Present and Future Computing needs in Simulation and Analysis of Reacting Flows

John B. Bell

1. Reacting Flow Simulation Overview

John Bell (PI) / LBNL

- Integrated approach to algorithm development for multiphysics applications
  - Mathematical formulation – exploit structure of the problems
  - Discretization – match discretization to mathematical properties of the underlying processes
  - Software / solvers – evolving development of software framework to enable efficient implementation of applications
  - Prototype applications – real world testing of approaches
- Application areas
  - Combustion
  - Subsurface flow
  - Astrophysics

1. Reacting Flow Simulation Overview – cont'd

John Bell (PI) / LBNL

- Combustion
  - Key collaborators: M. Day, M. Lijewski, R. Cheng, P. Glarborg, C. K. Law
  - Resources: NERSC + INCITE (OLCF)
  - Current focus – gaseous flames
    - Emissions from low swirl burner
    - Detailed categorization of flame dynamics as a function of turbulent intensity in an idealized configuration
    - High-pressure flames in an idealized configuration
  - Future focus – gaseous flames
    - Flame behavior in low swirl burner at high pressure
    - Emissions at high pressure
    - More complex fuels

1. Reacting Flow Simulation Overview – cont'd

John Bell (PI) / LBNL

- Subsurface flow
  - Key collaborators: M. Day, M. Lijewski, K. Pruess, ASCEM
  - Resources: NERSC
  - Current focus – incompressible models
    - Simplified model for mixing at CO_2 / brine interface
    - Reactive transport in vadose zone
  - Future focus – more thermodynamically realistic models
    - Compressible multiphase, multicomponent with mass transfer between phases
    - Non-isothermal effects
    - Multiscale models for heterogeneous subsurface flow

1. Reacting Flow Simulation Overview – cont'd

John Bell (PI) / LBNL

- Astrophysics
  - Key collaborators: A. Almgren, P. Nugent, S. Woosley, M. Zingale, A. Burrows, A. Heger, D. Kasen, J. Neimeyer, W. Schmidt, M. White, …
  - Resources: NERSC + INCITE (OLCF)
  - Current focus
    - "Standard" SNIa model
      - Convection leading up to ignition
      - Deflagration
      - Detonation
    - Cosmology
      - Dark matter + hydro + self-gravity simulation of Lyman-alpha forest
  - Future focus
    - Alternative SNIa models
      - Merging white dwarfs
      - Sub-Chandra models
    - Cosmology simulation with higher fidelity physics (detailed hydrogen and helium physics), radiation transport and models for structure formation

2. Current HPC Methods -- Overview

- Core algorithm technology
  - Finite volume discretization methods
  - Geometric multigrid
  - Block-structured AMR
- Implemented in BoxLib framework
  - Class structure to support development of structured AMR methods
  - Manages data distribution and load balancing
  - Efficient metadata manipulation
- Hybrid parallelization strategy
  - Distribute patches to nodes using MPI
  - Thread operations on patches using OpenMP

