TACC: SITE UPDATE

John Cazes
Director, High Performance Computing
The University of Texas at Austin

IXPUG, October 2020

WHAT IS TACC?

Grendel, 1993

The Texas Advanced Computing Center, at UT Austin is a (primarily) NSF-funded center to provide and apply large scale computing resources to the open science community.

Frontera, 2019

TACC AT A GLANCE - 2020

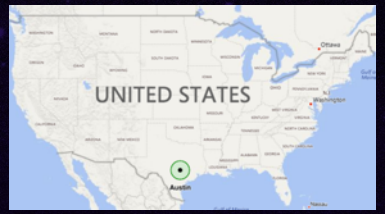

Personnel

185 Staff (~70 PhD)

Facilities

12 MW Data center capacity
Two office buildings, Three Datacenters, two visualization facilities, and a chilling plant.

Systems and Services

>Seven Billion compute hours per year
>5 Billion files, >100 Petabytes of Data, NSF Frontera (Track 1), Stampede2 (XSEDE Flagship), Jetstream (Cloud), Chameleon (Cloud Testbed) system

Usage

>15,000 direct users in >4,000 projects, >50,000 web/portal users, User demand 8x available system time.
Thousands of training/outreach participants annually

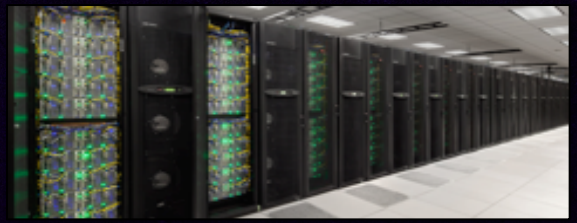

TACC ECOSYSTEM

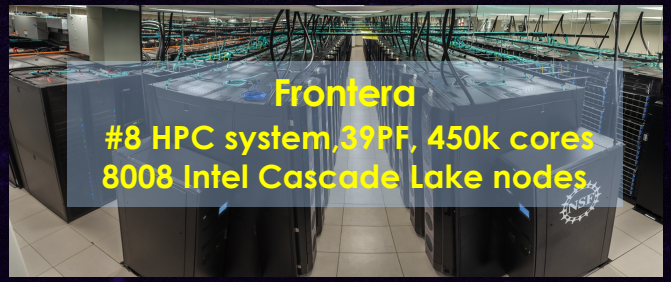

Frontera
#8 HPC system,39PF, 450k cores
8008 Intel Cascade Lake nodes

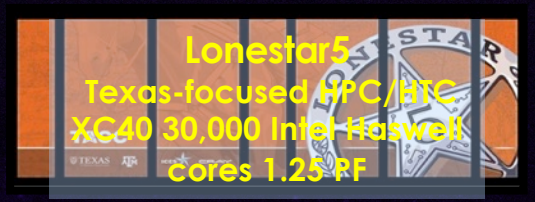

Lonestar5
Texas-focused HPC/HTC
XC40 30,000 Intel Haswell cores 1.25 PF

Longhorn
GPU and ML
96 Power 9 nodes
384 NVIDIA V100s
2.8 PF

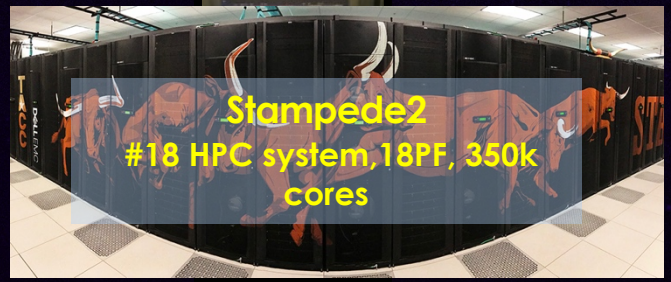

Stampede2
#18 HPC system,18PF, 350k cores

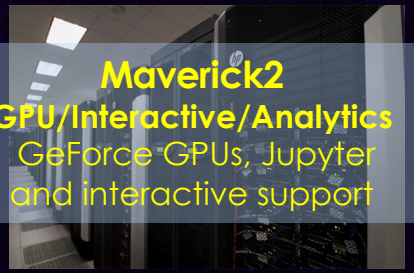

Maverick2
GPU/Interactive/Analytics
GeForce GPUs, Jupyter and interactive support

