Cori: Enabling World-Changing Science

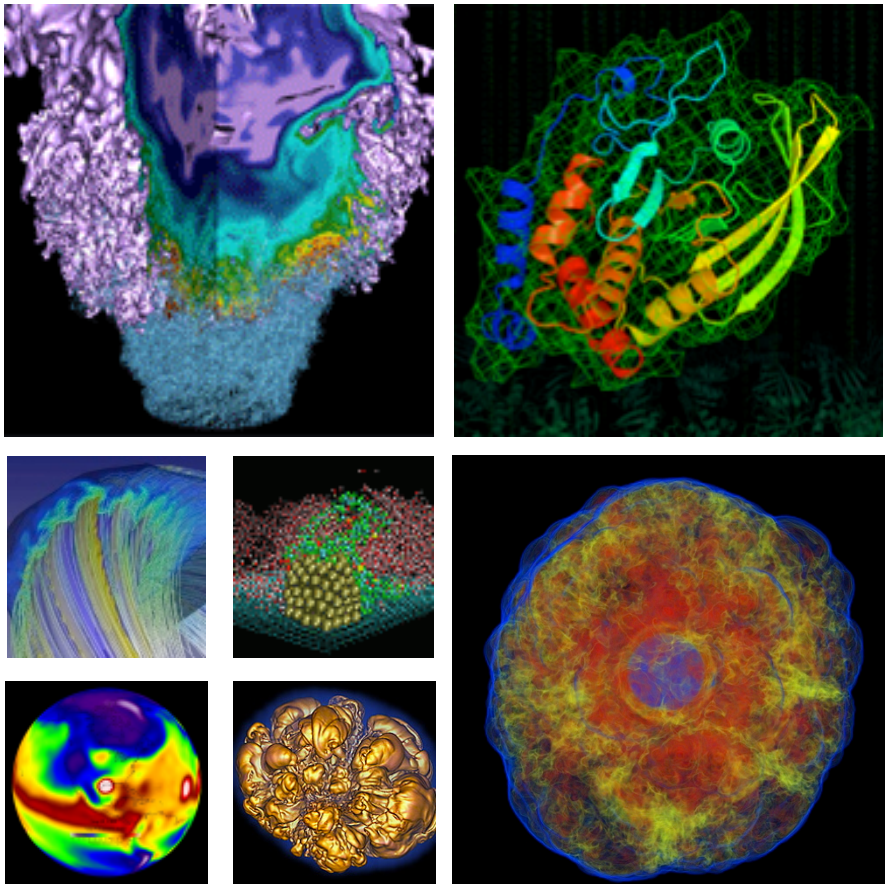

Richard Gerber

NERSC Senior Science Advisor

June 22, 2016

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

NERSC: Mission Computing for the U.S. Department of Energy Office of Science

U.S. DEPARTMENT OF ENERGY | Office of Science

Largest funder of physical science research in U.S.

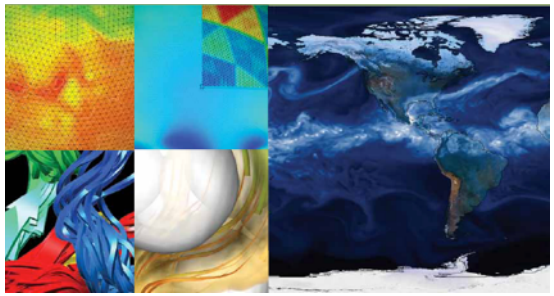

Bio Energy, Environment

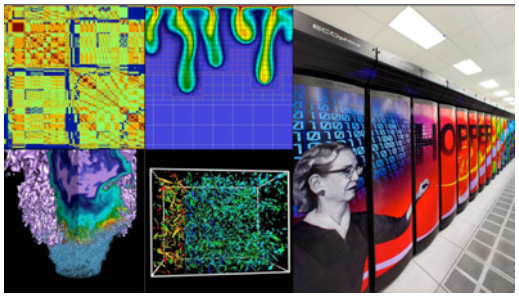

Computing

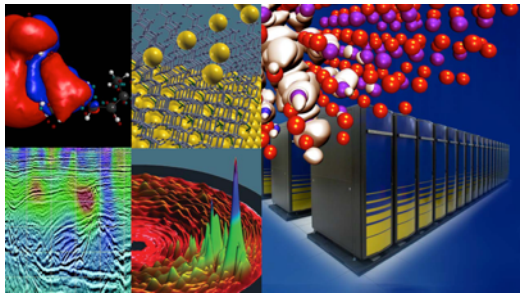

MatSci, ChemSci, Geophysics

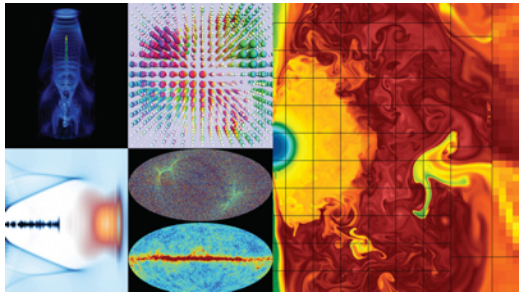

Particle Physics, Astrophysics

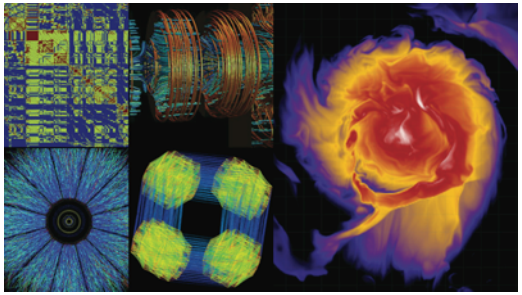

Nuclear Physics

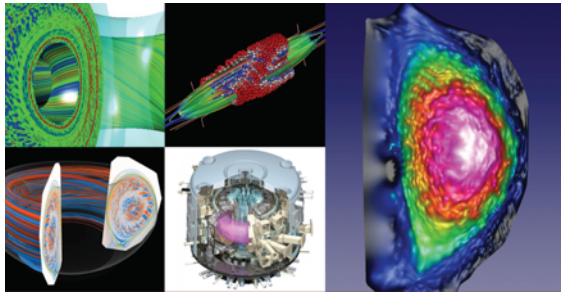

Fusion Energy, Plasma Physics

NERSC – Science First

- **Strong focus on science**
- **~2,000 referred publications per year**
- **First of a kind systems for scientific research**
- **6,000 users**
- **Capability and high throughput computing**
  - ~70% of hours use over 16K cores on Edison
- **Diversity of algorithms (~600 codes)**
- **Extreme scale computing and data analysis**