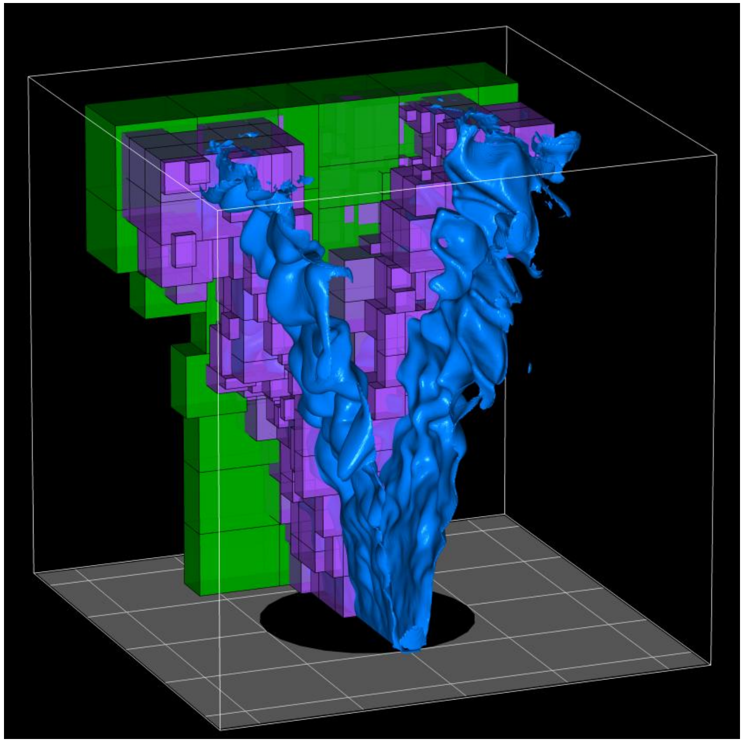

2. Current HPC Methods -- Combustion

- Formulation
  - Low Mach number model derived from asymptotic analysis
  - Removes acoustic wave propagation
  - Retains compressibility effects due to thermal processes
- Numerics
  - Adaptive projection formulation
  - Operator split treatment of chemistry
  - Dynamic estimation of chemistry work for load balancing
- LMC

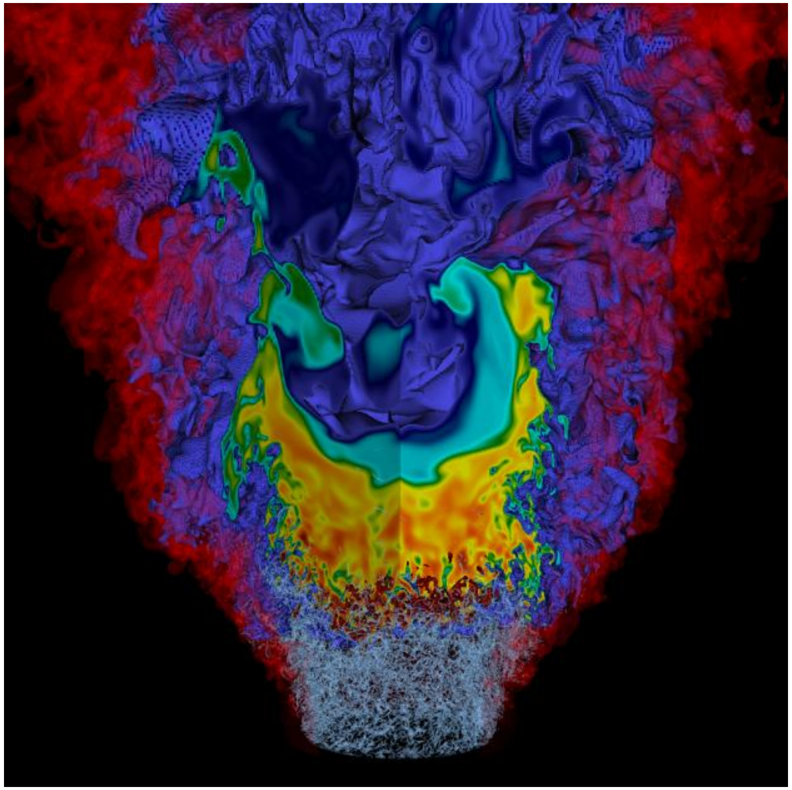

Simulation of NOx emissions in a low swirl burner fueled by hydrogen. Effective resolution is 4096^3

2. Current HPC Methods – Subsurface Flow

- Fractional step scheme
  - Total velocity formulation
  - Decouple parabolic pressure from essentially hyperbolic component conservation equations
- Numerics
  - Godunov methods for nonlinear hyperbolic component conservation equations
  - Operator split treatment of geochemistry
  - Integrated into AMR framework
- PMAMR

