Edison Phase I Early User Demographics and Science Results

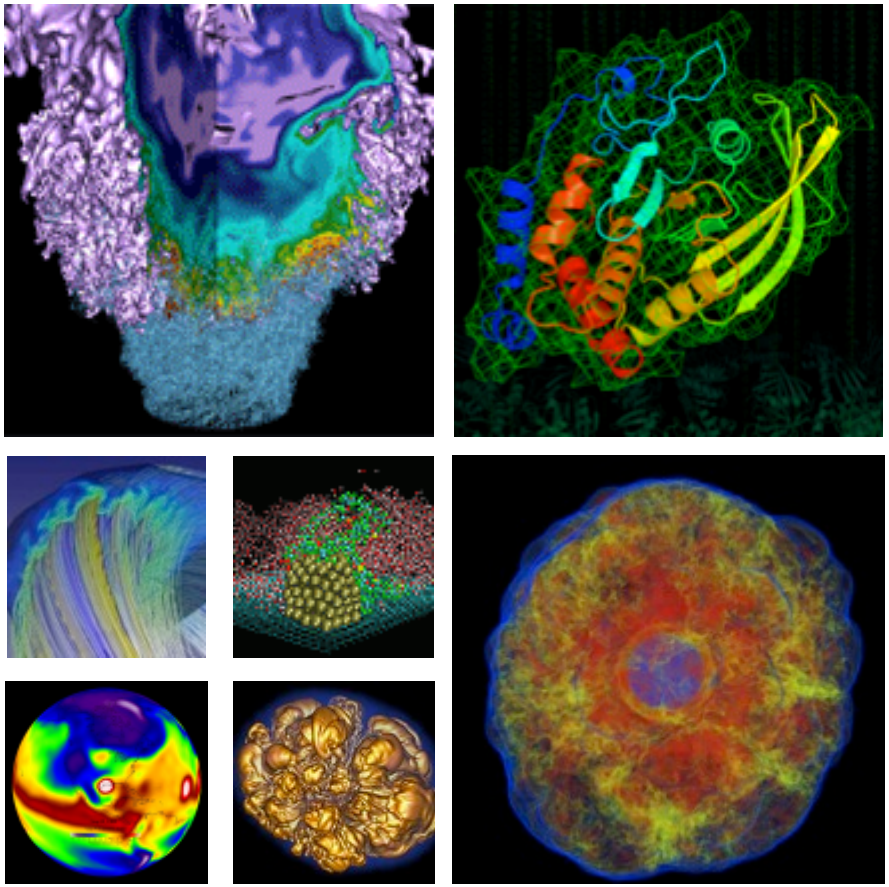

Zhengji Zhao, and Richard Gerber

NERSC User Services

July 22, 2013

Edison Phase I – Cray XC30

- System was delivered to NERSC November 29, 2013.
- System was opened to early users on February 5, 2013.
- All users enabled on March 3.
- Compute portion was accepted on March 28.
- System was decommissioned on June 24, 2013 for Phase II installation.

Edison Phase 1 Configuration

Compute node	Two 8 core Intel Sandy Brid
System size	664 nodes, 10624 cores; 21248 with Hyper-Threading
Cores per node	16; 32 with Hyper-Threading New
Clock speed	2.6GHz
Memory	64GB DDR3 1600 MHz 4GB/core
NUMA Domain	2
Interconnect	Aries with dragonfly network topology New
Lustre file system	1.6 PB with 36GB/s bandwidth
IO Servers (OSSs)	18
OSTs	72
Storage file system	Cray Sonexion 1600 New
Login nodes	quad-core, quad-socket (16 total cores) 2.0 GHz Intel "Sandy Bridge" processors with 512 GB memory.
Batch system	External Batch Server New

Intel programing environment was the default

Baseline performance

NERSC-6 Application Benchmarks

Application	CAM	GAMESS	GTC	IMPACT-T	MAESTRO	MILC	PARATEC
Concurrency	240	1024	2048	1024	2048	8192	1024
Streams/Core	2	2	2	2	1	1	1
Edison P1	260.04	1083.20	736.99	477.37	954.15	418.37	194.00
Speedup[1)]	**1.3**	**1.3**	**1.8**	**1.3**	**2.0**	**2.2**	**1.8**
Hopper[2)]	348	1389	1338	618	1901	921	353

[1)] Speedup=Time(Hopper)/Time(Edison)

[2)] Hopper, Cray XE6, has 2 twelve-core AMD 'MagnyCours' 2.1-GHz processors per node running at a frequency of 2.1GHz, interconnected with Cray's Gemini Network.

Average system utilization was 84% (2/5-6/24)

