Cosmological Simulations for Large-Scale Sky Surveys

Katrin Heitmann

OLCF User Meeting, July 23, 2014

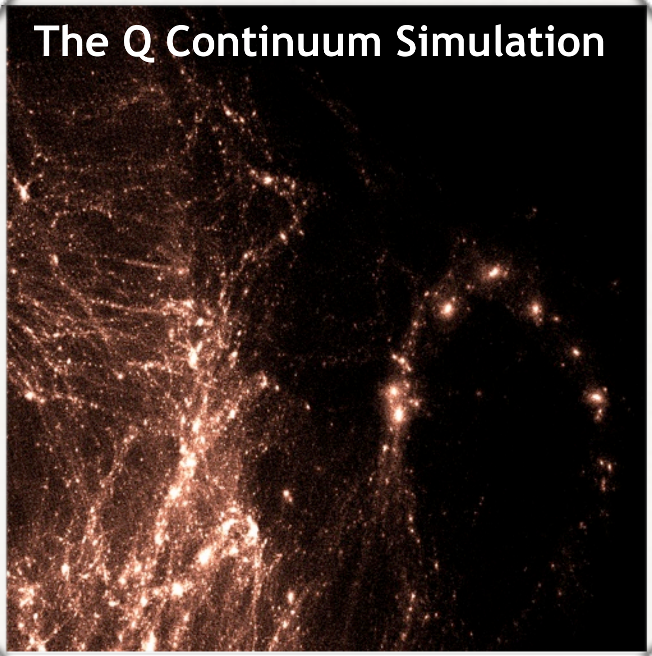

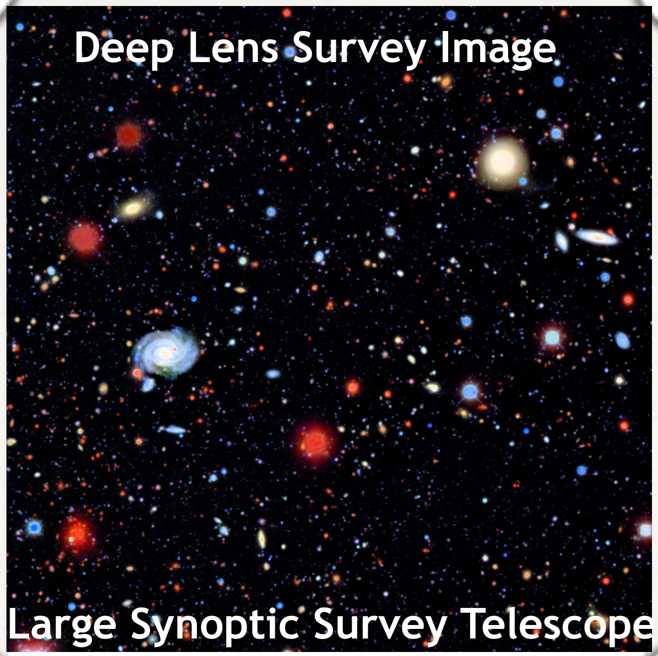

Thanks to many collaborators!

- **The HACC* Team (in alphabetical order with main focus):**

  David Daniel (Poisson solver, FFT), Patricia Fasel (MPI framework), Hal Finkel (Tree, BG/Q, I/O), Nick Frontiere (GPU, FFT), Salman Habib (Overall design, algorithms), Katrin Heitmann (Validation), Vitali Morozov (BG/Q optimization), Adrian Pope (Cell, P3M, BG/Q)

- **The Analysis Team (in alphabetical order with main focus):**

  Chung-Hsing Hsu (Halo finder), Joe Insley (Visualization), Eve Kovacs (Synthetic skies), Juliana Kwan (Synthetic skies, emulators), Nan Li (Lensing), Steve Rangel (Feature tracking, density estimators), Tom Peterka (Density estimator), Chris Sewell (GPU halo finder), Martin White (Synthetic skies, BAO), Nikhil Padmanabhan (BAO), George Zagaris (CosmoTools)

*** HACC = Hardware/Hybrid Accelerated Cosmology Code**

Modern Cosmology and Sky Maps

- Modern cosmology is the story of mapping the sky in multiple wavebands
- Maps cover measurements of objects (stars, galaxies) and fields (temperature)
- Maps can be large (Sloan Digital Sky Survey has~200 million galaxies, many billions for planned surveys)
- Statistical analysis of sky maps
- All precision cosmological analyses constitute a statistical inverse problem: **from sky maps to scientific inference**
- Therefore: *No* cosmology without (large-scale) computing

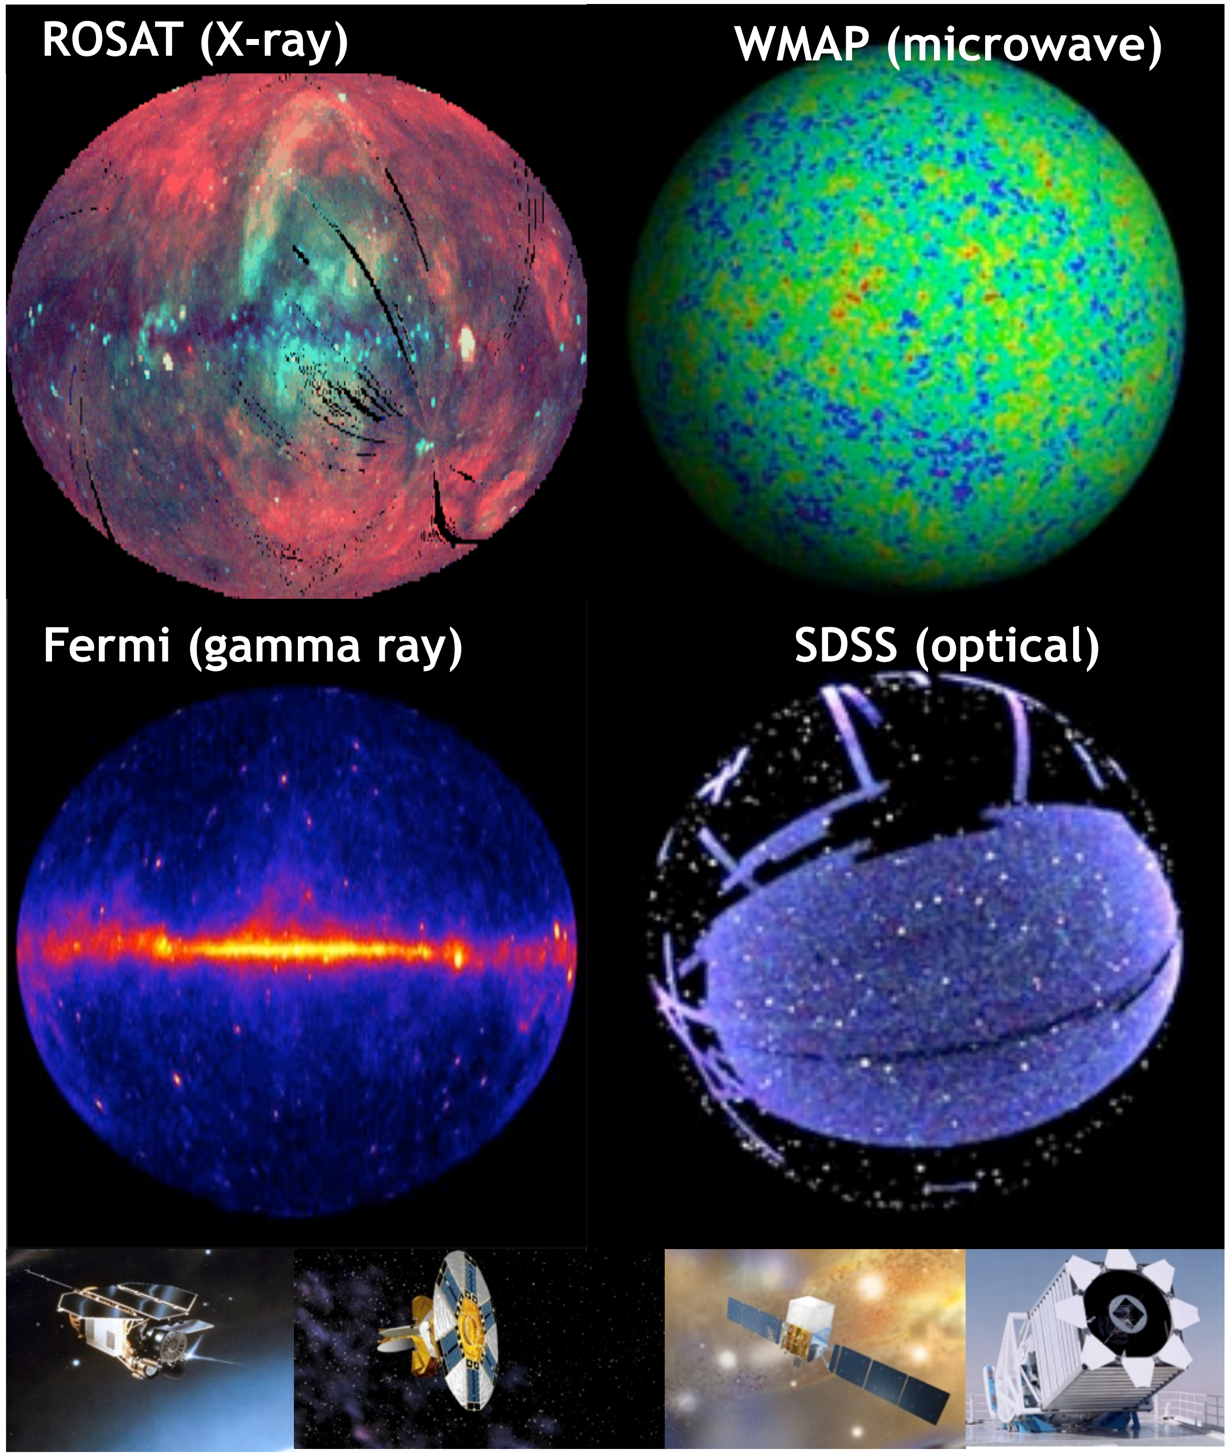

Observations from the SDSS: positions of 1,000,000 galaxies with redshifts (and therefore distance) leading to a 3-D map

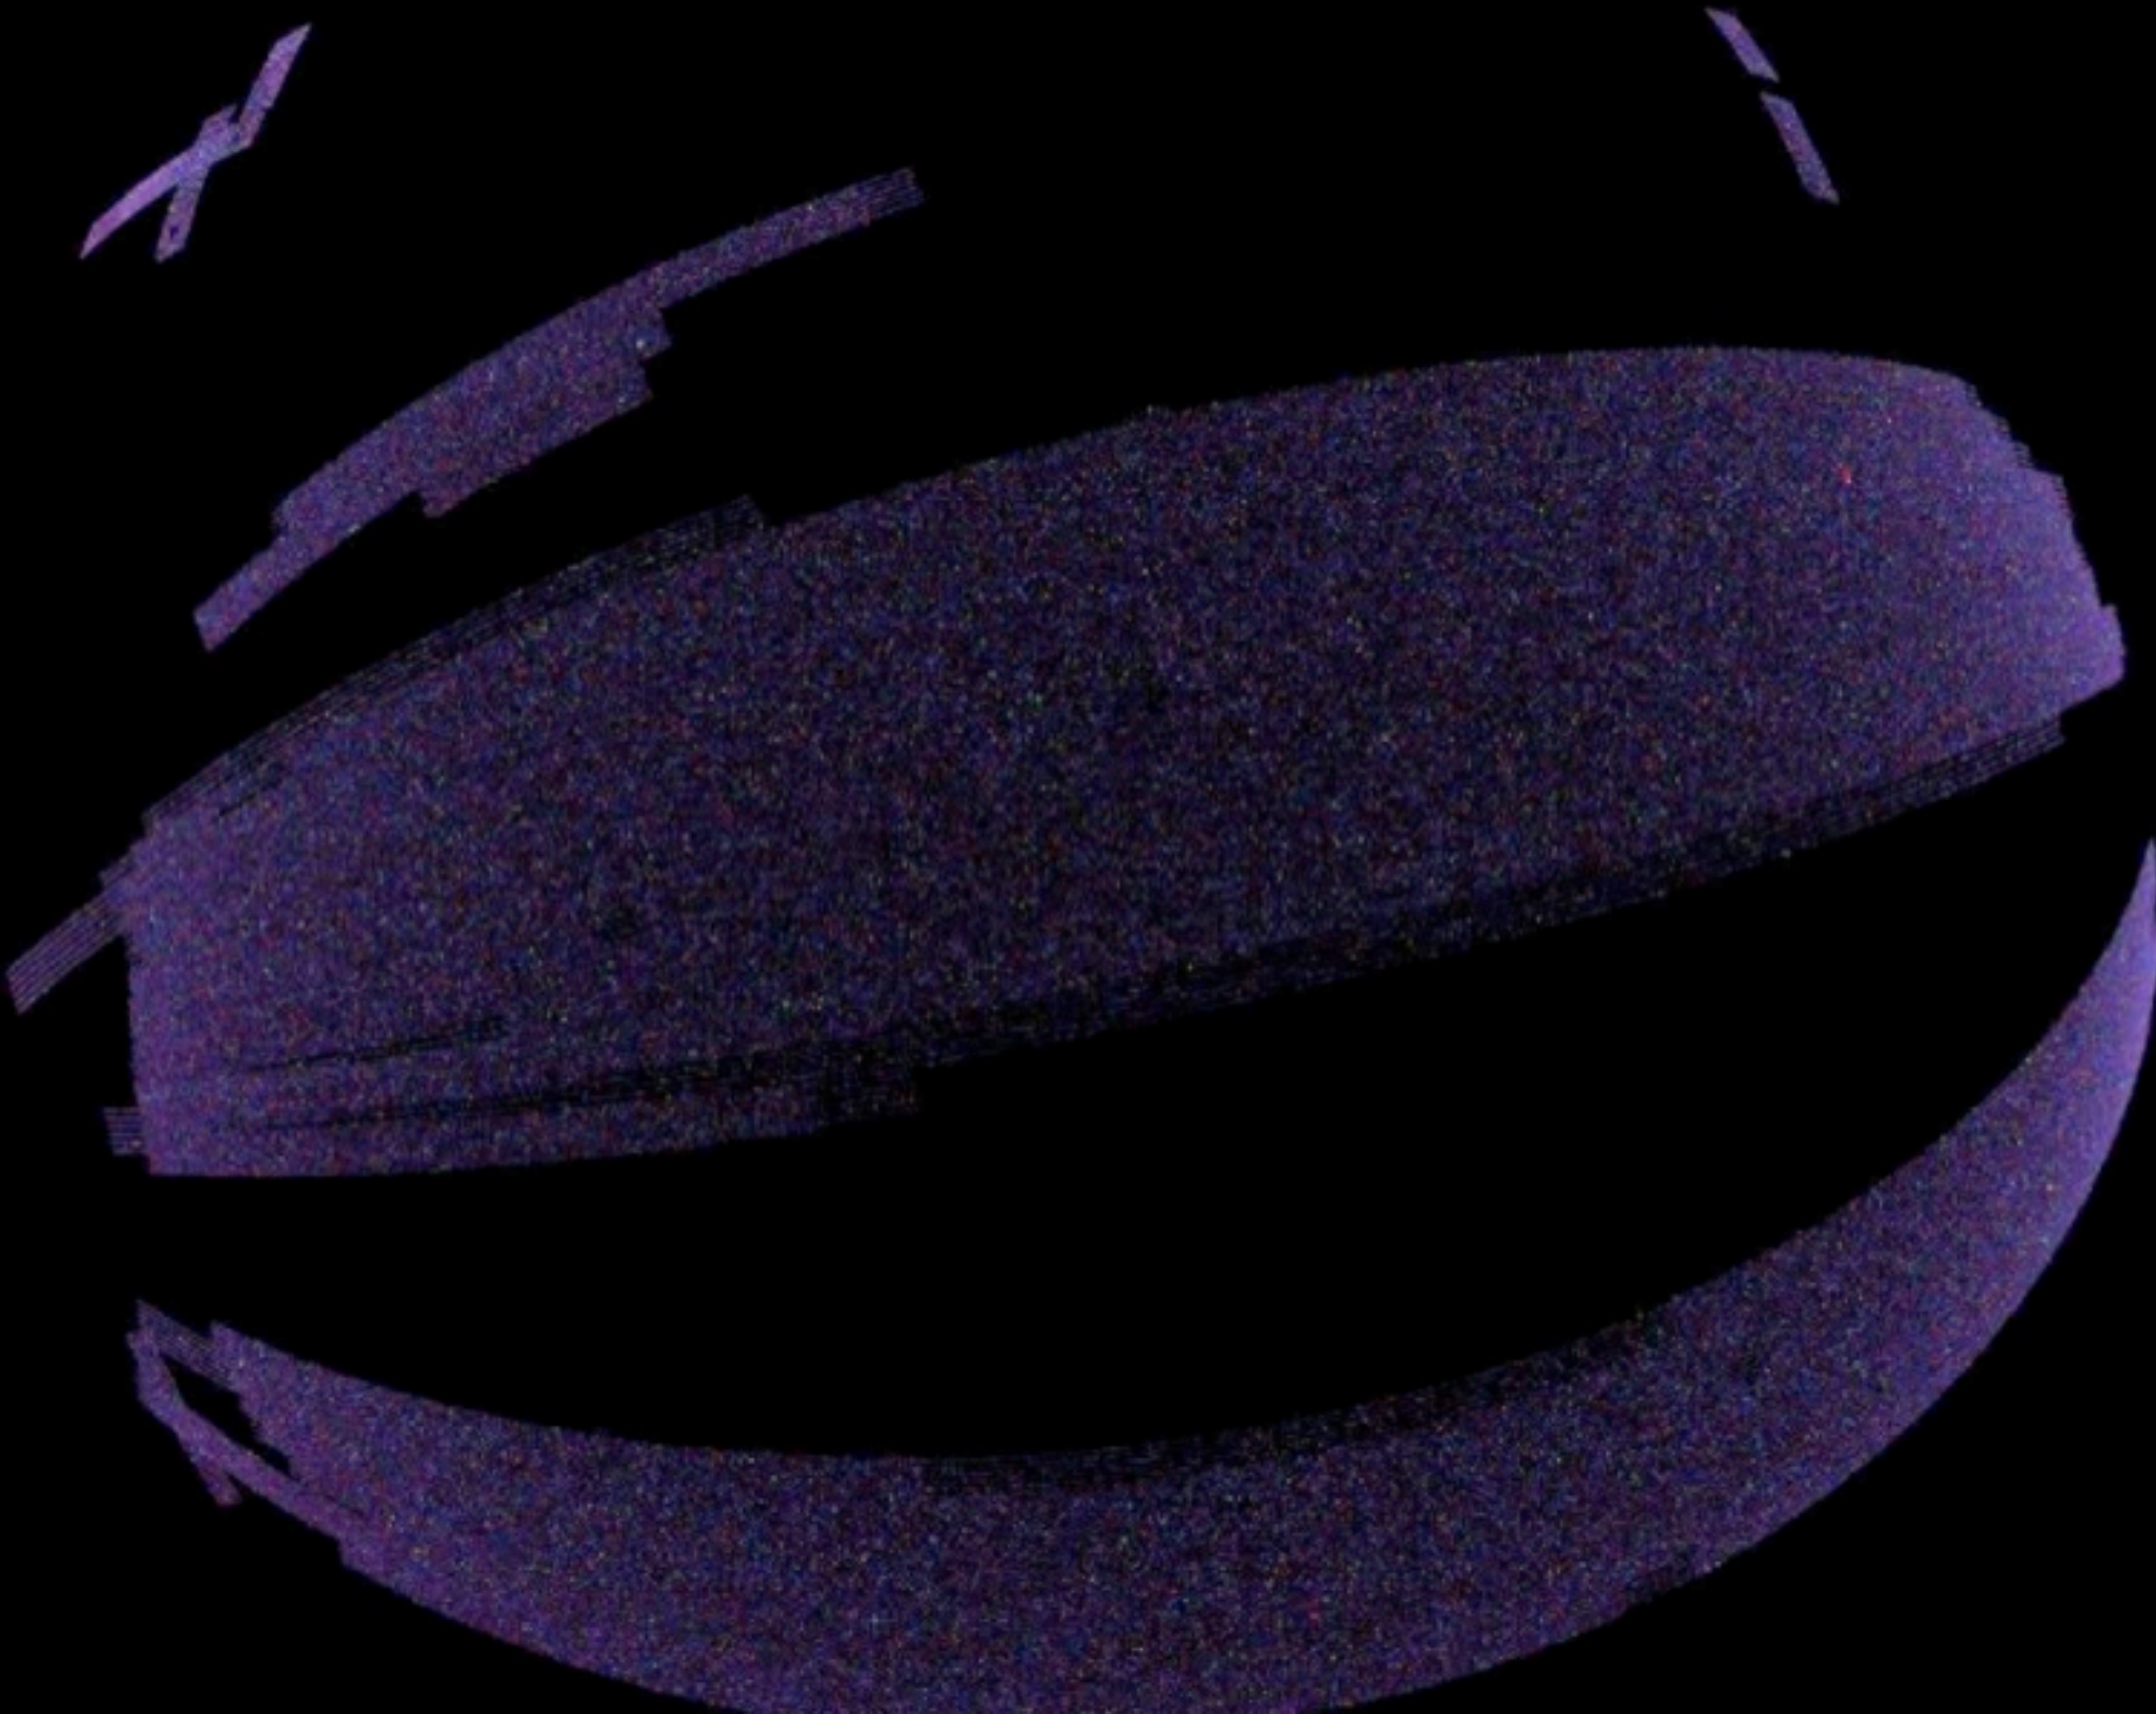

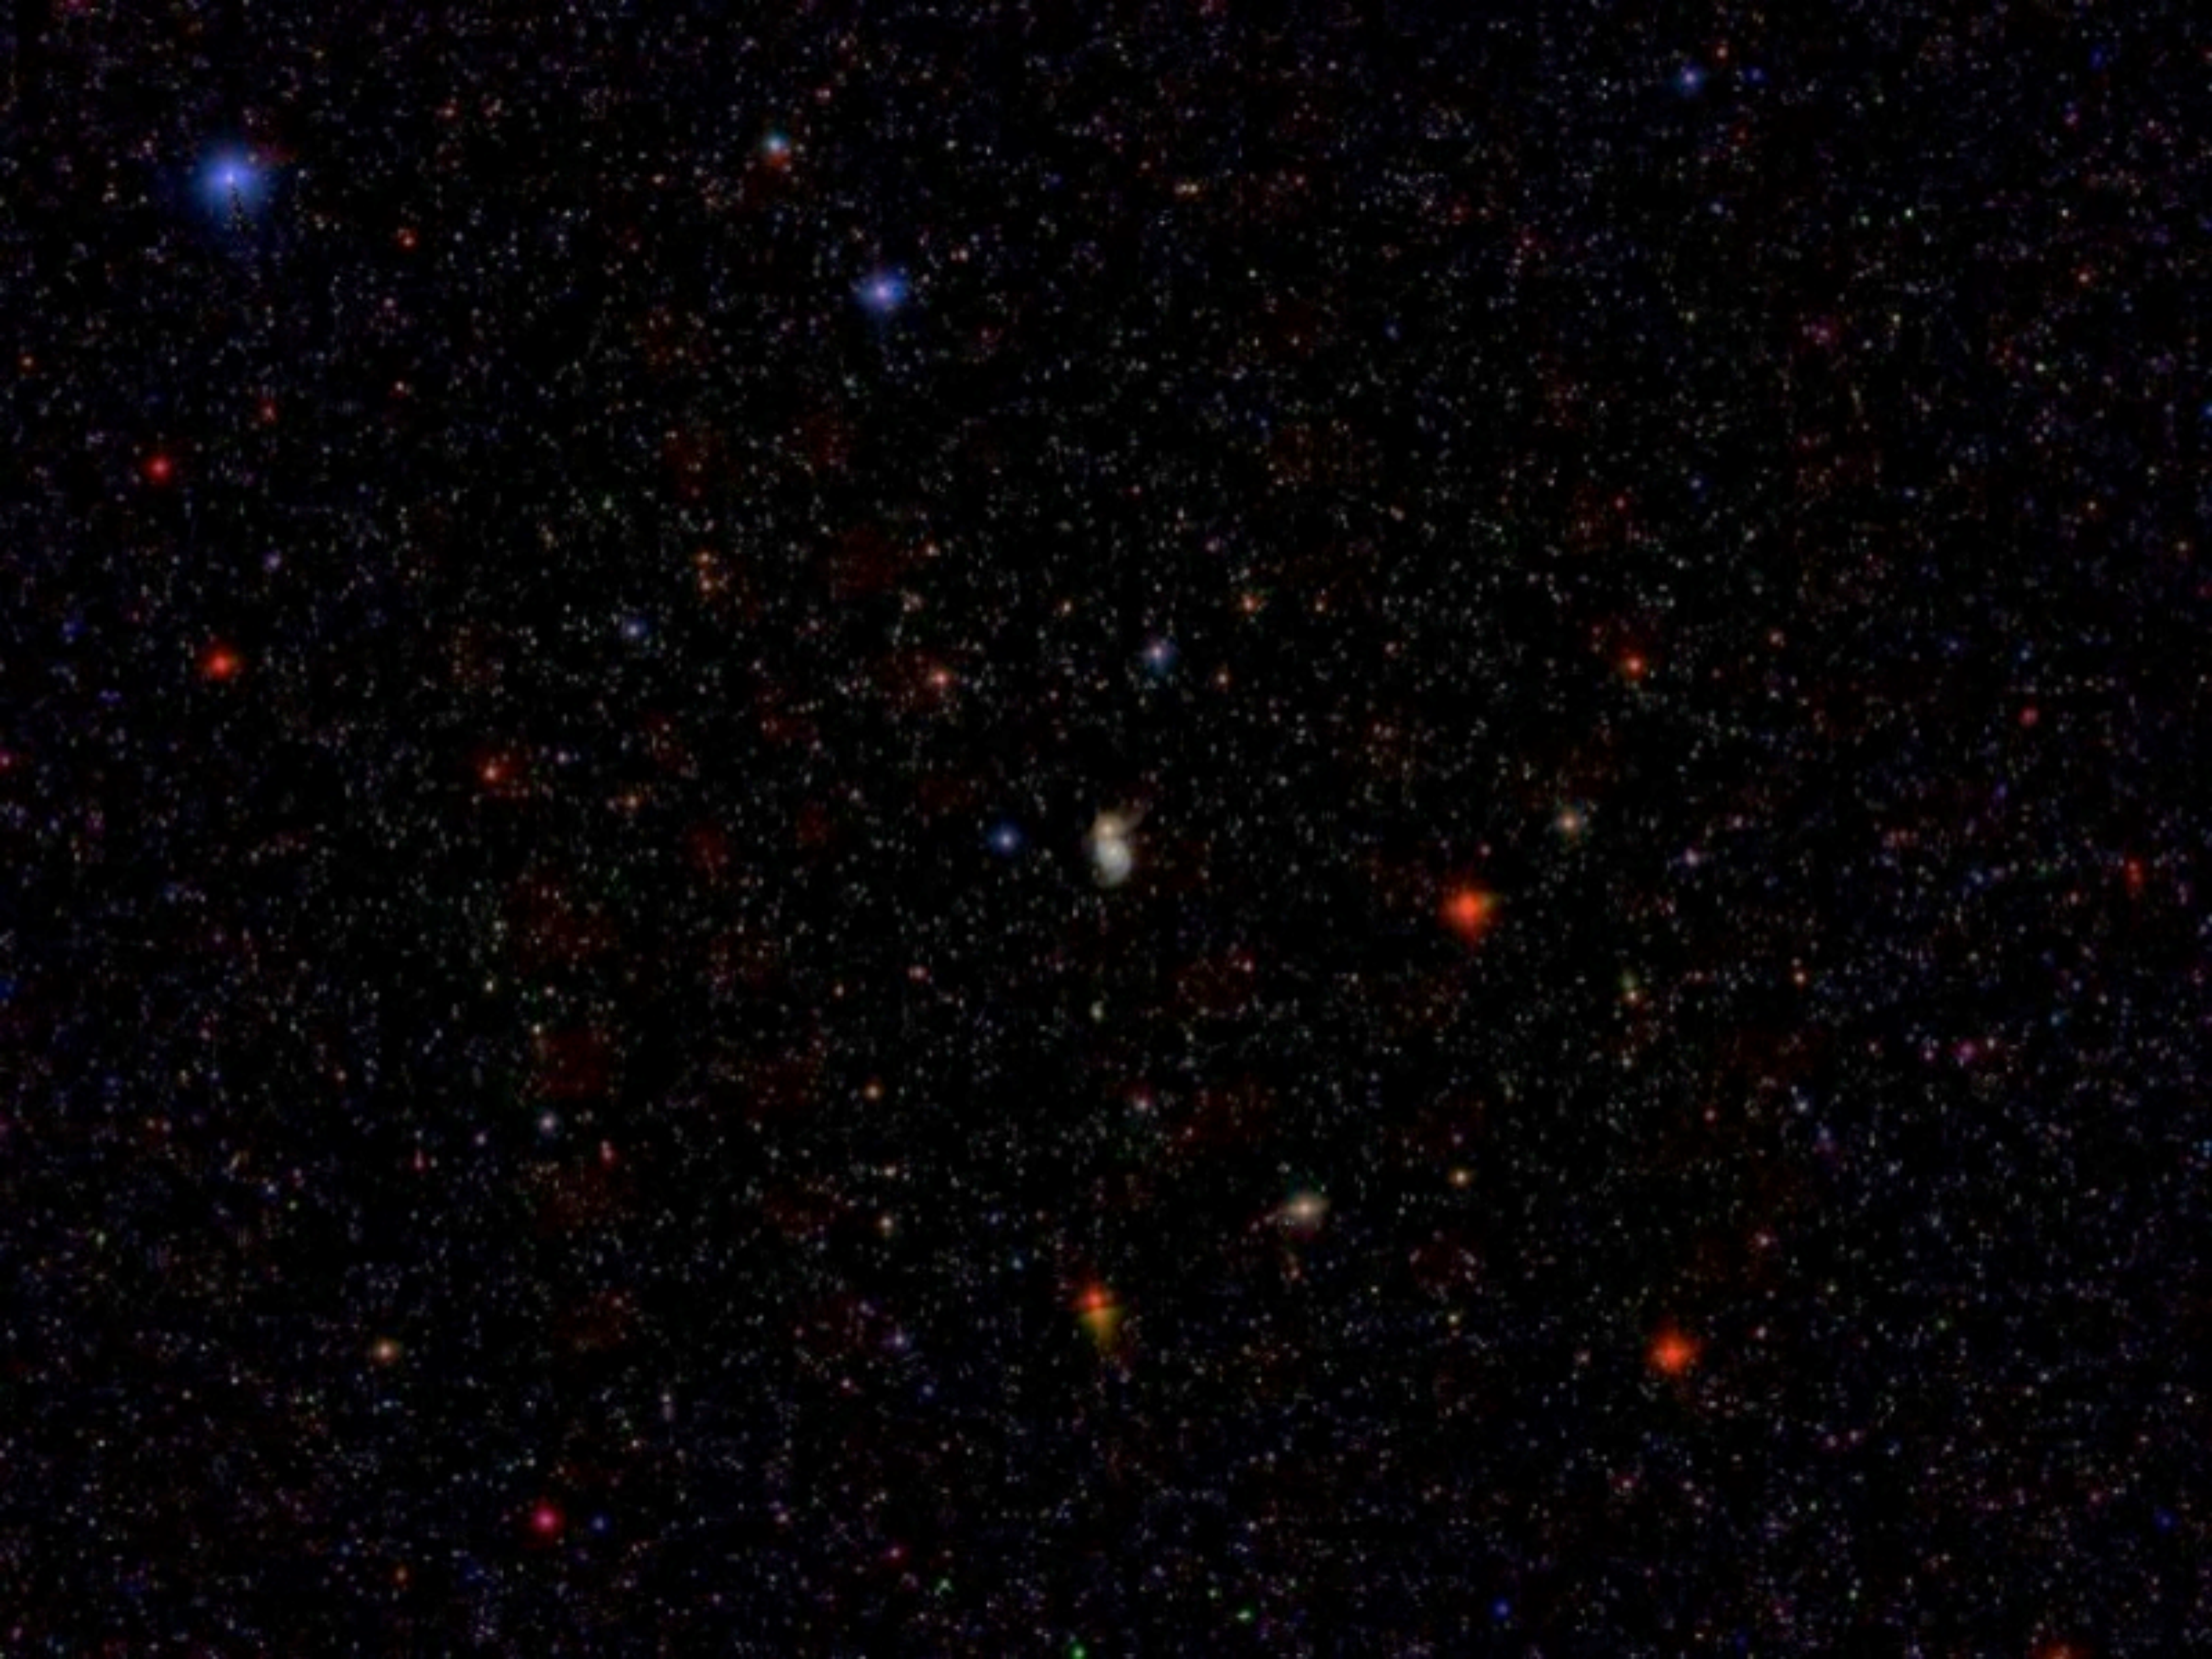

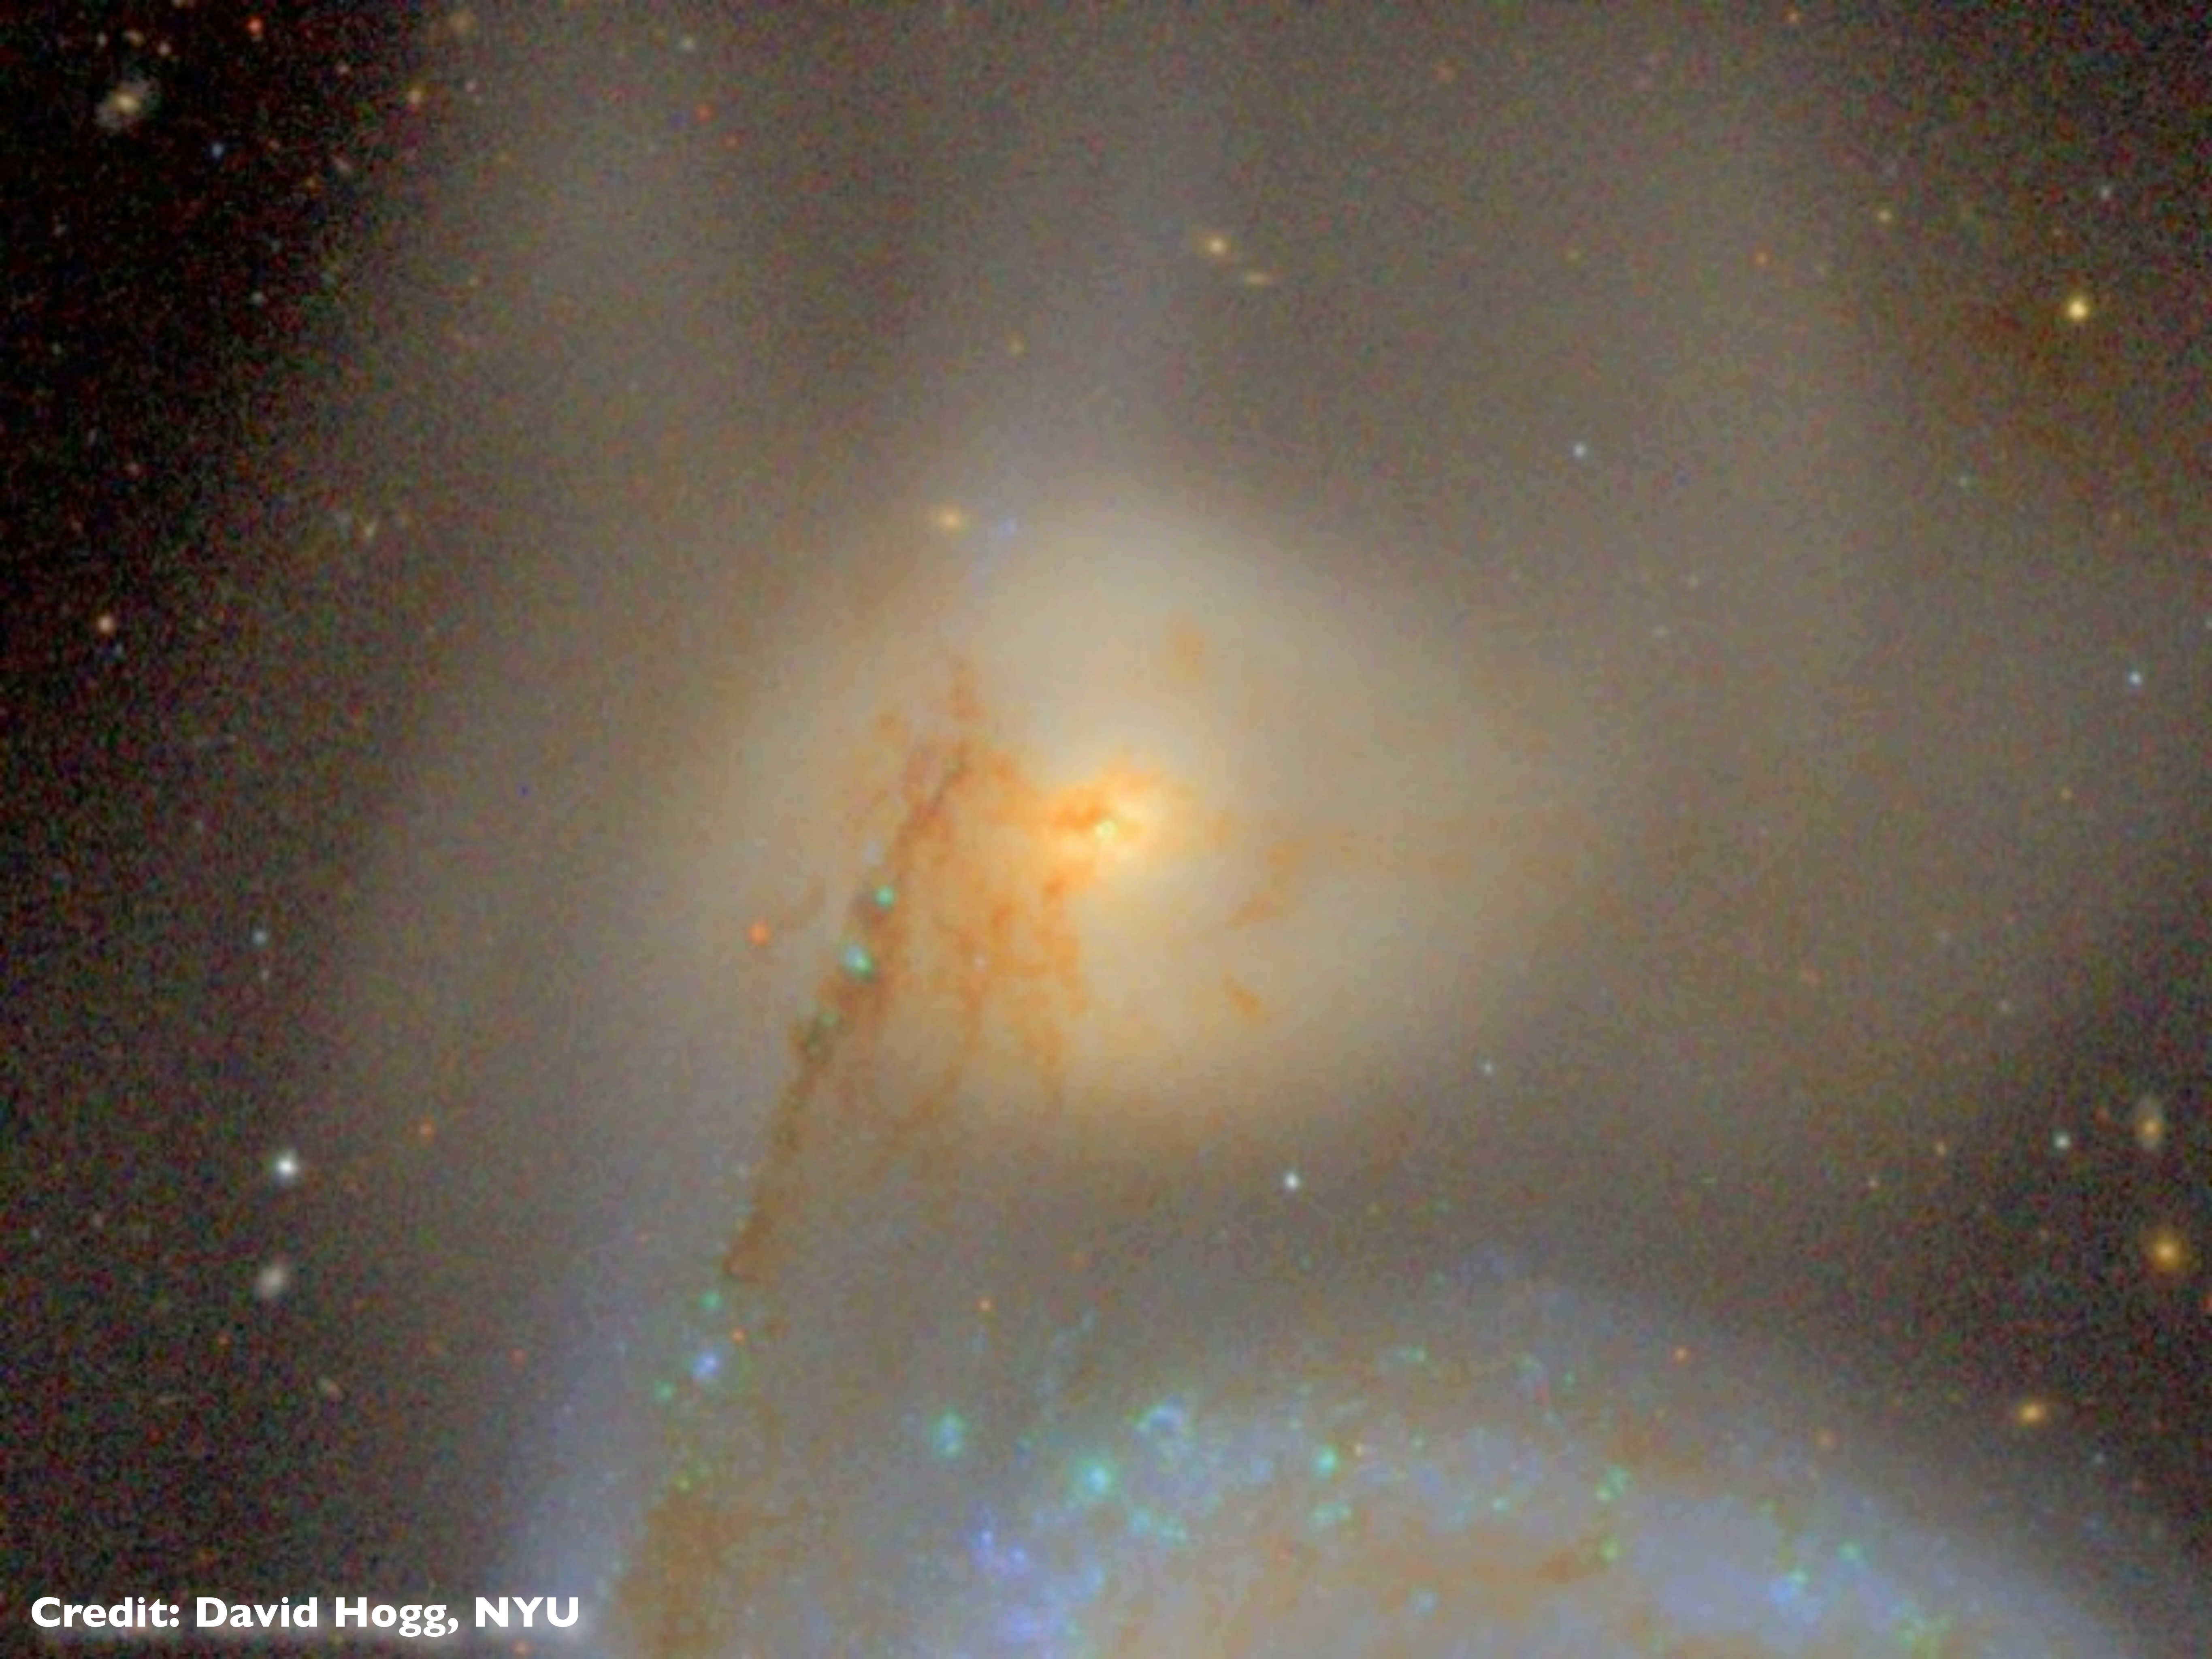
Credit: David Hogg, NYU

The Content of the Universe: It's dark!

- **Dark Energy:** Multiple observations show that the expansion of the Universe is accelerating (first in 1998, Nobel prize 2011)
- Questions: What is it? Why is it important now? Being totally ignorant, currently our main task is to characterize it better and exclude some of the possible explanations
- Independent of what we find, we will learn new, fundamental physics, this is not just the hunt for a couple numbers!
- **Dark Matter:** Observations show that ~27% of the matter in the Universe is "dark", i.e. does not emit or absorb light
- So far: indirect detection, aims: characterize nature of dark matter and detect the actual dark matter particle

~95% of the Universe is "dark" -- we do not understand the nature and origin of dark energy and dark matter.

Cosmology with HACC: The Nature of the Dark Universe

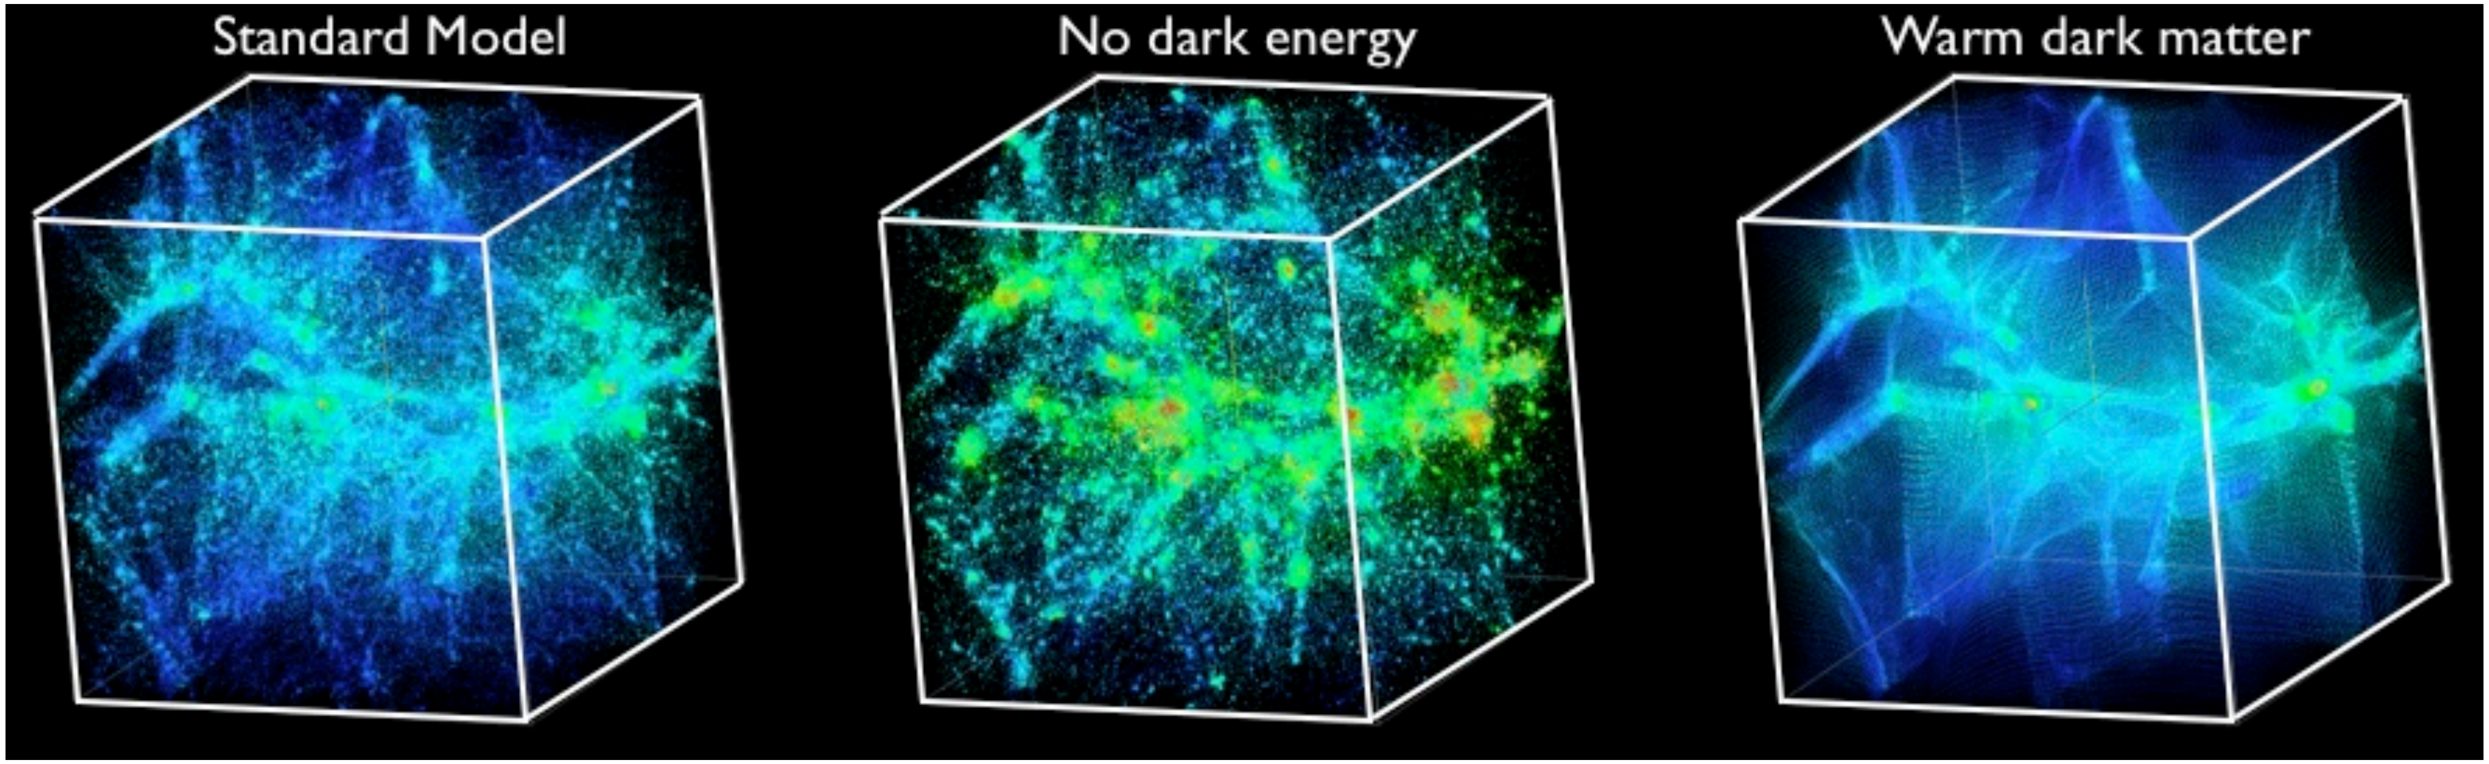

- **Exploration of different dark energy models:** How would a dark energy model beyond Einstein's cosmological constant alter the distribution of matter (and galaxies) in the Universe?
- **Exploration of dark matter and neutrinos in the Universe:** What can cosmology tell us about different matter components in the Universe?
- **Cosmic Emulation:** How can we explore different cosmological models in an efficient way if the simulations are so expensive (ask me after the talk ...)

Structure Formation: The Basic Paradigm

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

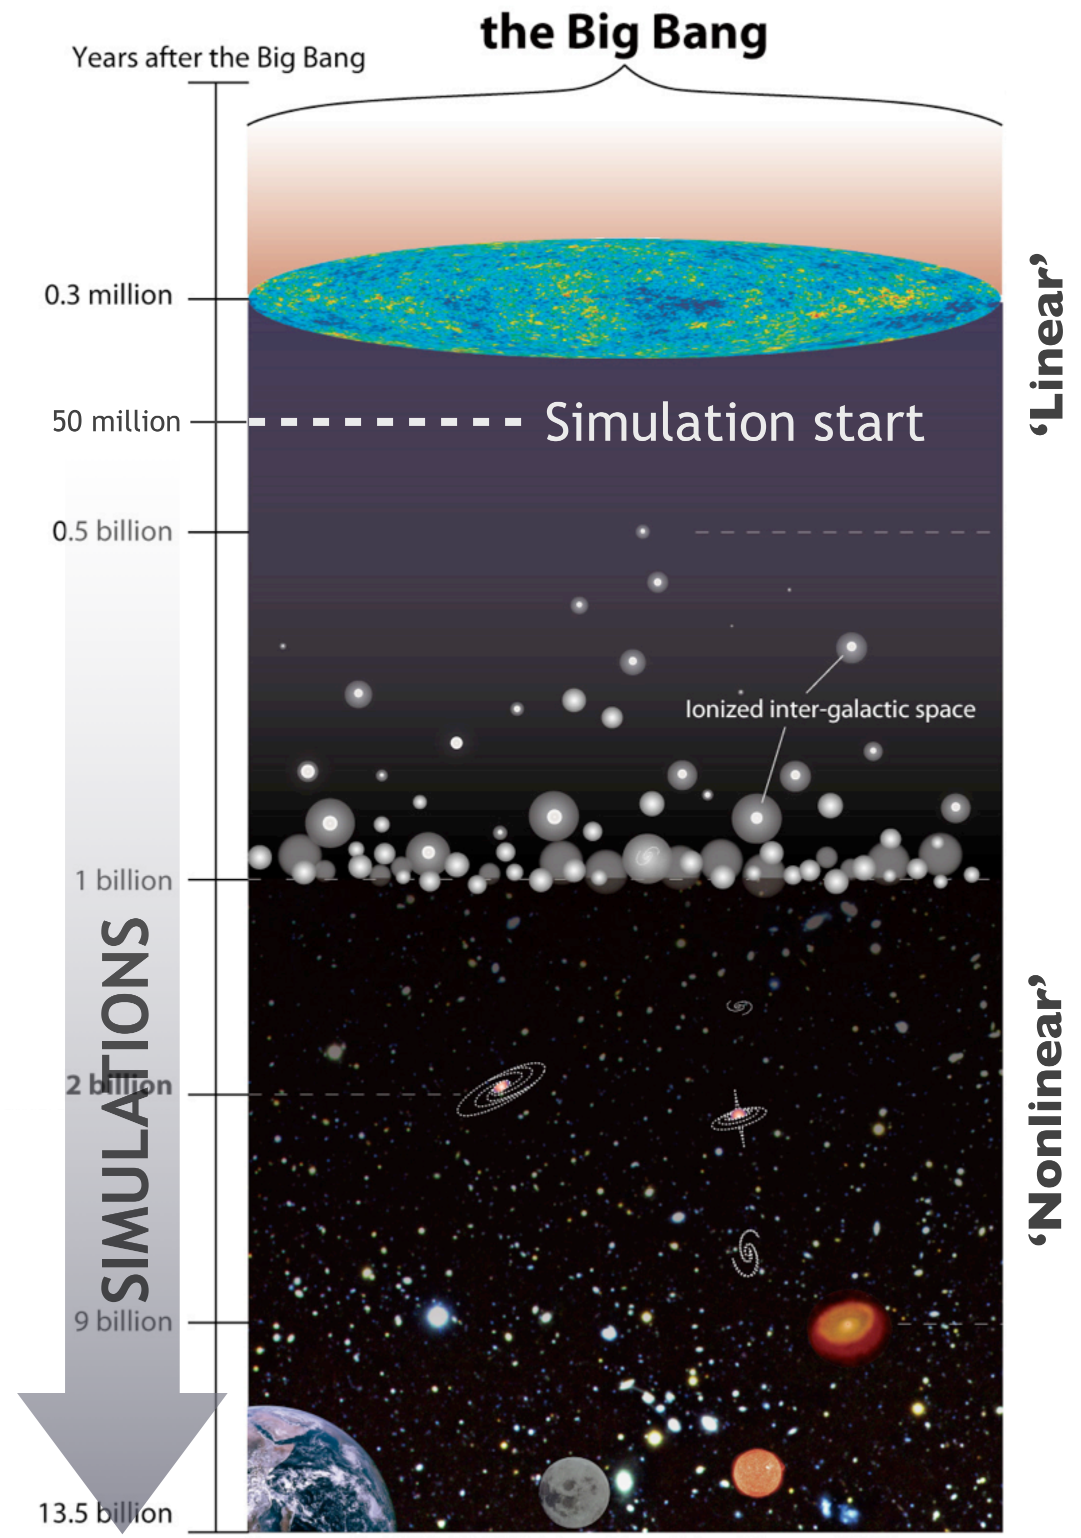

Connecting Theory and Observations

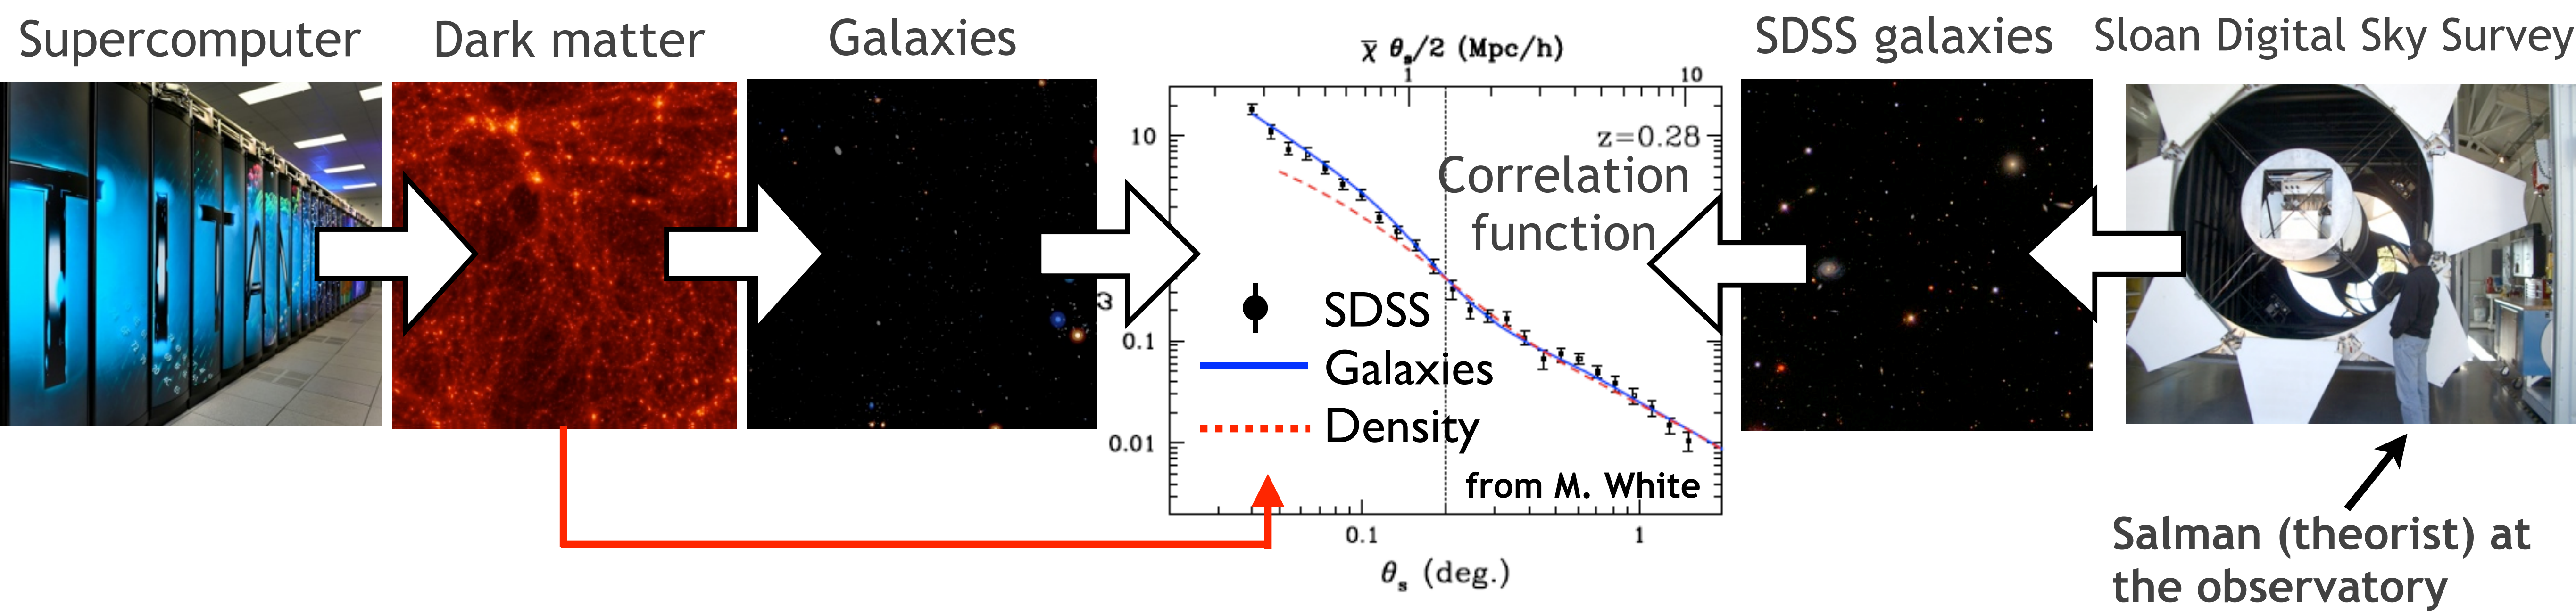

- Simulate the formation of the large scale structure of the Universe via dark matter tracer particles, taking dark energy into account in expansion history
- Measure the high-density peaks (dark matter halos) in the mass distribution
- "Light traces mass" to first approximation, therefore populate the halos with galaxies, number of galaxies depends on mass of halo (constraints from observations)
- Galaxy population prescription (hopefully) independent of cosmological model
- Future challenges: Error bars on measurements are shrinking, surveys cover more volume and resolve fainter and fainter galaxies

Connecting Theory and Observations

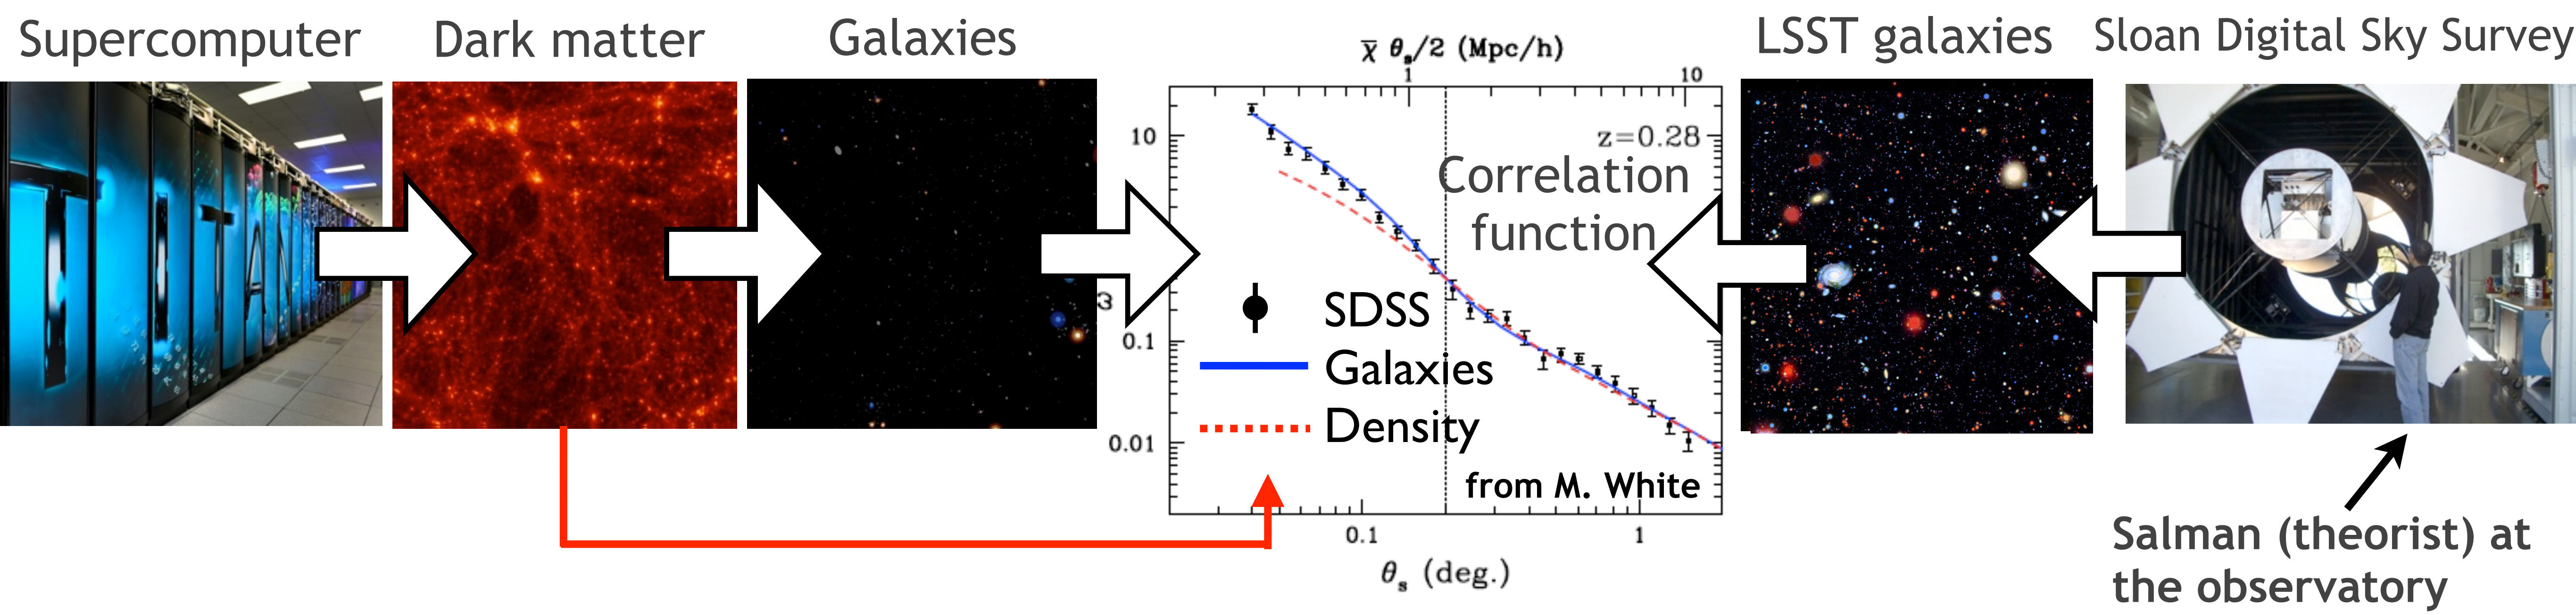

- Simulate the formation of the large scale structure of the Universe via dark matter tracer particles, taking dark energy into account in expansion history
- Measure the high-density peaks (dark matter halos) in the mass distribution
- “Light traces mass” to first approximation, therefore populate the halos with galaxies, number of galaxies depends on mass of halo (constraints from observations)
- Galaxy population prescription (hopefully) independent of cosmological model
- Future challenges: Error bars on measurements are shrinking, surveys cover more volume and resolve fainter and fainter galaxies

Computing the Universe

- Gravity dominates at large scales, key task: solve the Vlasov-Poison equation (VPE)
- VPE is 6-D and cannot be solved as PDE, therefore N-body methods
- Particles are tracers of the dark matter in the Universe, mass typically at least ~10^9 $M_\odot$

$$m_p \sim V/n_p$$

- At smaller scales, add gas physics, feedback etc., sub-grid modeling inevitable

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

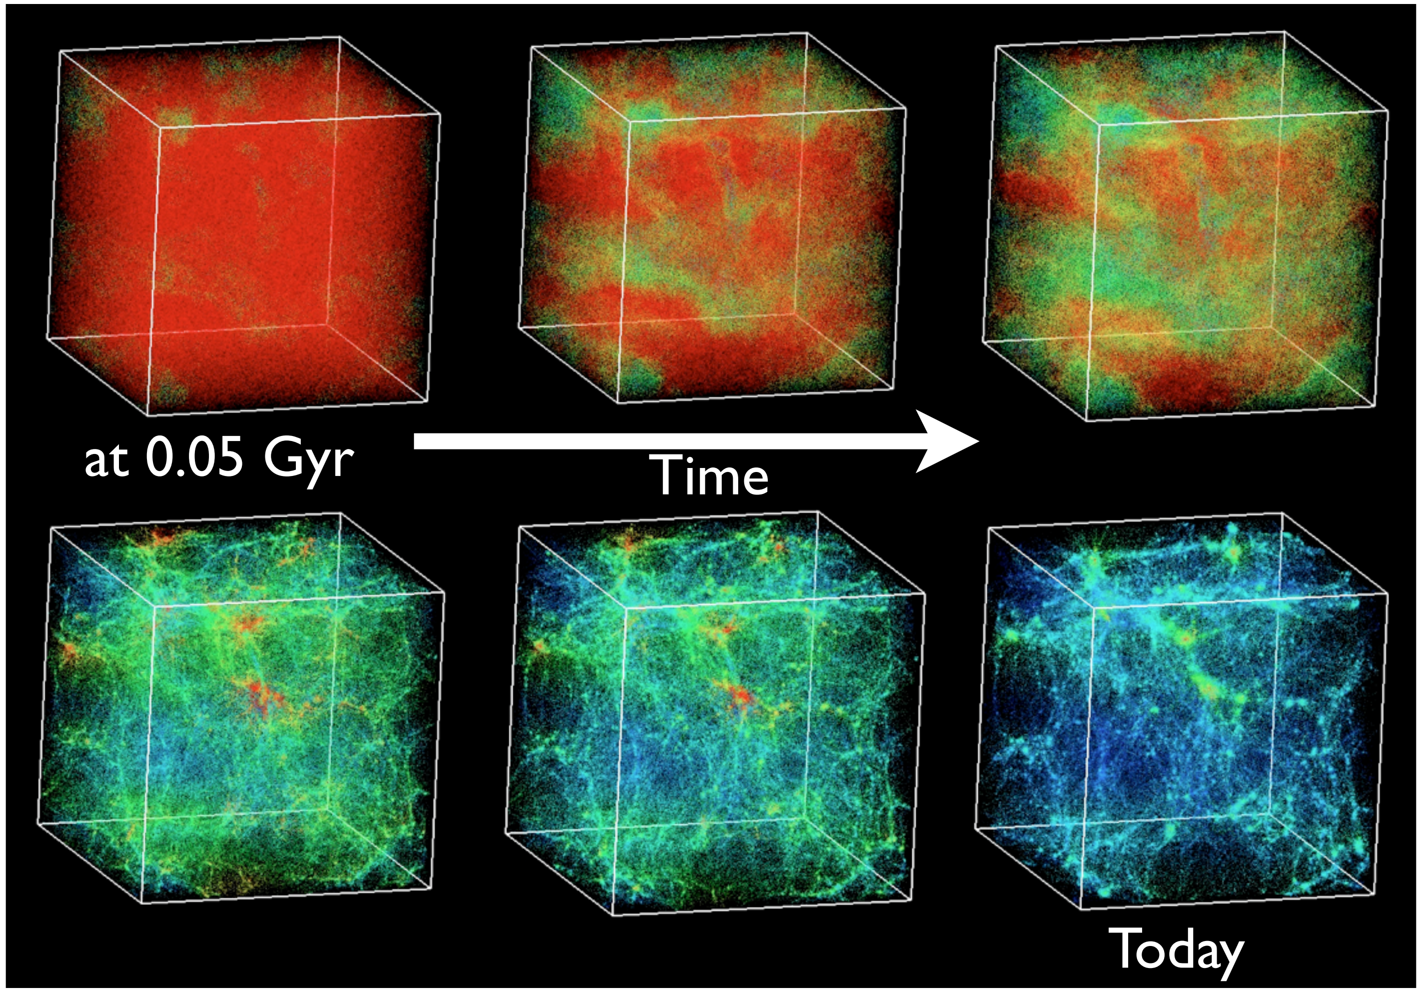

> "The Universe is far too complicated a structure to be studied deductively, starting from initial conditions and solving the equations of motion."
> *Robert Dicke (Jayne Lectures, 1969)*

The HACC Story Begins ...

Andrew White Dec 7, 2007 + What if you had a petaflop/s

- **... with an email:** Los Alamos National Lab offers the opportunity to run open science projects on the fastest supercomputer in the world for the first six months of the machine's existence: **Roadrunner**
- **Roadrunner:** First machine to achieve Petaflop performance via Cell-acceleration, CPU/Cell hybrid architecture (more details later) (equivalent to ~200,000 laptops)
- **The Challenges:**
  - The machine has a "crazy" architecture, requiring major code re-designs and rewrites (we ended up writing a brand **new code**)
  - Roadrunner probably **one of a kind**, code-design needs to be flexible and **portable** to other future architectures
- **Cosmologists are poor -- so we took on the challenge!**
- **Outcome:** **MC^3 (Mesh-based Cosmology Code on the Cell) which later morphed into HACC, N-body code to simulate large-scale structure formation in the Universe**

The HACC Story Begins ...

☐ Andrew White   Dec 7, 2007   + What if you had a petaflop/s

- **... with an email:** Los Alamos National Lab offers the opportunity to run open science projects on the fastest supercomputer in the world for the first six months of the machine's existence: **Roadrunner**
- **Roadrunner:** First machine to achieve Petaflop performance via Cell-acceleration, CPU/Cell hybrid architecture (more details later) (equivalent to ~200,000 laptops)
- **The Challenges:**

  Andy White: "forward-looking"

  - The machine has a "~~crazy~~" architecture, requiring major code re-designs and rewrites (we ended up writing a brand **new code**)
  - Roadrunner probably **one of a kind**, code-design needs to be flexible and **portable** to other future architectures
- **Cosmologists are poor -- so we took on the challenge!**
- **Outcome:** **MC³ (Mesh-based Cosmology Code on the Cell) which later morphed into HACC, N-body code to simulate large-scale structure formation in the Universe**

So we started thinking --

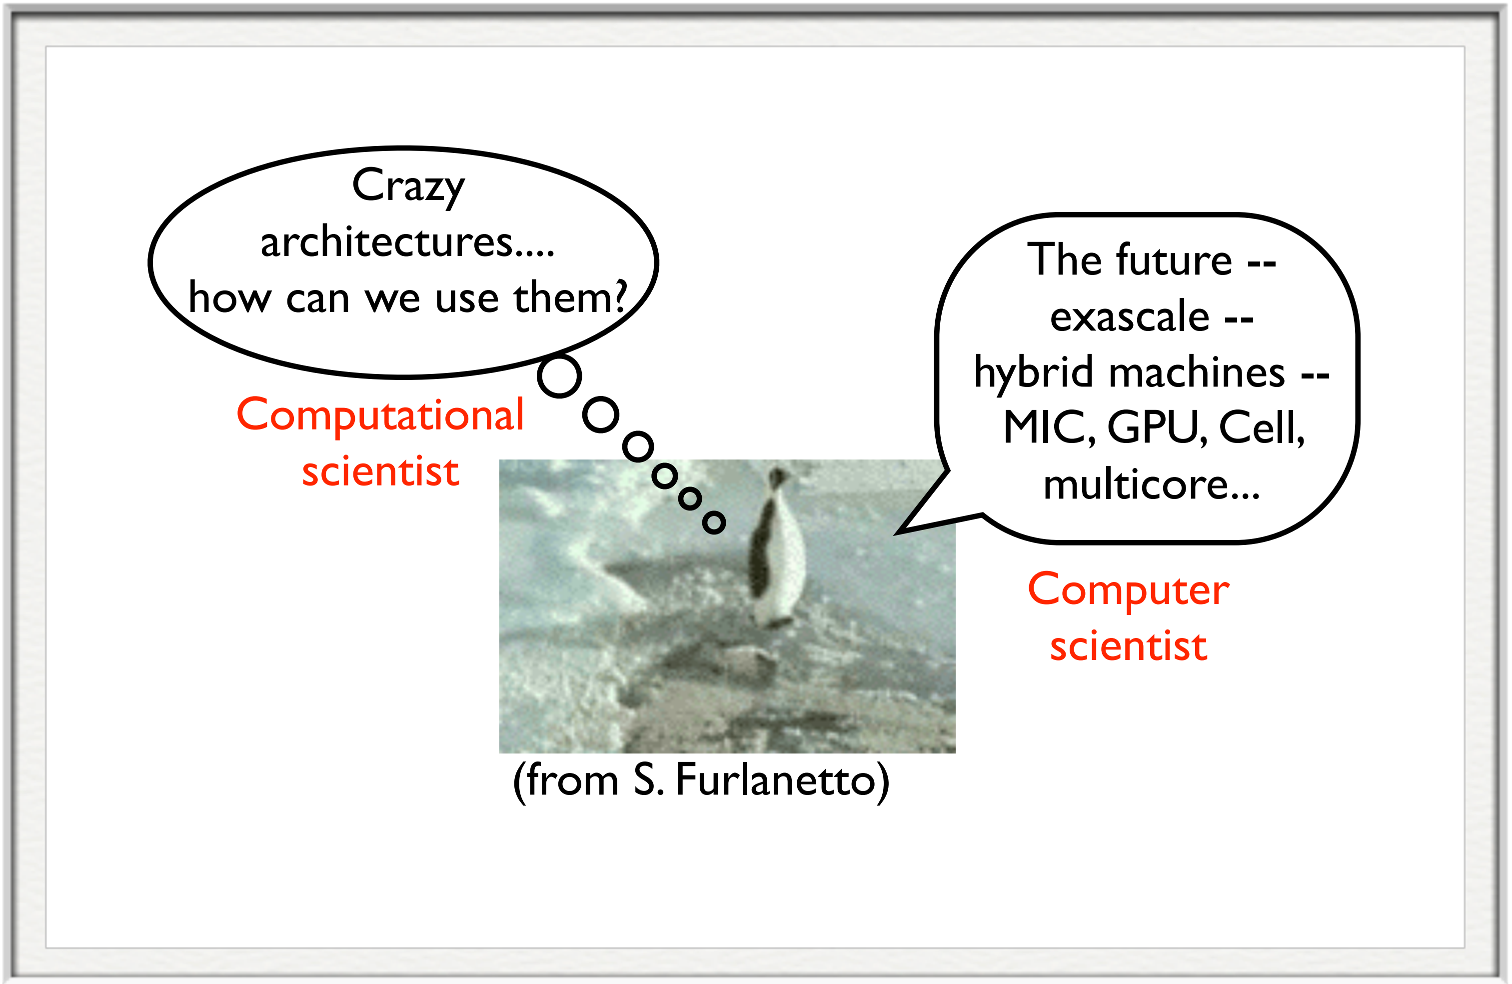

(from S. Furlanetto)

The Roadrunner Architecture

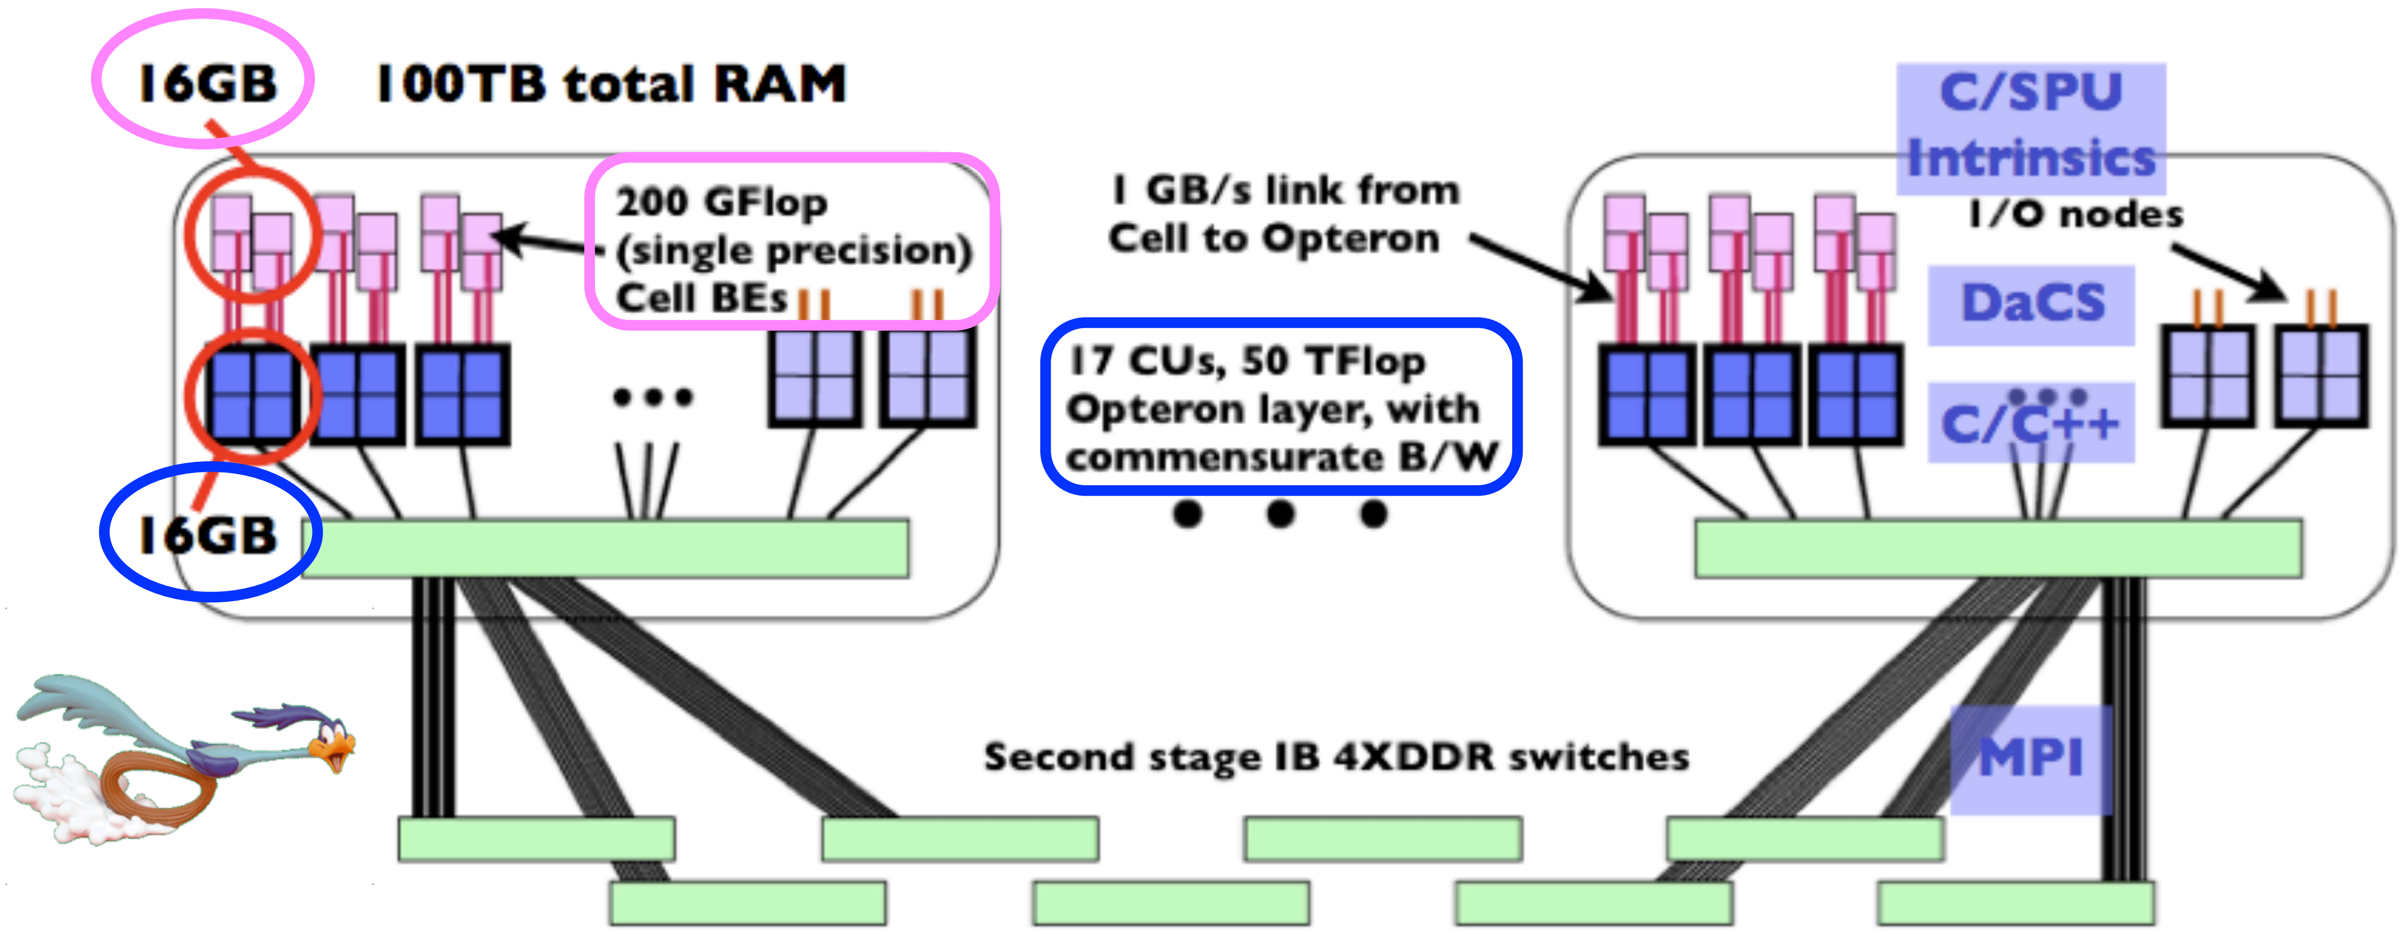

- **Opterons** have little compute (5% of total compute) but half the memory and balanced communication, for N-body codes, memory is limiting factor, so want to make best use of CPU layer
- **Cells dominate** the compute but communication is poor, 50-100 times out of balance (also true for CPU/GPU hybrid systems)
- **Multi-layer programming model:** C/C++/MPI (Message Passing Interface) for Opterons, C/Cell-intrinsics for Cells (OpenCL or Cuda for GPUs)

Design Challenges and Solutions for MC^3

- **Challenges (summarized from last slide):**
  - Opterons have half of the machine's memory, balanced communication, but not much compute, standard programming paradigm, C/C++/MPI
  - Cells have other half of machine's memory, slow communication, **lots** of compute (95% of machine's compute power), new language required
- **Design desiderata:**
  - Distribute memory requirements on both parts of the machine (different on GPUs!)
  - Give the Cell lots of (communication limited) work to do, make sure that Cell part is easy to code and later on easy to replace by different programming paradigm
- **Our Solution: Long-range/short range force splitting + overload concept**
  - **Long-range:** Particle-Mesh solver, C/C++/MPI, **unchanged for different architectures**, FFT performance dictates scaling (custom pencil decomposed FFT)
  - **Short-range:** **Depending on node architecture** switch between tree and particle-particle algorithm; tree needs "thinking" (building, walking) but computationally less demanding (BG/Q, X86), PP easier but computationally more expensive (Cell, GPU), P3M and TreePM implementations
  - **Overload concept: Allows for easy swap of short-range solver and minimization of communication (reassignment of passive/active in regular intervals)**

‘HACC In Pictures’

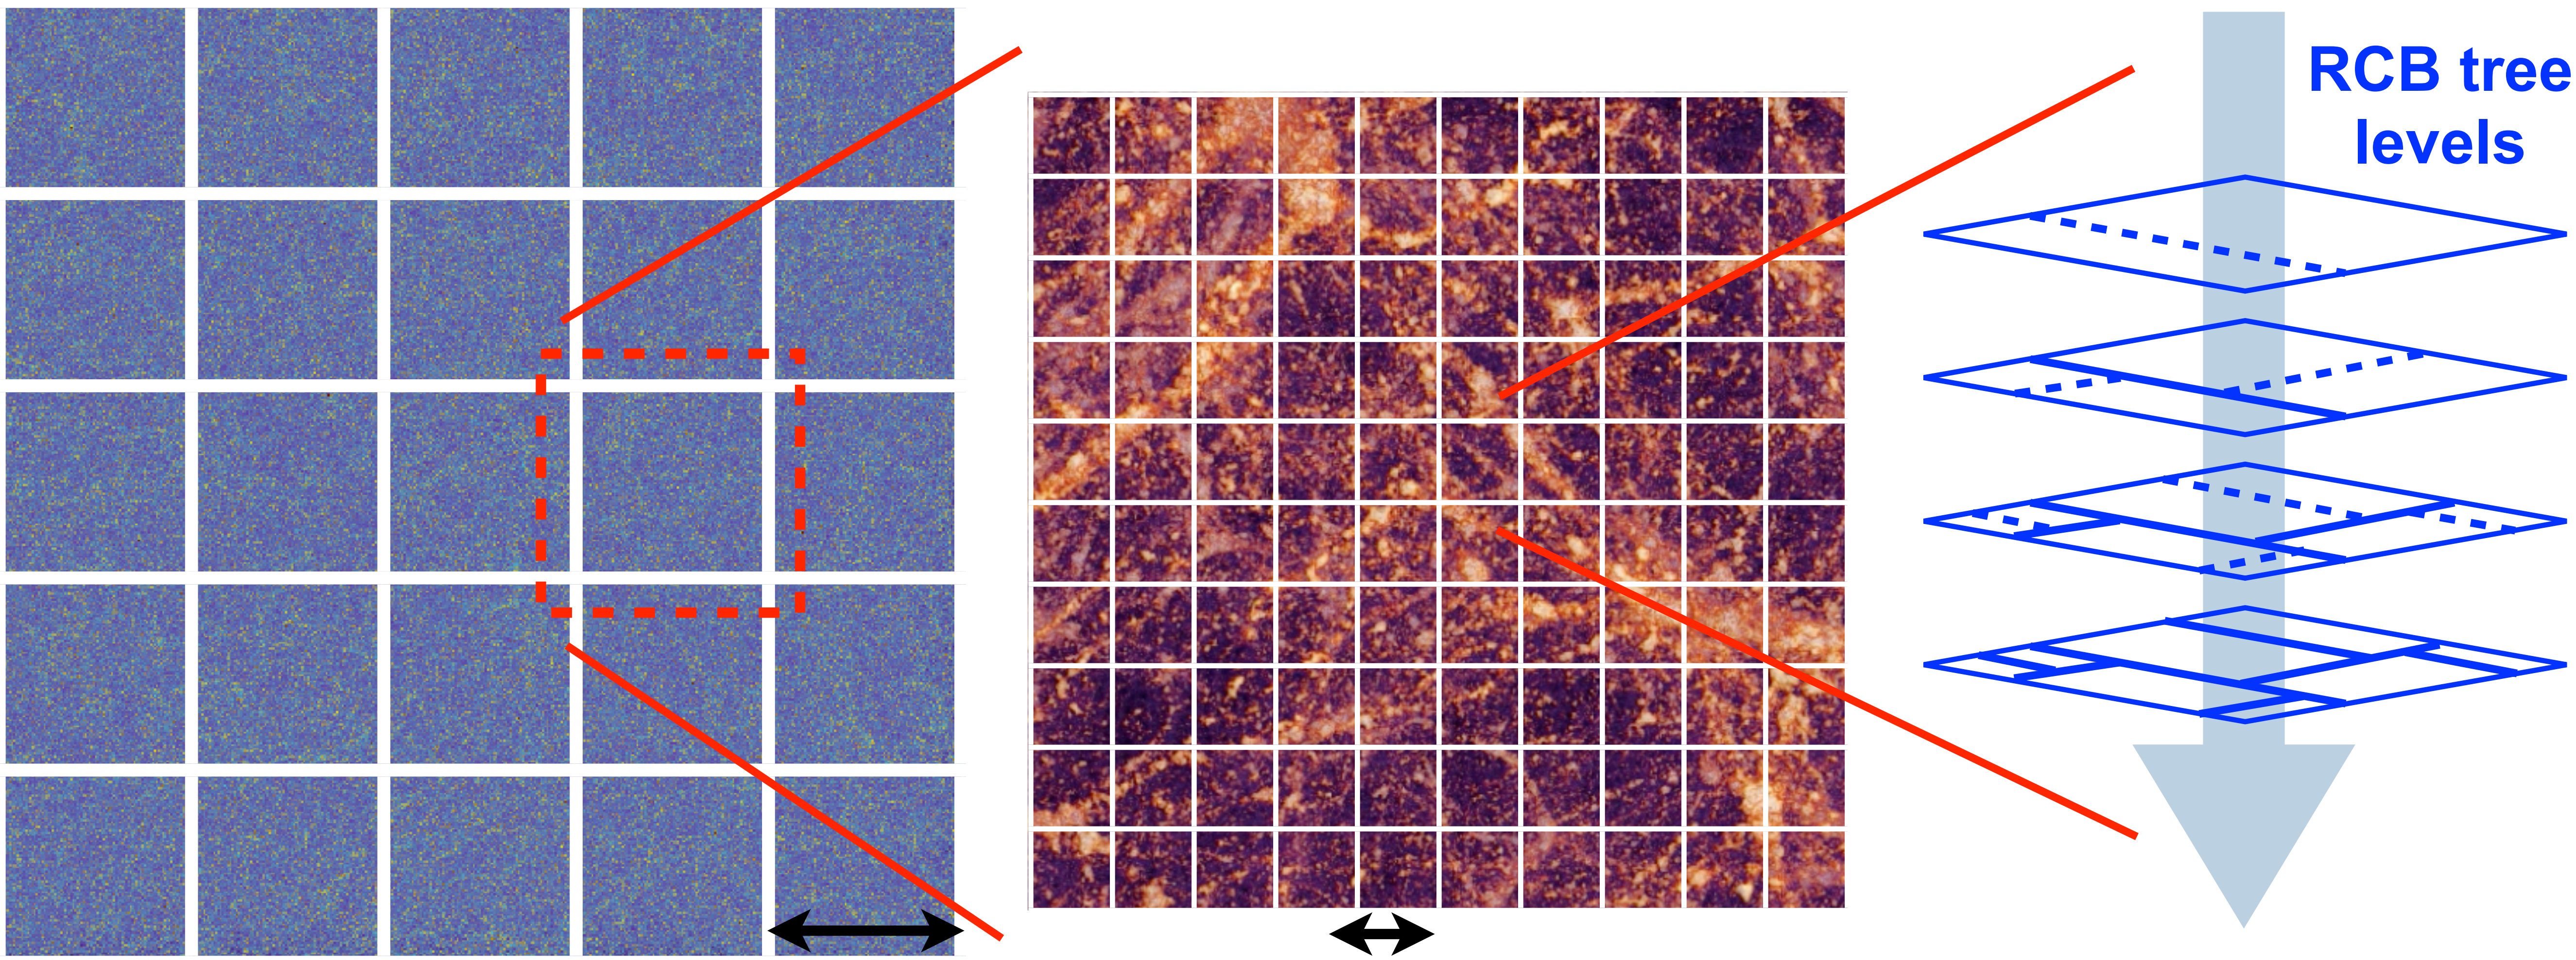

HACC Top Layer:
3-D domain decomposition with particle replication at boundaries (‘overloading’) for Spectral PM algorithm (long-range force)

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

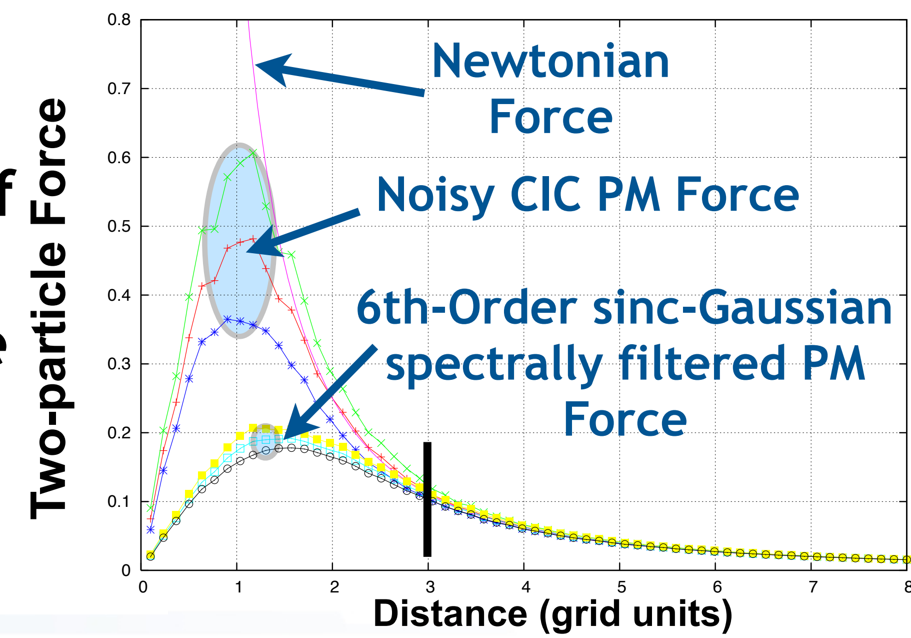

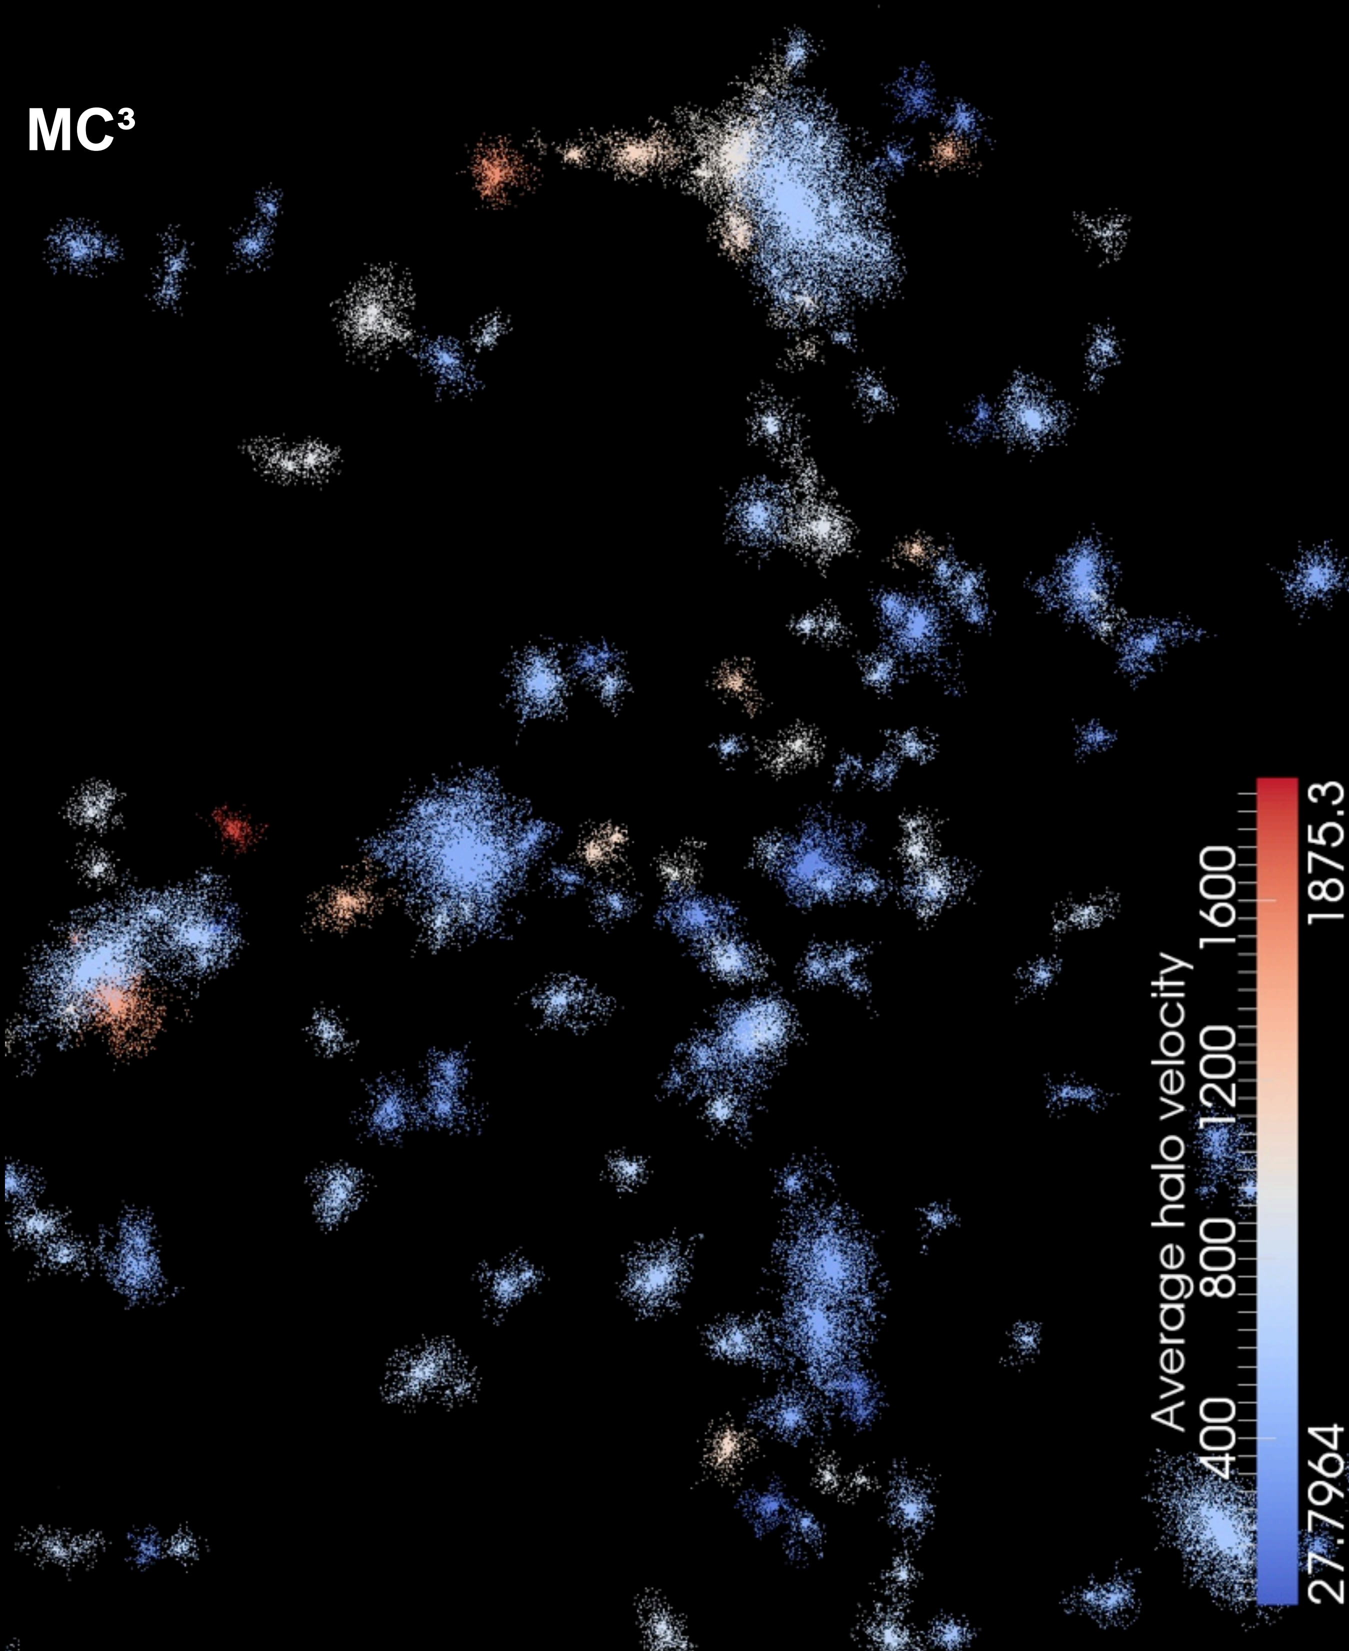

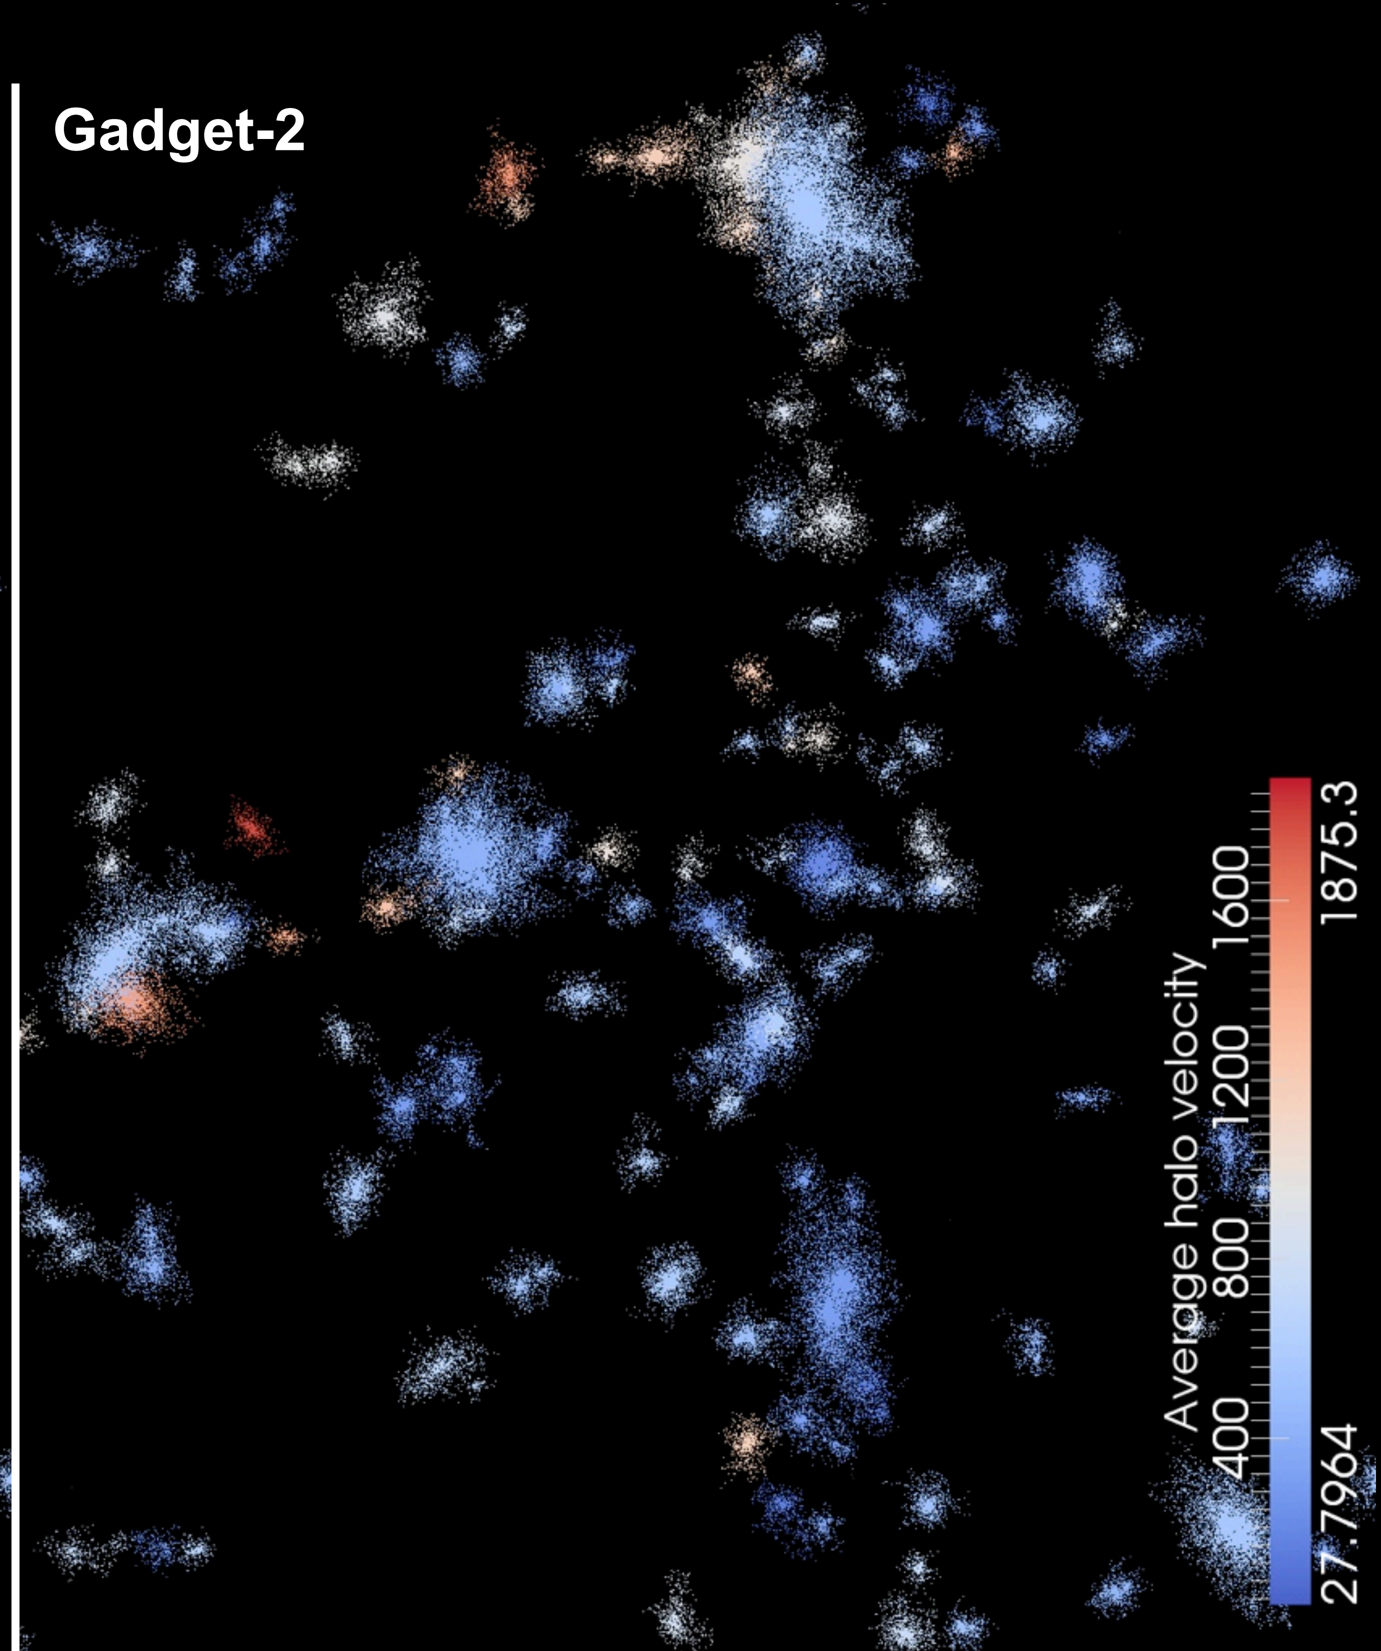

Snapshot from Code Comparison simulation, ~25 Mpc region; halos with > 200 particles, b=0.15
Differences in runs: P³M vs. TPM, force kernels, time stepper: MC³: a; Gadget-2: log(a)
Power spectra agree at sub-percent level

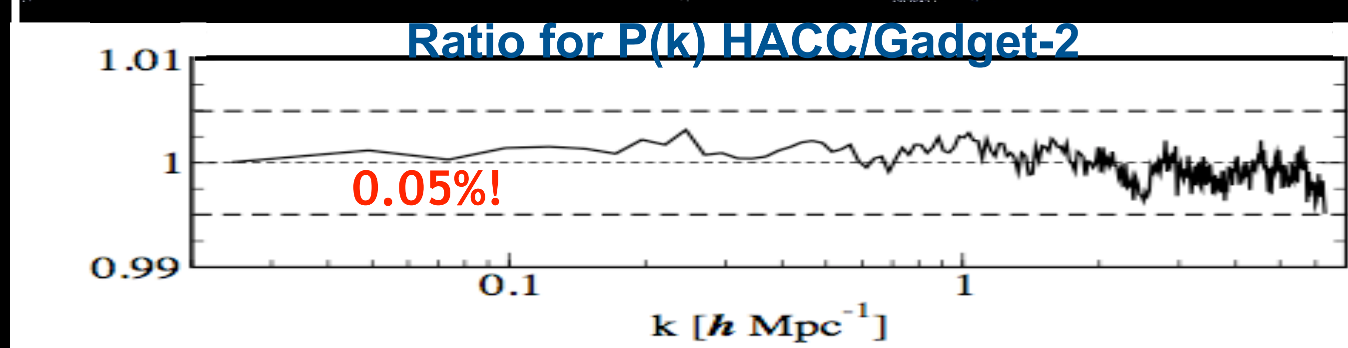

The Story continues, MC³ becomes HACC: CPU+GPU

- **Proof of concept for easy portability:** replace Cell part by GPU implementation
- Paul Sathre (CS undergraduate at the time) in summer 2010 successfully ports code within weeks (with guidance from Adrian Pope)
- **New challenges:**
  - CPU/GPU performance and communication out of balance **AND** unbalanced memory (CPU/main memory dominates)
  - New programming language on GPU, OpenCL
- **With the arrival of Titan in 2013 (GPU accelerated supercomputer at Oak Ridge National Lab):**
  - Nick Frontiere rewrote and optimized P3M short-range solver for GPUs
  - NVIDIA's Justin Luitjens and Argonne's Vitali Morozov wrote tree short-range solver

Weak and Strong Scaling Results

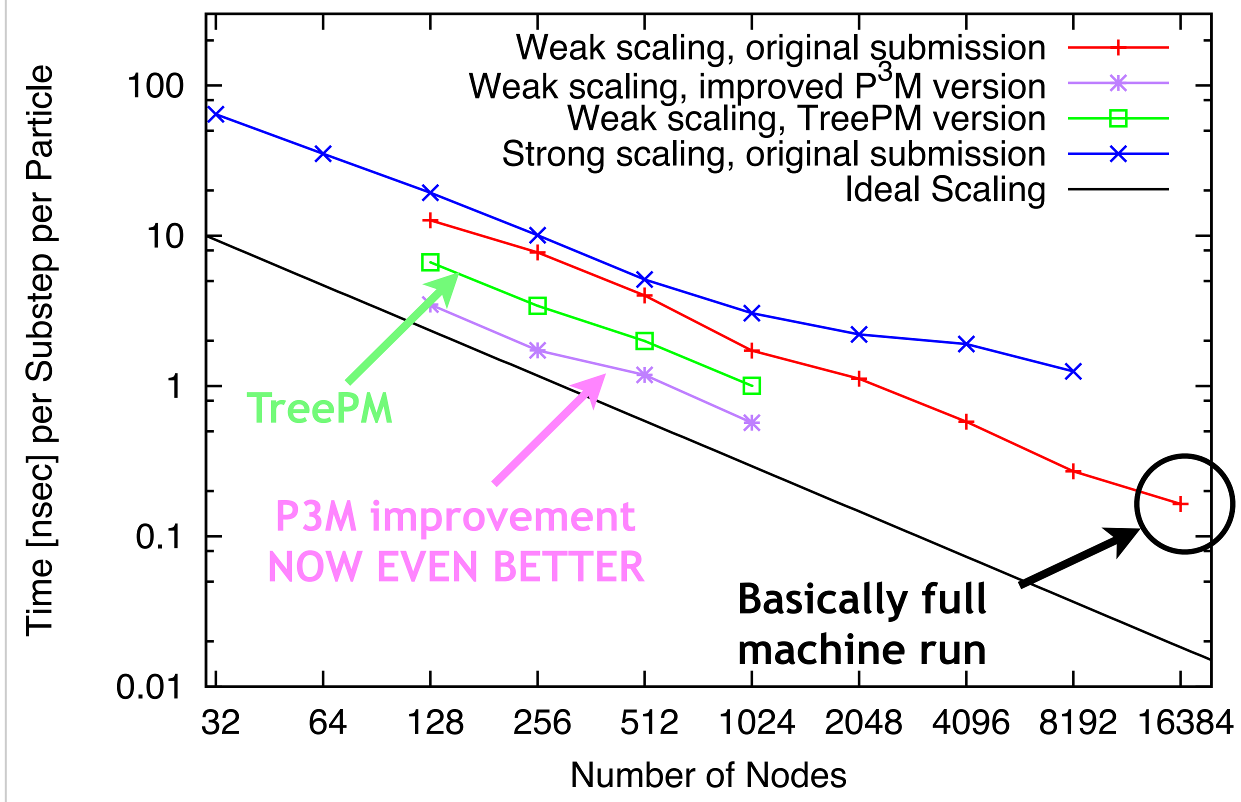

S. Habib et al. 2013: SuperComputing13, Gordon Bell Finalist

- **20.54 Pflops peak performance evolving 1.23 trillion particles in test run on ~75% of machine (full machine not available at the time)**

The Story continues, MC³ becomes HACC: CPU+GPU

- **Proof of concept for easy portability:** replace Cell part by GPU i
- Paul Sathre (CS undergrad time) in summer 2010 successfully ports code within weeks (with guidance from Adrian Pope)

Code is ready but no big machine available!

- **New challenges:**
  - CPU/GPU performance and communication out of balance **AND** unbalanced memory (CPU/main memory dominates)
  - New programming language on GPU, OpenCL
- **With the arrival of Titan in 2013 (GPU accelerated supercomputer at Oak Ridge National Lab):**
  - Nick Frontiere rewrote and optimized P3M short-range solver for GPUs
  - NVIDIA's Justin Luitjens and Argonne's Vitali Morozov wrote tree short-range solver

Weak and Strong Scaling Results

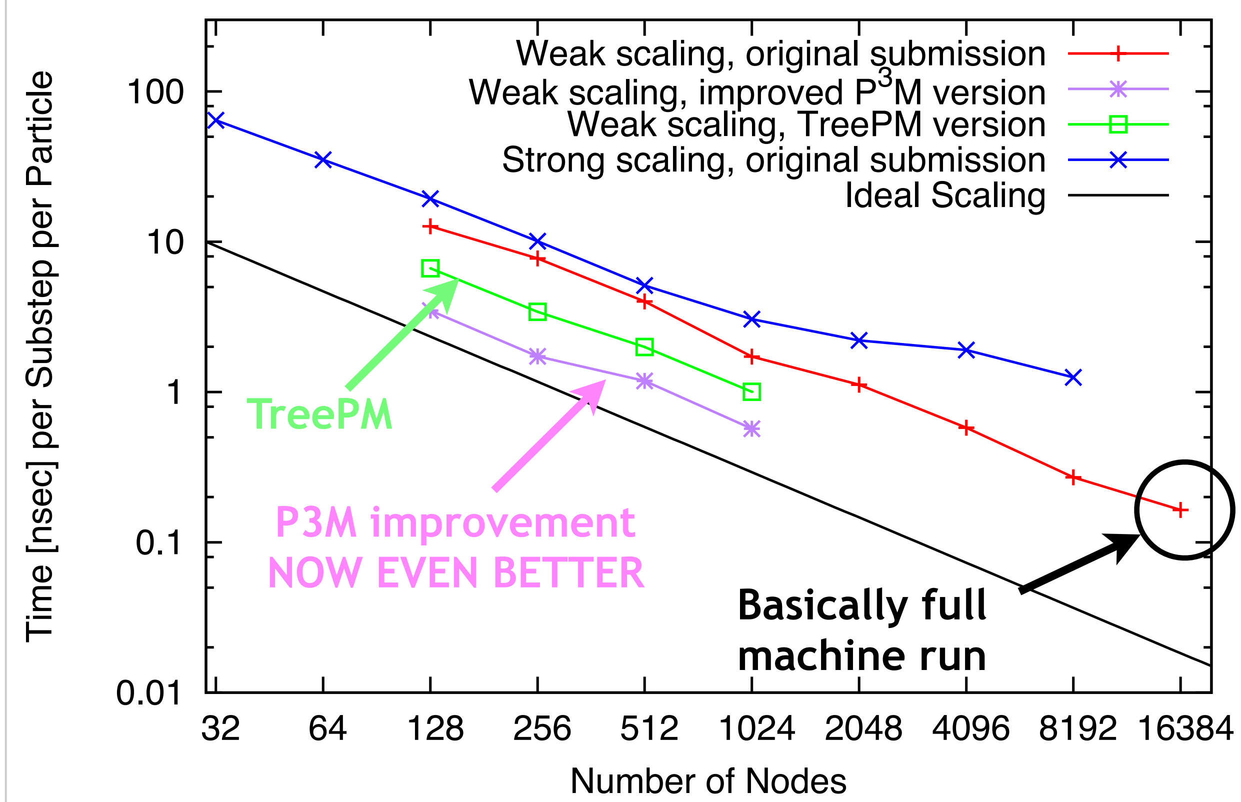

S. Habib et al. 2013: SuperComputing13, Gordon Bell Finalist

- **20.54 Pflops peak performance evolving 1.23 trillion particles in test run on ~75% of machine (full machine not available at the time)**

Some more details about HACC on the GPU

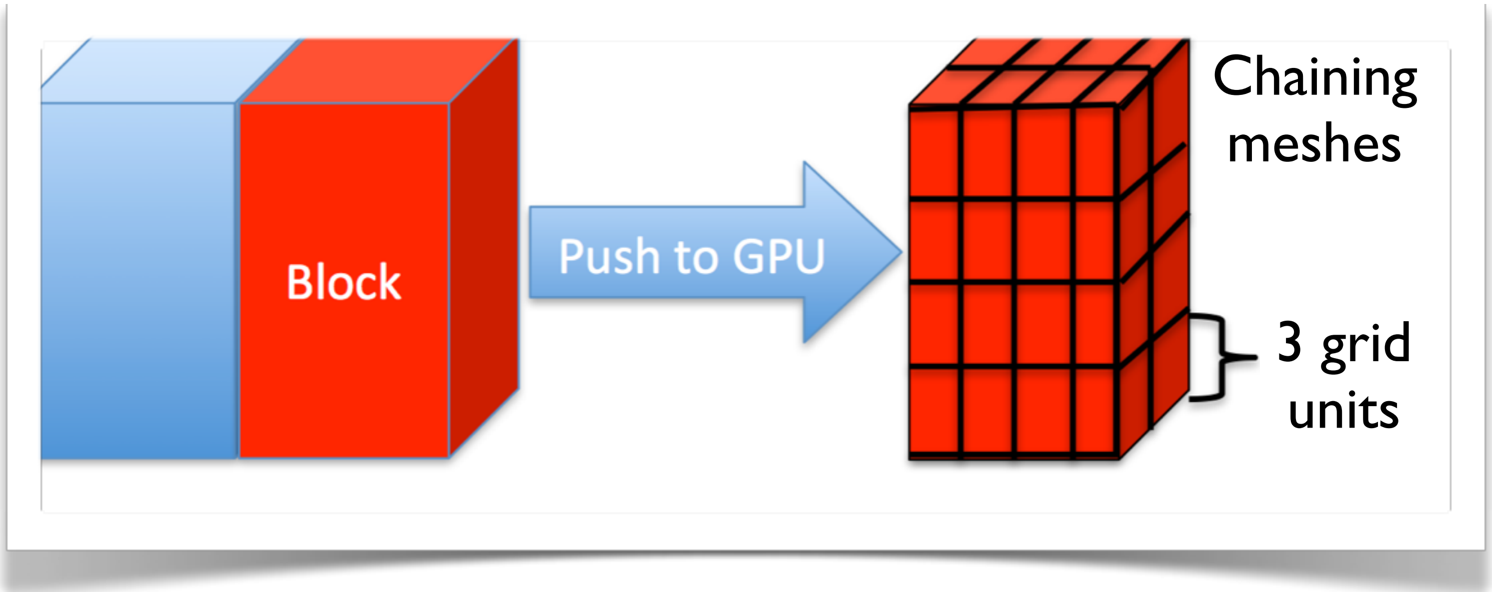

- Memory imbalance: 32GB on host side, 6GB on GPU
- Now need to stream particles through the GPU, have to keep GPU busy!

- **Adaptive load balancing:**
  - At early times: Distribution very smooth, each node takes roughly same time
  - At late time: Large dark matter clumps have formed with up to millions of particles on some nodes, out of balance by up to a factor of 10
  - Use overload concept again on node itself, keep blocks on most nodes as big as possible, reconcile after each step
  - Measure timing for each node, if node takes twice as much as the average node, subdivide into 2 blocks and hand over 1 block to less busy node
  - Overhead due to extra communication small

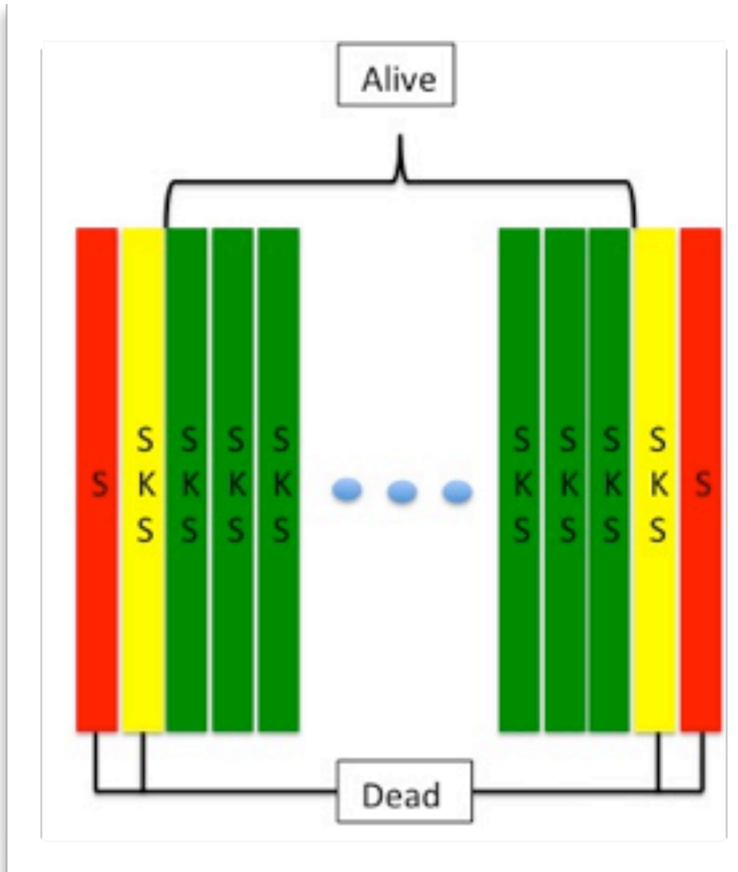

The Q Continuum Simulation, finished July 13 on ~90% of Titan under INCITE, evolving more than half a trillion particles. Shown is the output from one node (~33 million particles), 1/16384 of the full simulation

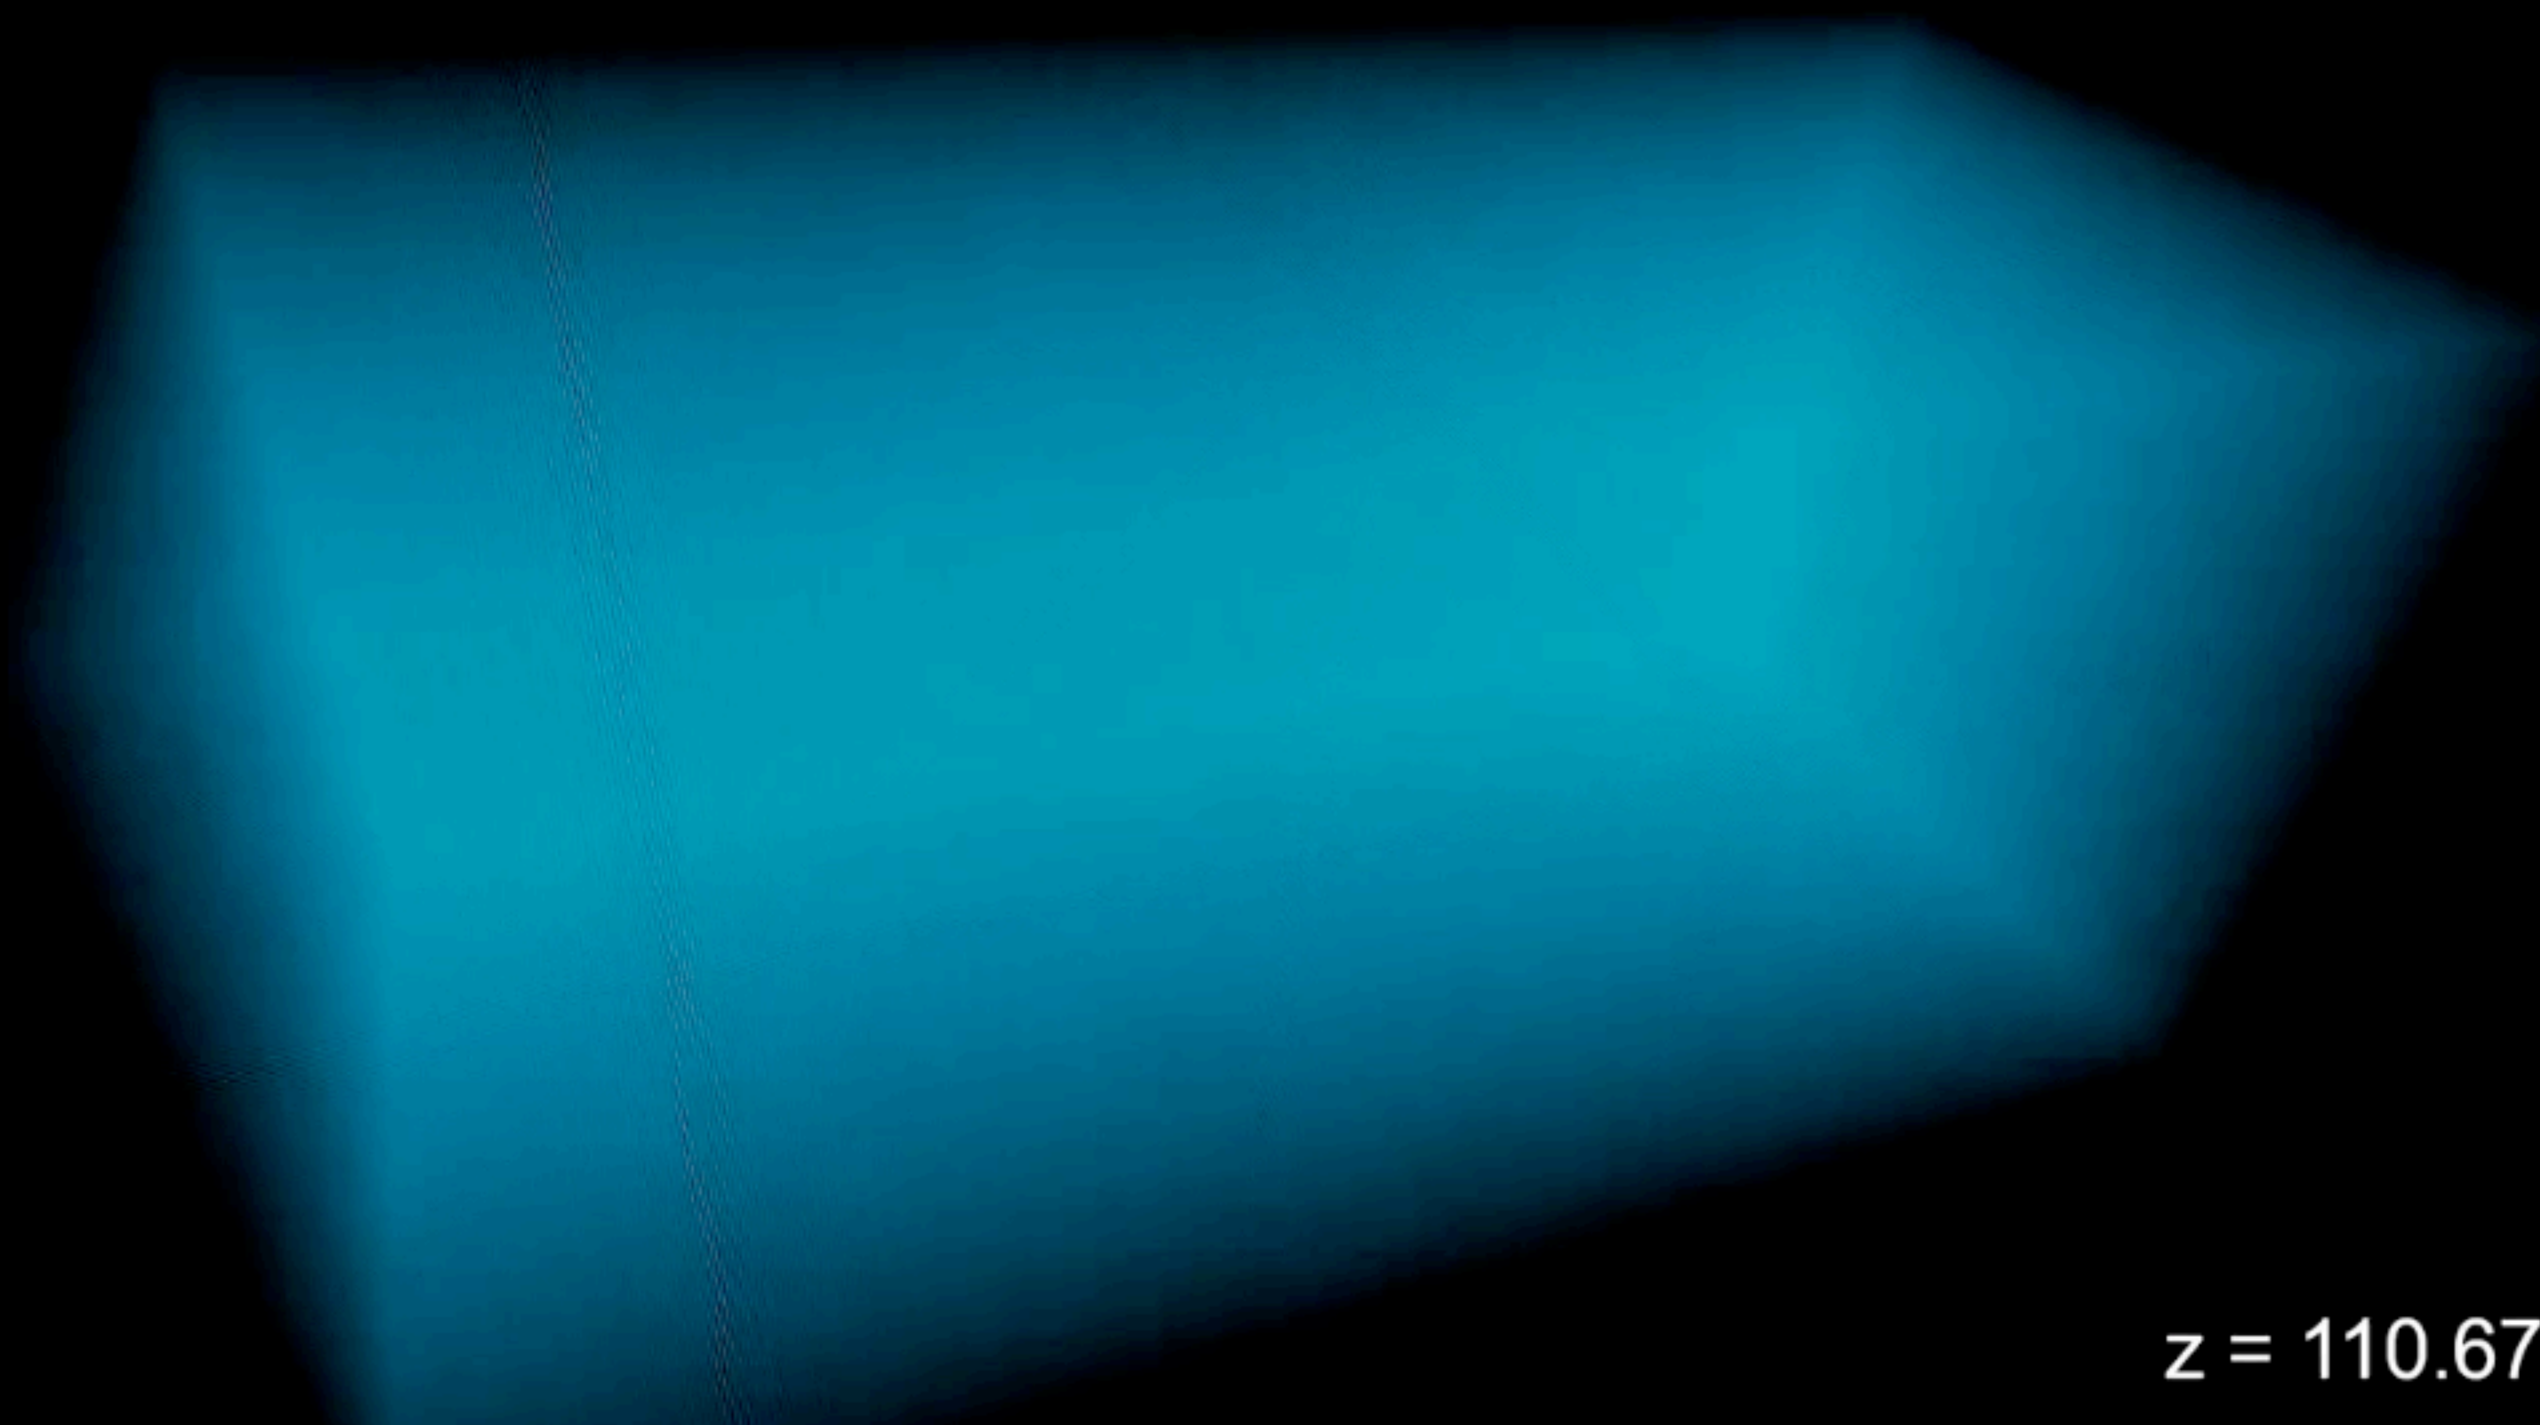

Q Continuum: Extradimensional plane of existence

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

The 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

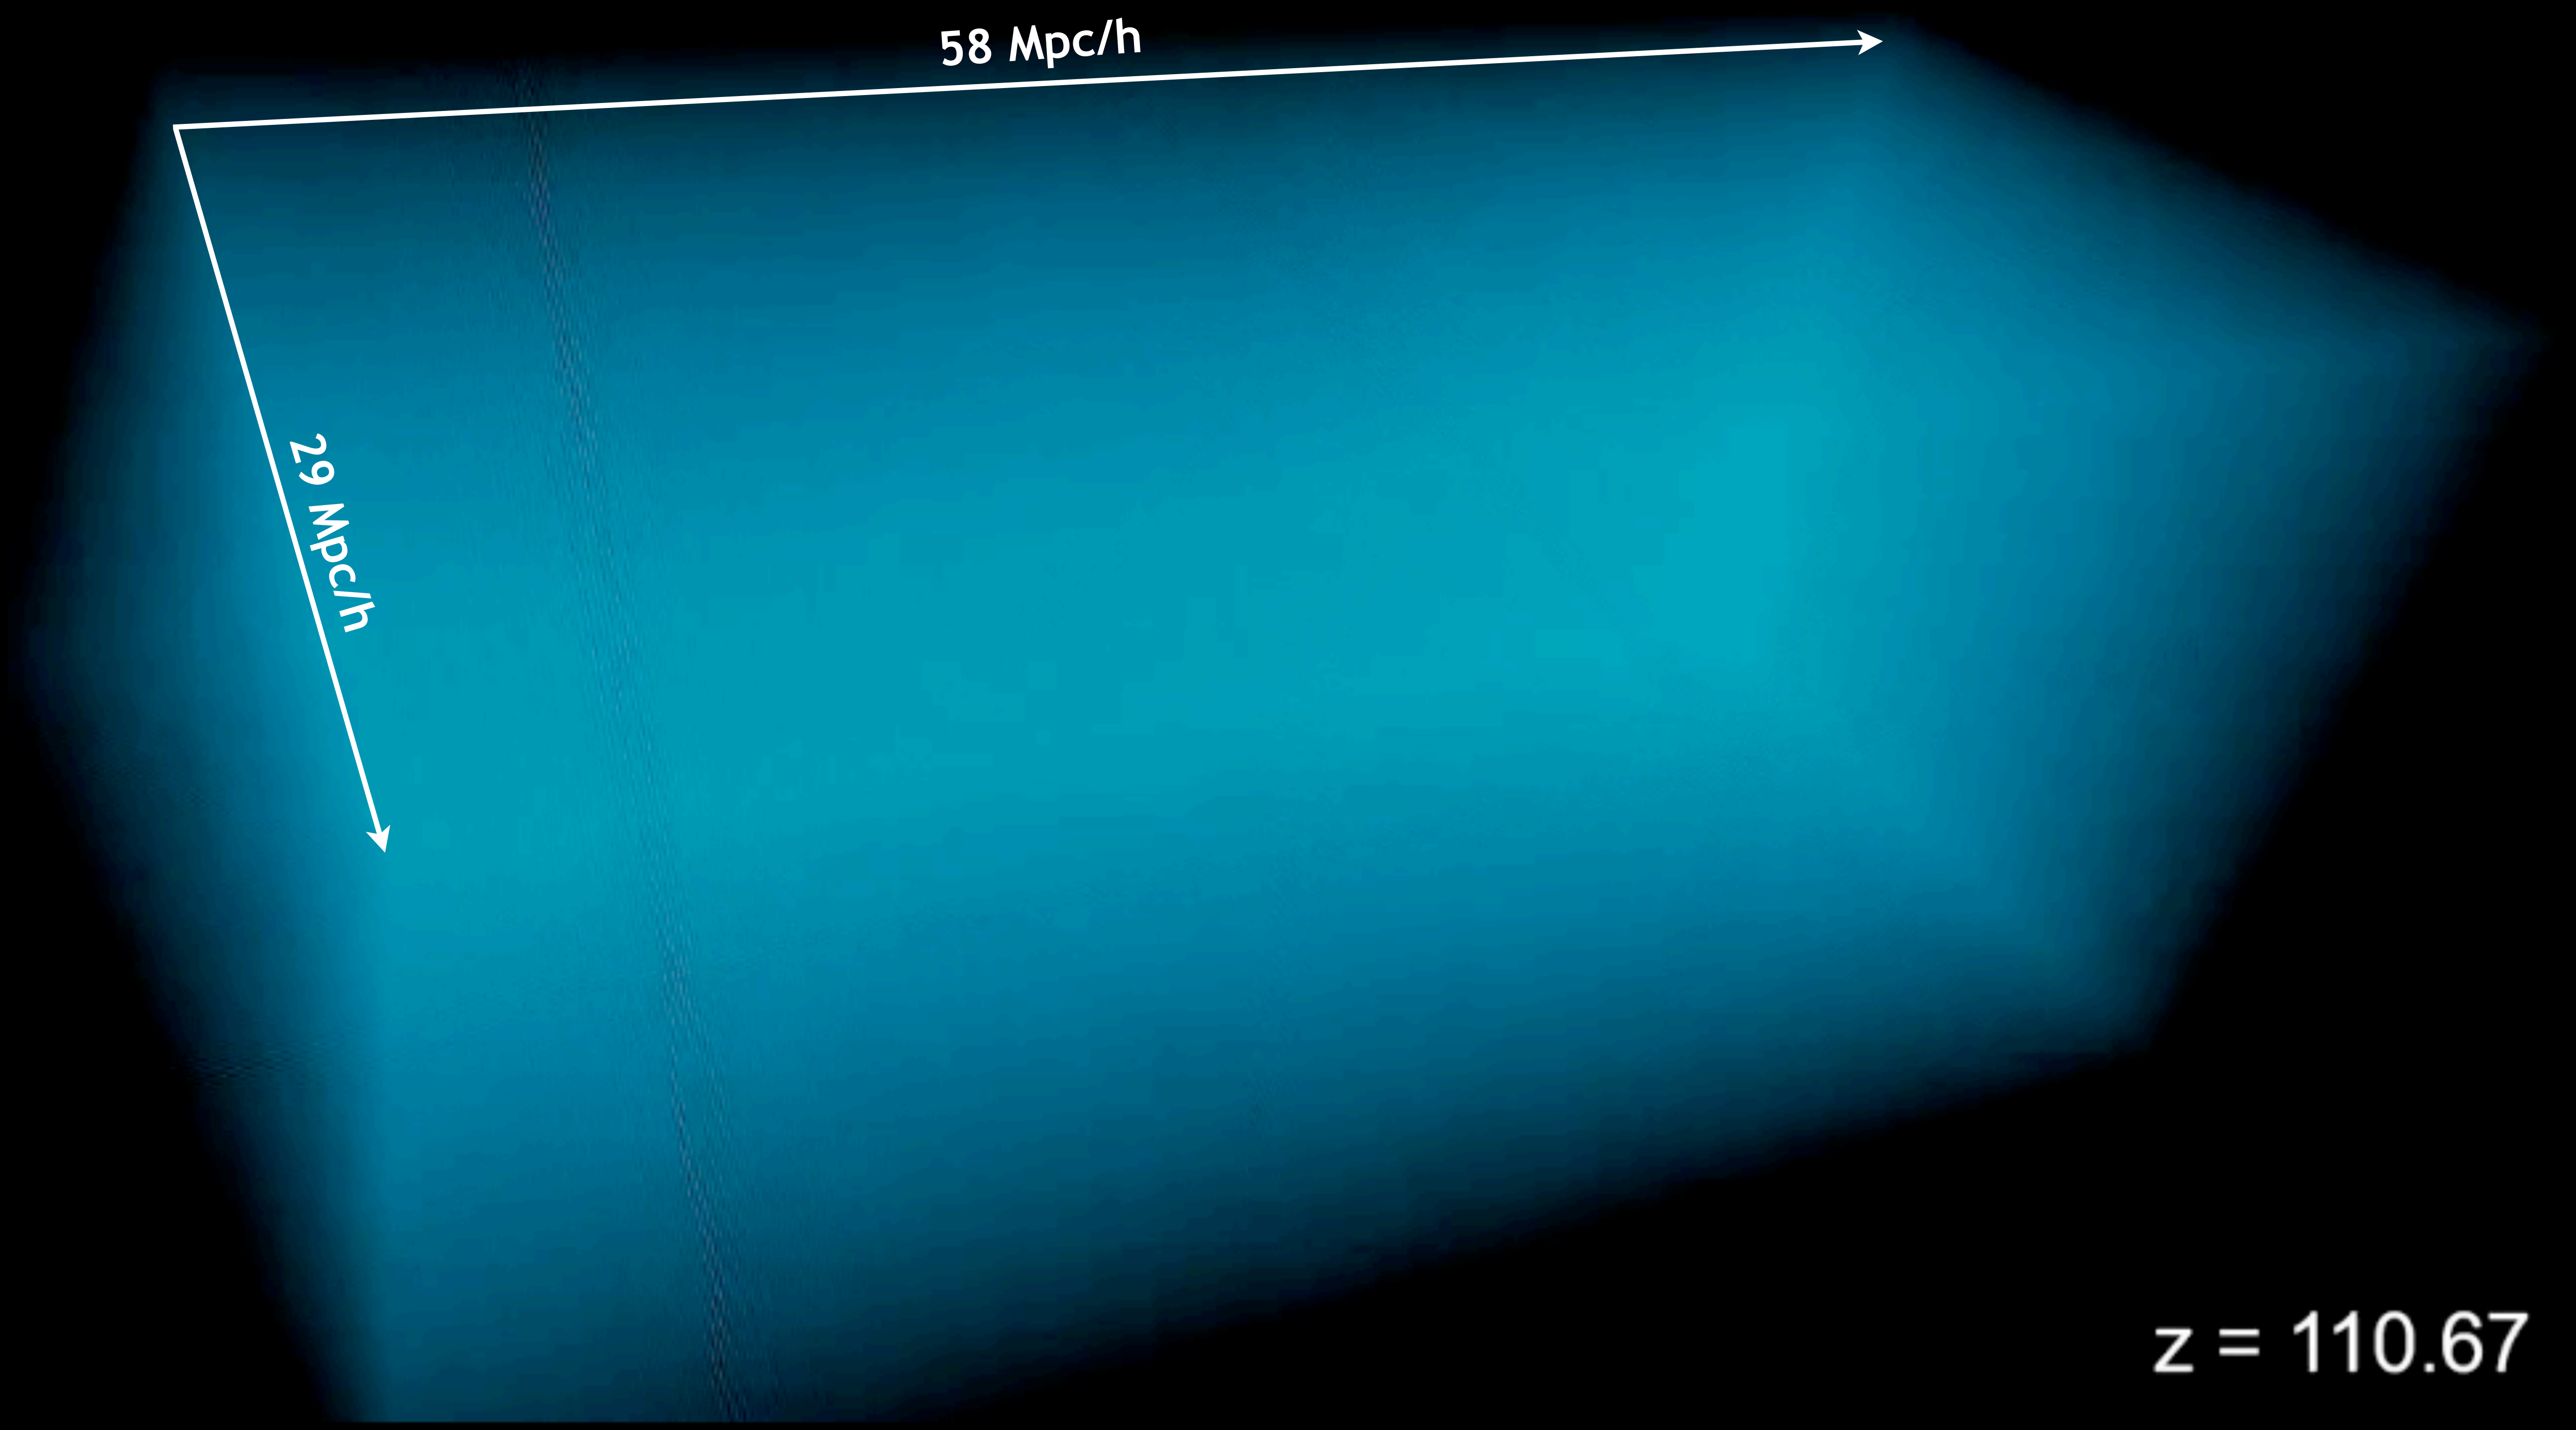

Q Continuum: Extradimensional plane of existence

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

The 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

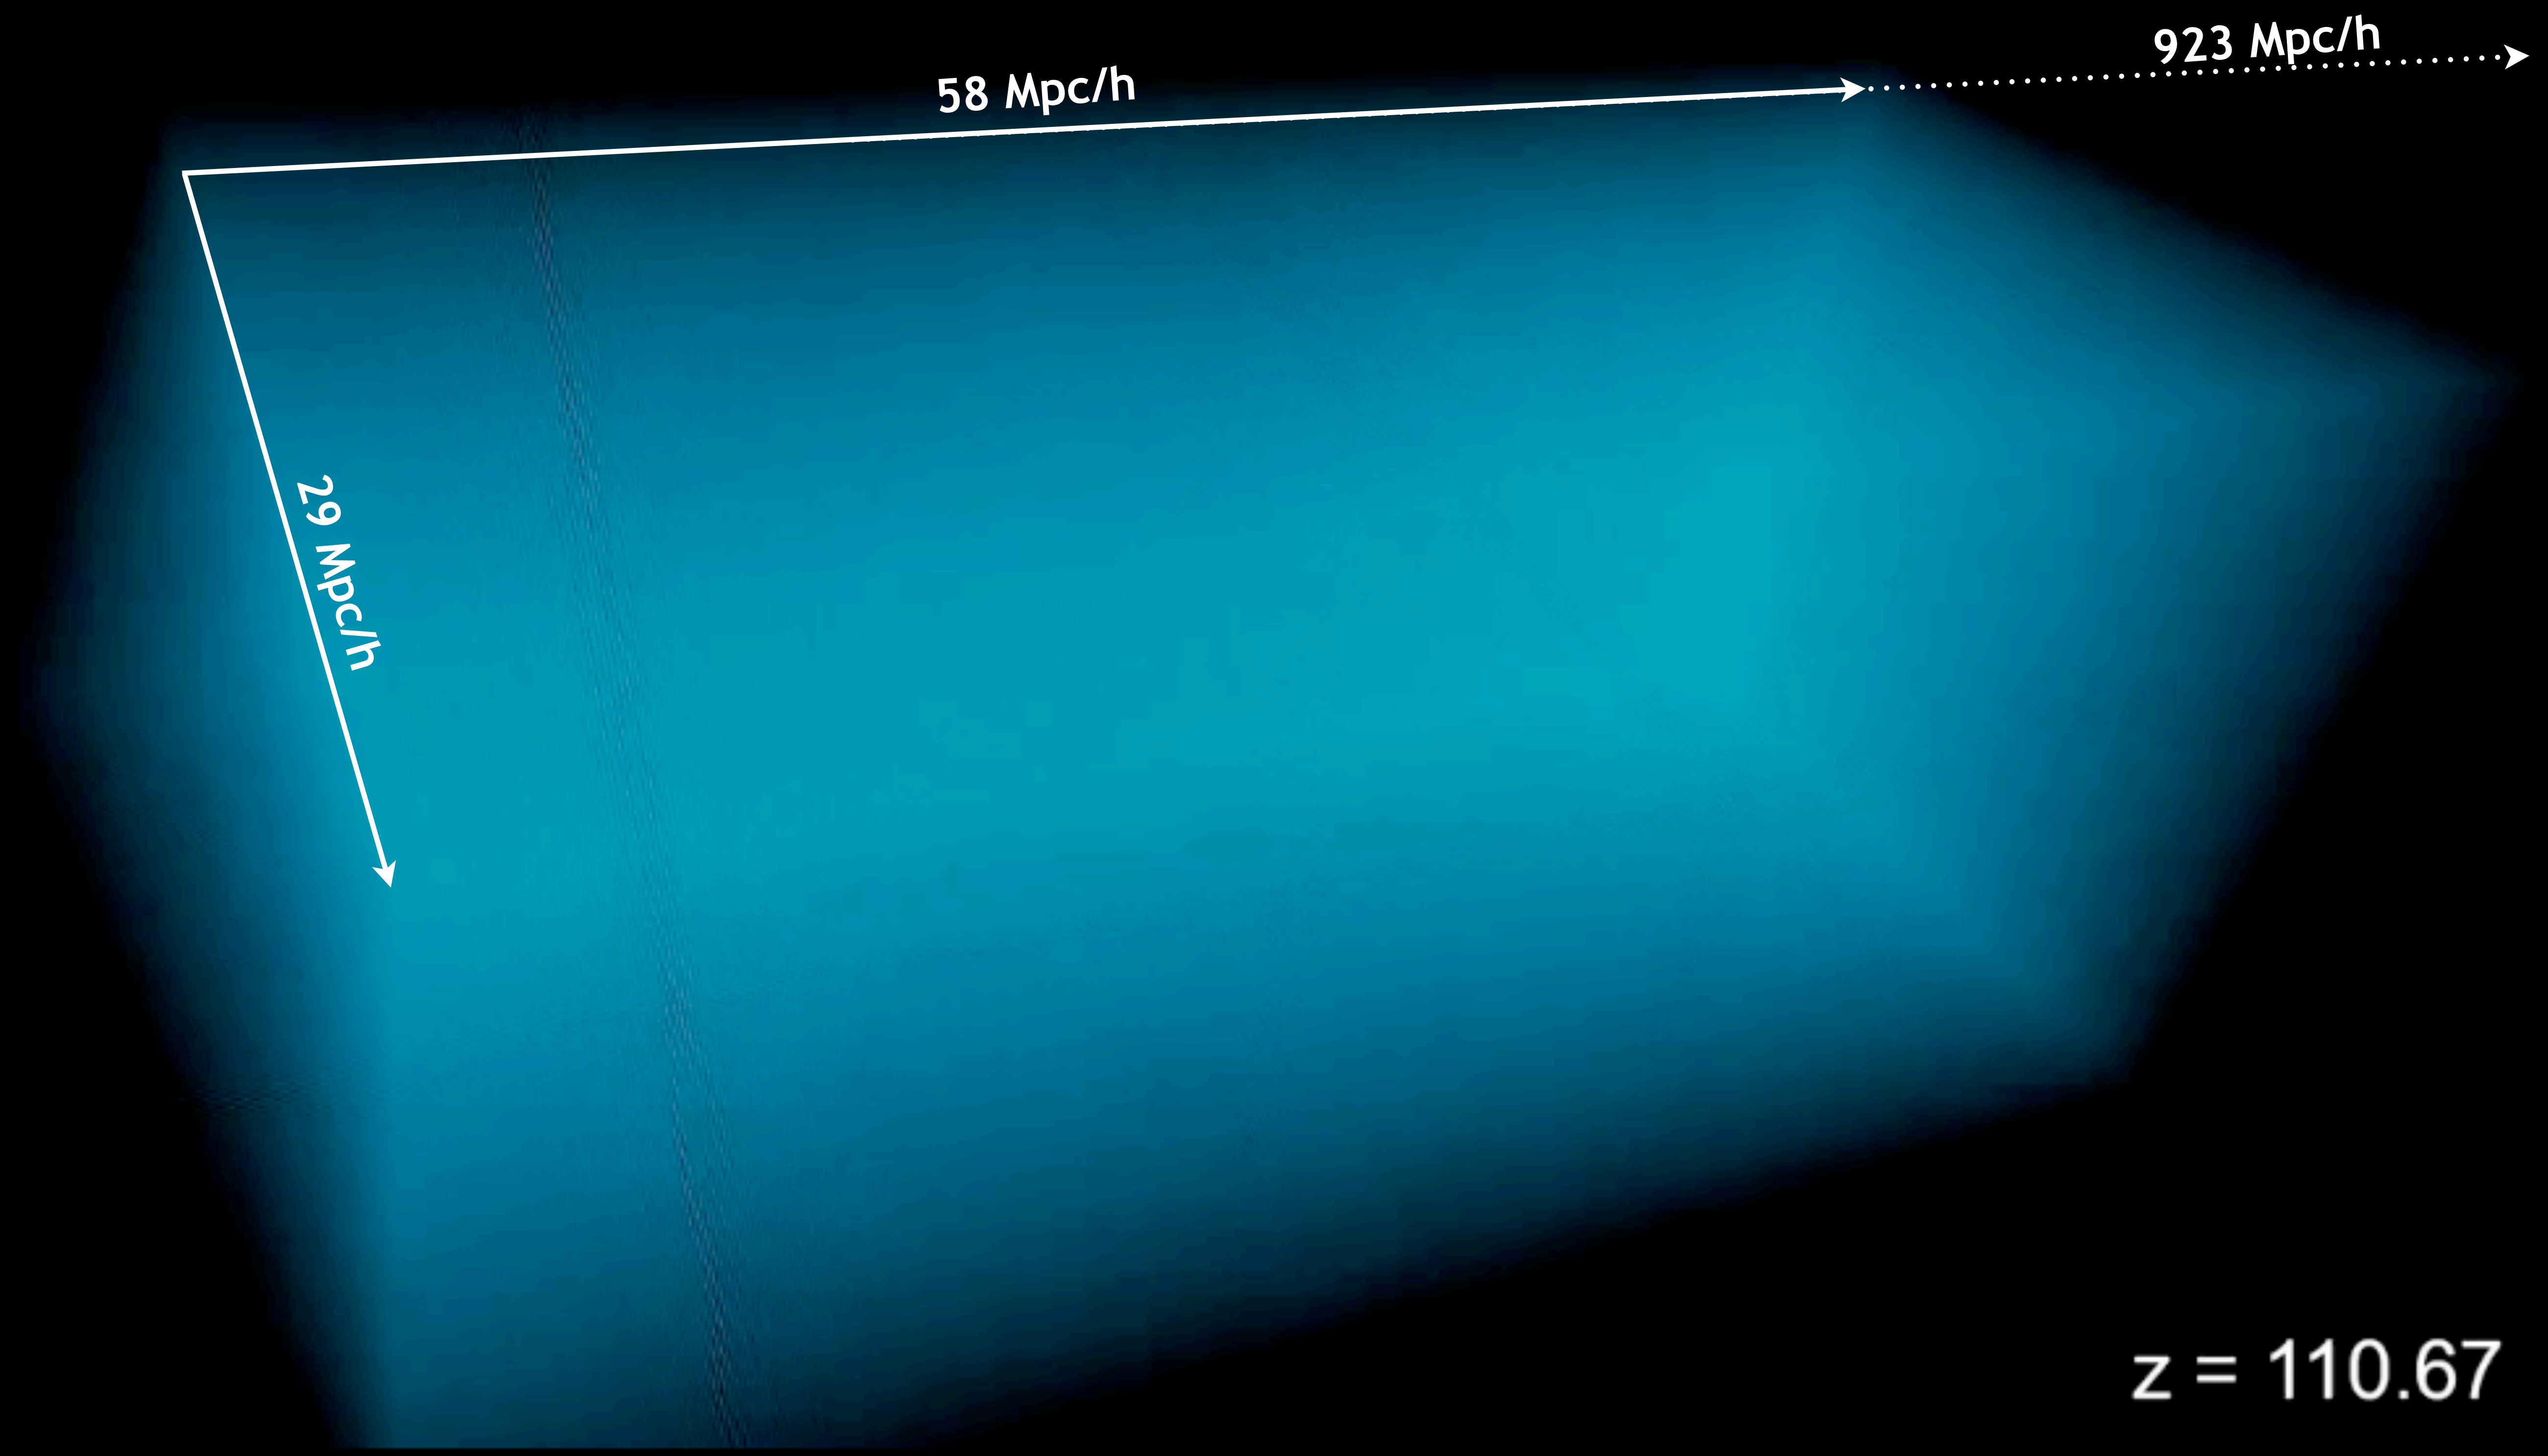

Q Continuum: Extradimensional plane of existence

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

The 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

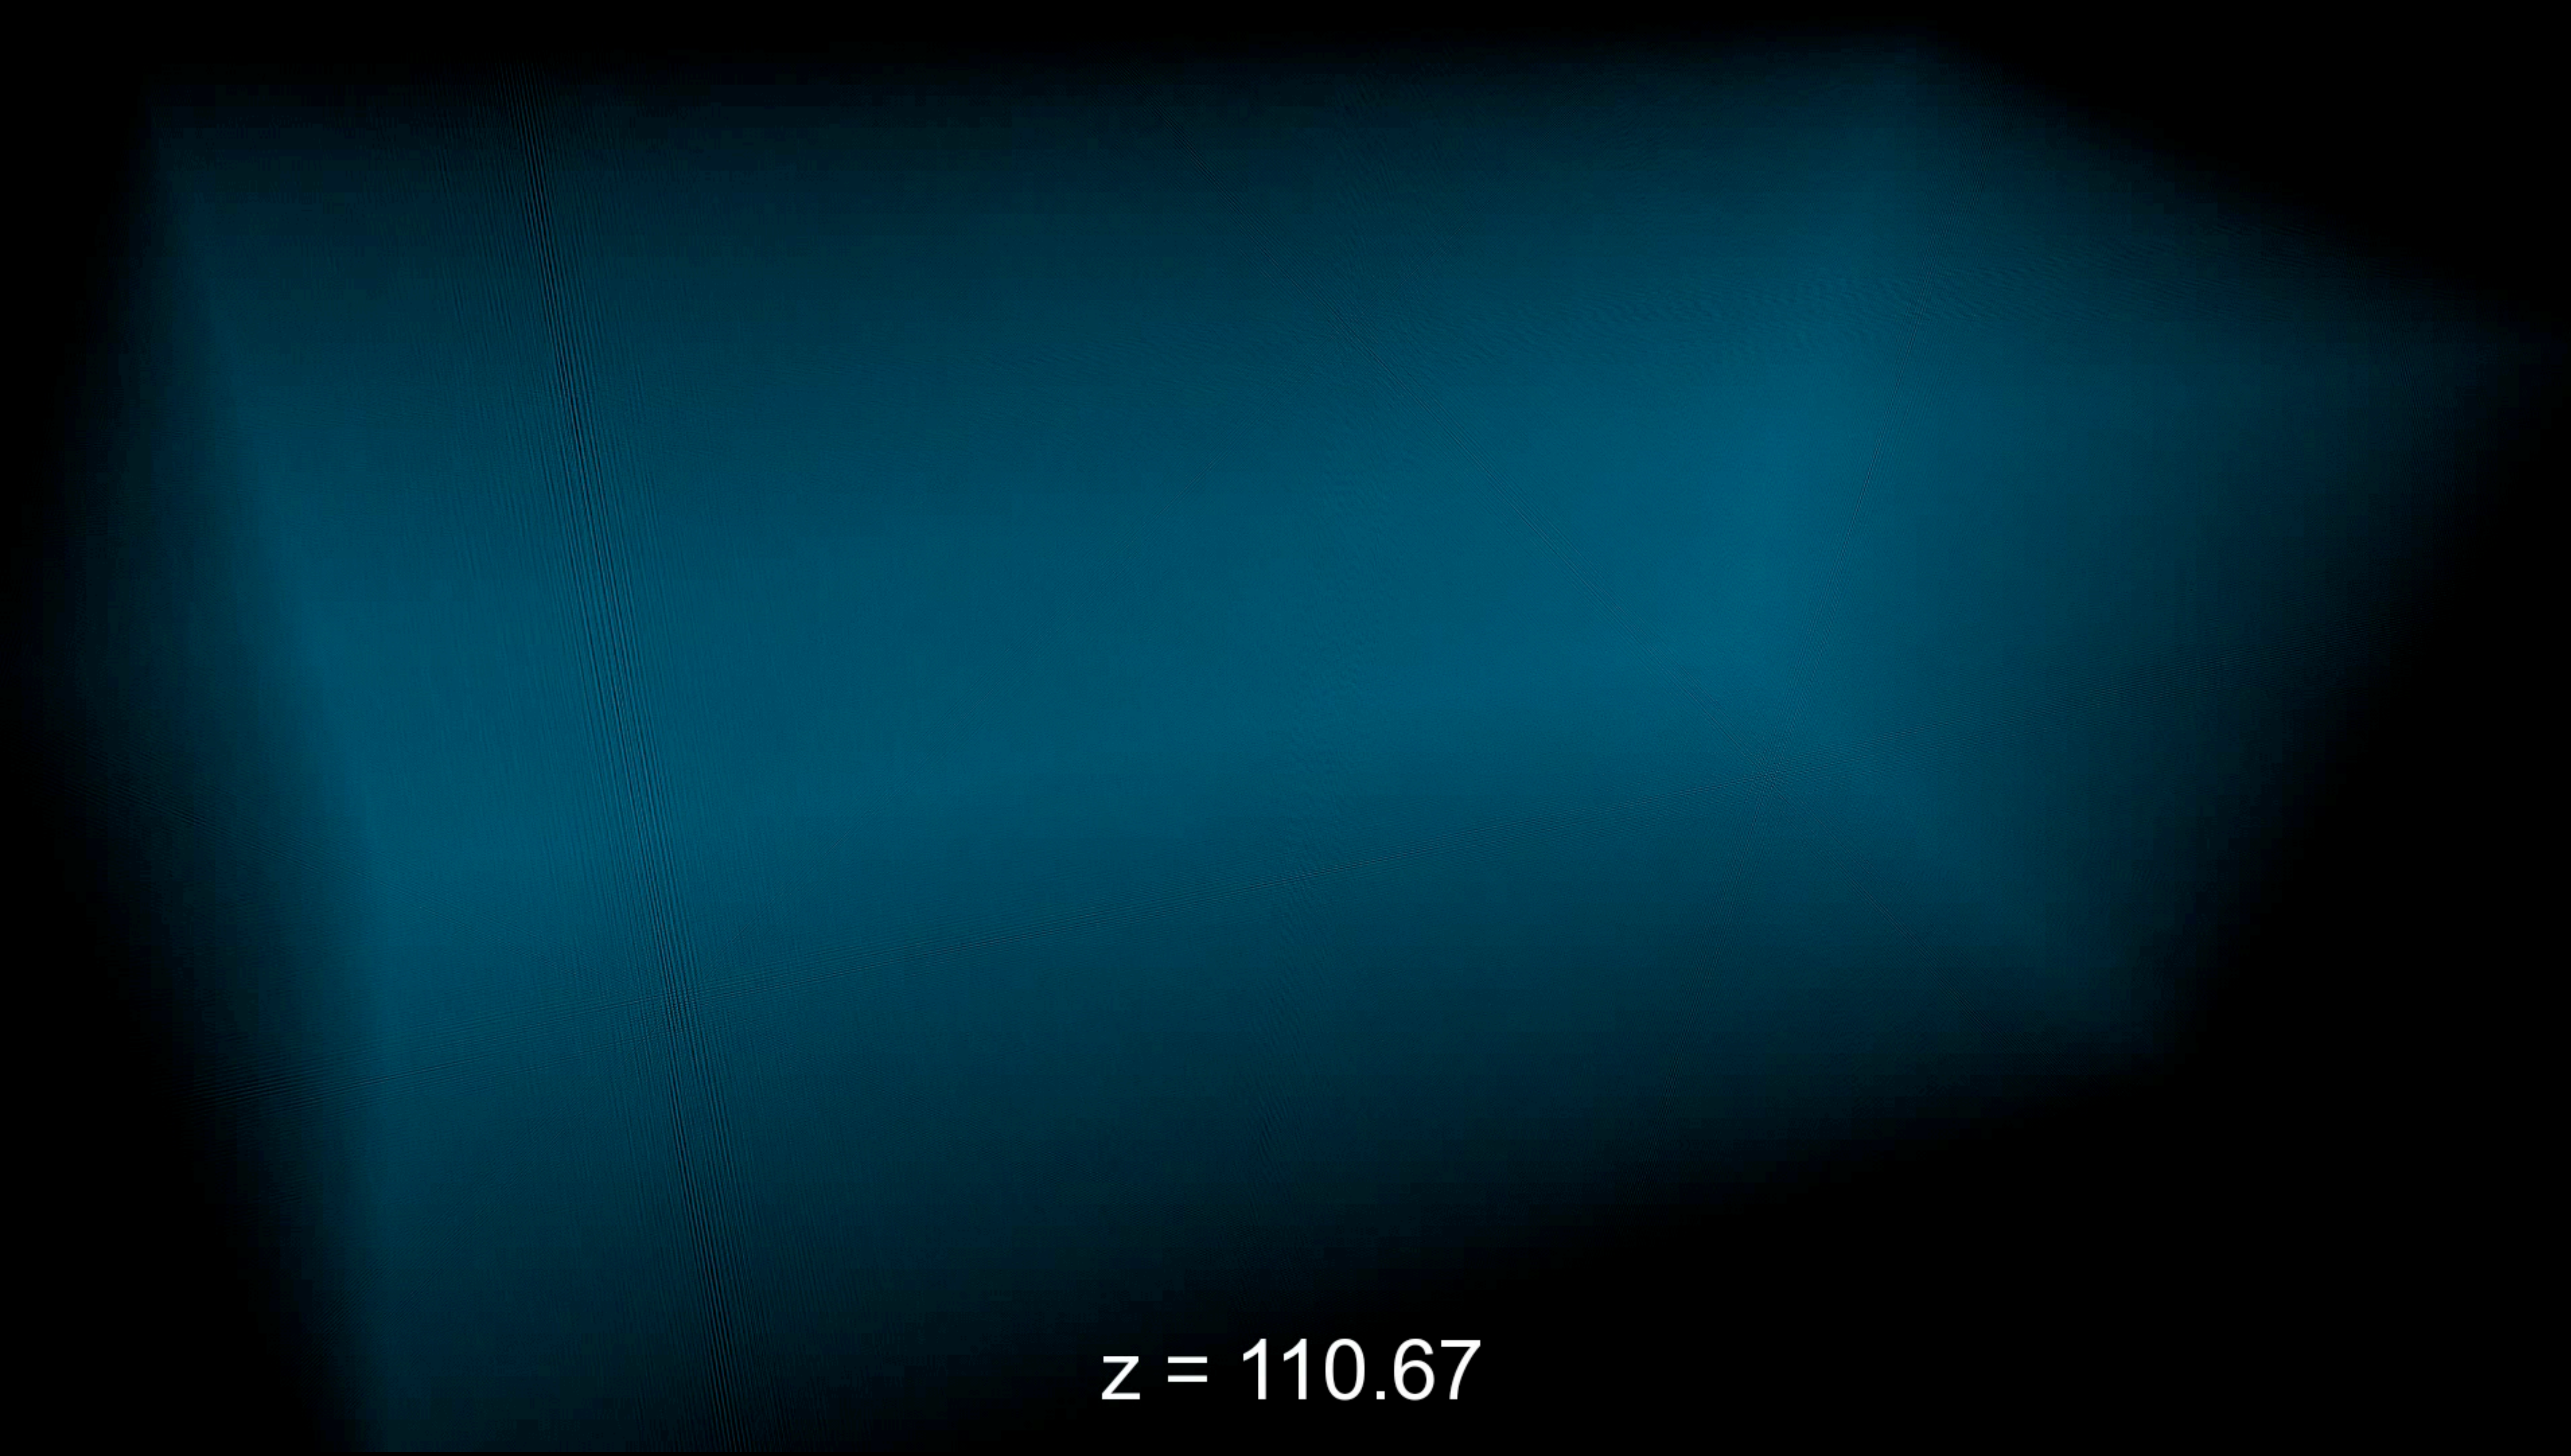

Q Continuum: Extradimensional plane of existence

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

Cosmology with HACC: The Q Continuum Simulation

- Simulation finished on Sunday, July 13! THANKS to OLCF for great support and excellent throughput
- Specifications:
  - (1300 Mpc)^3 volume, 0.55 trillion particles; mass resolution of $m_p \sim 10^8 M_\odot$
  - Ran on 16,384 nodes of Titan (~88% of full machine) for overall ~2 weeks (mostly 24 hour chunks)
  - 100 time snapshots saved for exquisite tracking of structure over time
- For comparison (only simulations relevant for cosmological surveys):
  - Millennium simulation (Springel et al. 2005): (500 Mpc)^3, 10 billion particles, *still used today because of related galaxy catalogs (1973 citations),* $m_p \sim 10^9 M_\odot$
  - Millennium XXL (Angulo et al. 2012): (3 Gpc)^3, 300 billion particles, $m_p \sim 10^{10} M_\odot$
  - Bolshoi simulation (Klypin et al. 2011): (256 Mpc)^3, 8 billion particles, $m_p \sim 10^8 M_\odot$
  - Dark Sky simulation (Skillman et al. 2014): (8 Gpc)^3, 1.1 trillion particles, $m_p \sim 4 \cdot 10^{10} M_\odot$
  - Outer Rim simulation (HACC run on Mira): (3 Gpc)^3, 1.1 trillion particles, $m_p \sim 10^9 M_\odot$
- Q Continuum: Unprecedented mass resolution in a cosmological volume (64x of Bolshoi)!

Cosmology with HACC: The Q Continuum Simulation

FIFA WORLD CUP Brasil 2014

- **Simulation finished on Sunday, July 13! THANKS to OLCF for great support and excellent throughput**
- **Specifications:**
  - (1300 Mpc)^3 volume, 0.55 trillion particles; mass resolution of $m_p \sim 10^8 M_\odot$
  - Ran on 16,384 nodes of Titan (~88% of full machine) for overall ~2 weeks (mostly 24 hour chunks)
  - 100 time snapshots saved for exquisite tracking of structure over time
- **For comparison (only simulations relevant for cosmological surveys):**
  - Millennium simulation (Springel et al. 2005): (500 Mpc)^3, 10 billion particles, *still used today because of related galaxy catalogs (1973 citations),* $m_p \sim 10^9 M_\odot$
  - Millennium XXL (Angulo et al. 2012): (3 Gpc)^3, 300 billion particles, $m_p \sim 10^{10} M_\odot$
  - Bolshoi simulation (Klypin et al. 2011): (256 Mpc)^3, 8 billion particles, $m_p \sim 10^8 M_\odot$
  - Dark Sky simulation (Skillman et al. 2014): (8 Gpc)^3, 1.1 trillion particles, $m_p \sim 4 \cdot 10^{10} M_\odot$
  - Outer Rim simulation (HACC run on Mira): (3 Gpc)^3, 1.1 trillion particles, $m_p \sim 10^9 M_\odot$
- **Q Continuum: Unprecedented mass resolution in a cosmological volume (64x of Bolshoi)!**

Cosmology with HACC: The Q Continuum Simulation

- **Simulation finished on Sunday, July 13! THANKS to OLCF for great support and excellent throughput**
- **Specifications:**
  - (1300 Mpc)^3 volume, 0.55 trillion particles; mass resolution of $m_p \sim 10^8 M_\odot$
  - Ran on 16,384 nodes of Titan (~88% of full machine) for overall ~2 weeks (mostly 24 hour chunks)
  - 100 time snapshots saved for exquisite tracking of structure over time
- **For comparison (only simulations relevant for cosmological surveys):**
  - Millennium simulation (Springel et al. 2005): (500 Mpc)^3, 10 billion particles, *still used today because of related galaxy catalogs (1973 citations),* $m_p \sim 10^9 M_\odot$
  - Millennium XXL (Angulo et al. 2012): (3 Gpc)^3, 300 billion particles, $m_p \sim 10^{10} M_\odot$
  - Bolshoi simulation (Klypin et al. 2011): (256 Mpc)^3, 8 billion particles, $m_p \sim 10^8 M_\odot$
  - Dark Sky simulation (Skillman et al. 2014): (8 Gpc)^3, 1.1 trillion particles, $m_p \sim 4 \cdot 10^{10} M_\odot$
  - Outer Rim simulation (HACC run on Mira): (3 Gpc)^3, 1.1 trillion particles, $m_p \sim 10^9 M_\odot$
- **Q Continuum: Unprecedented mass resolution in a cosmological volume (64x of Bolshoi)!**

Cosmology with the Q Continuum

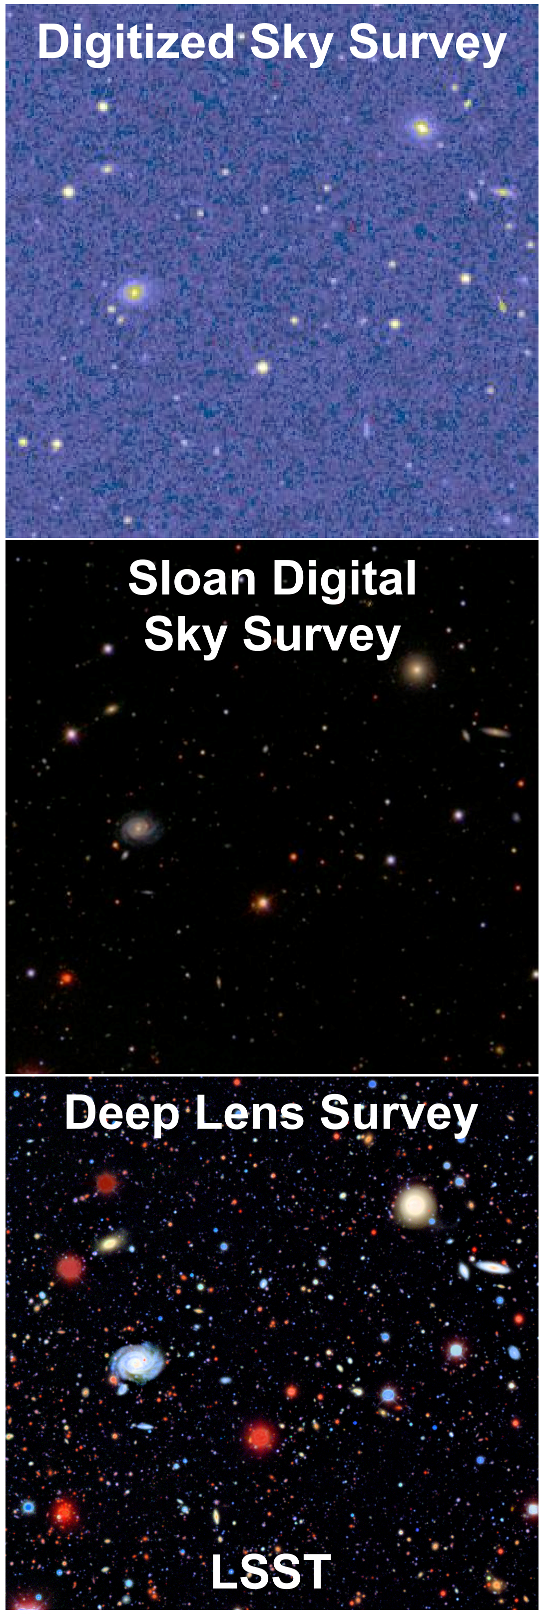

- **Previously for SDSS:** Identify halos and populate them with galaxies, enables us to model the bright galaxies on the sky
- **Now:** Mass resolution enables us to identify ***halos within halos, so-called sub-halos*** (track halo particles over time even after they have become part of another halo)
- **Use semi-analytic code (Galacticus, developed by Andrew Benson) to model galaxy population within subhalos (more details next)**
- **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

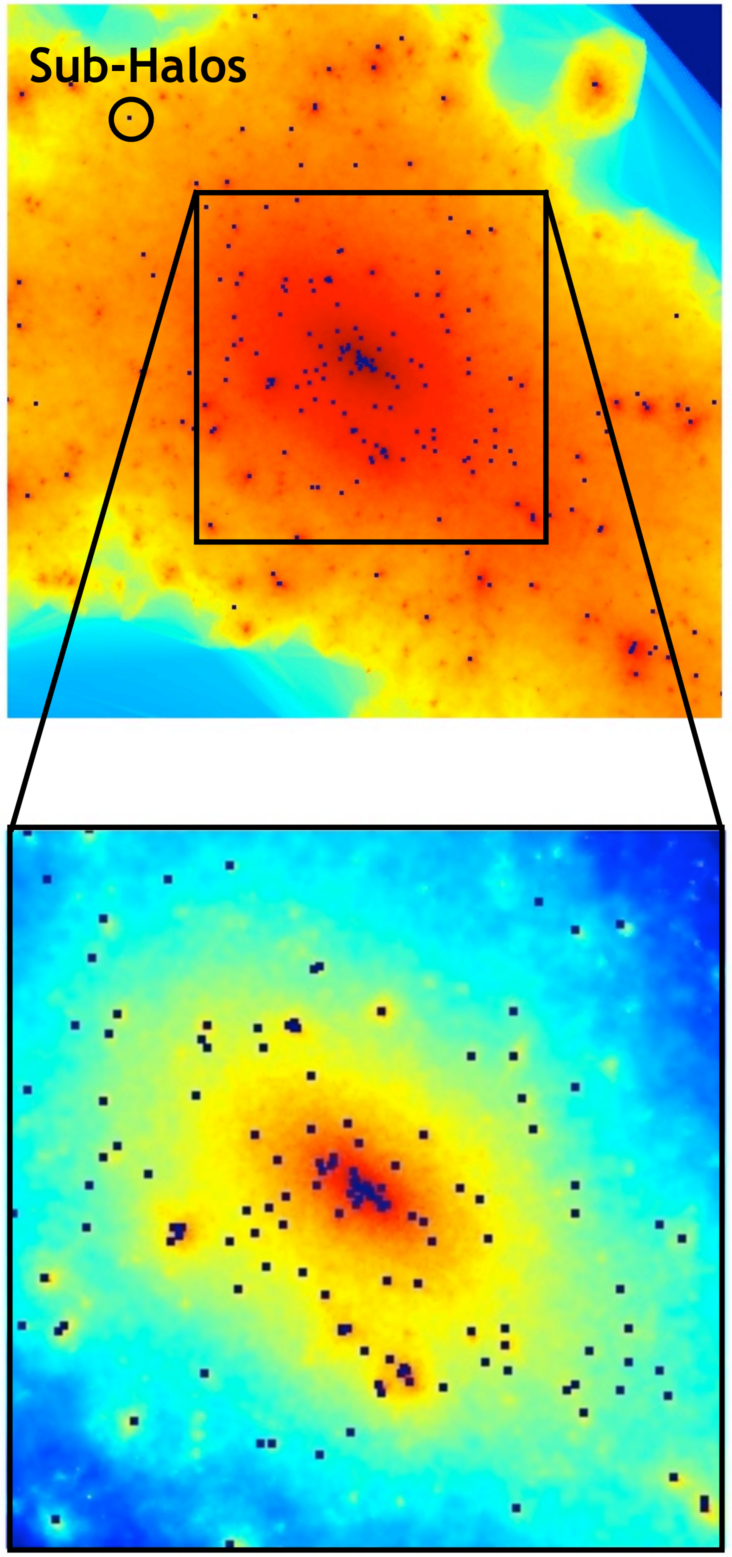

Cosmology with HACC: Synthetic Sky Maps

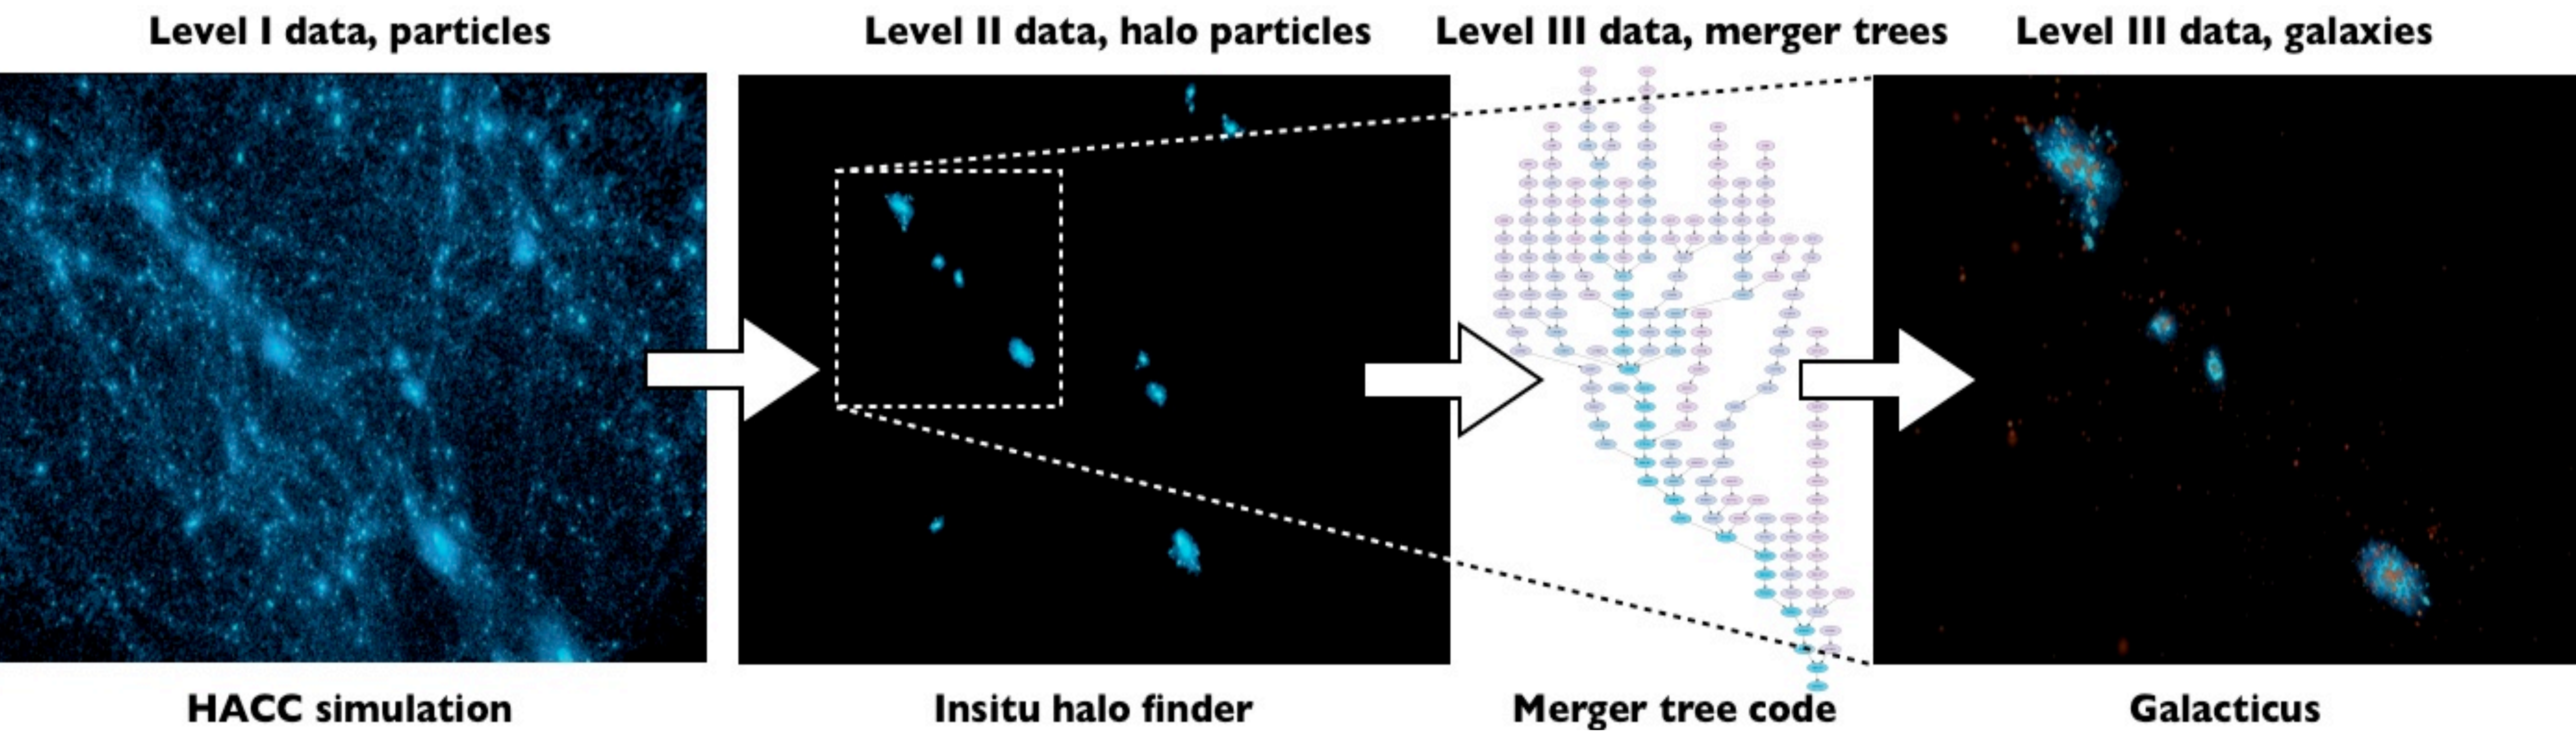

- Downscaled simulation (ran also on Titan) to develop full analysis pipeline
- Extract halos and subhalos from the simulation and track their *complete* formation history over time (develop merger tree for every single halo)
- Run Galacticus (set of ODEs) on top of merger trees to determine galaxies and their properties that live in the halos and subhalos, compare to observations
- Building currently proto-type database for galaxies

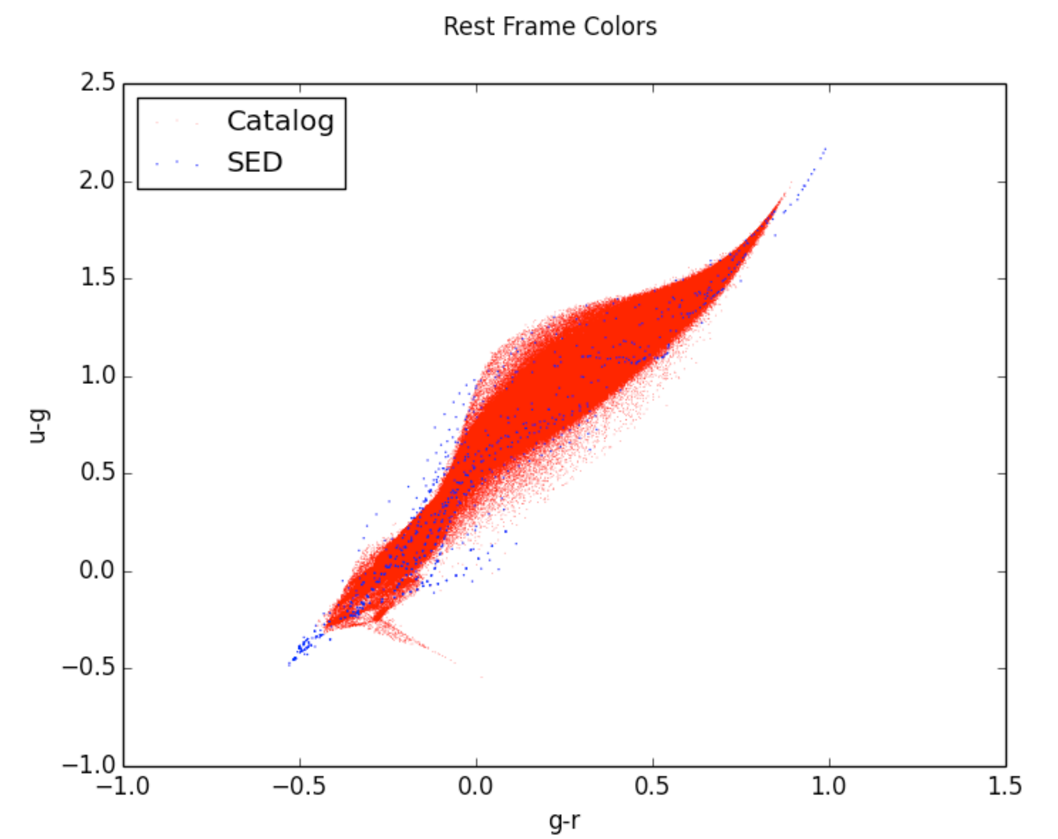

Cosmology with HACC: Exquisite Statistics

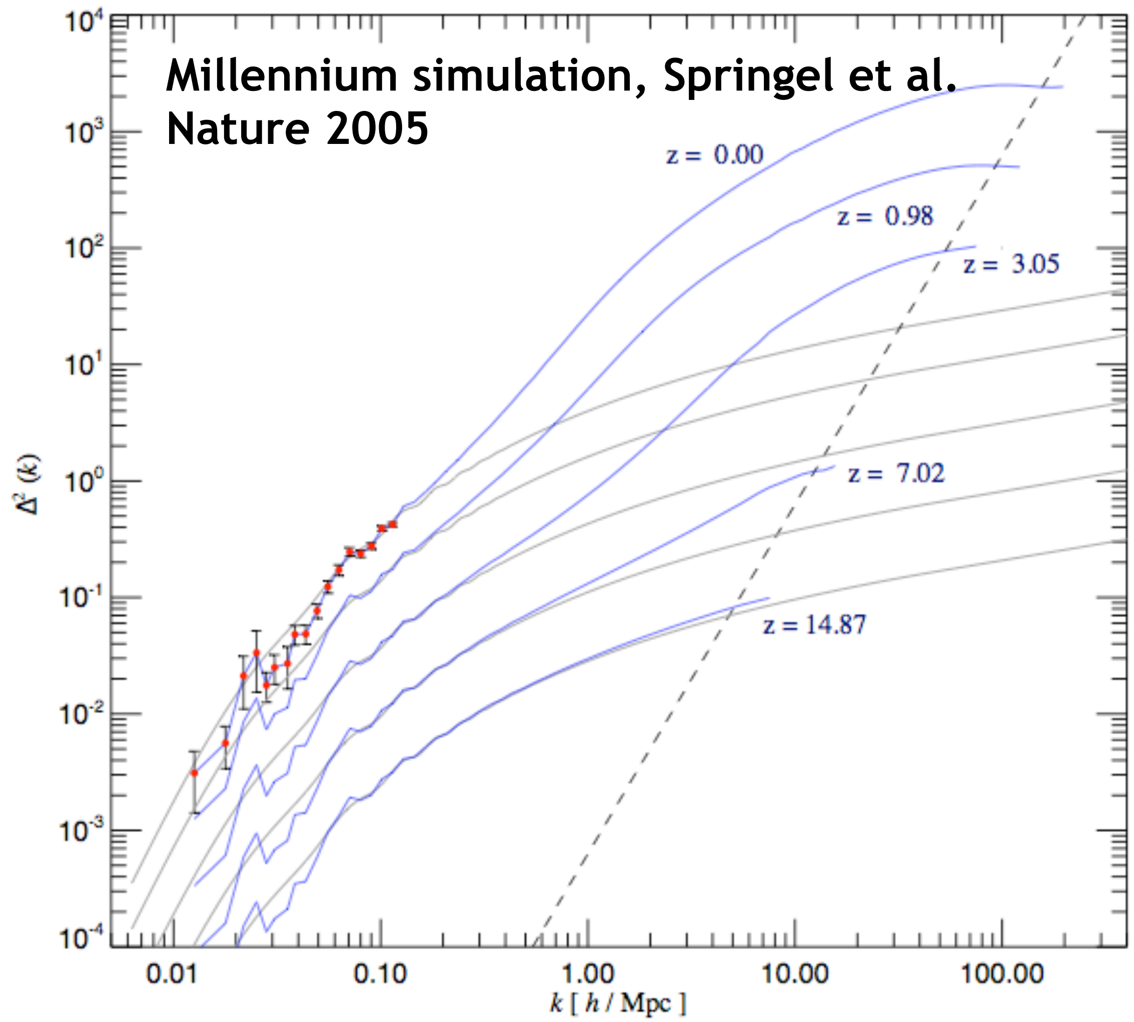

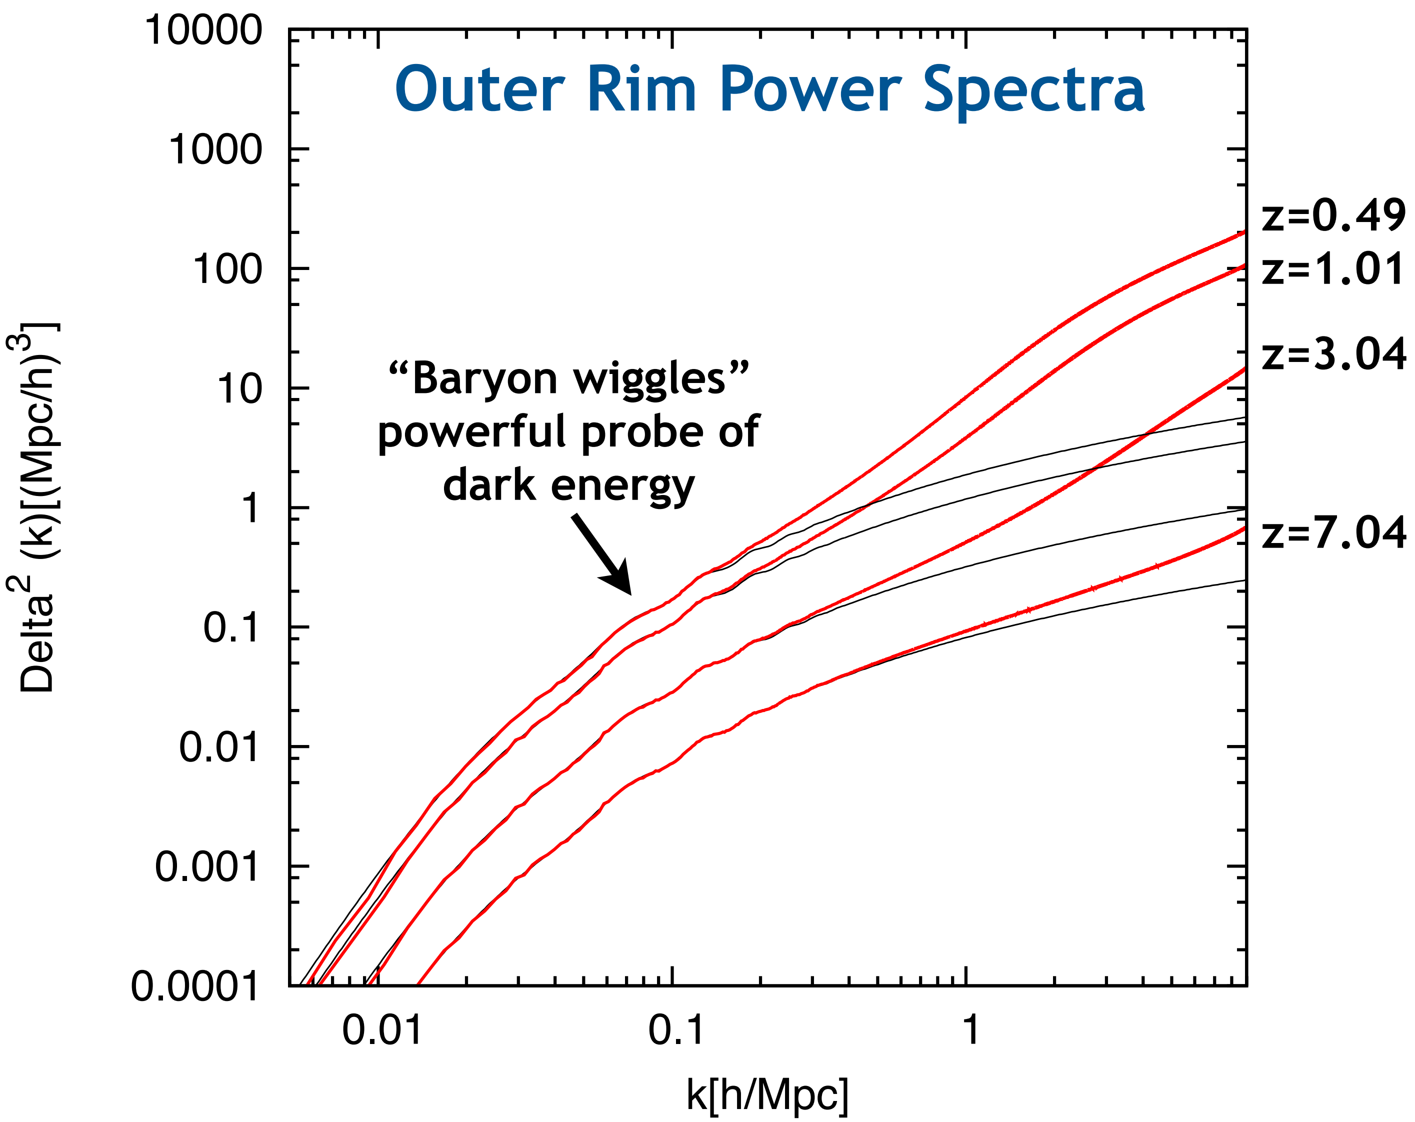

- Mass resolution of Millennium simulation and Outer Rim run very similar (~ $10^9 M_\odot$ particle mass), but volume different by a factor of 216 (Outer Rim volume = Millennium XXL, but 7 times lower mass resolution)
- Exceptional statistics at high resolution enable many science projects

Cosmology with HACC: Strong Lensing

Level II data, halo particles

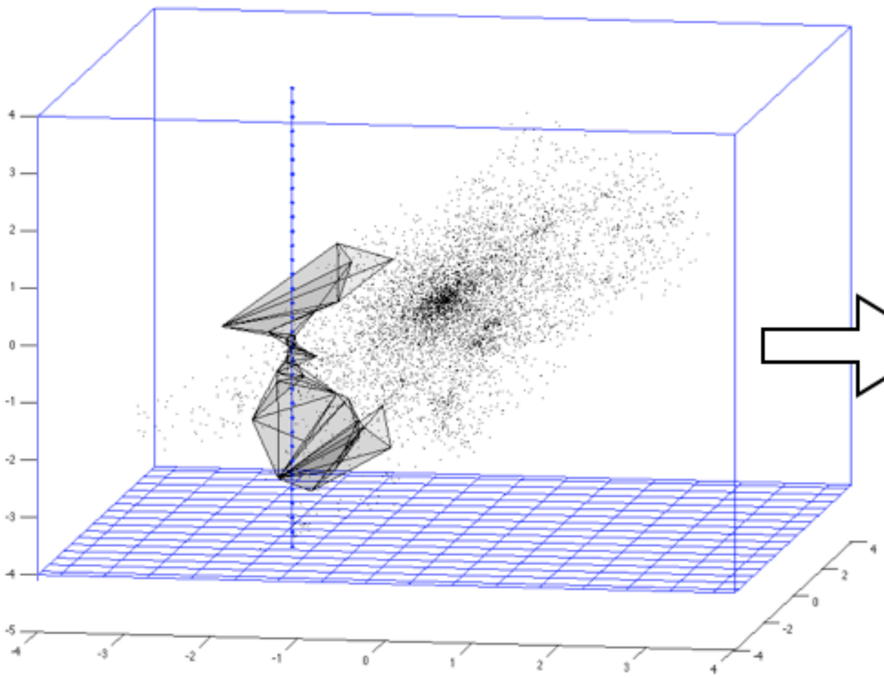

Level II data, projected density

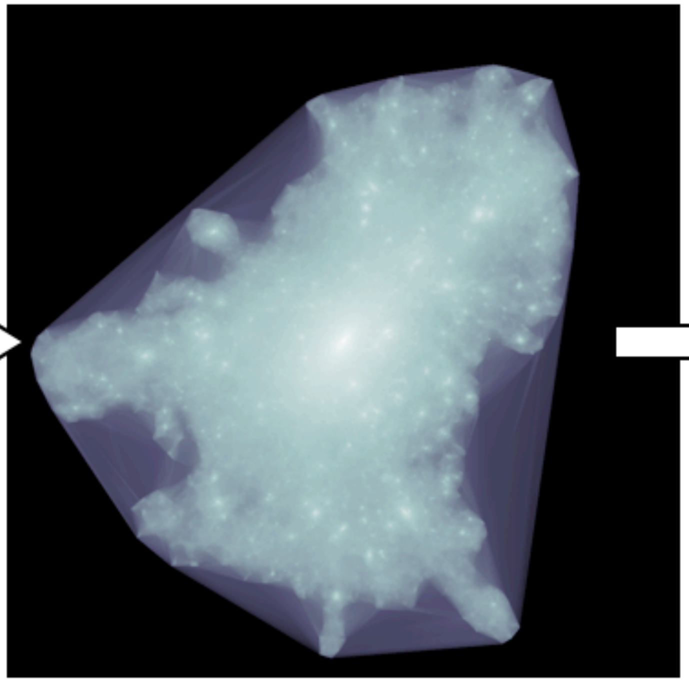

Level III data, simulated image

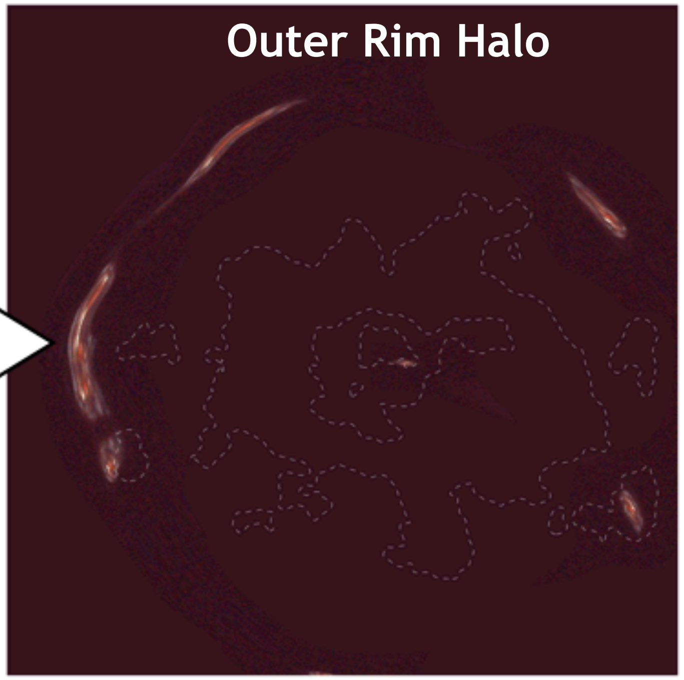

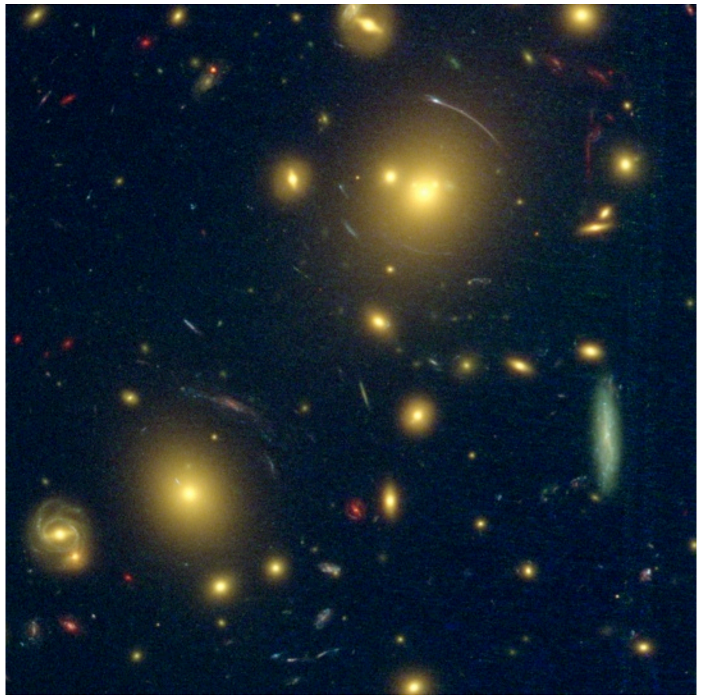

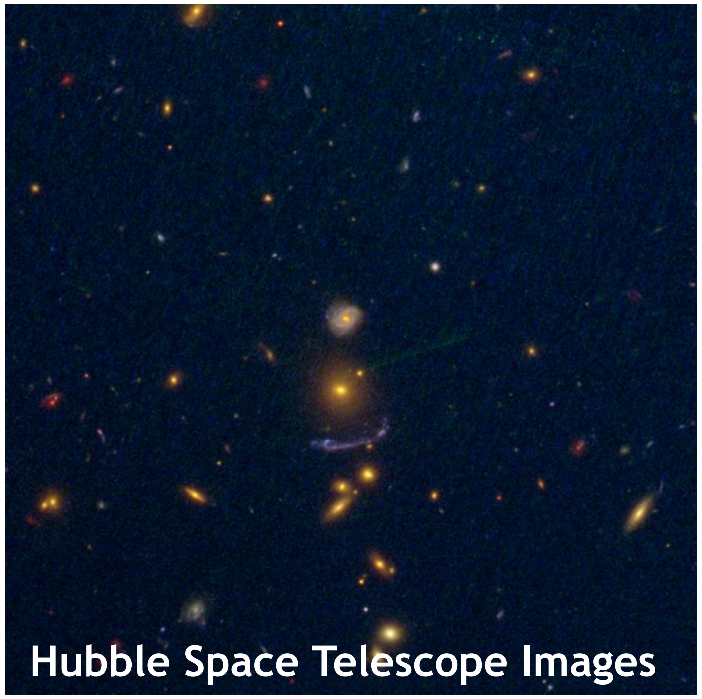

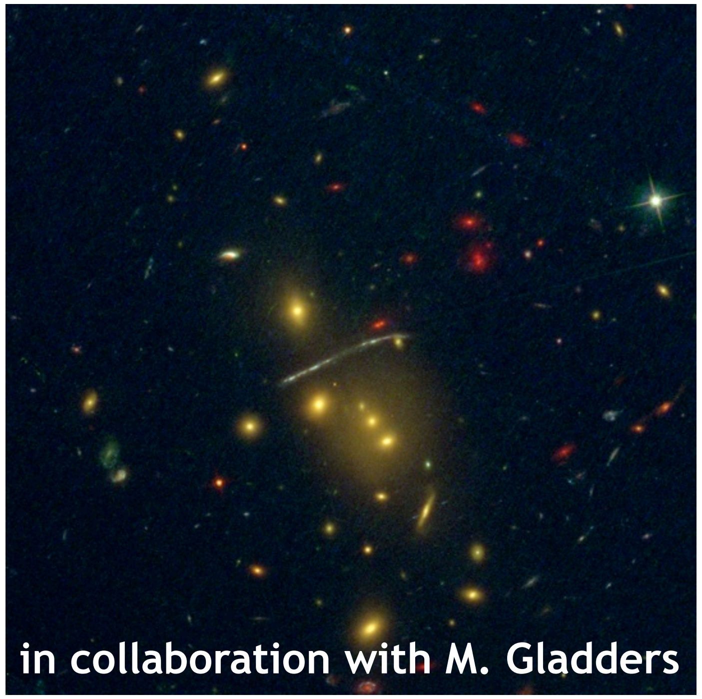

Cosmology with HACC: Future Computational Challenges

Cosmology with HACC: Future Computational Challenges

Cosmology with HACC: Future Computational Challenges

- **CORAL machines: 100-250 PFlops (~10x of Titan) and 4000 TB of memory (~5x of Titan)**
- **On Titan CPU/GPU memory: 32/6~5**
  - If the ratio goes down (more memory on GPU):  remember that we need to be able to "feed" the GPU
  - If the ratio stays the same, we are ok, since the code is working right now
  - If the ratio goes down (too many GPUs per CPU):  it will be hard to get enough work quickly to the GPUs
- **Need good interconnect to get good FFT performance**
- **Prefer "fat" nodes for two reasons:**
  - Reduce overall communications because fewer nodes
  - Reduce overload burden
- **Weak scaling (more particles and more volume = same mass resolution):** factor of 2x speed up (performance/memory ratio, problem size 5 times as big, performance 10x better)
- **Strong scaling:** 10x speed up (for fat nodes, overload improvement)
-

Cosmology with HACC: Future Challenges

- **Ultimate goal: resolve dwarf galaxies with 100 particles each in large cosmological volume**
  - (3 Gpc)^3 volume with $m_p \sim 10^6 M_\odot$ = 1 quadrillion particles = 40 PB per snaphot, we will keep going far beyond exascale!
- **Hydrodynamics, feedback, stars need to be treated correctly within cosmological volumes**
- **More near term future: Analysis challenge**
  - Analysis bottlenecks: data sets become larger and larger, analysis tools have to scale with the code
  - Analysis often done by students, codes cannot be too complicated and tools cannot be optimized for every new architecture (see next talk)
  - Workflows become ever more involved as the simulations are getting closer to reality
- **More near term future: Data challenges**
  - Very few groups can run these large simulations, but large community wants to use the results, how can we make them efficiently available?
  - Cosmology simulations very science rich: We have to keep the raw data! Where?

Summary and Outlook

- **Exciting times for cosmology!**
- **Analysis of current and future cosmological data requires large-scale precision simulations, exploring different models**
- **Current and future supercomputers allow for such simulations but are challenging to use at the same time!**
- **HACC is a large scale computational tool to address both challenges**
    - Designed with very high performance but still ease of portability in mind
    - Runs on all currently available supercomputing architectures
    - In addition to the code itself, extensive work on analysis tools (in-situ as well as post-processing) to create data as close to observations as possible
- **At Argonne: KH lead for LSST DESC and co-lead for DES simulation WG, Salman Habib co-lead for DESI simulation WG -- this ensures that the simulations will be used by the major dark energy surveys in the best possible way!**

EXTRA SLIDES

Cosmology with HACC: Mapping the Matter Distribution

Lyman-alpha Forest

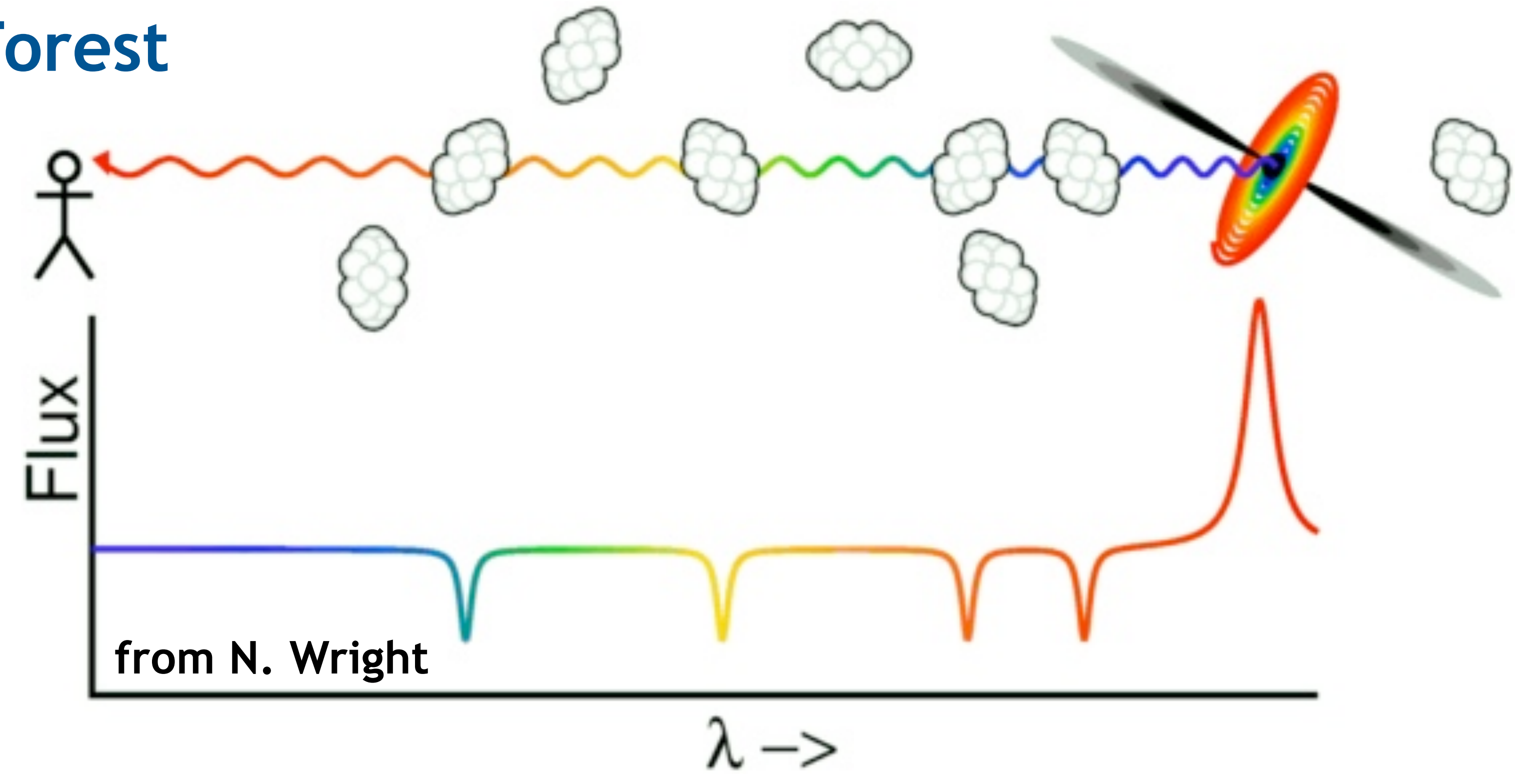

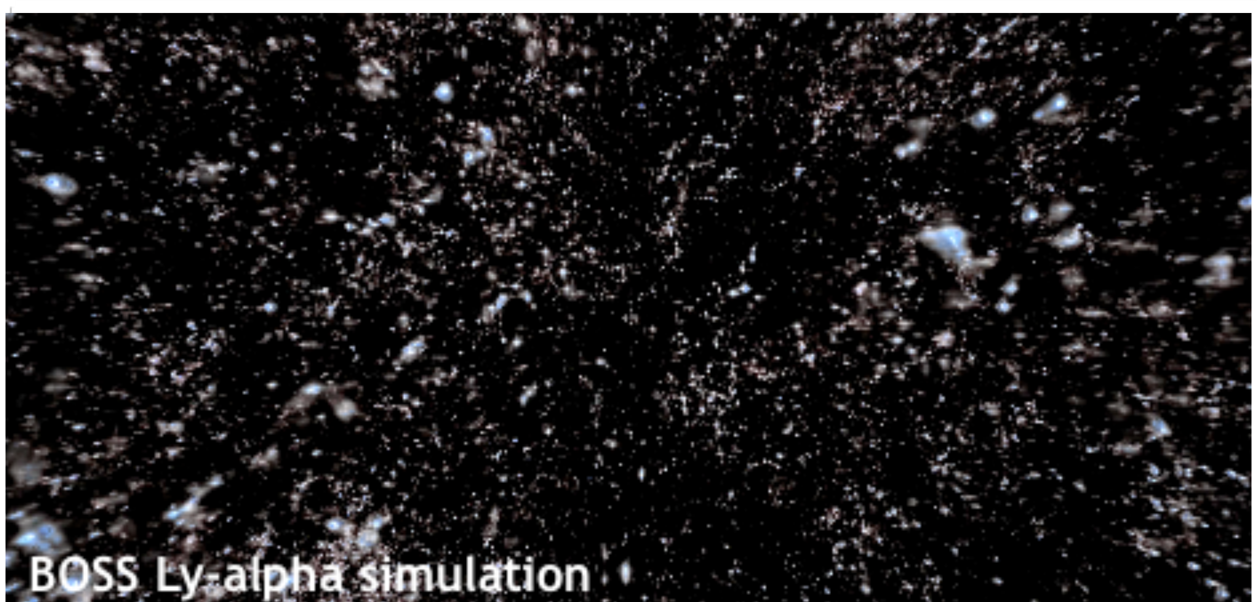

Roadrunner view (halos) of the Universe at z=2 from a 64 billion particle run (9 runs on one weekend)

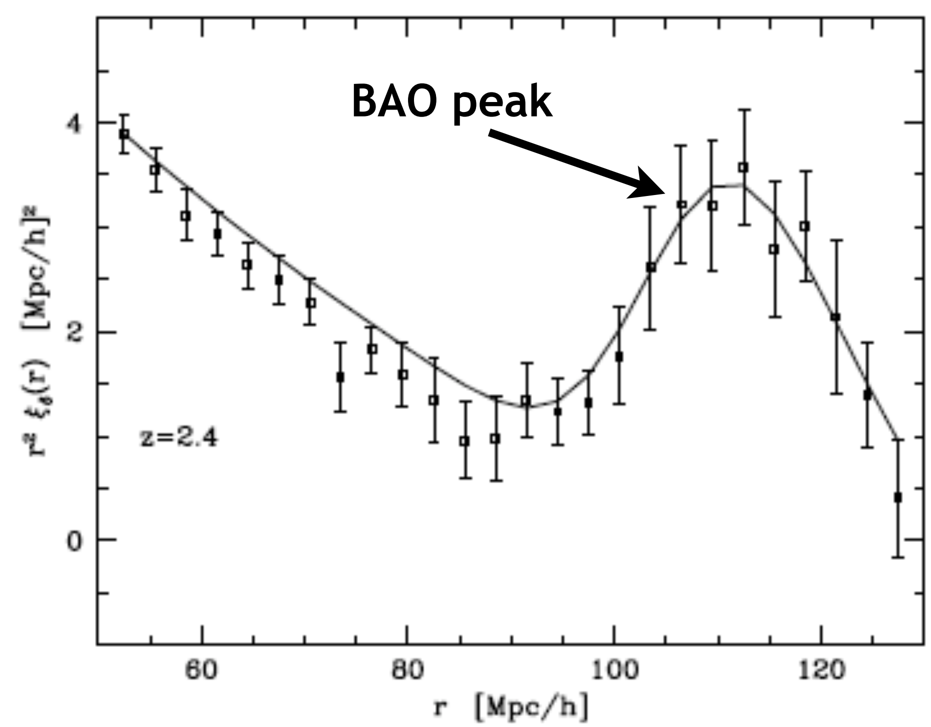

M. White et al. 2010

Computing the Universe: Simulating Surveys

- **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, ~$10^8 M_\odot - 10^{10} M_\odot$, N~$10^{11} - 10^{12}$
- **Force Resolution:** ~kpc, yields a (global) spatial dynamic range of **10**^{6}
- **Throughput:** Large numbers of simulations required (**100**'s **--1000**'s), development of analysis suites, and emulators; peta-exascale computing exploits
- **Computationally very challenging!** HACC is aimed to meet these requirements

