Combustion Exascale Co-Design Center

Director: Jacqueline Chen (SNL)

Deputy Director: John Bell (LBNL)

Research Area Leads: Curtis Janssen (SNL), Scott Klasky (ORNL), Patrick McCormick (LANL), Robert Moser (University of Texas at Austin), Daniel Quinlan (LLNL), John Shalf (LBNL)

Goals and Objectives

There is an urgent need to increase our understanding and ability to reliably predict combustion phenomena. The U.S. is at a critical crossroads where urgent issues in energy security, climate change, and economic competitiveness are converging. Aggressive federal mandates for reducing greenhouse gas emissions by 80 percent by 2050 will require the design of new high-efficiency, low-emissions combustion systems for transportation, power generation and industrial processes. Further mandates to reduce our petroleum usage by 25 percent by 2020, have prompted research and development in developing alternative carbon neutral renewable fuels and placed additional requirements for fuel flexibility on new combustion systems.

The goals of this project are to develop an exascale combustion modeling capability that combines simulation and analysis, to develop the necessary computer science tools to facilitate the development of these applications, and to quantify hardware constraints for an effective system. Exascale computing will enable combustion scientists to perform first principles direct numerical simulations with sufficient physics fidelity in appropriate flow regimes to answer fundamental questions in order to meet pollutant and greenhouse gas emissions targets, reduce dependence on petroleum and promote economic competitiveness.

Co-Design Strategy

Current evolutionary trends in computer architecture design are not sustainable. Extrapolating current technology trends leads to machines that have unrealistic power requirements, are intractable to build and maintain, and are too expensive. Power and cost constraints lead to billion-way concurrencies and many-core architectures with reduced memory per core and high costs for data movement. Effective utilization of these types of architecture for combustion applications requires that we fundamentally rethink how we study combustion computationally. We need to take a holistic approach, i.e. simultaneously consider the algorithms used to model the relevant physics, the programming models used to implement those algorithms and how different architectural features affect the performance of the methodology. Key issues that need to be addressed include:

- Significant growth in explicit parallelism within a hierarchical many-core architectural framework.
- Reduced memory per core and an increase in the cost of data movement in terms of both performance and power relative to floating point operations.
- A growth in machine complexity to the point that fault tolerance must be an essential component throughout the software stack.
- An increased disparity between I/O speed and compute speed.

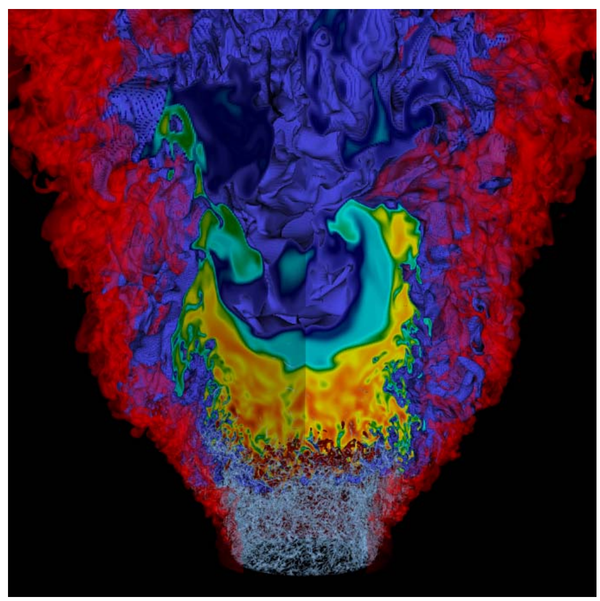

Figure 1 NO_x emission from a high-fidelity simulation of a low swirl injector fueled by lean premixed hydrogen.

Current experience with petascale combustion simulation provides some guidance on the structure of an exascale combustion simulation capability. Analysis of spatial resolution requirements and temporal scales indicates that the core simulation methodology needs to support custom formulations designed for specific flow/flame regimes and support adaptive mesh refinement. Both fully compressible and low Mach number formulations using a block-structured adaptive mesh refinement approach will be supported. The block-structured paradigm provides a natural model for hybrid parallelization models.

Combustion simulations generate enormous amounts of data. At the exascale, current approaches based on writing out data for subsequent analysis will not be viable. However, the data are also highly complex with a large dimensional "phase space" corresponding to the detailed chemical description of the problem. Thus, a key element of combustion co-design is the development of an *in situ* data management and analysis system that enables a significant reduction in the amount of output data while preserving the ability to explore detailed relationships associated with turbulence chemistry interactions.

Finally, as we move toward a predictive capability for turbulent combustion, we need to embed uncertainty quantification as an integral part of the simulation. We need to be able to assess how uncertainties in chemical parameters affect simulation results. We also need to be able to solve inverse problems to identify ranges of parameters that are consistent with experimental data. For turbulent combustion, the expense of the simulations makes simple "black box" approaches infeasible. As part of this project we will investigate several alternative strategies for embedding uncertainty quantification into the simulations and enabling uncertainty calculations through co-design hardware and software elements.

These three areas--basic simulation methodology, *in situ* analysis and uncertainty quantification--form the main themes of the combustion exascale co-design center. The center is organized around vertically-integrated co-design efforts in each of these areas drawing on computer science, applied math, and combustion science expertise. In each case, we will examine core algorithms and data structures, investigate alternative programming models to express the algorithms and explore how the algorithms map to alternative architectures using a mixture of performance modeling, autotuning and hardware simulation. For each of these areas, our strategy is to examine the problems at a range of granularities ranging from global skeletal app's that provide an overview of data movement and flow control to compact app's that represent critical computational components of the algorithm to key computational kernels characteristic of the algorithm. Capturing the principal features of combustion simulation across this range of granularities enables us to consider global network topology between nodes, data movement and heterogeneity within a node and core-level performance issues. This combined perspective also ensures coherence in the overall software design.

Although combustion is an ideal candidate for exascale computing, new approaches in terms of both algorithms and programming models are needed to effectively utilize emerging exascale architectures. Furthermore it is only through integrated evaluation of these approaches with different architectural variations that we will be able to identify viable hardware designs.

Specific Research Topics

Here we identify some of the specific research topics within our main thematic areas in more detail.

<u>PDE Methodology</u>:

- Examine basic PDE discretization approaches from the perspective of minimizing memory requirements and data movement instead of optimizing floating point efficiency
- Investigate hierarchical scalable approaches to AMR that minimize global communication and data movement
- Investigate alternative programming models that facilitate expressing parallelism and controlling data layout in structured grid discretizations
- Explore the effect of architecture variation on performance of stencil operation kernels using a combination of modeling, autotuning and hardware simulations

<u>In Situ Data Analysis</u>:

- Analyze strategies for integrating analysis and data reduction techniques with simulations
- Investigate programming models for topological analysis, visualization and data reduction techniques
- Investigate hybrid staging areas and asynchronous execution models for optimizing analytics and reducing energy consumption
- Identify key architectural features needed for efficient execution of core analysis kernels

<u>Uncertainty Quantification</u>:

- Investigate data structures and data management approaches needed to integrate UQ into simulation framework.
- Explore different data-staging approaches to facilitate computation of derivative information from time-dependent simulations
- Investigate *in situ* analytics support needed for integrated UQ
- Explore potential hardware support needed for instrusive UQ algorithms such as polynomial chaos expansions

Contacts:

Jacqueline Chen (jhchen@sandia.gov)

John Bell (jbbell@lbl.gov)