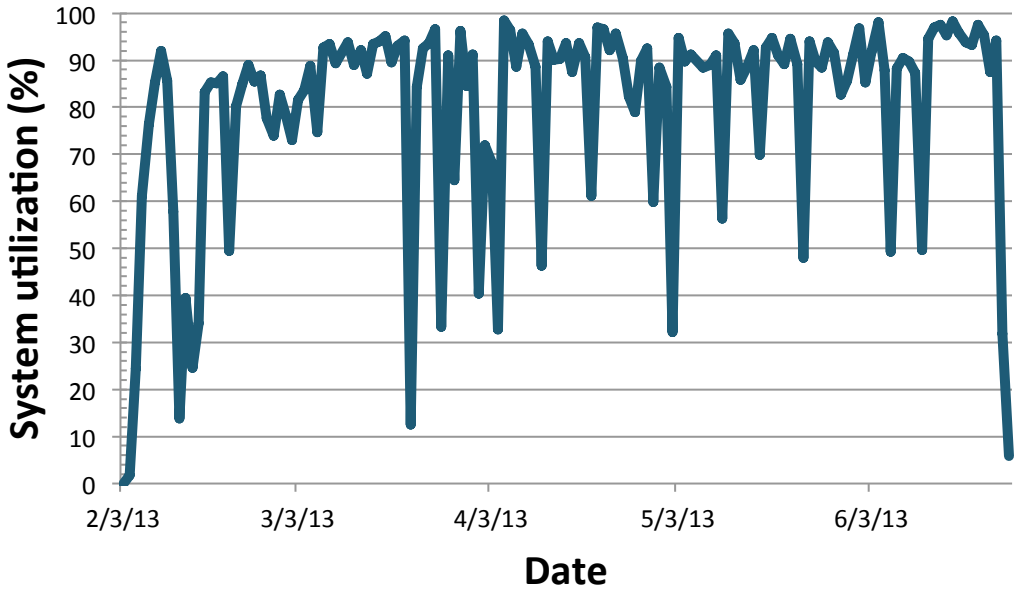

30 M raw machine hours used from 2/5 through 6/24 - 60M Hopper equivalent hours

Used by 273 projects; more than 455 users compiled; 577 users ran jobs

DARPA Mission Partners running

Edison outages during 2/5-6/24

Duration (h)	Description	SWO
3.3	Logins available, batch jobs not running.	No
77.1	Unavailable	Yes
172.5	Scheduled maintenance	Yes

There were 18 outages in total (2/4-6/24) including one center wide outage.

Effort needed for a successful installation, implementation and deployment

46 Cray cases and 84 bugs were filed for Edison

676 tickets created for Edison by 170 users and staff

User trainings and special talks for Edison NERSC User Group meetings in addition to providing web docs

NERSC users were able to quickly start using Edison to satisfy some of the large demand for computing cycles.

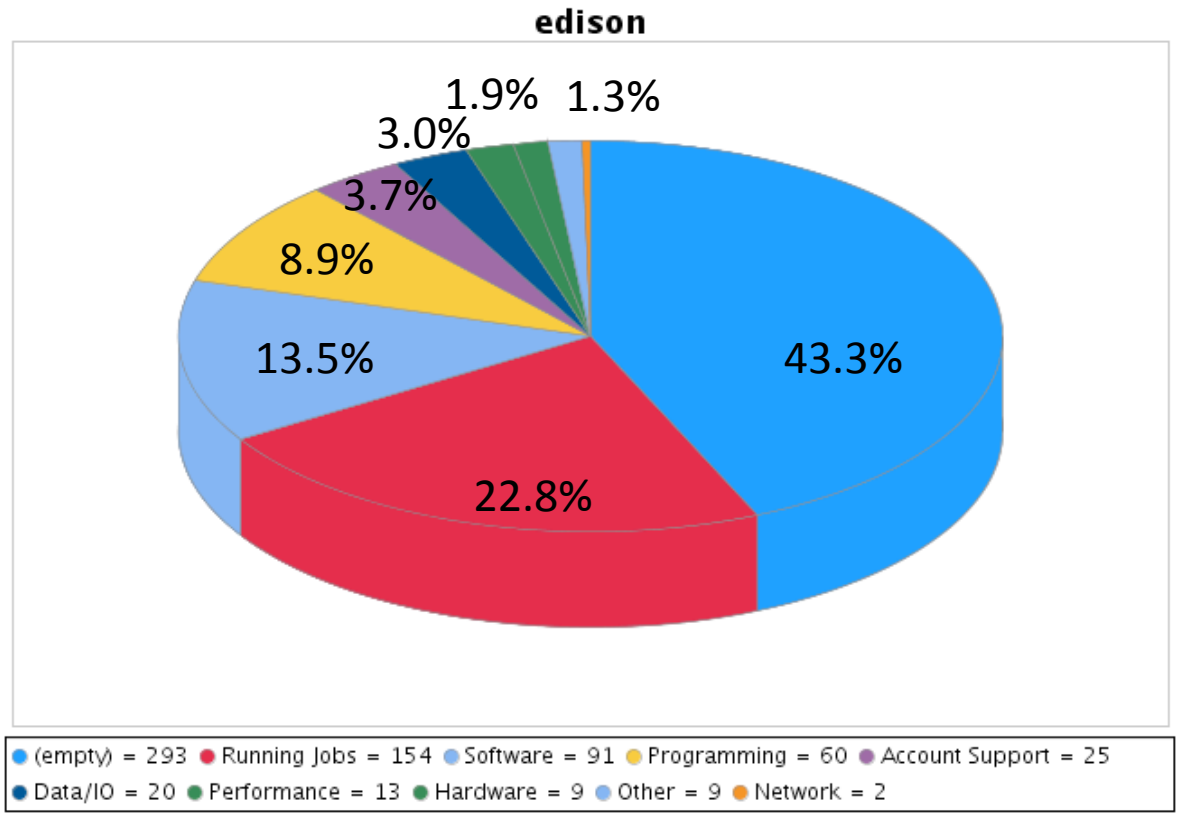

U.S. DEPARTMENT OF ENERGY | Office of Science

All DOE Science Offices are represented during Edison Phase I time (2/5-6/24)

Breakdown of Edison hours by DOE science offices (2/5-6/24)

