Next-Generation Cosmological Simulations

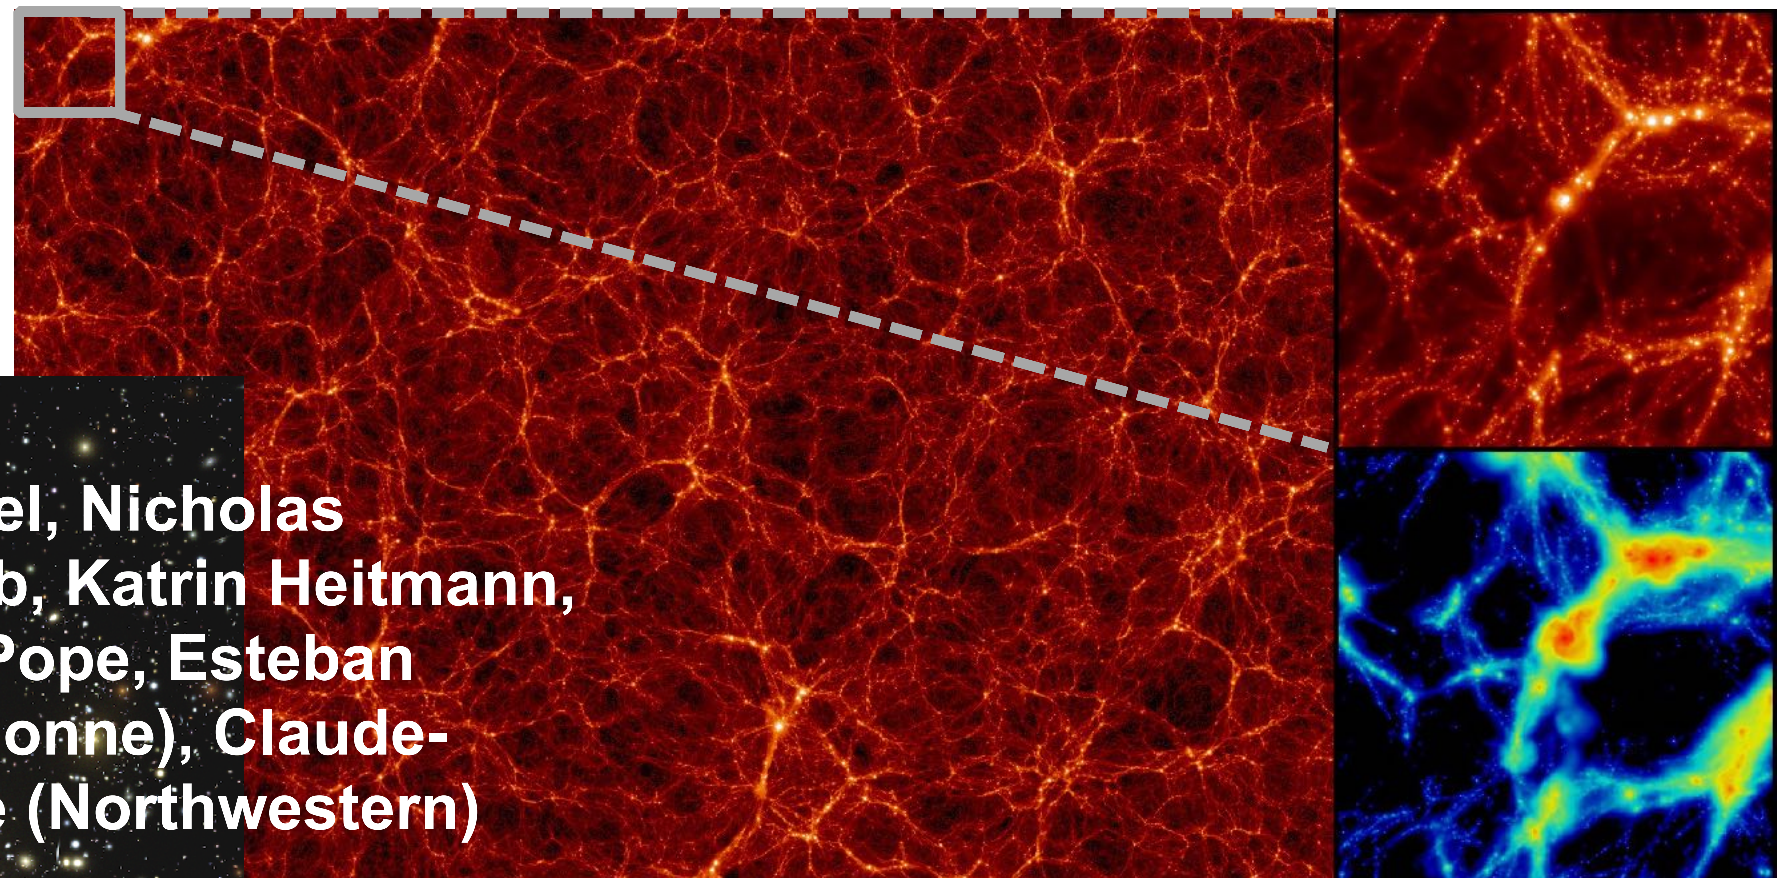

JD Emberson, Hal Finkel, Nicholas Frontiere, Salman Habib, Katrin Heitmann, Vitali Morozov, Adrian Pope, Esteban Rangel, Tom Uram (Argonne), Claude-Andre Faucher-Giguere (Northwestern)

Salman Habib
CPS and HEP Divisions
Argonne National Laboratory

CoSyne: Cosmological Synergies in the Upcoming Decade
Paris, December 9, 2019

Role of Cosmological Simulations

- **Theory of Structure Formation:** The Vlasov-Poisson equation is very poorly understood, especially in the case of gravity ("wrong-sign" Poisson equation)
  - No shielding unlike (non-neutral) plasmas
  - Formation of arbitrarily small scales over time
  - Chaotic particle trajectories
  - Perturbation theory only valid over a (pathetically) small spatial dynamic range
  - Only way to add realistic astrophysical mechanisms
  - Results from simulations often have the sense of a "discovery" (e.g., NFW)
  - A lot of intuition in building simple structure formation models comes from simulation results (halo models)
- **Modeling Surveys:** Cosmology is an observational science, simulations are very important for checking the robustness of inference methods
  - Can create classes of "close to observed" universes
  - Synthetic catalogs useful for testing analysis methods and optimization of survey strategies
  - End-to-end simulations allow for testing a vast number of potential systematics
  - New forward modeling approaches for generating galaxy catalogs

Cosmological Simulations: Good News and Bad News

- **Cosmological simulations: Bad News**
  - Need large amounts of memory
  - Need large spatial dynamic range everywhere
  - Need large dynamic range in mass
- **Comparison with other fields**
  - Dynamic range in time is not extreme wrt to molecular dynamics or protein folding (picosec to sec; femtosec to millisec)
  - Computational work can be parallelized efficiently (closest to plasmas and unlike, say, dynamical systems)
  - Complexity is not extreme (unlike, say, climate simulations)
- **Good News:**
  - Next-generation/exascale systems well-suited for cosmological simulations

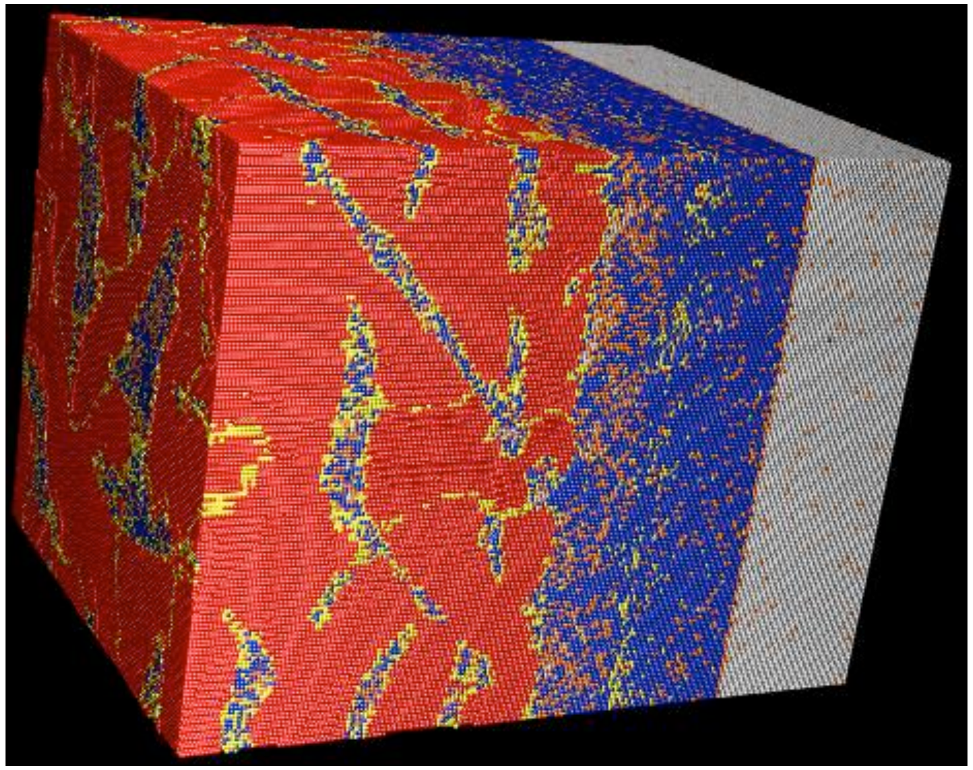

Molecular dynamics simulation of shock-induced material transformation (~9ps)

Equations/Physics

- Gravity dominates at large scales, need to solve the Vlasov-Poisson equation, a ***6-dimensional*** PDE
- Very complex solutions, reduction to lower dimensions not possible because of multi-streaming
- 6-D nature precludes grid-based methods (other problems too)
- Use N-body methods with the understanding that f(x,p) must be a smooth function
- Gas dynamics coupled in using Euler equations; astrophysical effects (heating, cooling, star formation enter as subgrid models)

Particle or Grid?

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

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

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

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

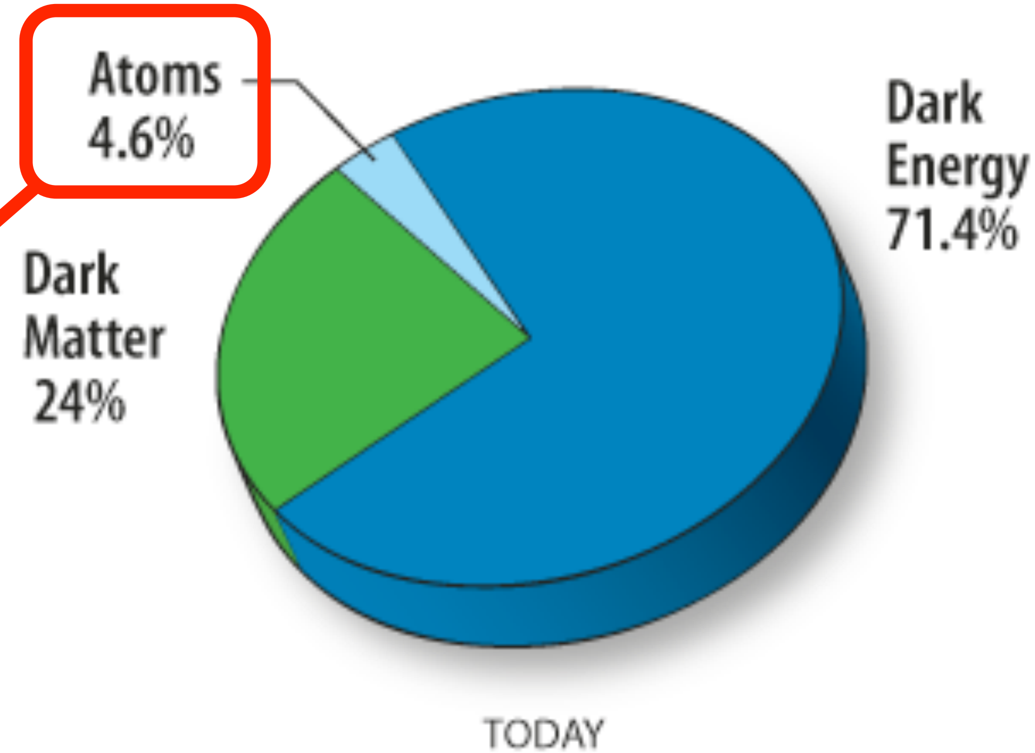

Current composition of the Universe [WMAP 9 result, courtesy NASA]