Journal Covers

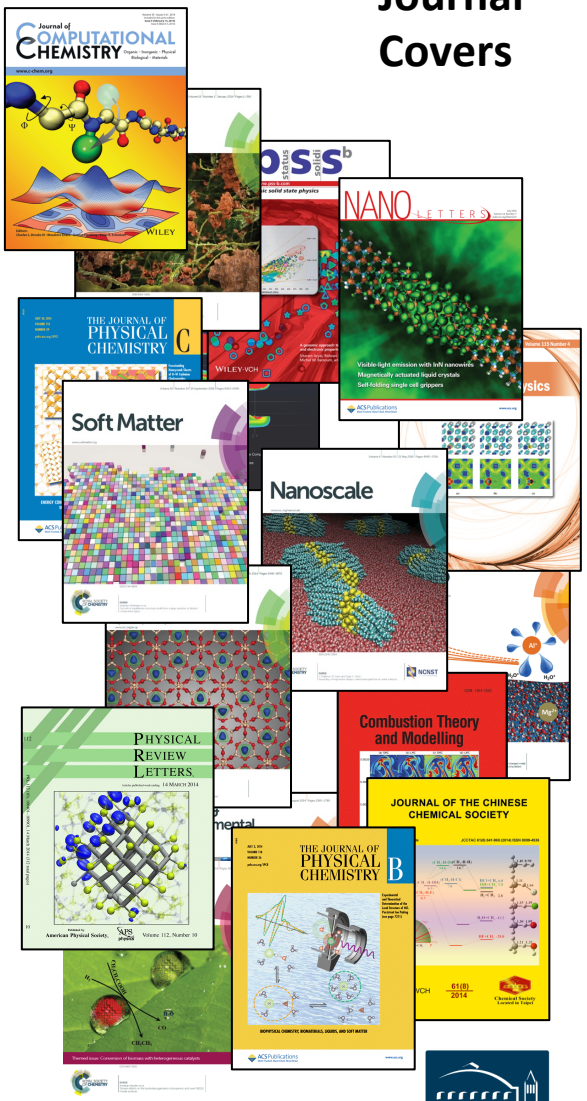

Cori (NERSC 8)

- **Cray XC system with 9,300 Intel Xeon Phi (codenamed: Knights Landing) compute nodes**
  - Intel Xeon Phi partition arriving Summer 2016
  - Self-hosted processor, 68 cores per node
  - On-package 16 GB high-bandwidth memory
- **Cori will support the broad Office of Science research community and begin to transition the workload to more energy efficient architectures**
- **Data Intensive Science Support**
  - 10 Intel Xeon (Haswell) processor cabinets (Phase 1) to support data intensive applications
  - NVRAM Burst Buffer to accelerate data intensive applications (1.5 PB, 1.5 TB/sec)
  - 28 PB of disk, >700 GB/sec I/O bandwidth
- **Robust Application Readiness Plan**
  - Deep engagements, outreach and training
  - Application deep dives with Intel and Cray
  - 8 postdocs integrated with key application teams

Partnership with Los Alamos & Sandia
Sister system to Trinity

Key Cori Advantages for the NERSC Workload

- **Increases NERSC compute capability by 3X**
- **Intel Xeon Phi single socket self-hosted processor**
  - (Relative!) ease of programming using portable programming models and languages (MPI+OpenMP)
- **Low-power manycore (68) processor with hardware threads**
  - Start NERSC community down the path to exascale
- **512b vector units**
  - Opportunity for 32 flops / clock
- **16 GB High bandwidth on-package memory**
  - Bandwidth ~5X that of DDR4 DRAM memory
  - Many scientific applications are memory-bandwidth bound
- **Integrated compute/data system**

To run efficiently on Cori users will have to optimize their codes to:

- **Manage Domain Parallelism**
  - independent program units; explicit
- **Increase Thread Parallelism**
  - independent execution units within the program; generally explicit
- **Exploit Data Parallelism**
  - Same operation on multiple elements
- **Improve data locality**
  - Cache blocking; Use on-package memory

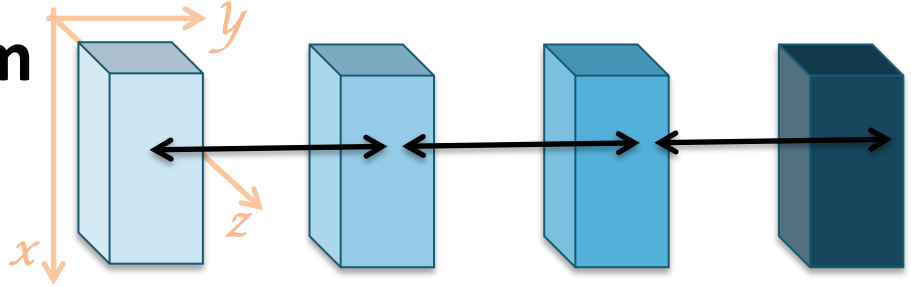

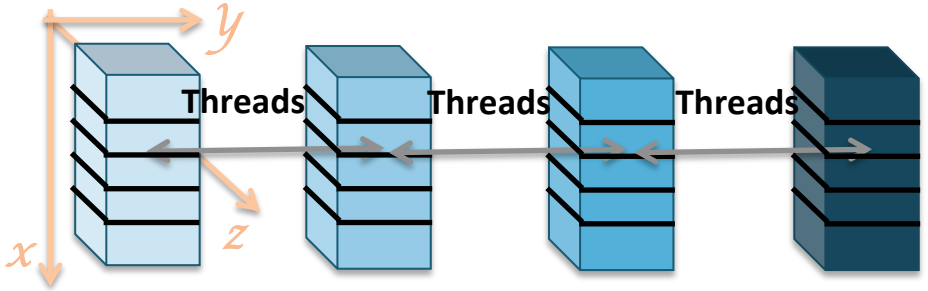


```
|--> DO I = 1, N
|       R(I) = B(I) + A(I)
|--> ENDDO
```

NERSC Exascale Science Application Program (NESAP)

- **Goal: Prepare DOE Office of Science user community for Cori manycore architecture**
- **Partner closely with ~20 application teams and apply lessons learned to broad SC user community**
- **NESAP activities include:**

