NERSC Workload Analysis on Hopper

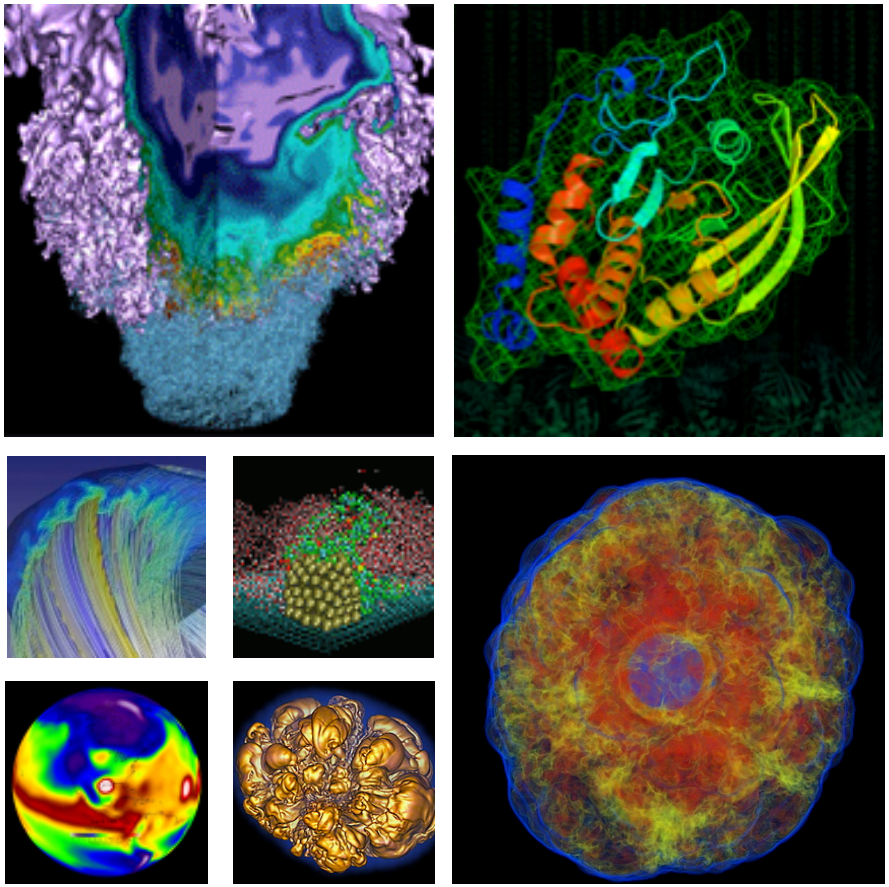

Katie Antypas, Brian Austin, Tina Butler, Richard Gerber, Cary Whitney, Nick Wright, Woo-Sun Yang, Zhengji Zhao

February 22, 2013

Understanding the NERSC workload is key to procuring productive, high performing systems for science

- **Conducted workload analysis to understand application requirements and guide future system procurements**
- **Important for understanding efforts needed to transition workload to future architectures**
- **Analyzed the workload by:**
  - Science area
  - Application code
  - 3rd party application usage
  - Algorithm
  - Job size
  - Memory usage
  - Threading usage
  - Library usage
- **I/O workload analysis is in progress and will be ready in Mid-March 2013**

Methods

- **Data collected in this presentation came from a variety of sources**
  - System accounting logs
  - NIM database
  - ALPS command line capture log
  - Automatic Library Tracking Database (ALTD)
  - Application Resource Usage (ARU) Tool

