The Opportunities and Challenges of Exascale Computing

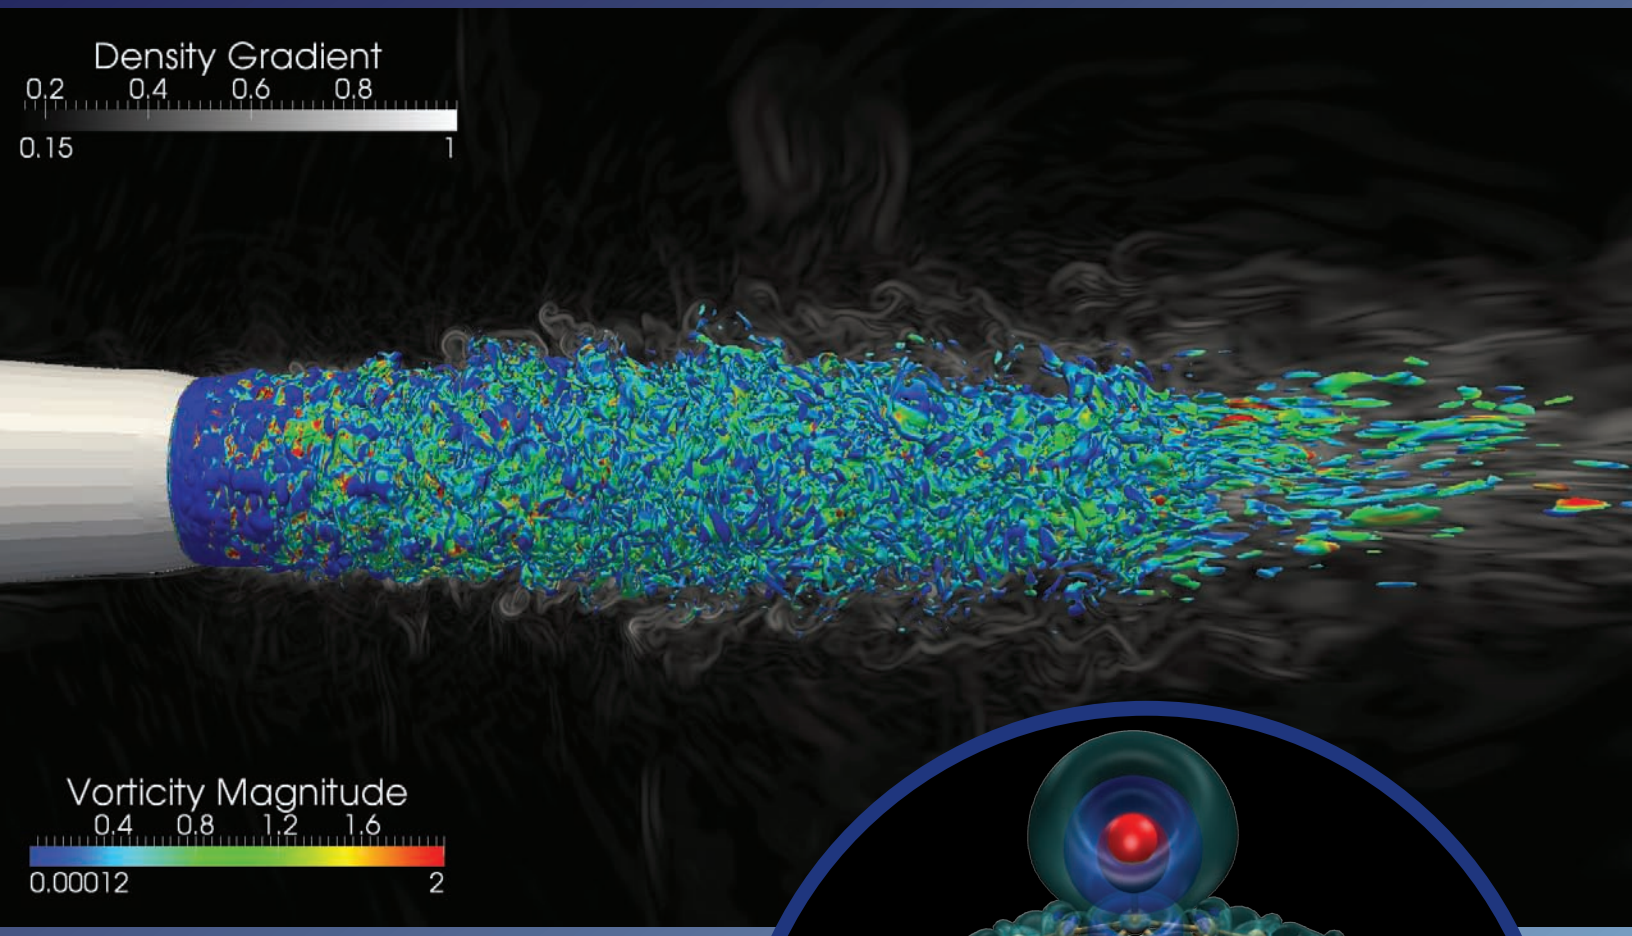

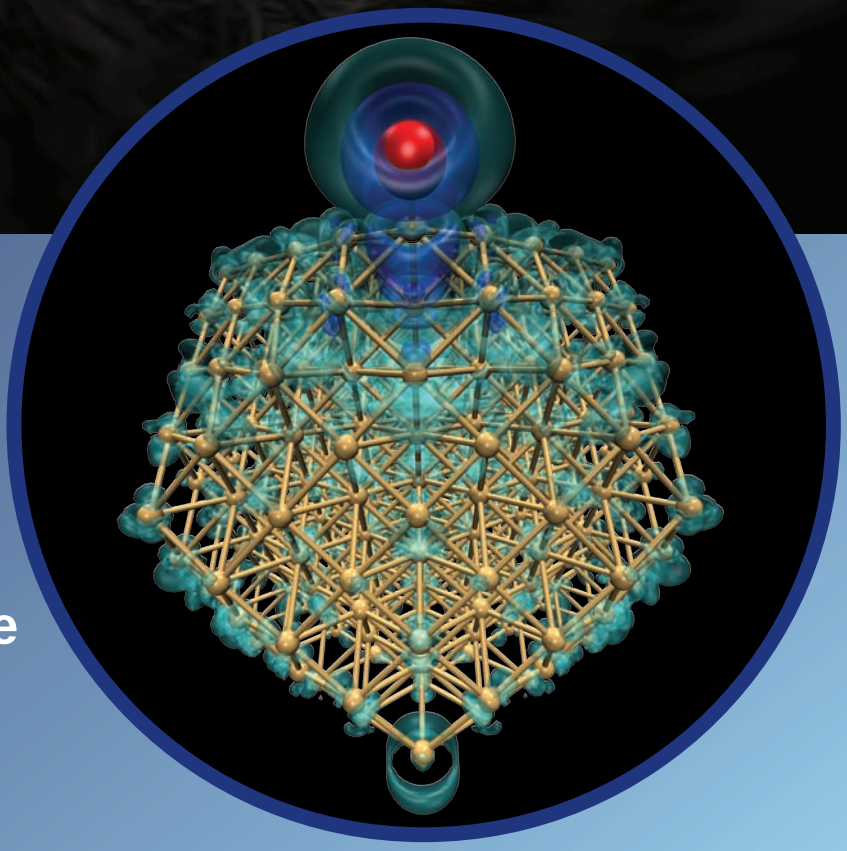

Summary Report of the Advanced Scientific Computing Advisory Committee (ASCAC) Subcommittee

Fall 2010

U.S. DEPARTMENT OF ENERGY | Office of Science

Report on Exascale Computing

The ASCAC Subcommittee on Exascale Computing[1]

Steve Ashby (PPNL), Pete Beckman (ANL), Jackie Chen (Sandia), Phil Colella (LBNL), Bill Collins (LBNL/UCB), Dona Crawford (LLNL), Jack Dongarra (Univ. of Tennessee), Doug Kothe (ORNL), Rusty Lusk (ANL), Paul Messina (ANL), Tony Mezzacappa (ORNL), Parviz Moin (Stanford Univ.), Mike Norman (UCSD), Robert Rosner (Univ. of Chicago/chair), Vivek Sarkar (Rice Univ.), Andrew Siegel (ANL), Fred Streitz (LLNL), Andy White (LANL), and Margaret Wright (NYU)

[1] The Subcommittee acknowledges the assistance of Rick Stevens and William Tang in the areas of biology and fusion energy science.

Executive Summary

The ASCAC Subcommittee on Exascale Computing has reviewed the exascale initiative as described in workshop reports and expert testimony. Computational modeling, simulation, prediction, and control at exascale offer the prospect of transformative progress in energy, national security, the environment, and our economy, and for fundamental scientific questions. Although the path to exascale necessarily involves numerous complex challenges, the almost-certain benefits far outweigh the costs.

Finding: ***The mission and science opportunities in going to exascale are compelling.***

For the growing number of problems where experiments are impossible, dangerous, or inordinately costly, extreme-scale computing will enable the solution of vastly more accurate predictive models and the analysis of massive quantities of data, producing quantum advances in areas of science and technology that are essential to DOE and Office of Science missions. For example, exascale computing will push the frontiers of

- Adaptation to regional climate changes such as sea level rise, drought and flooding, and severe weather patterns;
- Reduction of the carbon footprint of the transportation sector;
- Efficiency and safety of the nuclear energy sector;
- Innovative designs for cost-effective renewable energy resources such as batteries, catalysts, and biofuels;
- Certification of the U.S. nuclear stockpile, life extension programs, and directed stockpile work;
- Design of advanced experimental facilities, such as accelerators, and magnetic and inertial confinement fusion;
- First-principles understanding of the properties of fission and fusion reactions;
- Reverse engineering of the human brain;
- Design, control and manufacture of advanced materials.

U.S. economic competitiveness will also be significantly enhanced as companies utilize accelerate development of superior new products and spur creativity arising from modeling and simulation at unprecedented speed and fidelity. An important ingredient is the boost to continued international leadership of the U.S.-based information technology industry at all levels, from laptop to exaflop.

Finding: ***Making the transition to exascale poses numerous unavoidable scientific and technological challenges.***

Every major change in computer architecture has led to changes, usually dramatic and often unanticipated, and the move to exascale will be no exception. At the hardware level, feature size in silicon will almost certainly continue to decrease at Moore's Law pace at least over the next decade. To remain effective in high-end computing systems and in consumer electronics, computer chips must change in radical ways, such as containing 100 times more processing elements than today. Three challenges to be resolved are:

- *Reducing power requirements*. Based on current technology, scaling today's systems to an exaflop level would consume more than a gigawatt of power, roughly the output of Hoover Dam. Reducing the power requirement by a factor of at least 100 is a challenge for future hardware and software technologies.

- *Coping with run-time errors*. Today's systems have approximately 500,000 processing elements. By 2020, due to design and power constraints, the clock frequency is unlikely to change, which means that an exascale system will have approximately one billion processing elements. An immediate consequence is that the frequency of errors will increase (possibly by a factor of 1000) while timely identification and correction of errors become much more difficult.
- *Exploiting massive parallelism*. Mathematical models, numerical methods, and software implementations will all need new conceptual and programming paradigms to make effective use of unprecedented levels of concurrency.

Finding: ***The benefits of going to exascale far outweigh the costs.***

With the necessity of redoing computer applications comes the opportunity to include new predictive and analytical capabilities that will ideally underpin all mathematical models going forward. Two such capabilities deserve emphasis because they involve fundamental advances: (1) sophisticated uncertainty quantification can be embedded within each application, providing an overall statement of accuracy for computational forecasts and predictions; and (2) the application framework may be designed from the beginning as multi-scale – say, from molecules to planet-level. The consistent presence of both these crosscutting technologies will represent a major advance, long discussed but never before possible, in the very character of simulations.

Exascale will have a broad and positive impact on U.S. industrial competitiveness. As already noted, exascale technology breakthroughs will affect leadership from laptop to exaflop systems, because programmers at every scale will be faced with issues of performance and programming. As with high-performance computing in the past, high-tech industries such as transportation, aerospace, nuclear energy, and petroleum will rapidly acquire exascale applications and technology, especially those that allow accurate representation of multiple scales. Science breakthroughs at exascale may also lead to exponential growth in new industries such as renewable energy and materials by design.

Finding: ***The exascale initiative as described in workshop reports and expert testimony portends an integrated approach to the path forward.***

The high-level implementation strategies planned for exascale bring together application centers that will focus on development of new capabilities in a specific mission or science application area with community efforts in applied mathematics and computer science, spanning the gulf between the requirements of applications and systems. Simultaneously supporting technology developments common to all exascale systems will be pursued along with Laboratory-industry partnerships to accomplish the research, development, delivery, and operation necessary for exascale systems.

Finally, the initiative proposes to pursue the innovative strategy of co-design. In the past, machines have been purchased that have run current workloads in the best way possible and have represented the future of technology. However, over the next decade, as applications are redesigned and reengineered, the programming languages and tools will change radically and the systems themselves will be transformed. In many contexts, co-design of applications, software, and hardware will lead to much better use of the opportunities of exascale.

Recommendation: DoE should ***proceed expeditiously with an exascale initiative so that it continues to lead in using extreme scale computing to meet important national needs.***

It is the considered recommendation of the ASCAC Subcommittee on Exascale Computing that DOE proceed expeditiously to execute an exascale initiative. Exascale computing will uniquely provide knowledge leading to transformative advances for our economy, security and society in general. A failure to proceed with appropriate speed risks losing competitiveness in information technology, in our industrial base writ large, and in leading-edge science.

1. Introduction

This Report from the ASCAC Subcommittee on Exascale Computing is designed to cover the main issues raised by 'going to the exascale', and to provide some guidance on the level of risk involved in pursuing – and not pursuing – this direction of high performance computing.[2] 'Going to the exascale' will mean a radical change in computing architecture – basically, vastly increasing the levels of parallelism to the point of millions of processors working in tandem – which will force radical changes in how hardware is designed (at minimum, driven by economic limitations on power consumption), in how we go about solving problems (e.g., the application codes), and in how we marry application codes to the underlying hardware (e.g., the compilers, I/O, middleware, and related software tools). Understanding the advantages to be gained by going to the exascale, and evaluating the risks involved by going down this path, requires both an evaluation of past experiences in moving from the megaflop era to the present petaflop era, as well as an assessment of the readiness of advanced applications to take transformative advantage of exascale computing. These are the issues discussed in our Report[3], at a somewhat more general level than the highly detailed discipline-oriented Grand Challenges Workshop reports sponsored the Department of Energy (see Appendix 2), upon which much of our discussions are based.

"Going to the exascale" is a challenging venture – as will be described in this Report in some detail – but as we also explain, this step is an essential component in maintaining the United States as the world-wide high technology leader. The challenges inherent in developing exascale computing as a practical endeavor are considerable, and significant investments will be needed to accomplish this. As a consequence, it behooves us to make sure that the return on these investments is commensurate with the scale of investment. These considerations – the need to expose both the risks *and* the benefits of proceeding to the exascale – heavily influenced the structure and contents of this Report. Thus, the heart of this Report contains four sections that address the central questions raised when considering exascale computing:

- *Why exascale computing?* Chapter 2 deals with the central tenet of exascale computing, namely its promise to transform the role of modeling and simulation in science and engineering, turning computing into a true science-based predictive discipline. We discuss the role that exascale computing can play in elucidating how complex systems work – and in the case of manmade complex systems, how their performance can be optimized. The impact of exascale computing in transitioning modeling and simulations into a robust science-based predictive science is a key element in transforming science and engineering disciplines, and in transforming the relationship between modeling and simulations and industrial/manufacturing practices.
- *What can history teach us?* Chapter 3 focuses on an examination of past impacts of modeling and simulation in selected scientific and engineering disciplines. As modeling and simulation evolved from the megaflop era of the late 1970s to the present era of petaflop computing[4], the expectation was that advances in the application domains would justify the concomitant investments in hardware and software; here we discuss these advances for a selected number of application disciplines, focusing in particular on disciplines that have served as barometers of the health of high technology in the United States.

---

[2] The formal charge to our subcommittee is provided in Appendix 1.

[3] This Report has been authored by members of the Subcommittee; the exceptions are the sections on biology (provided by R. Stevens/ANL) and fusion (provided by W.M. Tang/PPPL); in both of these cases, the subcommittee did not have the expertise to authoritatively author the sections on these topics.

[4] 1 megaflop = 10^6 floating point operations per second; 1 petaflop = 10^{12} floating point operations per second. The deskside Digital Equipment Corp. *MicroVAX* was a megaflop computer; the Cray Corp. *Jaguar* massively-parallel computer at ORNL is a petaflop-scale machine.

- *What important applications will be transformed?* Chapter 4 examines the potential impact that exascale computing will have on a selected set of key science and engineering disciplines. Rather than being exhaustive, we focus on a limited number of areas for which we can establish in what ways exascale computing will transform these application disciplines, again targeting disciplines whose robustness has traditionally been seen as hallmark of US prominence in high technology industries. It is the transformations of these disciplines – driving science-based predictive modeling and simulation as a robust research paradigm for science and engineering – that provide the ultimate return on the investments that will be needed to get to the exascale.
- *What are the computational/technical challenges that must be overcome?* Chapter 5 considers in some detail the variety of challenges that will be faced in making exascale computing a practical reality. These challenges arise both in the hardware realm and in the software, and will call for potentially revolutionary changes in the ways we build and use high performance computers. These challenges are sufficiently complex, and inter-related, that a new methodological approach is called for in how the various research disciplines – from computer engineering to applied mathematics and computer science, and ultimately to the applications – interact as they pursue their own research agendas: We will argue that we will need to develop and deploy a 'co-design' methodology, in which the design of hardware, algorithms, programming models, and software tools is carried out in a coupled and iterative fashion.

2. Why Exascale?

The most obvious question – the key question really – is of course: why go to the exascale? This question is not meant in the trivial sense that one would pose for any expenditure whatsoever in leading edge computing technologies, but rather is motivated by the fact that the transition from current petascale computing to the exascale will involve investments across the board – from hardware to fundamental algorithms, programming models, compilers, and application codes – that will dwarf previous levels of investment made as computer architectures have evolved in the past. That is, we recognize that the values to society extracted from this change in computing paradigm has to be commensurate with the costs of developing this type of computing – and given the substantial costs, we need to be sure that the extracted values are similarly substantial. We will make the argument in the following that the extracted values are in fact very large – but will do so in two stages, first by making some general points about the present frontiers of computing – independent of discipline – and then by focusing on a few example disciplines to illustrate the more general point.

2.1 Transformation, not Evolution

As a first step, it is useful to be explicit about the kinds of arguments that will – and will not – justify large investments in next-generation computing. In general, one can divide forefront computing problems into three general categories:

- *Incrementally advanced computing.* In traditional scientific computing, many computational problems are characterized by a benign connection between advances in computing capabilities and the benefits gained by such advances: Colloquially speaking, such problems have the characteristic that computing a bit better gains results that are a bit better. Sometimes, such progress can be important – for example, it enables more robust predictability of computations – but rarely if ever transforms a discipline.
- *Voracious computing.* In some problem areas of computational sciences, there is no a priori reason to believe that increased computing capabilities will lead to increased knowledge, even in an incremental way. A classic example is homogeneous isotropic incompressible hydrodynamic turbulence: having resolved the viscous subrange, there is currently no reason to believe that larger computing domains (meaning, higher spatial and temporal resolution) will gain us a better

understanding of fundamental turbulence physics. Thus, such problems are voracious, in the sense that one cannot give an a priori reason for limiting the needed computing resources at any fixed level.

- *Transformational computing.* In some cases, computational problems have the property that they can be fully solved with sufficiently large computations, where 'sufficiently large' means that one can specify quantitatively in advance what 'large enough' means. Turbulence in a finite domain (such as in a tea cup, or in a jet engine combustion chamber) is a problem of this sort: A 'sufficiently large' computation will be capable of fully solving the problem. Such computing is transformational in the sense that a smaller calculation does not solve the underlying problem, but a sufficiently large calculation can solve the problem completely.[5]

The essence of this argument is that the level of expenditures going to the exascale is such that it can only be justified if important disciplinary areas can be identified in which exascale computing will be transformational, in the sense just discussed. As we will show, exascale computing will enable such transformations, transformations that touch many disciplines.

2.2 The Complexity Challenge

The great frontier of computational physics and engineering is in the challenge posed by high-fidelity simulations of real-world systems, that is, in truly transforming computational science into a fully predictive science. Real-world systems are typically characterized by multiple, interacting physical processes ('multi-physics'), interactions that occur on a wide range of both temporal and spatial scales. At the petascale, we are just beginning to carry out such fully coupled calculations – for example, coupling radiation and matter flows – but these tend to be limited in their spatial and temporal dynamic range. This is particularly true for 'direct numerical simulations' (i.e., simulations in which no physical processes are parameterized).

Not only would we expect the complexity – and thus the physical fidelity – of exascale computing to increase substantially beyond what is possible today, but it will also be possible to affect a transition to 'predictive' science and engineering calculations, that is, to apply modern methods of uncertainty quantification (UQ) in order to provide quantitative bounds on the applicability of exascale calculations. Such methods have been in development for the past decade, but relatively rarely applied because the required computational resources can substantially exceed that of the calculations subjected to uncertainty quantification. (Thus, for computations that are at the 'bleeding edge' of what is possible to compute today, reliable UQ becomes a virtually impossible task.)

Finally, one can ask whether there exist specific examples in which the kinds of transformations discussed above have taken place as we moved in the past from the megaflop era of the early 1980s to the current petaflop era. We will not go into any details in the present section (greater detail will be found below in Sections 3 and 4, as well as in far greater detail in the DOE/OASCR discipline-oriented workshop reports referred to in Appendix 2), but as part of the following discipline-oriented discussion, will briefly mention some examples in which such transformations have taken place.

2.2.1 *Aerospace, Airframes and Jet Turbines*

Firms such as Boeing, GE and Pratt & Whitney have a long history of following the computational trends outlined above, linking high fidelity simulations with reduced order computational engineering design tools. As a result, we see today tools such as 'numerical wind tunnels' – and reduced use of most low Mach number wind tunnels aimed at full or half-scale model exploration. Wind tunnel experiments, especially at large Mach numbers, are still carried out for the purpose of code and model development and

[5] It is important to note here that there are numerous examples of transformational advances in computing that were not the result of massive increases in computing speed or capacity, but rather emerged as a result of, for example, new algorithmic insights. Because the target of our report is on the possible benefits of going to the exascale, we do not focus on such advances here.

validation, but the industry has already seen cost savings from the reduced reliance on such experiments. When combined with modern UQ methodology, one can expect further cost savings (as experiments are expensive, but will be in increasingly lower demand) and improvements in design cycle time (as simulations can be carried out far more expeditiously than experiments). There is excellent historical data on progress of this kind over the past two decades, as (for example) airframe firms such as Boeing have evolved both high fidelity and reduced order engineering codes as the leading edge hardware moved from serial megaflop machines to today's massively parallel petaflop systems. (For further details, see immediately below in our discussion of high performance computing in industry.)

2.2.2 *Astrophysics*

Transformational calculations are in the offing in at least three distinct areas of computational astrophysics: coupled N-body/hydrodynamics/radiation transport codes for structure formation in the early universe; modeling of both Type Ia (thermonuclear) and Type II (core collapse) supernovae; magnetic dynamo theory. Based on what we know can be accomplished on petascale machines, we can with considerable confidence predict that in all three areas, it will be possible to carry out predictive calculations that connect directly with observational data; and put us in a position to definitively confirm (or invalidate) extant theoretical models. Indeed, in the area of N-body calculations over the past decade, we have directly witnessed the transformational power of such simulations to exclude entire categories of cosmological models, based on direct comparisons between N-body simulations and large-scale astronomical surveys of distant galaxies and quasars.

2.2.3 *Biological and Medical Systems*.

In biology, the challenges of modeling at multiple scales—from atomic, through genomic and cellular, to ecosystems—are already pushing beyond the petascale class of computing systems coming online. For example, a computational approach to understanding cellular systems and their related genes and biochemical pathways, referred to as *systems biology*, aims to develop validated capabilities for simulating cells as spatially extended mechanical and chemical systems in a way that accurately represents processes such as cell growth, metabolism, locomotion, and sensing. Modeling and simulation provide only a local view of each process, without interactions between modalities and scales. Exascale computing and new simulation management tools are needed to represent the various macroscopic subsystems and to enable a multiscale approach to biological modeling.

2.2.4 *Climate and Weather*.

Models that synthesize our observations and theories of the Earth system as accurately as possible are central to research on climate change. Over the last thirty years, these models have advanced considerably in the representation of key processes and in spatial and temporal resolution. Thanks to recent advances in computing and climate science, the scientific community is now developing Earth System Models (ESMs) that could simulate the co-evolution of the physical, chemical, and biogeochemical states of the climate system. These models will be capable of predicting how anthropogenic pollutants and land-surface alterations interact with natural chemical and ecological processes. ESMs will include the primary feedbacks among physical climate change, the chemistry of the oxidant cycle, and the biogeochemistry of the carbon cycle. One of the main scientific challenges in this activity is the integration of the terrestrial and ocean ecosystems with the rest of the ESM.

The spatial resolution of both climate and Numerical Weather Prediction (NWP) models have advanced by factors of 4 to 5 between the first and fourth Assessment Reports of the Intergovernmental Panel on Climate Change (IPCC) issued in 1990 and 2007. Several groups are running ocean models at "eddy-resolving" scales of 0.1 degrees, sufficient to resolve the interactions of upper ocean waters with their neighboring continents and with atmospheric weather systems such as hurricanes. Fully coupled models are capable of simulating the climate at scales of 25 km – a resolution comparable to the size of an average U.S. county.

2.2.5 *Combustion*

Closely coupled experiments and simulations of combustion processes have already had a major impact on the efficiency of extant combustion motors, from reciprocating engines to turbines. As in the case of airframes and jet turbines, we can expect reduced reliance on expensive experimental test benches, and increased reliance on fully coupled combustion hydrodynamic and structural mechanics codes, both for high fidelity direct numerical simulations and for reduced order engineering design calculations. Again, one can extrapolate from recent progress to expect substantial time and cost savings as we move to the exascale.

2.2.6 *Materials Science*

One of the great challenges of materials science is to bridge the gaps between the atomistic, molecular, meso, and macroscopic levels of descriptions. This is one of the classic 'grand challenges' of computational science, and progress has been very slow. However, recent petascale-level calculations of phenomena as varied as atomistic descriptions of Kelvin-Helmholtz instability of counter-flowing liquid metals and solidification of heterogeneous melts are beginning to show explicitly how definitive results can be obtained for sufficiently large simulations – meaning, one can now explicitly predict the size of calculations needed to definitively address (for example) the spectrum of domain sizes expected as a result of specific melt solidification processes. Thus, such calculations need no longer be regarded as inherently 'voracious'.

2.2.7 *Fusion Energy*

Worldwide energy consumption has risen 20-fold during the 20th century, and this growth rate shows no sign of decreasing at present. Of the current 15 TW load, 80-90% is derived from fossil fuels, but with peak oil imminent and climate change accelerating, a sea change in energy production mechanisms is vital. Fusion is the power source of the sun and other stars, which occurs when certain isotopes of the lightest atom, hydrogen, combine to make helium in a very hot (100 million degrees centigrade) plasma. In addition to being an attractive, long-term form of nuclear energy, it can have a major impact on climate change challenges such as the reduction of carbon dioxide emission by an order of magnitude. Major progress achieved to date [10 million Watts (MW) of power sustained for 1 second with a gain of order unity] has been dramatic and has led to the ITER Project – an international $20B class burning plasma experiment in Cadarache, France that is supported by seven nations (including the US) that represent over half of the world's population. ITER is designed to produce 500 million Watts of heat from fusion reactions for over 400 seconds with gain exceeding 10 – thereby demonstrating the scientific and technical feasibility of magnetic fusion energy. While many of the technologies used in ITER will be the same as those required in an actual demonstration power plant (DEMO), further advances in science and technology are needed to achieve the 2500 MW of continuous power with a gain of 25 in a device of similar size and field. Accordingly, strong R&D programs are needed to harvest the scientific knowledge from ITER and leverage its results. Advanced computations in tandem with experiment and theory are essential in this mission. The associated research demands the accelerated development of computational tools and techniques that aid the acquisition of the scientific understanding needed to develop predictive models which can prove superior to extrapolations of experimental results. This is made possible by access to leadership class computing resources that allow simulations of increasingly complex phenomena with greater physics fidelity. Exascale computations of key plasma physics and materials science processes in fusion devices will help to design, run, and interpret expensive large-scale experiments like ITER and DEMO, and are expected to significantly accelerate progress toward the realization of fusion power plants.

2.2.8 *National Security*

In the realm of national security, the development and design of nuclear weapons has been centrally involved with forefront numerical simulation virtually from their beginnings. Some areas of

computational science – such as particle transport – owe much of their intellectual basis to the early work on neutron transport (or "neutronics"); and many modern computational algorithms, especially in the domain of compressible fluid dynamics, were developed in the context of weapons research. Within the past 15 years or so, since the advent of the Comprehensive Test Ban Treaty, modeling and simulations have become yet more central to weapons work: for example, simulations have played an increasingly important role in certifying the performance and reliability of the U.S. nuclear stockpile. This represented a radical transformation of the role of simulations within the weapons program, and naturally coupled to a concurrent transformation of computing itself, as massively parallel computer architectures became the dominant computing paradigm of today. Given that actual testing of nuclear weapons is increasingly unlikely as Russia and the U.S. work towards the asymptotic goal of 'going to zero', assuring the reliability of the remaining (smaller) nuclear stockpile will become increasingly important – and hence modeling and simulations will play yet more critical roles in the certification program, as well as in the associated life extension programs. Thus, the requirements for science-based predictive capabilities of the weapons codes will become much more stringent, driving advances in multiphysics simulation capabilities and placing increasing demands on the ability to carry out calculations on broad spatial and temporal scales.

