# Computation-Driven Discovery for the Dark Universe

Salman Habib
HEP and MCS Divisions
Argonne National Laboratory

PIs: K. Heitmann (ANL), A. Slozar (BNL), S. Dodelson (FNAL), P. Nugent (LBNL), J. Ahrens (LANL), R. Wechsler (SLAC)


*ASCR*
*HEP*

# ‘Precision’ Cosmology

- Instrumentation Advances
- Cosmic Acceleration
- Nature of Dark Matter
- Primordial Fluctuations
- Neutrinos
- Cosmic Structure Formation


Boyle, Smith

Perlmutter, Riess, Schmidt


Mather, Smoot

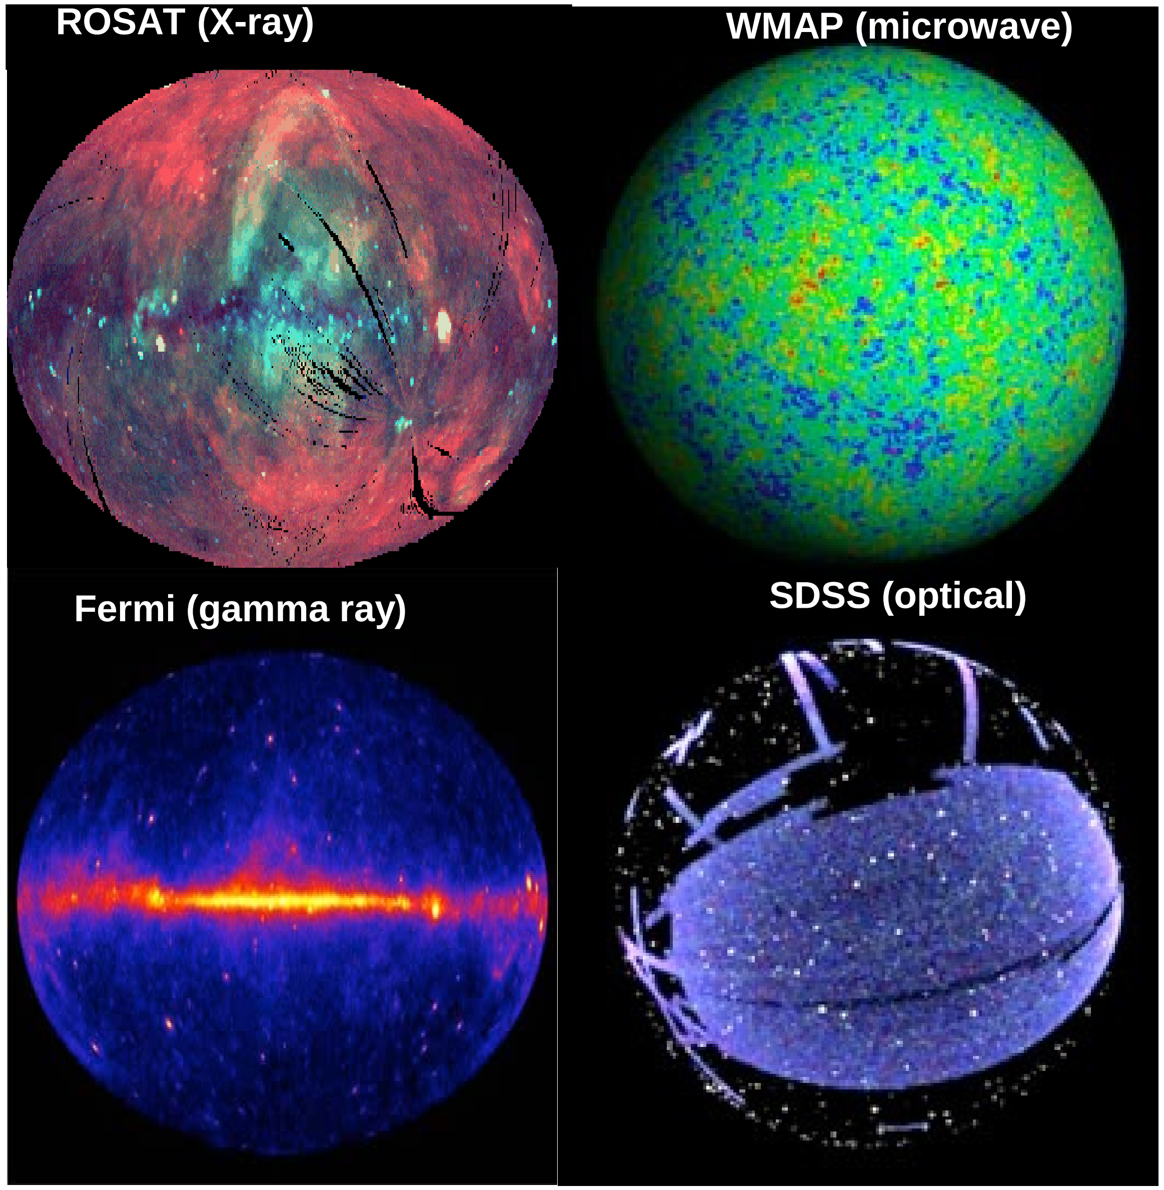


The Source of Knowledge: Sky Surveys

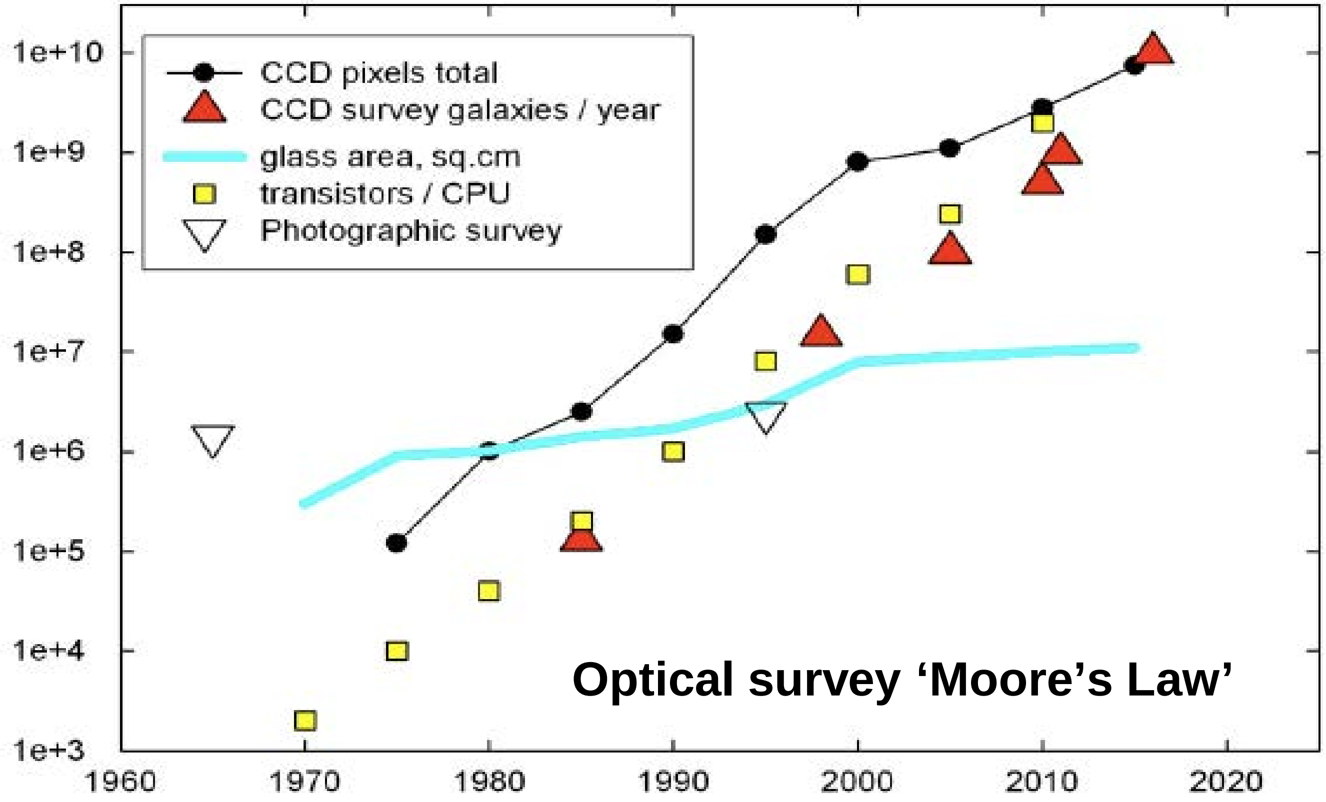


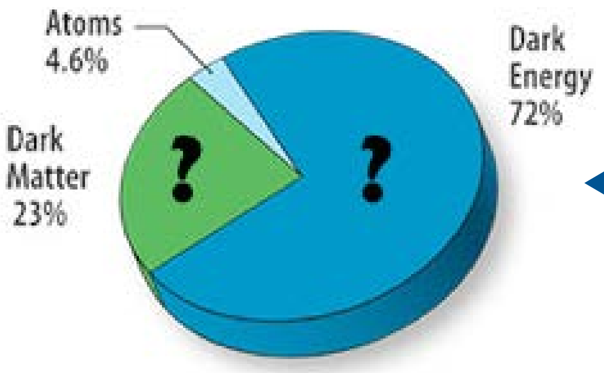


The Cosmic Puzzle: Who ordered the rest of it?

~1% level

# Key Tools: The Correlation Function and Power Spectrum

- All structure formation probes of cosmology in some way study density fluctuations
- Hence desire **robust** ways to characterize clustering statistics of the underlying mass field and its tracers (e.g., galaxies)
- The (2-point) **correlation function** is the excess probability of finding an object pair separated by a distance $r_{12}$ compared to that for a random distribution:
  $dP=n^2(1+\xi(r_{12}))dV_1dV_2$
  where $n$ is the mean density; the power spectrum $P(k)$ is the Fourier transform of the correlation function
- The primordial fluctuations, as best known currently, are Gaussian, and completely specified by 2-point statistics
- Nonlinear structure formation induces non-zero higher point correlation functions

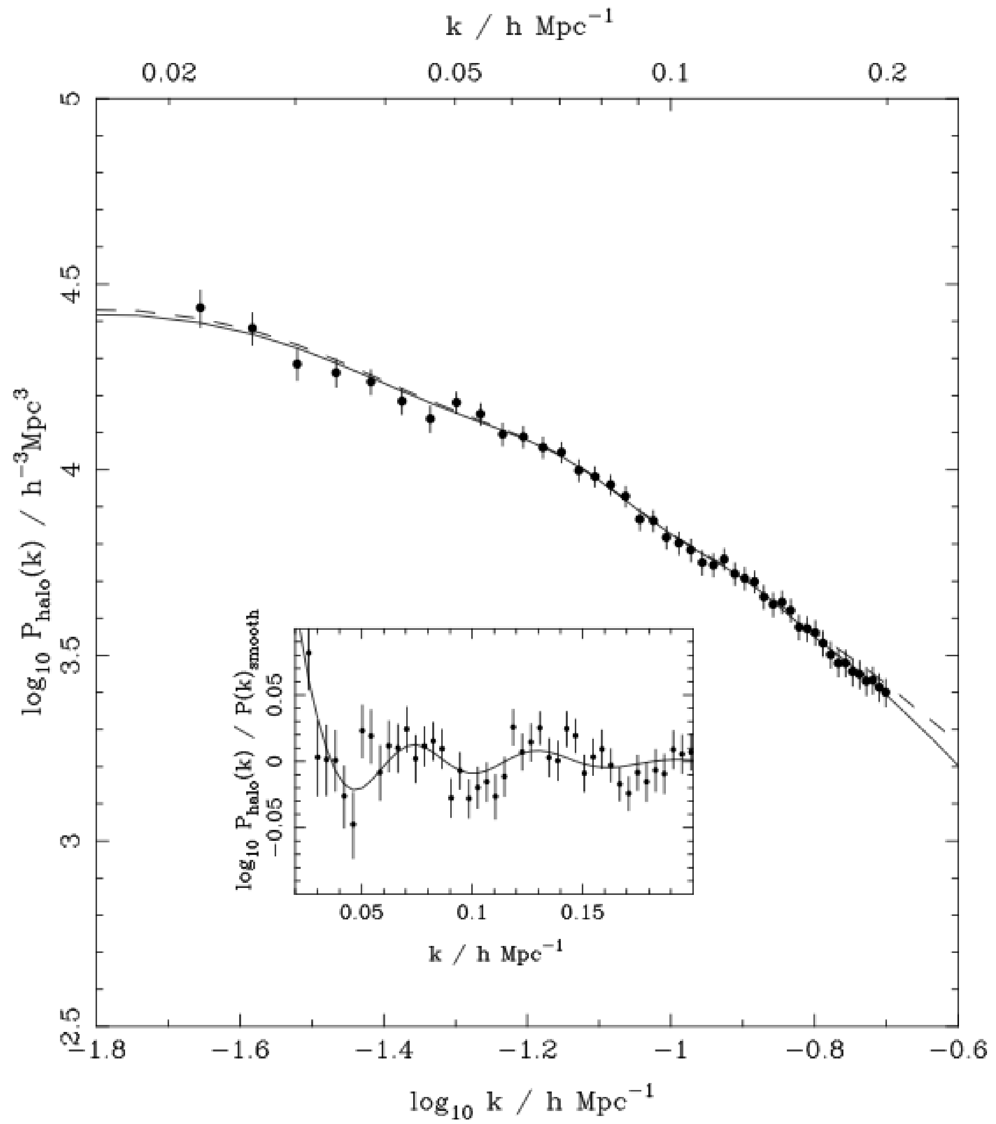


**SDSS DR7 galaxy power spectrum**

# The Precision Cosmology Revolution

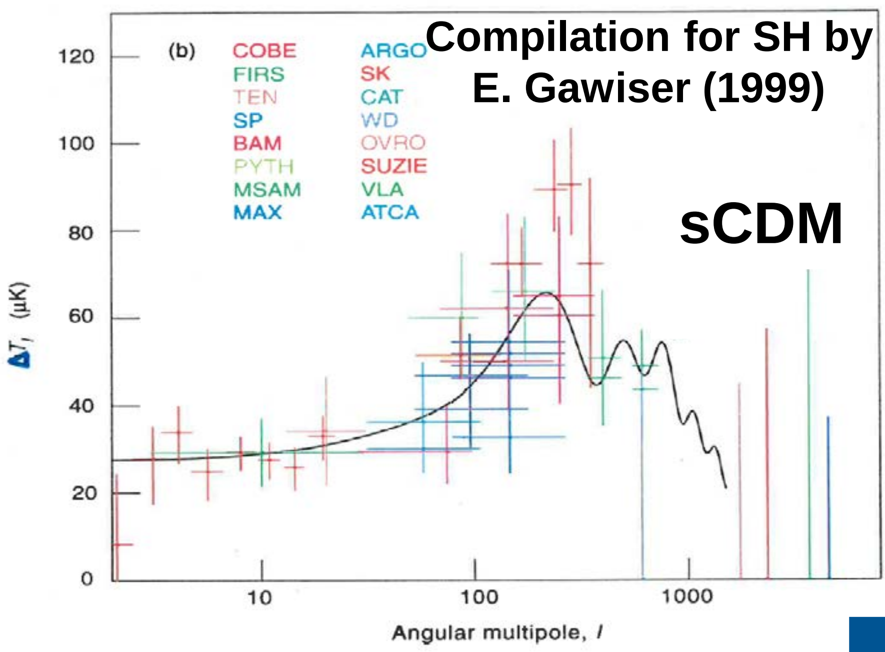


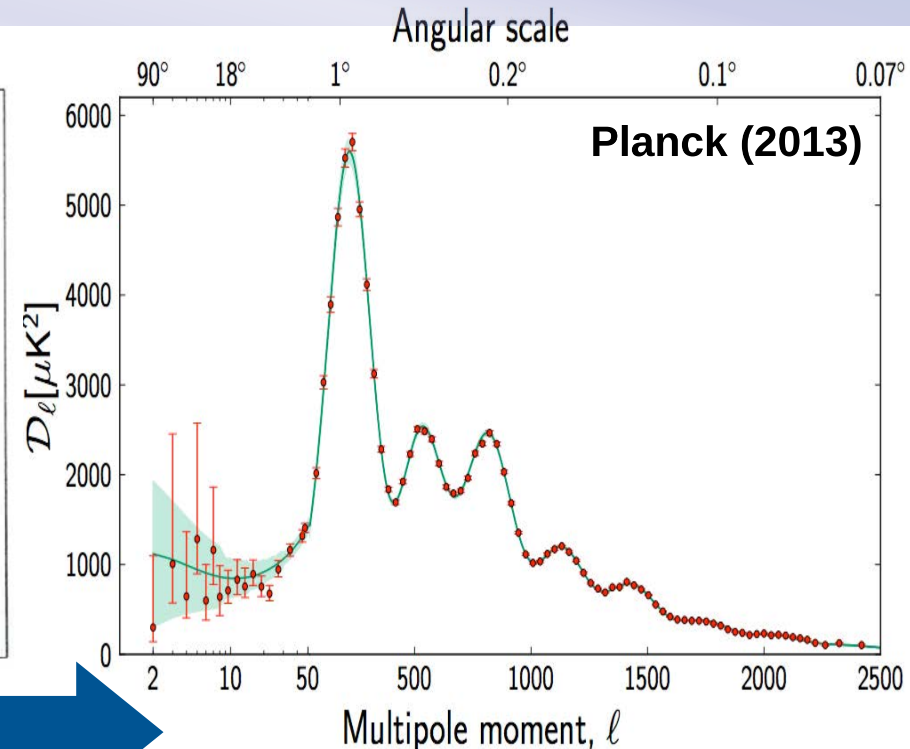


CMB

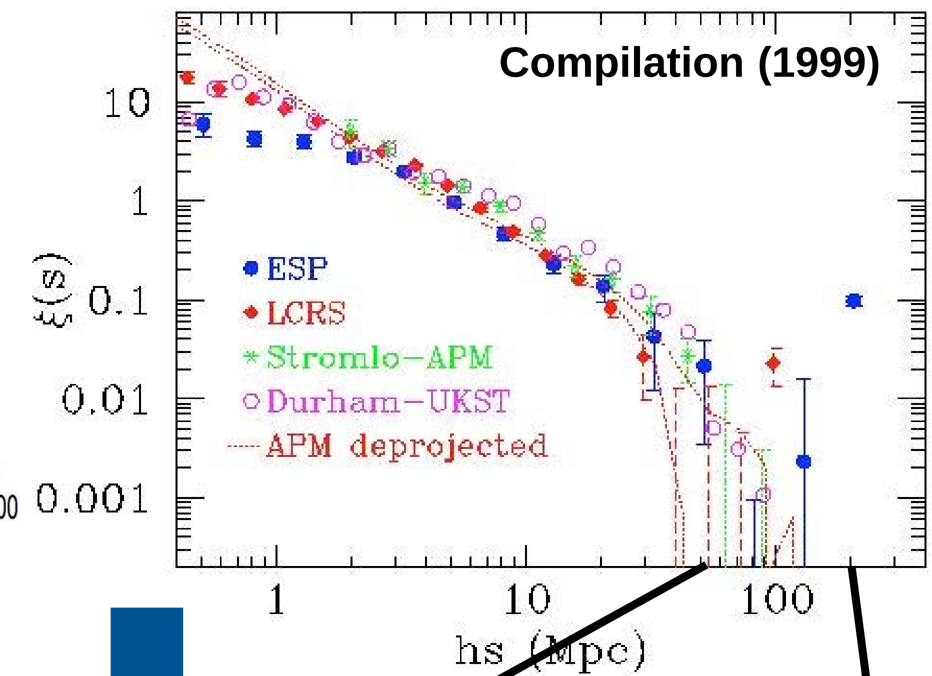


LSS

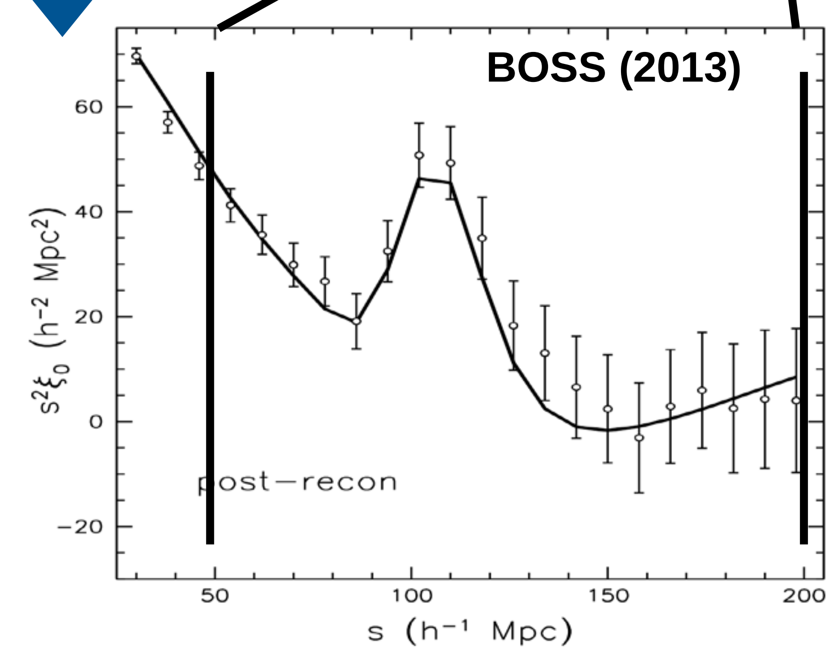


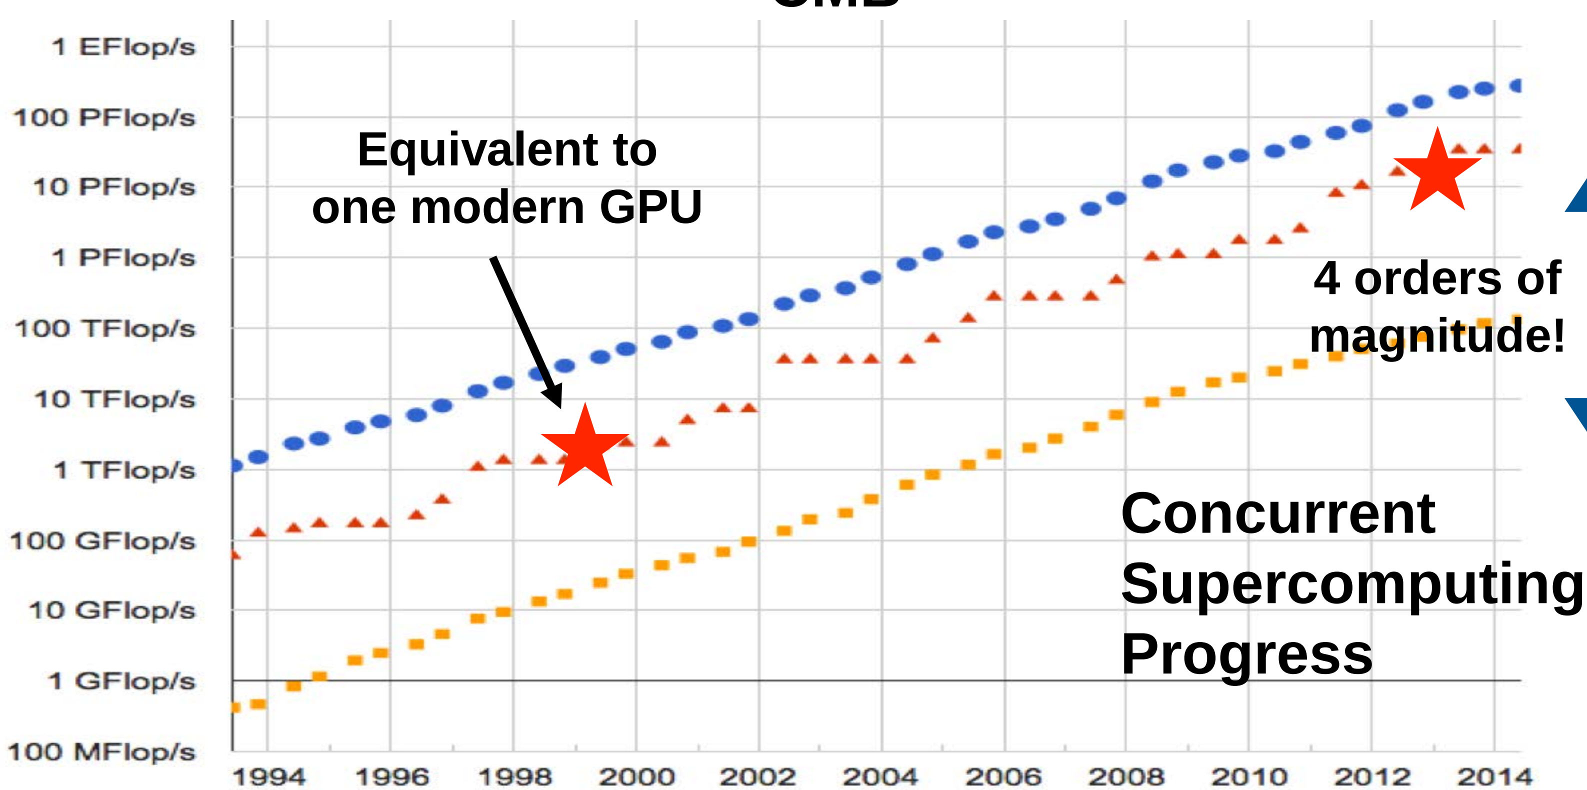

# Data ‘Overload’: Observations of Cosmic Structure

- Cosmology=Physics+Statistics
  - Mapping the sky with large-area surveys across multiple wave-bands, at remarkably low levels of statistical error

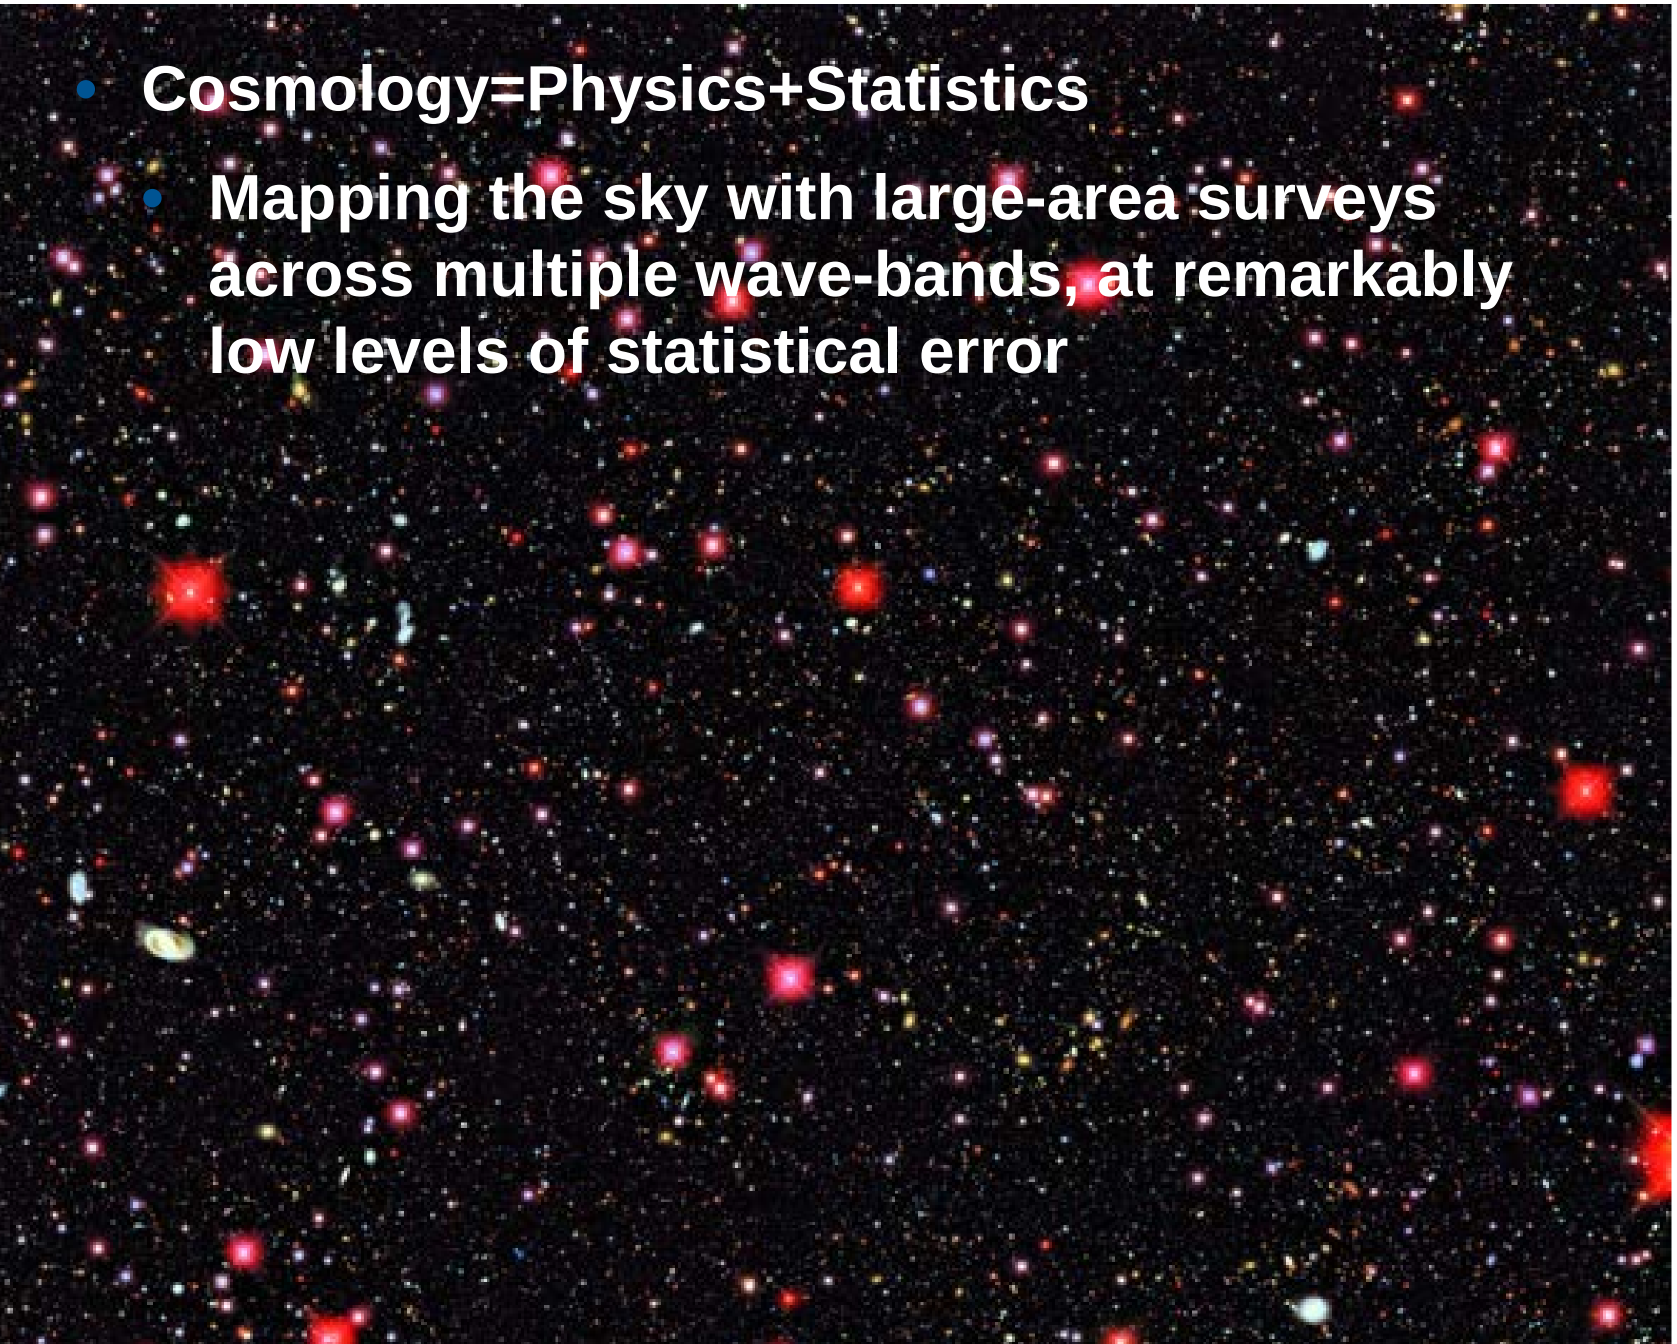

Galaxies in a moon-sized patch (Deep Lens Survey). LSST will cover 50,000 times this size (~400PB of data)


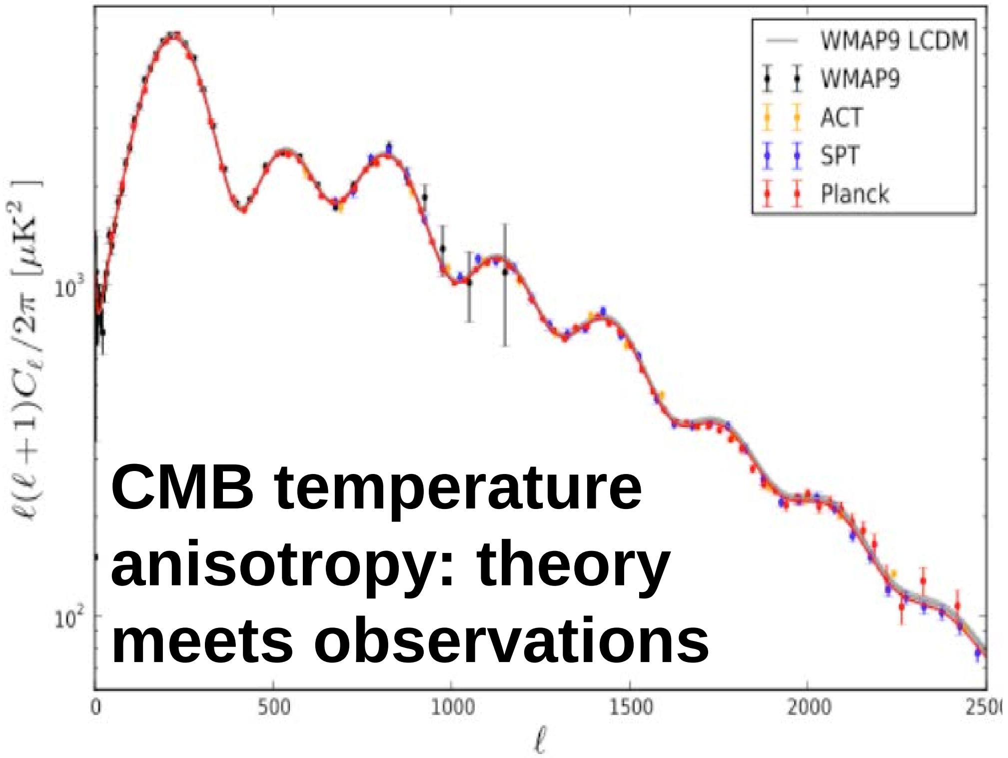


CMB temperature anisotropy: theory meets observations

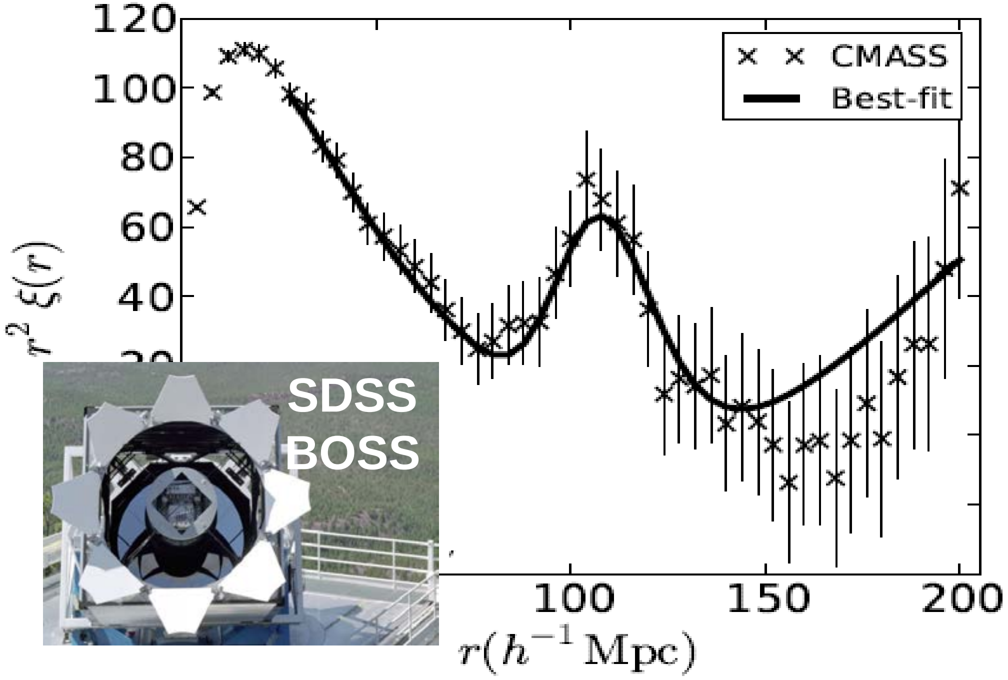


The same signal in the galaxy distribution

# Computational Cosmology

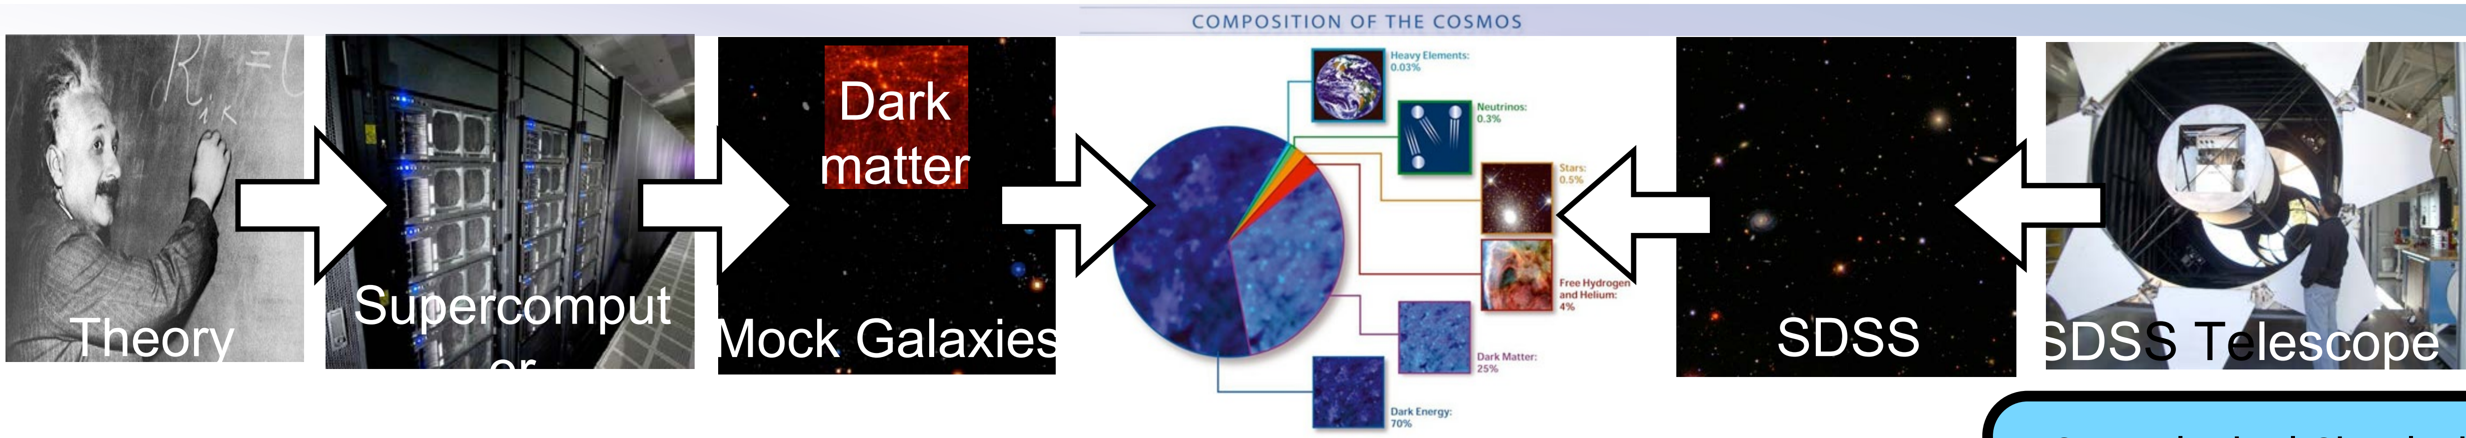


- **Three Roles of Cosmological Simulations**
  - Basic theory of cosmological probes
  - Production of high-fidelity ‘mock skys’ for end-to-end tests of the observation/analysis chain
  - Essential component of analysis toolkits
- **Extreme Simulation and Analysis Challenges**
  - Large dynamic range simulations; control of subgrid modeling and feedback mechanisms
  - Design and implementation of complex analyses on large datasets; new fast (approximate) algorithms
  - Solution of large statistical inverse problems of scientific inference (many parameters, ~10-100) at the ~1% level

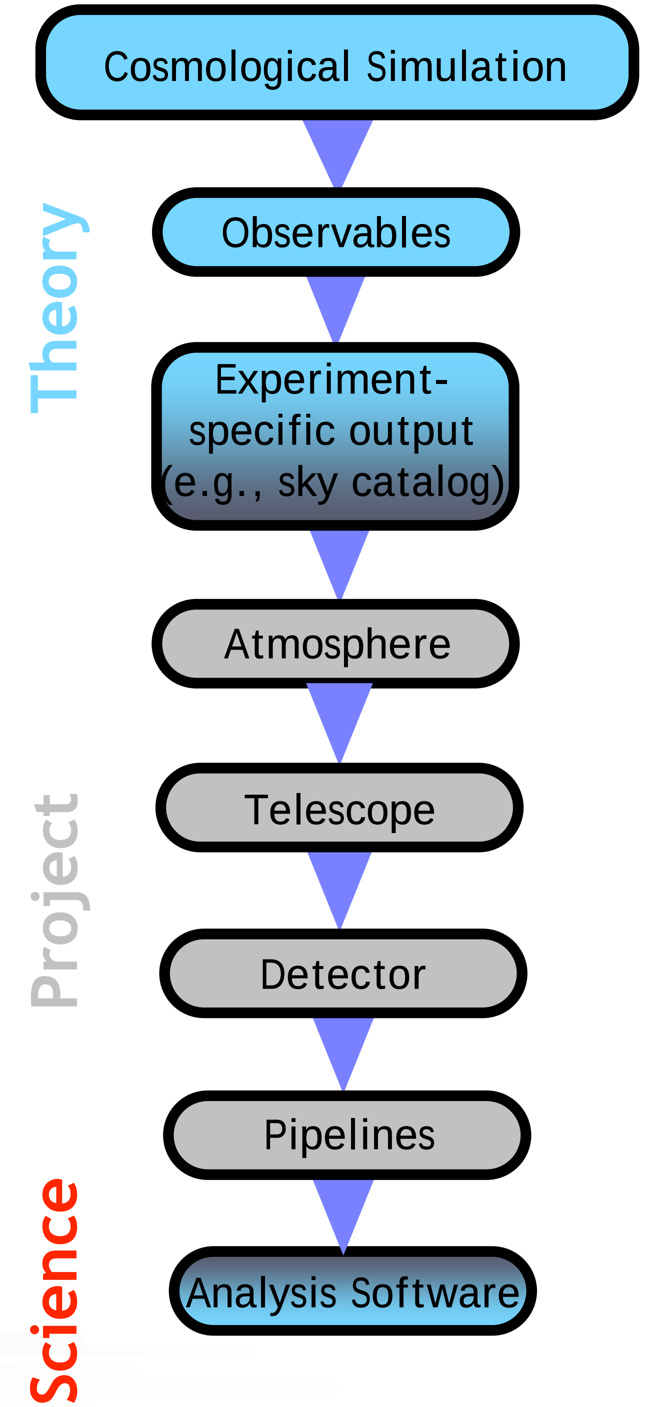

# Connecting Theory and Observations: Challenges & Opportunities

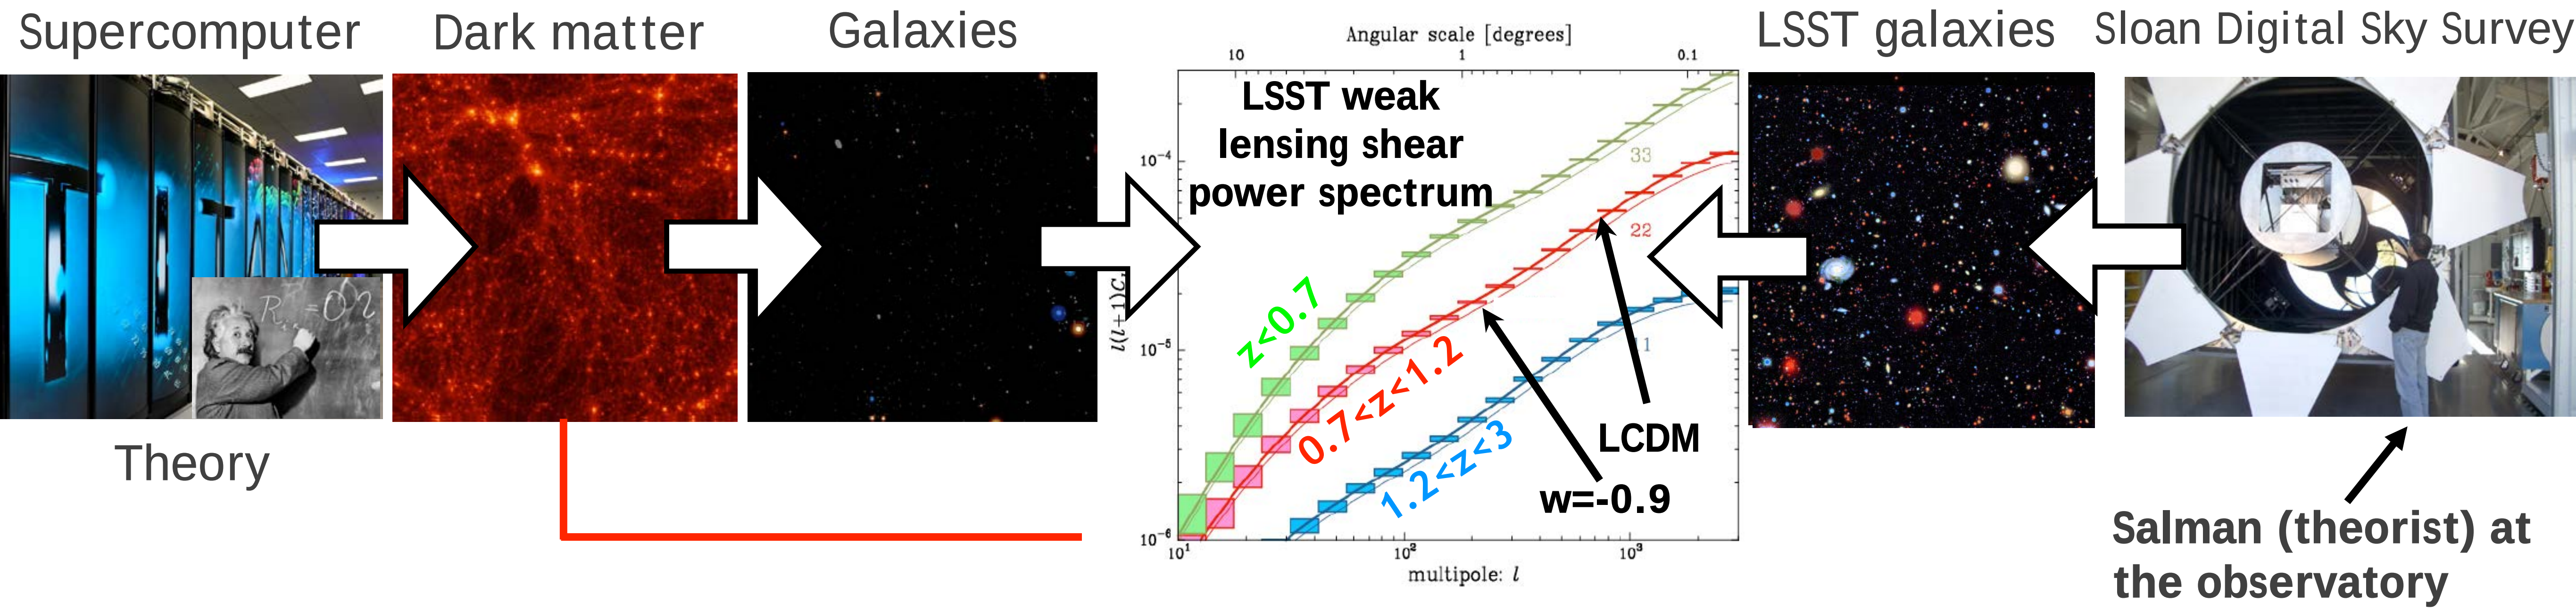


- Error bars will shrink dramatically
  - Predictions have to be accurate at the sub-percent level
  - Modeling and understanding of systematics becomes ever more important (e.g. baryonic effects)
  - We can go beyond LCDM and explore new fundamental physics: neutrinos, modified gravity, dynamical dark energy, self-interacting dark matter …
- Surveys will become deeper and resolve much fainter galaxies
  - Synthetic sky map making becomes more difficult, more physics
  - Much higher resolution simulations will be required
  - New cosmological probes, cross-correlations will be available

# Precision Cosmology: ‘Big Data’ Meets Supercomputing

**Supercomputer Simulation Campaign**

**Major stats + ML+ sampling + optimization collaboration**

**Simulations + CCF**

**Emulator based on Gaussian Process Interpolation in High-Dimensional Spaces**

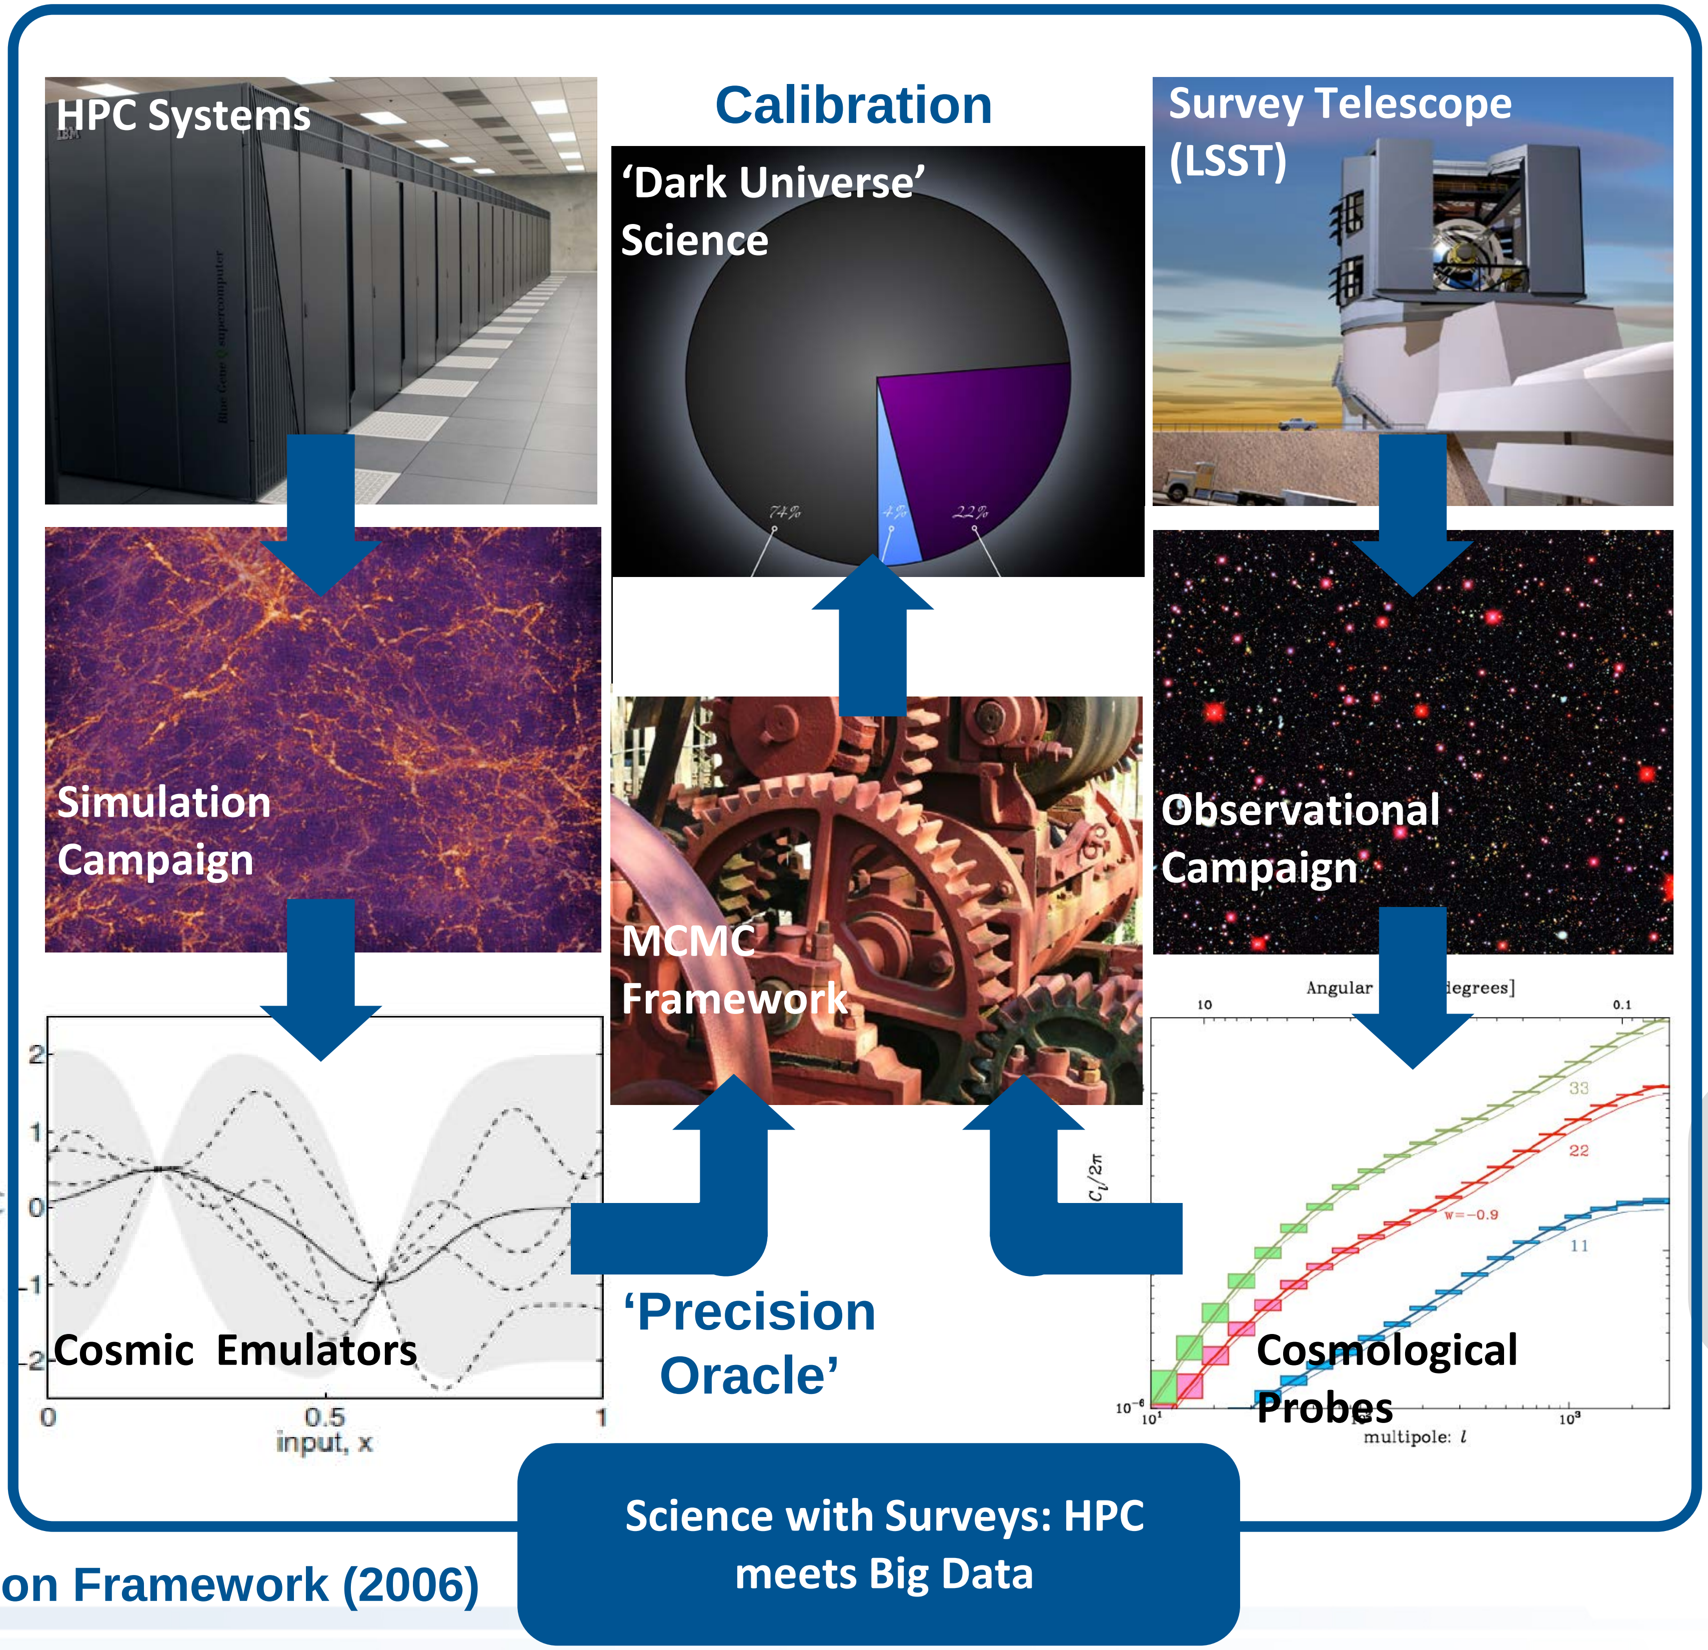


**Mapping the Sky with Survey Instruments**

**Observations: Statistical error bars will ‘disappear’ soon!**

**CCF= Cosmic Calibration Framework (2006)**

# Example: Analytics/Workflow Complexity

**Gaussian Random Field Initial Conditions** → **High-Resolution N-Body Code (HACC)** → **Multiple Outputs Halo/Sub-Halo Identification** → **Halo Merger Trees**

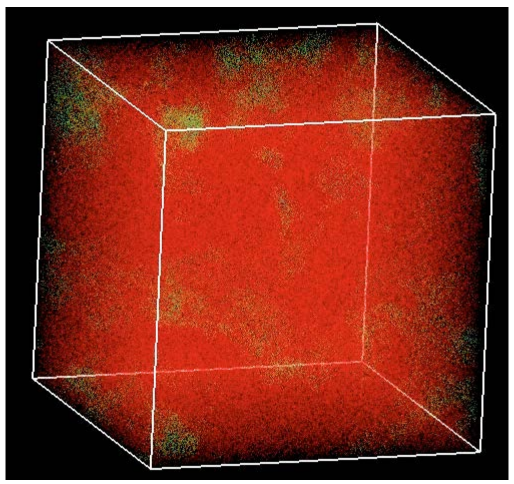

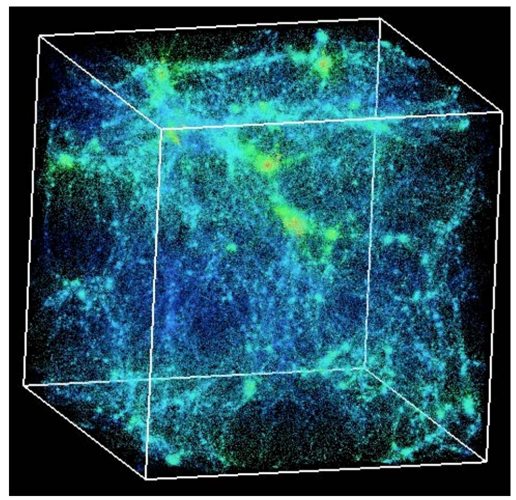

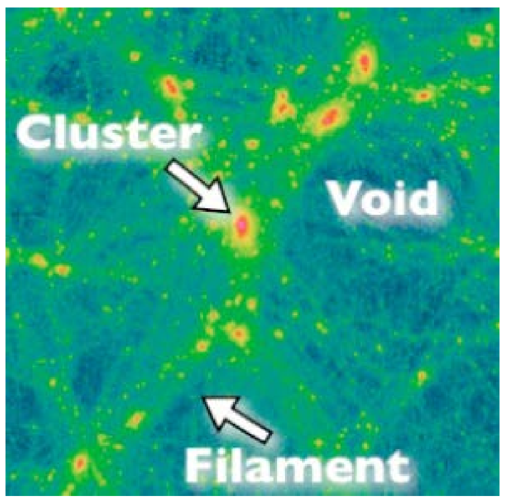


**Semi-Analytic Modeling Code (Galacticus)** → **Galaxy Catalog** → **Realistic Image Catalog** → **Atmosphere and Instrument Modeling** → **Data Management Pipeline** → **Data Analysis Pipeline** → **Scientific Inference Framework**

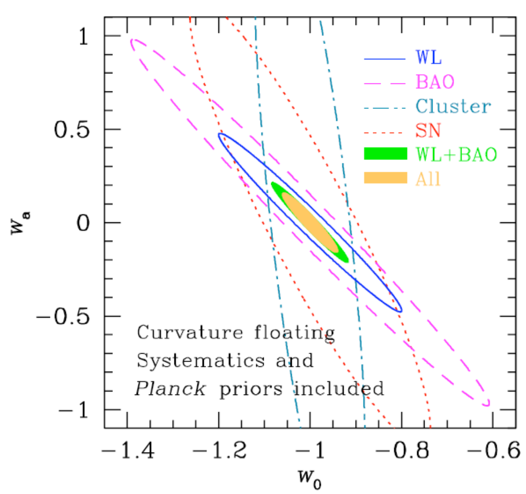


- **Simulation Campaigns:** Statistics of virtual universes; construction of emulators
- **Modeling the Measurement:** End-to-End modeling necessary to understand crucial systematic errors
- **PDACS:** Custom workflow system under development
- **Data-Intensive Computing:** New architectures needed

# Large Scale Structure: Vlasov-Poisson Equation

$$\frac{\partial f_i}{\partial t} + \dot{\mathbf{x}}\frac{\partial f_i}{\partial \mathbf{x}} - \nabla\phi\frac{\partial f_i}{\partial \mathbf{p}} = 0, \qquad \mathbf{p} = a^2\dot{\mathbf{x}},$$

$$\nabla^2\phi = 4\pi G a^2(\rho(\mathbf{x},t) - \langle\rho_{\rm dm}(t)\rangle) = 4\pi G a^2 \Omega_{\rm dm}\delta_{\rm dm}\rho_{\rm cr},$$

$$\delta_{\rm dm}(\mathbf{x},t) = (\rho_{\rm dm} - \langle\rho_{\rm dm}\rangle)/\langle\rho_{\rm dm}\rangle),$$

$$\rho_{\rm dm}(\mathbf{x},t) = a^{-3}\sum_i m_i \int d^3\mathbf{p} f_i(\mathbf{x},\dot{\mathbf{x}},t).$$

**Cosmological Vlasov-Poisson Equation**

- **Properties of the Cosmological Vlasov-Poisson Equation:**
  - **6-D PDE with long-range interactions, no shielding, all scales matter, models gravity-only, collisionless evolution**
  - **Extreme dynamic range in space and mass (in many applications, million to one, ‘everywhere’)**
  - **Jeans instability drives structure formation at all scales from smooth Gaussian random field initial conditions**

# Large Scale Structure Simulation Requirements

- **Force and Mass Resolution:**
  - Galaxy halos ~100kpc, hence force resolution has to be ~kpc; with Gpc box-sizes, a dynamic range of a million to one
  - Ratio of largest object mass to lightest is ~10000:1
- **Physics:**
  - Gravity dominates at scales greater than ~0.1 Mpc
  - Small scales: galaxy modeling, semi-analytic methods to incorporate gas physics/feedback/star formation
- **Computing 'Boundary Conditions':**
  - Total memory in the PB+ class
  - Performance in the 10 PFlops+ class
  - Wall-clock of ~days/week, in situ analysis

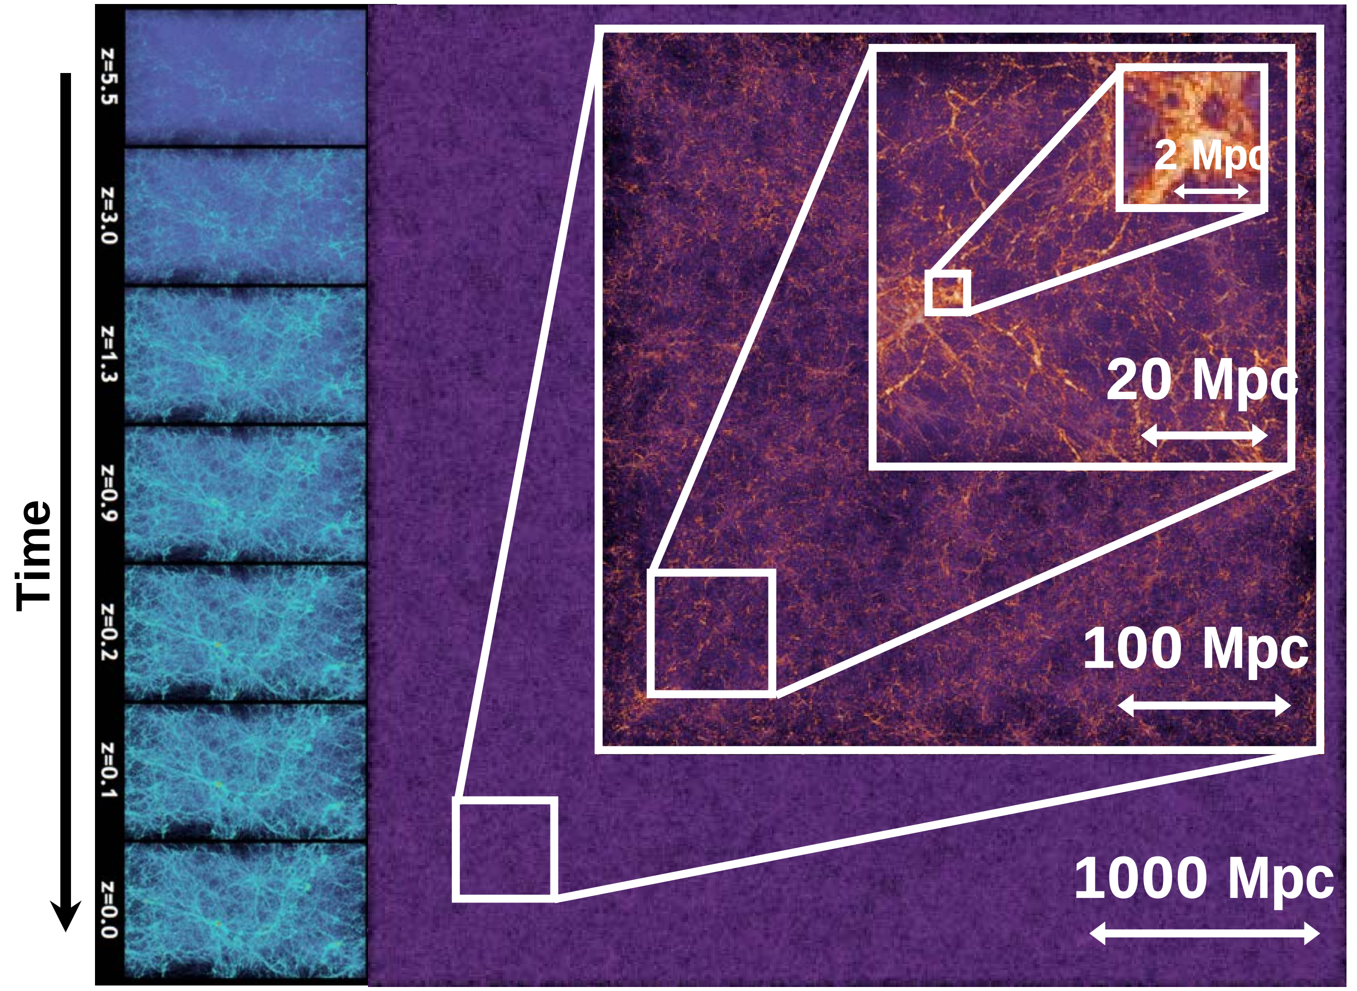


Gravitational Jeans Instability: 'Outer Rim' run with 1.1 trillion particles

**Key motivation for HACC: Can the Universe be run as a short computational 'experiment'?**

# Combating Architectural Diversity with HACC

- **Architecture-independent performance/scalability:** ‘Universal’ top layer + ‘plug in’ node-level components; minimize data structure complexity and data motion
- **Programming model:** ‘C++/MPI + X’ where X = OpenMP, Cell SDK, OpenCL, CUDA, --
- **Algorithm Co-Design:** Multiple algorithm options, stresses accuracy, low memory overhead, no external libraries in simulation path
- **Analysis tools:** Major analysis framework, tools deployed in stand-alone and in situ modes

