Climate Applications Support at NERSC

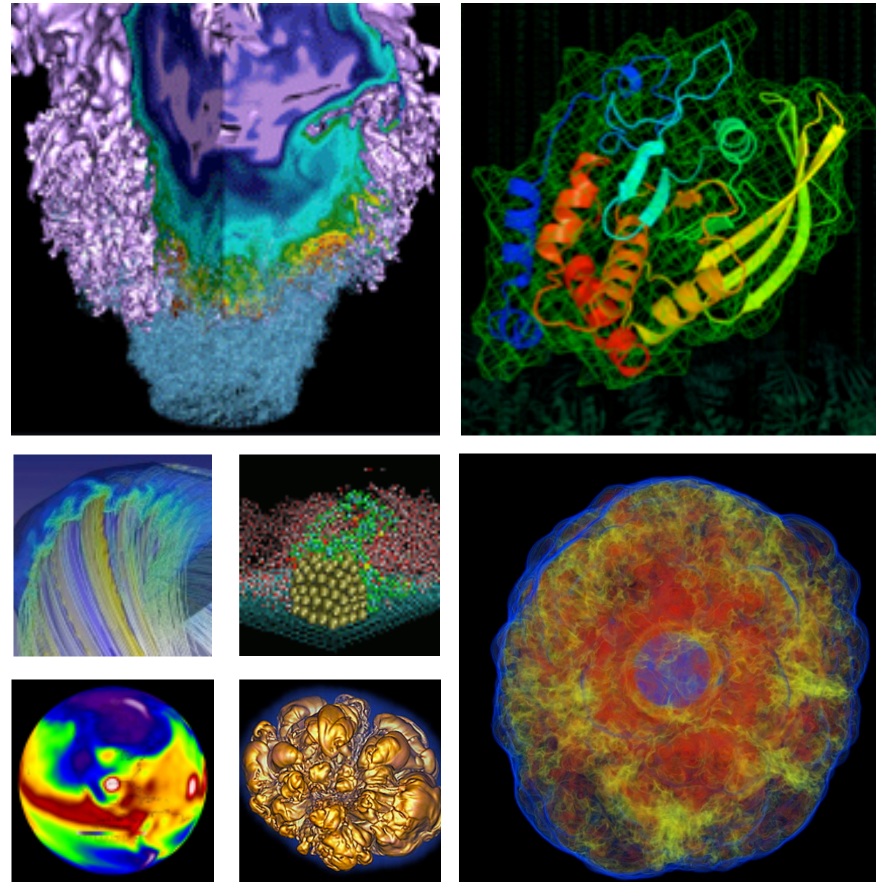

NERSC 40 YEARS at the FOREFRONT 1974-2014

Helen He

NERSC Climate PIs Telecon
March 16, 2016

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Climate Projects

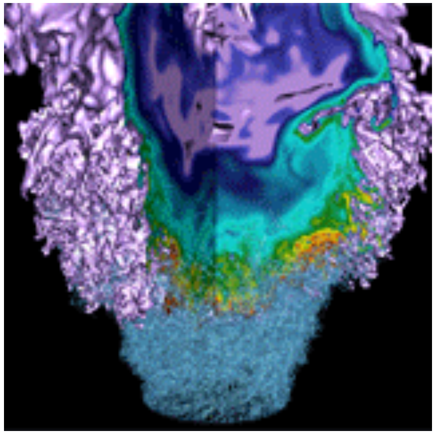

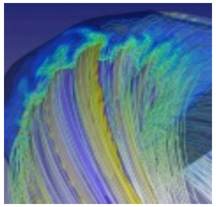

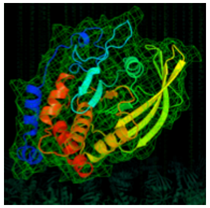

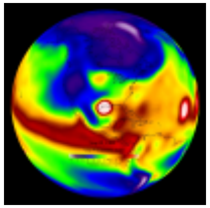

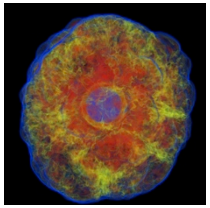

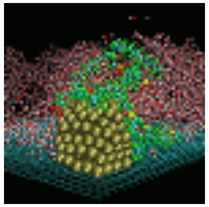

NERSC 40 YEARS at the FOREFRONT 1974-2014

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

NERSC Workload Diversity

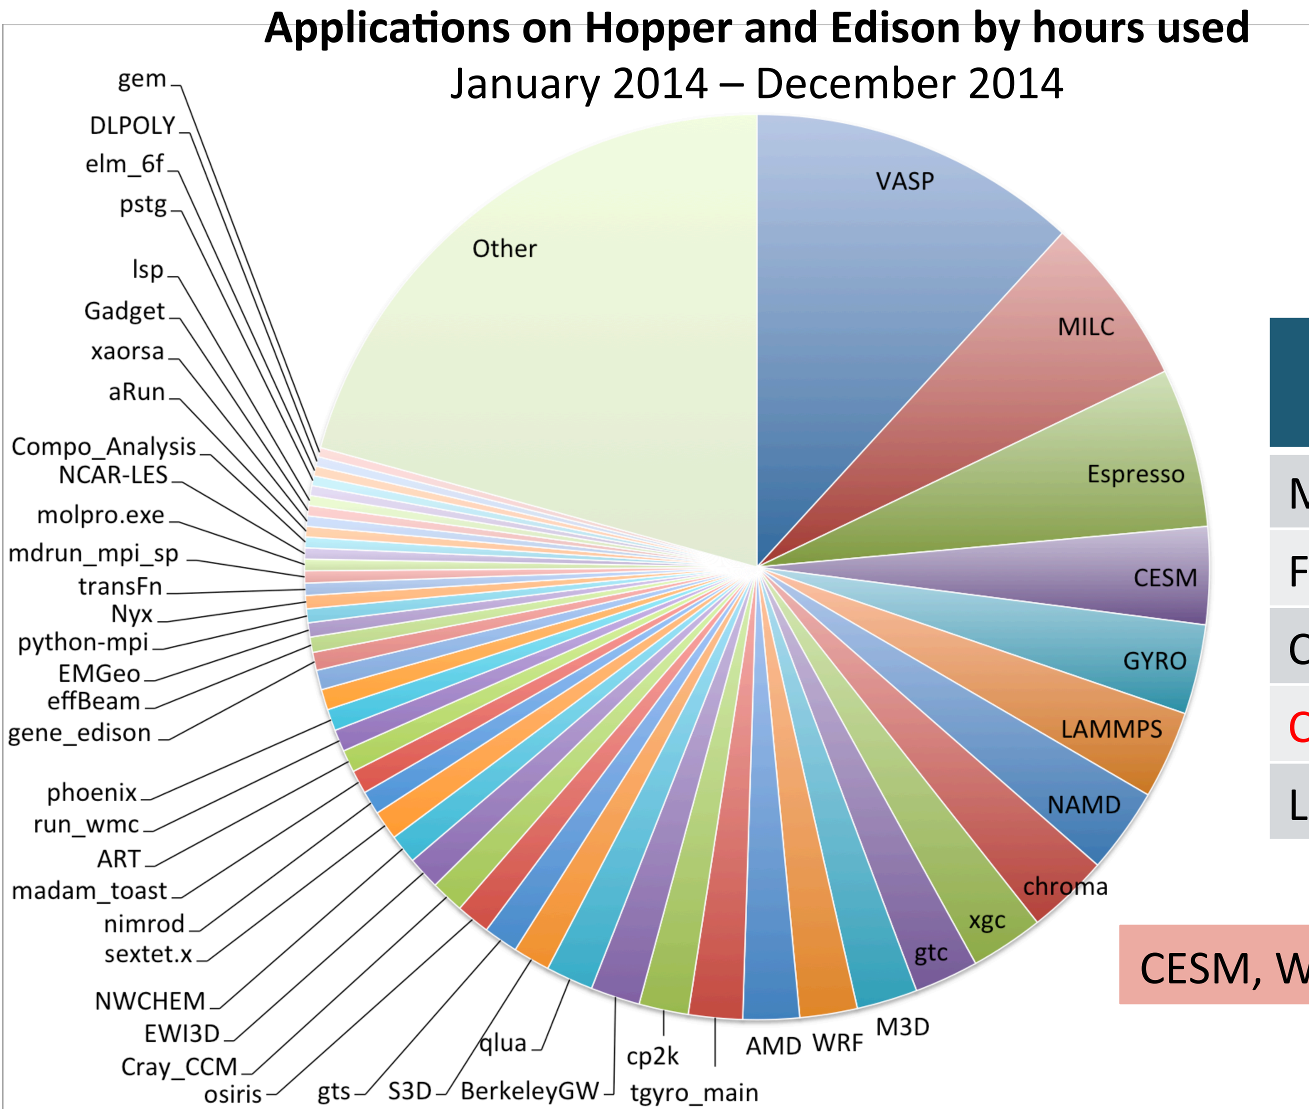

6,000 users

850 projects

Top 5 Science Categories by allocation (2014)	
Materials Science	20%
Fusion Energy	18%
Chemistry	12%
Climate Research	11%
Lattice QCD	11%

CESM, WRF, NCAR-LES among shown

70 Climate Projects at NERSC (AY2016)

- **Awards are published at:**
  - https://www.nersc.gov/users/accounts/awarded-projects/2016-allocation-awards/
- **See the detailed list of projects in "Extra Slides" (ordered by allocation hours).**
- **30 projects use CESM/ACME or CESM components. 254 active users**
- **17 projects use WRF. 85 users (36 active).**

AY16 Climate Allocations by PI's Institutions

