LLNL-PRES-????
SANDSAND2011-8531C

Computational Science Panel
Grace Hopper Celebration of Women in Computing

Carol S. Woodward
Center for Applied Scientific Computing
Lawrence Livermore National Laboratory

Janine C. Bennett
Scalable Modeling and Analysis Systems
Sandia Livermore National Laboratories

11/11/11

The mission of the Energy Department is to ensure America's security and prosperity by addressing its energy, environmental and nuclear challenges through transformative science and technology solutions

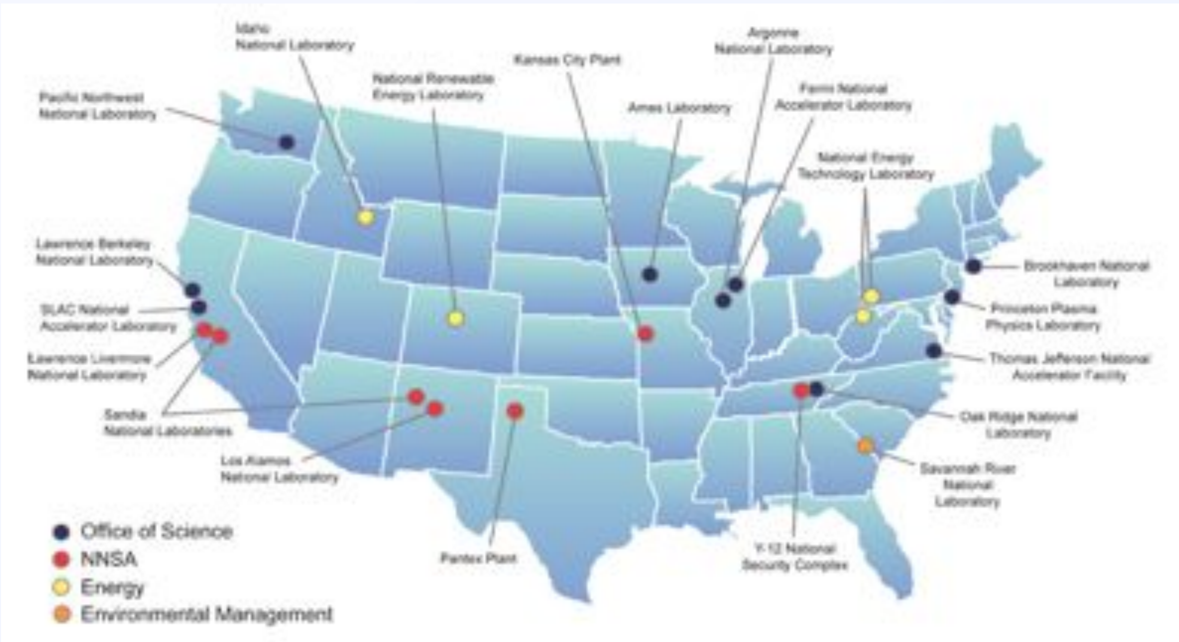

- Maintain a vibrant U.S. effort in science and engineering as a cornerstone of our economic prosperity
  - Biological science
  - Chemical science
  - Computing
  - Environmental science
  - Materials science
  - Physics
- Enhance nuclear security through defense, nonproliferation, and environmental efforts

Two main types of DOE labs:

- National Nuclear Security Administration: Lawrence Livermore, Los Alamos, Sandia [2 locations], …
- Office of Science: Argonne, Oak Ridge, Lawrence Berkeley, Pacific Northwest, …

Many important U.S. national missions depend on high-performance computing and simulation

- Nuclear and conventional weapons
- Global and regional climate
- Nonproliferation
- Counter-terrorism
- Cryptanalysis
- Energy Security
- Drug design
- Engineering design for economic competitiveness
- Data management for massive experiments
- Cyber security

Supercomputing is key to the U.S. National Nuclear Security Administration's *Advanced Simulation and Computing* Program

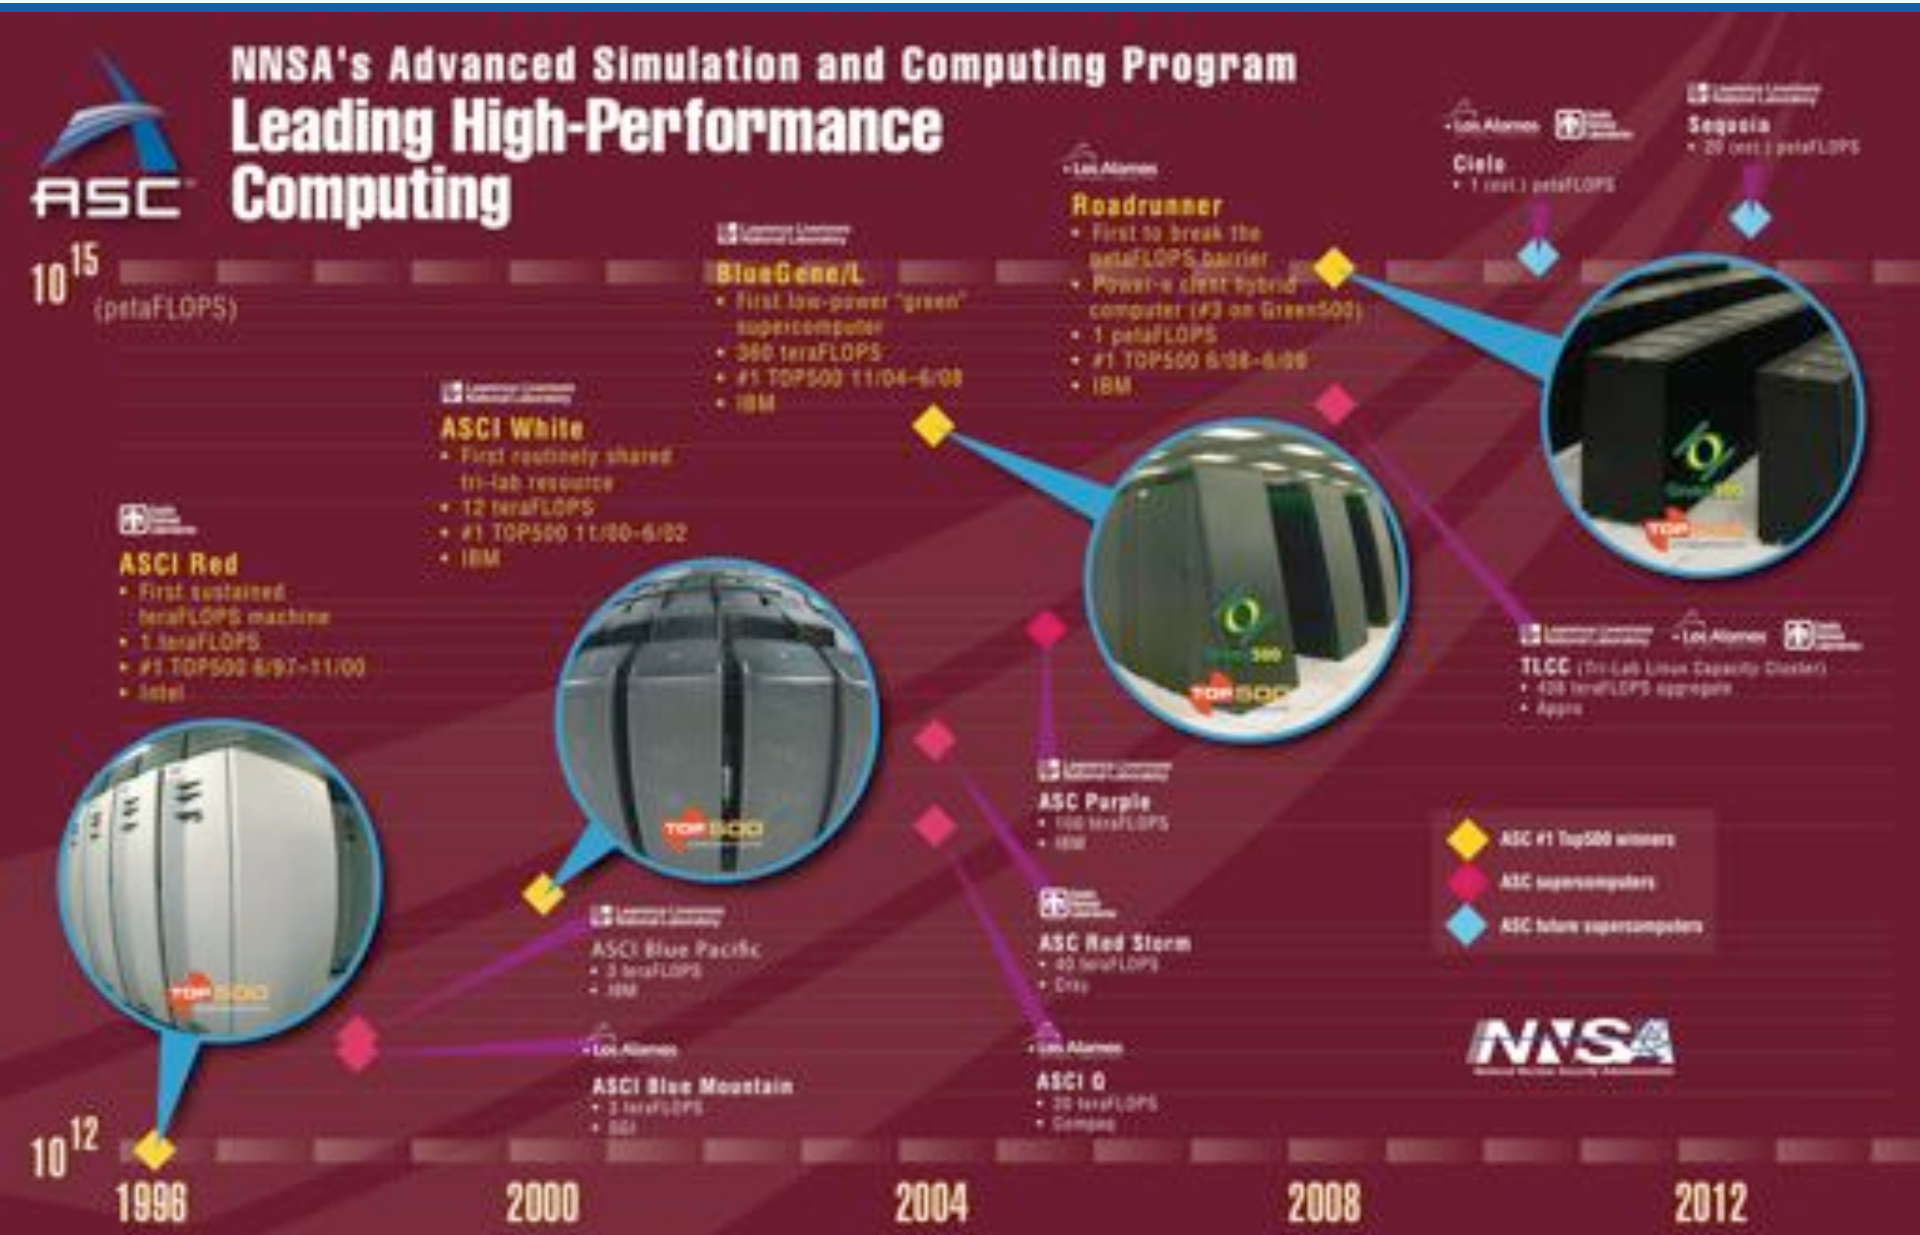

DOE Office of Science Leadership Class Facilities look to the future of HPC for breakthrough scientific discovery