Cosmology: Data-Driven Progress

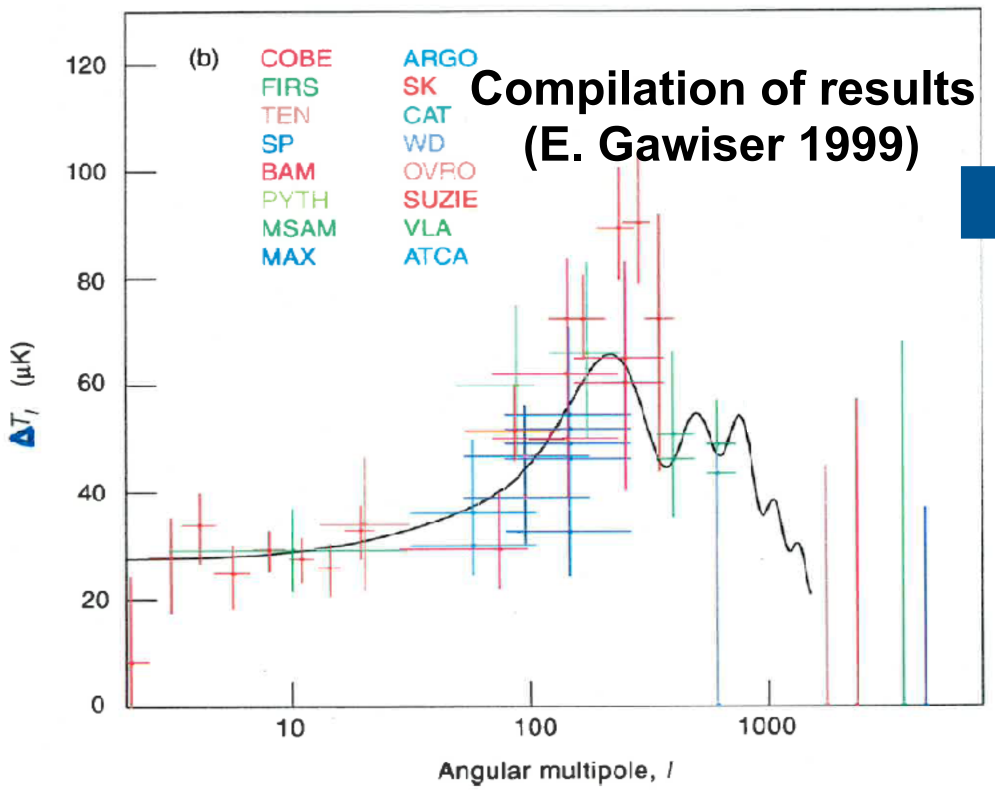

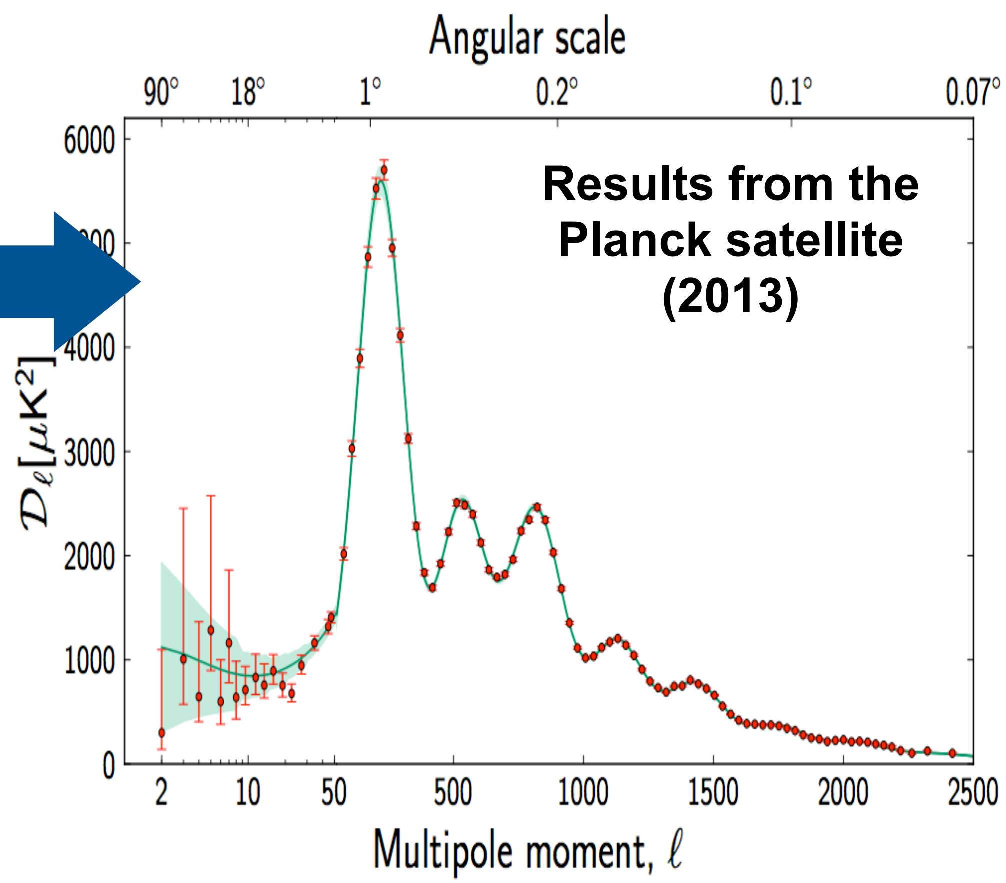

- Massive increase in sensitivity of cosmic microwave background (CMB) observations
- Cross-correlation with galaxy surveys
- New era of CMB modeling/simulations

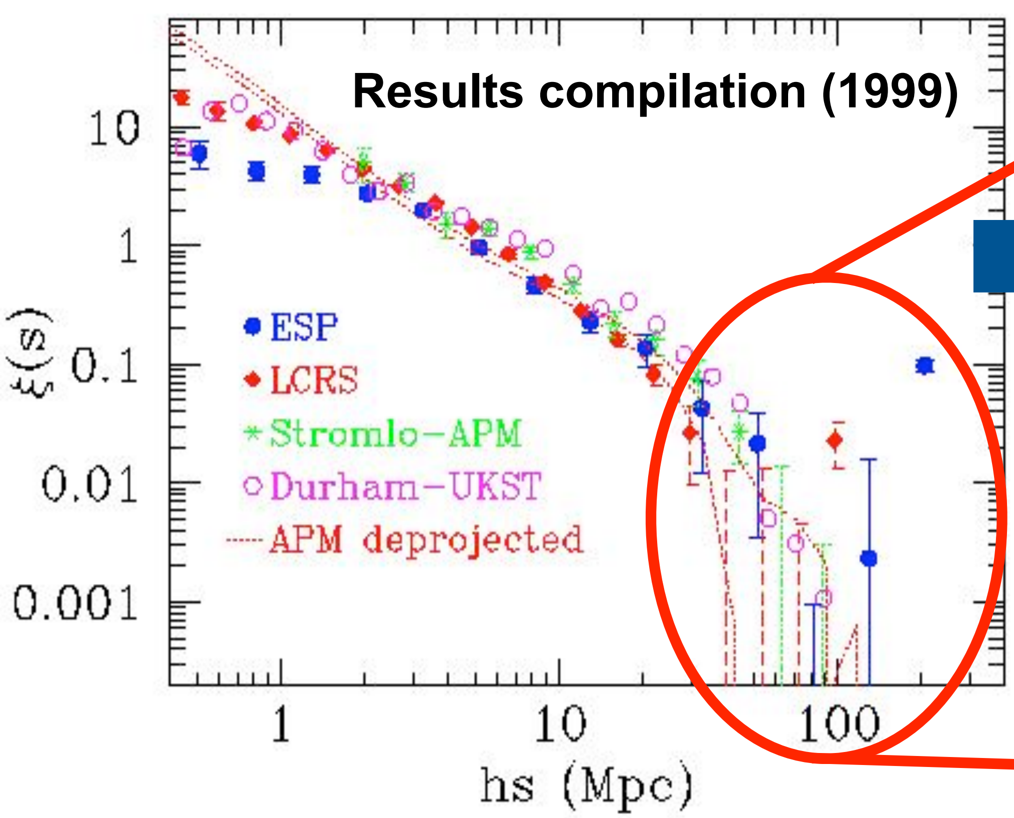

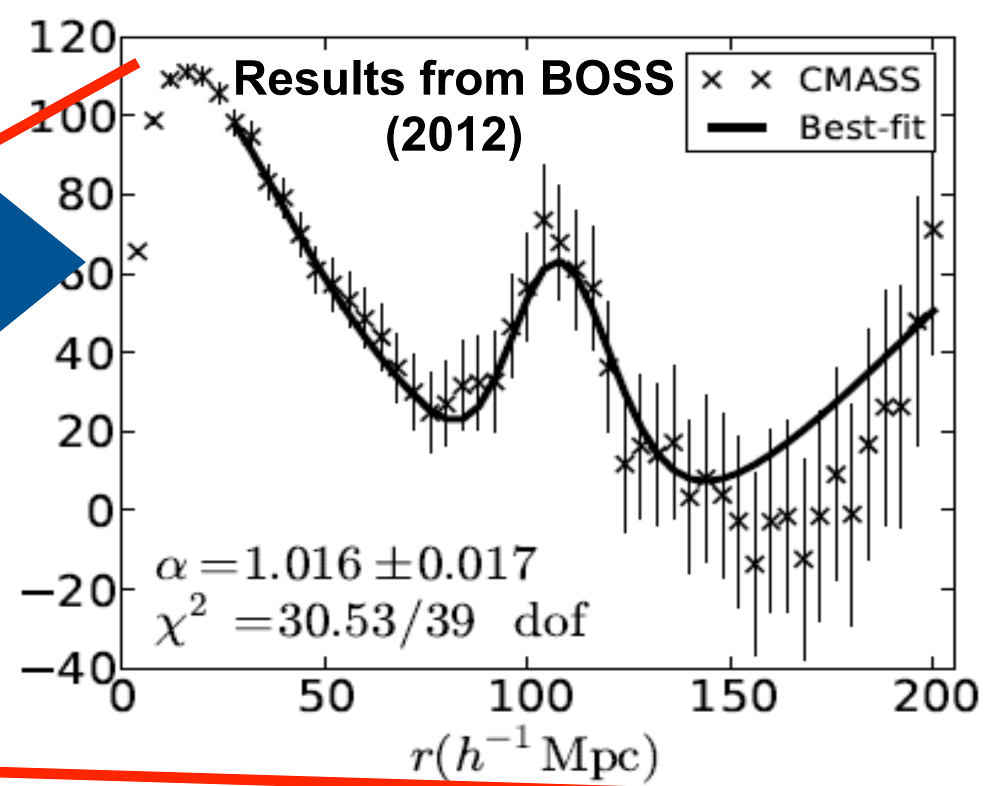

- Massive increase in volume of galaxy surveys
- Next-generation galaxy clustering simulations
- Multi-physics codes needed to meet accuracy requirements

Nature of Cosmological Simulations

- **Foundations and "Smoking Guns":** Qualitative (in some sense) rather than quantitative predictions:
  - Dark matter annihilation signals
  - Halo substructure
  - Features in the primordial power spectrum
  - Foundations of structure formation
  - Modified gravity theories
- **Cosmological Astrophysics:** Melding of cosmology and astrophysics problems
  - Galaxy formation
  - Cluster astrophysics
  - Galaxy-halo connection
  - Black hole formation
- **Precision Cosmology:** Simulations with physics, accuracy, and statistical control matched to precision cosmological probes:
  - Gravitational lensing
  - Galaxy clustering
  - Cluster cosmology
  - Ly-alpha forest
  - Neutrino masses
  - Primordial non-Gaussianity
  - Cross-correlations
  - Investigations of systematics
  - 21cm and other intensity mapping surveys