2.2.9 *Nuclear Engineering*

Commercial nuclear power plants (NPPs) generate over 20% of the electricity currently produced in the United States. Historically, the design of these plants reflected the fact that modeling and simulation (M&S) has always been an integral part of nuclear engineering analysis, safety, and design. Computational analyses have been used to predict quantities such as the aging of structures, power distributions in cores, transient safety behavior, etc.

From a modeling and simulation (M&S) perspective, nuclear energy does not exactly have the same drivers as other programs, e.g., national security, climate change, or the aircraft industry, to name a few. Unless M&S activities are driven by requirements set by industry and regulatory agencies, approaches in which better can become the enemy of "good enough" could prevail and defocus efforts. Verification and validation (V&V) activities will be crucial. The stakeholders are: (1) the nuclear power industry, which is regulated, capital intensive, endures lots of public security, is risk adverse, and only uses common off-the-shelf (COTS) technologies; (2) nuclear technology vendors, which exist to provide services, are dominated by proprietary data and tools, and have not bought into HPC; (3) the Nuclear Regulatory Commission (NRC), which is conservative and slow-moving; (4) the DOE, which is now focusing on nuclear energy more so than in the recent past, and (5) professional societies.

The existing (legacy) tools for application to material property determination, spent fuel reprocessing, fuel performance, reactor safety and design, and waste forms and storage are based on a large experimental database and hence will be insufficient moving forward. Experimental testing will be extremely expensive, protracted, and in some cases unfeasible. Furthermore, the existing experimental database is insufficiently documented and often has inadequate precision to support a modern validation process. Complementing or replacing testing with high-fidelity computer simulation will make it possible to collect simulated data that can, in conjunction with a sound experimental validation program, be used to understand fundamental processes that affect facility efficiency, safety, and cost.

2.3 Connecting to Industry

Here the obvious questions are: What does industry need – and what does industry want? The particular area we focus on is the use of predictive computing as an industrial design tool – and it is important to realize that, by and large, industry does not generally rely upon forefront computing technologies as part of its design processes, whether one considers product development or the development of processes (e.g., modeling and simulation of assembly lines, fabrication facilities, etc.). Instead, industrial design relies heavily upon reduced order descriptions (meaning, invariably highly parameterized engineering

software), involving codes that run on modest-scale computing systems (at times, laptops or workstations), but run on them continuously as a design phase space is mapped out and explored for greater efficiencies and lowered costs (and, in some case, greater performance and safety properties). Of course, these parameterizations often rely on a combination of experiments and high fidelity direct (e.g., unparametrized) numerical simulations, but these are used sparingly: the first because the required experimentation is very expensive, and the second because current direct simulation tools are of limited utility and are also relatively expensive to develop (and to deploy on the requisite computing hardware). As a result, current computational design tools can be regarded only marginally as truly predictive, since their domain of robust predictability is largely constrained by the validity of the parameterizations inherent in the reduced order descriptions. This means that design choices that depart significantly from existing designs generally cannot be easily explored using existing design software.

It is however important to note that in some industries – the airframe and jet turbine industries, as well as the heavy equipment manufacturers, perhaps most notably – there has been a long history of code evolution: As leading edge high fidelity computations have become more and more capable, their successful deployment as calibration tools for reduced order simulation codes has led to a corresponding increase in capability of these reduced order engineering codes. This lock-step, or bootstrap, process linking high fidelity and reduced order engineering codes can be expected to continue as we reach towards the exascale – so that as high fidelity codes move to exascale platforms, the engineering codes can be expected to migrate to giga, tera, and petascale platforms.

Industry is also faced with new government mandates, for example to reduce petroleum use by 25 percent by 2020 and greenhouse gas emissions by 80 percent by 2050, requiring significant retooling of all aspects of our nation's energy use. Achieving such aggressive usage and emission goals will require, for example, the automobile and truck sector to significantly shorten its product development cycle for cleaner, more efficient engine technologies. The transportation sector alone accounts for two-thirds of the nation's oil use and one-quarter of the nation's greenhouse gas emissions. Compounding these challenges, fuels will also be evolving, adding another layer of complexity and further underscoring the need for efficient product development cycles. However, under current Edisonian cut-and-try methodology, design cycles and market penetration of new engine technologies simply takes too long. These challenges present a unique opportunity to marshal U.S. leadership in supercomputing to develop predictive science and engineering simulation tools with the required fidelity working in close connection with key U.S. industries.

2.4 Why now?

As in any case involving the potential for large expenditures, it is incumbent upon us to consider the question: 'why now?' That is, does there exist urgency for proceeding to the exascale? What are the risks involved if we do not proceed? In fact, the importance for moving to exascale is hard to overstate. A leadership role for the US in exascale computing will be accompanied by the collateral benefits of leadership and innovation in multiple industries, including aerospace, automotive, energy, health care, information technology, and manufacturing, as well as the job creation that accompanies such leadership. Since the aerospace industry has long been an early adopter of high performance computing (HPC) for simulating airframes and jet turbines (Section 3.1 below), exascale capability in the U.S. will enable U.S. aerospace companies (such as Boeing) to gain a competitive advantage over foreign aerospace companies. In contrast, the automotive industry has been a more recent adopter of HPC; but given the competitive challenges for fuel efficiency, access to exascale simulations of internal combustion engines (Section 3.4) could be a key differentiator for US automotive companies. These companies are poised to use advanced computational fluid dynamics to significantly shorten their design cycle, in order to meet the federal mandated 80% reduction of greenhouse gas emissions by 2050 and the 25% reduction in petroleum usage by 2020. Access to exascale simulations of combustion will of course have an even bigger impact on the energy industry, since over 85% of the world's energy is generated by burning fossil fuel; further,

exascale capability will also help increase the fraction of power generated by carbon-free nuclear power plants (Section 3.6). In health care, exascale computing can help not only in advancing the scientific research necessary for preventing and curing cancer and other challenging diseases, but also in lowering the cost and increasing the effectiveness of health care services through, for example, the use of techniques such as compressive sensing to obtain lower radiation dosages[6]. Exascale computing can enable the development of novel materials in the U.S., along with infrastructure for high-value manufacturing processes involving these materials as was done in the past for materials such as semiconductors and polymers (Section 4.6). For the information technology industry, leadership in hardware and software innovation for exascale computing will be accompanied by leadership in the full spectrum of computing from embedded computers and handheld mobile devices to laptop and desktop servers to departmental servers, all of which includes the use of high levels of intra-chip multithreaded parallelism. There is a critical need to find co-designed hardware and software solutions in this full spectrum (e.g., the power challenges for mobile devices bear a lot of similarity to those for nodes in an exascale systems) and whichever country leads in exascale computing will also enjoy technology leadership in the IT sector in general. Finally, leadership in exascale computing is also critical to national security, both because of the need for simulation as a replacement for testing of nuclear weapons and because of nuclear test ban treaties and because of needs in the intelligence communities for mining large data sets (Sections 2.2.8 and 4.8).

If the U.S. chooses to be a follower rather a leader in exascale computing, we must be willing to cede leadership in all the industries listed above. In addition, since scientific advances will be gated by access to exascale computing, a non-trivial fraction of our research budget for science (including astrophysics, climate modeling, fusion energy, computational biology, materials science, nuclear science) will be spent on gaining access to leadership exascale systems, thereby using our funds to subsidize the R&D costs of whichever country is first to reach exascale computing (currently, the closest contender is China[7]). This approach will of course not be feasible for any classified codes. In addition, there is the danger of compromising our intellectual property in unclassified codes if they are to be run on foreign computer systems.

3. What Can History Teach Us?

As discussed in Section 2, assessing the wisdom of proceeding to the exascale must of necessity be informed by our historical experience in exploiting dramatic shifts in available computing power, shifts that occurred as we moved from megaflops to petaflops over the past 40 years. This section is devoted to an examination of some of these past experiences. We do not aim to be comprehensive: the series of workshops devoted to exascale computing in the broad range of disciplines in which modeling and simulation matters (see Appendix 2) provides the full picture of what has been accomplished to date[8]. Here we aim for a more concise description, with the intent of illustrating the breadth of accomplishments achieved in transitioning from the megaflop era of the 1970s to the present petaflop era.

3.1 Aerospace, Airframes and Jet Turbines

The aerospace industry was one of the early pioneers to apply computational fluid dynamics (CFD) to the design process. The primary motivation has always been to reduce expensive physical testing (e.g., wind

[6] See, e.g., Mao, D. 2009. "Equally-Sloped Tomographic Reconstruction and Its Application to Radiation Dose Reduction." In *Compressive-Sensing Workshop*, L Carin & G Arnold (Duke Univ., Feb 25-26, 2009).

[7] See, e.g., Markoff, J. 2010. "Chinese Supercomputer Is Ranked World's Second-Fastest, Challenging U.S. Dominance", New York Times, May 31, 2010.

[8] For example, reaching petascale-level computational capabilities has had major impact on both QCD and nuclear structure theory, the details of which can be found in the respective disciplinary Exascale Workshop reports.

tunnel tests) and the number of prototypes built during the design cycle. In the early seventies, computations (using linearized inviscid approximations of flow equations) of transonic wings convincingly demonstrated the resulting cost savings of CFD as well as the potential limitations imposed by ground-test facilities, such as wind tunnel walls. Successively refined approximations to the fundamental governing equations were applied to simulations of complex three-dimensional aircraft beginning with non-linear inviscid approximations and culminating with Reynolds-Averaged Navier-Stokes (RANS) methods in the 1990's. Each successive stage of approximation allowed a simulation to capture a new class of physics. Effective use of each successive stage of refinement, however, became practical only when the requisite computer power for that stage became available[9].

As can be seen from Figure 1 below, the predictive capability of the mainstream CFD in the design cycle seems to have reached a plateau. Owing to the use of CFD, the number of wind tunnel tests in the development of Boeing jet liners reduced dramatically from 70-80 tests in the 1980's (B757 and B767) to 11 in the 1990's (B777)[10]. This trend has not continued with the development of the newest Boeing model (B787) where the required number of tests has remained at 11 (earlier Boeing projection for this figure was 5). This is in spite of the fact that the available computer power in the same period has increased by nearly four orders of magnitude. That is, although the design cycle time (a highly coveted factor in aerospace design) has decreased significantly, the computations have not achieved the state of predictability and accuracy required to reduce the required number of wind tunnel tests.

The RANS approach remains the industry standard in CFD for aerospace applications. In this method, the engineers solve for the mean flow quantities, which are then used to obtain quantities of engineering interest such as lift and drag forces. Although the governing flow equations, the Navier-Stokes equations, are derived from first principles, the averaging process introduces unclosed terms to account for flow turbulence, which are modeled phenomenologically, and not based on first principles. There is a long and rich history in turbulence research aimed at calibrating RANS models for engineering applications using

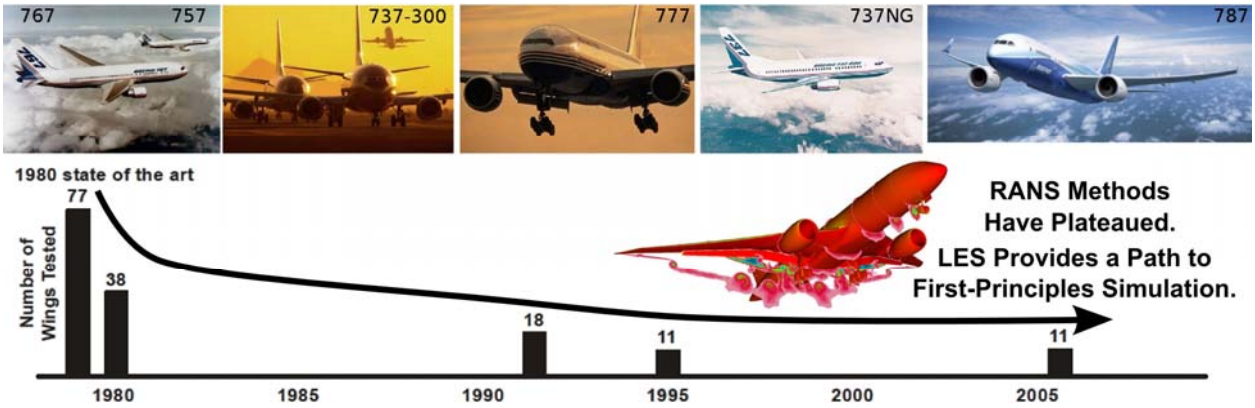

Figure 1: Timeline of the development of Boeing aircraft showing how the number of wings tested for each design has remained fixed at 11 for more than a decade indicating that existing RANS methods have plateaued in terms of predictive performance. LES methods provide a new path to first-principles simulation that, when combined with robust optimization strategies on exascale platforms, are expected to further reduce the number of required wing prototype tests and also provide higher-fidelity data, such as aerodynamic noise sources, that enable more robust design.

---

[9] Chapman DR. 1979. "Dryden Lecture: Computational Aerodynamics Development and Outlook." *AIAA Journal* **17**(12): 1293-1313; Mavriplis DJ, D Darmofal, D Keyes, & M Turner. 2007. "Petaflops Opportunities for the NASA Fundamental Aeronautics Program." AIAA Paper 2007-4084 (18th AIAA Computational Fluid Dynamics Conference, June 25-28, 2007).

[10] Garrett, M. 2006. Before the United States Senate Committee on Commerce, Science and Transportation Subcommittee on Technology, Innovation, and Competitiveness. Testimony of Michael Garrett Director, Airplane Performance Boeing Commercial Airplanes, July 19, 2006.

experimental as well as direct numerical simulation (DNS) data, but the models have not proven to be sufficiently universal and predictive when applied to situations outside the domain of calibration data.

The RANS approach suffers from the fundamental (epistemic) uncertainties in the functional forms of the turbulence closure models, which at least for the current models used in engineering practice, cannot be made more accurate by mere calibration using experimental or DNS data. Another important limitation of the RANS approach is the absence of a path to higher accuracy with increasing computer power. That is, the predictive accuracy of the RANS method is not necessarily improved by increasing the spatial and temporal resolution of the computational grid; it is limited by the (irreducible) accuracy of the turbulence closure models.

A different and transformational approach to computation of turbulent flows, that has received recent attention in the aerospace industry, is the large-eddy simulation technique. In this method one computes the large-scale field directly by solving the low-pass filtered Navier-Stokes and mass conservation equations, and incorporates phenomenological models to account for the effects of the small unresolved scales on the large scale field. In contrast to the RANS approach, LES has a well-defined path to increased predictive capability through increased computer power. Through increased spatial and temporal resolution, afforded by advances in computer technology, the LES equations approach the true flow equations derived from first principles as the effect of phenomenological subgrid scale models diminish. There is, however, one important caveat in the promise of LES that remains an important obstacle for its use, especially for external aerodynamics. Turbulence near a smooth wall (e.g., aircraft wings) consists of small eddies that are important to the overall flow dynamics and are responsible for a large fraction of viscous drag. Resolving these structures is not feasible in LES, as it requires enormous grid resolution comparable to that for DNS. Due to this limitation, hybrid RANS/LES methods are being considered for applications in external aerodynamics.

This improved CFD capability has also directly helped reduce the number of expensive high-pressure rig tests required for certification of jet engines by 50% at Pratt & Whitney in a similar timeframe wherein Boeing's use of CFD led to fewer wind tunnel tests. However, as in the case of Boeing, Pratt & Whitney has seen only a 10% further reduction in the required engine tests in the past decade[11]. Flow in gas turbine engines used for aircraft propulsion is more complicated than that encountered in external aerodynamics. A typical engine consists of thousands of moving parts throughout the turbomachinery components, multi-phase reacting flows in the combustor, and intricate cooling passages inside its turbine blades.

The gas-turbine industry also uses a RANS approach to compute the flow in turbomachinery passages. However, to simulate the complex turbulent mixing and multi-physics phenomena (liquid fuel break-up, mixing of hot gases and chemical reactions) in the combustor, higher-fidelity approaches such as LES appear to be necessary. The large scale turbulent mixing of the reactants inside the combustor is particularly suitable for LES solution and unsuitable for RANS. As a result, there has been an intense effort to develop the LES methodology for simulation of combustion in complex configurations. Figure 2 shows the progress made in the LES of turbulent combustion over the course of a decade. Given the computational resources of the time, the first LES of turbulent combustion was limited to gas phase combustion in a simplified coaxial research combustor chamber[12]. By the end of the decade[13], a realistic (i.e., a sector of a PW6000 engine) combustor was simulated and integrated with the RANS simulations

[11] Kim WW & S Syed. 2004. "Large-Eddy Simulation Needs for Gas Turbine Combustor Design." In *Proceedings of the 42nd AIAA Aerospace Sciences Meeting and Exhibit*, pp. 3162-3175, Reston, Virginia.

[12] Pierce CD & P Moin. 1998. "Large eddy simulation of a confined coaxial jet with swirl and heat release." AIAA Paper 1998-2892.

[13] Moin P & G Iaccarino. 2007. "Complex effects in large eddy simulations." *Lecture Notes in Computational Science and Engineering*, **56**: 1-14.

of the turbomachinery. The model included break-up of the injected liquid fuel into droplets, Lagrangian tracking and evaporation and burning of the fuel droplets.

Each new generation of computer has provided a path to increasingly first-principles, simulation-based design. The current industry standard simulation tools have plateaued in their ability to resolve the critical technical challenges that arise as the operability and performance thresholds of aerospace vehicle and jet engines are driven to their limits. To transform the aerospace and jet engine industries, a new suite of first-principles, simulation-based engineering design tools, such as LES and hybrid RANS/LES, must become a standard component in the engineering design cycle. This increased level of simulation fidelity has the potential to expand the current design envelope, eschew industrial stagnation and augment human creativity with first-principles, simulation-based engineering science tools. As history has shown, effective use of this new technology becomes practical when a new level of computing power becomes widely available.

Figure 2. A Motivating Example: Evolution of First-Principles Computational Capability

State-of-the-art in 1997	**State-of-the-art in 2007**
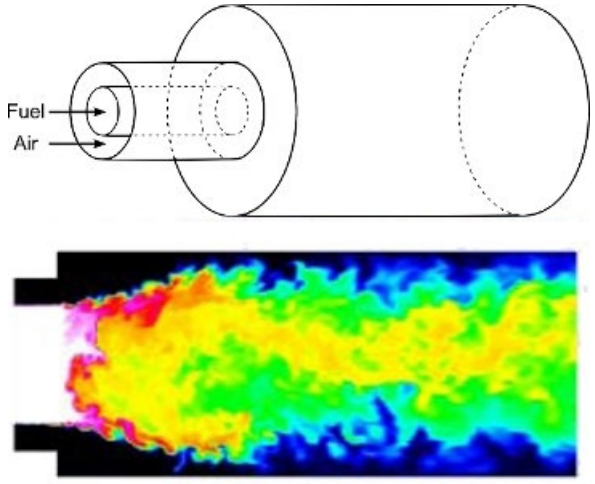	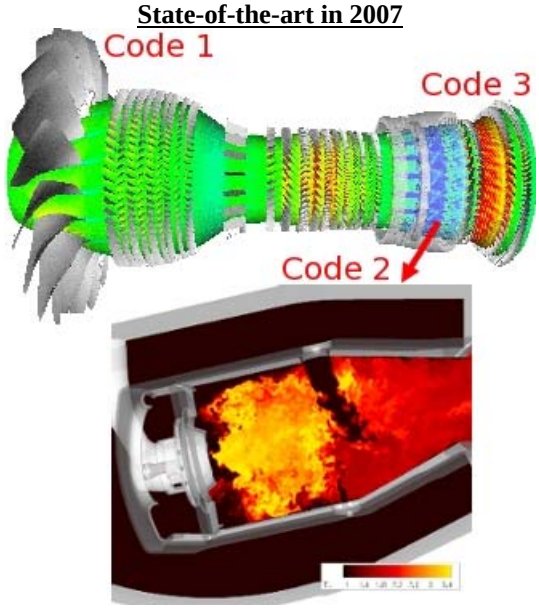
Shown are a schematic (top) of a coaxial combustor with air-methane co-flow and an instantaneous snapshot of the mixture fraction (bottom) from the first LES of turbulent combustion in a research combustor performed at Stanford University. The simulations used a flamelet and progress-variable approach with over two million degrees of freedom computed on the ASCI Red platform at Sandia National Laboratories. Note that the highest-fidelity simulations of reactive flow in 1997 were restricted to simple geometries and gas-phase only physics.	Comprehensive simulation of the flow through an entire jet engine performed by Stanford University, under the NNSA ASCI Program, and Pratt & Whitney. The compressor (Code 1) and the turbine (Code 3) were computed using RANS while the combustor (Code 2) used LES. Data is transferred across these interfaces using a coupled, multi-code, multi-physics integration environment. Contours of entropy in the compressor and turbine sections are shown while the flame in the combustor is visualized via contours of temperature in a mid-plane section from a realistic PW6000 engine.

3.2 Astrophysics

Astrophysics: Thermonuclear Supernovae

Type Ia supernovae (SNe Ia) are thermonuclear explosions of massive carbon-oxygen white dwarf stars in binary systems. They produce a significant fraction of the heavy elements and most of the iron in the universe. They are also of great importance in cosmology: observations using SNe Ia as "standard

candles" revealed that the expansion rate of the universe is accelerating and led to the discovery of dark energy[14]. Understanding dark energy ranks among the most compelling problems in all of physical science.

Currently, the calibration of SNe Ia as standard candles is entirely empirical. It uses a relation between the peak luminosity of SNe Ia and the rate at which they fade[15]. The accuracy of this calibration must be improved from its present value of ~15% to better than 1% in order to study quantitatively the behavior of dark energy with redshift (i.e., with the age of the universe) and therefore distinguish among the various explanations of it that have been proposed[16]. Providing a better means to calibrate SNe Ia as standard candles is a major goal of three-dimensional simulations of SNe Ias. An independent, reliable method of calibrating SNe Ia would make sure that any evolution of their properties with redshift does not confound their use as standard candles to determine the properties of dark energy accurately.

Comparison of computer simulations and observations led to the conclusion that SNe Ia explosions most likely involve two stages: a buoyancy-driven turbulent nuclear combustion phase that expands the white dwarf star, followed by a detonation phase that incinerates the star and causes it to explode[17]. However, a full understanding of the explosion mechanism does not yet exist. This is a consequence of the enormous disparity between the width of the nuclear flame, the detonation wave, the size of the star, the complexity of buoyancy-driven turbulent nuclear combustion, and uncertainty about how the detonation is triggered.

The increasing computational resources available to astrophysicists in recent years have produced deeper insights into buoyancy-driven turbulent nuclear combustions[18]. This increase in computational resources has also made possible multiphysics three-dimensional simulations of the explosion[19]. These simulations led to the discovery of an entirely new explosion mechanism[20], demonstrating that such simulations are necessary in order to capture with sufficient fidelity the physical processes that are involved.

Astrophysics: Core-Collapse Supernovae

Core-collapse supernovae are among the most violent events in the universe. They mark the death of massive stars (i.e., stars larger than roughly 10 times the size of our own sun) and the birth of neutron stars and stellar-mass black holes. Core-collapse supernovae serve both to synthesize new elements and to disperse elements synthesized in massive stars during their lifetimes. In the end, they are the dominant

---

[14] Riess AG, et al. 1998. "Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant." *AJ* **116**: 1009-1038; Perlmutter S, et al. 1999. "Measurements of Ω and Λ from 42 High-Redshift Supernovae." *ApJ* **517**(2): 565-586.

[15] Phillips MM. 1993. "The Absolute Magnitudes of Type Ia Supernovae." *ApJL*, **413**(2, Part 2): L105-108.

[16] Frieman JA, MS Turner, & D Huterer. 2008. "Dark Energy and the Accelerating Universe." *Ann. Rev. Astron. Ap* **46**: 385-432.

[17] Hoyle F & WA Fowler. 1960. "Nucleosynthesis in Supernovae." *ApJ* **132**(3): 565-590; Arnett WD. 1969. "A Possible Model of Supernovae: Detonation of ^{12}C." *A&A* **5**(2): 180-212; Hansen CJ and JC Wheeler. 1969. "A Calculation of a White Dwarf Supernova." *A&A* **3**(3): 464-474; Nomoto K, D Sugimoto, & S Neo. 1976. "Carbon Deflagration Supernova, an Alternative to Carbon Detonation." *A&A* **39**(2): L37-42; Khokhlov AM. 1991. "Delayed Detonation Model for Type IA Supernovae." *A&A* **245**(1): 114-128; Gamezo VN, et al. 2003. "Thermonuclear Supernovae: Simulations of the Deflagration Stage and Their Implications." *Science* **299**(5603): 77-81.

[18] Khokhlov AM. 1995. "Propagation of Turbulent Flames in Supernovae." *ApJ* **449**(2, Part 1): 695-713; Zingale M, et al. 2005. *ApJ*, **632**: 1021; Zhang J, OEB Messer, AM Khokhlov, & T Plewa. 2007. "On the Evolution of Thermonuclear Flames on Large Scales." *ApJ* **656**(1): 347-365.

[19] Gamezo VN, AM Khokhlov, & ES Oran. 2005. "Three-Dimensional Delayed-Detonation Model of Type Ia Supernovae." *ApJ* **623**: 337-346; Röpke FK & W Hillebrandt. 2005. "Full-Star Type Ia Supernova Explosion Models." *A&A* **431**(2): 635-645.

[20] Plewa T, AC Calder, & DQ Lamb. 2004. "Type Ia Supernova Explosion: Gravitationally Confined Detonation." *ApJL*, **612**(1, Part 1): L37-40; Jordan GC, et al. 2008. "Three-Dimensional Simulations of the Deflagration Phase of the Gravitationally Confined Detonation Model of Type Ia Supernovae." *ApJ* **681**(2): 1448-1457.

source of elements between oxygen and iron. They are also one of the only sites in the modern universe where neutrino interactions with matter have macroscopic, dynamic consequences. All of these facts combine to make core-collapse supernovae remarkable cosmic laboratories for nuclear and neutrino physics.

The explosion mechanism of core-collapse supernovae was among the first applications in the history of computational science. The first attempts to form a theory of such supernovae[21] led to the earliest approaches to simulate these events on computers[22]. These attempts were soon followed by the first generation of truly multiphysics simulations[23]. However, the complexity of the problem has meant that, for close to a half century, supernova modelers have struggled to determine the precise nature of the explosion mechanism and to use that understanding to produce quantifiable predictions of the consequences of these events.

In the past few years, as advances in computational power have enabled new levels of simulation, scientists have begun to peel back the layers of feedback-laden uncertainty in their understanding of core-collapse supernovae. Various nucleosynthesis calculations imply that the composition and distribution of the inner nickel-rich ejecta is very sensitive to the details of the explosion mechanism[24]. Simulations have shown that the neutrino-driven mechanism cannot work for spherical symmetry[25], and that multidimensional effects are advantageous for shock revival in the context of the delayed mechanism[26]. Along with the computational discovery of the standing accretion shock instability (SASI)[27], all of these findings lead to the conclusion that multiphysics simulations must be carried out in three spatial dimensions to achieve the requisite physical fidelity.

3.3 Climate modeling

Over the past decade, it has become increasing clear that clouds have a significant effect on the Earth's heat budget, and that changes in clouds affect the temperature change in global warming. This is called cloud feedback, and has posed the largest uncertainty in the study of climate sensitivity for almost twenty years[28]. Change that occurs in low clouds represents perhaps the largest uncertainty[29], although feedbacks

[21] Hoyle F. 1946. "Note on the Origin of Cosmic Rays." *MNRAS* **106**(5): 384-389.

[22] Colgate SA & RH White. 1966. "Hydrodynamic Behavior of Supernovae Explosions." *ApJ* **143**(3):626-681.

[23] Arnett WD. 1969. "A Possible Model of Supernovae: Detonation of ^{12}C." *A&A* **5**(2): 180-212; Arnett WD. 1967. "Mass Dependence in Gravitational Collapse of Stellar Cores." *Can. J. Phys.* **45**(5): 1621-1641.

[24] Fröhlich C, et al. 2006. "Composition of the Innermost Core-Collapse Supernova Ejecta." *ApJ* **637**(1): 415-426; Pruet J, et al. 2006. "Nucleosynthesis in Early Supernova Winds. II. The Role of Neutrinos." *ApJ* **644**(2): 1028-1039; Kifonidis K, et al. 2006. "Non-Spherical Core Collapse Supernovae – II. The Late-Time Evolution of Globally Anisotropic Neutrino-Driven Explosions and Their Implications for SN 1987 A." *A&A* **453**(2): 661.

[25] Liebendörfer M, A Mezzacappa, F-K Thielmann, OEB Messer, WR Hix, & SW Bruenn. 2001. "Probing the Gravitational Well: No Supernova Explosion in Spherical Symmetry with General Relativistic Boltzmann Neutrino Transport." *PR D* **63**(10): 103004-1–103004-13.

[26] Herant M, et al. 1994. "Inside the Supernova: A Powerful Convective Engine." *ApJ* **435**(1, Part 1): 339-361; Burrows A, J Hayes, & BA Fryxell. 1995. "On the Nature of Core-Collapse Supernova Explosions." *ApJ* **450**: 830-850; Janka HT & E Müller. 1996. "Neutrino Heating, Convection, and the Mechanism of Type-II Supernova Explosions." *A&A* **306**: 167.