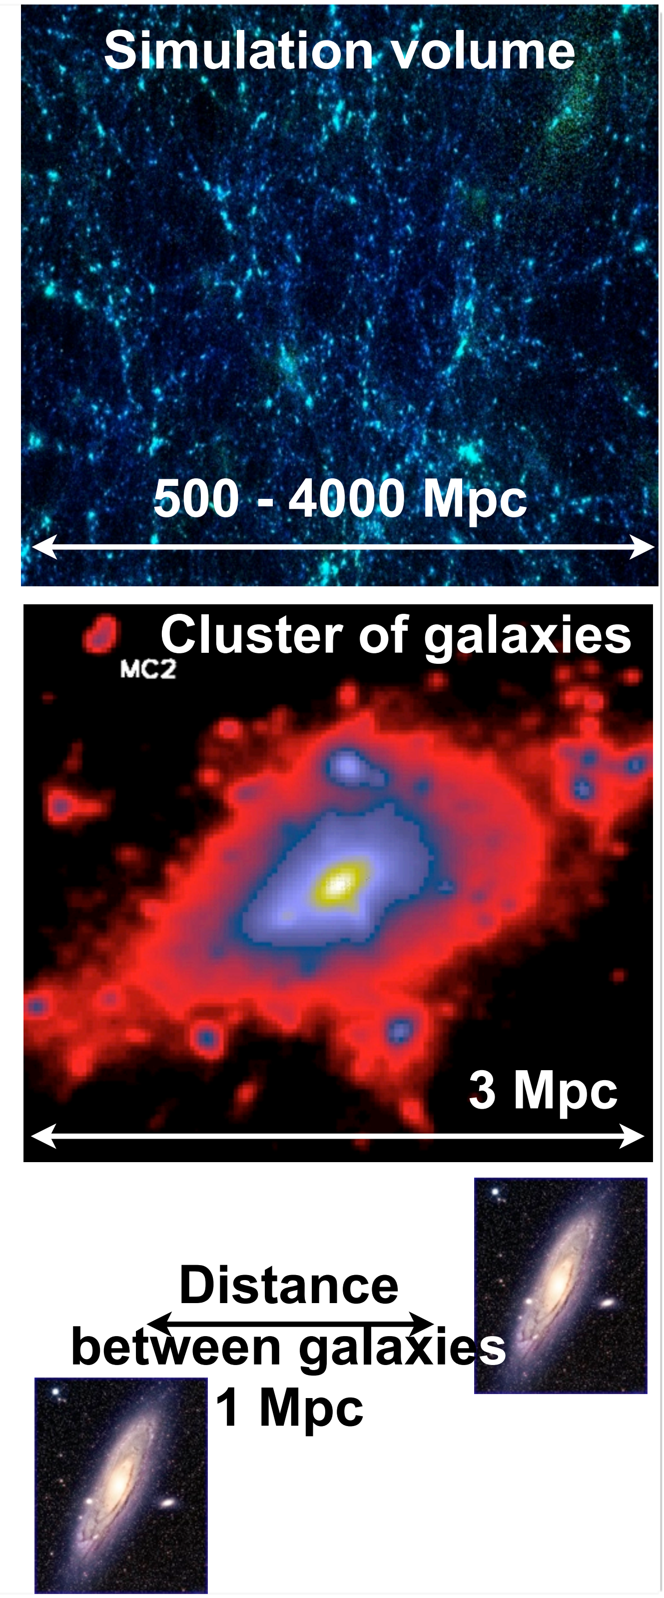

HACC (MC³) Performance on Roadrunner

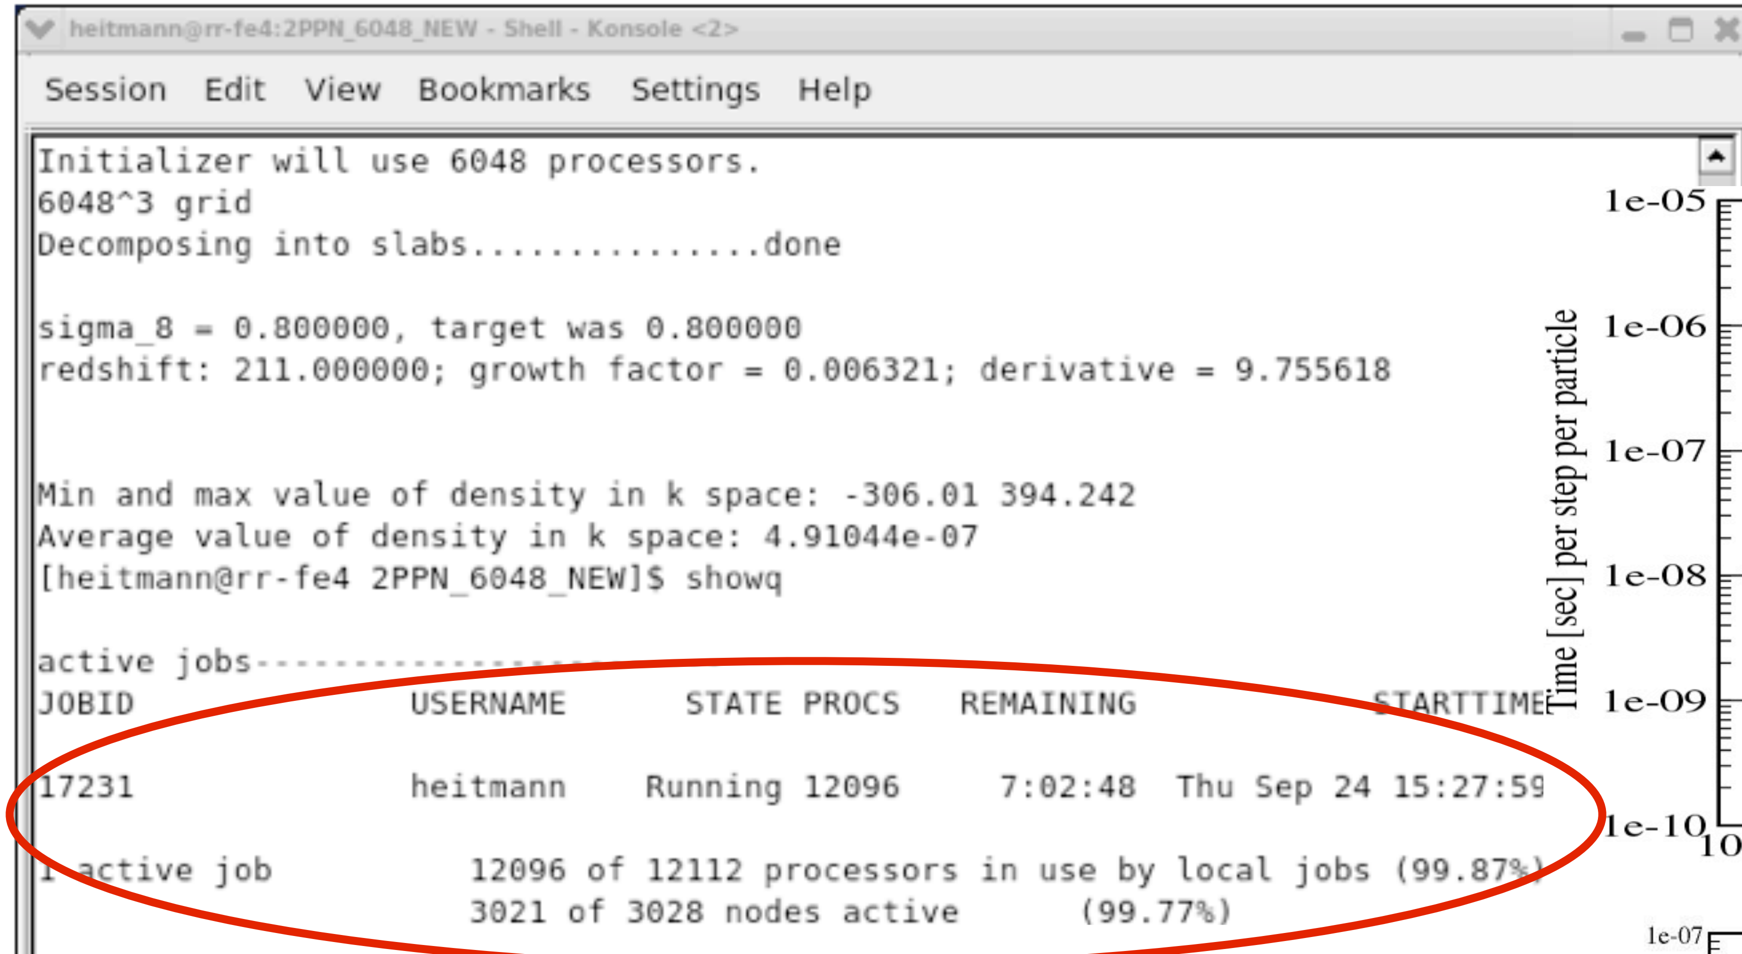

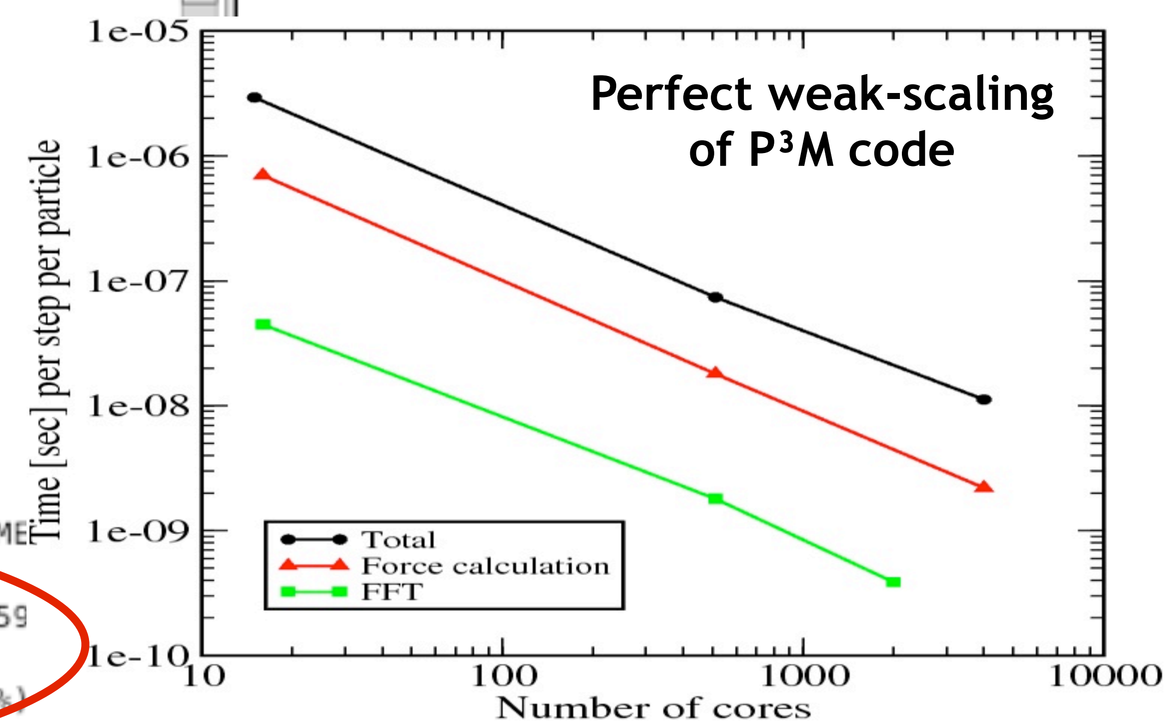

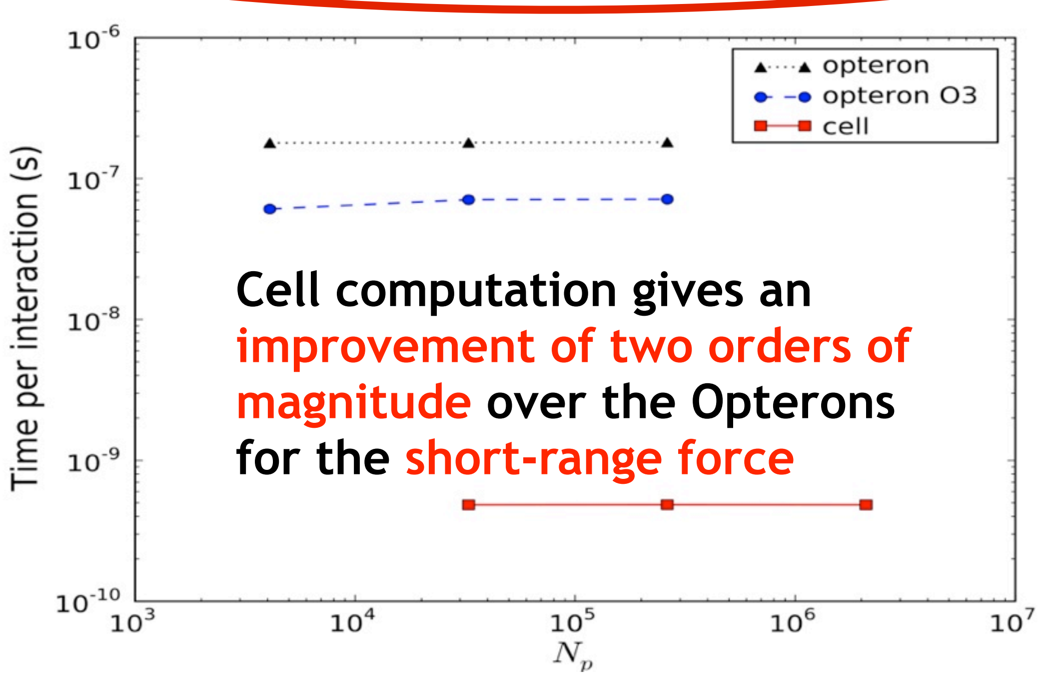

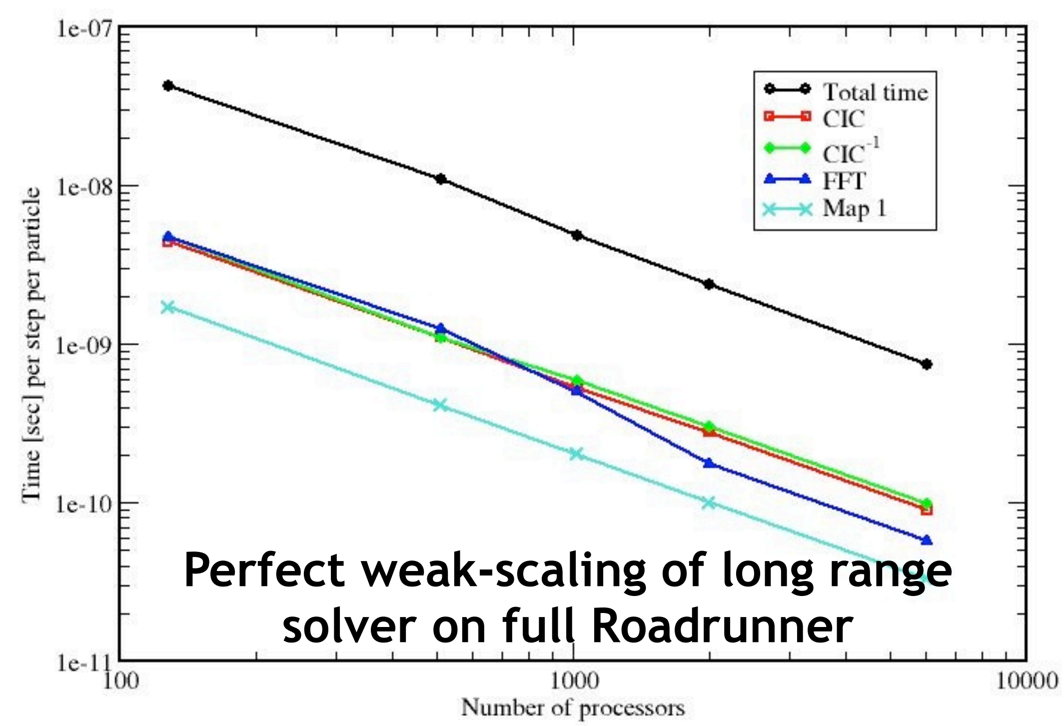

HACC In Situ Analysis

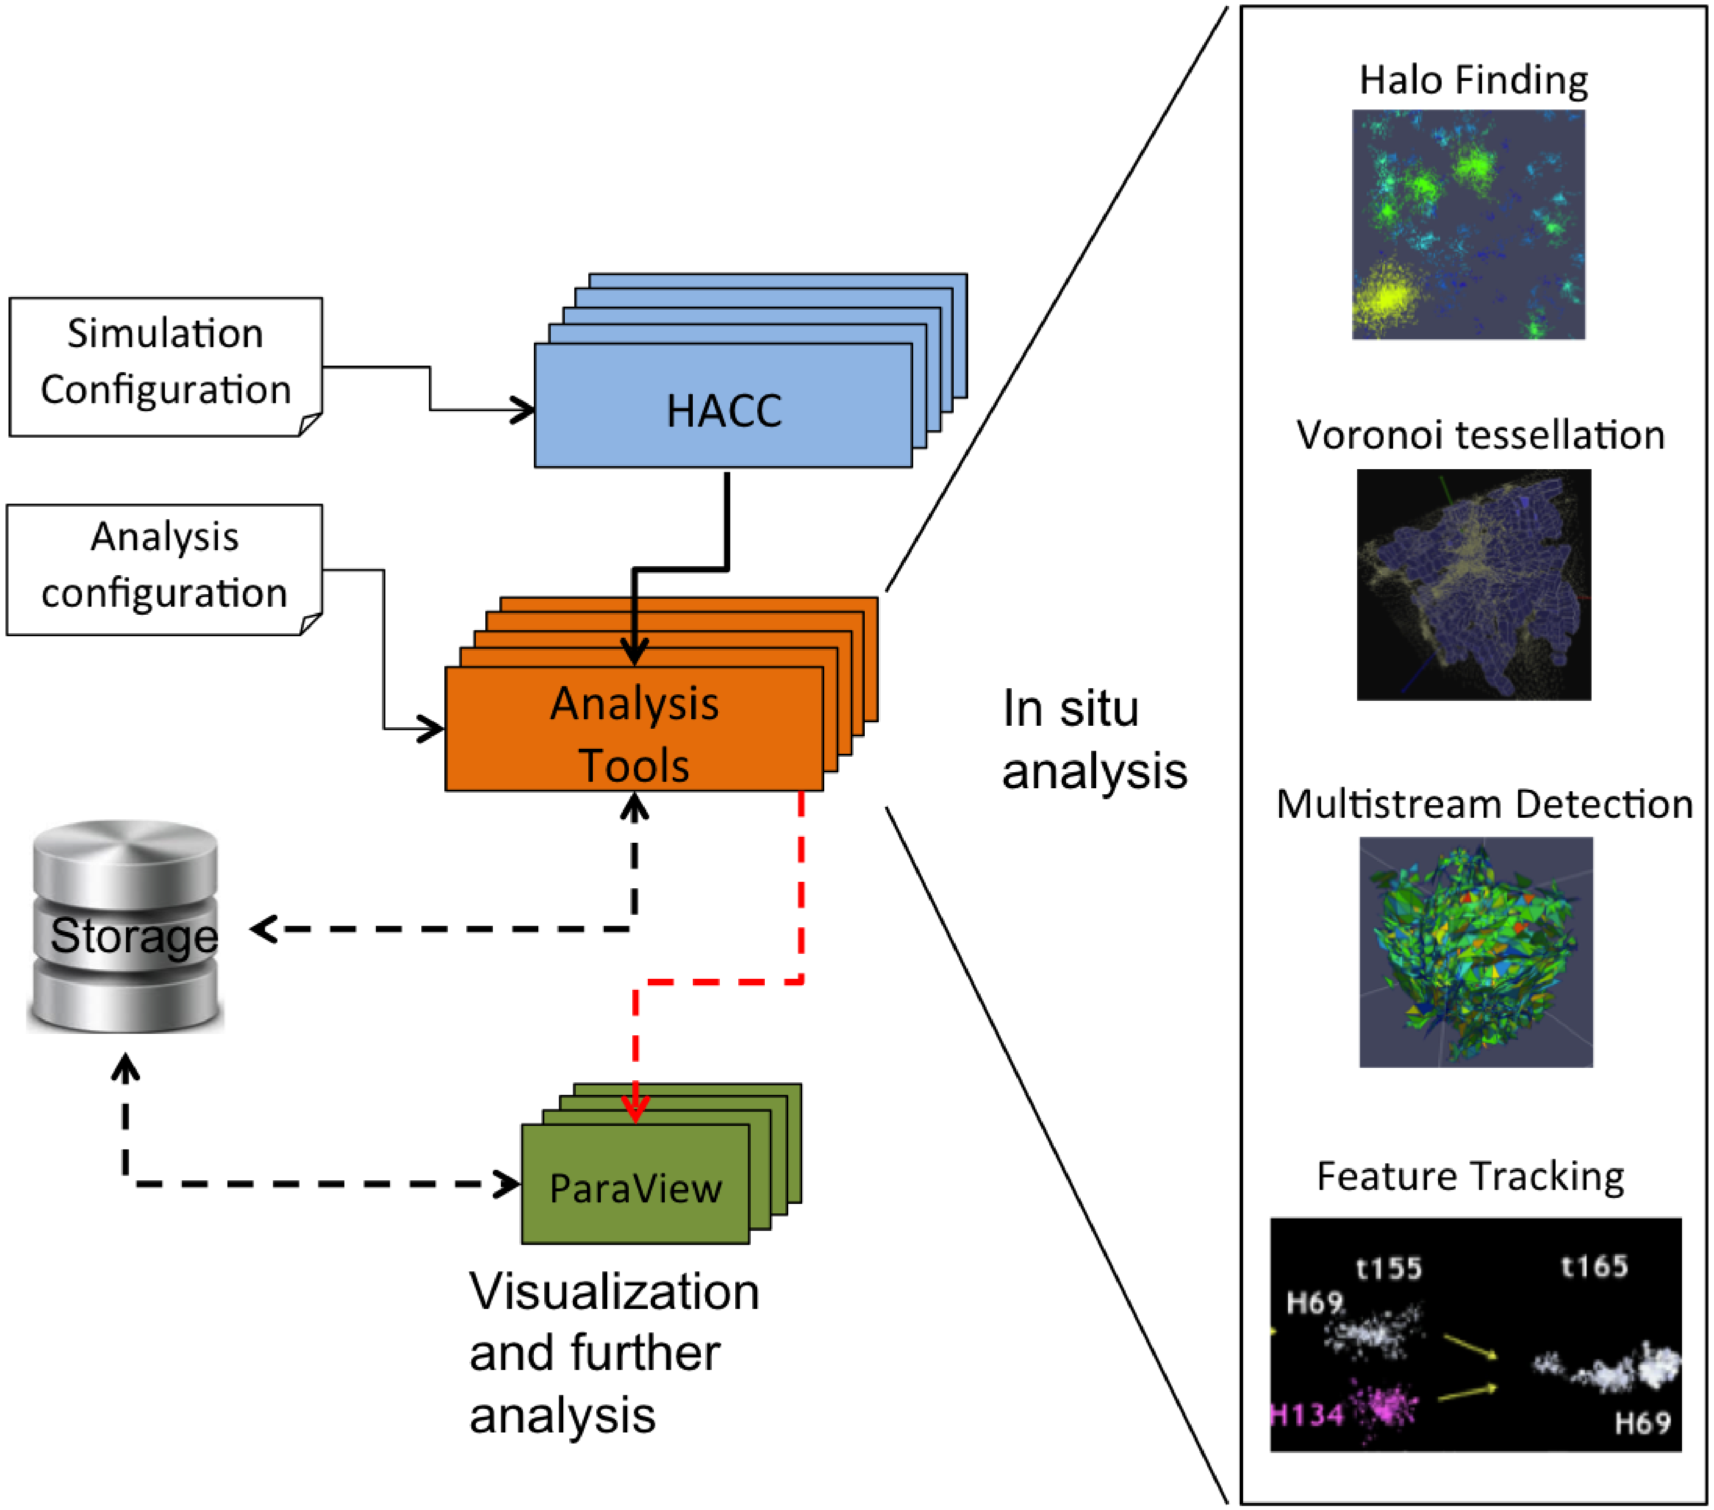

The Story Continues … and makes it into “Die Süddeutsche”

Newspaper I read every morning

sueddeutsche.de

Politik | Wirtschaft | Geld | Kultur | Sport | Leben | Karriere | München & Region | Bayern | M

Home > Digital   Supercomputer - Rasend schnell

Supercomputer

Rasend schnell

“He is the fastest calculator in the world:”

Er ist der schnellste Rechner der Welt: der amerikanische Supercomputer "Roadrunner" hat die Petaflop-Grenze geknackt. Sein Job: die Simulation von Atombombenexplosionen.

Twittern 0 | Empfehlen | Senden | +1 0

Ein Rechner der US-Regierung schafft erstmals mehr als eine Billiarde Operationen in der Sekunde (Petaflops) und ist damit nun der schnellste Computer der Welt. Das berichten das US-Energieministerium und der Hersteller IBM am Montag.

Der Computer namens Roadrunner wurde am Los Alamos National Laboratory (LANL) in New Mexico installiert. Er wird zuvorderst für die Forschung an US-Atomwaffen rechnen. Der neu konstruierte Roadrunner ist auf einen Schlag mehr als doppelt so schnell wie der bisherige Spitzenreiter der "Top 500"-Liste der Supercomputer.

Anfangs soll der Roadrunner aber vor allem wissenschaftliche Probleme lösen. Beispielsweise sind Tests von Klimamodellen vorgesehen, doch rechnet das LANL mit Anwendungen in diversen Bereichen, darunter die Kosmologie, die Entwicklung von Antibiotika oder die Astrophysik. Danach wird der Supercomputer laut LANL militärischen Aufgaben zugeteilt und unter Geheimhaltung Explosionen nuklearer Waffen simulieren, um physikalische Modelle zu verbessern und das Vertrauen in das nukleare Arsenal der USA ohne tatsächliche Atomtests zu erhalten.

Supercomputer mit Vorbildfunktion “leads by example”

"Für uns und die HPC-Community ist es hoch erfreulich, dass es ein System gibt, das diese Marke geknackt hat", sagt Thomas Lippert, Leiter des Jülich Supercomputing Centre. Dadurch werde dem Supercomputing berechtigte Aufmerksamkeit zuteil.

Technologisch dürfte Roadrunner Vorbildwirkung haben. "Es zeichnet sich ab, dass Hybrid-Technologie auf jeden Fall Zukunft haben", meint Lippert. Damit sind Systeme gemeint, die klassische CPUs mit Beschleunigern wie beispielsweise den Cell-Chips oder Grafikprozessoren kombinieren.

“technology of the future”

HACC in a Nutshell

- **Long-range/short range force splitting:**
  - **Long-range:** Particle-Mesh solver, C/C++/MPI, **unchanged for different architectures**, FFT performance dictates scaling (custom pencil decomposed FFT)
  - **Short-range: Depending on node architecture** switch between tree and particle-particle algorithm; tree needs “thinking” (building, walking) but computationally less demanding (BG/Q, X86), PP easier but computationally more expensive (Cell, GPU)
- **Overload concept to allow for easy swap of short-range solver and minimization of communication (reassignment of passive/active in regular intervals)**

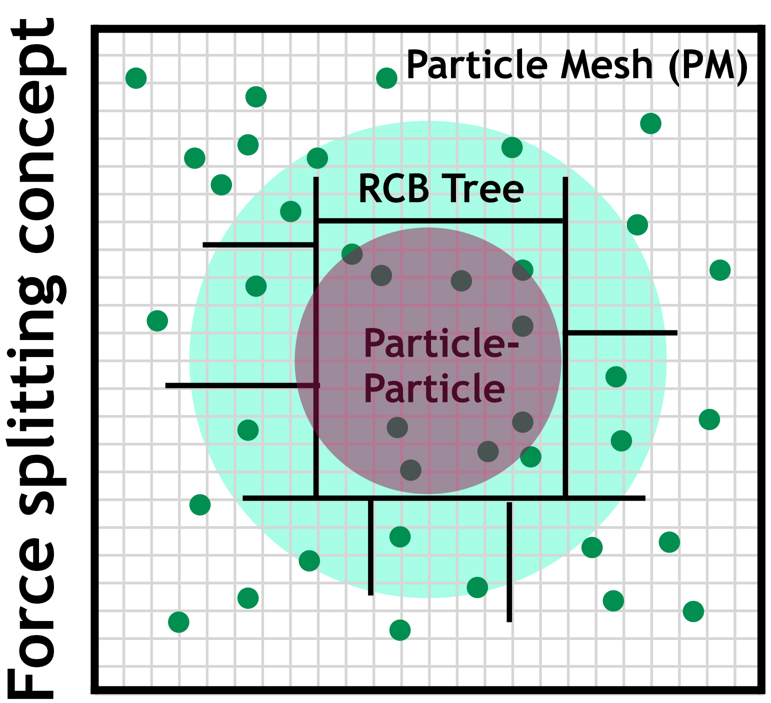

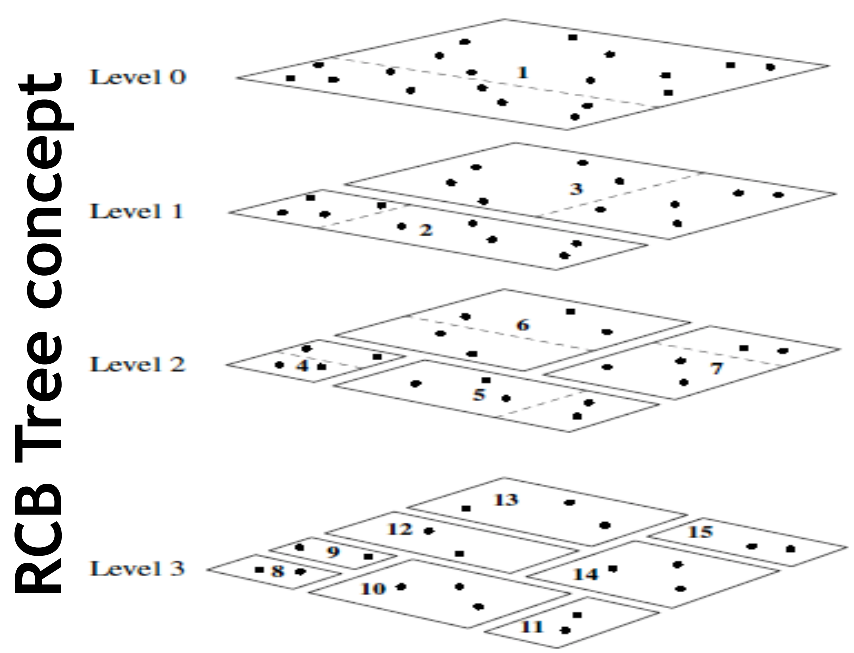

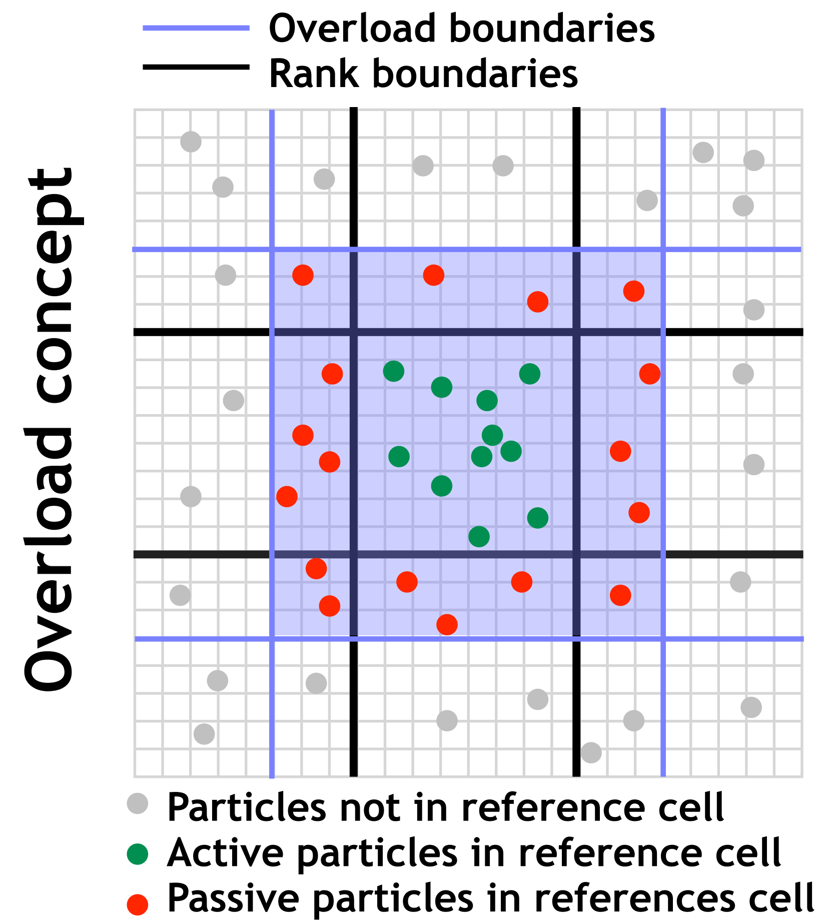

Splitting the Force: The Long-Range Solver

- **Spectral Particle-Mesh Solver:** Custom (large) FFT-based method -- uses (i) 6-th order Green function, (ii) 4th order spectral Super-Lanczos gradients, (iii) high-order spectral filtering to reduce grid anisotropy noise
- **Short-range Force:** Asymptotically correct semi-analytic expression for the difference between the Newtonian and the long-range force; uses a 5th order polynomial
- **Pencil-decomposed Parallel 3-D FFT:** Fast 3D-to-2D combinatorics, HACC scalability depends entirely on FFT performance
- **Time-stepping uses Symplectic Sub-cycling**

Long-range/short range split

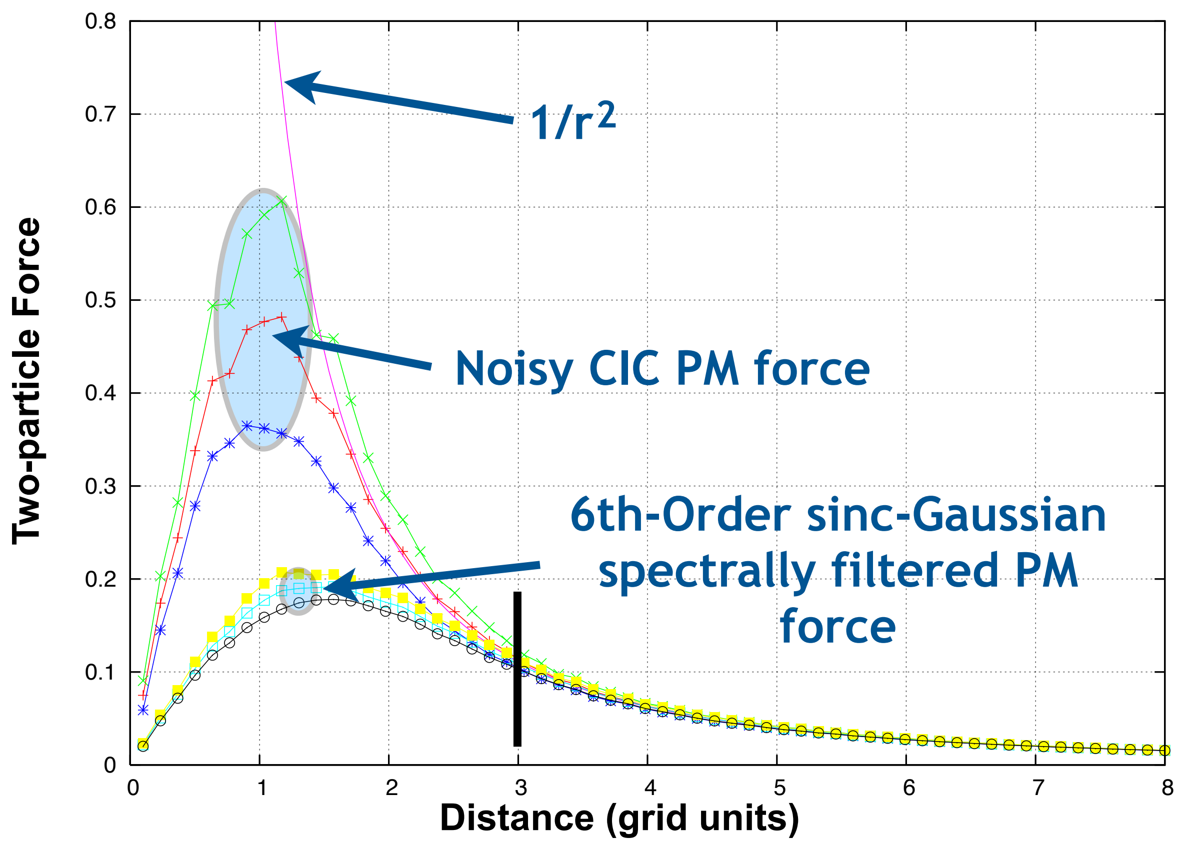

Decomposition