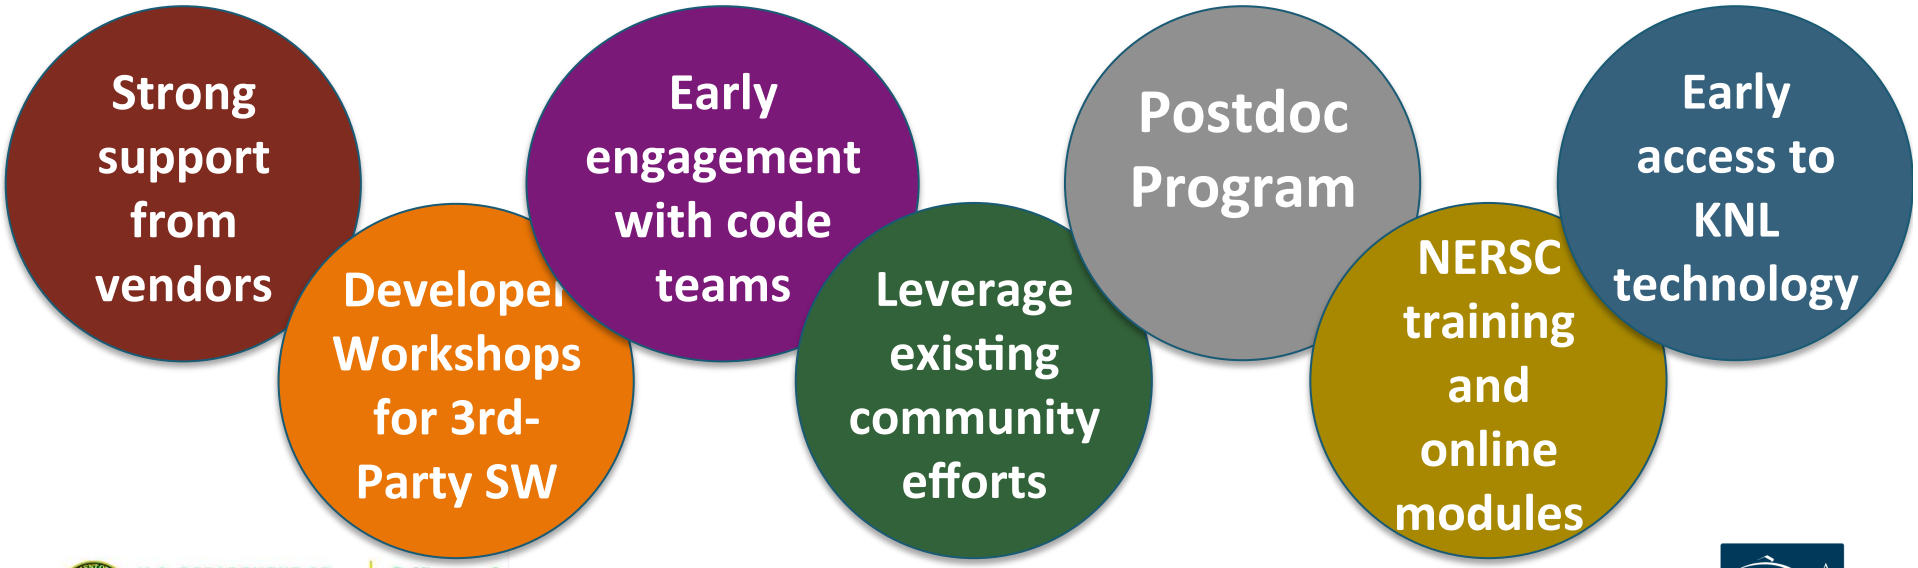

We are initially focusing on 20 codes

Breakdown of Application Hours on Hopper and Edison 2013

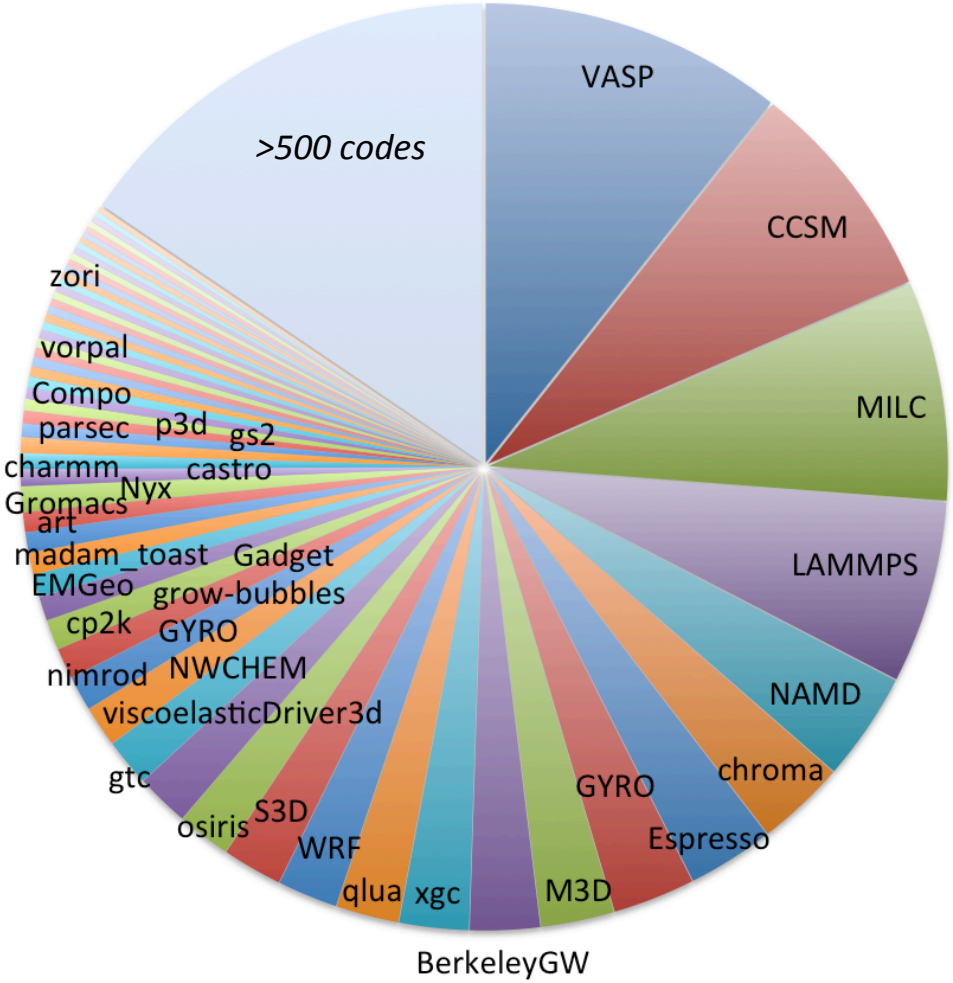

- 10 codes make up 50% of the workload
- 25 codes make up 66% of the workload
- Training and lessons learned will be made available to all application teams

NESAP Codes

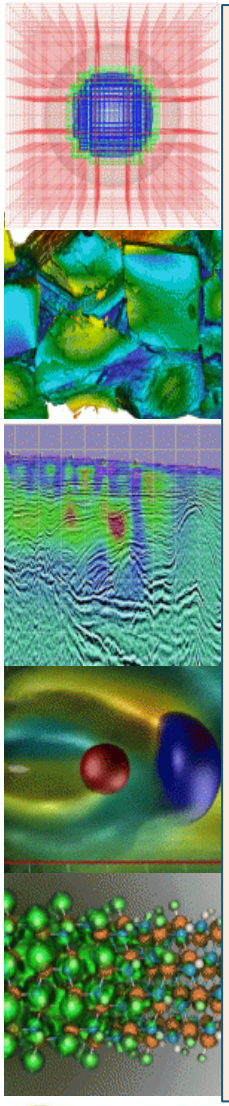

Advanced Scientific Computing Research

Almgren (LBNL) **BoxLib AMR Framework**

Trebotich (LBNL) **Chombo-crunch**

High Energy Physics

Vay (LBNL) **WARP & IMPACT**

Toussaint(Arizona) **MILC**

Habib (ANL) **HACC**

Nuclear Physics

Maris (Iowa St.) **MFDn**

Joo (JLAB) **Chroma**

Christ/Karsch (Columbia/BNL) **DWF/HISQ**

Basic Energy Sciences

Kent (ORNL) **Quantum Espresso**

Deslippe (NERSC) **BerkeleyGW**

Chelikowsky (UT) **PARSEC**

Bylaska (PNNL) **NWChem**

Newman (LBNL) **EMGeo**

Biological and Environmental Research

Smith (ORNL) **Gromacs**

Yelick (LBNL) **Meraculous**

Ringler (LANL) **MPAS-O**

Johansen (LBNL) **ACME**

Dennis (NCAR) **CESM**

Fusion Energy Sciences

Jardin (PPPL) **M3D**

Chang (PPPL) **XGC1**

Resources for Code Teams

- **Early access to hardware**
  - Access to Babbage (Intel Xeon Phi KNC coprocessor) and early "white box" KNL test systems
  - Early access and significant time on the full Cori system
- **Technical deep dives**
  - Access to Cray and Intel staff on-site staff for application optimization and performance analysis
  - Multi-day deep dive ('dungeon' session) with Intel staff at Oregon Campus to examine specific optimization issues
- **User training sessions**
  - From NERSC, Cray and Intel staff on OpenMP, vectorization, application profiling
  - Knights Landing architectural briefings from Intel
- **NERSC staff as code team liaisons (hands on assistance)**
- **Strong connection with IXPUG users group**
- **New NERSC Application Performance Group (Now Hiring!)**
- **8 Postdocs (2 openings!)**

