One decade of GPUs for cosmological simulations (in Strasbourg) : fortunes & misfortunes

---

Dominique Aubert

with P. Ocvirk, J. Chardin, J. Lewis, N. Deparis (Strasbourg, F) & N. Gillet (SNS Pisa, It)
CODA Collaboration - CLUES Collaboration - Kulkarni G., Haenhelt, M. (Cambridge, UK)

Université
de Strasbourg

The epoch of Reionization

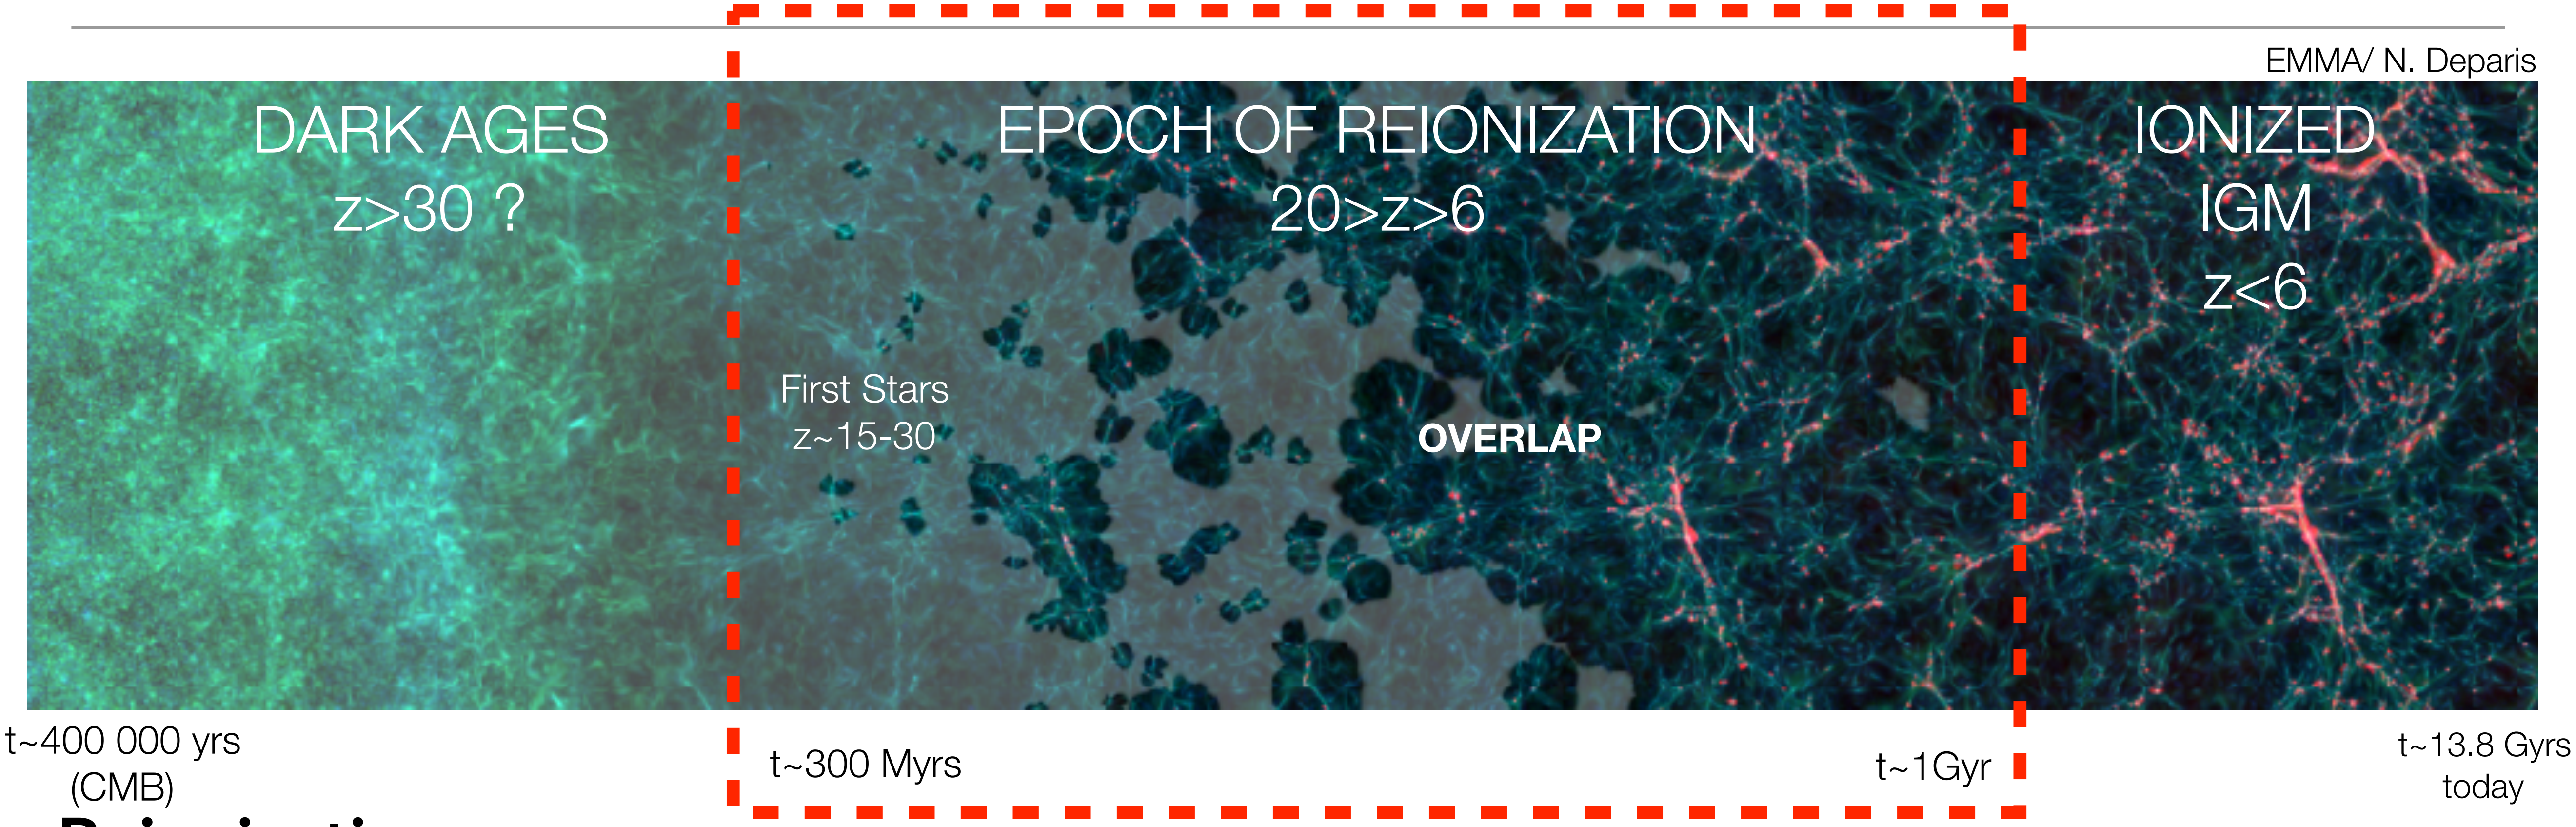

EMMA/ N. Deparis

Reionization

= end by **z ~6**// t~1 Gyr

= a great and rapid cosmological transition (few **100s Myrs**)

=Driven by IONIZING LIGHT and the first sources (and **dominated by stars**)

= network of HII regions (**ionized bubbles**)

= ionization **(0.9999**) & heating (~**10 000 K**) of IGM (by stars and quasars)