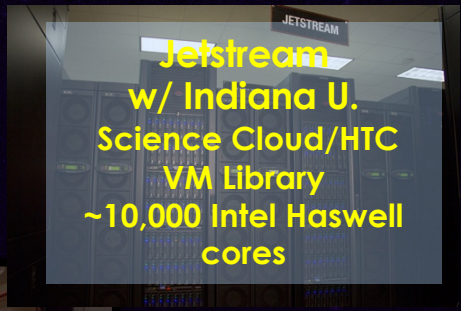

Jetstream
w/ Indiana U.
Science Cloud/HTC
VM Library
~10,000 Intel Haswell cores

Stockyard
Shared Storage Across TACC
30PB, Lustre

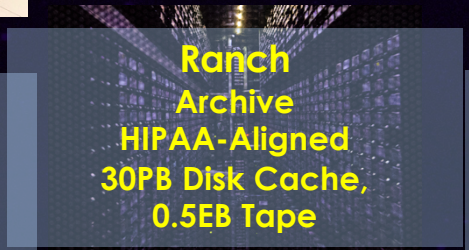

Ranch
Archive
HIPAA-Aligned
30PB Disk Cache,
0.5EB Tape

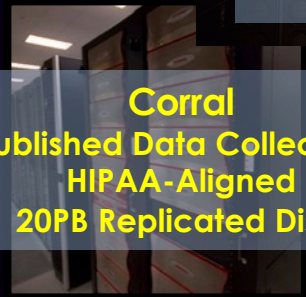

Corral
Published Data Collections
HIPAA-Aligned
20PB Replicated Disk,

Stallion

Lasso

CORONAVIRUS RESPONSE
COVID19 HPC CONSORTIUM

The first phase of response is always Pandemic Modeling and Supplies Response:
Lauren Meyers + CDC

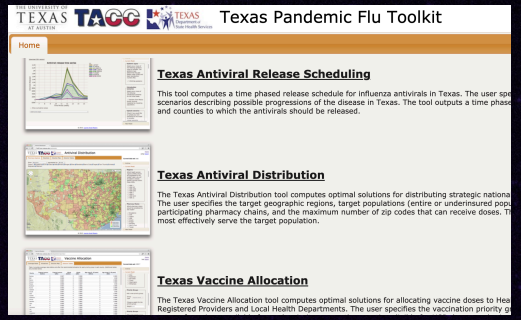

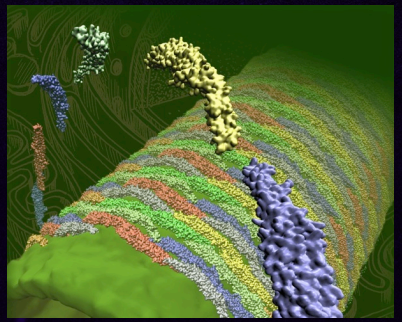

- *Then protein folding/structure*
- Researchers from the University of Illinois at Champaign-Urbana used **Stampede** to unravel how a newborn protein folds and to design novel enzymes. On **Stampede**, scientists can perform simulations to observe long biological events that were hard to access before, and to relate molecular structures and interactions with their biological functionalities.
- Voth, U. Chicago and Aksimentiev, Illinois.

In each case, we start fast because we have years-long working relationships with the scientists

Then Whole Viron Simulation
Rommie Amaro (UCSD), et al. studying whole-viron models of influenza – Retooling for CoVID now (on **Frontera**)

Modeling surface protein "stalks" which control interactions with cells and drugs

200M atom, all-atom, fully solvated model.

