From Supercomputers to GPUs

What a physicist should know about current computational capabilities

Craig Rasmussen (Research Support Services, University of Oregon)

Which one?

Gordon Bell Prize: Price Performance

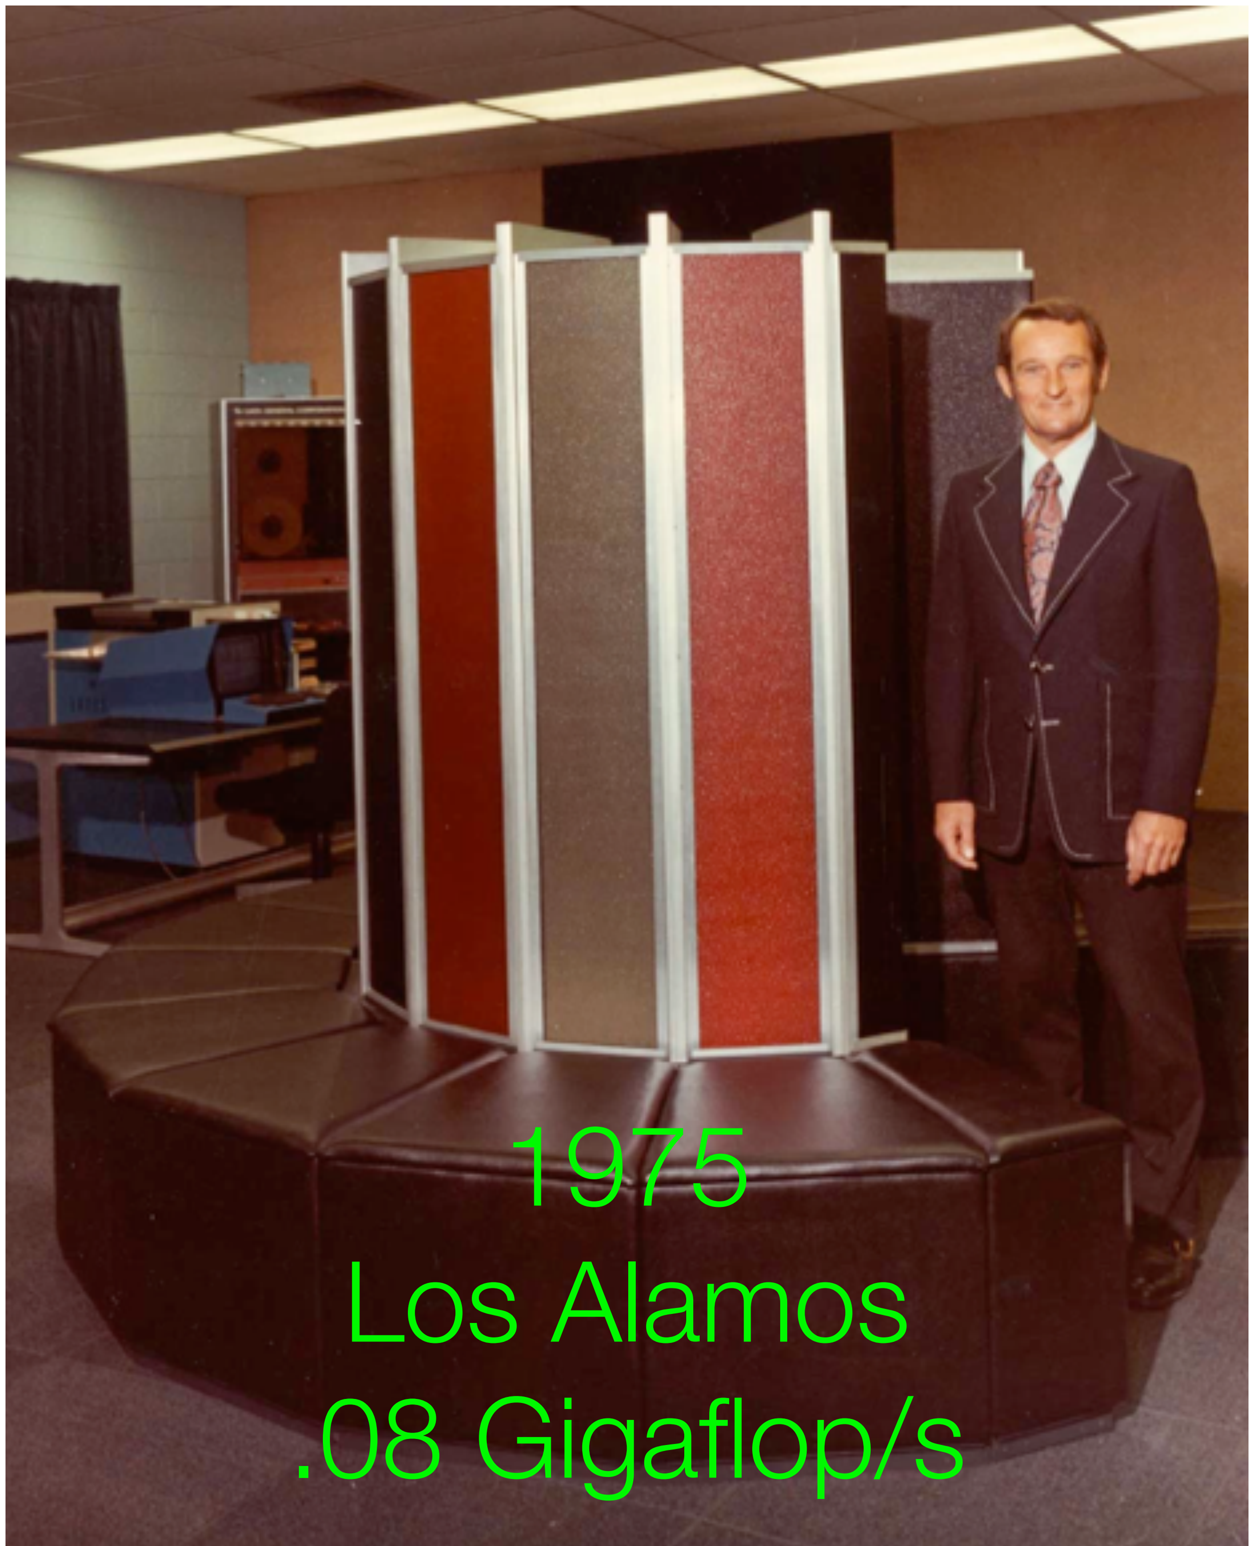

Year	Gigaflop/s Performance
1988	1
1998	1,020
2008	1,350,000

Year	Gigaflop/s Cost
1989	$2,500,000
1999	$6,900
2009	$8

The range and scale of computational resources is breathtaking

Mega	Giga	Tera	Peta	Exa
10^6	10^9	10^{12}	10^{15}	10^{18}
Displays	Networks & Storage bandwidths	Operations per second	Operations per second	Operations per second

- Notice the discrepancy between what you can see (display), store, and compute! *(should have referenced source)*

I just reviewed a proposal requesting 180,000,000 core hours of computer time for the DOE

- Interested in the physics?

I just reviewed a proposal requesting 180,000,000 core hours of computer time for the DOE

- Nuclear Physics
  - Domain Wall Fermions and Highly Improved Staggered Quarks for Lattice QCD
  - Chroma Lattice QCD Code Suite
  - Weakly Bound and Resonant States in Light Isotope Chains Using MFDn
- High Energy Physics
  - Hardware/Hybrid Accelerated Cosmology Code for Extreme Scale Cosmology
  - The MILC Code Suite for Quantum Chromodynamics (QCD) Simulation and Analysis
  - Advanced Modeling of Particle Accelerators

180,000,000 core hours enables a lot of science

- Fusion Energy Physics
  - Understanding Fusion Edge Physics Using the Global Gyrokinetic XGC1 Code
  - Addressing Non-Ideal Fusion Plasma Magnetohydrodynamis Using M3D-C1
- Basic Energy Science
  - Parsec: A Scalable Computational Tool for Discovery and Design of Excited State Phenomena in Energy Materials
  - BerkeleyGW: Massively Parallel Quasiparticle and Optical Properties Computation for Materials and Nanostructures
  - Materials Science using Quantum Espresso
  - Large-Scale 3-D Geophysical Inverse Modeling of the Earth
  - Large-Scale Molecular Simulations with NWChem

180,000,000 core hours seems like a lot!

- Advanced Scientific Computing Research
  - Optimization of the BoxLib Adaptive Mesh Refinement Framework . . .
  - High-Resolution CFD and Transport in Complex Geometries . . .
- Biological and Environmental Research
  - CESM Global Climate Modeling
  - Multi-Scale Ocean Simulation for Studying Global to Regional Climate Change
  - Gromacs Molecular Dynamics (MD) Simulation for Bioenergy and Environmental Biosciences
  - Meraculous: a Production de novo Genome Assembler for Energy-Related Genomics Problems

Who are the players?

