Climate Applications at NERSC

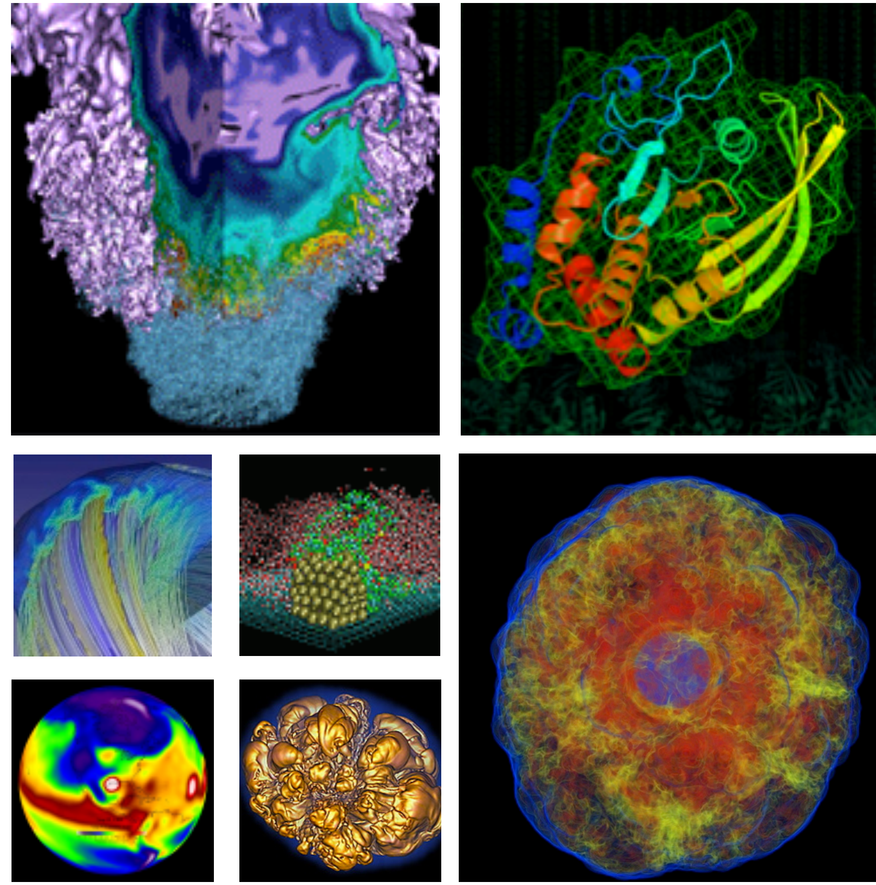

NERSC 40 YEARS at the FOREFRONT 1974-2014

Helen He

NERSC Climate PIs Telecon
Dec 4, 2015

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

Climate Projects

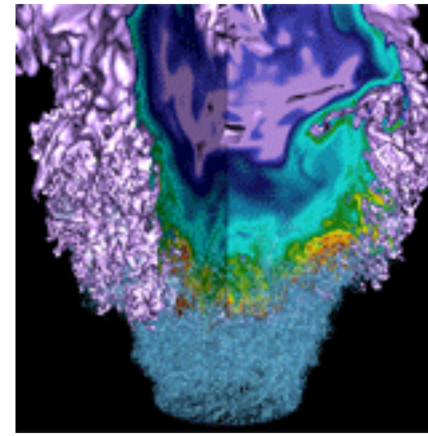

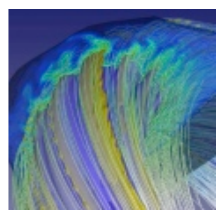

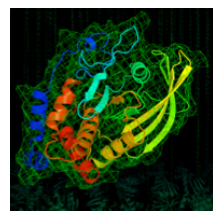

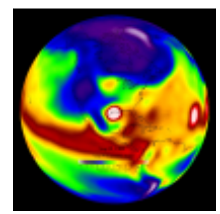

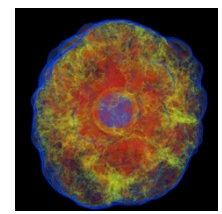

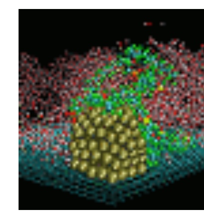

NERSC 40 YEARS at the FOREFRONT 1974-2014

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

75 Climate Projects at NERSC (AY2015)

- **Awards are published at:**
  - https://www.nersc.gov/users/accounts/awarded-projects/2015-allocation-awards/
- **Or you can search in NIM**

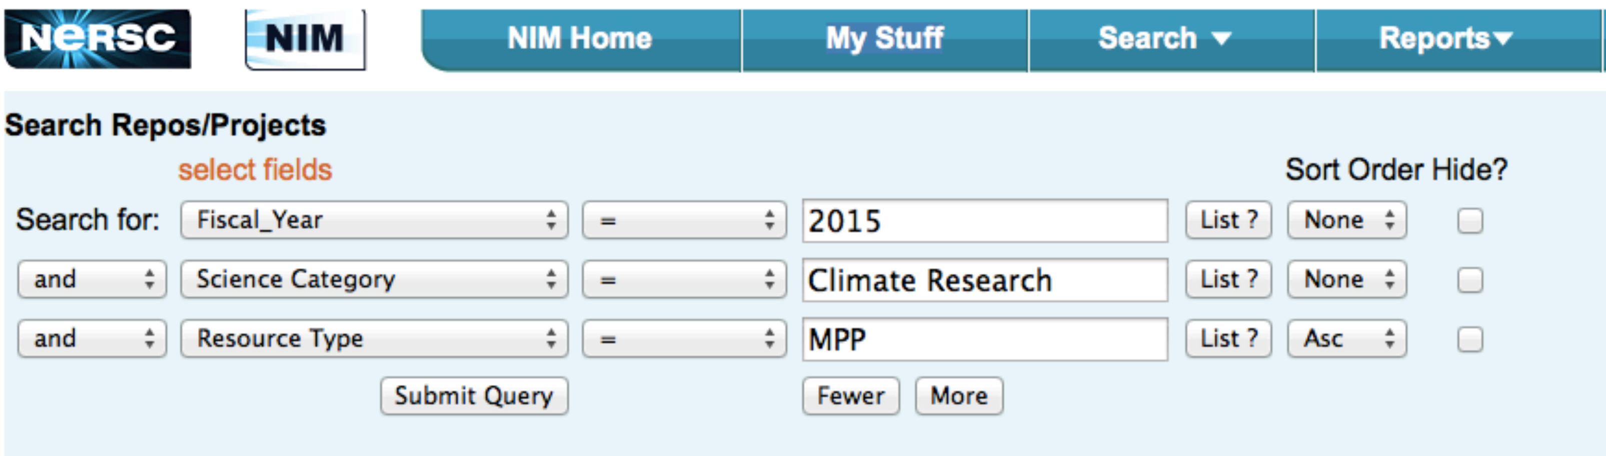

- **29 projects use CESM or CESM components. 247 users**
- **16 projects use WRF. 36 users.**

