ChaNGa

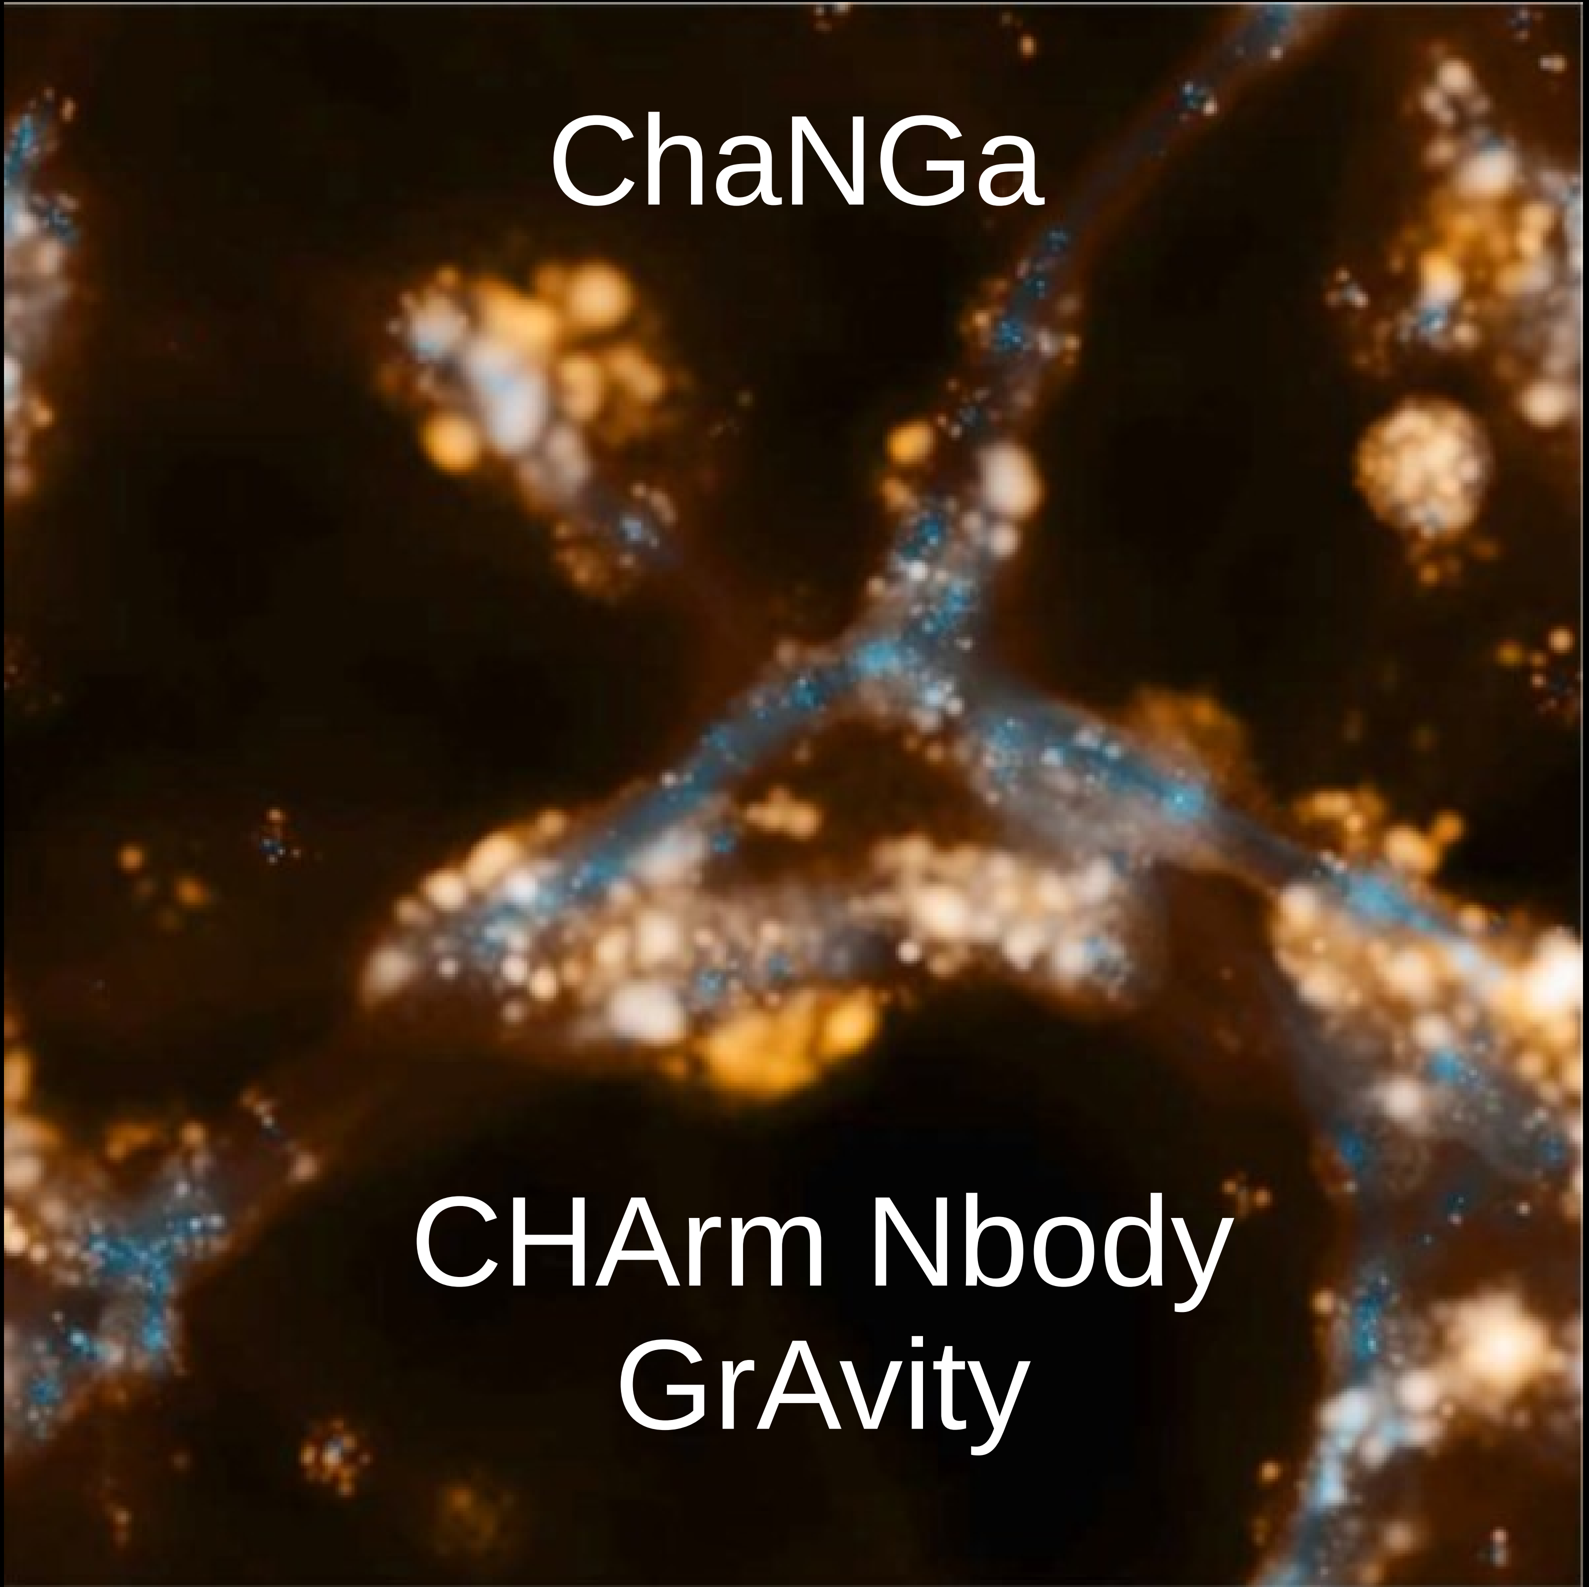

CHArm Nbody GrAvity

PPL PARALLEL PROGRAMMING LAB
DEPT OF COMPUTER SCIENCE, UNIVERSITY OF ILLINOIS

Laxmikant Kale
Filippo Gioachin
Pritish Jetley
Celso Mendes
Amit Sharma
Lukasz Wesolowski
Gengbin Zheng
Edgar Solomonik
Harshitha Menon
Orion Lawlor

Thomas Quinn
Graeme Lufkin
Joachim Stadel
James Wadsley
Greg Stinson

Funding:
NSF
NASA

Outline

- Overview of computational cosmology
- Specific Challenges for Blue Waters
- ChaNGa design
- Recent scaling results
- Recent science results
- Future

History of the Universe

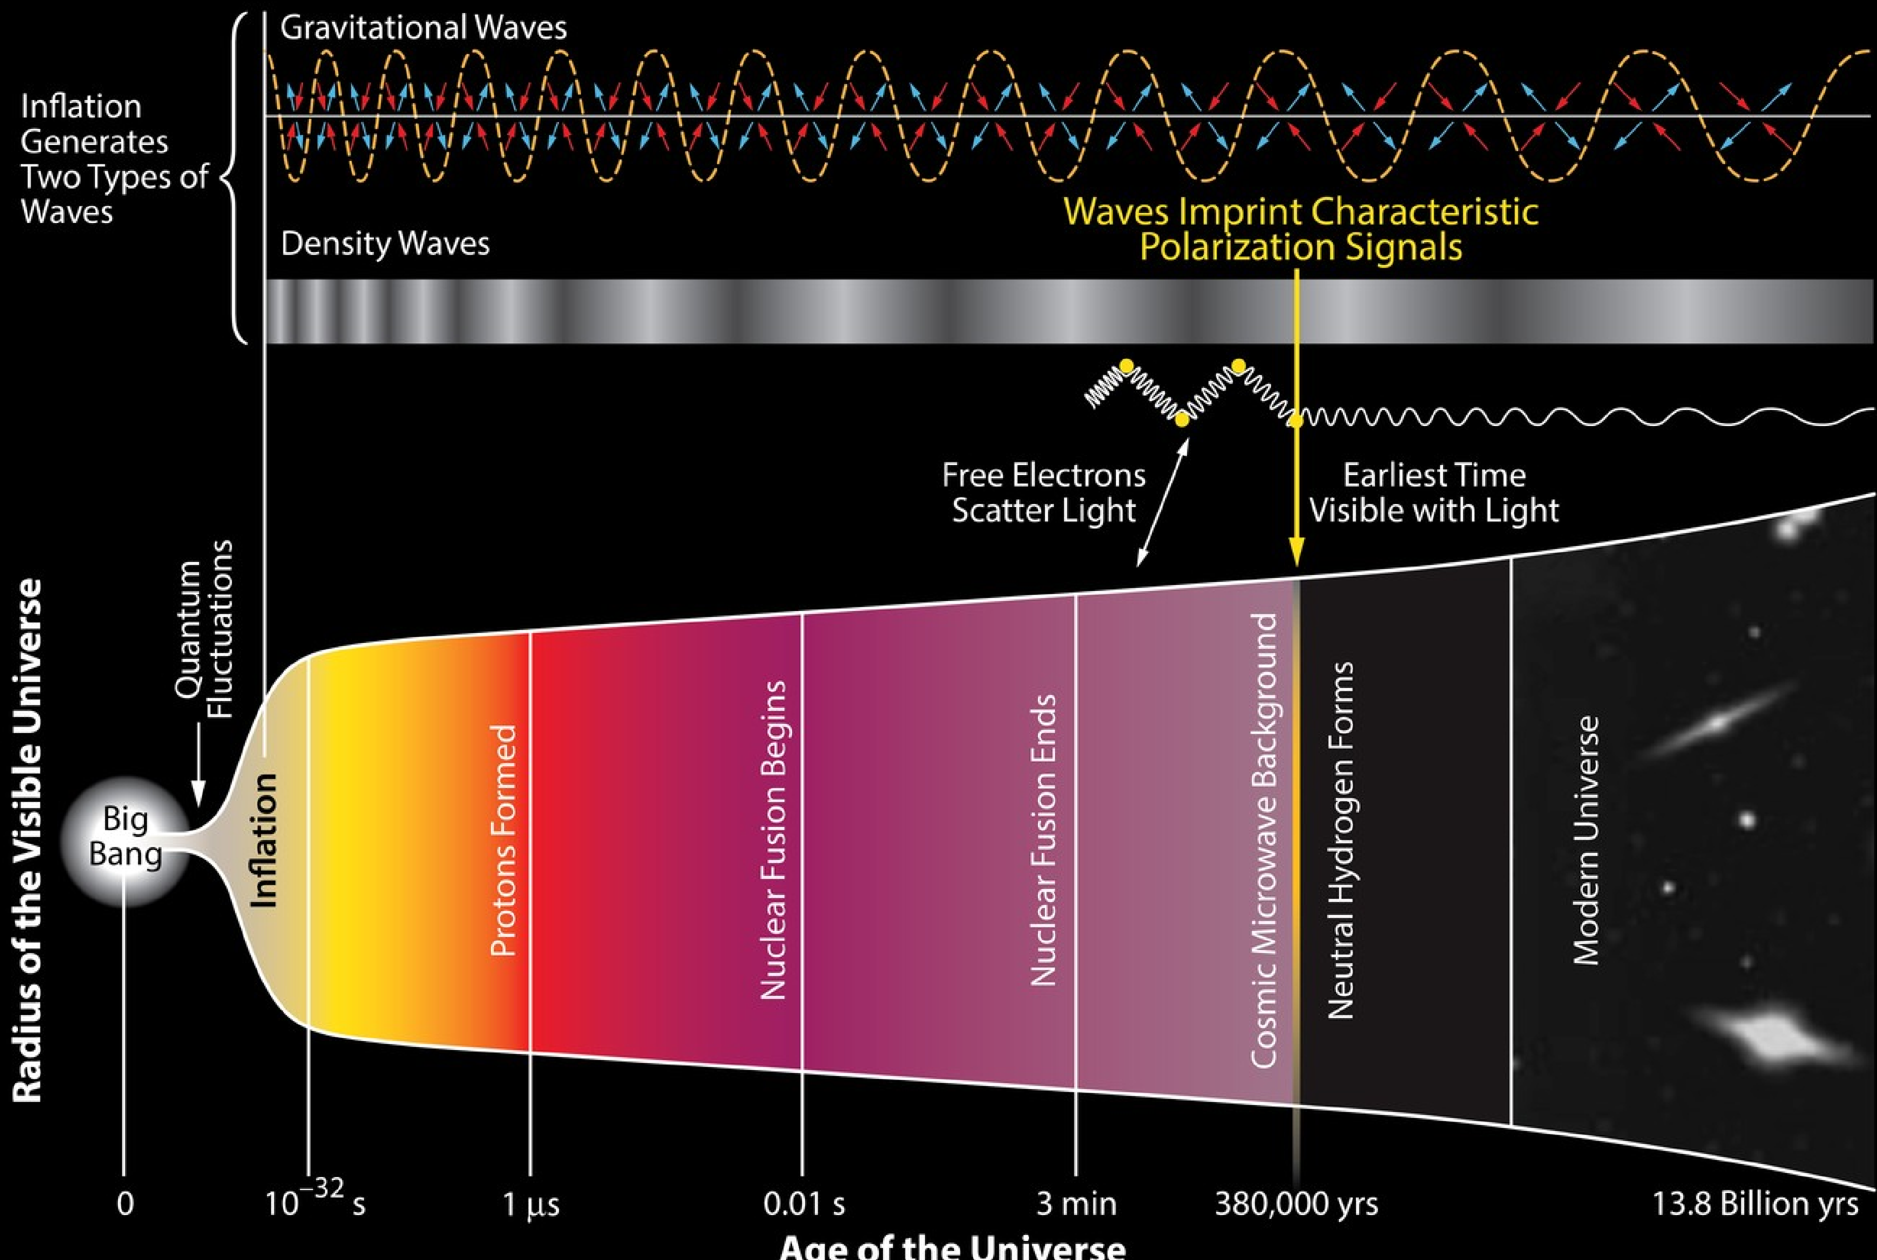

Cosmology at 13.6 Gigayears

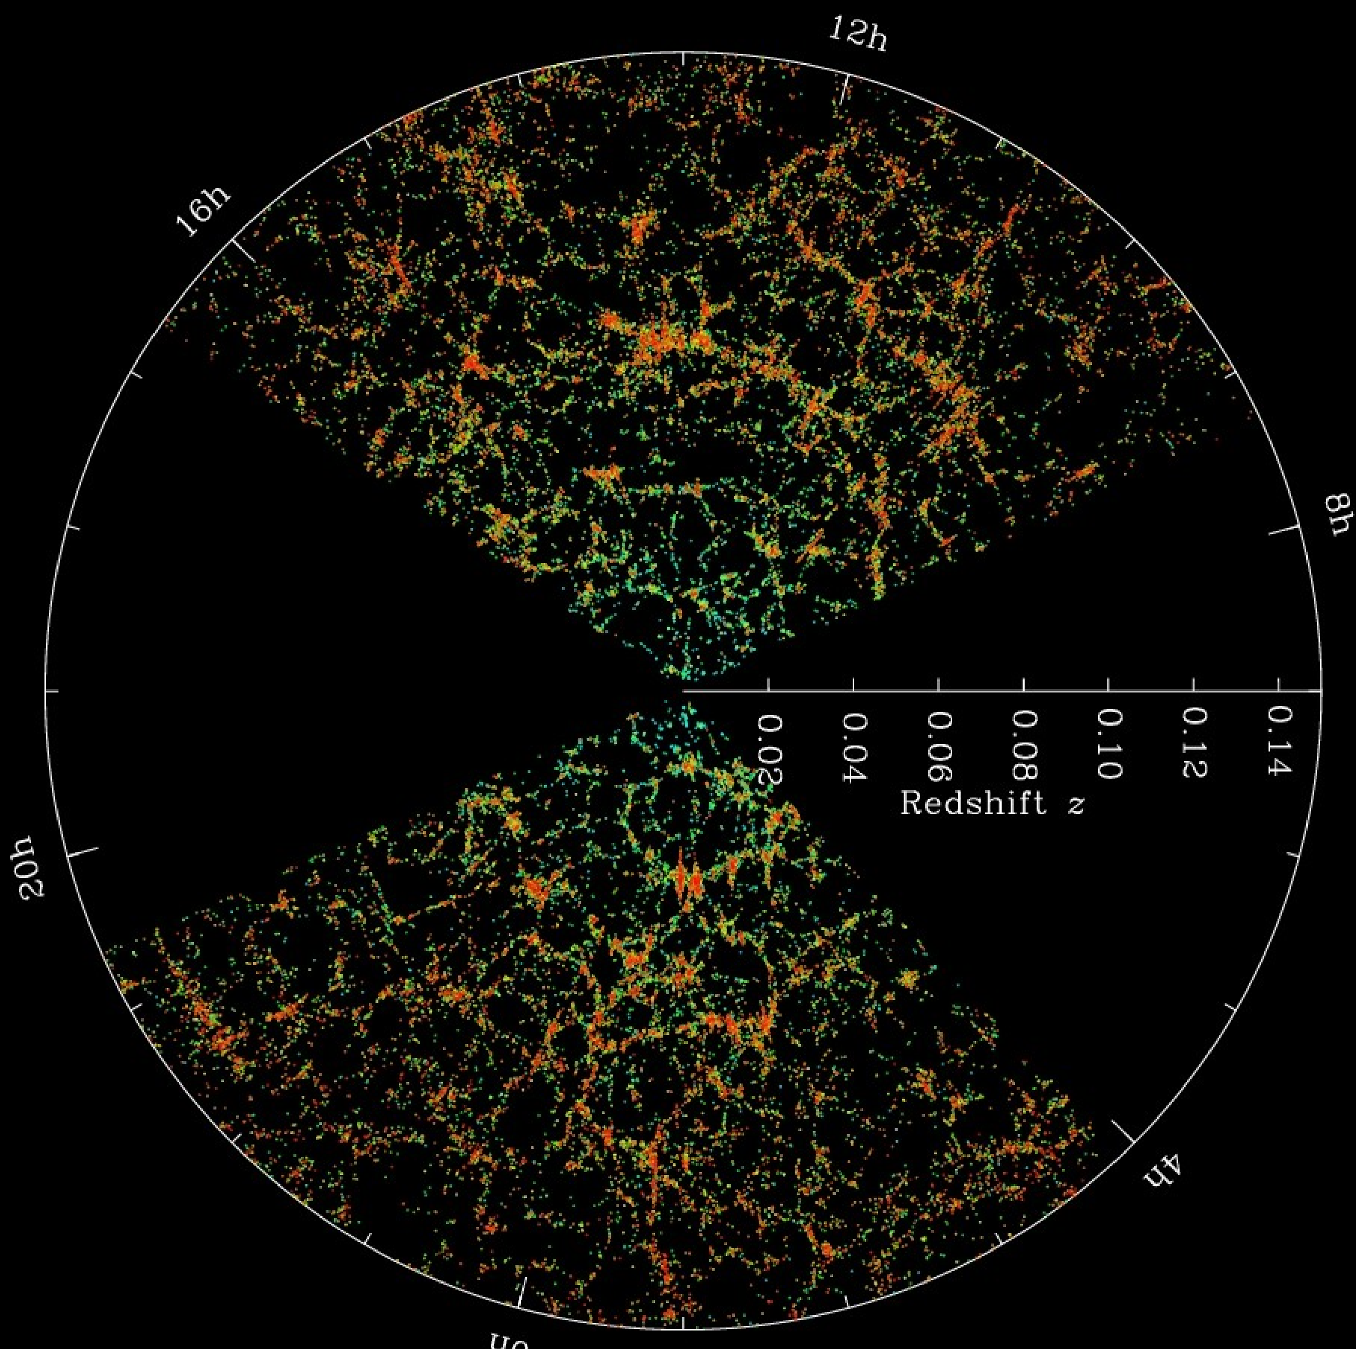

... is not so simple

Fundamental Problem: Dark Matter and Energy: What is it?