Cosmology: Simulation Frontiers

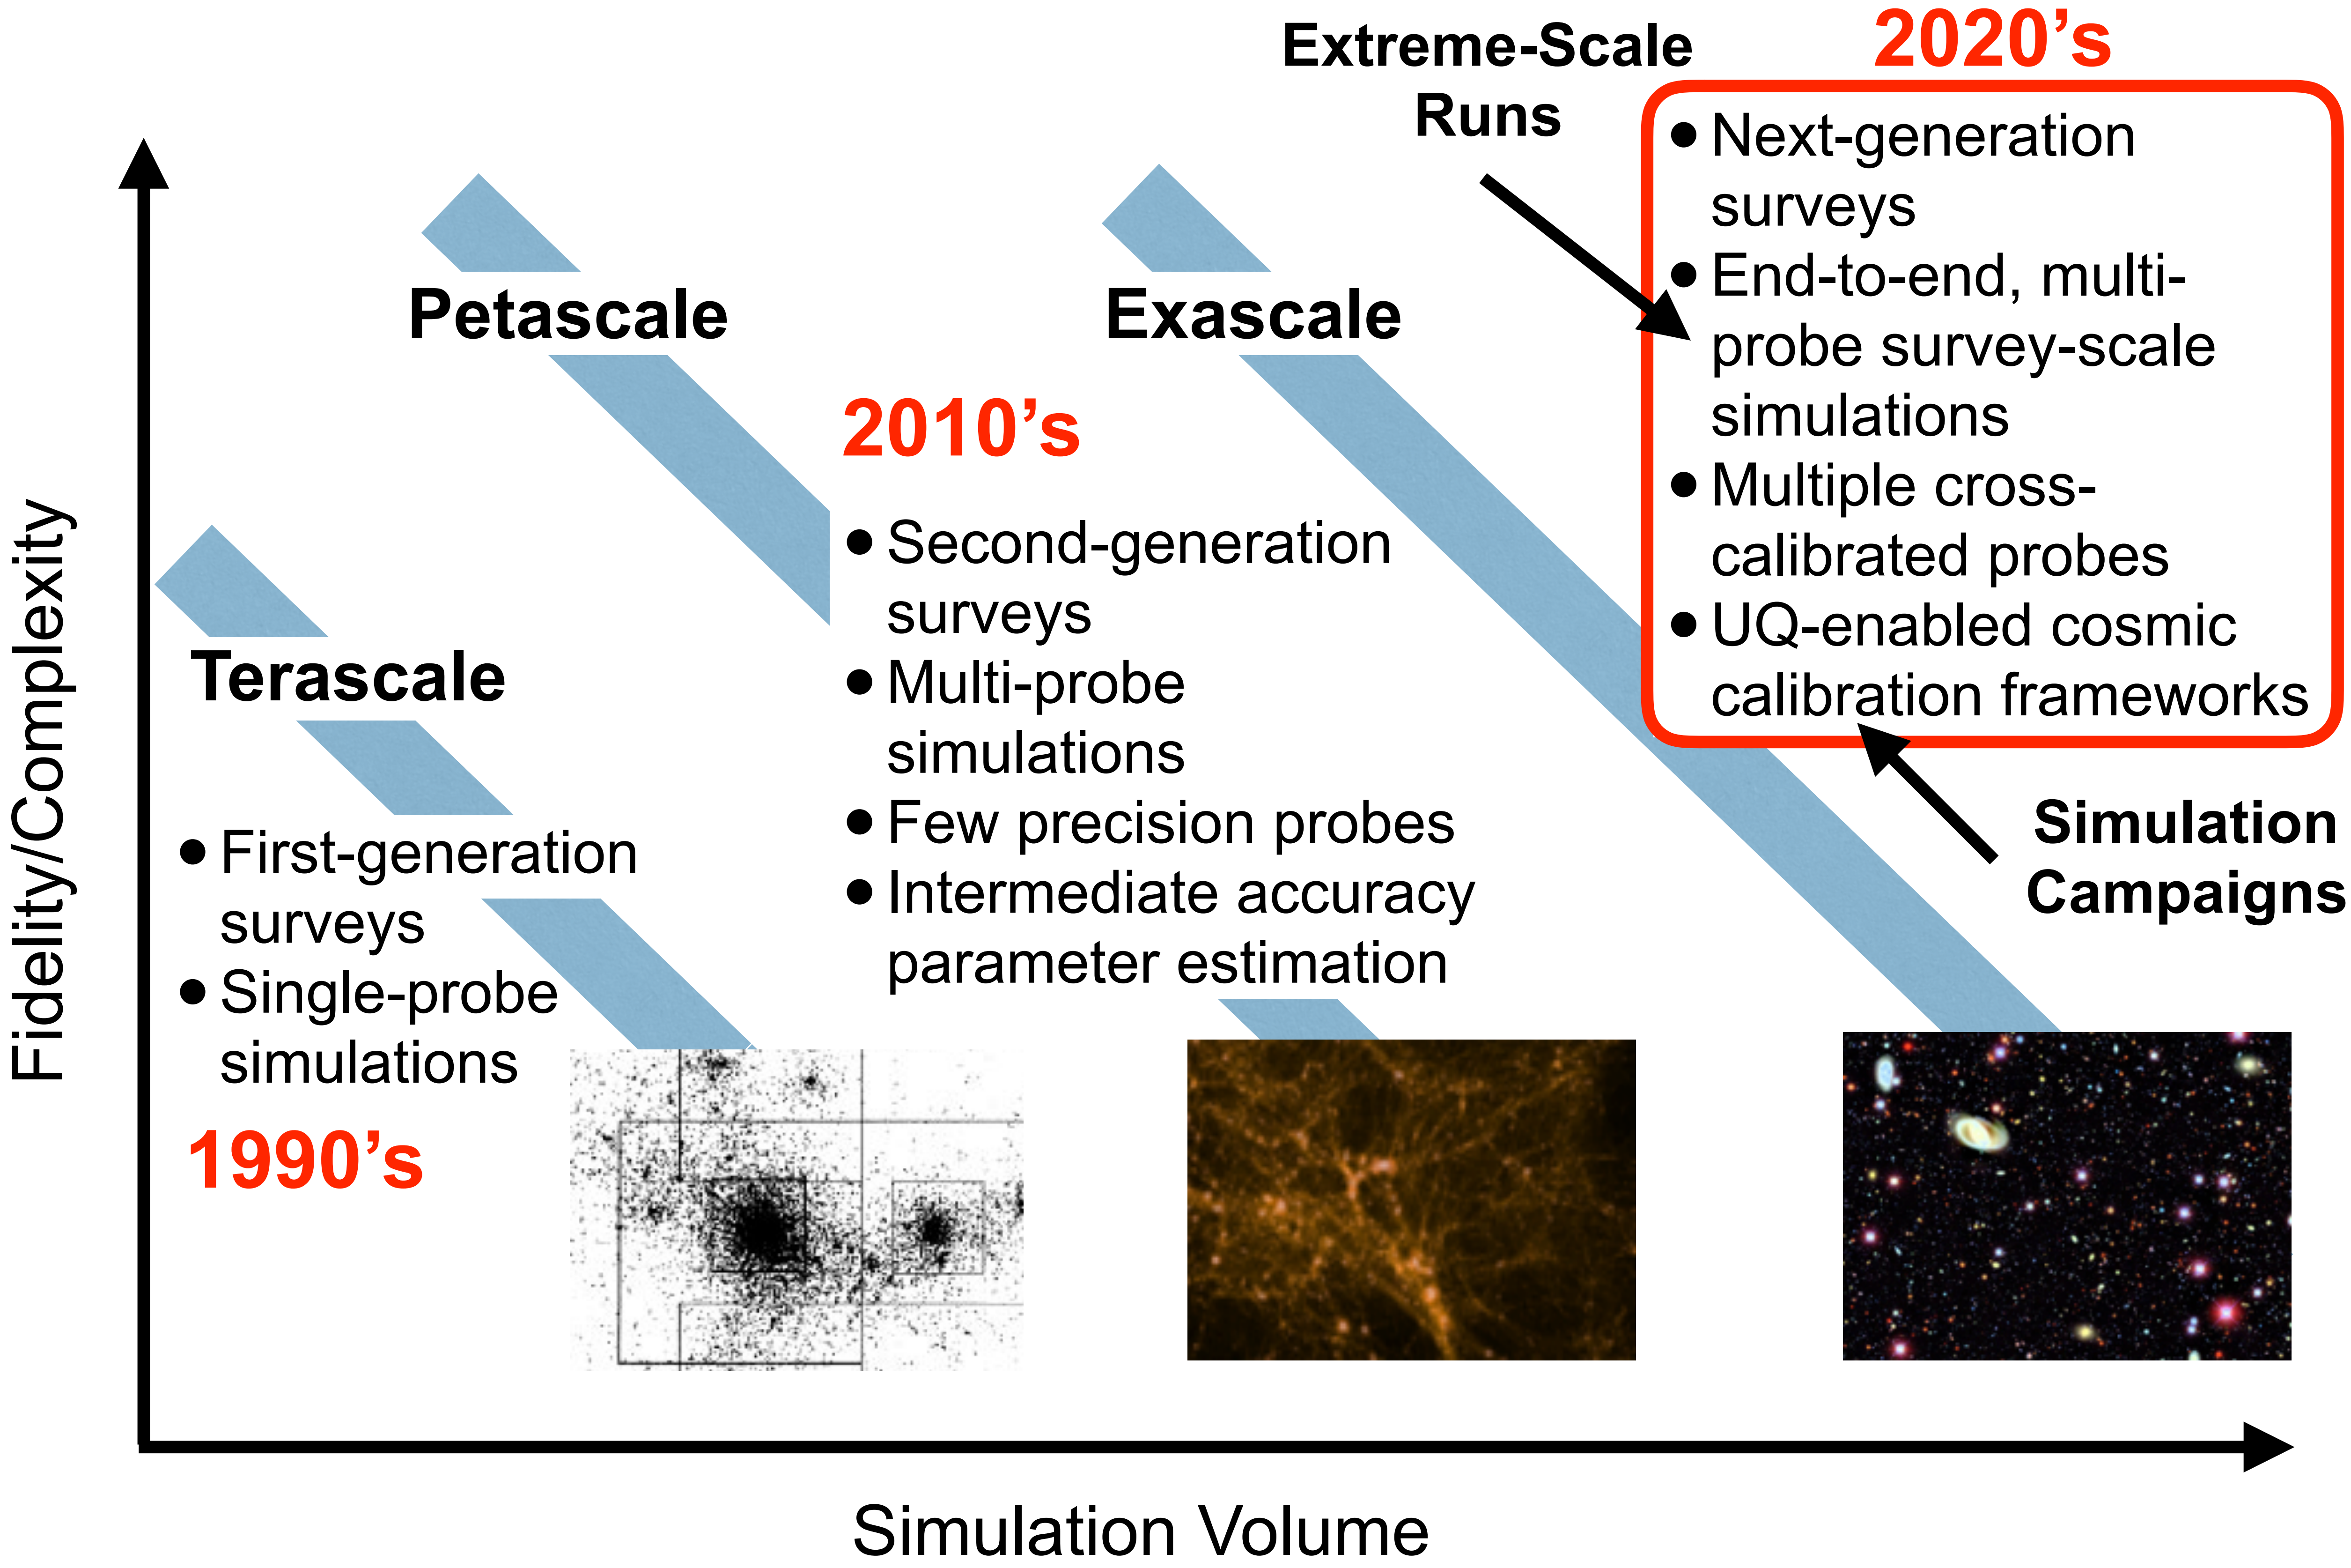

ExaSky: Challenge Problem

- **Exascale Imperatives**
  - Sky surveys can measure cosmological statistics at the **1%** level or better — including cross-correlations
  - Scientific inference seriously limited by shortcomings in theoretical modeling (nonlinear gravitational effects, baryonic effects on the matter distribution, —)
  - Scales of interest cover ~1-10s of Mpc
    - Galaxy spatial and velocity bias
    - Baryonic effects on the matter power spectrum (cosmic shear)
    - kSZ and tSZ simulations
    - Ly-alpha forest
    - CMB lensing/delensing
    - Cross-correlations

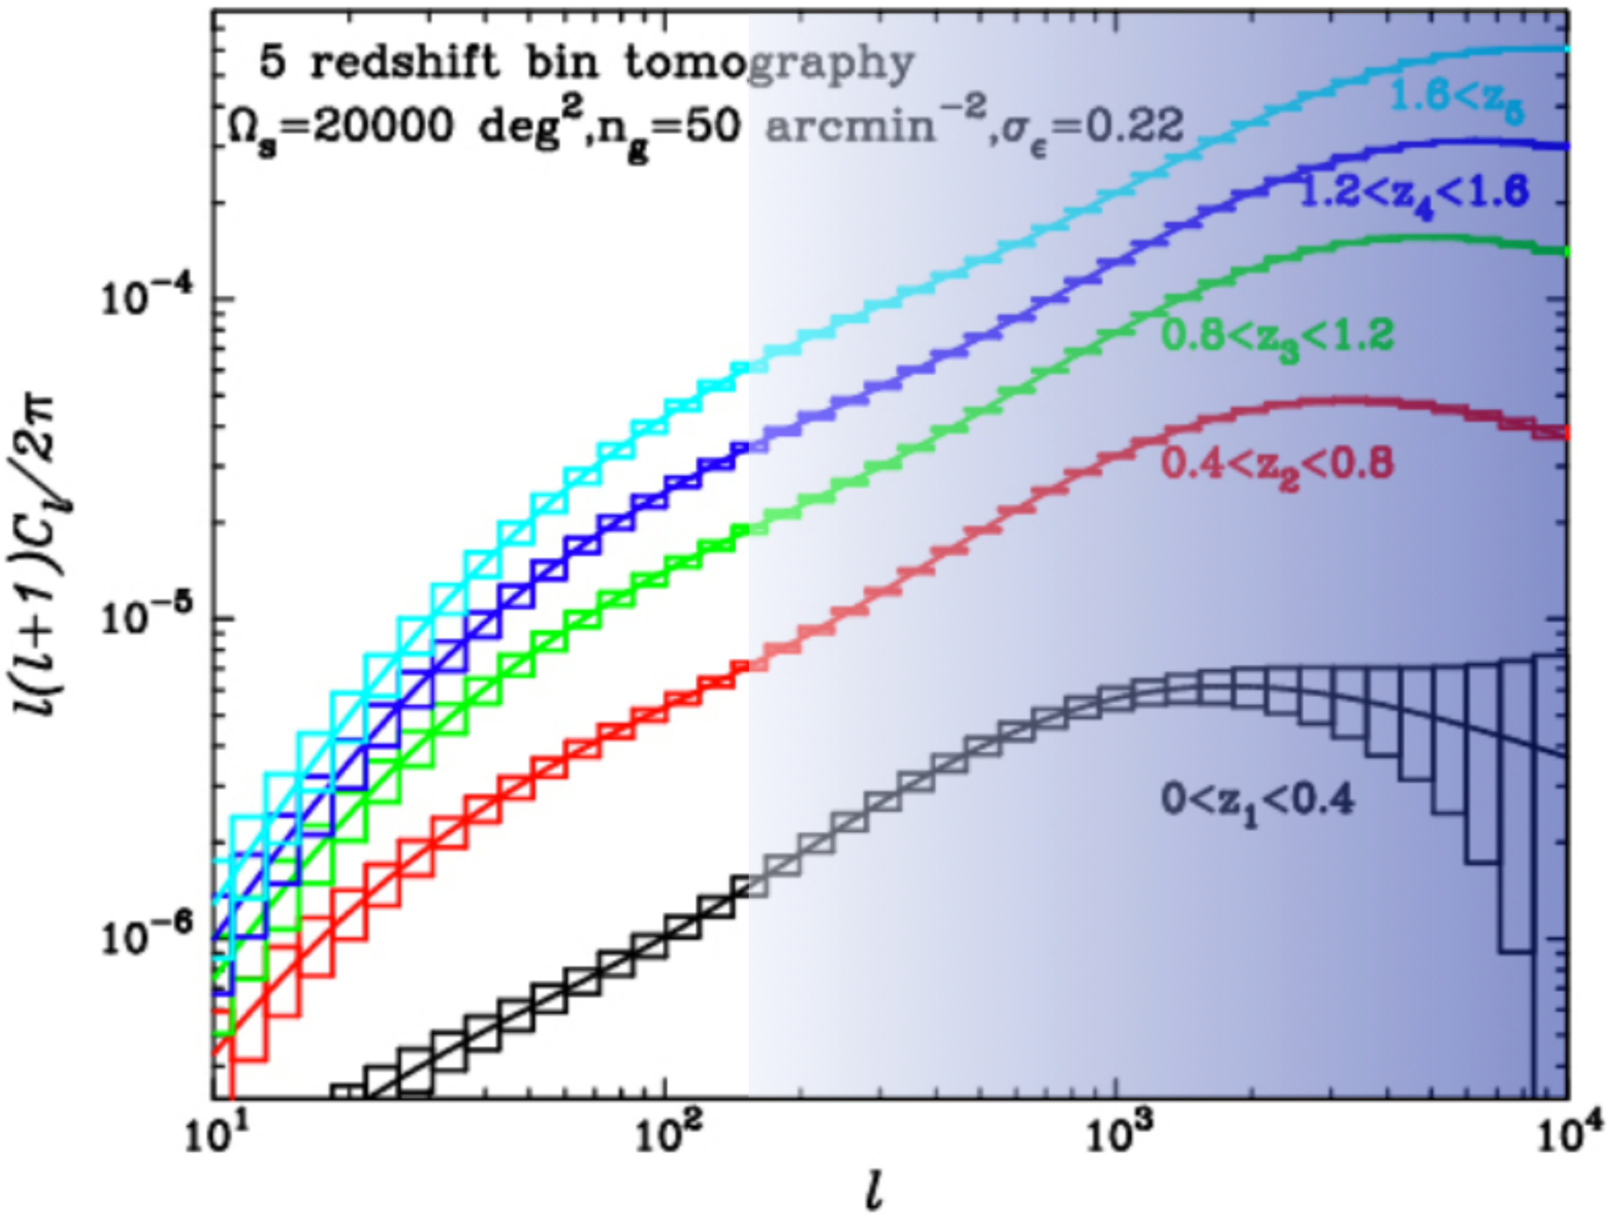

Forecast for LSST tomographic weak lensing shear power spectrum (Takada et al 2005); the nonlinear contribution (dark band) starts roughly around l~100, quickly becoming dominant

An Early Simulation

- Suite of 300 (and less) particle simulations
- Run on a CDC 3600, ~1Mflops, 32KB+ at LANL
- Is twelve orders of magnitude improvement in both performance and memory good enough for precision cosmology?

