows rapidly with an $O(10^3)$ variation in its thickness, and that the computational mesh needs to cover wide wings (aspect ratio ~ 10) and a long fuselage render full resolution of relevant length scales close to the wall infeasible. Such direct approaches are deemed unreachable during the lifetimes of the current researchers. However, the national demand to enable superior aircraft designs at reduced cost and risk has led researchers to explore alternative computational approaches that are predictive yet affordable. Large-eddy simulation (LES) is one such approach where only the energy-containing scales are resolved directly and the scales unresolvable are modeled. The current state-of-the-art in aircraft LES is to deploy subgrid-scale (SGS) stress models in conjunction with wall models at (or close to) the wall to augment the turbulent shear stress, which otherwise is underpredicted on coarse grids. The LES grid now needs to resolve only the energetic eddies in the outer portion of BL, the sizes of which scale with the local BL thickness deemed resolvable. This approach includes wall-modeled LES (WMLES) and hybrid RANS/LES, and they are considered as the only viable ways of conducting aircraft LES with reasonable accuracy.

Several estimates of the cost of WMLES to date generally predict that the aircraft LES at flight Reynolds number can be done with $O(10^9 \sim 10^{10})$ grid points. NASA in their CFD Vision 2030 report (Slotnick *et al.* 2014) estimated, with $O(10^8)$ timesteps integration for 100 convective timescales and using operation counts from typical explicit-in-time flow solvers, that about 10^5 exaflops are required for an aircraft LES calculation. This was considered unattainable with the petascale machines available at that time, and therefore NASA considered LES of powered aircraft configuration across the full flight envelope a "grand-challenge problem", which may not be solved routinely by and beyond 2030, but, if achieved, would prove to be a game-changer in engineering design capability. However, with the scheduled introduction of first exascale machines in the next decade, we would be able to solve the same problem within just $O(1)$ day (or less!). This would be a major breakthrough in aerospace industry and defense technologies, where the idea of using LES has long been discouraged due to cost considerations, and instead low-fidelity methods which fail to capture critical flow phenomena (e.g. flow separation) have been used. Accurate prediction of mean-flow quantities (e.g., lift, drag, mean velocity) with unprecedentedly short turnaround times will fundamentally change the aircraft design cycle. Access to the time-resolved velocity and pressure fields in realistic flows will pave the new way for flow physics investigation which has been restricted to canonical geometries (pipe, channel, flat plate) at low Reynolds numbers for a long time.

At this moment, a technical barrier for researchers to prepare for the forthcoming exascale-computing era is the uncertainty in type of architecture that will be supported stably in the future. As specific composition of the next generation exascale computers is not yet clear, and because the programming paradigm is vastly different in Intel and NVIDIA based architectures, only teams with sufficient resource can invest early on code porting and optimization. Guidance from HPC leadership agencies and disclosure of information related to the types of new machines well in advance will be helpful to researchers. Investment on and open education for the domain-specific language (DSL), which helps scientists to develop mesh-based applications portable across different HPC computer architectures, will be helpful as well.

Grand Challenge problem

Binary black hole waveforms for LISA using numerical relativity

Hari Sundar, University of Utah

Gravitational waves are unique messengers that carry detailed information from deep within some of the most extreme astrophysical sources, virtually unchanged, across the cosmos. The field of gravitational wave astronomy began with the spectacular observations by the Laser Interferometer Gravitational Wave Observatory (LIGO) of a merging binary black hole (BBH) in 2015. Nonetheless, such ground-based detectors are limited by size and terrestrial noise sources. Thus one of the most exciting prospects for this new field is the Laser Interferometer Space Antenna (LISA). At over 500,000 times the size of LIGO, LISA will enable unprecedented observations of gravitational wave sources including black holes, neutron stars, white dwarfs and even early universe cosmology. Such measurements will be extensive and remarkably precise, indeed much more so than current gravitational wave observations. Their analysis will require equally precise theoretical models.

The full scientific potential of LISA can only be realized if the gravitational wave community can produce waveforms and templates that match the high quality of expected LISA data. Current codes in numerical relativity can not be scaled up to provide waveforms of sufficient quality, duration in time, and for the variety of sources (spin and eccentricity) expected for LISA.

