Cosmological N-Body Simulations and Galaxy Surveys

Adrian Pope, High Energy Physics, Argonne National Laboratory, apope@anl.gov

CScADS: Scientific Data and Analytics for Extreme-scale Computing, 30 July 2012

ANL: HEP (Salman Habib, Katrin Heitmann), ALCF, MCS

LANL: CCS-6, CCS-7

LBL: P, CR

UC Berkeley, Northwestern, Rutgers, VA Tech

Structure Formation: The Basic Paradigm

- **Solid understanding of structure formation; success underpins most cosmic discovery**
  - Initial conditions laid down by inflation
  - Initial perturbations amplified by gravitational instability in a dark matter-dominated Universe
  - Relevant theory is gravity, field theory, and atomic physics ('first principles')
- **Early Universe:**
  - Linear perturbation theory very successful (CMB)
- **Latter half of the history of the Universe:**
  - Nonlinear domain of structure formation, impossible to treat without large-scale computing

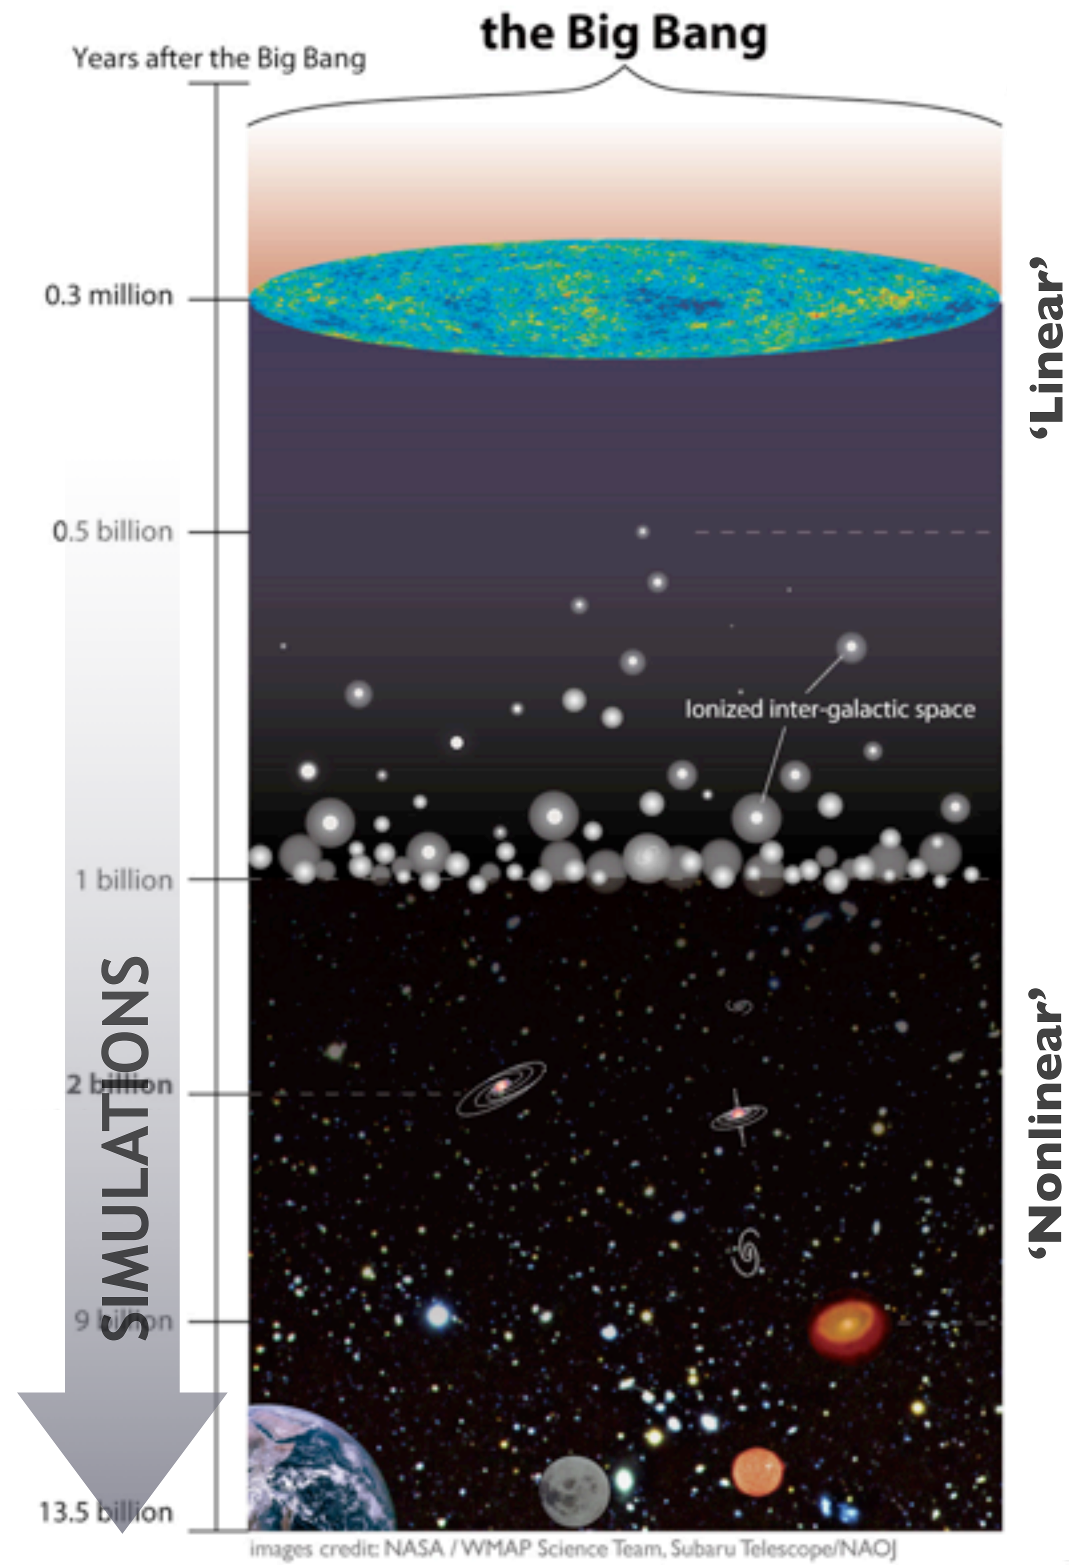

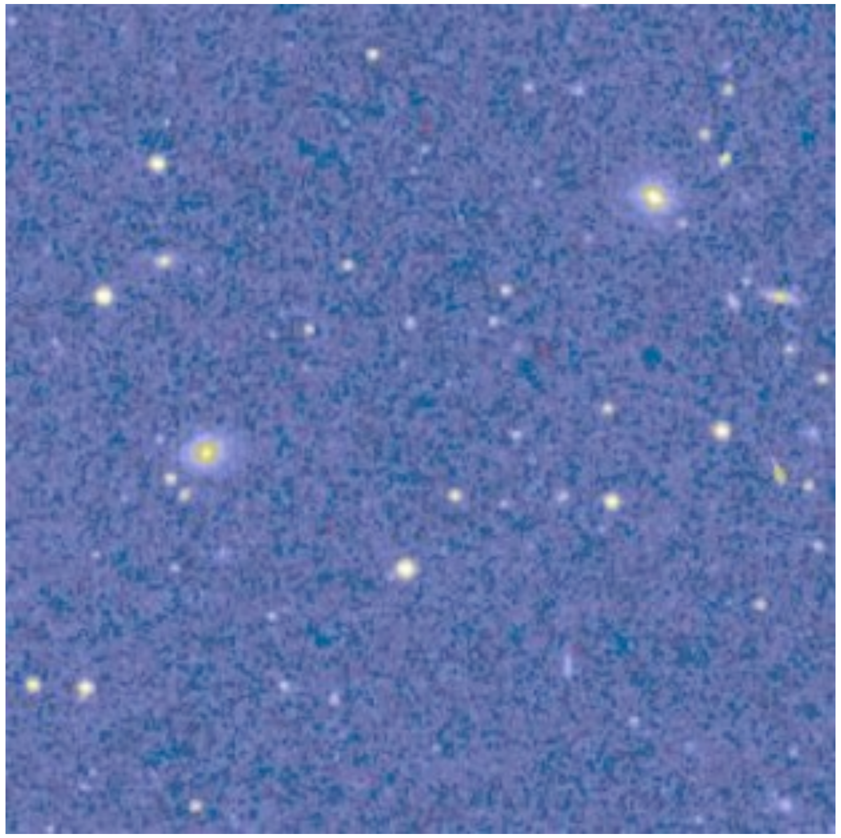

Digitized Sky Survey
1950s-1990s

Sloan Digital Sky Survey
2000-2008

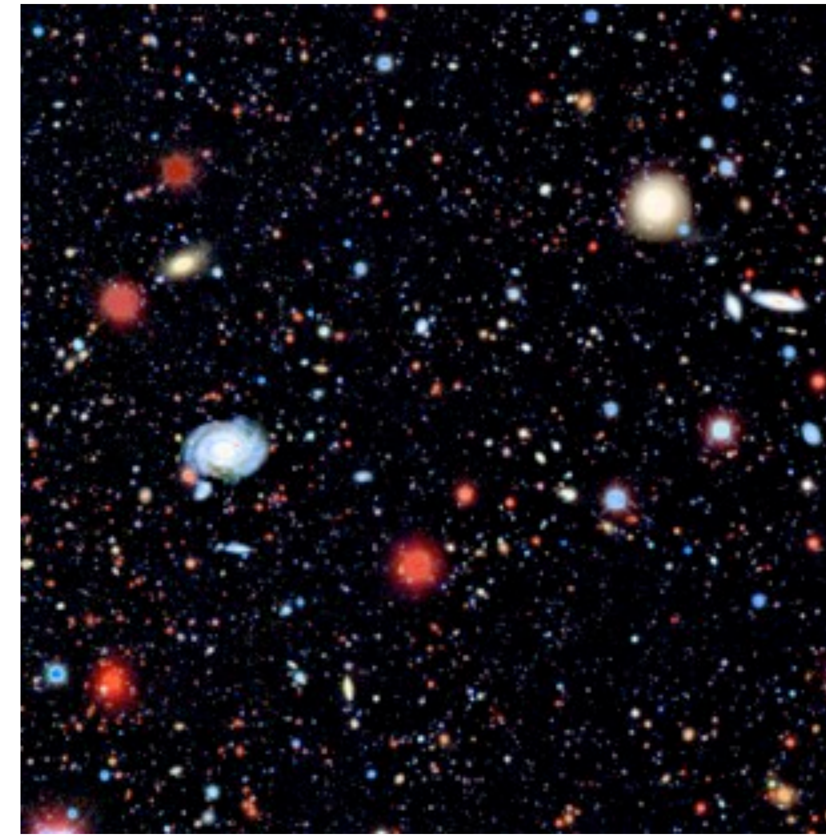

Large Synoptic Survey Telescope
2020-2030
(Deep Lens Survey image)

