

Title: Exascale Class simulations -- A leap forward in Cosmological Hydrodynamics

Speakers: Nick Frontiere

Collection/Series: Cosmic Ecosystems

Subject: Cosmology

Date: July 28, 2025 - 1:45 PM

URL: <https://pirsa.org/25070015>

Abstract:

Exascale computing is enabling a new generation of cosmological simulations, spanning both gravity-only and full hydrodynamics at unprecedented scale. The Frontier Exascale Simulation is the largest hydrodynamic run to date by over an order of magnitude, evolving more than 4 trillion particles in a 4.6 Gpc volume down to redshift zero. The simulation is well suited for predictive comparisons with multi-wavelength observations and for constructing full-sky mock surveys for upcoming observatories. In addition, a new suite of gravity-only simulations is now reaching beyond 10 trillion particles, producing survey-encompassing mocks ideal for large-scale structure analysis and tests of primordial non-Gaussianity. This talk will highlight these capabilities and their role in advancing predictive cosmology and survey science.



EXASCALE CLASS SIMULATIONS

A TRILLION-PARTICLE LEAP IN COSMOLOGICAL HYDRODYNAMICS

NICHOLAS FRONTIERE
CPS Division

Cosmic Ecosystems 2025
July 28th 2025



Argonne National Laboratory is a
U.S. Department of Energy laboratory
managed by UChicago Argonne, LLC.



Cosmological Simulations At Argonne

▪ Role of Cosmology Simulations

- Simulations produce **full sky synthetic catalogs**, characterize systematics, and provide constraints of LCDM by solving the inverse problem
- **Major impact** on multi-billion dollar DOE investment in ground and space-based sky surveys for planning, calibration, and validation
- Requires **large-scale**, controlled accuracy, multi-physics simulations at high dynamic range

▪ HACC (Hardware-Hybrid Accelerated Cosmology Code)

- **Hybrid** approach -- MPI + X (OpenMP, CUDA, HIP, SYCL, ...)
- Spectral Particle Mesh solver with a 3D FFT scalable to millions of ranks capable of processing grids **25000³+ at scale**
- Short range force solvers individually optimized to perform on both multi-core and **GPU accelerated** hardware
- Capable of **extreme scale simulations** on all modern supercomputing hardware



Simulation Characteristics

Domain

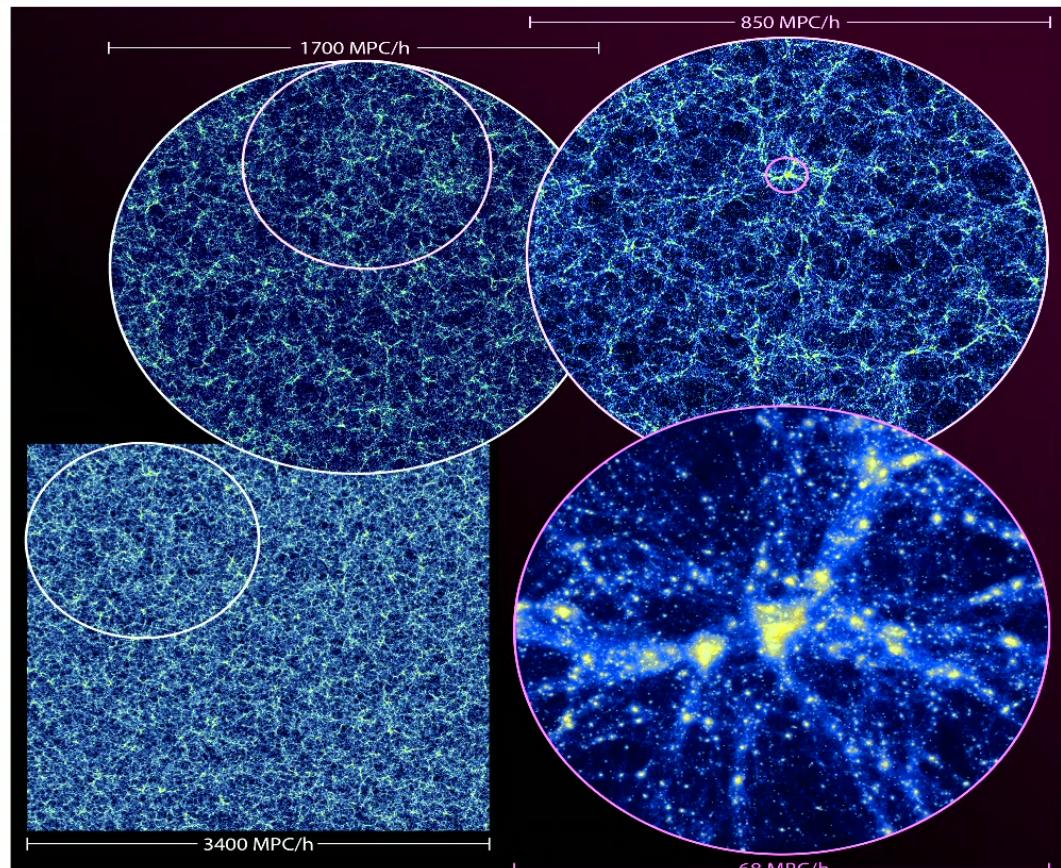
- 1-10 Gpc scale simulations
- Dynamic range of million to one
- Mass resolutions $10^7\text{-}10^9 M_{\odot} h^{-1}$

Gravity

- Vlasov-Poisson equation
- Dominates scales $> \sim 1 \text{ Mpc}$
- N-body

Gas Physics

- High order SPH solver, CRKSPH
- UVB Heating and Atomic Cooling
- Supernovae and Star formation
- Chemical Enrichment
- AGN Feedback
- Emulated and calibrated to observations



Last Journey Simulation



Cosmological Simulations At Argonne

▪ Role of Cosmology Simulations

- Simulations produce **full sky synthetic catalogs**, characterize systematics, and provide constraints of LCDM by solving the inverse problem
- **Major impact** on multi-billion dollar DOE investment in ground and space-based sky surveys for planning, calibration, and validation
- Requires **large-scale**, controlled accuracy, multi-physics simulations at high dynamic range

▪ HACC (Hardware-Hybrid Accelerated Cosmology Code)

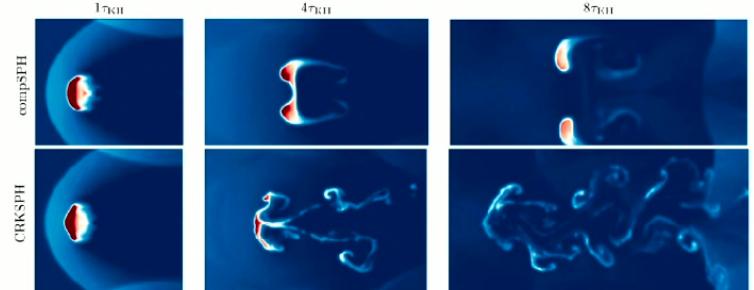
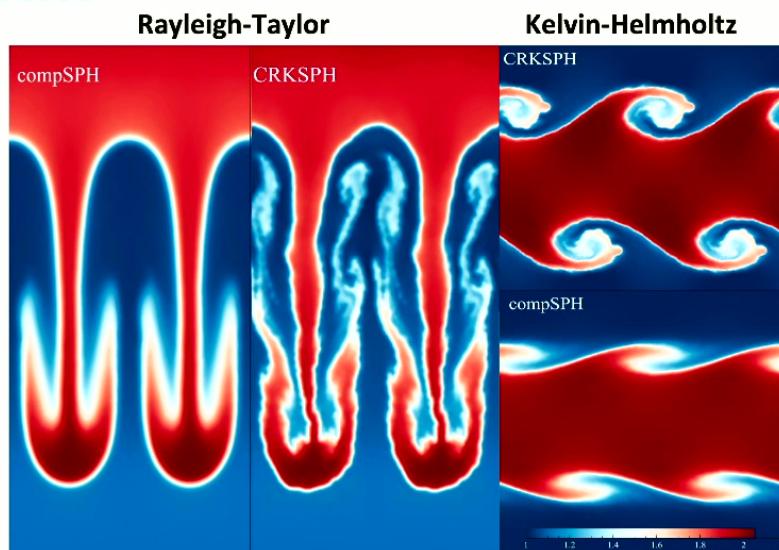
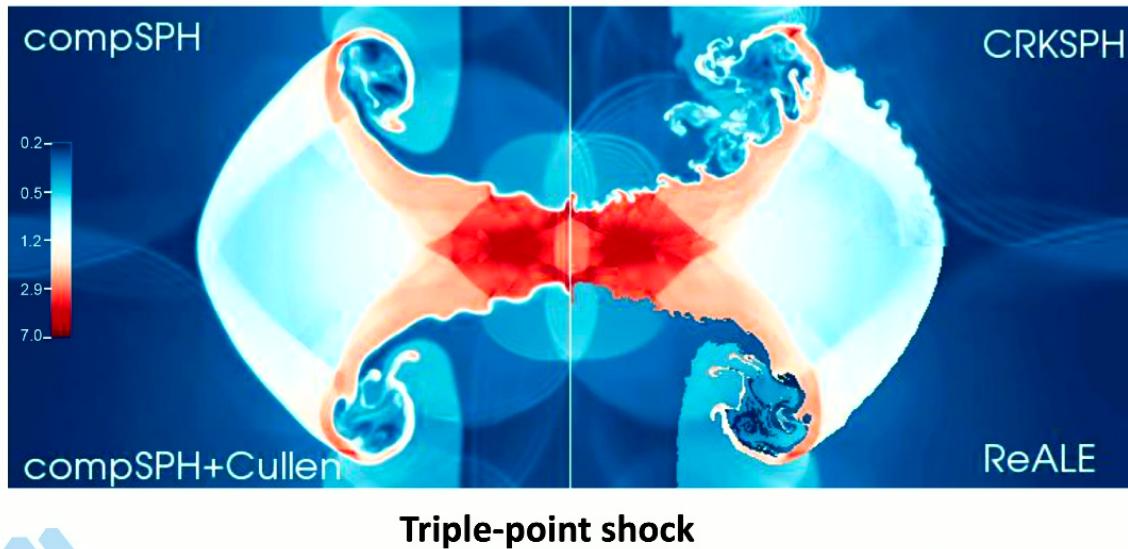
- **Hybrid** approach -- MPI + X (OpenMP, CUDA, HIP, SYCL, ...)
- Spectral Particle Mesh solver with a 3D FFT scalable to millions of ranks capable of processing grids **25000³+ at scale**
- Short range force solvers individually optimized to perform on both multi-core and **GPU accelerated** hardware
- Capable of **extreme scale simulations** on all modern supercomputing hardware