- Not baryons
- Simulations show: not known neutrinos
- Candidates:
  - Sterile Neutrinos
  - Axions
  - Lightest SUSY Particle (LSP)

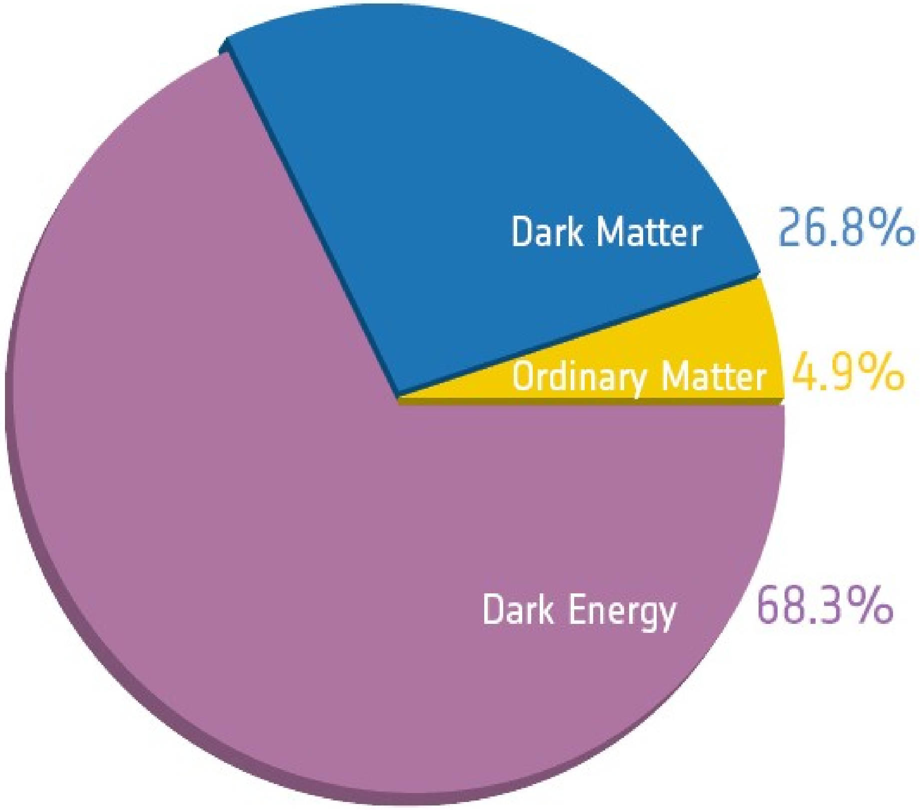

Computational Cosmology

- CMB has fluctuations of 1e-5
- Galaxies are overdense by 1e7
- It happens (mostly) through **Gravitational Collapse**
- Making testable predictions from a cosmological hypothesis requires
  - Non-linear, dynamic calculation
  - e.g. **Computer simulation**

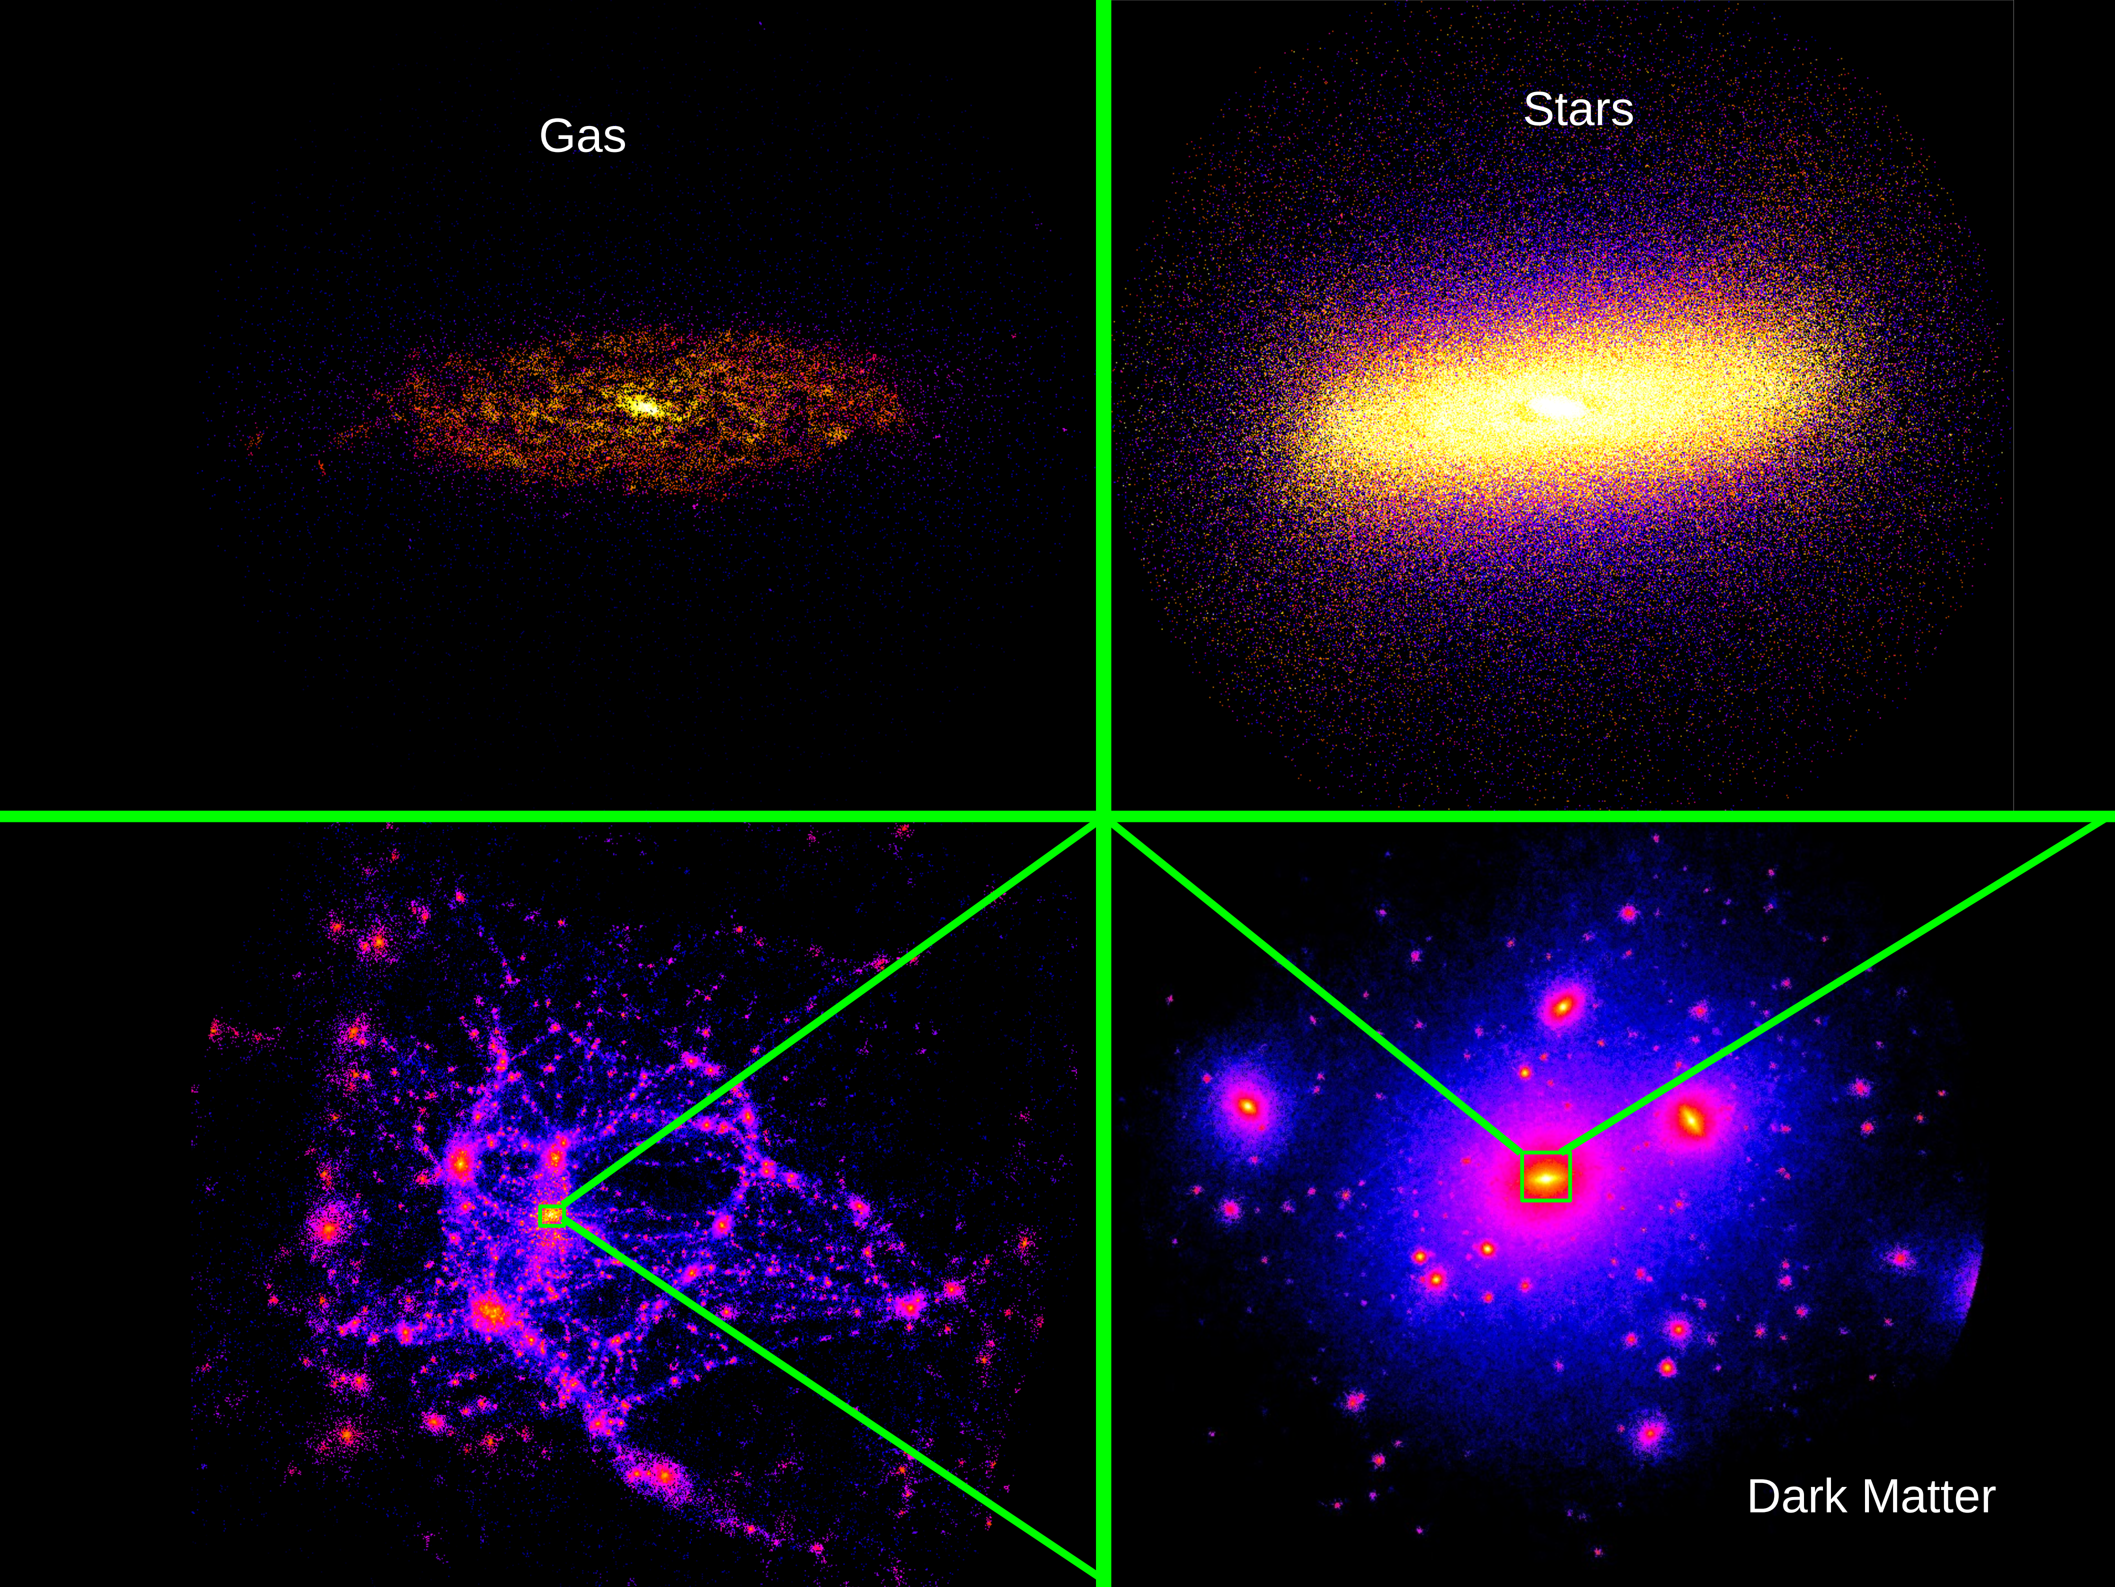
Gas
Stars
Dark Matter

Substructure down to 100 pc

Computational Challenges

- Large spacial dynamic range: > 100 Mpc to < 1 kpc
  - Hierarchical, adaptive gravity solver is needed
- Large temporal dynamic range: 10 Gyr to < 1 Myr
  - Multiple timestep algorithm is needed
- Gravity is a long range force
  - Hierarchical information needs to go across processor domains

Light vs. Matter

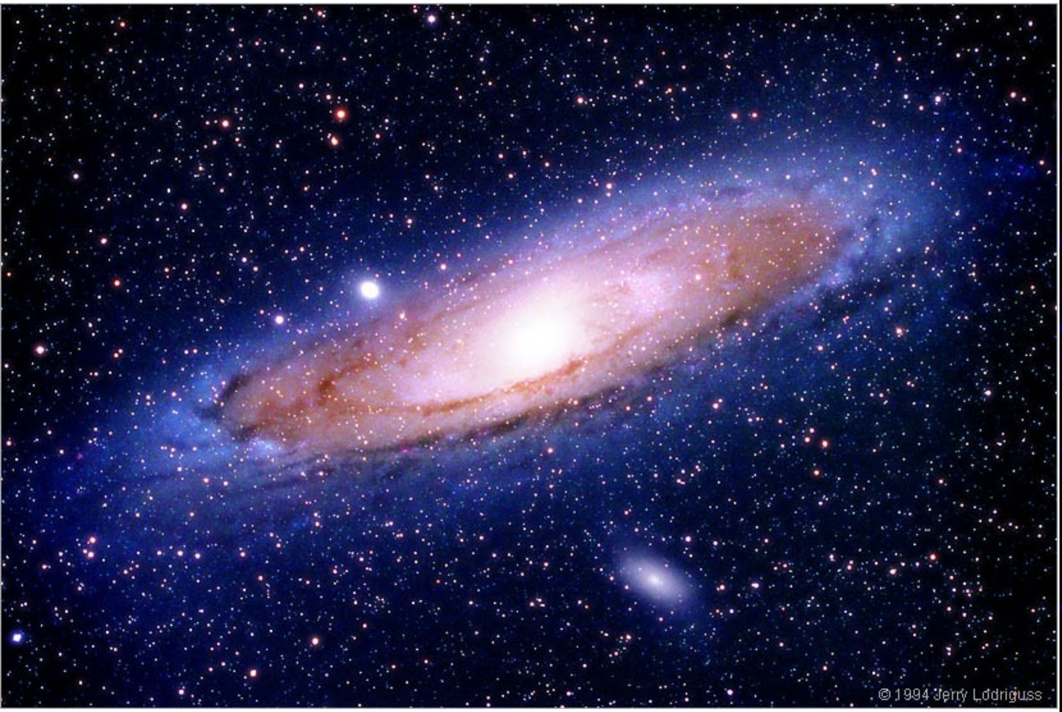

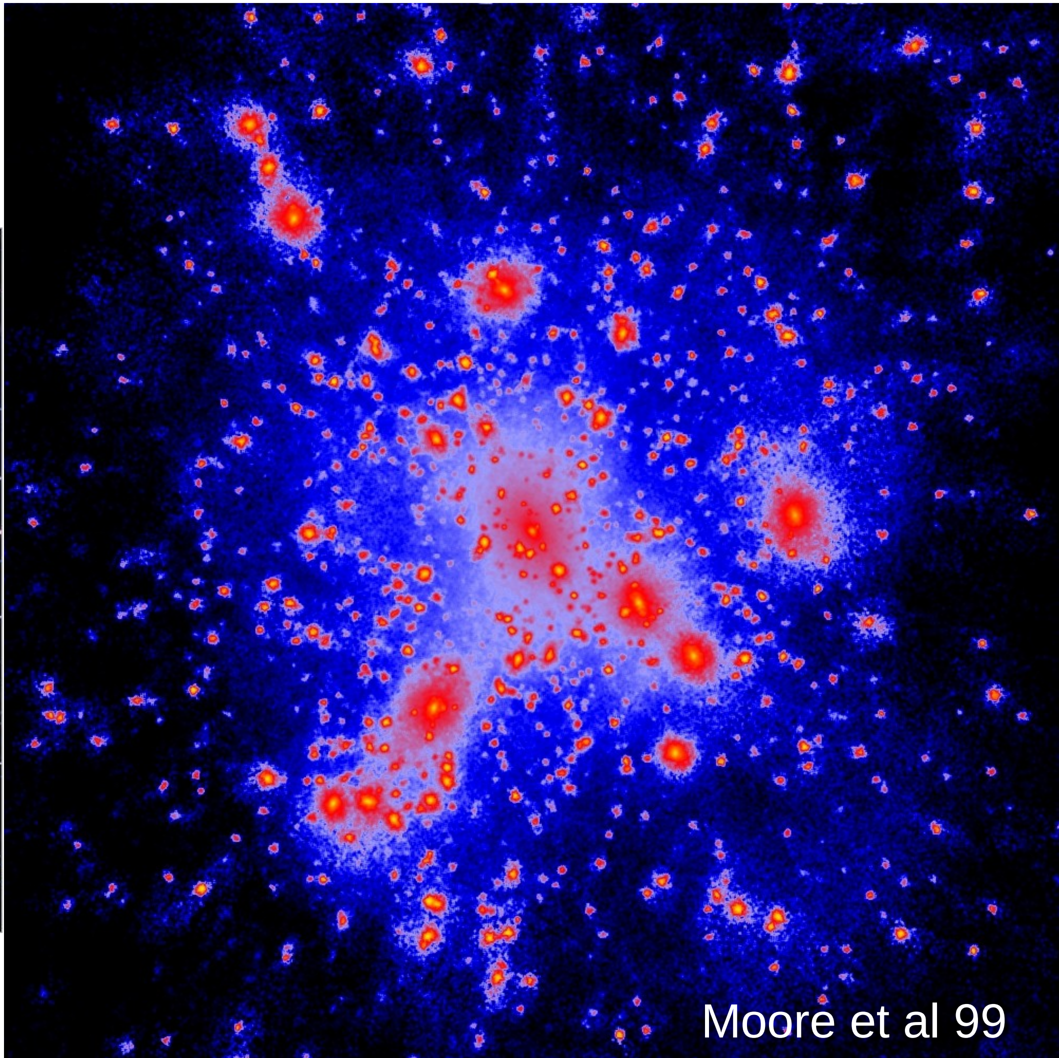

Smooth Particle Hydrodynamics

- Making testable predictions needs Gastrophysics
  - High Mach number
  - Large density contrasts
- Gridless, Lagrangian method
- Galilean invariant
- Monte-Carlo Method for solving Navier-Stokes equation.
- Natural extension of particle method for gravity.

