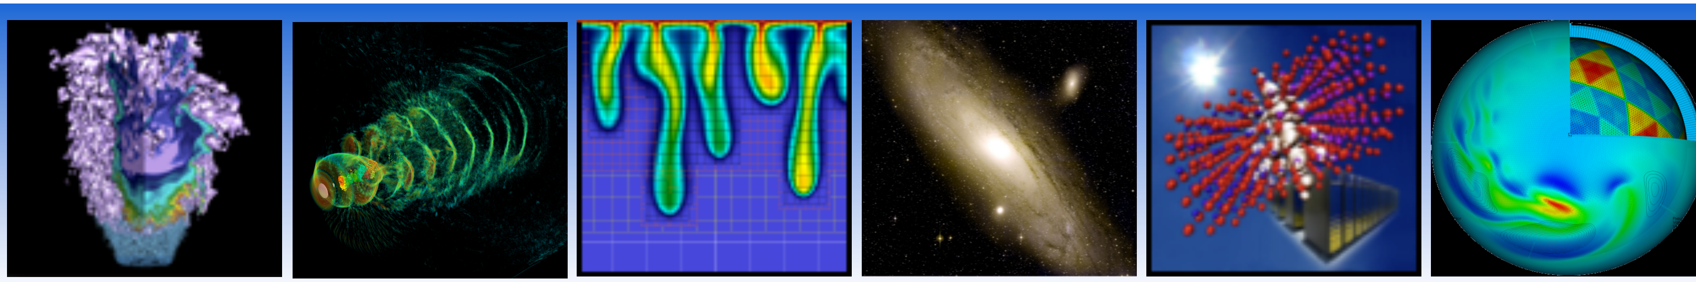

Meeting Goals & Process

Large Scale Computing and Storage Requirements for Biological and Environmental Research: Target 2017

Harvey Wasserman

NERSC User Services Group

Lawrence Berkeley National Laboratory

September 11-12, 2012

U.S. DEPARTMENT OF ENERGY | Office of Science

NERSC

BERKELEY LAB

Logistics: Schedule

- Agenda on workshop web page
  - http://www.nersc.gov/science/requirements/BER/meeting-agenda/
- Mid-morning / afternoon break, lunch
- Self-organization for dinner
- 3 science areas, one workshop
  - Science-focused but crosscutting discussion
  - Explore areas of common need (within BER)
- Wednesday: overview, review, key findings

Logistics: Final Report

- Final reports from 2009-2011 workshops (Target: 2014) on web
  - http://www.nersc.gov/science/requirements
- PI case studies + NERSC summary
- Target deadline: November 15
  - PI review prior

Logistics: Remote Presentation

- Need your view graphs in advance

Why is NERSC Collecting Computational Requirements?

- NERSC is science driven.
- Your input helps create the science-based justification for
  - acquiring the resources and
  - implementing the services that you need to reach your research goals
- Help NERSC make informed decisions for technology and services
  - guide procurements, staffing, and to improve the effectiveness of NERSC services.
- Result: NERSC can better provide what you need for your work

- Different from ERCAP:
  - Longer term focus
  - Not what you think you can get, but what you need

Science Areas

- Climate science
- Subsurface science / biogeochemistry
- Biological systems: molecular dynamics, protein-genome binding, DOE Systems Biology Knowledgebase (Kbase)

Science Areas

- Climate
  - Tom Bettge (NCAR)
  - William Collins (LBNL)
  - Stephen Price (LANL)
  - Ruby Leung, Jin-Ho Yoon (PNNL)
  - Christian Stan (GMU/COLA)
  - David Bader (LLNL)
  - Gilbert Compo (U. Colorado)
- Subsurface science / biogeochemistry
  - Tim Scheibe (PNNL)
- Biological Systems Science
  - Victor Markowitz, David Goodstein (LBNL/JGI)
  - Loukas Petridas (ORNL)
  - Mohammed AlQuraishi (Stanford (Harvard?))
  - Tom Brettin (ORNL)

Final Thoughts

- Purpose is not to justifiy science or approach
- We seek requirements not encumbered by "policy."
- Storage:
  - Scratch: output from runs
  - Project: shared code, data
  - HPSS Archive
- Benchmark code/mini-app/problem set to represent science area

Final Thoughts

- Key is to tie expected science outcome to computational need – as specifically as possible.
- We seek requirements not encumbered by "policy."
- Storage:
  - Scratch: output from runs
  - Project: shared code, data
  - HPSS Archive
- Mutually beneficial.

