IN SITU VISUALIZATION HACKATHON 2018

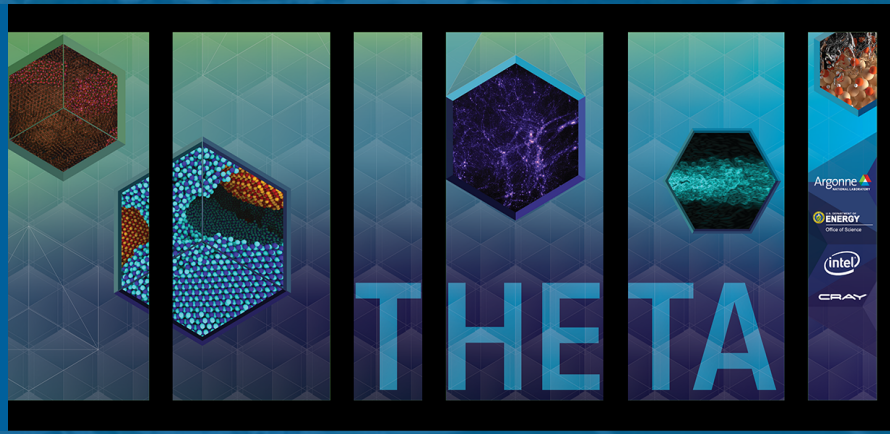

web.alcf.anl.gov/visinternal/MISC/insitu_2018_agenda.html

July 10-12, 2018

IXPUG
INTEL EXTREME PERFORMANCE USERS GROUP

Argonne
NATIONAL LABORATORY

ALCF

User Facility

- Deliver cycles to support computational science
- Partner with community to produce science
- Partner with community on R&D in hardware and software

In 2017

Delivered:

- ***8 Billion*** core-hours of total compute time

Supported:

- ***356*** active projects
- ***976*** facility users

Producing:

- ***211*** publications

3 **Supercomputers in the Top500 list**

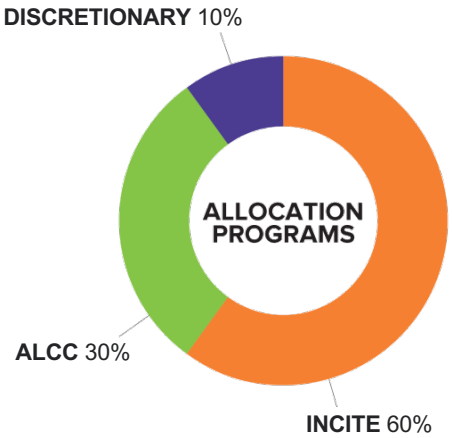

THETA

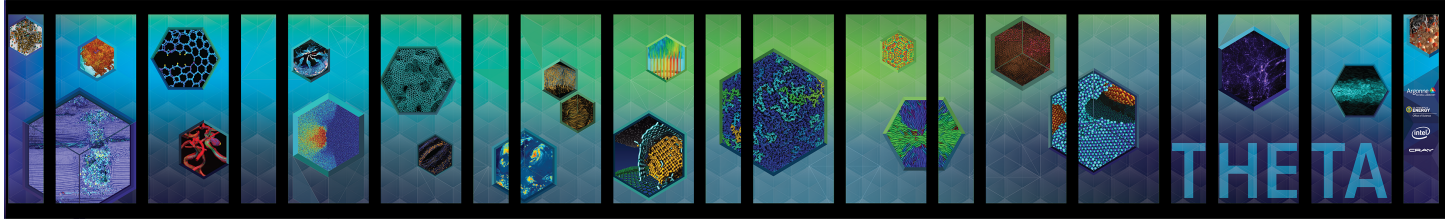

Pre-exascale system: Simulation, Data, and Learning

- Production July 1, 2017
- Serves as a bridge between Mira and Aurora, transition and data analytics system
- Cray XC40 system. Runs Cray software stack
- 11.69 PF peak performance
- 2nd Generation Intel® Xeon Phi™ processor
  - Knights Landing, 7230 SKU 64 cores 1.3GHz
- 24 racks
- 4,392 nodes
- 281,088 cores
- 70.27.TB MCDRAM
- 843.26 TB DDR4
- 562.17 TB SSD
- Cray Aries high speed interconnect in dragonfly topology
- 10PB Lustre file system, 200 GB/s throughput