CRKSPH Algorithm

▪ Conservative Reproducing Kernel SPH

- Higher-order reproducing kernels exactly reproduce constant and linear order fields
- Conservative reformulation of the dynamical equations maintains exact conservation of energy and momentum
- New artificial viscosity that capitalizes on the increased accuracy of the velocity gradients, reduces excessive SPH artificial diffusion
- No artificial conductivity required



'Blob' test with CRK-SPH showing blob disruption due to Kelvin-Helmholtz and Rayleigh-Taylor instabilities

Age of Exascale

▪ Next generation of supercomputers

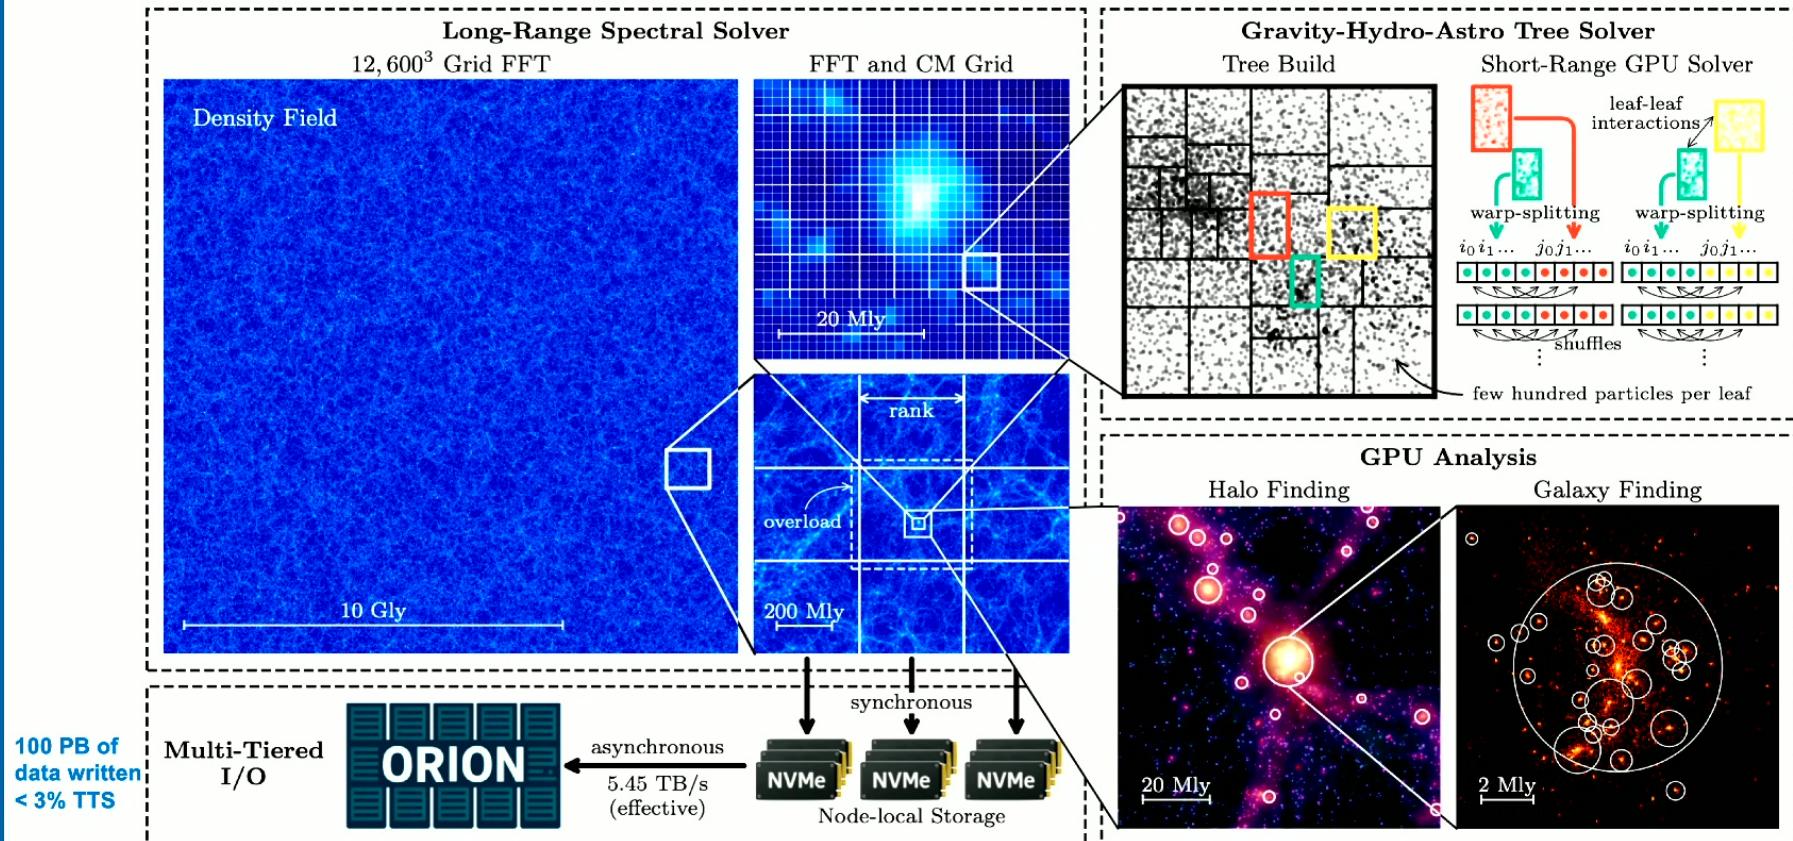
- Frontier ORNL 1.353 Eflop/s
- Aurora ANL 1.012 Eflop/s
- El Capitan LLNL 1.742 Eflop/s
- GPU accelerated
- Mixed precision

▪ Exascale class capabilities

- 10-50x improvement
- Hydro simulations are typically 10-30x slower than GO
- Previous state-of-the-art GO simulations can now include gas physics. Trillion(s) particles
- Next-generation GO simulations in the 10+ trillion particle class.



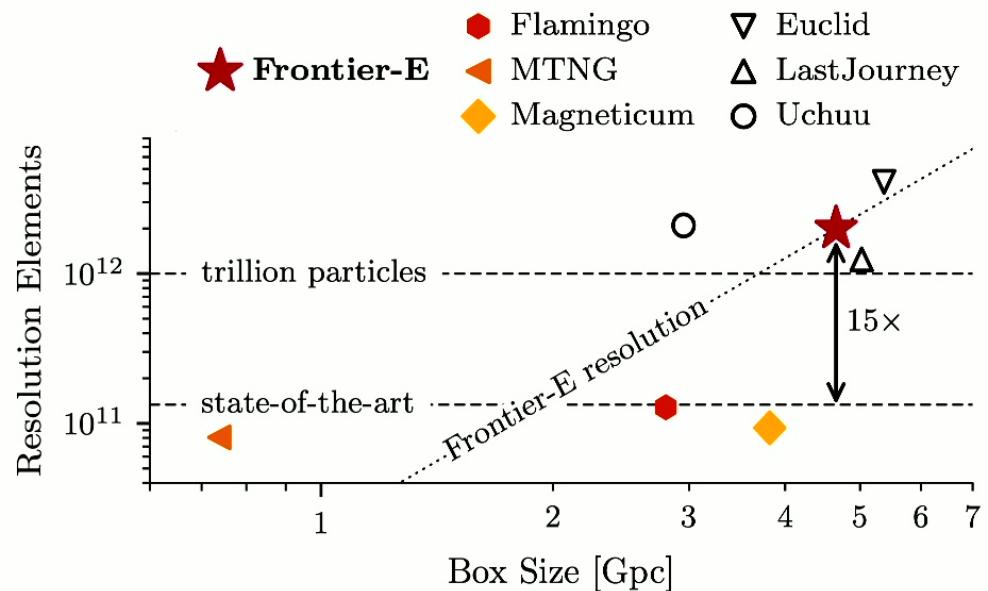
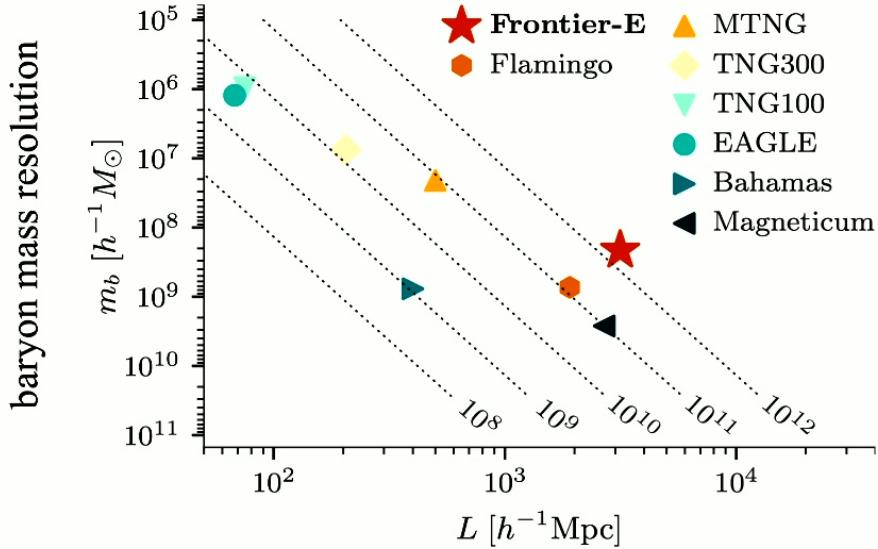
CRK-HACC Runtime Overview



All compute intensive solvers and analysis are offloaded to the **GPU**

At scale on Frontier achieves **513 PFLOPs** of performance evaluating **46.6 billion** particles/s