[27] Blondin JM, A Mezzacappa. & C DeMarino 2003. "Stability of Standing Accretion Shocks, With an Eye Towards Core-Collapse Supernovae." *ApJ* **584**: 971.

[28] Bony, S, R Colman, VM Kattsov et al. 2006. "How well do we understand and evaluate climate change feedback processes?" *J. Climate* **19**: 3445–3482; Soden, B. J., and I. M. Held. 2006. "An assessment of climate feedbacks in coupled ocean–atmosphere models." *J. Climate* **19**: 3354–3360.

[29] Webb, MJ, ACA Senior, DMH Sexton et al. 2006. "On the contribution of local feedback mechanisms to the range of climate sensitivity in two GCM ensembles." *Climate Dyn.* **27**: 17–38.

from high clouds[30] also represent a major source of uncertainty[31]. All of the hypotheses relate the changes in high clouds to those in deep convection. Thus it is essential to understand cumulus convection to discuss the high cloud's feedback. The climate community has recognized the uncertainty of cumulus convection for a long time. The spatial resolution of current global climate models (GCMs) for climate simulations is on the order of 50 to 100 km, and this is in inadequate to resolve cumulus convection. As a result, GCMs contain parameterizations for clouds and convective processes, and systematic analyses of the uncertainties in GCMs have highlighted the role of these parameterizations.

This analysis highlights the need to transition to climate models that can simulate clouds and ocean eddies at their native scales, namely at global resolutions of 1 to a few kilometers.

In global simulations of climate change with current generation models, the forecast perturbations in column-integrated cloud amount are inconsistent across the multi-model ensemble for most regions on the globe. The models do not concur on either the magnitude or sign of changes in global mean cloud radiative effect, and the size of the cloud radiative effect changes relative to the present-day value is generally less than 10%; this is comparable to various forcings and current uncertainties in the top-of-atmosphere planetary energy budget. Thus, the model projections suggest that observational verification of global cloud radiative effect changes will be a significant challenge. The sign of radiative feedbacks from low clouds and deep convective systems is not consistent across the IPCC multi-model ensemble[32]. Tropical clouds, mixed-phase clouds and cloud phase, as well as intermodel differences in meridional shifts in storm tracks contribute the divergences in cloud responses[33].

Current parameterizations contain large numbers of free parameters that govern the evolution of micro- and macrophysical cloud properties in relation to the meteorological environment. As is well known, simulated climate radiative feedbacks and climate responses are quite sensitive to variations in these free parameters[34]. It has proven difficult to exclude parameter settings that produce large positive cloud feedbacks and climate sensitivities to in greenhouse gas concentrations. Many of the parameters are only loosely constrained by observations or process-oriented modeling and are connected to the physics of cloud formation through scaling arguments.

3.4 Combustion

Burning fossil fuel generates over 85 percent of the world's energy. Therefore, to minimize pollution, optimize efficiency and to prevent further climate change, combustion must be understood and controlled. Although it is widely recognized that diversified energy sources must be developed, optimizing combustion is still the most significant route today to controlling global change and will remain the primary source of energy for at least the next fifty years. Combustion design engineers must understand more precisely how combustion occurs in engines and gas turbines for transportation and power

---

[30] Lindzen, RS. 1990. "Some coolness concerning global warming." *Bull. Amer. Meteor. Soc.* **71**: 288–299; Ramanathan, V, & W Collins. 1991; "Thermodynamic regulation of ocean warming by cirrus clouds deduced from observations of the 1987 El Niño." *Nature* **351**: 27–32; Lindzen, R. S., M.-D. Chou, & A. Y. Hou. 2001. "Does the Earth have an adaptive infrared iris?" *Bull. Amer. Meteor. Soc.* **82**: 417–432.

[31] Lau, K-M, C-H Sui, M-D Chou, & W-K Tao. 1994. "An inquiry into the cirrus-cloud thermostat effect for tropical sea surface temperature." *Geophys. Res. Lett.* **21**:1157–1160; Lin, B, BA Wielicki, LH Chambers, Y Hu, & KM Xu. 2002. "The iris hypothesis: A negative or positive cloud feedback?" *J. Climate* **15**: 3–7.

[32] Bony, S, & J-L Dufresne. 2005. "Marine boundary layer clouds at the heart of tropical cloud feedback uncertainties in climate models." *Geophys. Res. Lett.* **32**: L20806.

[33] IPCC. 2007. Climate Change 2007: "The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change." ed. S. Solomon, D. Qin, M. Manning et al. New York: Cambridge Univ. Press.

[34] Senior, CA, & JFB Mitchell. 1993. "Carbon dioxide and climate: The impact of cloud parameterization." *J. Climate* **6**:393–418; Murphy, JM, DM Sexton, & DN Barnett et al. 2004. "Quantification of modeling uncertainties in a large ensemble of climate change simulations." *Nature* **429**: 768–772.

generation, and high performance simulation is currently the most efficient way to investigate these phenomena. The complexity of these flows is such that up until recently, simulations of combustion in complex geometry chambers were either impossible or not predictive. For the first time, petascale computing has provided results that can be directly used in industry to build gas turbine engines or to modify their design to prevent combustion instabilities. HPC is beginning to having a direct impact on industry design processes and providing more complete understanding of mixing and combustion in internal combustion engines. For example, recently Cummins used engineering computational tools to bring a diesel engine to market with minimal testing, resulting in a reduction of 10-15% in development time and cost. While the incremental computational tools were based on Reynolds-averaged engineering approximations, the automotive and fuels industries will benefit enormously from more advanced large-eddy simulation approaches in the future that capture cycle-to-cycle variability. Large eddy simulation has been more widely developed and adopted by both the aerospace and aircraft engine manufacturers, for example, for understanding the coupling between acoustics, heat release rate, and engine geometry in 'acoustic instability' described below (see Section 3.1). Similar advances can be made in the automotive industry to optimally design novel engines that are more fuel efficient utilizing alternative fuels – carbon neutral renewable bio-fuels.

For example, for the first time, a simulation of a complete gas turbine chamber (Figure 3) including the explicit resolution of unsteady turbulent motions and acoustic waves propagating in such engines was performed using petascale computers (BlueGene/P located at ALCF at Argonne National Laboratory). The simulation technique used (LES) was able to reproduce a phenomenon called ‘combustion instability’ which has been observed to occur in real engines (an helicopter chamber). Combustion instabilities pose serious challenges in helicopter engines but also in industrial furnaces, rocket, aircraft or industrial turbines: they prevent engine operation at low levels of emissions, lead to catastrophic engine failures and accidents and were impossible to predict prior to the advent of petascale computing. Having the ability to simulate these instabilities using LES on petascale machines enables engineers to design control methods to prevent their occurrence.

LES and other engineering level combustion simulations require subgrid models representing combustion and turbulent transport processes. Complementing LES, at the microscale, direct numerical simulation (DNS) is a first principles approach to investigate interactions between turbulence and chemical kinetics associated with burning petroleum and renewable fuels. The resulting benchmark data not only provide fundamental insight, but also provide unique validation data for model development, in many cases unattainable by experimentation due to the harsh combustion environments.

The advent of petascale computing power has made it possible to glean fundamental physical insight into fine-grained ‘aero-thermo-chemical’ interactions in small laboratory-scale turbulent flames with detailed chemistry using three-dimensional DNS. The instantaneous governing equations are solved without averaging or filtering; all relevant continuum scales are resolved on the grid with no closure models for turbulence or combustion using accurate numerical methods. Such simulations are costly, requiring tens

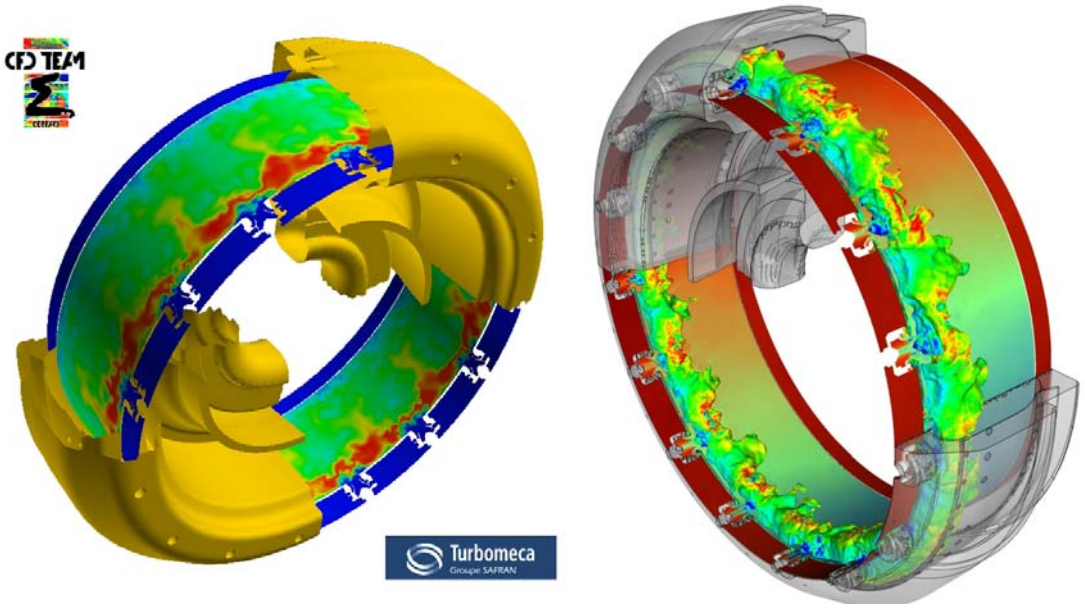

Figure 3: (Left) Instantaneous temperature field on a cylinder view plane passing through the injectors of the helicopter chamber. Red zones correspond the hot gases and blue zones to cold air coming from the compressor. (Right) Instantaneous pressure field in the same view plane with isosurfaces of temperature.

of million of cpu-hours on a petascale computer, up to several billion grid points, and generating 100's of terabytes of raw data. High performance computing at the petascale and beyond is providing access to direct simulation of a large dynamic range of scales [$O(10^4)$] and the transport of ever increasing numbers of species representative of hydrocarbon fuels of practical interest.

For example, a recent direct numerical simulation (DNS) of the near field of three-dimensional spatially-developing turbulent ethylene jet flame in highly-heated coflow was performed on Jaguar, a CrayXT5 at Oak Ridge National Laboratory with a reduced mechanism transporting 22 species to investigate the stabilization mechanism, which may have common features to diesel engine lifted stabilization. The DNS was performed at a jet Reynolds number of 10,000 with over 1.29 billion grid points. Figure 4 shows instantaneous images of several scalar fields – scalar dissipation rate (i.e., local mixing rate), mixture fraction, and hydroperoxy, methyl and hydroxyl intermediates – obtained from the DNS data. The results

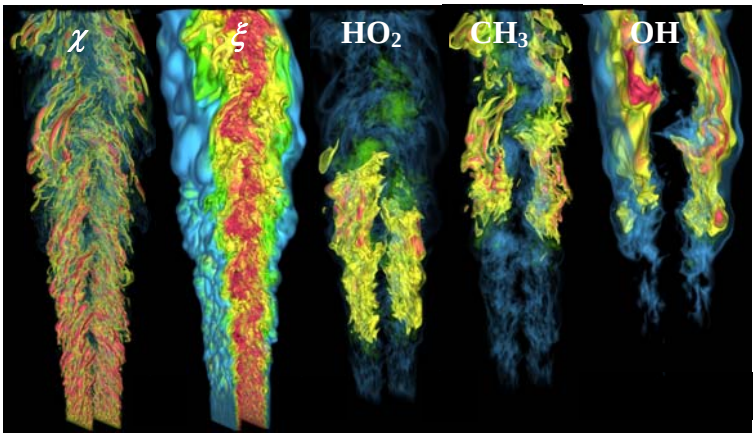

Figure 4: Instantaneous volume rendering of scalar dissipation rate χ, mixture fraction ξ, HO_2, CH_3 and OH radicals in a lifted ethylene/air jet flame (Re=10,000) in a hot coflow.

show that auto-ignition in a fuel-lean mixture at the flame base is the main source of stabilization of the lifted jet flame.

This was verified by applying several diagnostics to the lifted jet flame DNS including the development and application of a chemical explosive mode (CEM) analysis, a Damköhler number defined by the CEM and the local scalar dissipation rate, and by Lagrangian tracking of key scalar and velocity statistics at the stabilization point. Many of these diagnostics are only available through detailed simulation such as presented here, as even the most advanced optical laser diagnostics are limited to measurements in lower-dimensions and of only a few species. Therefore, DNS complements experimentation in providing a more complete picture of combustion processes, and validation data required for model development.

3.5 Fusion Energy

Fusion energy research has historically been a leading HPC application domain. Significant progress in both particle and fluid simulations of fine-scale turbulence and large-scale dynamics in magnetically-confined plasmas have been enabled by the combination of access to powerful supercomputing resources at the petascale level, together with innovative advances in analytic and computational methods for developing reduced descriptions of physics phenomena spanning a huge range in time and space scales. In particular, the fusion community has made excellent progress in developing advanced codes for which computer run-time and problem size scale well with the number of processors on massively parallel machines. A good example is the effective usage of the full power of multi-teraflop to petaflop systems to produce three-dimensional, general geometry, nonlinear particle simulations which have accelerated progress in understanding the nature of plasma turbulence in fusion-grade high temperature plasmas. These calculations, which typically utilize multi-billions of particles for thousands of time-steps, would not have been possible without access to such powerful modern supercomputers together with advanced diagnostic and visualization capabilities to help interpret the results. This is significant because one of the central plasma physics problems on the road to creating a working fusion power plant is understanding, predicting, and controlling instabilities caused by unavoidable plasma inhomogeneities. One consequence is the occurrence of turbulent fluctuations ("microturbulence") which can significantly increase the transport rate of heat, particles, and momentum across the confining magnetic field – an overall effect that severely limits the energy confinement time for a given machine size and therefore the performance and economy of a tokomak device[35]. Accelerated progress on this critical issue is especially important for ITER because the size and cost of a fusion reactor is determined by the balance between such loss processes and the self-heating rates of the actual fusion reactions[36]. Petascale computational modeling and simulation is irreplaceable in dealing with such challenges because of the huge range of temporal and spatial scales that must be taken into account. Existing particle-in-cell (PIC) techniques have clearly demonstrated excellent scaling on current terascale leadership class supercomputers. For

[35] Tang, WM & VS Chan. 2005. "Advances and Challenges in Computational Plasma Science," *Plasma Phys. Control. Fusion* **47**(2): R1-R34.

[36] Batchelor, DA, et al. 2007. "Simulation of Fusion Plasmas: Current Status and Future Direction." *Plasma Science and Technology*. **9**: 312-387; Tang, WM, et al. 2010. "Scientific Grand Challenges: Fusion Energy Sciences and the Role of Computing at the Extreme Scale." DoE-SC Peer-reviewed report on major workshop held March 18-20, 2009, Final Publication, 212pp. DOE. 2010. Scientific Grand Challenges: Fusion Energy Sciences and the Role of Computing at the Extreme Scale: March 18-20, 2009, Washington D.C. PNNL-19404; production support provided by Pacific Northwest National Laboratory for the U.S. Department of Energy's Office of Advanced Scientific Computing Research and the Office of Fusion Energy Sciences, Washington, D.C. Available at http://www.er.doe.gov/ascr/ProgramDocuments/Docs/FusionReport.pdf

example, the Gyrokinetic Toroidal Code (GTC)[37], a mature PIC code, has demonstrated excellent scaling on virtually all of the current leadership class facilities worldwide utilizing MPI and OpenMP[38].

In general, reliable predictive calculations of turbulent transport can only be achieved through experimentally validated simulations, as it fills the space between empirical observations – which are expensive and difficult to acquire – and theory, which cannot handle the nonlinearity of the problem under realistic experimental conditions. To do this, the majority of researchers world-wide uses a gyrokinetic (*ab initio*) approach. For example, PIC codes solve the nonlinear equations underlying gyrokinetic theory, with excellent scaling to more than 100,000 cores having already been demonstrated. However, in order to move in a timely manner to producing simulations with the highest possible physics fidelity, it is expected that computing at the exascale will be necessary. The associated software/algorithmic advances needed must be compatible with the evolving hardware and therefore developed in a true "co-design" sense. With the unprecedented resolution in a multi-dimensional phase-space enabled by access to HPC platforms at the petascale and beyond to the exascale, these advanced kinetic simulation capabilities are expected to have direct relevance to existing experimental devices as well as to ITER. The impressive advances achieved by such codes give great promise of delivering scientific discoveries appropriate for "path to exascale" grand challenges.

In common with general PIC approaches, the gyrokinetic particle-in-cell method[39] consists of moving particles along the characteristics of the governing equation – here the 5-dimensional GK-equation. The equation increases in complexity due to the fact that the particles are subjected to forces from an externally imposed (equilibrium) magnetic field and from internal electromagnetic fields generated by the charged particles. In the "scatter" phase of a PIC simulation, a grid is used to map the charge density at each point due to the particles in the vicinity; the Maxwell equations governing the fields are then solved for the forces that are then gathered back to the particles' positions during the "gather" phase of the simulation. Finally, during the "push" phase of the simulation, this information is used for advancing the particles by solving the equations of motion. The original parallel scheme implemented in GTC consists of a one-dimensional domain decomposition in the toroidal direction and a particle distribution within these domains[40]; each process is in charge of a domain and a fraction of the number of particles in that domain, and inter-process communications are handled with Message Passing Interface (MPI) calls. This method scales extremely well to a large number of processors, but is eventually dominated by communications as more particles move in and out of the shrinking domains at each time step. Tests on the dual core BG-L system were conducted in co-processor mode, where one BG-L core is used for computation and the second is dedicated to communication[41]. Additional tests were then conducted in virtual mode node, where both cores participate in both computation and communication. Impressive results were obtained in virtual node mode with a per-core efficiency of over 96%.

---

[37] Lin, Z, et al. 1998. "Turbulent transport reduction by zonal flows: Massively parallel simulations," *Science* **281**, 1835-1837; Lin, T., el al. 2000. "Gyrokinetic simulations in general geometry and applications to collisional damping of zonal flows," *Phys. Plasmas* **7**(5): 1857-1862.

[38] Ethier, S, et al. 2010. "Large Scale Gyrokinetic Particle Simulation of Microturbulence in Magnetically Confined Fusion Plasmas," to appear in the IBM Journal of Research and Development; Oliker, L, *et al*. 2007. "Scientific Application Performance on Candidate PetaScale Platforms," *In Proc. IPDPS'07*, Long Beach, CA, March 24-30; Ethier, S, WM Tang, & Z Lin. 2005. "Gyrokinetic particle-in-cell simulations of plasma microturbulence on advanced computing platforms," *J. Phys.: Conf. Series* **16**: 1-15; Oliker, L, *et al*. 2005. "Leading computational methods on scalar and vector HEC platforms," In *Proc. SC2005*, Seattle, WA, Nov 12-18; Oliker, L, et al. 2004. "Scientific Computations on Modern Parallel Vector Systems," In *Proc. SC2004*, Pittsburgh, PA, Nov 6-12.

[39] Lee, WW. 1983. "Gyrokinetic approach in particle simulation," *Phys. Fluids* **26**(2): 556-562; Lee, WW. 1987. "Gyrokinetic Particle Simulation Model," *J. Comput. Phys.* **72**: 243-269.

[40] Ethier, S, WM Tang, & Z Lin. 2005. "Gyrokinetic particle-in-cell simulations of plasma microturbulence on advanced computing platforms," *J. Phys.: Conf. Series* **16**: 1-15.

[41] Ethier, S, et al. 2010. "Large Scale Gyrokinetic Particle Simulation of Microturbulence in Magnetically Confined Fusion Plasmas," to appear in the IBM Journal of Research and Development

The greater success of the PIC codes as compared to others in the Fusion Energy Sciences (FES) portfolio provide important "lessons learned" of how modern algorithmic approaches can expedite progress toward needed new insights. Motivated by the strong shift to multi-core environments extending well beyond the quad-core level, PIC studies have focused on improving current programming frameworks – such as systematically testing the proposed two-level hybrid programming model. In making full use of the multiple cores on a node, it is necessary to avoid being constrained to an MPI process on each core. Since some arrays get replicated on all these processes, a memory limit is encountered for the larger problem sizes of interest. The "hybrid programming" method discussed in Section 5 provides a much more viable way to mitigate this problem. Results supporting the efficacy of such an approach have been successfully obtained on the quad-core Leadership Class Facilities – Intreprid (IBM Blue-Gene-P) at Argonne National Laboratory (ALCF) and Jaguar CRAY XT-4 at Oak Ridge National Laboratory (OLCF).

As a specific example, the full deployment of the more efficient hybrid OpenMP/MPI algorithm together with a new radial domain decomposition feature in the GTC-P code operating on the Intrepid LCF at ANL. This capability greatly facilitates examining the key question of how plasma microturbulence properties might be affected as the plasma size increases from that of existing experiments to the future very large plasmas characteristic of ITER. In particular, it is now possible to efficiently examine the global turbulence properties in devices ranging from current scale experiments, which exhibit an unfavorable "Bohm-like" scaling with plasma size to the ITER scale plasma which is expected to exhibit a more favorable "gyro-Bohm" scaling of confinement. Multi-scale global kinetic simulations of plasma microturbulence, which take into account both the micro and meso scales of interest, can systematically investigate the turbulent transport properties of a magnetically-confined thermonuclear plasma spanning the size of existing machines today to the ITER device which will be more than 3 times larger in minor radius. Gaining improved understanding of the expected improvement confinement for these much larger ITER-scale plasmas is critically important because the microturbulence-driven nonlinear diffusion of particles and energy works against the plasma self-heating produced by the fusion reactions. The "scientific discovery" aspect of such studies is that while the simulation results can be validated against present-day tokamaks, there are no existing devices today that have a minor radius that is even one-third the size of ITER. Accordingly, the role of high physics fidelity predictive simulations takes on an even more important role – especially since the expected improvement in confinement for ITER-sized devices cannot be experimentally validated until after it is constructed and operational. Such advanced kinetic simulations provide also help illustrate how computational capabilities at the extreme scale can expedite the delivery of key scientific insights appropriate for grand challenges at the petascale and beyond. The "weak" scaling of GTC-P (to over 130,000 processors) enabled by the parallel algorithm implemented with MPI and OpenMP and developed specifically for the IBM-Blue Gene-P system at the ALCF is illustrated in Figure 5 below.

Overall, the excellent scaling of fusion PIC turbulence codes on the LCF's provide encouragement that moving forward to the multi-petascale (e.g., BG-Q at ANL) and eventually to the exascale is an achievable goal that will enable incorporation of significantly improved fusion energy science physics fidelity. In general, new insights gained from advanced simulations provide great encouragement for being able to include increasingly realistic dynamics to enable deeper physics understanding of plasmas in both natural and laboratory environments.

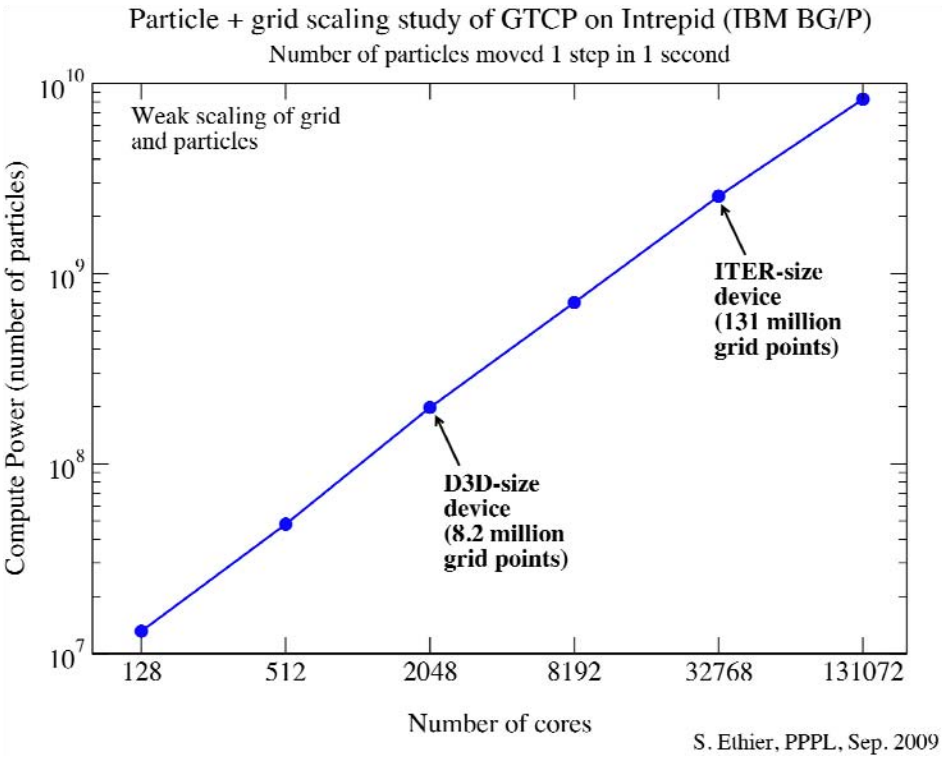

Figure 5: Speedups of the GTC-P code on the quad-core IBM Bkye-Gene-P (courtesy of S. Ethier)

3.6 Nuclear Engineering

With the growing interest in operating the existing fleet of NPPs beyond their original design lifetimes, in constructing new NPPs, and in developing and deploying advanced nuclear energy systems to meet the rising demand for carbon-free energy, there are significant opportunities for the application of petascale and exascale computing. Challenges currently confronting the nuclear industry include: running plants at higher temperature, higher power ("uprates"), greater burn-up, and increased efficiencies; having a diversity of plant sizes; affordable high temperature process heat; and reduced product development time. Computational thrusts include material properties, fuel cycle design and optimization, reactor performance (macroscopic, microscopic fuel performance), radiochemical separation, and waste form behavior.

Basic material properties, including nuclear (neutron and gamma reactions), thermophysical (e.g. thermal conductivity, phase diagrams), mechanical (e.g. tensile properties, fracture toughness) and chemical (e. g. corrosion rates) have to be determined under static and dynamic conditions. More accurate and reliable estimations of these properties will result from insight gained by high-fidelity first-principles atomistic and mesoscale models, with a potential for a huge return on investment in areas such as fuel process selection. Reprocessing was abandoned in the 1970s as an option within the current nuclear fuel cycle, so there is great opportunity for development of advanced processes. Current reprocessing models provide only qualitative descriptions of process behavior. Empirical models of chemical behavior for major components are used to provide overall descriptions of various reprocessing strategies. Fuel development and performance evaluation is currently an empirical process that takes decades. For acceptance, new fuels must be fabricated, placed in a test reactor over multiple cycles, tested under multiple accident scenarios, undergo post-irradiation examinations, and finally be placed in an operational reactor for several cycles. Fuel performance simulation tools can help to accelerate current fast and thermal reactor fuel qualification practices by helping decrease the time required to qualify new fuels, with the goal being

a reduction by a factor of three in the current 10- to 15-year qualification time. The tools must accommodate all relevant fuel types in both normal operating (quasi steady state) and transient conditions. Because current models are based on empirical curve fits to fuel behavior in common nuclear reactor environments, they typically cannot be relied upon to predict the behavior of fuels under conditions outside their narrow range of calibration. Reactor safety and design simulation tools require models for thermal hydraulics, neutronics, and structural mechanics. Such "reactor core codes" have been in existence for decades, but need improved physical, numerical, and geometric fidelity. Codes developed at the time used lumped parameter models for predictions of neutronics, thermal hydraulics, and structural mechanics quantities. These simple codes were calibrated against a very large experimental database, developed over the years for specific projects by calibrated against principally integral data. Neutronics modeling has traditionally relied on both stochastic (Monte Carlo) simulations and deterministic transport and diffusion theory approaches. Monte Carlo techniques incorporate the basic physics at the level of stochastic particle tracking with the general system geometry and material cross sections governing the particle track histories. Monte Carlo offers the strong conceptual advantage of keeping a close (essentially exact) correspondence between computational and physical geometric and cross section energy dependence models. Nevertheless, Monte Carlo can become computationally impractical for several different classes of problems. Existing reactor core codes typically employ the traditional single-channel or sub-channel approach to model reactor core heat transfer in the flowing coolant. Traditionally, separate thermal hydraulic code systems are used to execute design and safety calculations. Structural mechanics software development has been driven by a breadth of applications over the last 30 years which include automotive, aerospace, national defense, civil infrastructure, and, in the 1970's and 80's, nuclear reactor technology. These developments have led to a number of finite element-based computer programs that have relatively mature element technologies and solution algorithms.