= **UV Background**

= the initial stages of galaxy formation

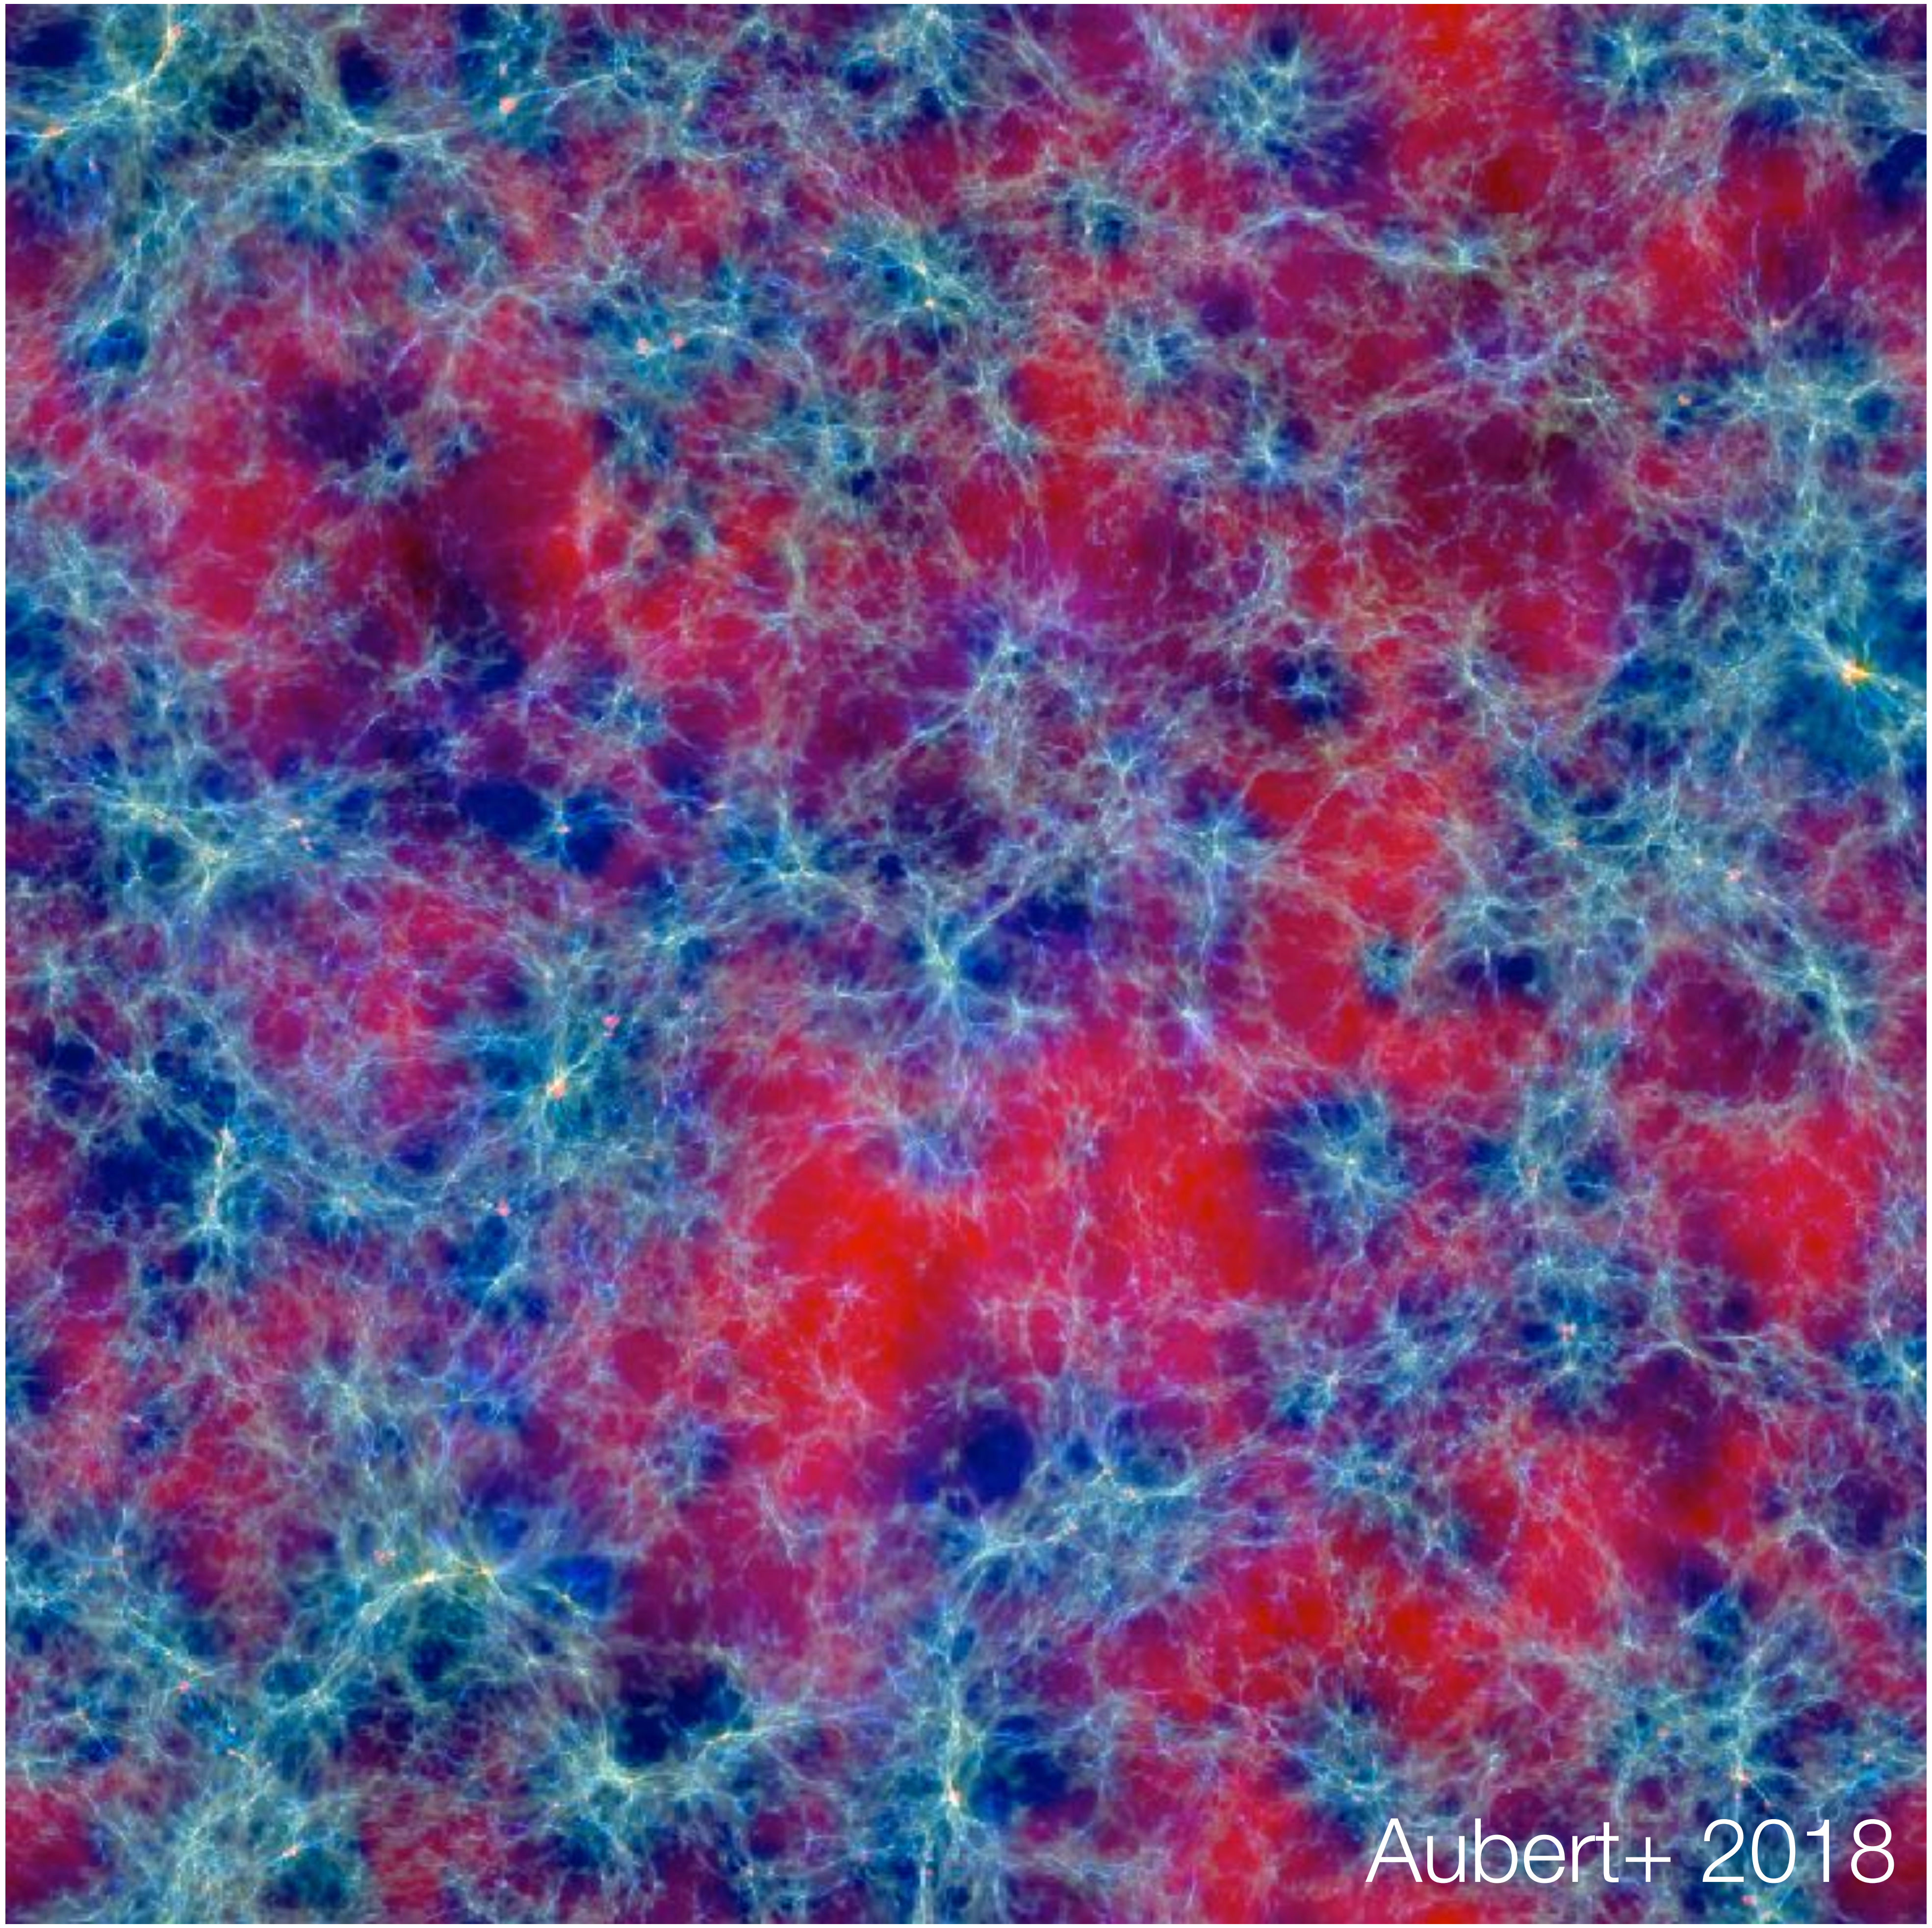

CODA I-AMR
(Aubert+ 18)
91 Mpc/2048^3
16 billions resolution elements with AMR
@ z=6

32768 cores+**4096 GPUs** on Titan(DOE/ORNL) using EMMA simulation code (Aubert+ 15)

20+ millions cpu hours
Jan.-Mar 2017

INCITE CODA
(PI : Shapiro)

WMAP5-CLUES ICs
spatial res.=500 pc
mass res~2e6 Msol
stellar res ~70 000 Msol

A few words about GPUs

- **G**raphics **P**rocessing **U**nits
- Large number of Multiprocessors (~100) + Scheduler

- **Fit to Large number of Independent, intensive and memory friendly tasks**

- x10 to x100 compared to CPU
- High-level interface with **CUDA (Nvidia)**, OpenCL (Kronos)
- High-end GPUs: up to 16 GB VRAM
- Volta GPUs (e.g. Summit/ORNL) :
  - unified CPU/GPU memory,
  - GPU/GPU comm,
  - GPU shared among CPUs
  - hardwired tensor operations

- Explicit conservative solver of **Radiative Transfer** equations + local thermo-chemistry: **ensures high load and local computations for cell update**
- also true for grid-based Hydrodynamics with non trivial Riemann solvers, reconstructions, etc..