Exascale Class Simulations

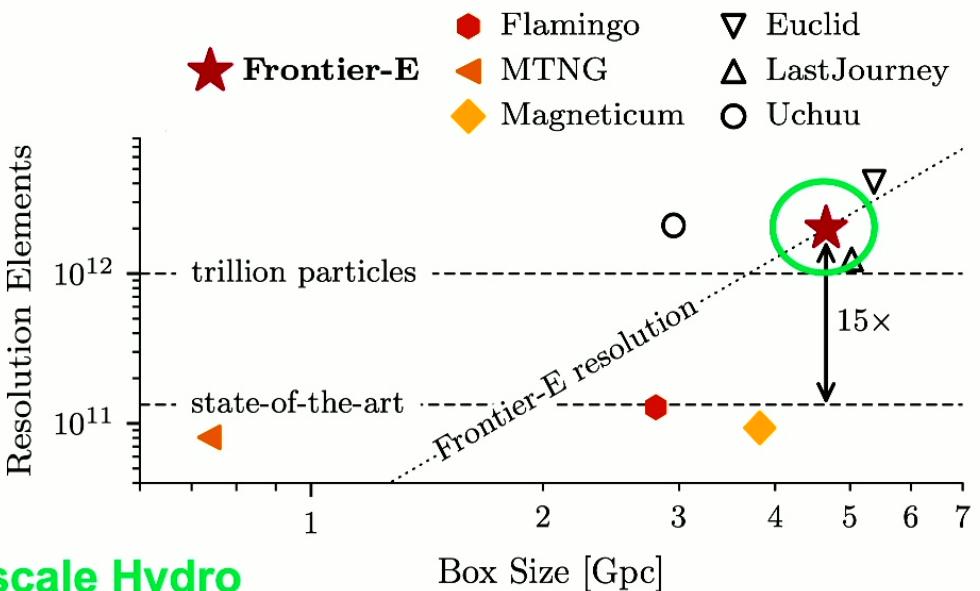
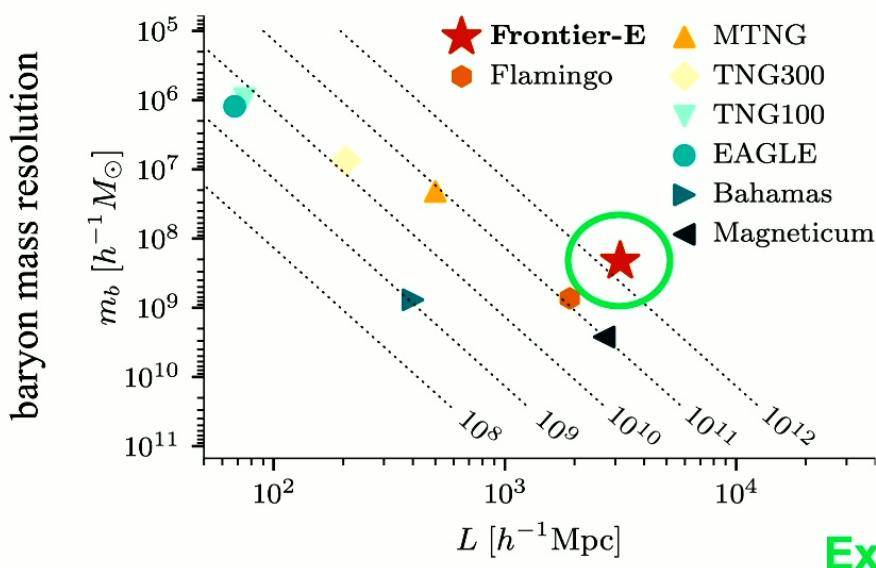


Exascale unlocks the ability to run previous **state of the art** gravity only simulations now including **hydrodynamics** in a **week** of machine time (~ 1 yr CPU)

Gravity only simulations break into the **10+ trillion** particle class



Exascale Class Simulations



Exascale unlocks the ability to run previous **state of the art** gravity only simulations now including **hydrodynamics** in a **week** of machine time (~ 1 yr CPU)

Gravity only simulations break into the **10+ trillion** particle class



Frontier Exascale Simulation

▪ Survey Scale

- $3.150 \text{ Gpc } h^{-1}$ domain
- 2×12600^3 particles – 4 trillion
- $2 \times 10^8 M_\odot h^{-1}$ baryon mass resolution

▪ Cluster Science

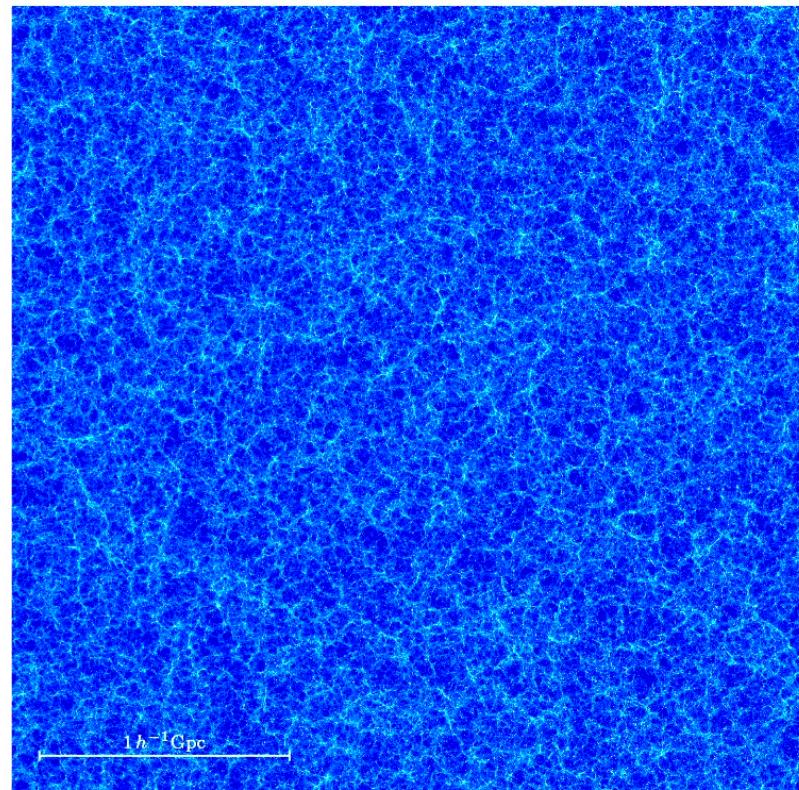
- 2000+ (13) clusters $> 10^{15} M_\odot h^{-1}$
- 570,000+ (75k) $> 10^{14} M_\odot h^{-1}$

▪ Full sky particle LC ($z < 3$)

- Consistent maps and catalogs
- Downscale LC out to $z < 5$

▪ Extensive Data Outputs (12 PBs)

- Full particles for objects $> 10^{13} M_\odot h^{-1}$
- Merger trees, catalogs and snapshots
- ~ 1 billion galaxies at $z = 0$



Evolved to $z = 0$, Frontier-E marks a new generation of survey-scale simulations with over an order-of-magnitude leap in scale

Frontier Exascale Simulation

Survey Scale

- 3.150 Gpc h^{-1} domain
- 2×12600^3 particles – 4 trillion
- $2 \times 10^8 M_\odot h^{-1}$ baryon mass resolution

Cluster Science

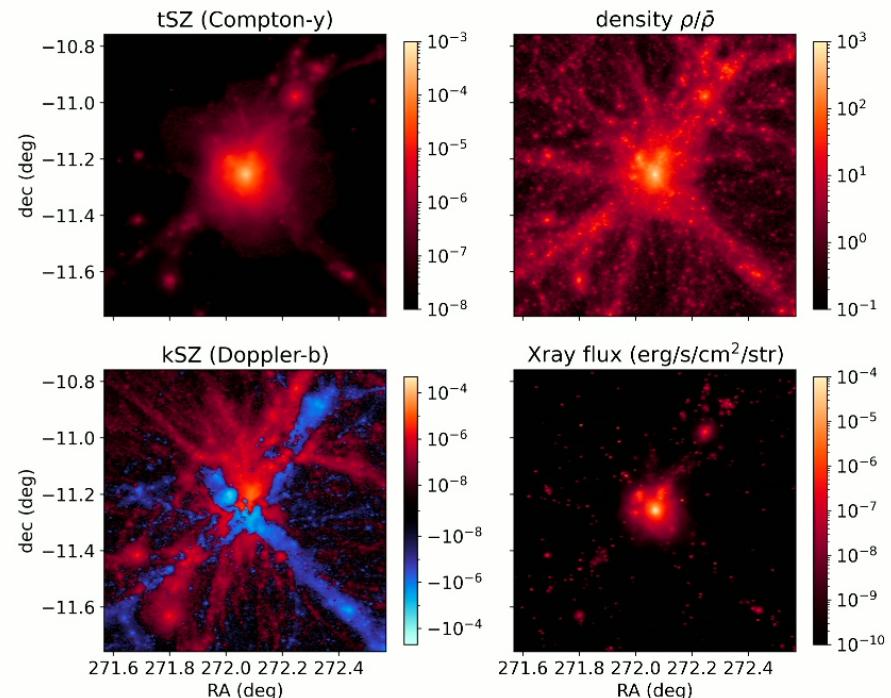
- 2000+ (13) clusters $> 10^{15} M_\odot h^{-1}$
- 570,000+ (75k) $> 10^{14} M_\odot h^{-1}$

Full sky particle LC ($z < 3$)

- Consistent maps and catalogs
- Downscale LC out to $z < 5$

Extensive Data Outputs (12 PBs)

- Full particles for objects $> 10^{13} M_\odot h^{-1}$
- Merger trees, catalogs and snapshots
- ~ 1 billion galaxies at $z = 0$



Mock observations of a cluster on the light cone. Thermal and kinematic Sunyaev-Zeldovich maps, the total matter density contrast, and the flux emission in a bolometric X-ray band



Frontier Exascale Simulation

Survey Scale

- 3.150 Gpc h^{-1} domain
- 2×12600^3 particles – 4 trillion
- $2 \times 10^8 M_\odot h^{-1}$ baryon mass resolution

Cluster Science

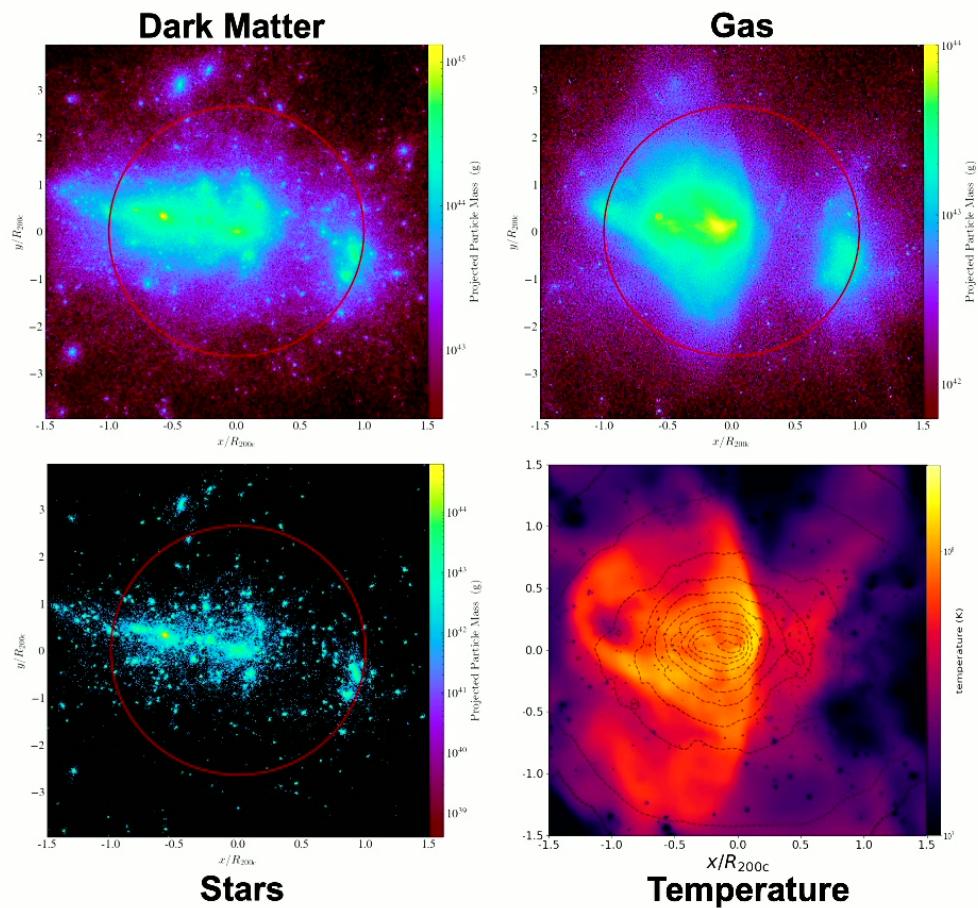
- 2000+ (13) clusters $> 10^{15} M_\odot h^{-1}$
- 570,000+ (75k) $> 10^{14} M_\odot h^{-1}$

Full sky particle LC ($z < 3$)

- Consistent maps and catalogs
- Downscale LC out to $z < 5$

Extensive Data Outputs (12 PBs)

- Full particles for objects $> 10^{13} M_\odot h^{-1}$
- Merger trees, catalogs and snapshots
- ~ 1 billion galaxies at $z = 0$



The large volume enables the identification and study of rare, dynamically complex systems



Frontier Exascale Simulation

Survey Scale

- 3.150 Gpc h^{-1} domain
- 2×12600^3 particles – 4 trillion
- $2 \times 10^8 M_\odot h^{-1}$ baryon mass resolution

Cluster Science

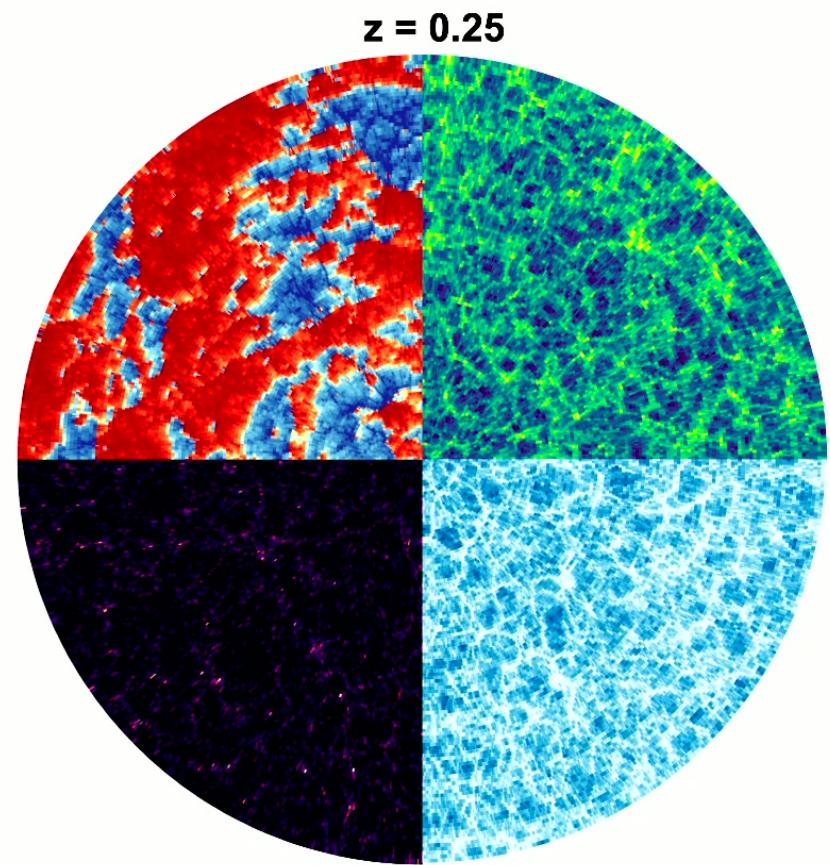
- 2000+ (13) clusters $> 10^{15} M_\odot h^{-1}$
- 570,000+ (75k) $> 10^{14} M_\odot h^{-1}$

Full sky particle LC ($z < 3$)

- Consistent maps and catalogs
- Downscale LC out to $z < 5$

Extensive Data Outputs (12 PBs)

- Full particles for objects $> 10^{13} M_\odot h^{-1}$
- Merger trees, catalogs and snapshots
- ~ 1 billion galaxies at $z = 0$



Slice through full-sky maps showing, clockwise from top left: kSZ, tSZ, X-ray, and density.



Frontier Exascale Simulation

- **Survey Scale**

- 3.150 Gpc h^{-1} domain
- 2×12600^3 particles – 4 trillion
- $2 \times 10^8 M_\odot h^{-1}$ baryon mass resolution

- **Cluster Science**

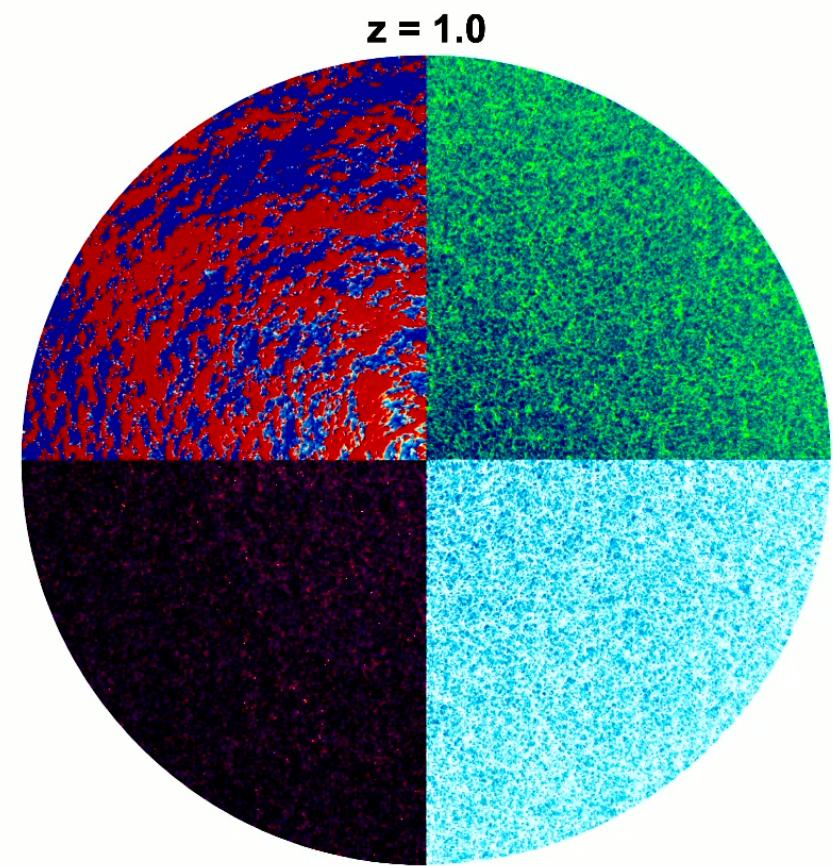
- 2000+ (13) clusters $> 10^{15} M_\odot h^{-1}$
- 570,000+ (75k) $> 10^{14} M_\odot h^{-1}$

- **Full sky particle LC ($z < 3$)**

- Consistent maps and catalogs
- Downscale LC out to $z < 5$

- **Extensive Data Outputs (12 PBs)**

- Full particles for objects $> 10^{13} M_\odot h^{-1}$
- Merger trees, catalogs and snapshots
- ~ 1 billion galaxies at $z = 0$



Slice through full-sky maps showing, clockwise from top left: kSZ, tSZ, X-ray, and density.



Frontier Exascale Simulation

Survey Scale

- 3.150 Gpc h^{-1} domain
- 2×12600^3 particles – 4 trillion
- $2 \times 10^8 M_\odot h^{-1}$ baryon mass resolution

Cluster Science

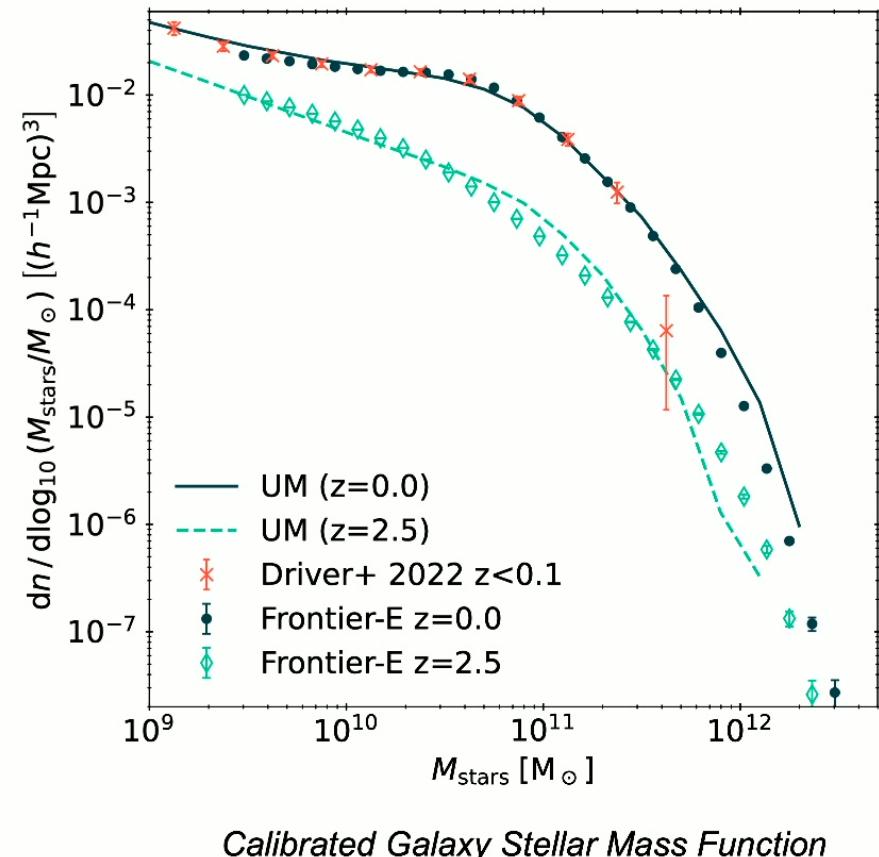
- 2000+ (13) clusters $> 10^{15} M_\odot h^{-1}$
- 570,000+ (75k) $> 10^{14} M_\odot h^{-1}$

Full sky particle LC ($z < 3$)

- Consistent maps and catalogs
- Downscale LC out to $z < 5$

Extensive Data Outputs (12 PBs)

- Full particles for objects $> 10^{13} M_\odot h^{-1}$
- Merger trees, catalogs and snapshots
- ~ 1 billion galaxies at $z = 0$



Frontier Exascale Simulation

Survey Scale

- 3.150 Gpc h^{-1} domain
- 2×12600^3 particles – 4 trillion
- $2 \times 10^8 M_\odot h^{-1}$ baryon mass resolution

Cluster Science

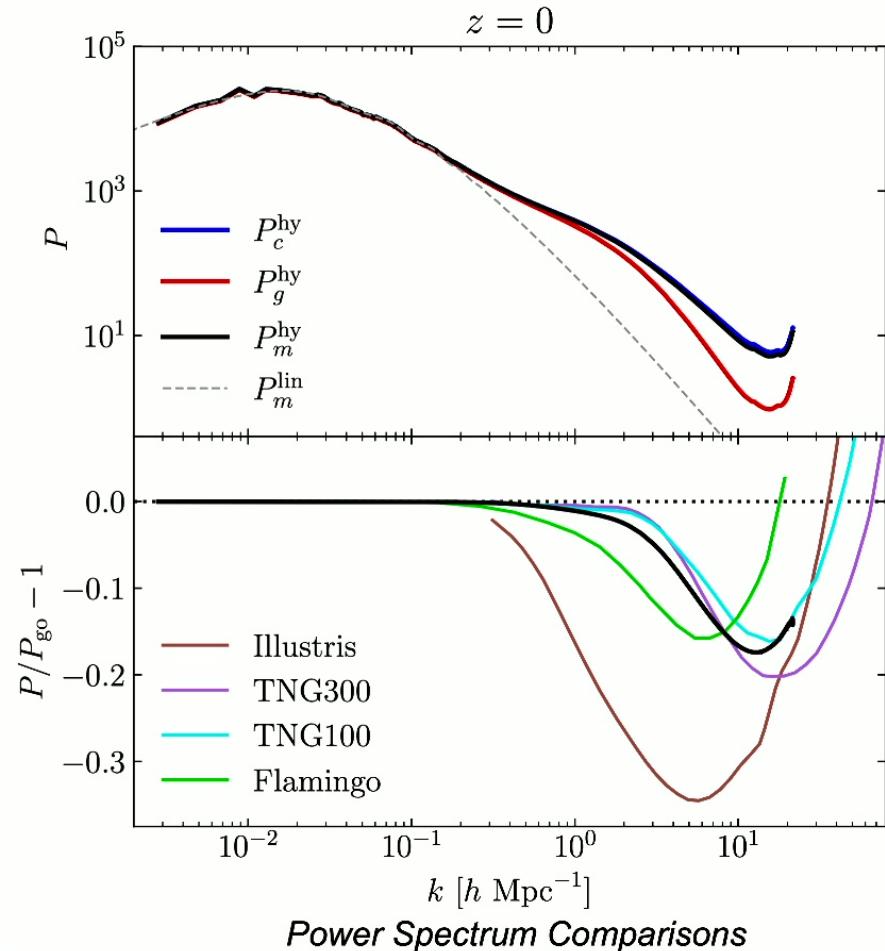
- 2000+ (13) clusters $> 10^{15} M_\odot h^{-1}$
- 570,000+ (75k) $> 10^{14} M_\odot h^{-1}$

Full sky particle LC ($z < 3$)

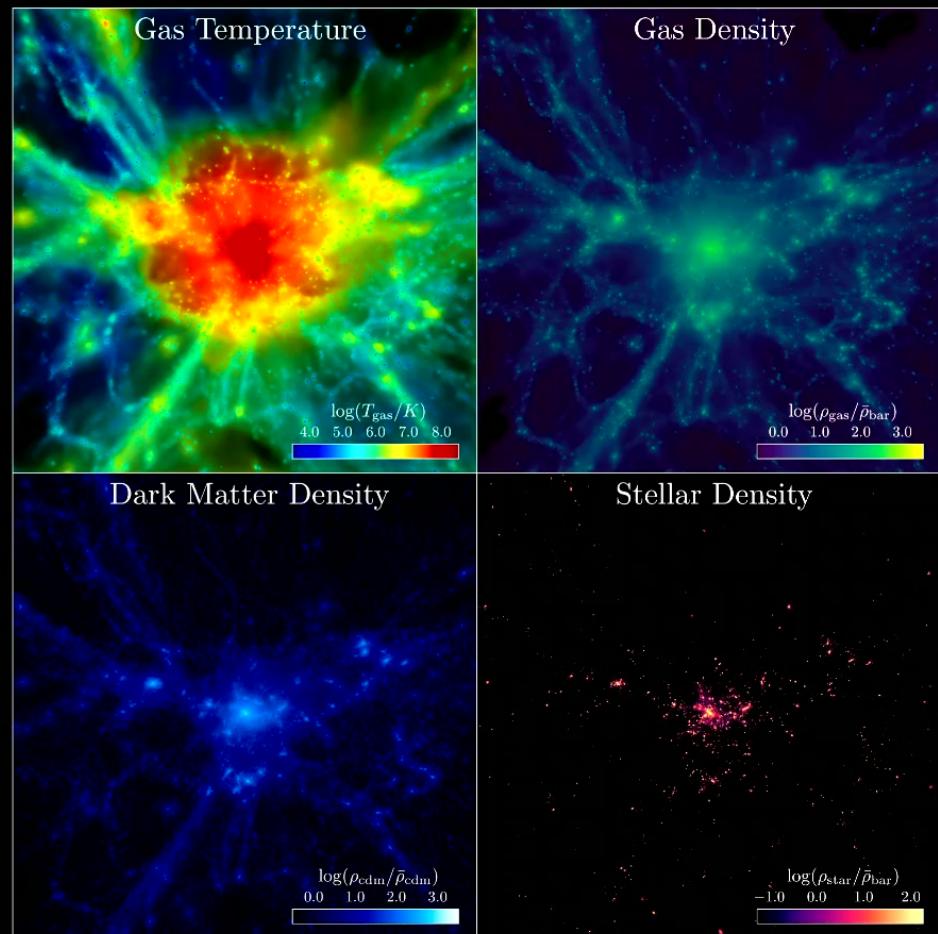
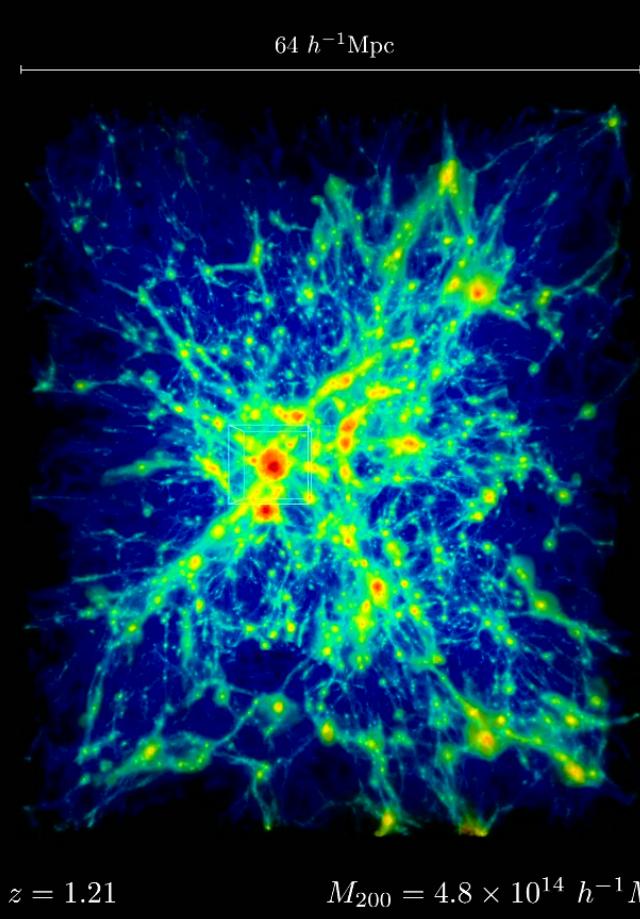
- Consistent maps and catalogs
- Downscale LC out to $z < 5$

Extensive Data Outputs (12 PBs)

- Full particles for objects $> 10^{13} M_\odot h^{-1}$
- Merger trees, catalogs and snapshots
- ~ 1 billion galaxies at $z = 0$



