# Evolution of the Small Galaxy Population

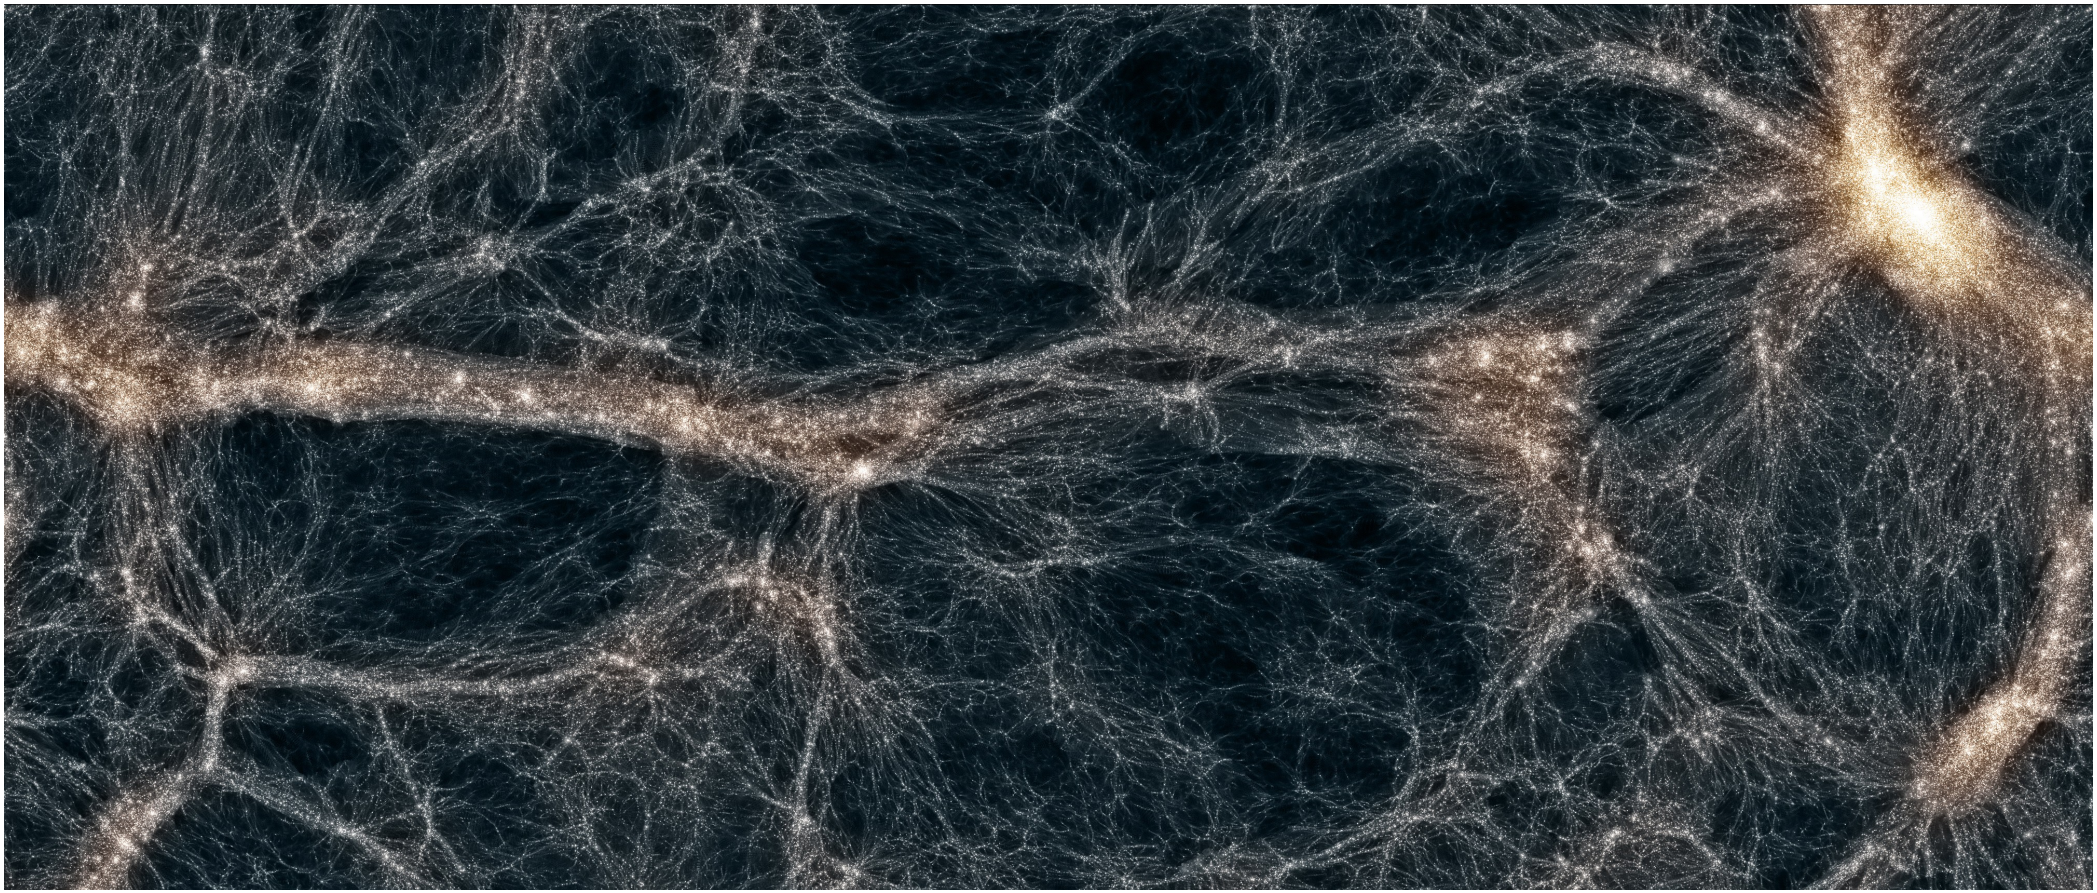

Thomas Quinn
University of Washington
NSF PRAC Award 1144357

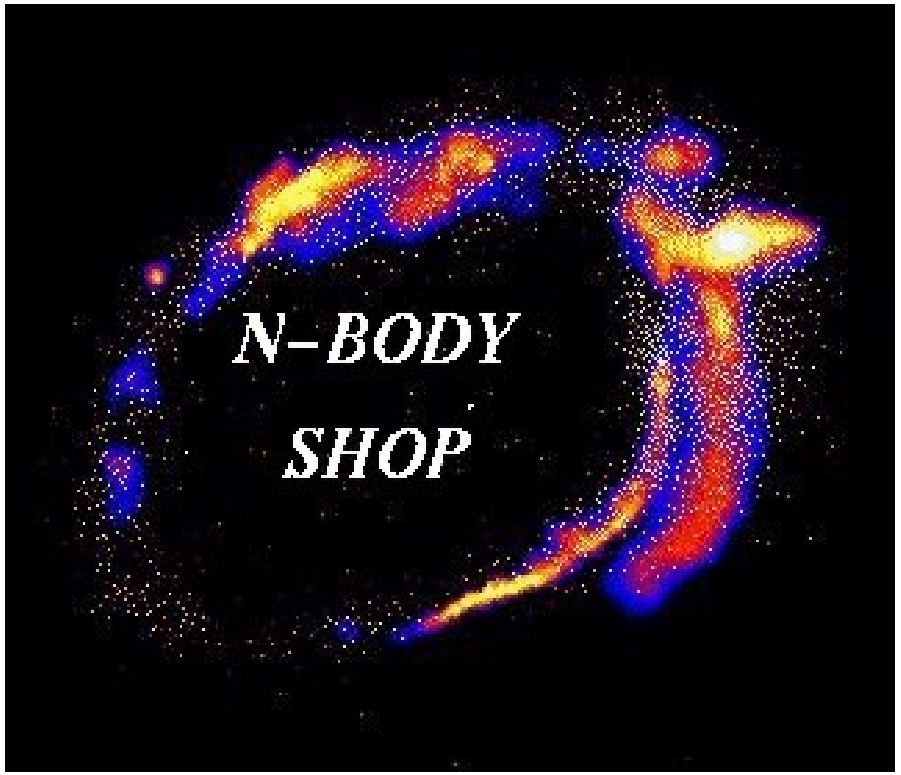


PPL PARALLEL PROGRAMMING LAB
DEPT OF COMPUTER SCIENCE, UNIVERSITY OF ILLINOIS

Fabio Governato
Lauren Anderson
Michael Tremmel
Charlotte Christensen
Joachim Stadel
James Wadsley
Greg Stinson

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

# Outline

- Scientific background (Why it matters)
- Need for high resolution (Key Challenges)
- Project goals (Why Blue Watters)
- MPI vs HLL, Charm++  and ChaNGa (Key Challenges)
  - Importance of BW group
- Preliminary results (Accomplishments)
- Data analysis challenges (Shared Data)

# Galaxy formation: can this...

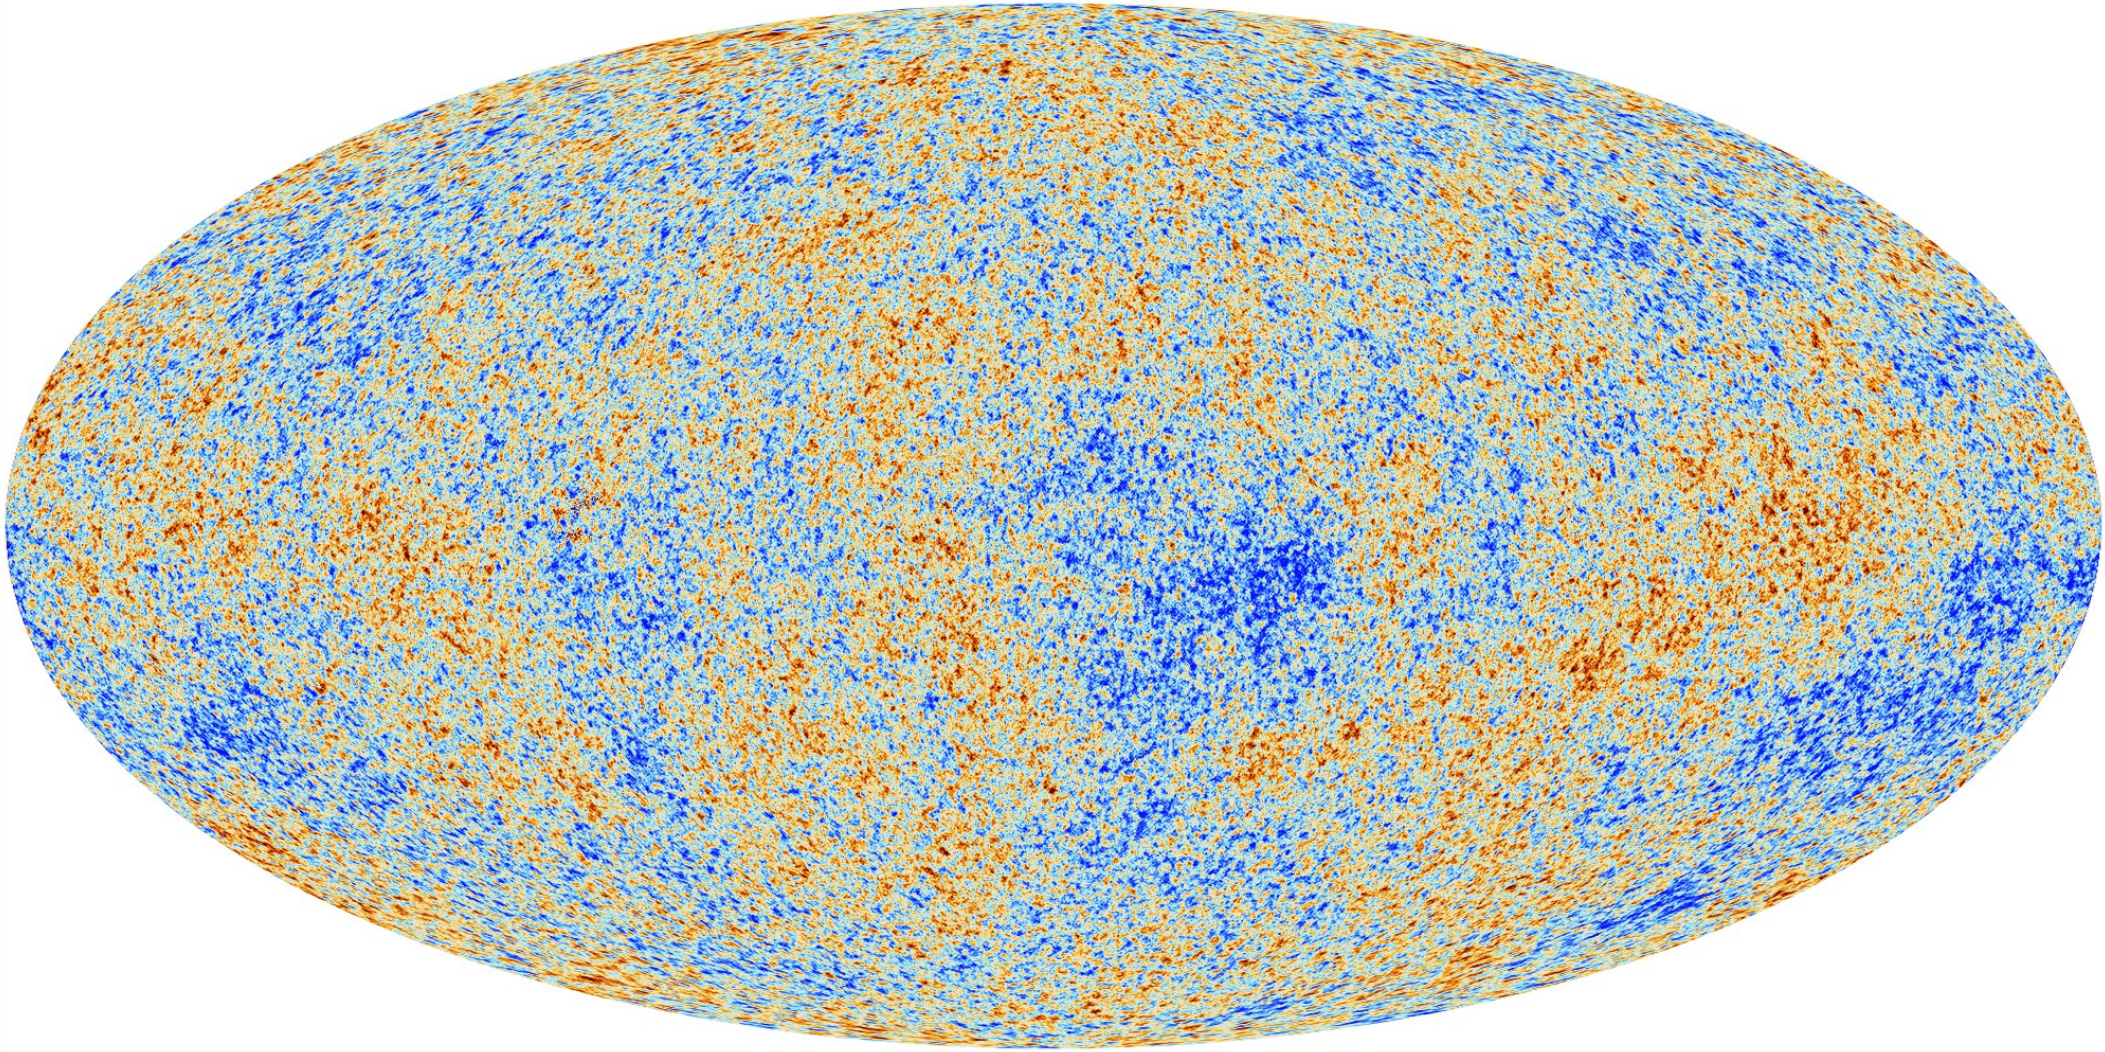

Image courtesy ESA/Planck

... turn into this?