Argonne Leadership Class Facility

- Intrepid (IBM BG/P)
- Surveyor (IBM BG/P)

Oak Ridge Leadership Class Facility

- Jaguar (Cray XT-5/XT-4)
  - *TOP500 #1 at 2.3 PF Peak*

LBNL Large-Scale Computing Systems

Franklin (NERSC-5): Cray XT4

- 9,532 compute nodes; 38,128 cores
- ~25 Tflop/s on applications; 356 Tflop/s peak

Hopper (NERSC-6): Cray XE6

- Phase 1: Cray XT5, 668 nodes, 5344 cores
- Phase 2: > 1 Pflop/s peak

Foundational mathematics and computer science research support our advances in high-performance computing

Scalable numerical algorithms for petascale scientific simulation

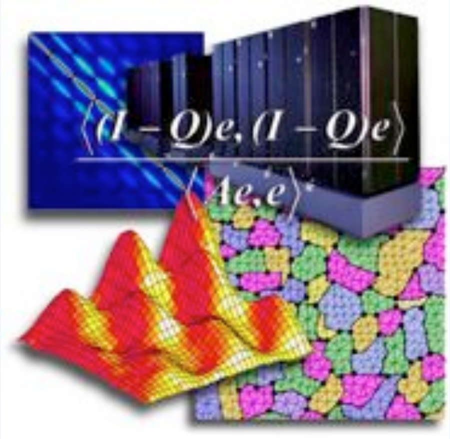


BlueGene/L

New tools and programming paradigms for portable performance on hundreds of thousands of processors

Algorithms and software R&D leverage platform investments

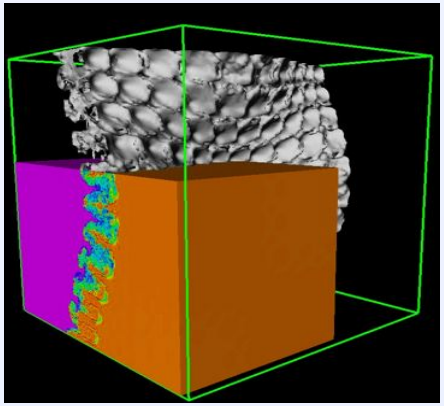

Interactive data exploration tools for petabyte-sized data sets

Some examples of simulations and the innovation with and through them

- **Aerodynamic drag reduction for heavy vehicles**
  - ***HPC simulations helped design and understand wind tunnel tests***
  - **Provided guidance to industry to improve the fuel economy of class 8 tractor-trailers through the use of aerodynamic drag reduction**
- **Core collapse supernova modeling**
  - ***Innovation in mathematics allowed never-before completed simulations***
- **Climate – hydrology coupling**
  - ***Increasing complexity of simulation and employing HPC helps us understand effects of climate change***

Class 8 tractor-trailers are responsible for 12-13% of total US petroleum consumption

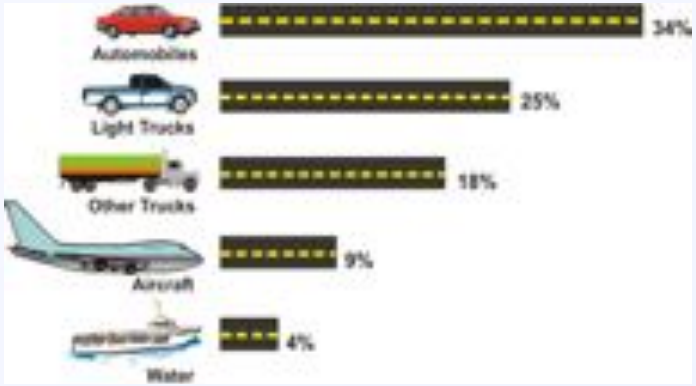

Usable energy losses come mainly from aerodynamic drag and rolling resistance

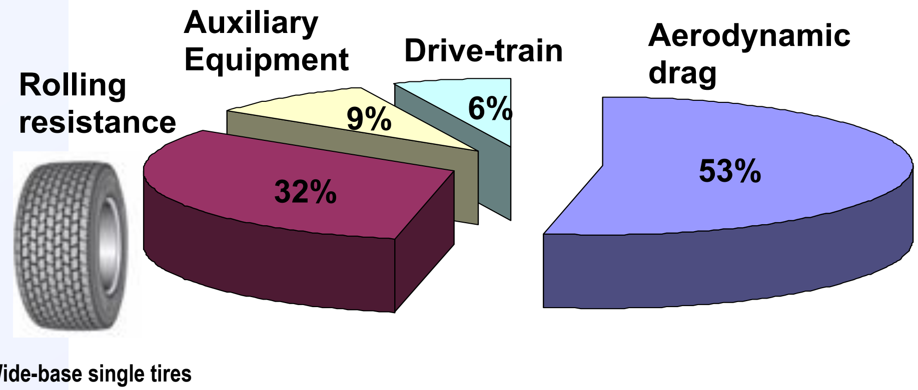

Aerodynamic drag breakdown on a typical truck

Losses in nearly all categories can be reduced through presently available technology

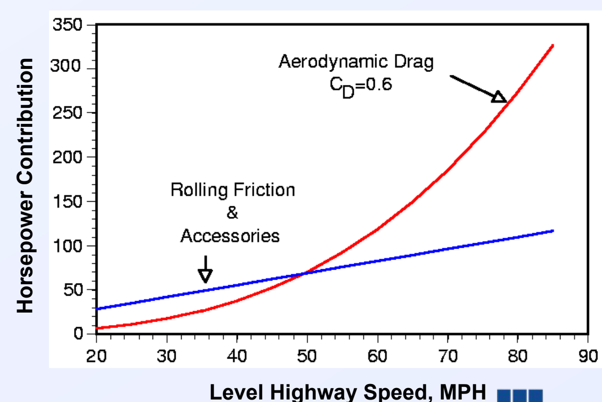

Mathematical models help us understand aerodynamic properties

Most drag is from pressure difference

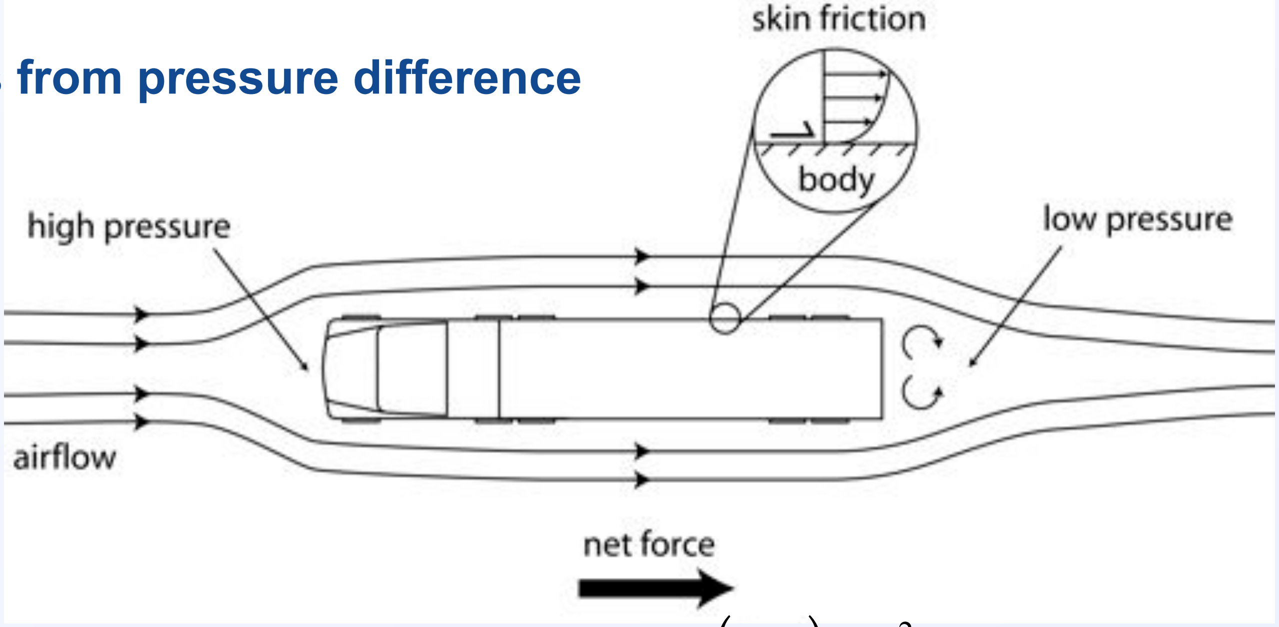

$$Drag = C_D \times S \times (1/2)\rho U^2$$

$$\frac{\Delta FuelConsumption}{FuelConsumption} = \eta \times \left(\frac{\Delta C_D}{C_D} + \frac{\Delta S}{S} + \frac{3\Delta U}{U} \right)$$

$\eta \approx$ **0.5-0.7** **shape** **cross-section** **speed**

Full-scale wind tunnel test at NASA NFAC facility: used to test drag reduction designs developed through simulation

- Different combinations of tractors and trailers were tested
  - Two tractors
  - Three trailers
- Performed 140 wind tunnel runs
- Twenty-three aerodynamic drag reduction devices/concepts were tested
- Performed acoustic measurements on selected add-on devices
- Couldn't test everything with wind tunnel:
  - Boundary layers in the tunnel are not as seen on the road

Design of full-scale wind tunnel test relied on high fidelity computational results