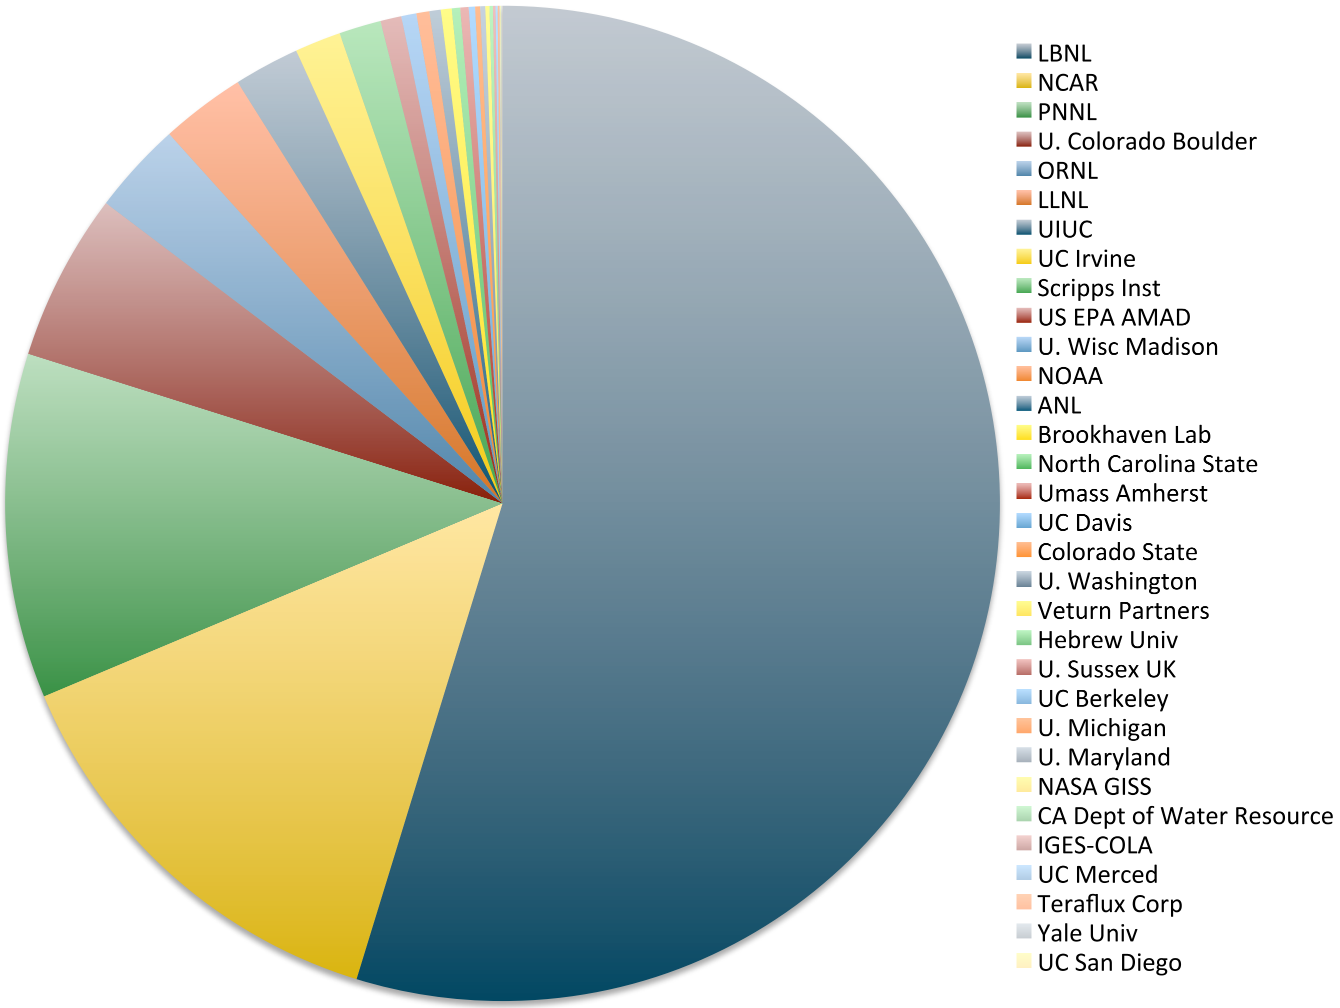

Climate Projects at NERSC (AY2016)-1

The last “Codes” column is obtained from answers to ERCAP request Question 20.

Repo	Project Title	PI	Organization	Codes
acme	Accelerated Climate for Energy	Bill Collins	LBNL	ACME
k_2	Cloud-System Simulations with a Multiscale Nonhydrostatic Global Atmospheric Model	Bill Collins	LBNL	ARAM
mp9	Climate Change Simulations with CESM: Moderate and High Resolution Studies	Gerald Meehl	NCAR	CESM
m958	Ocean Atmosphere Reanalyses for Climate Applications (OARCA) 1831-2015	Gil Compo	U Colorado, Boulder	Forecast Model, Ensemble Filter, …
m1704	Multiscale Methods for Accurate, Efficient, and Scale-Aware Models of the Earth System	Bill Collins	LBNL	CESM, Chombo AMR Dycore
m1867	Water Cycle and Climate Extremes Modeling (WACCEM)	Ruby Leung	PNNL	MPAS, WRF-CLM
m1199	High-Latitude Application and Testing (HiLAT) of Global and Regioanl Climate Models	Philip Rasch	PNNL	CESM
m2467	Quantifying and Reducing Biogeochemical and Aerosol Feedback Uncertainties	Forrest Hoffman	ORNL	ACME, CESM
m1041	Predicting Ice Sheet and Climate Evolution at Extreme Scales	Esmond Ng	LBNL	FELIX, CISM, BISICLES, POP2x
mp193	Program for Climate Model Diagnosis and Intercomparison	His-Yen Ma	LLNL	CESM, CAM5, ACME, SAM