# Modeling Star Formation: it's hard

- Gravitational Instabilities
- Magnetic Fields
- Radiative Transfer
- Molecular/Dust Chemistry
- Driven at large scales: differential rotation
- Driven at small scales: Supernovea and Stellar Winds
- Scales unresolvable in cosmological simulations

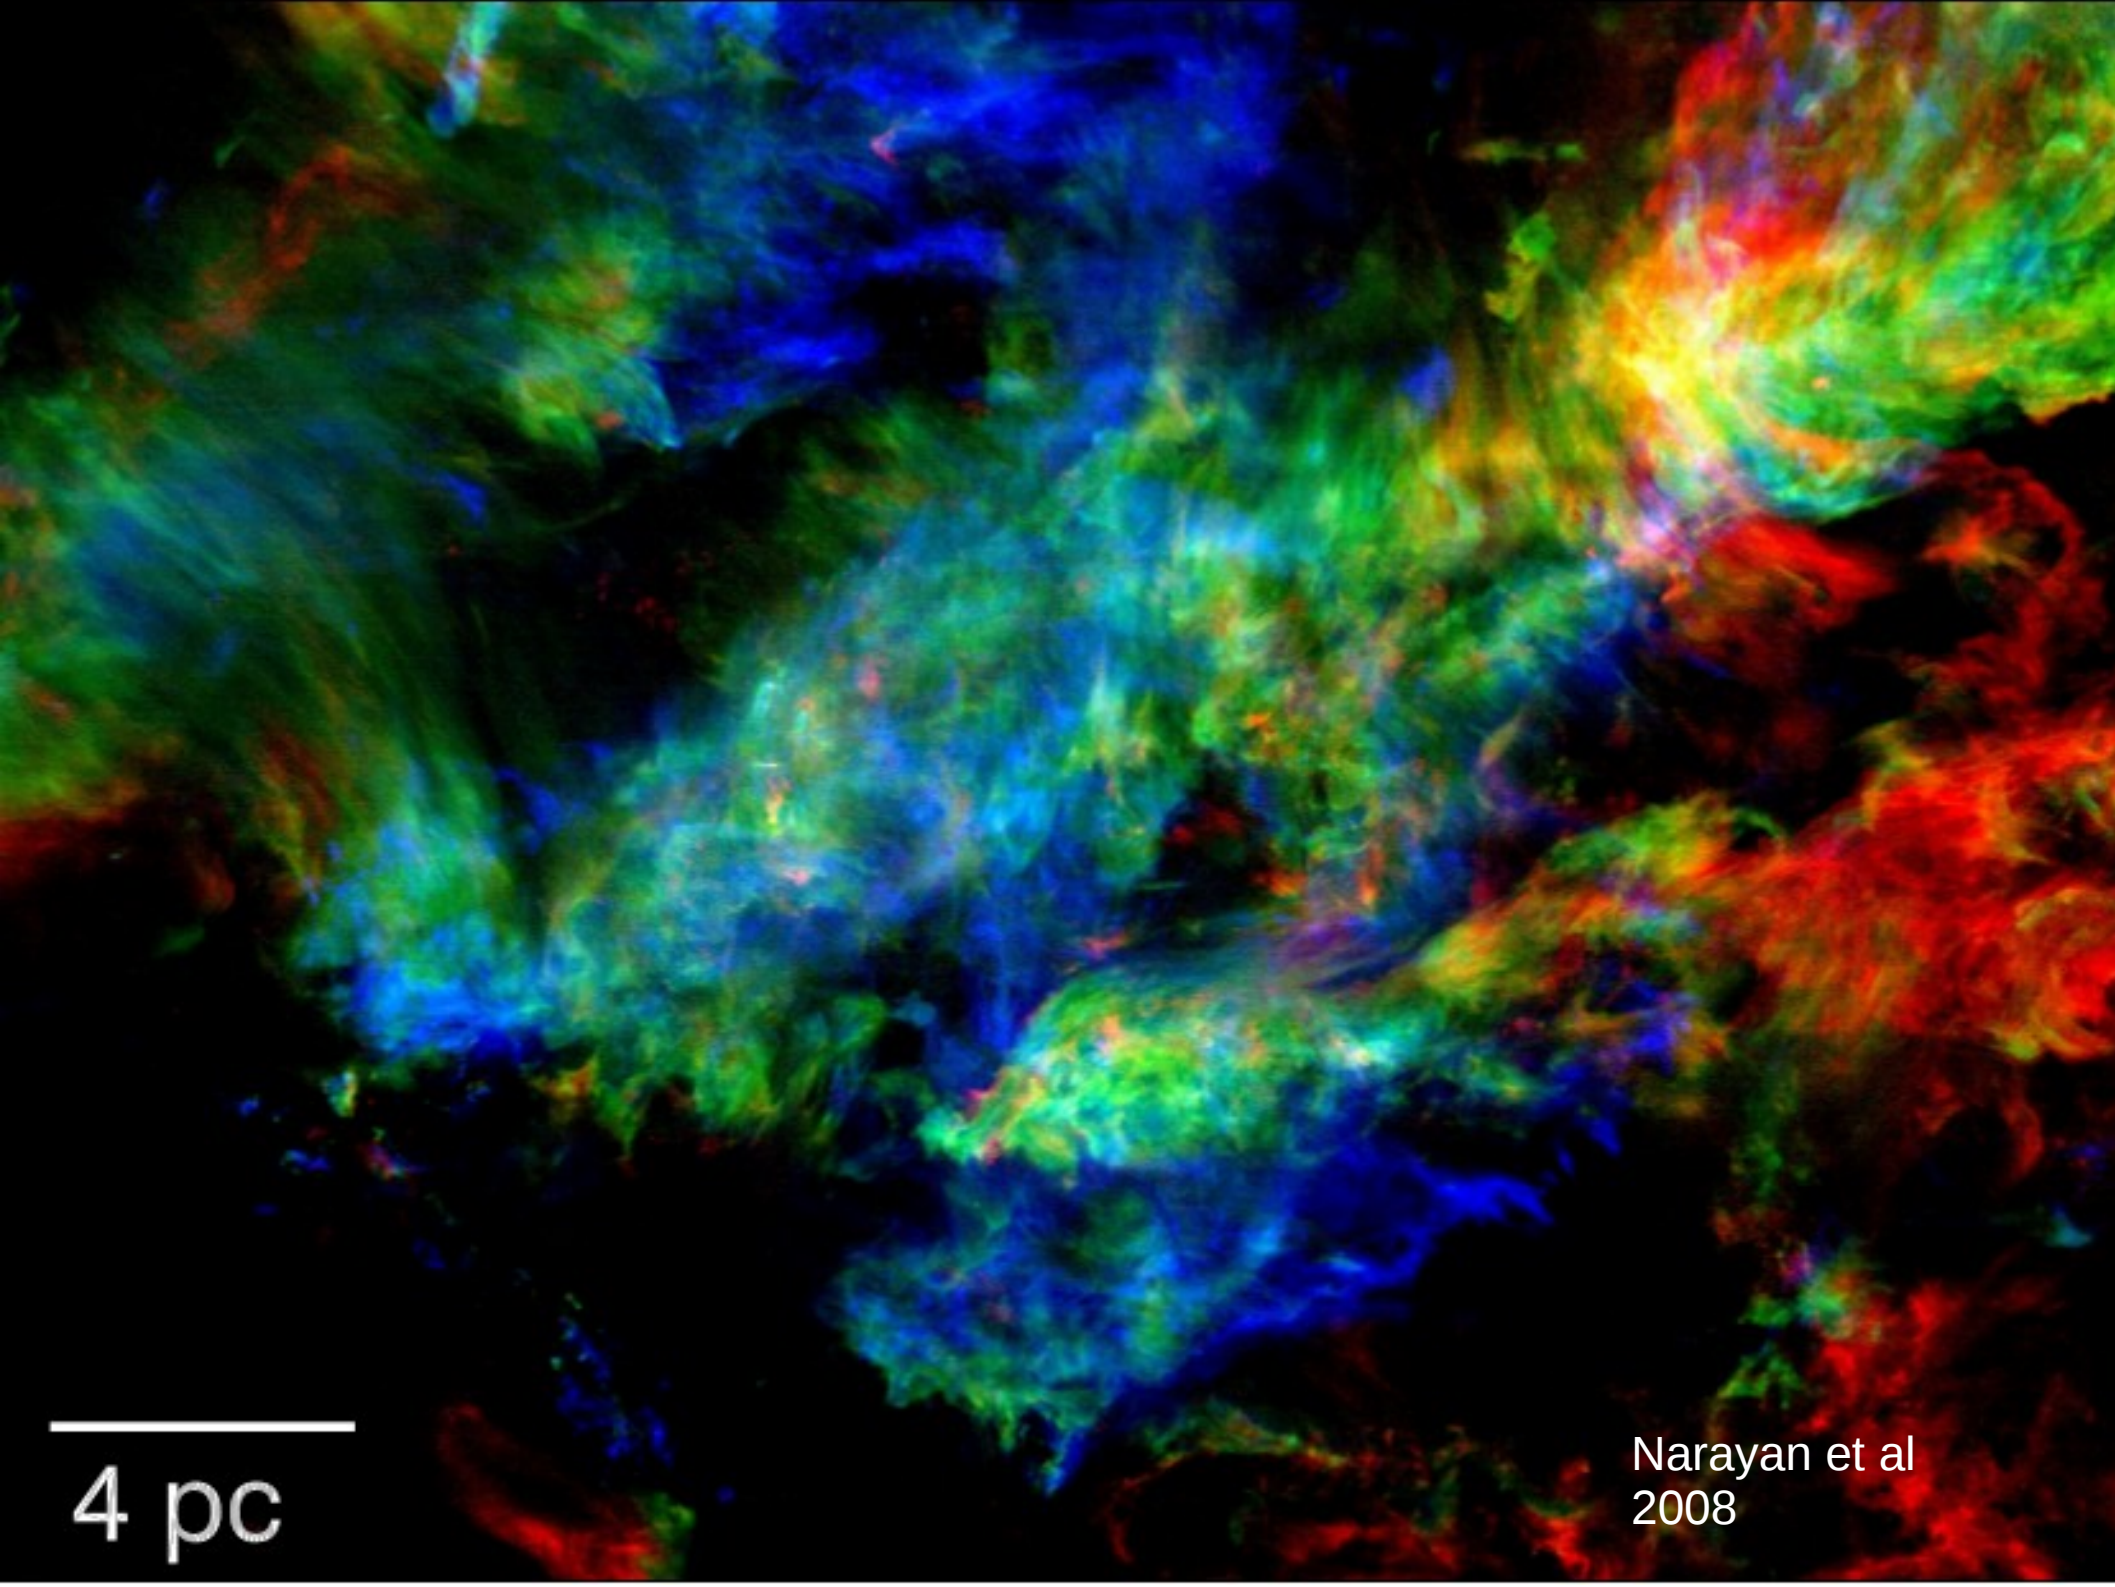
4 pc
Narayan et al
2008

# Resolution and Subgrid Models

- Maximize Simulation Resolution
  - Capture tidal torques/accretion history (20+ Mpc)
  - Adapt resolution to galaxy (sub-Kpc)
- Capture Star Formation in a sub-grid model
  - Stars form in high density environments
  - Supernovea/stellar winds/radiation regulate star formation
  - Mitigate issues with poor resolution (overcooling)
  - Tune to match present day stellar populations

# Blue Waters: High Redshift Galaxies

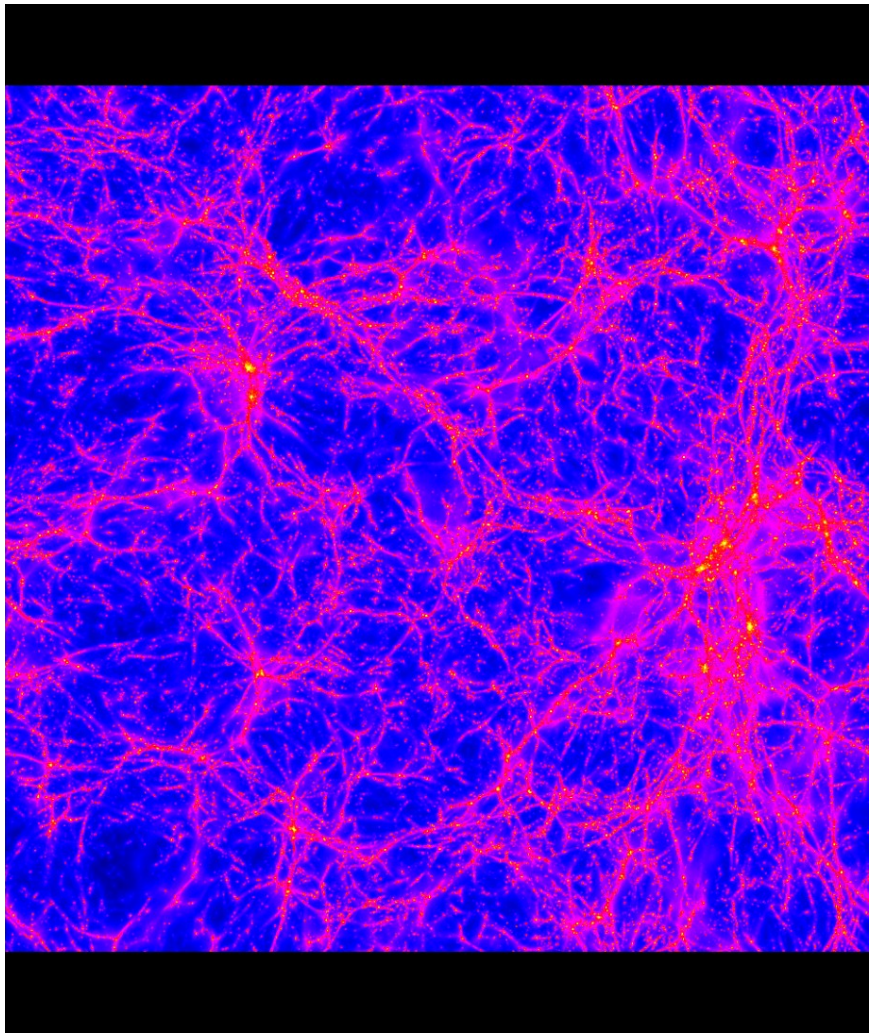

- 25 Mpc Volume
- Few million particles/galaxy
- Goals:
  - Models to compare with HST Frontier fields
  - Physical properties of high z galaxies and connection to the present day

HPC is dying, and MPI is killing it - Mozilla Firefox

13th Annual ... | 21901_centauru... | HPC is dying,... | SC_Tangram ... | http://lo...hboard/8 | How do I tak...

www.dursi.ca/hpc-is-dying-and-mpi-is-killing-it/

ubuntu screenshot

HOME ABOUT CONTACT

# HPC is dying, and MPI is killing it

Jonathan Dursi

hpc

Jonathan Dursi

Toronto

# Charm++

- C++-based parallel runtime system
  - Composed of a set of globally-visible parallel objects that interact
  - The objects interact by asynchronously invoking methods on each other
- Charm++ runtime
  - Manages the parallel objects and (re)maps them to processes
  - Provides scheduling, load balancing, and a host of other features, requiring little user intervention


# Charm Nbody GrAvity solver

- Massively parallel SPH
- SNe feedback creating realistic outflows
- SF linked to shielded gas
- SMBHs
- Optimized SF parameters