NERSC serves a broad range of science disciplines for the DOE Office of Science

2012 Allocation Breakdown

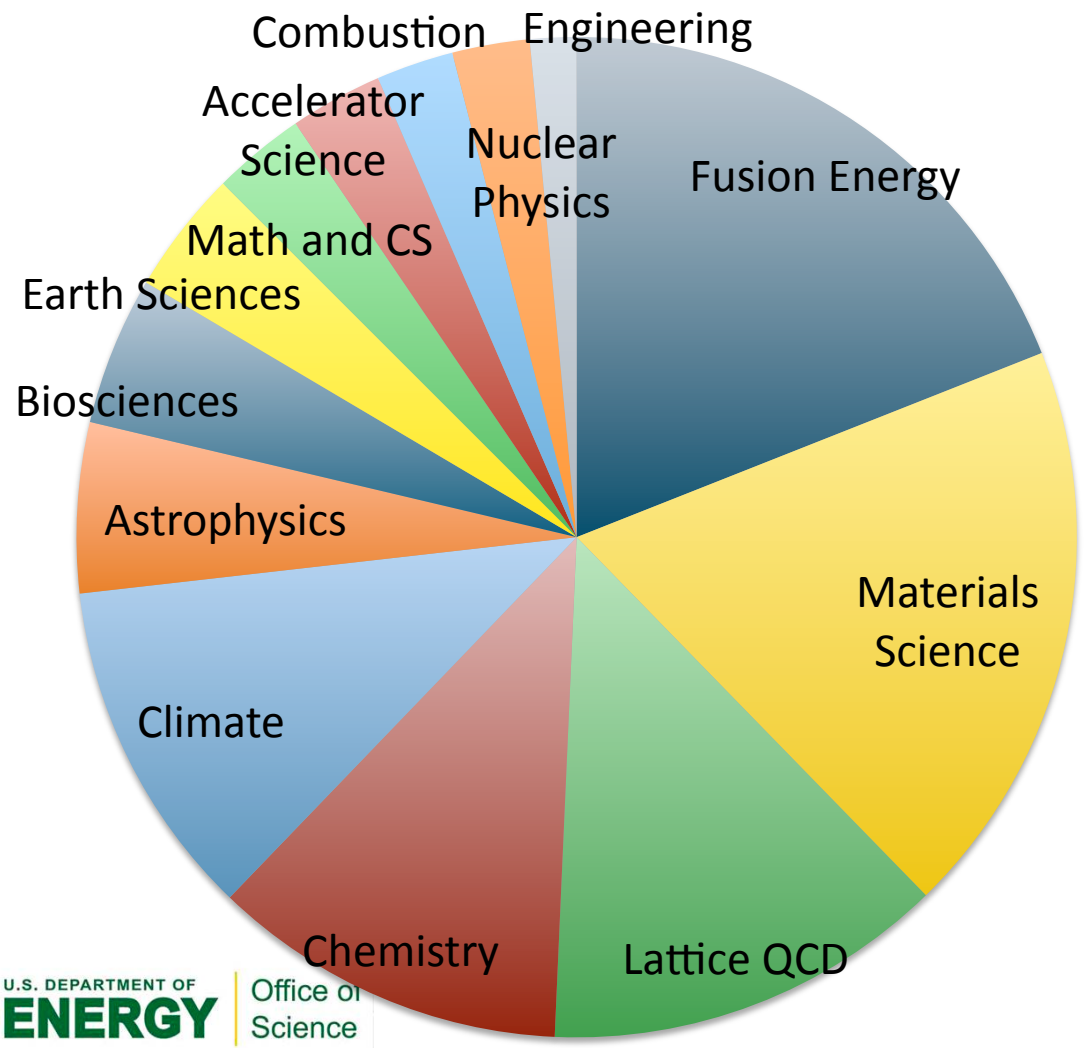

NERSC serves:

- **Over 4500 users**
- **Over 650 projects**

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

Over 650 applications run on NERSC resources

Top Application Codes on Hopper by Hours Used

Jan – Nov 2012

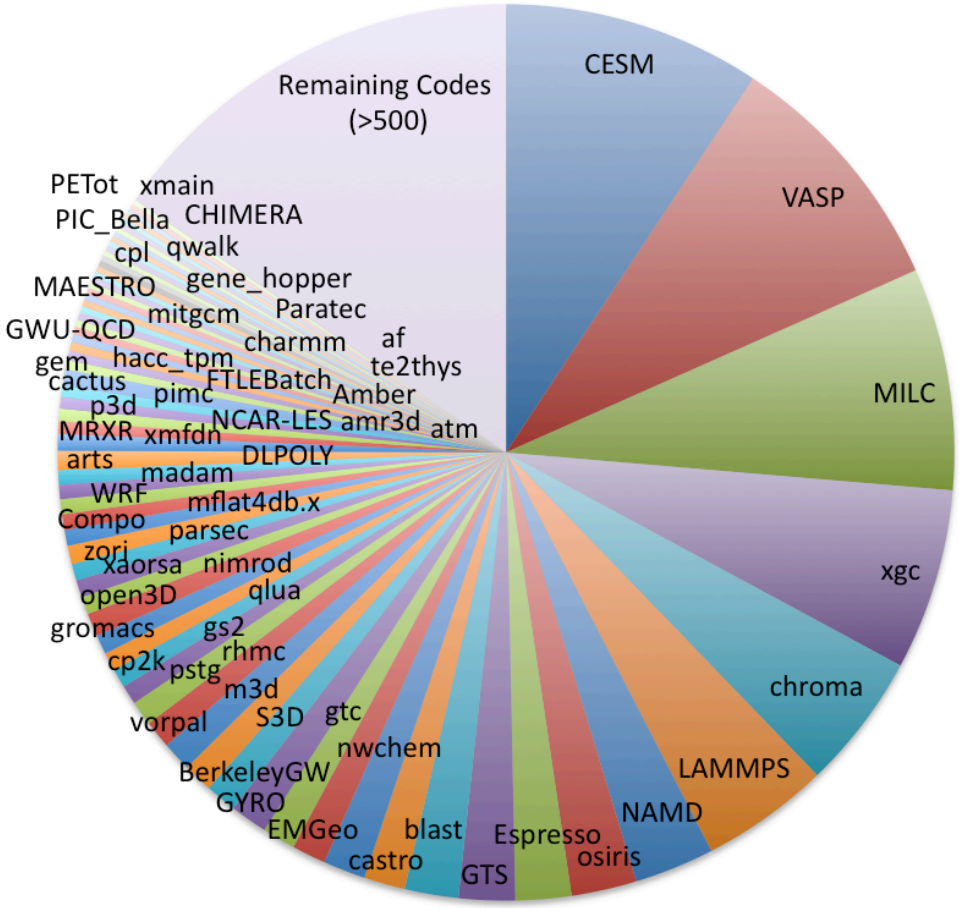

- **10 codes make up 50% of workload**
- **25 codes make up 66% of workload**
- **75 codes make up 85% of workload**
- **remaining codes make up bottom 15% of workload**

Approximately 80% of the workload needs to transfer to NERSC-8
(20% can remain on Edison)

NERSC installs and directly supports software used by over 20% of workload

Percent of workload using NERSC installed and supported application codes

(On Hopper by Number of Hours)