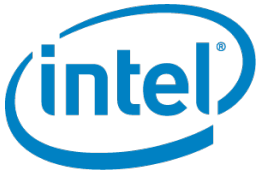

Cori: Advancing Science

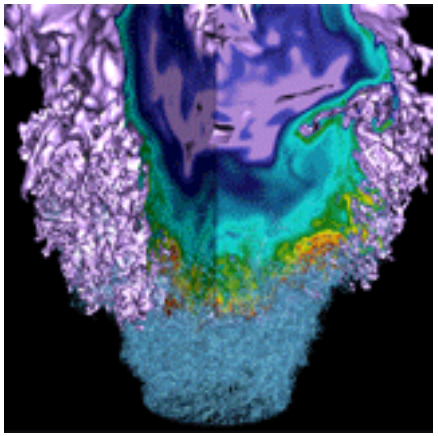

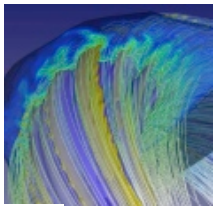

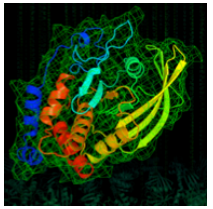

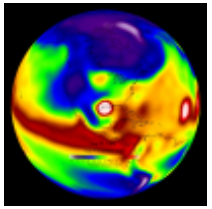

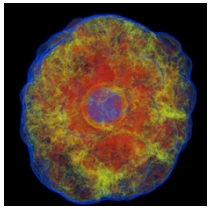

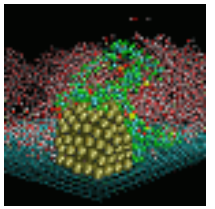

NERSC 40 YEARS at the FOREFRONT

Origin and Growth of Structure in the Universe

- **Large-scale cosmological simulations are vital to extracting knowledge from observations**
  - Connect theory and observations (precision cosmology)
  - "Tool of discovery" for analyzing large data sets
  - Modeling and control of systematics
- **Cori is expected to enable the largest high-resolution cosmology simulation ever**
  - Will support DESI (Dark Energy Spectroscopic Instrument) science projects
  - The HACC code from Argonne National Laboratory has been tuned for excellent performance on the Intel Xeon Phi
  - Extreme data and compute requirements requirements are a good fit for Cori's data (Intel Xeon) & compute (Intel Xeon Phi) partitions

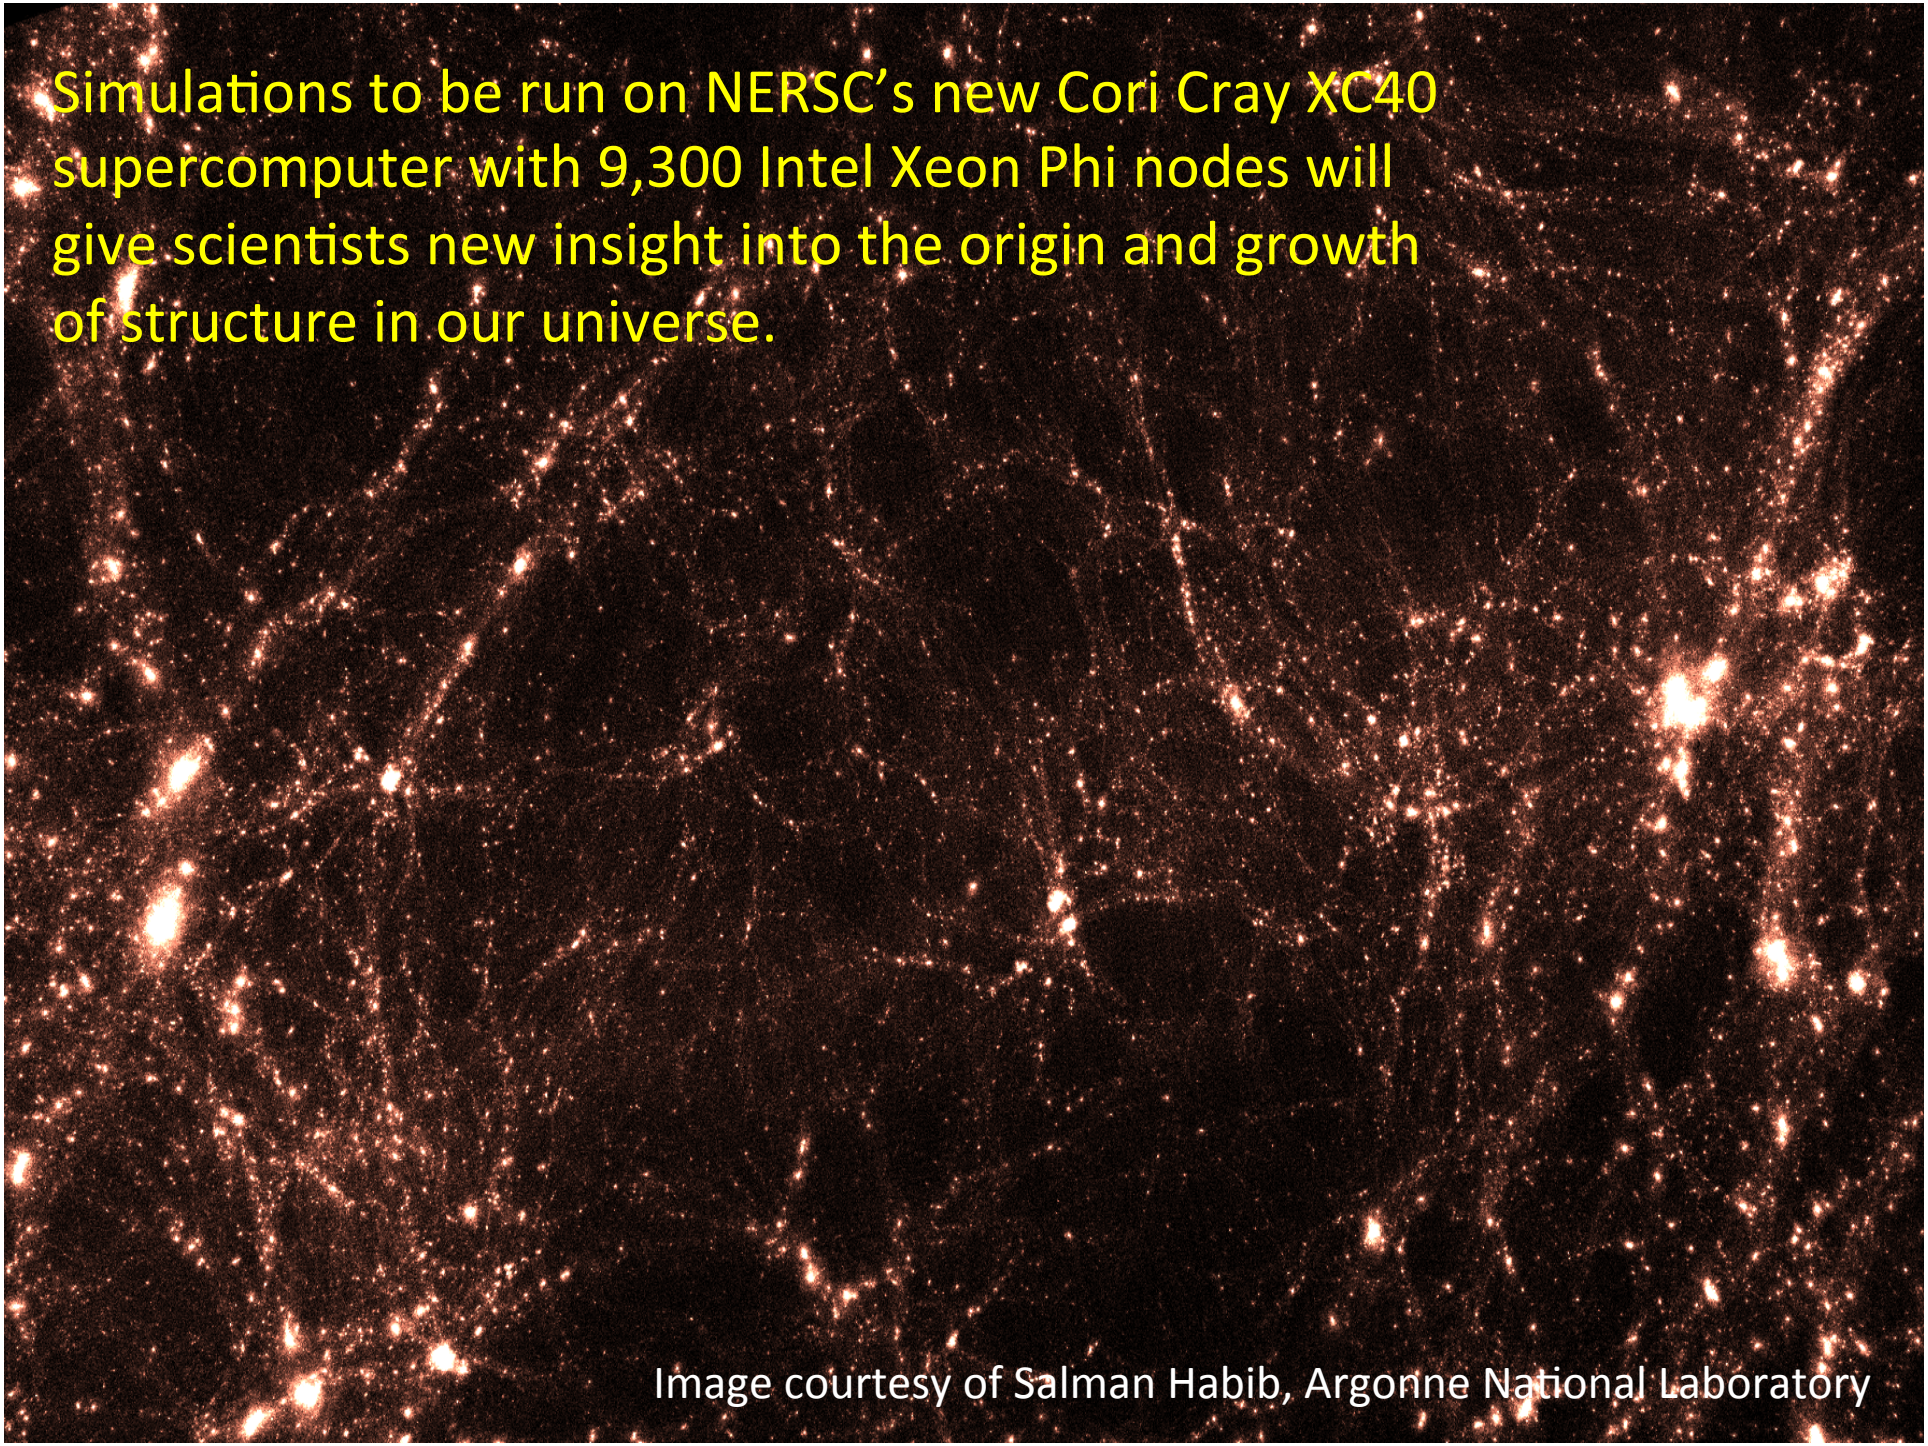
Simulations to be run on NERSC's new Cori Cray XC40 supercomputer with 9,300 Intel Xeon Phi nodes will give scientists new insight into the origin and growth of structure in our universe.
Image courtesy of Salman Habib, Argonne National Laboratory