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Climate Projects at NERSC (AY2015)-1

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

Repo	Project Title	PI	Organization	Codes
mp9	Climate Change Simulations with CESM: Moderate and High Resolution Studies	Gerald Meehl	NCAR	CESM
m1517	Calibrated and Systematic Characterization Attribution and Detection of Extremes	Travis O'Brien	LBNL	CESM
m958	Ocean Atmosphere Reanalyses for Climate Applications (OARCA) 1831-2015	Gil Compo	U Colorado, Boulder	Forecast Model, Ensemble Filter, …
k_2	Cloud-System Simulations with a Multiscale Nonhydrostatic Global Atmospheric Model	Bill Collins	LBNL	ARAM
m1204	Center at LBNL for Integrative Modeling of the Earth System (CLIMES)	Bill Collins	LBNL	CESM
m1867	Water Cycle and Climate Extremes Modeling (WACCEM)	Ruby Leung	PNNL	MPAS, WRF-CLM
m1704	Multiscale Methods for Accurate, Efficient, and Scale-Aware Models of the Earth System	Bill Collins	LBNL	CESM, Chombo AMR Dycore
acme	Accelerated Climate for Energy	Bill Collins	LBNL	ACME
m1657	Environmental effects in the lifecycle of convective clouds	Mikhail Ovchinnikov	PNNL	WRF, SAM
mp193	Program for Climate Model Diagnosis and Intercomparison	His-Yen Ma	LLNL	CESM, CAM5, ACME, SAM

Climate Projects at NERSC (AY2015)-2

Repo	Project Title	PI	Organization	Codes
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
m1199	High-Latitude Application and Testing (HiLAT) of Global and Regioanl Climate Models	Philip Rasch	PNNL	CESM
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	---
m1041	Predicting Ice Sheet and Climate Evolution at Extreme Scales	Esmond Ng	LBNL	FELIX, CISM, BICICLES, POP2x
m1231	Evaluation of the Interactions Among Tropospheric Aerosol Loading, Radiative Balance, Clouds, and Precipitation	Jonathan Pleim	US EPA AMAD	WRF-CMAQ, WRF
m1795	Projections of Ice Sheet and Ocean Evolution	Stephen Price	LANL	POPCICLES, CISM, BICICLES
m1642	Evolution in Cloud Population Statistics of the MJO	Samson Hagos	PNNL	ARWRF
m1540	Dynamic Downscaling of Climate Projections for DoD Installations	Rao Kotamarthi	ANL	NRCM (WRF), RegCM4
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

CESM

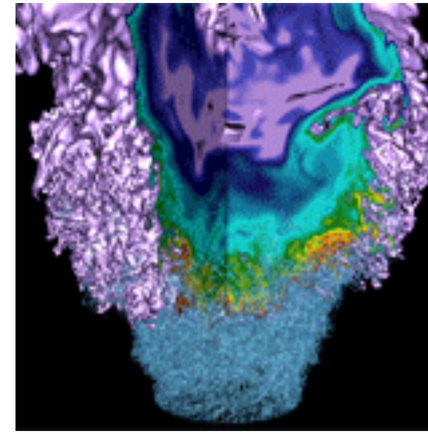

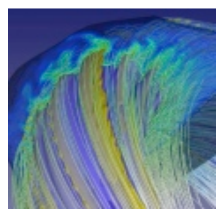

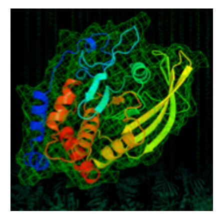

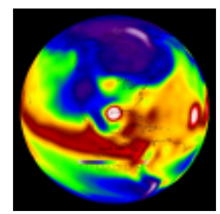

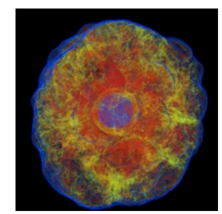

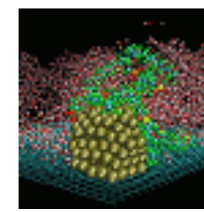

CESM Support

- 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+ 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. Bundle jobs, dependency jobs. Optimize job sizes and walltime requests.

CESM Readiness on Cori Phase 1

- The CESM/1.2.2 version ported by NCAR/NERSC to Cori Phase1.
- Source code distribution and scripts: /project/projectdirs/ccsm1/collections/cesm1_2_2/. Use similar scripts as before to create new cases by using the "-machine corip1" option.
- This is a functional port only, not tuned for optimal pe-layouts or performance. Based on popular demands 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.
- Also older CESM versions are not ported to Cori P1. You can get the 3 *.corip1 files from the cesm1_2_2 Machines directory along with the corip1 entry in "config_machines.xml" and copy them into your version.
- Users with Developers access can use cesm1_5_beta02 or newer. More cases there with pe-layout tuned.

CESM Features

- **Load balances between components**
- **Number of threads per component**
- **Limitation of parallelism in certain decompositions due to number of latitude layers and space resolution.**
- **Limited scaling due to global reductions in certain phases.**
- **Climate validation is essential. Sensitive to specific compilers, MPI libraries, math and I/O library versions, etc.**
- **Need long simulations.**
  - Jobs need to be broken up due to limited batch queue wall time.
  - Use dependency jobs.

CESM Survey, May 2014

- **In 2014, approximately 30 projects using CESM/CCSM**
- **Total hours allocated to these projects are about 190M (out of about 2 billion total NERSC hours). About 100M hours used so far.**
- **Biggest projects are at NCAR and LBL; others include PNNL, LLNL, LANL. 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**
- **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**

Top 12 Projects Using CESM/CCSM

Repo	Project Title	Inst.	Hours Used
m1040	Investigation of the Magnitudes and Probabilities of Abrupt Climate TransitionS (IMPACTS)	LBNL	10M
mp9	Climate Change Simulations with CESM: Moderate and High Resolution Studies	NCAR	18M
m1704	Multiscale Methods for Accurate, Efficient, and Scale-Aware Models of the Earth System	LBNL	7.5M
cascade	CASCADE: CAibrated and Systematic Characterization, Attribution and Detection of Extremes	LBNL	6.5M
m1199	Interactions of Aerosol, Clouds and Precipitation (ACP) in the Earth System	PNNL	6.7M
m1517	Attribution of extreme weather risk to anthropogenic emissions	LBNL	3.2M
m1204	Center at LBNL for Integrative Modeling of the Earth System (CLIMES)	LBNL	1.8M
m1795	Projections of Ice Sheet Evolution	LANL	1.4M
mp193	Program for Climate Model Diagnosis and Intercomparison	LLNL	1.4M
m328	Global cloud modeling	Colo. StateU	1.1M
m1374	Multiscale Modeling of Aerosol Indirect Effects on Decadal Timescales	PNNL	794K
m1576	Reducing Uncertainty of Climate Simulations Using the Super-Parameterization	COLA	789K