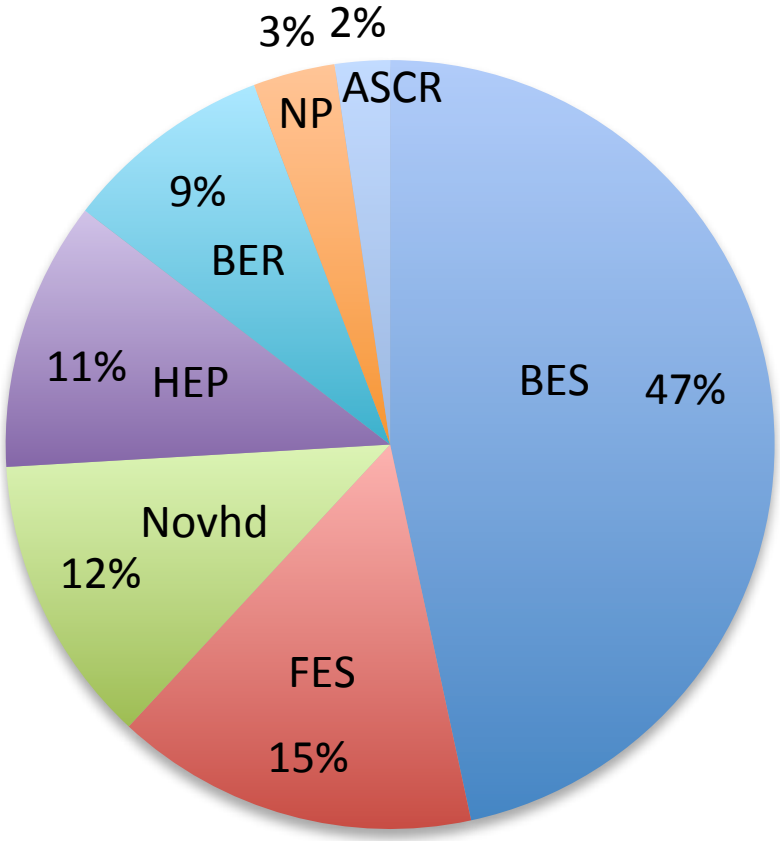

Fair share weight: (3/15 – 6/24)

- DARPA: 25%
- NERSC: 75%
  - ASCR: 4%
  - BER: 13%
  - BES: 22%
  - FES: 13%
  - HEP: 10%
  - NP: 8%
  - Overhead (staff): 5%

We deployed the Fairshare component in Moab to balance the usage between DOE offices during 3/15-6/24. Still need further experiment with fairshare.

All science areas are represented during Edison Phase I time with highest usage from materials science

Breakdown of Edison hours by science area (2/5-6/24)