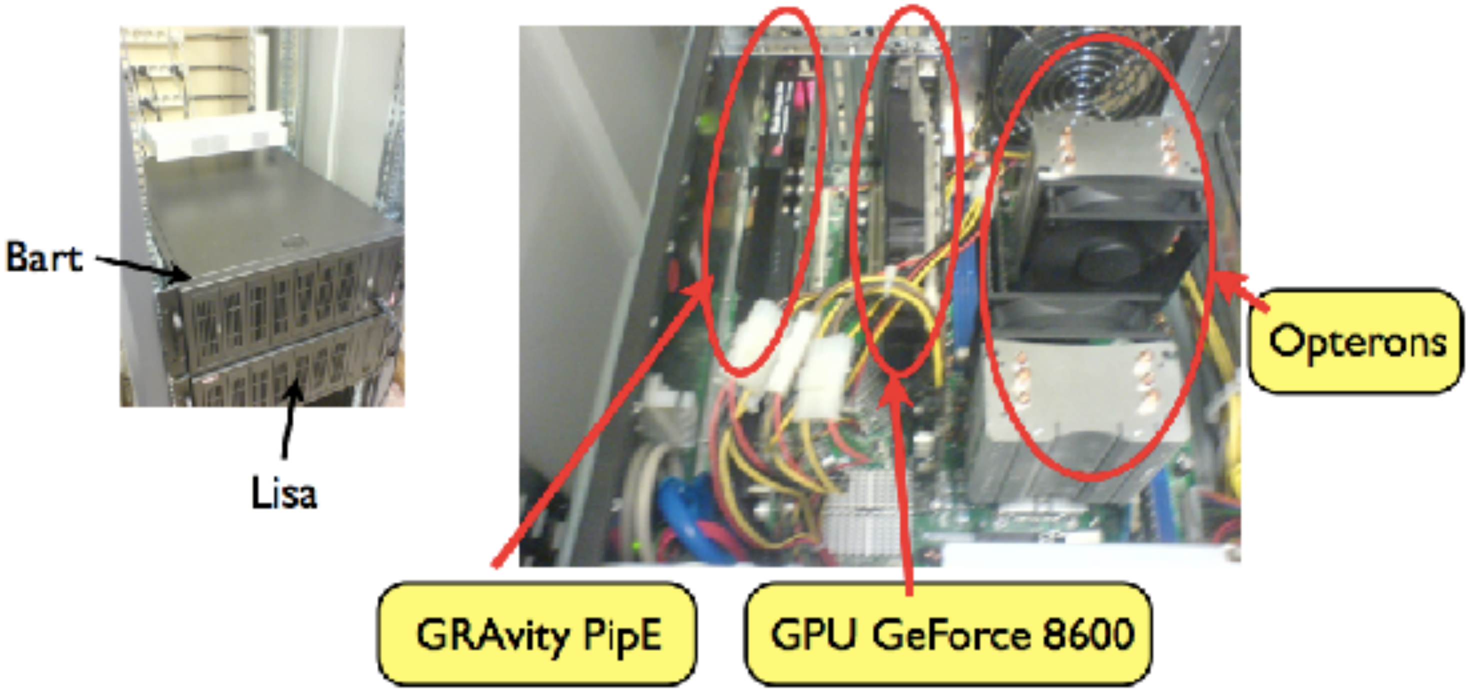

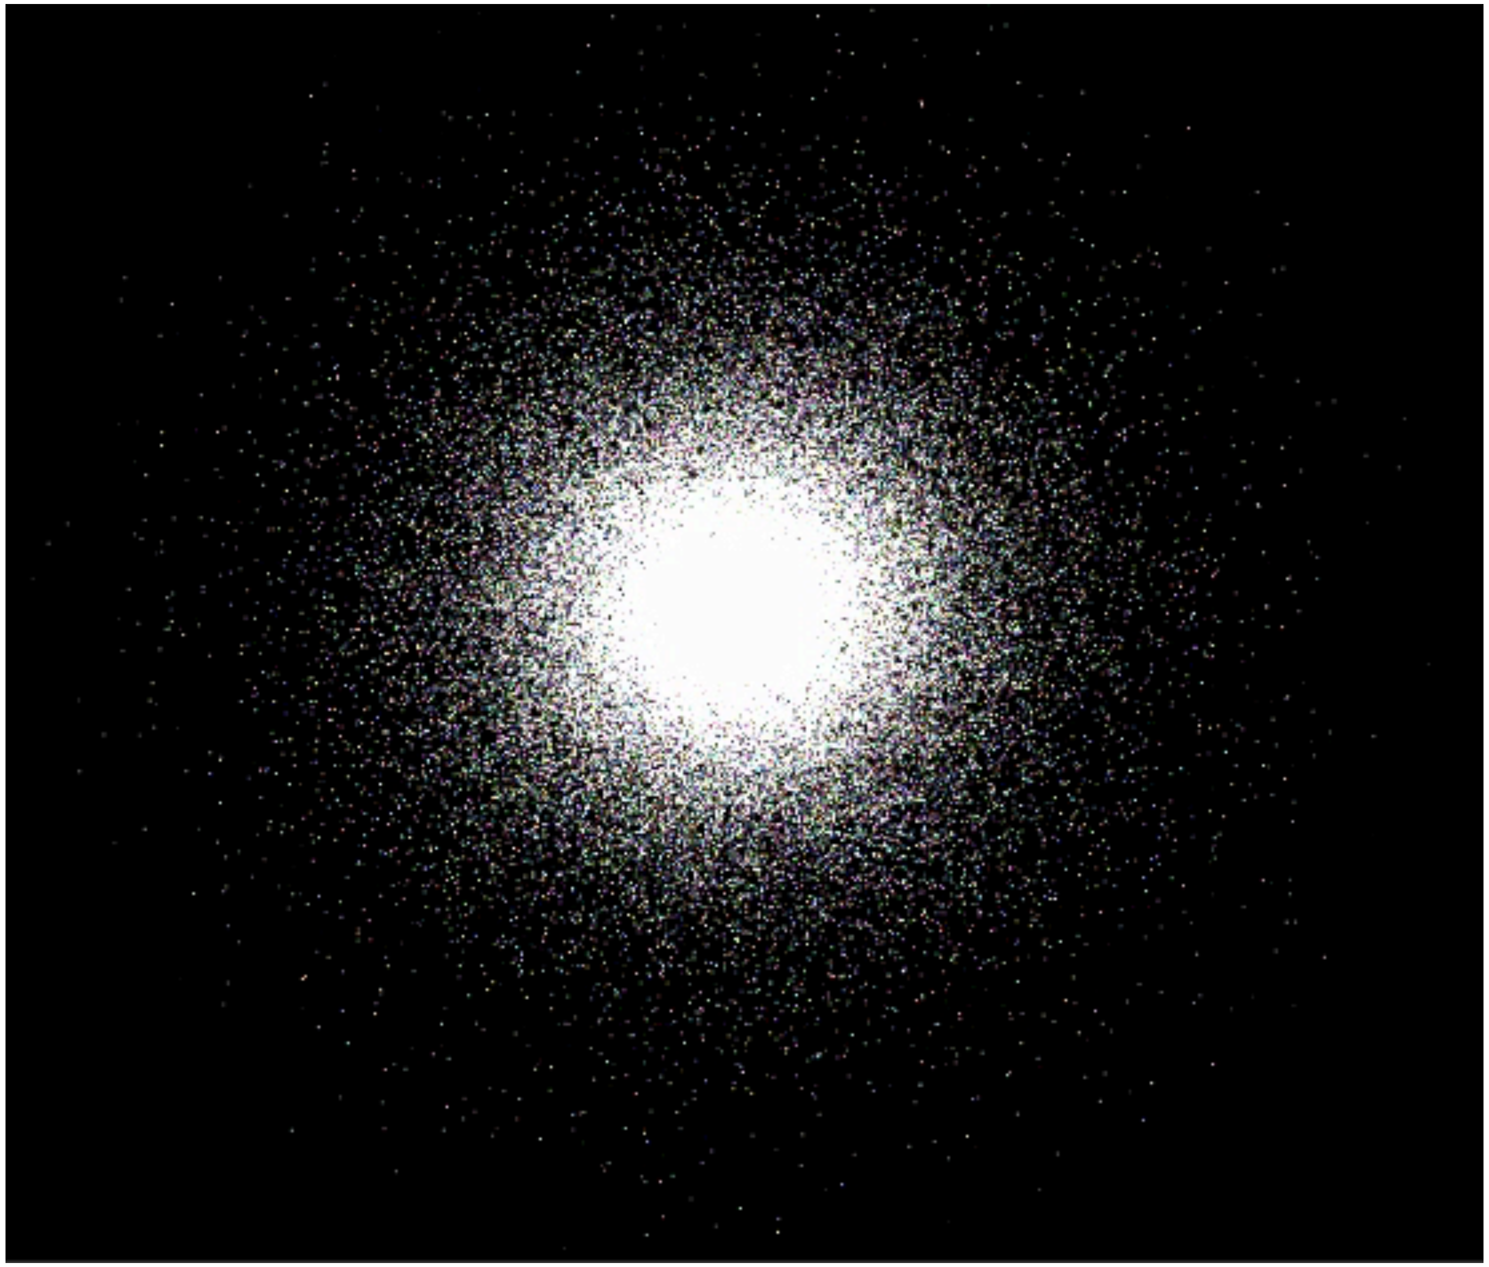

2008 Direct N-Body Calculation (With Christian Boily & Romaric David, Strasbourg)

100k bodies

similar throughput as
GRAPE boards
(with hardwired $1/r^2$ calculators)
with
GPUs GeForce 8800 GTX (with a cost 10x smaller !)

Gabrielle Allen Jaroslaw Nabrzyski
Edward Seidel Geert Dick van Albada
Jack Dongarra Peter M.A. Sloot (Eds.)

