Enzo Lectures

Mike Norman, Matt Turk
Laboratory for Computational Astrophysics
UC San Diego

	Morning	Afternoon
Mon.	Introduction to Enzo	Enzo projects
Tue.	Setting Up and Running Enzo	Introduction to YT
Wed.	Enzo Algorithms	Lab session
Thu.	Applications to First Stars, First Galaxies, and Reionization	Lab session
Fri.	What’s New in Enzo 2.0?	Q & A

Enzo Resources

- Enzo website (code, documentation)
  - http://lca.ucsd.edu/projects/enzo
- 2010 Enzo User Workshop slides
  - http://lca.ucsd.edu/workshops/enzo2010
- yt website (analysis and vis.)
  - http://yt.enzotools.org
- Jacques website (analysis and vis.)
  - http://jacques.enzotools.org/doc/Jacques/Jacques.html

Introduction to Enzo: Topics

- What is Enzo?
- Enzo Physics
- Structured AMR
- Algorithms (Overview)
- Parallelism and Scalability
- Application Sampler

What is Enzo?

- Started as an AMR hydro cosmology code developed by Greg Bryan in 90s, but is now more general purpose with many developers and application areas
- Multiscale, multiphysics, parallel
- Open source community code
  - Over 100 users, 30 developers
- Major new release (V2.0) coming very soon

Example of What Enzo Can Do

AMR galaxy formation and mergers

Enzo Physics (V2.0)

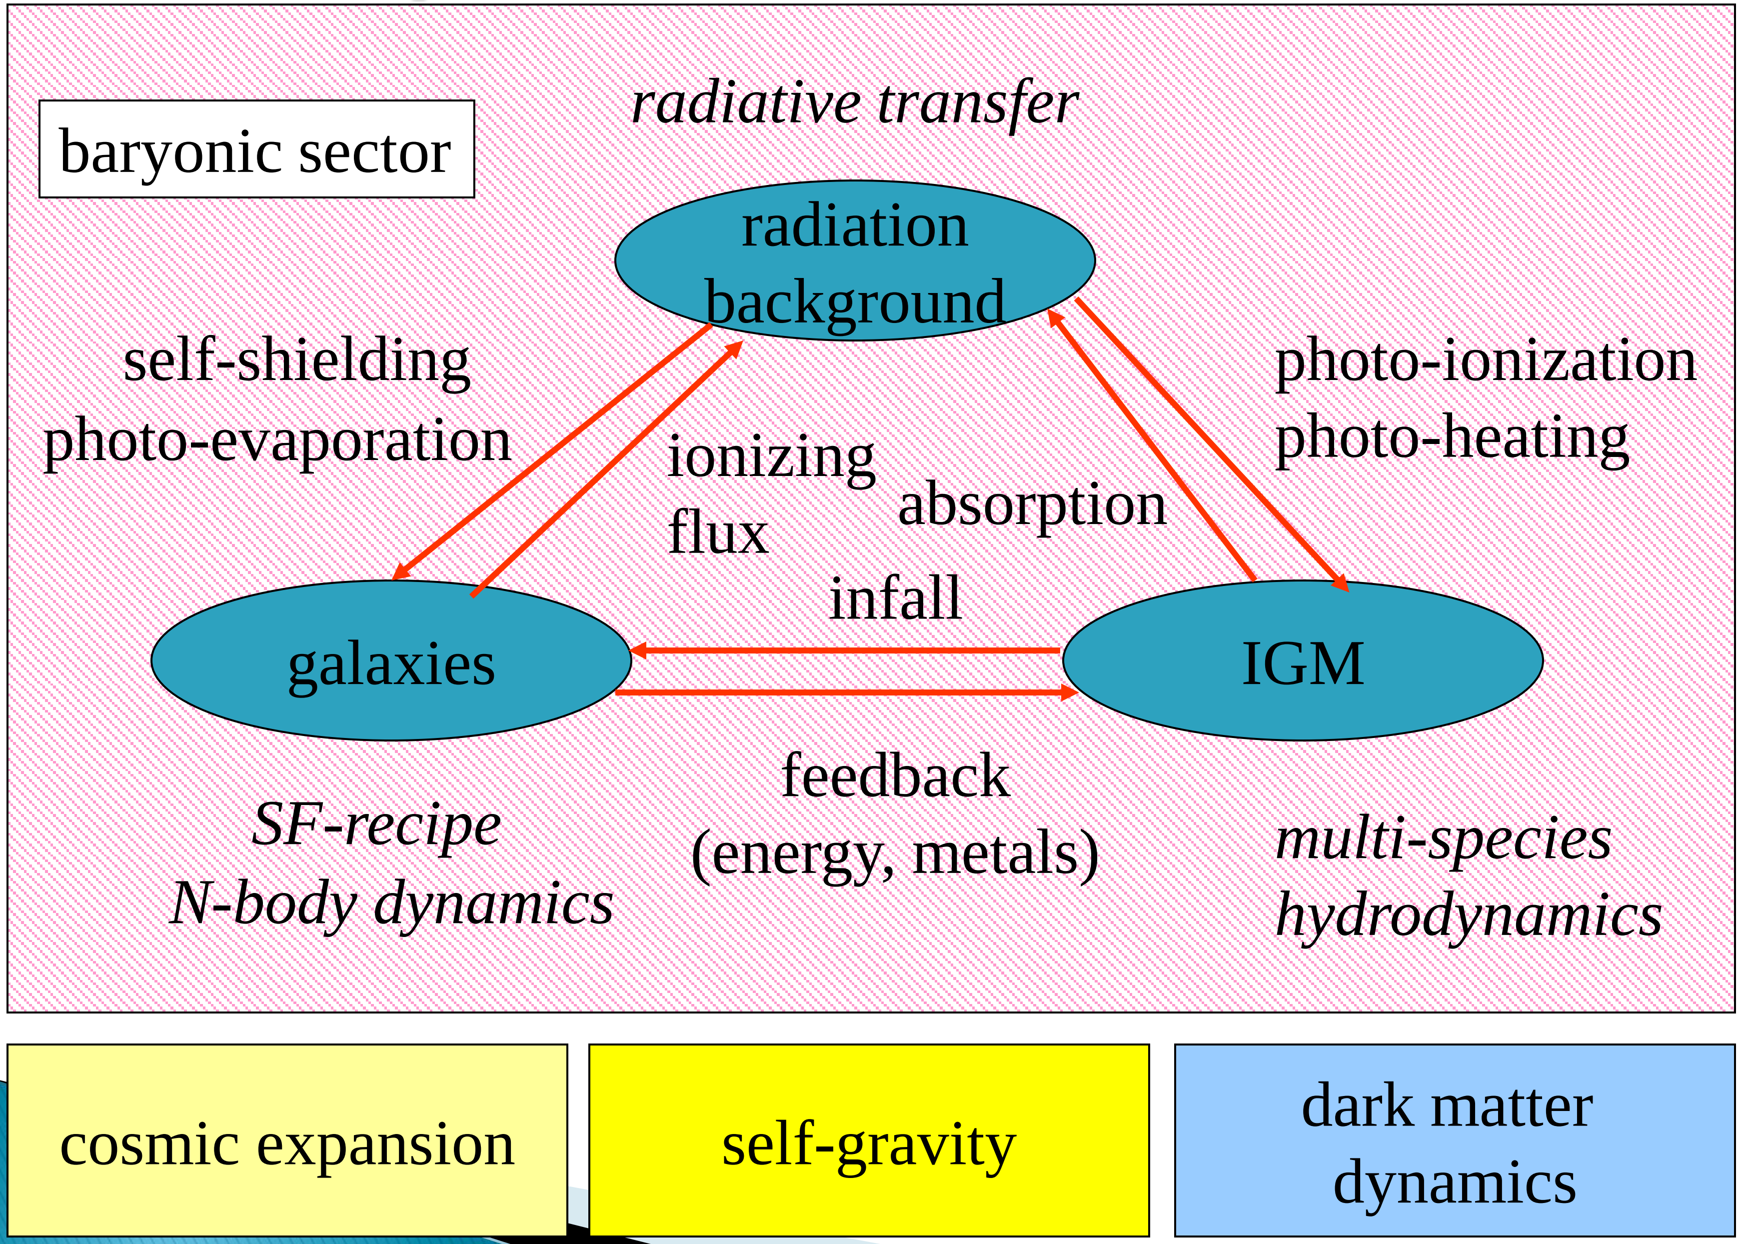

Resolving Galaxy Formation: Multiscale Challenge

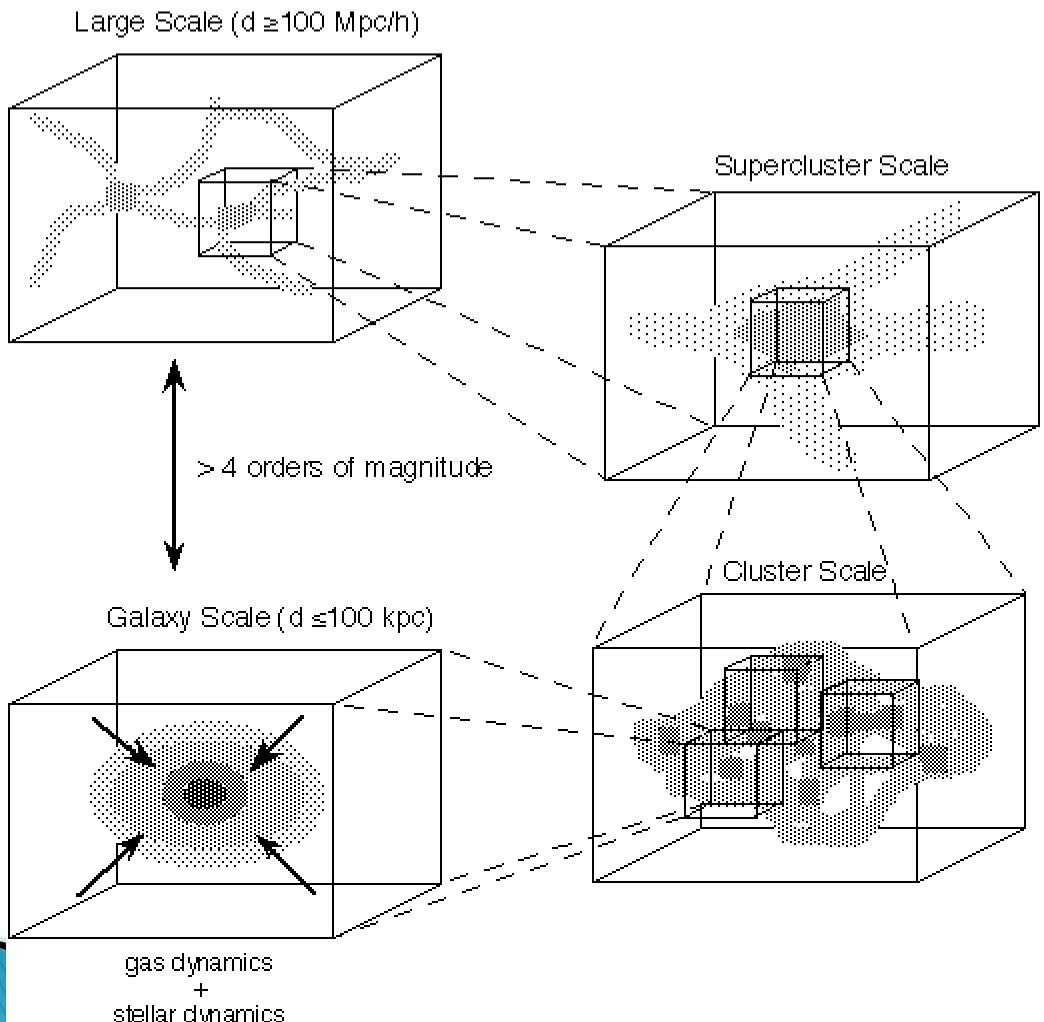

dynamic range requirement:

> 10^4 spatial

> 10^9 mass

Structured Adaptive Mesh Refinement
(Berger and Colella 1989)

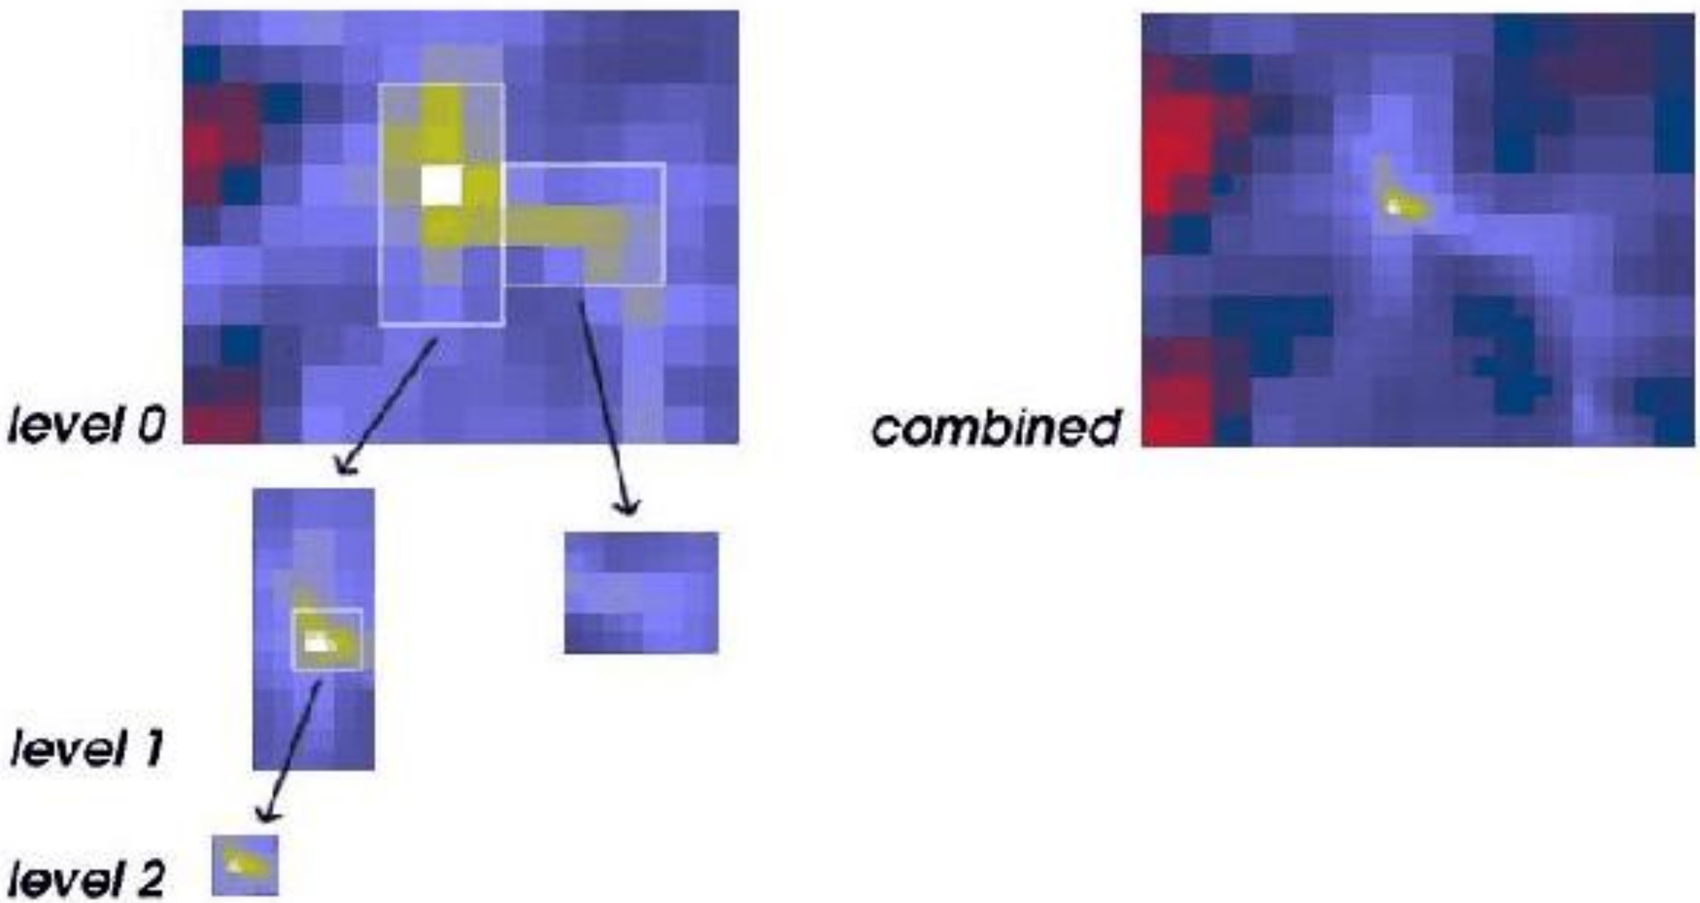

How Mesh Refinement is Done

- Flag all cells on a given level L “true” if they meet some physically motivated condition, or boolean of conditions
- Construct bounding boxes whose side lengths are between *min* and *max*
- Create a level L+1 subgrid for each box refined by a factor of R relative to parent
- Initialize data by interpolating from overlapping parent grid