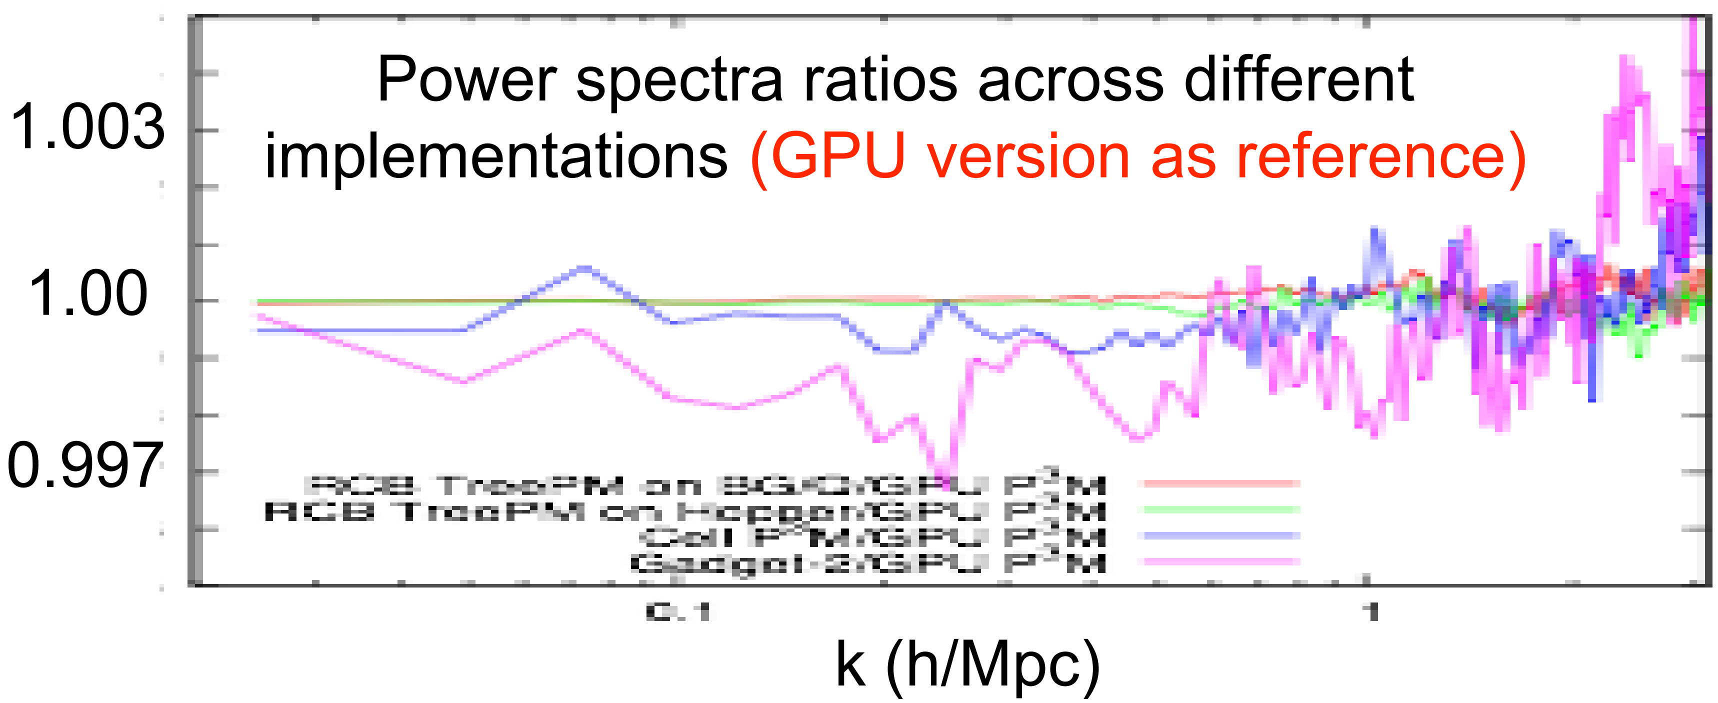


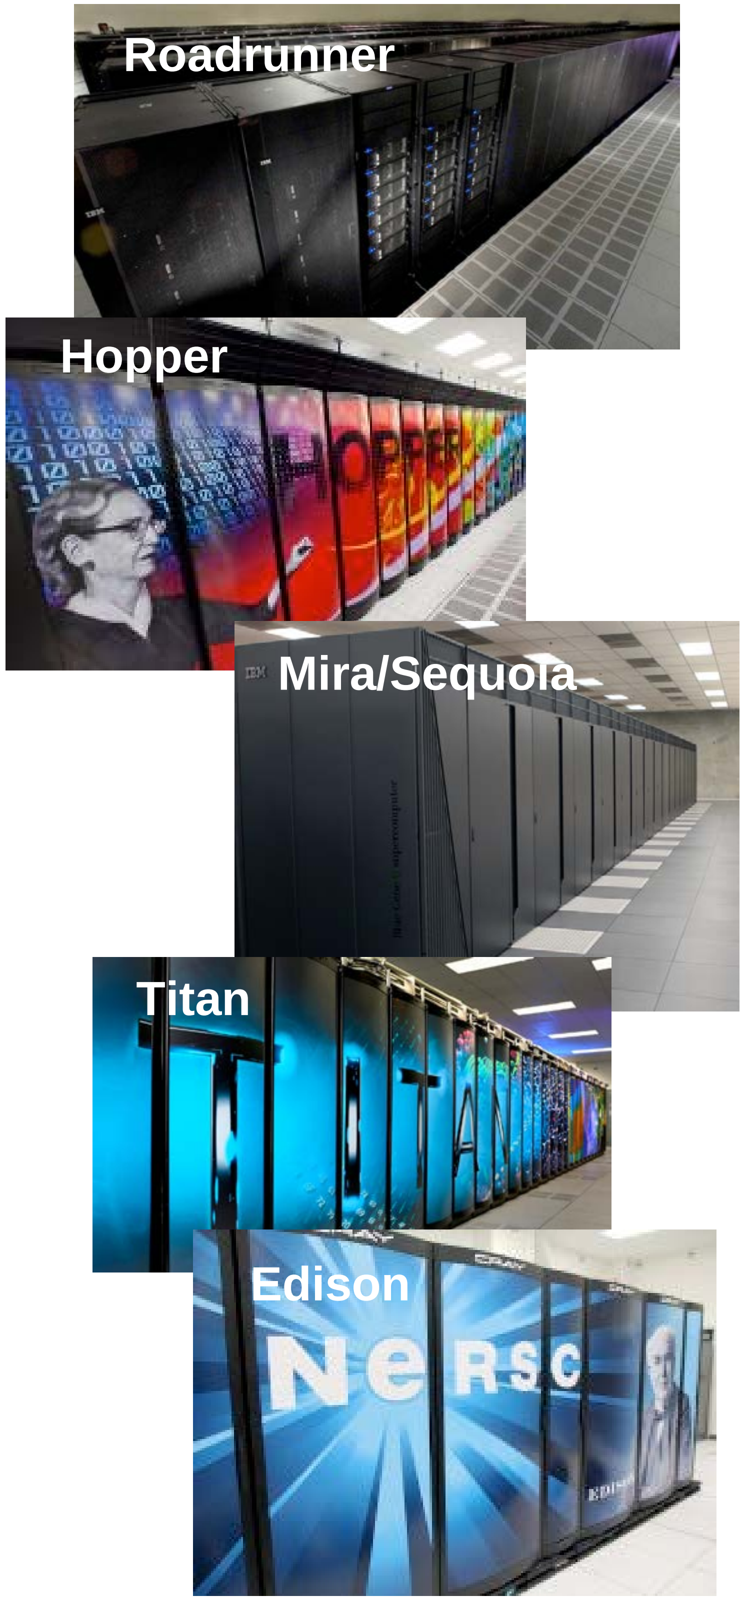

# ‘HACC In Pictures’

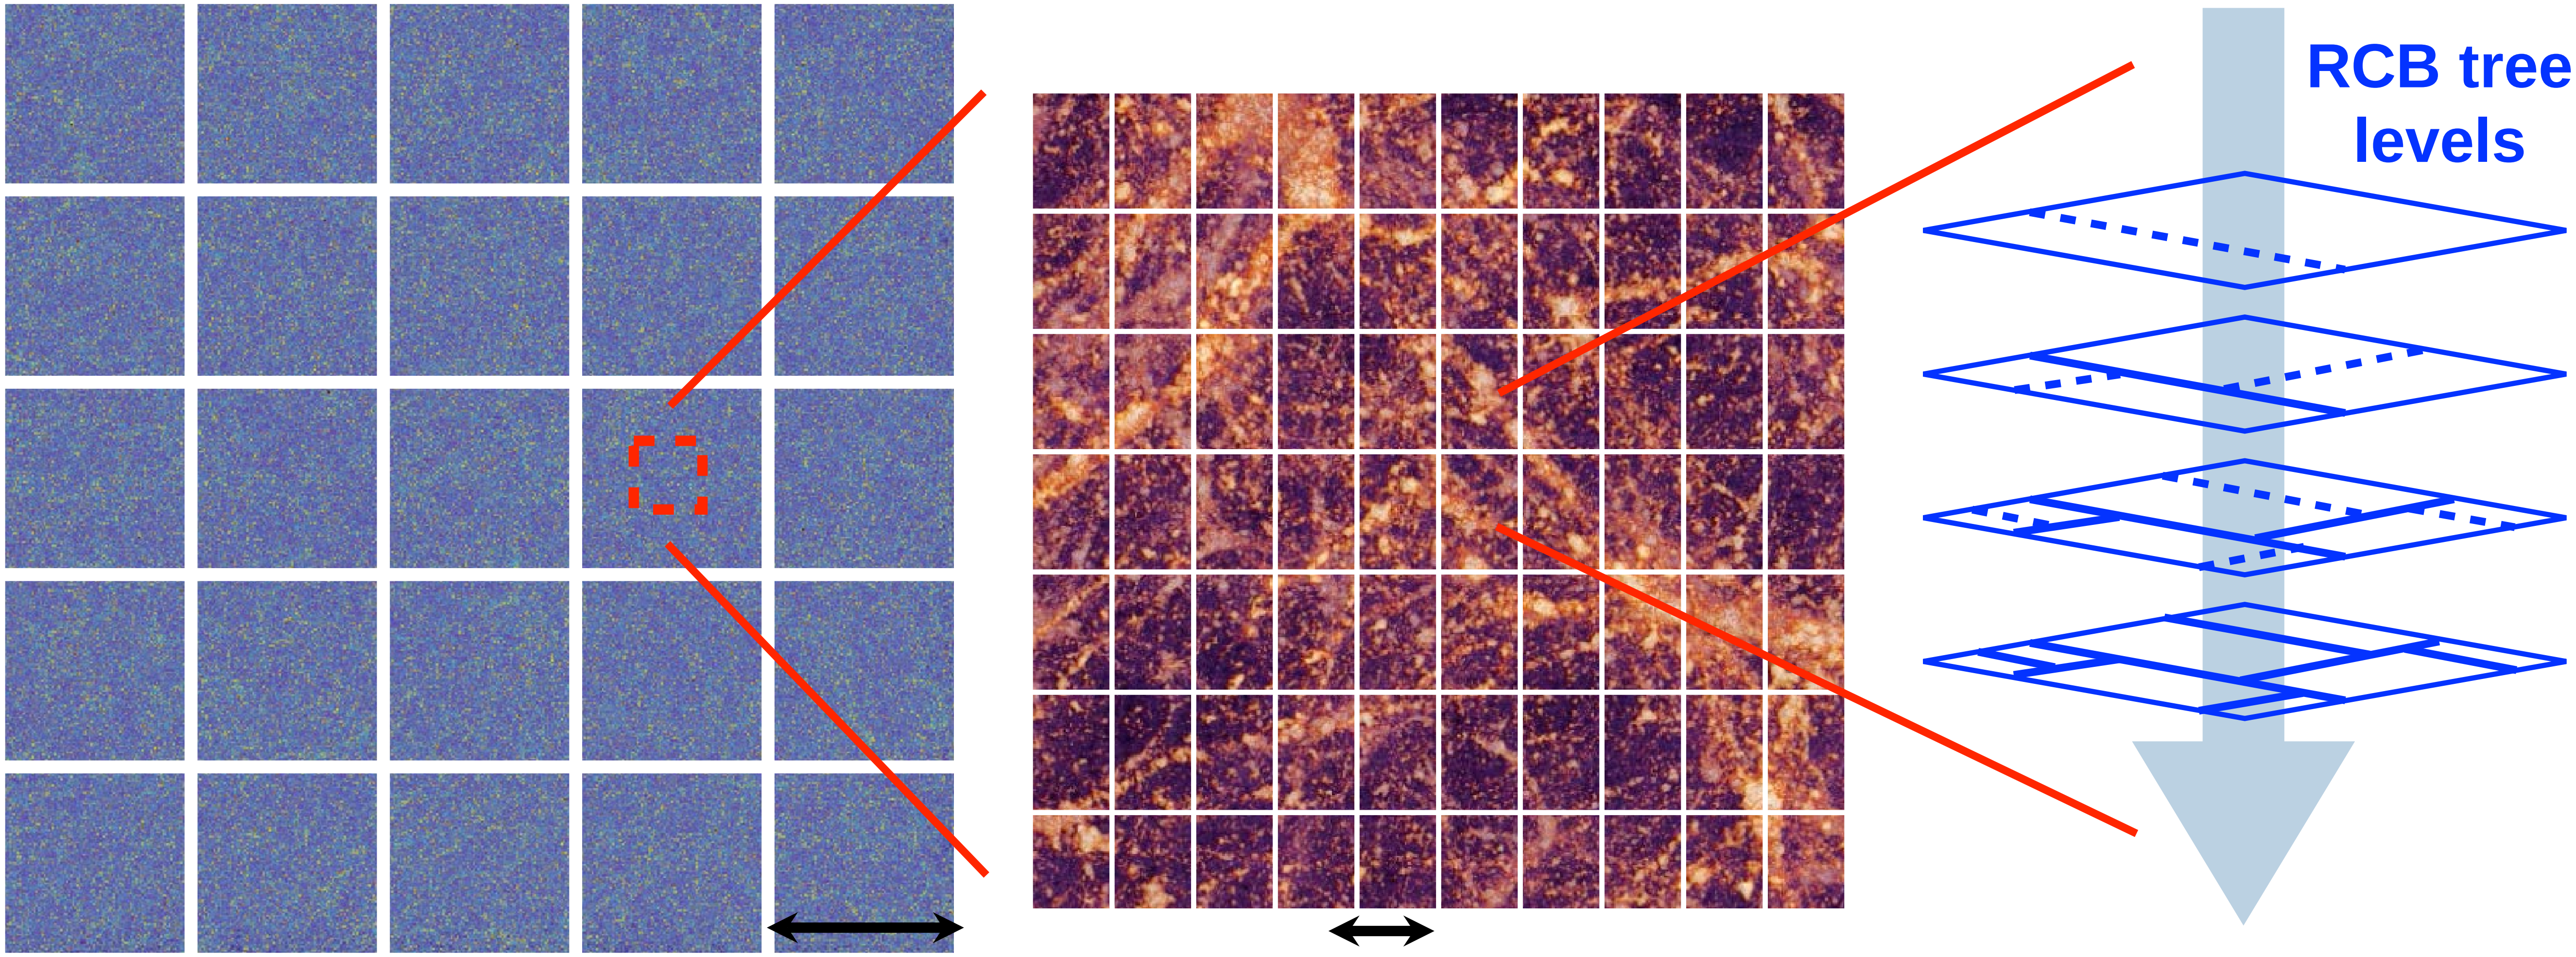


~50 Mpc

~1 Mpc

**HACC Top Layer:**
**3-D domain decomposition with particle replication at boundaries (‘overloading’) for Spectral PM algorithm (long-range force)**

**Host-side**

**HACC ‘Nodal’ Layer:**
**Short-range solvers employing combination of flexible chaining mesh and RCB tree-based force evaluations**

**GPU: two options, P3M vs. TreePM**

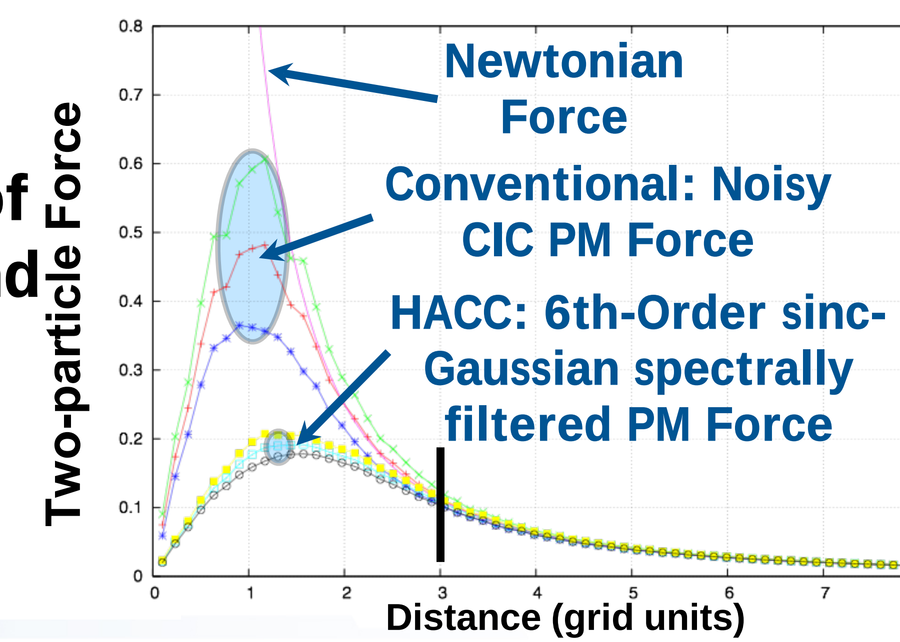

# Accelerated Systems: Specific Issues — Titan

## Imbalances and Bottlenecks

- Memory is primarily host-side (32 GB vs. 6 GB) (against Roadrunner's 16 GB vs. 16 GB), important thing to think about (in case of HACC, the grid/particle balance)
- PCIe is a key bottleneck; overall interconnect B/W does not match Flops
- There's no point in 'sharing' work between the CPU and the GPU, performance gains will be minimal -- GPU must dominate
- The only reason to write a code for such a system is if you can truly exploit its power (2 X CPU is a waste of effort!)

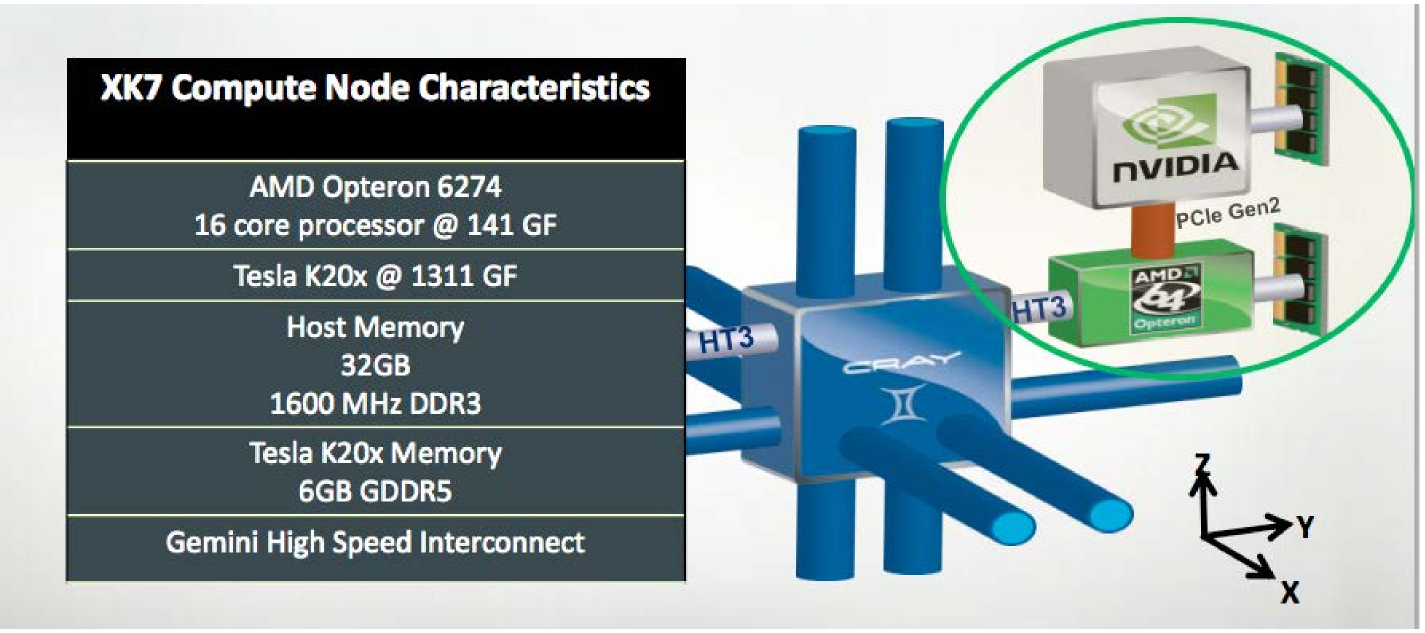


## Strategies for Success

- It's (still) all about understanding and controlling data motion
- Rethink your code and even approach to the problem
- Isolate hotspots, and design for portability around them (modular programming)
- Like it or not, pragmas will never be the full answer

# HACC: Algorithmic Features

- **Fully Spectral Particle-Mesh Solver:** 6th-order Green function, 4th-order Super-Lanczos derivatives, high-order spectral filtering, high-accuracy polynomial for short-range forces
- **Custom Parallel FFT:** Pencil-decomposed, high-performance FFT (~15K^3)
- **Particle Overloading:** Particle replication at ‘node’ boundaries to reduce/delay communication (intermittent refreshes), important for accelerated systems
- **Flexible Chaining Mesh:** Used to optimize tree and P3M methods
- **Optimal Splitting of Gravitational Forces:** Spectral Particle-Mesh melded with direct and RCB (‘fat leaf’) tree force solvers (PPTPM), short hand-over scale (dynamic range splitting ~ 10,000 X 100); pseudo-particle method for multipole expansions
- **Mixed Precision:** Optimize memory and performance (GPU-friendly!)
- **Optimized Force Kernels:** High performance without assembly
- **Adaptive Symplectic Time-Stepping:** Symplectic sub-cycling of short-range force timesteps; adaptivity from automatic density estimate via RCB tree
- **Custom Parallel I/O:** Topology aware parallel I/O with lossless compression (factor of 2); 1.5 trillion particle checkpoint in 4 minutes at ~160GB/sec on Mira

# HACC on Titan: GPU Implementation Performance

- P3M kernel runs at 1.6TFlops/node at 40.3% of peak (73% of algorithmic peak)
- TreePM kernel was run on 77% of Titan at 20.54 PFlops at almost identical performance on the card
- Because of less overhead, P3M code is (currently) faster by factor of two in time to solution
- New load balancing method

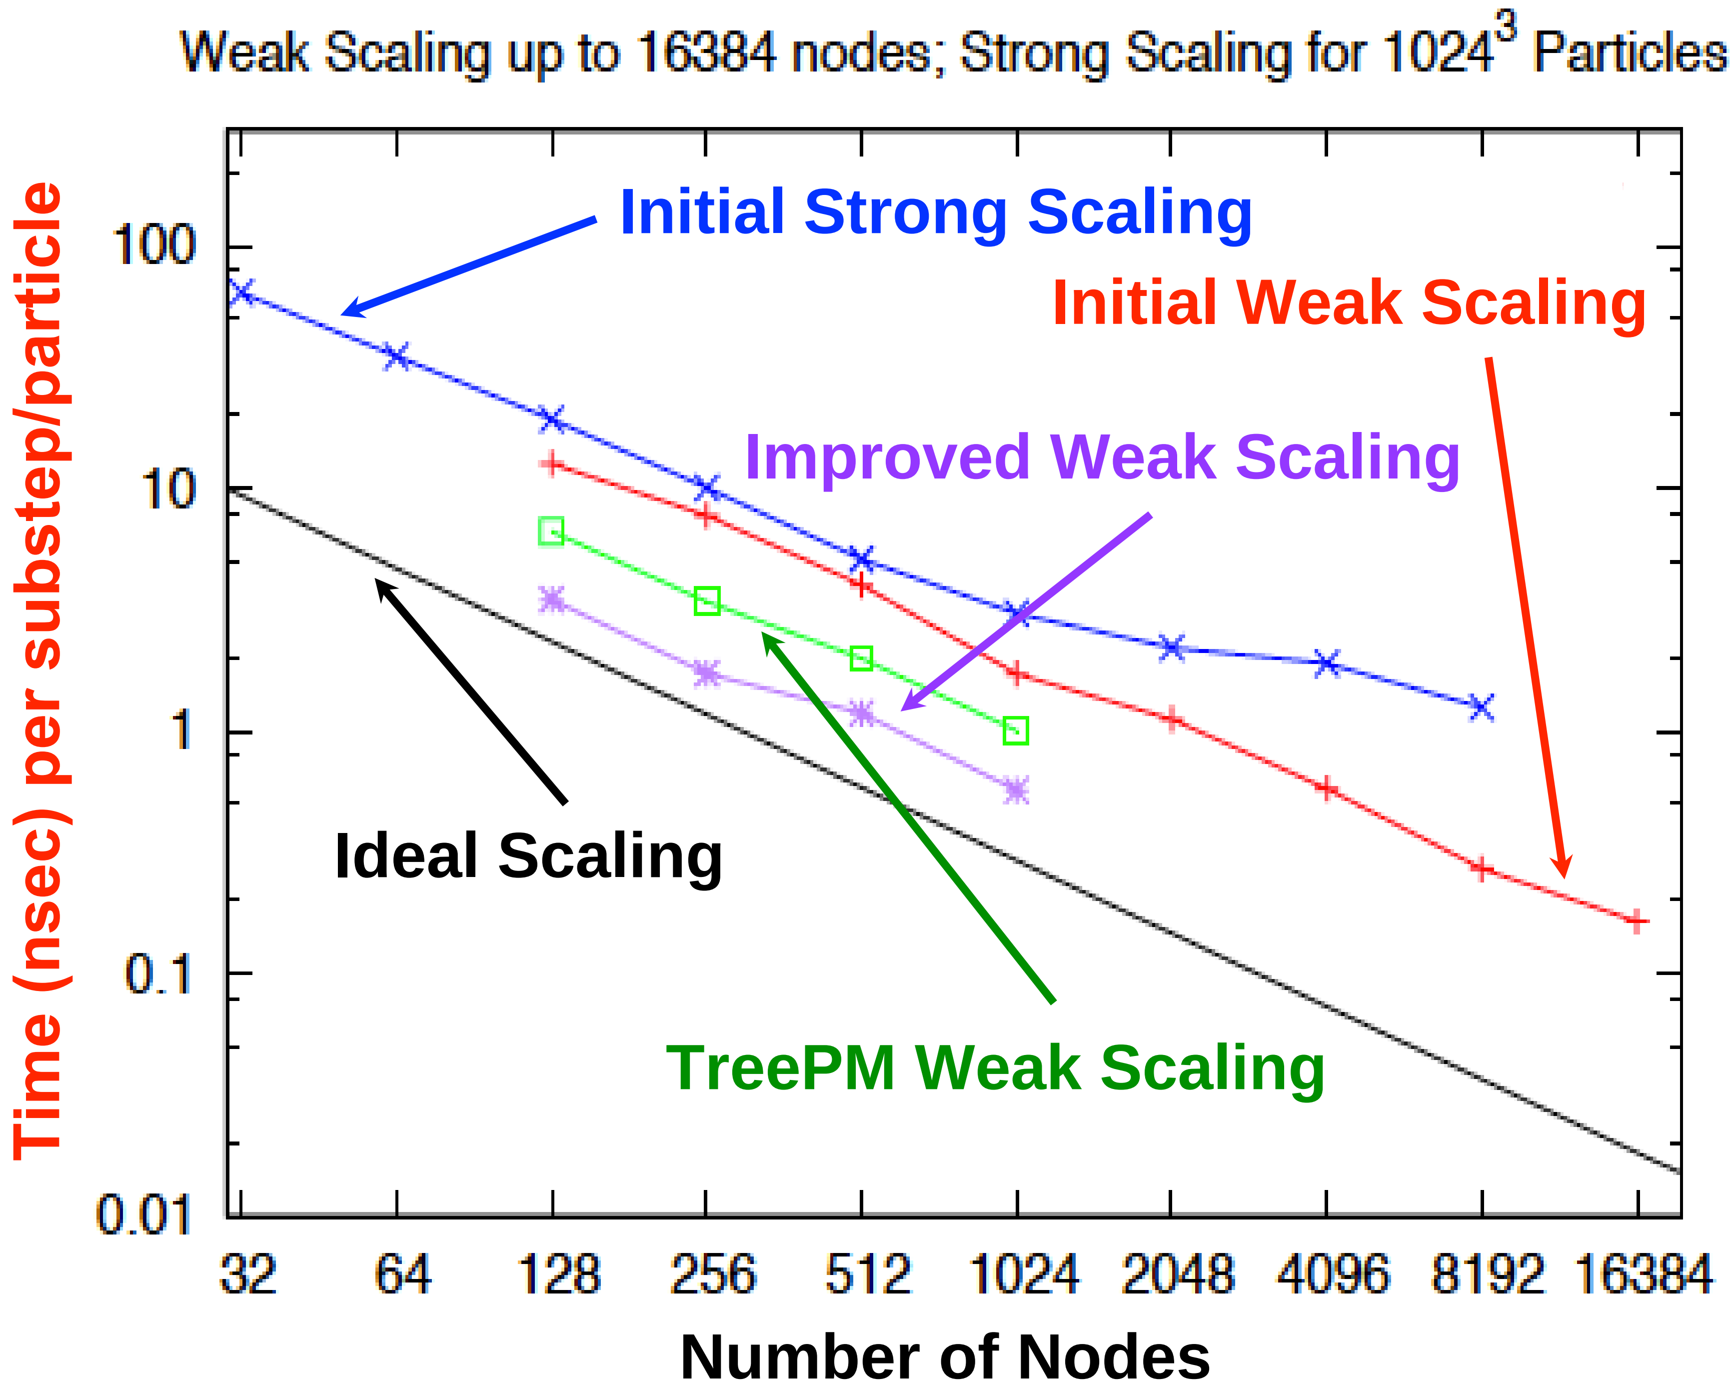


99.2% Parallel Efficiency

# The Q Continuum and the Outer Rim Simulations

## Simulating the LCDM Universe with Unprecedented Volume and Resolution

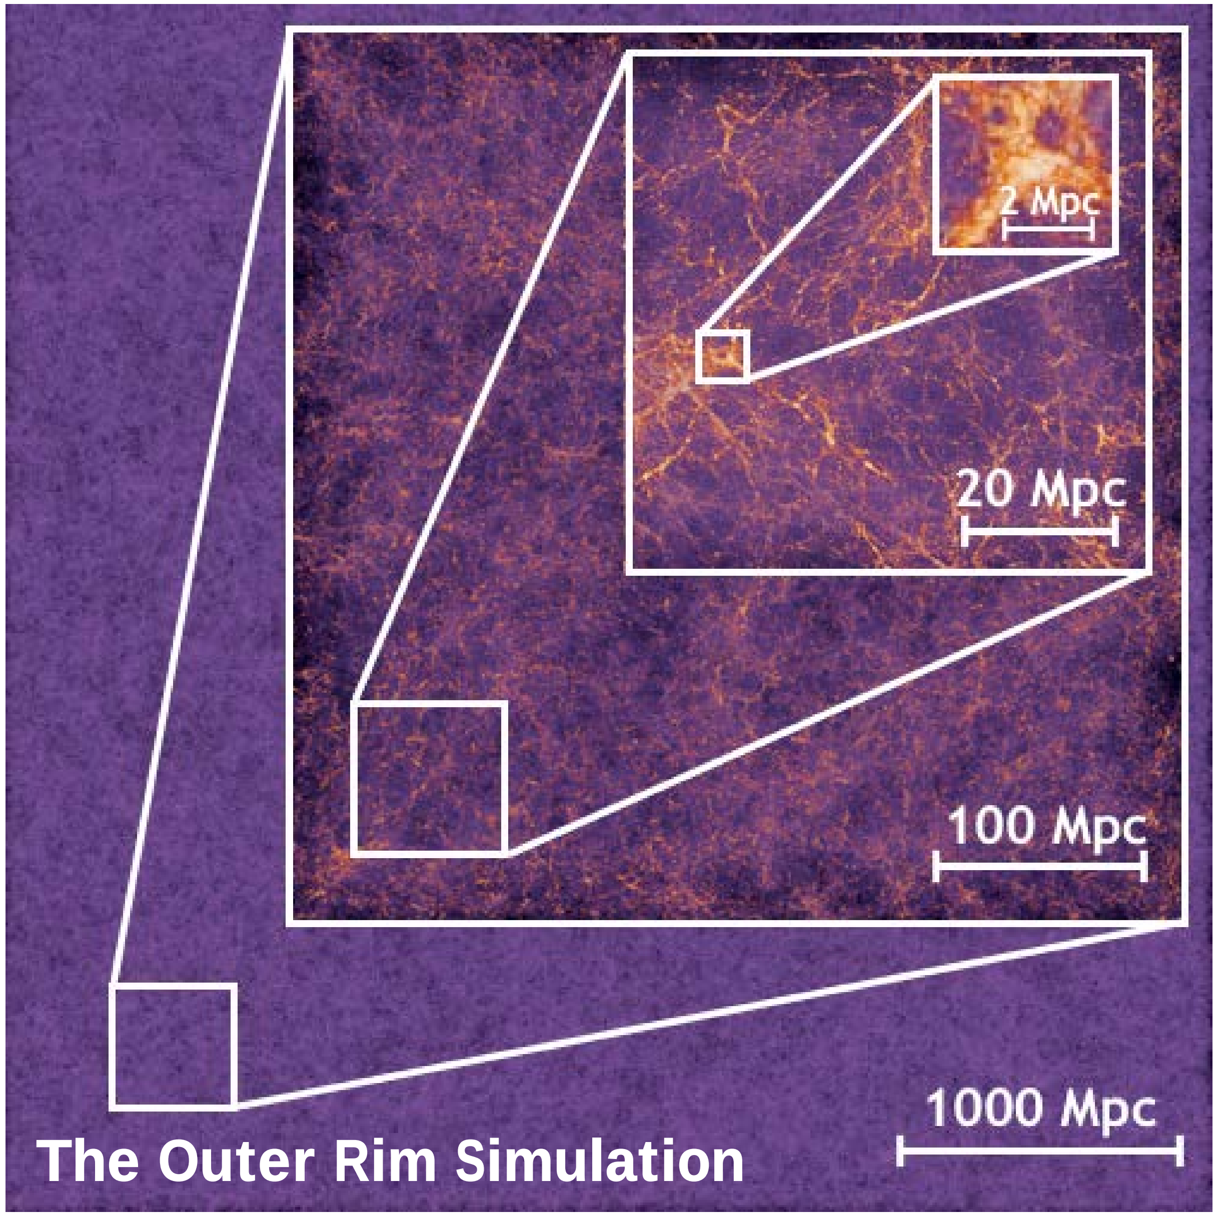


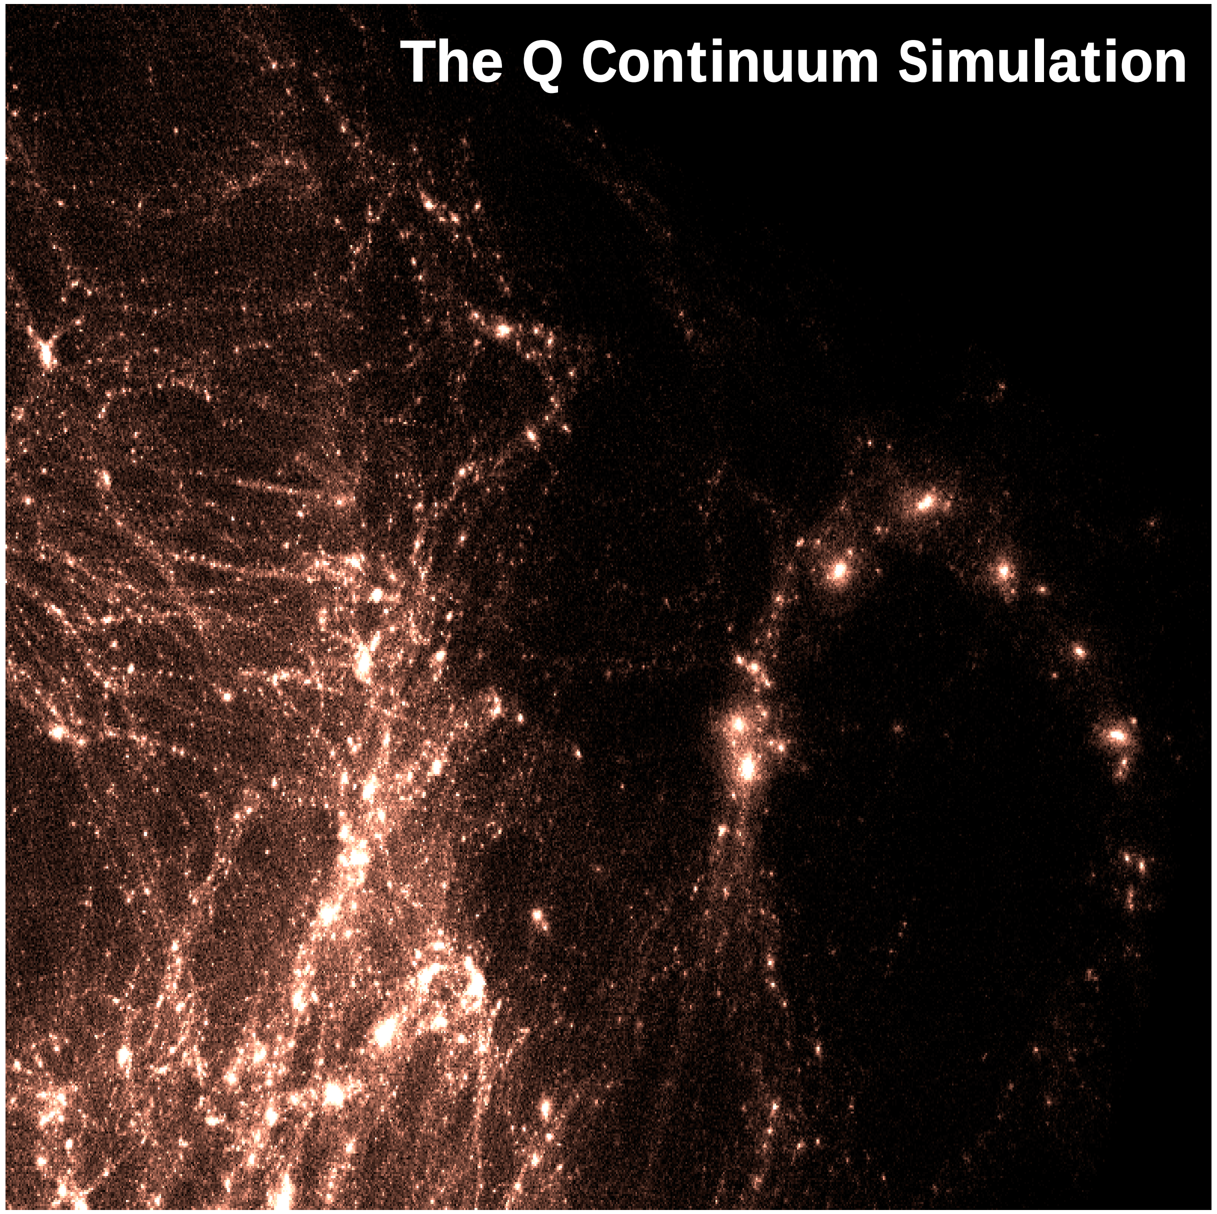


(4225 Mpc)$^3$ volume, 1.07 trillion particles carried out on ~67% of Mira at Argonne, 4PB of data, 216x Millennium simulation

(1300 Mpc)$^3$ volume, 0.55 trillion particles carried out on ~90% of Titan at Oak Ridge, 2PB of data, 64x Bolshoi simulation

**The high resolution Q Continuum Simulation, finished July 13 on ~90% of Titan under INCITE, evolving more than half a trillion particles. Shown is the output from one node (~33 million particles), 1/16384 of the full simulation**

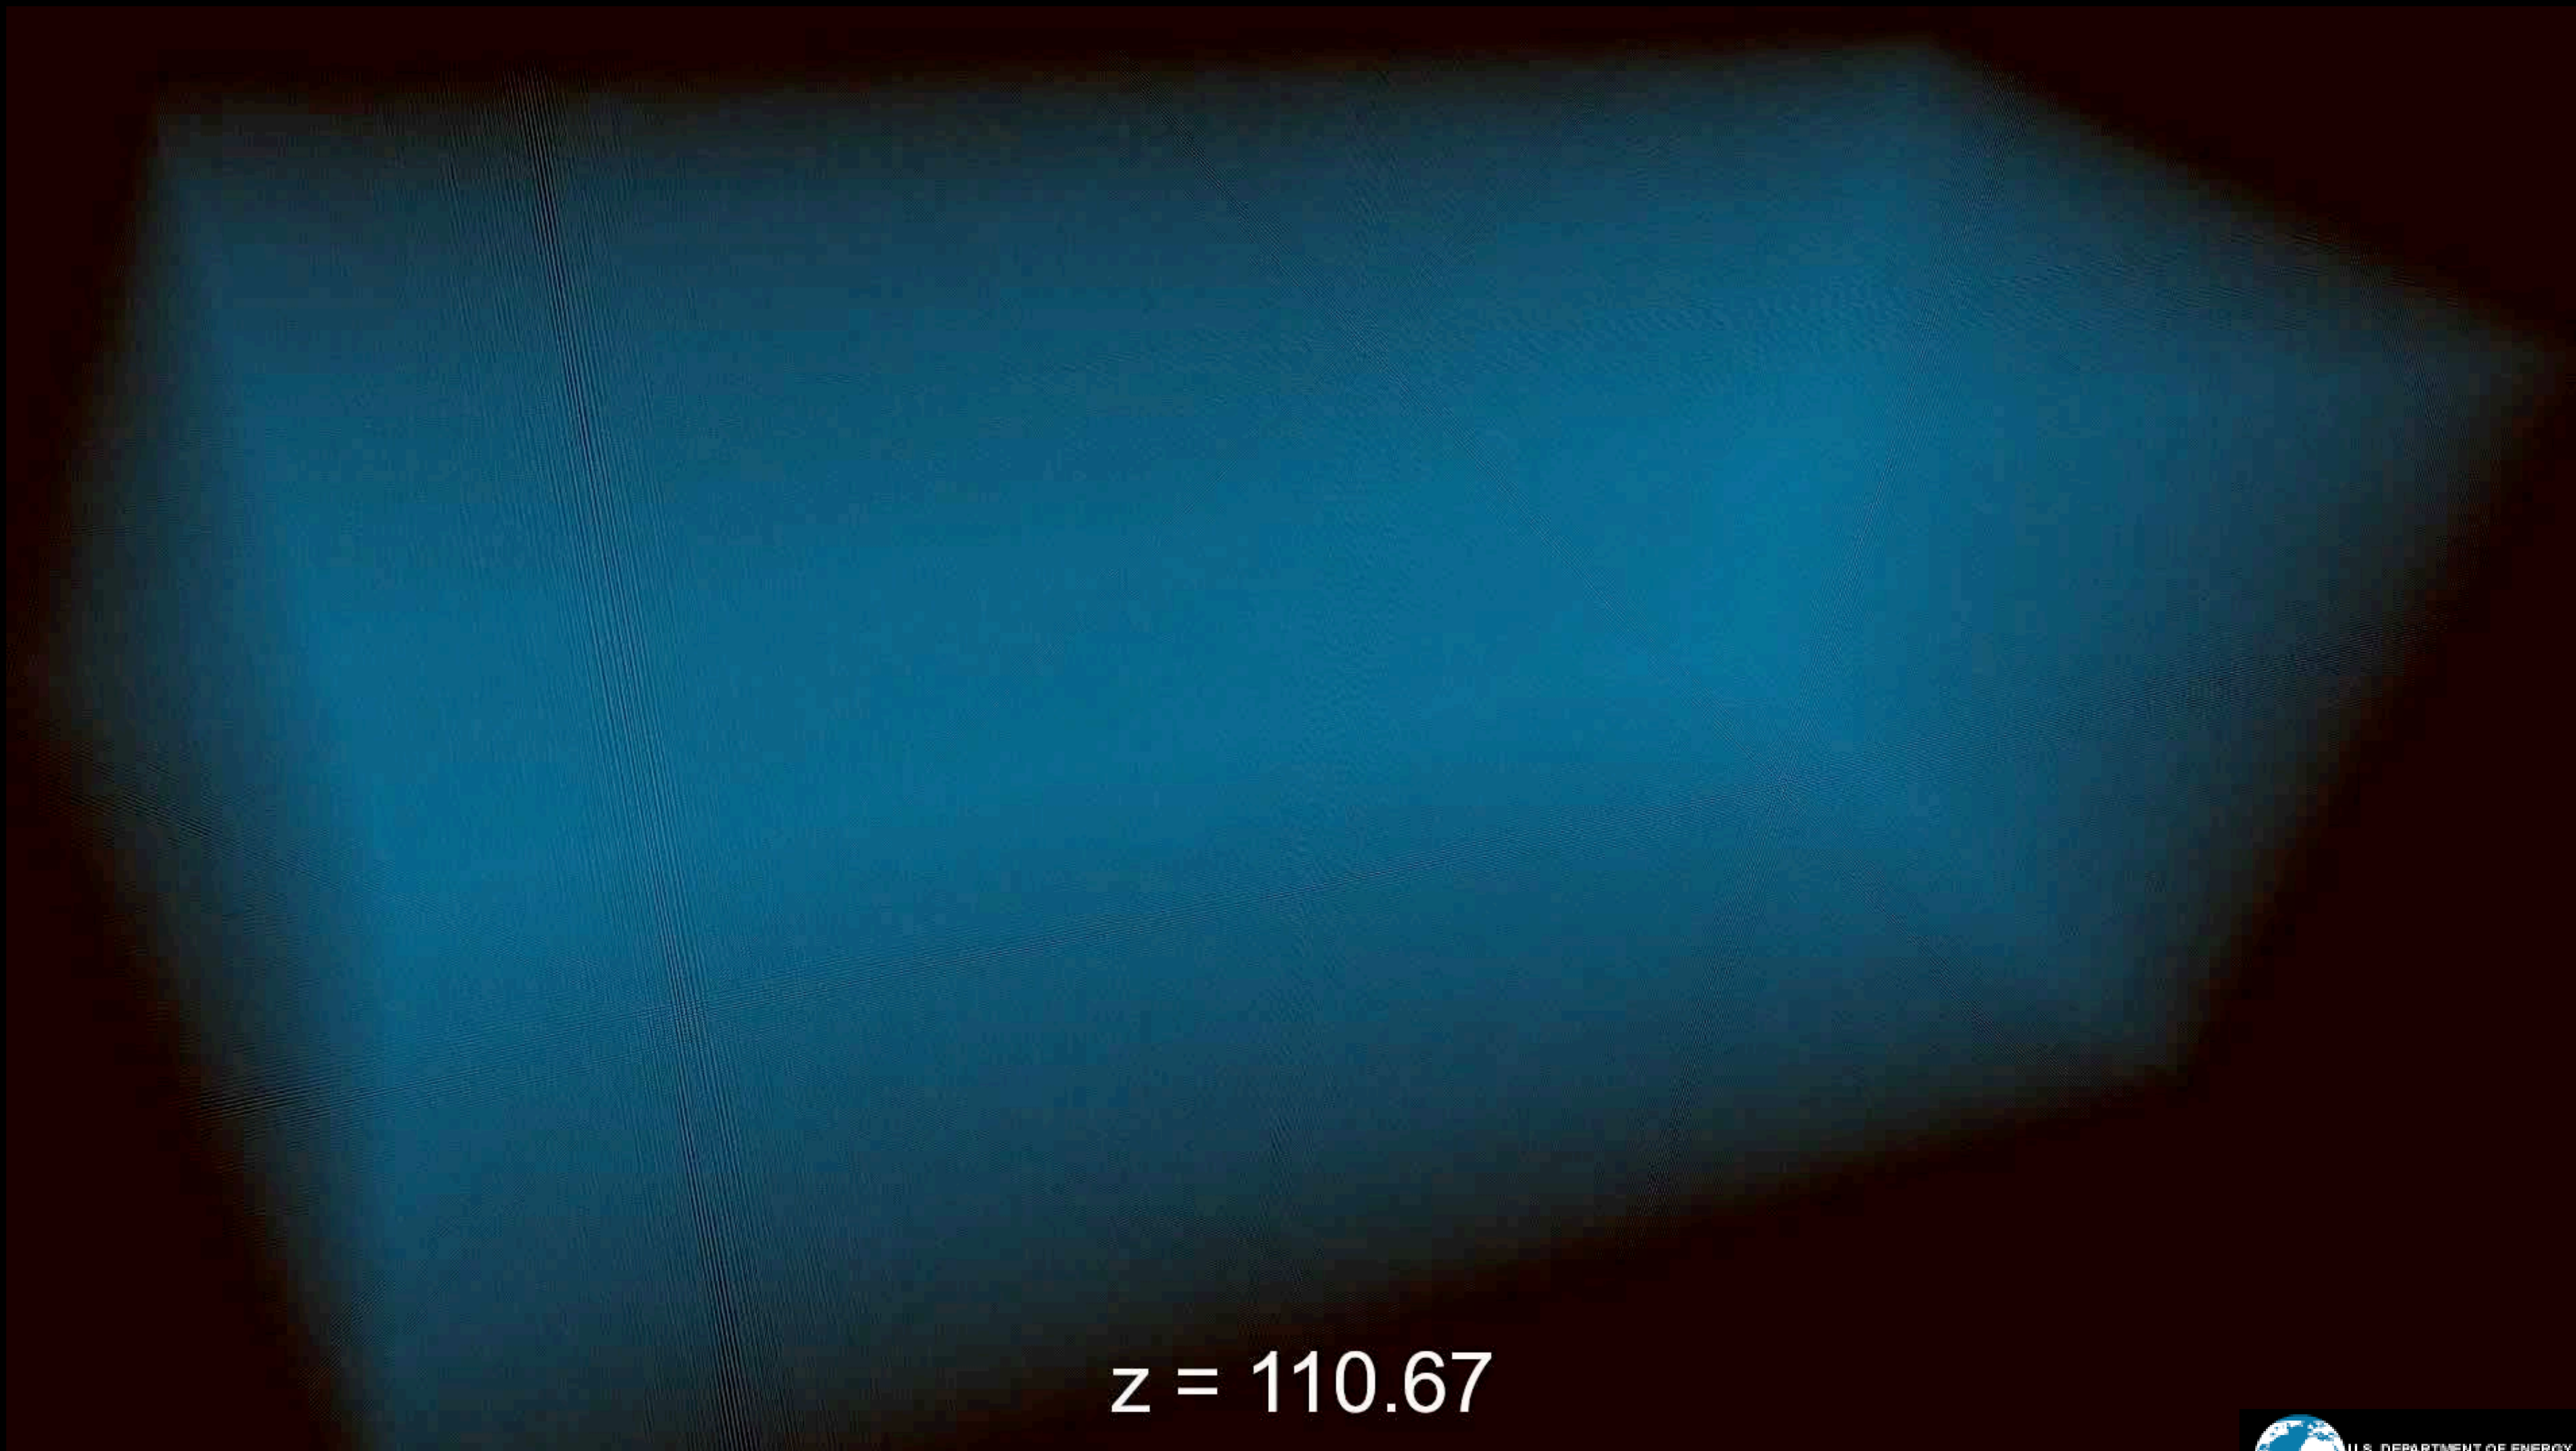


U.S. DEPARTMENT OF ENERGY
INCITE
LEADERSHIP COMPUTING

**Q Continuum: Extradimensional plane of existence**

**Visualization: Silvio Rizzi, Joe Insley et. al., Argonne**

# Merger Tree

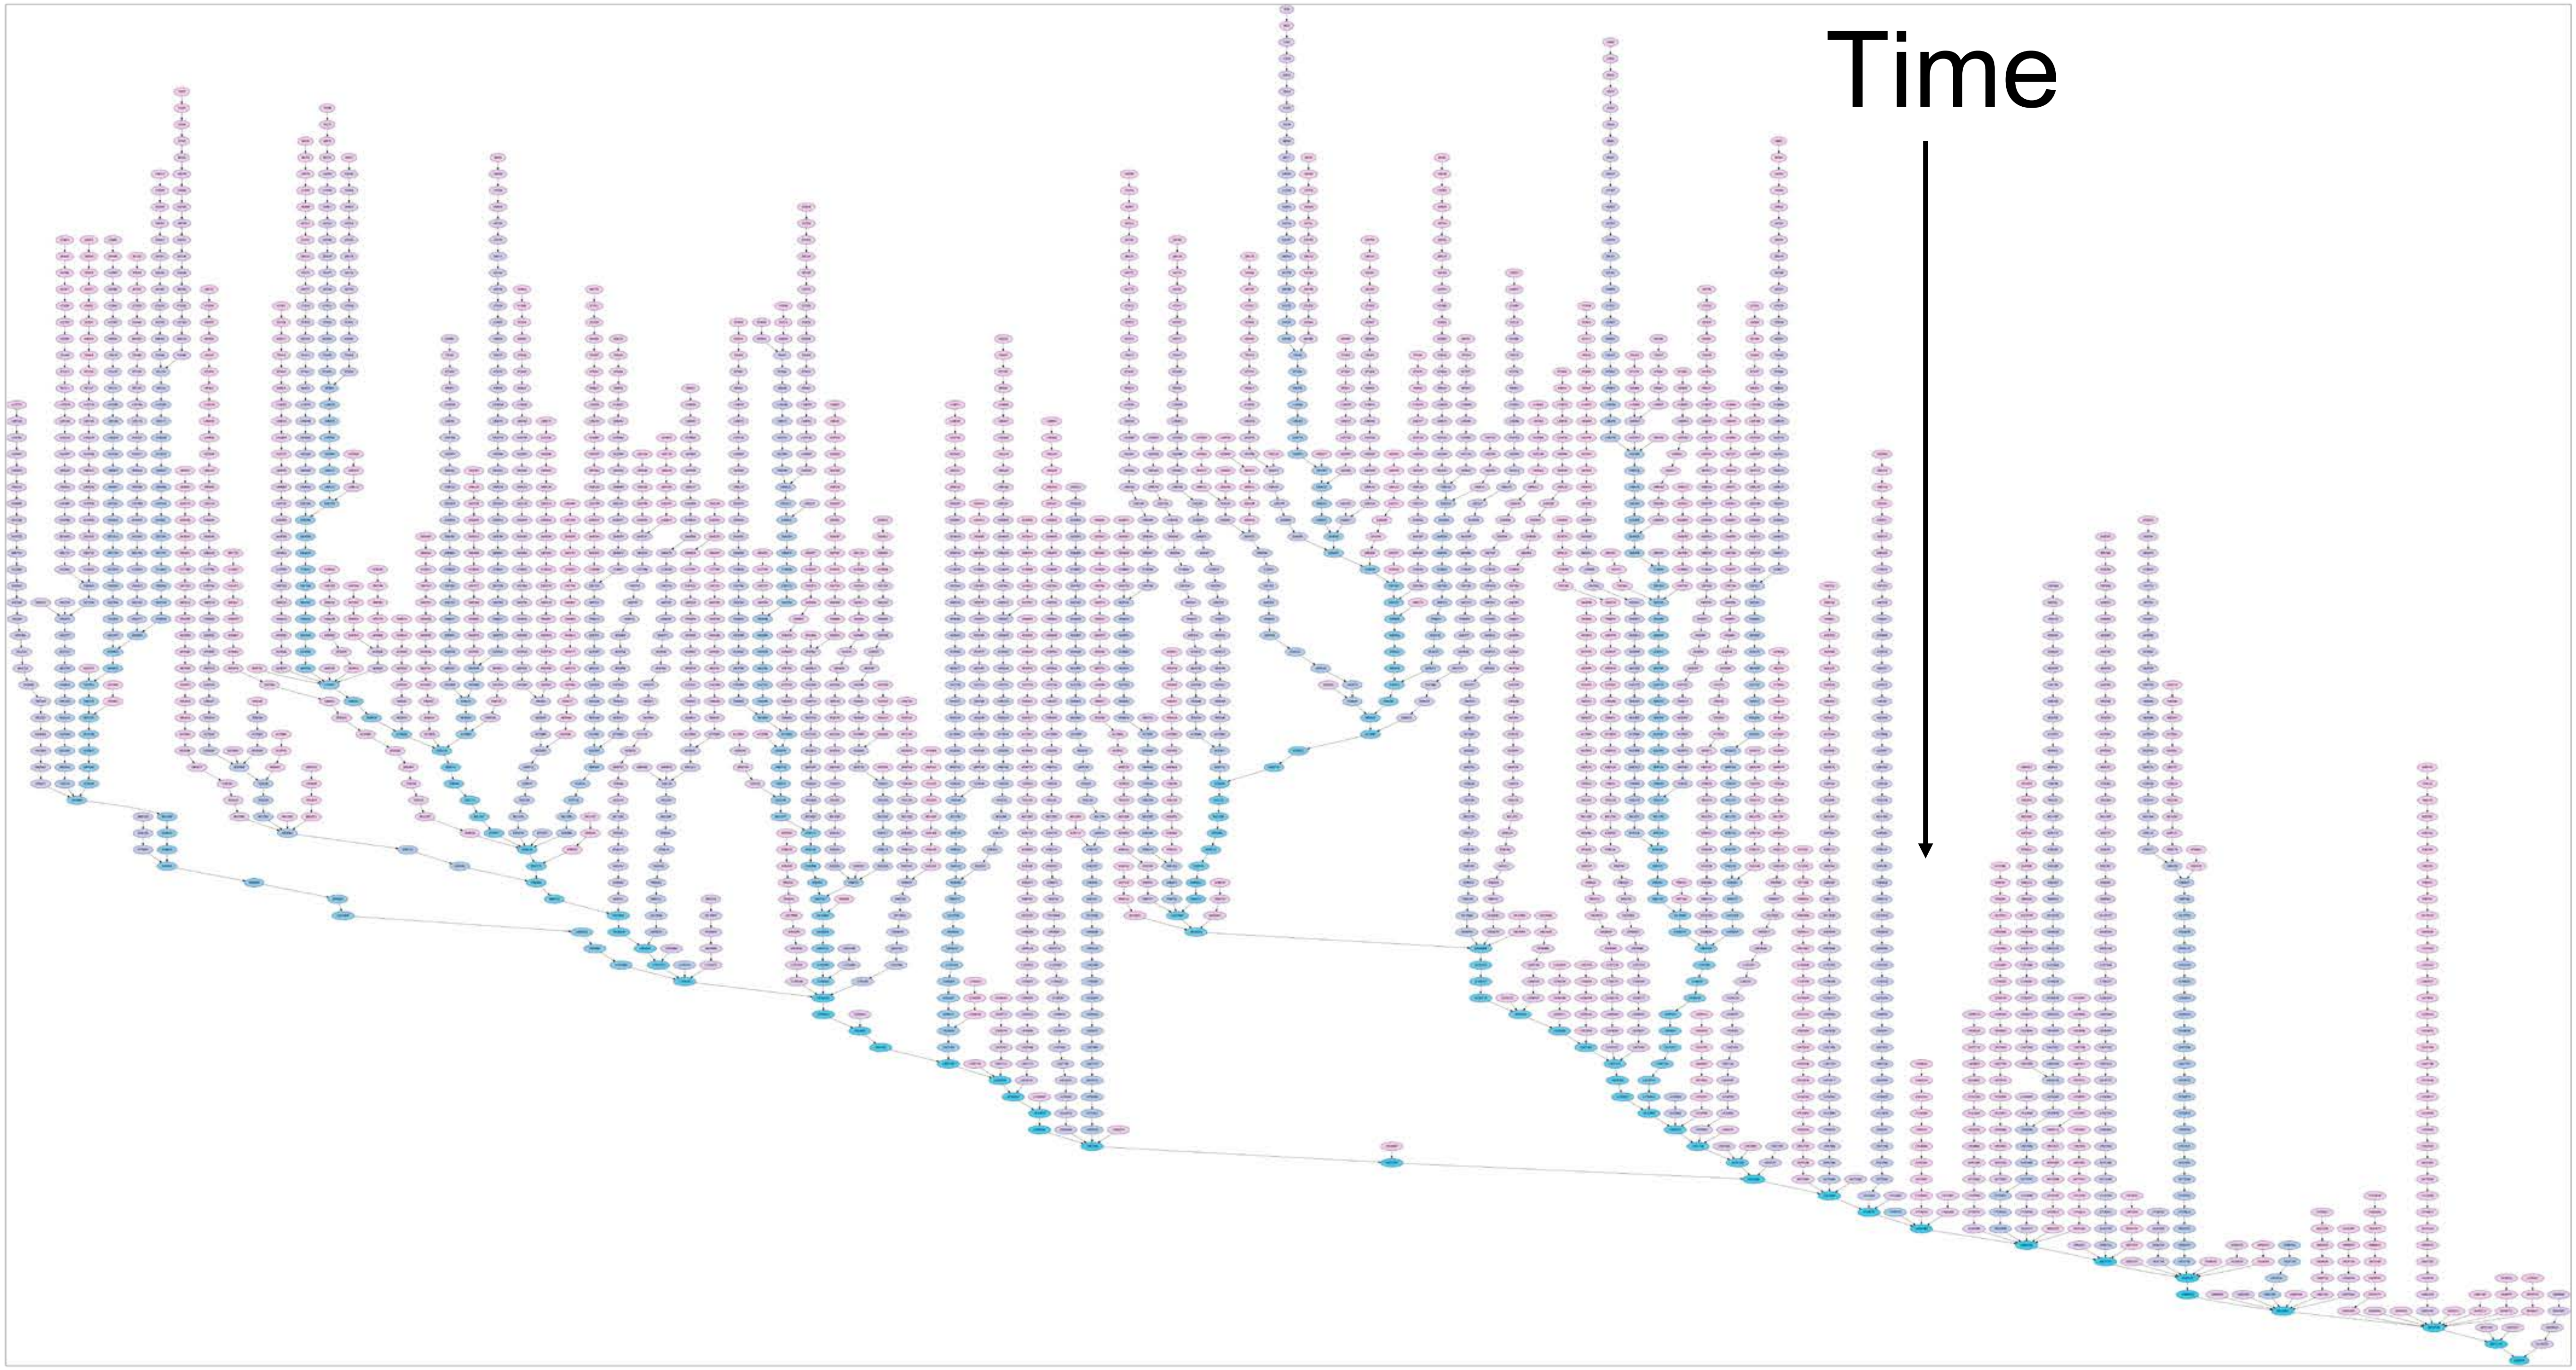

# Cosmology with the Q Continuum

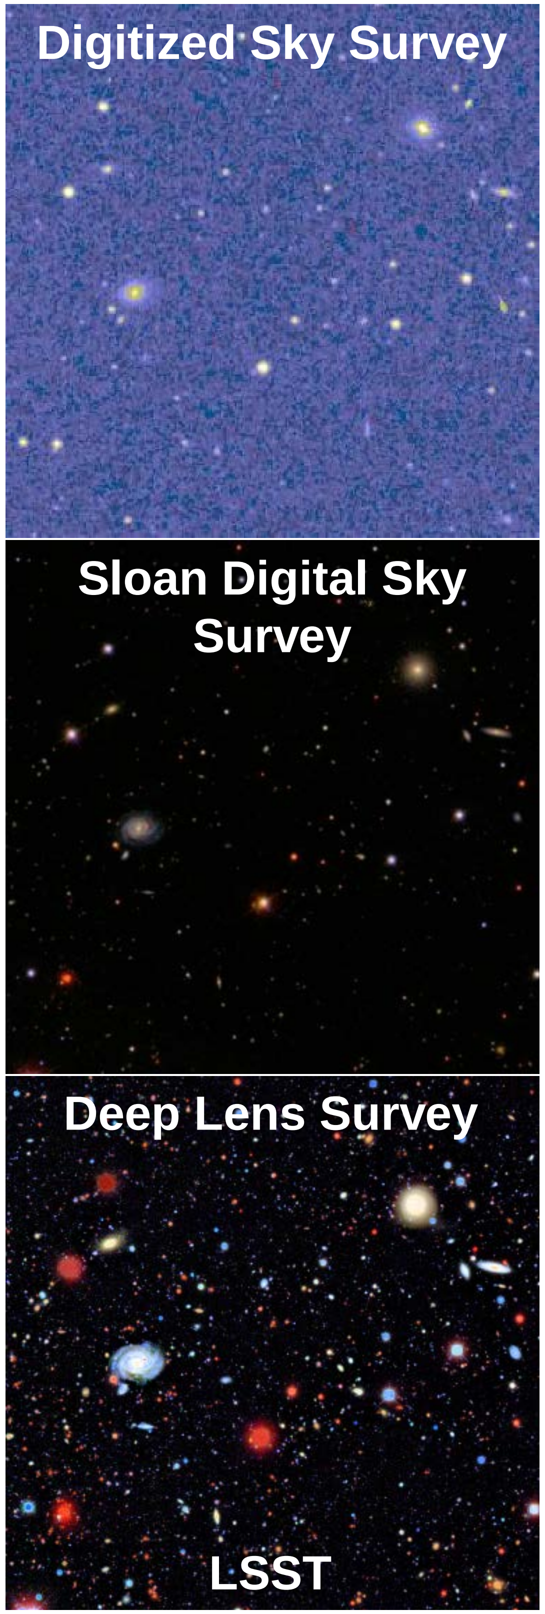


- **Previously for SDSS:** Identify halos and populate them with galaxies, enables us to bright galaxy clustering
- **Now:** Mass resolution enables us to identify ***halos within halos, so-called sub-halos*** (track halo particles over time even after they have become part of another halo)
- **Use semi-analytic code (Galacticus, developed by Andrew Benson) to model galaxy population within subhalos**
- **A few Q Continuum project examples:**
  - Replace Millennium simulation in end-to-end simulation pipeline for the full LSST survey by the end of this year
  - Galaxy-galaxy lensing for the Dark Energy Survey (DES)
  - Cluster lensing for DES
  - Strong lensing for HST

**2-d projected halo density**

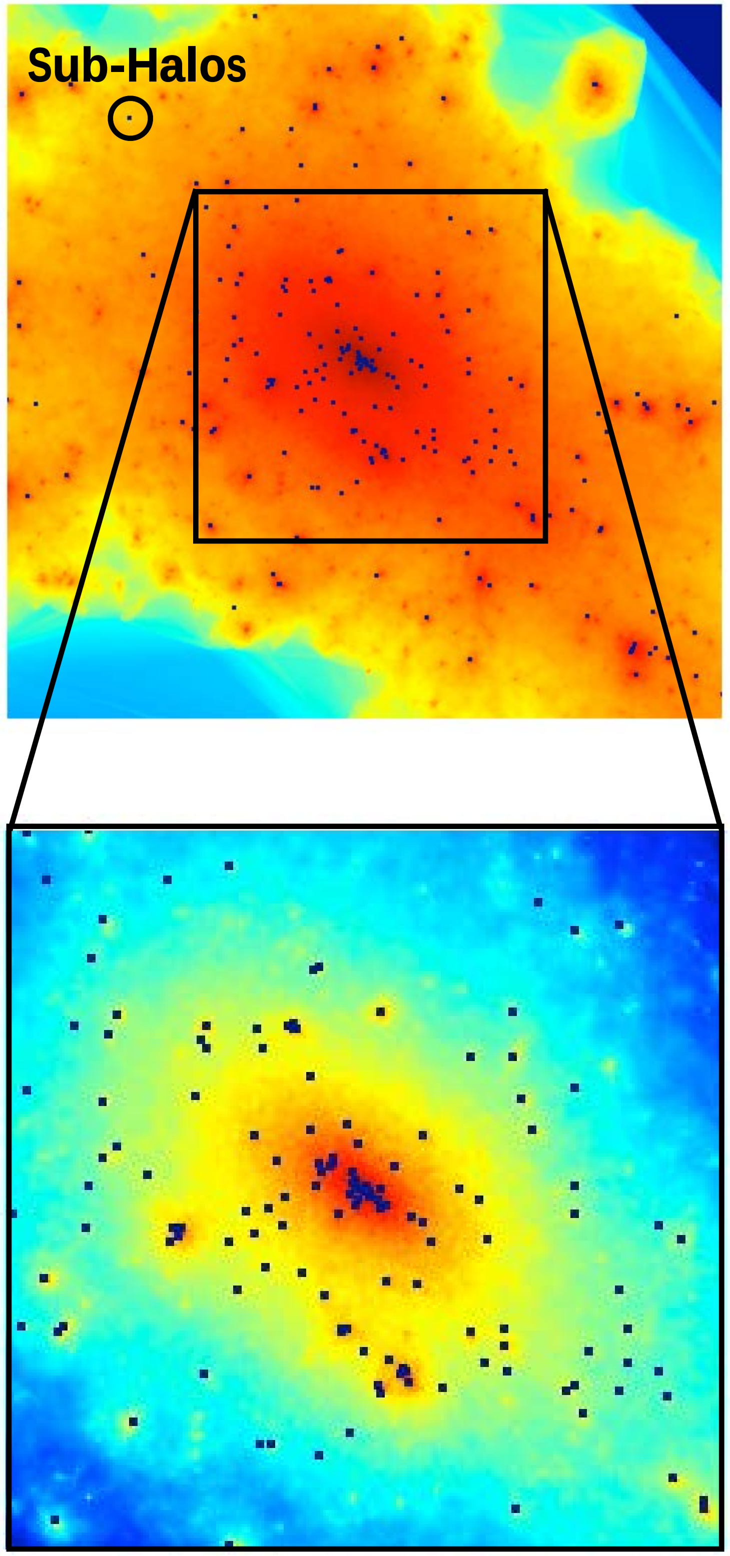

# Cosmology with HACC: Exquisite Statistics

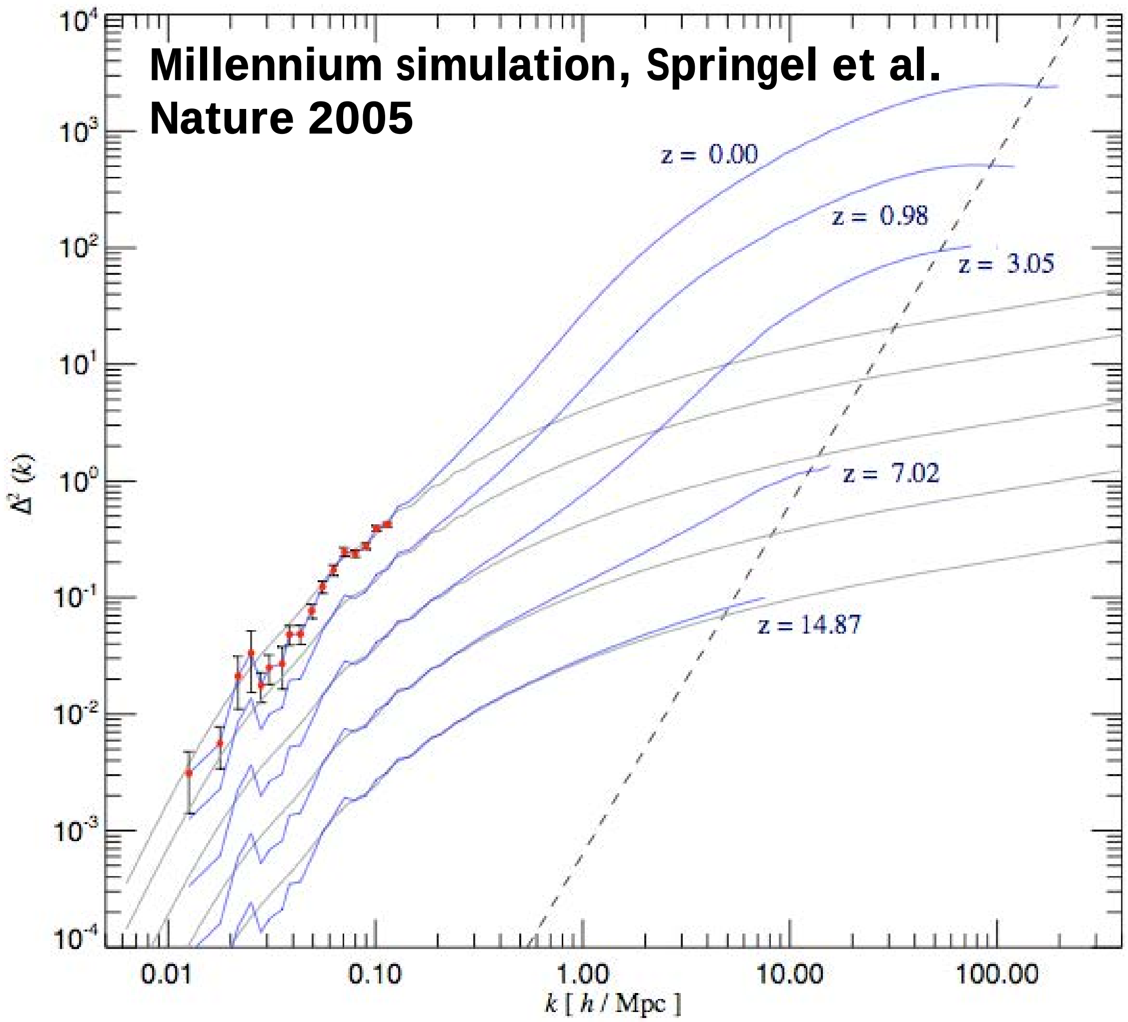


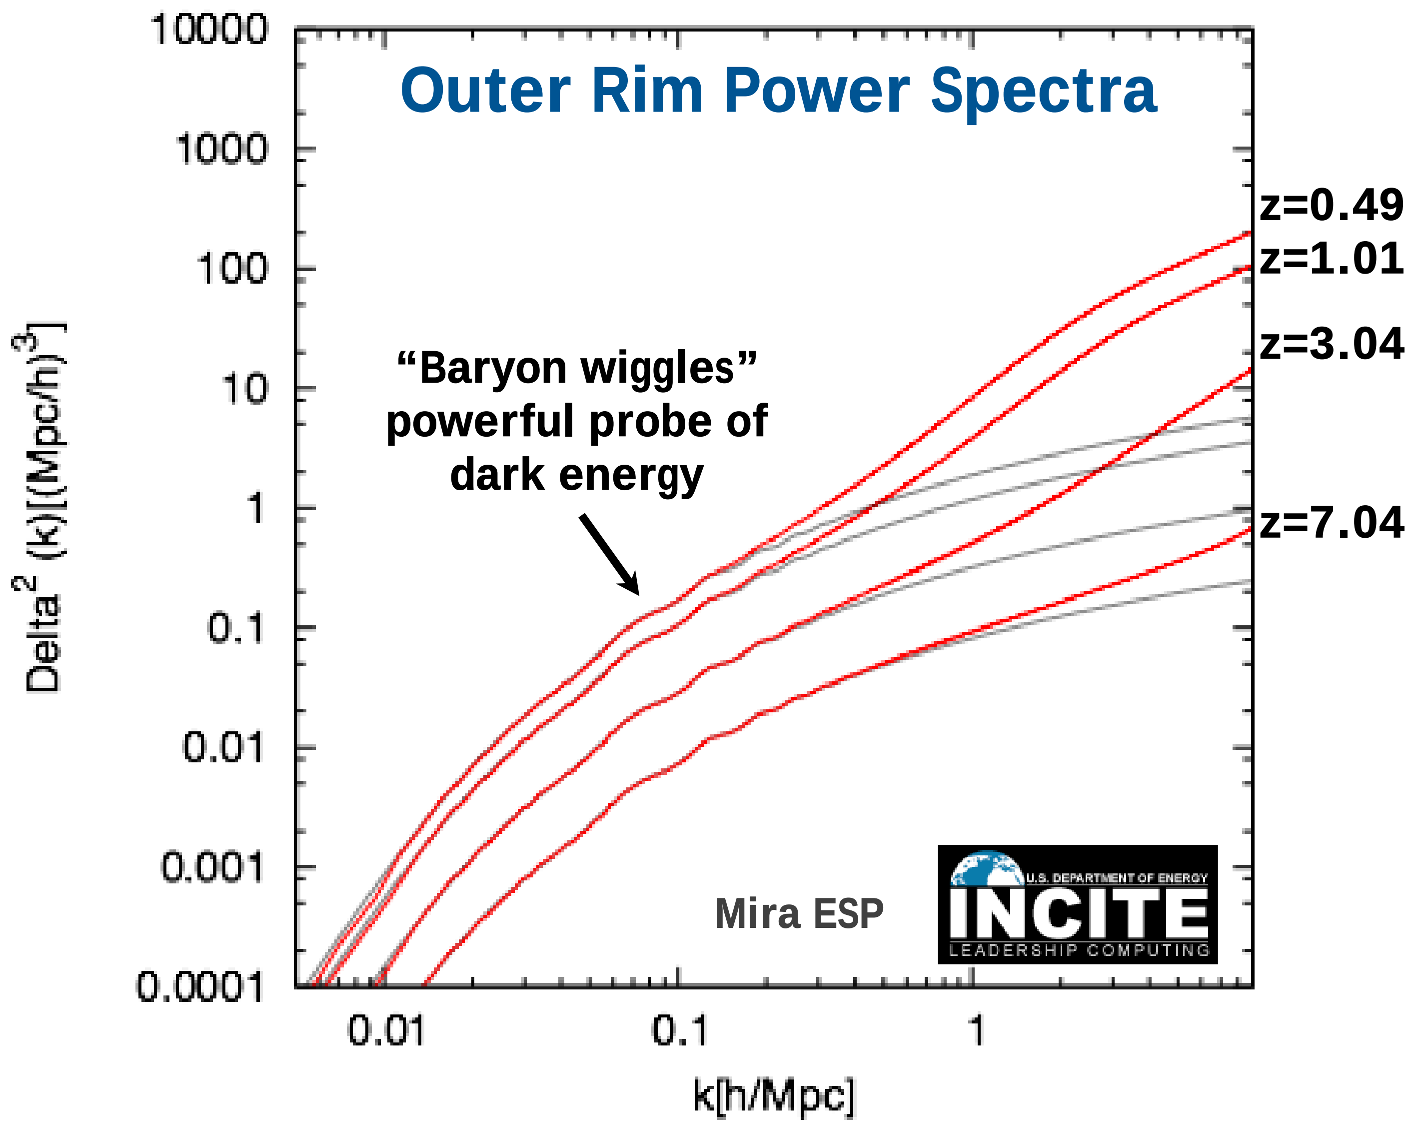


- Mass resolution of Millennium simulation and Outer Rim run very similar (~$10^9 M_\odot$ particle mass), but volume different by a factor of 216 (Outer Rim volume = Millennium XXL, but with 7 times higher mass resolution)
- Exceptional statistics at high resolution enable many science projects

# Cosmic Calibration: Solving the Inverse Problem

- **Challenge:** To extract cosmological constraints from observations in non-linear regime, need to run Markov Chain Monte Carlo code; input: 10,000 - 100,000 different models
- **Current strategy:** Fitting functions for e.g. P(k), accurate at 10% level, not good enough!
- **Brute force:** Simulations, ~30 years on 2000 processor cluster…
- **Only alternative:** emulators