Star Formation/Feedback

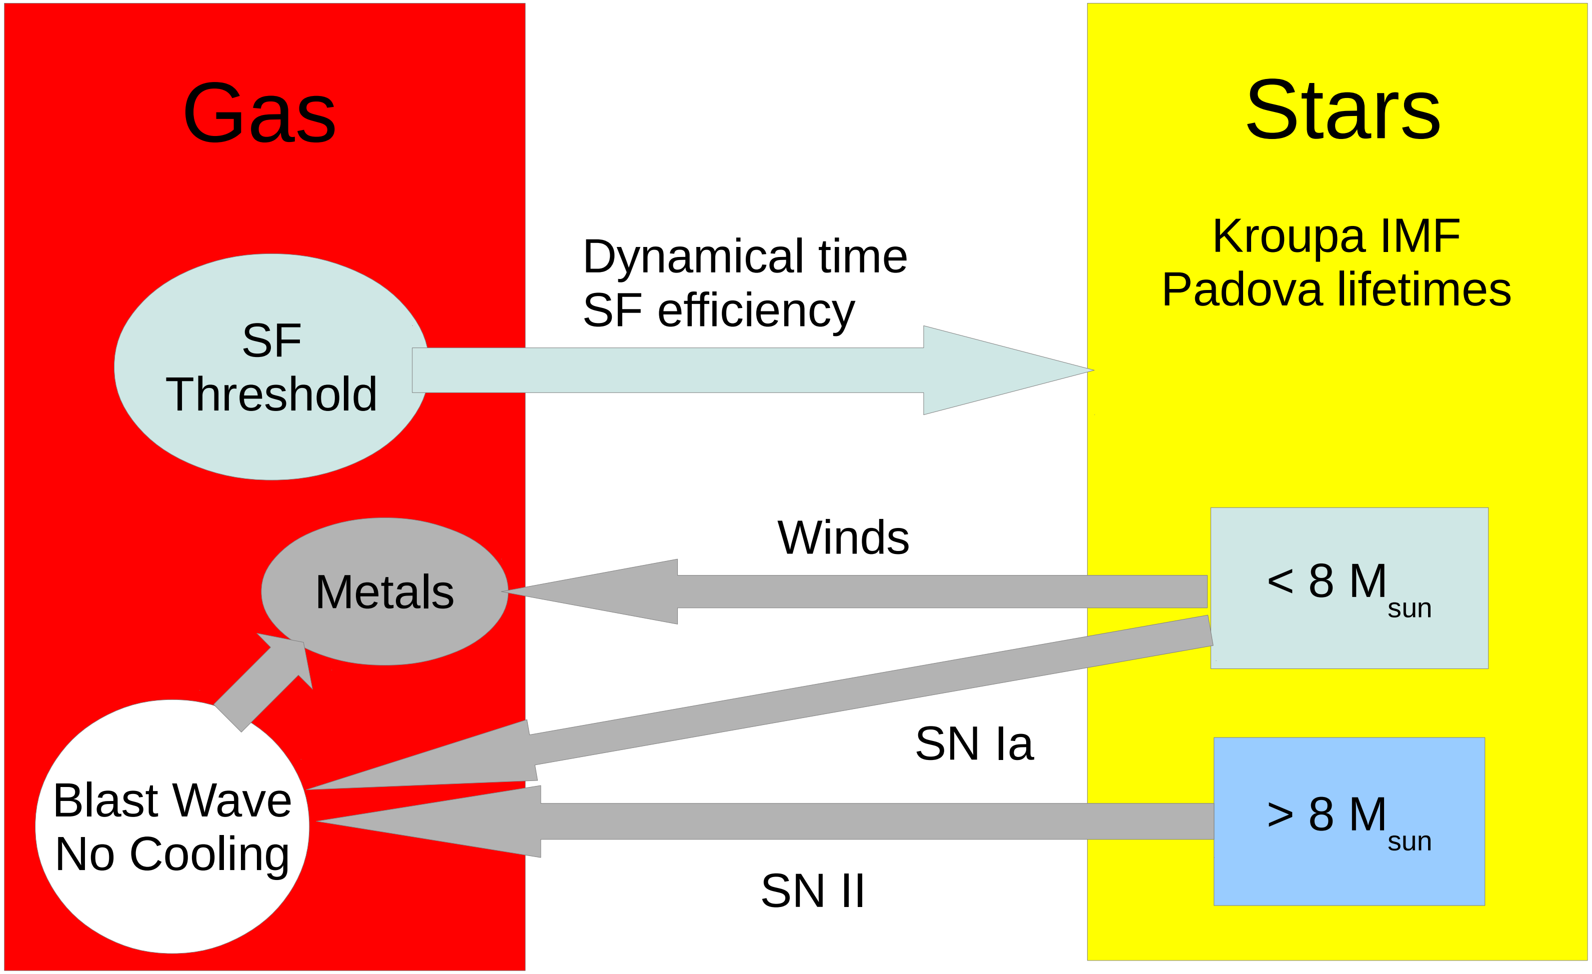

Stinson et al 2006

Galaxies simulated to the present

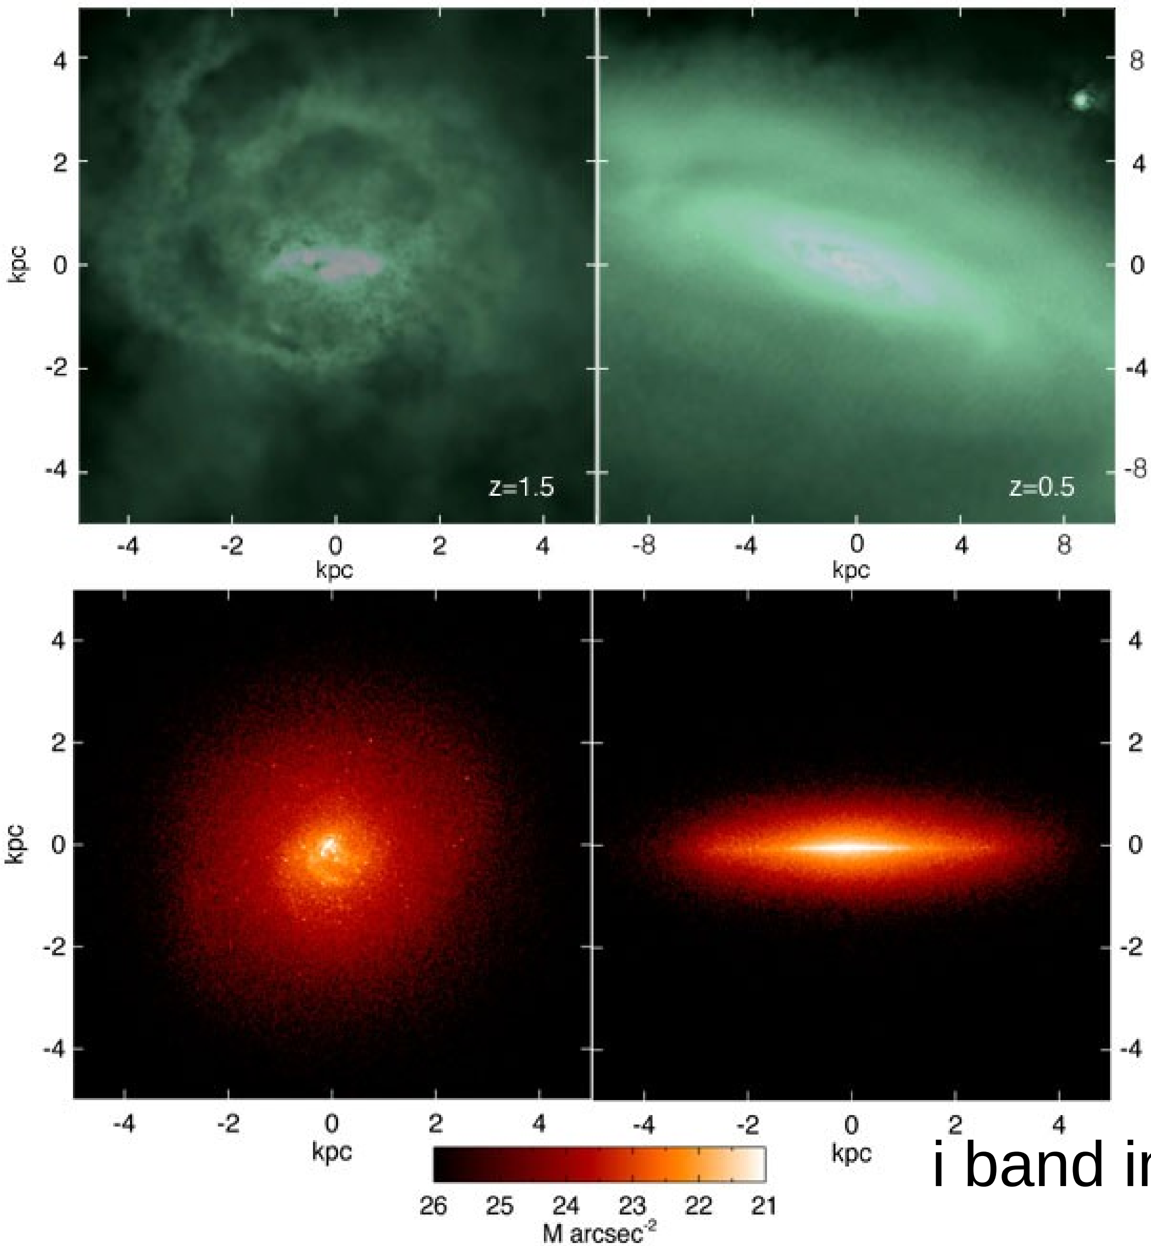

i band image

Reproduces:
* Light profile
* Mass profile
* Star formation
* Angular momentum

The Hubble Ultra Deep Field

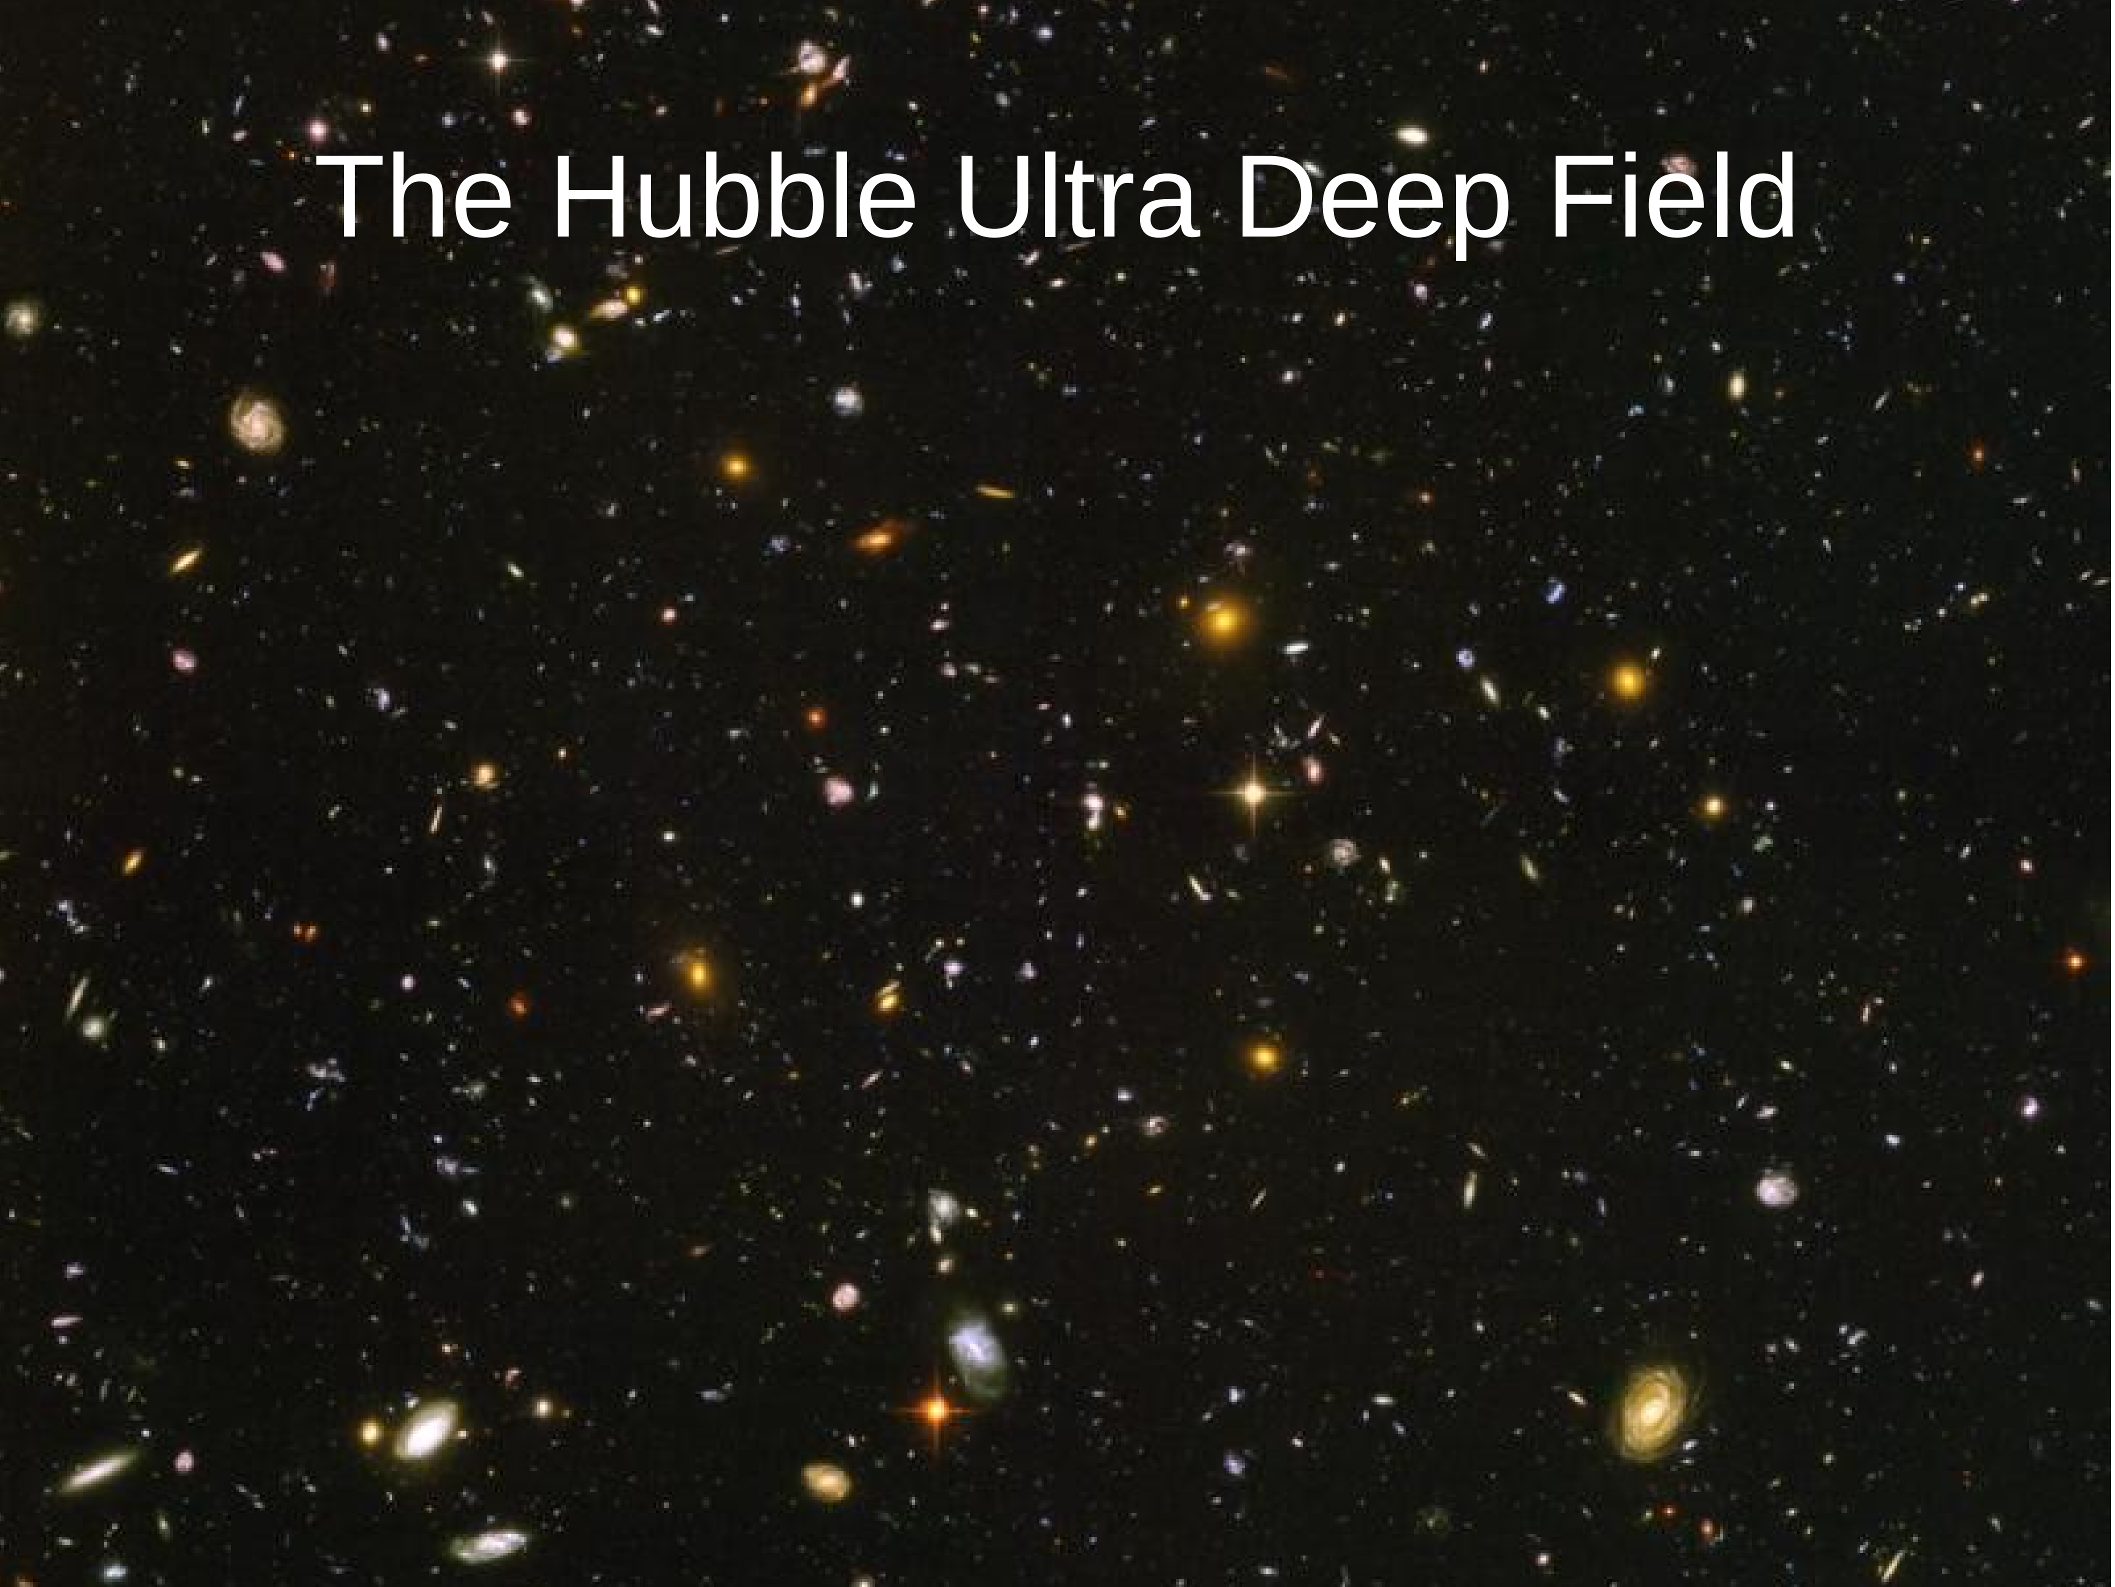

High Redshift Galaxies

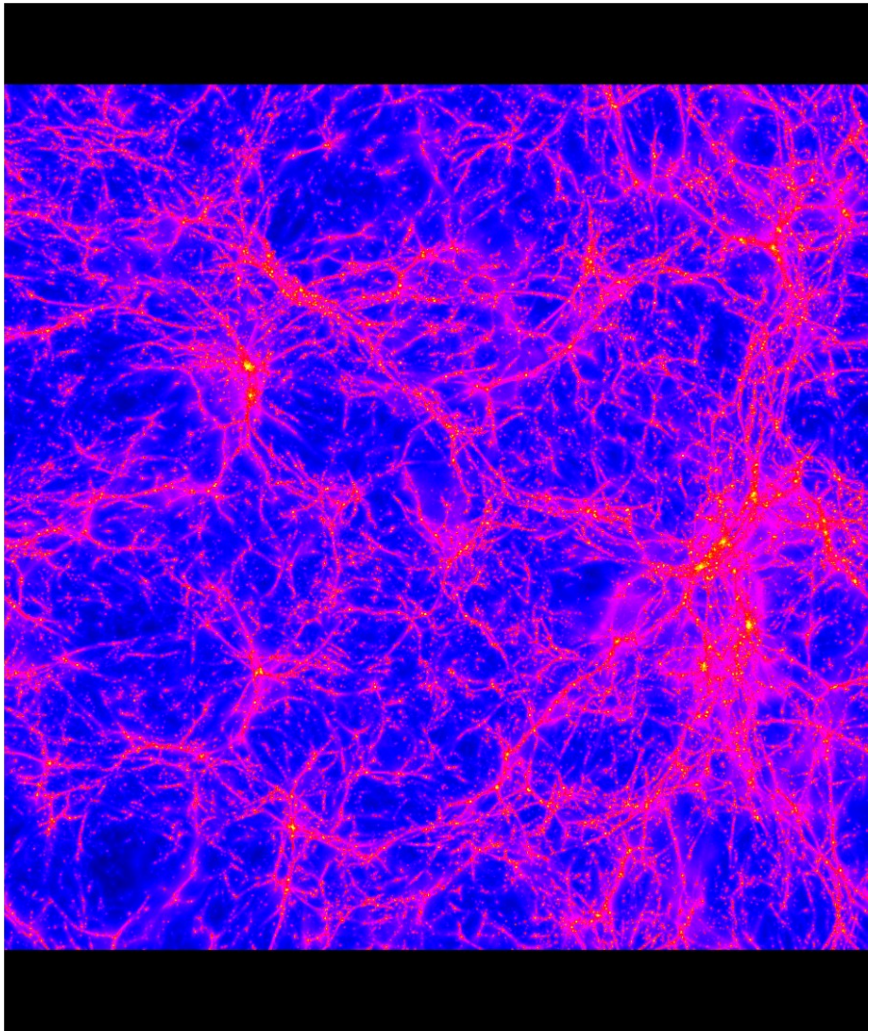

- Galaxies seen by Hubble 12 Gyr ago.
- How do they relate to the Milky Way?
- What is their formation history?
- 300M core-hours on Bluewaters

Cosmo25

- 2 billion particles