- Lawrence Berkeley National Laboratory (6)
- Oak Ridge National Laboratory (2)
- Argonne, Pacific Northwest, Los Alamos National Laboratories
- Brookhaven National Laboratory
- Jefferson National Accelerator Facility
- National Center for Atmospheric Research
- Princeton (2)
- Universities of Texas, Arizona, Iowa State

CESM Global Climate Modeling: *Increase in the state-level extreme precipitation by 2050s*

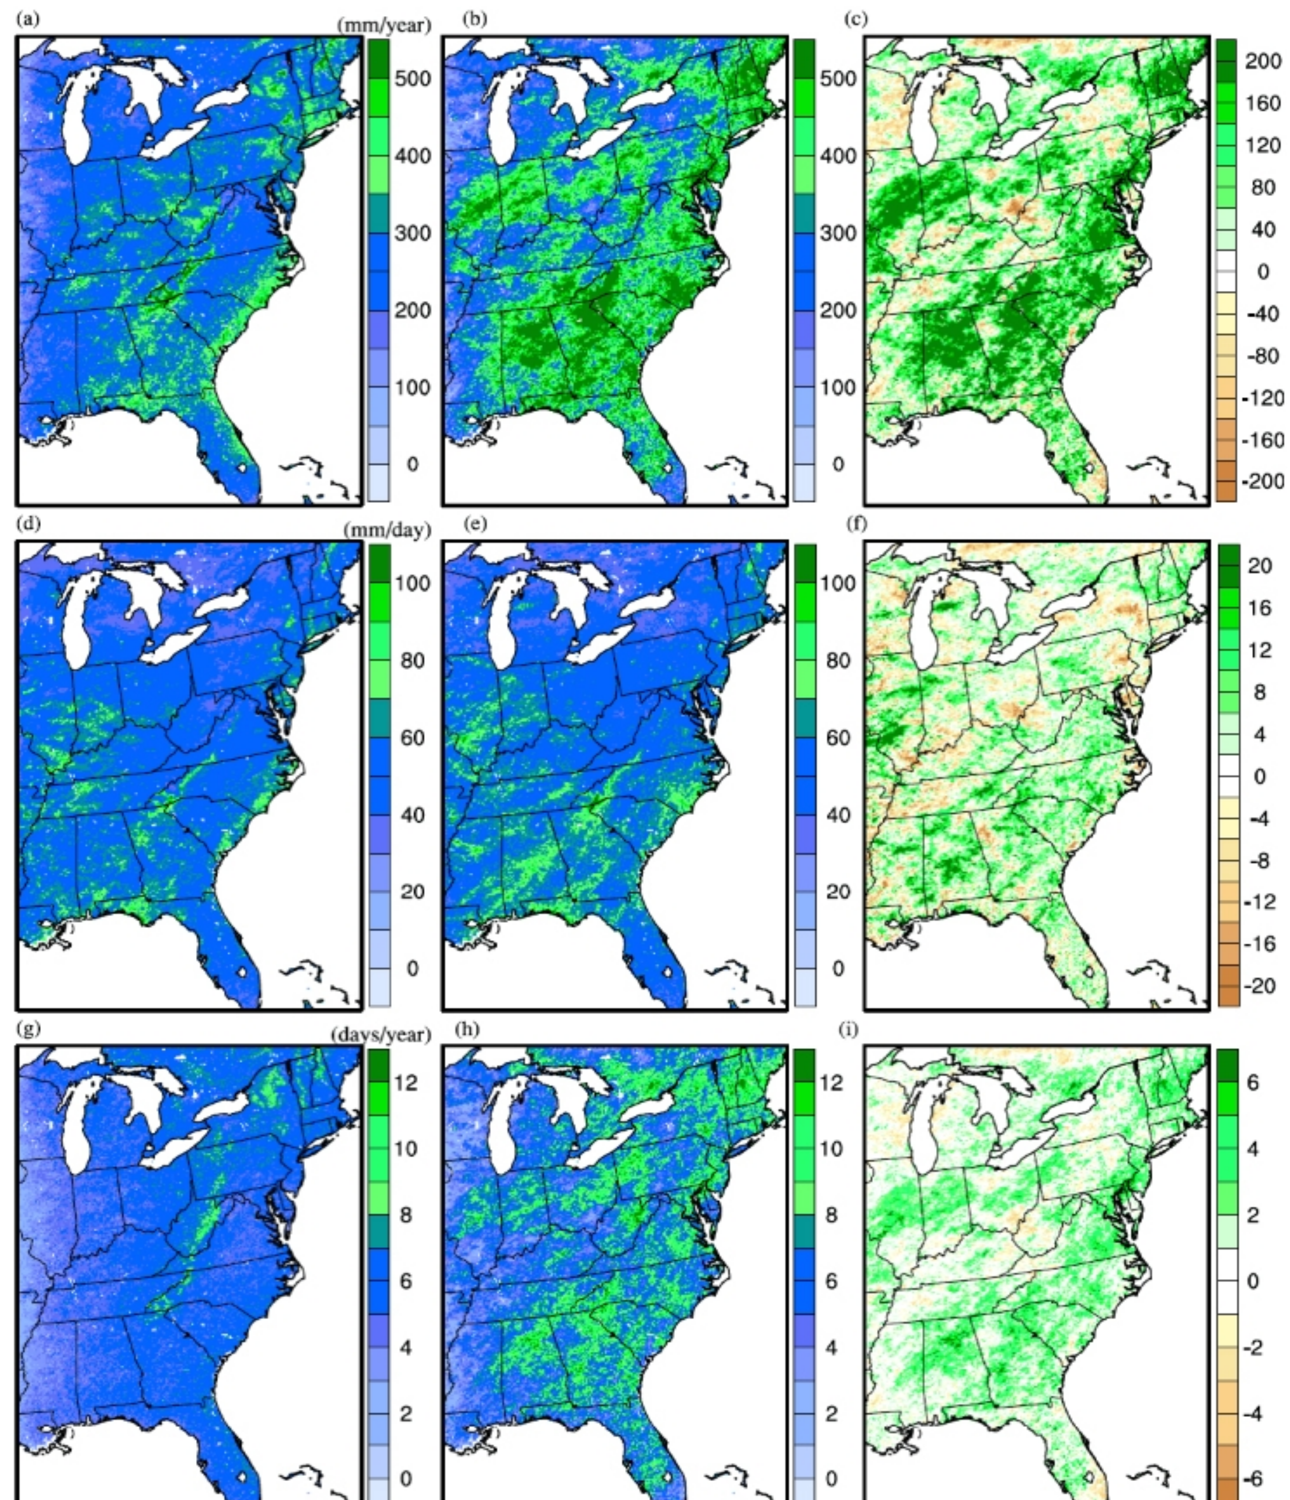

The range and scale of computational resources is breathtaking

Mega	Giga	Tera	Peta	Exa
10^6	10^9	10^{12}	10^{15}	10^{18}
Displays	Networks & Storage bandwidths	Operations per second	Operations per second	Operations per second

- But what is the limit?

What’s the limit?

- Frequency?
- Transistors?

Heat is a concern!