LNCS 5544

Computational Science – ICCS 2009

9th International Conference
Baton Rouge, LA, USA, May 2009
Proceedings, Part I

Part I

Springer

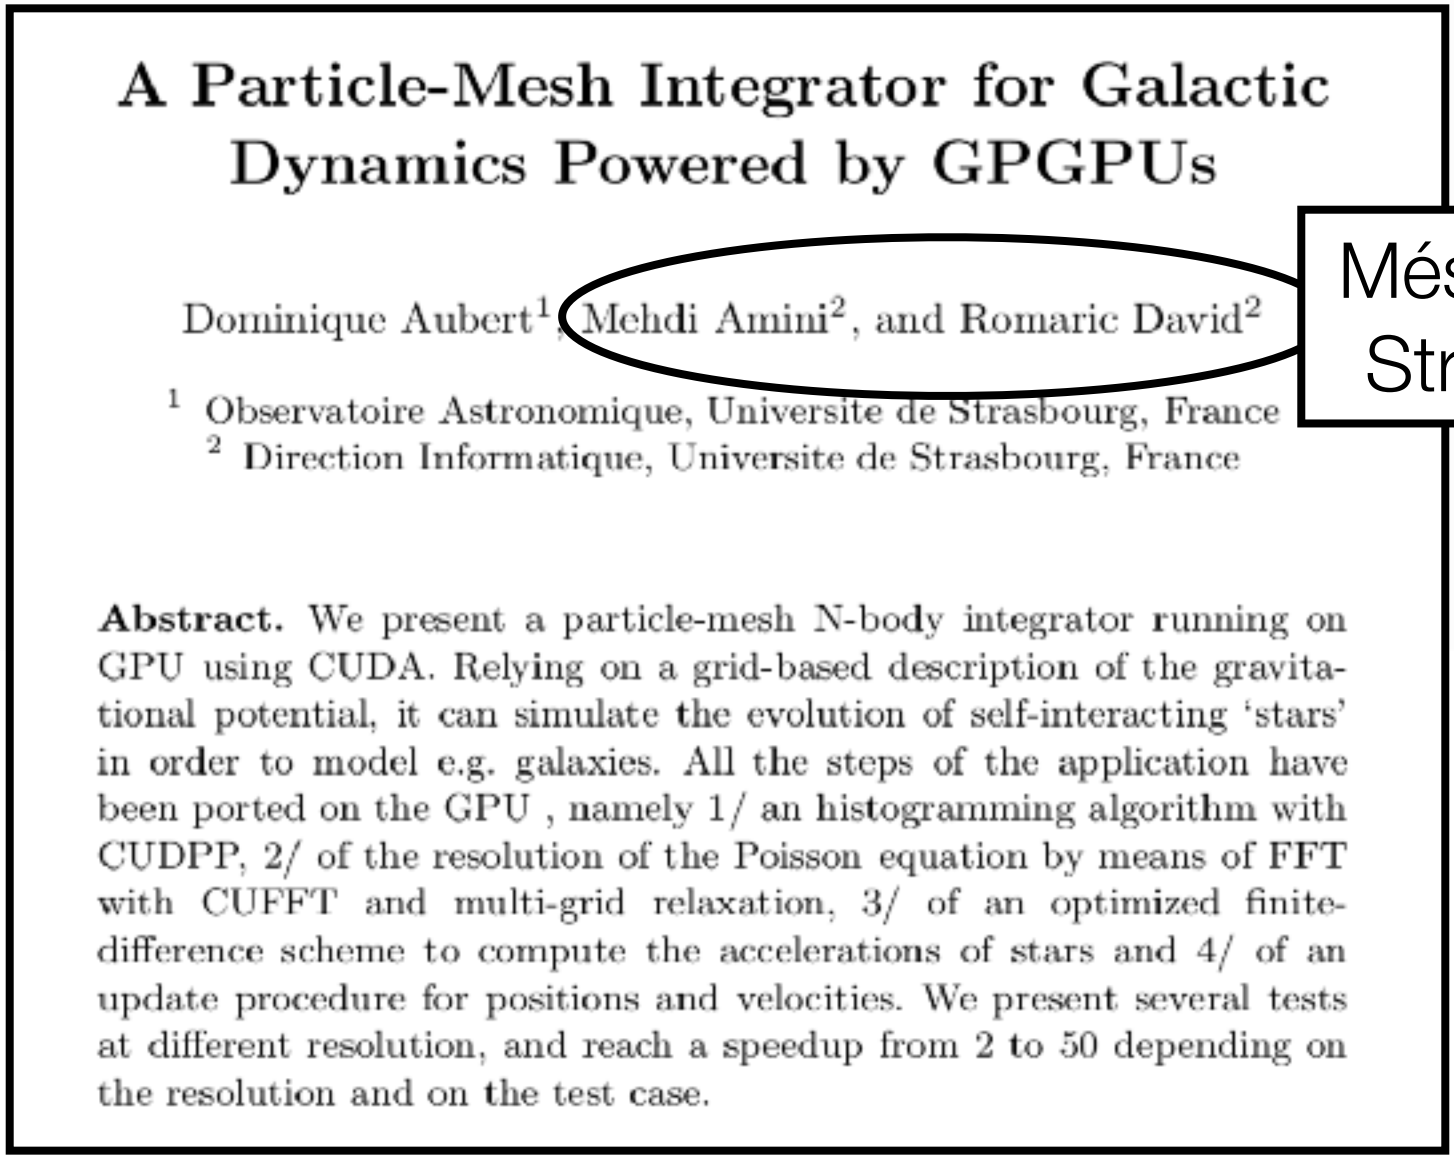

A Particle-Mesh Integrator for Galactic Dynamics Powered by GPGPUs

Dominique Aubert[1], Mehdi Amini[2], and Romaric David[2]

[1] Observatoire Astronomique, Universite de Strasbourg, France
[2] Direction Informatique, Universite de Strasbourg, France

Abstract. We present a particle-mesh N-body integrator running on GPU using CUDA. Relying on a grid-based description of the gravitational potential, it can simulate the evolution of self-interacting 'stars' in order to model e.g. galaxies. All the steps of the application have been ported on the GPU , namely 1/ an histogramming algorithm with CUDPP, 2/ of the resolution of the Poisson equation by means of FFT with CUFFT and multi-grid relaxation, 3/ of an optimized finite-difference scheme to compute the accelerations of stars and 4/ of an update procedure for positions and velocities. We present several tests at different resolution, and reach a speedup from 2 to 50 depending on the resolution and on the test case.

