# Bigger is Better

## Trends in super computers, super software, and super data

**Michael L. Norman, Director**
**San Diego Supercomputer Center**
**UC San Diego**

# Why are supercomputers needed?

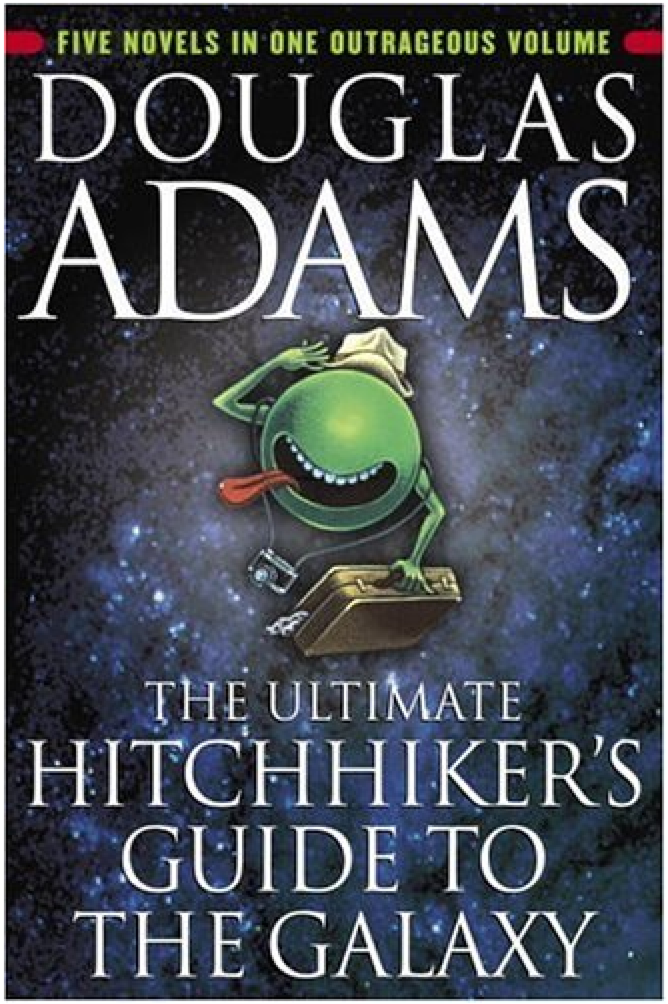


*The universe is famously large….*
Douglas Adams

# Complexity and beauty on a vast range of scales

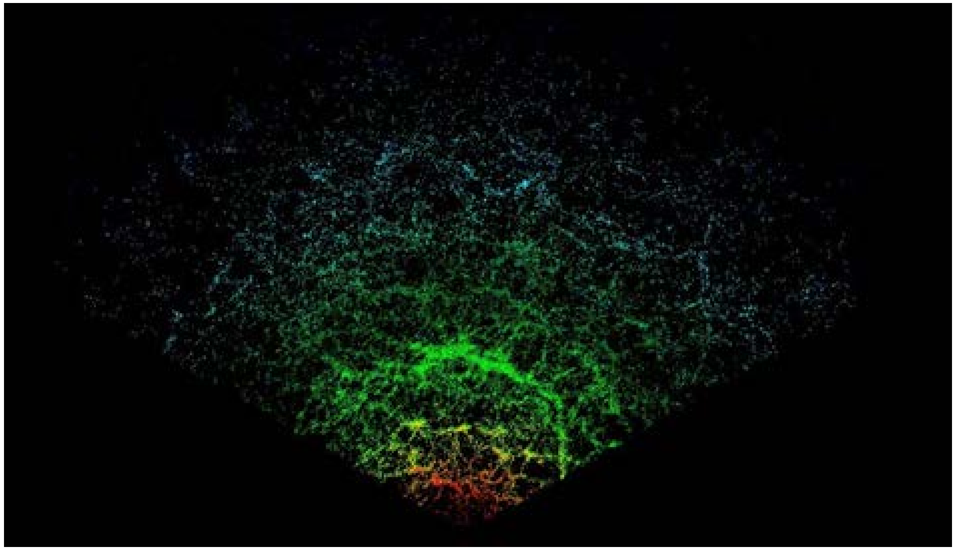

*How can we possibly understand all that?*

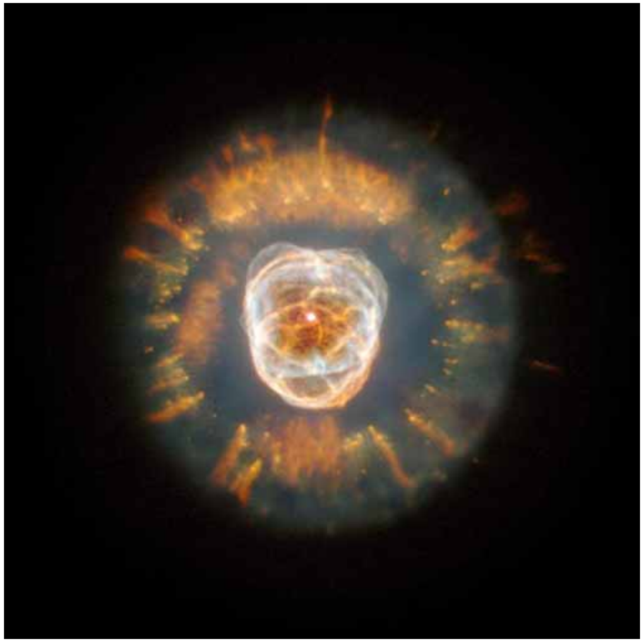

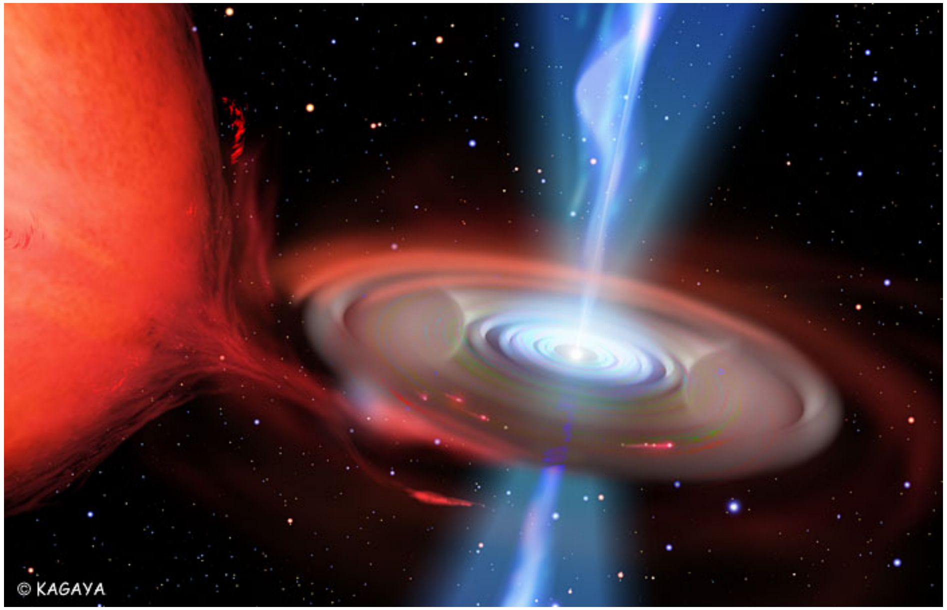


DON'T PANIC

# Equations of astrophysics fluid dynamics (non-relativistic)

**Conservation of Mass**

$$\frac{D\rho}{Dt} + \rho\nabla\cdot\mathbf{v} = 0; \quad (1)$$

**Conservation of Momentum**

$$\rho\frac{D\mathbf{v}}{Dt} = -\nabla\mathrm{p} + \left(\frac{\chi}{c}\right)\mathbf{F} + \frac{1}{4\pi}(\nabla\times\mathrm{B})\times\mathrm{B} - \rho\nabla\Phi; \quad (2)$$

**Conservation of Gas Energy**

$$\rho\frac{D}{Dt}\left(\frac{\mathrm{e}}{\rho}\right) + \mathrm{p}\nabla\cdot\mathbf{v} = c\kappa_{\mathrm{E}}\mathrm{E} - 4\pi\kappa_{\mathrm{P}}\mathrm{B}_{\mathrm{p}}; \quad (3)$$

**Conservation of Radiation Energy**

$$\rho\frac{D}{Dt}\left(\frac{\mathrm{E}}{\rho}\right) + \nabla\cdot\mathbf{F} + \nabla\mathbf{v}:\mathbf{P} = 4\pi\kappa_{\mathrm{P}}\mathrm{B}_{\mathrm{p}} - c\kappa_{\mathrm{E}}\mathrm{E}; \quad (4)$$

**Conservation of Magnetic Flux**

$$\frac{\partial\mathrm{B}}{\partial t} = \nabla\times(\mathbf{v}\times\mathrm{B}). \quad (5)$$

**Newton's law of Gravity**

$$\nabla^2\Phi = 4\pi\mathrm{G}\rho. \quad (15)$$

**Microphysics**

$$p(\rho, e), \kappa_P, \kappa_E, \chi$$

# *Is it Real or Memorex?*

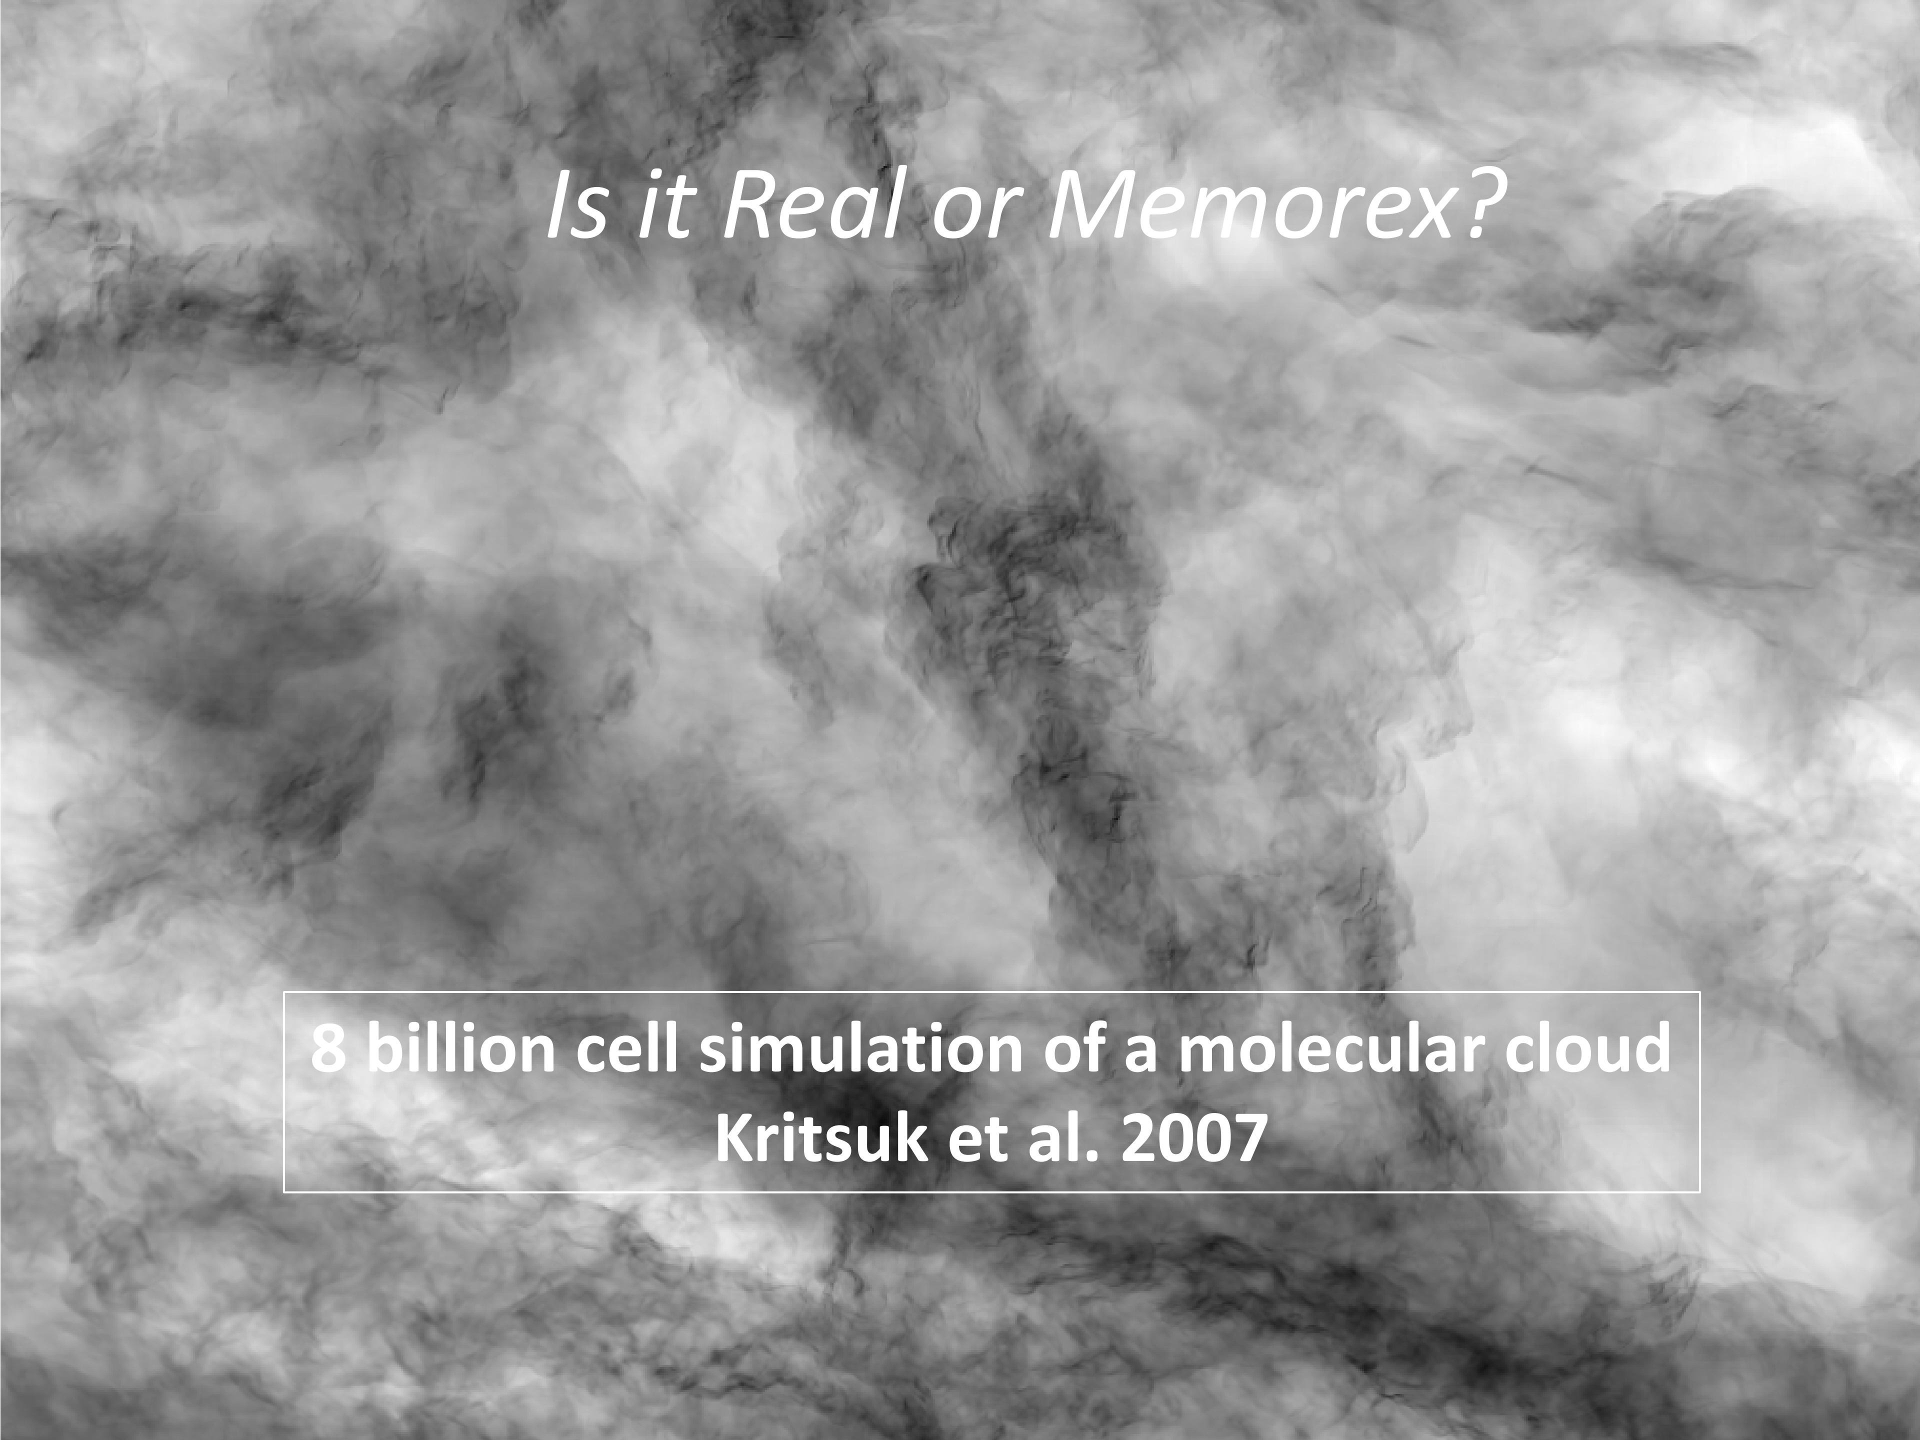

**8 billion cell simulation of a molecular cloud**
**Kritsuk et al. 2007**

# Outline

- Astrocomputing and supercomputing
- A bit about computational methodology
- Supercomputing technology trends
- Exploring cosmic Renaissance with supercomputers

# Astrocomputing and Supercomputing

- Astrophysicists have always been at the vanguard of supercomputing
  - Martin Schwarzschild used LASL's ENIAC for stellar evolution calculations (40s 50s)
  - Stirling Colgate, Jim Wilson pioneering simulations of core collapse supernovae (late 60s)
  - Larry Smarr 2-black hole collision (mid 70s)

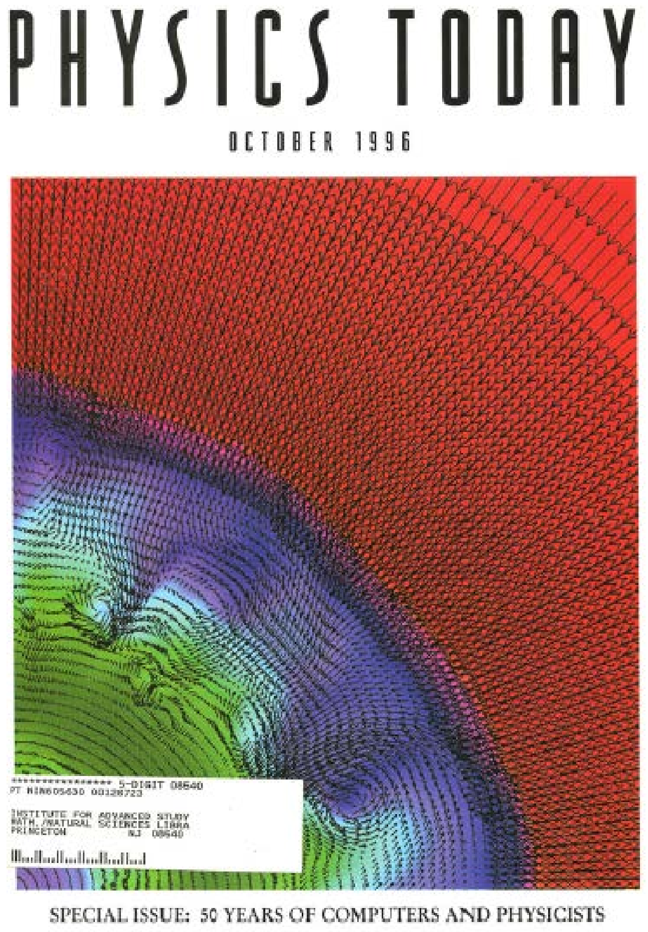


"Probing Cosmic Mysteries Using Supercomputers", Norman (1996)

Cosmological N-body simulations

*

The *Millenium Simulation*

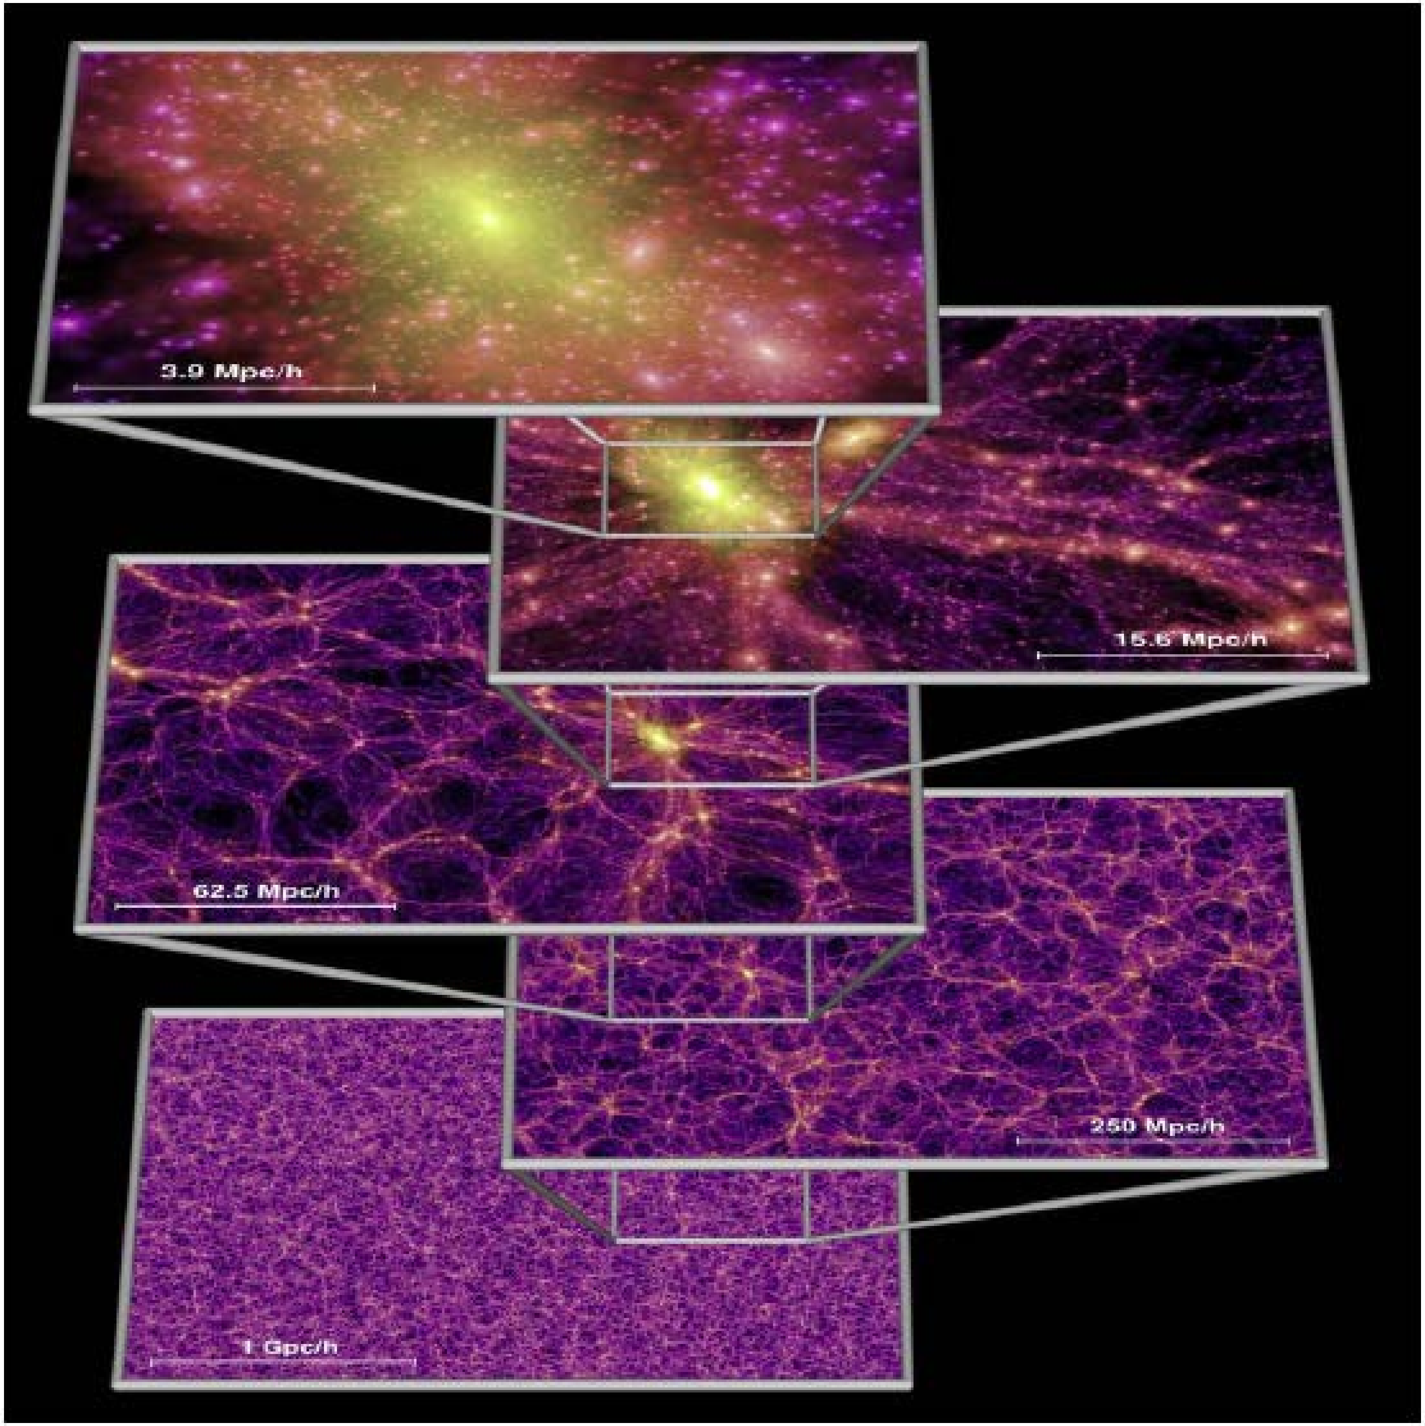


Springel et al. (2005)

# Gravitational N-body simulations

(N=$10^{12}$, 2012)