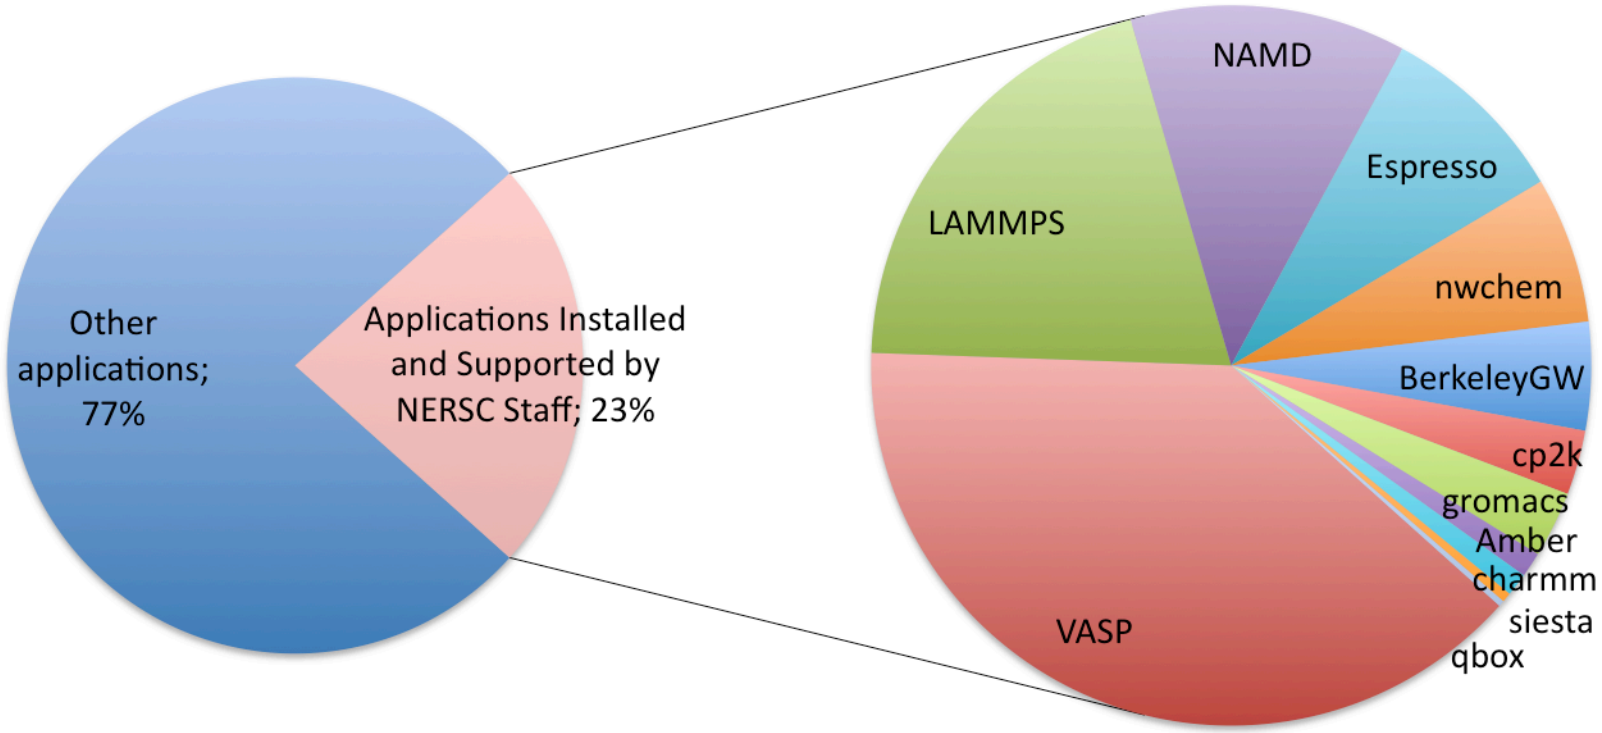

In addition to science area and code diversity, NERSC supports a broad range of algorithms

