Cosmological Simulations

Volker Springel

- Simulation predictions for the dark sector of ΛCDM
- Can we falsify ΛCDM with simulations?
- Exaflop computing
- Beyond the dark sector: Hydrodynamic simulations
- Challenges for galaxy formation simulations

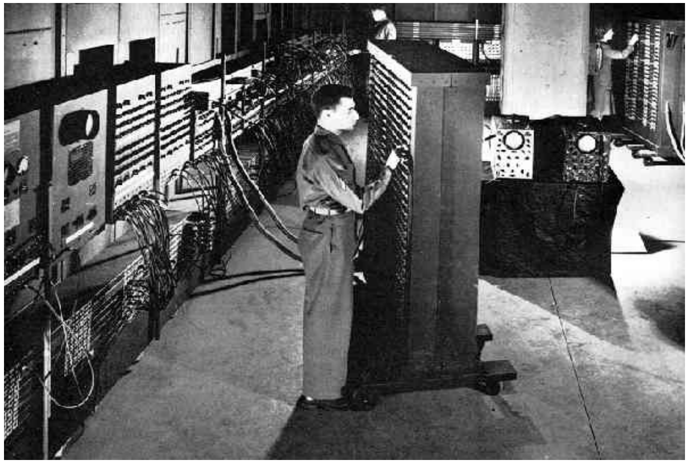

Heidelberg Institute for Theoretical Studies
HITS
Universität Heidelberg

Unsolved Problems in Cosmology and Astrophysics
Benasque, February 2011

The first slice in the CfA redshift survey

1100 galaxies in a wedge,
6 degrees wide and 110 degrees long

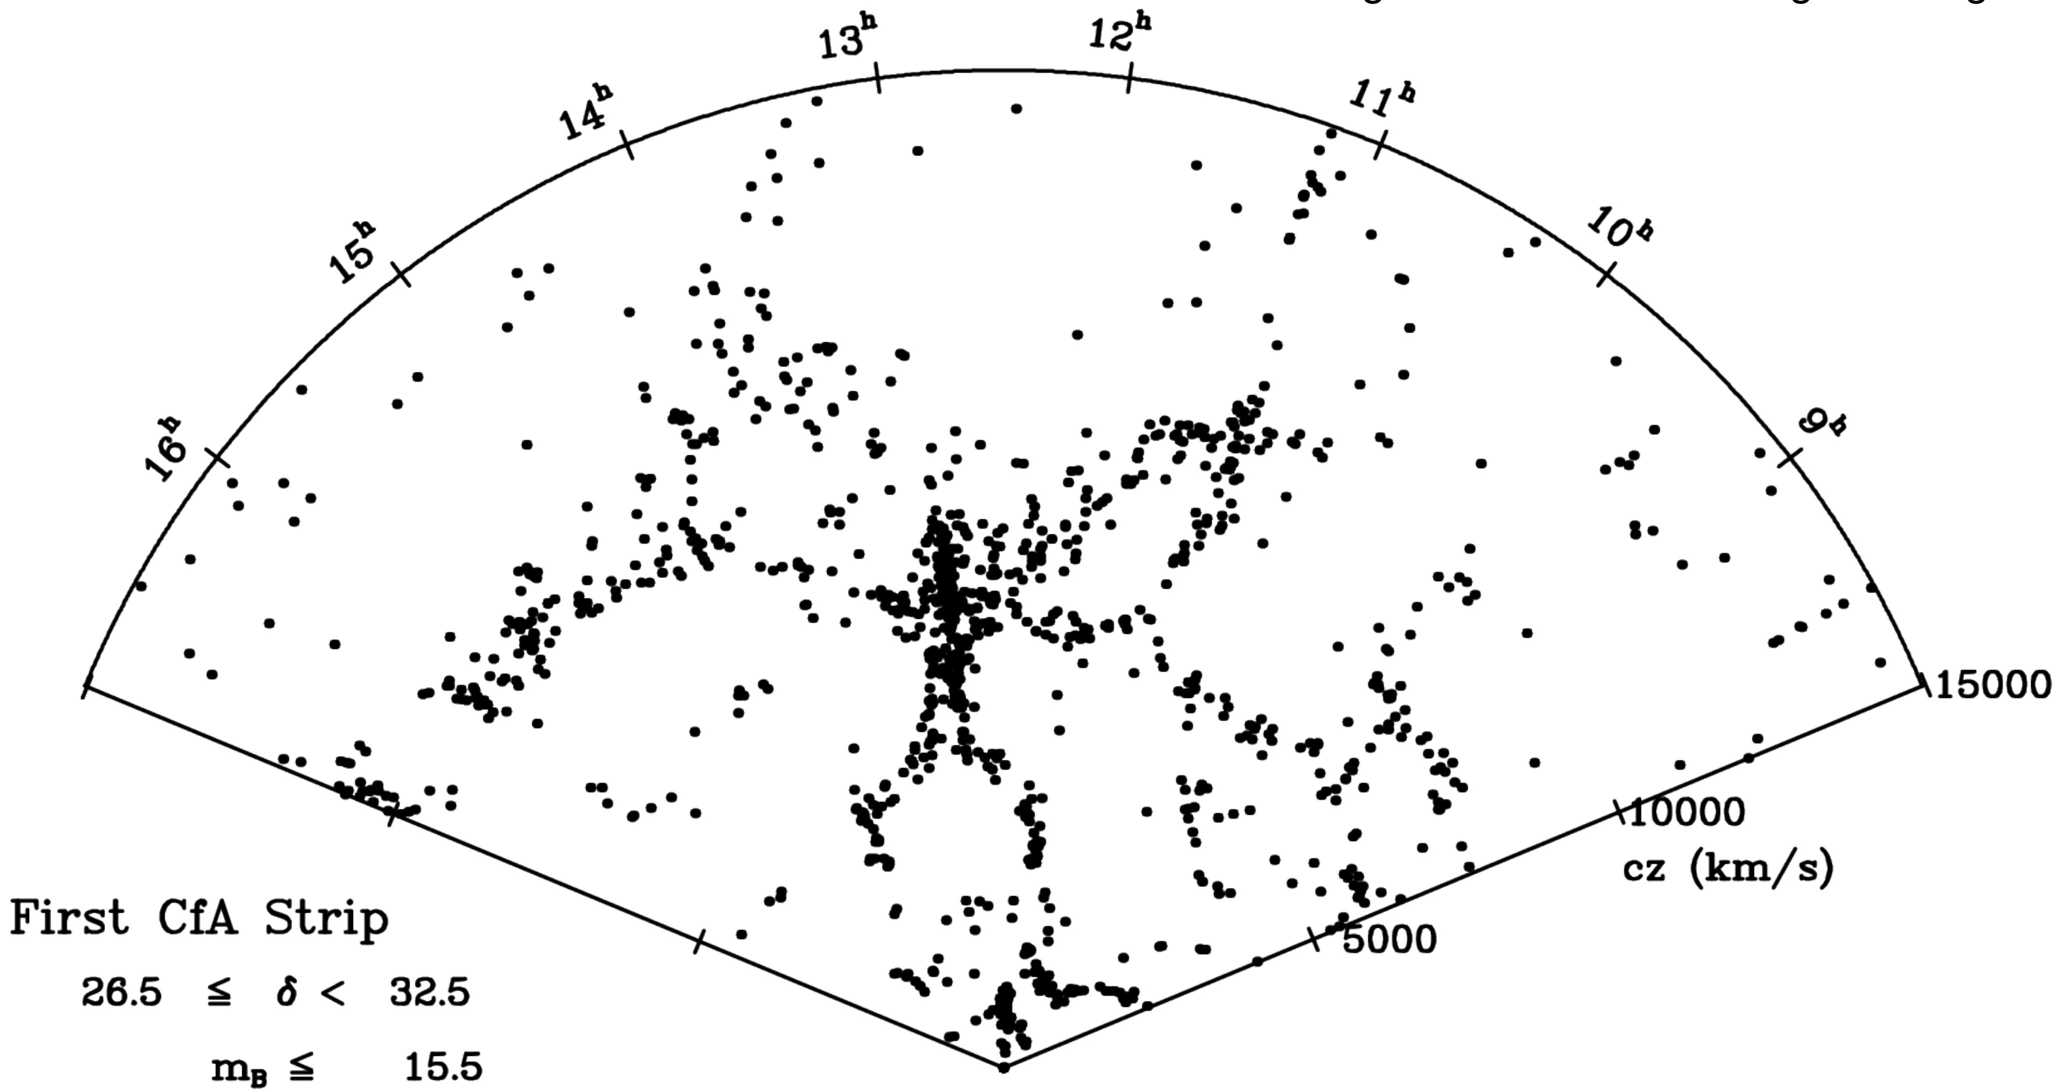

de Lapparent, Geller, Huchra (1986)

Current galaxy redshift surveys map the Universe with several hundred thousand galaxies

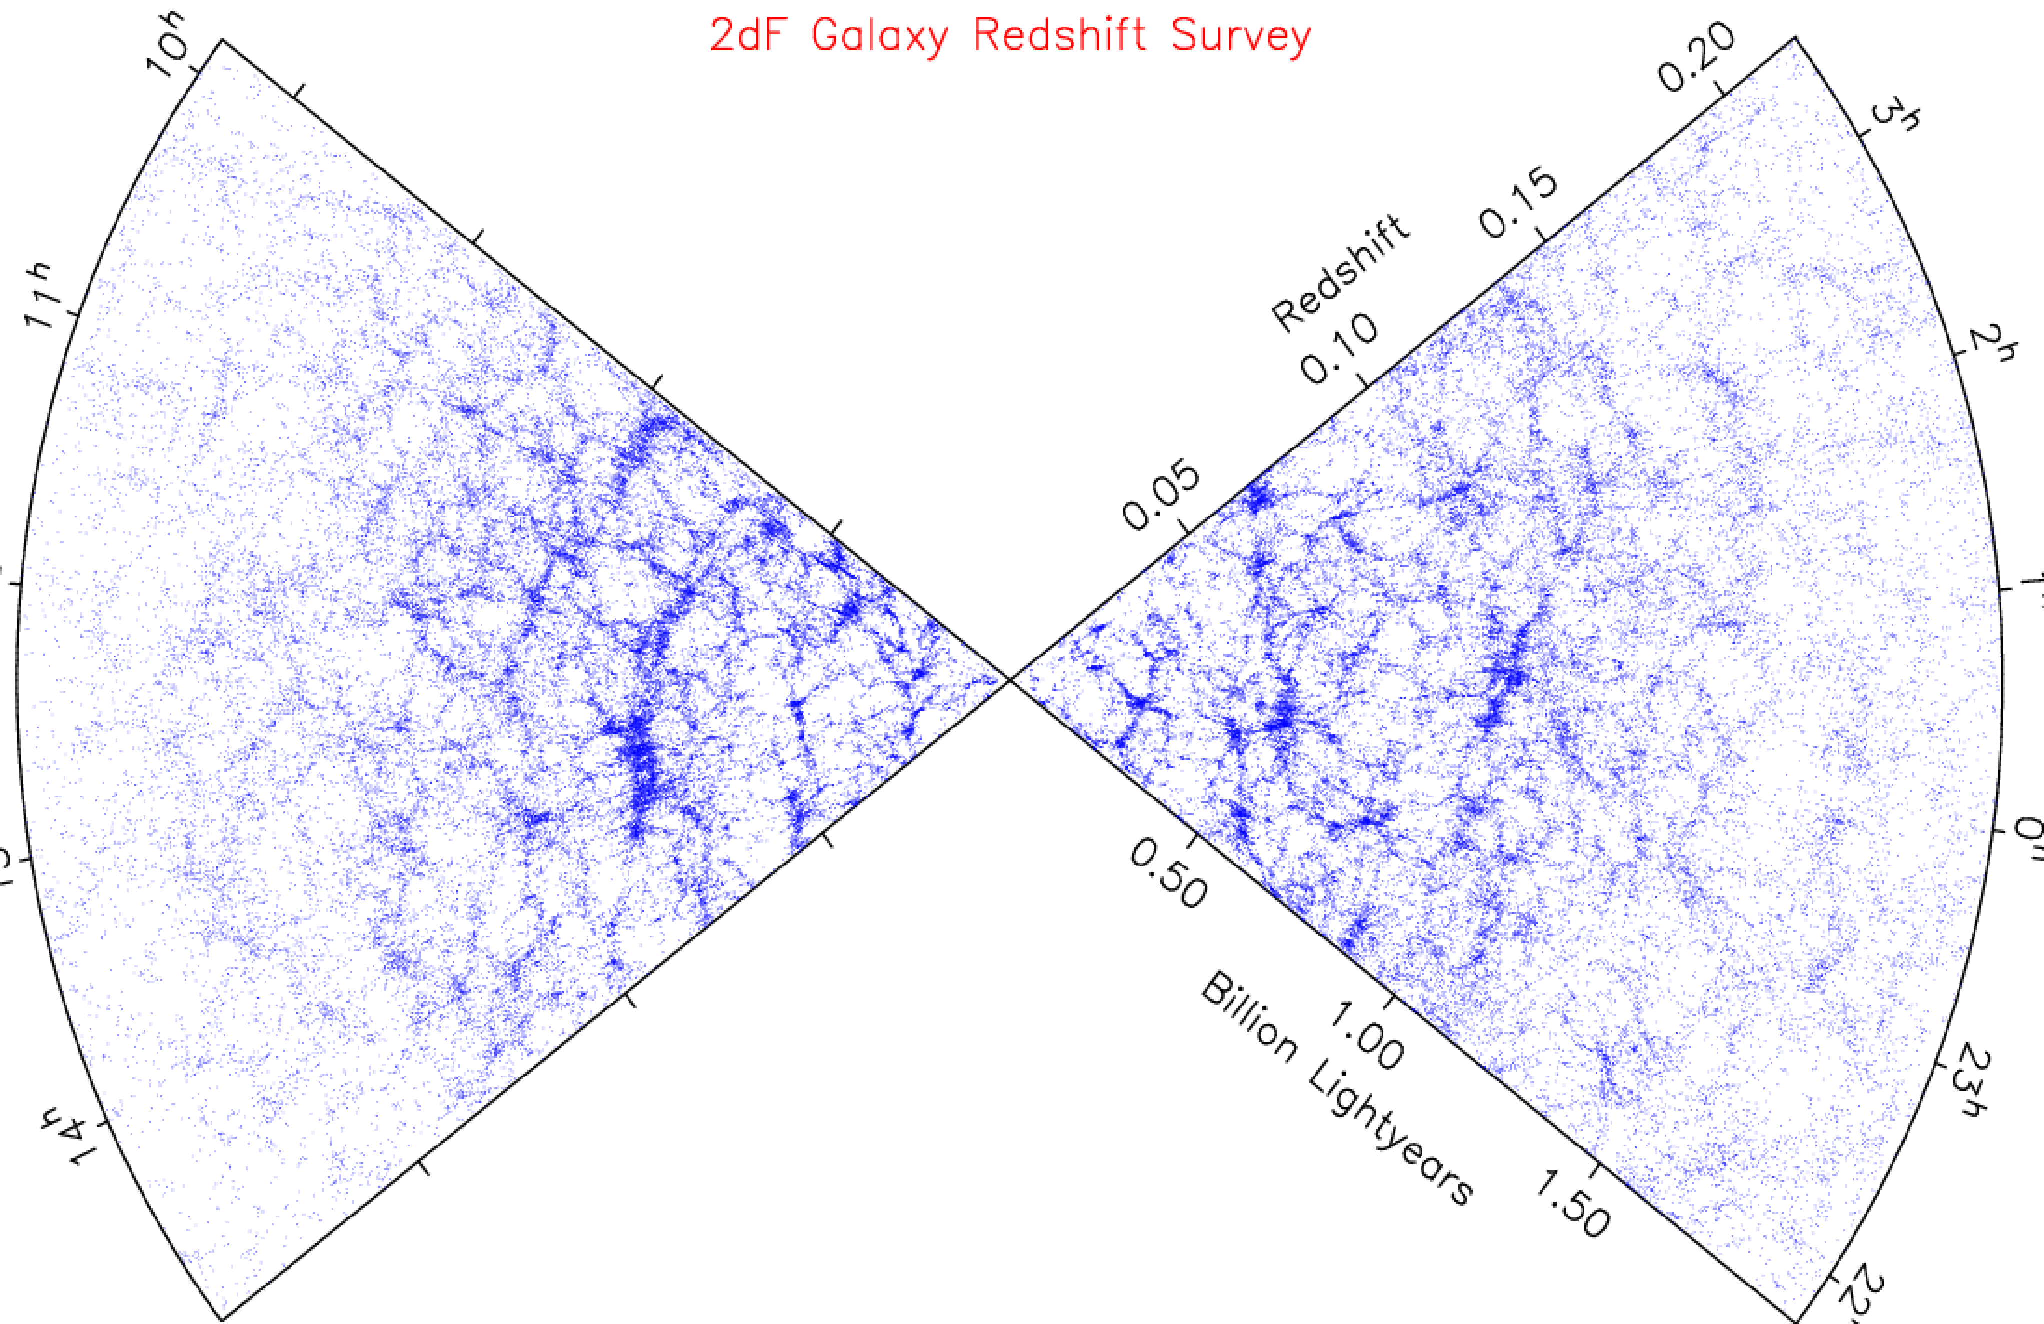

The initial conditions for cosmic structure formation are directly observable

THE MICROWAVE SKY

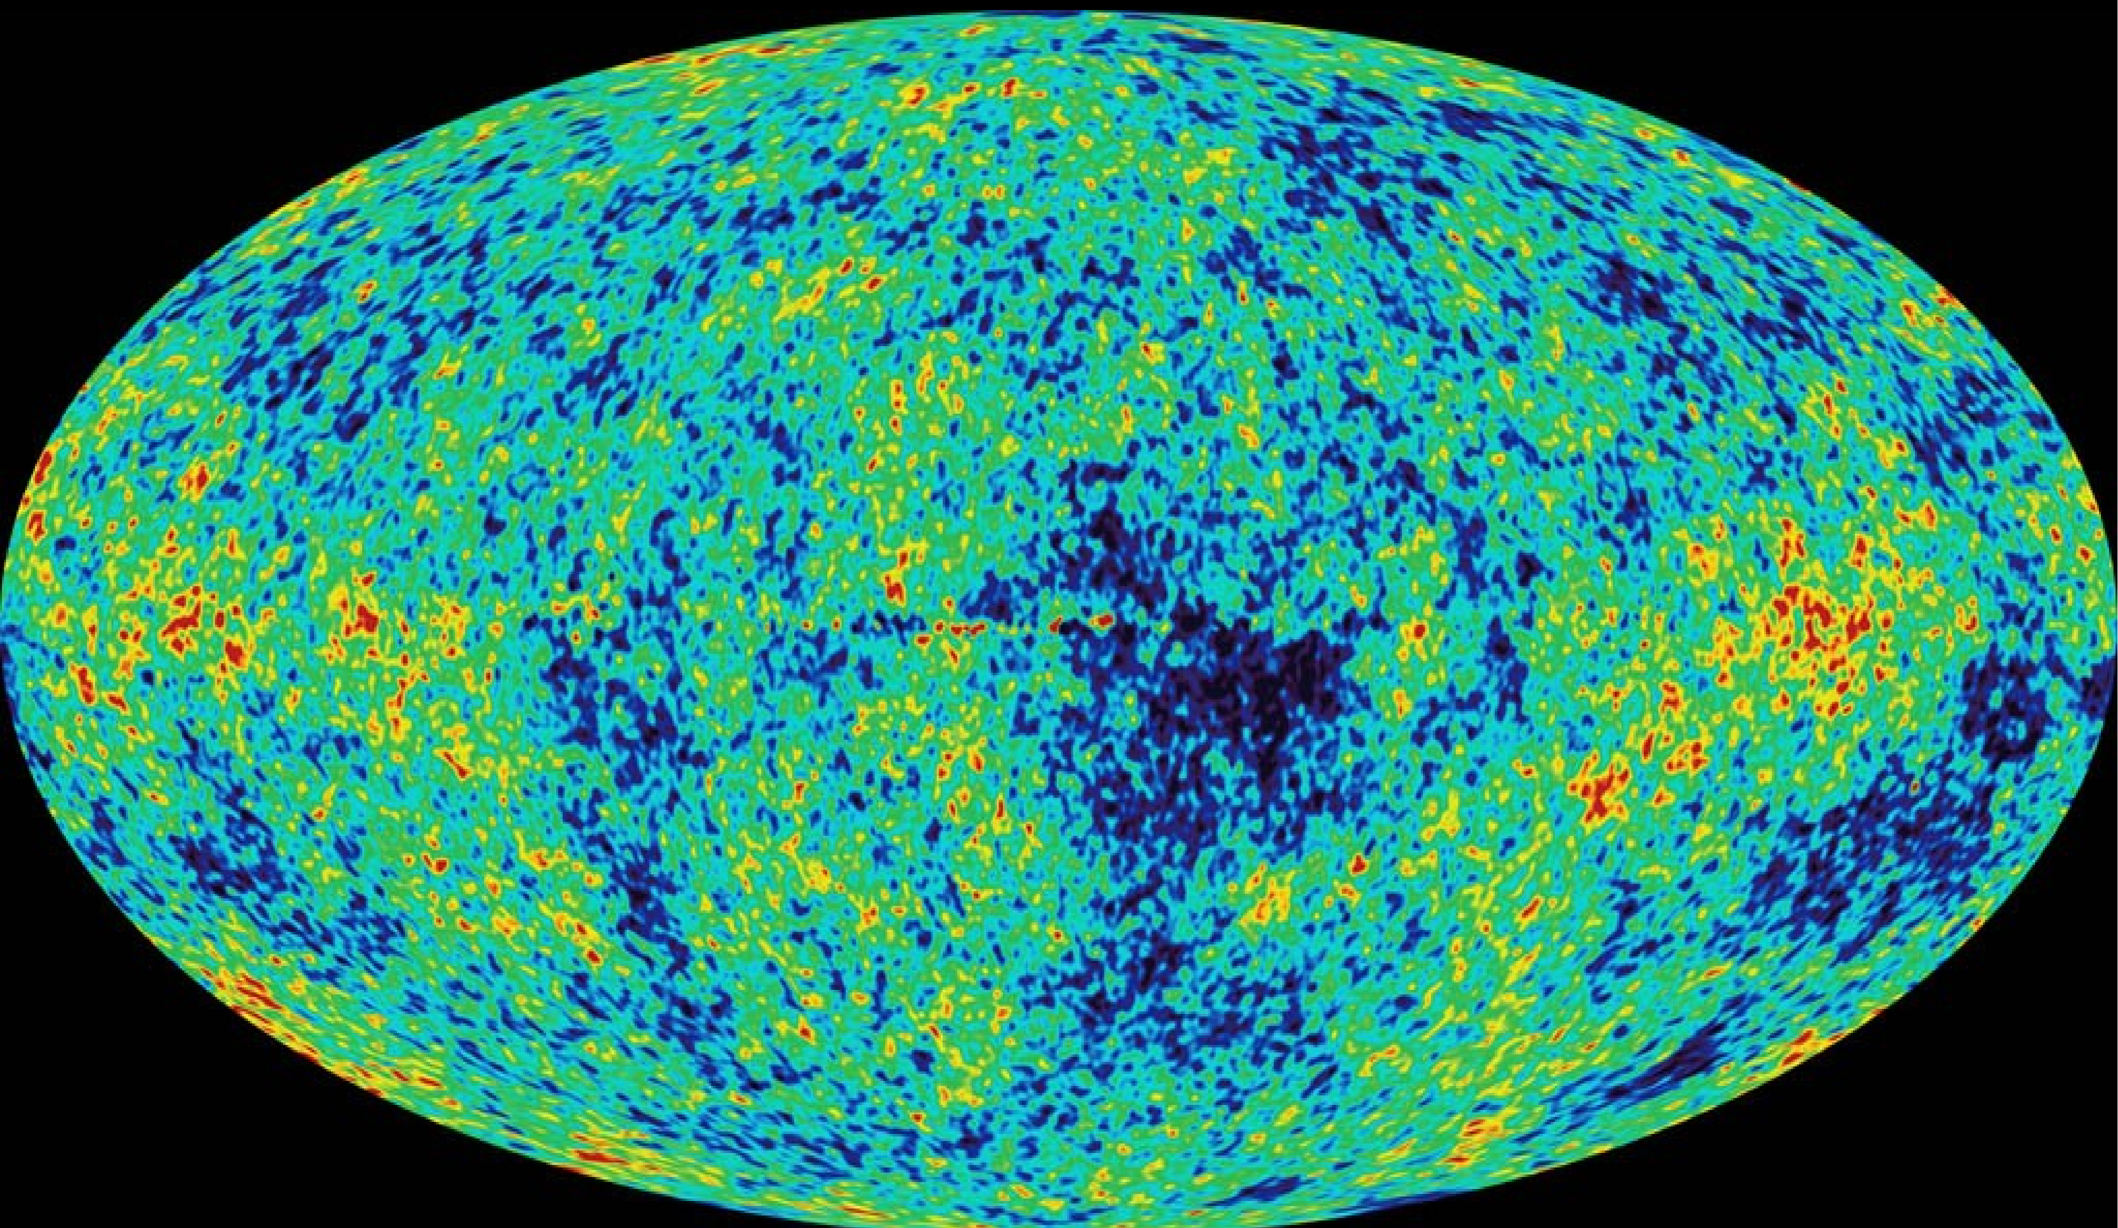

WMAP Science Team (2003, 2006, 2008, 2010)

The basic dynamics of structure formation in the **dark matter**

BASIC EQUATIONS AND THEIR DISCRETIZATION

Gravitation
(Newtonian approximation to GR in an expanding space-time)

Dark matter is collisionless

Monte-Carlo integration as **N-body System**

3N **coupled**, non-linear differential equations of second order

Friedmann-Lemaitre model

$$H(a) = H_0 \sqrt{a^{-3}\Omega_0 + a^{-2}(1-\Omega_0-\Omega_\Lambda) + \Omega_\Lambda}$$

Collisionless Boltzmann equation with self-gravity

$$\frac{\mathrm{d}f}{\mathrm{d}t} \equiv \frac{\partial f}{\partial t} + \mathbf{v}\frac{\partial f}{\partial \mathbf{x}} - \frac{\partial \Phi}{\partial \mathbf{r}}\frac{\partial f}{\partial \mathbf{v}} = 0$$

$$\nabla^2\Phi(\mathbf{r},t) = 4\pi G \int f(\mathbf{r},\mathbf{v},t)\mathrm{d}\mathbf{v}$$

Hamiltonian dynamics in expanding space-time

$$H = \sum_i \frac{\boldsymbol{p}_i^2}{2\,m_i\,a(t)^2} + \frac{1}{2}\sum_{ij}\frac{m_i m_j\,\varphi(\boldsymbol{x}_i-\boldsymbol{x}_j)}{a(t)}$$

$$\nabla^2\varphi(\boldsymbol{x}) = 4\pi G\left[-\frac{1}{L^3} + \sum_{\boldsymbol{n}}\tilde{\delta}(\boldsymbol{x}-\boldsymbol{n}L)\right]$$

Problems:

- N is very large
- All equations are coupled with each other

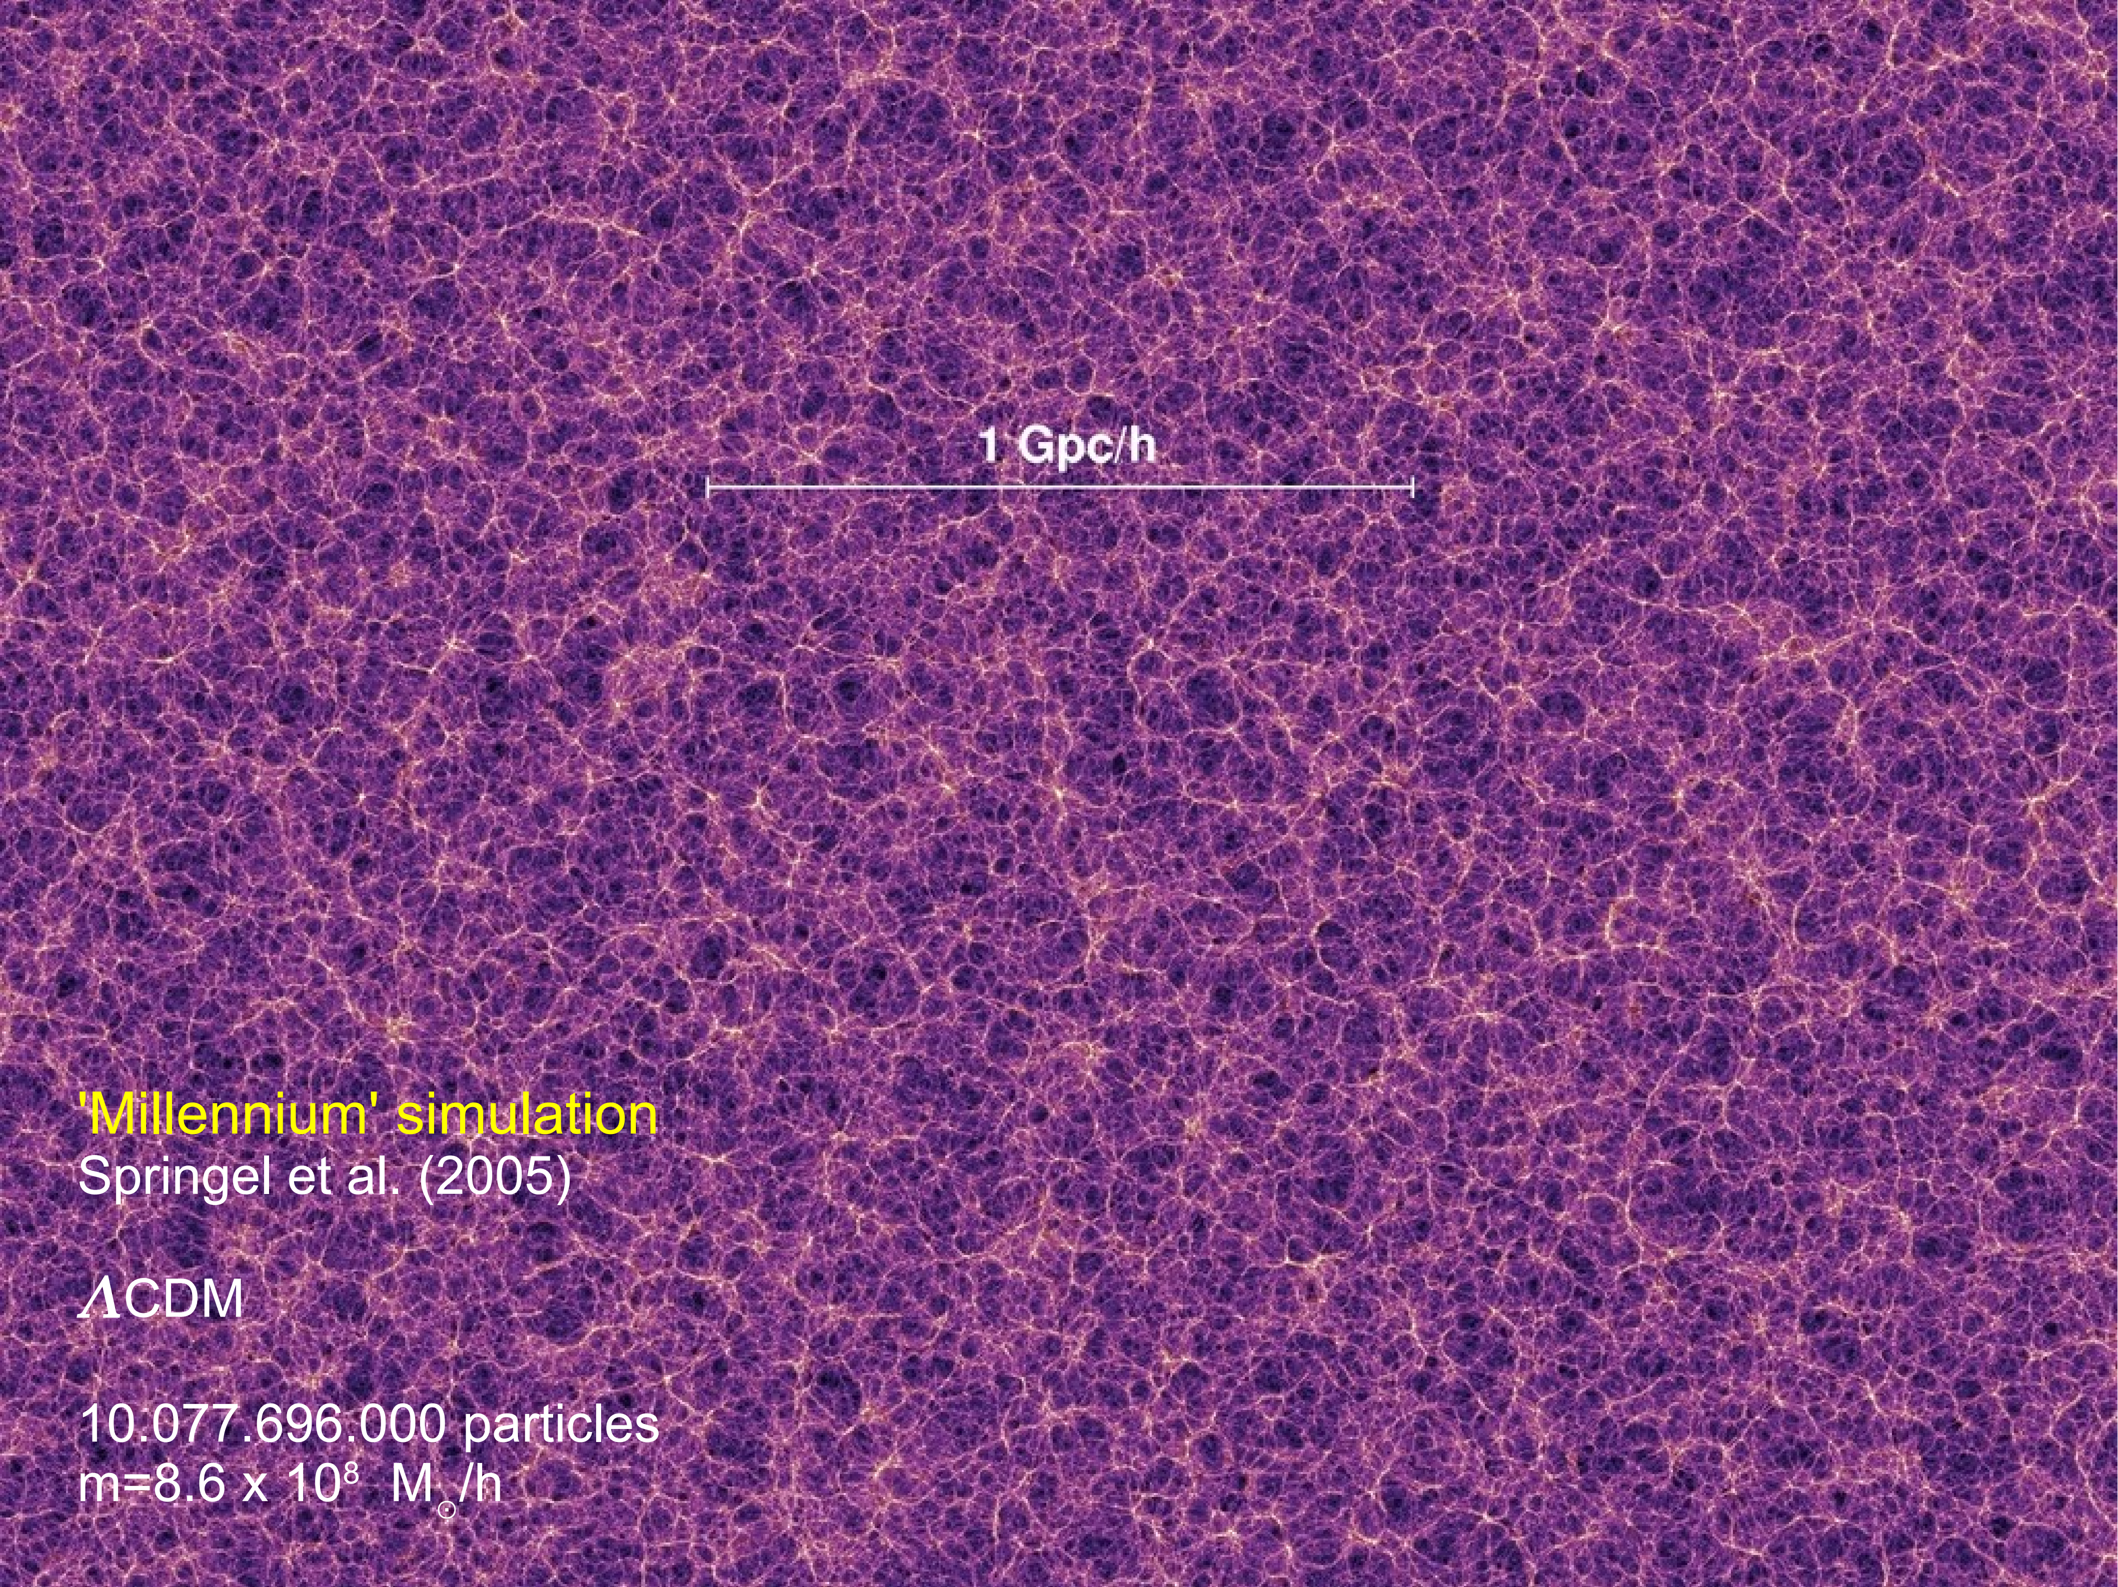
1 Gpc/h
'Millennium' simulation
Springel et al. (2005)
ΛCDM
10.077.696.000 particles
m=8.6 x 10^8 $M_\odot$/h

Why are **cosmological simulations** of structure formation **useful for studying the dark universe**?

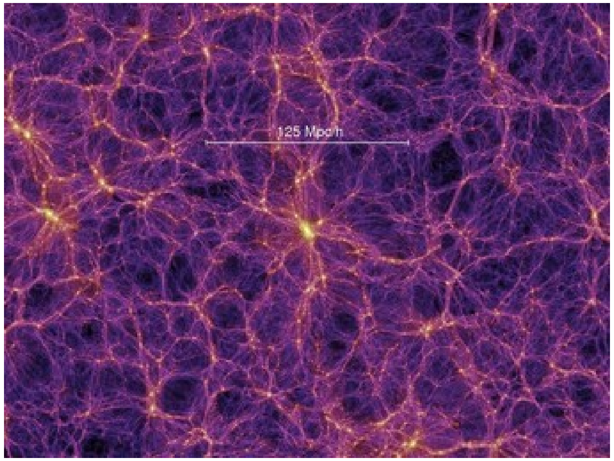

Simulations are the theoretical tool of choice for calculations in the non-linear regime.

They connect the (simple) cosmological initial conditions with the (complex) present-day universe.

Predictions from N-body simulations:

- Abundance of objects (as a function of mass and time)
- Their spatial distribution
- Internal structure of halos (e.g. density profiles, spin)
- Mean formation epochs
- Merger rates
- Detailed dark matter distribution on large *and* fairly small scales
- Galaxy formation models
- Gravitational lensing
- Baryonic acoustic oscillations in the matter distribution
- Integrated Sachs-Wolfe effect
- Dark matter annihilation rate
- Morphology of large-scale structure (“cosmic web”)
- ....

Simulations provide accurate measurements for halo abundance as a function of time

CONVERGENCE RESULTS FOR HALO ABUNDANCE

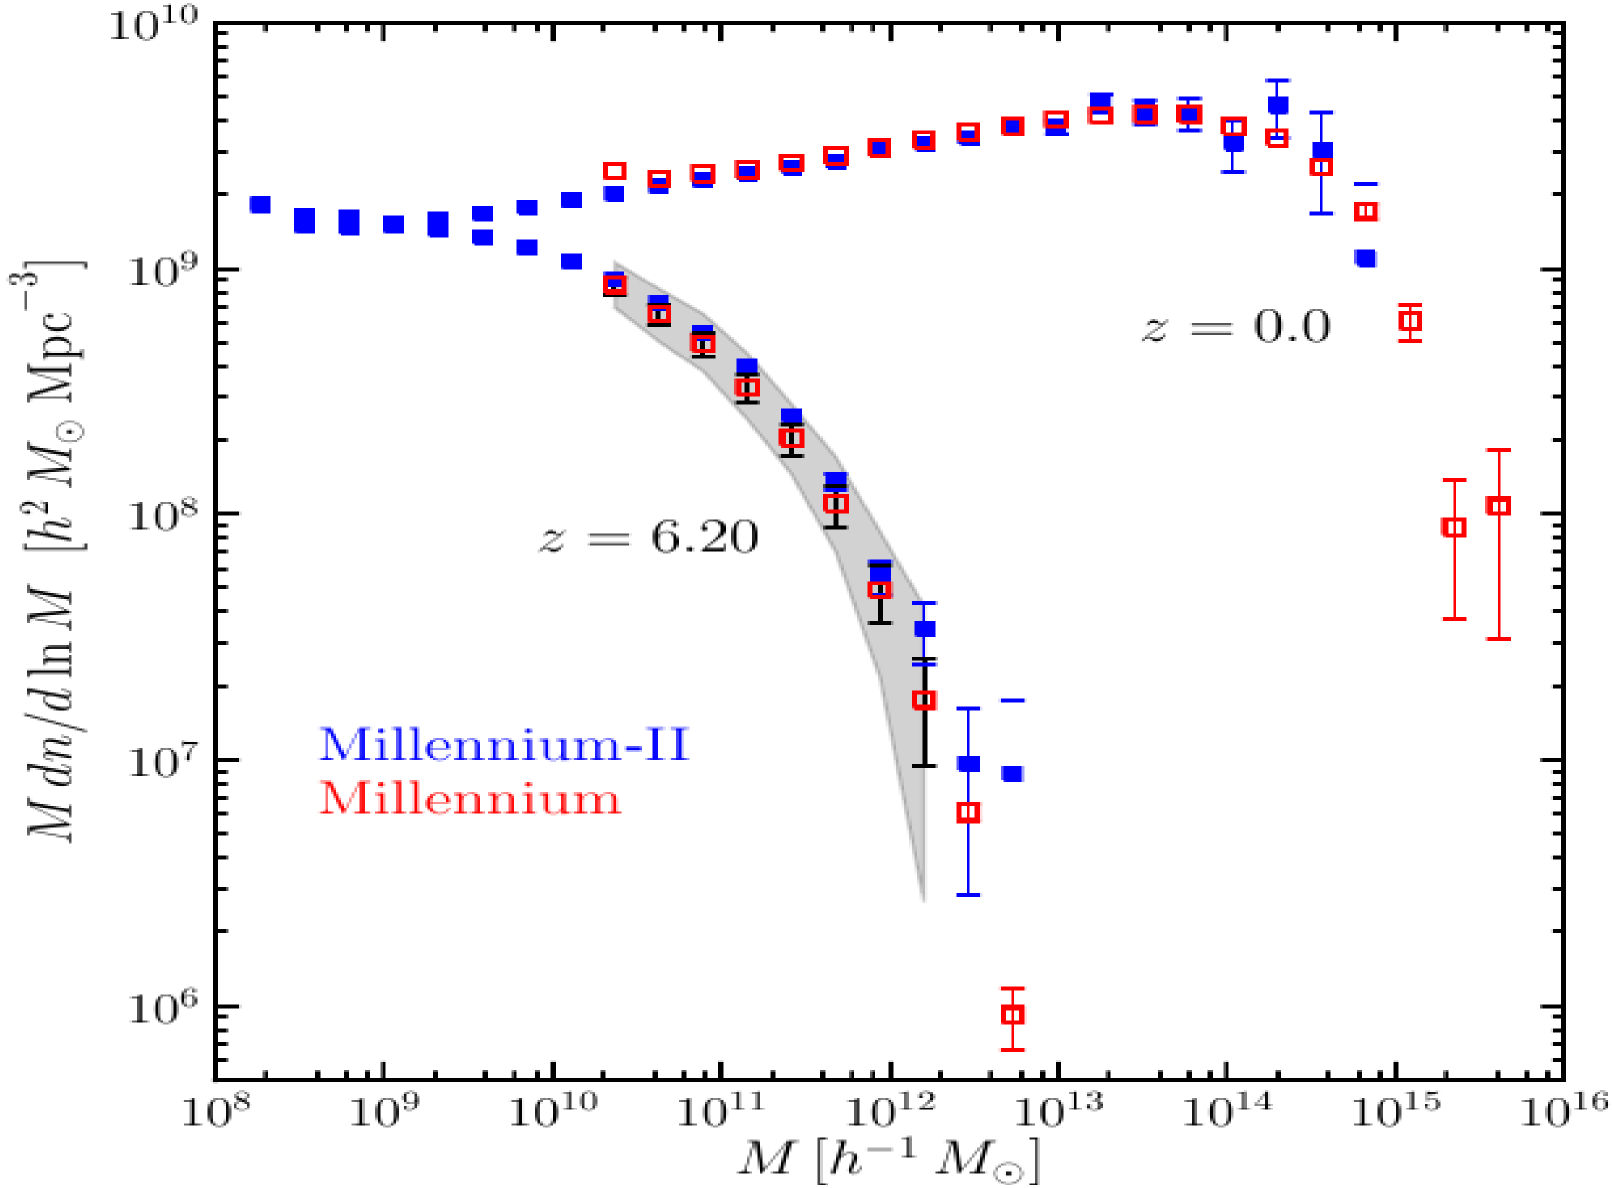

Boylan-Kolchin, Springel, White, et al. (2009)

Simulated and observed large-scale structure in the galaxy distribution

MOCK PIE DIAGRAMS COMPARED TO SDSS, 2DFGRS, AND CFA-2

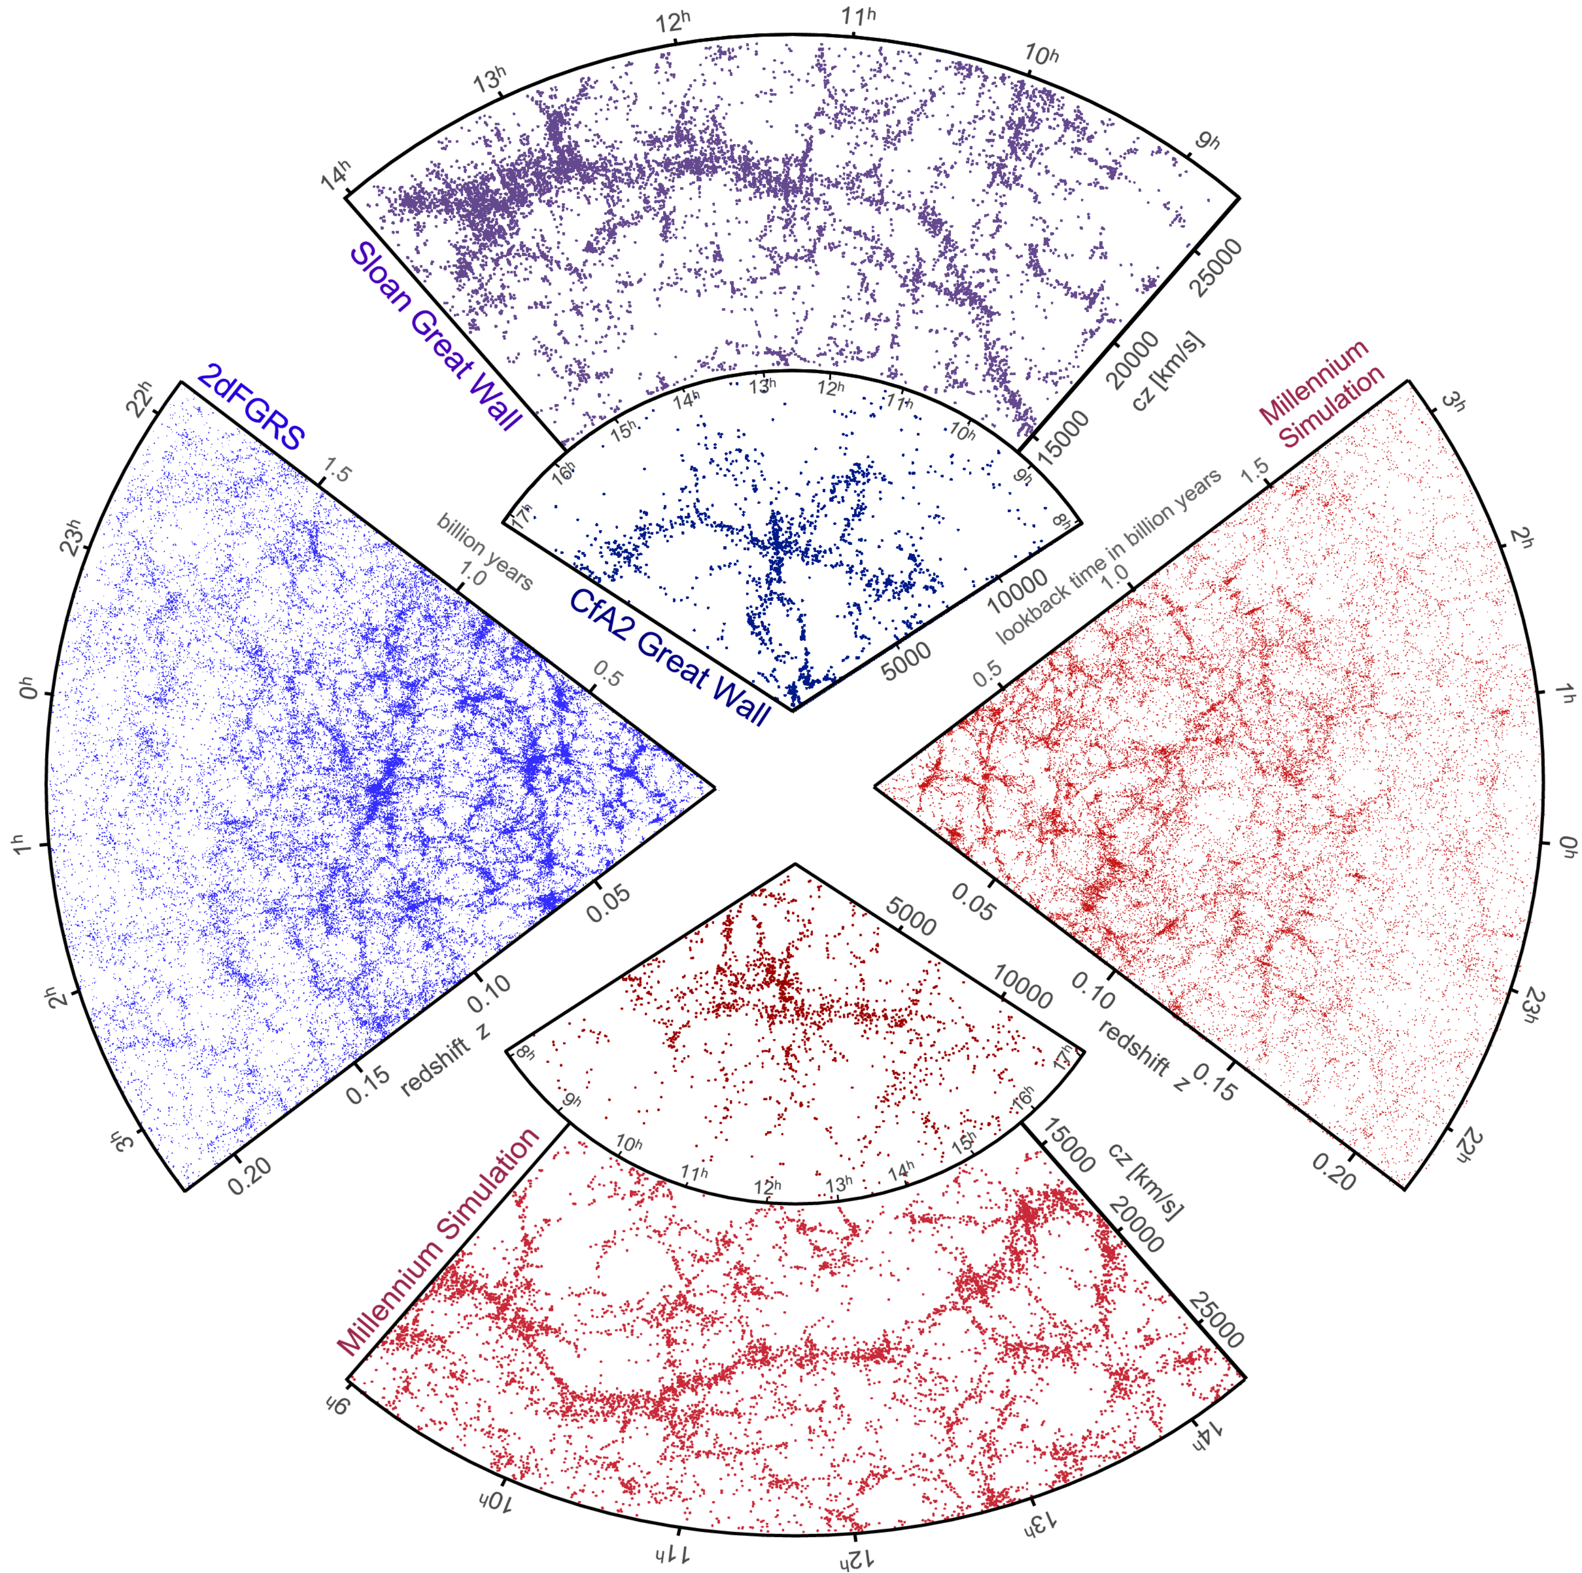

Springel et al. (2006)