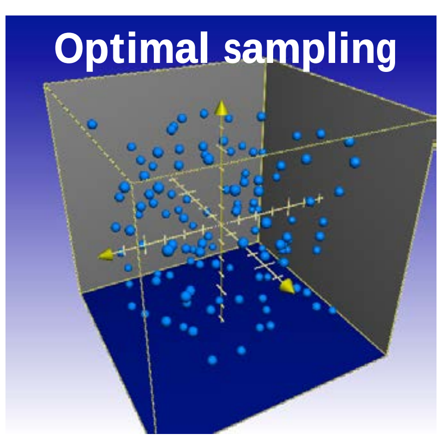



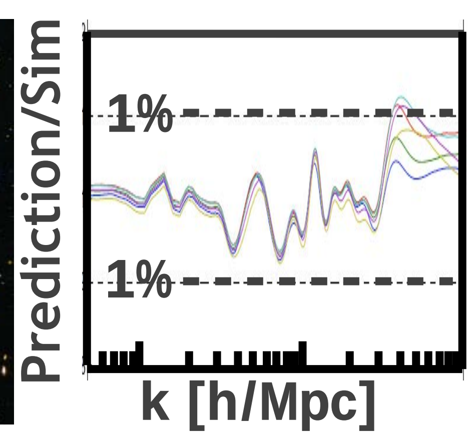


**Heitmann et al. 2006, Habib et al. 2007**

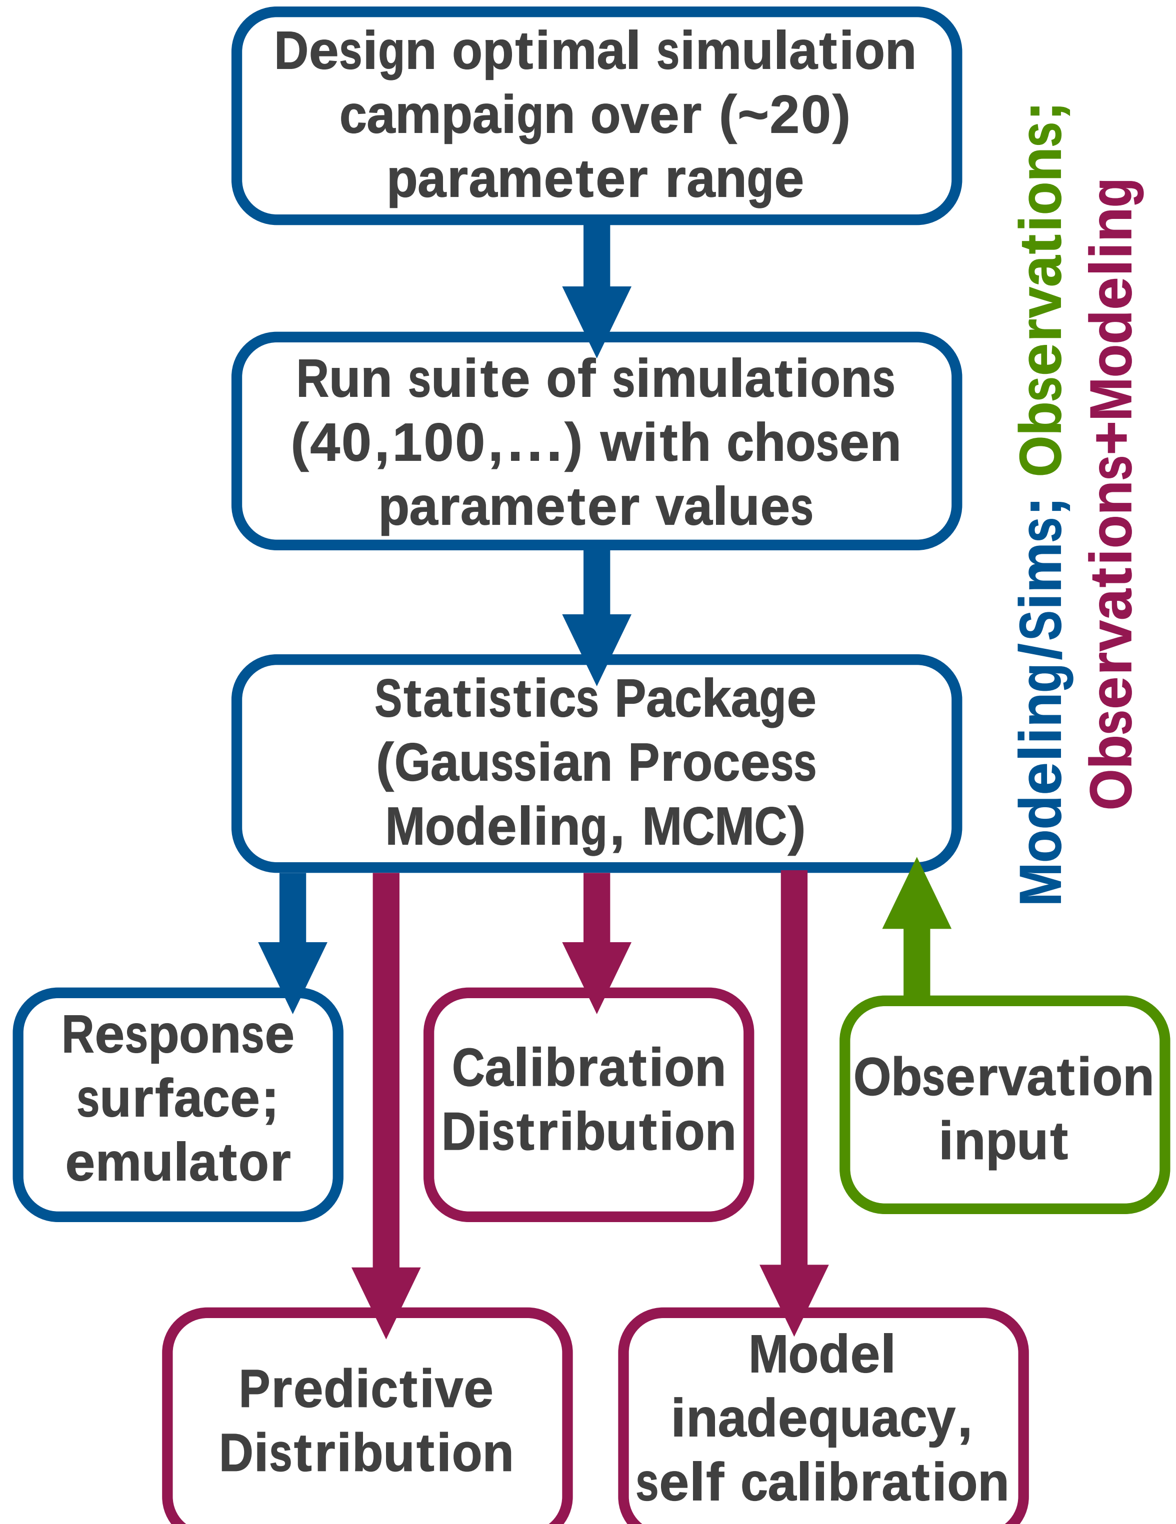

# The Coyote Universe: Emulator Science

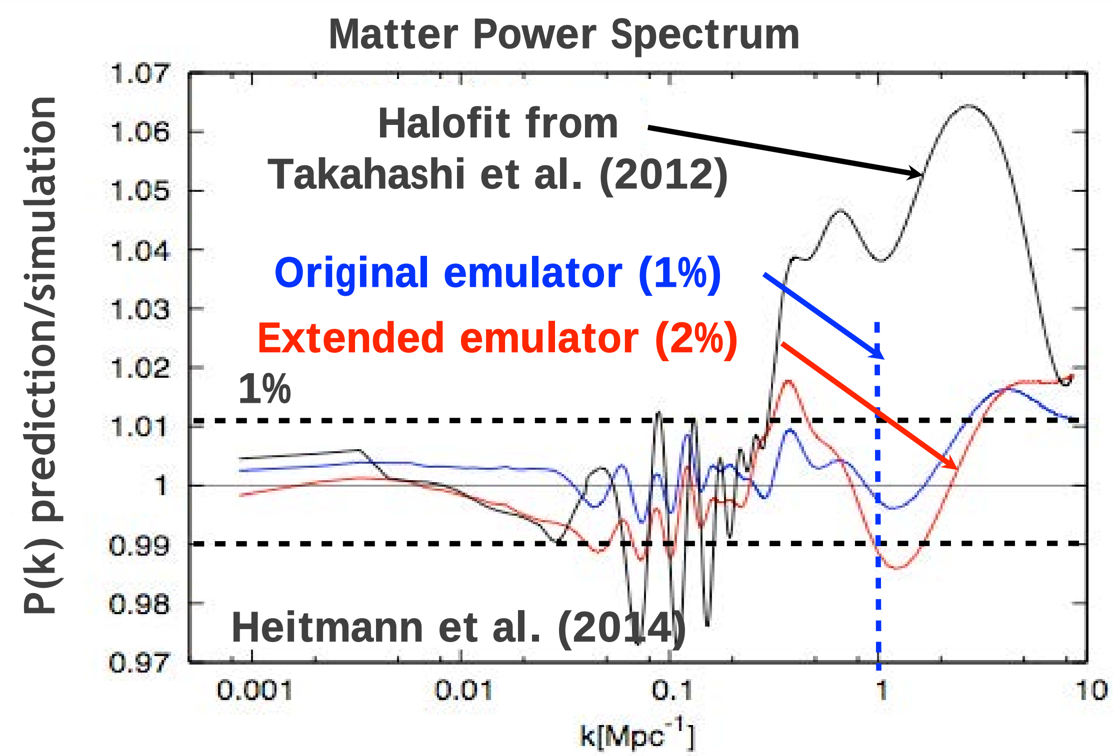


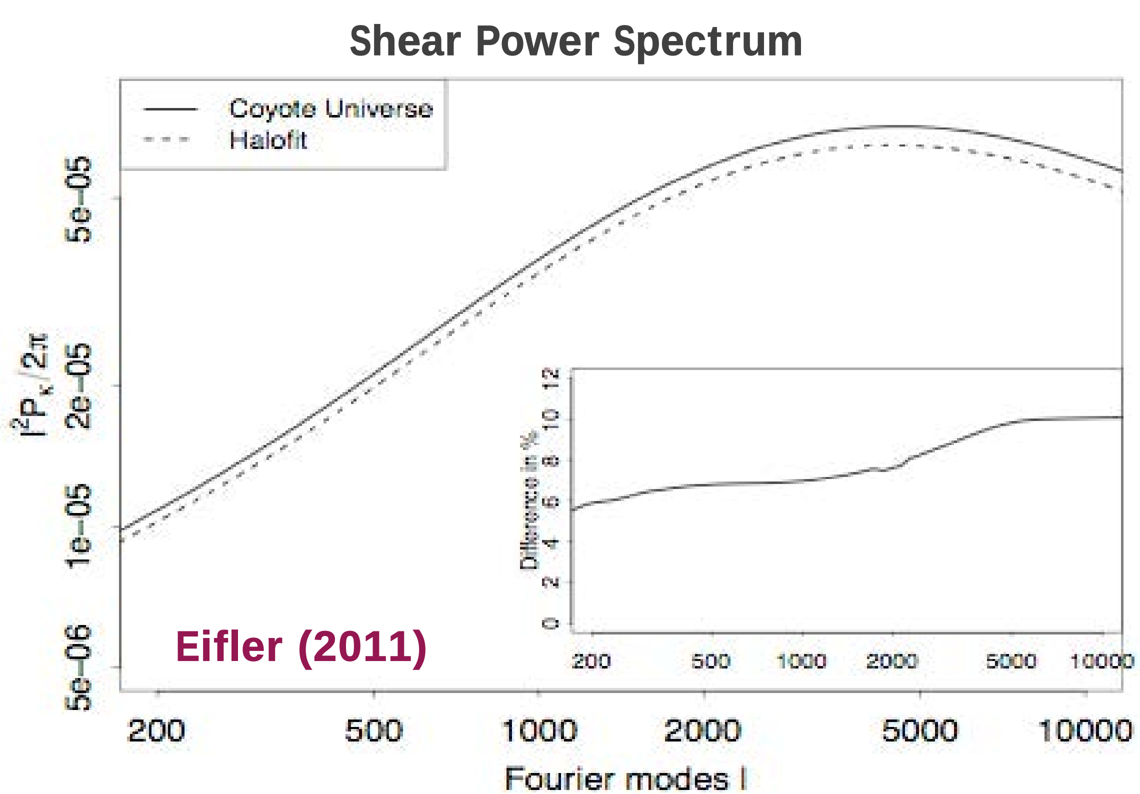


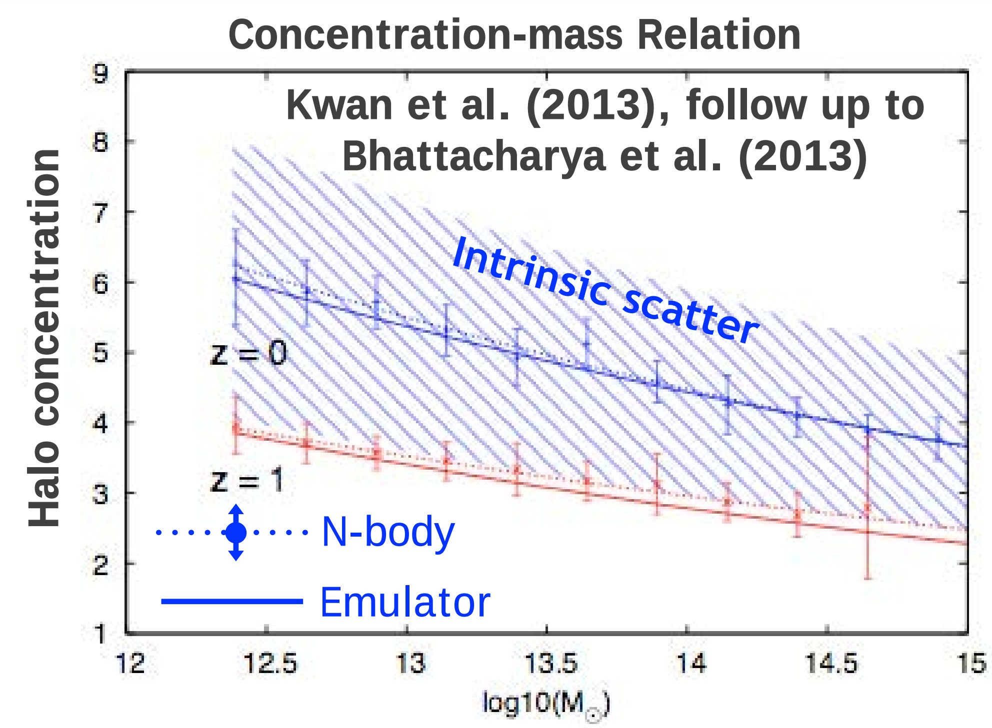


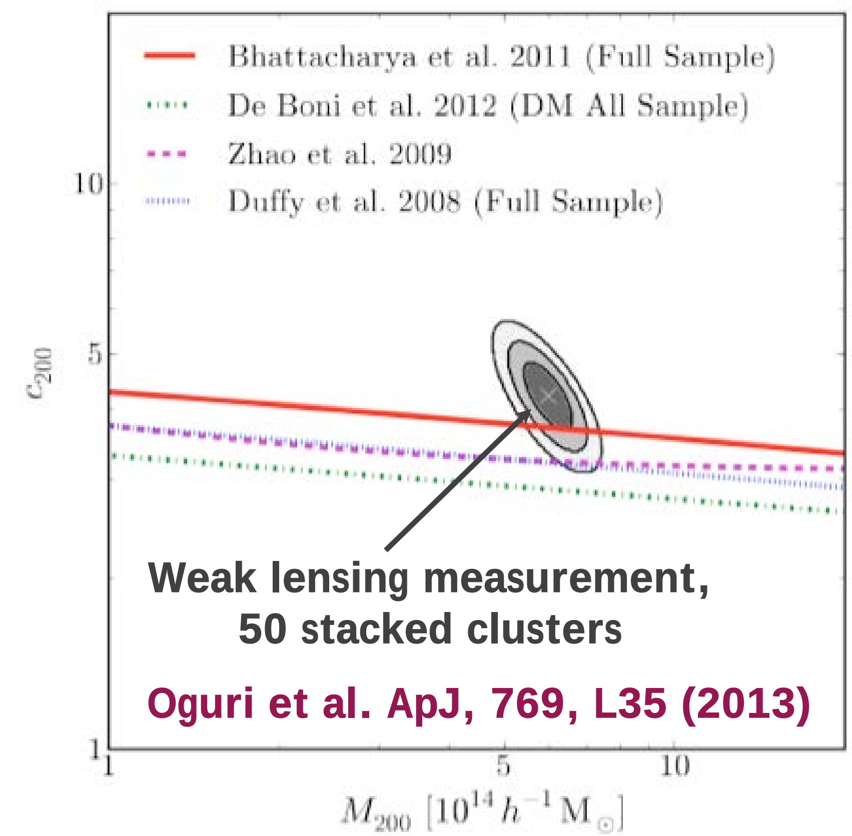

# Nyx

- 3-D Cartesian grid, finite volume representation
- Evolve dark matter as collisionless Lagrangian fluid
- Evolve baryons as ideal gas using unsplit, Godunov-type methodology
- Adaptive mesh refinement (AMR) to extend dynamic range
- Uses BoxLib software framework developed at LBL
- Code paper: ApJ, 765, 39 (2013)

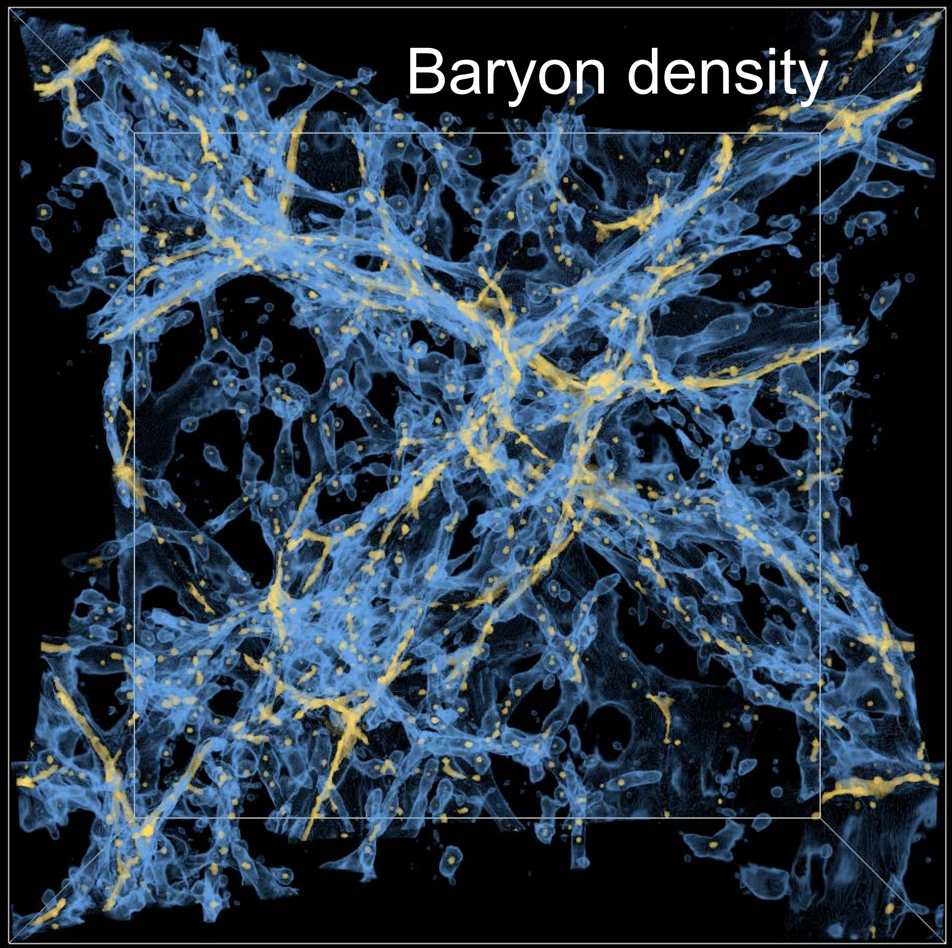


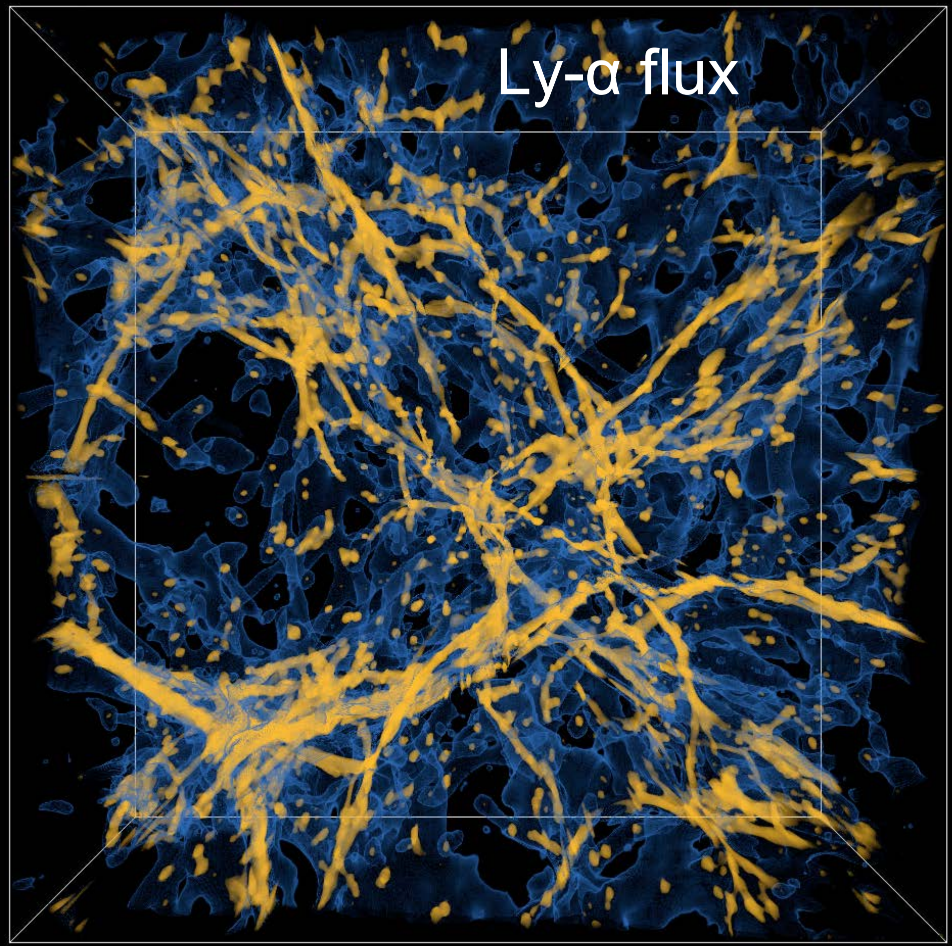

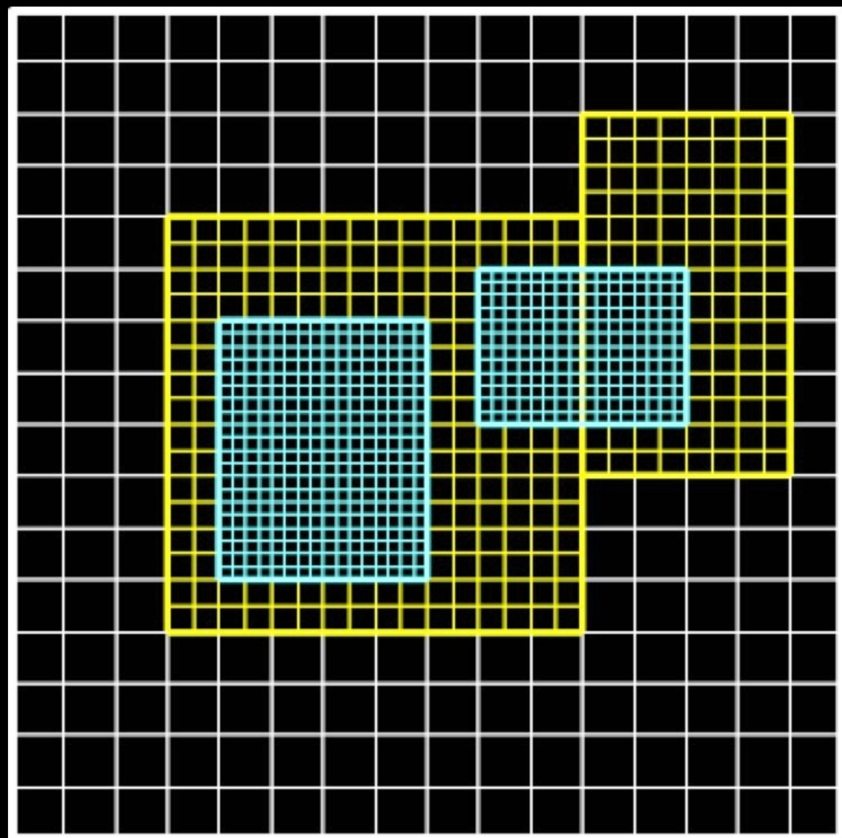

- AMR: patch-based refinement, with jump up to a factor of 4.

- Hydro: unsplit finite volume scheme better characterizes fluid flow.

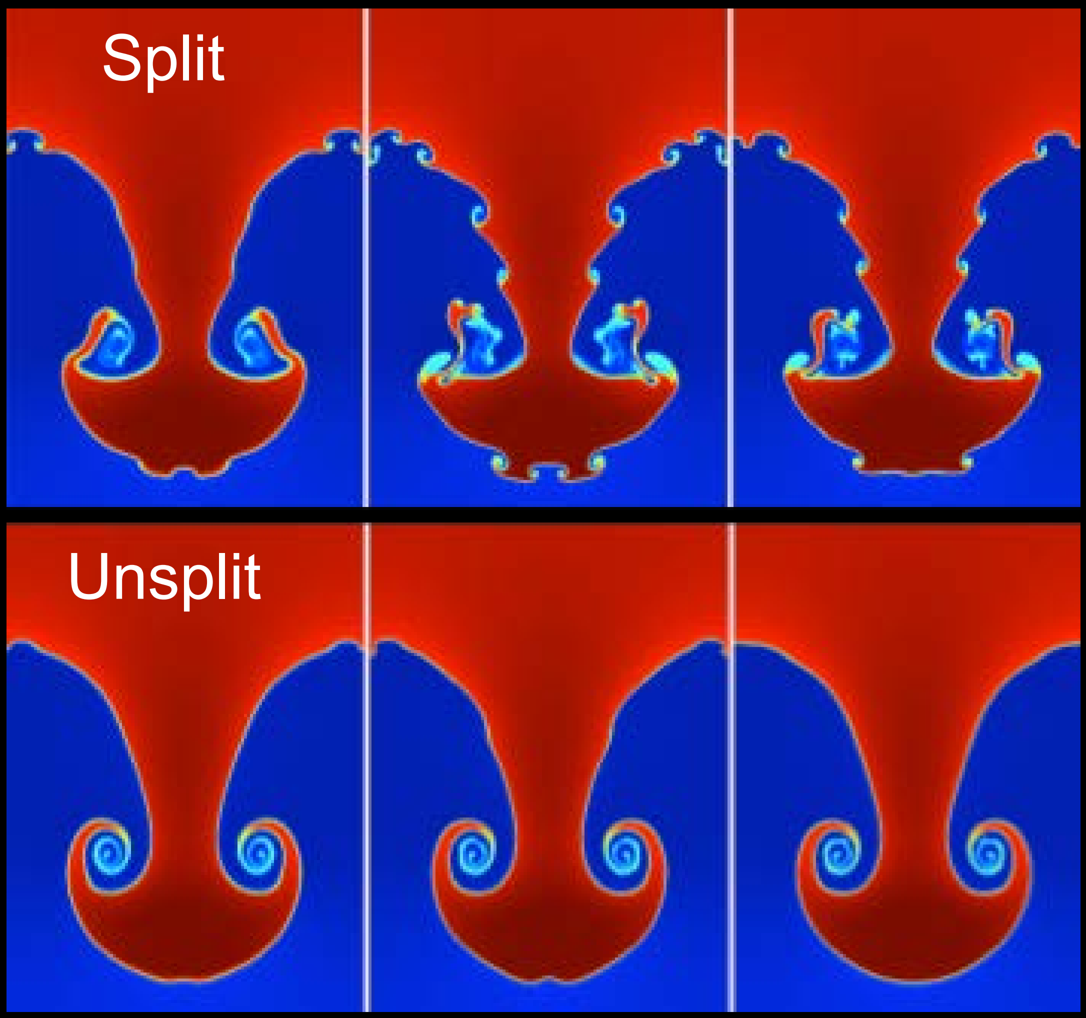

# Excellent scaling

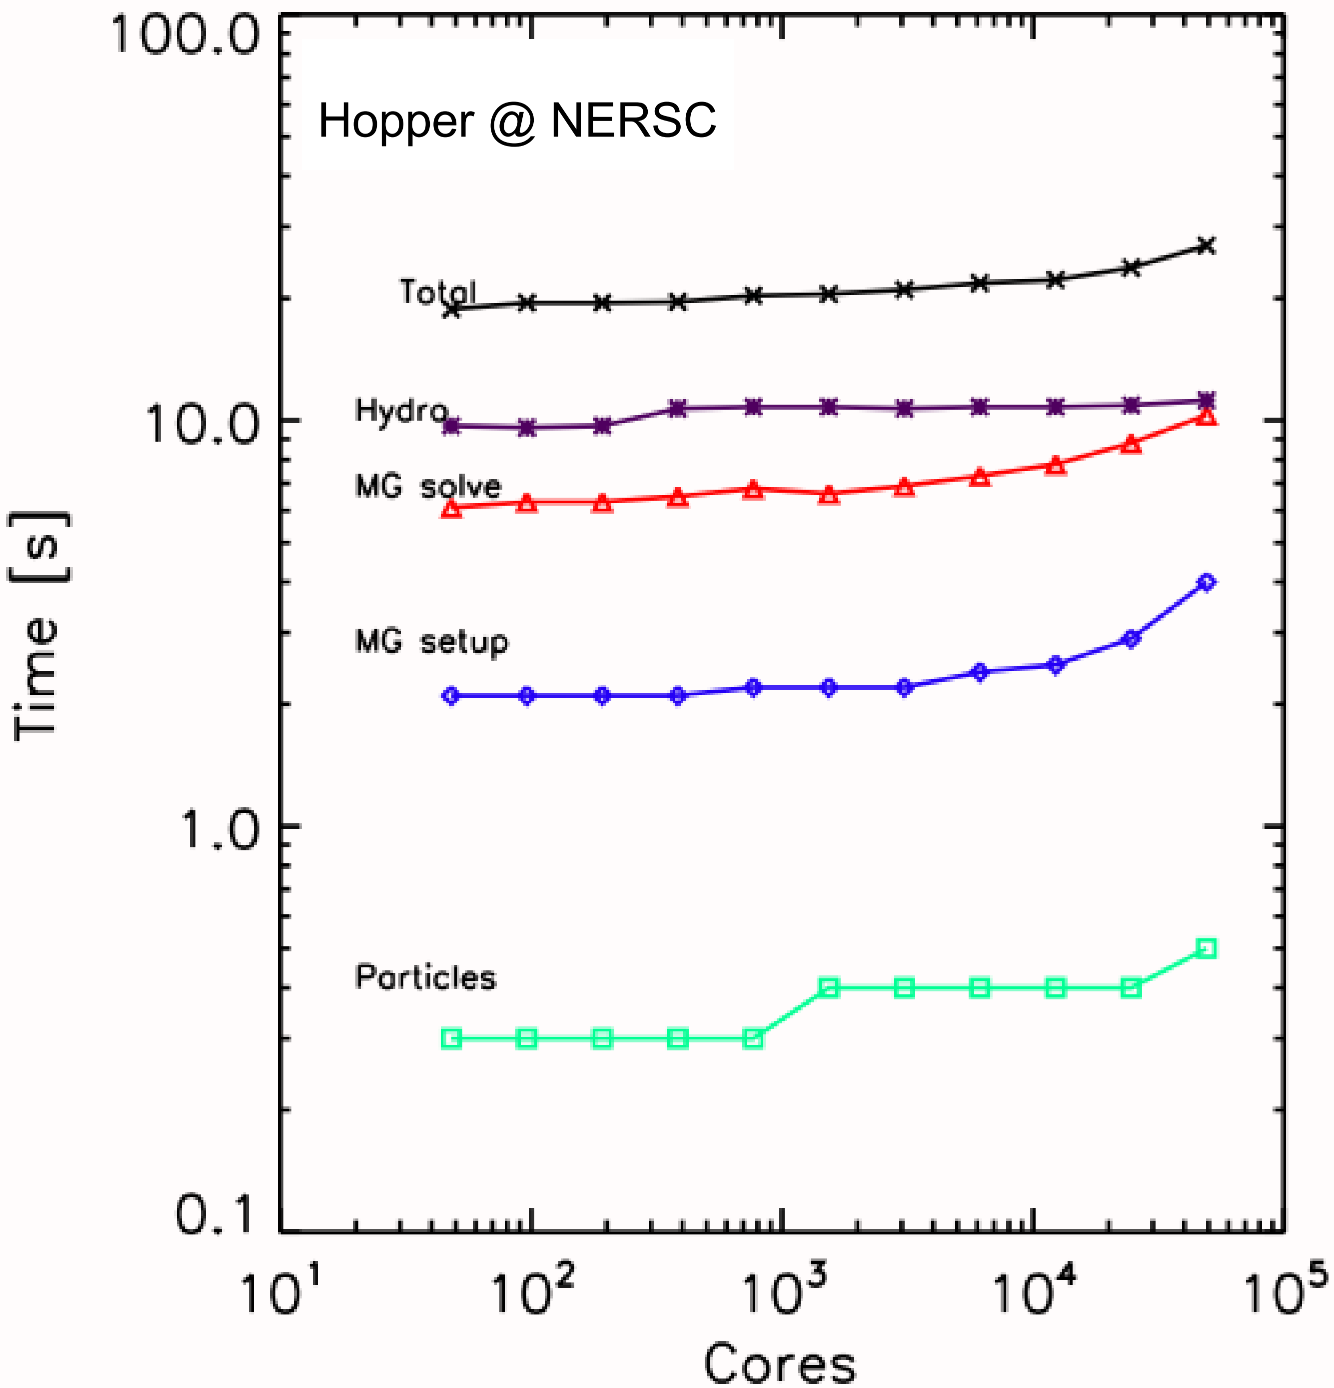


- Currently we are using NERSC resources under ALCC allocation.
- Mostly running $2048^3$ and $4096^3$ runs.
- Hopper/Edison: standard cluster architecture, 24 cores on a node, 32/64GB per node, ~5,000 nodes.
- Analysis pipeline on par with simulations.