White Dwarf Mergers

- **Binary white dwarf star systems, the embers of stars like our Sun, can merge by giving off gravitational radiation and may produce Type 1a supernovae.**
  - Gravitational waves first detected this year! (likely Nobel prize)
  - Type 1a supernova were key to Nobel-prize winning discovery of accelerating expansion of the universe
- **The increased capability provided by Cori will allow a range of studies never before possible**
  - Low power Intel Xeon Phi cores provide more capability and capacity than would other be available.
  - Simulations performed using the *Castro* code, which makes use of the BoxLib AMR framework , which is being optimized for the Intel Xeon Phi at NERSC

Scientists at Stony Brook University will be able to use Cori to gain new understanding into the dynamics of the merger of two white dwarf stars and how that could lead to a supernova explosion.

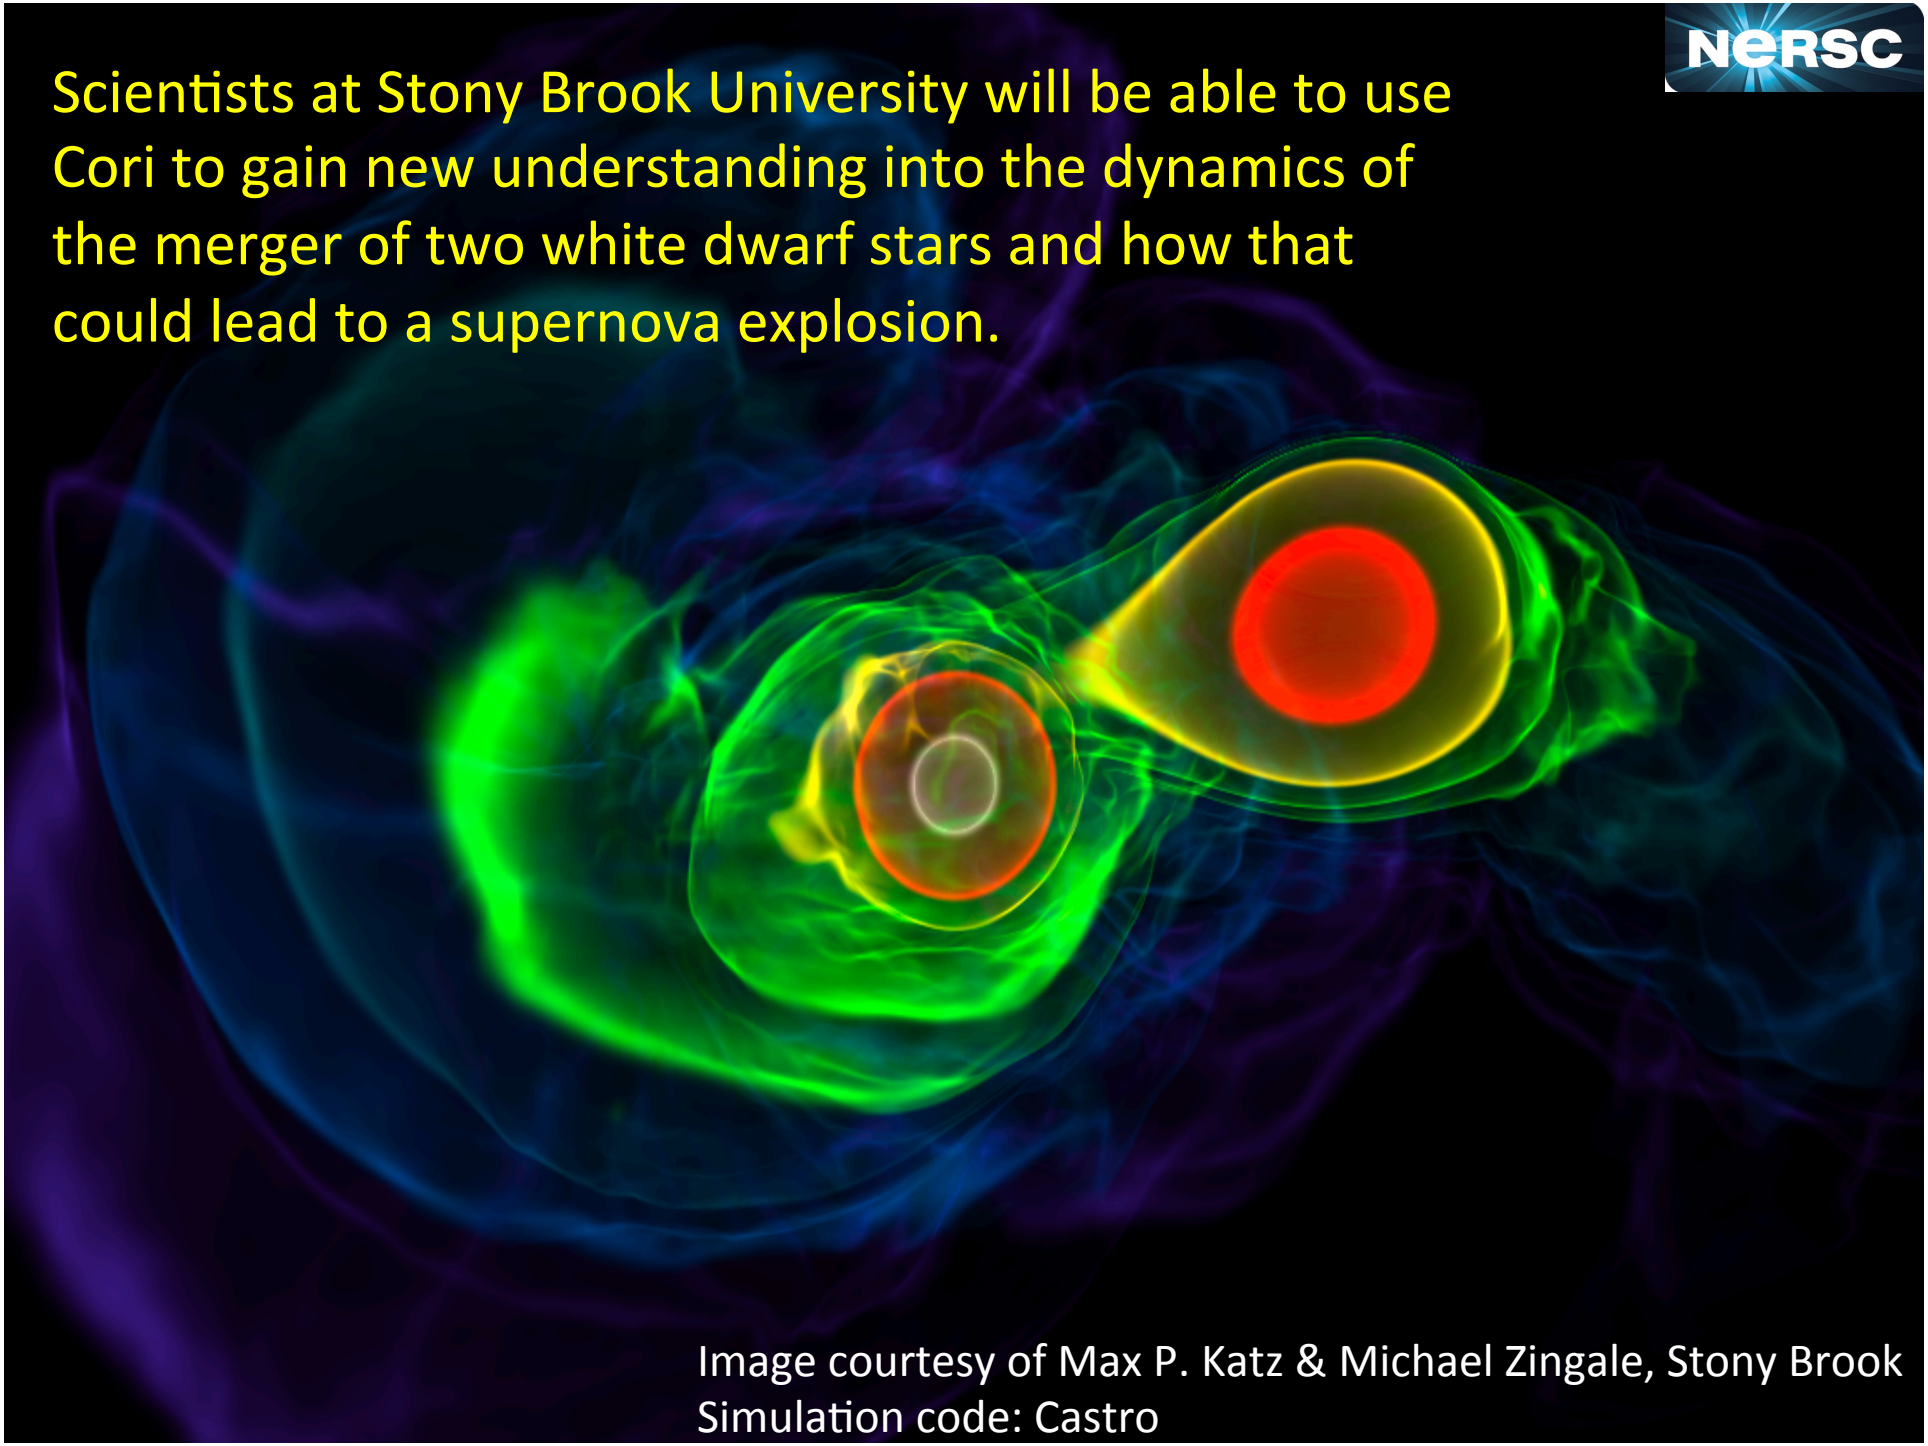

