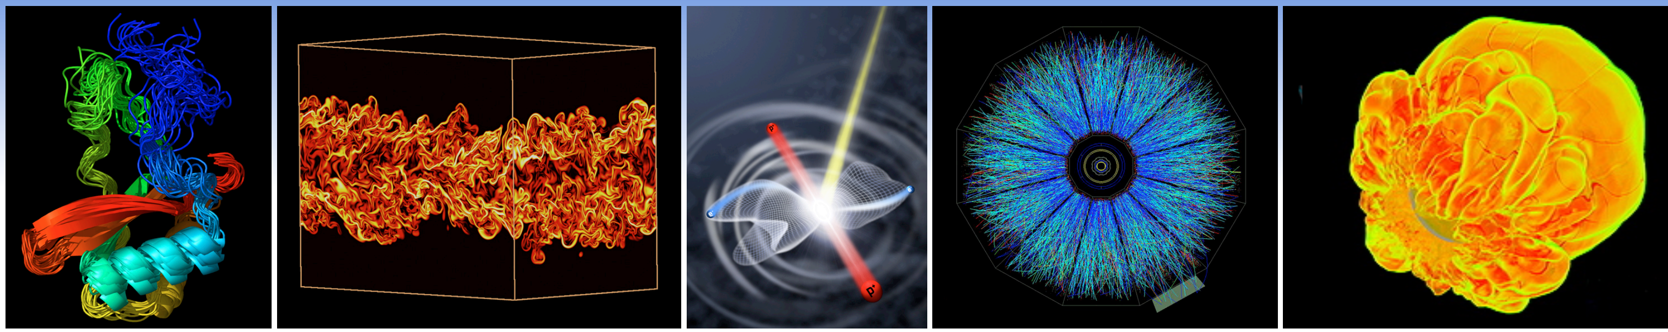

Workshop Goals & Process

Large Scale Computing and Storage Requirements for Biological and Environmental Research
Joint BER / ASCR / NERSC Workshop

Harvey Wasserman

NERSC Science Driven System Architecture Group

Lawrence Berkeley National Laboratory

May 7-8, 2009

U.S. DEPARTMENT OF ENERGY | Office of Science

NERSC

BERKELEY LAB

Logistics: Schedule

- Agenda on workshop web page
  - http://www.nersc.gov/projects/science_requirements/BER/agenda.php
- Mid-morning / afternoon break, lunch
- Self-organization for dinner
- 3 science areas, one workshop
  - Science-focused but cross-science discussion
  - Explore areas of common need (within BER)
- Breakout sessions Friday AM in one room

Logistics: Case Studies

- Two co-leads (for each science area)
  - help roll up discussions into major case studies
- Case Studies:
  - Narrative describing science & NERSC reqmts
  - Initial set based on discussions with co-leads
    - Minimum set to capture BER mission and unique NERSC requirements
    - Actual number may vary
  - Co-leads suggested discussion leader
  - Encourage participation by all; roundtable

Logistics: Templates

- Web templates: web “Reference Material”
  - Based on NERSC info
  - Summary of projects as we know them
  - Good point of departure
    - A framework for discussion
    - But not necessarily the entire story

Logistics: Final Report Content

- Format similar to ESnet
  - But NERSC requirement space much broader than Esnet
  - See "Reference Material" on web site
  - Contents
    - Executive summary,
    - ~2-page case study reports,
    - NERSC synthesis of all results

Logistics: Final Report Schedule

- Revised case studies due to NERSC .. May 21
- NERSC draft report …………………….. June 21
- Participants review period ……………….. July 7
- NERSC Near final .... July 21
- BER AD approval .... August 4
- NERSC Revisions .... August 15
- Final Report posted on Workshop Webpage
  .... August 16

Why is NERSC Collecting Computational Requirements?

- Help ASCR and NERSC make informed decisions for technology and services.

- Input is used to guide procurements, staffing, and to improve the effectiveness of NERSC services.

  – Includes hardware, software, support, data, storage, analysis, work flow

- Result: NERSC can better provide what you need for your work.