Number of Users and Jobs (apruns)

	CESM		CCSM		CAM		CLM		POP	
	#user	#job	#user	#job	#user	#job	#user	#job	#user	#job
Hopper	40	4590	26	5274	17	3384	1	38	1	140
Edison	26	2454	4	83	10	175	0	0	1	18434

Typical User CESM Jobs at NERSC

- Following slides show samples of how jobs are run
- aprun command launches the job
- aprun –j2 uses hyperthreading on Edison
- aprun –N#: number of MPI tasks per node
- aprun –d#: number of OpenMP threads per MPI
- The CESM task/thread layout for each model component is then shown; some timing profiles, too
- Parallelism typically driven by model component efficiencies, queue time, allocation size, and other considerations

Typical User Jobs (1/6)

- **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

- **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

Typical User Jobs (2/6)

- **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

    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 User Jobs (3/6)

- **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

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

CESM NESAP

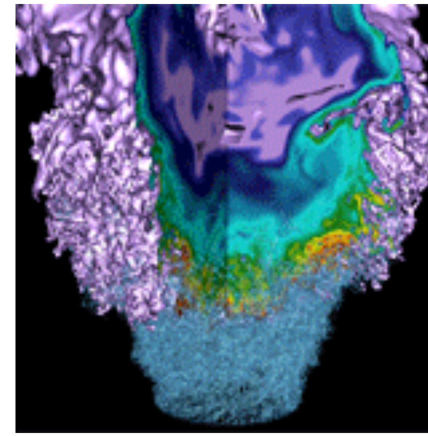

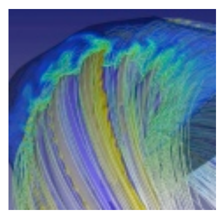

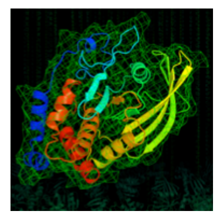

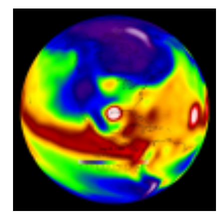

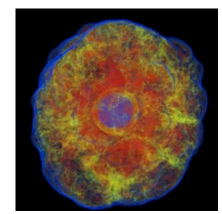

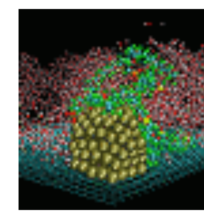

CESM MG2

MG2: CESM kernel for radiation transfer workload

- Typically takes 10% of CESM time
- Compute bound
- Very little vectorization: pipeline dependencies, heavy use of math intrinsics

Sample Code Dependencies and Vectorization Prototype from Intel Dungeon Session

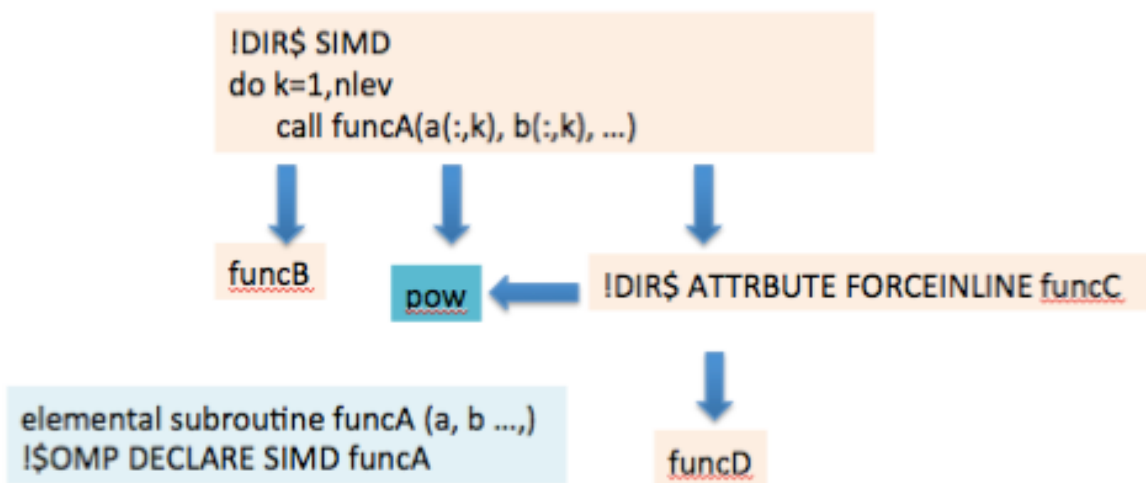

MG2: Sample Code Change

- Remove 'elemental' attribute and move the 'mgncol' loop inside routine

Before change:

```
elemental function
wv_sat_svp_to_qsat(es, p)
result(qs)

  real(r8), intent(in) :: es  !
SVP
  real(r8), intent(in) :: p
real(r8) :: qs

  ! If pressure is less than SVP,
set qs to maximum of 1.
  if ( (p - es) <= 0._r8 ) then
     qs = 1.0_r8
  else
     qs = epsilo*es / (p -
omeps*es)
  end if

end function wv_sat_svp_to_qsat
```

After change:

```
function wv_sat_svp_to_qsat(es, p,
mgncol) result(qs)
  integer,
intent(in) :: mgncol
  real(r8), dimension(mgncol),
intent(in) :: es  ! SVP
  real(r8), dimension(mgncol),
intent(in) :: p
  real(r8), dimension(mgncol) :: qs
  integer :: i
  do i=1,mgncol
  if ( (p(i) - es(i)) <= 0._r8 ) then
     qs(i) = 1.0_r8
  else
     qs(i) = epsilo*es(i) / (p(i) -
omeps*es(i))
  end if
  enddo
end function wv_sat_svp_to_qsat
```

CESM MG2: OMP SIMD ALIGNED

- **Using the "ALIGNED" attribute achieved 8% performance gain when the list is explicitly provided.**
- **However, the process is tedious and error-prone, and often times impossible in large real applications.**
  - !$OMP SIMD ALIGNED added in 48 loops in MG2 kernel *(by Christopher Kerr)*, many with list of 10+ variables

!$OMP SIMD ALIGNED	!$OMP SIMD	!dir$ VECTOR ALIGNED	-align array64byte	-openmp	Time per iteration (usec) on Edison
x			x	x	**444**
x				x	446
	x		x	x	**484**
	x			x	482
		x	x		452
		x			456
					473