CESM Support

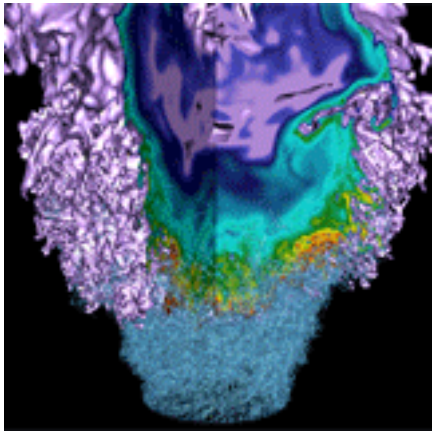

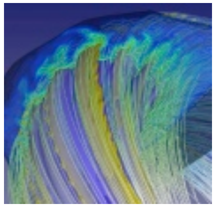

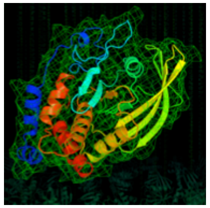

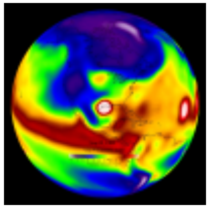

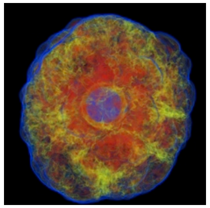

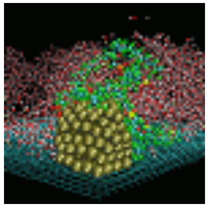

CESM Support

- Proactively work with NCAR to prepare CESM port and configuration changes for upcoming NERSC systems, and Programming Environment updates on existing systems.
- Send to 240+ active CESM users multiple times a year on recommended configurations to run on NERSC systems (Hopper, Edison, Cori)
- Investigate build and run issues for individual users. Help users to adapt to new NERSC system environment.
- Help to improve workflow for best throughput. Bundle jobs, dependency jobs. Optimize job sizes and walltime requests.

CESM on Cori P1 and Edison (1)

- The CESM/1.2.2 version ported by NCAR/NERSC to Cori P1 (new machine) and Edison with SLURM batch scheduler.
- Source code distribution and scripts: /project/projectdirs/ccsm1/collections/cesm1_2_2/. Use similar scripts as before to create new cases by using the "-mach corip1"or "-mach edison" option.
- Functional port on Cori P1 only, not tuned for optimal pe-layouts or performance. Based on popular demand and resources, NCAR/NERSC can help to tune pe-layouts for a few specific model configurations.
- You can tune the settings in "env_mach_pes.xml" yourself by adjusting the NTASKS_xxx and NTHRDS_xxx values and then redo the cesm_setup steps afterwards.

CESM on Cori P1 and Edison (2)

- Older CESM versions are not ported to Cori P1 or Edison. Using Edison as an example, you can get “mkbatch.edison”, “env_mach_specific.edison” from the cesm1_2_2 Machines directory along with the edison entry in "config_machines.xml" and copy them into your version.
- Users with Developers access can use cesm1_5_beta04 or newer. More cases there with pe-layout tuned.
- Several problems related to model run crashes with specific model configurations have been worked out with users and NCAR. CESM source codes in NERSC /project have been modified.
- Work with users on compilation and run time issues with specific model configurations on case by case basis.

CESM Survey, May 2014

- Survey conducted to help NERSC and NCAR understand CESM usage model for system procurement.
- Biggest projects are at NCAR and LBNL; others include PNNL, LLNL. Wide range of science goals, including IPCC production runs, extreme weather events prediction, model development, abrupt climate change simulation, and more
- Job profile as a function of model component varies: CAM ~50-85%.
- Thread parallelism also varies, using 1, 2, 3, 4, and 6 threads per MPI per run for CAM and POP
- Parallelism typically driven by model configuration, component efficiencies, queue time, allocation size, and other considerations
- Key point: due to the large # of projects and diverse science goals, it is impossible to distinguish any one simulator that is dominant at NERSC

Typical CESM User Jobs, May 2014 (1/6)

- **Repo mp9, user aaa, IPCC AR5**
  - Typical 104-node job on Hopper:
    - aprun -n 828 -N 8 -d 3 ./cesm.exe
    - 2603 MB/task
    - #   cpl ntasks=128 nthreads=1 rootpe= 0 ninst=1
    - #   cam ntasks=768 nthreads=3 rootpe= 0 ninst=1
    - #   clm ntasks=128 nthreads=3 rootpe= 0 ninst=1
    - #   cice ntasks=640 nthreads=3 rootpe=128 ninst=1
    - #   pop2 ntasks= 60 nthreads=3 rootpe=768 ninst=1
    - #   sglc ntasks=768 nthreads=3 rootpe= 0 ninst=1
    - #   swav ntasks=768 nthreads=3 rootpe= 0 ninst=1
    - #   rtm ntasks=128 nthreads=3 rootpe= 0 ninst=1

To run on Edison with SLURM now:

```
#SBATCH -N 104
…
setenv OMP_NUM_THREADS 3
srun -n 828 -c 3 ./cesm.exe
```

TOT	Run	Time:	10413.674 seconds	28.531 seconds/mday	8.3 myears/wday
LND	Run	Time:	422.124 seconds	1.157 seconds/mday	204.68 myears/wday
ICE	Run	Time:	756.239 seconds	2.072 seconds/mday	114.25 myears/wday
ATM	Run	Time:	8673.825 seconds	23.764 seconds/mday	9.96 myears/wday
OCN	Run	Time:	2268.196 seconds	6.214 seconds/mday	38.09 myears/wday
GLC	Run	Time:	0 seconds	0 seconds/mday	0 myears/wday
CPL	Run	Time:	1379.231 seconds	3.779 seconds/mday	62.64 myears/wday

Task distribution reflects optimal load balance among model components.

NERSC Exascale Science Application Program (NESAP)

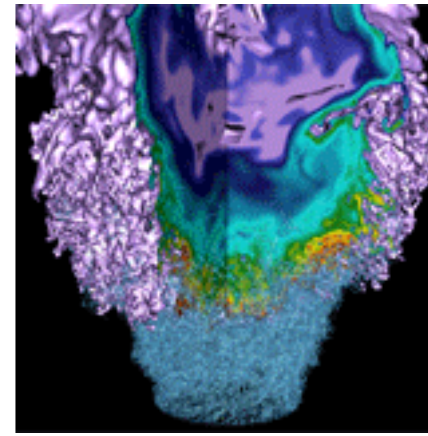

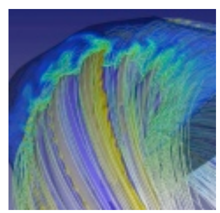

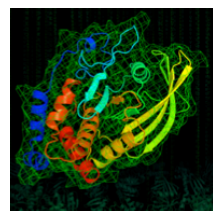

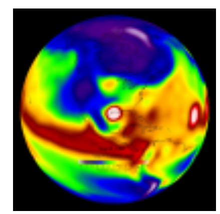

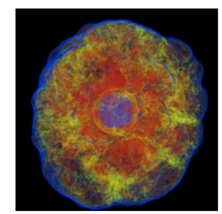

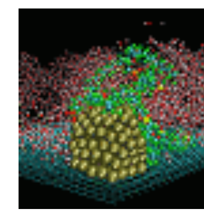

NERSC 40 YEARS at the FOREFRONT 1974-2014

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

NESAP Climate Projects

- **The goal for NESAP is to prepare DOE Office of Science user community for Cori Phase 2, the Intel Xeon Phi Knights Landing (KNL) manycore architecture.**
- **Selected Climate NESAP Projects are:**
  - ACME (PI: Hans Johansen, LBNL)
  - CESM (PI: John Dennis, NCAR)
  - MPAS (PI: Todd Ringler, LANL)
  - WRF (PI: John Michalakes, NOAA)
- **We have postdoc positions available to work on NESAP projects. Please help to spread the word.**
  - Position description at: http://cs.lbl.gov/careers-and-fellowships
    Choose "NERSC Only", Requisition ID is 81356.