Examples of Information Sought

- Type of simulation, #, reason for #, algorithms, solver
- Parallelism: method, weak or strong scaling, implementation, concurrency, limits
- Key physical parameters and their limits:
  – spatial resolution, # of atoms/energy levels, integration range, …
- Representative code
- Key science result metrics and goals

Examples of Information Sought

- Typical science process (workflow)
- Data: amount stored / transferred for input, results, and fault mitigation
- How all of this is
  - Driven by the science
  - Likely to change and why

Climate Science

- Lawrence Buja (NCAR), David Randall (Colo. State)
- Role of Eddies in Meridional Overturning Circulation
  - Christopher Wolfe (Scripps/UCSD)
- Coupled High-Res Modeling of the Earth System
  - Christopher Kerr (GFDL)
- km-scale cloud resolving model
  - Dave Randall
- CCSM moderate/high-res studies
  - L. Buja
- Vegetation-air exchanges / regional climate models
  - Ned Patton (NCAR)

Environmental Science

- Timothy Scheibe, Bruce Palmer (PNNL)
- Subsurface science / biogeochemistry
- Climate science
  - Ned Patton (NCAR)

Biological Science

- Teresa Head-Gordon, Victor Markowitz (LBNL)
- Biological pathways and networks:
  – Costas Maranas (Penn State):
- Molecular dynamics:
  – Teresa
- Bioinformatics, database, data management, JGI, JBEI:
  – Victor Markowitz, (LBNL)
- Proteomics, clustering, Metabolomics
  – Lee Ann McCue (PNNL)

Final Thoughts

- LBNL will try to record – could use help
- Requirements characterization process is not complicated.
- Mutually beneficial.

Scaling Science

Inspired by **P. Kent**, "*Computational Challenges in Nanoscience: an* ab initio *Perspective*", Peta08 workshop, Hawaii (2008) and **Jonathan Carter** (NERSC).

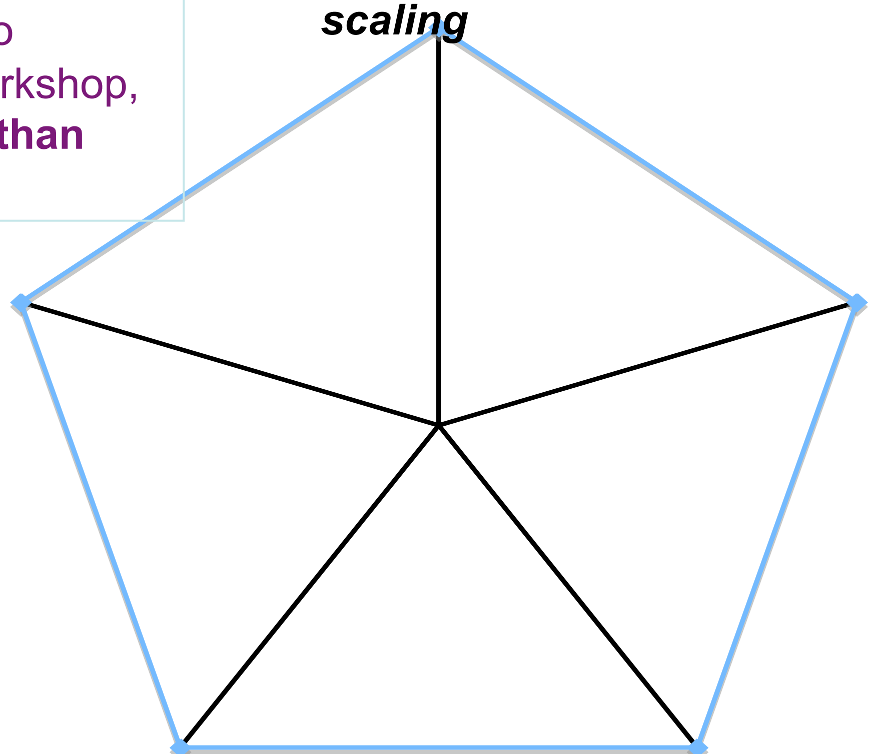

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

BACKUP SLIDES

Workload Analysis

- Ongoing activity within NERSC SDSA*
- Effort to drill deeper than this workshop
  - Study representative codes in detail
- See how the code stresses the machine
  - Help evaluate architectural trade-offs