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
- Use profile-guided optimization (PGO)

Version 3 (Intel compiler only)

- Use !$OMP SIMD ALIGNED to force vectorization

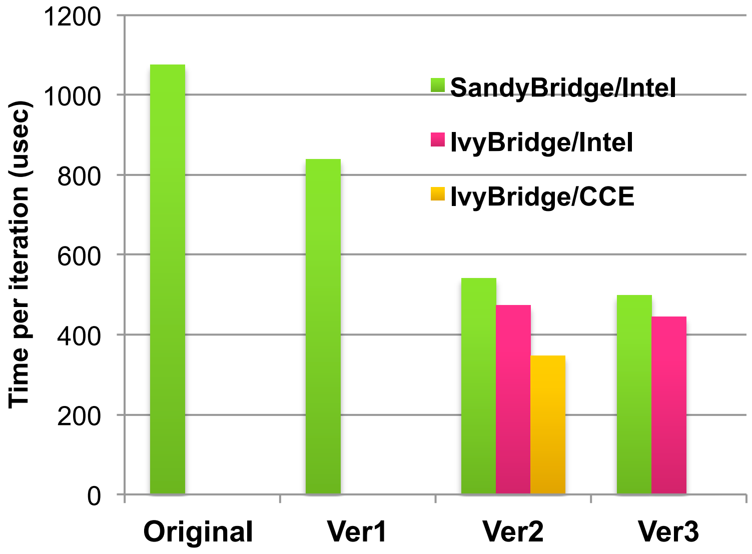

MG2 Optimization Summary

- Directives and flags are helpful. Not a replacement for code modifications.
- Break up loops and push loops into functions where vectorization can be dealt with directly. Try different compilers.
- Incremental improvements not necessary a BIG win from any one thing. Accumulative results matter.
- Performance and portability: use !$OMP SIMD is beneficial but very hard to use: need to provide the aligned list manually.
- Requested optional alignment declaration in Fortran language standard.

MPAS-O NESAP

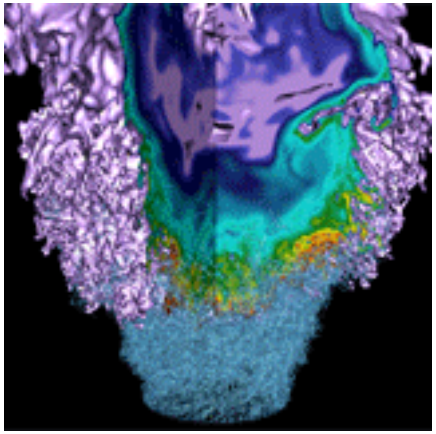

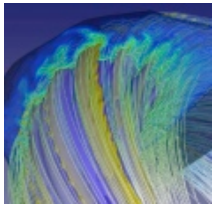

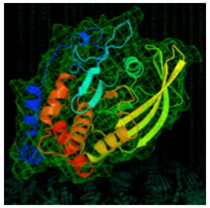

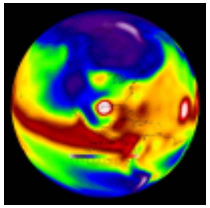

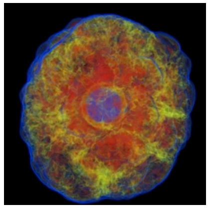

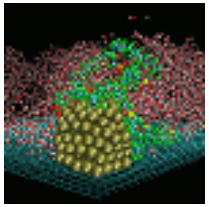

MPAS-O: Tasking

- **MPAS-O model uses unstructured meshes, data stored in memory is unstructured. Next contiguous element in an array may not be a neighbor of the previous element. Elements are decomposed into blocks.**
- **Threaded Block Loops: OpenMP Tasking**

```
block => domain % blocklist

do while (associated(block))

   call compute_block(block)

   block => block % next

end do
```

→

```
block => domain % blocklist
do while (associated(block))
   block_d = block
   !$omp task
   firstprivate(block_d)
   call compute_block(block_b)
   !$omp end task
   block => block % next
end do
!$omp task wait
```

Courtesy of Douglas Jacobsen et. al., NCAR Multi-Core 2015 Workshop

MPAS-O: Threaded Element Loops

- **Use Pre-computed decompositions (SPMD)**

```
eleStart = get_ele_start(iThread)
eleEnd = get_ele_end(iThread)

do iElement = eleStart, eleEnd
  … compute on elements ...
end do
```

- **Use OpenMP Directives (loop parallelism)**

```
!$omp do private(...)
do iElement = 1, nElements
  … compute on elements ...
end do
!$omp end do
```