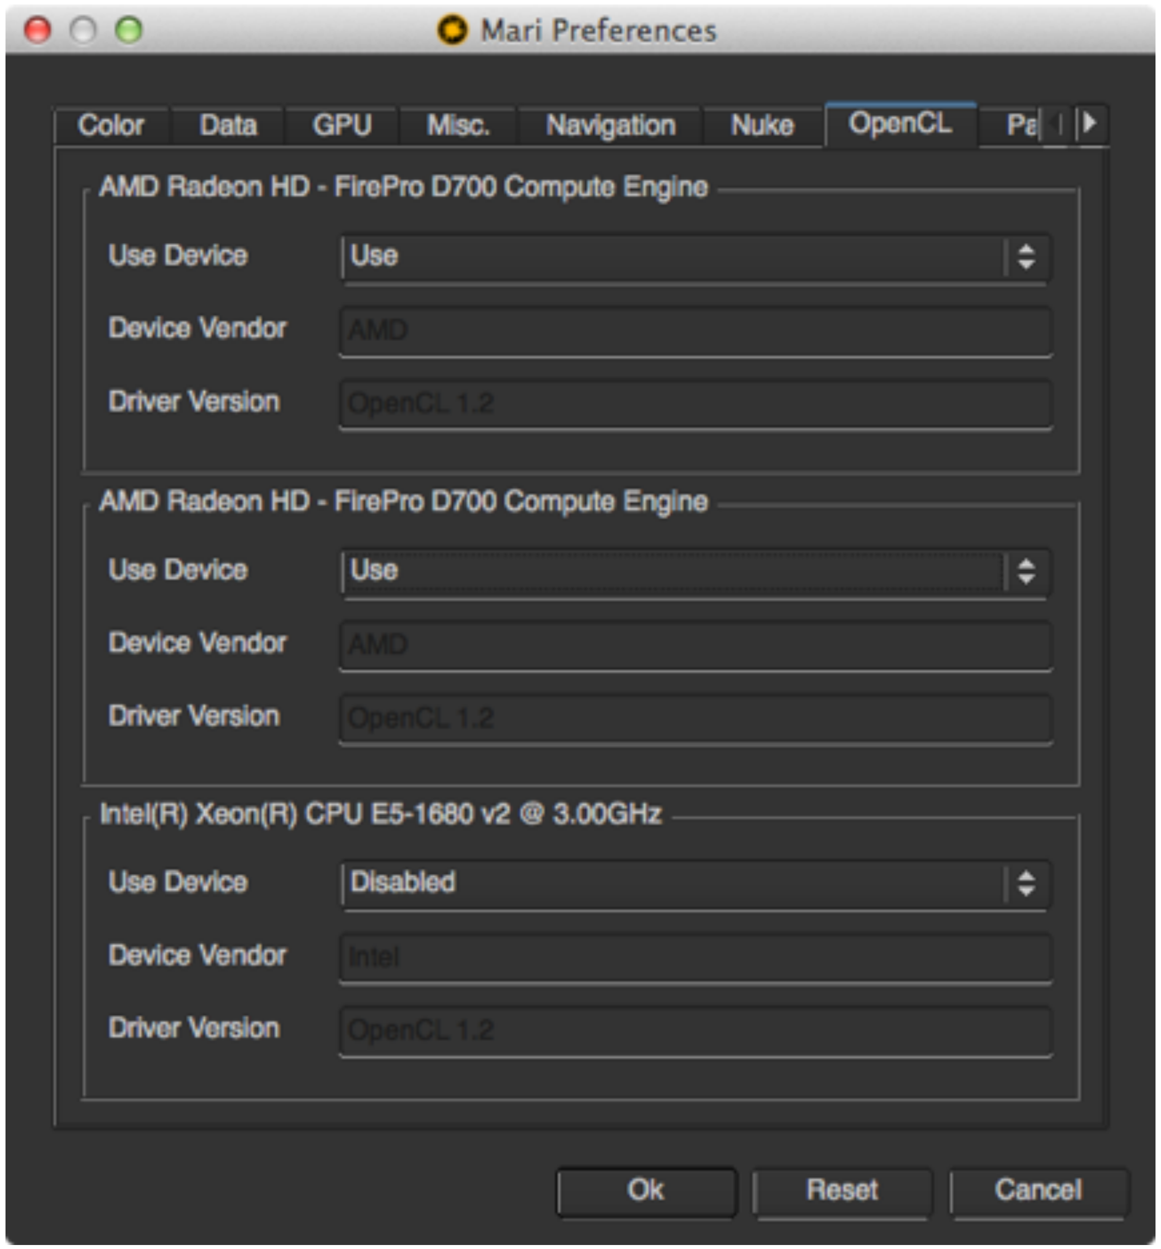

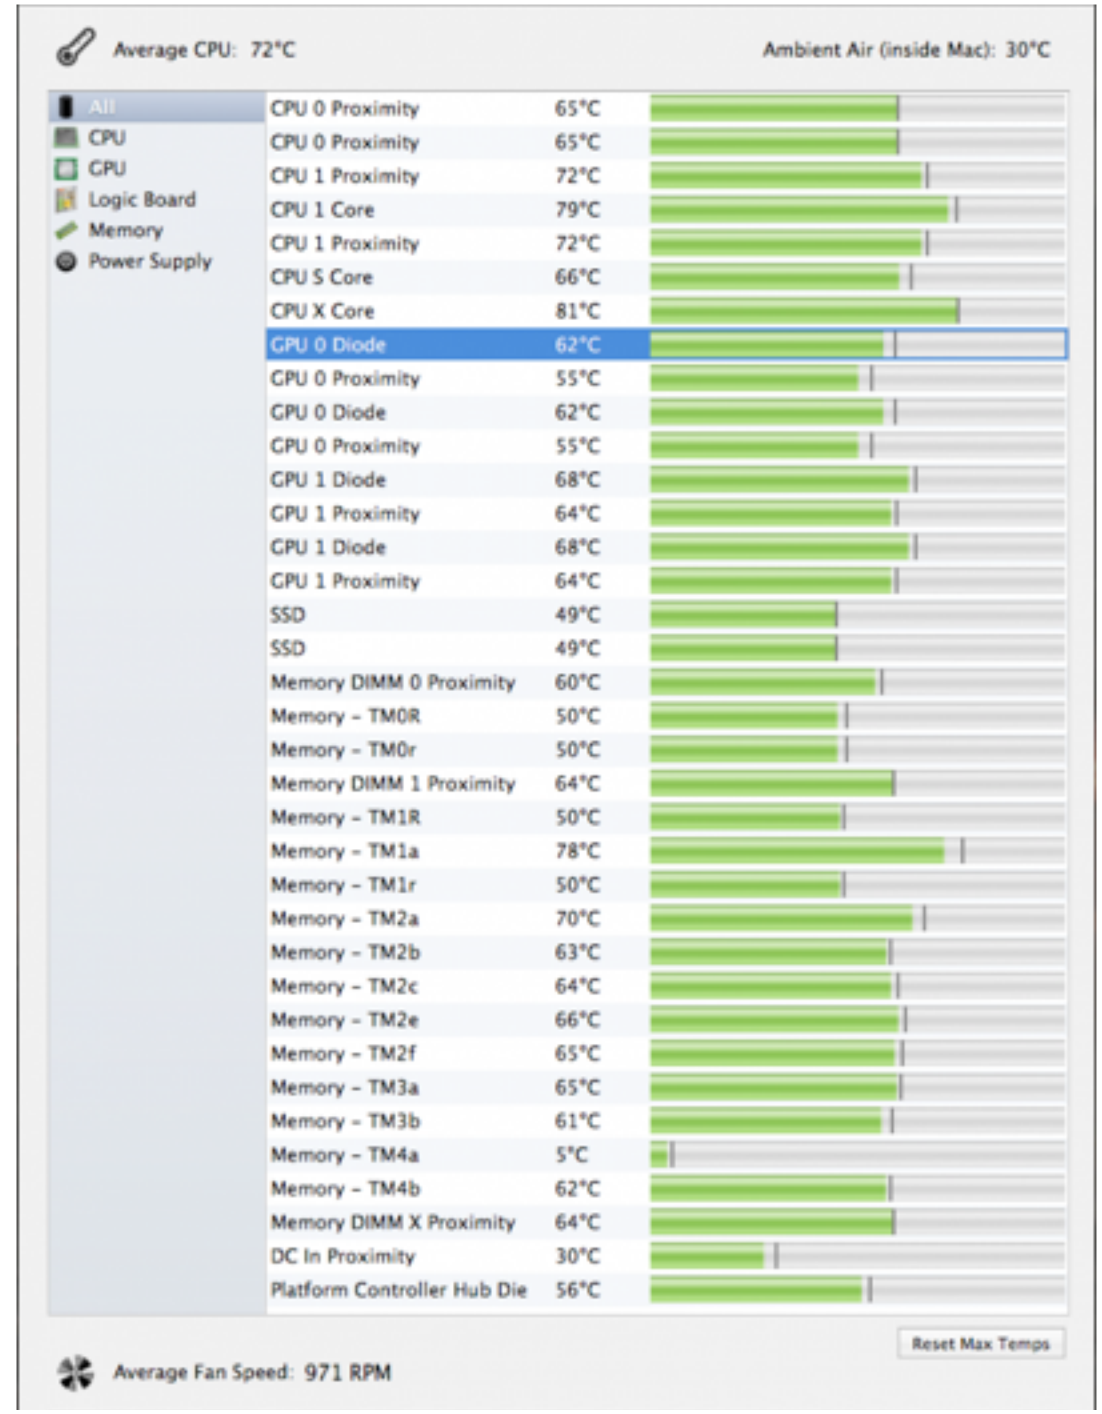

When’s the limit?

- Frequency
- Transistors

When did the music die?

I gotta retire by 2020!

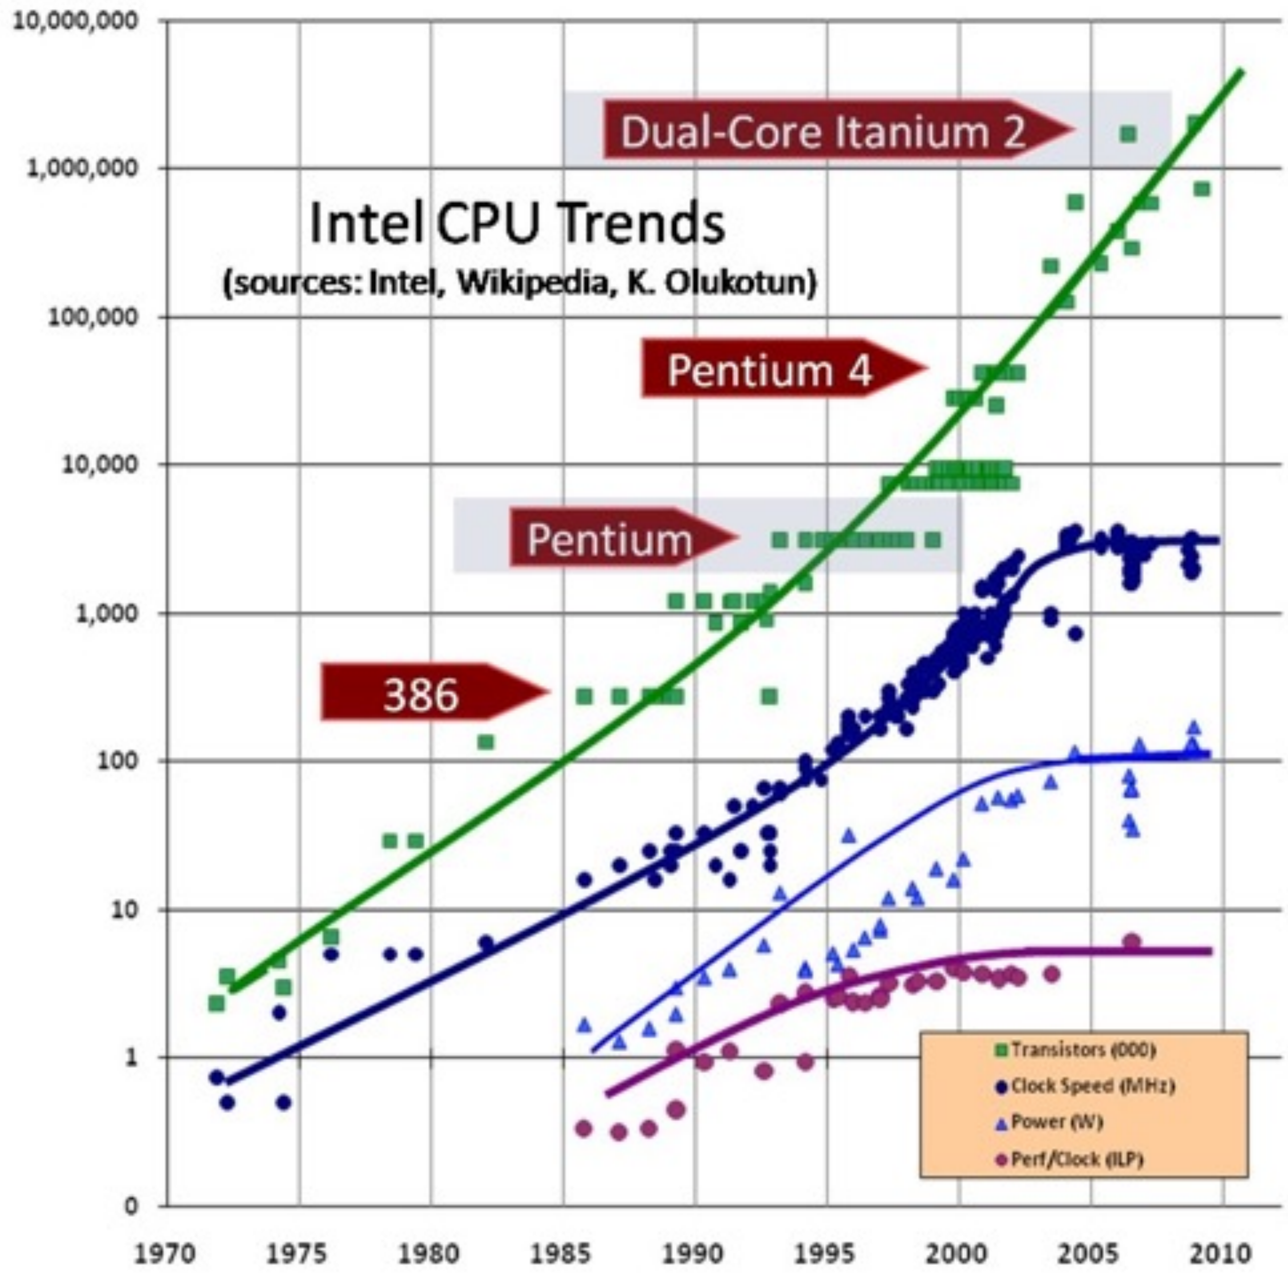

Machines are getting really really really big

- What can we dream about doing with them?
- Think big
  - because if we don't think big we can be assured that someone else is
  - need to stay competitive