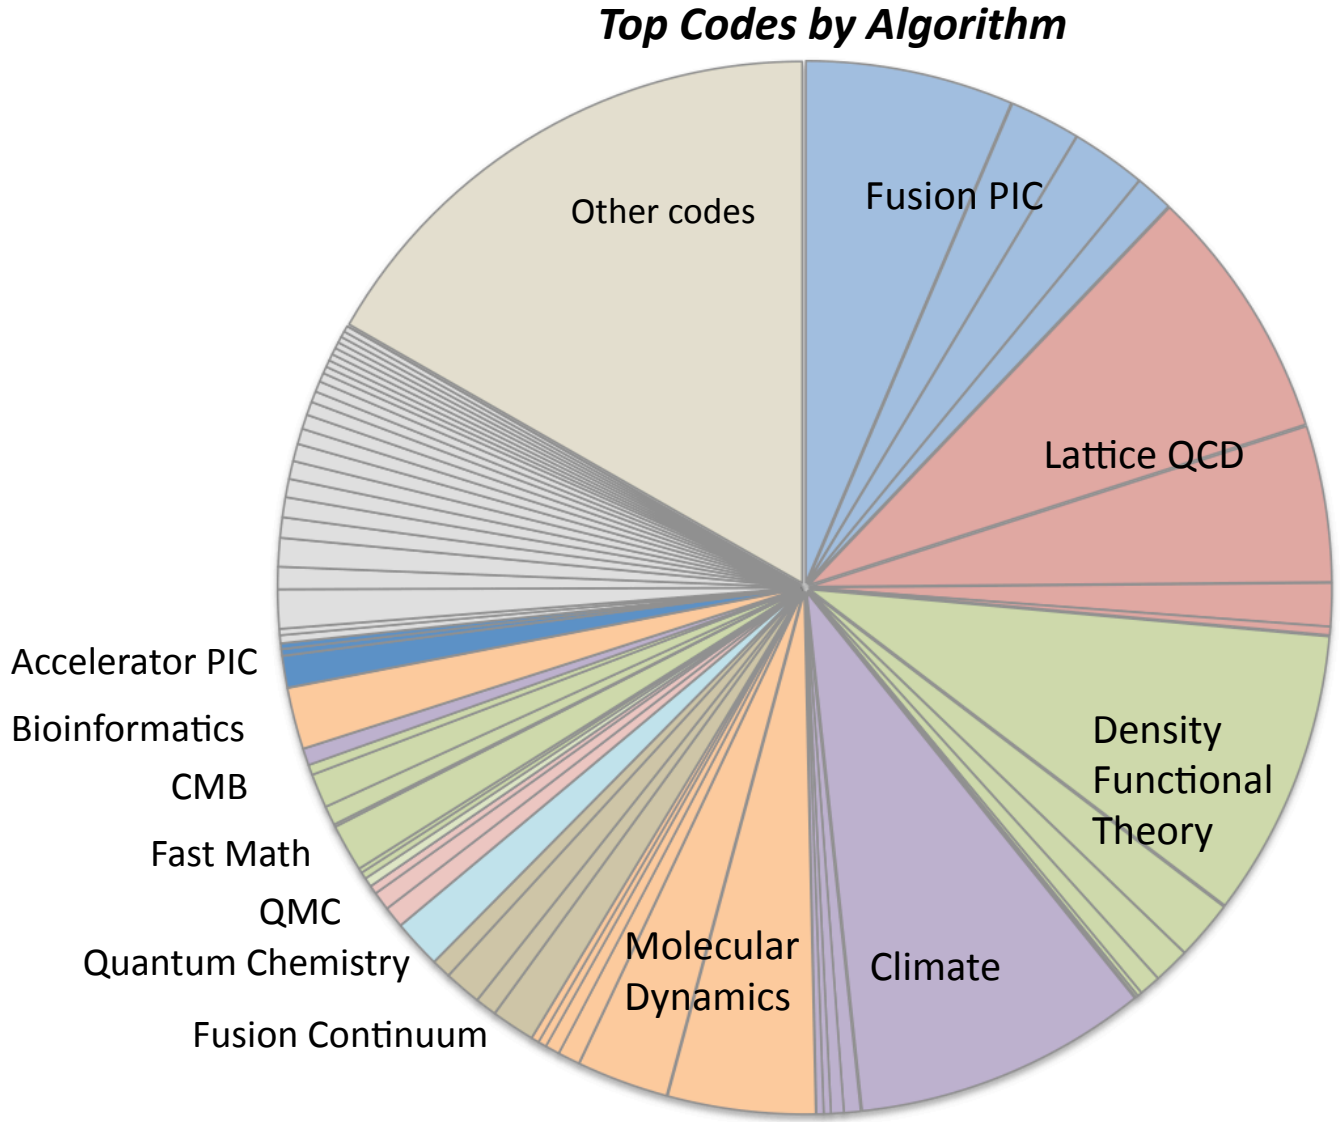

NERSC users run applications at every scale to conduct their research

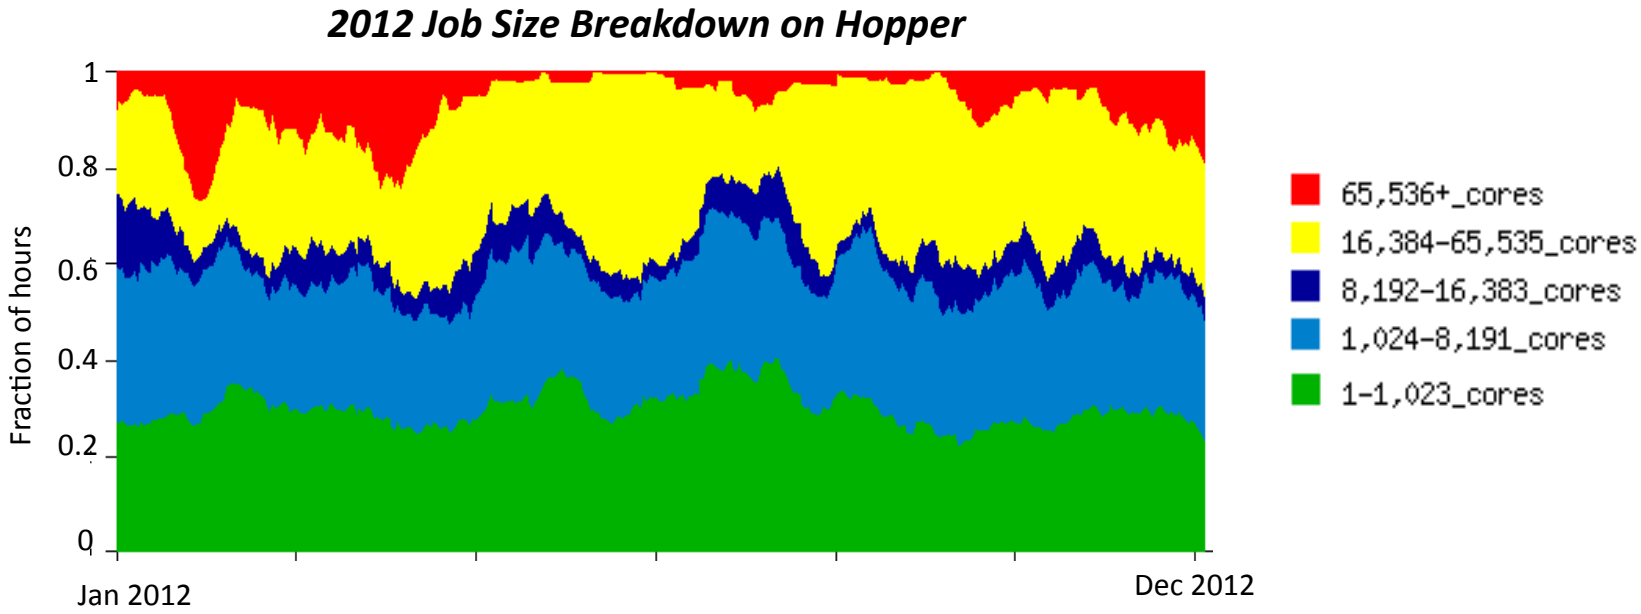

Approximately 15% of computational cycles use more than 40% of the compute nodes ***and*** 40% of cycles use more than 10% of the system.