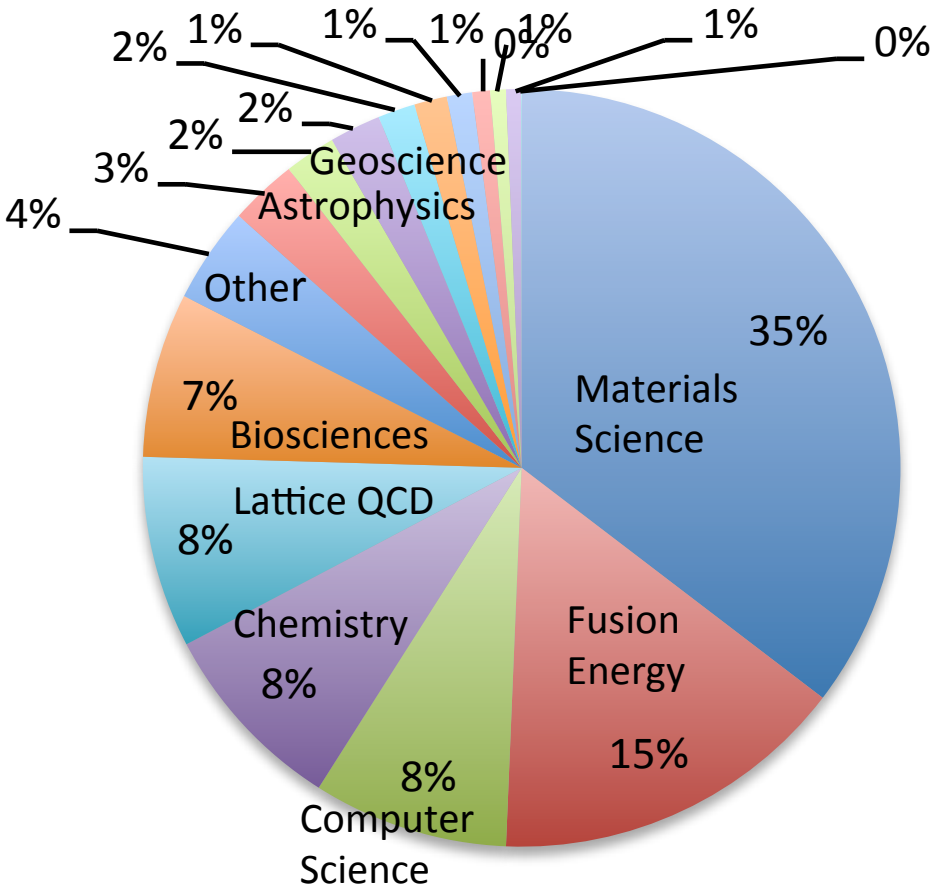

Top applications on Edison

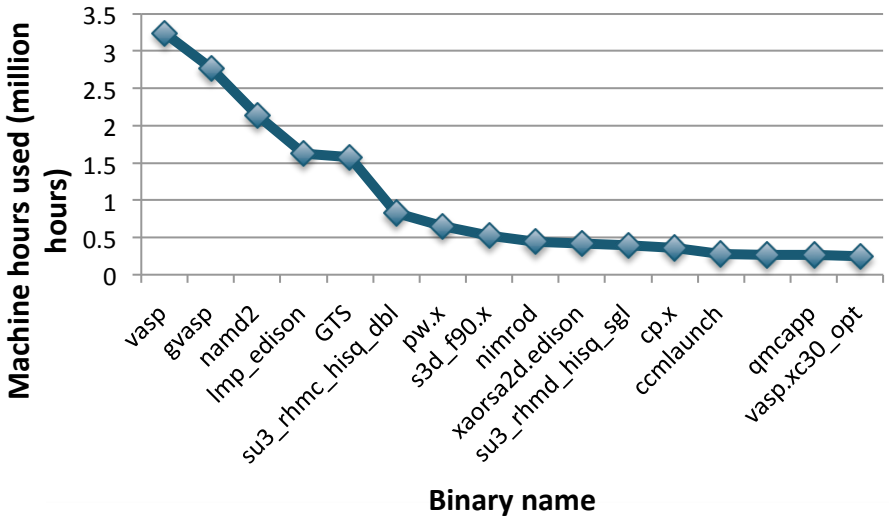

More than 3500 binaries ran on the system

Largest Project Usage of Edison Phase 1

PI/ Project	Institution	Office	Project	Hours Used
Yai-Yim Ching mp250	Univ. Missouri-Kansas City	BES	Electronic Structures and Properties of Complex Ceramic Crystals and Novel Materials	4.4M
darpa			DARPA HPCS Mission Partner Access	2.4M
Doug Toussaint mp13	Univ. of Arizona	HEP	Quantum Chromodynamics with four flavors of dynamical quarks	1.7 M
Wei-Li Lee m19	PPPL	FES	Turbulent Transport and Multiscale Gyrokinetic Simulation	1.6 M
Ivaylo Ivanov m1254	Georgia State Univ.	BER	Integrative Modeling of Protein/DNA Complexes at the Replication Fork	1.2M
Gary Grest m1334	Sandia Labs	BES	Controlling Nanoparticle Assembly to Engineer New Materials	1.1M

Applications were run with all concurrencies during Edison Phase I time (2/5-6/24)

Breakdown of Edison hours by the number of cores used (2/5-6/24)

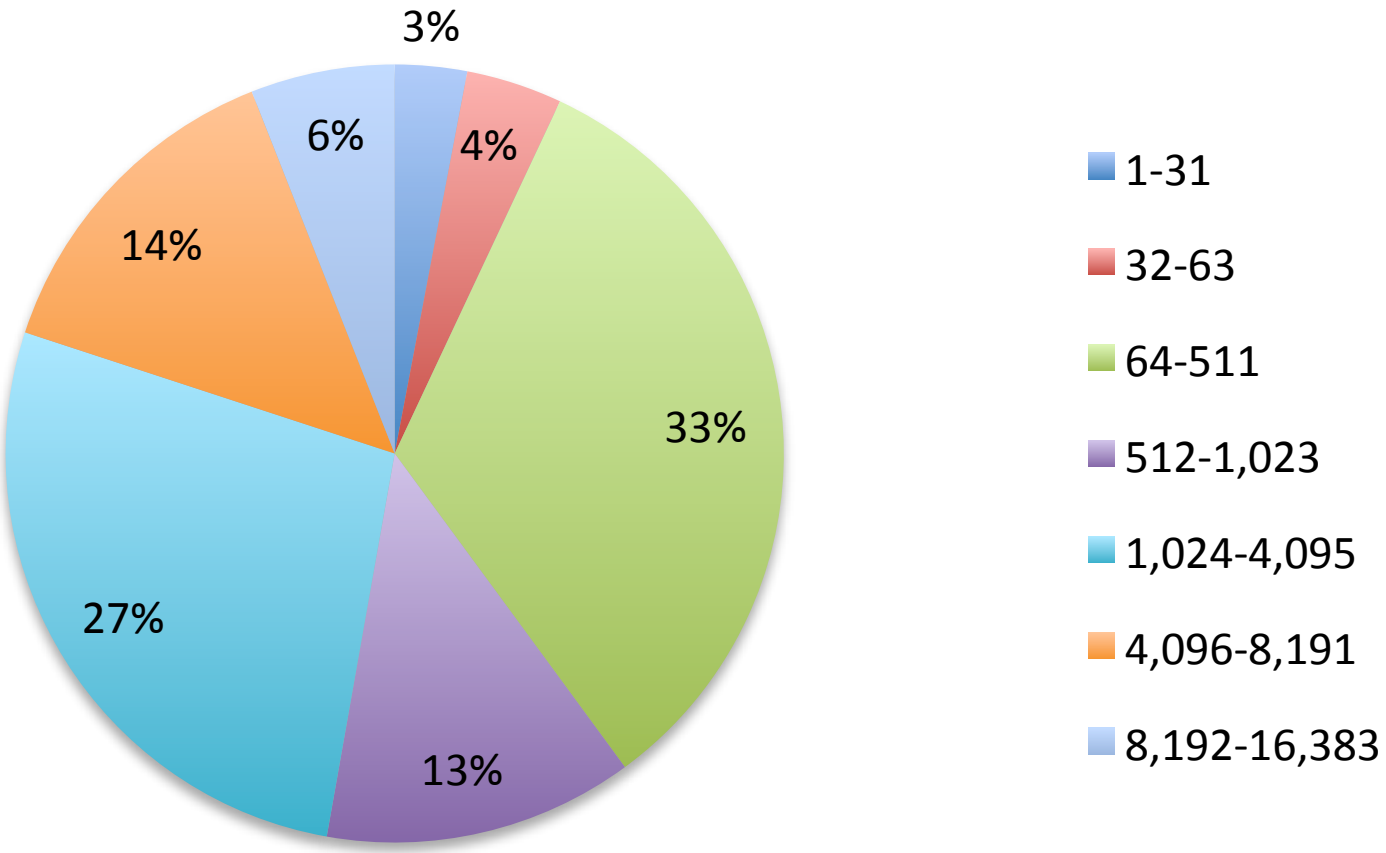

Edison Phase 1 Performance users reported

User Erwin Jaeger from ORNL:

"AORSA uses ScaLAPACK to solve for the RF wave fields, plasma heating, and current drive in a tokamak plasma.
With 4096 processors, a 400x400 grid takes 5.0 hours on Hopper and just 2.5 hours on Edison. (This is the largest grid that will fit in the Hopper memory).
A 512x512 grid (and larger) should easily fit in the Edison memory, but the present turnaround is very long (weeks). "

User Wai-Yim Ching from UMKC (used VASP code-materials science code):

"Edison is 3 times faster than NERSC's Hopper supercomputer for this project, allowing simulations of greater complexity and duration."