Scaling Science

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

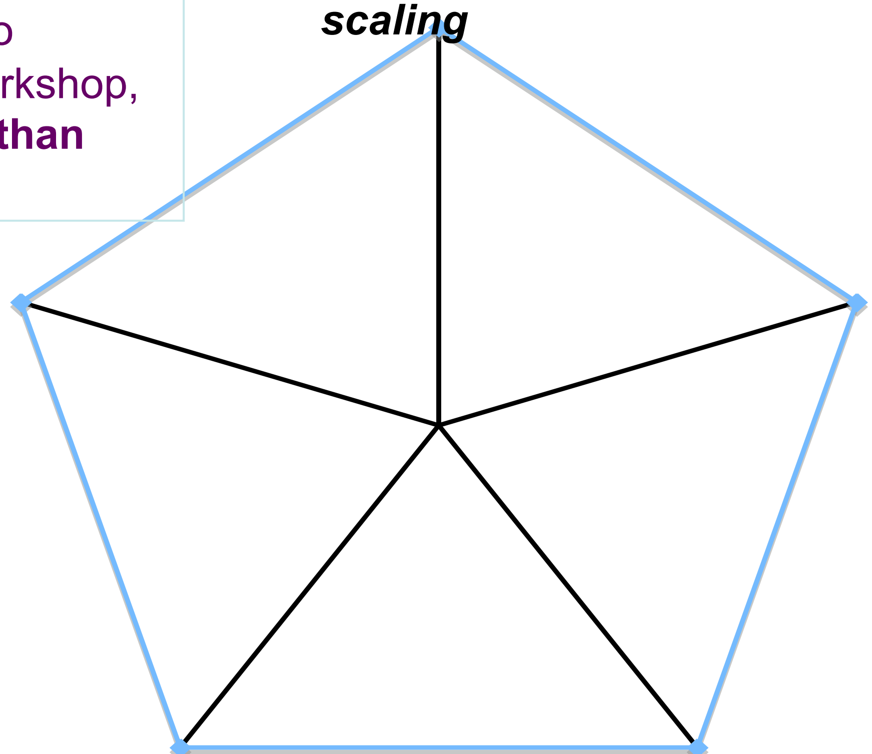

Scaling Science

HPC dimensions of Climate Prediction

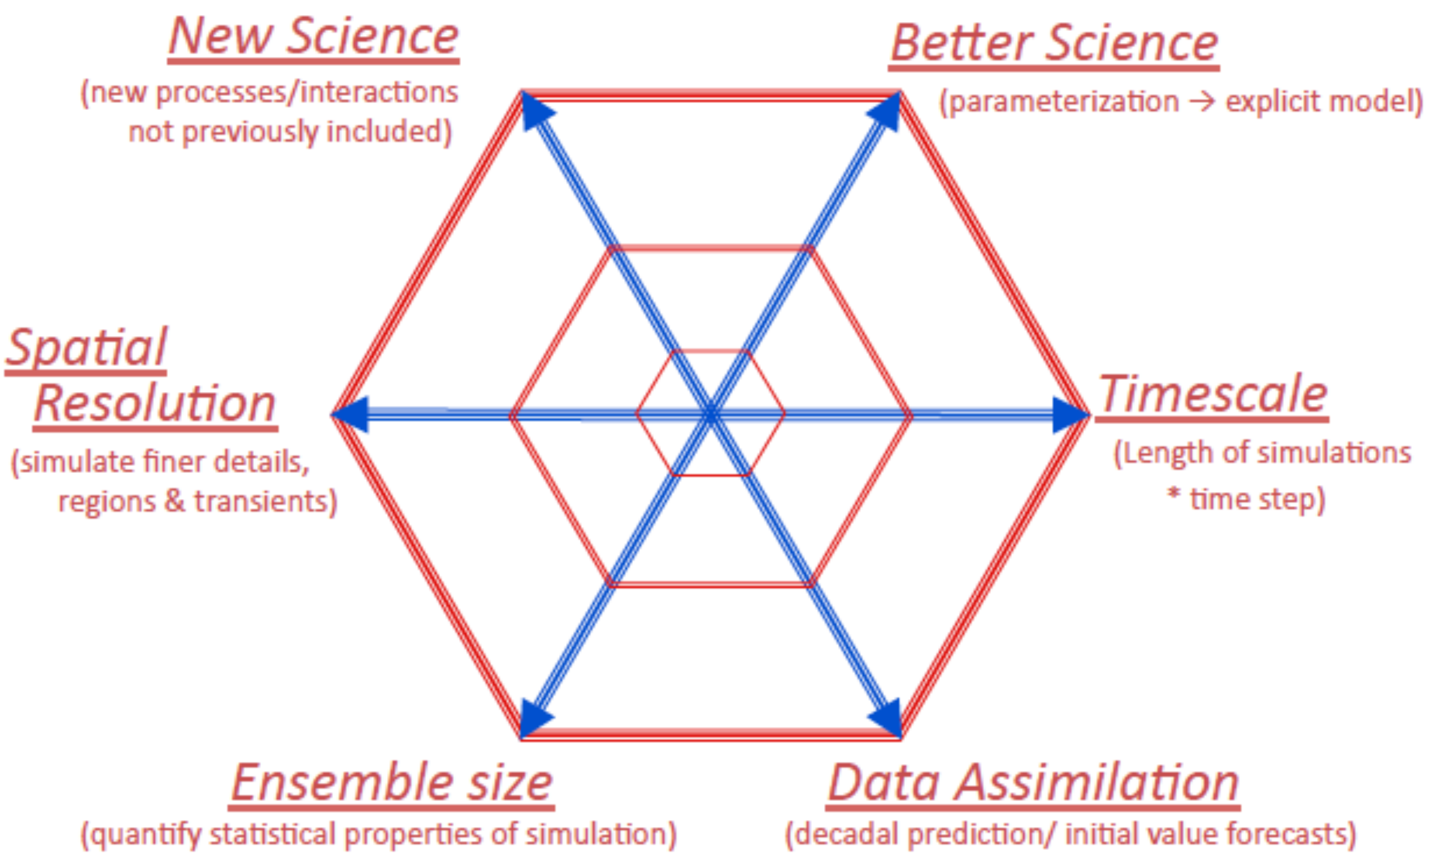

Lawrence Buja (NCAR) / Tim Palmer (ECMWF)

U.S. DEPARTMENT OF ENERGY | Office of Science

NERSC

BERKELEY LAB

Current Usage

PI / Project	Repo	2011	2012
Washington / CESM	mp9	47.6	30.3
Collins / CLIMES	m1204	4.2	2.2
Collins / IMPACTS	m1040	10.8	7.1
Collins / Ice Sheet	m1343	2.1	2.3
Leung / Frameworks	m1178	5.4	2.4
Stan / Multiscale	m1441	5.3	5.2
Stan / Super-Parameterization	m1576	-	9.6