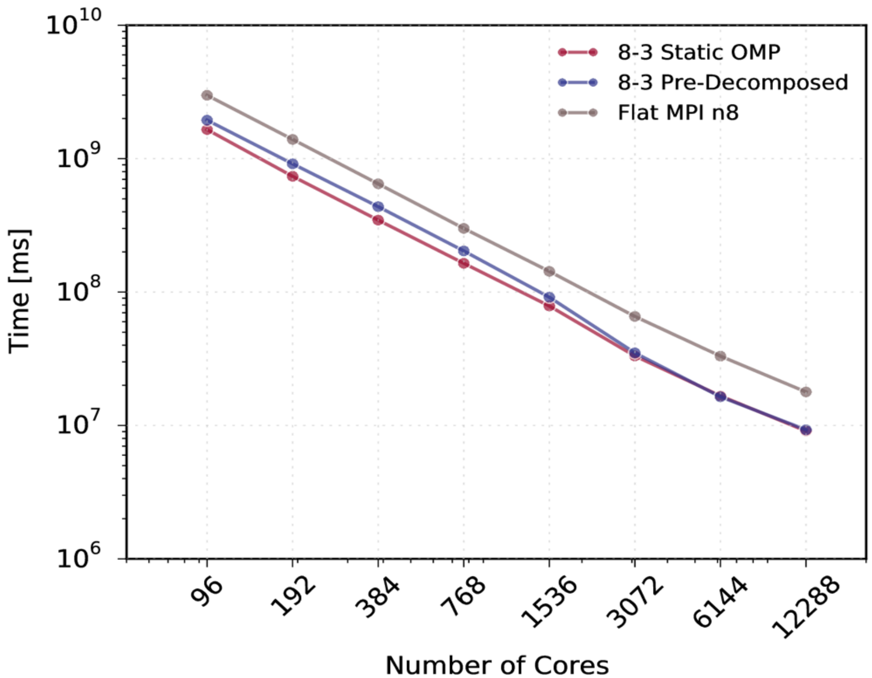

- **Loop parallelism better than SPMD**
- **Both better than pure MPI**

Courtesy of Douglas Jacobsen et. al., NCAR Multi-Core 2015 Workshop

MPAS-O: Compare Schedules and SIMD

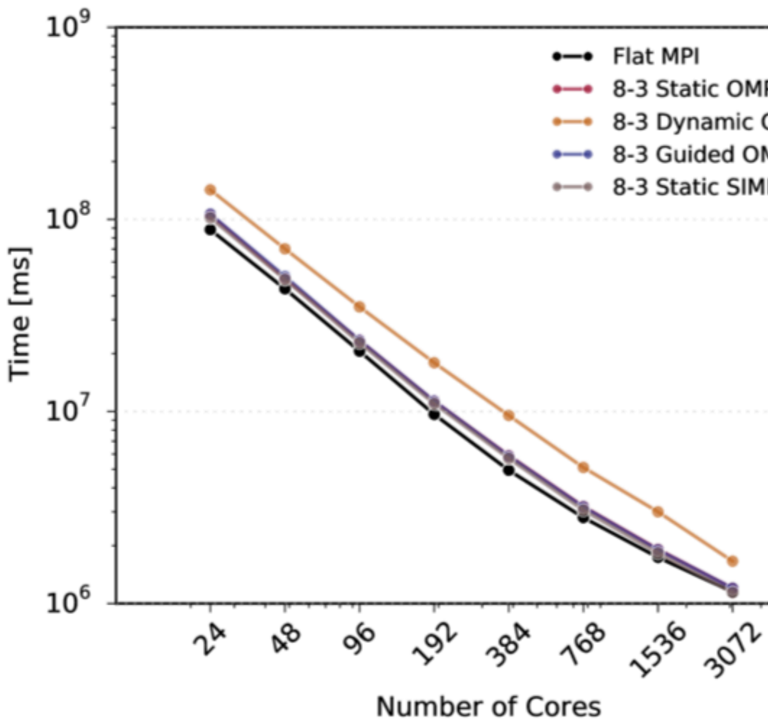

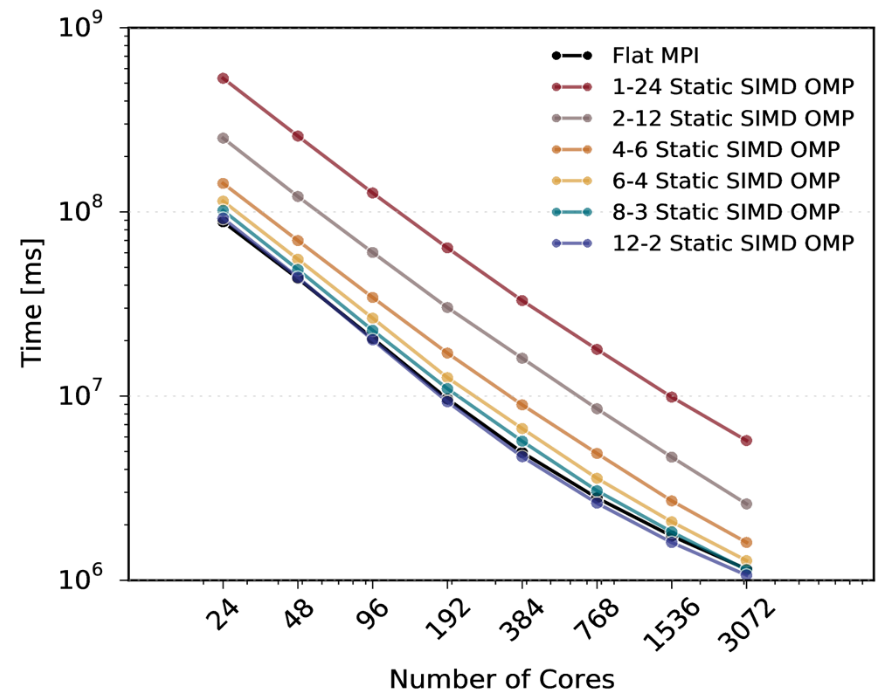

- **Good to explore different OpenMP schedules**
- **Good to experiment with different combinations of MPI tasks and OpenMP threads to find a sweet spot. "2-12 Static SIMD OMP" is the best in this case.**
- **SIMD directive helps a little, to vectorize loops compilers can not auto-vectorize.**

Courtesy of Douglas Jacobsen et. al., NCAR Multi-Core 2015 Workshop

MPAS-O: Strong Scaling with Full Code

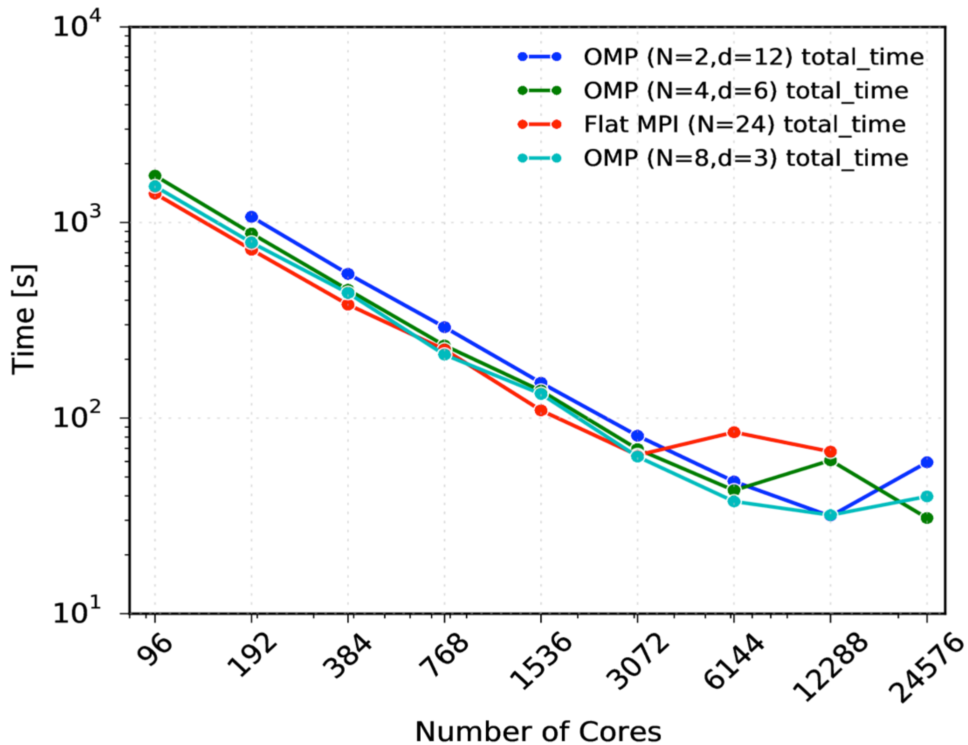

- **OpenMP helps scaling for larger core counts**
- **"OMP (N=8,d=3)" is the best in this case**