THE ASTRONOMICAL JOURNAL VOLUME 75, NUMBER 1 FEBRUARY

Structure of the Coma Cluster of Galaxies*

P. J. E. Peebles†
Palmer Physical Laboratory, Princeton University, Princeton, New Jersey
(Received 7 October 1969)

In some cosmologies, a cluster of galaxies is imagined to be a gravitationally bound system which, in analogy with the formation of the Galaxy, originated as a collapsing protocluster. It is shown that a numerical model based on this picture is consistent with the observed features of the Coma Cluster of galaxies. The cluster mass derived from this model agrees with previous values; however, an analysis of the observational uncertainty within the framework of the model shows that the derived mass could be consistent with the estimated total mass provided by the galaxies in the cluster.

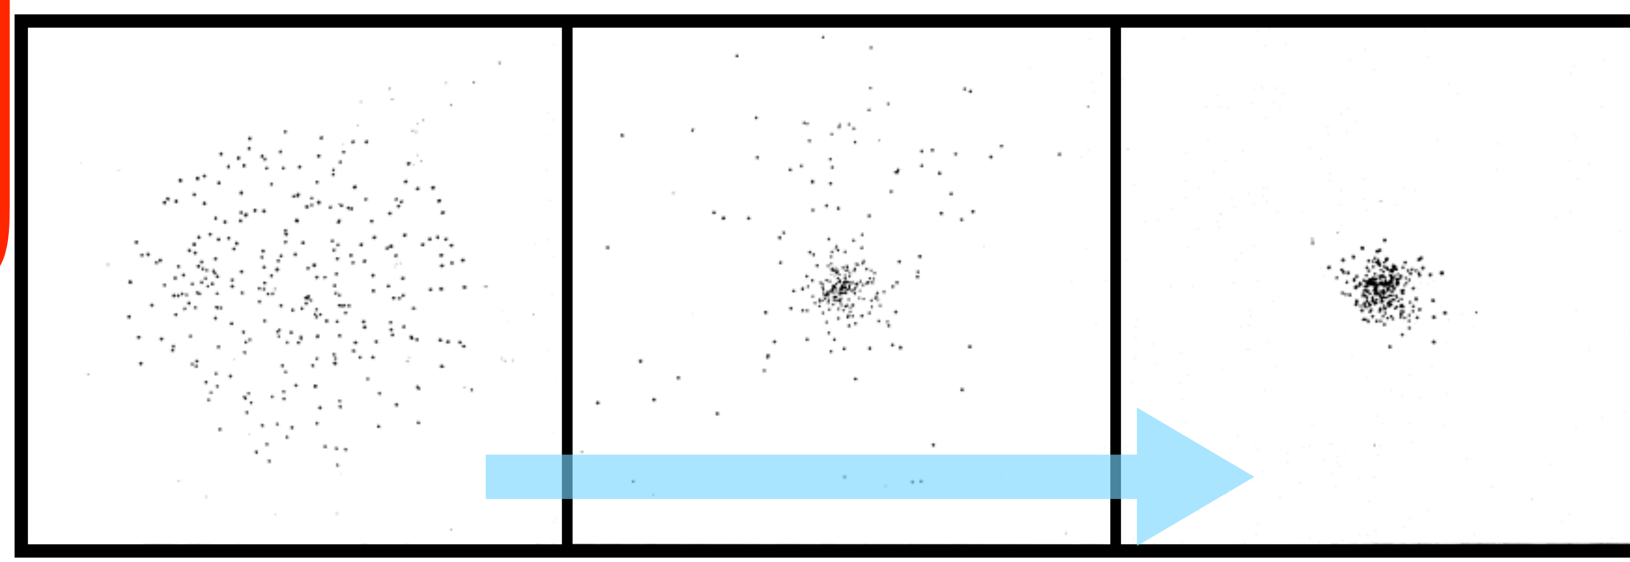

"The Universe is far too complicated a structure to be studied deductively, starting from initial conditions and solving the equations of motion."

Robert Dicke (Jayne Lectures, 1969)

Cosmological Simulations: Current Status

- **Gravity-Only:** The situation is still not as good as one would like; it would be great if all codes agreed to ~0.1% for a "nice" observable like the power spectrum
  - Most simulation issues appear to be understood, but hand-tuning still exists
- **Hydro Sims:** The situation is getting better as improved SPH, AMR, and moving mesh codes are compared
  - Agreement on test cases still an order of magnitude worse than gravity-only
  - Subgrid models are (mostly) uncontrolled

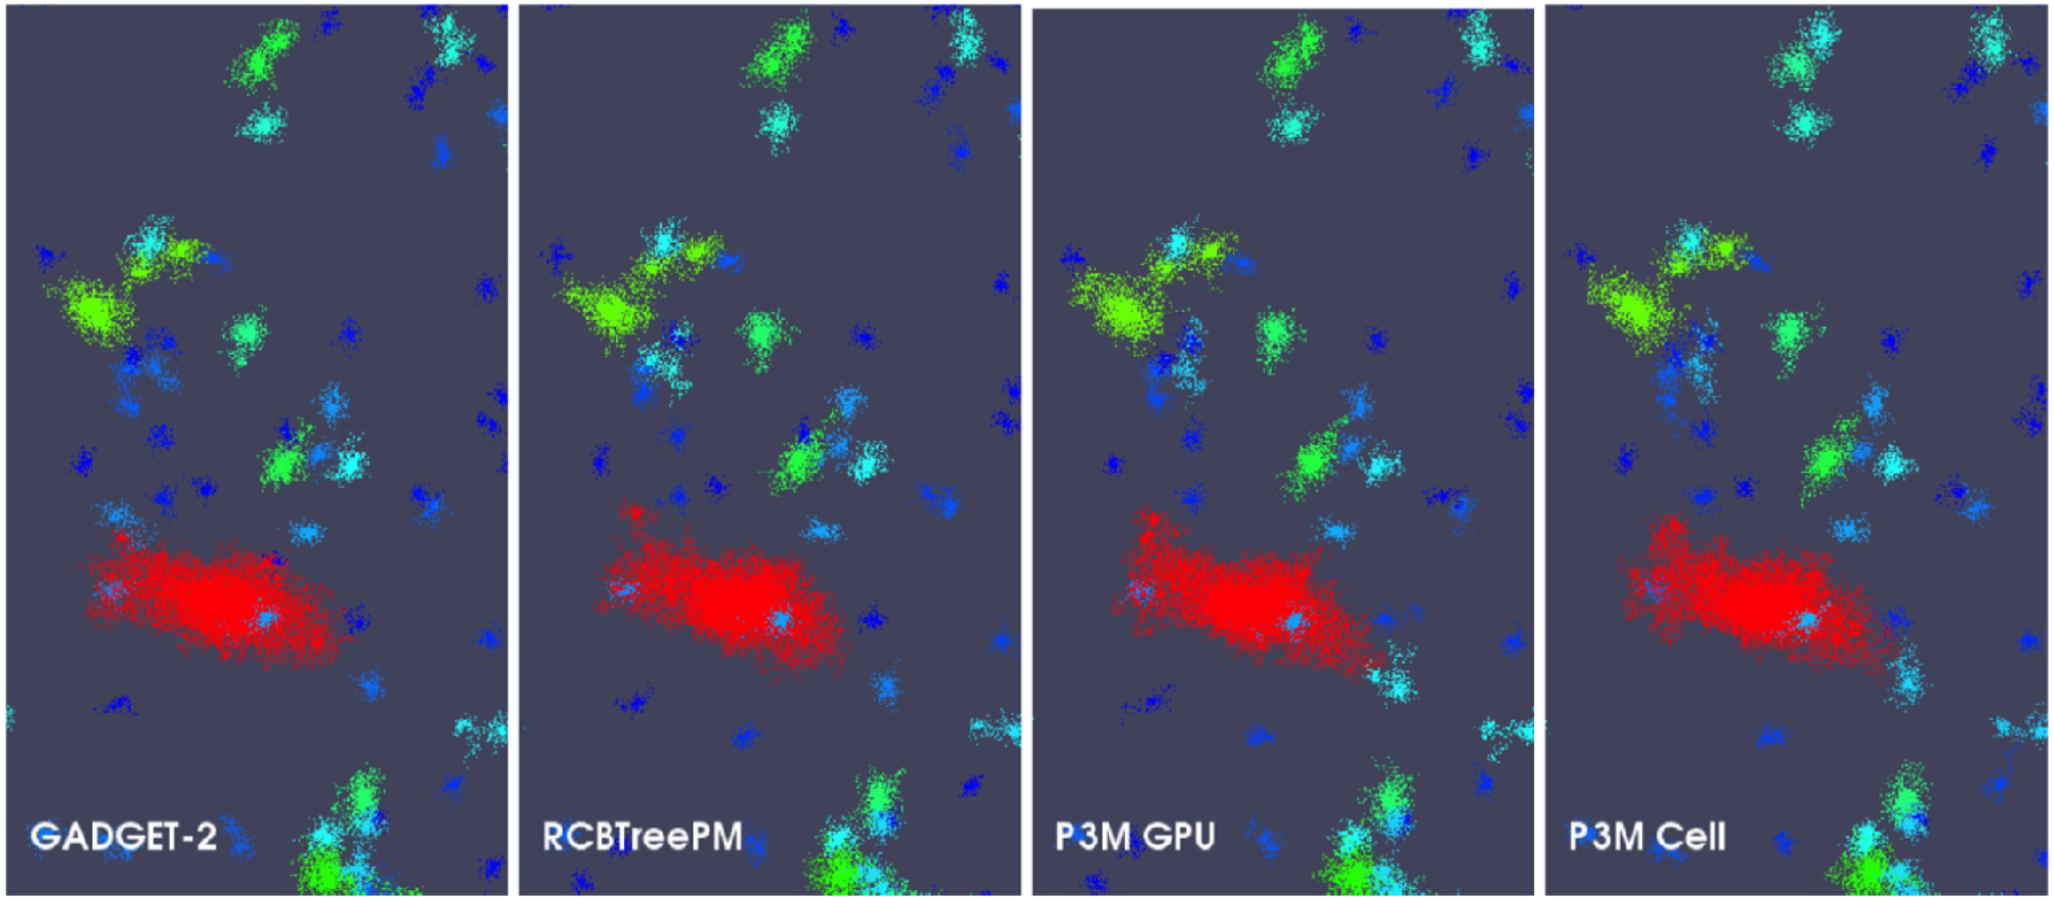

Comparison of multiple algorithms on the same problem, halos are color-coded according to the number of particles (blue~100, green~1000, red~10,000

Evolution of Computing Power

- **Supercomputer Evolution:**
  Supercomputing has come a very long way since 1988, here are some numbers for the Cray-YMP:
  - 8 nodes, 2.67 GFlops (peak)
  - Total RAM, 512MB + 4GB SSD

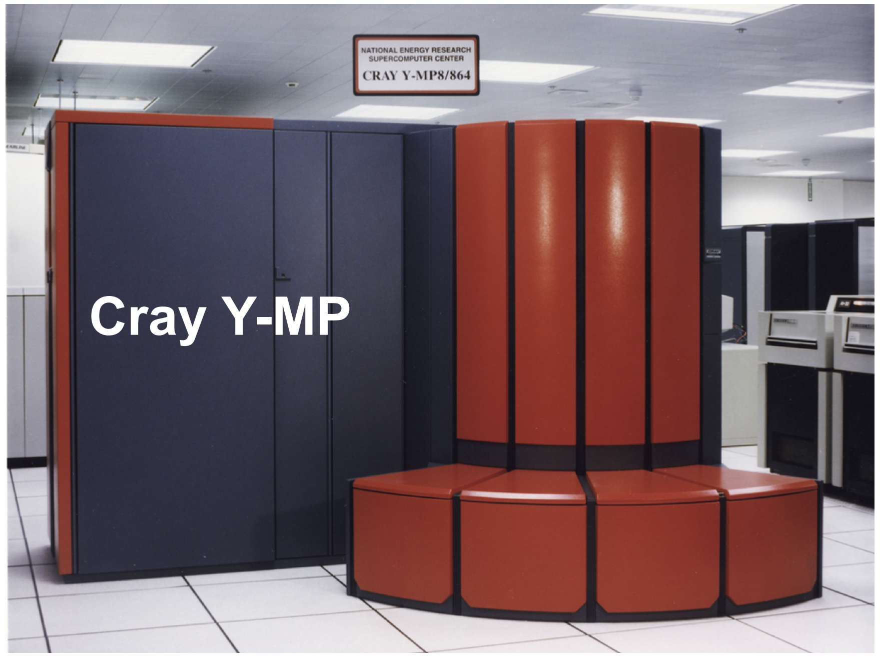

- **Example Current System:** Summit, CPU+GPU
  - 4608 nodes, 194 PFlops (peak)
  - Total RAM, 2.8 PB
  - In Dec 2018, three ~2 trillion particle simulations were run on Summit using HACC in one week

And exascale is on its way (2021/22) — machines ***a billion times faster*** wrt 1988!