Image courtesy of Max P. Katz & Michael Zingale, Stony Brook
Simulation code: Castro

Revolutionary Particle Accelerators

Particle accelerators are essential tools in modern life that power scientific discovery, cure cancer, secure our borders, and help create a wide range of products.

Existing accelerators are huge and costly.

- New Rochester Mayo Clinic Proton Therapy Center: $188 M
- Heidelberg Proton & Carbon Therapy Center, €119M, 670 Tons
- LHC, $10 B, 27 km, 150 MW

Laser-Plasma accelerators: compact particle acceleration at potentially a fraction of the price.

High-resolution three-dimensional modeling is needed to capture all the possible physical effects in the full range of spatial scales.

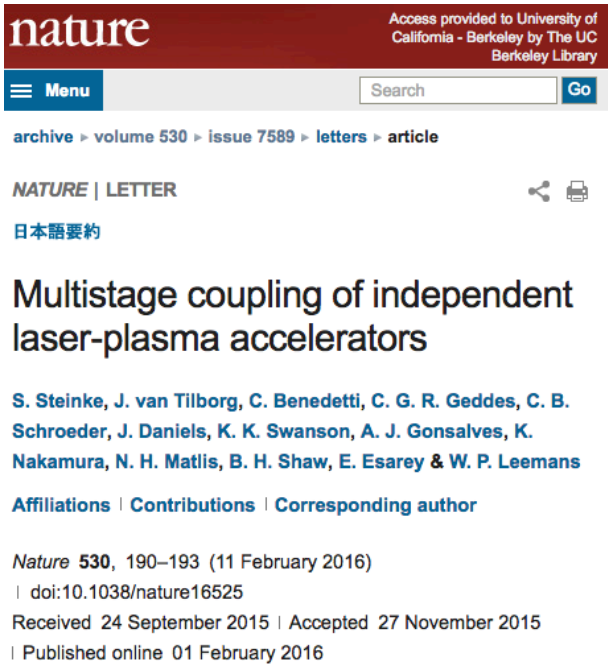

High-resolution three-dimensional modeling is not possible on Edison or Cori phase 1 (Intel Xeon) but will be accessible on Cori phase 2 (Intel Xeon Phi partition).