Menon+ 2014, Governato+ 2014

# Scaling to .5M cores

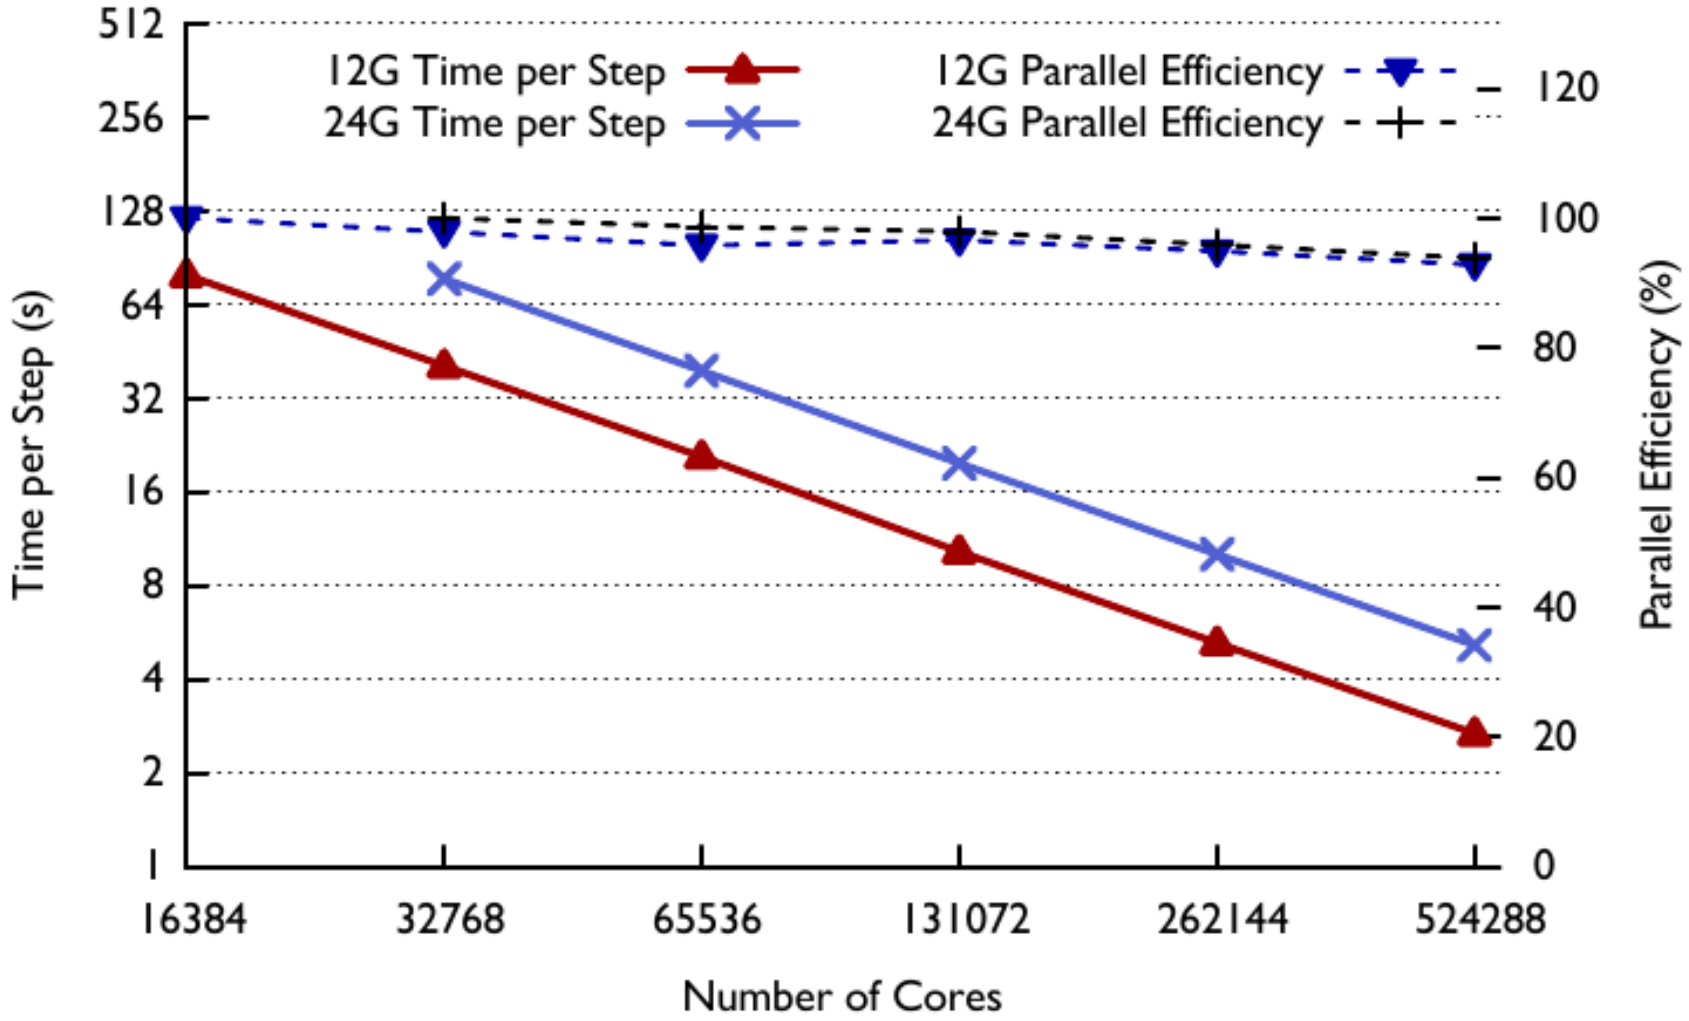

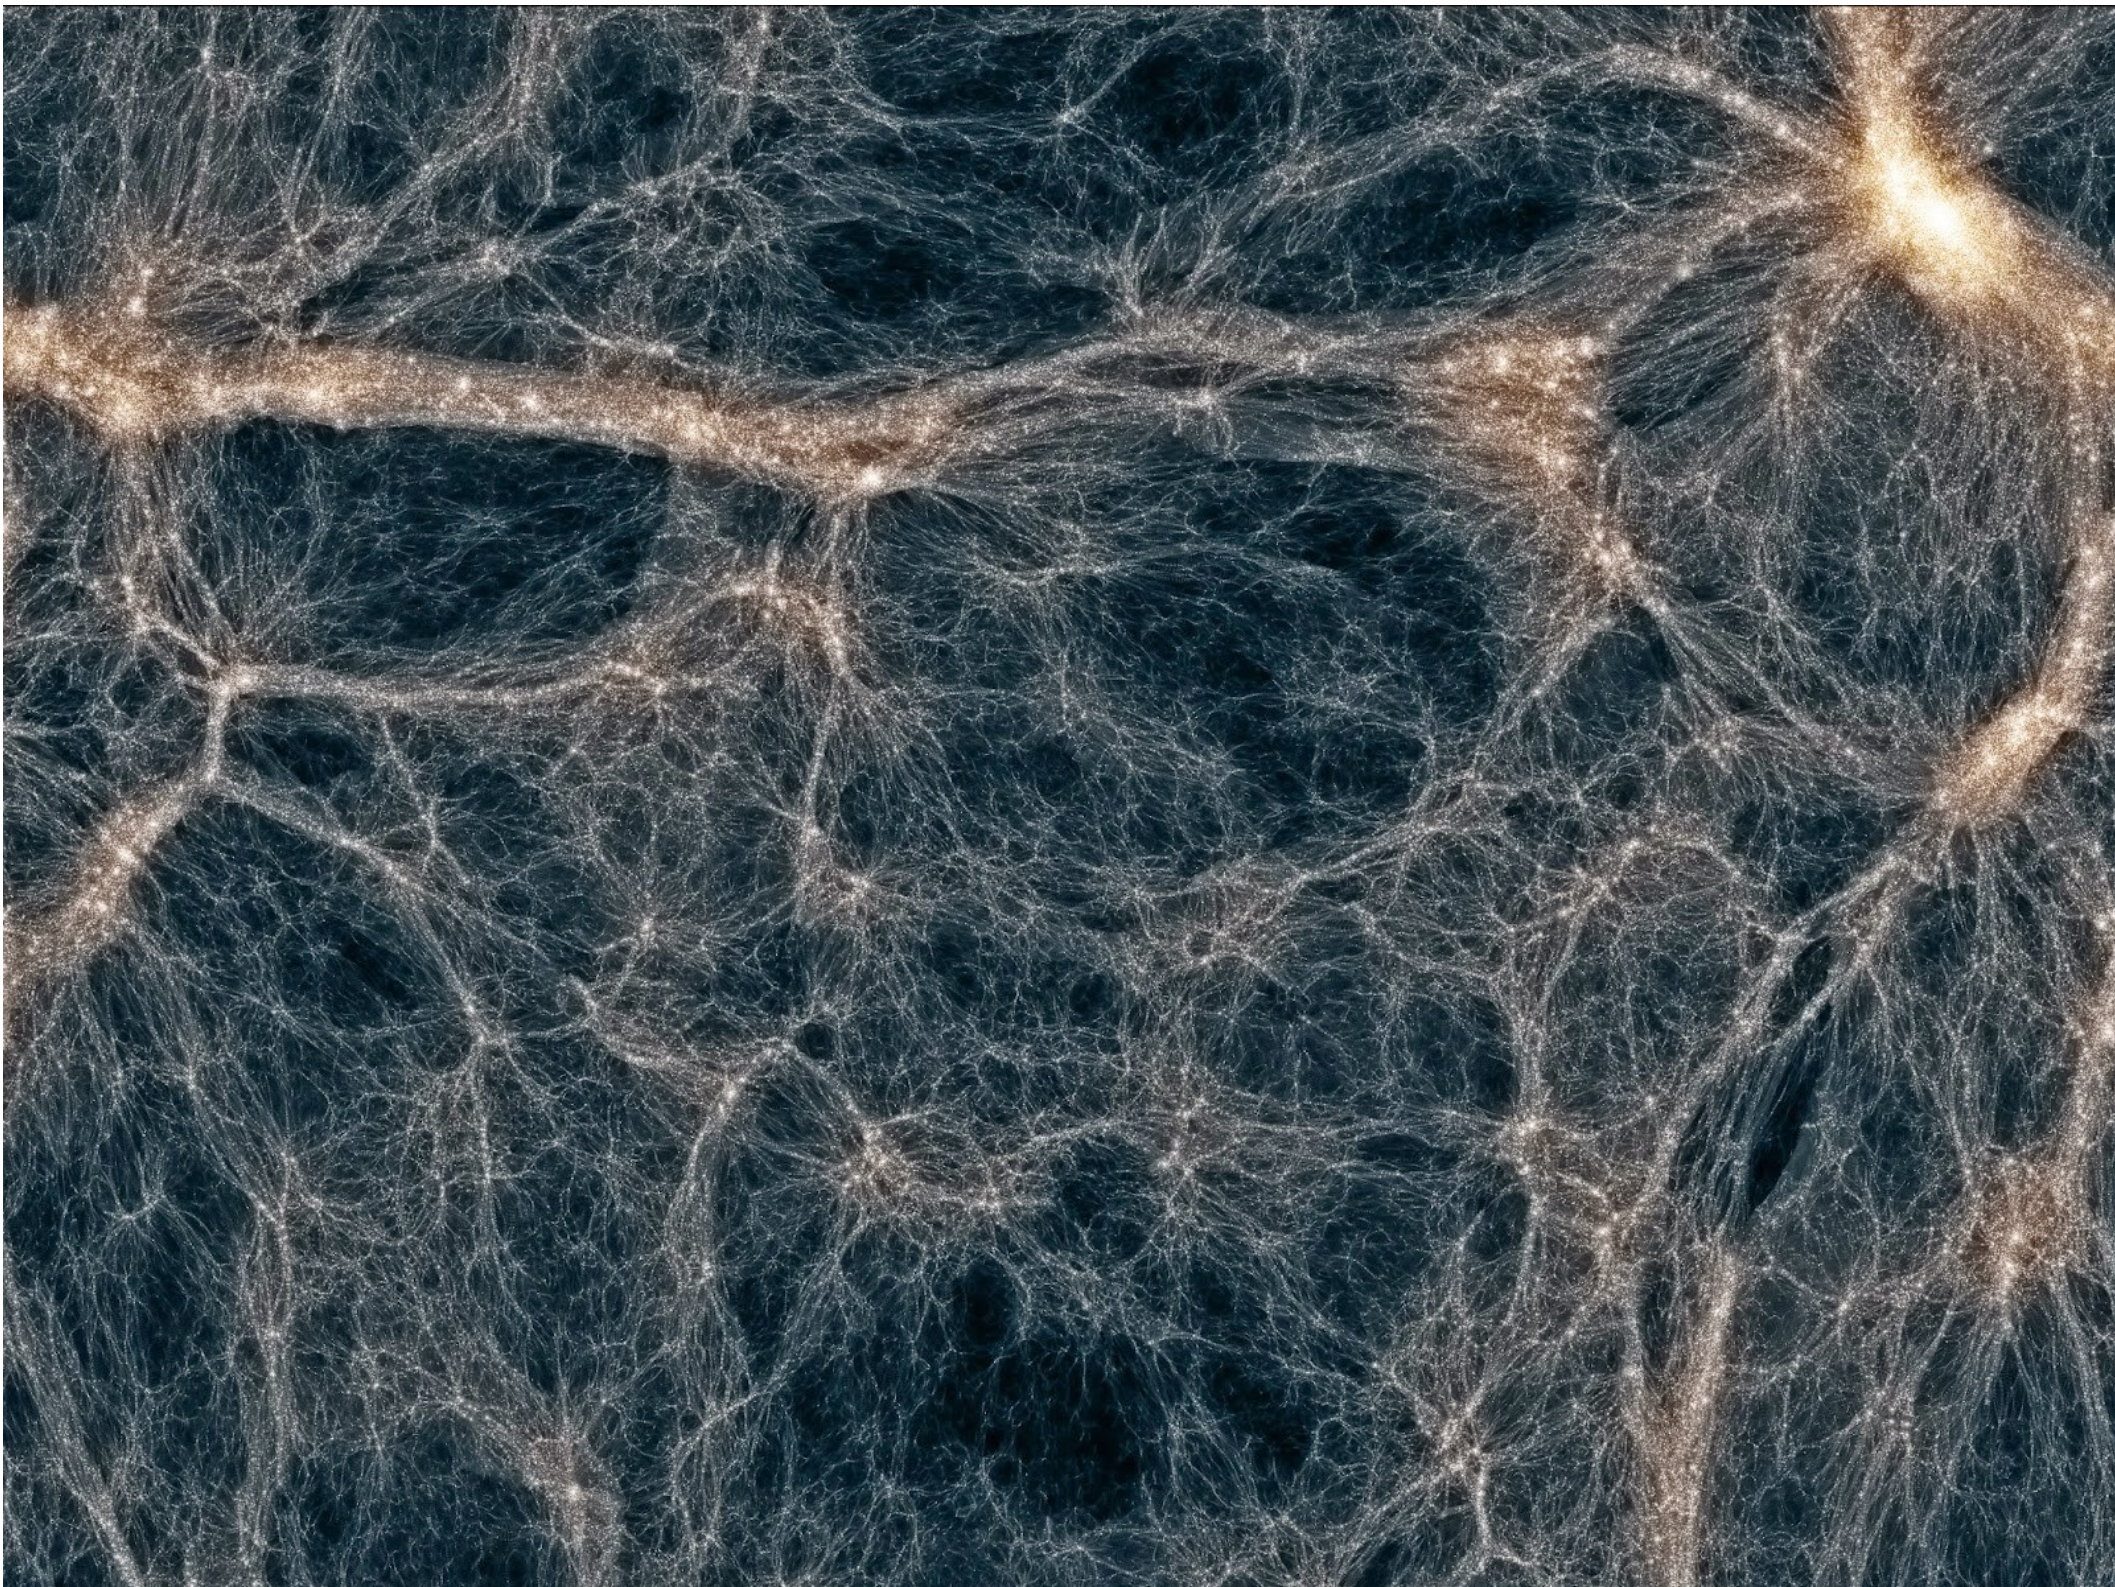

# Clustered/Multistepping Challenges

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