$4096^3$ hydro simulation

Blue: F~0; Red: F~1

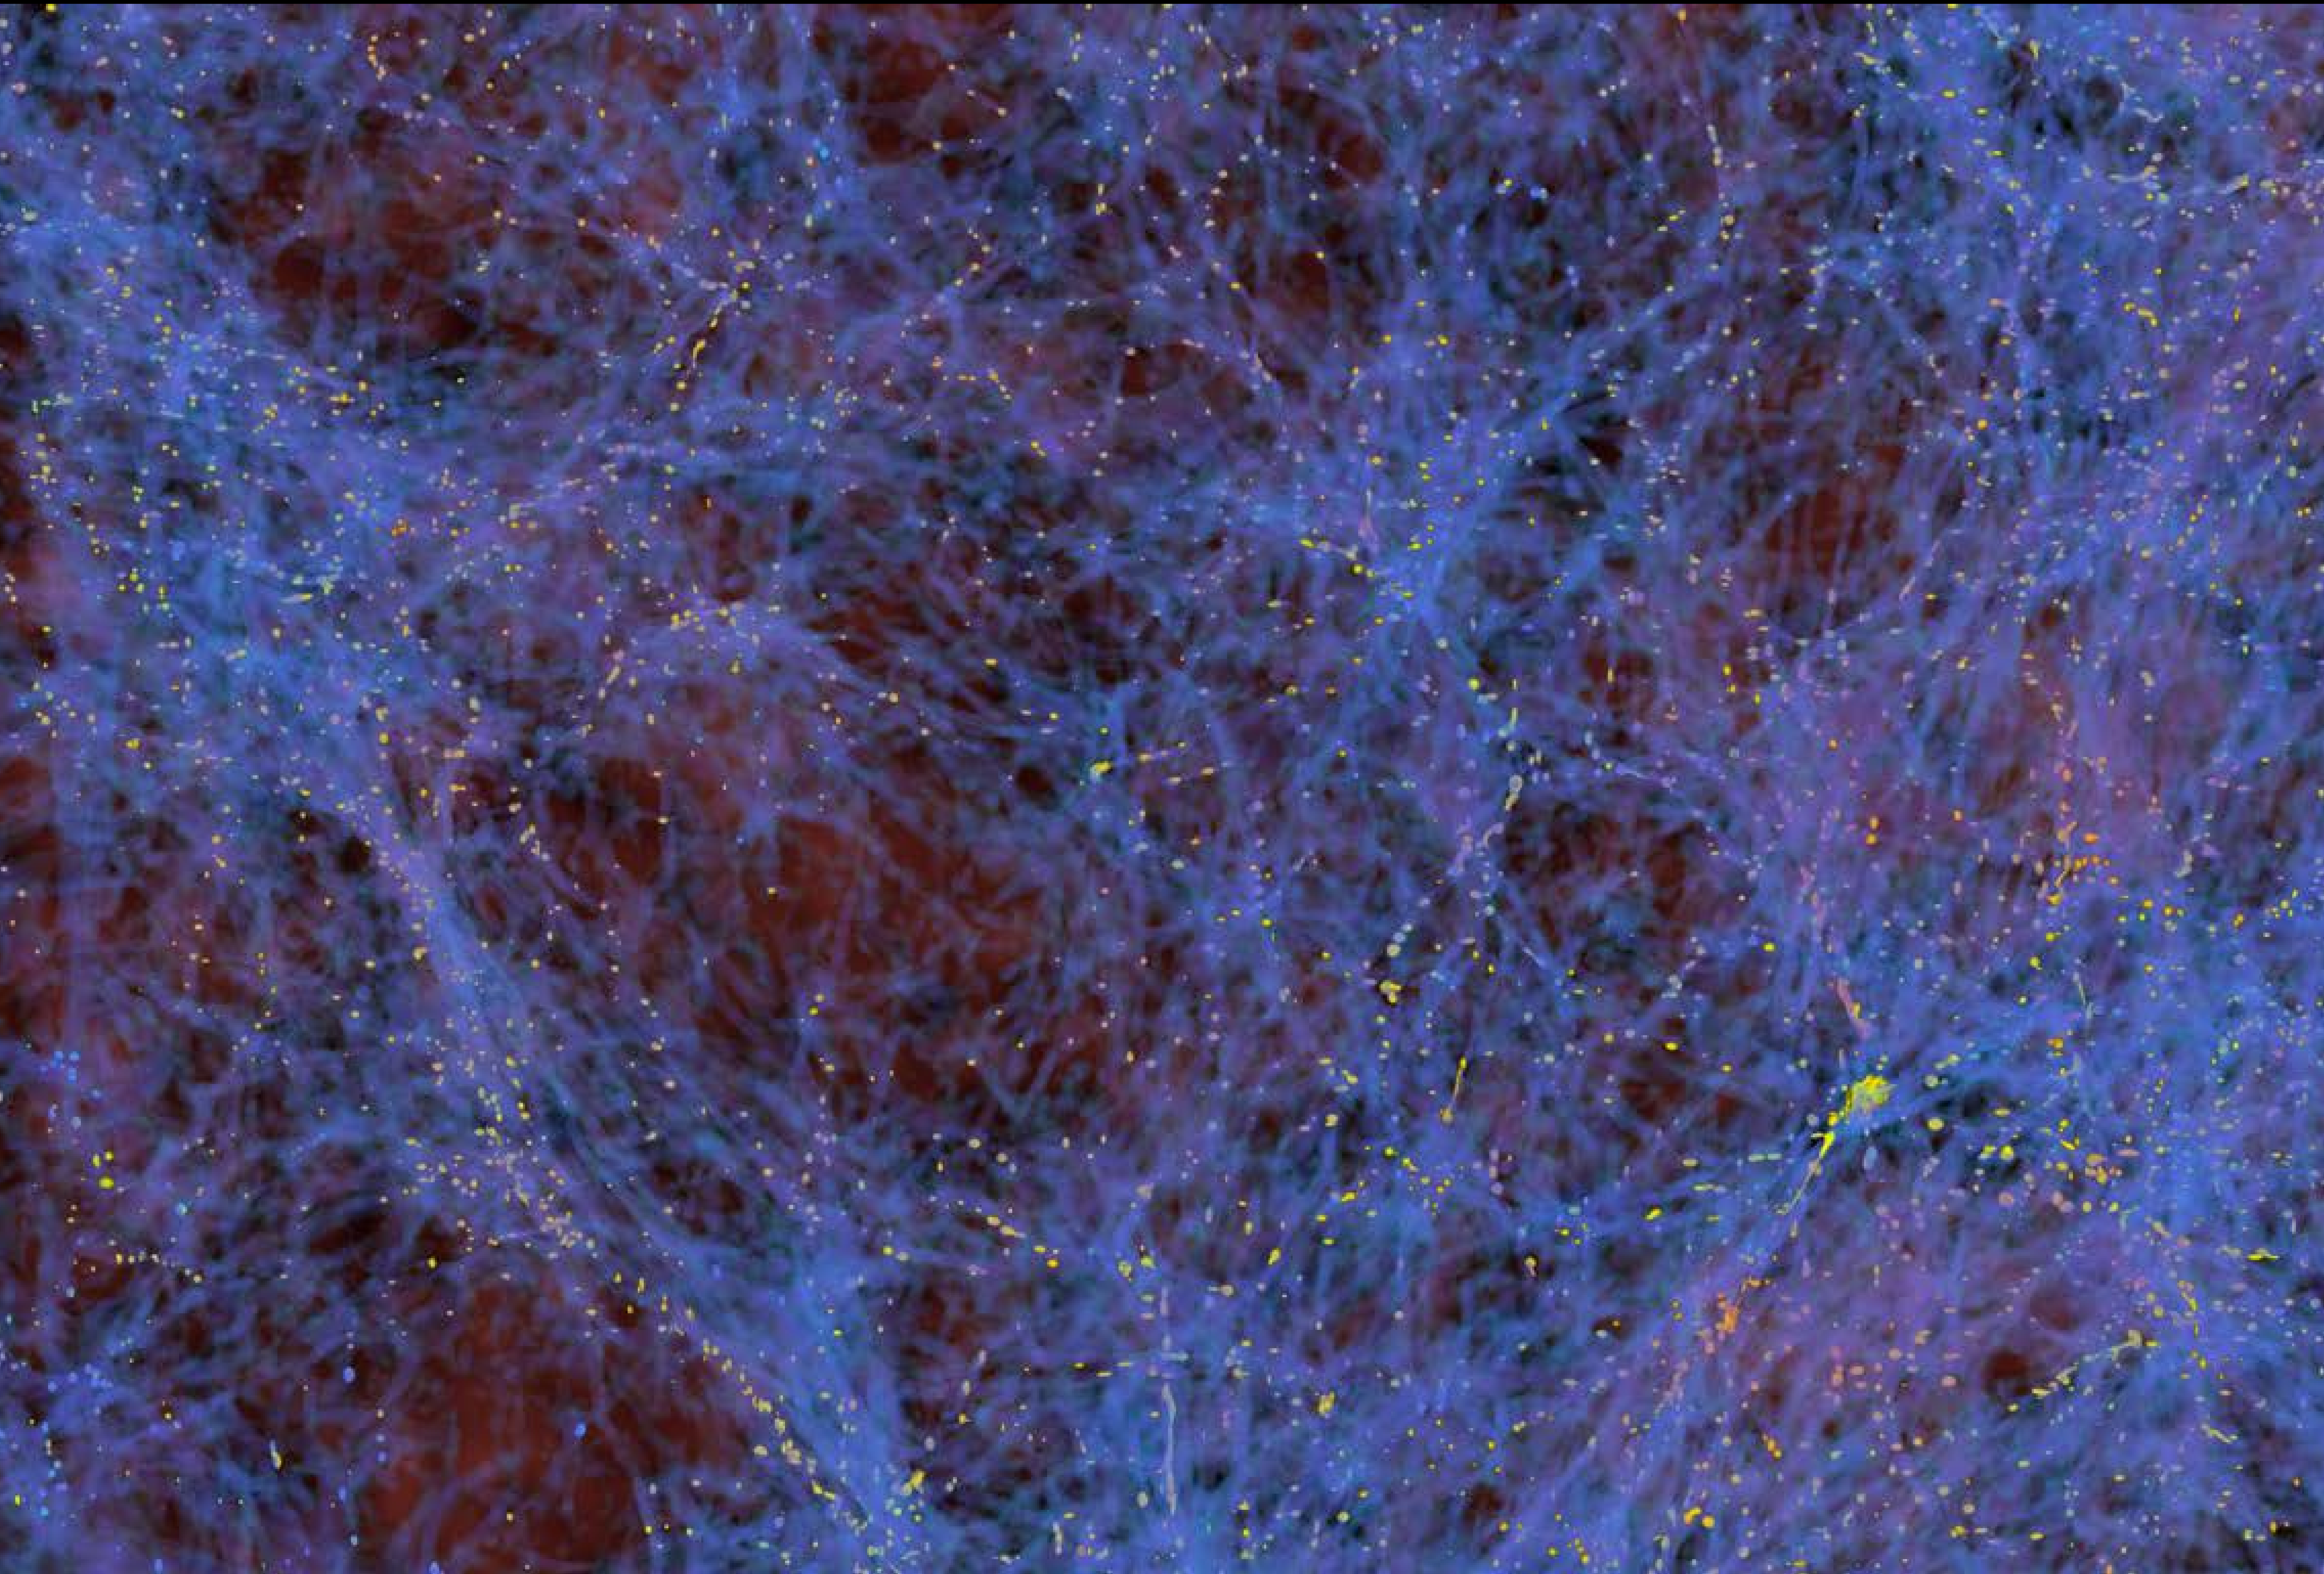

# optically-thin hydro simulations

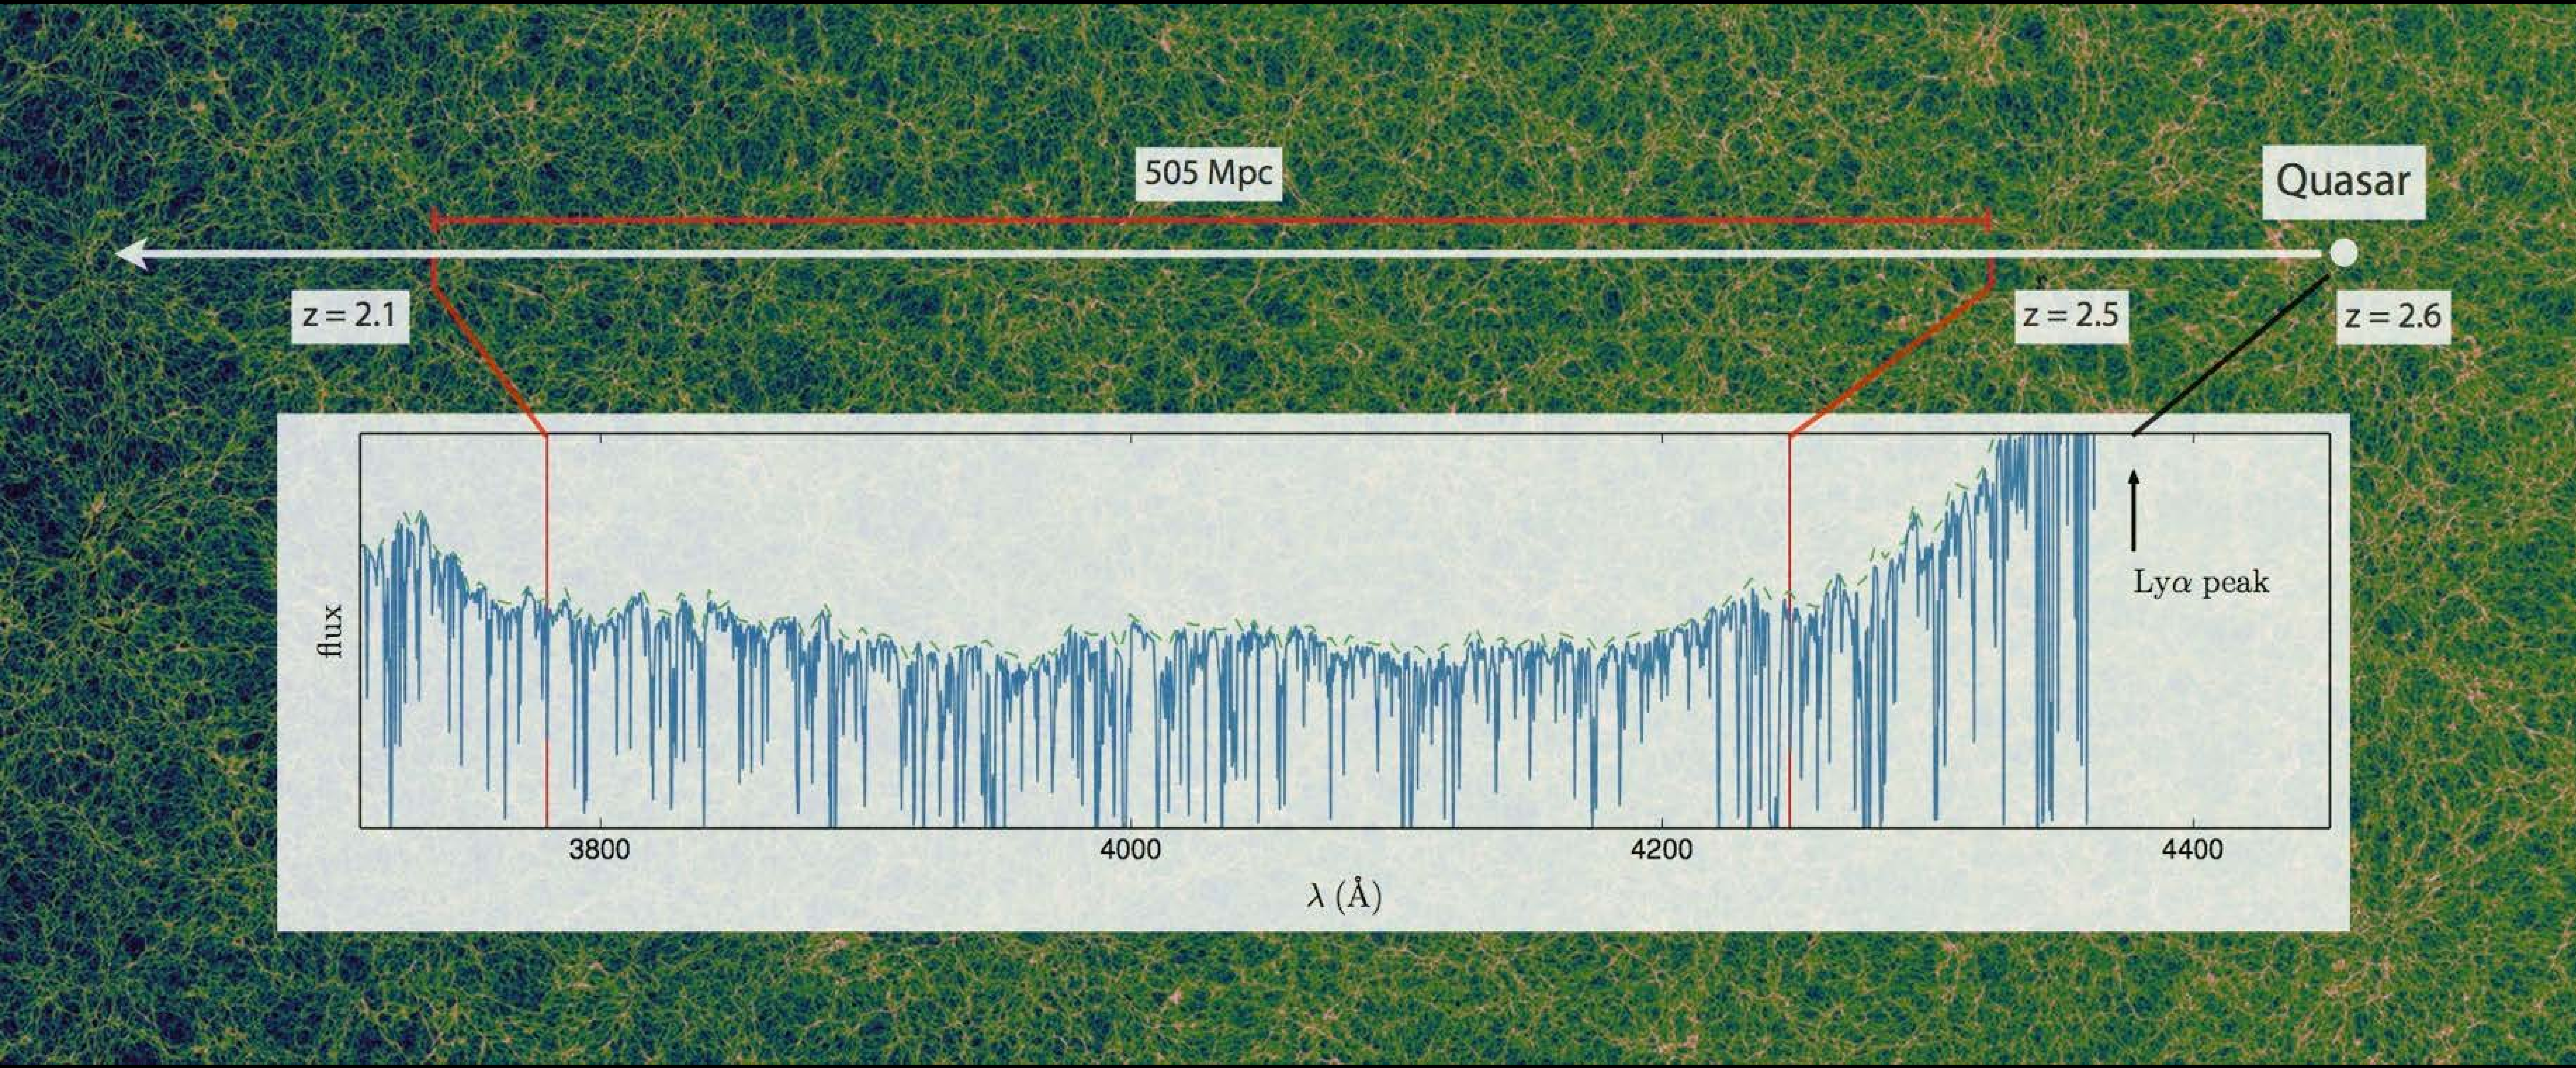


Zarija Lukić, Lawrence Berkeley National Laboratory

(Casey Stark, Peter Nugent, Martin White, Avery Meiksin, Ann Almgren)

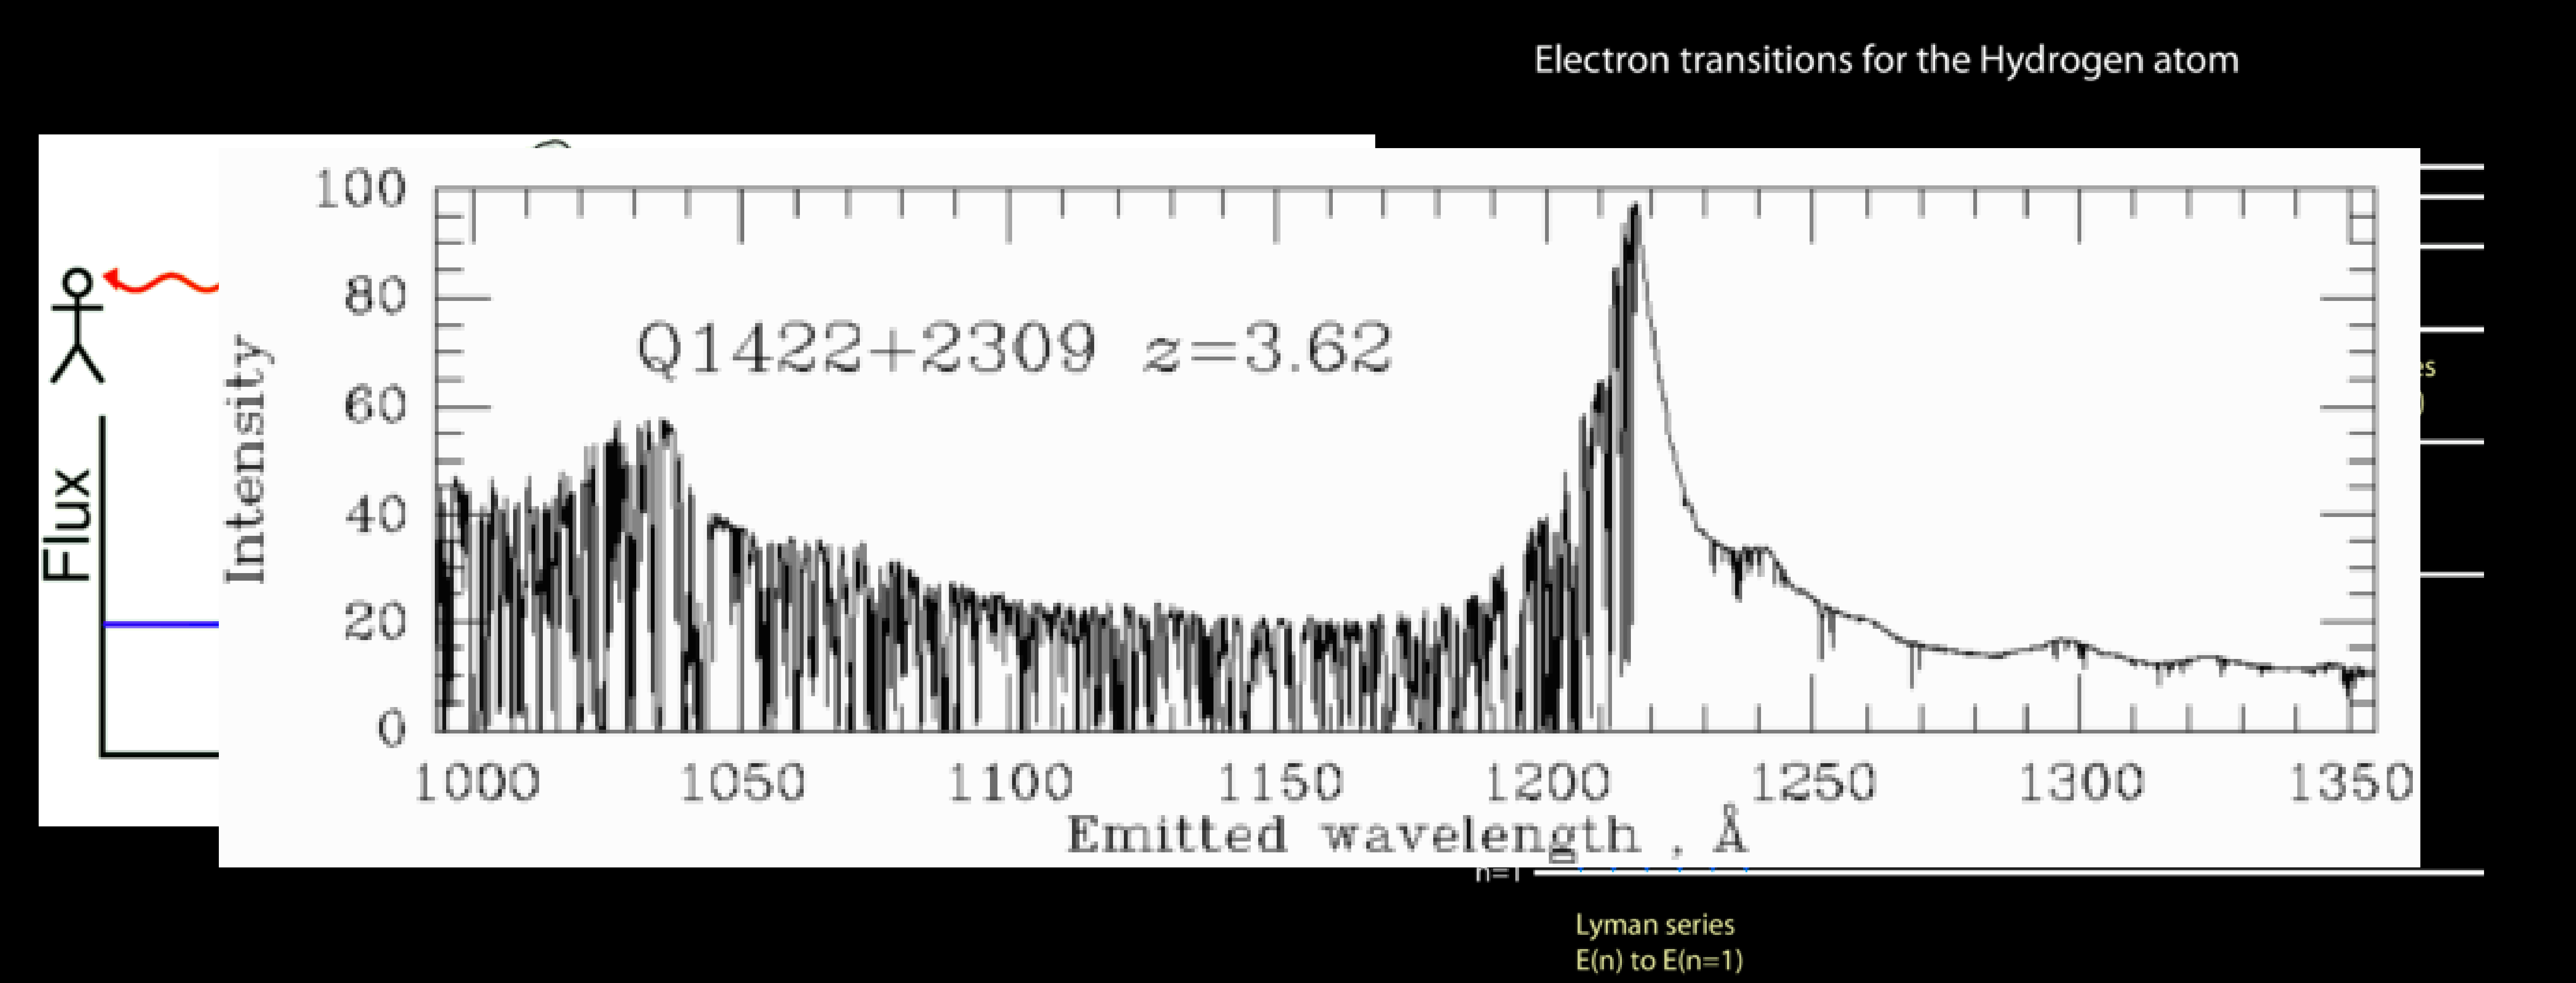

Electron transitions for the Hydrogen atom
Flux
Intensity
Q1422+2309 z=3.62
100
80
60
40
20
0
1000
1050
1100
1150
1200
1250
1300
1350
Emitted wavelength , Å
Lyman series
E(n) to E(n=1)


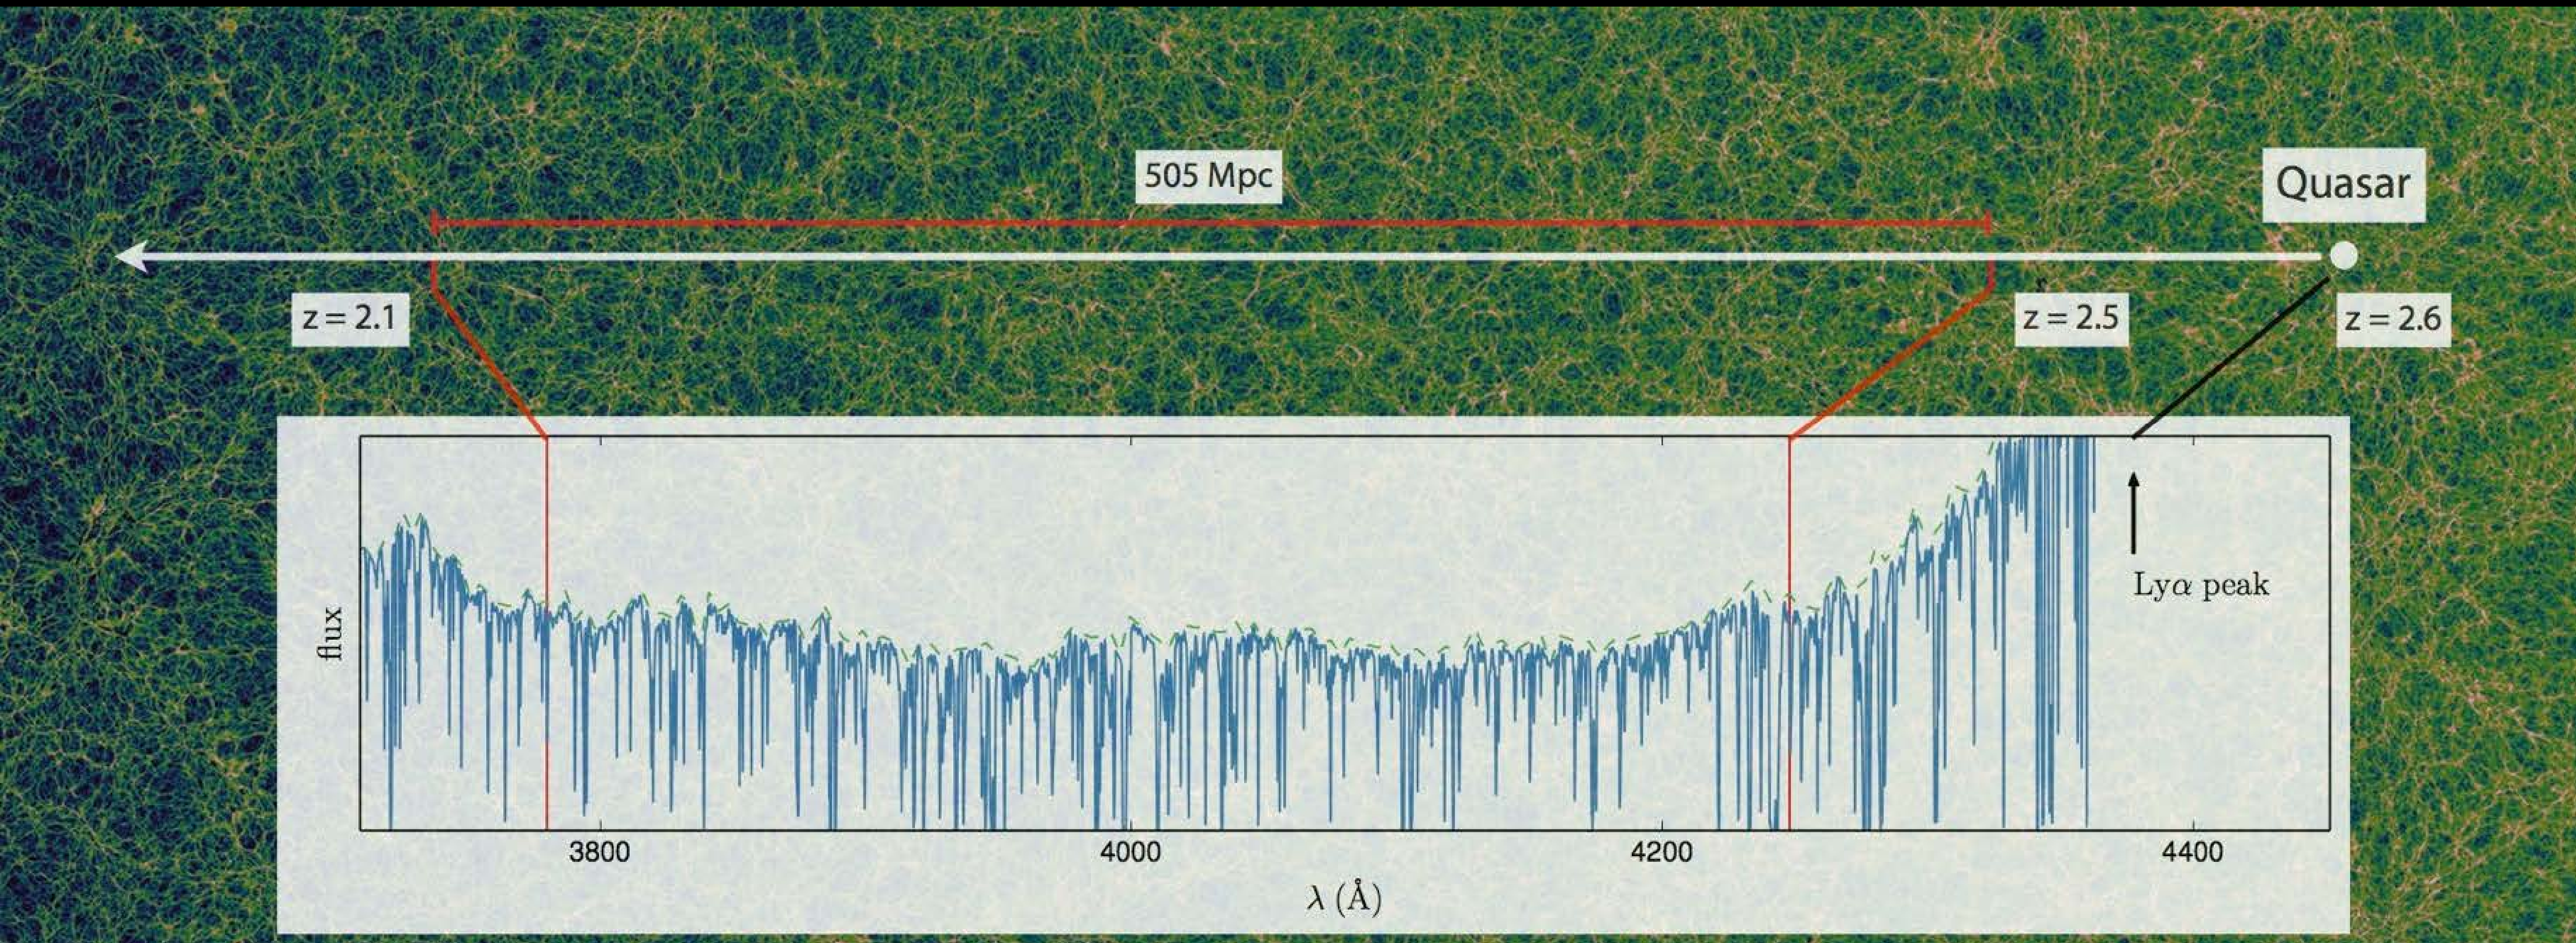

505 Mpc
Quasar
z = 2.1
z = 2.5
z = 2.6
Lyα peak
flux
3800
4000
4200
4400
λ (Å)

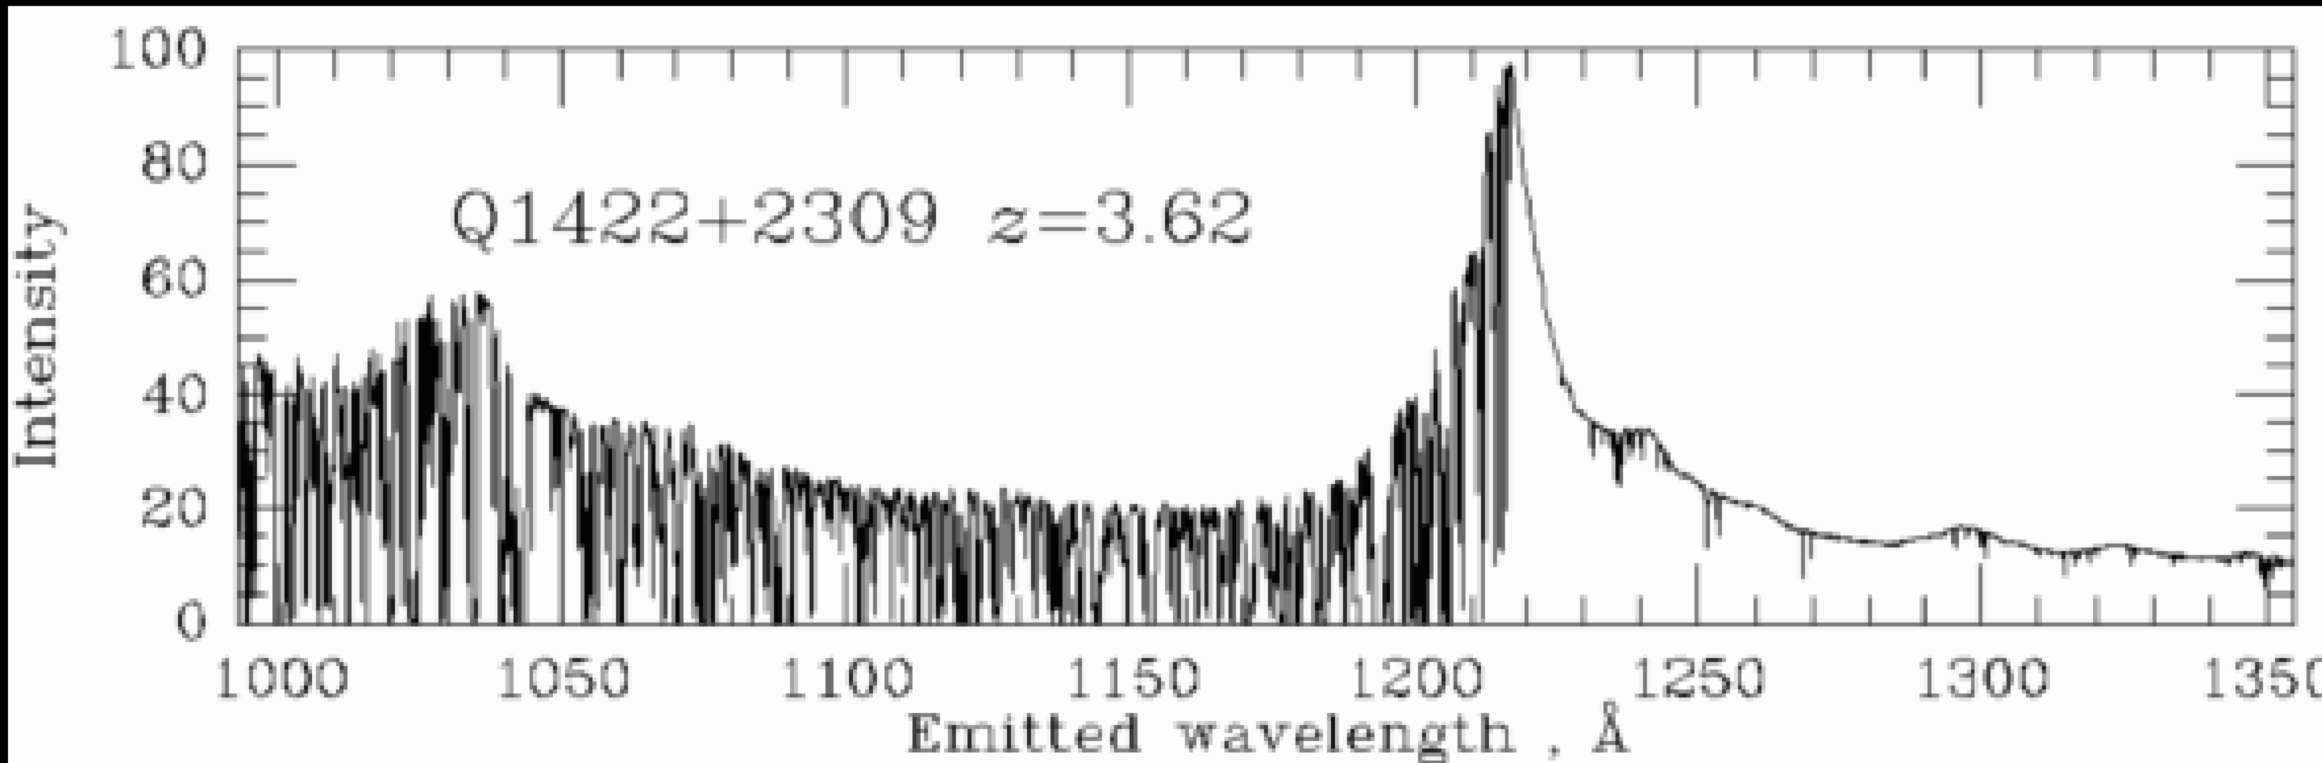
Q1422+2309 z=3.62
Intensity
Emitted wavelength , Å
0
20
40
60
80
100
1000
1050
1100
1150
1200
1250
1300
1350

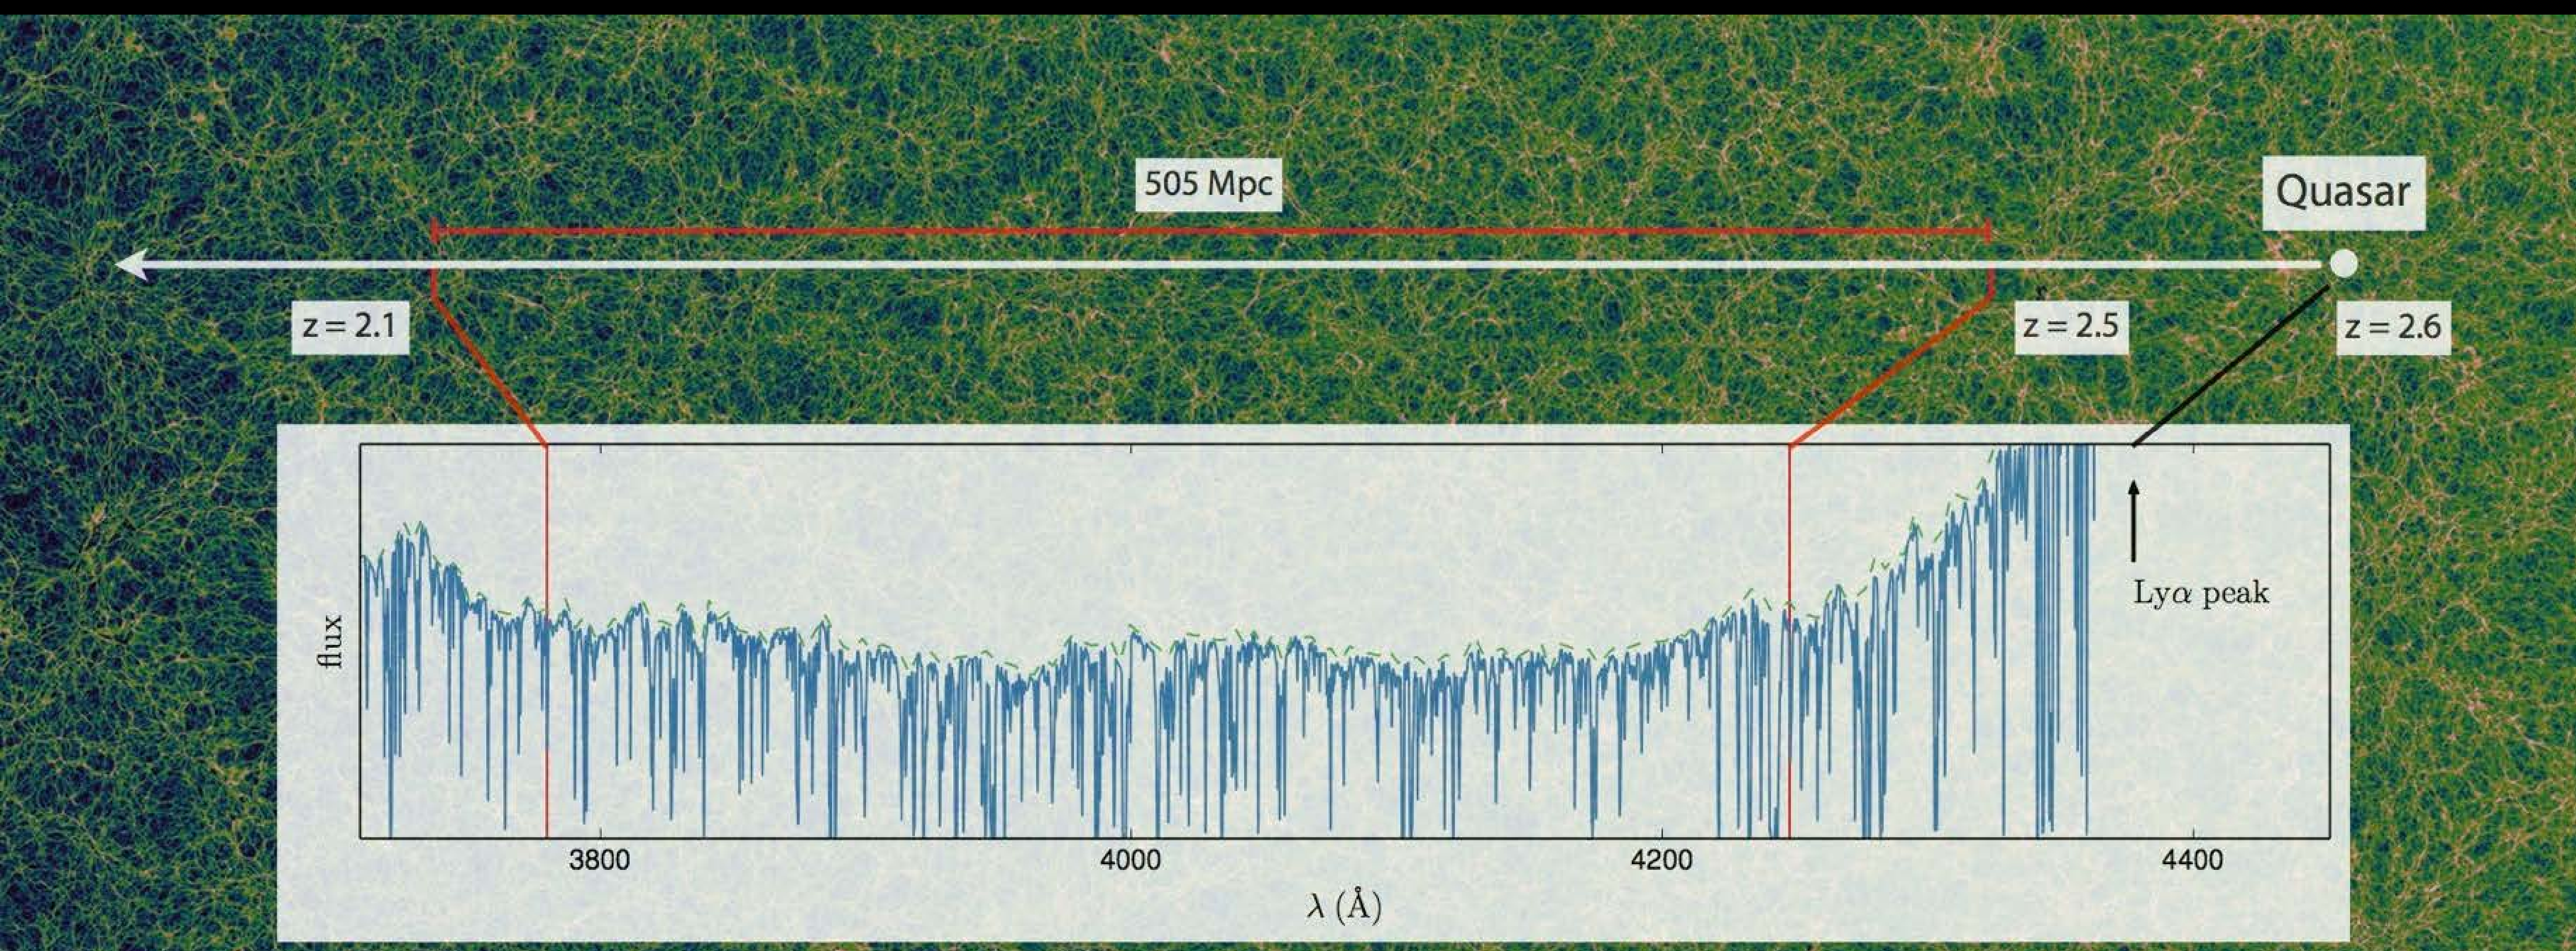
505 Mpc
Quasar
z = 2.1
z = 2.5
z = 2.6
Lyα peak
flux
3800
4000
4200
4400
λ (Å)