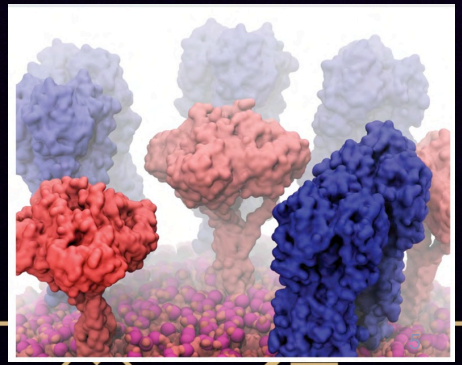

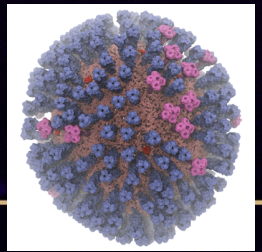

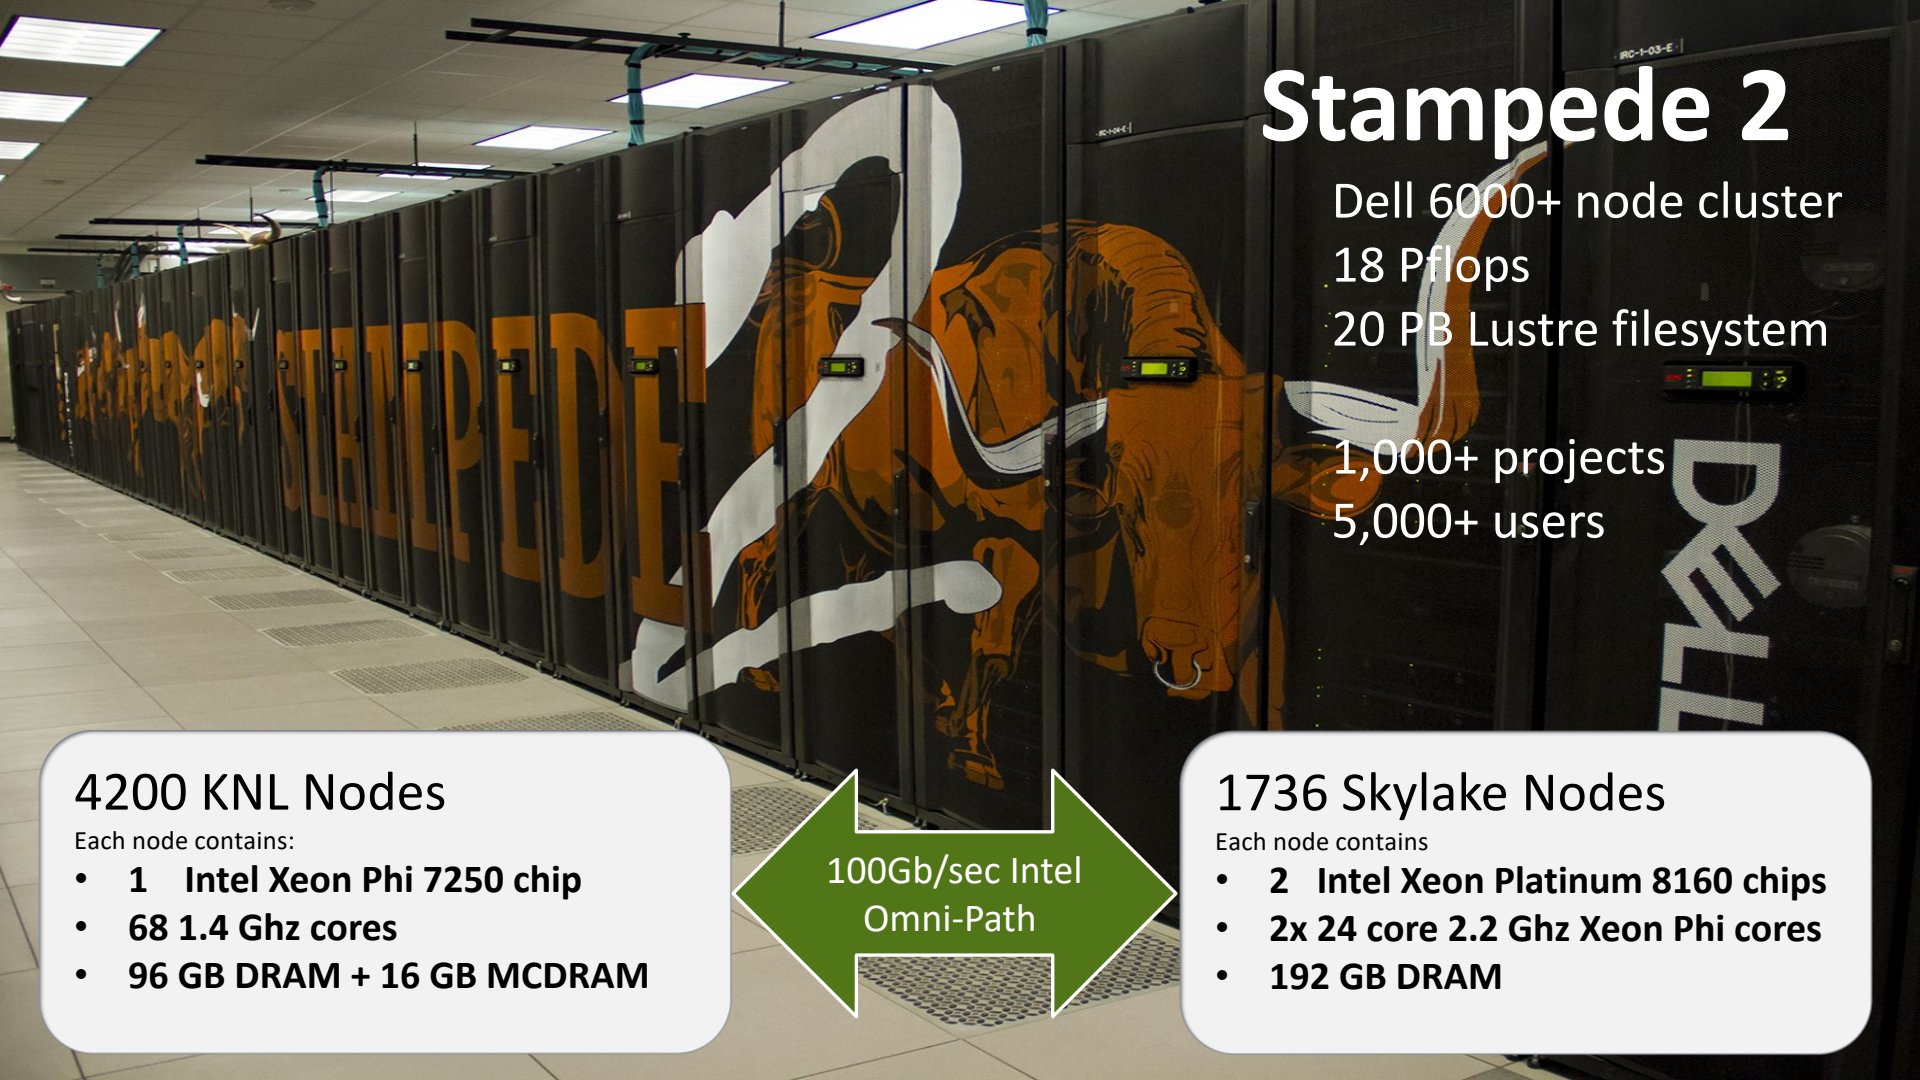

Stampede 2

Dell 6000+ node cluster
18 Pflops
20 PB Lustre filesystem

1,000+ projects
5,000+ users

4200 KNL Nodes

Each node contains:

- **1 Intel Xeon Phi 7250 chip**
- **68 1.4 Ghz cores**
- **96 GB DRAM + 16 GB MCDRAM**

100Gb/sec Intel Omni-Path

1736 Skylake Nodes

Each node contains

- **2 Intel Xeon Platinum 8160 chips**
- **2x 24 core 2.2 Ghz Xeon Phi cores**
- **192 GB DRAM**