Méso-centre
Strasbourg

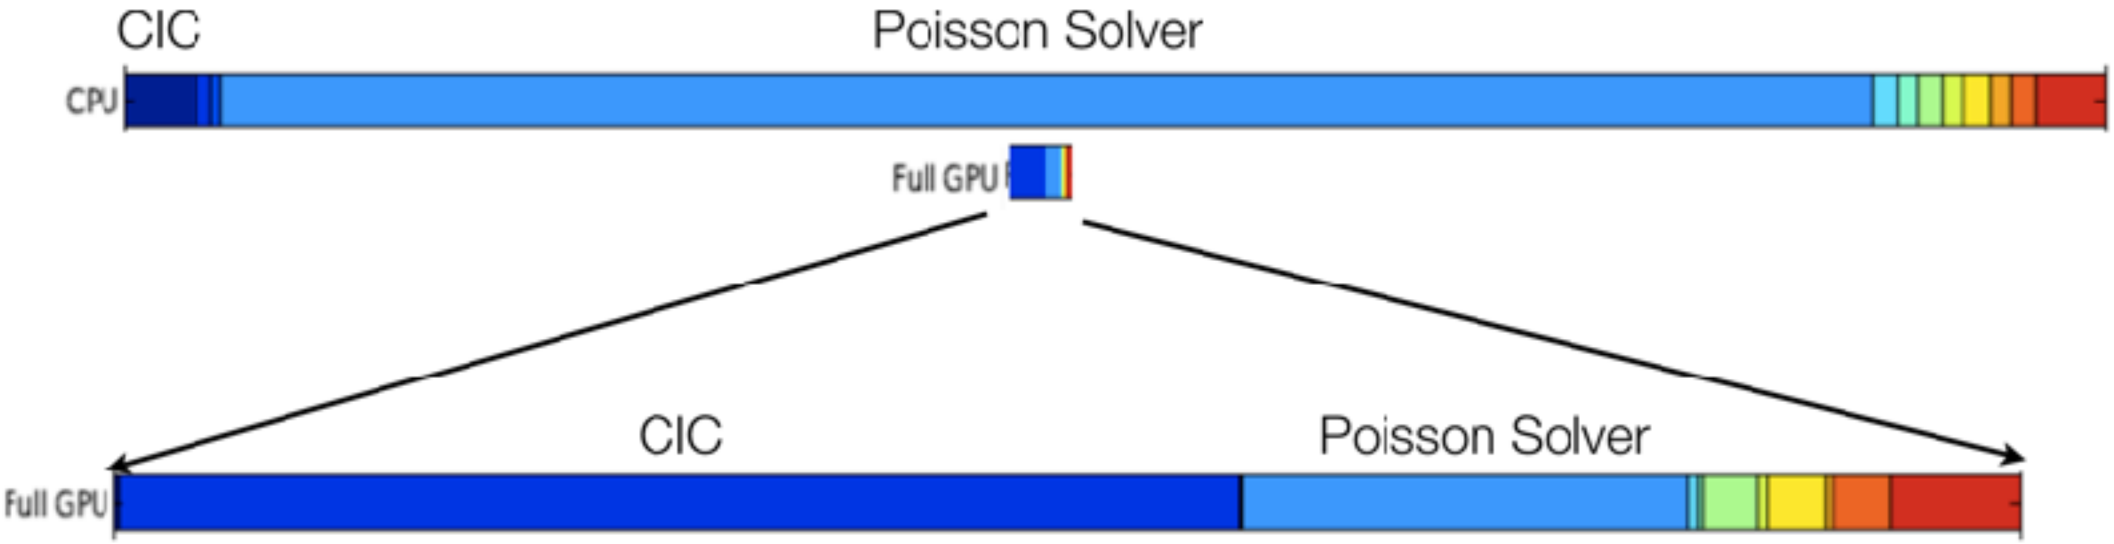

Multi-GPU PM

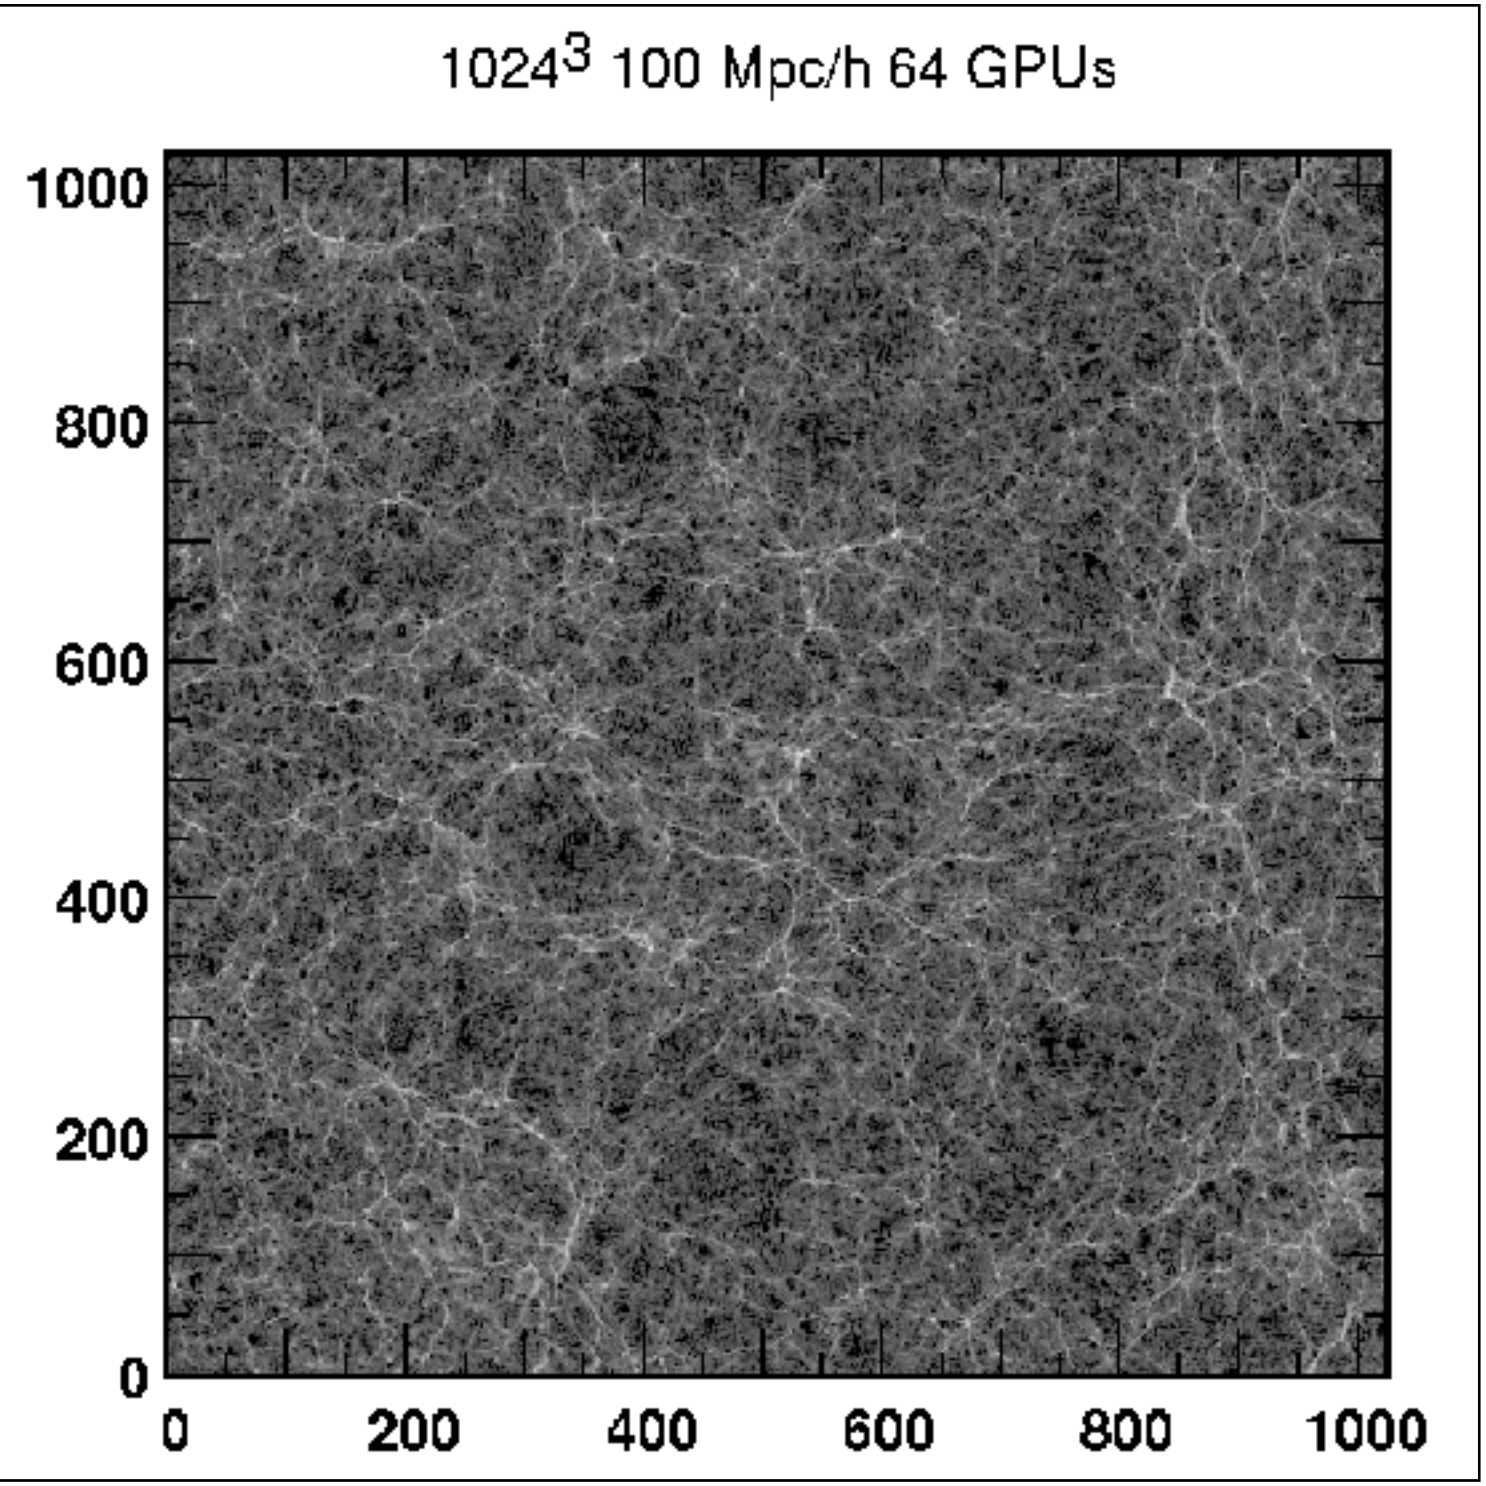

1.2 billion particles
(1024^3 real particles
+2 10^8 ghosts)
8 sec/tstep on 64 Teslas
with 25 % spent in
communications

with sort optimisation
we may expect 6 sec/
tstep
communication~40%

asynchronous coms ?