As an example of how M&S in nuclear energy has evolved in the petascale era, consider the seven-dimensional phase-space (location, velocity [energy + angle], and time) solutions required of a Boltzmann transport model to deliver a reliable assessment of neutron spectra and distribution throughout a nuclear reactor core. The scale of the problem is enormous: 5 orders of magnitude in space and 10 in neutron energy. A transport solver that incorporates a resolved discretization for all scales using current discrete models would require $10^{17\text{-}21}$ degrees of freedom (DOF) for a single time step, which are beyond even exascale computational resources. This problem size precludes, then (for the time being), a single integrated ab initio computational approach. Instead, variations of current "multilevel techniques" will have to continue, with the new objective being to make the process more consistent. Present reactor transport methods use an inconsistent "three-level" homogenization approach[42], often utilizing distinctly different simulation codes at each level, in modeling neutron transport in the core of a nuclear reactor. The spatial and energy domains of this three-step approach typically follows this prescription: (1) a fine mesh in 1-D cylindrical geometry of an approximate small subset (pincell) of the reactor core with a first-principles representation of the energy spectrum; (2) a coarser mesh with a 2-D transport solution in a larger subset (lattice) of the core with grouped representation of the energy spectrum provided by the previous step; and, (3) a very coarse mesh in a 3-D diffusive transport of neutrons in the full homogenized core of the reactor with a very coarse representation of the energy spectrum provided by the previous step. The first two steps require $10^{7\text{-}8}$ DOF with $10^{2\text{-}3}$ independent calculations, each on single-processor machines. Recent work on petascale resources [Evans ref] has demonstrated that this three-level approach can at least be reduced to an "inconsistent two-step" approach, where step (A) utilizes the energy fidelity of step (1) with the spatial domain of step (2) and step (B) uses the energy fidelity of step (2) and the spatial domain of step (3). In this approach, now possible at the petascale, each step requires $10^{11\text{-}13}$ DOF per step with 10^2 independent high-order (A) calculations for every time step.

[42] Evans, TM, GG Davidson, & RN Slaybaugh. 2010. "Three-Dimensional Full Core Power Calculations for Pressurized Water Reactors." *SciDAC 2010 Conference* (July, 2010, Chattanooga, TN)

4. What Applications May Be Transformed By Going to the Exascale?

As discussed earlier, a key question to be addressed in considering going to the exascale is the readiness of key applications to take this step, as well as the likelihood that taking this approach will lead to transformative changes in these application areas. This question is addressed in the present section, focusing once again on a selection of disciplines to illustrate the breadth of applications that are ready for this transition; a much more complete and through picture can be obtained from the detailed disciplinary exascale reports listed in Appendix 2.

4.1 Aerospace, Airframes and Jet Turbines

The impact of the aerospace industry on the U.S. economy is significant, with an estimated $170 billion in sales and approximately 625,000 people employed in 2005. The aerospace industry, however, faces several challenges, such as a national air traffic system that, in its present form, cannot handle expected future increases in demand, an aging aerospace workforce, heightened regulatory requirements, and an increasingly competitive global market. In this regard, transformative shifts in the design and construction of aerospace systems are crucial to maintain a healthy aerospace industry presence in the US. Computing at an extreme scale will have transformational effects on several key applications in the aerospace industry. As pointed out in Section 3.1, the move from RANS to LES as the industry standard and its use in the design cycle represents a paradigm shift for the aerospace industry. In addition, there are several outstanding scientific problems in these sectors that can be understood and hopefully controlled using extreme scale computing.

As pointed out in Section 3.1, the accuracy achieved with the RANS approach for prediction of quantities of engineering interest in the airframe industry has reached a plateau owing to the epistemic uncertainties inherent in such turbulence models. As a result, the design of aircraft and propulsions systems relies on an iterative process where several expensive prototypes are constructed and tested in wind tunnels. Hybrid RANS/LES approaches with grounding in the first principles can overcome the limitations of RANS and enhance the predictive capability of CFD beyond the present seemingly stagnant state of speculative trial-and-error in design. In addition, building a complete flight-envelope characterization (accounting for irreducible uncertainties, e.g. angle-of-attack, flight conditions, and geometry) will only be possible with computing at the Exascale and beyond. Such a design framework for aerodynamically optimized vehicles and propulsion systems is a critical resource for the design and construction of next-generation aircraft and propulsion systems. Figure 6 provides estimates of the computing requirements to meet these design goals to address several Grand Challenges in aerospace systems.

One of the major problems confronting the aircraft industry is the aerodynamic noise generated by engine exhaust jets and airframes, particularly during take-off and landing approaches. The noise problem has been a major issue for high-speed commercial aircraft (e.g., the Concorde, and the abandoned supersonic transport program in the U.S.) and more recently for military aircraft, both for impact on communities surrounding airports and military bases, and on the crew stationed on aircraft carrier decks. In 2009, the Veterans Administration spent nearly a billion dollars for hearing loss compensation. It is known that turbulence is a major contributor to aircraft noise. Unfortunately, modern optical diagnostic techniques are far from adequate in measuring the spatio-temporal data needed to reveal the mechanics of aerodynamic noise; only high-fidelity simulation techniques, such as LES, are capable of predicting both the far-field noise as well as details of the noise generating turbulent eddies.

Exascale computing would have transformational impact on the discovery of the mechanics of noise generation, and would be instrumental in designing noise mitigation strategies. Figure 7 shows the turbulent flow from a supersonic exhaust jet obtained from a breakthrough state of the art LES computation in 2010. This first-of-a-kind “hero”-calculation lacks high-fidelity representation of the flow inside the nozzle, and the agreement with the measured noise data is only fair, presumably due to this

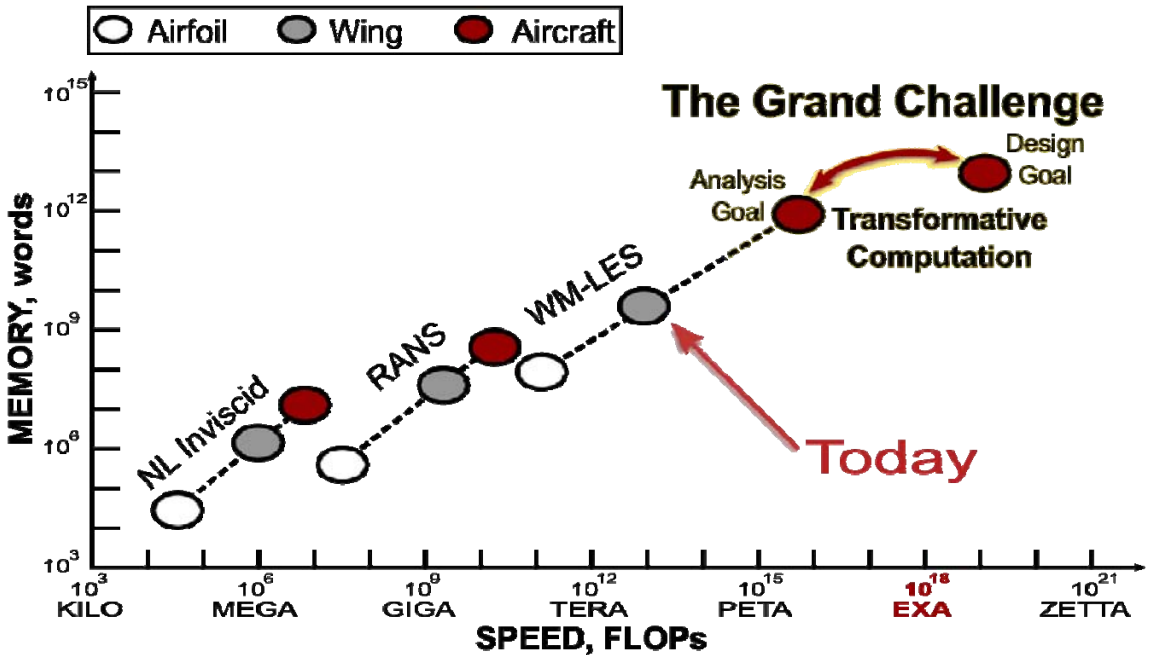

Figure 6: Computer speed and memory requirements for analysis and design of airfoils, wings, and complete aircraft for three different stages of approximation. Wall-modeled LES and hybrid RANS-LES methods provide a clear path to first-principles design of next-generation aircraft as exascale computing arrives. Transitioning this technology to future exascale platforms will have a transformative impact upon simulation-based engineering design, making possible the design of aerodynamically optimized vehicles including integrated effects of propulsion, structures, and active controls, a "Grand Challenge" of aerodynamic design

inadequate grid resolution. As exascale computing tools become available, high-fidelity tools would not only be used to understand and predict flow-generated noise, they will be used to learn how to control it. It has been demonstrated that the tools of optimal control theory[43] and high-fidelity CFD codes can be coupled to achieve flow control[44]. Such demonstration calculations have been extremely compute-intensive, and limited to very simple flows. Exascale computing would be the enabling technology for complex flow control and shape optimization (e.g., of aircraft wings and nozzle exits), potentially leading to a major transformational effect on the aerospace industry.

One of the outstanding technical problems in the gas-turbine industry is the migration of hot fluid parcels from the combustor to the turbine. The hot-streak migration is a limiting factor in the design of turbines, as turbine blades, designed based on mean flow temperatures, are damaged severely when encountering the migrating hot-spots. High-fidelity simulation of the flow inside the combustor of a jet engine is a daunting task due to the multi-physics phenomena present. Even in the modern LES computations of combustors using petascale class computers (see Figure 6), reduced order models are used for critical phenomena such as primary atomization of the injected liquid fuel into micron size droplets, the evaporation process of the droplets and the chemical mechanisms involved. Exascale computing would be the enabling technology for simulation of the jet engine combustors based on first principles, which in turn promises to facilitate the discovery of mitigating strategies for the suppression of the hot-streak migrations.

[43] Bewley TR. 2001. "Flow control: New Challenges for a New Renaissance." *Progress in Aerospace Sciences*, **37**: 21-58.

[44] Marsden AL, M Wang, JE Dennis Jr. & P Moin. 2004. "Optimal aeroacoustic shape design using the surrogate management framework." *Optimization and Engineering*, **5**(2): 235-262.

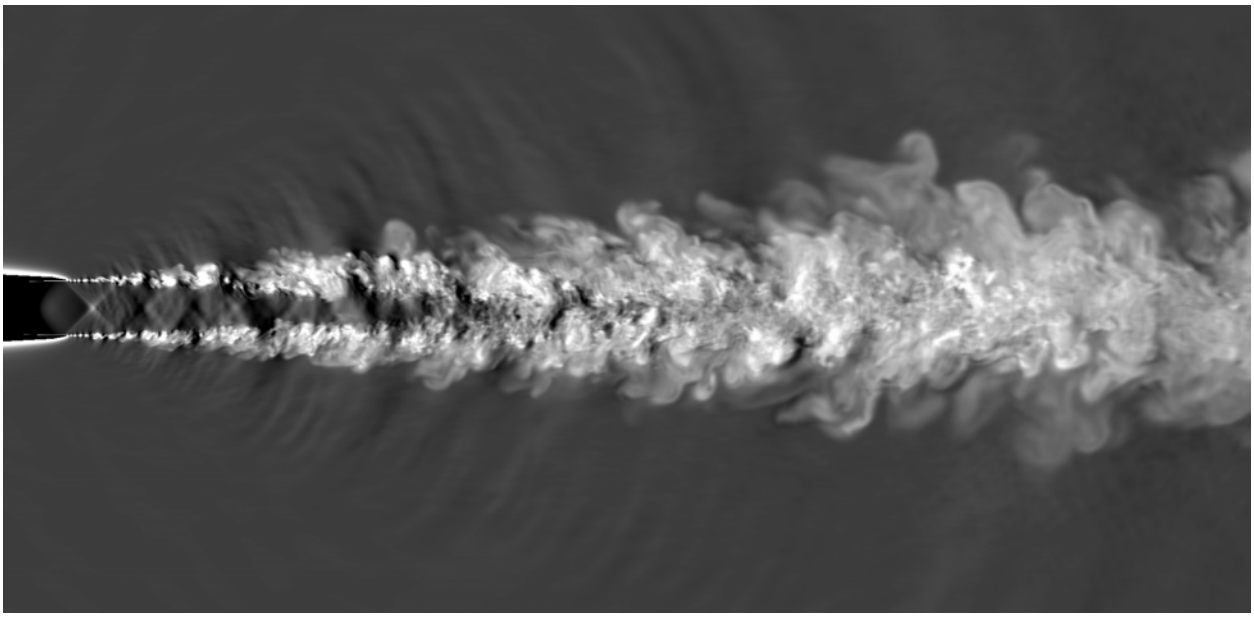

Figure 7: A supersonic jet engine nozzle (M=1.7) rapidly accelerates high-pressure gas into the atmosphere. A complex shock train forms inside the jet plume generating large-scale eddies and fine-scale turbulence further downstream, all of which contribute to the radiated noise seen propagating downstream of the nozzle. Courtesy of the Center for Turbulence Research, Stanford University.

The expansion and widespread adoption of high-performance computing technology in government, academia and industry provides the path to first-principles design and discovery. The aerospace and jet-engine industries have a significant motivation to apply these new tools to resolve the higher-fidelity details of turbulent mixing, chemical reaction, vibrations, multiphase spray dynamics, and species transport present in fully coupled engine and aircraft operation. This improved physics understanding, combined with refined geometrical and physical fidelity, enables the development and adaptation of reduced-order models for parametric studies of such systems. From this, a streamlined optimization methodology for modern aerospace systems will be possible, guided by the need to appropriately apply physics-based models of the required level of fidelity and thereby accurately predict and quantify the uncertainty in these virtual experiments. The transformational impact such a design tool will have upon these industries would be a direct result of the investment in extreme-scale computing.

As complex engineering systems are driven to their fundamental limits, more accurate computational analysis tools to reduce the cost and schedule during the design, operation, and maintenance of these systems will be needed by industry. Adopting physics-based design tools, combined with reduced-order models for engineering design optimization, are needed as soon as such tools and models are practical. The development and deployment of extreme-scale computing architectures will facilitate a fully coupled, system-level optimization and engineering design cycle that will enhance industrial competitiveness.

4.2 Astrophysics

Astrophysics has seen advances driven by the huge improvements in modeling and simulation capabilities over the past forty years like perhaps no other physical discipline. As a science based virtually entirely on remote observation, modeling and simulations have become an essential tool for assimilating the enormous and complex data sets generated by modern astronomical observations, for modeling these

observations, and for reaching an understanding of astrophysical processes by using simulations as a virtual experimental laboratory. These advances have touched essentially every aspect of astrophysics, from solar and stellar astrophysics to galactic astronomy, extragalactic astronomy and cosmology. For example, in the latter-most case, the combination of new observational tools and techniques and powerful simulation tools based on N-body and hydrodynamic codes have led to the present era of precision cosmology. While computations have had a remarkable transformational effect on astrophysics, brevity of space dictates that we only focus on one area of astrophysics where computations have allowed the kind of progress that simply would not have occurred in the absence of high-performance computations – this is the area of exploding stars, one of the most dramatic endpoints of stellar evolution.

4.2.1 Thermonuclear Supernovae

Explosions of thermonuclear supernovae (usually referred to as Type Ia supernovae, or SNe Ia) involve hundreds of nuclei and thousands of nuclear reactions. These explosions also involve complex hydrodynamic phenomena taking place in degenerate matter and strong gravitational fields (rendering terrestrial experiments of limited utility). Buoyancy-driven turbulent nuclear combustion during the deflagration phase dominates the early part of the explosion and drives an expansion and pulsation of the star. A deflagration-to-detonation transition (DDT) and propagation of the resulting detonation wave through the star has been posited to explain the observed nucleosynthesis and its distribution in space and velocity[45]. In the alternative gravitationally confined detonation (GCD) model, fluid flow triggers a detonation that sweeps through the star, producing the observed abundances, spatial distribution, and velocities of the elements. All of this takes place in around 3 seconds, followed by rapid free expansion of the star at velocities of 10,000 – 25,000 km s^{-1}. These phenomena involve spatial scales from ~10^{-3} cm – 10^{9} cm and temporal scales from ~10^{-10} s – 10 s, making simulations of SNe Ia a manifestly exascale problem. Advances are needed in both the speed at which the problem can be addressed and the scale (physical size) of the system that can be handled.

Several key physical processes in SNe Ia are not fully understood, and consequently the understanding of the explosion mechanism is uncertain. These physical processes include the smoldering phase, which precedes the explosion phase and is thought to determine the number of points where ignitions occur and their location(s). The buoyancy-driven, turbulent nuclear combustion phase—or deflagration phase, which releases nuclear energy and expands the star—also represents a frontier, as the understanding of reactive turbulence in strong gravity is incomplete. Finally, the origin of the detonation wave that incinerates the star and causes it to explode is uncertain. Whether the physical conditions necessary for a DDT are achieved in the deflagration phase of SNe Ia is unclear. The alternative, in which fluid flow during the deflagration phase triggers the detonation, is not fully understood.

Exascale computing, combined with exascale memory, will produce breakthroughs in the understanding of these physical processes, transforming scientists' ability to simulate SNe Ia. It will enable the qualitative improvement of scientists' understanding of the smoldering phase, thereby reducing the uncertainty in the initial conditions for simulations of the explosion phase. It will make possible studies of buoyancy-driven turbulent nuclear combustions—which include capturing this physical process by simulations that resolve length scales only 3-4 decades below the largest physical scales and that use a self-similar subgrid model[46] if needed—that could verify current expectations. If these studies do not verify these expectations, exascale computing will determine that the process is more complicated and

[45] Nomoto K, F-K Thielemann, & K Yokoi. 1984. "Accreting White Dwarf Models for Type I Supernovae. III. Carbon Deflagration Supernovae." *ApJ* **286**(2, Part 1): 644-658; Khokhlov AM. 1991. "Delayed Detonation Model for Type IA Supernovae." *A&A* **245**(1): 114-128; Gamezo VN, AM Khokhlov, & ES Oran. 2005. "Three-Dimensional Delayed-Detonation Model of Type Ia Supernovae." *ApJ* **623**: 337-346.

[46] Khokhlov AM. 1995. "Propagation of Turbulent Flames in Supernovae." *ApJ* **449**(2, Part 1): 695-713; Zhang J, OEB Messer, AM Khokhlov, & T Plewa. 2007. "On the Evolution of Thermonuclear Flames on Large Scales." *ApJ* **656**(1): 347-365.

provide the data needed to construct an appropriate subgrid model. Finally, exascale computing will also make possible studies that verify whether buoyancy-driven turbulent nuclear burning in a white dwarf star produces the physical conditions needed for a DDT to happen.

Advances in SNe Ia modeling during the next decade will most likely come from a combination of high-resolution simulations of the key physical processes described above and whole-star simulations of SNe Ia. Sustained petascale computing will enable verification studies of buoyancy-driven turbulent nuclear burning that will dramatically improve the understanding of this key physical process and will make possible whole-star SNe Ia simulations to treat buoyancy-driven turbulent nuclear combustion over a larger range of scales, providing new insights into the energy cascade and instabilities produced by this physical process. With sustained exascale computing, it may be possible to attempt first-principle simulations of SNe Ia from ignition through the deflagration phase (i.e., the buoyancy-driven turbulent nuclear burning phase), a difficult problem.

A key component of studies at both the petascale and the exascale will be global validation of the models using large numbers of SNe Ia simulations. The need to perform large ensembles of simulations means the average time to perform simulations of adequate resolution will have to be reasonably short to allow for several such simulations to be performed in a given real time. Thus, a careful mix of a few high-fidelity and many low-fidelity simulations will be required. Even so, it means that high-capacity as well as high-capability exascale computing platforms will be needed.

Achieving the promise of exascale computing for SNe Ia simulations presents several challenges:

- SNe Ia simulation codes need to exhibit strong scaling and run efficiently on platforms with millions of cores and/or that exploit accelerators. Weak scaling will be insufficient because the computational demand scales as the fourth power of the resolution.
- SNe Ia simulations, in common with core-collapse supernova and stellar evolution simulations require many physical variables per computational cell (e.g., fluid variables, flame variables, nuclear species variables, and radiation transport variables). Thus, the smaller memory per core of future platforms will require the development of new algorithms for efficient domain decomposition and load balancing.
- New parallel I/O algorithms need to be developed. These include those that can handle files of many terabytes and beyond, along with mass stores that can accommodate exabytes of data. The turbulent nature of the deflagration phase demands high-temporal resolution in the retained data sets. This leads to the production of remarkable data volumes (easily many petabytes and even exabytes).
- New algorithms for scientific data analysis, including visualization, need to be developed to handle petabytes and exabytes of data, along with data archiving techniques that can deal with up to exabytes of data and allow for comparative analyses to be performed between huge data sets.

The major scientific outcomes of SNe Ia simulations at the exascale will be as follows: (1) ascertaining the explosion mechanism, (2) calibrating SNe Ia as standard candles to an accuracy sufficient to study quantitatively the behavior of dark energy with redshift (i.e., with the age of the universe), and (3) understanding the contribution of SNe Ia to nucleosynthesis.

Understanding the explosion mechanism will also impact ideas about the interaction of reactive flow and turbulence. The deflagration phase is ultimately a straightforward problem in combustion, trading many of the complications of terrestrial burning (e.g., geometry of devices, unmixed fuels, soot production, etc.) for far more fundamental ones (e.g., extremely strong gravity, huge Reynolds numbers, and remarkably stiff reaction kinetics). As such, SNe Ia simulations represent unique numerical laboratories in which to explore basic ideas in reactive turbulent flow. The production of realistic SNe Ia simulations will require advances in this basic area.

An understanding of the explosion mechanism will make possible simulations that can predict correlations among the observed properties of SNe Ia. This will allow them to be better calibrated as standard candles, enabling them to be used to study quantitatively the behavior of dark energy with redshift, and thus to have a strong impact on scientists' understanding of dark energy.

SNe Ia simulations also predict the nucleosynthetic yields for various elements and isotopes, yields that can be tested by observations. These yields are intimately connected with the physical processes that occur during the explosion phase. Consequently, comparisons of nucleosynthetic predictions with observations provide indirect information on these processes, and therefore on the explosion mechanism.

With carbon and oxygen burning being followed by silicon burning and, in the deep interior, an extended period in nuclear statistical equilibrium, SNe Ia simulations are voracious consumers of the nuclear data that govern these burning processes, including binding energies; partition functions; and strong, electromagnetic, and weak interaction reaction rates[47]. Important reactions, such as the burning of carbon to produce oxygen [$^{12}C(\alpha,\gamma)^{16}O$], are the target of ongoing efforts to better measure their reaction rates. Of particular importance are the weak interaction rates for isotopes of iron peak elements, which determine the neutron richness of the simulated ejecta. These continued improvements in the nuclear data improve the nucleosynthetic predictions from SNe Ia simulations, thereby strengthening the constraints on them that are imposed by observations of their ejecta and of solar abundances.

4.2.2 Core-Collapse Supernovae

Iron-core collapse and bounce are governed by the interplay of general relativistic gravity with the weak and strong nuclear interactions at extremes of neutron richness and density (e.g., $> 10^{14}$ g/cm^3). The subsequent evolution of the event involves neutrino radiation hydrodynamics and nuclear kinetics among other physical processes. The experimental fact of nonzero neutrino masses means scientists must ultimately solve a macroscopic-scale problem in quantum kinetics as well, directly computing the dynamic, nucleosynthetic, and other observational consequences of flavor oscillations in situ as part of fully integrated simulations.

There are profound consequences to this complexity. These simulations make use of a variety of computational algorithms and implementations and stress essentially all facets of a modern, general-purpose computer — the input/output (I/O), memory size and latency, processor performance, communication bandwidth and latency, and more — in a manner shared with only a handful of other computational problems. These simulations will stress all facets of a general-purpose supercomputer.

The ability to simulate core-collapse supernovae realistically will depend on the development of discrete representations of the underlying nonlinear partial differential and integro-partial differential equations governing their evolution. This will require efficient and scalable-solution algorithms of the resultant nonlinear algebraic equations, as well as computer codes based on these solution algorithms that can take advantage of the memory and central processing unit (CPU) capabilities of petascale to exascale architectures. Advances in each of these areas will be required, along with considerable work devoted to enhancements of the computational ecosystem surrounding these machines. Core-collapse supernova codes produce prodigious volumes of simulation data over long wall clock times. Efficiently writing these data and managing and analyzing them after they are written are as important to producing meaningful science through supernova simulation as is any algorithmic or implementation improvement that might be made for the computational step itself.

[47] See, e.g., Calder AC, et al. 2007. "Capturing the Fire: Flame Energetics and Neutronization for Type Ia Supernova Simulations." *ApJ* **656**(1): 313-332; Seitenzahl IR, DM Townsley, F Peng, & JW Truran. 2009. "Nuclear Statistical Equilibrium for Type Ia Supernova Simulations." *Atomic Data and Nuclear Data Tables* **95**(1): 96-114.

Using current petascale platforms and their immediate successors, scientists may be able to determine the general nature of the explosion mechanism itself by performing three-dimensional radiation-magnetohydrodynamics simulations with spectral neutrino transport. As machines capable of peak speeds of 100 petaflops emerge, significant quantitative statements concerning the details of explosive nucleosynthesis in the event and the neutrino emission can be expected. At the exascale, scientists will finally be able to determine precisely how supernovae explode by undertaking transformative numerical experiments that incorporate quantum kinetics on macroscopic scales with nuclear physics components realistic enough to accurately predict the isotopic output of these events. These kinds of simulations are utterly unimaginable on current platforms but promise to be accessible at the exascale. This is truly applying quantum mechanics, a theory of the smallest things known, to some of the most "macroscopic" bodies in the universe.

The multiphysics nature of core-collapse simulations will require new computational techniques ranging from scalable linear algebra to methods to solve coupled ordinary differential equations (ODEs). The high number of degrees of freedom at each spatial grid point (e.g., neutrino flavors, energies, and angles, as well as nuclear species) currently represents a large amount of unrealized parallelism in modern supernova codes. Methods to handle these calculations concurrently on multicore platforms and platforms incorporating accelerators of various kinds will likely determine the efficacy of future codes.

Several of the major questions posed in the 2007 Nuclear Science Long Range Plan (DOE 2007) are germane to core-collapse supernova simulation:

What are the phases of strongly interacting matter, and what roles do they play in the cosmos?

What is the nature of neutron stars and dense nuclear matter?

The nature of dense nuclear matter formed at the center of a supernova explosion provides a unique opportunity to explore the low-temperature, high-density region of the QCD phase diagram. Knowledge obtained from observation and simulation in this region will complement the better-studied, high-temperature (e.g., quark-gluon plasma [QGP]) regions of the phase diagram, which are presently accessible to terrestrial experiment.

What is the origin of the elements in the cosmos?

One of the most important and distinctive observables from core-collapse supernovae is their pattern of nucleosynthesis. The creation and transmutation of a wide variety of intermediate- and high-mass species in the event is a nonlinear phenomenon. Supernova nucleosynthesis has a dynamic effect on the explosion mechanism, ultimately rendering post-processing of simulation results to be of only qualified utility. The subsequent dissemination of the produced species enriches the interstellar medium, setting the stage for successive generations of star formation and death.

Nuclear physics experiments at the Facility for Rare Isotope Beams (FRIB) will constrain temperature, density, timescales, and neutrino fluxes at the r-process nucleosynthesis site from observations of elemental abundances[48]. Simulations of core-collapse supernovae will be the essential ingredients in tying these experimental measurements to the astrophysical site of the r-process, because a self-consistent determination of all of these conditions can only be achieved through computation at scales beyond those currently possible.

What is the nature of the neutrinos, what are their masses, and how have they shaped the evolution of the universe?

Core-collapse supernovae are, from an energetics point of view, neutrino events. They represent the only instance in the modern universe where neutrino interactions have a discernible, macroscopic effect on the dynamics of baryonic matter. Spectral neutrino transport is required to accurately model the event, and

[48] RIA Working Group 2006

the resulting neutrino templates will be invaluable in interpreting and calibrating detections in terrestrial experiments. Comparing future observations to simulation results will be vital to interpreting those observations and using them to constrain the properties of neutrinos.

Accurate and precise knowledge of the characteristics of neutron-rich matter at high density is a prerequisite for understanding core-collapse supernovae. Precise data for electron-capture processes on progressively larger nuclei is a fundamental need for the simulations, a need that can only be filled by advances in nuclear structure theory. Conversely, core-collapse supernova simulations provide the crucial link in testing these theoretical results, as it is only at the extremes of density and neutron richness realized in these simulations where these predictions are manifest. As nuclei in the collapsing core make the transition from an ensemble of nuclei to nuclear matter, exotic forms of matter are expected[49]. The details of this transition region are of considerable importance in determining accurate neutrino spectra, again providing a unique link between fundamental theory and physical observables.

In addition, core-collapse supernovae are prodigious sources of gravitational waves (GW)[50]. Because the signal-to-noise ratio for GW detectors presents a serious complication for detection, the production of useful templates for detectors like the Laser Interferometer Gravitational-Wave Observatory (LIGO) and VIRGO is essential for meaningful data analysis. Furthermore, as nonaxisymmetric oscillations are required for GW production, multidimensional, fully integrated simulations are the only path to producing these signal templates. Therefore, the only path forward to interpreting possible future GW wave detections from core-collapse supernovae relies wholly on simulations providing the requisite context.

4.3 Climate Modeling

Although substantial uncertainty exists as to the degree and impacts of future climate change, especially at local and regional scales, it is generally agreed that significant adaptation will be required. Furthermore, the magnitude of climate change later in the century depends upon the near- and intermediate-term mitigation strategies used to reduce the emission of greenhouse gases. These strategies also must satisfy an increasing energy demand of a growing global population experiencing an improvement in its standard of living. Predicting these future climate changes and evaluating the effects of mitigation strategies require Earth system models (ESMs) that are far more accurate and comprehensive than those in use today. Integrated assessment models provide the framework for climate predictions by defining the emissions scenarios and elucidating the relationships among the natural and human systems that are at the core of climate change studies. In the next decade, integrated assessment and comprehensive ESMs will probably be combined into a single system that could be used to investigate scientific issues and to formulate policy options for adaptation and mitigation.