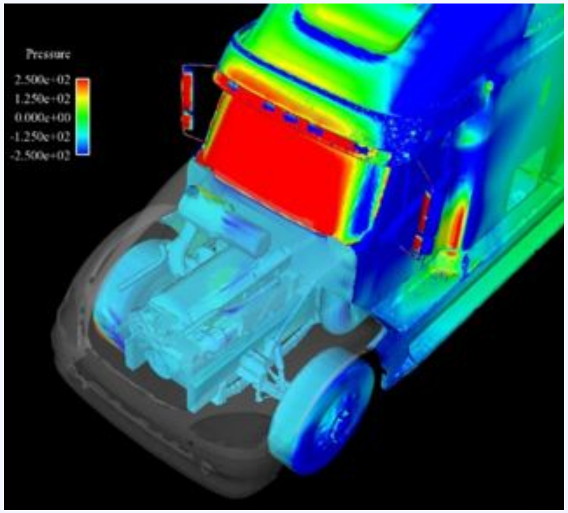

Freightliner Columbia tractor

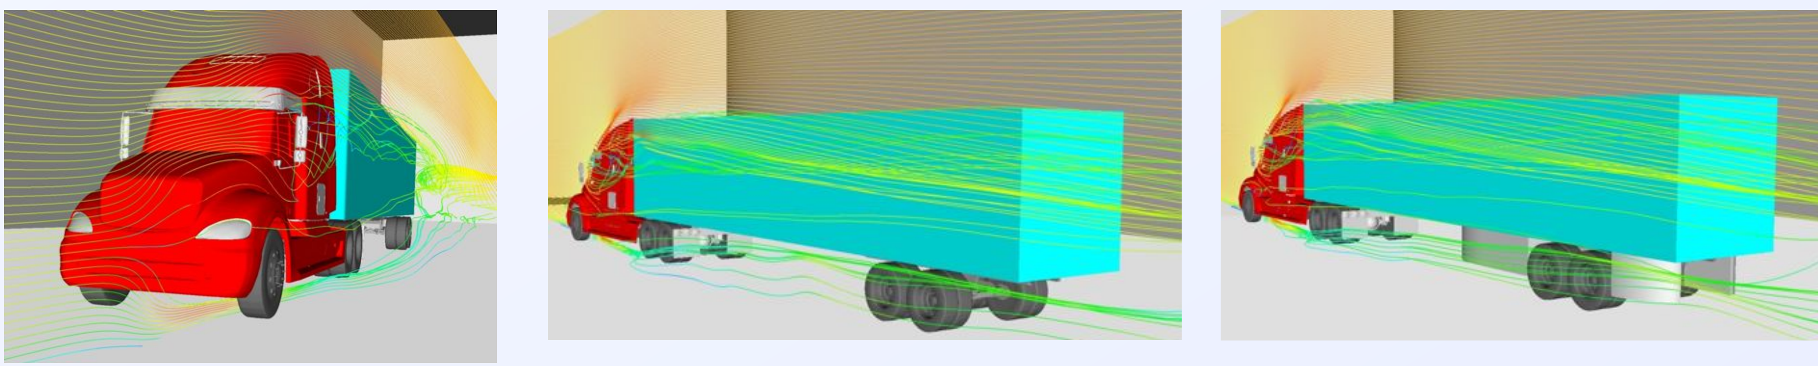

Velocity streamlines about a class 8 heavy vehicle in the wind tunnel

Performance of aerodynamic devices

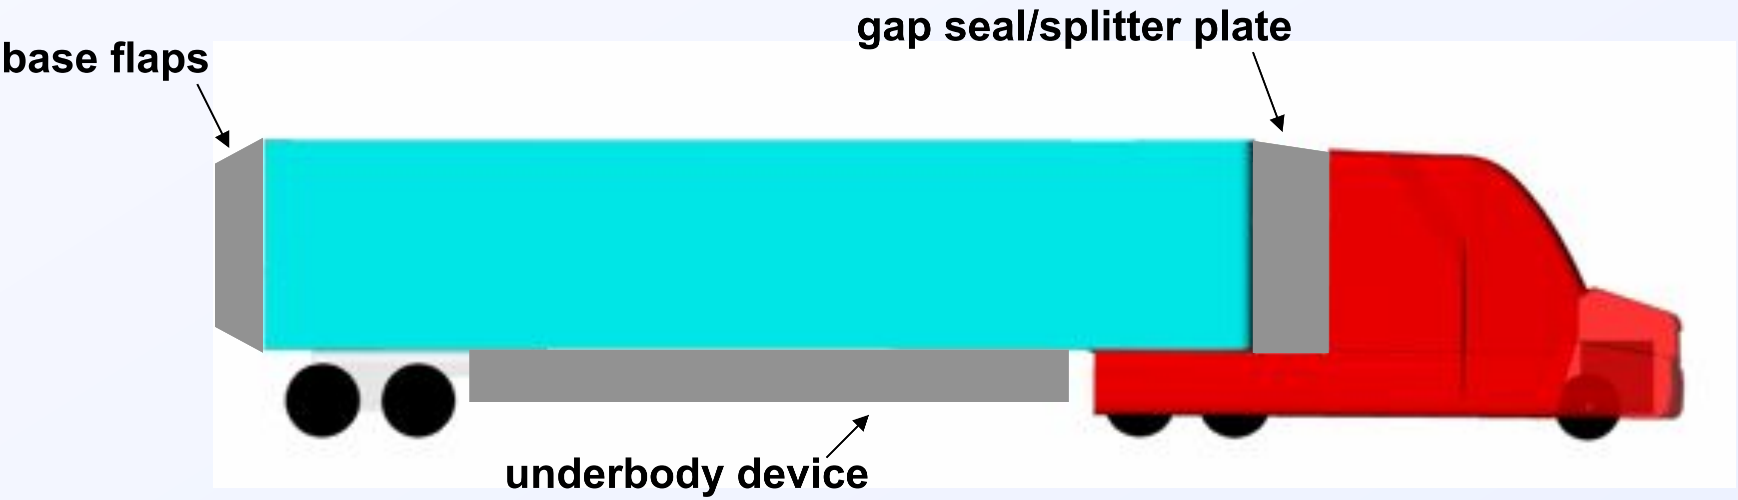

- Base flaps: 4-7% FEI (Fuel Economy Improvement)
- Underbody devices: 5-7% FEI
- Gap devices: 1-2% FEI
- Wide-base single tires: 4-5% FEI

Kambiz Salari and Jason Ortega (LLNL), Ronald Schoon and Andrea Brown (Navistar International Corporation), Ralph Hulseman and Victor Abarotin (Michelin Americas Research Company), Mark Betzina and Christopher Hartley (NASA)

Work sponsored by U.S. DoE Energy Efficiency and Renewable Energy Vehicle Technologies Program

Stellar collapse simulations must deliver high accuracy over long time periods

- Evolve by fusion burning, iron core collapse, blow up, and neutrino radiation cooling
- Stellar collapse, supernova explosion, and neutron star formation is one of the most energetic, *but poorly understood*, phenomena in astrophysics
- Models: radiation transport, hydrodynamics & self-gravity
- Neutrinos carry off ~99% of energy; 1% is observed (have to model this)

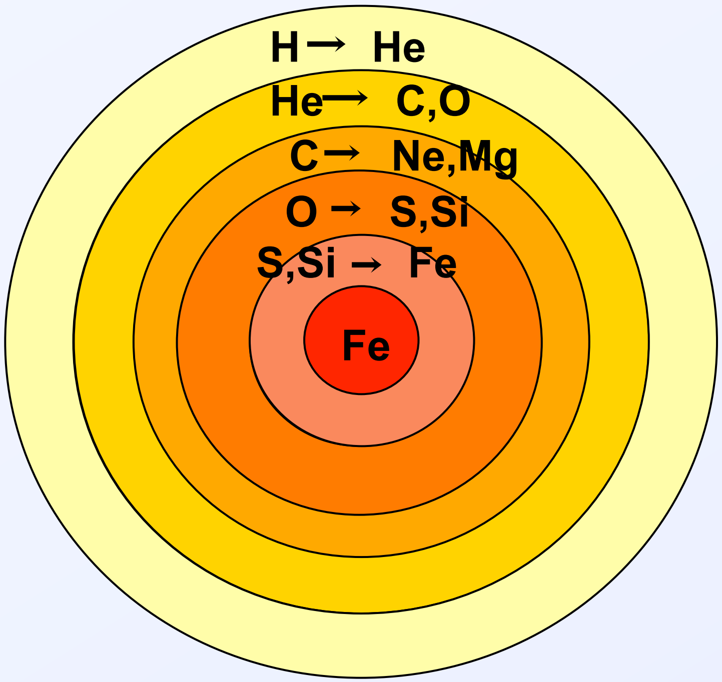

Onionskin-like structure

$\Delta t_{CFL,core}(10^{-7}\ s) < \text{desired step} < \Delta t_{CFL,shock}(\sim 10^{-4}\ s) < t_{cooling}(10s)$

We model the hydrodynamics with a Lagrangian formulation and pose the discrete system as a Differential-Algebraic Equation

Lagrangian Eqns (fluid flow) + Self-gravity + Conservation form

=>

Stellar collapse hydrodynamics model

$$\frac{V(U_n) - V(U_{n-1})}{\Delta t} = \theta F(U_n) + (1-\theta)F(U_{n-1}),$$

$$G(U_n) = 0, \quad \text{self} - \text{gravity}$$

Implicit approach works for Eulerian forms as well

We apply Newton-Krylov to the space-discretized form of this DAE system using the SUNDIALS KINSOL package

w/. D. Reynolds (SMU) & D. Swesty (SUNY SB)

We formulated a novel implicit approach to provide stability for longer time steps

Mathematical innovation was required:

- **Non-smooth discretization method -> Freeze parameters in the solver to prevent oscillations**
- **Solver convergence was slow -> Developed an effective local accelerator**
- **Unknowns and equations at differing scales can cause difficulties with stopping criteria -> Use scale factors in convergence criteria**
- **Certain variables have positivity constraints which can be violated in nonlinear updates or differencing -> Apply a transform to ensure quantities remain positive**

Implicit approach enabled radiation-hydrodynamic modeling of an entire proto-neutron star cooling

- **Initial central density: 5 x 10^{14} g/cm^3**
- **Initial radiation distribution contributes ½ of pressure support in the star center and diminishes radially**
- **Star contracts as neutrinos diffuse out**
- ***Implicit for neutrino diffusion as well as hydrodynamics***

Mass contours show contraction

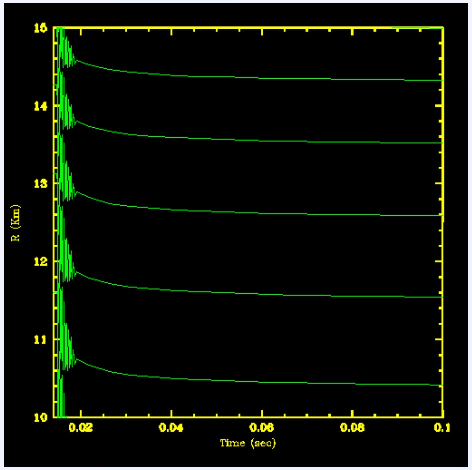

Neutrino cooling signal decays over 15s timescale

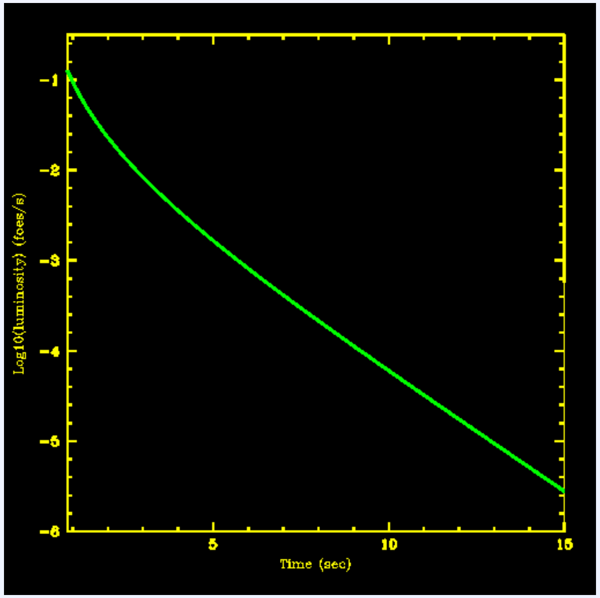

Explicit CFL restriction:

$\Delta t \sim 2.5 \times 10^{-8}$ s => 10^9 steps

Implicit used $\Delta t \sim 2.5 \times 10^{-5}$ s

0.1% of the number of explicit steps!

