Computation-Driven Discovery for the Dark Universe

SDAV

FASTMATH

QUANTIFICATION OF UNCERTAINTY IN EXTREME SCALE COMPUTATIONS
QUEST

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)

Argonne NATIONAL LABORATORY
BERKELEY LAB
BROOKHAVEN NATIONAL LABORATORY
Fermilab
SLAC NATIONAL ACCELERATOR LABORATORY

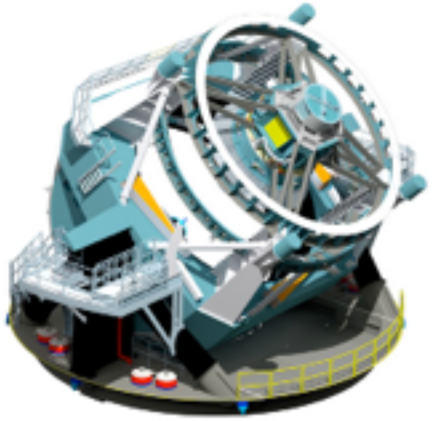

ASCR
HEP

ASCR Leadership Computing Challenge
ALCC

U.S. DEPARTMENT OF ENERGY
INCITE
LEADERSHIP COMPUTING

SciDAC
Scientific Discovery through Advanced Computing

Science at the Cosmic Frontier: ‘Precision’ Cosmology

- **Instrumentation Advances:** Wide/deep sky coverage, multi-wavelength observations (near-term target, ~1% errors on cosmological parameters or better)
- **Precision Cosmology Science:**
  - Nature of cosmic acceleration (physics of structure formation)
  - Nature and interactions of dark matter
  - Primordial fluctuations and Early universe
  - Probes of fundamental physics (e.g., neutrino sector)

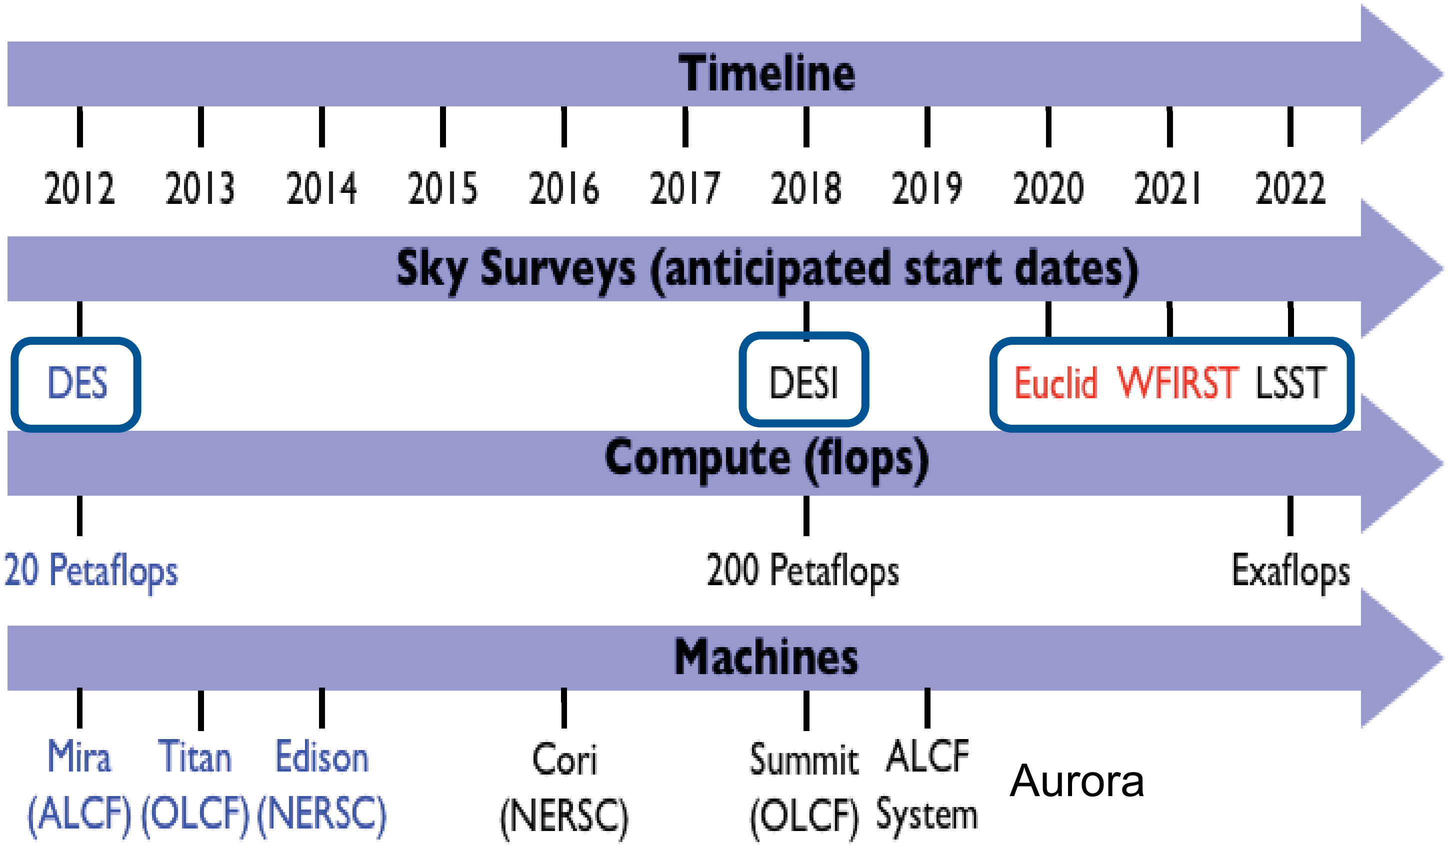

Timeline

2012 2013 2014 2015 2016 2017 2018 2019 2020 2021 2022

Sky Surveys (anticipated start dates)

DES

DESI

Euclid WFIRST LSST

Compute (flops)

20 Petaflops

200 Petaflops

Exaflops

Machines

Mira (ALCF)

Titan (OLCF)

Edison (NERSC)

Cori (NERSC)

Summit (OLCF)

ALCF System

Aurora

- **Theory/Modeling:** Predictive theory and modeling becoming critically important
- **Computational Cosmology:** Primary theoretical/modeling approach to bridge the the **“theory/modeling gap”**

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

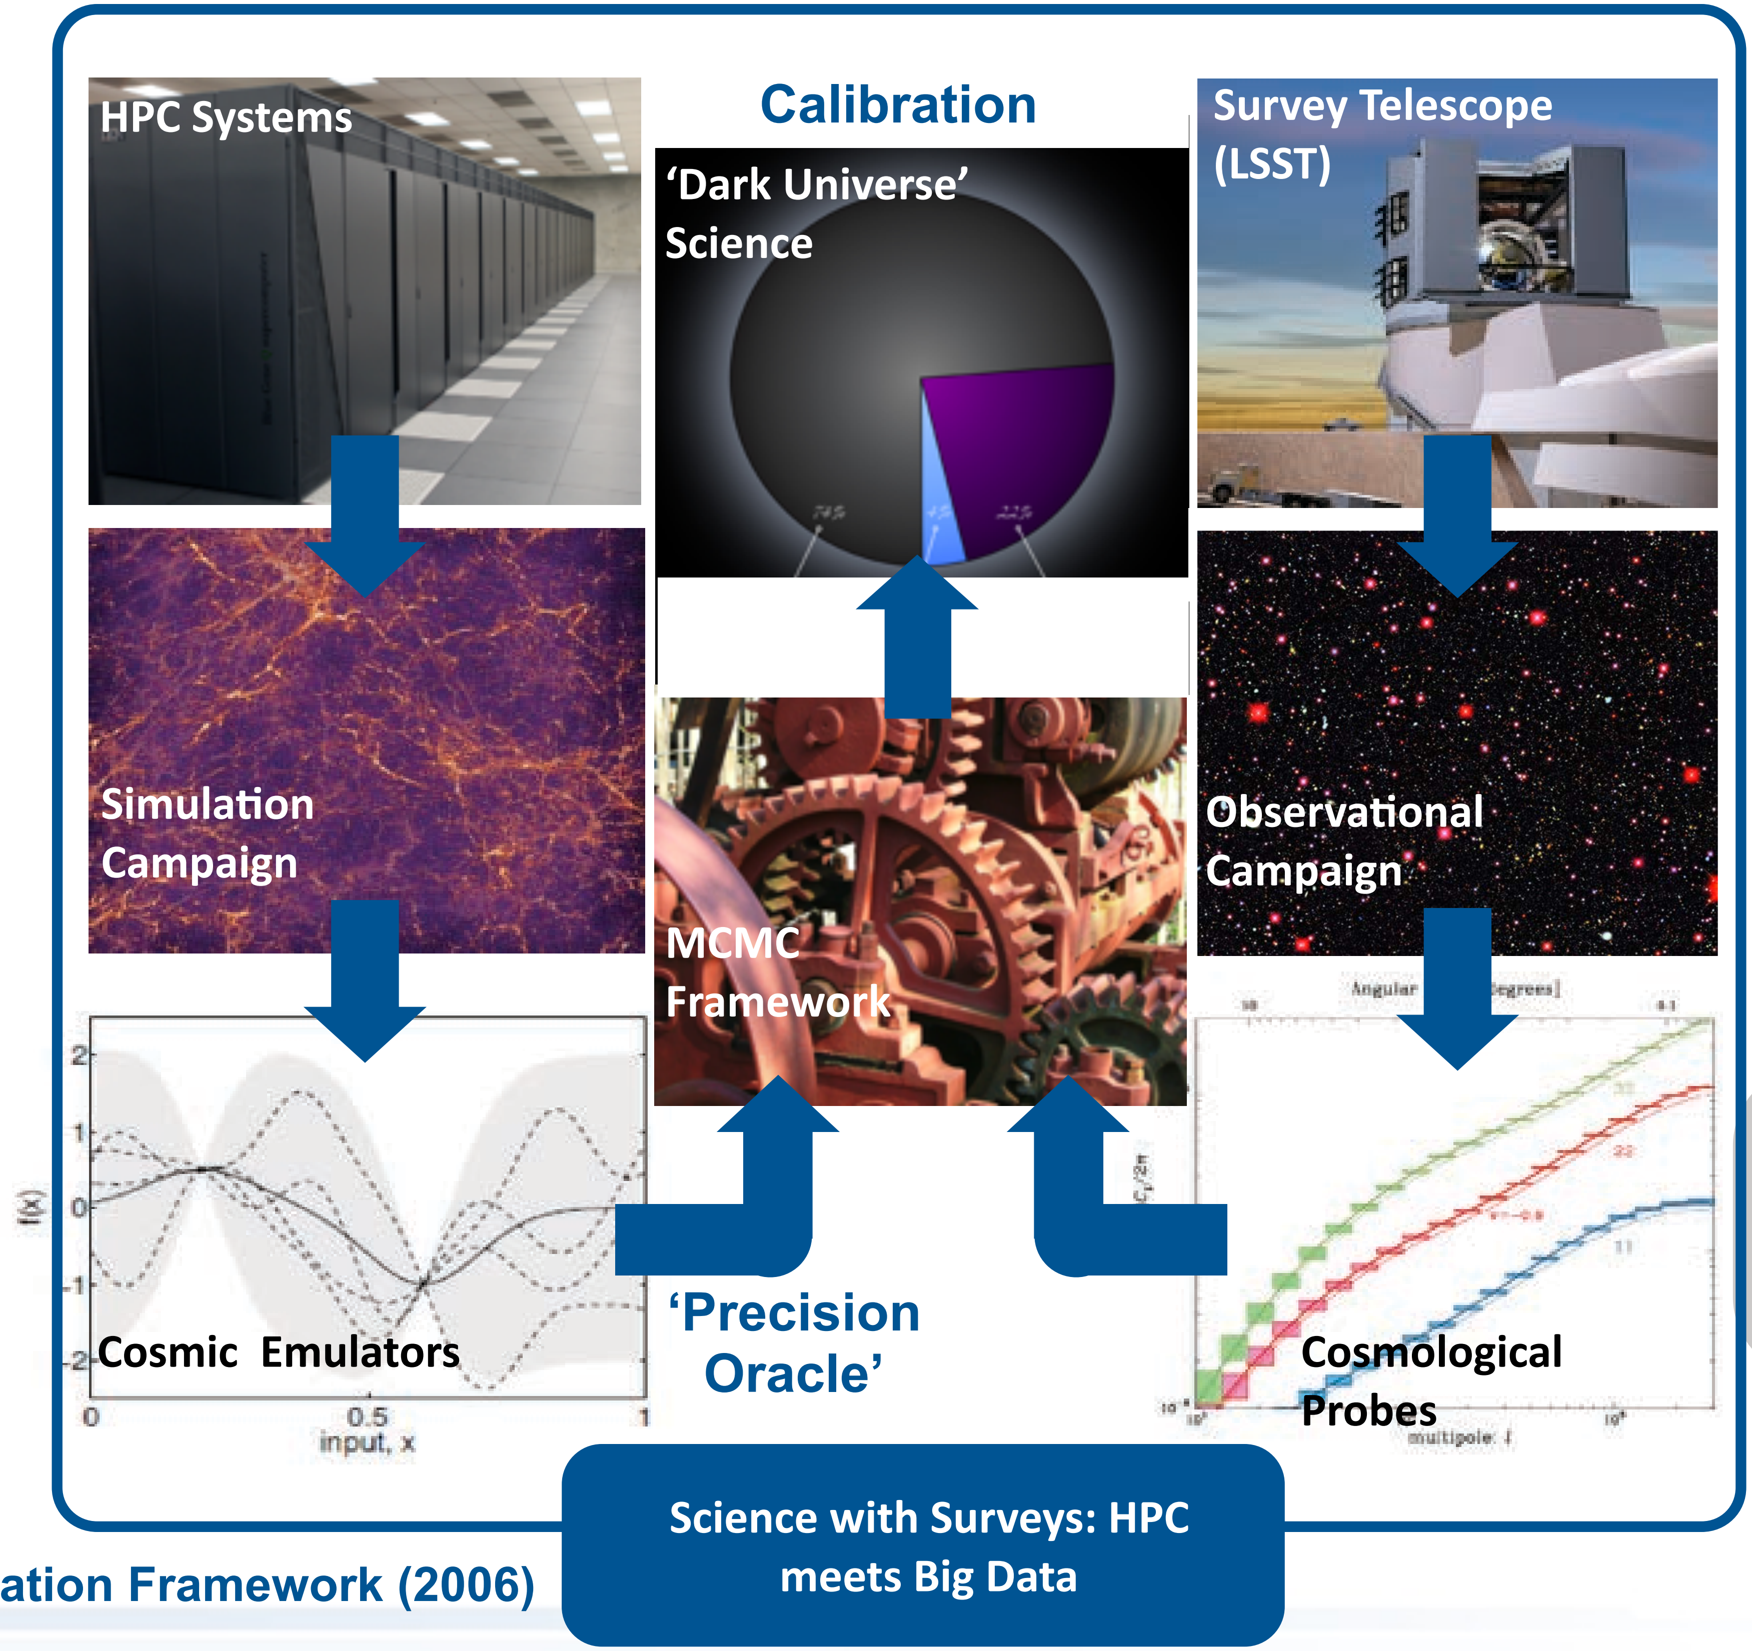

Science with Surveys: HPC meets Big Data

Mapping the Sky with Survey Instruments

Observations: Statistical error bars will ‘disappear’ soon!

CCF= Cosmic Calibration Framework (2006)

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

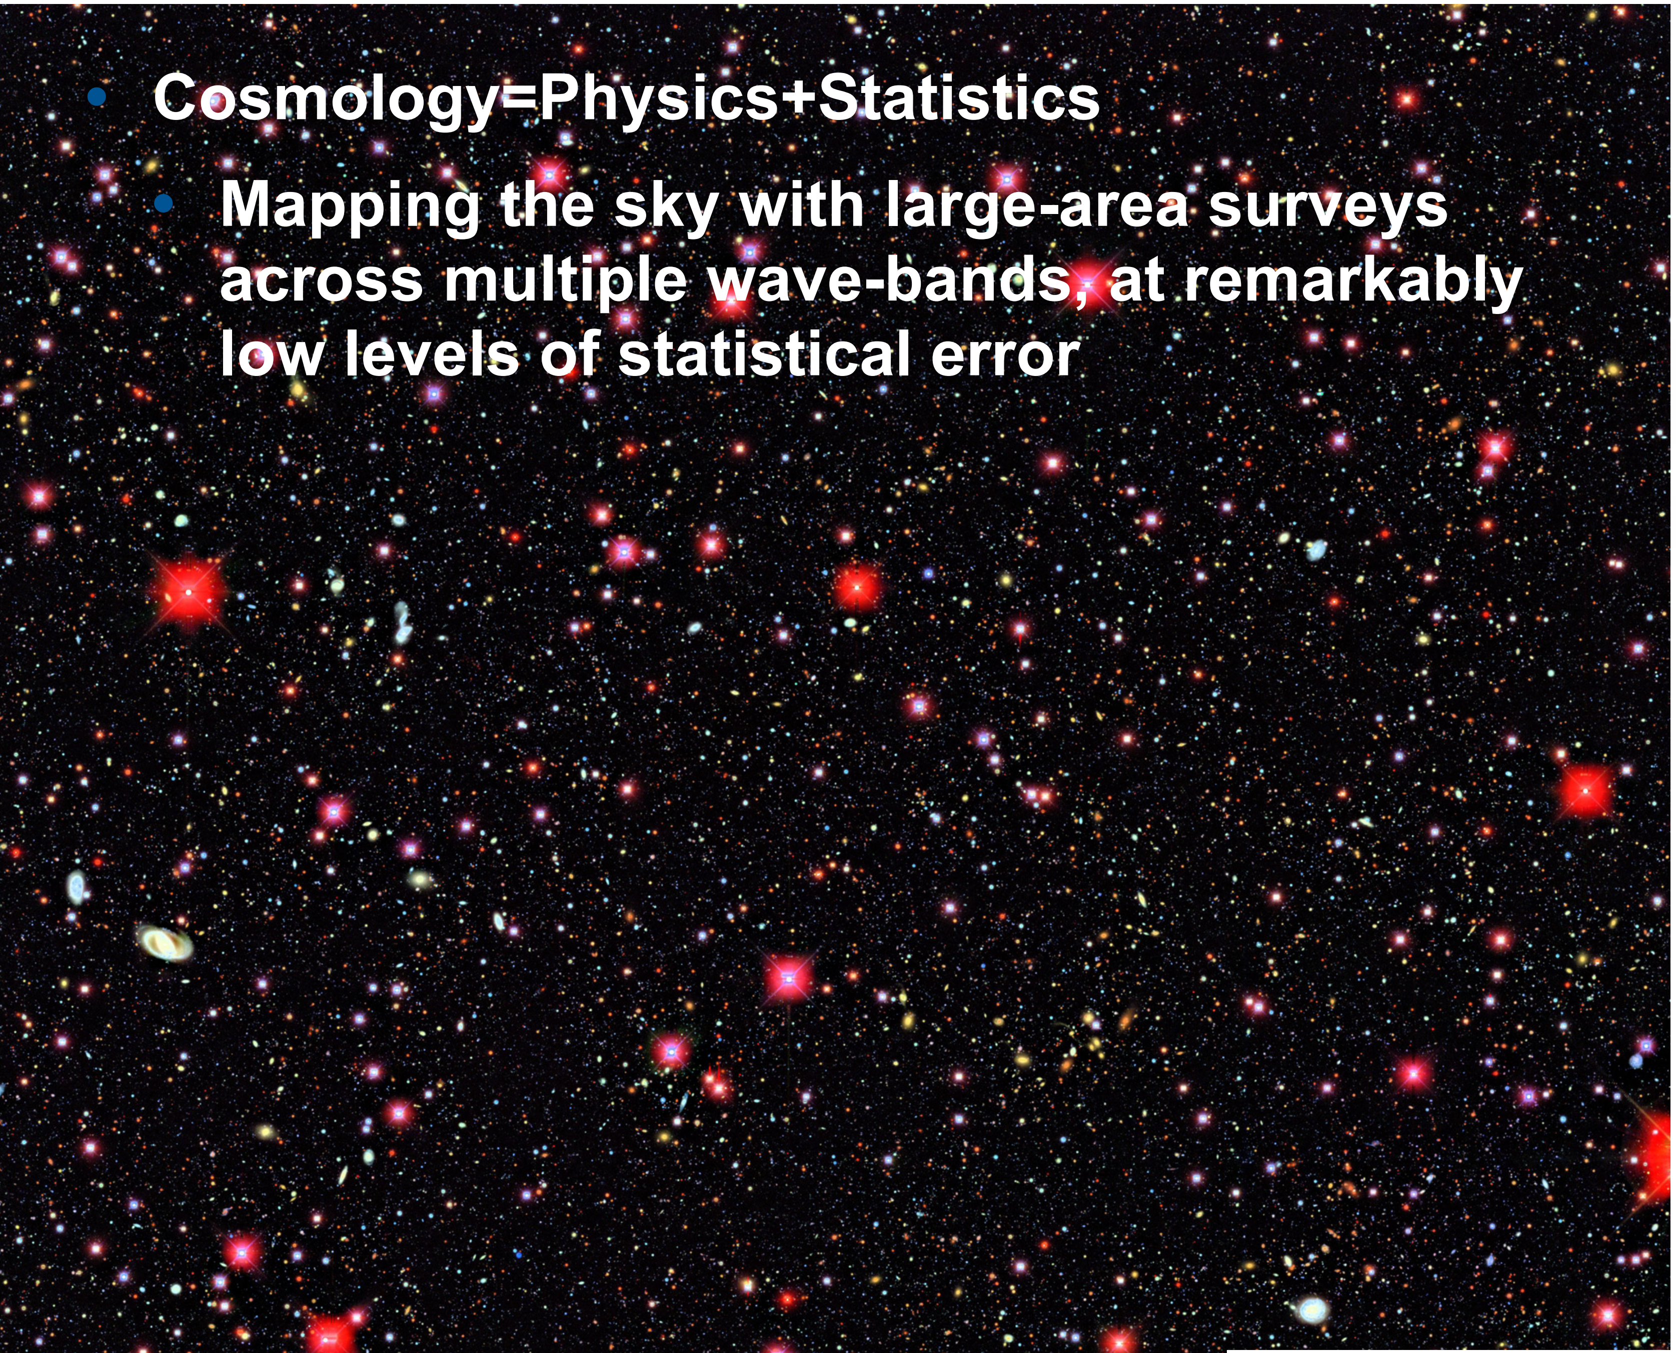

Galaxies in a moon-sized patch (Deep Lens Survey). LSST will cover 50,000 times this size (~400PB of data)

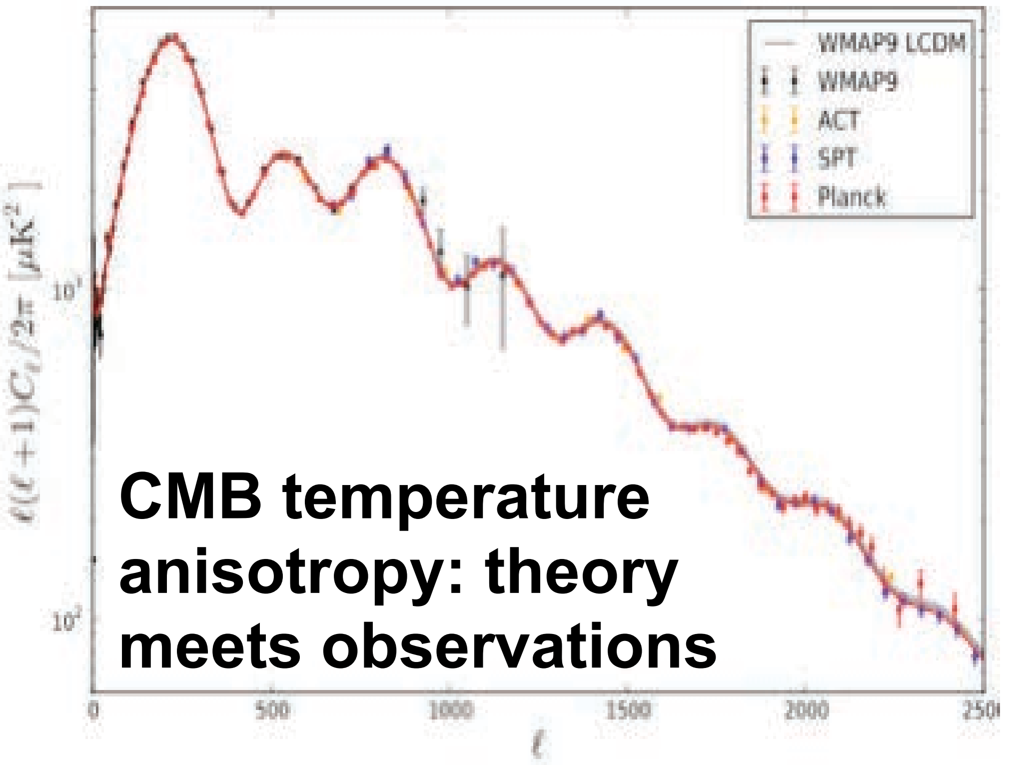

CMB temperature anisotropy: theory meets observations

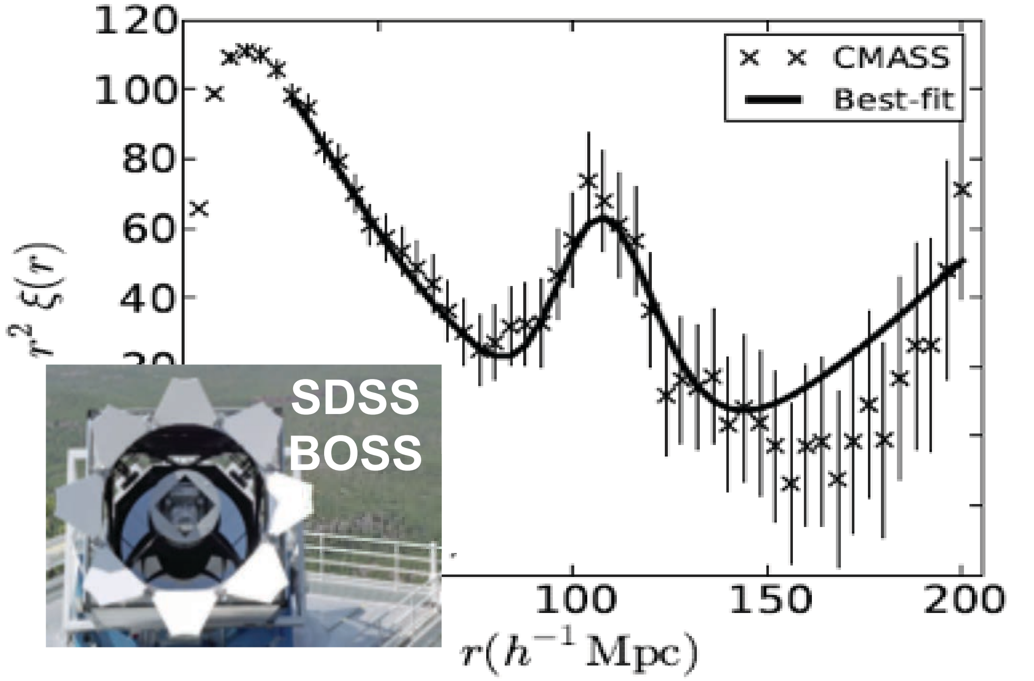

The same signal in the galaxy distribution

The Precision Cosmology Revolution

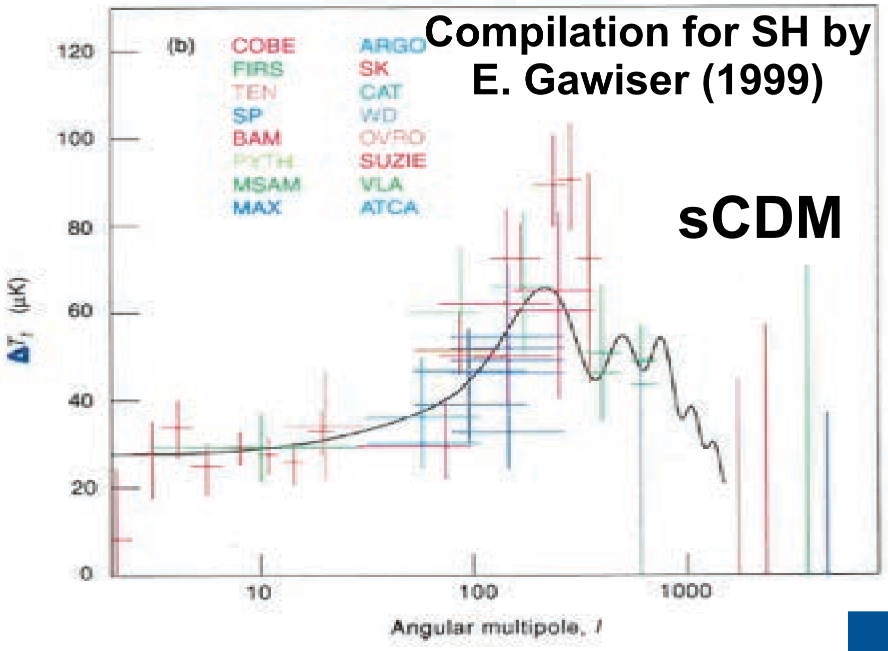

Compilation for SH by E. Gawiser (1999)

sCDM

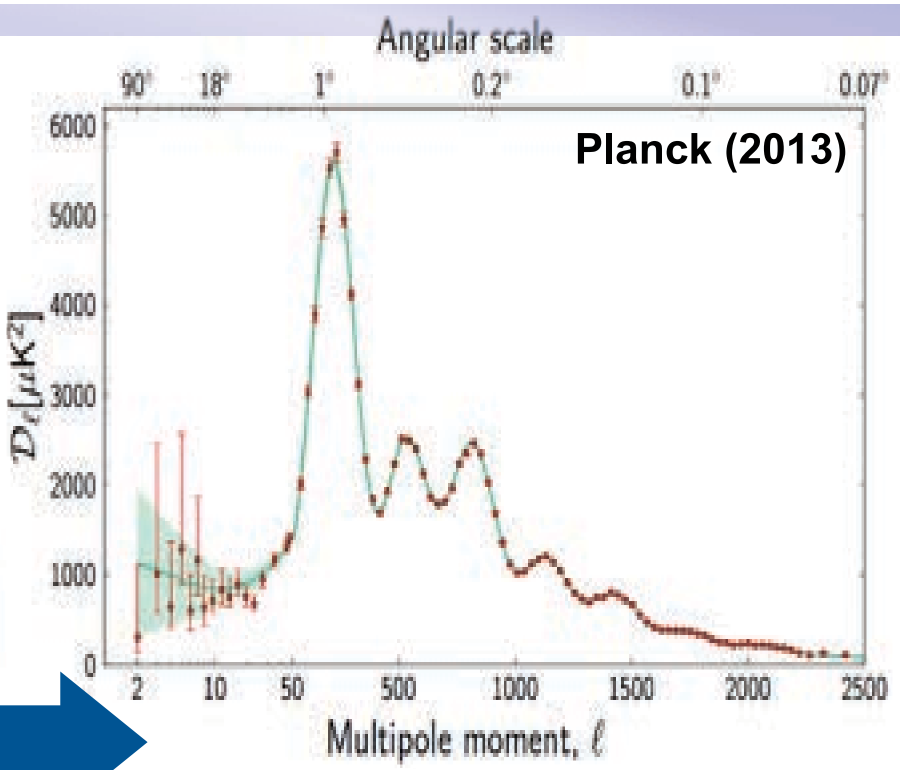

Planck (2013)

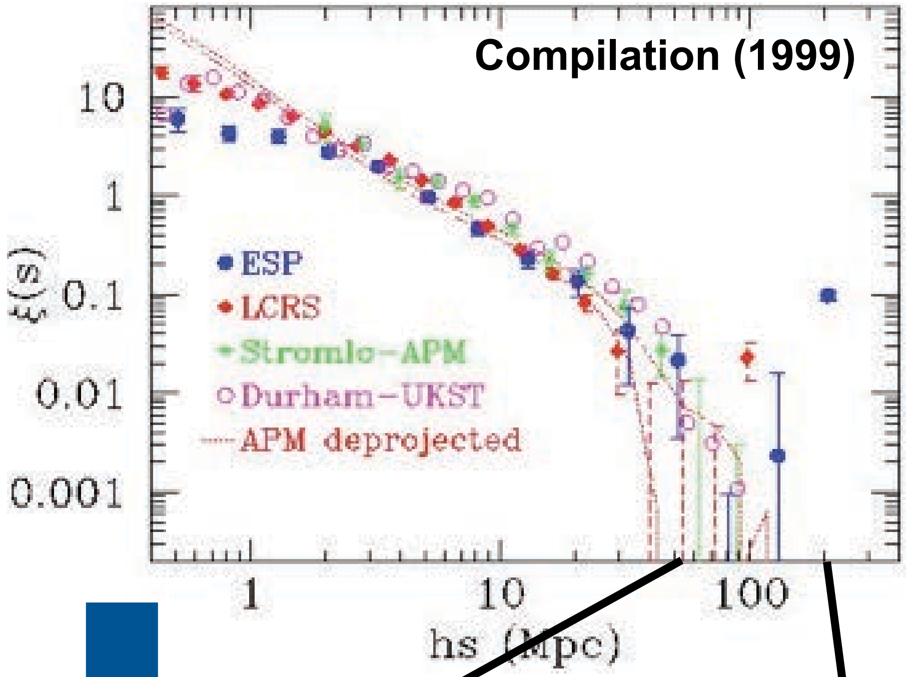

Compilation (1999)

CMB

LSS

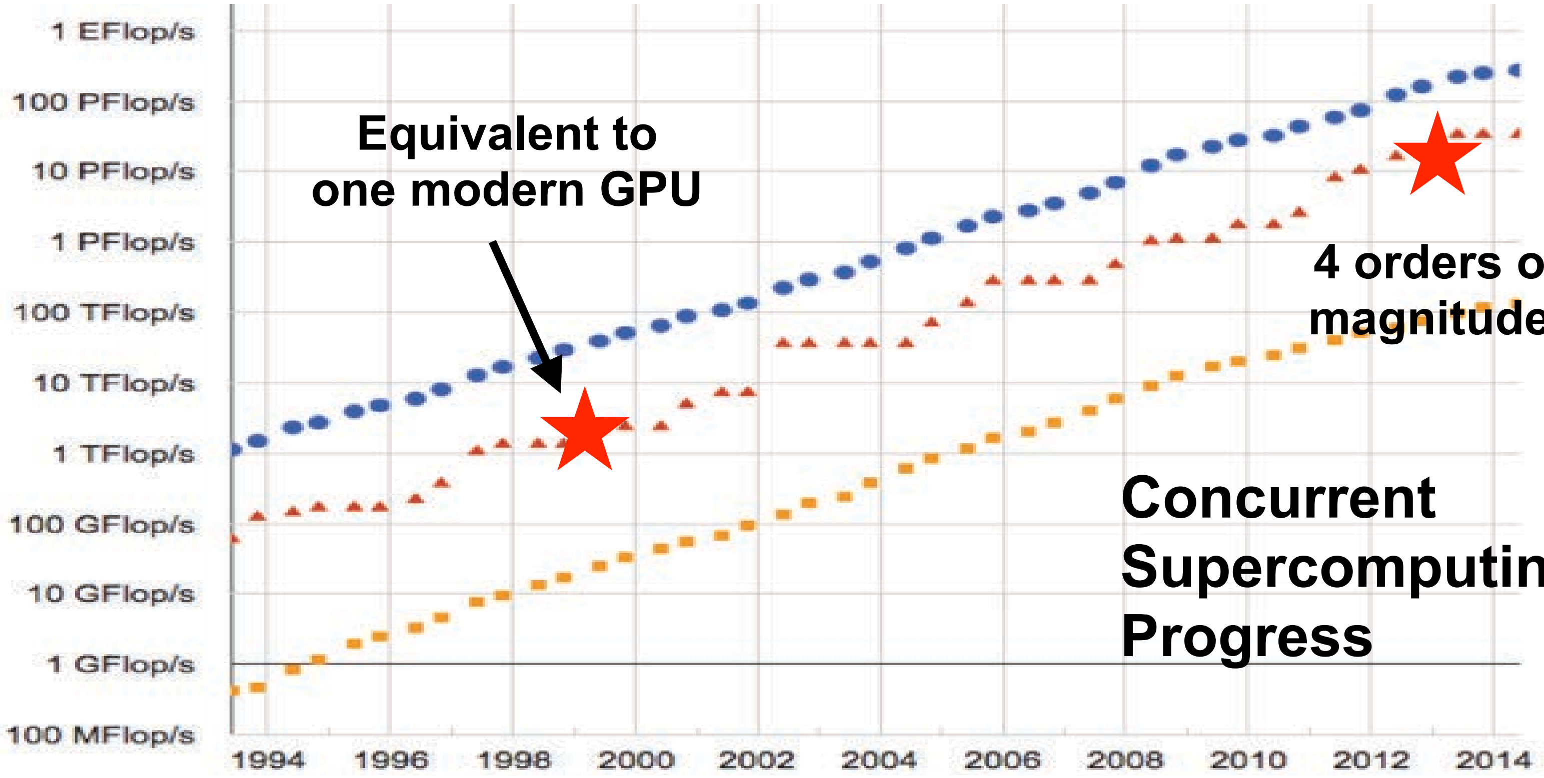

Equivalent to one modern GPU

4 orders of magnitude!

Concurrent Supercomputing Progress

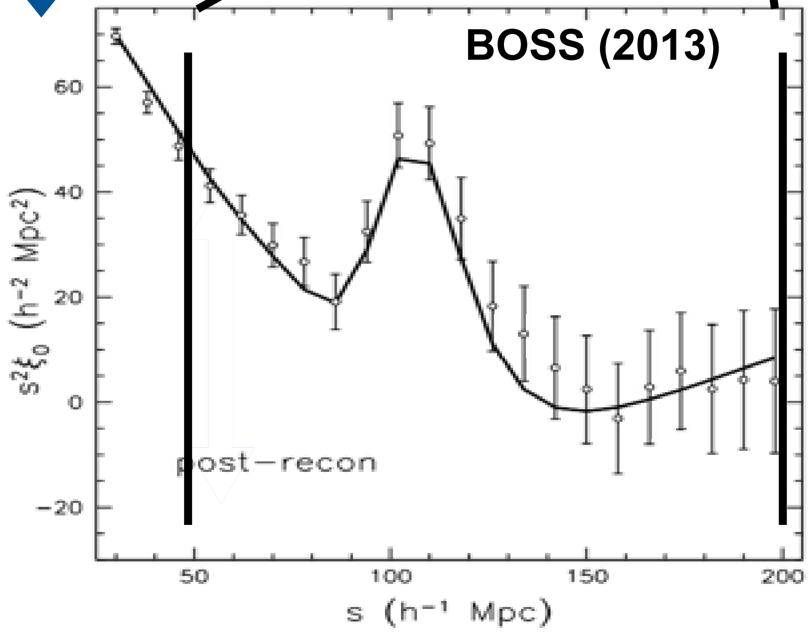

BOSS (2013)

Connecting Theory and Observations: Challenges & Opportunities

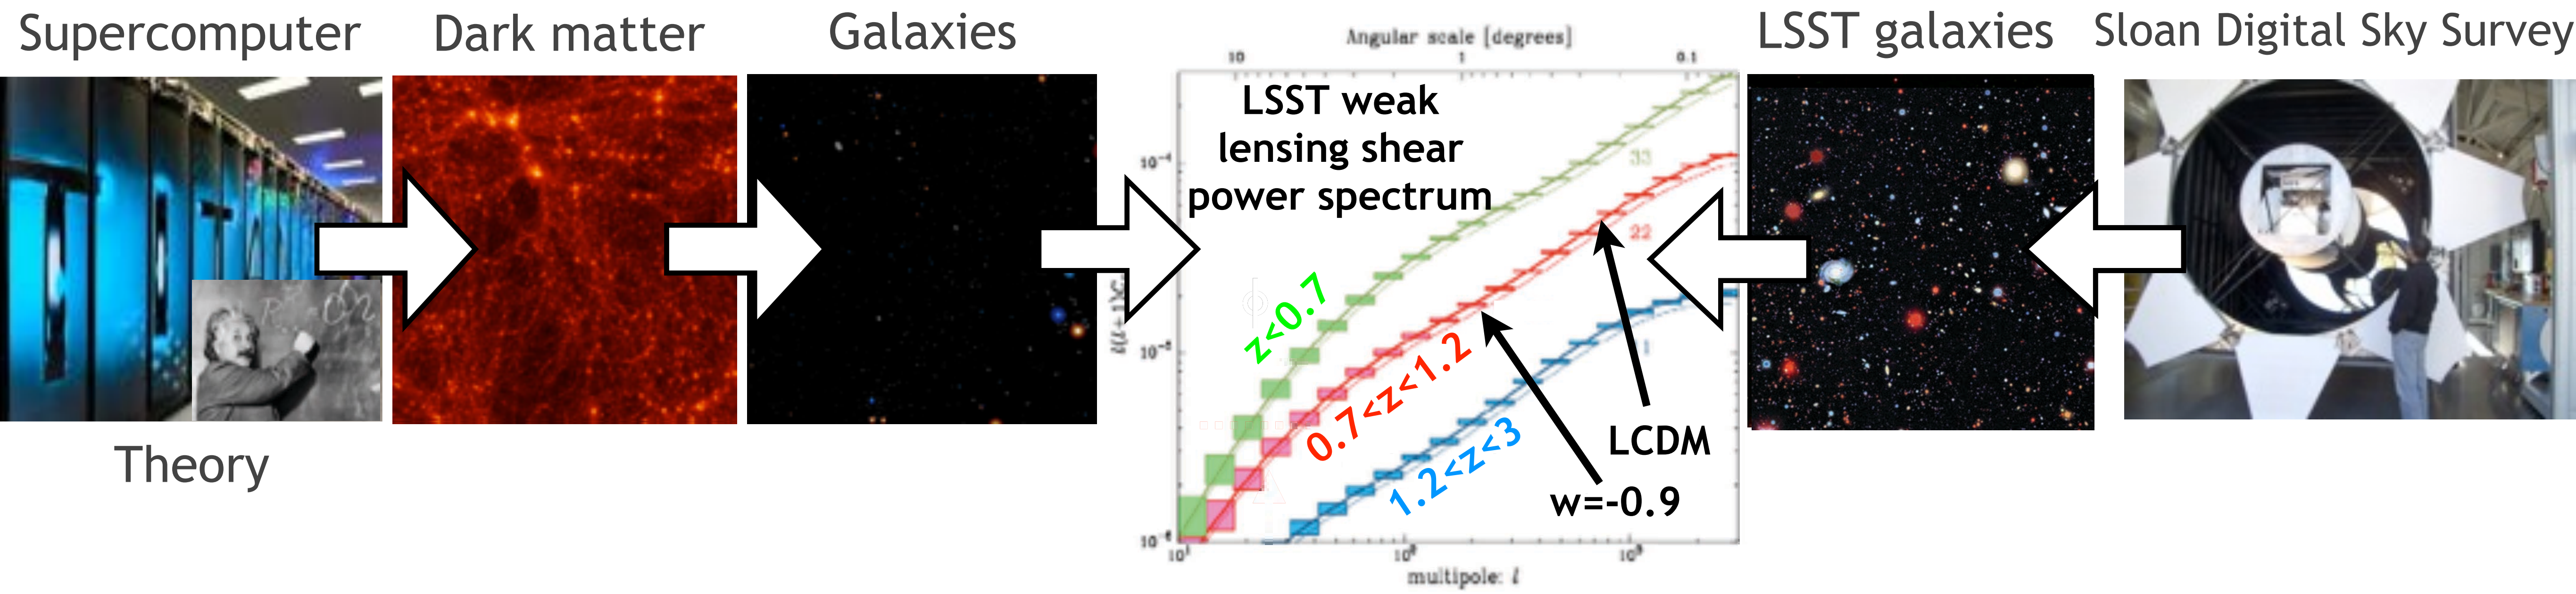

Supercomputer

Dark matter

Galaxies

LSST weak lensing shear power spectrum

LSST galaxies

Sloan Digital Sky Survey

Theory

- **Error bars are shrinking dramatically**
  - Many predictions have to be accurate at the sub-percent level
  - Modeling and understanding of systematics is becoming 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 are going deeper and will target/resolve fainter/different galaxies**
  - Synthetic sky map making becomes more difficult, more physics needed
  - Significantly higher resolution simulations will be required
  - New cosmological probes, cross-correlations across multiple probes

Example: Analytics/Workflow Complexity

Gaussian Random Field Initial Conditions

High-Resolution N-Body Code (HACC)

Multiple Outputs Halo/Sub-Halo Identification

Halo Merger Trees

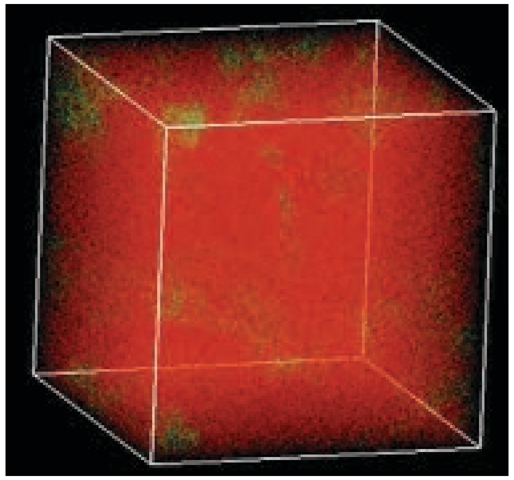

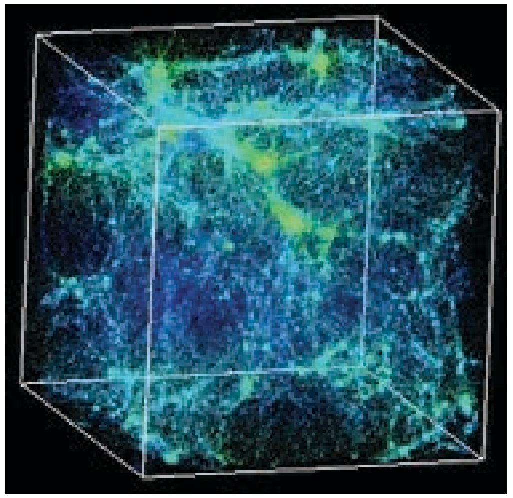

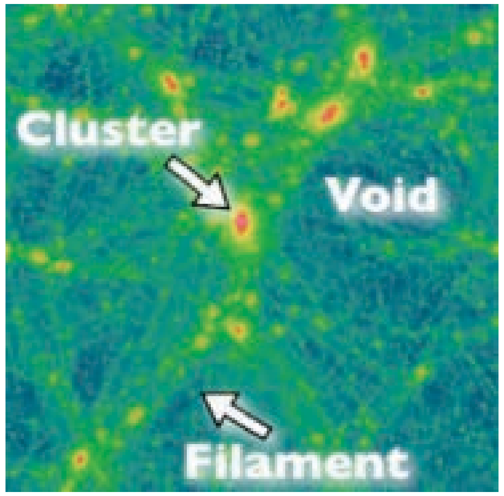

Semi-Analytic Modeling Code (Galacticus)

Galaxy Catalog

Realistic Image Catalog

Atmosphere and Instrument Modeling

Data Management Pipeline

Data Analysis Pipeline

Scientific Inference Framework

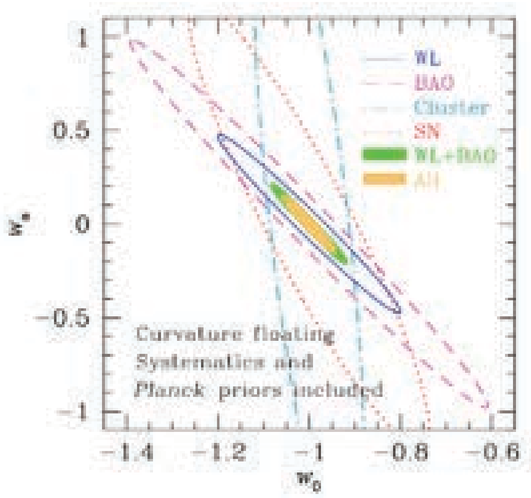

- **Simulation Campaigns:** Statistics of virtual universes; construction of emulators
- **Modeling the Measurement:** End-to-End modeling necessary to understand crucial systematic errors
- **PDACS:** Custom workflow system under development
- **Data-Intensive Computing:** New architectures needed

Large Scale Structure: Vlasov-Poisson Equation

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

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

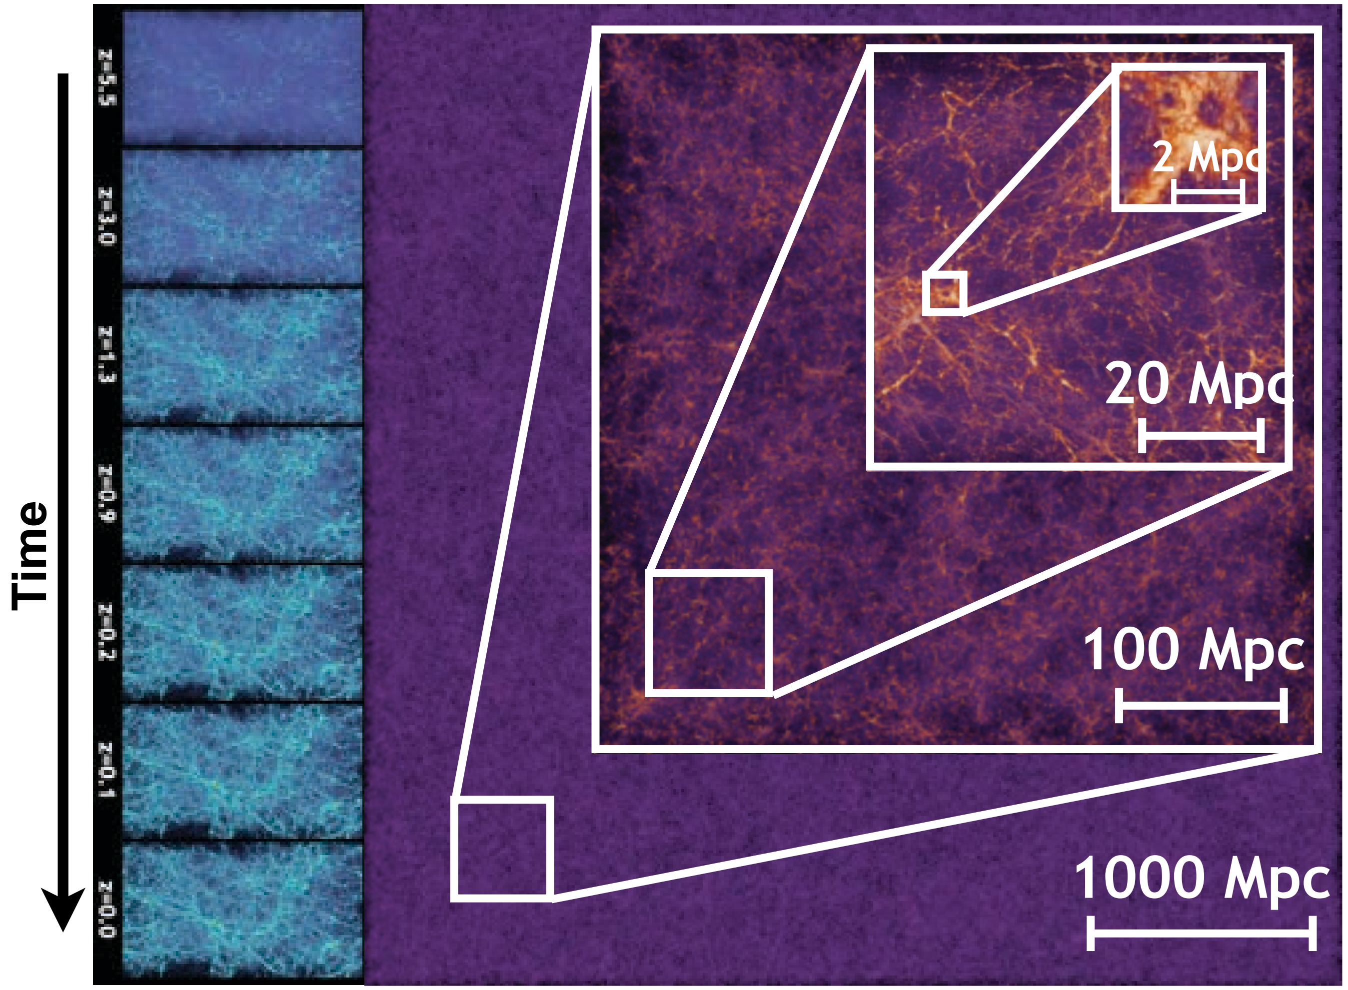

Gravitational Jeans Instability: 'Outer Rim' run with 1.1 trillion particles

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
- **Performance:** First production science code to break 10PFlops sustained, runs at full scale on all current DOE supercomputing systems (Gordon Bell Finalist 2012/2013, benchmark code for CORAL procurement)
- **Load Balancing:** New, highly efficient, task-based load balancing implemented on Titan

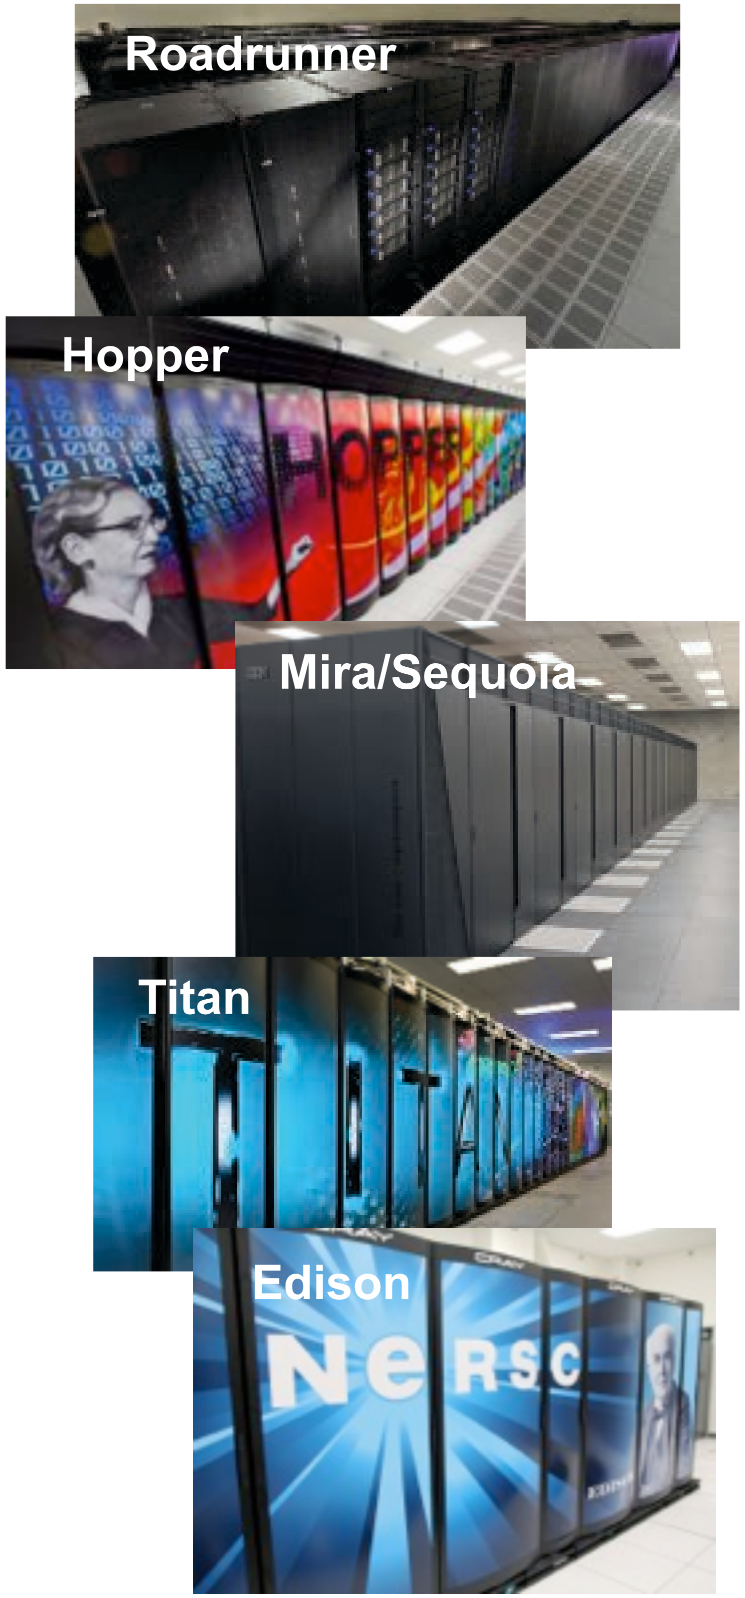

S. Habib et al., arXiv:1410.2805 (in press)

The Q Continuum and the Outer Rim Simulations

Simulating the LCDM Universe with Unprecedented Volume and Resolution

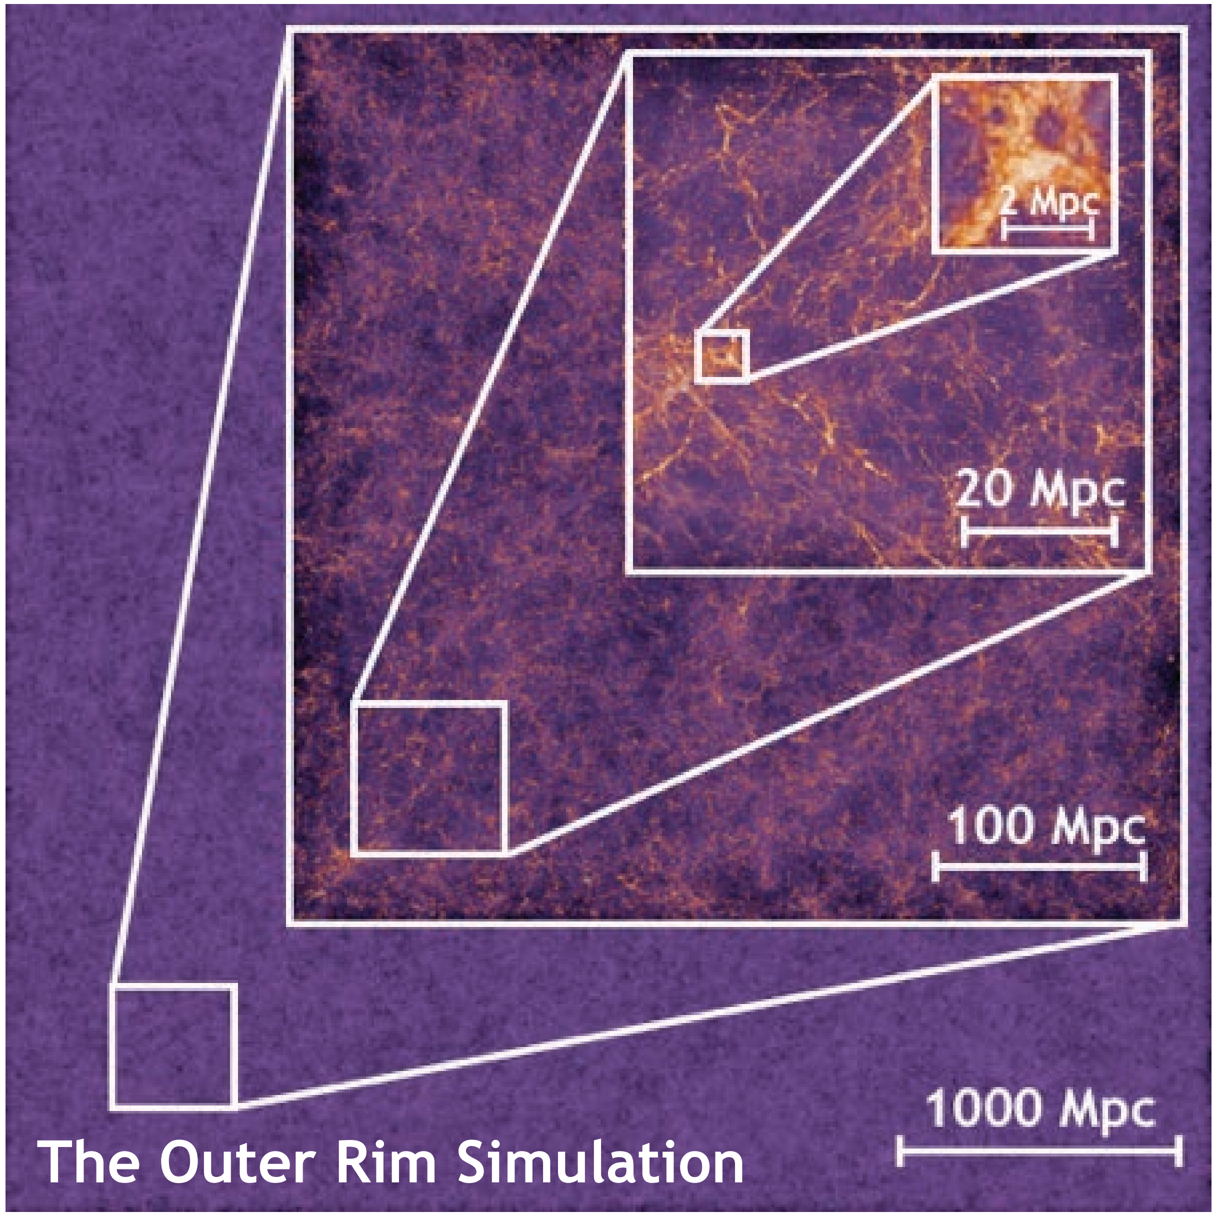

(4225 Mpc)^3 volume, 1.07 trillion particles carried out on ~67% of Mira at Argonne, 4PB of data, 216x Millennium simulation

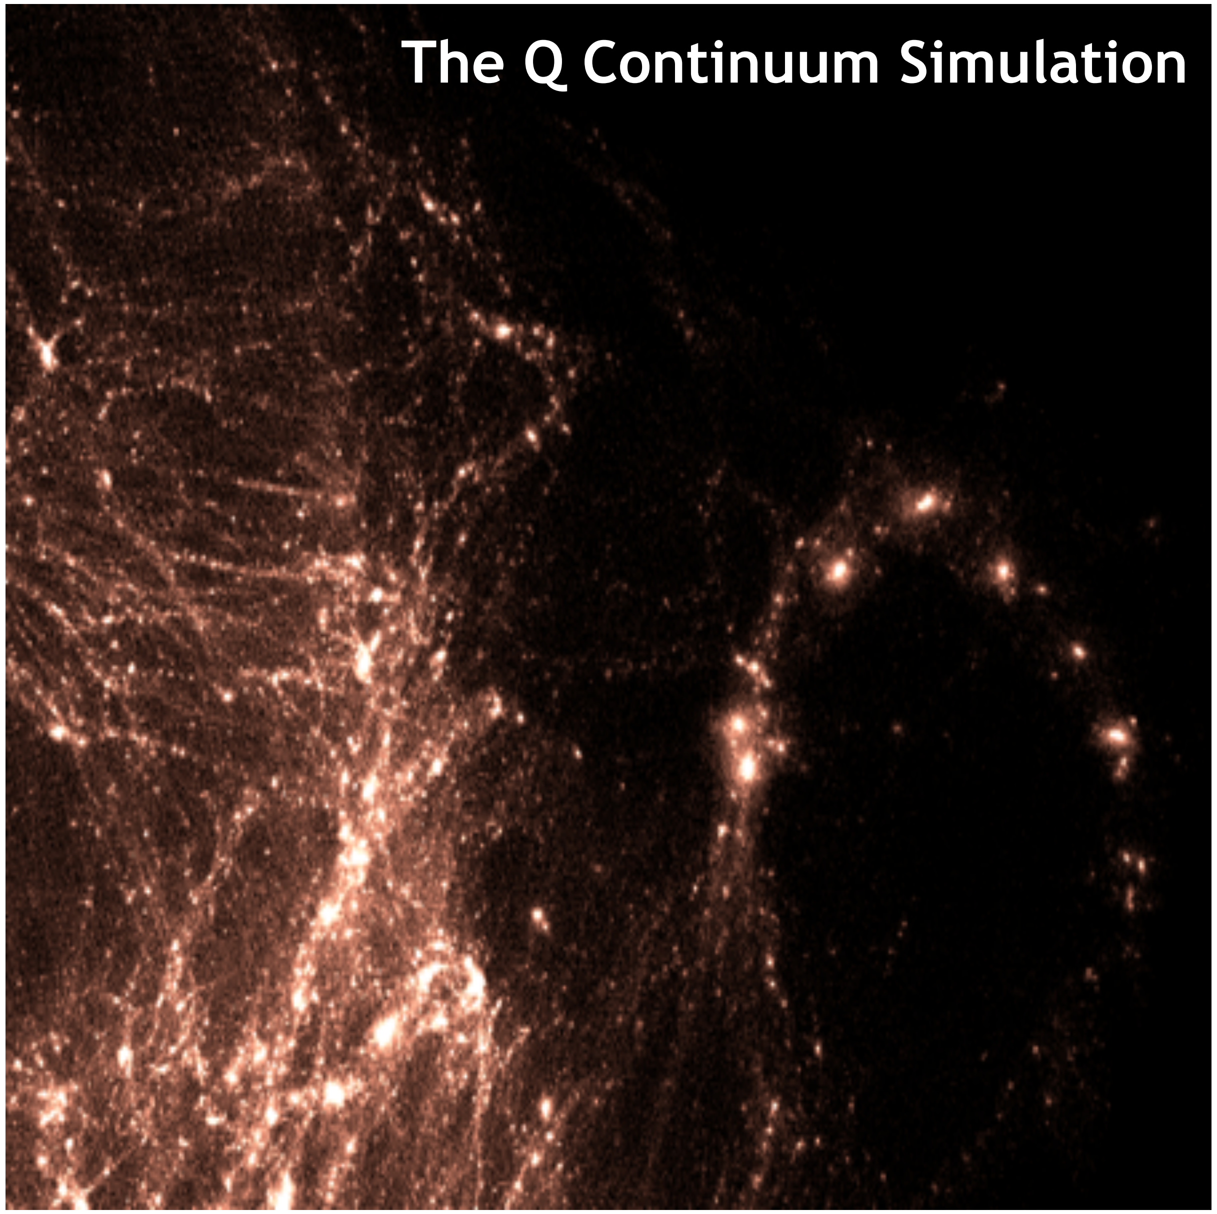

(1300 Mpc)^3 volume, 0.55 trillion particles carried out on ~90% of Titan at Oak Ridge, 2PB of data, 64x Bolshoi simulation

Cosmology with the Q Continuum Run

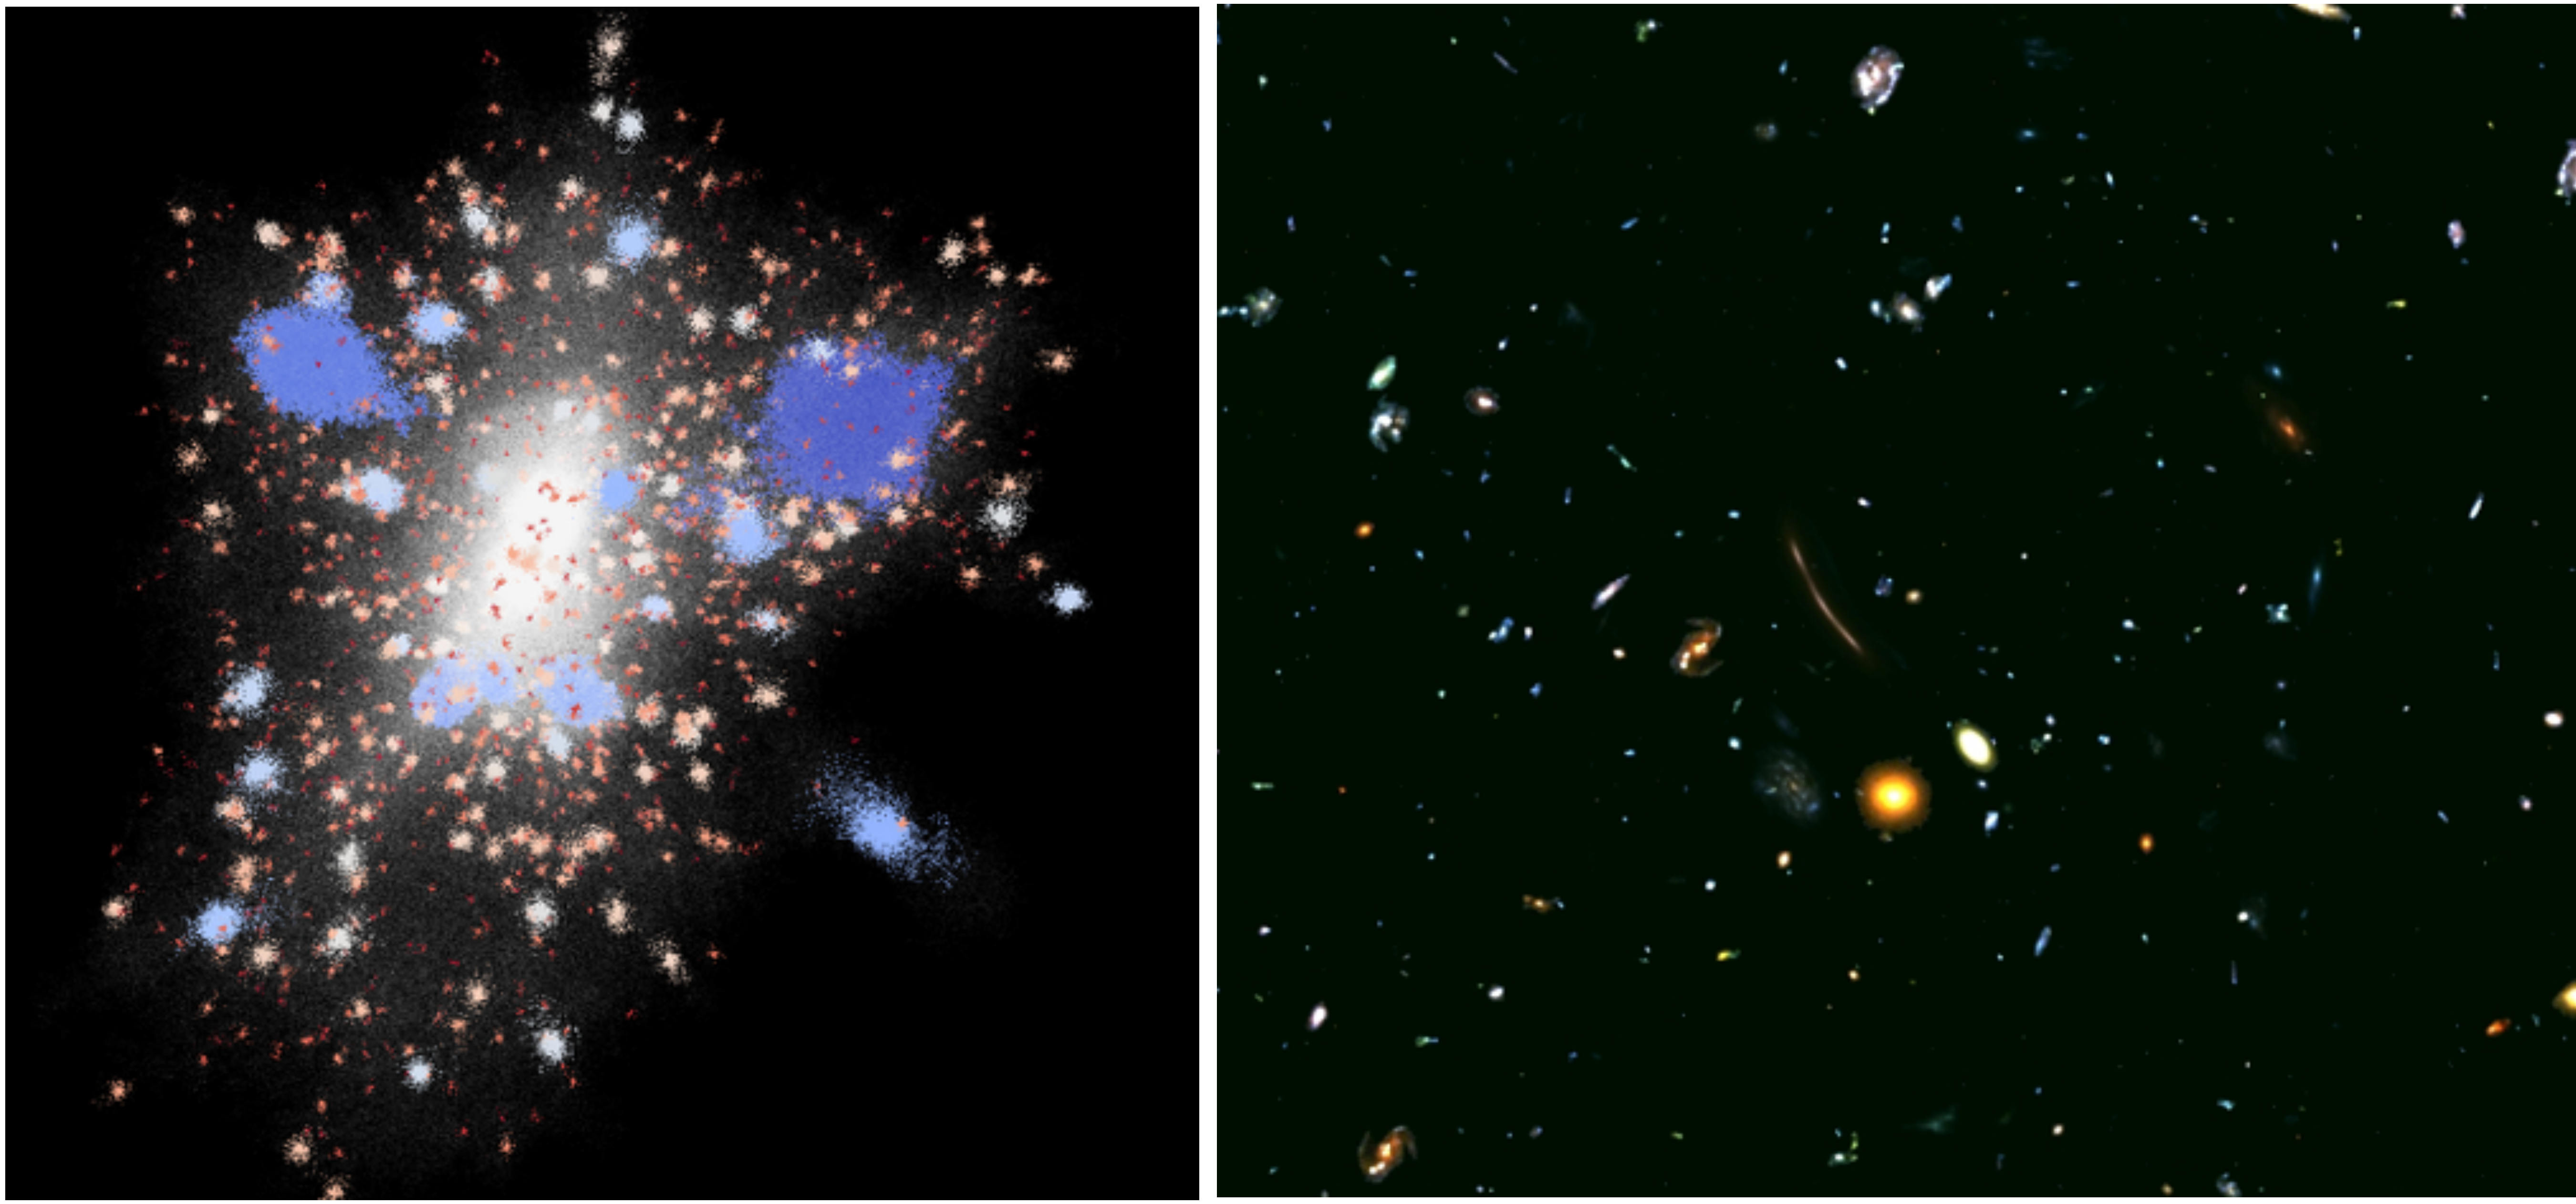

- Many cosmological statistics available at high resolution
- Highly resolved cluster-scale halos used for strong lensing predictions (left, halo with ~1400 subhalos, right, background galaxies lensed by a simulated cluster)

Heitmann et al. 2015 (in press)

Cosmology with HACC: Exquisite Statistics

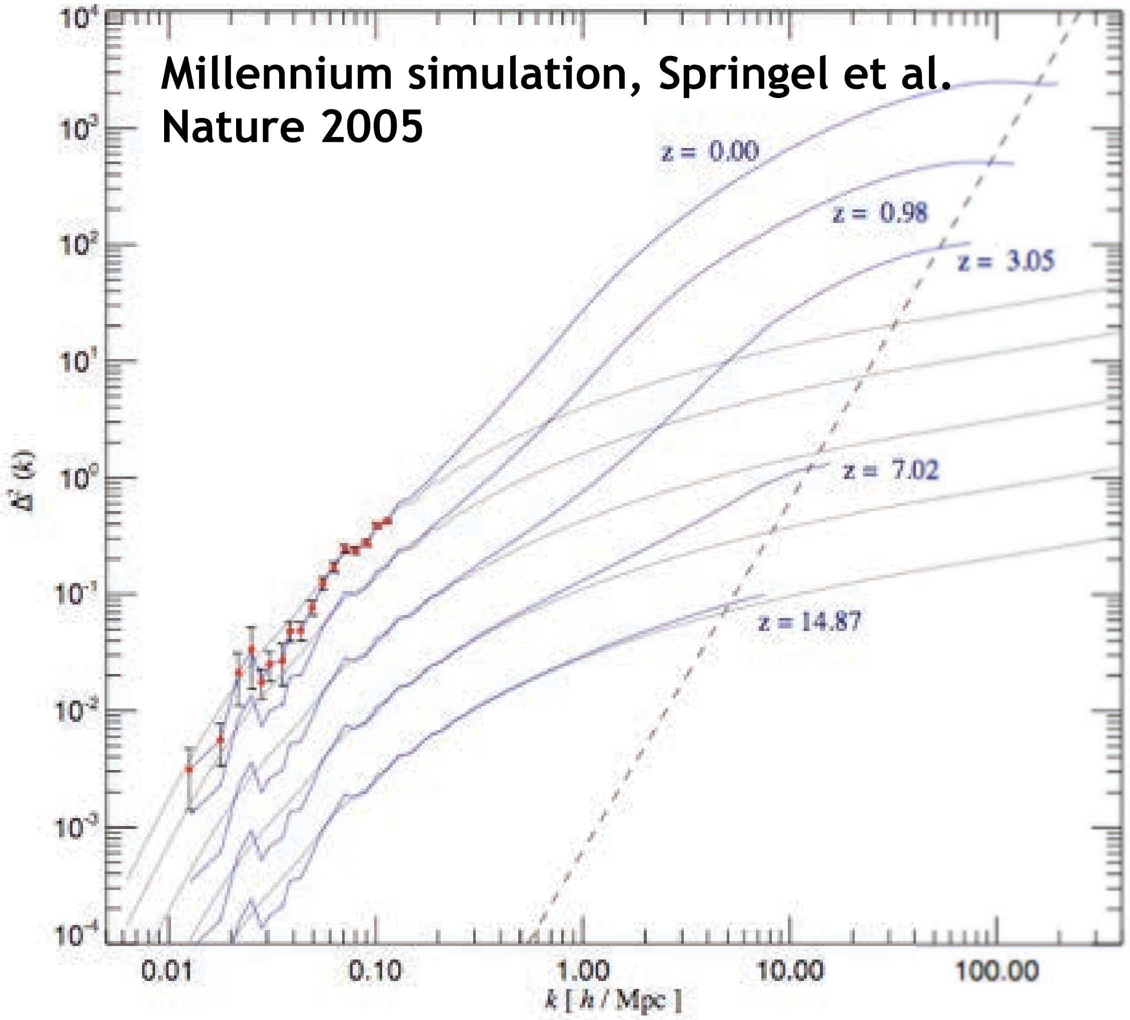

Millennium simulation, Springel et al. Nature 2005

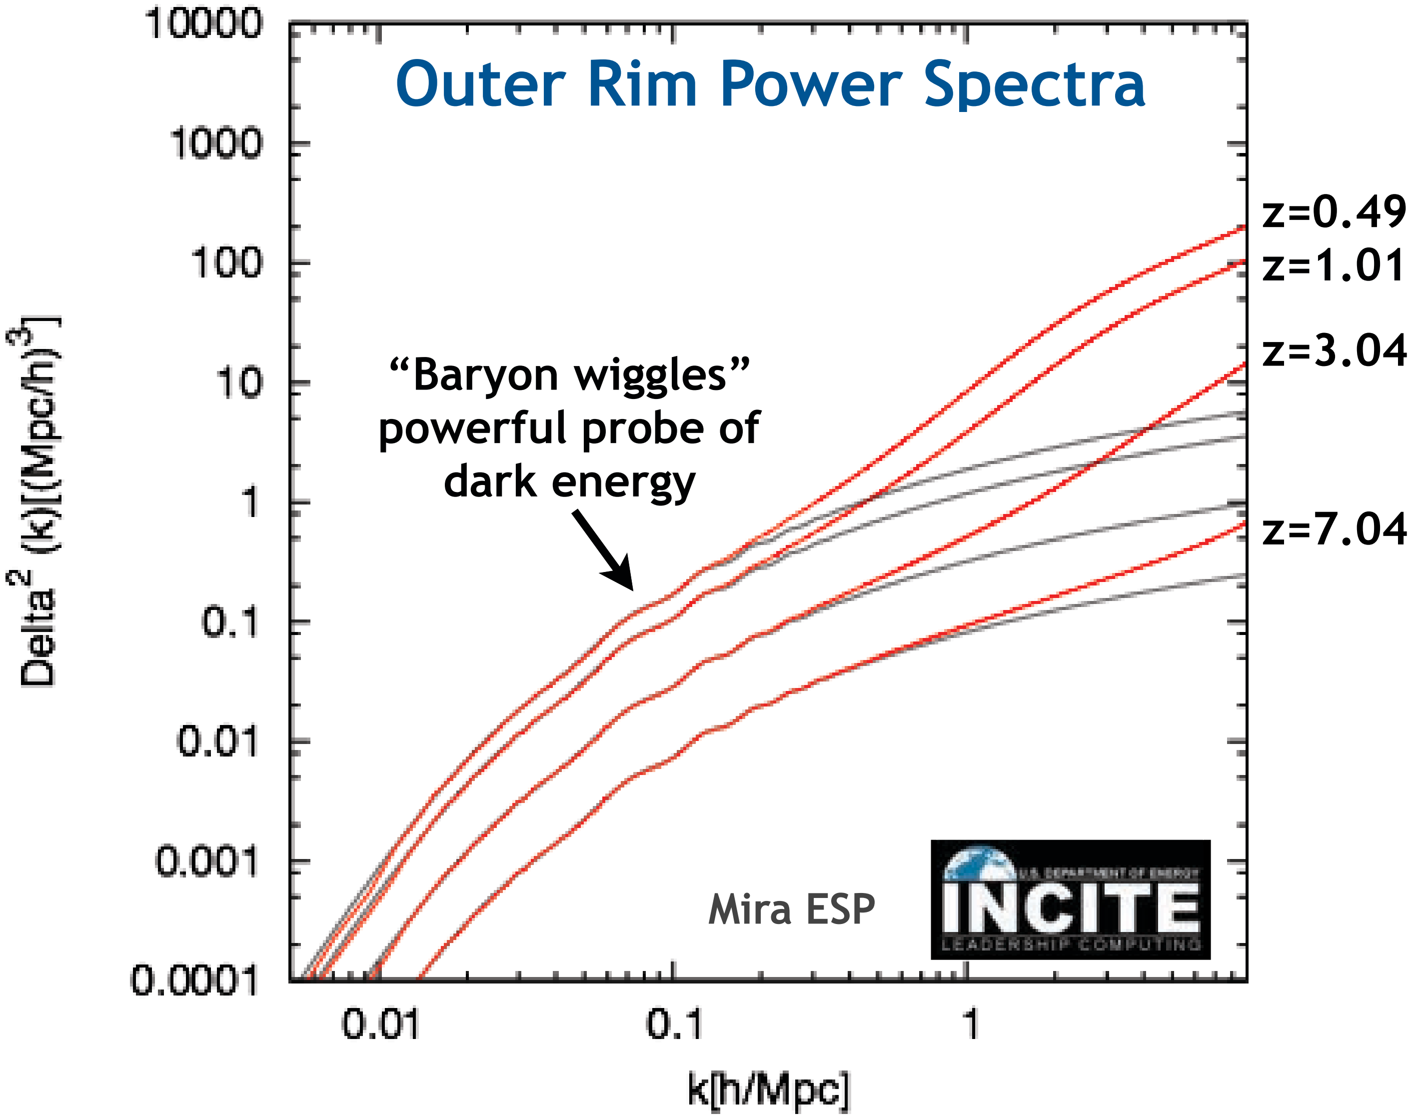

Outer Rim Power Spectra

"Baryon wiggles" powerful probe of dark energy

Mira ESP

- Mass resolution of Millennium simulation and Outer Rim run very similar (~ $10^9 \mathrm{M}_\odot$ 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

Habib et al. 2014 (in press)

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:** Hopeless –
- **Solution:** Emulators

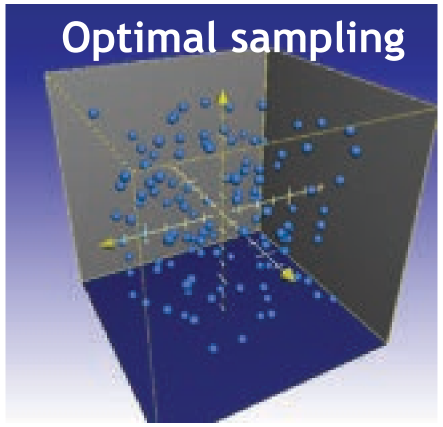
Optimal sampling


CosmicEmu publicly available

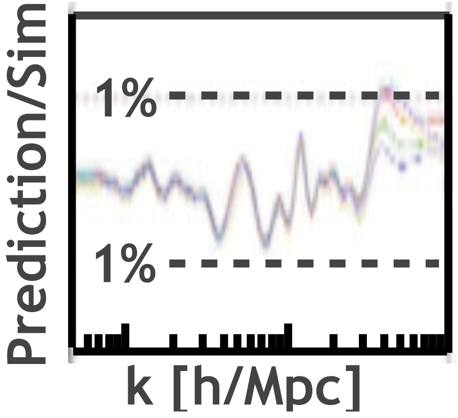

Heitmann et al. 2006, Habib et al. 2007

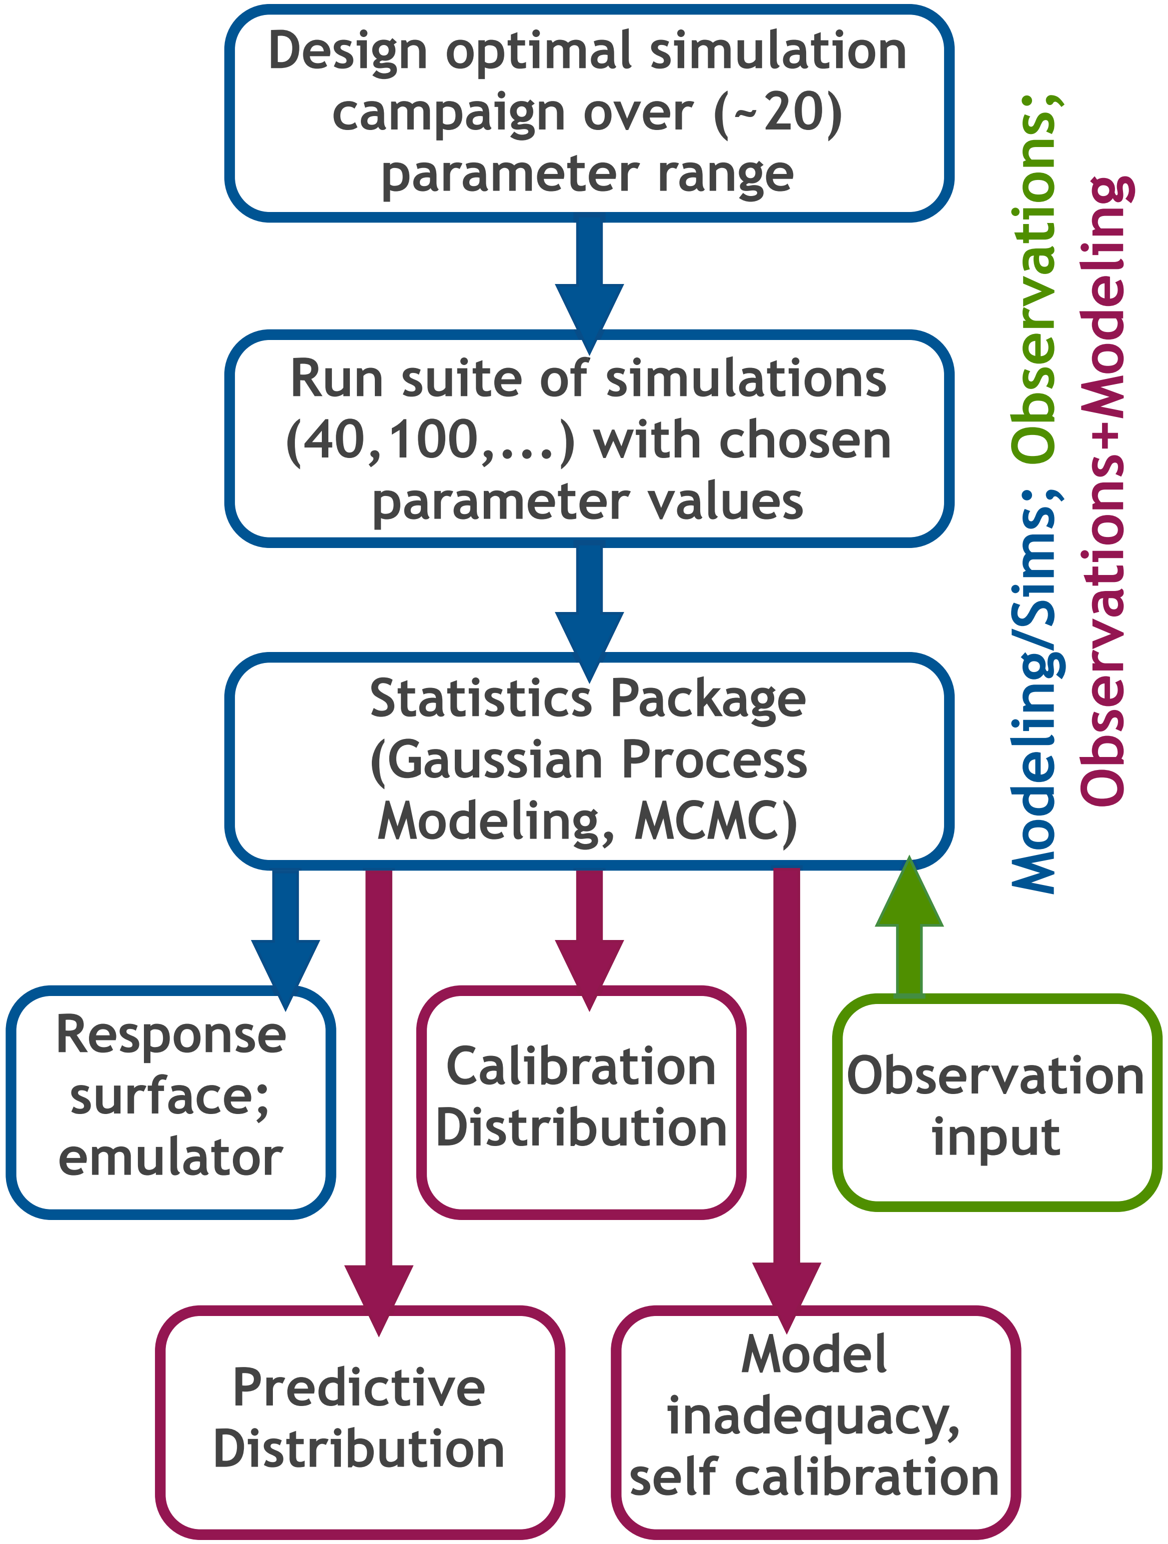
Design optimal simulation campaign over (~20) parameter range

Run suite of simulations (40,100,...) with chosen parameter values

Statistics Package (Gaussian Process Modeling, MCMC)

Response surface; emulator

Calibration Distribution

Observation input

Predictive Distribution

Model inadequacy, self calibration

Modeling/Sims; Observations; Observations+Modeling

Emulator Science

- **Previous/Current Results:**
  - Emulator for the matter power spectrum (current state of the art); Heitmann et al. 2014
  - Emulators for galaxy statistics, Kwan et al. 2015 (in press)
  - Emulators for halo profiles (c(M) relation), Kwan et al. 2013
  - Emulators for covariances, in prep.
- **Titan-Mira Universe Suite:**
  - Increase number of dimensions to 8, adding neutrinos and dynamical dark energy
  - Introduce new nested lattice sampling method for increase of accuracy; good accuracy with only 26 simulation runs

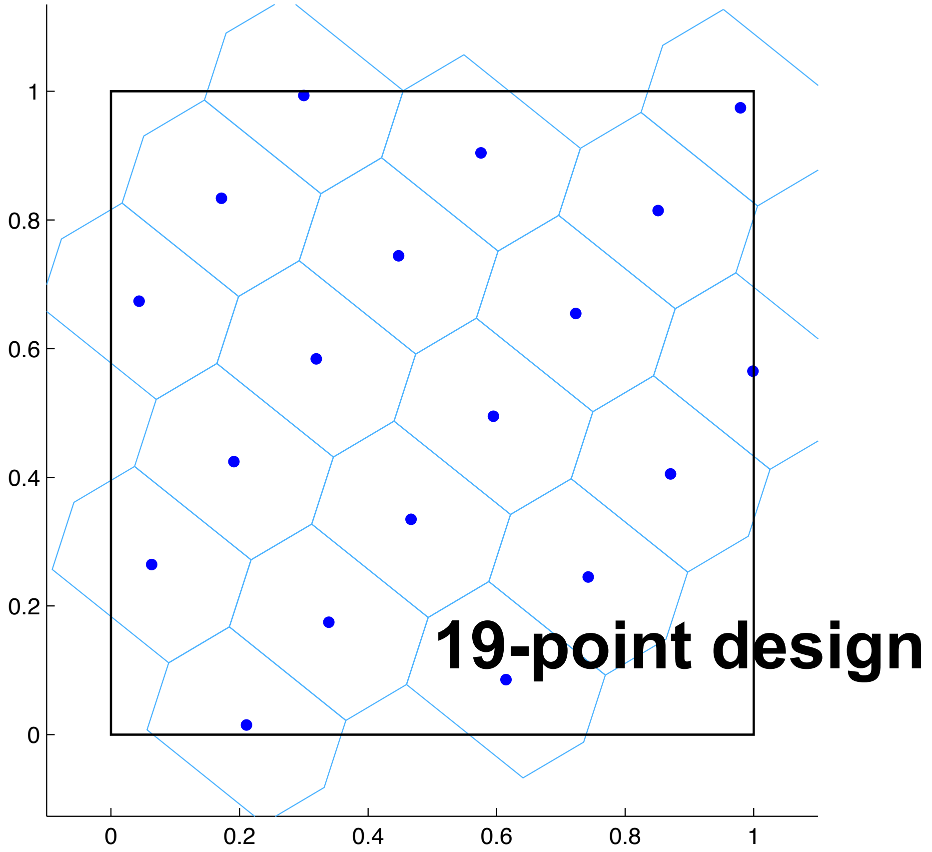

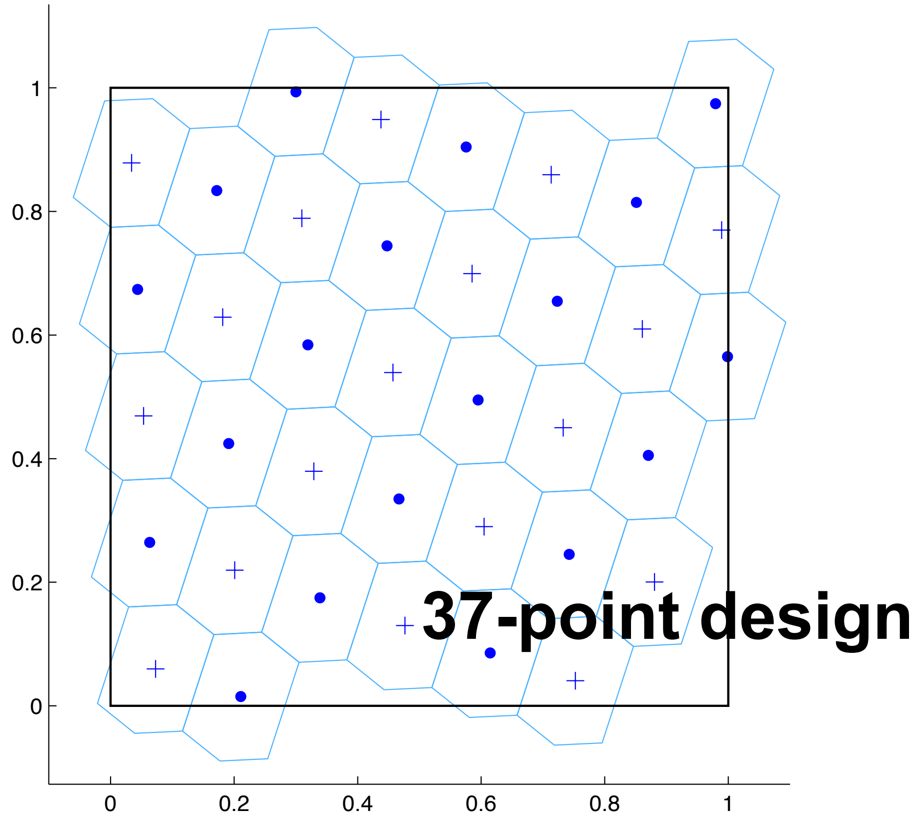

Heitmann et al. 2015 (in prep)

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)