***Science Driven System Architecture Team, http://www.nersc.gov/projects/SDSA/**

Workload-Driven Characteristics

- Memory requirements as *f*(algorithm, inputs)
- Memory-to-floating-point operation ratio
- Memory access pattern
- Interprocessor communication pattern, size, frequency
- Parallelism type, granularity, scaling characteristics, load balance
- I/O volume, frequency, pattern, method, desired percent of total runtime
- How science drives workload scaling: problem size, data set size, memory size

Example: Climate Modeling

- CAM dominates CCSM3 computational requirements.
- FV-CAM increasingly replacing Spectral-CAM in future CCSM runs.
- Drivers:
  - Critical support of U.S. submission to the Intergovernmental Panel on Climate Change (IPCC).
  - V & V for CCSM-4
- 0.5 deg resolution tending to 0.25

Climate Without INCITE

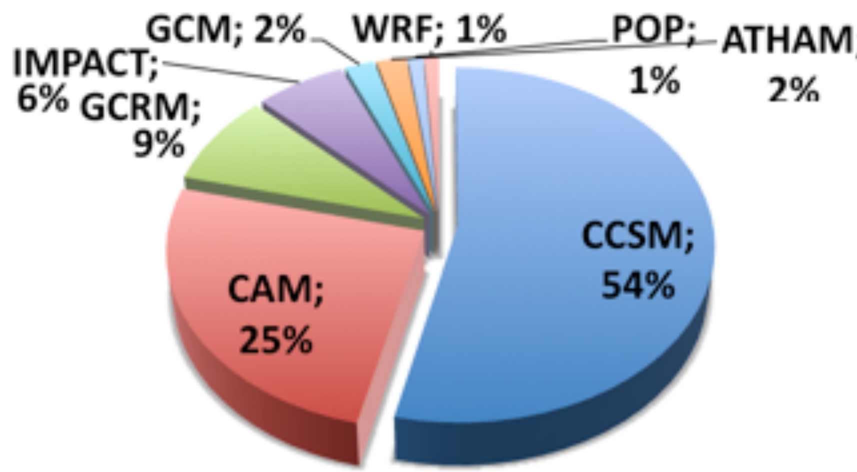

- Focus on ensemble runs - 10 simulations per ensemble, 5-25 ensembles per scenario, relatively small concurrencies.