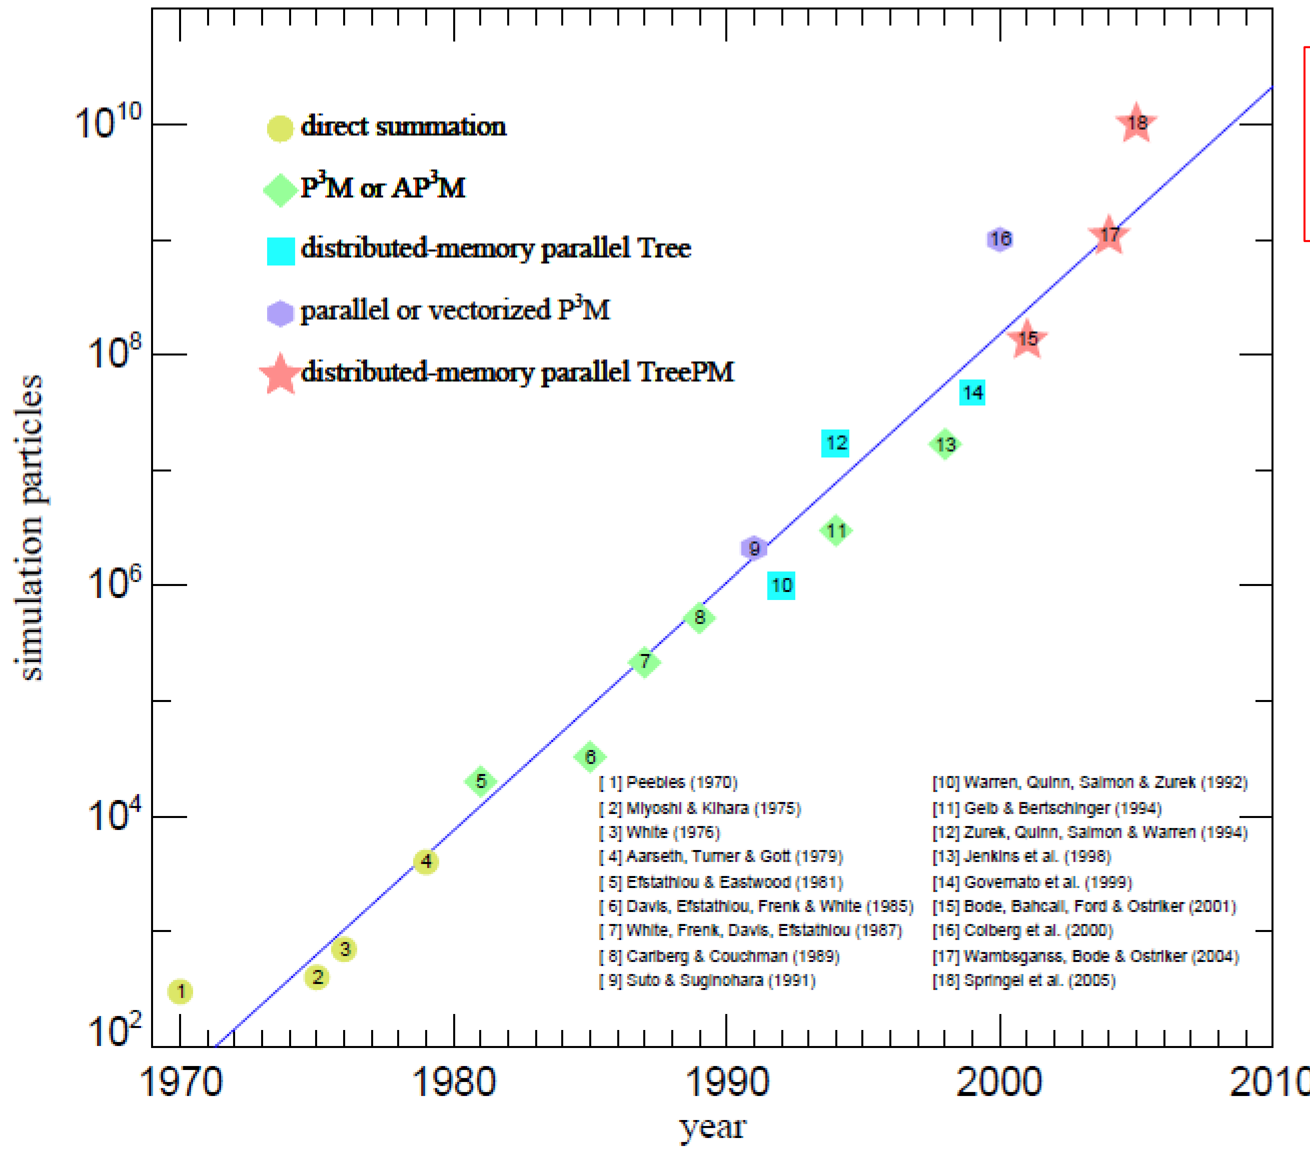


2012 ACM Gordon Bell prize finalist

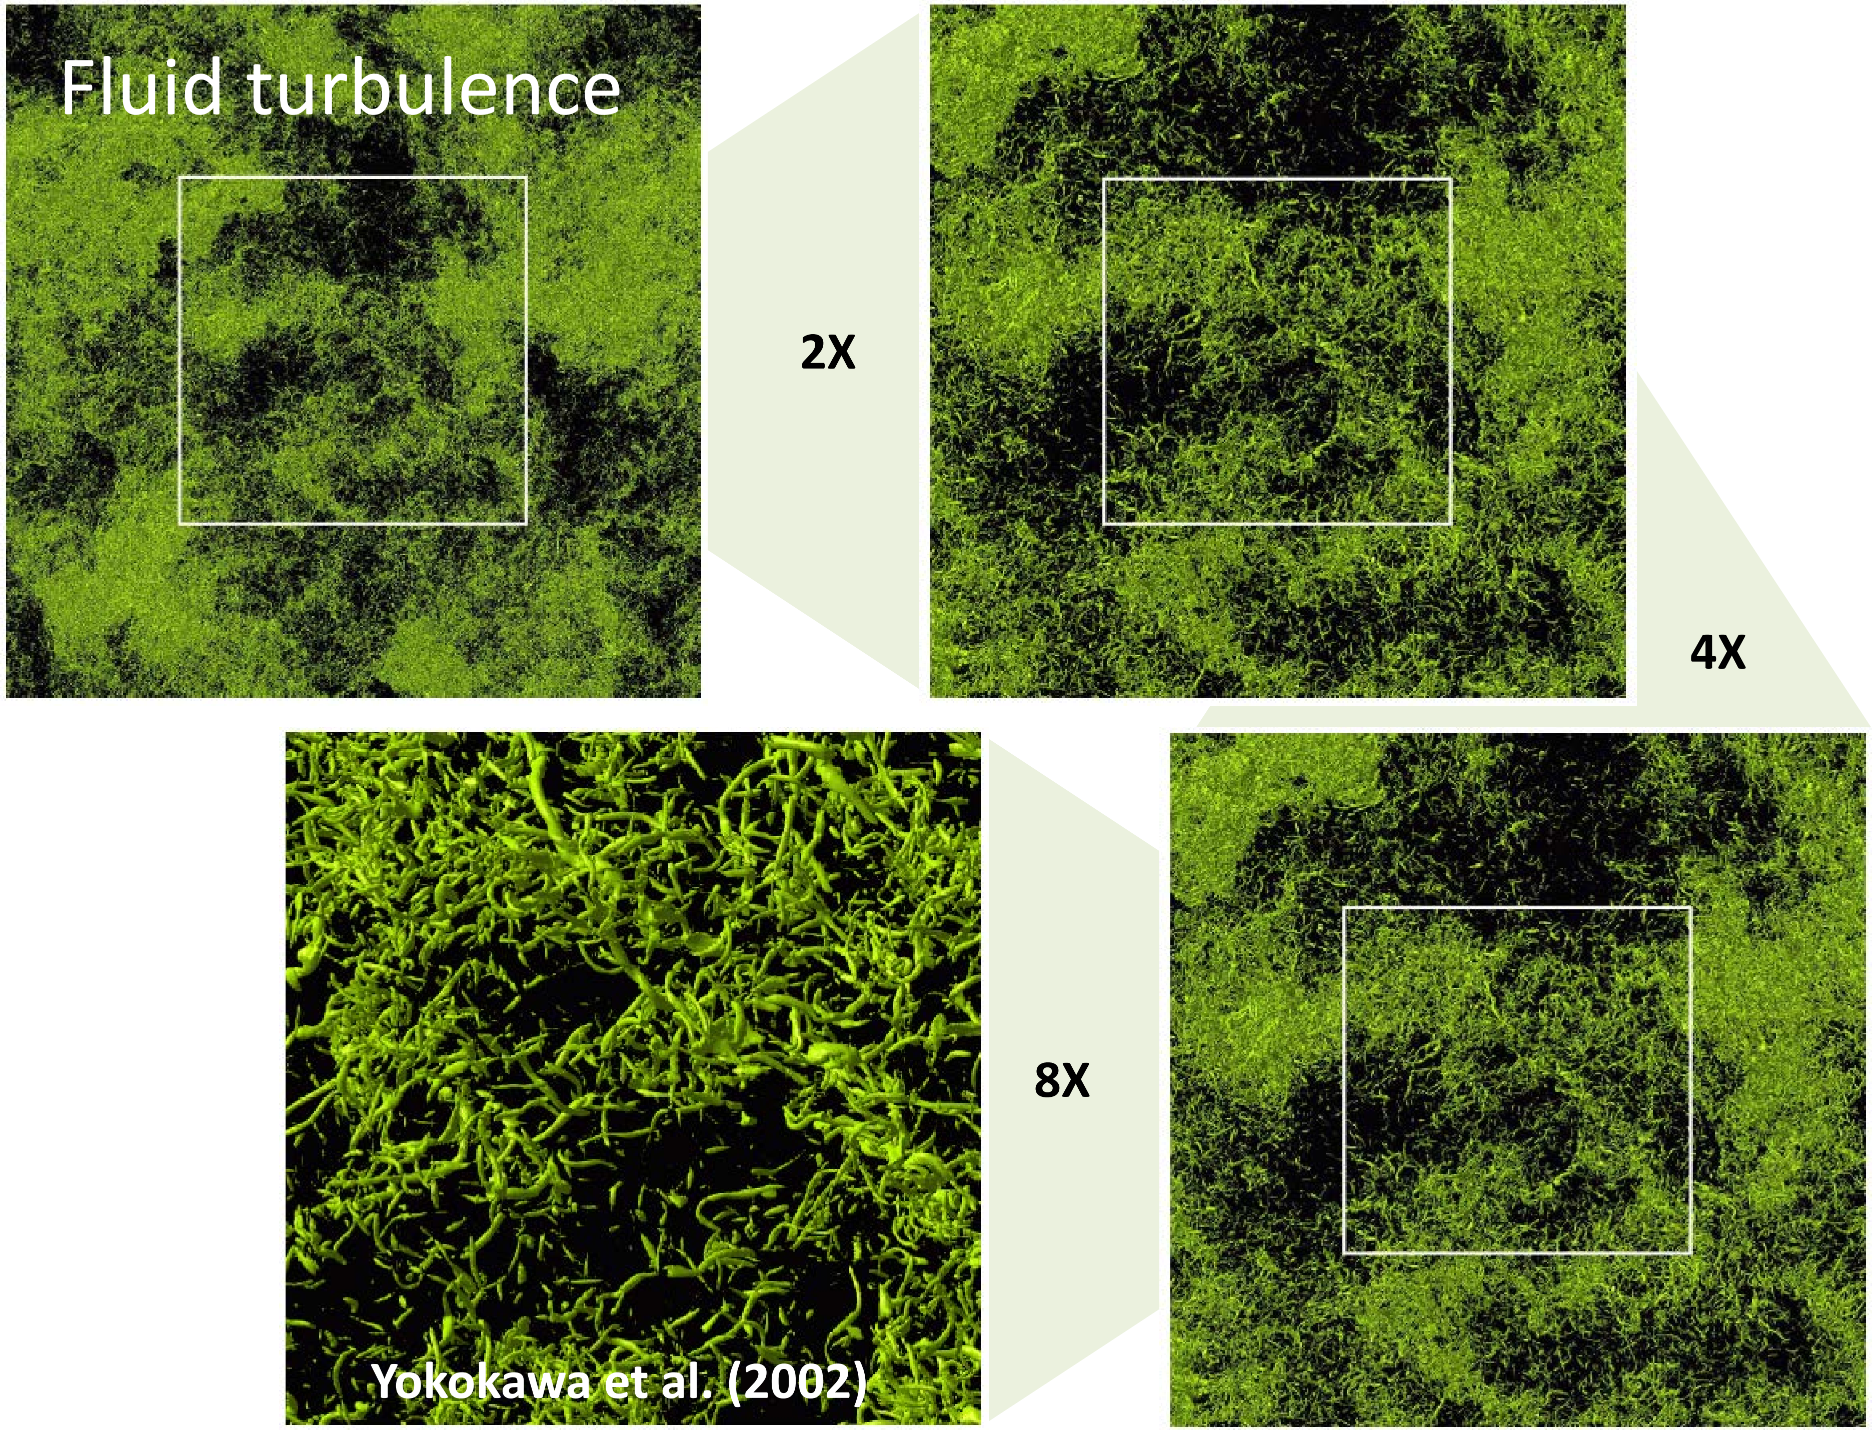
Fluid turbulence
2X
4X
8X
Yokokawa et al. (2002)

# Astrocomputing and Data computing

- Astronomers have always been at the vanguard of digital data explosion
  - VLA radio telescope
  - Hubble Space Telescope
  - Sloan Digital Sky Survey

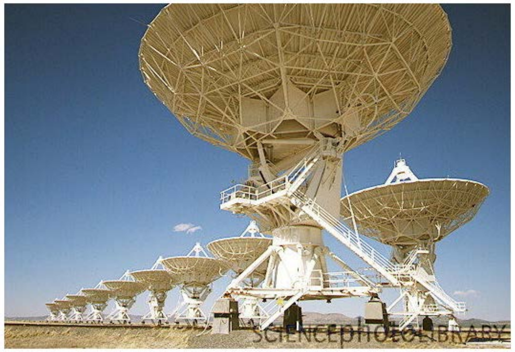

# Sloan Digital Sky Survey

- **“The Cosmic Genome Project”**
- Two surveys in one
  - Photometric survey in 5 bands
  - Spectroscopic redshift survey
- Data is public
  - 2.5 Terapixels of images
  - 40 TB of raw data => 120TB processed
  - 5 TB catalogs => 35TB in the end
- Started in 1992, finished in 2008
- Database and spectrograph built at JHU (SkyServer)

*The University of Chicago*
*Princeton University*
*The Johns Hopkins University*
*The University of Washington*
*New Mexico State University*
*Fermi National Accelerator Laboratory*
*US Naval Observatory*
*The Japanese Participation Group*
*The Institute for Advanced Study*
*Max Planck Inst, Heidelberg*

*Sloan Foundation, NSF, DOE, NASA*

Slide courtesy of Alex Szalay, JHU

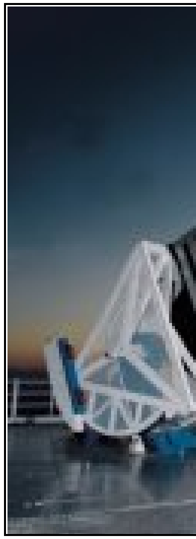

SD
2.4

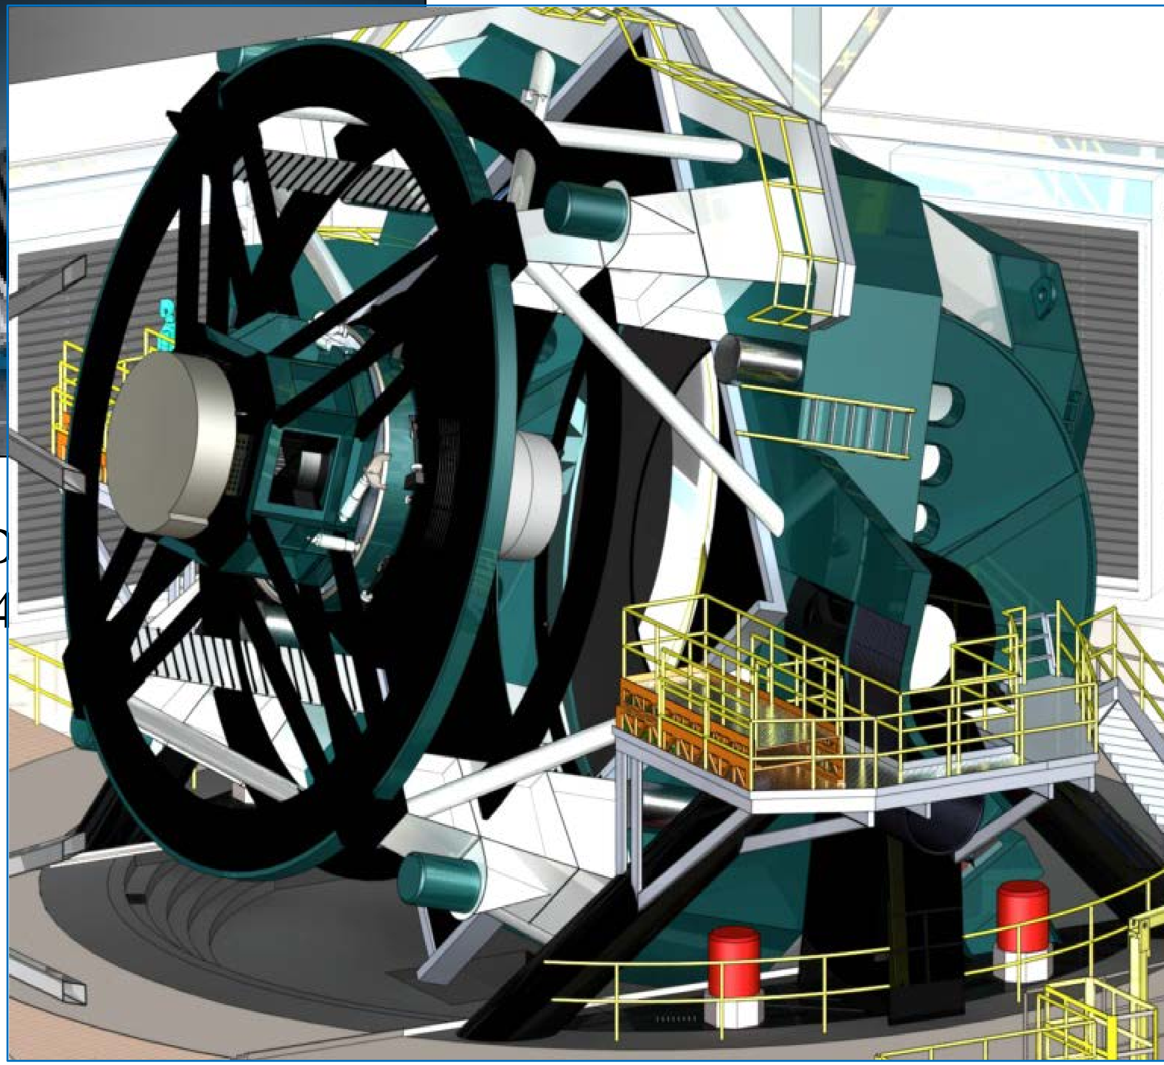

LSST
8.4m  3.2Gpixel

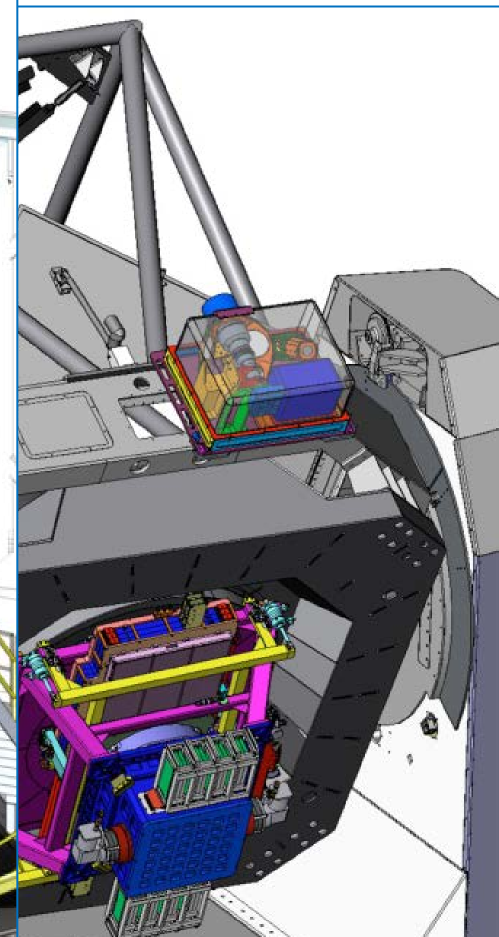

PanSTARRS
1.8m  1.4Gpixel

# Galaxy Survey Trends

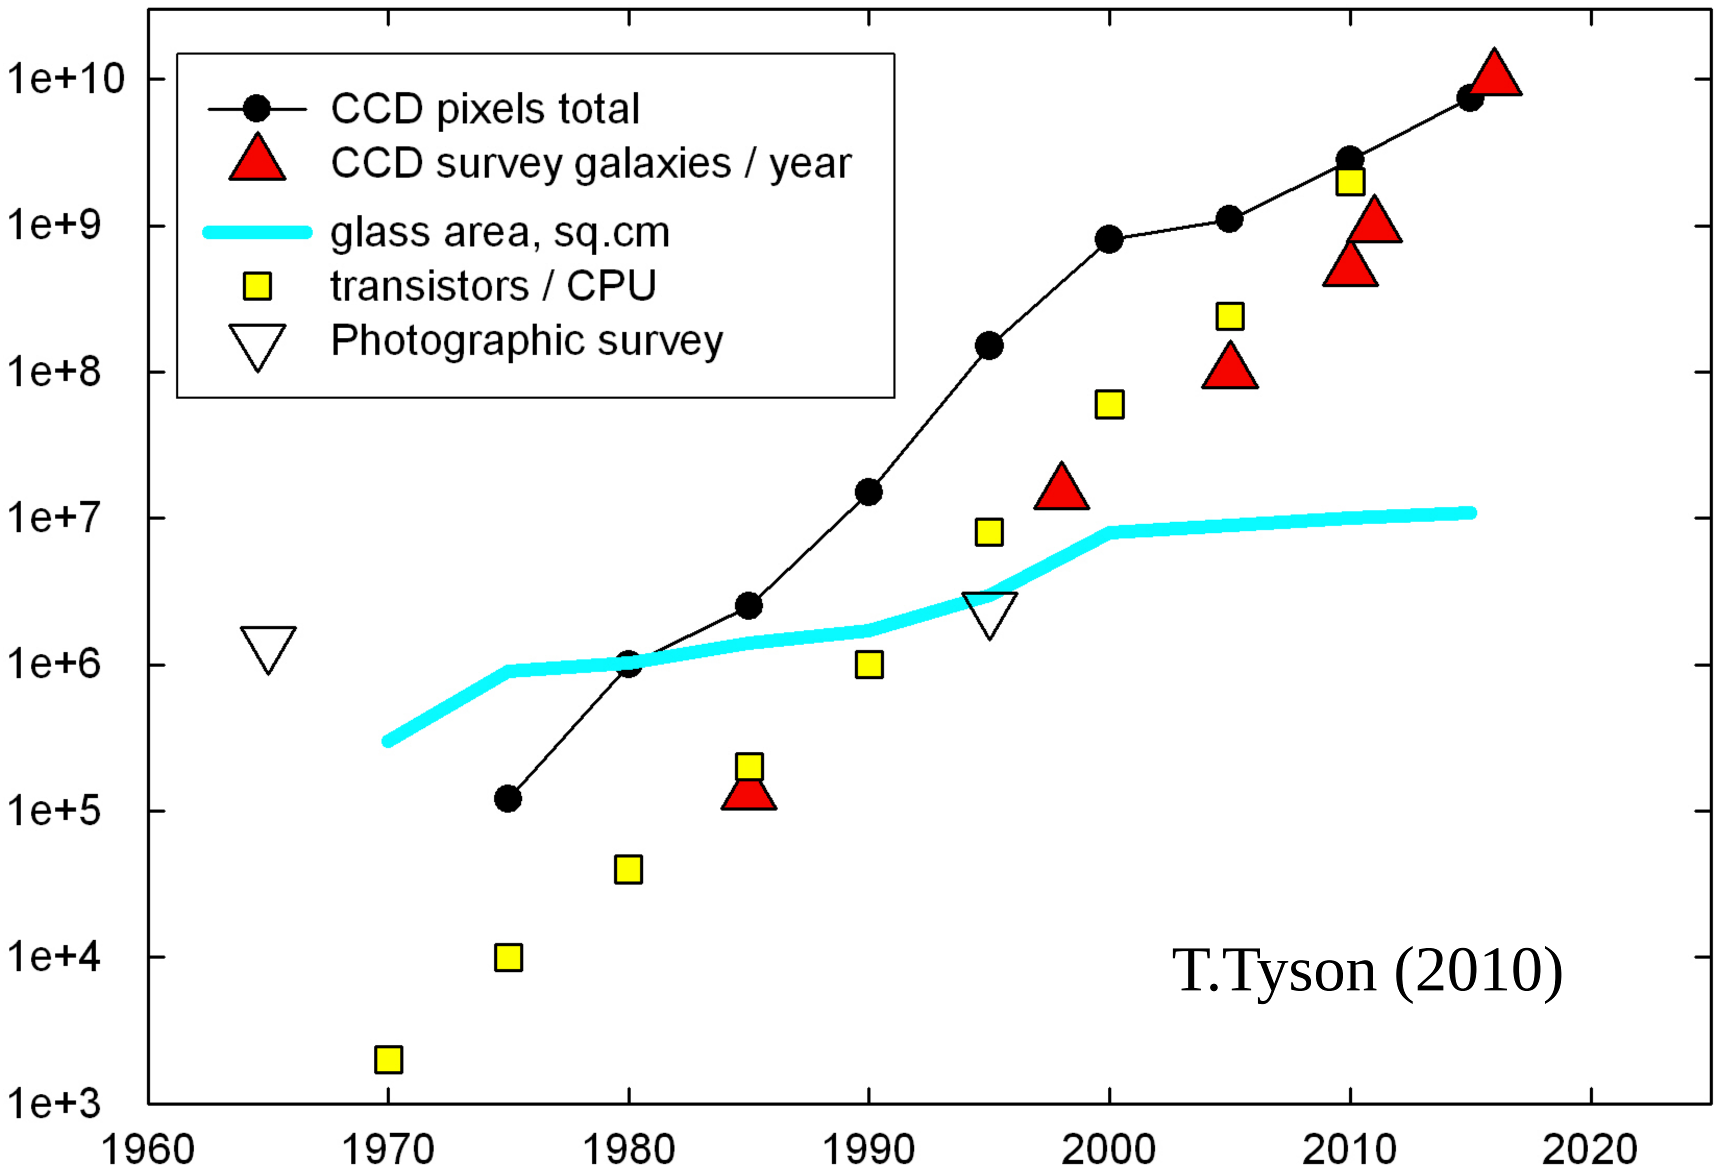

How are supercomputers used?

# A BIT ABOUT COMPUTATIONAL METHODOLOGY

*12th 'Kingston meeting': Computational Astrophysics*
*ASP Conference Series Vol. 123, 1997*
*David A. Clarke & Michael J. West (eds.)*

# Computational Astrophysics: The "New Astronomy" for the 21st Century

Michael L. Norman

*Laboratory for Computational Astrophysics, Astronomy Department, and NCSA, University of Illinois at Urbana-Champaign, Urbana, IL 61801, U.S.A.*

**Abstract.** I discuss the role computer simulation has played in astronomical research, reviewing briefly the origins of the field only to place into perspective the enormous strides which have been achieved in recent decades. I will highlight areas where computational astrophysics has already made a scientific impact, and attempt to discover the conditions which lead to real progress. Finally, I will prognosticate on what the future may hold in store for the second "New Astronomy" revolution already well underway.

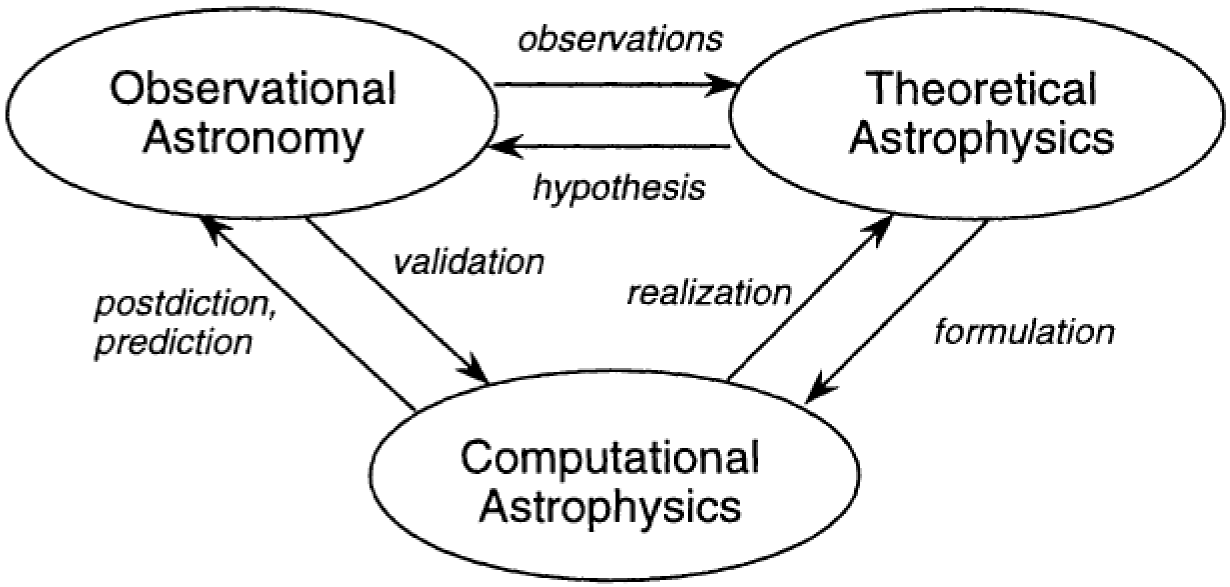


Figure 2. Synergies between observational, theoretical and computational astrophysics.