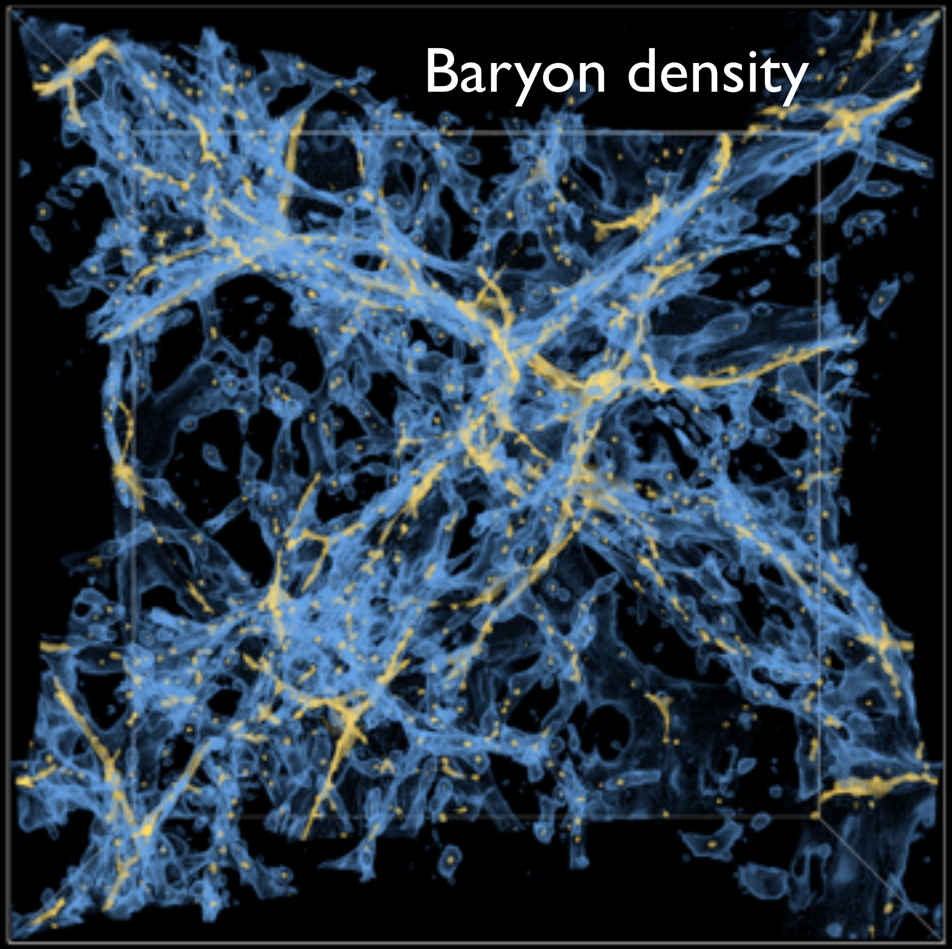

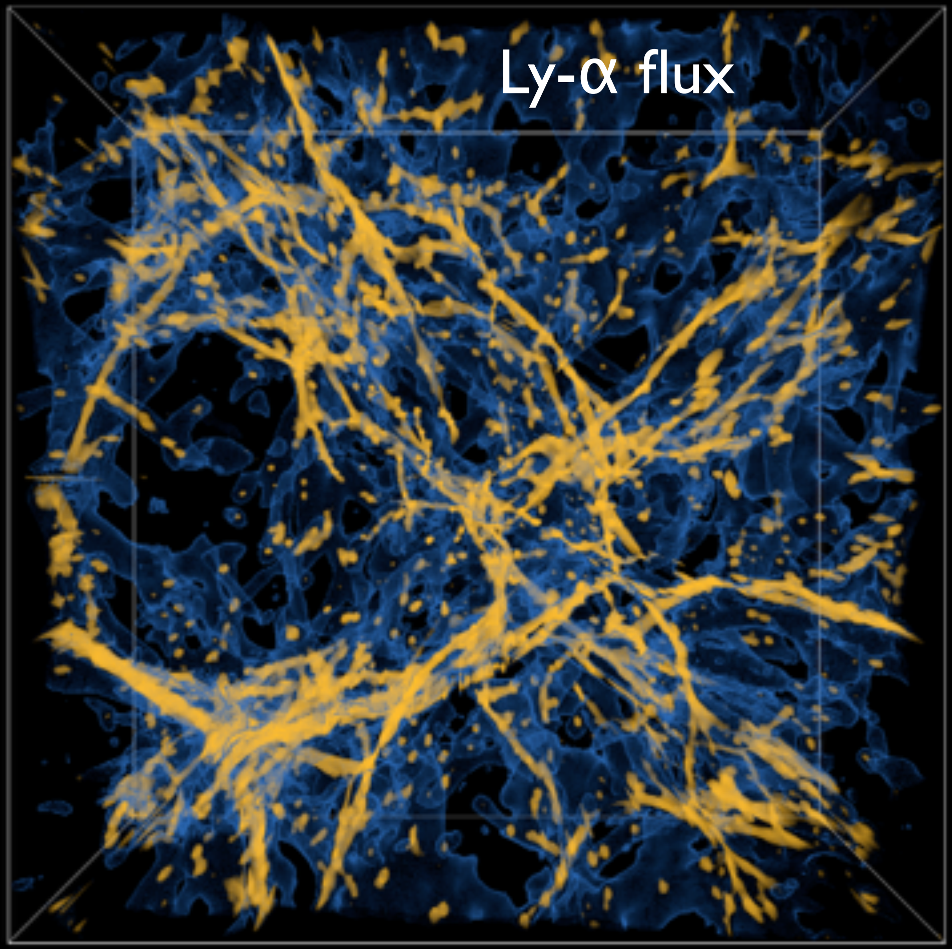

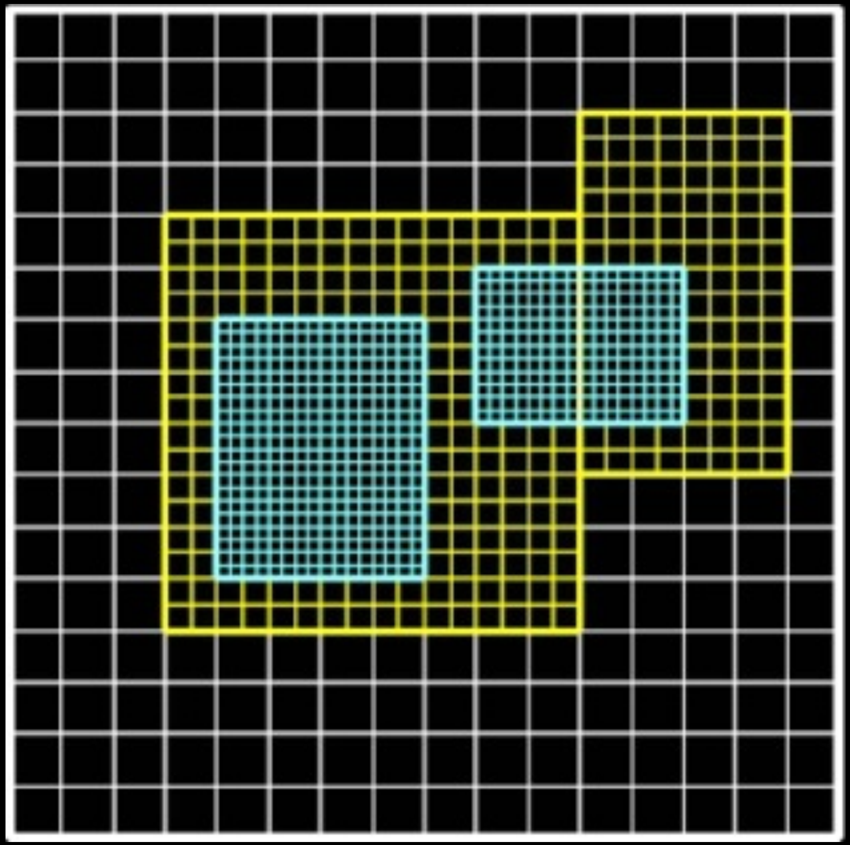

- AMR: patch-based refinement, with jump up to a factor of 4.

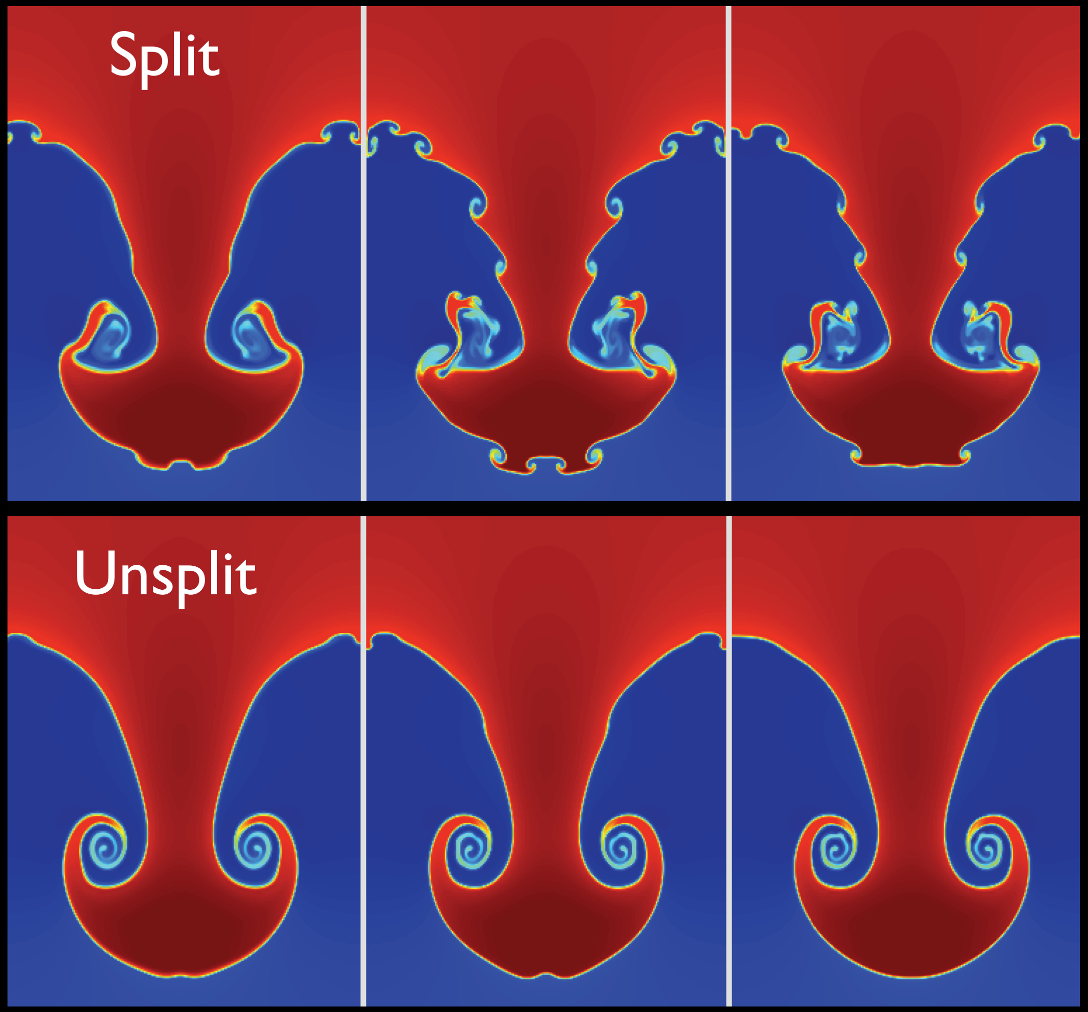

- Hydro: unsplit finite volume scheme better characterizes fluid flow.

Excellent scaling

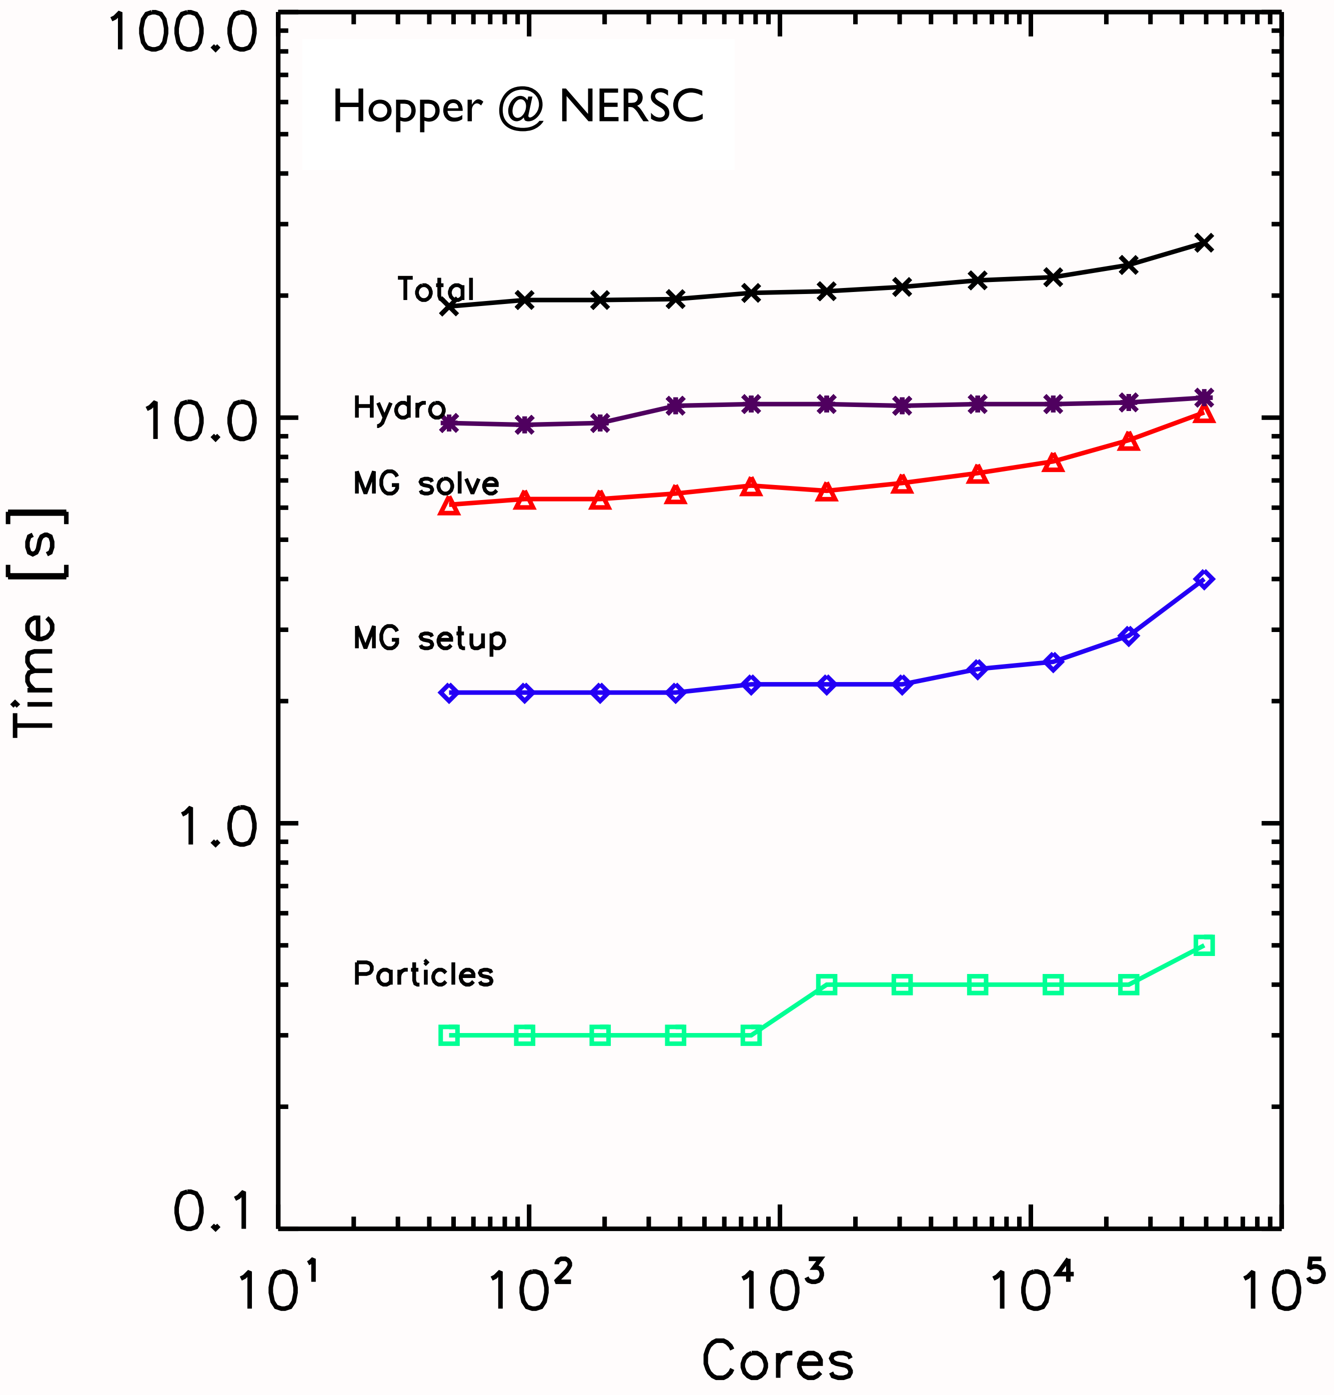

- Currently we are using NERSC resources under ALCC allocation.
- Mostly running 2048^3 and 4096^3 runs.
- Hopper/Edison: standard cluster architecture, 24 cores on a node, 32/64GB per node, ~5,000 nodes.
- Analysis pipeline on par with simulations.