CESM NESAP Case Study

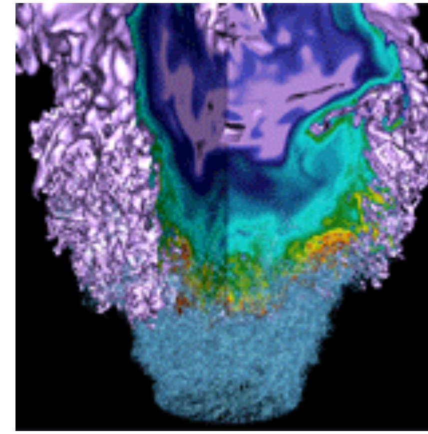

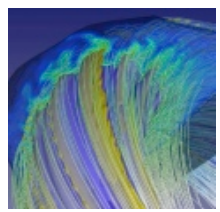

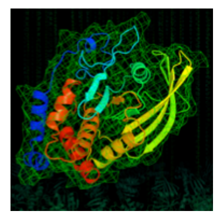

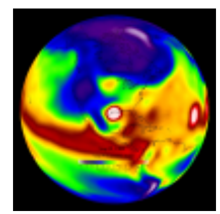

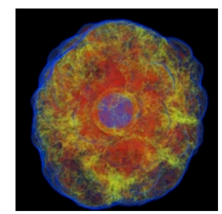

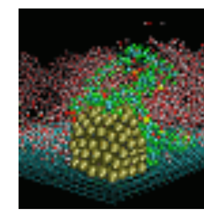

CESM NESAP Team Members

NCAR: John Dennis (PI), Chris Kerr, Sean Santos

Cray: Marcus Wagner

Intel: Nadezhda Plotnikova, Martyn Corden

NERSC Liaison: Helen He

MG2 Kernel

- **MG2 is a kernel for CESM that represents version 2 of the Morrison-Gettleman micro-physics package. Typically consumes about 10% of CESM run time.**
- **Kernel is core bound**
  - Not bandwidth limited at all
  - Shows very little vectorization
    - Some loop bounds are short (e.g. 10)
    - Dependent sequence of instructions
  - Heavy use of math instrinsics that do not vectorize
- **Kernel has long complex loops with interleaved conditionals and elemental function calls.**

MG2 Vectorization Prototype

- **Use compiler report to check and make sure key functions are vectorized (and all functions on the call stack are vectorized too)**
- **Add !$OMP DECLARE SIMD and !DIR$ ATTRIBUTE FORCEINLINE when needed.**

Example call stack for vectorization and inlining

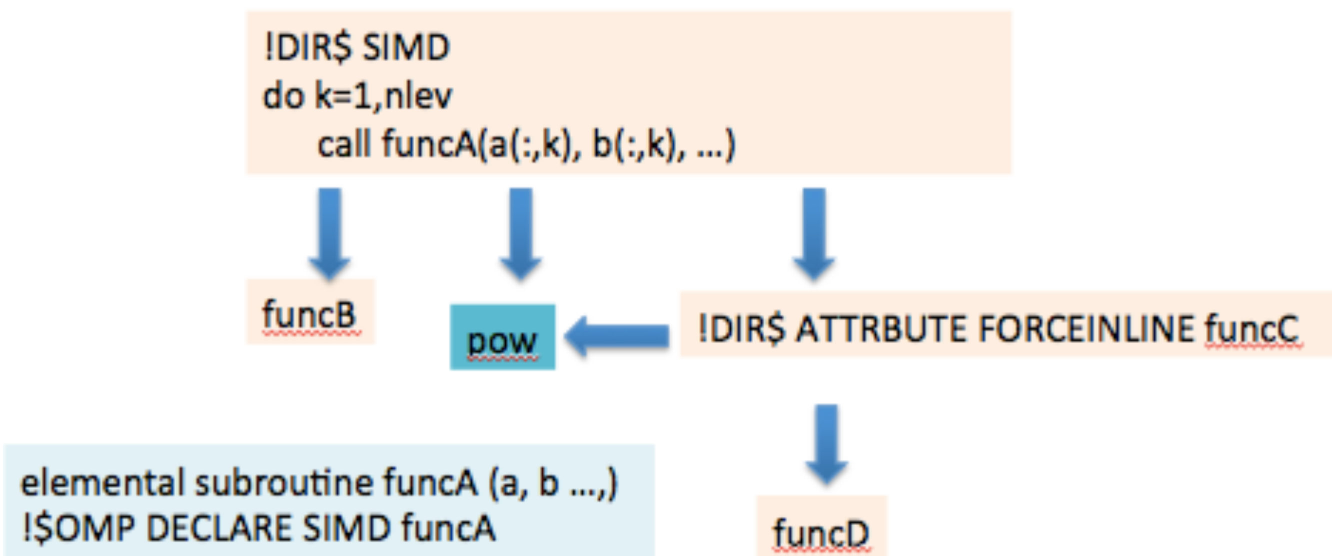

MG2 Optimization Steps

Version 1

- Simplify expressions to minimize #operations
- Use internal GAMMA function

Version 2

- Remove “elemental” attribute, move loop inside.
- Inline subroutines. Divide, fuse, exchange loops.
- Replace assumed shaped arrays with loops
- Replace division with inversion of multiplication
- Remove initialization of loops to be overwritten later
- Use more aggressive compiler flags. Try different compilers.
- Use profile-guided optimization (PGO)

Version 3 (Intel compiler only)

- Use !$OMP SIMD ALIGNED to force vectorization

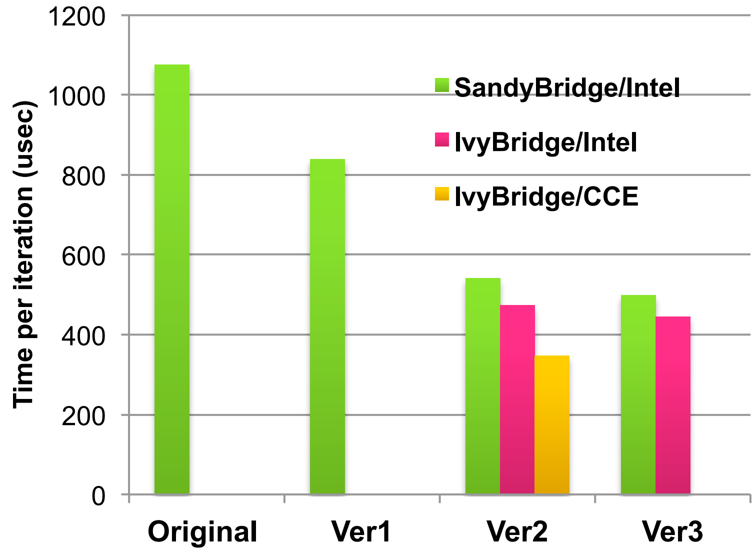

MG2 Summary

- Directives and flags can be helpful, however not a replacement for programmers' work on code modifications.
- Break up loops and push loops into functions where vectorization can be dealt with directly and can expose logic to compiler.
- Incremental improvements not necessary a BIG win from any one thing. Accumulative results matter.
- Performance and portability is a major goal: use !$OMP SIMD proves to be beneficial but very hard to use regarding the need of providing the aligned list.
- Requested optional alignment declaration in Fortran language standard.