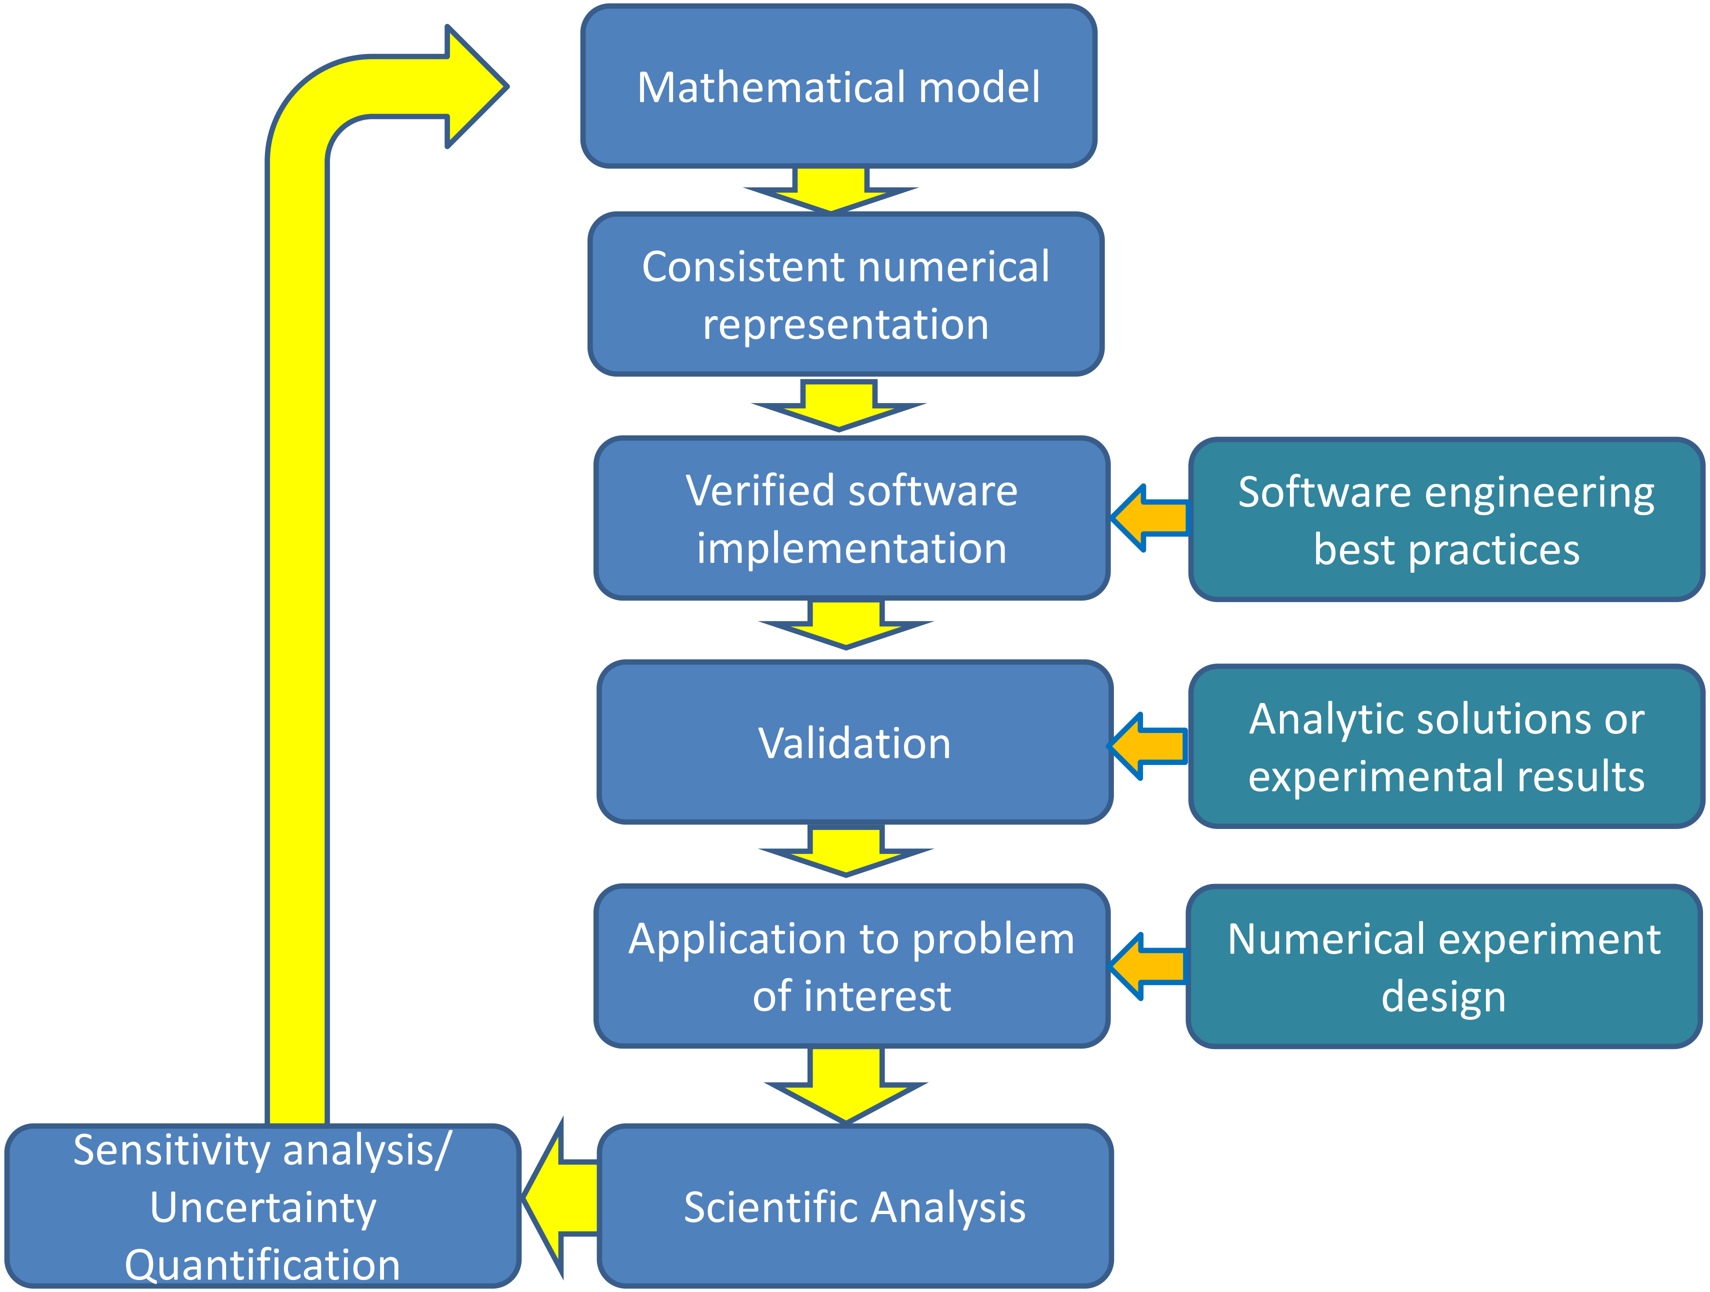
Mathematical model
Consistent numerical representation
Verified software implementation
Software engineering best practices
Validation
Analytic solutions or experimental results
Application to problem of interest
Numerical experiment design
Scientific Analysis
Sensitivity analysis/ Uncertainty Quantification

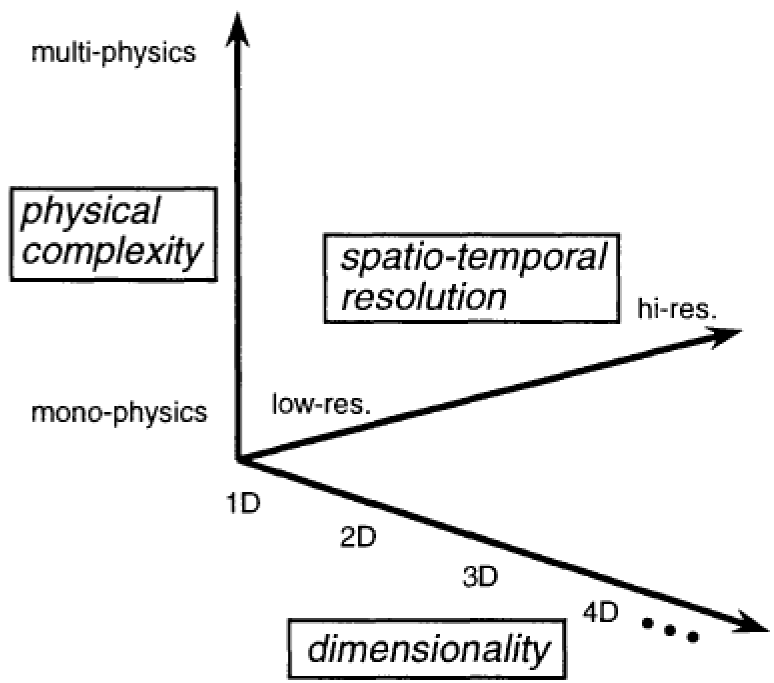


Figure 4. Progress in numerical modeling occurs along (at least) three axes in a conceptual phase space: dimensionality, spatio-temporal resolution, and physical complexity.

# Effect of Increased Resolution

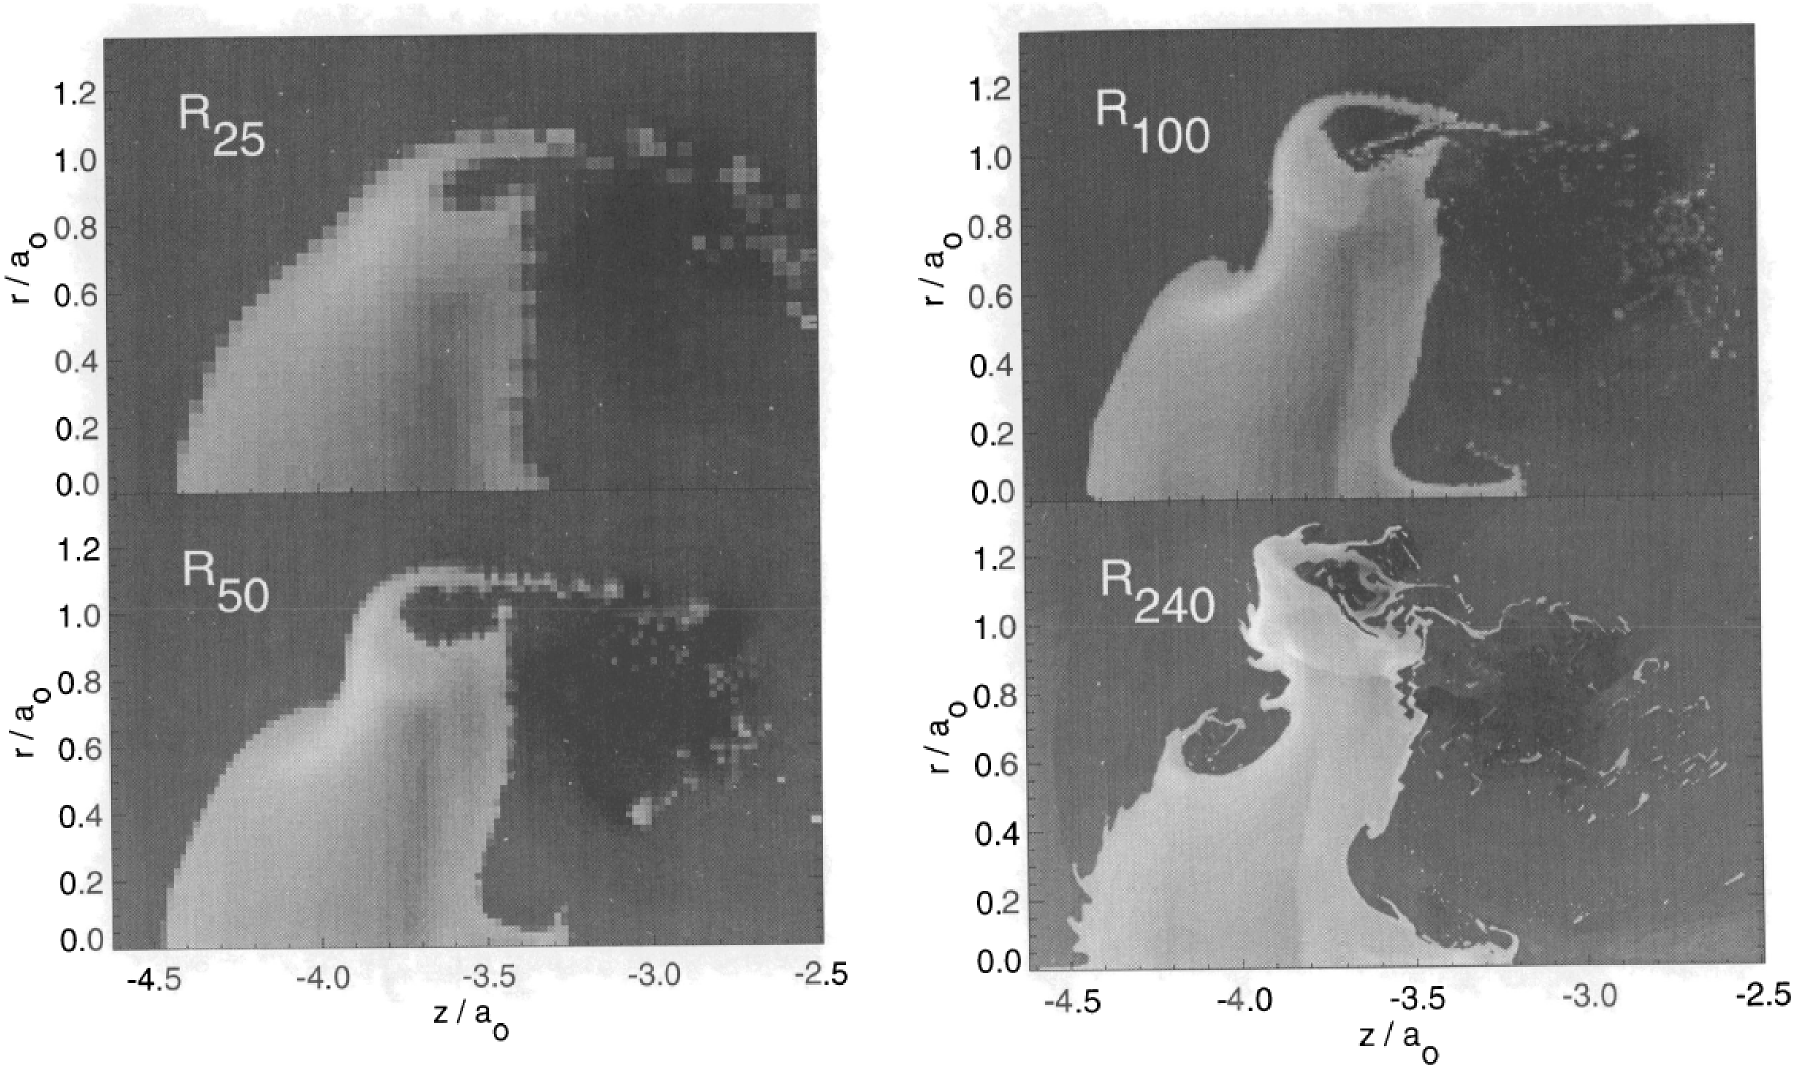


**MacLow et al. (1994)**

# Effect of Additional Physics

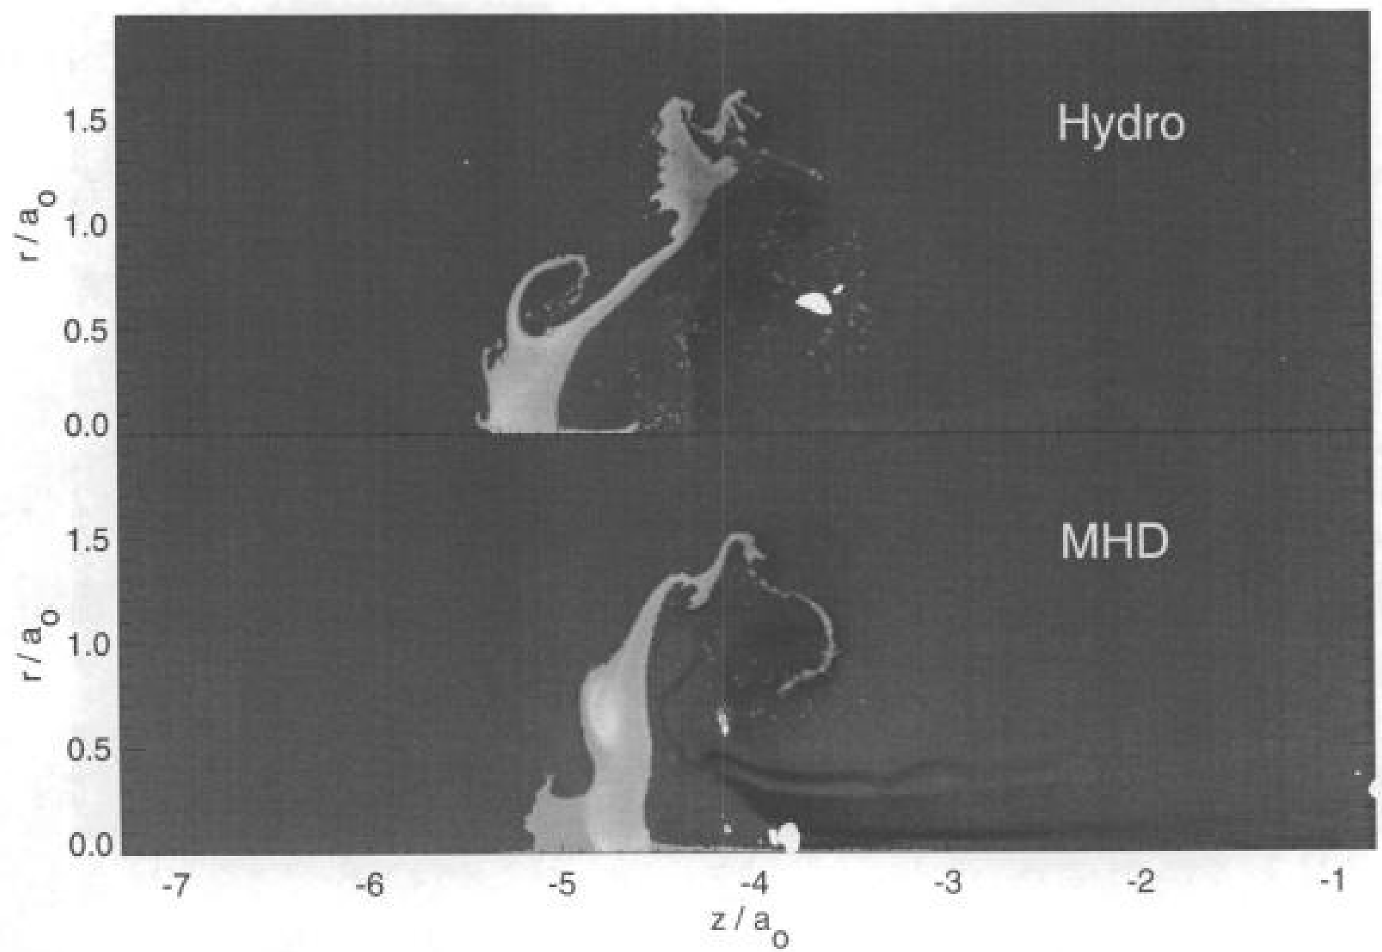

# Effect of Increased Dimensionality

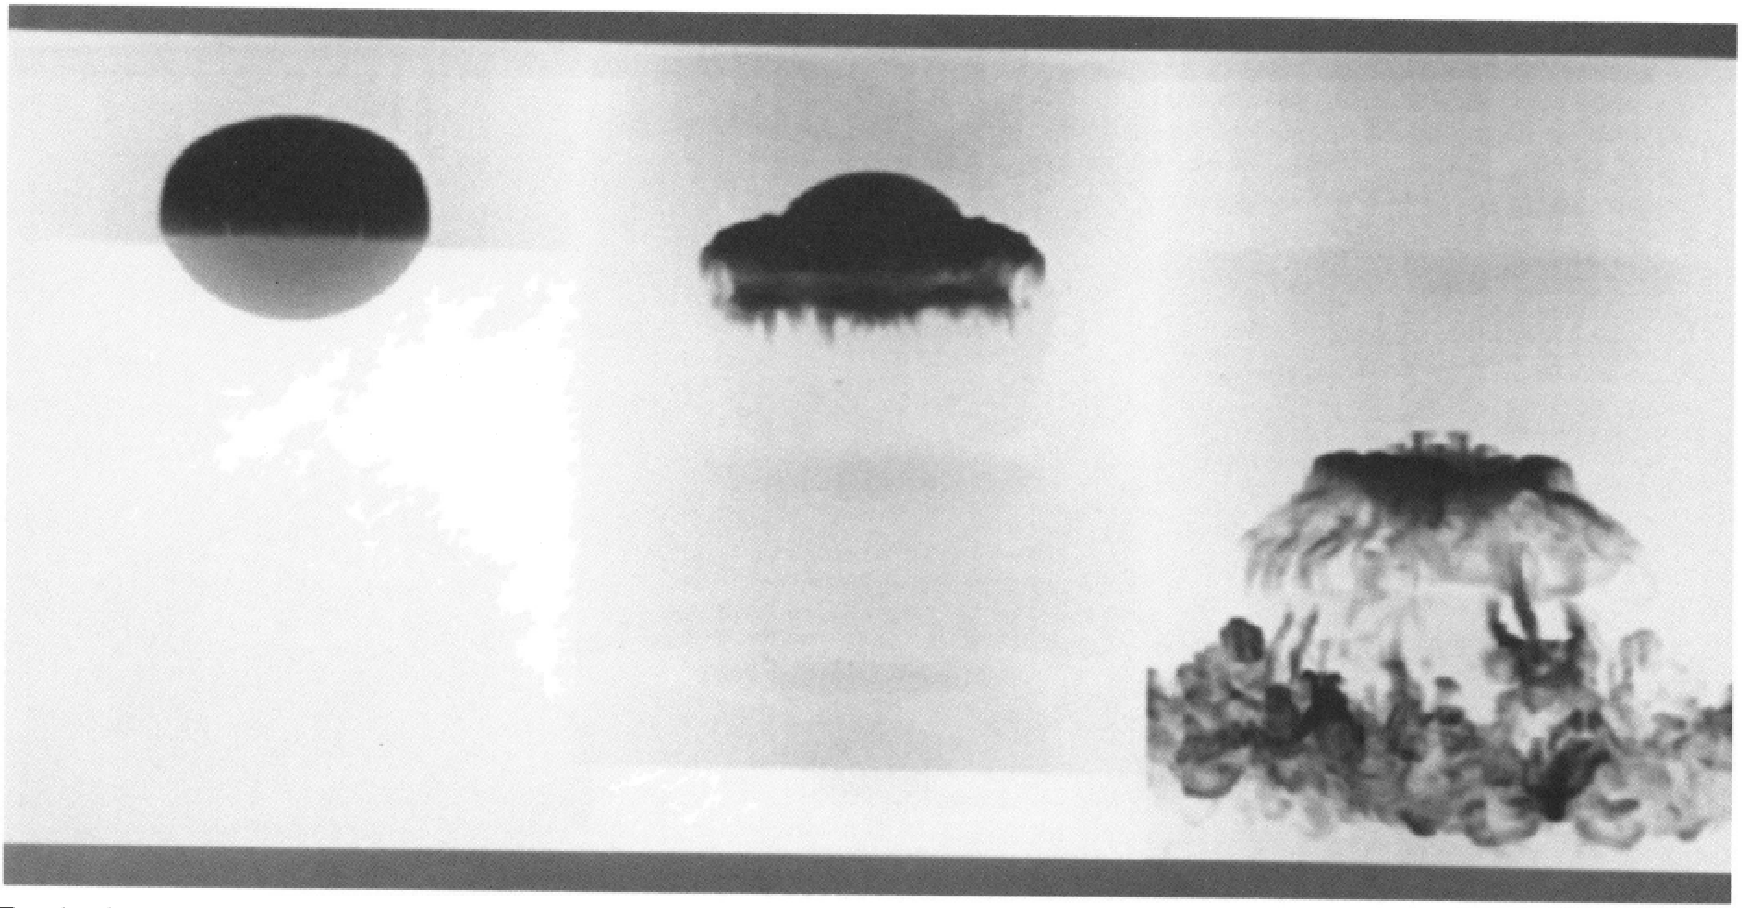

FIG. 1.—A time series of gray-scale images of the three-dimensional distribution of the logarithm of the density taken at $t = 0.5t_c$ (*left*), $t = 2.0t_c$ (*center*), and $t = 4.5t_c$ (*right*).

STONE & NORMAN (see 390, L18)

**Stone and Norman (1992)**

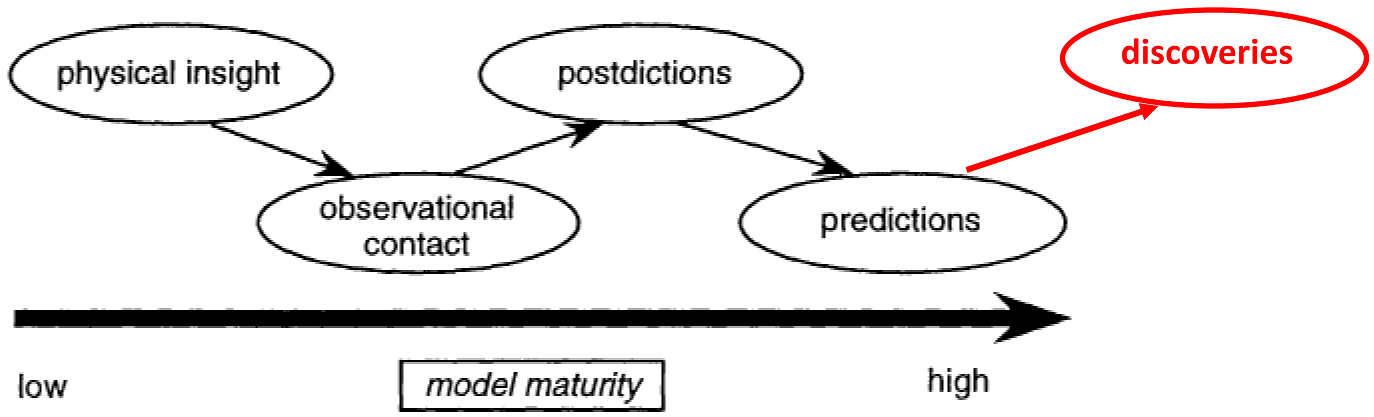


Figure 8. The primary contribution of a numerical model changes as it matures.

# TRENDS IN SUPERCOMPUTERS

# Top500 #3 Cray XT5 Jaguar (Oak Ridge, USA)


37,360 AMD Operton CPUs, 6 cores/CPU ➔224K cores
2.3 Pflops peak speed
3D torus interconnect

# Top500 #2 Tainhe-1A (Tianjin, China)

Hybrid CPU/GPU cluster (XEON/NVIDIA) 186K cores
4.7 Pflops peak speed
Proprietary interconnect

# Top500 #1 Fujitsu K Computer (Riken, Japan)


88,000 Sparc64 CPUs, 8 cores/CPU ➔700K cores
11.28 Pflops peak speed
Tofu interconnect (6D torus = 3D torus of 3D tori)

# US retakes supercomputing crown with 16-petaflops Sequoia; China promises 100 petaflops by 2015

By Sebastian Anthony on June 20, 2012 at 2:05 pm | 3 Comments

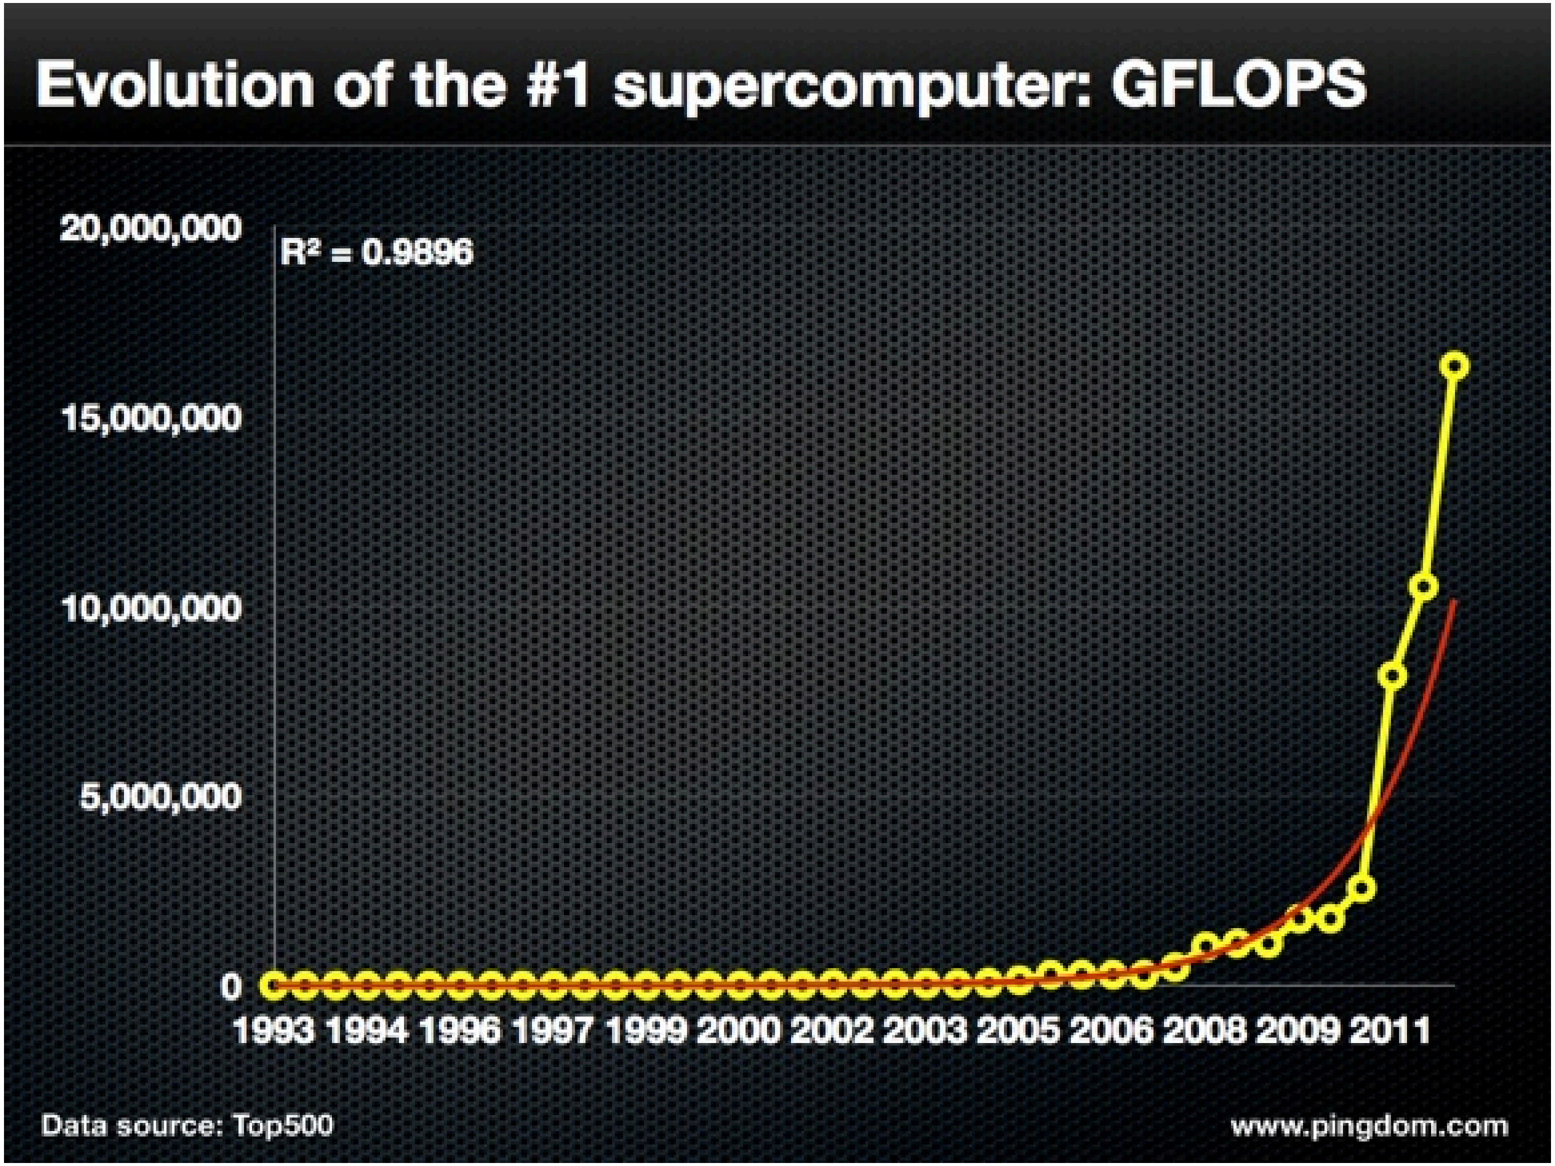
Evolution of the #1 supercomputer: GFLOPS
20,000,000
15,000,000
10,000,000
5,000,000
0
R² = 0.9896
1993 1994 1996 1997 1999 2000 2002 2003 2005 2006 2008 2009 2011
Data source: Top500
www.pingdom.com

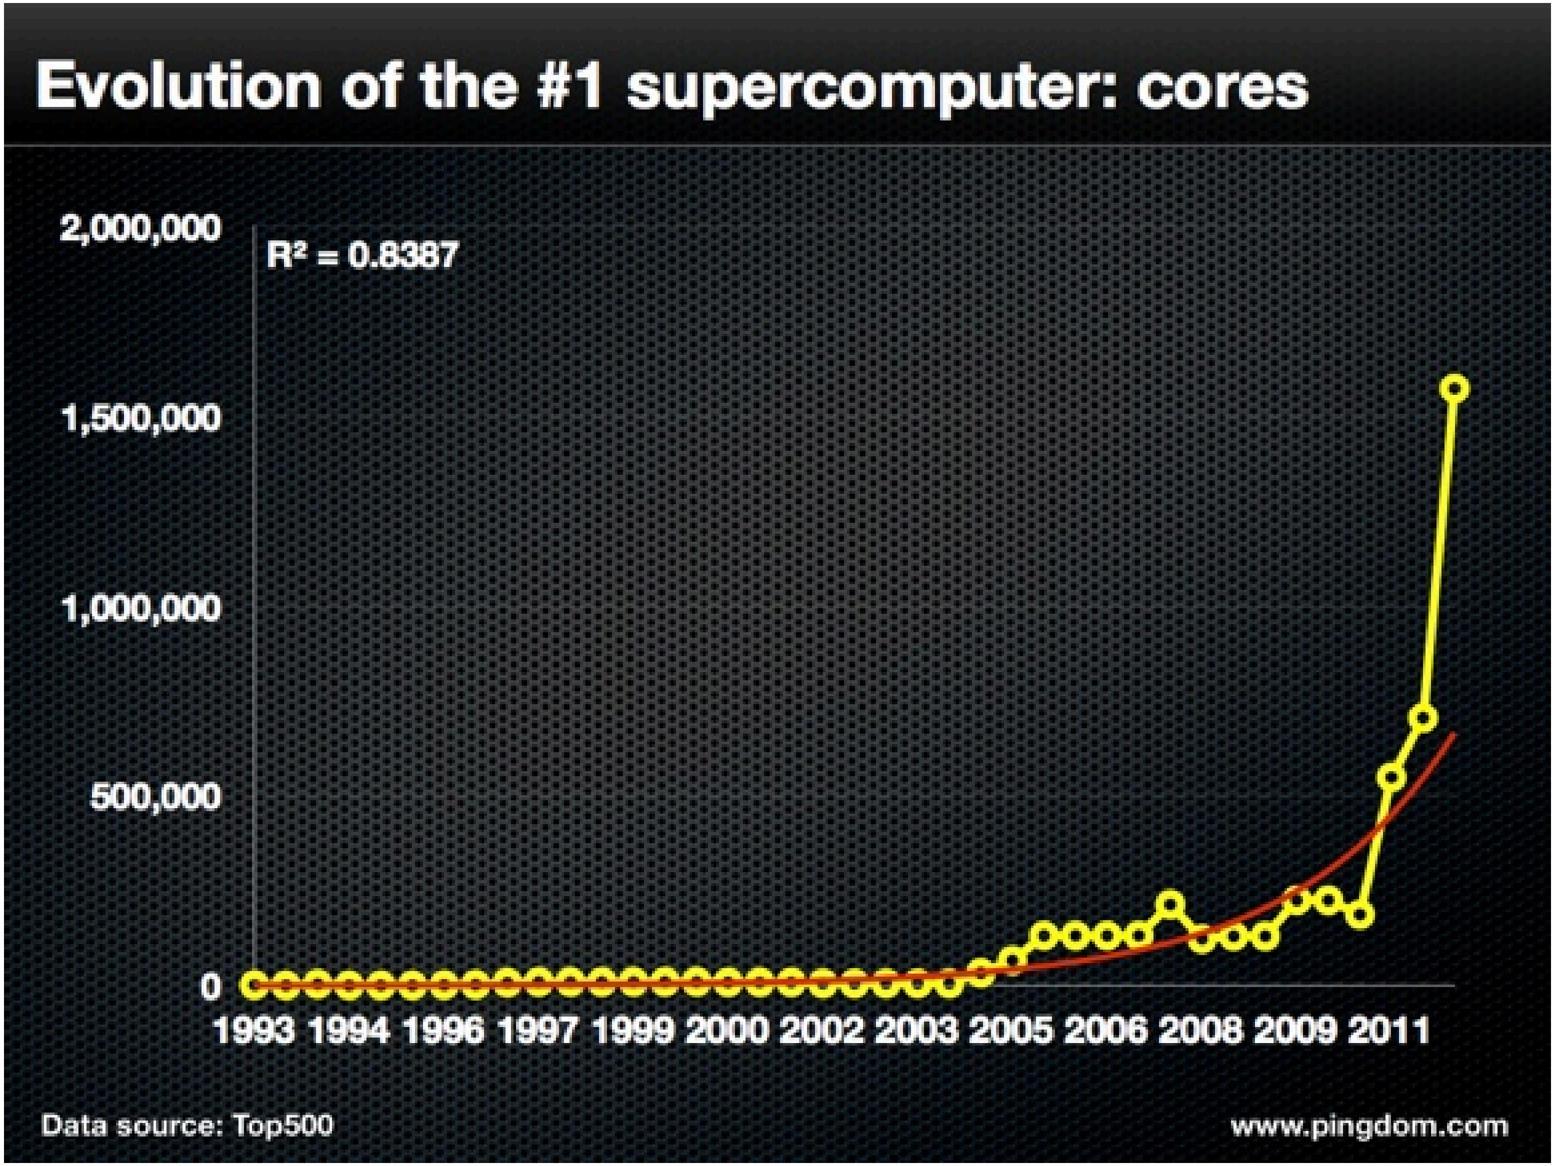
Evolution of the #1 supercomputer: cores
2,000,000
1,500,000
1,000,000
500,000
0
R² = 0.8387
1993 1994 1996 1997 1999 2000 2002 2003 2005 2006 2008 2009 2011
Data source: Top500
www.pingdom.com

# It’s all about the cores

**Cores come in many forms**

- Multicore CPUs
- Many core CPUs
- GPUs

**How you access them is different**

- On the compute node
- Attached devices (GPUs, FPGAs,

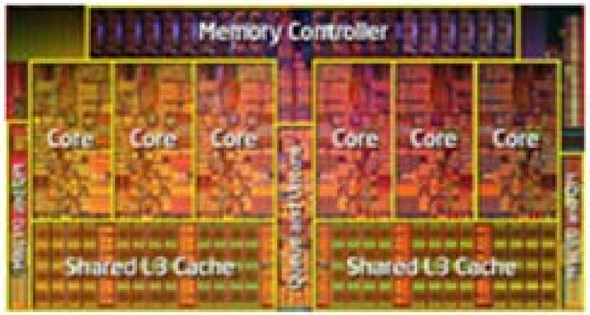


Intel 6-core CPU

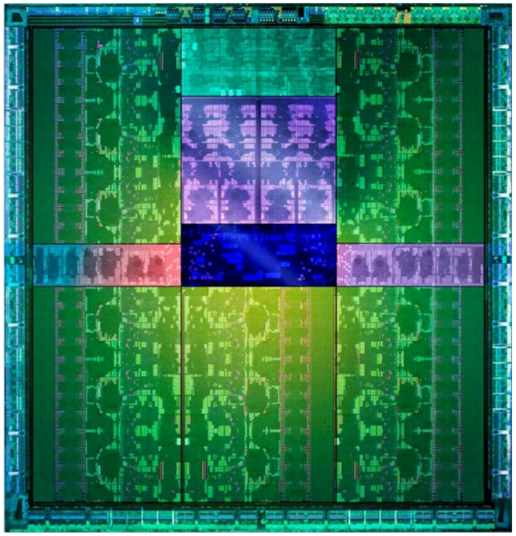

NVIDA GPU

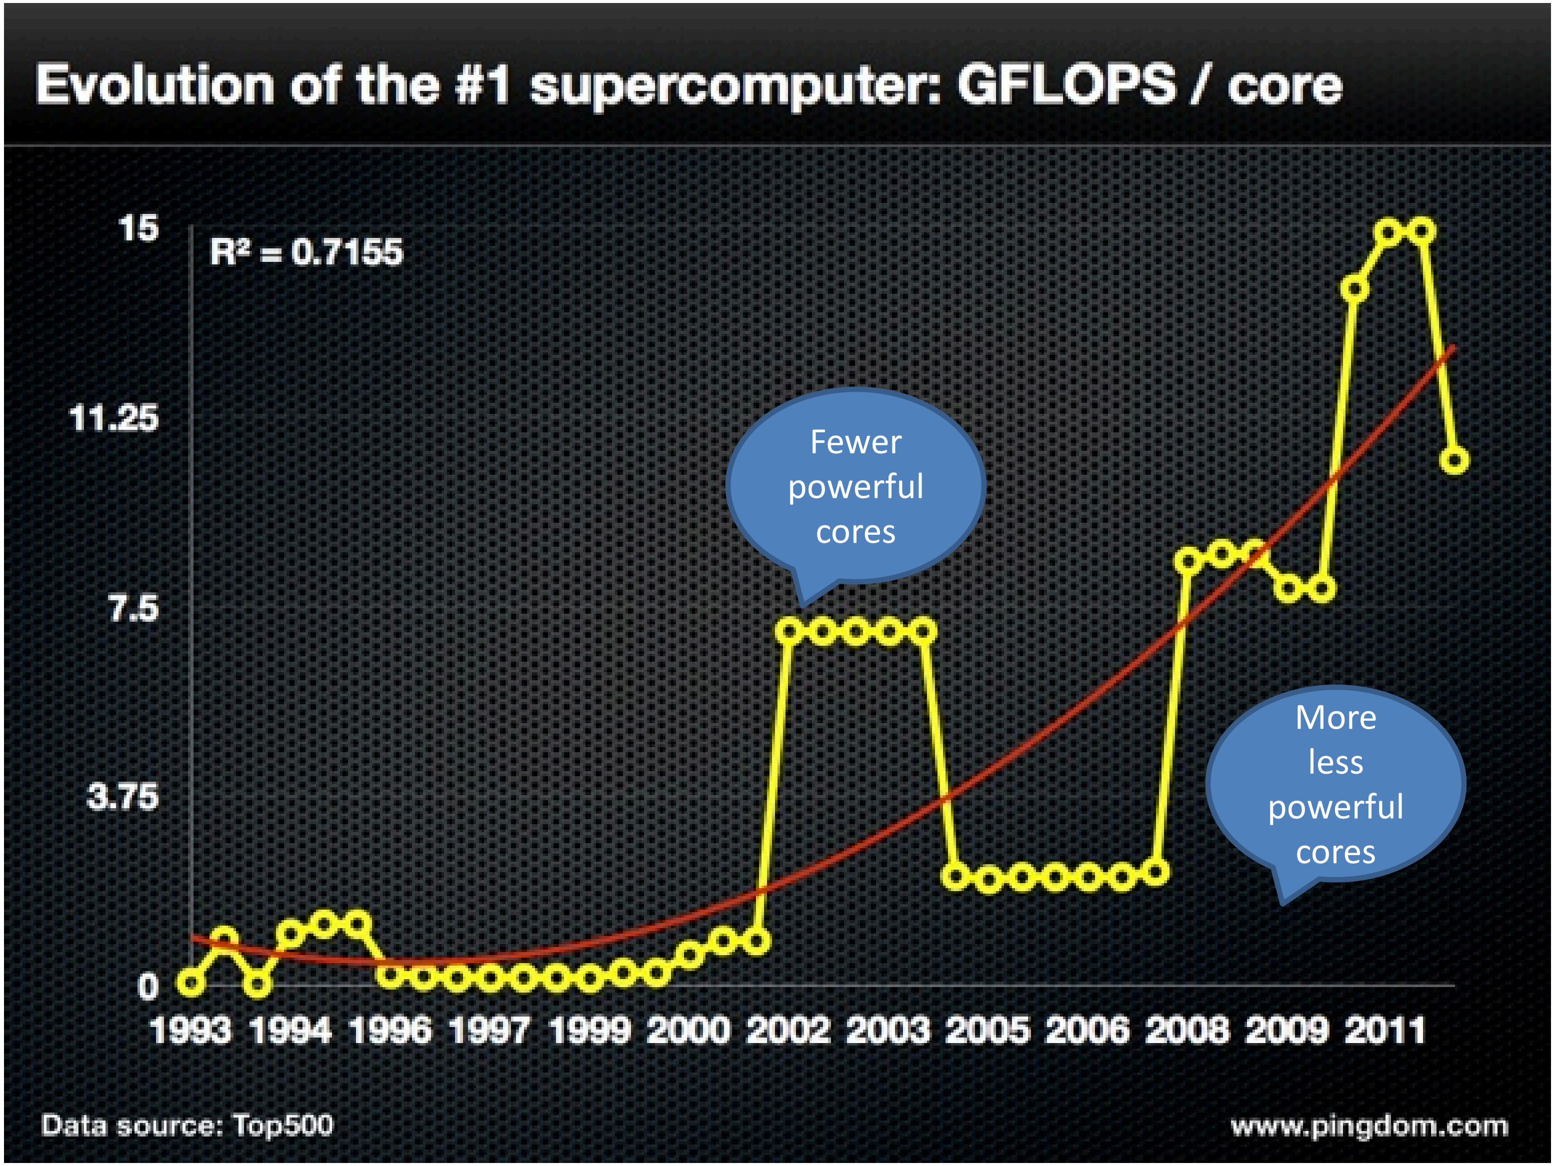
Evolution of the #1 supercomputer: GFLOPS / core
R² = 0.7155
15
11.25
7.5
3.75
0
1993 1994 1996 1997 1999 2000 2002 2003 2005 2006 2008 2009 2011
Fewer powerful cores
More less powerful cores
Data source: Top500
www.pingdom.com

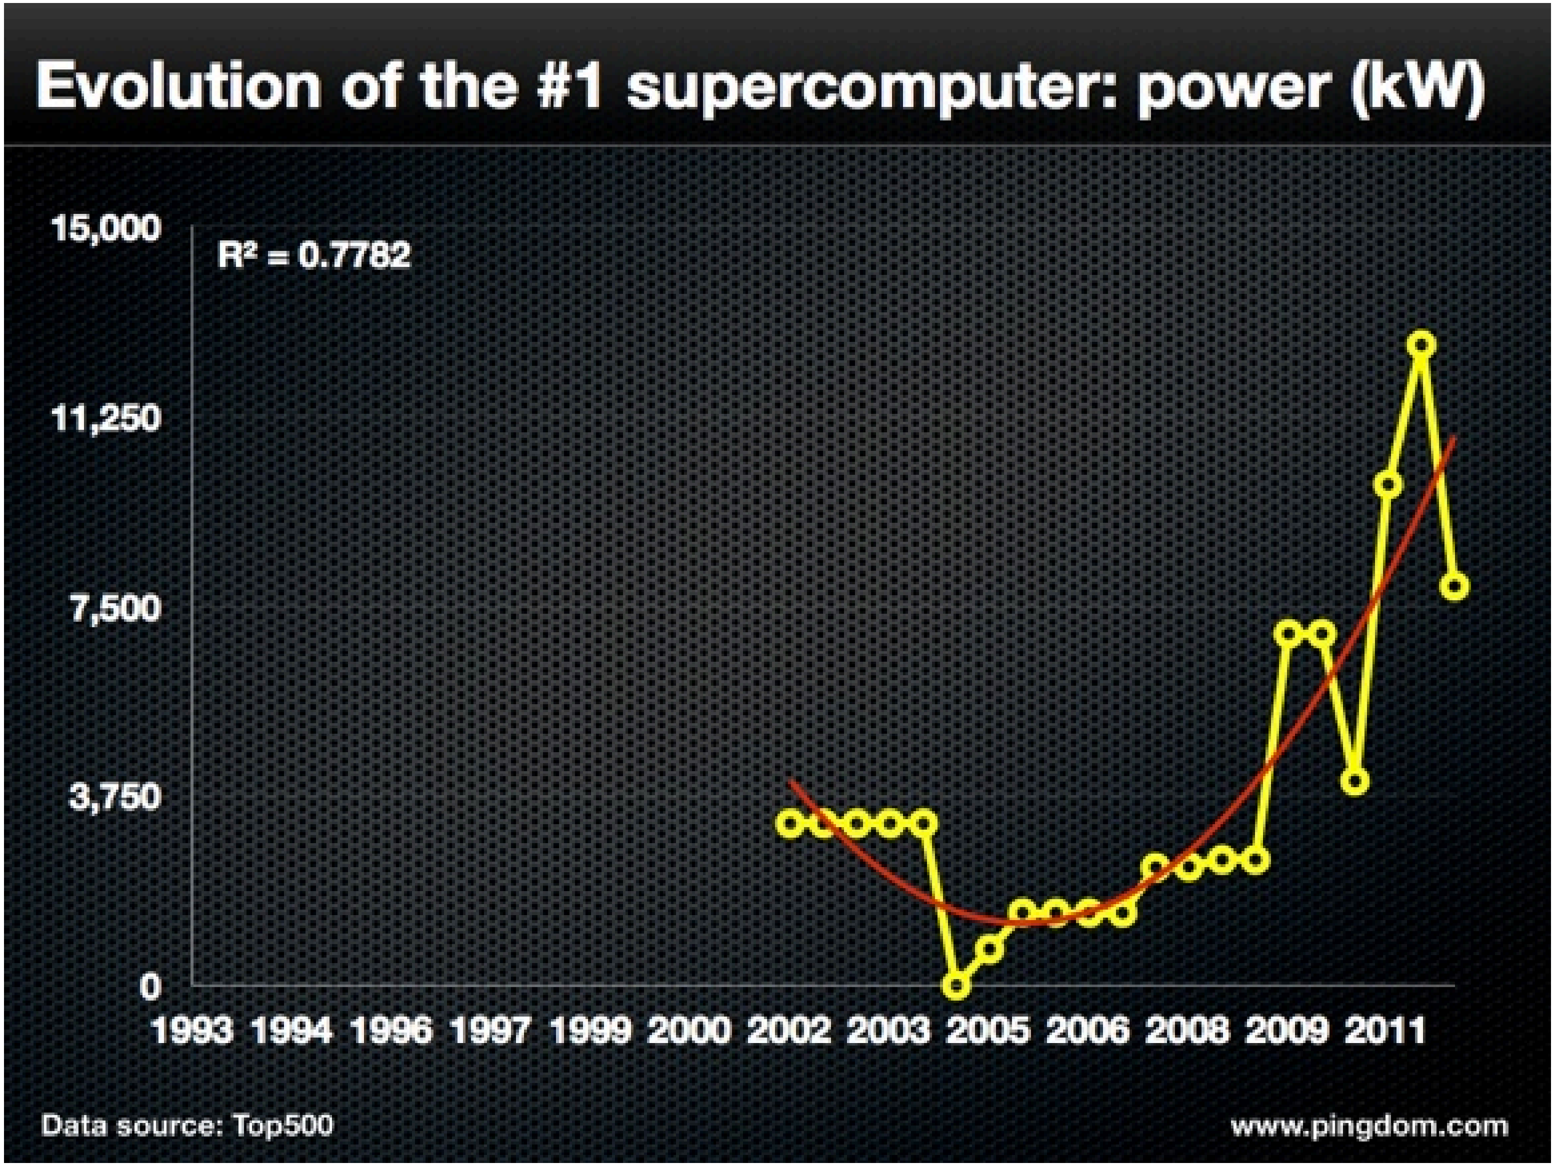
Evolution of the #1 supercomputer: power (kW)
15,000
11,250
7,500
3,750
0
R² = 0.7782
1993
1994
1996
1997
1999
2000
2002
2003
2005
2006
2008
2009
2011
Data source: Top500
www.pingdom.com