4096^3 hydro simulation

Blue: F~0; Red: F~1

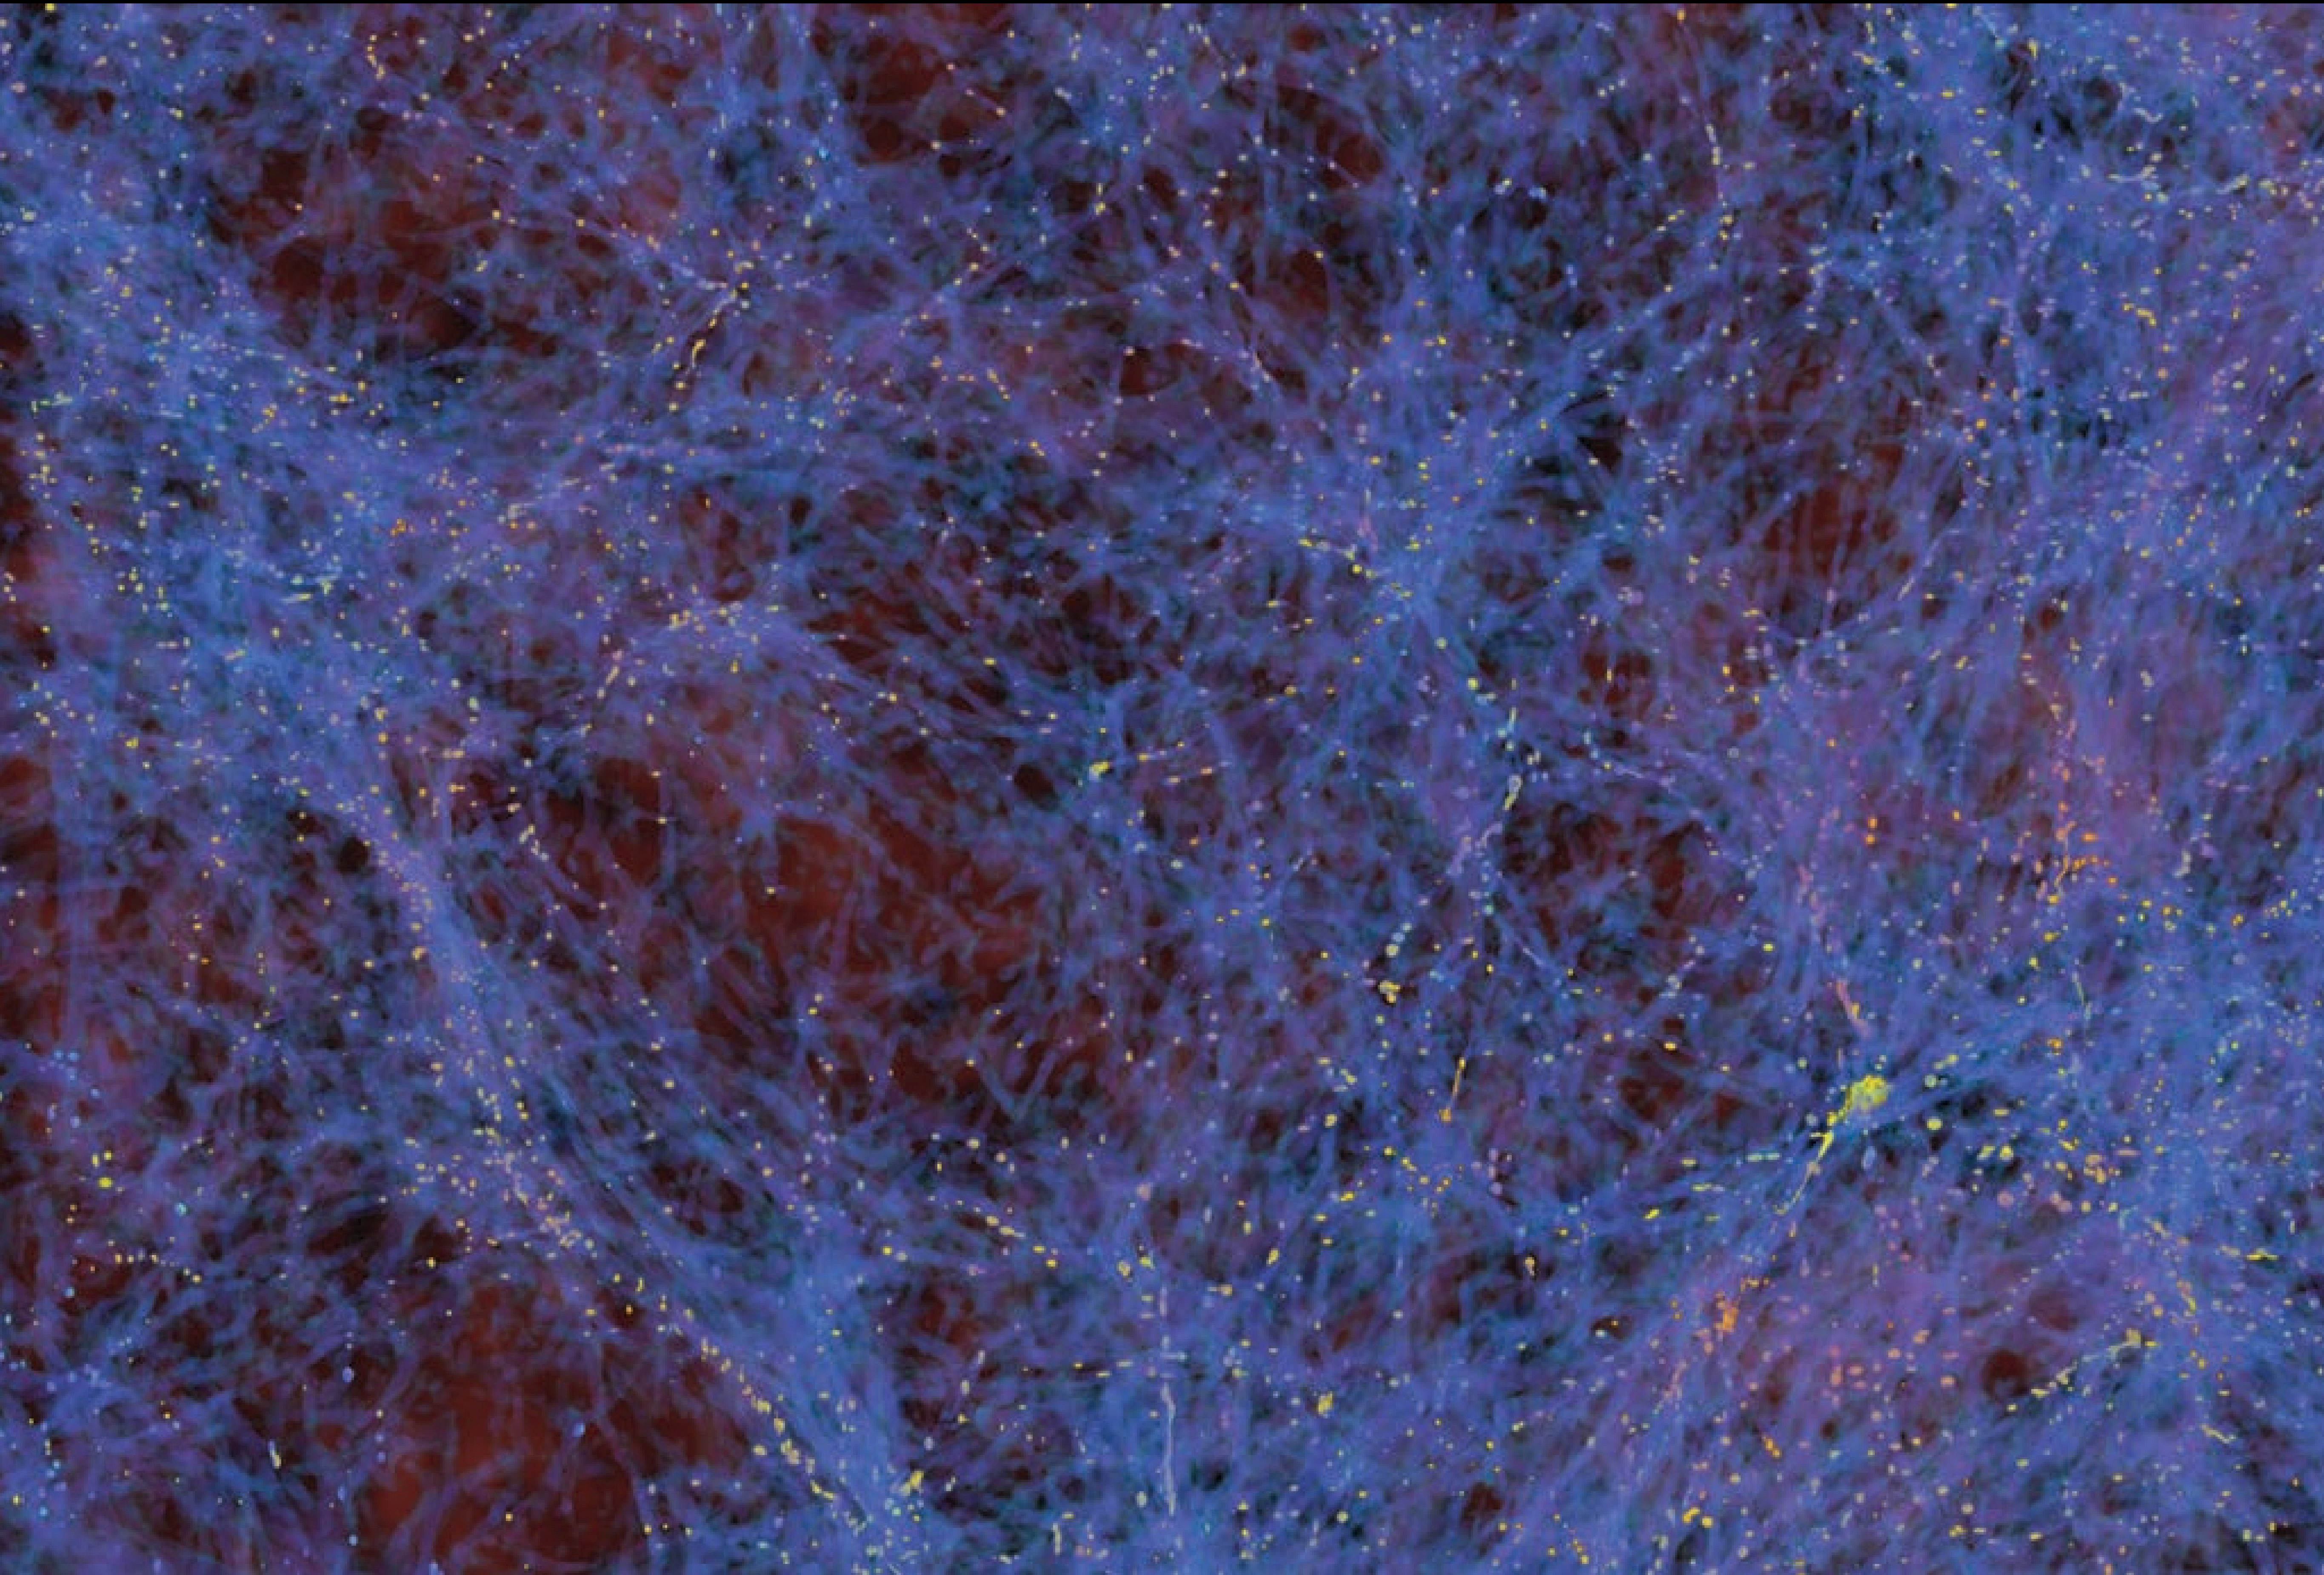

The Lyman-α forest in optically-thin hydro simulations

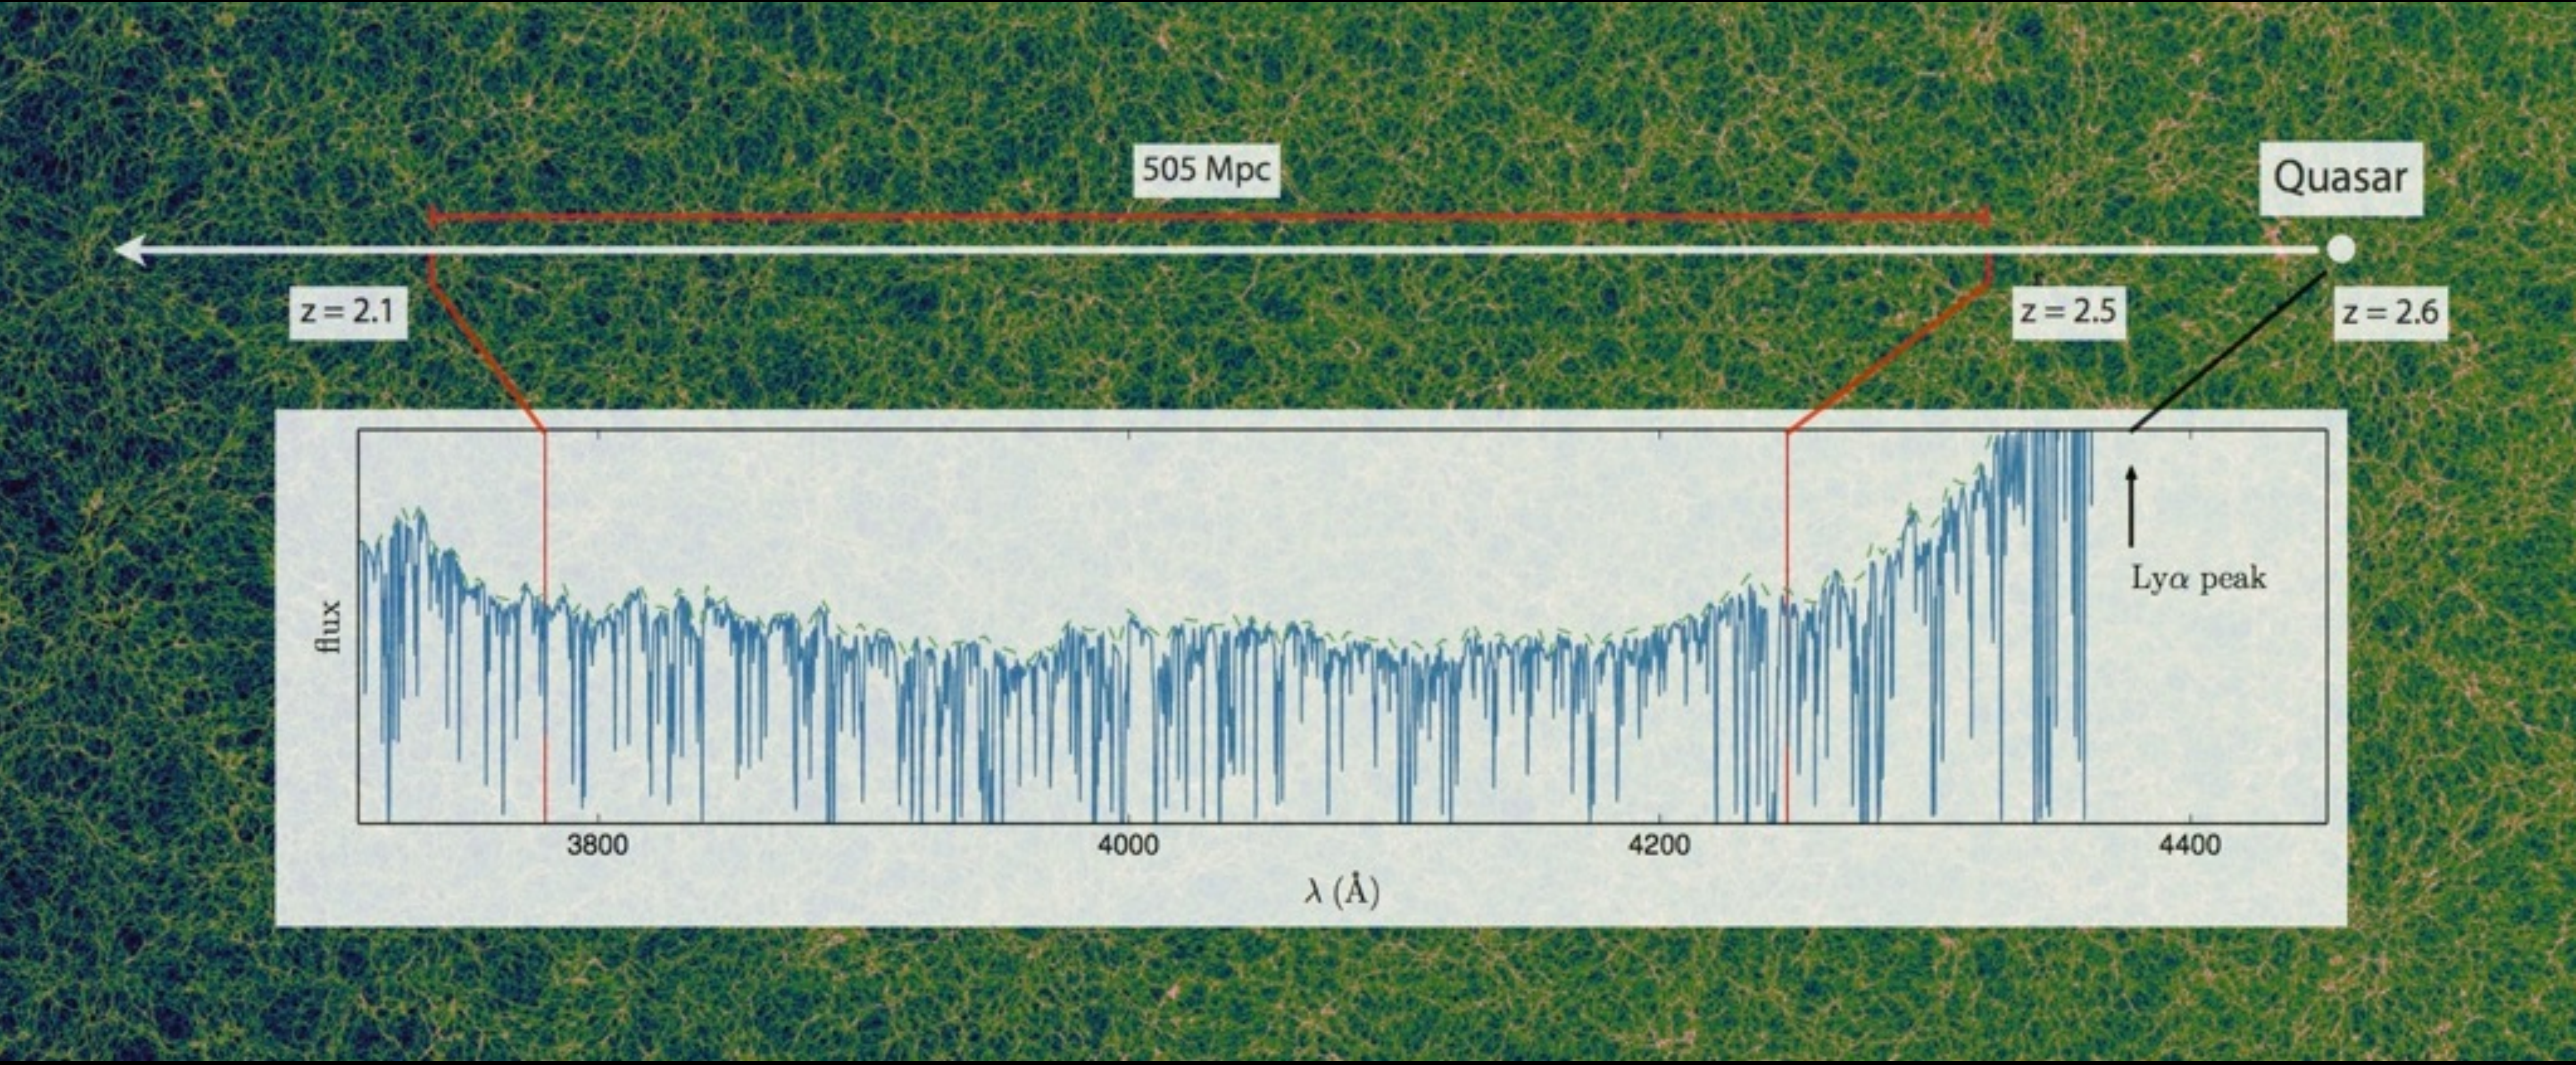

Zarija Lukić, Lawrence Berkeley National Laboratory

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

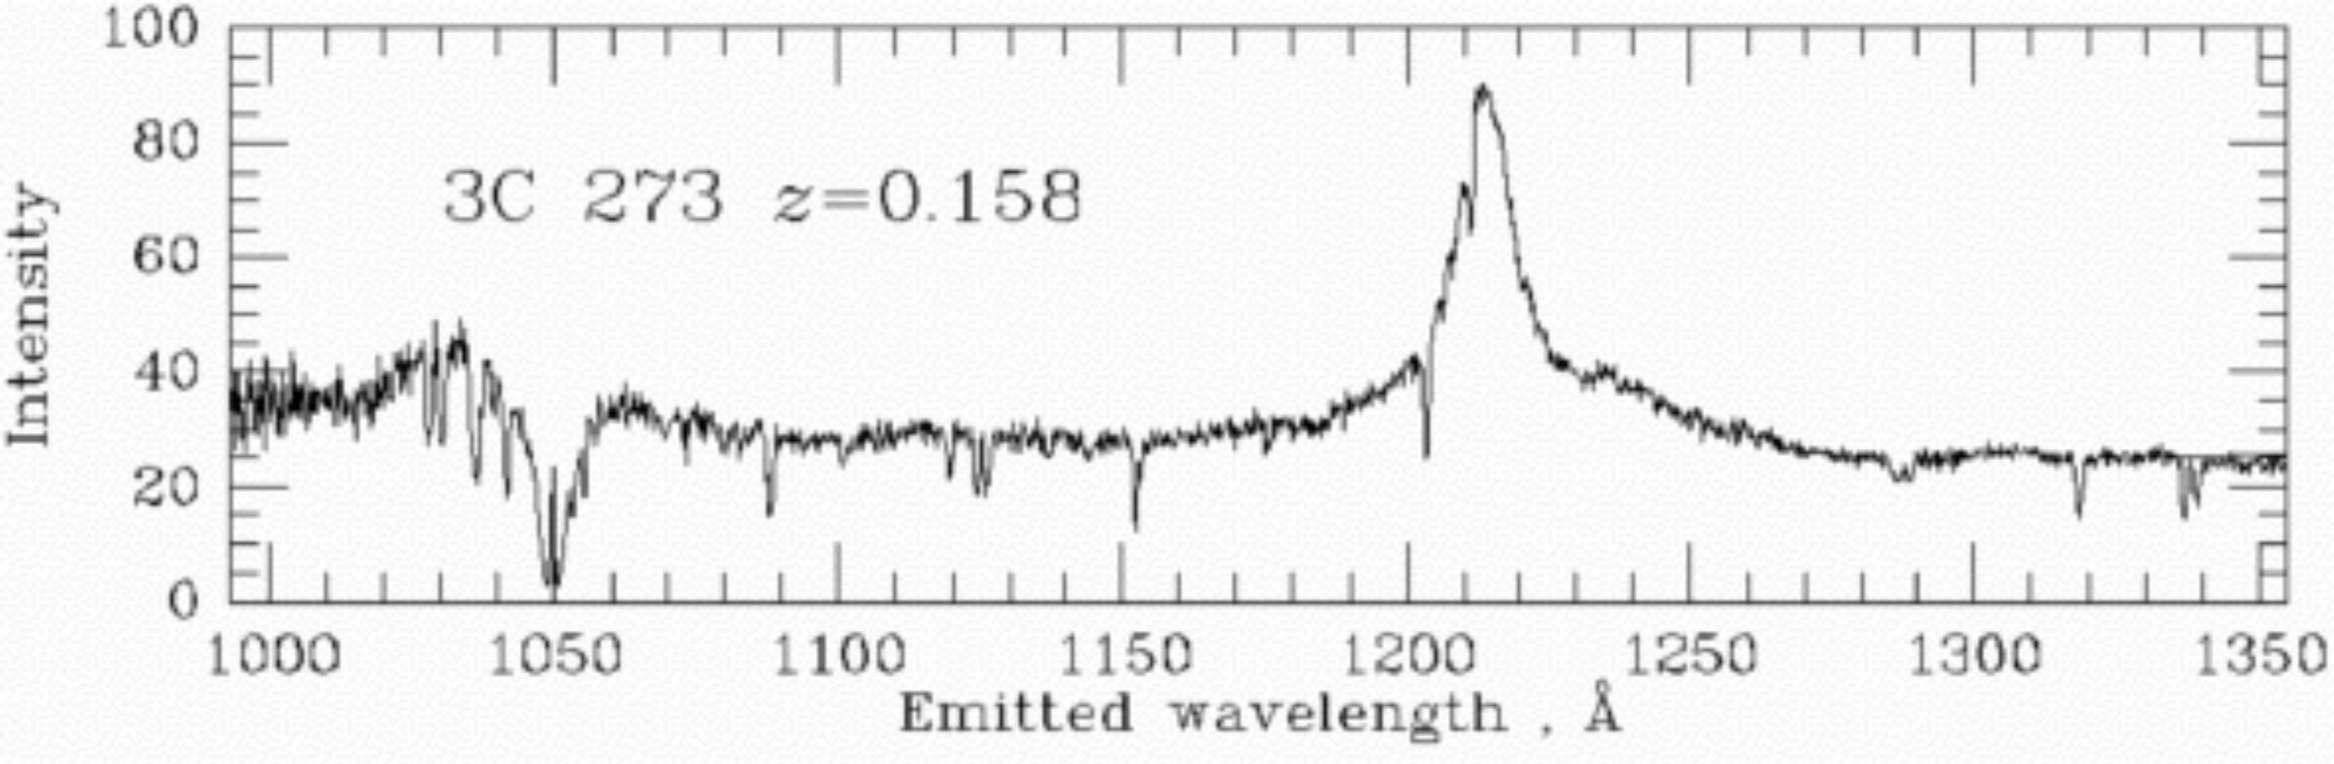
3C 273 z=0.158
Intensity
0
20
40
60
80
100
1000
1050
1100
1150
1200
1250
1300
1350
Emitted wavelength , Å

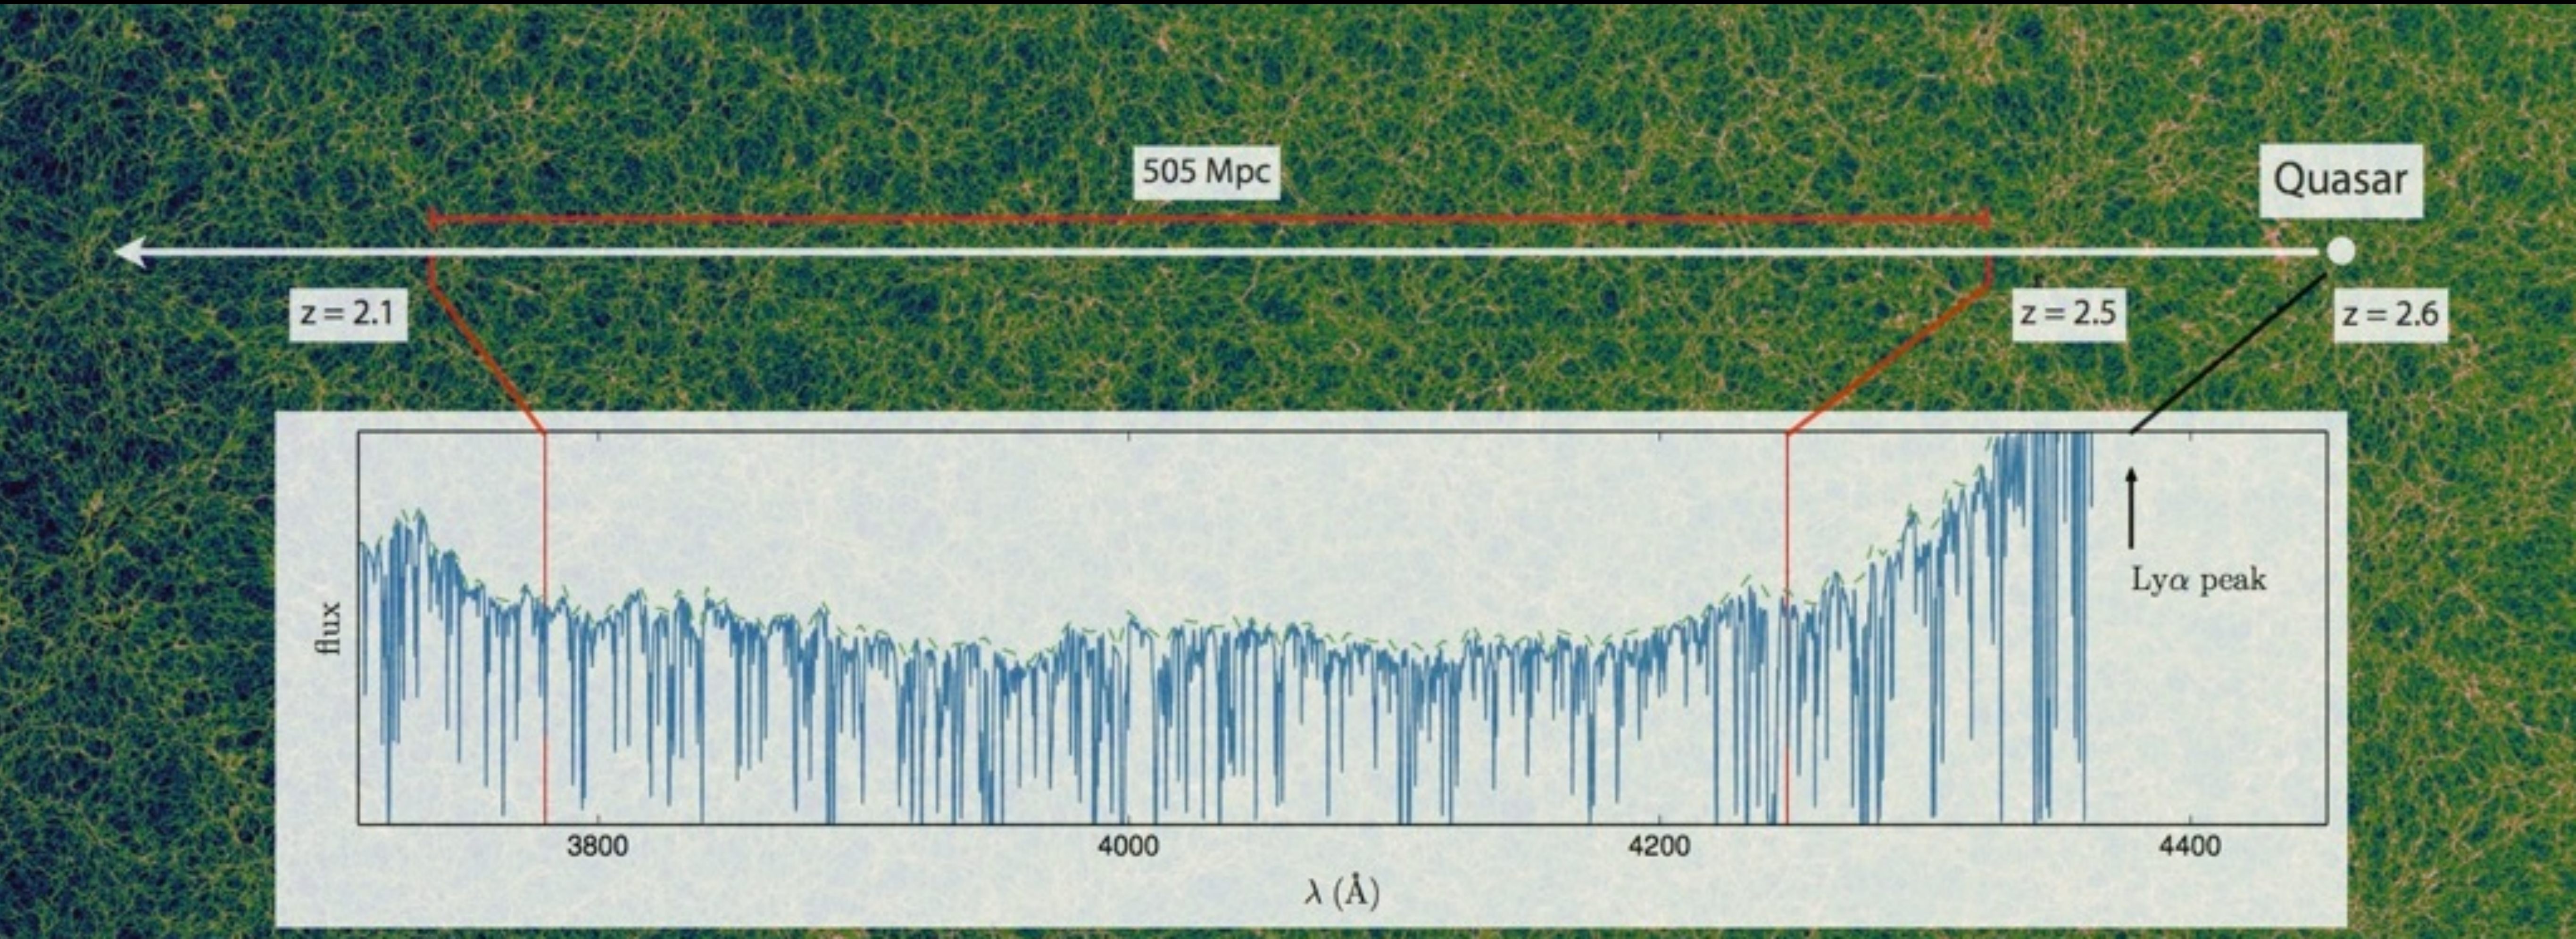
505 Mpc
Quasar
z = 2.1
z = 2.5
z = 2.6
Lyα peak
flux
3800
4000
4200
4400
λ (Å)