The predictions from integrated ESMs will be most credible if the important processes in the climate system, for example mixing by ocean eddies, are simulated at their native spatial and temporal scales. Critical organized features in the atmosphere and ocean including clouds and eddies have characteristic sizes of 1 to 10 km. Some of the major sources of uncertainty in climate predictions from existing models are associated with the aggregate effects of these phenomena. Experience with current climate models suggests that simulation of climate change with a model with 10-km grid resolution is inherently a petascale problem. In fact, even higher resolution is required to resolve these features with sufficient fidelity to the physical principles underlying their formation and evolution. Since the computational cost increases nonlinearly with higher resolution, it is likely that predictions of societal and environmental change at 1-km resolution would require truly extreme scale computers.

---

[49] Ravenhall DG, CJ Pethick, & JR Wilson. 1983. "Structure of Matter Below Nuclear Saturation Density." *PRL* **50(**26): 2066-2069.

[50] Ott CD. 2009. "Topical Review: The Gravitational-Wave Signature of Core-Collapse Supernovae." *Classical and Quantum Gravity* **26**(6): Art. No. 063001. DOI:10.1088/0264-9381/26/6/063001. Abstract accessed May 10, 2009, at http://www.iop.org/EJ/abstract/0264-9381/26/6/063001.

4.4 Combustion

The United States is at a critical juncture where urgent issues in energy security, climate change, and economic competitiveness are converging. Aggressive national goals for reducing petroleum use by 25 percent by 2020 and greenhouse gas emissions by 80 percent by 2050 will require major improvements in all aspects of our nation's energy use. At the same time, the U.S. transportation industry is under tremendous pressure from international competitors and challenging economic conditions. Achieving usage and emission goals will require the automobile and truck sector to significantly shorten its product development cycle for cleaner, more efficient engine technologies. The transportation sector alone accounts for two-thirds of the nation's oil use and one-quarter of the nation's greenhouse gas emissions. Concurrently, fuels will also be evolving, adding another layer of complexity and further underscoring the need for efficient product development cycles. However, under the current Edisonian cut-and-try approach, design cycles and market penetration of new engine technologies simply takes too long. These challenges present a unique opportunity to marshal U.S. leadership in supercomputing to develop predictive combustion science and engineering simulation tools used by both the basic and applied DOE Offices – Office of Basic Energy Sciences (BES), the Energy Efficiency and Renewable Energy (EERE), the Office of Vehicle Technologies (OVT), and the Combustion Energy Frontier Research Center (CEFRC), working in close connection with the U.S. transportation and fuels industries. Simulations that reliably predict efficiency and pollutant emission for new engines and new fuels require higher fidelity modeling than is possible with current computing resources or their straightforward evolution.

Reliable prediction requires, for example, the incorporation of heterogeneous kinetics with quantified uncertainties in turbulent combustion simulations for processes such as soot formation/burnout and increased fidelity coupling of high-pressure, low-temperature chemistry with turbulent transport – and these vital enhanced modeling techniques will only be feasible at exascale computing performance levels.

In particular, combustion scientists must focus on the science underlying the development of non-petroleum-based fuels, including carbon-neutral biofuels, and their optimal use in transportation. This science intrinsically involves chemistry with transport at conditions far from equilibrium and at extreme pressures and a coordinated multi-scale approach for understanding and predicting combustion in turbulent environments.

Combustion in practical devices covers a myriad of time and length scales, from the scale of the electron to those of the largest scales of turbulence dependent upon the geometry of the device. To tackle this daunting challenge and complexity, a multi-scale approach is adopted wherein experiments, theory and direct computation are brought to bear on a limited range of scales (4-5 decades) and fundamental physical insights gained are encapsulated in reduced-order parameterizations that are used to upscale knowledge to bridge the scales. Several high-fidelity computational approaches in both the atomistic and continuum regimes utilize petascale computing. Exascale computing would greatly facilitate higher fidelity or access to more practically relevant parameter regimes (e.g., higher pressure, higher turbulence levels, and more complex fuels).

In the continuum regime where turbulence scales interact with flame, ignition, and mixing scales turbulence-chemistry interactions are important. Virtually all combustion devices operate under turbulent environments due to enhanced mixing and greater efficiency. Many of the fundamental turbulence-chemistry interactions are amenable to investigation by first principles direct numerical simulation (DNS) and high-fidelity large-eddy simulation (LES) of building block, laboratory scale flows. Whereas DNS focuses on the fully resolving the fine-grained physics, LES resolves the energy-containing end of the turbulence spectrum down to a specified cut-off in the inertial or dissipative end of the spectrum and the unresolved subgrid scales are modeled. As such these methods are complementary. Both DNS and LES require the horsepower of high-performance supercomputing at the exascale and beyond to resolve all relevant flow and chemical scales. Exascale simulations are required, for example, to understand the coupling between low-temperature ignition kinetics and turbulent mixing at high pressure that determines

lifted flame stabilization, ignition timing, rate of combustion, and emissions characteristics (e.g. sooting behavior). Understanding complex low-temperature high pressure kinetics of alternative fuels and its coupling with turbulent transport at high pressure requires much greater resolution and the transport of large numbers of reactive scalars only afforded by extreme scale computing power. Moreover, in-situ reduction strategies for accurate and computationally affordable inclusion of heterogeneous kinetics with quantified uncertainties in DNS and LES are required. The insights gained from exascale simulations will enable the development of predictive multi-scale models to optimally design future evolving fuels and engines.

Motivating Example: HCCI Combustion with Alternative Fuels – A Multi-scale Problem

Homogeneous charge compression ignition (HCCI) is one of several Low Temperature Combustion strategies being considered by automotive manufacturers. It has the potential to deliver high efficiencies similar to diesel, but with orders of magnitude smaller pollutant emissions. It achieves this by burning a highly homogeneous mixture with a low fuel concentration, avoiding the high fuel concentrations responsible for soot and NOx emissions. These mixtures are too fuel lean to propagate an ordinary flame; to work the fuel chemistry and conditions must be right for spontaneous combustion to occur in proper phase with the piston motion. This makes the engines much more difficult to control, and much more sensitive to fuel chemistry, than conventional engines. Many fundamental questions remain about low-temperature ignition chemistry, particularly for alternative fuels. For example, the properties of oxygenated hydrocarbon molecules in fuels biomass-derived fuels, the nitrogen-containing molecules in oil shale, and the many cyclic compounds in non-conventional fuels are poorly characterized, and even less is known about the ignition chemistry of fuel blends. To operate at high loads, some inhomogeneities are needed to moderate the rate of pressure increase. These lead to sequential ignition front propagation down the temperature gradients, producing combustion modes ranging from homogeneous explosion to unsteady deflagration, seen in Fig. 5 above.

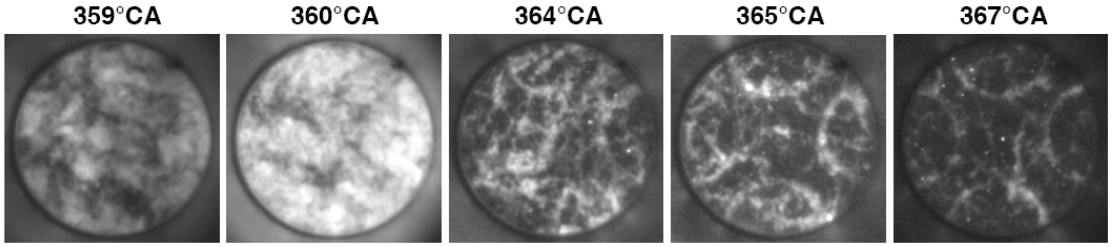

Figure 8: Chemiluminescence images of stratified sequential autoignition in an HCCI engine, courtesy John Dec, Sandia National Laboratories.

These new modes operate far from chemical equilibrium – with highly transient, intermittent ignition occurring at multiple sites, and extinction in some regions induced by turbulent strain or by encountering compositions beyond the flammability limits. At the low pre-ignition temperature, mixing can compete with chemical reactions, resulting in strong chemistry-turbulence interactions. These complexities pose challenges, but also opportunities, as the strong sensitivities to both the fuel chemistry and to the fluid flows provide multiple control options. Theoretical models, experimentation, and first-principles simulations from electronic structure at the atomic/molecular scale to direct numerical simulation at the continuum scales are required to predict the behavior of alternative fuels in HCCI and other novel combustion devices, and so facilitate design of an optimal combined engine-fuel system. The first principle simulations require exascale computing horsepower to resolve the chemical and flow scales at high pressure.

At the forefront of advanced internal combustion engine concepts are low temperature, high pressure, compression ignited engines. This is the only engine strategy that can meet future government-mandated

standards on decreased NO_x/soot[51] emissions and increased gasoline mileage. However, prototype engines of this design are hard to control because of sensitivity to cycle-to-cycle variations within the cylinder. This sensitivity comes from the low temperature and the spark-less compression ignition, both of which are dependent on turbulent mixing of fuel and oxidizer and the details of the fuel chemistry. Simulation currently can play only a minor role in influencing engine design because the coupling of turbulence with detailed fuel chemistry over multiple engine cycles exceeds both our scientific understanding and our computational resources. Computing at the exascale, along with increased scientific investments, offers the first opportunity where turbulence with sophisticated chemistry can finally be applied to real world problems in combustion.

For example, pioneering multi-cycle simulations of one particular low-temperature, high-pressure engine design, the thermally stratified HCCI engine could be performed on an exascale computer. LES simulations of the full engine geometry, validated with related engine measurements, would be able to capture the essential aspects of thermal stratification and the mechanisms responsible for producing it. DNS simulations of turbulent combustion at the micro-scale in an HCCI environment can reveal the mechanisms controlling spontaneous autoignition of mixture pockets subjected to varying degrees of scalar inhomogeneities and provide unique benchmark data for developing and testing subgrid combustion models that capture multi-regime ignition and flame phenomena at high pressure. Adaptive and heterogeneous kinetics would enable the practical inclusion of detailed fuel chemistry by automatically adjusting to the stochastic chemical variability of the mixture pockets. Using the combination of LES, DNS, and detailed chemistry, it would be possible to develop both diagnostic and predictive criteria that characterize the heat release modes of the engine as a function of the level of stratification and turbulence parameters. The validated ability to provide both scientific understanding and engineering information would have a major impact on the development of engines and fuels for the future.

In summary, future predictive simulation tools running on exascale computing systems will enable deep understanding of underlying chemical and combustion science processes, enhance combustion engine design and performance, and ultimately yield a dramatic reduction in engine development timescales, time to market, and development costs, while ensuring the timely achievement of energy security and emissions goals, and enhancing the competitiveness of U.S. engine manufacturers and fuel producers.

4.5 Computational Biology

The ultimate goal of exascale computing applications to challenges in modern biology is to go from atoms to organs or from microbes to ecosystems: for example, to enable an understanding of how the brain works as an energy efficient, biologically-based information system, or to understand microbial processes and their impact on the geosphere. In the process, these newly enlarged scales of computing will resolve unfathomably complex research issues in a host of fields as diverse as neuroscience and microbial metagenomics.

At exascale, new scalable tools that admit a variety of time, space and trajectory sampling methods (and fully exploit the hundreds of millions of cores of an exascale machine) will enable long time integrations, implicit solvation conditions, and mixed molecular mechanics and quantum mechanics models, to allow *breakthrough science*. For example, a large biochemical network within a full-scale model of a eukaryotic cell could be modeled in the span of a few hours.

It is important to note that the first million-atom simulation in biology was conducted just five years ago – an all-atom simulation of the ribosome conducted at Los Alamos National Laboratory. This million-particle simulation milestone had already been achieved a decade prior in materials science and

[51] NO_x is a short-handed way of referring to oxides of nitrogen, such as NO (nitric oxide) and NO_2 (nitrogen dioxide), especially in the context of combustion.

cosmology (computational scientists in both these fields now perform multibillion-particle simulations). While biology researchers have achieved impressive methodological advances that permit the modeling of the largest assemblies in the cell, it is only for short periods of time. And, these simulations are unlikely to scale to the size of a single cell, *even a small bacterium*, for relevant times such as minutes or hours – even if researchers can employ computers capable of achieving 1,000 petaflops/s.

Today, researchers are currently limited to the microsecond timescale for protein folding required by the huge number of intermolecular, interaction computations. Scientists also lack rigorous coarse-grained models that permit the scaling up of macromolecular pathways and supramolecular cellular processes. Similarly, systems biology methods lack the dynamic resolution needed for coupling genomic and other data in order to fully map cellular networks, to predict their functional states, and to control the time-varying responses of living cells. Nor can current kinetics models adequately analyze the dynamics of complex living systems. Exascale computing will be needed to achieve those capabilities.

Within the next decade, scientists expect to have the complete genome sequence of more than 10,000 bacteria and archaea and other single-celled microbes. Exascale computing platforms will make it possible in principle to systematically reconstruct the metabolic networks of all sequenced microbes through automated comparative analysis, to reconstruct their regulatory networks by integrating a variety of data sources, and to combine these reconstructions into functional models of cellular states. Exascale computing will be critical to make this a routine class of computation such that it can become part of the standard way we analyze genomes in the future.

Metabolic model reconstruction, is something that is now possible to achieve using today's computational capabilities. It is also a prerequisite for any meaningful reconstruction of transcriptional regulatory networks: the metabolic model is the scaffold on which scientists can assemble and implement regulatory networks. With exascale computing it will become possible to use metabolic models as key computational constraints to drive the self-consistent annotation and curation of all sequenced genomes.

Another related area likely to benefit from exascale is the prediction of feasible *parameter values for dynamic models of metabolism* that would enable scientists to design organisms that would perform a variety of tasks. These models might also contribute to the development of treatments for emerging types of infections.

Understanding macromolecular complexes is crucial to drug discoveries that counter certain biothreats as well as advancing bioenergy applications. Exascale opportunities in macromolecular proteins and protein complexes – the "nonmachines" involved in virtually all elementary processes of life – fall within four areas of biological systems and problems in terms of their increasing complexity of temporal and spatial dimensions. These opportunities are:

1. Macromolecular folding, including protein folding and RNA folding
2. Biochemical binding and reaction mechanisms, such as enzyme catalysis and protein/ligand interactions
3. Macromolecular pathways, including DNA replication and repair fidelity, protein synthesis, chromatin organization, and RNA editing
4. Supramolecular cellular processes, such as protein signaling networks, plant cell-wall formation, and endocytosis

Building multi-scale whole cell models for cellular populations of 10 billion cells or greater – a distinctly exascale computational problem – would build up from the abstract molecular networks, via molecular interaction models, through course graining of these models to mesoscale process models of the major elements of the cell.

Surveying the challenges presented by population, communities, ecosystems, and evolutional dynamics (genomics and metagenomics) emphasizes on understanding bacteria as well as microbial cells with a

nucleus or other organelles in them, for example, eukaryotes. Ultimately this research will inform our understanding of ecosystems and how humans are integrated into these systems.

Here, four primary challenges confront the biology community. These are: developing integrated ecological observatories and modeling them; modeling microbial macro- and microevolutionary dynamics and processes; modeling and engineering complex multispecies biological systems for the lab environment; and developing accurate analysis and annotation of genome and metagenome sequences.

New modeling innovations moving this field ahead can be divided into two categories: *ecological* and *computational* drivers. Ecologically, scientists need to comprehend the relevant measures of the organisms – genomic, metabolic, and proteomic. Computational drivers are needed to create better models of organisms that react to changes in the local environment, including complex models based on the genomic information in an ecological system. These models should include a linkage between chemical, environmental, geographical, and physical data.

Coupling models of terrestrial and marine microbial communities with the exchanges in the geosphere and with relevant atmospheric and oceanographic ecosystem processes is a multi-disciplinary, multi-scale problem that also requires advances in the underlying scientific understanding. However, the construction of integrated models is critical to exploring parameters and advancing our understanding of global climate, and the global carbon, nitrogen and phosphorus cycles.

The second major challenge here is modeling microbial macro- and microevolutionary dynamics and processes that can enhance scientists' understanding of how diversity in the biosphere reflects evolutionary processes. While engineering systems require a working knowledge of our actions, biological systems are self-replicating. Thus, researchers cannot model biological systems using the same techniques used to model chemical and physical systems, such as partial differential equations.

Coupling models of terrestrial and marine microbial communities with the exchanges in the geosphere and with relevant atmospheric and oceanographic ecosystem processes is a multi-disciplinary, multi-scale problem that also requires advances in the underlying scientific understanding. However, the construction of integrated models is critical to exploring parameters and advancing our understanding of global climate, and the global carbon, nitrogen and phosphorus cycles.

The second major challenge here is modeling microbial macro- and microevolutionary dynamics and processes that can enhance scientists' understanding of how diversity in the biosphere reflects evolutionary processes. While engineering systems require a working knowledge of our actions, biological systems are self-replicating. Thus, researchers cannot model biological systems using the same techniques used to model chemical and physical systems, such as partial differential equations.

A third significant challenge is the modeling and engineering of complex multispecies biological systems for the lab and environment. This effort is likely to involve bioreactor dynamics, host-viral interactions, host-predator processes, and process manipulation, and will also include an associated evaluation of theories, such as community assembly, functional redundancy, and parameter sweeps.

The fourth challenge is developing accurate annotation and analysis of genes in genomes and metagenomes. A number of sequence similarity-based tools perform functional assignments in current annotation pipelines. Consequently, the functions of new proteins are extrapolations of existing annotations. As a result, there are two interrelated parts to the problem of accurate gene annotation: accurately annotating a set of genes or genomes and extending the annotation to new genes, genomes, and metagenomes.

Exascale computing capabilities open up considerable possibilities for addressing the complexity of biological systems and offer the promise of moving towards a predictive science of systems biology. However, to leverage this capability the biological sciences community will need to embark on a push for next generation codes and tools explicitly targeting the requirements of exascale systems. A key

Motivating Example: Large-scale Simulation of Ion Channels

Voltage-gated ion channels, or Kv channels, are involved in the generation and spread of electrical signals in neurons, muscle, and other excitable cells. In order to open the gate of a channel, the electric field across the cellular membrane acts on specific charged amino acids that are strategically placed in the protein in a region called the voltage sensor. In humans, malfunction of these proteins, sometimes owing to the misbehavior of only a few atoms, can result in neurological diseases. A wealth of experimental data exists from a wide range of approaches, but its interpretation is complex. One must ultimately be able to visualize atom-by-atom how these tiny mechanical devices move and change their shape as a function of time while they perform. Researchers are using a tight integration of experiment, modeling, and simulation to gain insights into Kv channels. Their studies serve as a roadmap for simulating, visualizing, and elucidating the inner workings of these nanoscale molecular machines. Because these channels are functional electromechanical devices, they could be used in the design of artificial switches in various nanotechnologies. The practical applications of this work are significant. For example, the research in ion channel mechanisms may help identify strategies for treating cardiovascular disorders such as long-QT syndrome, which causes irregular heart rhythms and is associated with more than 3,000 sudden deaths each year in children and young adults in the United States. Moreover, the studies may help researchers find a way to switch or block the action of toxins – such as those emitted by scorpions and bees – that plug the ion channel pores in humans.

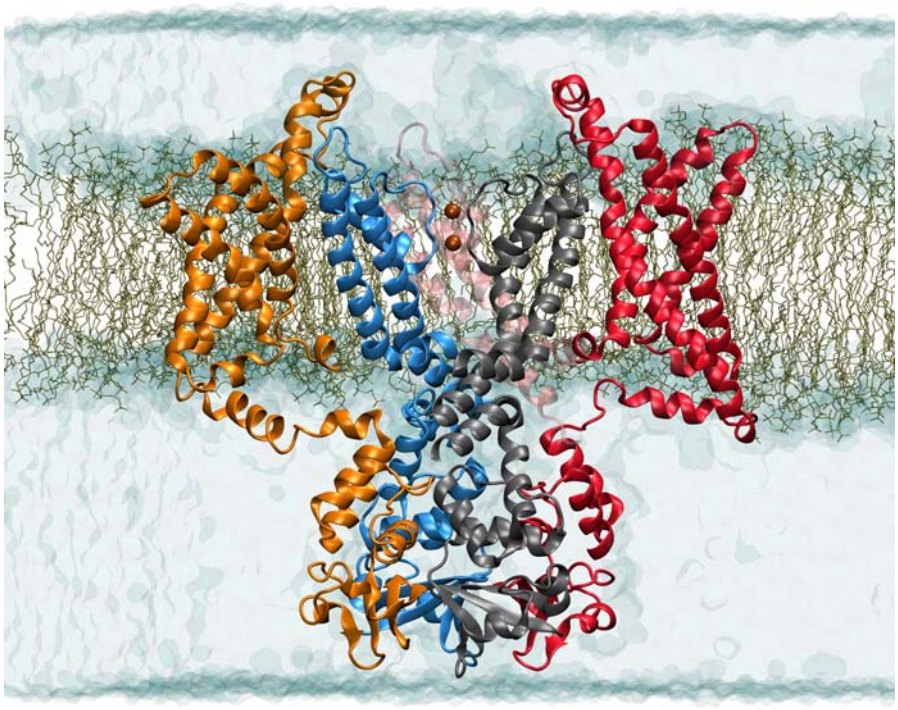

Figure 9: Complete model of the Kv1.2 channel assembled using the Rosetta method. The atomic model comprises 1,560 amino acids, 645 lipid molecules, 80,850 water molecules and ~300K+ and Cl- ion pairs. In total, there are more than 350,000 atoms in the system. The simulations were generated by using NAMD on the Cray X-T (Jaguar) at Oak Ridge National Laboratory and the Blue Gene/P at Argonne National Laboratory. Image courtesy of Benoit Roux, Argonne National Laboratory and University of Chicago

observation is that exascale computing capabilities will for the first time enable a truly multiscale attack on biological problems, coupling molecular level details to cellular networks, to populations and communities and coupling multiple communities and ecosystems to the global biosphere.

4.6 Materials Science

Materials innovations are central to many of the technological advances responsible for the quality of life and prosperity that Americans currently enjoy. In fact, many of the disruptive technological advances since the turn of the last century – modern transportation, medical treatments and prosthetics, space exploration, global communication, computers and the electronics industry – used advances arising from every corner of the materials world: metals, ceramics, semiconductors, polymers, and novel combinations of these. Materials establish and support entire industries, and tens of millions of manufacturing jobs depend on the availability of these advanced materials at affordable costs. The ability of the U.S. to not only develop new materials but transition them into the design and delivery of new products has translated into a decades-long competitive advantage in the global economy[52].

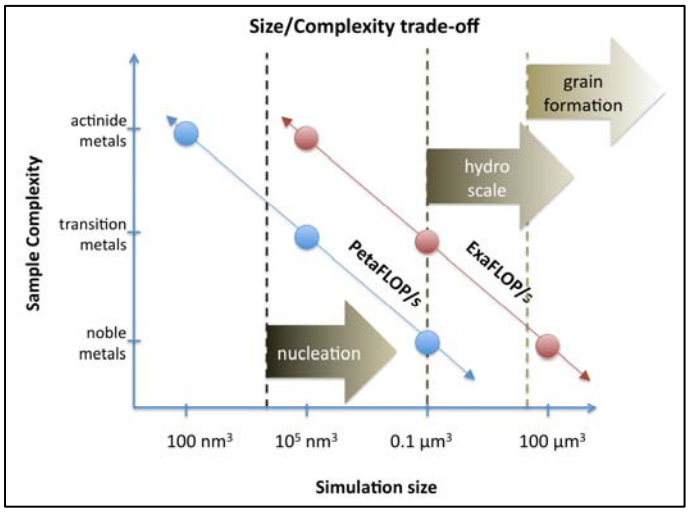

Figure 10: Illustrating the trade-offs made in a molecular dynamics simulation between model complexity (represented here by the choice of material) and simulation size for a given run time. Also shown are size ranges needed to model certain phenomenon, and the practical limit of a hydrodynamics simulation.

A quantifiable understanding of novel materials and their response is central as well to the technological challenges facing our country. Whether it is ceramics for high-efficiency automobiles, photovoltaics for next-generation solar power or smart alloys for efficient building construction, the nation requires the development of advanced materials with superior properties that will drive the next generation of technologies. In the highly competitive global marketplace that we find ourselves, minimizing time to solution and time to market is crucial. Computational modeling of subsystems (and entire systems) has been proven to dramatically shorten the time from concept to design, leading to well documented advances in the aeronautics, automotive and communications industries, to name just a few. These simulations are typically based on continuum equations, with materials defined through an equation of state and constitutive model. Currently, mesoscale heterogeneity (the defects, microstructure and multiple phases that would be present in a real material) is either ignored or treated in an ad hoc or averaged way. In many cases, the governing

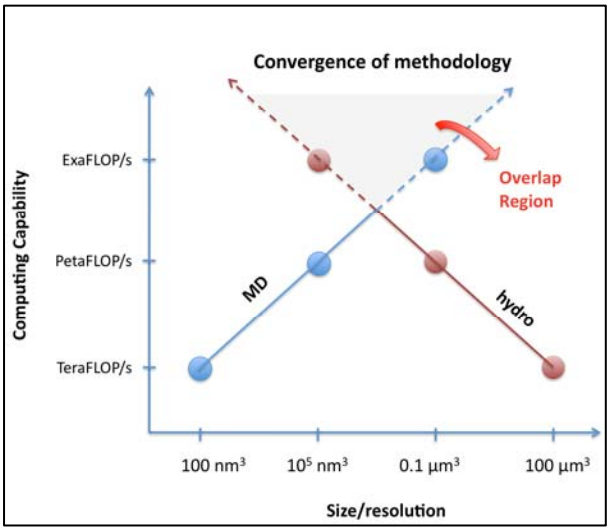

Figure 11: Schematic of the convergence with increasing computational capability between two common materials modeling techniques: molecular dynamics and hydrodynamics.

[52] As an example, consider Apple's iPhone – an iconic symbol of innovation and consumerism. From the use of amorphous metal technology to solder the engineered aluminosilicate glass face in the special alloy case to the miniature electronics and optics that fill the inside, this now-ubiquitous device would have been impossible to manufacture only a decade ago, and difficult to even imagine two decades ago.

constitutive equations are not actually known, and empirical models are derived from experimental data.

What is lacking is the ability to model real systems, taking into account the full range of structure and dynamics, including defects, surfaces, multiple phases and disorder. The very materials properties that we are attempting to improve, such as yield strength, electrical and thermal transport, fatigue and corrosion resistance, and magnetic hysteresis are greatly influenced by those defects. Nor is it sufficient in general to treat these effects in a mean-field way by averaging – it is often necessary to understand both the static and dynamic interaction between the defects as well as the impact of the defects on the properties of interest.

Why can we not do this with our current computers? It is instructive to consider two workhorse techniques for materials modeling – hydrodynamics and molecular dynamics – and examine the reasons why a simulation might fail to provide sufficiently useful information. Molecular dynamics simulations are characterized by a force field or potential, involving many adjustable parameters, which describes the interactions between atoms. There are no parameters required to describe the response of the materials, however – all the constitutive response emerges naturally from the interaction potentials. Such calculations are currently limited in size to fractions of a cubic micron simulated for 10's of nanoseconds, even on the largest computers. Hydrodynamics, by comparison, involves many adjustable parameters describing both interaction and the materials response. However, there is no real size or time limit in the simulation. There is a practical lower limit on resolution, as it makes no sense to model a atomically-sized region of space using continuum equations. At a given level of computing, computational scientists using either method encounter two common barriers to success: (a) the largest (or most finely resolved) simulation possible is still too small (or too poorly resolved) to capture the relevant behavior of interest, or (b) the most complex, compute-intensive simulation that can be solved in a reasonable time is still too simple or approximate to adequately describe the physics of interest. In many cases both (a) and (b) are true – which is particularly damning, since it prevents the investigator from performing the traditional trade-off between these two constraints: very often, one makes simplifying approximations to enable a larger simulation or investigates smaller systems in order to perform a more complicated calculation. For example, a simulator attempting to model nucleation using molecular dynamics on a petascale computer might choose to model molybdenum (a transition metal) rather than uranium (an actinide) (see Figure 10 above). On the other hand, investigating grain formation using molecular dynamics may not be possible, even in the simplest metals on today's computers.

The availability of an exascale platform will move the location of the constraints, allowing quite generally more detailed calculations of more complex materials. Of particular note, however, is that our current capability is tantalizingly close to enabling a direct overlap between these two important modeling methodologies (see Figure 11). State-of-the-art calculations involving billions of atoms have been performed that demonstrate the ability to model macroscopic (i.e., continuum) materials behavior with an atomistic model that makes no assumptions about the cooperative response (Figure 12). With the development of an exascale computer it is possible that such a calculation (which was heroic on a petascale computer) could be performed **on demand** during a hydrodynamics calculation, determining, for example, the equation of state for a mixed region at precisely the temperature, pressure and composition that was required. By tabulating this information as it is generated, one can envision

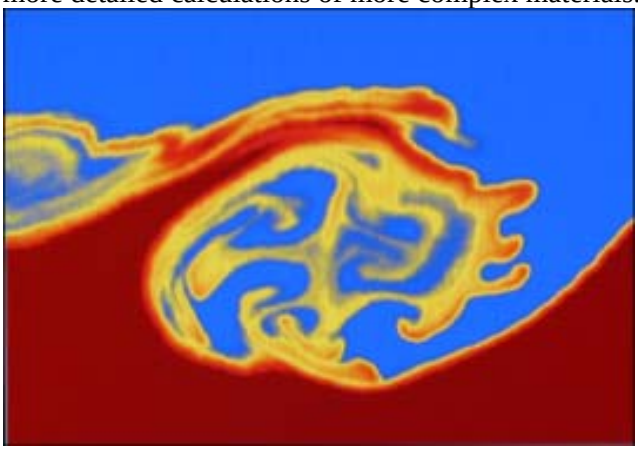

Figure 12: Detail from a 9-billion-atom molecular dynamics simulation of a developing Kelvin-Helmholtz instability at the sheared interface between aluminum and copper.

that such a simulation would teach itself as it runs, learning only those regions of this three dimensional phase space that is needed.