# Energy cost to reach Exaflop

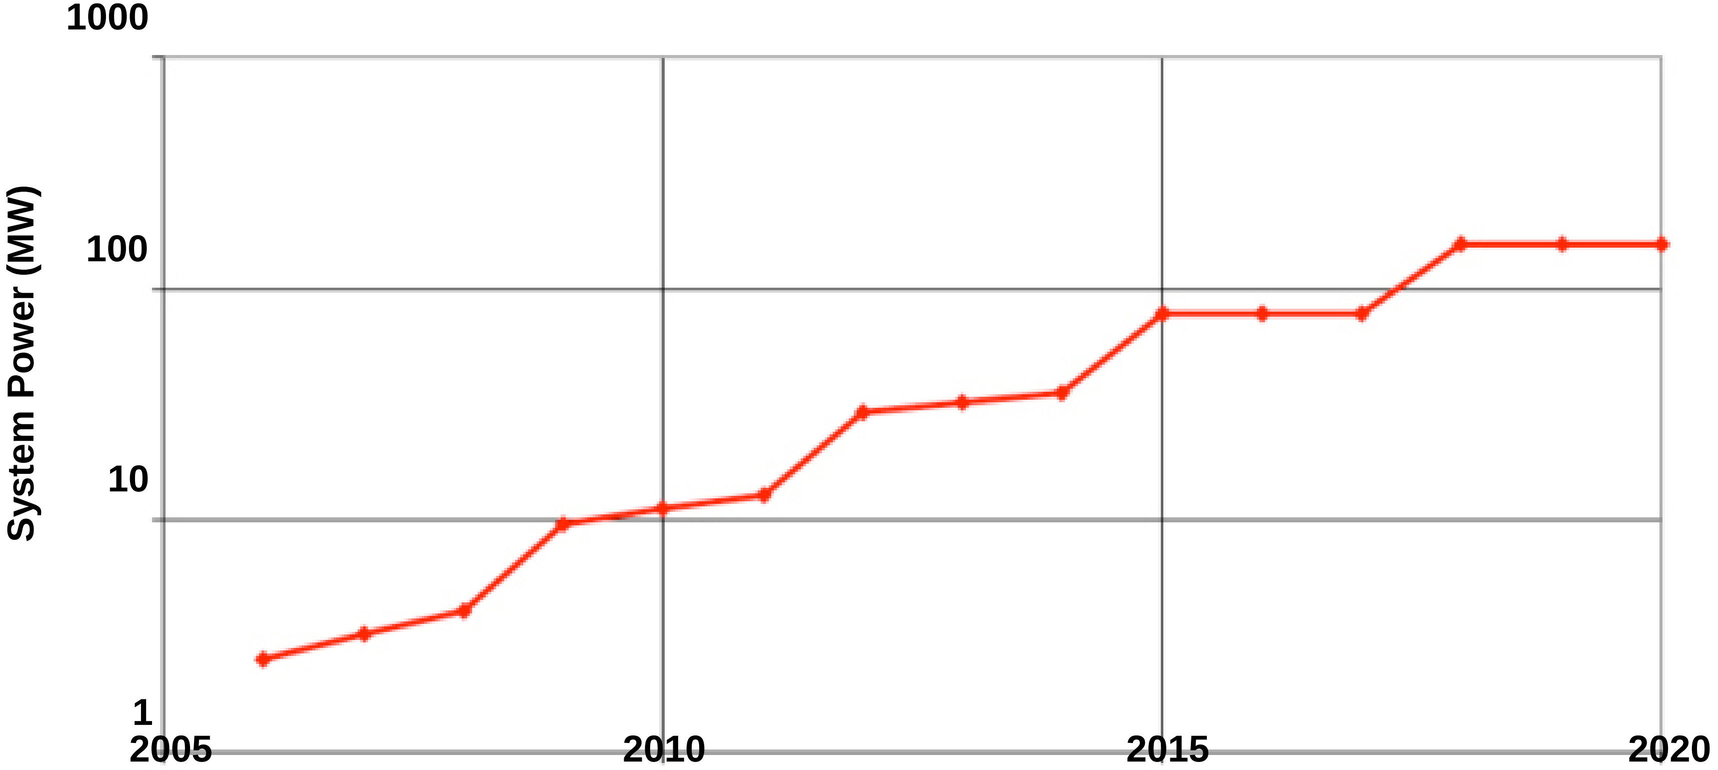


From Peter Kogge,
DARPA Exascale Study

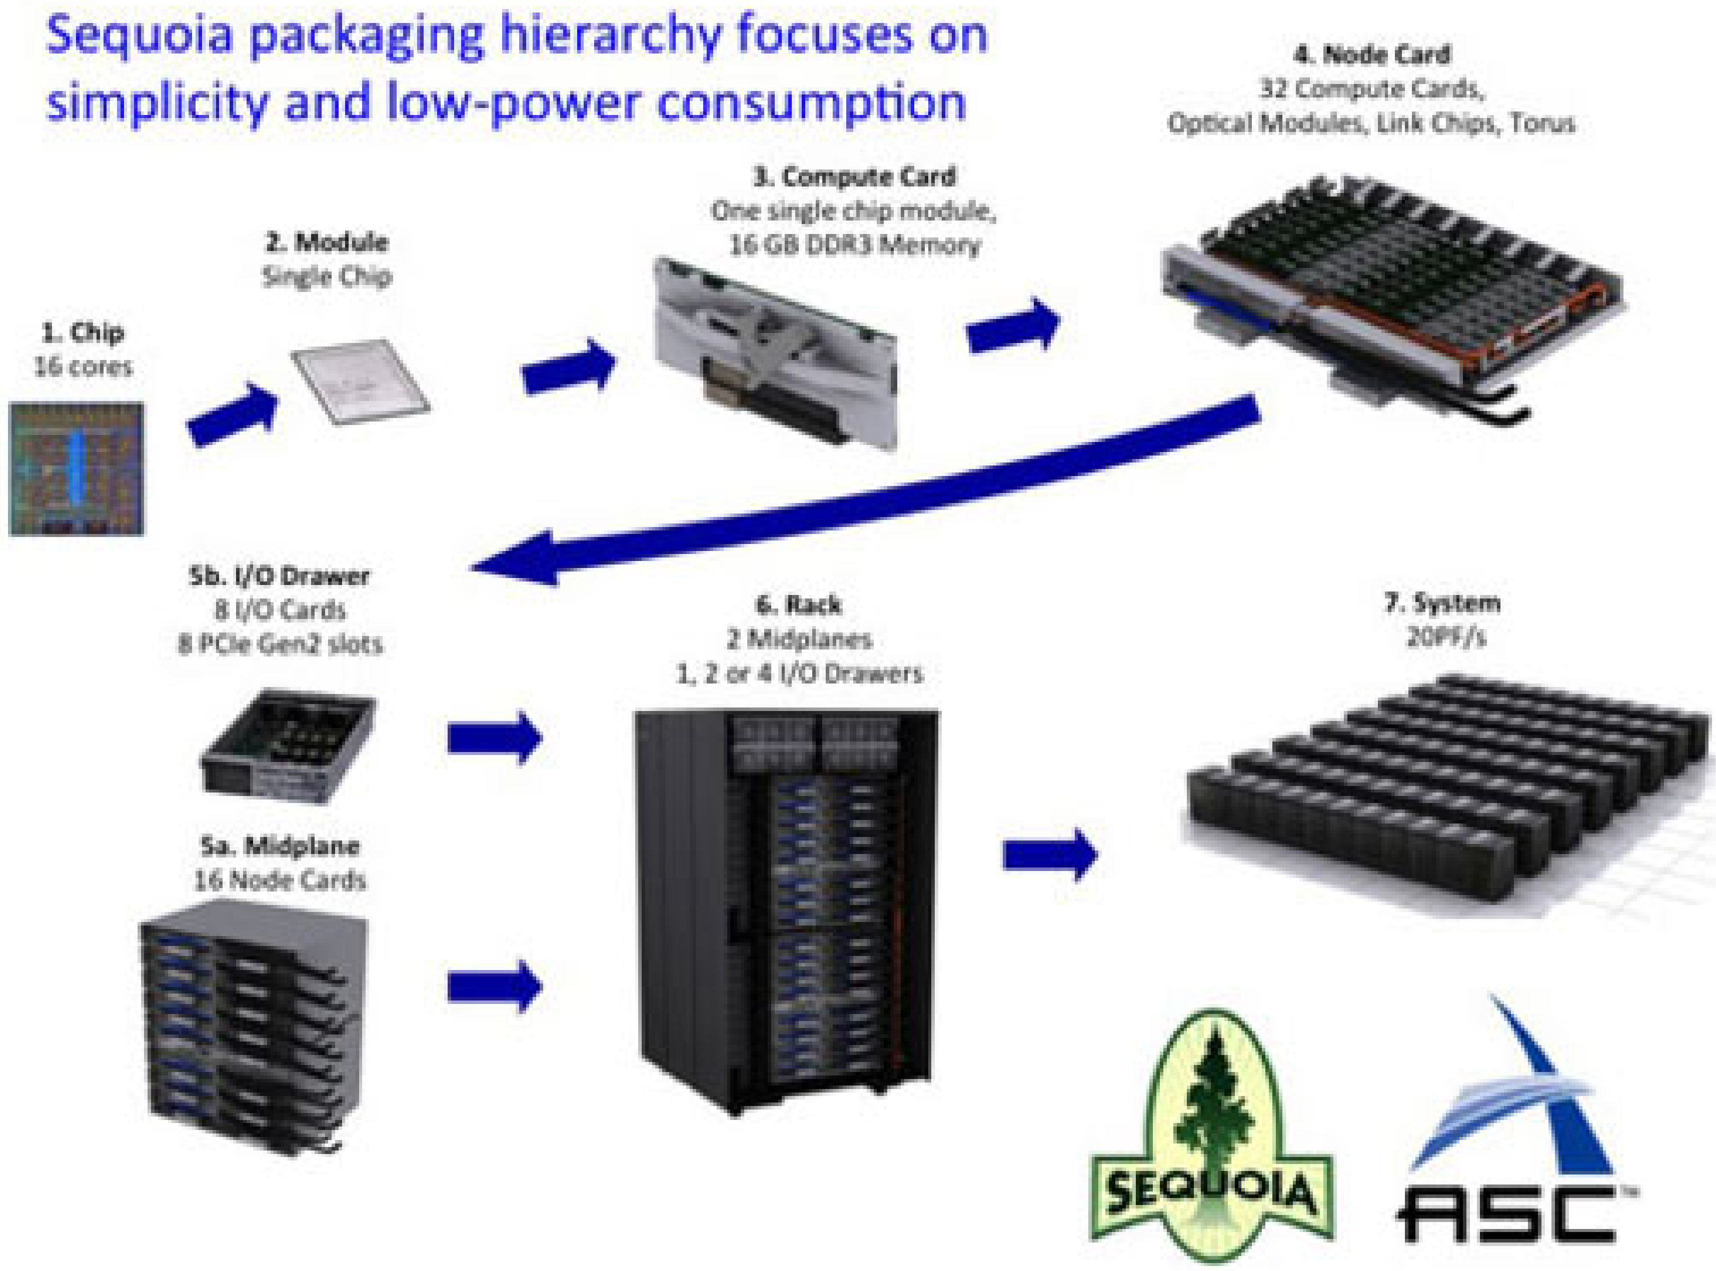
Sequoia packaging hierarchy focuses on simplicity and low-power consumption
1. Chip
16 cores
2. Module
Single Chip
3. Compute Card
One single chip module,
16 GB DDR3 Memory
4. Node Card
32 Compute Cards,
Optical Modules, Link Chips, Torus
5b. I/O Drawer
8 I/O Cards
8 PCIe Gen2 slots
5a. Midplane
16 Node Cards
6. Rack
2 Midplanes
1, 2 or 4 I/O Drawers
7. System
20PF/s
SEQUOIA
ASC

# TRENDS IN SUPER DATA

# The Data Deluge in Science

High energy physics

genomic
medicine

earth sciences

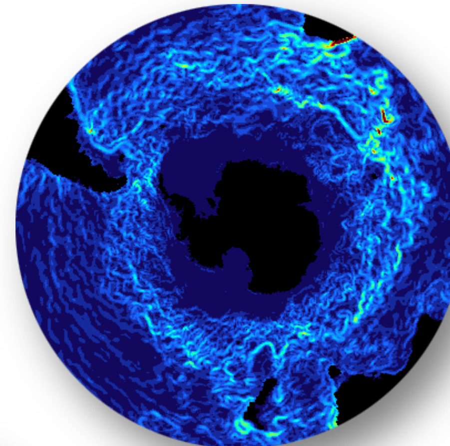

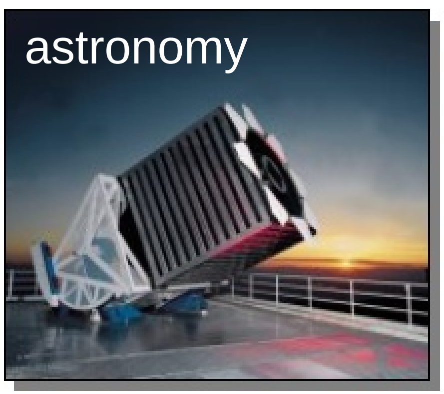



social sciences

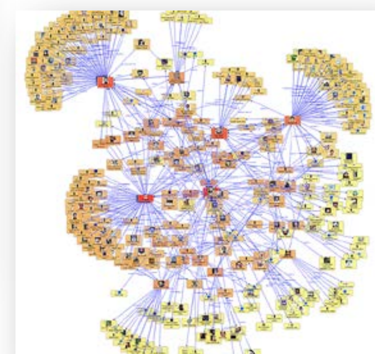

# *Why is scientific research becoming data-intensive?*

- **Capacity to generate, store, transmit digital data is growing exponentially**
  - digital sensors follow Moore's Law too
- **New fields of science driven by high-throughput gene sequencers, CCDs, and sensor nets**
  - genomics, proteomics, and metagenomics
  - astronomical sky surveys
  - seismic, oceanographic, ecological "observatories"
- **Emergence of the Internet (wired and wireless)**
  - remote access to data archives and collaborators
- **Supercomputers are prodigious data generators**

# *Cosmological Simulation Growth (M. Norman)*

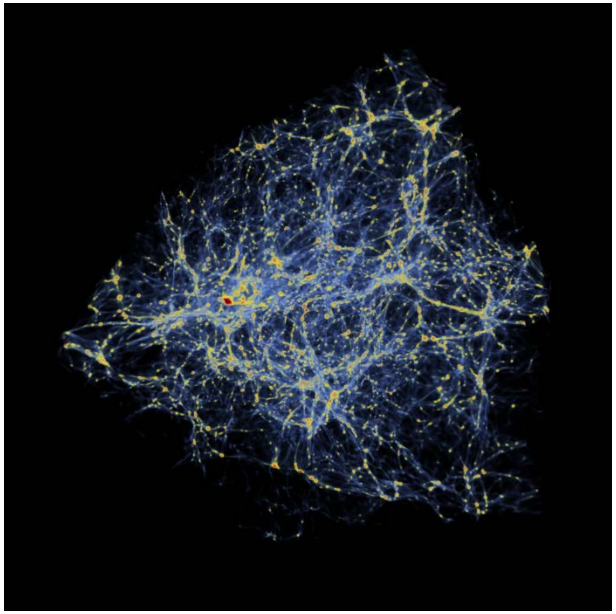

| Year | Ngrid | Ncell (B) | Ncpu | Machine |
|---|---|---|---|---|
| 1994 | $512^3$ | 1/8 | 512 | TMC CM5 |
| 2003 | $1024^3$ | 1 | 512 | IBM SP3 |
| 2006 | $2048^3$ | 8 | 2048 | IBM SP3 |
| 2009 | $4096^3$ | 64 | 16K | Cray XT5 |
| 2010 | $6400^3$ | 262 | 93K | Cray XT5 |

- **Increase of >2000 in problem size in 16 years**
- **2x every 1.5 years ➔ Moore's law for supercomputers**

# Coping with the data deluge

- **Density of storage media keeping pace with Moore's law, but *not I/O rates***
  - Time to process exponentially growing amounts of data is growing exponentially
  - Latency for random access limited by disk read head speed
  - Key insight: flash SSD reduces read latency by 100x

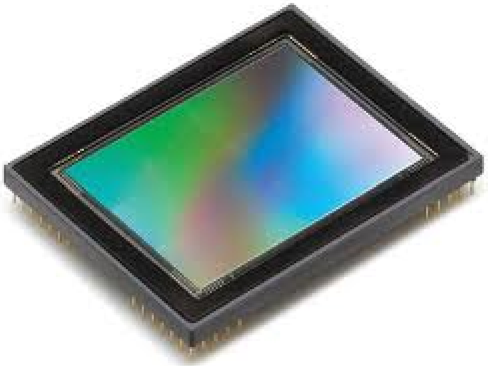

# *2012: Era of Data Supercomputing Begins*

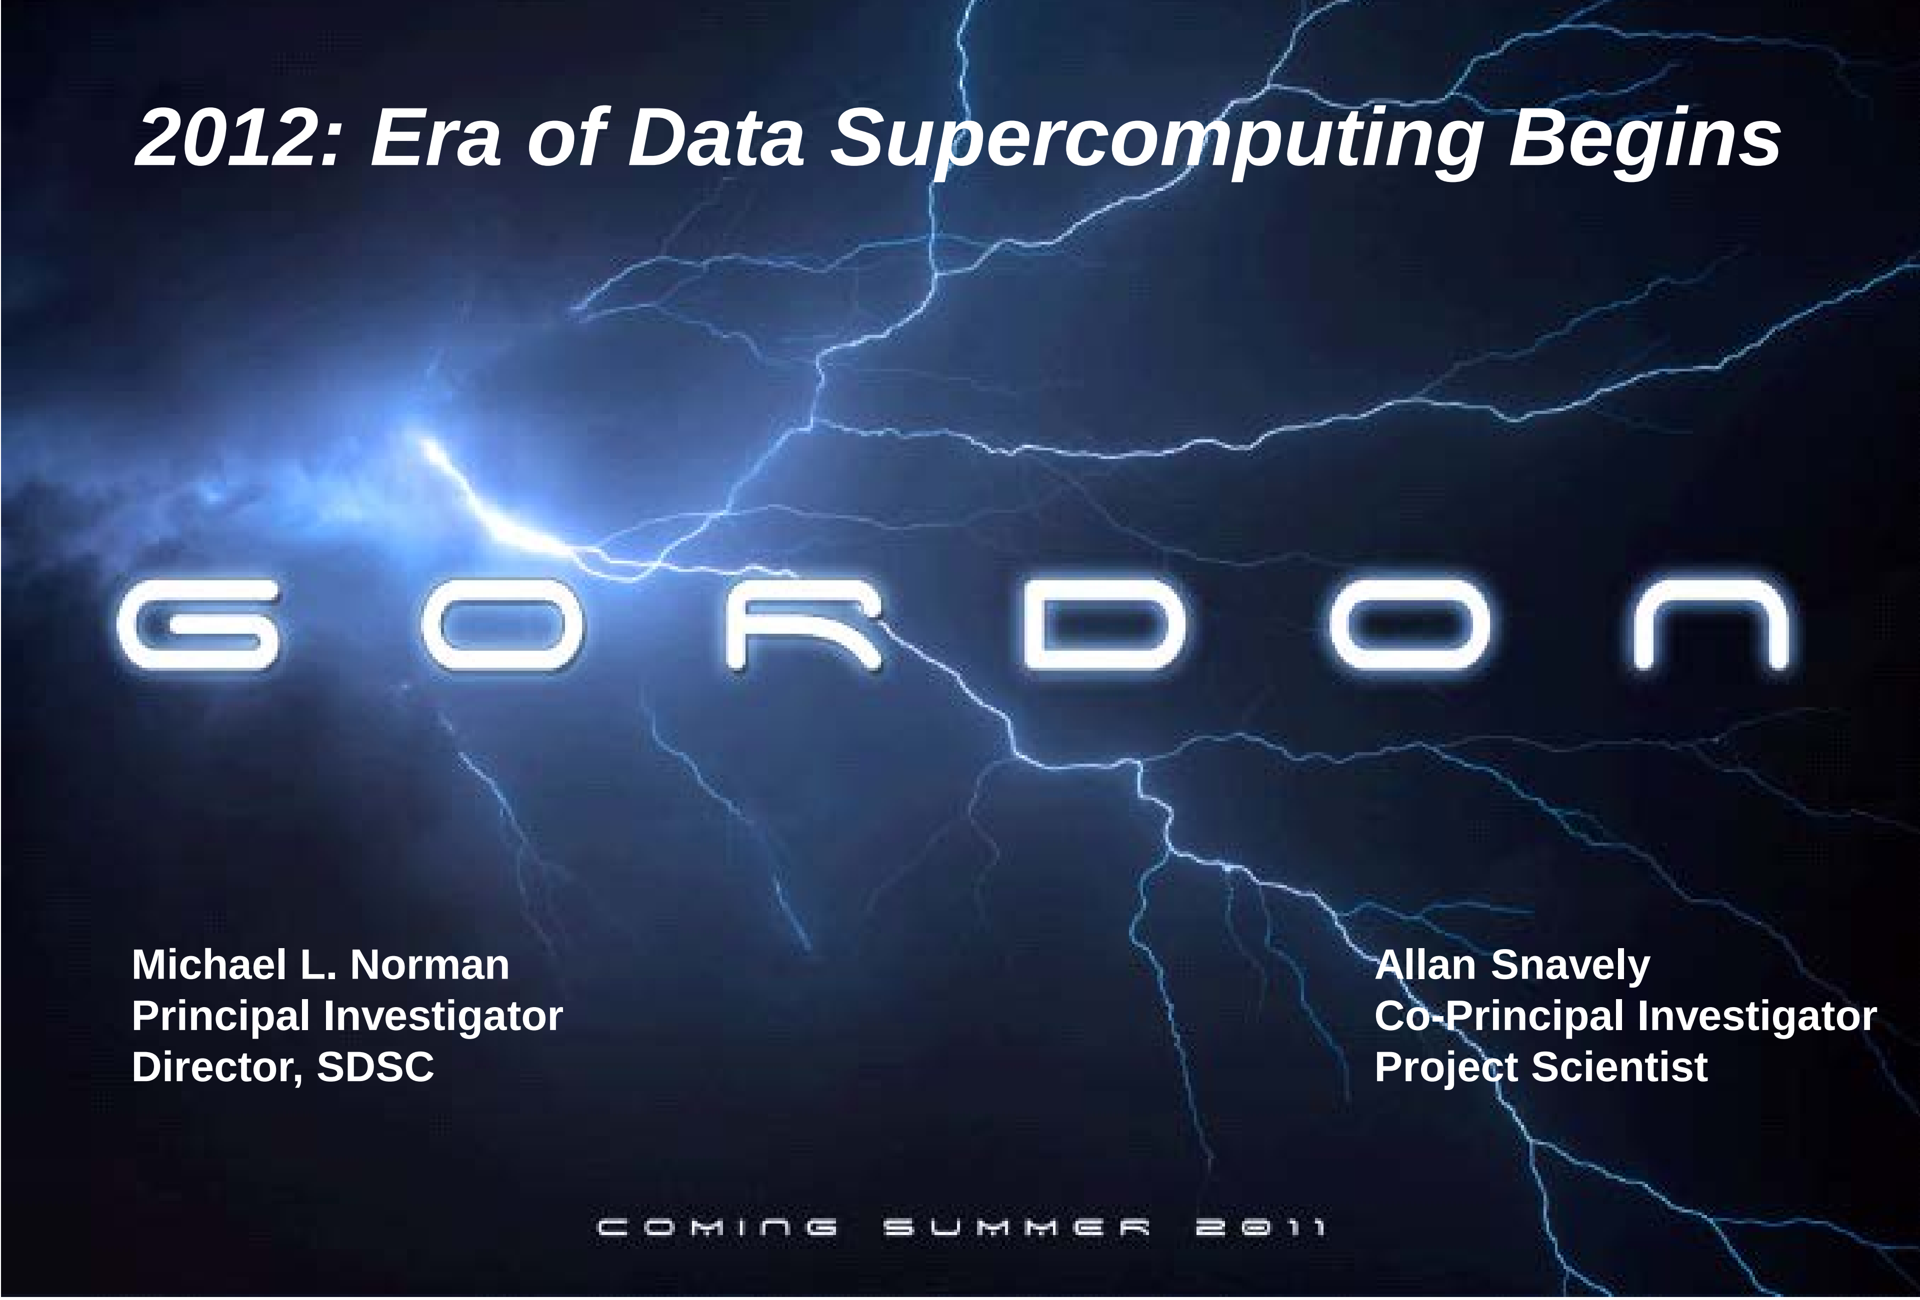


**Michael L. Norman**
**Principal Investigator**
**Director, SDSC**

**Allan Snavely**
**Co-Principal Investigator**
**Project Scientist**

# What is Gordon?

- **A "data-intensive" supercomputer based on SSD flash memory and virtual shared memory**
  - *Emphasizes MEM and IO over FLOPS*
- **A system designed to accelerate access to massive amounts of data being generated in all fields of science, engineering, medicine, and social science**
- **Went into production Feb. 2012**
- **Funded by the National Science Foundation and available as to US researchers and their foreign collaborators thru XSEDE**

# *2012: First Academic Data-Supercomputer "Gordon"*

- **16K cores/340 TF**
- **64 TB DRAM**
- **300 TB of flash SSD memory**
- **software shared memory "supernodes"**
- **Designed for "Big Data Analytics"**

# Gordon Design: Two Driving Ideas

- **Observation #1: Data keeps getting further away from processor cores ("red shift")**
  - Do we need a new level in the memory hierarchy?

- **Observation #2: Data-intensive applications may be serial and difficult to parallelize**
  - Wouldn't a large, shared memory machine be better from the standpoint of researcher productivity?
  - ➔Rapid prototyping of new approaches to data analysis