80% of the Hopper workload uses less than or equal to 1 GB of memory per task

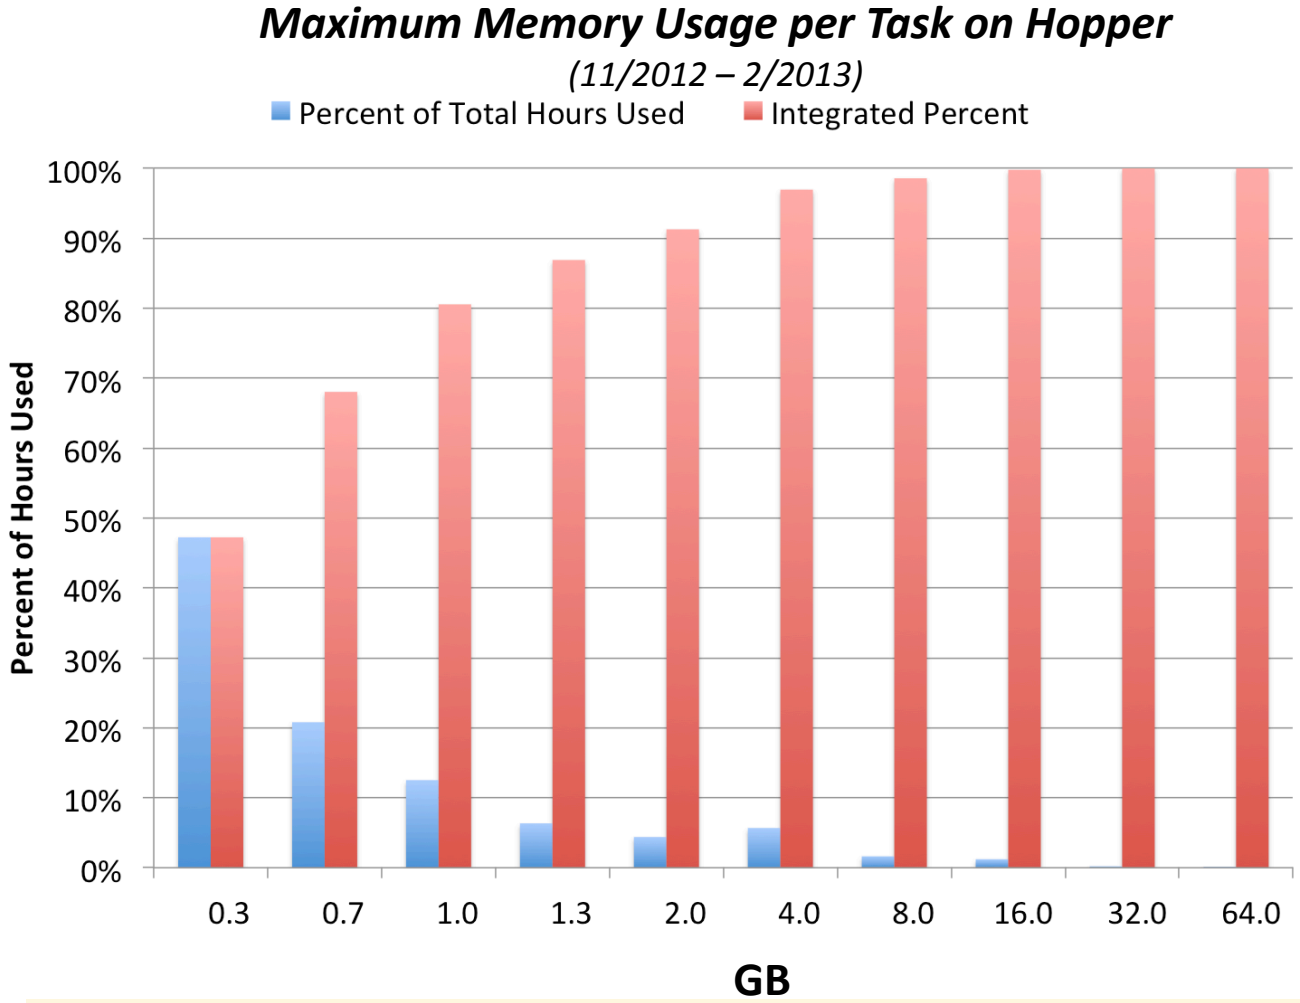

And >90% of the Hopper workload uses 2GBs or less per task

Over 20% of hours used on Hopper run with more than a single thread

% of Hours Used by Thread Count

(Hopper system 11/2012-2/2013)

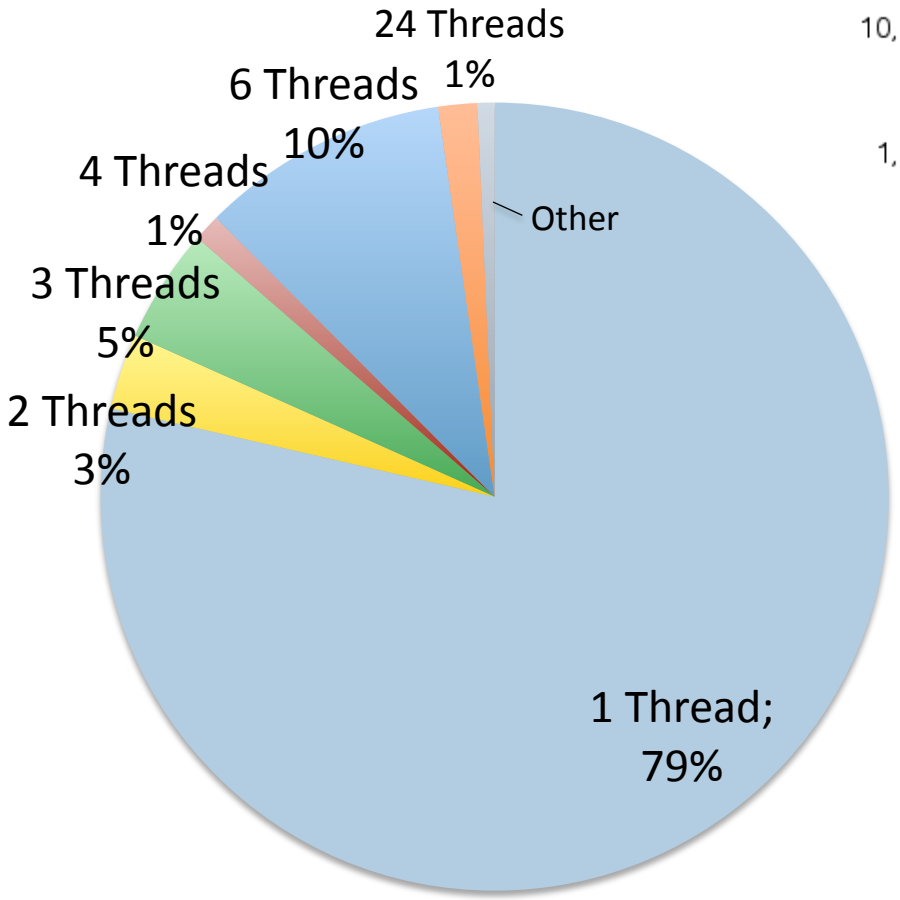

Node Hours used by Application Thread Count

(Hopper system 11/2012-2/2013)