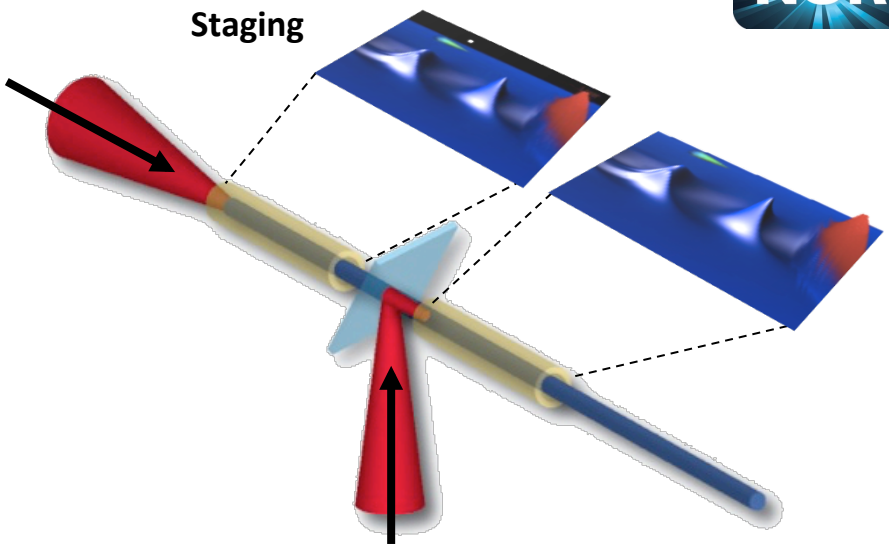

Two-color injection of ultra-high quality beam

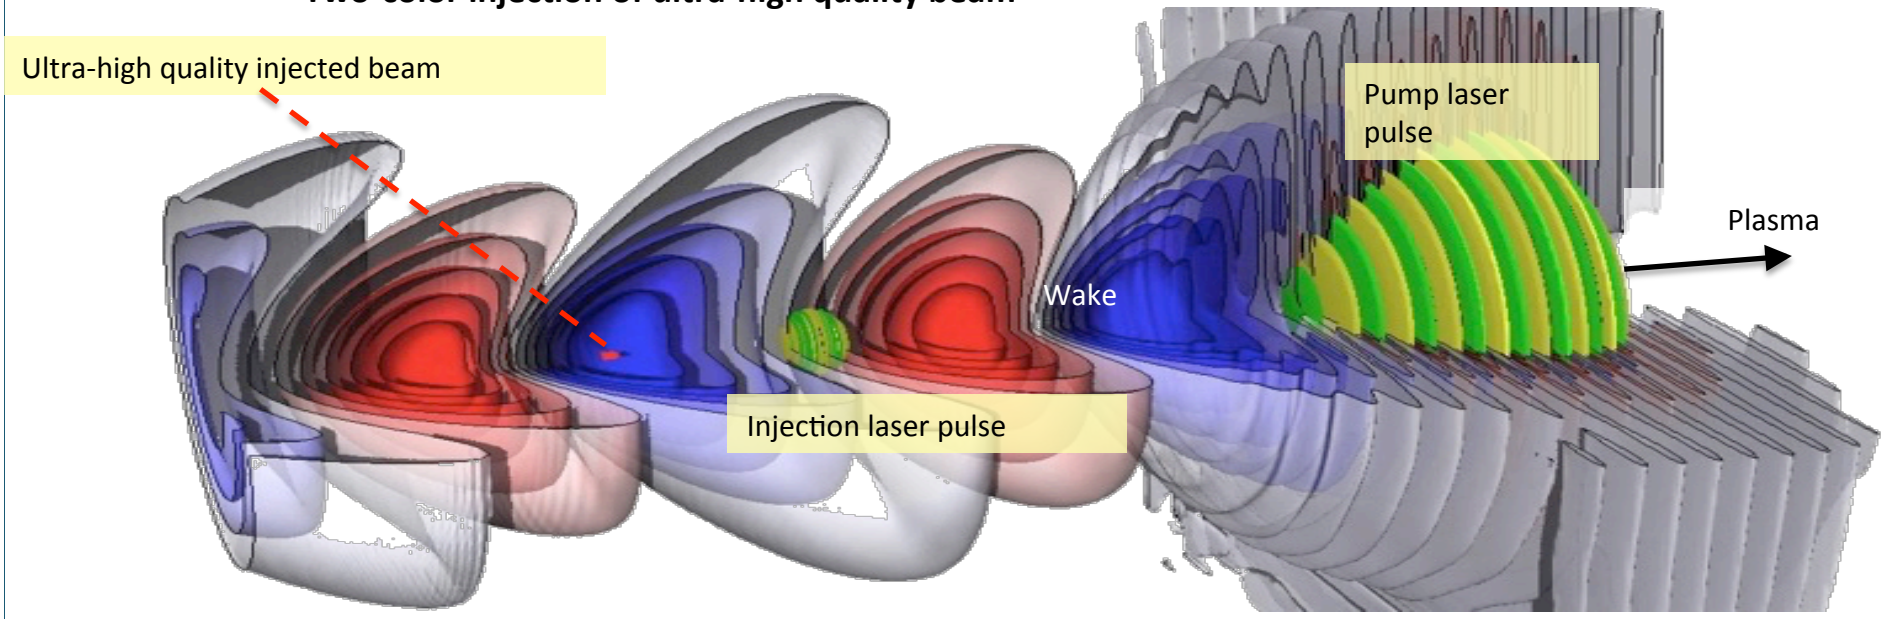

S. Steinke, et al., ``Multistage coupling of independent laser plasma accelerators,'' Nature (2016).
L.-L. Yu et al., "Two-color laser ionization injection," Phys. Rev. Lett. (2014)