Such a direct coupling of modeling techniques is extremely powerful. In this example, it would result in the extension of classical molecular dynamics into continuum length and time scales. One readily imagines a similar direct multiscale approach being used to connect, e.g., Hartree-Fock calculations with molecular dynamics simulations of biological systems, or orbital-free density functional theory calculations with a phase field model.

An inadequate representation of mesoscale heterogeneities intrinsic to real materials is only one failing of current materials models. The effects of corrosion and other interfacial chemistry phenomena also cannot be predicatively modeled, nor can the effects of aging. The development of exascale resources will introduce a true paradigm shift in materials simulation–it will offer the opportunity to perform simulations on multiple concurrent scales, moving us substantially closer to a true materials modeling capability.

4.7 Fusion Energy

Over the next decade, a computational initiative called the Fusion Simulation Program (FSP), which will be led by the Office of Fusion Energy Sciences (OFES) in partnership with the Office of Advanced Scientific Computing Research (OASCR), is currently being developed with the primary objective being to produce a world-leading, experimentally-validated predictive simulation capability that is important to ITER and relevant to major current and planned toroidal fusion devices. This is expected to be a "transformational" capability that will demand the development over the next decade of advanced software designed to use leadership class computers (at the petascale and beyond) for carrying out unprecedented multi-scale physics simulations to provide information vital to delivering a realistic integrated fusion simulation modeling tool. Modules with much improved physics fidelity will enable integrated modeling of fusion plasmas in which the simultaneous interactions of multiple physical processes are treated in a self-consistent manner. The associated comprehensive modeling capability will be developed in close collaboration with experimental researchers and validated against experimental data from the major U.S. facilities and also internationally. Since each long-pulse shot in ITER is expected to cost over $1M, this new capability promises to be a most valuable tool for discharge scenario modeling and for the design of control techniques under burning plasma conditions.

Some specific examples of expected advances that are needed to enable a comprehensive integrated modeling capability include:

- The effective coupling of state-of-the-art codes for the plasma core and the plasma edge region.
- The effective coupling of state-of-the-art codes for MHD dynamics and auxiliary heating of the plasma via RF waves.
- The development of more realistic reduced models based on results obtained from the DNS-type (direct numerical simulation) major codes that use petascale capabilities.
- The development of advanced frameworks and workflow management needed for code coupling.

Some specific examples of outstanding challenges in the fusion energy science application area include:

- Efficient scaling of MHD codes beyond terascale levels to enable higher resolution simulations with associated greater physics fidelity.
- Efficient extension of global PIC codes into fully electromagnetic regimes to capture the fine-scale dynamics relevant not only to transport but also to help verify the physics fidelity of MHD codes in the long-mean-free-path regimes appropriate for fusion reactors.
- Mastery of data management to help with the development (including debugging) of advanced integrated codes.

- Development of innovative data analysis and visualization to deal with increasingly huge amount of data generated in simulations at the petascale and beyond.

If proper investments in research efforts are made, specific fusion energy science state-of-the-art codes (such as GTC, GTC-P, GTS, and XGC1) can be transformed to versions using a hybrid programming framework (see Section 5 below) and then systematically exercised to demonstrate and to test the effectiveness of this proposed paradigm. If such methods for optimally utilizing many-core systems prove to be truly effective, then the expected outcome would be that such codes can proceed to realize their high potential for achieving new scientific discoveries, first with multi-petascale and then exascale computing power.

Other examples of expected outcomes if dedicated efforts are properly supported include:

- Clear demonstration of the ability to effectively integrate (beginning with two-way coupling) advanced codes to deliver new physics insights.
- Significant progress on the ability to reliably predict the important Edge Localized Modes (ELM) behavior at the plasma periphery.
- Significant improvement in the physics fidelity of "full device modeling" of ITER along with significant progress in achieving reliable predictive capability for ITER.

The new Fusion Simulation Program (FSP) is anticipated to deal with many of these challenges; and we can already envisage that the associated research progress enabled by compute power at the extreme scale will demand much greater allocations of HPC time than has been available to date. For example, a single XGC1 production run carried out at present to provide new insights into nonlinearly-coupled core-edge plasma dynamics requires around 100K cores * 240 hours = 24M CPU hours. If additional dynamics (such as the modeling of the RF auxiliary heating) were also included, then the need for computational resources at the exascale would be a reasonable expectation. In assessing the promise and importance of the FSP, U.S. Energy Secretary Steven Chu as recently stated in an "all hands meeting" (September 27, 2010) at the Princeton Plasma Physics Laboratory: "The world's energy challenge requires a strong continued commitment to plasma and fusion science. … Progress in fusion has to be grounded in validated predictive understanding: the DoE is clearly interested in your planning and progress for a strong Fusion Simulation Program (FSP)."

What risks arise if the dedicated investments described in the preceding are not made? Most obviously, the key risk is that the ability of the U.S. to maintain a world leadership role in predictive simulations of fusion energy systems would be damaged. This is a critical issue in that the U.S. can be viewed as being "facilities poor" in that major international investments in fusion facilities – including ITER and dynamic new programs in China and Korea – have moved past the U.S., where no new facilities have been built in the 21st century. Primacy in validated predictive fusion simulation capability can in fact be maintained as the U.S. leverages its strengths in HPC facilities, considerable experience in interdisciplinary computational science via programs such as SciDAC, and possible new programs for co-design at the exascale.

At the fusion energy HPC operational level, leading fusion codes (especially the MHD codes with their formidable scaling challenges) would run the risk of not being able to effectively utilize the large number of processors at the exascale. As discussed in detail in Section 5 below, the development of effective mathematical algorithms for integration/coupling is very difficult and could prove difficult to achieve within the next decade without the needed support today. If the fusion energy science applications were only able to effectively utilize a small fraction of the cores on a CPU, the U.S. could fall behind international efforts in developing innovative new methods for per processor performance. Without a position of strength in either advanced simulation capability or in a vibrant experimental facilities program, the strategic position of the U.S. with respect to developing and benefiting from fusion energy could be severely compromised – a significant risk that should be avoided.

To conclude, reliable whole-device modeling capabilities in Fusion Energy Science will surely demand computing resources at the petascale range and beyond to address ITER burning plasma issues. Even more powerful exascale platforms will be needed to meet the future challenges of designing a demonstration fusion reactor (DEMO). Whole device integrated modeling, which can achieve the highest possible physics fidelity, is a most worthy exascale-relevant project for producing a world-leading realistic predictive capability for fusion. This should prove to be of major benefit to U.S. strategic considerations for a carbon-free energy future, ecological sustainability, and U.S. energy security.

4.8 National Security

In one particular mission area of the Department of Energy, the need for significant continued improvements in computational capabilities has been long recognized: This mission relates to U.S. national security, and as a particular exemplifier, the nuclear weapons stockpile managed by the National Nuclear Security Administration (NNSA). As already discussed in Section 2.2.8 above, computing at the leadership level has long characterized modeling and simulation related to the stockpile; and it is broadly understood that continued performance of the stockpile will require continued advances in modeling and simulation capabilities for stockpile certification.

Modeling and simulations in the national security realm ranges from fundamental science, including especially the behavior of materials under extreme physical conditions to modeling of weapons at both the subsystems and the fully-integrated systems levels, and to – most important in the context of the certification process – uncertainty analysis focused on expected weapons performance. The details (at the unclassified level) are provided in the Workshop Report "Scientific Grand Challenges in National Security: The Role of Computing at the Extreme Scale" referenced in Appendix 2. This report provides a broad view of the needs for high-performance computing within the national security realm, ranging from the weapons program itself to programs that deal with (for example) nuclear nonproliferation; for present purposes, we focus here only on the main points related to the weapons program, focused in particular on the expected transformational role that computing at the exascale is expected to play in this realm.

The essence of ensuring that nuclear weapons systems function as intended centers on a quantitative assessment of their functionality. Until the effective implementation of nuclear test bans, these quantitative assessments were carried out by a combination of detailed inspection and actual (underground) tests. Subsequent to the nuclear test ban, such assessments have increasingly relied on a combination of detailed inspection ('surveillance'), subsystem tests and experimentation, and modeling and simulation. Since the expected performance degradation of weapons systems is commonly expected to increase with time since the last underground tests, the scope of the simulations has progressively increased; and one can expect that in the not-too-distant future, there may be the need for a capability to model entire weapons systems.

Tackling the study of extremely complex phenomena as exemplified by nuclear weapons systems is a classical 'grand challenge' problem, in which 'emergent' phenomena exist that do not trivially derive from an understanding of the system components (no matter how detailed that understanding is) but are instead a consequence of nonlinear interactions between these components. Thus, the 'grand challenge' here is very much a multi-physics, multi-scale (in both space and time) problem; and the kinds of questions that arise in other complexity problems – the origins of this complexity, how to characterize it, and how to harness its consequences – arise here as well.

To date, modeling and simulations have proved to be highly successful in furthering our understanding of the workings of weapons systems at the subsystem level; and a key ambitious step over the next decade is to extend this understanding to the fully integrated system. It is broadly believed that taking this step will require a new paradigm for integrating research at a broad range of levels – from the fundamental

sciences[53] to the computing disciplines – and thus the need (and the opportunity) to move beyond better observation, measurement, and simulation capabilities to a true "co-design" strategy for accelerated discovery, prediction, and control: That is an exciting grand challenge. To quote the aforementioned report, "… through co-design, the evolution of each capability will be guided by that of others so that observation, measurement, and simulation will be well matched to jointly tackle the study of complex phenomena. Achieving this qualitatively higher level of interdisciplinary cooperation and integration of major assets is essential if scientists are to meet the great national security and societal challenges of our time—health, energy, defense, and information management." These needs (and the associated opportunities) both drive and justify NNSA's consequent strong focus on extreme computing applications; and these are extensively documented in the NNSA Advanced Simulation and Computing (ASC) strategic plans aimed towards resolving critical science and engineering issues; the ultimate aim is to move to validated predictive systems capability in a non-underground test environment – one in which validated uncertainty quantification provides the level of assurance needed to establish weapons system certification in the absence of underground testing, and thus ensures the safety, security, and effectiveness of the U.S. nuclear weapons stockpile.

4.9 Nuclear Engineering

Recent DOE studies have reviewed the status and basic science challenges, opportunities, and research needs for advanced nuclear energy systems, with specific attention to the role of predictive modeling and simulations (M&S) in addressing the difficulties posed by the radioactive materials and harsh environments found in these systems:

- Computational M&S offers the opportunity to accelerate nuclear energy development by simulating complex systems to evaluate options and predict performance, thus narrowing the technology path and optimizing testing requirements.
- Today's high-performance computational systems are capable of modeling complete reactor systems and related technologies; the availability of exascale systems will enable high-fidelity M&S that can further improve the performance of existing reactors and have a significant positive impact on both the design and the operation of future reactors.

Simulation has the potential for addressing the critical needs of advanced nuclear energy systems by providing the tools necessary for safety assessments, design activities, cost, and risk reduction. One can, for example, imagine virtual prototyping of reactor cores yielding data that leads to more accurate identification of design margins, allows early experimentation with novel design concepts, and ultimately significantly reduces plant certification timelines. In other areas, such as advanced fuel fabrication, atomistic fuel simulations could ultimately make it possible to target a small subset of promising candidate fuel types for further experimentation, greatly reducing the number of experiments to be performed. A simulation-based methodology is within reach with exascale computers.

The scope of the M&S tools needed to support the design, analysis and engineering of next-generation nuclear energy systems is daunting:

1. Integrated 3D reactor core simulations with rigorous propagation of uncertainty;

[53] The range of basic research topics central to this task is huge: It involves a more sophisticated microscopic description of fission and understanding of neutron reactions (ultimately leading to modeling of integrated nuclear systems), material science issues such as interfacial chemomechanics and atomistic simulation on engineering time scales, the ability to predict constitutive representations in multiphase materials, and more powerful electronic structure theory and simulation methods for nonadiabatic and strongly correlated systems. In the realm of chemistry, the goals include advancing the speed and accuracy of quantum chemistry modeling, enabling predictive actinide chemistry, following reactive molecular dynamics with quantified accuracy, and developing the methodology (and capability) for seamlessly modeling multiscale chemical reactivity.

2. Coupled thermal hydraulic and primary loop simulation;
3. Advanced fuel design and performance;
4. Fuel behavior engineering;
5. Advanced secondary loop and balance of plant engineering and analysis;
6. Advanced fuel cycle design;
7. Separations facility engineering optimization;
8. Repository design including seismic, geological, chemical, and thermal modeling and simulation; and
9. Overall nuclear energy systems model development suitable for alternative economic analysis.

Spent fuel aqueous reprocessing is very complicated (more than just solvent extraction) with a large number of different materials: multiple pathways must be considered; waste streams must be treated; improve coupling between computations and experiments must occur. Reprocessing occurs at high temperature, and is in dire need of better multi-scale M&S. The opportunities for impact on reprocessing with exascale M&S abound. These include developing new separation agents (can they be "engineered" to get desired results?), full-scale plant simulations using first principles (some codes are adequate, but others need work), integrating multiple codes, and separations simulations. Empirical understanding does not lead to appropriate scale-up – it will instead require exascale computing. Some of the payoffs for exascale computation include: reduced R&D cost and time; improved/accelerated design; process scale-up; reduced facility cost; opportunity for major change; and waste form design.

Many challenges confront viable and useful (predictive) M&S of fuel performance. These include the ability to reduce fuel development and qualification time, assess life cycle performance, address safety concerns, predict fuel rod behavior in design basis accident (DBA), and predict current and advanced (e.g., transuranic) fuel behavior. Important effects and requirements to incorporate include material properties, swelling, microstructural phase change, thermal properties, crack formation and mechanical property change. High-fidelity modeling of fuel performance is inherently multiscale, e.g., the effects of point defects and fission products must be considered. Exascale platform requirements drivers in fuel performance can be quantified. For example, a high-resolution simulation of a bundle of 40 rods (approximately 300 million elements having a ten micron size), with only thermo-mechanics and no coupling with other multi-physics models, is estimated to require about a half day on 20 PF platform. An even more rigorous simulation (approximately 1 billion elements with a one micron scale size) of one fuel pellet, again with only thermo-mechanics, would require approximately one day on a 1 PF platform. Incorporating all relevant additional physics, such as neutronics, fluid flow, and fundamental material science, will easily multiply these requirements by a factor of one thousand, bringing high-fidelity fuel performance simulation requirements to the exascale.

Opportunities for exascale M&S of existing and future advanced reactors include eliminating unrealistic assumptions that drive to more conservative designs and thus higher installation cost, helping to achieve higher power efficiencies, a reduction of learning curves to get efficiencies, helping to reduce the required number of repositories, improving safety posture, optimizing design of the power grid and the fuel cycle and better (more efficient) operations, including in-line monitoring and operator training. There are numerous issues confronting advanced reactor M&S today. The core is a coupled physics problem (not currently being done very well today) and the full system needs to be analyzed in one tool. Current reactor designs are excessively conservative. For simulation to be "on par with experiment", it will be possible in areas where the physics are well known and data exists and very difficult when it does not.

For the existing fleet of reactors, simulation can reduces margins, thereby increasing performance. For example, for 100 reactors @ $1M/day, a 5% power increase results in a $2B/yr return for 20+ years. For advanced new reactors, optimized designs can result without any "learning curve" (years in operations).

For the current fleet, for example, the difference between 30 years @ 60% capacity or 10 years @ 90% is unrealized potential equating to $328 billion in 2007 dollars over 30 years. M&S can help to ensure that only one nuclear waste repository is required, i.e., is it desirable to avoid or delay a second repository. There are also large potential payoffs from reduction in modeling bias and uncertainty on improved assessment of safety margins, favorable economics (enhanced operation of plants), reduction in the development time for advanced reactor technologies (new and improved fuels for reactors), and improved optimization of plant designs.

M&S can have a large impact on facility design: to study fuel cycle and resource issues: (once thru, plutonium recycle, actinide burning, waste disposal, etc.); to develop and optimize nuclear plant and facility designs (systems design, physical layout, materials requirements, cost, economics); to demonstrate fundamental safety issues (defensive systems, passive safety, establish licensing basis); to avoid (reduce) costly prototype construction and operations (critical assemblies, prototypes, intermediate scale plants); to accelerate optimum fuel selection (fuel irradiation is almost certainly required); and to characterize spent fuel related requirements (on-site storage and criticality, shipping cask designs, and repository requirements). M&S can also have a large impact on plant operations: to optimize performance in a mixed utility electrical grid (cycle length, fuel resource requirements, economics, outages: analyze 100,000's of core loading options); to demonstrate cycle-specific safety requirements (static and transient calculations: 10,000's of coupled neutronics/thermal-hydraulic/systems computations); to provide reactor operator support (optimize startup and power maneuvers: 1,000's of calculations per cycle); to provide on-line monitoring functions (real-time surveillance of safety margins: 1000's of calculations per cycle); and to enhance operator training (real-time simulation on full-scope simulators).

Simulation of advanced nuclear fuel cycles will require a hierarchy of models of vastly different physical systems across a wide range of space-time scales, from detailed molecular dynamics of new materials to systems level simulation of the entire cycle. The final goal will be optimization in the presence of modeling and input uncertainty in order to design safe, reliable, economical, and socially acceptable end-to-end solutions for nuclear energy production. While there have been many advances in fundamental enabling technologies in mathematics and computer science in the past, additional research and development will undoubtedly be required to tackle a problem of this scale. At each level, new enabling technologies will be required to enhance predictive capability, understand and propagate uncertainties, model unresolved physics, and couple multiple physical models at a wide range of space-time scales. Likewise, new research and development is required to analyze, visualize, and optimize the results of such large simulations, and to do so in a way that is useful to designers and decision makers, who must be fully aware of the limitations of the computational predictions and the uncertainties inherent in the simulated results, due to the inevitable uncertainties of input parameters and modeling assumptions. Associated with this is the stringent need to establish careful protocols for simulation code verification and validation.

Enhanced use of M&S for nuclear energy will lead to improvements in knowledge and reduction of uncertainties that will produce cost savings in current reactor operations and substantially reduce the cost of future reactors. Such improvements could also provide a basis for innovative designs that will reduce the need for excessively conservative (and costly) system specifications, improve efficiency and performance, enhance safety and reliability, and extend operating lifetimes.

5. Challenges in Going to the Exascale

5.1 Motivating our approach to meeting the challenges of going to the exascale

Element #1: Overall value proposition – applications outcome oriented, "reproducible"

Creating an exascale computer capable of effectively running the applications just described will require significant R&D breakthroughs. The previous section laid out the case for the wide range of scientific

and technical advances that could be made with an exaflop computer. This section discusses the challenges required to make that three order of magnitude jump in technology. In this type of discussion, it is often far too easy to talk about that jump as some quantitative steps in an evolutionary process, when in fact the jump implies significant qualitative changes in the way solutions must be approached. Consider the following simple table[54] that illustrates three orders of magnitude in change:

Technology	*Quantitative Rate*	*Qualitative Change*
Marathon Runner	10 mph	Explore a town
Car	100 mph	Explore a country
Jet	1,000 mph	Explore a world
Space Craft	10,000 mph	Explore the solar system

The analogy to computing challenges is not quite the same, because we do not have to explore totally different technologies to make the leap in three orders of magnitude. However, just like we would not think of asking a marathon runner to explore the solar system, we cannot use current technology to produce an exaflop system. An exaflop system made entirely out of today's technology would probably cost $100B, require $1B per year to supply the needed power, and its own dedicated power plant to produce that power. At the other extreme, it would be possible to build a relatively inexpensive exaflop computer that would be totally incapable of solving any of the applications already described. The cost of the actual processing units in current computers is so inexpensive they are often considered "free." But an exaflop computer without memory, networking, disks, and all of the other essential components of a complete exascale system is useless. So, the true challenge of "going to the exascale" is a complex interaction of R&D efforts to drive the costs down to something feasible, while increasing our capabilities to use the system we can afford to build to meet the needs described by the various applications. The R&D efforts span a wide range of areas, including hardware, systems software, application algorithms, and computer science. In addition to these challenges, we must also face the challenge of building a skilled workforce to use exascale systems. We need renewed efforts to recruit and train the next generation needed to exploit these capabilities.

Element #2: Maintain US leadership in HPC industry and competitiveness

Understandably, American leadership in applied energy technologies will be important for the future of the American economy. In the face of fierce competition from low cost-of-labor countries intent on extending their advantage, America must continue to invest aggressively in advanced technology areas where it is more of a challenge for emerging economies to dominate. Low cost labor countries are not resting on their laurels. They are not ignoring technology development. As evidence, the graph below highlights the rise of Chinese supercomputing. In 2010, China has produced a #1 system[55], using U.S.

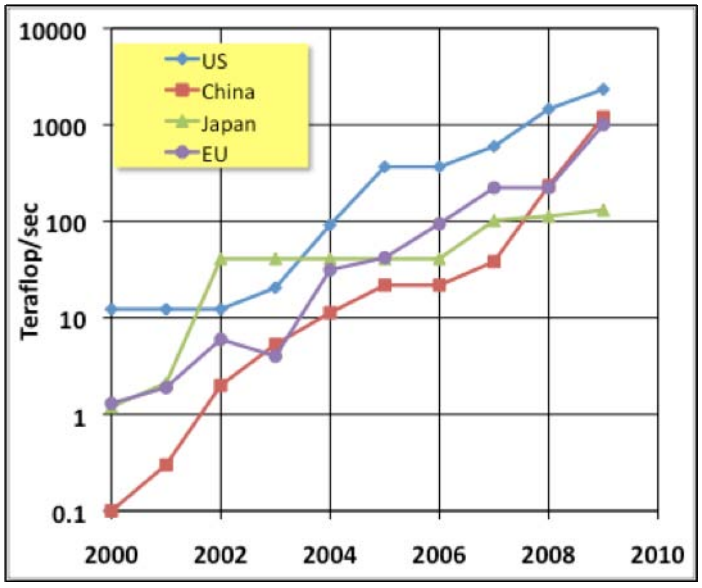

Figure 13: Rapid Rise of China in Top 500 Ranks

[54] Adapted from a talk given by Mark Seager, Lawrence Livermore National Laboratory.
[55] New York Times, Oct. 28, 2010.

processors. The country is currently developing a processor and a platform to support it, and will not be content to leverage American technology over the long term. This presents a challenge, because these systems will not remain the inheritance of research institutions alone. China's drive for "indigenous innovation" is our generations "Sputnik"– this time the race to an Exascale System. They will rapidly find their way into supporting local industry and national competitiveness.[56]

High Performance Computing is one area of remaining advantage to the United States today. America continues to lead in HPC because DoE (including NNSA), DARPA, NSF and other areas of government have made aggressive investments *in the past*. In particular, NNSA has led in this, having recognized over a decade ago that integral to meeting the requirement of the zero testing in Nevada was the *balanced and coordinated development* of advanced applications, underlying science, verification and validation, uncertainty quantification technologies, and, of course, the development and siting the reliable and powerful computers. All of this was underwritten by the expertise that evolved at the national laboratories as these investments bore fruit. Together these produced actionable information for decision makers.[57]

5.2 The Hardware Challenges

5.2.1 Exaflop hardware needs major R&D progress

The Institute of Advanced Architectures and Algorithms (IAA) published a recent report in the International Journal of Distributed Systems and Technologies, in which thirteen of their members proclaimed "The architectural challenges for reaching exascale are dominated by power, memory, interconnection networks and resilience."[58] Table 1 below compares current HPC designs with potential exascale designs from the DOE Exascale Initiative Roadmap. The baseline we need is a factor of 500 change in peak system performance. The difference in factor changes for the various components show where simple scaling of systems (e.g., buying 500 2 Pf/s systems) will be inadequate. Take for example, the power line in the table. While the peak speed goes up by 500, the power cost cannot go up by more

	2010	2018	Factor Change
System peak	2 Pf/s	1 Ef/s	500
Power	6 MW	20 MW	3
System Memory	0.3 PB	10 PB	33
Node Performance	0.125 Gf/s	10 Tf/s	80
Node Memory BW	25 GB/s	400 GB/s	16
Node Concurrency	12 cpus	1,000 cpus	83
Interconnect BW	1.5 GB/s	50 GB/s	33
System Size (nodes)	20 K nodes	1 M nodes	50
Total Concurrency	225 K	1 B	4,444
Storage	15 PB	300 PB	20
Input/Output bandwidth	0.2 TB/s	20 TB/s	100

Table 1: Potential Exascale Computer Design for 2018 and its relationship to current HPC designs. [59]

[56] *Exascale Workshop Panel Meeting Report*, DOE, ASCR, January 19-20, 2010, Washington D.C. p. 14.
[57] Conceptual Design Document, LVOC, LLNL (2010).
[58] *International Journal of Distributed Systems and Technologies*, April-June 2010, **1**: 1-22.
[59] DOE Exascale Initiative Roadmap, Architecture and Technology Workshop, San Diego, December, 2009.

than a factor of 3. That means that the power solution for an exaflop system has to be over 150 times more efficient than current technology. That is a huge challenge. Looking through the other entries, the table clearly echoes the sentiments of the IAA, and highlight key features that must be addressed in hardware or downstream in software.

5.2.2 Power Challenge

All of the technical reports on exascale systems identify the power consumption of the computers as the single largest hardware research challenge. Today, power costs for the largest petaflop systems are in the range of $5-10M[60] annually (depending on the best rates that can be negotiated with energy companies). So, to achieve an exascale system using current technology, the annual power cost to operate the system would be above $2.5 B per year. The power load would be over a gigawatt (more than many power plants currently produce). The target of 20 megawatts, identified in the DOE Technology Roadmap, is primarily based on keeping the operational cost of the system in some kind of feasible range.

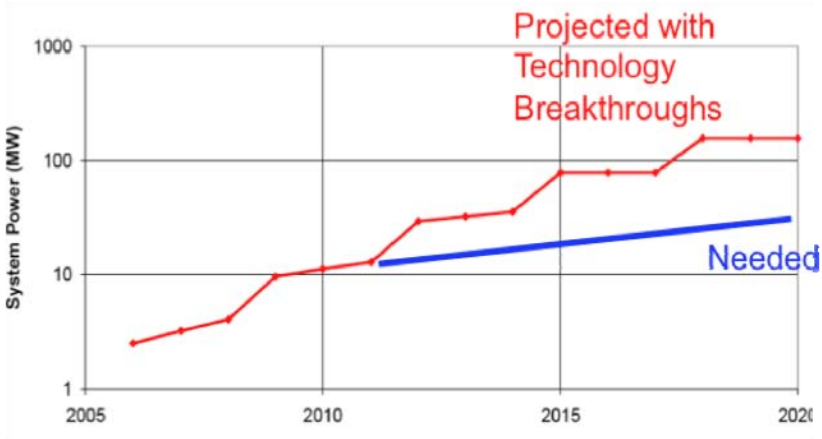

Figure 14: Possible Leadership class power requirements, from Peter Kogge (on behalf of Exascale Working Group), "*Architectural Challenges at the Exascale Frontier*", June 20, 2008

5.2.3 System Memory Challenge

Achieving this new power goal for exascale systems creates serious research challenges. There is overwhelming consensus that "power is the one area in which significant breakthroughs will be the most difficult to achieve.[61]" Figure 14 demonstrates projected technology with breakthroughs in red, and what is needed to get to ~20 MWs in blue. So, even with the optimistic expectations of current R&D activities, there is at least a factor of five gap between what we must have and what current research can provide. To get the additional factor of 5 improvement in power efficiency over projections, a number of technical areas in hardware design for exascale systems are currently being explored. These include: energy efficient hardware building blocks (CPU, memory, interconnect), novel cooling and packaging, Si-Photonic communication, and power-aware runtime software and algorithms.

From a scientist's perspective, the ratio of memory to processor is critical in determining the size of the problem that can be solved. The processor dictates how much computing can be done; the memory

[60] Lawrence Livermore National Laboratory, budget reports for petaflop systems.

[61] "Exascale: The Beginning of the Great HPC Disruption," *The Exascale Report*, http://theexascalereport.com, Mike Bernhardt, August 2010.

dictates the size of the problem that can be handled. In the Exascale design described above, there is 500 times more compute power, however only 30 times the memory, so applications cannot just scale to the speed of the machine. Scientists and computer scientists will have to rethink how they are going to use these systems. This factor of >10 loss in memory size/compute power means potentially totally redesigning the current application codes. Additionally, starting this process now will be critical to having the application codes ready when the exascale systems are in place in 2018. This process will lead to the co-design process discussed later.

As seen in Figure 15 below, the problem is compounded by the fact that memory access time has not improved with CPU cycle time, so it is taking far longer to get data to a CPU than it takes to use that data.