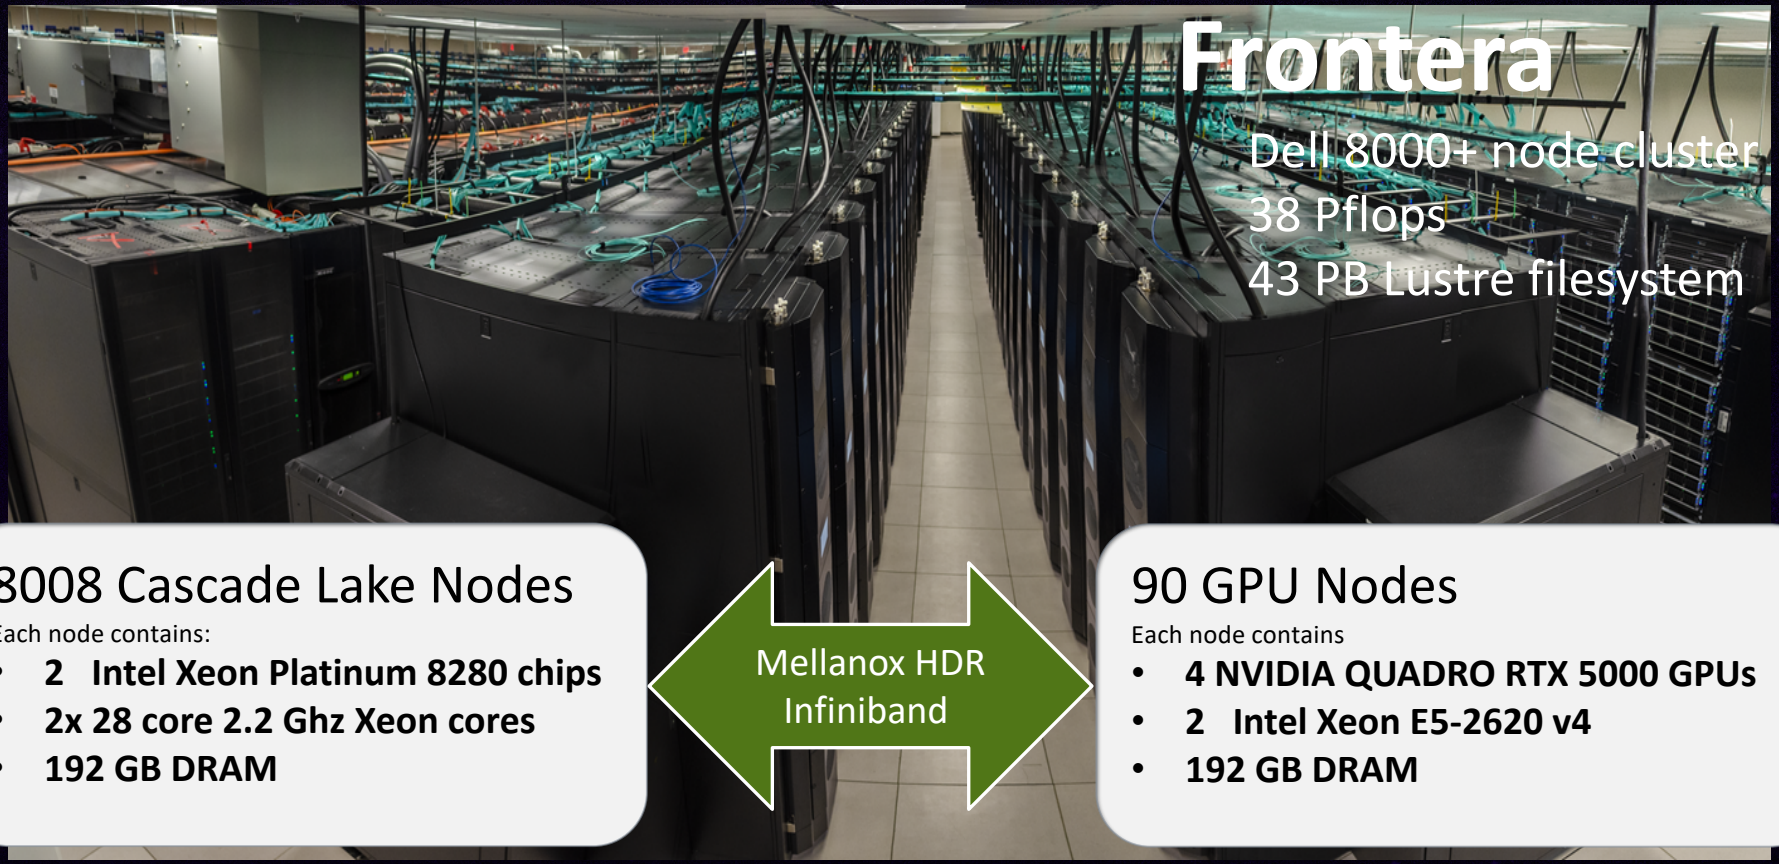

Frontera

Dell 8000+ node cluster

38 Pflops

43 PB Lustre filesystem

8008 Cascade Lake Nodes

Each node contains:

- **2 Intel Xeon Platinum 8280 chips**
- **2x 28 core 2.2 Ghz Xeon cores**
- **192 GB DRAM**

Mellanox HDR Infiniband

90 GPU Nodes

Each node contains

- **4 NVIDIA QUADRO RTX 5000 GPUs**
- **2 Intel Xeon E5-2620 v4**
- **192 GB DRAM**

FRONTERA SYSTEM --- INFRASTRUCTURE

- Frontera consumes almost 6 Megawatts of Power at Peak
- Direct water cooling of primary compute racks (CoolIT/DellEMC)
- Oil immersion Cooling (GRC)
- Solar, Wind inputs.
- Frontera is 5x more efficient in Ops/Watt than the system it replaced.
- Next system will be likely 3x this (for 10x capability).

TACC Machine Room Chilled Water Plant

FRONTERA TODAY (BRAGGING SLIDE)

- Frontera is the #8 ranked system in the world – and the fastest at any university in the world.
- Fastest primarily Intel-based system
- Highest ranked Dell system ever.
- Frontera and Stampede2 are #1 and #2 among US Universities

FRONTERA – UNIQUE FEATURES

- RTX queue
  - 90 nodes
  - 4 NVIDIA Quadro 5000 RTX cards per node
  - 16 GB and 11 TFlop single-precision performance per card
- NVDIMM queue
  - 2 TB Memory per node
  - 4 local 833 GB partitions per node (3.2 TB local storage)
  - 112 Cascade Lake cores per node