Image courtesy of Jean-Luc Vay, Berkeley Lab

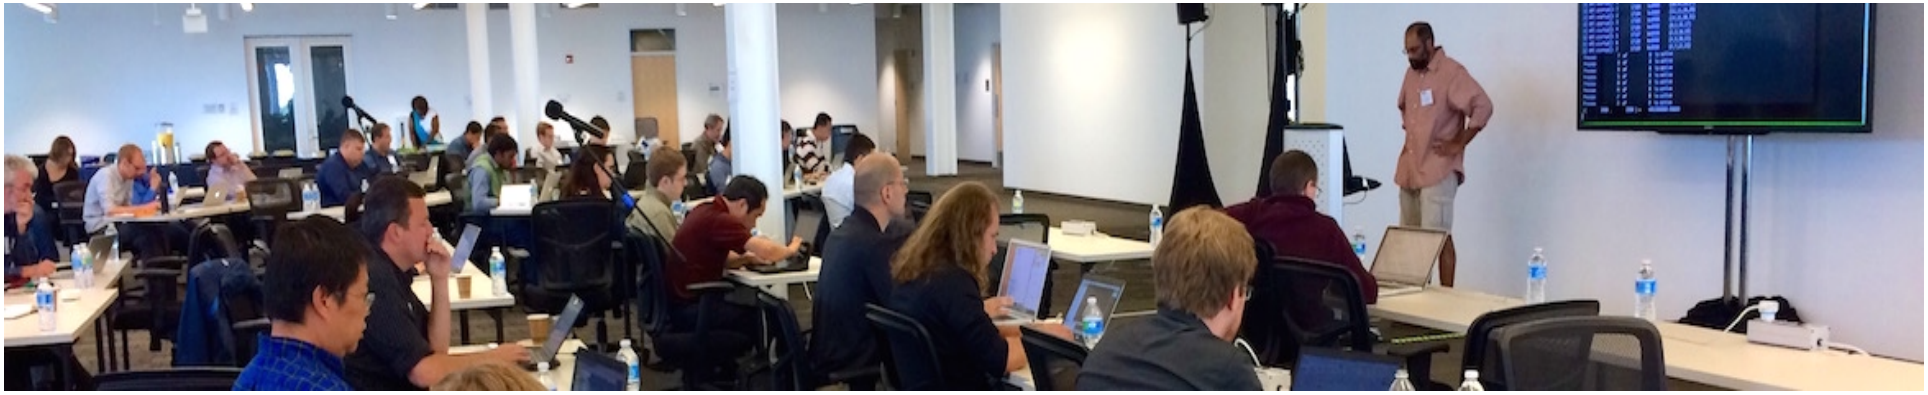

IXPUG

THE INTEL® XEON PHI™ USERS GROUP

IXPUG is a *community* of developers and performance experts sharing experiences and best practices in order to create applications optimized for Intel Xeon Phi processors.

BOF: Gearing Up Application Performance for Intel Xeon Phi (KNL) Supercomputers

Location: Frankfurt am Main, Germany
Date: Wednesday, June 22, 2016, 8:30-9:30 am (Frankfurt time)
Venue: Substanz 1+2, Forum

Workshop: Application Performance on Intel Xeon Phi – Being Prepared for KNL and Beyond"

Location: Frankfurt am Main, Germany
Date: Thursday, June 23, 2016, 8:30am-6:00pm
Venue: Marriott Frankfurt Hotel

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB
Lawrence Berkeley National Laboratory

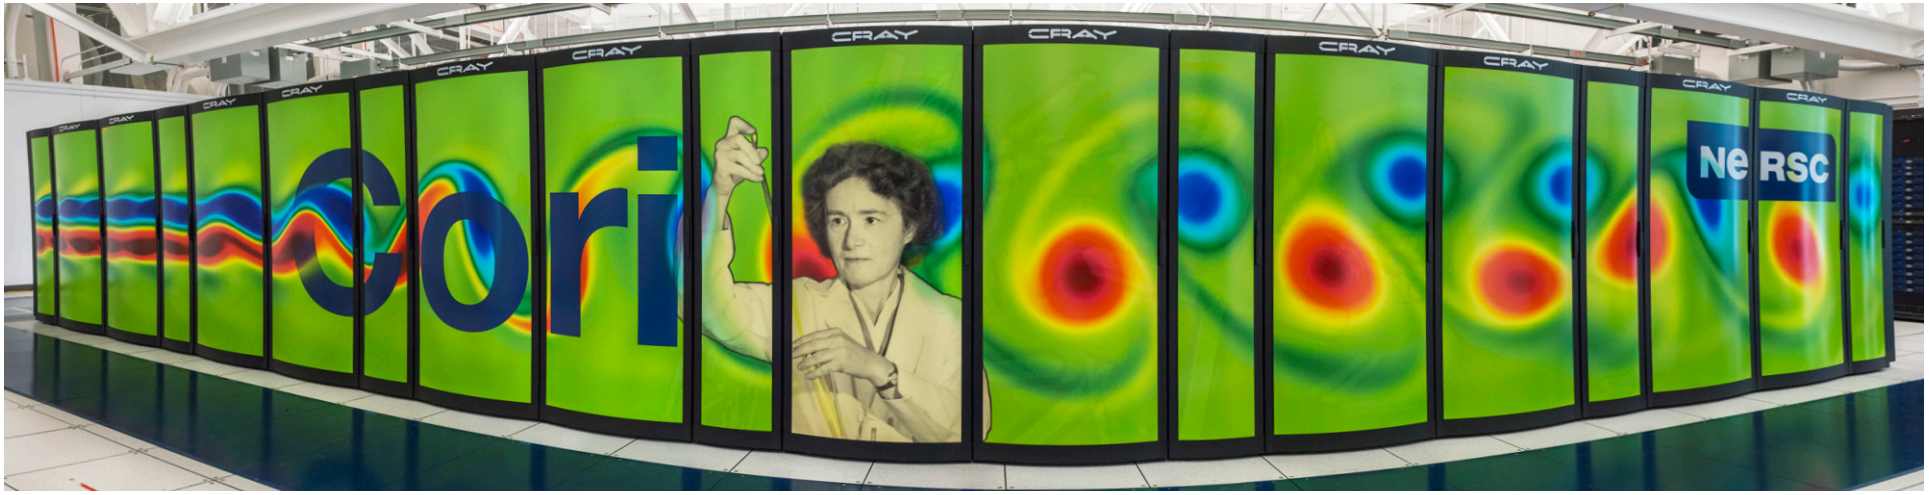

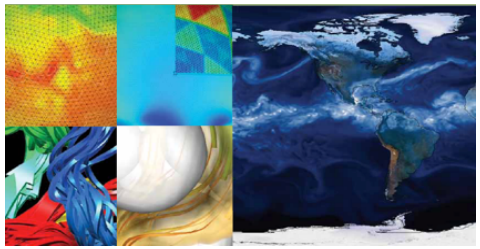

Bio Energy, Environment

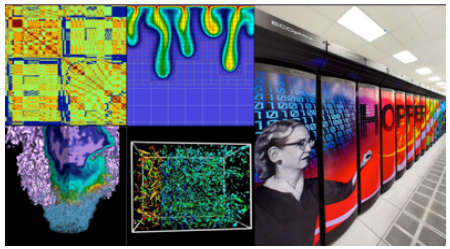

Computing

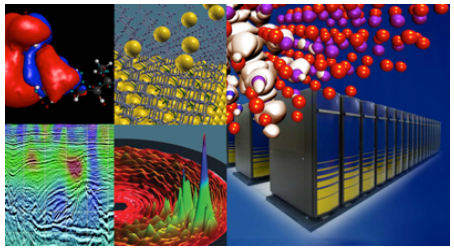

Materials, Chemistry, Geophysics

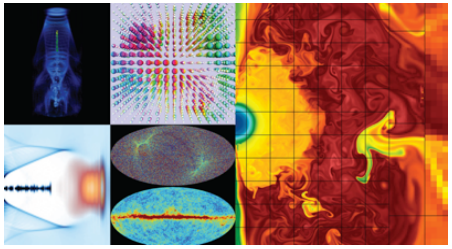

Particle Physics, Astro

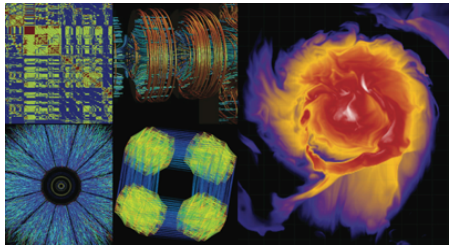

Nuclear Physics

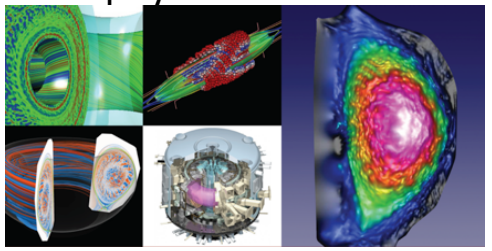

Fusion Energy, Plasma Physics

U.S. DEPARTMENT OF ENERGY | Office of Science

CRAY THE SUPERCOMPUTER COMPANY

NERSC

intel

BERKELEY LAB


NERSC