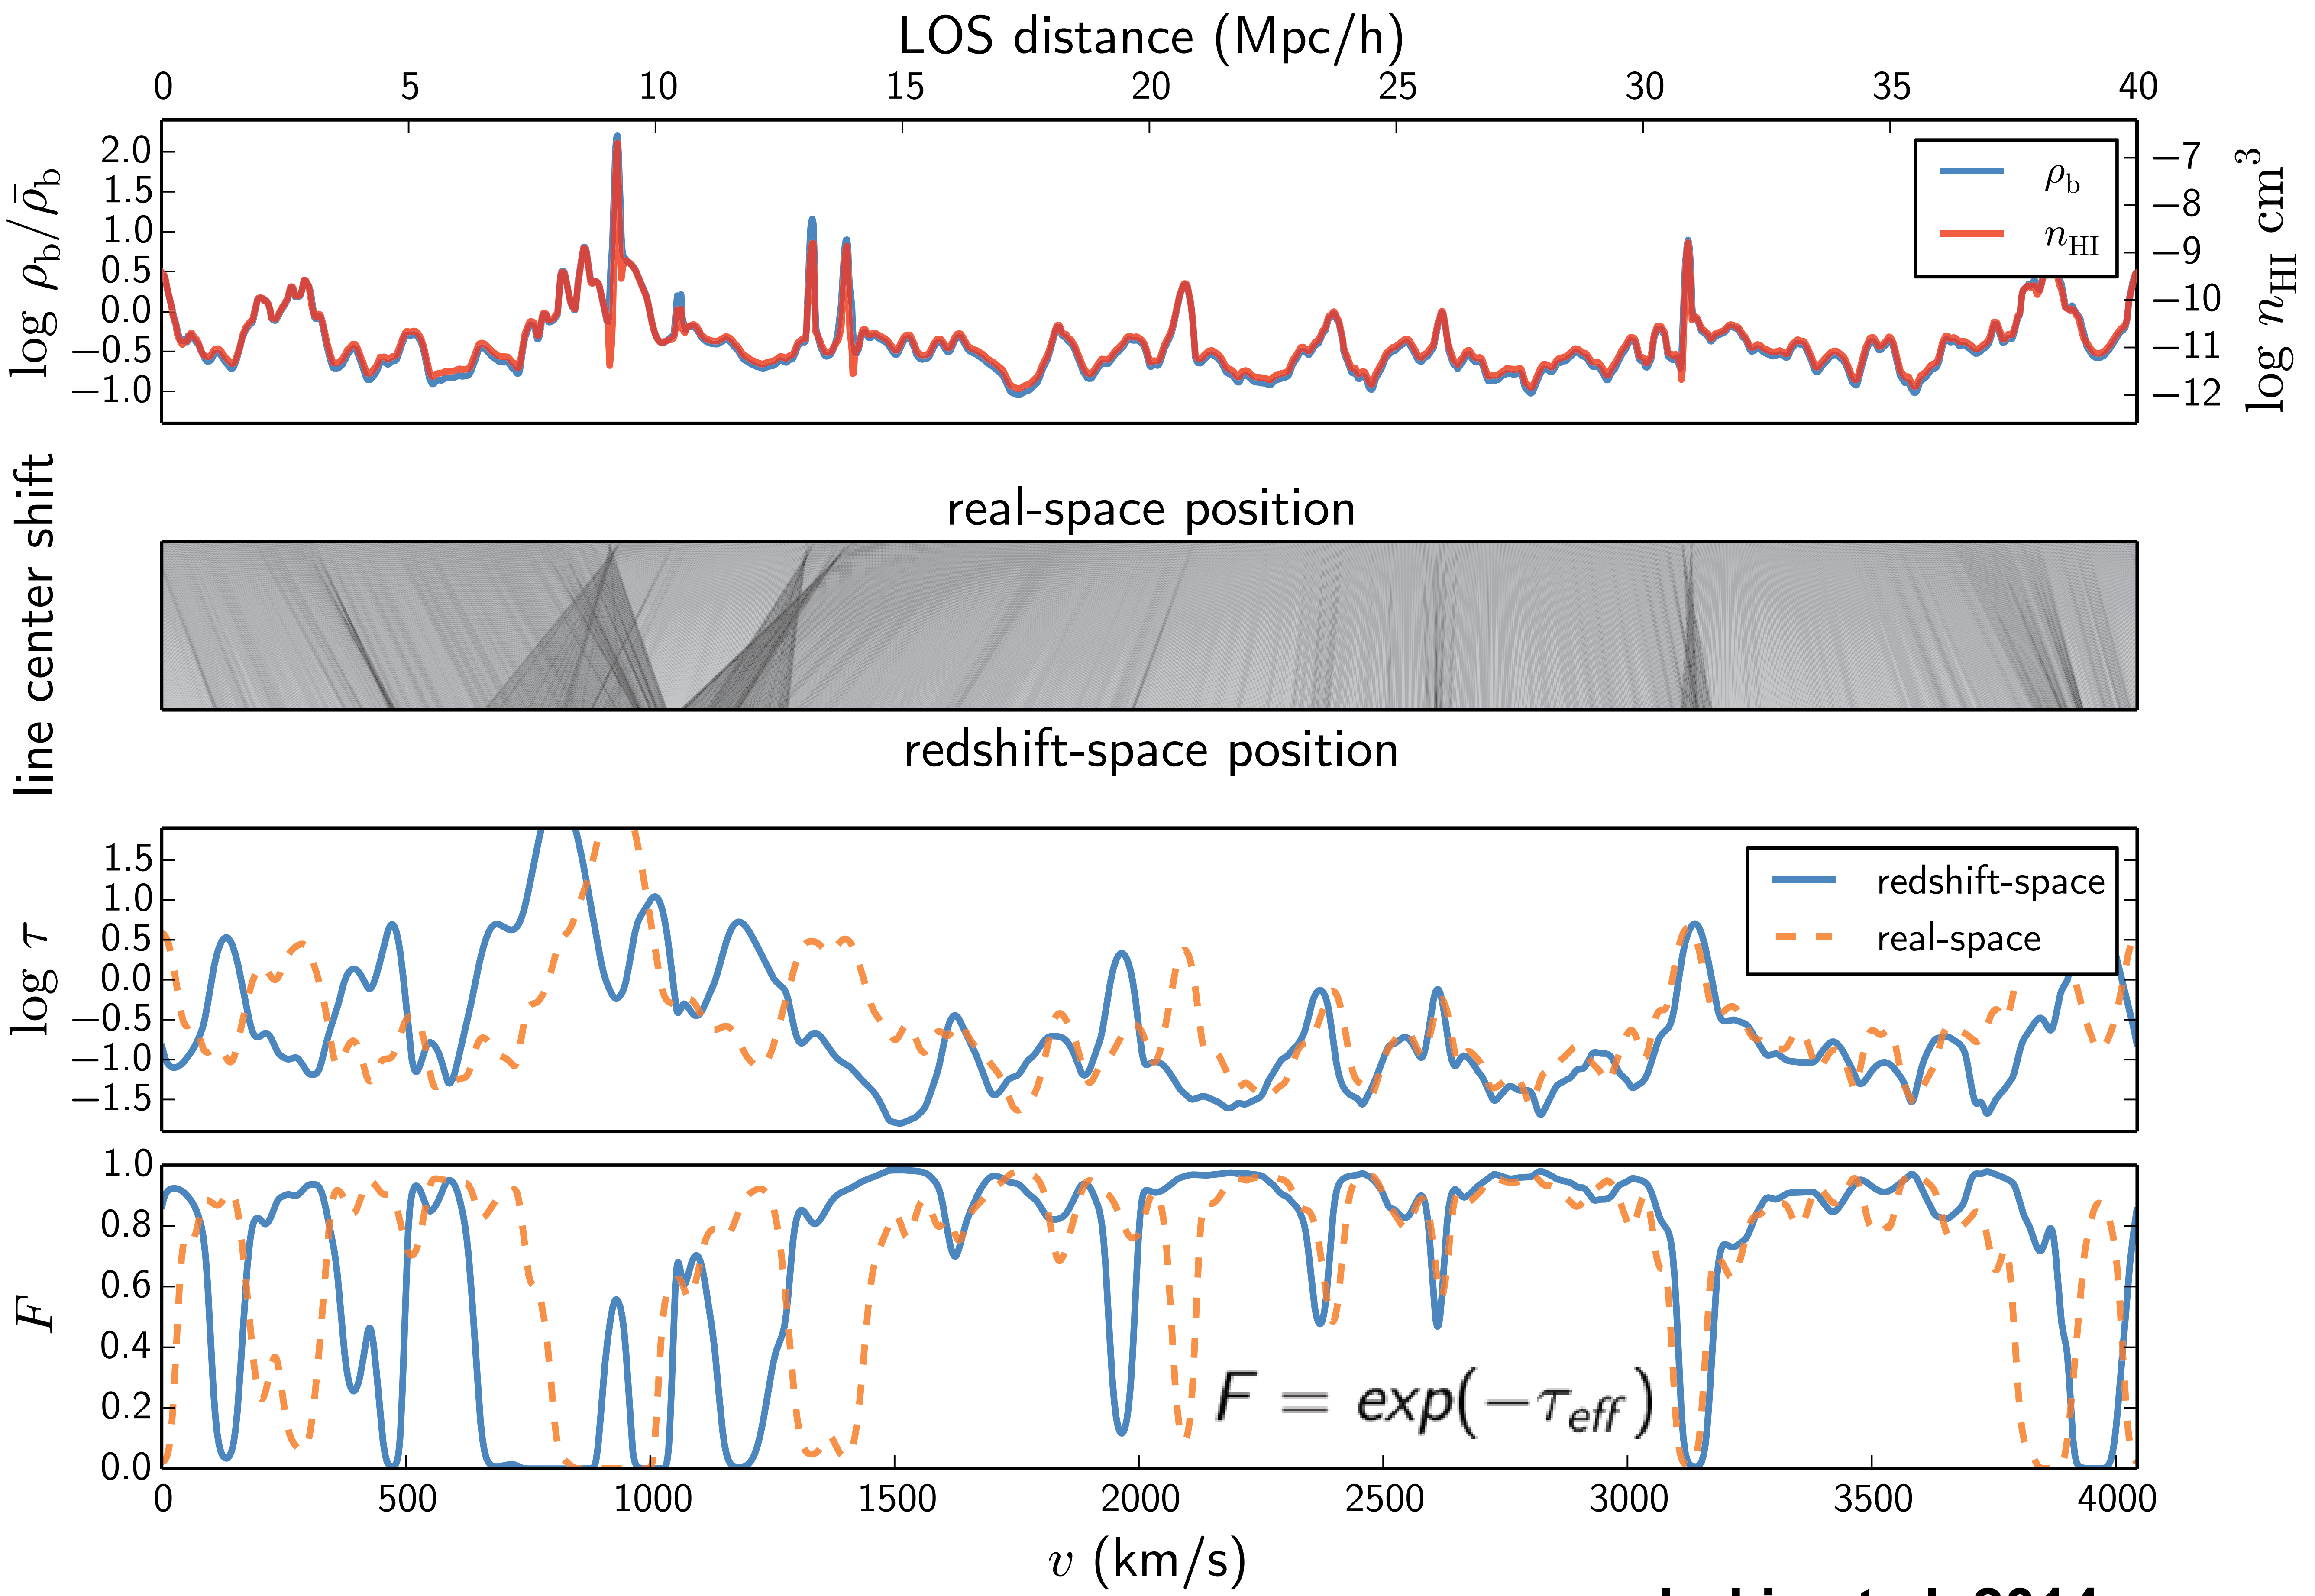

$F = exp(-\tau_{eff})$

Lukic et al. 2014

“Equation of state”

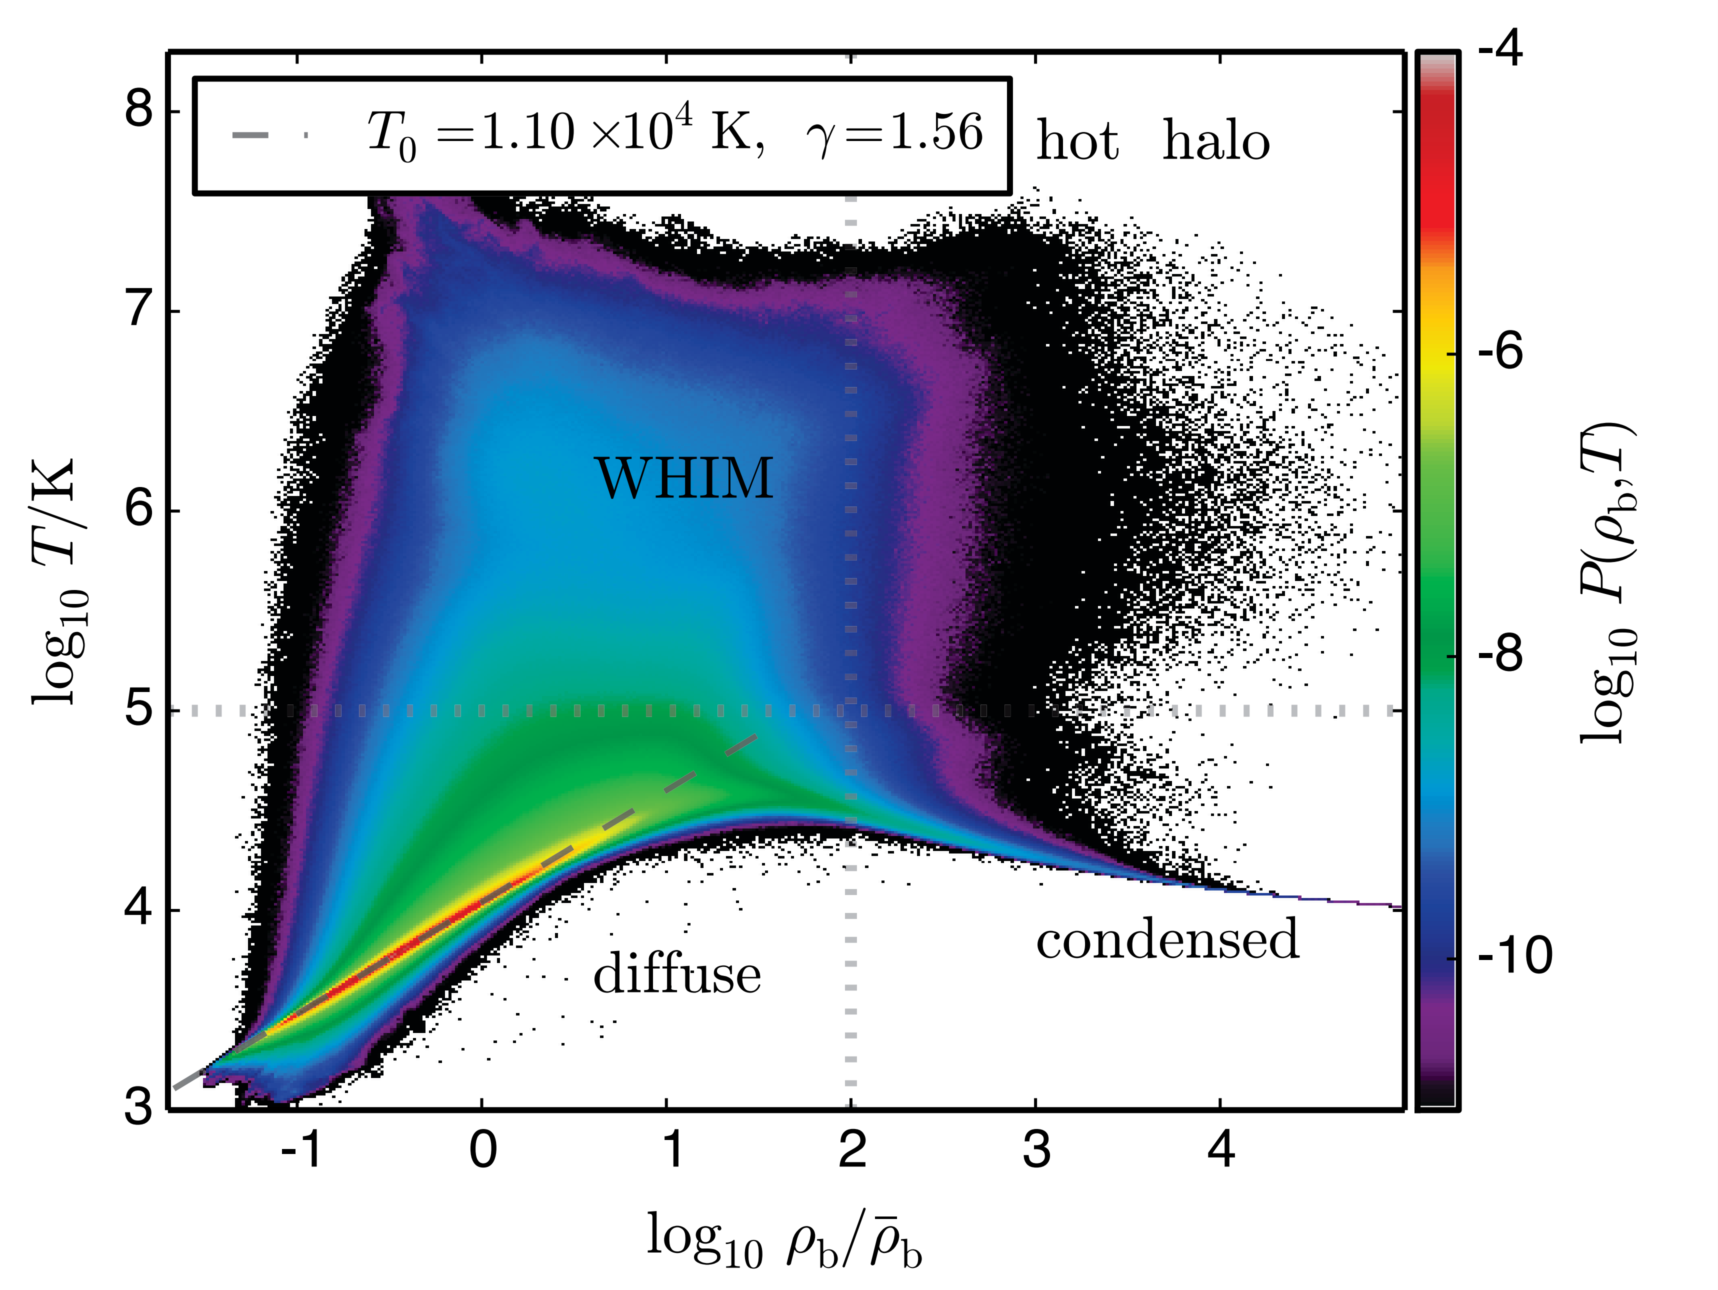

Nyx result from Lukic et al. 2014

$$T = T_0 \left(\frac{\rho}{\rho_0} \right)^{\gamma - 1}$$

- 4 phases of gas in simulations: “diffuse” relevant for the forest. (Tight density-temperature relation in this regime.)

Density - temperature

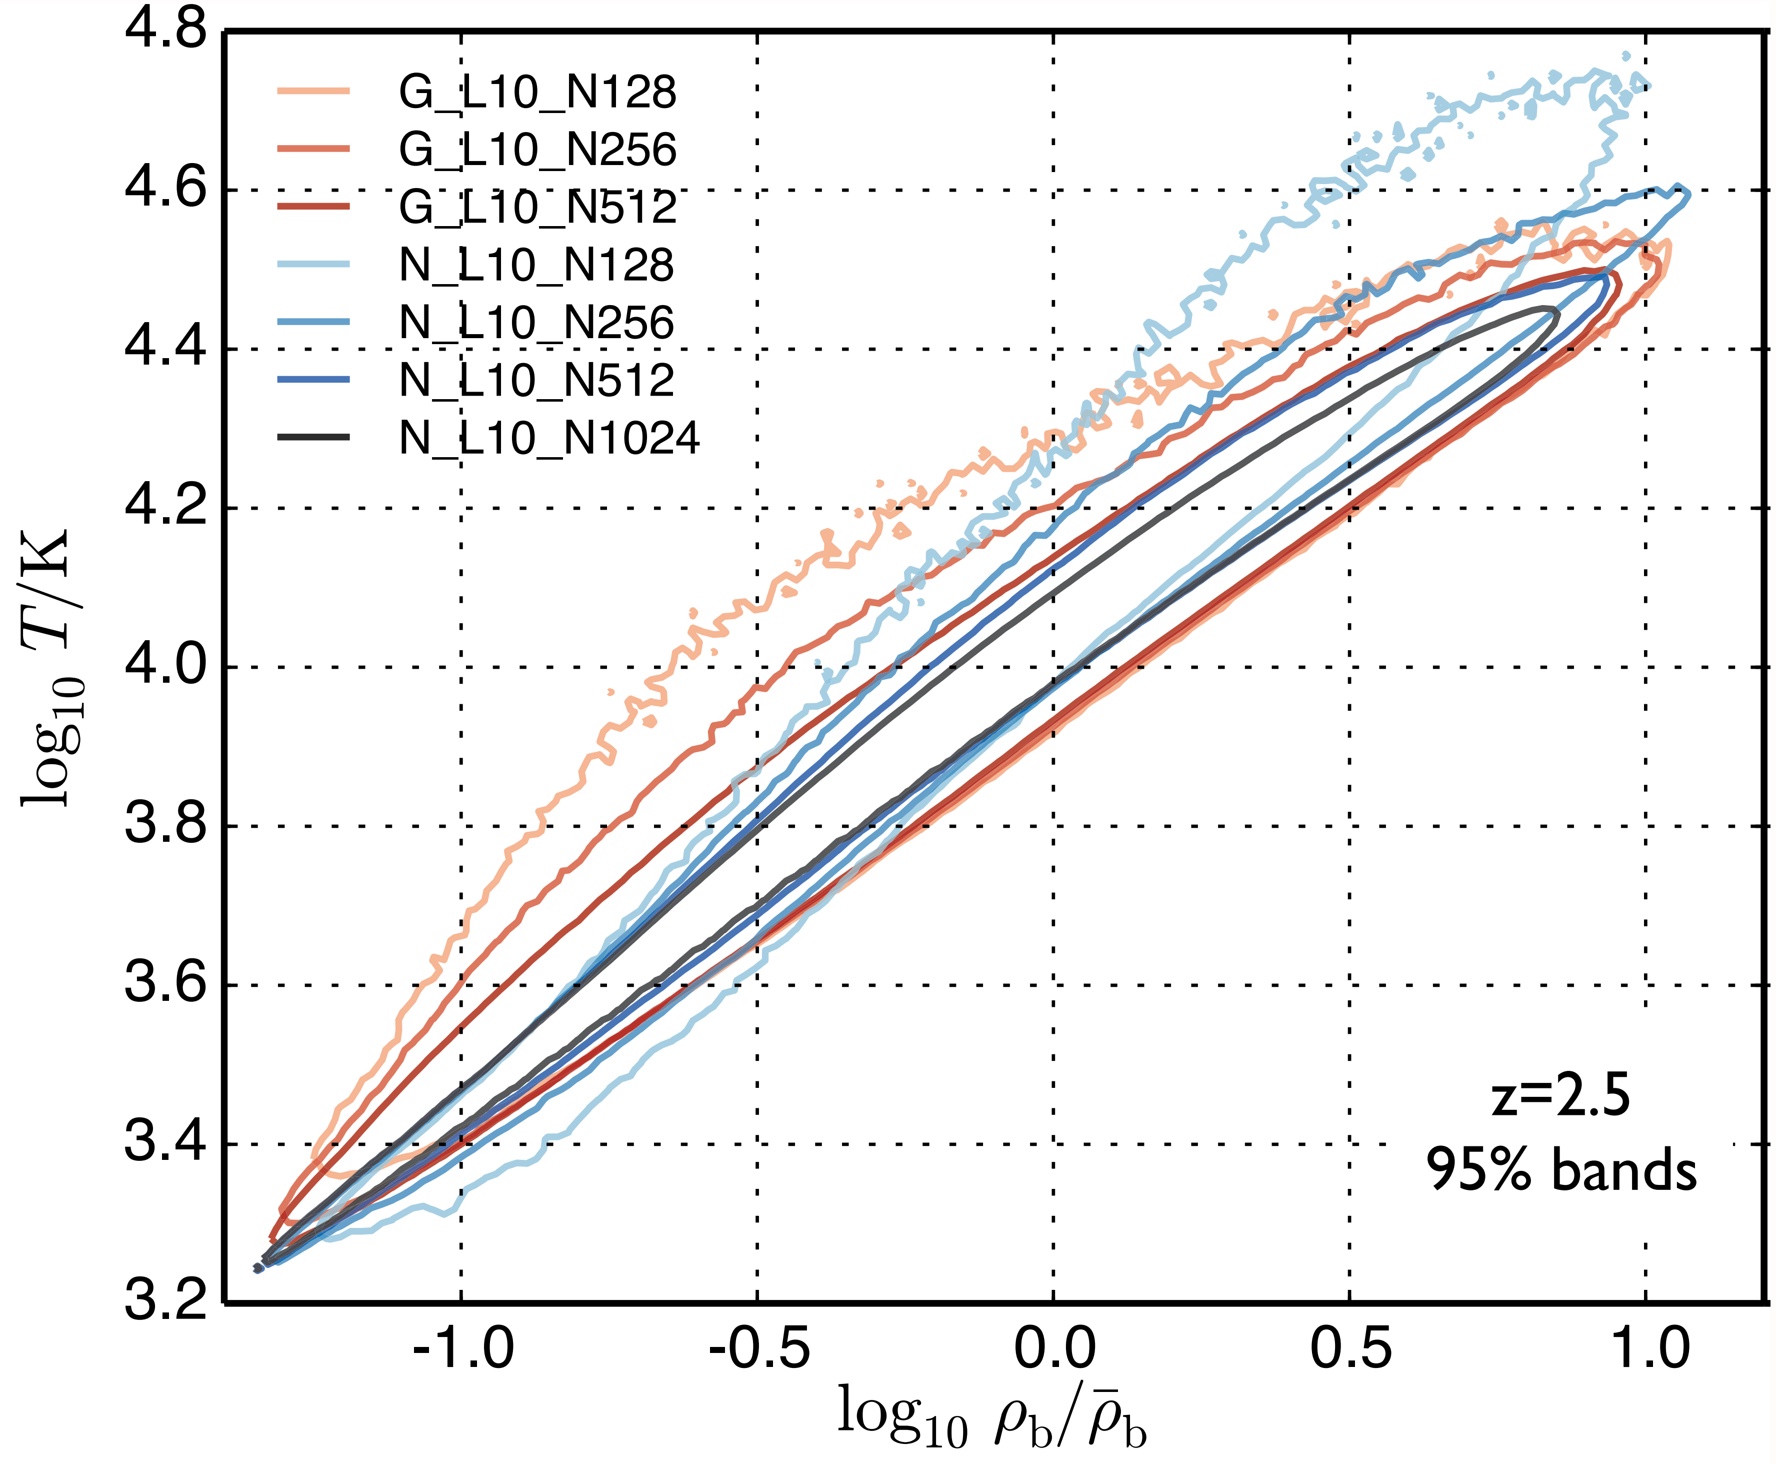

- SPH (Gadget) vs. Eulerian (Nyx) code.