The two-point correlation function of galaxies in the Millennium run is a very good power law

GALAXY TWO-POINT FUNCTION COMPARED WITH **2dFGRS**

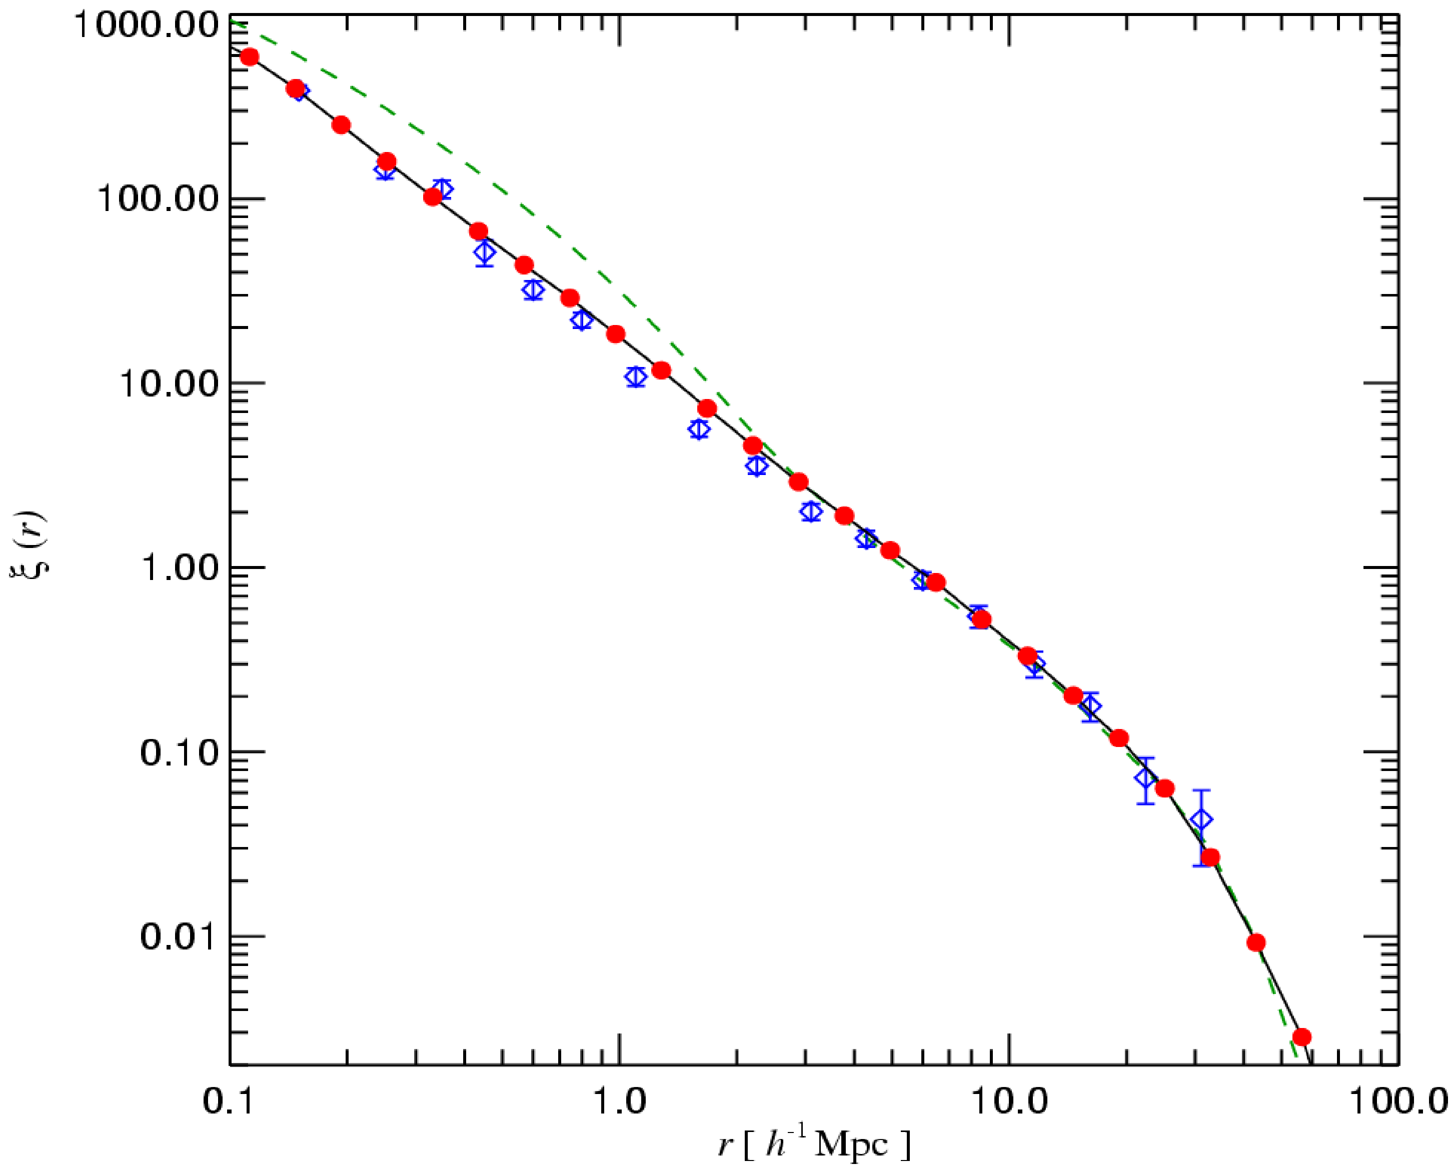

The galaxy distribution is **biased** with respect to the mass distribution

GALAXY AND MASS CLUSTERING AT DIFFERENT EPOCHS

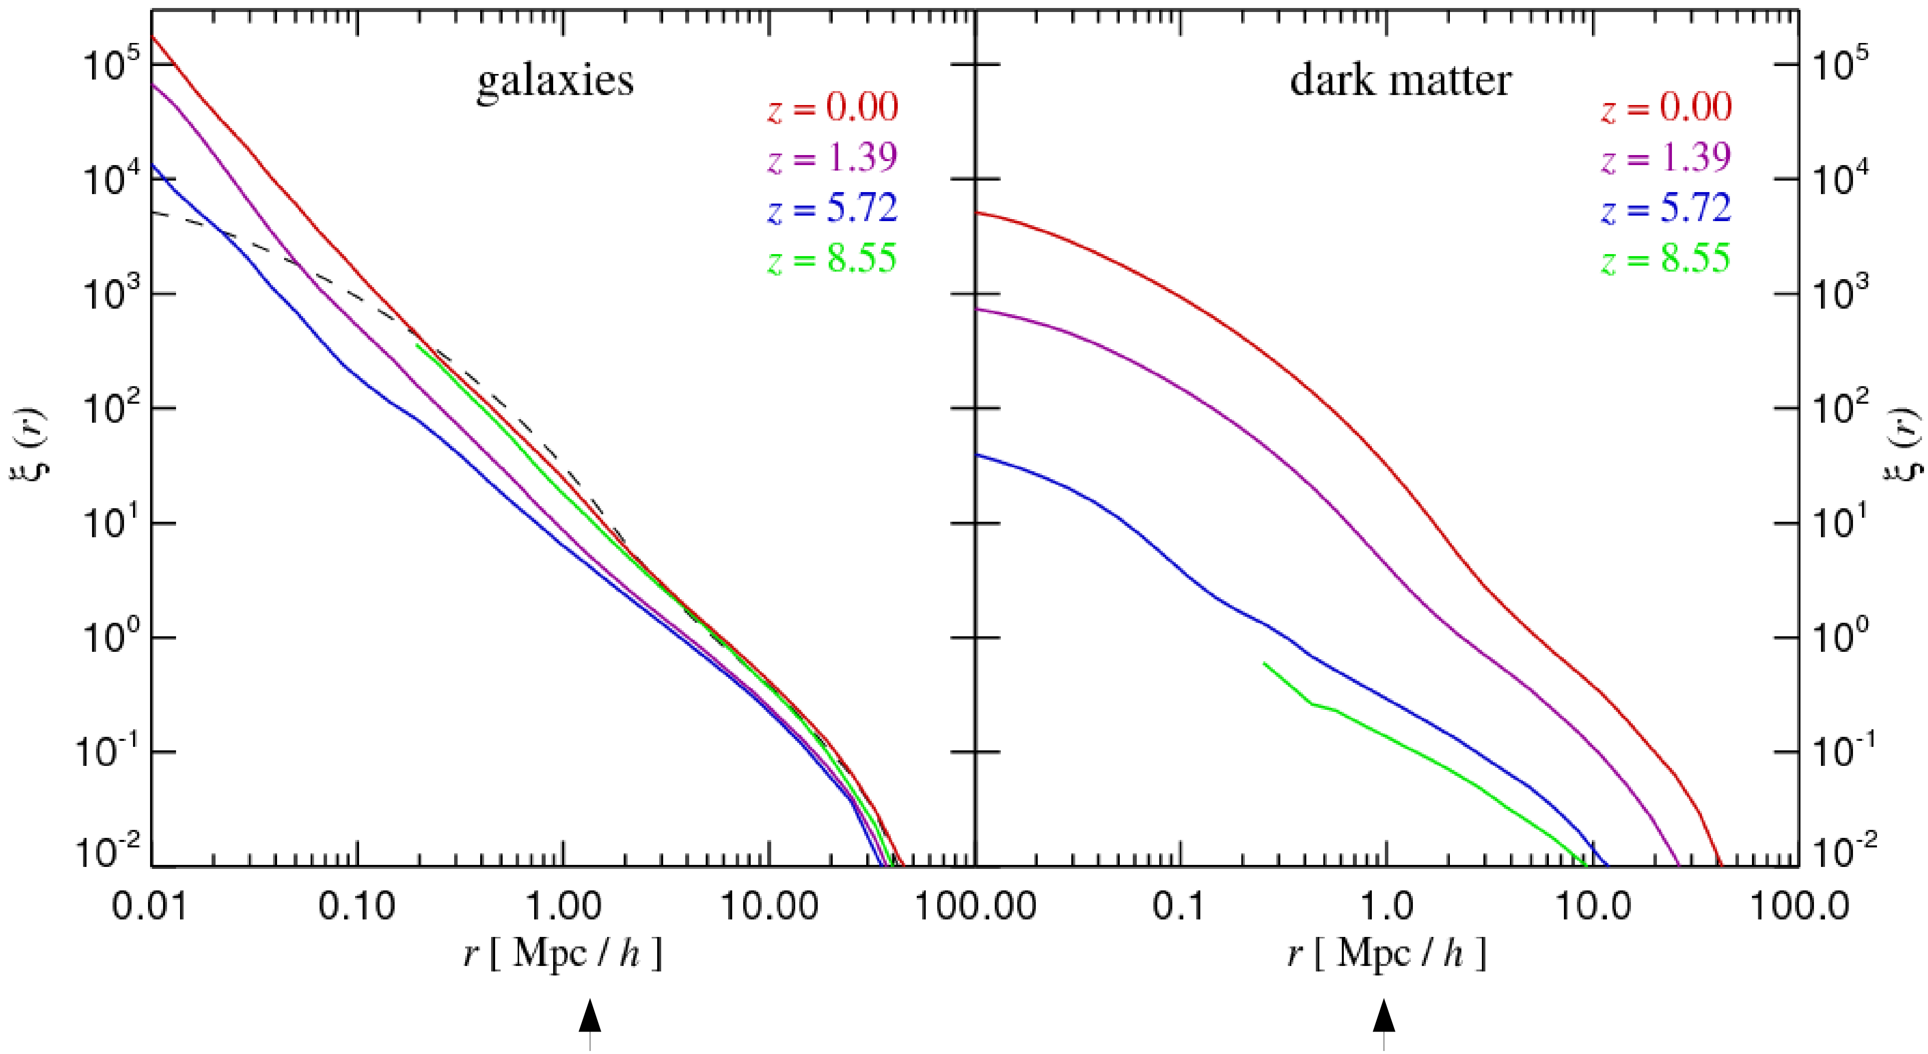

evolves little in time

evolves strongly in time

The large-scale clustering pattern of halos and galaxies is already imprinted on the initial conditions

TIME EVOLUTION OF THE MATTER AND GALAXY DISTRIBUTION

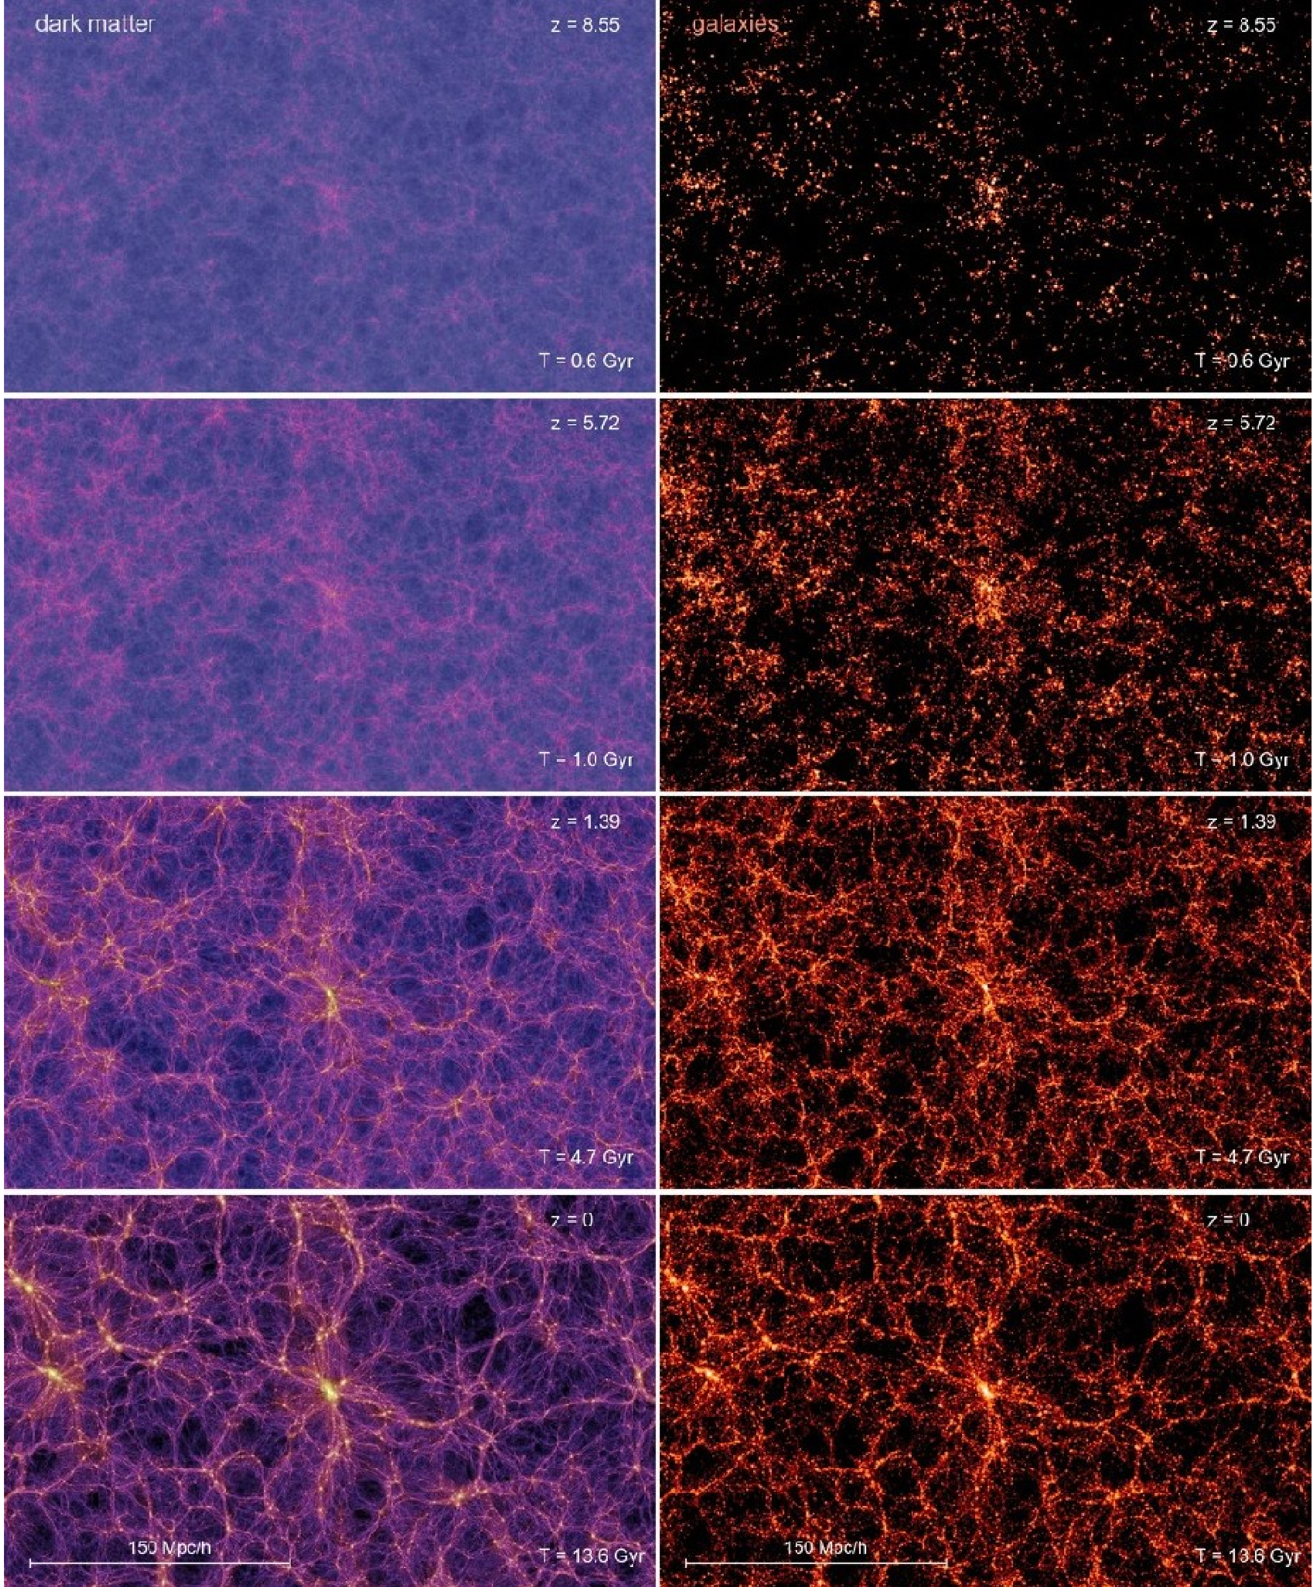

The baryonic wiggles remain visible in the galaxy distribution down to low redshift and may serve as a "standard ruler" to constrain dark energy

DARK MATTER AND GALAXY POWER SPECTRA FROM THE MILLENNIUM SIMULATION IN THE REGION OF THE WIGGLES

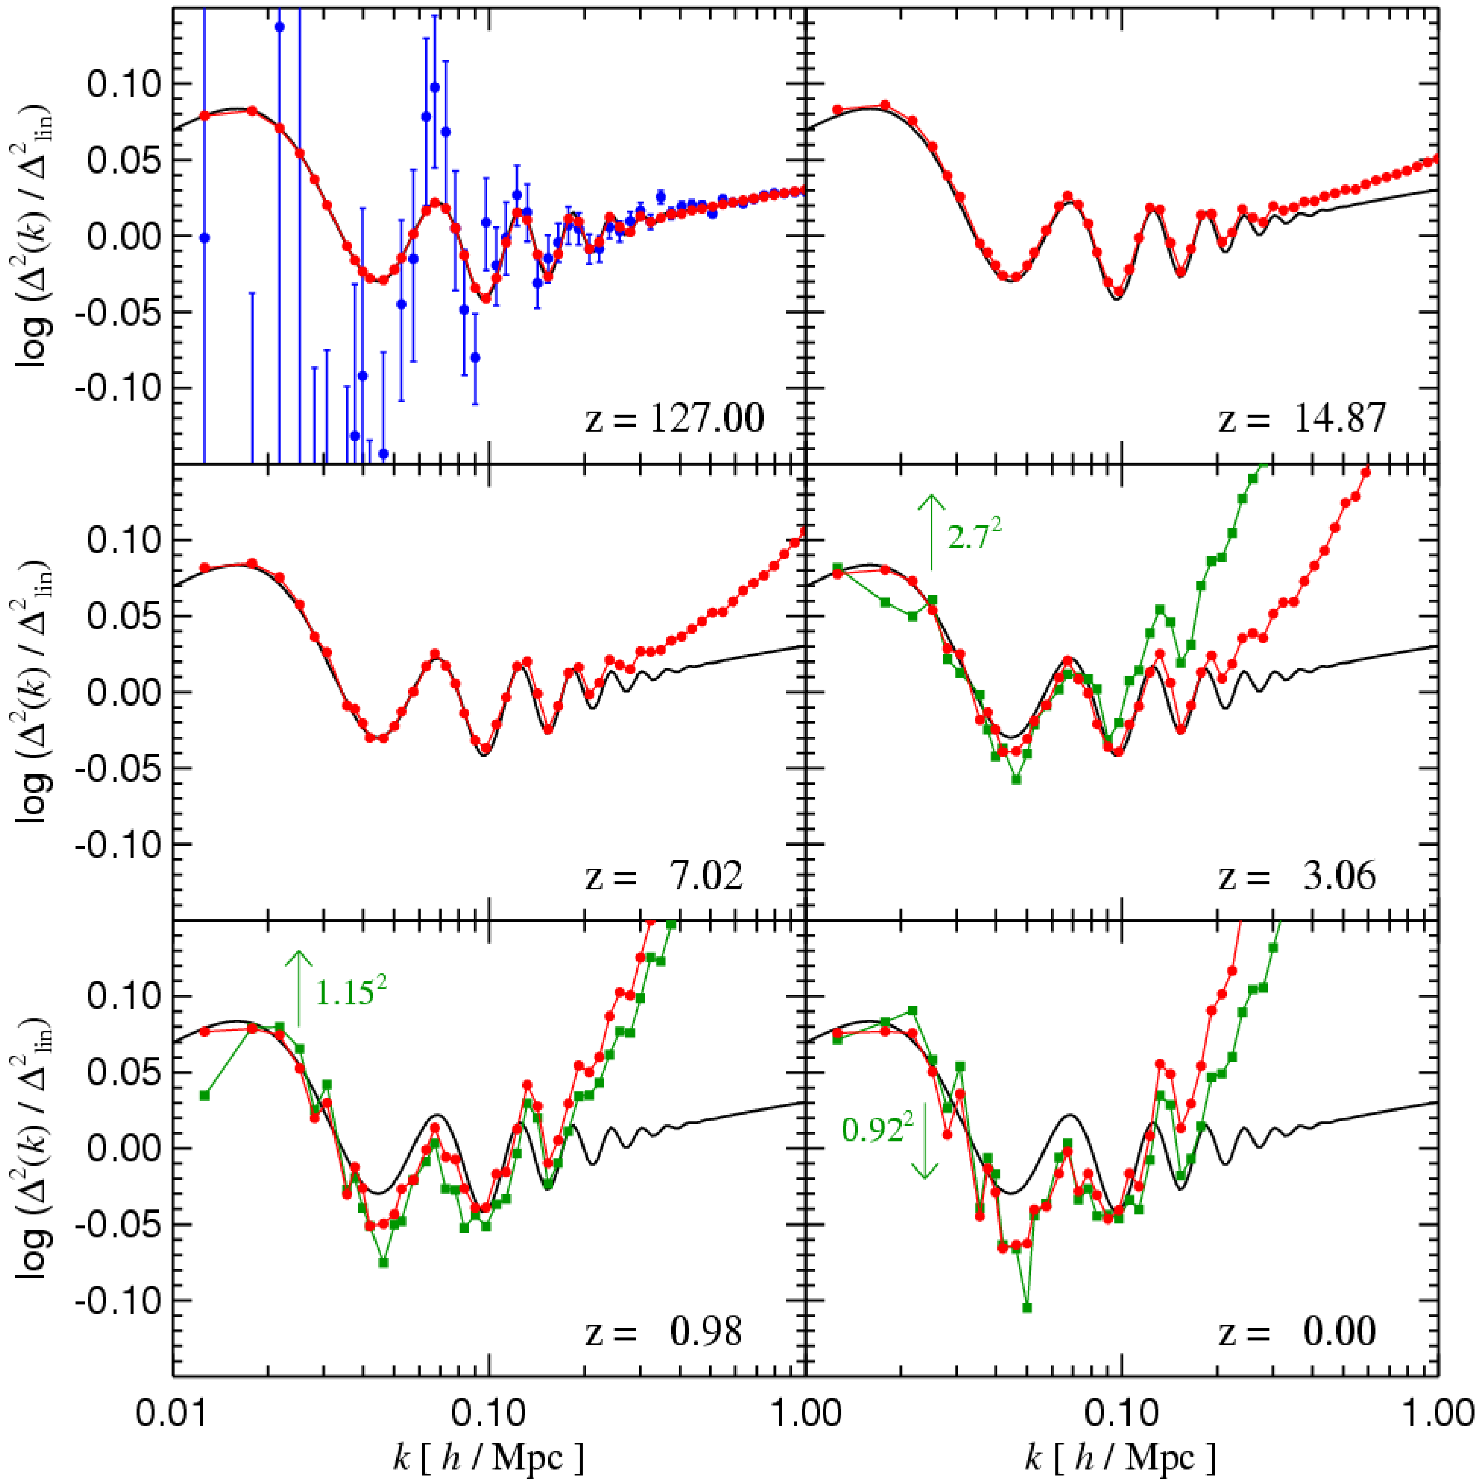

Springel et al. (2005)

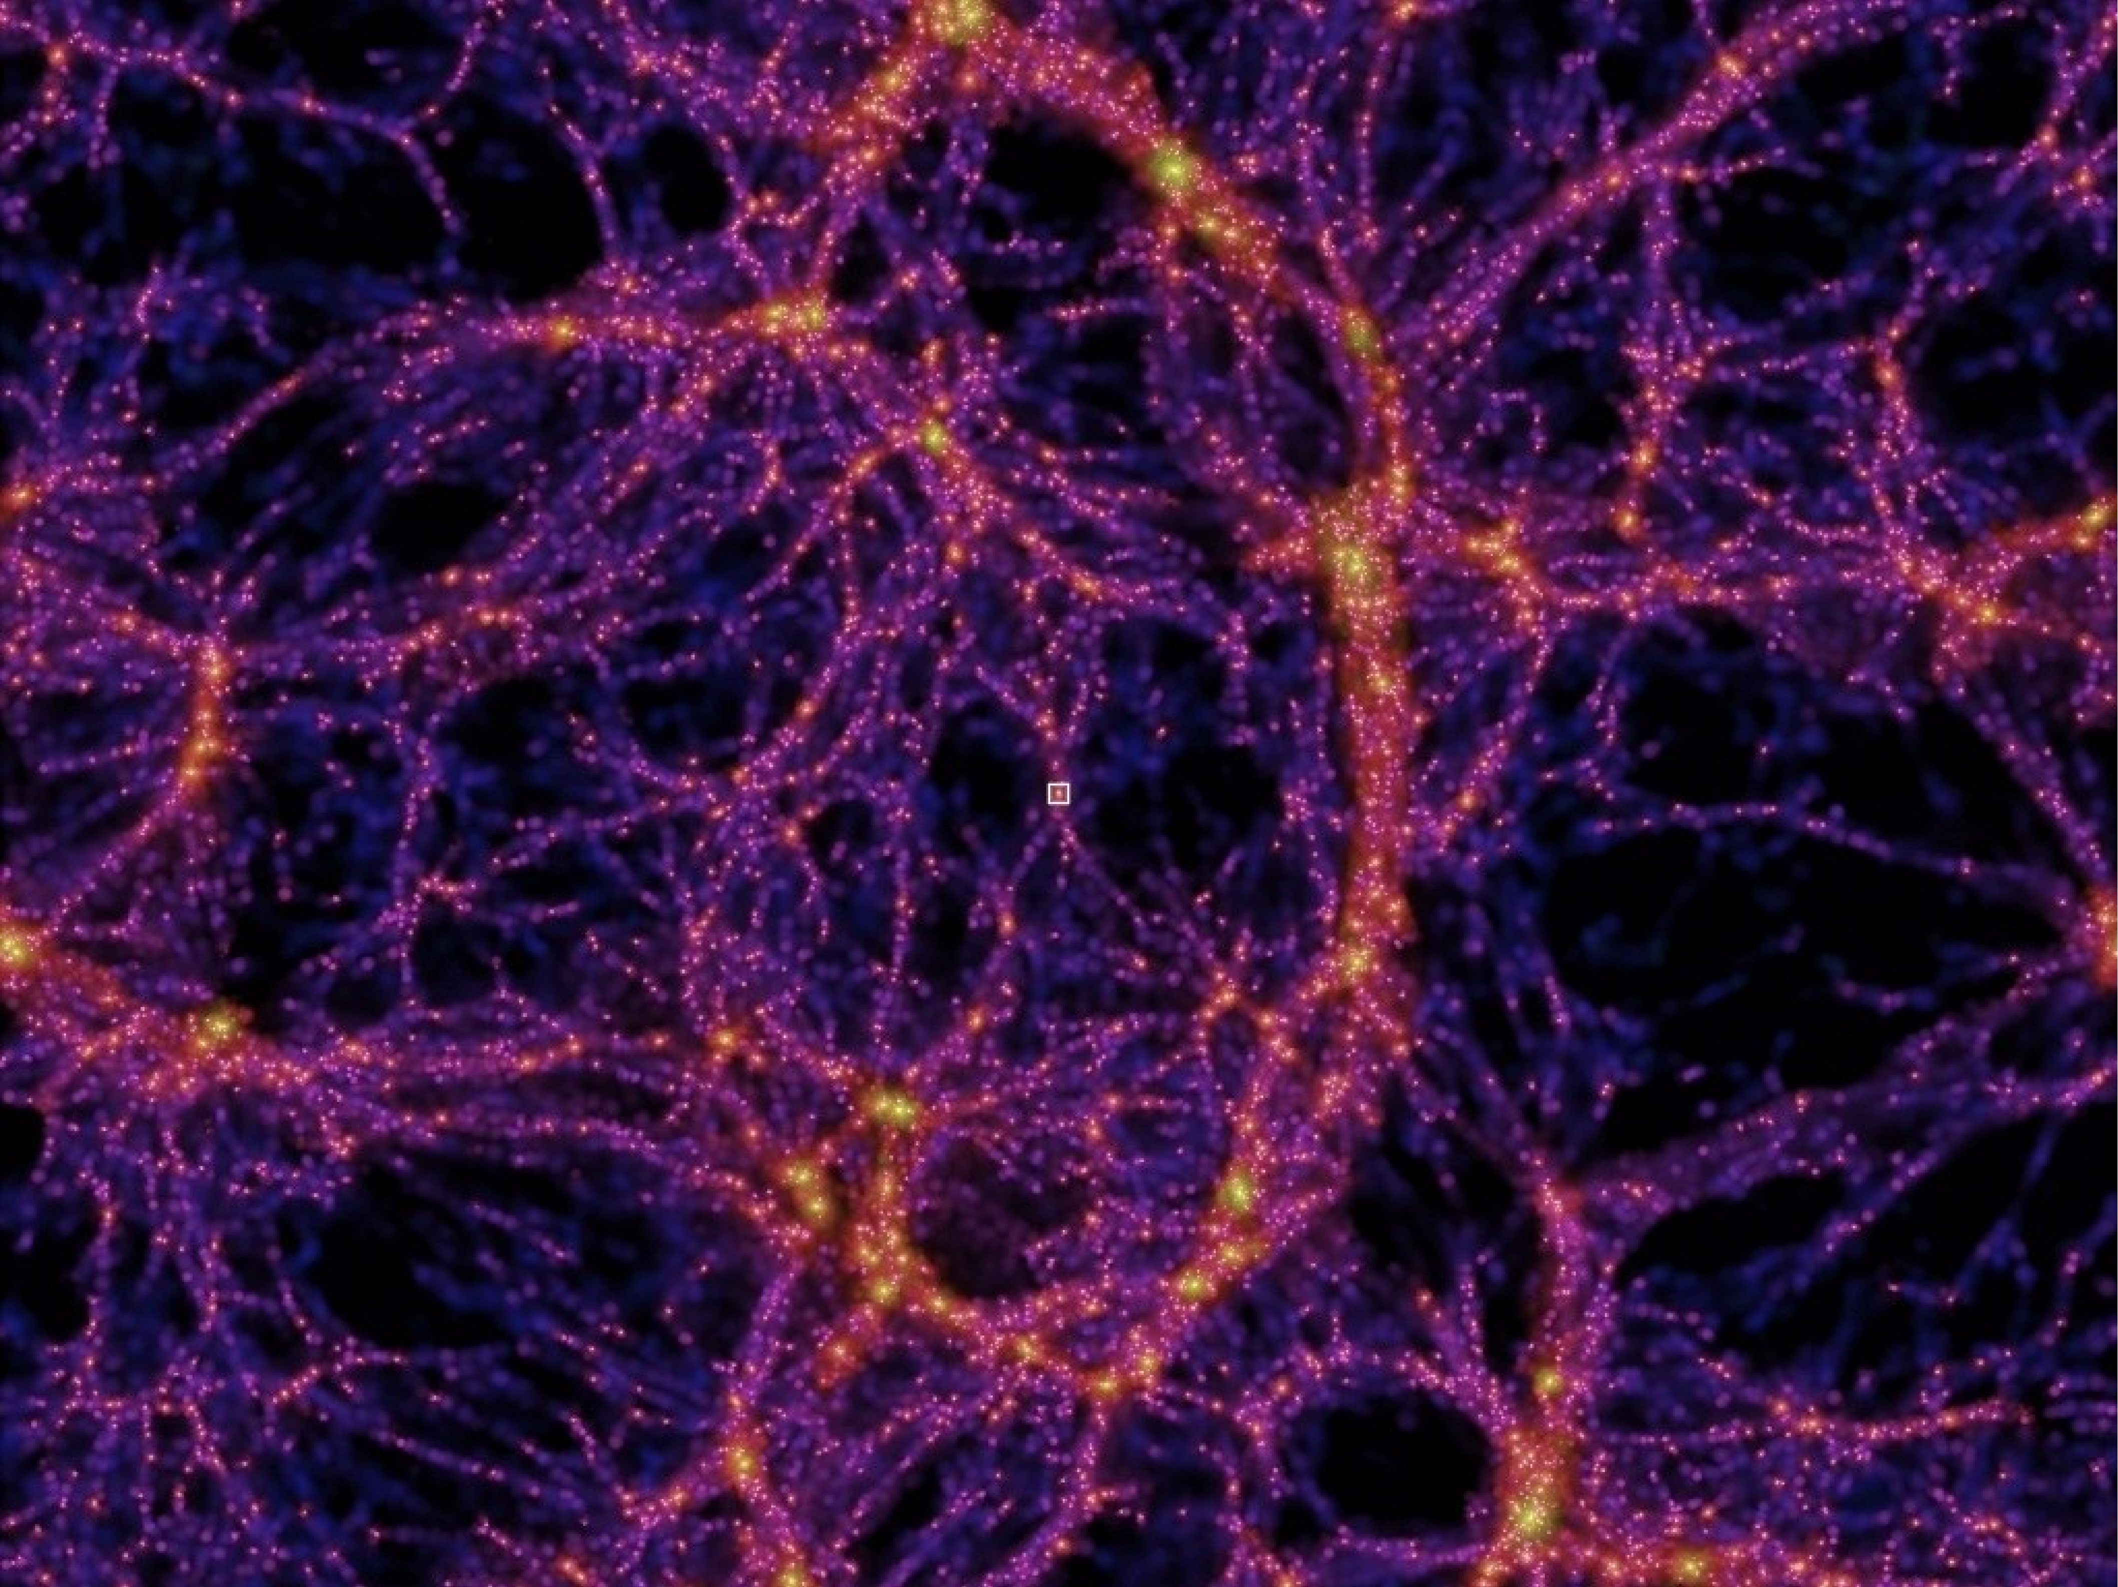

Structure of the central cusp

Spherically averaged density profiles of dark matter halos have a nearly universal shape

DENSITY PROFILE AS A FUNCTION OF RADIUS

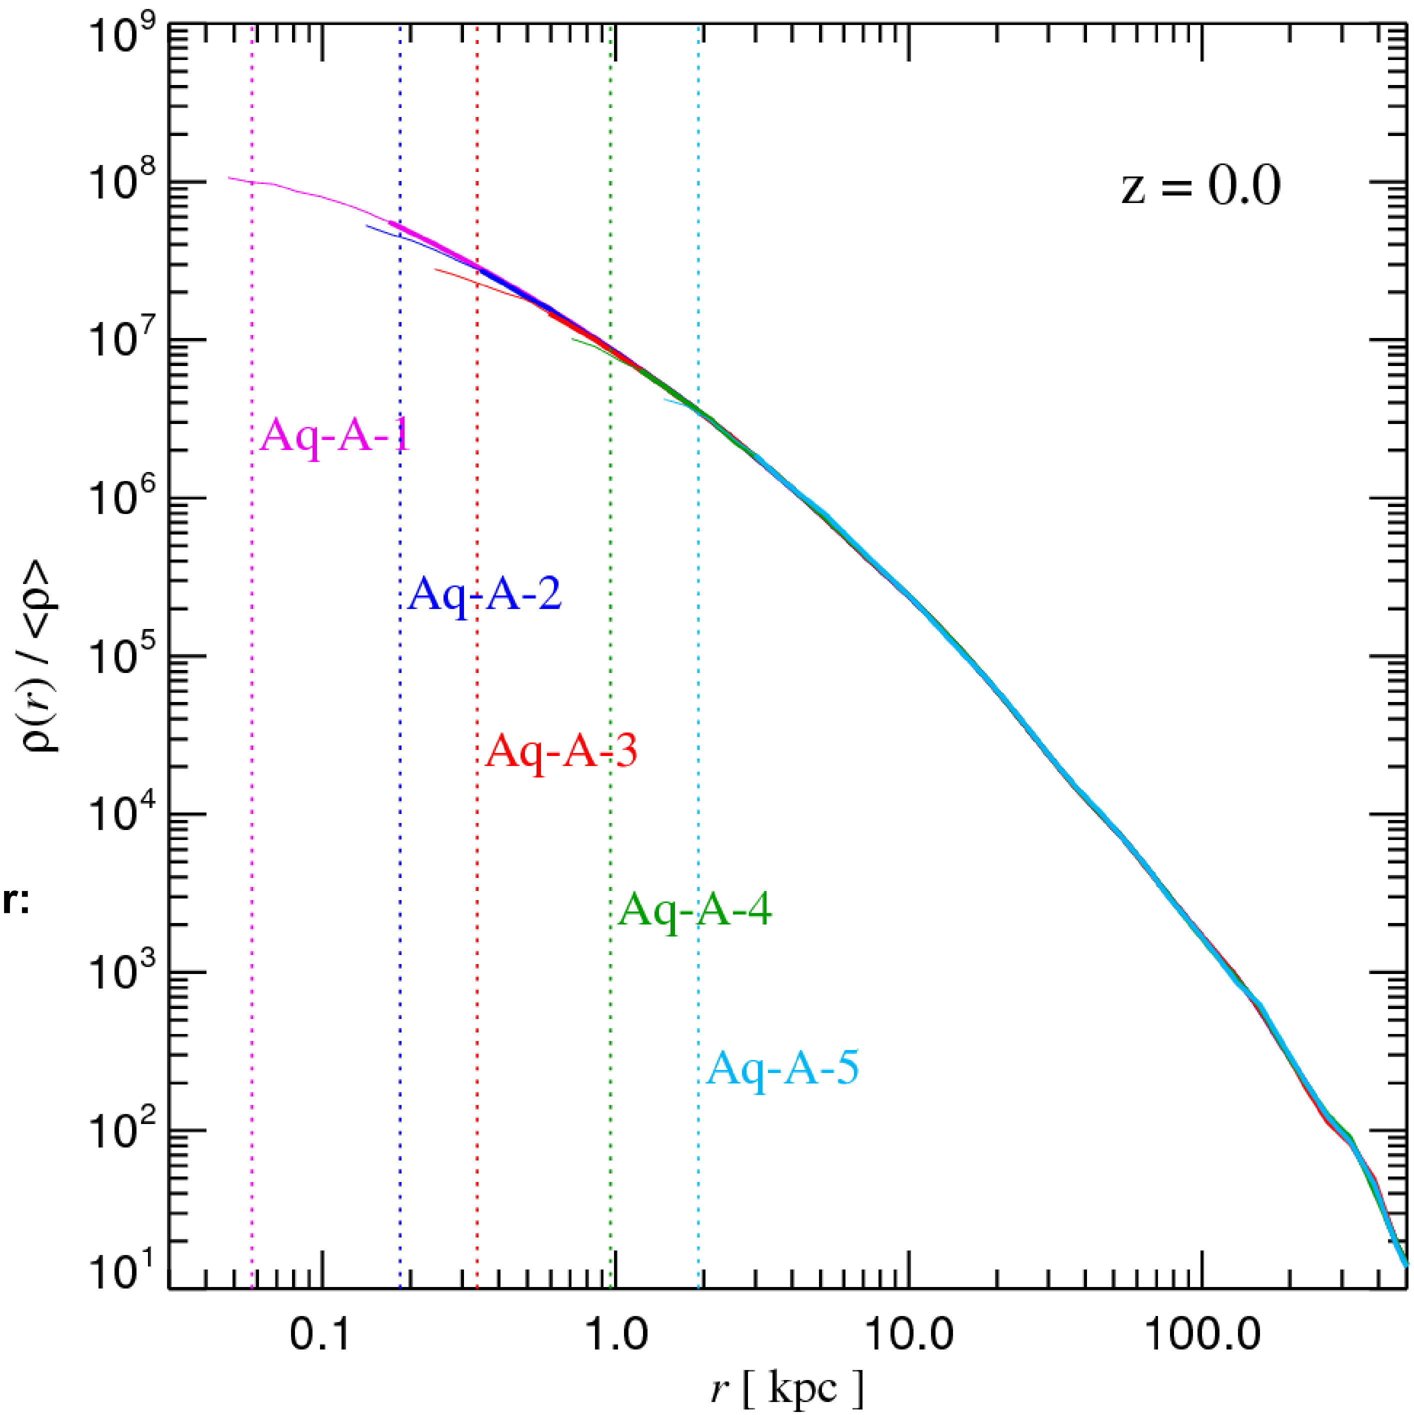

Fundamental importance for:

- Rotation curve of galaxies
- Internal structure of galaxy clusters
- Gravitational lensing
- DM annihilation
- Galaxy mergers

The logarithmic slope of the density profile does not show asymptotic behavior towards the core

SLOPE OF THE DENSITY PROFILE AS A FUNCTION OF RADIUS

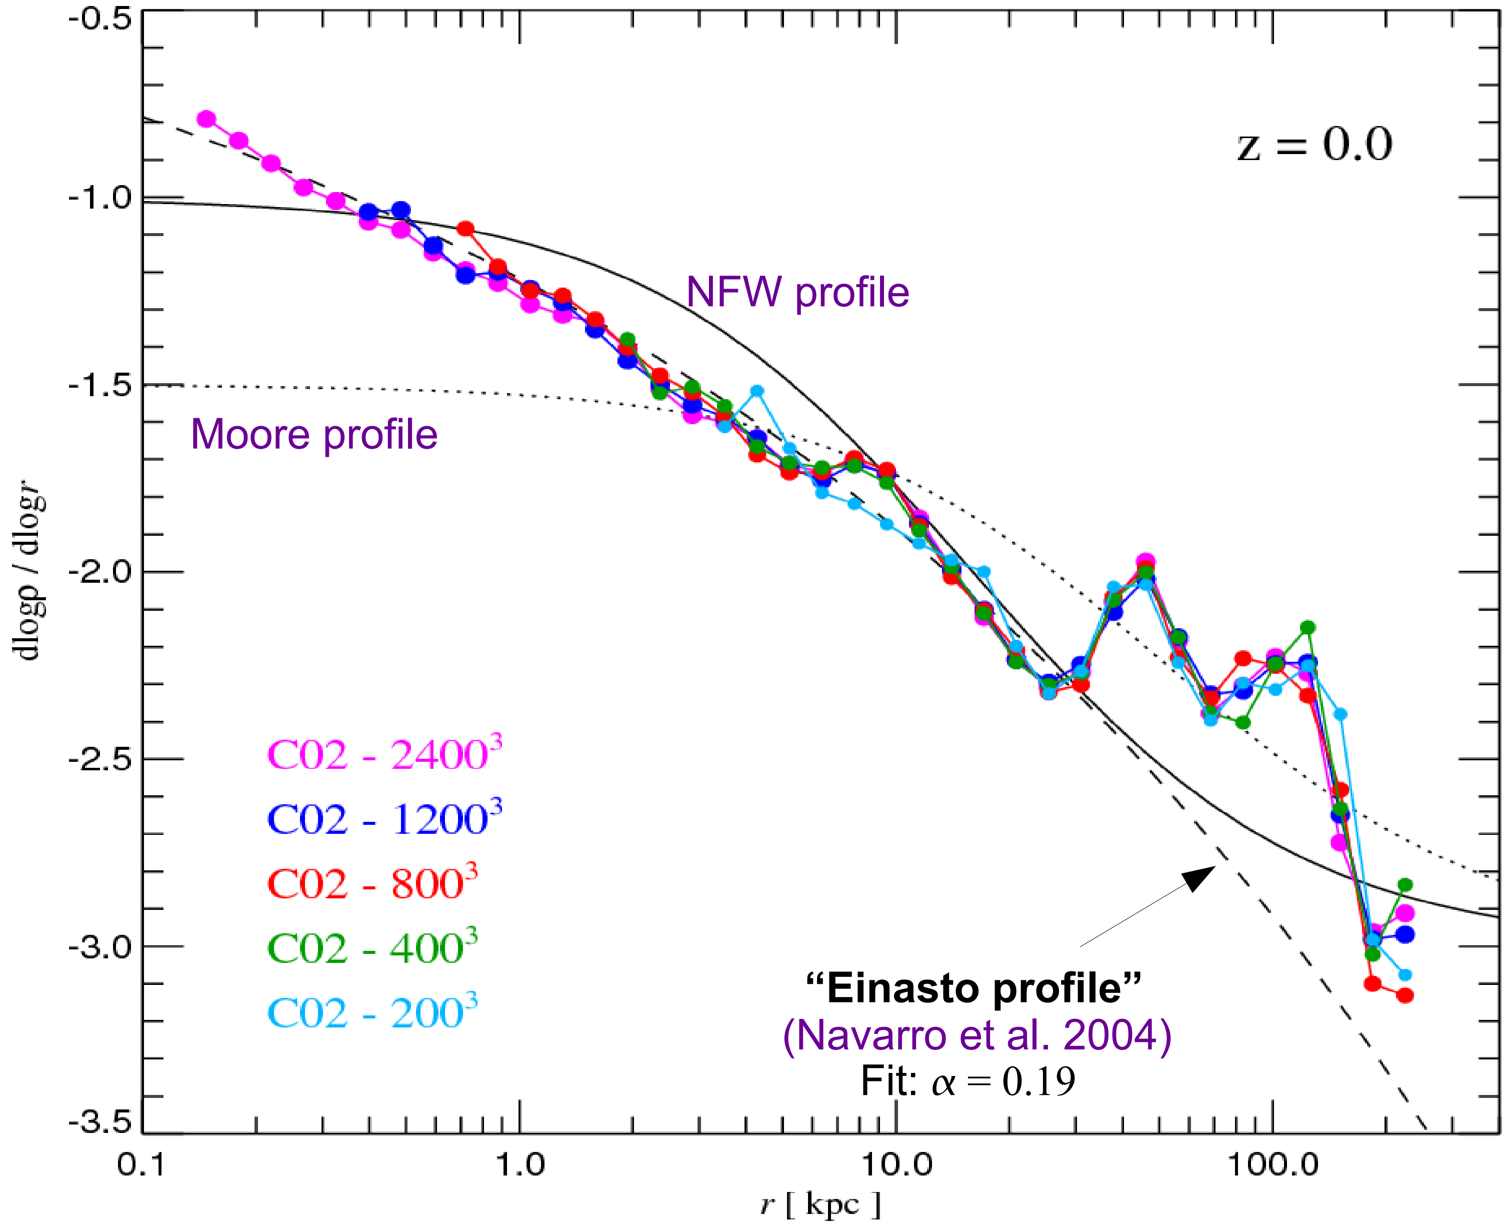

$$\frac{d\log\rho}{d\log r} = -2\left(\frac{r}{r_{-2}}\right)^{\alpha}$$

A consensus on the central structure of the cups seems to be emerging

RECENT RESULTS FROM THE 'GHALO' SIMULATION OF THE ZURICH GROUP

Stadel et al. (2009)

“The logarithmic slope of the radial density profile is close to a power law, gradually turning over to a slope of -0.8 at our innermost resolved point.”

“The Einasto profile also provides an excellent fit to the density profiles of the two simulations.”