User Jagannathan T.K. from Columbia Univ.:

LAMMPS runs 1.5-2 times faster compared to Hopper –summarized from his slides

Performance comments from users

Edison's per core performance is twice as fast as Hopper's. –user dks

VASP code: It seems at least from this test that Edison is almost 3 times faster than the hopper.– user aryal

QE: I've used 128 cores on Carver and 256 cores on Edison, but the calculation on Edison is almost 5x faster. – user mmentink

VASP: The Edison system has superior performance. Considering that the CPU frequency is similar, the speedup from Carver and Hopper is significant. The network connection on Edison is also better. It doesn't show a bending-down behavior up to 32 nodes. – user yiyang

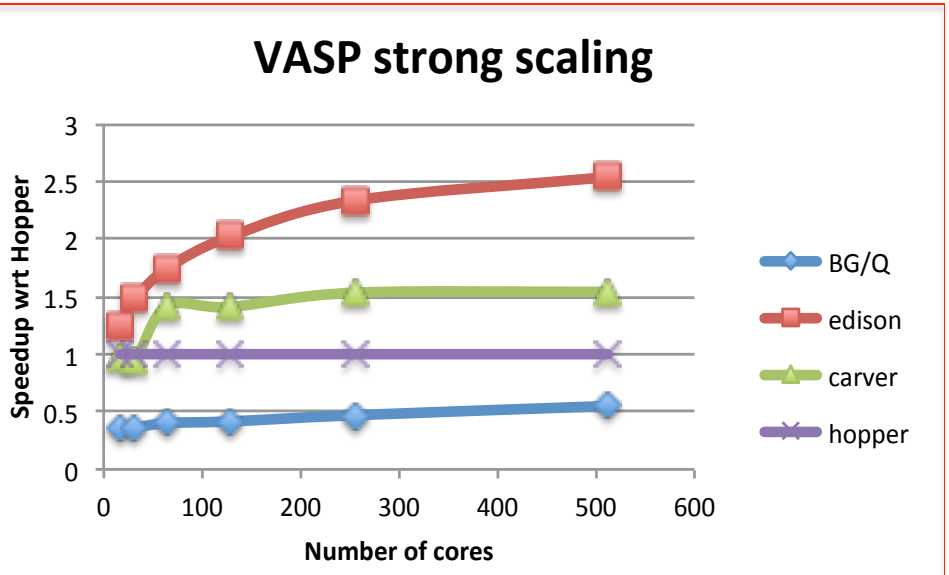

Edison Phase 1 Performance users reported- Co-Array performance (CAF)

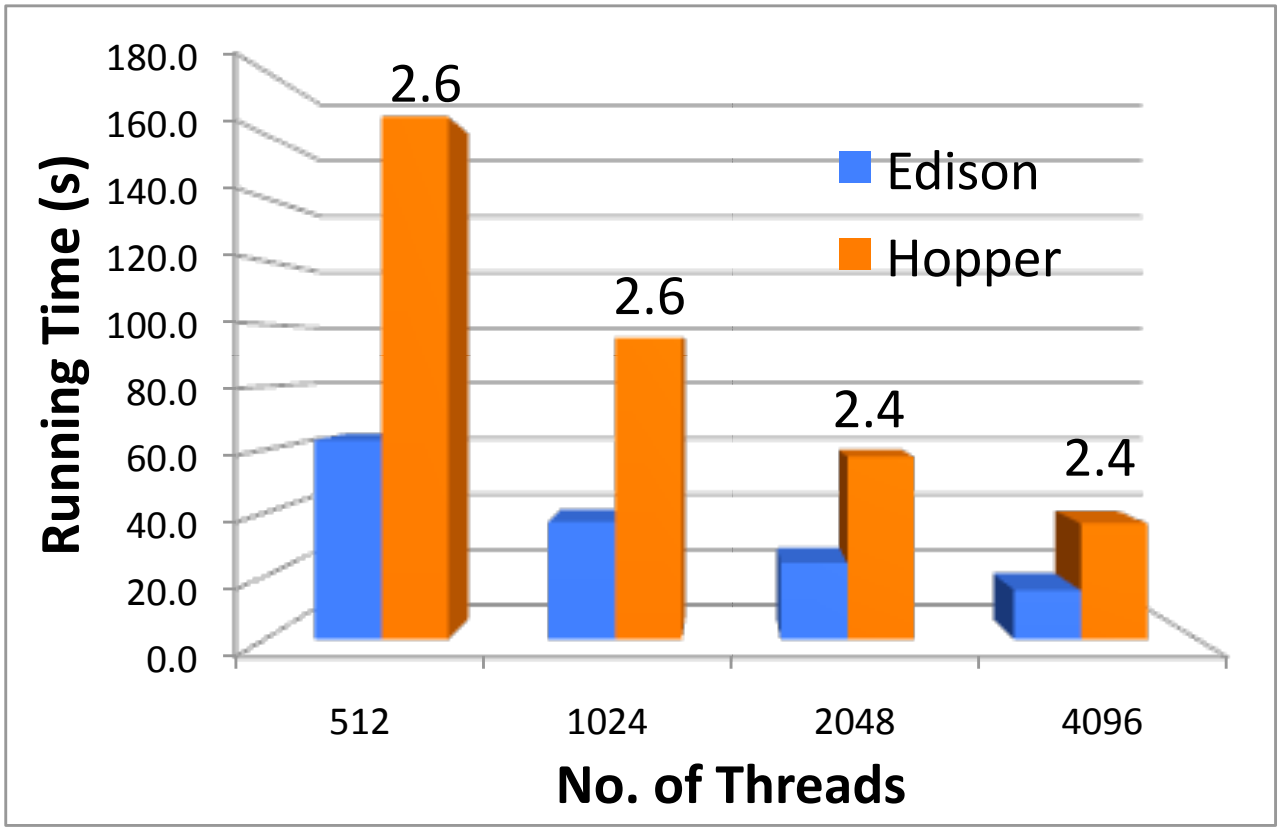

Impact-T

- Simulate the dynamics of a beam of charged particles within a particle Accelerator
- 3D Computational Domain, Partitioned in Y and Z direction
- Dominant communication pattern: Alltoallv within subgroups