WRF Support

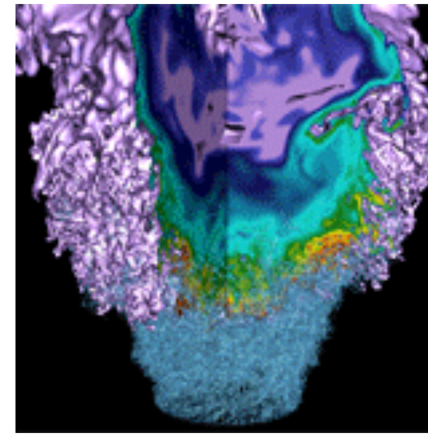

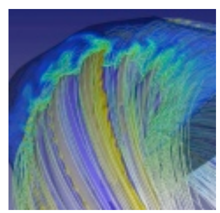

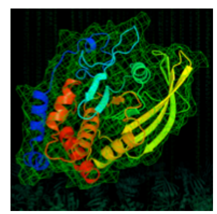

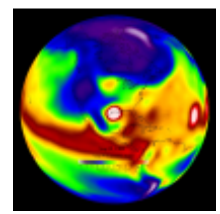

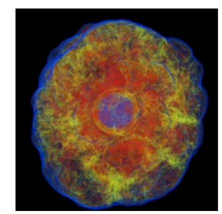

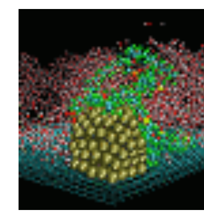

WRF Support

- User survey conducted in Oct 2015. 14 responses (~40% from active users). Most users use version 3.5.1 and up, some has own customizations, 100 to 500 cores with pure MPI. 2 users use MPI/OpenMP.
- Helped individual users to build different WRF versions with different compilers on Hopper and Edison.
- WRF is a Tier3 NESAP application.
- WRF team has configuration for KNC; Will release a KNL version to use for Cori.
- The KNL version will be optimal with hybrid MPI/OpenMP.
- Encourage users to explore hybrid (dm+sm) on Edison and Cori now.

What NERSC Can Help

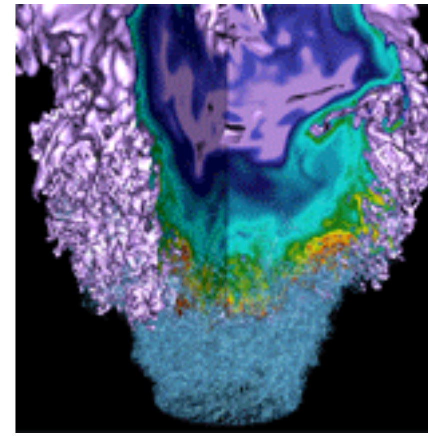

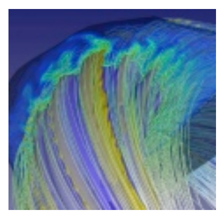

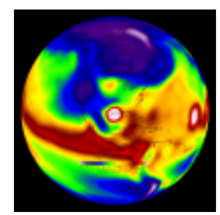

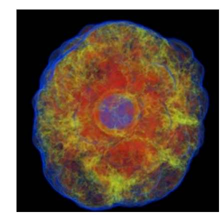

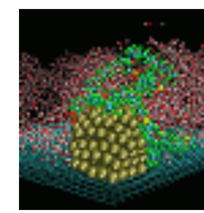

We can offer help on

- **Porting and optimizing applications**
- **Workflow optimization and queue analysis, job bundle advice etc. for better job turnaround**
- **Compute reservations**
  - https://www.nersc.gov/users/queues/reservation-form/
- **Queue boost and use special queue**
- **Network tuning for large data transfer**
- **Science gateway (portals) for climate data**
- **Large data storage at NERSC HPSS archive system**

20th Century Reanalysis Project

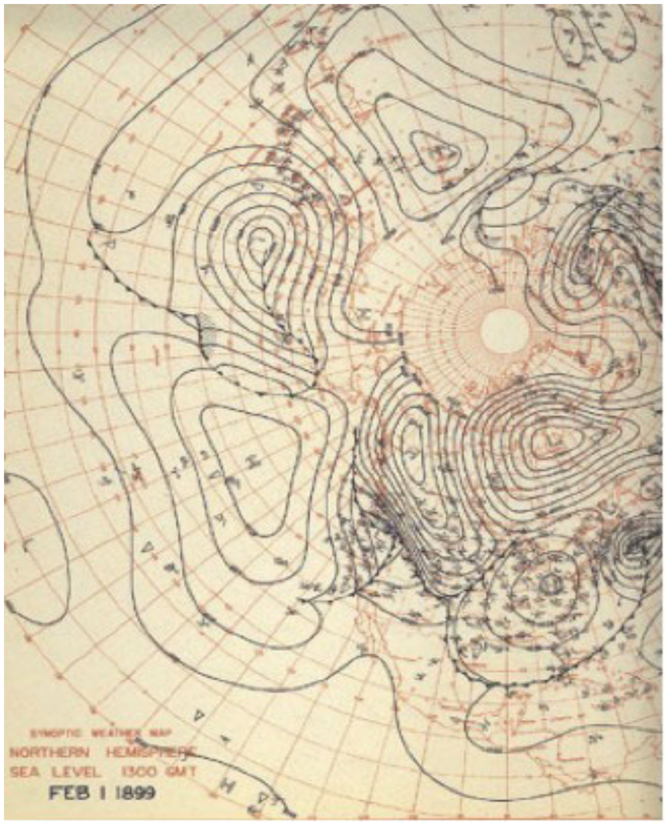

Global Tropospheric Circulation Maps

- **PI: Gil Compo, University of Colorado, Boulder**
- **NERSC helped on:**
  - Provide options for data transfer between NERSC and NCAR.
  - Tuning and debugging the application
  - Adding multi-level parallelism to bundle several associated parallel jobs together to obtain high throughput and simplify data handling
  - Provide special queue access and increase max job limit on machines.

Global Ensemble Reforecast Project

- **PI: Tom Hamill, NOAA.**
- **NERSC helped on:**
  - Data transfer from NERSC to NOAA
  - Porting from IBM to Cray systems
  - Optimizing an MPI reduction function (run time reduced from 24 min to a few seconds)
  - Suggesting post-processing workflow optimizations
  - Details at https://www.nersc.gov/news-publications/nersc-news/nersc-center-news/2012/end-to-end-network-tuning-sends-data-screaming-from-nersc-to-noaa/
  - Web Gateway for Global Ensemble Reforecast Data
    - http://portal.nersc.gov/project/refcst/v2/
    - This portal allows users to download selected days of the full Reforecast model output

Do You Have a "realtime" Use Case?

- **Example user cases are:**
  - Realtime data processing with experiments or events.
  - Realtime data analysis from users of science gateway portals. Often short jobs that need an immediate response
  - Software development/test of workflow systems with realtime queue
  - Visualization jobs - this would replace the functionality traditionally offered by a vis node and allows for interactive visualization tasks
- **SLURM on Cori now supports a "realtime" partition.**
  - Users can request a small number of on-demand nodes if their jobs have special needs that cannot be accommodated through the regular batch system.
  - Request Form at: https://www.nersc.gov/users/queues/realtime-queue-request-form/
  - Requires DOE approval to use "realtime" on Cori

We Welcome Science Highlights

- **Sample science highlights can be found at:**
  - https://www.nersc.gov/news-publications/publications-reports/science-highlights-presentations/
  - A couple of examples also shown in “Extra Slides”
- **A new web form to submit your research to us at:**
  - https://www.nersc.gov/science/share-your-research/

Tools and Libraries

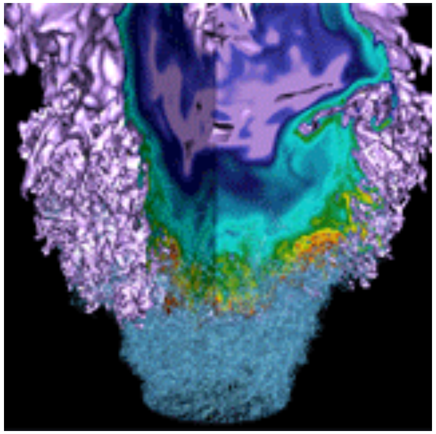

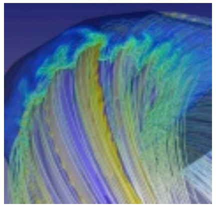

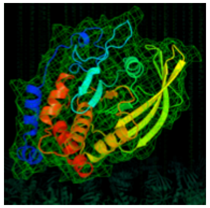

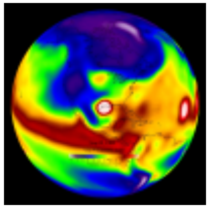

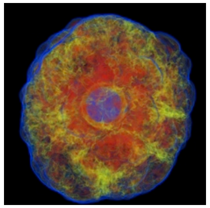

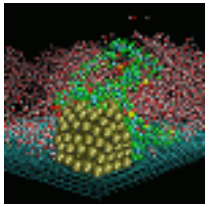

Available Tools and Libraries for Climate Applications

- nco
- cdo
- ncview
- ncl (ncar)
- nccmp
- netcdf4-python
- cray-netcdf, cray-parallel-netcdf
- cray-hdf5, cray-parallel-hdf5
- esmf