Stark et al. in prep.

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

Simulation Inputs

Analysis Tools Configuration

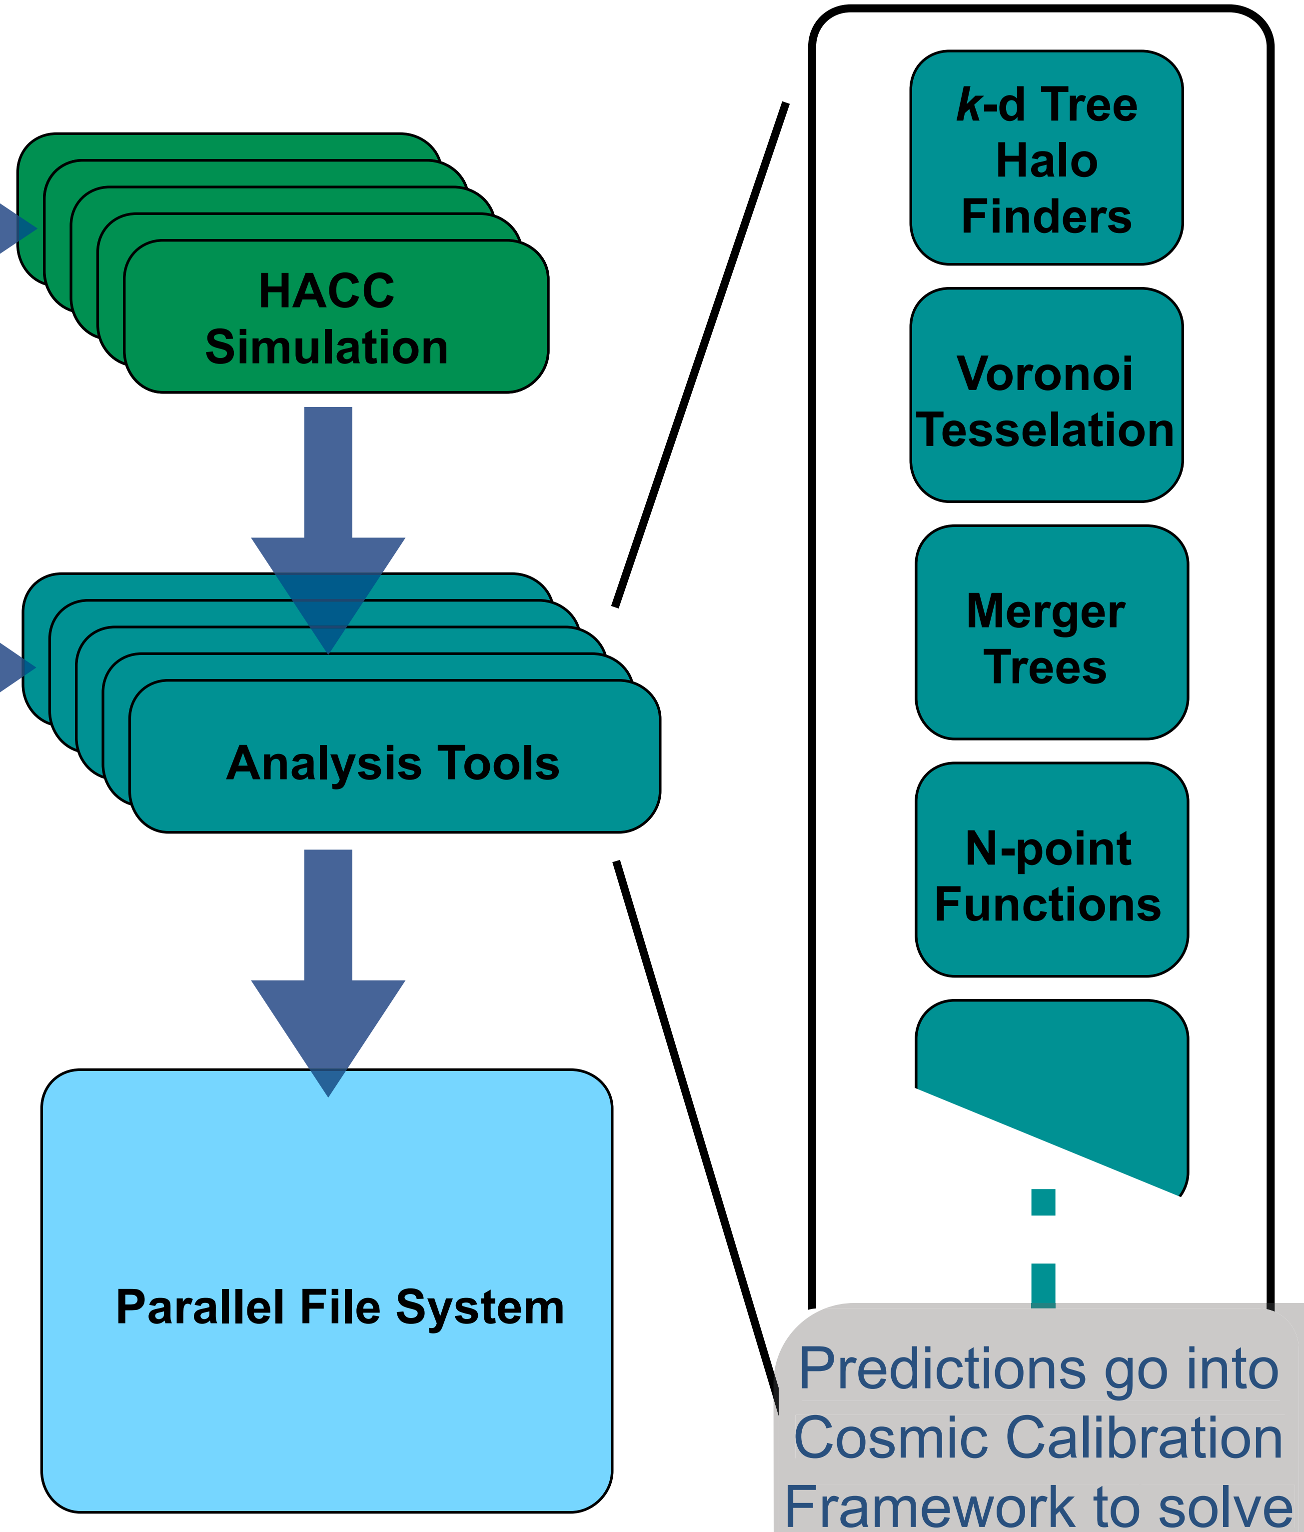

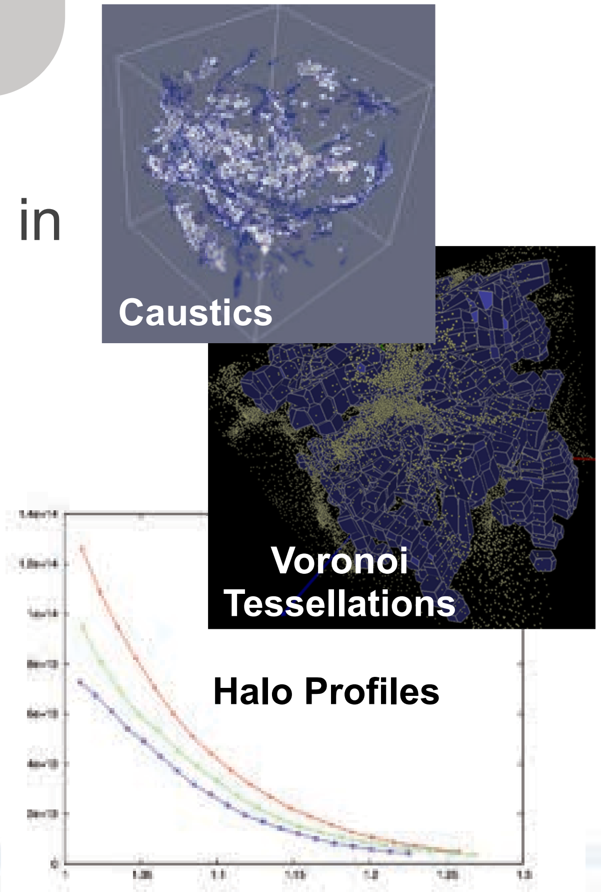

Predictions go into Cosmic Calibration Framework to solve the Cosmic Inverse Problem

In Situ Analysis and Co-Scheduling

- **Analysis Dataflows:** Analysis dataflows are complex and any future strategy must combine elements of in situ and off-line approaches
- **CosmoTools Test:** Test of coordinated off-line analysis ("co-scheduling")
- **Portability:** Analysis routines implemented using PISTON (part of VTK-m, built on NVIDIA's Thrust library)
- **Example Case (Titan):** Large halo analysis (strong scaling bottleneck) offloaded to alternative resource using a listener script that looks for appropriate output files

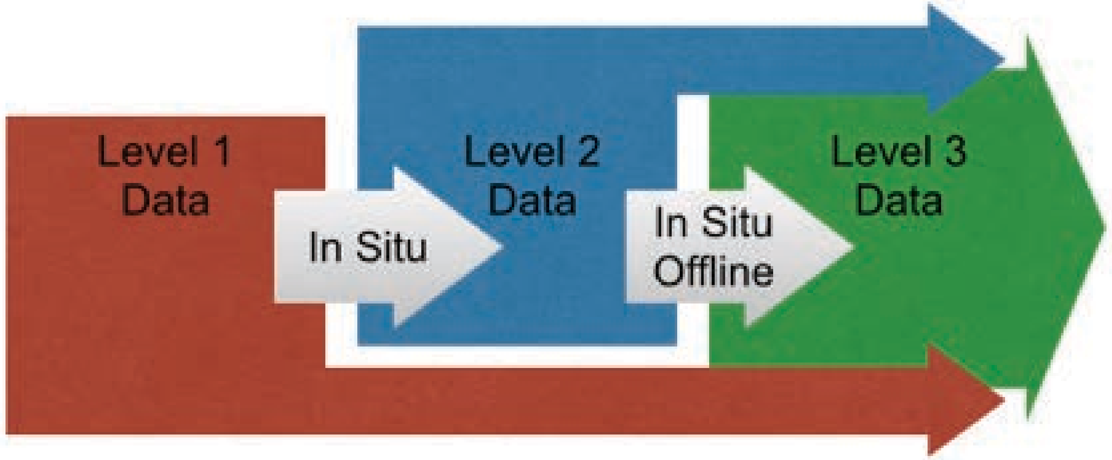

Level 1 Data

In Situ

Level 2 Data

In Situ Offline

Level 3 Data

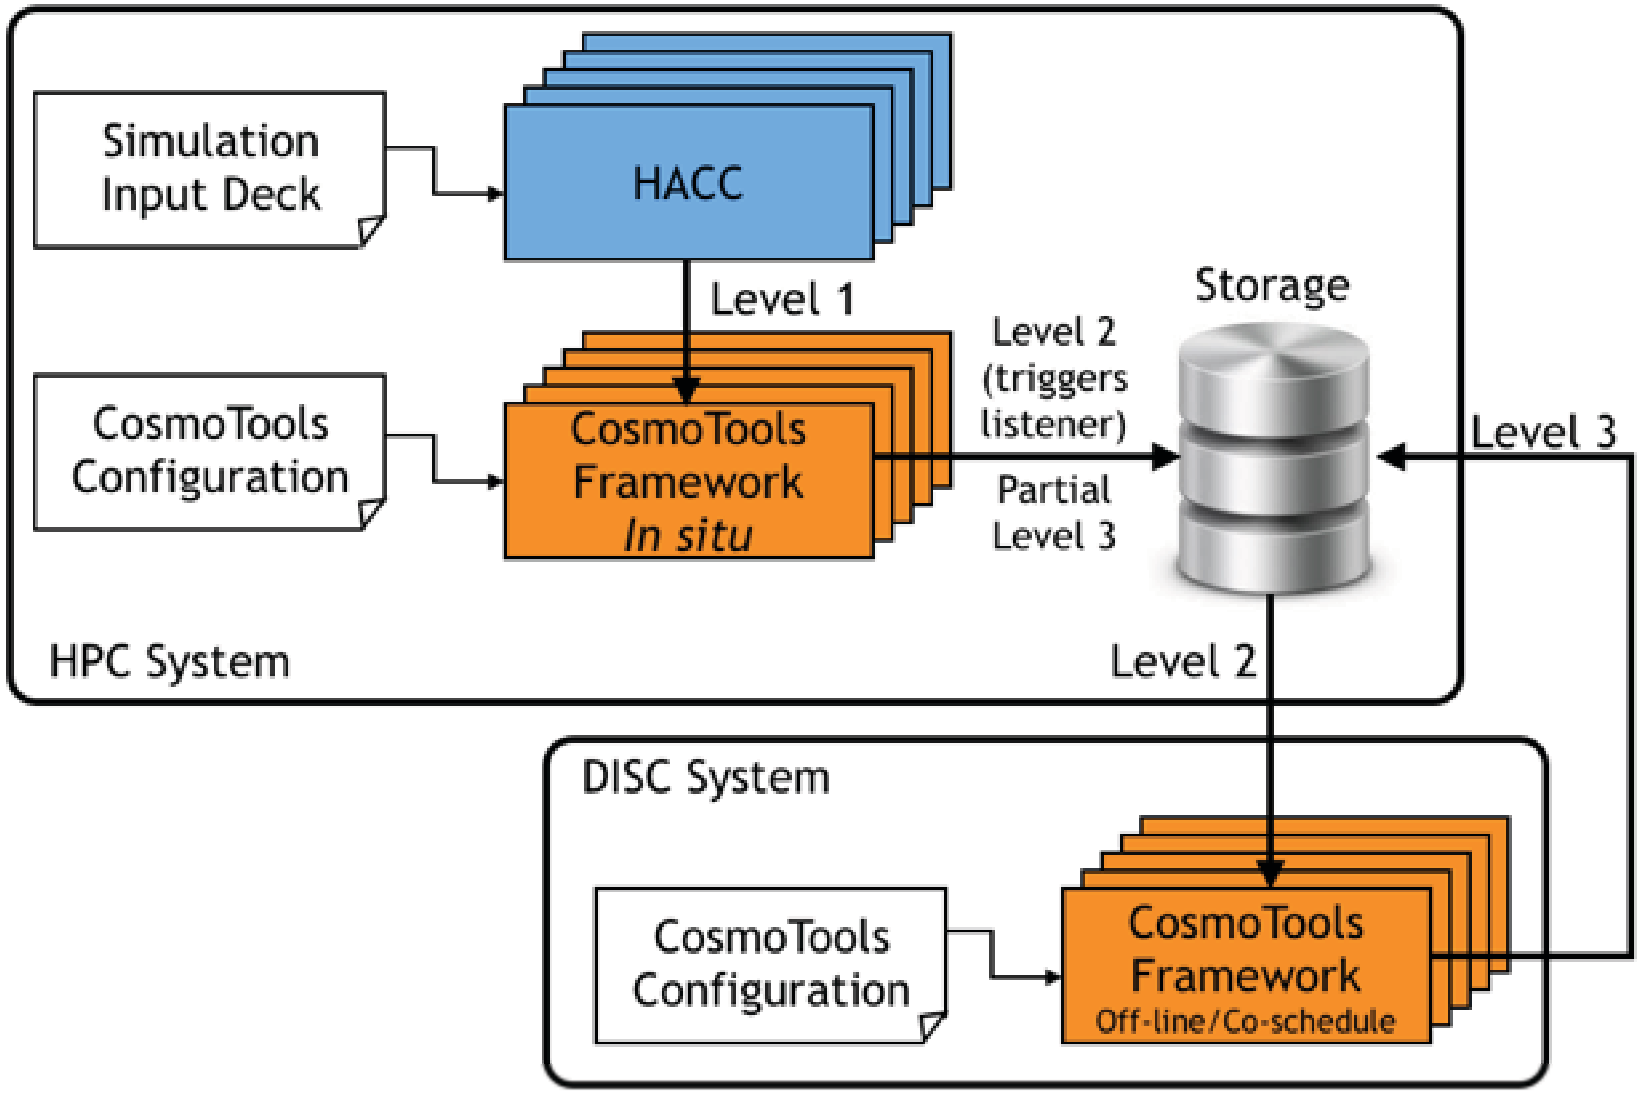

Simulation Input Deck

HACC

Level 1

CosmoTools Configuration

CosmoTools Framework *In situ*

Level 2 (triggers listener)

Partial Level 3

Storage

Level 3

HPC System

Level 2

DISC System

CosmoTools Configuration

CosmoTools Framework Off-line/Co-schedule

Sewell et al. 2015, SC15 (to appear)

SciDAC-3++: Computing the Sky — Simulation and Analysis for Cosmological Surveys

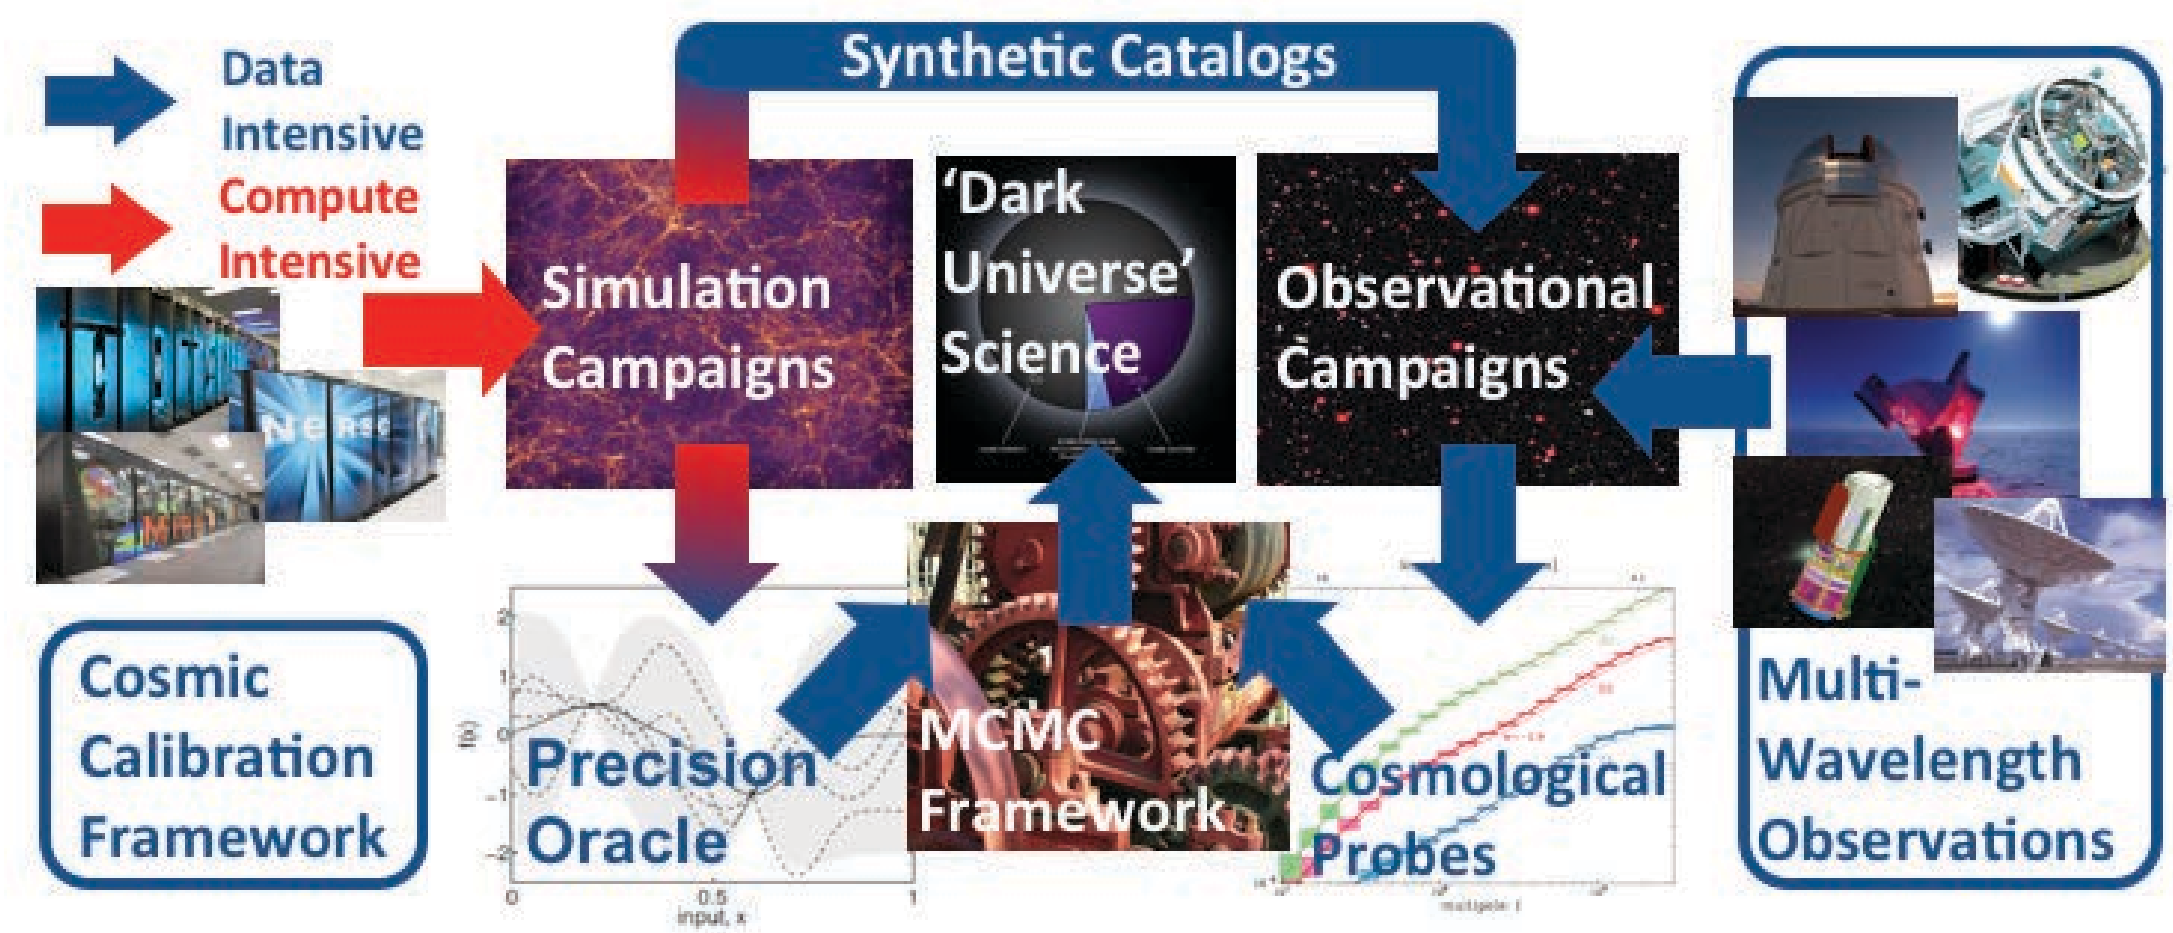

Data Intensive

Compute Intensive

Synthetic Catalogs

Simulation Campaigns

'Dark Universe' Science

Observational Campaigns

Cosmic Calibration Framework

Precision Oracle

MCMC Framework

Cosmological Probes

Multi-Wavelength Observations

- **Highlights:**
  - Next-generation emulators (including covariances), in situ visualization framework, merger trees in CosmoTools, new HACC algorithms for Summit
  - Add neutrino capability to Nyx, implement approaches for next-gen architectures, add in situ capability to Nyx
  - In situ data reduction schemes for ART, refactoring of ART I/O framework