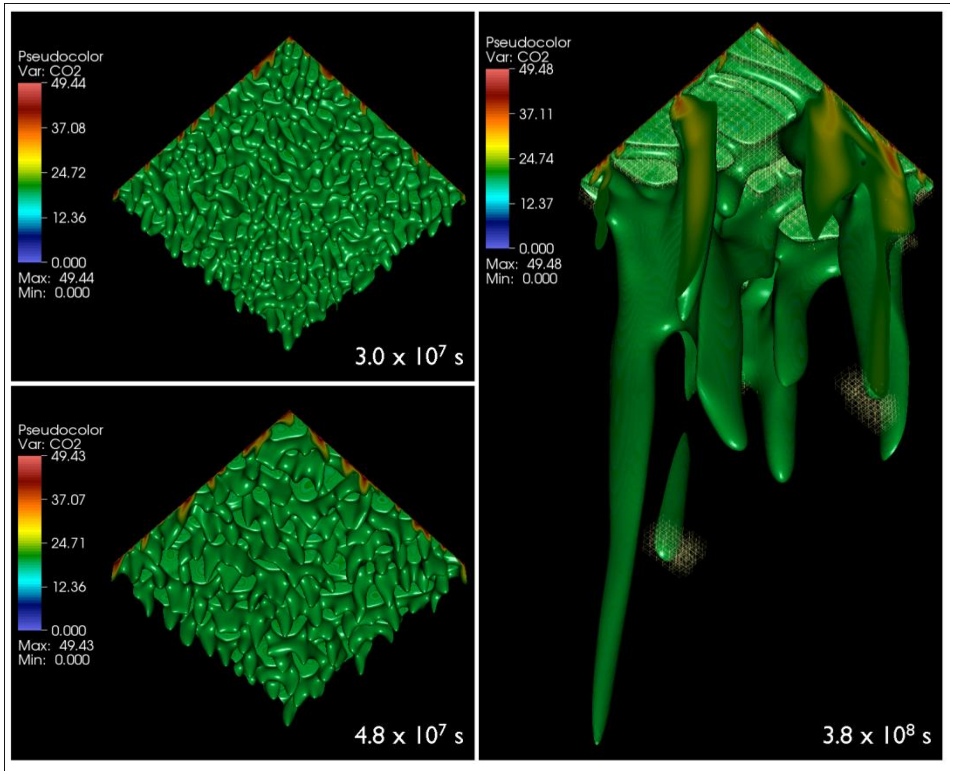

Simulation of mixing-induced by diffusion of CO2 into brine. AMR used to resolve the diffusive length scale

2. Current HPC Methods -- Astrophysics

- Formulation
  - MAESTRO
    - Low Mach number formulation
    - Removes acoustic wave propagation
    - Retains compressibility effects due to stratification and thermal processes
  - CASTRO
    - Compressible flow formulation
    - Self-gravity
    - Models for turbulent flame propagation
  - NYX
    - CASTRO + collision-less particles to represent dark matter