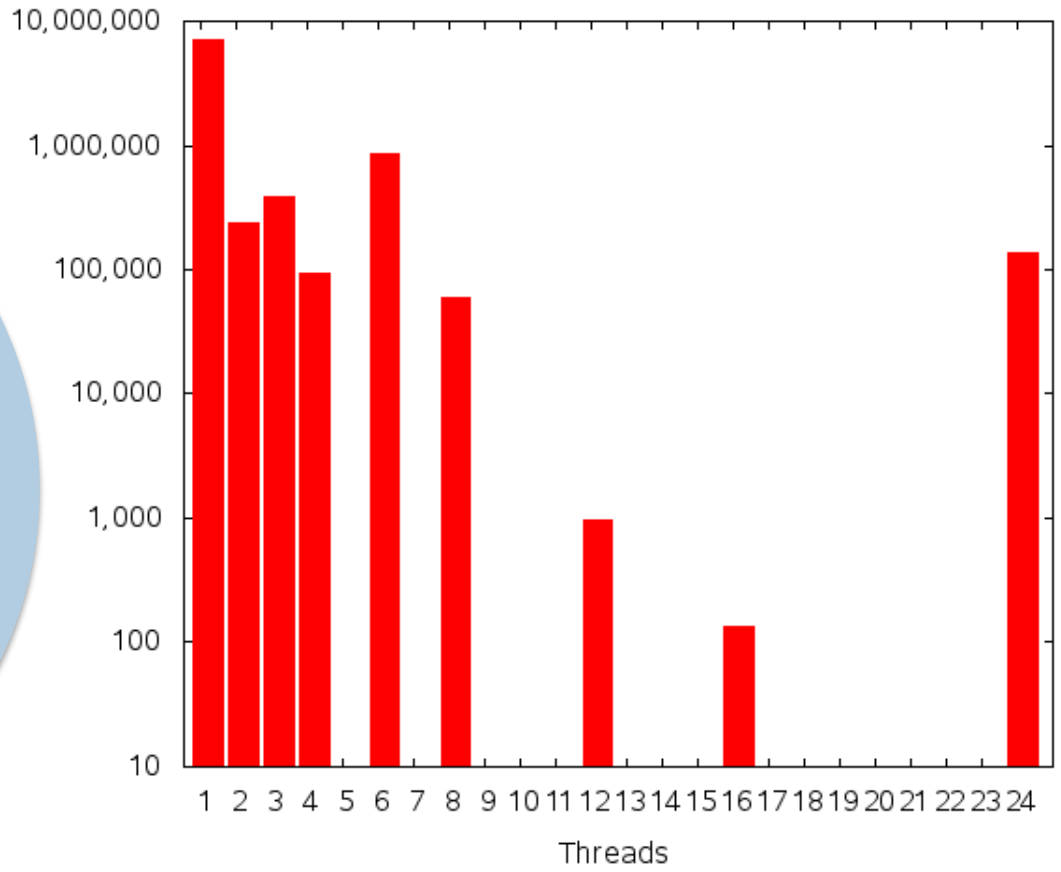

NERSC users require an array of software to conduct their research

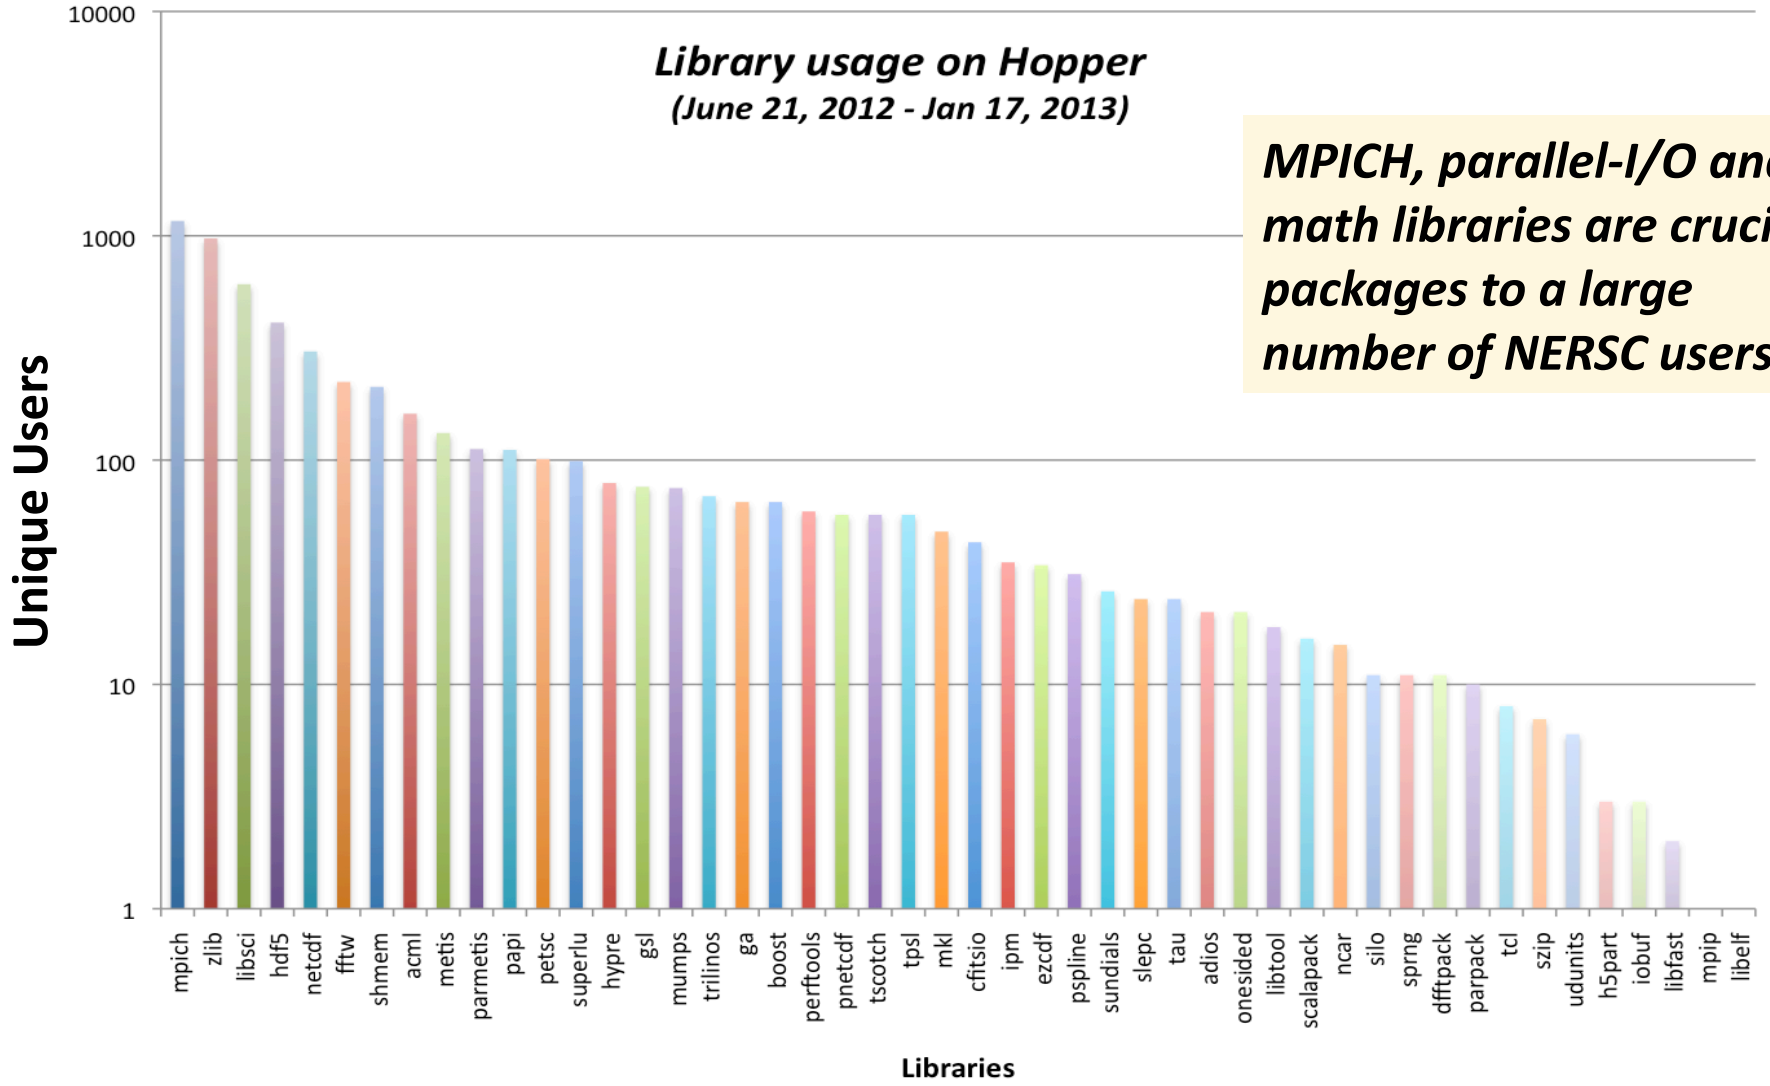

MPICH, parallel-I/O and math libraries are crucial packages to a large number of NERSC users

Summary

- NERSC has new tools to capture memory, threading, and library usage on Cray systems
- New tools will allow workload characterization on both Hopper and Edison systems
- Understanding workload is key to quantifying the effort that will be required to transition applications to future architectures

- Stay tuned for I/O workload analysis

Journal Cover Stories from NERSC-Enabled Research 2012