First steps with Enzo
Enzo User Workshop
June 28-30, 2010

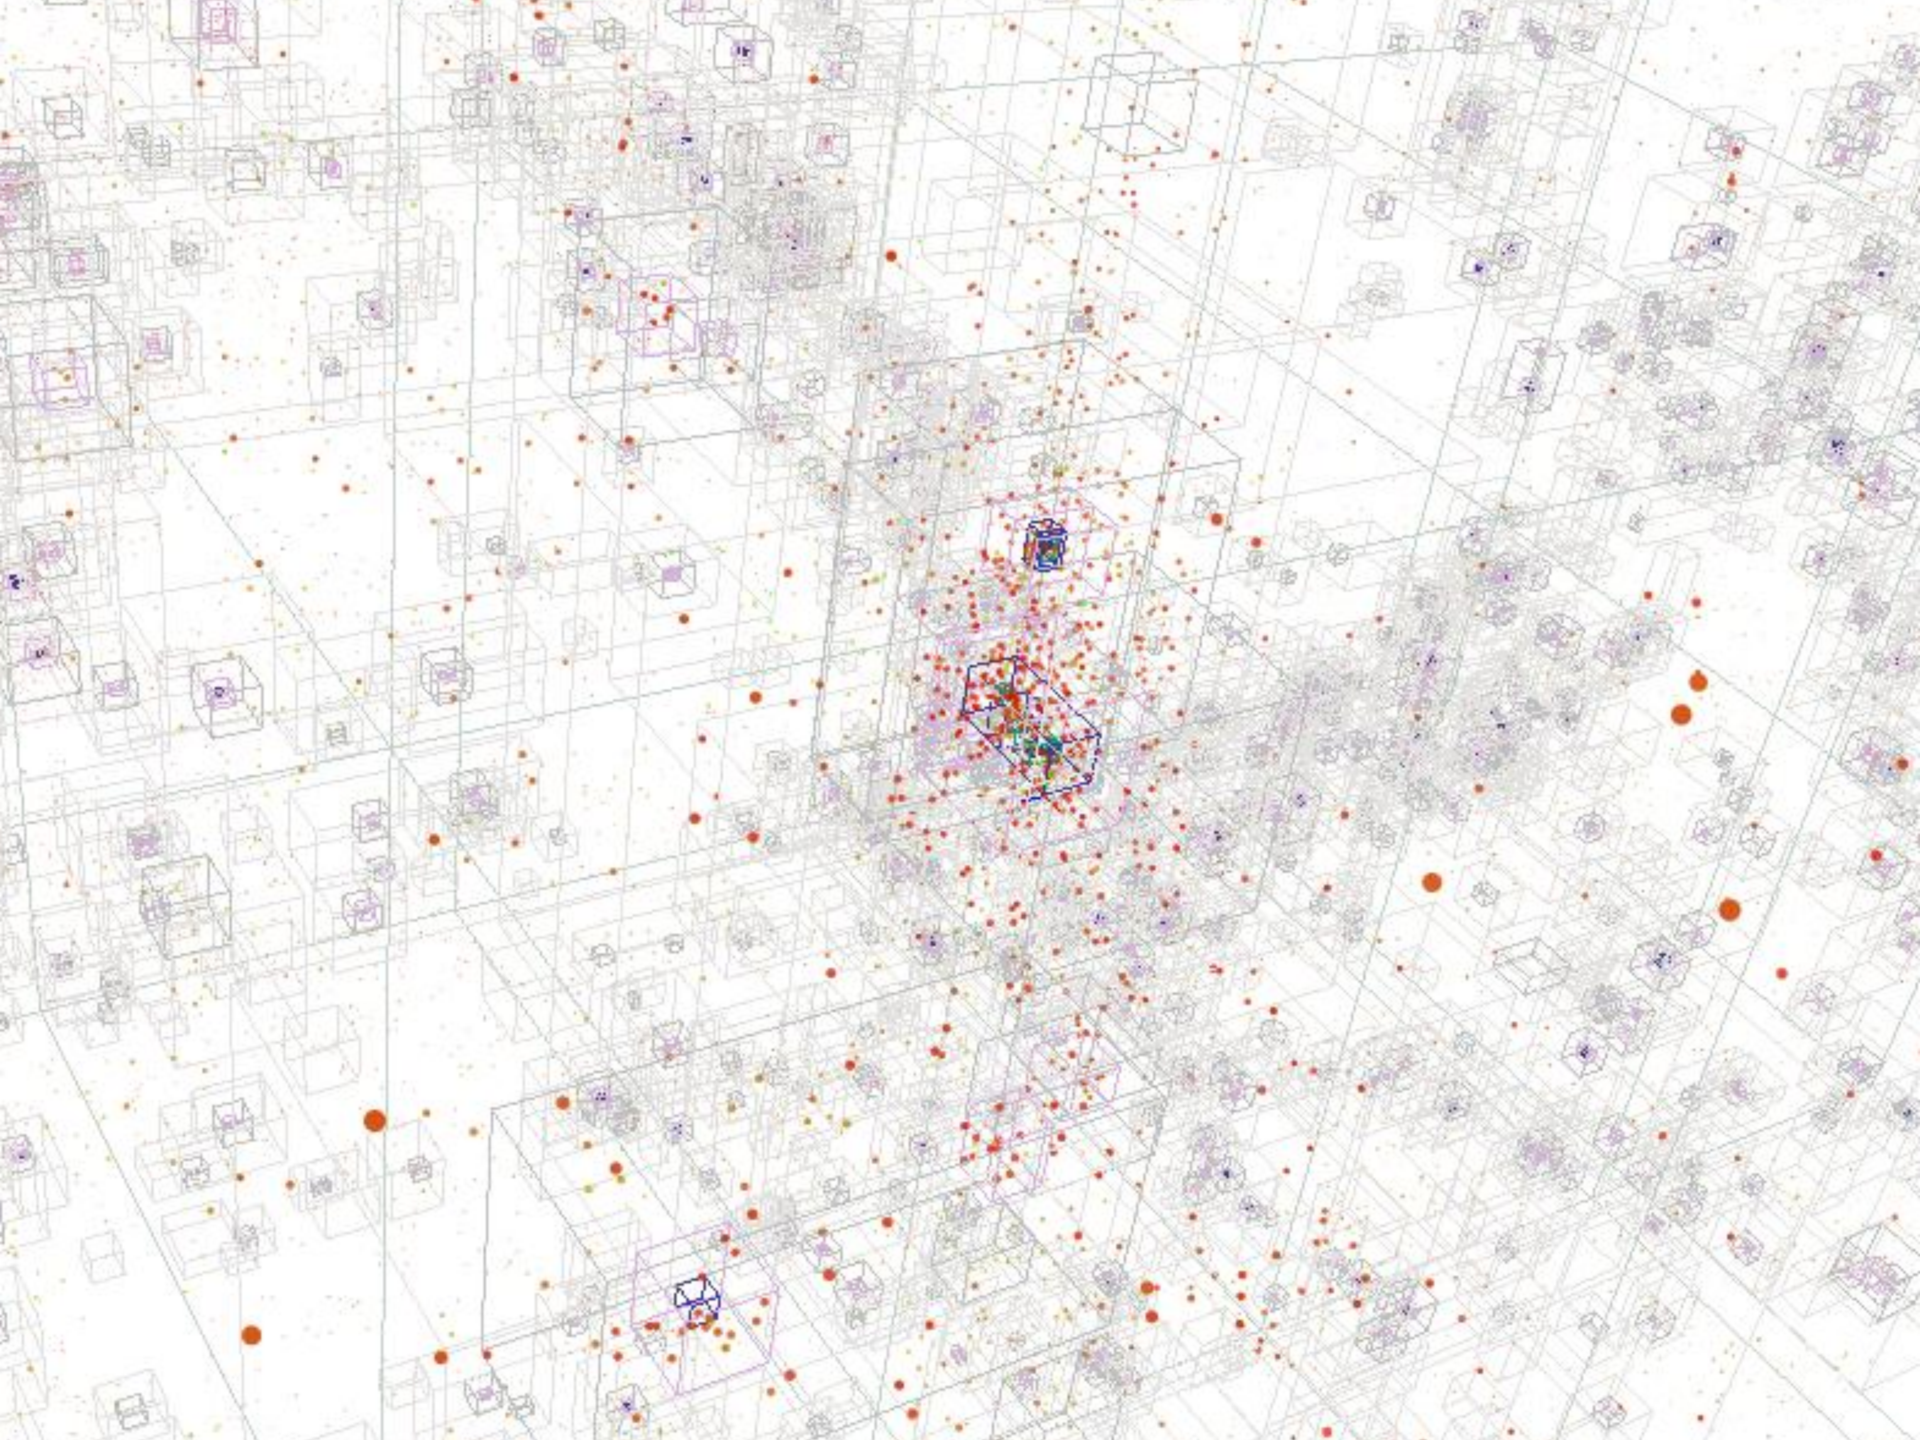

Hierarchical Timestepping

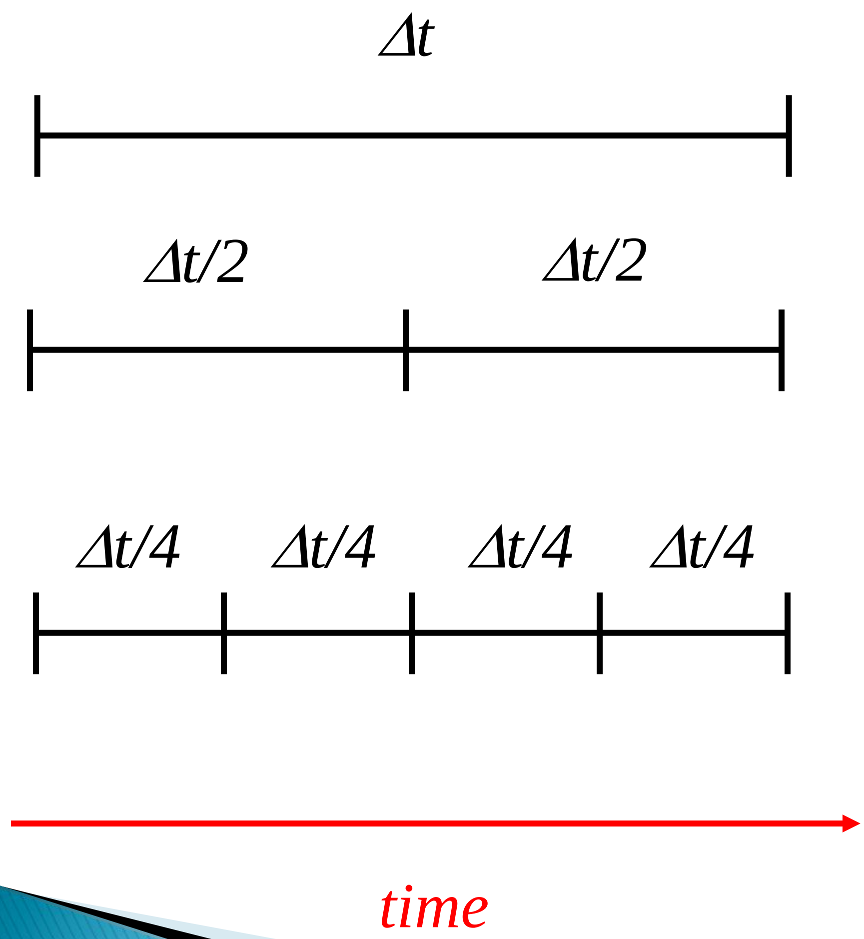

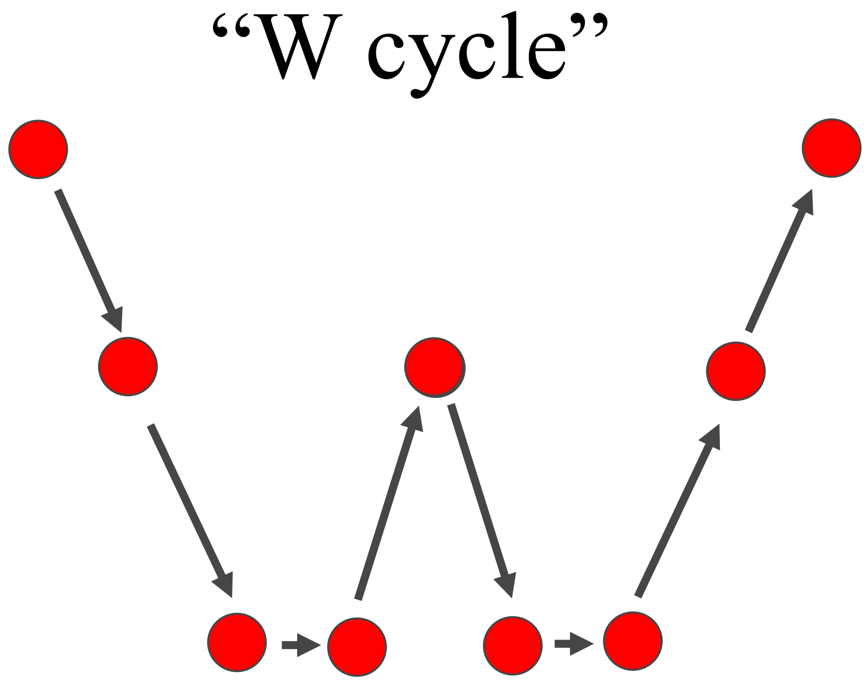

--2nd order-accurate

-- opposite to ART

Main Loop of Hydrodynamic Cosmology Code

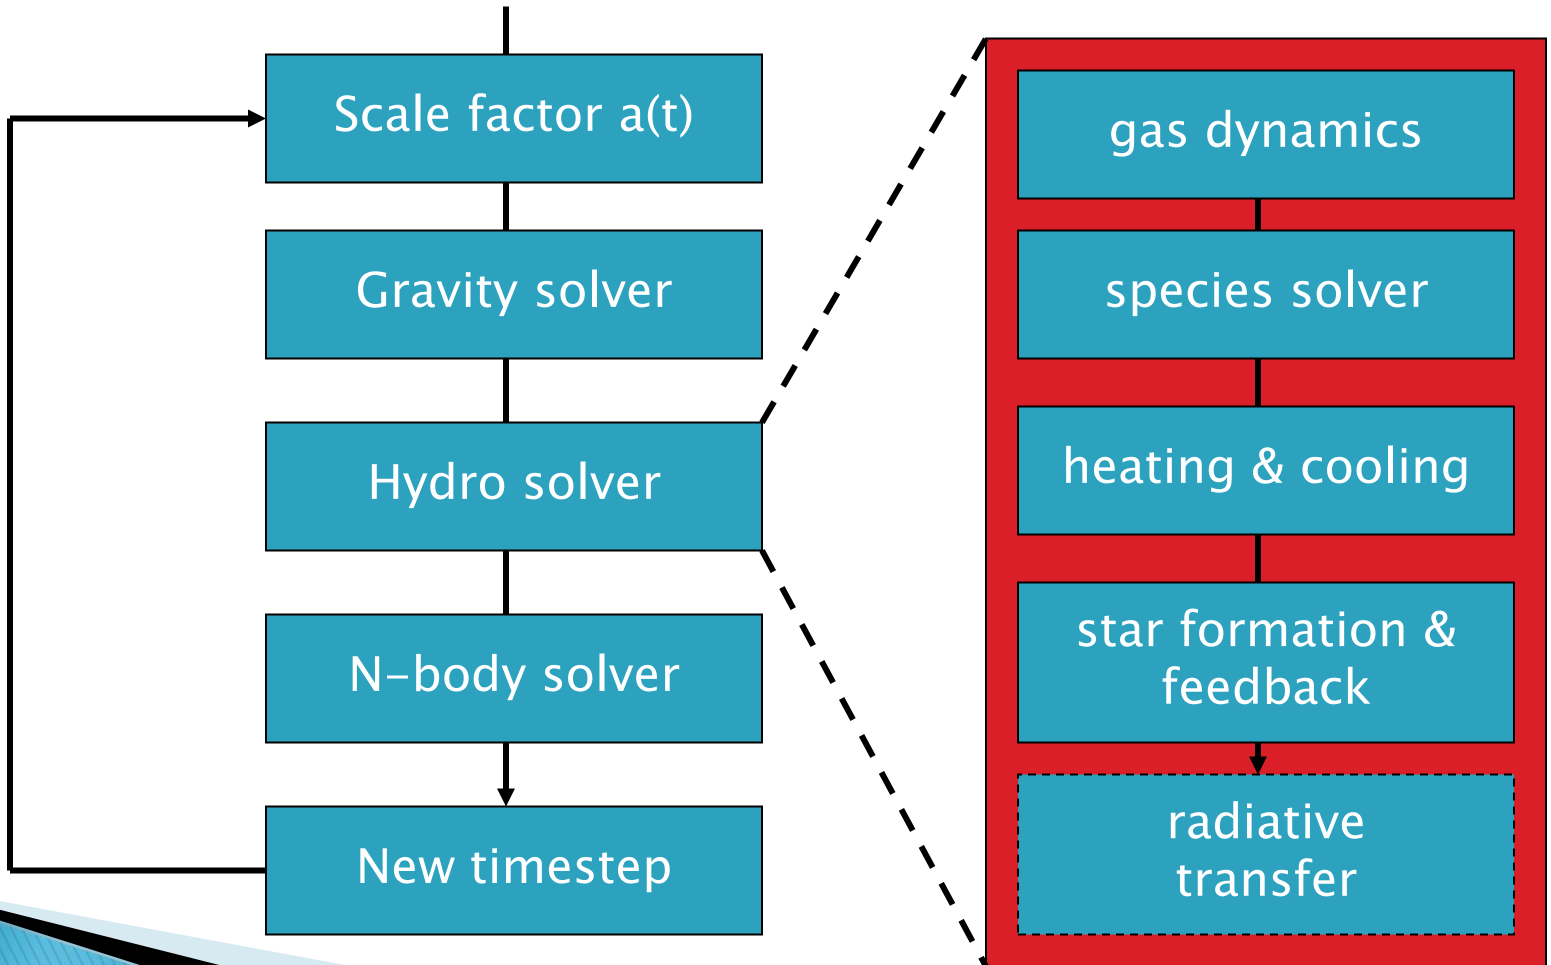

Enzo Dark Matter Solver: AMR-Particle Mesh (Dark matter, stars)

- [1] Mass assignment
  - Mass assigned at every level of the mesh hierarchy with CIC cloud whose size is proportional to Δx(level)
- [2] Field solve
  - *Root grid*: 3D FFT
  - *Subgrid*: local multigrid with boundary potentials interpolated from parent grid, or copied from sibling grids
- [3] Particle push
  - Particle belongs to the finest grid which contains it
  - CIC force interpolation: resolution ~2Δx(level)
  - Leapfrog integration with local timestep Δt(level)

Hydro algorithms requirements posed by cosmological structure formation

- Must be resolution-matched to the dark matter
  - adaptive, multi-resolution methods
- High order-accurate
  - Need to follow perturbation growth accurately from linear to nonlinear regime
  - CDM exhibits power on all scales, including grid scale
- Monotonic
  - Gravity amplifies density perturbations; numerical oscillations disastrous
- Shock capturing
  - virialization of collapsing perturbations involve very strong shocks ($M>100$)

Enzo Euler Solvers (V1.5)

- Piecewise Parabolic Method (Colella & Woodward 1984; Bryan et al. 1995)
  - Comoving frame equations
  - DE, LR formulations
  - Directionally split
  - Dual energy formalism
  - Self-gravity
  - variety of Riemann solvers
- ZEUS hydro (Stone & Norman 1992a)
  - Staggered finite difference
  - Van Leer monotonic advection
  - Artificial viscosity

Baryons: The Tail of the Dog

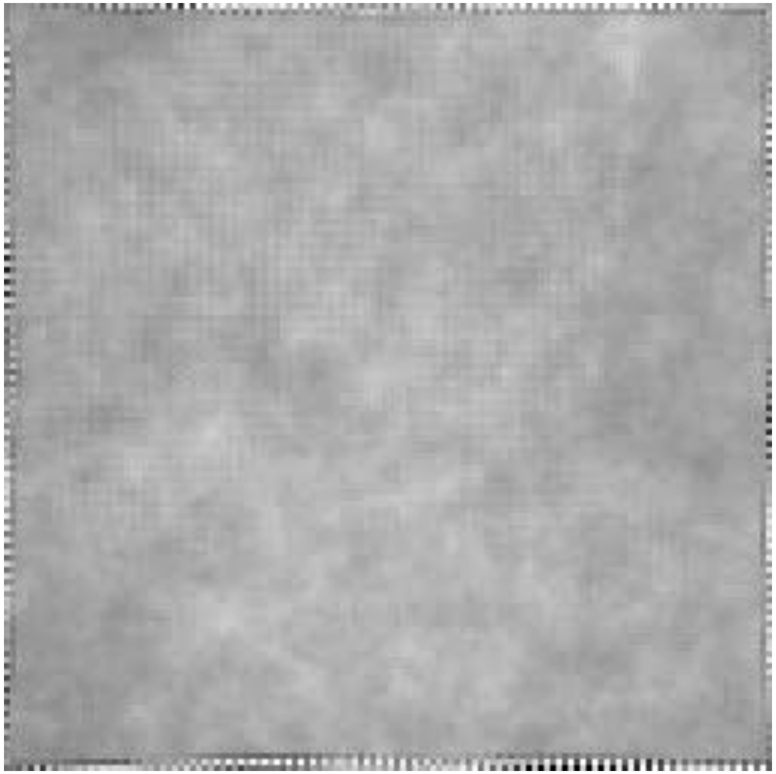

Dark matter

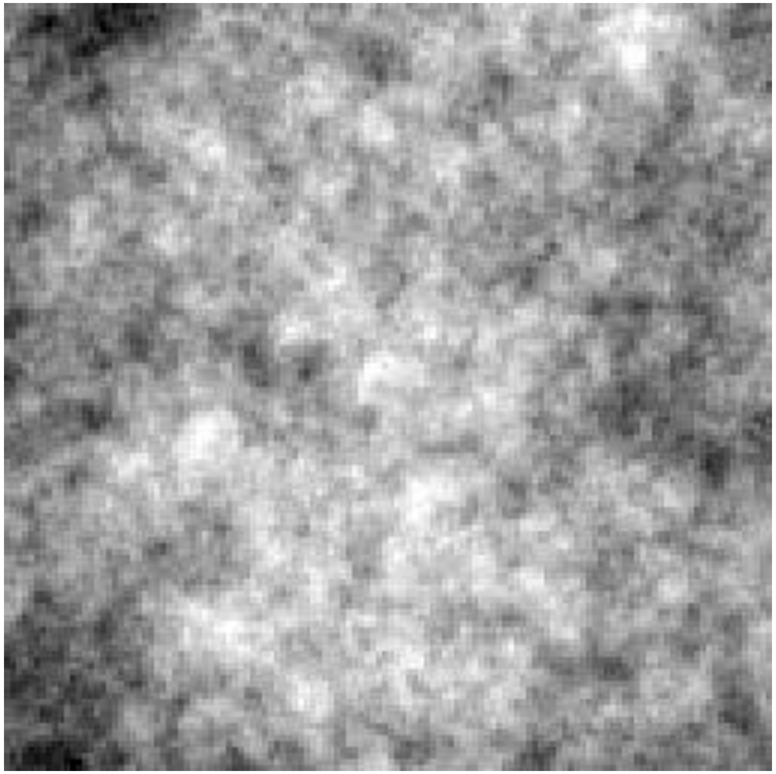

Gas

Just Shocking: Temperature

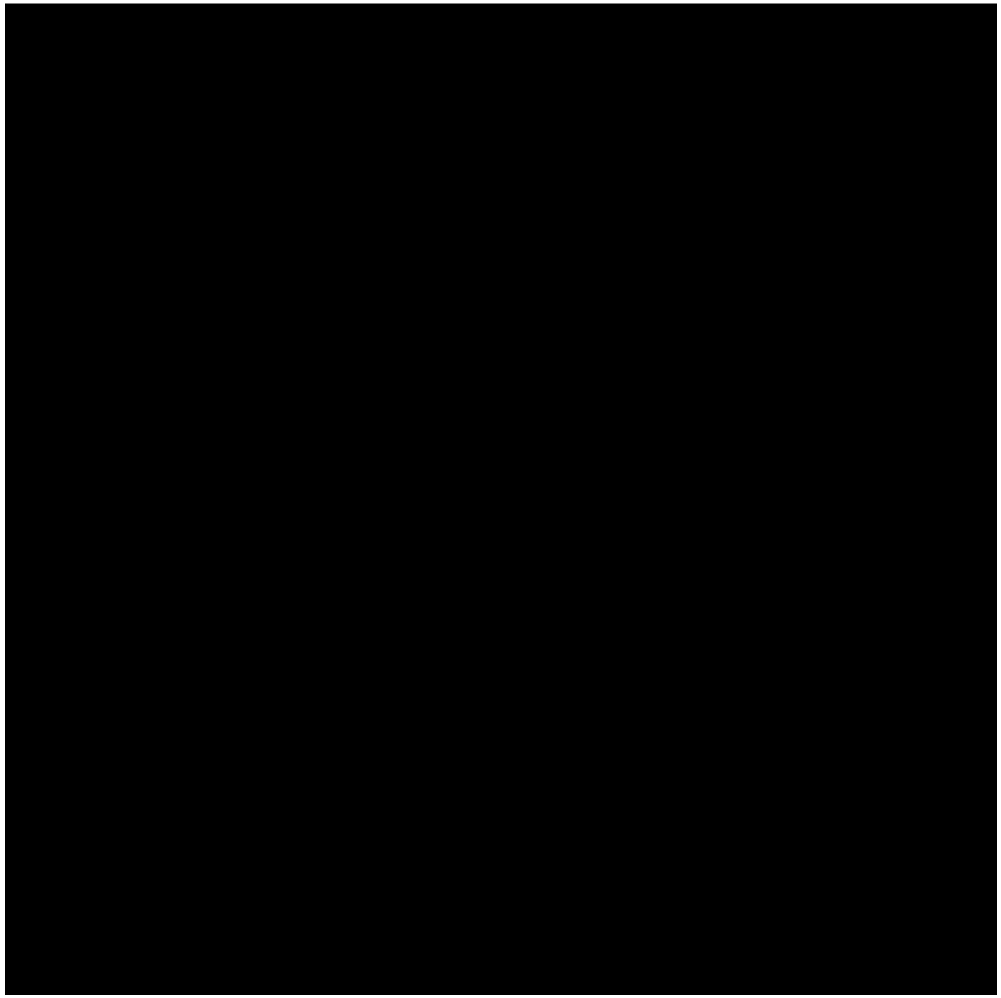

Comparison of Enzo with GADGET: Tree+SPH (O'Shea et al. 2005)

Dark Matter: Tree vs. AMR-PM

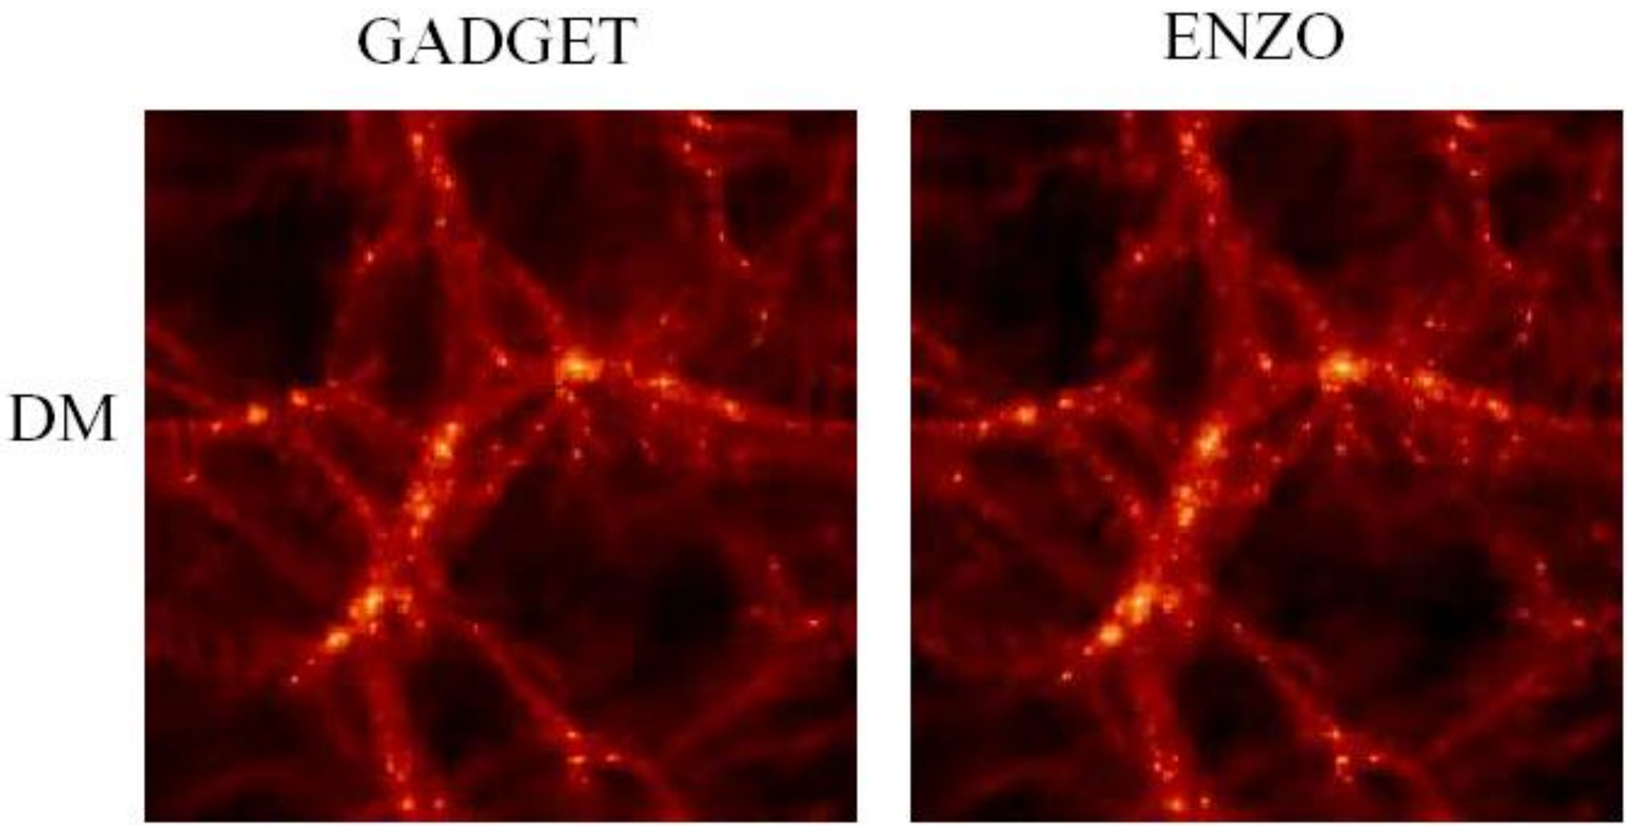

$N = 64^3$

$M = 128^3, N = 64^3$

DM Power spectrum

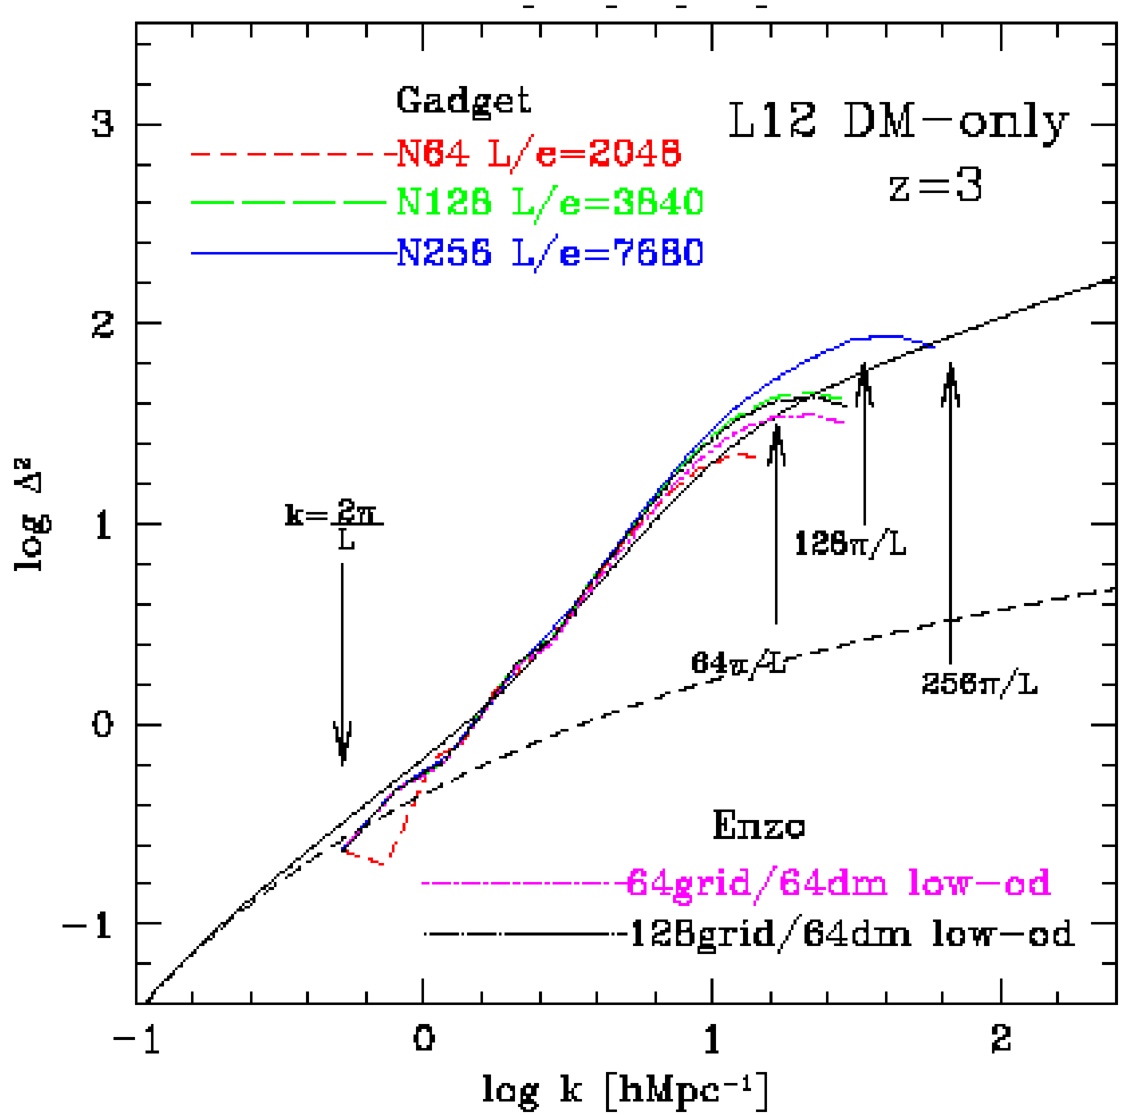

DM Halo mass function

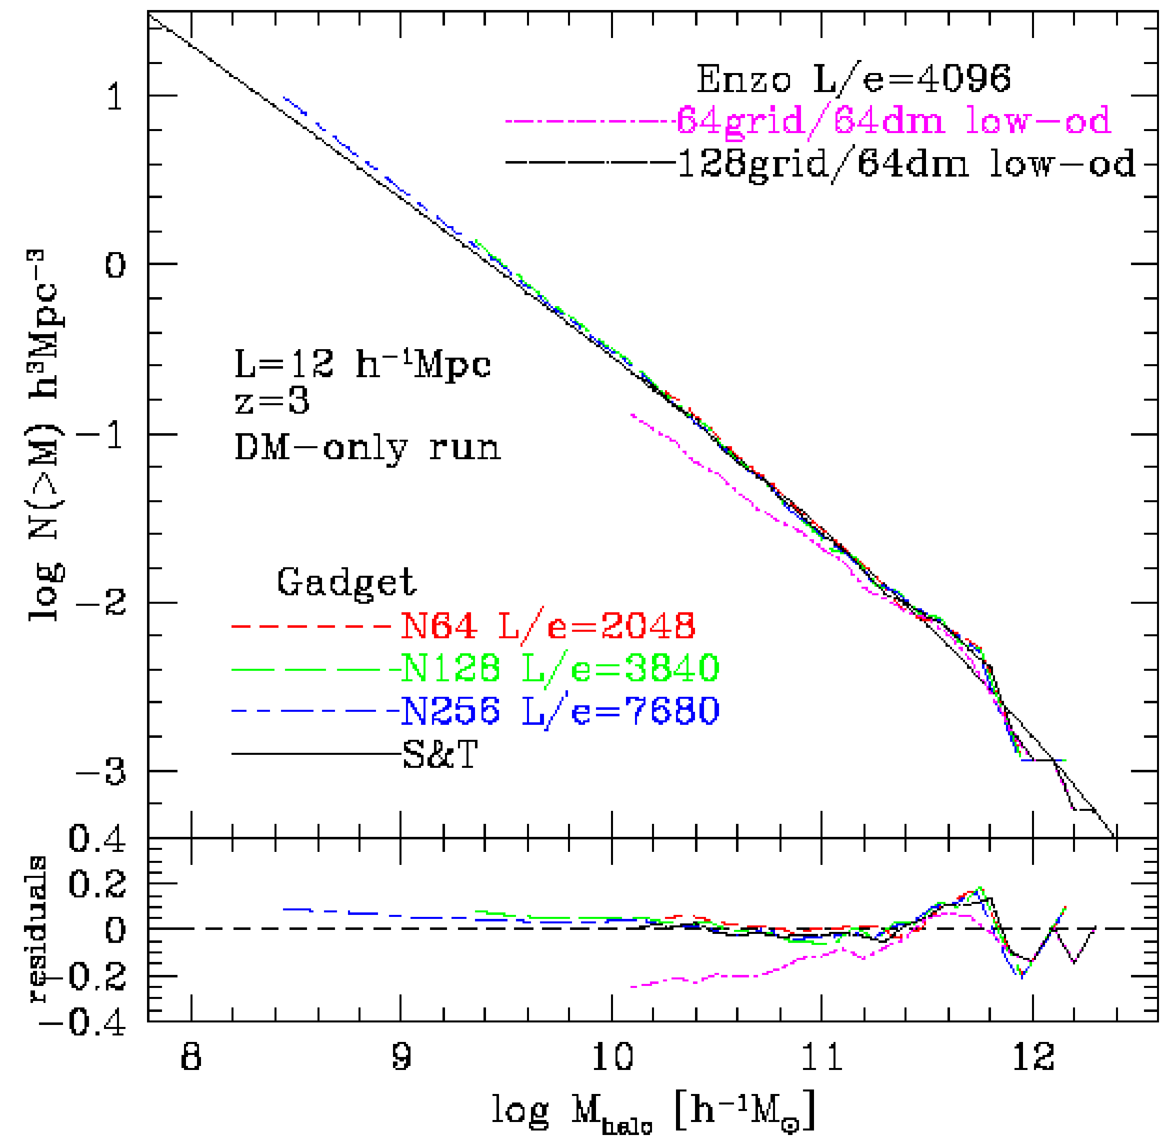

Baryons: SPH vs. AMR-PPM

nonradiative gas “adiabatic”

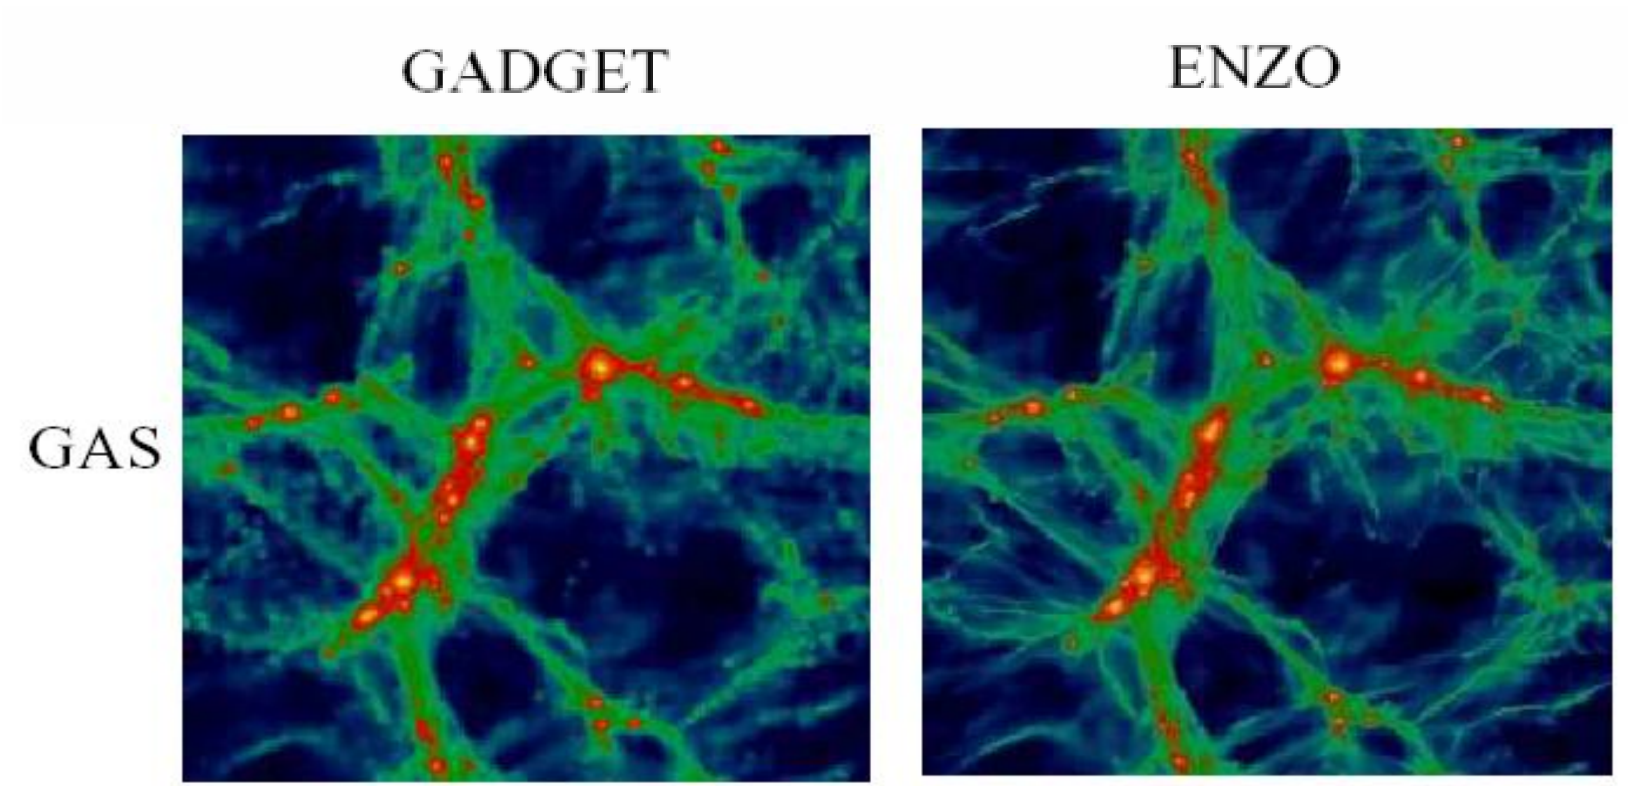

$N = 64^3$

$M = 128^3, N = 64^3$

Halo Baryon Mass Functions

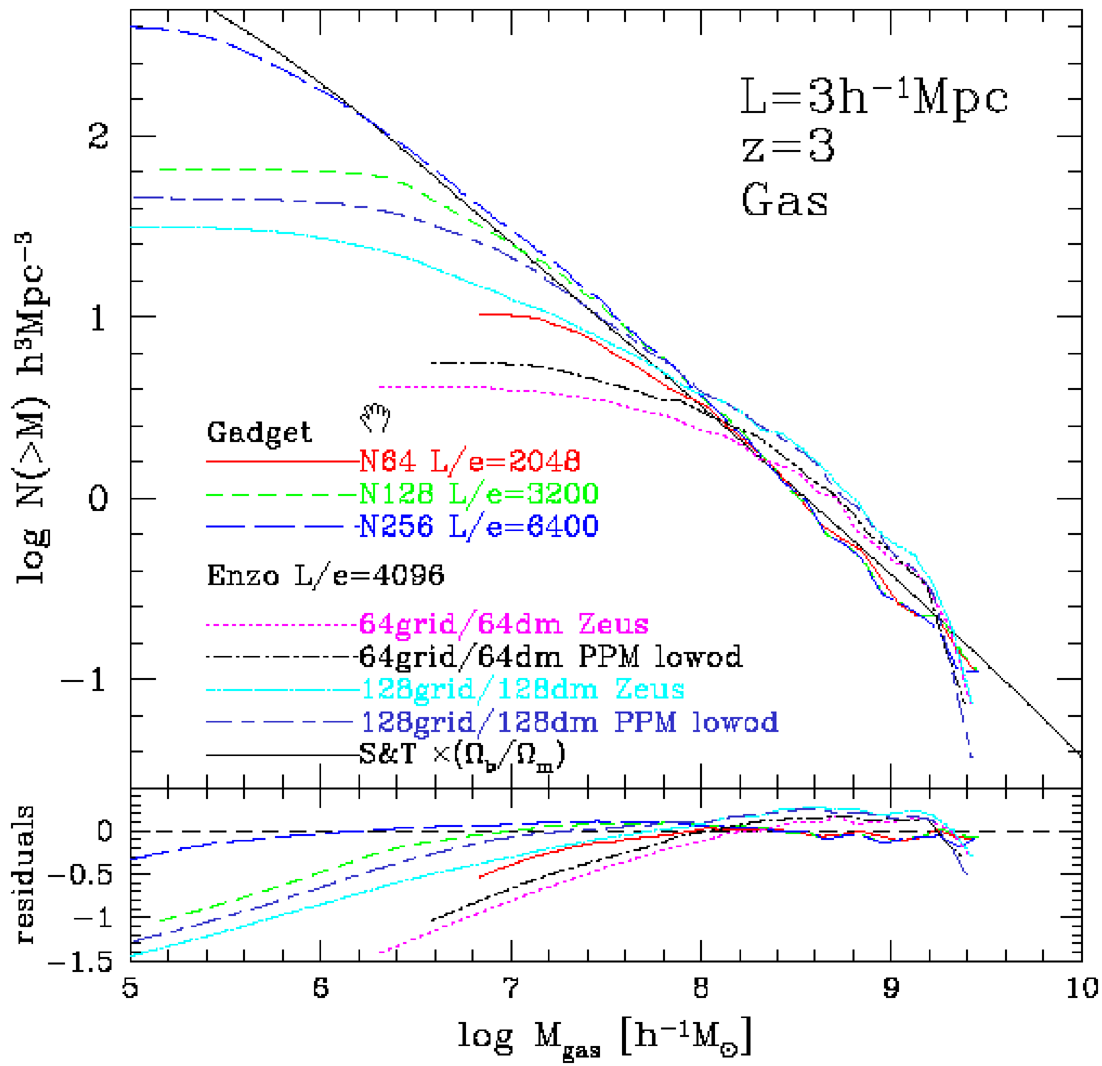

Baryon Fractions in Halos

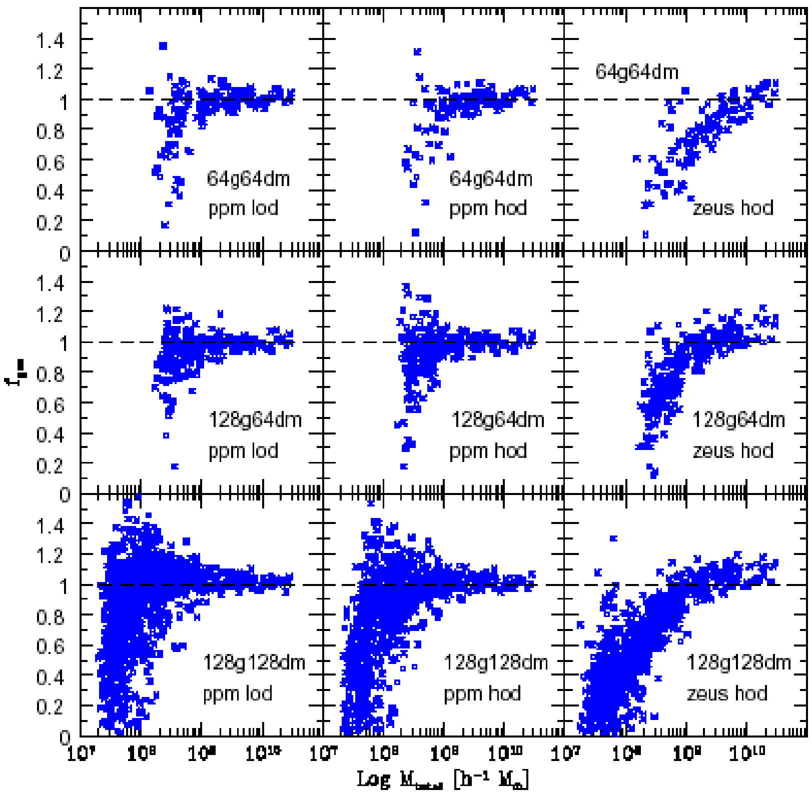

Enzo

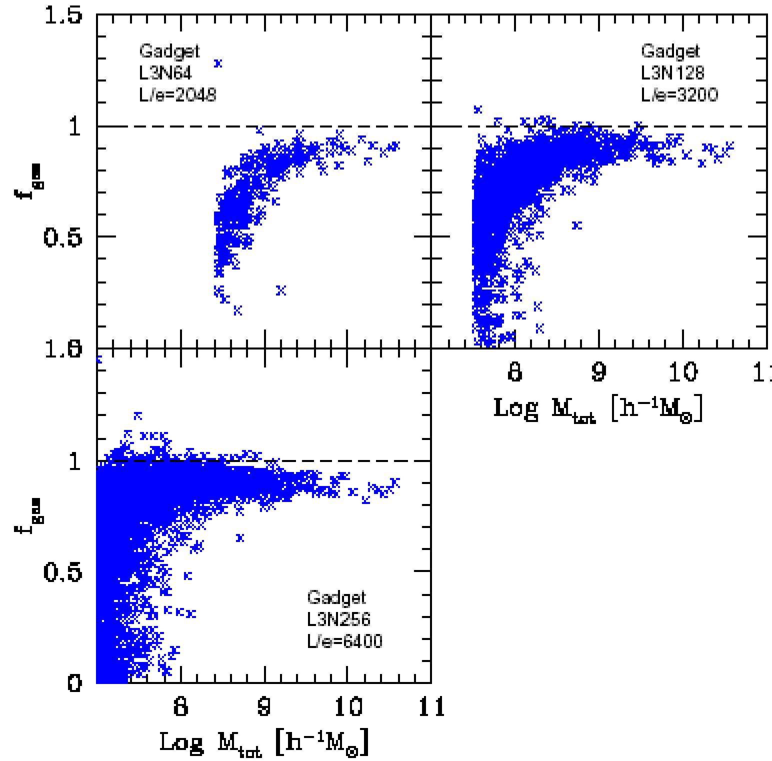

GADGET

Code Comparison: Conclusions

- DM quantities agree over entire range of halo masses if AMR uses low overdensity threshold and 2x finer force grid
  - GADGET gravity solver much faster than Enzo's
- Baryonic properties agree somewhat less well; sensitive to resolution and method
  - Preheating in AV methods modifies entropy distributions, baryon fractions
  - PPM asymptotes to ~ 1.0 cosmic mean
  - SPH asymptotes to ~ 0.9 cosmic mean

Enzo Additional Physics (V1.5)

- multi-species Euler solver for cosmic gas
- nonequilibrium primordial chemistry solver for H, H^+, He, He^+, He^{++}, e^-, H^-, H_2^+, H_2, D, D^-, DH
- Various UV and X radiation backgrounds
- radiative heating and cooling (inverse Compton, line and continuum, etc.)
- variety of parameterized star formation & feedback recipes
- metallicity fields
- tracer particles

Parallelism and Scaling

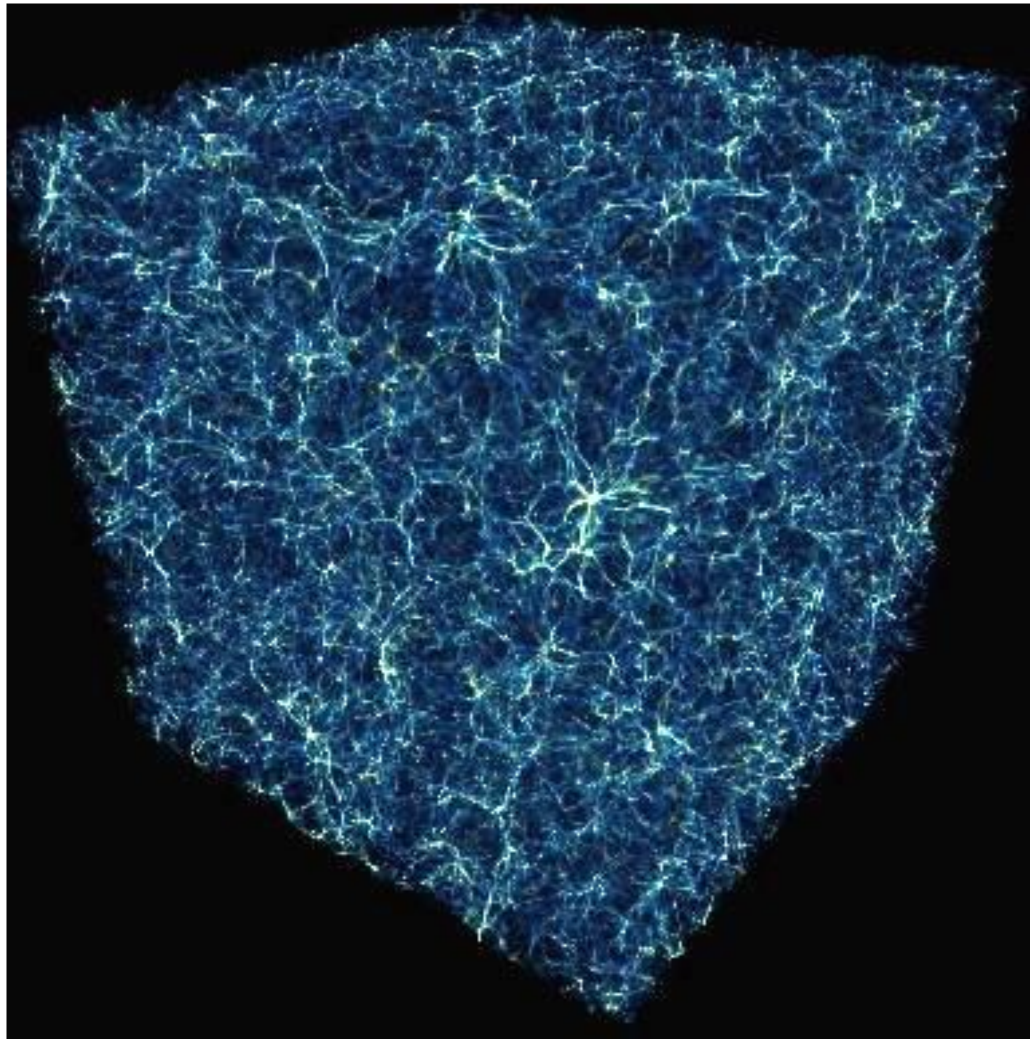

512^3, 7-level
256 cpus
>300,000 grids

Enzo Usage Modes

- Unigrid=uniform, nonadaptive Cartesian
  - Largest to date: 6400^3
- AMR everywhere
  - Largest to date: 1024^3, 7-level
- AMR “zoom in”
  - Largest to date: 1024^3 effective, 42-level

AMR = collection of grids (patches); each grid is a C++ object

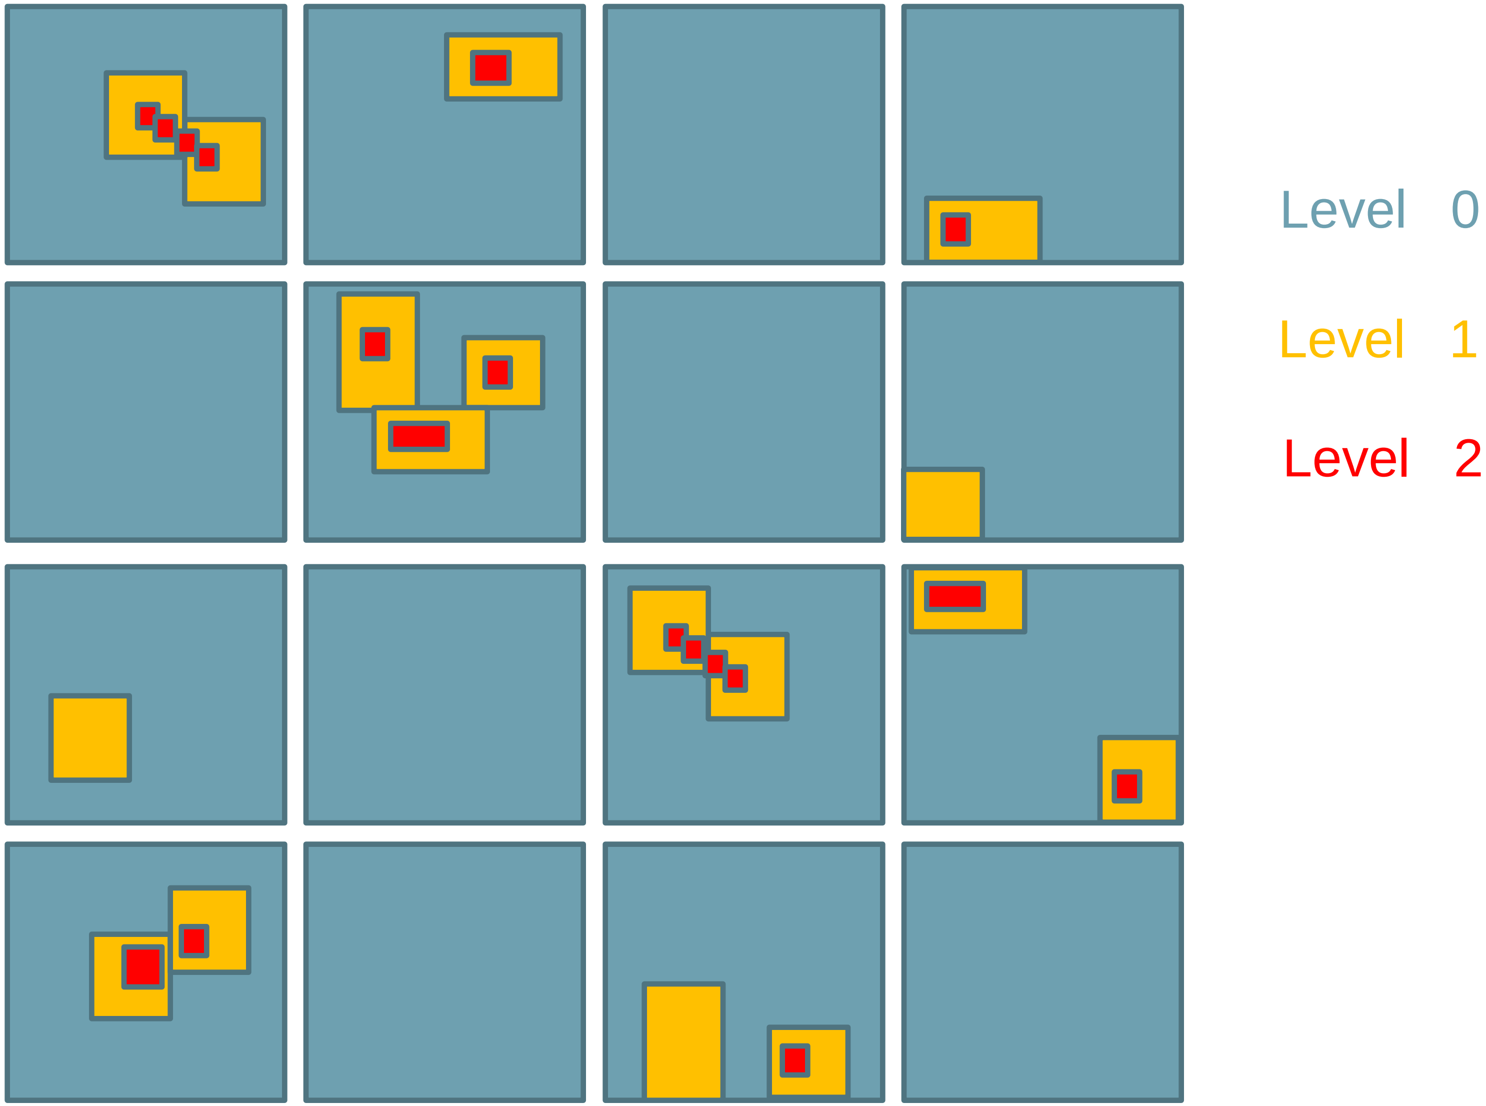

Unigrid = collection of Level 0 grid patches

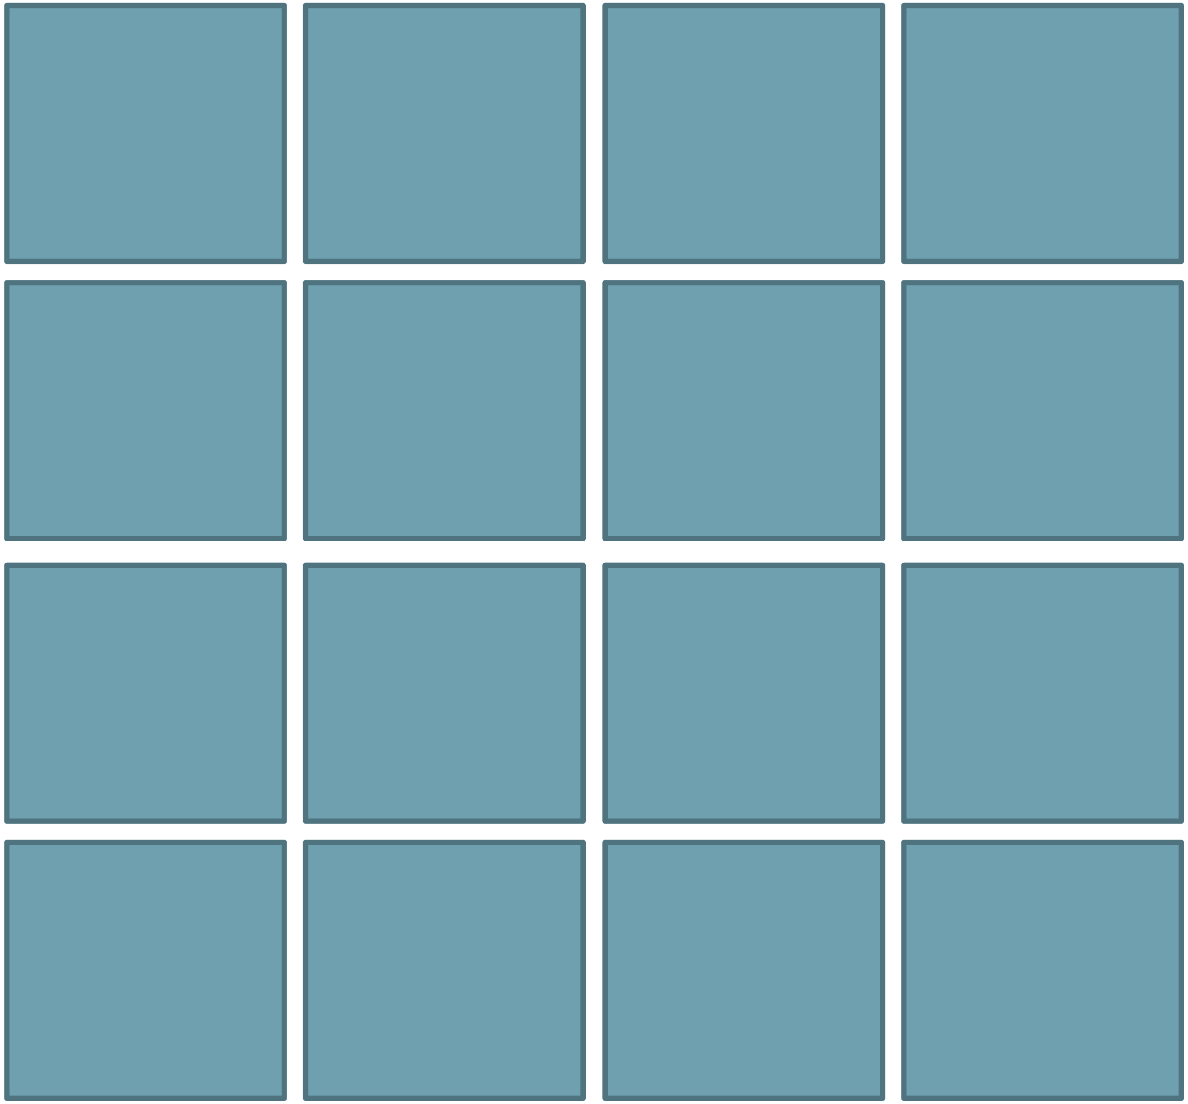

Evolution of Enzo Parallelism

- Shared memory (PowerC) parallel (1994–1998)
  - SMP and DSM architecture (SGI Origin 2000, Altix)
  - Parallel DO across grids at a given refinement level including block decomposed base grid
  - O(10,000) grids
- Distributed memory (MPI) parallel (1998–2008)
  - MPP and SMP cluster architectures (e.g., IBM PowerN)
  - Level 0 grid partitioned across processors
  - Level >0 grids within a processor executed sequentially
  - Dynamic load balancing by messaging grids to underloaded processors (greedy load balancing)
  - O(100,000) grids

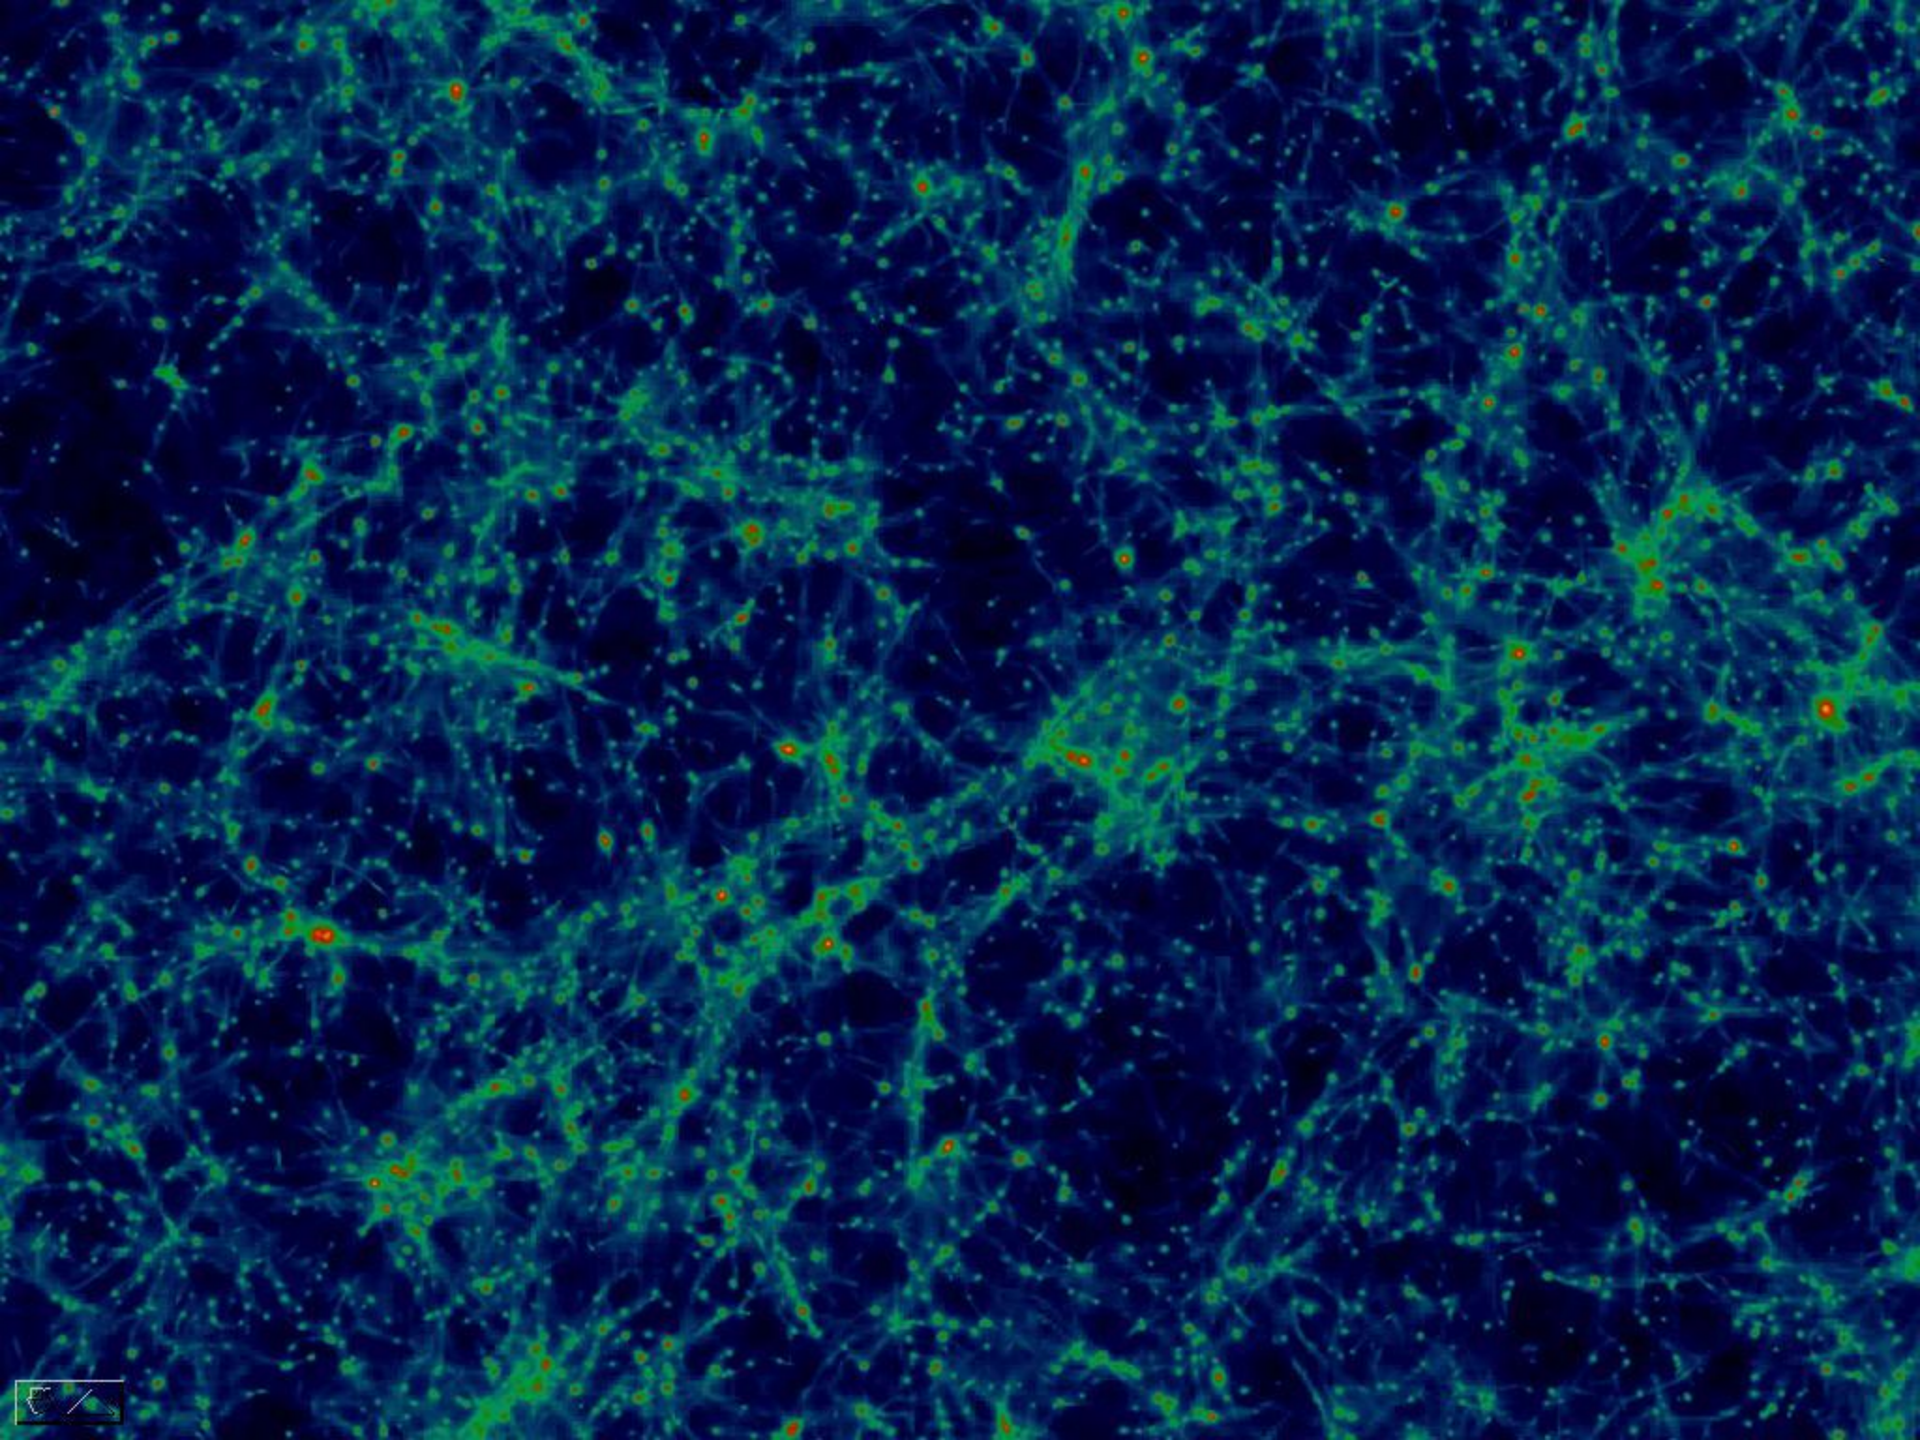

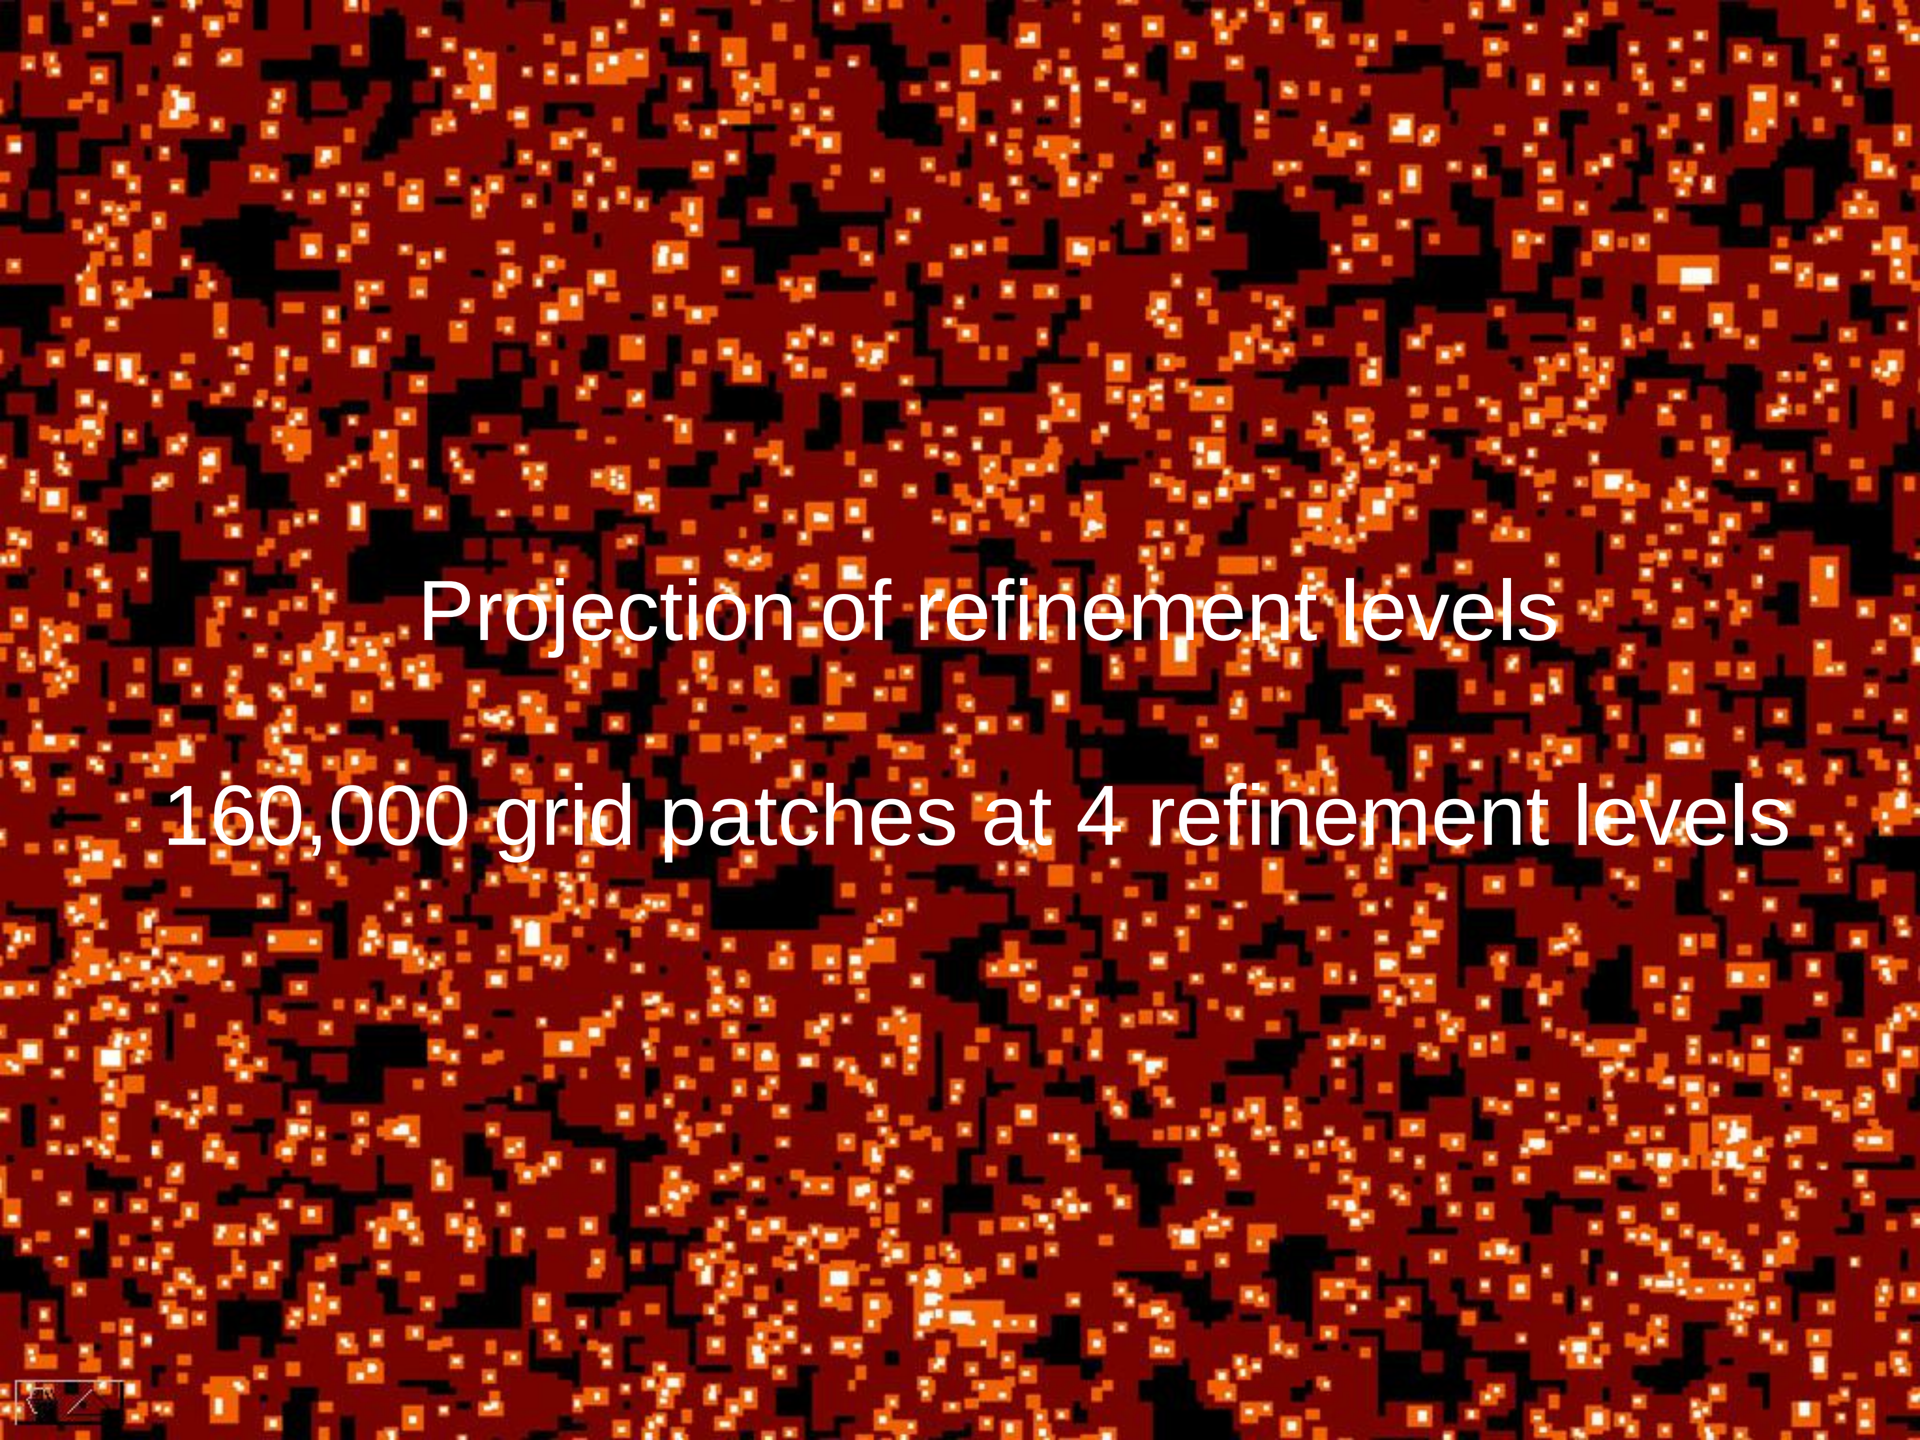
Projection of refinement levels
160,000 grid patches at 4 refinement levels

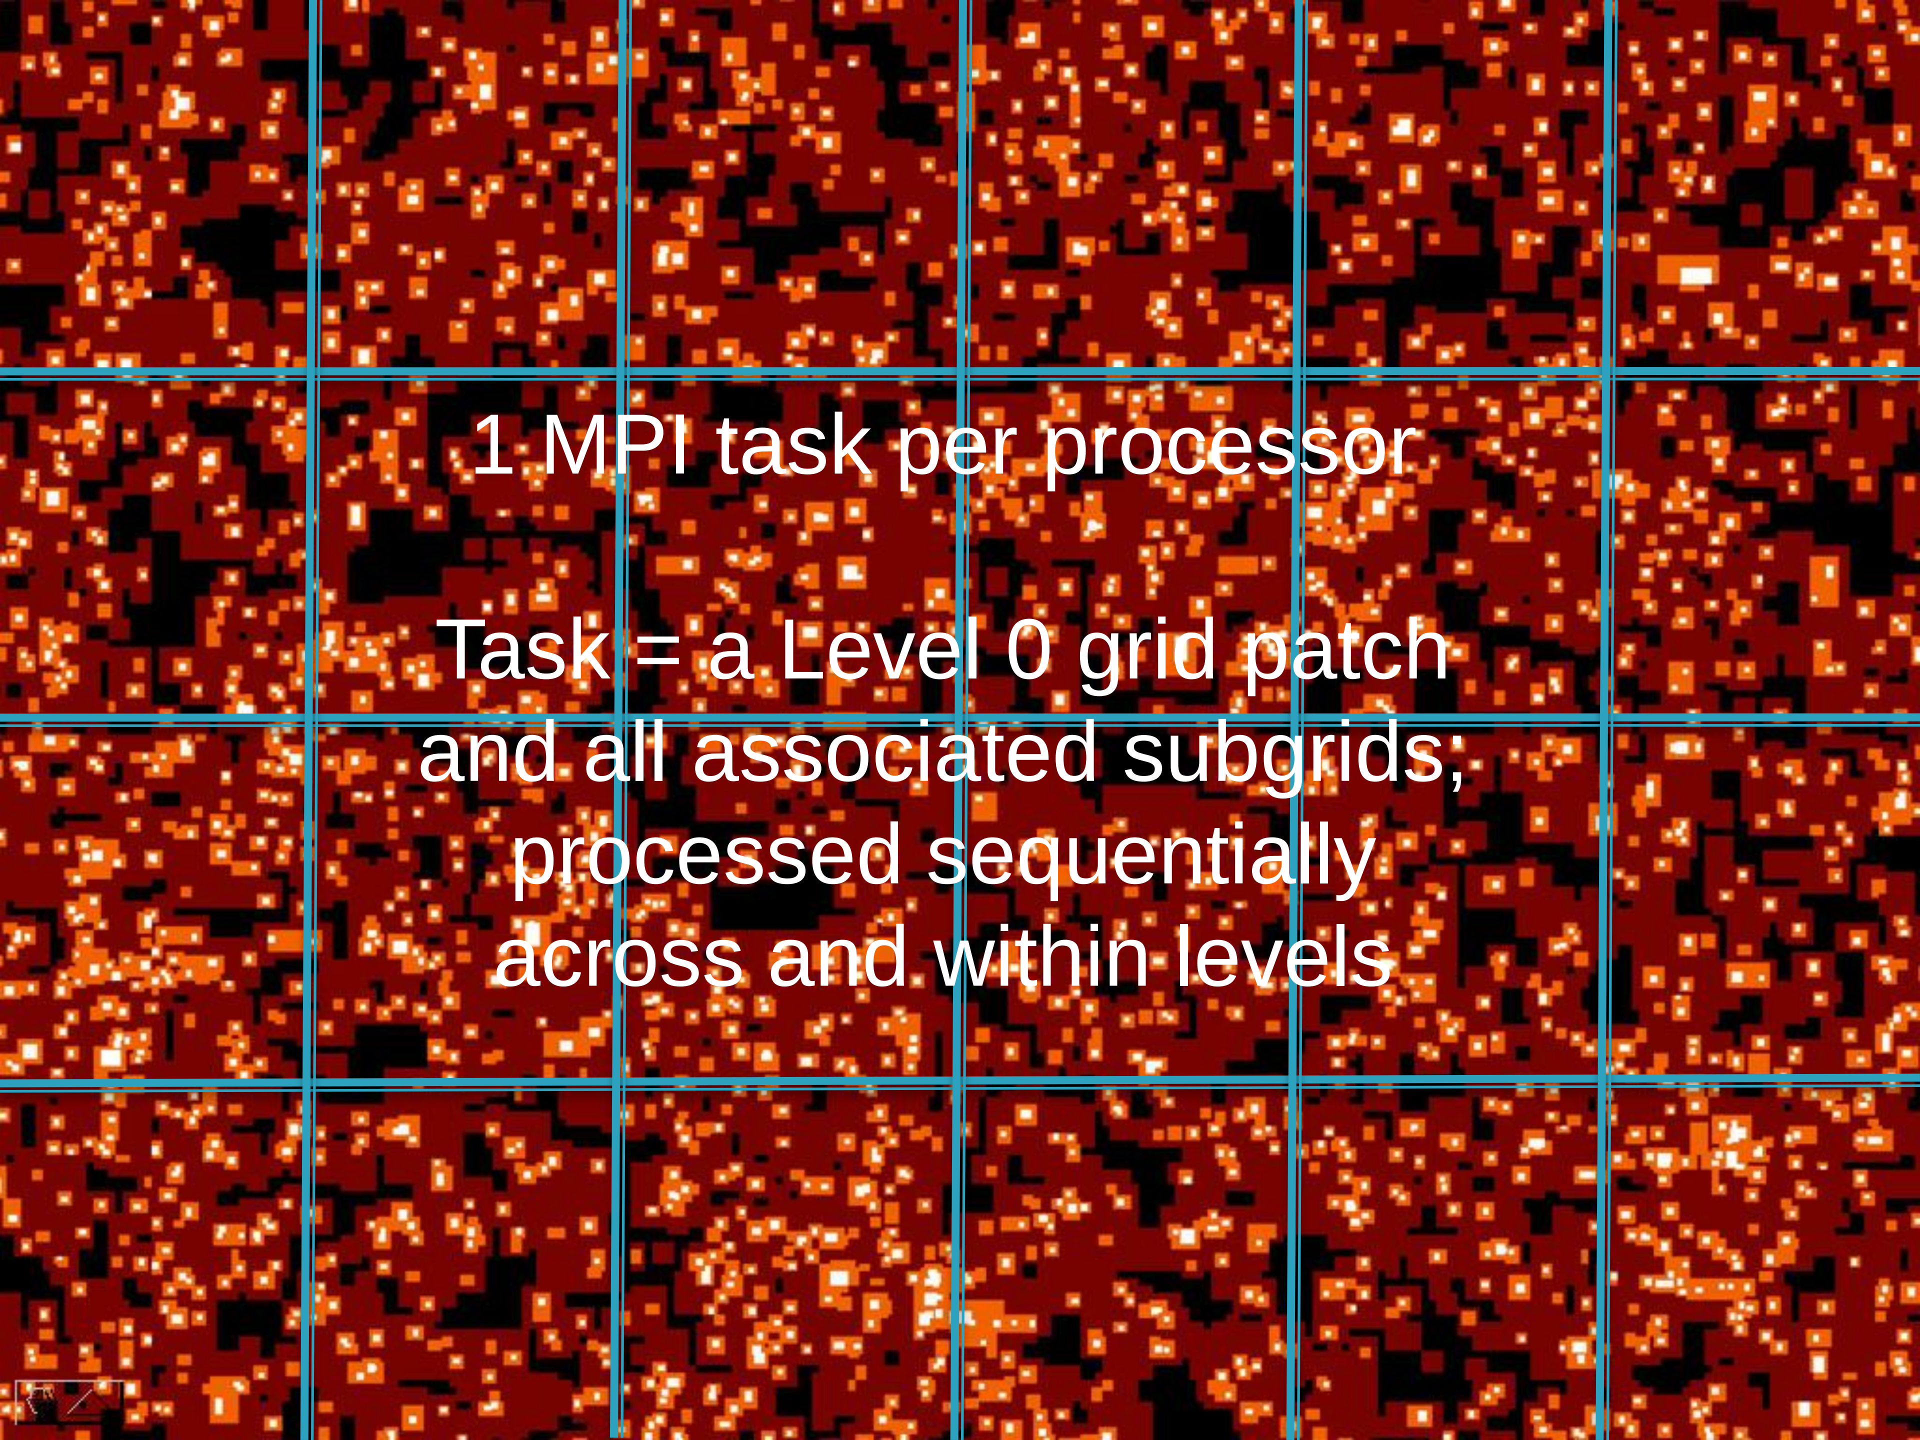
1 MPI task per processor
Task = a Level 0 grid patch
and all associated subgrids;
processed sequentially
across and within levels

Evolution of Enzo Parallelism

- Hierarchical memory (MPI+OpenMP) parallel (2008-)
  - SMP and multicore cluster architectures (SUN Constellation, Cray XT4)
  - Level 0 grid partitioned across shared memory nodes/multicore processors
  - Parallel DO across grids at a given refinement level within a node
  - Dynamic load balancing less critical because of larger MPI task granularity (statistical load balancing)
  - O(1,000,000) grids

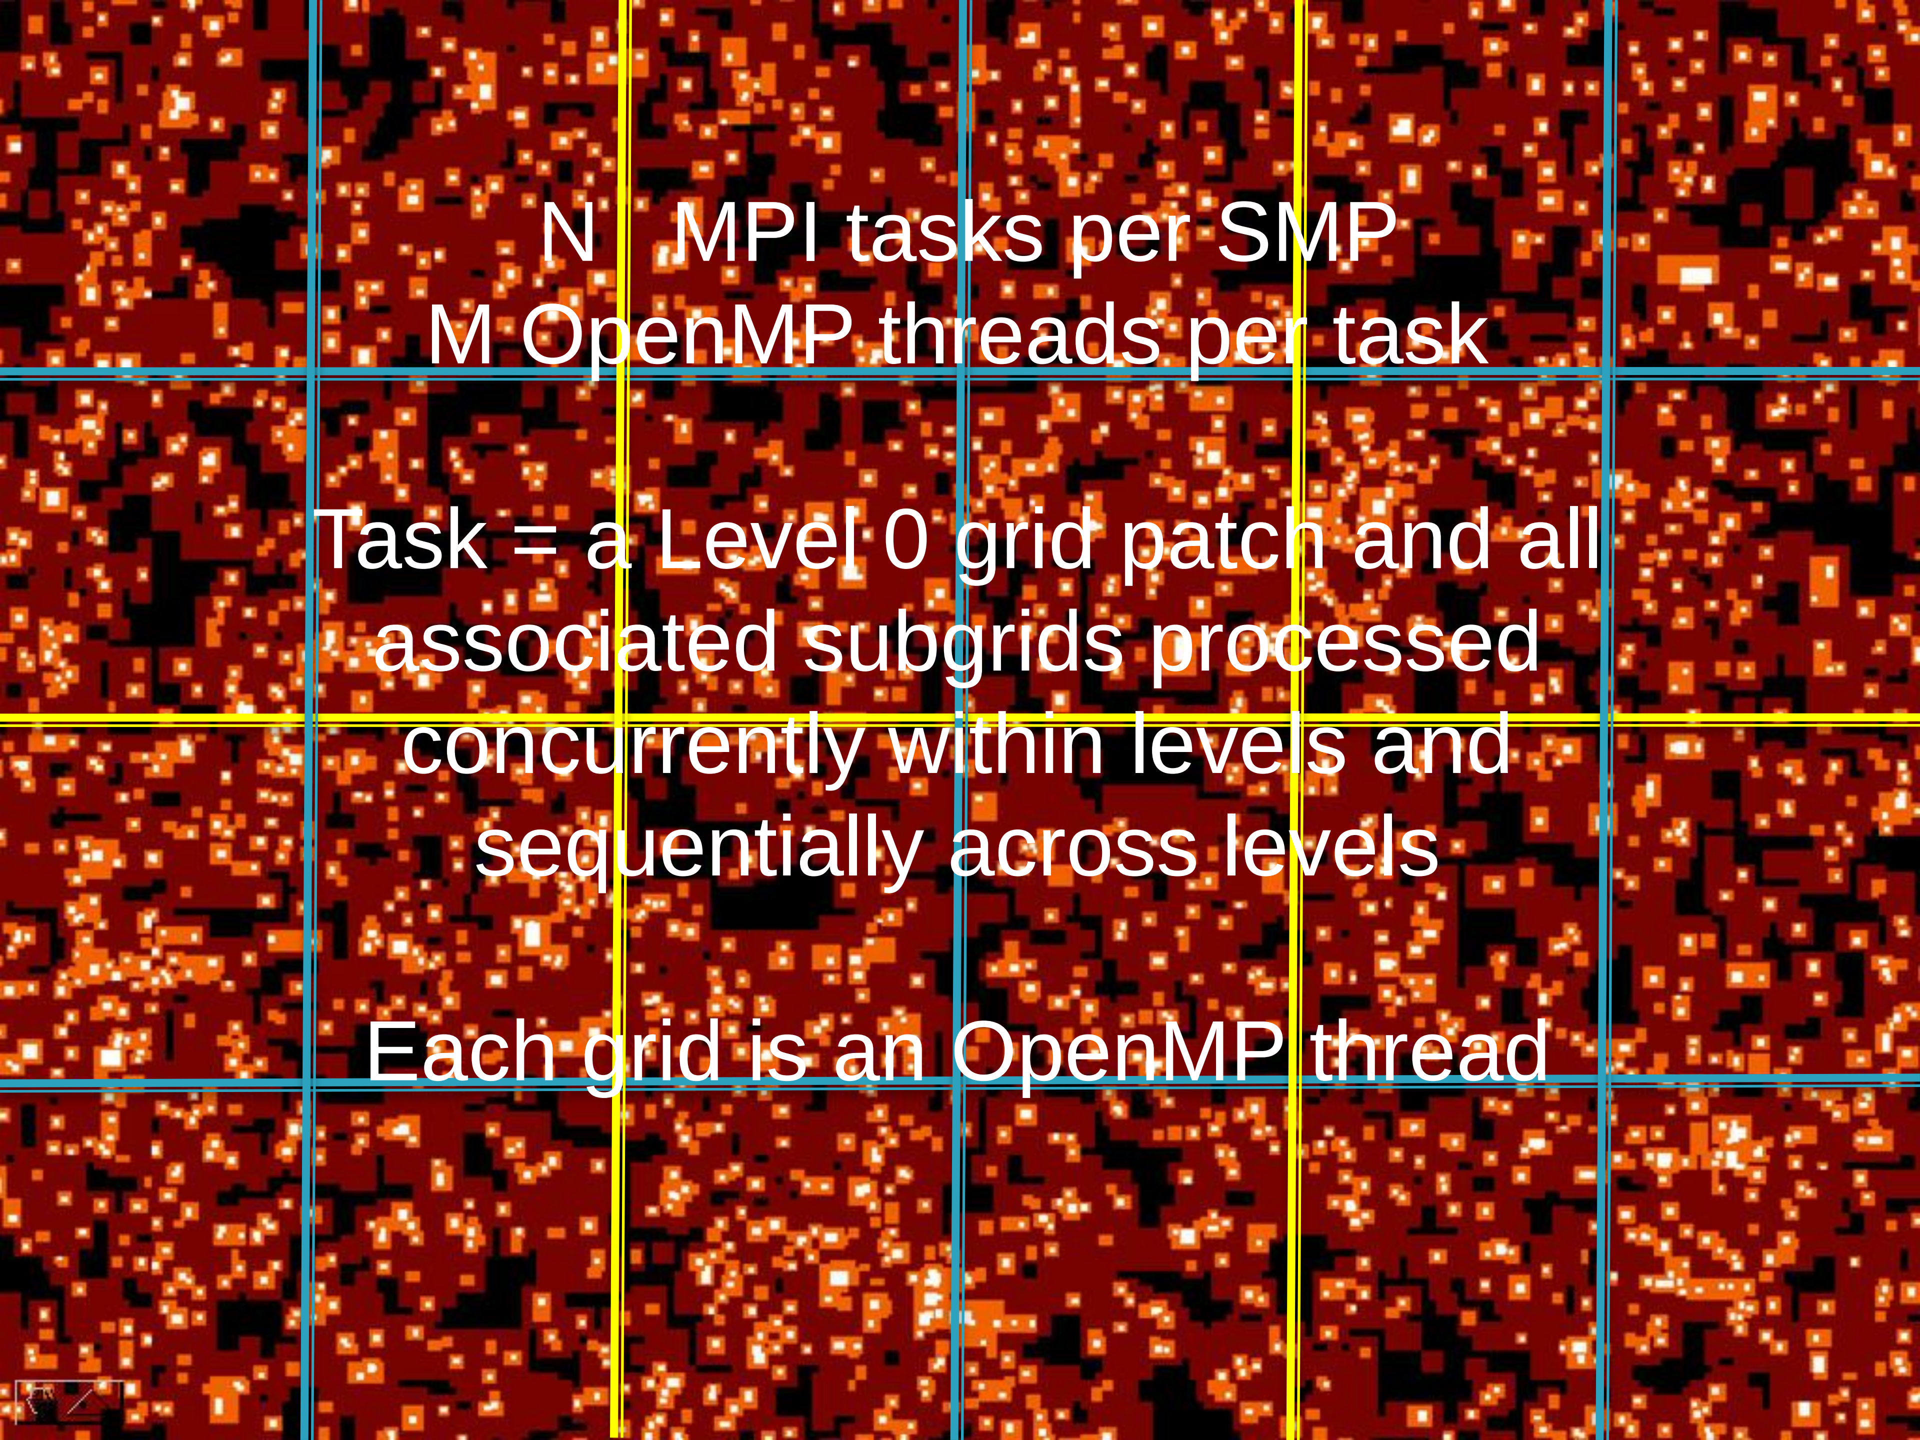
N MPI tasks per SMP
M OpenMP threads per task
Task = a Level 0 grid patch and all associated subgrids processed concurrently within levels and sequentially across levels
Each grid is an OpenMP thread

Scaling the Base Grid Size

- Enzo was "born" as an AMR code
- Level 0 grids were treated no differently than other grids
  - Size, shape, and location could be arbitrary
  - Sibling search procedure determined adjacency
- Optimizations
  - Set level 0 topology map using *MPI_Dims_Create* for optimal interprocessor communication
    - Huge speedup for 3D FFT, hydro/MHD
    - Beneficial for moving particles between processors

Other optimizations

- 64 bit integers for addressing > 2 billion particles
- Asynchronous message-passing and I/O buffering
  - Irecv/Isend
- Embarrassingly parallel I/O
  - each processor reads only the data it owns
  - parallel particle sort by Level 0 grid patch
- Static task to node mapping for particle load balancing
  - For machines with high- and low-memory nodes
- Scalable external boundary conditions
  - Each processor stores only its piece of the boundary condition
- Inline analysis and visualization
  - Power spectra, structure functions, projections

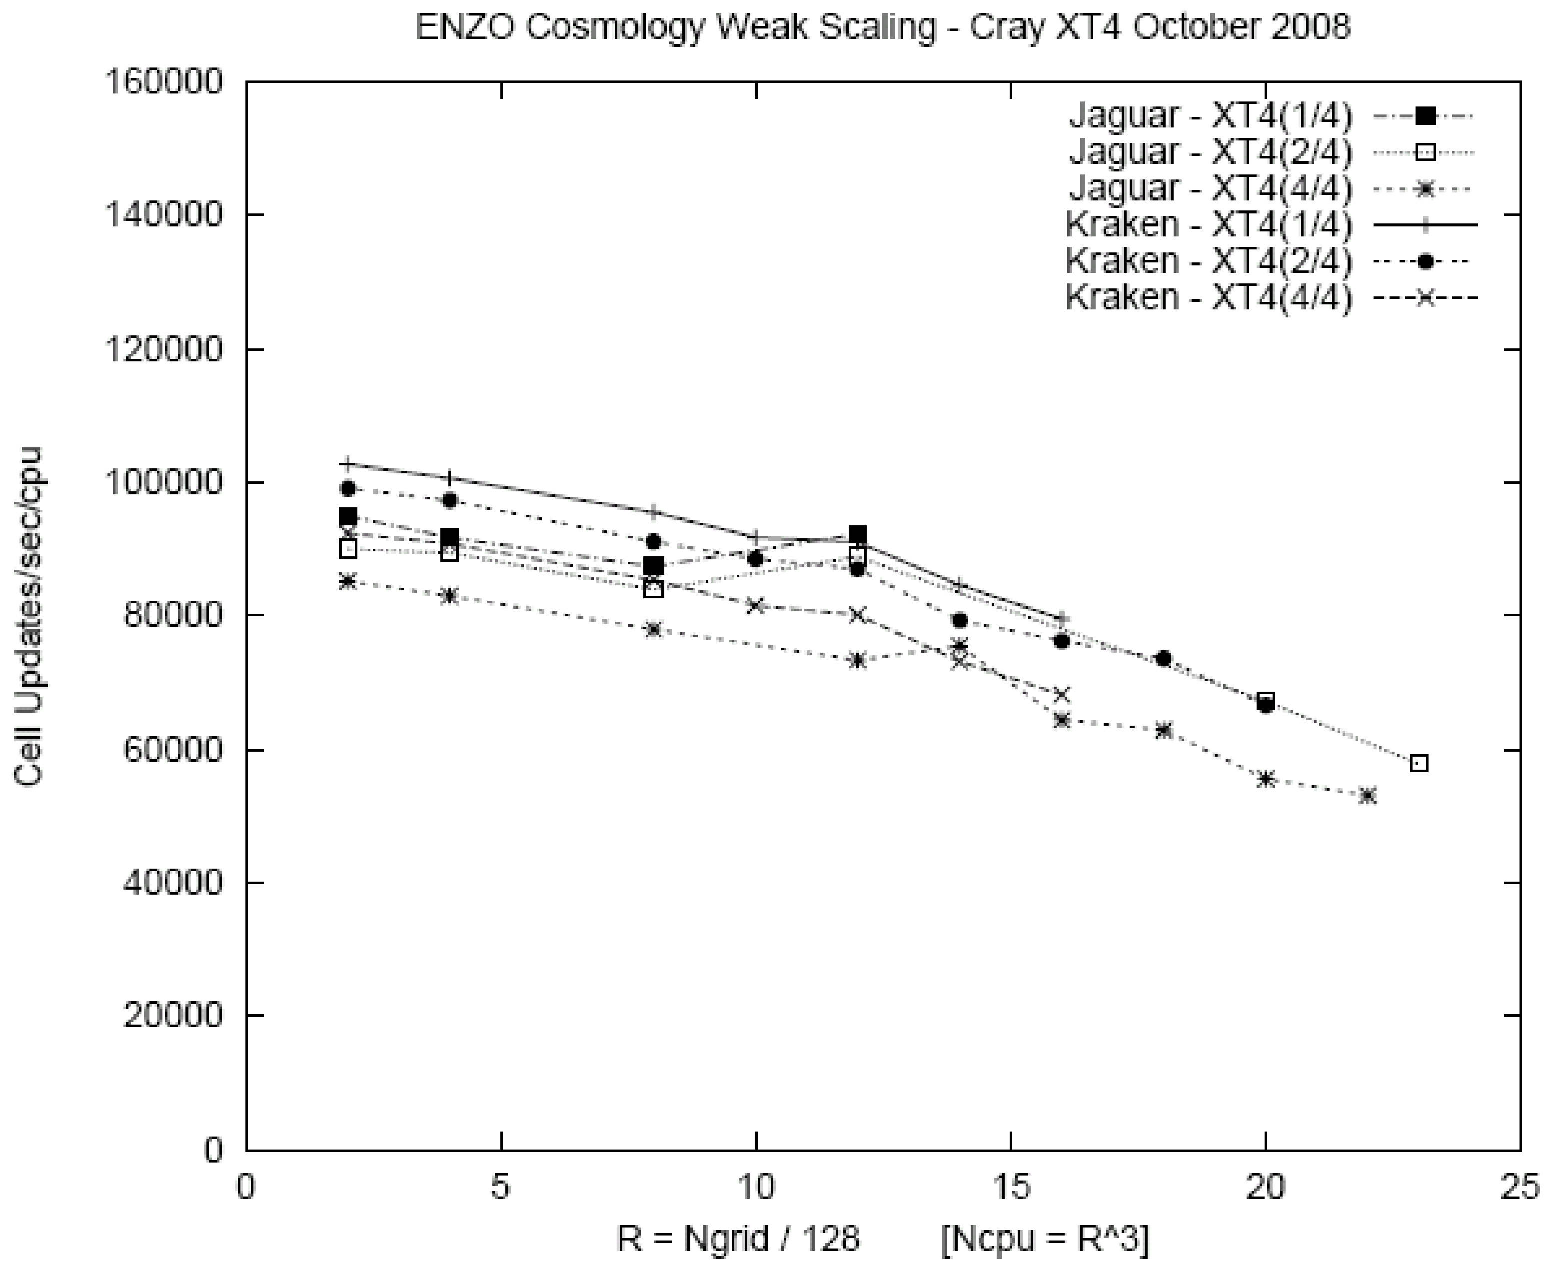
ENZO Cosmology Weak Scaling - Cray XT4 October 2008
Jaguar - XT4(1/4)
Jaguar - XT4(2/4)
Jaguar - XT4(4/4)
Kraken - XT4(1/4)
Kraken - XT4(2/4)
Kraken - XT4(4/4)
Cell Updates/sec/cpu
Cell Updates/sec/task
R = Ngrid / 128 [Ncpu = R^3]
160000
140000
120000
100000
80000
60000
40000
20000
0
0
5
10
15
20
25

A Huge Unigrid: 6400^3 *Enzo*

M. Norman, R. Harkness

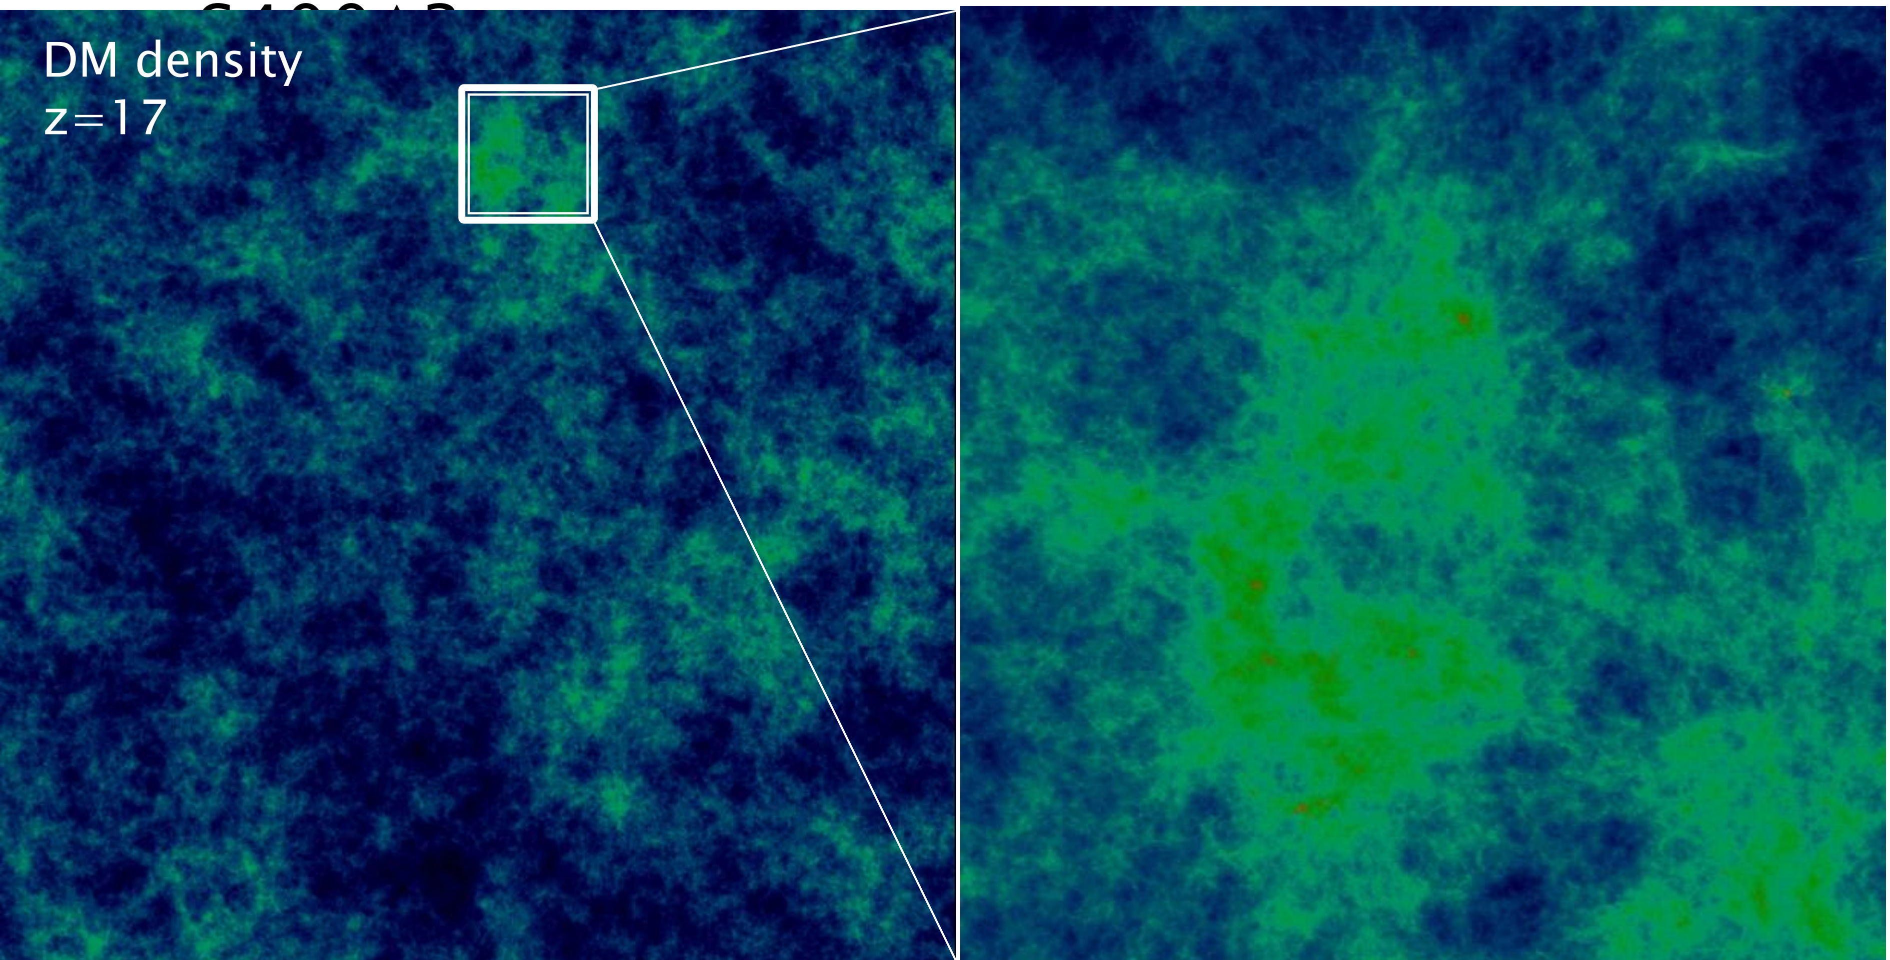

6400^3 cells/particles, 80 Mpc box, DM+Gas+SF/FB

93,000 cores

AMR Scaling Results

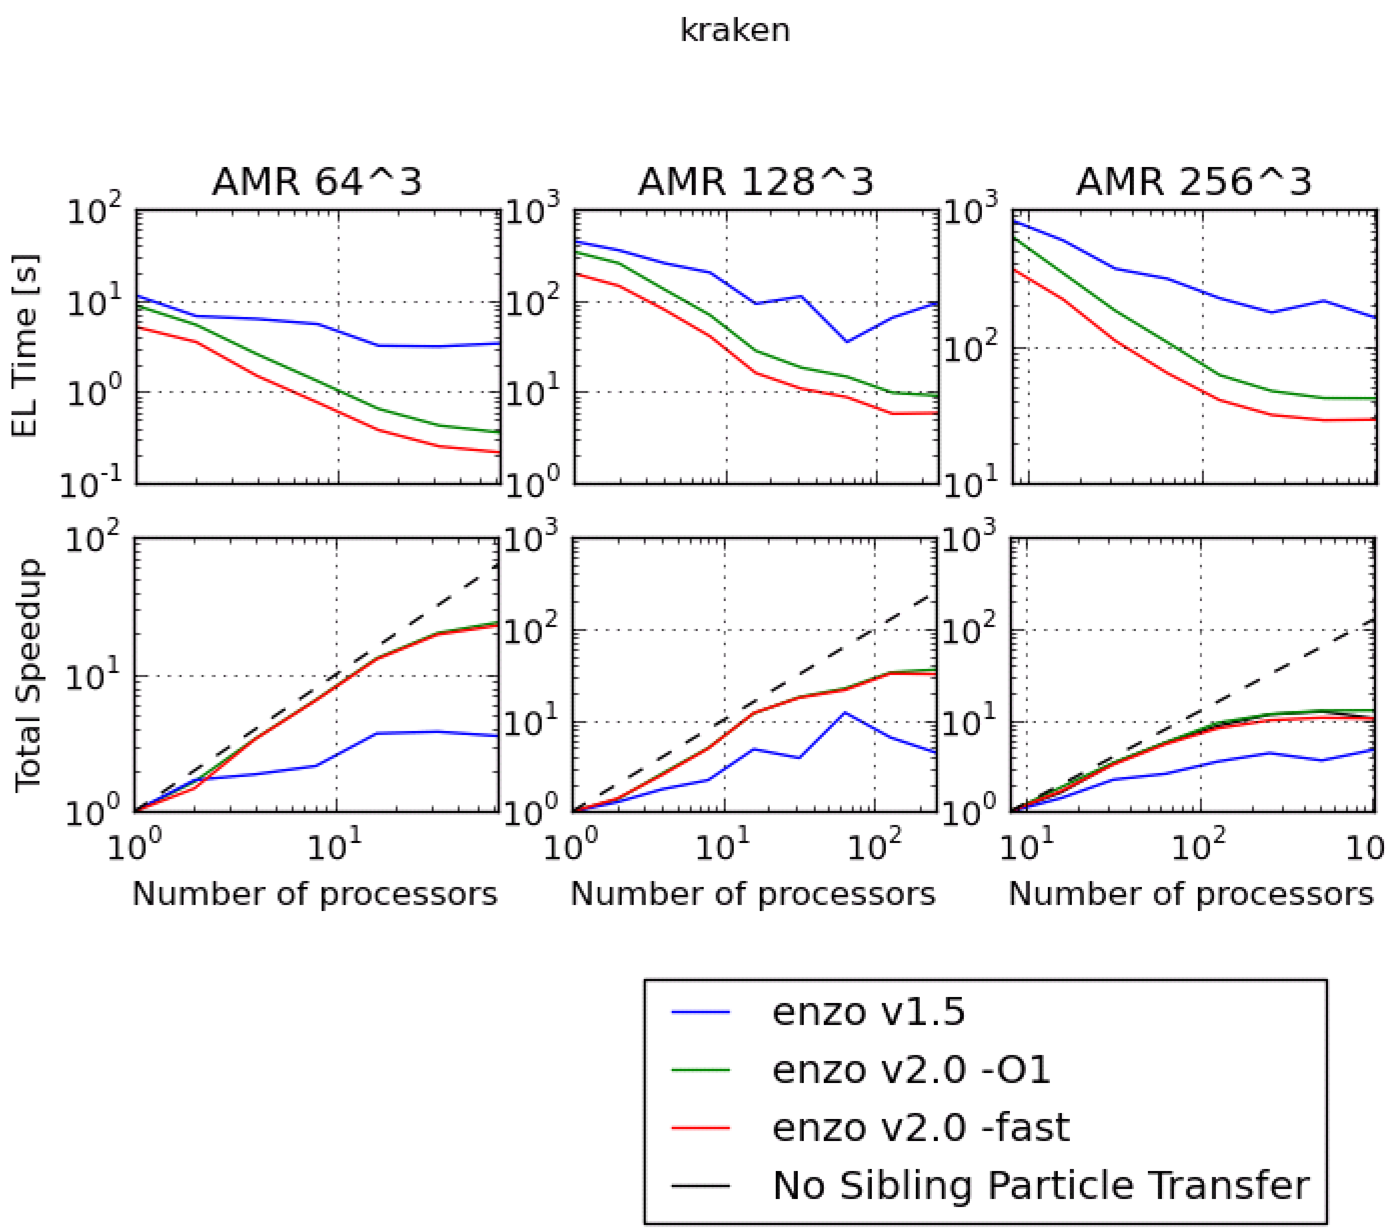

Application sampler

Simulating X-ray Clusters
Loken et al. 2002, Motl et al. 2004, Hallman et al. 2005, 2006