- (25 Mpc)^3
- Forces ~ 350pc
- SPH ~ 40 pc
- 100s of galaxies
- 5 TB dataset

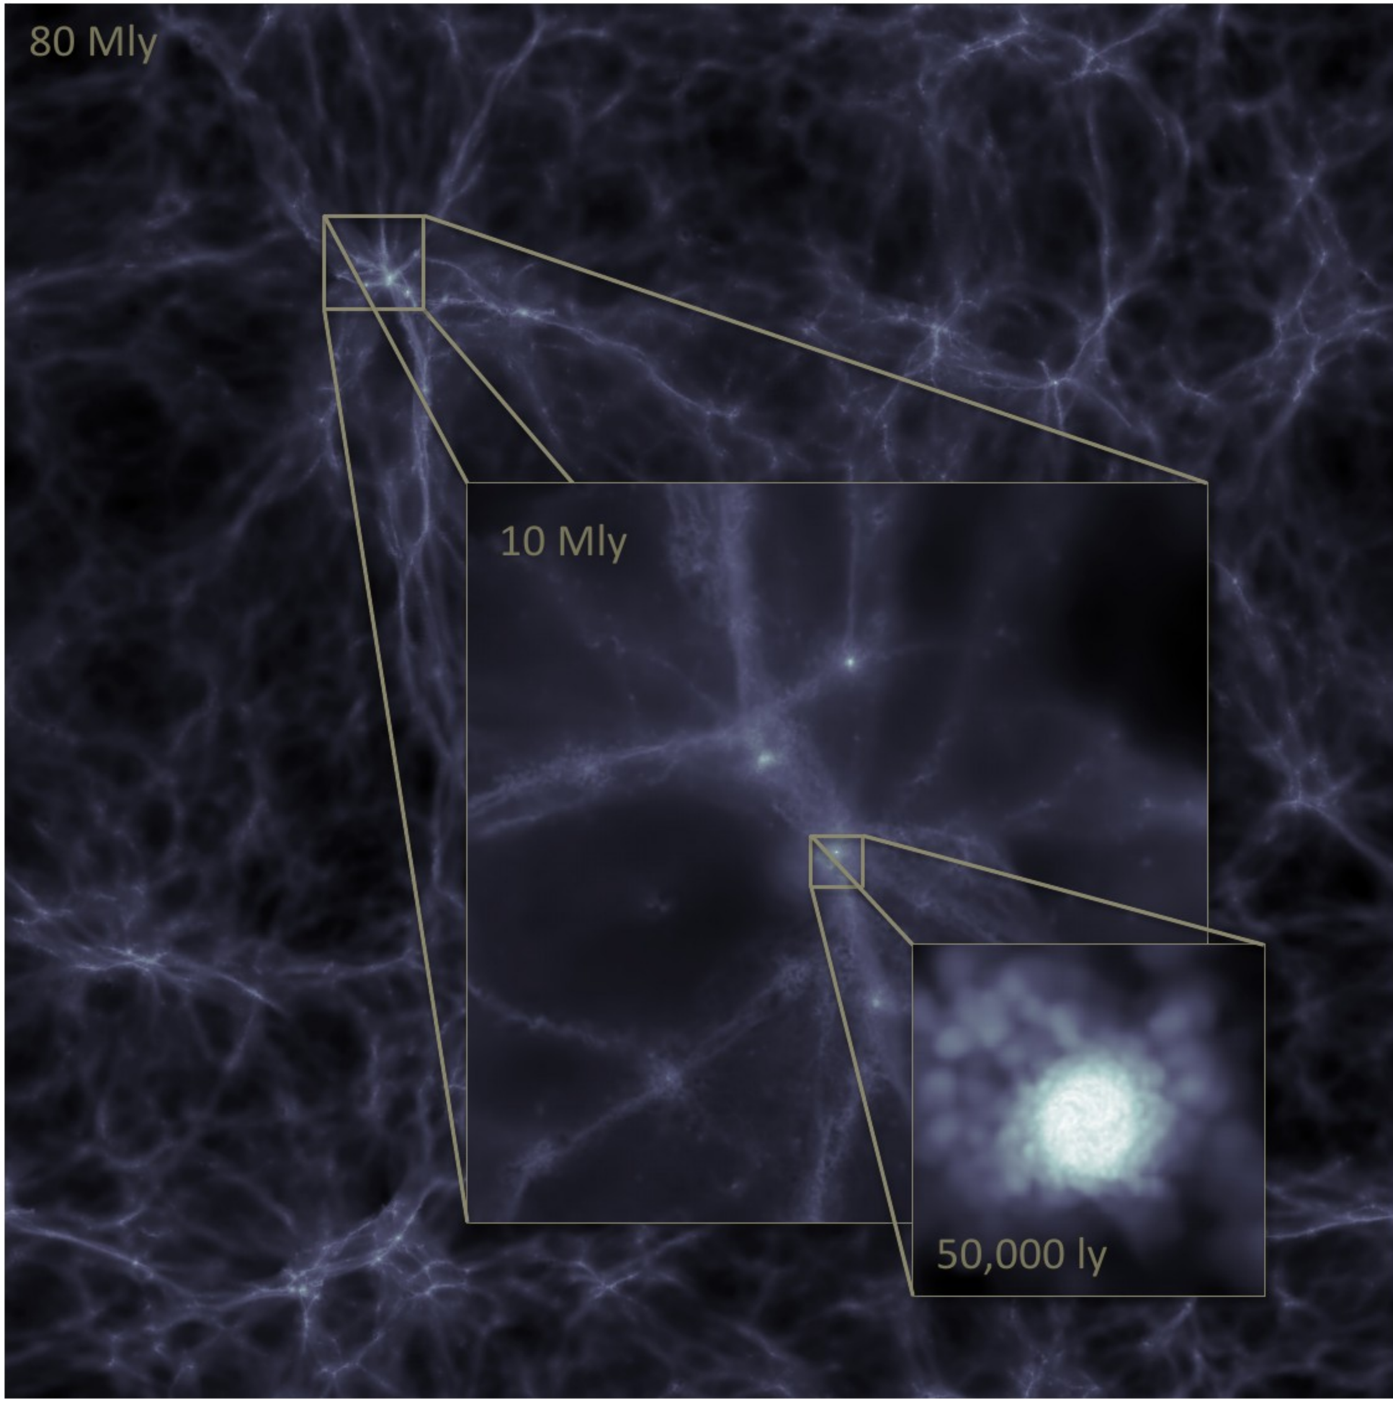

Simulations

	First Stage	Near Future
	Vulcan	Enterprise
Timeline	February 2014	Summer 2014
Size	$(25\ \mathrm{Mpc})^3$	$(25\ \mathrm{Mpc})^3$
Nparticles	2 billion	25 billion
Duration in z	100-4	100-0
Force Resolution	350 pc	175 pc
Morphologies	5e10 M_{tot} (1e9 M_*)	5e9 M_{tot}
Size	5 TB	500 TB
Extra Physics		Black hole feedback H2 regulated star formation

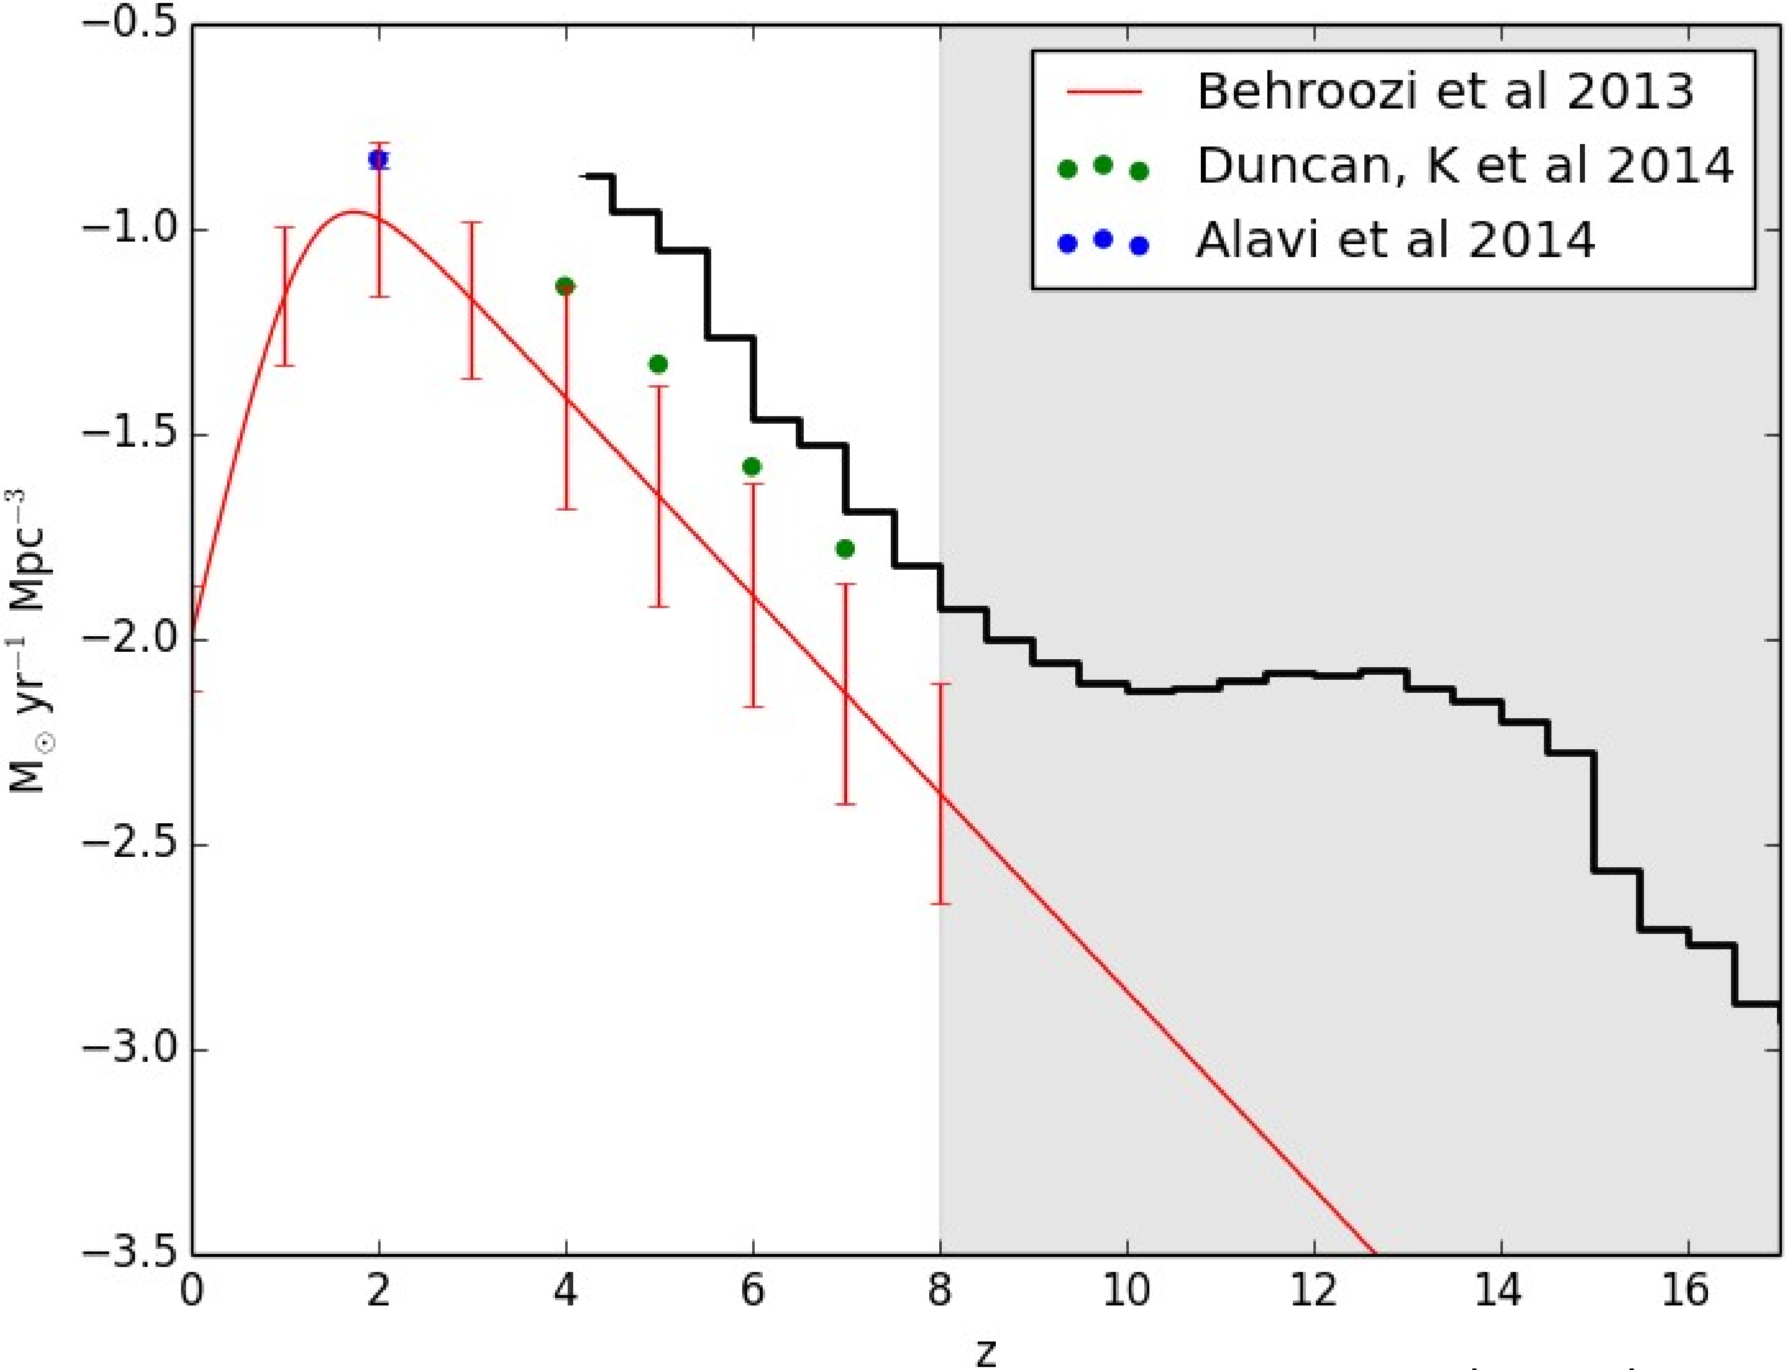

Lauren Anderson, in prep.