ALCF PRODUCTION RESOURCES

Mira – IBM BG/Q

- 49,152 nodes
- 786,432 cores
- 786 TB RAM
- 10 PF

Cetus – IBM BG/Q

- 4,096 nodes
- 65,536 cores
- 64 TB RAM
- 836 TF

Vesta – IBM BG/Q

- 2,048 nodes
- 32,768 cores
- 32 TB RAM
- 419 TF

COOLEY

- Analytics/Visualization cluster
- Peak 223 TF
- 126 nodes; each node has
  - 2 -Intel Xeon E5-2620 Haswell 2.4 GHz 6-core processors
  - NVIDIA Tesla K80 graphics processing unit (24GB)
  - **384 GB of RAM**
- Aggregate RAM of 47 TB
- Aggregate GPU memory of ~3TB
- Cray CS System
- 216 port FDR IB switch with uplinks to our QDR infrastructure
- Mounts the same GPFS file systems as Mira, Cetus

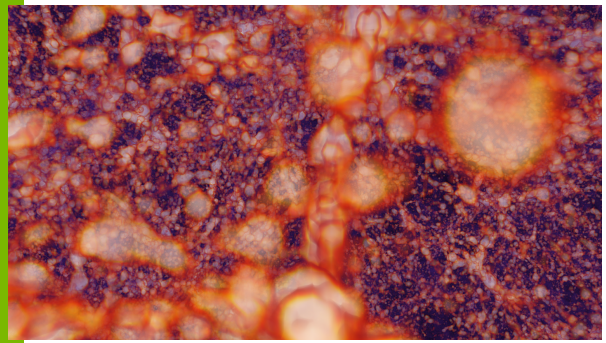

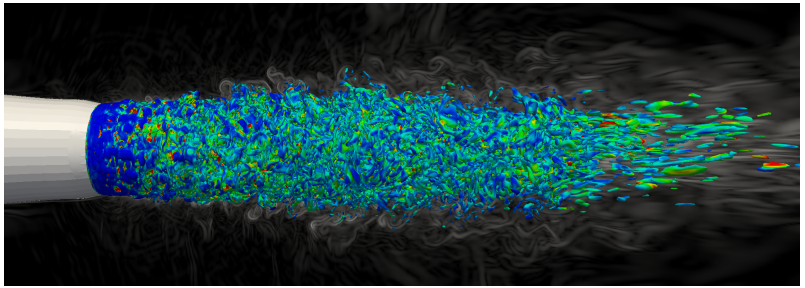

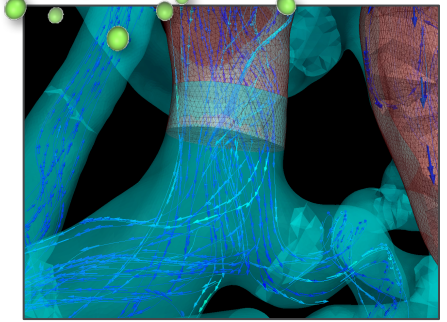

AURORA 2021

Argonne's Future Exascale Supercomputer

- One quintillion (10^{18}) calculations per second.

Expected Architectural Features:

- More than 50,000 nodes
- Nodes will have high single-thread performance
- Nodes will enable exceptional performance with codes that have concurrency at a modest scale
- More than 5 petabytes of memory
- All memory will be high performance
- Memory in a node will be coherent
- Compute hardware will have equal access to all resources (memory, fabric, etc.)

LARGE-SCALE COMPUTING AND VISUALIZATION ON THE CONNECTOMES OF THE BRAIN

Objectives:

- development of imaging and analytical pipelines for full mammalian brains at the level of individual cells, axons and blood vessels
- integration on large-scale computing systems

Imaging technique:

- X-Ray extended tomography (or Mosaic Tomography) with 1micron resolution done at the beamline 32-ID-C on the Advanced Photon Source

Segmentation :

- Tensor flow based segmentation to extract features like cell bodies, myelinated axons and blood vessels

Science: Narayanan (Bobby) Kasthuri and team
Slide courtesy Rafael Vescovi, Hanyu Li

Argonne NATIONAL LABORATORY

RESOURCE ACCESS

Theta and Cooley

- See Haritha for help with OTP tokens
- www.alcf.anl.gov/user-guides
  - Onboarding Guide
- Login to:
  - theta.alcf.anl.gov
  - cooley.alcf.anl.gov
- Storage:
  - Home directory
    - Code, etc.
  - /projects/InsituVis2018
    - Shared code
    - Data
- Reservations:
  - July 10, 2:00 – 6:00pm
    - Theta: 32 nodes
    - Cooley: 32 nodes
  - July 11, 9:00am – 9:00pm
    - Theta: 64 nodes
    - Cooley: 32 nodes
  - July 12, 9:00am – 2:00pm
    - Theta: 64 nodes
    - Cooley: 32 nodes
  - Submit with –q R.InSituVis2018