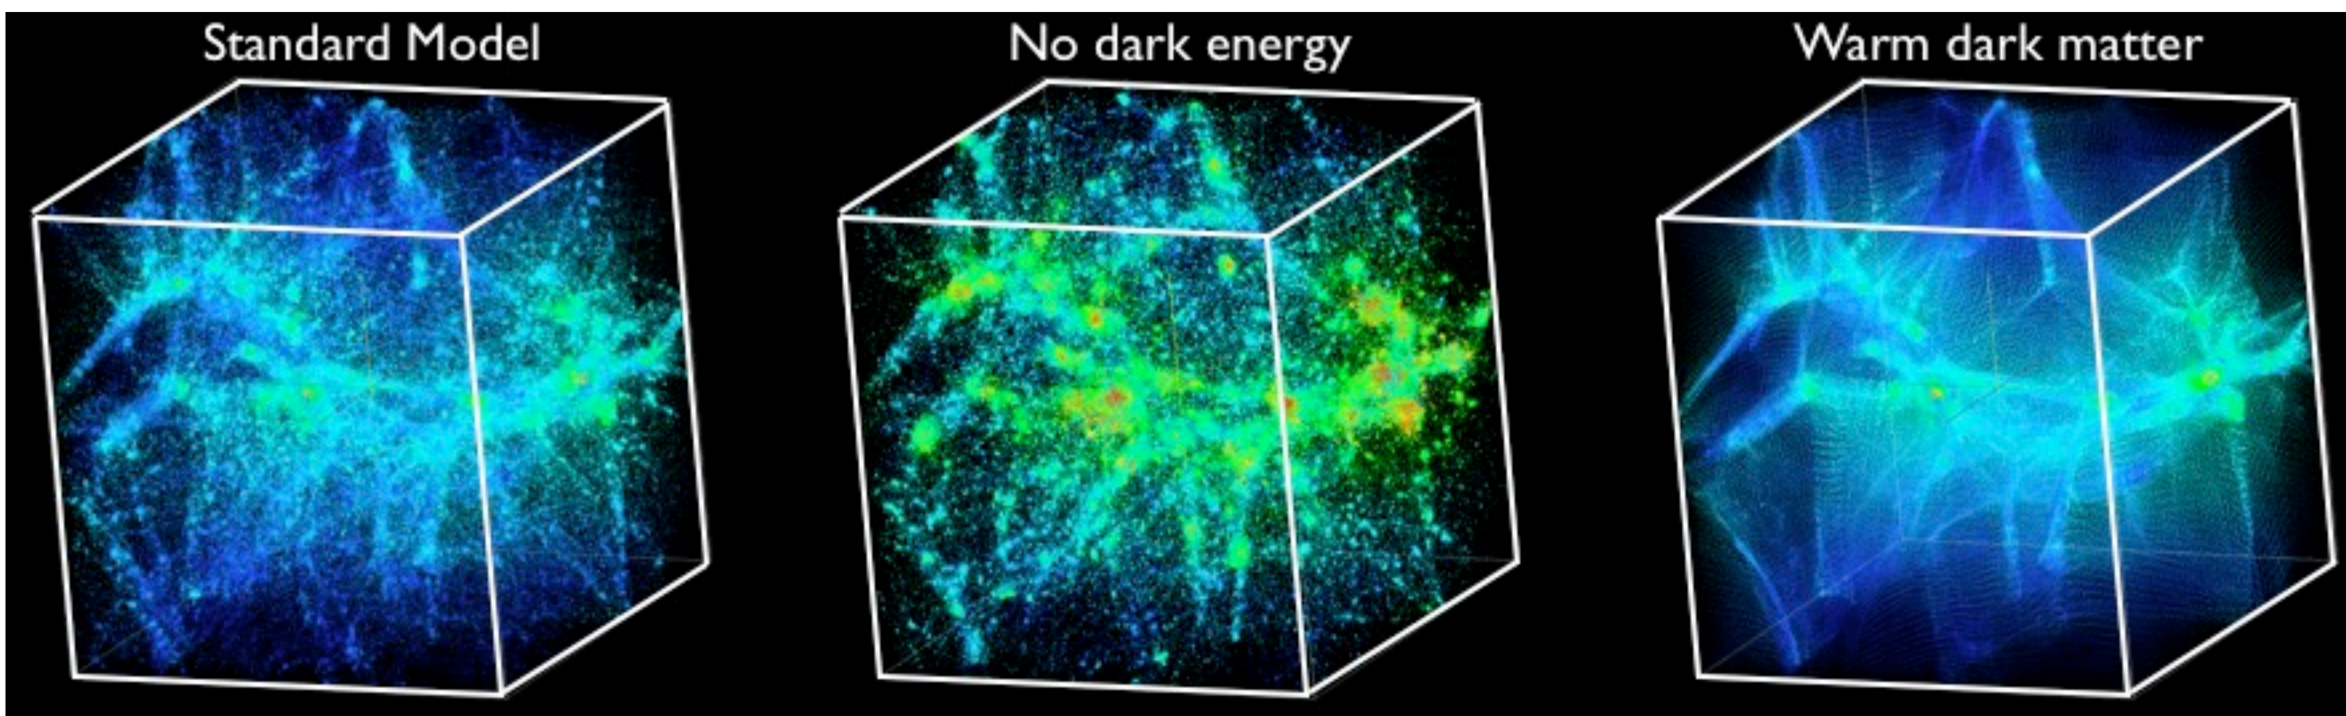

Gravity in an Expanding Universe

- **Evolution**
  - Initial conditions: Gaussian random field of small density contrast
  - Structure formation: gravitational collapse
    - Vlasov-Poisson equation in expanding space
    - Expansion slows collapse from exponential to power law, structures on many scales
- **Accuracy**
  - Observations will soon require percent-level theoretical predictions
  - N-body methods use tracer particles to evolve phase-space distribution (no 6D PDE)
- **Inverse problem**
  - Compare forward model of structure formation with galaxy surveys
  - Build emulators from limited number of simulations in parameter space

$$
\begin{aligned}
\frac{\partial f_i}{\partial t} &+ \dot{\mathbf{x}}\frac{\partial f_i}{\partial \mathbf{x}} - \nabla\phi\frac{\partial f_i}{\partial \mathbf{p}} = 0, \qquad \mathbf{p} = a^2\dot{\mathbf{x}}, \\
\nabla^2\phi &= 4\pi G a^2(\rho(\mathbf{x},t) - \langle\rho_{\rm dm}(t)\rangle) = 4\pi G a^2 \Omega_{\rm dm}\delta_{\rm dm}\rho_{\rm cr}, \\
\delta_{\rm dm}(\mathbf{x},t) &= (\rho_{\rm dm} - \langle\rho_{\rm dm}\rangle)/\langle\rho_{\rm dm}\rangle), \\
\rho_{\rm dm}(\mathbf{x},t) &= a^{-3}\sum_i m_i \int d^3\mathbf{p} f_i(\mathbf{x},\dot{\mathbf{x}},t).
\end{aligned}
$$

Simulations of Structure Formation in the Universe

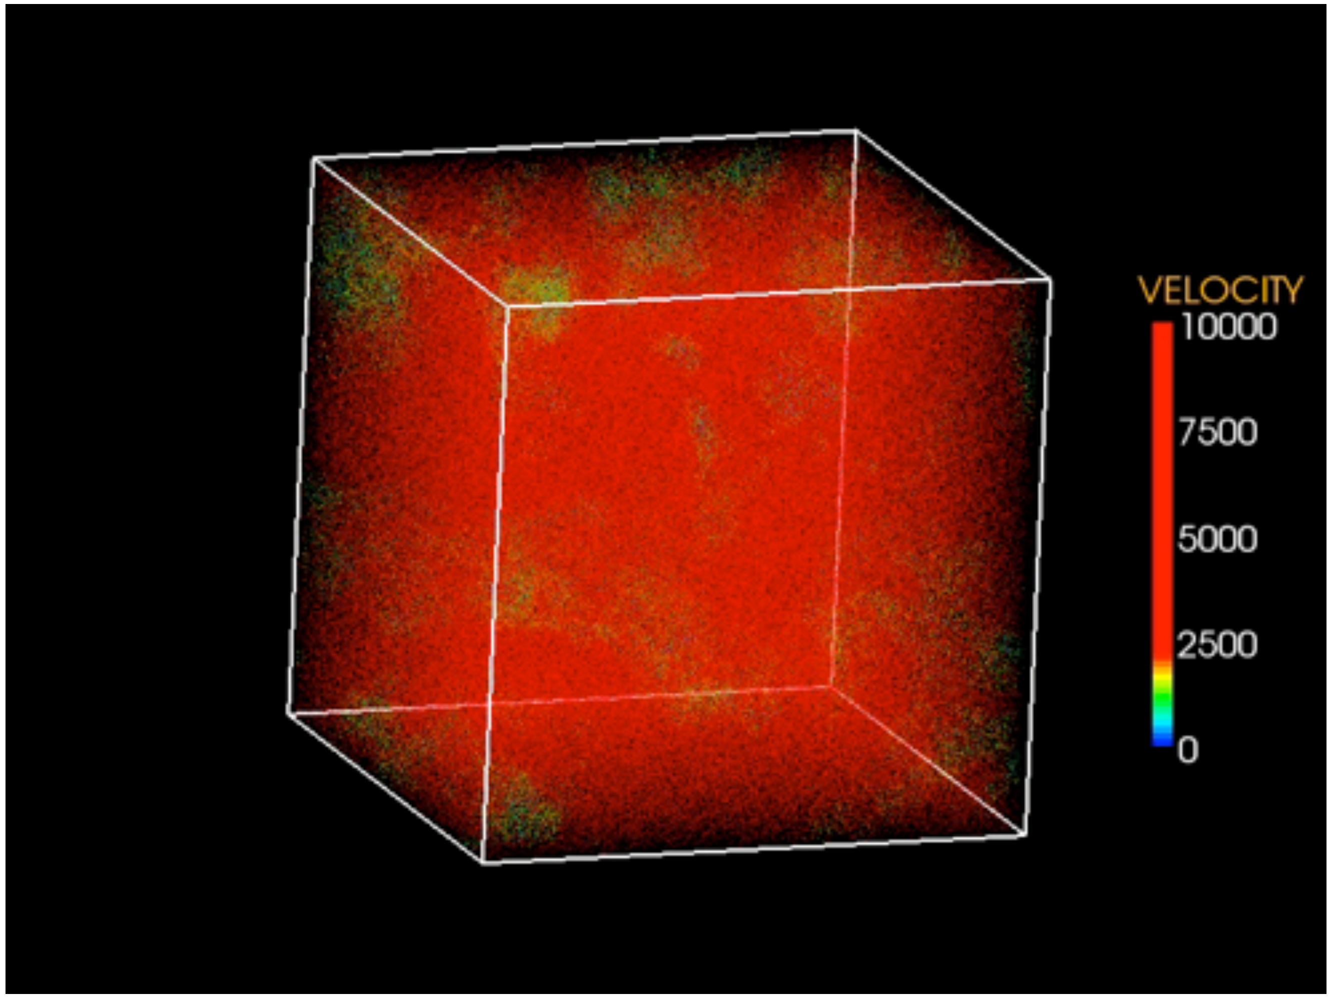

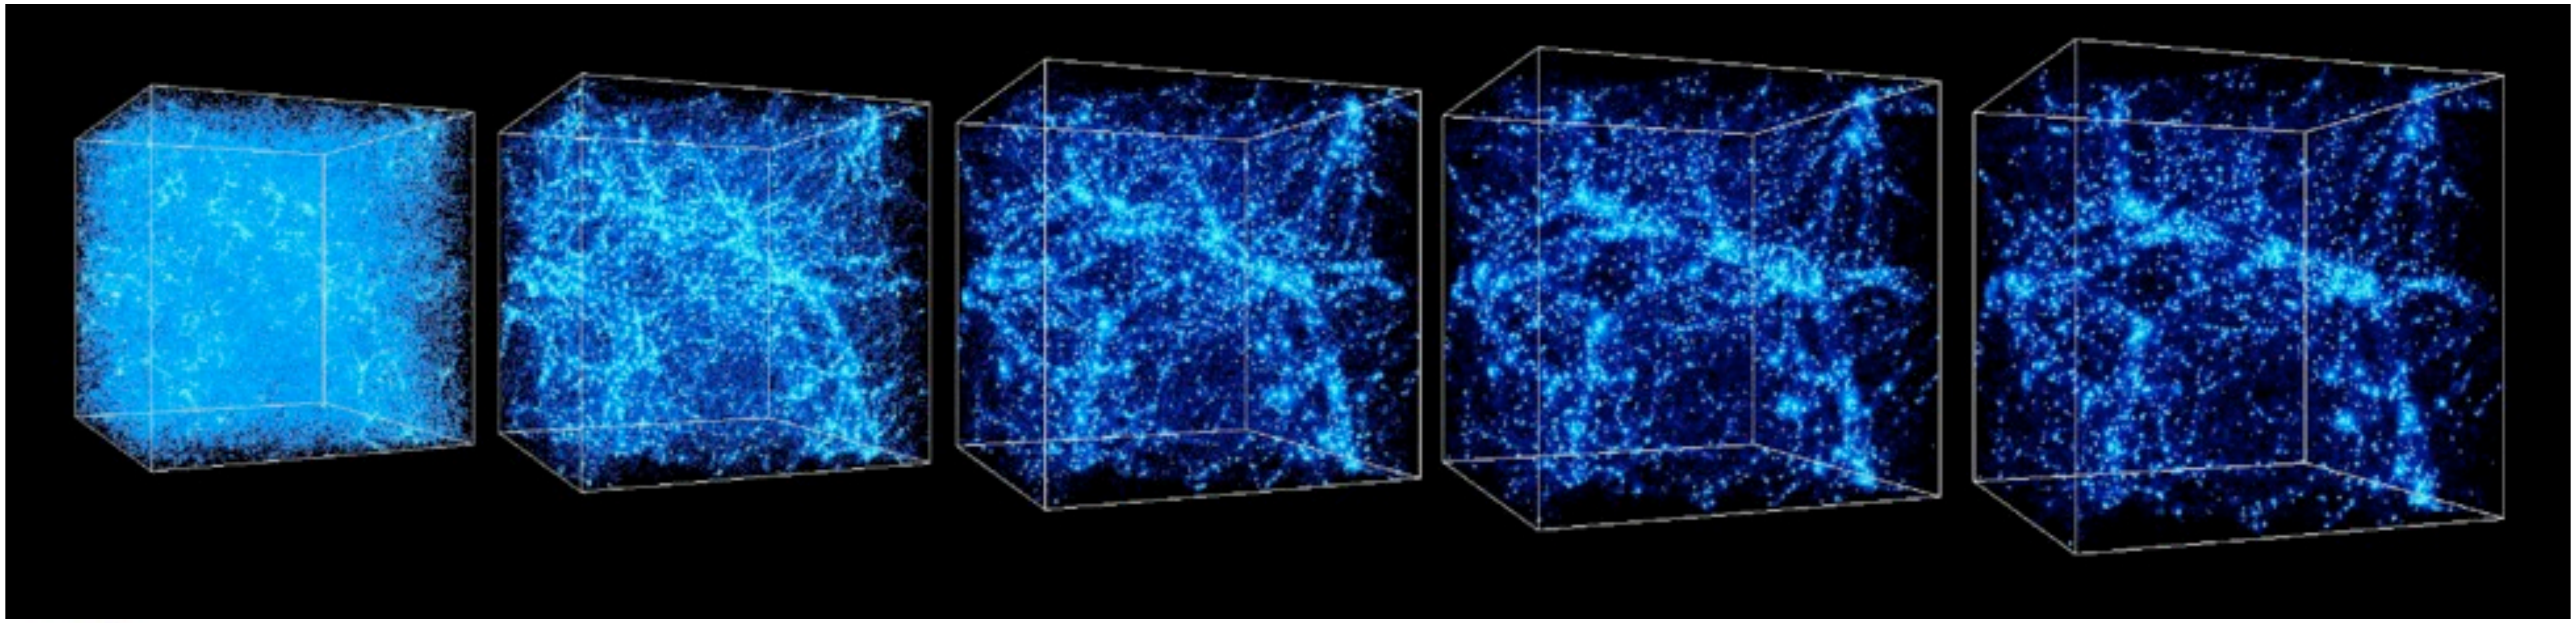

HACC: Hybrid Accelerated Cosmology Code

- **Requirements**
  - Particles: 10^{10} standard, 10^{11} state-of-the-art, 10^{12} very soon
  - Throughput: many runs for inverse problems, days for each run
- **Supercomputers**
  - Parallelism: code must scale
    - Many nodes: message passing, weak scaling, Mira = 48k nodes
    - Many cores: local threading, strong scaling, Mira = 16 cores x 4 hardware threads per node
  - Architectures: code must adapt
    - CPU, IBM Cell, IBM Blue Gene, GPU, Intel MIC
    - Different programming paradigms and optimizations
- **Design**
  - Split gravitational force calculation
    - Long-range: spectral (FFT) particle-mesh (PM) methods using MPI, portable
    - Short-range: algorithm choice and optimizations for particular architectures, modular, high intensity
- **Versions**
  - P^3M (particle-particle, particle-mesh): IBM Cell, GPU
  - TreePM: CPU, IBM Blue Gene

Data Challenges: At Runtime

- **Checkpoints**
  - Up to (roughly) half of machine memory to disk, Mira = 100s TB (without compression)
  - Must be fast to reduce time taken from allocation
  - Stressful for network-attached storage
  - Robust persistent local storage would be nice
- **Data reduction**
  - Cannot store enough full particle outputs to do all analysis in post-processing
  - On-the-fly analysis to produce reduced outputs
    - Halo finding (galaxies, clusters), 2D mass projection (weak lensing), sub-samples of particles, etc.
  - Reduced outputs must be sufficient for scientific analysis
- **Architectures**
  - More varied analysis code is more difficult to optimize for different HPC systems

Data Challenges: Sharing Data Products

- Rich simulation outputs
  - Many different scientific queries on the same simulations
  - Insufficient bandwidth to disperse complete (reduced) data products
- How should scientists interact with the outputs?
  - Move analysis local to data
  - Further reductions before network transfer
- What will that require?
  - What kind of hardware?
    - HPC, DB, DISC (Data-Intensive SuperComputer)
  - What kind of software? Who will write/maintain? What can we build on?
    - SQL, NoSQL
  - Have worked with Globus Online
  - Looking at Galaxy (computational biology)

Data Challenges: Visualization

- **Particle number**
  - Far more particles than pixels
    - Some experiments with particle sampling methods
  - Volume rendering seems to be more common than direct particle methods
  - Visualize reduced outputs
- **Dynamic range**
  - Scales: homogeneous on large scales, complex structures everywhere on small scales
  - Huge density contrasts

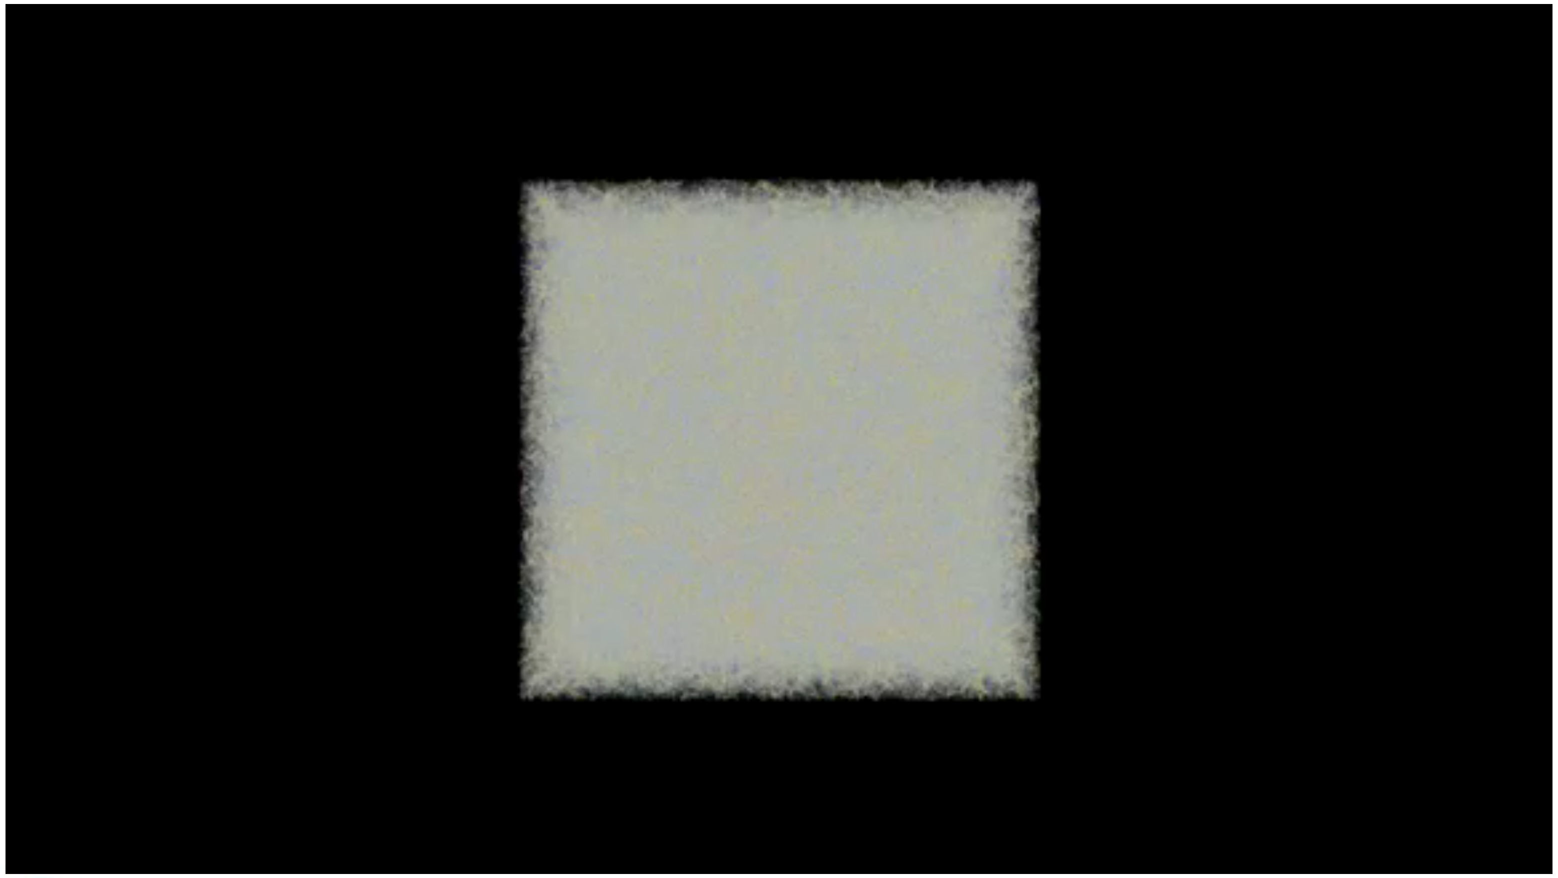