presented
@ GPU technology conference 2010,
unpublished

Simulations of the Cosmic Reionization

Explicit treatment of a fluid-like radiation

RT Moment **benefits** from grid based hydrodynamics methods

$$\begin{aligned}\frac{\partial N_\nu}{\partial t}+\frac{\partial \mathbf{F}_\nu}{\partial \mathbf{r}} &= S_\nu-\kappa_N N_\nu,\\ \frac{\partial \mathbf{F}_\nu}{\partial t}+c^2\frac{\partial \mathbf{P}_\nu}{\partial \mathbf{r}} &= -\kappa_F \mathbf{F}_\nu.\end{aligned}$$

RT Moment must also **satisfy the same constraints** as grid based hydrodynamics methods

We use an explicit solver:
Because of the speed of light, RT can be quite expensive compared to collisionless and hydro- dynamics

$$\Delta t < \frac{\Delta x}{v}$$

Courant Condition

Typically **200-600 rad. steps per dyn. step**

One solution (among others) : **Hardware acceleration**

Radiative Post-Processing with ATON (Aubert & Teyssier 2010)

128 GPUs (Titane/ CEA /CCRT) 100 Mpc - 1024^3

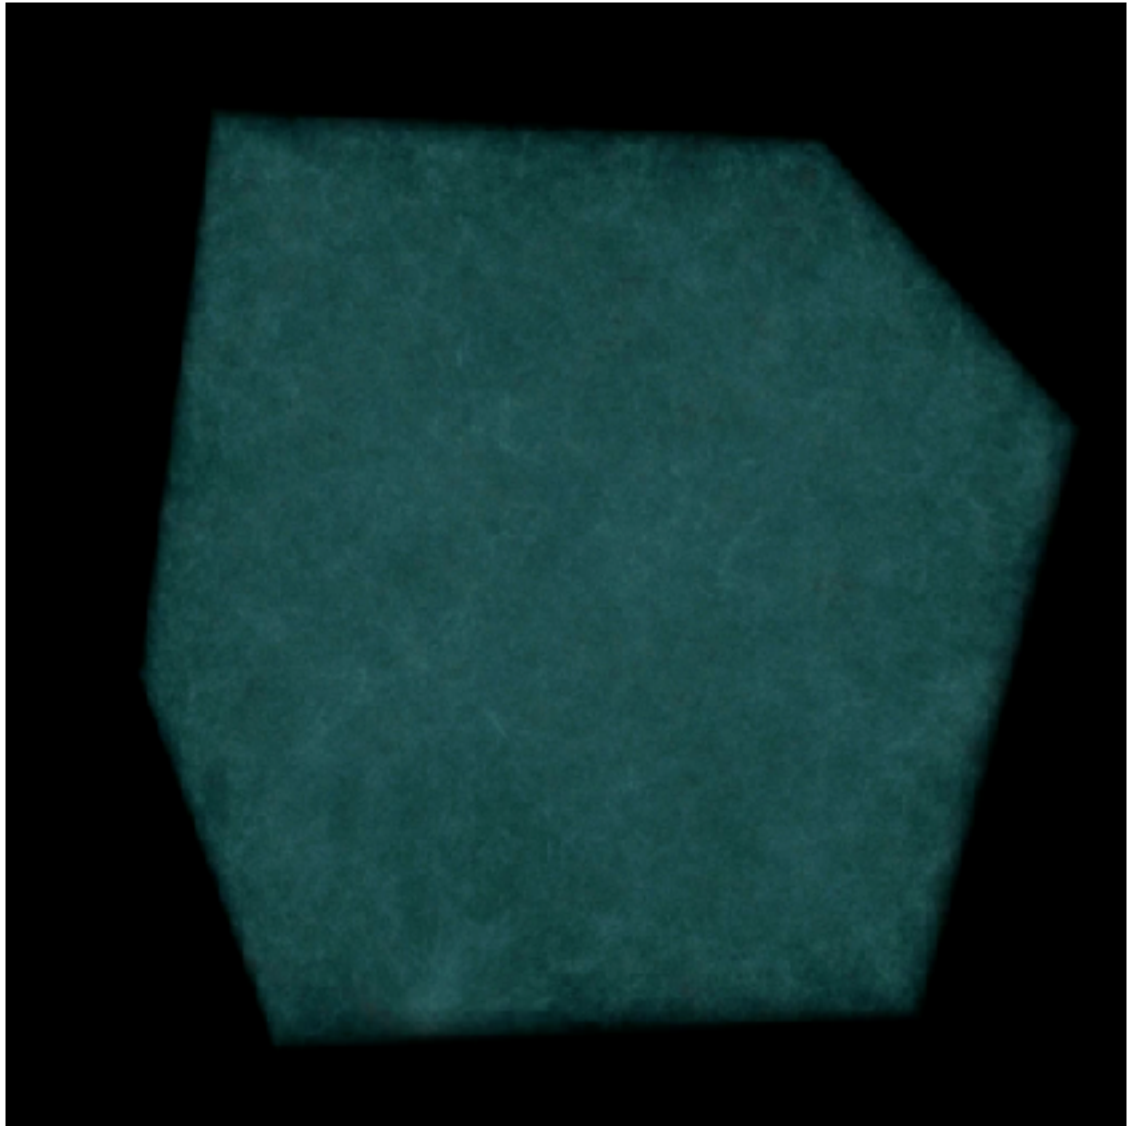

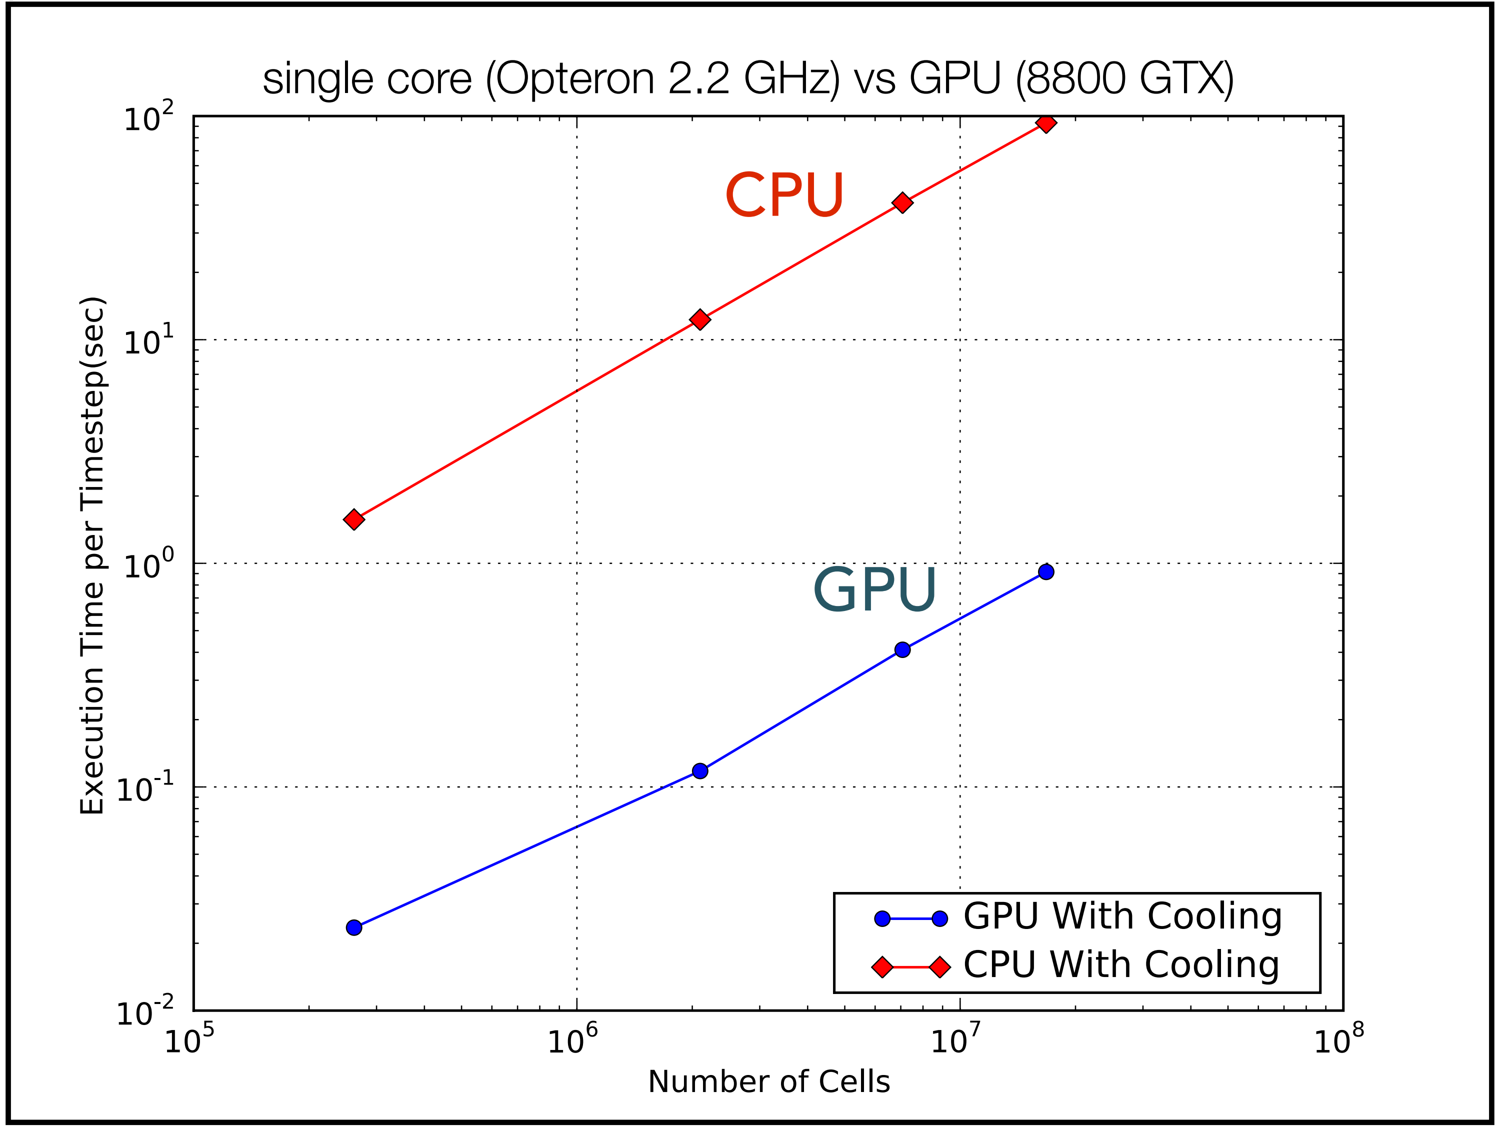