Summary

- **Facilitate information exchange among NERSC climate users.**
- **Keep the communication channels open between users and NERSC/DOE.**
- **Please let us know:**
  - your science and computational accomplishment
  - your need for running climate applications and performing climate data analysis at NERSC

Extra Slides

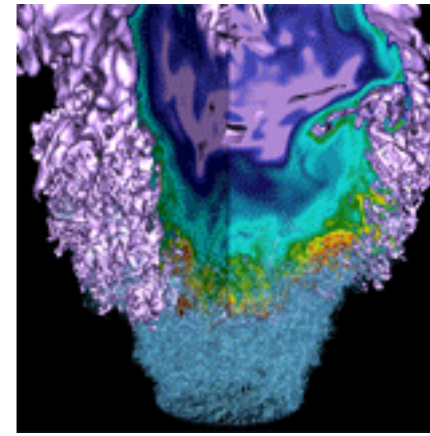

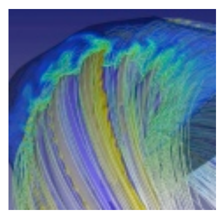

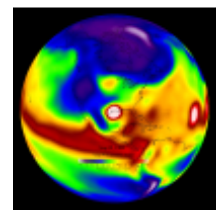

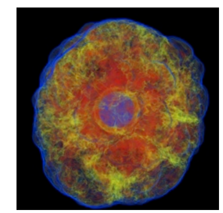

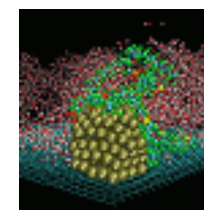

Climate Projects at NERSC (AY2016)-2

Repo	Project Title	PI	Organization	Codes
m1517	Calibrated and Systematic Characterization Attribution and Detection of Extremes	Travis O'Brien	LBNL	CESM
m1355	Analysis of Global Coupled 0.1-deg POP/CICE in the CESM Framework	Julie Mcclean	Scripps Inst	CESM
m1795	Projections of Ice Sheet and Ocean Evolution	Stephen Price	LANL	POPCICLES, CISM, BICICLES
m2071	Determining impacts of small ice crystals on bulk scattering properties of ice clouds to improve representations of ice clouds on satellite retrieval algorithms and numerical models.	Junshik Um	U. Illinois U-C	Amsterdam discrete dipole approximation
m2222	Expanding the computational frontier of multi-scale atmospheric simulation to advance understanding of low cloud / climate feedbacks	Michael Prtichard	UC Irvine	SPCESM
sobl	Southern Ocean Uptake in the MPAS-Ocean Model	Edward Patton	NCAR	NCAR-LES
m2098	Calculations of single-scattering properties of small ice crystals to improve in-situ measurements, satellite retrieval algorithms, and numerical models	Greg Mcfarquhar	U. Illinois U-C	ADDA
m1657	Environmental effects in the lifecycle of convective clouds	Mikhail Ovchinnikov	PNNL	WRF, SAM
m2410	Development of the LES ARM Symbiotic Simulation and Observation (LASSO) Workflow	William Gustafson	PNNL	SAM, WRF

Climate Projects at NERSC (AY2016)-3

Repo	Project Title	PI	Organization	Codes
m1204	Center at LBNL for Integrative Modeling of the Earth System (CLIMES)	Bill Collins	LBNL	CESM
m1231	Evaluation of the Interactions Among Tropospheric Aerosol Loading, Radiative Balance, Clouds, and Precipitation	Jonathan Pleim	US EPA AMAD	WRF-CMAQ, WRF
m2320	Consequences of cloud super-parameterization for land-atmosphere coupling physics	Michael Pritchard	UC Irvine	SPCAM, UPCAM
m1642	Evolution in Cloud Population Statistics of the MJO	Samson Hagos	PNNL	ARWRF
m2402	Next-Generation Intensity-Duration-Frequency Curves Considering Spatiotemporal Non-Stationarity in Climate, Intense Precipitation Events, and Snowmelt	Ruby Leung	PNNL	WRF
m1196	Interactions of Clouds, Convection, and Climate	David Romps	LBNL	DAM
m2067	Evaluation of the large-scale and regional climatic response across North Africa to natural variability in oceanic modes and terrestrial vegetation among the CMIP5 models	Michael Notaro	U. Wisc. Madison	CESM
m1660	Regional Modeling of Land-Ocean-Atmosphere Interactions	Larry Beng	PNNL	WRF, WRF-Chem, CLM
m1602	Investigating the Role of Biogeochemical Processes in the Northern High Latitudes Using an Earth System Model	Atul Jain	U. Illinois U-C	Integrated Science Assessment Model

Climate Projects at NERSC (AY2016)-4

Repo	Project Title	PI	Organization	Codes
m2190	Next Generation Global Prediction System (NGGPS) Benchmarking	John Michalales	NOAA	NGGPS benchmark suite
m2136	Evaluation and improvement of Convective Parameterizations in ACME model	Wuyin Lin	BNL	ACME
m1929	A Multiscale Reduced-Order Method for Integrated Earth System Modeling	George Pau	LBNL	CESM-PFLOTRAN-ROM, PROMEs
m1637	Reducing the Uncertainty in Simulating Aerosol Chemistry over Multiple Spatial Scales	Jerome Fast	PNNL	WRF
m1372	High resolution model development to quantify the impact of icebergs on the stability of the Atlantic Meridional Overturning Circulation	Alan Condron	UMass Amherst	MITgcm
m1439	Earth System Modeling	Yang Zhang	North Carolina State	WRF/CAM5, WRF/CMAQ,WRF/Chem, WRF/Chem-ROMS
m2250	Major improvements on the longwave radiative interactions between surface and clouds in the Polar Regions in atmospheric global circulation model (GCM)	Daniel Feldman	LBNL	CESM, LBLRTM-DISORT
m1540	Dynamic Downscaling of Climate Projections for DoD Installations	Rao Kotamarthi	ANL	NRCM (WRF), RegCM4
m2249	Integrating Climate Change into Air Quality Modeling	Ajith Kaduwela	UC Davis	CMAQ, WRF

Climate Projects at NERSC (AY2016)-5

Repo	Project Title	PI	Organization	Codes
m328	Global cloud modeling	David Randall	Colorado State	CESM, GCRM, SAM, Super-CAM
m1822	Using AMIE data to study cloud processes within the Madden-Julian Oscillation	Robert Houze	U. Washington	WRF
m2451	Integrated Predictive Systems for Solar Energy with Modeling, Post Processing and Machine Learning	Cameron Whiteman	Vertum Partners	WRF
m2478	Linked Rainfall and Runoff Intensity-Duration-Frequency in face of Climate Change and Uncertainty	Eugene Yan	ANL	WRF
m1481	Improvement of Representation of the Cloud Aerosol Interaction in Large-Scale Models	Alexander Khain	Hebrew Univ Jerusalem	WRF
m726	Simulating the tropical cloud-climate interactions and MJO during the AMIE and DYNAMO field experiments using the NCAR WRF model	Yi Wang	U. Sussex UK	AMIE-WRF, CESM-UM
m2076	Role of Forest Ecosystems in Mitigating Climate Change Impacts through Land-Surface Energy and Water Processes	Guangshan Chen	U. Wisc. Madison	RegCM4-CNDV
m1865	Predictability of the carbon-climate system on seasonal to decadal time scales	Inez Fung	UC Berkeley	CCSM
mp79	Parallel modeling of climate, chemistry, and physics affecting the global atmosphere	Donald Wuebbles	U. Illinois U-C	CESM, CAM-Chem, WACCM