Issues with Next-Generation Architectures

Problems:

- Next-gen architecture characteristics
  - ***designed for massive (local) concurrency***
  - ***multiple memory hierarchies***
  - ***difficult to program***
  - ***high flops to byte ratio desired***
  - ***future is very hard to predict***
- How to write a code you don't have to keep rewriting? (And not rely on somebody else?)
- Systems are designed for crunching the simulation but not analyzin the data, this makes end-to-end modeling activities difficult to carry out
- The US HPC community has possibly a billion lines of code — dealing with next-gen computing is a ***problem faced by everyone***

Precision Cosmology Context

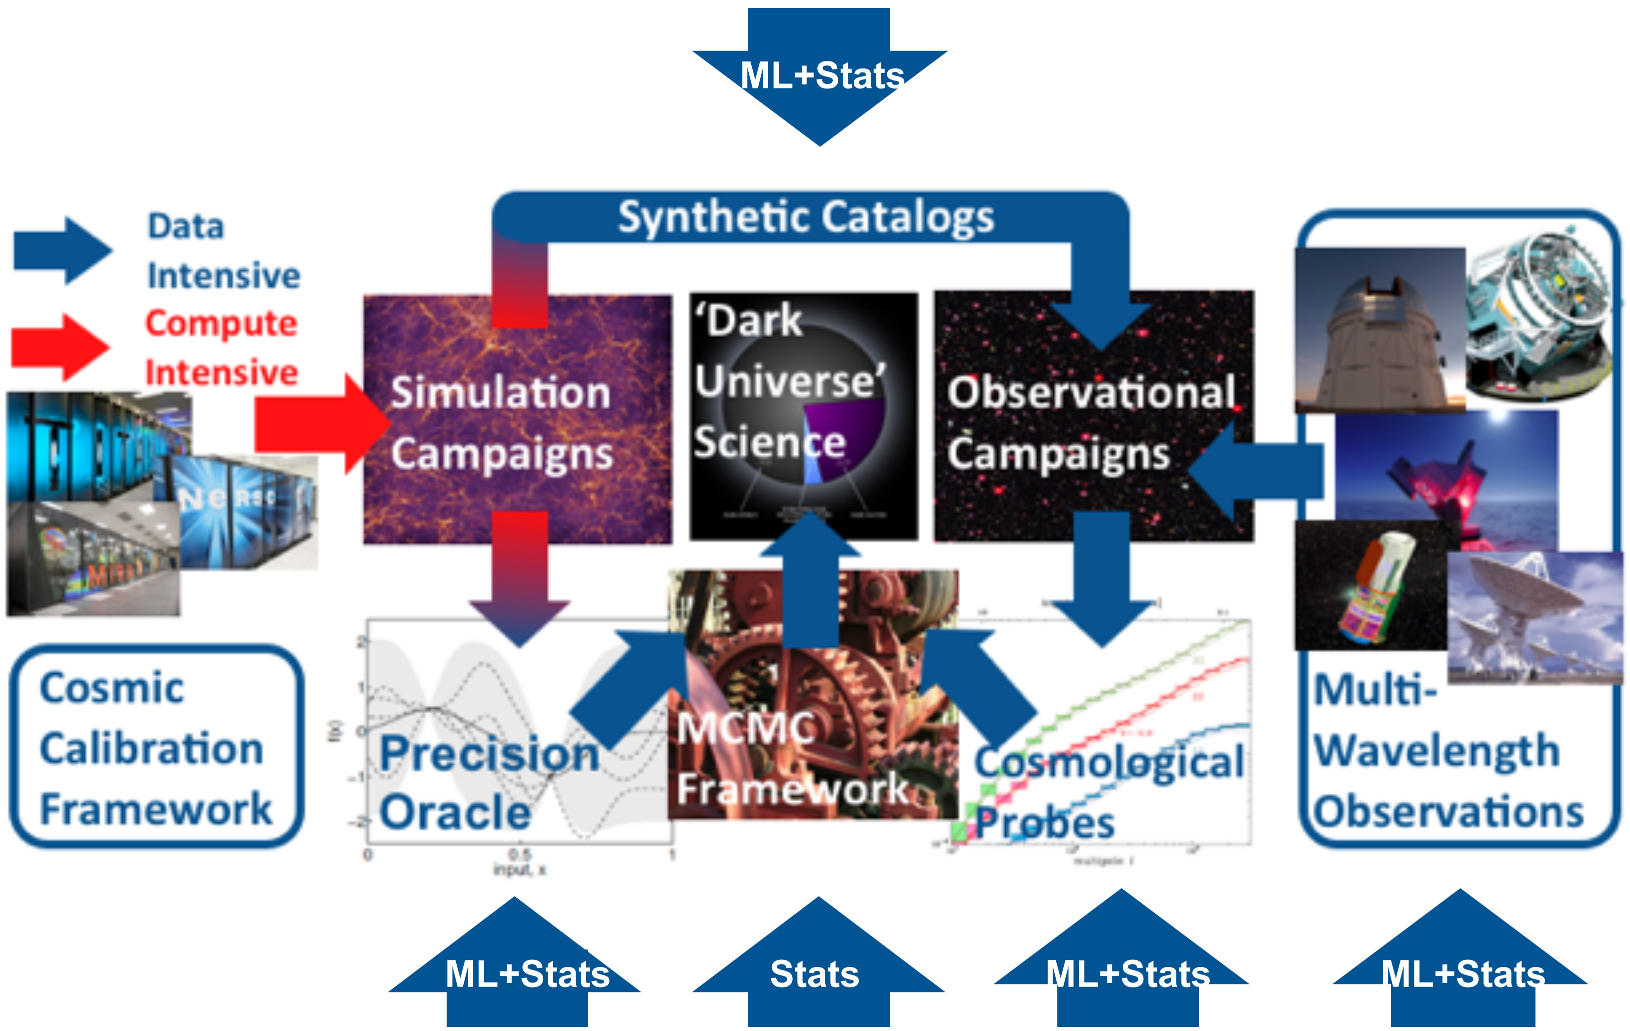

- **Cosmological Probes:** Distribution of mass and light to study the 'Dark Universe'
- **Modeling + Observation:** Large-scale data-intensive computing
- **'AI at Scale':** Need to speed up current state of the art by ***many orders of magnitude*** to enable dealing with datasets in the exascale era
- **Applications:** Image classification, gravitational lens characterization, fast sky catalog/image generation, fast prediction of summary statistics, fast likelihood estimation, cosmological parameter estimation, —

End-To-End Workflow Complexity

Gaussian Random Field Initial Conditions → High-Resolution N-Body (+ Hydro) simulations → Multiple Outputs Halo/Sub-Halo Identification → Halo Merger Trees

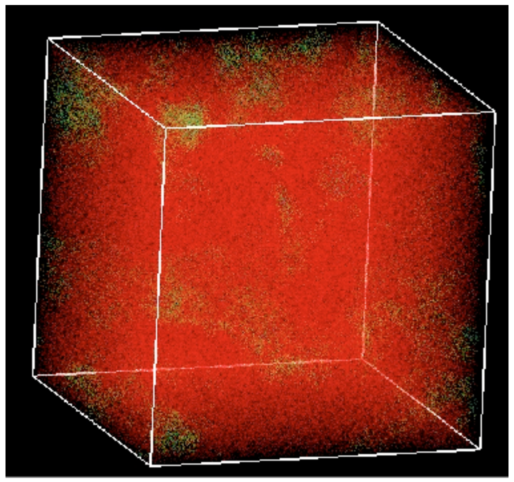

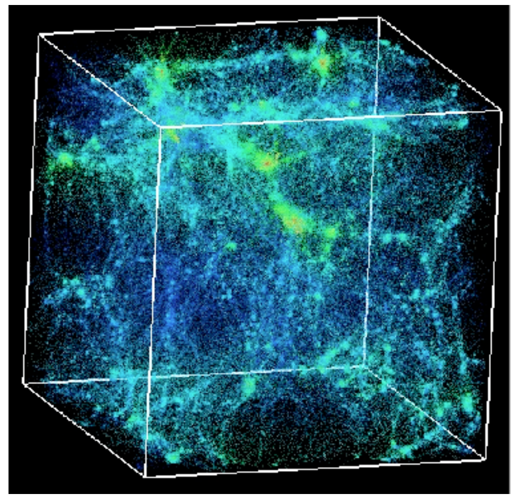

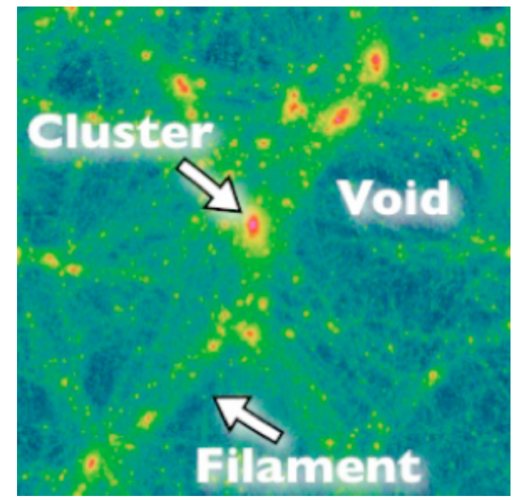

Semi-Analytic + Empirical Galaxy Modeling

Galaxy Catalog

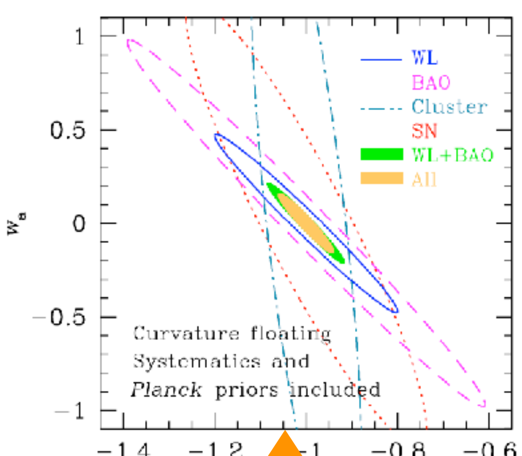

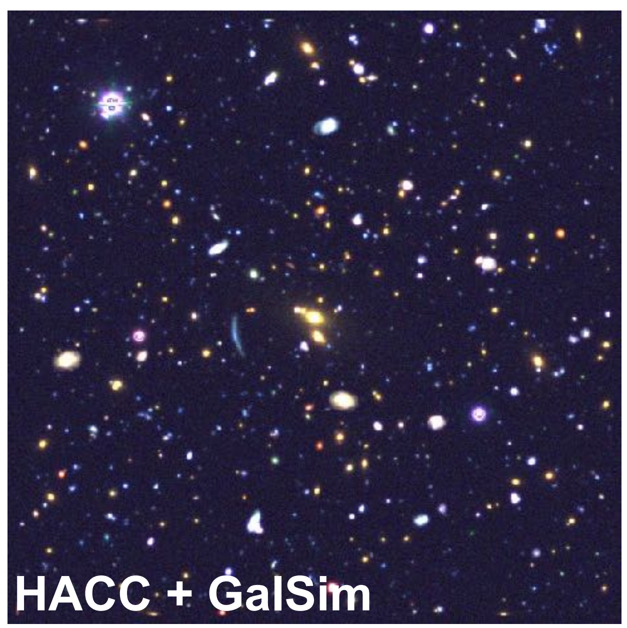

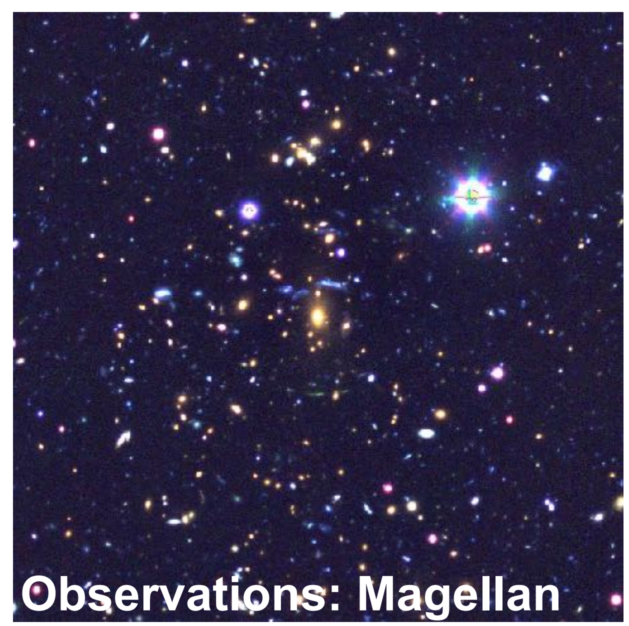

Realistic Image Catalog

Atmosphere and Instrument Modeling → Data Management Pipeline → Data Analysis Pipeline → Scientific Inference Framework

Simulating Surveys: "PIC+MD"

- **Simulation Volume:** Large survey sizes impose simulation volumes **~(4** Gpc)^3, memory required **~100**TB **- 1**PB
- **Number of Particles:** Mass resolution depends on ultimate object to be resolved, $\sim 10^8 \mathrm{M}_\odot - 10^{10} \mathrm{M}_\odot$ (subhalos to halos), N~$10^{11} - 10^{12}$
- **Easy to remember:** ~(1000 Mpc)^3 and ~(1000)^3 particles lead to a mass resolution of $\sim 10^{11} \mathrm{M}_\odot$
- **Therefore:** ~(10,000)^3 particles in ~(1000 Mpc)^3 leads to $\sim 10^8 \mathrm{M}_\odot$, (5Gpc)^3 at this mass resolution would require 125 trillion particles = 5PB per snapshot! (Largest systems currently have ~2PB of RAM)
- **Force Resolution:** ~kpc, yields a (global) spatial dynamic range of **~10^6**

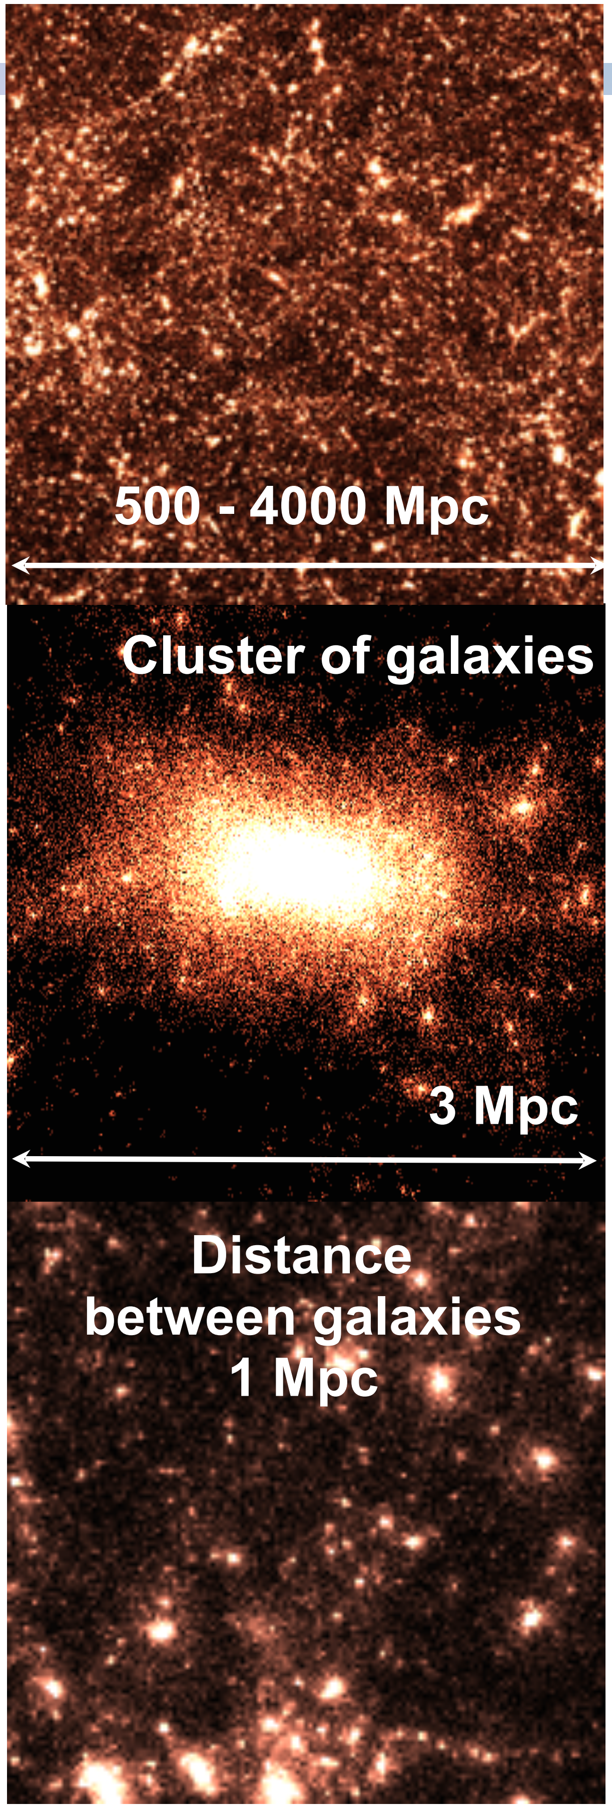

Meeting the Challenge: HACC

- **Pre-History:** MC2, a parallel PM cosmology code written by modifying a beam physics Vlasov-Poisson solver
- **History:** December 2008, Los Alamos decides to build the world's first PFlops system (Roadrunner, based on CPU + IBM Cell)
- **Roadrunner Universe Project:** We design and implement a high-resolution, hybrid, accelerated (Spectral P3M) gravity-only cosmology code
- **HACC on the BG/Q and Titan:** Code modified to run at extreme performance levels (first production code to break 10 PFlops sustained performance); 'Outer Rim' and 'Q Continuum' simulations
- **CRK-SPH:** New SPH scheme that solves many known problems of SPH developed and integrated within HACC
- **HACC on Summit:** Three of the world's largest simulations done in a week (~2 trillion particles each) during system acceptance

HACC on Next-Generation Systems: ECP

- **ECP:** ExaSky — cosmological simulation component of the US DOE Exascale Computing Project
- **Codes:** HACC and Nyx (AMR hydro) aimed for challenge problems on the 2021/2022 timescale
- **Exascale Systems:** Aurora at Argonne, Frontier at Oak Ridge (pre-exascale systems — Summit at Oak Ridge and Perlmutter at LBNL)
- **Exascale/Post-Exascale Environment:** Performance in the 10's of Eflops, file systems in the EB range, major AI/ML capabilities
- **Challenge Problem:** Perform multi-Gpc N-body and hydro simulations with robust control of errors and subgrid modeling consistent with code resolution

Public Specs on Aurora and Frontier:

- Performance: >1EF
- System interconnect: Cray Slingshot
- Compute node: CPU+GPU (1:4; 1:3)
- System Memory: ~10PB
- Storage: >200PB at ~10TB/s
- # Cabinets: ~100-200

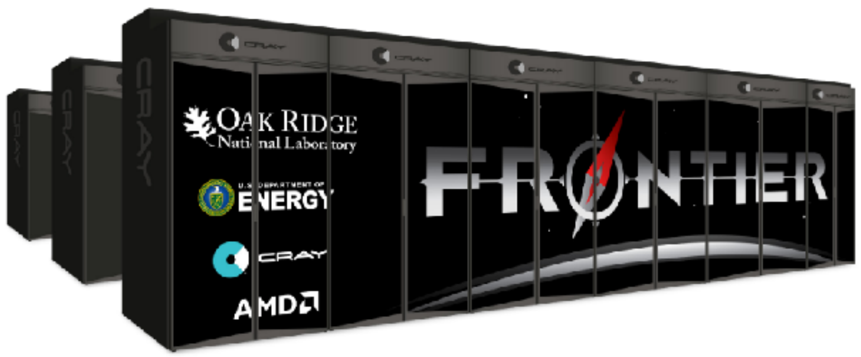

HACC In Pictures (“Million to One” Dynamic Range)

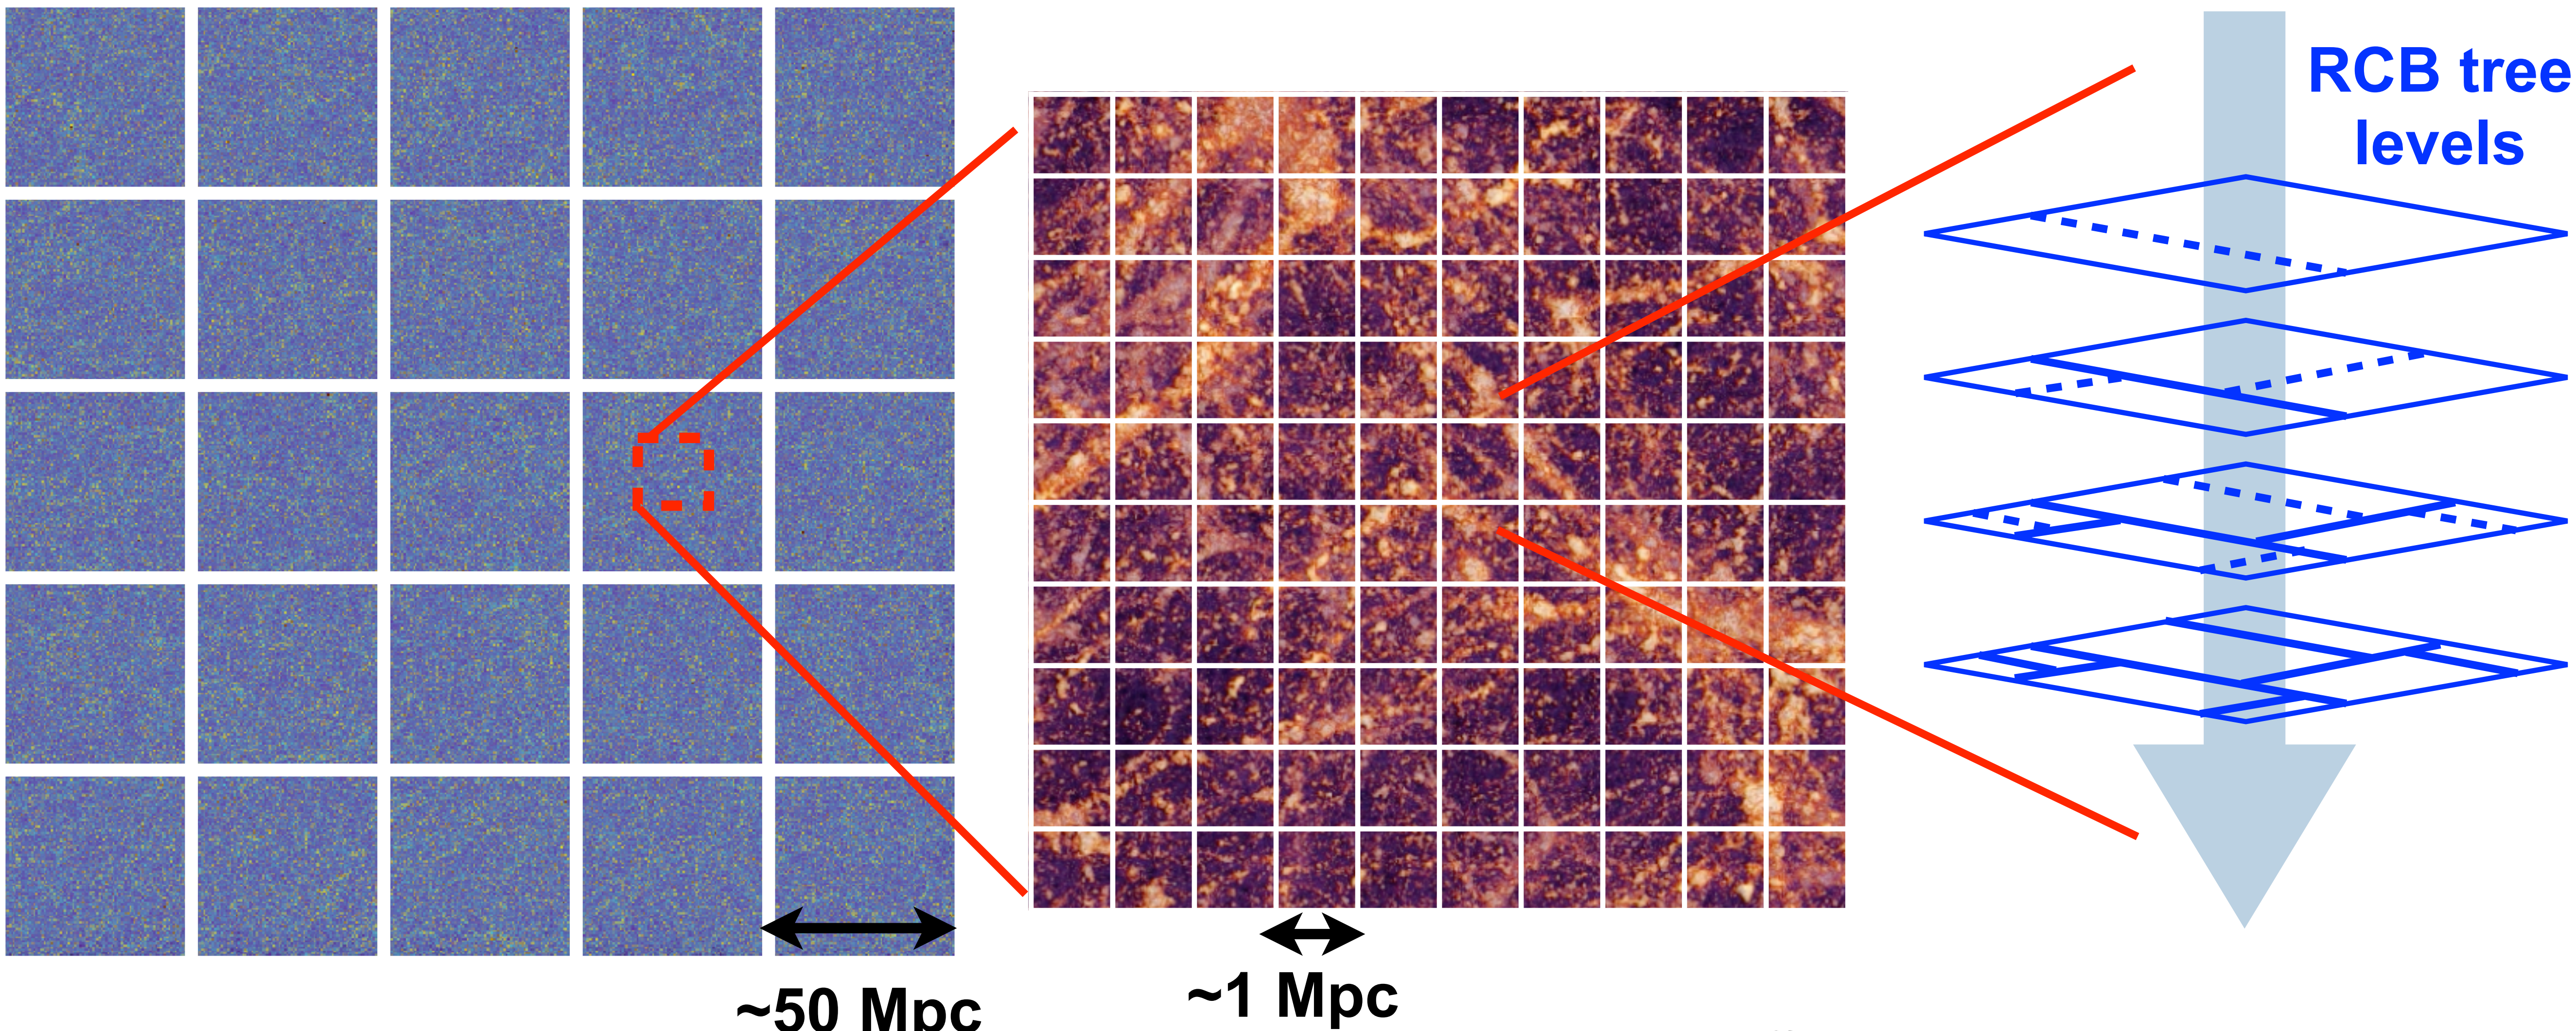

HACC Top Layer:
3-D domain decomposition with particle replication at boundaries (‘overloading’) for Spectral PM algorithm (long-range force)