Courtesy of Douglas Jacobsen et. al., NCAR Multi-Core 2015 Workshop

ACME NESAP

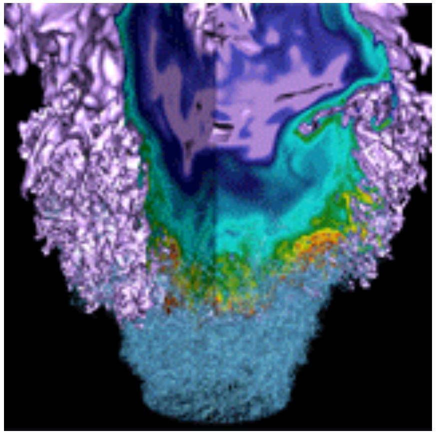

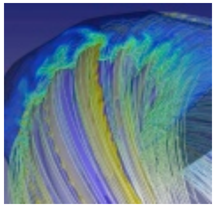

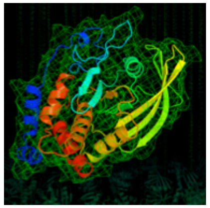

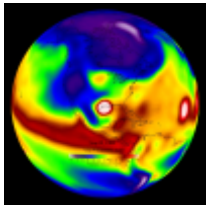

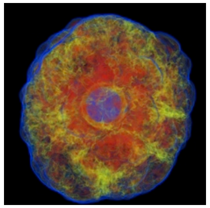

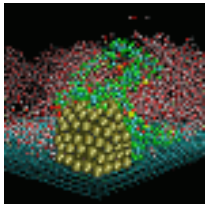

Current Status

- ACME v1 running at 5 SYPD. This model could be 2-3 times more expensive than ACME v0, and will need to run 3-4 times faster, requiring a 6-12x increase in computing power
- Good high level profile already collected by Patrick Worley using ACME's built in timers.
- Most expensive component is the atmosphere model, where we must get both tracer transport and dynamics running well.
- Atmosphere physics may or may not require much work. Still can be a lot of work due to the large size of code base.

Work Plan

- **Need to look at individual subcomponents, propose from small to large:**
  - Transport mini-app
  - HOMME (dynamics + transport)
  - ACME runnign in aqua planet model (CAM only, no other components)
  - ACME FC5 compsets (atmosphere + land + data ocean/sea ice)
  - ACME watercycle prototype

WRF

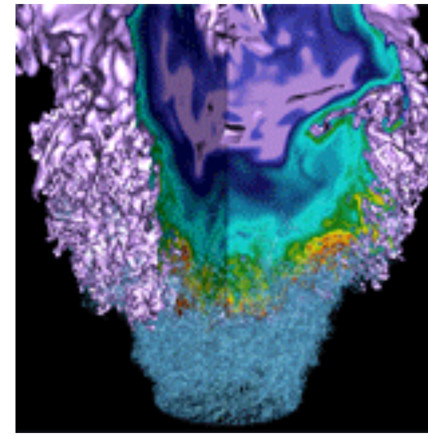

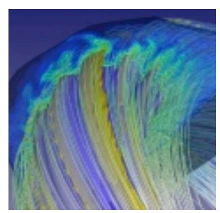

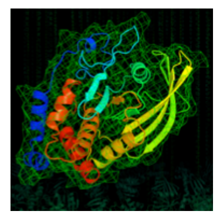

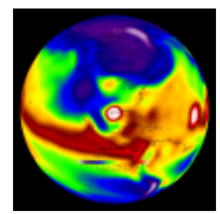

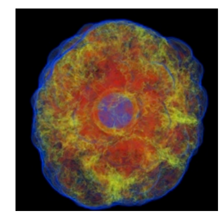

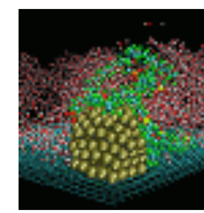

WRF Survey, Oct 2015

- **Which versions of WRF do you use?**
- **Which compilers and what versions of NETCDF do you use to build wrf.exe?**
- **Which WRF configure options do you choose? Please list the details of the options, such as:**
  - 54. Cray XC30 CLE/Linux x86_64, Xeon ifort compiler (dmpar)
  - 55. Cray XC30 CLE/Linux x86_64, Xeon ifort compiler (dm+sm)
- **Do you run with pure MPI or hybrid MPI/OpenMP?**
- **What model configurations and cases do you run? (some descriptions here would be good)**
- **How many nodes do you usually use? What is your typical aprun line?**
- **What is the typical wall time of your jobs?**
- **Do you have any specific questions or requests for us?**

WRF Support

- 14 responses. 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.