Climate Projects at NERSC (AY2016)-6

Repo	Project Title	PI	Organization	Codes
mp231	Three-dimensional Global Atmospheric Aerosol and Chemistry Modeling	Joyce Penner	U. Michigan	IMPACT, GCE cloud model
m2352	Marine boundary layer clouds simulated by SPCAM5-CLUBB	Steven Ghan	PNNL	SPCAM5
m1590	Regional Aerosol Emission and Its Direct and Cloud-Related Forcing Estimates	Tami Bond	U. Illinois U-C	CESM
m1006	Dimention Reduction of the Cloud-Aerosol-Radiation (CAR) Ensemble Modeling System	Xin-Zhong Liang	U. Maryland	CWRF/CAR
m411	Interaction of Atmospheric Chemistry and Aerosols with Climate	Philip Cameron-Smith	LLNL	CESM, IMPACT
m843	Integrating observations and simulations of shallow warm clouds to advance understanding of drizzle formation and evolution	Ann Fridlind	NASA GISS	DHARMA
m640	Influences of the Boundary Layer Flow on Vegetation-Air Exchanges of Energy, Water and Greenhouse Gases	Xuhui Lee	NCAR	NCAR-LES
m1576	Reducing Uncertainty of Climate Simulations Using the Super-Parameterization	Cristiana Stan	IGES - COLA	SP-CCSM
m2319	Using carbonyl sulfide to quantify regional terrestrial biological carbon fluxes	John Campbell	UC Merced	STEM

Climate Projects at NERSC (AY2016)-7

Repo	Project Title	PI	Organization	Codes
m997	Using the California Central Valley Groundwater-Surface Water Simulation Model to study impacts of extended drought on groundwater levels, surface water flows, and agricultural management practices.	Charles Brush	CA Dept of Water Resource	IWFM, PEST
refcst	A multi-decadal reforecast data set to improve weather forecasts for renewable energy applications	Thomas Hamill	NOAA	GFS
m2492	Paleo-megadroughts and Abrupt Climate Changes in the Speleothem Records	Inez Fung	UC Berkeley	CAM
m2420	Next Generation Ecosystem Experiments -- Tropics	Charles Koven	LBNL	CLM, CESM, ACME
m184	High-Spatial Resolution Models of Atmospheric Chemistry, Aerosols for Regional Scale Climate Simulations	Rao Kotamarthi	ANL	CESM, CLM, CMAQ, GEOS-Chem
m1178	Development of Frameworks for Robust Regional Climate Modeling	Ruby Leung	PNNL	WRF-CLM-ROMS, MPAS-A/MPAS-O, HOMME-POP-ROMS
m1357	CAM5 Parameter Sensitivity and Uncertainty Quantification for CSSEF	Yun Qian	PNNL	CAM5.1
m612	Continuous Dynamic Grid Adaptation in Regional and Global Atmospheric Models	Joseph Prusa	Teraflux Corp	CEU, EULAG
m1269	A multi-decadal reforecast data set to improve weather forecasts for renewable energy applications	Nigel Quinn	LBNL	HydroGeosphere

Climate Projects at NERSC (AY2016)-8

Repo	Project Title	PI	Organization	Codes
m2095	Influence of the Boundary Layer Flow on Vegetation-Air Exchange of Energy, Water and Carbon Dioxide	Jianping Huang	Yale Univ	NCAR-LES
m1861	Towards parameterization of root-rock hydrologic interactions in the Earth System Model	Inez Fung	UC Berkeley	CLM4/CLM4RR
m1374	Multiscale Modeling of Aerosol Indirect Effects on Decadal Timescales	Steven Ghan	PNNL	CESM, SPCAM
m543	River Transport and Hydrology in CCSM	Marcia Branstetter	ORNL	CESM
m238	Theoretical Calculation of Water Vapor Continuum Absorption	Andrew Lacis	NASA GISS	atm
m1060	North Pacific Mesoscale Coupled Air-Ocean Simulations Compared with Observations	Ivana Cerovecki	UCSD	POP2, WRF

Typical CESM User Jobs, May 2014 (2/6)

- **Repo m1040, ccc, IPCC runs**
- **304 jobs on Edison, 6 to 104 nodes, up to 6 hrs**
- **Typical 52-node job:**
  - aprun -j 2 -n 828 -N 8 -d 3 ./cesm.exe
  - 1214 MB/task
  - # cpl ntasks=128 nthreads=1 rootpe=0 ninst=1
  - # cam ntasks=768 nthreads=3 rootpe=0 ninst=1
  - # clm ntasks=128 nthreads=3 rootpe=0 ninst=1
  - # cice ntasks=640 nthreads=3 rootpe=0 ninst=1
  - # pop2 ntasks=60 nthreads=3 rootpe=768 ninst=1
  - # sglc ntasks=1 nthreads=1 rootpe=0 ninst=1
  - # swav ntasks=128 nthreads=1 rootpe=0 ninst=1
  - # rtm ntasks=128 nthreads=3 rootpe=0 ninst=1

To run on Edison with SLURM:

#SBATCH -N 52
…
setenv OMP_NUM_THREADS 3
srun -n 828 -c 3 ./cesm.exe

- **Typical 21-node job:**
  - aprun -n 504 -N 24 -d 1 ./cesm.exe
  - # cpl ntasks=480 nthreads=1 rootpe=0 ninst=1
  - # datm ntasks=24 nthreads=1 rootpe=480 ninst=1
  - # clm ntasks=480 nthreads=1 rootpe=0 ninst=1
  - # sice ntasks=480 nthreads=1 rootpe=0 ninst=1
  - # socn ntasks=480 nthreads=1 rootpe=0 ninst=1
  - # sglc ntasks=480 nthreads=1 rootpe=0 ninst=1
  - # swav ntasks=480 nthreads=1 rootpe=0 ninst=1
  - # rtm ntasks=480 nthreads=1 rootpe=0 ninst=1

To run on Edison with SLURM:

#SBATCH -N 21
…
setenv OMP_NUM_THREADS 1
srun -n 504 ./cesm.exe

Typical CESM User Jobs, May 2014 (3/6)

- **Repo mp9, user aaa, IPCC AR5**
  - 116 jobs on Hopper, 86 to 104 nodes, up to 36 hrs
  - Typical 86-node job:
    - aprun -n 1680 -N 24 -d 1 ./ccsm.exe : -n 192 -N 12 -d 2 ./ccsm.exe
    - 579 MB/task
    - # total number of tasks = 1872
    - # maximum threads per task = 2
    - # cpl ntasks=1680 nthreads=1 rootpe=0
    - # cam ntasks=1664 nthreads=1 rootpe=0
    - # clm ntasks=400 nthreads=1 rootpe=1280
    - # cice ntasks=1280 nthreads=1 rootpe=0
    - # pop2 ntasks=192 nthreads=2 rootpe=1680
    - # sglc ntasks=1 nthreads=1 rootpe=0

MPMD using different num_threads for different executables is not yet supported in SLURM.