Slide was provided by Hongzhang Shan at CRD/LBNL

Long-Range Interactions of the Cancer Drug Doxorubicin with Biomolecules

PI: Wai-Yim Ching
Univ. Missouri-Kansas City (BES)

Doxorubicin is a chemotherapy drug used to treat leukemia and other cancers.

Understanding interactions with its environment can help improve efficacy and reduce toxicity.

Researchers at the University of Missouri-Kansas City have used ***4.4 M Edison Phase I hours*** on this and related research.

Without this additional time, Doxorubicin simulations would not have been conducted.

Edison is **3 *times faster*** than NERSC's Hopper supercomputer for this project, allowing simulations of greater complexity and duration.

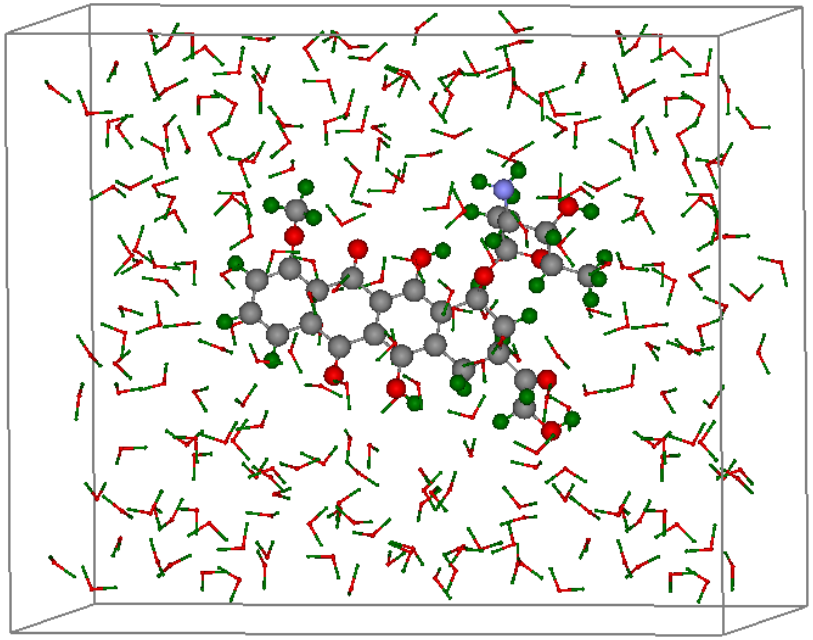

Simulation of doxorubicin in water to investigate the effect of changes in partial charge distribution of biomolecules on long range interactions.

W.Y. Ching, R. Sakidja, S. Aryal, C. Dharmawardhana, N. Li
University of Missouri-Kansas City

Edison Phase I Early User Science

U.S. DEPARTMENT OF ENERGY | Office of

Slide was provided by Wai-Yim Ching

QCD simulations on Edison

PI: Doug Toussaint
Univ. Arizona
(HEP-Lattice Gauge Theory)

Goals: understand strong interactions, find fundamental parameters, help experiments like dark matter searches.

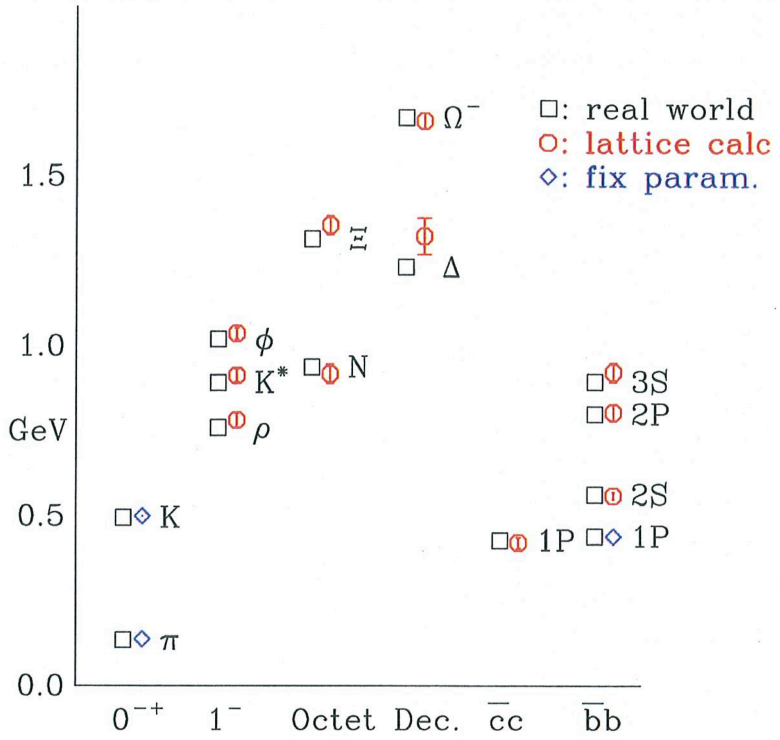

Calculated spectrum results compared to real world

Performance of Edison vs Hopper

cores	core-hours
Edison	
1024	2651
1536	2643
2048	2750
Hopper	
6144	22000
12288	23000
18432	20000

Core-hours for one molecular dynamics trajectory (smaller is better)

2 processes per core on Edison (32 per node), 1 process per core on Hopper (24 per node)

User Doug Toussaint used 1.7 million raw machine hours on Edison with MILC code.