Host-side

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

GPU (compute-intensive): two options, P3M vs. TreePM

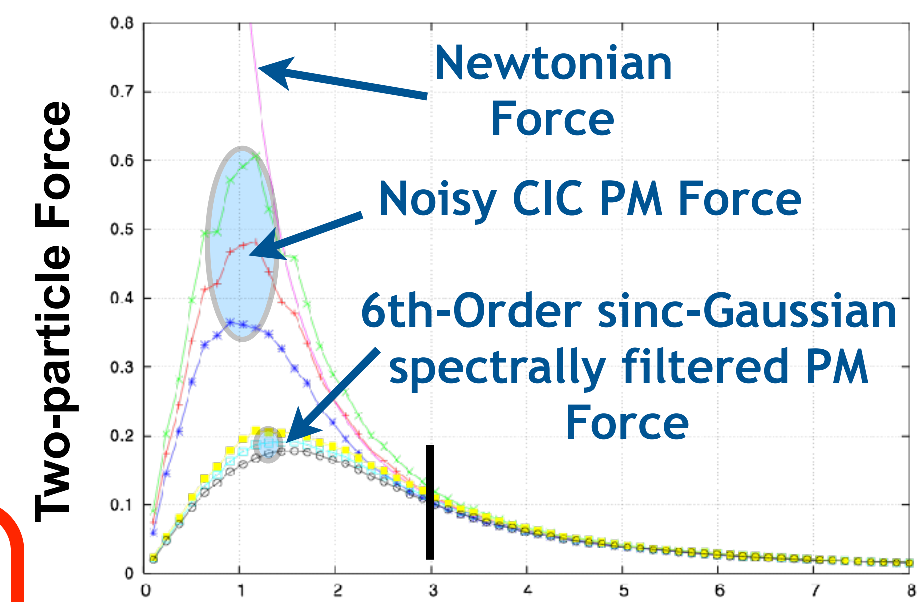

HACC: Algorithmic Features

- **Gravity — Hybrid Grid/Particle:** 6-th order spectral Poisson solver; 4-th order super-Lanczos spectral derivatives; short-range forces matched via spectral filters (high-accuracy polynomial fits), custom parallel 3D FFT (SWFFT)
- **Gasdynamics — CRK-SPH:** Higher-order SPH scheme solves known SPH issues in dealing with mixing, tracking instabilities, etc.
- **Chaining Mesh/Local Trees:** Data structures optimize local force solvers (tree/fast multipole/P3M)
- **Adaptive (Symplectic) Time-Stepping:** 2nd-order split-operator method; sub-cycling based on the RCB tree depth; implicit solver for subgrid model

Habib et al., New Astron. 42, 49 (2016); Comm. ACM 60, 97 (2017) [Research Highlight]

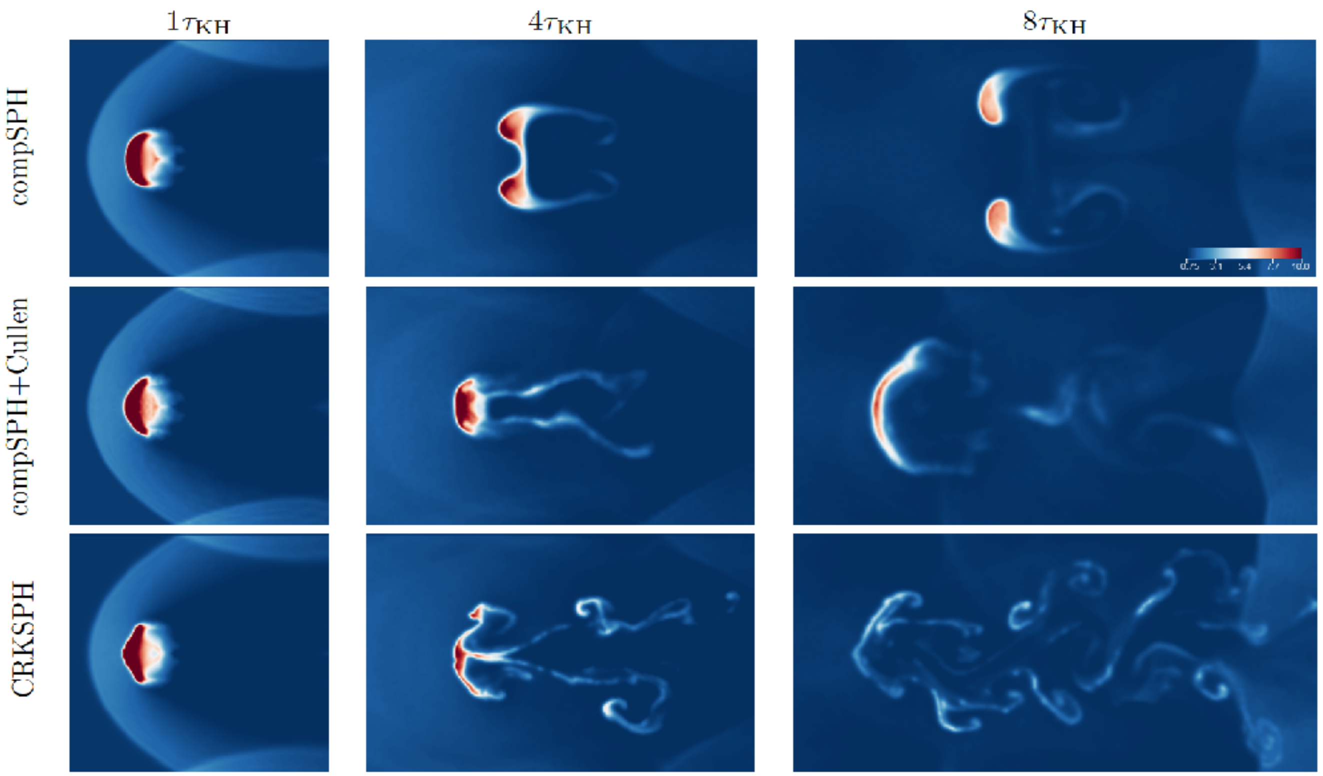

Frontiere et al., JCP 332, 160 (2017)

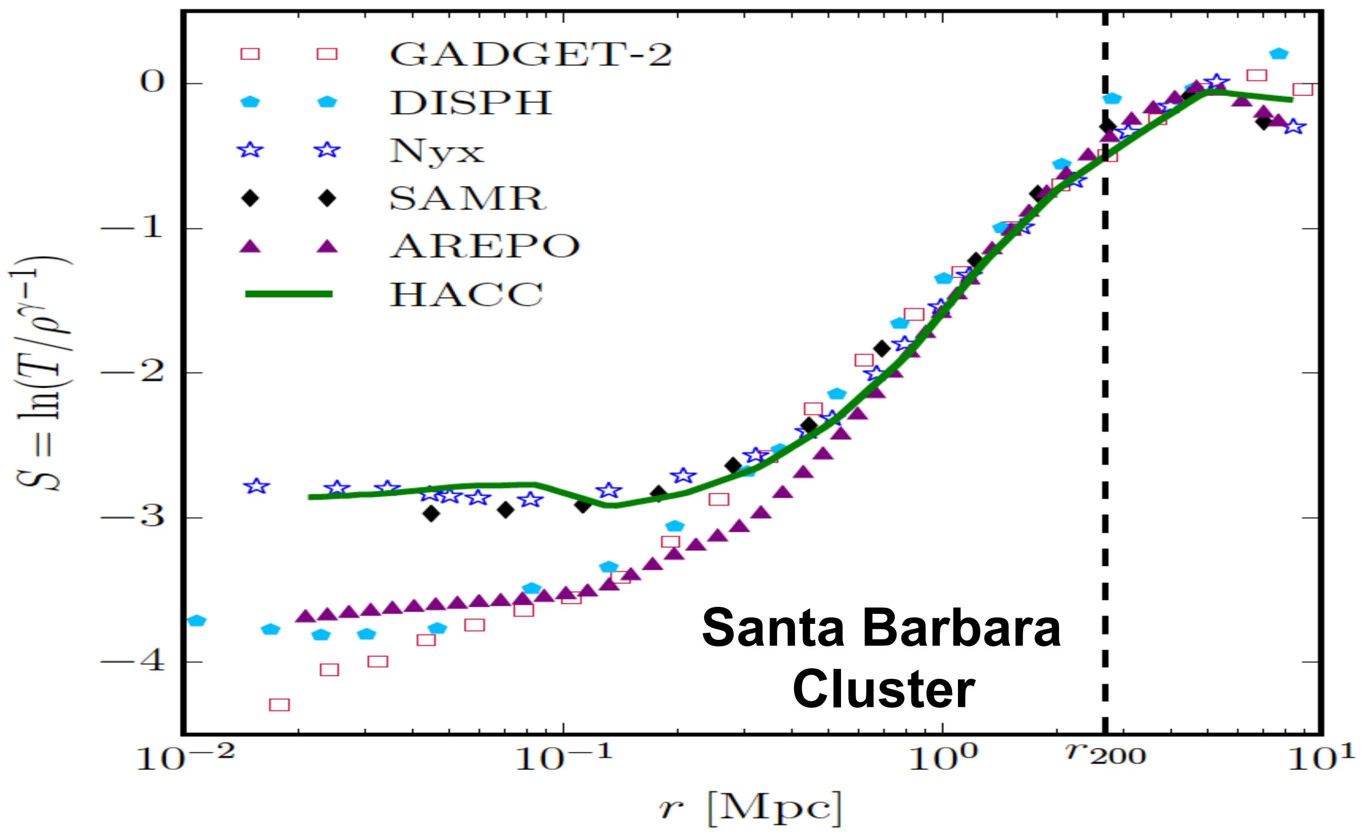

PICS: Generating Strong Lensing "Observations"