There are several reasons for the lack of scalability of numerical relativity codes on modern HPC resources. We highlight these, as these will have to be addressed to ensure the success of the LISA project. The first challenge is the high-levels of adaptivity required for handling large mass ratios and for varied spin and eccentricities. This causes highly non-uniform load across the domain and makes it particularly challenging to partition and load balance. In particular, attempting to load balance uniformly based on overall work, results in different processes operating in different regions of the strong scaling curve i.e., the grain size is large on certain processes while being extremely small on others. Potentially, having a cluster with of nodes each having optimal efficiency at different grain sizes might be appropriate for this application, but clearly is not an efficient design for a general purpose HPC resource. A second challenge is the requirement for generating waveforms corresponding to hundreds of orbits. Such long running simulations will either need extremely fast machines that can reduce the runtime per timestep of simulation or will need us to develop parallel-in-time approaches. Finally, the overall complexity of the formulation makes it difficult to optimize, debug and verify numerical relativity codes. It will be essential to have tools to generate, debug and verify codes for future architectures to ensure that the resources can be effectively utilized before they become obsolete.

INCORPORATION OF DATA SCIENCE METHODS IN LARGE-SCALE HPC

TIMOTHY J. WILLIAMS

Introduction

The DOE-managed Leadership Computing Facility (LCF) has as of 2017 an expanded mission: support the three"pillars" of simulation, data, and learning for large-scale scientific computing.[1] In addition to traditional simulation based HPC, LCF is to support methodologies from data science such as machine learning (ML). Development projects such as those in the Argonne sites Aurora Early Science Program are working toward incorporating these methods into exascale computing. This presents a grand challenge in refining the methods for scientific discovery and scaling up their execution on current and future leadership-class HPC architectures.

Scope of Data Science Methods

Commercial/internet-focused Big Data work leading to successes in software and methods such as Spark and the deep learning (DL) revolution have led to an explosion of efforts to apply them to computational science. More traditional statistical and supervised ML approaches have a longer history in areas such as cosmology.[2] More recently, DL has been increasingly employed in a range of science application areas, for example in using tokamak data to train models to predict disruptions.[3] Other newly improved AI approaches such as natural language processing have been employed, for example, in automated searching of chemistry journal articles to guide organic dye selection for solar windows.[4] In situ/streaming data techniques such as intelligent compression and reduced precision[5] are also at play.

Data Sources and Derived Models

Experimental/observational as well as simulation data are sources for data-derived models. Some efforts such as neural tissue mapping from electron-microscopy[6] align well with mainstream ML successes in image classification. Others, such as deconstructing detector events[7] and tokamak profile data[3] demand different kinds of neural networks. Deriving predictive reduced models from simulation data, such as turbulence[8][9] and lattice QCD[10] requires quantitative validation against theory or more expensive "hero" simulations. Interpretability and quantifying uncertainty in the models are essential for scientific rigor.[11]

Next-Generation HPC System Requirements

To justify the use leadership-class computing, these new scientific workflows must run at scale (for Aurora, this means exascale) and meet threshold expectations of performance. They must make effective use of LCF hardware, which for at least the next generation will be accelerator-based. Hardware support for ML operations and associated reduced precision comes naturally from GPGPU accelerators (NVIDIA (Summit, Perlmutter), AMD (Frontier), Intel (Aurora)) as they support the commercial market. High-level, optimized single-GPU programming software such as Tensorflow and PyTorch is assumed by developers in this space. What exascale HPC needs beyond this is scalable software: distributed memory parallelism across thousands of nodes, as well as I/O and data ingestion/networking to accommodate large data volumes/rates. Also needed is support for complex workflows to couple data-driven methods with simulation.

Conclusion

Given the new prevalence ML/DL/AI in computational science, and growth of experimental/observational data production, the DOE LCF needs to incorporate data science methods are surely not unique for next-generation HPC systems. Accommodating data science methods in those systems hardware with software that is consistent across scales of computing and capitalizes optimally on mainstream high-level programming environments is a significant challenge. Meeting the challenge would impact a wide spectrum of computational science research.