FIELDS OF SCIENCE

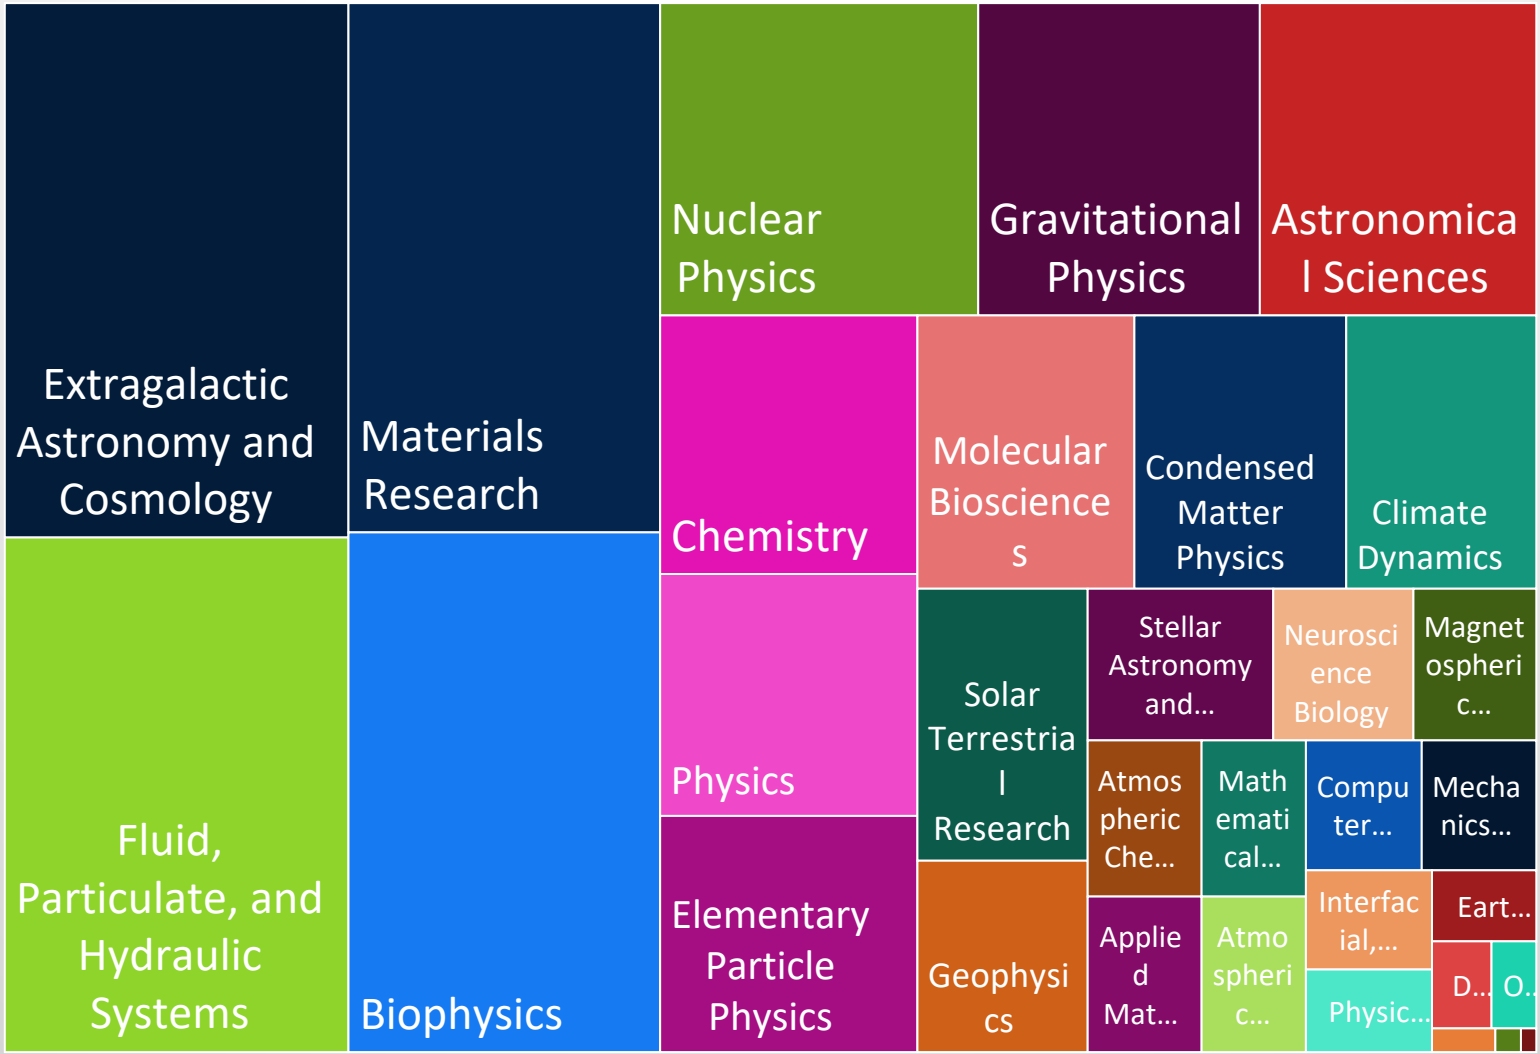

TACC TEXAS

TEXASCALE DAYS

Frontera dedicated to large scale runs from half to whole system

- Node counts from 3800 – 7900 nodes
- Schedule after a maintenance to avoid draining queues
- Projects must have successfully executed on 2048 nodes
- Each project has 24 hours on half-system(~3900 nodes) or full-system(~7900 nodes)
- Completed 3 series in October 2019, March 2020, and September 2020
- Supported 14 projects
- Liaison directly with projects via slack to work through initial issues

3-D STELLAR HYDRODYNAMICS

PAUL WOODWARD
UNIVERSITY OF MINNESOTA

- The project's goal is to study the process of Convective Boundary Mixing (CBM) and shell mergers in massive stars.
- The computational plan includes a sequence of brief three-dimensional simulations alternating with longer one-dimensional simulations.
- Ran on 7,300+ nodes for more than 80 hours during Frontera large-scale capability demonstration.
- Saw 588 GFlop/s/node — or 4 Petaflops of sustained performance — for more than 3 days!