A Few Other Items

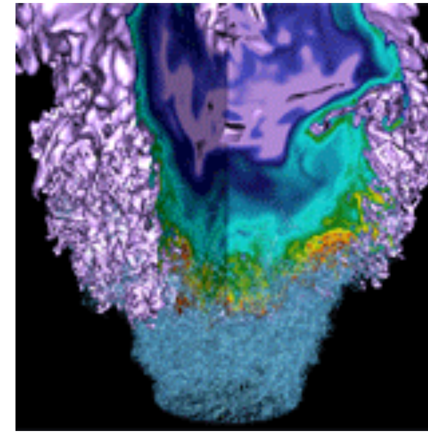

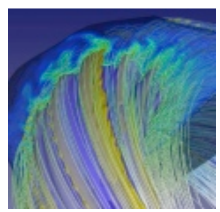

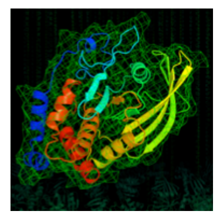

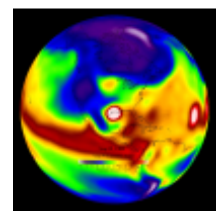

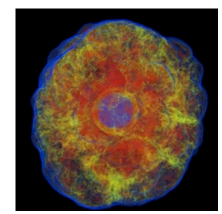

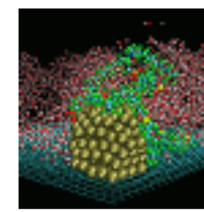

NERSC 40 YEARS at the FOREFRONT 1974-2014

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

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

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**
- **…**

Global Tropospheric Circulation Maps

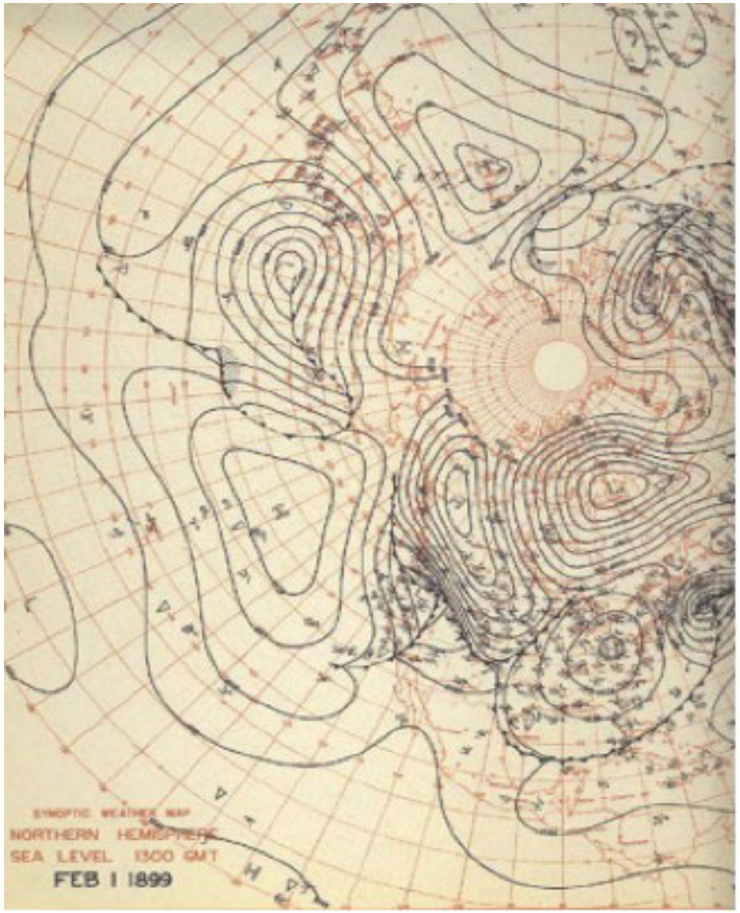

- Gil Compo, University of Colorado, Boulder – INCITE 2007 Project
- Provide options for data transfer between NERSC and NCAR.
- Tuning and debugging the application
- Adding multi-level parallelism to bundle several associated parallel jobs together to obtain high throughput and simplify data handling
- 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, reduce run time 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 new web form to submit your research to us at:**
  - https://www.nersc.gov/science/share-your-research/

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.

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.

Capps and Whiteman, 15th Annual WRF Users' Workshop

Work was performed at LBNL using NERSC

Extra Slides

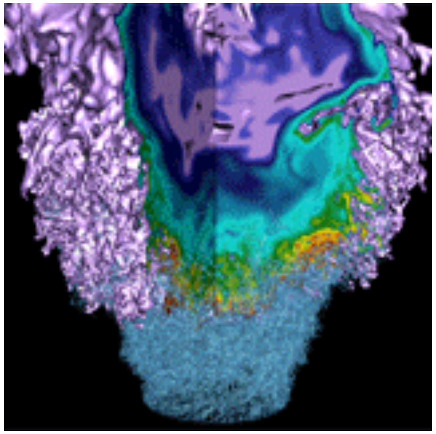

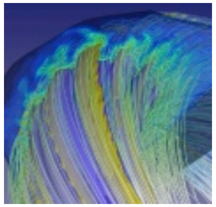

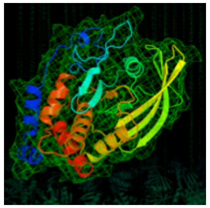

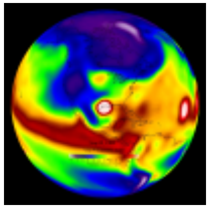

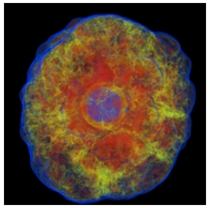

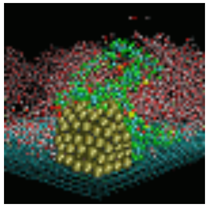

NERSC 40 YEARS at the FOREFRONT 1974-2014

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Climate Projects at NERSC (AY2015)-3

Repo	Project Title	PI	Organization	Codes
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
m2190	Next Generation Global Prediction System (NGGPS) Benchmarking	John Michalales	NOAA	NGGPS benchmark suite
m2320	Consequences of cloud super-parameterization for land-atmosphere coupling physics	Michael Pritchard	UC Irvine	SPCAM, UPCAM
m1374	Multiscale Modeling of Aerosol Indirect Effects on Decadal Timescales	Steven Ghan	PNNL	CESM, SPCAM
m2261	Greening the Grid USAID/India	Avi Purkayastha	LBNL	WRF
sobl	Southern Ocean Uptake in the MPAS-Ocean Model	Edward Patton	NCAR	NCAR-LES
m1439	Earth System Modeling	Yang Zhang	North Carolina State	WRF/Chem,WRF/CAM5, WRF/CMAQ, CESM
m2082	Multi-Scale Global to Cloud-Permitting Simulations for MJO Events	Jimy Dudhia	NCAR	MPAS-Atmosphere
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

Climate Projects at NERSC (AY2015)-4

Repo	Project Title	PI	Organization	Codes
m1576	Reducing Uncertainty of Climate Simulations Using the Super-Parameterization	Cristiana Stan	IGES - COLA	SP-CCSM
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
m1929	A Multiscale Reduced-Order Method for Integrated Earth System Modeling	George Pau	LBNL	CESM-PFLOTRAN-ROM, PROMEs
m1660	Regional Modeling of Land-Ocean-Atmosphere Interactions	Larry Beng	PNNL	WRF, WRF-Chem, CLM
m2136	Evaluation and improvement of Convective Parameterizations in ACME model	Wuyin Lin	BNL	ACME
m1481	Improvement of Representation of the Cloud Aerosol Interaction in Large-Scale Models	Alexander Khain	Hebrew Univ Jerusalem	WRF
m1196	Interactions of Clouds, Convection, and Climate	David Romps	LBNL	DAM
m2249	Integrating Climate Change into Air Quality Modeling	Ajith Kaduwela	UC Davis	CMAQ, WRF
m328	Global cloud modeling	David Randall	Colorado State	CESM, GCRM, SAM, Super-CAM
m2222	Expanding the computational frontier of multi-scale atmospheric simulation to advance understanding of low cloud / climate feedbacks	Michael Prtichard	UC Irvine	SPCESM