SciDAC: Scientific Discovery through Advanced Computing

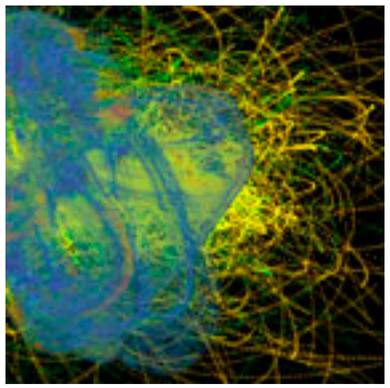

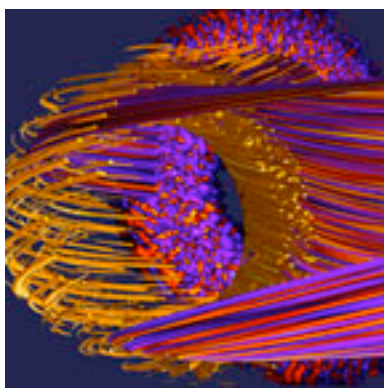

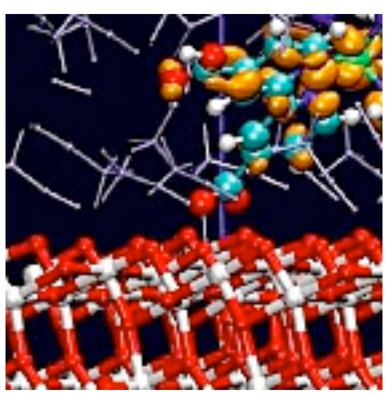

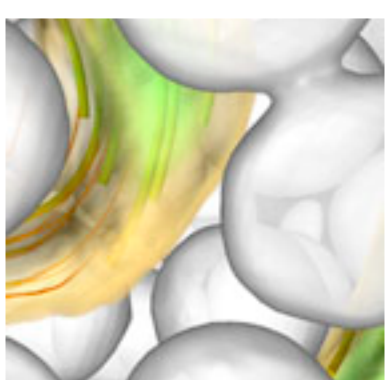

- Program areas include:
  - Climate science, fusion research, high energy physics, nuclear physics, astrophysics, material science, chemistry, particle accelerators, biology, and subsurface flow of contaminants through groundwater
- Analgren - From Convection to Explosion: End-to-End Simulation of Type Ia Supernovae
- Berry - Parallelization in Time: Applications to Plasma Turbulence
- Drut - From Lattice QCD to Ultracold Atoms and Graphene: accelerating the Monte Carlo approach to many-fermion physics
- Ferraro - Fluid Modeling of Fusion Plasmas with M3D-C1
- Iacovella - Flexible order parameters for quantifying the rate-dependent energy release mechanism of Au nanowires
- Jones - MHD Turbulence in a Cosmic Structure Context

VisIT Example (Dave Pugmire, Hank Childs, ...)