Bader / PCMDI	mp193	6.8	3.0
Bader / CSSEF		-	-
Compo / SIRCA	m958	5.4	9.2
Scheibe / Subsurface	m749	3.9	2.4
JGI / Production	m342	5.6	4.6
JGI / IMG	m1045	10.4	18.1
Smith / MD Protein Dynamics	m906	12.0	5.6
McAdams / Transcription Factors	m926	2.1	0.5

BACKUP SLIDES

Workload Analysis

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

***Advanced Technologies Group**

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

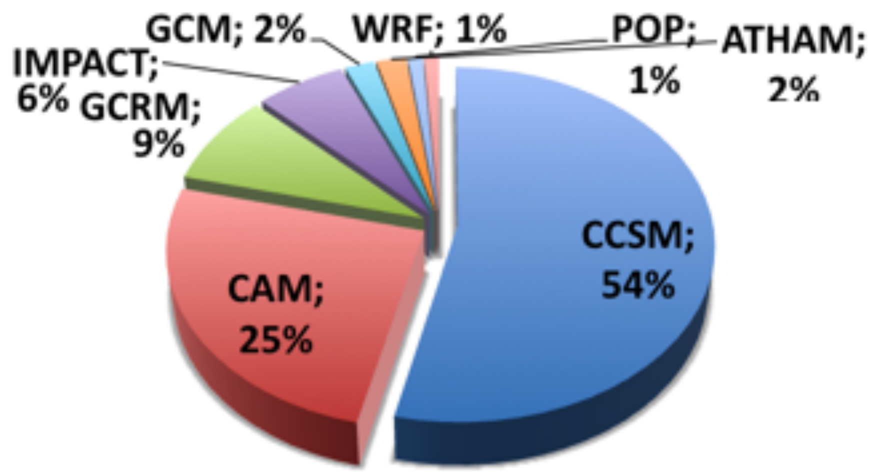

- Focus on ensemble runs - 10 simulations per ensemble, 5-25 ensembles per scenario, relatively small concurrencies.