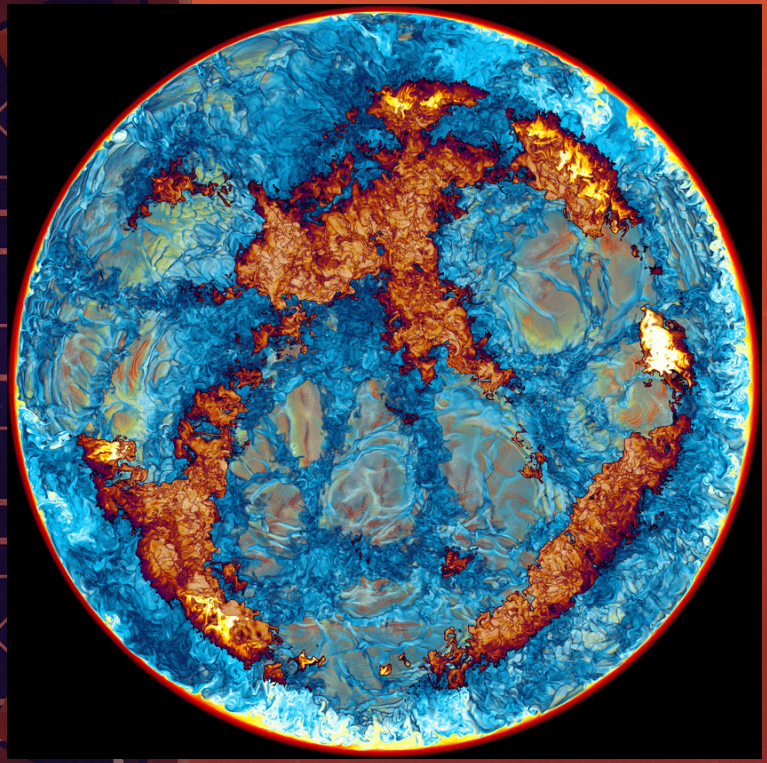

PREDICTION AND CONTROL OF TURBULENCE-GENERATED SOUND

DANIEL BODONY
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

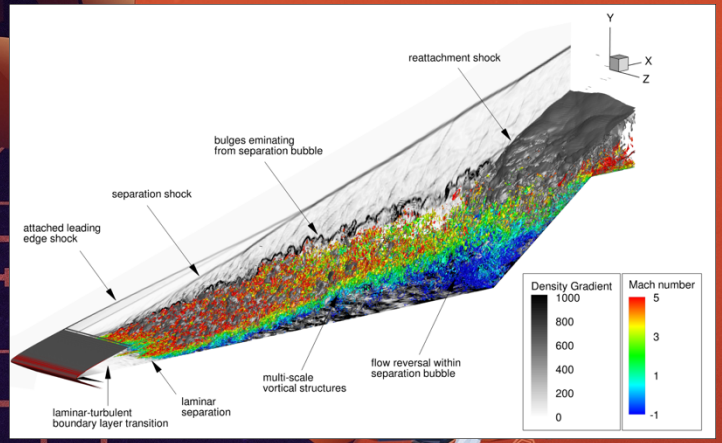

- Simulated fluid-structure interactions relevant to hypersonic vehicles.
- Simulations replicated a companion experiment performed at NASA Langley in their 20-inch Mach 6 tunnel.
- Frontera runs used 2 MPI ranks per node (one per socket) and 26 OpenMP threads per MPI rank.
- Saw superlinear speedup on up to 2,000+ nodes by fitting into cache rather than fetching from main memory.
- Linear speedup up to 4,000 nodes.

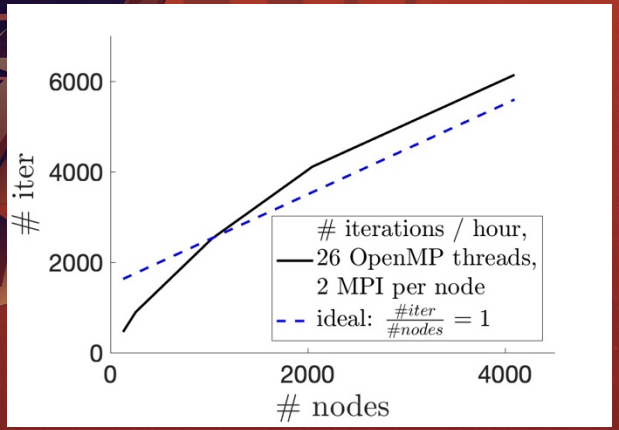

CENTER FOR THE PHYSICS OF LIVING CELLS

ALEKSEI AKSIMENTIEV
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

- The nuclear pore complex serves as a gatekeeper, regulating the transport of biomolecules in and out of the nucleus of a biological cell.
- To uncover the mechanism of such selective transport, the Aksimentiev lab at UIUC constructed a computational model of the complex.
- The team simulated the model using memory-optimized NAMD 2.13, 8tasks/node, MPI+SMP.
- Ran on up to 7,780 nodes on Frontera.
- One of the largest biomolecular simulations ever performed.
- Scaled close to linear on up to half of the machine.
- Plan to build a new system twice as large to take advantage of large runs