Reynolds, Swesty, and W., 2008

LLNL is working on the science for the next generation uncertainty-quantified climate system models

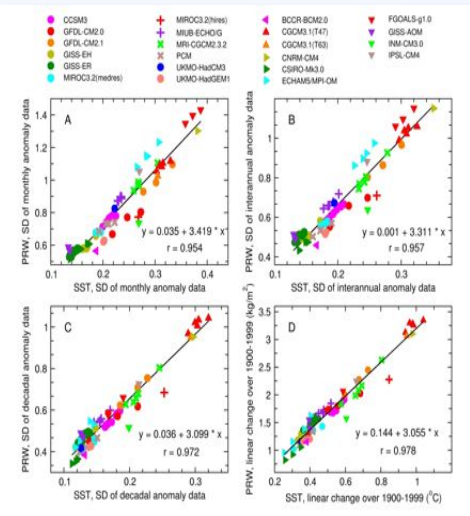

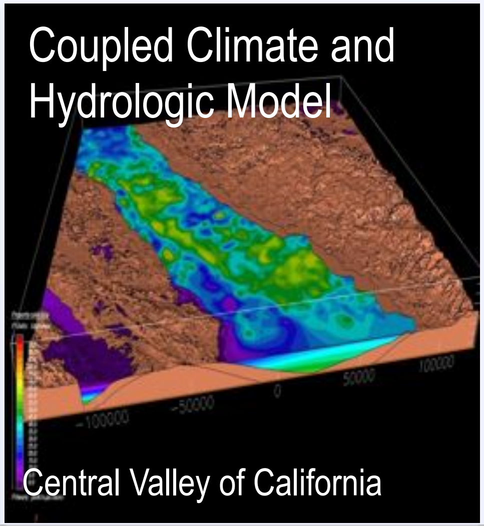

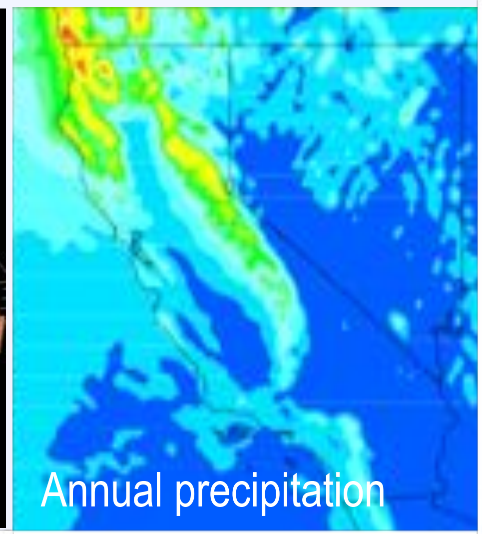

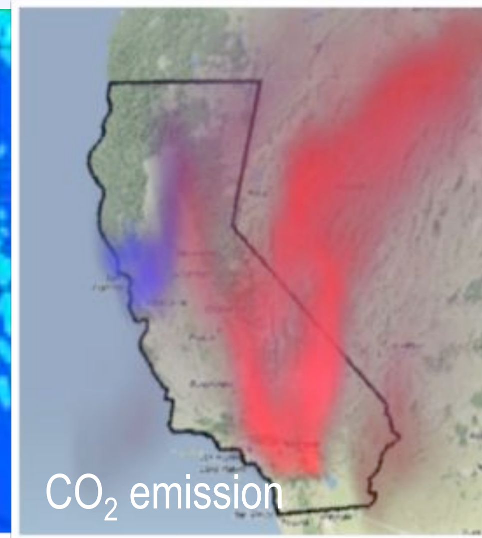

Uncertainty Quantification	Improved Physics Base	Local/Regional Resolution	Data at Scale (GHGIS)
Understand parameter sensitivity in multi-dimensional space Understand feedbacks in the carbon cycle	Improved cloud, aerosol, **coupled climate-hydrology** and other capabilities	Novel algorithms for ultra-high resolution global climate on peta-exa scale platforms	Radiocarbon and satellite measurements for full earth system models and global validation of carbon emissions

The hydrologic cycle is usually simulated with disconnected processes each run separately

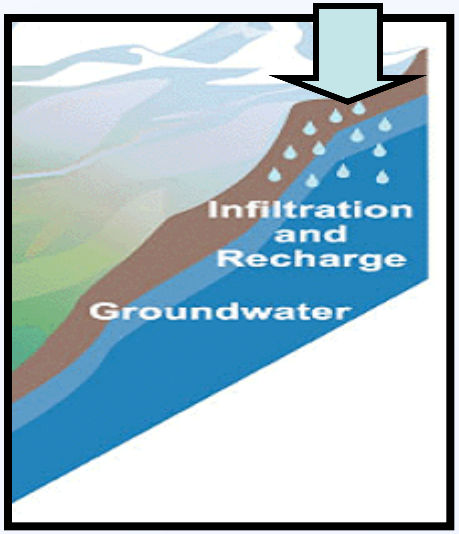

Groundwater / Vadose Model

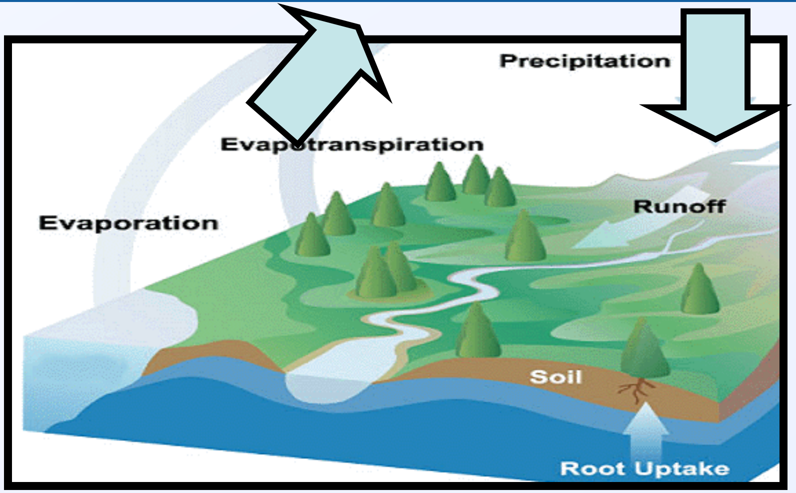

Land Surface Model

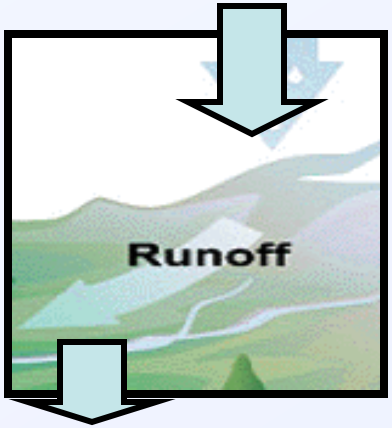

Surface Water Model

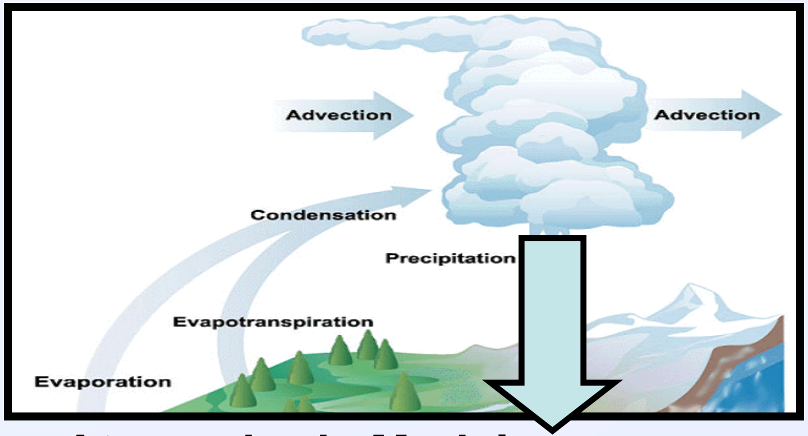

Atmospheric Model

We are developing a coupled simulation framework

- **Variably saturated subsurface flow through ParFlow package – developed for HPC frameworks from the start, incorporated innovative solvers for very fast solution**

- **Coupled overland flow and routing from surface boundary**
  - ***Ran on > 16,000 processors of Julich IBM supercomputer and demonstrated ability to conduct watershed size simulations with resolutions of 1-10 m laterally and 0.01 - 0.1 m vertically***

- **Recent work focused on coupling ParFlow and WRF**
  - **WRF: NCAR's Weather Research & Forecasting system**

- ***Targeting watershed-scale coupled subsurface-surface-climate simulations***

Accurate coupled simulations are critical to wind power prediction and use as well as water resources management

- **Intermittent nature of wind power => accurate wind forecasting is essential**
- **Low-level winds significantly depend on land surface energy and moisture fluxes**
- **Experiments show these fluxes are sensitive to soil moisture – must couple winds to surface**

- **In the USA, we depend on climate for all renewable water resources**
- **Changing climate potentially affects**
  - **Water demands & water quality**
  - **Amount, rate, and location of water inputs to the land surface and outputs to the atmosphere**
- **Nature of water use feeds back to atm:**
  - **Ex: One acre of corn will transpire 4,000 gal of water per day**

Climate forcing scenarios can be used to understand watershed response under climate change

1. Control simulation based on WY1999 “CNTL”

2. “Hot” case H: T=T+2.5K

3. “Hot+Wet” case HW: T=T+2.5K, P=P*1.2

4. “Hot+Dry” case HD: T=T+2.5K, P=P*0.8

Maxwell and Kollet; Nature Geo (2008)

Initial coupled runs show definite areas of subsurface flow influence

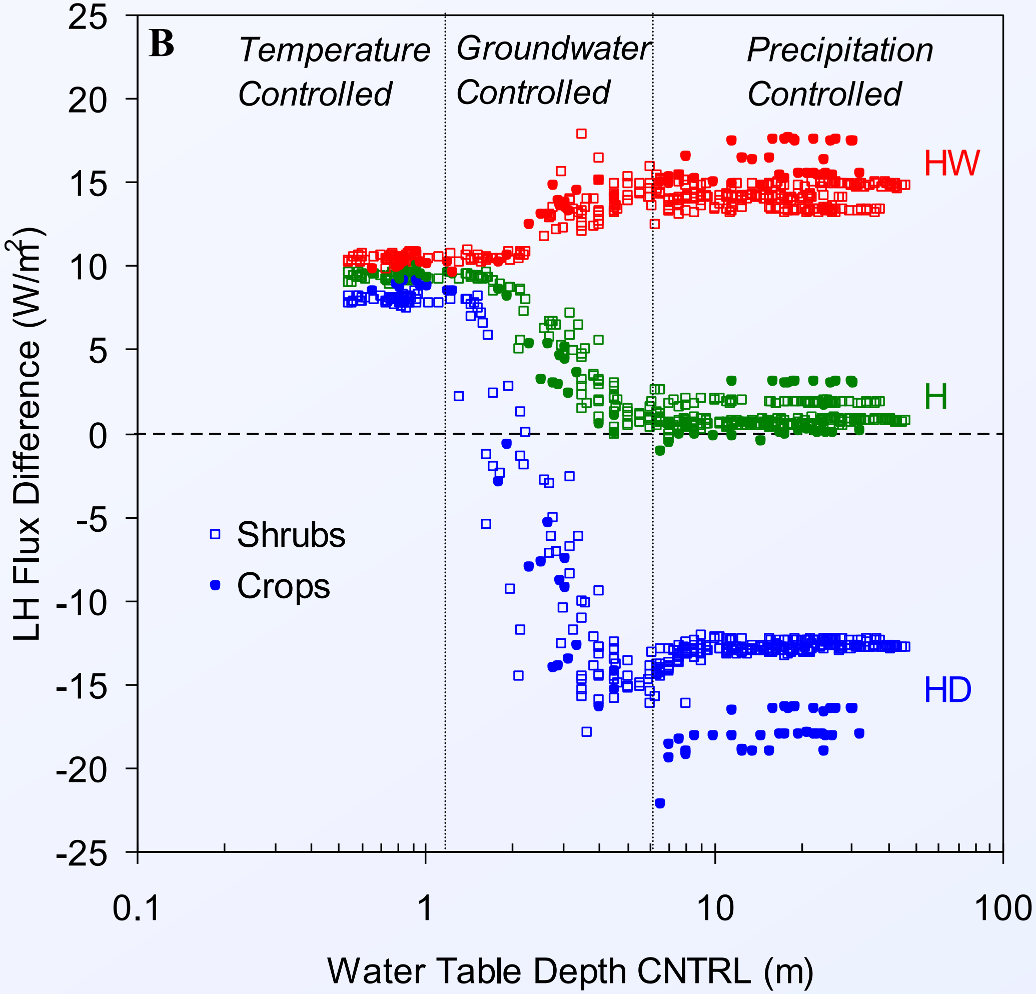

For latent heat:

- **Deep water table: precipitation is primary control**
- **Very shallow water table: temperature is primary control**
- **For 1-10m water table: groundwater is significant**

Maxwell and Kollet *Nature Geo* (2008)

Low-level winds are slower with coupled model due to cooler predicted surface temperatures

- 36 hr simulation
- Morning start
- 6 runs each
  - 3 microphysics
  - 2 temp. initializations
- Wind speed daily cycle
  - Max early evening
  - Decrease at night
- Daytime – convective heating dominates flow

Domain-averaged 40 m wind velocities

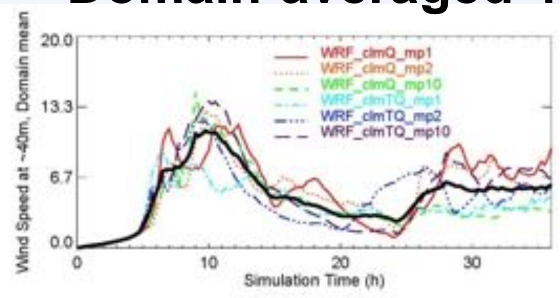

Weather simulation only

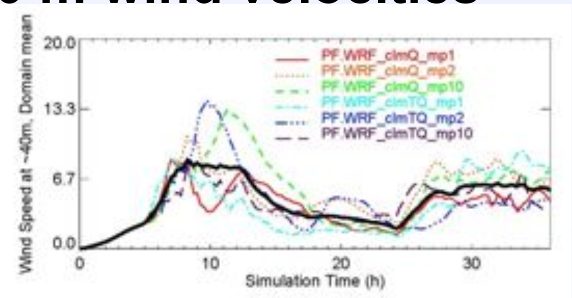

Weather & subsurface

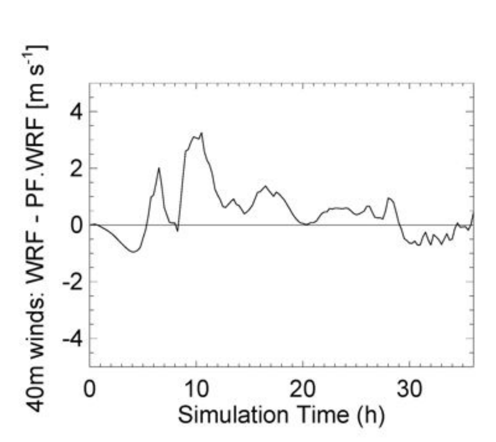

Difference in 40 m winds averaged over microphysics and domain

Weather simulation-alone soil temps are ~5 deg higher giving stronger convective plumes and stronger horizontal winds

At the petascale performance increases through core count not clock rate: *we have reached a critical turning point in HPC design*

- New Constraints
  - 15 years of *exponential* clock rate growth has ended
- Moore's Law reinterpreted:
  - How do we use all of those transistors to keep performance increasing at historical rates?
  - Industry Response: #cores per chip doubles every 18 months *instead* of clock frequency!

Need new programming models

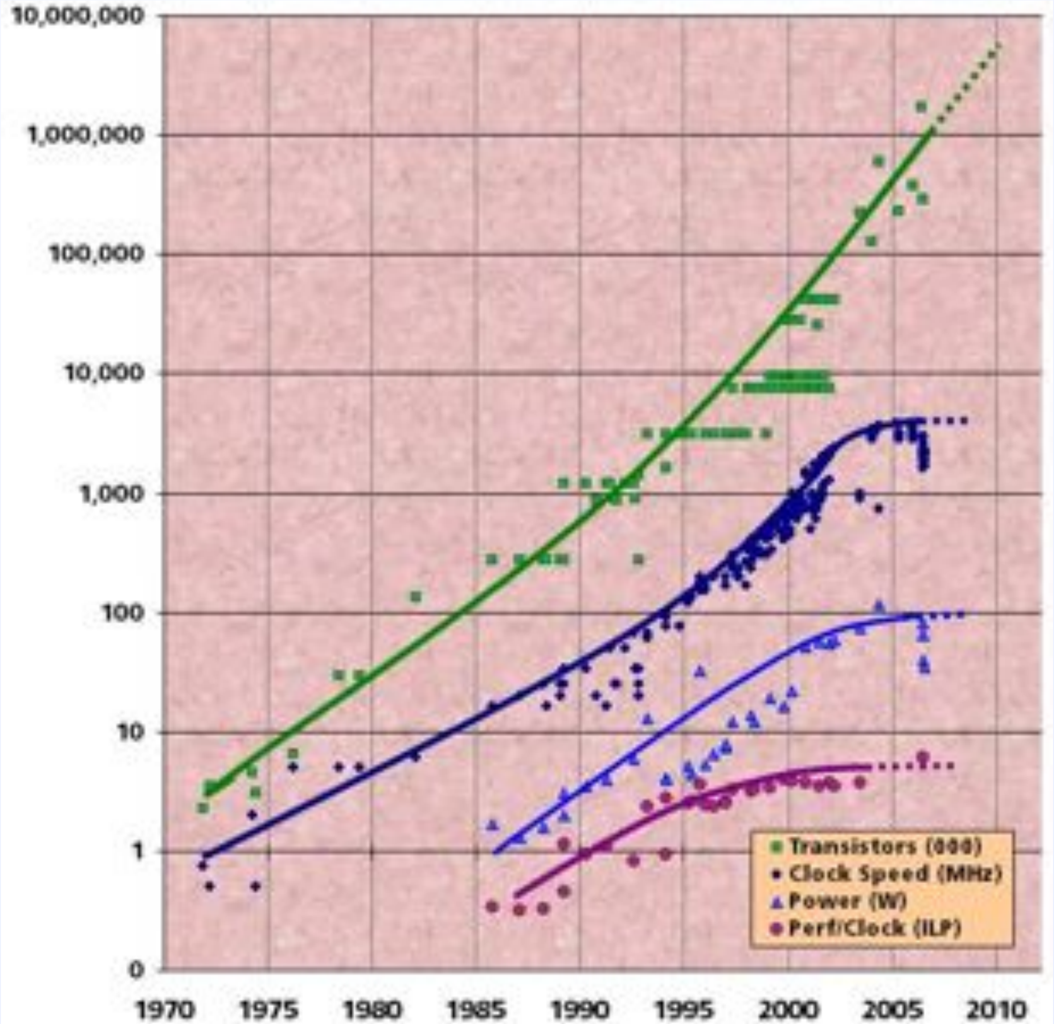

Acknowlegement: Kunle Olukotun, Lance Hammond, Herb Sutter, and Burton Smith

Challenge to the applied math and computer science communities: Exascale ≠ Petascale X 1000

- Traditionally, PDE-based applications have expected 10x increase in resolution with each 1000x increase in compute capability, *but not this time:*
  - We won't have 1000x the memory available
  - The processors won't be 10x faster
  - Proportionally, we won't be able to move as much data on or off each processor
  - Introduction of massive parallelism at the node level is a significant new challenge (MPI is only part of the solution)

- However, exascale computing is an opportunity for…
  - ***More Fidelity:*** Incorporate more physics instead of increased resolution
  - ***Greater Understanding:*** Develop uncertainty quantification (UQ) to establish confidence levels in computed results and deliver predictive science

What are critical exascale technology investments?

- **System power** is a first class constraint on exascale system performance and effectiveness.
- **Memory** is an important component of meeting exascale power and applications goals.
- **Programming model.** Early investment in several efforts to decide in 2013 on exascale programming model, allowing exemplar applications effective access to 2015 system for both mission and science.
- **Investment in exascale processor design** to achieve an exascale-like system in 2015.
- **Operating System strategy for exascale** is critical for node performance at scale and for efficient support of new programming models and run time systems.
- **Reliability and resiliency** are critical at this scale and require applications neutral movement of the file system (for check pointing, in particular) closer to the running apps.
- ***HPC co-design strategy and implementation*** *requires a set of a hierarchical performance models and simulators as well as commitment from apps, software and architecture communities.*

Acknowledgement: Horst Simon, LBNL

Now, Janine Bennett will speak on her experiences at SNL

Why do we need exascale computing?

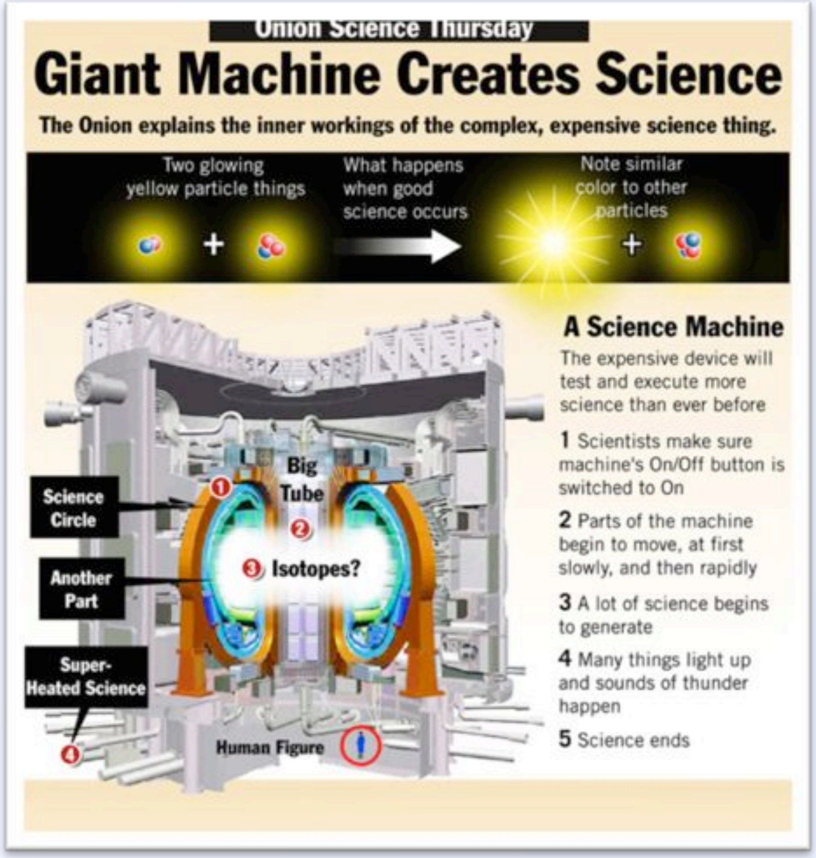

http://www.theonion.com/articles/scientists-ask-congress-to-fund-50-billion-science,2291/