# Load Variance

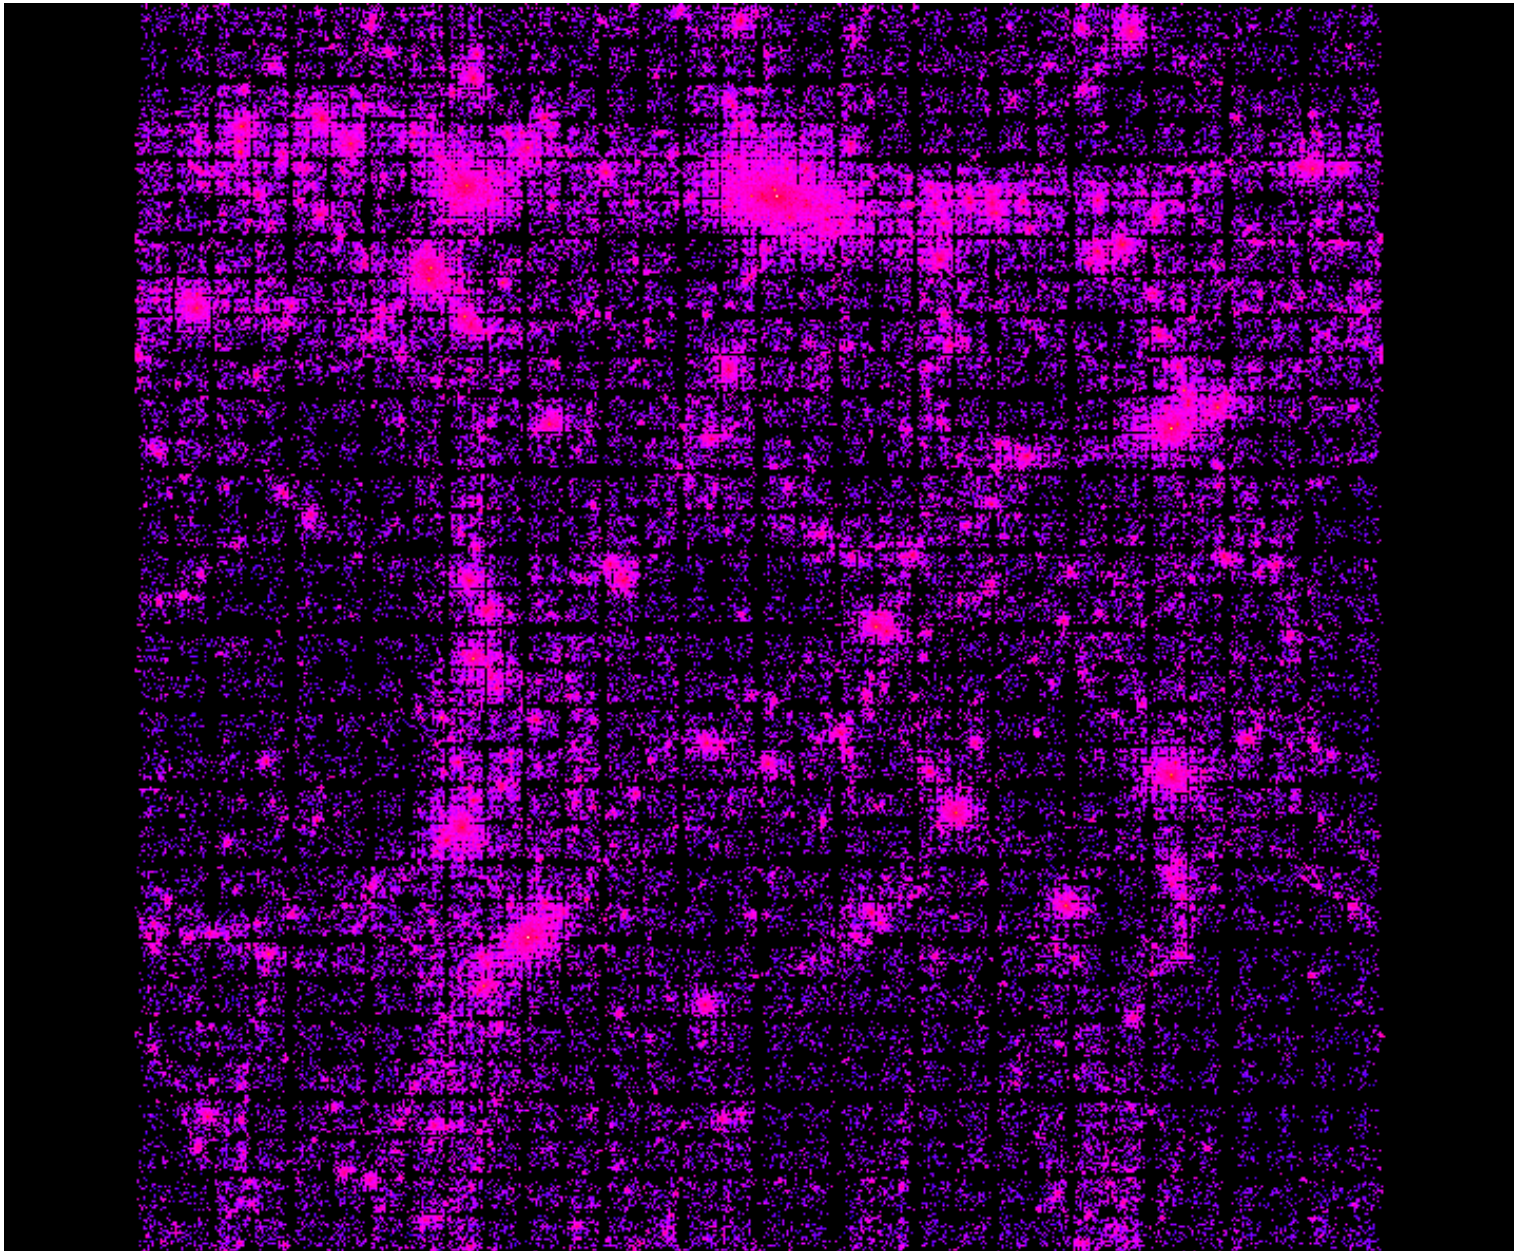

# ORB Load Balancing

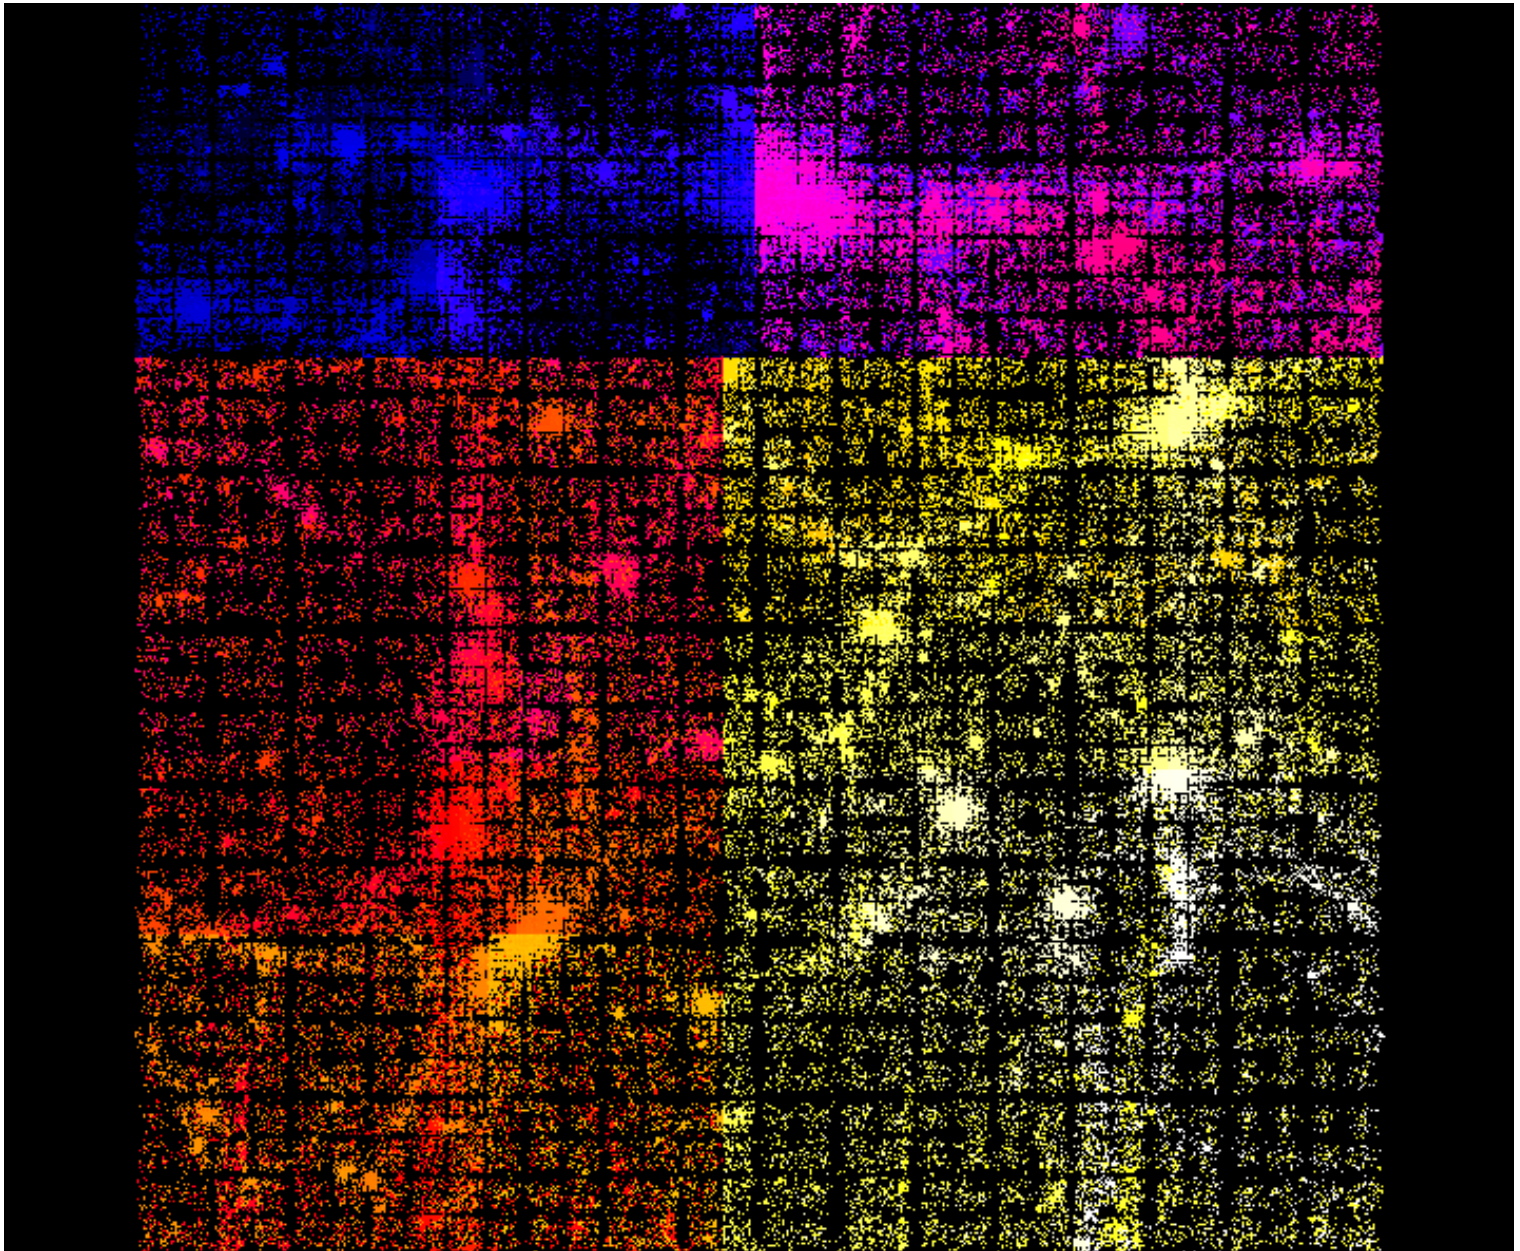

# Multistep speedups for 2 billion clustered particles

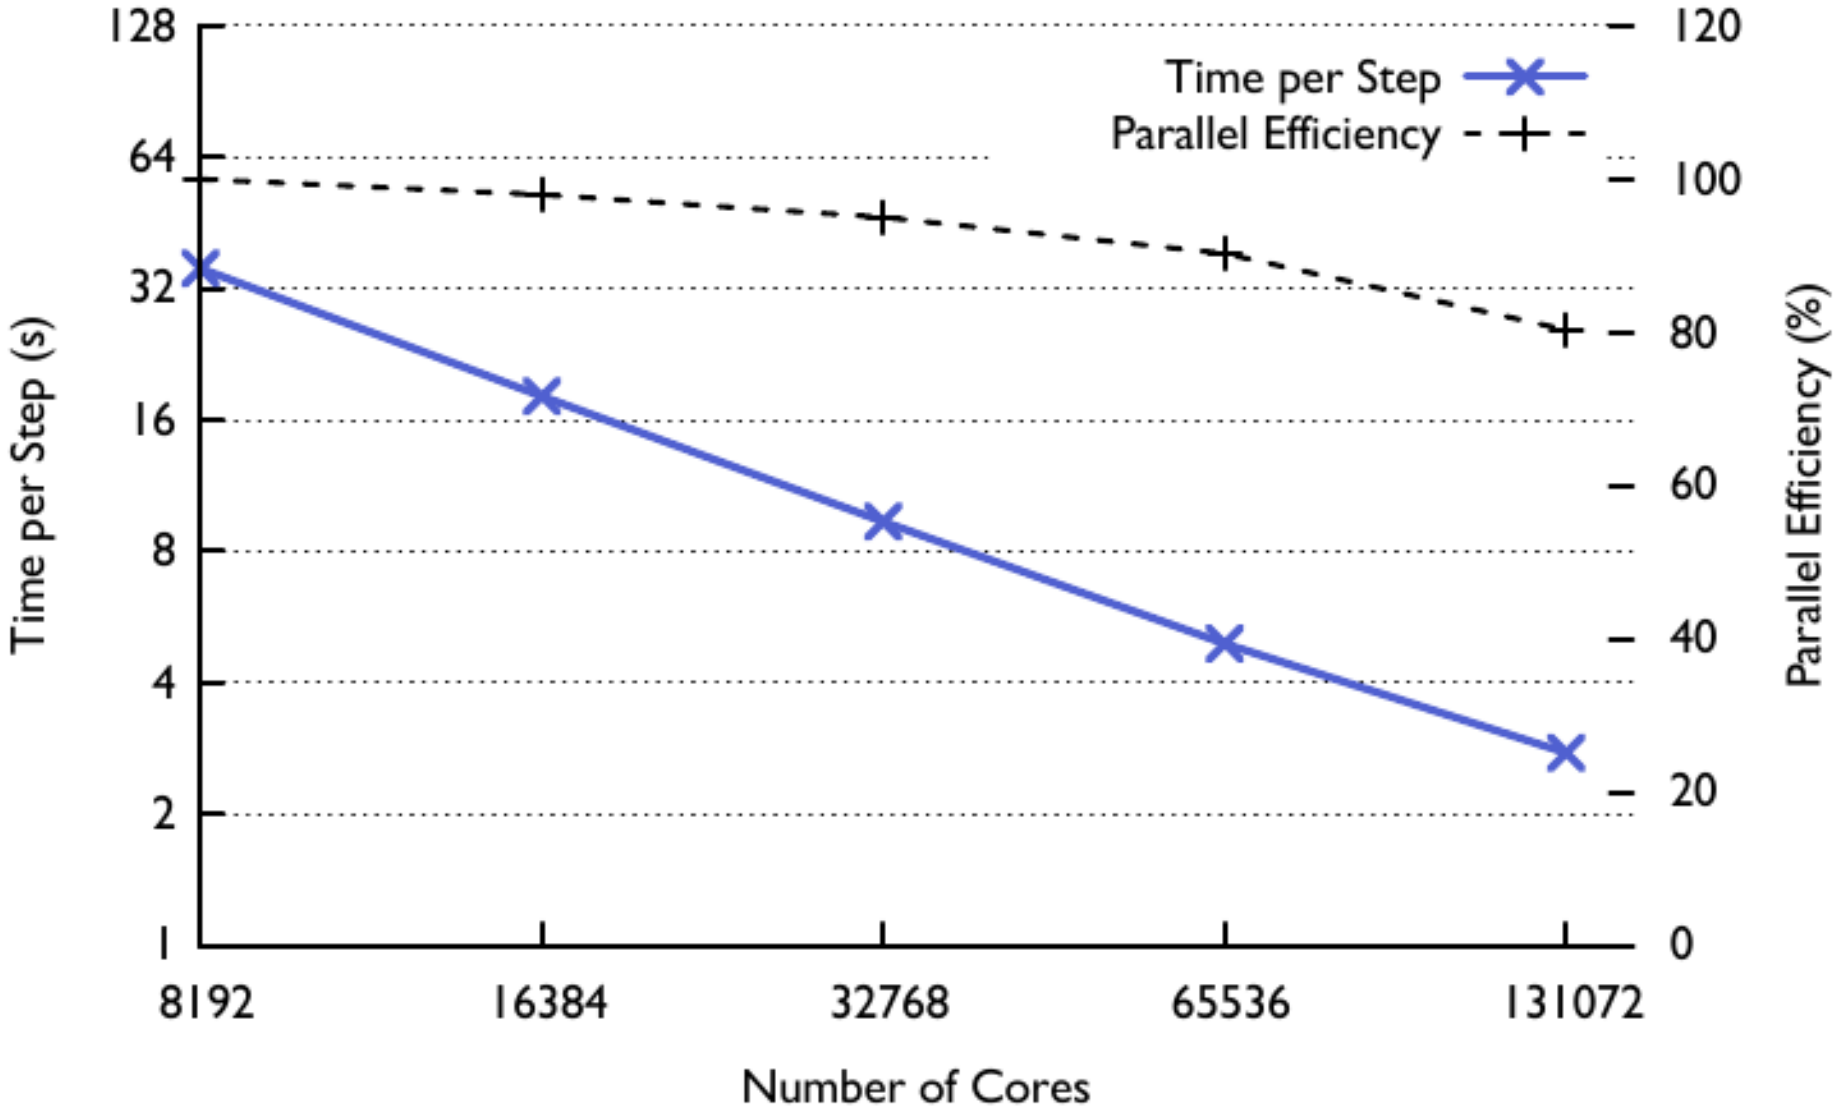

# The Vulcan

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