U.S. DEPARTMENT OF ENERGY | Office of

Slide was provided by Doug Toussaint at Univ. Aizona

Turbulent Transport and Multiscale Gyrokinetic Simulation

PI:Wei-li Lee
PPPL (FES)

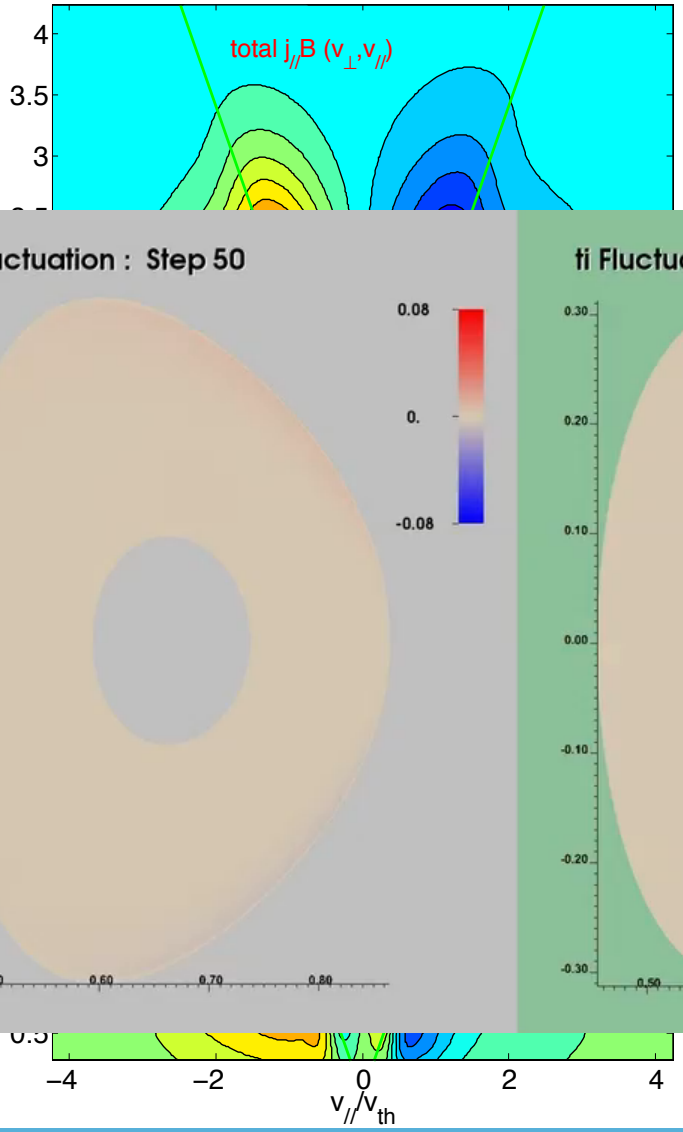

Can turbulence drive plasma current or change bootstrap current?

Slide contents were provided by Weixing Wang at PPPL

Controlling Nanoparticle Assembly to Engineer New Materials

PI: Gary S. Grest
Sandia Labs (BES)

Polymer brushes are surface-coating macromolecules with wide-ranging applications in medicine, biotechnology and nanotechnology.

Researchers from Sandia National Labs used 850 K hours on Edison Phase I in two months to extend their understanding of the interaction between polymer brushes and nanoparticles.

The goal is to control the assembly of nanoparticles into arrays that serve as building blocks for devices like solar cells, LEDs, and memory storage.

Because the code, called LAMMPS, runs ***2.6 times faster*** on Edison than on NERSC's previous-generation supercomputer, Edison has supplied an additional ***1.2 M raw machine hours*** to this ASCR Leadership Computing Challenge project.

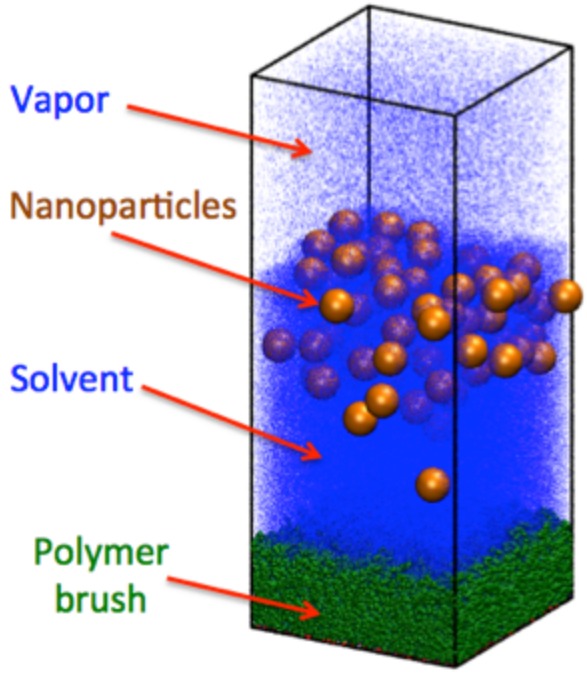

Shengfeng Cheng, Mark J. Stevens, and Gary S. Grest
Sandia National Labs

Edison Phase I Early User Science

U.S. DEPARTMENT OF ENERGY | Office of

Slide was provided by Gary S. Grest

Science at scale for CO2 sequestration

PI: David Trebotich
CRD/LBNL (ASCR)