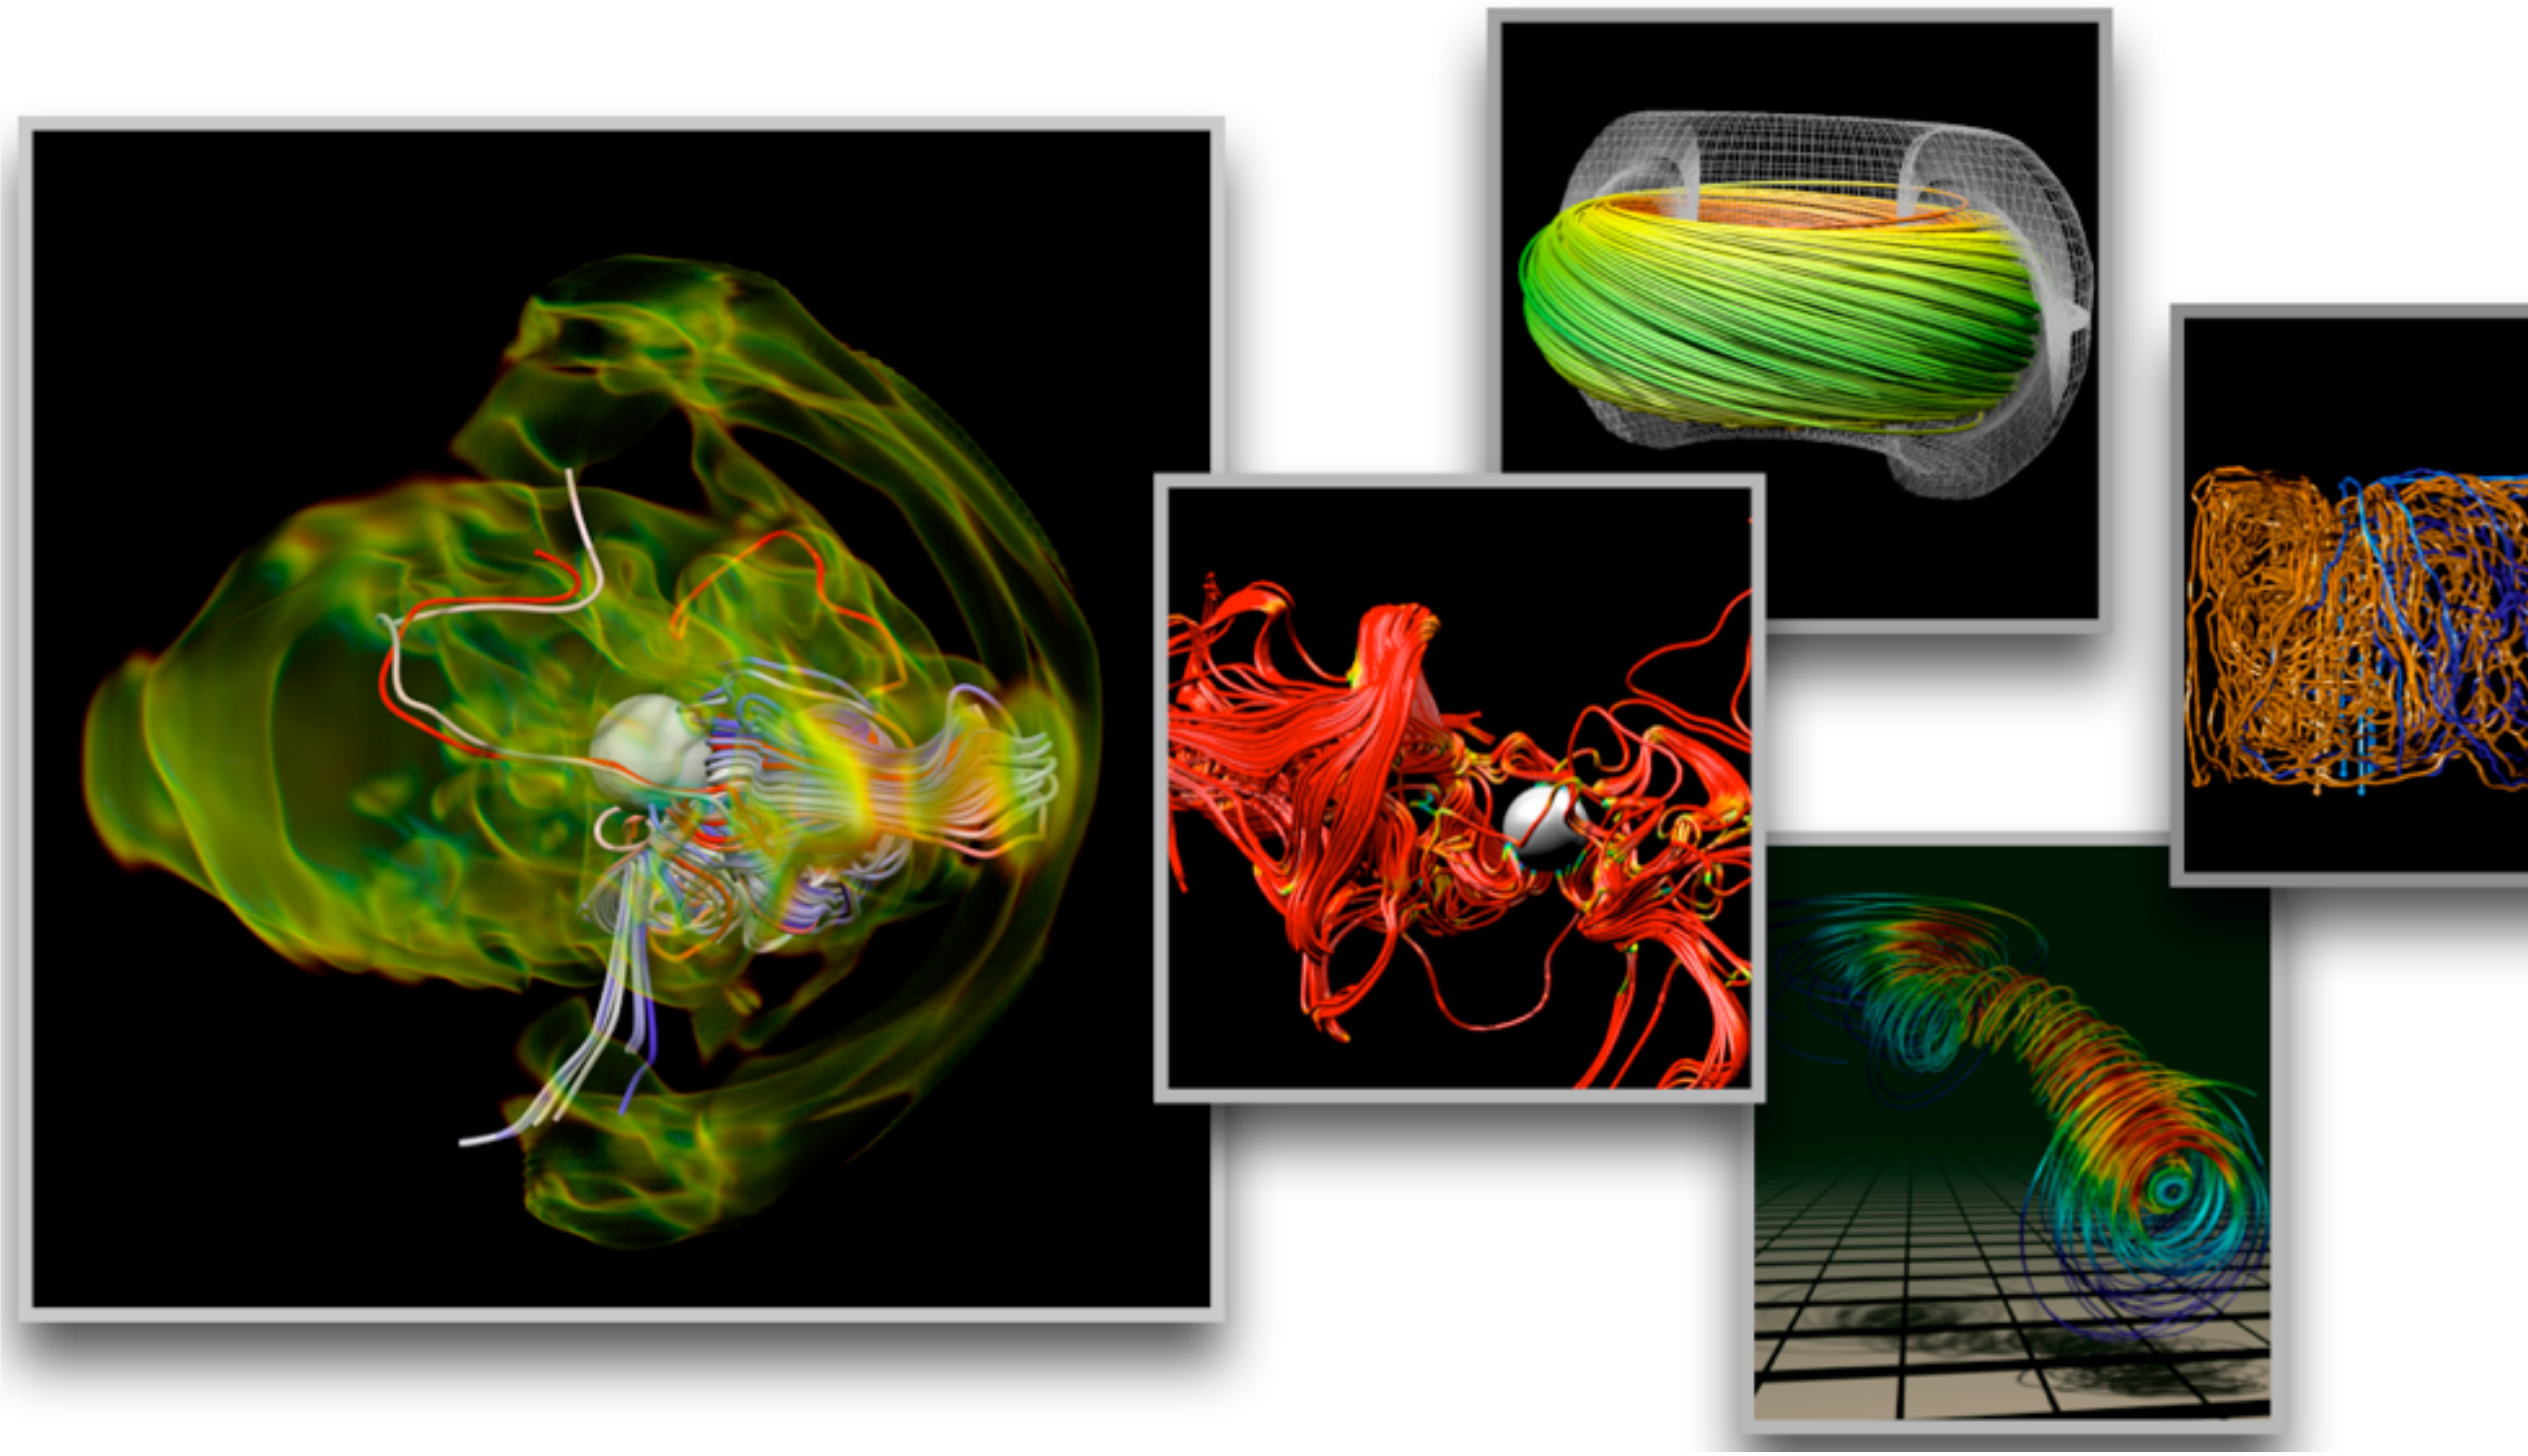

ParaView
Examples
Astrophysics

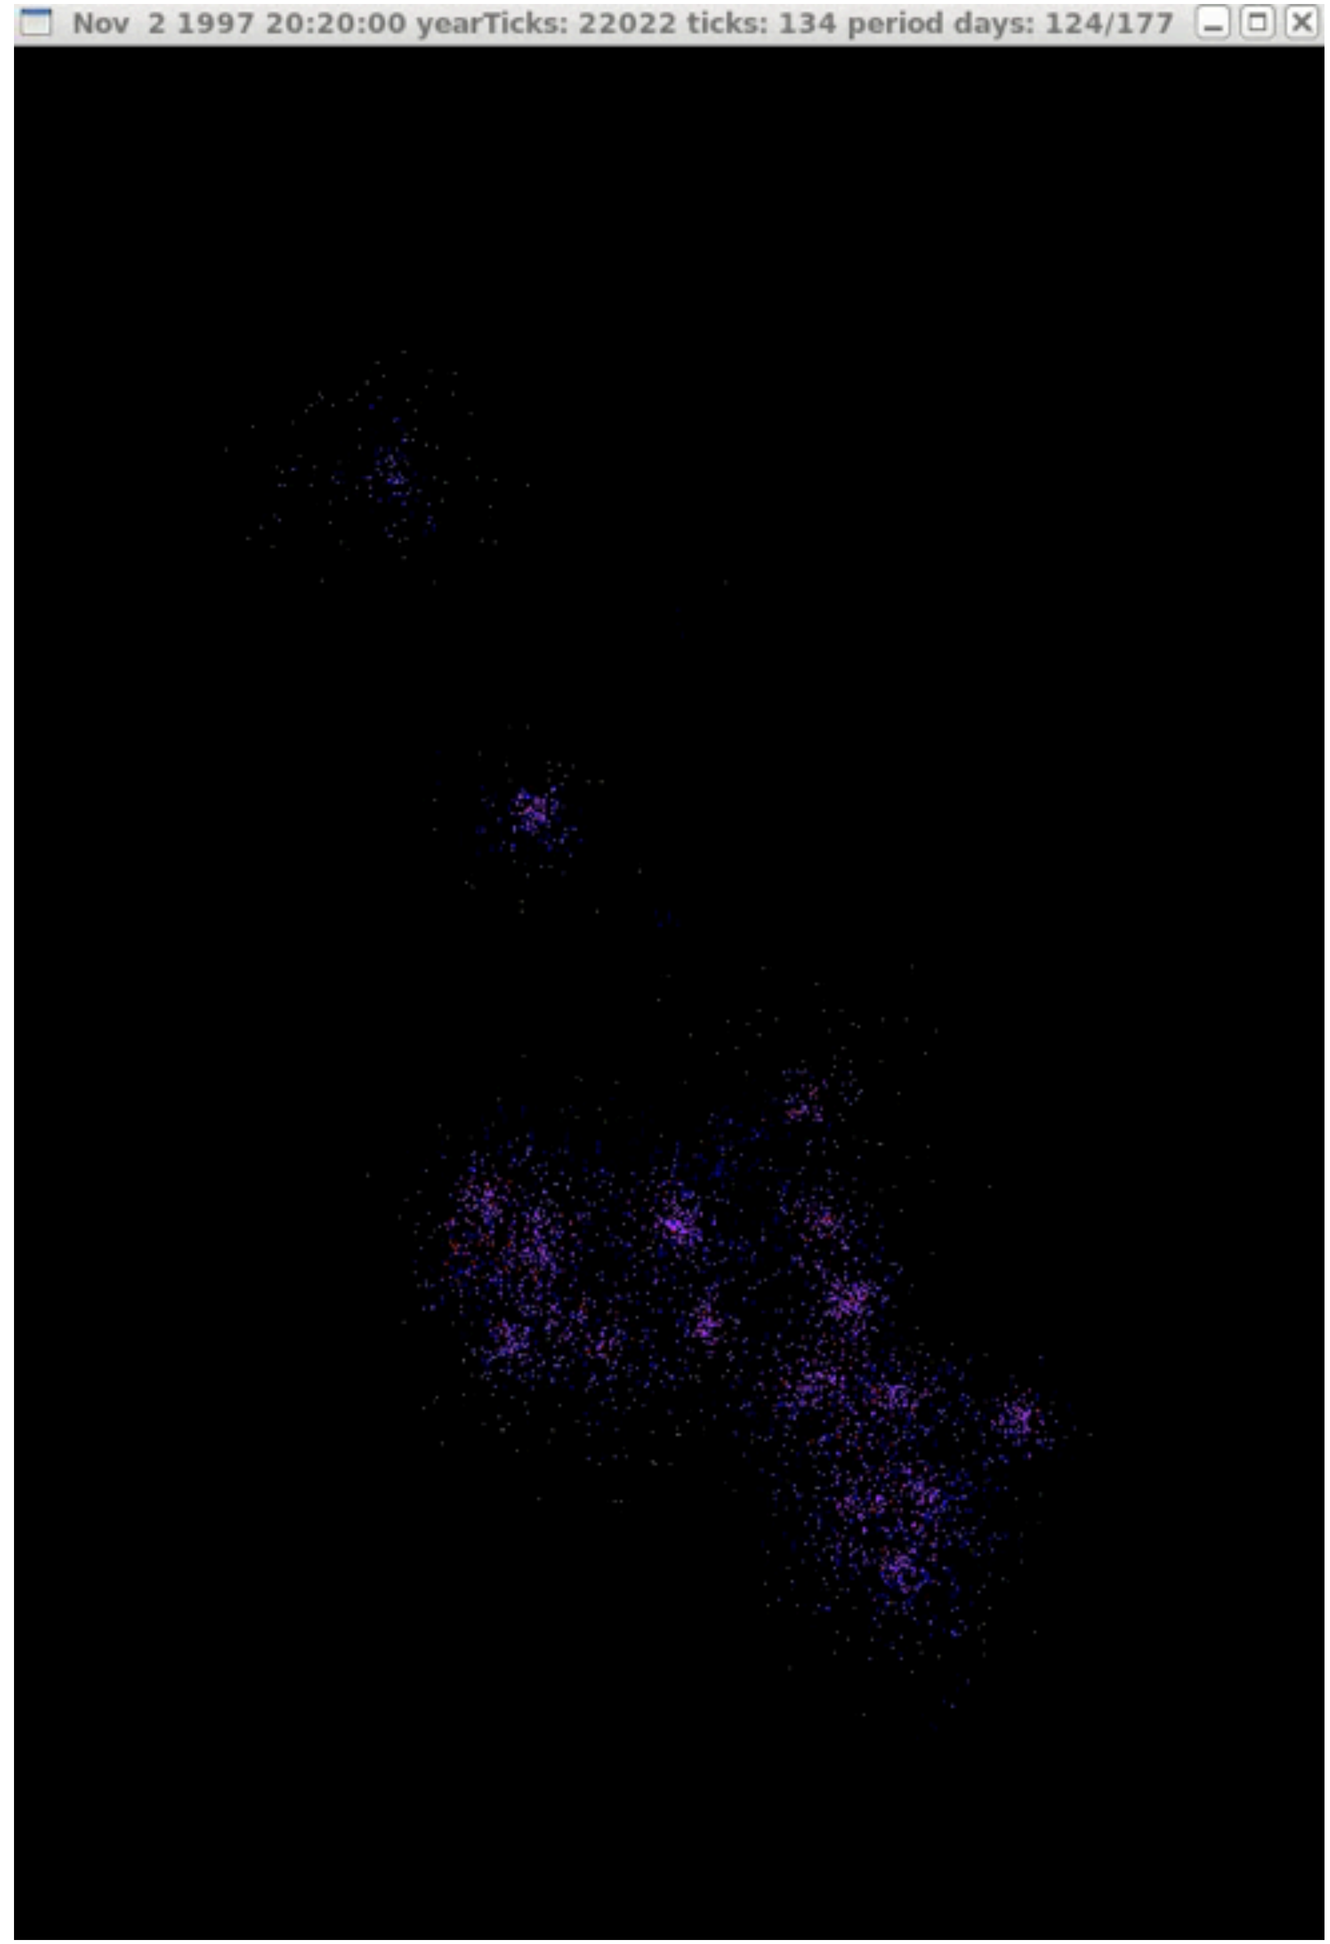

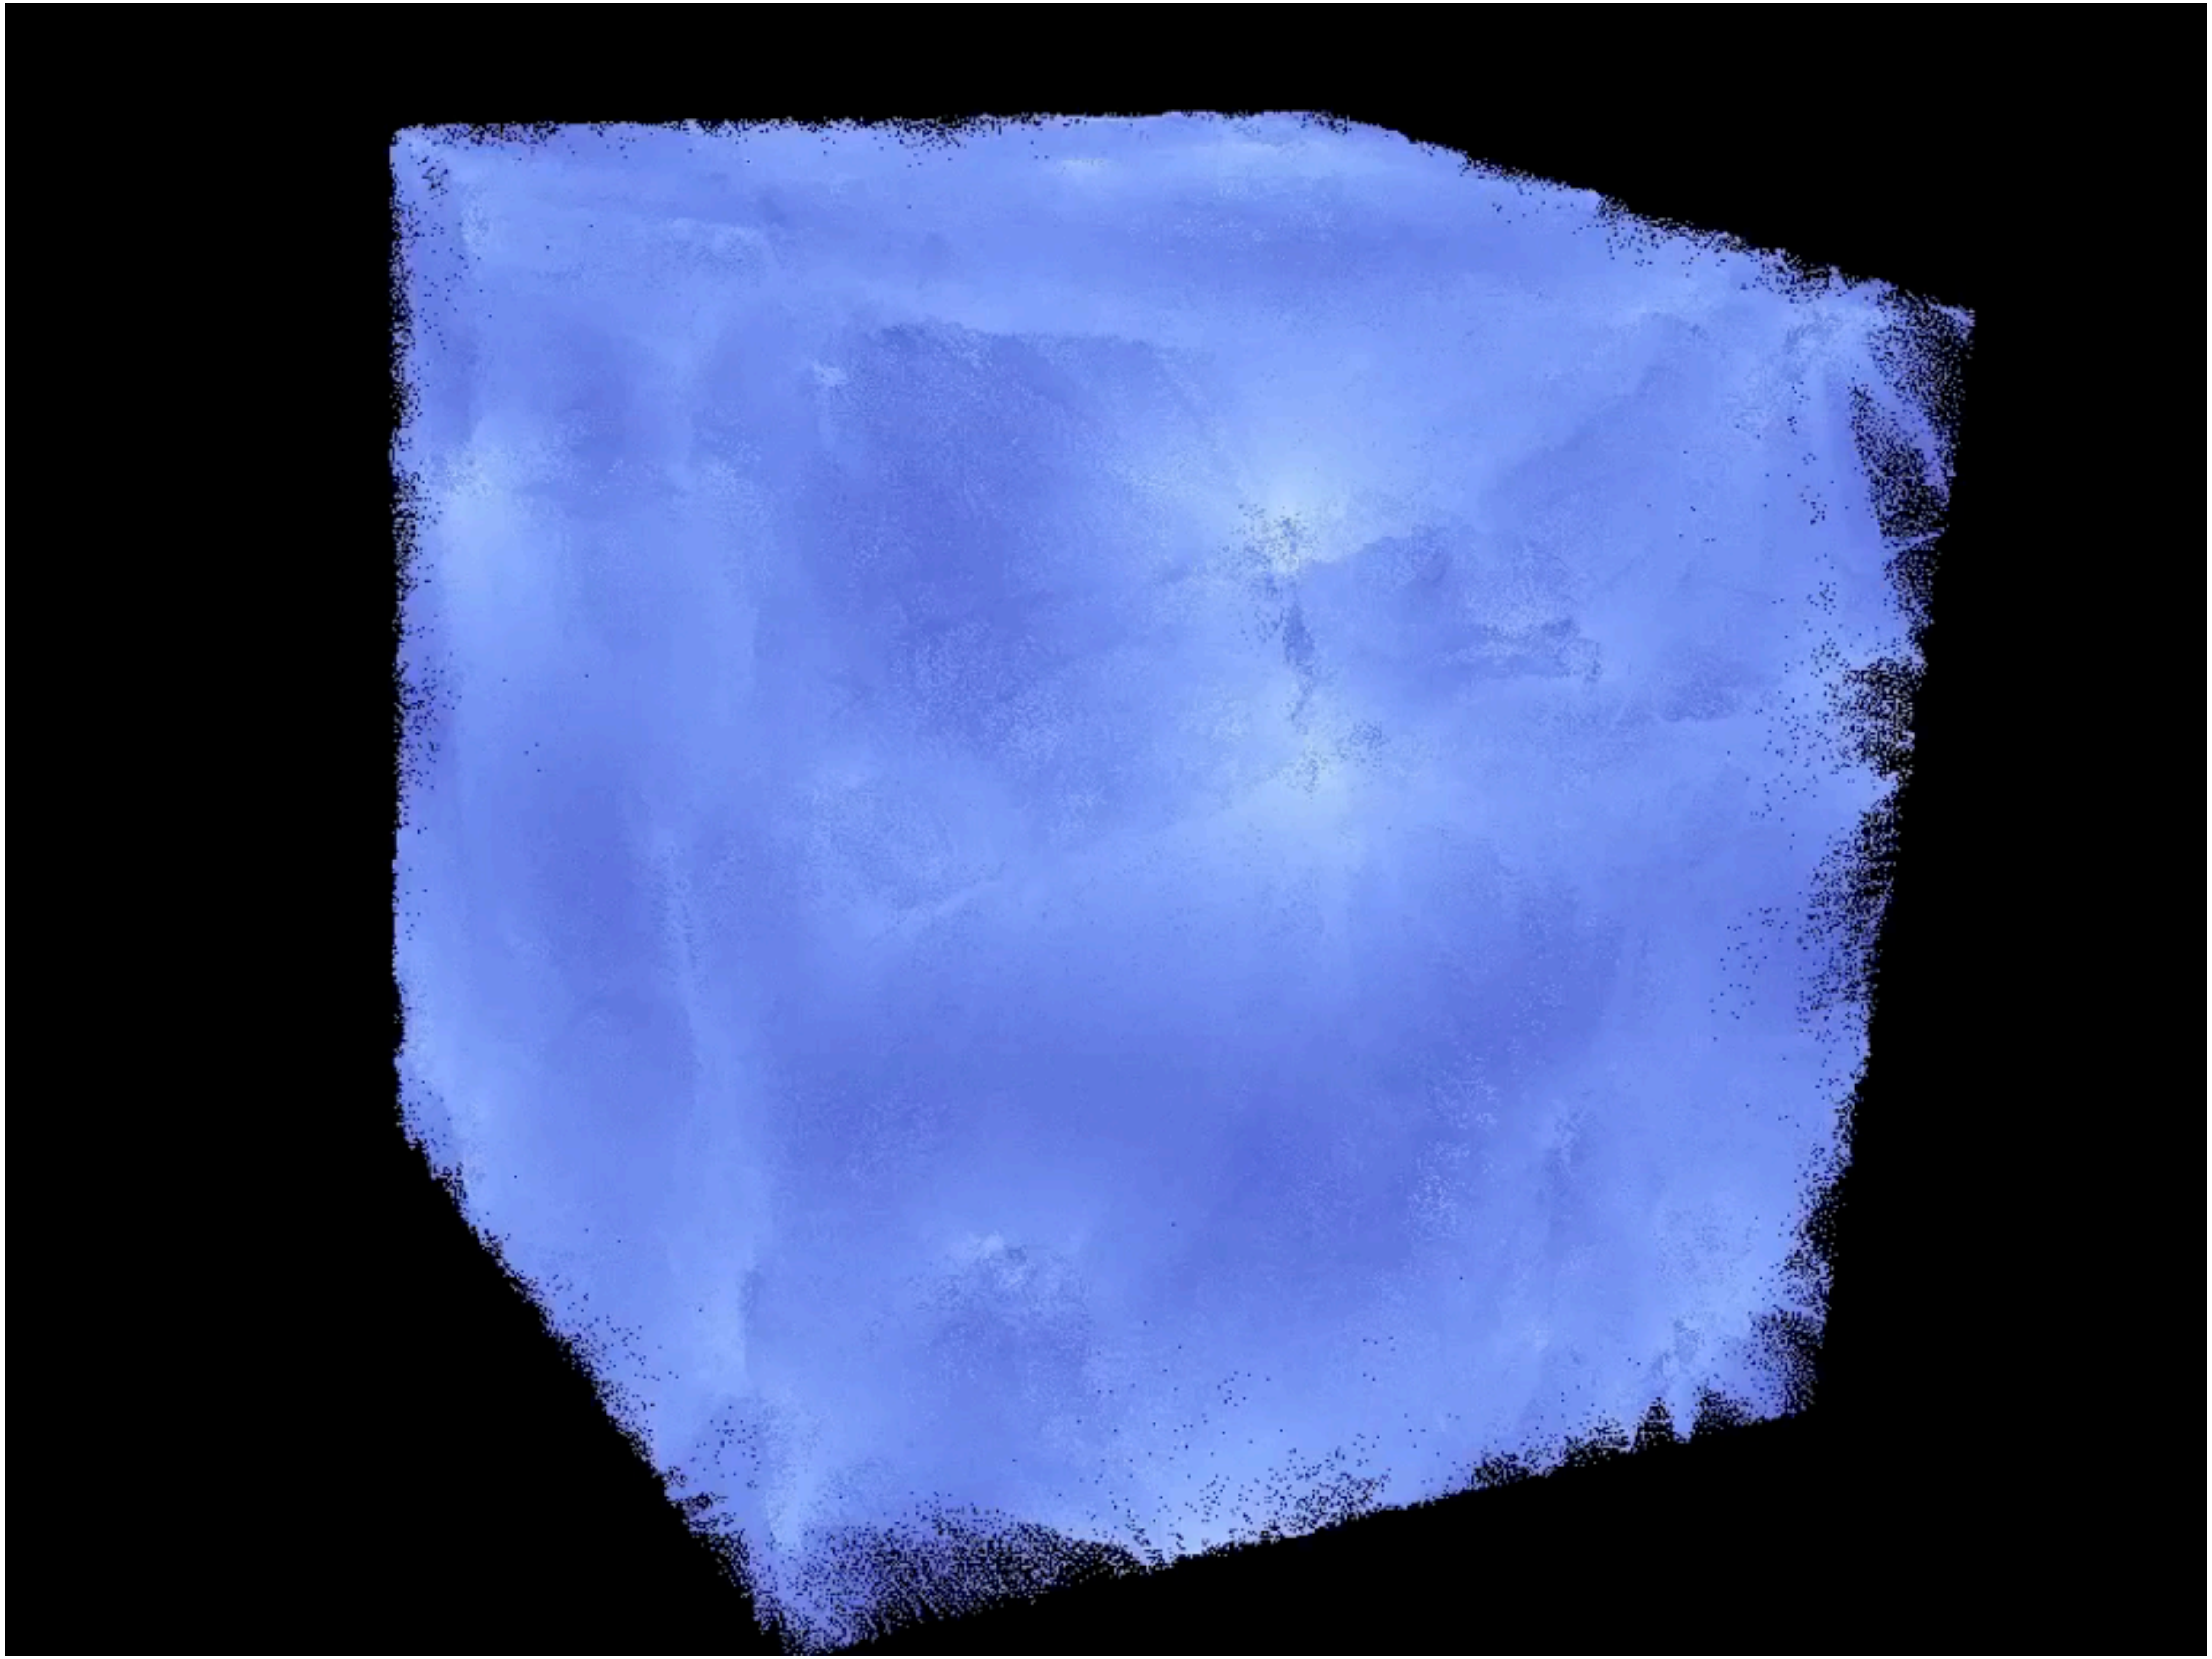

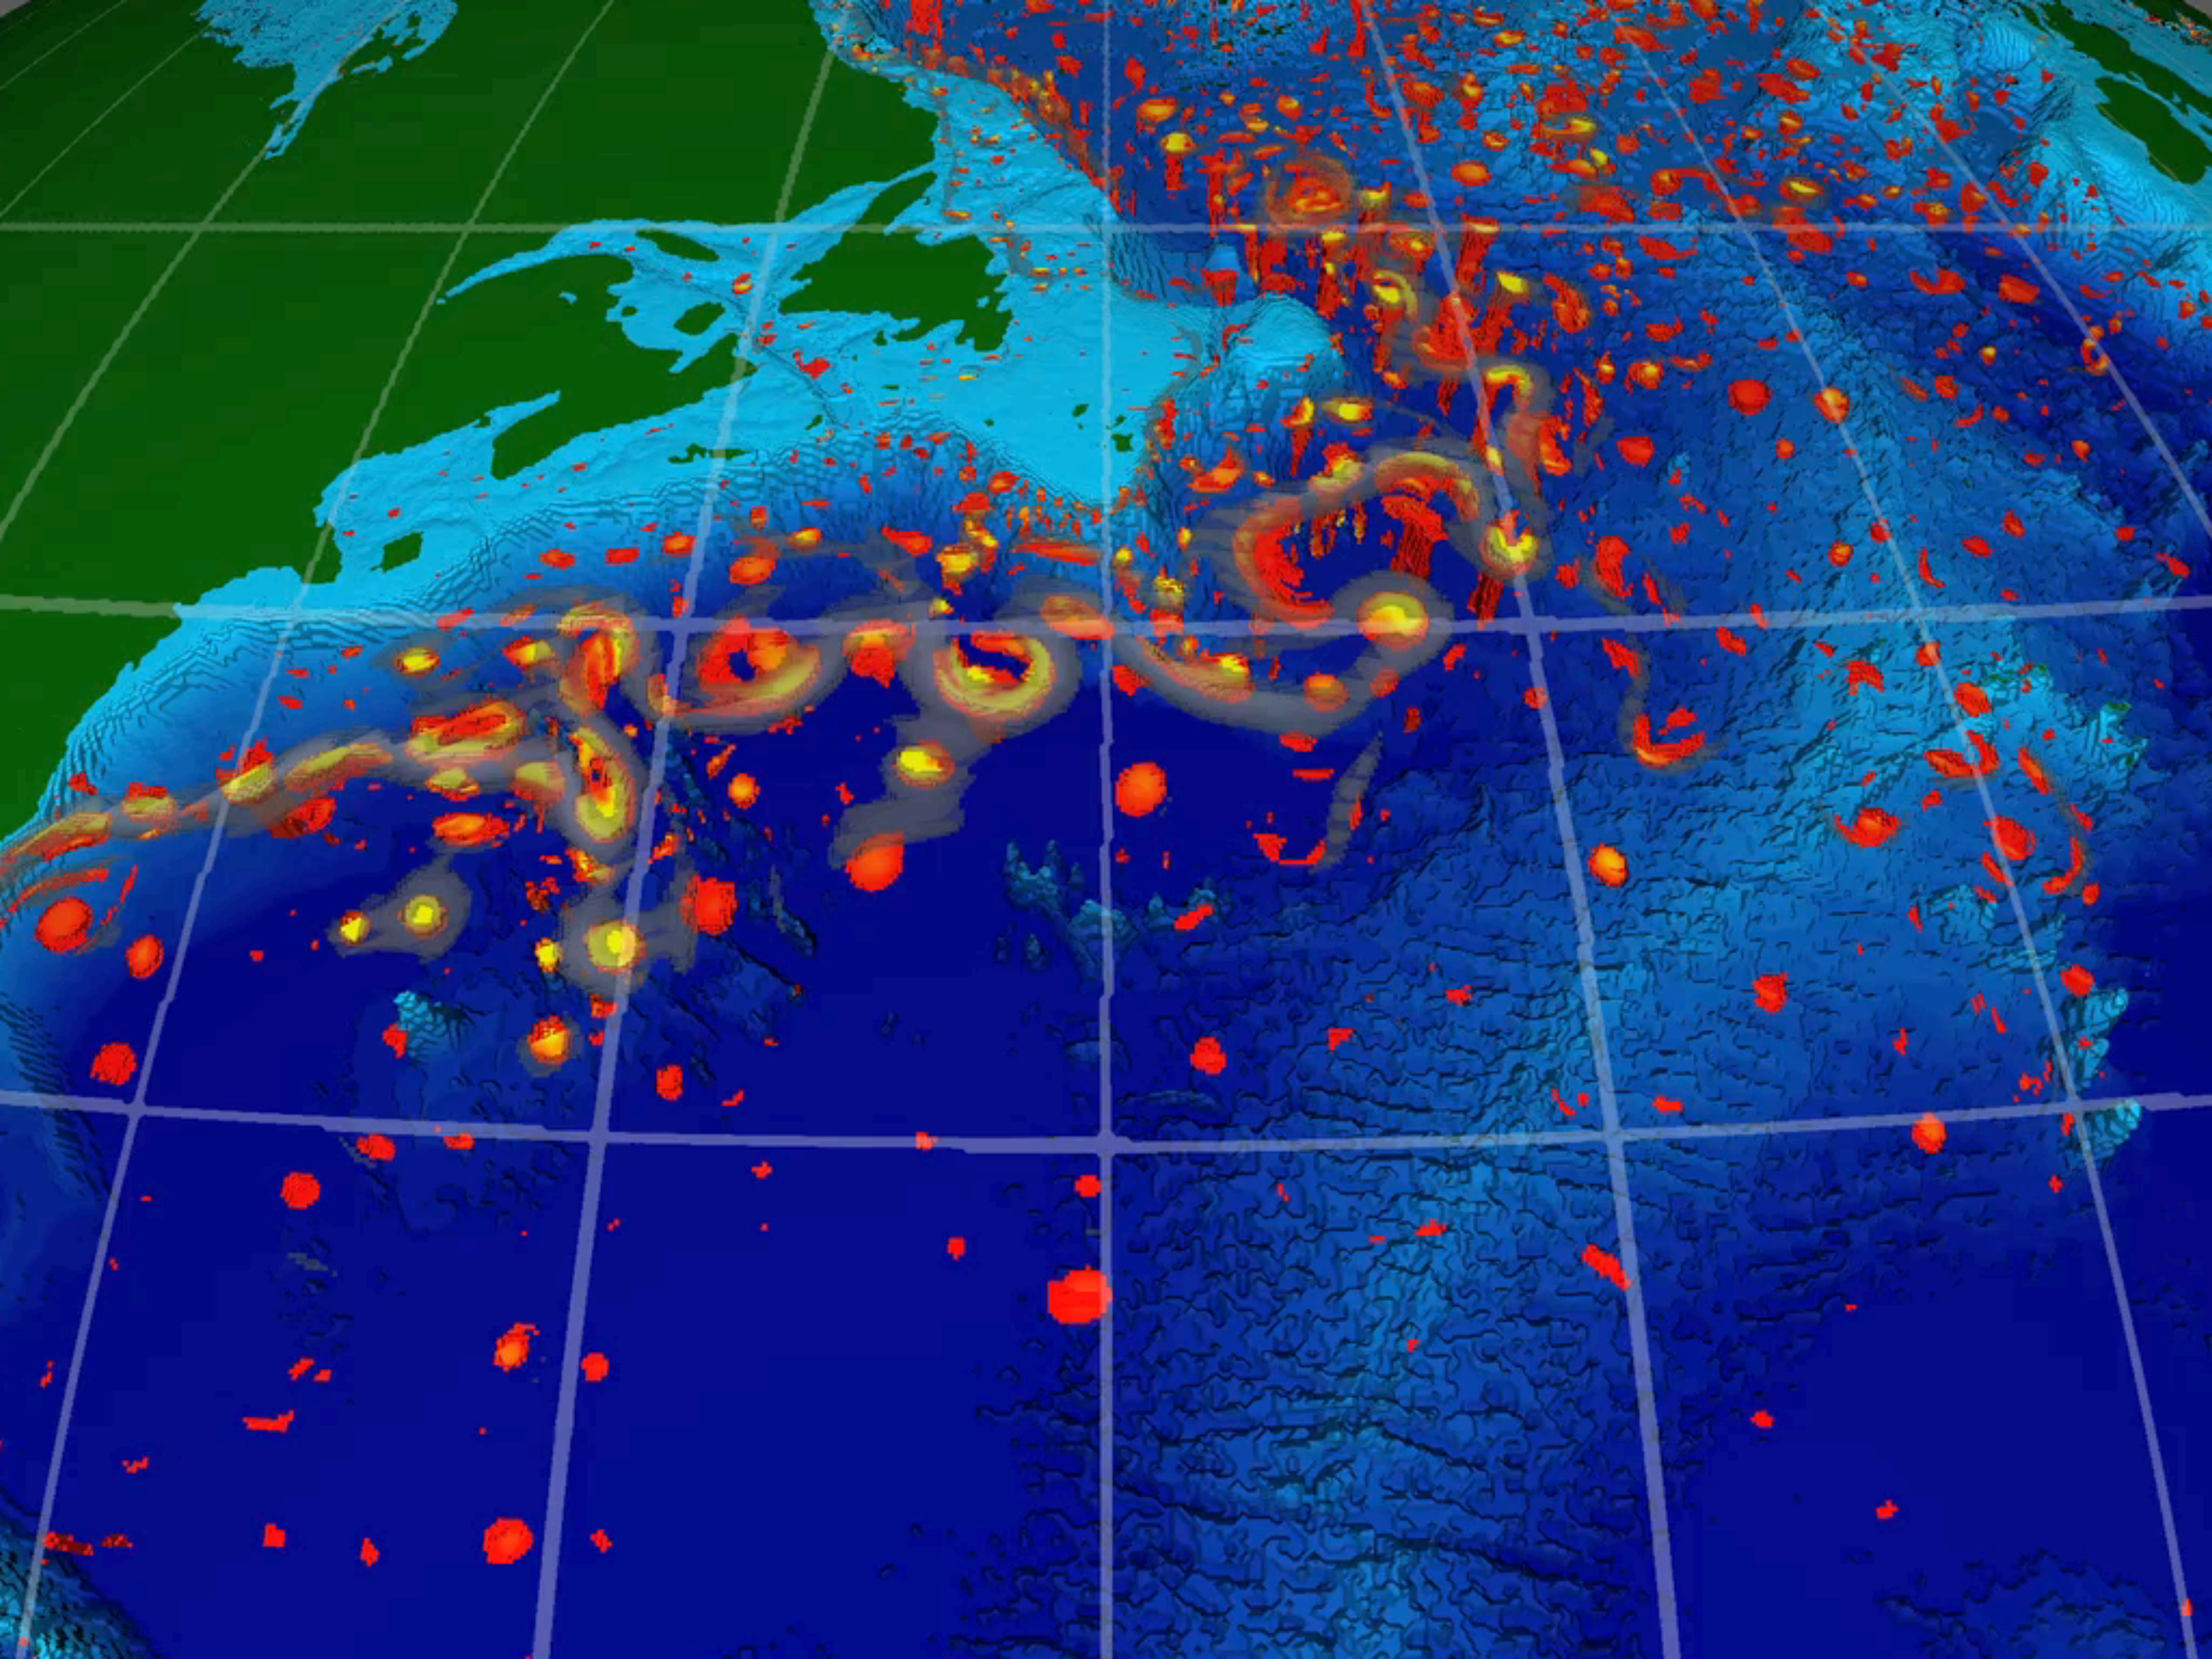

How do I get access to large computational facilities?

- Local clusters within research groups and departments
- ACISS cluster at UO
- INCITE program within the Department of Energy

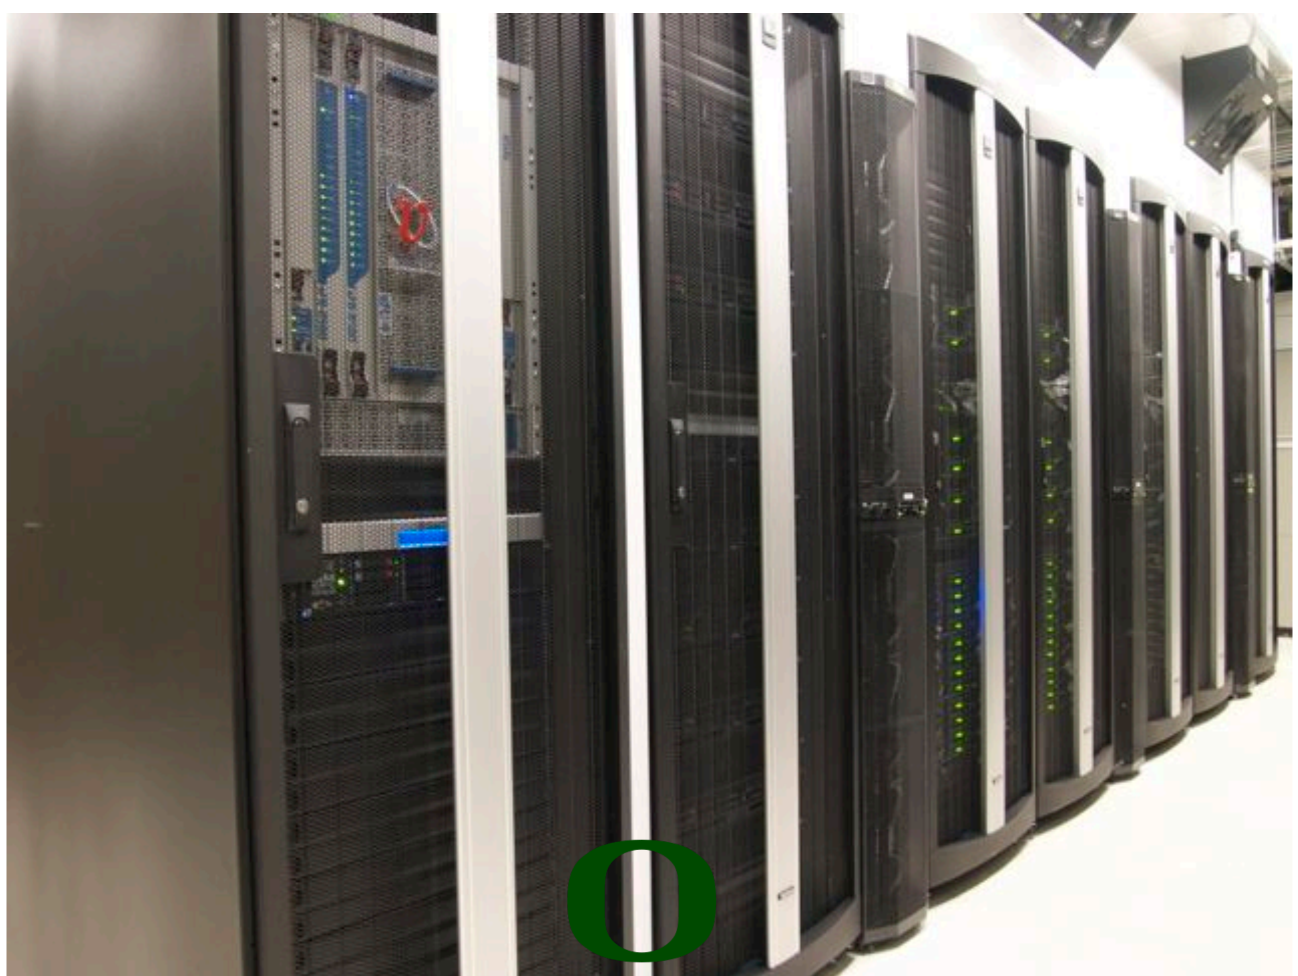

ACISS System - Compute Resources