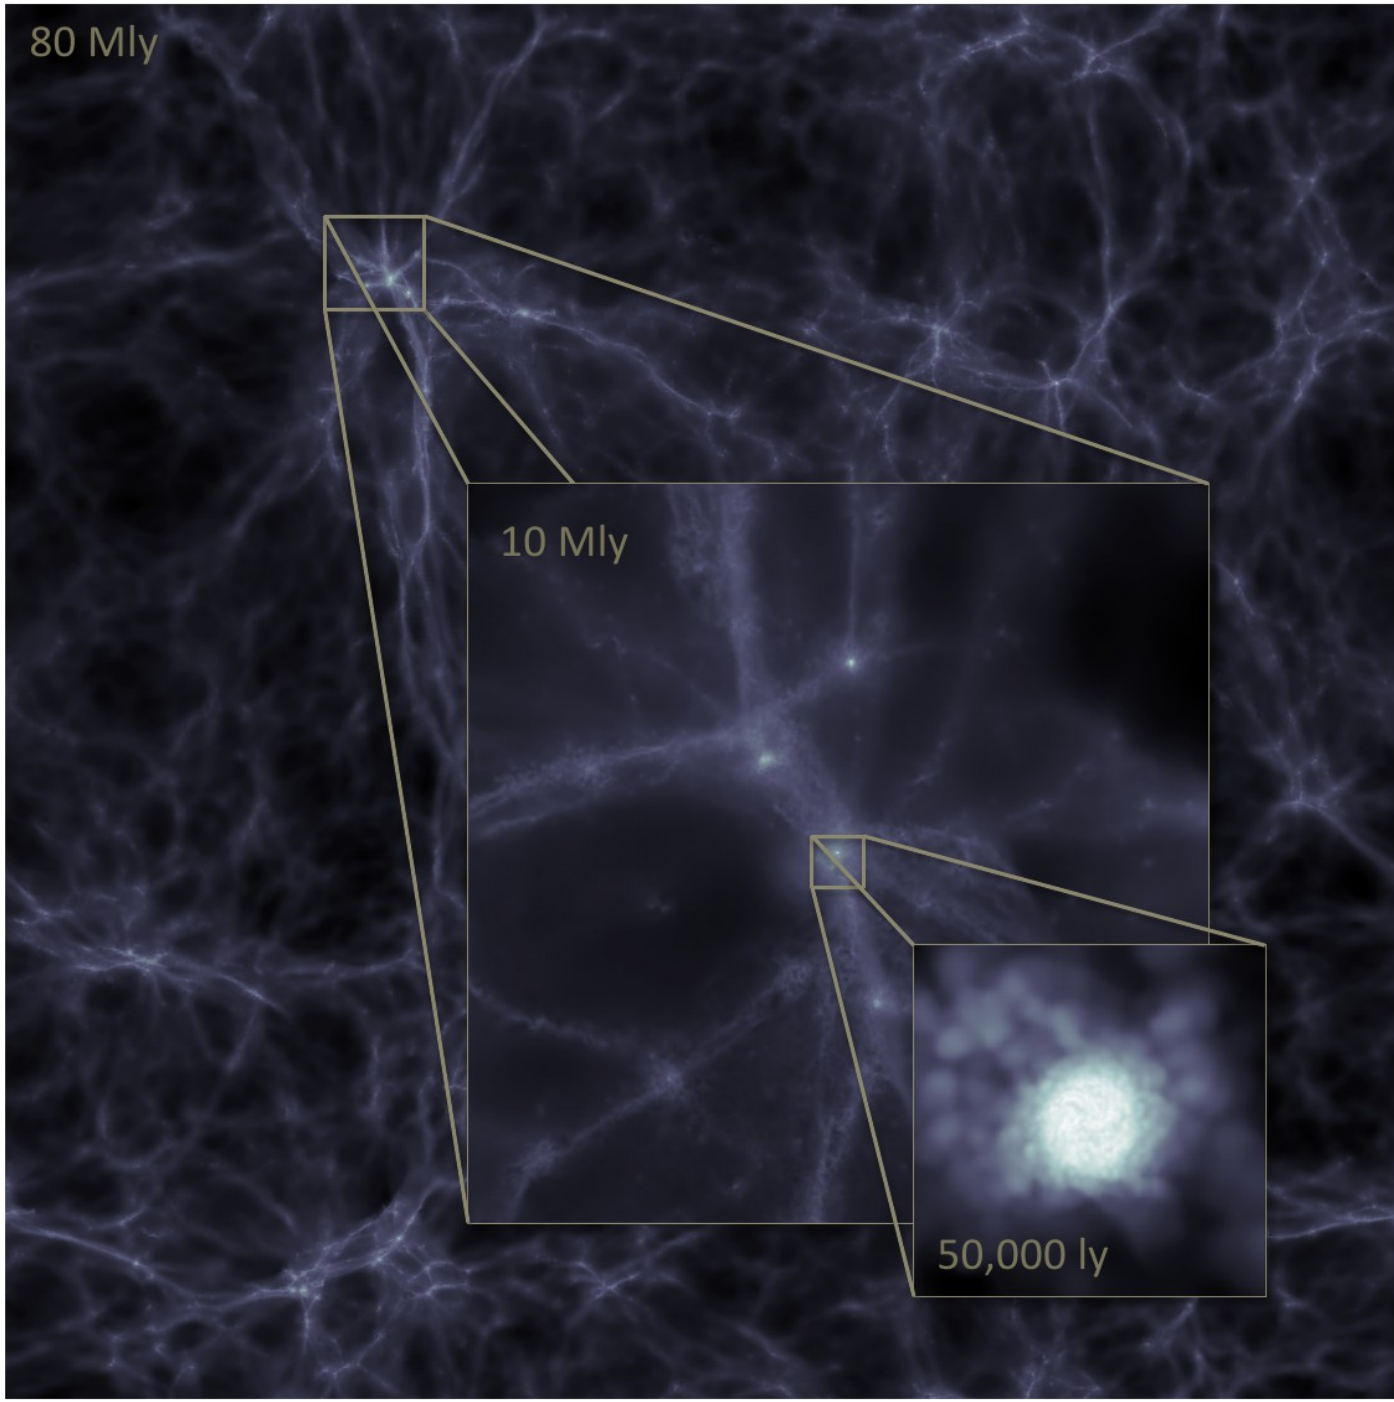

# Luminosity Function

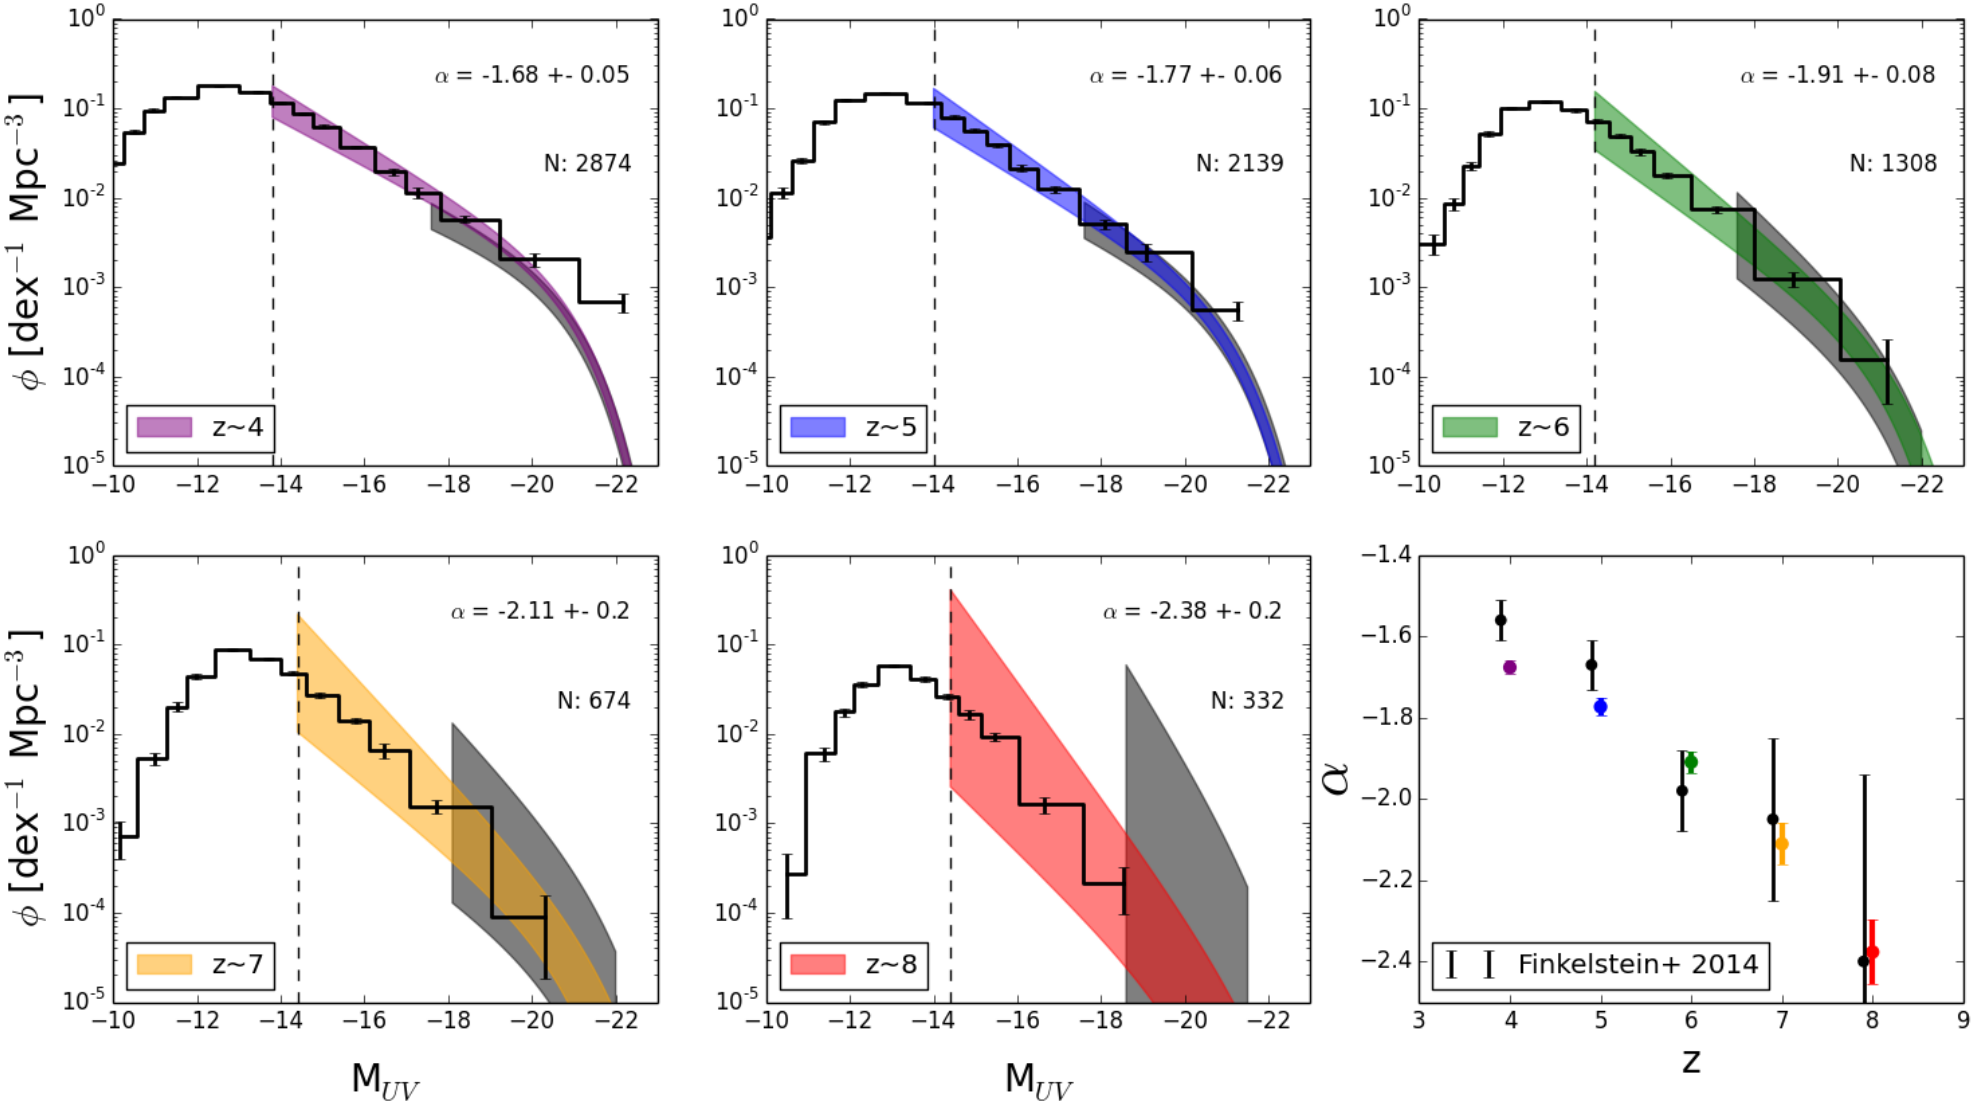


Comparison with Finkelstein+ 2014

L. Anderson+ 2015

$f_{esc} \sim 1\%$ $z \sim 4\text{-}6$

$f_{esc} < 1\%$ $z \sim 8$

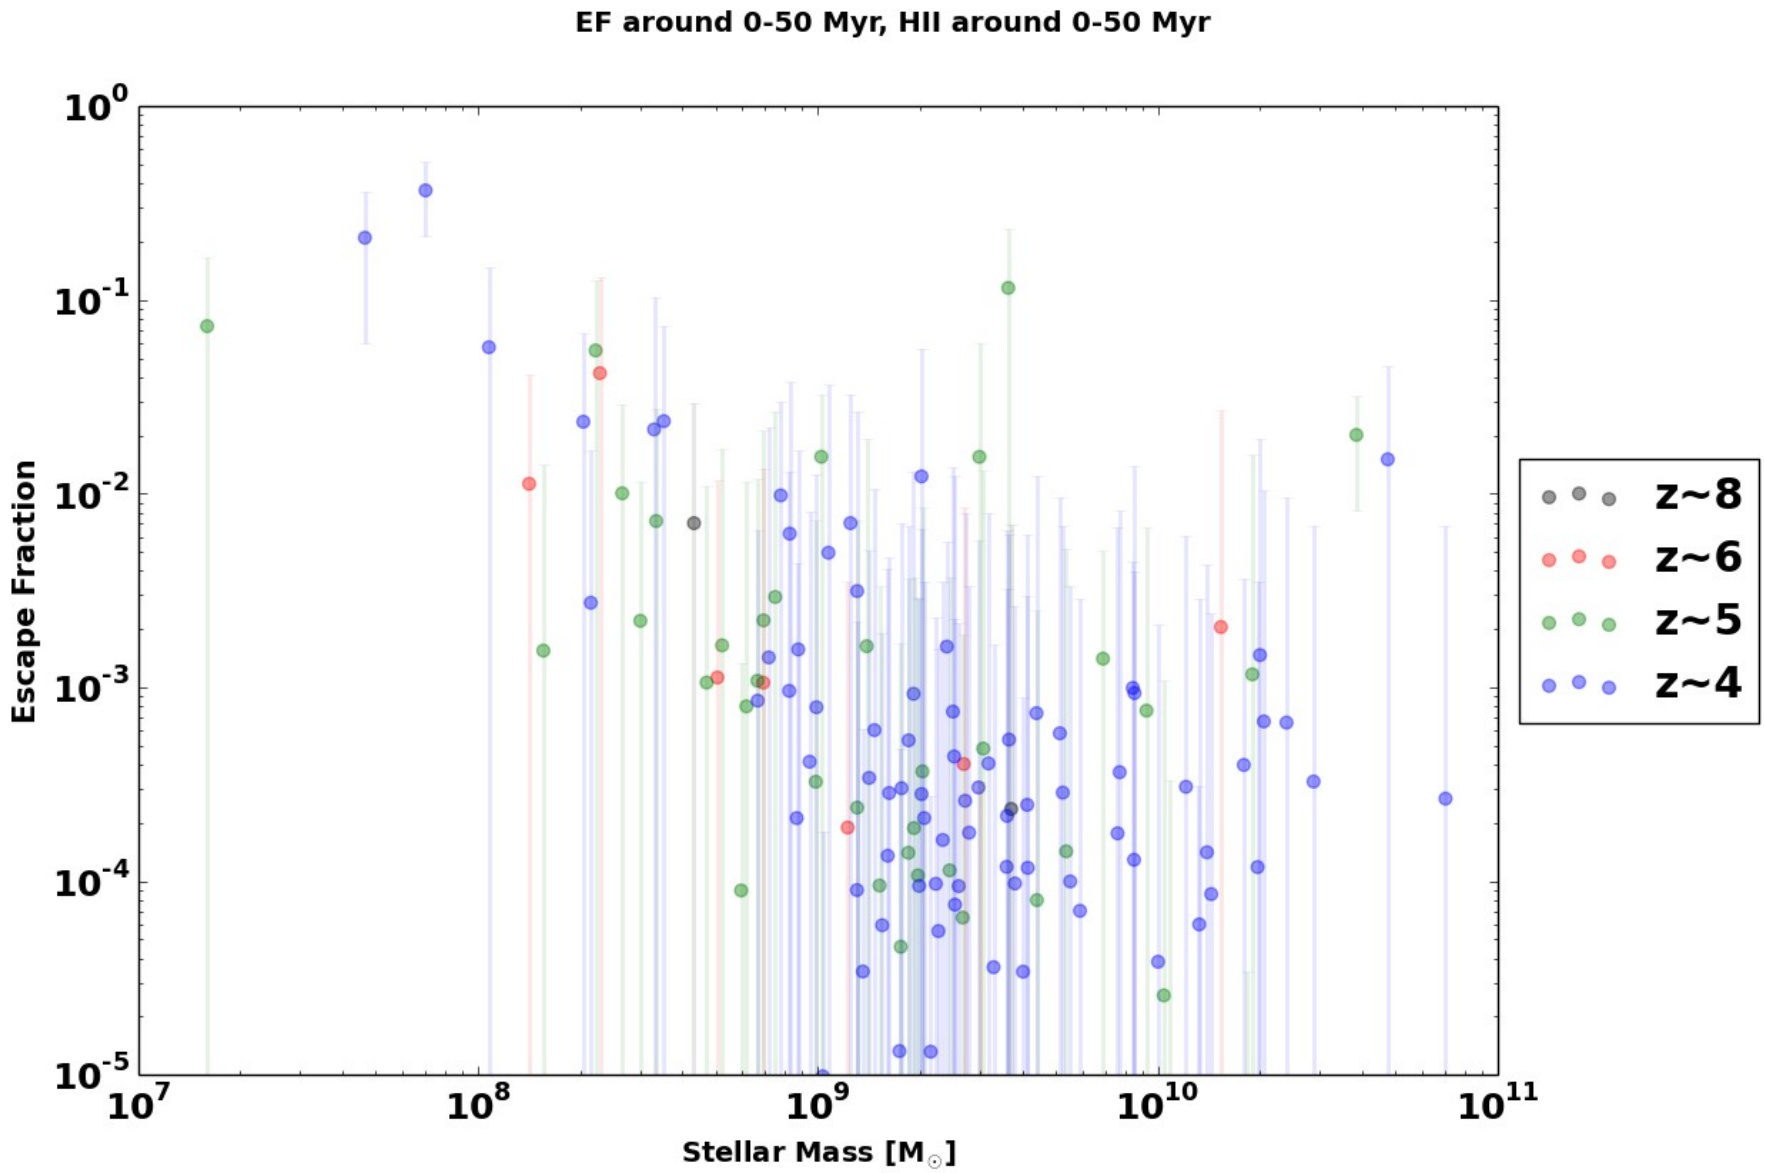