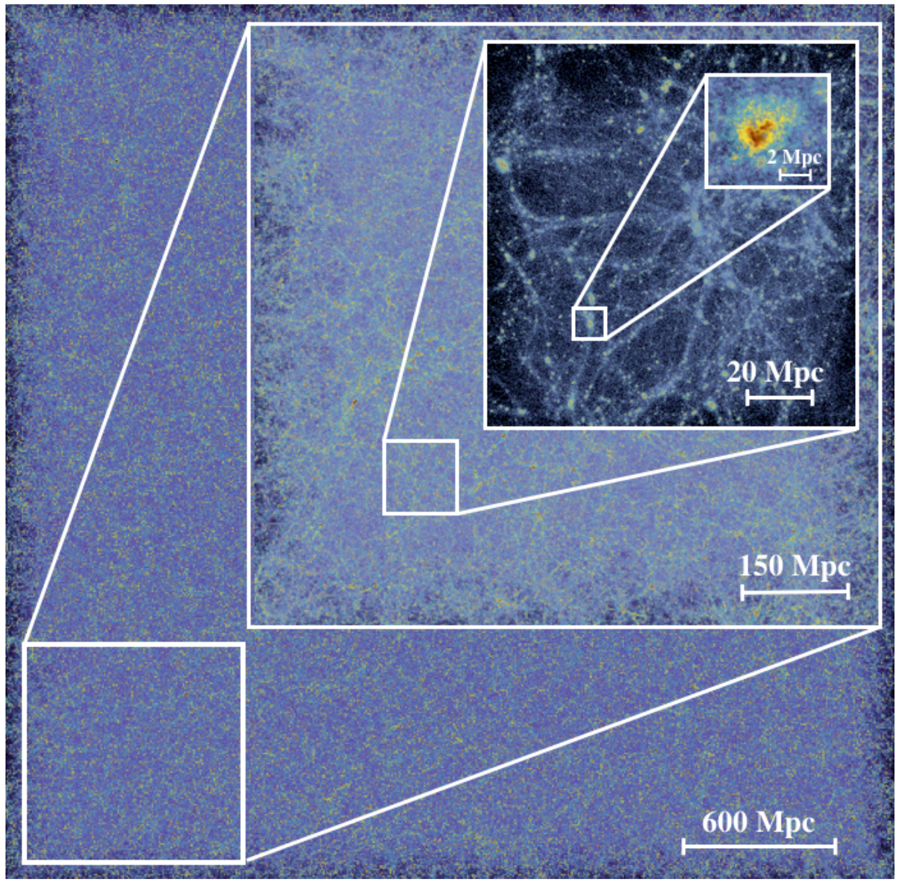

Preliminary Slice of a 230 Billion Particle Simulation

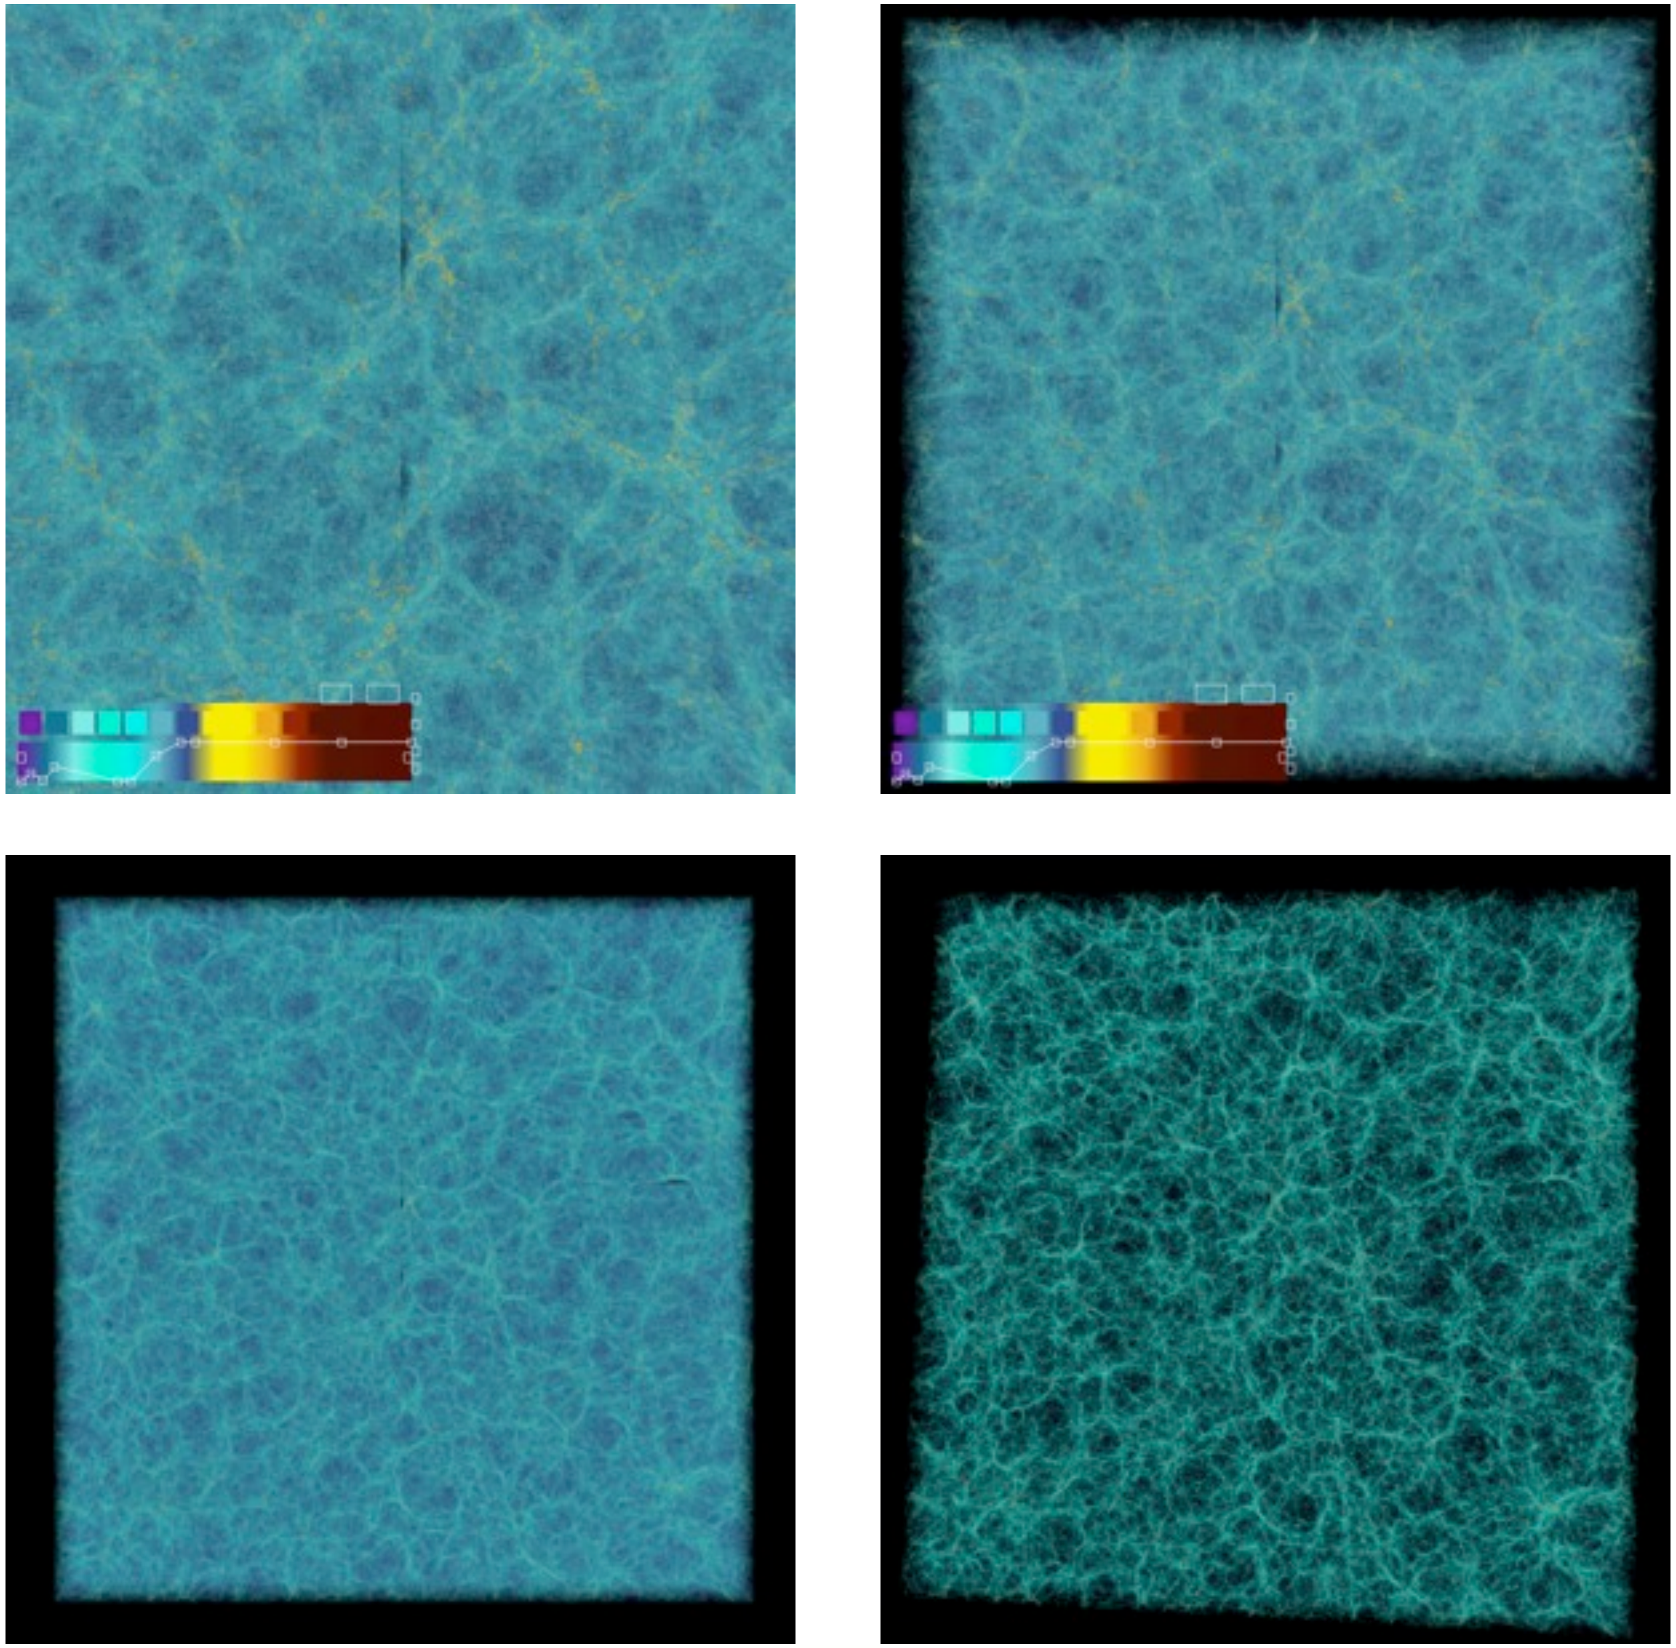

Data Challenges: Putting It All Together

- Observational data volume will not be the challenge
  - But these data sets are very rich and scientists will want to query it in many different ways
- Numerical theory can produce (almost arbitrarily) large volumes of data
  - Also rich
- How do we put these together?
  - Data-Intensive SuperComputers (DISC)