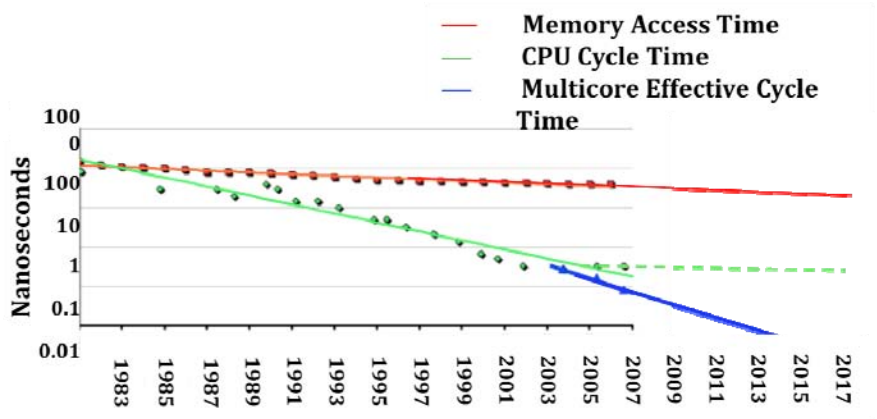

Figure 15: CPU and memory cycle trends, 1983 projected forward to Exaflop systems.

5.2.4 Data Movement

The system memory challenge is only one important aspect of a broader Data Movement challenge for Exascale systems. Figure 16[62] shows the conceptual organization of the hierarchy of components that store and move data in HPC systems. This hierarchy has been designed to provide and store the data needed by processing units. At the top of the pyramid are registers and cache. These are the fastest, smallest, and most expensive memory modules in the system. Main memory is much slower (processors can perform 100 operations in the time it takes to get one word of data from memory), and much larger than cache. External storage is far larger, but also operates at extremely slow rates relative to processor speeds. Since all these levels of the pyramid are needed to run scientific applications, the challenge is providing as much capacity at each level (consistent with an overall plausible machine cost) and providing the most effective methods for moving data among the levels as dictated by the needs of the various applications. The two key barriers to providing adequate capacity at each level are:

- Achieving adequate rates of data transfer, or bandwidth, and reducing time delays, or latency, between the levels, and
- Per-disk performance, failure rates, and energy efficiency no longer improving.

Neither bandwidth nor latency in the entire memory pyramid have improved at rates comparable to Moore's Law for processing units. On current petaflop systems, memory access at all levels is the limiting factor in almost all applications, so the situation for exaflop systems will be critical. Research in optical interconnects and routers can provide critical improvements in latency and bandwidth. But work will also be needed in improving the efficiency of data when and how it needs to move. Research options include better data analysis to anticipate needed data before it is requested (thus hiding latency), determining when data can be efficiently recomputed instead of stored (reducing demands for bandwidth),

[62] *ExaScale Software Study: Software Challenges in Extreme Scale Systems*, DARPA IPTO, Sept. 14, 2009, p. 26.

and improved data layouts (to maximize the use of data when it is moved between levels. Other hardware research, including non-volatile memory gap fillers and advanced packaging (chip stacking) will help fill the current latency gap and provide enough overall storage capacity to run exascale applications.

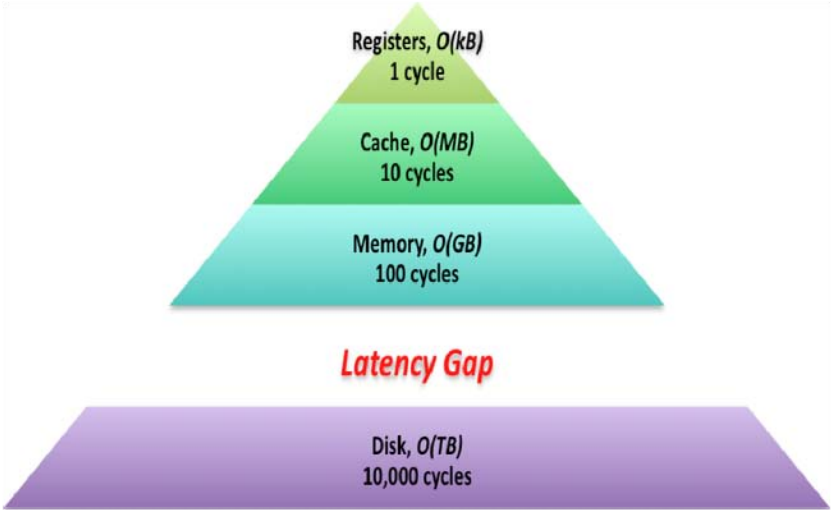

Figure 16: The data access/storage hierarchy.

5.2.5 System Resiliency Challenge

As we move from petascale systems to exascale systems, the number of system components will be increasing faster than component reliability, with projections in the minutes or seconds for exascale systems. From the current knowledge and observations of existing large systems, it is anticipated that exascale systems will experience various kind of faults many times per day. Increasing evidence points to a rise in silent errors (faults that never get detected, or get detected long after they generated erroneous results), causing havoc which will only get more problematic as the number of components rise with exascale systems. Systems running 100 million cores will continually see core failures and the tools for dealing with them will have to be rethought.[63]

The current approach for resilience, which relies on automatic or application level checkpoint/ restart, will not work because the time for checkpointing and restarting will exceed the mean time to failure (MTTF) of a full system. This set of projections presents a difficult challenge: finding new approaches to run applications until their normal termination, despite the projected unstable nature of exascale systems. The ability for a scientist to make forward progress will be difficult unless alternative methods to fault recovery are provided that do not involve checkpoint/restart. [64]

Currently, there is technical progress started in several areas. These include improving hardware and software reliability, better understanding of the root cause of RAS collection and analysis and, additionally, fault resilient algorithms and applications to assist the application developer[65], and local recovery and migration. The goal of this research is to improve the mean time for interrupts (MTTI) by >100x, so that applications can run for many hours. An additional goal is to improve by a factor of 10X the hardware reliability and improve by a factor of 10X the local recovery and migration of data.

---

[63] Thibodeau, P. 2009. *Supercomputers with 100 million cores coming by 2018, COMPUTERWORLD,* http://www.computerworld.com, November 16, 2009.

[64] Cappello, F, et al. 2009. "Toward Exascale Resilience." *International Journal of High Performance Computing Applications* **23**: 374-388.

[65] *System Software Breakout Session,* Cross-cutting Technologies for Computing at the Exascale, February 2-4, 2010

5.3 The "Co-Design" Challenge

Co-design is a critical component of the exascale initiative. In fact, this process will both integrate all components of the exascale initiative and illuminate the trade-offs that must be made within this complex undertaking.

Classically, co-design refers to the boundary between hardware and software implementations in embedded chip design[66]. This field has a long history with specific courses being taught in academic and industrial settings. The following discussion[67] of the basic tenets of co-design is an instructive viewpoint in broadening the context to exascale,

> "Software/Hardware codesign can be defined as the simultaneous design of both hardware and software to implement in a desired function. Successful codesign goes hand in hand with co-verification, which is the simultaneously verification of both software and hardware and in what extent it fits into the desired function. In today's world it is necessary to incorporate codesign in the early system design phase and put software-hardware integration downstream because traditional methodologies aren't effective any longer. Today, we try to put the foundation to incorporate codesign and co-verification into the normal product development in place. Especially when products incorporating embedded systems are involved."

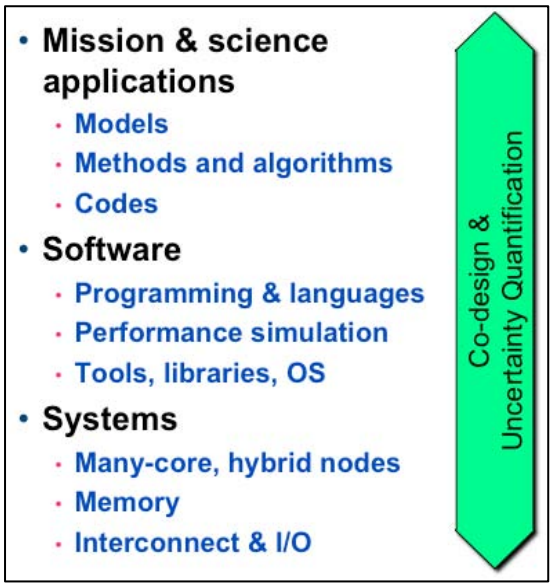

Exascale co-design is a very complex undertaking. Nowhere is this complexity more evident than in the requirement for exascale to enable a variety of predictive mission and science outcomes. Examples are anticipation and mitigation of regional climate change events, effective design of new, carbon-aware energy sources, certification of the U.S. stockpile, and understanding of the world around us, from the formation of nucleons to dark energy and dark matter.

System and hardware vendors are another important piece of this puzzle. With clock rates flat at roughly several Gigahertz, systems will require more than one billion concurrent operations to achieve exascale levels of performance. Most of this increase in concurrency will be within the compute node. Other critical issues on the node designers' plate are power (roughly a factor of 1000x better than today), resiliency to faults or data corruption (roughly a factor of 100x on the node), memory size and speed and at least constant performance relative to peak. Historically, the high-level system design process relies on important, fixed software kernels or snippets of code that can be used to measure hardware design trade-offs. Unfortunately, application code, methods and models will be also be adapting, both to new requirements for exascale applications (e.g. new non-hydrostatic models of atmosphere and ocean) and to the changing computing landscape. Thus, a simultaneous or iterative co-design process must be developed and applications and system developers must commit to it.

The key is the system software layer that mediates the interaction between applications and hardware. This also will be changing, for example, a new programming or execution model will likely be required to effectively manage on-node concurrency and data movement. In fact, the exascale plan is early focus and

[66] See for example *Hardware-Software Co-design of Embedded Systems* (1997)
[67] http://www.npd-solutions.com/swcodesign.html

decision with respect to the programming models, so as to facilitate the interactions between applications and hardware. Co-design requires the ability to accurately simulate, at several levels, the performance of key application components. For understanding of overall power and performance, a high level simulation of nodes, interconnects and I/O is required. Cycle-accurate simulations of chip power, resilience and performance are necessary for detailed hardware designs. These simulators are often proprietary because they contain detailed vendor-specific innovations, adding yet another level of complexity that co-design needs to manage.

To design productive, cost-effective systems, we must move from simple rule-of-thumb design choices for key design points such as system bandwidth, programming model, and I/O to co-designed systems that efficiently balance design parameters. Co-design is critical to understanding and optimizing the trade-offs among applications, software and hardware. It must be built into the exascale plan, not as an afterthought, but as an overarching methodology necessary for success.

5.4 The Applied Mathematics Challenges

The need for a robust program of applied mathematics research is pervasive in every description of the potential transformative effects of exascale computing (see Section 4). Although specific "wish lists" differ from application to application, they are strikingly similar in emphasizing the need for better mathematical models (for, e.g., resolving fine structure in a rapidly changing environment, defining a self-teaching adaptive model, evaluating options in designing human-made systems, and analyzing the structure of uncertainties arising from imperfect measurements and incomplete knowledge of physical or biological phenomena). In addition, increasingly subtle and sophisticated models necessarily demand new numerical methods that can cope with features such as huge dimensionality, nonlinear and/or coupled structures, mixtures of continuous and discrete variables, and drastically different scales.

The applied mathematics component of an exascale program should include attention to activities with time horizons ranging from medium-term to very long term, where both ends of the time scale are essential. The description "medium-term" is deliberate because experience in adapting to new computational modalities shows that short-term, one-off strategies (except for an overriding and urgent national need) are likely to be wasteful. As Knuth, the guru of modern computer science, famously said: "We should forget about small efficiencies, say about 97% of the time: premature optimization is the root of all evil (or at least most of it) in programming"[68].

Even though much remains unknown about the details of exascale systems, a clear medium-term priority is the definition and implementation of algorithms that are scalable at very large levels of parallelism (such as on million-core machines) and that remain sufficiently fast under different (as yet unmade) hardware decisions about bandwidth and latency; see Section 5.2. Scalability should be modeled and analyzed mathematically, using abstractions that represent key architectural ingredients. Simulations and experiments that indicate the effects of hardware and software perturbations on algorithmic efficiency can then guide the definition of methods that retain scalability under a variety of hardware scenarios. Improvements in overall solution time are what matters most in this context, not artificial flop numbers.

With respect to the long term, the most dramatic transformational mathematical models and numerical algorithms require an extended period of time, comprising both focused effort and unpredictable surprises, to move from the first "gleam in the eye" to robust and efficient software – in large part, precisely because of the disruptive nature of a genuine innovation. For example, the fast multipole method of Greengard and Rokhlin[69], an immediately obvious breakthrough, was published in 1987, but its widespread use in applications arguably needed 15 years of effort by numerous researchers to work out crucial details and develop general-purpose software implementations. In this spirit, the strategies for

[68] Knuth, D. 1974. "Structured programming with GoTo statements." *Computing Surveys* **6**: 268.
[69] Greengard, L & V Rokhlin. 1987. "A Fast Algorithm for Particle Simulations." *JCP* **73**(2): 325-348.

applied mathematics in exascale science will require sustained support over time for people-intensive activities, early identification of the hardest (and least straightforward) research problems, and built-in flexibility to pursue unexpected and promising new directions as they arise.

5.4.1. Mathematical modeling

It is natural for those developing mathematical models of practical problems to limit themselves to formulations that can be solved numerically using currently available methods. Although essential when the problem needs to be solved in the short term, an *ab initio* focus on feasibility can create a too-rigid environment in which non-standard or "blue-sky" formulations are either never thought about or else summarily rejected. For example, a problem formulation that represents many real-world problems yet tends to be avoided because of its known intractability is constrained nonlinear optimization with a mixture of continuous, discrete, and categorical variables. But the prospect of massive increases in computational power means that modeling ideas previously dismissed as impossible or impractical may well become realistic, and should be carefully examined and analyzed.

Creative rethinking of mathematical models is an essential strategy to address the challenges of exascale science. The highly desired "transformational" changes flowing from exascale computing are most likely to come from new formulations that change the way we think about problems, rather than from applying more resources to an existing formulation to obtain a more accurate solution or to solve a larger problem. Mathematical models are inherently an approximation of reality, and an exascale initiative provides an opportunity to loosen the grip of, or even remove, computationally-imposed simplifications. The major challenge is to devise models that capture the important details of physical and engineered systems as they really are. This will almost certainly generate much harder subproblems and/or much more data, but the gains are likely to be eminently worthwhile.

5.4.2. Numerical algorithms

The need for scalable algorithms in an exascale initiative has already been stressed. Another essential feature, highlighted in a 2009 talk by Kathy Yelick called *Ten ways to waste a parallel computer*, is a "back to basics" approach to reformulation.[70] Without careful analysis of both new models and new numerical methods, there is the risk of significant inaccuracy or large computational overhead in unexpected parts of the overall solution process, as illustrated in the following two examples related to numerical methods for partial differential equations:

1. All indications are that memory will become the rate-limiting factor along the path to exascale, and investments should accordingly be made in designing algorithms with reduced memory requirements. Examples where this work is appropriate include: (i) algorithmically scalable matrix-free methods (e.g., multigrid) for sparse systems of equations, where "algorithmically scalable" means that the total resources needed to solve the problem (flops plus memory) are proportional to the resources needed to evaluate the associated operator; (ii) high-order methods that perform more computation to obtain greater accuracy for each computational degree of freedom; and (iii) adaptive models/methods designed to use the smallest possible number of degrees of freedom to obtain the needed level of accuracy.
2. Many calculations related to DOE missions involve models that depend on both space and time. In current methods, obtaining better spatial resolution typically requires a comparable reduction in the time step. A frequent argument for exascale science is that it will allow much finer spatial resolution in numerous application domains, with (for example) meshes reduced in size by a factor of ten. Unfortunately, simply reducing mesh spacing by a factor of ten could lead to a ten-fold increase in the time for solution, even with perfect weak scaling. Several strategies, all in the spirit of rethinking, should be explored to avoid this inefficiency. For example, models can be made more implicit to

[70] See isca09.cs.columbia.edu/ISCA09-WasteParallelComputer.pdf

avoid restrictive time-step conditions arising from stiff processes that rapidly relax to equilibrium (e.g., in the context of low Mach-number fluid flows). A further strategy is aggressive use of subcycling in time for processes that are fast, but either are localized in physical space or involve only a small subset of the variables in state space. A motivating example here is advection in the jet stream in atmospheric modeling.

Approaches of this flavor across the spectrum of numerical methods will almost certainly lead to increased algorithmic complexity, in addition to the daunting software-related challenges discussed in Section 5.5 below. The substantially greater work needed to devise exascale numerical methods and software leads us to observe that, for decades, there has been, roughly speaking, a dichotomy in the wish list for the mathematical software used to solve scientific and engineering problems. On one hand, many DOE scientists have neither time nor inclination to become experts in numerical methods and software techniques, preferring to leave software development to mathematicians and computer scientists. On the other hand, some scientists and engineers want to become deeply involved in producing domain-specific methods and software to attain the highest possible efficiency for their particular problem. An exascale science program needs to address the needs of both these groups.

For the first, "professional" mathematical software and libraries (meaning software developed by mathematicians and computer scientists for relatively generic problems such as solving linear systems or eigenvalue problems) should be developed for increasingly broad problem categories as we move toward exascale. In this way, domain scientists will be able to use state-of-the-art software components that can be shared across multiple application domains, e.g., through SciDAC Centers for Enabling Technologies (CETs). Since writing software is universally recognized to be time-consuming and error-prone, scientists and engineers throughout DOE will benefit from availability of software that they can use off the shelf while experimenting with domain-specific challenges rather than (for example) writing their own sparse matrix package. For the second group, specific scientific case studies should be identified that require significant involvement of domain scientists, mathematicians, and computer scientists in end-to-end software development (using the co-design paradigm described in Section 5.3 above).

5.4.3. Mathematics for massive data

It is, of course, impossible for DOE to support every form of mathematics just in case it turns out to be useful. But two recently developed fields in the mathematical sciences — machine learning and compressive sensing — are now unarguably important for DOE science and engineering because they provide demonstrable means for understanding, representing, and learning from very large sets of data. Research in both should be part of a DOE exascale science initiative.

5.4.3.1. Machine learning

This inherently interdisciplinary topic is variously located in departments of mathematics, computer science, statistics, electrical engineering, and business. We do not propose to define the field, except to note that its essence is to design algorithms that "learn" (detect patterns and relations) from large data sets, and then make intelligent decisions based on data. Several of DOE's activities involve analysis of massive quantities of data of multiple varieties (numerical, qualitative, categorical), which means that machine learning is a natural candidate for inclusion in exascale science. Machine learning is becoming ever more important in real-time control of data collection systems, which are also of interest to DOE.

5.4.3.2. Compressive sampling

For slightly more than five years, the new field of "compressive sampling", started in the mid-2000s[71], has emerged with the seemingly miraculous property that signals can be reproduced exactly based on many fewer samples than the dimension of the signal. The key is that the signals are assumed to be sparse, in the sense of having fewer degrees of freedom than their dimension. Thus the signal is sampled

[71] Donoho, D. 2004. *Compressed sensing*, Technical report, Stanford.

(and simultaneously compressed) at a greatly reduced rate while achieving significant fidelity for sparse signals. Compressive sensing is a "hot topic", and is already leading to a rethinking of sensing and data collection mechanisms. Although there is, as yet, no direct connection with extreme-scale computing, DOE should ensure that research is underway to develop effective compressive sensing techniques along the path to exascale.

5.4.4. Symbolic computing

Symbolic computing relates to manipulation of mathematical symbols (expressions, equations, etc.) in symbolic rather than numerical form. Traditional DOE applied mathematics has tended to avoid symbolic computation in favor of numerical methods, with the exception of sparse matrix structure analysis and automatic differentiation, for two reasons: symbolic computing has been considered (and was) inordinately costly, and, until recently, symbolic and numerical software were difficult to blend. Both of these drawbacks have been greatly reduced or eliminated, and many applied mathematical problems, particularly in modeling, can benefit from the option of switching between symbolic and numerical environments. Modeling languages in optimization, for example, derive significant efficiencies from algebraic simplification of non-linear constraints. Compute-intensive tasks that combine numerical and symbolic computation will provide a useful element in applied mathematics research related to exascale computing.

5.5 The Algorithmic Challenges

Advancing science in key areas requires development of next-generation physical models to satisfy the accuracy and fidelity needs for targeted simulations. The impact of these simulation fidelity needs on requirements for computational science is twofold. First, more complex physical models must be developed to account for more aspects of the physical phenomena being modeled. Second, for the physical models being used, increases in resolution for key system variables, such as numbers of spatial zones, time steps or chemical species, are needed to improve simulation accuracy, which in turn places higher demands on computational hardware and software.

Application models represent the functional requirements that drive the need for certain numerical algorithms and software implementations. The choice of model is in part motivated by the science objectives, but it is also constrained by the computer hardware characteristics attainable in the relevant time frame. The choice and specification of system attributes (e.g., peak speed or node memory capacity) tend to constrain the functional attributes able to be employed in a given physical model on that system.

Science priorities lead to science models, and models are implemented in the form of algorithms. Algorithm selection is based on various criteria, such as appropriateness, accuracy, verification, convergence, performance, parallelism and scalability.

Models and associated algorithms are not selected in isolation but must be evaluated in the context of the existing computer hardware environment. Algorithms that perform well on one type of computer hardware may become obsolete on newer hardware, so selections must be made carefully and may change over time.

Moving forward to exascale will put heavier demands on algorithms in at least two areas: the need for increasing amounts of data locality in order to perform computations efficiently, and the need to obtain much higher factors of fine-grained parallelism as high-end systems support increasing numbers of compute threads. As a consequence, parallel algorithms must adapt to this environment, and new algorithms and implementations must be developed to extract the computational capabilities of the new hardware.

As with science models, the performance of algorithms can change in two ways as application codes undergo development and new computer hardware is used. First, algorithms themselves can change, motivated by new models or performance optimizations. Second, algorithms can be executed under

different specifications, e.g., larger problem sizes or changing accuracy criteria. Both of these factors must be taken into account.

Significant new model development, algorithm re-design and science application code reimplementation, supported by (an) exascale-appropriate programming model(s), will be required to exploit effectively the power of exascale architectures. The transition from current sub-petascale and petascale computing to exascale computing will be at least as disruptive as the transition from vector to parallel computing in the 1990's.

Uncertainty quantification will permeate the exascale science workload. The demand for predictive science results will drive the development of improved approaches for establishing levels of confidence in computational predictions. Both statistical techniques involving large ensemble calculations and other statistical analysis tools will have significantly different dynamic resource allocation requirements than in the past, and the significant code redesign required for the exascale will present an opportunity to embed uncertainty quantification techniques in exascale science applications.

5.5.1 New multicore-friendly and multicore-aware algorithms

Scalable multicore systems bring new computation/communication ratios. Within a node data transfers between cores is relatively inexpensive, but temporal affinity is still important for effective cache use. Across nodes, the relative cost of data transfer is growing very large. The development of new algorithms that take these issues into account can often perform very well, as do communication-avoiding algorithms that increase the computation/communication ratio or algorithms that support simultaneous computation/communication, or algorithms that vectorize well and have a large volume of functional parallelism.

5.5.2 Adaptive Response to Load Imbalance

Adaptive multiscale algorithms are an important part of the DOE portfolio because they apply computational power precisely where it is needed. However, they introduce challenging computational requirements because they introduce dynamically changing computation that result in load imbalances from static distribution of tasks. As we move towards systems with billions of processors, even naturally load-balanced algorithms on homogeneous hardware will present many of the same daunting problems with adaptive load balancing that are observed in today's adaptive codes. For example, software-based recovery mechanisms for fault-tolerance or energy-management features will create substantial load-imbalances as tasks are delayed by rollback to a previous state or correction of detected errors. Scheduling based on a Directed Acyclic Graph (DAGs) also requires new approaches to optimizing for resource utilization without compromising spatial locality. These challenges require development and deployment of sophisticated software approaches to rebalance computation dynamically in response to changing workloads and conditions of the operating environment.

5.5.2 Multiple precision algorithms/software

Algorithms and applications are becoming increasingly adaptive and we have seen that various adaptivity requirements have become an essential, key component of their roadmap to exascale computing. Another aspect of this quest to adaptivity is related to the development of libraries that recognize and exploit the presence of mixed precision mathematics. A motivation comes from the fact that, on modern architectures, the performance of 32-bit operations is often at least twice as fast as the performance of 64-bit operations. Moreover, by using a combination of 32-bit and 64-bit floating point arithmetic, the performance of many linear algebra algorithms can be significantly enhanced while maintaining the 64-bit accuracy of the resulting solution. This can be applied not only to conventional processors but also to other technologies such as GPUs, and thus can spur the creation of mixed precision algorithms that more effectively utilize heterogeneous hardware.

Mixed precision algorithms can easily provide substantial speedup for very little code effort by mainly taking into account existing hardware properties. Earlier work has shown how to derive mixed precision versions for various architectures and for a variety of algorithms for solving general sparse or dense linear systems of equations. Typically, a direct method is first applied in single precision in order to achieve a significant speedup compared to double precision. Then an iterative refinement procedure aims at retrieving the lost digits. Iterative refinement can also be applied for eigenvalue and singular value computations.

Of current interest is to extend and incorporate this approach in applications that do not necessary originate from linear algebra, and to study the robustness of mixed precision algorithms on large-scale platforms. Indeed, the convergence of the mixed precision iterative refinement solvers strongly depends on the condition number of the matrix at hand. The conditioning can be determined at run time and proper precision can be selected. Ideally, the user could specify the required precision for the result and the algorithm would choose the best combination of precision on the local hardware in order to achieve it. The actual mechanics would be hidden from the user.

5.5.3 Fast implicit solvers

Carefully analyzing complex problems, and adapting preconditioners to the underlying problem physics is how most of the progress in this area is being made. However, it is typically the case that advanced preconditioners are composed of standard algebraic components such as advanced multigrid/multilevel methods, incomplete factorizations and basic smoothers. Furthermore, we need to renew our focus on basic iterative methods in an attempt to address bottlenecks due to collective operations (e.g., dot-products) and poor kernel performance. Emphasis on block methods, recycling methods, s-step like methods and mixed precision formulations will be necessary to address the next generation of problems.

5.5.4 Communication avoiding

Algorithmic complexity is usually expressed in terms of the number of operations performed rather than the quantity of data movement to memory. This is antithetical to the true costs of computation where memory movement is very expensive and operations are nearly free. To address the critical issue of communication costs, there is a need to investigate algorithms that reduce communication to a minimum. One needs to derive bandwidth and latency lower bounds for various dense and sparse linear algebra algorithms on parallel and sequential machines, e.g., by extending the well-known lower bounds for the usual $O(n^3)$ matrix multiplication algorithm. Then discover new algorithms that attain these lower bounds in many cases. Second, for Krylov subspace methods like GMRES, CG and Lanczos, one should focus on taking k steps of these methods for the same communication costs as a single step.

Today we are at the extreme, with the next generation of machines having billions of potentially parallel tasks. We are being challenged to devise algorithms and software that can effectively exploit the parallel hardware systems that are being developed. When solving very large problems on parallel architectures the most significant concern becomes the cost per iteration of the method—typically on account of communication and synchronization overheads. This is especially the case for preconditioned Krylov methods, which are the most popular class of iterative methods for large sparse systems.

5.5.4 Auto-tuning

Libraries need to have the ability to adapt to the possibly heterogeneous environment in which they have to operate. The adaptation has to deal with the complexity of discovering and implementing the best algorithm for diverse and rapidly evolving architectures. This calls out for automating the process, both for the sake of productivity and for correctness. Here, productivity refers both to the development time and for the user's time to solution. The objective is to provide a consistent library interface that remains the same for users independent of scale and processor heterogeneity, but which achieves good performance and efficiency by binding to different underlying code, depending on the configuration. The diversity and rapid evolution of today's platforms mean that auto-tuning of libraries such as BLAS will be

indispensable to achieving good performance, energy efficiency, load balancing, etc., across this range of systems. In addition, the auto-tuning has to be extended to frameworks that go beyond library limitations, and are able to optimize data layout (such as blocking strategies for sparse matrix/SpMV kernels), stencil auto-tuners (since stencils kernels are diverse and not amenable to library calls) and even tuning of optimization strategy for multigrid solvers (optimizing the transition between the multigrid coarsening cycle and bottom-solver to minimize runtime). Adding heuristic search techniques and combining them with traditional compiler techniques will enhance the ability to address generic problems extending beyond linear algebra.

5.5.5 Scheduling and memory management for heterogeneity and scale

Extracting the desired performance from environments that offer massive parallelism, especially where additional constraints (e.g., limits on memory bandwidth and energy) are in play, requires more sophisticated scheduling and memory management techniques than have heretofore been applied to linear algebra libraries. Another form of heterogeneity comes from confronting the limits of domain-decomposition in the face of massive explicit parallelism. Feed-forward pipeline parallelism can be used to extract additional parallelism without forcing additional domain-decomposition, but exposes the user to dataflow hazards. Ideas relating to a data flow-like model, expressing parallelism explicitly in directed acyclic graphs (DAGs), so that scheduling tasks dynamically, support massive parallelism, and apply common optimization techniques to increase throughput. Approaches to isolating side-effects include explicit approaches that annotate the input arguments to explicitly identify their scope of reference, or implicit methods such as using language semantics or strongly typed elements to render code easier to analyze for side-effects by compiler technology. New primitives for memory management techniques are needed that enable diverse memory management systems to be managed efficiently and in coordination with the execution schedule.