# Optical Depths Seen by One Individual Stellar Particle

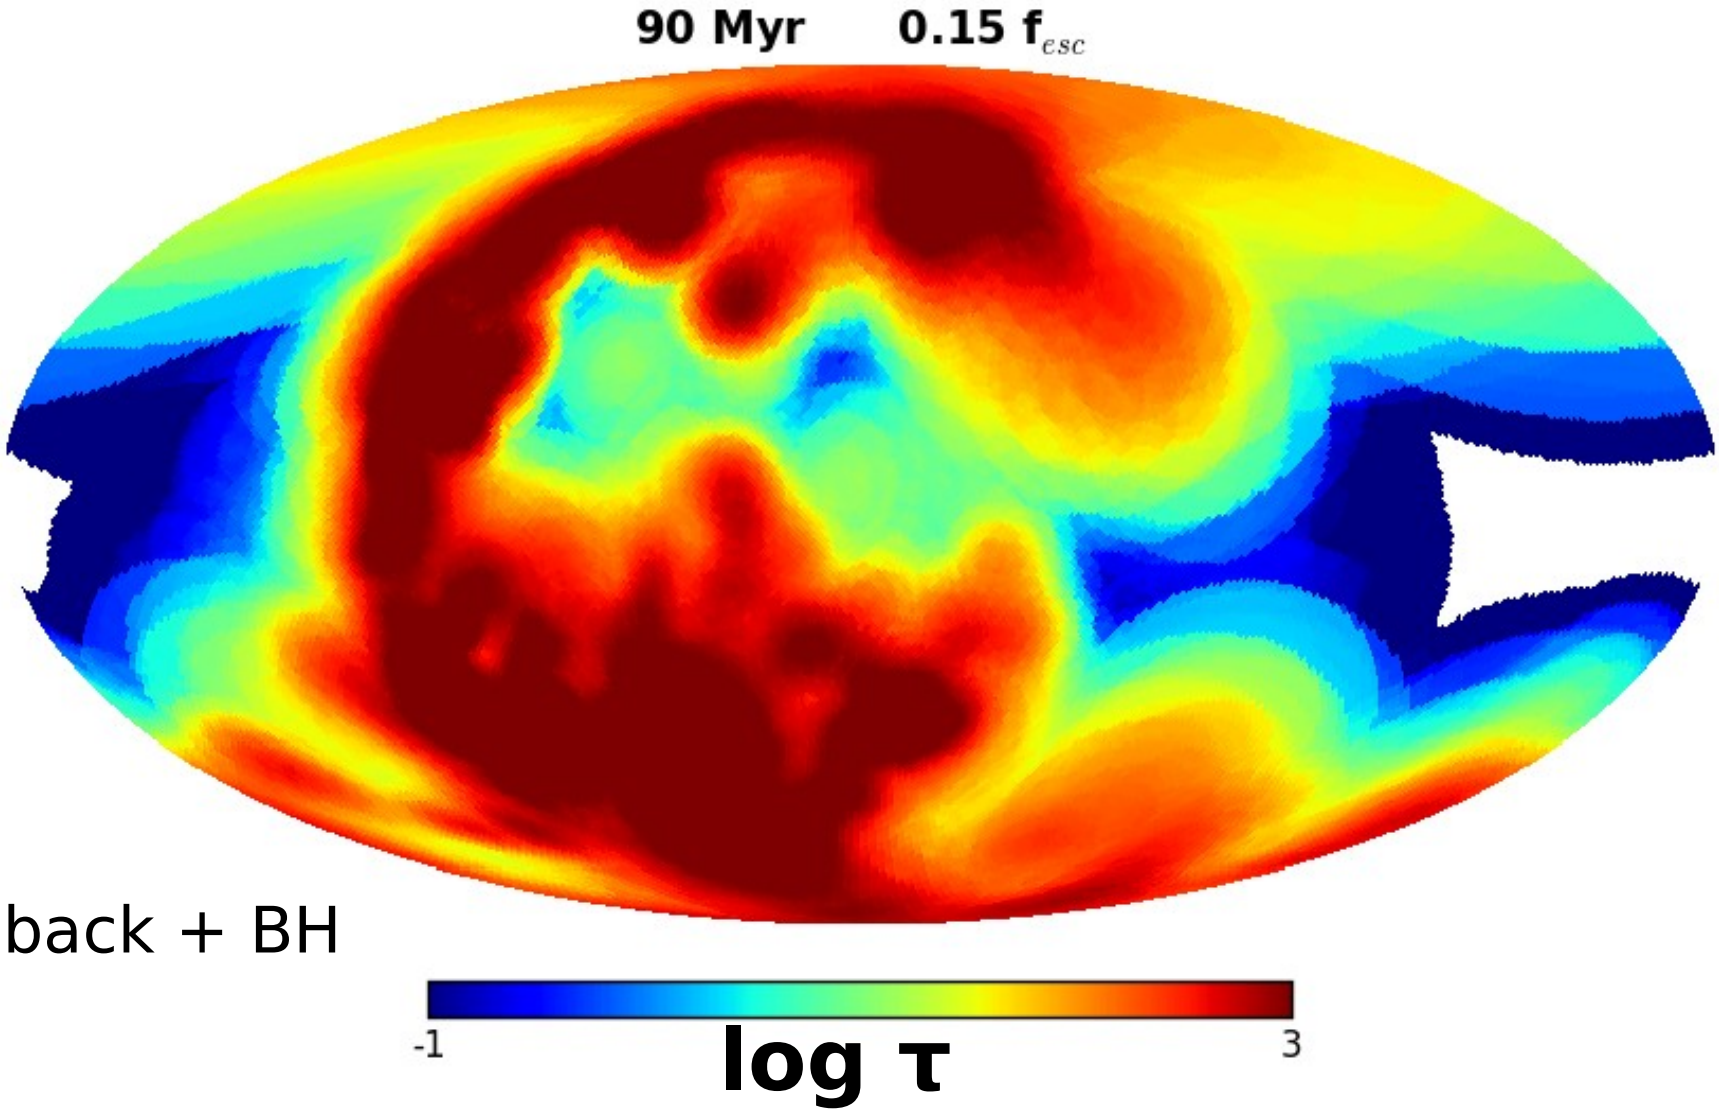


Stellar Feedback + BH

# Let’s add Black Holes!

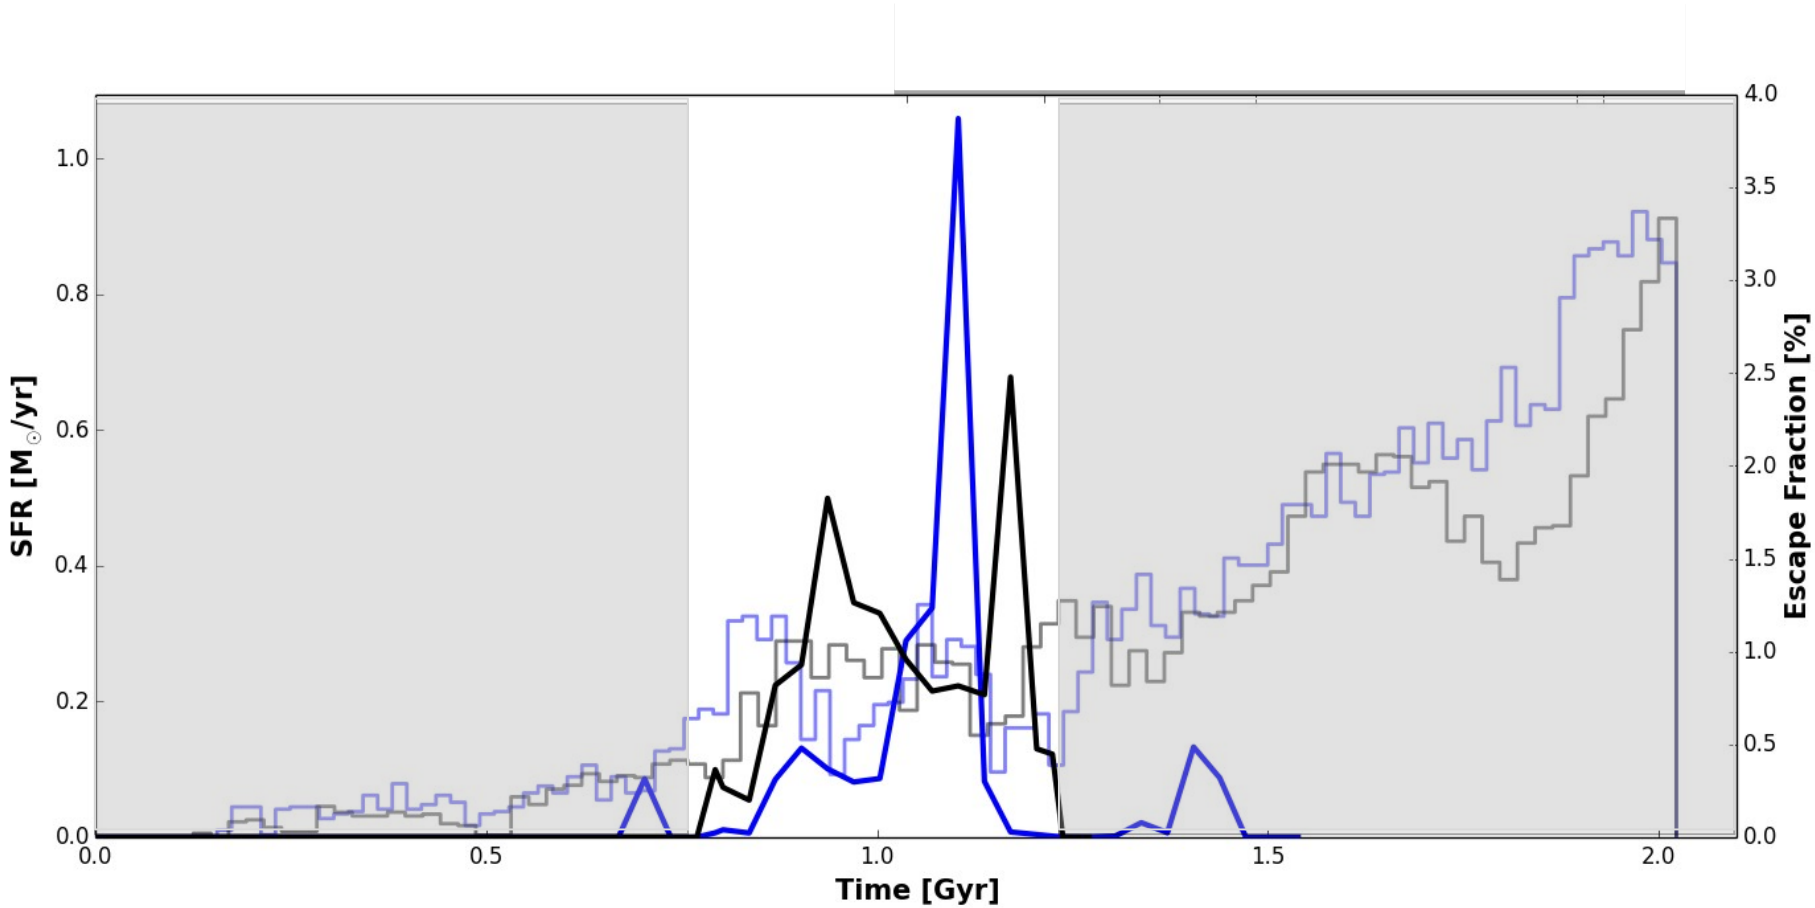


Future Work: Correlation with BH activity?