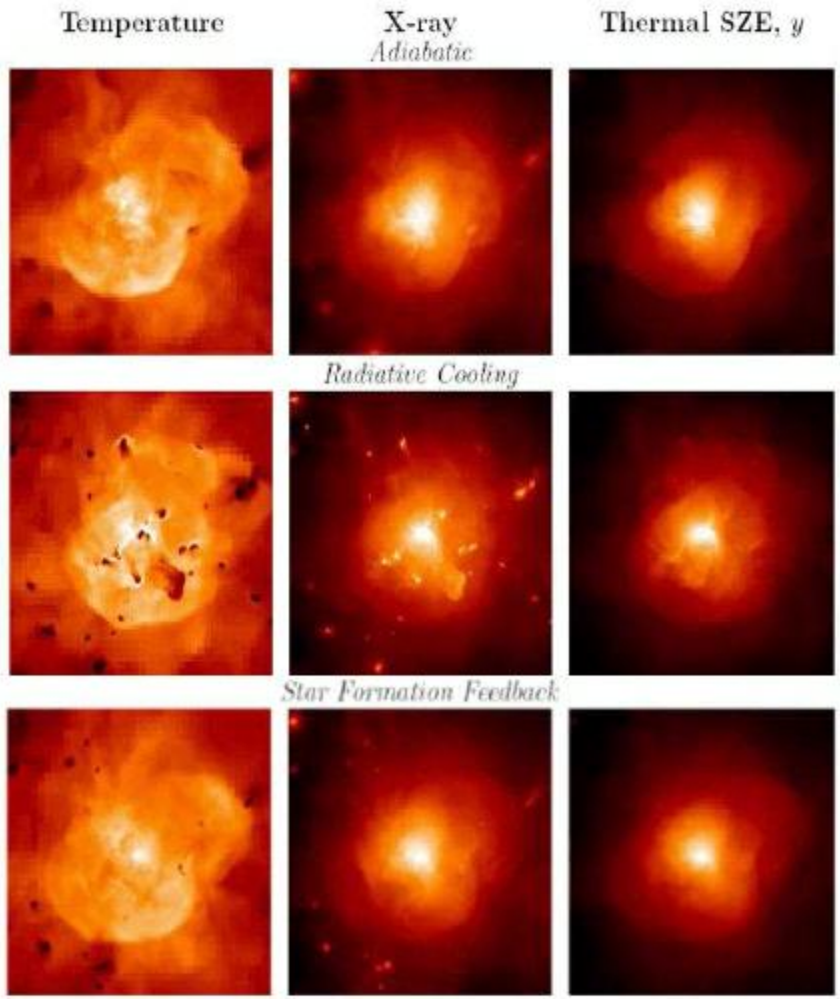

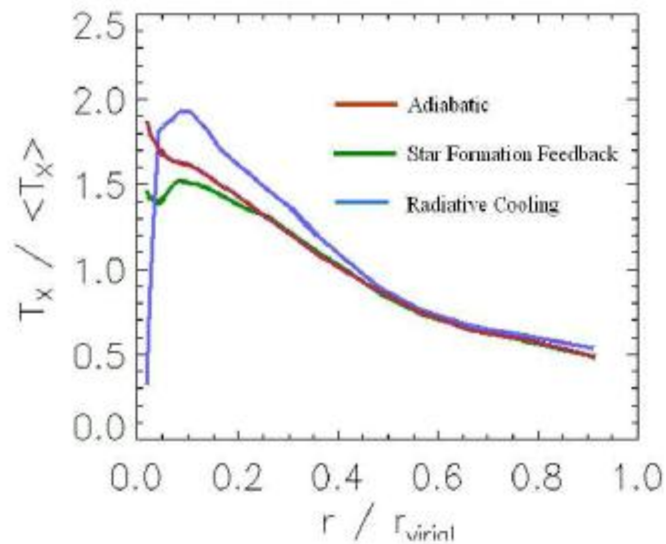

Shocks and Turbulence in X-ray Clusters
Vazza et al. 2008

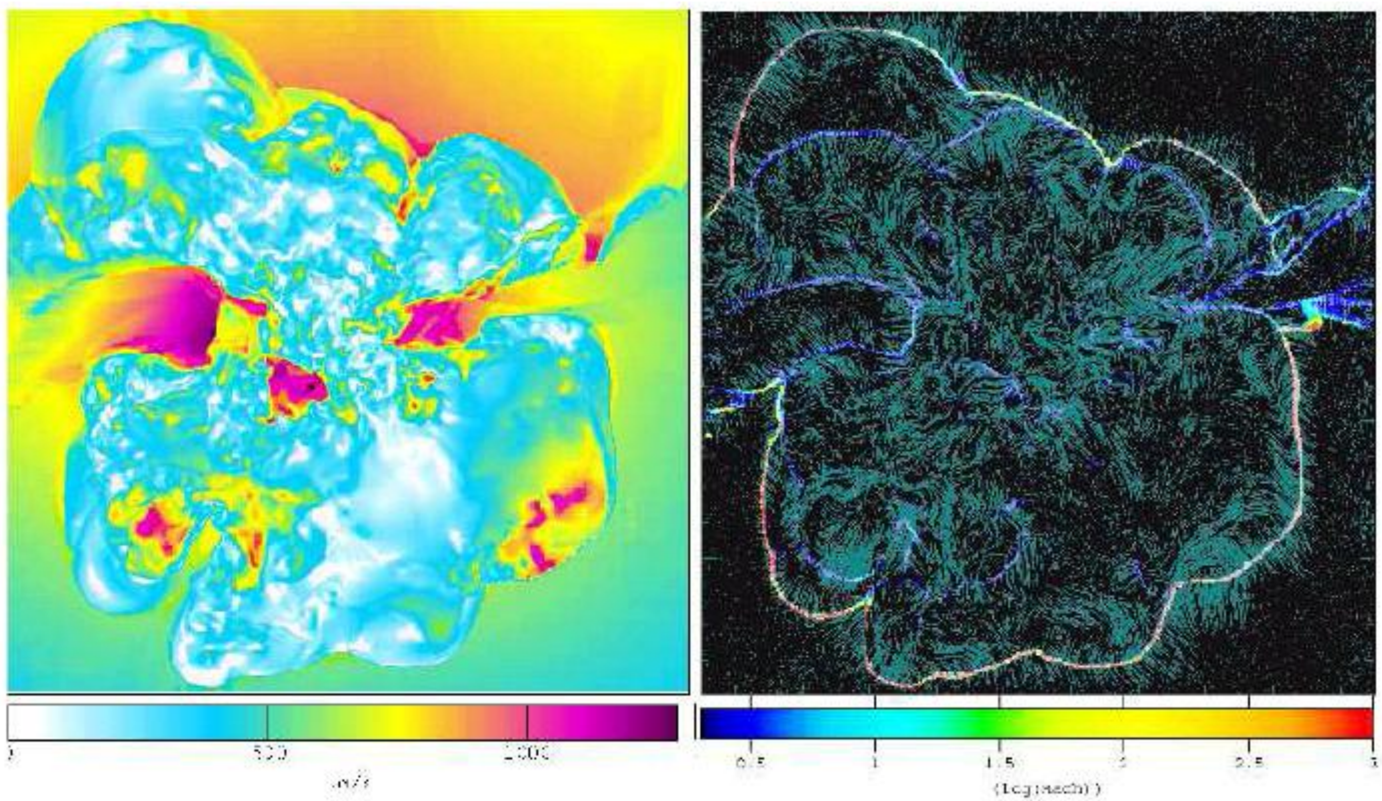

Lyman α Forest

Zhang et al. 1995, 1997, 1998; Bryan et al. 1999; Jena et al. 2005

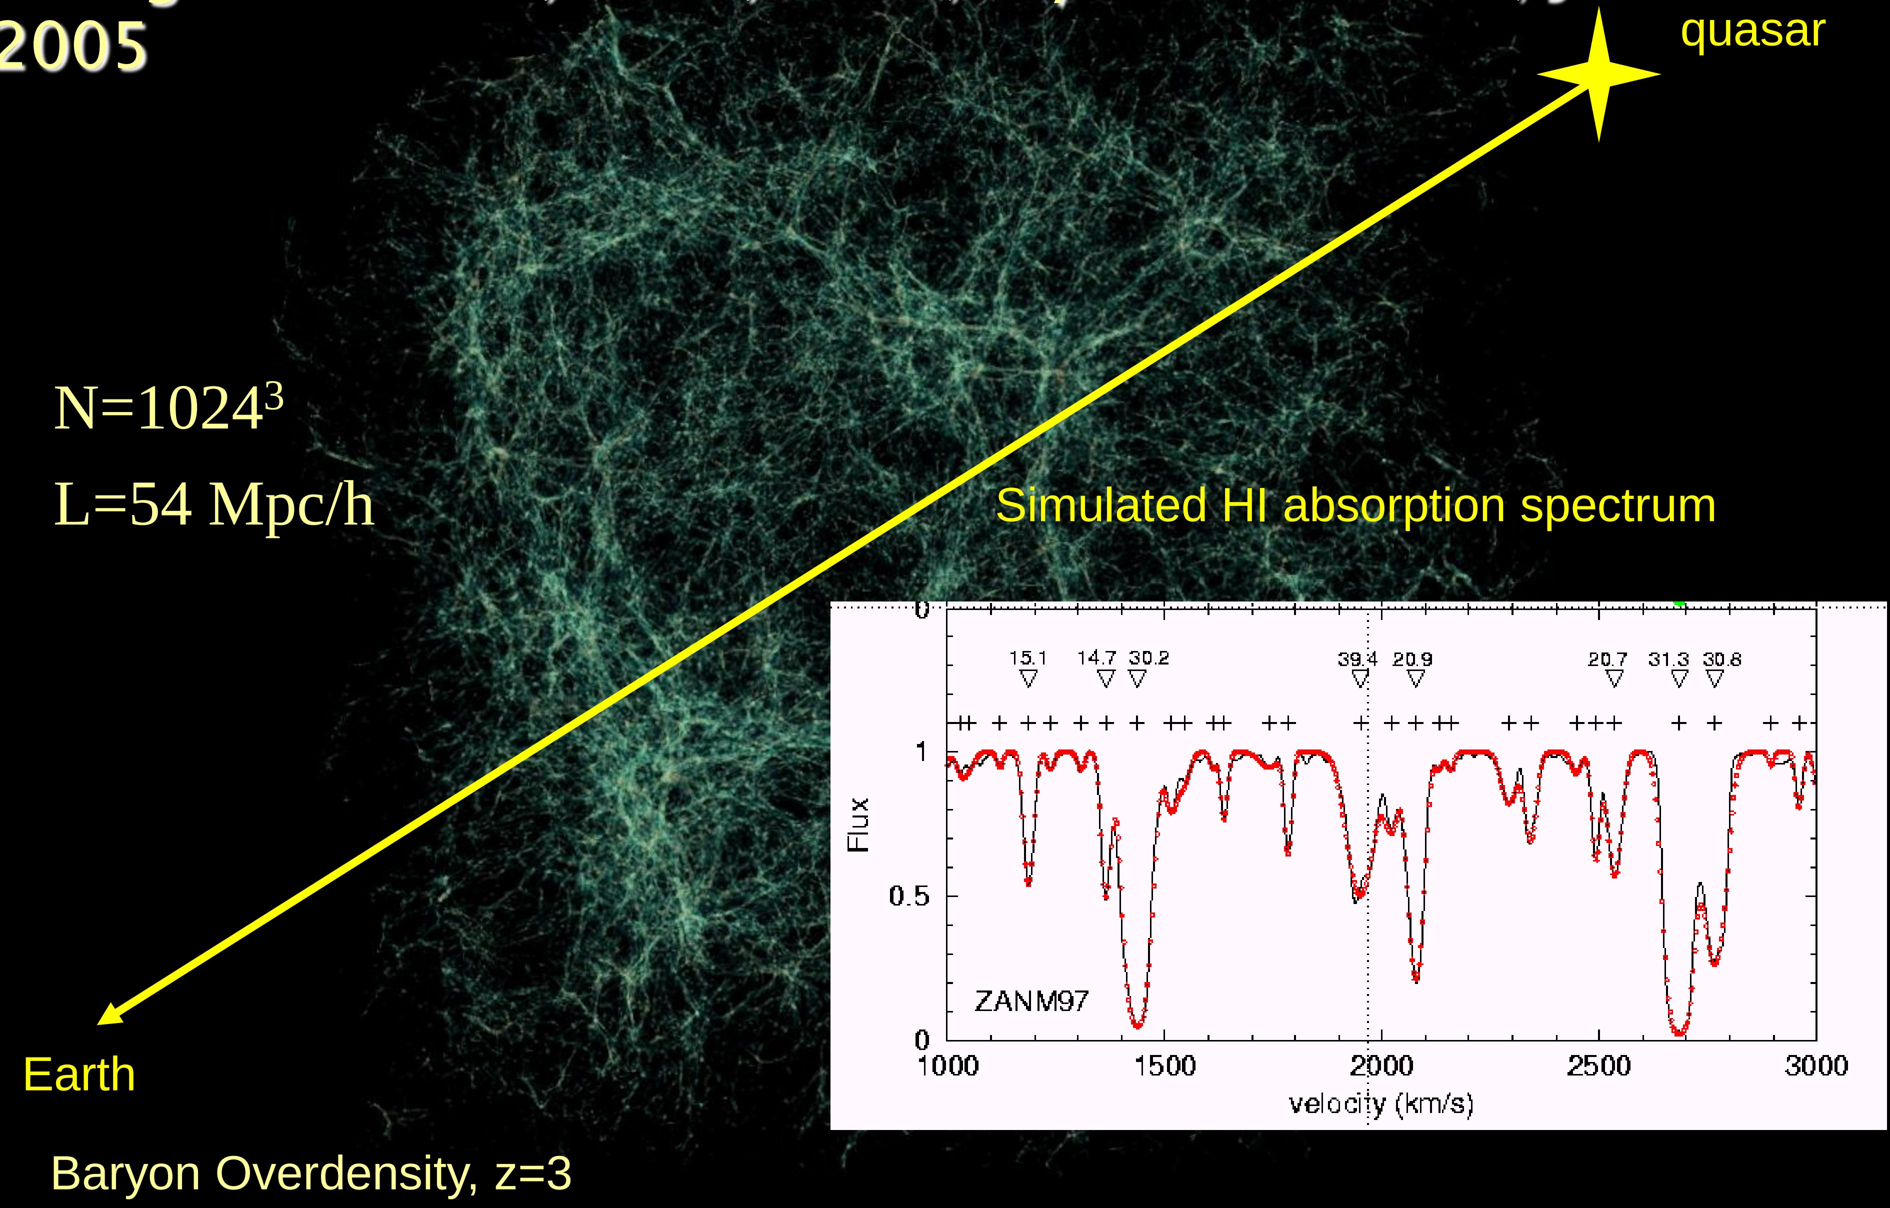

BAO in the Lyman α Forest

Norman et al. 2009

4096^3 =
68.7 billion cells
and particles

16,384 processors

2 million CPU-hrs
NICS Kraken

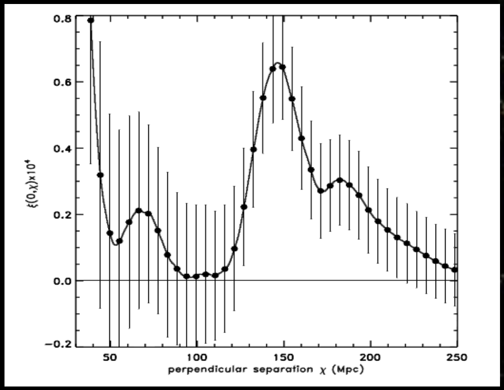

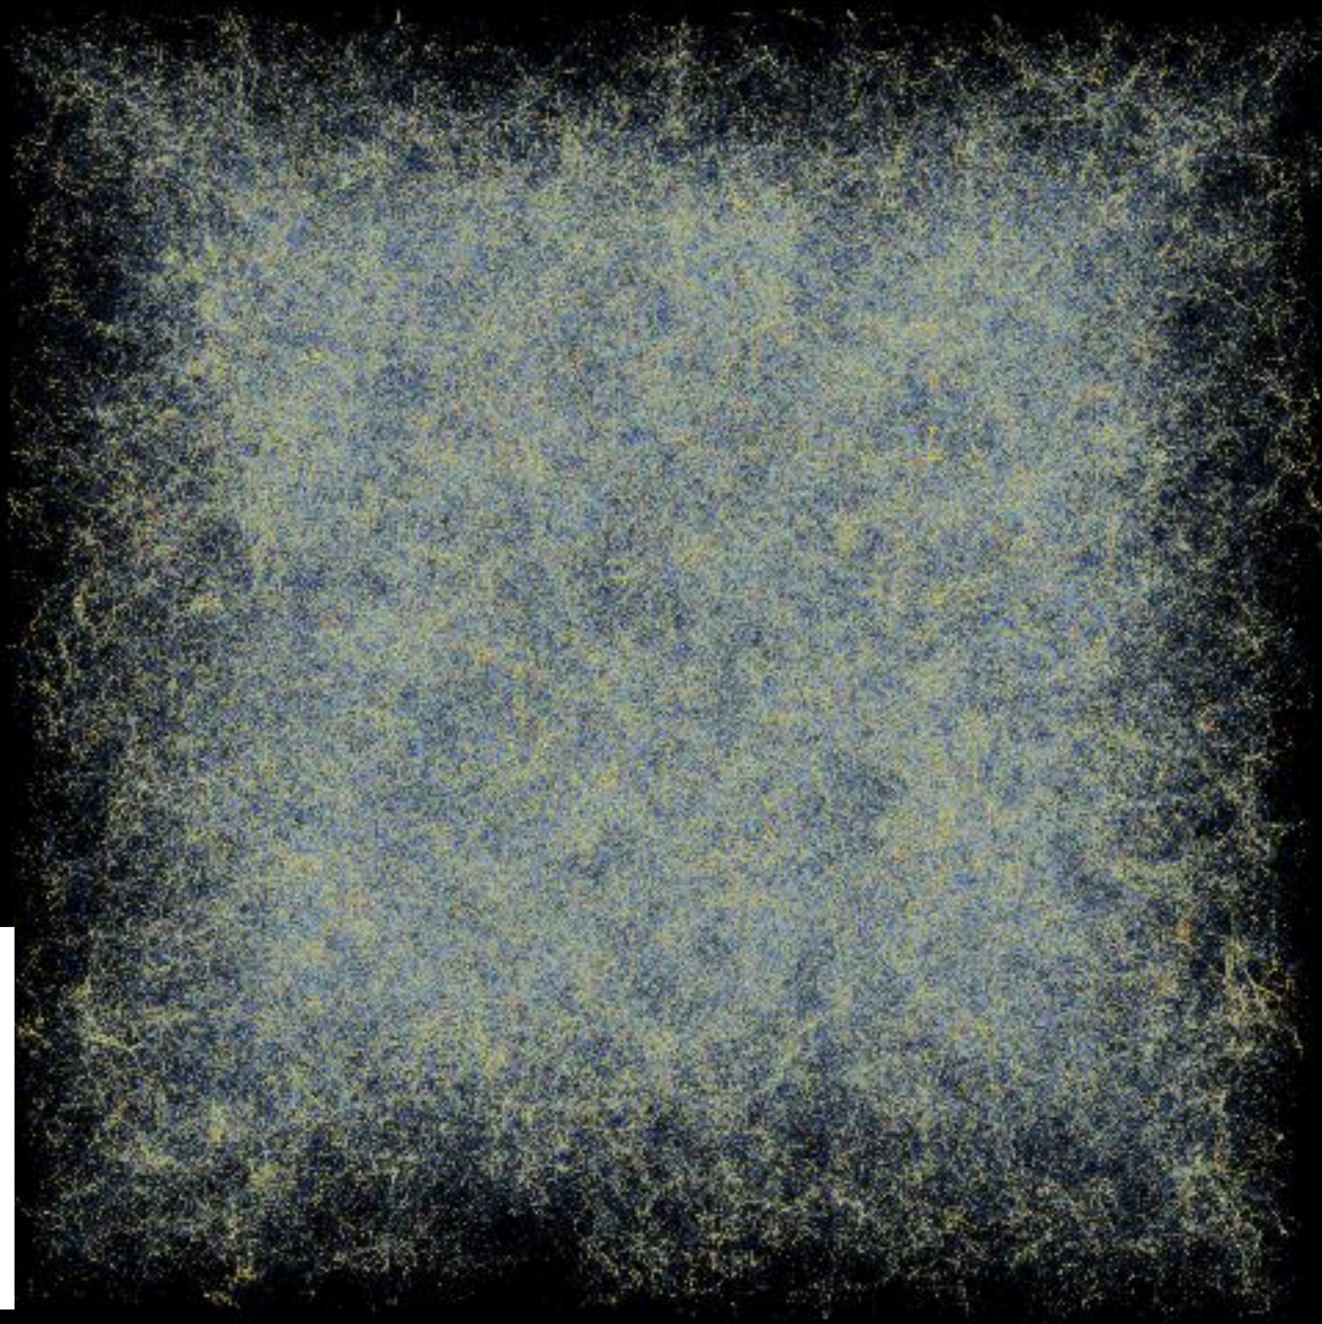

600 Mpc

Formation of First Stars

Abel, Bryan & Norman (2001)

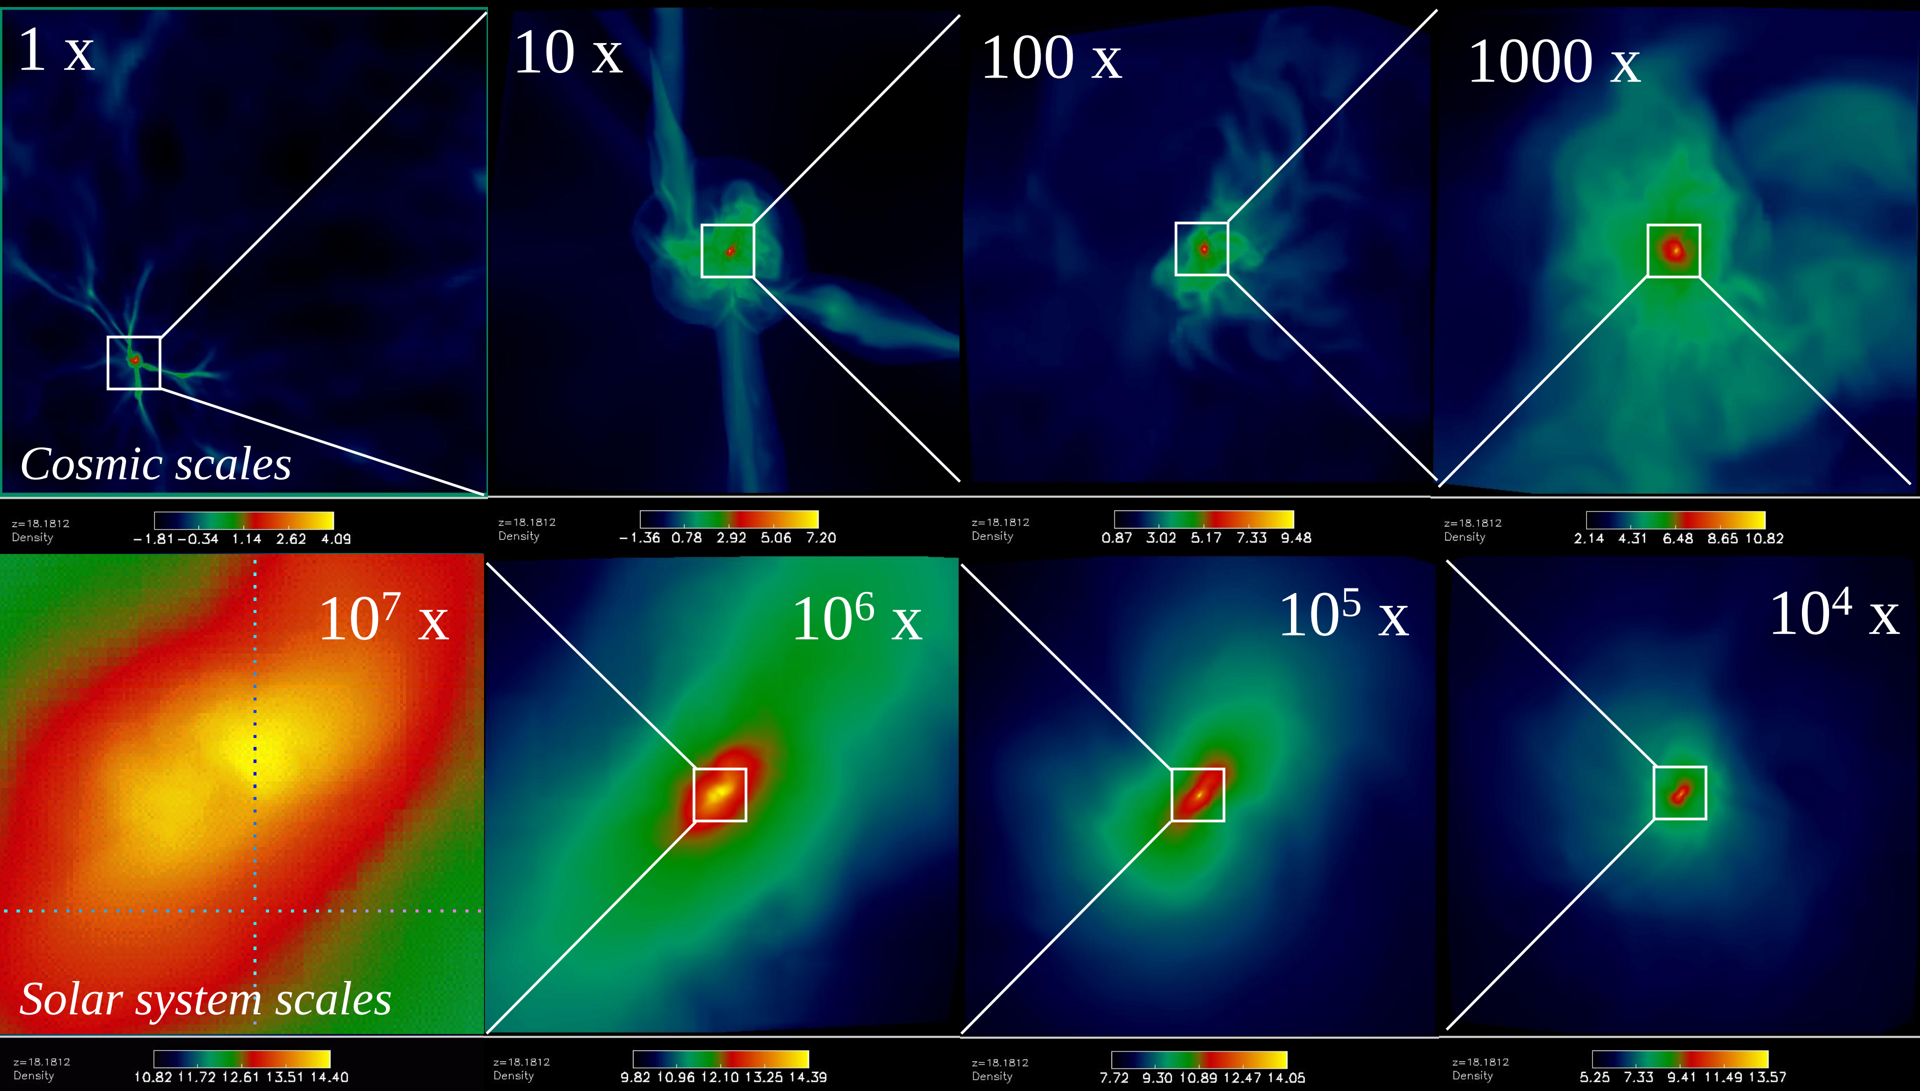

Pop III Binarity: *Princeton Twist* Survey

Turk et al. in prep

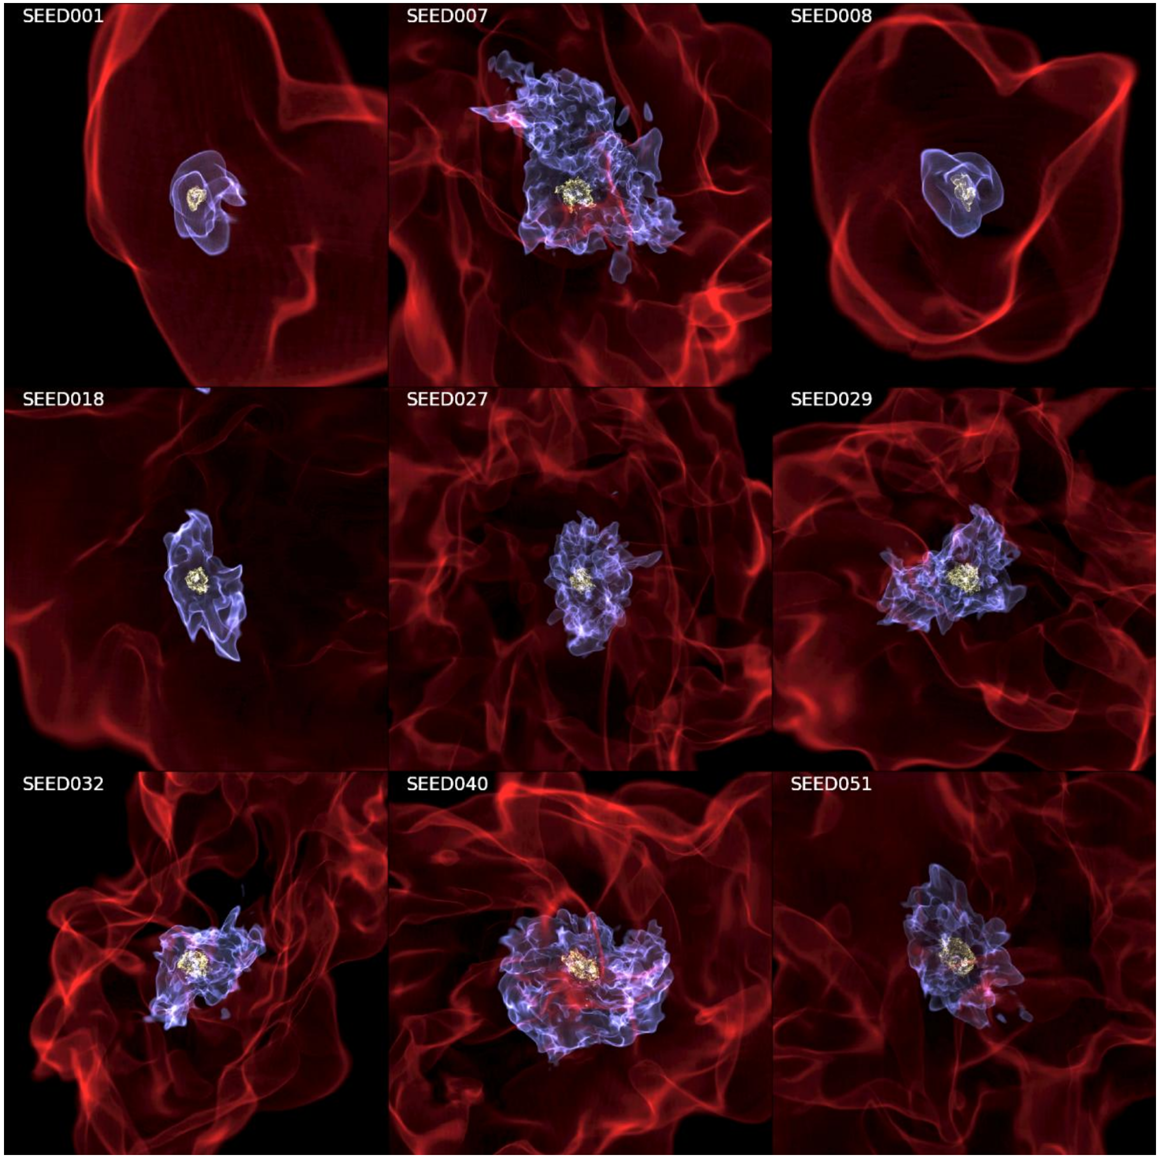

Pop III Star Formation – John Wise (Princeton)

512^3 AMR everywhere, 1 Mpc/h box (z=14.5)

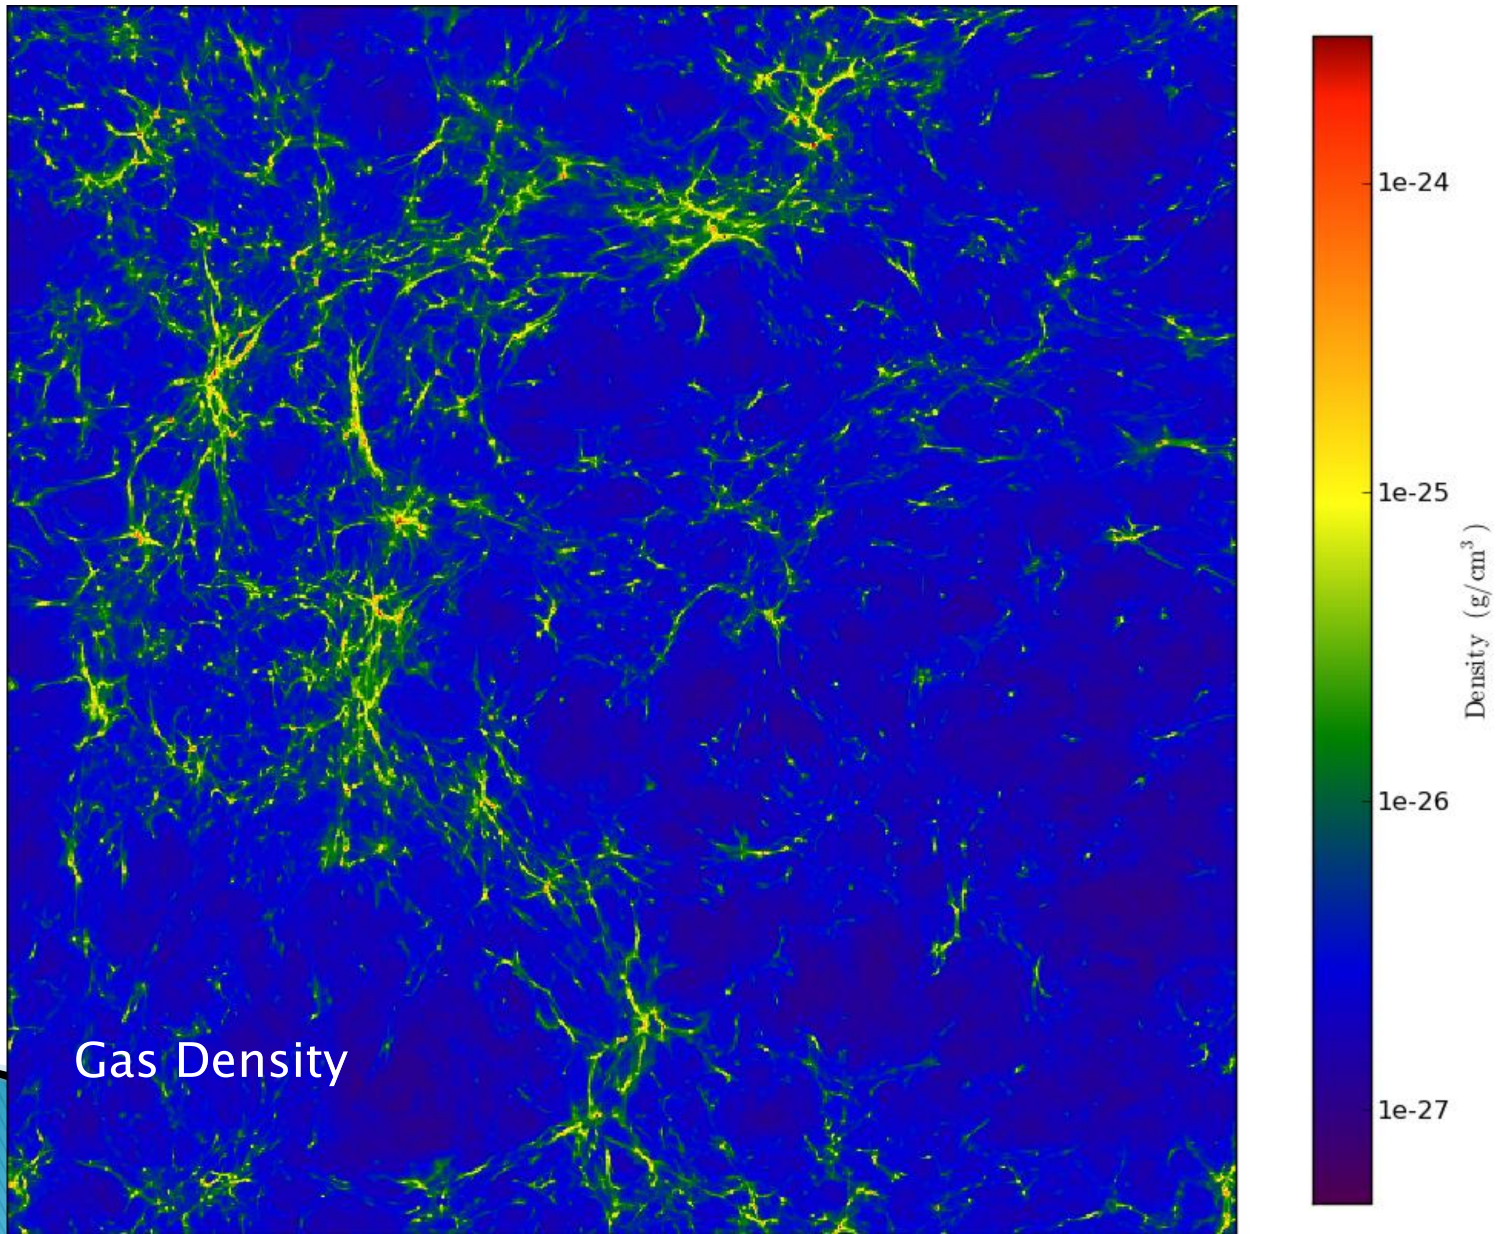

Pop III Star Formation – John Wise (Princeton)

512^3 AMR everywhere, 1 Mpc/h box (z=14.5)