ChaNGa Features

- Tree-based gravity solver
- High order multipole expansion
- Periodic boundaries (if needed)
- Individual multiple timesteps
- Dynamic load balancing with choice of strategies
- Checkpointing (via migration to disk)
- Visualization

Latency hiding strategies

- Multiple “treepieces”/core (over decomposition)
- Division into multiple work units (*all concurrently*)
  - Off processor gravity treewalk
  - SPH treewalk
  - Local gravity treewalk
  - Ewald summation
- Method prioritization
  - Data requests get high priority

Overall Algorithm

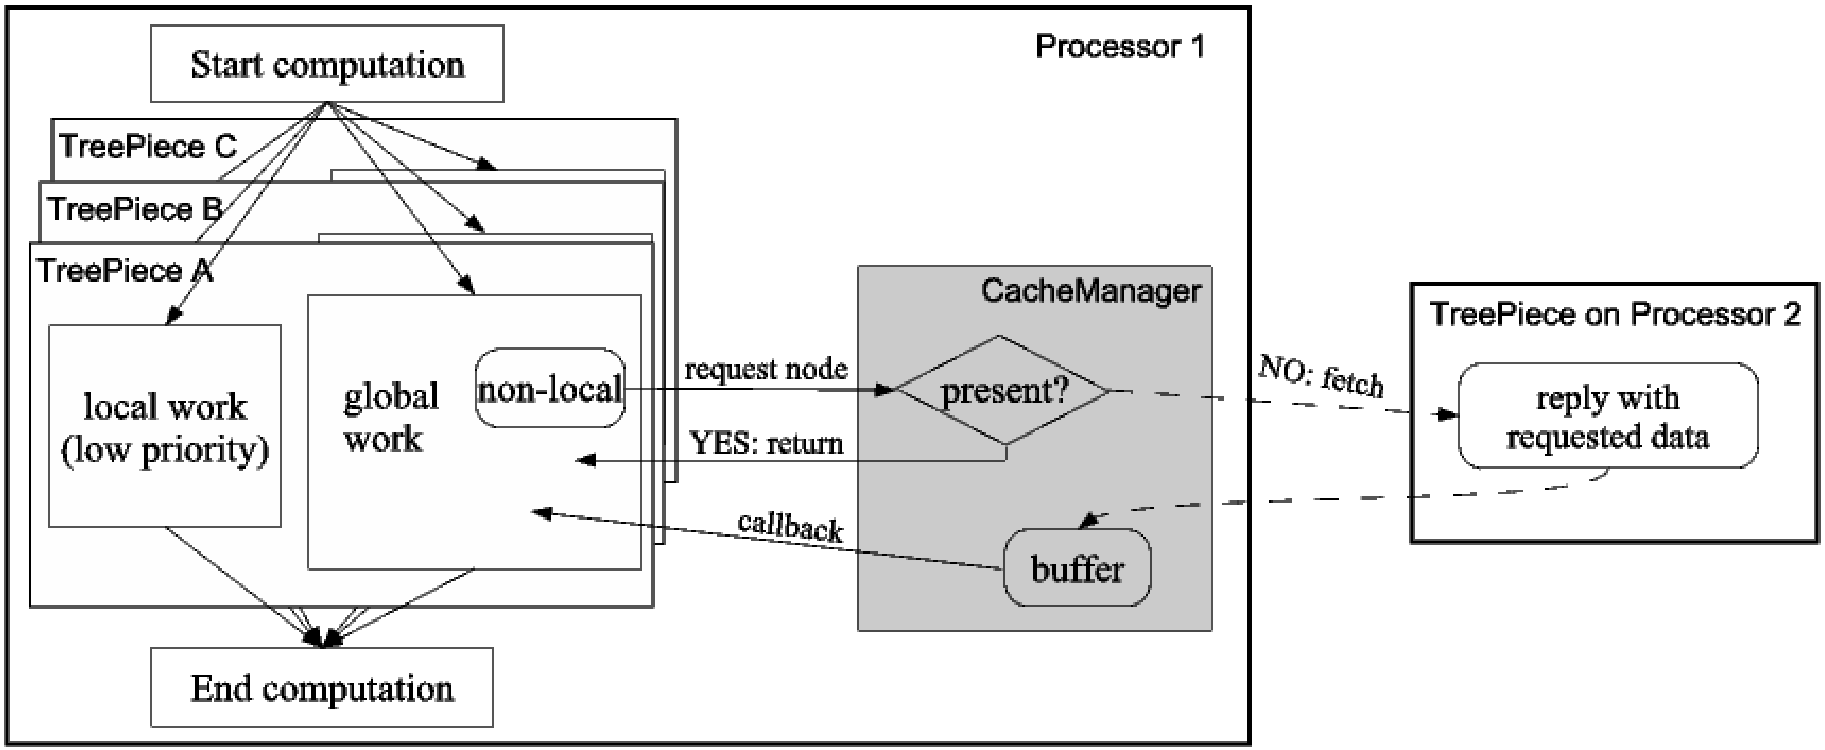

Overlap of Phases

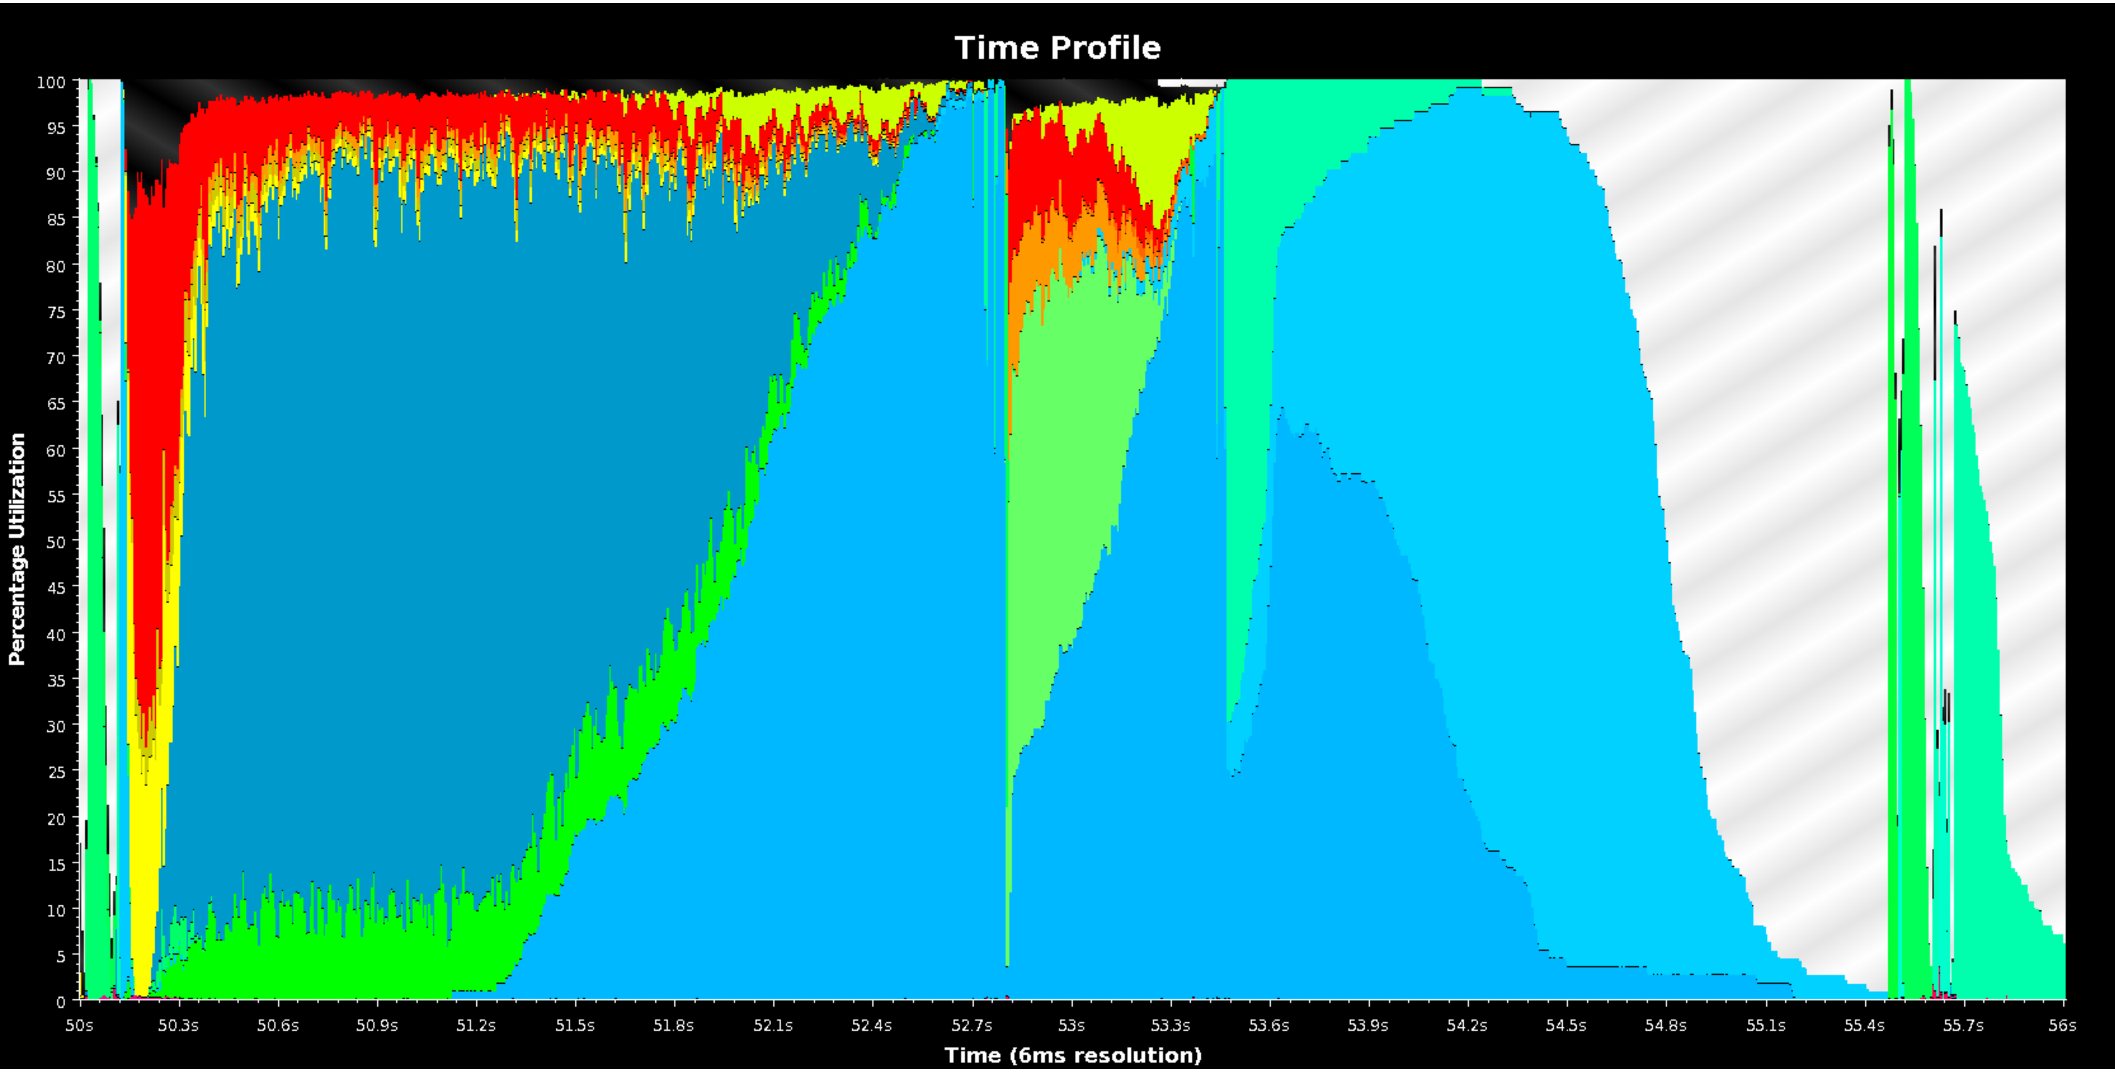

Gravity Hydrodynamics

Scaling to .5M cores

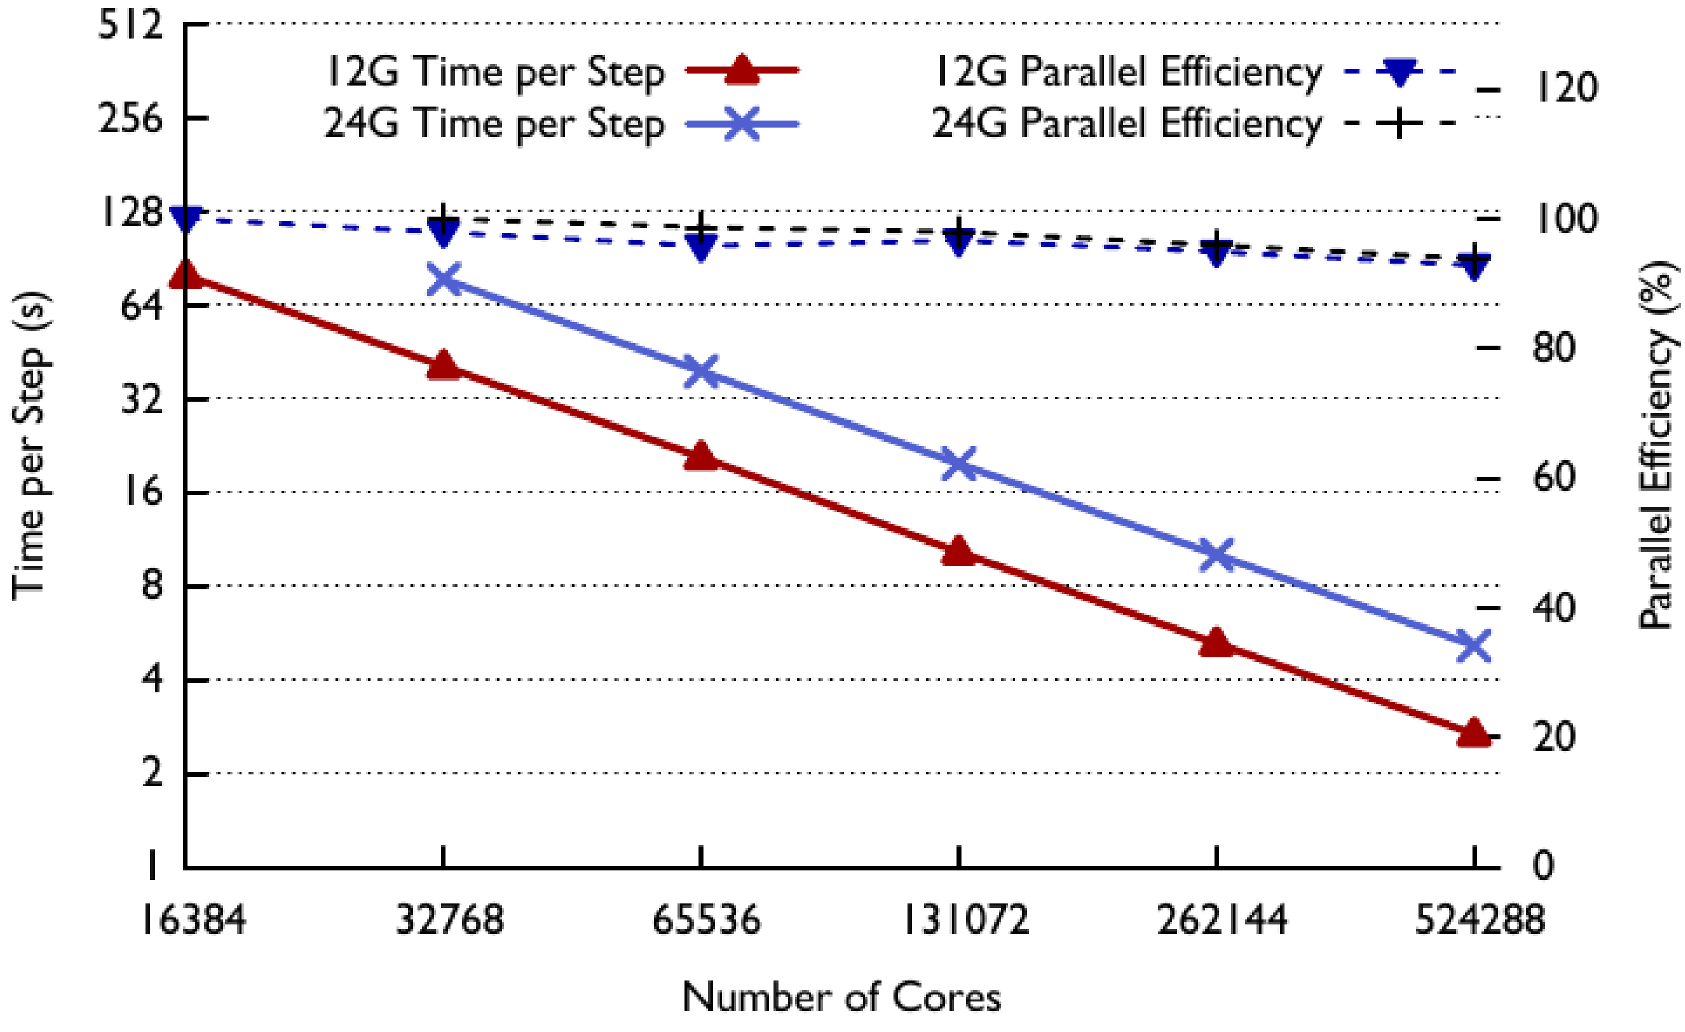

Optimizations for Large Core Count

- Domain Decomposition
  - Reuse previous domain information
  - Only re-decompose when necessary
  - Optimize sort
  - Quiescence detection for particle migration
- Hierarchical Load Balancing
- Treebuilding and approximate remote node location

Clustered/Multistepping Challenges

- Load/particle imbalance
- Communication imbalance
- Fixed costs:
  - Domain Decomposition
  - Load balancing
  - Tree build