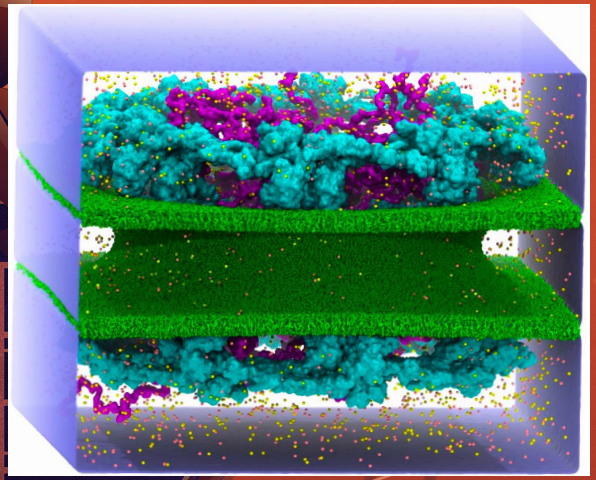

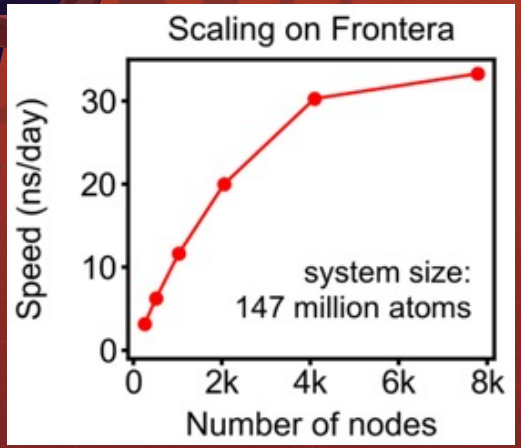

FRONTIERS OF COARSE-GRAINING

GREGORY VOTH
UNIVERSITY OF CHICAGO

- Mature HIV-1 capsid proteins self-assemble into large fullerene-cone structures.
- These capsids enclose the infective genetic material of the virus and transport viral DNA from virion particles into the nucleus of newly infected cells.
- On Frontera, Voth's team simulated a viral capsids containing RNA and stabilizing cellular factors in full atomic detail for over 500 ns.
- First molecular simulations of HIV capsids that contain biological components of the virus within the capsid.
- The team ran on 4,000 nodes on Frontera.
- Measured the response of the capsid to molecular components such as including genetic cargo and cellular factors that affect the stability of the capsid.

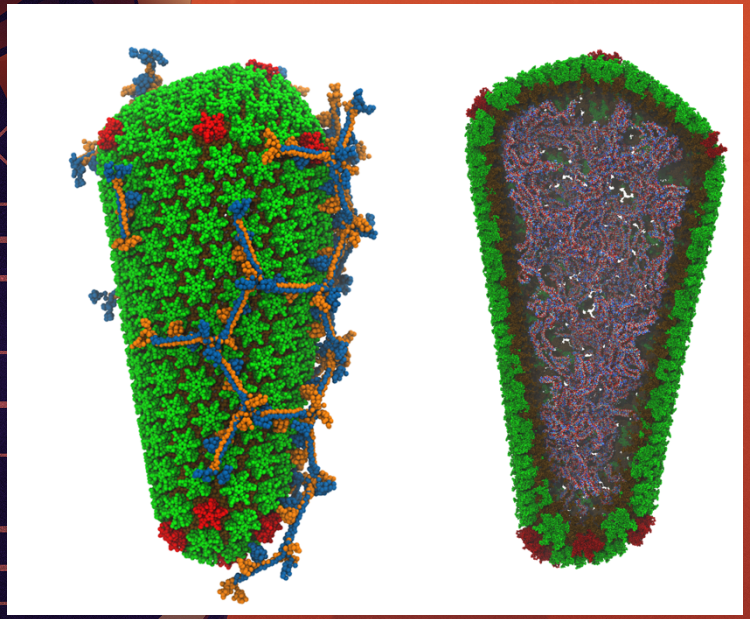

"State-of-the-art supercomputing resources like Frontera are an invaluable resource for researchers. Molecular processes that determine the chemistry of life are often interconnected and difficult to probe in isolation. Frontera enables large-scale simulations that examine these processes, and this type of science simply cannot be performed on smaller supercomputing resources."

-Alvin Yu, Postdoctoral Scholar in Voth Group

cazes@tacc.utexas.edu