# The Memory Hierarchy of a Typical Supercomputer

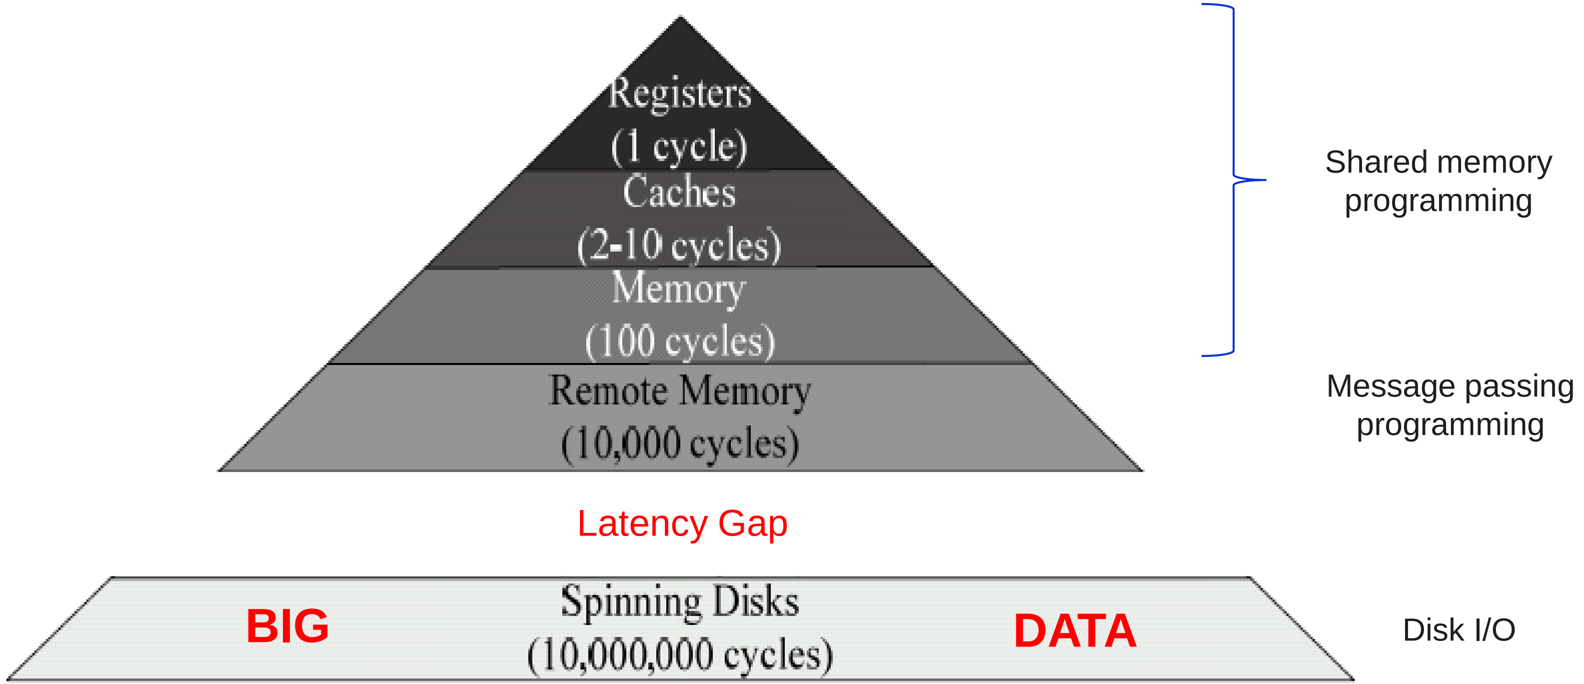

# The Memory Hierarchy of Gordon

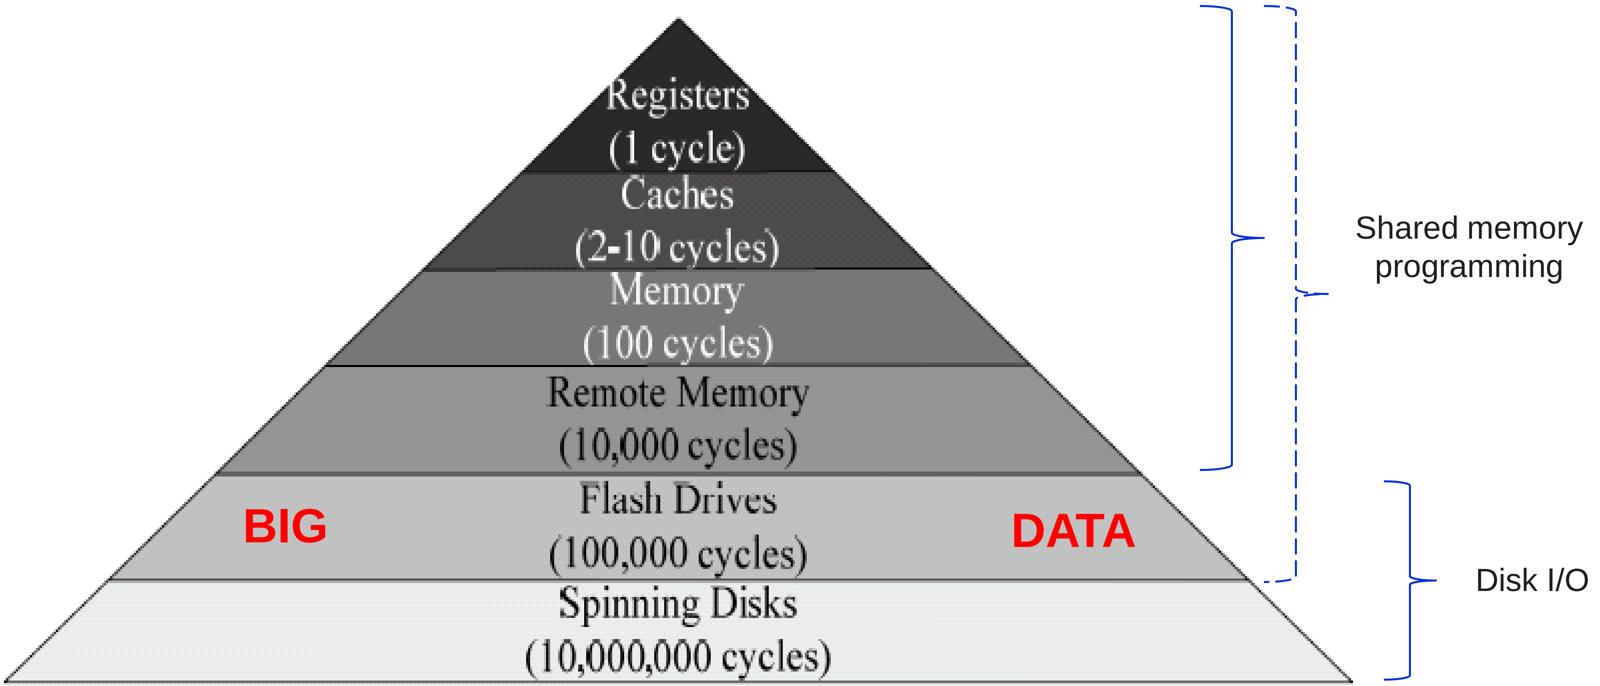

# Gordon 32-way Supernode

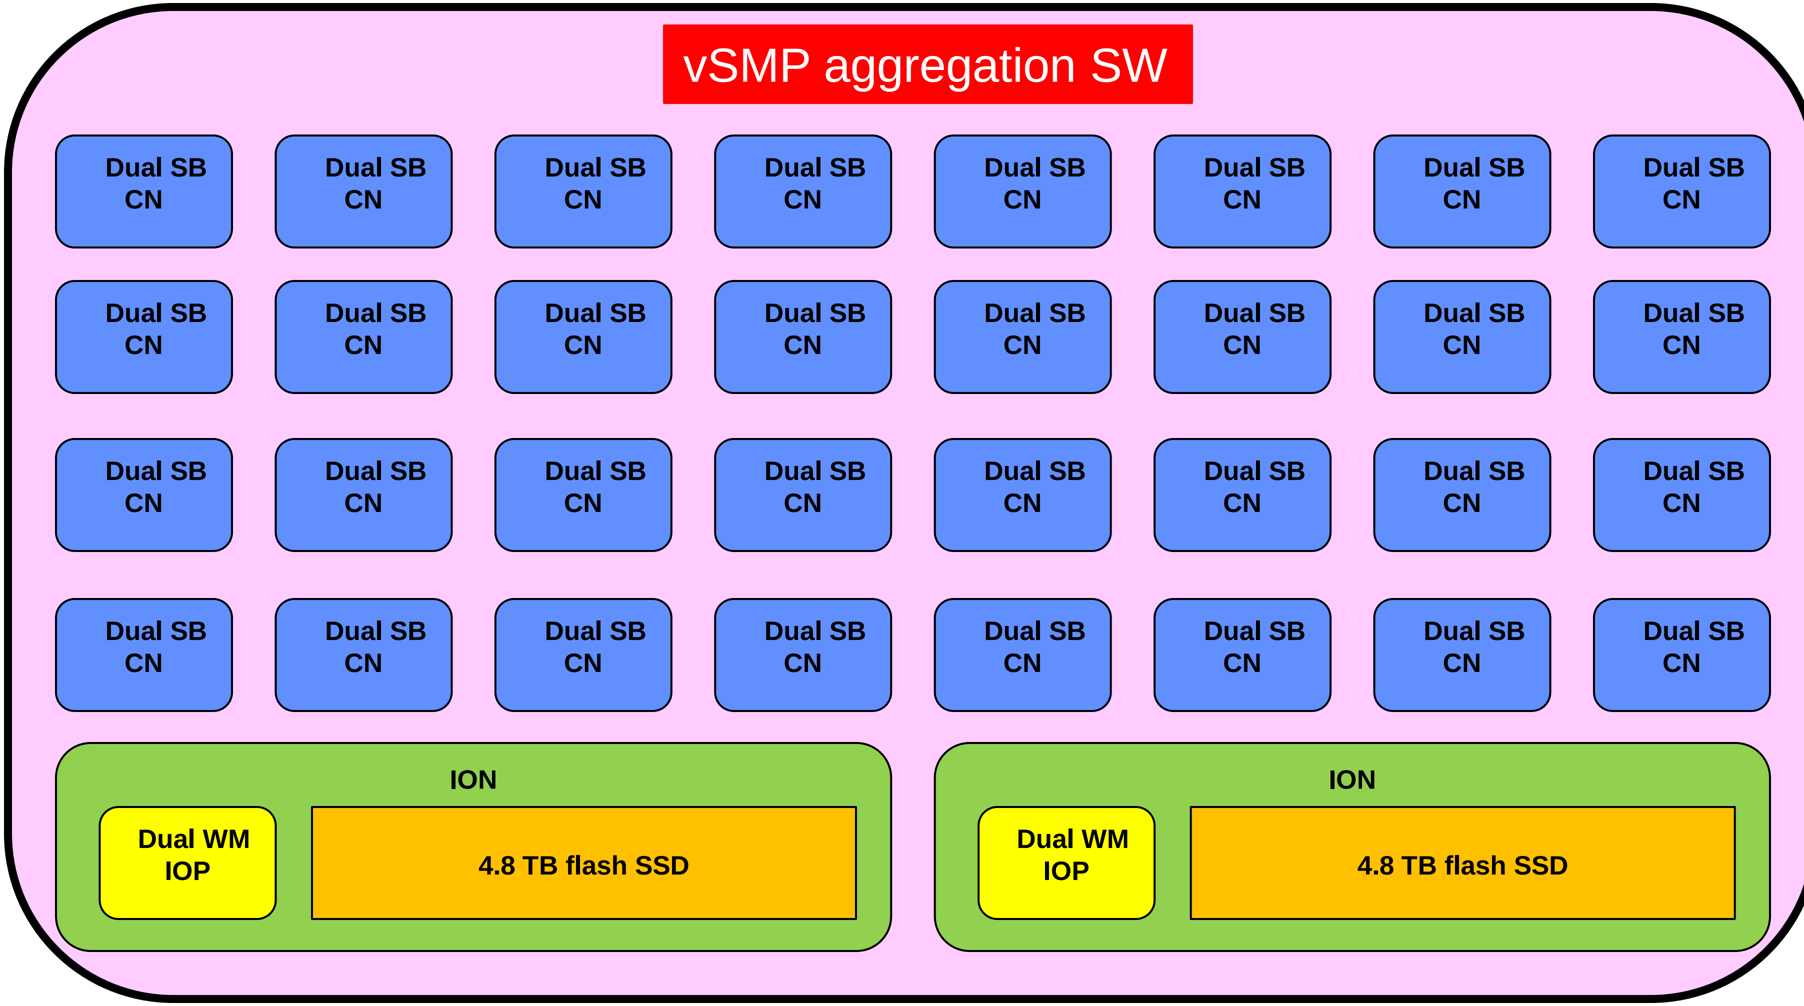

# Gordon 32-way Supernode

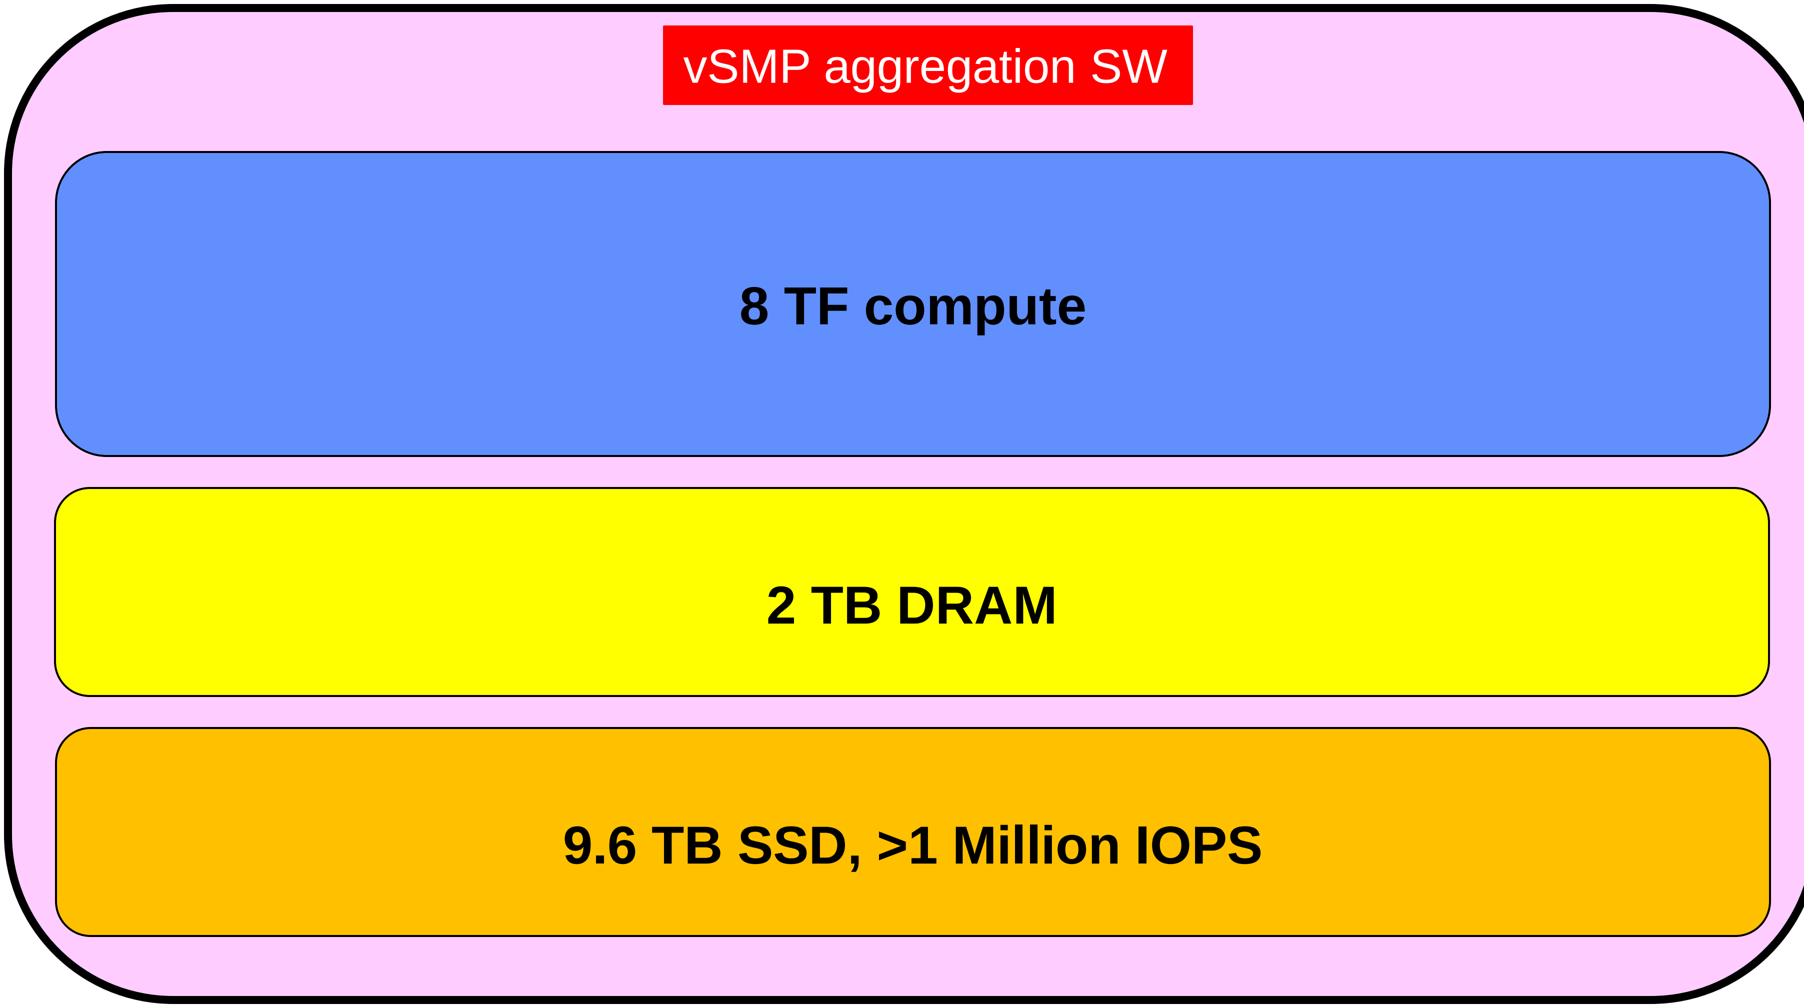

# Gordon Architecture: Full Machine

- **32 supernodes = 1024 compute nodes**
- **Dual rail QDR Infiniband network**
  - 3D torus (4x4x4)
- **4 PB rotating disk parallel file system**
  - >100 GB/s

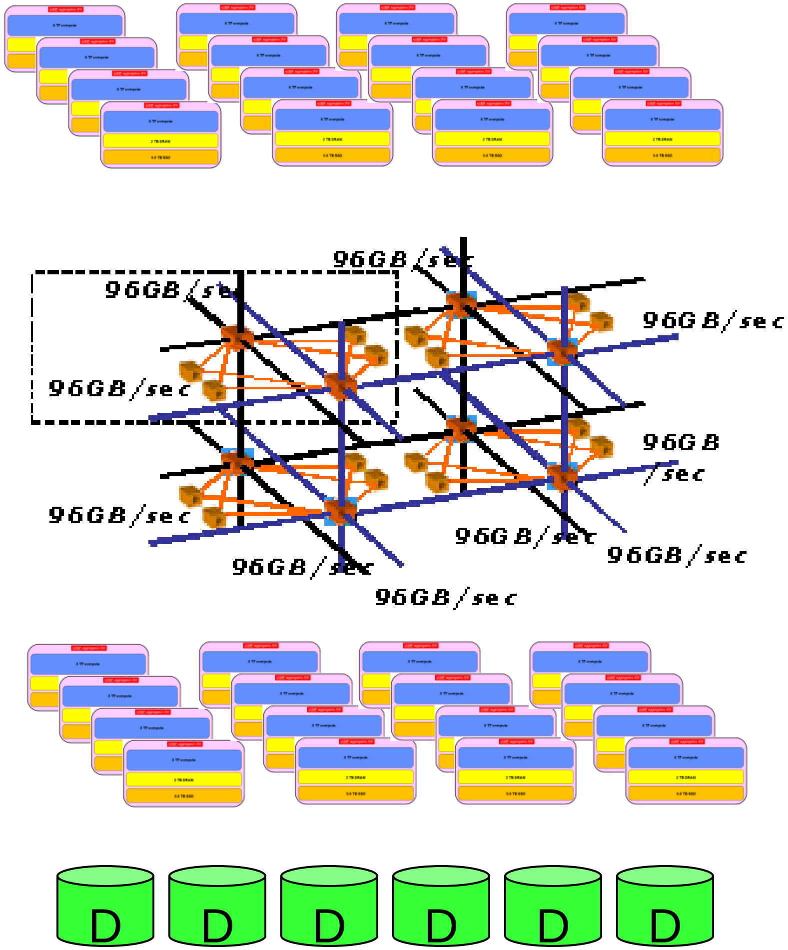

# Probing Cosmic Renaissance by Supercomputer

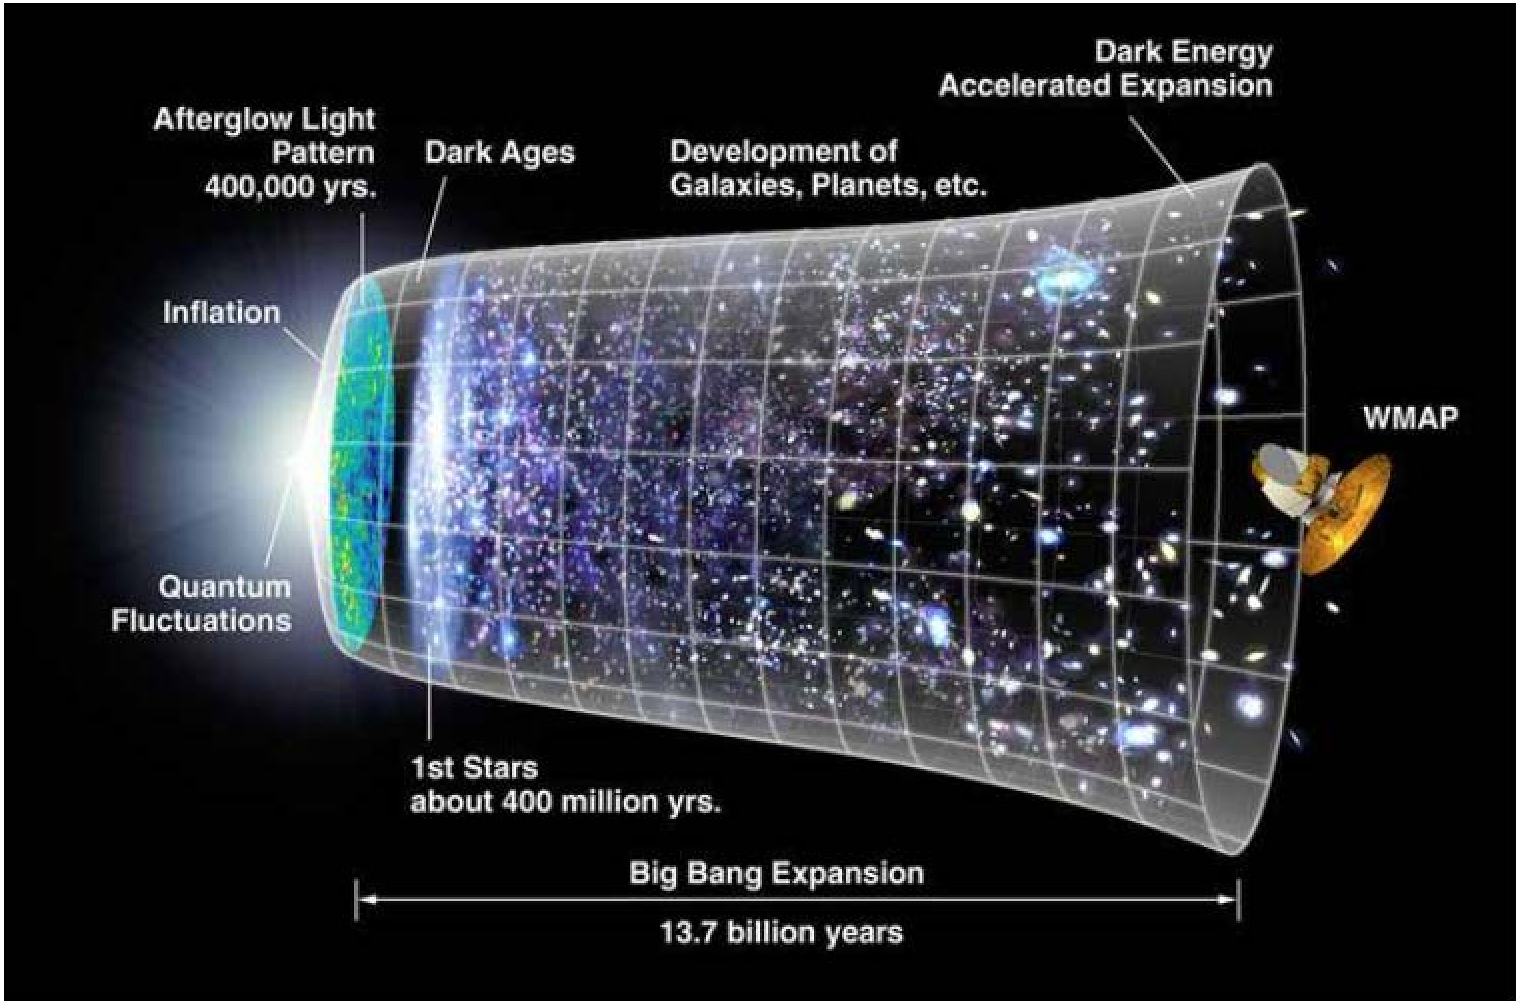

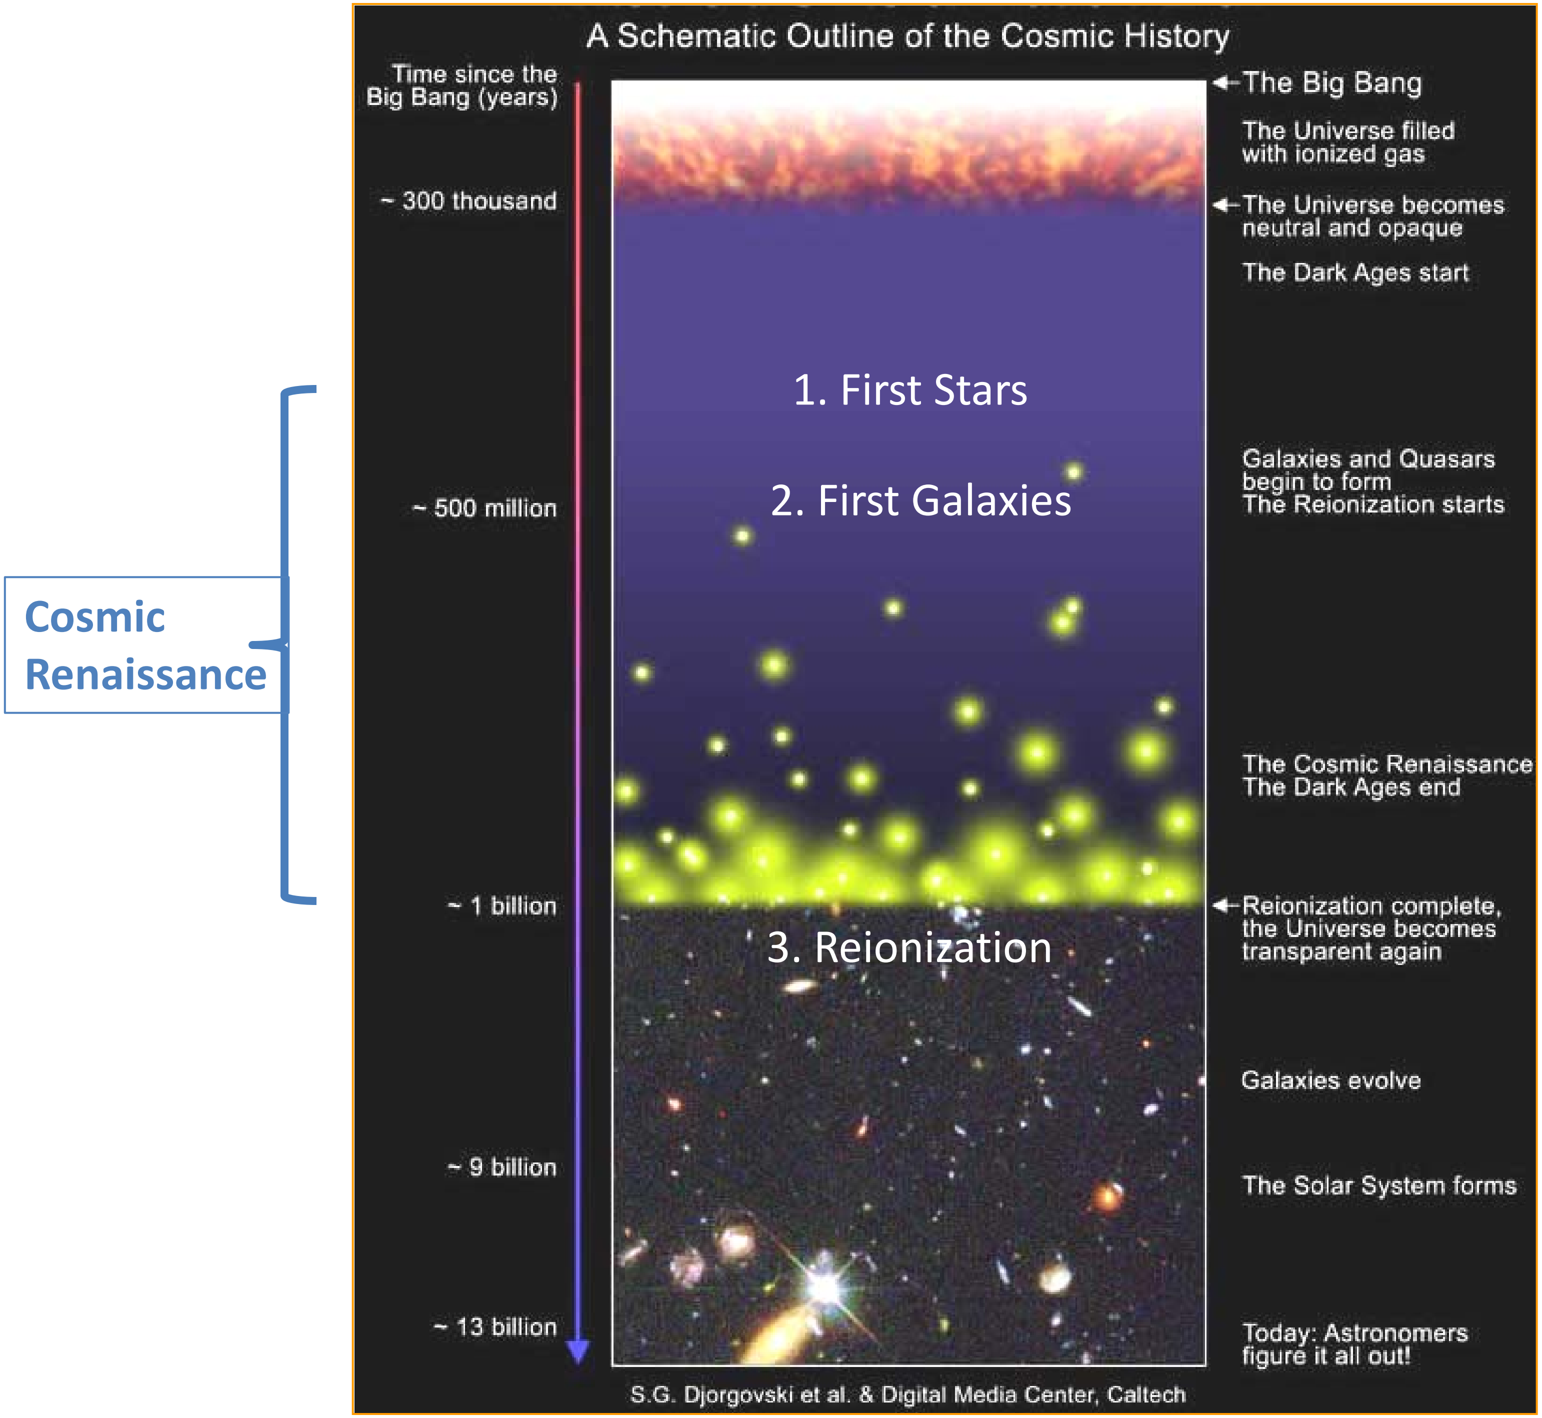

A Schematic Outline of the Cosmic History
Time since the Big Bang (years)
~ 300 thousand
~ 500 million
~ 1 billion
~ 9 billion
~ 13 billion
1. First Stars
2. First Galaxies
3. Reionization
The Big Bang
The Universe filled with ionized gas
The Universe becomes neutral and opaque
The Dark Ages start
Galaxies and Quasars begin to form
The Reionization starts
The Cosmic Renaissance
The Dark Ages end
Reionization complete, the Universe becomes transparent again
Galaxies evolve
The Solar System forms
Today: Astronomers figure it all out!
S.G. Djorgovski et al. & Digital Media Center, Caltech
Cosmic Renaissance

# Simulating the first generation of stars in the universe

- If large objects form via *mergers* of smaller objects……..*Where did it all begin*?
  - *What kind of object is formed*?
  - What is their significance?