- Basic nodes (17 Teraflops)
  - 128 ProLiant SL390 G7
  - Two Intel X5650 2.66 GHz 6-core CPUs per node (1,536 total cores)
  - 72 GB DDR3 RAM per basic node
- Address requirements for more compute cores available for running many jobs simultaneously and larger jobs

ACISS System - Compute Resources

- GPU nodes (156 Teraflops)
  - 52 ProLiant SL390 G7
  - Two Intel X5650 2.66 GHz 6-core CPUs per node (624 total cores)
  - 3 NVIDIA M2070 GPUs per node (156 total GPUS)
  - 72GB DDR3 per GPU node
- Address needs for science problems requiring greater computational horsepower.

ACISS System - Compute Resources

- Fat nodes
  - 16 ProLiant DL 580 G7
  - Four Intel X7560 2.266 GHz 8-core CPUs per node (512 total cores)
  - 384 GB DDR3 RAM per fat node
- Address requirement for scientific problems needing very large memory.

INCITE Seeks Research Proposals to Accelerate Scientific Discoveries and Technological

U.S. DEPARTMENT OF ENERGY
INCITE
LEADERSHIP COMPUTING

- The DOE INCITE program provides substantial computational resources every year via a competitive proposal process
  - http://hpc.science.gov/

Fortran is not dead: long live Fortran

- Convolution filter applied to image of Lena Soojblom

```
pure CONCURRENT subroutine convolve(Image, Filter)
   real, intent(in out), HALO(:,:) :: Image
   real, intent(in) :: Filter(-3:3,-3:3)

   Image(0,0) = sum(Filter * Image)

end subroutine convolve
```