Load Variance

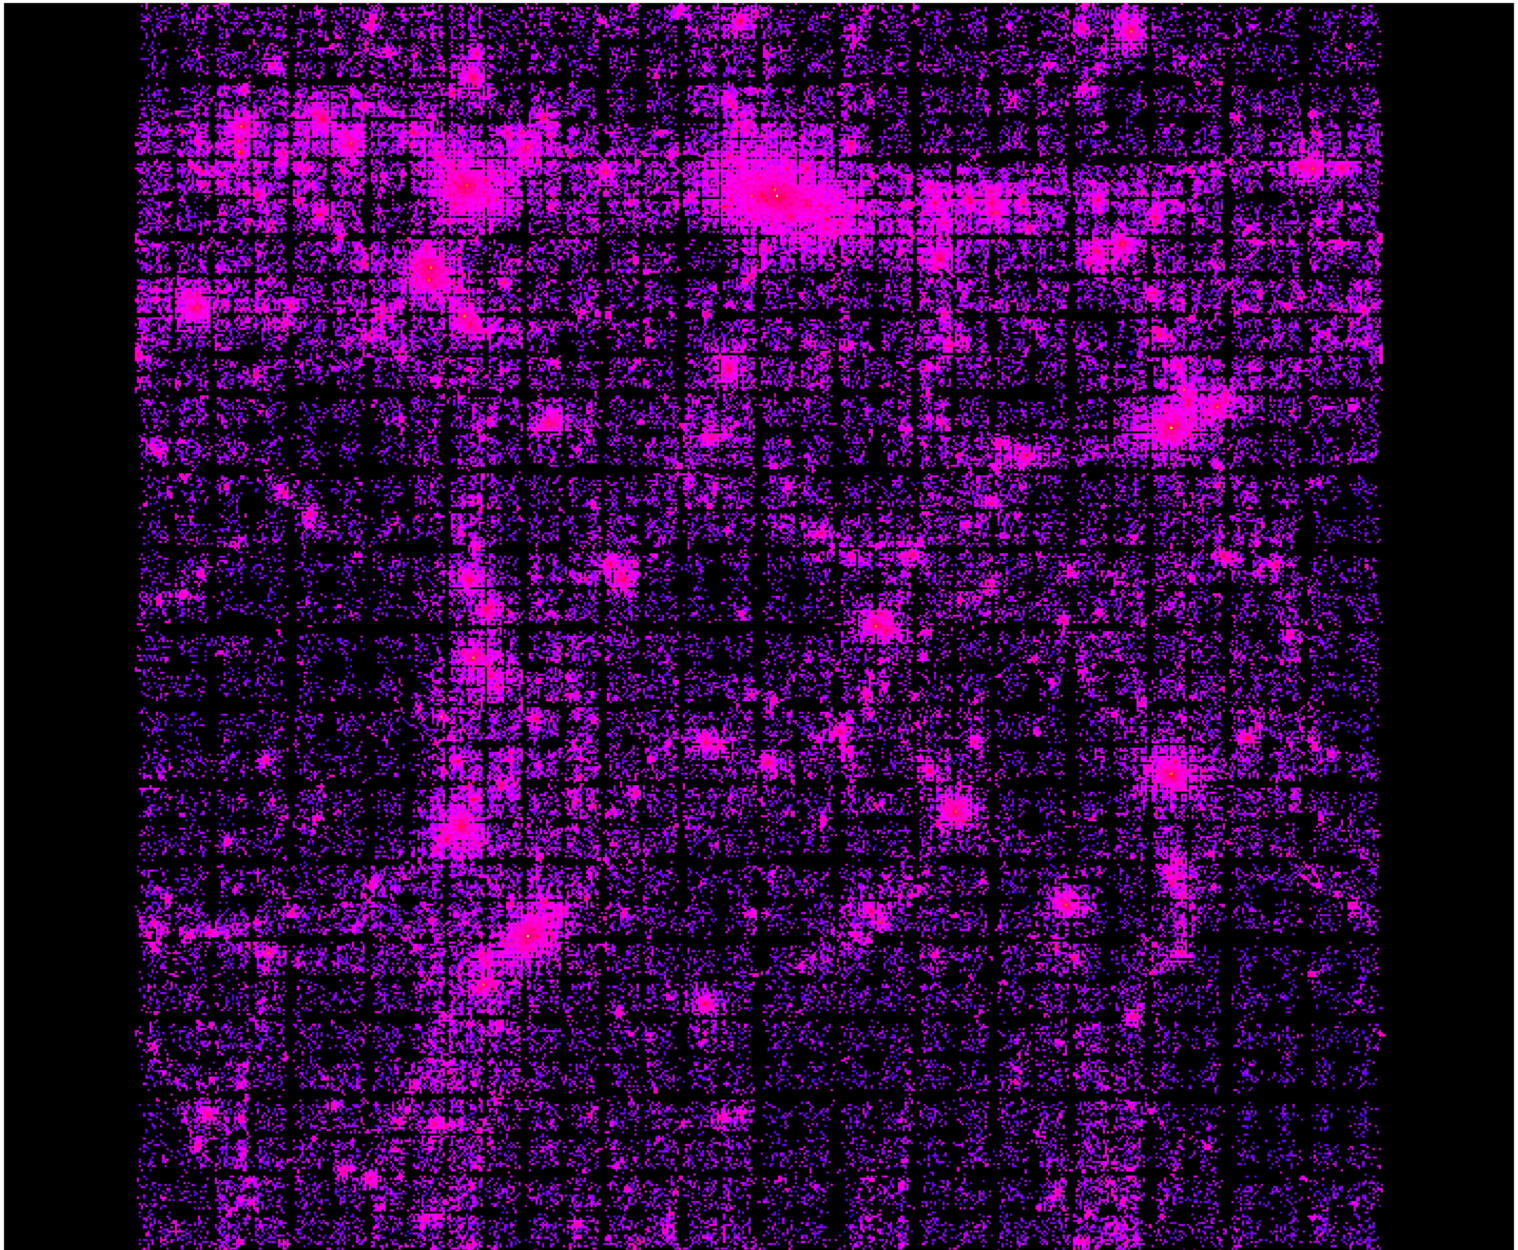

ORB Load Balancing

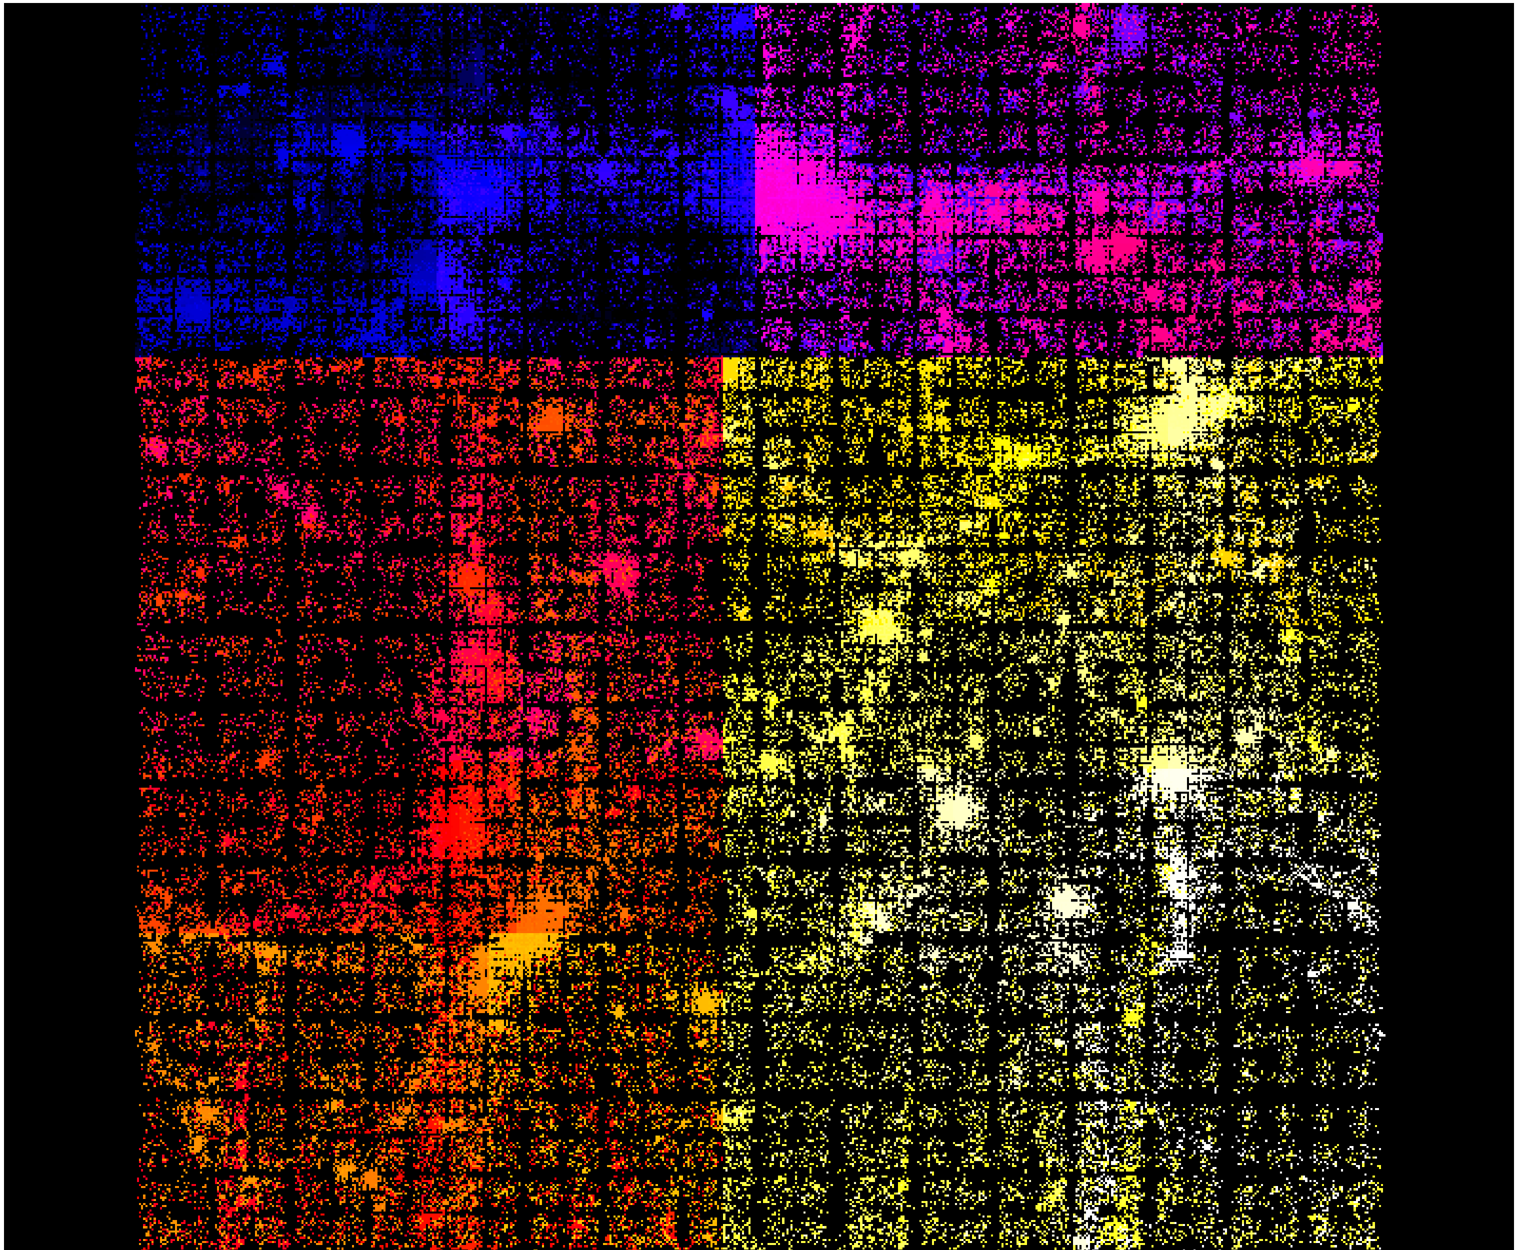

Load distributions

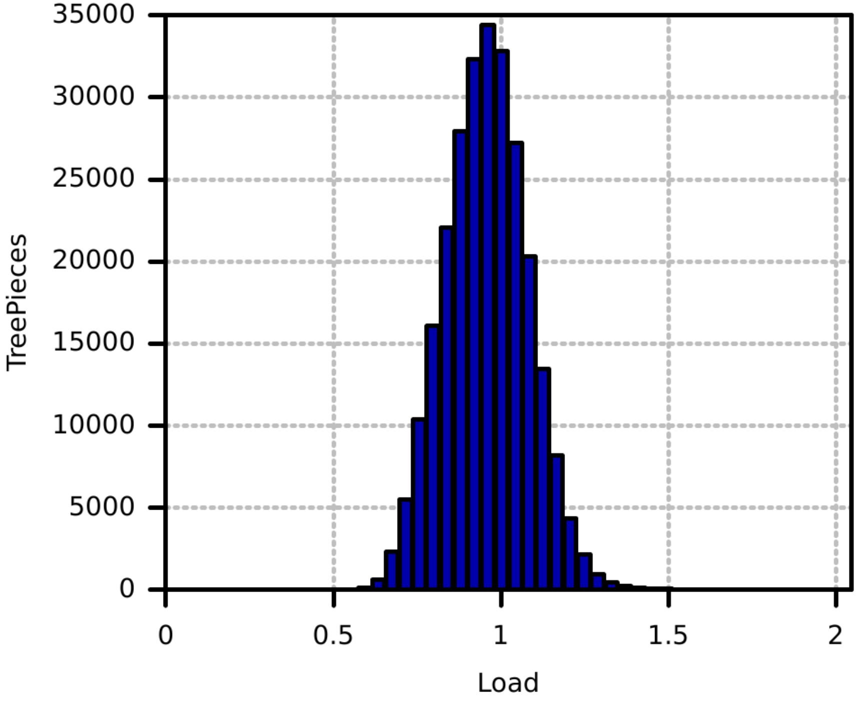

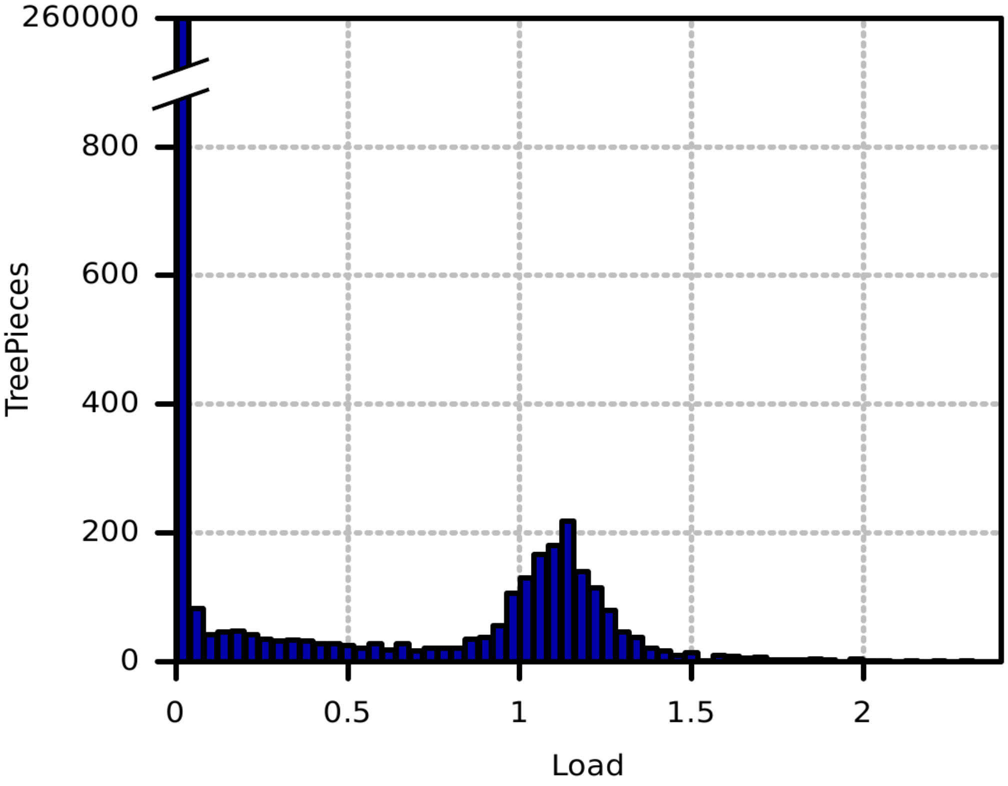

Intra-node work balancing

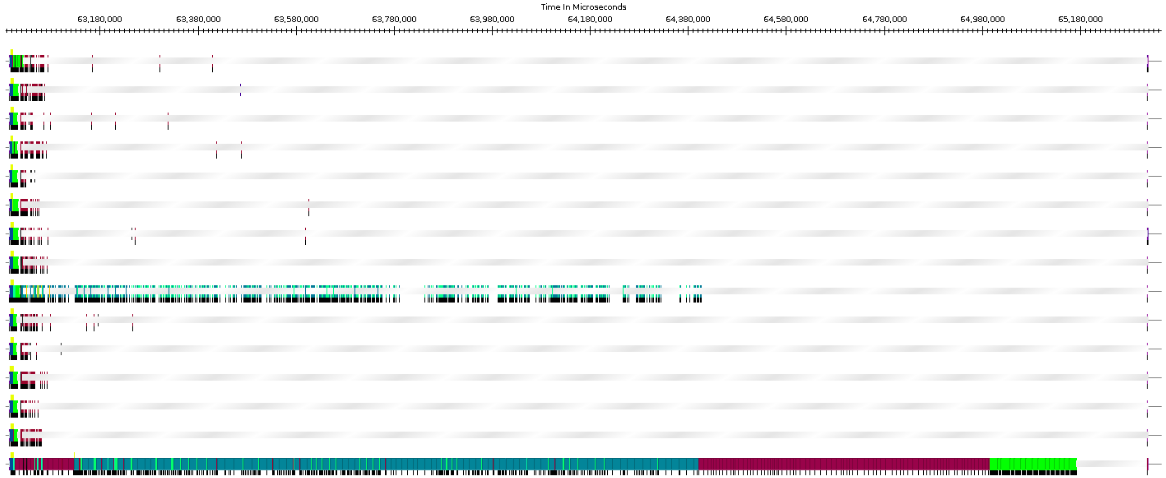

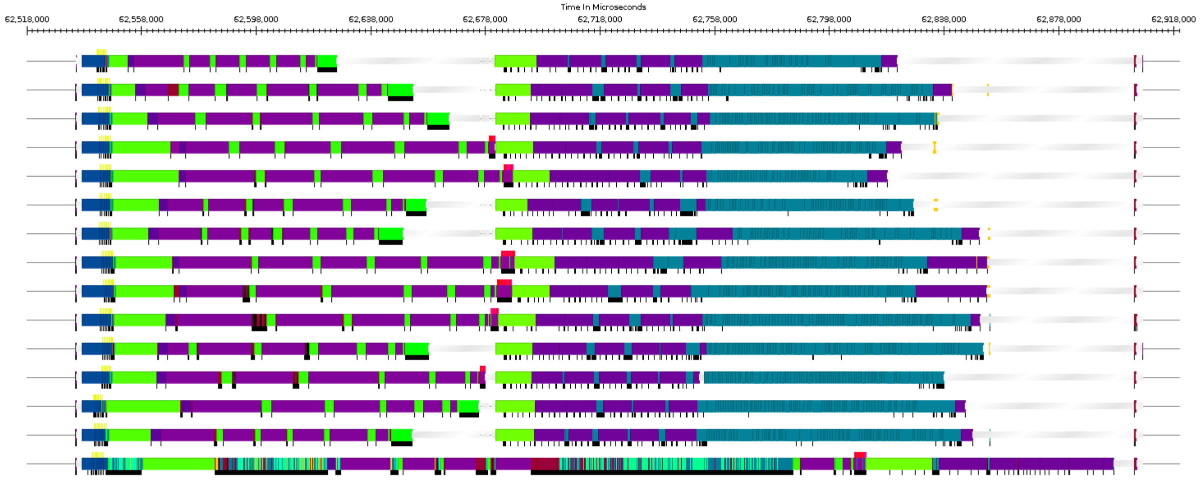

Communication-Load Imbalance

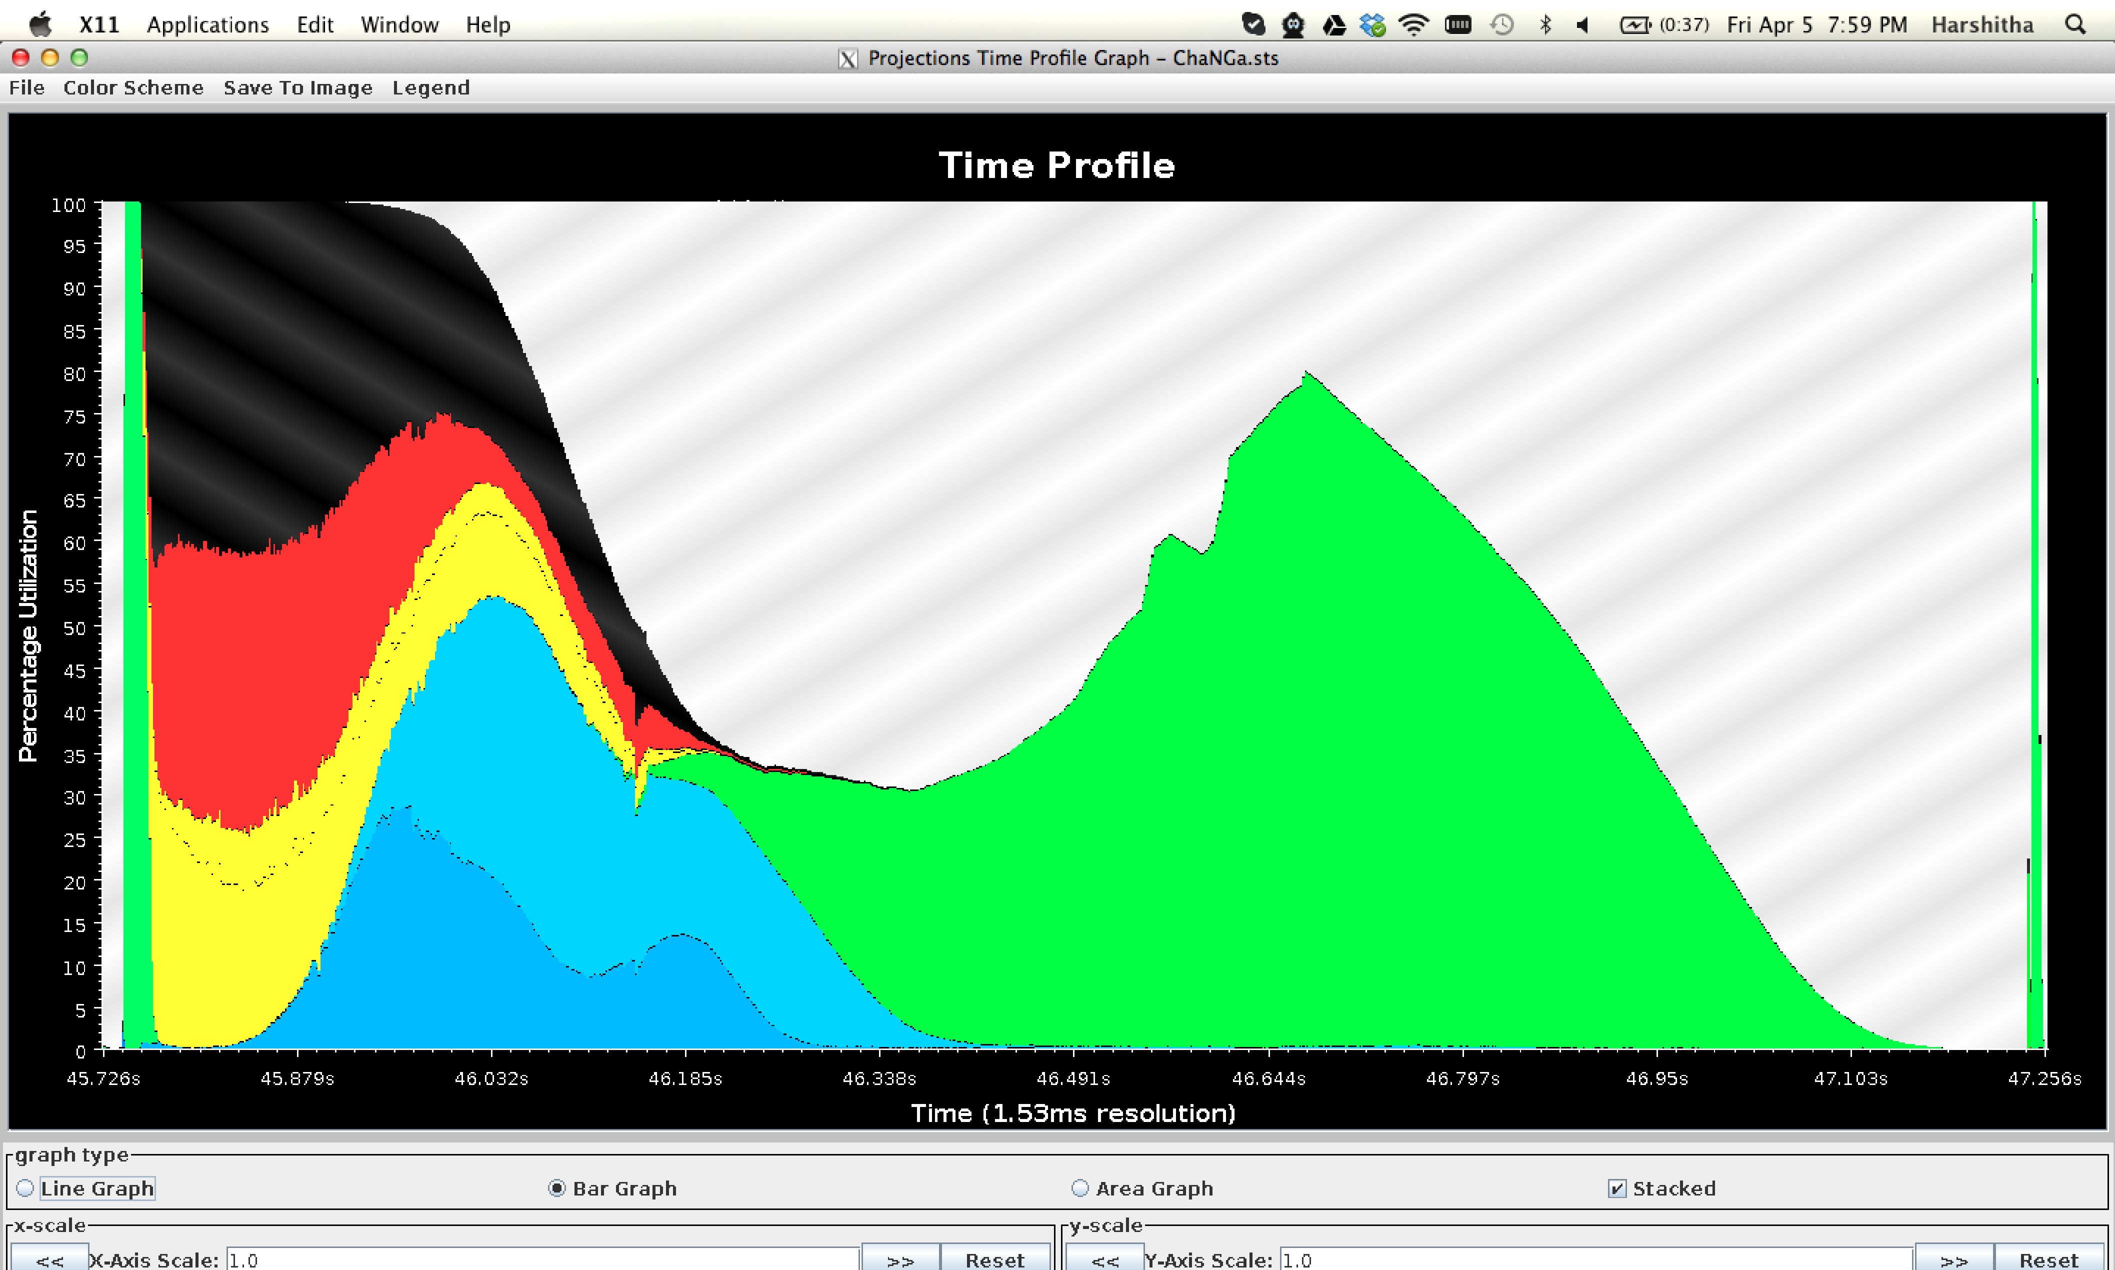

Replicating Data to Balance Load

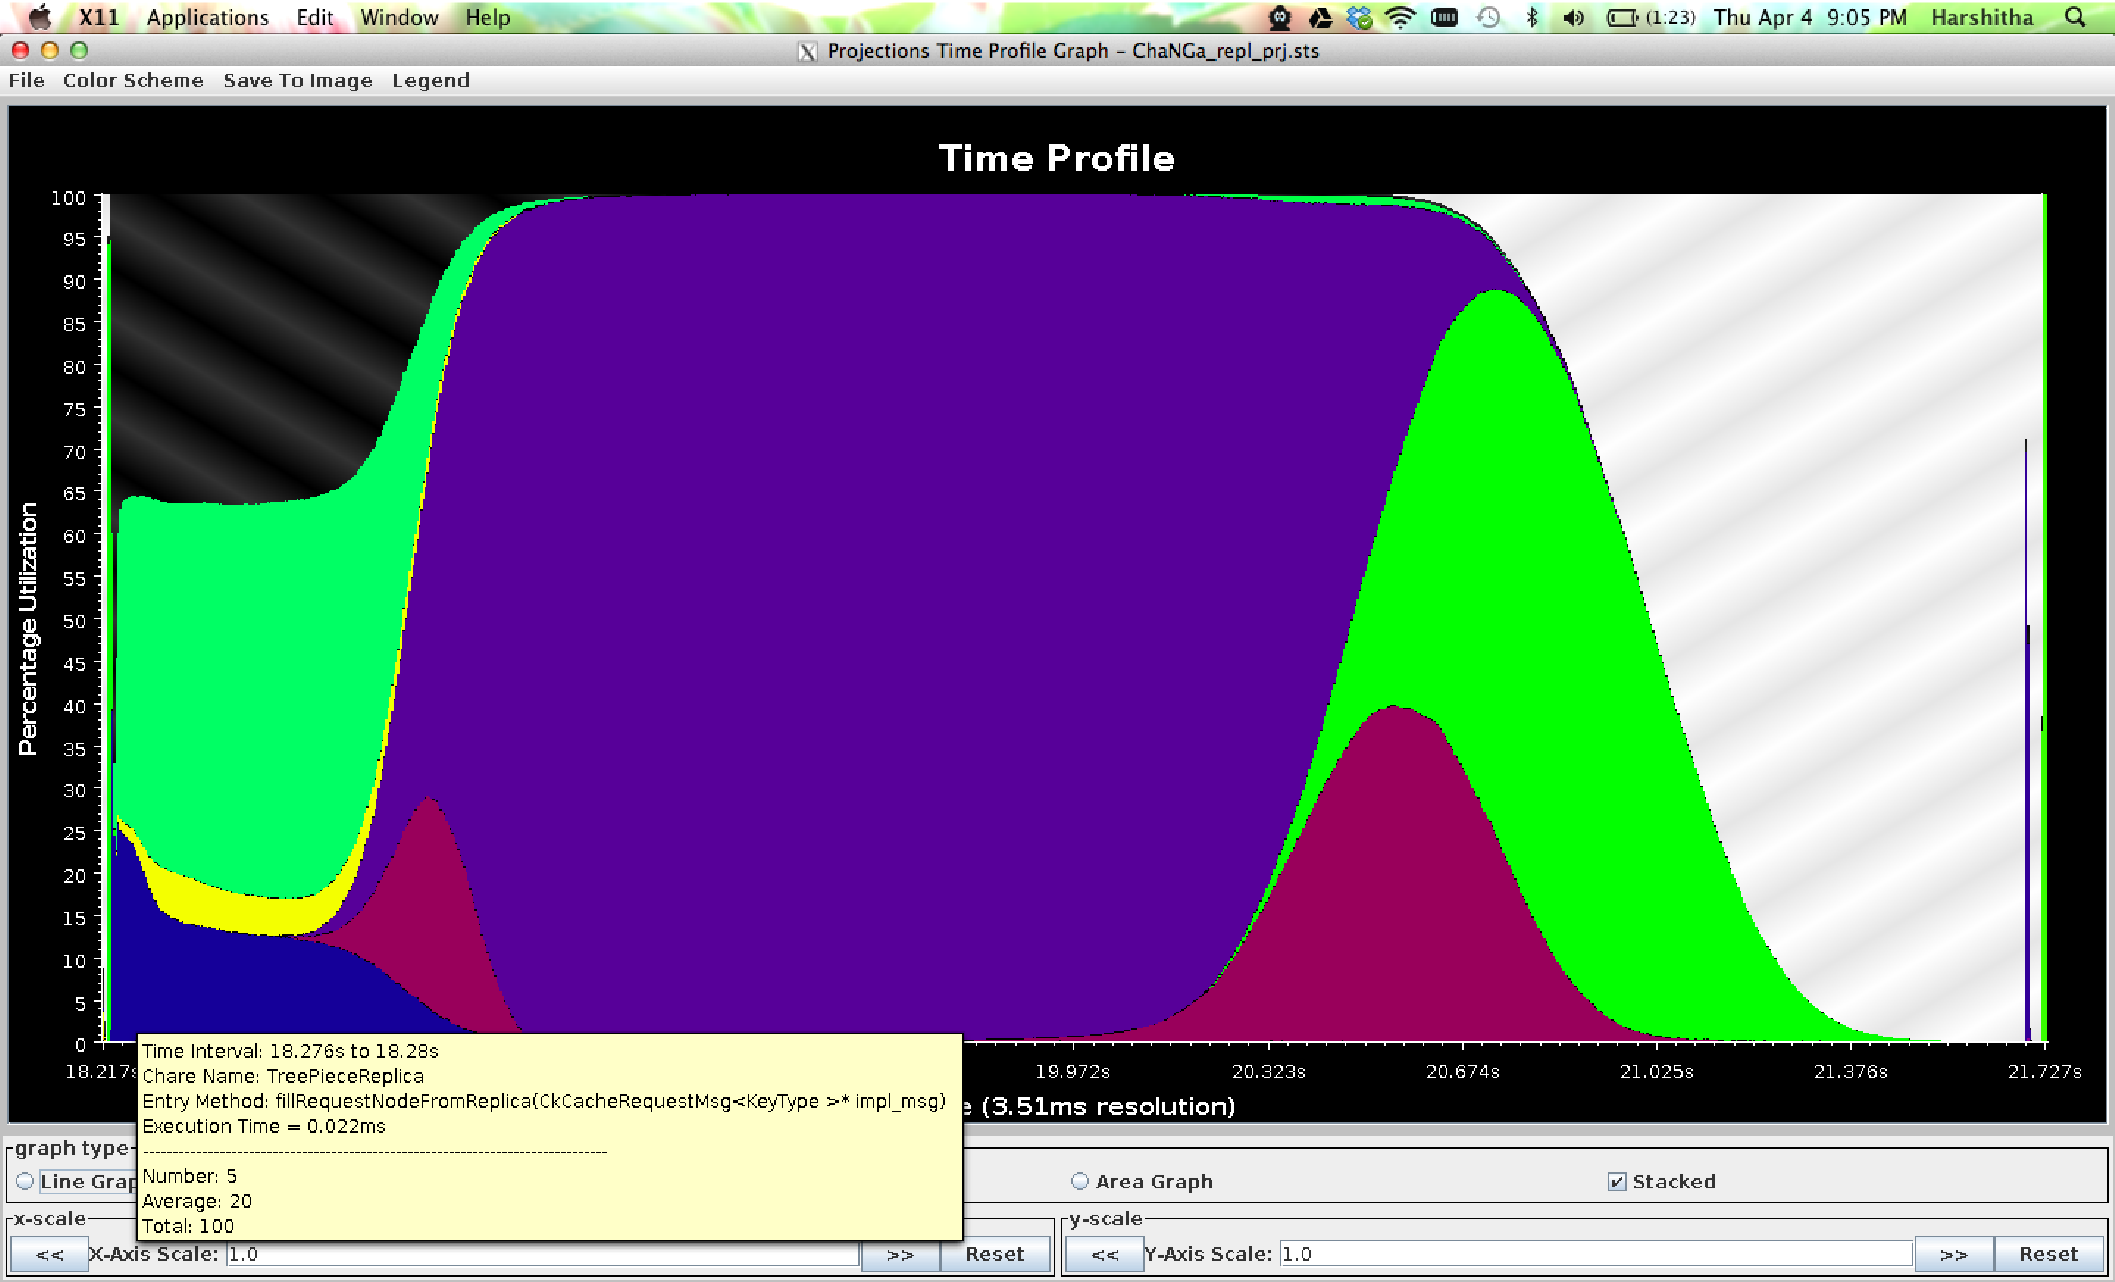

Multistep speedups

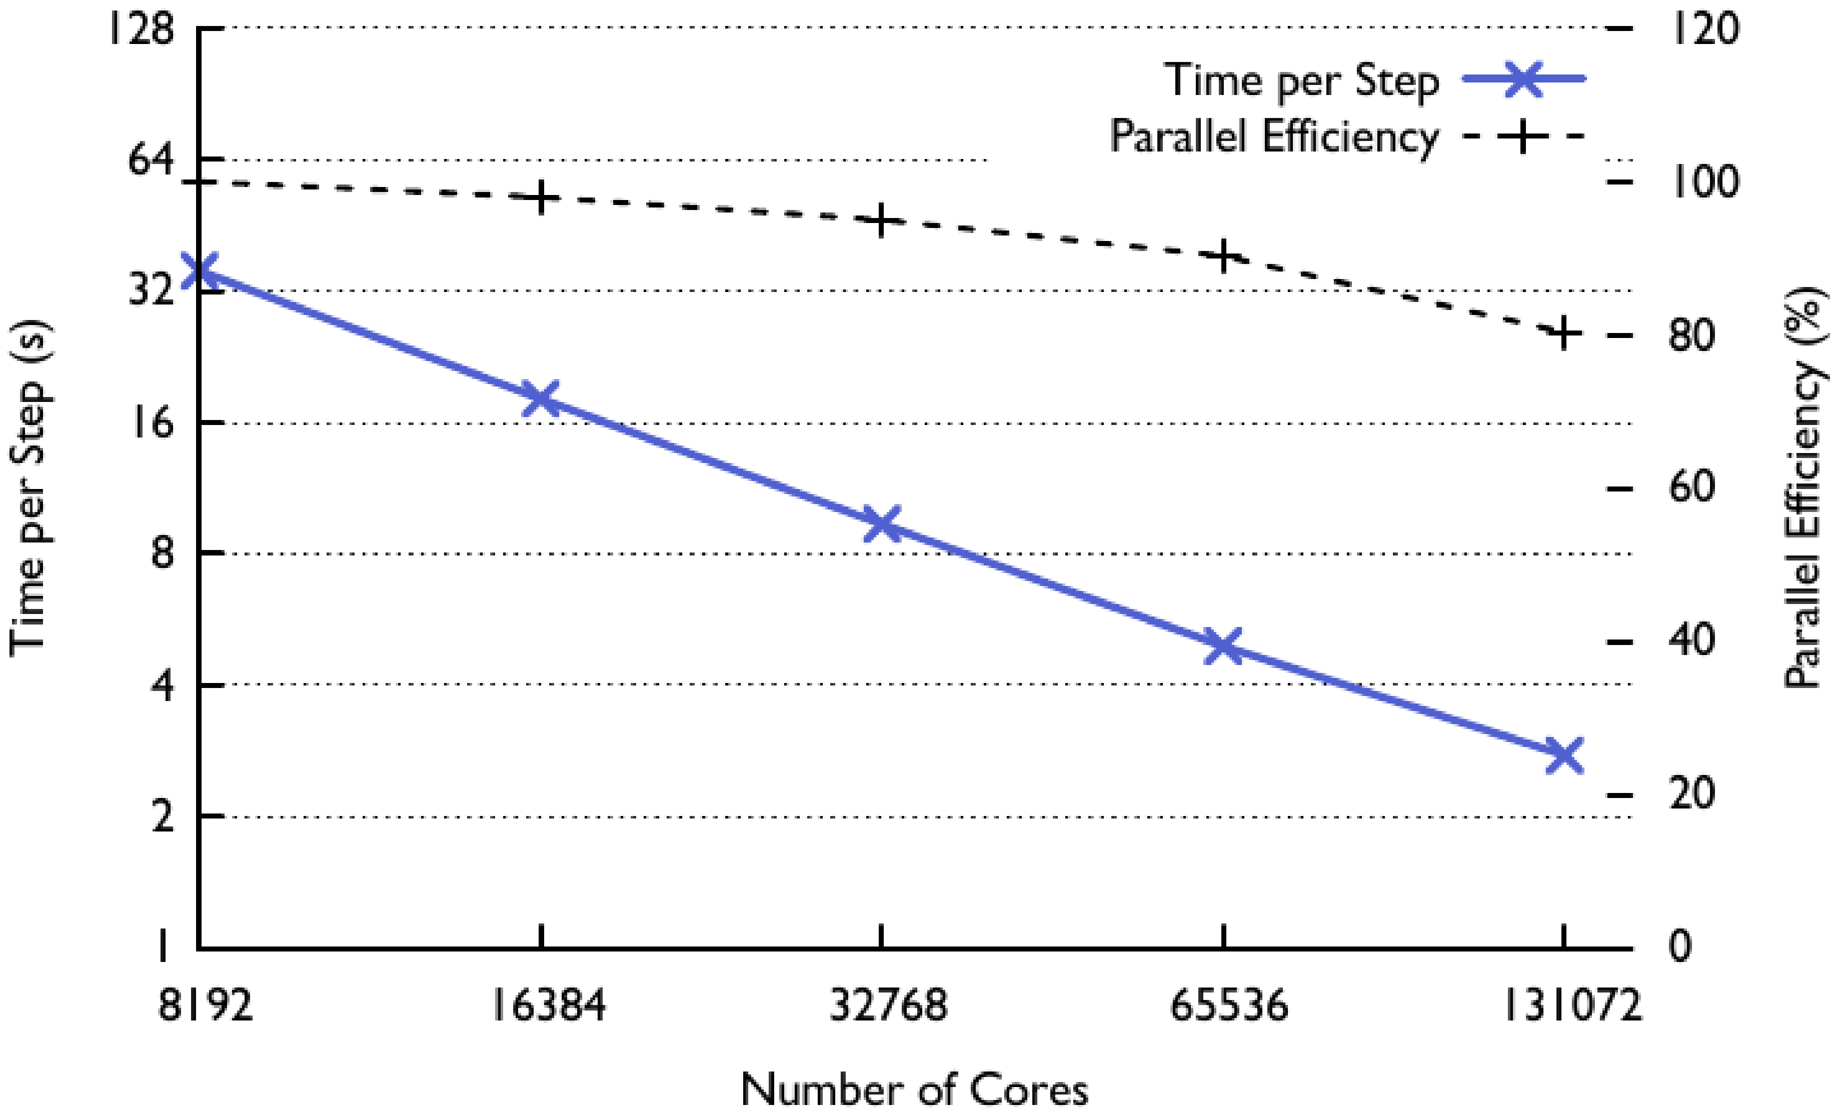

Multistep speedups

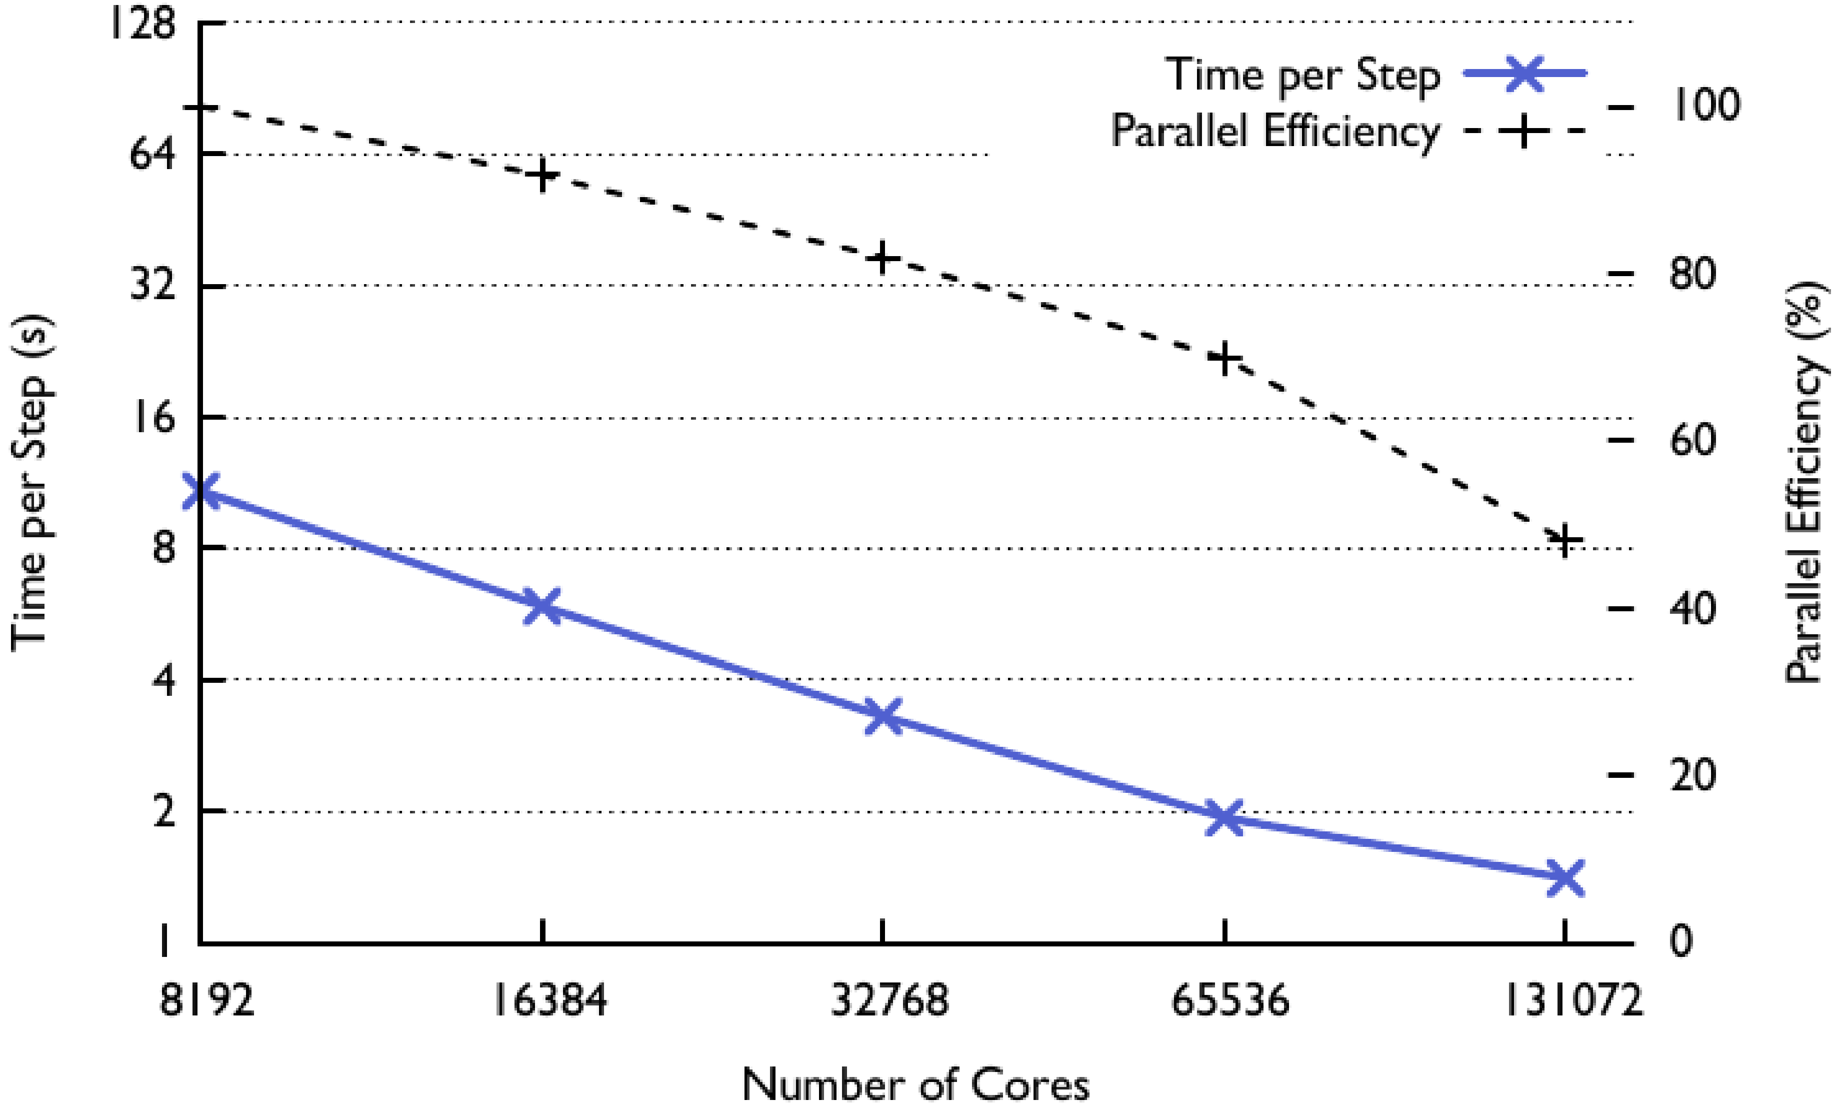

Clusters of Galaxies

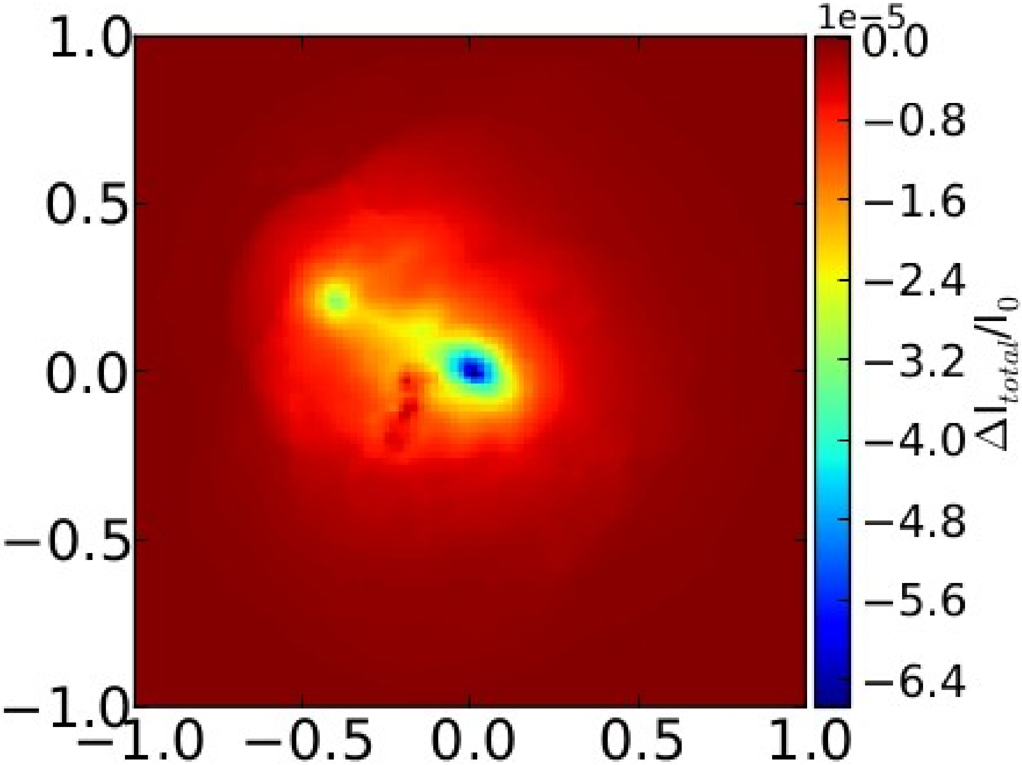

- Largest bound objects in the Universe