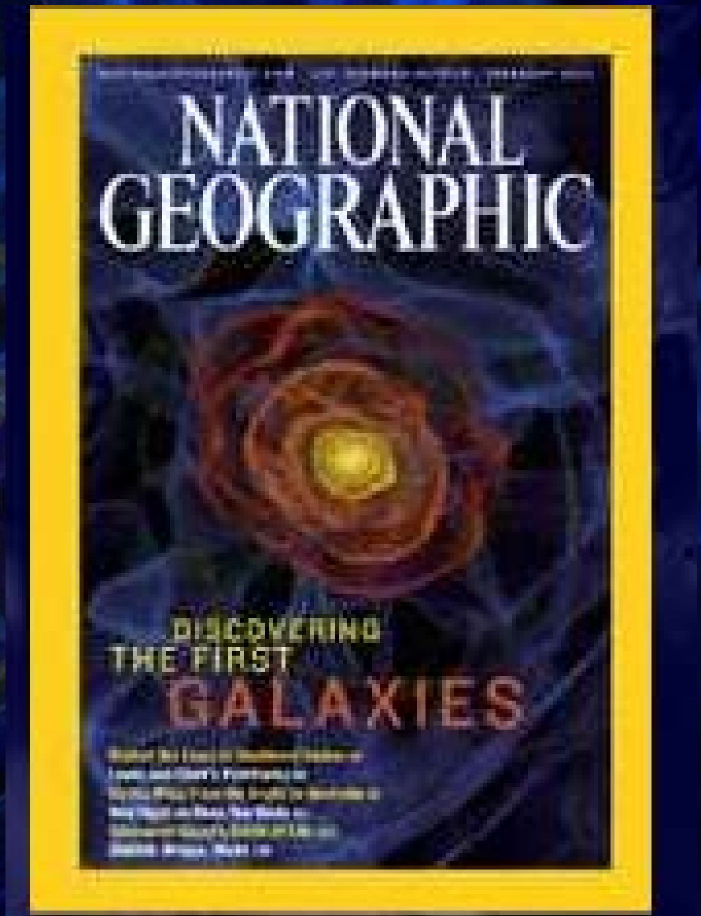


February 2003

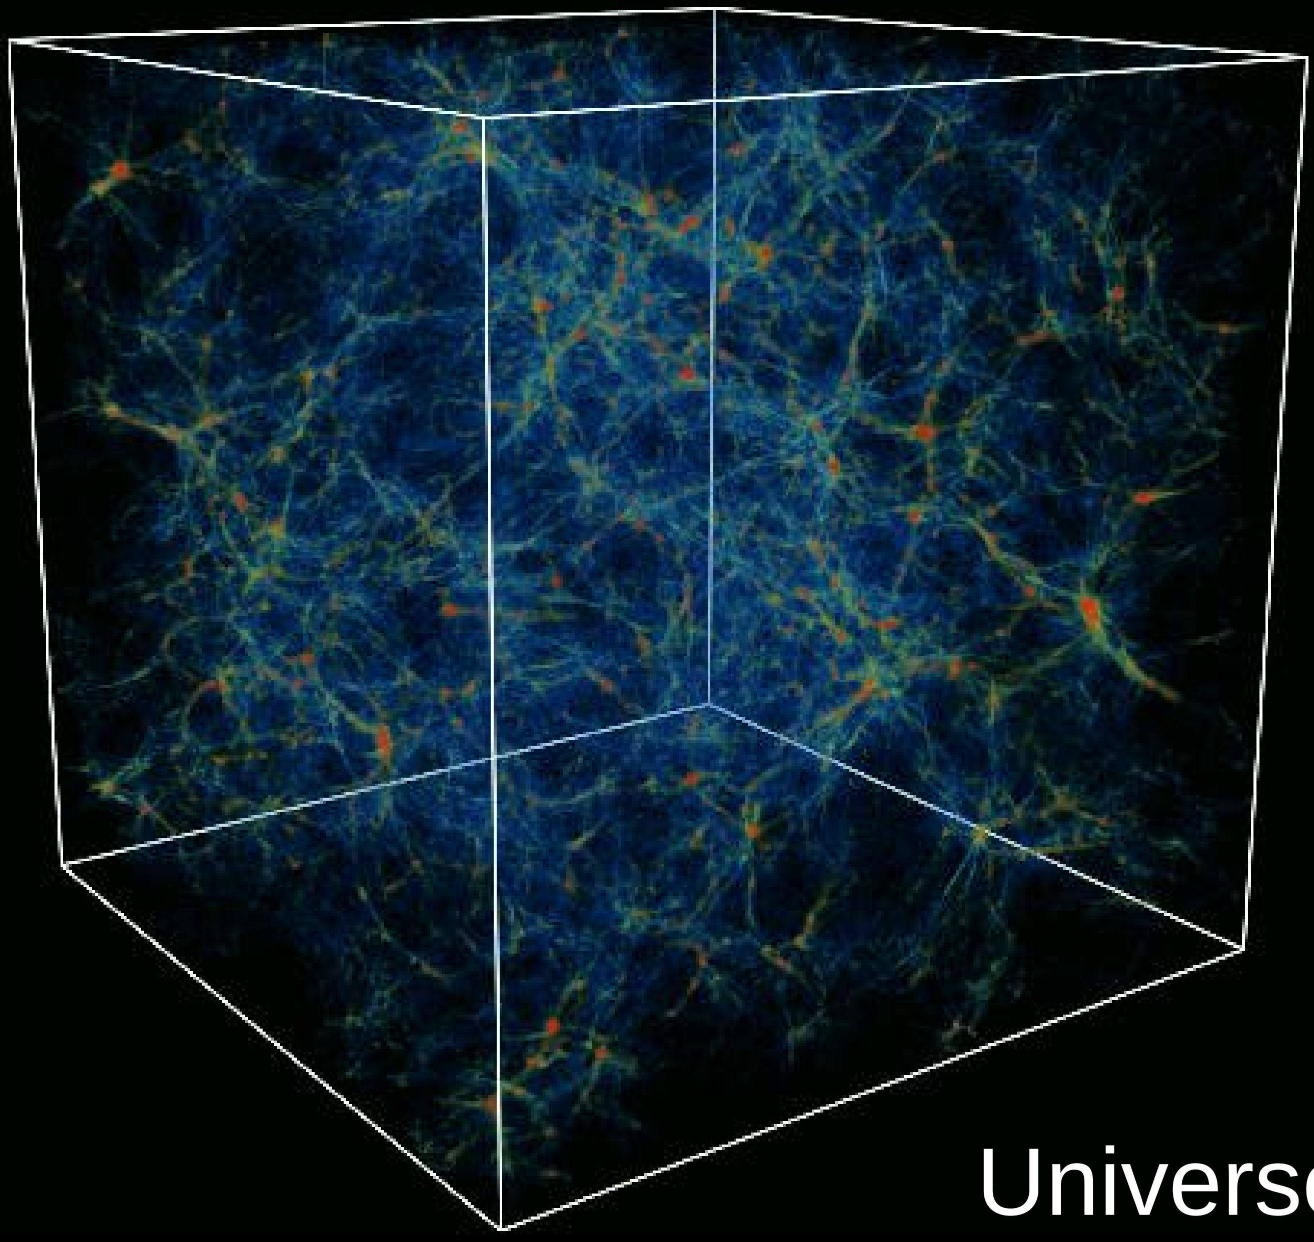

Universe in a Box

# The Universe is an IVP suitable for computation

- *Globally*, the universe evolves according to the Friedmann equation

$$H(t)^2 \equiv \left(\frac{\dot{a}}{a}\right)^2 = \frac{8\pi G}{3}\rho - \frac{k}{a^2} + \frac{\Lambda}{3}$$

Hubble parameter

scale factor a(t)

mass-energy density

spacetime curvature

cosmological constant

# The Universe is an IVP...

- *Locally**, its contents obey:
  - Newton's laws of gravitational N-body dynamics for stars and collisionless dark matter
  - Euler or MHD equations for baryonic gas/plasma
  - Atomic and molecular processes important for the condensation of stars and galaxies from diffuse gas
  - Radiative transfer equation for photons

(*scales << horizon scale ~ ct)

# Gridding the Universe

- Transformation to *comoving coordinates* x=r/a(t)
- Triply-periodic boundary conditions