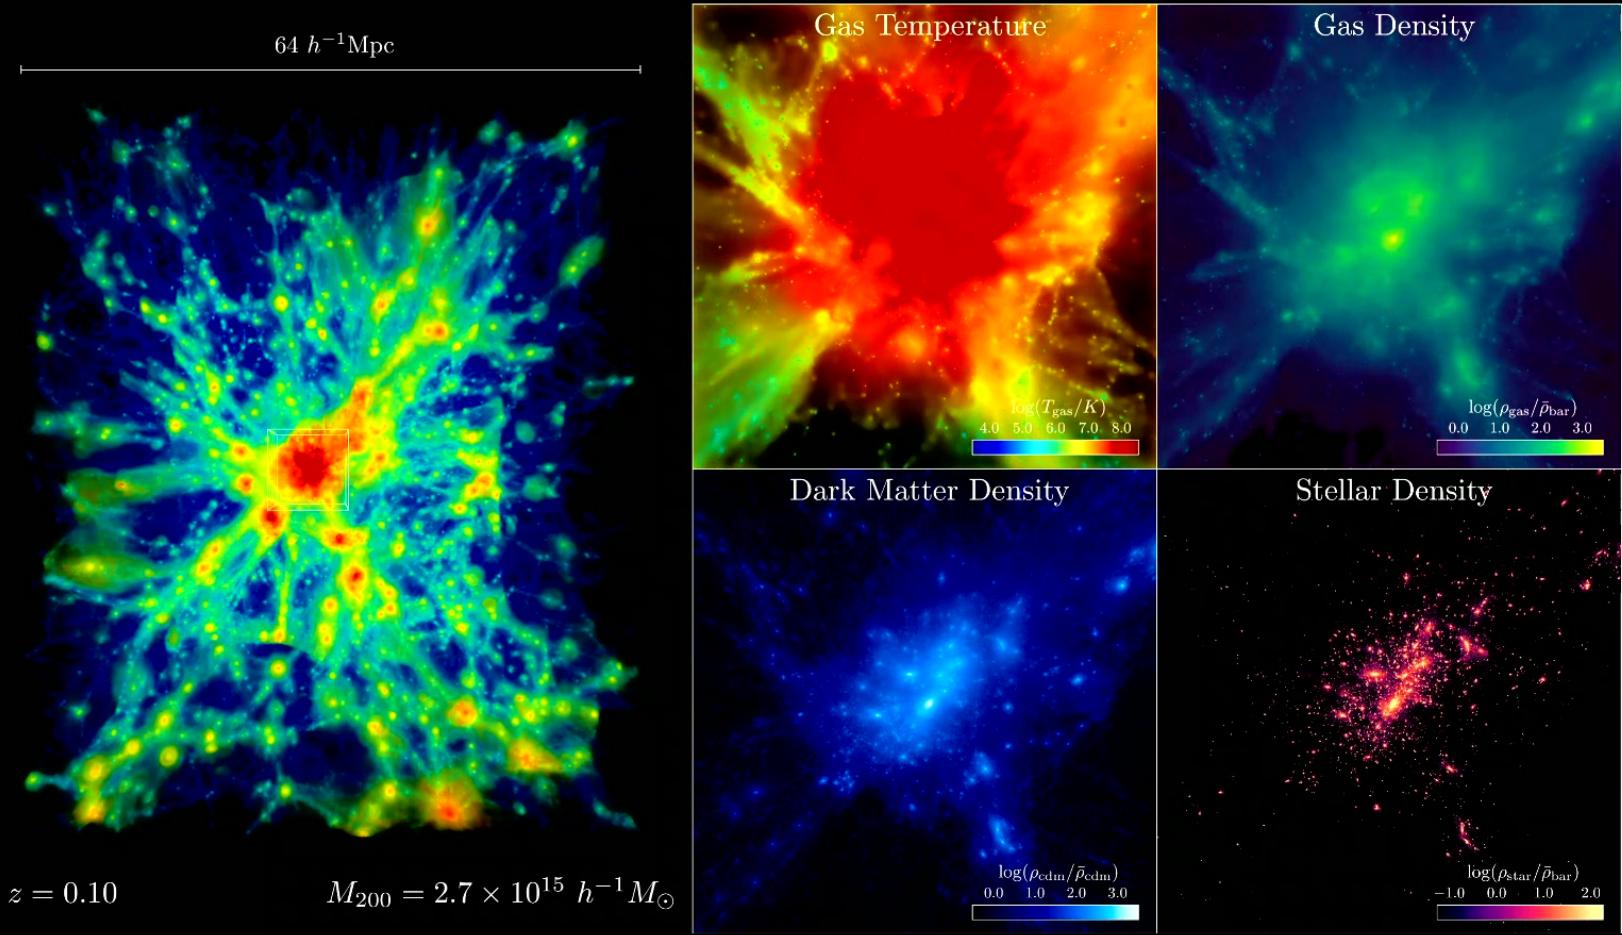
Ember Cluster



1/100,000th subvolume largest cluster zoom-in



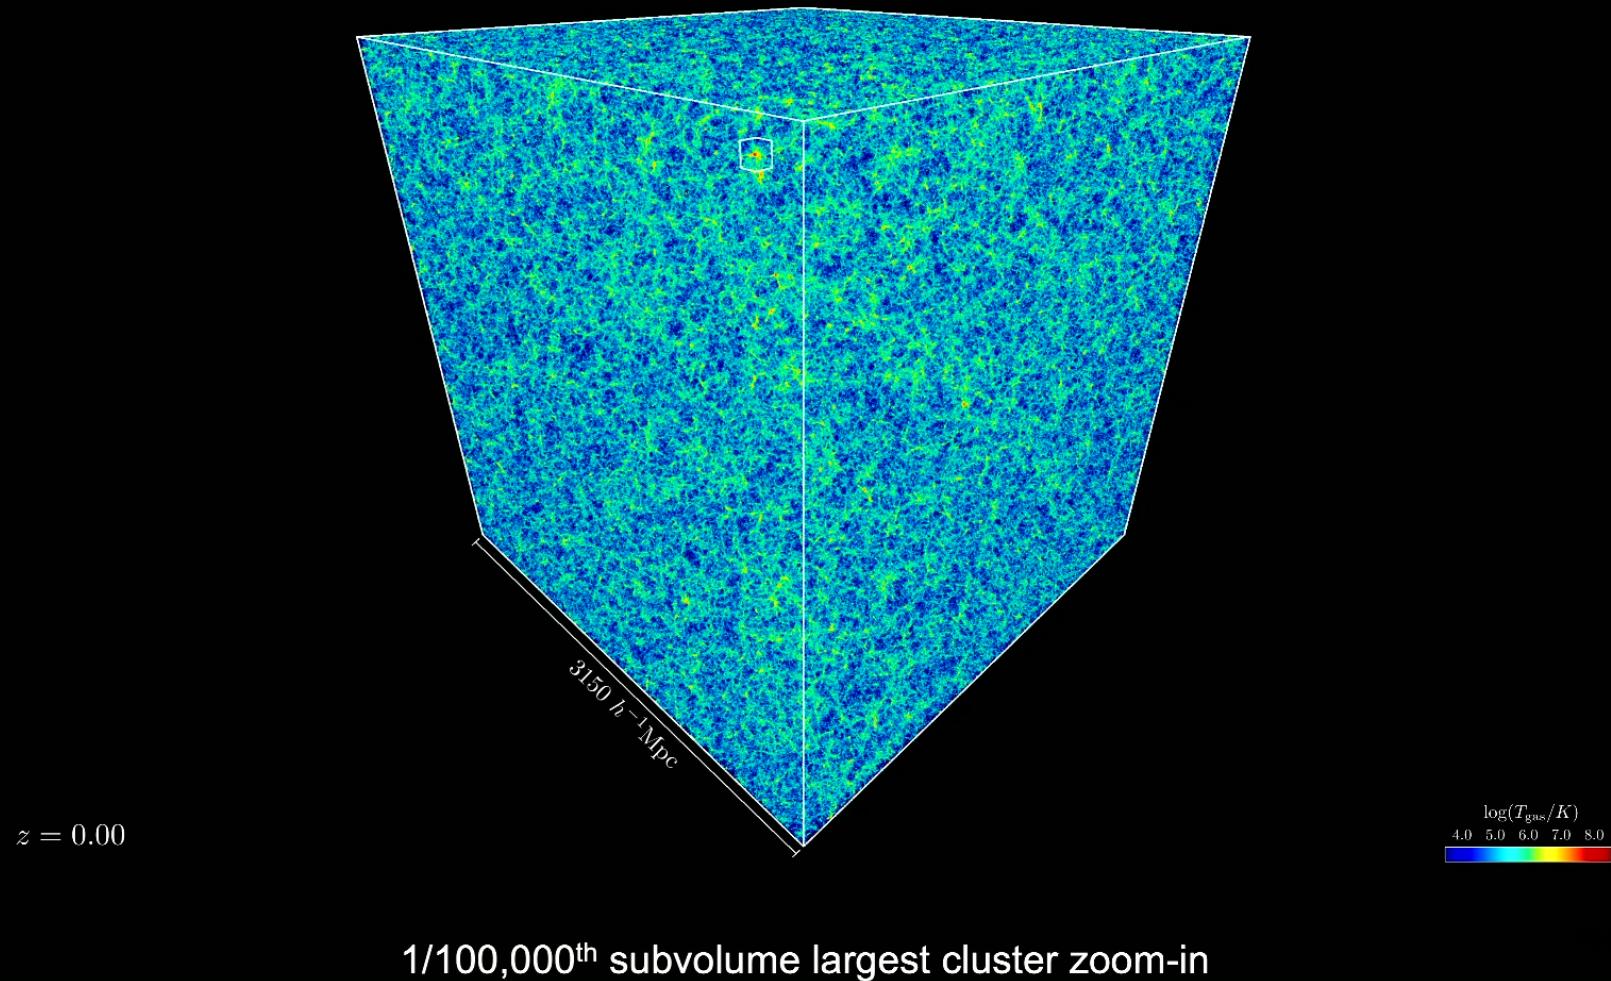
Ember Cluster



1/100,000th subvolume largest cluster zoom-in

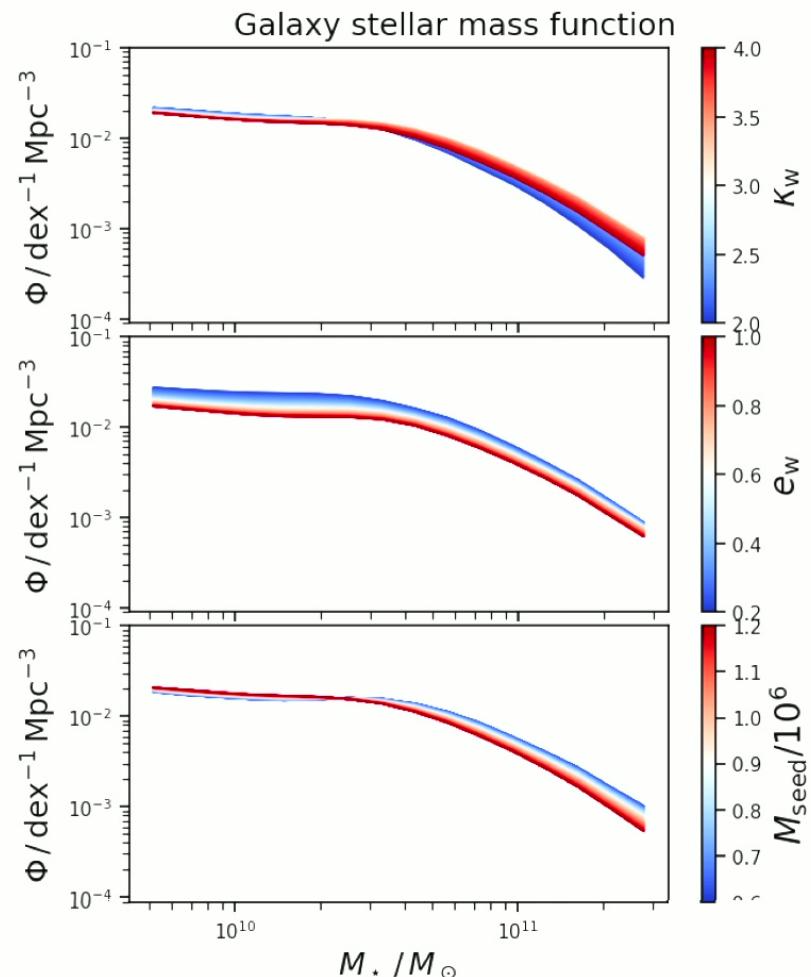


Ember Cluster



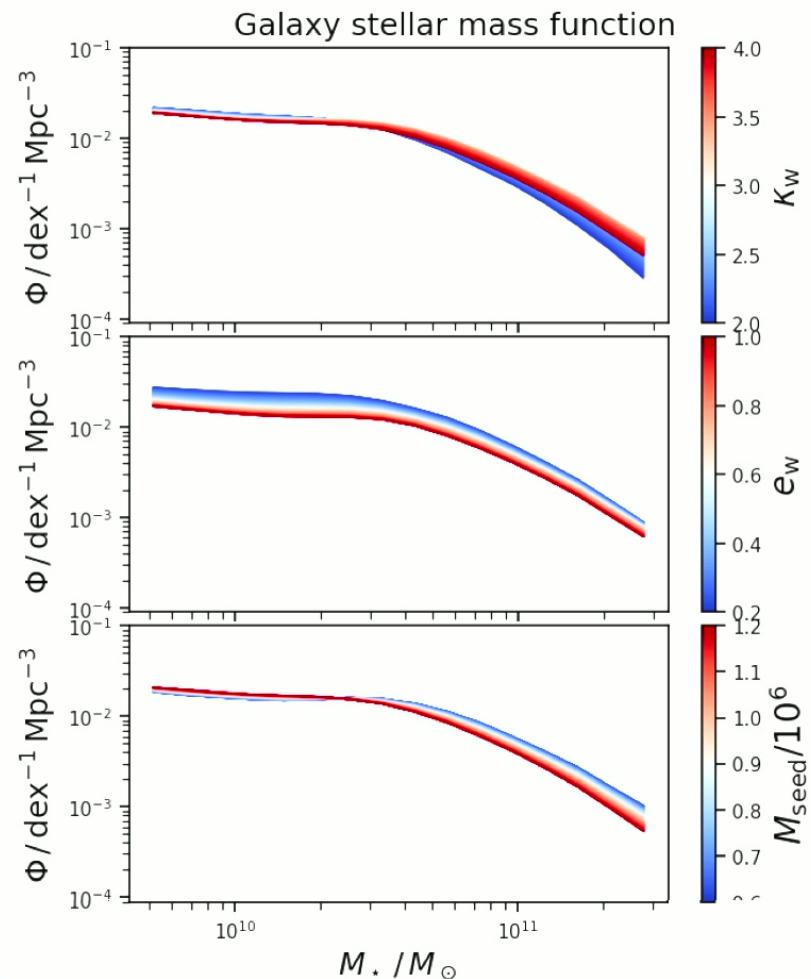
Frontier Exascale Calibration

- **Simulation Suites**
 - 128 Mpc h^{-1} initial calibration set (64)
 - 256 Mpc h^{-1} cluster calibration (16)
- **Targets**
 - Multi-redshift Galaxy Stellar Mass Function
 - Cluster Gas Density Profiles
- **Emulation**
 - Varying across 5 active parameters
 - Two wind, three for AGN
- **Exascale Throughput**
 - Exploration at larger volumes
 - Parameter expansion – hydro + cosmology



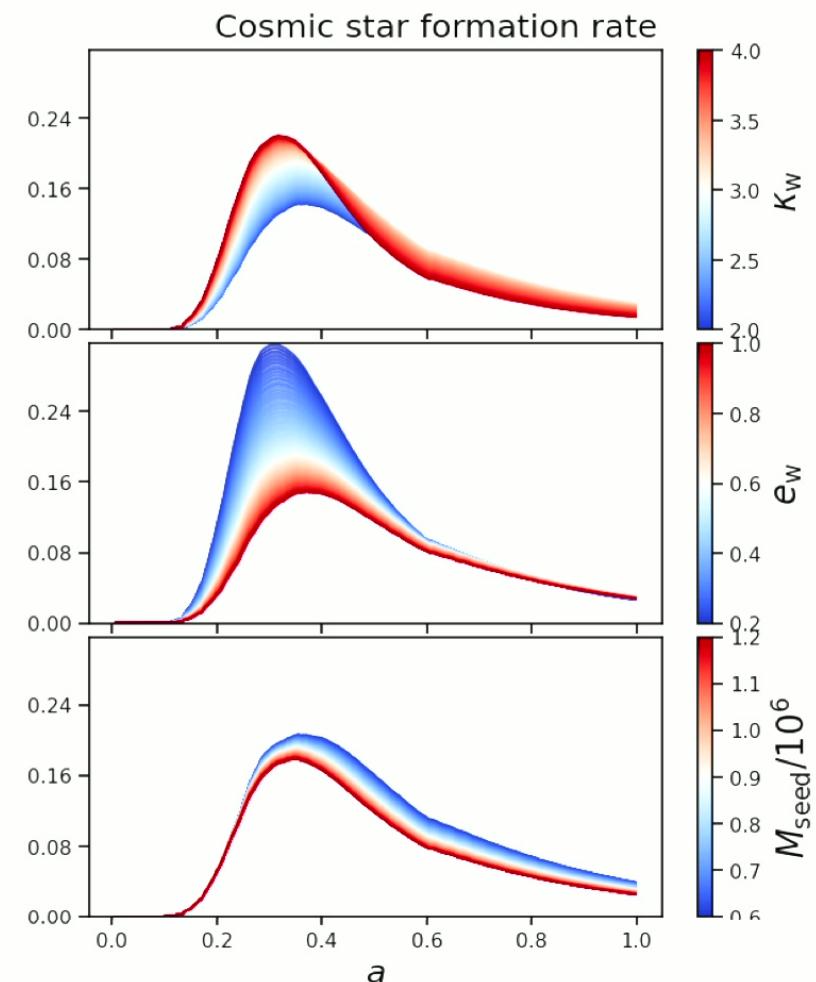
Frontier Exascale Calibration

- **Simulation Suites**
 - 128 Mpc h^{-1} initial calibration set (64)
 - 256 Mpc h^{-1} cluster calibration (16)
- **Targets**
 - Multi-redshift Galaxy Stellar Mass Function
 - Cluster Gas Density Profiles
- **Emulation**
 - Varying across 5 active parameters