- Used to study evolution of Dark Energy
- Need 1 kpc resolution in 600 Mpc volume

John Ruan, et al 2013

Dwarf Galaxies and the Milky Way Disk

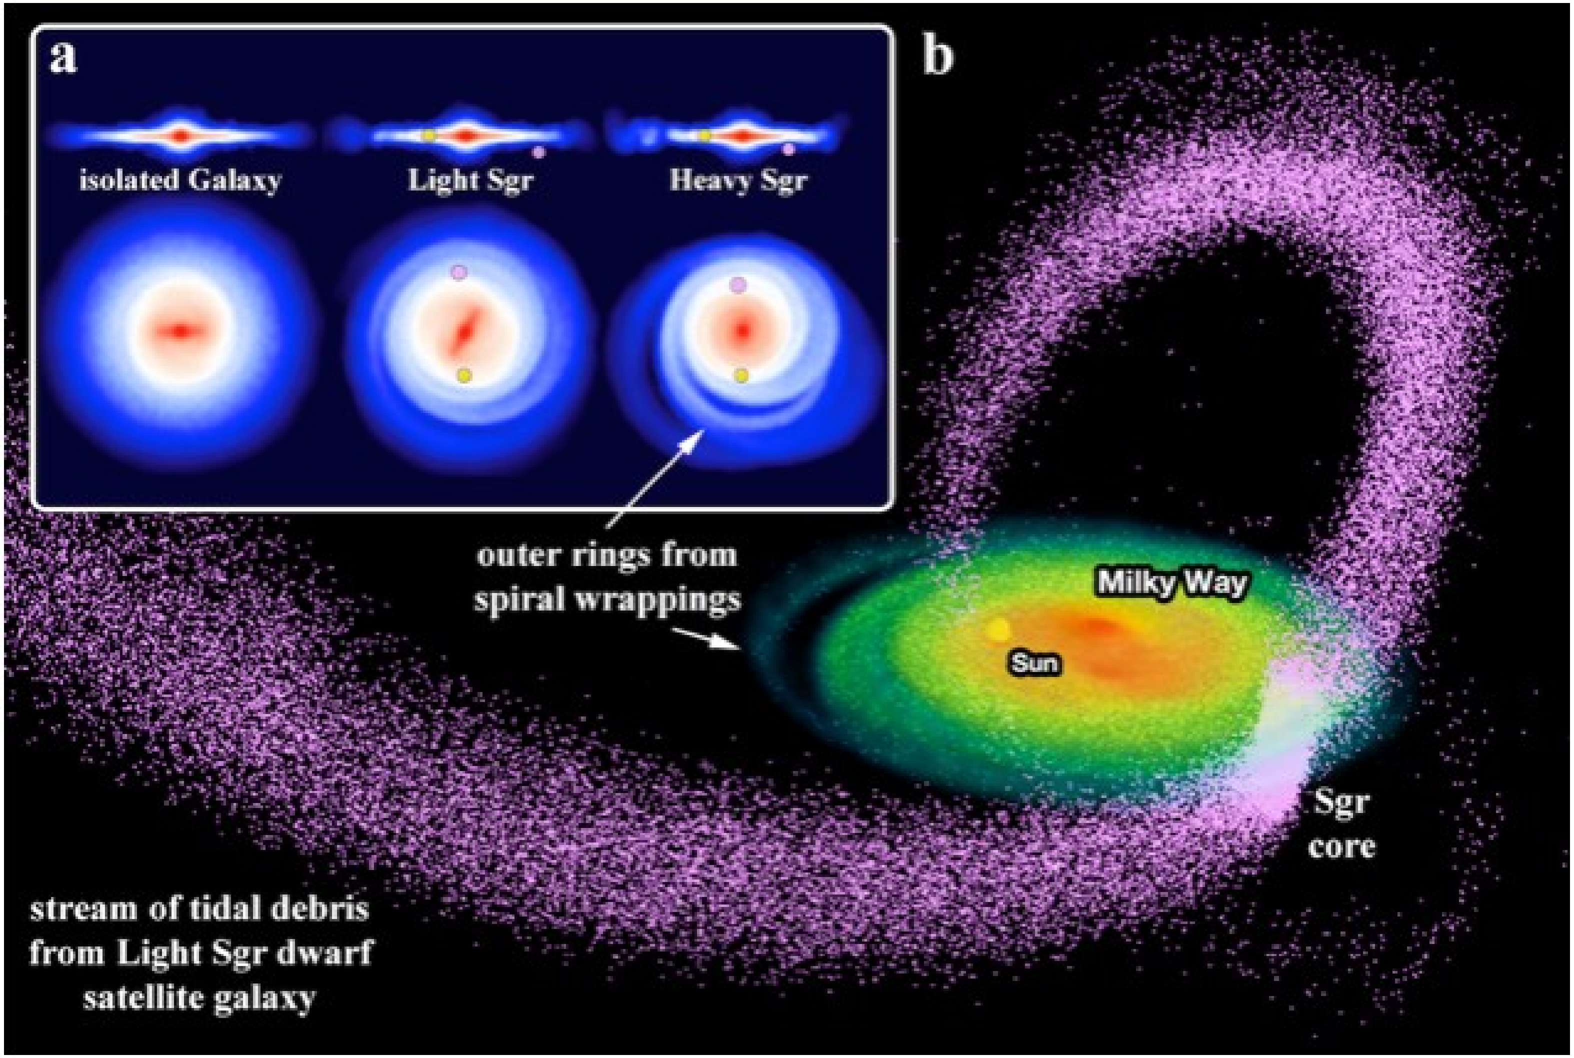

Purcell et al, Nature 2011

Future

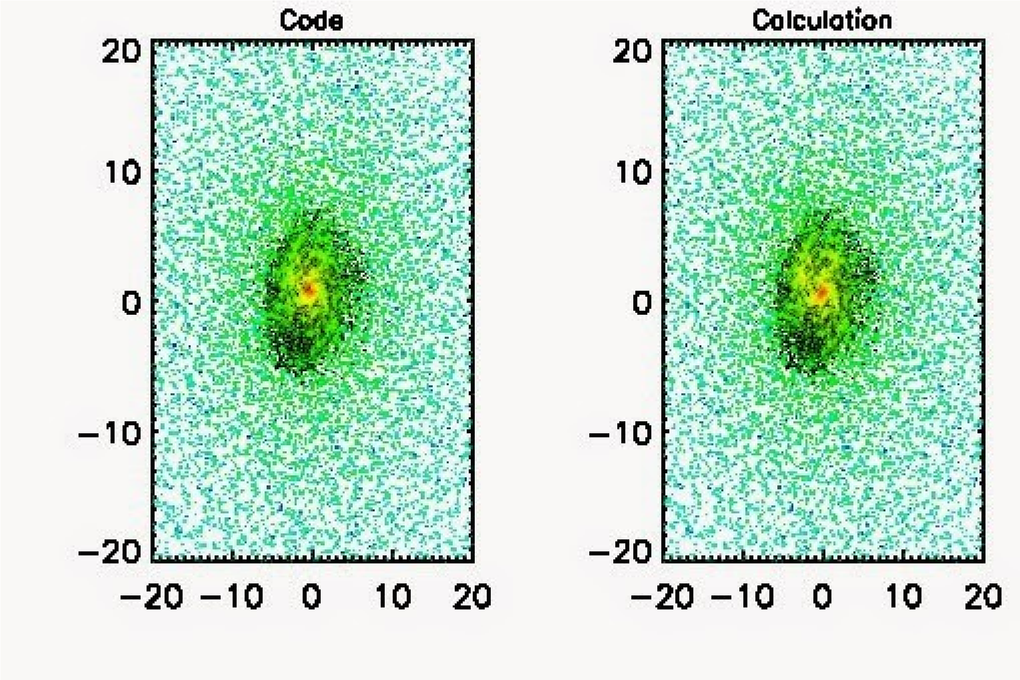

- More Physics
  - Massive Black Holes
  - Radiative transfer
  - Self-Interacting DM
  - Reuse of legacy code
- Better gravity algorithms
  - Fast Multipole Method
  - Heterogeneous machines
- Other Astrophysical problems
  - Planet formation/Planetary Rings

Galactic structure in the local Universe: What's needed

- 1 Million particles/galaxy for proper morphology/heavy element production
- 25 Mpc volume
- 800 M core-hours
- Necessary for:
  - Comparing with Hubble Space Telescope surveys of the local Universe
  - Interpreting HST images of high redshift galaxies

Large Scale Structure: What's needed

- 700 Megaparsec volume for "fair sample" of the Universe
- 18 trillion core-hours (~ exaflop year)
- Necessary for:
  - Interpreting future surveys (LSST)
  - Relating Cosmic Microwave Background to galaxy surveys

Summary

- Cosmological simulations provide a challenges to parallel implementations
  - Non-local data dependencies
  - Hierarchical in space and time
- ChaNGa has been successful in addressing this challenges using Charm++ features
  - Computation/Communication overlap
  - Message priorities
  - New load balancers