References

[1] Barbara Helland. View from germantown. `https://science.osti.gov/-/media/ascr/ascac/pdf/meetings/201709/HellandASCAC170926.pdf`. Accessed: 2020-01-05.

[2] Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, Laurent Daudet, Maria Schuld, Naftali Tishby, Leslie Vogt-Maranto, and Lenka Zdeborová. Machine learning and the physical sciences. *Rev. Mod. Phys.*, 91:045002, Dec 2019.

[3] Julian Kates-Harbeck, Alexey Svyatkovskiy, and William Tang. Predicting disruptive instabilities in controlled fusion plasmas through deep learning. *Nature*, 568(7753):526–531, April 2019.

[4] Edward J. Beard, Ganesh Sivaraman, Ivaro Vzquez-Mayagoitia, Venkatram Vishwanath, and Jacqueline M. Cole. Comparative dataset of experimental and computational attributes of UV/vis absorption spectra. *Scientific Data*, 6(1):307, December 2019.

[5] Sheng Di and Franck Cappello. Fast error-bounded lossy hpc data compression with sz. *2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 730–739, 2016.

Leadership resolution for the turbulence problem

P.K Yeung (Georgia Tech), D.A. Donzis (Texas A&M), K.R. Sreenivasan (New York Univ)

Fluid motions in a turbulent state, with disorderly fluctuations in time and space, play a dominant role in a vast variety of applications including air-breathing propulsion, pollutant dispersion, wind power, extreme weather, oceanography, astrophysics, etc. The "turbulence problem", i.e., the task of understanding and predicting the nature and effects of these fluctuations, is a scientific endeavor of huge complexity, with a breadth of challenges mirroring the diversity of applications in which this phenomenon arises. In principle, direct numerical simulation based on exact equations for the conservation of mass, momentum and energy is the best numerical approach for enhancing physical understanding. Since the simulations must be capable of representing a wide range of scales at high accuracy, they are extremely resource-intensive, especially if flow conditions of high physical realism are desired.

Turbulent flows being very diverse, to maximize the science impact of our work, it is appropriate to put the highest priority on computing power to a simplified problem that has highest general importance, and then proceed with other more complex problems of fundamental interest. The focus on nonlinear energy transfer suggests that this primary configuration is incompressible isotropic turbulence at high Reynolds number, which has indeed been a focus of some of the most ambitious simulations worldwide, involving over 1 trillion grid points in spatial resolution. In particular, this flow configuration is ideal for delineating the nature of fine-scale intermittency, which occurs in all turbulent flows in diverse applications. However, recent analyses also indicate that even more computing power is needed than previously thought for resolving the emergence of extreme fluctuations, the mixing of heterogeneous substances, and the transport of pollutants.

We expect that, with the appropriate hardware and software, a next-generation leadership-class supercomputer will be capable of supporting high-impact advances into the turbulence problem, including high-Reynolds-number isotropic turbulence as noted above, and several flows of greater complexity. For example, if the flow speed approaches or exceeds the speed of sound, thermodynamic fluid property variations as well as acoustic and shock wave phenomena introduce much additional complexity. Conditions approaching hypersonic flight also require consideration of departures from the standard "perfect gas" behavior as well as chemical reactions in high-temperature gas dynamics. In other contexts, turbulence with density stratification due to vertical mean profiles of temperature and/or salinity is of great interest from environmental and ocean science perspectives. Salinity has a very low molecular diffusivity which makes both measurement and computation very difficult, and the nature of double-diffusive convection where temperature and salinity have opposing effects is still not well understood. Finally, magnetohydrodynamic turbulence where the velocity field is modified by, and possibly also modifies, the magnetic field — is a broad subject of its own, with relevance in engineering (nuclear reactors, metallurgy, etc) and other problems such as convection inside planets (including the Earth) and stars (including the Sun), the solar wind (responsible for space weather), etc; again, computation is a powerful tool of investigation.

We emphasize that turbulence is a grand challenge problem in science, engineering and computing, and one where advances in fundamental understanding can lead to diverse payoffs.