5.5.6 Fault tolerance and robustness for large-scale systems

Modern PCs may run for weeks without rebooting and more data servers are expected to run for years. However, because of their scale and complexity, today's supercomputers run for only a few days before rebooting. Exascale systems will be even more complex and have millions of processors in them. The major challenge in fault tolerance is that faults in extreme scale systems will be continuous rather than an exceptional event. This requires a major shift from today's software infrastructure. Every part of the exascale software ecosystem has to be able to cope with frequent faults; otherwise applications will not be able to run to completion. The system software must be designed to detect and adapt to frequent failure of hardware and software components. On today's supercomputers every failure kills the application running on the affected resources. These applications have to be restarted from the beginning or from their last checkpoint. The checkpoint/restart technique will not be an effective way to utilize exascale systems, because checkpointing stresses the I/O system and restarting kills 999,999 running tasks because 1 fails in a million task application. With the potential that exascale systems will be having constant failures somewhere across the system, application software isn't going to be able to rely on checkpointing to cope with faults. A new fault will occur before the application could be restarted, causing the application to get stuck in a state of constantly being restarted. For exascale systems, new fault tolerant paradigms will need to be developed and integrated into both existing and new applications.

Research in the reliability and robustness of exascale systems for running large simulations is critical to the effective use of these systems. New paradigms must be developed for handling faults within both the system software and user applications. Equally important are new approaches for integrating detection algorithms in both the hardware and software and new techniques to help simulations adapt to faults.

5.5.7 Building energy efficiency into algorithms foundations

It is widely recognized (see section 5.2 on *The Hardware Challenges*) that emerging constraints on energy consumption will have pervasive effects on HPC; power and energy consumption must now be

added to the traditional goals of algorithm design, viz. correctness and performance. The emerging metric of merit becomes performance per watt. Consequently, we believe it is essential to build power and energy awareness, control and efficiency into the foundations of our numerical libraries. First and foremost this will require us to develop standardized interfaces and APIs for collecting energy consumption data, just as PAPI has done for hardware performance counter data. Accurate and fine-grained measurement of power consumption underpins all tools that seek to improve such metrics (anything that cannot be measured cannot be improved). Secondly, we must use these tools to better understand the effects that energy saving hardware features have on the performance of linear algebra codes. Finally, we must identify parameters and alternative execution strategies for each numerical library that can be tuned for energy efficient executions, and to enhance our schedulers for better low-energy execution.

5.5.8 Sensitivity analysis

Many areas of modeling and simulation are still pushing to reach high-fidelity solutions to a given set of input conditions. However, as performance and fidelity improves, it becomes possible and imperative to study the sensitivity of a model to parameter variability and uncertainty, and to seek an optimal solution over a range of parameter values. The most basic form, the forward method for either local or global sensitivity analysis, simultaneously runs many instances of the model or its linearization, leading to an embarrassingly parallel execution model. The adjoint sensitivity method, with its powerful capabilities for efficiently computing the sensitivity of an output functional with respect to perturbations in a great many parameters, is a workhorse algorithm in weather prediction and in engineering design such as shape optimization. It requires the simulation of the forward and the adjoint problem; hence its parallelization will depend on the capability for highly efficient simulation.

5.5.9 Multiscale/multiphysics modeling

Engineering is increasingly operating at the micro- and nano-scales to achieve objectives at the macroscale. Models of these processes are intrinsically multiscale and multiphysics. For example, electrochemically reactive surfaces play a central role in the fabrication as well as the functional capabilities of an enormous variety of technological systems. Precise control of surface processes during fabrication is required in applications including on-chip interconnections between transistors, decorative and industrial coatings, batteries for electric vehicles, thin film photovoltaic solar devices, magnetic materials, and patterned deposits for sensors. Surface processes are occurring at the nano-scale and must be modeled by Kinetic Monte Carlo (KMC) methods, whereas reactions and diffusion in the electrolyte can be modeled by deterministic (PDE) methods. The two computations must be dynamically linked. Such a computation is very demanding and is currently consuming huge numbers of cycles on NCSA's supercomputers, with only modest resolution of the problem domain. Simulation is only the tip of the iceberg of this type of problem, where parameter estimation and optimal design are the ultimate goals and require orders of magnitude more computation time.

Cell biology is another area where processes operating at the microscale yield change at the macroscale (phenotypical change). In microscopic systems formed by living cells, the small numbers of some reactant molecules can result in dynamical behavior that is discrete and stochastic rather than continuous and deterministic. An analysis tool that respects these dynamical characteristics is the stochastic simulation algorithm (SSA), which applies to well-stirred (spatially homogeneous) chemically reacting systems. Usually, a large ensemble of SSA simulations is used to estimate the probability density functions of important variables in the system. This leads to an embarrassingly parallel implementation. At the same time, cells are not spatially homogeneous. Spatio-temporal gradients and patterns play an important role in many cellular processes. The modeling of stochastic diffusive transfers between subvolumes is an important challenge for parallelization.

5.5.10 Summary

The move to extreme-scale computing will require tools for understanding complex behavior and for performance optimization to be based on a knowledge-oriented process. Performance models and expectations will be used to drive knowledge-based investigation and reasoning. It will raise the level at which tools interoperate and can be integrated with the application development and execution environment. The challenges for performance analysis and tuning will grow as performance interactions and factor analysis must involve a whole system perspective.

The co-design methodology is iterative, requiring frequent interactions among hardware architects, systems software experts, designers of programming models, and implementers of the science applications that provide the rationale for building extreme-scale systems. As new ideas and approaches are identified and pursued, some will fail. As with past experience, there may be breakthroughs in hardware technologies that result in different micro and macro architectures becoming feasible and desirable, but they will require rethinking of certain algorithmic and system software implementations.

5.6 The Computer Science Challenges[72]

5.6.1 Overview

The coming transition in computer architectures as peak capability approaches the exascale offers both challenges and opportunities. The challenges involve a paradigm shift in programming methodologies. Existing technologies for writing parallel scientific applications have sustained HPC application software development for the past decade and have been successful for Petascale computing, but were architected for coarse-grained concurrency largely dominated by bulk-synchronous algorithms. Future hardware constraints and growth in explicit on-chip parallelism will likely require a mass migration to new algorithms and software architecture that is as broad and disruptive as the migration from vector to parallel computing systems that occurred 15 years go. The applications and algorithms will need to rely increasingly on fine-grained parallelism, strong scaling, and fault resilience. Addressing these challenges opens up a renewed opportunity to introduce a higher level of software engineering into current fusion application subsystems that will enhance the modularity, portability, and performance of codes while extending their capabilities to new levels. At the same time, past sound investments must be protected, and a migration path from current to future environments must be elaborated.

5.6.2 Programming Models

A programming model[73] effort is a critical component of a program to build effective exascale computing systems, because with clock speeds projected to be flat or even dropping to save energy, all performance improvements within a chip will come from increased parallelism. The amount of memory per arithmetic functional unit will drop significantly, implying the need for fine-grained parallelism and a programming model other than the currently used message passing or coarse-grained threads (e.g., PThreads/OpenMP). It would be premature to rule out any of the architectural models for increasing on-chip parallelism, yet history suggests that a programming model specialized to a single architecture is doomed to fail, and even if architectures become somewhat specialized to a class of applications, the programming model must be portable across all viable architectures. Thus, the exascale software program needs to allow researchers to pursue multiple hardware solutions, while programming models support a range of possible solutions.

The timing of the programming model effort is important: if the machines arrive with no viable programming model, it will significantly delay science impact or result in limited domains of impact; if the programming model is developed prior to the machine design and without regard to their features,

[72] Much of the material in this section has been taken from the Exascale Fusion Workshop Report, which had a specific computer science section, originally written by Kathy Yelick, V. Balaji, and Al Malony, as well as by Rusty Lusk.

[73] "Programming model" is used in a very general sense as a set of languages, libraries and tools used to develop application programs.

they may not be suitable to the hardware. Thus, co-design is essential – the programming model effort must be tightly coupled with multiple architecture efforts. The programming model developers need to have intimate knowledge of the proposed hardware designs and must be able to influence those designs. At the same time, the programming model development must be responsive to the needs of applications and must support the kinds of algorithms that will be used on exascale machines.

Challenges abound for programming model designers, including support for multiphysics applications, support for both fine-grained and course grained parallelism, interoperability, portability, scalability, latency-tolerance, support for energy efficient programming, the need for performance feedback, and reducing concurrency errors such as race conditions. Some of these challenges are specific to scientific programming or large-scale parallelism and some will cross over into more general programming areas. However, a program with too many goals and lack of prioritization is unlikely to succeed. Therefore, the DOE exascale effort should focus on the two most critical problems: support for fined-grained parallelism within a chip (including data parallel hardware, heterogeneous processors, and locality control) and support for fault-tolerant programming between chips, i.e., allowing programmers to write applications that tolerate hardware failures at the level of a single chip. The parallelism and scaling problems between chips are important, but the message-passing model as realized by the MPI interface and its implementations provide a viable solution that (with some investments) can be made to scale, whereas a viable programming model for massive on-chip parallelism does not exist. Similarly, there are only preliminary research efforts on fault tolerance programming, while current checkpoint approaches are likely to be inadequate, because the frequency of component failures in an exascale system may be close to the time to checkpoint an application to disk, in which case no forward progress is possible. Techniques to build redundancy into algorithms and software will become increasingly important and must be supported in the programming model.

To realize co-design, some application teams may work more closely with one hardware team or may work equally with multiple teams, but the programming model must be responsive to needs of all applications and hardware efforts. The programming model effort should include representatives from vendor teams, who may not have access to all of the hardware requirements from other teams but will see the interactions of their design decisions on the programming model design process. Application and algorithm experts should be represented as well, if not on a day-to-day basis then at regular intervals to influence the programming model team. While the programming model team should respond to requirements above and below, they may make conscious decisions against fully supporting a particular class of algorithms or hardware, i.e., some things may not run optimally if it would affect the viability of a coherent model for the more common applications.

5.6.3 I/O

The number of devices in exascale systems, both in terms of the number of clients and the number of devices in the storage system itself, creates distinct computer science challenges. The number and failure characteristics of these devices mean that exascale I/O systems must operate in a much more fault-prone environment than current HPC systems. Tolerating client failures and network data loss, and distributing data and reconstructing after loss, will be critical. Data placement will become more critical as well, including coordinating the placement of data with where in the system jobs are executed, as well as using mixed resources such as solid-state storage as part of the I/O system. Facilitating all this, new computational science data model support must provide a high-productivity interface for scientists and inform the storage system about the nature of the data it is storing.

5.6.4 Getting There from Here

Computational scientists have a large investment in currently successful codes. As new architectures require new programming approaches to fully exploit their power, a chicken-and-egg problem arises. If the new architectures only support a programming model radically different from current approaches, applications will not be able to convert their large codes in a timely manner. On the other hand, if only

current models are supported on the new machines, much of their capabilities will be wasted. The solution is for the new architectures, new programming models, and applications to go forward together: large application codes need an incremental migration path into the future. The new architectures must be capable of running existing scalable applications from the beginning of their deployment, even if at significantly less than peak performance, with only modest modifications to the code. Then tuning for performance can begin. For any radically new approach that will require significant code redesign and redevelopment, it is critical that an environment that adequately hosts the new programming model be available on current highly parallel machines.

What does this mean in the current situation? All scalable codes rely heavily on the message-passing model, expressed in the syntax and semantics of MPI, for their high-level structure. For some codes, this is not their highest-level structure, because they use libraries supporting a higher level of abstraction, but those libraries are in turn implemented in MPI. Some codes rely on multithreading within each MPI process to save memory; these codes are the ones most ready for the next generation of architectures. The multithreading is usually expressed using OpenMP, because the OpenMP and MPI standards work well together by explicit design. While MPI is a stable and (almost) adequate mechanism for expressing parallelism among separate address spaces, the choice of programming model for expressing parallelism within a single address space (shared-memory parallelism), whether expressed in a language or a library, is far less clear. OpenMP has the disadvantage that it is difficult to program for peak performance because of its lack of mechanism for expressing locality. Other shared-memory approaches, such as the Partitioned Global Address Space (PGAS) languages, may do better, but they need further development in the area of interactions with MPI. Similarly, brand-new approaches for programming heterogeneous processors will be of most use if they explicitly take into account that they may be used in the context of a larger distributed-memory computation. Very abstract parallel languages such as those developed as part of DARPA's High-Productivity Computing Systems (HPCS) project, may also play a role in a migration of existing applications to advanced architectures. Here the critical components are 1) robust, scalable implementations on current large machines, so that application may begin experimenting with them, 2) good implementations on new architectures, to demonstrate the performance potential of the new approaches, and 3) the capability of combining the new languages with existing ones, so that progress can be made incrementally.

Finally, MPI itself needs to be continually refreshed, in order to respond to the challenges of new, very large machines and to interoperate with new languages and libraries, since MPI is likely to remain a critical component of large-scale codes[74]. The MPI-3 Forum is currently addressing issues related to scalability and interoperability.

5.6.5 Tools

The promise of extreme-scale computing systems lies in the delivery of scaled simulations and other computations that optimize the performance potential of these machines. The processes of parallel program understanding and performance engineering are at the heart of achieving optimization goals. However, these processes and the tools applied to support them will need to change along with the programming models and frameworks in response to architectural and system evolution. Traditional techniques for parallel debugging, performance diagnosis, and tuning will become intractable as the factors of scale, software complexity, application sophistication, and hardware integration continue to increase. While interactions among these factors create the need to observe application performance across the whole system hierarchy, it is the model-oriented knowledge of the computational semantics of the application and of performance expectations with respect to peta/exa-scale systems architectures and

[74] Palaji, P, A Chan, R Thakur, W Gropp, & E Lusk. 2009. "Toward Message Passing for a Million Processes: Characterizing MPI on a Massive Scale Blue Gene/P." *Proceedings of the International Supercomputing Conference*.

capabilities that ultimately must be incorporated into tools to better focus and automate correctness/performance problem identification and guide tuning decisions.

A fundamental challenge for tools to understand application program behavior and performance is the need to observe parallel application execution in the target system environment. In general, the more that can be learned from the execution via measurements, the higher value of information a tool will have to assess behavior and problems. However, scale complicates observation and analysis by amplifying the amount and complexity of measured data, as well as the effects of measurement on the execution. Whereas these concerns can be addressed in part by enhancing current petascale tools with more scalable infrastructure and analysis methods, extreme scale will ultimately force a more intelligent methodology to optimize observation value versus measurement cost. It is important that such an approach be rooted in knowledge about the application – its structure, computational model/domain, algorithms – so that what is observed can be related back at a higher level to application-specific concerns.

The problem is more than just a matter of scale. Application behavior and, in particular, performance, will be determined by a complex interplay of the program code, processor, memory, interconnection network, and I/O operation. Achieving extreme scale performance requires an optimized orchestration of these components and a whole system view in order to understand root causes of inefficiencies. While measurement of code execution will be more difficult due to increased processing heterogeneity, there has been very little heretofore in tools to observe other system aspects. Multi-level observation will be needed to understand behavior and performance *in toto*, and this support must be integrated with the runtime environment, operating system, I/O system, and even the machine hardware.

Greater levels of multi-core parallelism and heterogeneous processors will additionally constrain observation because certain events of interest are less (not) visible and harder to measure. More highly integrated accelerator devices might not allow any access to internal parallel operation. These limitations together with the sheer massive parallelism will make it impossible to observe all concurrent operations in the system. Furthermore, the increased importance of the memory system for extreme scale will expand observation focus to understanding data transfer behavior. Measurement infrastructure is woefully lacking for these tasks.

5.6.6 Fault Tolerance

The next step in extreme-scale computing is not a mere scaling up of solutions based on lessons learned at tera- and petascale. Some of the components out of which such systems will be constructed are operating at or near the limit that they can hope to reach (e.g., processor size and clock speed) and it is clear that we can expect a dramatic increase in the number of constituent components in an exascale system. It is becoming clear that an exascale computational system will comprise a massively parallel fabric connecting a million or so computational nodes; each will be able to handle 100-1000 concurrent functional units; and each will contain a multi-level memory hierarchy globally consisting of a million or so memory modules. The corresponding exascale storage system will consist of file systems striped across millions of spinning disks.

To underline the need for a pervasive fault resilience approach encompassing every stage, imagine a test computation on such an exascale platform, with a known result. The computation is executed on the many nodes of this massively parallel system, touching millions of memory modules, communicating across a million sockets, executing on hundreds to thousands of functional units at each node, and finally writing the result to exascale storage striped across a million rotating disks. At this point, if an unexpected result is reported, it is quite certainly too late for any remedial measures. A pervasive approach intervenes at every stage in this process and must seamlessly handle faults in hardware and software.

Addressing fault tolerance at individual points in the system will be insufficient unless it is addressed at other points as well. Fault resilience will be needed in the memory subsystem, in storage media, in the

communication fabric, and on the computational node itself. Software that integrates multiple types of fault notifications and responses from multiple parts of the system is in its infancy.

5.7 The Educational Challenges

Major challenges in exascale science include the building of understanding and awareness among groups with high prestige in both academia and industry, and the dearth of highly competent young scientists in this field, two issues that are not entirely unrelated. Many of the reasons for these problems are reasonably well understood, but not easily dealt with. Two of the most important reasons (which are indeed related) include:

1. Migration of talented computer science graduate students to research areas not directly related to computational science. This migration has been underway for well over a decade, and is driven not only by disparities in financial rewards, but also by the growing academic disconnect between computer science on the one hand and applied mathematics and science/engineering disciplines on the other.
2. A lack of respect and appreciation in all directions among groups whose contributions are needed to produce state-of-the-art tools for exascale science. Application scientists who focus primarily on building computational tools are sometimes regarded by their scientific community as not being "real" scientists. This phenomenon is particularly noticeable in both physics and chemistry, reflecting in part the penetration of "community codes". From the opposite perspective, high-level software designers and programmers may not welcome or appreciate the contributions made by scientific disciplines to building state-of-the-art computational tools.

DOE's SciDAC program has helped enormously to address the second problem among people associated, directly or indirectly, with DOE. More broadly, however, these problems involve the difficult-to-change attitudes of academic communities.

On the bright side, interest in computational science and engineering (CSE) worldwide has measurably increased during the past 15 years. For example, the Society for Industrial and Applied Mathematics (SIAM) has a thriving activity group that holds popular research conferences every year, and the SIAM Journal on Scientific Computing recently reorganized its structure so that there is an explicit section devoted to CSE. In addition, the extremely successful DOE Computational Science Graduate Fellowship (CSGF)[75] program has produced, over nearly 20 years, a cadre of outstanding students, fulfilling its mission of building a community of leaders in computational science. The CSGF program has been successful for several reasons: a distinguished committee of scientists selects the winning fellows and reviews them individually every year until completion of the PhD; CSGF fellows have taken positions at DOE labs, in industry, and in high-prestige academic departments, in the last case conveying an appreciation of DOE to their students; and a light touch is exercised to generate student involvement with DOE problems through a summer practicum at a DOE lab.

DOE should continue to support programs, with CSGF the prime example, which visibly involve participation by top academics in applied mathematics, computer science, and application disciplines. Researchers tend to retain attitudes engendered in graduate school by their most respected professors, even when those attitudes are mistaken or out of date. DOE should therefore think of creative new ways to make leading academics aware of the intellectual challenges associated with exascale science. This could start, for example, by convening a distinguished committee charged with defining explicit strategies.

The two difficulties mentioned earlier do not constitute the only problems facing education and training

[75] Note that the CSGF program should not be confused with the much more recent Office of Science Graduate Fellowship Program [SCGF].

related to exascale science. Almost no universities, even those with faculty working on computational science and engineering, have, or are likely to develop, a curriculum that focuses on issues associated with exascale science. In addition, as our subcommittee has noted already, many of the issues in exascale science are not yet understood, which means that a straightforward program of training in the usual sense is impossible. Exascale hardware and its features will keep changing, so that training people too early to think about specific hardware configurations is a bad idea. However, it is important to start soon to lay the foundations for future thinking about exascale science.

An appealing strategy for DOE would be a program consisting of (1) workshops for researchers, with widespread dissemination of high-level summaries of the latest trends; (2) short courses for those who will need to use, or who want to understand, exascale computing; and (3) preparation of web-based materials and demonstrations suitable for varying levels of expertise. Given the unavoidable disruptive shifts in practice when moving to exascale, as well as the large-scale waste of time and money from lack of preparation to use exascale computing efficiently, our subcommittee recommends that DOE make a generous investment in education and training. "Generous" here means that it should be visibly worthwhile for the best researchers and educators to spend their time designing curricula, preparing materials, and acting as instructors. Substantial use should be made of web-based instruction, as in the parallel computing course at Berkeley's Par Lab (see below).

Before embarking on such a program, investigation should be undertaken of the experiences (positive and negative) of previous programs of a related nature, such as the Education, Outreach, and Training program of NSF's Partnerships for Advanced Computational Infrastructure (EOT-PACI). Another, more recent, example, is the Par Lab at UC Berkeley, a multidisciplinary research project funded by Intel and Microsoft, which hosted a short course in 2009 on parallel programming, with participants both at Berkeley and online.[76]

To be successful, an exascale science education and training program needs to be devised and managed with creative flair, not business as usual. One ideal feature, which is true of the CSGF program, is that leading researchers would be involved in developing the program structure and contributing to its content (as with the Par Lab short course). A second desirable feature, again in the mold of CSGF, is that a DOE exascale science education program emphasizing primarily the excitement of the science could capture the interest of undergraduate and even high school students throughout the country, attracting them to an interest in DOE applications.

6. Summary

In this report, we have focused first on answering a central question: Are there scientific disciplines in whose domain exascale computing may arguably lead to transformational changes in these disciplines? We believe we have answered this question in the affirmative: From biology to nuclear engineering, computing at the exascale promises dramatic advances in our capabilities to model and simulate complex phenomena, at levels of fidelity that have the potential to dramatically change both our understanding and our ability to comprehend. Thus, there are almost certain to be great benefits to going to the exascale. Next, we turned to the questions directly related to "going to the exascale": What will be involved, and what will be the challenges? And how difficult will the transition from the present era of tera- and peta-scale computing to a new era of exascale computing be? As our discussion in Section 5 should make plain, the challenges to be overcome are both numerous and difficult; and one cannot diminish the consequent needs for substantial investments in research and development, across the board – from the details of the hardware, to fundamental algorithms and programming models, to error and fault detection and management at both the hardware and software ends, to I/O and checkpointing, and finally to the

[76] See parlab.eecs.berkeley.edu/2009bootcamp

tools that make it possible to effectively compute: the (optimizing) compilers, the load balancers, the performance diagnostic tools, and of course the programming languages themselves.

These challenges couple hardware design and computing itself intimately, to a degree not seen since the early days of computing; and for this reason, the notion of 'co-design' – discussed in detail in Section 5 – should not be taken lightly as simply the latest fashion, but rather as a 'call to arms' to the applied math, the computational and application sciences, the computer science, and the hardware community (including the hardware vendors) to closely couple their efforts as we enter this new era. It is to be especially noted that we are talking about a transition that will take us into the latter part of this decade, and the early part of the next – and this means that this 'call to arms' needs to be especially noted by our academic colleagues who are now training the researchers of the coming decade, as well as the institutions that provide the support for this education.

We finally note that computing in general, and leadership computing in particular, is a key element in the U.S. national strategy for industrial competitiveness. In a number of the application areas discussed in our report, the international nature of competition contributes an important additional element of urgency: There is a concerted push by both Europeans and Asians in the direction of high-performance computing, and in its industrial applications; and it behooves us to recall that our past preeminence in this realm has been an important component of our past successes in making the U.S. a leading manufacturing and exporting nation.

Appendix 1: The Subcommittee Charge Letters

Department of Energy
Office of Science
Washington, DC 20585
October 29, 2009

Office of the Director

Professor Roscoe Giles, Chair
Department of Electrical & Computer Engineering
Boston University
8 St. Mary's Street
Boston, MA 02215

Dear Professor Giles:

Over the last few years, several workshops and subcommittee reports have identified and described the scientific opportunities for high performance computing. By this letter I am charging the Advanced Scientific Computing Advisory Committee (ASCAC) to assemble a sub-committee to look at the results of these activities and to analyze the opportunities and challenges for the Office of Advanced Scientific Computing Research (ASCR) and the Office of Science associated with exascale computing. Specifically, I would like the sub-committee to deliver a report that:

- Assesses the opportunities and challenges of exascale computing for the advancement of science, technology, and Office of Science missions.
- Identifies strategies that ASCR can use to address the challenges and deliver on such opportunities.

We would appreciate the committee's preliminary comments by March 30, 2010, and a final report by August 15, 2010. I appreciate ASCAC's willingness to undertake this important activity.

If you have any questions regarding this matter, please contact either Michael Strayer, the Associate Director of the Office of Science for ASCR, or Christine Chalk, the Designated Federal Official for the ASCAC.

Sincerely,

W. F. Brinkman
Director, Office of Science

Under Secretary for Science
Washington, DC 20585

November 2, 2009

Professor Roscoe Giles
Chair, Department of Electrical & Computer Engineering
Boston University
8 St. Mary's Street
Boston, MA 02215

Roscoe

Dear Prof. Giles,

As your subcommittee begins to work on Dr. Brinkman's the new charge, I would like for you to consider the following.

Modeling and simulation at the extreme scale have the potential to span the entire Department and to forge lasting interlocks between applied and basic science, technology, and engineering. I would therefore ask the sub-committee to acknowledge and comment upon the broader issue of opportunities and challenges for exascale computing to advance *Department of Energy* missions. I can imagine that will entail including among the subcommittee members familiar with both NNSA and various applied programs that are amenable extreme scale computing.

We live in times when the Nation is faced with important issues involving energy, environment, and national security, yet resources are constrained. It is therefore very important to have clear justification for how new endeavors in extreme scale computing (ultimately reaching the exascale) will affect science, technology, and society.

I look forward with great interest to your report.

Sincerely,

Steven E. Koonin
Under Secretary for Science

Appendix 2: Previous Workshop Reports on Exascale Computing

The following exascale computing-focused workshop reports, sponsored by the DOE Office of Advanced Scientific Computing Research (OASCR) and the 'applications', i.e., the various scientific and engineering disciplines within the Department of Energy that draw on advanced modeling and simulations, provide a detailed accounting of the issues discussed in this Report. Much of our material has been drawn from thèse workshop reports, which form the formal documentary basis for our Committee's délibérations.

1. **Modeling and Simulation at the Exascale for Energy and the Environment Town Hall Meetings Report** Workshop report from the Grand Challenges Series of Workshops at Lawrence Berkeley National Laboratory, April 17-18, 2007.
2. **Challenges in the Climate Change Science and the role of Computing at the Extreme Scale** Workshop report from the Grand Challenges Series of Workshops in Washington DC November 6-7, 2008.
3. **Challenges for the Understanding the Quantum Universe and the Role of Computing at the Extreme Scale** Workshop report from the Grand Challenges Series of Workshops at SLAC in Menlo Park, CA, December 9-11, 2008.
4. **Forefront Questions in Nuclear Science and the Role of Computing at the Extreme Scale** Workshop report from the Grand Challenges Series of Workshops in Washington DC, January 26-28, 2009.
5. **Fusion Energy Sciences and the Role of Computing at the Extreme Scale** Fusion Energy Workshop report from one of the Grand Challenges Series of Workshops in Washington, DC, March 18-20, 2009.
6. **Science Based Nuclear Energy Systems Enabled by Advanced Modeling and Simulation at the Extreme Scale** Workshop report from the Grand Challenges Series of Workshops in Washington DC May 11-12, 2009.
7. **Opportunities in Biology at the Extreme Scale of Computing** Workshop report from the Grand Challenges Series of Workshops in Chicago, IL, August 17-19, 2009.
8. **Scientific Grand Challenges: Exascale Workshop Panel Meeting Report** Workshop report from a panel of 12 scientists and engineers with experience in government, universities, national labs and industry. The panel met on January 2010 in Washington, D.C. to review a collection of reports prepared to document the need for a new generation of extreme-computing capability for the DOE's missions.
9. **Scientific Grand Challenges in National Security: The Role of Computing at the Extreme Scale** The NNSA workshop report from the Grand Challenges Series of workshops, held October 9-12, 2010 in Washington D.C.

In addition, several other reports – separate from the above workshop reports – focus on various aspects of exascale computing relevant to our délibérations, and are listed below.

1. **Visualization and Knowledge Discovery: Report from the DOE/ASCR Workshop on Visual Analysis and Data Exploration at Extreme Scale** The developed visualization and data exploration tools have served admirably with gigabyte and even terabyte datasets, but at the peta- and exascale levels, those tools will no longer suffice. Scientists and researchers met under the auspices of ASCR in Salt Lake City on June 7-8, 2007 to discuss the coming "data tsunami" and issues involved in data exploration, data understanding, and data visualization at the petascale and beyond.
2. **Advanced Scientific Computing Research: Delivering Computing for the Frontiers of Science - Facilities Division Strategic Plan for High Performance Computing Resources** The strategic vision for High Performance Computing (HPC) resources in the Facilities Division of the Office of Advanced Scientific Computing Research (ASCR) program in the Department of Energy's Office of Science for the next 10 years, archived August 13, 2007.
3. **Computational Research Needs in Alternative and Renewable Energy** Final report from the workshop on Computational Research Needs in Alternative and Renewable Energy held in Rockville, Maryland September 19 and 20, 2007.
4. **Scientific Impacts and Opportunities in Computing Workshop Report** The workshop was conducted for the Office of Advanced Scientific Computing Research to identify high impact opportunities in computing for investment in research to maintain the nation's preeminence in scientific discovery and competitiveness (May 2008).
5. **International Exascale Software Project Roadmap.** This draft report (posted at the DOE/OASCR web site on 1/27/10) provides a technology roadmap developed over the course of roughly one year by the international software community for high end scientific computing.