FV-CAM Characteristics

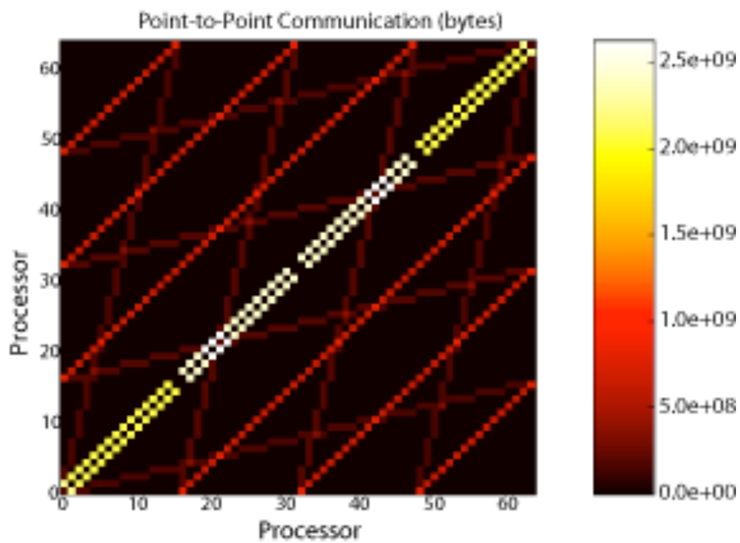

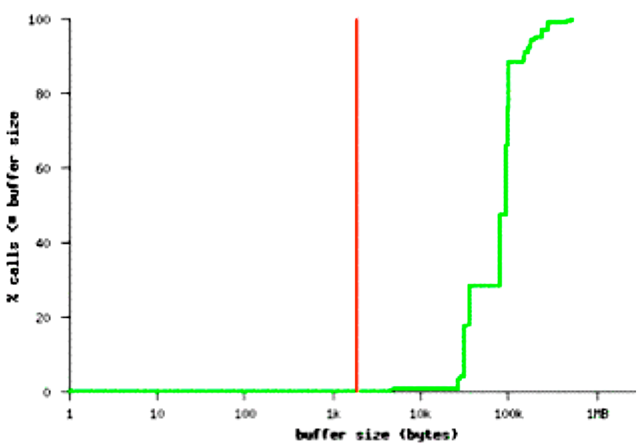

- Unusual interprocessor communication topology – stresses interconnect.
- Relatively low computational intensity – stresses memory subsystem.
- MPI messages in bandwidth-limited regime.
- Limited parallelism.

How Science Drives Architecture

Science areas \ Algorithm	Dense linear algebra	Sparse linear algebra	Spectral Methods (FFTs)	Particle Methods	Structured Grids	Unstructured or AMR Grids	Data Intensive
Accelerator Science		X	X	X	X	X	
Astrophysics	X	X	X	X	X	X	X
Chemistry	X	X	X	X			X
Climate			X		X	X	X
Combustion					X	X	X
Fusion	X	X		X	X	X	X
Lattice Gauge		X	X	X	X		
Material Science	X		X	X	X		
BioScience			X	X			X

Machine Requirements

Algorithm / Science areas	Dense linear algebra	Sparse linear algebra	Spectral Methods (FFT)s	Particle Methods	Structured Grids	Unstructured or AMR Grids	Data Intensive
Accelerator							
Astrophysics							
Chemistry							
Climate							
Combustion							
Fusion							
Lattice Gauge							
MatSci							
BioScience							
	High Flop/s rate	High performance memory system	High bisection bandwidth	High performance memory system	High flop/s rate	Low latency, efficient gather/scatter	Storage, Network Infrastructure

NERSC
NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

1km-Scale Global Climate Model Requirements

1km-Scale required to resolve clouds

- Simulate climate 1000x faster than real time
- 10 Petaflops sustained per simulation (~200 Pflops peak)
- 10-100 simulations (~20 Exaflops peak)
- DOE E3SGS report suggests exaflop requires 180MW

Computational Requirements:

- Advanced dynamics algorithms: icosahedral, cubed sphere, reduced mesh, etc.
- ~20 billion cells → 100 Terabytes of Memory
- Decomposed into ~20 million total subdomains → massive parallelism

200km (now)

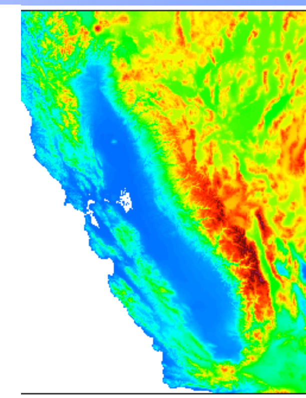

1km

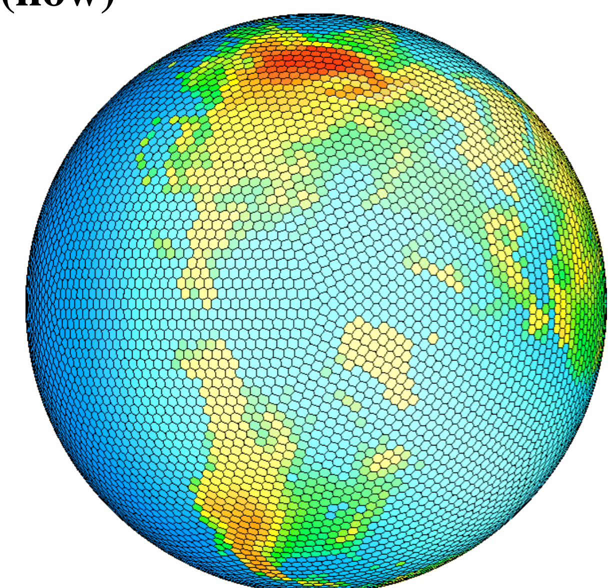

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB