Exploring the Dark Universe

Katrin Heitmann

ATPESC Dinner Talk, August 2, 2019

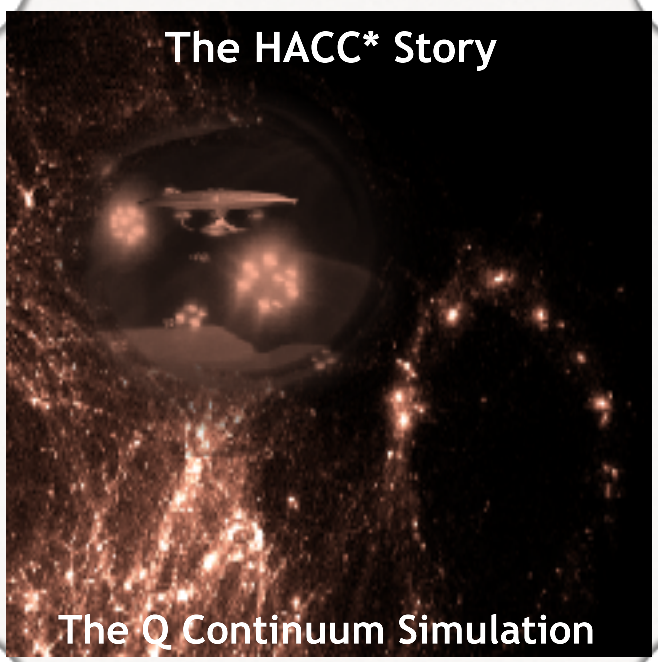

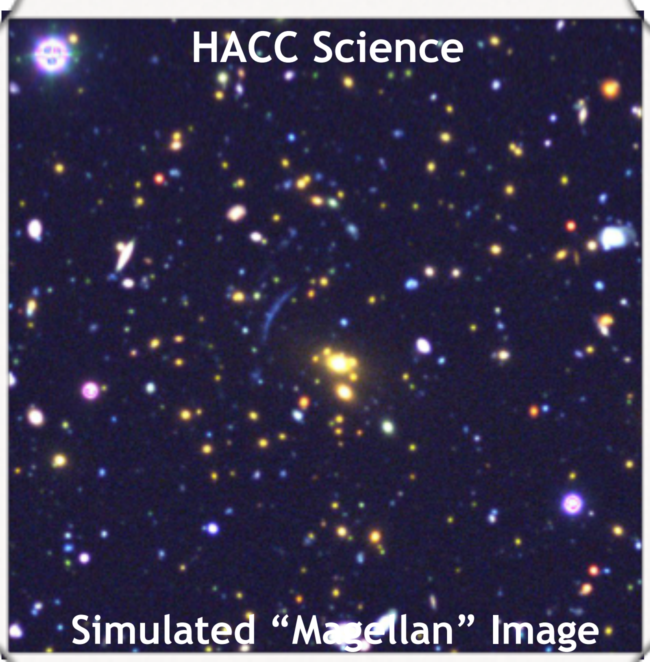

* HACC = Hardware/Hybrid Accelerated Cosmology Code

Thanks to many collaborators!

Physics/Astrophysics (Cosmology, Field Theory, Lattice QCD) • Computer Science • Statistics

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 recorded >500 million photometric objects, 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

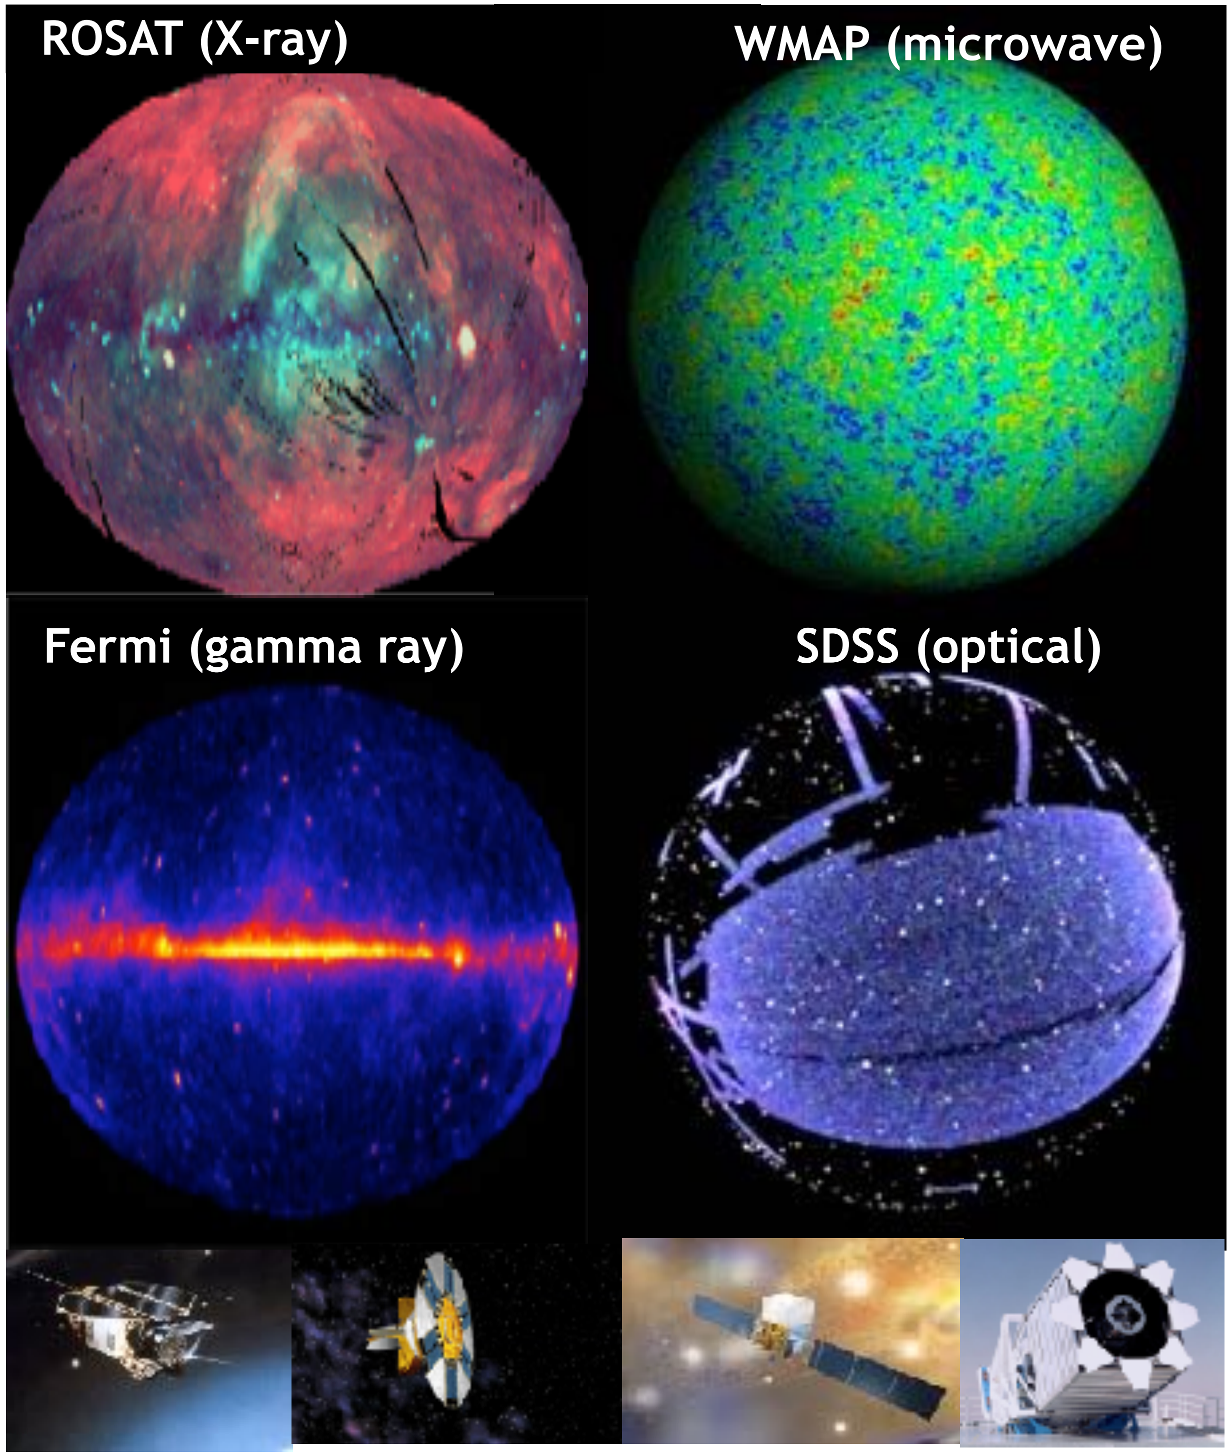

Observations from the SDSS: positions of 1,000,000 galaxies with redshifts (and therefore distance) leading to a 3-D map

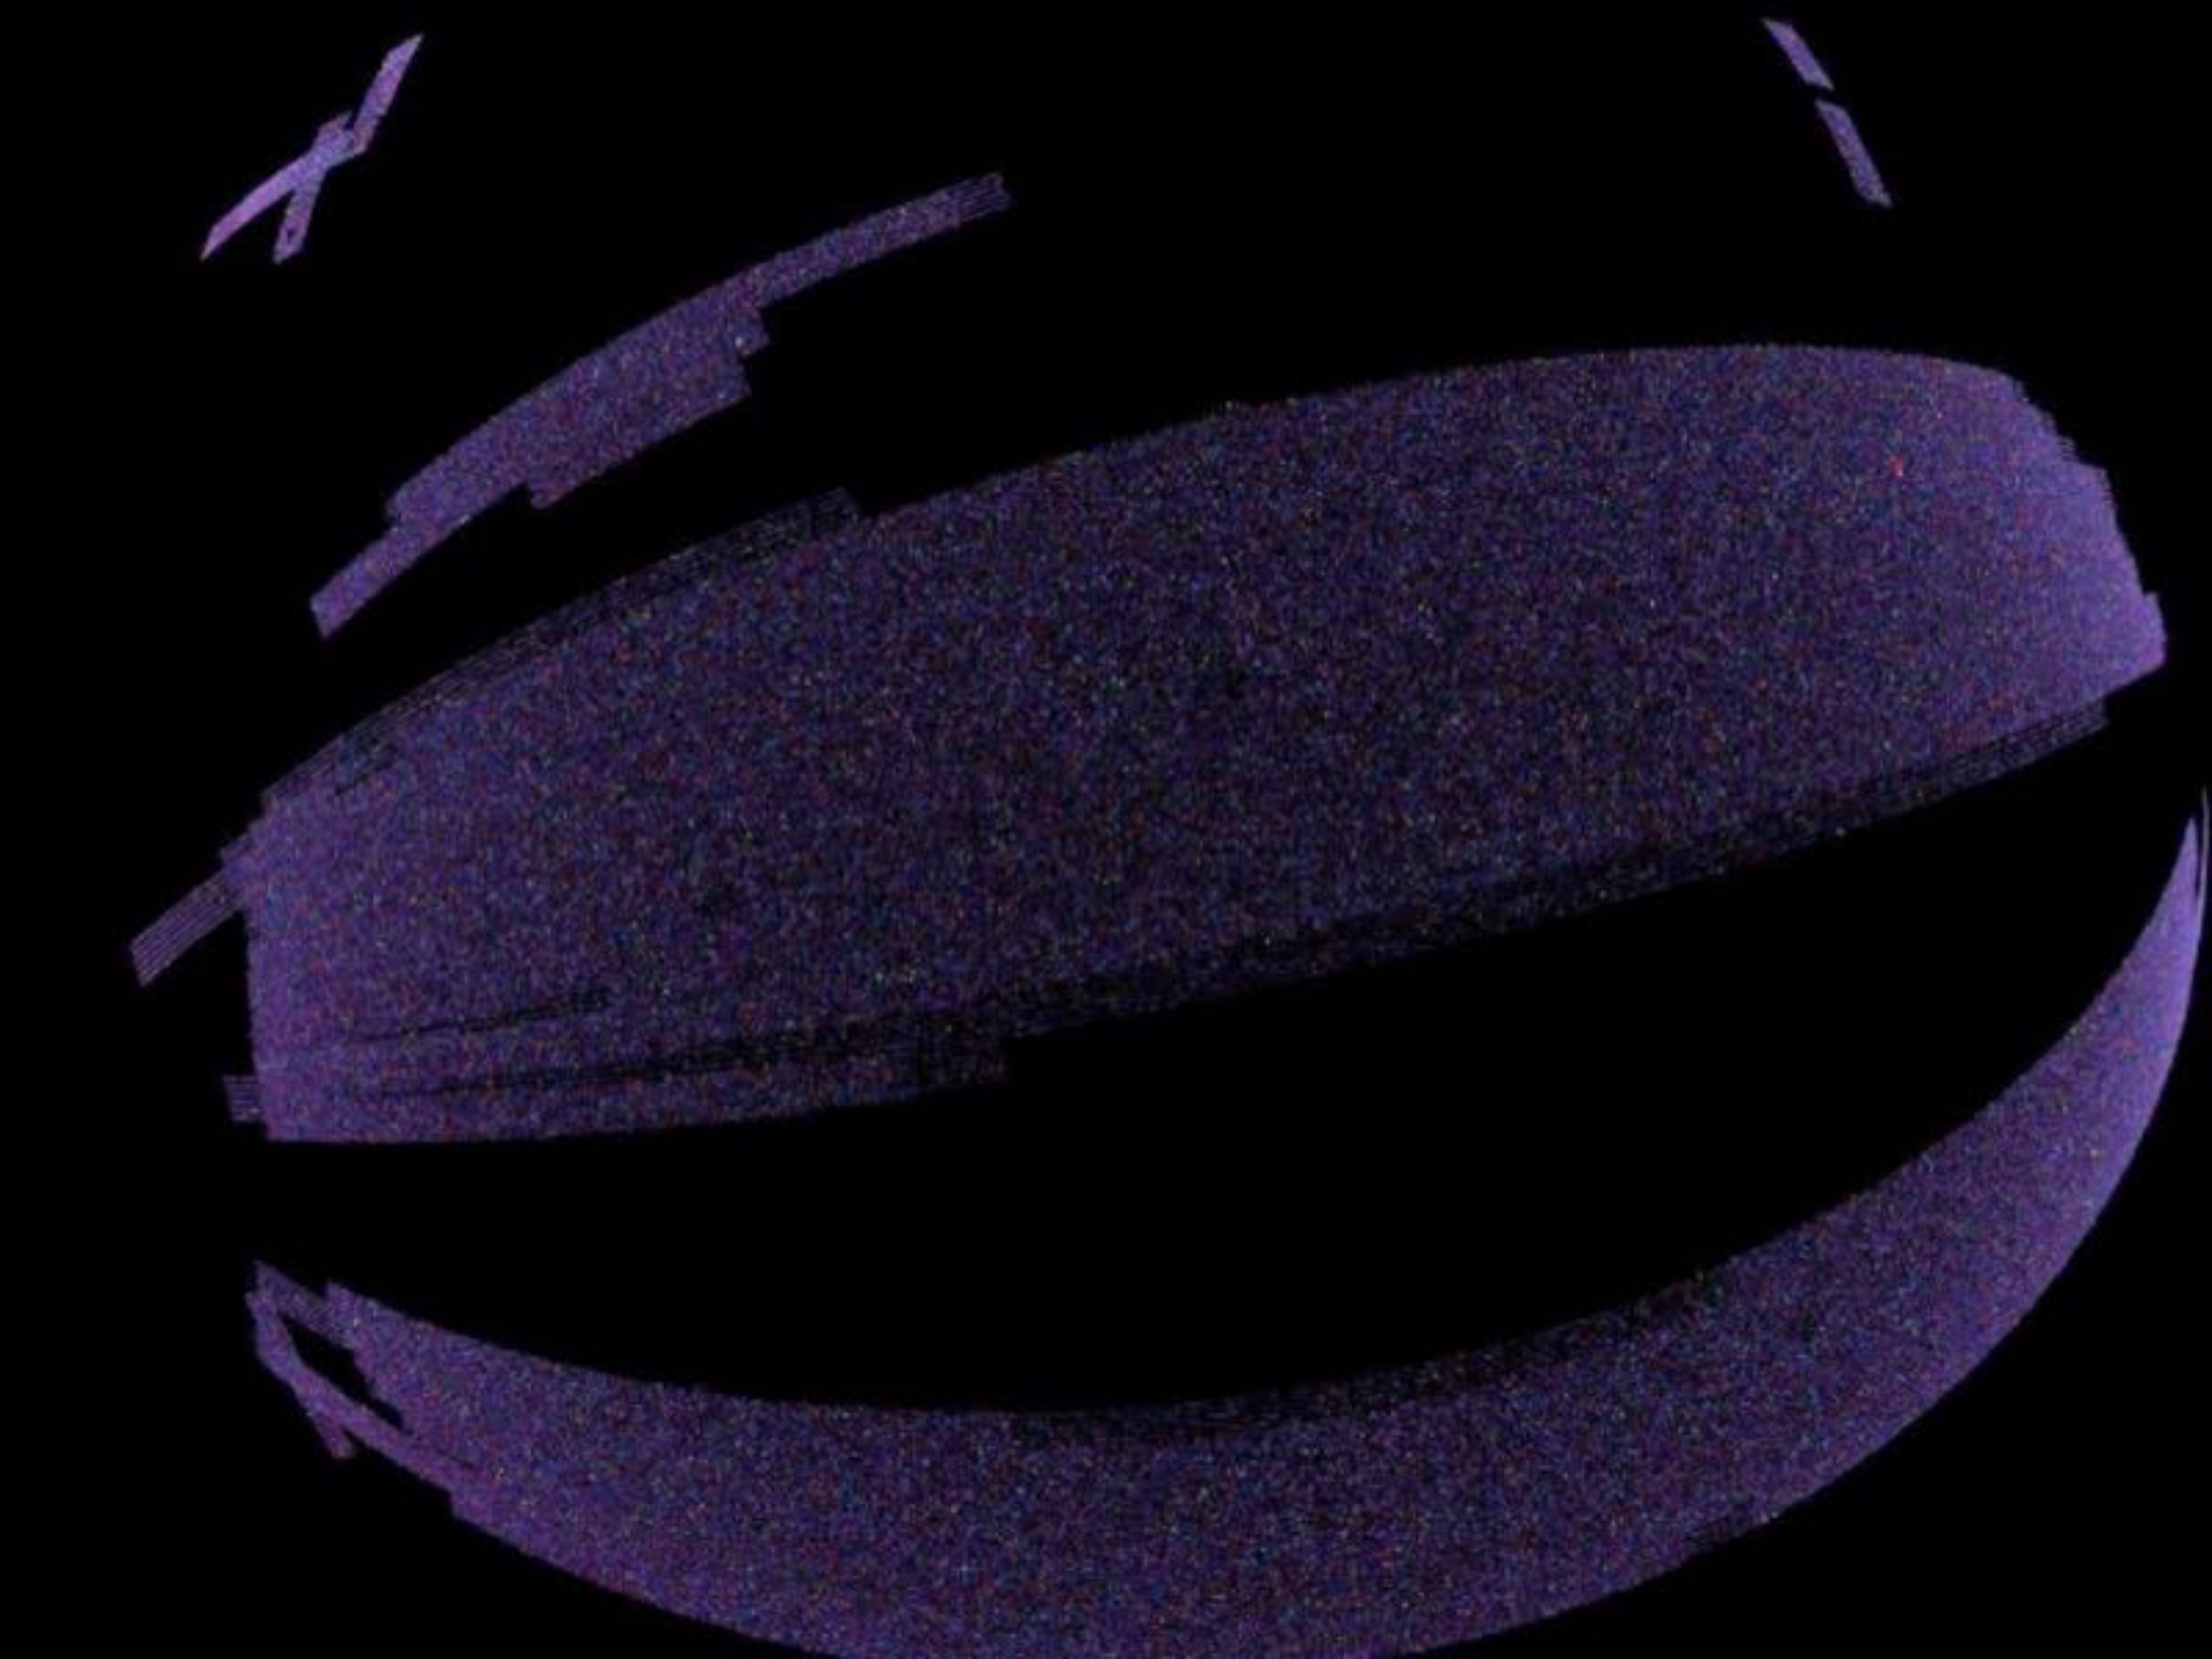

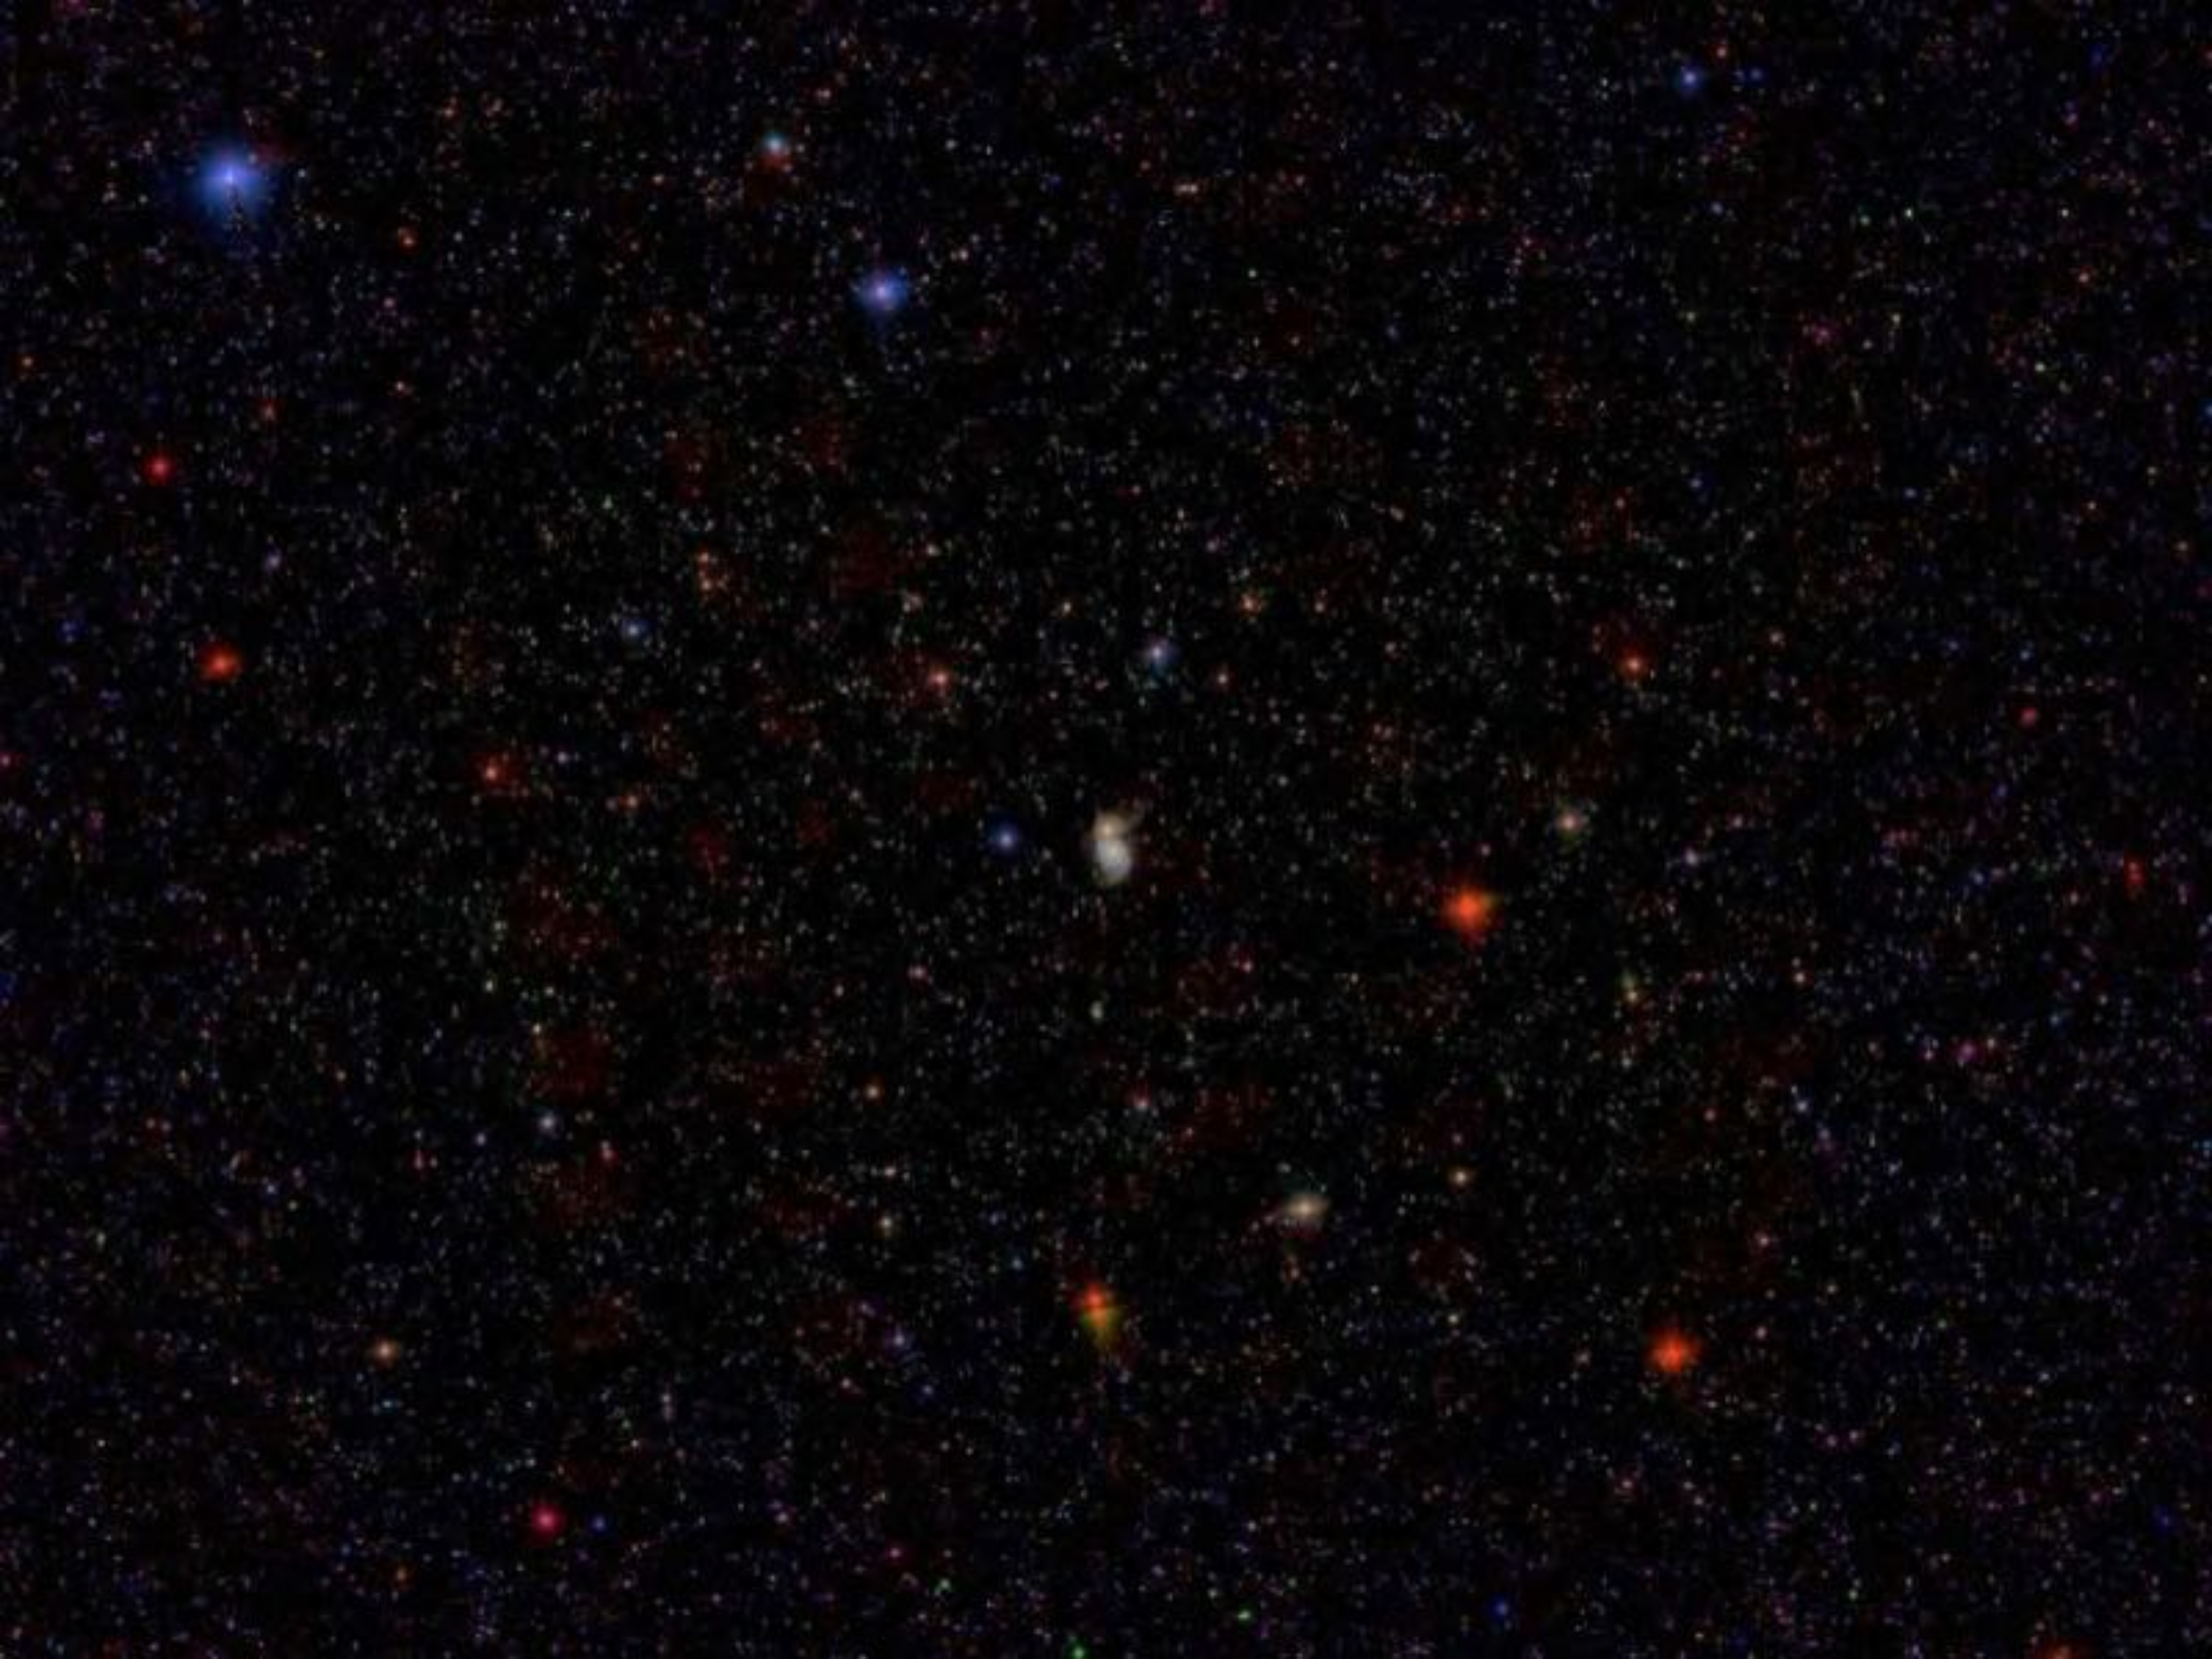

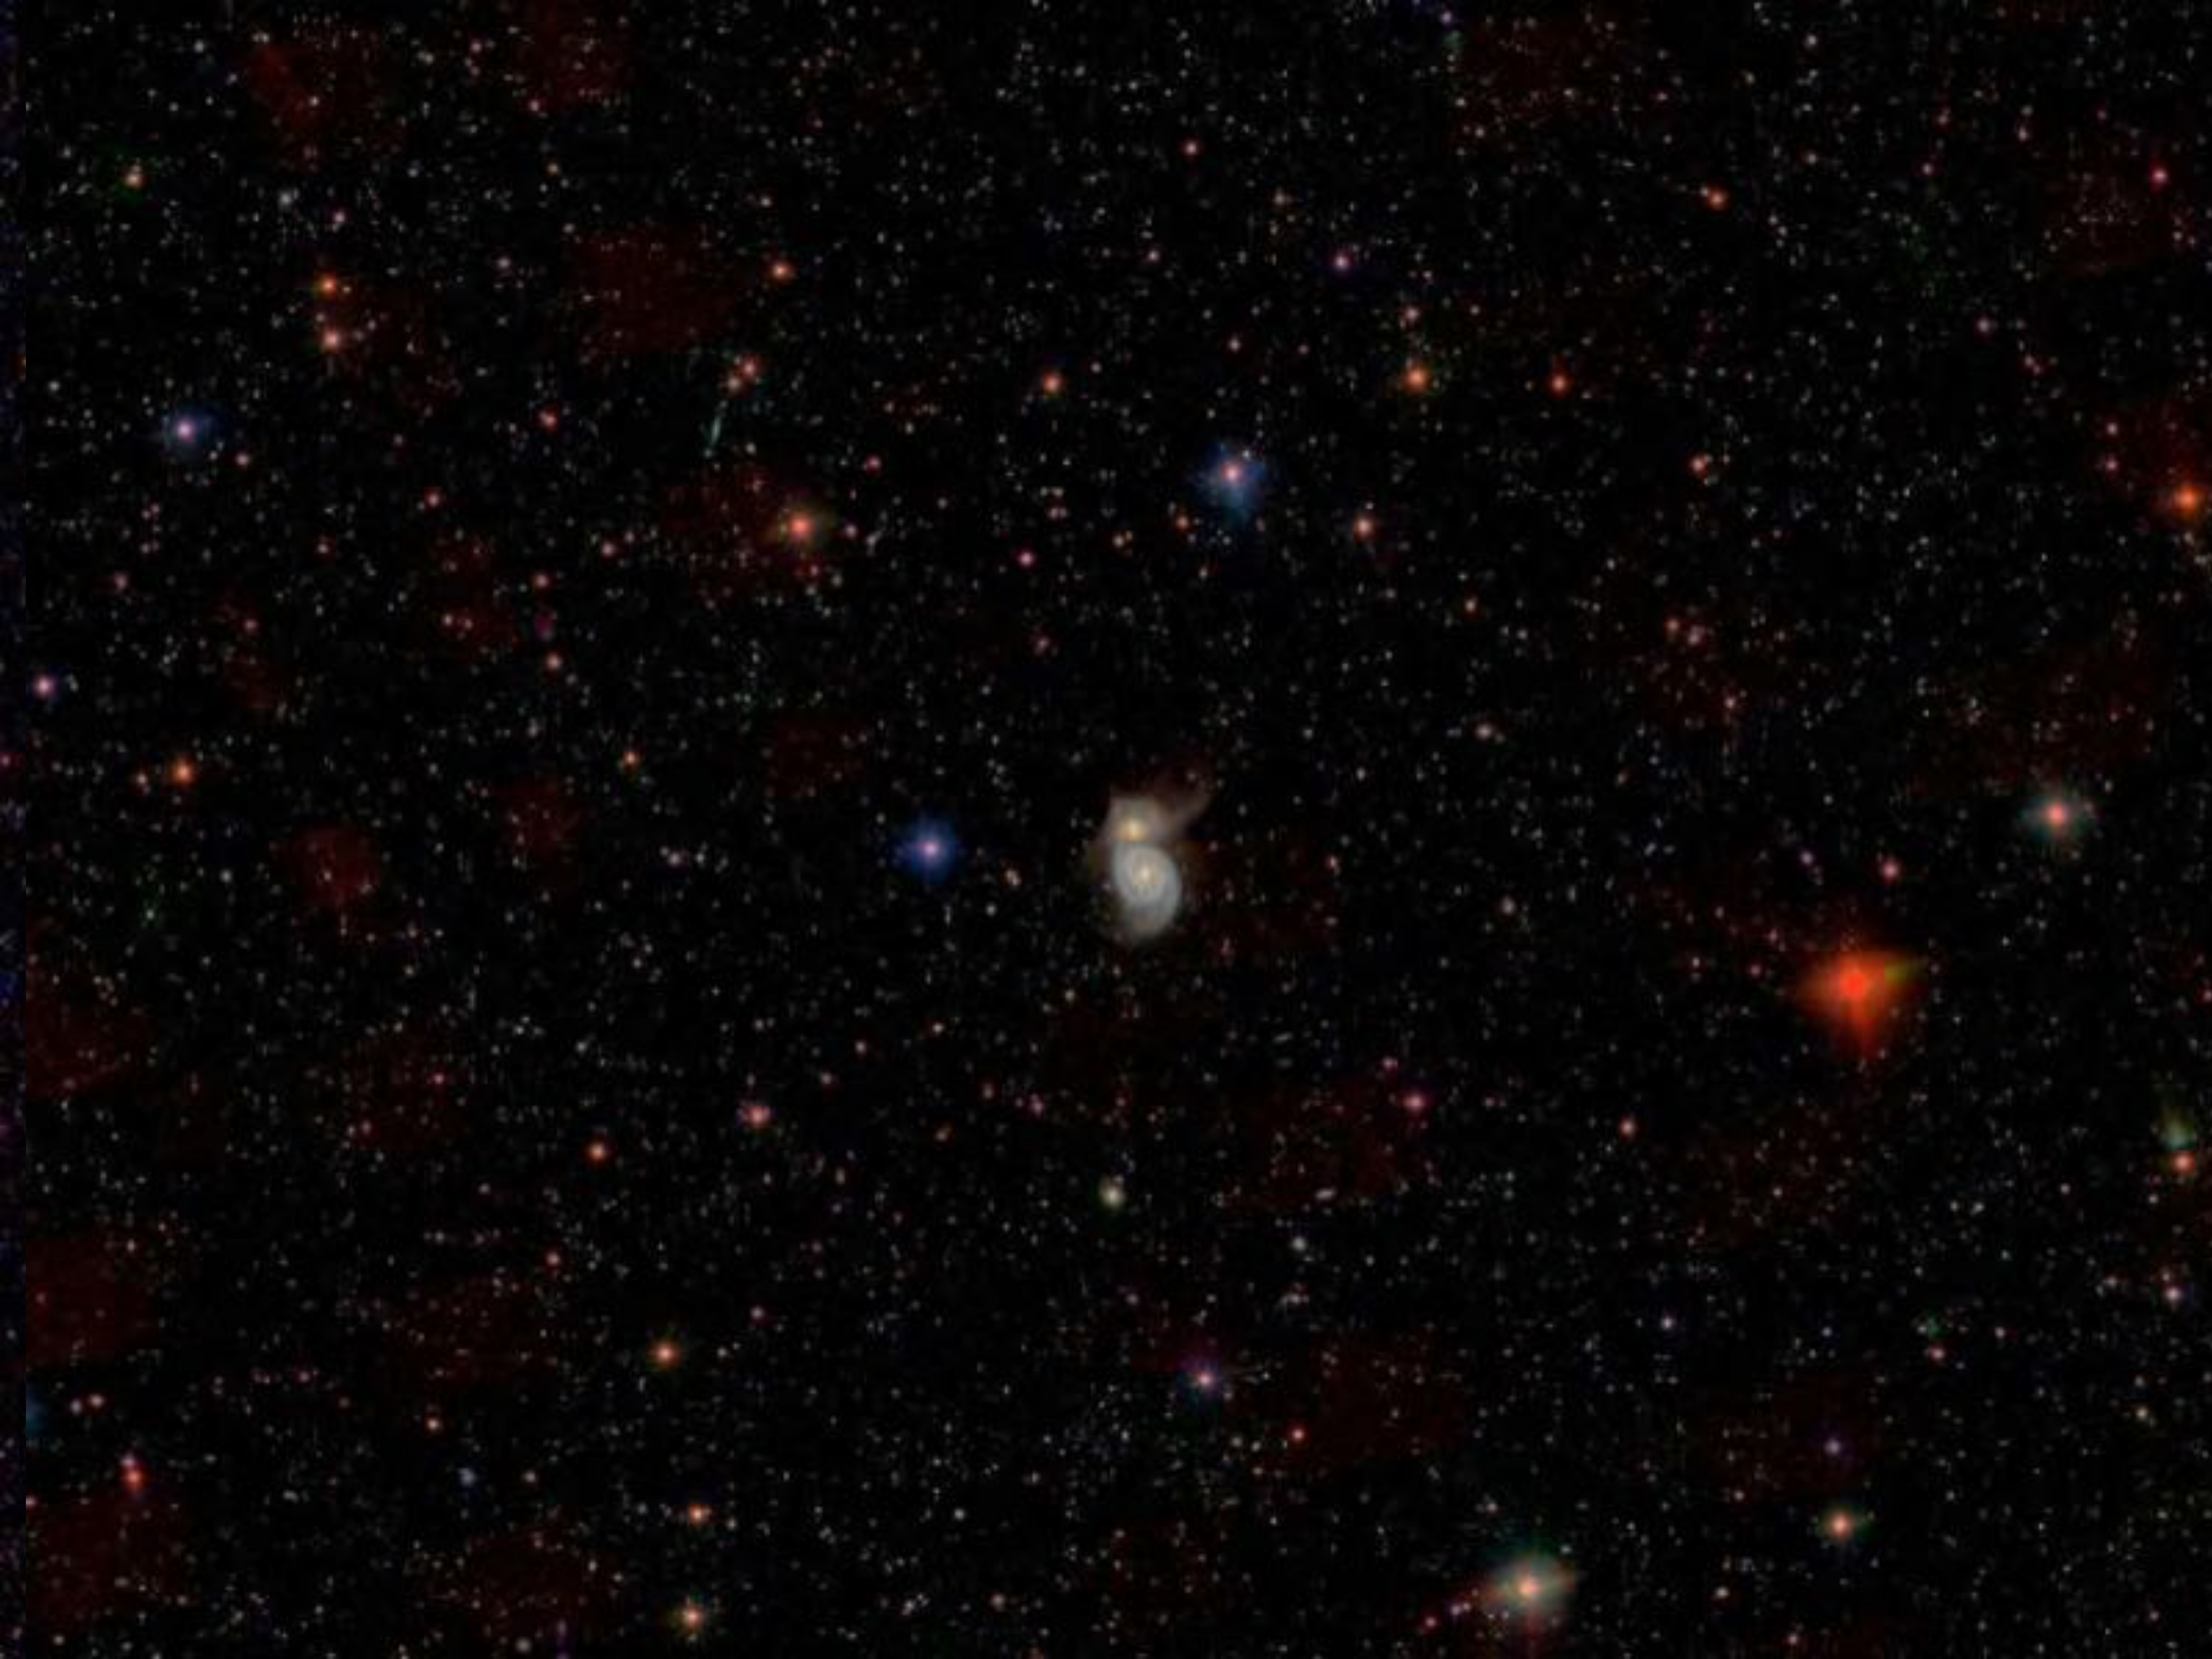

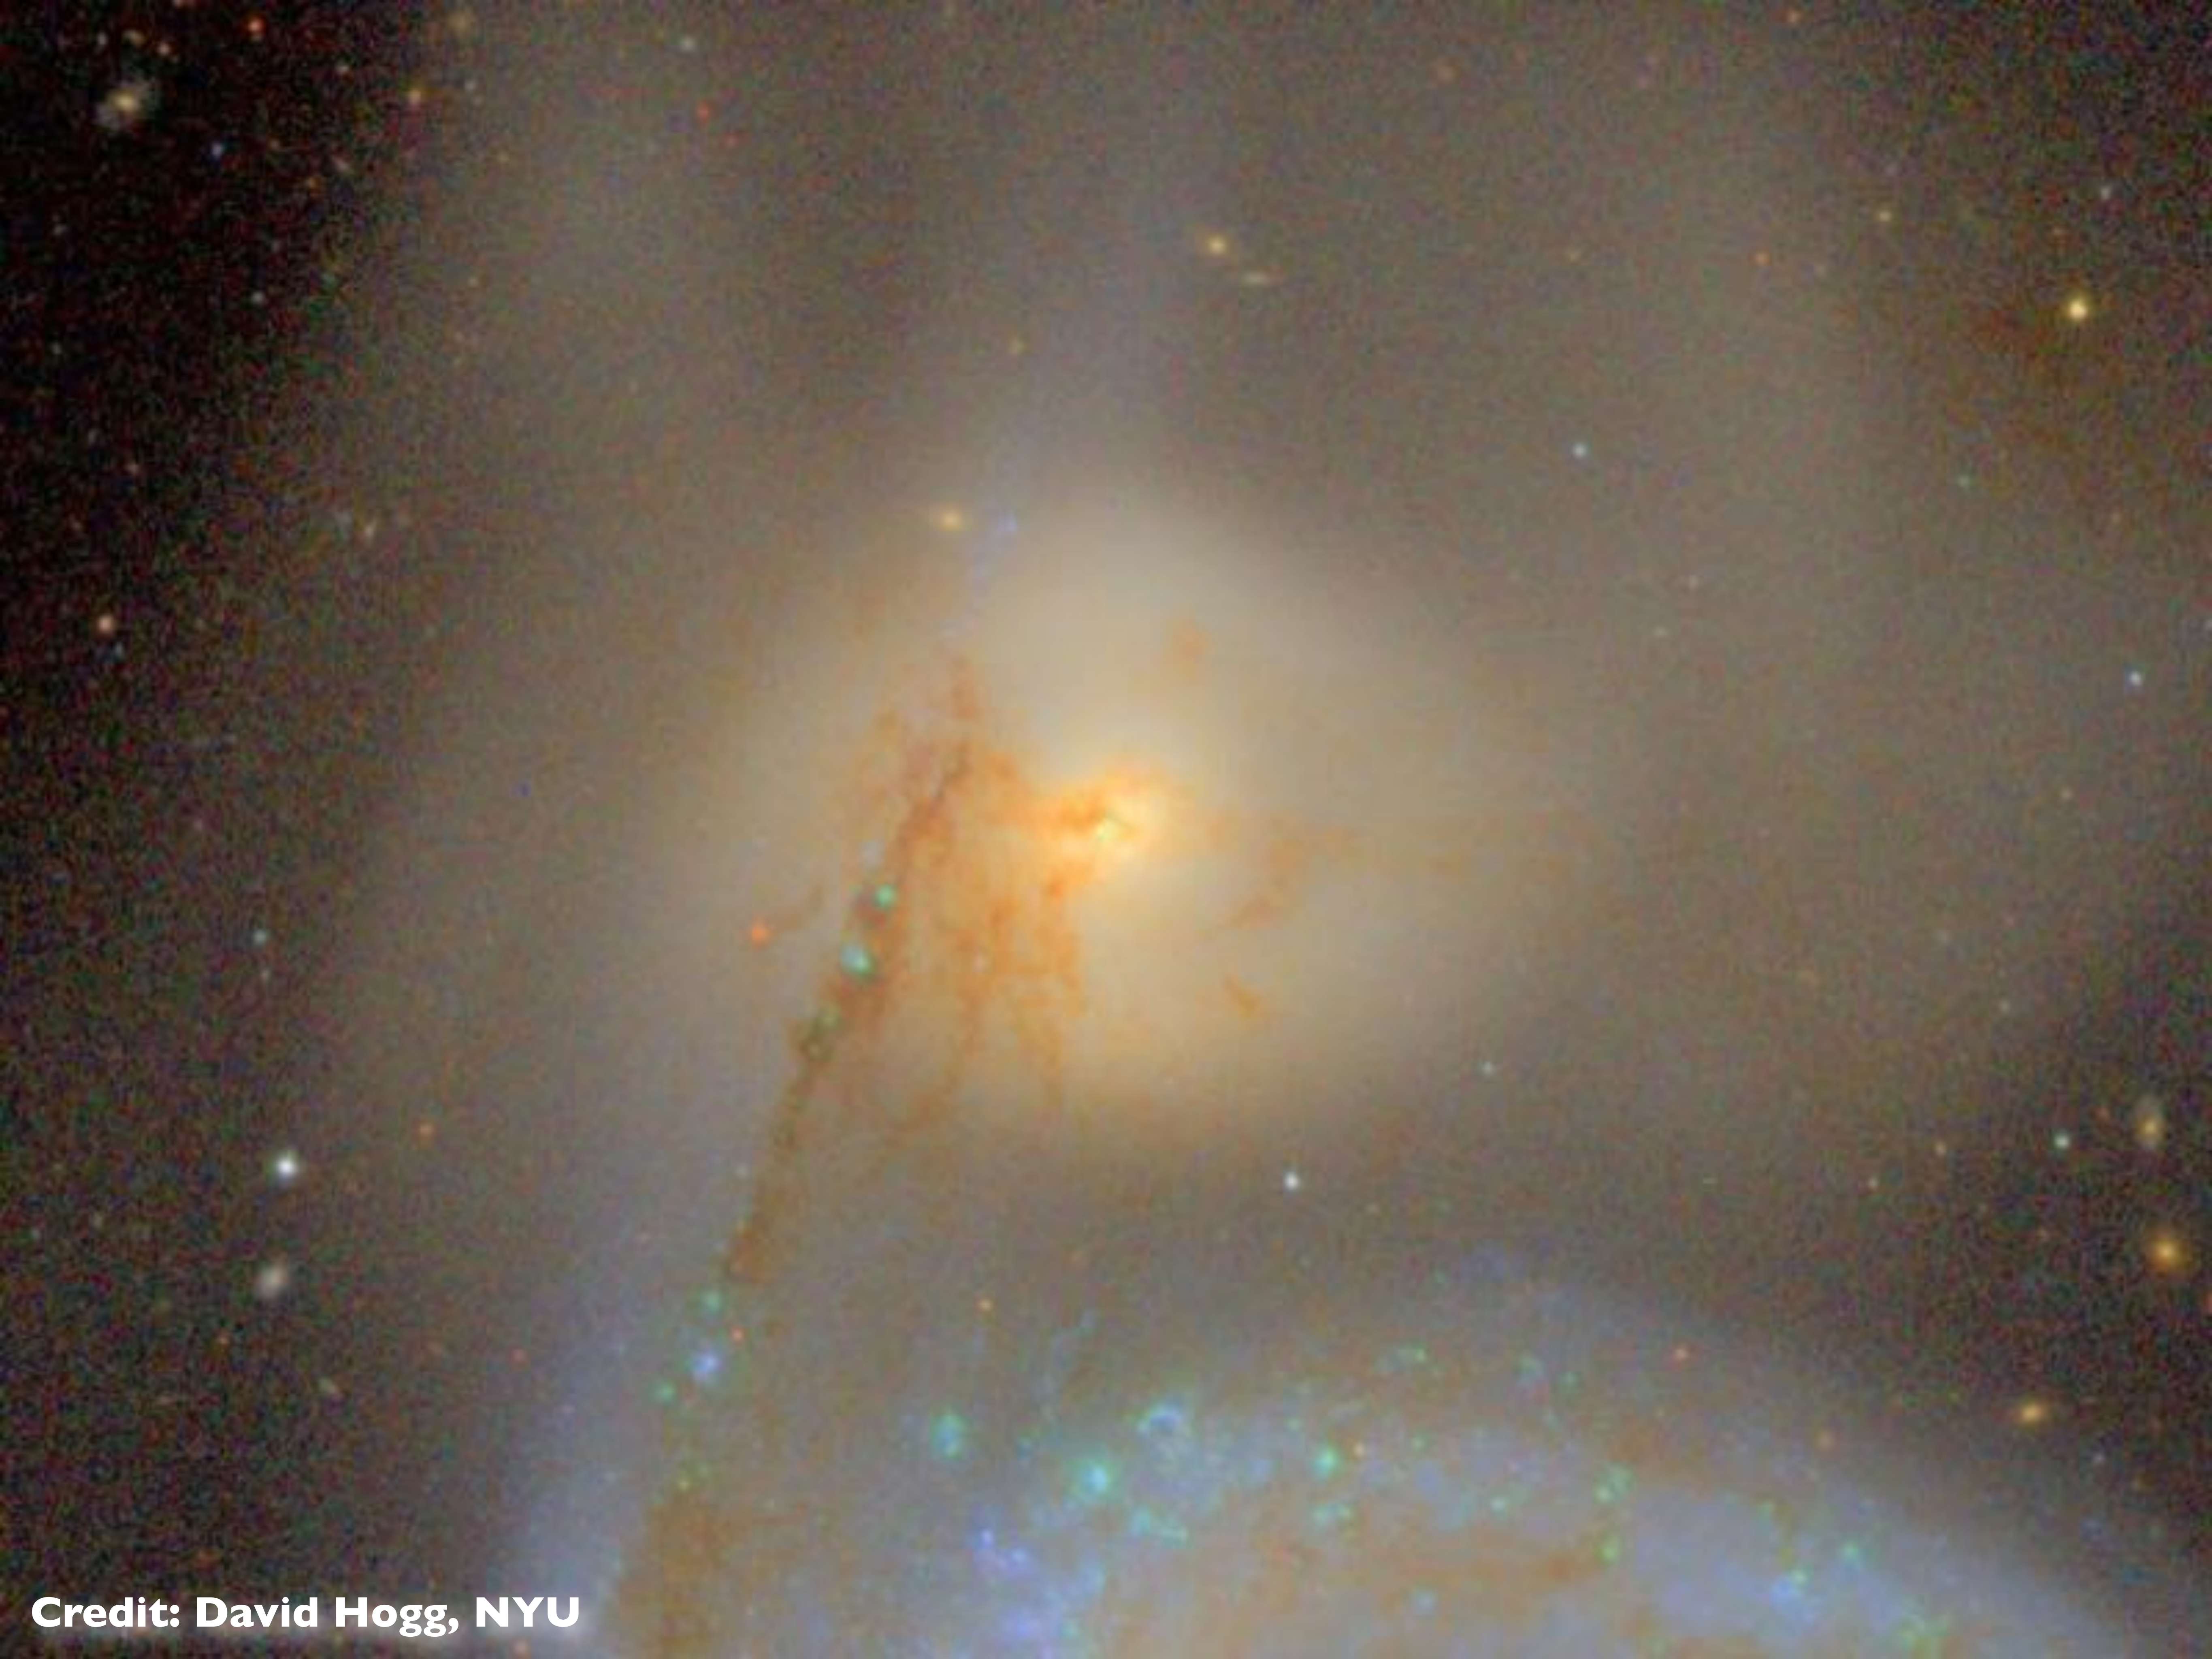
Credit: David Hogg, NYU

Relevant Numbers for Optical Surveys

- Survey volume
  - Surveys cover enormous observational volumes, ~(4 Gpc)^3, of order 12 billion light-years in linear scale
- Number of Galaxies
  - Number of galaxies in surveys can be in the tens of billions, typical intergalactic separation is ~1 Mpc
- Typical galaxy mass
  - The mass associated with a typical bright galaxy in surveys ranges from 10^{11}-10^{13} $M_{\odot}$
- Typical galaxy size
  - The length scale of the mass distribution is approximately 100 kpc

***1 parsec = 3.26 light-years = 31 trillion kilometers**

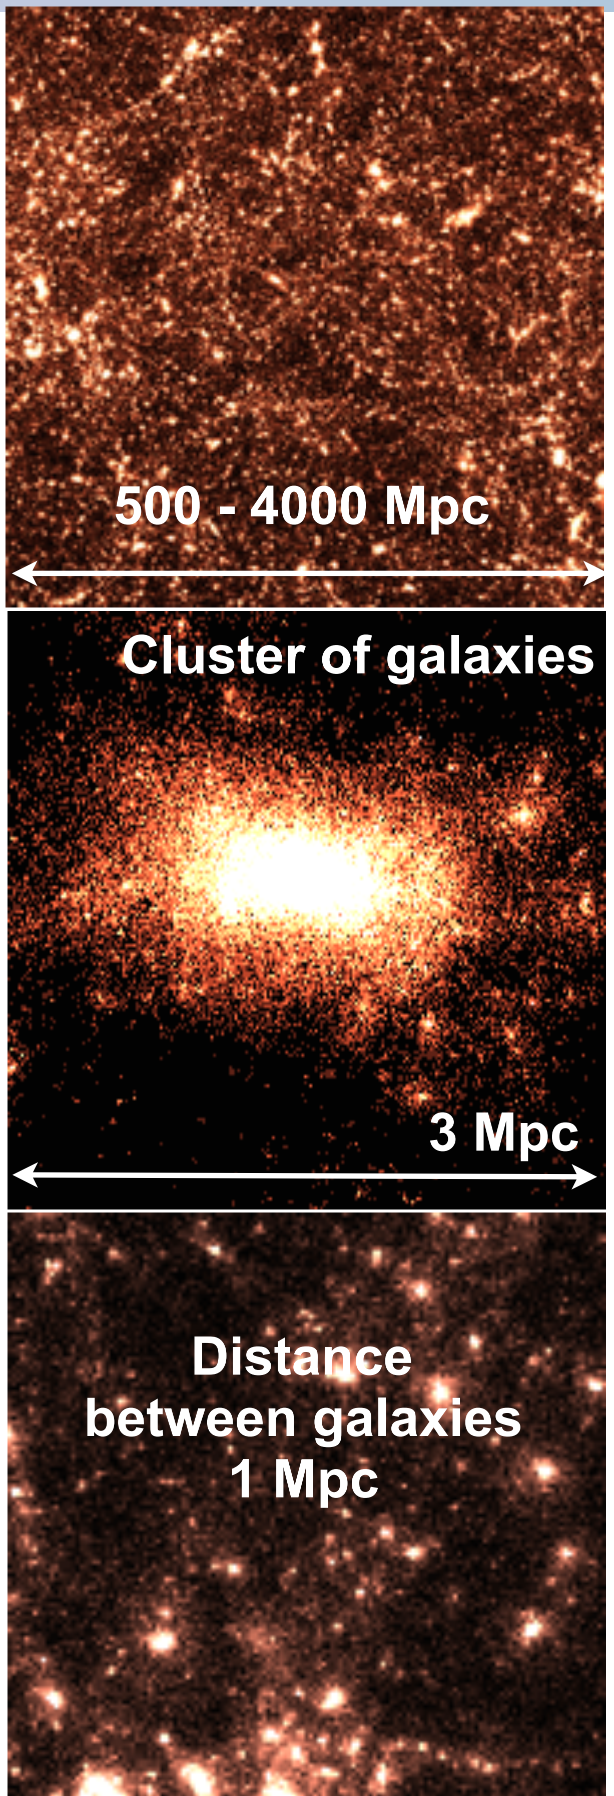

Relevant Numbers for Optical Surveys

- Survey volume
  - Surveys cover enormous observational volumes, ~(4 Gpc)^3, of order 12 billion light-years in linear scale
- Number of Galaxies
  - Number of galaxies in surveys can be in the tens of billions, typical intergalactic separation is ~1 Mpc
- Typical galaxy mass
  - The mass associated with a typical bright galaxy in surveys ranges from 10^{11}-10^{13} $M_\odot$
- Typical galaxy size
  - The length scale of the mass distribution is approximately 100 kpc

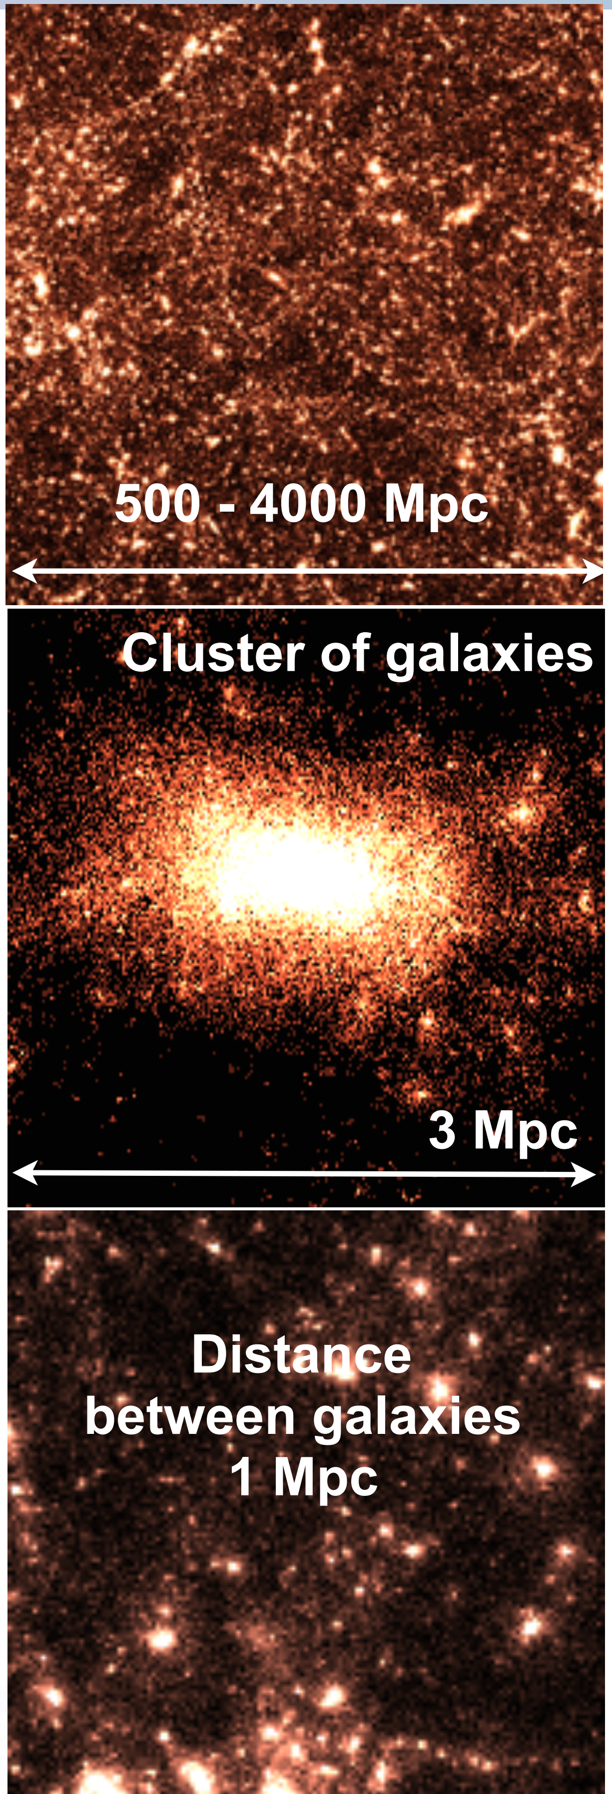

***1 parsec = 3.26 light-years = 31 trillion kilometers**

The Content of the Universe: It's dark!

- **Dark Energy:** Multiple observations show that the expansion of the Universe is accelerating (first measured in 1998, Nobel prize 2011)
- Imagine you throw a ball in the air and instead of coming down it flies upwards faster and faster!
- 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
- **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

Observer's

Theorist's Universe

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

The Evolution of the Universe: Structure Formation

- **Solid understanding of structure formation; success underpins most cosmic discovery**
  - Initial conditions determined by primordial fluctuations, measured from the cosmic microwave background
  - 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
- **Latter half of the history of the Universe:** Nonlinear domain of structure formation, impossible to treat without large-scale computing

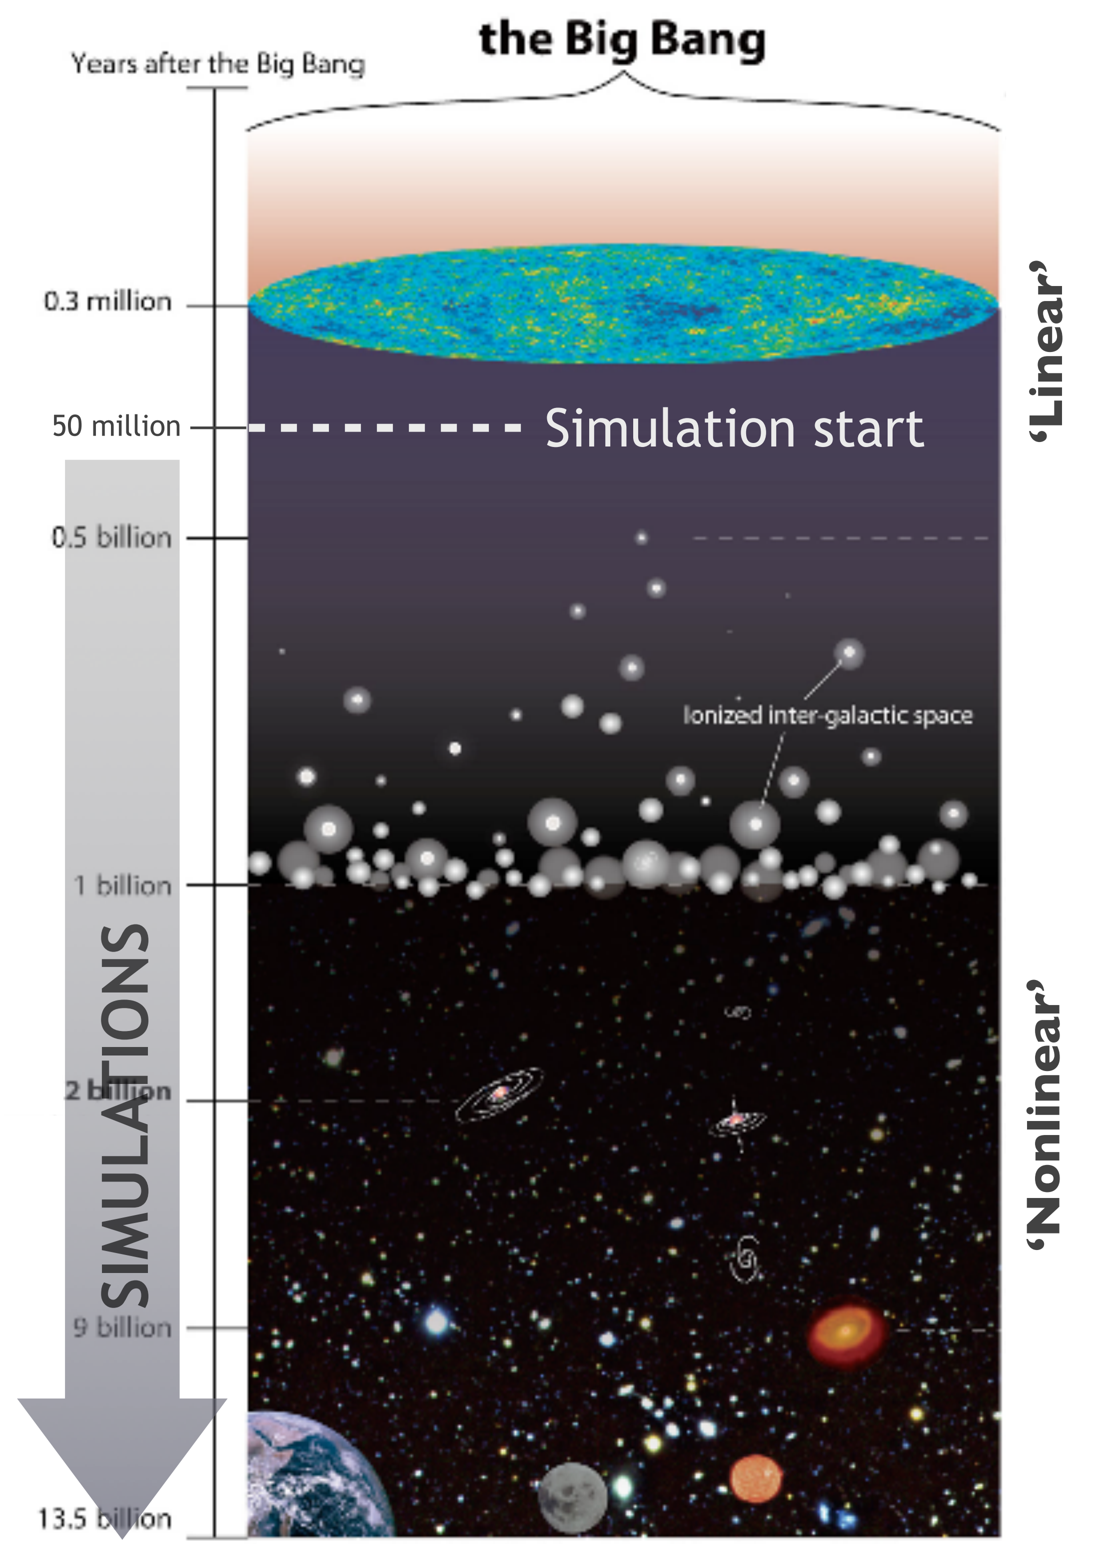

Computing the Universe

- Gravity dominates on large scales, use Monte Carlo sampling of density with tracer particles
- Particles are tracers of the dark matter in the Universe, mass typically at least ~10^9 M☉

$$m_p \sim V/n_p$$

- Simulate galaxy size objects ($v^2/c^2 << 1$), Newtonian description accurate
- At smaller scales, add gas physics, feedback etc., sub-grid modeling inevitable

> "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)*

$$\ddot{\mathbf{x}} + 2\frac{\dot{a}}{a}\dot{\mathbf{x}} = -\frac{\nabla\Phi}{a^2}$$ Equation of motion for tracer particles in expanding Universe

$$\frac{\dot{a}}{a} = H = \frac{H_0}{a^{3/2}}\sqrt{\Omega_{tot} + a^3\Omega_\Lambda}$$ CDM + baryons + DE

$$\nabla^2\Phi(\mathbf{x}) = 4\pi G a^2[\rho(\mathbf{x},t) - \rho_b(t)]$$ Poisson equation

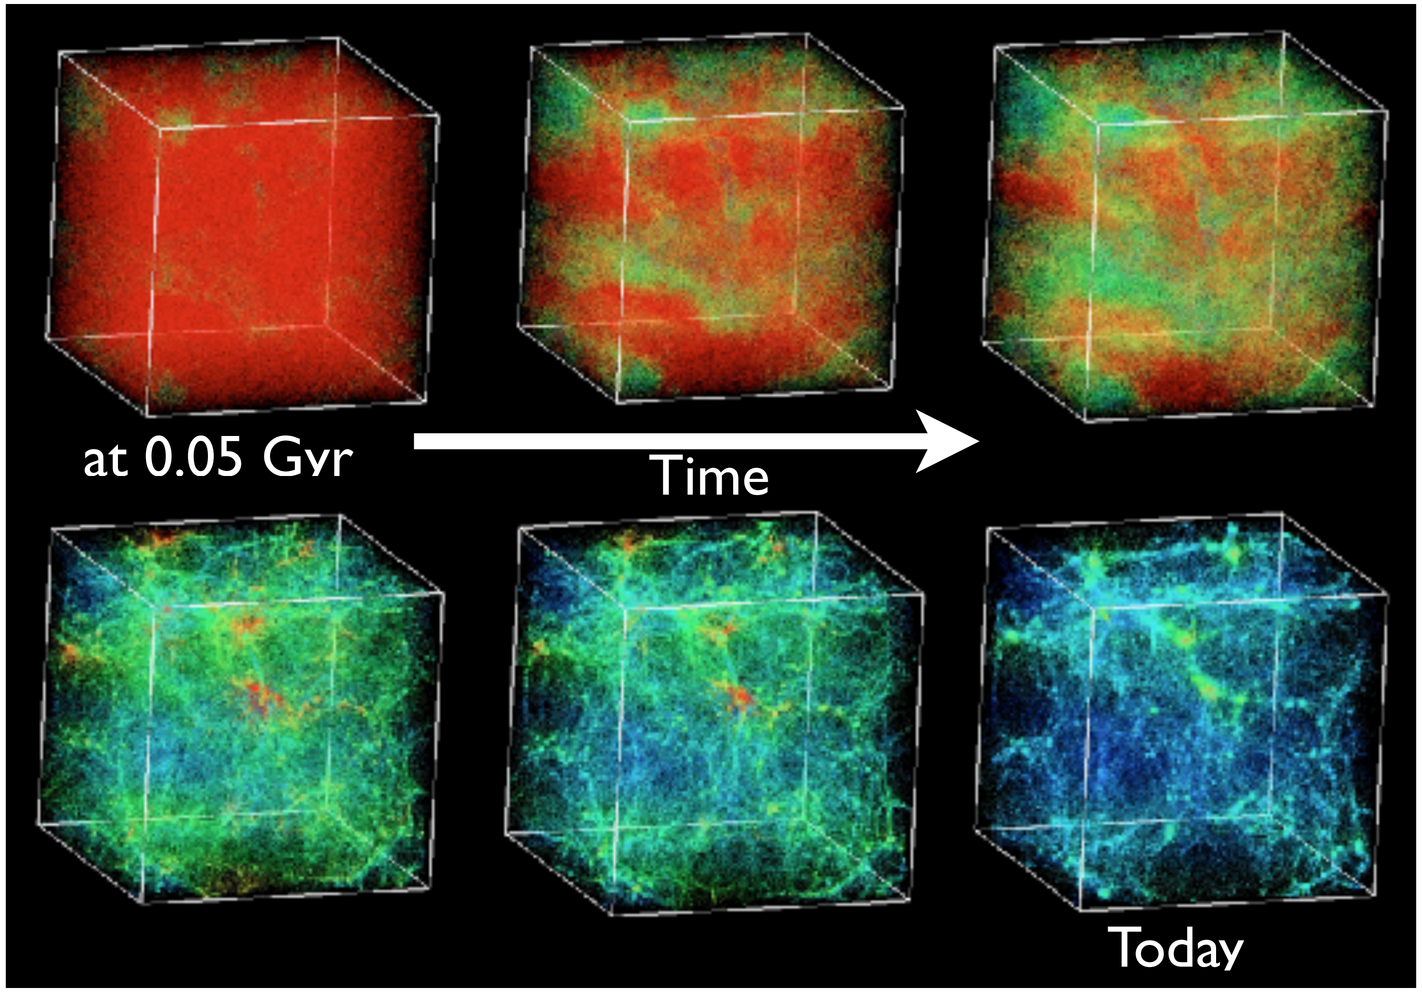

The Ingredients to Create your Own Universe

- Constituents
- Initial conditions
- Dynamical rules
- Computer (big!)

Structure Formation in the Universe: The Delta Quadrant Simulation

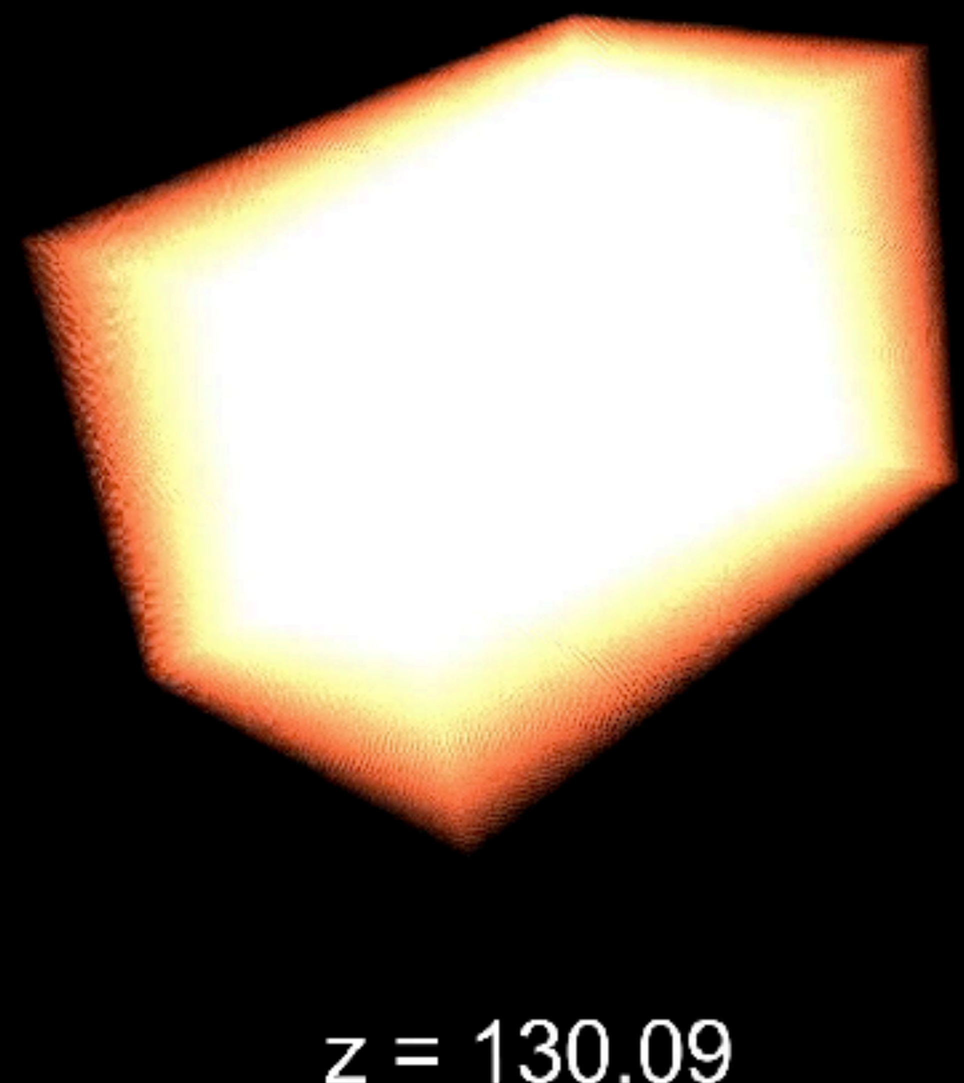

$(500\ \mathrm{Mpc}/h)^3$ box, 3072^3 particles, saved 750 snapshots, ~800TB of data,
Movie from one rank (out of 1152 ranks), ~(41Mpc/h x 41Mpc/h x 62.5Mpc/h)

Structure Formation in the Universe: The Delta Quadrant Simulation

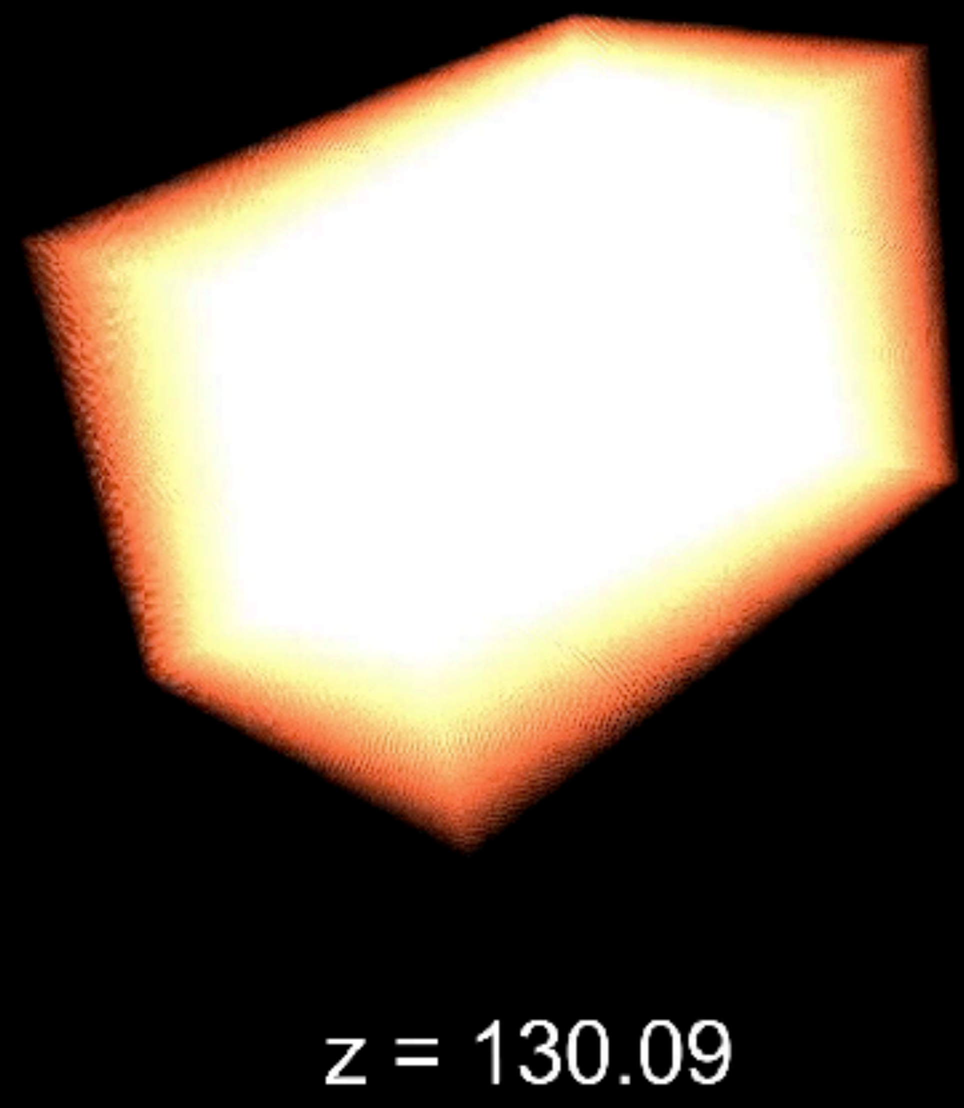

$(500\ \mathrm{Mpc}/h)^3$ box, 3072^3 particles, saved 750 snapshots, ~800TB of data,
Movie from one rank (out of 1152 ranks), ~(41Mpc/h x 41Mpc/h x 62.5Mpc/h)

Changing the Constituents

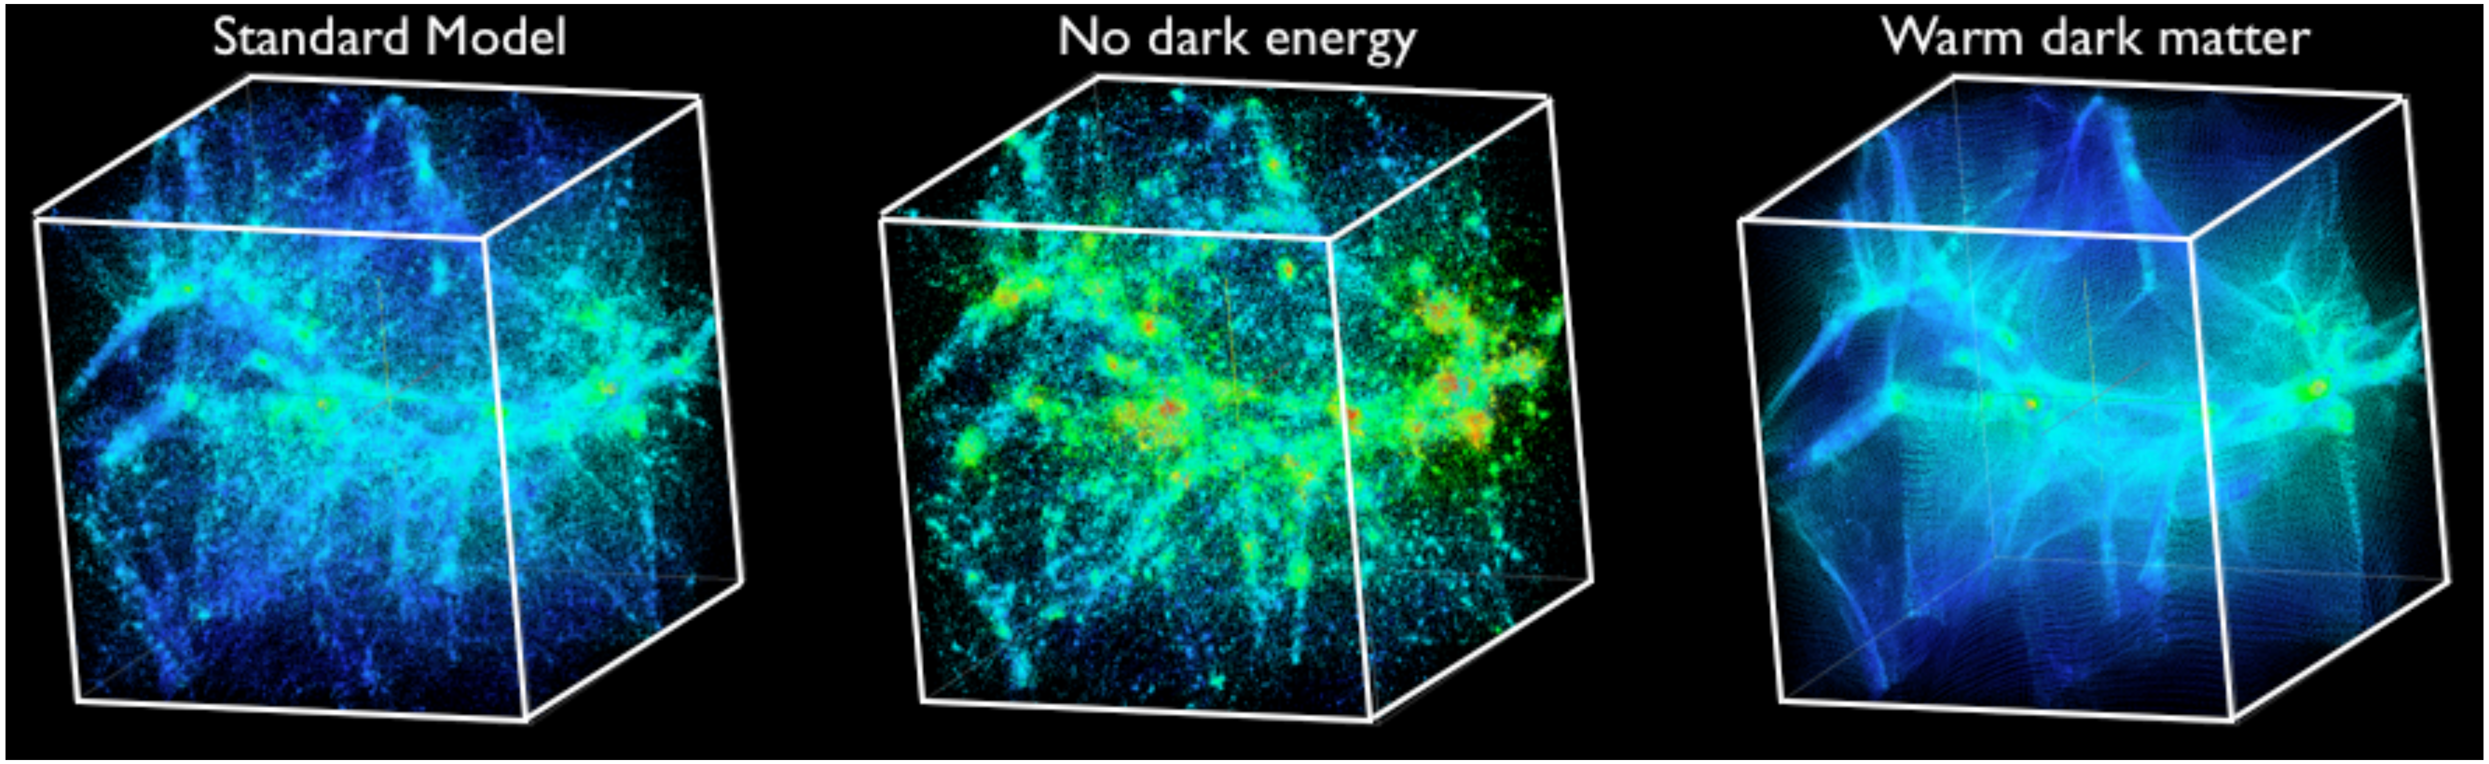

- **Exploration of different dark energy models:** How would a dark energy model beyond Einstein's cosmological constant alter the distribution of matter (and galaxies) in the Universe?
- **Exploration of dark matter and neutrinos in the Universe:** What can cosmology tell us about different matter components in the Universe?

Connecting Theory and Observations

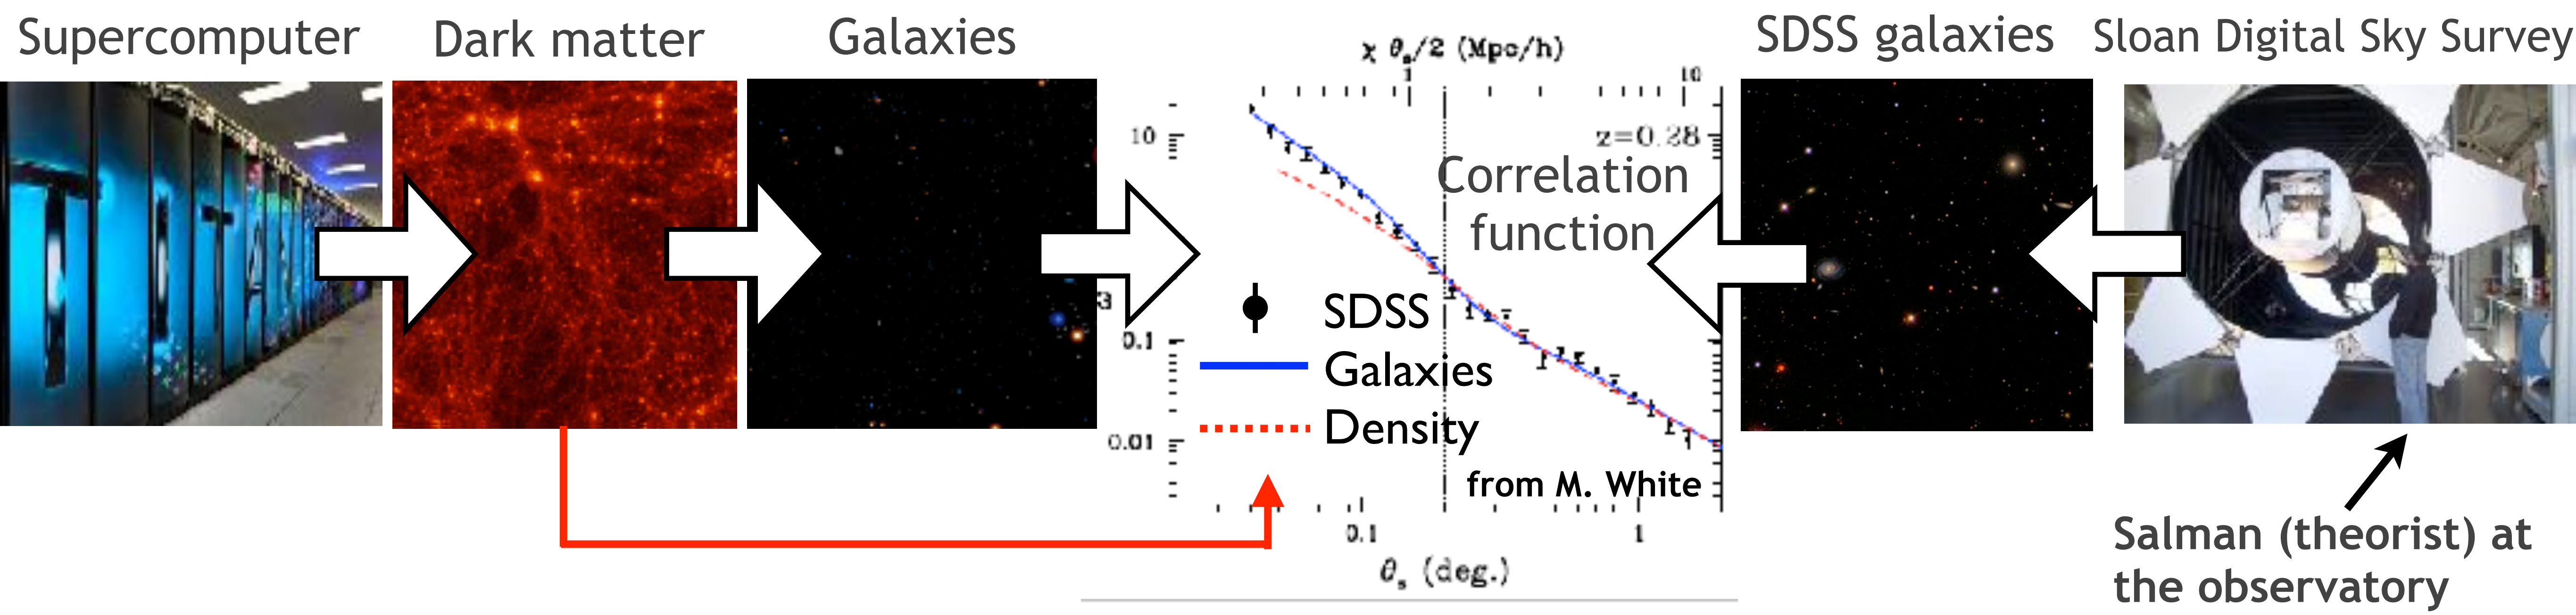

- Simulate the formation of the large scale structure of the Universe via dark matter tracer particles, taking dark energy into account in expansion history
- Measure the high-density peaks (dark matter halos) in the mass distribution
- "Light traces mass" to first approximation, so populate the halos 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

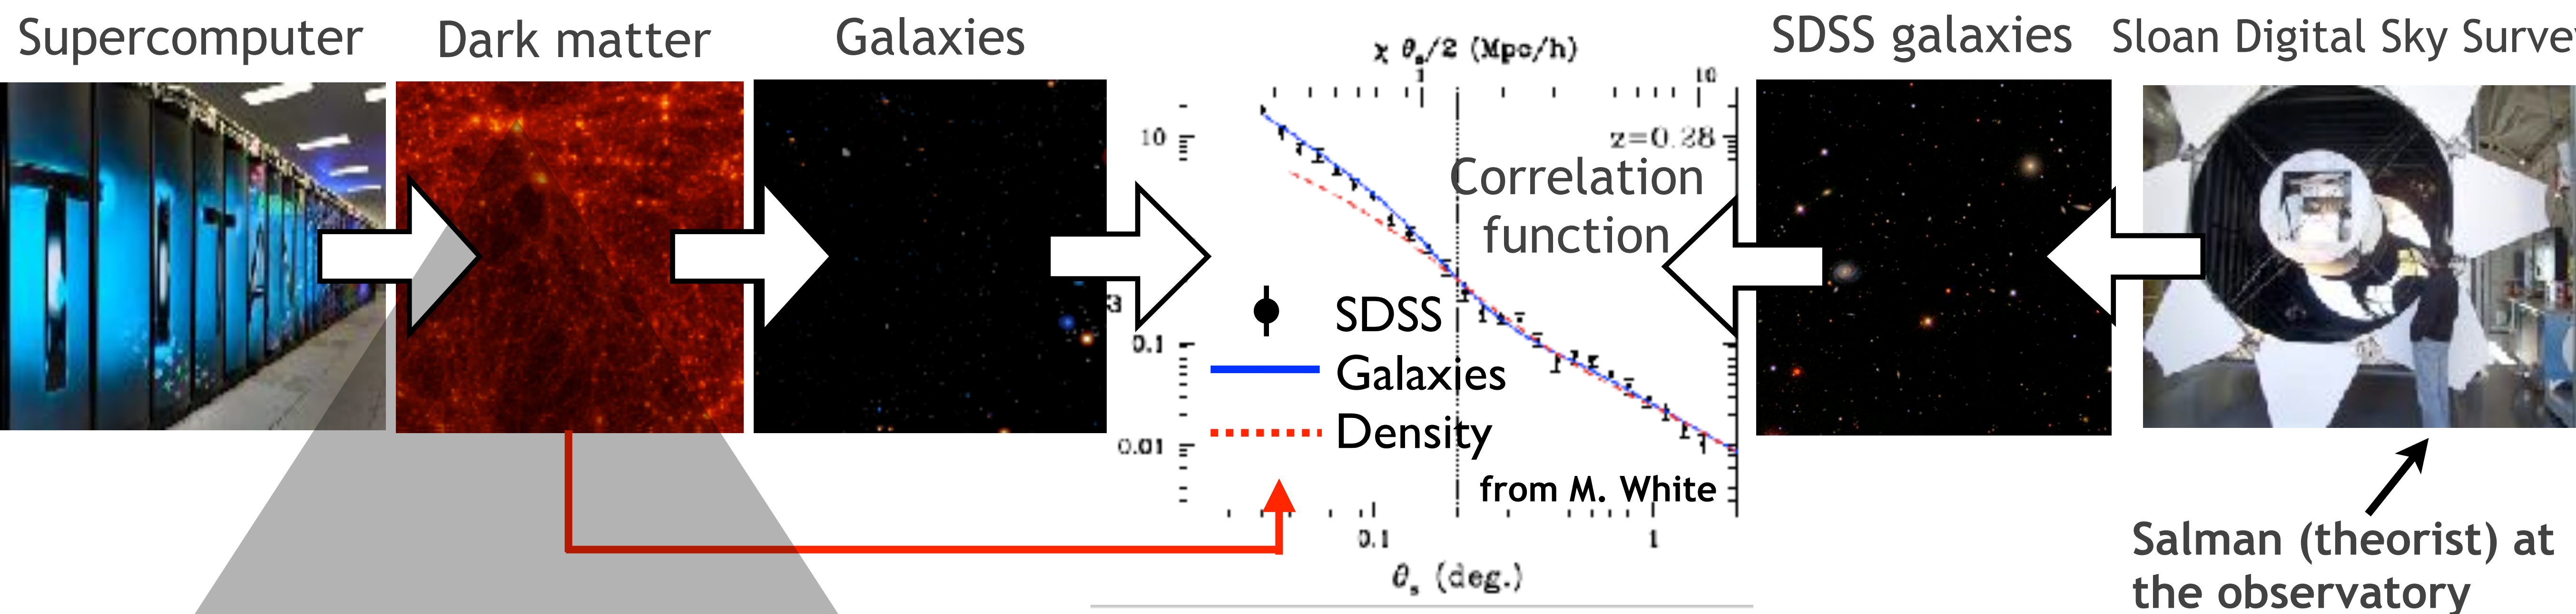

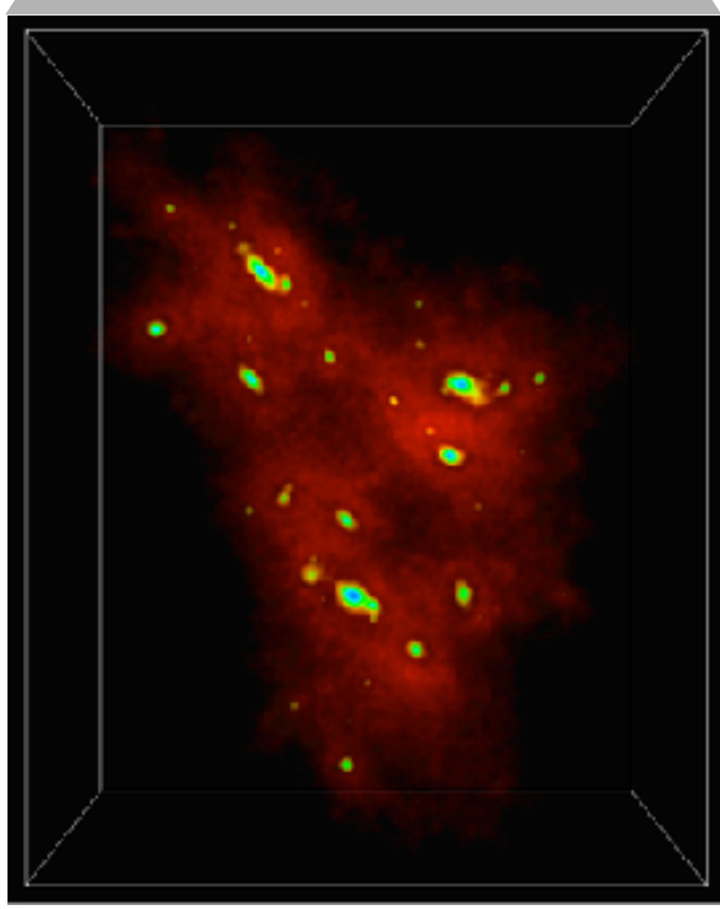

ation of the large scale structure of the Universe via dark
icles, taking dark energy into account in expansion history

density peaks (dark matter halos) in the mass distribution

s" to first approximation, so populate the halos with galaxies,
s depends on mass of halo (constraints from observations)

prescription (hopefully) independent of cosmological model

: Error bars on measurements are shrinking, surveys cover
resolve fainter and fainter galaxies

Connecting Theory and Observations

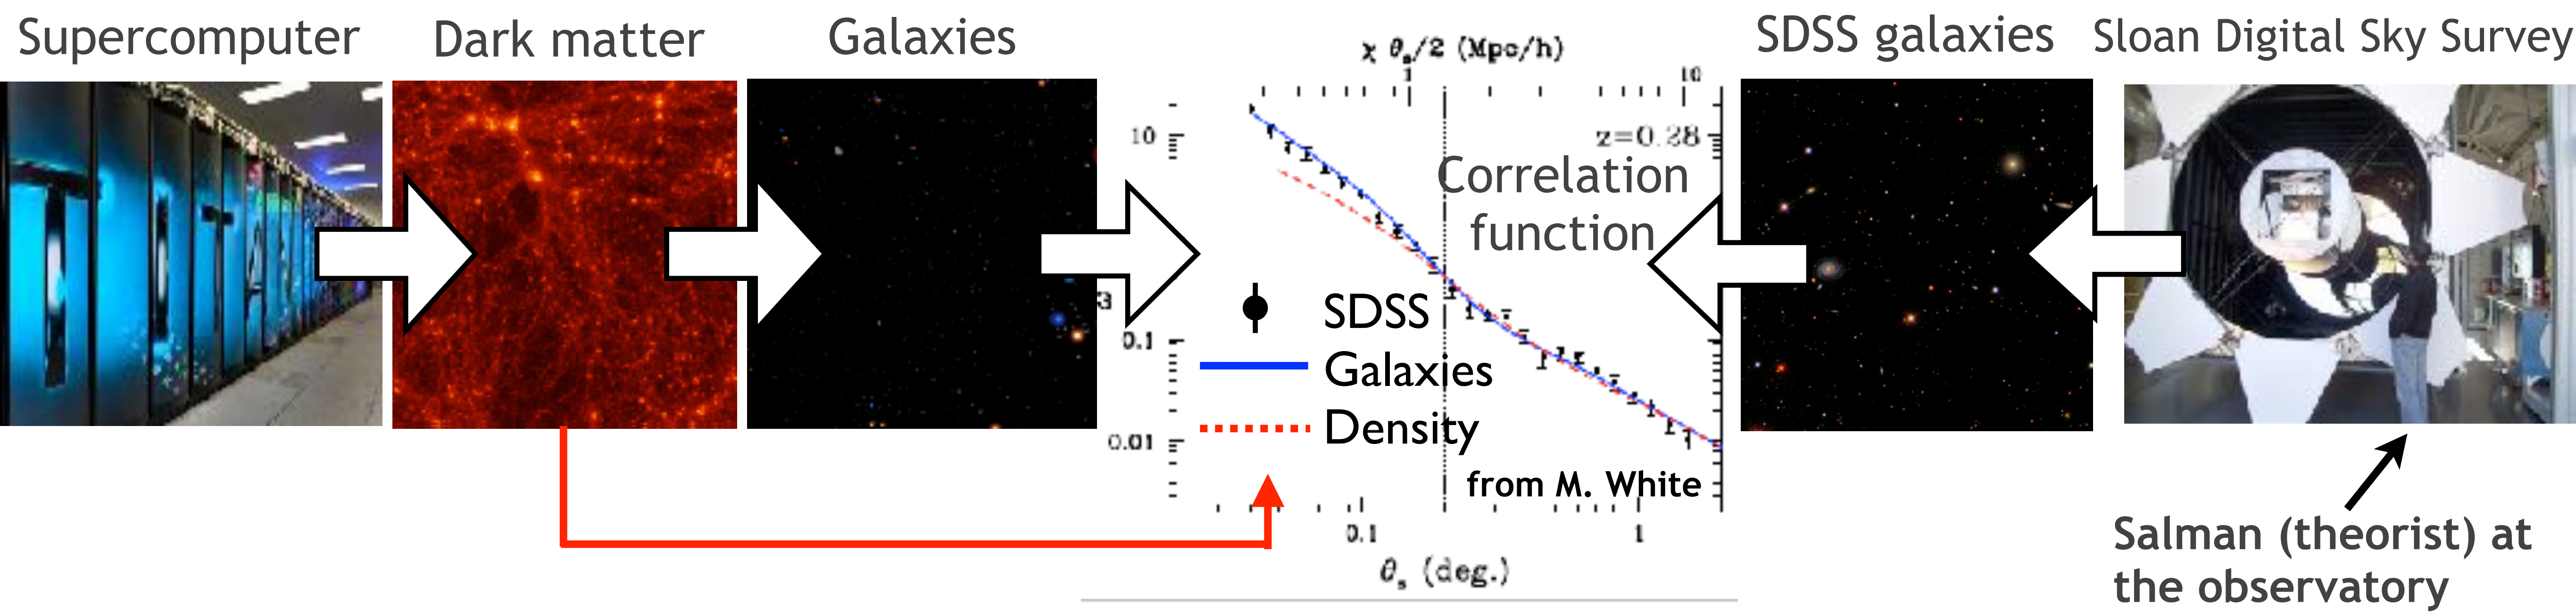

- Simulate the formation of the large scale structure of the Universe via dark matter tracer particles, taking dark energy into account in expansion history
- Measure the high-density peaks (dark matter halos) in the mass distribution
- “Light traces mass” to first approximation, so populate the halos 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

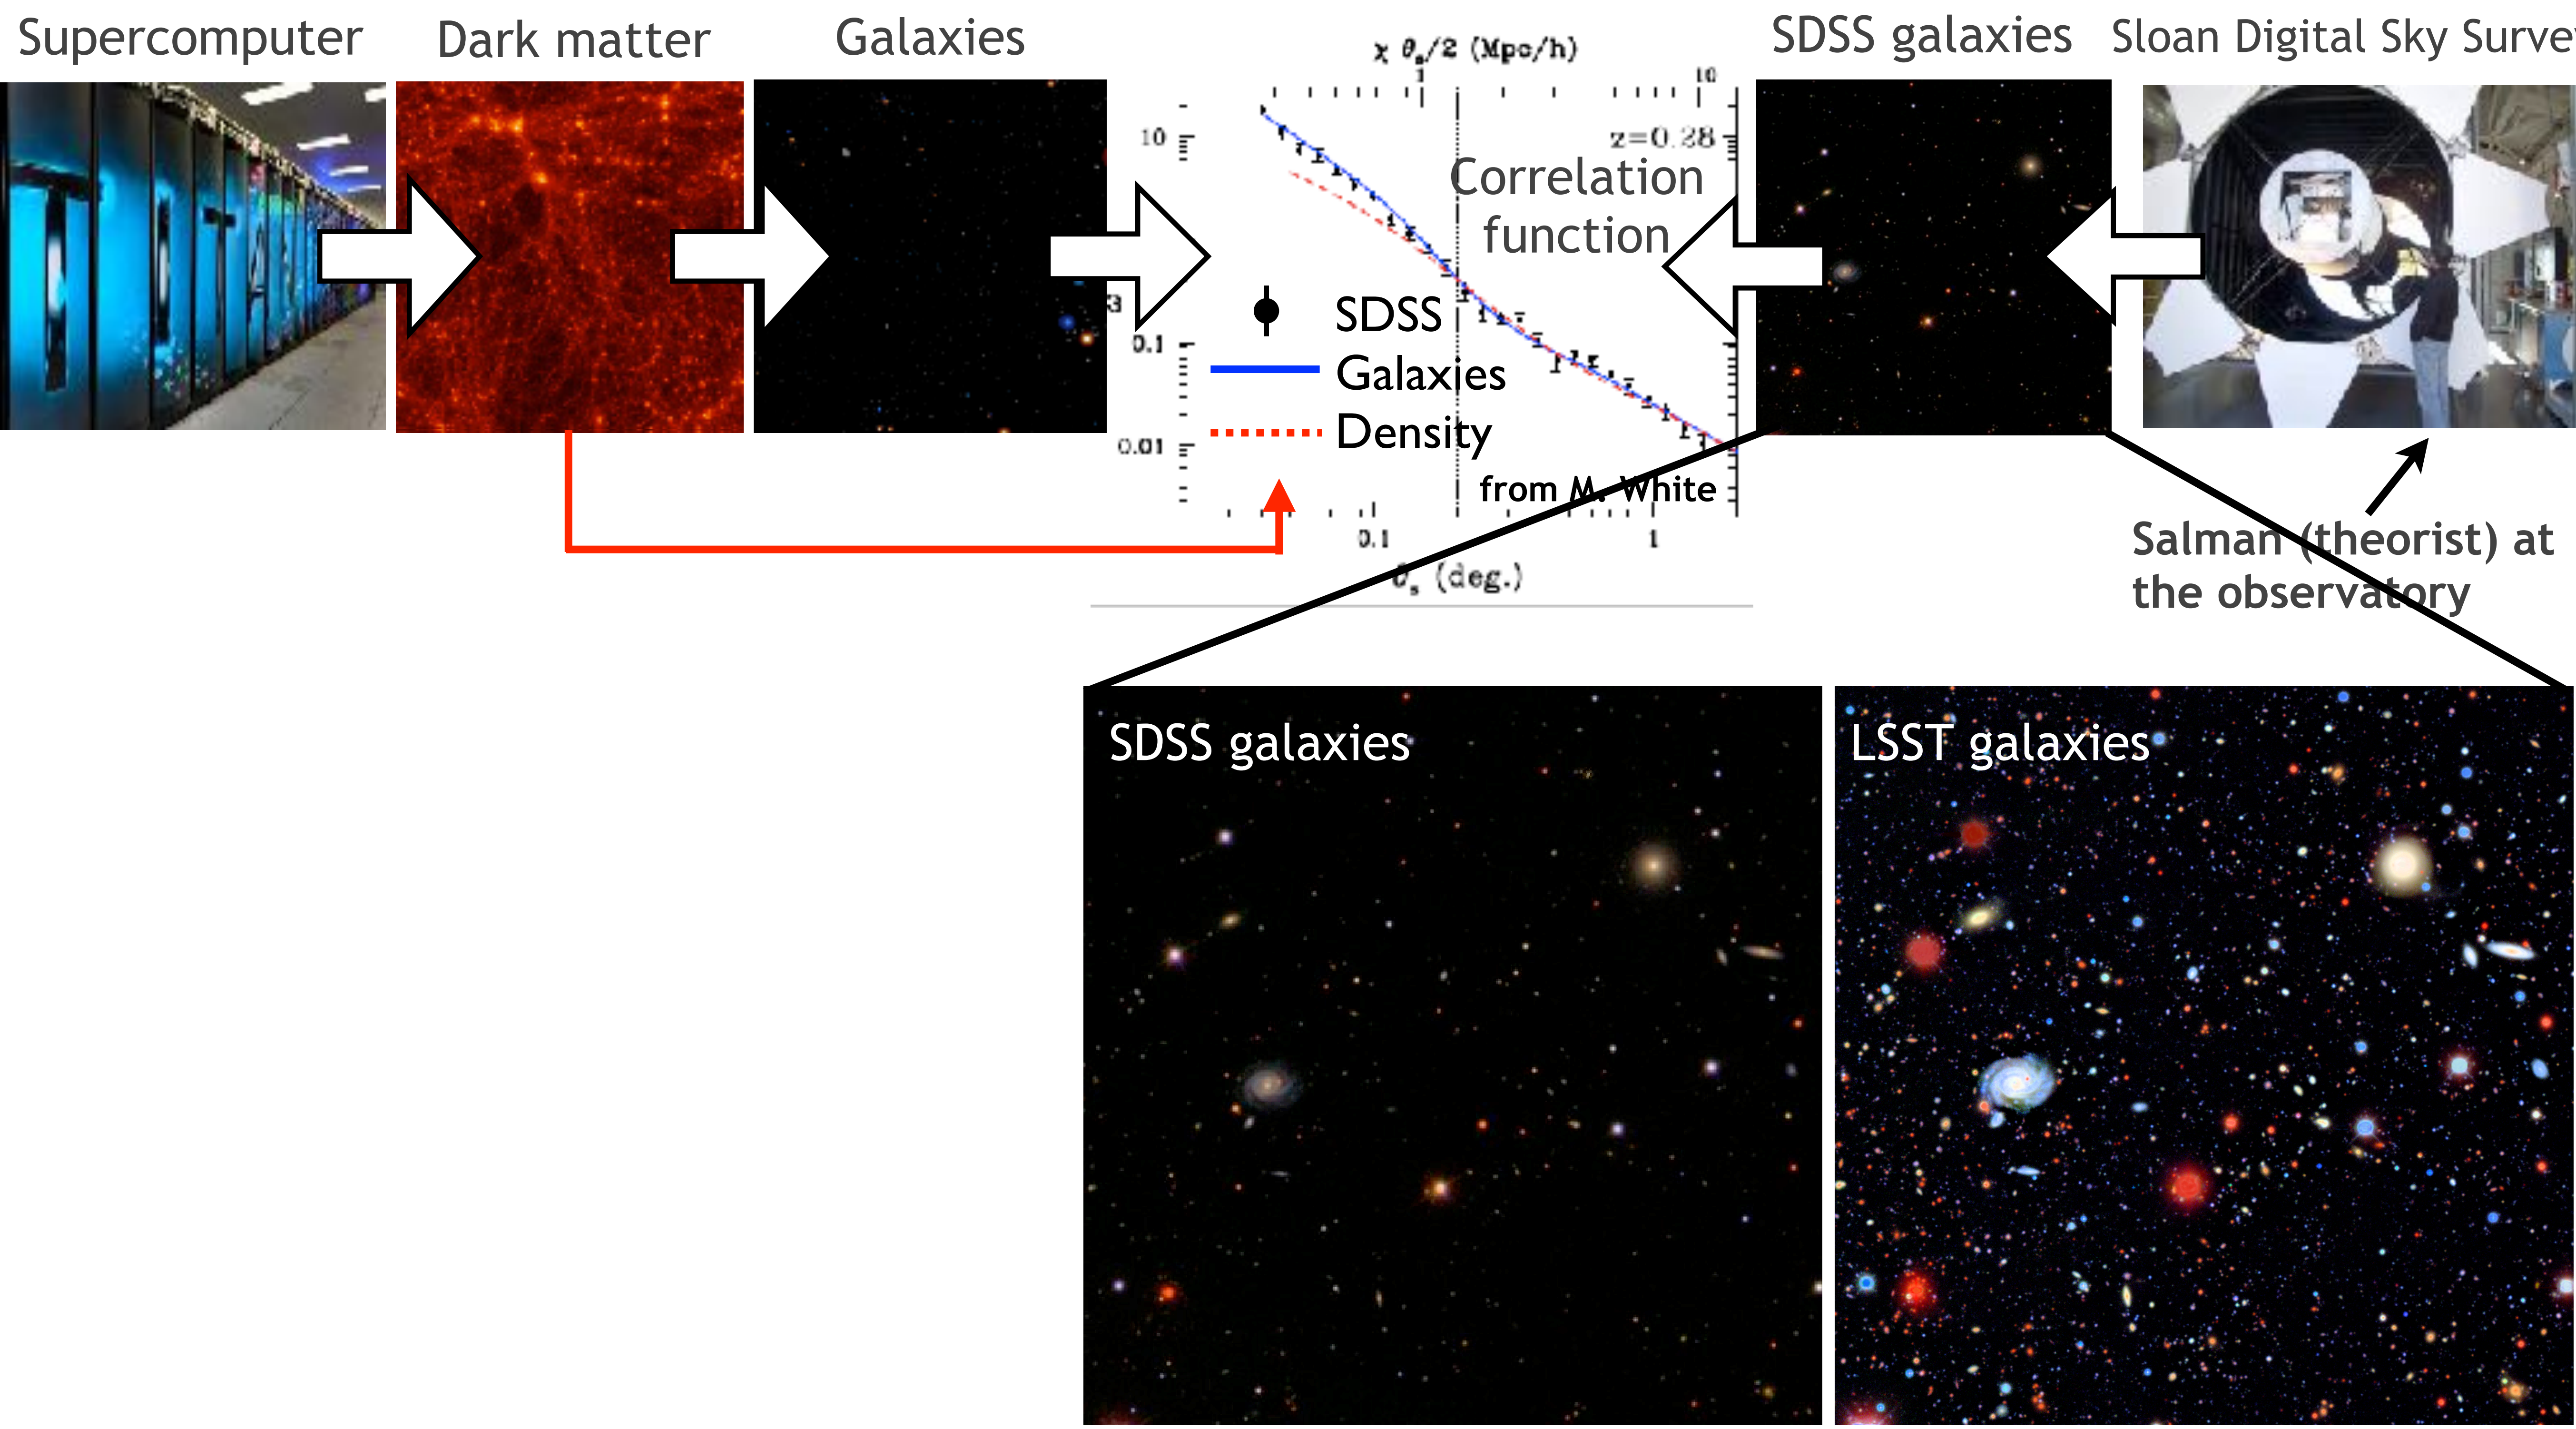

Connecting Theory and Observations

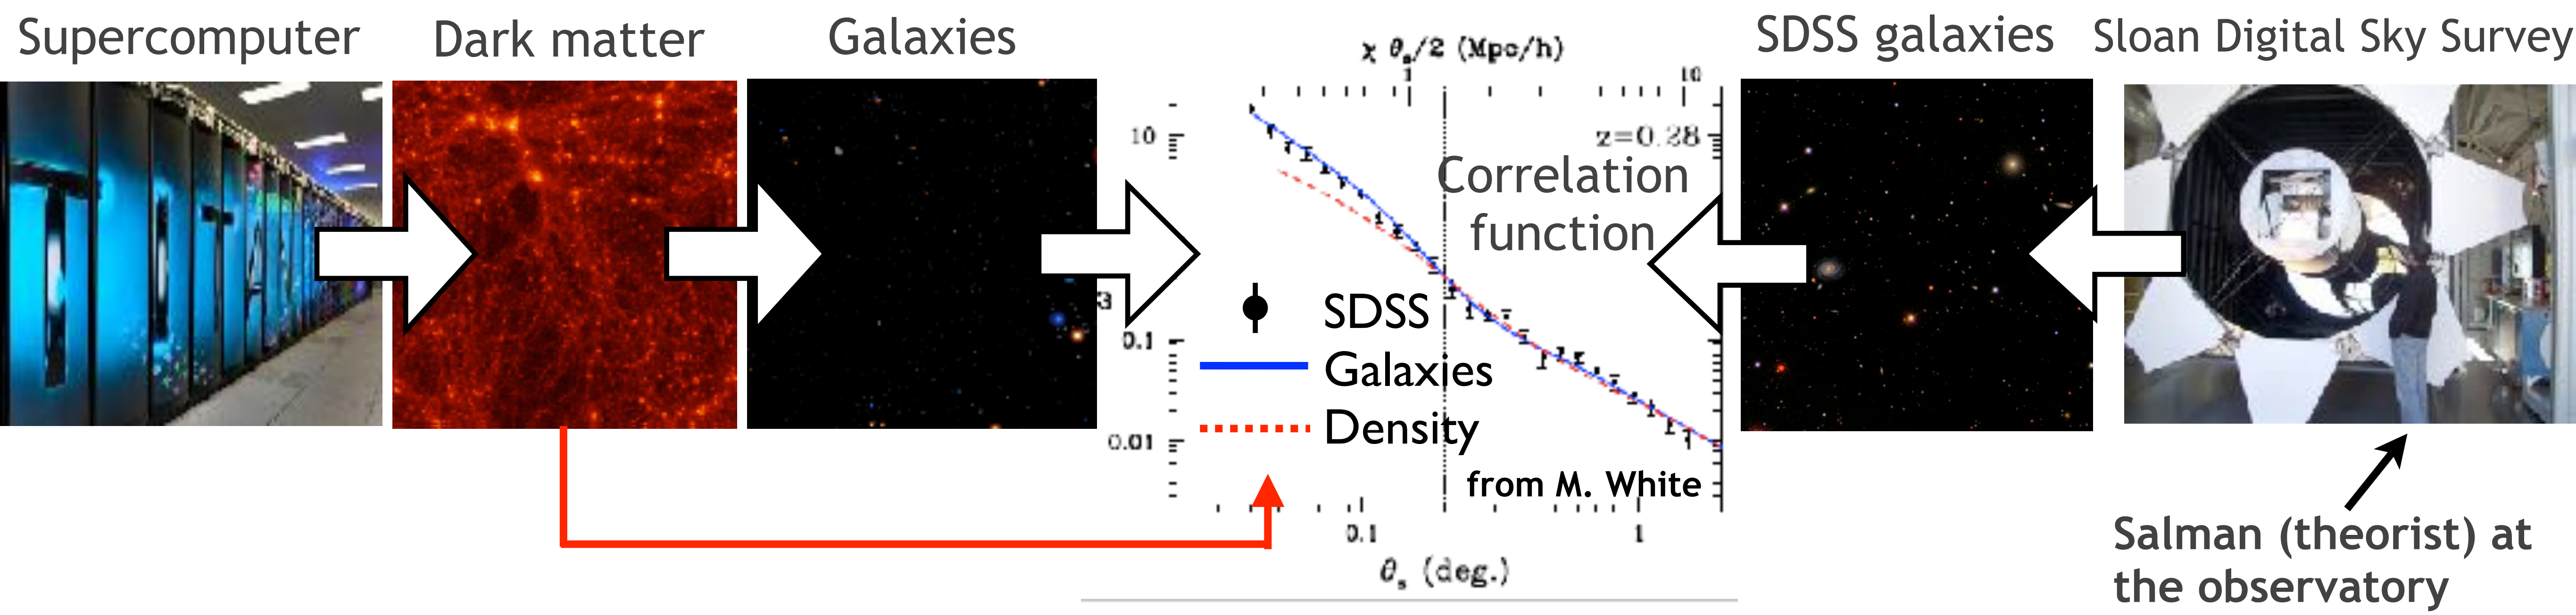

- Simulate the formation of the large scale structure of the Universe via dark matter tracer particles, taking dark energy into account in expansion history
- Measure the high-density peaks (dark matter halos) in the mass distribution
- “Light traces mass” to first approximation, so populate the halos 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

The HACC Story

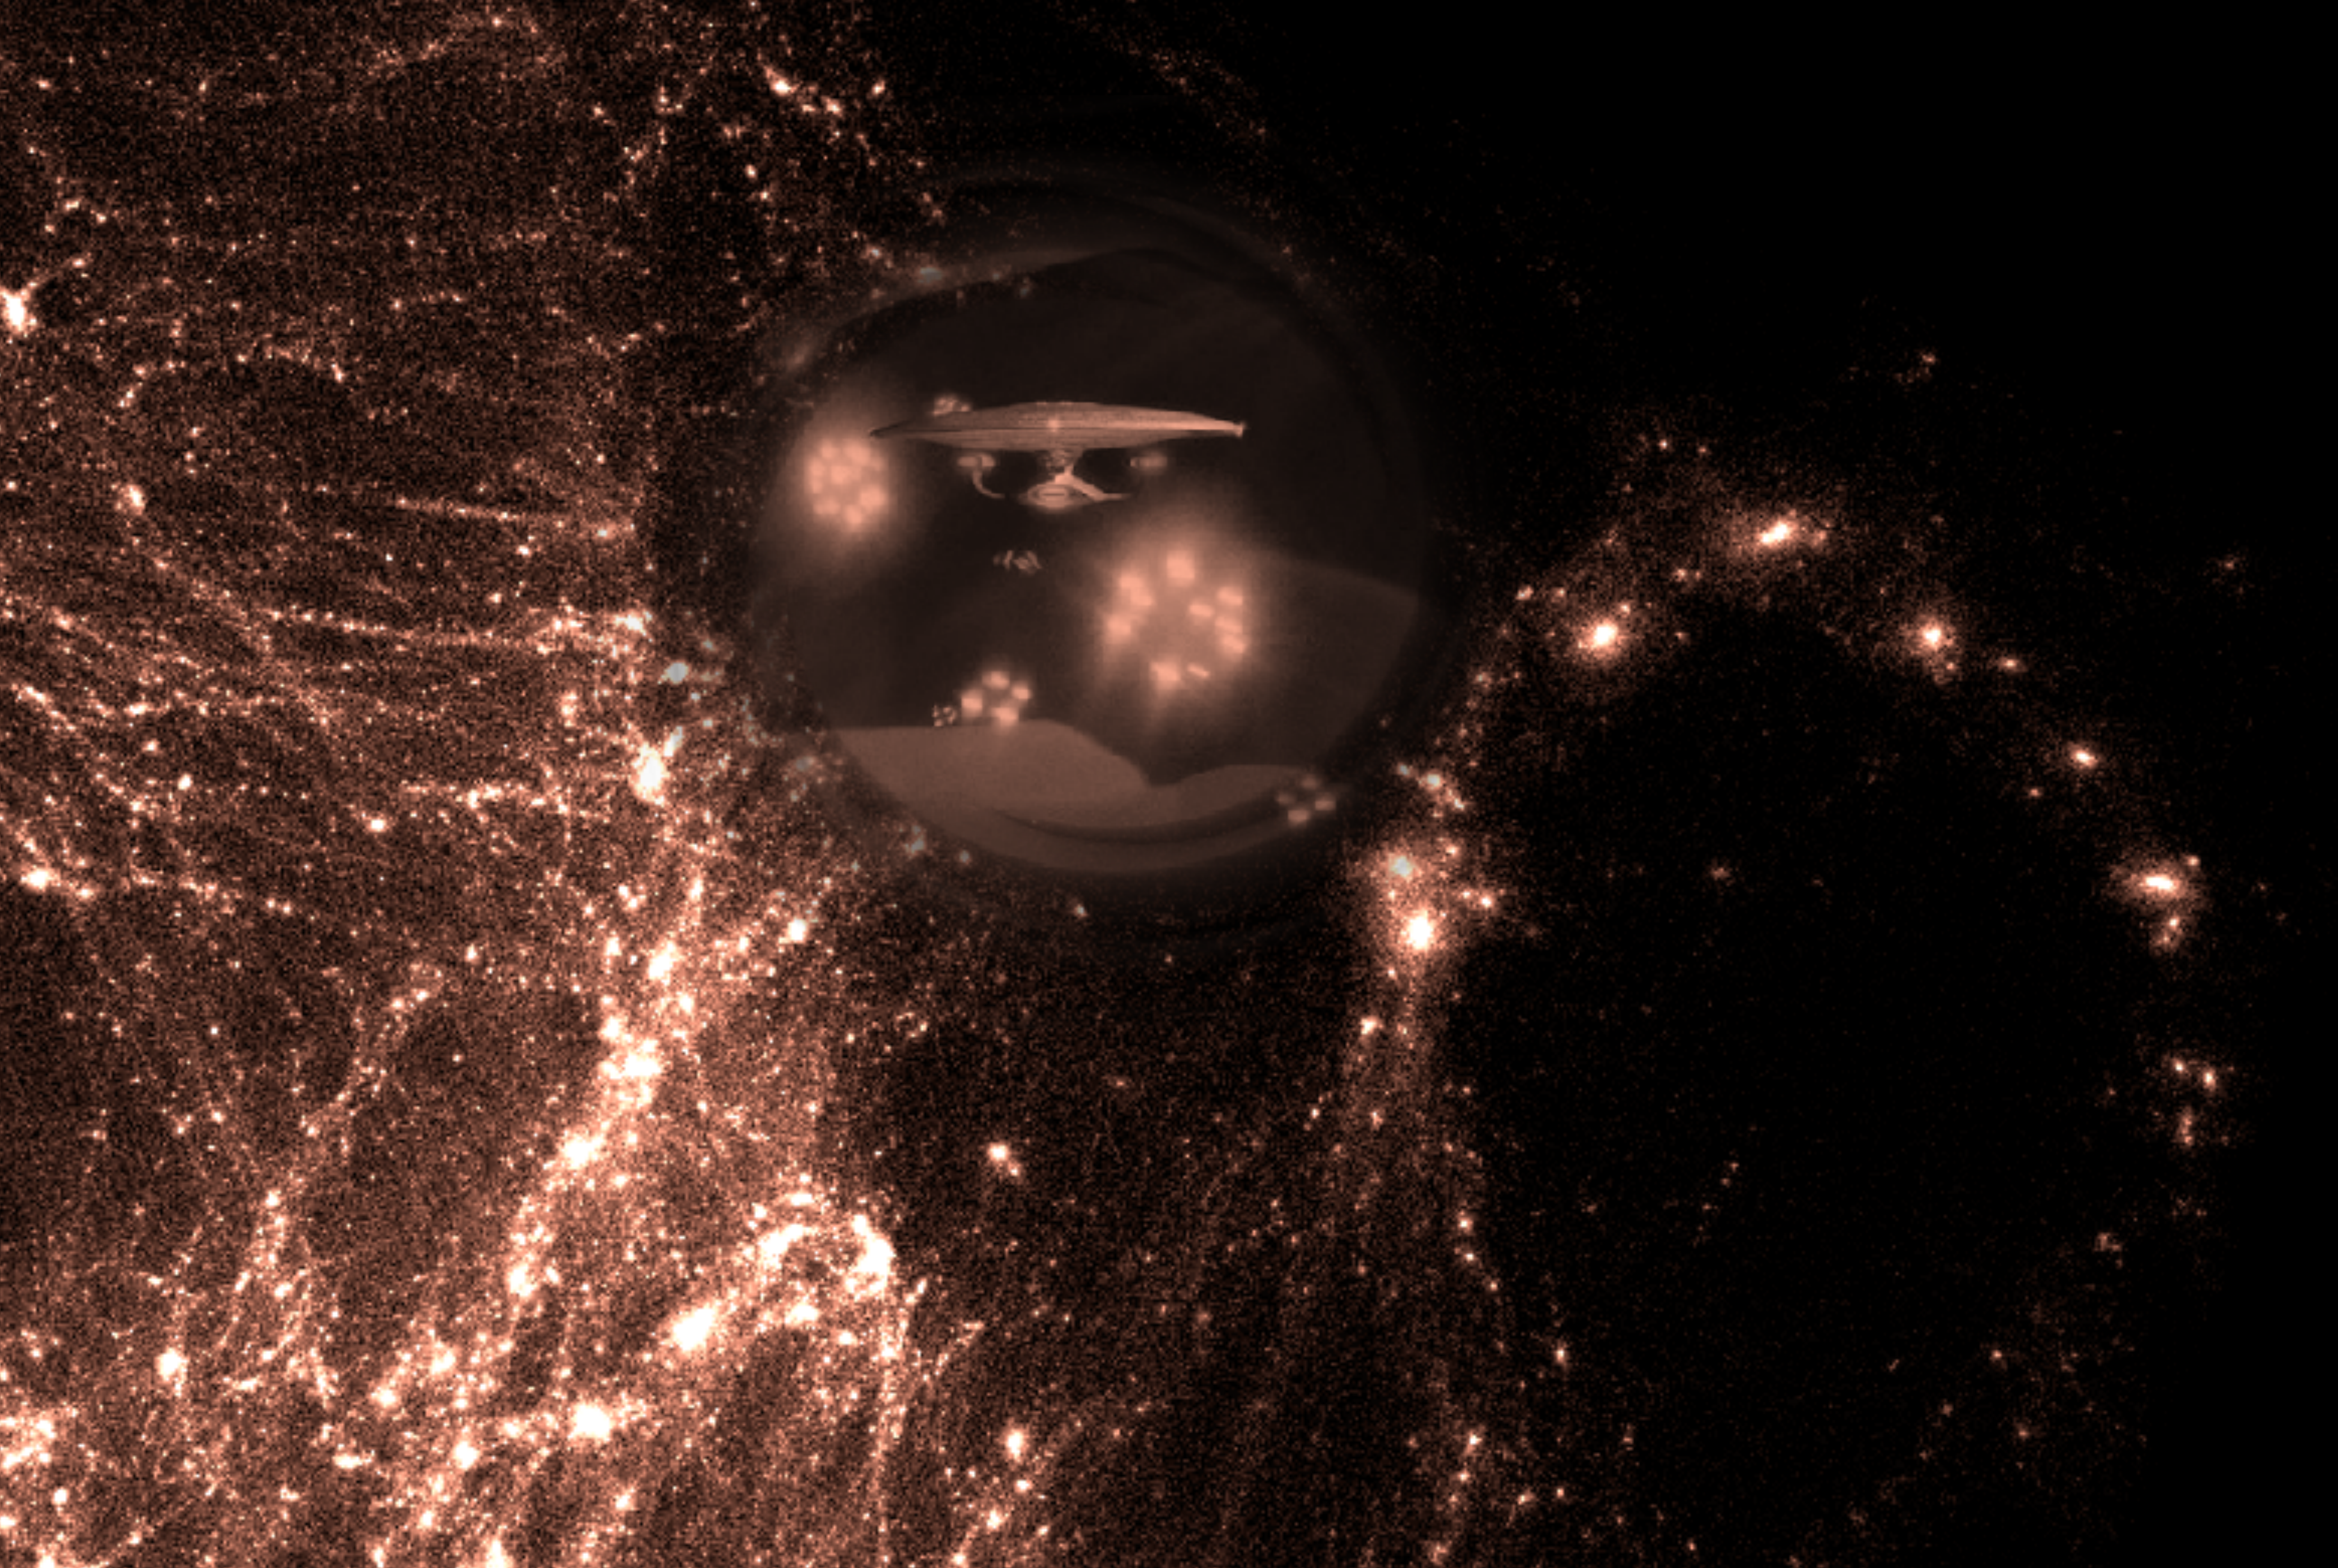

The HACC Story

Chapter 1 — The Southwest

The HACC Story

Chapter 1 — The Southwest

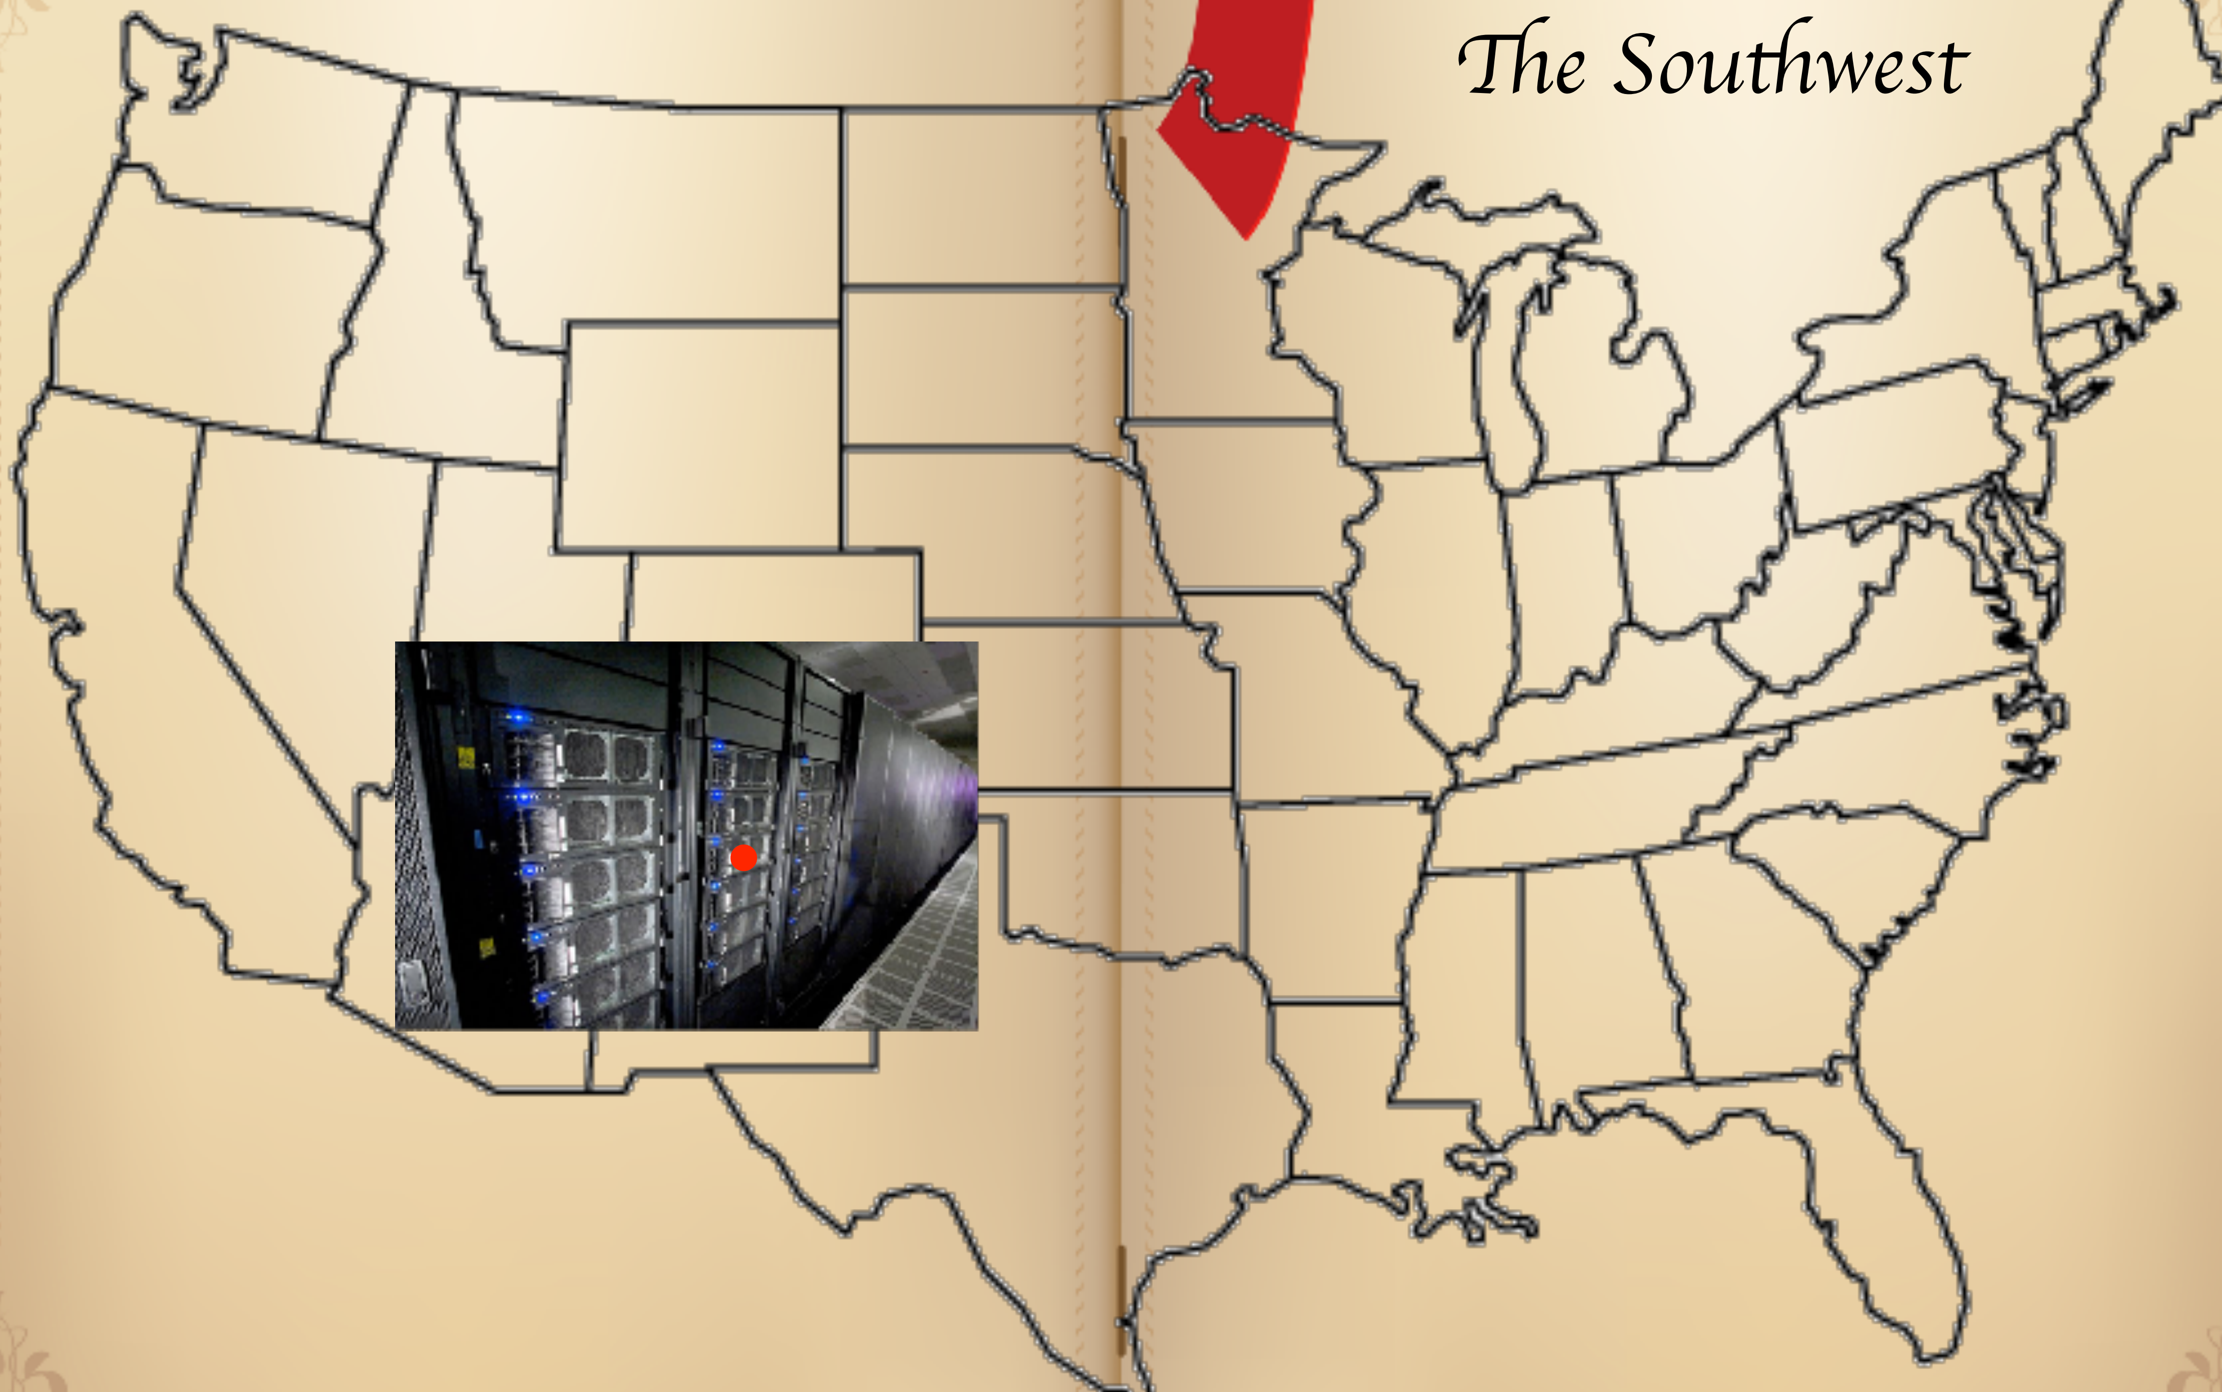

The HACC Story

Chapter 1 — The Southwest

The 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:** **MC3 (Mesh-based Cosmology Code on the Cell, based on MC2, a PM code written in High-Performance Fortran) which later morphed into HACC, N-body code to simulate large-scale structure formation in the Universe**

The 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, based on MC2, a PM code written in High-Performance Fortran) which later morphed into HACC, N-body code to simulate large-scale structure formation in the Universe**

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"

So we started thinking --

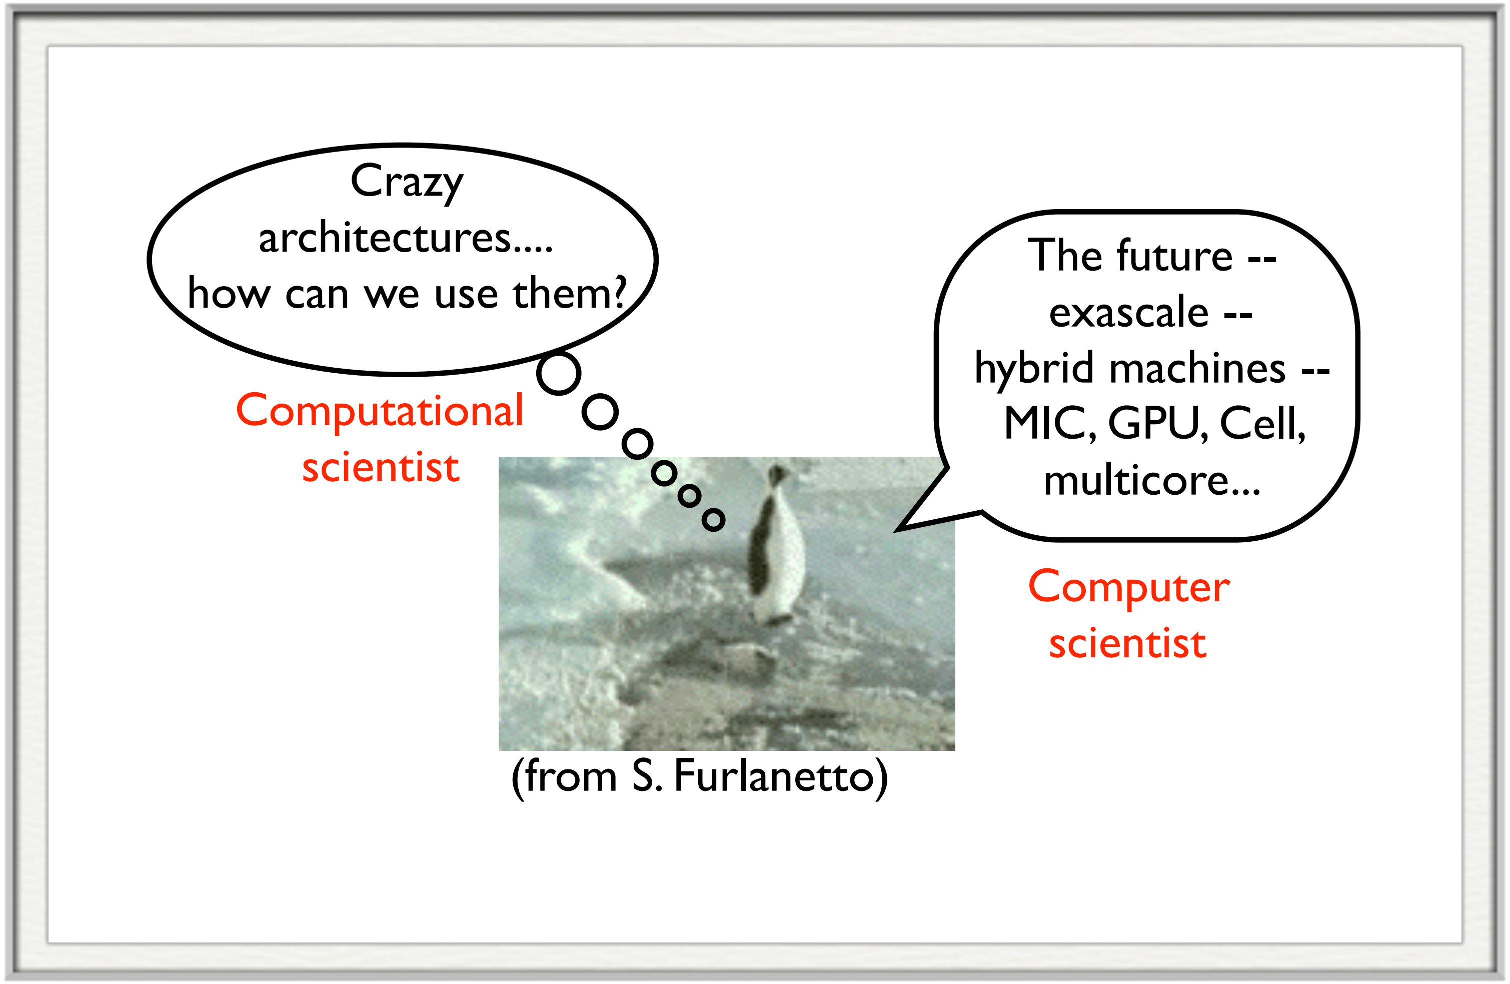

(from S. Furlanetto)

So we started thinking --

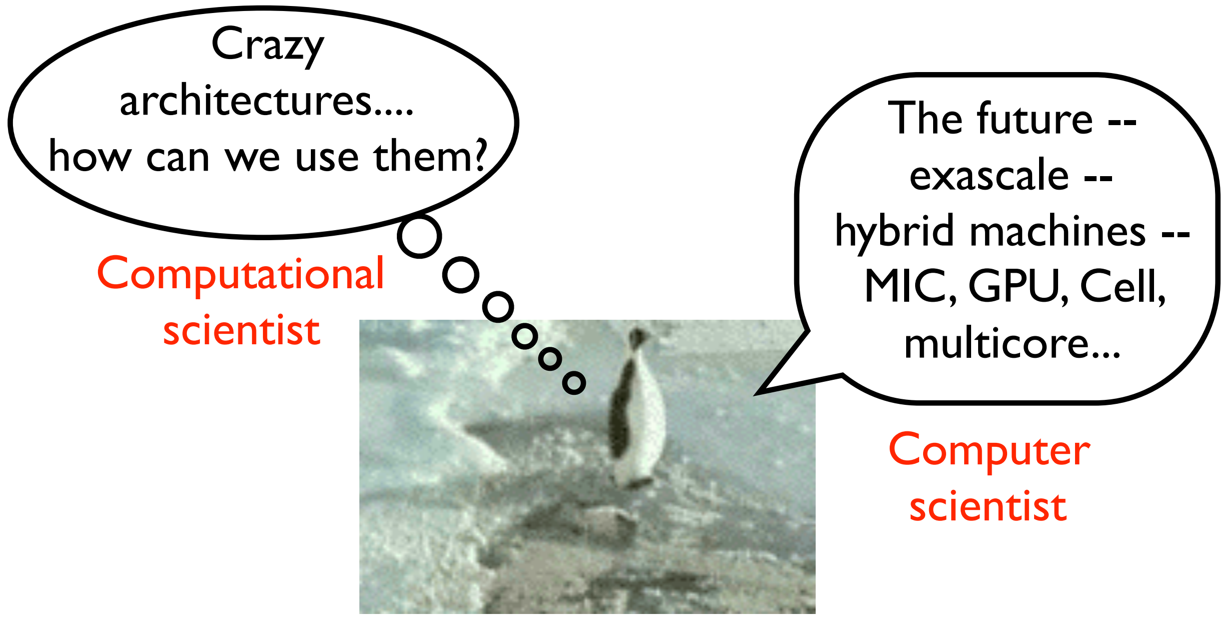

(from S. Furlanetto)

The Roadrunner Architecture

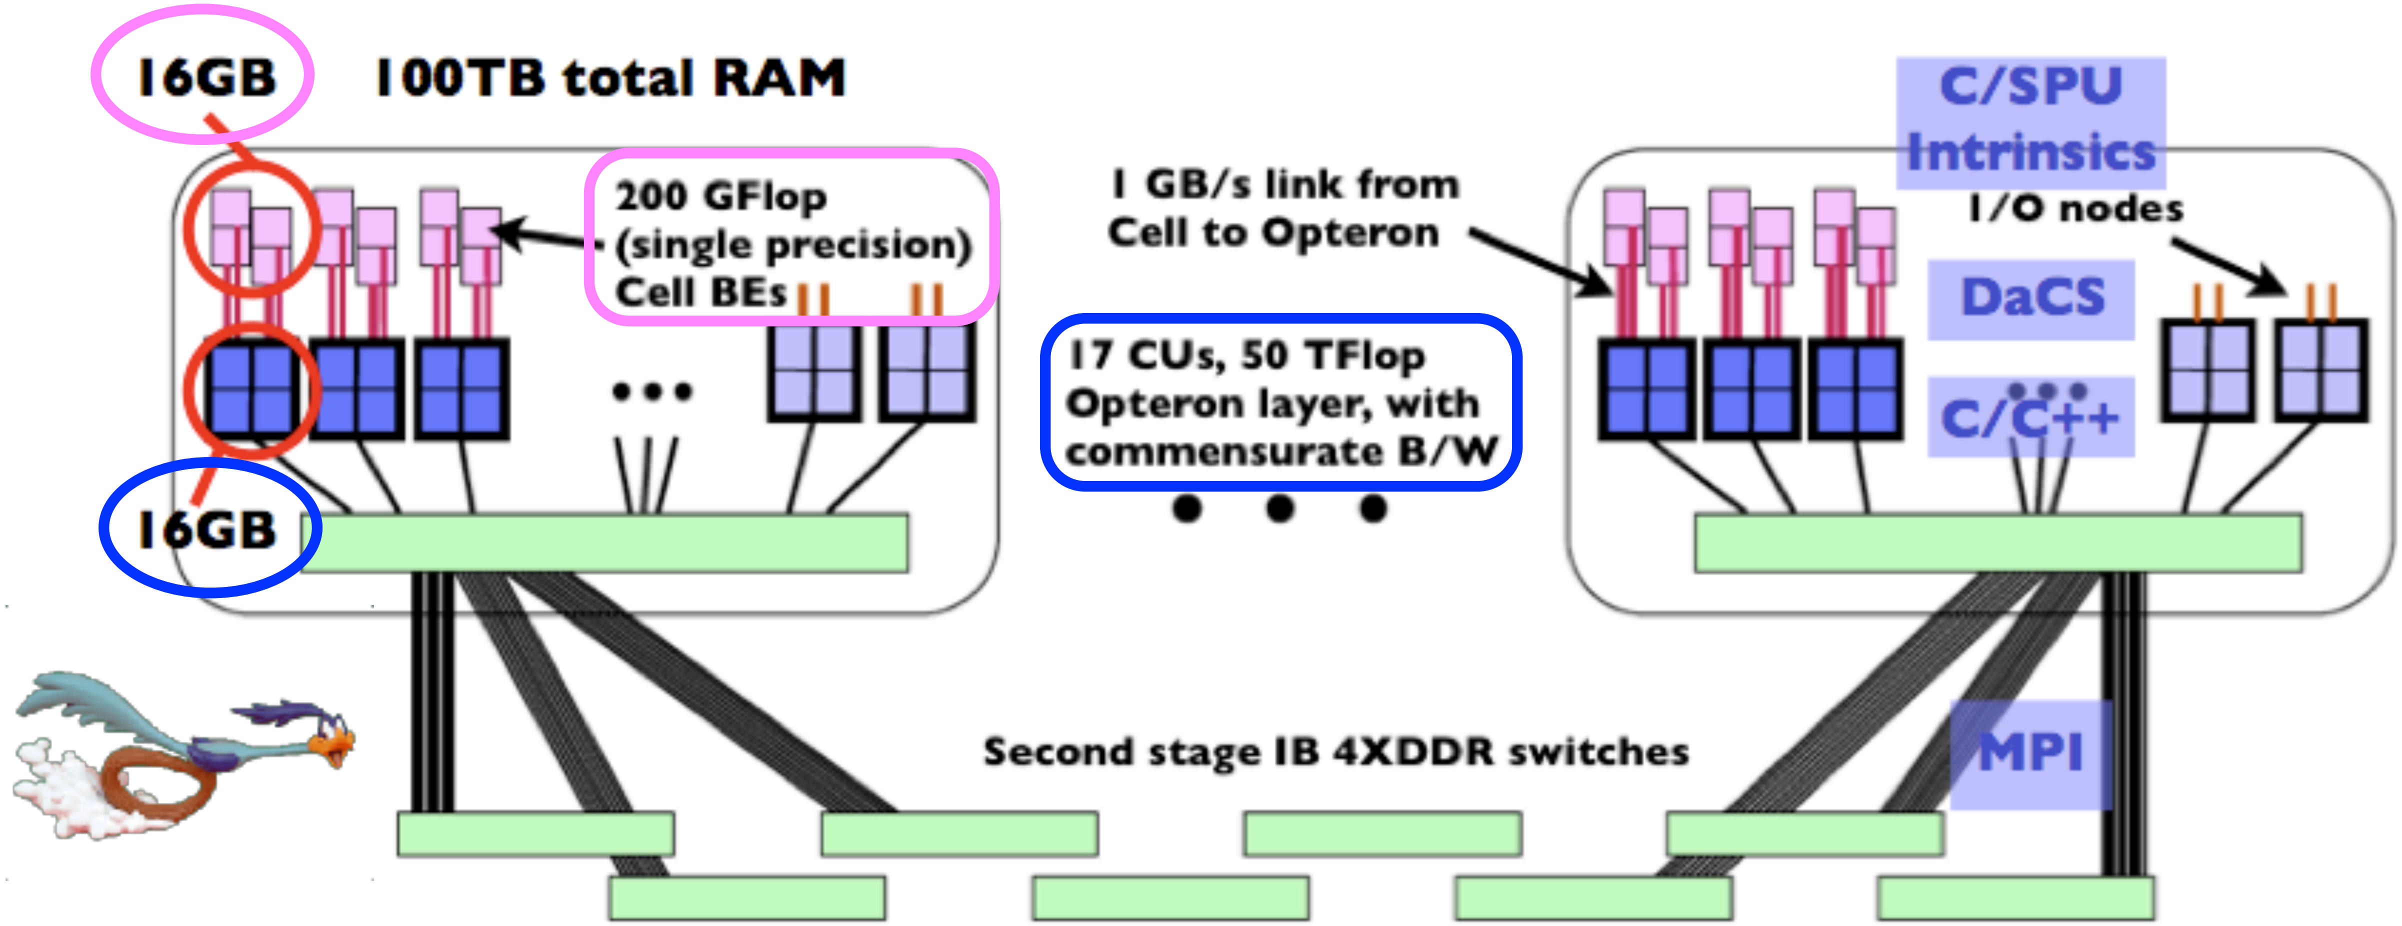

- **Opterons** have little compute (5% of total compute) but half the memory and balanced communication, for N-body codes, memory is limiting factor, so want to make best use of CPU layer
- Cells **dominate** the compute but communication is poor, 50-100 times out of balance (also true for CPU/GPU hybrid 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

HACC in a Nutshell

S. Habib et al. 2016, New Astronomy

- **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 (GPU)
- **Overload concept to allow for easy swap of short-range solver and minimization of communication (reassignment of passive/active in regular intervals)**
- **Adaptive time stepping, analysis on the fly, mixed precision, custom I/O, ...**

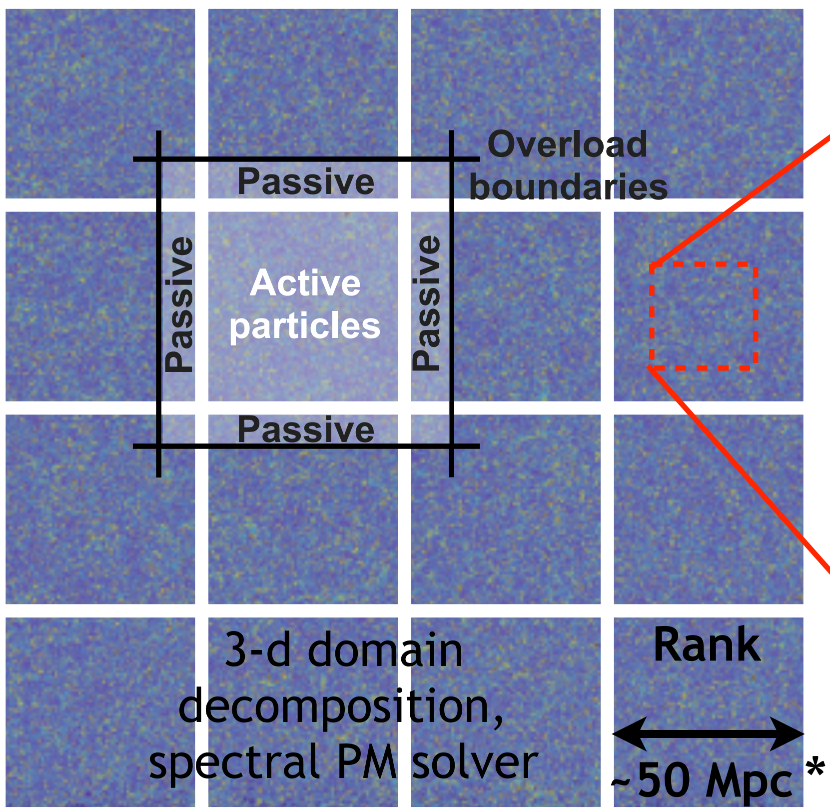

HACC Top Layer

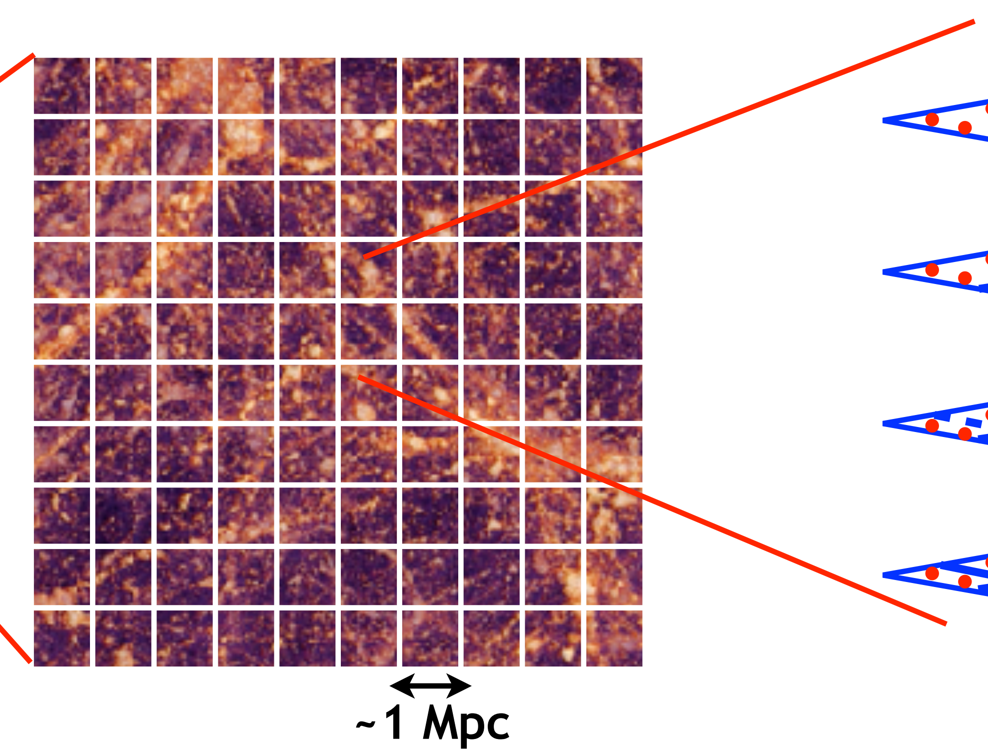

HACC 'Nodal' Layer

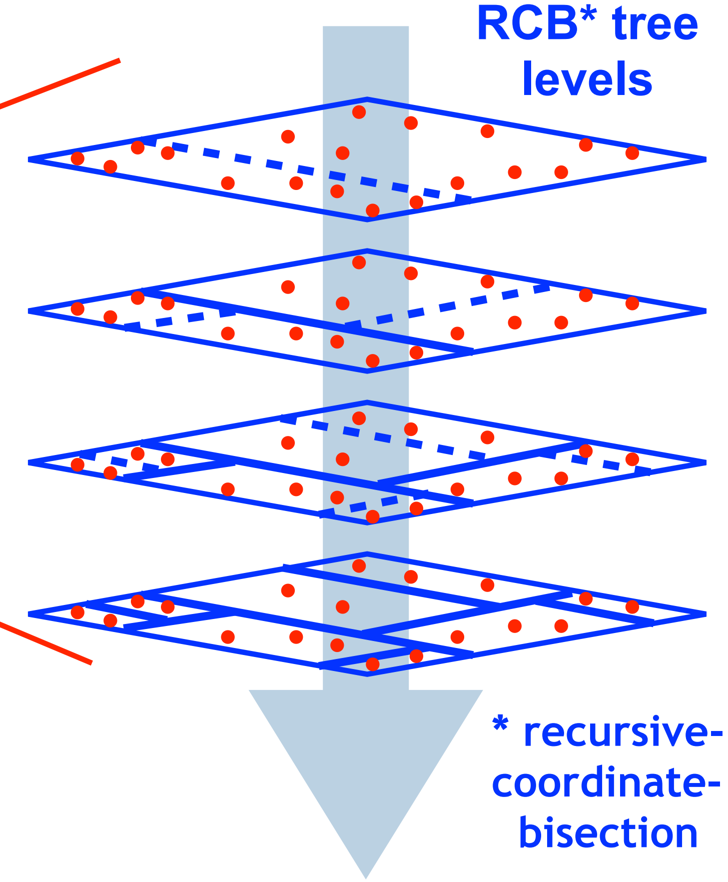

* Mpc~3M light years, distance between 2 galaxies

HACC (MC^3) Performance on Roadrunner

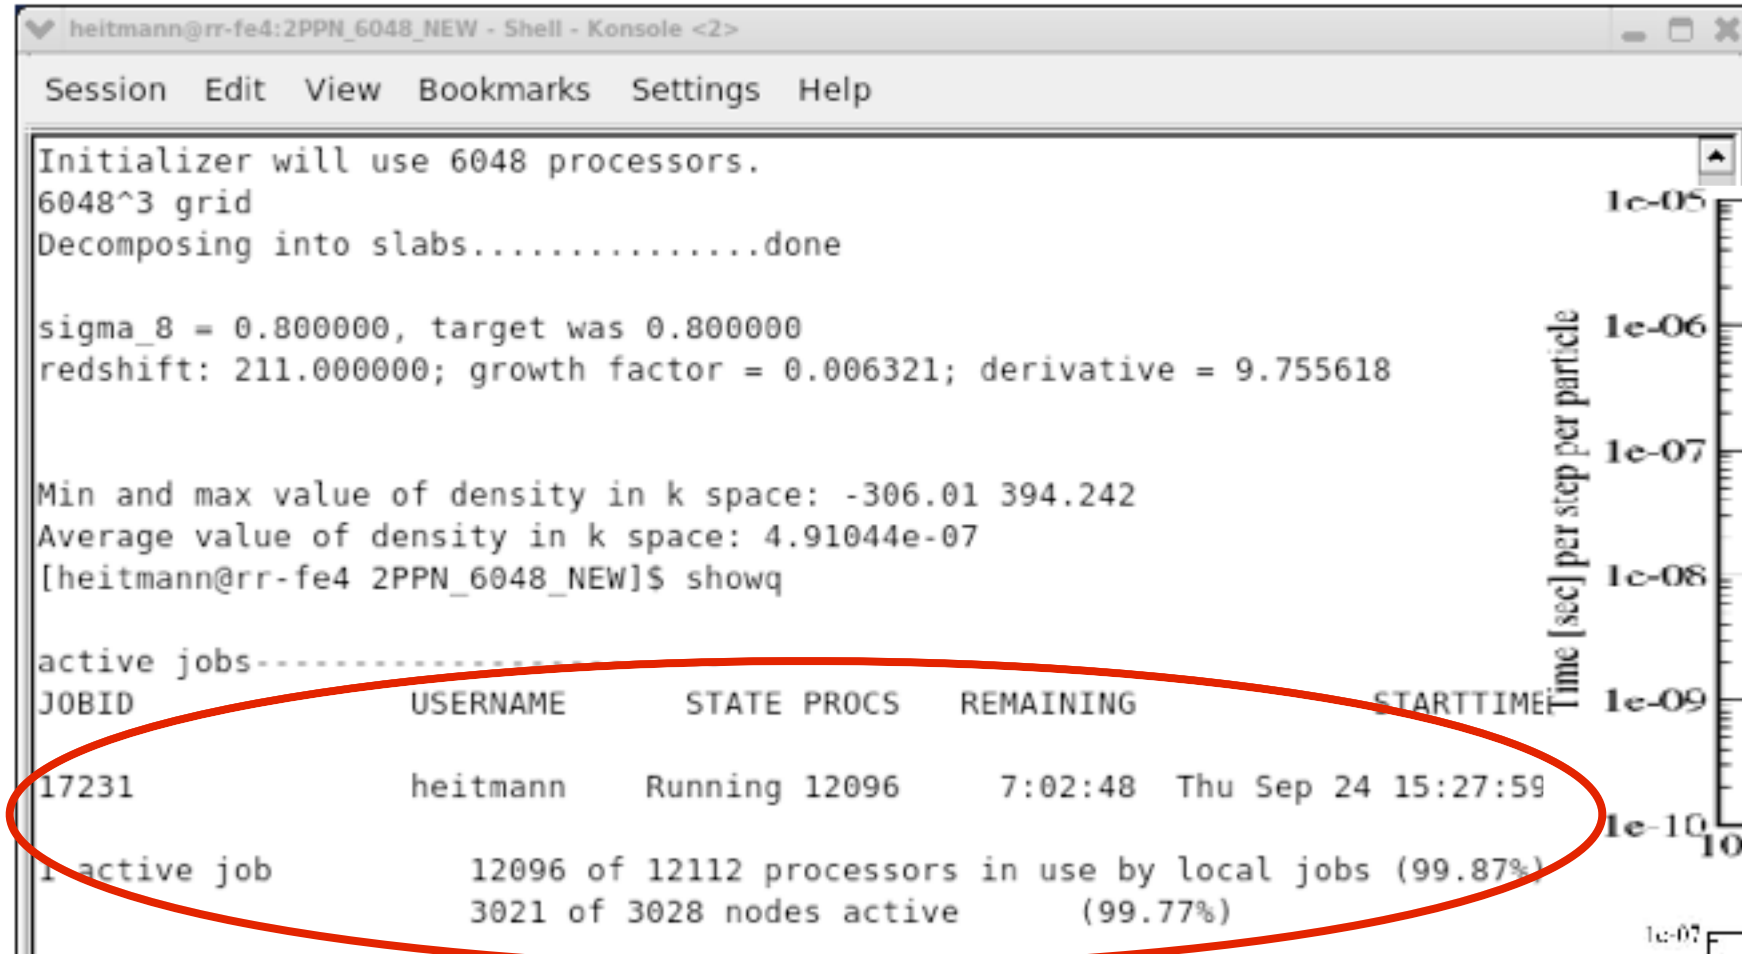

```
heitmann@rr-fe4:2PPN_6048_NEW - Shell - Konsole <2>
Session  Edit  View  Bookmarks  Settings  Help
Initializer will use 6048 processors.
6048^3 grid
Decomposing into slabs...............done

sigma_8 = 0.800000, target was 0.800000
redshift: 211.000000; growth factor = 0.006321; derivative = 9.755618


Min and max value of density in k space: -306.01 394.242
Average value of density in k space: 4.91044e-07
[heitmann@rr-fe4 2PPN_6048_NEW]$ showq

active jobs------------------------
JOBID              USERNAME      STATE  PROCS   REMAINING              STARTTIME

17231              heitmann    Running 12096     7:02:48   Thu Sep 24 15:27:59

1 active job          12096 of 12112 processors in use by local jobs (99.87%)
                       3021 of 3028 nodes active       (99.77%)
```

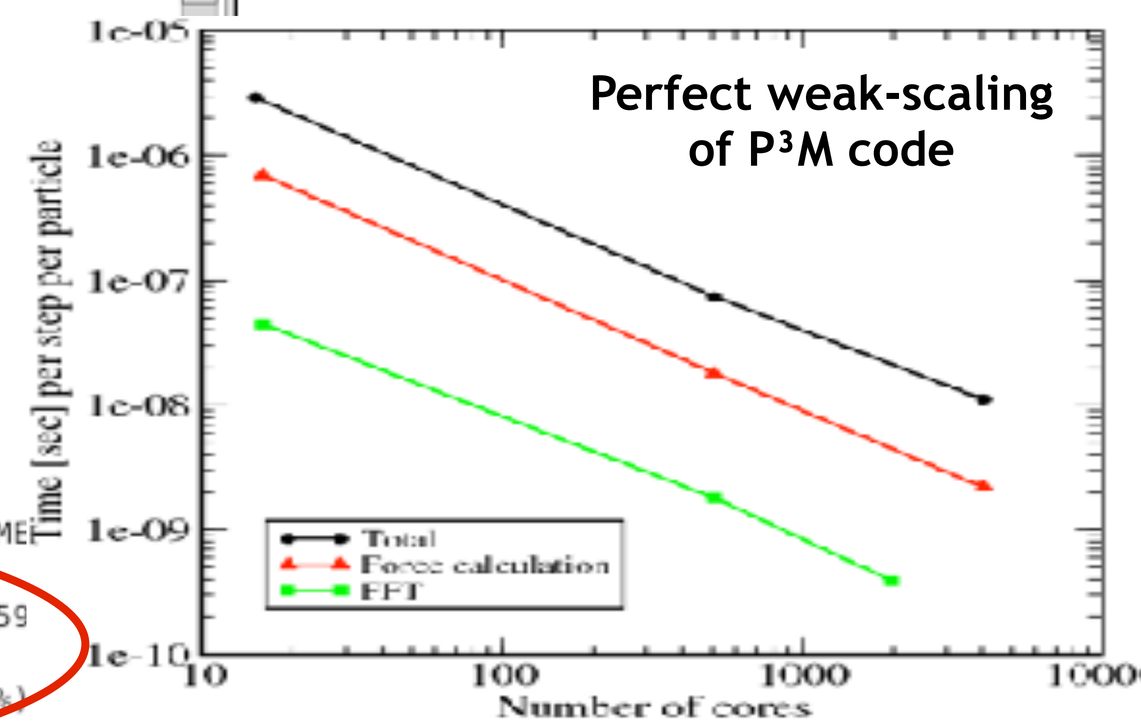

Perfect weak-scaling of P^{3}M code

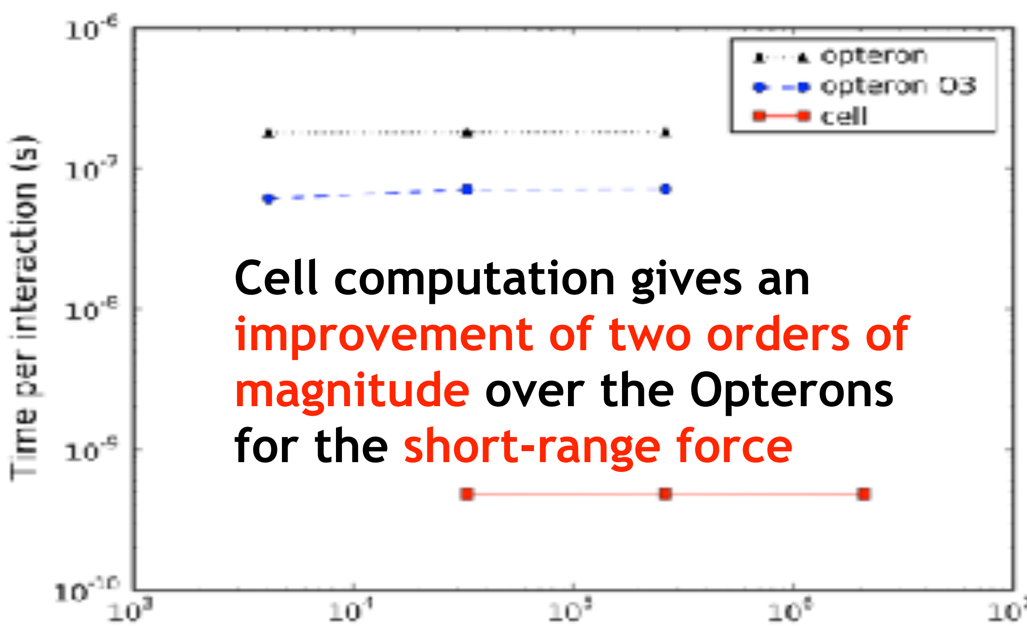

Cell computation gives an improvement of two orders of magnitude over the Opterons for the short-range force

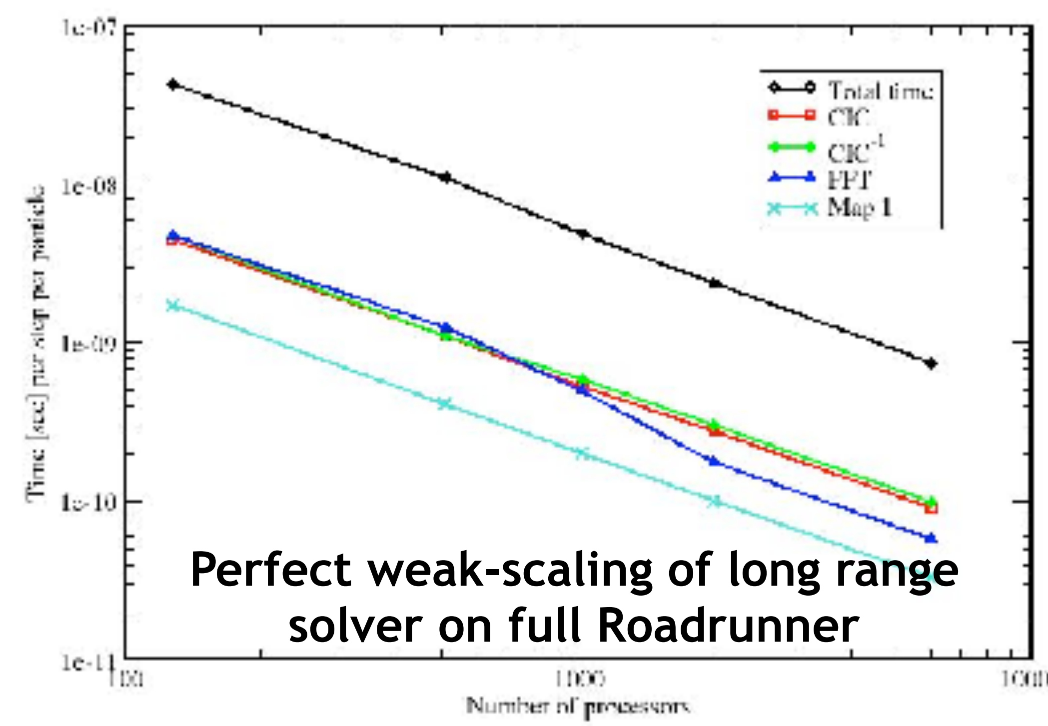

Perfect weak-scaling of long range solver on full Roadrunner

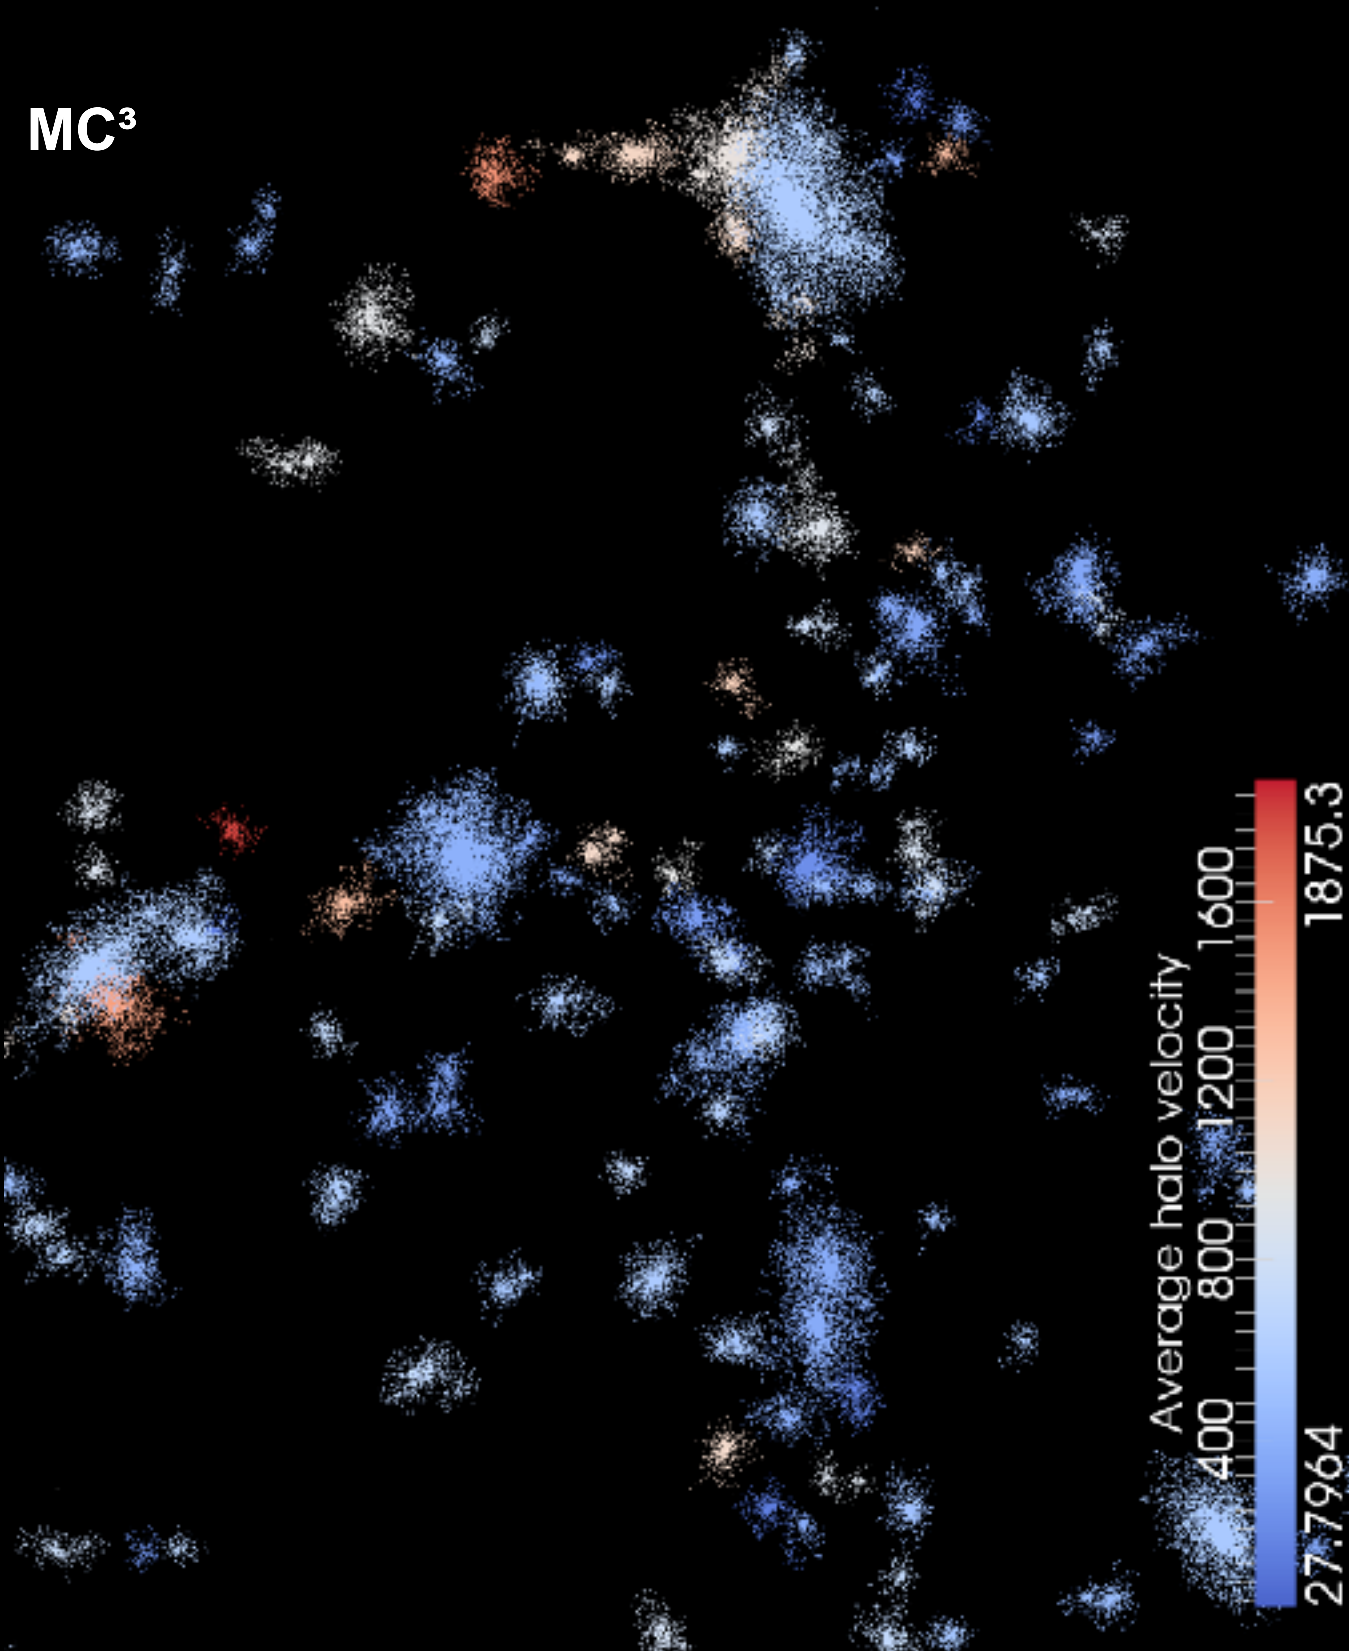

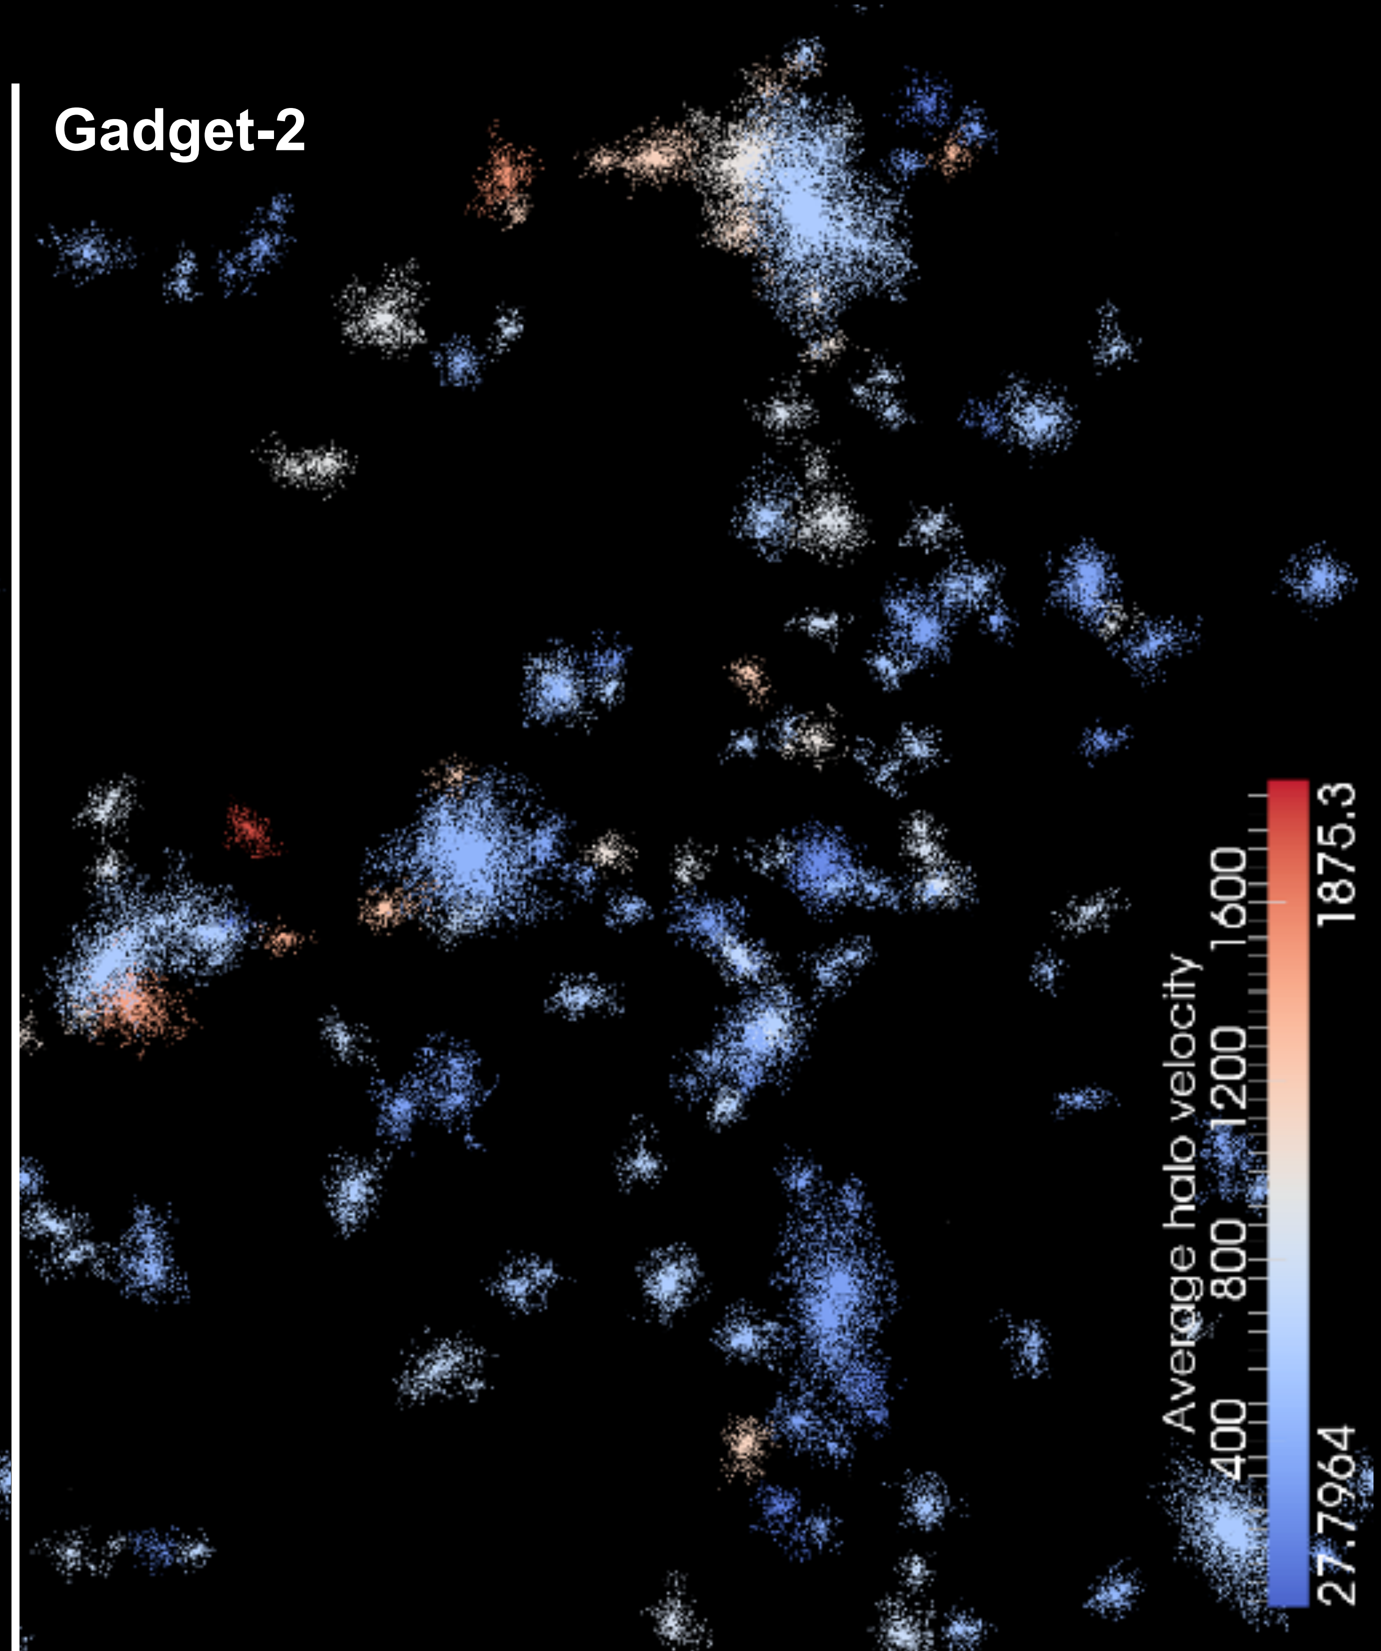

Snapshot from Code Comparison simulation, ~25 Mpc region; halos with > 200 particles, b=0.15
Differences in runs: P³M vs. TPM, force kernels, time stepper: MC³: a; Gadget-2: log(a)
Power spectra agree at sub-percent level

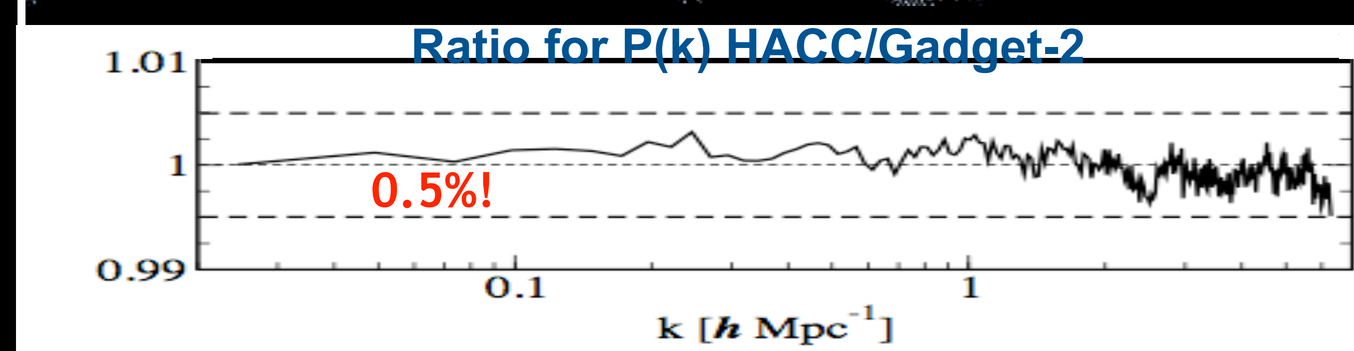

The HACC Story

Chapter 2 — The Midwest

The HACC Story

Chapter 2 — The Midwest

The HACC Story

Chapter 2 — The Midwest

The HACC Story

Chapter 2 — The Midwest

The HACC Story

Chapter 2 — The Midwest

The Story continues, Multi-core systems, BG/Q

- **Proof of concept for "easy" portability: IBM Blue Gene (BG) systems**
  - BG/Q Mira at Argonne: 10 PFlops, arrived in 2012, 750,000 cores, 16GB per node
  - BG/Q Sequoia at Livermore: twice as large
- **New challenges:**
  - BG/Q systems have many cores but no accelerators
  - Slab-decomposed FFT does not scale well on very large number of cores
- **Solutions:**
  - Particle-particle interaction now replaced by tree, OpenMP node parallel
  - Pencil decomposed FFT
  - Adaptive time stepping
- **Achieved 13.94PFlops on Sequoia, 90% parallel efficiency on 1,572,864 cores**
- **3.6 trillion particle benchmark run**

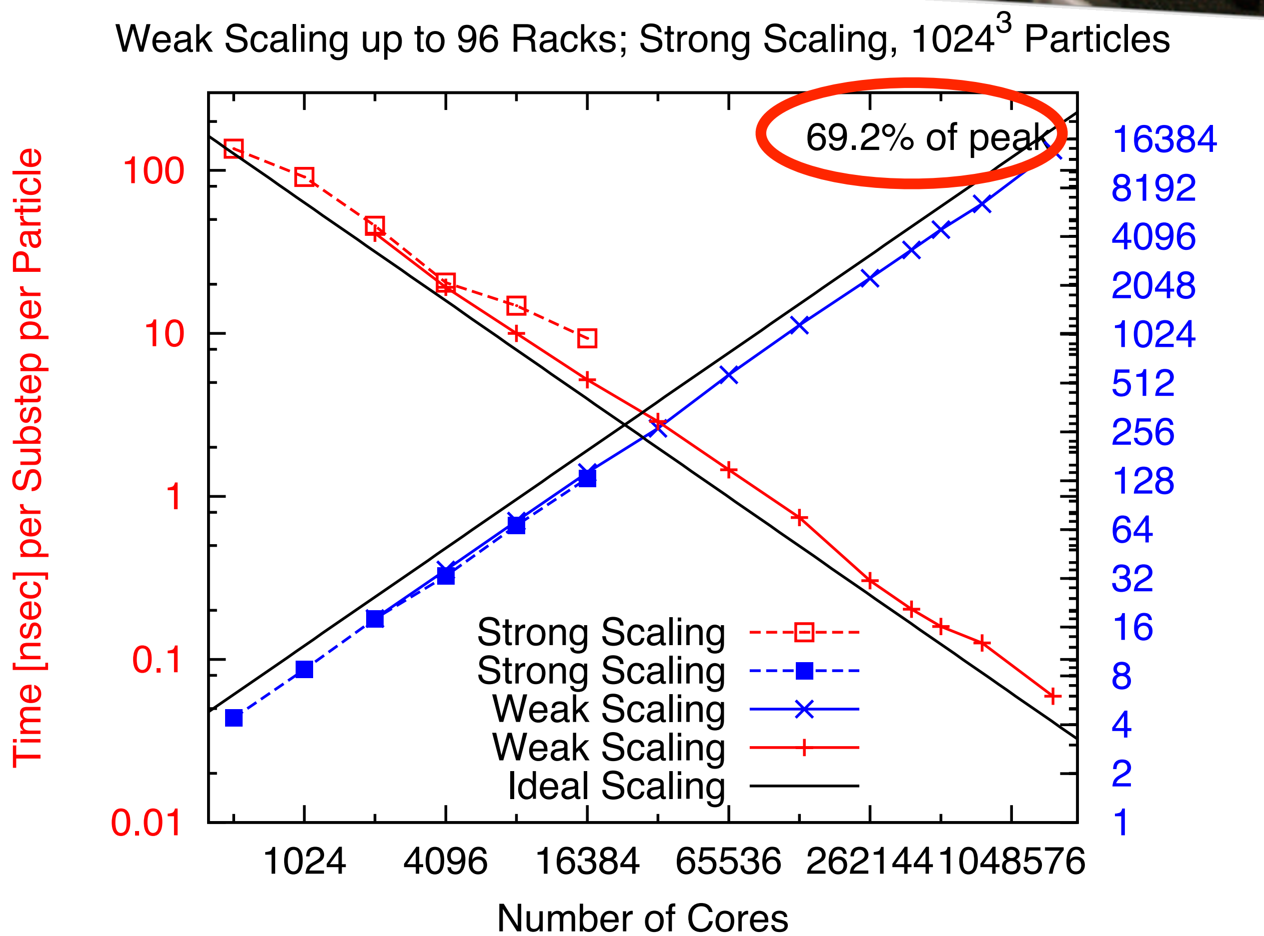

Habib et al. SC12, Gordon Bell Finalist

The HACC Story

Chapter 3 — The South

The HACC Story

Chapter 3 — The South

The HACC Story

Chapter 3 — The South

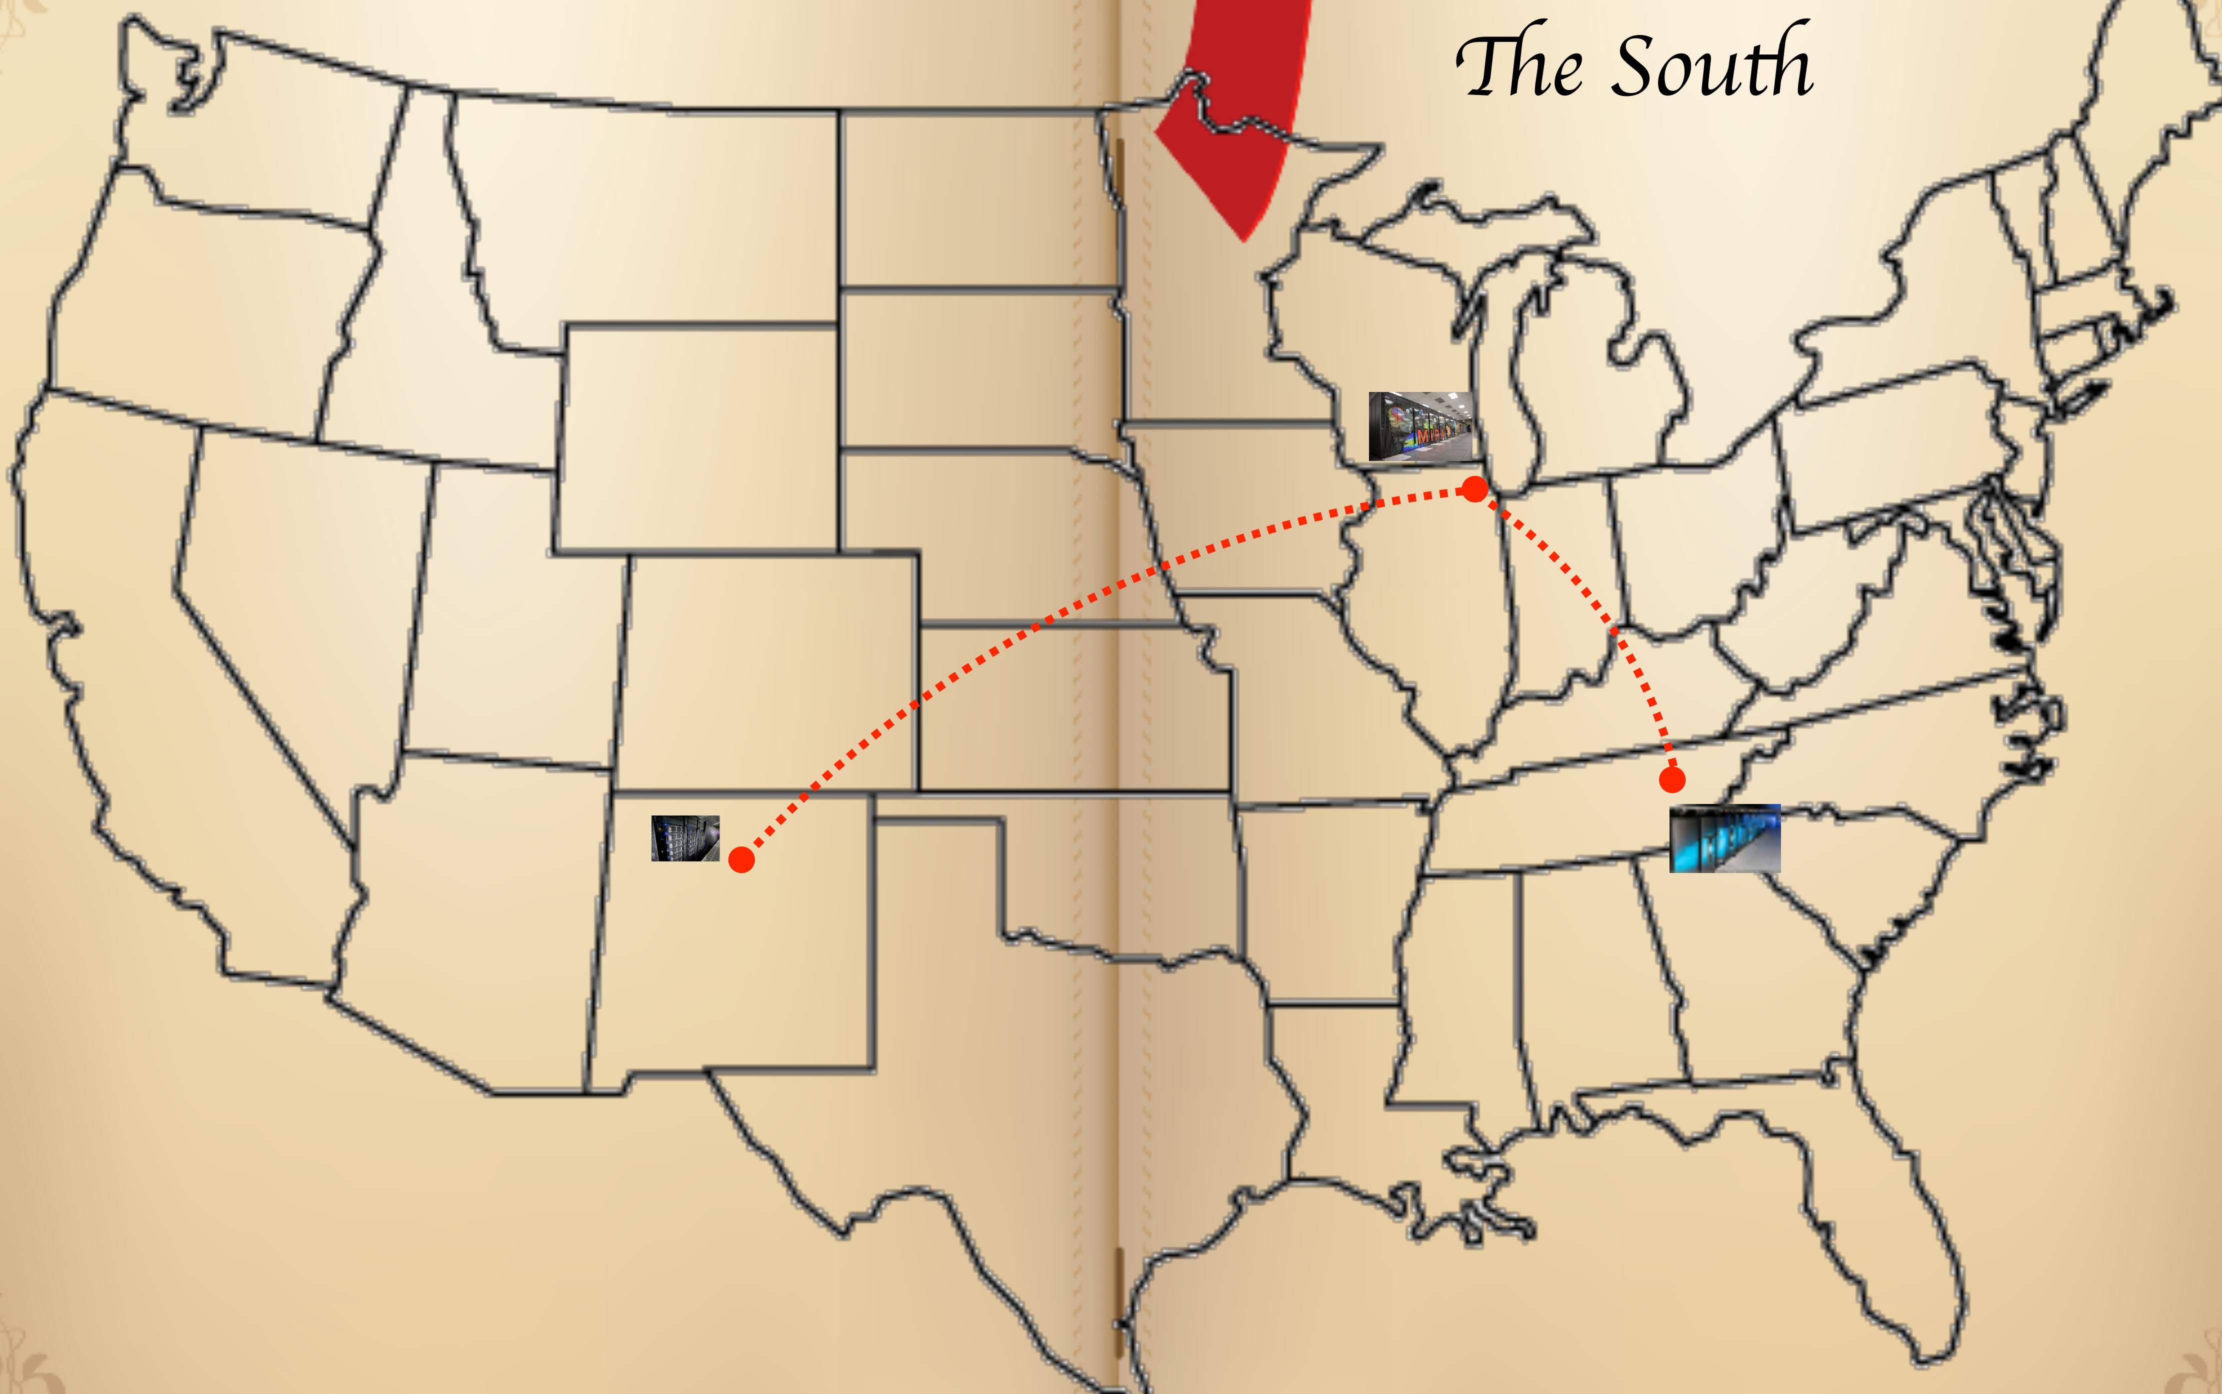

The HACC Story

Chapter 3 — The South

The HACC Story

Chapter 3 — The South

HACC Accelerated Again: CPU+GPU

- **Proof of concept for easy portability:** Go back to Roadrunner short-range solver, replace Cell part by GPU implementation (first version already existed in 2010)
- **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, now in CUDA
  - New load-balance scheme
  - Now running on Summit

Weak and Strong Scaling Results

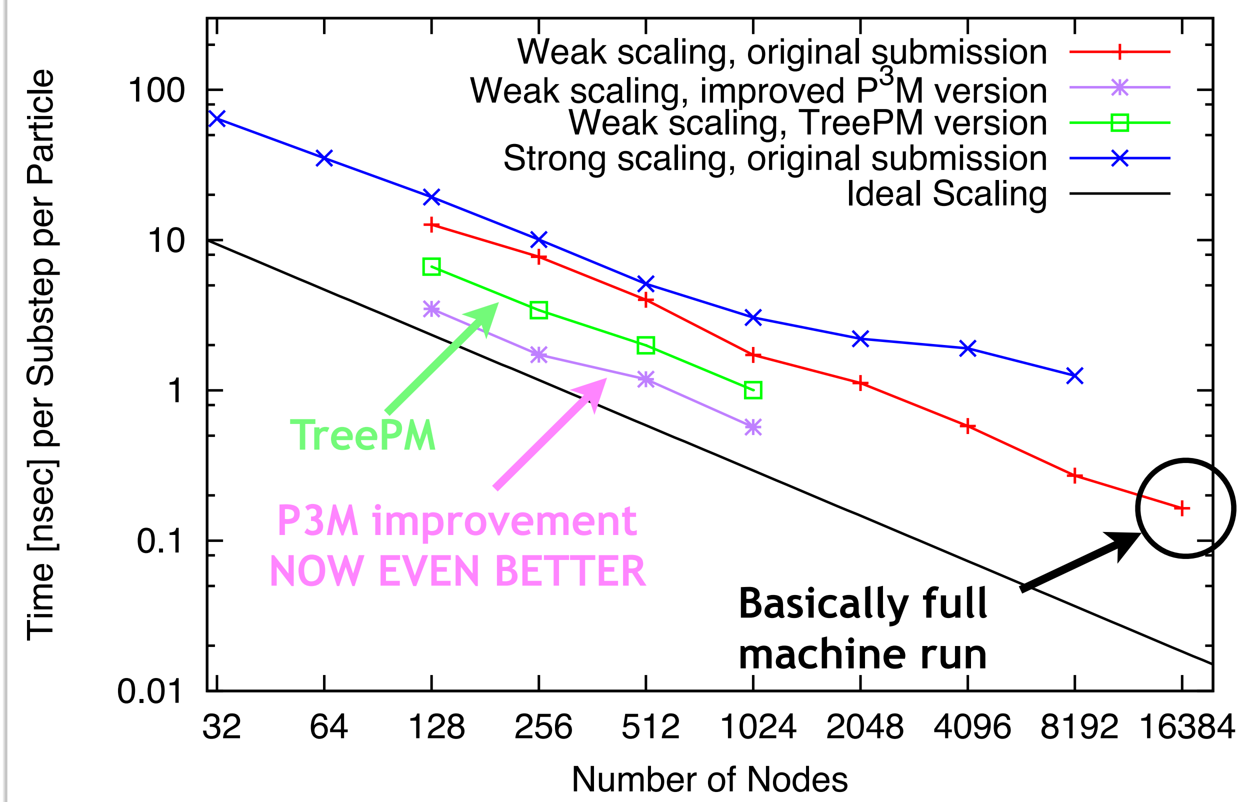

S. Habib et al. 2013: SuperComputing13, Gordon Bell Finalist

- 20.54 Pflops peak performance evolving 1.23 trillion particles in test run on ~75% of machine (full machine not available at the time)
- Summitdev: Speed up of short-range solver by ~3.2x (with no changes) on NVIDIA Pascals

The HACC Story

Chapter 4 –

The West and back Home

The HACC Story

Chapter 4 — The West and back Home

The HACC Story

Chapter 4 — The West and back Home

The HACC Story

Chapter 4 — The West and back Home

The HACC Story

Chapter 4 — The West and back Home

The HACC Story

Chapter 4 —

The West and back Home

The HACC Story

Chapter 4 — The West and back Home

The HACC Story

Chapter 4 — The West and back Home

The HACC Story

Chapter 4 — The West and back Home

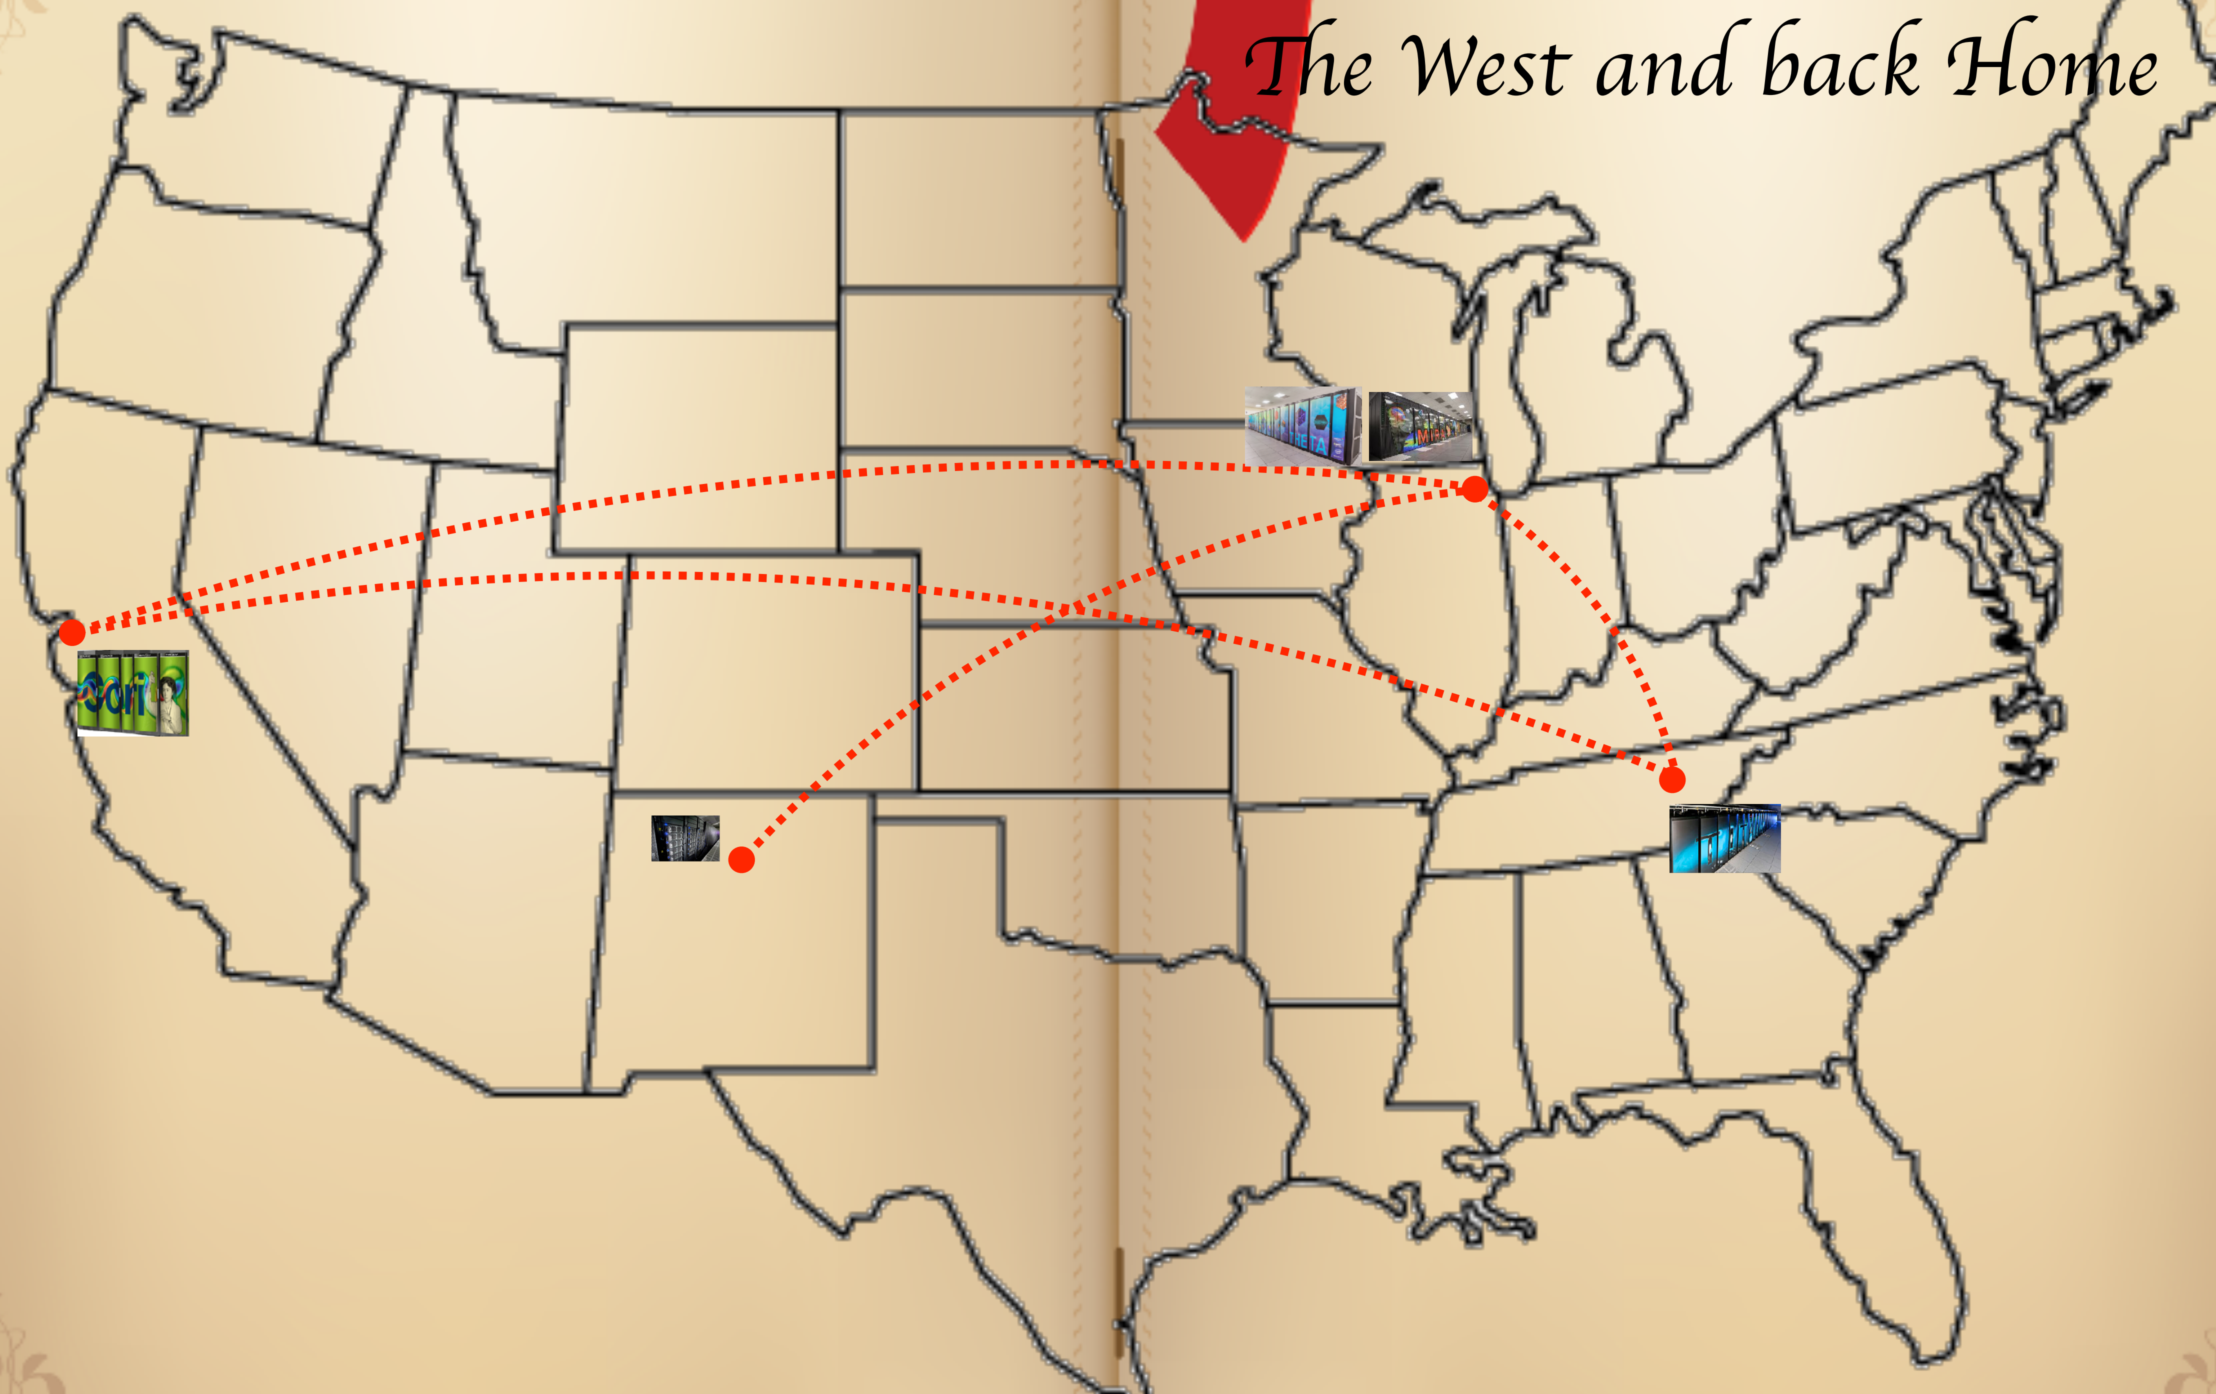

HACC on KNL: Successful Port and New Physics

- Proof of concept for easy portability: KNL
- Short-range solver again tree-based, minor modifications to the short-range solver, obtained expected performance for gravity-only part
- And: New physics! Supercomputers are getting faster but memory increase doesn't keep up, so instead of bigger, more physics: baryons
- HACC is particle-based code, therefore use Smoothed Particle Hydrodynamics (SPH) approach with new developments (Conservative reproducing kernel) to overcome traditional SPH short-comings = CRK HACC

Weak scaling up to 3072 nodes

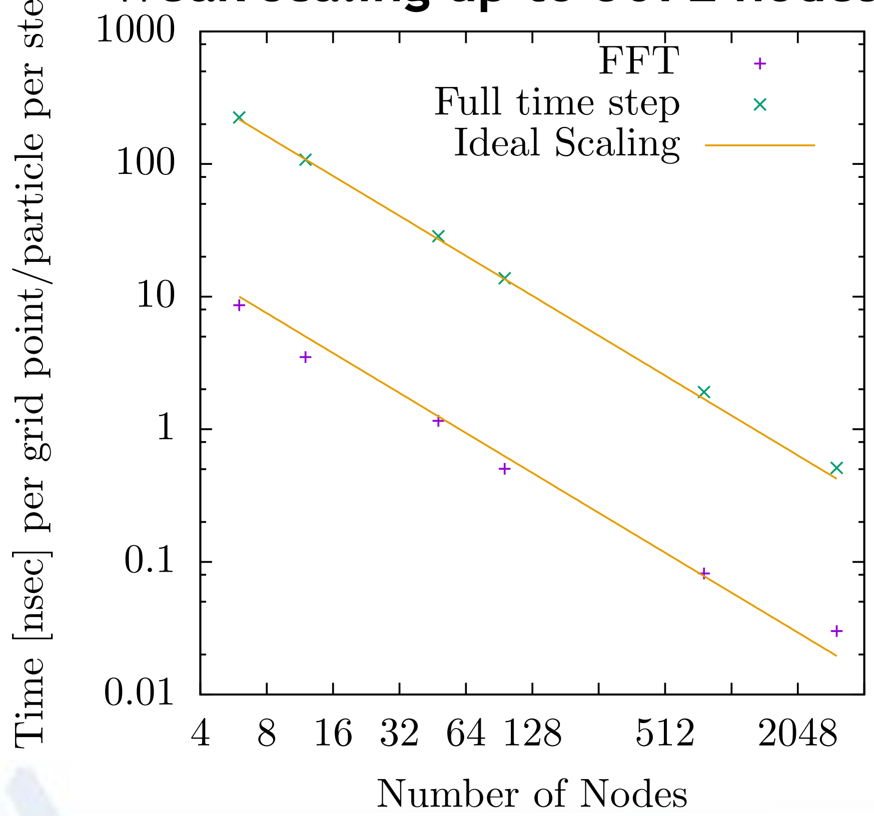

Dark matter density

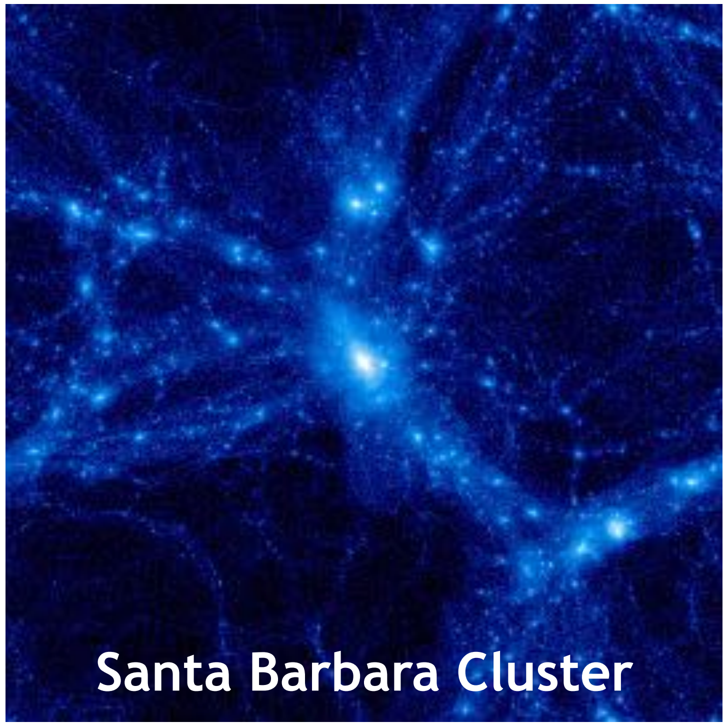

Baryon density

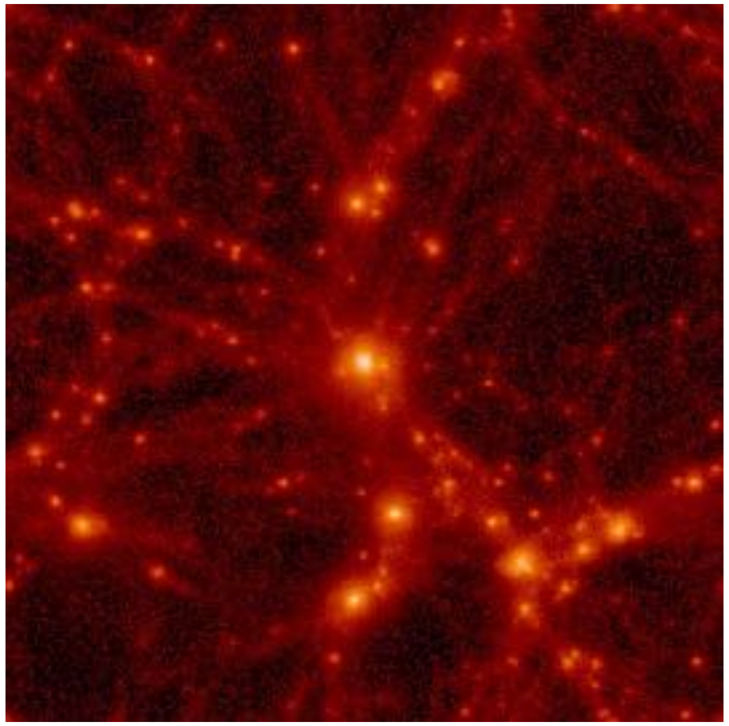

Borg Cube Simulation carried out with CRK-HACC on Theta

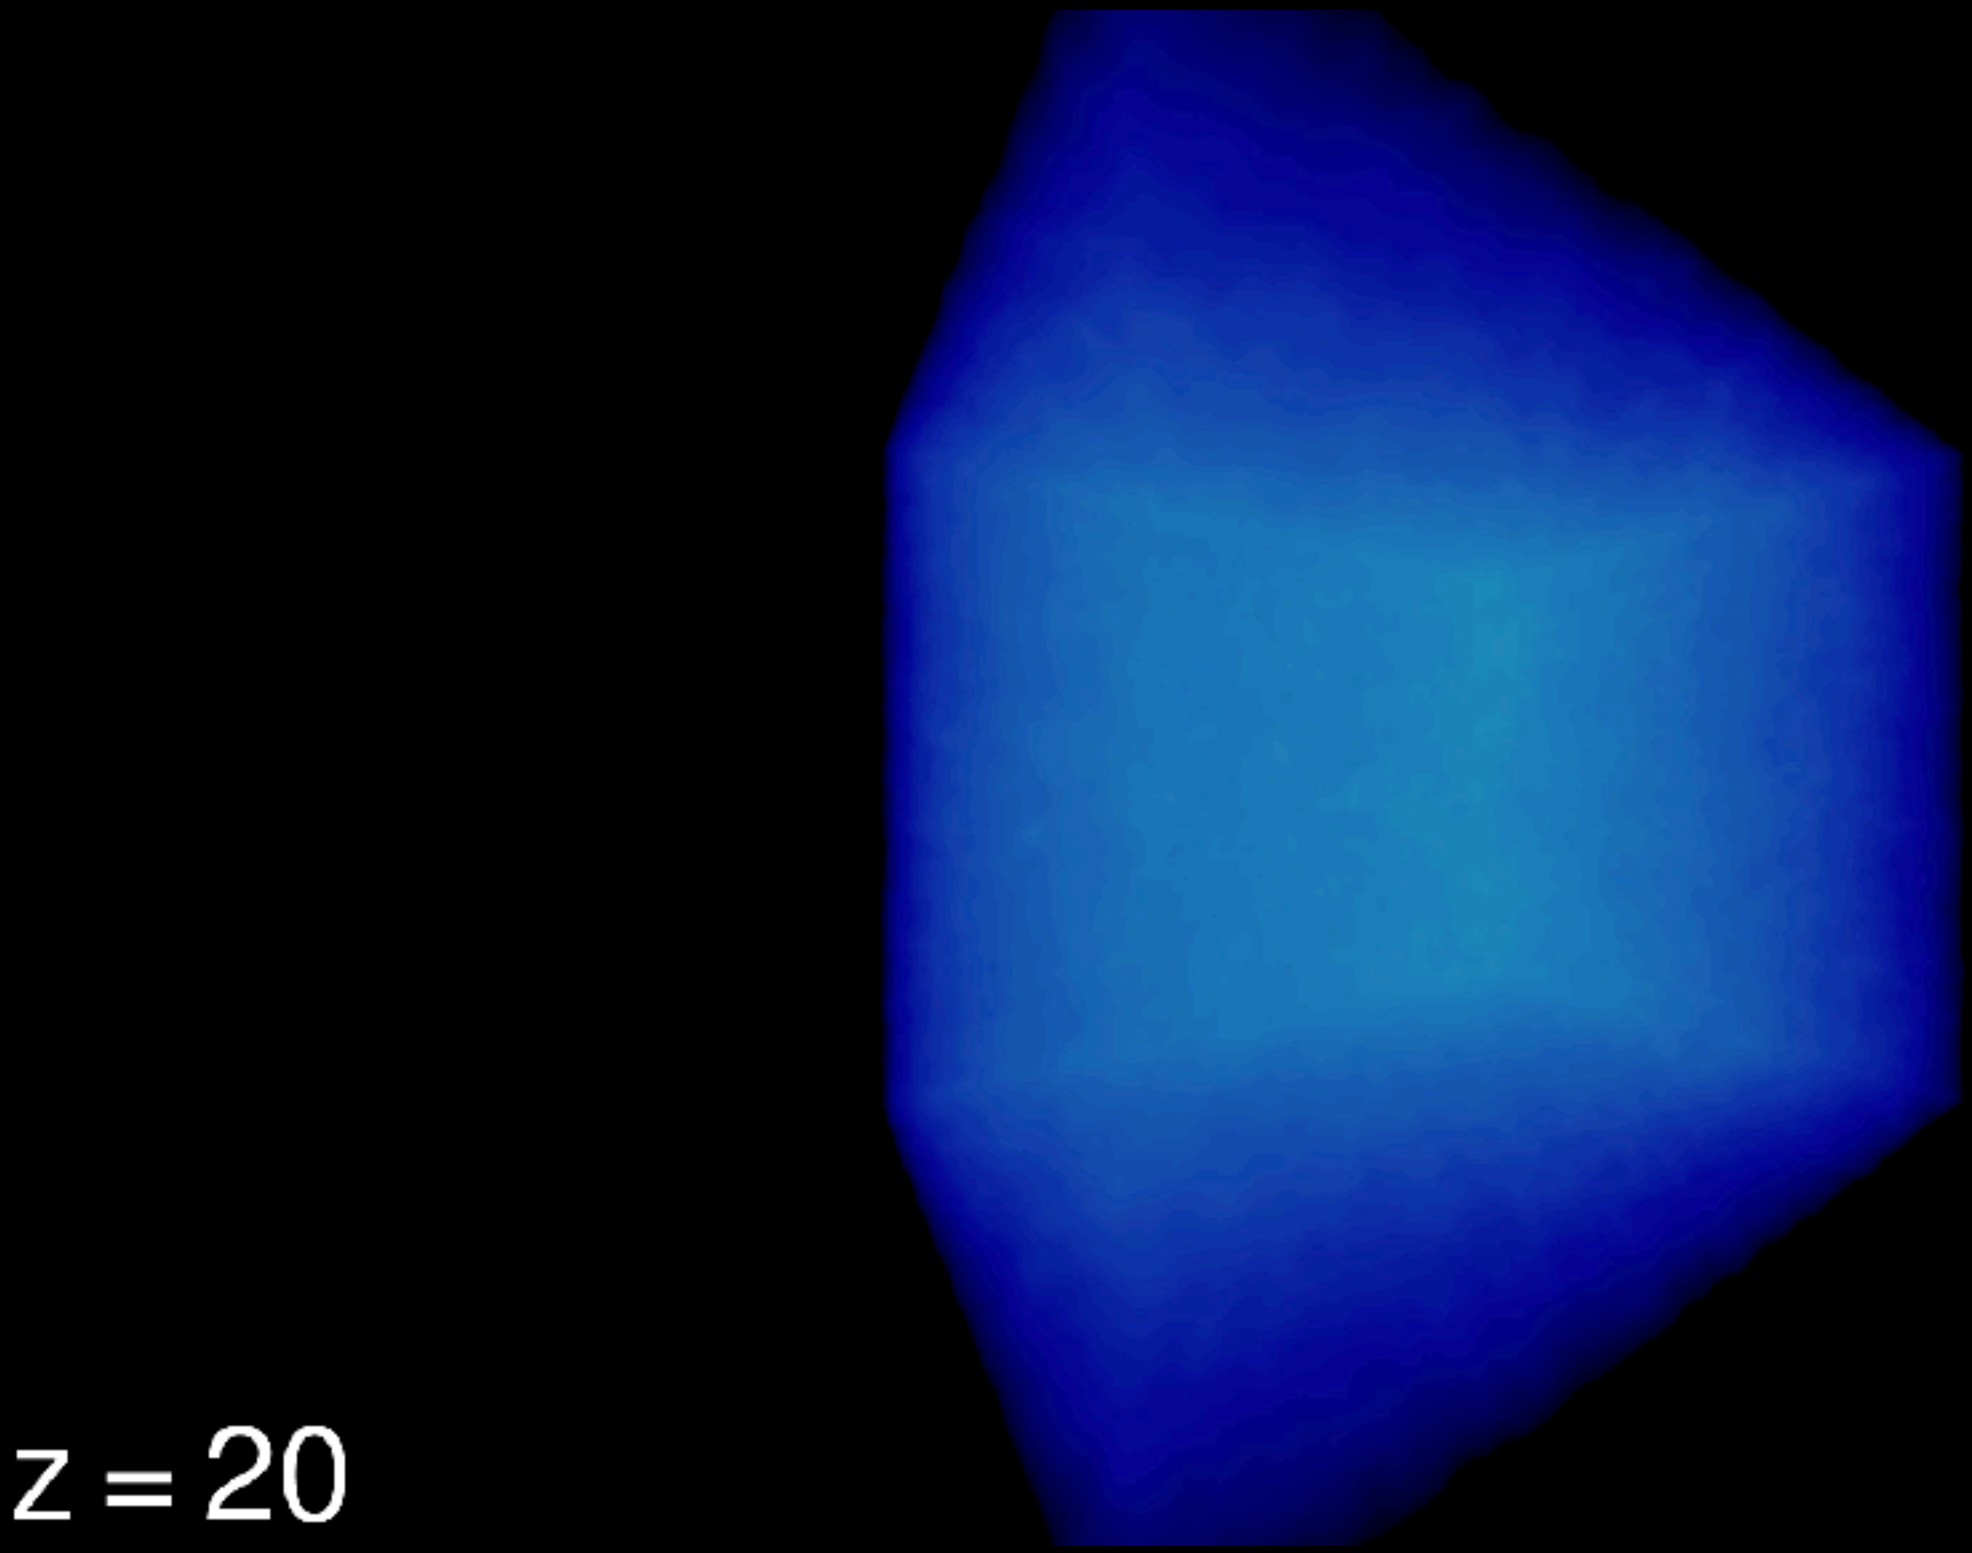

Emberson et al., ApJ 2019. Color: Baryon temperature, white: density.
Full size: 1126 Mpc, 24 billion particles, movie: Zoom in to (35x35x47)Mpc volume.

Borg Cube Simulation carried out with CRK-HACC on Theta

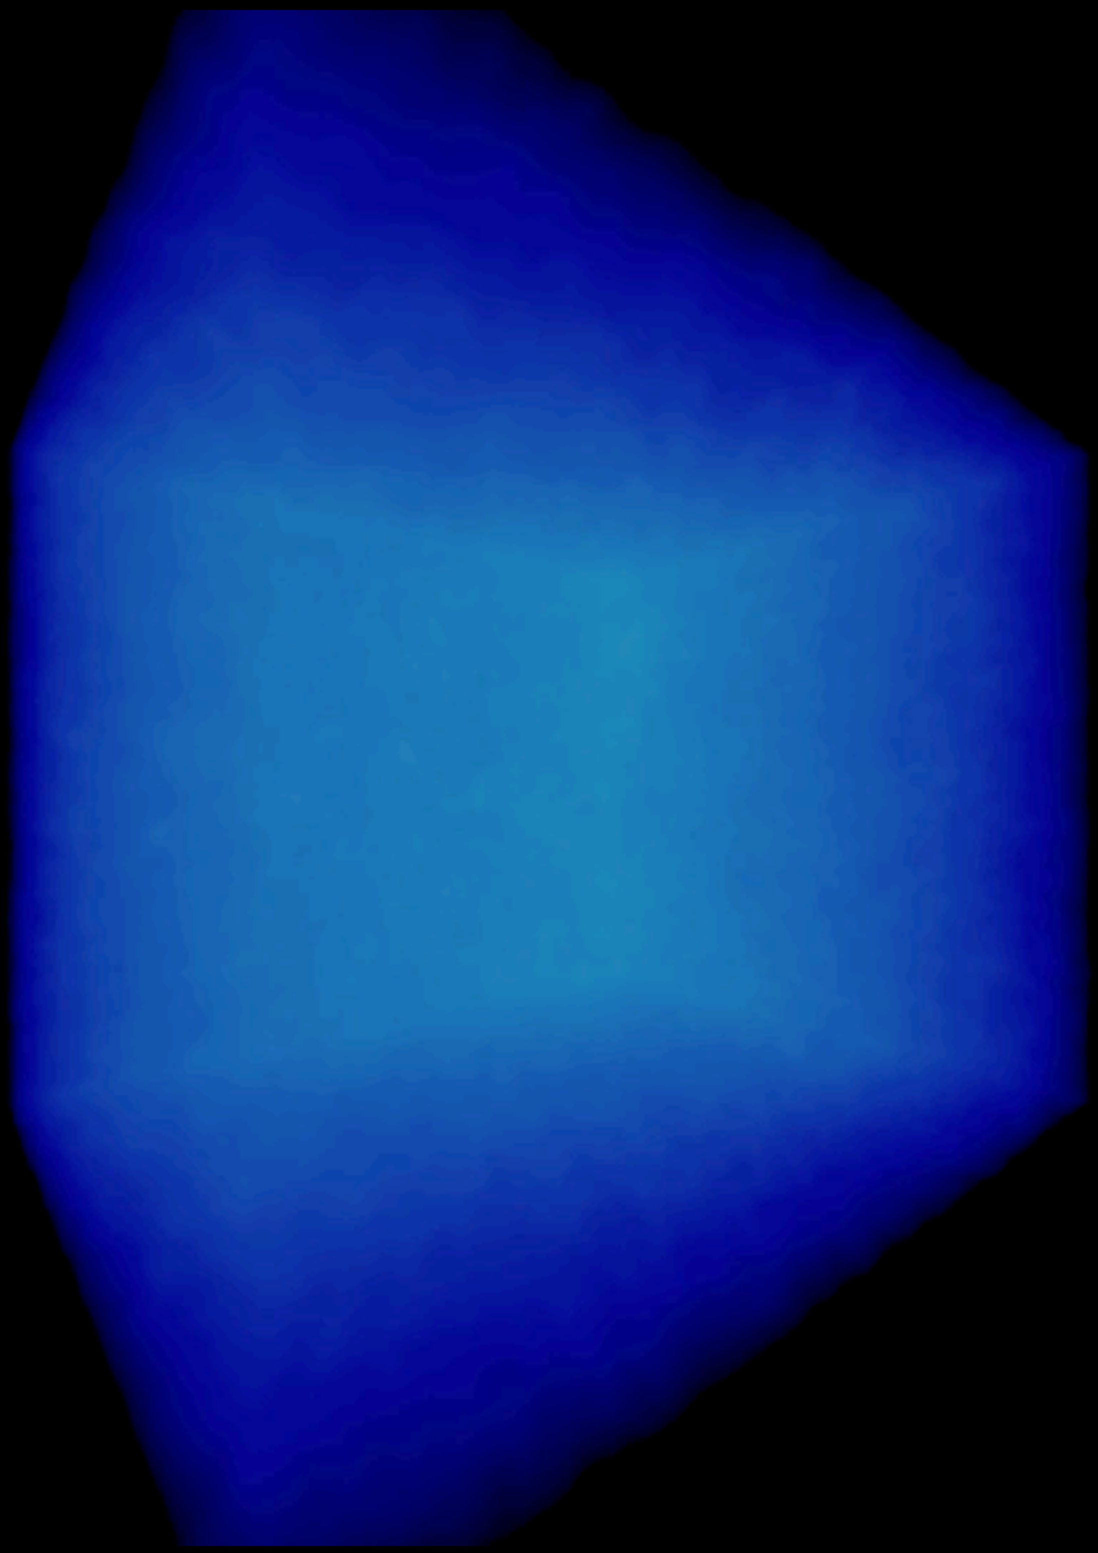

z = 20

Emberson et al., ApJ 2019. Color: Baryon temperature, white: density.
Full size: 1126 Mpc, 24 billion particles, movie: Zoom in to (35x35x47)Mpc volume.

The HACC Story

Chapter 5 —

Back to the South

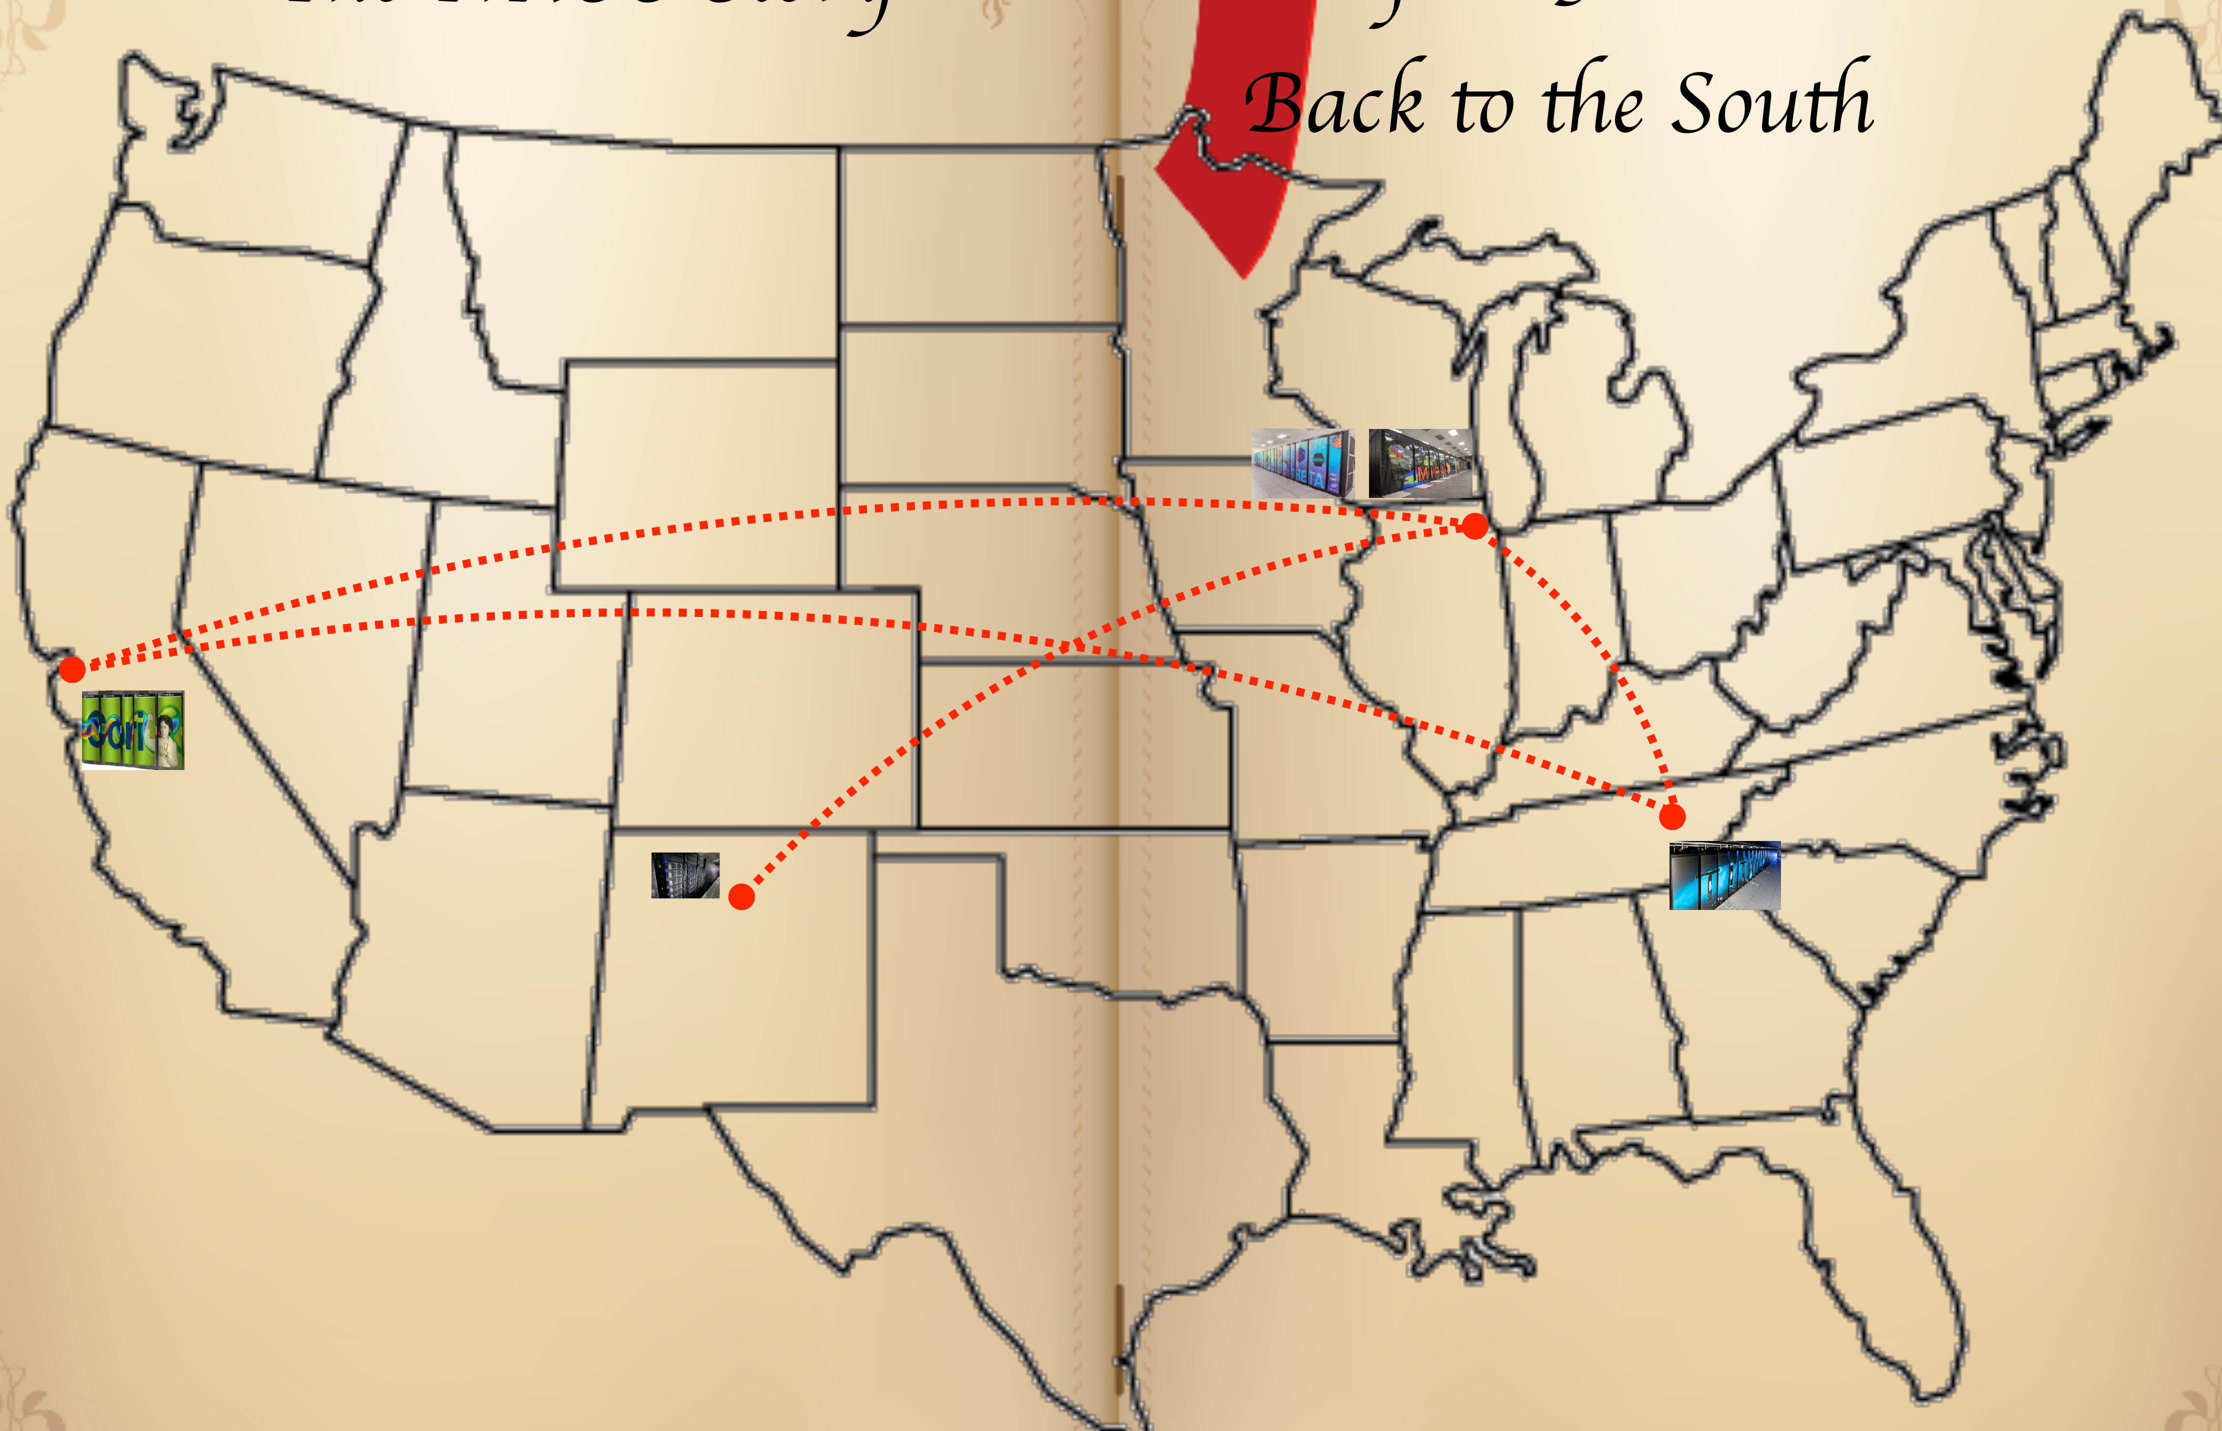

The HACC Story

Chapter 5 — Back to the South

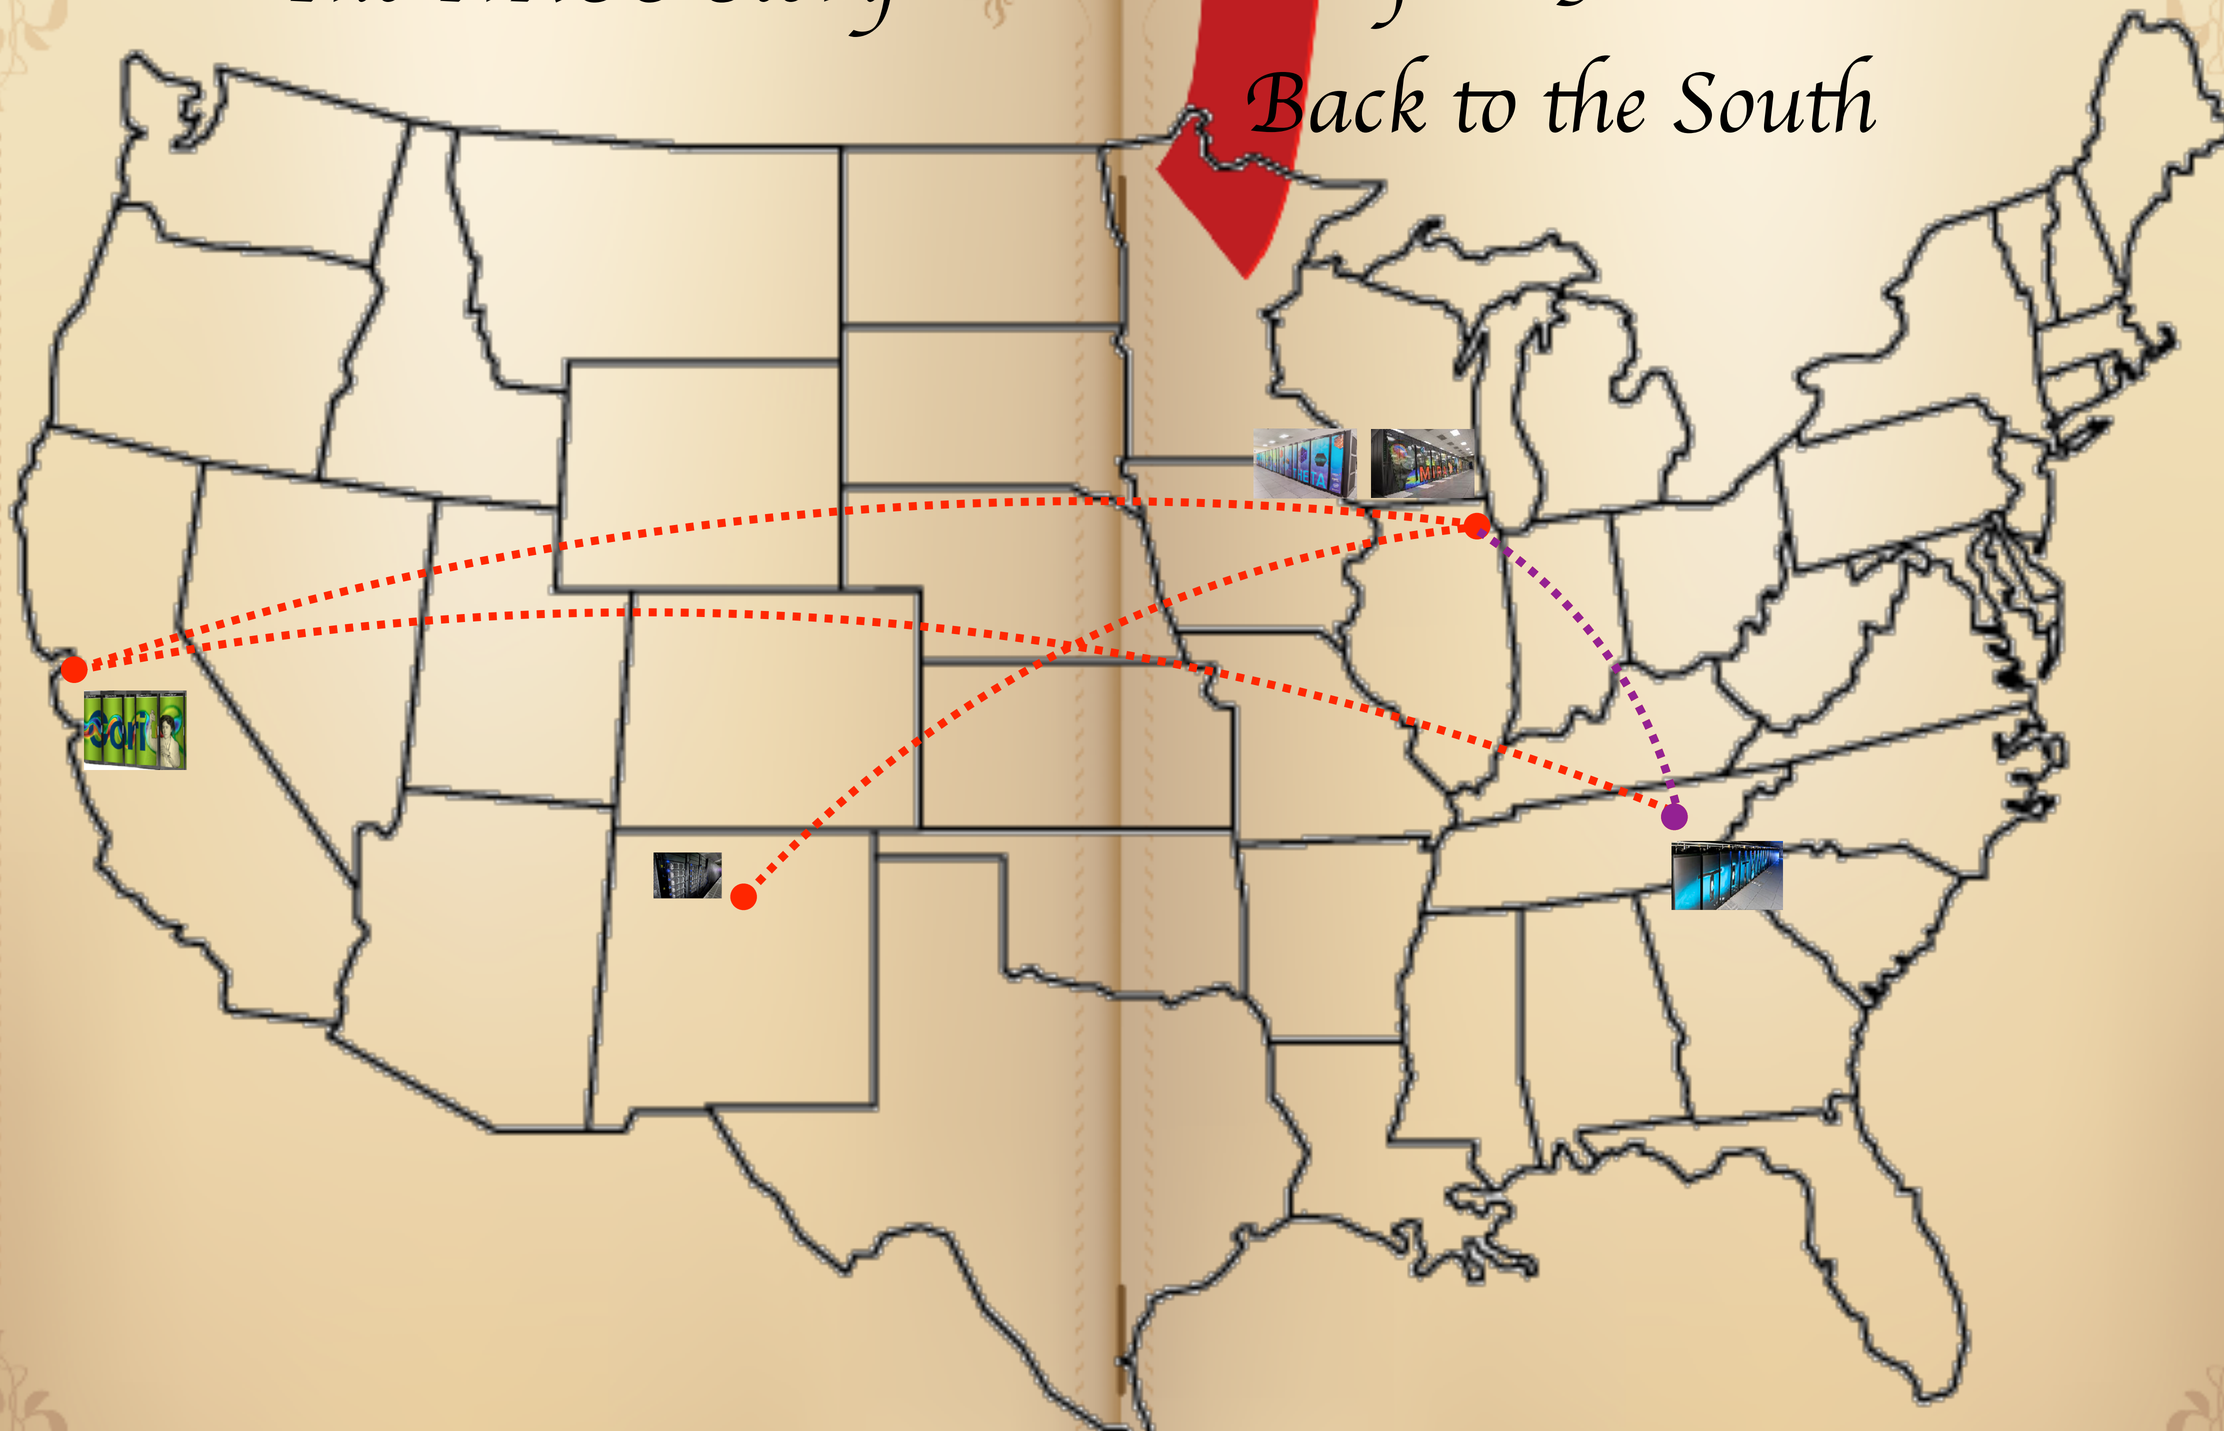

The HACC Story

Chapter 5 — Back to the South

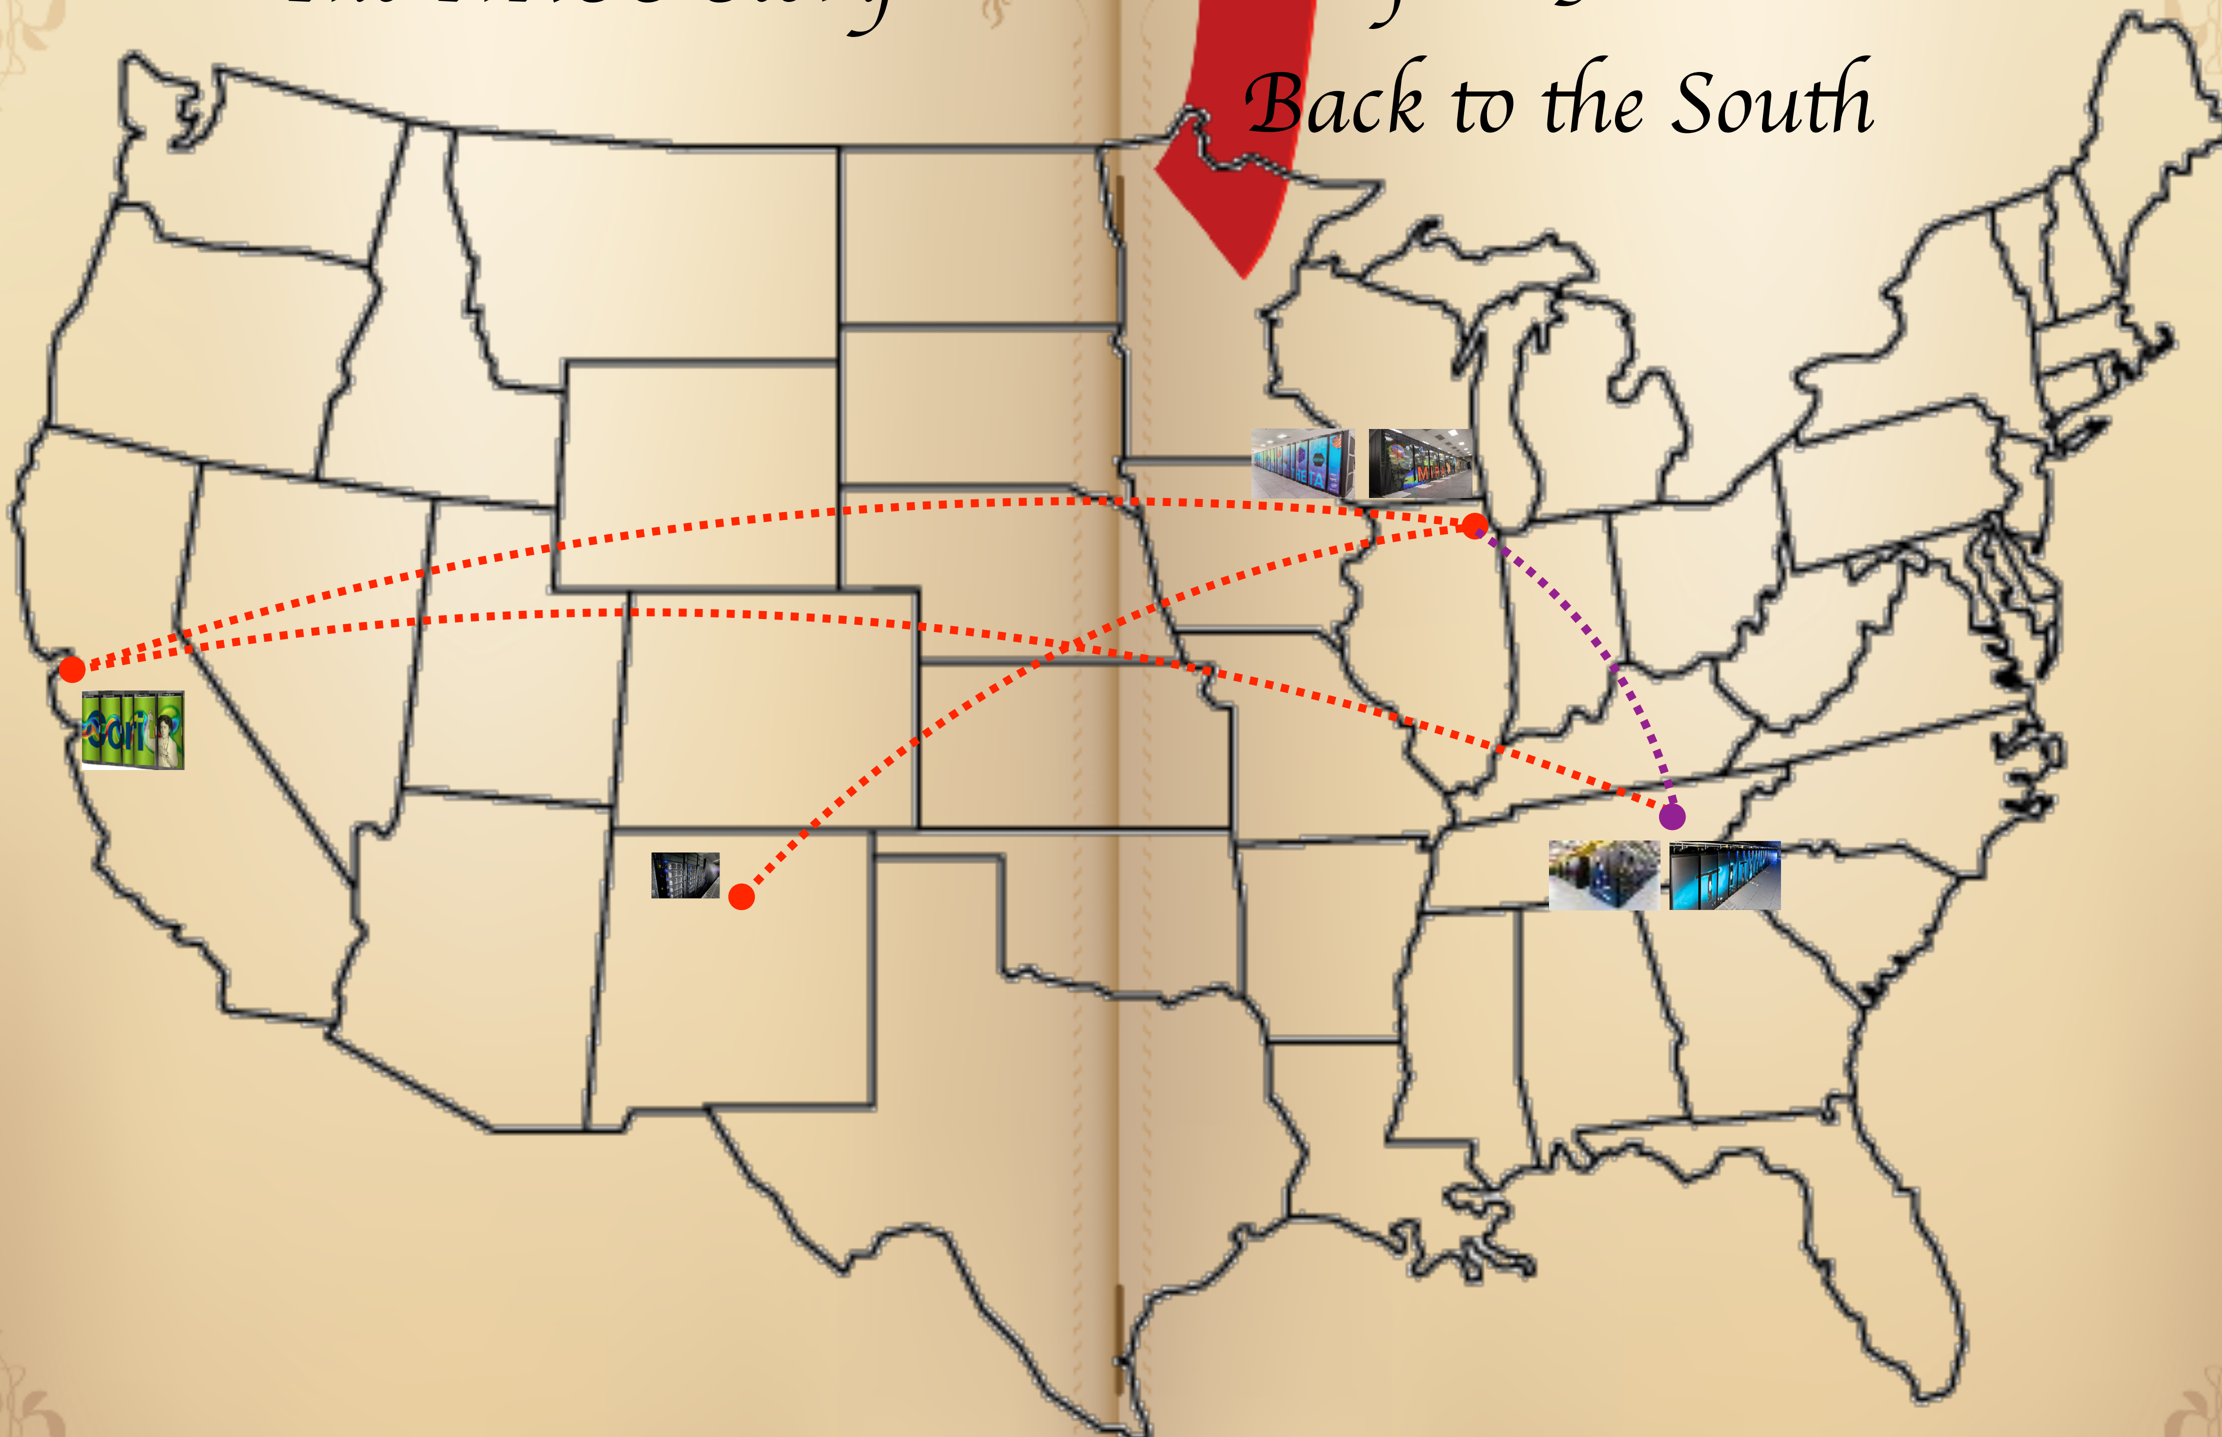

The HACC Story

Chapter 5 — Back to the South

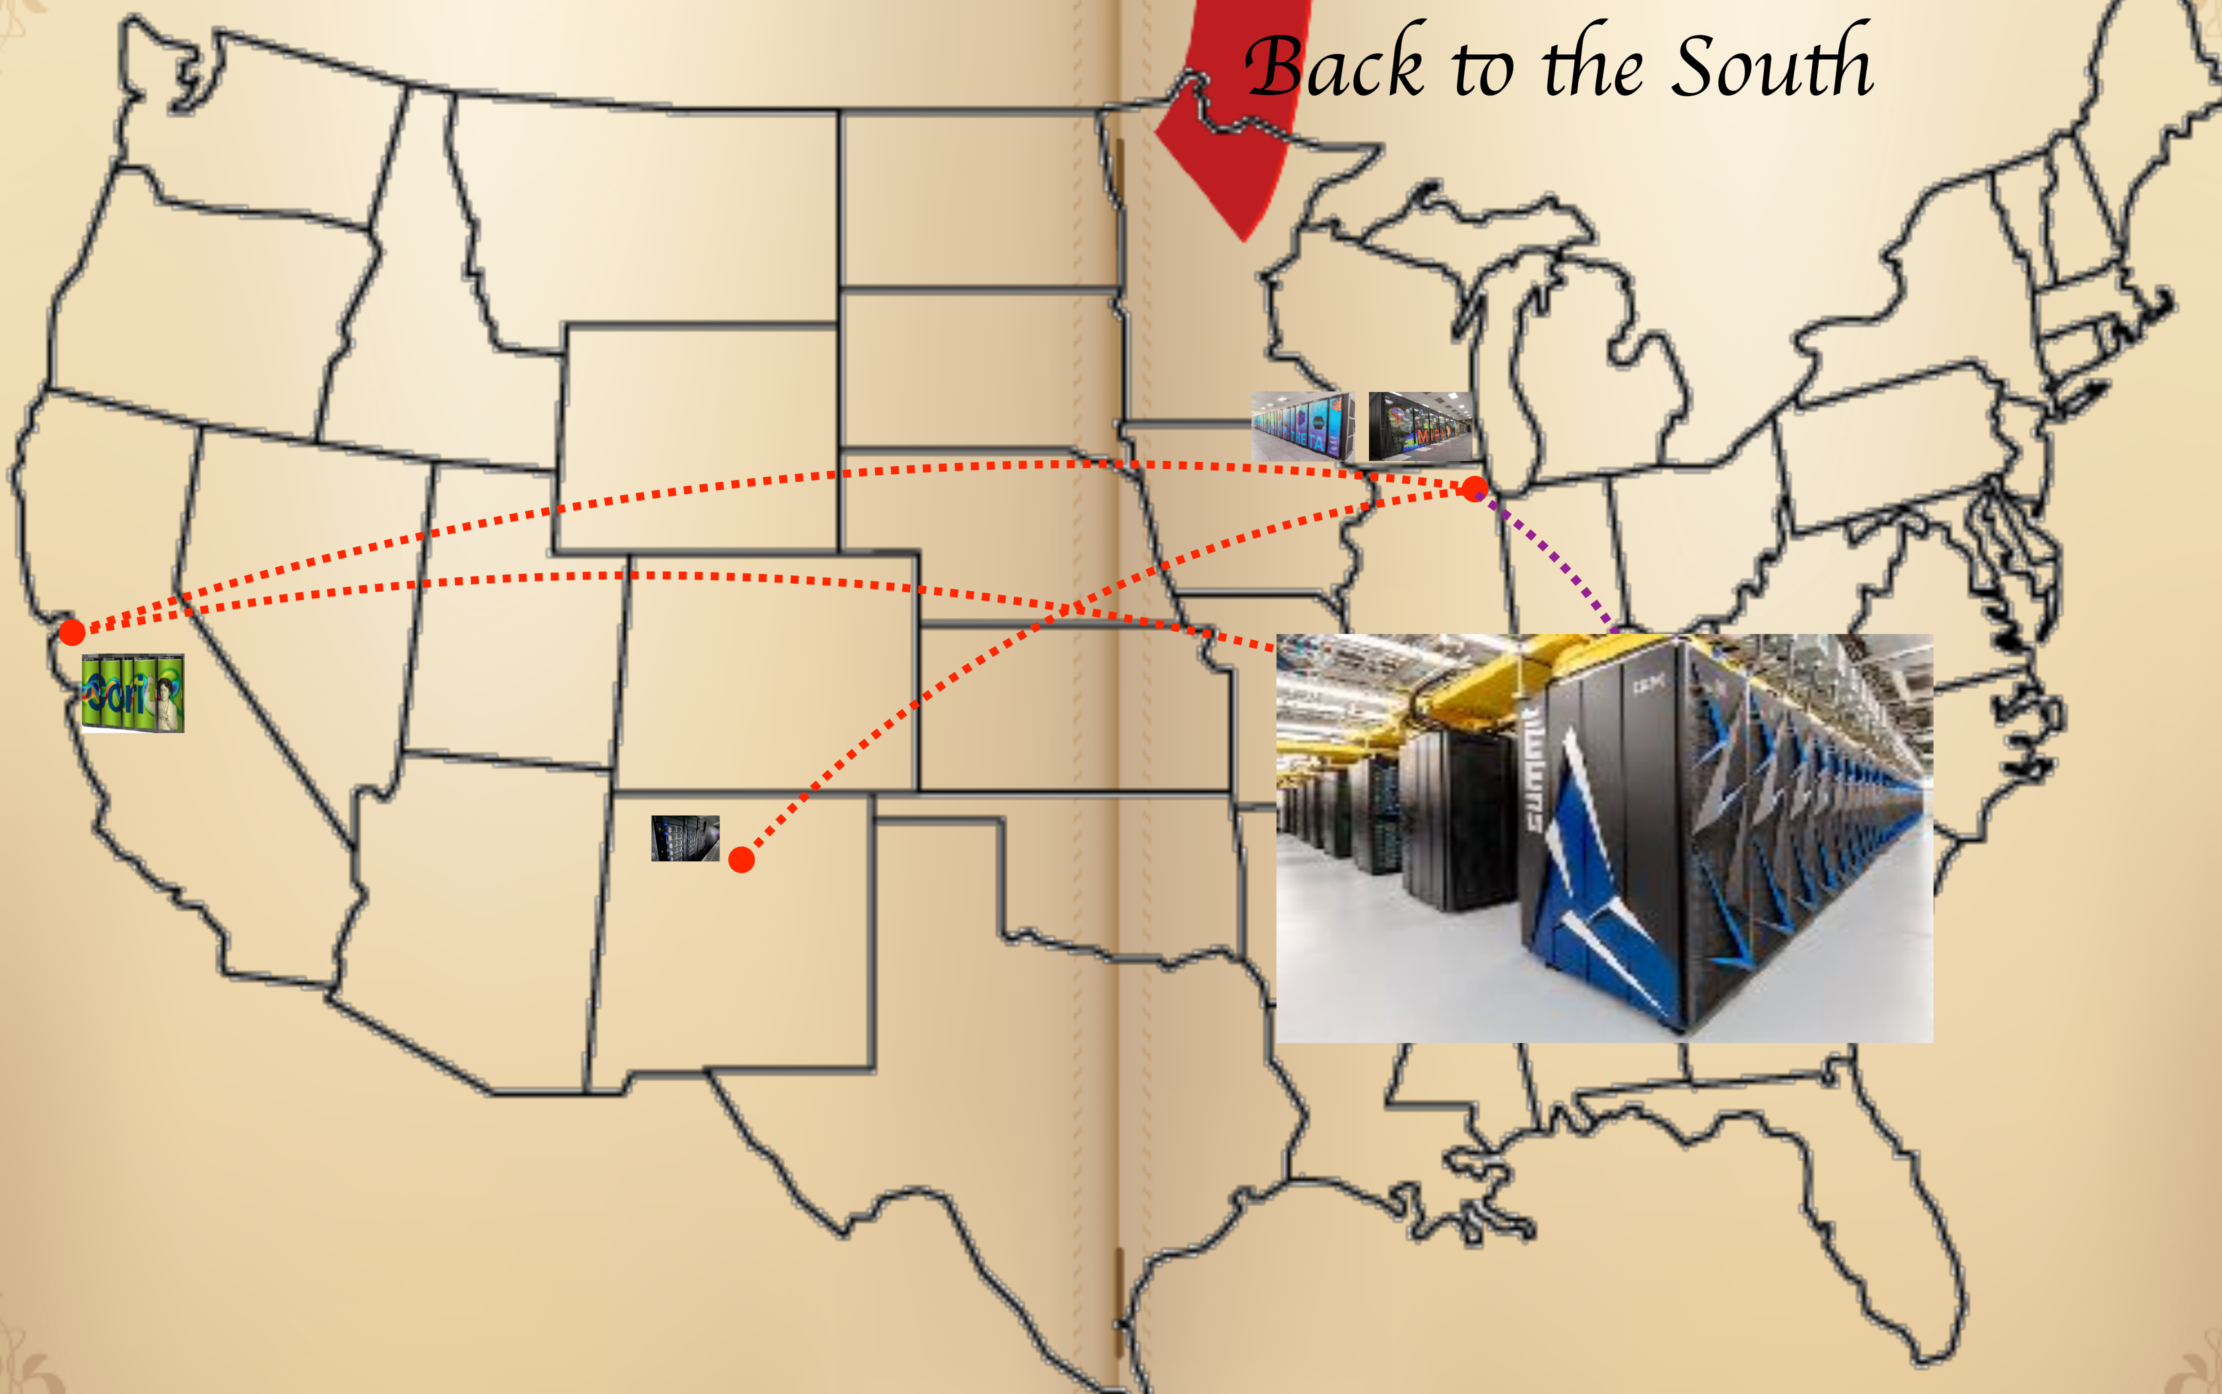

The HACC Story

Chapter 5 — Back to the South

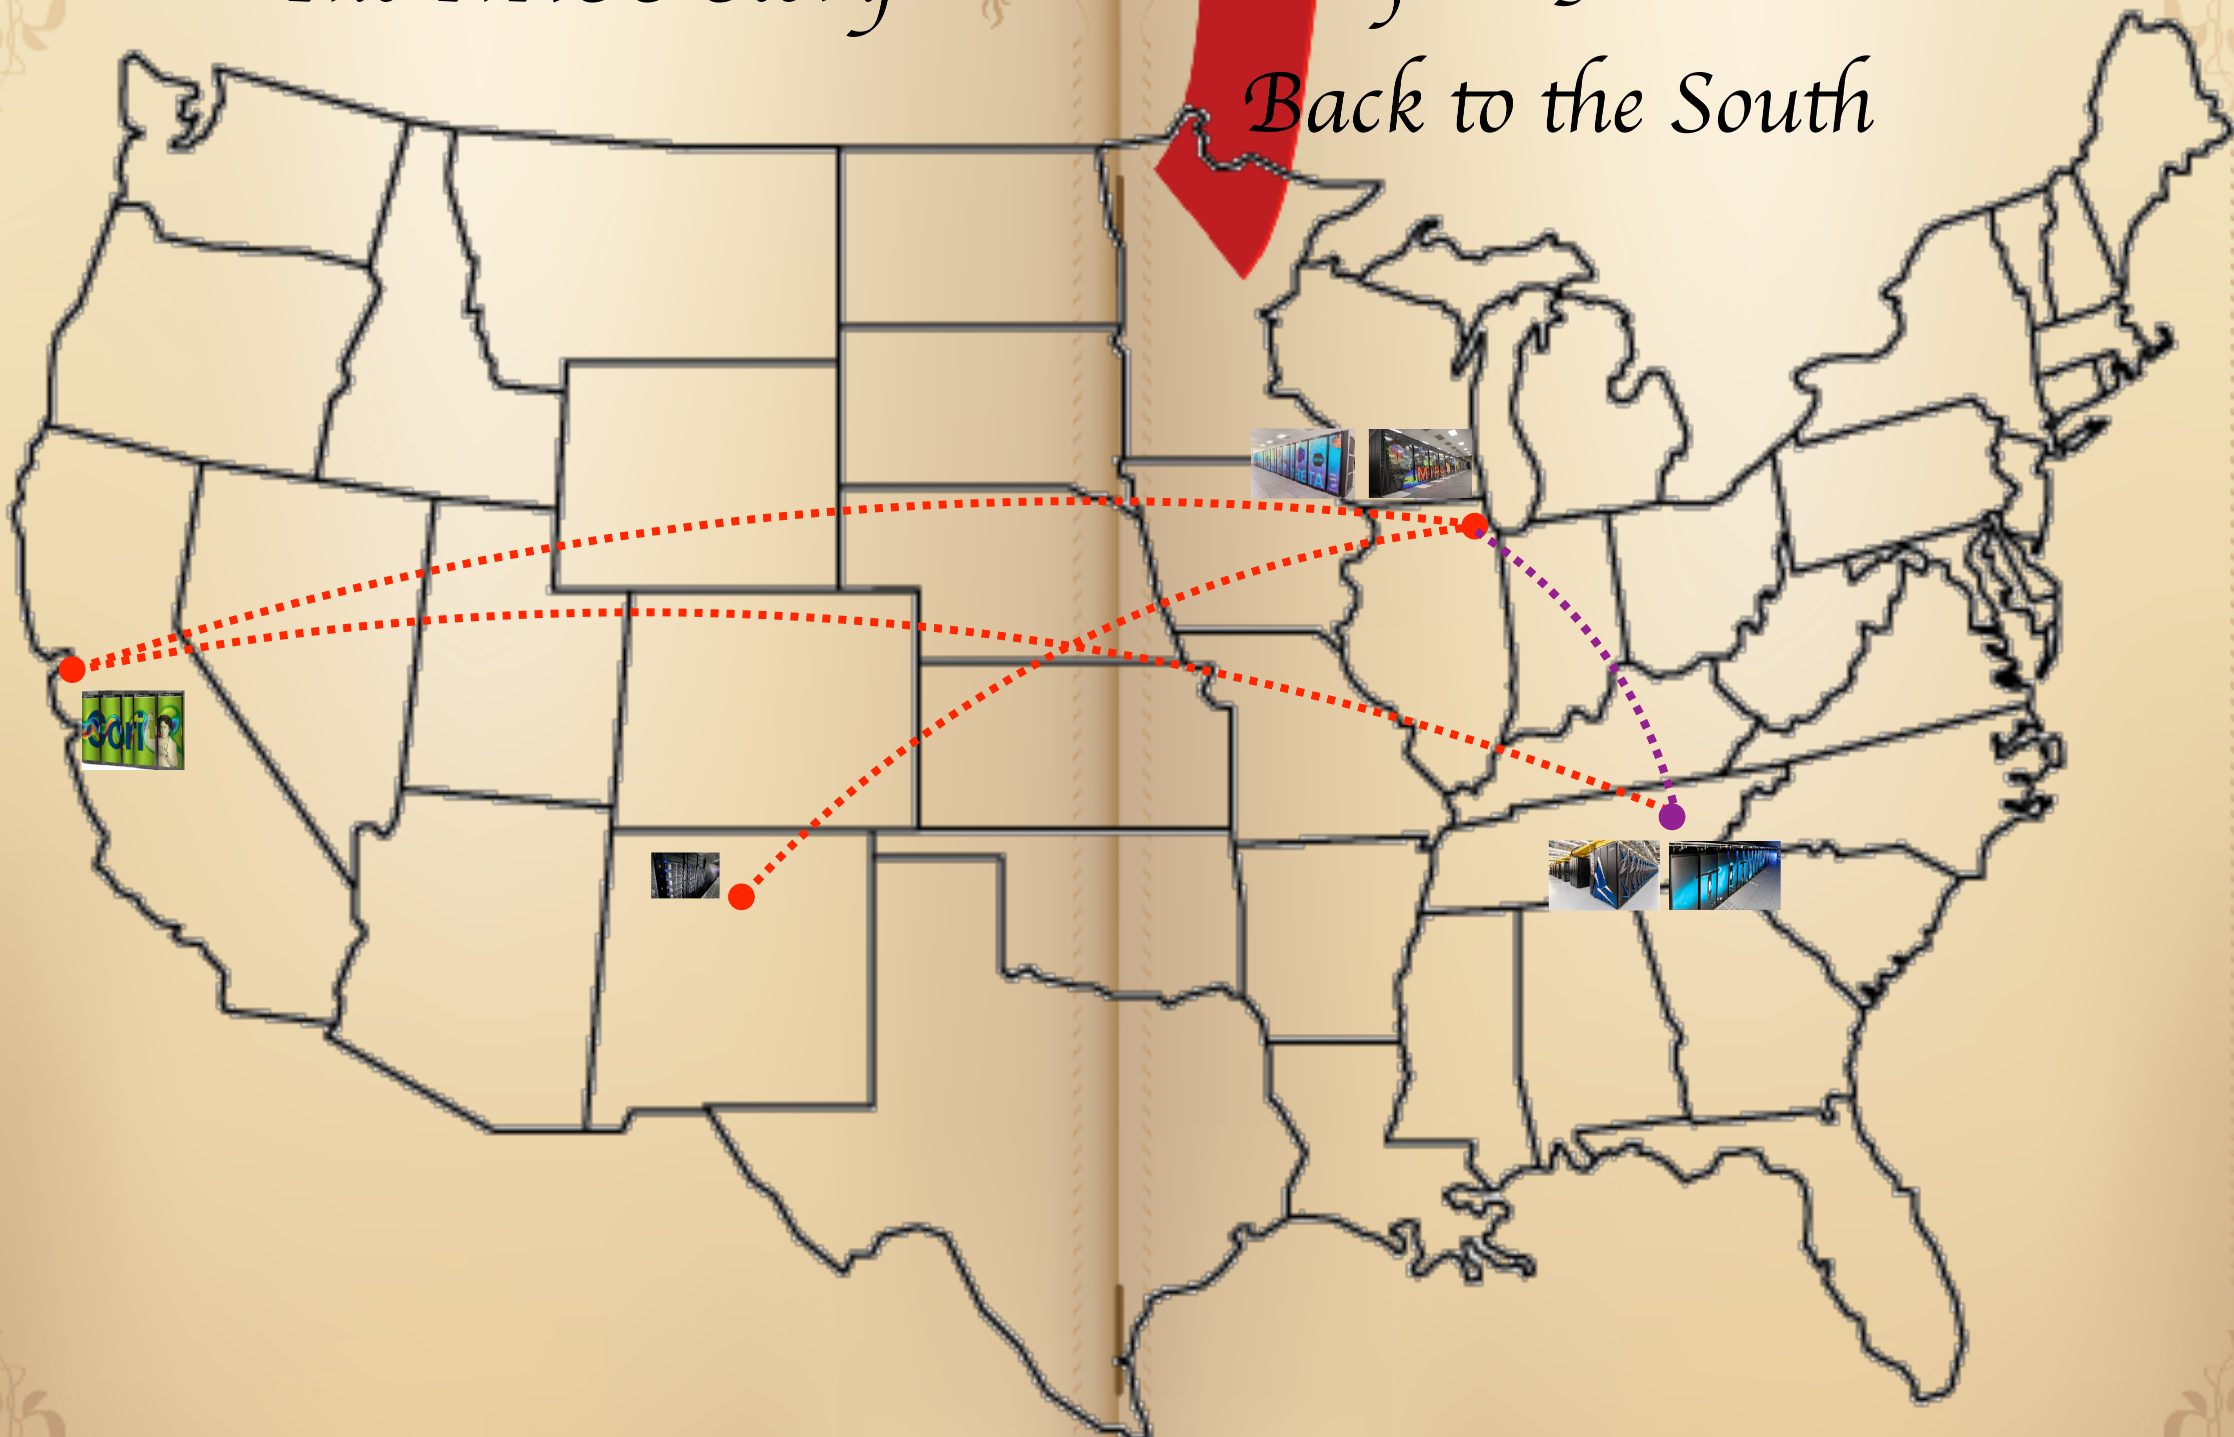

The HACC Story

Chapter 5 — Back to the South

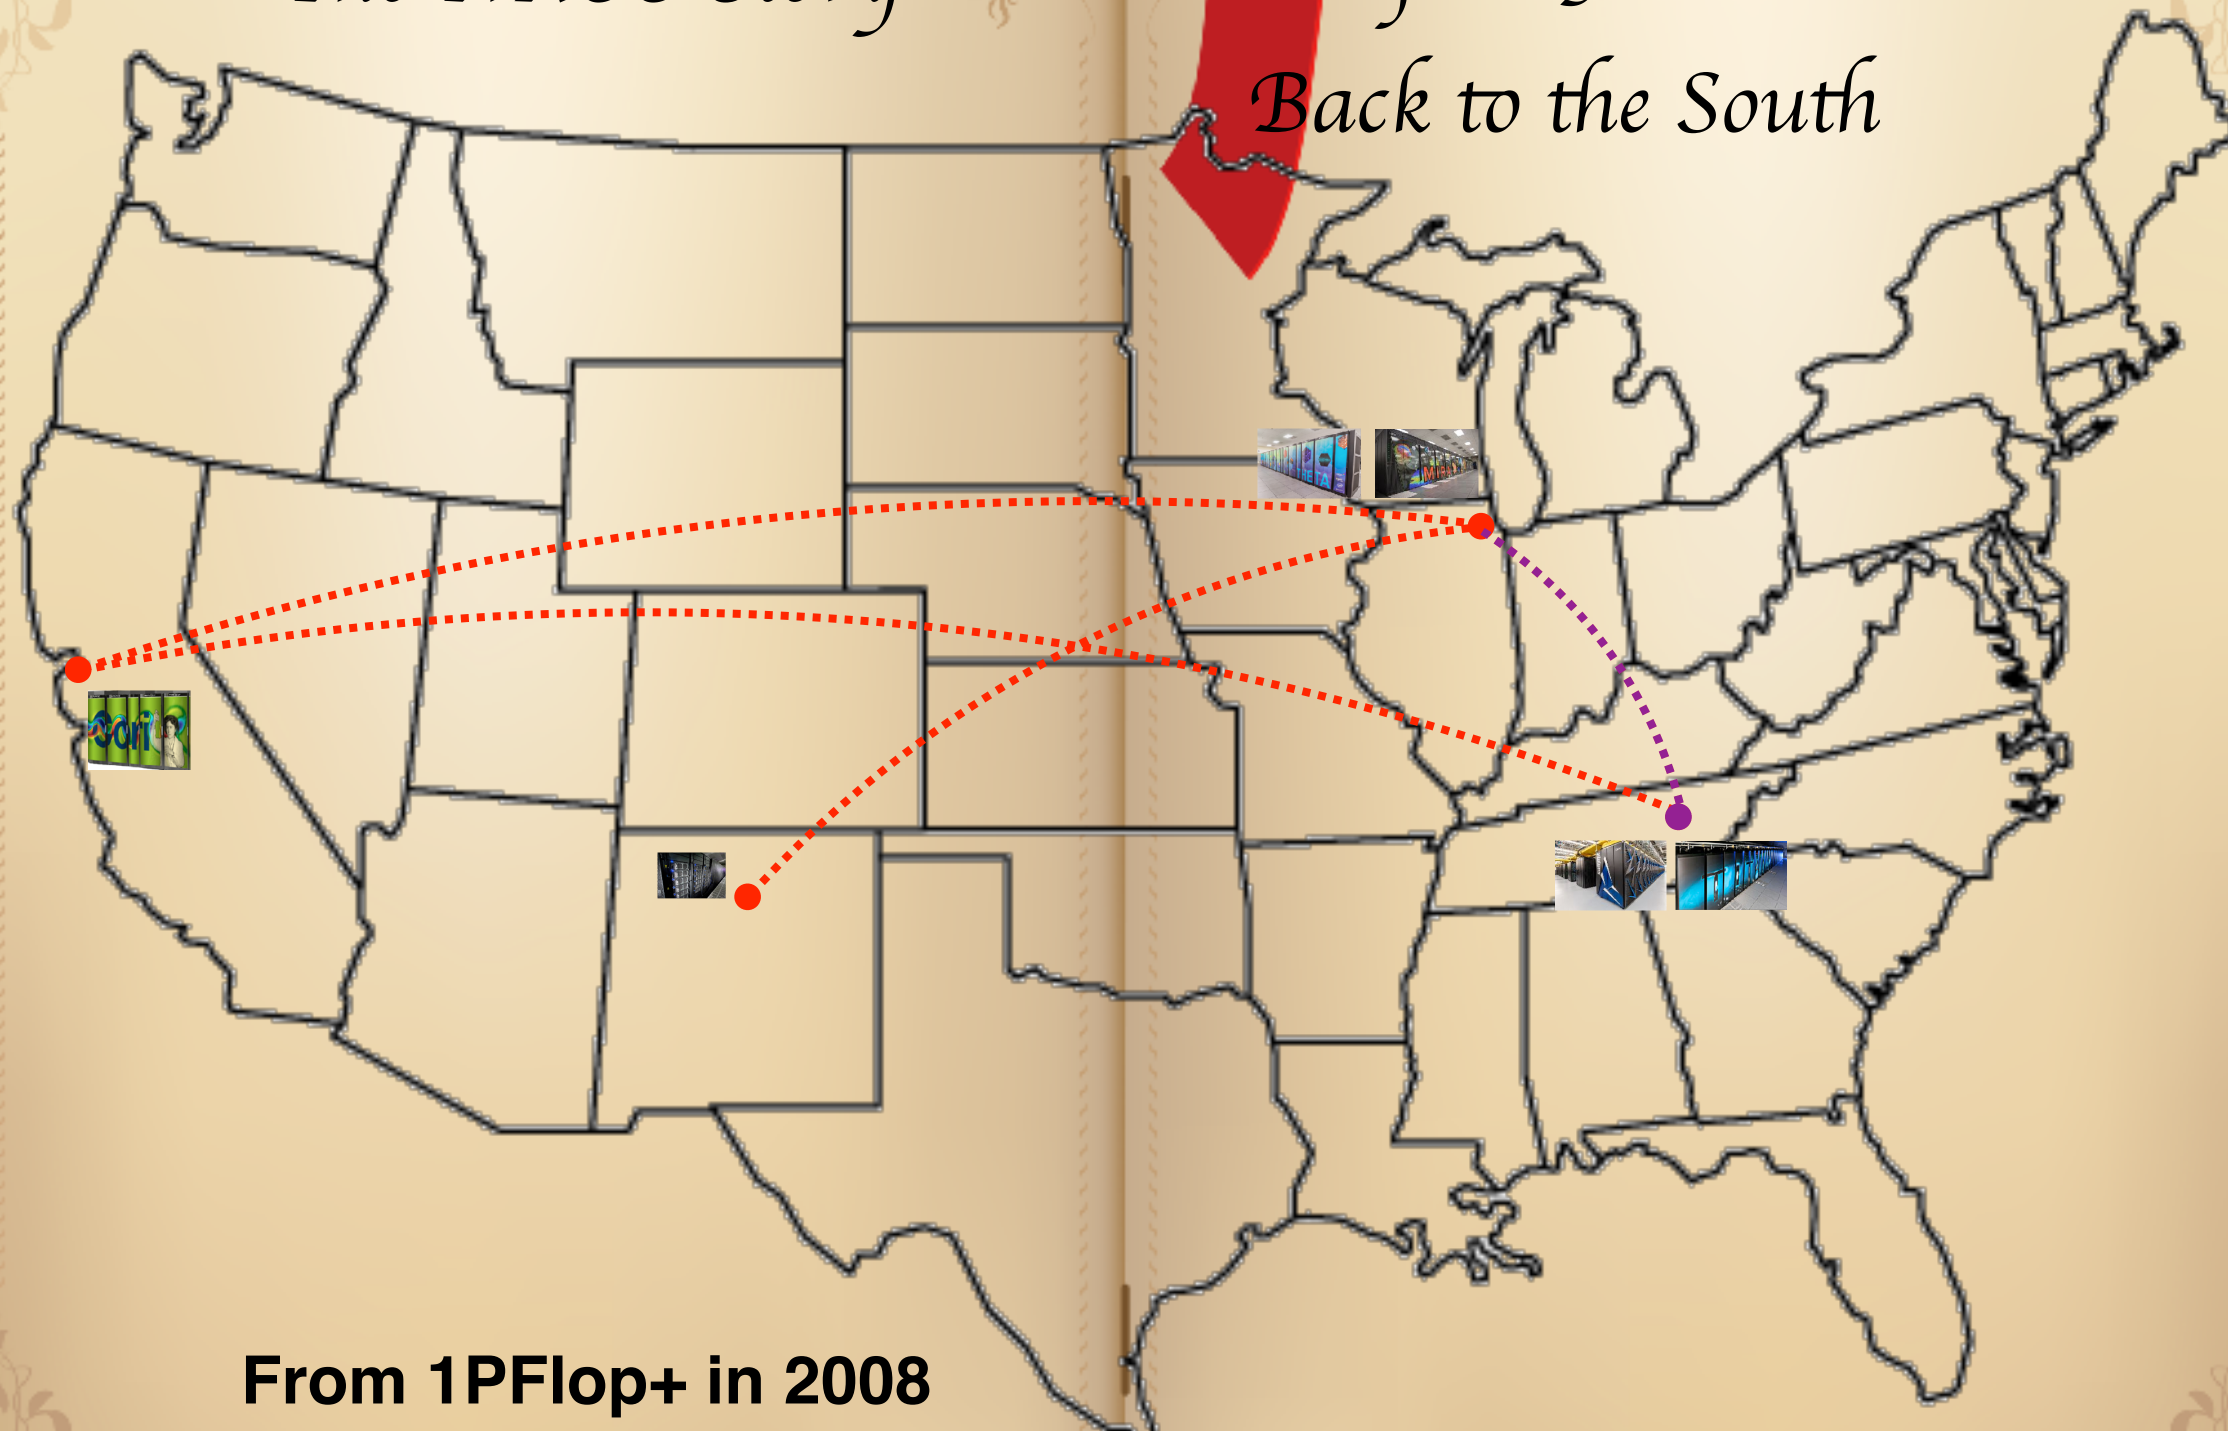

From 1PFlop+ in 2008

The HACC Story

Chapter 5 — Back to the South

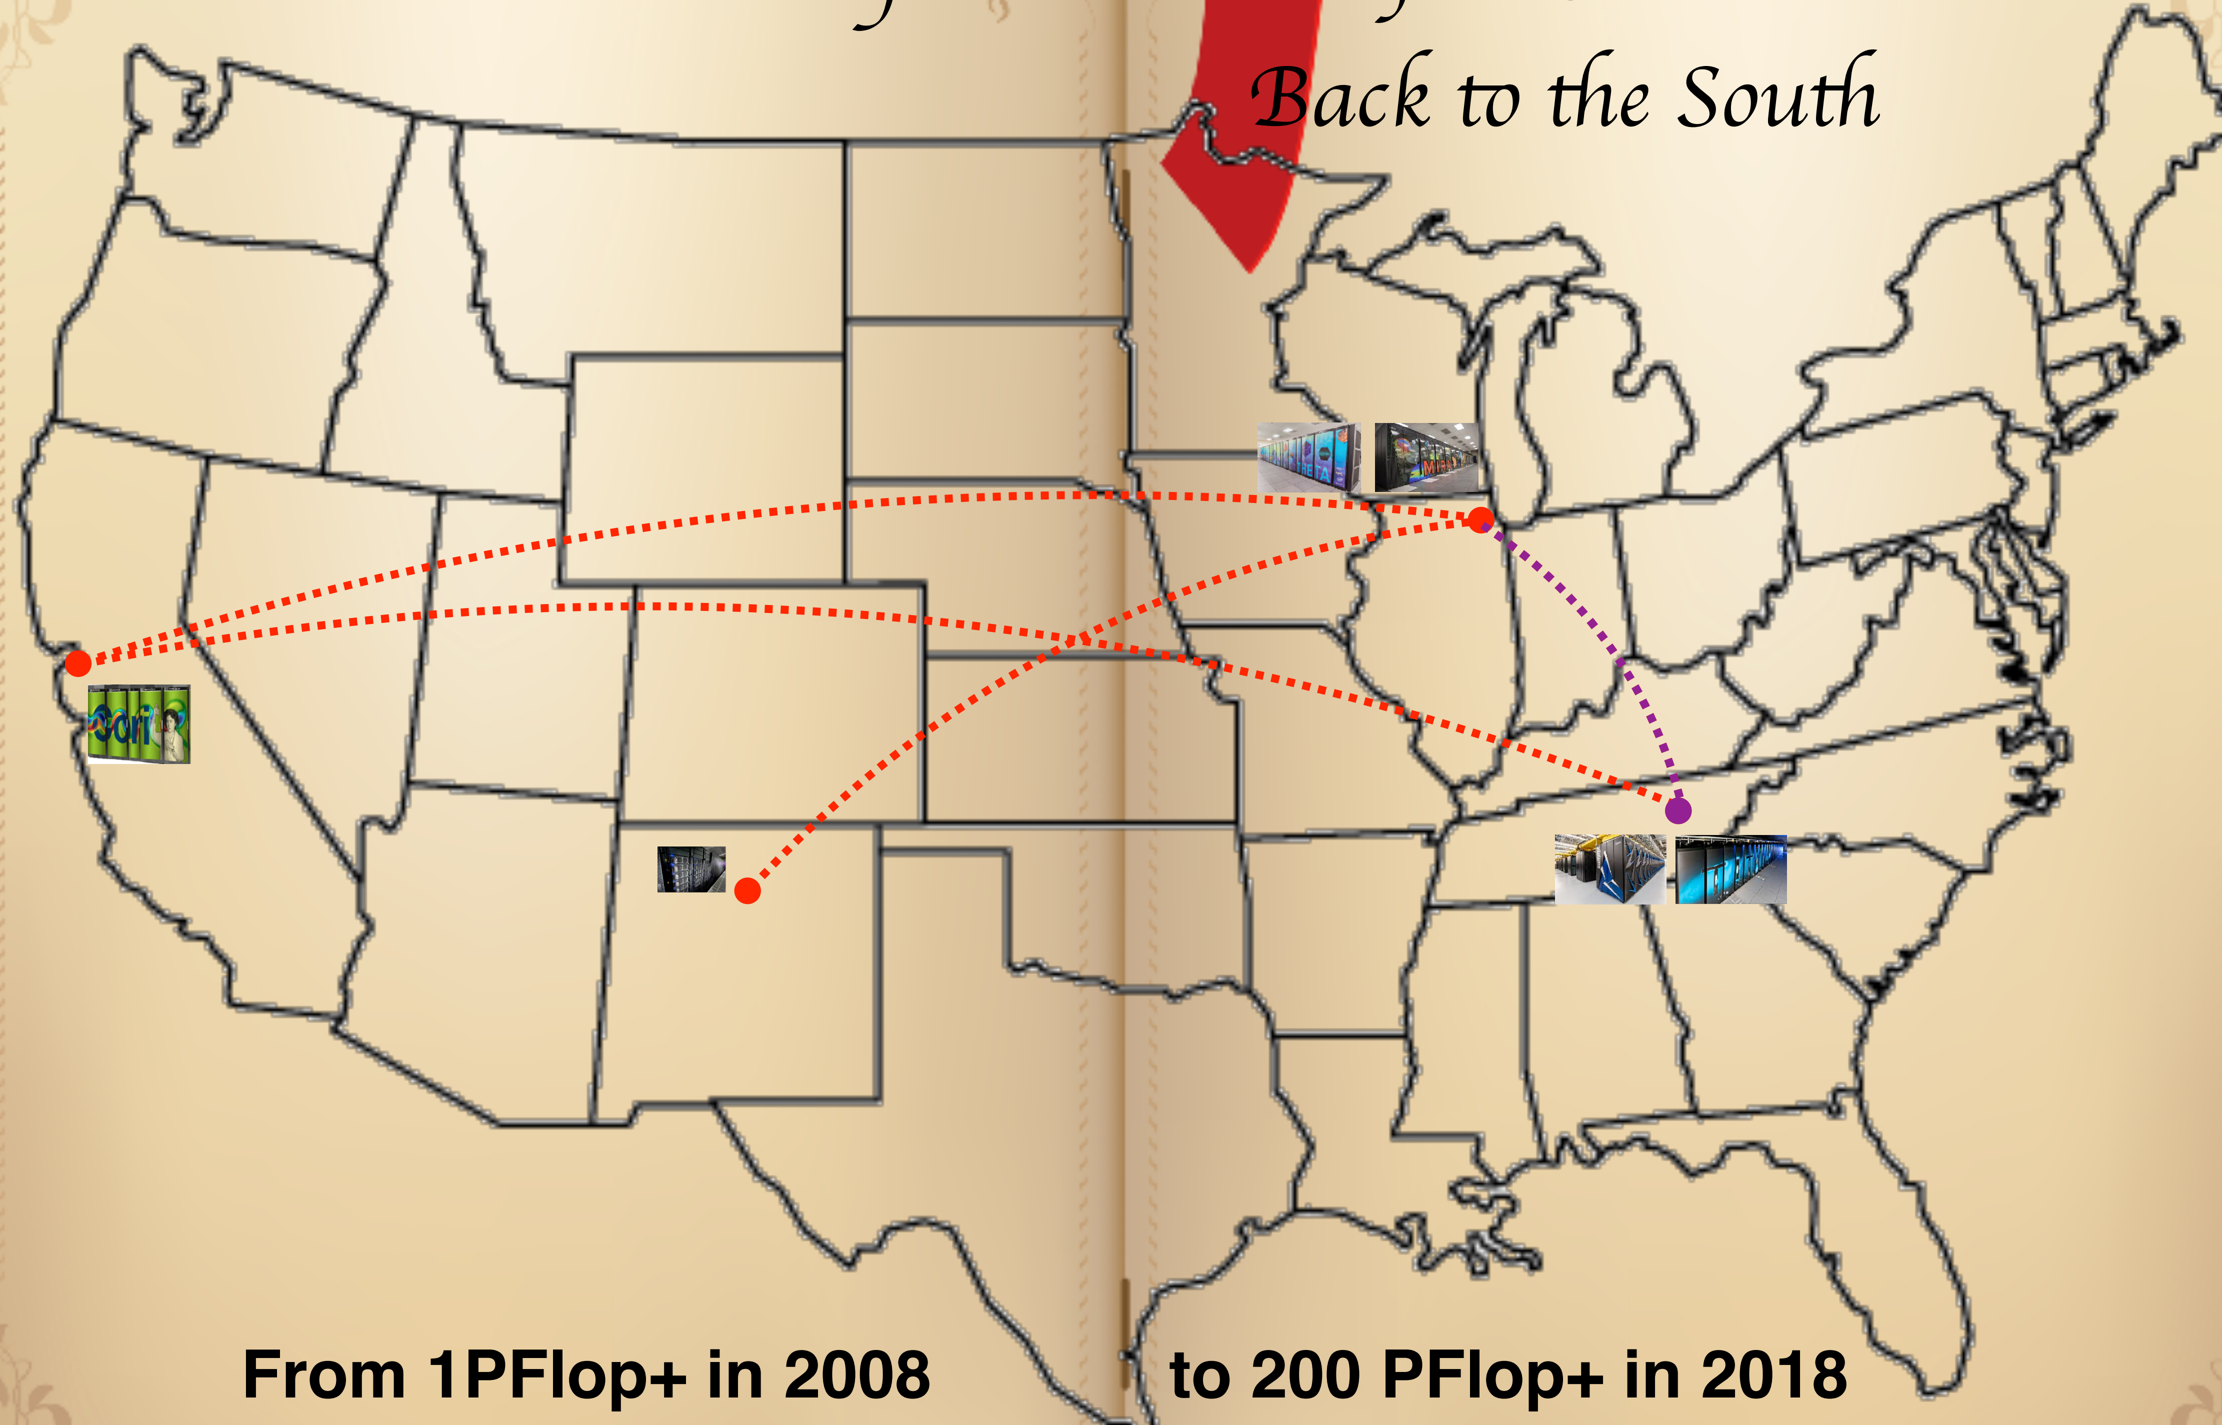

From 1PFlop+ in 2008 to 200 PFlop+ in 2018

Qo’noS simulation carried out with HACC on Summit

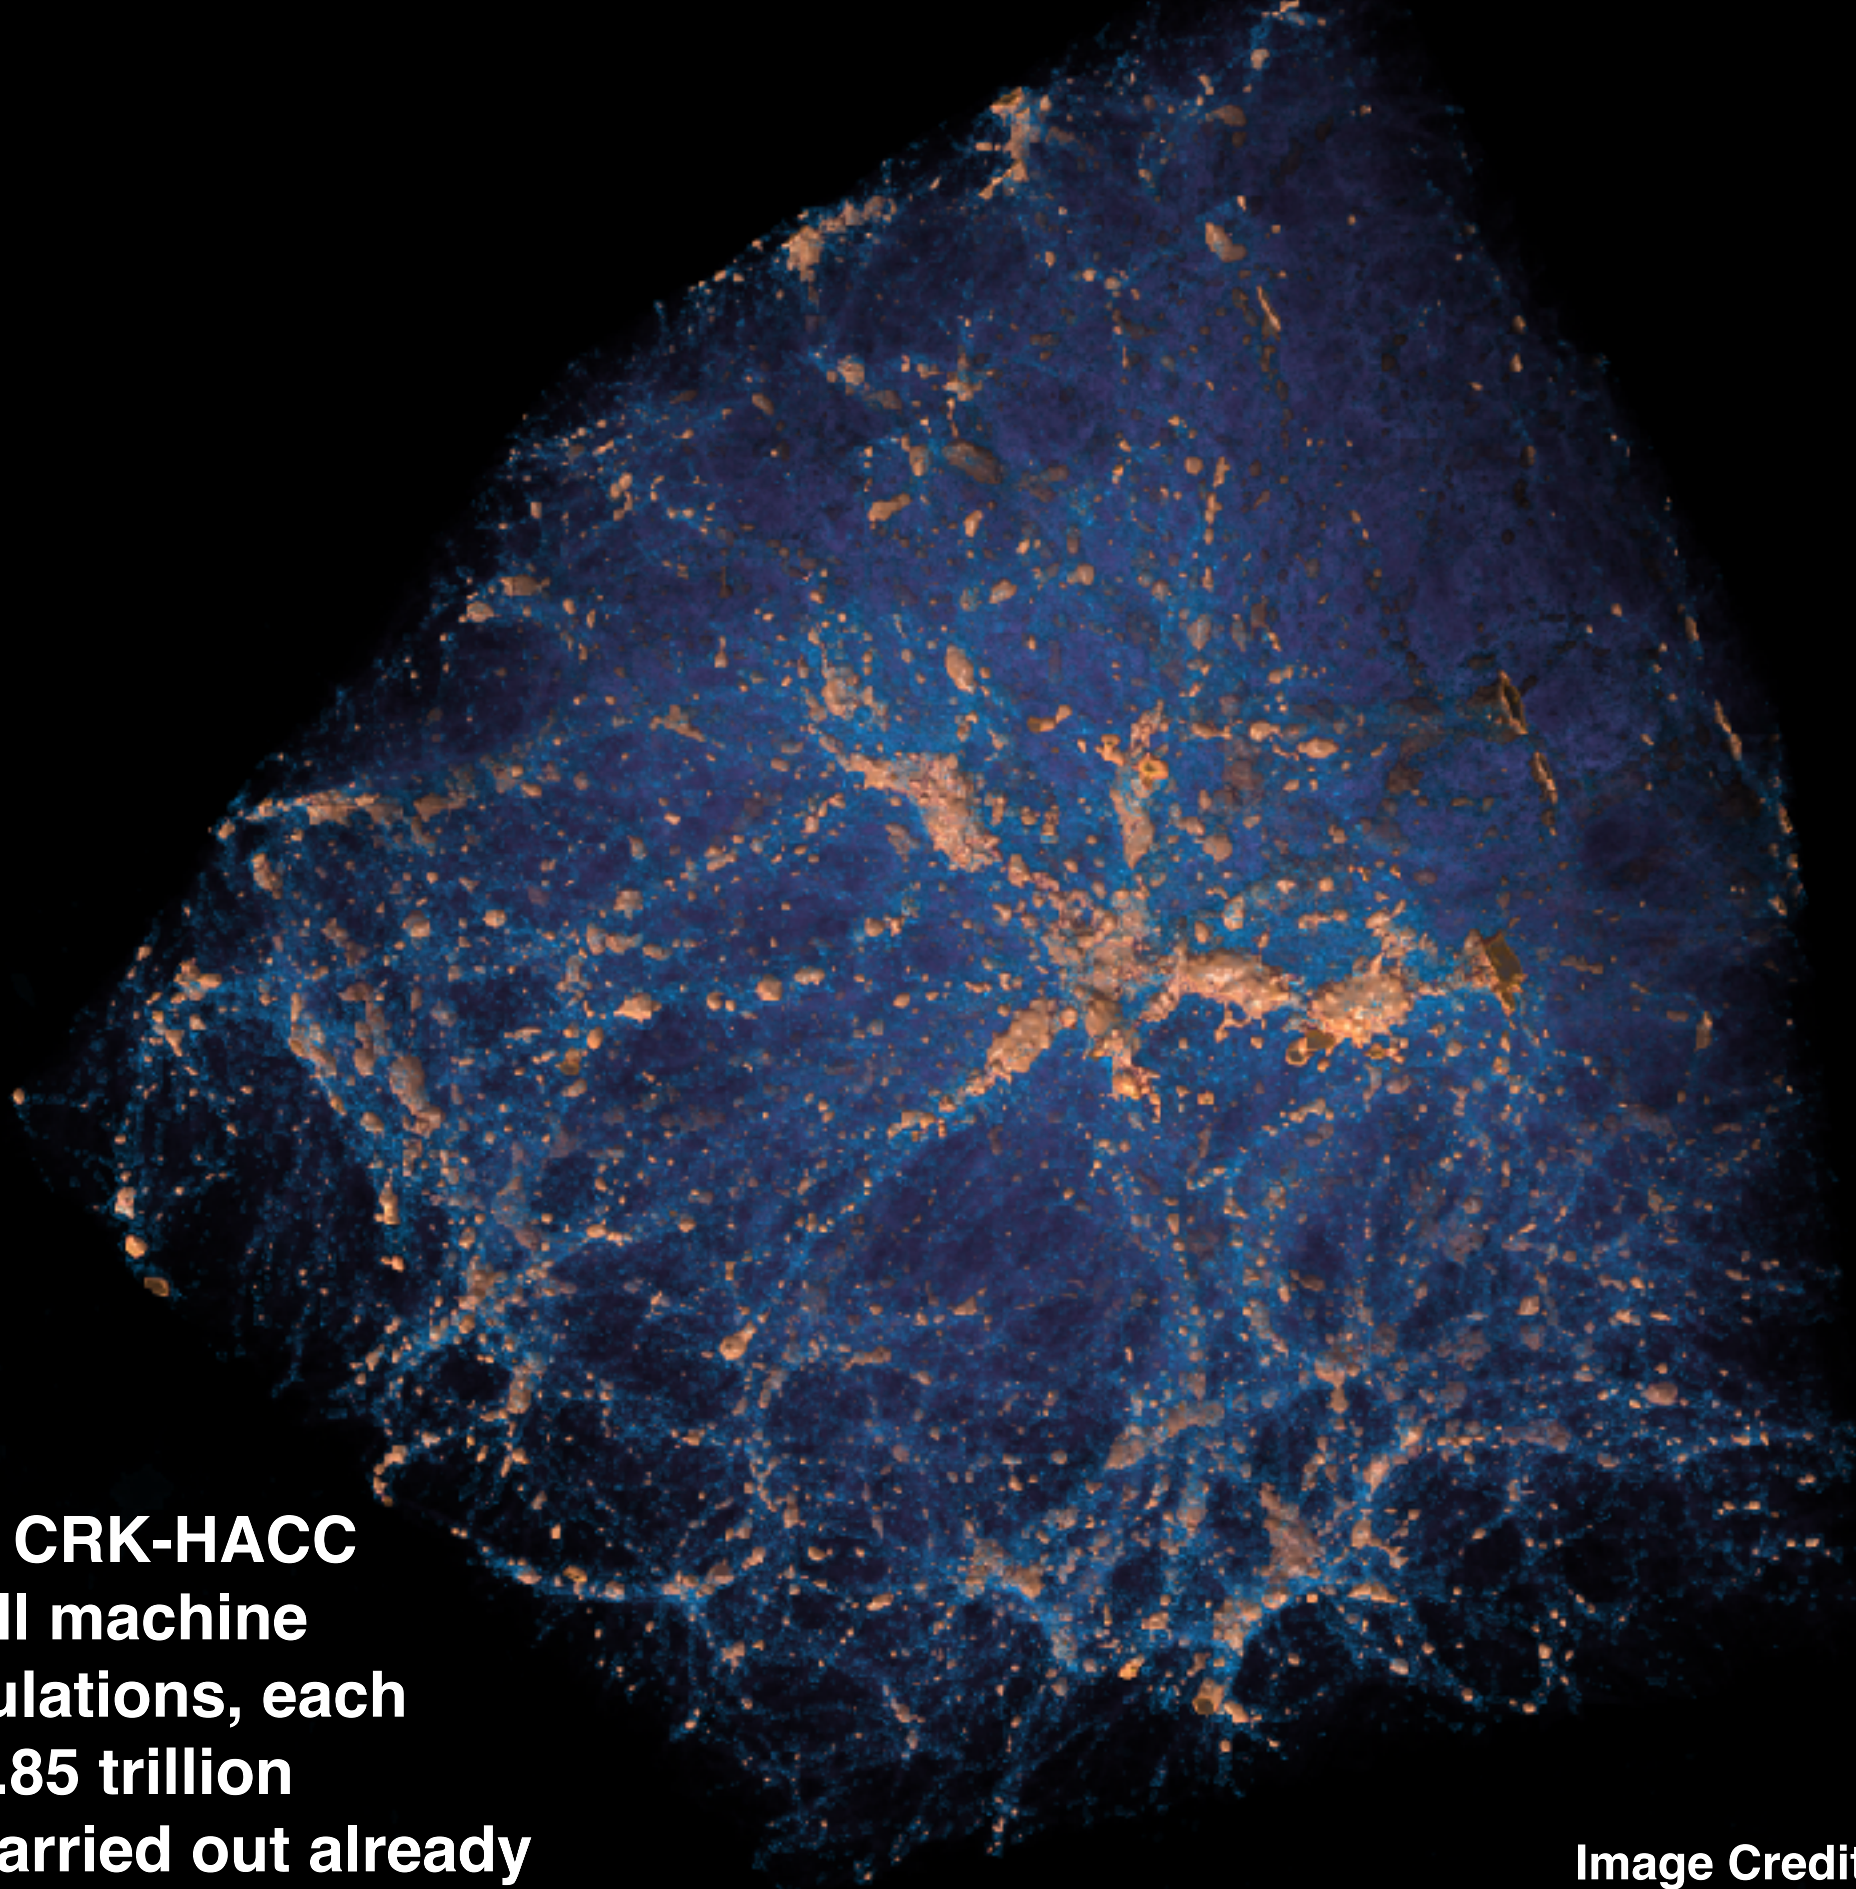

- HACC and CRK-HACC scale to full machine
- Three simulations, each evolving 1.85 trillion particles carried out already

Image Credit: S. Rizzi, J. Insley

HACC in Pictures

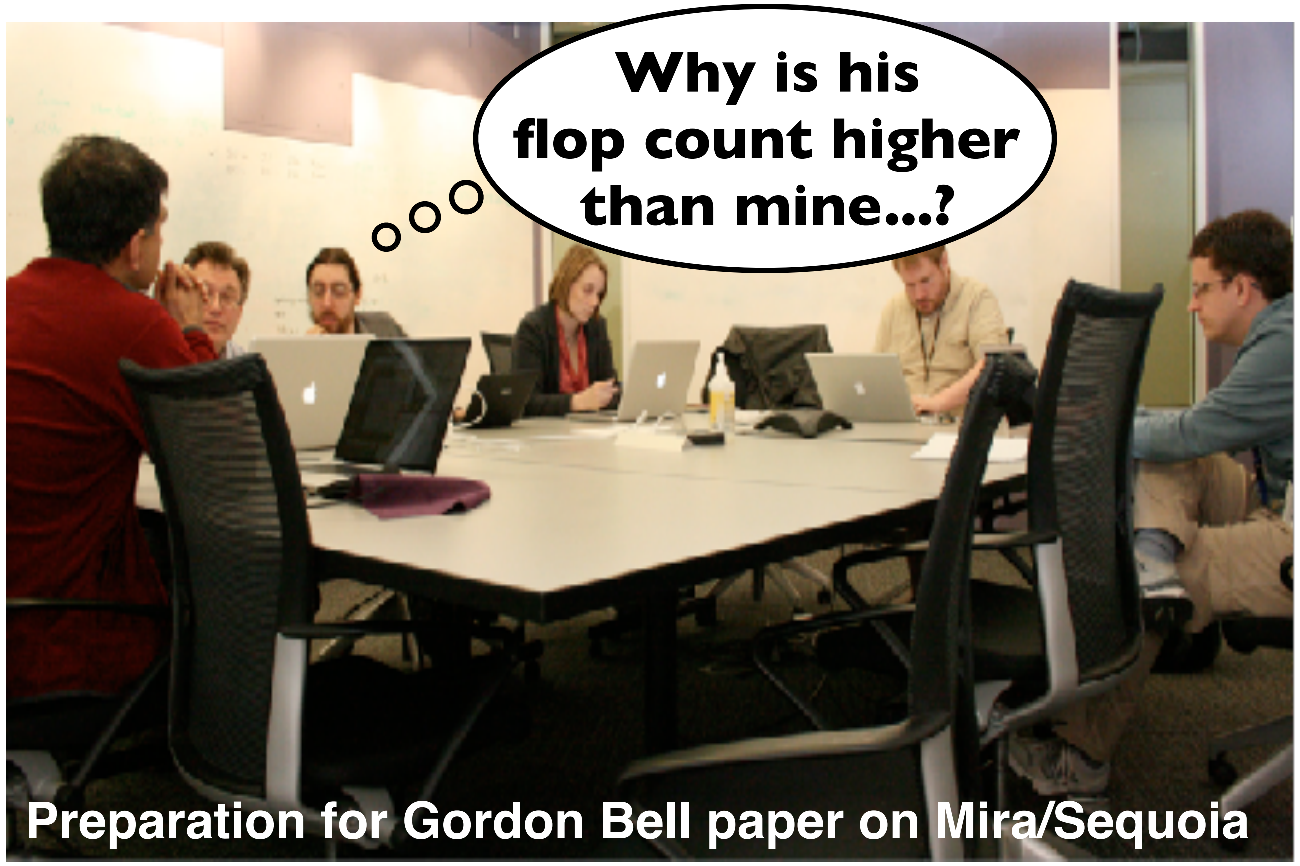


Edison


MC2


SUMMIT
Edison
TITAN
MIRA
5
HACC

HACC Science

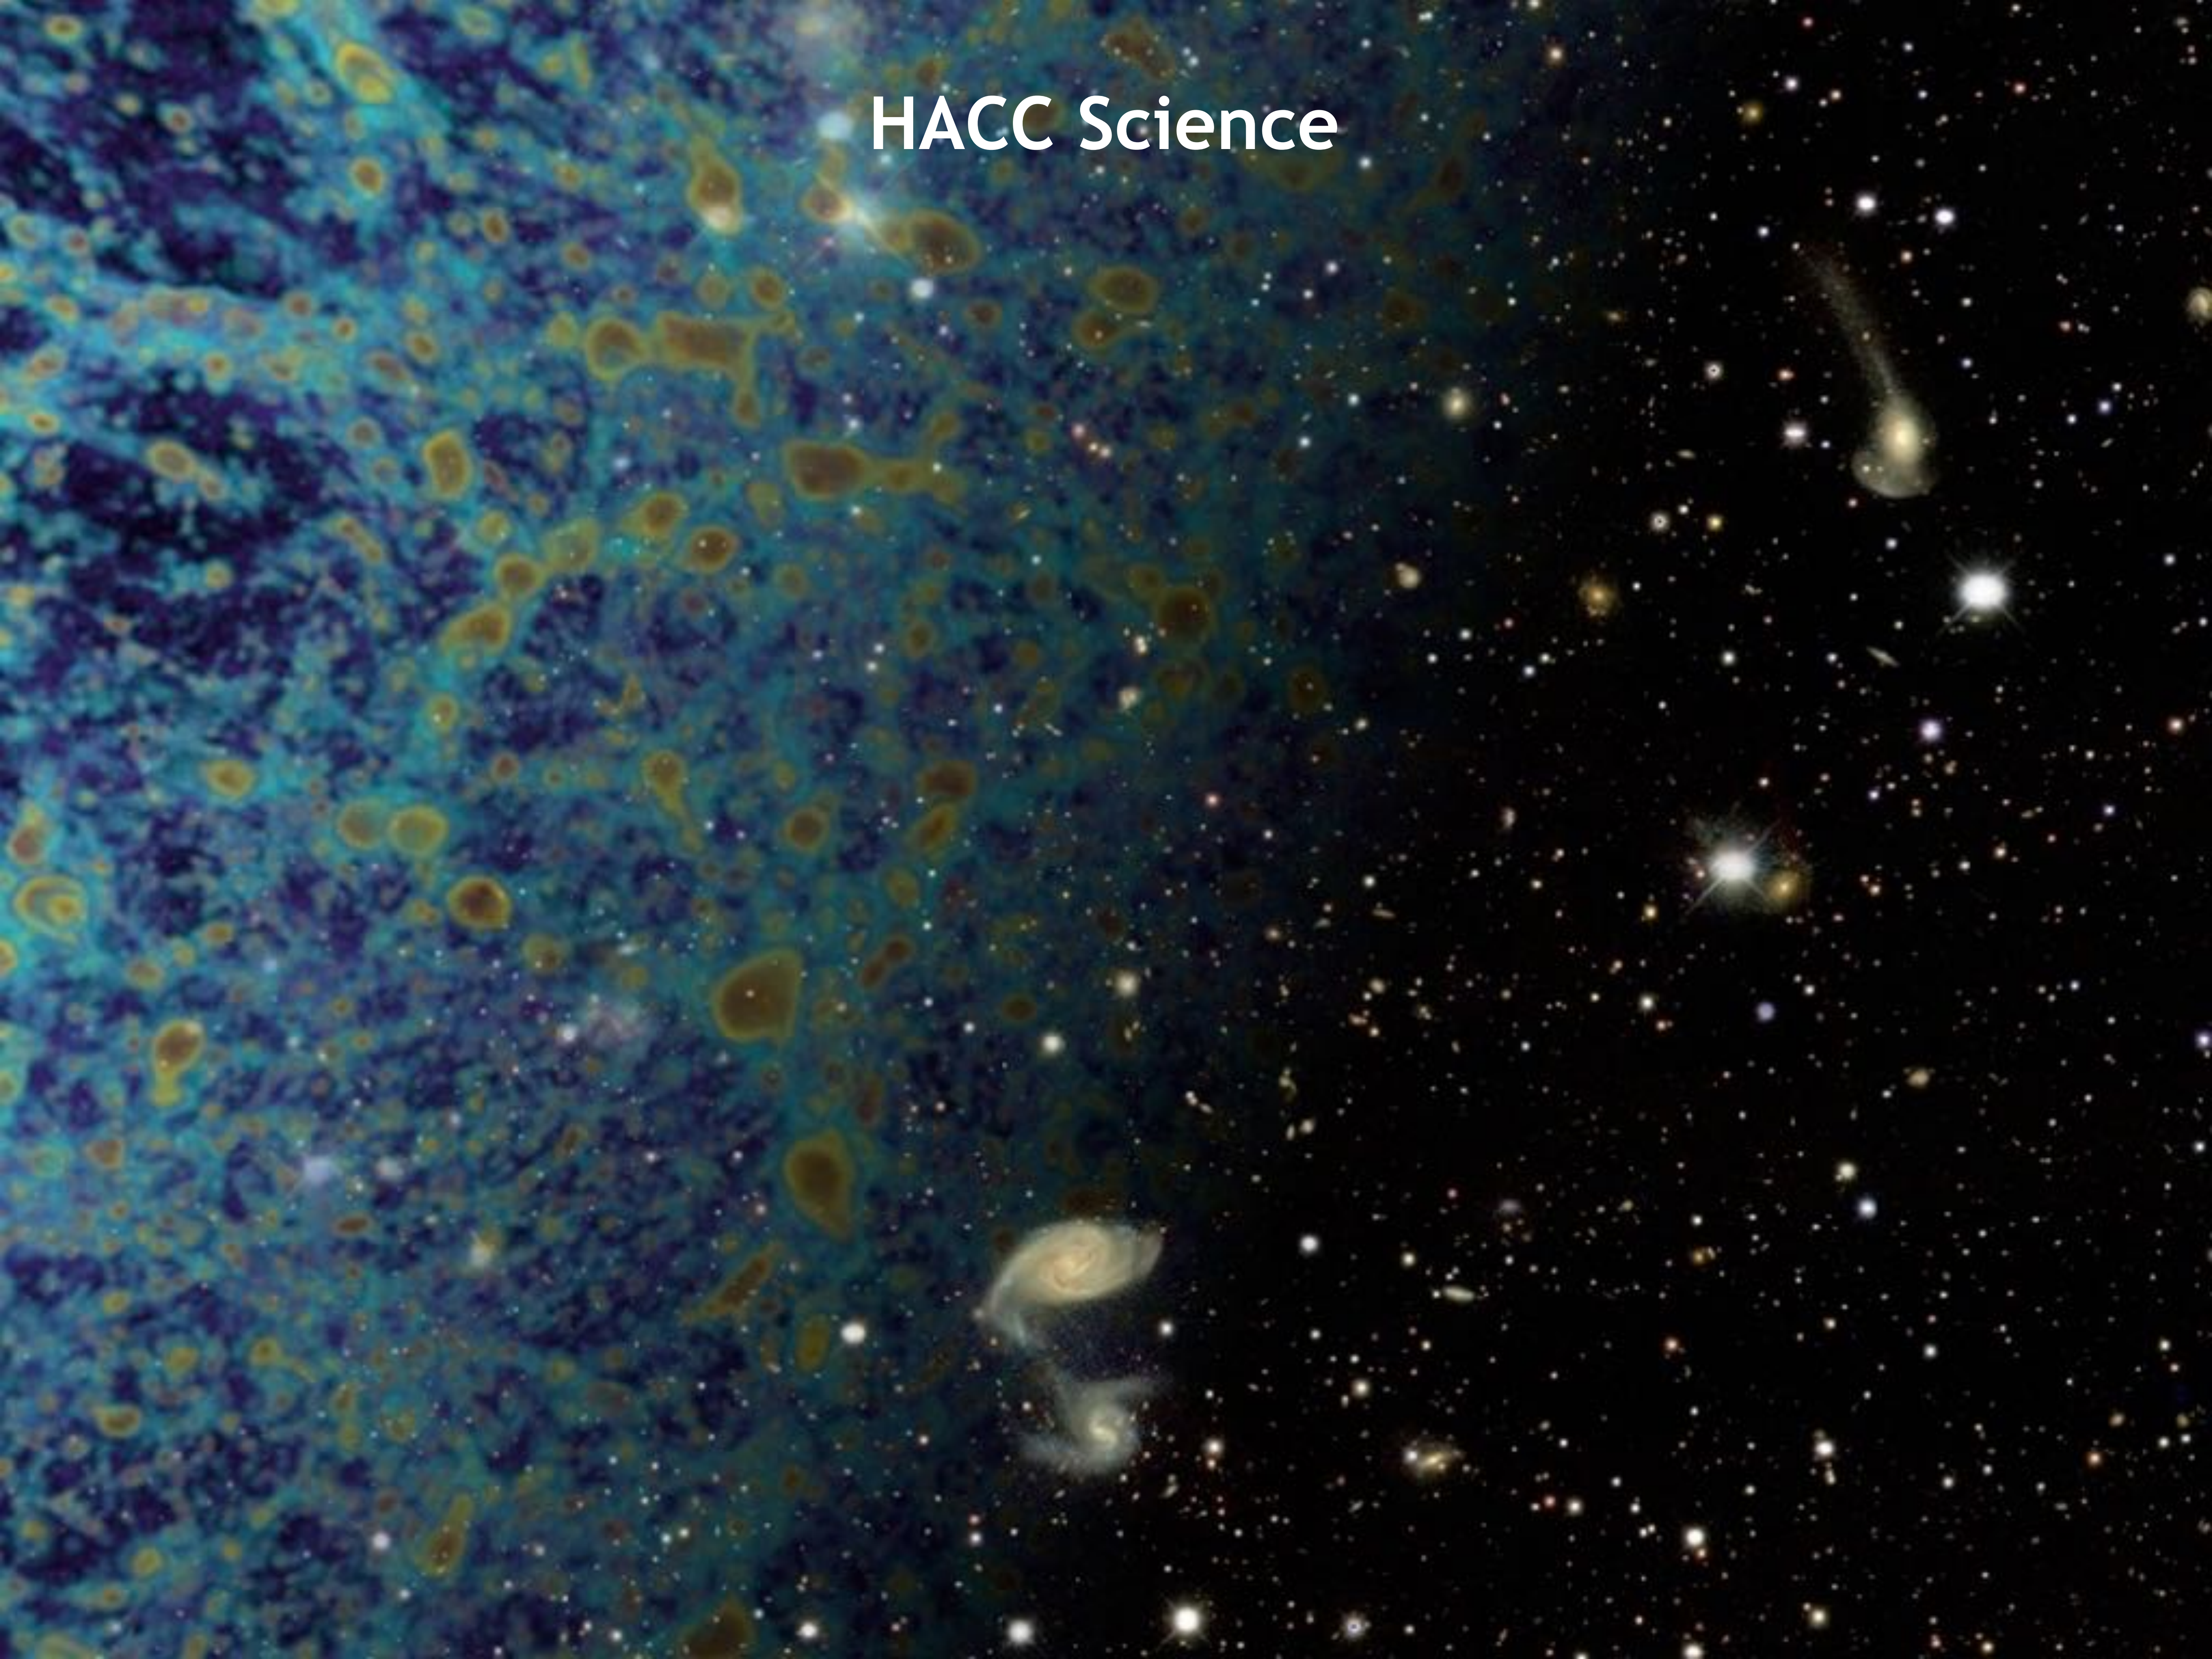

HACC Science

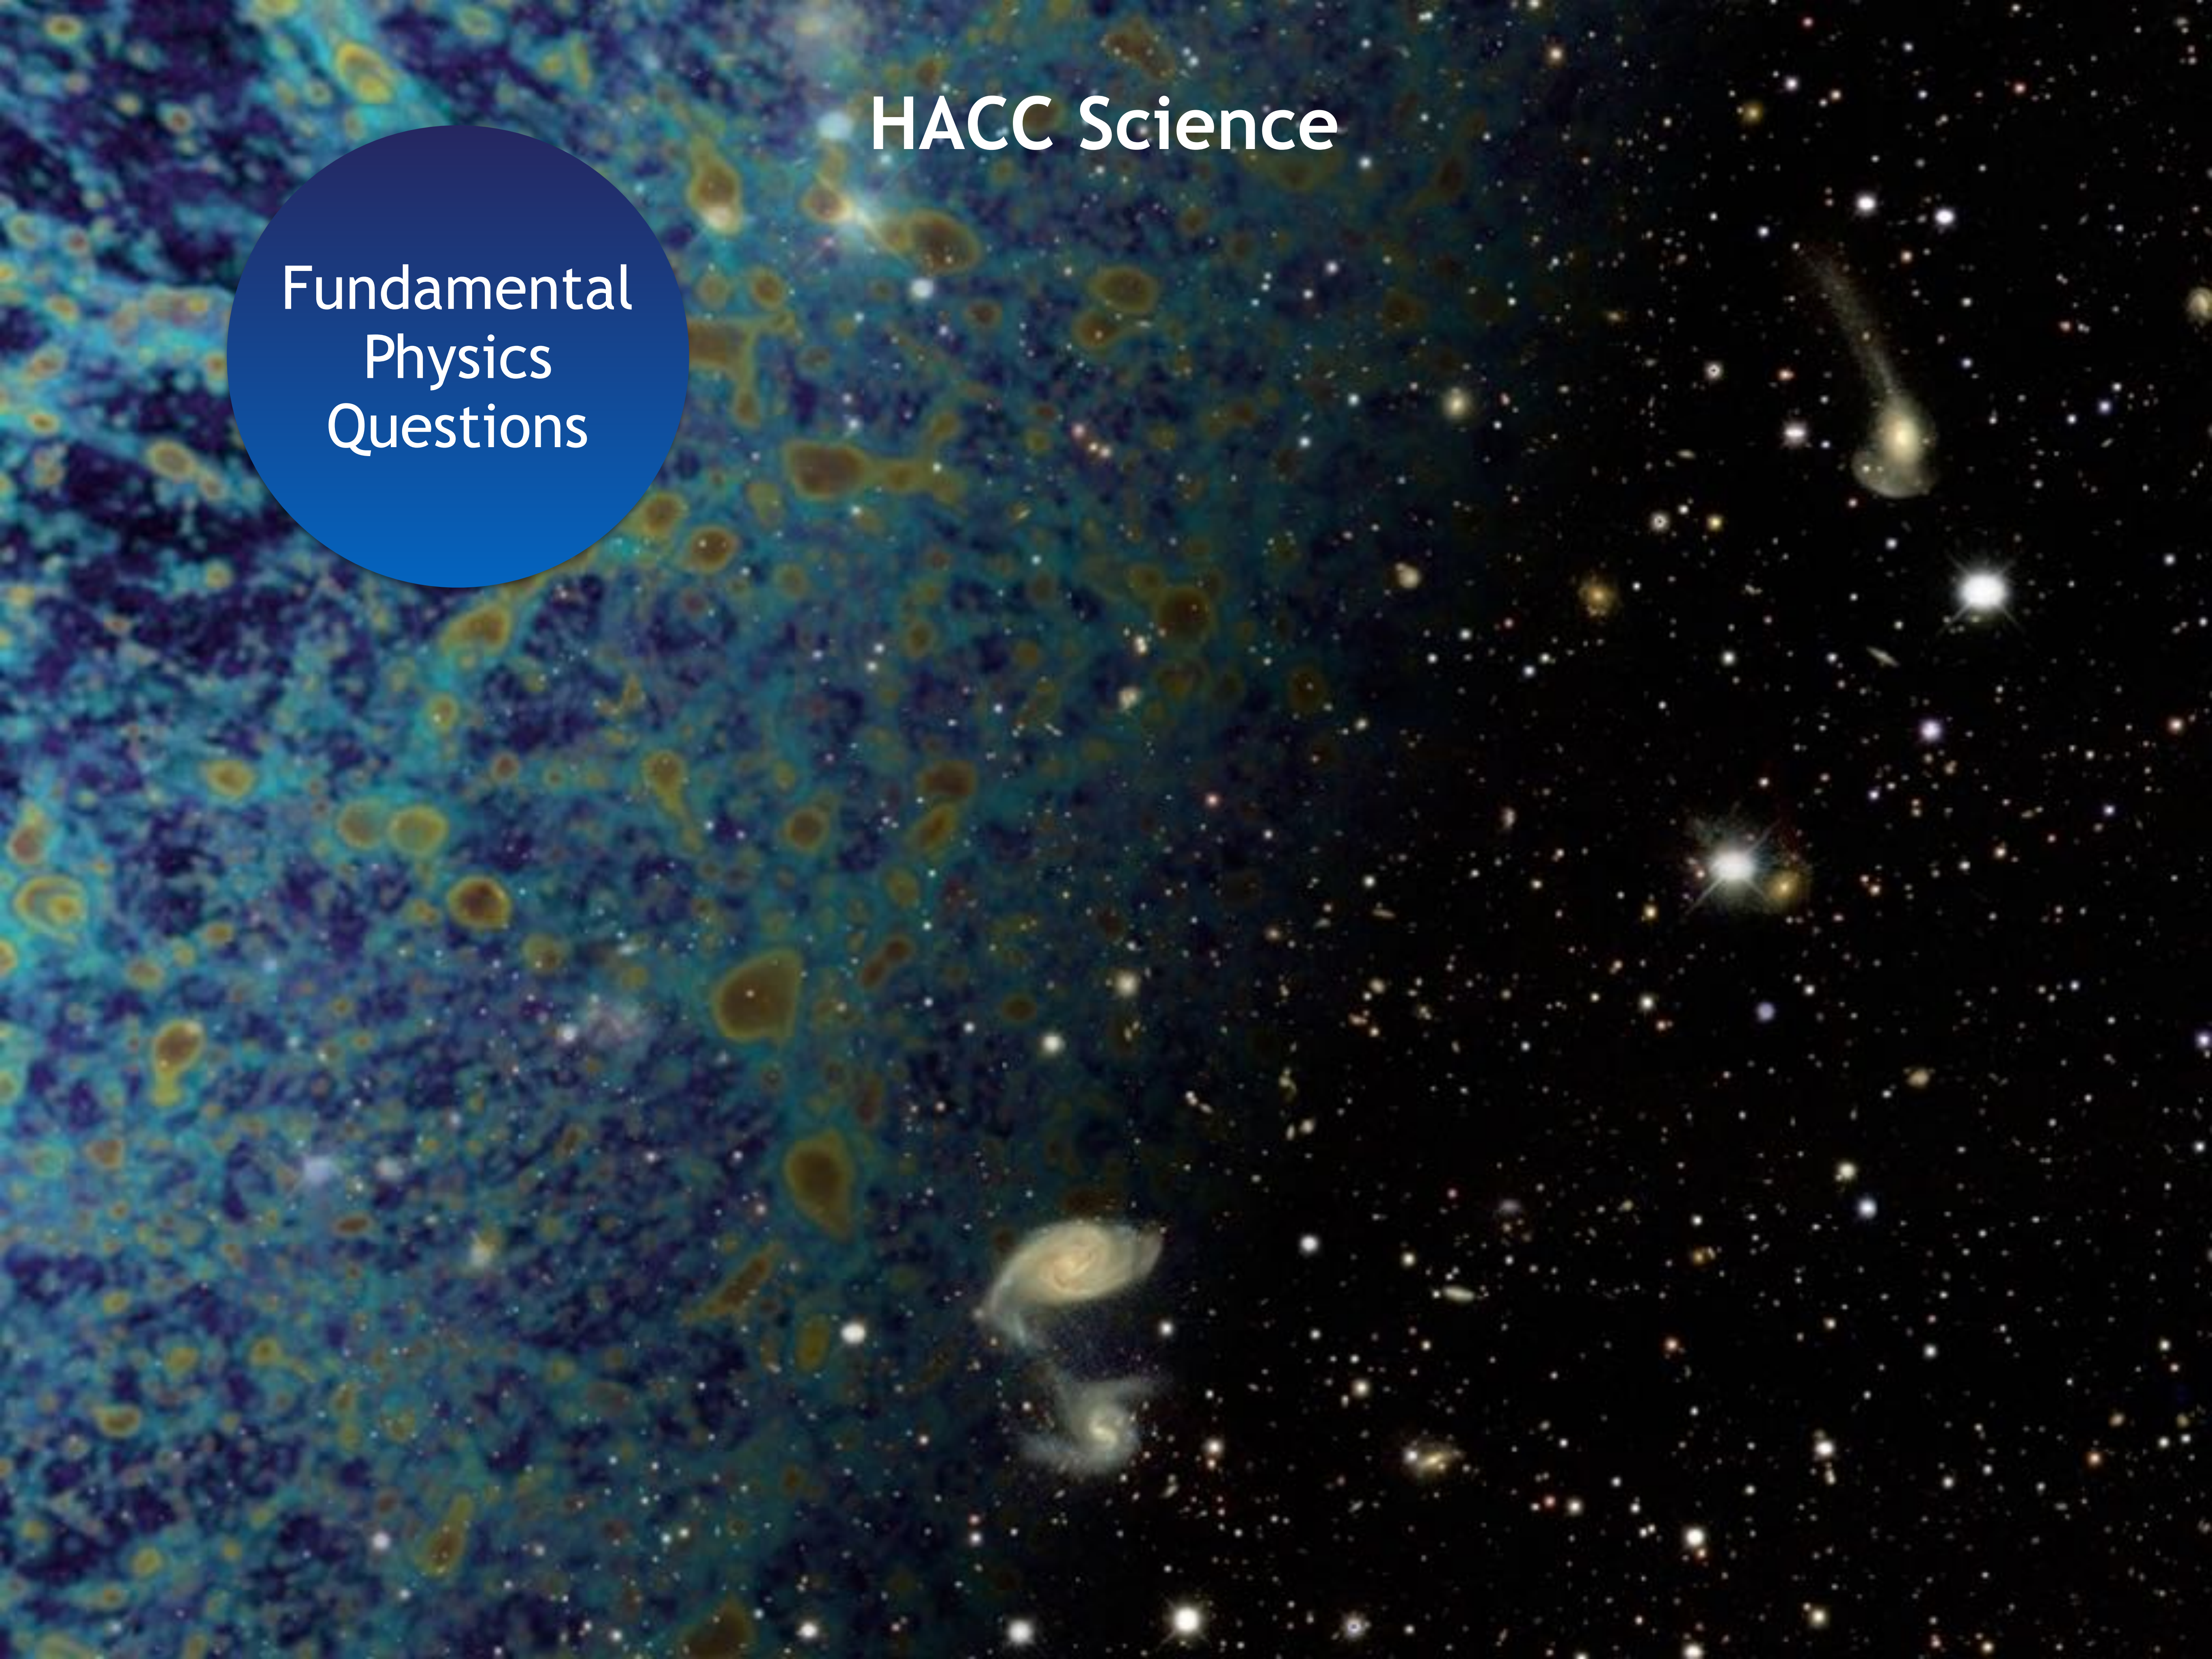

HACC Science

Fundamental Physics Questions

New Probes to Understand the Universe

HACC Science

Fundamental Physics Questions

New Probes to Understand the Universe

Survey Science

The Large Synoptic Survey Telescope

- 8.4-m mirror
- 37 billion stars and galaxies
- 10 year survey of the sky
- 10 million alerts, 15 TB of data, every night!
- First light ~2020, first science data in ~2022
- In the meantime: LSST DESC (Dark Energy Science Collaboration) is generating simulations as close to the real data as possible

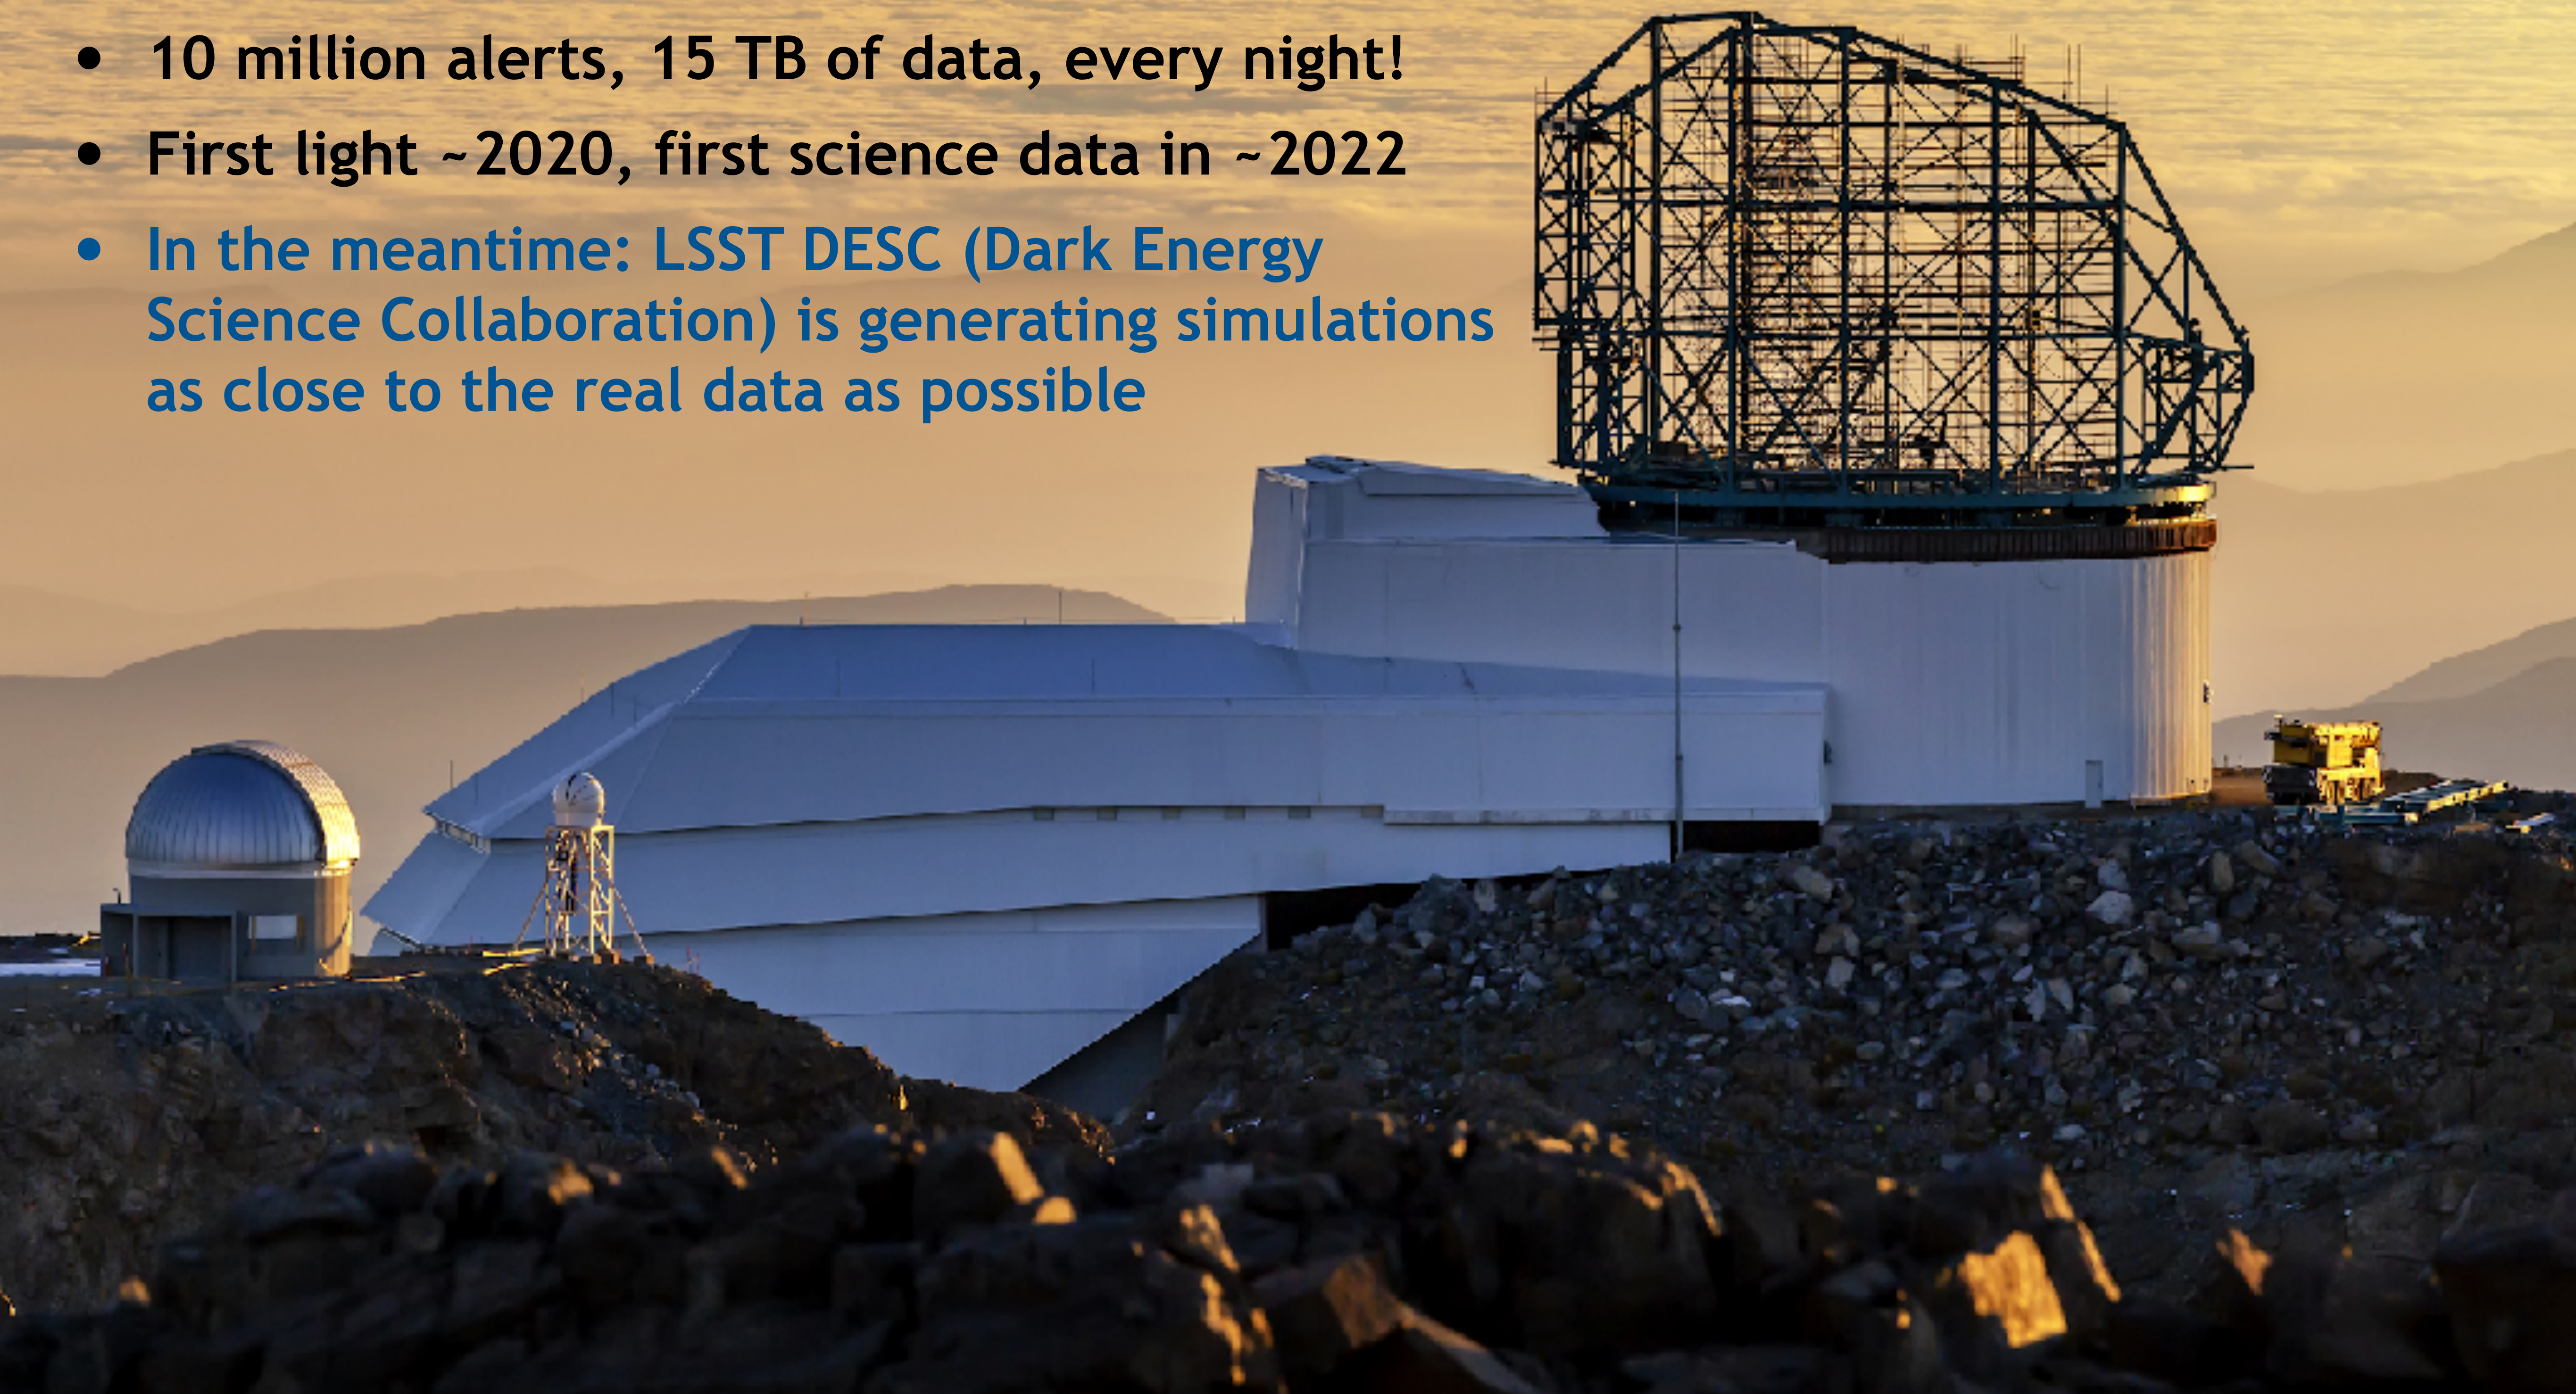

The LSST Survey and Data

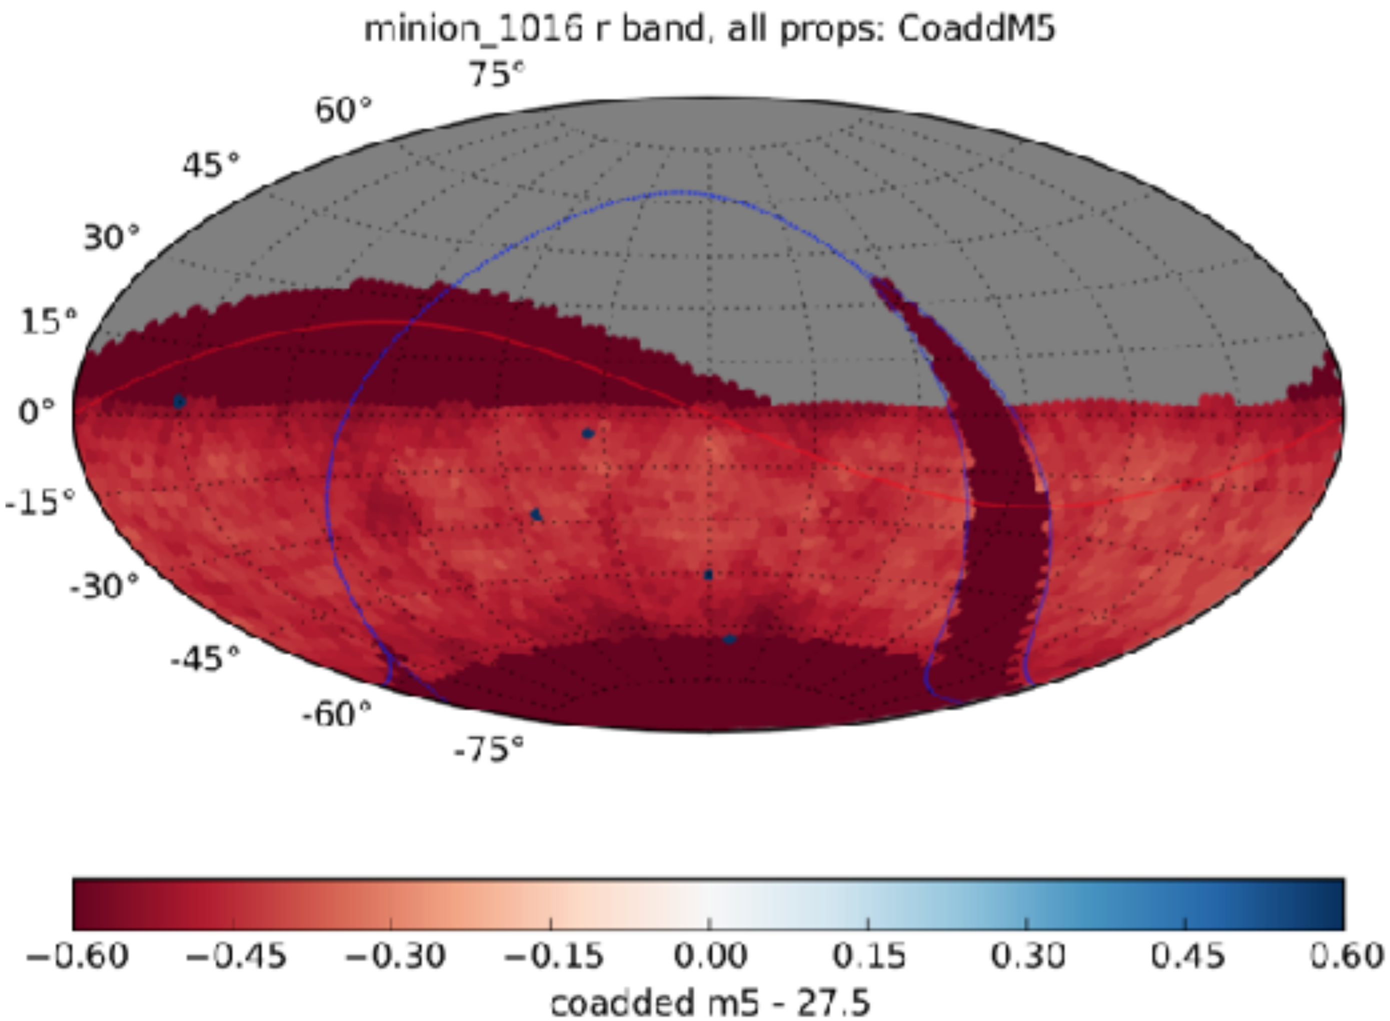

Main survey: 18,000 sq deg of southern sky, "Deep Drilling Fields", 10 sq deg each, plus further "mini-surveys" to support special science cases

- Worldwide "Alerts" released nightly (with minimal info)
- Annual data releases: Images, Object and Source tables
- LSST DESC "Data Challenge": To prepare for the arrival of data, simulate a patch of the survey as realisticly as possible

The LSST Survey and Data

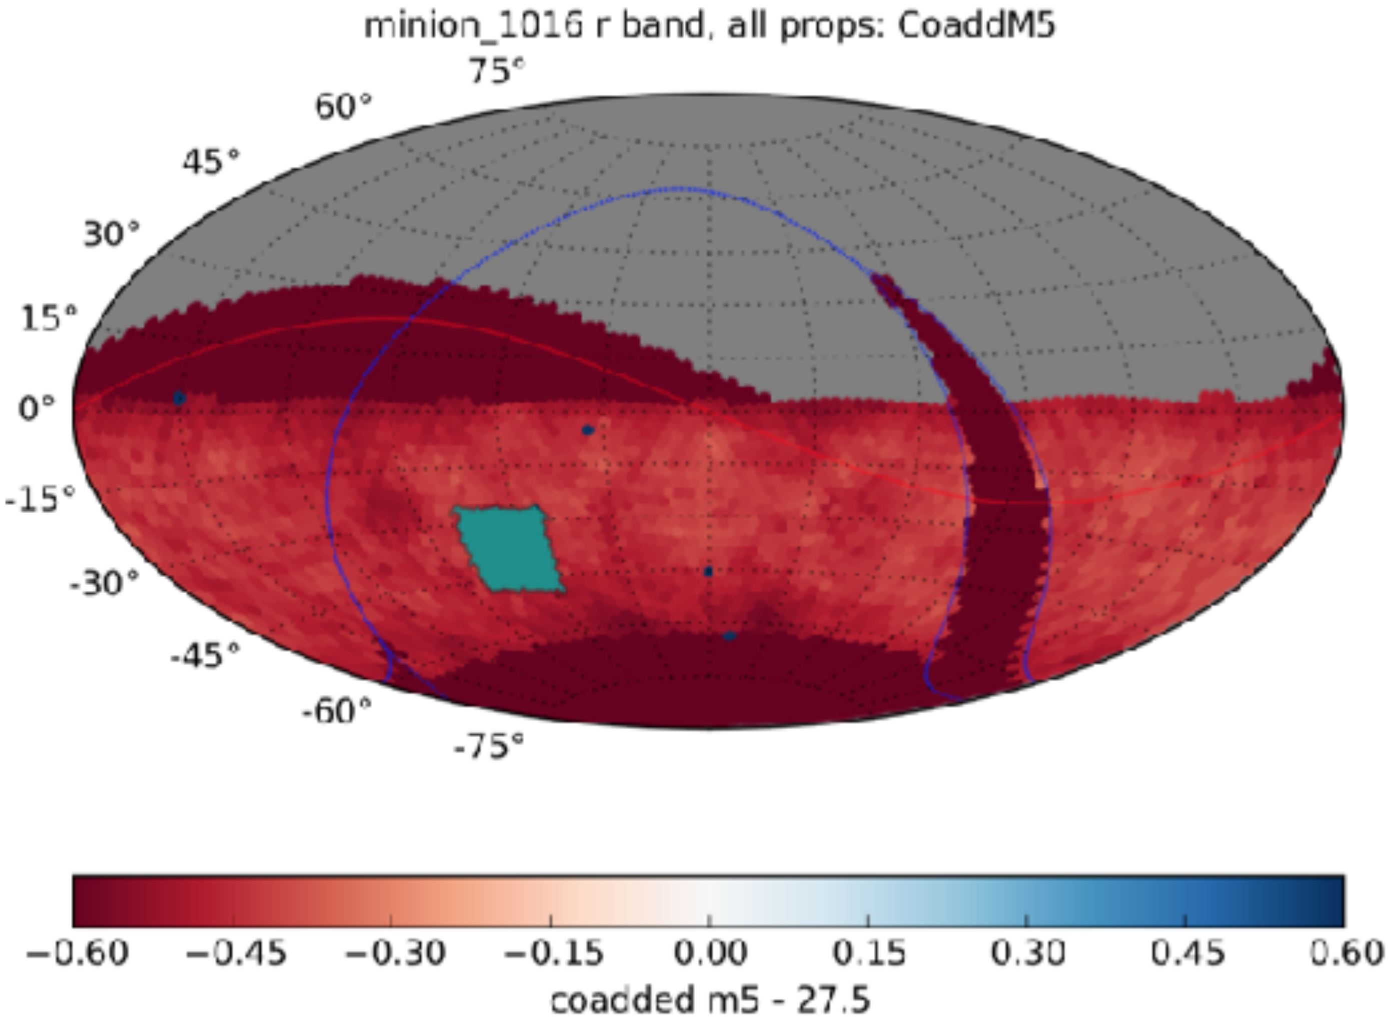

Main survey: 18,000 sq deg of southern sky, “Deep Drilling Fields”, 10 sq deg each, plus further “mini-surveys” to support special science cases

- Worldwide “Alerts” released nightly (with minimal info)
- Annual data releases: Images, Object and Source tables
- LSST DESC “Data Challenge”: To prepare for the arrival of data, simulate a patch of the survey as realisticly as possible

Data Challenge Design

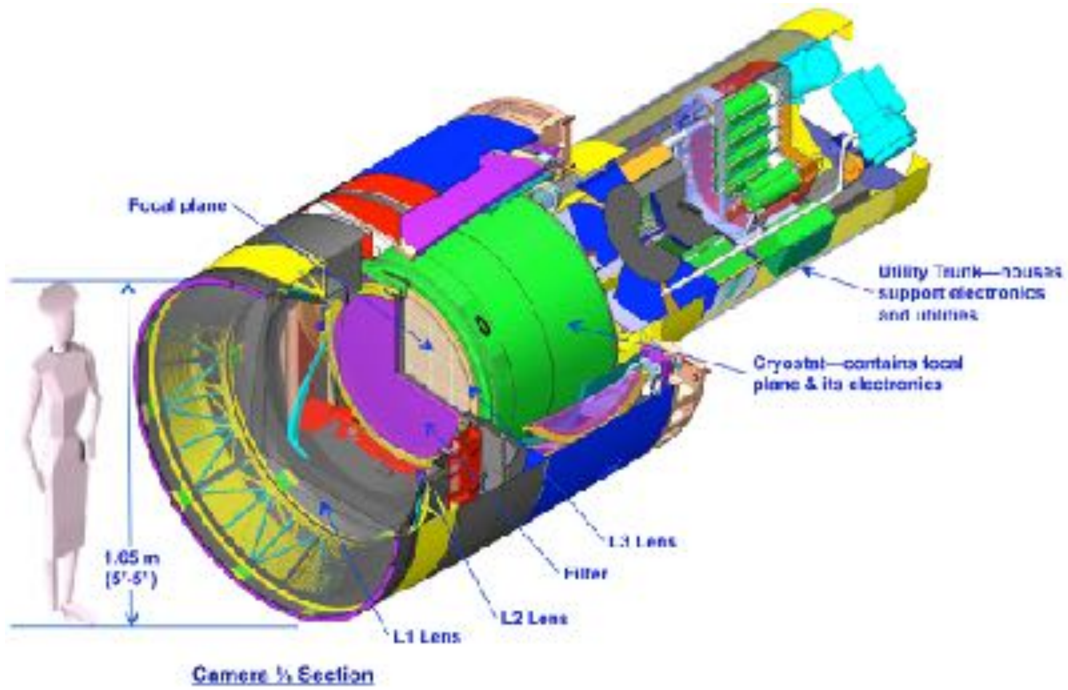

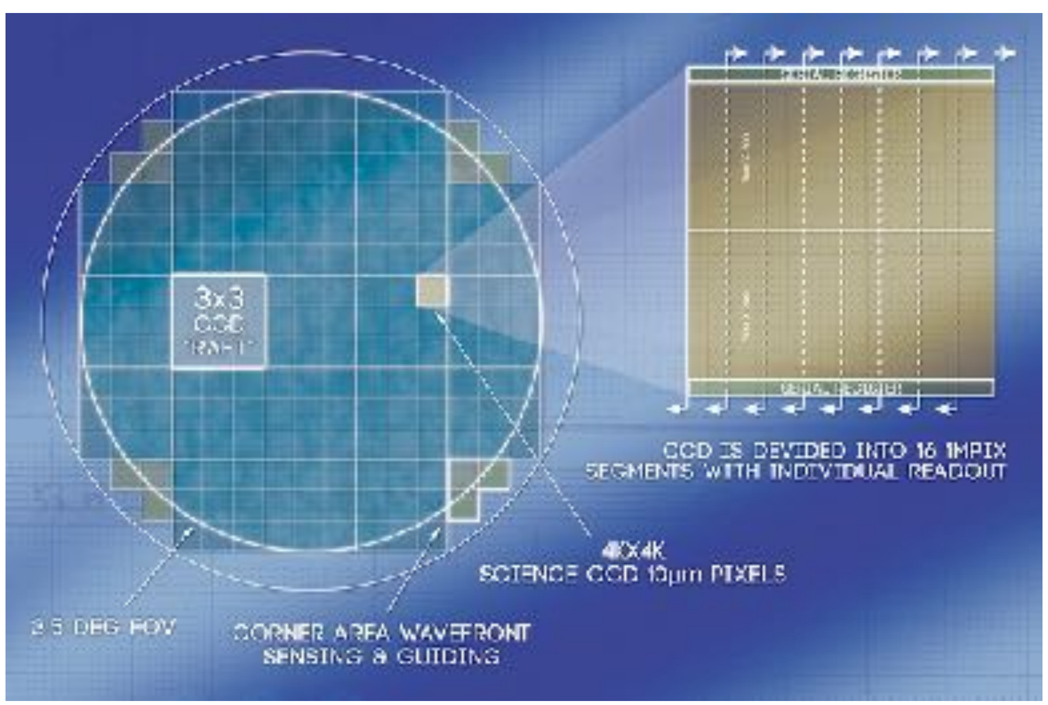

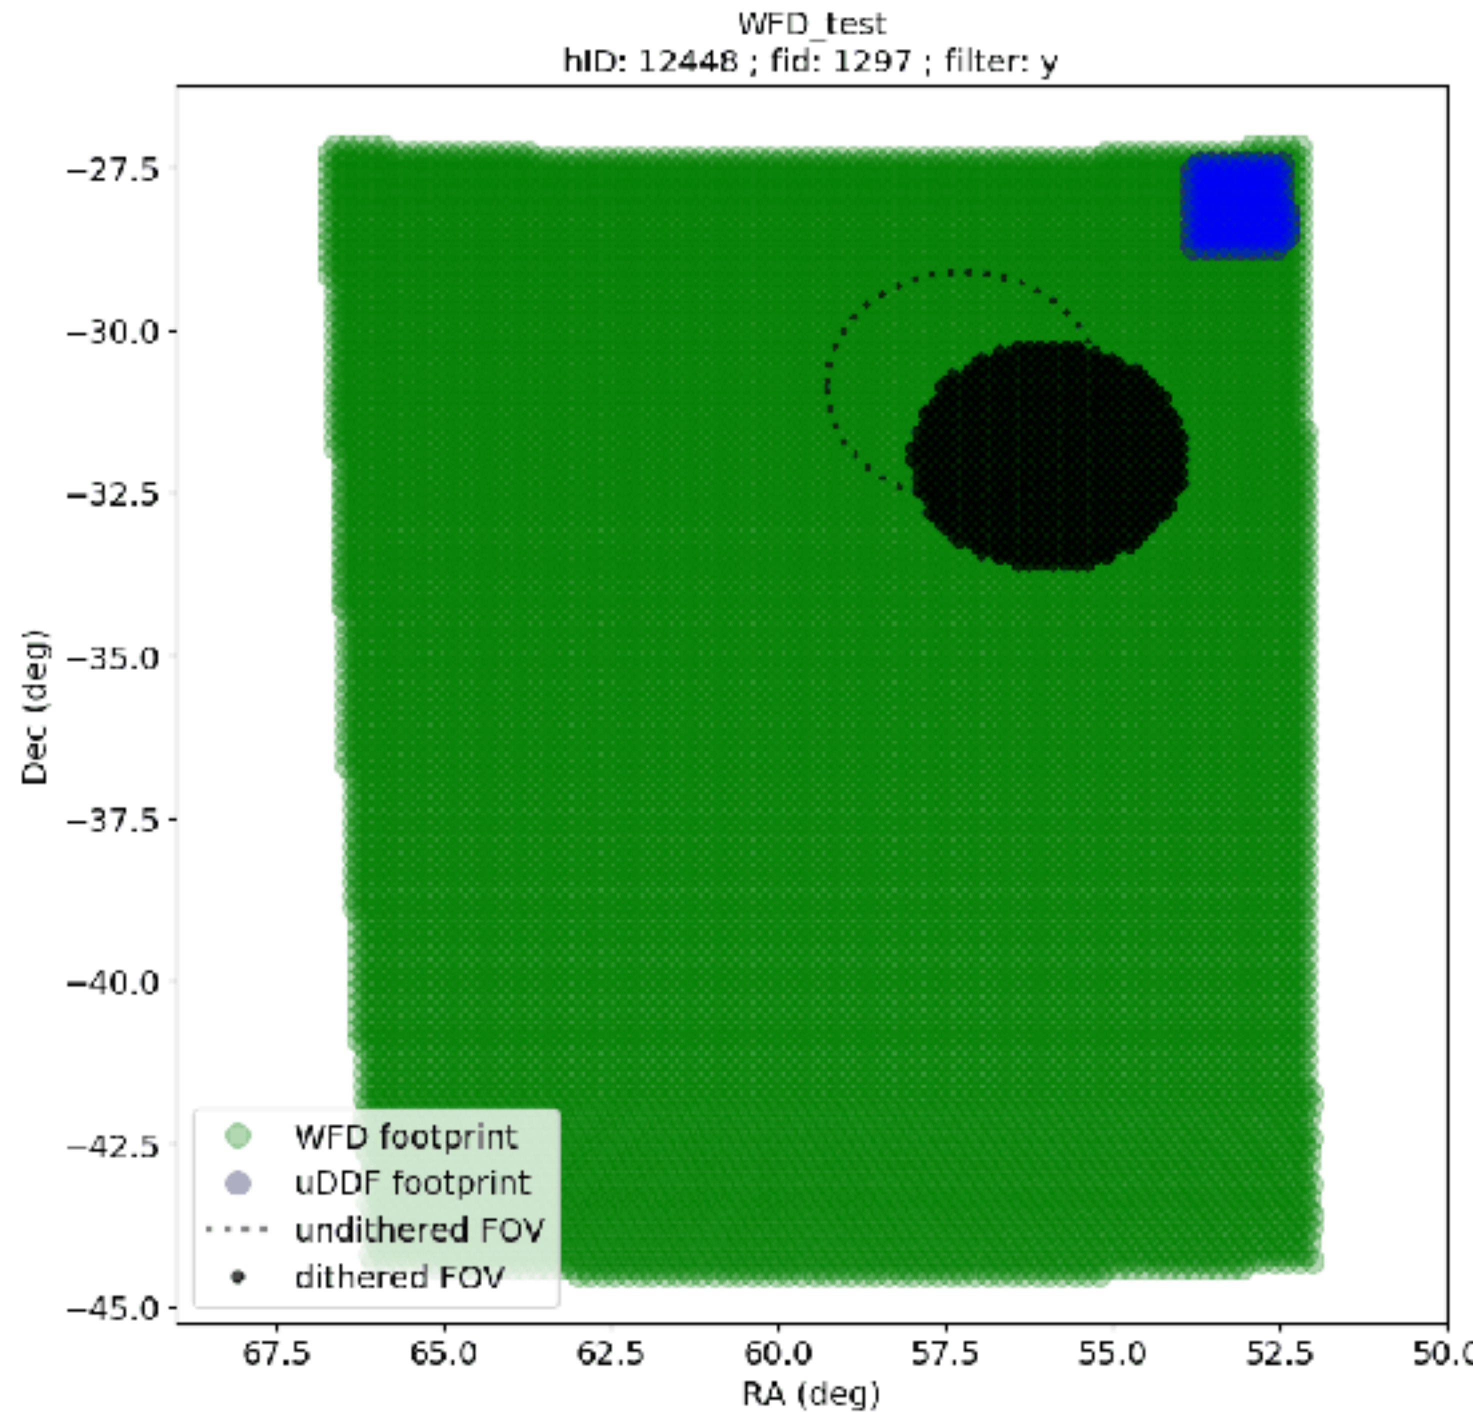

Image Credit: H. Awan, Rutgers U.

Data Challenge Design

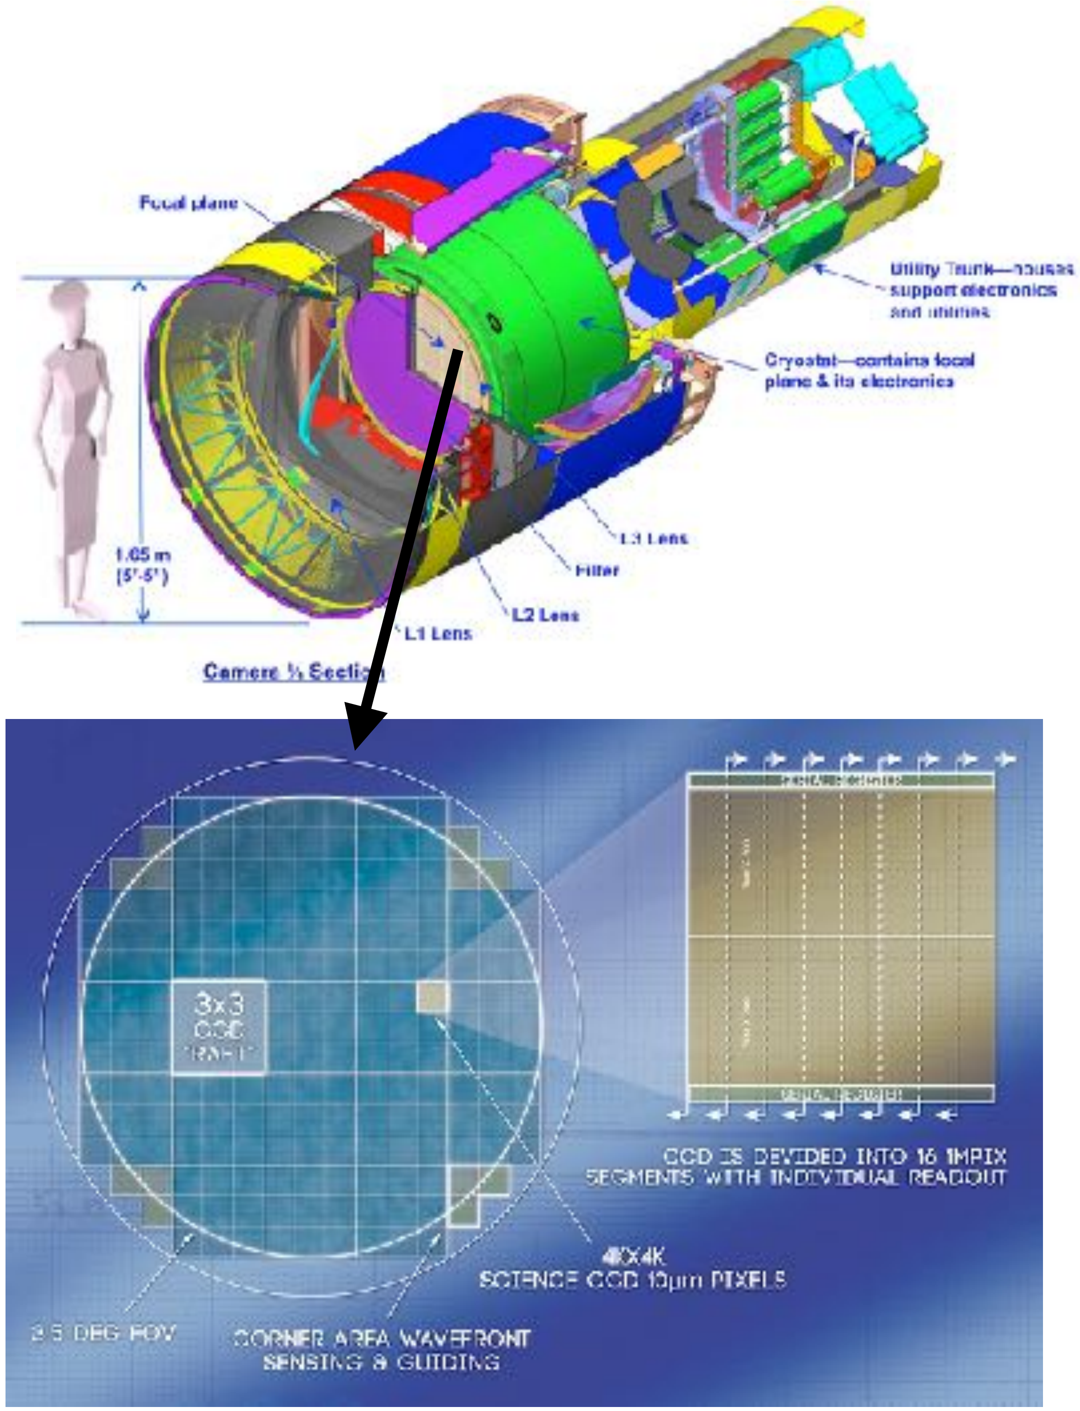

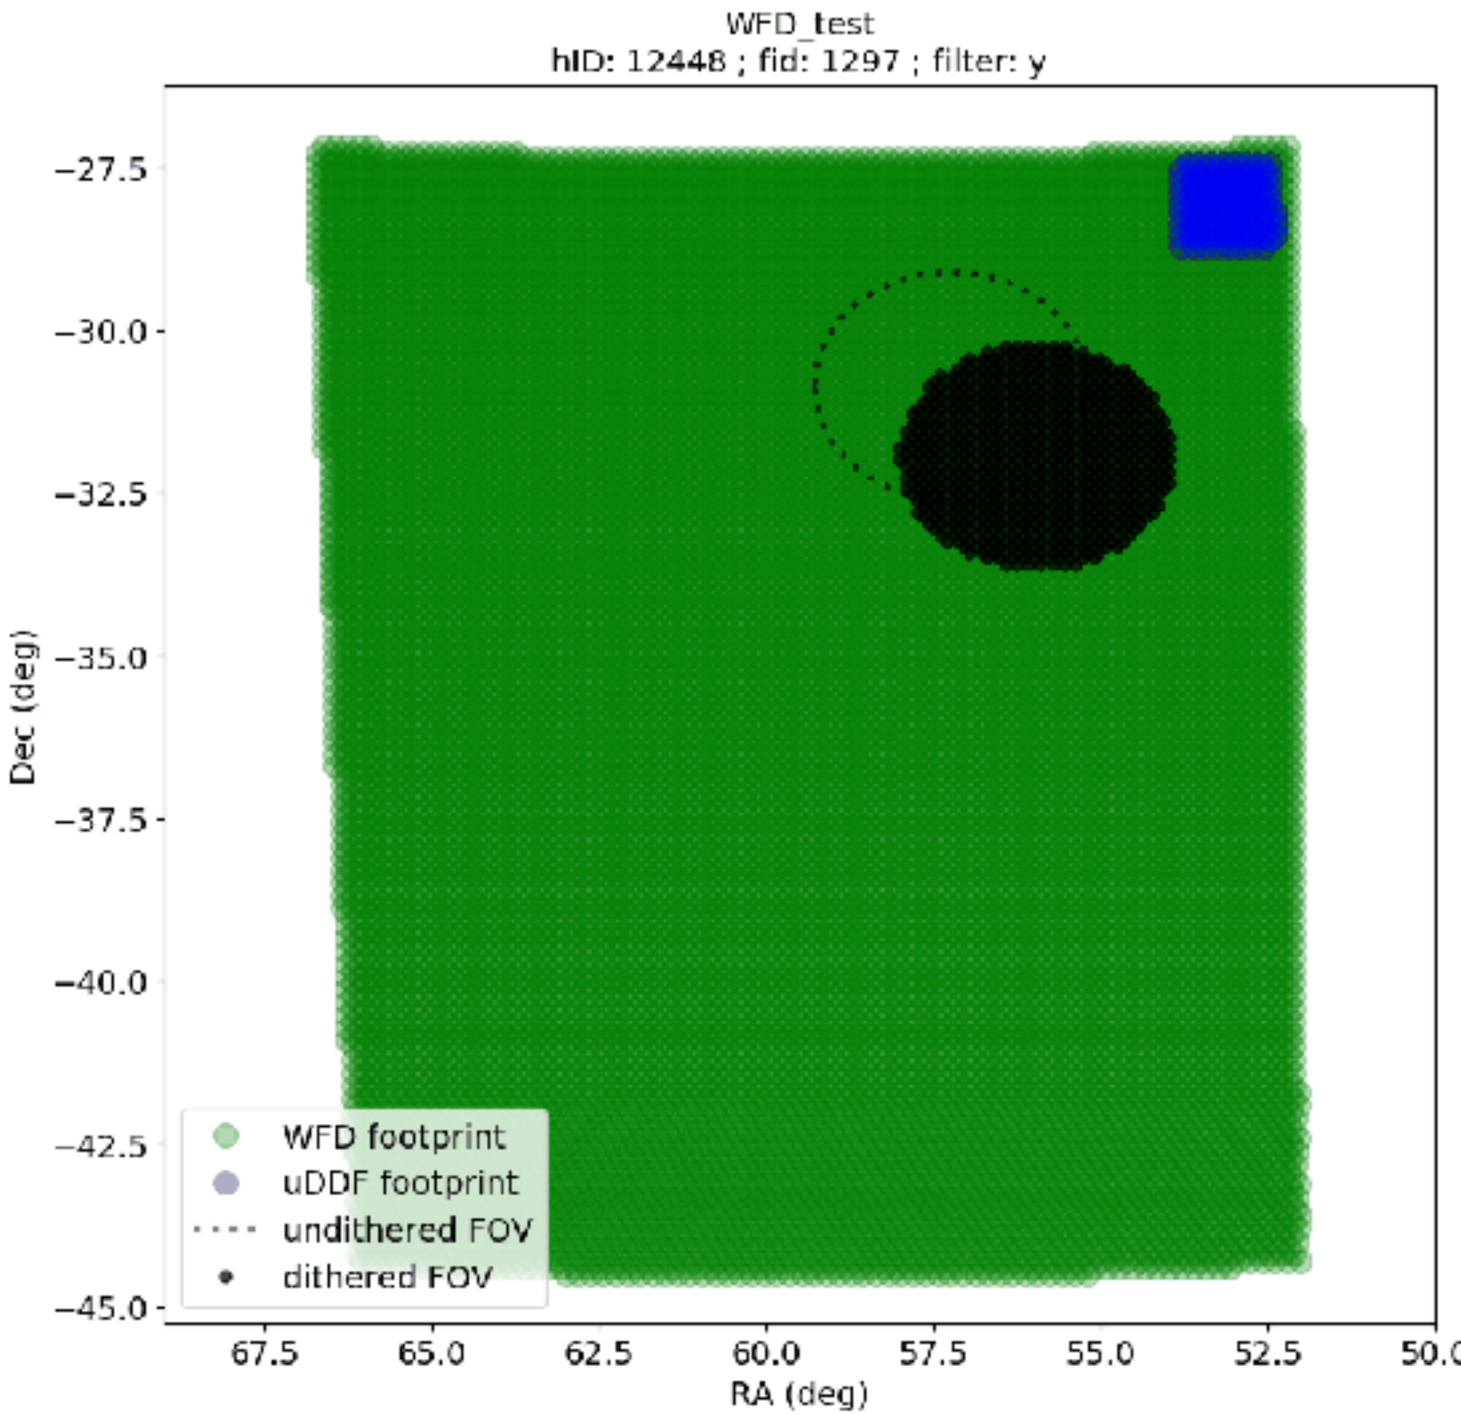

Image Credit: H. Awan, Rutgers U.

Data Challenge Design

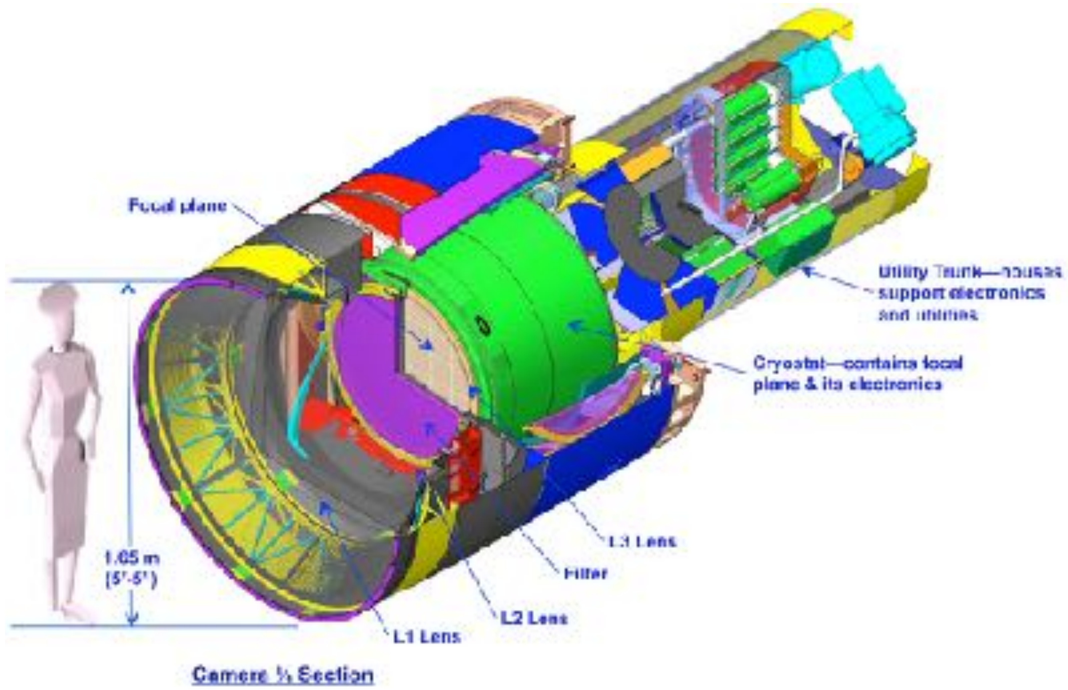

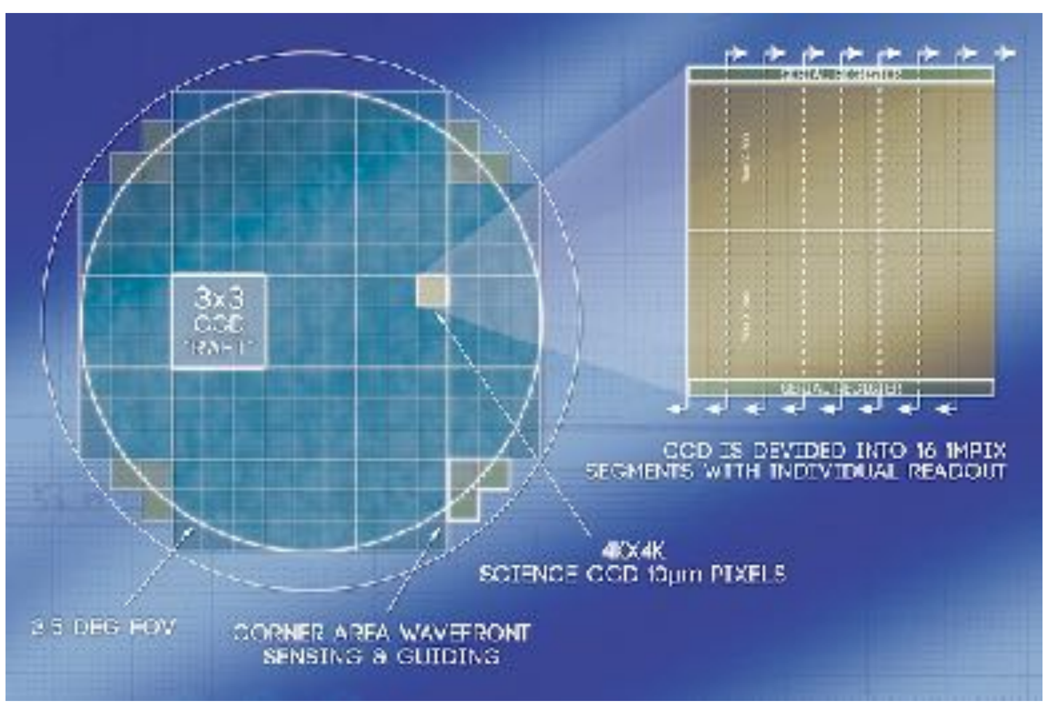

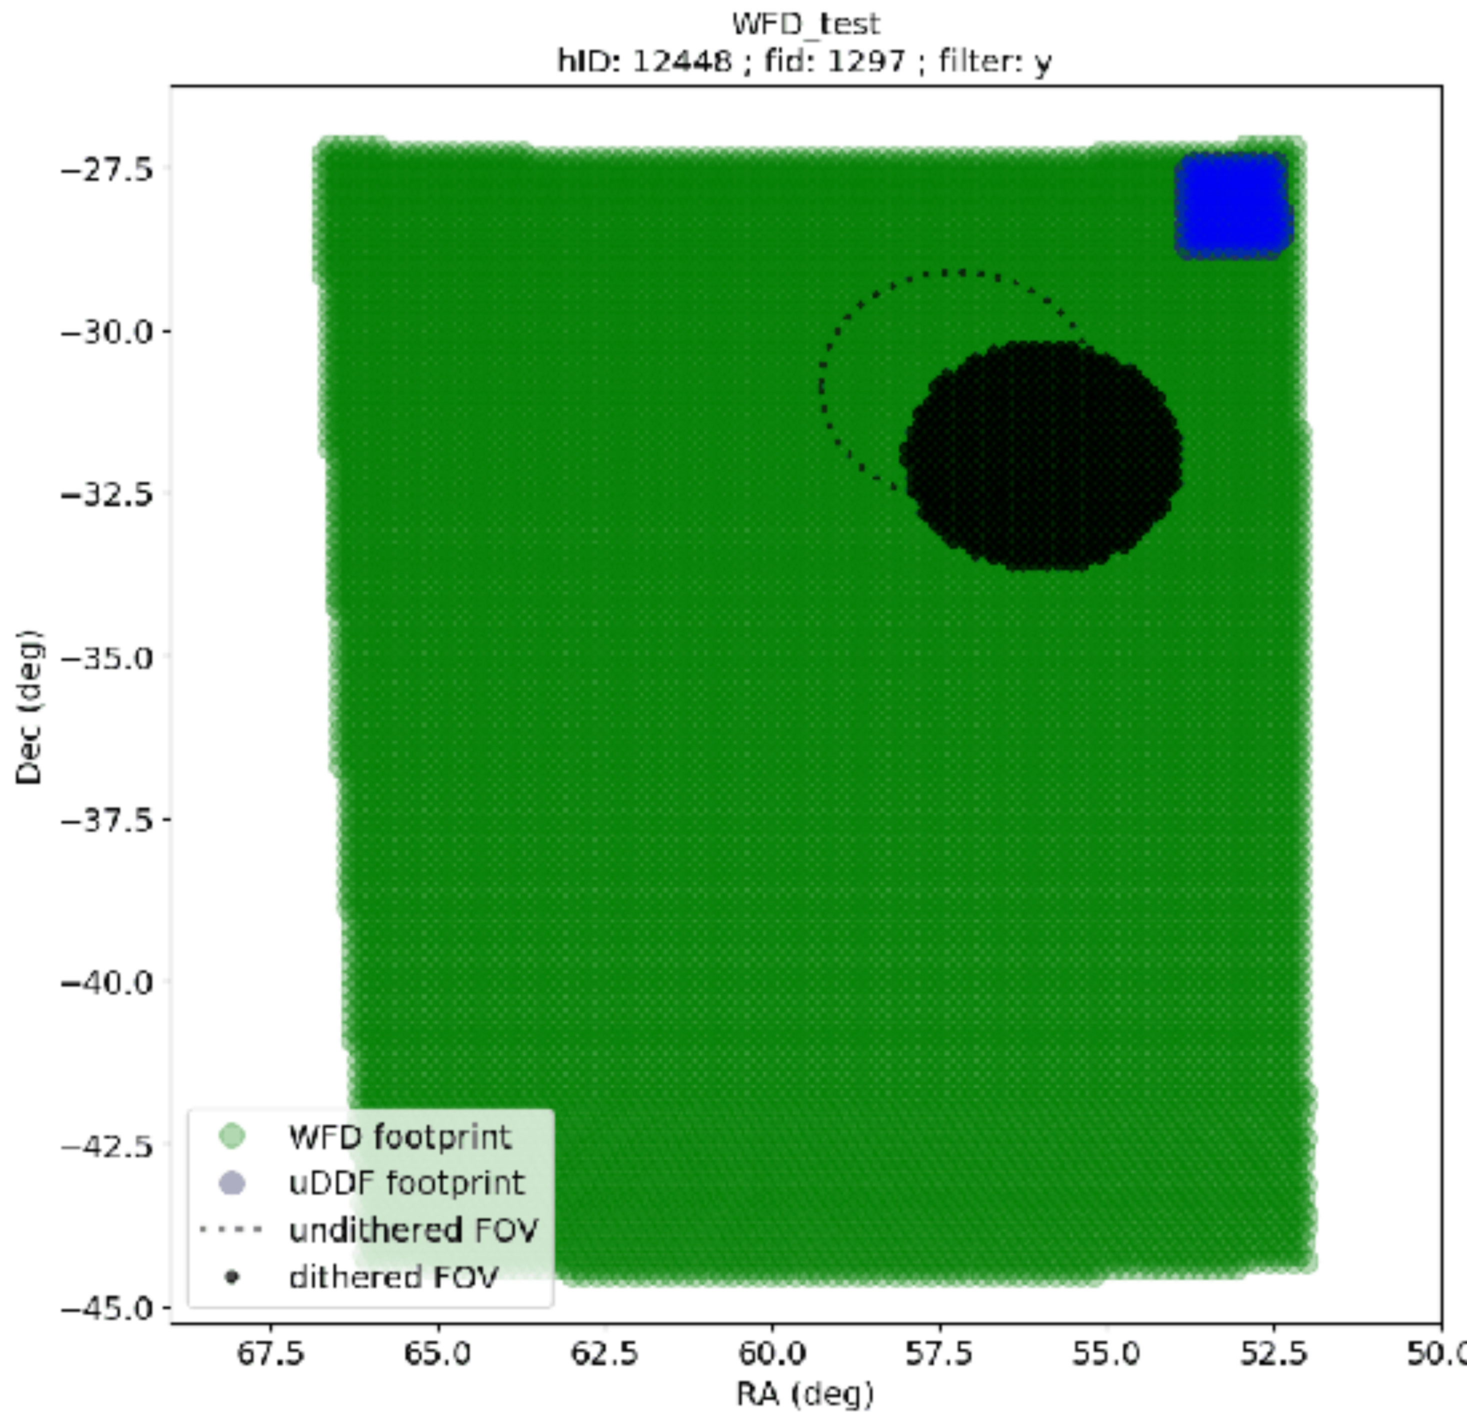

Image Credit: H. Awan, Rutgers U.

Data Challenge Design

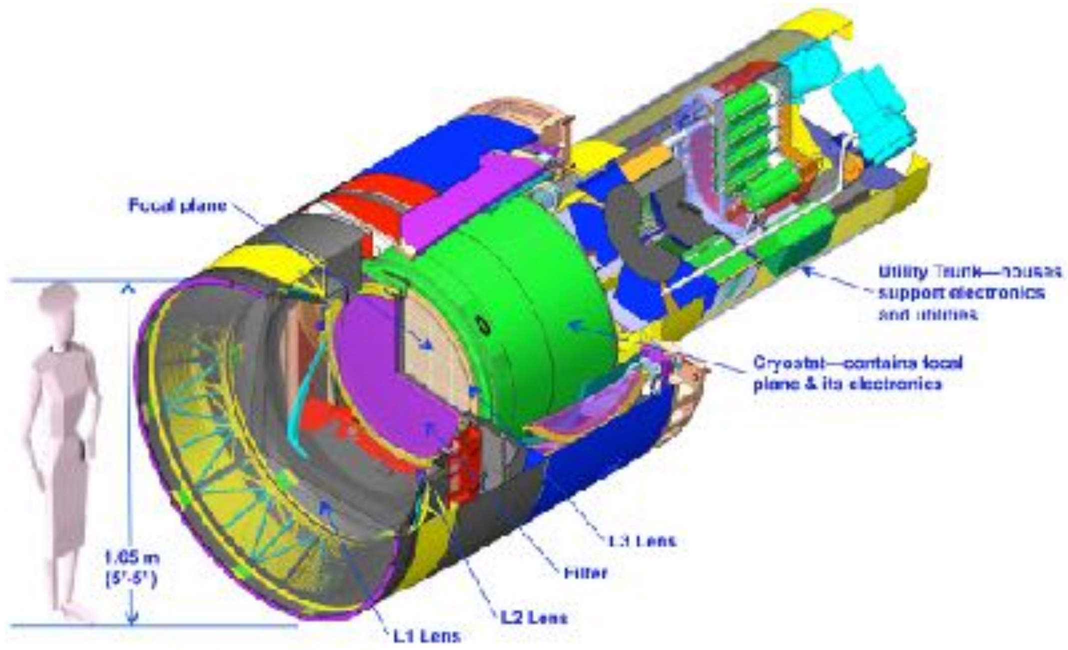

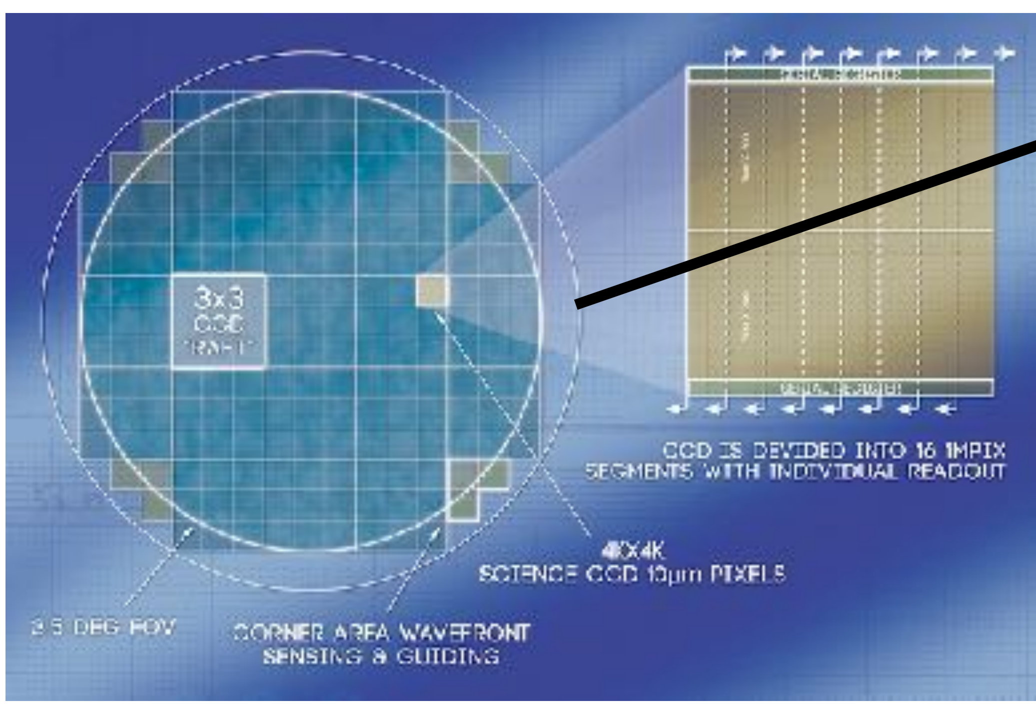

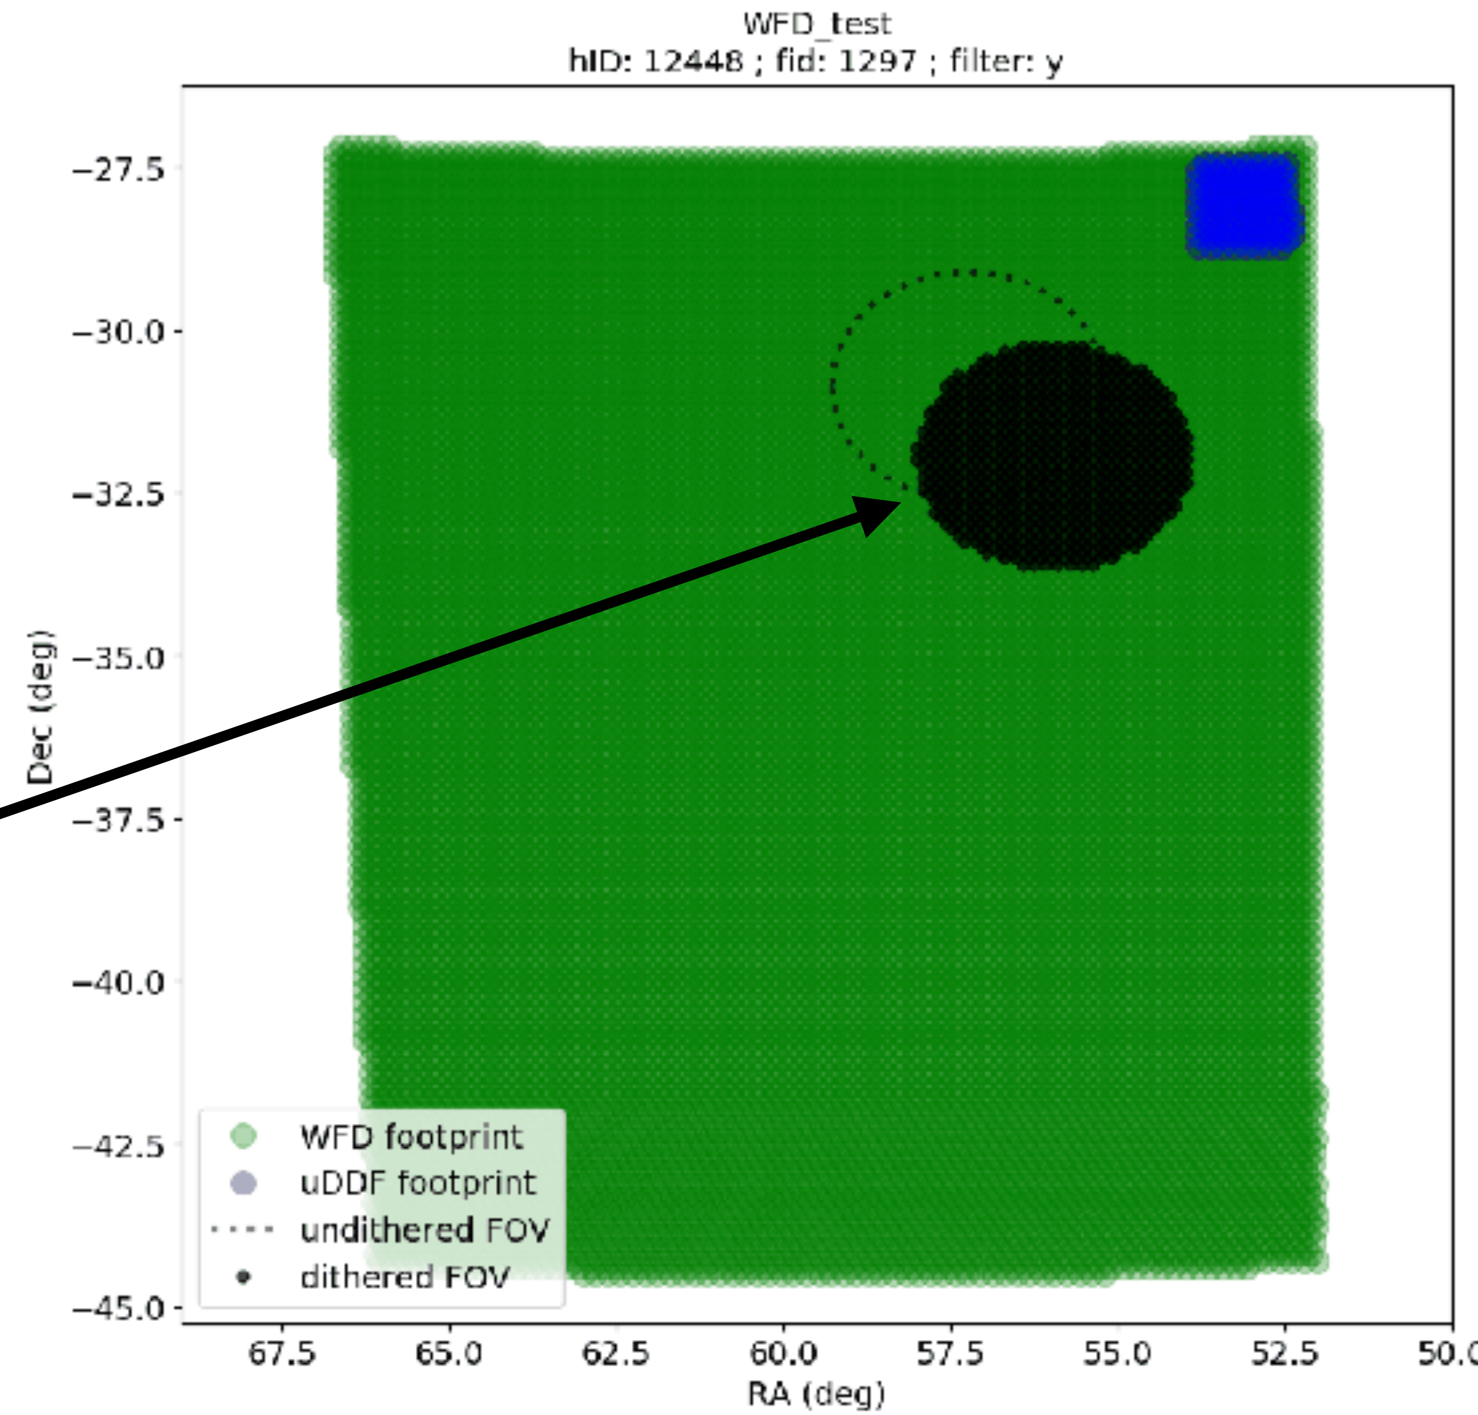

Image Credit: H. Awan, Rutgers U.