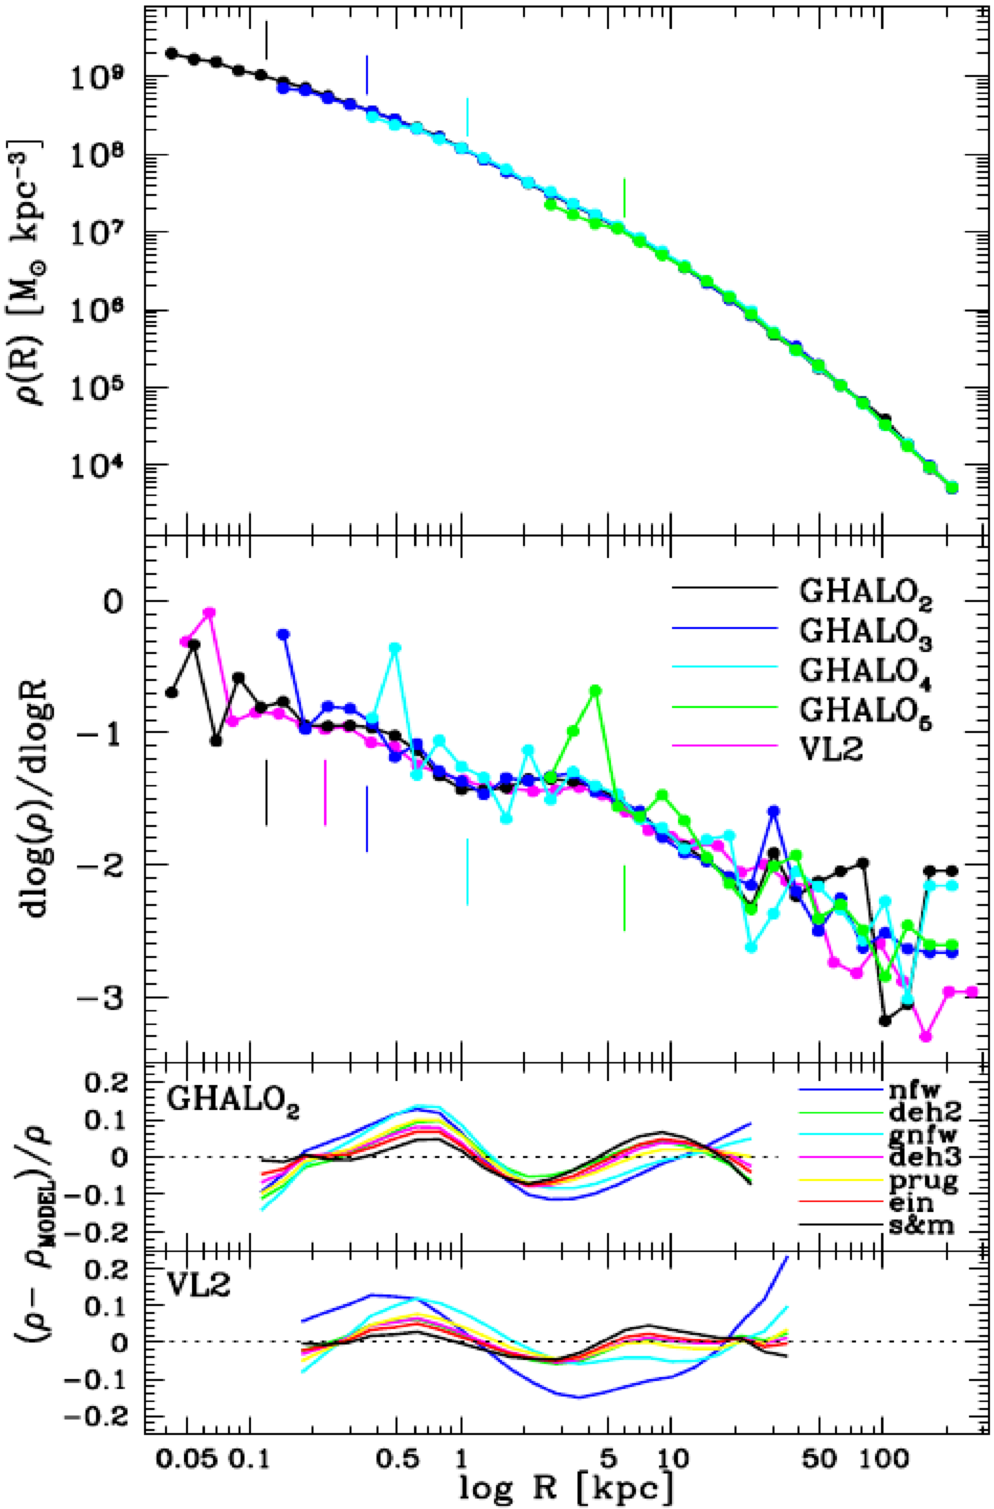

Our simulations allow us to study the convergence of **subhalo density profiles**

SPHERICALLY AVERAGED DENSITY PROFILES IN THE AQ-A HALO AT DIFFERENT RESOLUTION

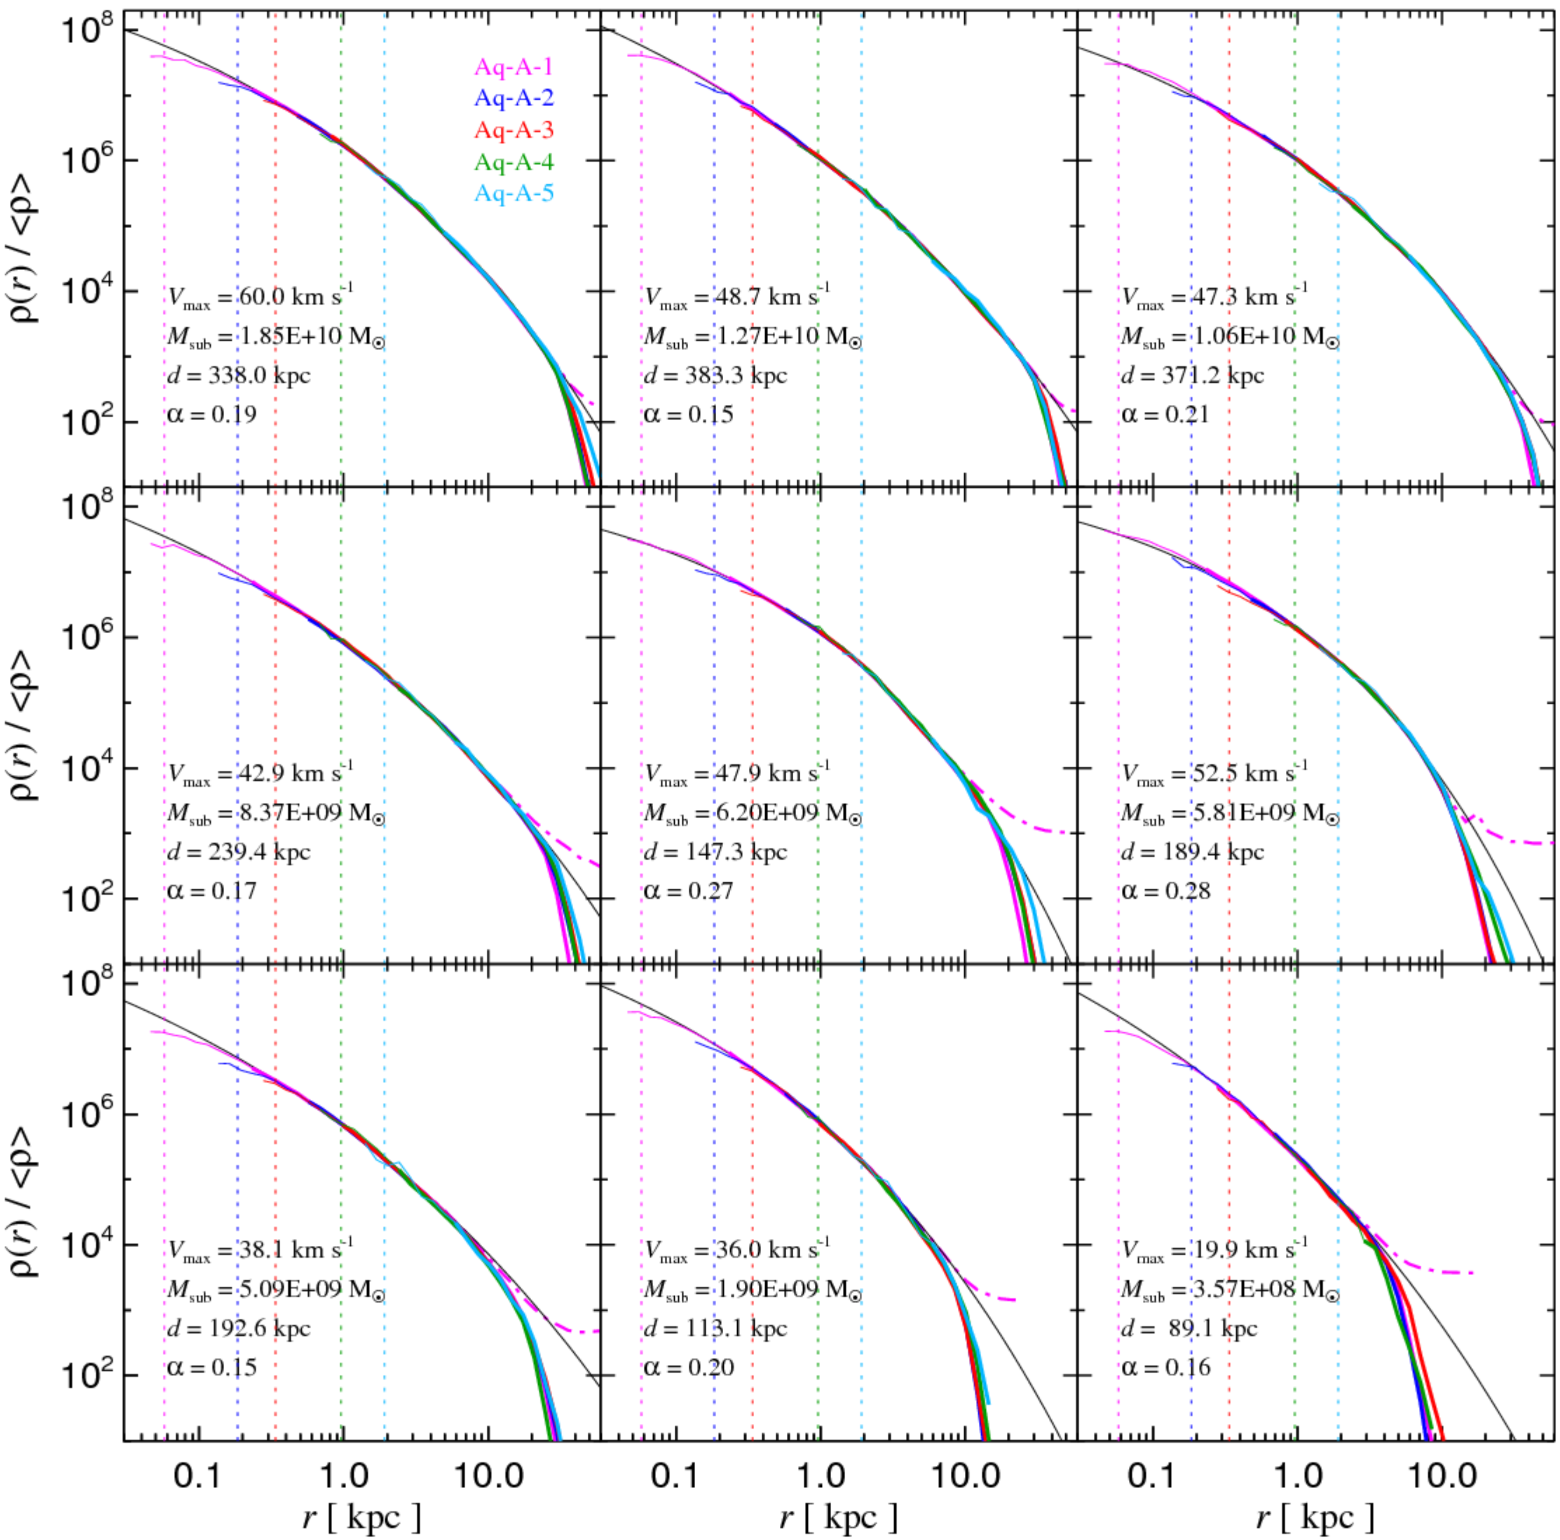

Dark matter substructure

Zooming in on dark matter halos reveals a huge abundance of dark matter substructure

DARK MATTER DISTRIBUTION IN A MILKY WAY SIZED HALO AT DIFFERENT RESOLUTION

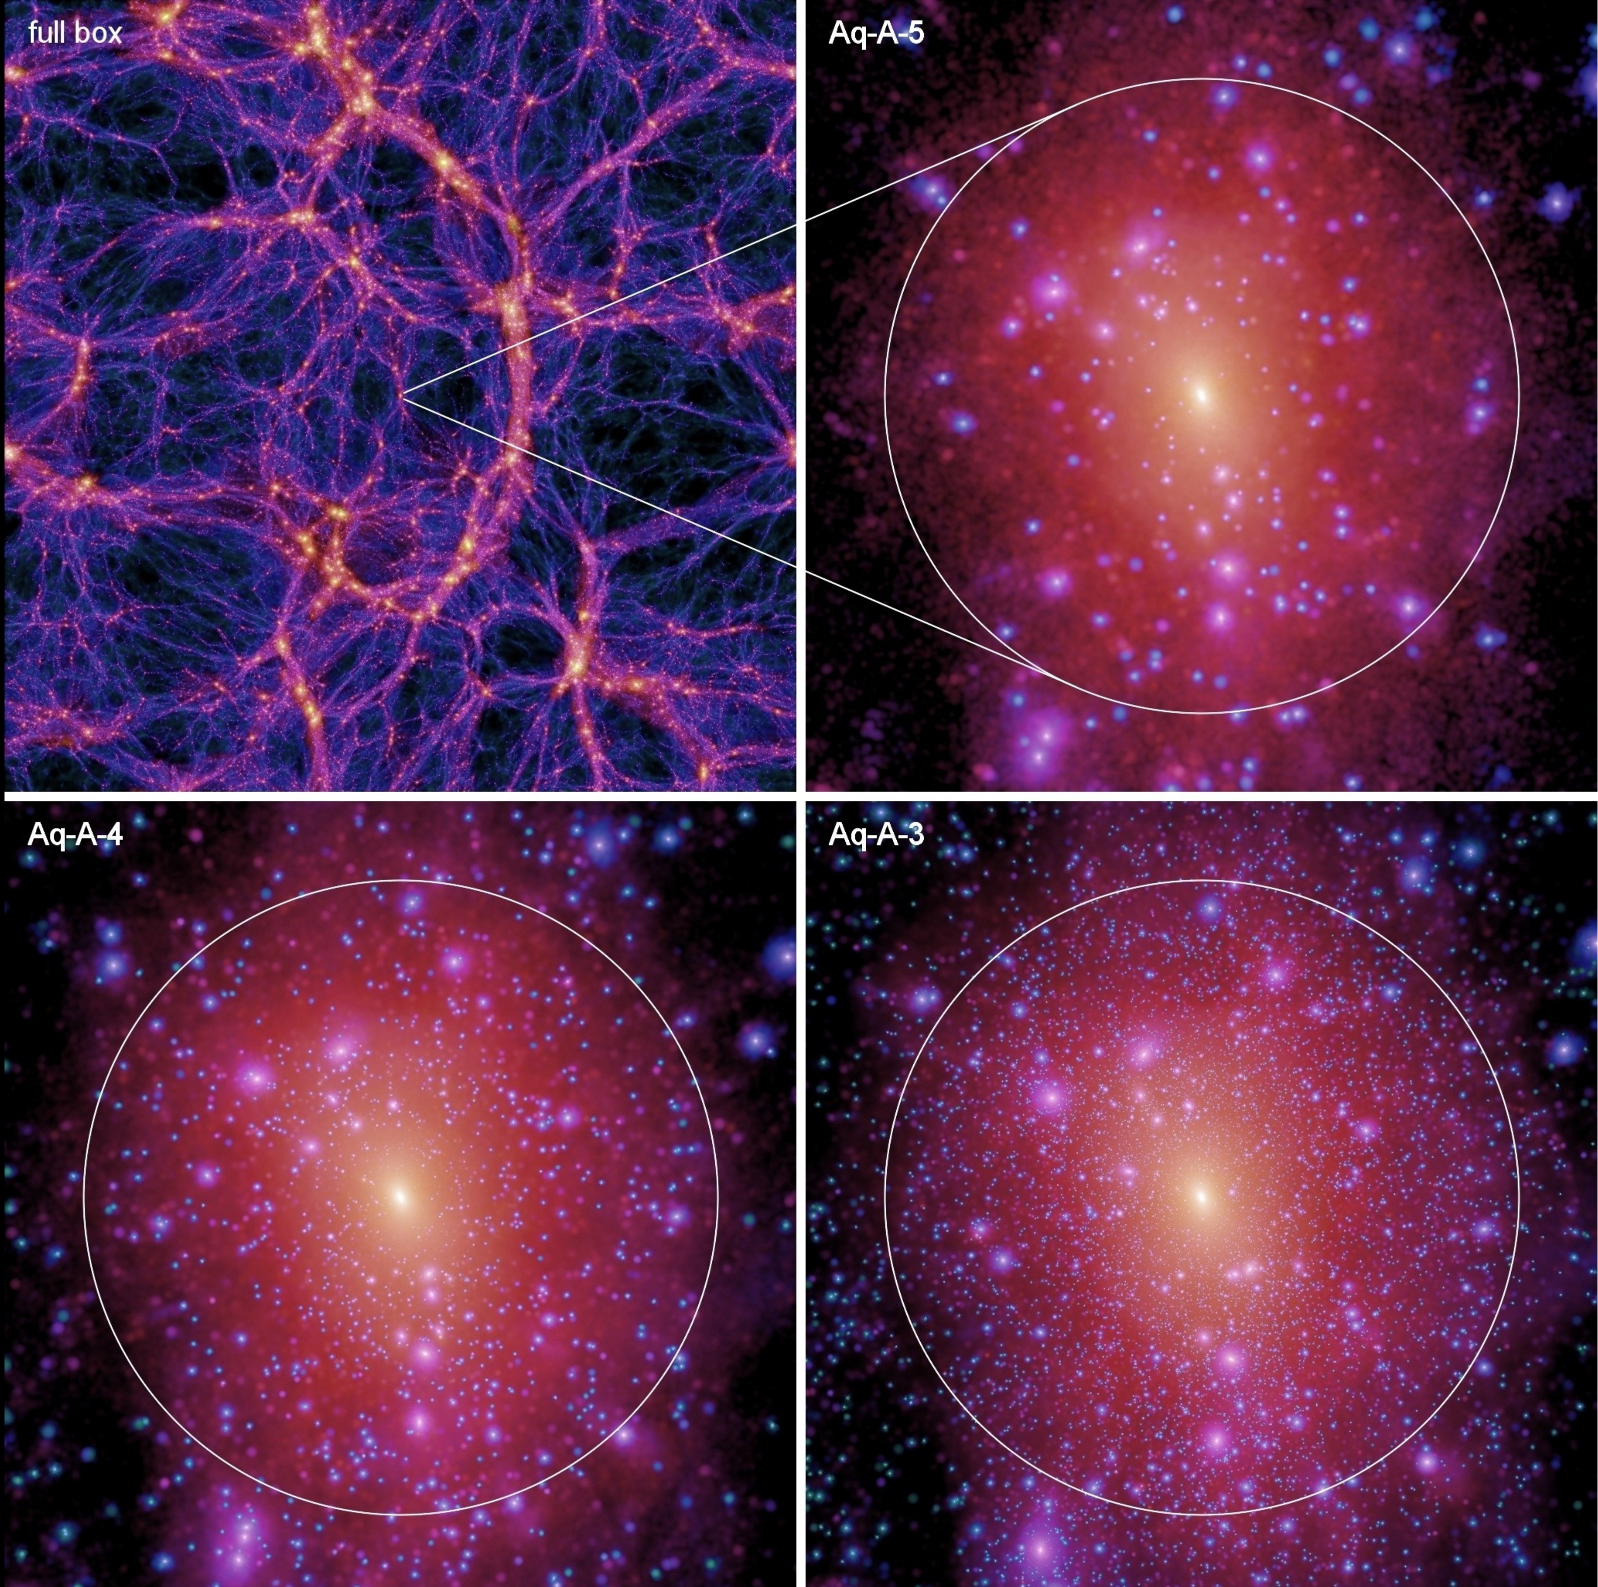

Zooming in on dark matter halos reveals a huge abundance of dark matter substructure

DARK MATTER IN A MILKY WAY SIZED HALO AT ULTRA-HIGH RESOLUTION

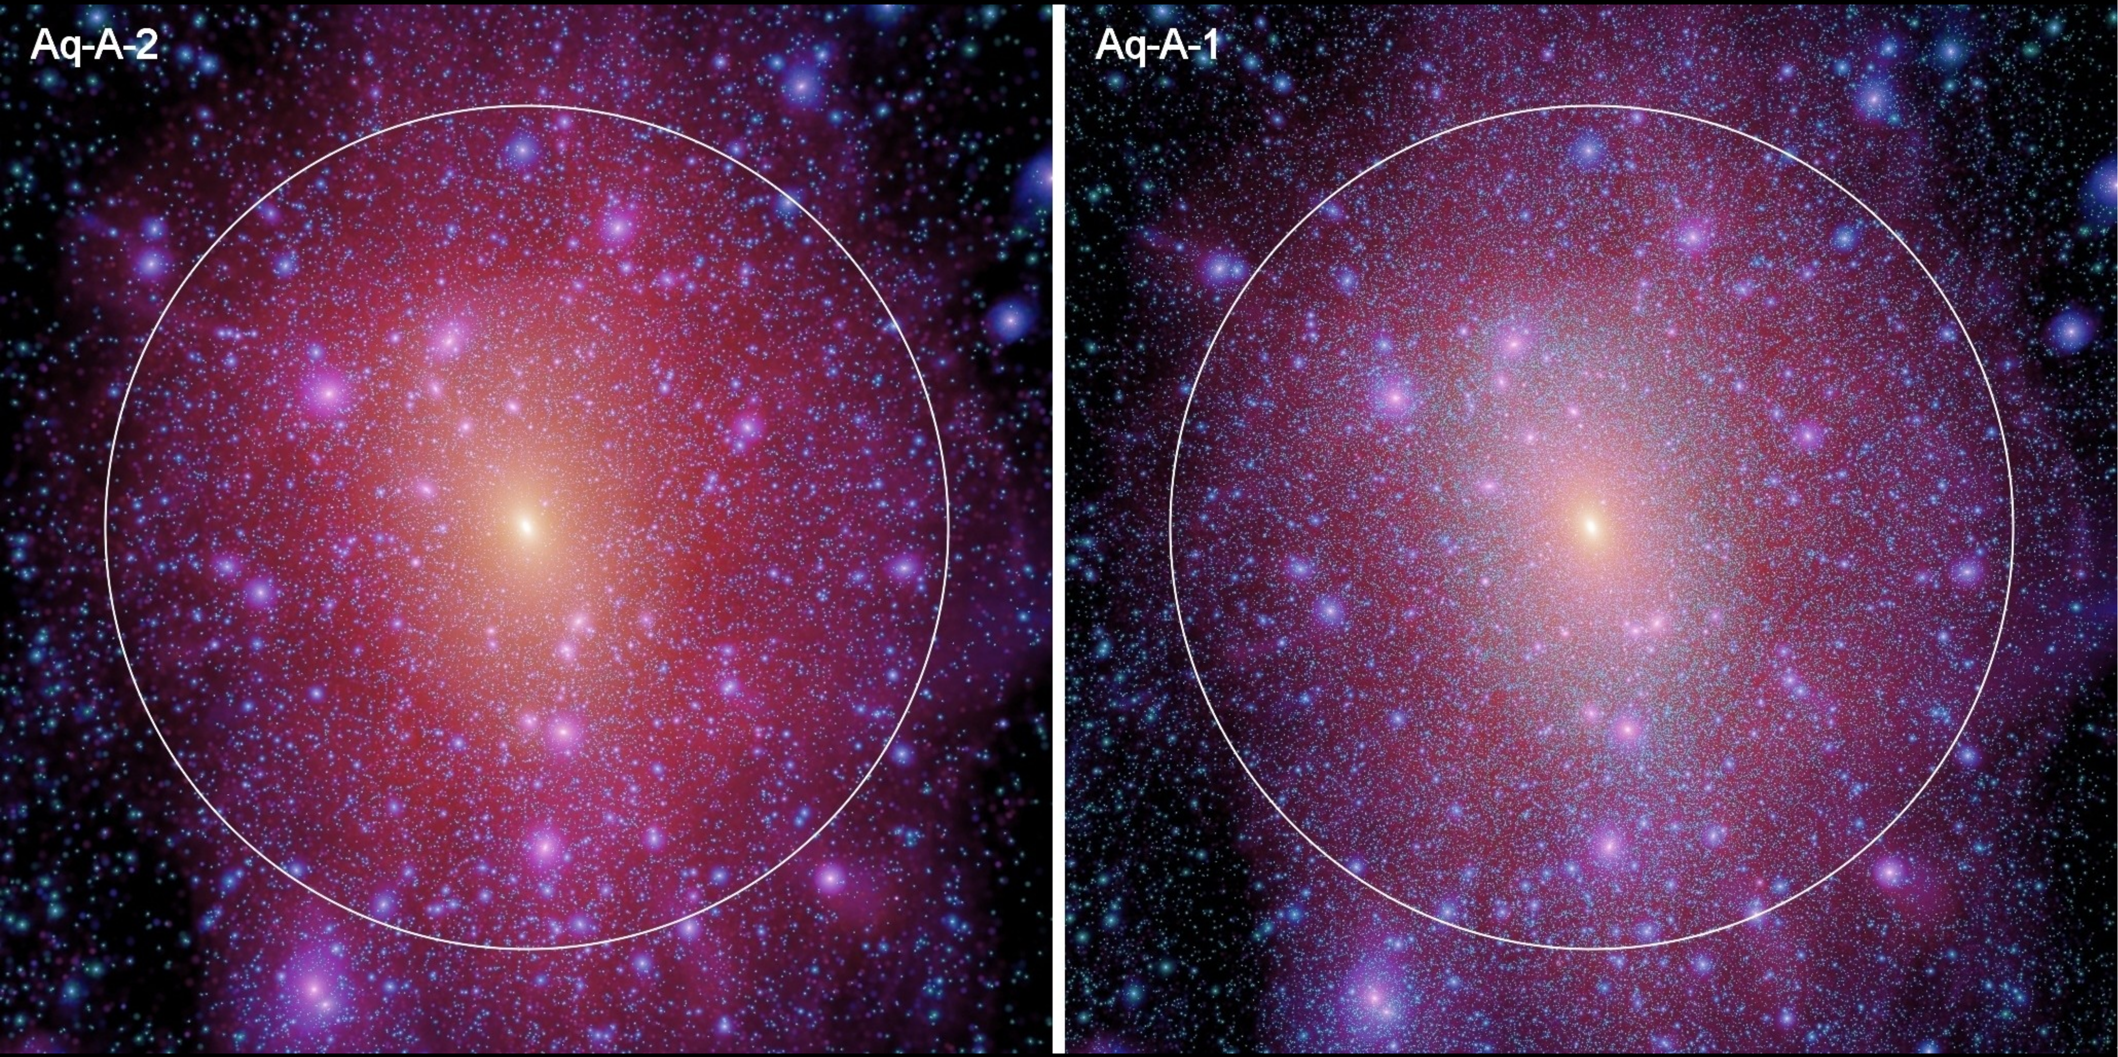

The subhalo abundance per unit halo mass is surprisingly uniform

VELOCITY FUNCTION IN OUR DIFFERENT HALOS

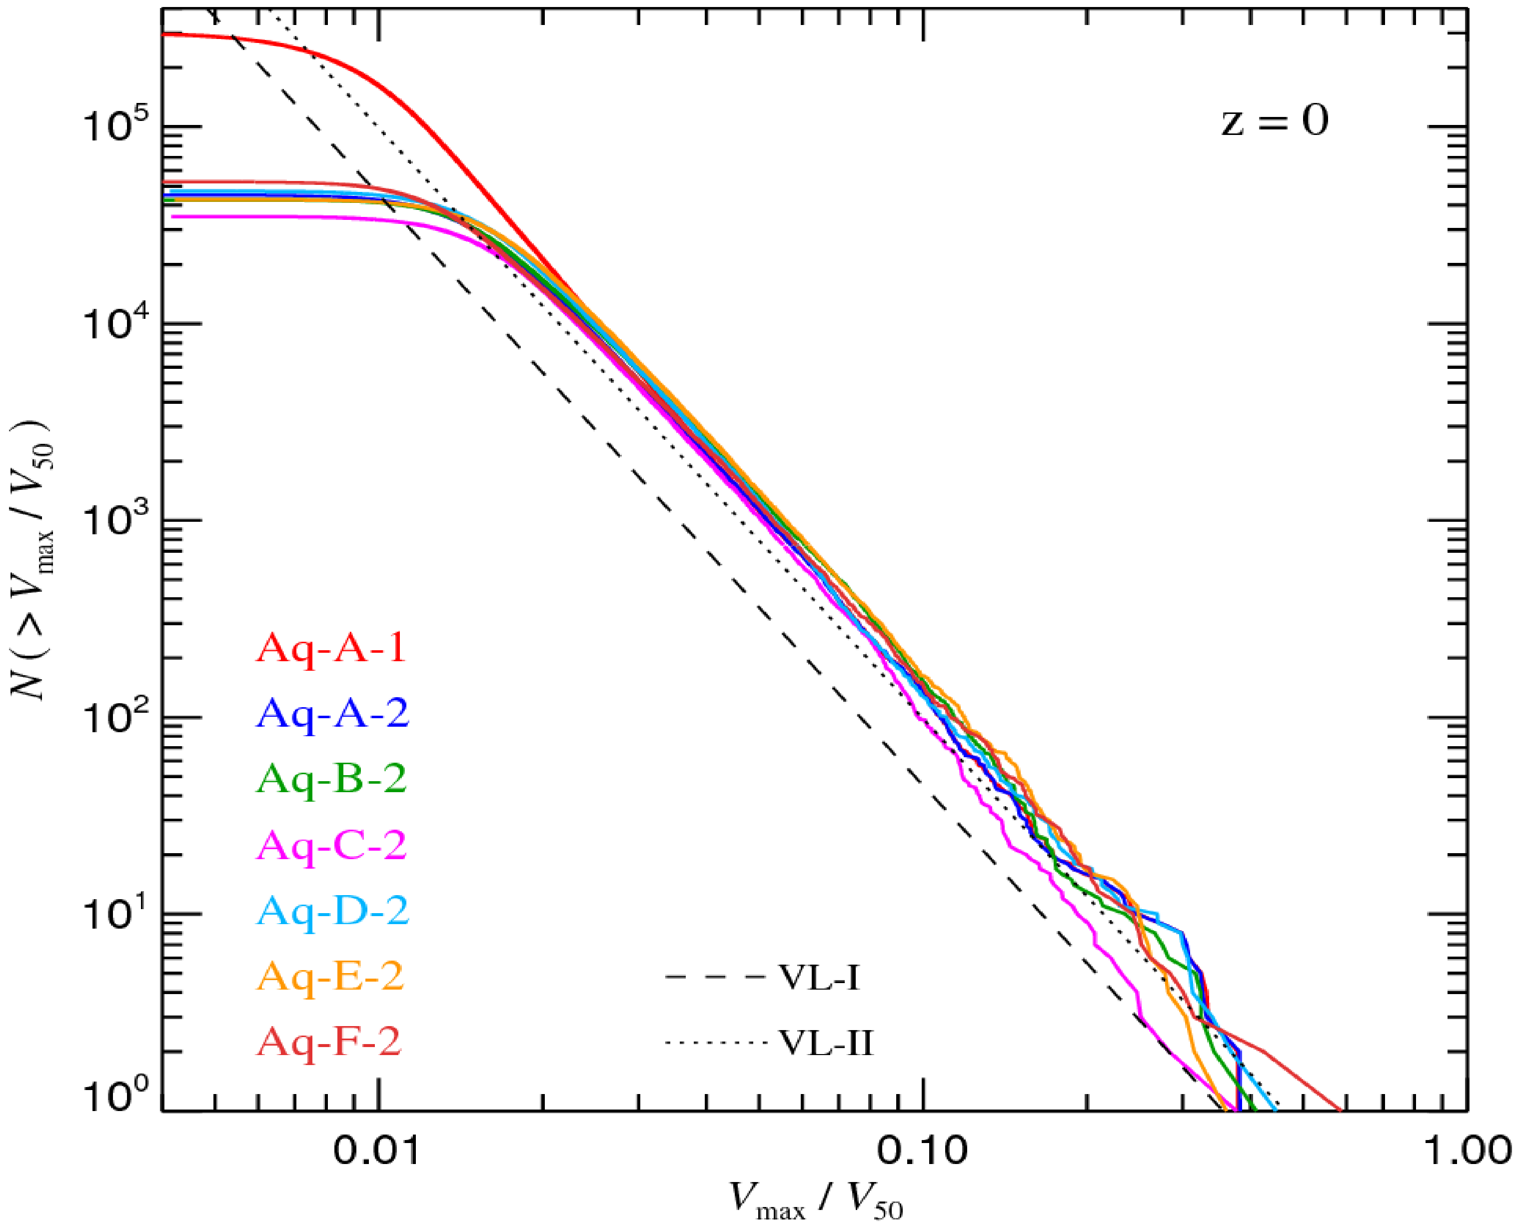

The cumulative mass fraction in **resolved substructures** reaches about 12-13%, we expect up to ~18% down to the thermal limit

FRACTION OF MASS IN SUBSTRUCTURES AS A FUNCTION OF MASS LIMIT

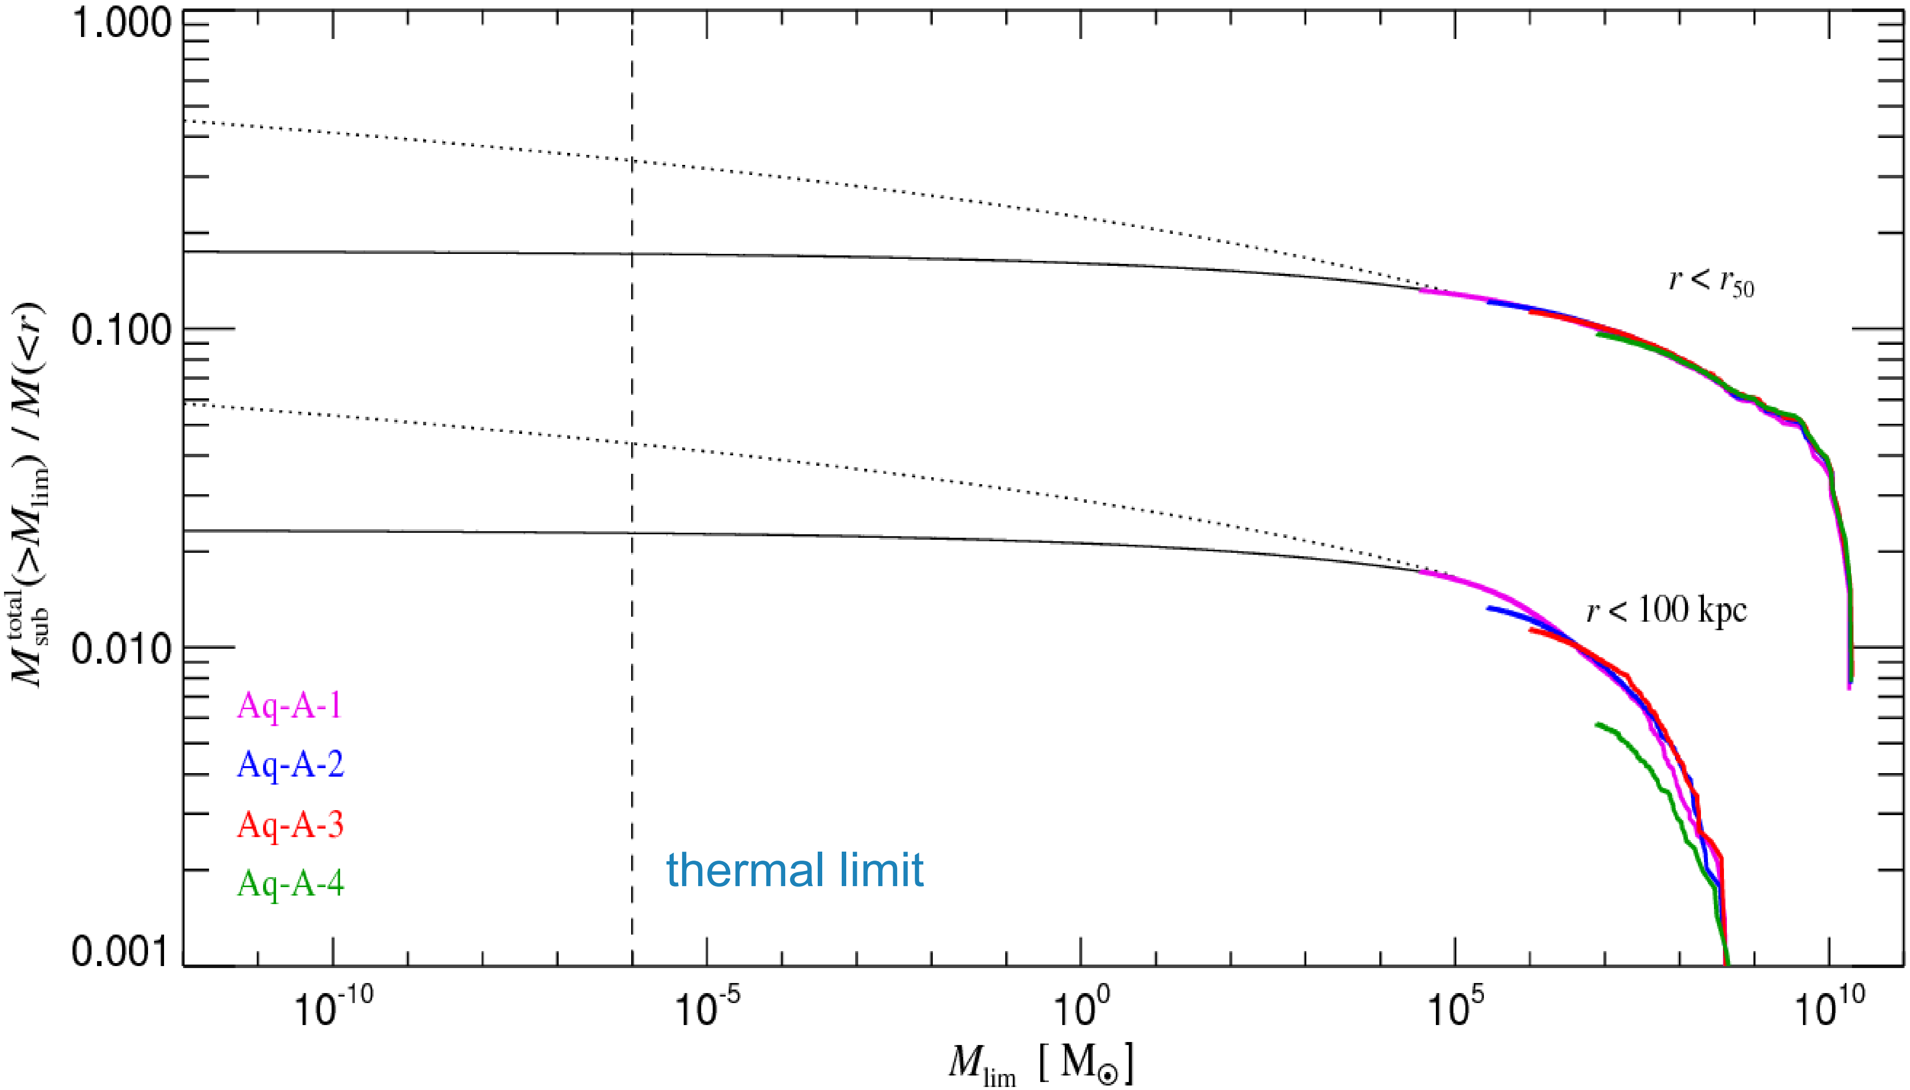

The radial distribution of substructures is strongly antibiased relative to all dark matter, and independent of subhalo mass

RADIAL SUBSTRUCTURE DISTRIBUTION IN Aq-A-1

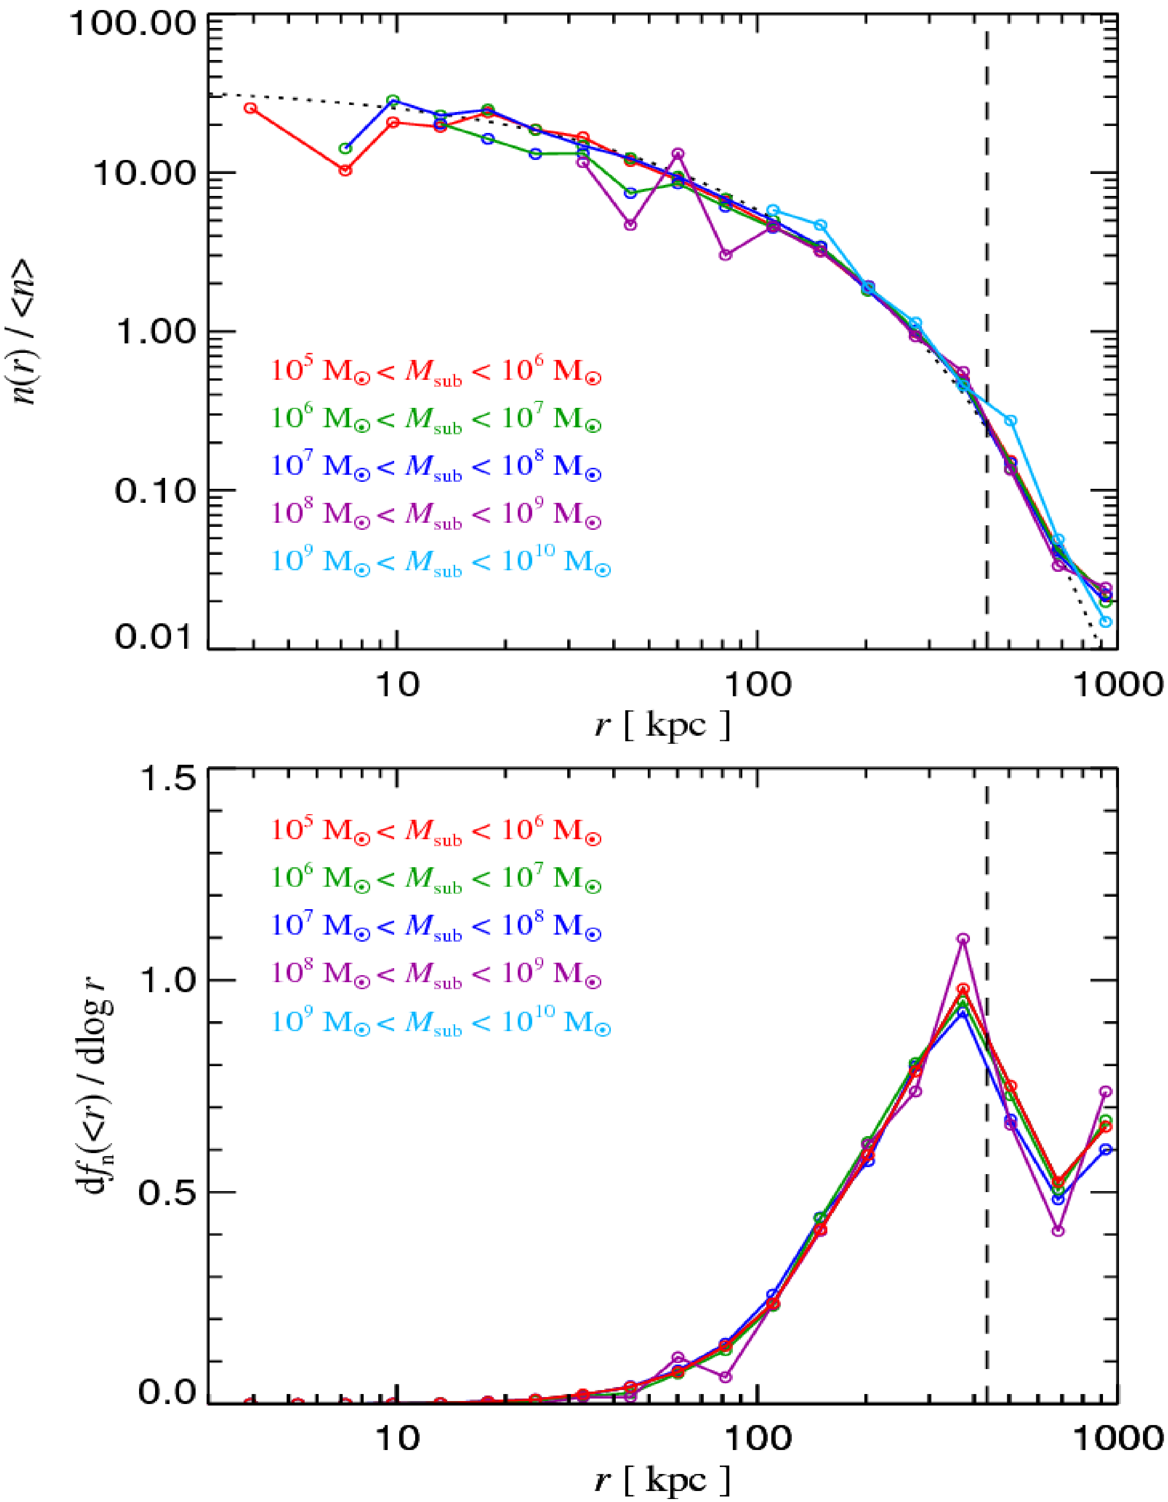

Most subhalos are at large radii, subhalos are more effectively destroyed near the centre

Subhalos are far from the Sun

see also Diemand et al. (2007, 2008)

The local mass fraction in substructures is a strong function of radius

MASS FRACTION IN SUBSTRUCTURES AS A FUNCTION OF RADIUS IN HALO AQ-A

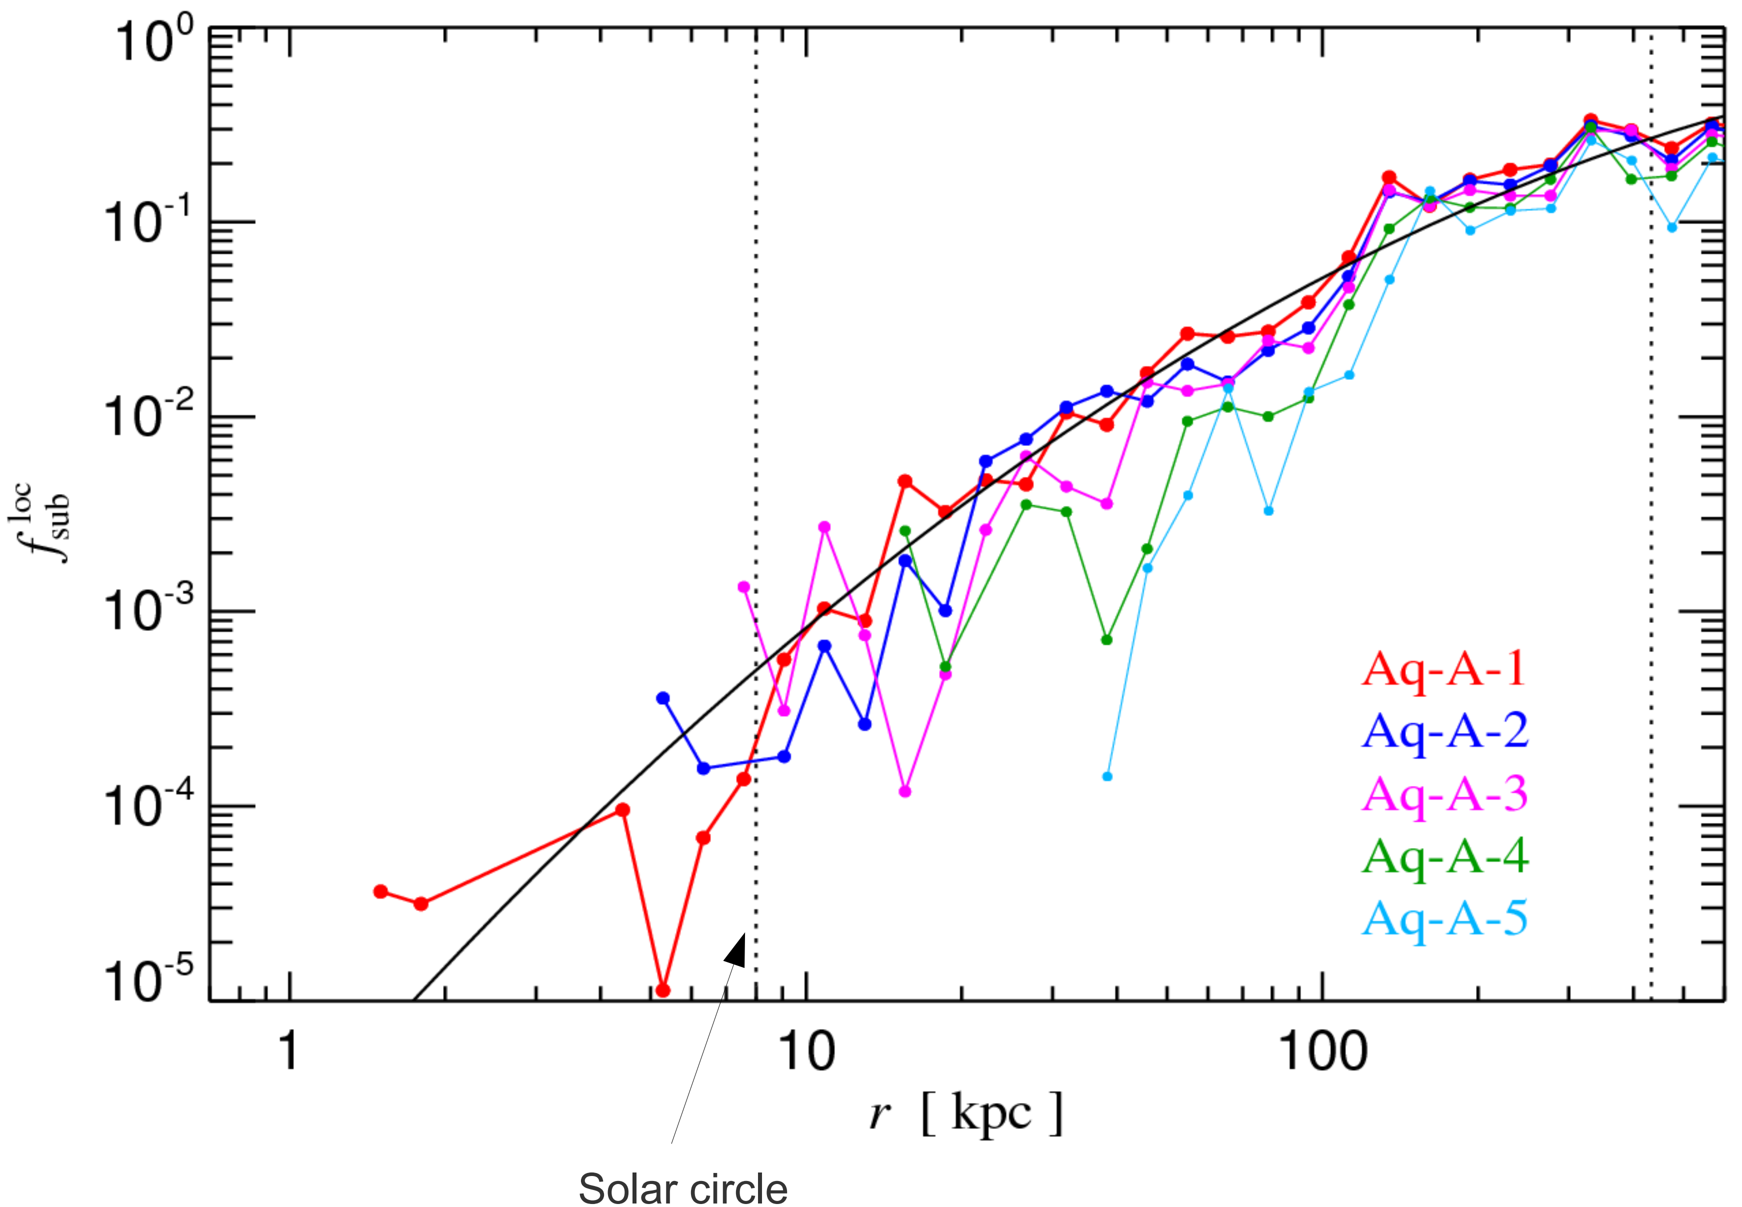

Dark matter annihilation predictions

Is the dark matter annihilation flux boosted significantly by dark matter substructures?