Across science and engineering, challenging problems drive the need for continued advances in high performance computing capability

Area		Applications	TeraFLOPS	PetaFLOPS	ExaFLOPS
	Health, Life Sciences	Viral evolution, advanced threat characterization, disease spread prediction, host–pathogen response, drug design, protein function	Molecular docking Macro molecular motions Nanosecond resolution	Molecular docking with few free energy calculations Macro molecular motions with 2,000-atom chemistry calculations Microsecond resolution	Free energy calculations Macro molecular motions with ~20,000-atom chemistry calculations Millisecond-scale resolution
	Environment: Agriculture, Natural Resources, Earthquake Analysis	Earth motion dynamics; structural safety of buildings and levees; ocean, atmosphere, and land climate models	Greenhouse gas effects 100–1,000 multicentury coupled physical climate simulations 50-km resolution	Energy technology consequences 100–1,000 century-long earth system model simulations 50-km resolution	High-resolution earth system 10 configurations of model, each run 100–1,000 times, century-long simulations 1-km resolution
	Energy Production and Conservation	Turbulent combustion, prediction of fuel efficiency, more efficient solar cells	Coupled hydrological, geomechanical, geochemical reservoir simulation	Full reverse-time, elastic, anisotropic seismic simulation 3D volume equivalents	Full reverse-time inversion, with coupled wave excitation, very large fold, in high resolution
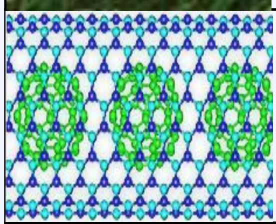	Nanoscience	Development of new materials: faster, lower power semiconductors; drug development; fuel cells; better airport screening	Band gaps in alloys, molecular behavior in simple fluids 10s of atoms with fully quantum mechanical, ab initio calculations	Molecular motion in confined geometries, kinetics of phase change 10s of thousands of atoms with ab initio codes	Interacting biological molecules, micron-scale machines 10s of millions of atoms with ab initio codes (would require algorithm development)
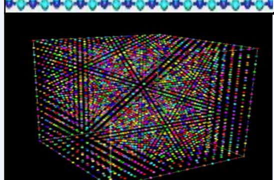	Basic Science, Fundamental Physics	Physics of early universe, nuclei and nuclear reactions, supernovae	QCD transition $N\pi$ reactions ^{6}Li 2D S-Nova	QCD plasma N-N reactions ^{12}C, ^{16}O 3D S-Nova	Electro-weak light nuclei ^{40}Ca Full 3D S-Nova

How will we achieve exascale computing?

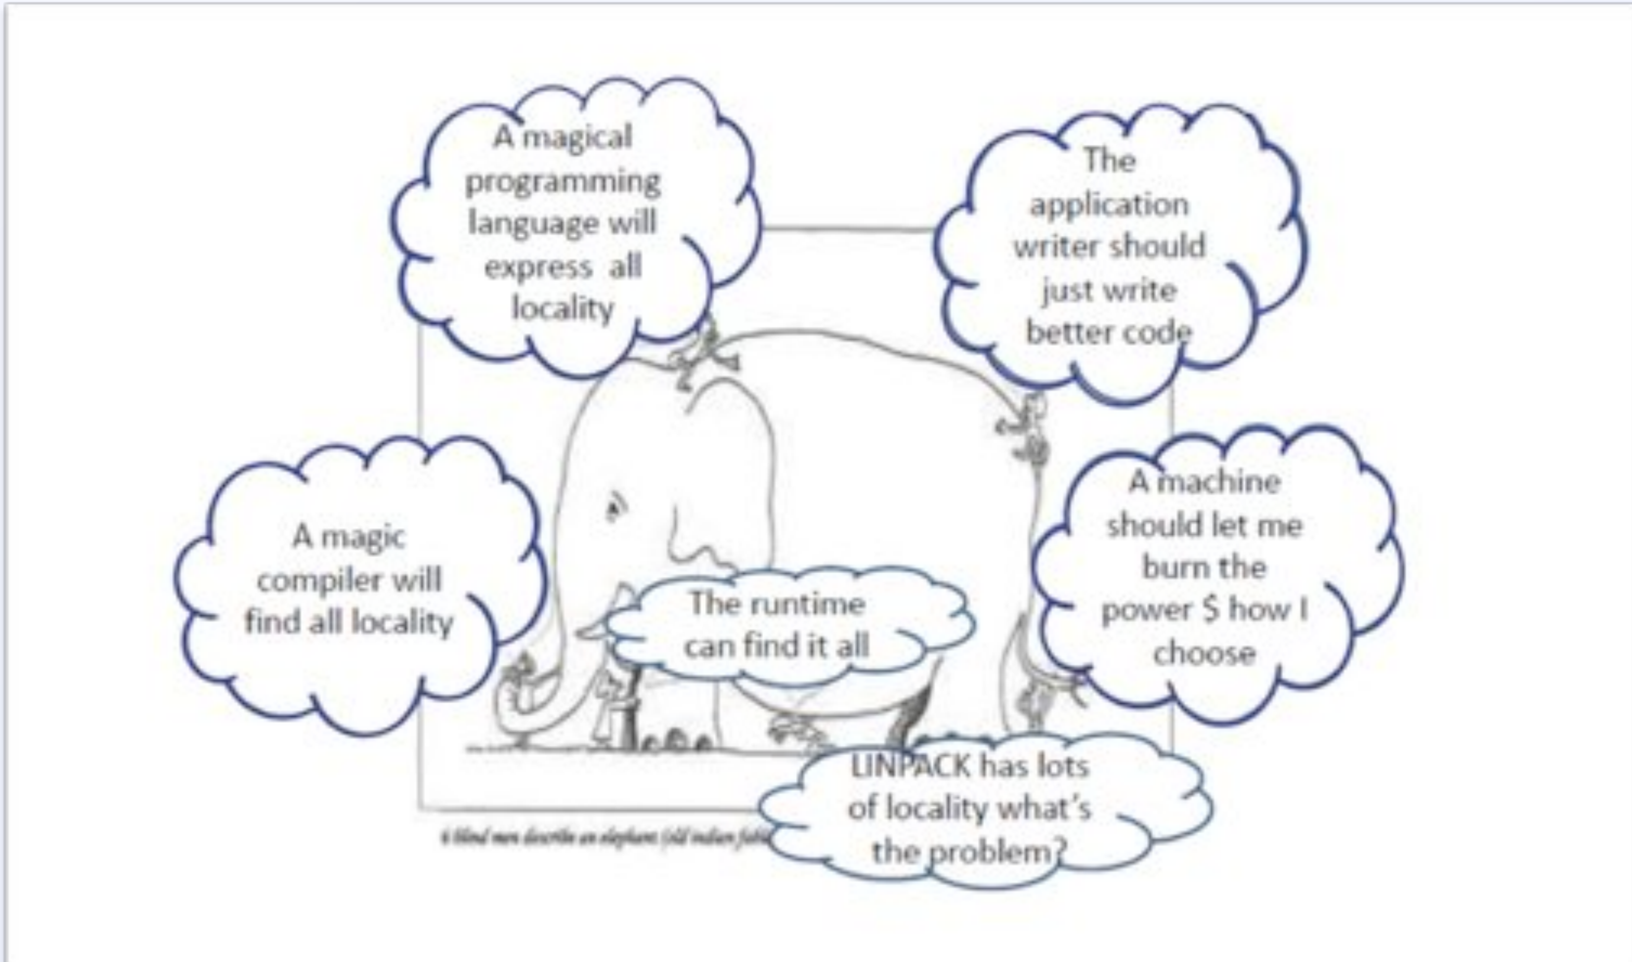

Acknowledgement: Allen Snavely, San Diego Super Computer Center

DOE SciDAC Exascale Co-design centers allow for unprecedented levels of collaboration between hardware and software researchers

- Create exascale-capable applications
- Inform hardware design
- Understand systems software requirements
- Devise new algorithms

- Center for Exascale Simulation of Advanced Reactors (CESAR)
- Exascale Co-Design Center for Materials in Extreme Environments (ExMatEx)
- Center for Exascale Simulation of Combustion in Turbulence (ExaCT)

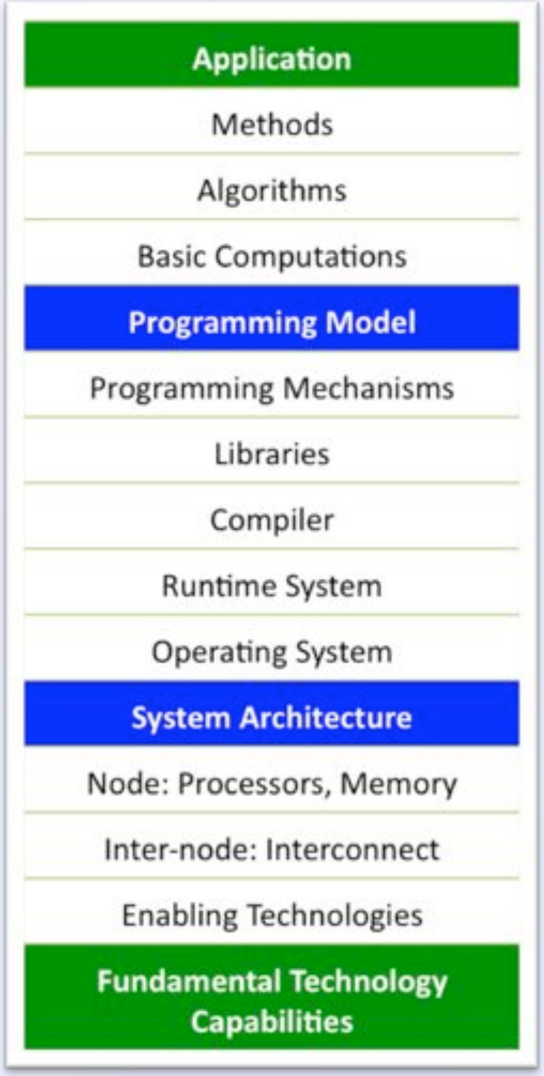

ExaCT will determine the optimal design of combustion applications and exascale hardware

Director: Dr. Jacqueline Chen, SNL

Deputy Director: Dr. John Bell, LBNL

Development goals:

- Combustion modeling & analysis capabilities
- Computer science tools
- Quantify hardware constraints at exascale

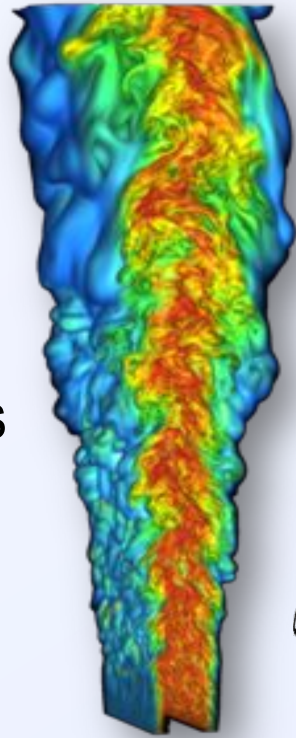

High-fidelity combustion simulations at the exascale are essential to gaining fundamental insights and providing validation for predictive models of 21st century fuels used in transportation, power generation, and industrial processes.