 - Two wind, three for AGN
- **Exascale Throughput**
 - Exploration at larger volumes
 - Parameter expansion – hydro + cosmology



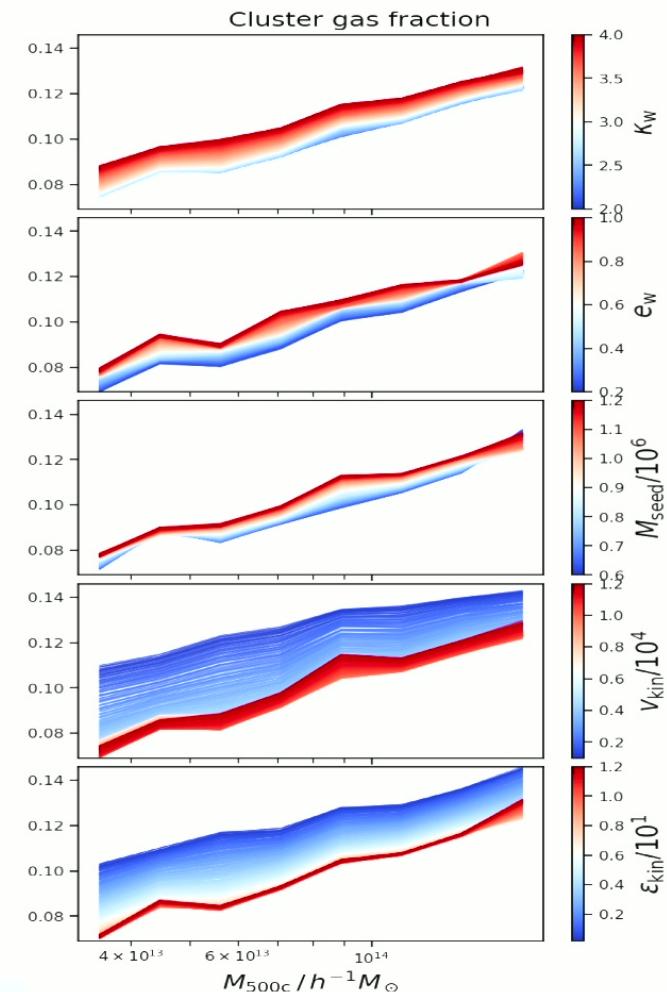
Frontier Exascale Calibration

- **Simulation Suites**
 - 128 Mpc h^{-1} initial calibration set (64)
 - 256 Mpc h^{-1} cluster calibration (16)
- **Targets**
 - Multi-redshift Galaxy Stellar Mass Function
 - Cluster Gas Density Profiles
- **Emulation**
 - Varying across 5 active parameters
 - Two wind, three for AGN
- **Exascale Throughput**
 - Exploration at larger volumes
 - Parameter expansion – hydro + cosmology



Frontier Exascale Calibration

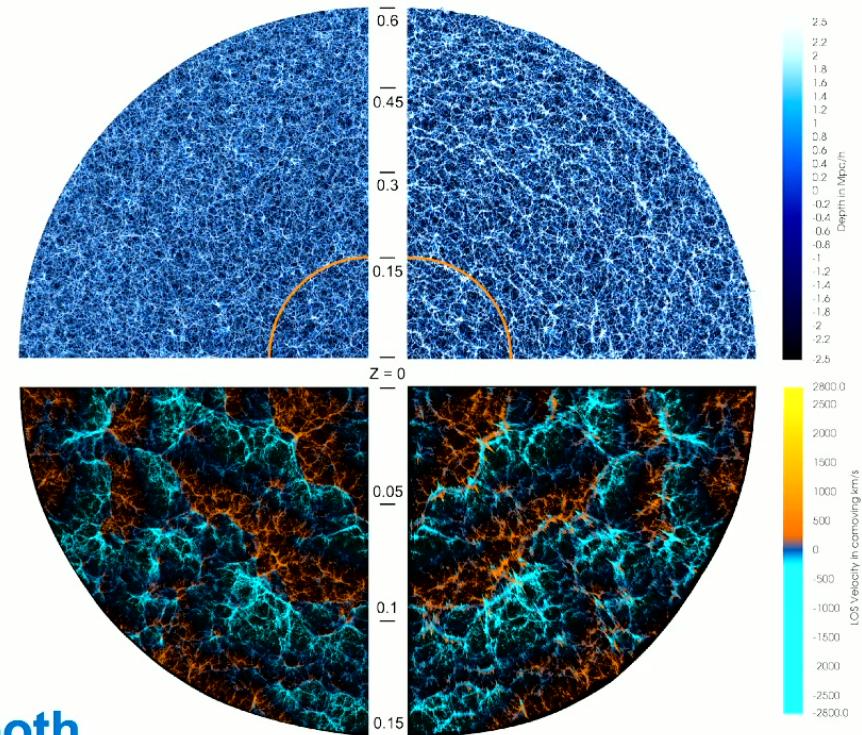
- **Simulation Suites**
 - 128 Mpc h^{-1} initial calibration set (64)
 - 256 Mpc h^{-1} cluster calibration (16)
- **Targets**
 - Multi-redshift Galaxy Stellar Mass Function
 - Cluster Gas Density Profiles
- **Emulation**
 - Varying across 5 active parameters
 - Two wind, three for AGN
- **Exascale Throughput**
 - Exploration at larger volumes
 - Parameter expansion – hydro + cosmology