SIMULATED ALL-SKY MAP OF THE DM ANNIHILATION FLUX AROUND THE SUN IN THE MILKY WAY

$L \propto \rho^2 \, dV$

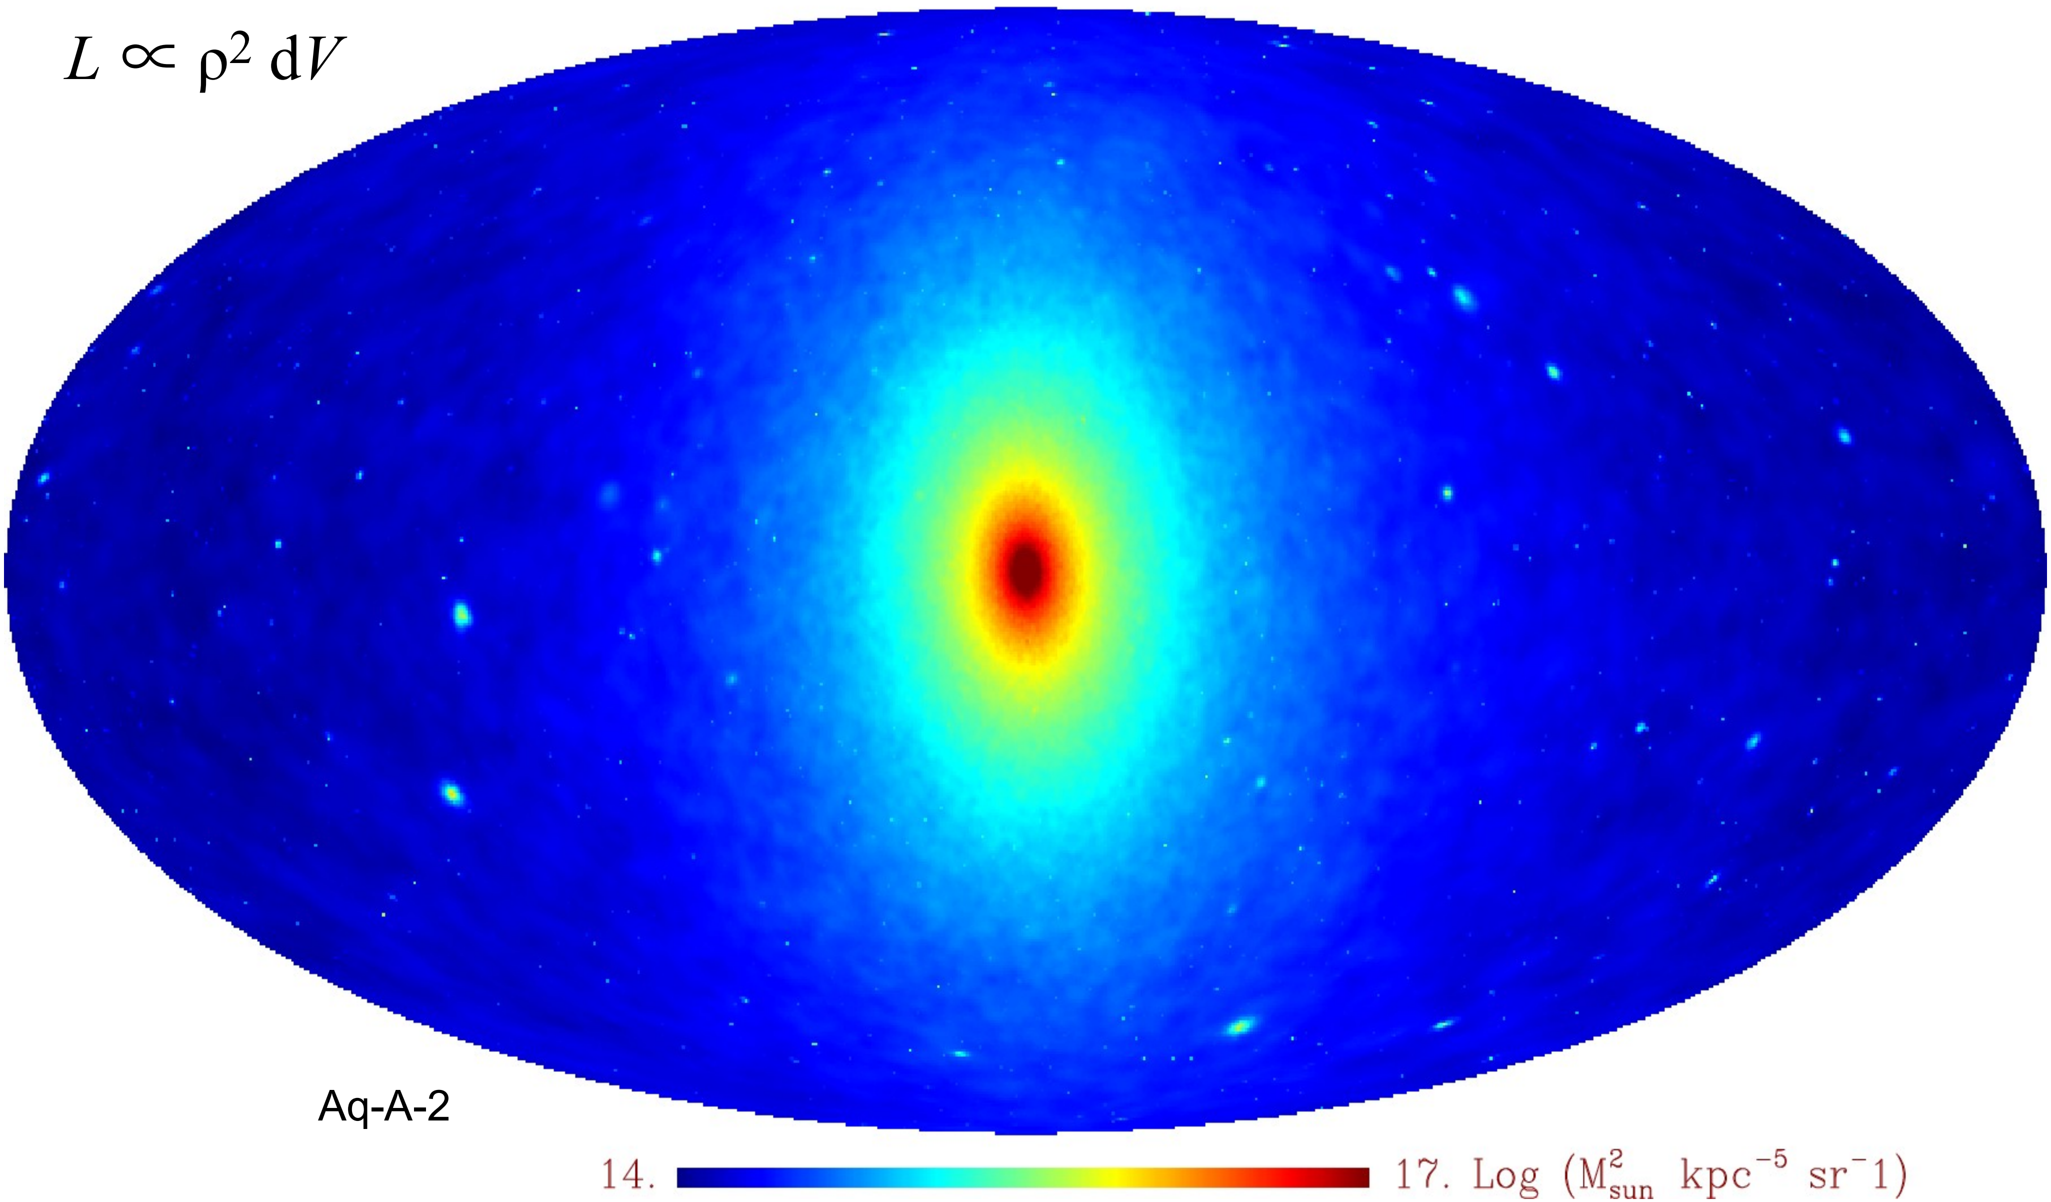

Aq-A-2

The annihilation luminosity from main halo and subhalos has a very different radial distribution

THE RELATIVE DISTRIBUTION OF MASS, MAIN HALO, AND SUBHALO LUMINOSITY

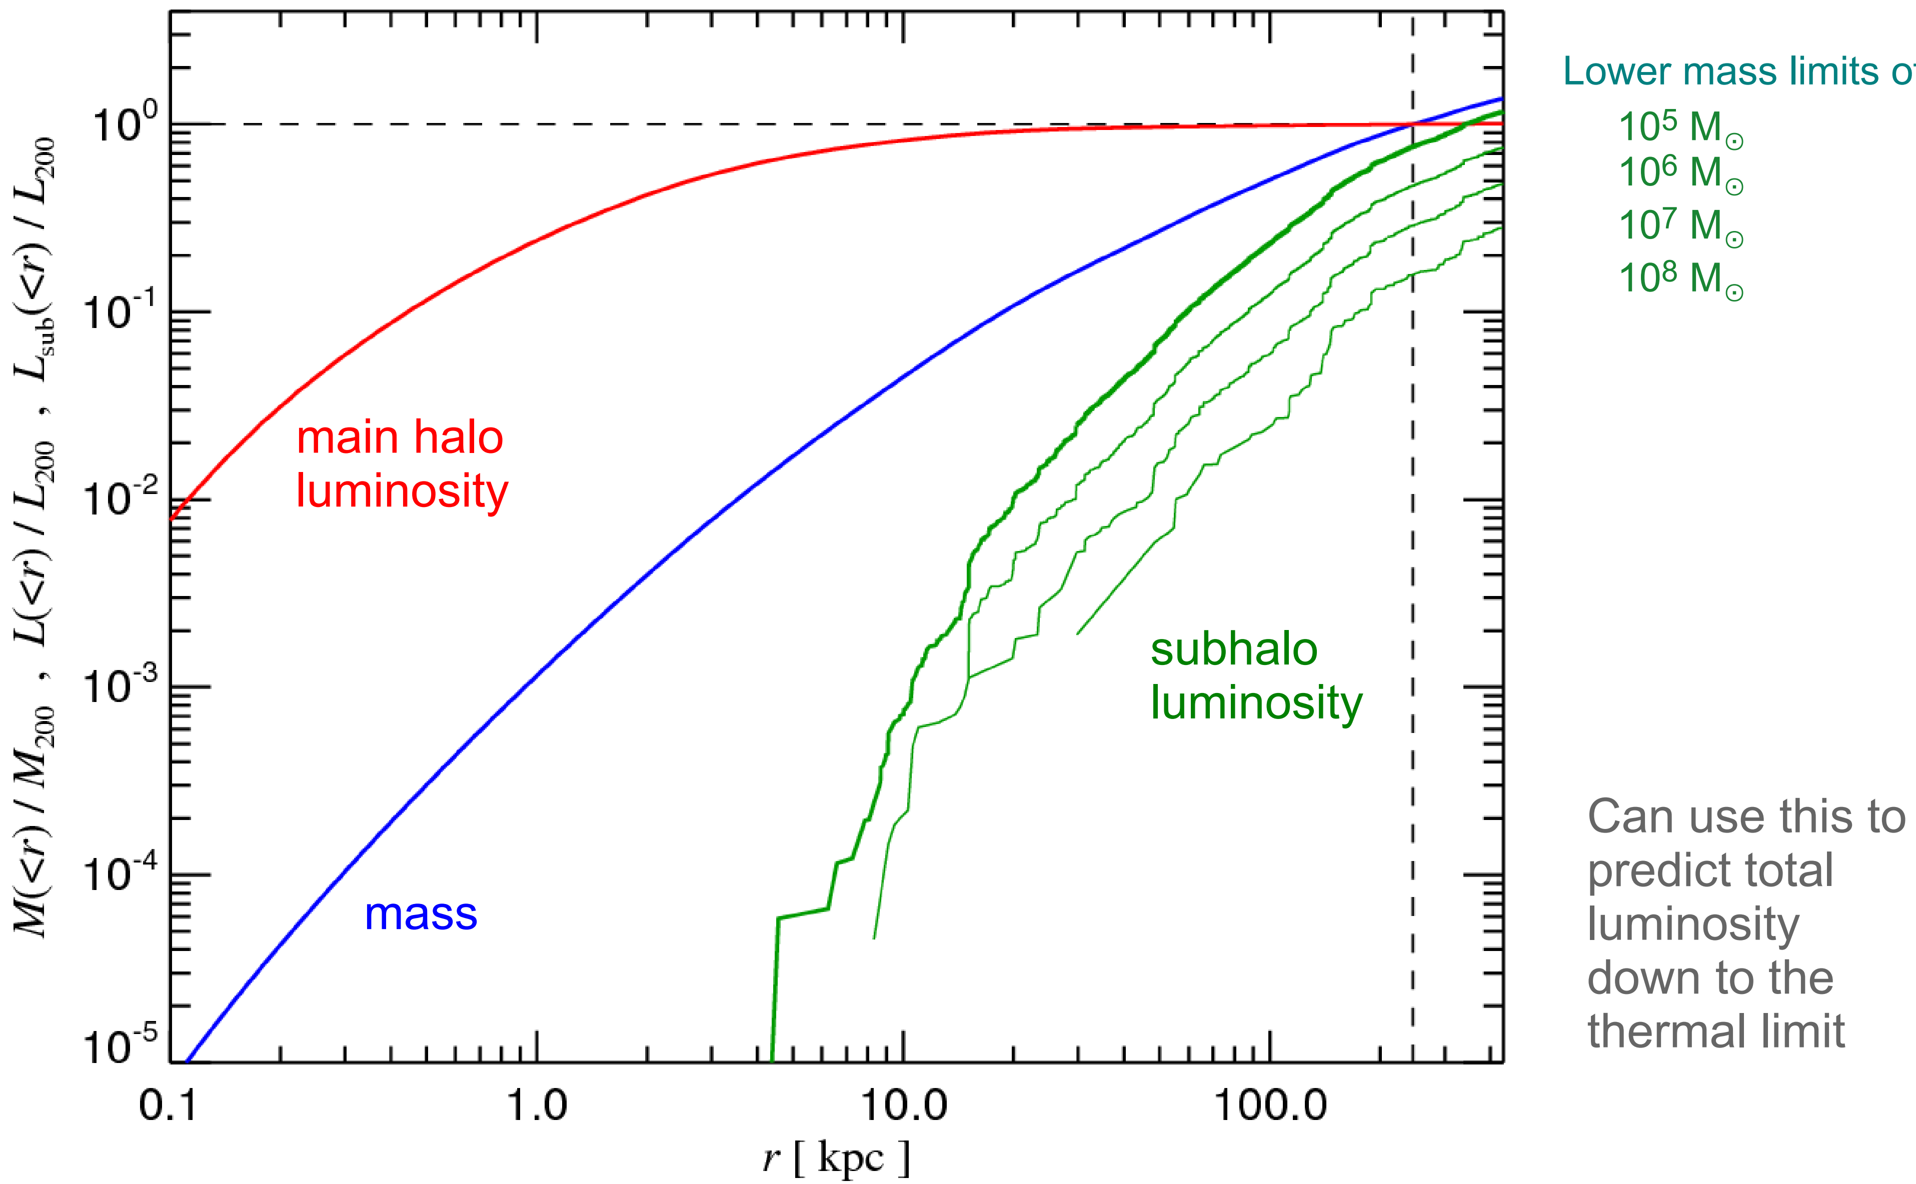

Can use this to predict total luminosity down to the thermal limit

Surface brightness profile of a typical subhalo with V_{max}=10 km/s at different distances from the galactic center

SURFACE BRIGHTNESS PROFILE OF DIFFERENT SUBHALO COMPONENTS

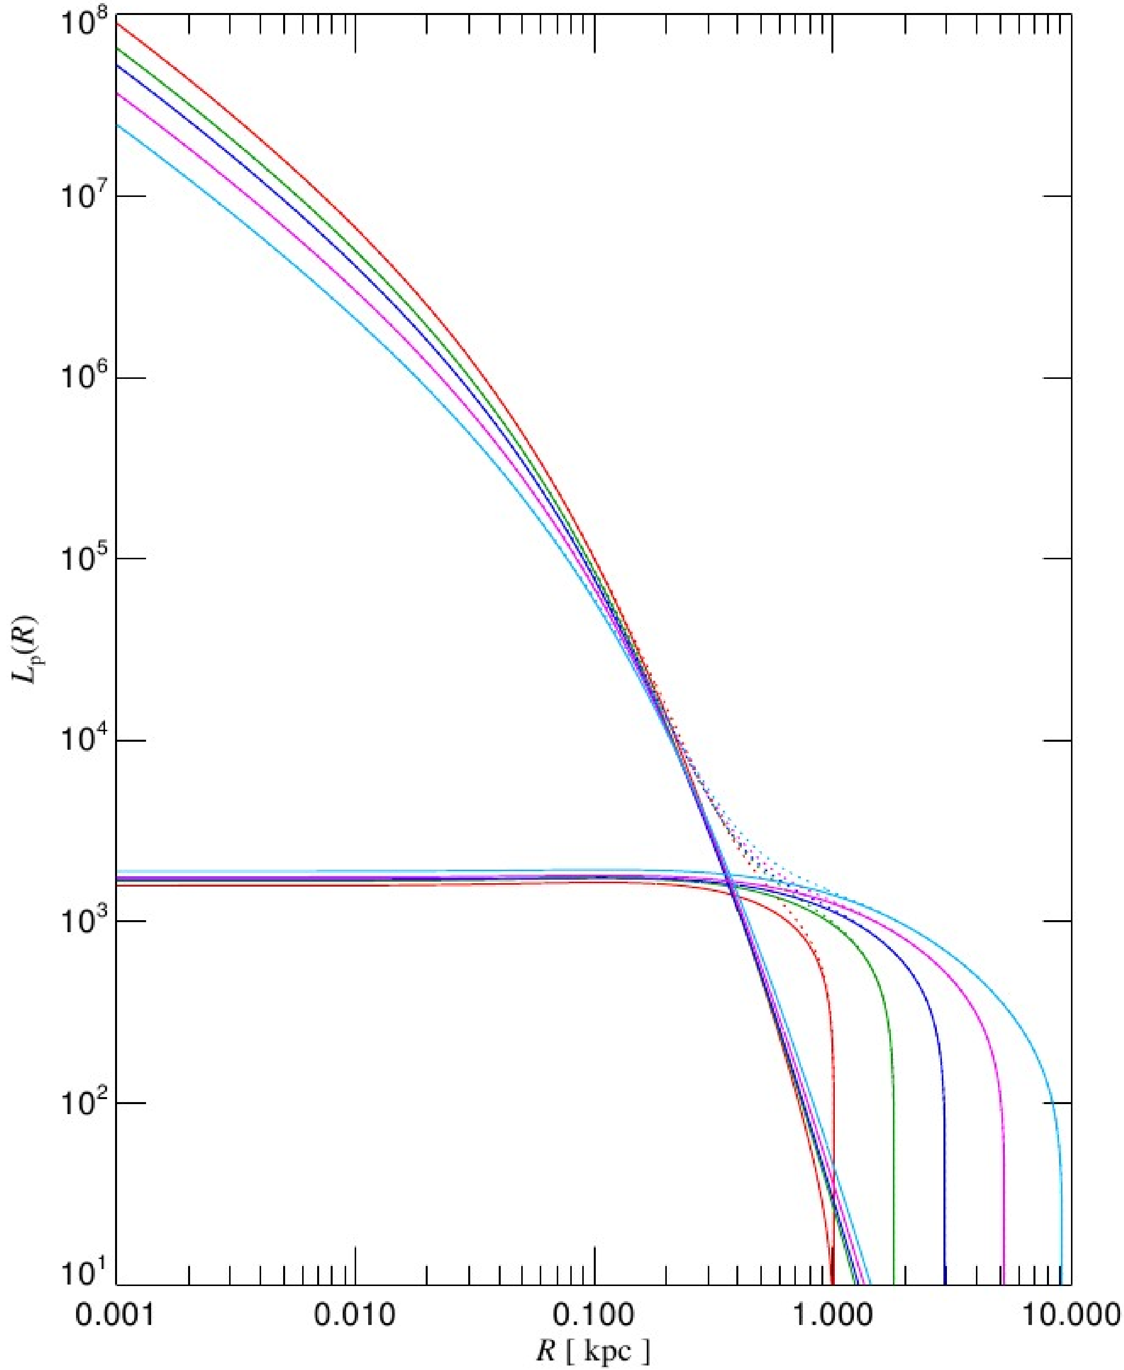

The sub-sub component appears as a (extended) “disk” on the sky

The central surface brightness of the smooth component actually increases with smaller distance (because the concentration increases)

Dark matter annihilation can be best discovered with an optimal filter against a bright background

THE SIGNAL-TO-NOISE FOR DETECTION WITH AN OPTIMAL FILTER

The optimal filter is proportional to the signal

signal

$$S/N = \sqrt{\tau A_{\rm eff}} \left[\int \frac{n_\gamma^2(\theta,\phi)}{n_\gamma(\theta,\phi) + b_\gamma(\theta,\phi)} \mathrm{d}\Omega \right]^{1/2}$$

background noise

The background dominates, then:

Main halo's smooth component:

$$(S/N)_{\rm MainSm} = f_{\rm MainSm} \left[\frac{\tau A_{\rm eff}}{b_\gamma} \right]^{1/2} \frac{F}{\theta_h}$$

Subhalo's smooth component:

$$(S/N)_{\rm SubSm} = f_{\rm SubSm} \left(\theta_h/\theta_{\rm psf} \right) \left[\frac{\tau A_{\rm eff}}{b_\gamma(\vec{\alpha})} \right]^{1/2} \frac{F}{(\theta_h^2 + \theta_{\rm psf}^2)^{1/2}}$$

Sub-substructure of a subhalo:

$$(S/N)_{\rm SubSub} = f_{\rm SubSub} \left(\theta_h/\theta_{\rm psf} \right) \left[\frac{\tau A_{\rm eff}}{b_\gamma(\vec{\alpha})} \right]^{1/2} \frac{F}{(\theta_h^2 + \theta_{\rm psf}^2)^{1/2}}$$

$$\boxed{\begin{array}{l} S/N \sim F/\theta \\ S/N \sim L/\, r_h\, d \end{array}}$$

Detectability of different annihilation emission components in the Milky Way

S/N for detecting subhalos in units of that for the main halo

30 highest S/N objects, assuming the use of optimal filters

$$S/N \propto C V_{\max}^4 / (r_{\rm half}^2 d)$$

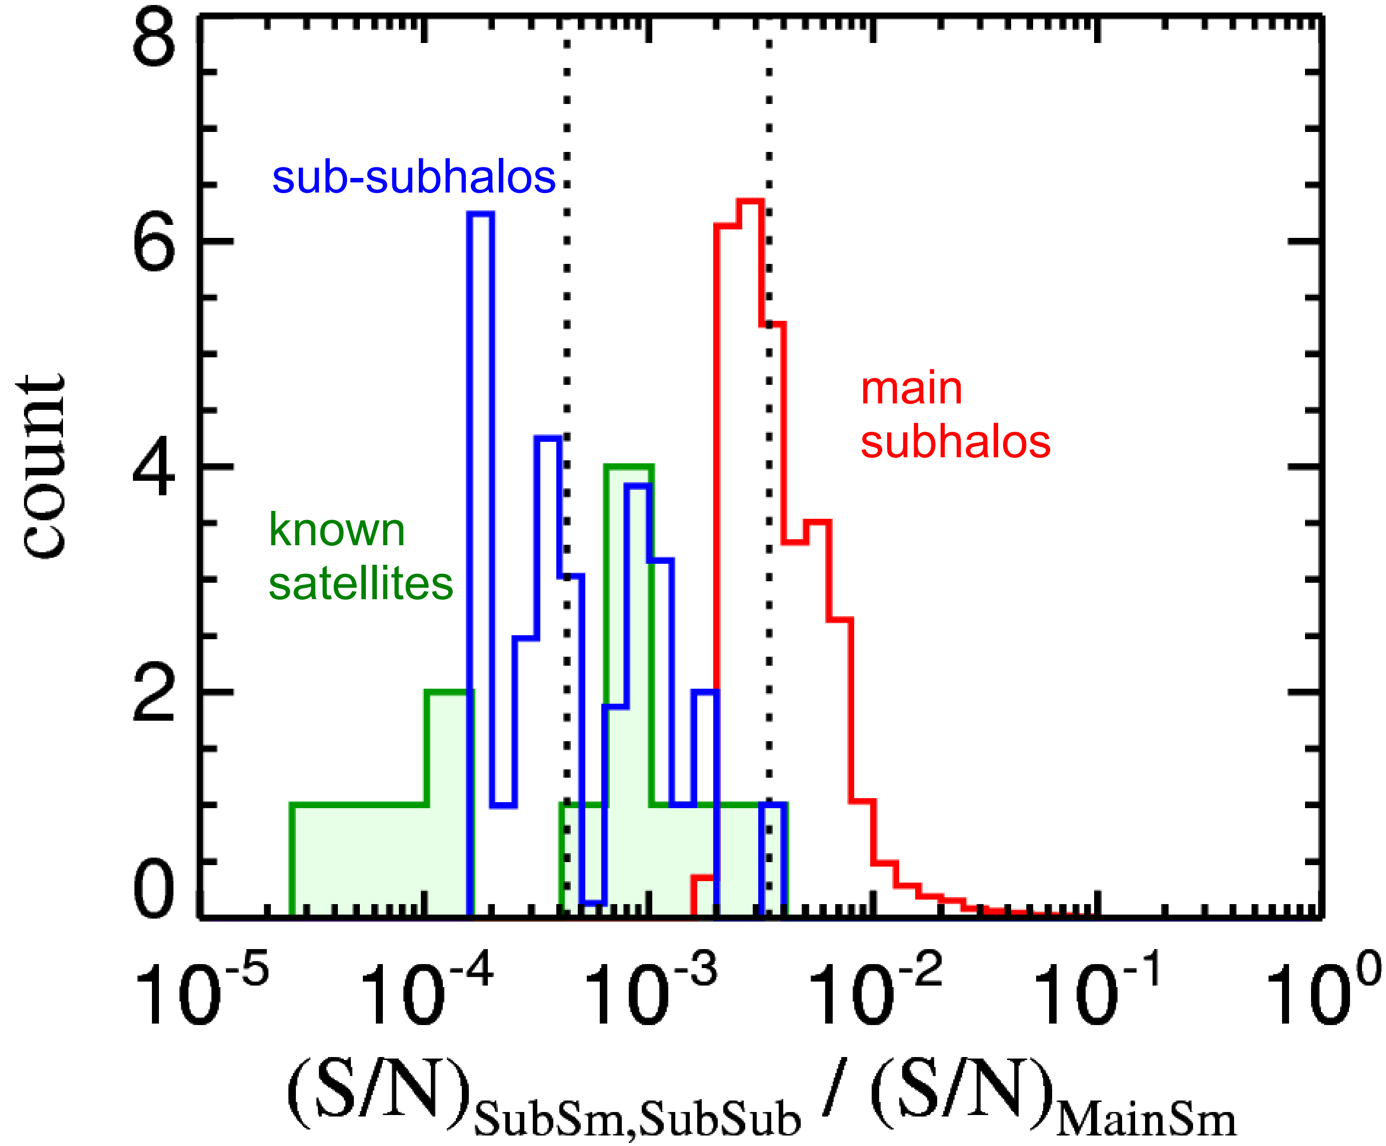

Highest S/N subhalos have 1% of S/N of main halo
Highest S/N subhalos have 10 times S/N of known satellites
Substructure of subhalos has no influence on detectability

The velocity distribution of dark matter at the Sun's position shows residual structure

DISTRIBUTION OF VELOCITY COMPONENTS AND VELOCITY MODULUS

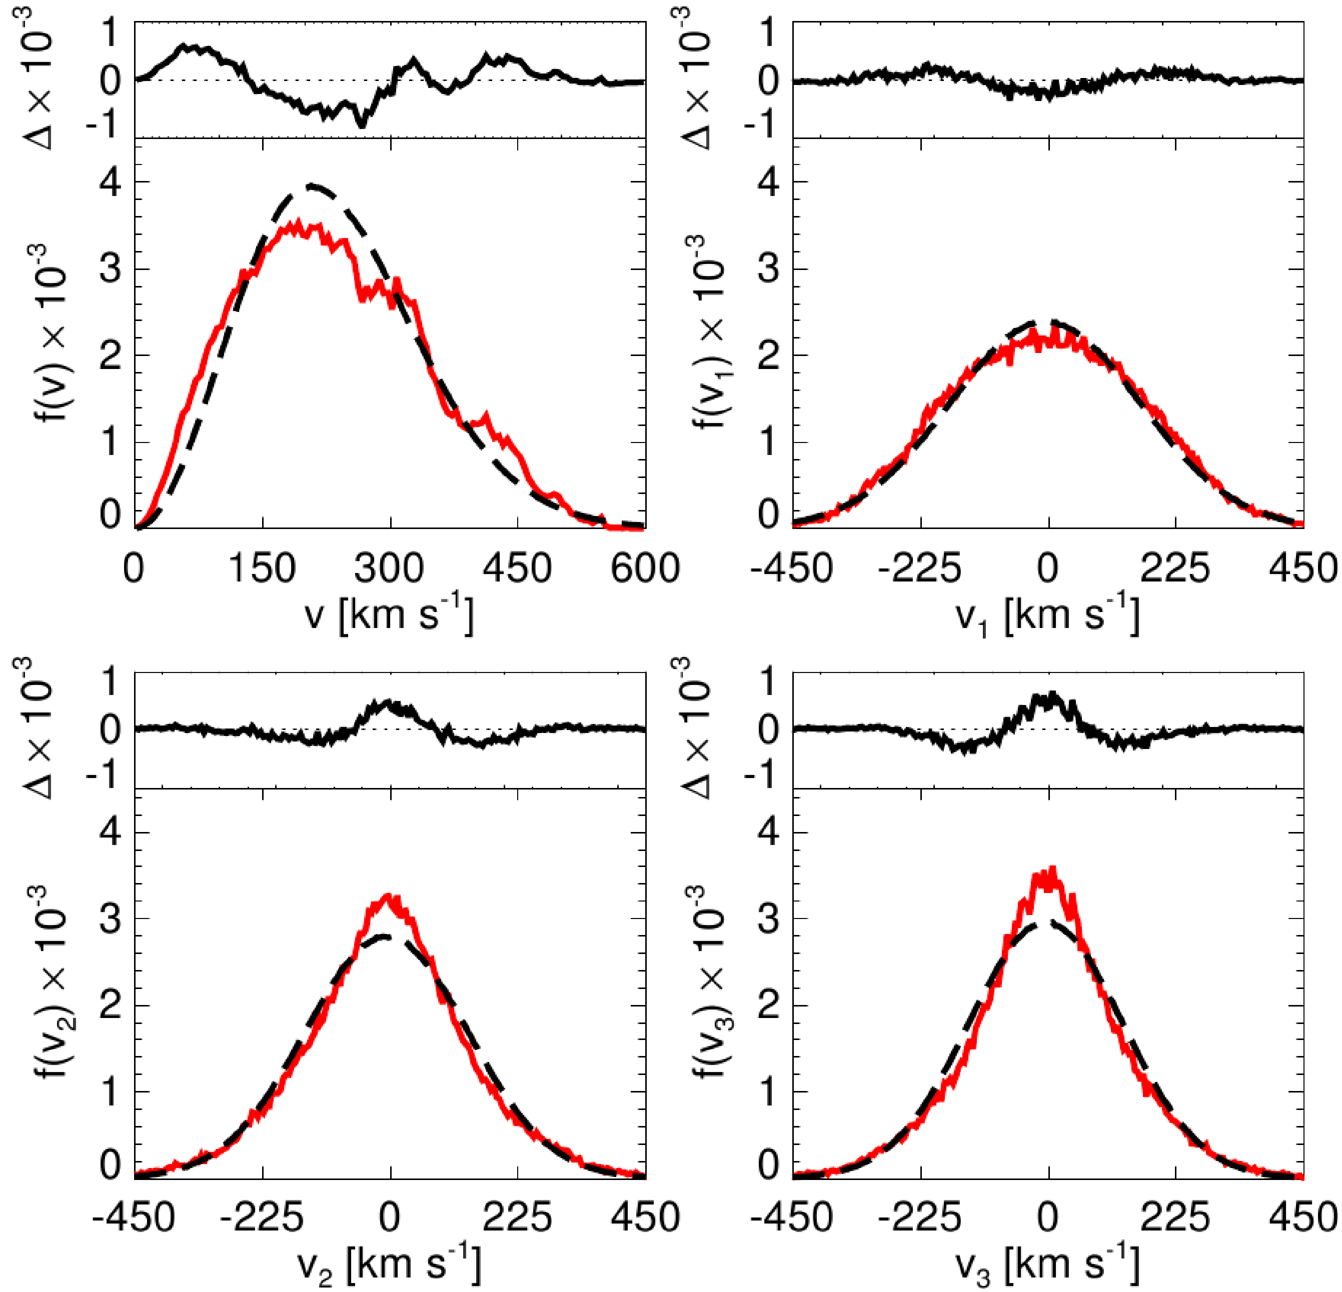

Vogelsberger et al. (2009)

Wiggles in the distribution of the modulus of the velocity point to residual structure in energy space

DISTRIBUTION OF THE VELOCITY MODULUS

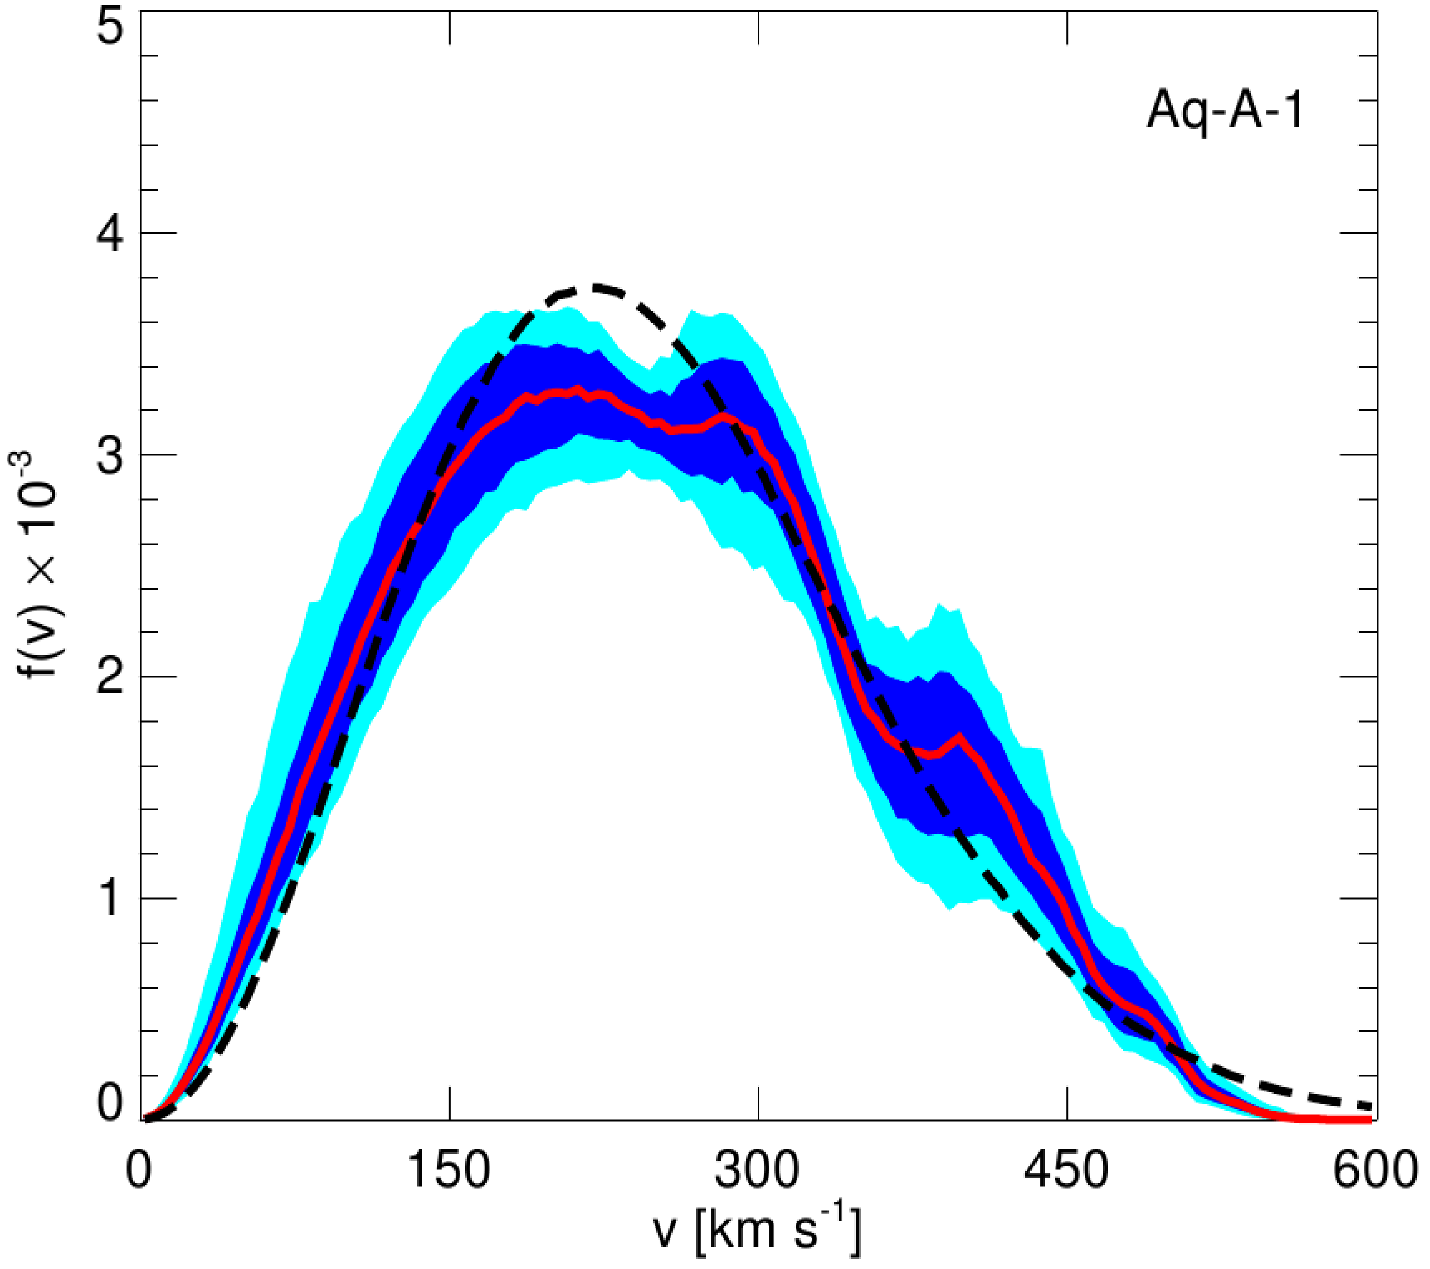

The wiggles are the same in well-separated boxes at the same radial distance, and are reproduced in simulations of different resolution

DISTRIBUTION OF THE VELOCITY MODULUS IN DIFFERENT WELL-SEPARATED BOXES

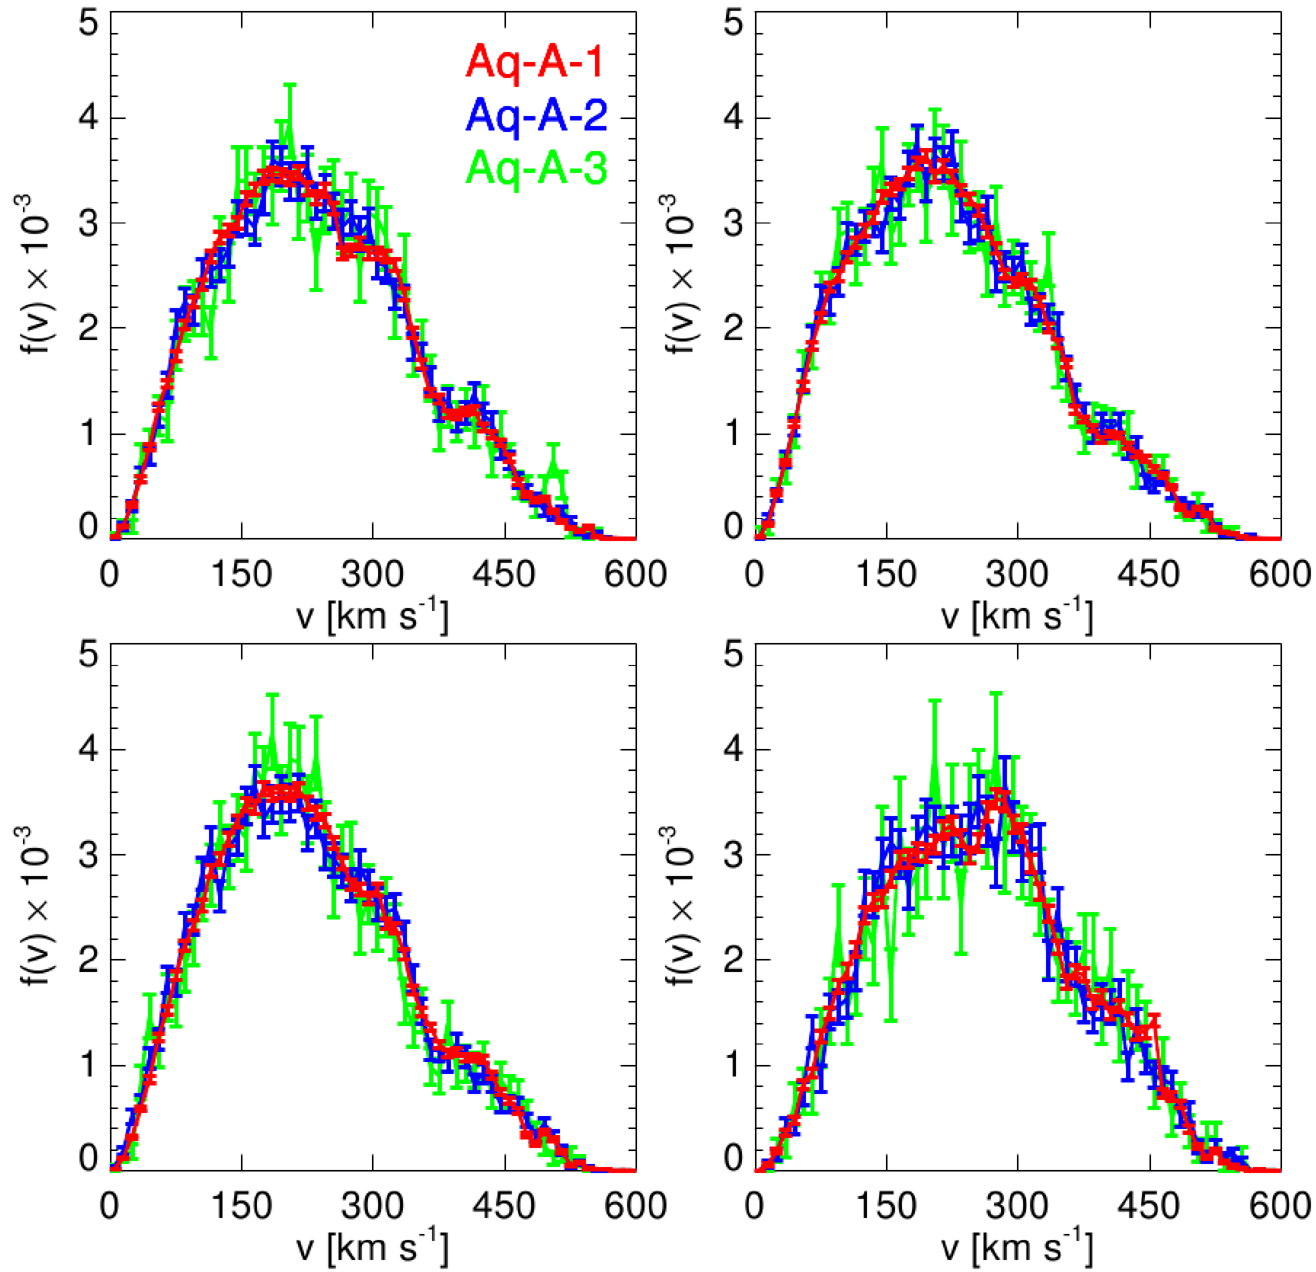

Cosmological N-body simulations have grown rap last four decades

Millennium-XXL Project

"**N**" AS A FUNCTION OF TIME

- Computers double their speed every 18 months (Moore's law)
- N-body simulations have doubled their size every 16-17 months
- Recently, growth has accelerated further.
  The Millennium Run should have become possible in 2010 – we have done it in 2004 !