Timing profile for one run:

TOT	Run	Time:	104327.625 seconds	11.433 seconds/mday	20.7 myears/wday
LND	Run	Time:	5020.151 seconds	0.55 seconds/mday	430.27 myears/wday
ICE	Run	Time:	10427.512 seconds	1.143 seconds/mday	207.14 myears/wday
ATM	Run	Time:	55907.355 seconds	6.127 seconds/mday	38.64 myears/wday
OCN	Run	Time:	80272.484 seconds	8.797 seconds/mday	26.91 myears/wday
GLC	Run	Time:	0 seconds	0 seconds/mday	0 myears/wday
CPL	Run	Time:	0 seconds	0 seconds/mday	0 myears/wday

Typical CESM User Jobs, May 2014 (4/6)

- **Repo mp9, user aaa, IPCC AR5**
  - 9 jobs on Edison, 16 to 84 nodes, up to 30 hrs
  - Typical 84-node job:
    - aprun -j 2 -S 6 --cc numa_node -n 1008 -N 12 -d 4 ./ccsm.exe
    - 1188 MB/task
    - #   cpl ntasks=960 nthreads=4 rootpe=0 ninst=1
    - #   cam ntasks=960 nthreads=4 rootpe=0 ninst=1
    - #   clm ntasks=48 nthreads=4 rootpe=0 ninst=1
    - #   cice ntasks=912 nthreads=4 rootpe=48 ninst=1
    - #   pop2 ntasks=48 nthreads=4 rootpe=960 ninst=1
    - #   sglc ntasks=1 nthreads=4 rootpe=0 ninst=1
    - #   swav ntasks=1 nthreads=4 rootpe=0 ninst=1
    - #   rtm ntasks=48 nthreads=4 rootpe=0 ninst=1

To run on Edison with SLURM:

#SBATCH -N 84
…
setenv OMP_NUM_THREADS 4
srun -n 1008 -c 4 ./cesm.exe

Typical CESM User Jobs, May 2014 (5/6)

- **Repos cascade, mp193, user mmm**
- **26 jobs on Hopper for cascade, 18 jobs for mp193, 1 to 854 nodes, up to 48 hrs**
  - Typical 854-node job:
    - aprun -n 20480 ./tstorms.exe filelist 16
  - Typical 320-node job
    - aprun -n 1280 -N 4 -d 6 ./ccsm.exe
    - 708 MB/task
    - # cpl ntasks=1280 nthreads=6 rootpe=0
    - # cam ntasks=1280 nthreads=6 rootpe=0
    - # clm ntasks=1280 nthreads=6 rootpe=0
    - # cice ntasks=1280 nthreads=6 rootpe=0
    - # docn ntasks=1280 nthreads=6 rootpe=0
    - # sglc ntasks=1280 nthreads=6 rootpe=0
  - 160 jobs on Edison, 84 nodes
    - aprun -j 2 -n 1008 -N 12 -d 4 -S 6 --cc numa_node ./cesm.exe
    - 1315 MB/task

To run on Edison with SLURM:

```
#SBATCH -N 320
…
setenv OMP_NUM_THREADS 6
srun -n 1280 -c 6 ./cesm.exe
```

Sample Edison Timing Profile

TOT	Run	Time:	16555.161 seconds	22.678 seconds/mday	10.44 myears/wday
LND	Run	Time:	2476.78 seconds	3.393 seconds/mday	69.77 myears/wday
ROF	Run	Time:	15.698 seconds	0.022 seconds/mday	11007.77 myears/wday
ICE	Run	Time:	1492.282 seconds	2.044 seconds/mday	115.8 myears/wday
ATM	Run	Time:	11063.755 seconds	15.156 seconds/mday	15.62 myears/wday
OCN	Run	Time:	10997.98 seconds	15.066 seconds/mday	15.71 myears/wday
GLC	Run	Time:	0 seconds	0 seconds/mday	0 myears/wday
WAV	Run	Time:	0 seconds	0 seconds/mday	0 myears/wday
CPL	Run	Time:	1806.178 seconds	2.474 seconds/mday	95.67 myears/wday
CPL	COMM	Time:	5557.865 seconds	7.614 seconds/mday	31.09 myears/wday

Typical CESM User Jobs, May 2014 (6/6)

- **Repo m1199, user jjj**
- **1203 ccsm/cesm jobs, 1826 atm jobs, 1 to 400 nodes**
- **Typical 400-node job:**
  - aprun -n 4800 -N 12 -d 2 ./cesm.exe
  - 668 MB/task
  - Wall time: 6 to 8 hrs
  - #   cpl ntasks=4800 nthreads=2 rootpe=0 ninst=1
  - #   cam ntasks=4800 nthreads=2 rootpe=0 ninst=1
  - #   clm ntasks=4800 nthreads=2 rootpe=0 ninst=1
  - #   cice ntasks=4800 nthreads=2 rootpe=0 ninst=1
  - #   docn ntasks=4800 nthreads=2 rootpe=0 ninst=1
  - #   sglc ntasks=4800 nthreads=2 rootpe=0 ninst=1
  - #   swav ntasks=1024 nthreads=1 rootpe=0 ninst=1
  - #   rtm ntasks=1024 nthreads=1 rootpe=0 ninst=1

To run on Edison with SLURM:

```
#SBATCH -N 400
…
setenv OMP_NUM_THREADS 2
srun -n 4800 -c 2 ./cesm.exe
```

Supercomputers and the Future of Wind Energy

Scientific Achievement

Success in simulations to assess impact of complex terrain on wind turbine efficiency is now leading to assessment of climate change and lightening effects on wind power production and revenue.

Significance and Impact

Electric utilities are keenly interested in evaluating the effects of climate change and extreme weather on wind energy.

Research Details

- Uses the WRF code to solve a 4D model of the atmosphere, including all basic dynamics and energy equations, while allowing for targeted parameterizations and interaction with the land and water bodies.
- NERSC resources allow evaluation of the sensitivity of the model results to different land surface types, atmospheric datasets and other forcings both internal and external to the model.

Capps and Whiteman, 15th Annual WRF Users' Workshop

This work has been using **4-D next-generation mesoscale numerical weather simulations** that model complex atmospheric data resolved to local wind farm scales. This will allow for optimal turbine site selection and evaluating the true climate change effects to wind energy.

Work was performed at LBNL using NERSC

Impact on Climate Simulations

- ***M. E. Maltrud and E. C. Hunke, Los Alamos National Laboratory; J. L. McClean, Naval Postgraduate School***

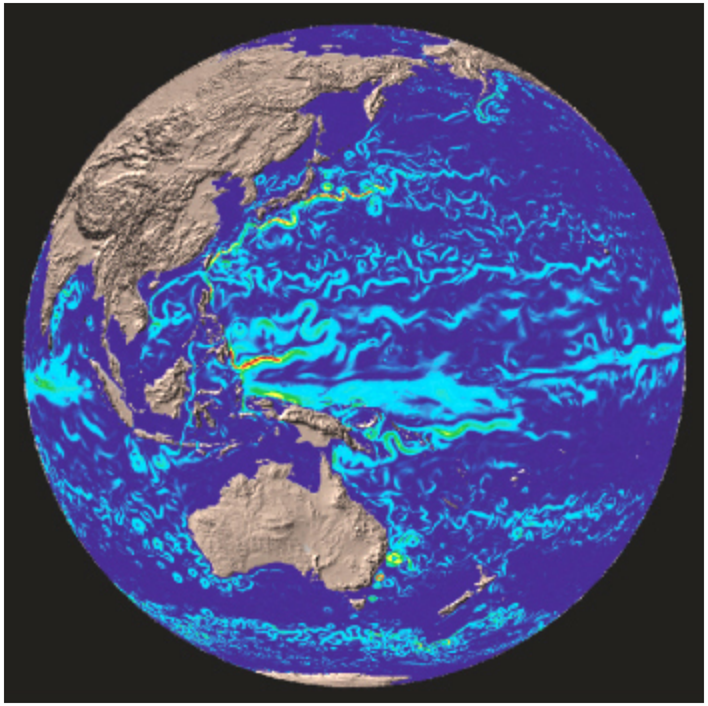

- **High-Resolution Global Coupled Ocean/ Sea Ice Modeling using POP**
- **NERSC provides**
  - **Environment for highly parallel runs**
  - **Service and stability**
  - **Special queue support**

Thank you.