Exascale Class Gravity Simulations

Upcoming Aurora Exascale Simulation Pair

- **Primordial non-Gaussianity GO suite**
 - 5.5 Gpc/h domain
 - $23,760^3$ particles – **13.4+ trillion**
 - $10^9 M_\odot h^{-1}$ mass resolution
- **Investigate PNG in $1 > k > 10^{-3} \text{ Mpc}^{-1}$**
 - Blinded analyses
 - SphereX full sky and DESI mocks
 - Assembly histories used for simulated sky observables
- **Synthetic sky catalogs generated using both galaxy-halo connection and ML methods**



Particle light-cone visualizations from the 'Farpoint' run (Frontiere et al., ApJS 2022), a high mass resolution, large-volume cosmological simulation run with HACC



Machine Learning Analyses

▪ SED Painting N-body Simulations

- Mass accretion histories, are connected with star formation histories and dust attenuation curves
- Produces billions of galaxies covering full sky, across survey sensitive redshift ranges

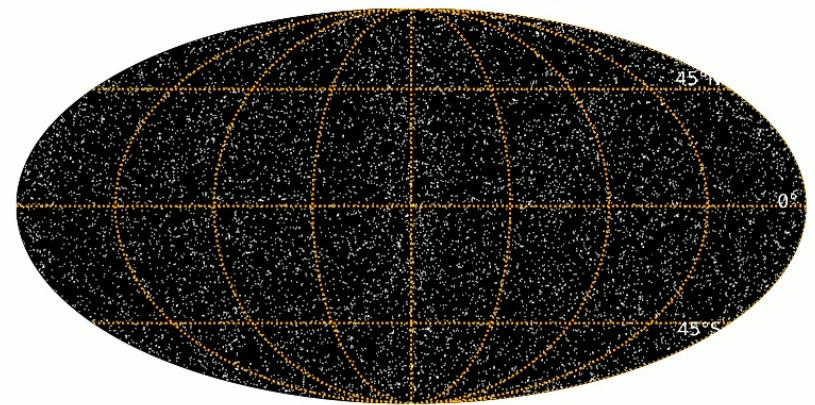
▪ Painting galaxies with colors

- Input stellar age, metallicity and SSP spectral emissions
- Output catalogs include luminosity profiles, galaxy SEDs along with photometry for Rubin, SPHEREx, COSMOS and other surveys

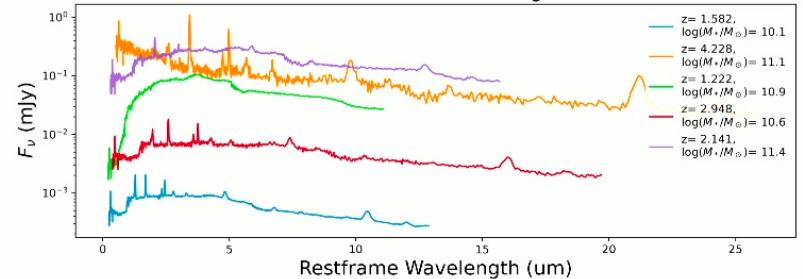
▪ Multi-modal, multi-task model training

- Scalar catalogs, time-series, SEDs, merger trees, galaxy images, density fields: multiple modalities representing multiple representations of the same object
- Foundation model connects simulation values with and observables

Sky Distribution of Galaxies in full sky from HACC-LastJourney simulation



Narrow band SEDs of a few galaxies



Machine Learning Analyses

▪ SED Painting N-body Simulations

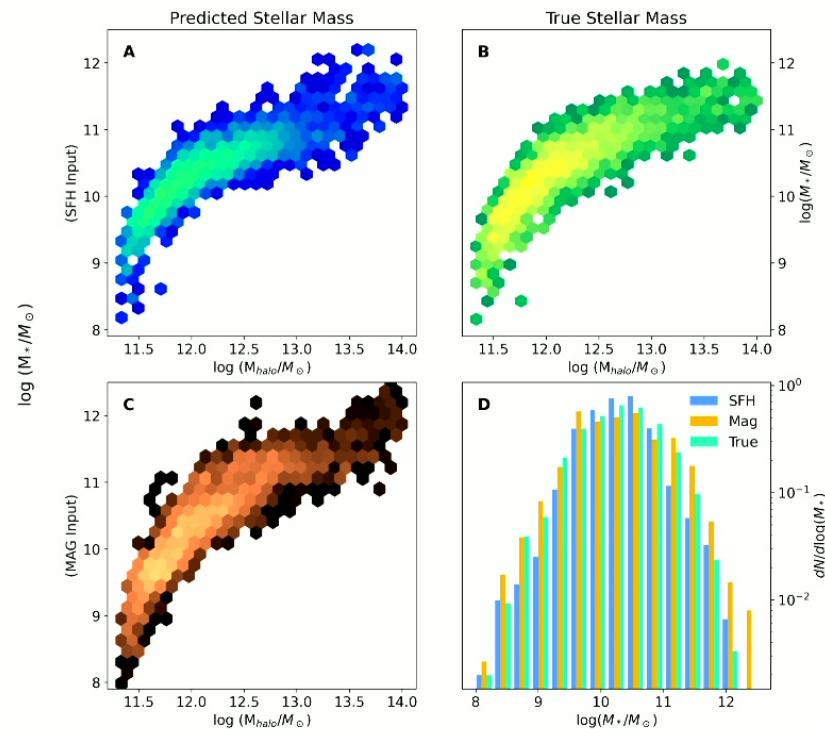
- Mass accretion histories, are connected with star formation histories and dust attenuation curves
- Produces billions of galaxies covering full sky, across survey sensitive redshift ranges

▪ Painting galaxies with colors

- Input stellar age, metallicity and SSP spectral emissions
- Output catalogs include luminosity profiles, galaxy SEDs along with photometry for Rubin, SPHEREx, COSMOS and other surveys

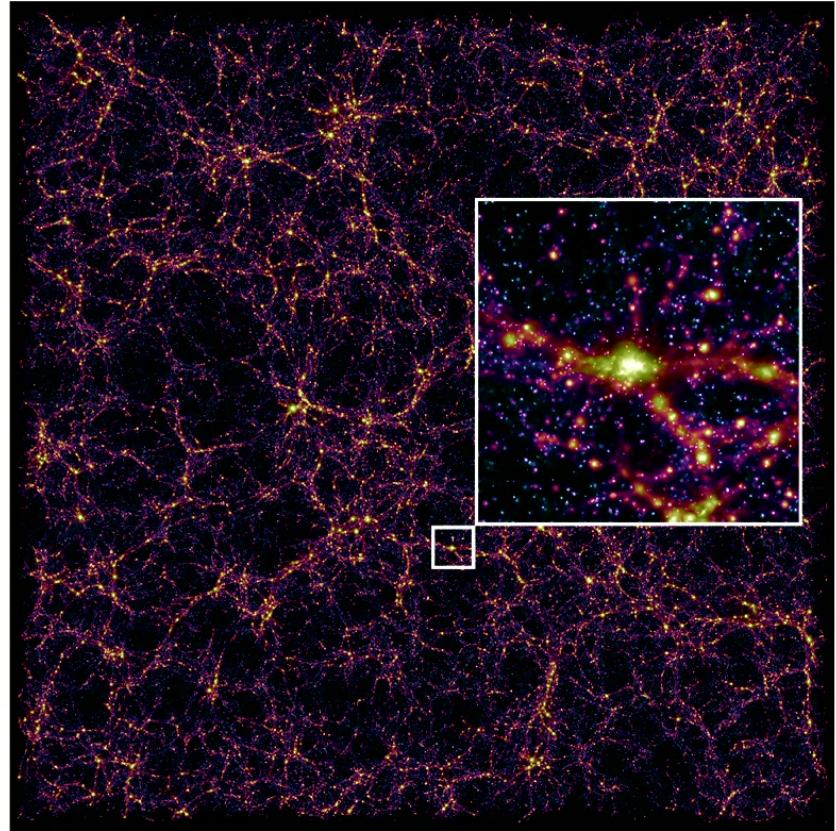
▪ Multi-modal, multi-task model training

- Scalar catalogs, time-series, SEDs, merger trees, galaxy images, density fields: multiple modalities representing multiple representations of the same object
- Foundation model connects simulation values with and observables



The Exascale Future is Bright

- Exascale machines provide a **powerful tool** to drive numerical cosmology science
- Enabling **>10x leap** in simulation scale: from trillion-particle hydro to 10+ trillion GO in ~week(s) of machine time.
- Generated datasets are massive, ideal for **ML training**
- Capable of running targeted **emulation suites** in new optimization domains
- Will significantly advance **predictive capabilities** and **simulation fidelity** necessary to validate and interpret results from ongoing and upcoming surveys



slice through downscaled Frontier-E simulation
raytracing visualization



Frontier Exascale Simulation

Survey Scale

- 3.150 Gpc h^{-1} domain
- 2×12600^3 particles – 4 trillion
- $2 \times 10^8 M_\odot h^{-1}$ baryon mass resolution

Cluster Science

- 2000+ (13) clusters $> 10^{15} M_\odot h^{-1}$
- 570,000+ (75k) $> 10^{14} M_\odot h^{-1}$

Full sky particle LC ($z < 3$)

- Consistent maps and catalogs
- Downscale LC out to $z < 5$

Extensive Data Outputs (12 PBs)

- Full particles for objects $> 10^{13} M_\odot h^{-1}$
- Merger trees, catalogs and snapshots
- ~ 1 billion galaxies at $z = 0$