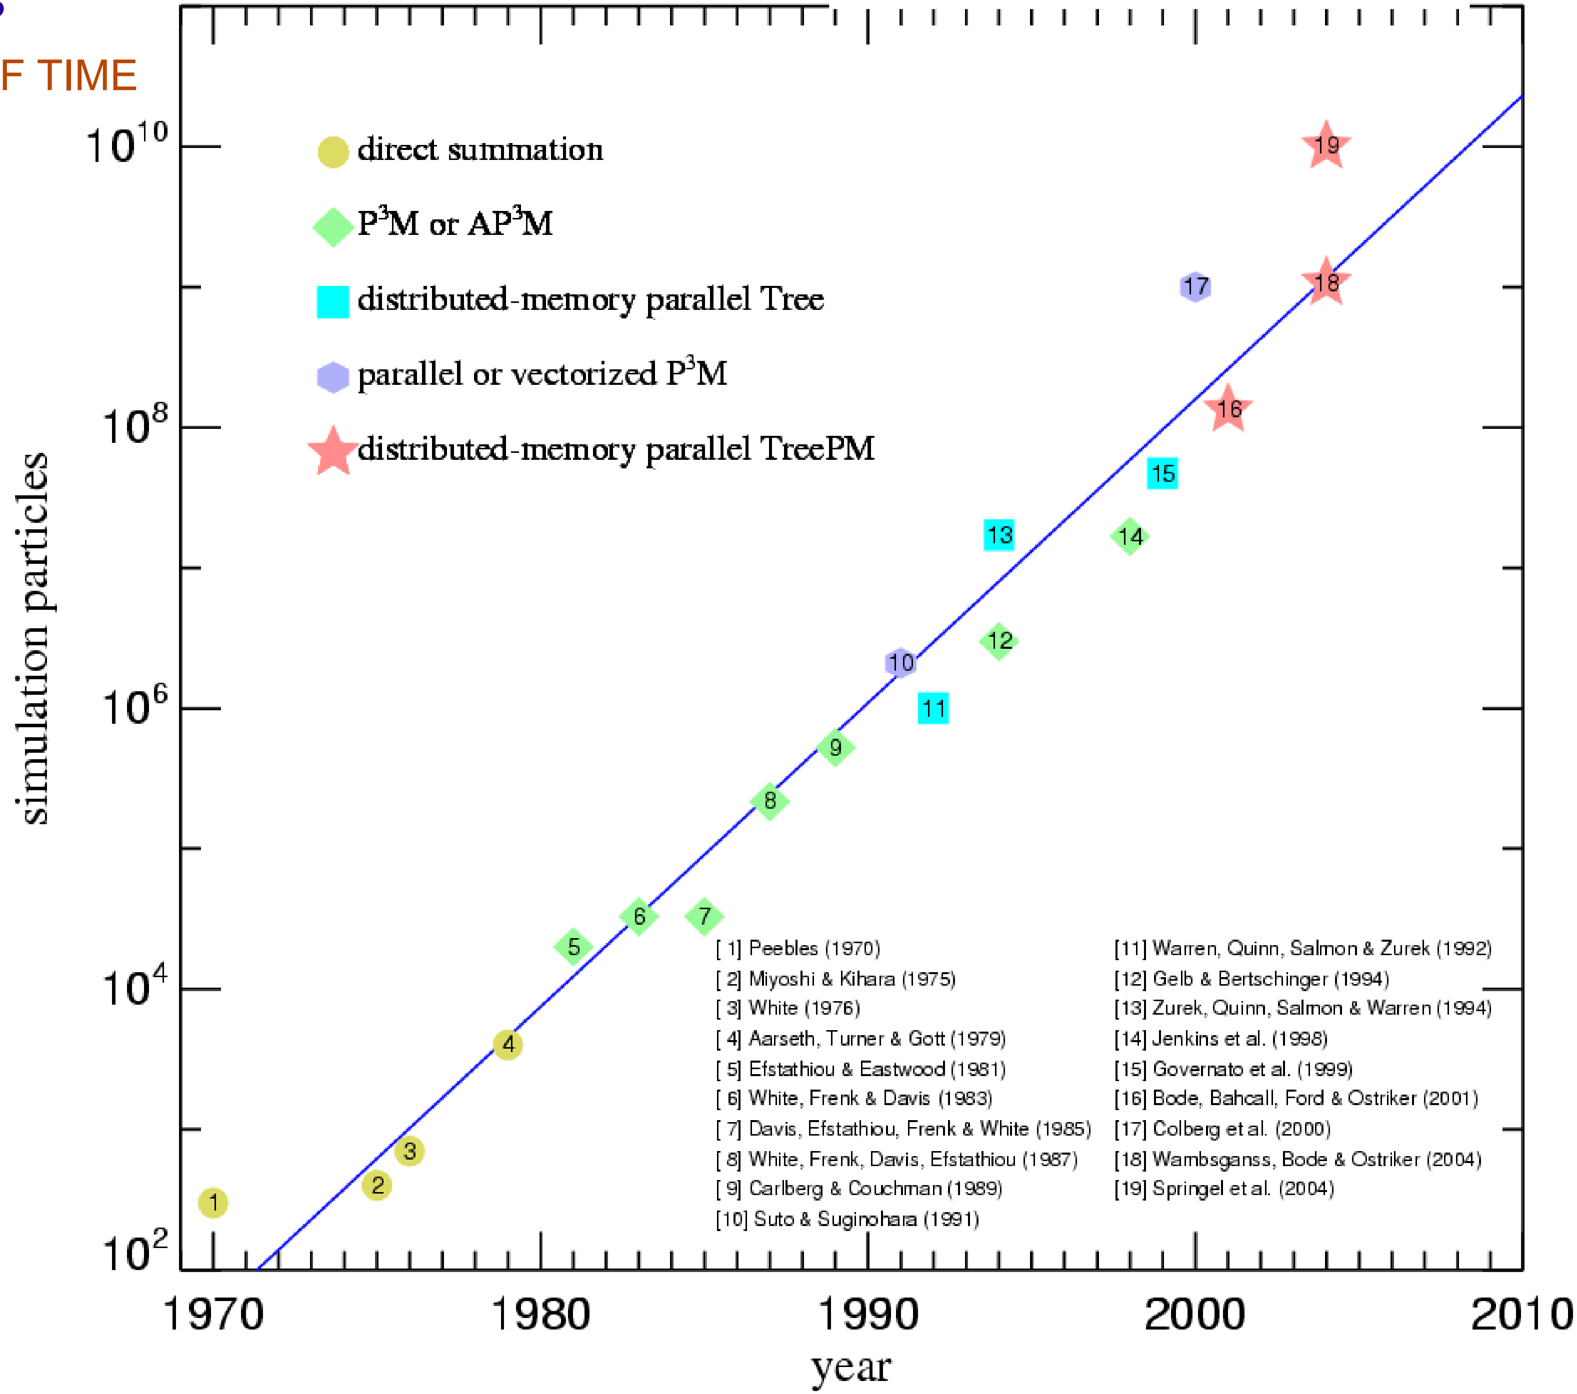

Millennium-XXL

Largest N-body simulation ever

303 billion particles

L = 3 Gpc/h

~700 million halos
at z=0

~25 billion (sub)halos
in mergers trees

$m_p = 6.1 \times 10^9\ M_\odot/h$

12288 cores,
30 TB RAM on
Supercomputer
JuRoPa in Juelich

2.7 million CPU-hours

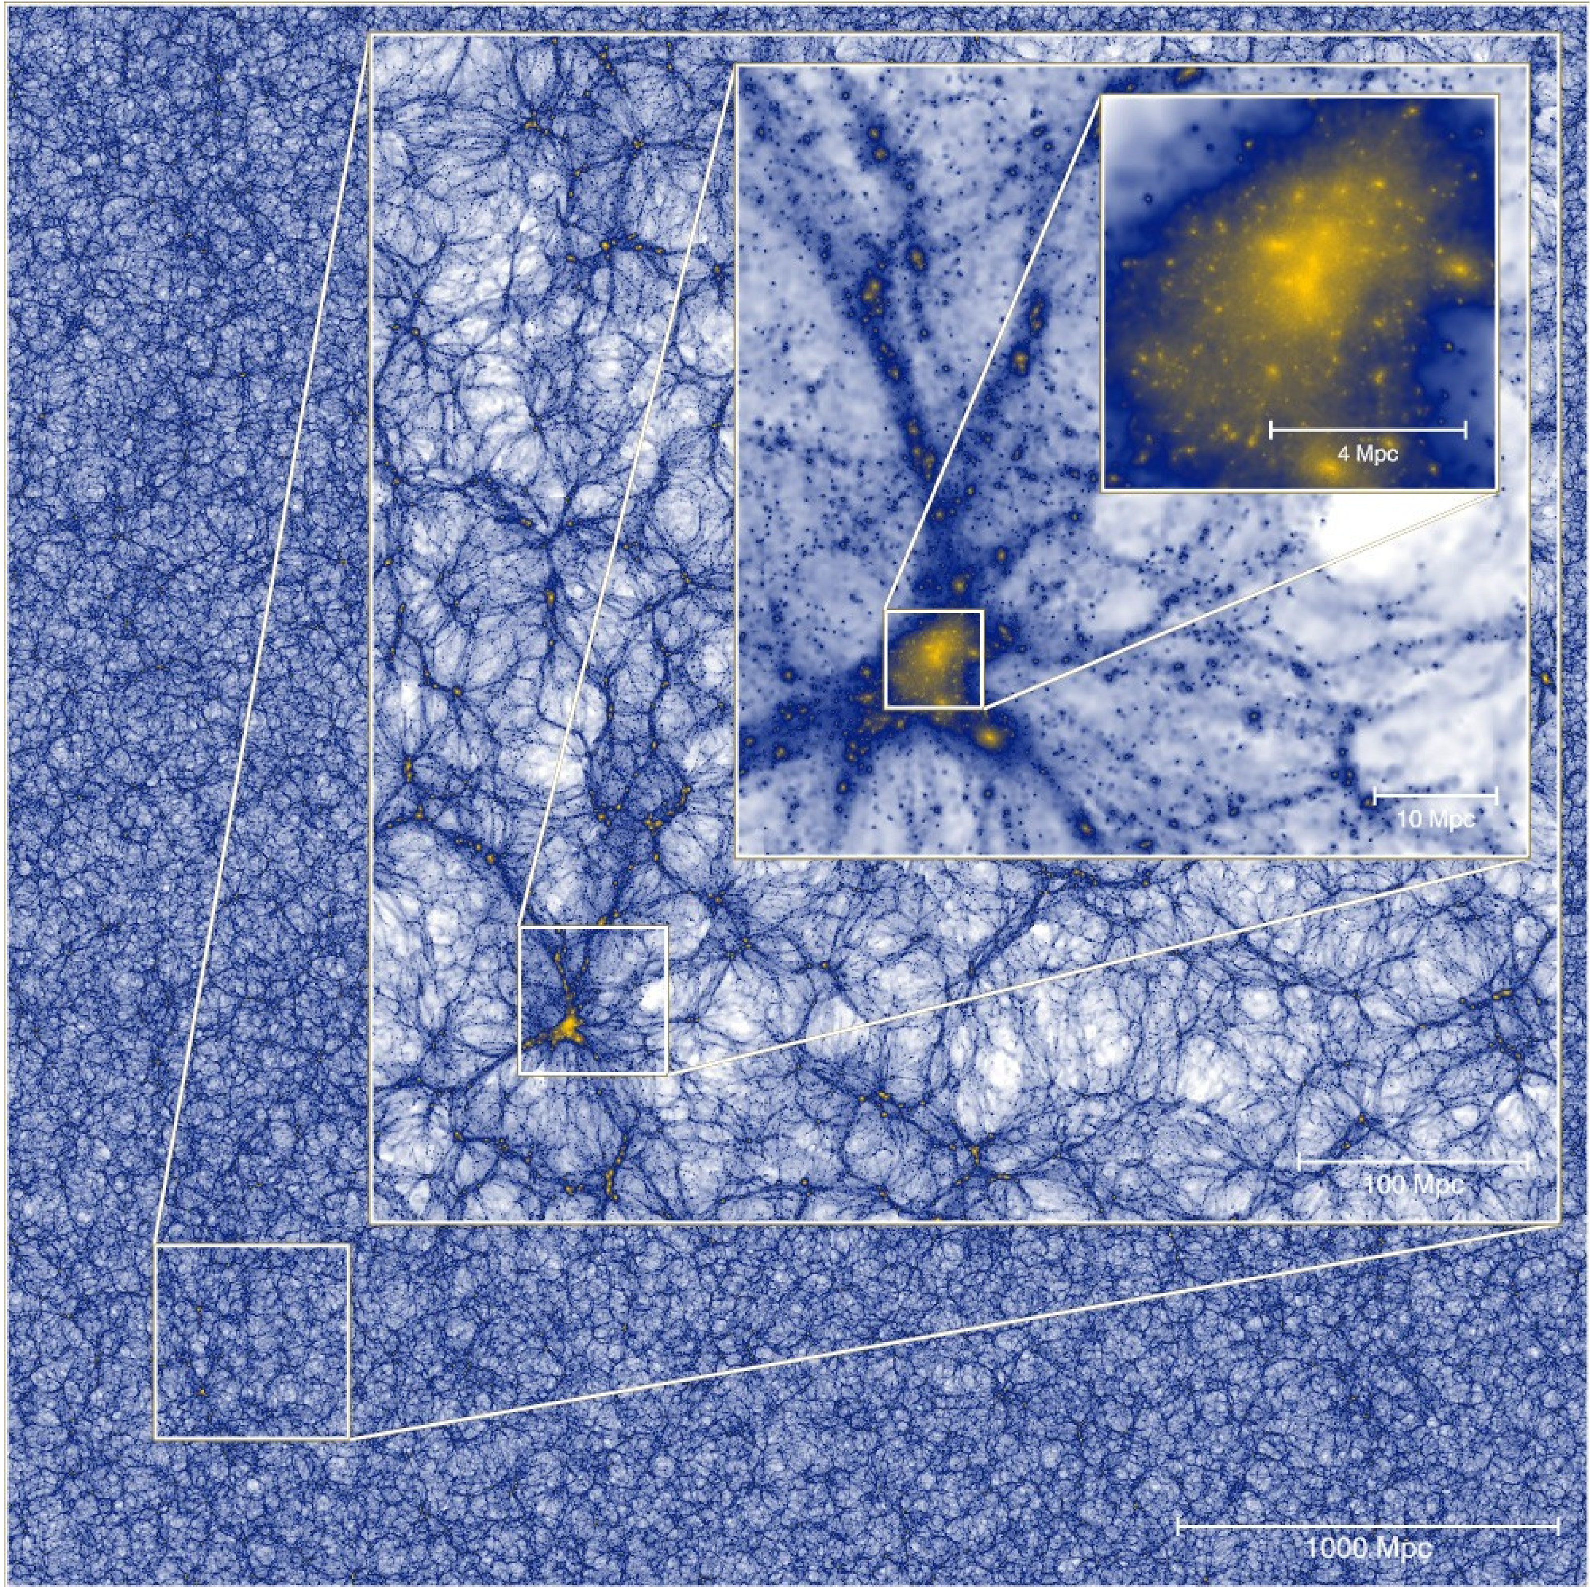

Are the presently known high-mass clusters still consistent with ΛCDM?

CLUSTER MASS FOR GIVEN ABUNDANCE AS A FUNCTION OF TIME

Detection of one violating cluster would invalidate ΛCDM

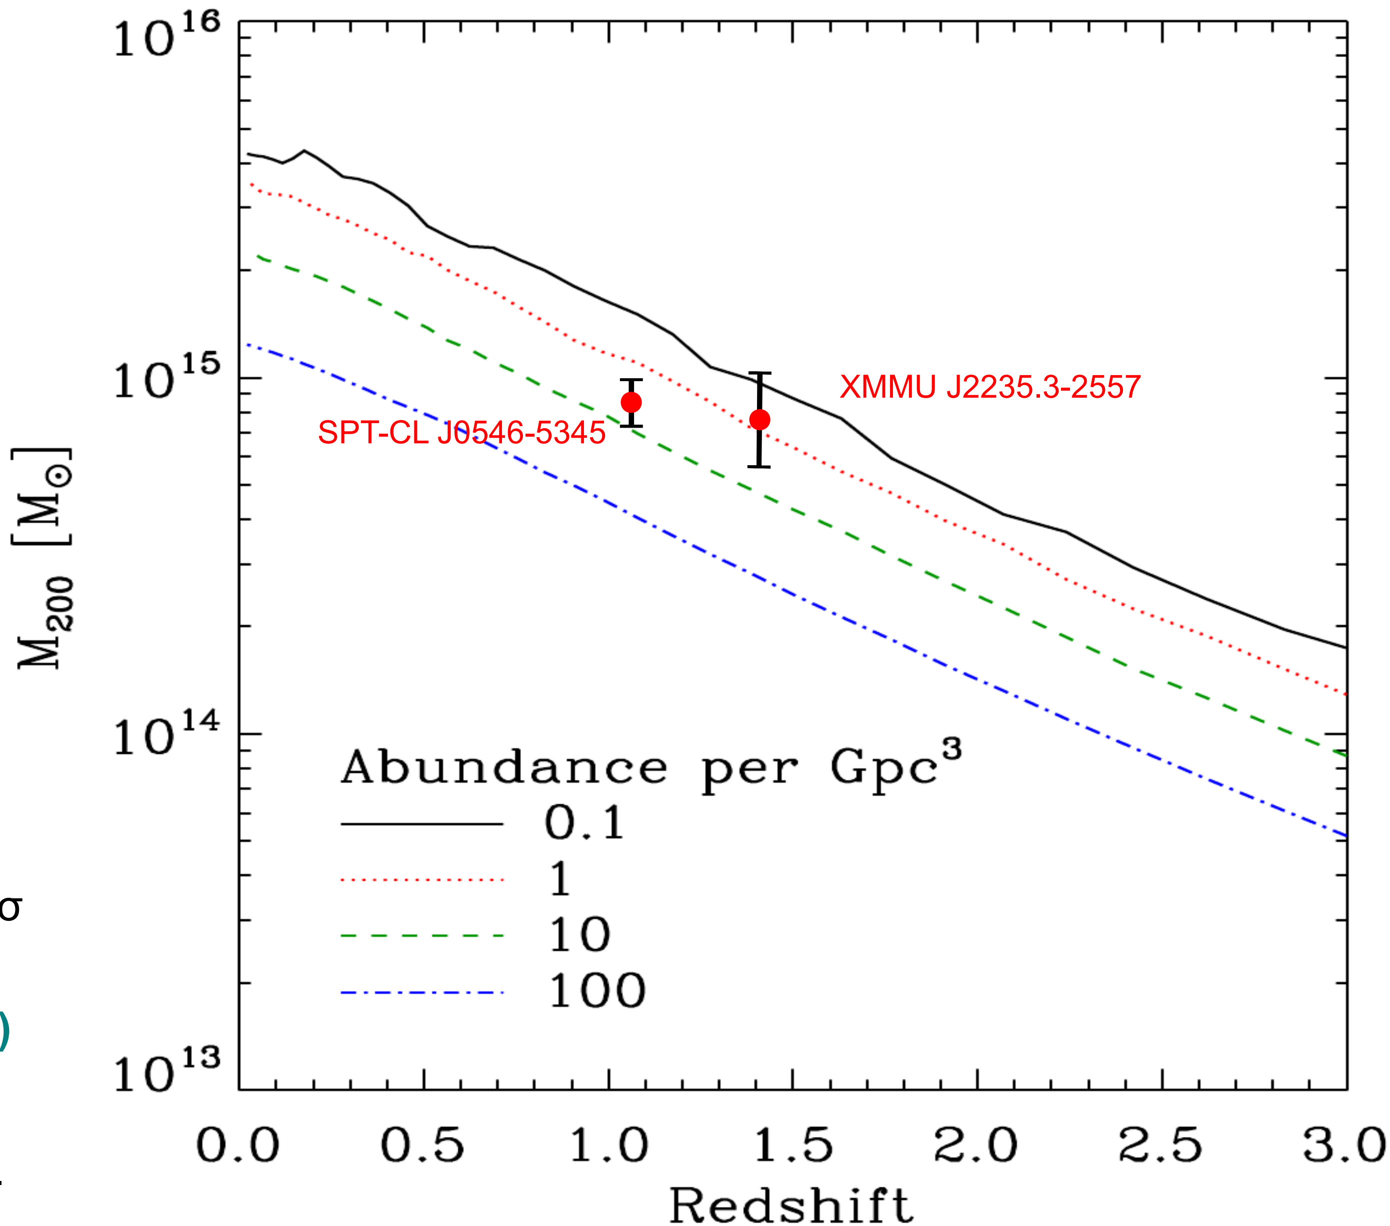

Holz & Perlmutter (2010) argued XMMU J2235.3-2557 to be inconsistent with ΛCDM at 3σ

Boyle, Jiminez & Verde (2010) argue that σ_8 would have to be ~4σ higher to accommodate massive clusters. Suggest non-Gaussian ICs as a solution.

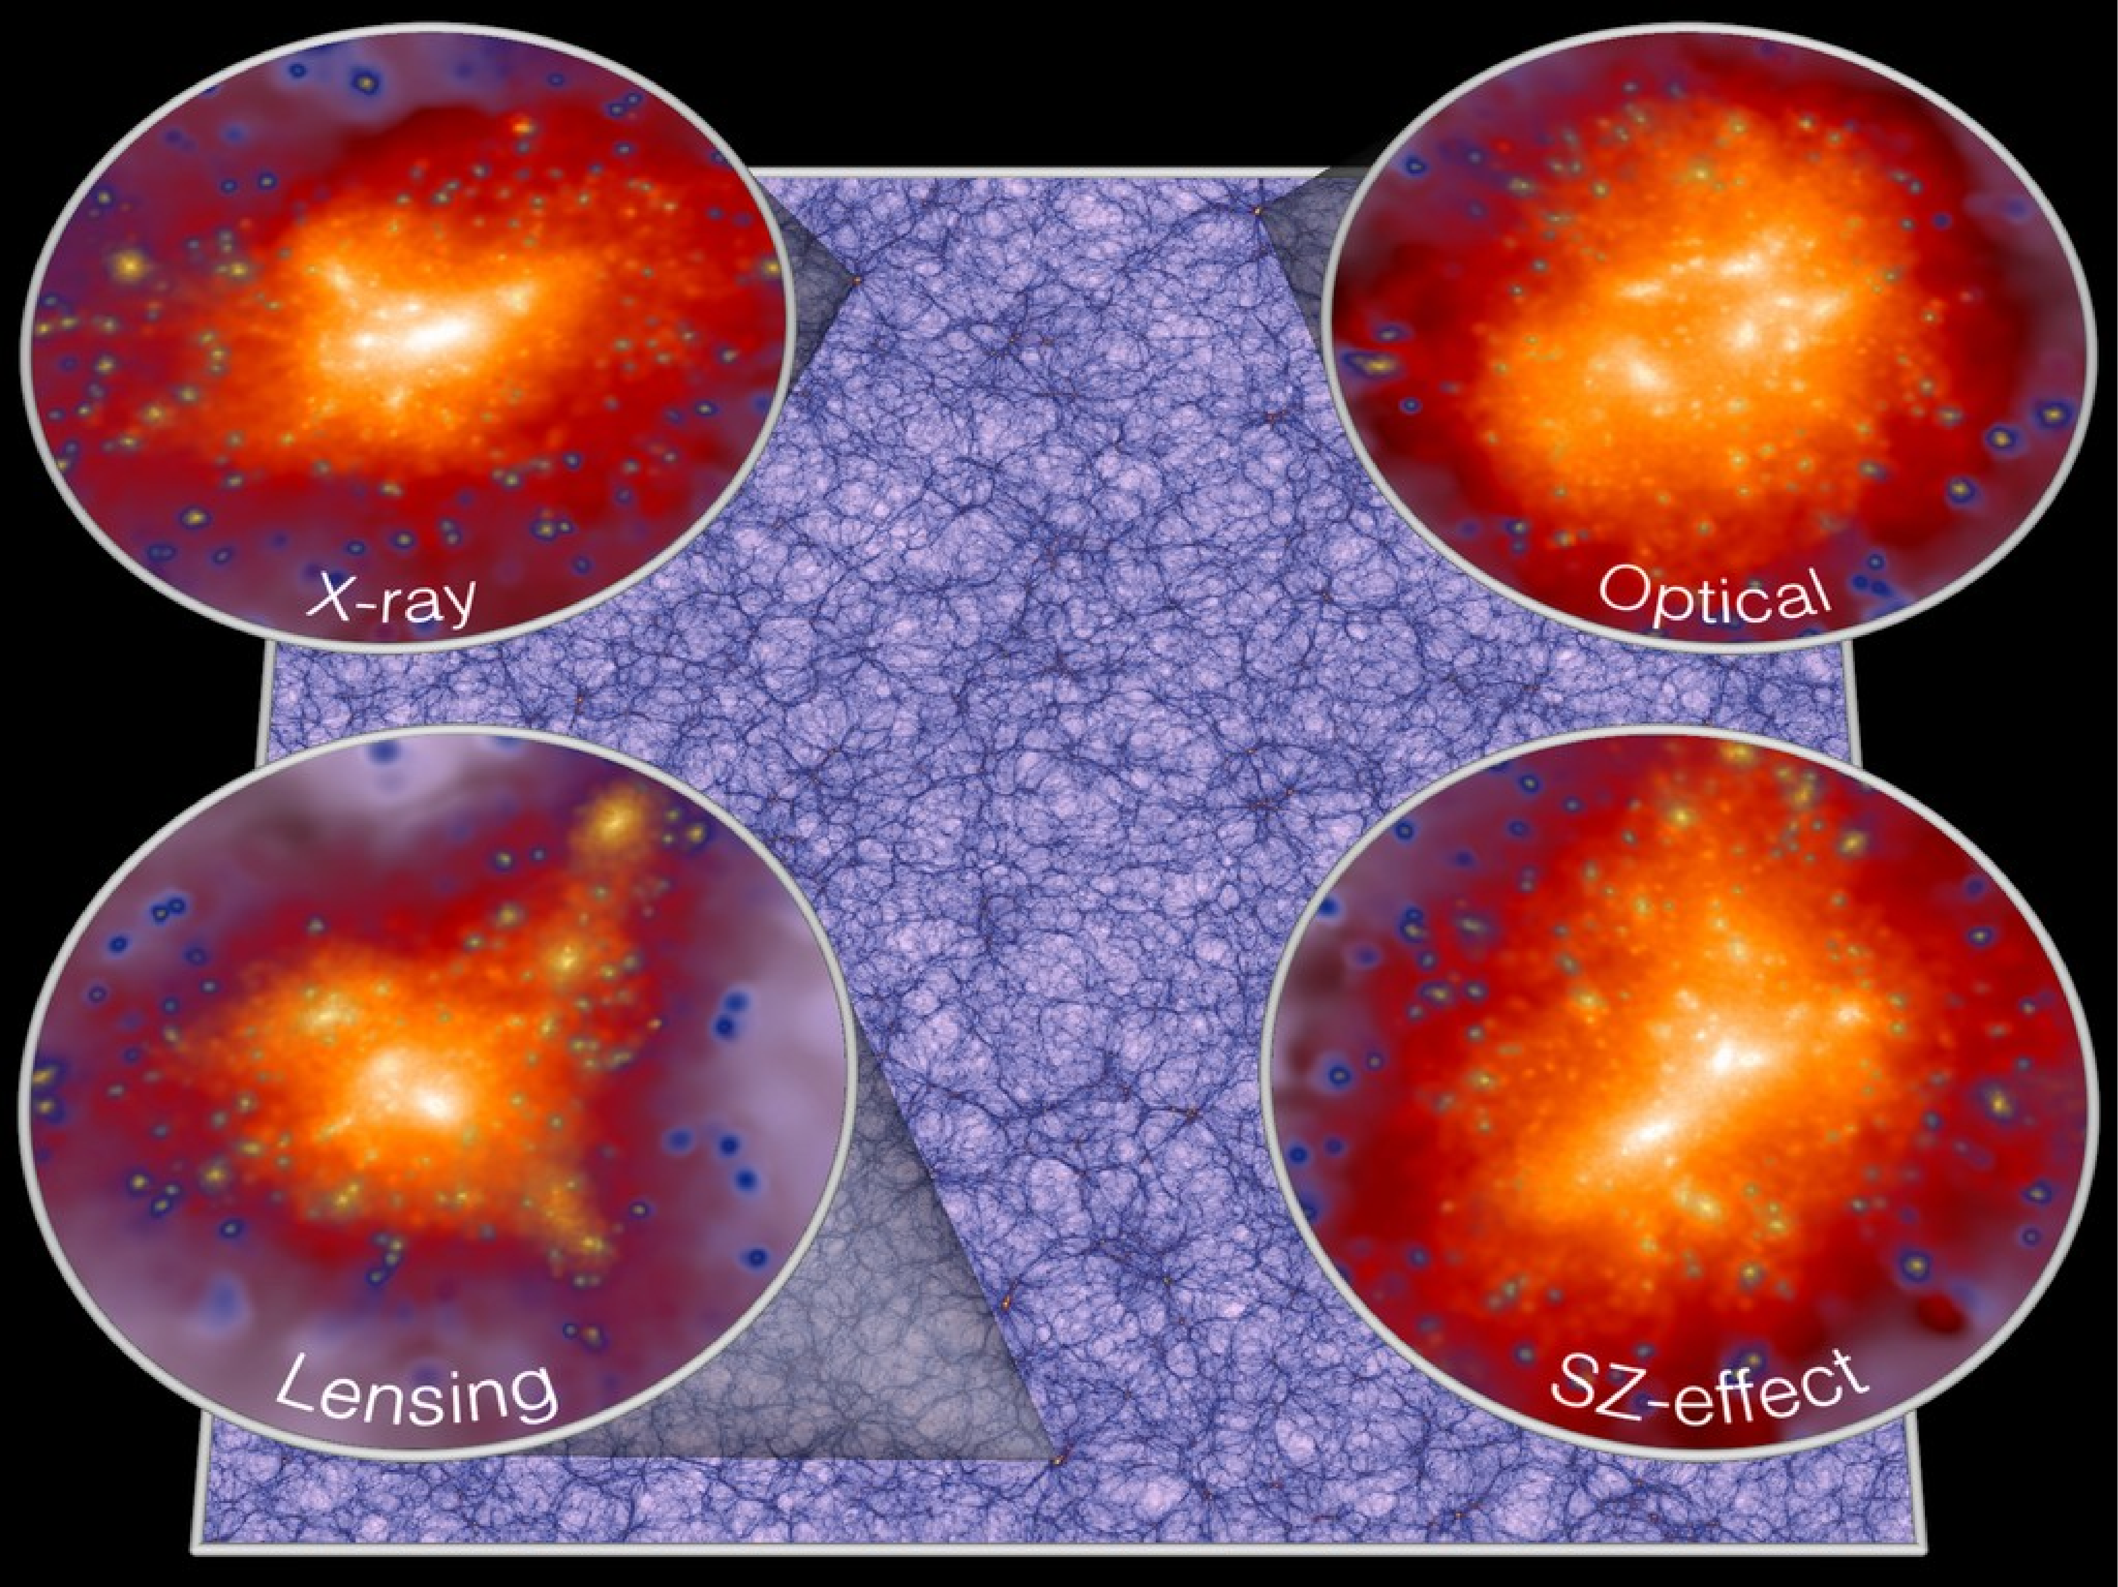
X-ray
Optical
Lensing
SZ-effect

Diversity in the Extreme

THE MOST MASSIVE AND RAREST CLUSTERS FOLLOW THE SCALING RELATIONS EXPECTED FROM MORE ABUNDANT SMALLER SYSTEMS

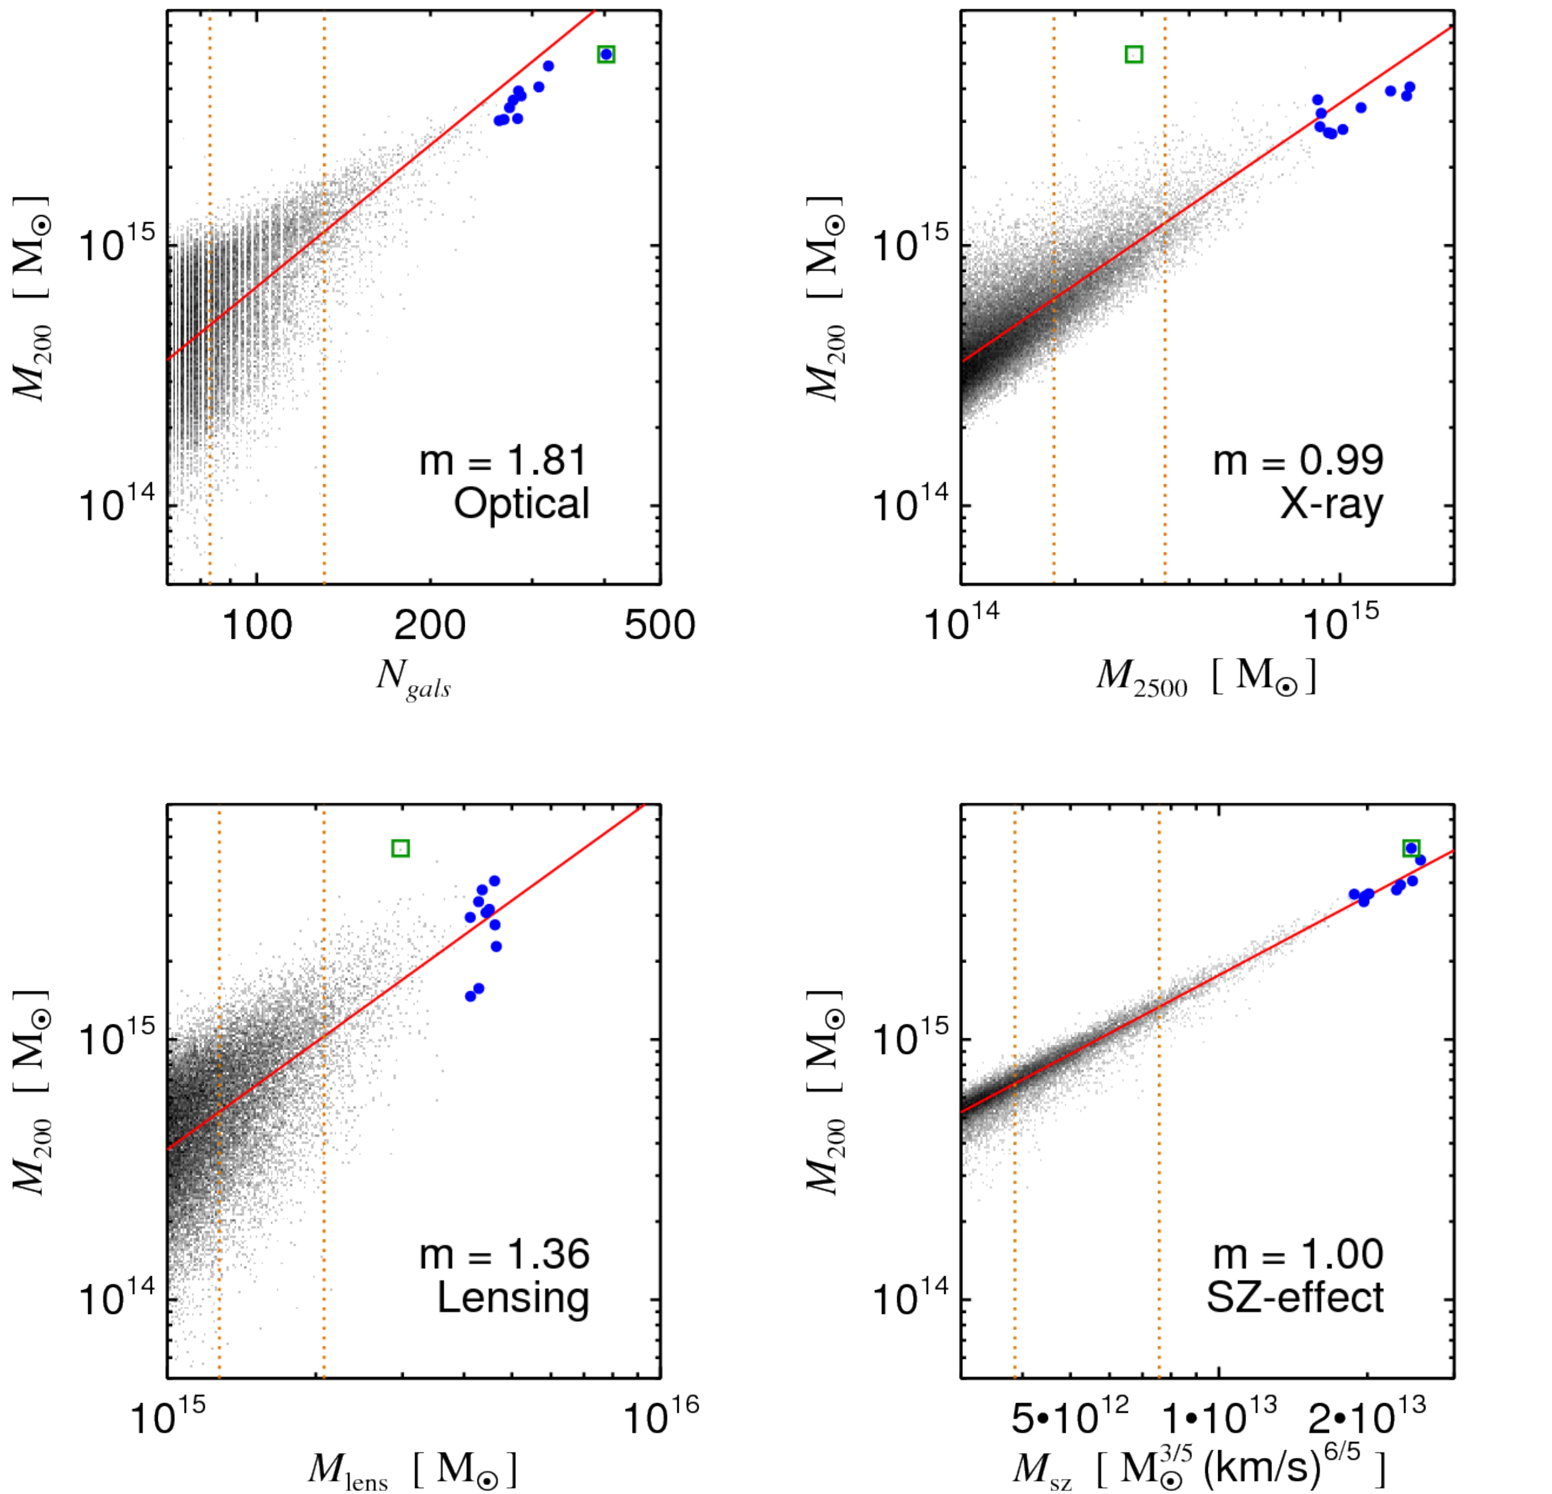

So far reported massive clusters not in conflict with ΛCDM (yet)

The existence of a firm upper limit for the properties of clusters is a strong and falsifiable prediction of ΛCDM

ABUNDANCE OF CLUSTERS USING DIFFERENT OBSERVATIONAL PROBES

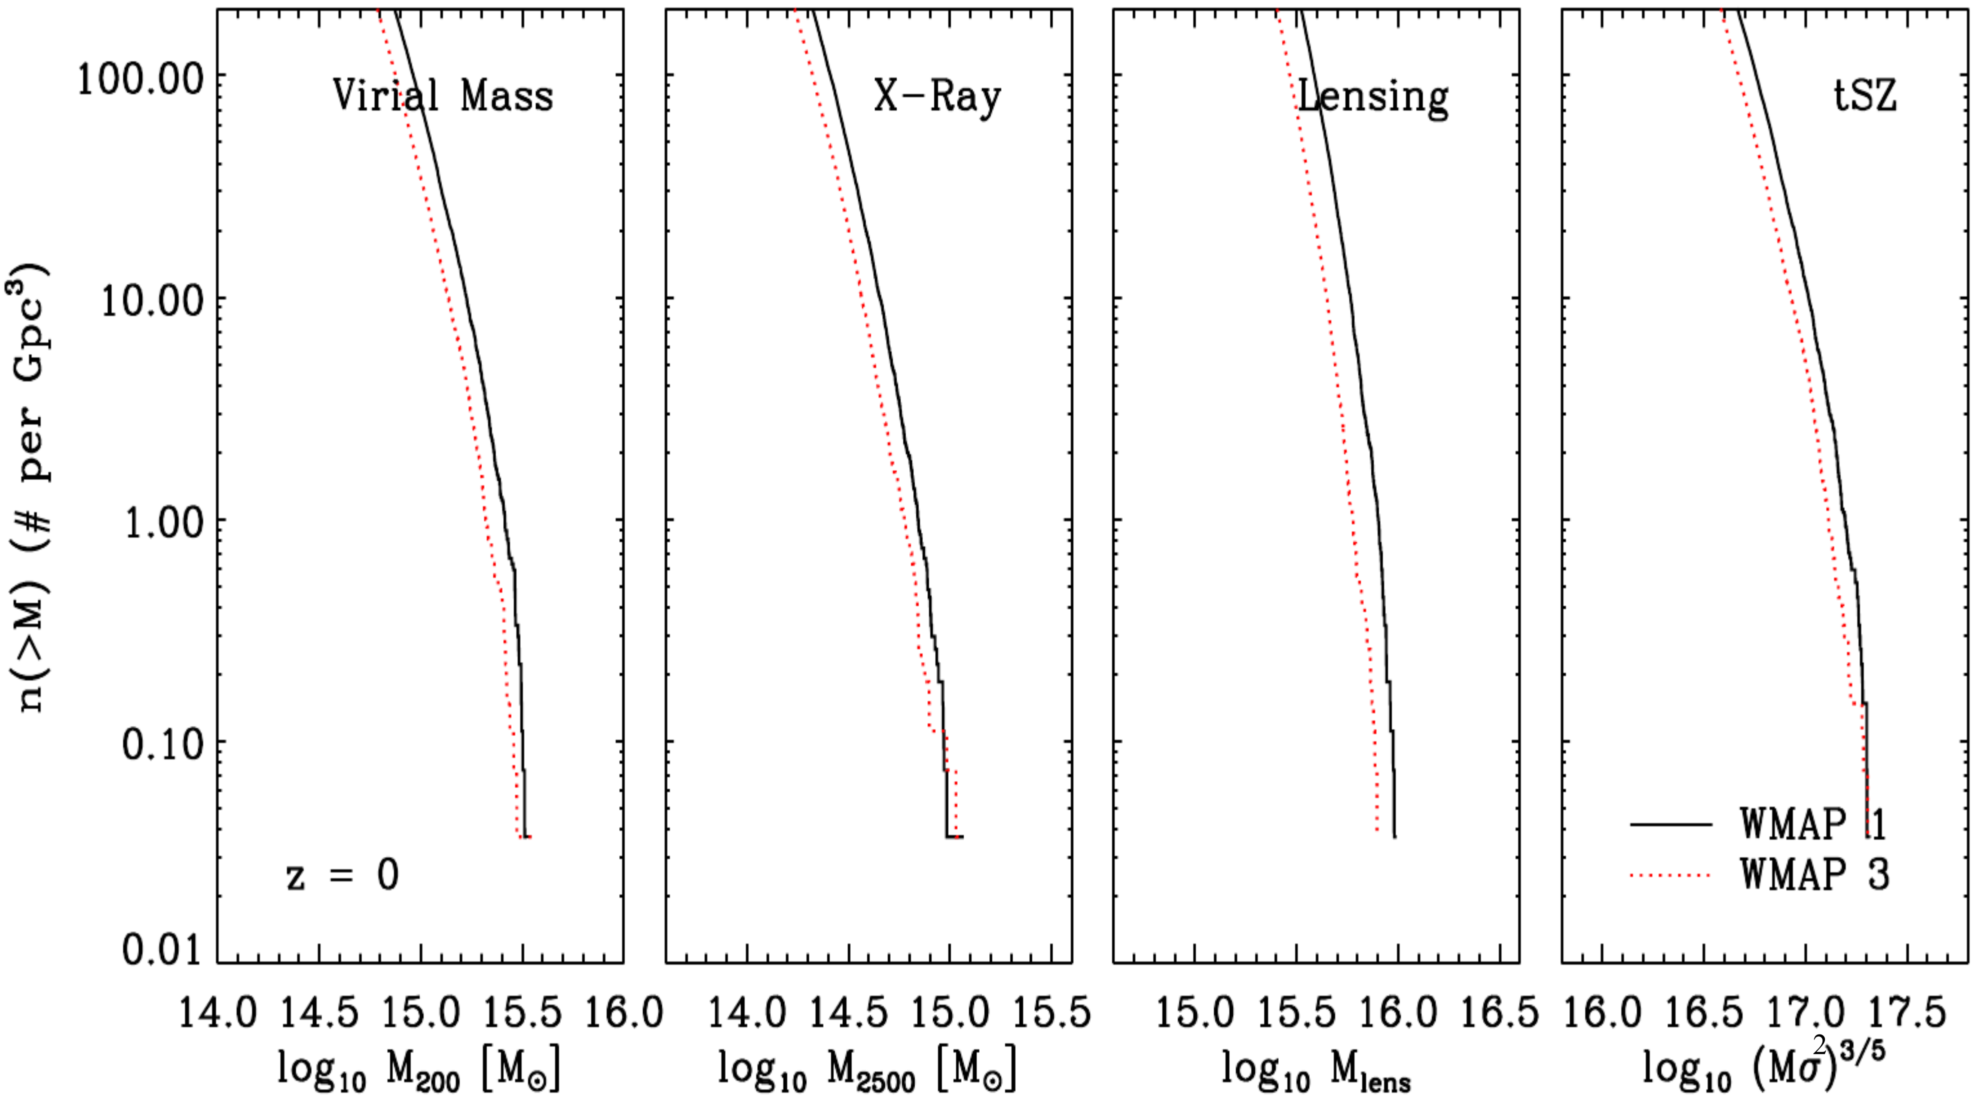

→ no objects above these limits should be detectable by any technique

Trouble ahead in the Exaflop regime

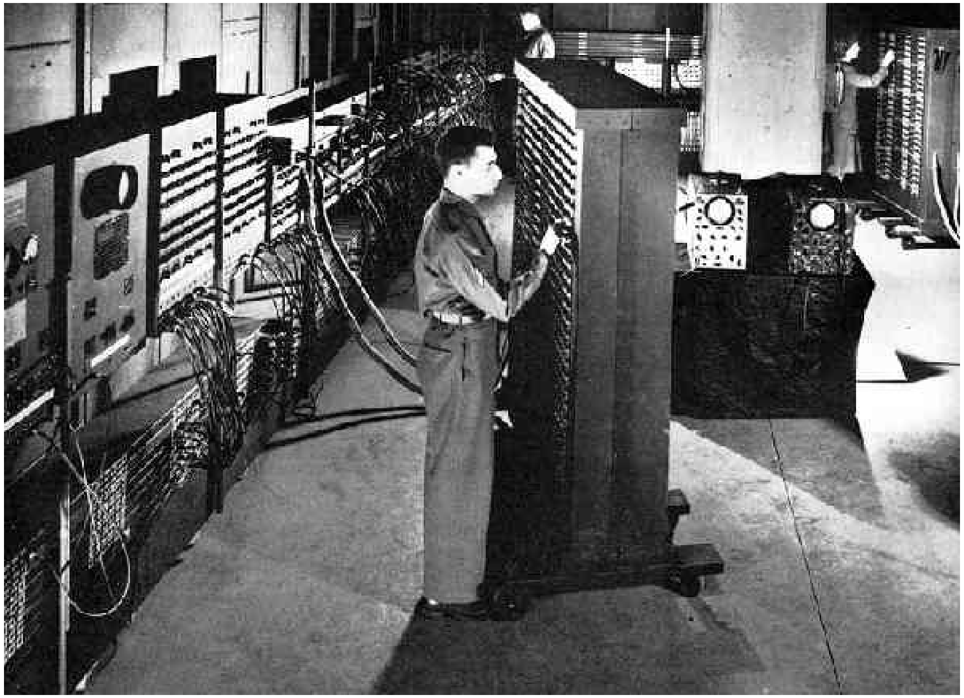

ENIAC, 1946

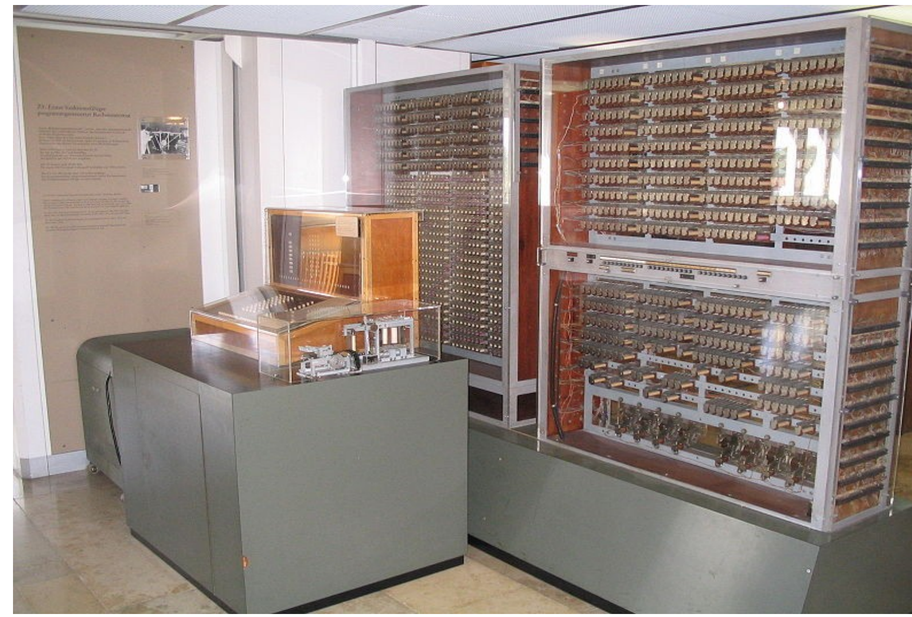

Zuse Z3, 1941

Performance: ~ 1000 Flops

Jaguar: Department of Energy Leadership computer
Designed for science from the ground up

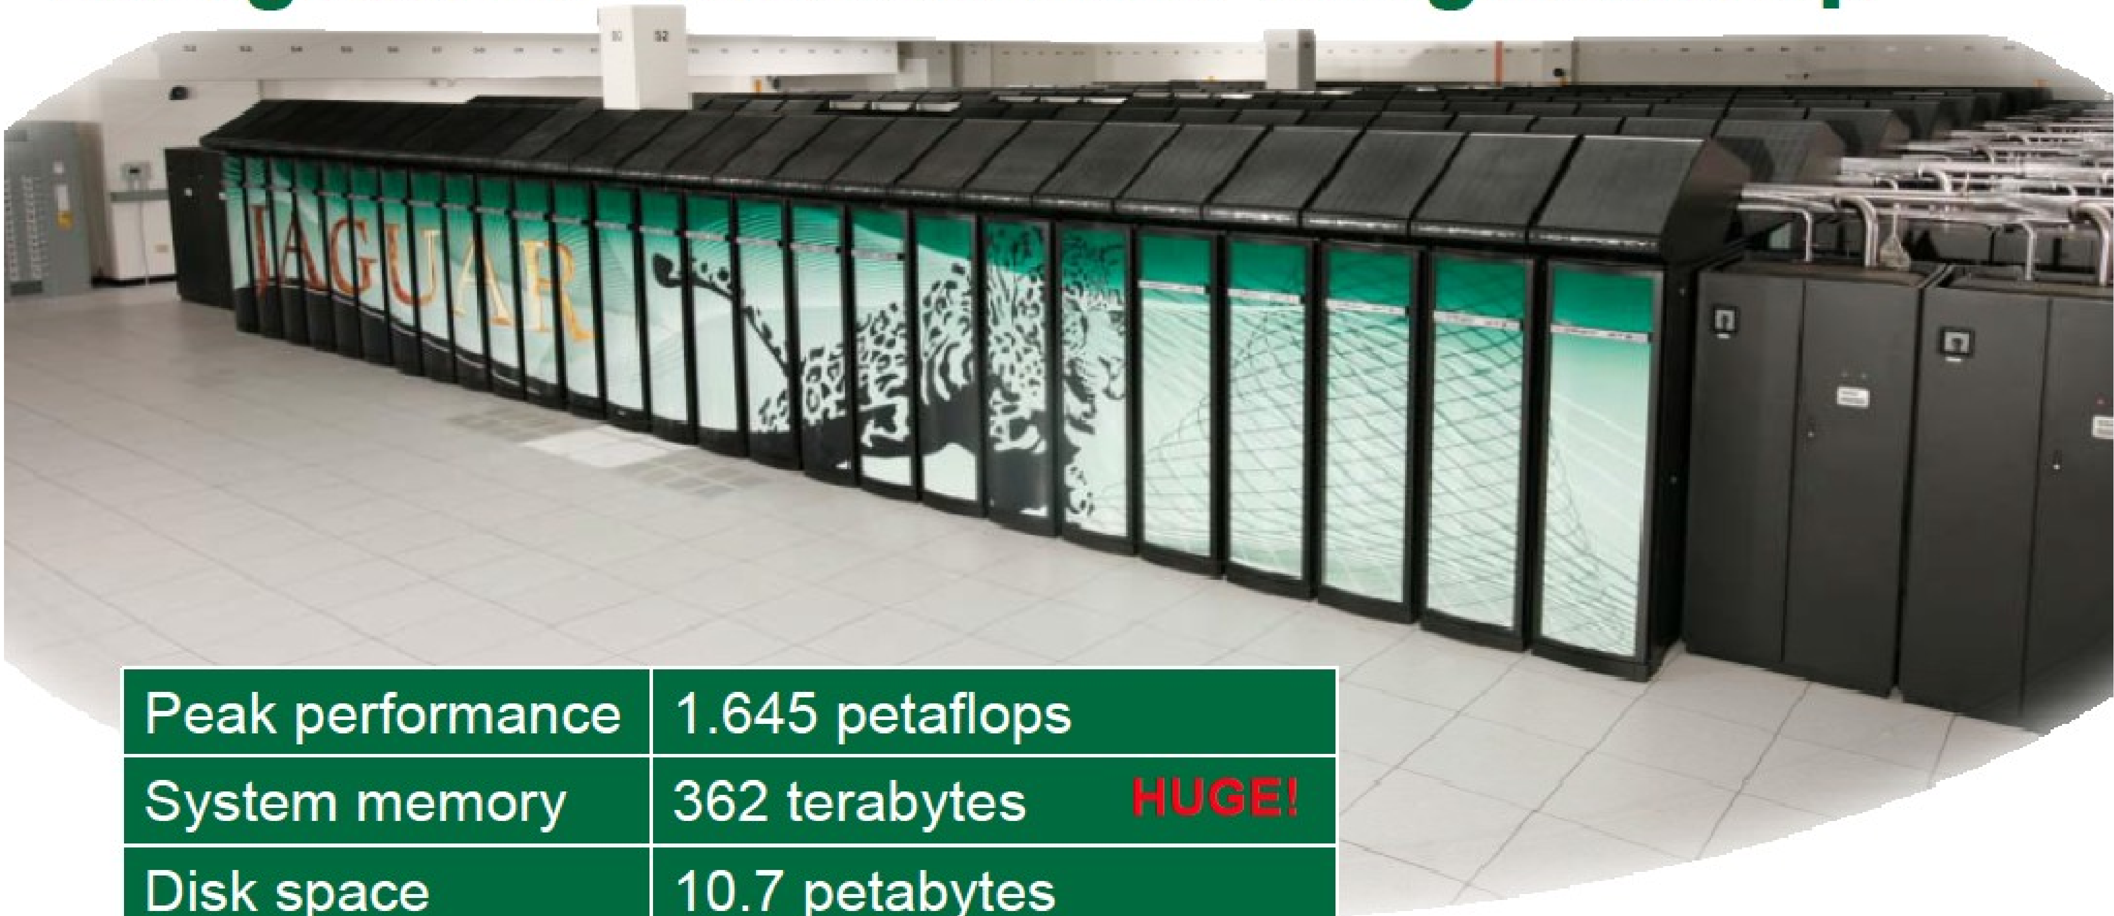

Peak performance	1.645 petaflops
System memory	362 terabytes **HUGE!**
Disk space	10.7 petabytes
Disk bandwidth	240 gigabytes/second
Interconnect BW	532 terabytes/second
Number cores	181504

60 years later - 10^{12} times faster

Currently the fastest supercomputers carry out about ~1 Petaflop, which are one thousand billion floating point operations per second

10^{18} in 2018

How long would the Millennium-XXL take on a Exaflop Supercomputer at peak performance?