PICS: Pipeline for Images of Cosmological Strong lensing

Simulated

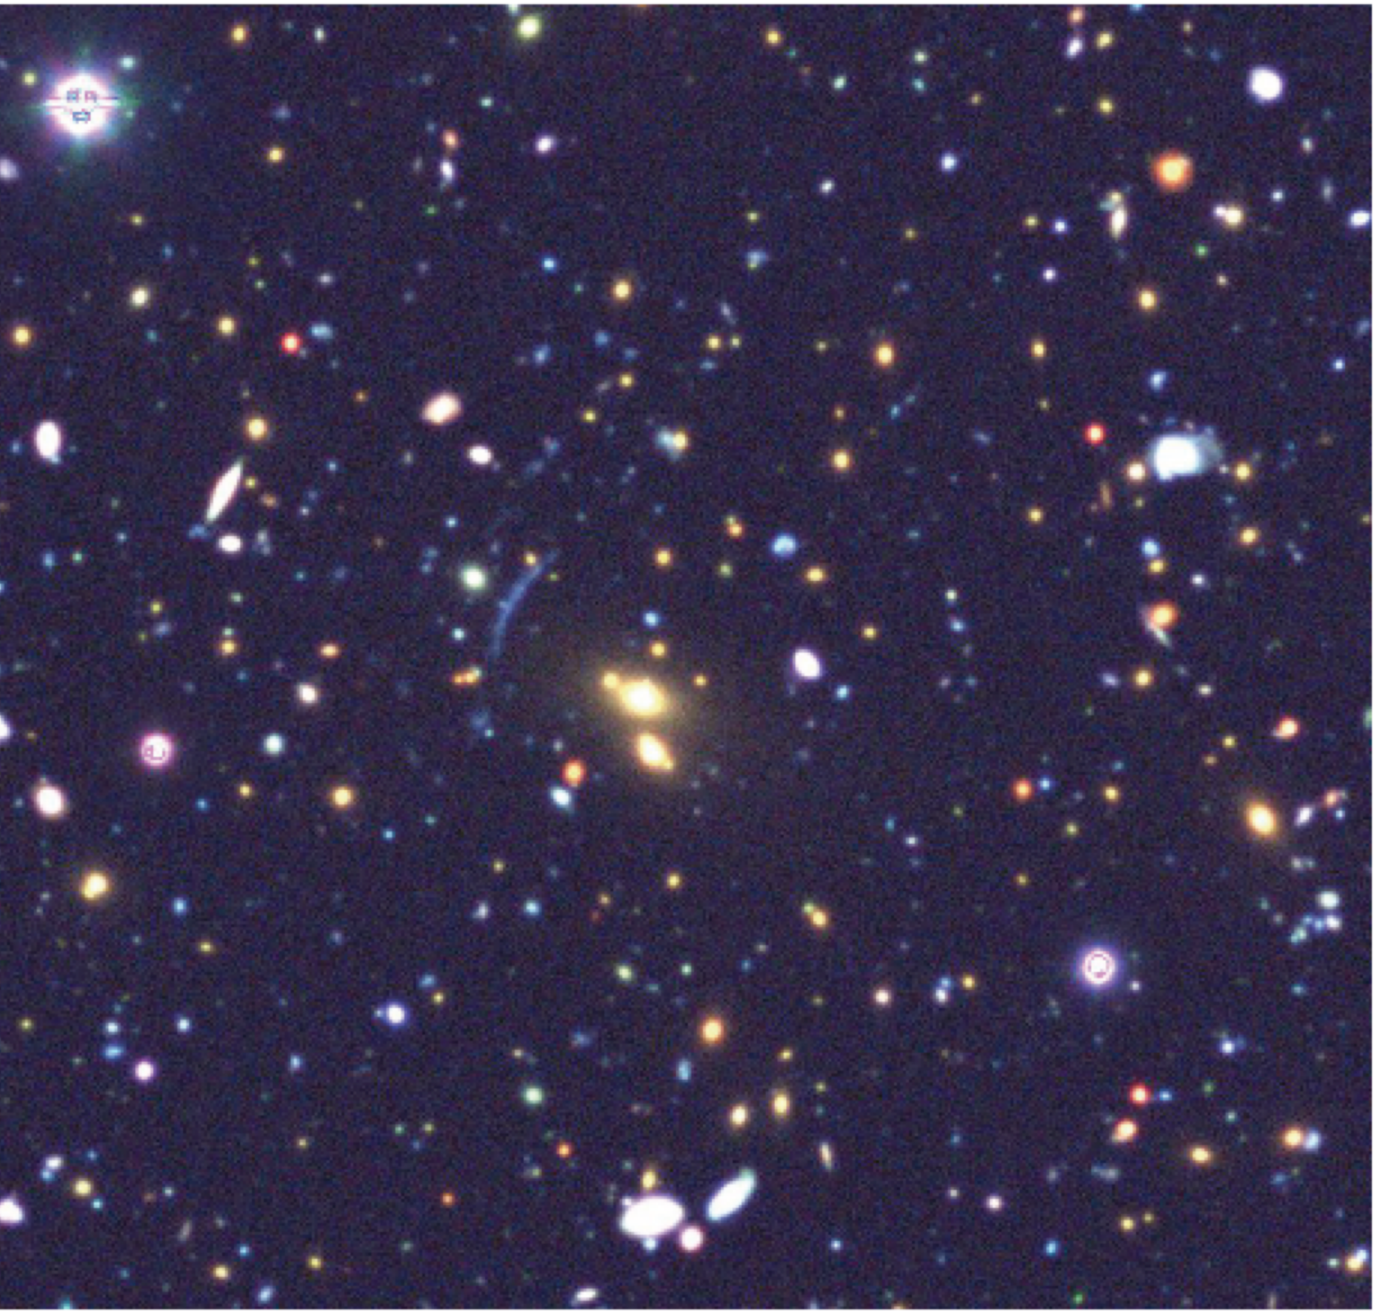

Real

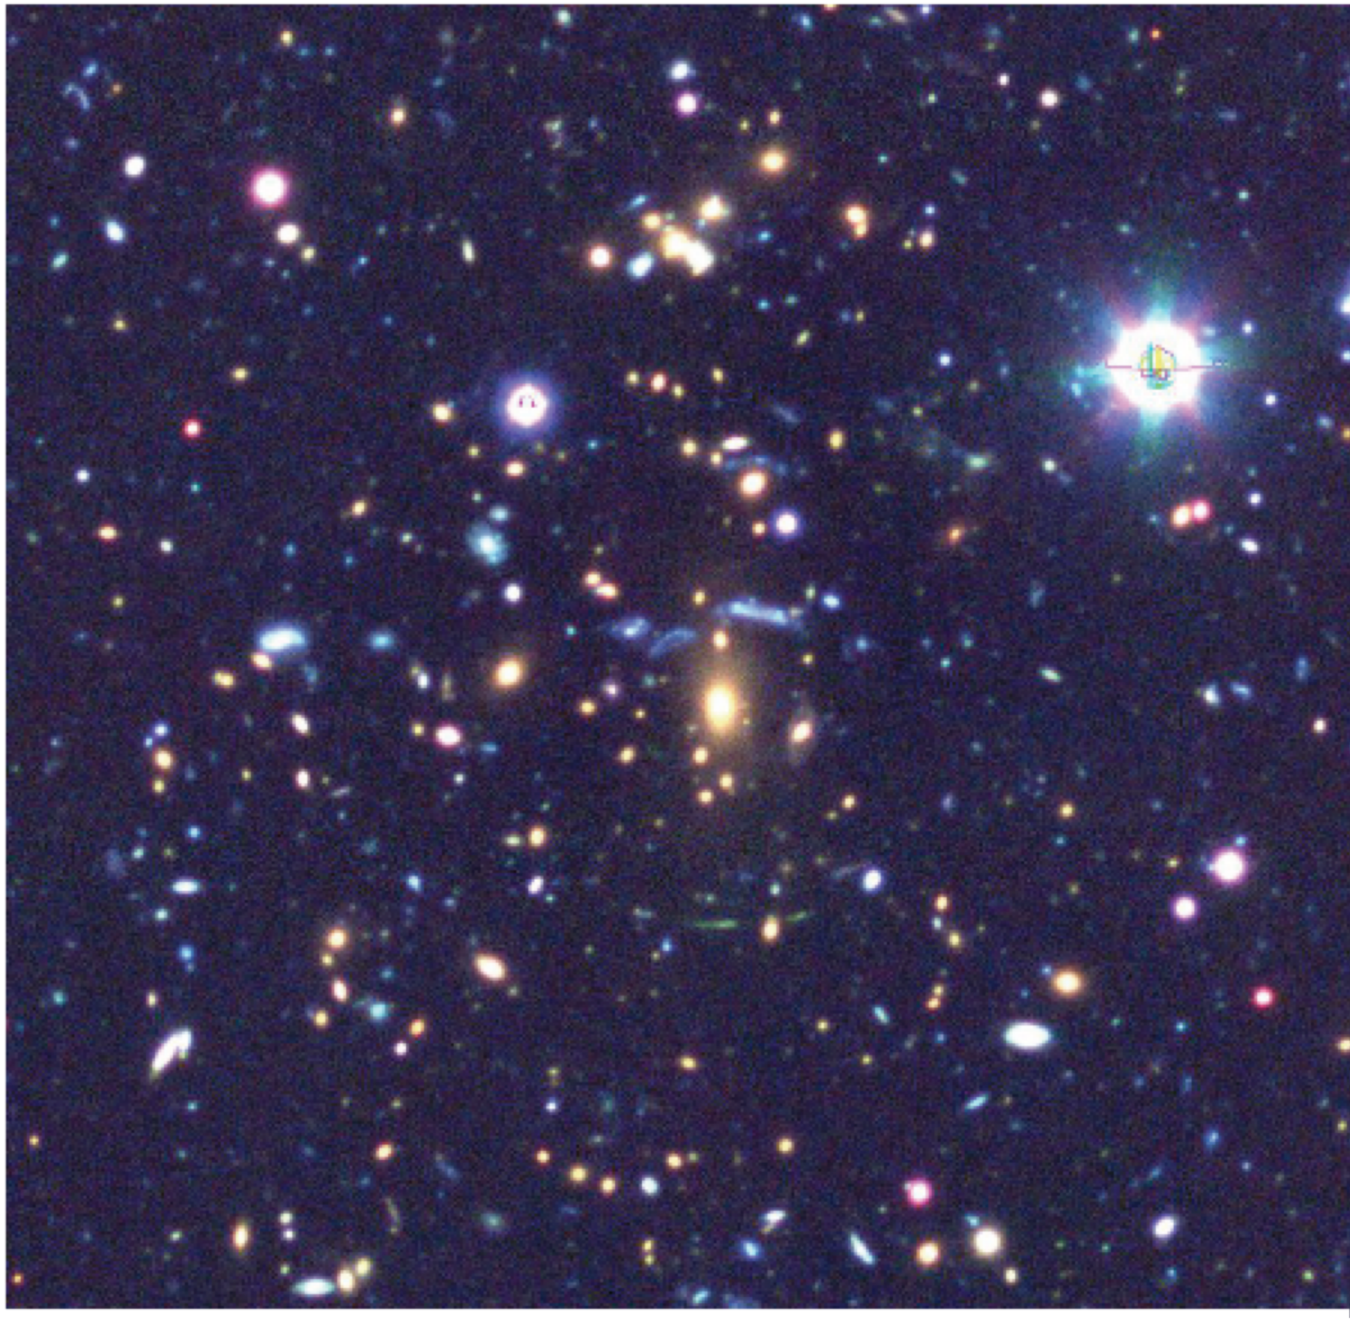

Simulated strong lens image to match SPT cluster observations taken with the MegaCAM camera on Magellan, in collaboration L. Bleem, M. Florian, S. Habib, M. Gladders, N. Li, S. Rangel

N. Li et al., arXiv:1511.03673


LSST DESC DC2 Synthetic catalog [Image credit: D. Boutigny (IN2P3)
and the LSST Dark Energy Survey Collaboration]

Science Drivers for Future HACC Development

- **Science Requirements:** Next-generation survey synthetic catalogs; detailed modeling for cosmological probes and cross-correlations; large-scale systematics studies
  - Weak lensing (WL) shear
  - Galaxy clustering (LSS)
  - Multiple LSS X WL cross-correlations
  - Cluster cosmology (mass calibration)
  - Secondary CMB anisotropies and backgrounds
  - Multiple CMB X LSS cross-correlations

Simulated sky maps for tSZ and kSZ observations (from Borg Cube)

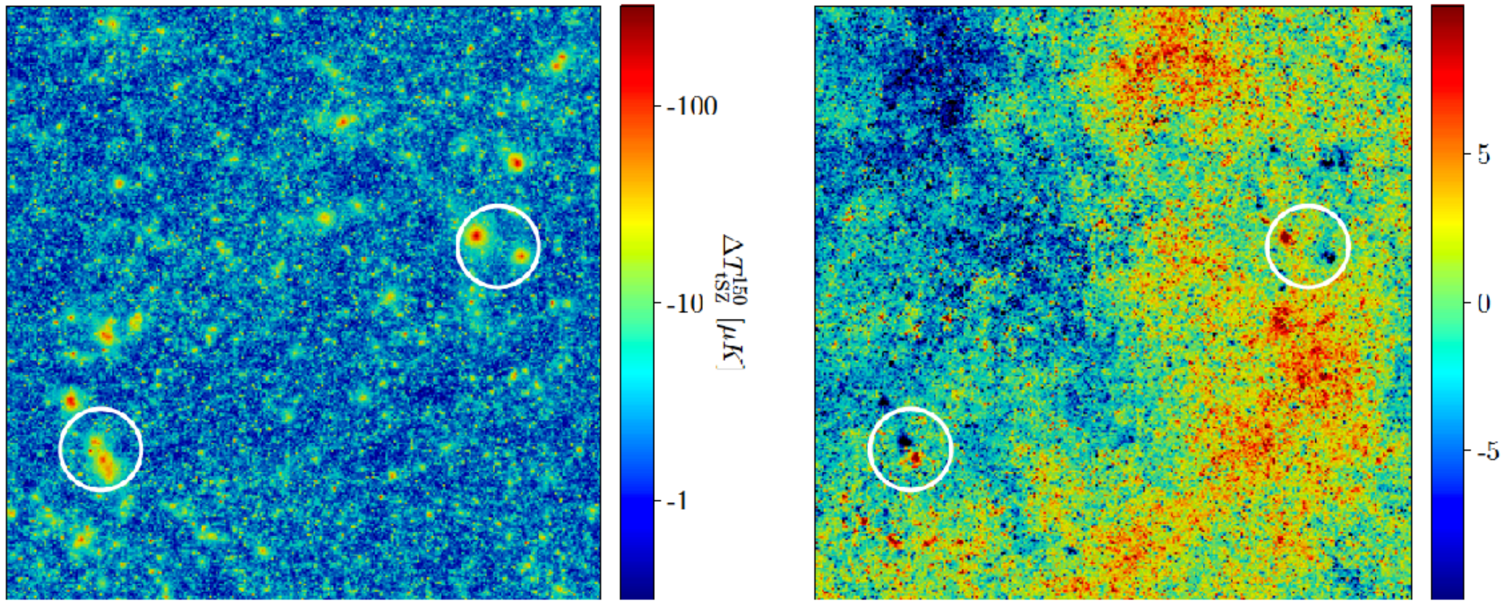

Many billions of objects, ~500 parameters each

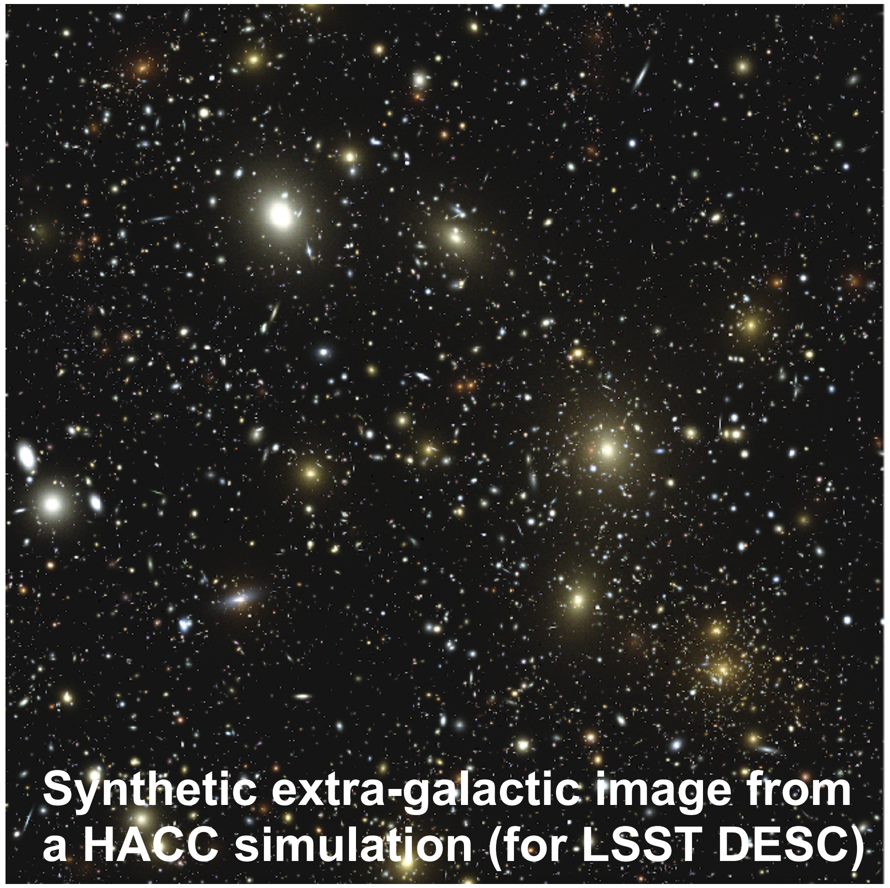
Synthetic extra-galactic image from a HACC simulation (for LSST DESC)

- **Hydro Developments:** New subgrid models for robust cosmological simulations
  - Standard subgrid models implemented
  - Gpc scale hydro runs at exascale

Cosmological Physics and Advanced Computing (+ Collaborators) at Argonne