Combustion research plays an important role in energy security

- Combustion accounts for a majority of energy used in the United States
- Computer simulations provide tools for design of efficient clean burning devices
- Sound scientific understanding is necessary to develop predictive, validated multi-scale models

Combustion simulations generate large, complex data

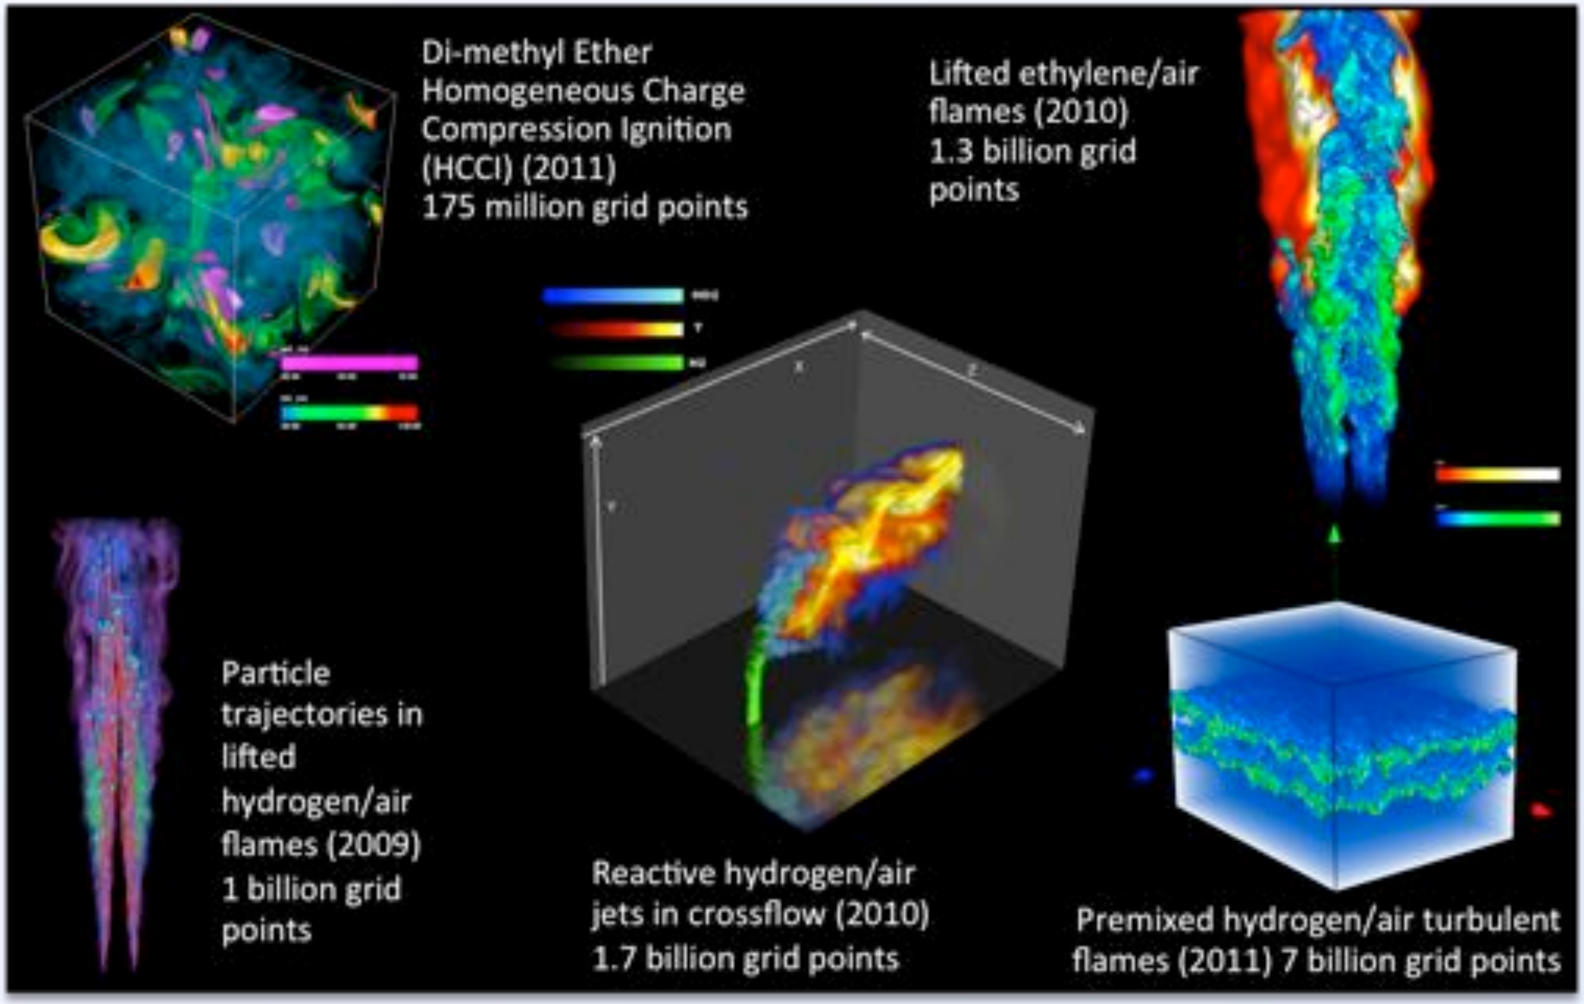

Recent direct numerical simulation (DNS) configurations performed using S3D, a DNS code written by Dr. Jacqueline Chen & her research group at the Combustion Research Facility, SNL.

We want to identify features of interest and analyze them in a variety of ways

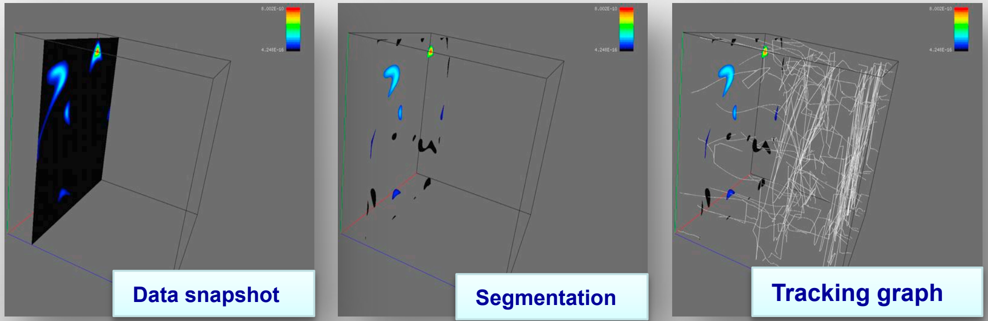

Tracking features in time

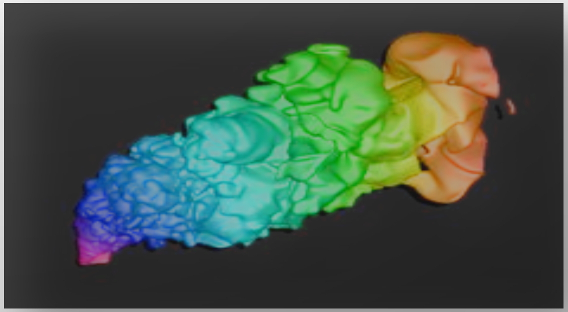

Jet-based coordinate systems allows for aggregation of statistics conditioned on bulk flame position

We want to identify features and characterize their shapes and analyze the behavior of other variables within these features

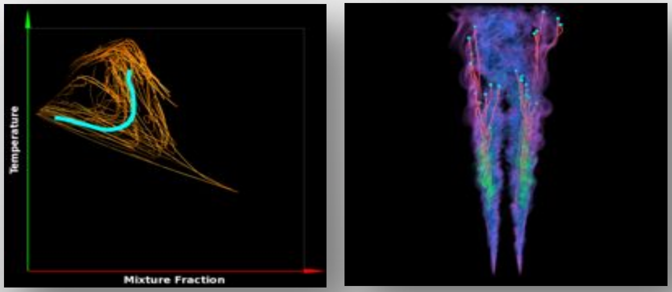

Studying particle trajectory information gives us insight into Lagrangian flow patterns

At exascale data analysis must be done *in situ* because I/O operations will be too expensive

- Historically analysis has been performed as a post-process
- Due to I/O restrictions analysis will need to be done in lock-step with the simulation

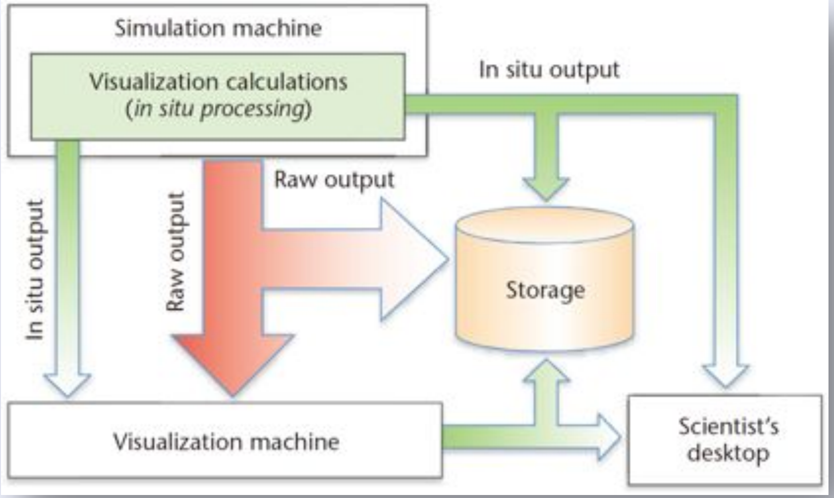

- *In situ* analysis challenges:
  - Sharing data structures between simulation and analysis codes
  - Developing scalable methods constrained to the simulation data partitioning
  - Developing methods that run at a very small fraction of simulation time

As we develop these new algorithms, we need a method to test their behavior on potential architectures

Structural Simulation Toolkit (SST) – create a multi-scale computer architecture for design and procurement of large-scale parallel machines as well as in the design of algorithms for these machines.