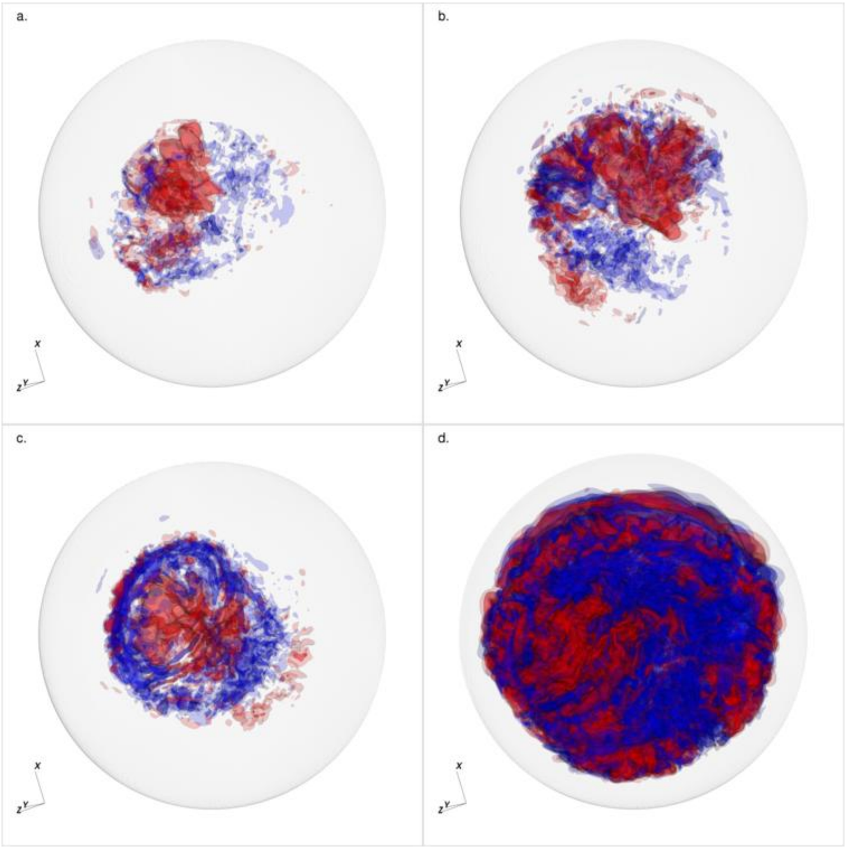

Simulation of advection leading up to ignition in a Chandrasakar white dwarf

- Numerics
  - Unsplit PPM
  - Multigrid
  - AMR

2. Current HPC Requirements

- Architectures currently used
  - Cray XT4,XT5,XE6
  - IBM power 6
  - Linux clusters
- Compute/memory load
  - Current simulations range from 2K – 24K cores
  - 1K – 10K GB
  - Minimum 0.5 GB / core
- Data read/written
  - 2 – 60 TB per run written
  - 100 – 500 GB checkpoint files
- Necessary software, services or infrastructure
  - C++ / F90
  - MPI + OpenMP (working together)
  - Htar
  - VisIt
- Known limitations/obstacles/bottlenecks
  - Hierarchical parallelization model scales to 200K processors
- Hours requested/allocated/used in 2010
  - Total usage during 2010 was approximately 40M hours. 10M at NERSC

3. HPC Usage and Methods for the Next 3-5 Years

- Upcoming changes to codes/methods/approaches
  - Development of higher-order methods
  - Development of *in situ* analysis methodology
- Changes to Compute/memory load
  - Anticipate need for 300M hour / year in 3-5 year (100+ M next yr)
  - 25K – 200K cores
  - 10K – 100K GB; 0.5 / core; could be reduced
- Changes to Data read/written
  - 10-100 TB
- Changes to necessary software, services or infrastructure
  - Need improved programming model
  - Fine-grained parallelism
    - Low overhead thread support
    - Support to express data locality
  - Support for *in situ* data analysis
- Anticipated limitations/obstacles/bottlenecks on 10K-1000K PE system.
  - Current approach should be extensible to 1000K cores
  - Improvements to programming model would help
- Key point is to directly link upcoming NERSC requirements to science goals
  - Limit of memory per core -> 0 hard to define
  - Need to maintain a reasonable level of memory per node (8 GB)

Strategy for New Architectures

- How are you dealing with, or planning to deal with, many-core systems that have dozens or hundreds of computational cores per node?
  - Block-structured AMR provides a natural model for hierarchical parallelism
    - MPI to distribute large patches to nodes
    - Thread based-model for fine-grained parallelism within a node
    - Need improved programming models

- How are you dealing with, or planning to deal with, systems that have a traditional processor augmented by some sort of accelerator such as a GPU or FPGA or similar?
  - Preliminary exploration of CUDA and OPENCL
  - Current programming model seems somewhat limited for our type of application
  - Expectation (hope) is that architecture will move toward heterogeneous models with "first-class" cores

4. Summary

- What new science results might be afforded by improvements in NERSC computing hardware, software and services?
  - Answered above
- Recommendations on NERSC architecture, system configuration and the associated service requirements needed for your science
  - Maintain system balance as much as possible
  - Keep (at least) memory per node fairly large
  - Aggressively pursue new programming models to facilitate intranode, fine-grained parallelization
  - Aggressively pursue programming model support *for in situ* analysis
- NERSC generally acquires systems with roughly 10X performance every three years. What significant scientific progress could you achieve over the next 3 years with access to 50X NERSC resources?
  - Major opportunity to integrate UQ with high-fidelity simulations
  - Combustion
    - High pressure flames for complex fuels
    - Develop models for engineering design
    - Assess role of uncertainty in rate data on fidelity of simulations; use experimental data to improve knowledge of rates
  - Subsurface flow
    - Quantify role of subsurface uncertainty of prediction of contaminant plumes and carbon sequestration scenarios
  - Astrophysics
    - High-fidelity simulation of a variety of supernova phenomena and comparison with observation
    - Detailed cosmological simulations compared to observation to reduce uncertainty in key cosmological parameters