# Black Hole Dynamics

- BHs are not assumed to always be stable at the center of their host galaxies

- Unresolved dynamical friction is applied as a sub-grid model
  Tremmel+2015

- High DM mass resolution avoids numerical noise
  Bellovary+ 2010, Tremmel+2015

Tremmel+ 2015

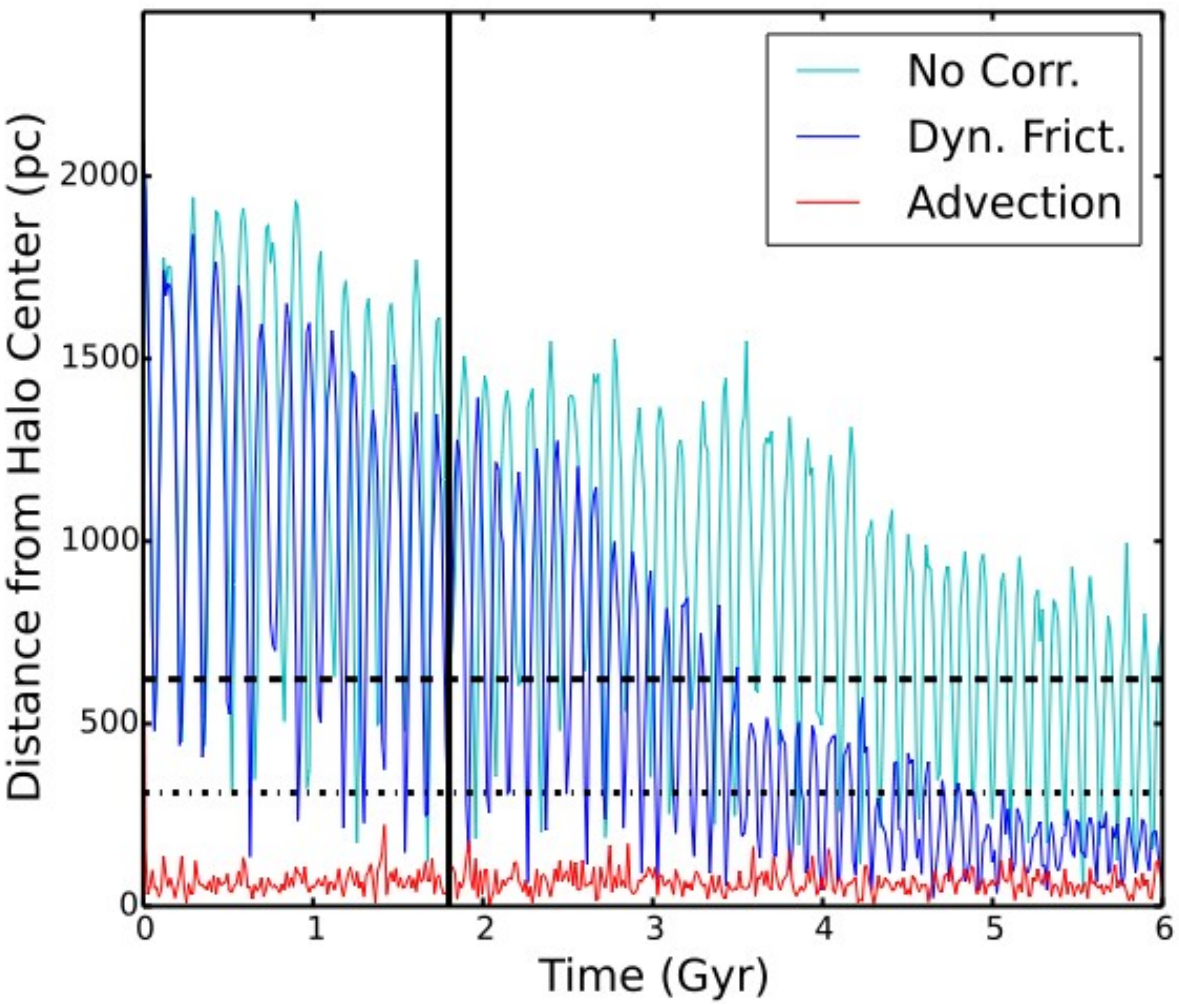

# Black Hole Feedback

# Black Hole Feedback

# First Results from Romulus

## Not All Galaxies Become Quenched

z = 3  z = 2  z = 1.2 z = 0.75 z = 0.5

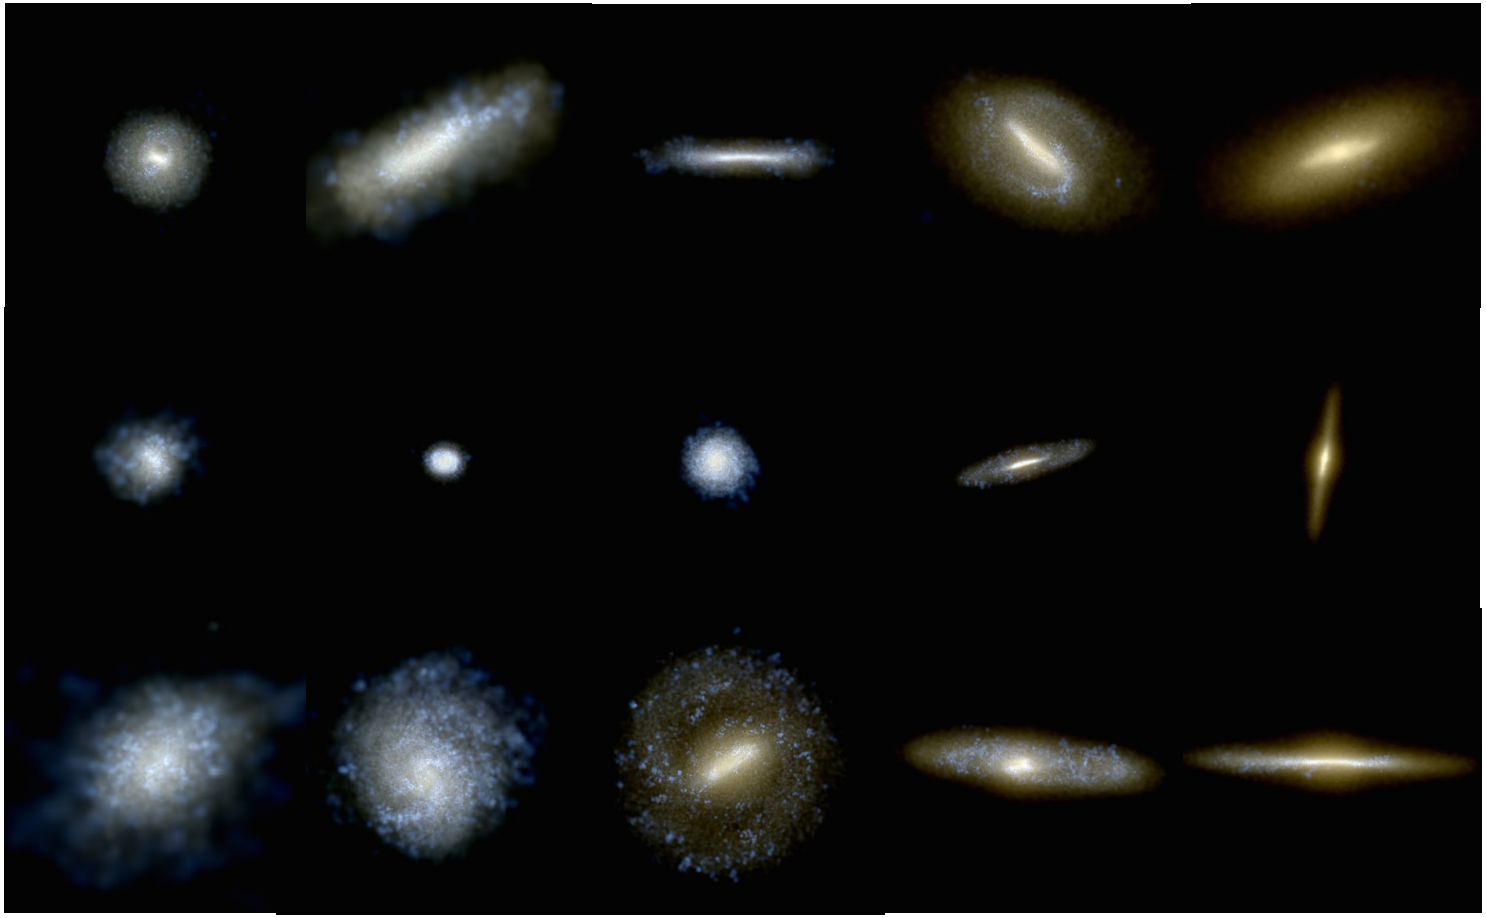

# Luminosity Function: Faint end slope

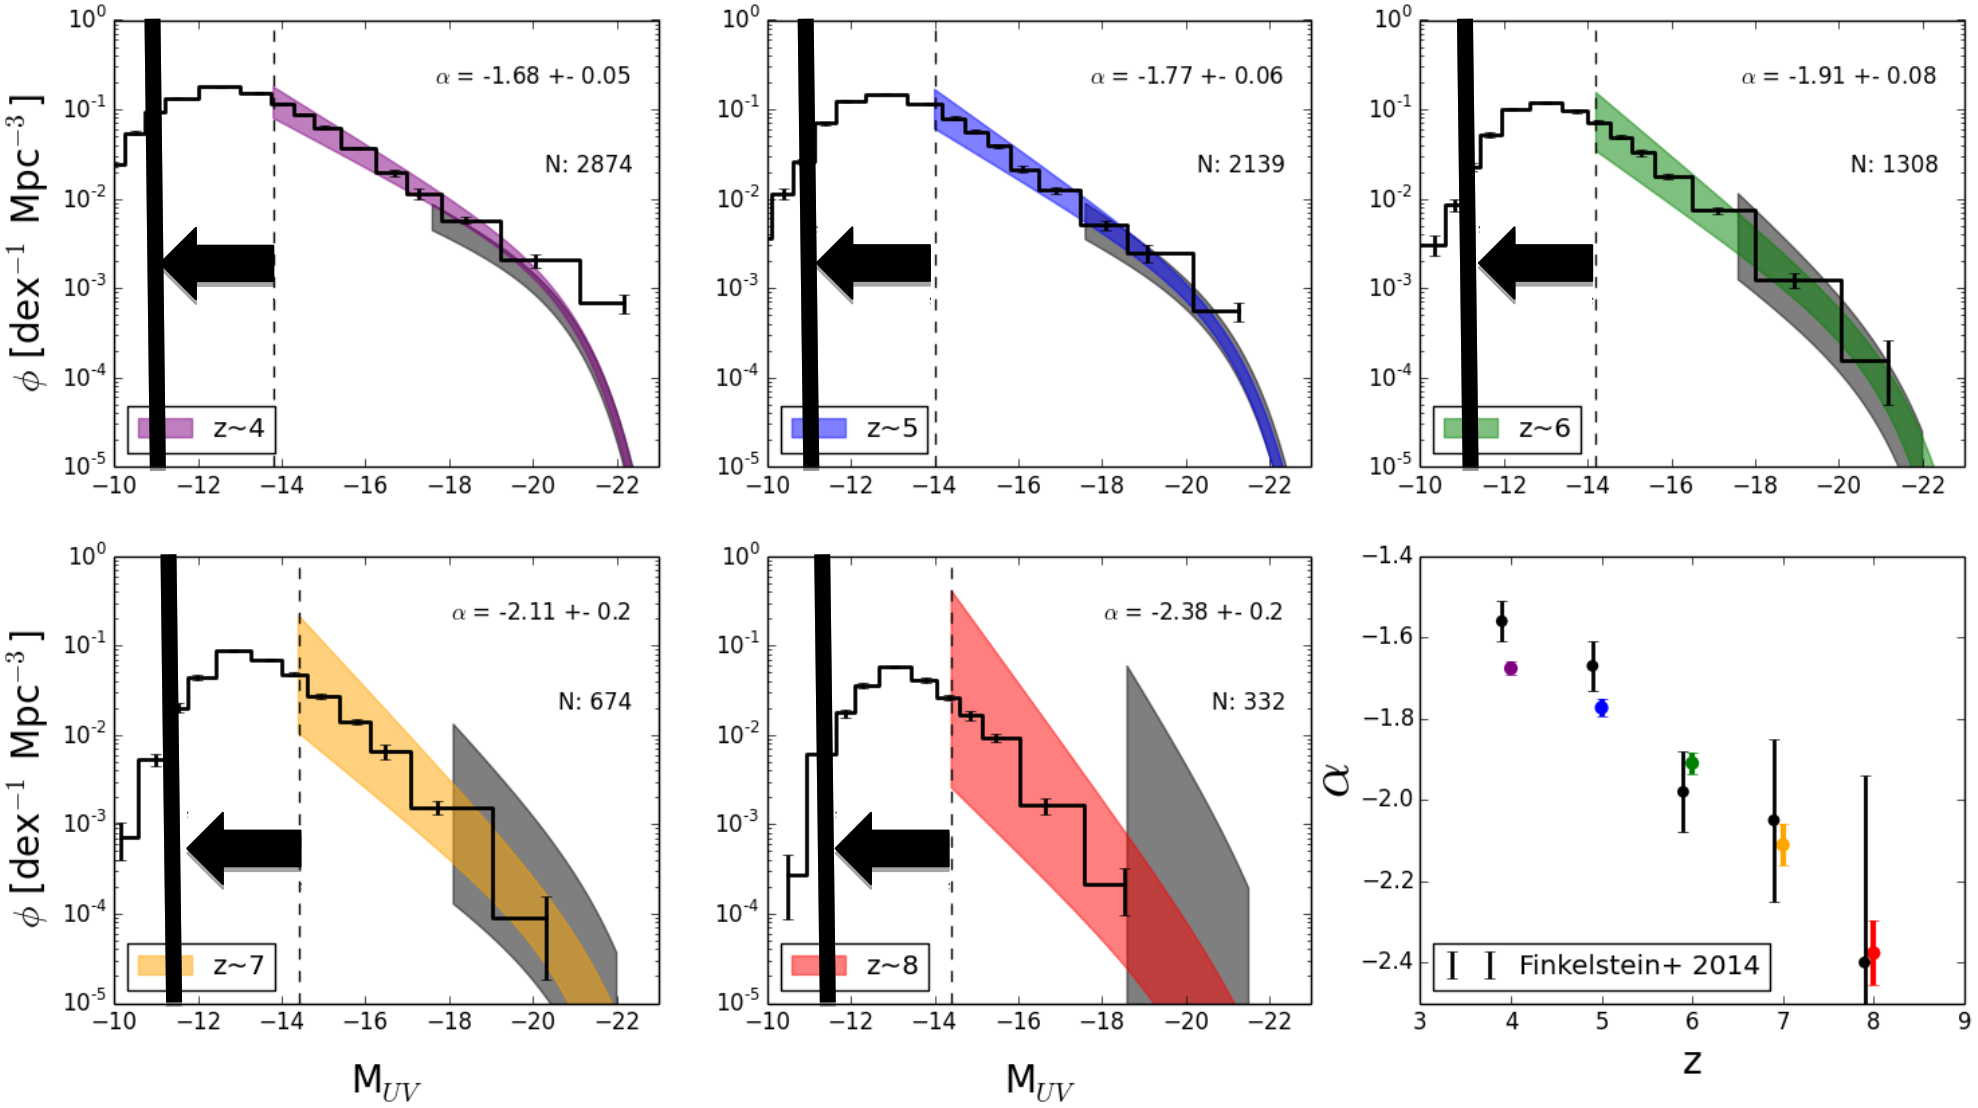

# Future Simulations

| | First Stage | Near Future |
|---|---|---|
| | Vulcan | Enterprise |
| Timeline | February 2014 | Autumn 2015 |
| Size | (25 Mpc)$^3$ | (25 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<br>H2 regulated star formation |