FV-CAM Characteristics

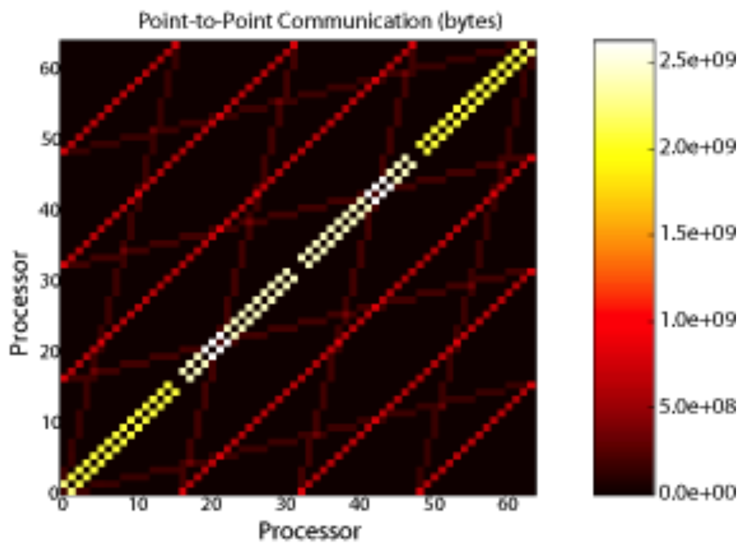

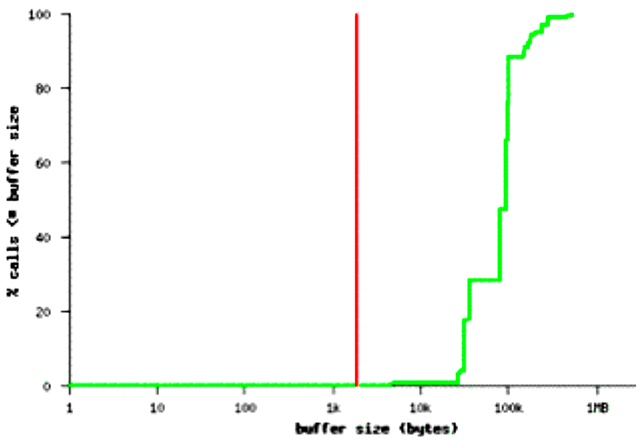

- Unusual interprocessor communication topology – stresses interconnect.
- Relatively low computational intensity – stresses memory subsystem.
- MPI messages in bandwidth-limited regime.
- Limited parallelism.

How Science Drives Architecture

Algorithm / Science areas	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

<table>
<tr><th>Algorithm / Science areas</th><th>Dense linear algebra</th><th>Sparse linear algebra</th><th>Spectral Methods (FFT)s</th><th>Particle Methods</th><th>Structured Grids</th><th>Unstructured or AMR Grids</th><th>Data Intensive</th></tr>
<tr><td>Accelerator</td><td rowspan="9">High Flop/s rate</td><td rowspan="9">High performance memory system</td><td rowspan="9">High bisection bandwidth</td><td rowspan="9">High performance memory system</td><td rowspan="9">High flop/s rate</td><td rowspan="9">Low latency, efficient gather/scatter</td><td rowspan="9">Storage, Network Infrastructure</td></tr>
<tr><td>Astrophysics</td></tr>
<tr><td>Chemistry</td></tr>
<tr><td>Climate</td></tr>
<tr><td>Combustion</td></tr>
<tr><td>Fusion</td></tr>
<tr><td>Lattice Gauge</td></tr>
<tr><td>MatSci</td></tr>
<tr><td>BioScience</td></tr>
</table>