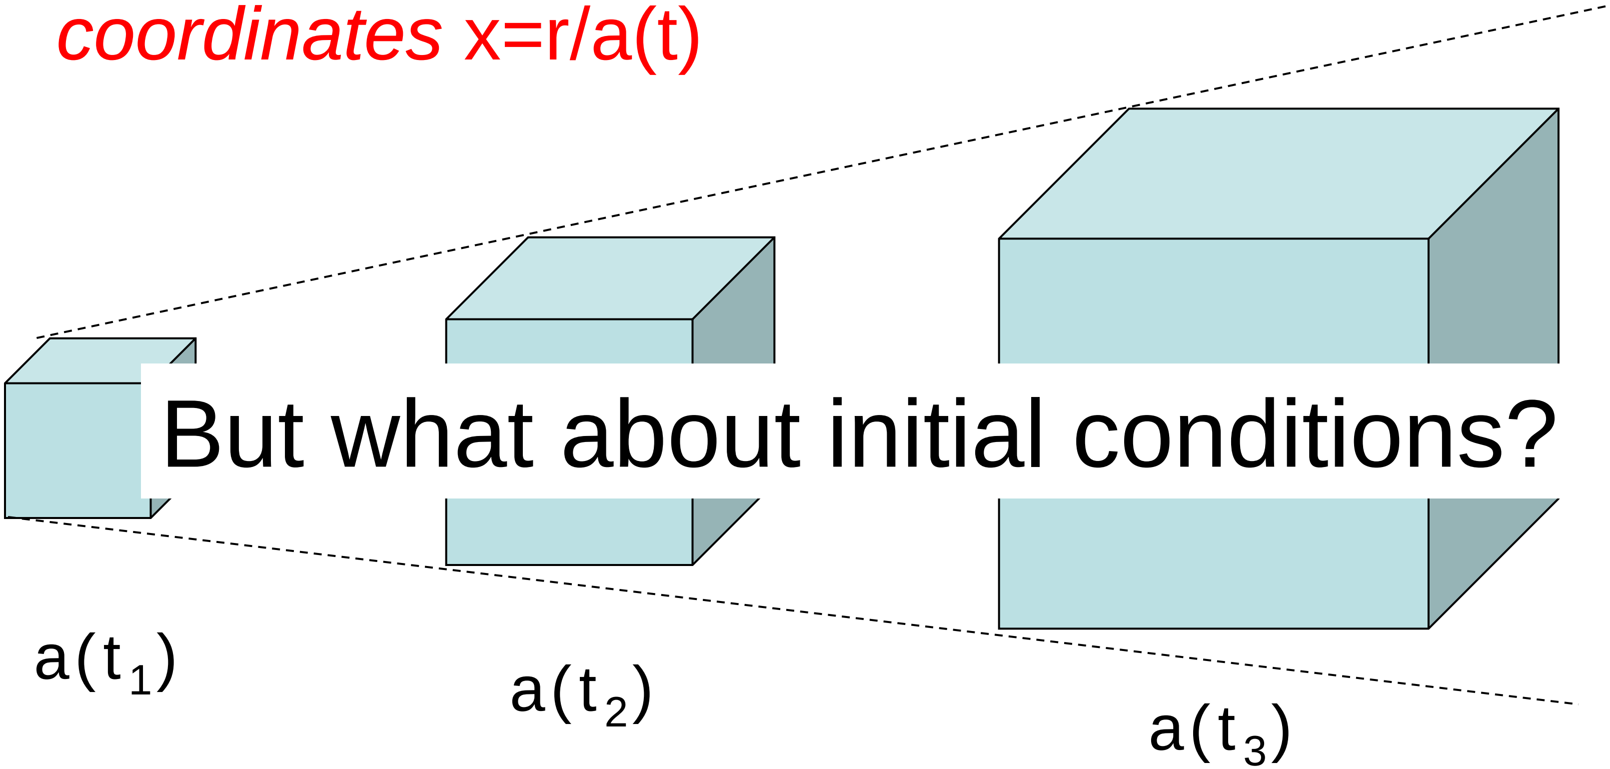

# Baby Picture of the Universe

## Image Shows Temperature Fluctuations in CMBR at 380,000 yr after BB

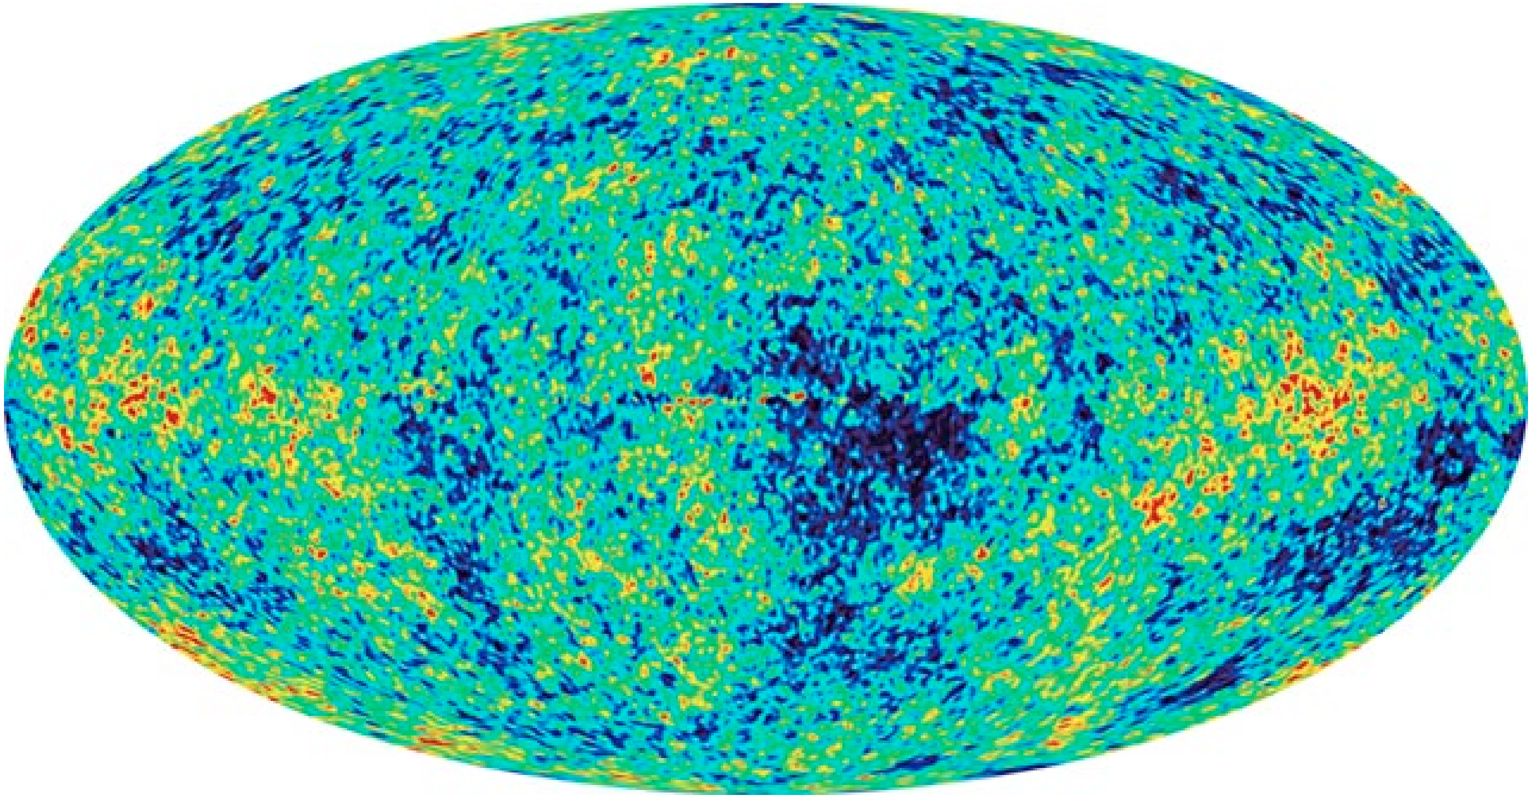

NASA/Princeton WMAP team

$\Delta T/T \sim 10^{-4}$

# Gravitational Instability: Origin of Cosmic Structure

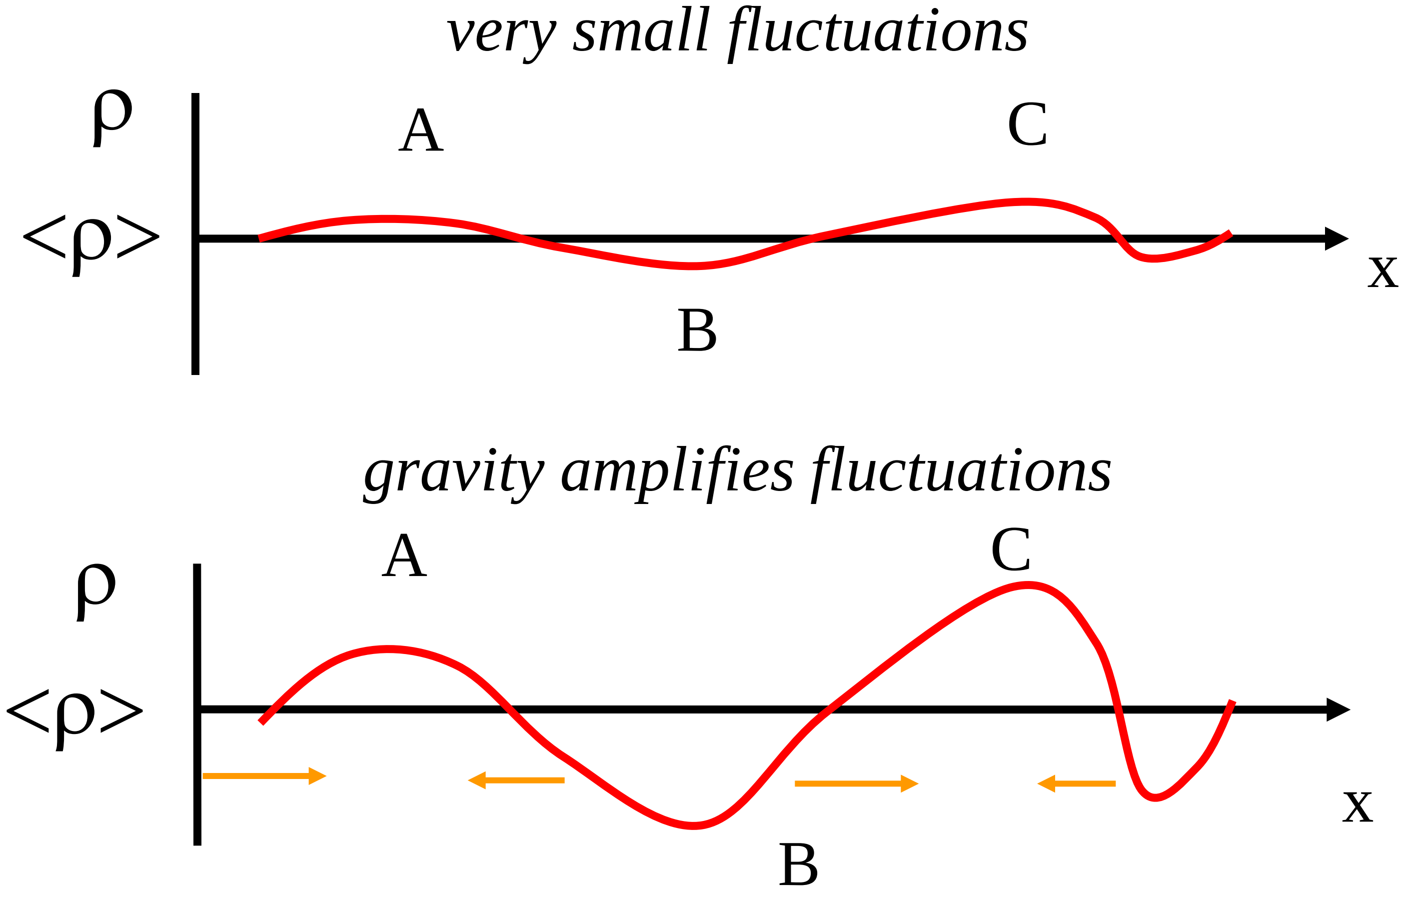

# Formation of First Bound Objects (Minihalos)

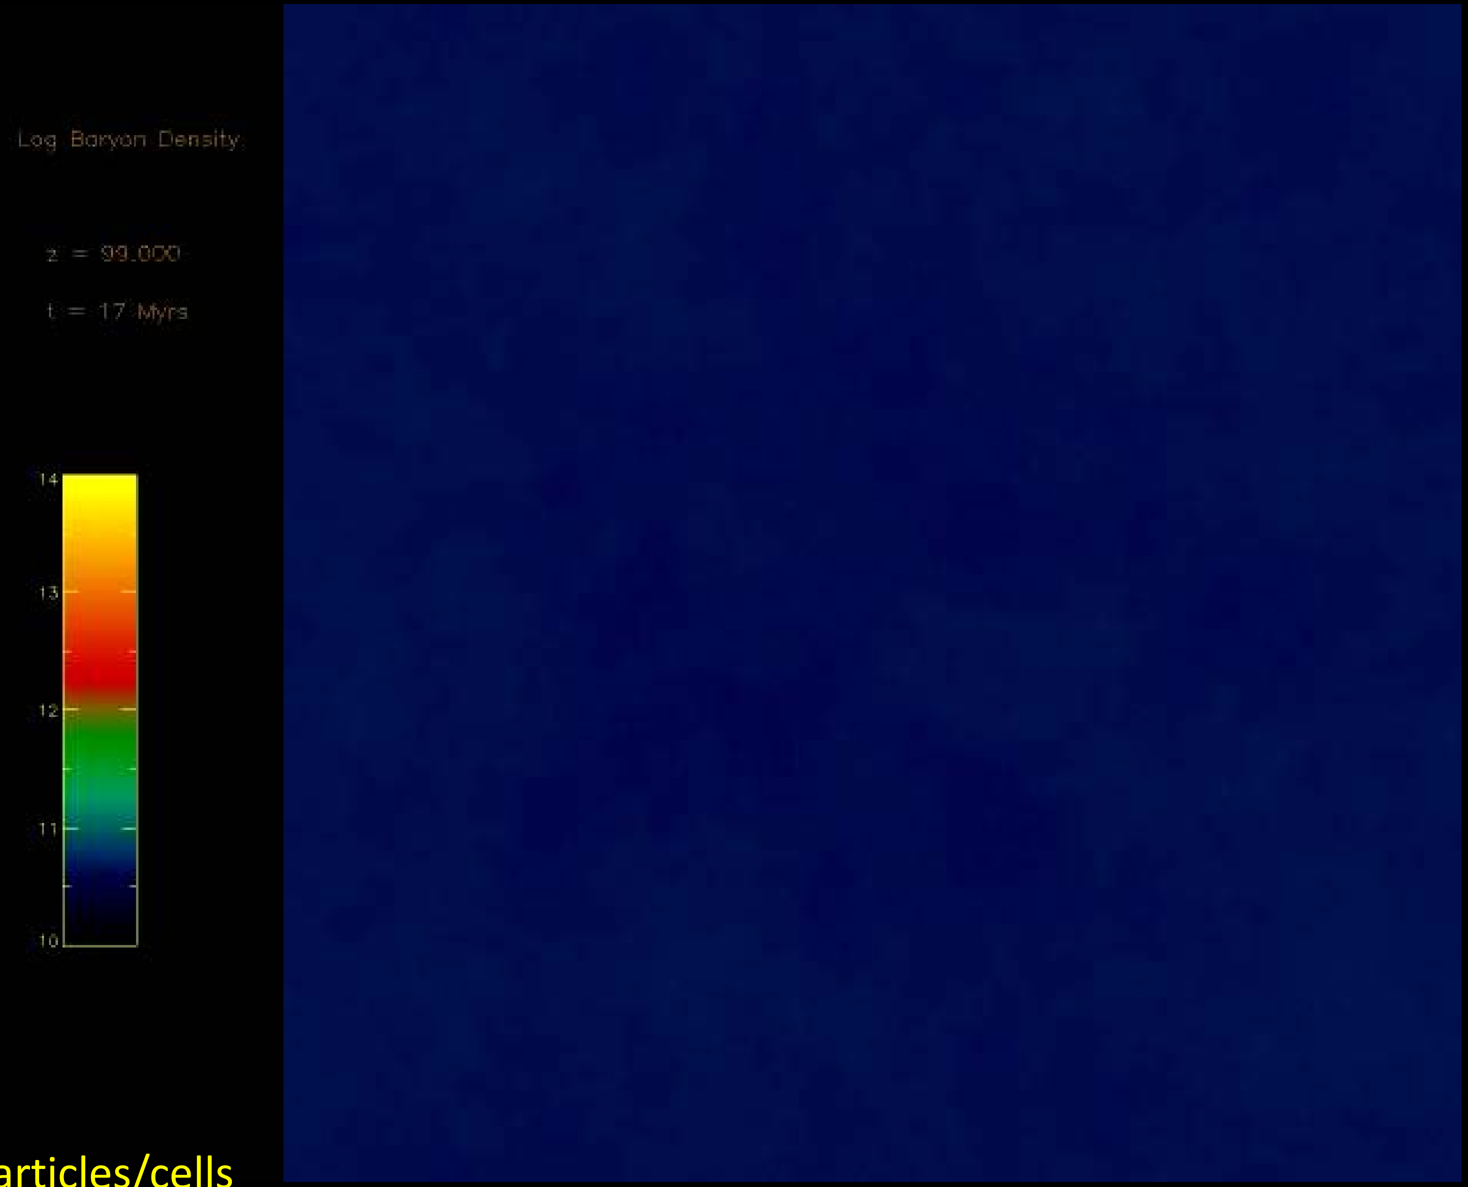


8 Mpc
1 billion particles/cells

# Formation of First Stars

Abel, Bryan & Norman (2001) *Science*

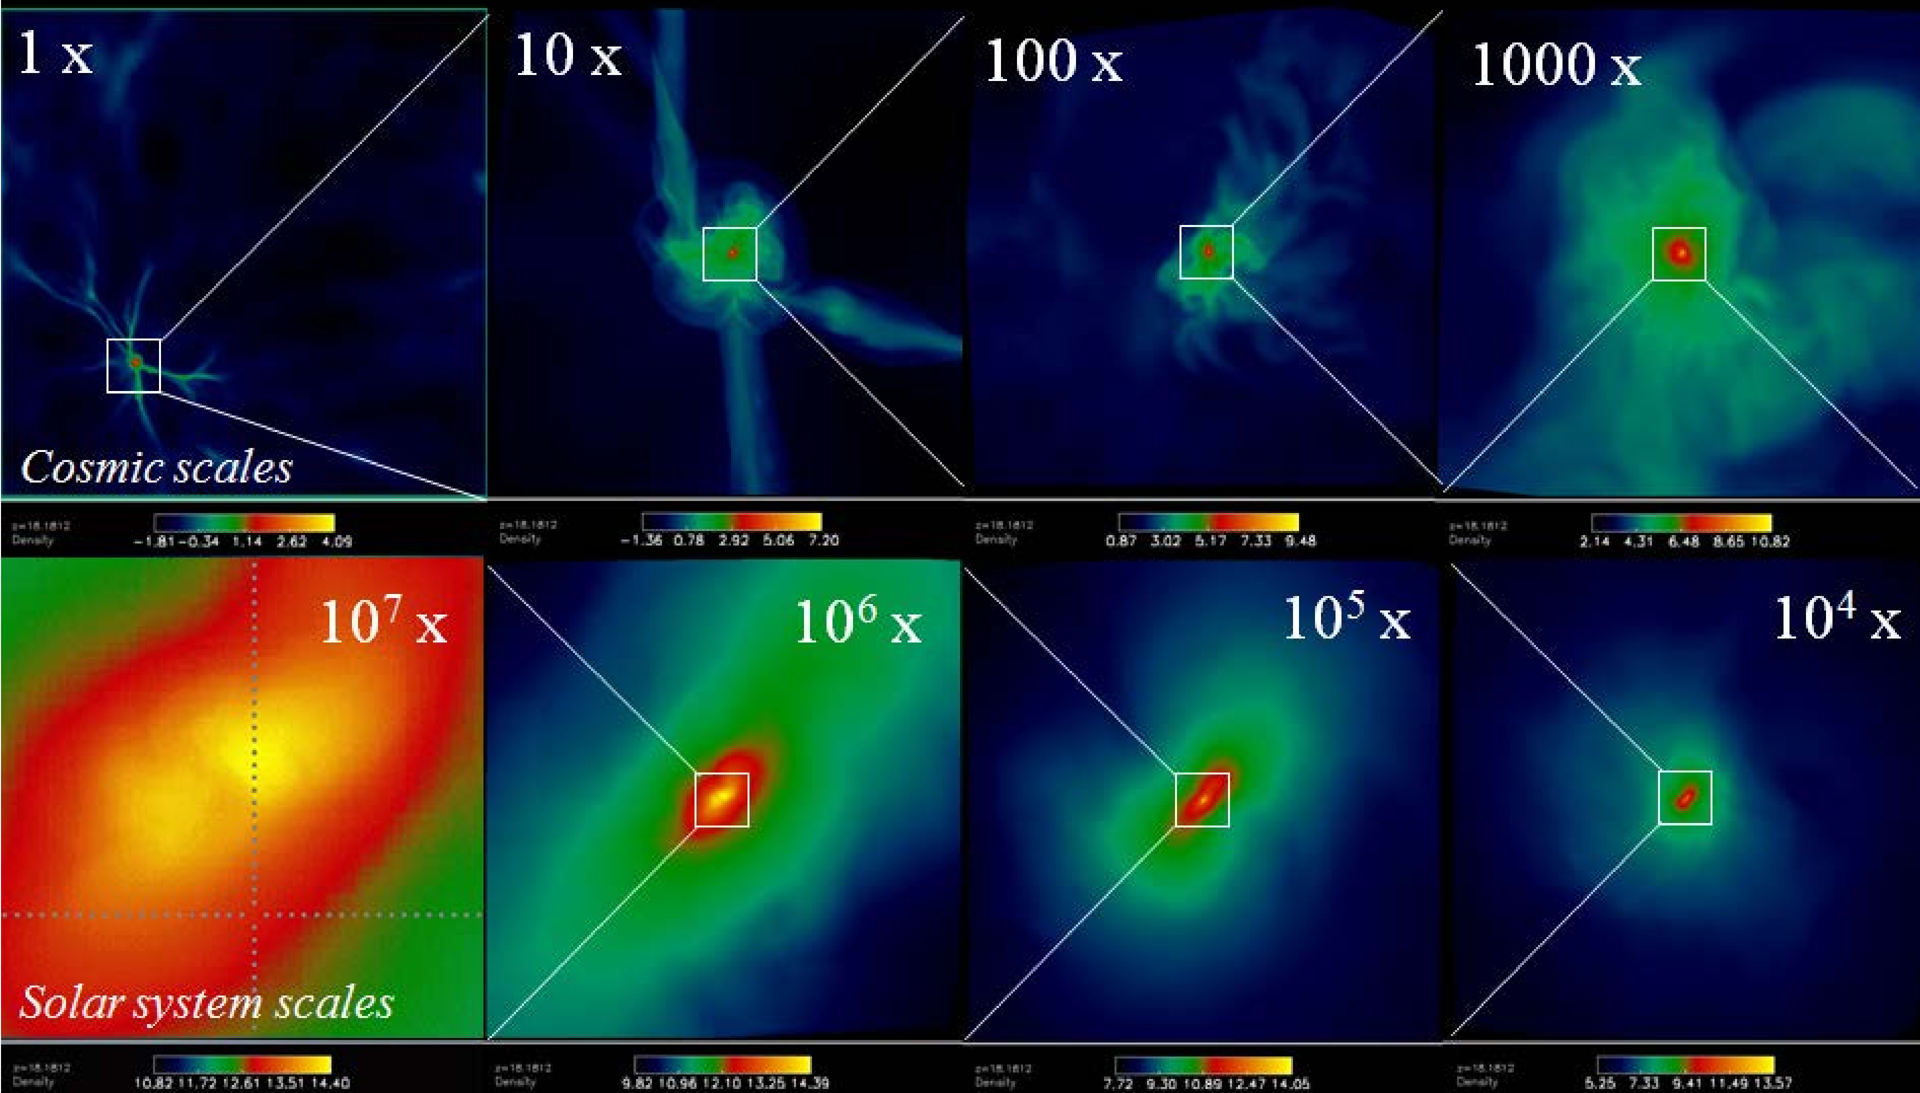

# Range of scales=5 x $10^7$



Fuld Hall
IAS, Princeton

# Formation of a First Generation Star
# (Zoom-in on one minihalo)

# Findings and Implications

Abel, Bryan & Norman (2002; *Science Express*)

- First stars are massive: ~100 M(solar)
- Only one star forms per *microgalaxy*
- They will be extraordinarily luminous and photo-ionize the intergalactic medium
- They will explode as supernovae, and seed the universe with heavy elements (C, N, O, Ca, Si, Fe…..)

# Making a First Galaxy (Protogalaxy)

- A first galaxy forms out of the debris of 100-1000 first stars pulled together by gravity
- Heavy elements produced by the first supernovae allow the gas to cool faster and produce the first "normal stars"
- Radiation from the first stars and galaxies ionizes and heats the intergalactic gas
- Ionized and chemically enriched gas gets ejected into space
- All of this physics needs to be simulated over a vast range of scales

# The Birth of a Galaxy

Wise et al. 2012a,b

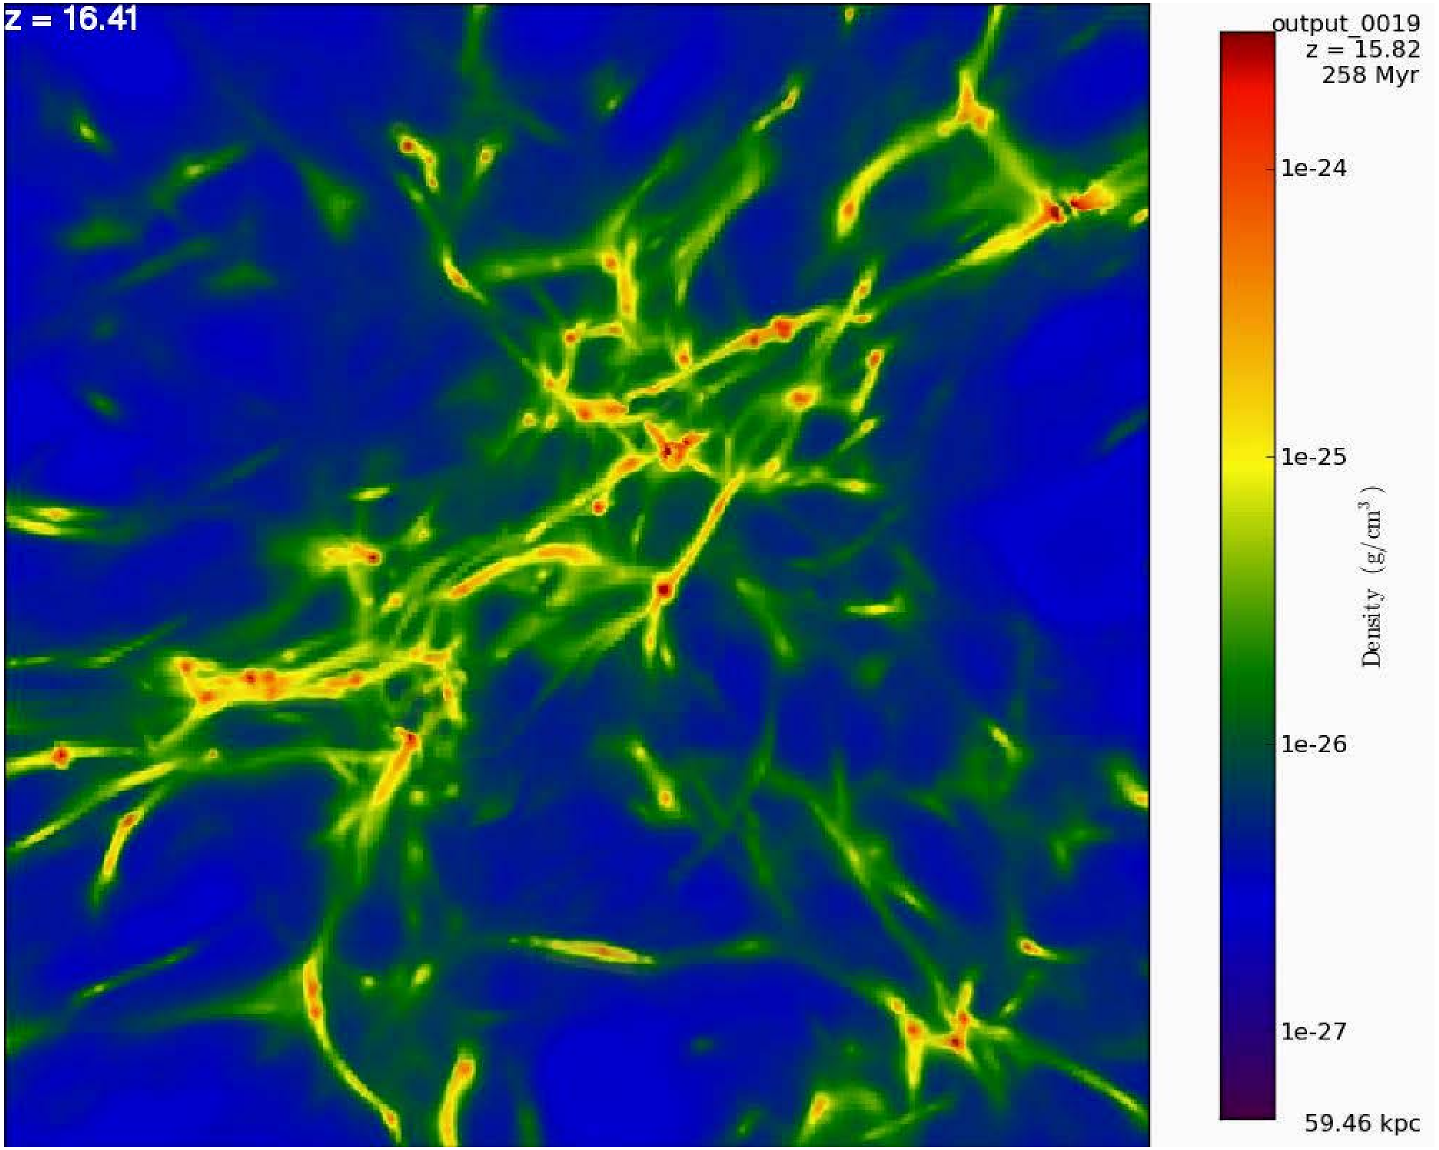

# The (Violent) Birth of a Galaxy

Wise et al. 2012a,b

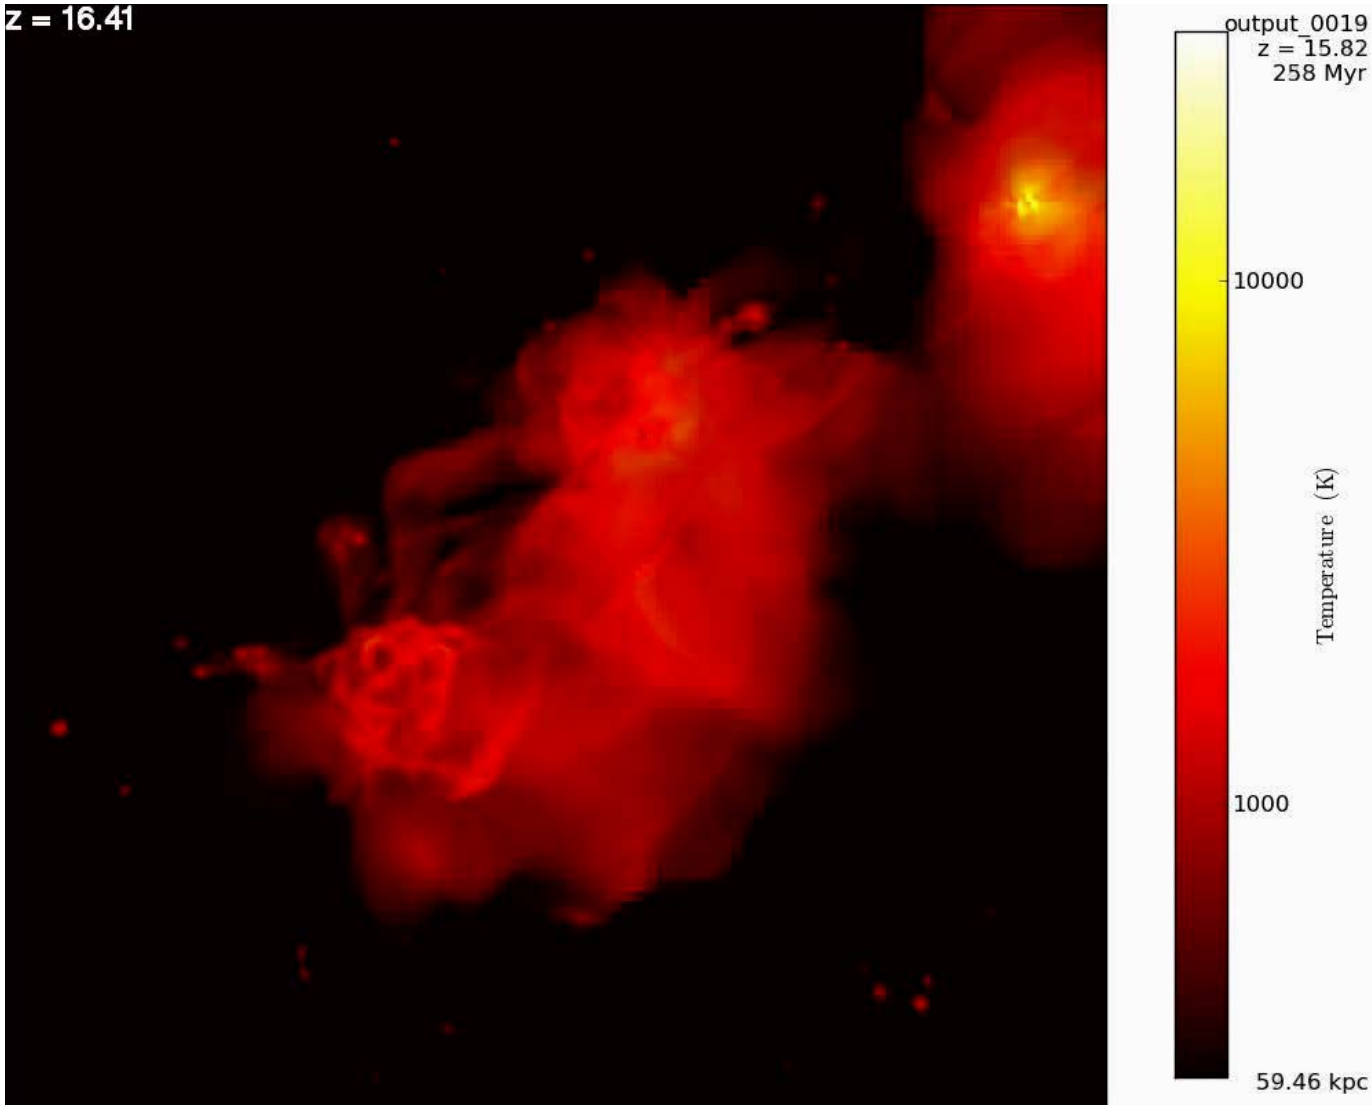

# The Birth of a Galaxy - Stars

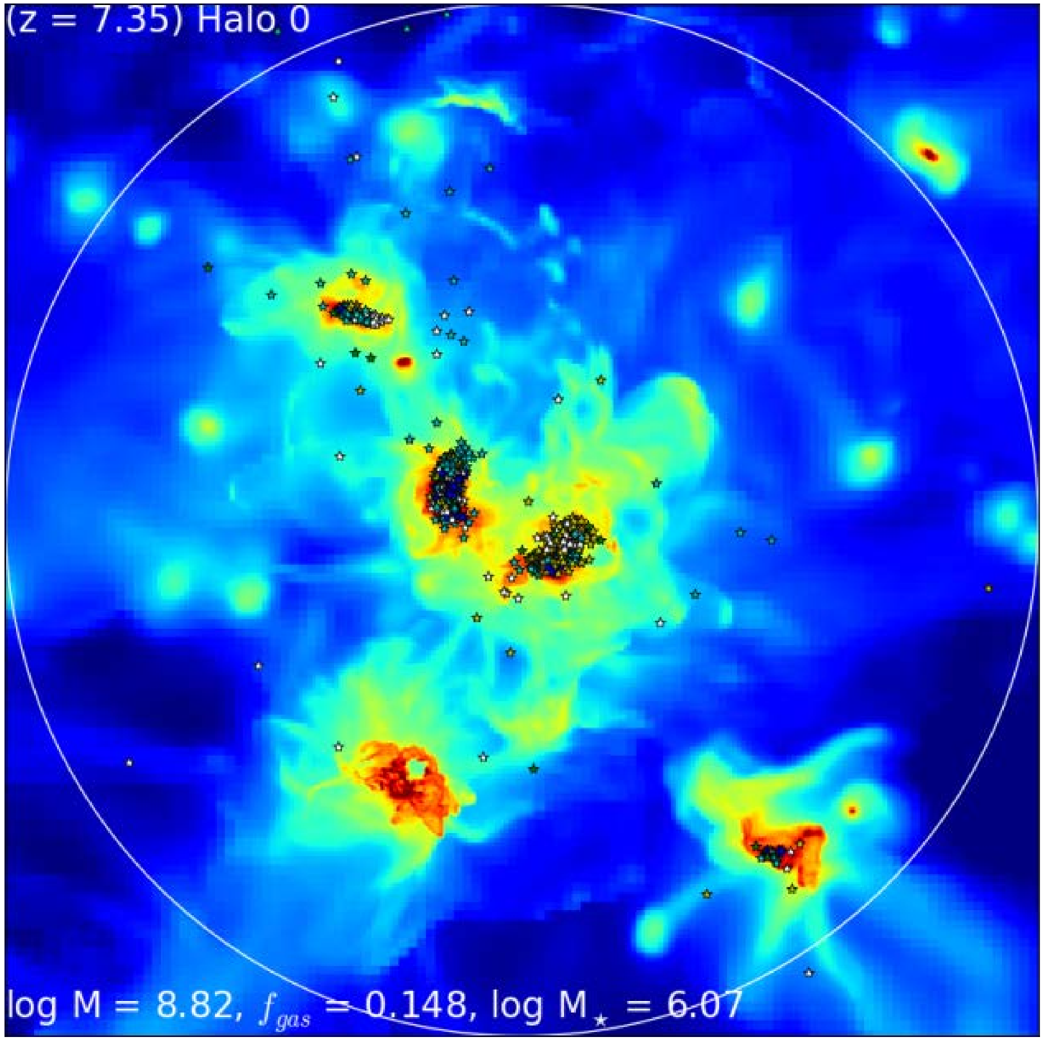

# First Galaxies and Reionization

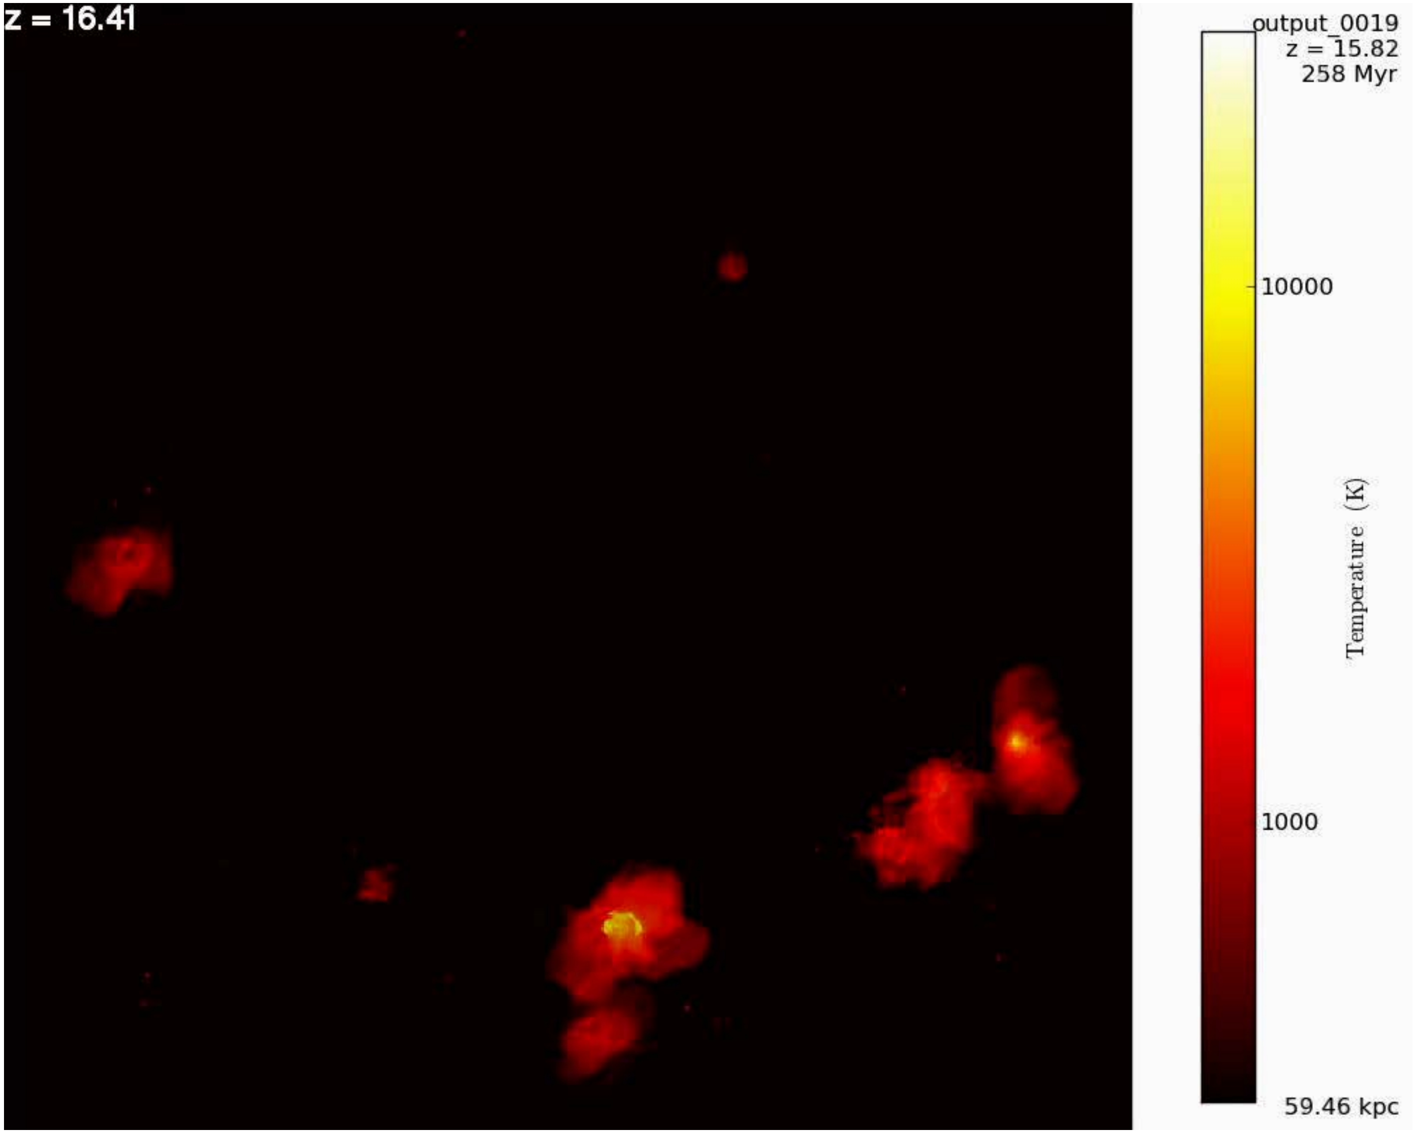

# Cosmology simulation matter power spectrum measurement using vSMP

We have run two large ($3200^3$ uniform grid) simulations, with and without radiation hydrodynamics, to measure the effect of the light from the first stars on the evolution of the universe. To quantitatively compare the matter distribution of each simulation, we use radially binned 3D power spectra.

- 2 simulations
- $3200^3$ uniform 3D grids
- 244GiB+ per field
- 15k+ files each

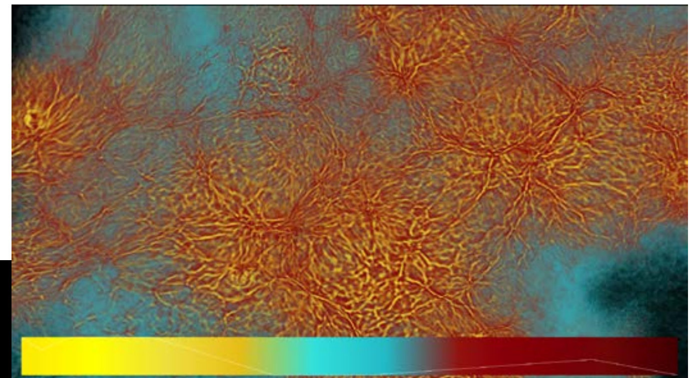

Difference

Individual simulations

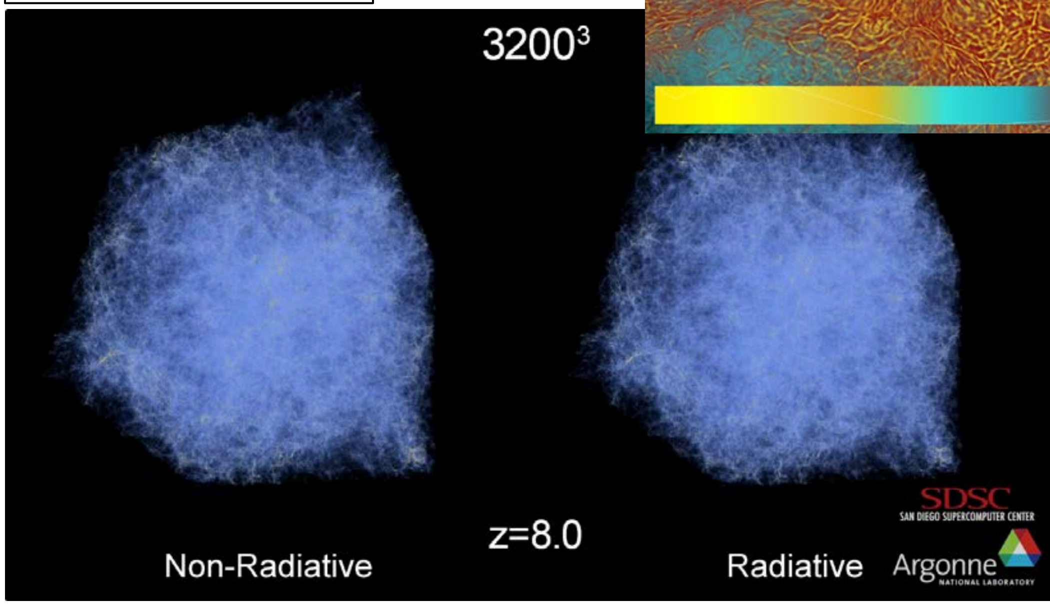


- Ran existing OpenMP-threaded code
- ~256GiB memory used
- ~5 ½ hours per field
- 0 development effort

Power spectra

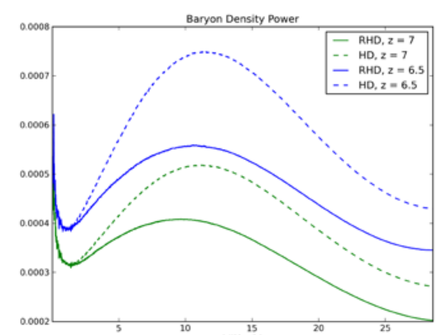


***Source:*** *Rick Wagner, Michael L. Norman. SDC. Used by permission. 2012*

# Key messages

- Astrocomputing and supercomputing
  - Astronomers have always been on the vanguard
  - Astronomy applications are voracious in their computing demands
- Technology trends
  - HW: Moore's law for supercomputing is alive and well (if not accelerating)
  - HW: Its all about the cores; different ways they are offered
  - SW: Efficient use requires heroic programming efforts
  - Data: new data-intensive architectures needed to cope with data deluge (Gordon)
- Applications to Cosmic Renaissance
  - First stars➔first galaxies➔reionization
  - Suppression of DM power by Jeans smoothing