# Future Results

Predict faint end slope of LF

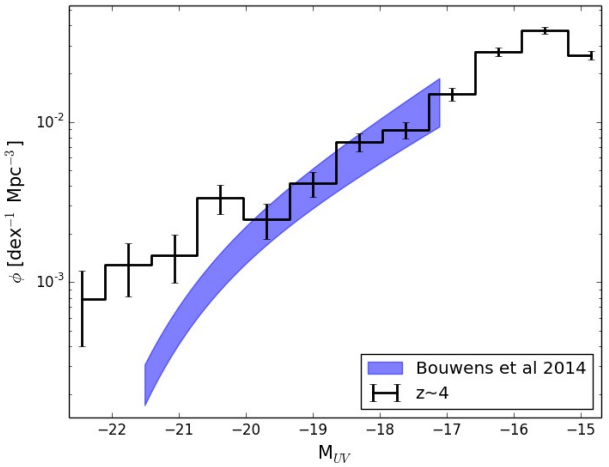


Measure escape fraction f(z, M, sfr, Z)

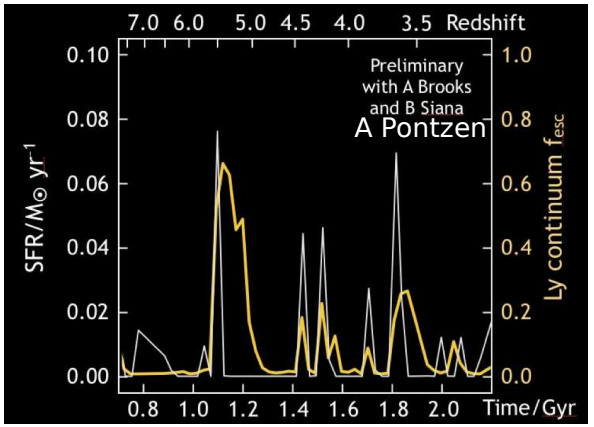


Morphologies of ~ 100 (1000) systems in Vulcan (Enterprise)

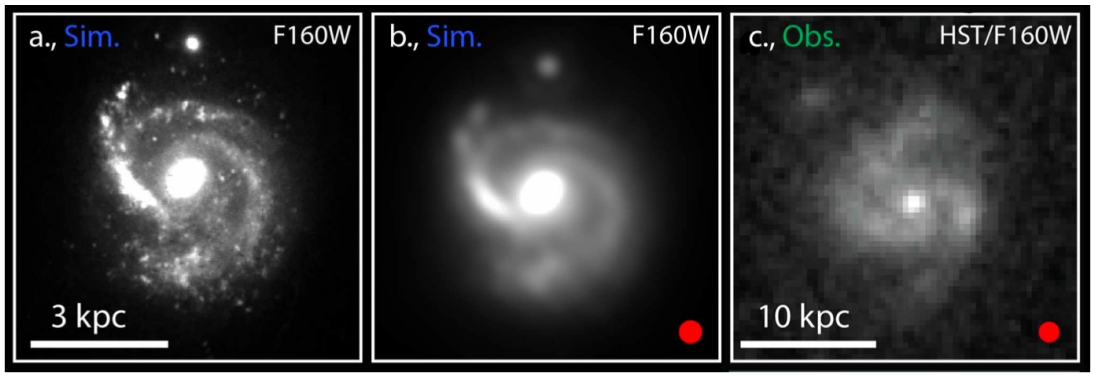


Law+ 2013

Evolution of SFR-$M_*$-Z relation

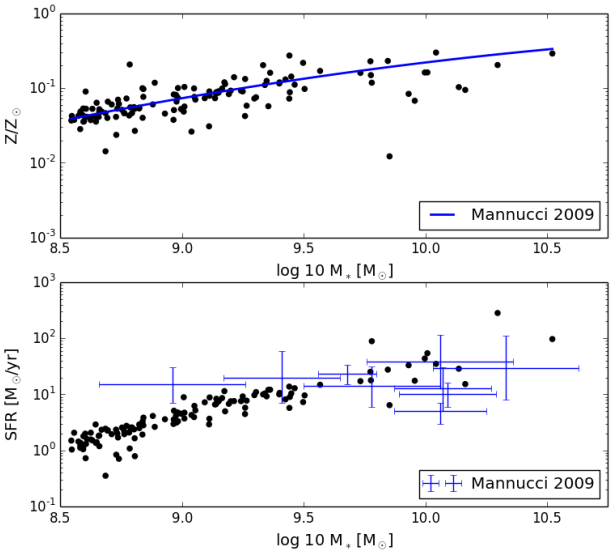


Stack for big data management and analytics

- A new big data mgmt & analytics system
  - Available as open source
  - Runs in shared-nothing clusters (Amazon EC2)
  - Also run in an HPC cluster at MIT
  - Think of it as Hive/Hadoop but faster
  - Think of it as Spark but faster
- An **operational service** deployed at UW

- Developed by the UW database group and eScience

# Myria Stack

Browser

MyMergerTree

Python

MyriaQ

Web UI

Language Parser

Query Optimizer

Google App Engine

MyriaX

Parallel Query Execution

# Acknowledgments


- NSF ITR
- NSF Astronomy
- NSF XSEDE program for computing
- BlueWaters Petascale Computing
- NASA HST
- NASA Advanced Supercomuting

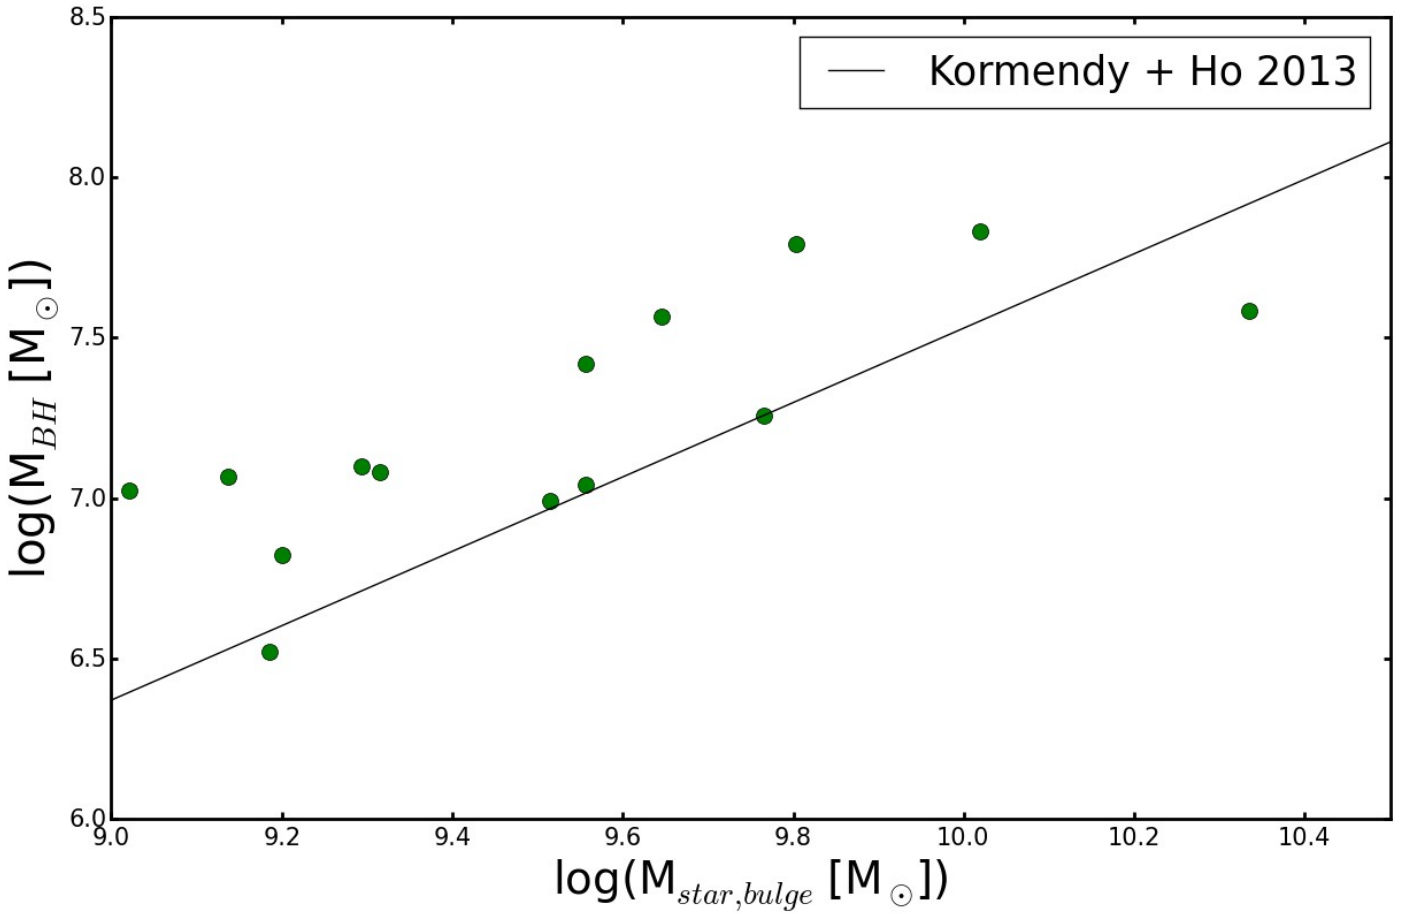
Kormendy + Ho 2013
log(M$_{BH}$ [M$_\odot$])
log(M$_{star,bulge}$ [M$_\odot$])
8.5
8.0
7.5
7.0
6.5
6.0
9.0
9.2
9.4
9.6
9.8
10.0
10.2
10.4

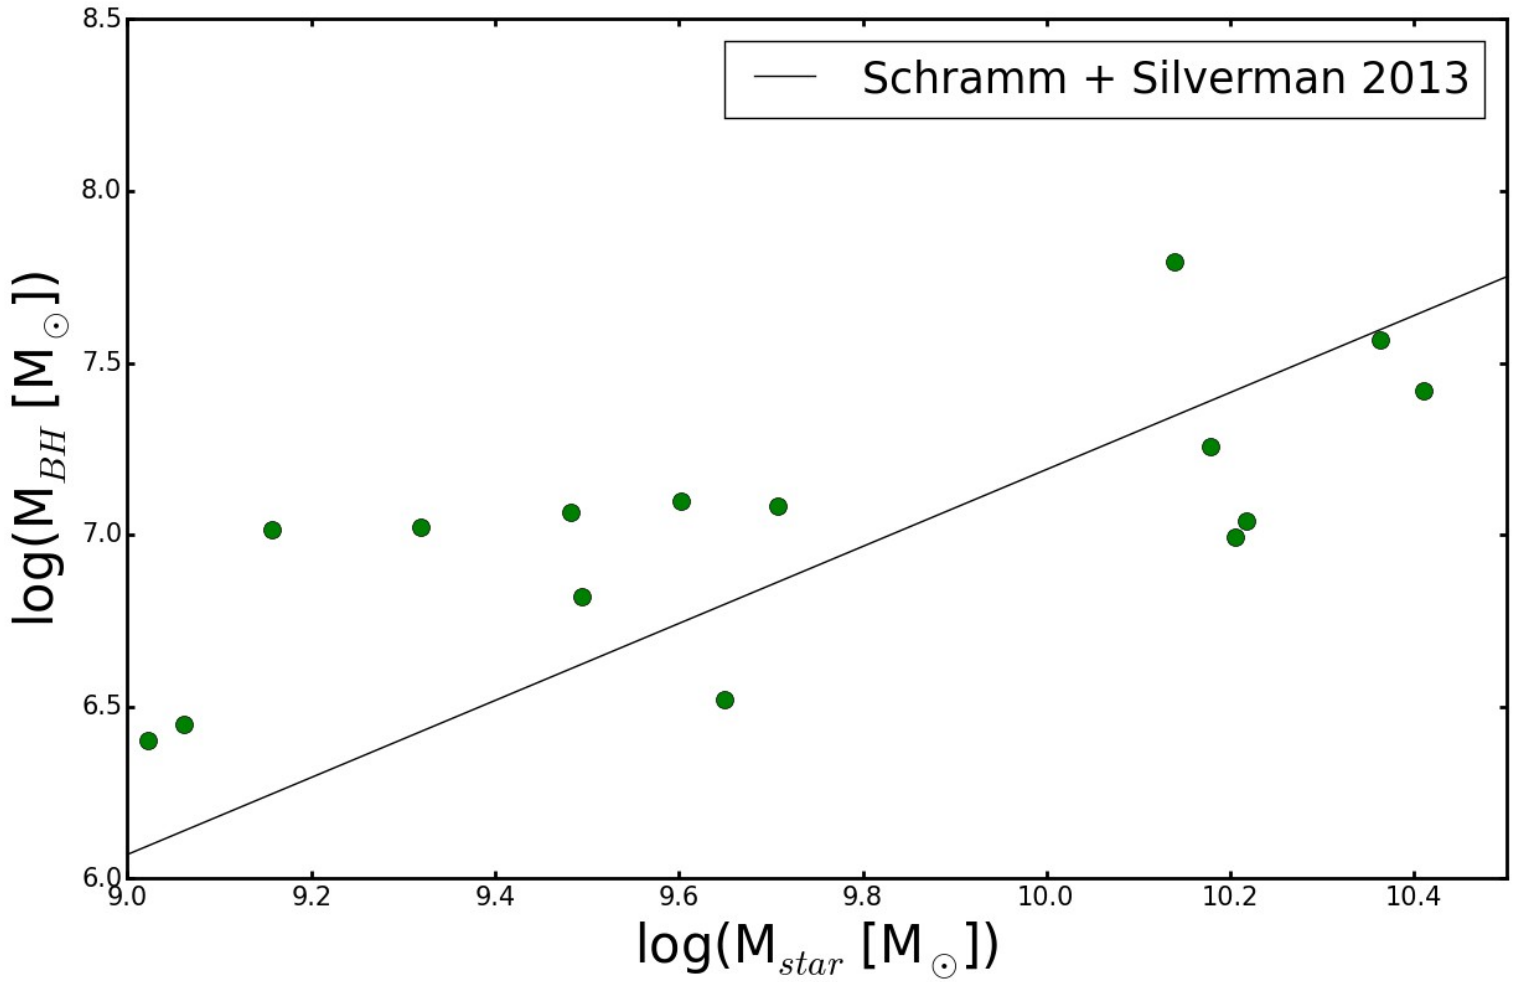
Schramm + Silverman 2013
log(M_BH [M_⊙])
log(M_star [M_⊙])
8.5
8.0
7.5
7.0
6.5
6.0
9.0
9.2
9.4
9.6
9.8
10.0
10.2
10.4