Challenges in Modeling Astrophysical Thermonuclear Explosions

Michael Zingale (Stony Brook)

The end states of stellar evolution result in exotic compact objects: white dwarfs, neutron stars, and black holes. When these remnants are in binary systems, interactions with the companion star can lead to a variety of different thermonuclear explosions: classical novae, X-ray bursts (XRBs), and thermonuclear (Type Ia) supernovae (SN Ia). SN Ia and novae are important sites of nucleosynthesis in the galaxy; XRBs allow us to probe the physics of dense nuclear matter; and SN Ia are important distance indicators for measuring the expansion history of the Universe. Simulations on these events have been ongoing for decades and allow us to understand the physics of the explosions, interpret the observations, and test different progenitor models for viability. Next generation supercomputers will allow us to push these simulations to unprecedented realism.

Modeling these events requires multiscale and multiphysics simulations. Spatial scales range from the size of the compact object down to the burning or dissipation scale. Temporal scales range from the millions of years (SN Ia) to hours (XRBs) of evolution that brought the star to the brink of explosion to the seconds-long explosion. Models must incorporate hydrodynamics, nuclear reactions, gravity, thermal conduction, magnetic fields, and radiation.

While a variety of approximations are used in astrophysics to model hydrodynamic flows, for stellar explosions, conservative finite-volume methods are the most popular. Adaptive mesh refinement (AMR) allows us to increase the range of length scales captured. Our group develops the AMReX Astrophysics Suite of open-source AMR multiphysics hydrodynamics codes[1]. Two codes in particular address stellar explosions: MAESTROeX is a low Mach number hydrodynamics code for modeling the highly subsonic stellar convection leading up to explosion and Castro is a fully compressible radiation hydrodynamics code that efficiently models the explosive phase. These codes share a common framework (AMR library, microphysics routines), are written in C++ and Fortran, and use MPI+OpenMP or MPI+CUDA.

Advances in our understanding of these events have been made both through the development of new algorithms (e.g. low-speed solvers) as well as new computer architectures. Today's supercomputers are dominated by GPU accelerators, as will the next generation of DOE supercomputers. Over the past year, we've ported the AMReX Astrophysics Suite to GPUs, enabling a large performance gain and opening new areas for scientific exploration. An important part of that process is performance portability. Our path allows us to use the exact same compute kernels on CPUs and GPUs, eliminating the burden of maintaining redundant code paths for different architectures.

To highlight the way that new computer architectures enable new science, we focus on our current work on XRBs. Our interest is in modeling the nucleosynthesis as a burning front propagates laterally through the accreted fuel layer on a neutron star. Due to the range of length and time scales, various approximations have been used in the literature. Our group started performing 2D hydrodynamical simulations of the flame propagating using the NERSC Cori and OLCF Titan (on CPUs) machines over a year ago. These calculations were the first to perform full hydrodynamics in both the lateral and vertical directions, allowing us to understand the complex dynamics of the flame front with realistic microphysics. The simulations were challenging, and we needed to make approximations to both the length scales (using artificially high neutron star rotation rates to compress the lateral scales) and time scales (boosting the flame speed). Taking advantage of GPUs, we are now able to eliminate these approximations in our current runs on the OLCF Summit machine. The same code is used, but running entirely on the GPUs, it is 5–10$\times$ faster than CPUs. We are on target to perform the first three-dimensional simulations this year as well as to explore the more interesting case of mixed H/He bursts, which involve more complicated nuclear reactions.

The present generation of supercomputers allow us to do simulations that were simply not possible with the previous generation. Future computers will give the same boost. With the next generation machines,

[1] `https://github.com/AMReX-Astro/`

we will be able to routinely do 3D simulations, cover a larger portion of the neutron star surface, increase the size of our reaction networks, and couple radiation.These will be the first global simulations of burning fronts on neutron stars that capture the realistic physics and allow us to connect directly to observations.

New architectures bring new programming challenges, and to date we've managed to keep pace with these developments. A key point is that funding needs to be provided to aid in preparing codes, not just for the science, especially for novel architectures. NSF should recognize this need. The software stack on the next generation machine is as important as the hardware. It should support the popular community-supported programming models today and have wide support for C/C++ and Fortran. To ensure the largest benefit to the community, application codes for these machines should be Open Source to the fullest extend allowed.