EFRC-NCGC validation experiment

- Reactive transport in capillary tube packed with calcite
  - 7 mm long x 500 μm diameter
  - 0.899 μm image resolution
  - 9.3 seconds residence time (steady-state)
- Chombo-Crunch simulation
  - 1.6B simulation grid points
  - 1.19 μm resolution
  - 700GB plot files, 600GB checkpoint files
  - 49,152 cores on Hopper Cray XE6
  - 100 time steps per hour
  - 1 time step = 0.000182 seconds real time

➔ Current rate law in reactive transport model overshoots steady-state effluent concentration

➔ Run coarser simulation (2μm) with new rate law, cf. Edison

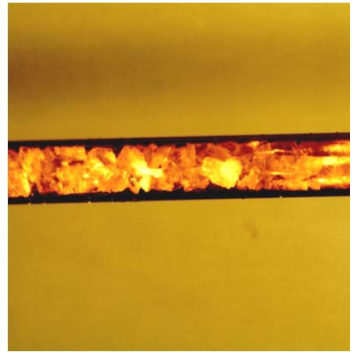

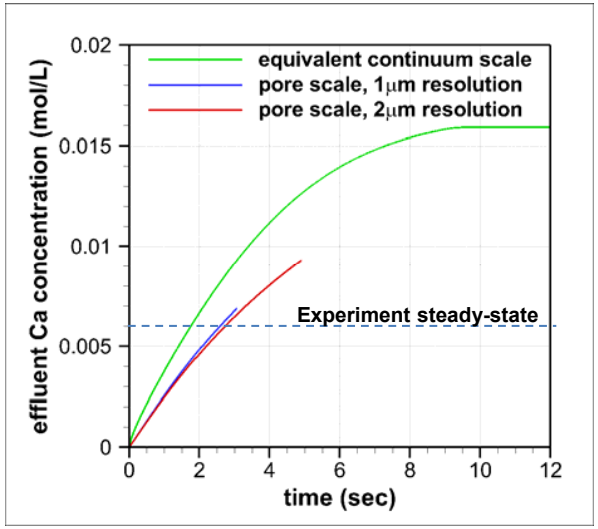

218k raw machine hours used

U.S. DEPARTMENT OF ENERGY | Office of

Slide was provided by David Trebotich

Mock Lyman-α Forest Datasets for the BOSS survey

PI: Peter Nugent
CRD/NERSC/LBNL
(HEP - Astrophysics)

The Baryon Oscillation Spectroscopic Survey (BOSS) is obtaining spectra of 160.000 distant quasars in order to map the distribution of matter in the Universe at an earlier epoch than that studied with regular galaxy surveys.

The goal is to measure the scale of Baryon Acoustic Oscillations (BAO) in the clustering of intervening gas (the "Lyman-α Forest"), and use it to study the geometrical properties of the Universe and its accelerated expansion.

In order to test our analysis methods, and to study the effect of potential systematics in our data, it is crucial to have a large number of synthetic datasets with well known properties.

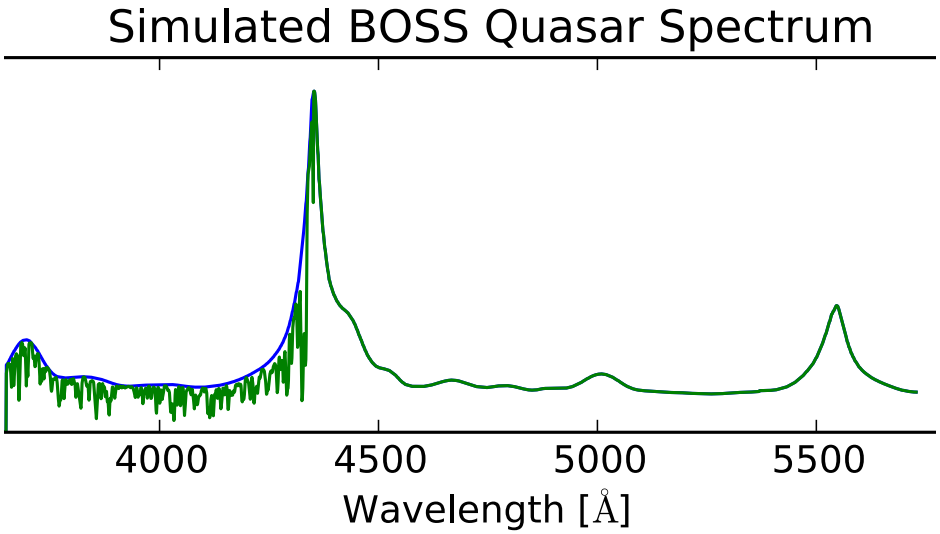

We used over 86k CPU hours (180 single node jobs with OpenMP threads) on Edison in order to generate 100 realizations of the Data Release 11. These data sets are already distributed within the BOSS collaboration and will be publicly released with the actual data.

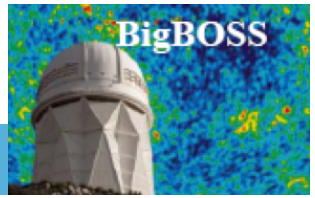

U.S. DEPARTMENT OF ENERGY | Office of

Slide was provided by A. Font-Ribera, S. Bailey

Conclusion

- **Edison Phase I was easy for users**
  - NERSC users were able to quickly start using Edison to satisfy some of the large demand for computing cycles.
- **Enabled effective application performance at scale, single node (high-throughput computing), and everything in between.**
- **Provided needed resources for DOE scientists**
- **NERSC/Cray are working to provide Phase II system as early as possible to meet our users' need.**
  - Queue turnaround will be largely improved

Acknowledgement

- NERSC users, Doug Toussaint, Weixing Wang,Gary Grest,David Trebotich, A. Font-Ribera, and S. Bailey, who provided the science stories and slides
- Katie Antypas, and other NERSC staff for valuable discussions and help.
- Randy Palmer for providing the number of Edison bugs.

National Energy Research Scientific Computing Center