Devanathan, PNNL: BES

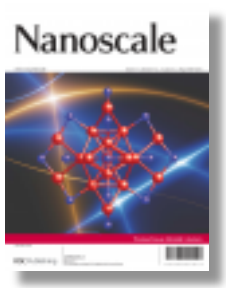

Jiang, ORNL: BES

Jiang, ORNL: BES

Daggett, U. Washington: BER

Dupuis, PNNL: BES

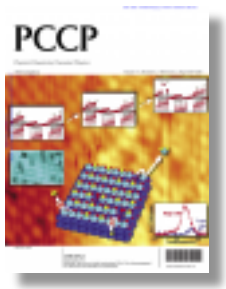

Petrik, PNNL: BES

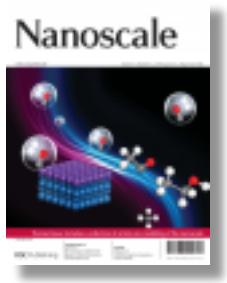

Jiang, ORNL: BES

Smit, UCB: BES

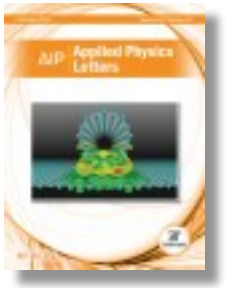

Varga, Vanderbilt: BES

Sugiyama, MIT: FES

Snurr, Northwestern: BES

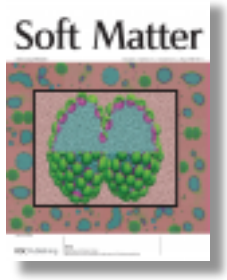

Striolo, U. Oklahoma: BES

Liang, U. Maryland: BER

U.S. DEPARTMENT OF ENERGY | Office of Science

BERKELEY LAB

In addition to science area and code diversity, NERSC supports a broad range of algorithms