15 min

One of the main problems:
Power Consumption

Petaflop Computer: **6 MW**

Exaflop Computer: **~ GW ?**

Need to get this down to **20-40 MW**

The number of cores on the top supercomputers grows exponentially

EXTREME GROWTH OF PARALLELISM

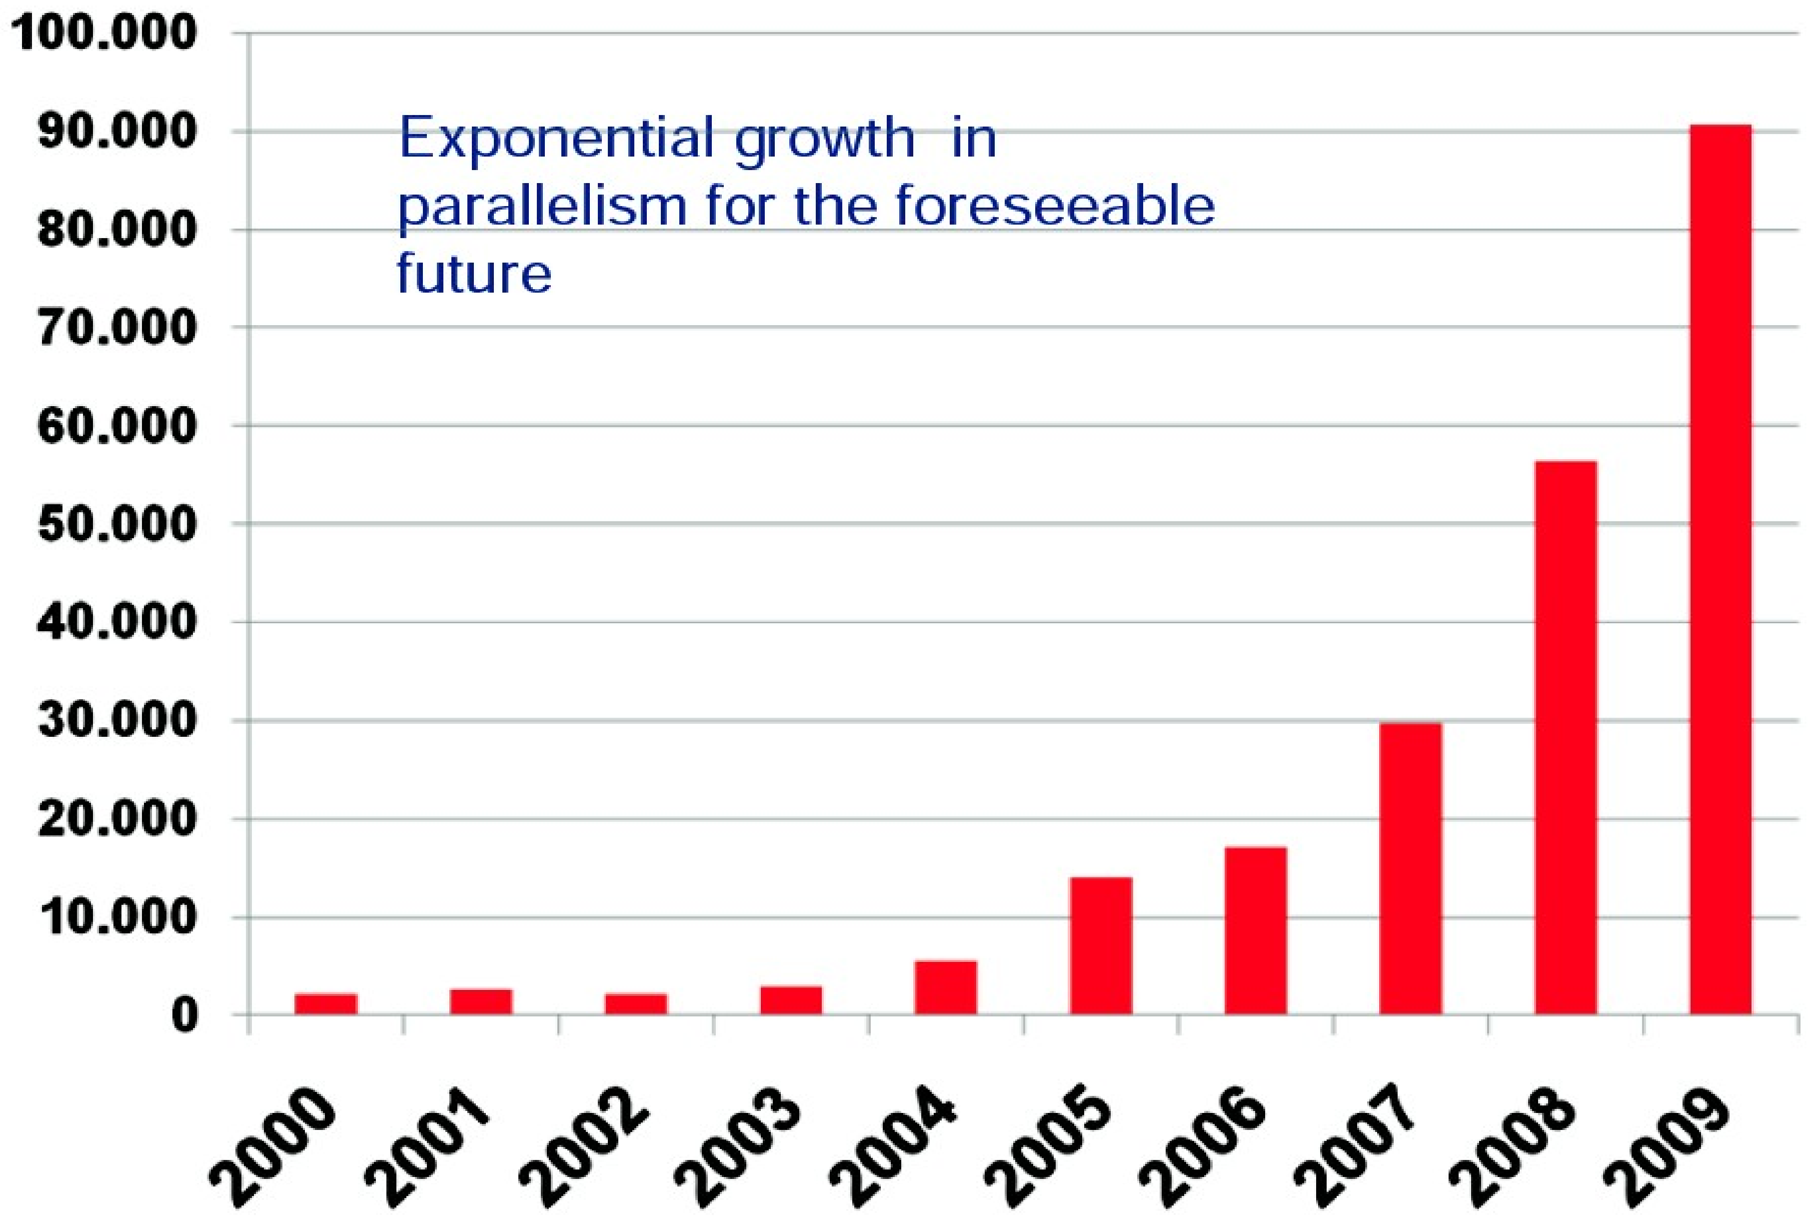

Challenges in exascale computing

CAN WE HAVE THE CAKE AND EAT IT?

- Memory per core decreasing.

- Applications need to deal with multiple hierarchies of memory.
  (especially on GPU-accelerated or hybrid systems)

- On systems with $>10^6$ cores, need fault-tolerant algorithms and codes.
  (resiliance has to be built into simulation codes)

- Cost of data access relative to floating point ops drastically increasing.

  Typical astrophysics codes run only at ~10% of the peak performance – and its getting worse with time.

None of our existing codes will survive and run on exascale platforms.

Astrophysics may be left behind in using these systems.

Hydrodynamical simulations

Dynamics of structure formation in baryonic matter

BASIC EQUATIONS

Astrophysical plasmas are extremely thin, with (usually) negligible viscosity

Euler equations of inviscid ideal gas dynamics

$$\frac{\partial \rho_c}{\partial t} + \frac{1}{a}\boldsymbol{\nabla}_c(\rho_c \boldsymbol{v}) = 0$$

$$\frac{\partial(\rho_c \boldsymbol{v})}{\partial t} + \frac{1}{a}\boldsymbol{\nabla}_c[(\rho_c \boldsymbol{v}\boldsymbol{v}^T + P_c)\boldsymbol{v}] = -H(a)\,\rho_c \boldsymbol{v} - \frac{\rho_c}{a^2}\boldsymbol{\nabla}_c \Phi_c$$

$$\frac{\partial(\rho_c e)}{\partial t} + \frac{1}{a}\boldsymbol{\nabla}_c[(\rho_c e + P_c)\boldsymbol{v}] = -2H(a)\,\rho_c e - \frac{\rho_c \boldsymbol{v}}{a^2}\boldsymbol{\nabla}_c \Phi_c$$

$$\nabla_c^2 \Phi_c = 4\pi G\,[\rho_c(\boldsymbol{x}) - \overline{\rho_c}]$$

Important hydrodynamical processes

Shock waves
Turbulence
Radiative transfer
Magnetic fields
Star formation
Supernova explosions
Black holes, etc...

Supersonic motion creates shock waves

SHOCK WAVES OF A BULLET TRAVELLING IN AIR

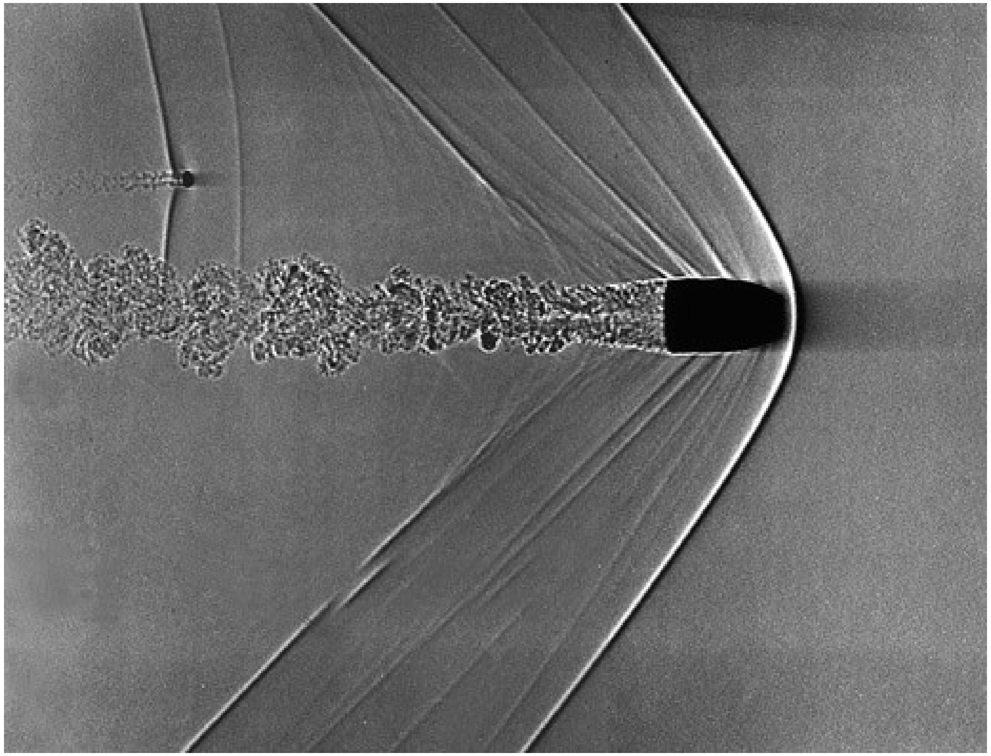

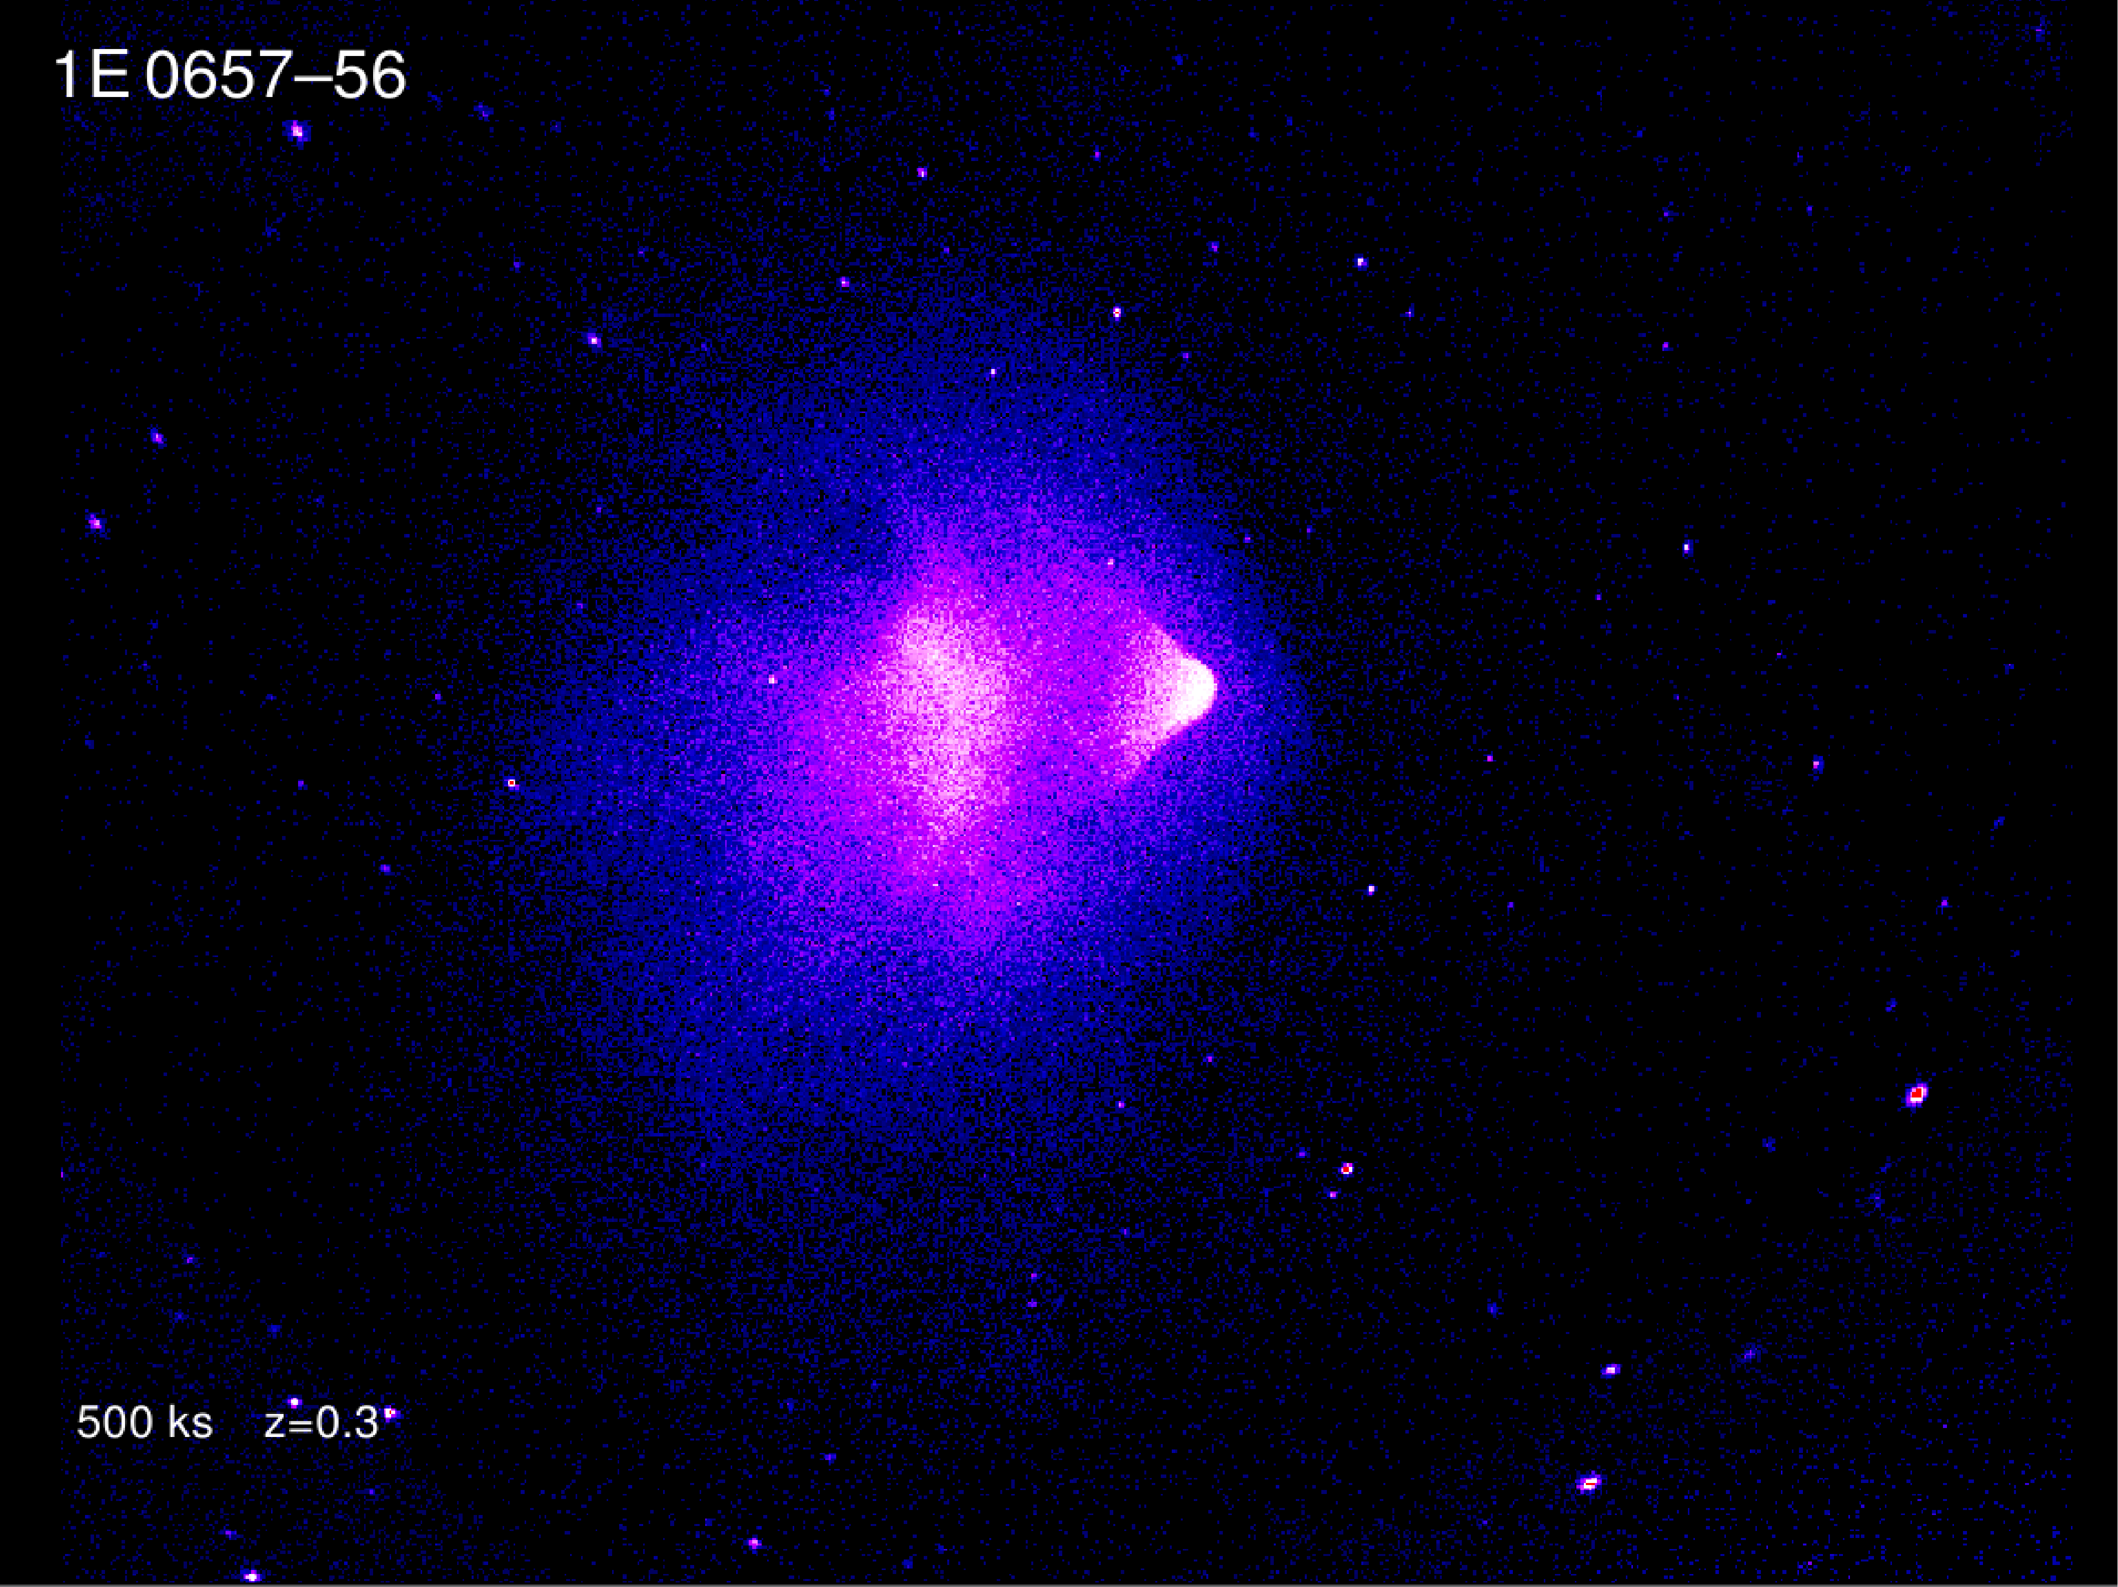
1E 0657–56
500 ks
z=0.3

Weak lensing mass reconstructions have confirmed an offset between mass peaks and X-ray emission

MASS CONTOURS FROM LENSING COMPARD TO X-RAY EMISSION

Clowe et al. (2006)

Magellan Optical Image

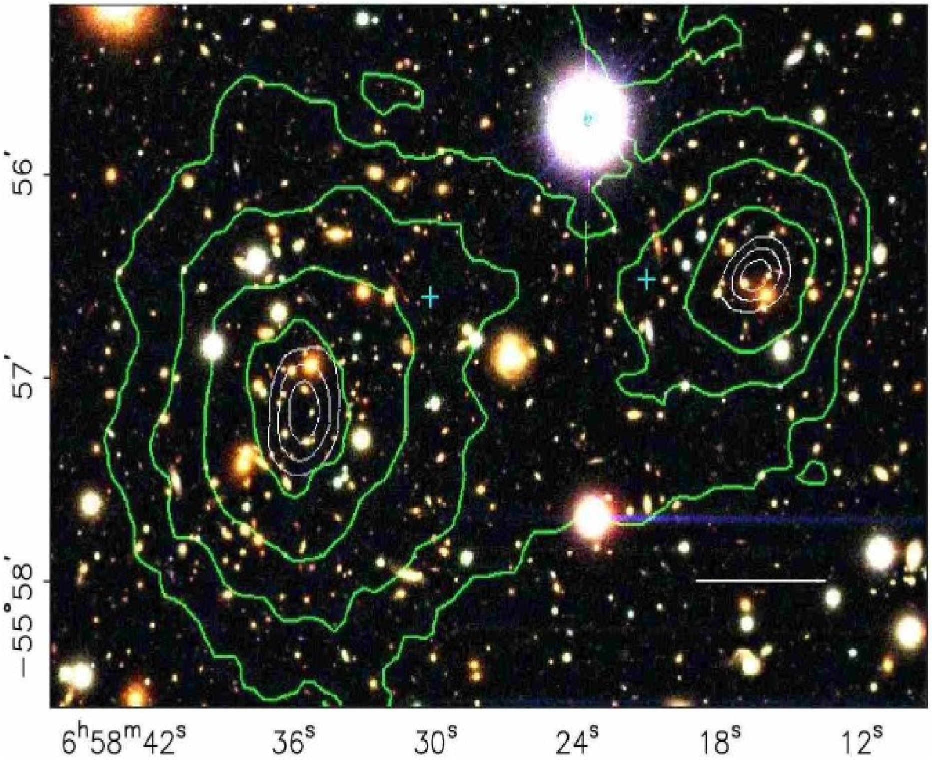

500 ksec Chandra exposure

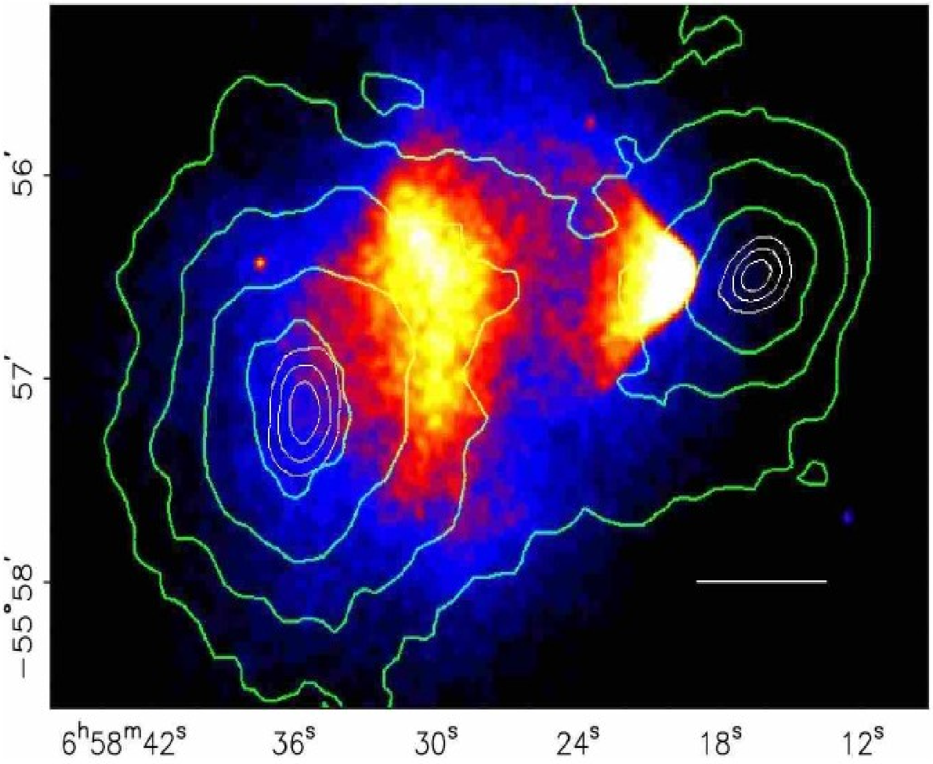

weak lensing mass contours overlaid

NASA Press Release Aug 21, 2006:

1E 0657-56: NASA Finds Direct Proof of Dark Matter

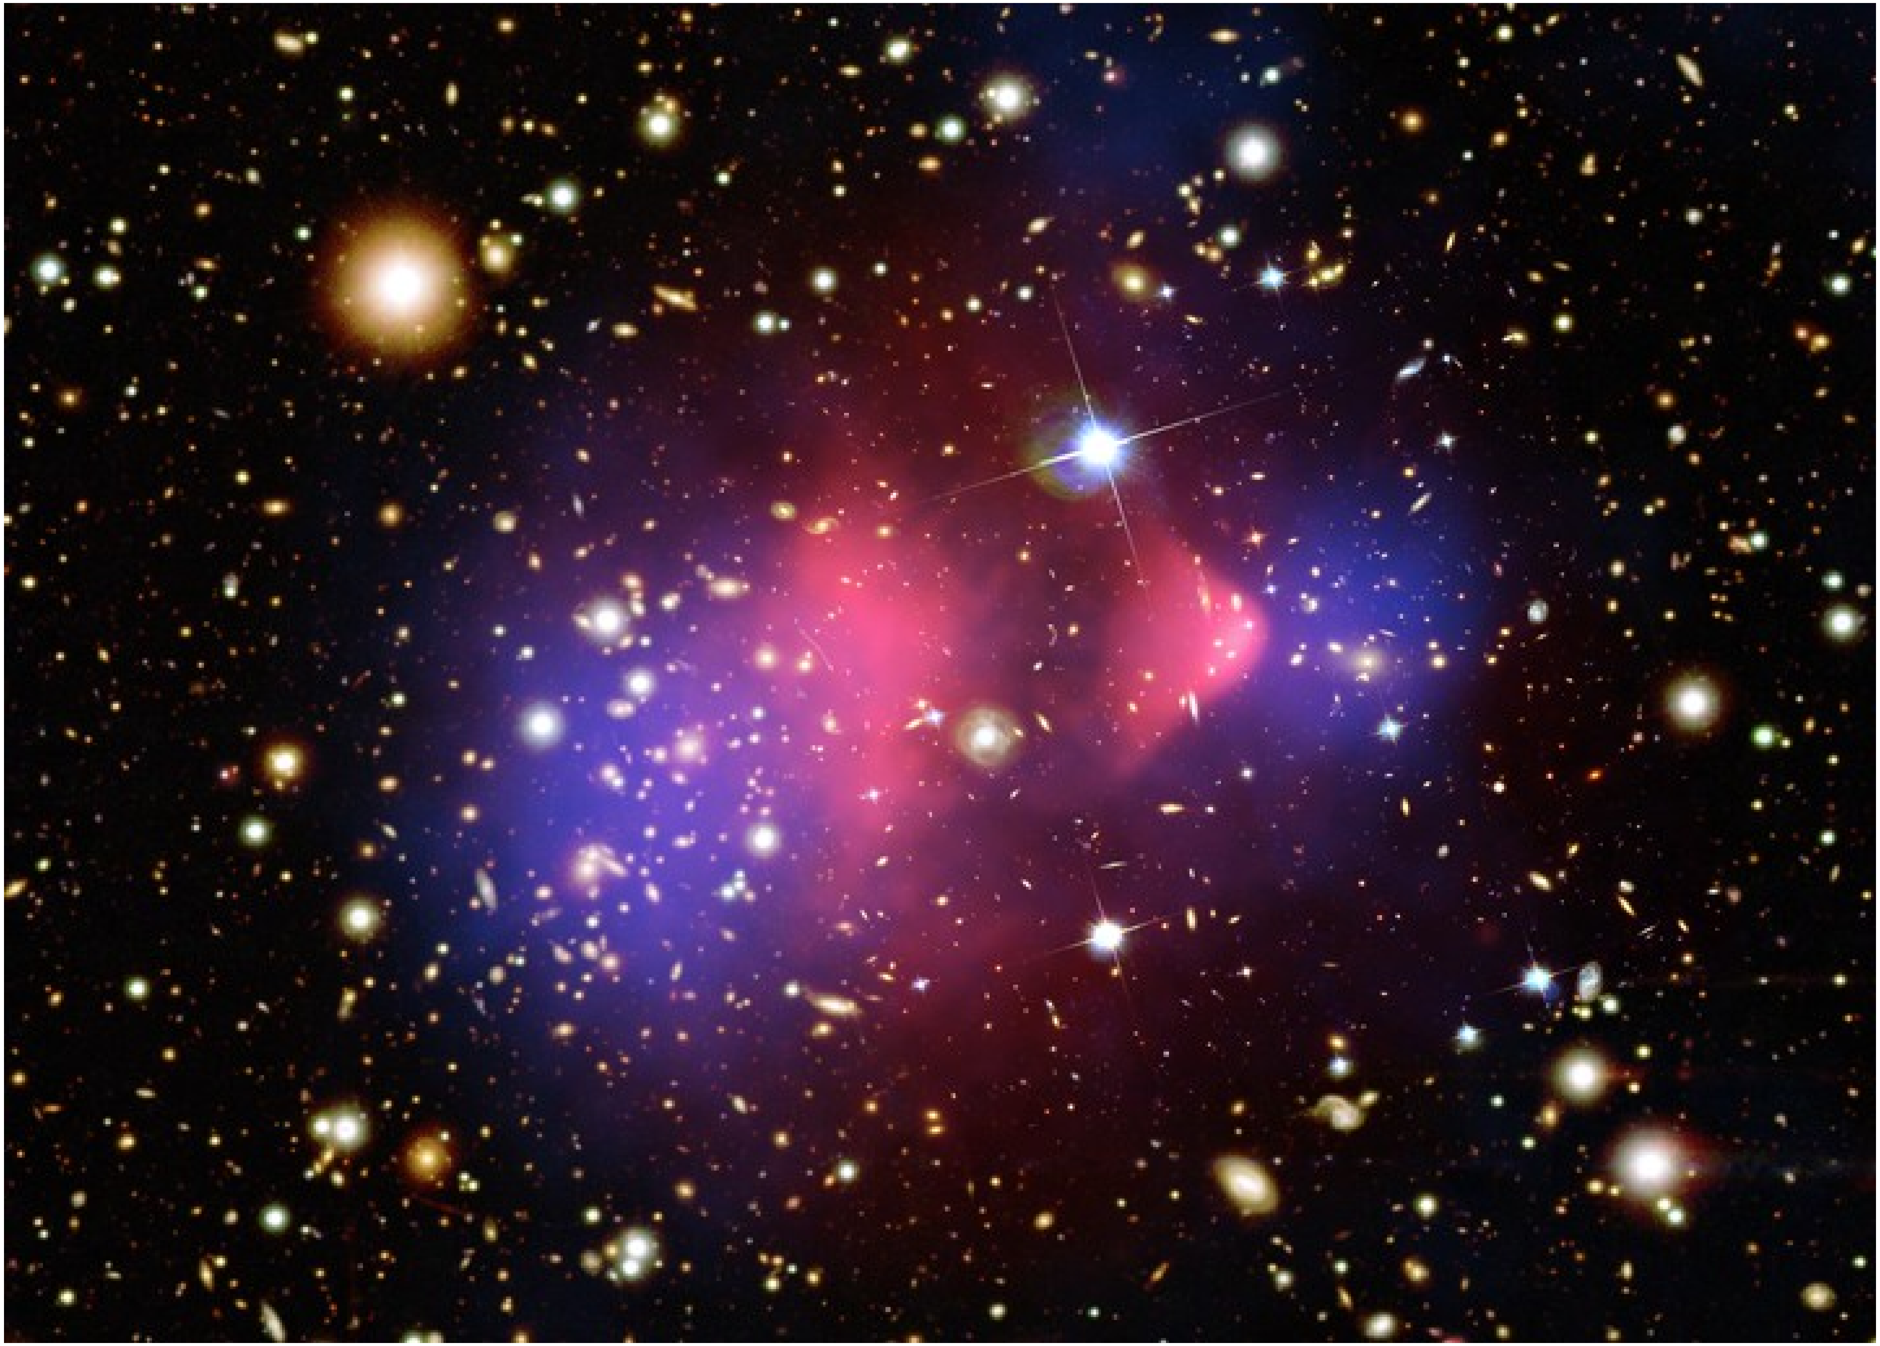

Fitting the density jump in the X-ray surface brightness profile allows a measurement of the shock's Mach number

X-RAY SURFACE BRIGHTNESS PROFILE

Markevitch et al. (2006)

shock strength:
M = 3.0 ± 0.4

shock velocity:
v_s = 4700 km/s

Usually, shock velocity has been identified with velocity of the bullet.

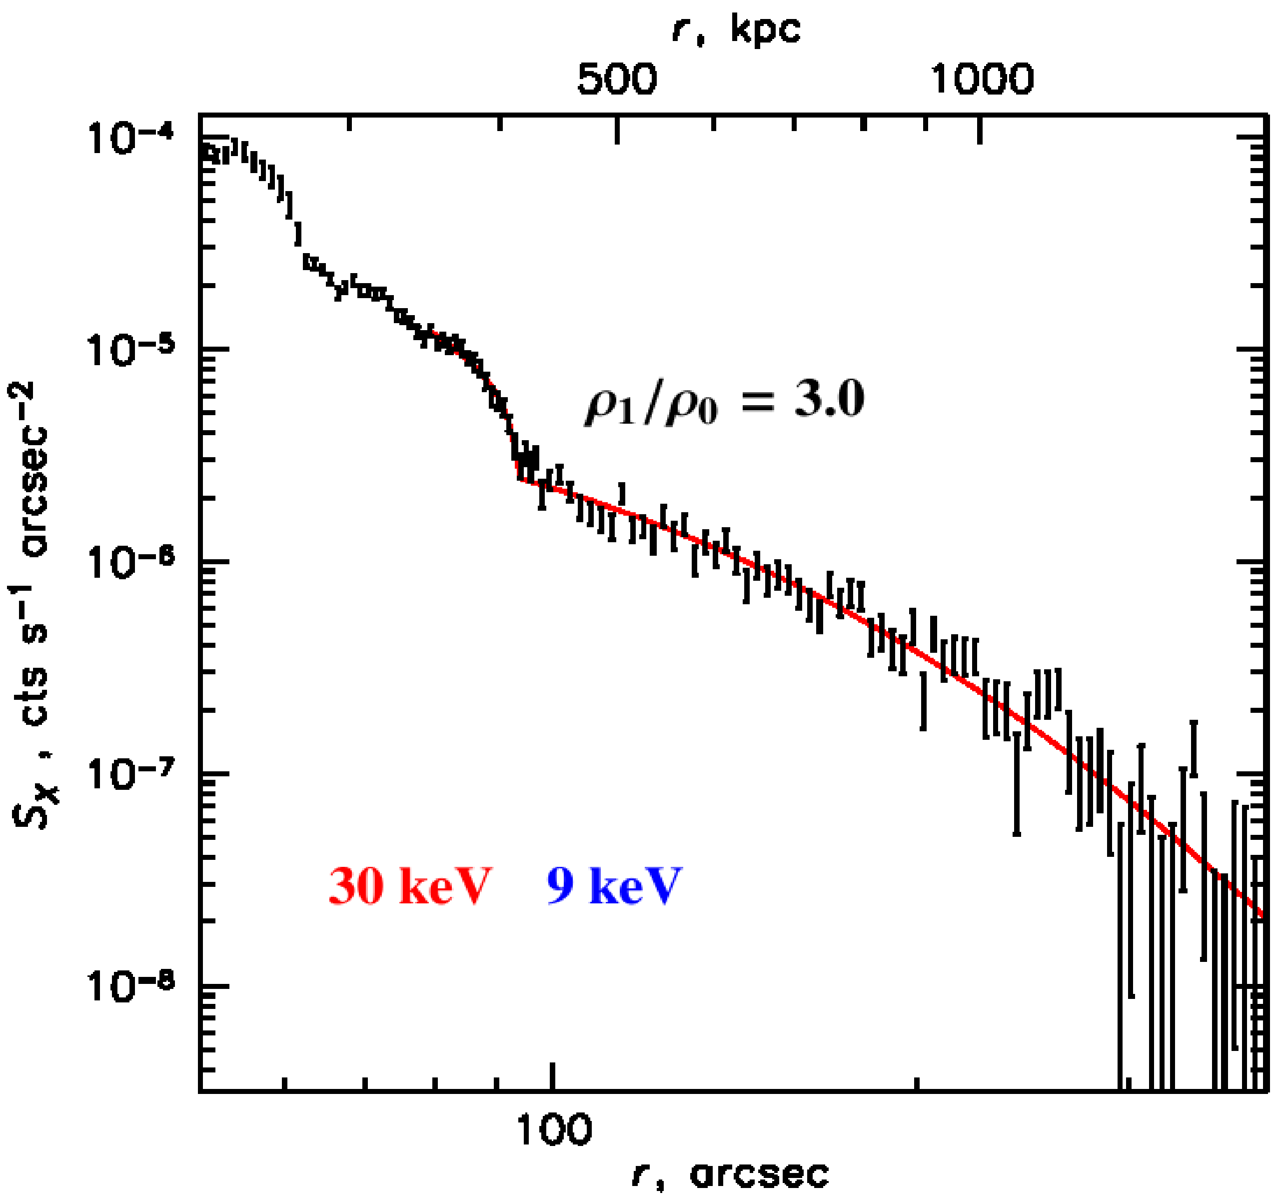

How rare is the bullet cluster?

DISTRIBUTION OF VELOCITIES OF THE MOST MASSIVE SUBSTRUCTURE IN THE MILLENNIUM RUN

Hayashi & White (2006)

Adopted mass model from Clowe et al. (2004):

NFW-Halo with:

M_{200} = 2.96 x 10^{15} $M_\odot$
R_{200} = 2.25 Mpc

V_{200} = 2380 km/sec
V_{shock} = 4500 km/sec

V_{sub}/V_{shock} = 1.9 chance: 10^{-2}

But, revised data from Clowe et al. (2006) and Markevitch el al. (2006):

M_{200} = 1.5 x 10^{15} $M_\odot$

V_{200} = 1680 km/sec
V_{shock} = 4740 km/sec

V_{sub}/V_{shock} = 2.8 chance: 10^{-7}

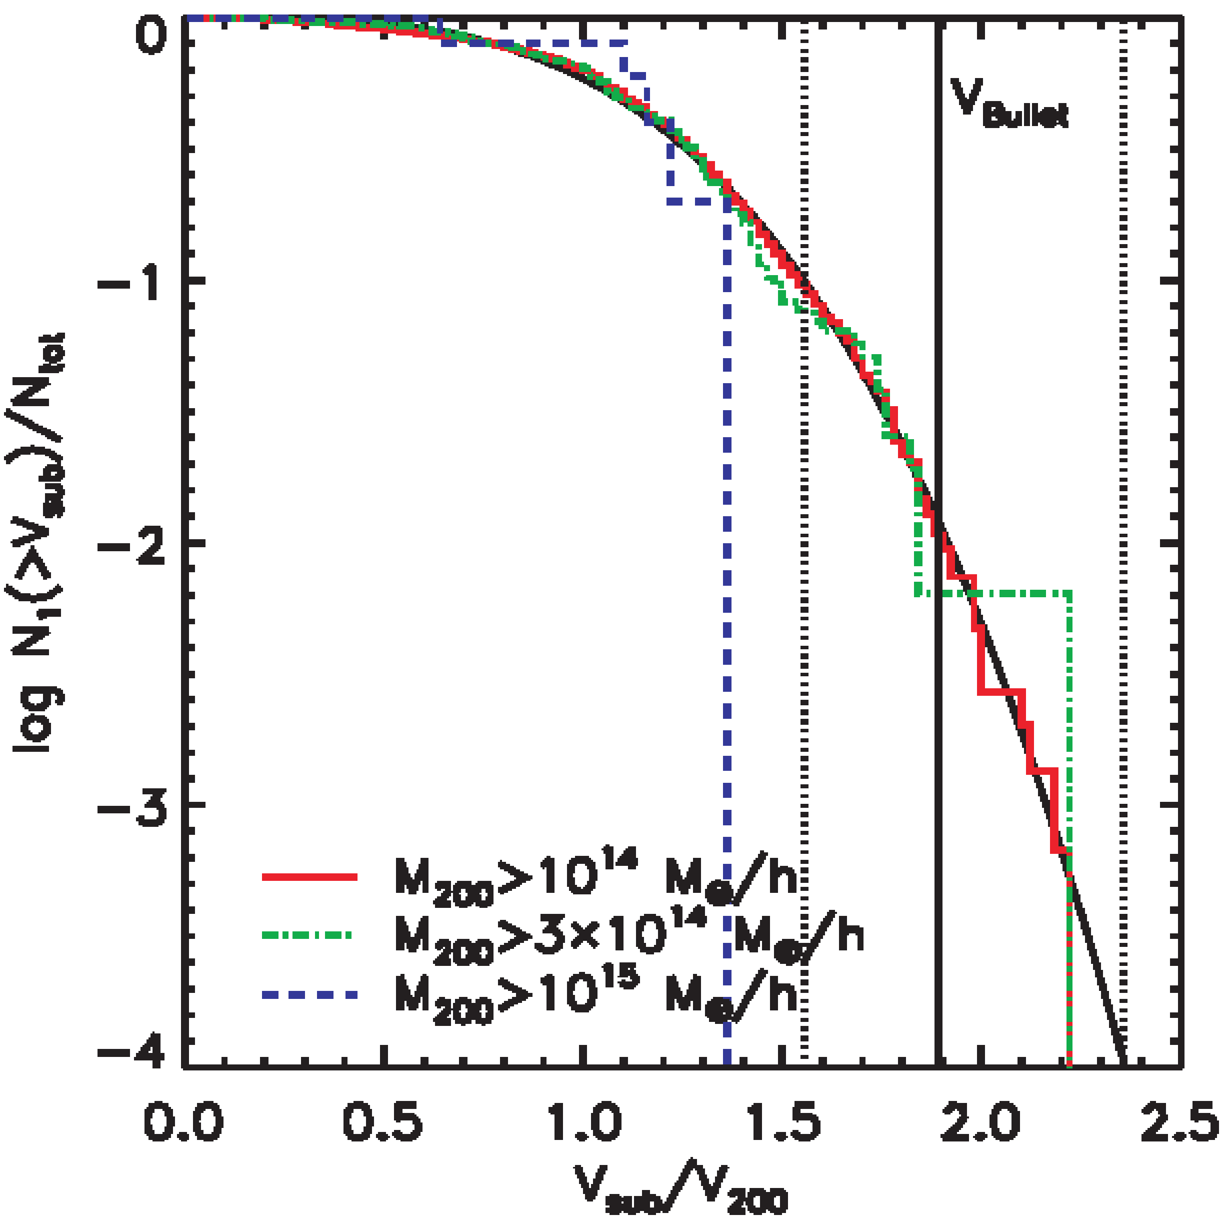

A simple toy merger model of two NFW halos on a zero-energy collision orbit

PARAMETERS OF A BASIC TOY MODEL

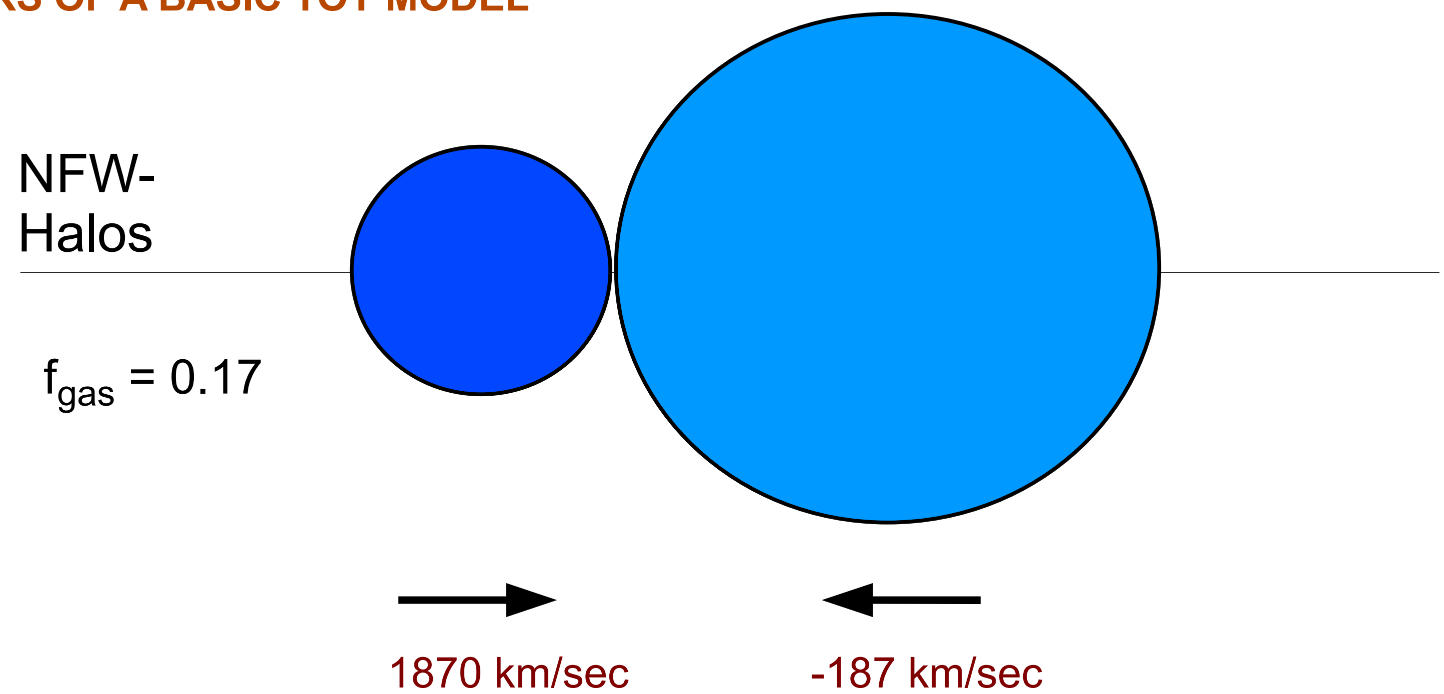

Mass model from Clowe et al. (2006):

$\mathbf{M_{200} = 1.5 \times 10^{14}\ M_\odot}$	$\mathbf{M_{200} = 1.5 \times 10^{15}\ M_\odot}$
R_{200} = 1.1 Mpc	R_{200} = 2.3 Mpc
c = 7.2	c = 2.0
V_{200} = 780 km/sec	V_{200} = 1680 km/sec