Climate Projects at NERSC (AY2015)-5

Repo	Project Title	PI	Organization	Codes
m1822	Using AMIE data to study cloud processes within the Madden-Julian Oscillation	Robert Houze	U. Washington	WRF
m1006	Dimention Reduction of the Cloud-Aerosol-Radiation (CAR) Ensemble Modeling System	Xin-Zhong Liang	U. Maryland	CWRF/CAR
m1178	Development of Frameworks for Robust Regional Climate Modeling	Ruby Leung	PNNL	WRF-CLM-ROMS, MPAS-A/MPAS-O, HOMME-POP-ROMS
m1637	Reducing the Uncertainty in Simulating Aerosol Chemistry over Multiple Spatial Scales	Jerome Fast	PNNL	WRF
m1355	Analysis of Global Coupled 0.1-deg POP/CICE in the CESM Framework	Julie Mcclean	Scripps Inst	CESM
mp79	Parallel modeling of climate, chemistry, and physics affecting the global atmosphere	Donald Wuebbles	U. Illinois U-C	CESM, CAM-Chem, WACCM
m411	Interaction of Atmospheric Chemistry and Aerosols with Climate	Philip Cameron-Smith	LLNL	CESM, IMPACT
m184	High-Spatial Resolution Models of Atmospheric Chemistry, Aerosols for Regional Scale Climate Simulations	Rao Kotamarthi	ANL	CESM, CLM, CMAQ, GEOS-Chem
m1626	High-Spatial Resolution Models of Atmospheric Chemistry, Aerosols for Regional Scale Climate Simulations	William Gustafson	PNNL	MPAS, WRF

Climate Projects at NERSC (AY2015)-6

Repo	Project Title	PI	Organization	Codes
m1590	Regional Aerosol Emission and Its Direct and Cloud-Related Forcing Estimates	Tami Bond	U. Illinois U-C	CESM
m1865	Predictability of the carbon-climate system on seasonal to decadal time scales	Inez Fung	UC Berkeley	CCSM
m2076	Role of Forest Ecosystems in Mitigating Climate Change Impacts through Land-Surface Energy and Water Processes	Guangshan Chen	U. Wisc. Madison	RegCM4-CNDV
m843	Integrating observations and simulations of maritime and continental boundary-layer clouds and mid- latitude deep convection to advance understanding of drizzle, rain, and ice formation	Ann Fridlind	NASA GISS	DHARMA
refcst	A multi-decadal reforecast data set to improve weather forecasts for renewable energy applications	Thomas Hamill	NOAA	GFS
m1269	A multi-decadal reforecast data set to improve weather forecasts for renewable energy applications	Nigel Quinn	LBNL	HydroGeosphere
m2420	Next Generation Ecosystem Experiments -- Tropics	Charles Koven	LBNL	CLM, CESM, ACME
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

Climate Projects at NERSC (AY2015)-7

Repo	Project Title	PI	Organization	Codes
m2319	Using carbonyl sulfide to quantify regional terrestrial biological carbon fluxes	John Campbell	UC Merced	STEM
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
mp231	Three-dimensional Global Atmospheric Aerosol and Chemistry Modeling	Joyce Penner	U. Michigan	IMPACT, GCE cloud model
m640	Influences of the Boundary Layer Flow on Vegetation-Air Exchanges of Energy, Water and Greenhouse Gases	Xuhui Lee	NCAR	NCAR-LES
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
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
m1861	Towards parameterization of root-rock hydrologic interactions in the Earth System Model	Inez Fung	UC Berkeley	CLM4/ CLM4RR
m2164	Studying the Madden-Julian Oscillation over a wide range of climates	Da Yang	UC Berkely	SPCAM

Climate Projects at NERSC (AY2015)-8

Repo	Project Title	PI	Organization	Codes
m2258	Scale-aware, Improved Hydrological and Biogeochemical Simulations of the Amazon Under a Changing Climate	Chaopeng Shen	Penn State	CLM+PAWS
m2492	Paleo-megadroughts and Abrupt Climate Changes in the Speleothem Records	Inez Fung	UC Berkeley	CAM
m2272	PARAFORCE	Carlo Lacagnina	SRON - Utrecht	Storage only
m1639	Assessing Climate Change Effects on Wind Energy	Cameron Whiteman	Vertum Partners	WRF
m1427	Collaborative Research: Wildfires and regional climate variability - Mechanisms, modeling, and prediction	Xiaohong Liu	Univ. Wyoming	CAM5
m543	River Transport and Hydrology in CCSM	Marcia Branstetter	ORNL	CESM
m2095	Influence of the Boundary Layer Flow on Vegetation-Air Exchange of Energy, Water and Carbon Dioxide	Jianping Huang	Yale Univ	NCAR-LES
m2232	Accelerated Weather Prediction	Yik-Kiong Hue	Intelligent Automation	WRF
m2478	PROJECT 51985	Eugene Yan	ANL	WRF
m2324	Collaborative Research: Understanding Long-Term Variability in the North Atlantic Subtropical High Driven by Evolving Tropical and Subtropical Large-Scale Heating	Wenhong Li	Duke Univ	---

Typical User Jobs (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

Typical User Jobs (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

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 User Jobs (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

Impact on Climate Simulations

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

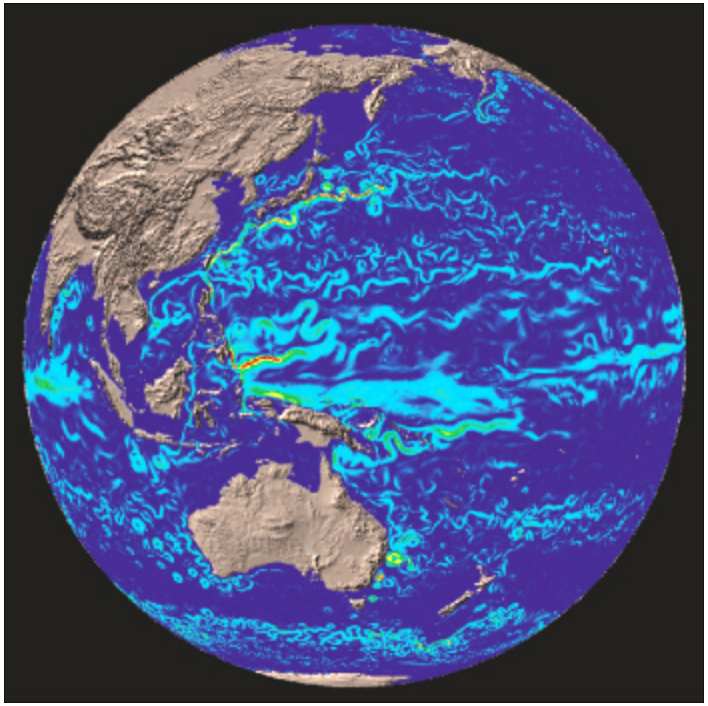

- **High-Resolution Global Coupled Ocean/ Sea Ice Modeling using POP**
- **NERSC provides**
  - **Environment for highly parallel runs**
  - **Service and stability**
  - **Special queue support**

NERSC 40 YEARS at the FOREFRONT

Thank you.