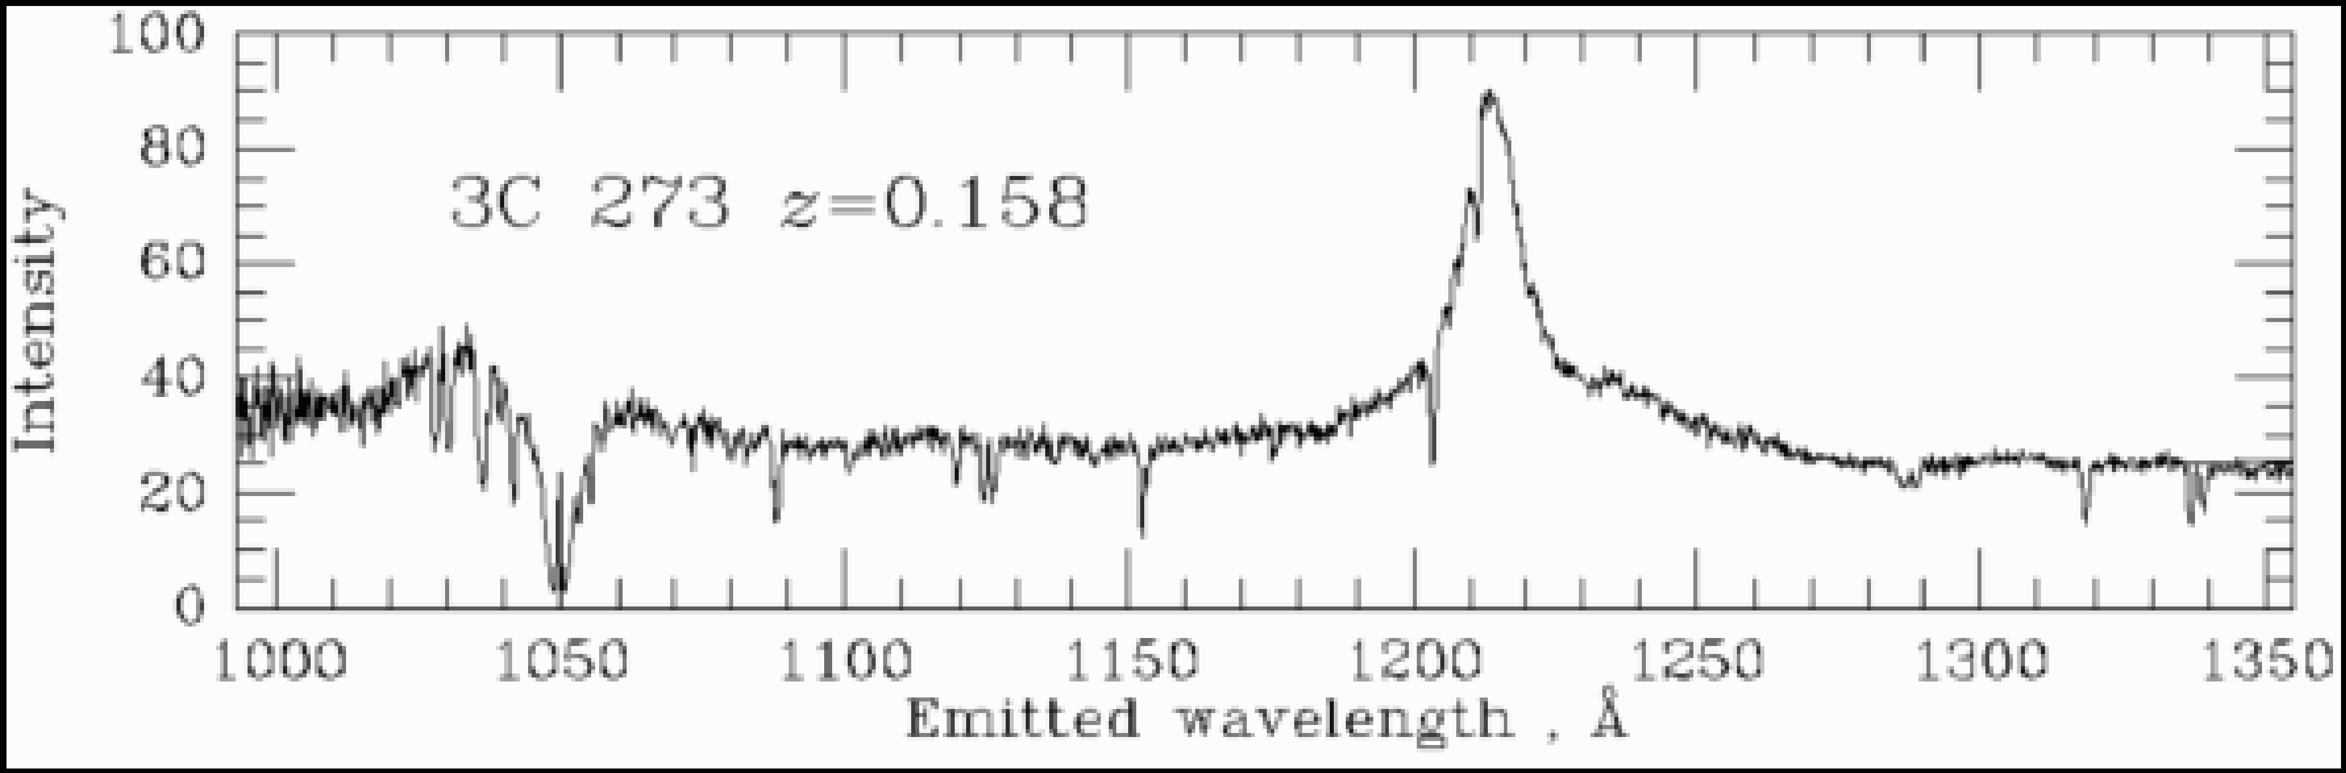
3C 273 z=0.158
Intensity
Emitted wavelength , Å
100
80
60
40
20
0
1000
1050
1100
1150
1200
1250
1300
1350

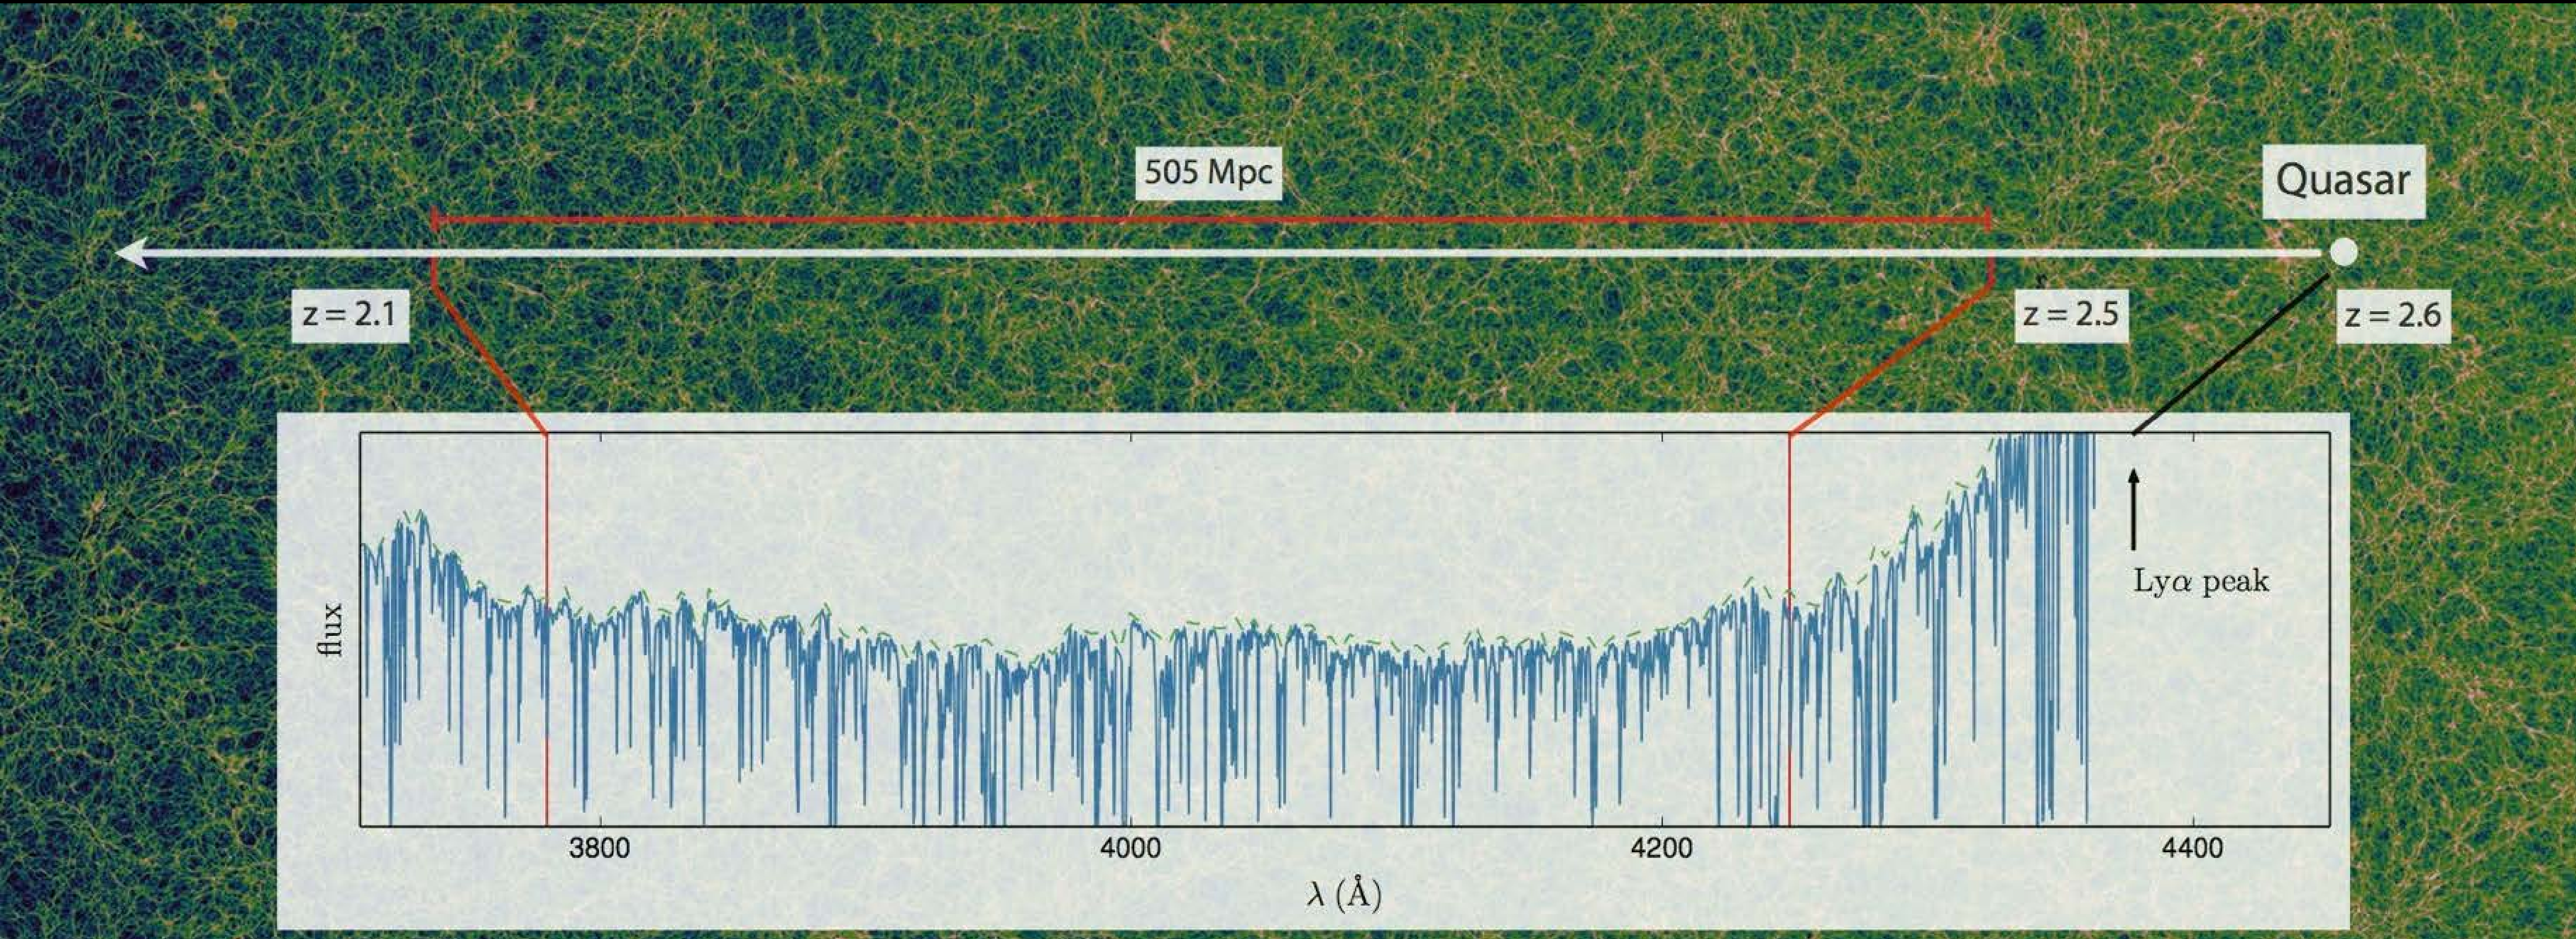
505 Mpc
Quasar
z = 2.1
z = 2.5
z = 2.6
Lyα peak
flux
3800
4000
4200
4400
λ (Å)

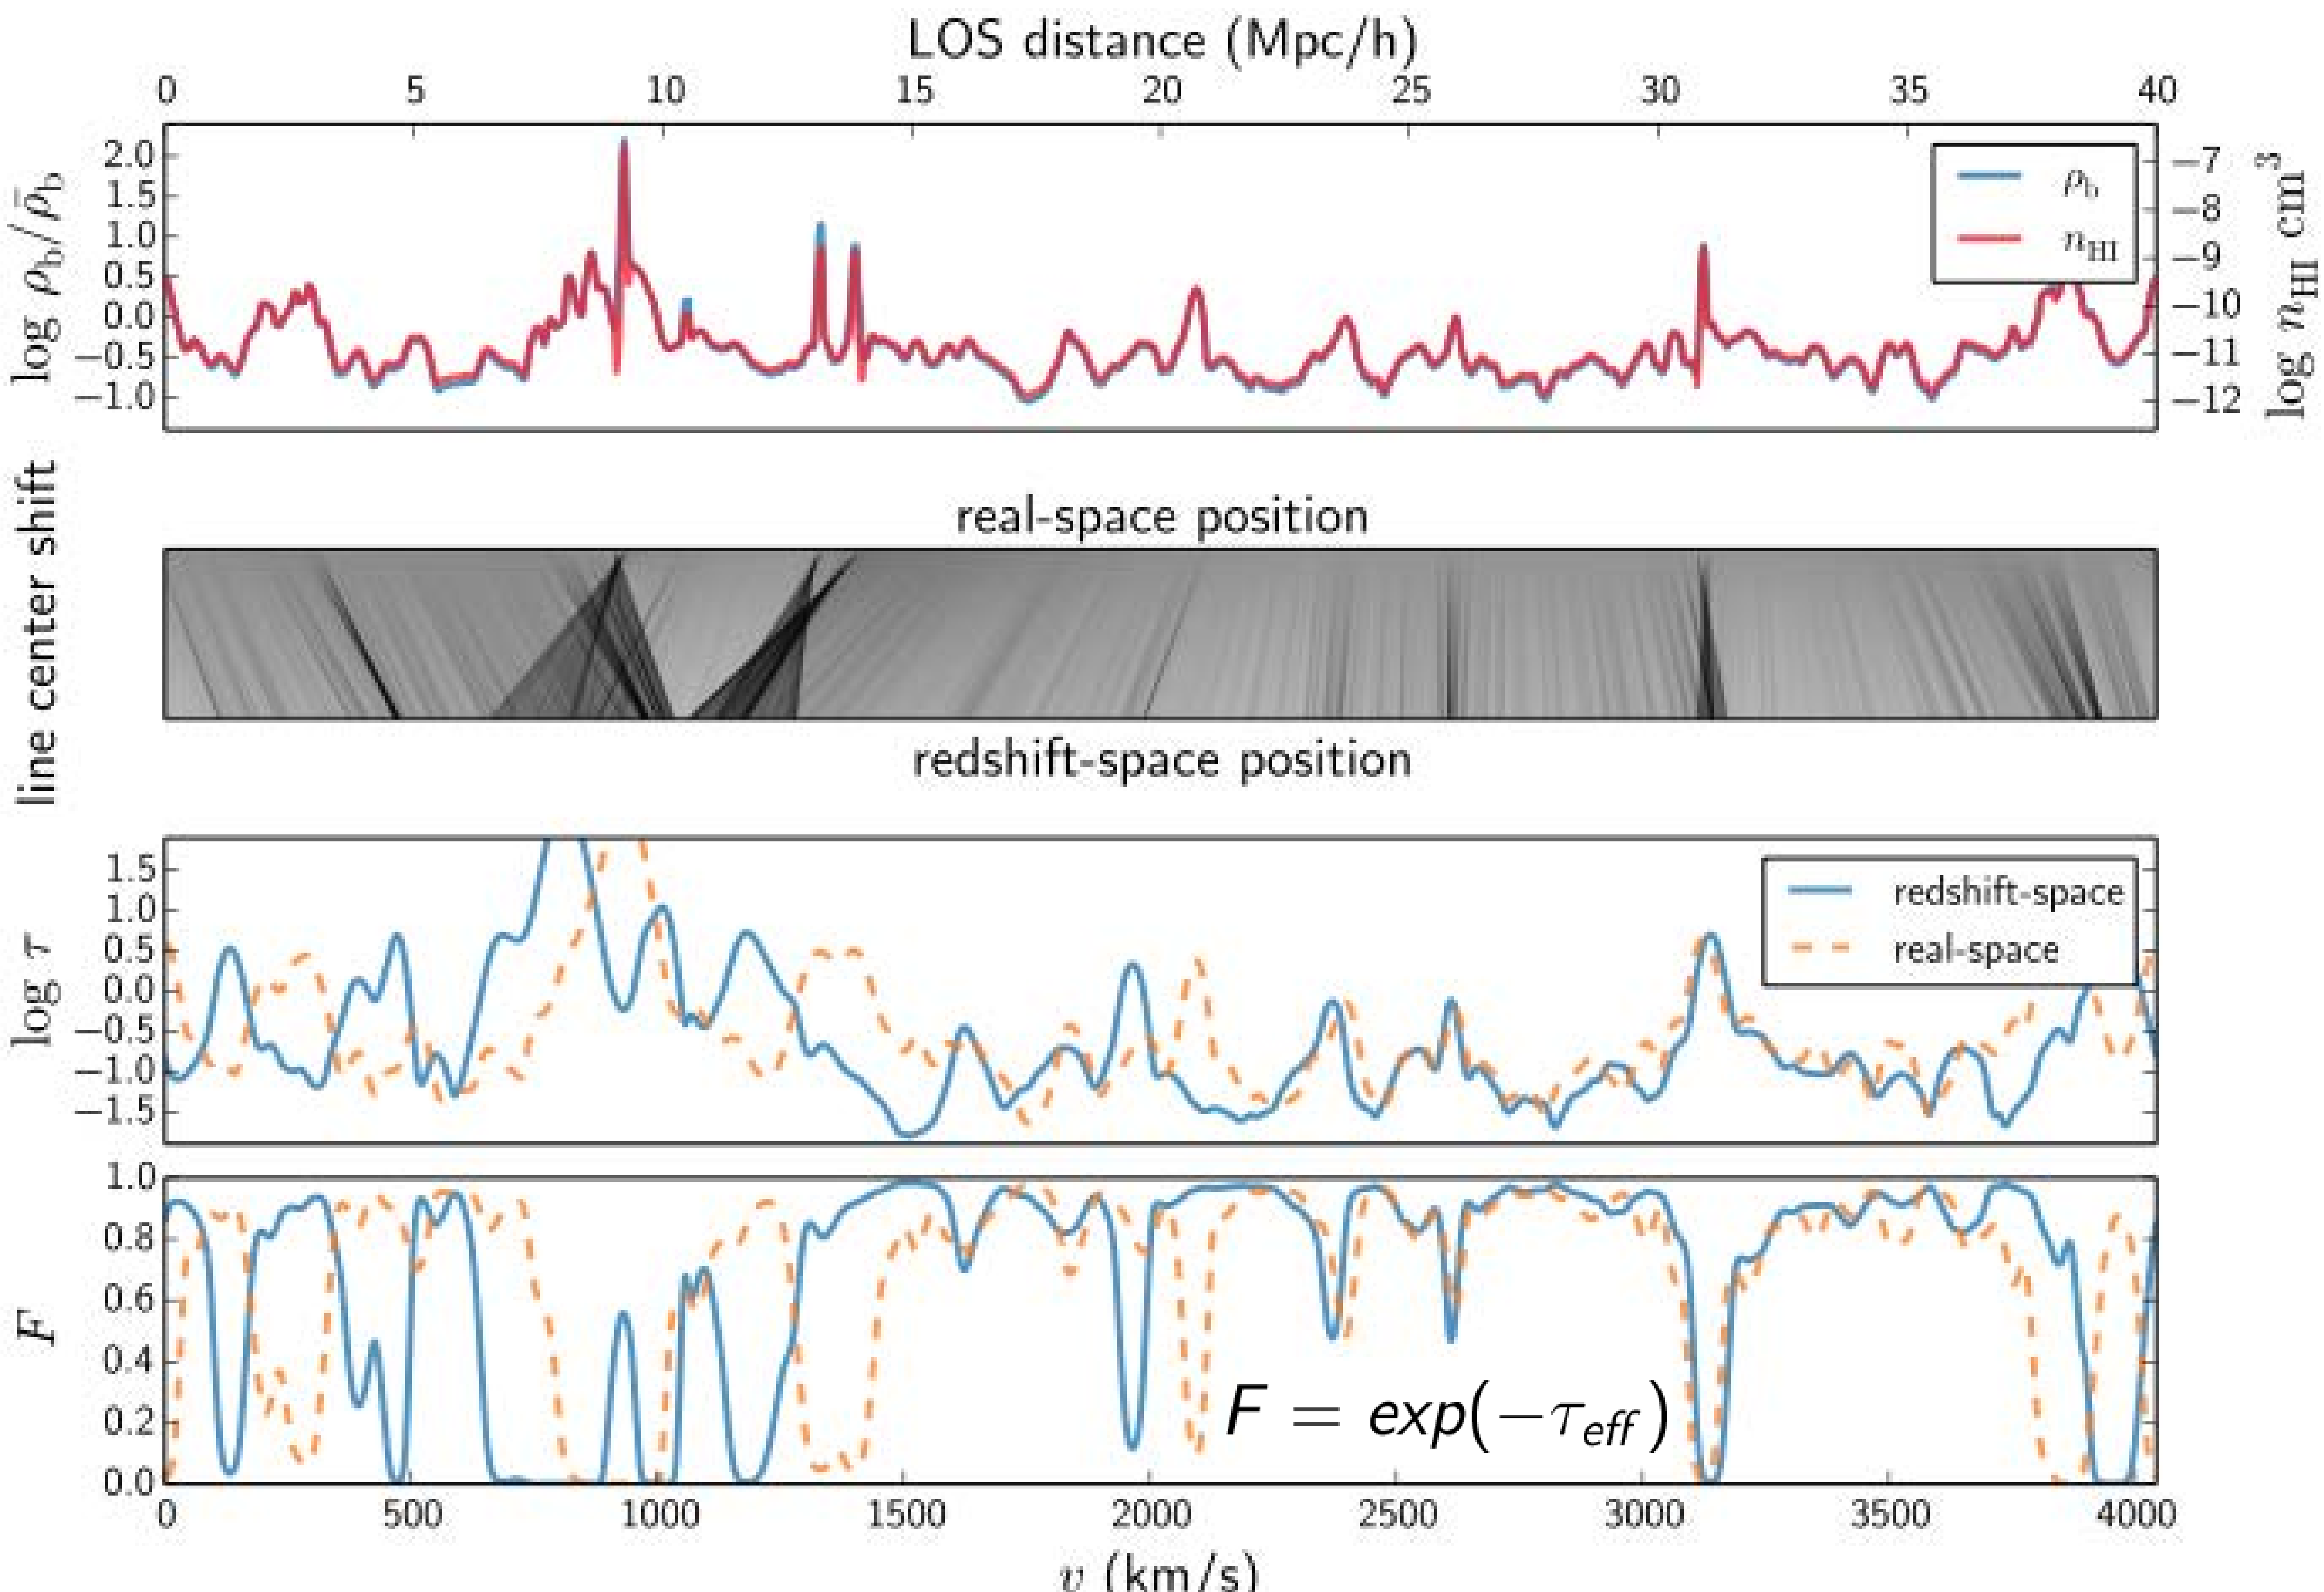

LOS distance (Mpc/h)
$\log \rho_b/\bar{\rho}_b$
$\log n_{HI}\ cm^3$
$\rho_b$
$n_{HI}$
real-space position
line center shift
redshift-space position
$\log \tau$
redshift-space
real-space
$F$
$F = exp(-\tau_{eff})$
$v$ (km/s)