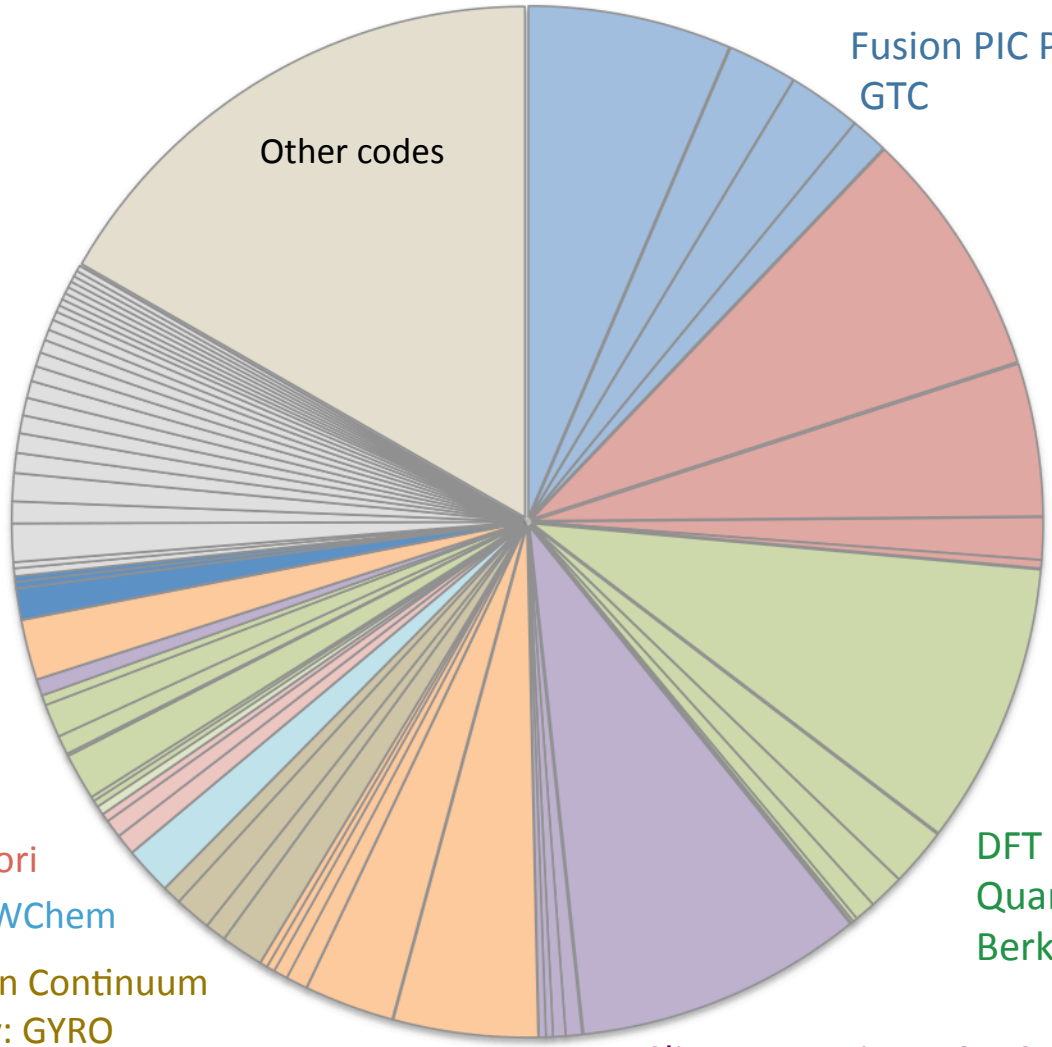

In addition to science area and code diversity, NERSC supports a broad range of algorithms

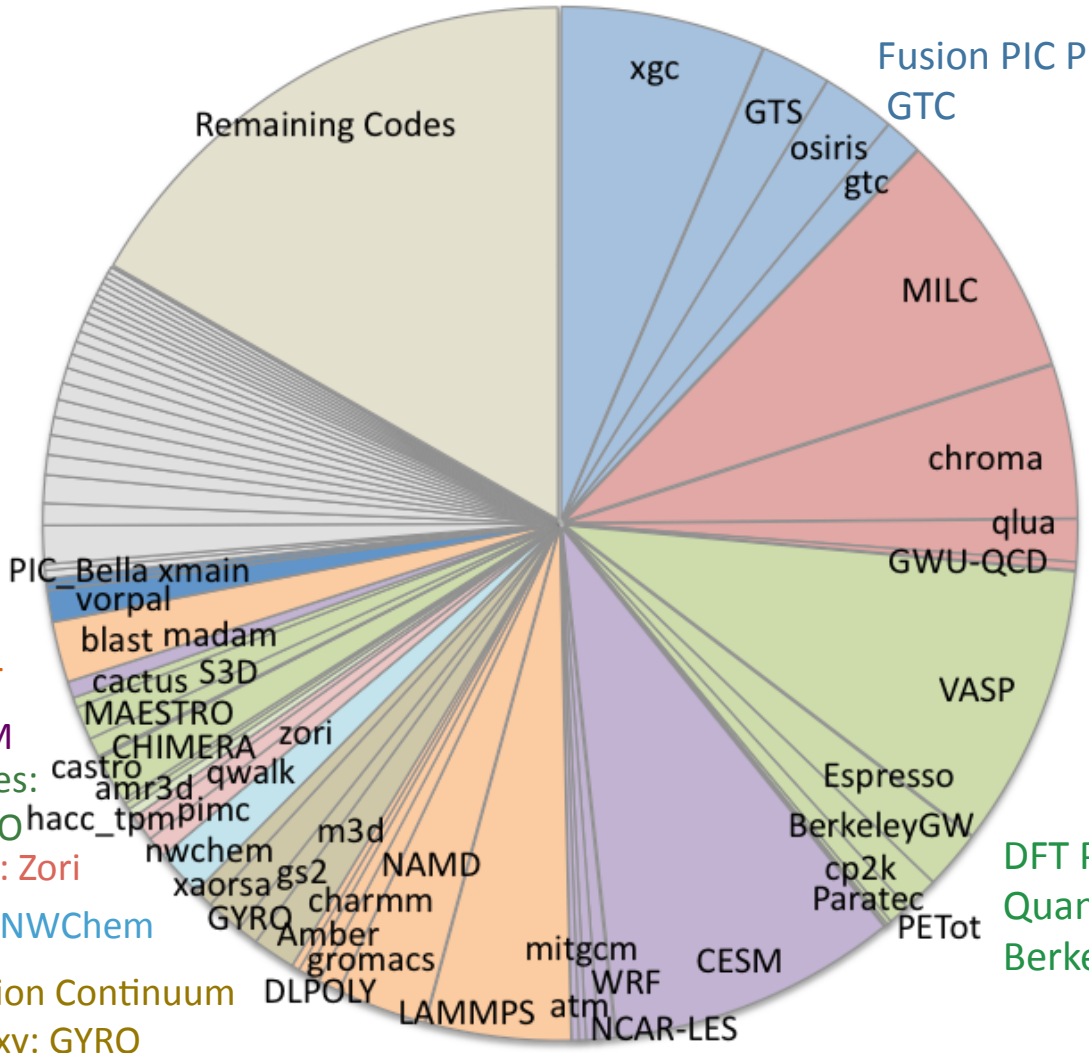

Fusion PIC Proxy: GTC

Lattice QCD Proxy: MILC

DFT Proxies: Quantum Espresso Berkeley GW

Climate Proxies: POP, SE-CAM, MPAS, WRF

Molecular Dynamics Proxies: NAMD, Amber

Fusion Continuum Proxy: GYRO

Quantum Chemistry Proxy: NWChem

QMC Proxy: Zori

Fast Math Proxies: FLASH, MAESTRO

CMB Proxy: MADAM

Bioinformatics Proxy: BLAST

Accelerator PIC: Impact