Data Challenge Design

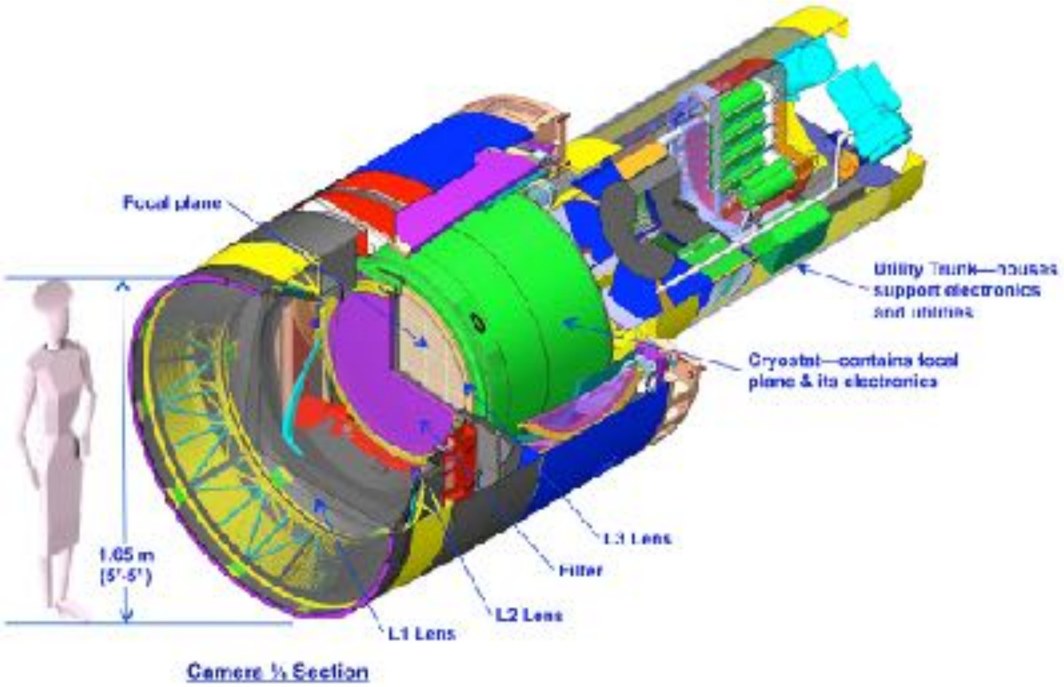

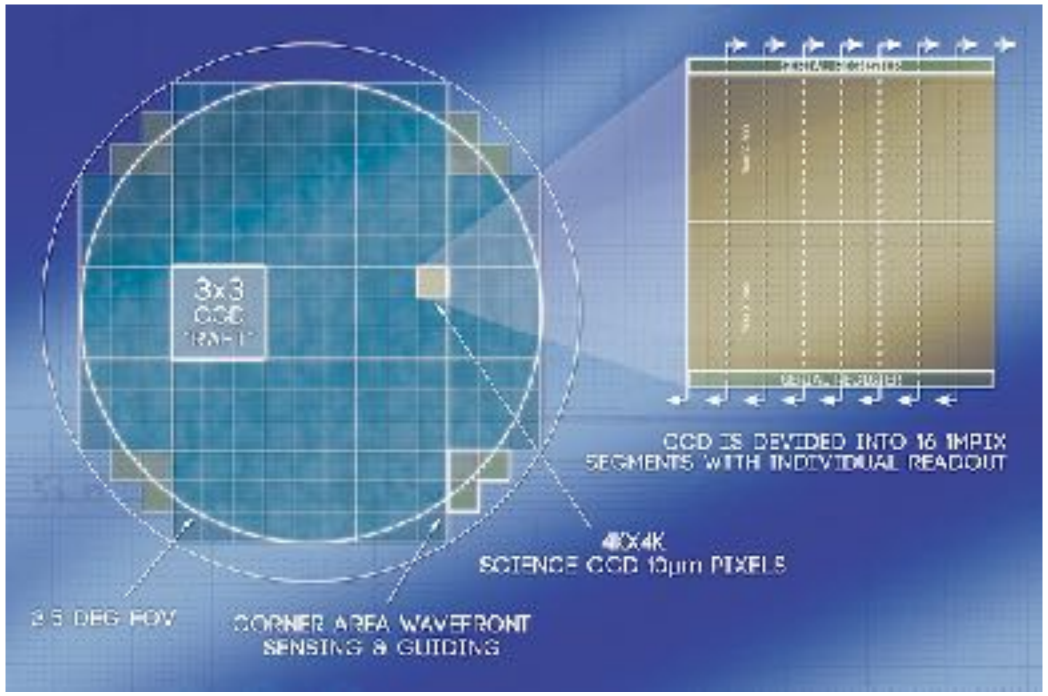

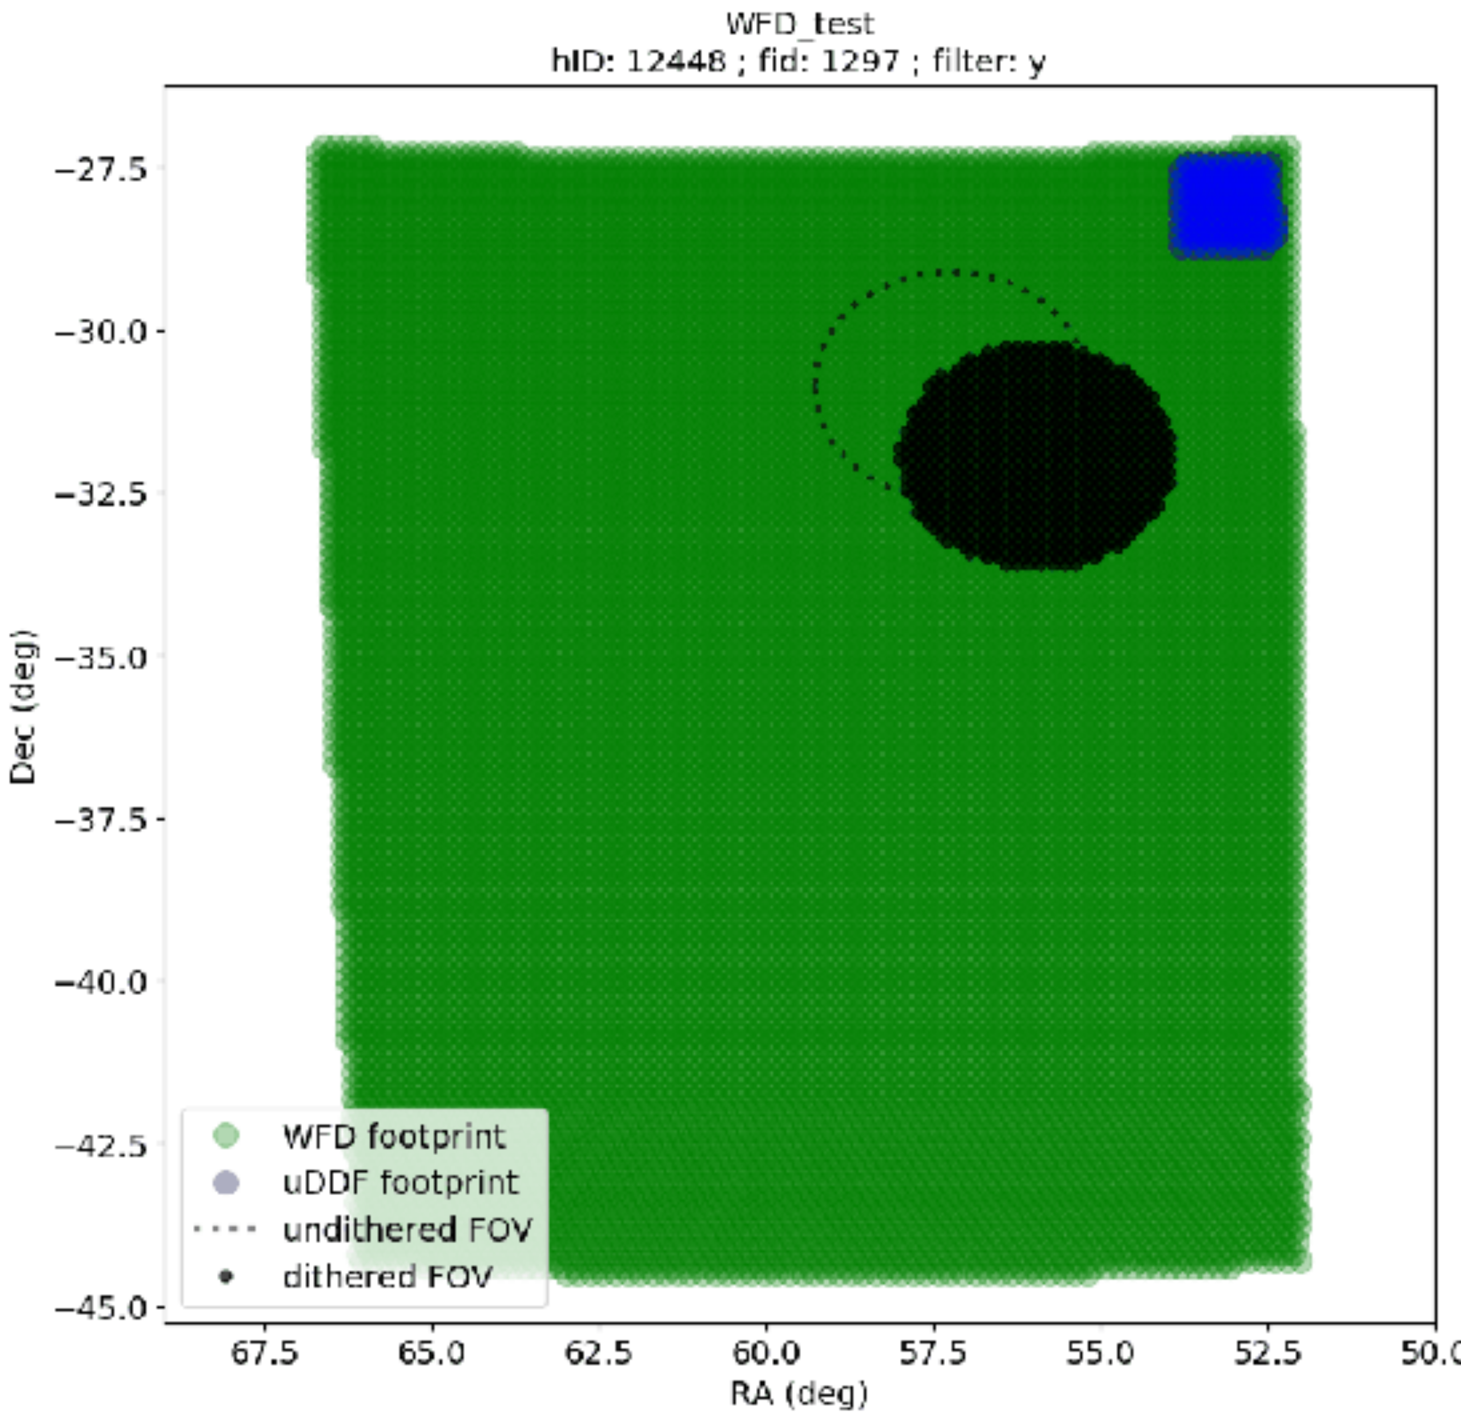

Image Credit: H. Awan, Rutgers U.

Data Challenge Design

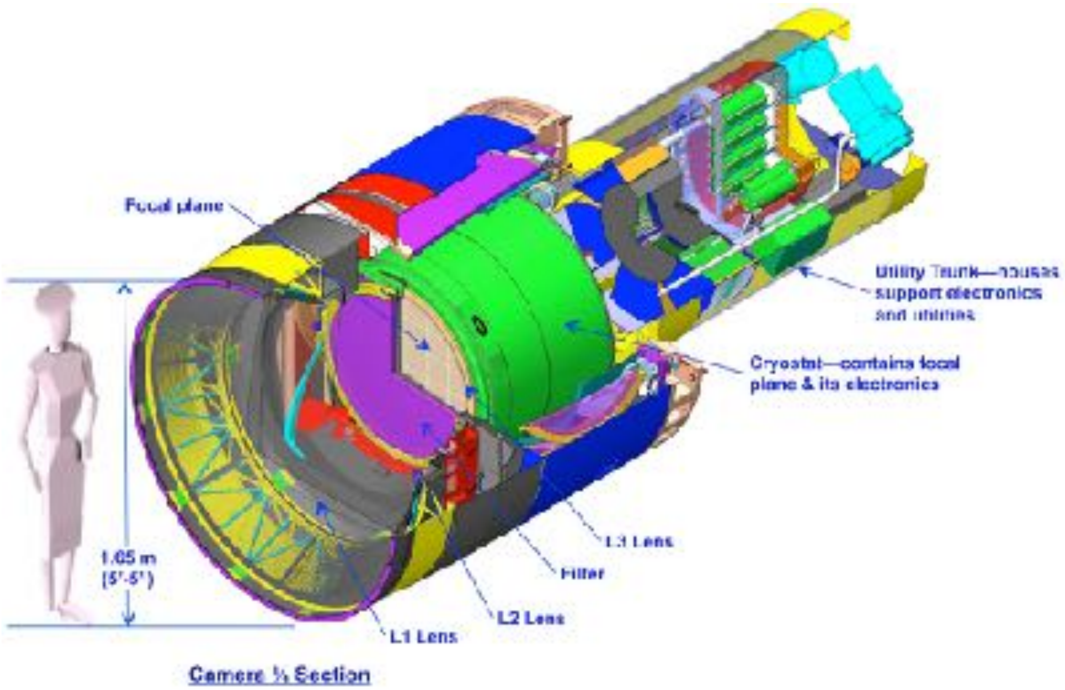

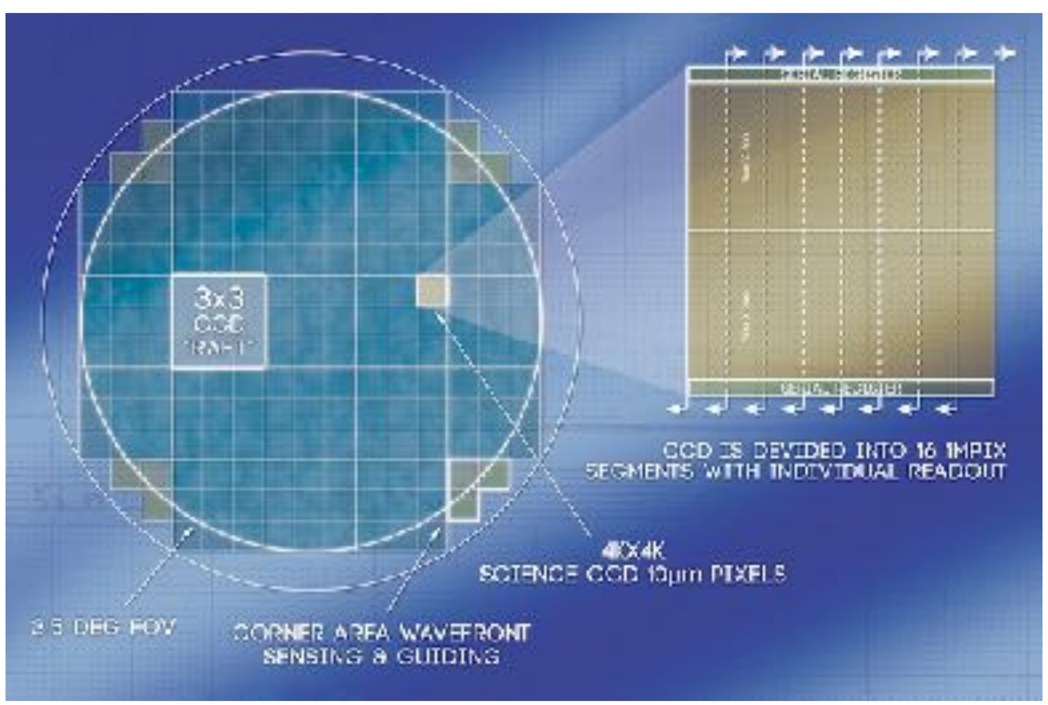

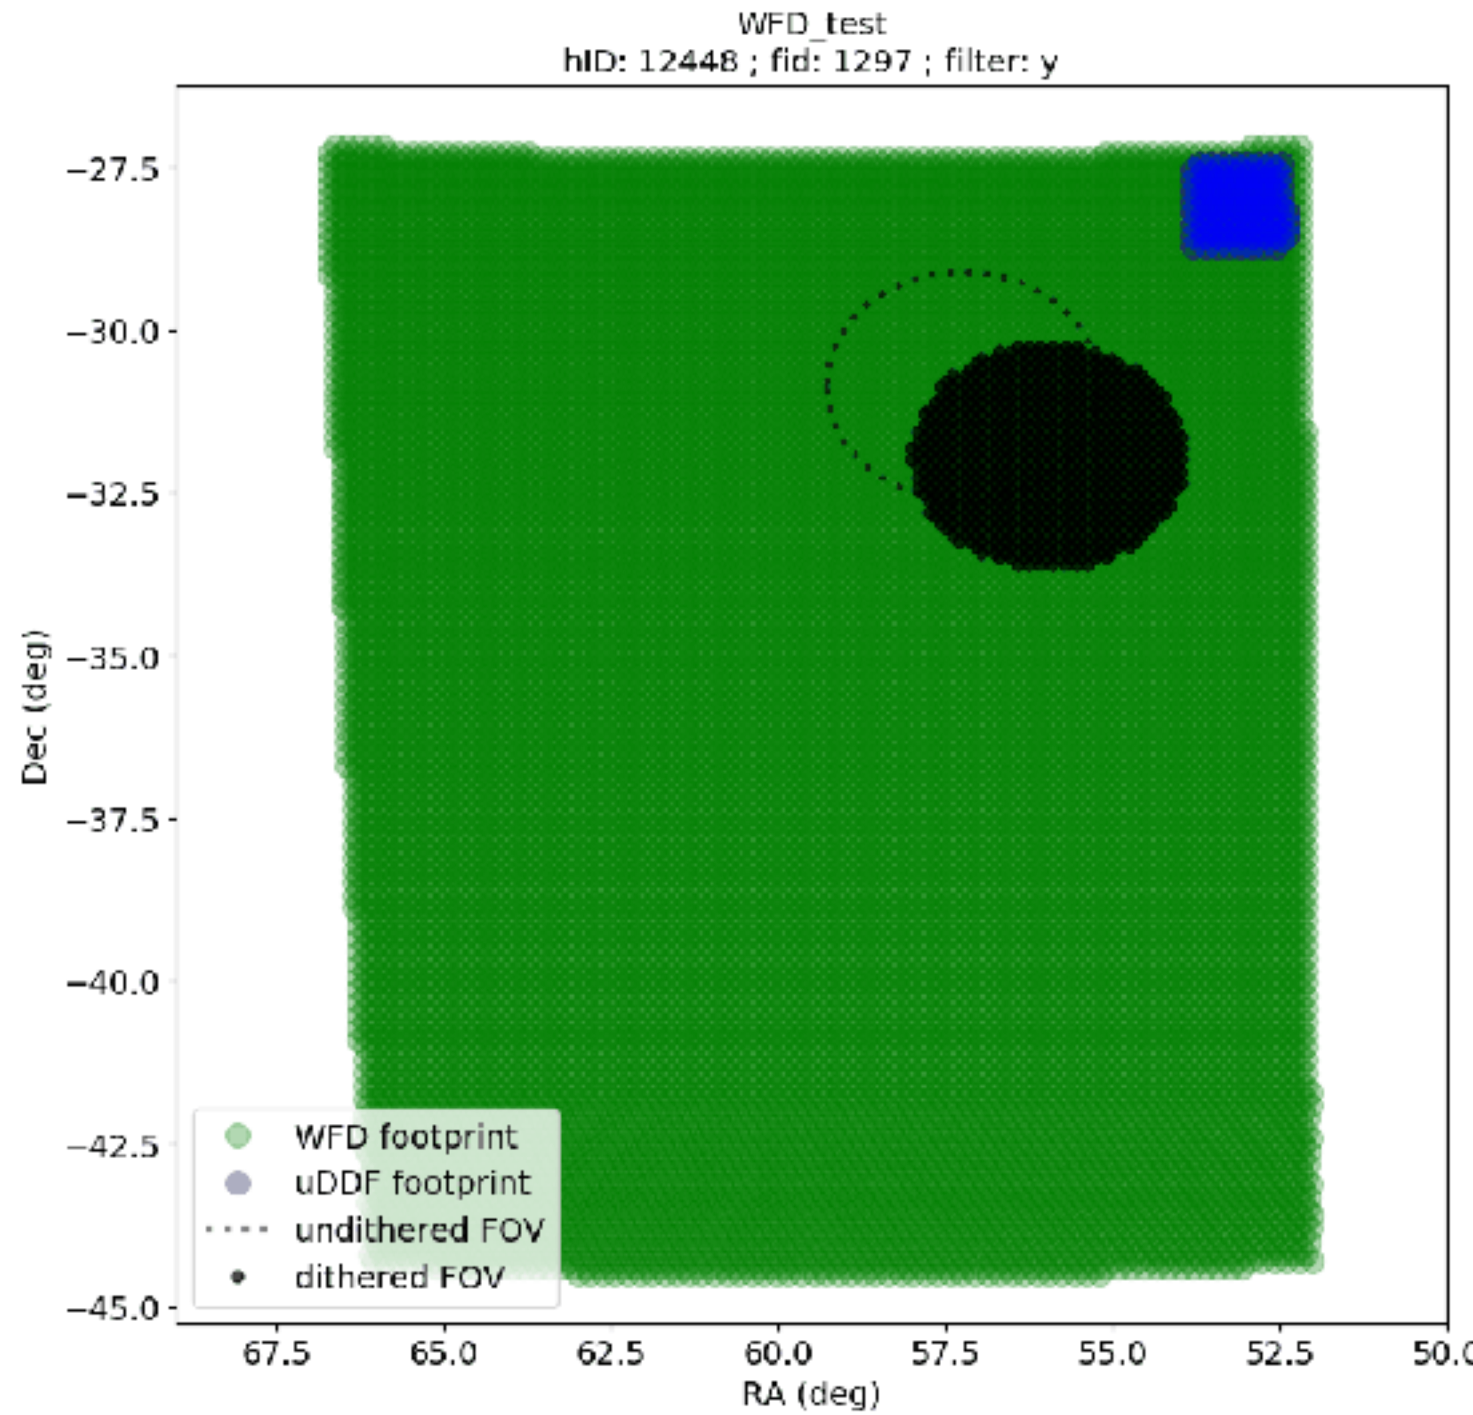

Image Credit: H. Awan, Rutgers U.

DESC End-to-End Simulation Workflow

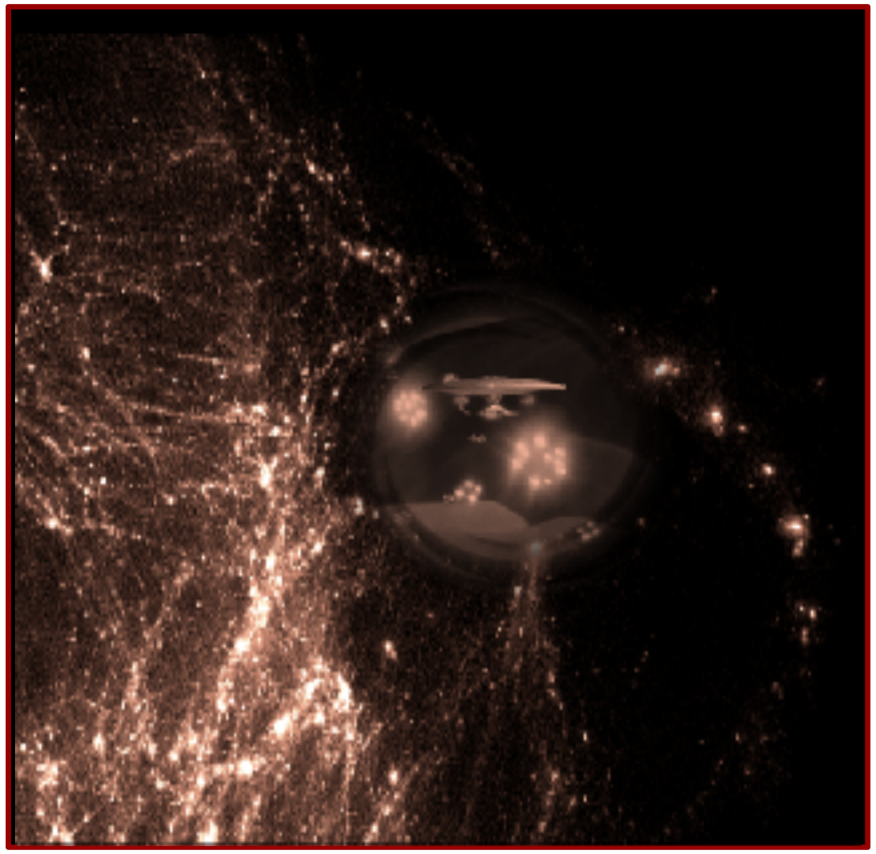

Extra-galactic catalog:
Galaxy shapes and colors, shears, dust, ….

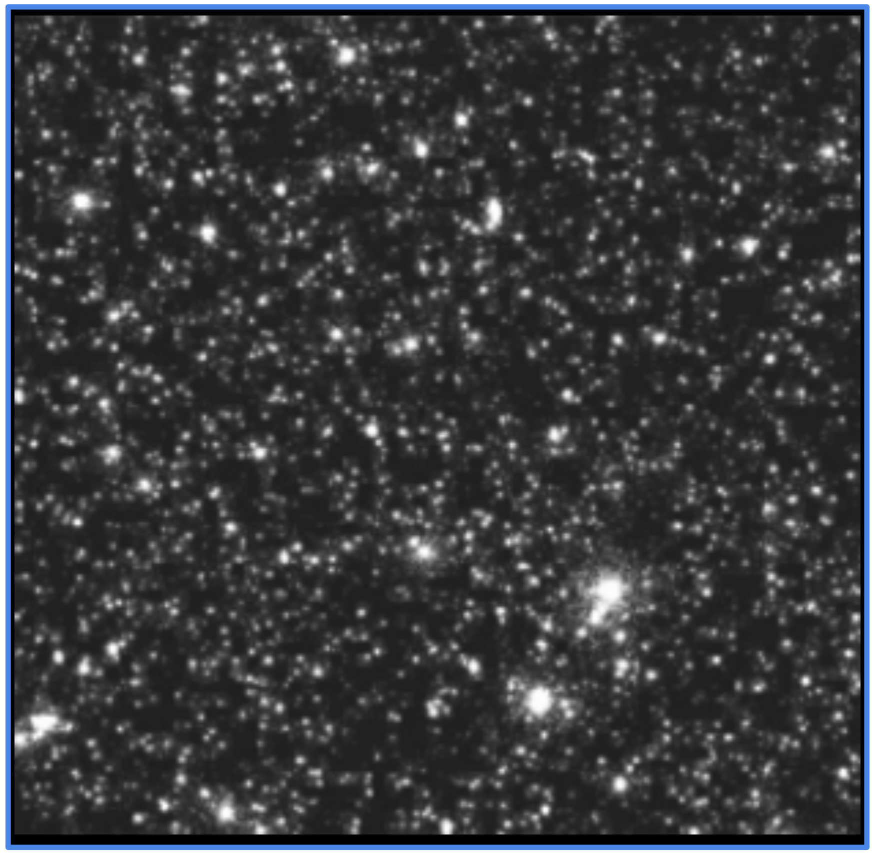

Image simulations:
Milky Way stars, dust, atmosphere and seeing conditions, telescope, CCD defects, cosmic rays, transients, …

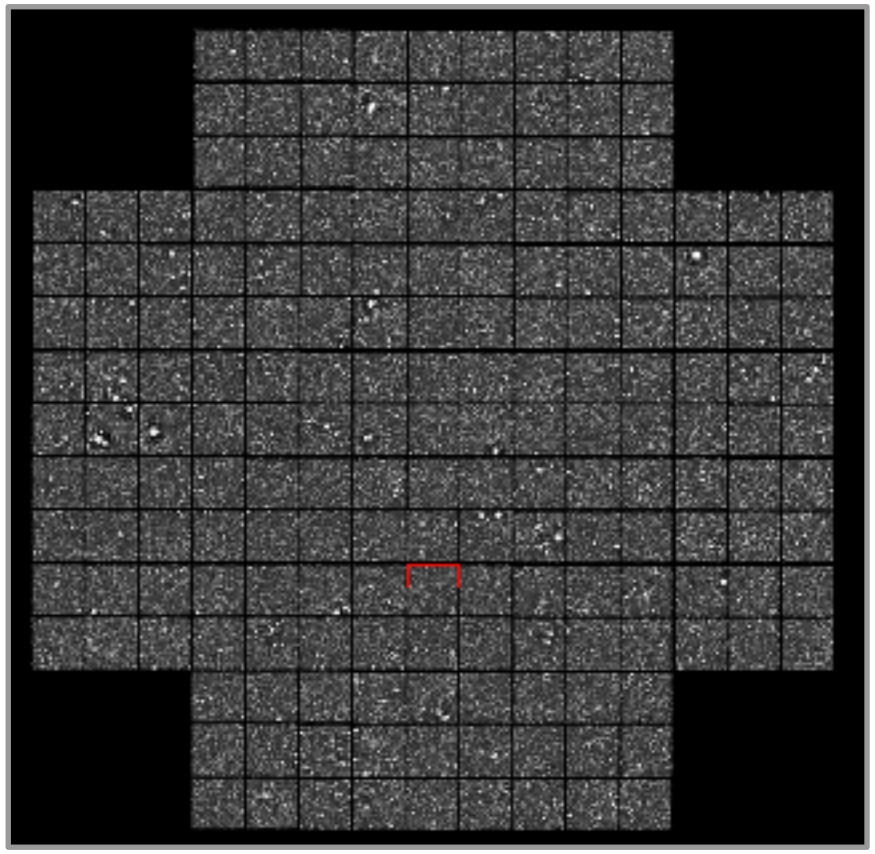

Processing with the LSST Science Pipeline:
Calibration, clean-up, …

CosmoDC2: The Extragalactic Catalog

Ingredients:

- A REALLY large simulation: HACC simulation with > 1 trillion particles, 4.5PB of data; 4.2Gpc volume, m_p~$2x10^9$Msun; carried out on 2/3 of Mira
- A physics-based model that allows us to determine which halos host what kind of galaxies
- Validation data and tests from past and ongoing surveys to confirm that our simulated data looks like the real Universe!
- Result: Catalog where each galaxy has more than 500 properties

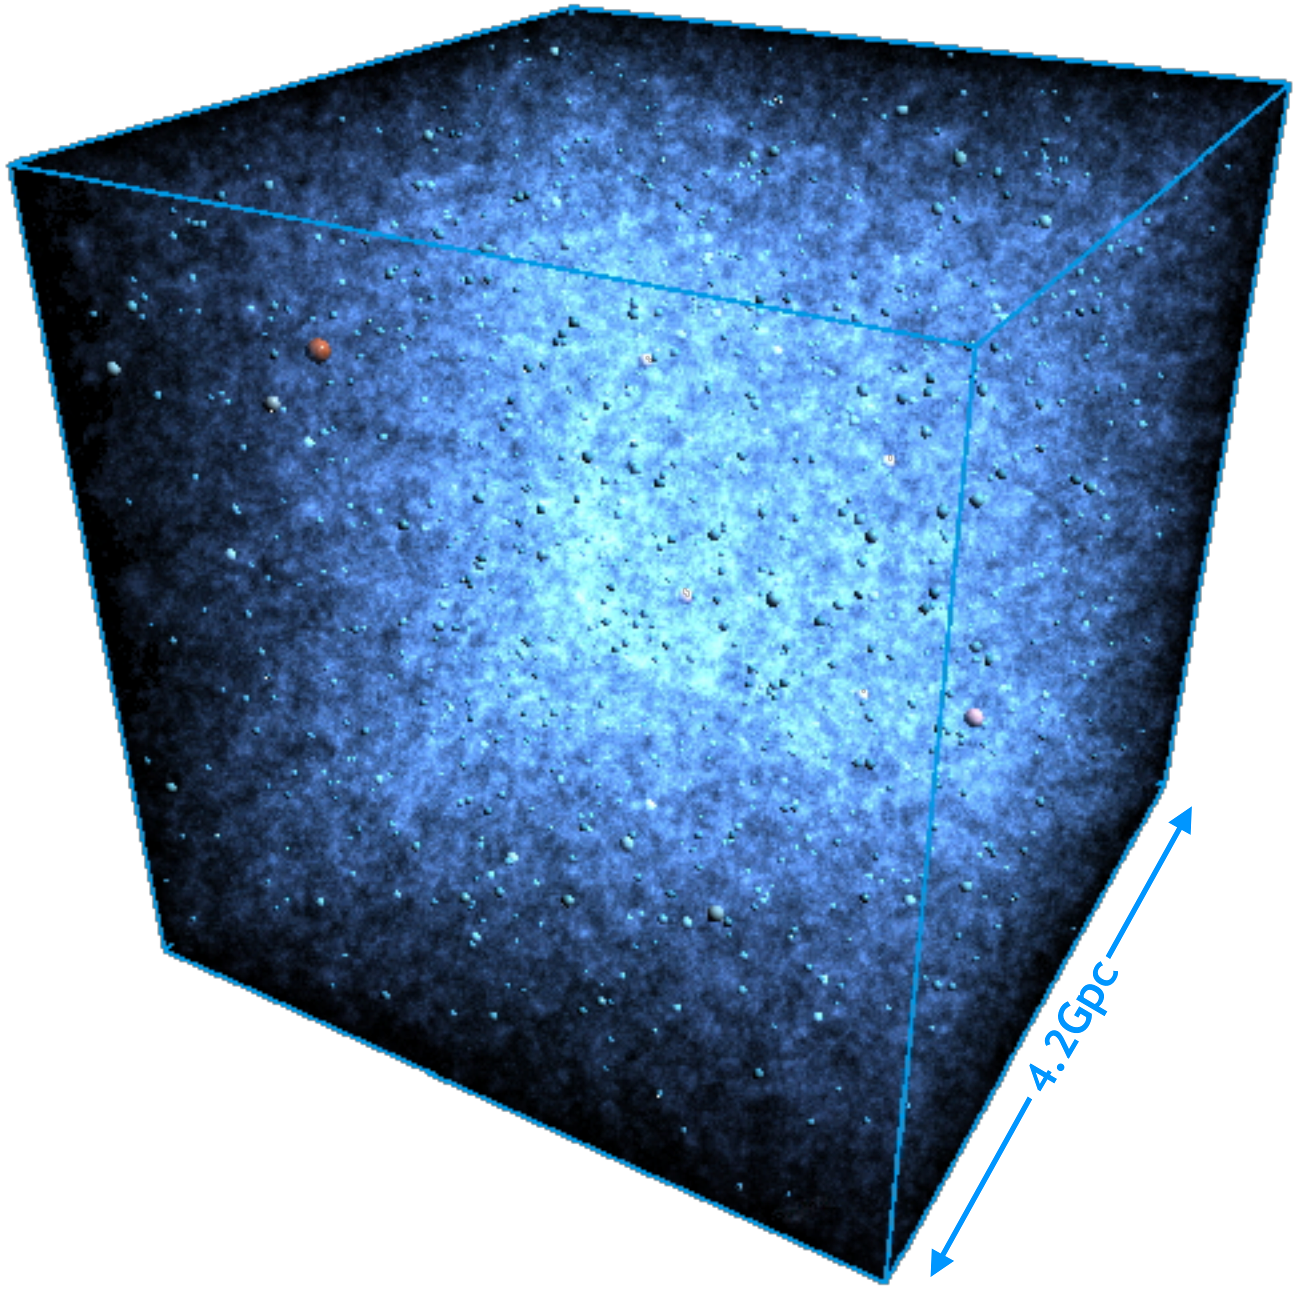

Large halos in the Outer Rim simulation, volume large enough to model future survey
Image Credit: S. Rizzi, J. Insley

Image Simulation Runs

- Image simulations rather costly, running at NERSC, Argonne (ALCF), and UK/France Grid
- Developed workflows for imSim using Parsl (ran on up to 4000 nodes at ALCF)
- 7.5M sensor visits need to be simulated and processed!

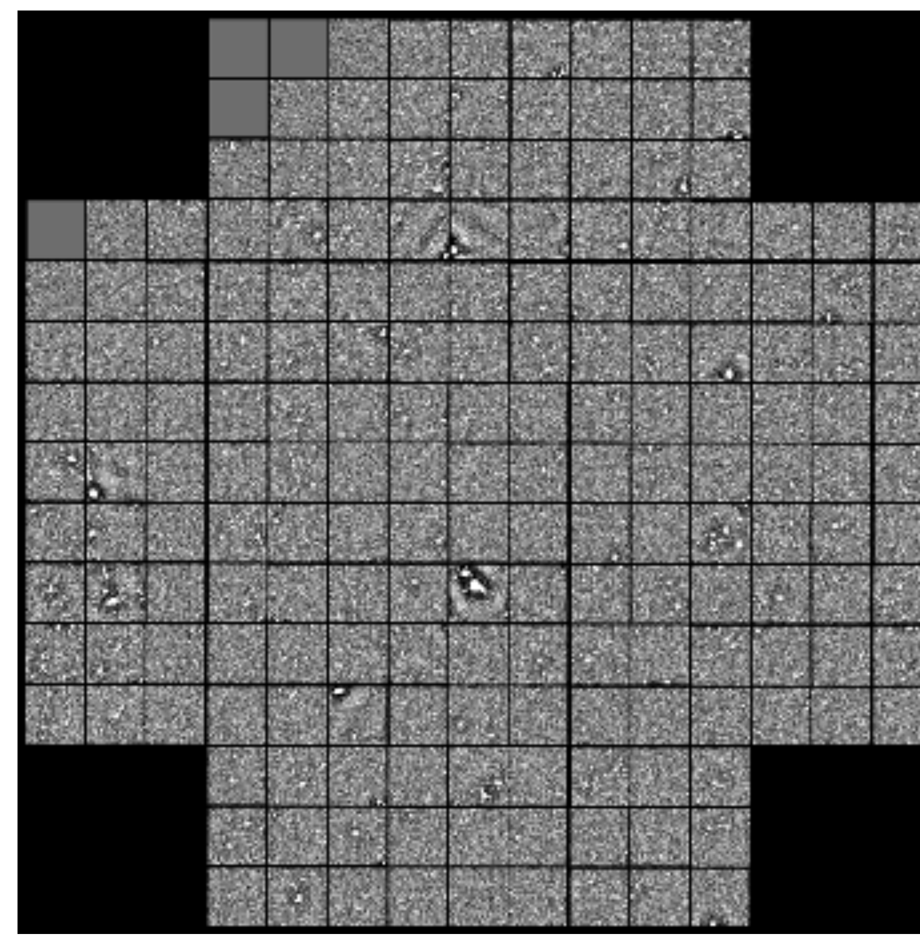

Image Credit: H. Kelly, SLAC

Processing with the LSST Science Pipeline

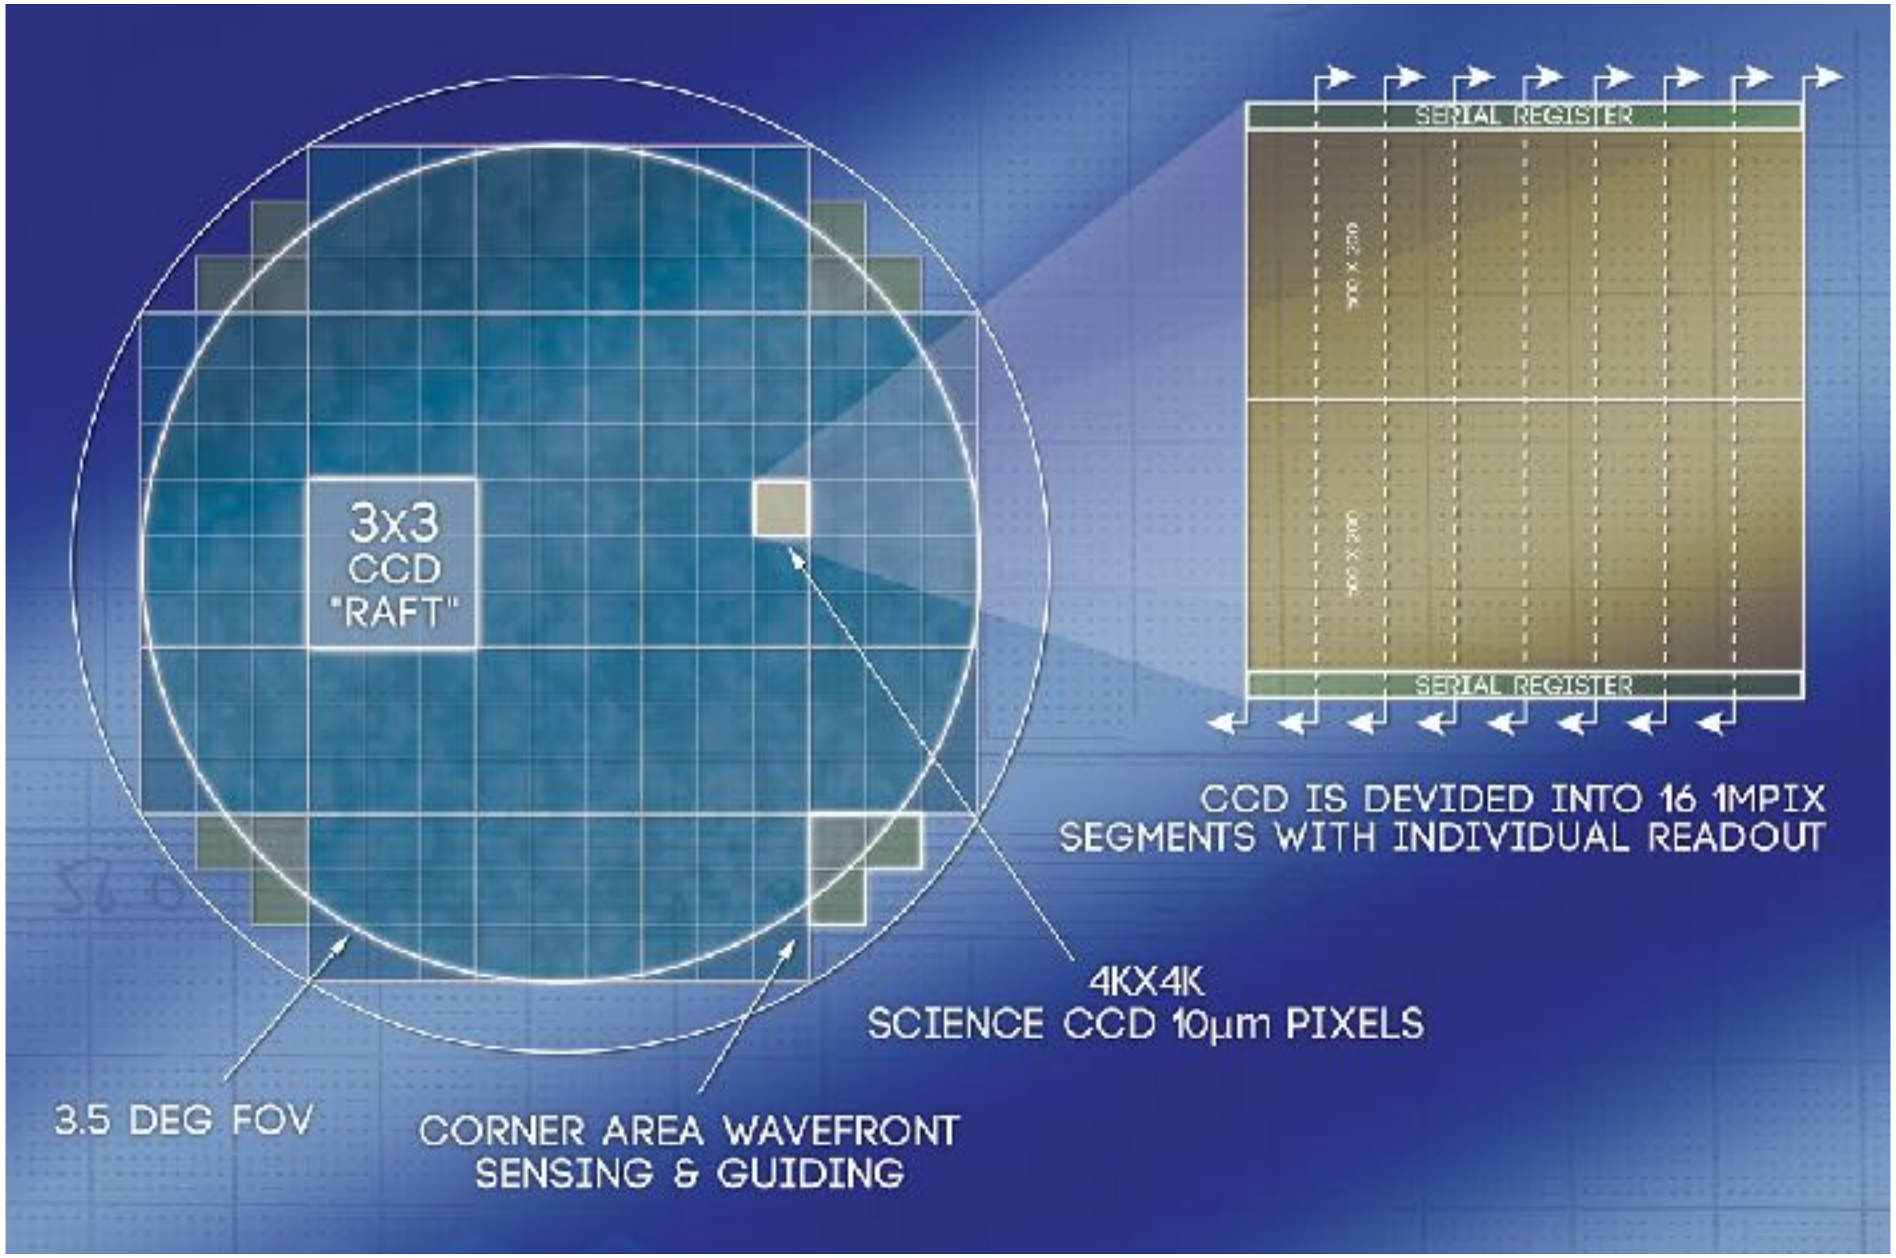

Processing with the LSST Science Pipeline

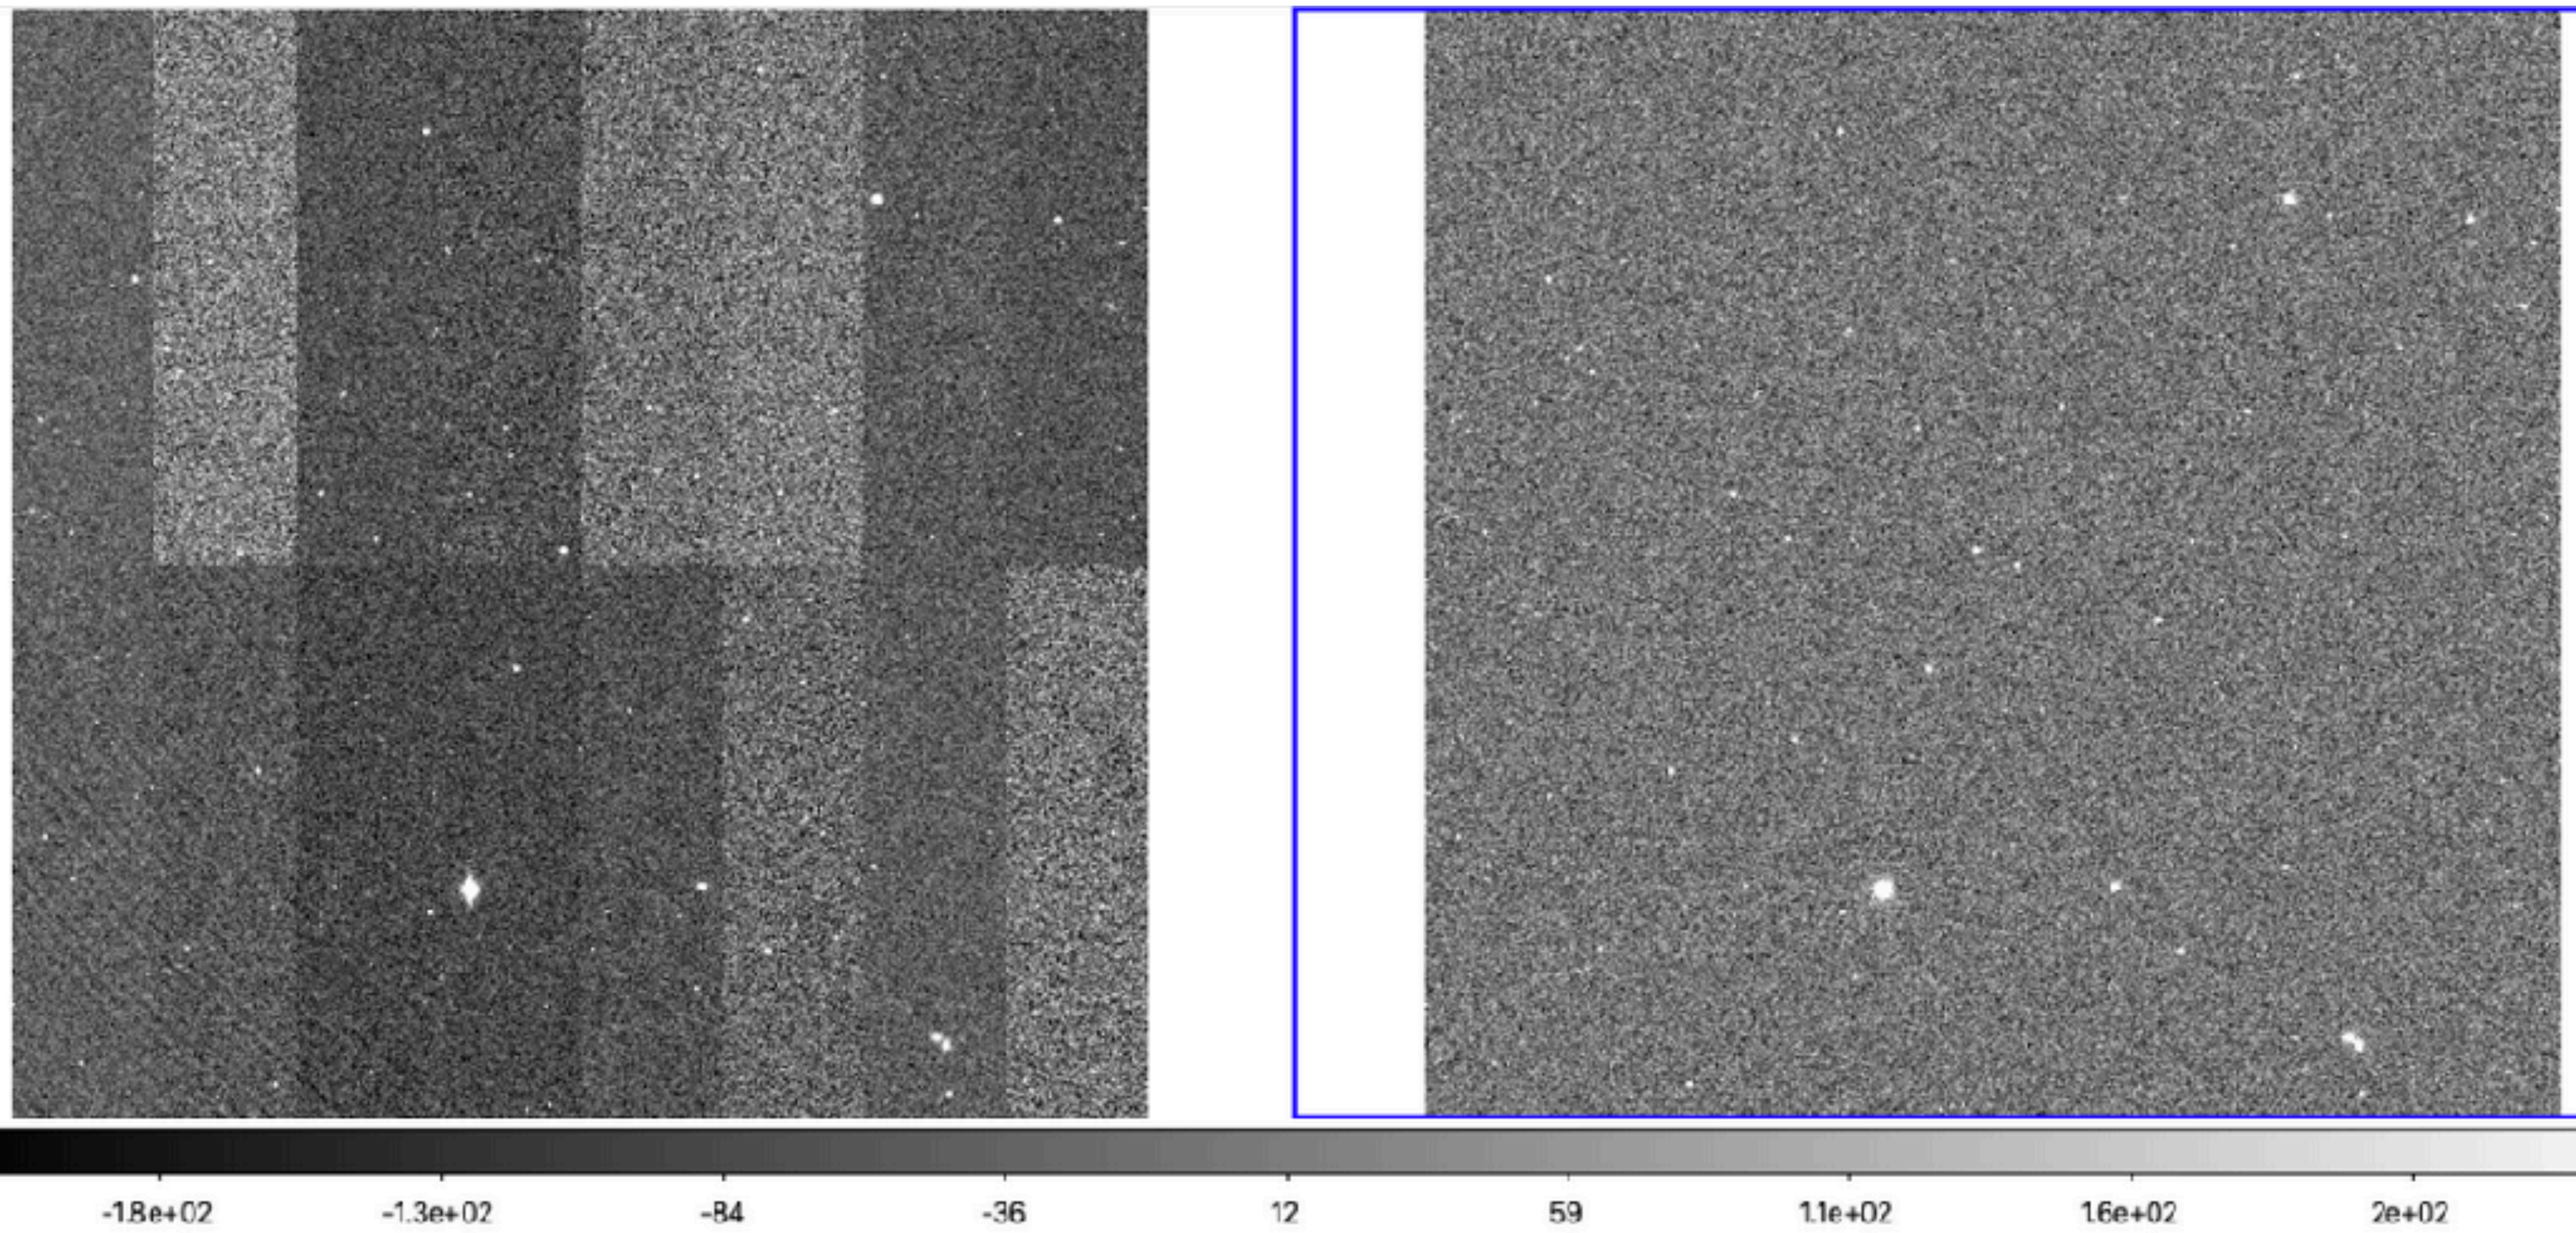

Image Credit: C. Walter, Duke U.

Image Credit: D. Boutingy and the DC2 Team of LSST DESC

Image Credit: D. Boutingy and the DC2 Team of LSST DESC