# Where the flux comes from

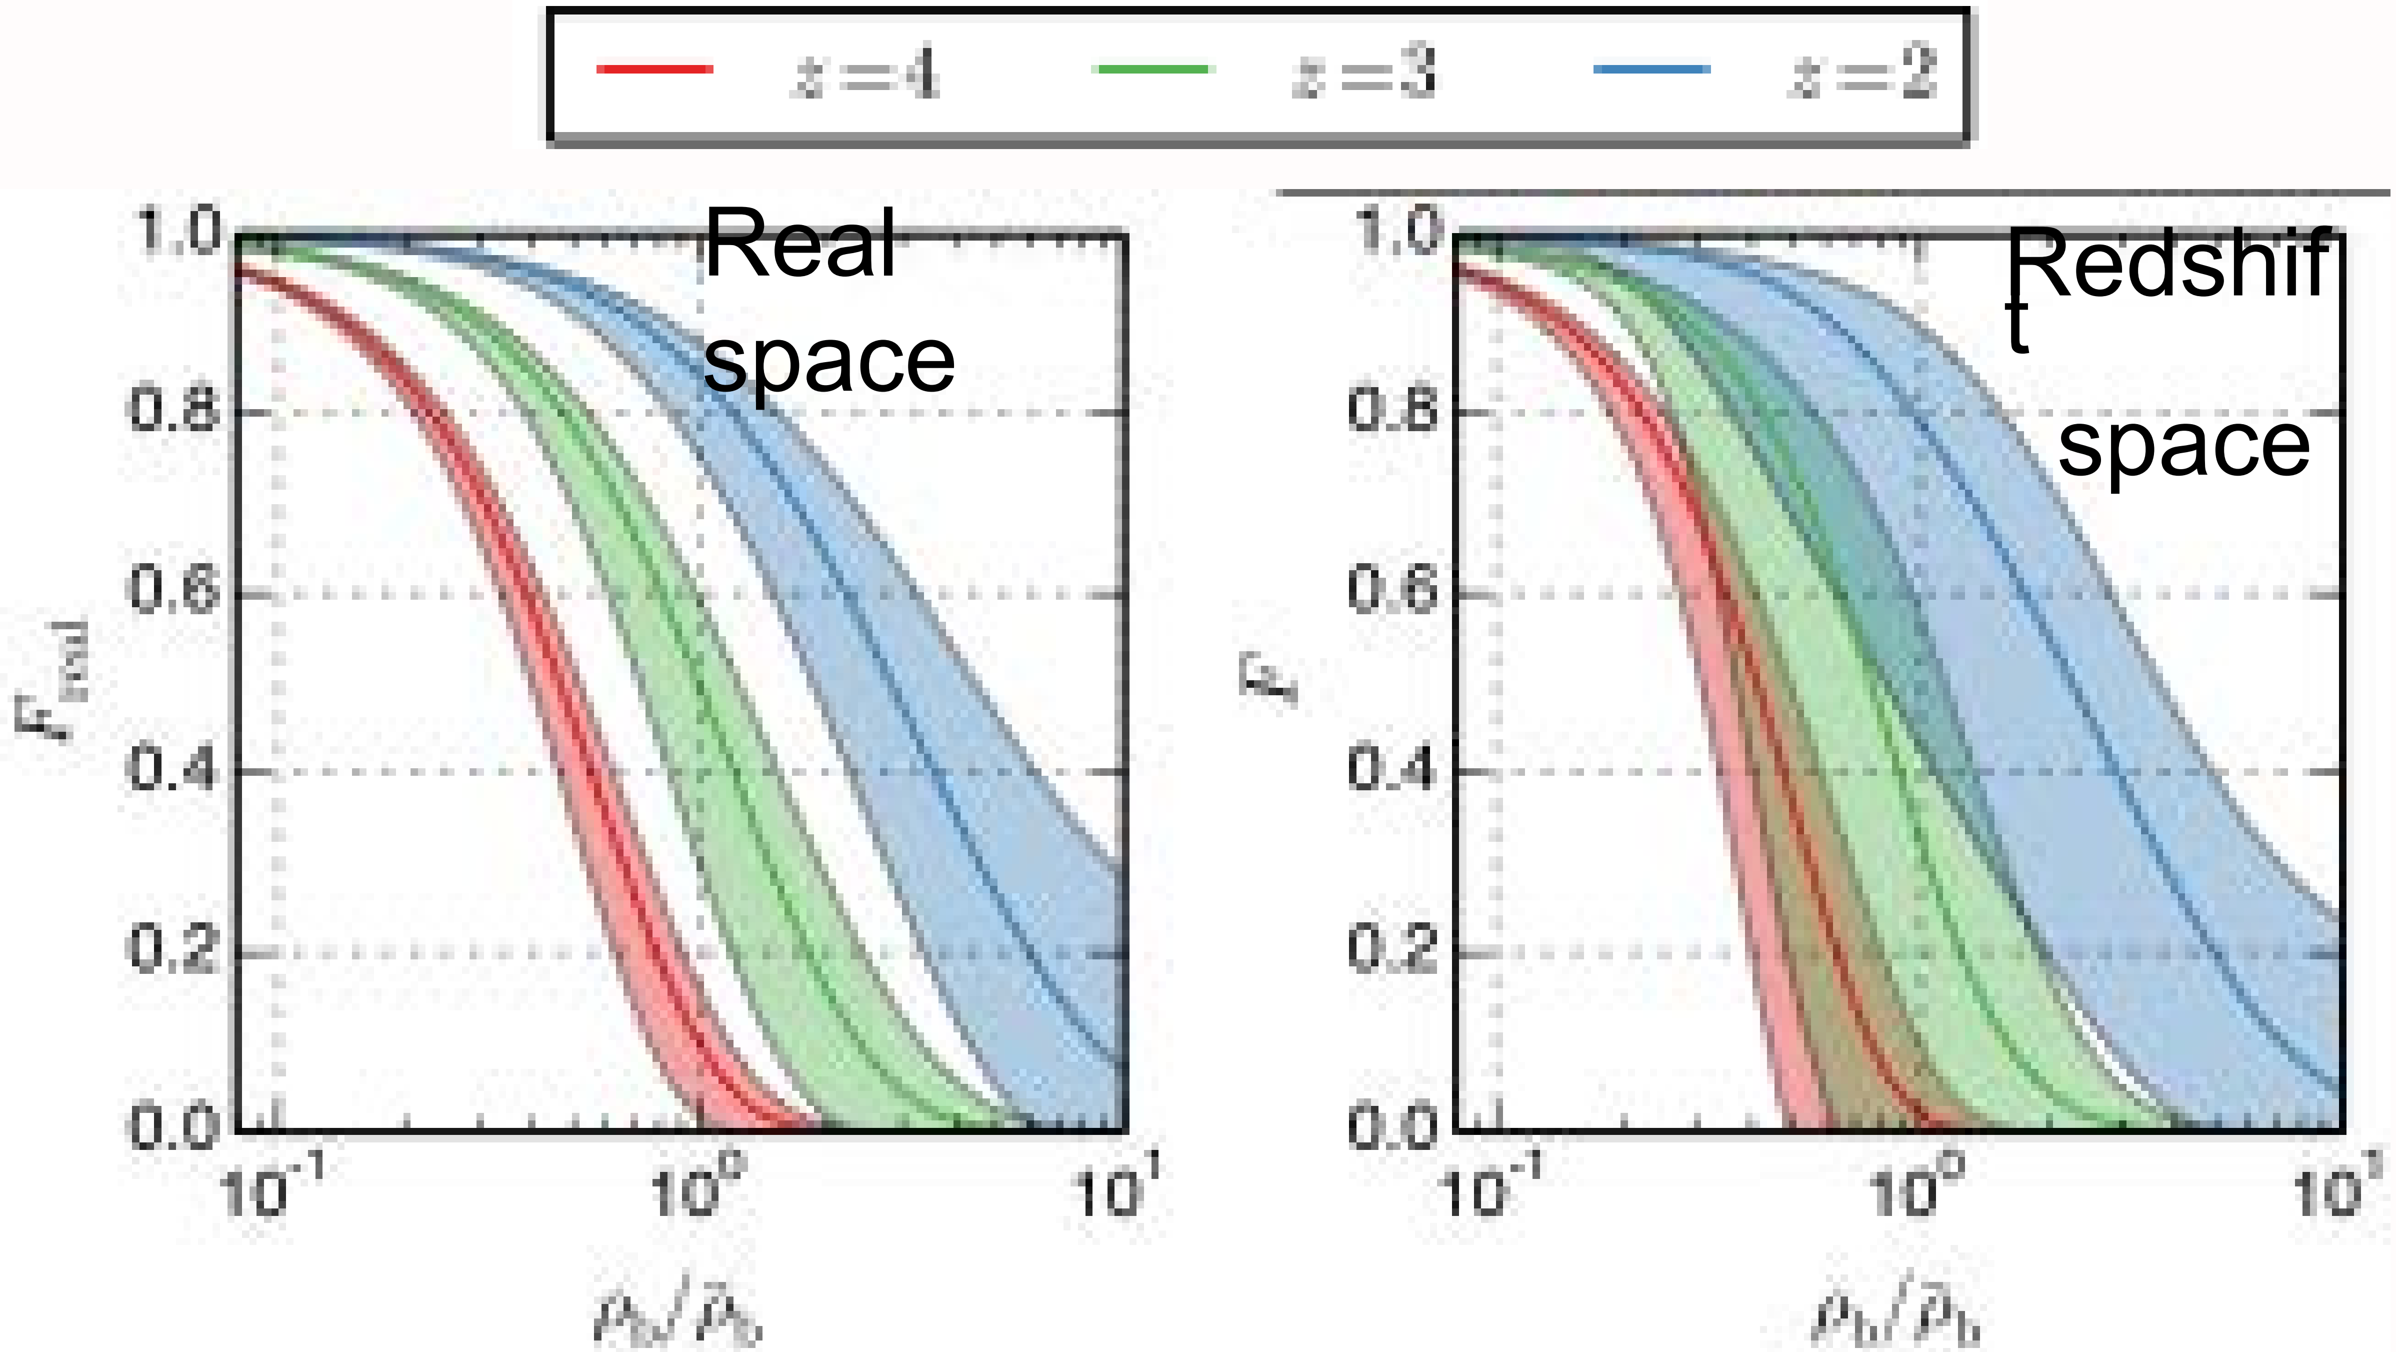


Lukić et al. 2014

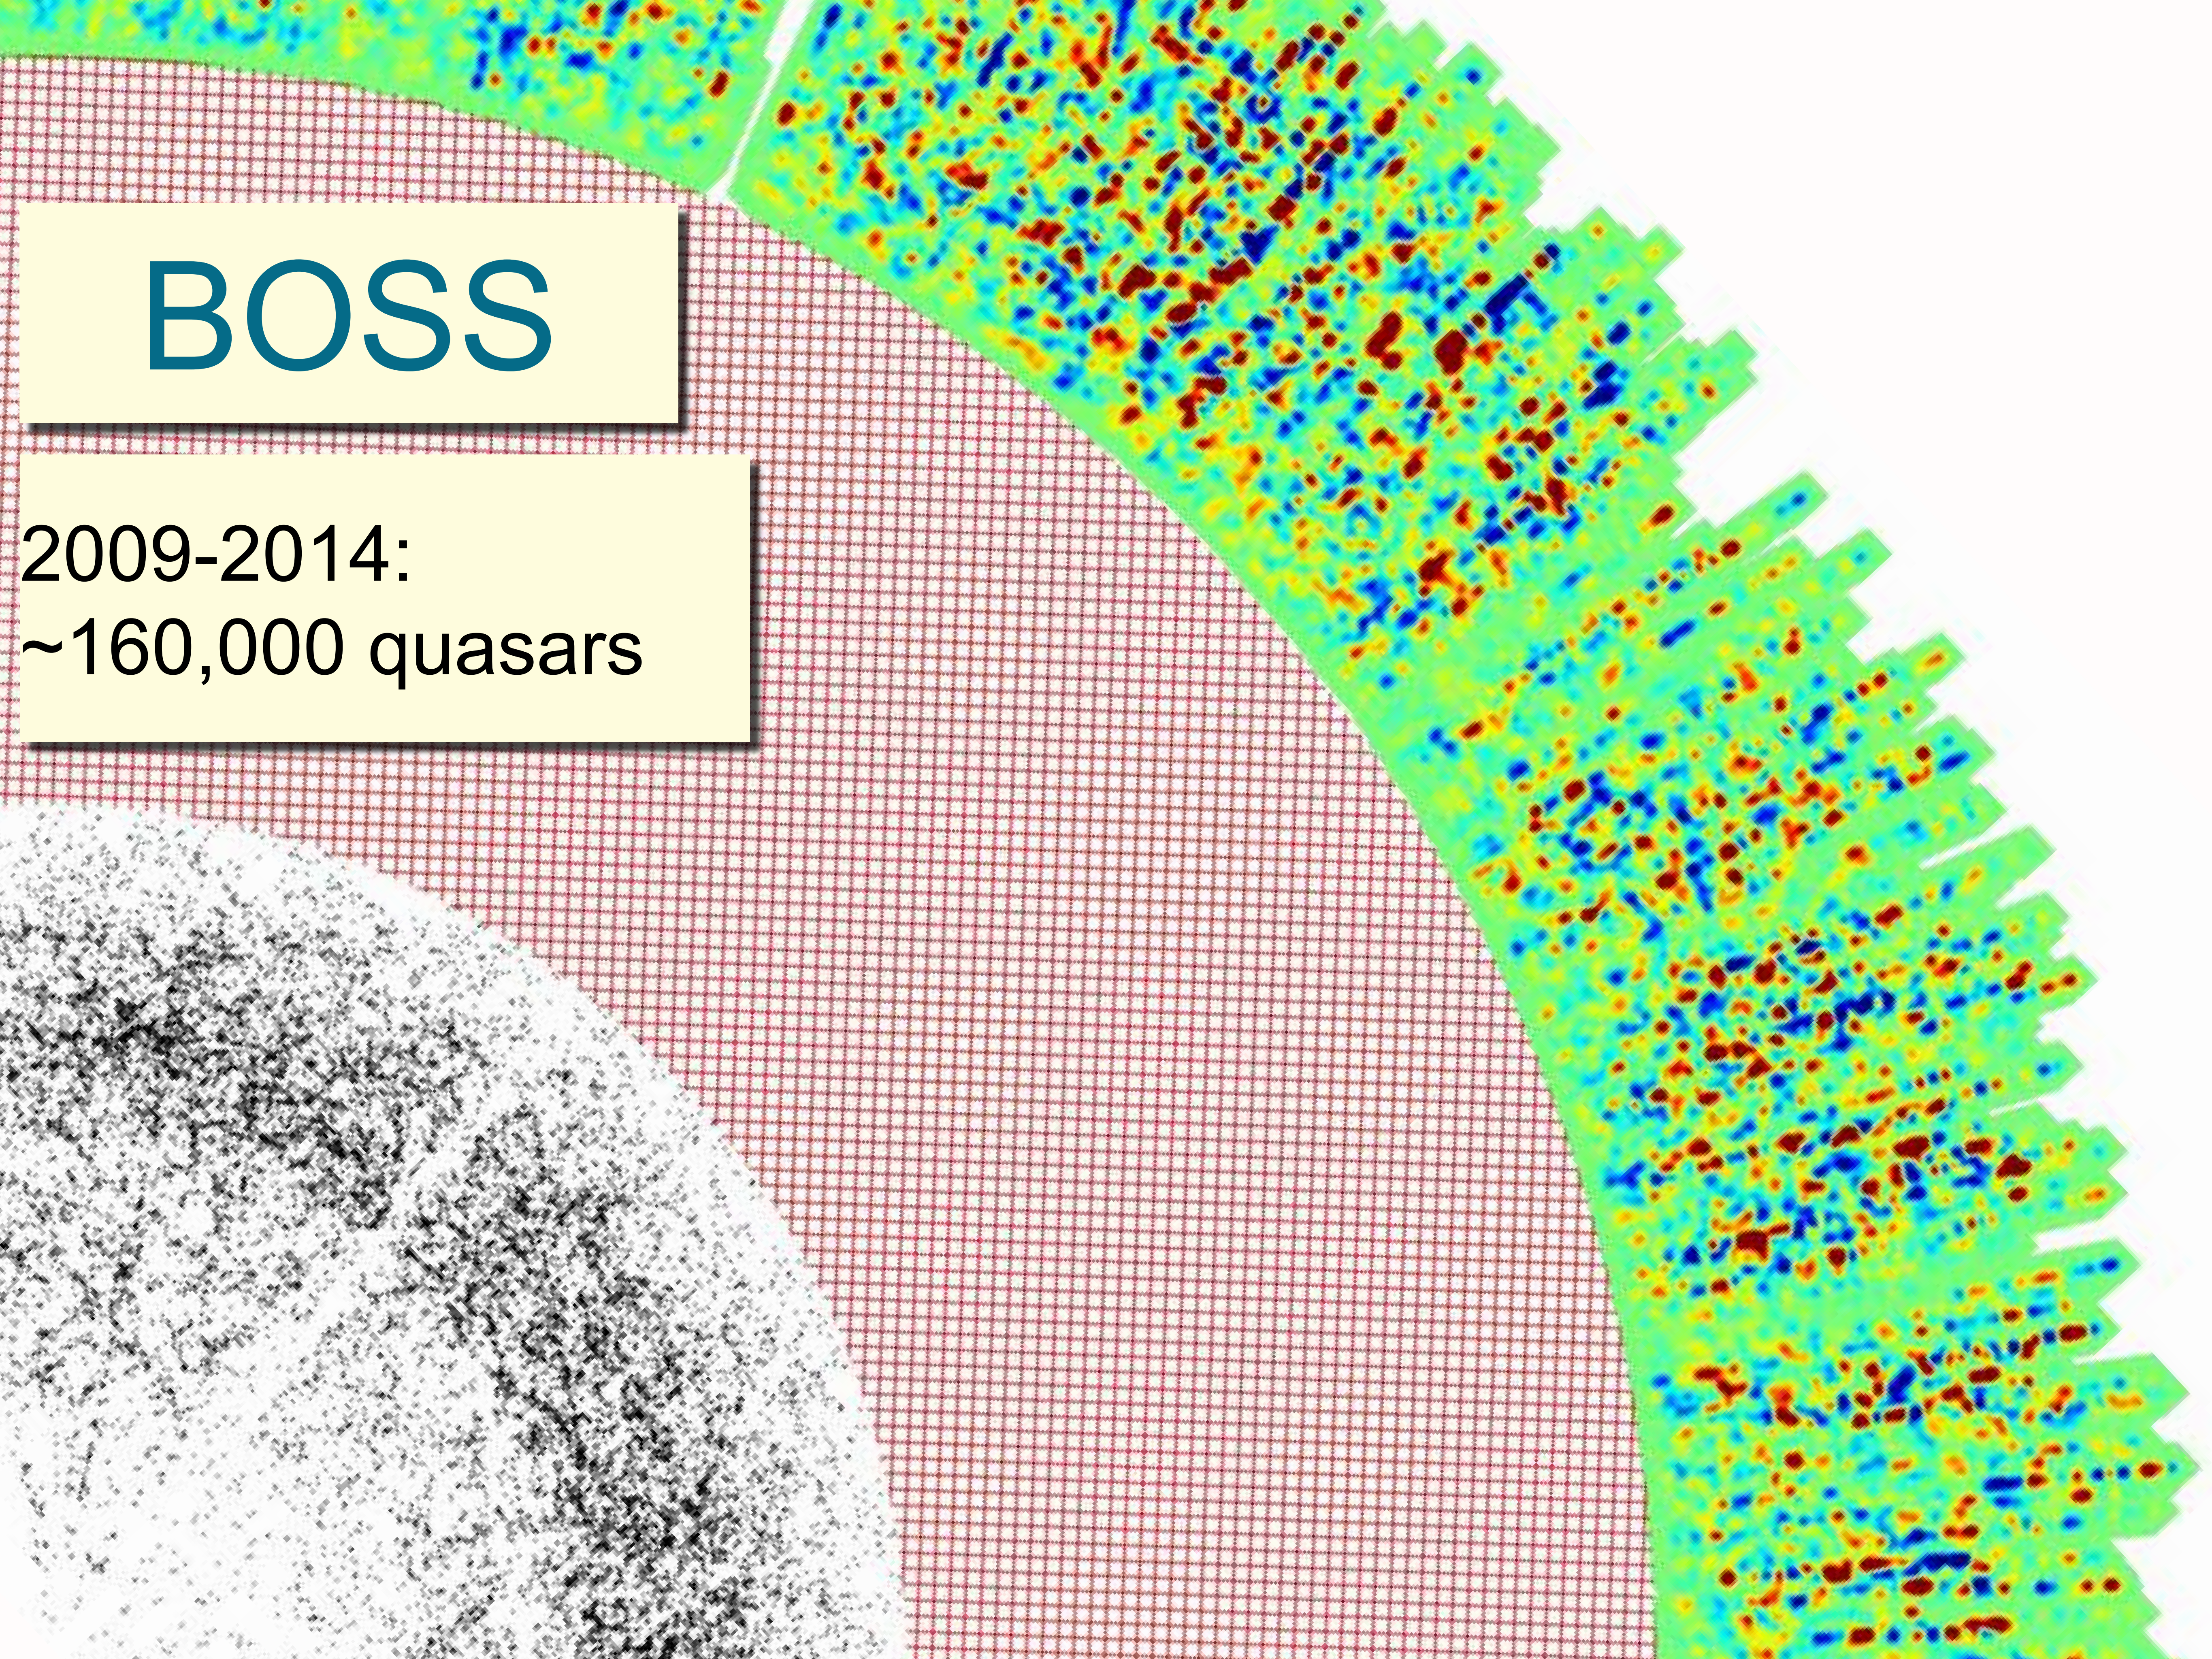
BOSS
2009-2014:
~160,000 quasars

# Density - temperature

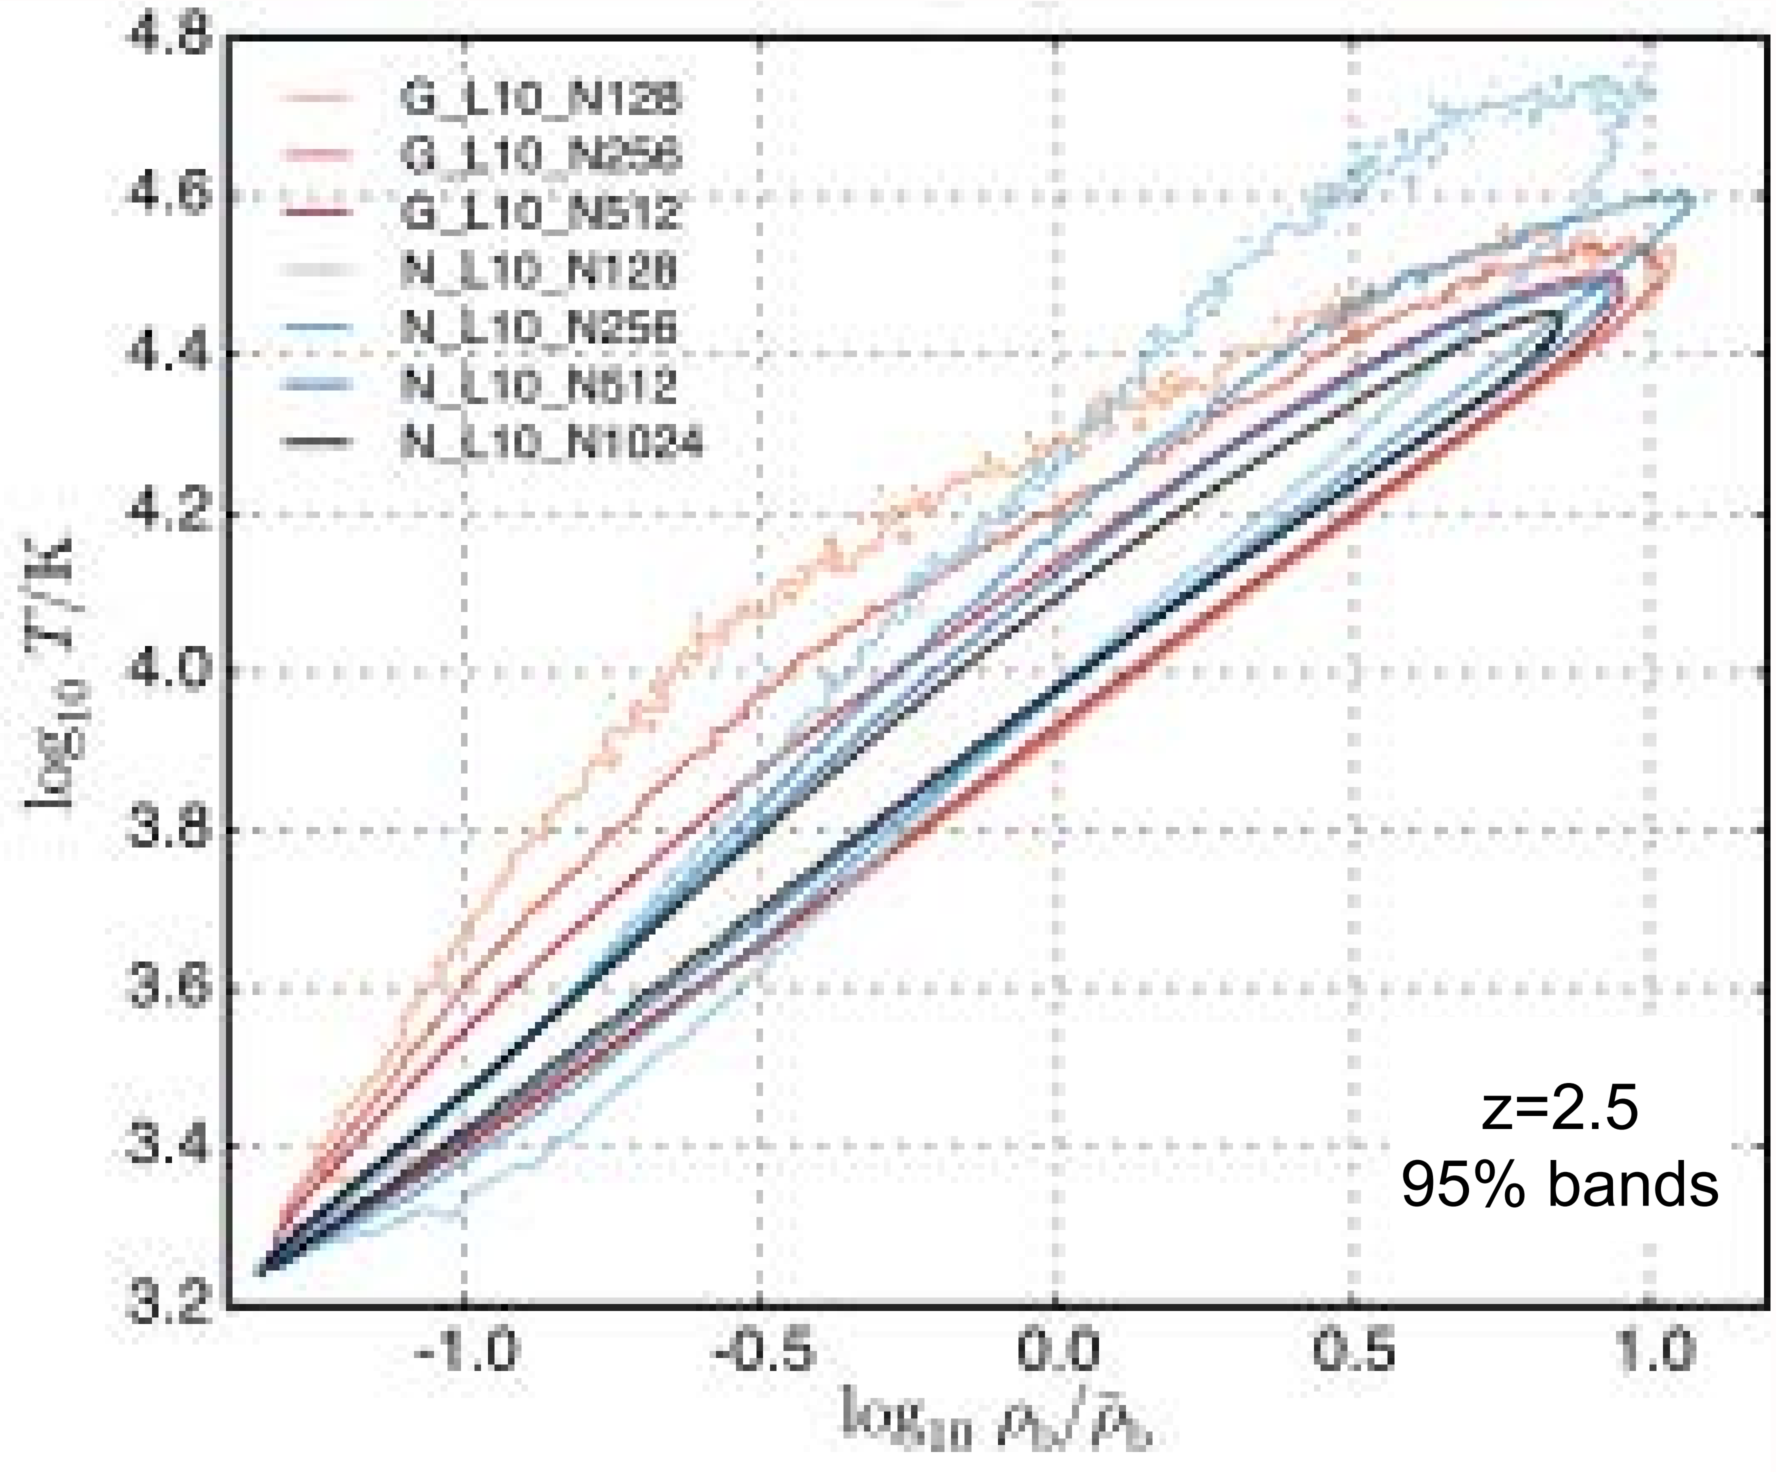


- SPH (Gadget) vs. Eulerian (Nyx) code.

Stark et al. in prep.

# The Cosmic Web in Ly-alpha

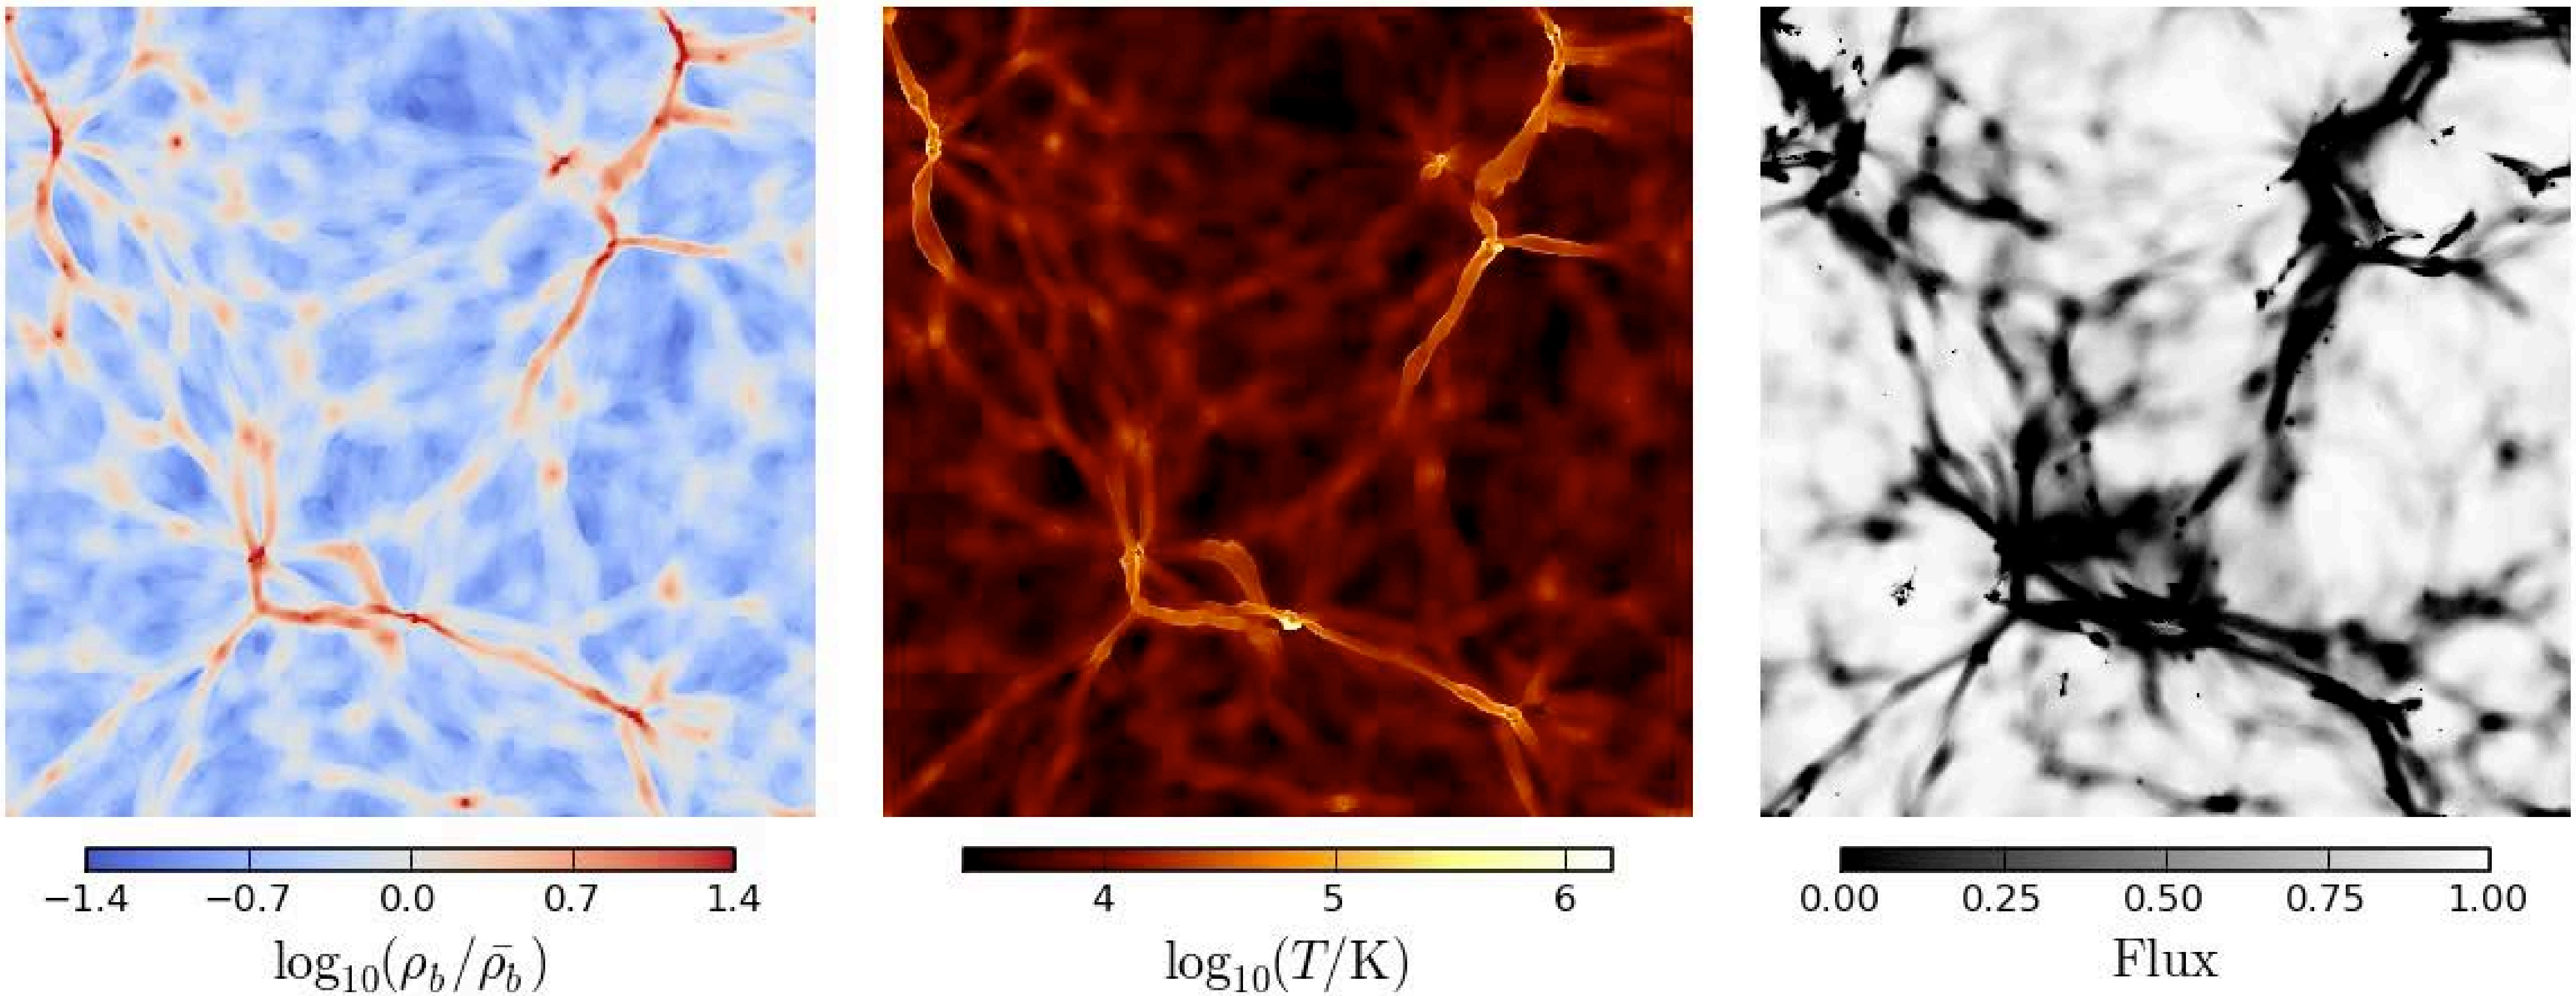


results coming soon

# In Situ Analysis

- **Data Reduction:** A trillion particle simulation with 100 analysis steps has a storage requirement of ~4 PB -- in situ analysis reduces it to ~200 TB
- **I/O Chokepoints:** Large data analyses difficult because I/O time > analysis time, plus scheduling overhead
- **Fast Algorithms:** Analysis time is only a fraction of a full simulation timestep
- **Ease of Workflow:** Large analyses difficult to manage in post-processing

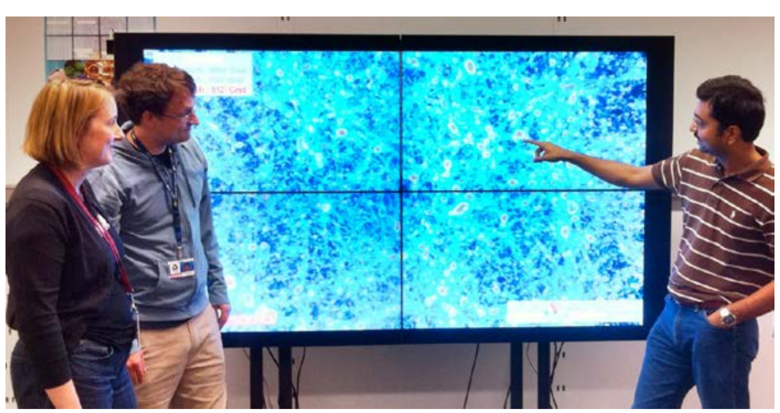

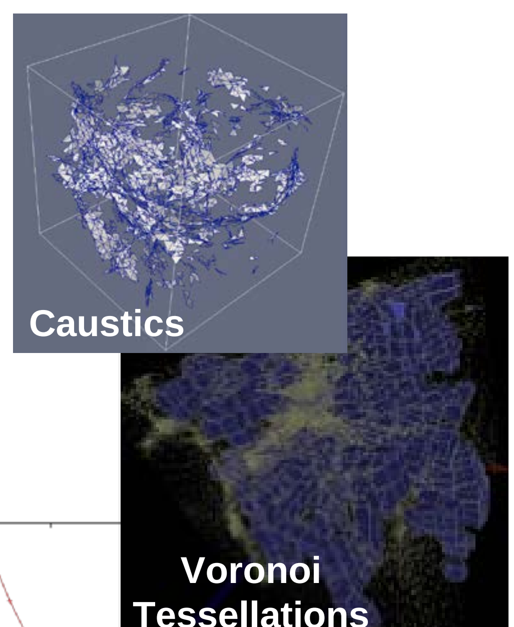


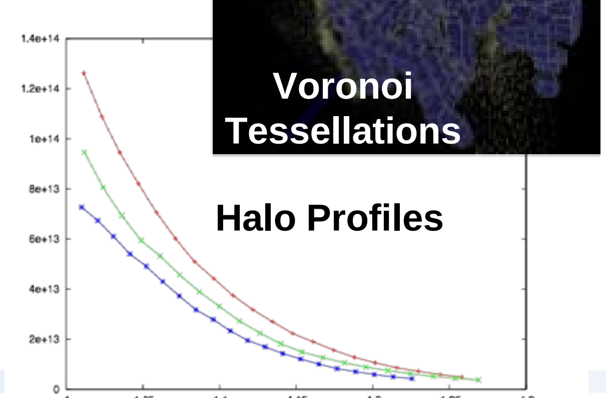


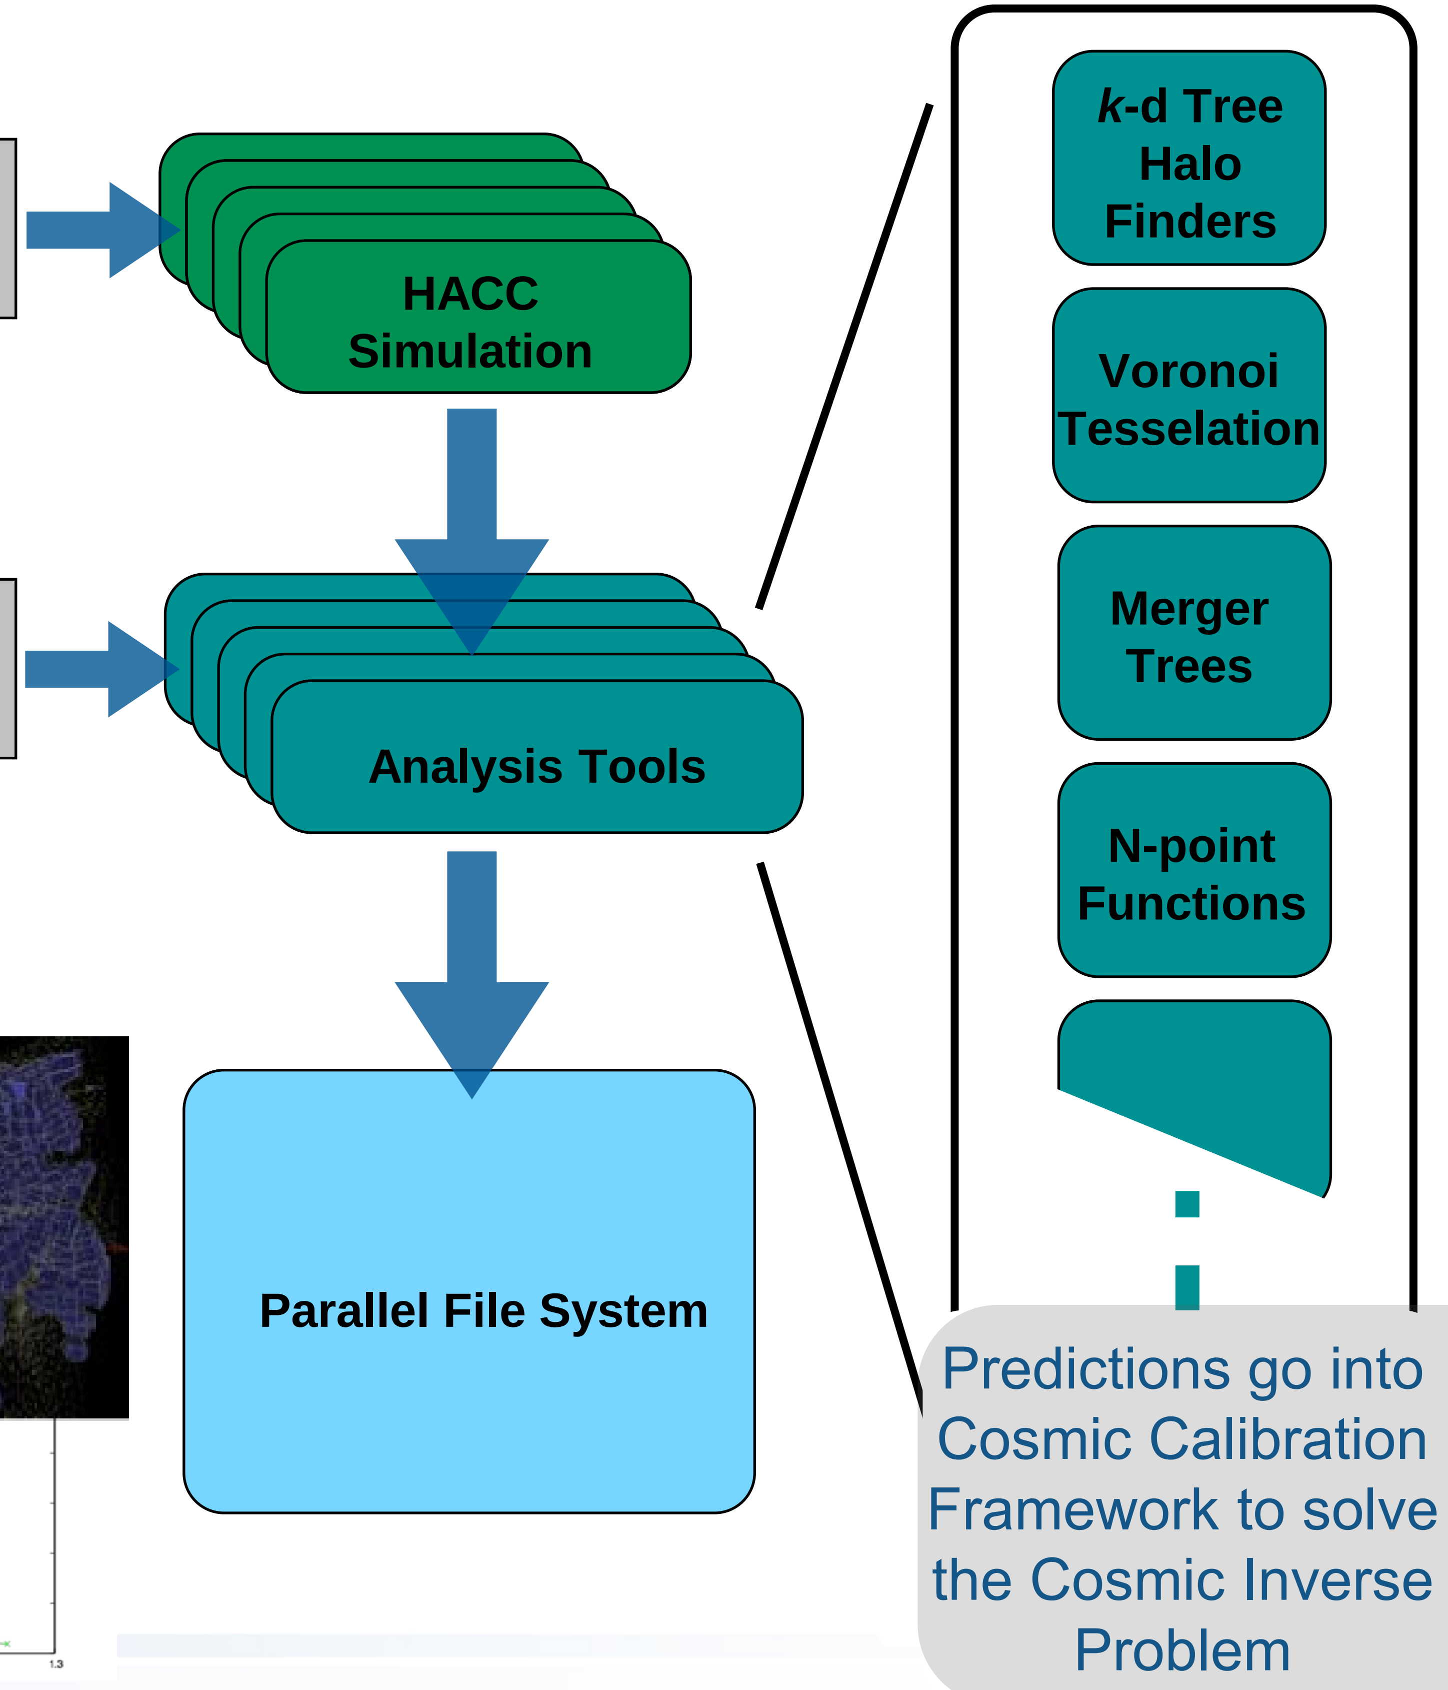

# Future Thoughts

- High Performance Computing ('PDEs')
  - Parallel systems with a fast network
  - Designed to run tightly coupled jobs
  - High performance parallel file system
  - Batch processing
- Data-Intensive Computing ('Analytics')
  - Parallel systems with balanced I/O
  - Designed for data analytics
  - System level storage model
  - Interactive processing
- The future of HPC is not 'HPC'!
  - HPC systems were meant to be balanced under certain metrics
  - These range from ~0.1 to ~0.001 on the same system and will get worse
  - A question of $$, not technology
- Data analysis is a major problem
  - When will analytics become truly interactive?