VIDEO OF THE TIME EVOLUTION OF A SIMPLE BULLET CLUSTER MODEL

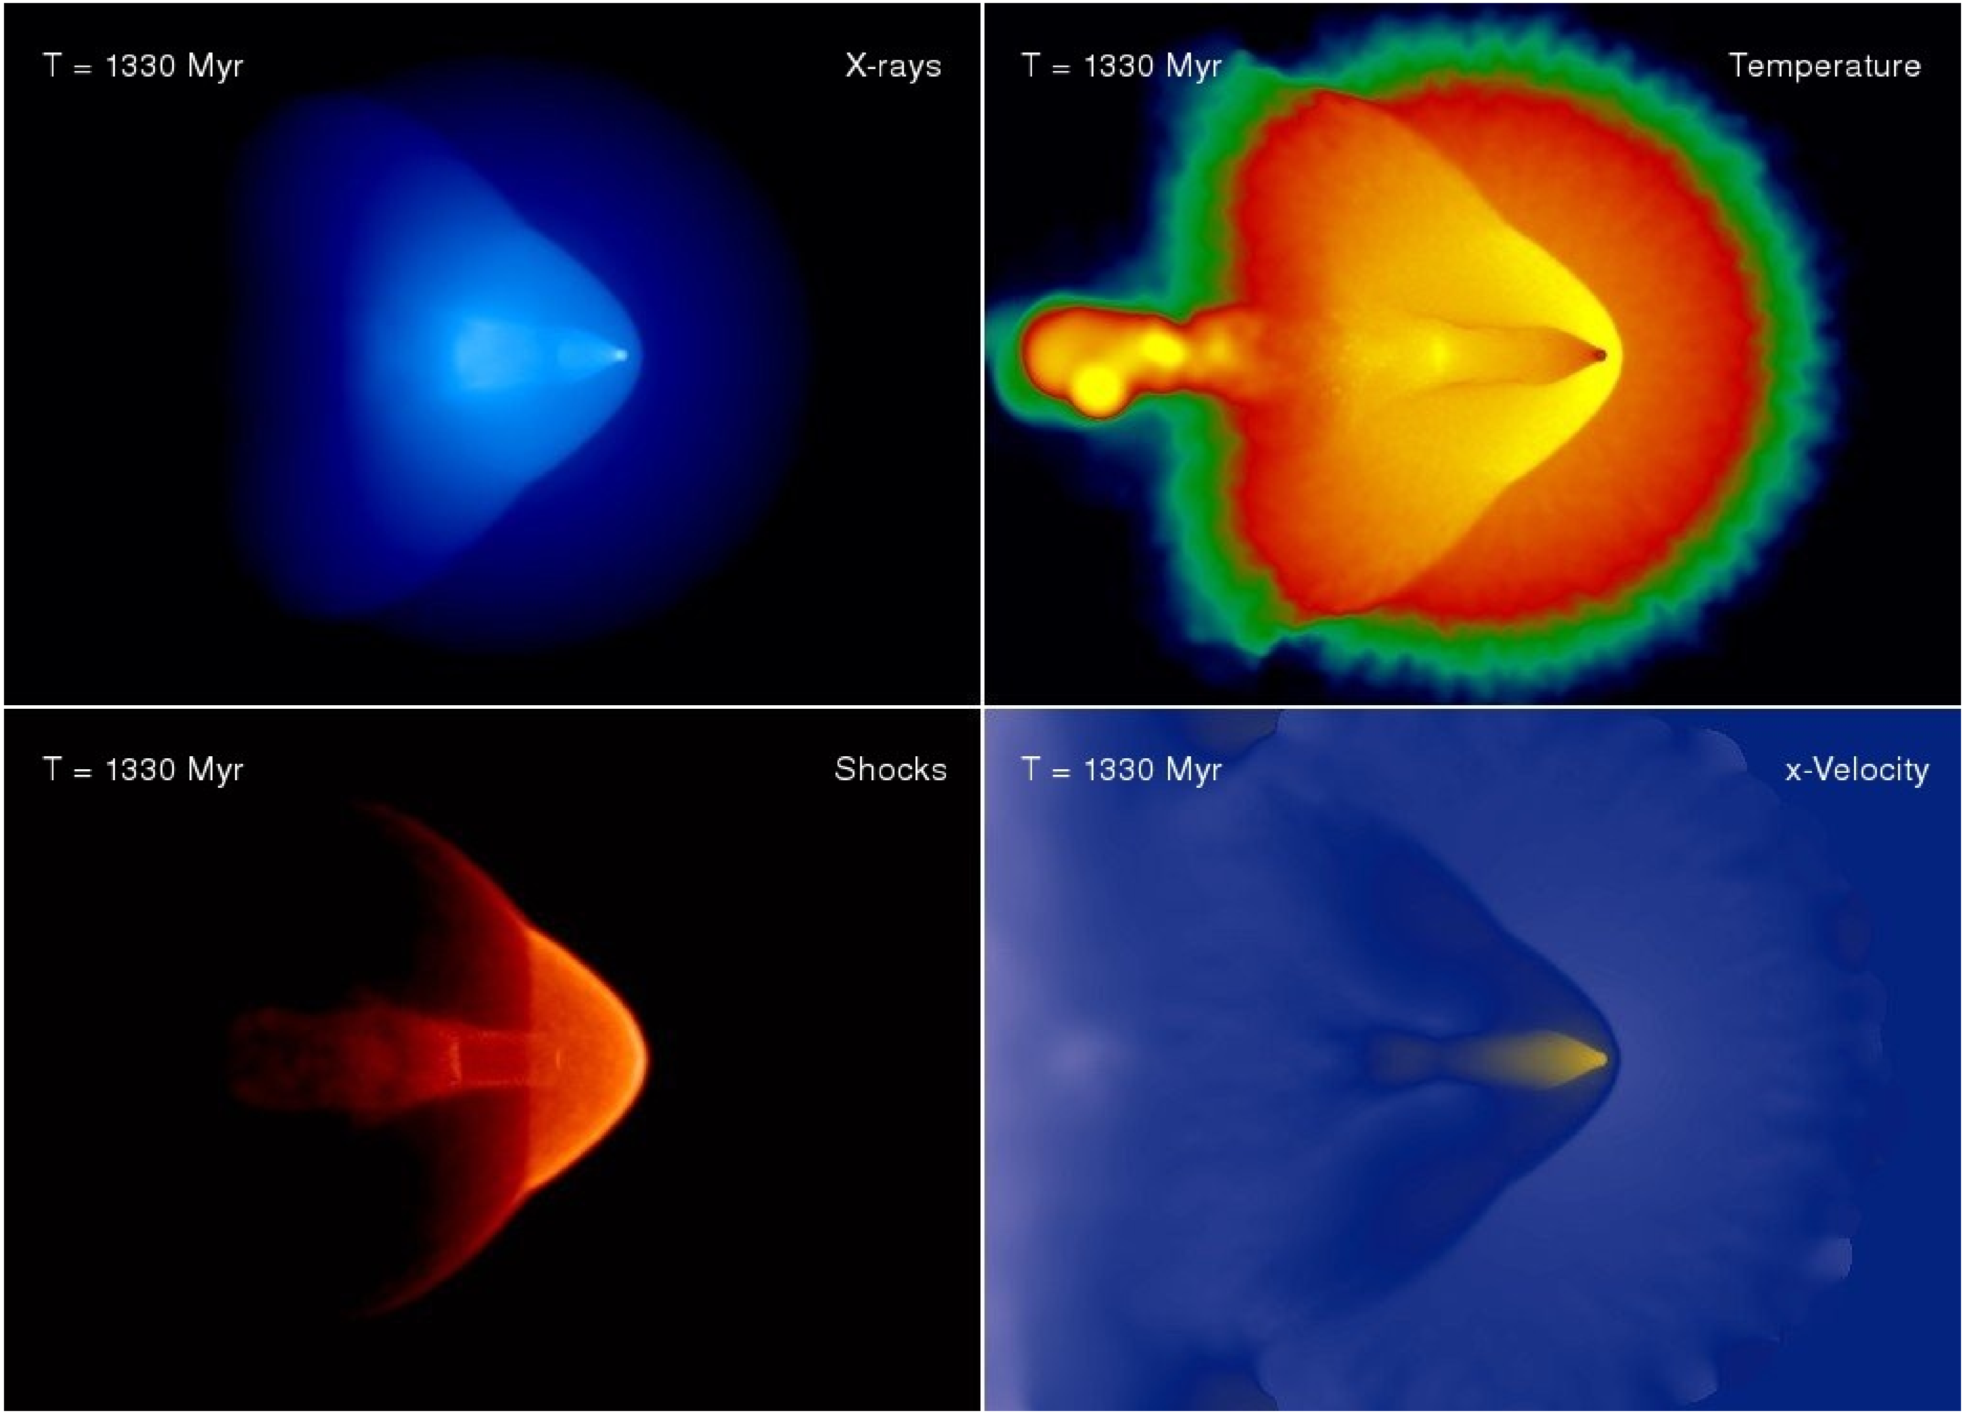

Drawing the observed X-ray map and the simulation images with the same color-scale simplifies the comparison

SIMULATED X-RAY MAP COMPARED TO OBSERVATION

Candra 500 ks image

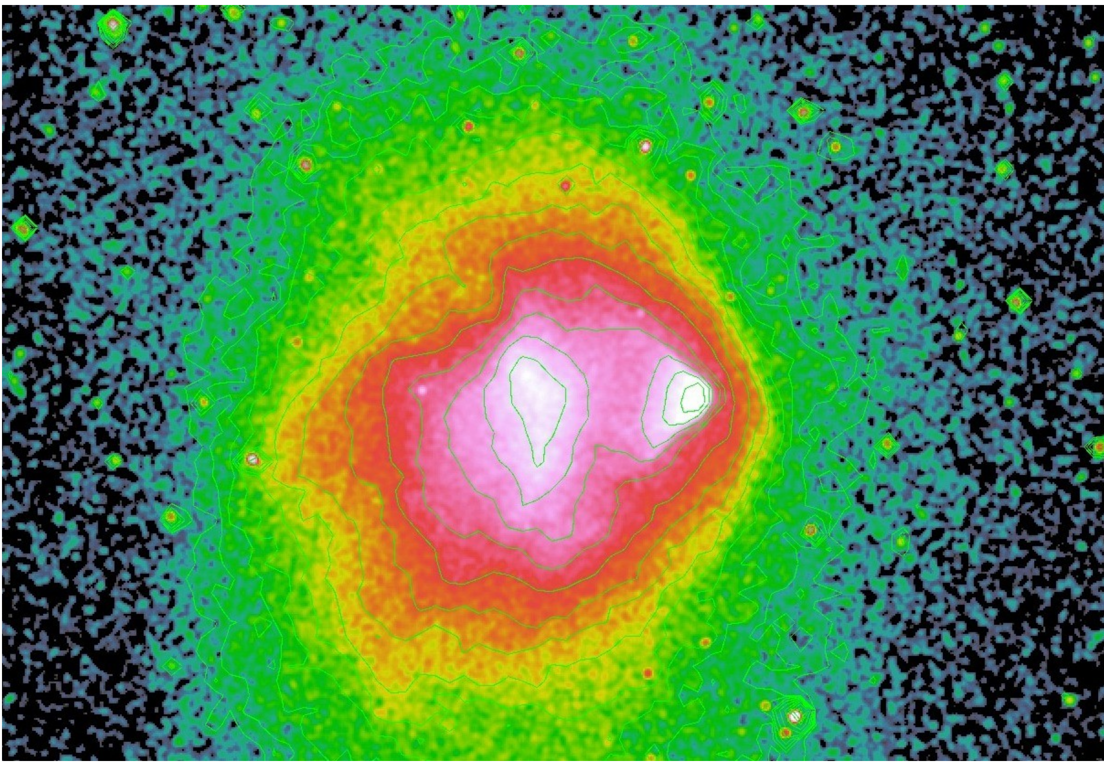

bullet cluster simulation

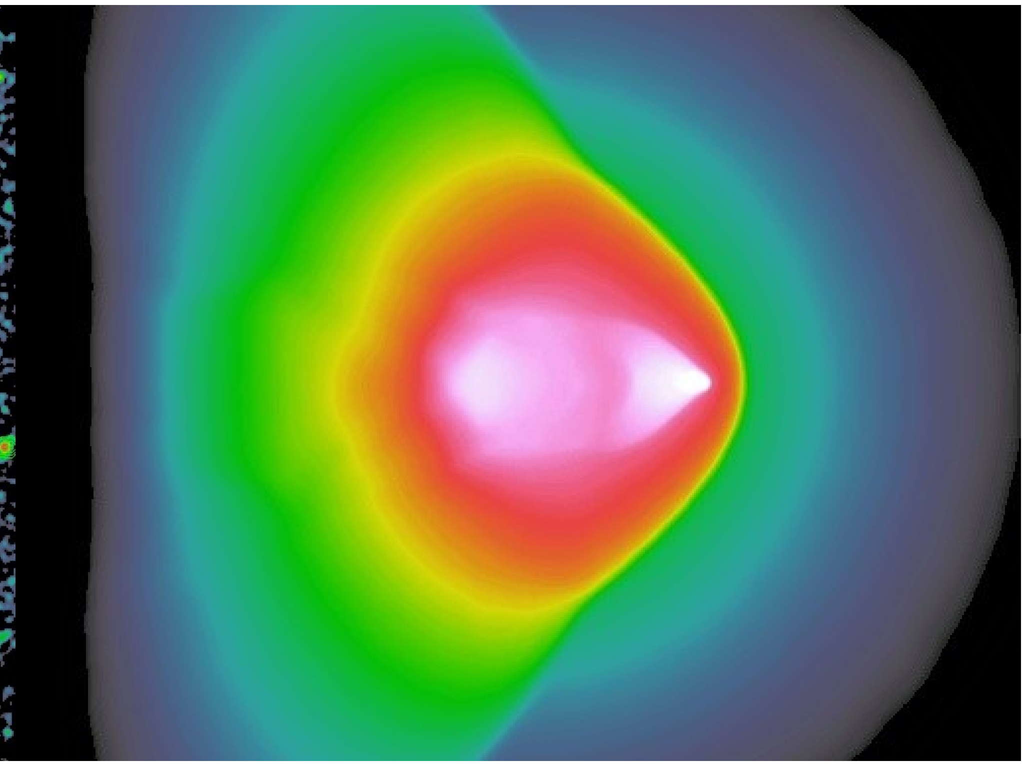

Springel & Farrar (2007)

The model also matches the observed temperature and mass profiles

COMPARISON OF SIMULATED TEMPERATURE AND MASS PROFILE WITH OBSERVATIONS

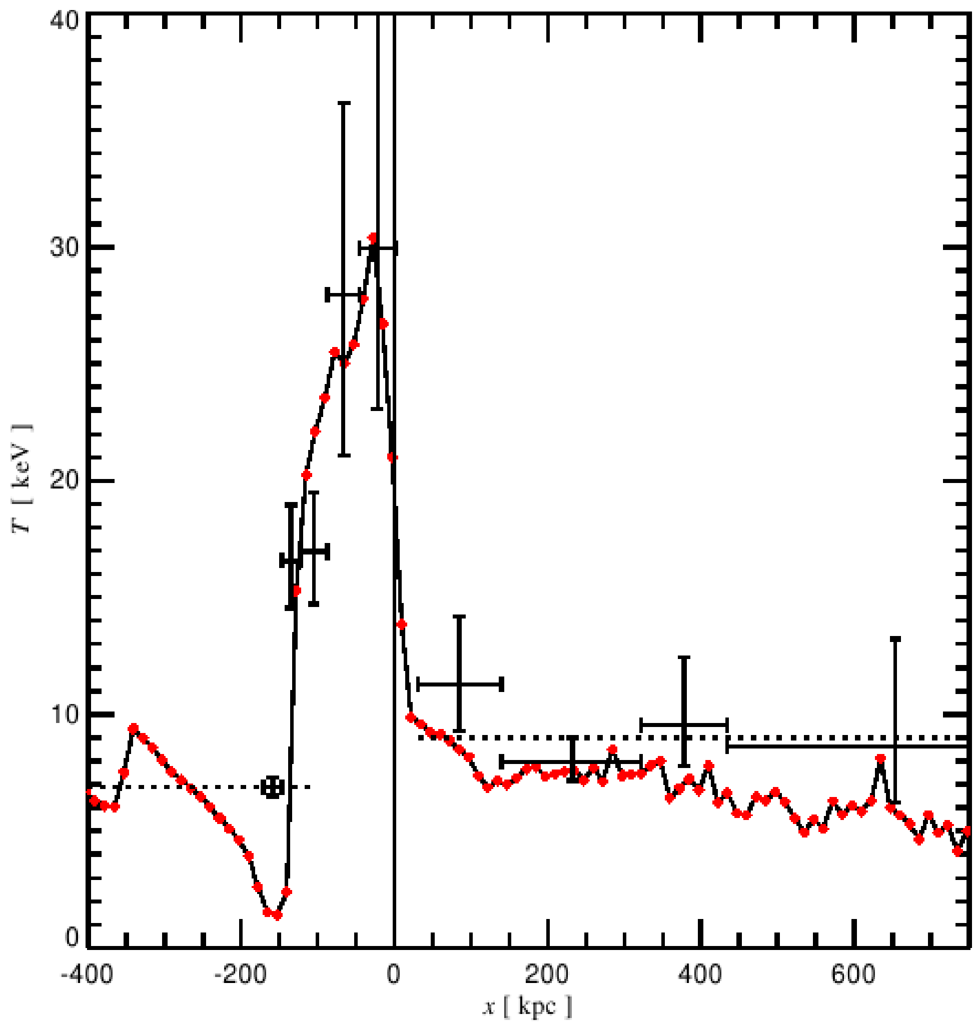

Data from **Markevitch et al. (2006)**

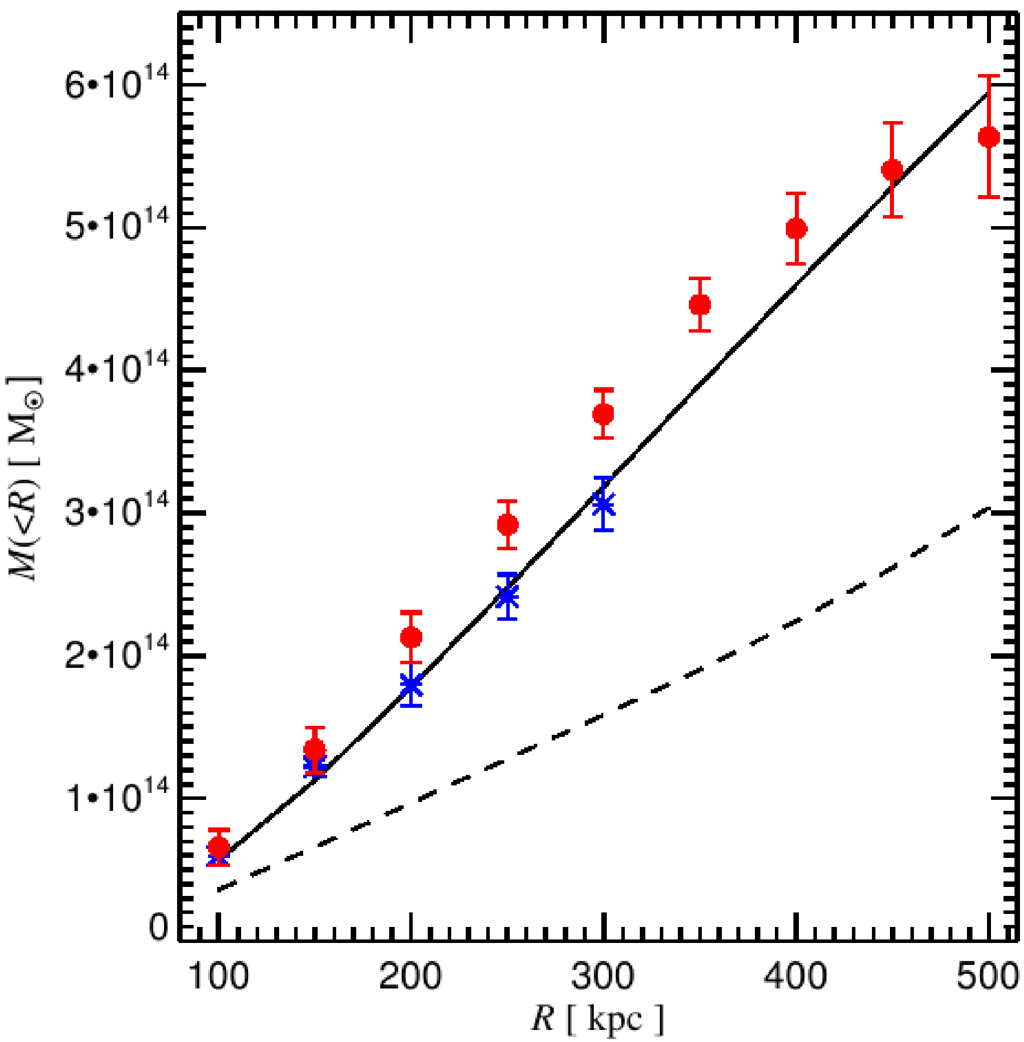

Data from **Bradac et al. (2006)**

Despite a shock speed of ~4500 km/s, the bullet moves considerably slower

VELOCITIES AND POSITIONS OF MAIN BULLET CLUSTER FEATURES AS A FUNCTION OF TIME

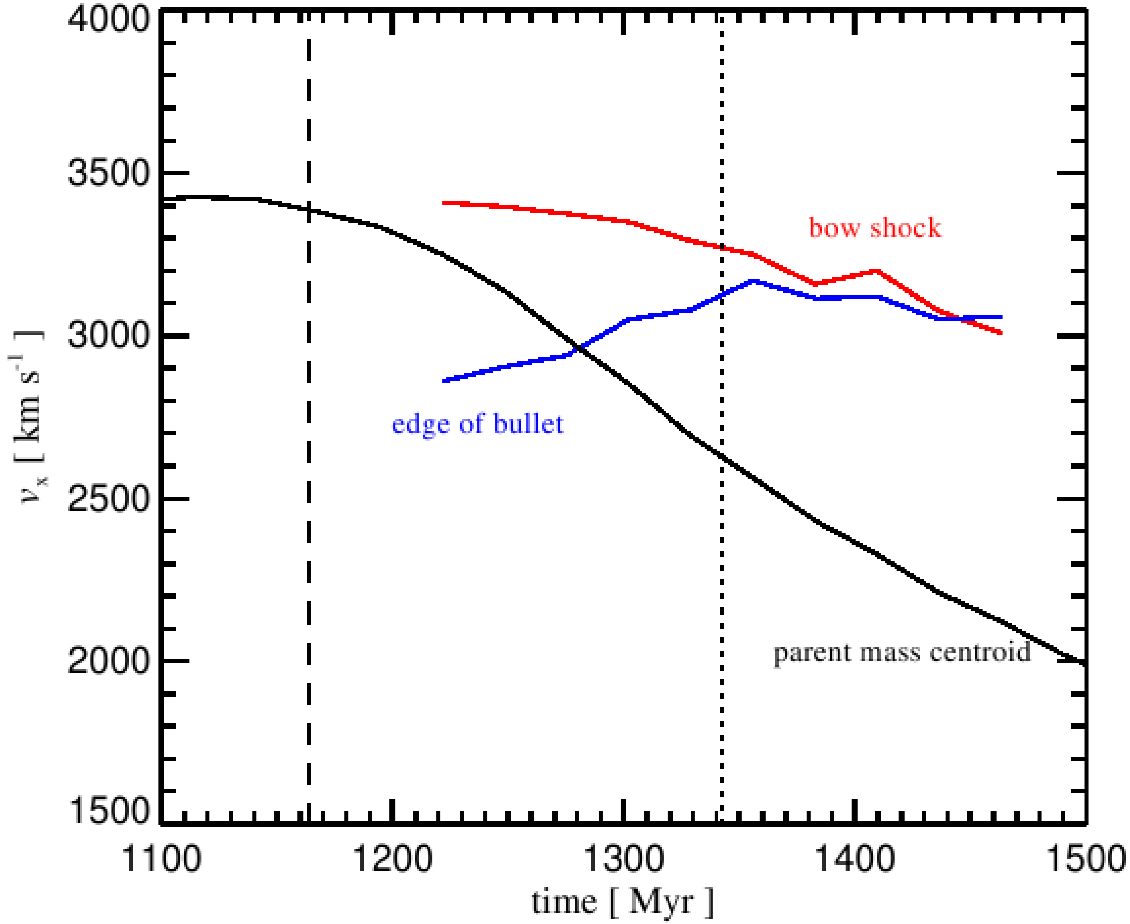

Shock speed: 4500 km/s

Pre-shock infall: -1100 km/s

Shock speed relative to bullet: -800 km/s

Speed of bullet: 2600 km/s

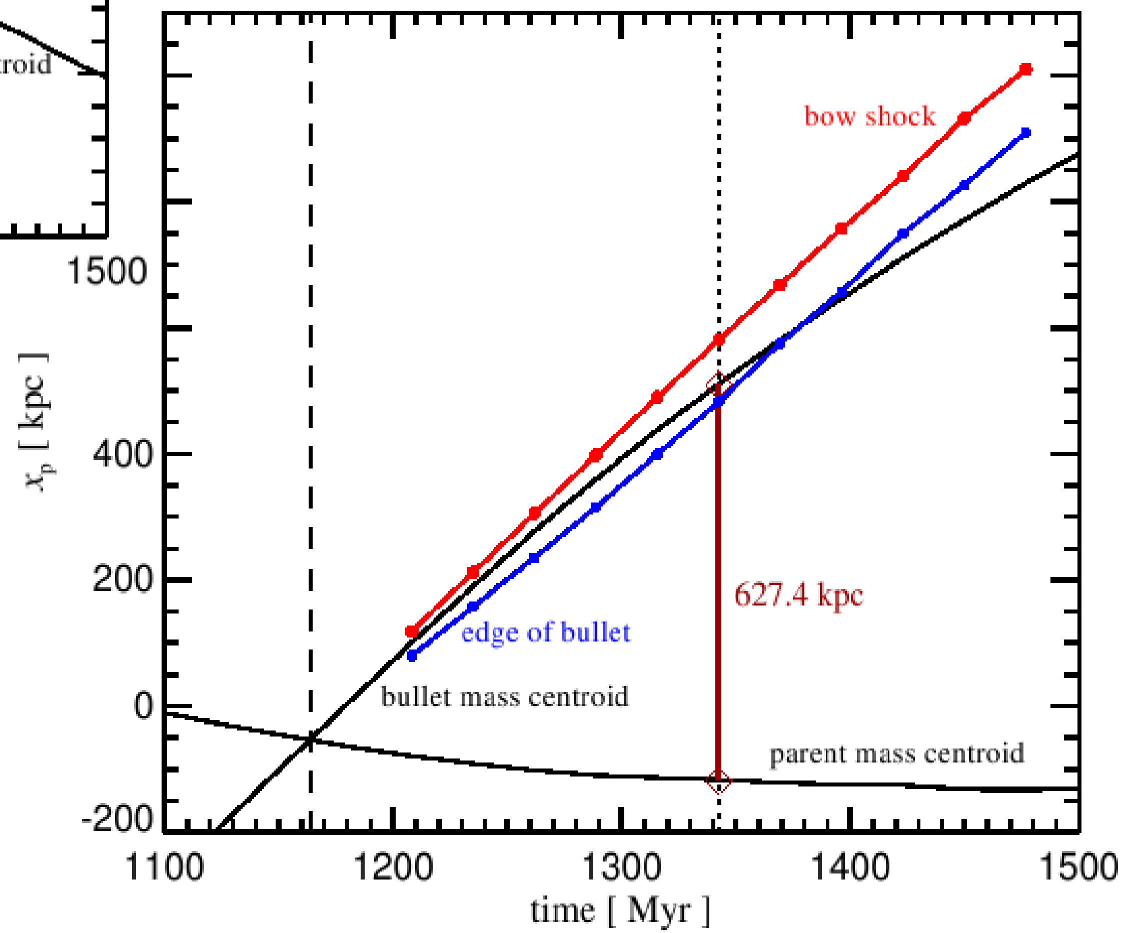

Uncertainties and errors in hydrodynamical numerical techniques

A cloud moving through ambient gas shows markedly different long-term behavior in SPH and Eulerian mesh codes

DISRUPTION OF A CLOUD BY KELVIN-HELMHOLTZ INSTABILITIES

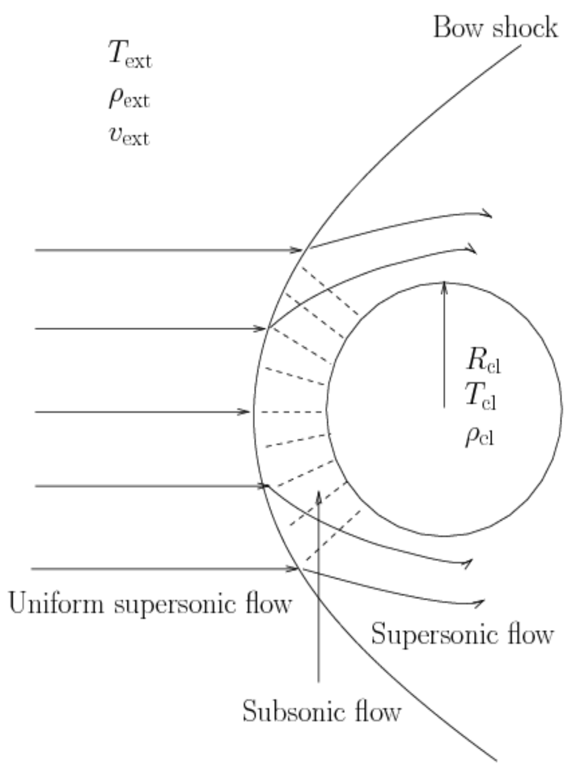

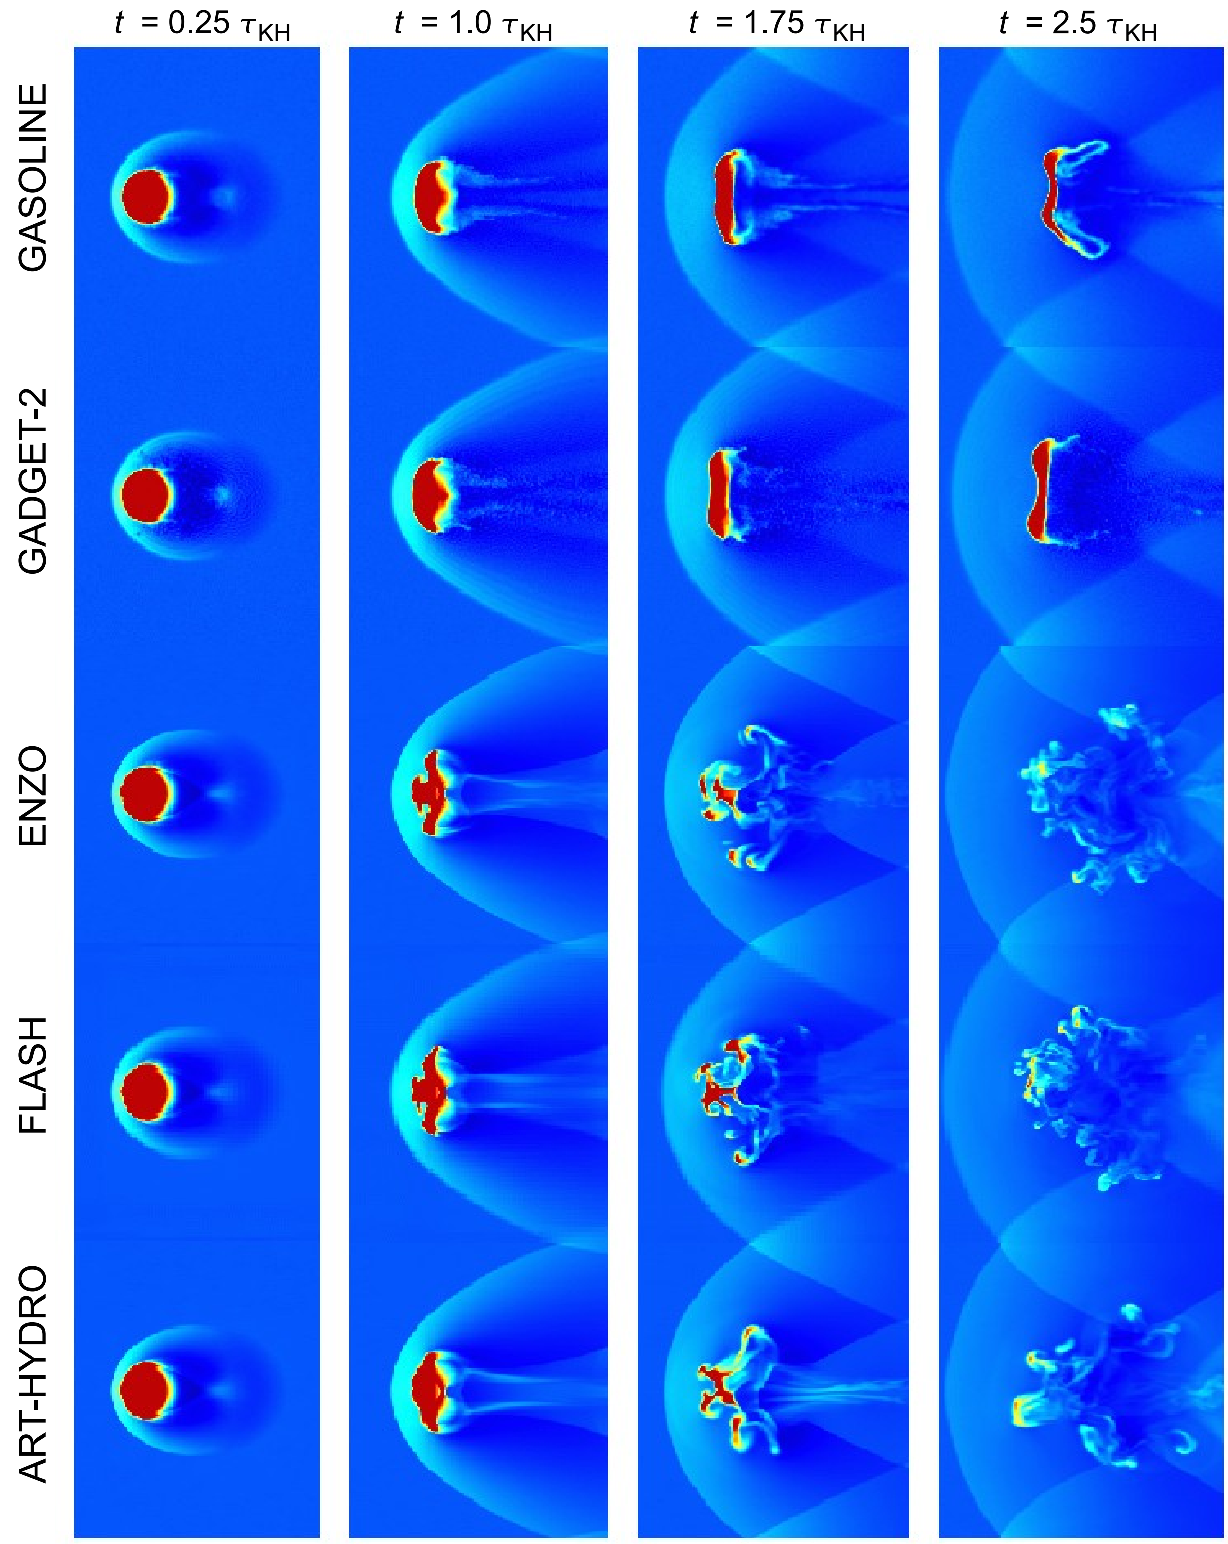

Agertz et al. (2007)

Moving-mesh hydrodynamics with AREPO

Volker Springel

Heidelberg Institute for Theoretical Studies

HITS

Universität Heidelberg

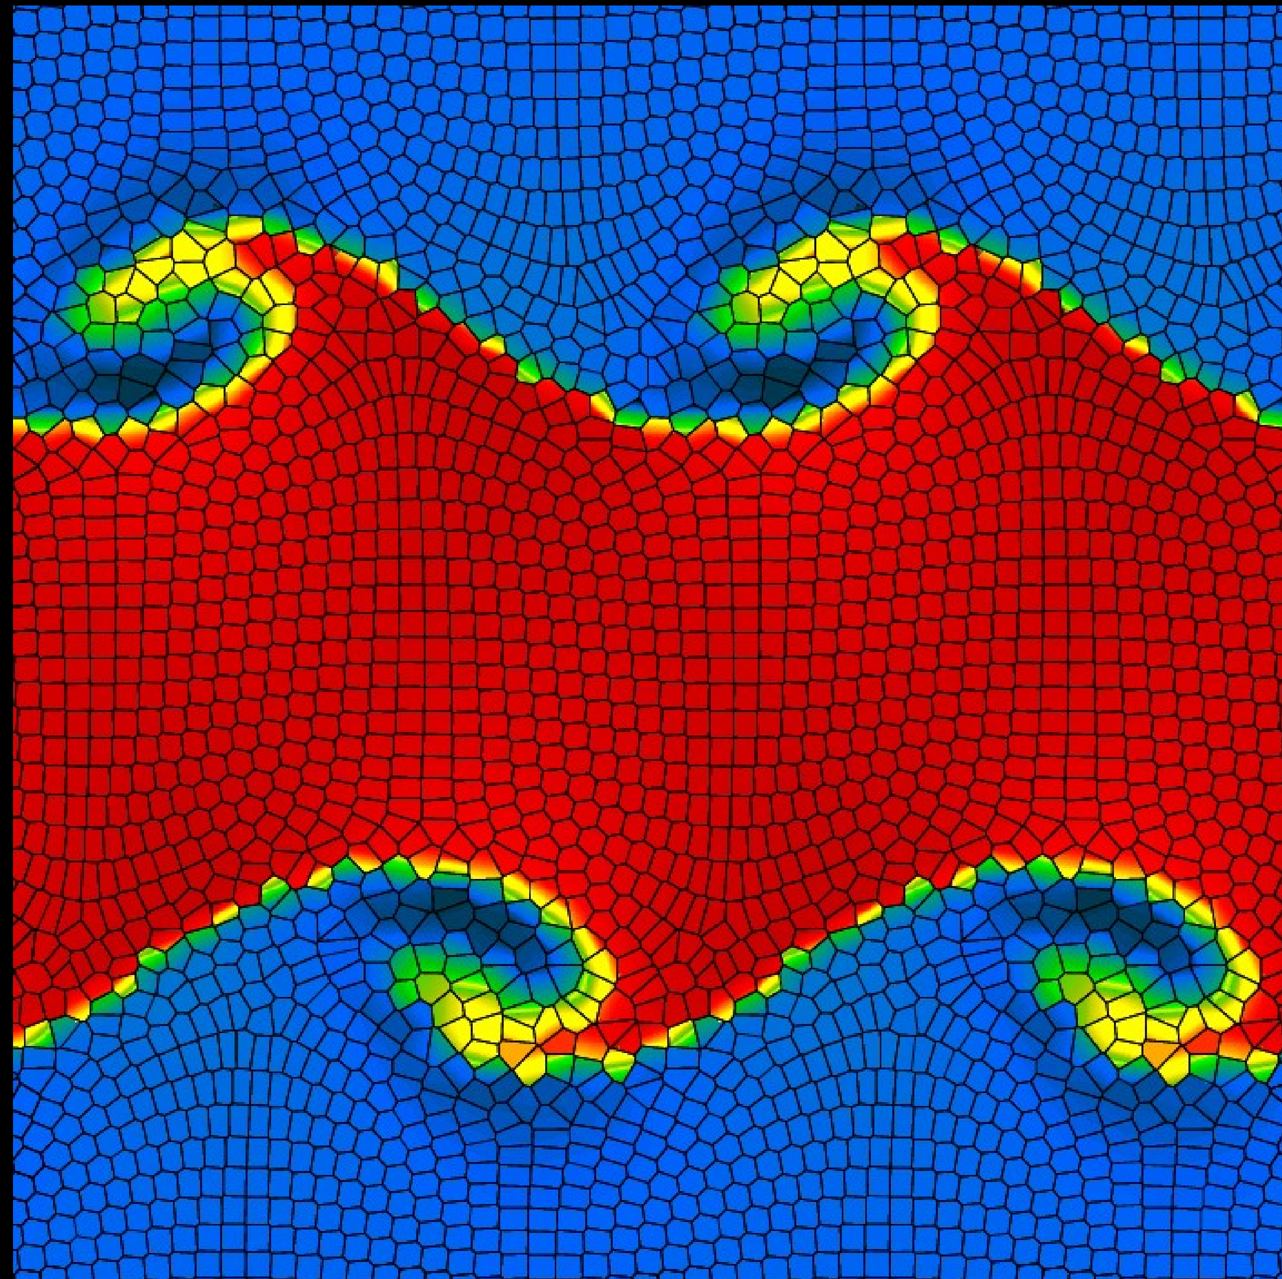

Voronoi and Delaunay tessellations provide unique partitions of space based on a given sample of mesh-generating points

BASIC PROPERTIES OF VORONOI AND DELAUNAY MESHES

Voronoi mesh

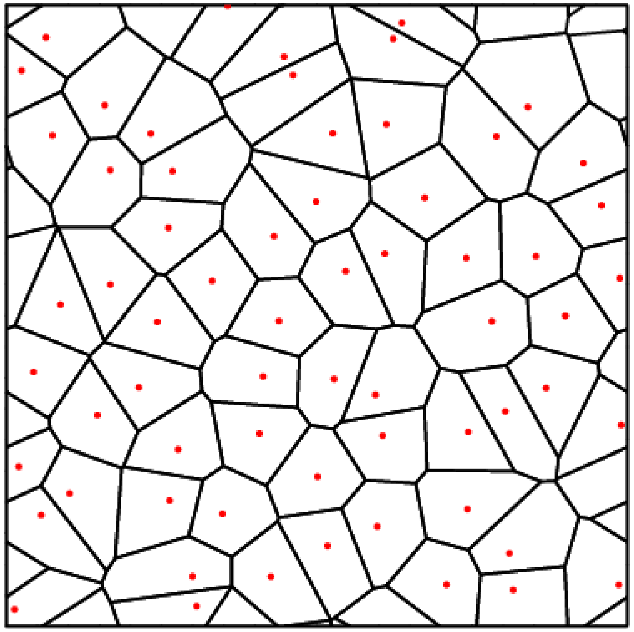

Delaunay triangulation

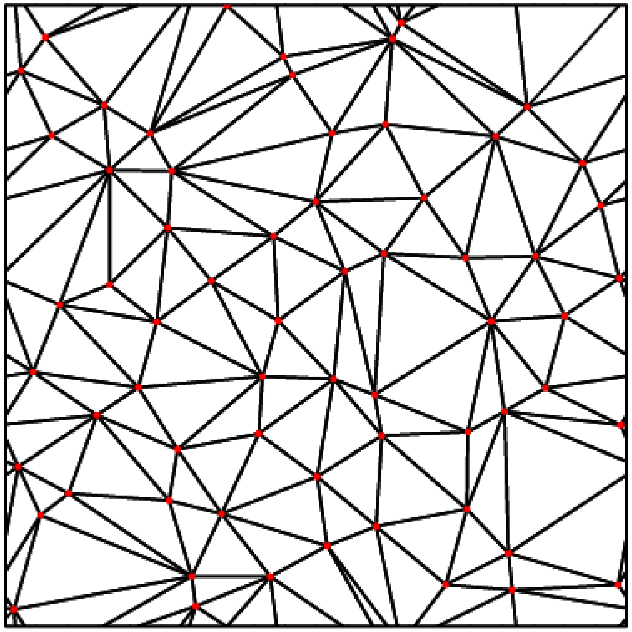

both shown together

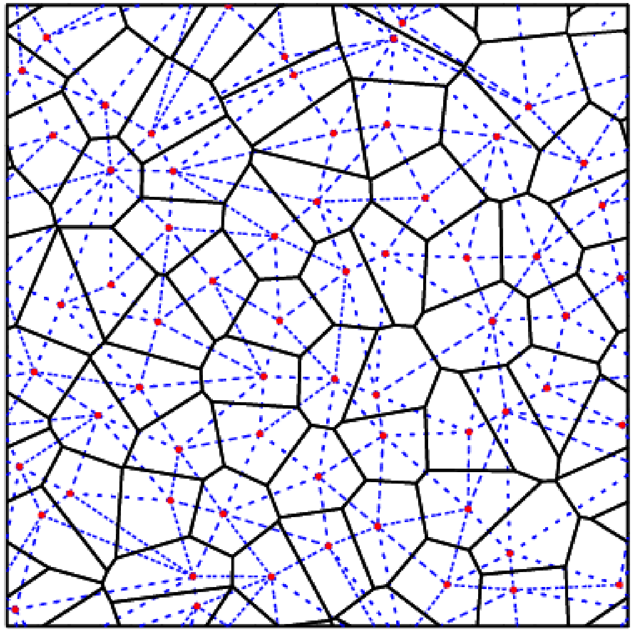

Each Voronoi cell contains the **space closest** to its generating point

The Delaunay triangulation contains only triangles with an **empty circumcircle**. The Delaunay tiangulation maximizes the minimum angle occurring among all triangles.

The centres of the circumcircles of the Delaunay triangles are the vertices of the Voronoi mesh. In fact, the two tessellations are the topological **dual graph** to each other.

The fluxes are calculated with an exact Riemann solver in the frame of the moving cell boundary

SKETCH OF THE FLUX CALCULATION

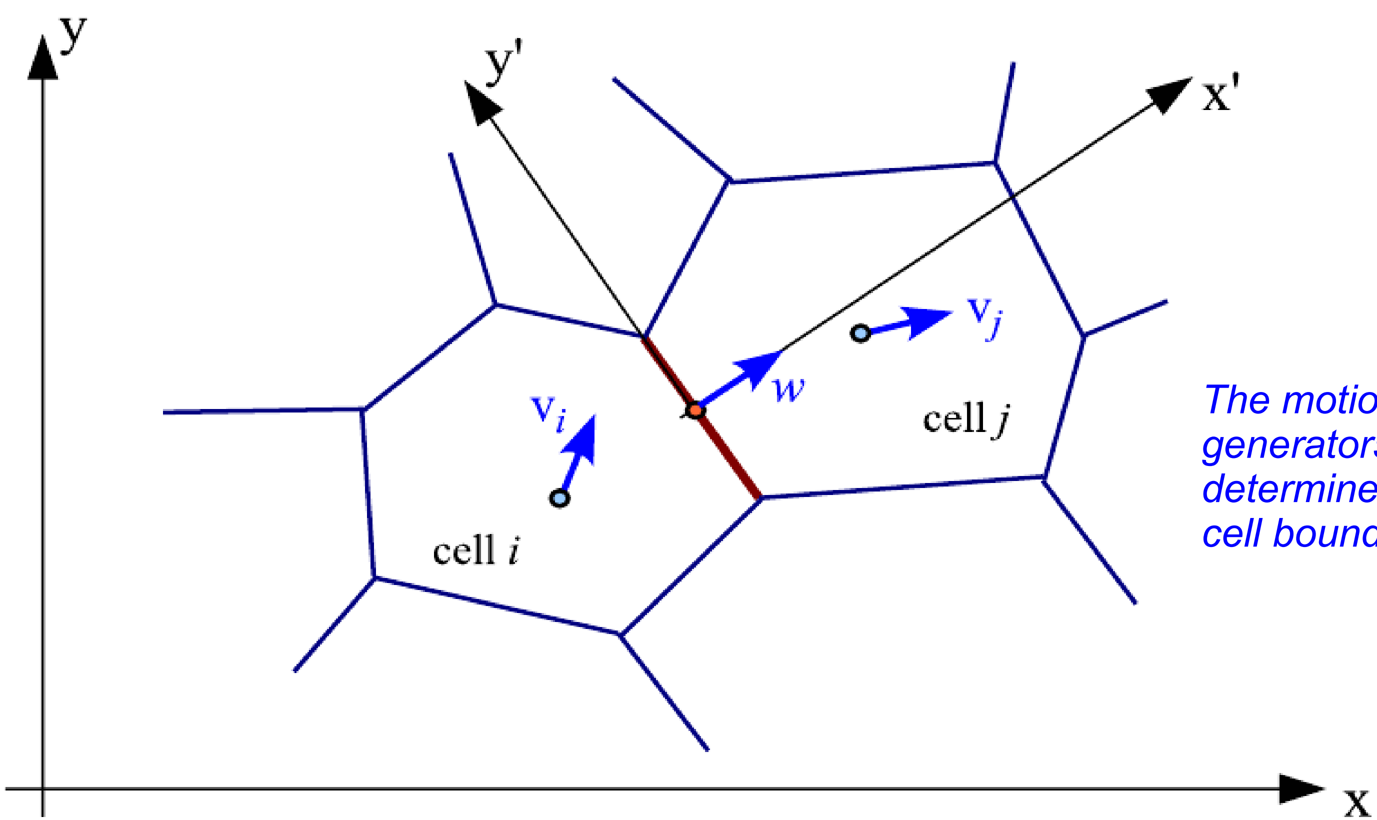

The motion of the mesh generators uniquely determines the motion of all cell boundaries

State left of cell face $\begin{pmatrix} \rho_L \\ \mathbf{v}_L \\ P_L \end{pmatrix}$

State right of cell face $\begin{pmatrix} \rho_R \\ \mathbf{v}_R \\ P_R \end{pmatrix}$

Riemann solver
(in frame of cell face)

$\rightarrow \begin{pmatrix} \rho \\ \mathbf{v} \\ P \end{pmatrix} \rightarrow \mathbf{F}(\mathbf{U})$

On large scales, the code produces very similar results as standard SPH techniques