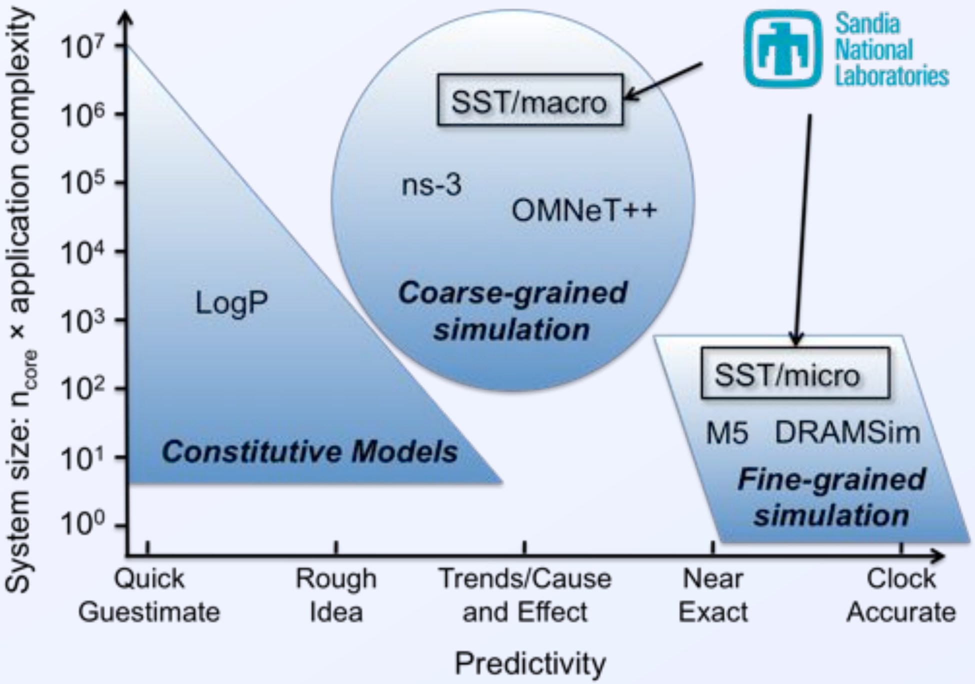

Acknowledgement: Curtis Janssen, SNL

The Structural Simulation Toolkit developed at Sandia permits the study of future HPC systems

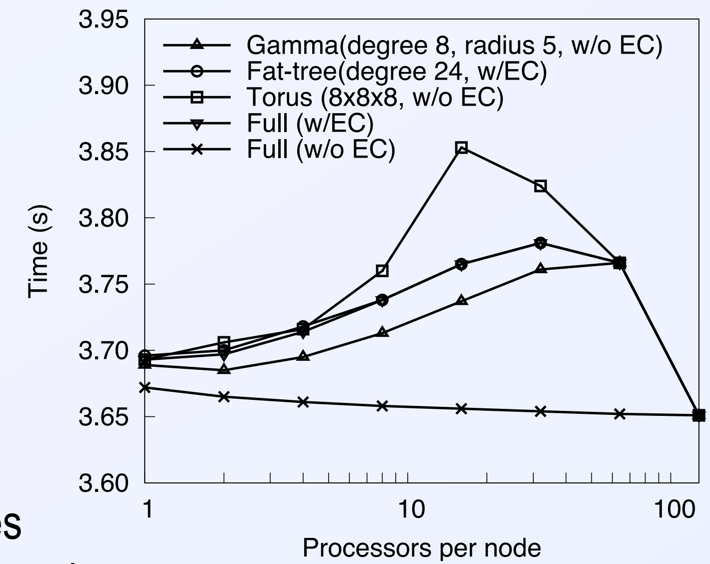

- Correctly identify causal relationships
  - Network topology
  - Node configuration
  - Noise/imbalance
  - Bandwidth
  - Latency
  - Resource contention
- Play "what if" games
  - Implementation effects for communication routines
  - Infinite performance in some components to stress others.
- Test changes to application, middleware, or resource management
  - Reordering code blocks, scheduling effects, etc.
- Test novel programming models
  - Fault-tolerant or fault-oblivious execution models
  - Alternatives to MPI, parallel runtime designs
  - Mixed programming models

Acknowledgement: Curtis Janssen, SNL

Networks are everywhere

U.S. Power Grid
[GENI]

Yeast-Protein
interactions
[Boardalier institute]

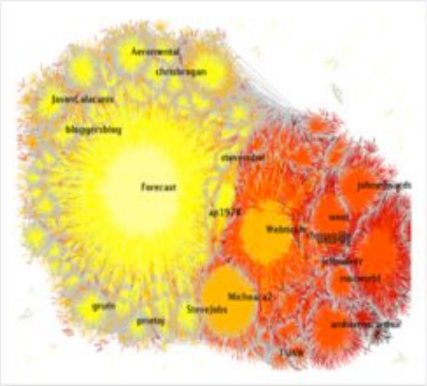

Twitter social network
[Akshay Java, 20077]

- Physical networks
  - Defined by physical connections
  - Clearly defined for each system
  - Sandia's interest: Power, water, communication networks
- Functional networks
  - Defined by well-defined functional dependencies between entities
  - Complete information is available; Needs abstraction
  - Sandia's interest: supply chains, chemical reaction networks
- Interaction networks
  - Defined by interactions between entities
  - Information is incomplete and noisy; Needs abstraction
  - Sandia's interest: cybersecurity, counter terrorism, counter intelligence, epidemics

Acknowledgement: Ali Pinar, SNL

Graph analysis capabilities are needed to effectively support national security needs

- We are building basic capabilities to:
  - Pinpoint network vulnerabilities
  - Identify substructures similar to a specified pattern
  - Predict how networks will evolve in time

- These enabling technologies will be used to:
  - Quantifiably secure design & operations for critical infrastructure
  - Discover potential threats to national security by analyzing massive volumes of data
  - Characterize events for rapid and effective response

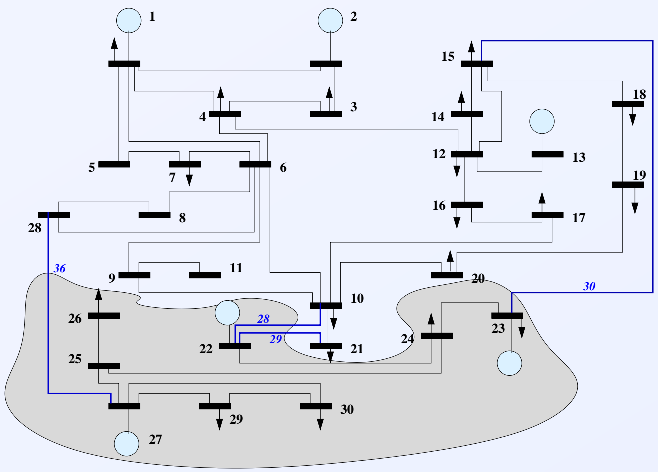

Vulnerability in the IEEE30-bus system

Acknowledgement: Ali Pinar, SNL

Good network models are essential for a variety of reasons

- Benchmarking algorithms and architectures
- Anonymization
  - Not all data can be shared
- Statistically significant data mining
- Predictive simulations
- Recovery of a network from limited observations

- Developing good models is extremely challenging task that requires combinatorics, linear algebra, optimization, and statistics

Acknowledgement: Ali Pinar, SNL

Research at Sandia is improving the Graph 500 benchmark

- Graph500 benchmark is designed to represent data-intensive, analytic workloads that are significantly different than those of traditional 3D physics codes.
- Benchmark used R-MAT to generate graphs
- Two major results of our analysis:
  - Identified and fixed a significant problem with the degree distribution
  - Showed that up to 75% of vertices were isolated in last year's benchmark

Graph500 will be updated based on our recommendations

Work funded by ASCR [Seshadri, Pinar, Kolda in preparation]

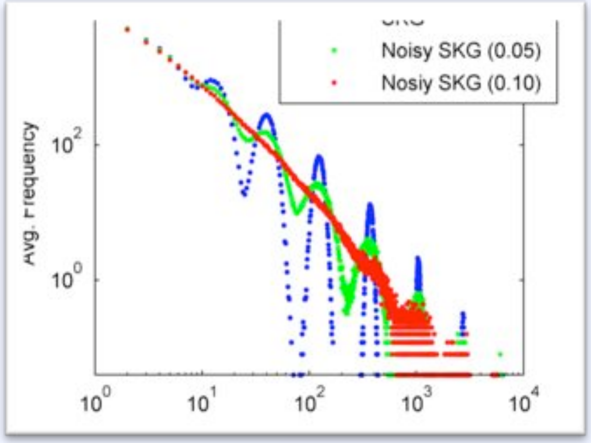

Degree distribution is smoothed by randomization

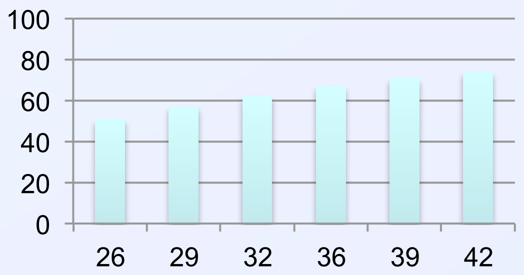

% of isolated nodes in Graph500 graphs

Acknowledgement: Ali Pinar, SNL

Come work with us at LLNL and Sandia!

We're looking for B.S., M.S., and Ph.D. candidates for summer, regular full-time, entry level, and post-doctoral positions:

- Scalable numerical algorithms for petascale scientific simulation
- New tools and programming paradigms for portable performance on thousands of processors
- Interactive data exploration tools for petabyte-sized data sets
- Advancing physics understanding by increasing simulation fidelity
- New paradigms for managing software complexity
- Computer security and information extraction for homeland security
- Database design for identifying new pathogens

Apply at https://careers.llnl.gov & www.jobs.sandia.gov

Internship opportunities

- Most are for the summer but start/end dates are negotiable
- Decisions typically made Feb – Apr timeframe
- You might generate interest in an internship position by applying to a job posting
- Opportunities to participate in workshops, seminars, brown-bag lunches, tours of Sandia/LLNL, social activities, and symposiums
- Interns contribute to the labs' missions by working on real-world, challenging projects
- Research findings may be showcased through summer symposium posters and presentations, presented at national conferences, and/or published in scientific journal articles

Full time employee benefits

Competitive salaries

On-site wellness programs & fitness facilities

401K with matching 6% + 3.5-6% automatic

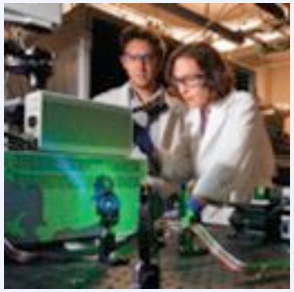

Fulfilling work environment

Flexible work schedules

Health beneftis

On-site wellness programs & fitness facilities

Career development

Education benefits

Community involvement

Do you have what it takes?

- Looking to build our workforce skills in these areas:
- Design and develop Oracle database applications using Oracle 11g, PL/SQL, and Oracle forms
- C and C++ programming skills (templating, inheritance, and polymorphism)
- Object oriented design, design patterns and scientific computing codes
- Algorithm development for parallel programming, using MPI and OpenMP
- Web application development using J2EE, JSP, JSF, Servlets, Tomcat, Ant/Maven, CVS/SVN, JUnite, JPA, Hiberate, Spring, Oracle, Eclipse, Hadoop, Lucene/Solr

Apply at https://careers.llnl.gov & www.jobs.sandia.gov

Questions?

Contact:

Carol S. Woodward
Lawrence Livermore National Laboratory
woodward6@llnl.gov

Janine C. Bennett
Sandia National Laboratories
jcbenne@sandia.gov

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

This work was performed under the auspices of the US Department of Energy (DOE) by the Lawrence Livermore National Laboratory under Contract Nos. DE-AC52-07NA27344.