GAS AND TEMPERATURE FIELDS IN A COSMOLOGICAL HYDRODYNAMIC SIMULATION

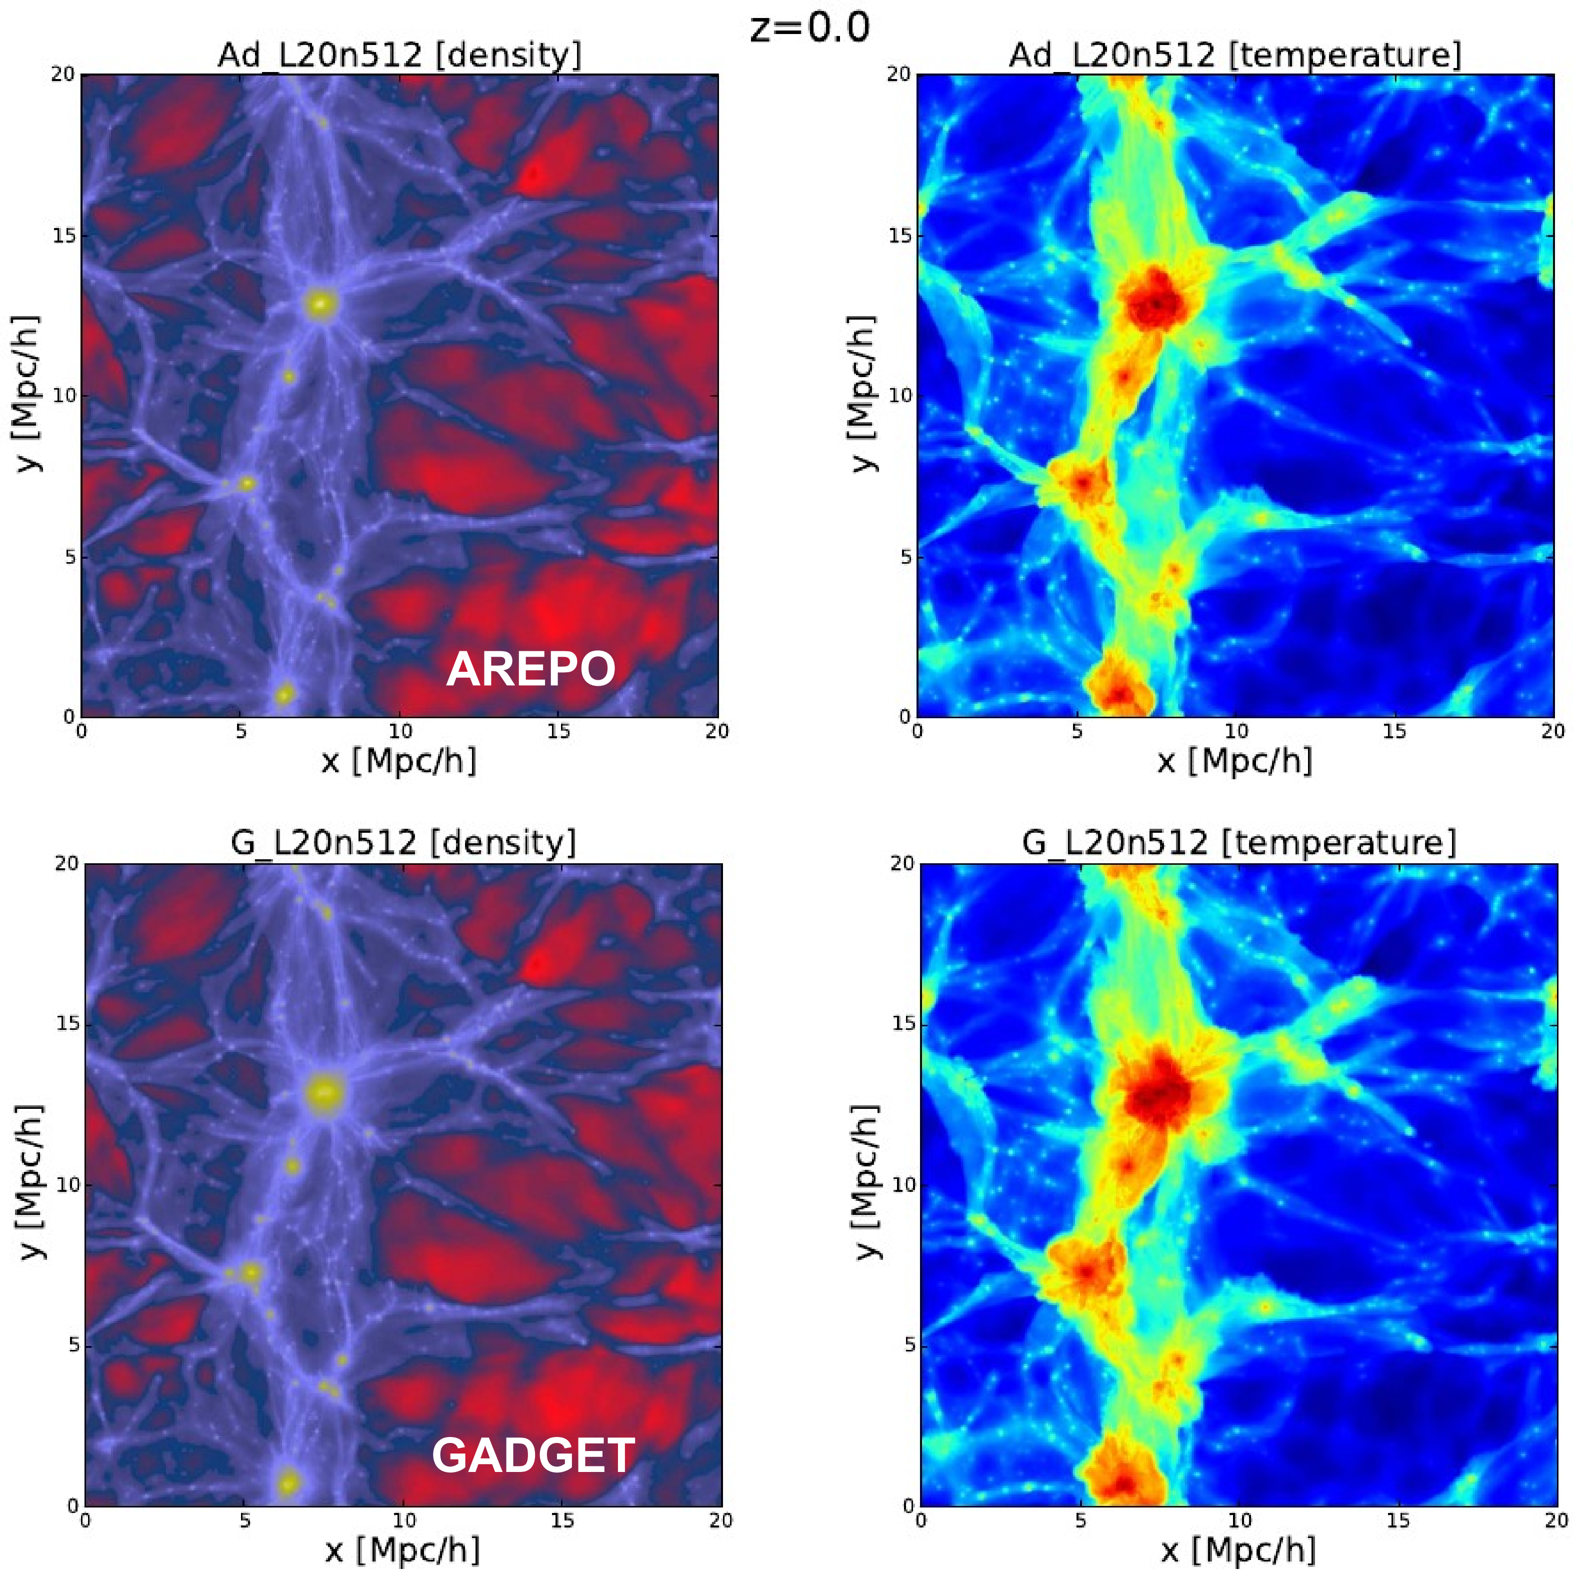

AREPO produces much better galaxy morphologies than SPH for identical initial conditions

GAS AND TEMPERATURE FIELDS IN A COSMOLOGICAL HYDRODYNAMIC SIMULATION

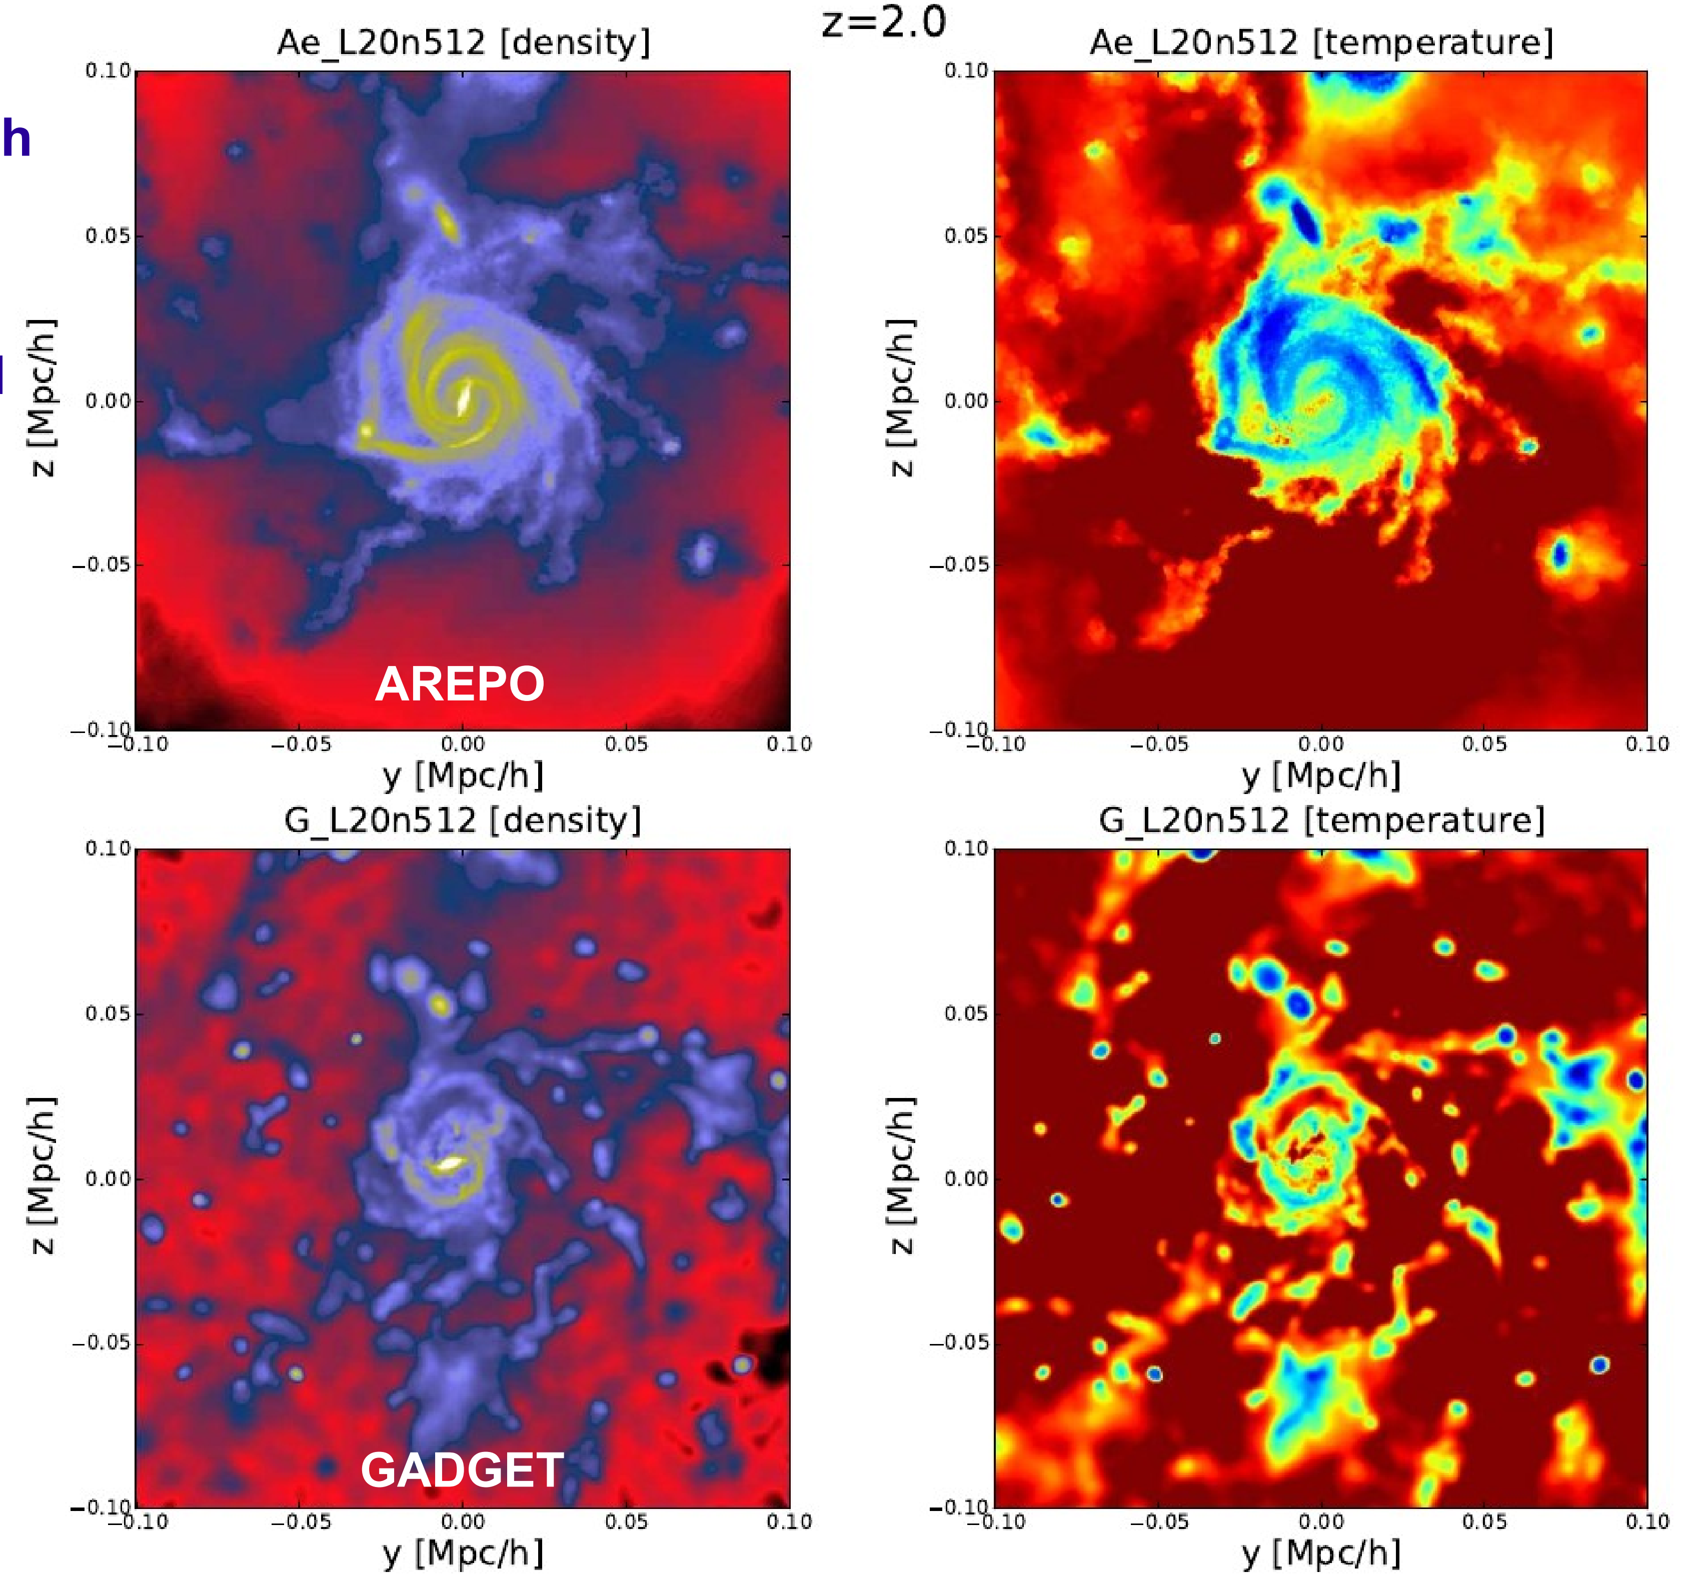

AREPO produces much better galaxy morphologies than SPH for identical initial conditions

GAS AND TEMPERATURE FIELDS IN A COSMOLOGICAL HYDRODYNAMIC SIMULATION

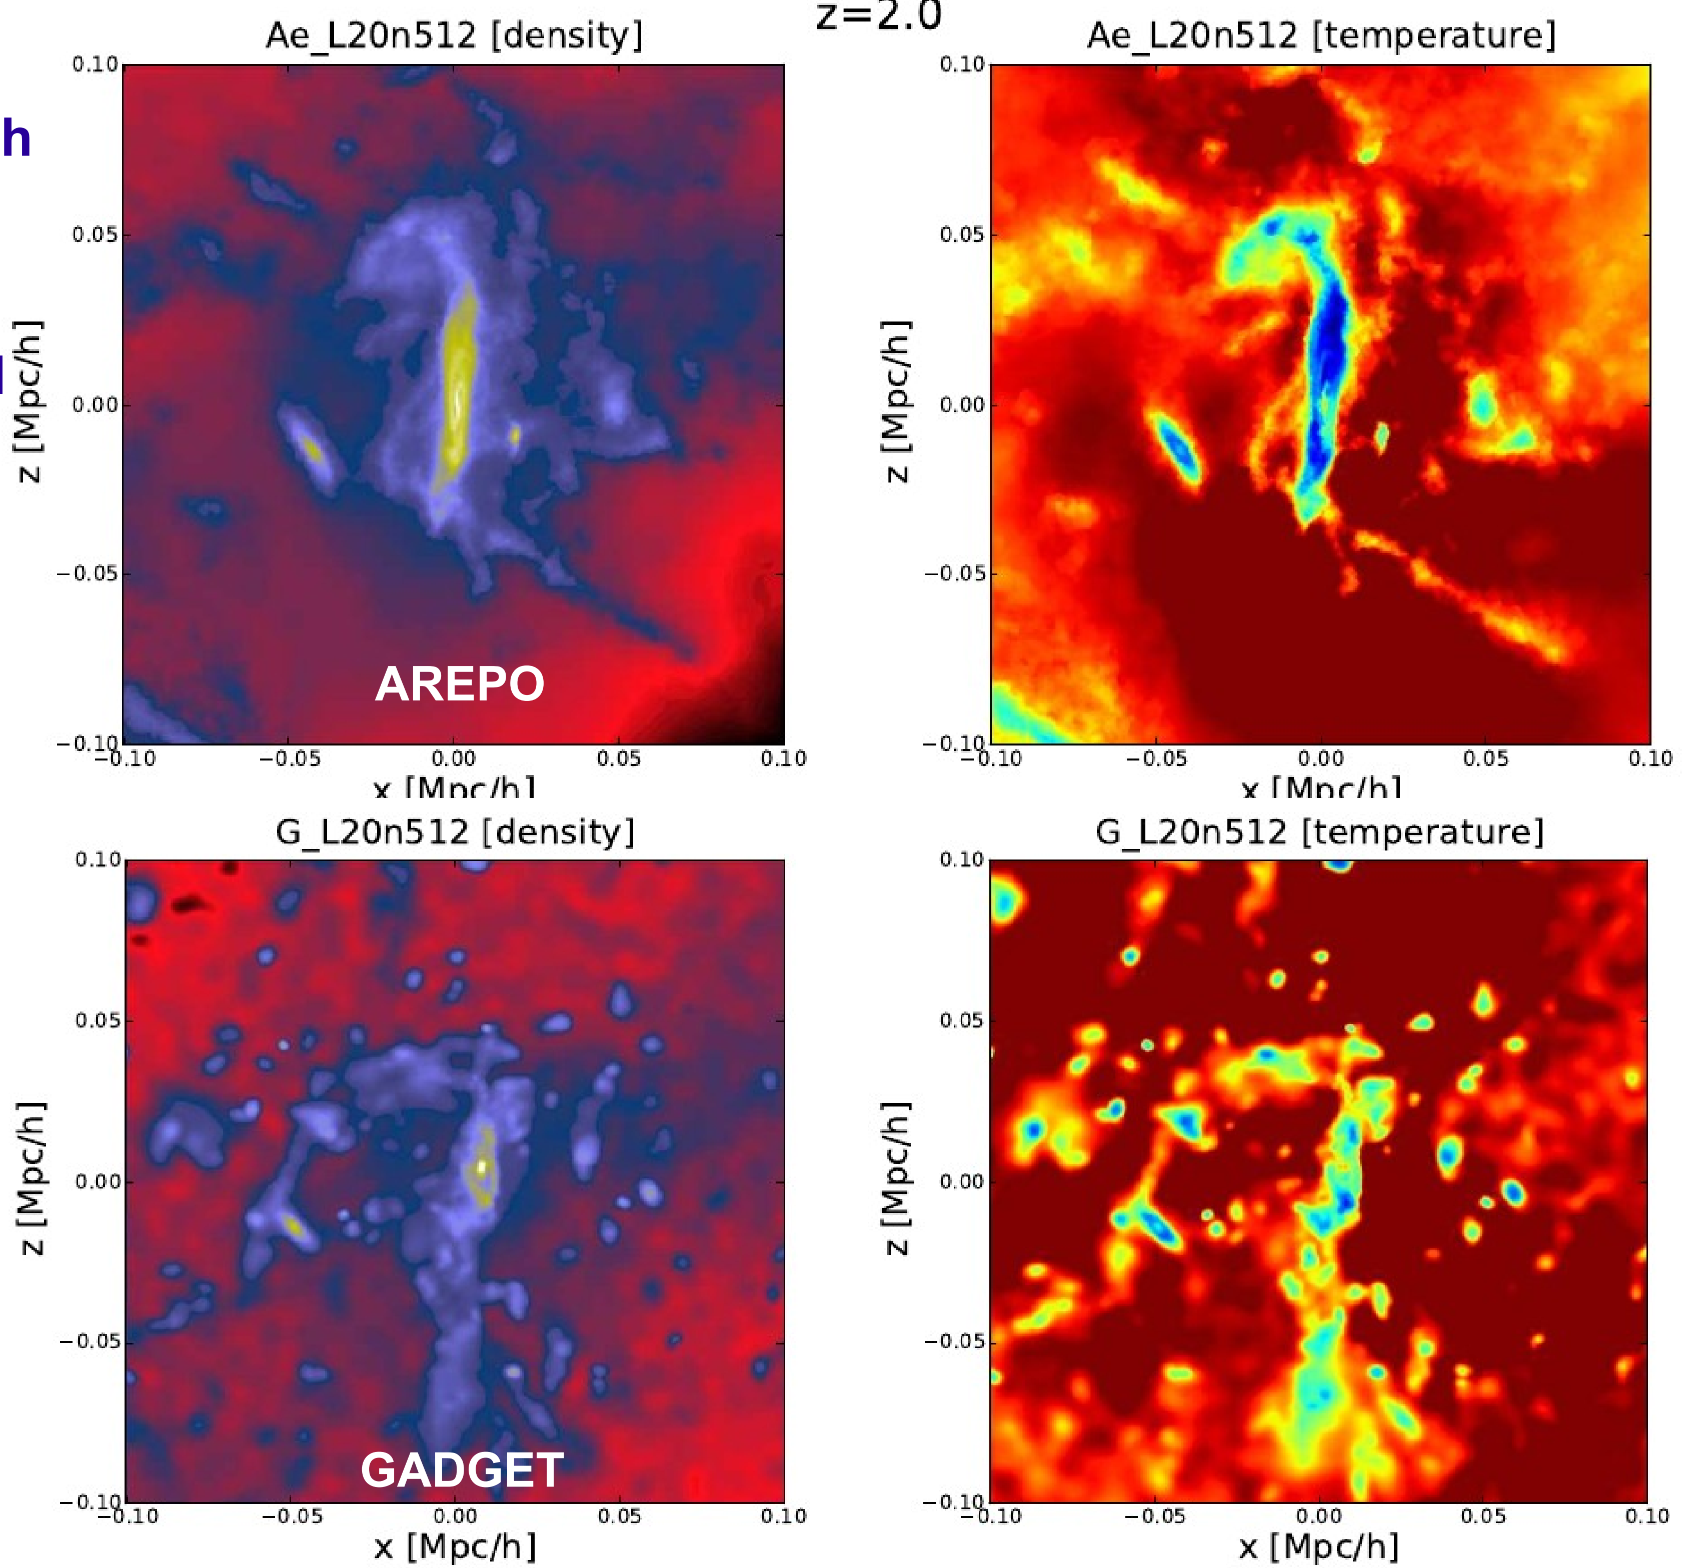

The moving-mesh approach can also be used to realize arbitrarily shaped, moving boundaries

STIRRING A COFFEE MUG

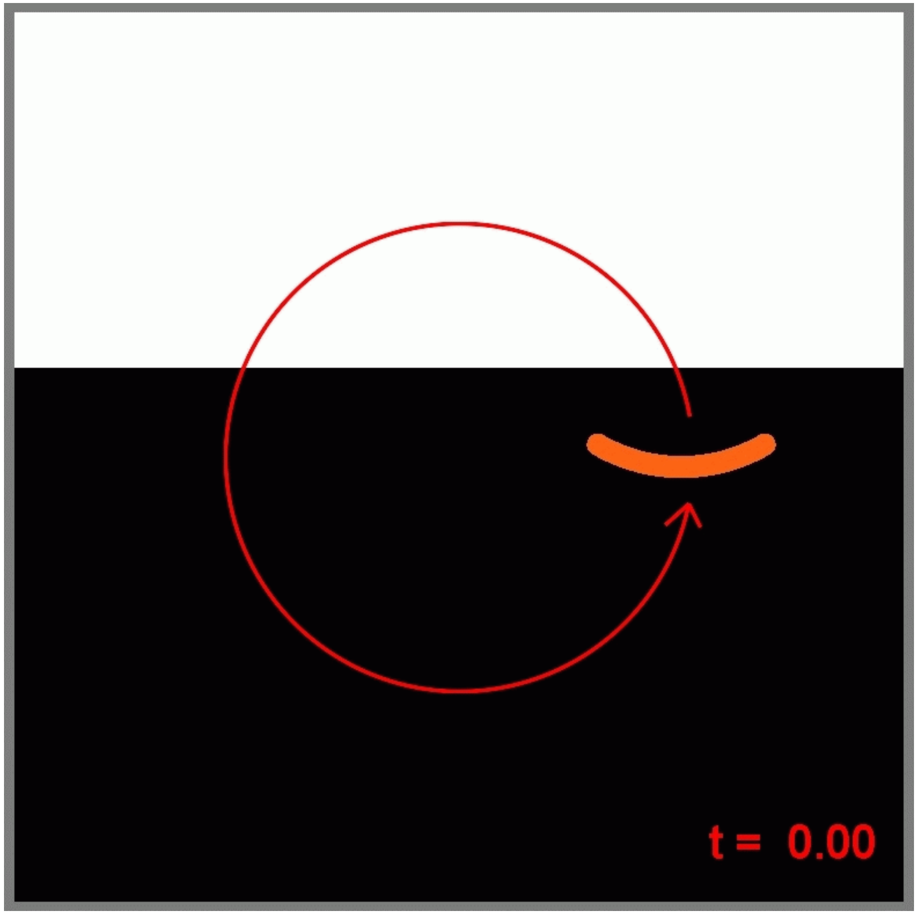

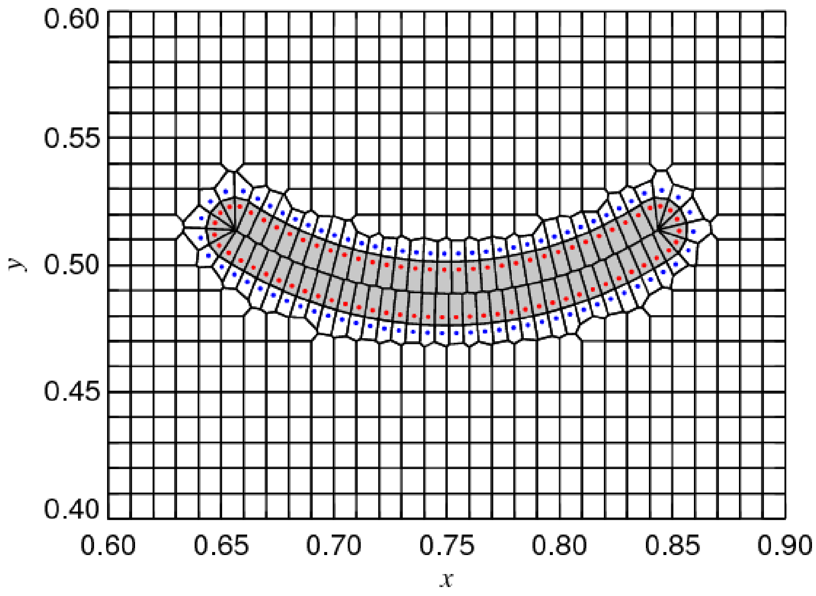

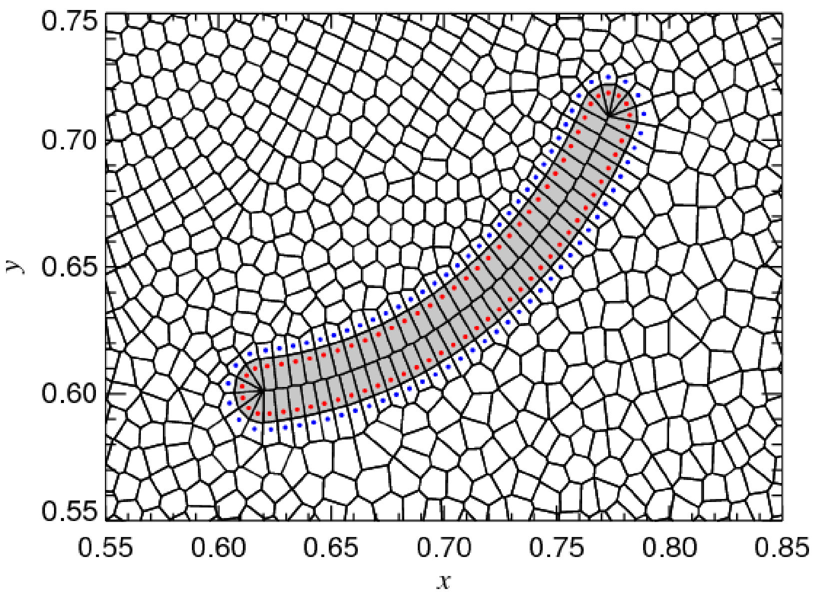

The challenge to simulate galaxy formation

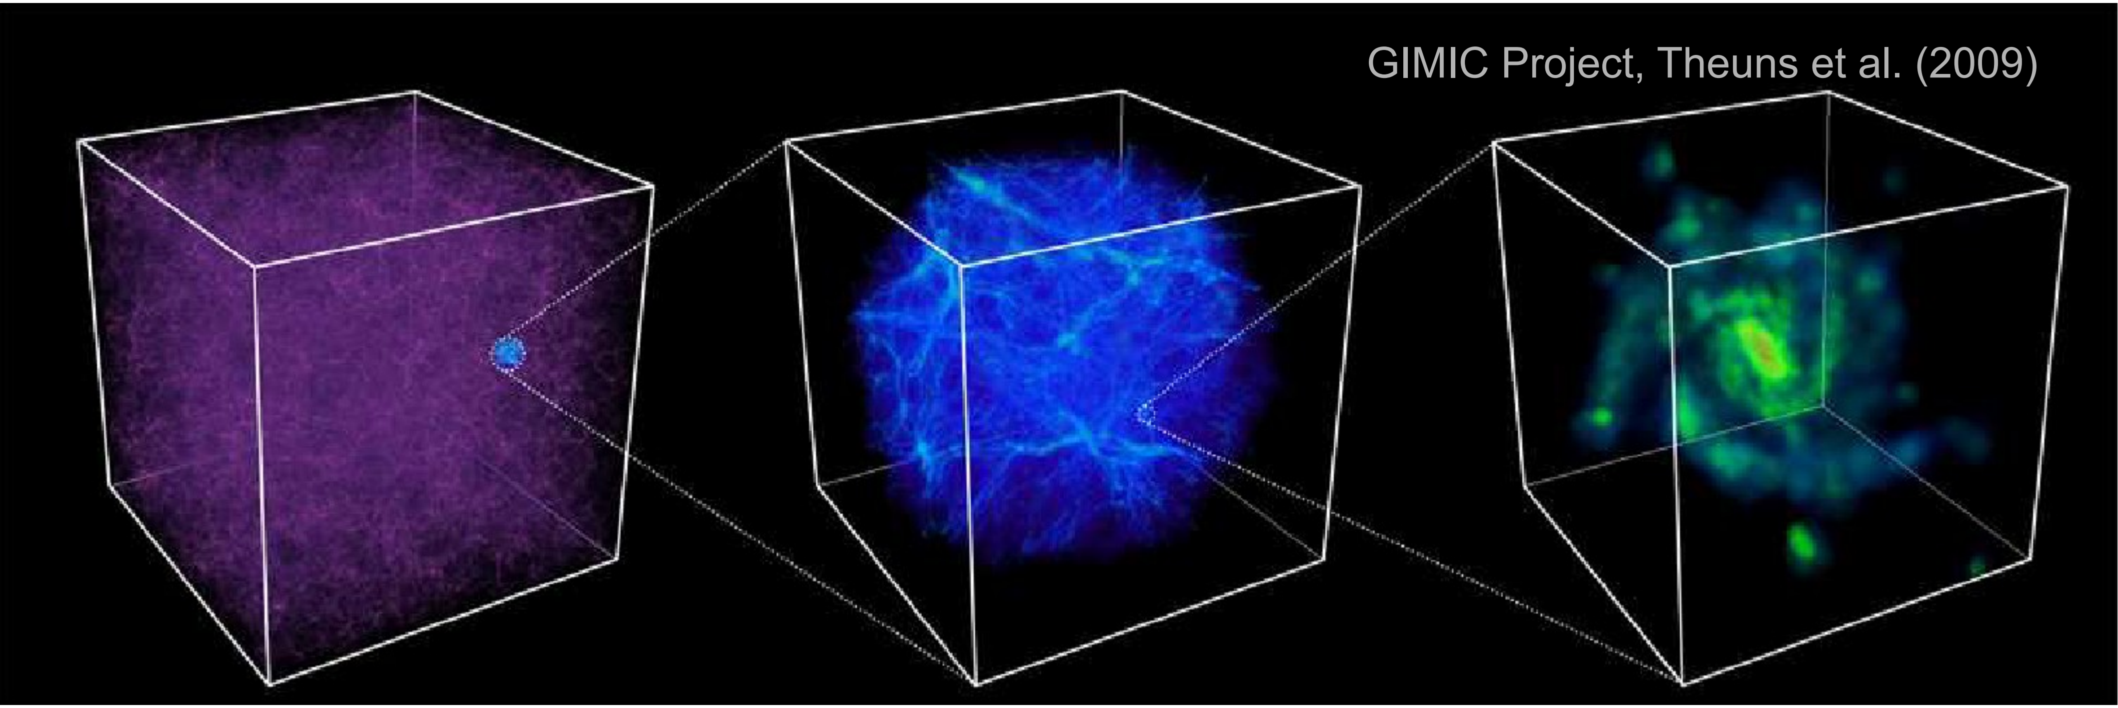

Hydrodynamical simulations aim to predict:

- Morphology of galaxies
- Fate of the diffuse gas, WHIM, metal enrichment
- X-ray atmospheres in halos
- Turbulence in halos and accretion shocks
- Large-scale regulation of star formation in galaxies through feedback processes from stars and black holes
- Transport processes (e.g. conduction)
- Radiative transfer
- Dynamical transformations (e.g. ram-pressure stripping)
- Magnetic fields

A long standing issue in galaxy formation theory: The shapes of the CDM halo mass function and the galaxy luminosity function are very different

THE OBSERVED LF COMPARED TO THE SHAPE OF THE CDM HALO MASS FUNCTION

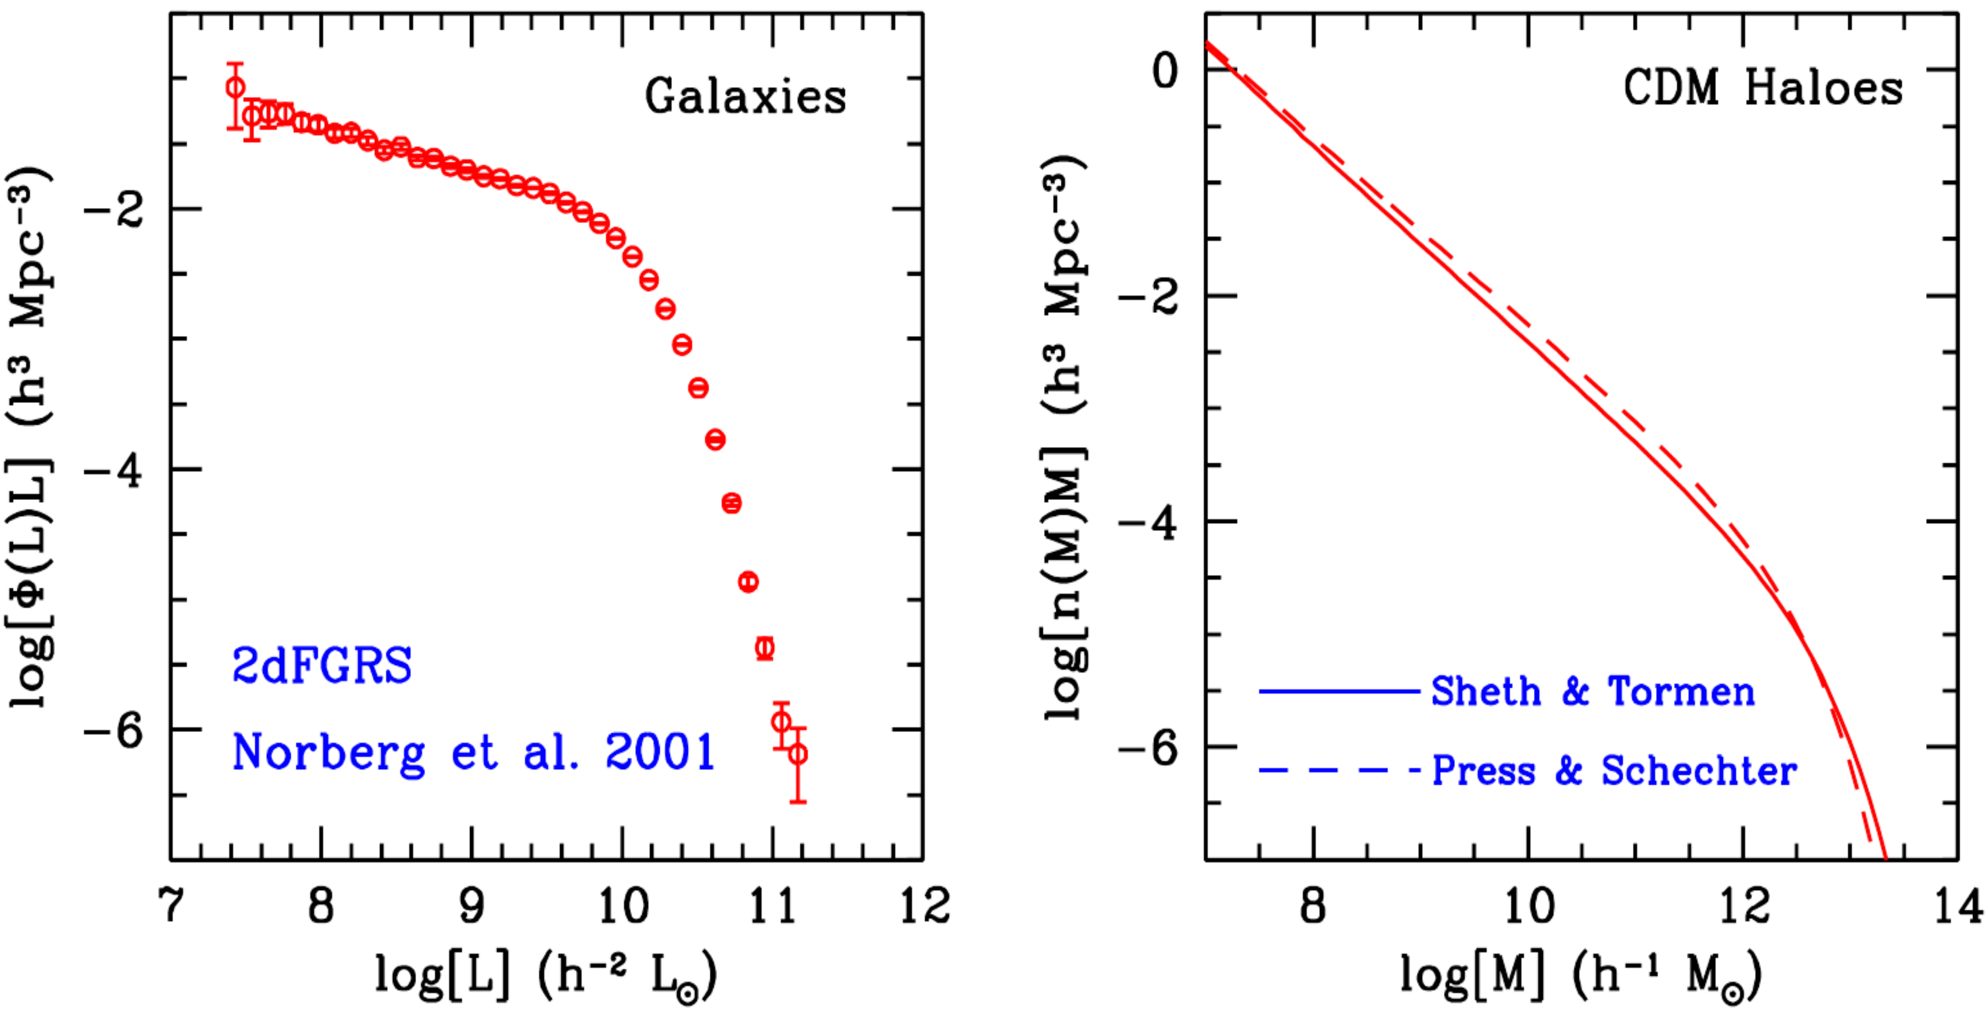

van den Bosch et al. (2004)

Abundance matching gives the expected halo mass – stellar mass relation in ΛCDM

STELLAR MASSES FROM SDSS/DR7 MATCHED TO ΛCDM SIMULATION EXPECTATIONS

Assumption:
Stellar mass is monotonically increasing with halo mass

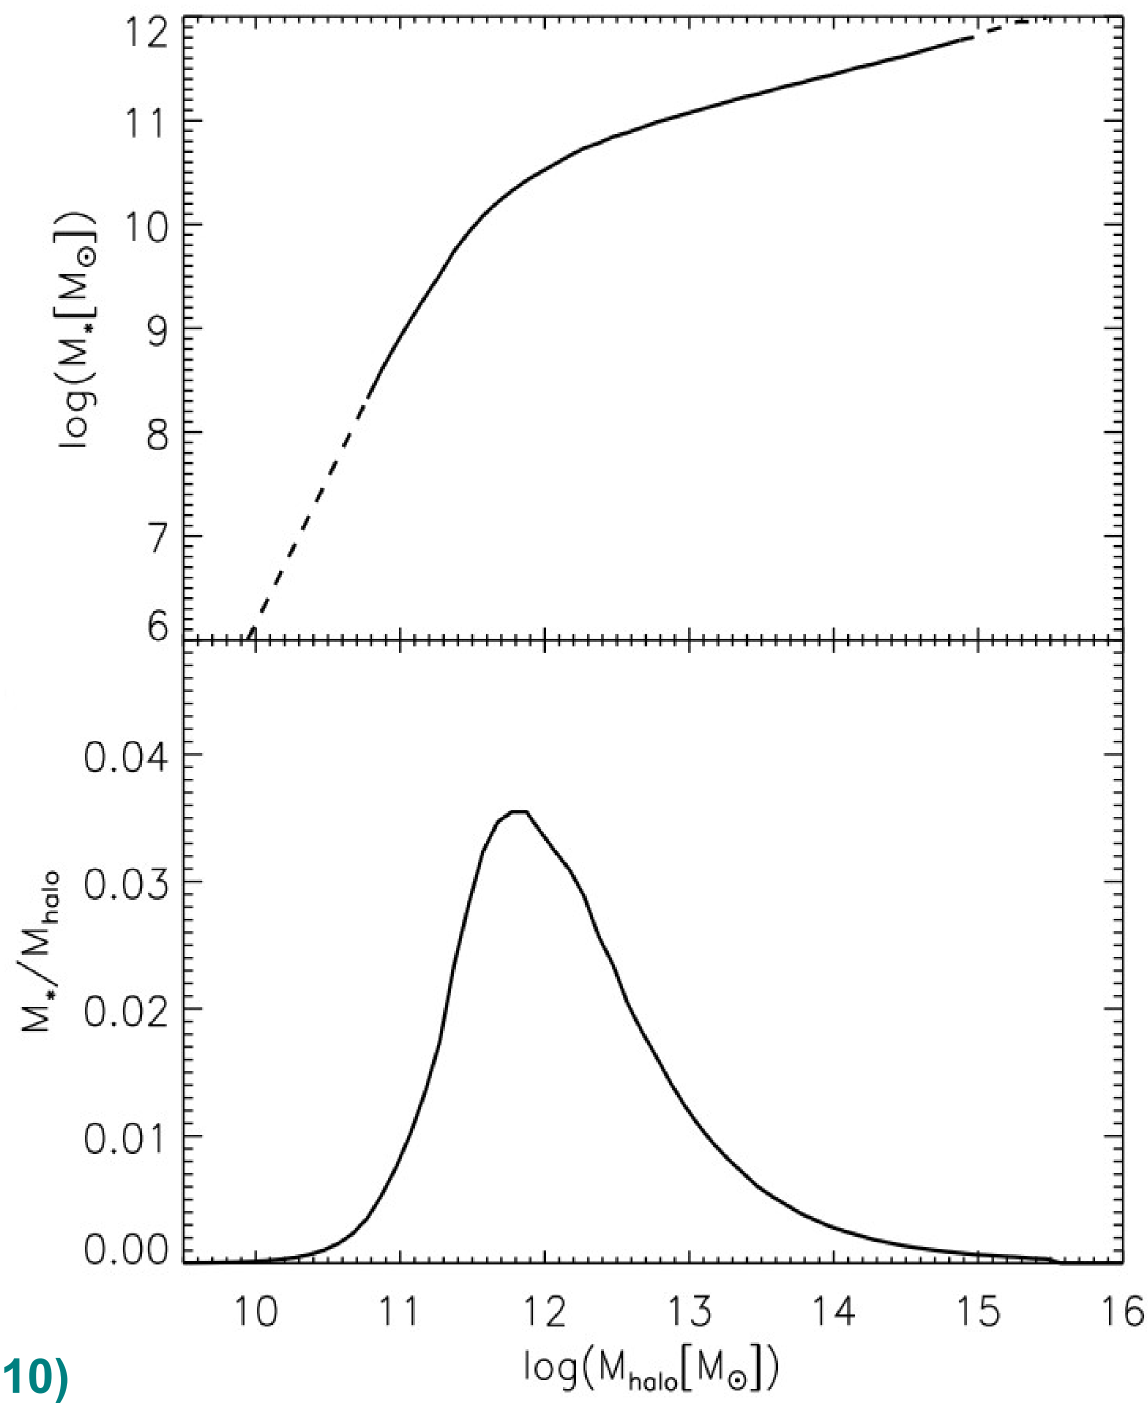

Guo, White & Boylan-Kolchin (2010)

Current cosmological hydrodynamic simulations have trouble to explain such a low galaxy formation efficiency

GALAXY FORMATION EFFICIENCY AS A FUNCTION OF HALO MASS

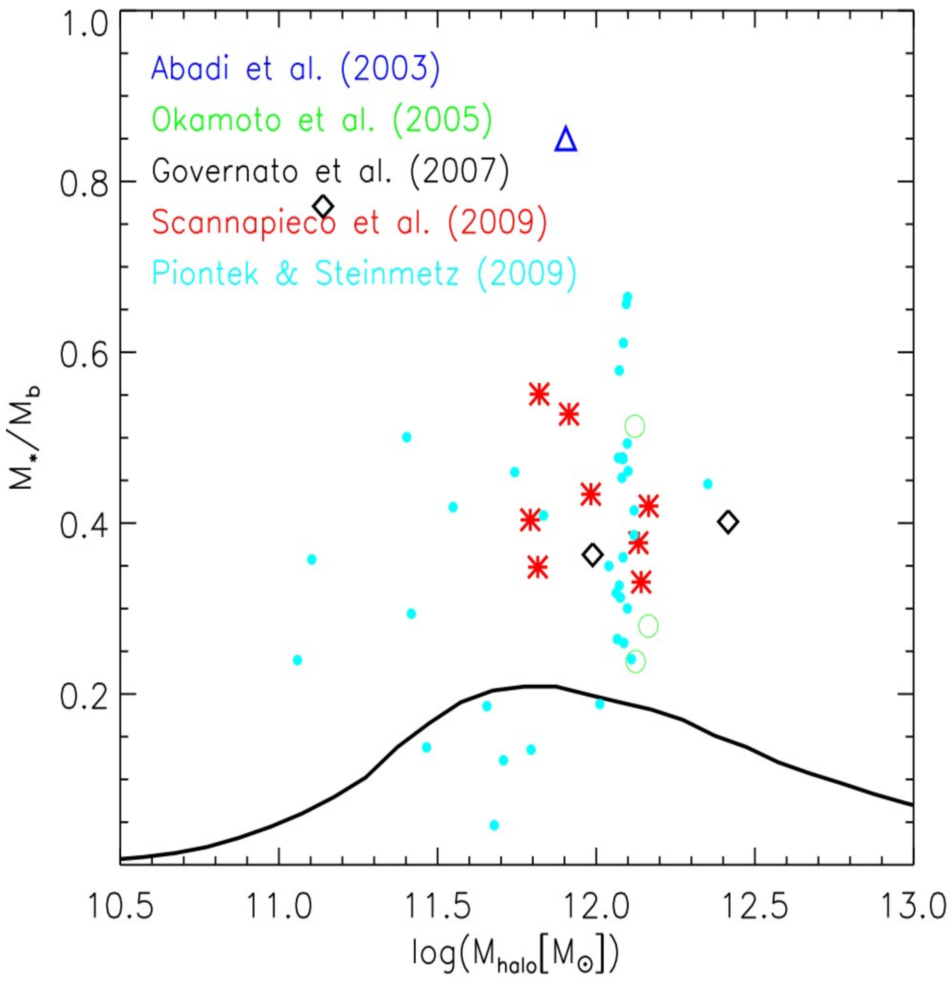

Guo, White & Boylan-Kolchin (2010)

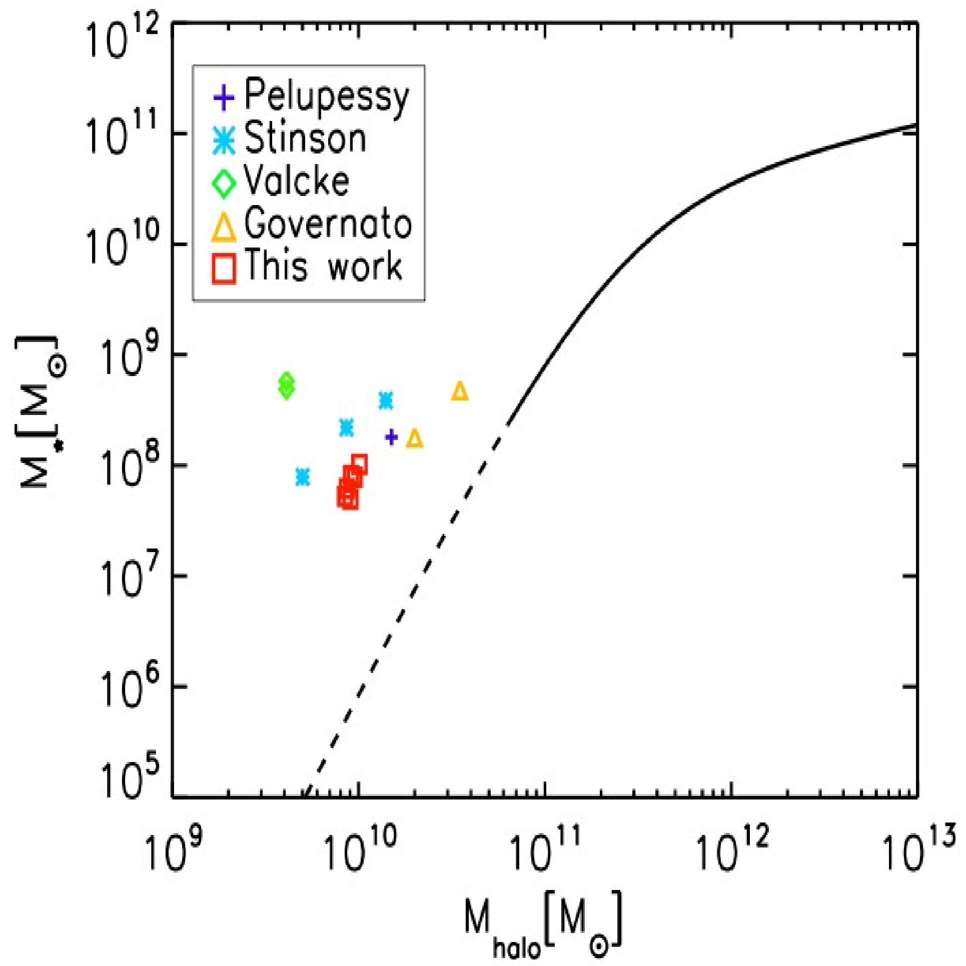

Sawala & White (2010)

Summary points

- Direct numerical simulations have become indispensable for studying the non-linear growth of structures in ΛCDM cosmologies.

- Current numerical techniques allow high-resolution simulations with an unprecedented dynamic range.

  One presently reaches $N>10^{11}$, with a dynamic range of $10^5 - 10^7$ in 3D.

- The future observation of a sufficiently massive cluster may easily rule out the ΛCDM model. The predicted satellite population may still be in tension with the observations.

- Understanding galaxy formation physics remains a serious challenge in ΛCDM, both at the faint and the bright end.

- Radiative magneto-hydrodynamics codes that follow structure formation still in their infancy.

- Exaflop computers arrive at the end of the decade, but it is highly questionable whether astrophysics can use them at scale.

Discussion points

Future codes cannot be written by lonely graduate students any more...

They require large, interdisciplinary teams with sustained funding.

Biggest challenges in arriving at better codes for galaxy formation:

- **Cope with huge dynamic range in time and space.**
- **Avoid work-load imbalance losses when on 10^3 to 10^7 cores.**
- **Code validation.**
- **Simulation data management.**
- **Dragging bright physics students into computational cosmology.**