Close-circuit GPU code without transfers between the CPU host and the devices
Static cartesian mesh : ensures predictability of computation and memory access patterns
x80 acceleration factor : **c=1**

	ATON (or CUDATON) (Aubert & Teyssier 2008,2010)	RAMSES-CUDATON 2013 Ocvirk, Gillet, Shapiro, Aubert et al. 2016	EMMA Aubert et al. 2015
Purpose	Rad. **Post-Processing** of a pre-existing hydro simulation	On the fly interaction of RAMSES (CPU) & CUDATON (GPU)	Multi-purpose cosmological simulation code with Radiative transfer
Radiative hydrodynamics	🟥	🟩	🟩
Adaptive Mesh refinement	🟥	🟥	🟩
GPU	🟩 x 80 (RT only)	🟩 (x80) RT almost free	🟩 **x5** Overall (Poisson+Hydro +RT)

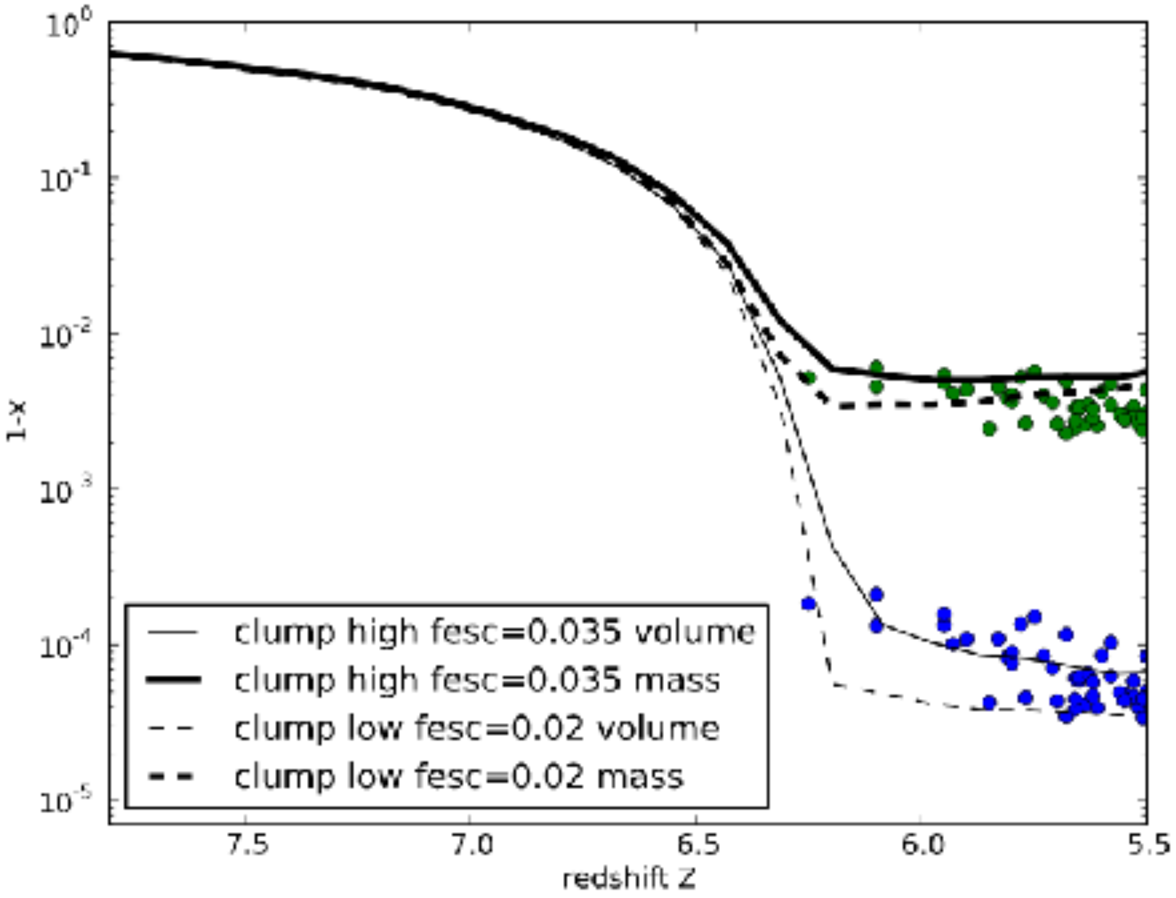

Figure 18. Evolution of the mass- and volume-averaged neutral fraction in the 100 Mpc h^{-1} box with a clumping factor assuming a high/low normalization (thin/thick lines). The values at $z = 6$ are consistent with measurements made by Fan et al. (2006) for both kinds of average methods (dots).

(A color version of this figure is available in the online journal.)

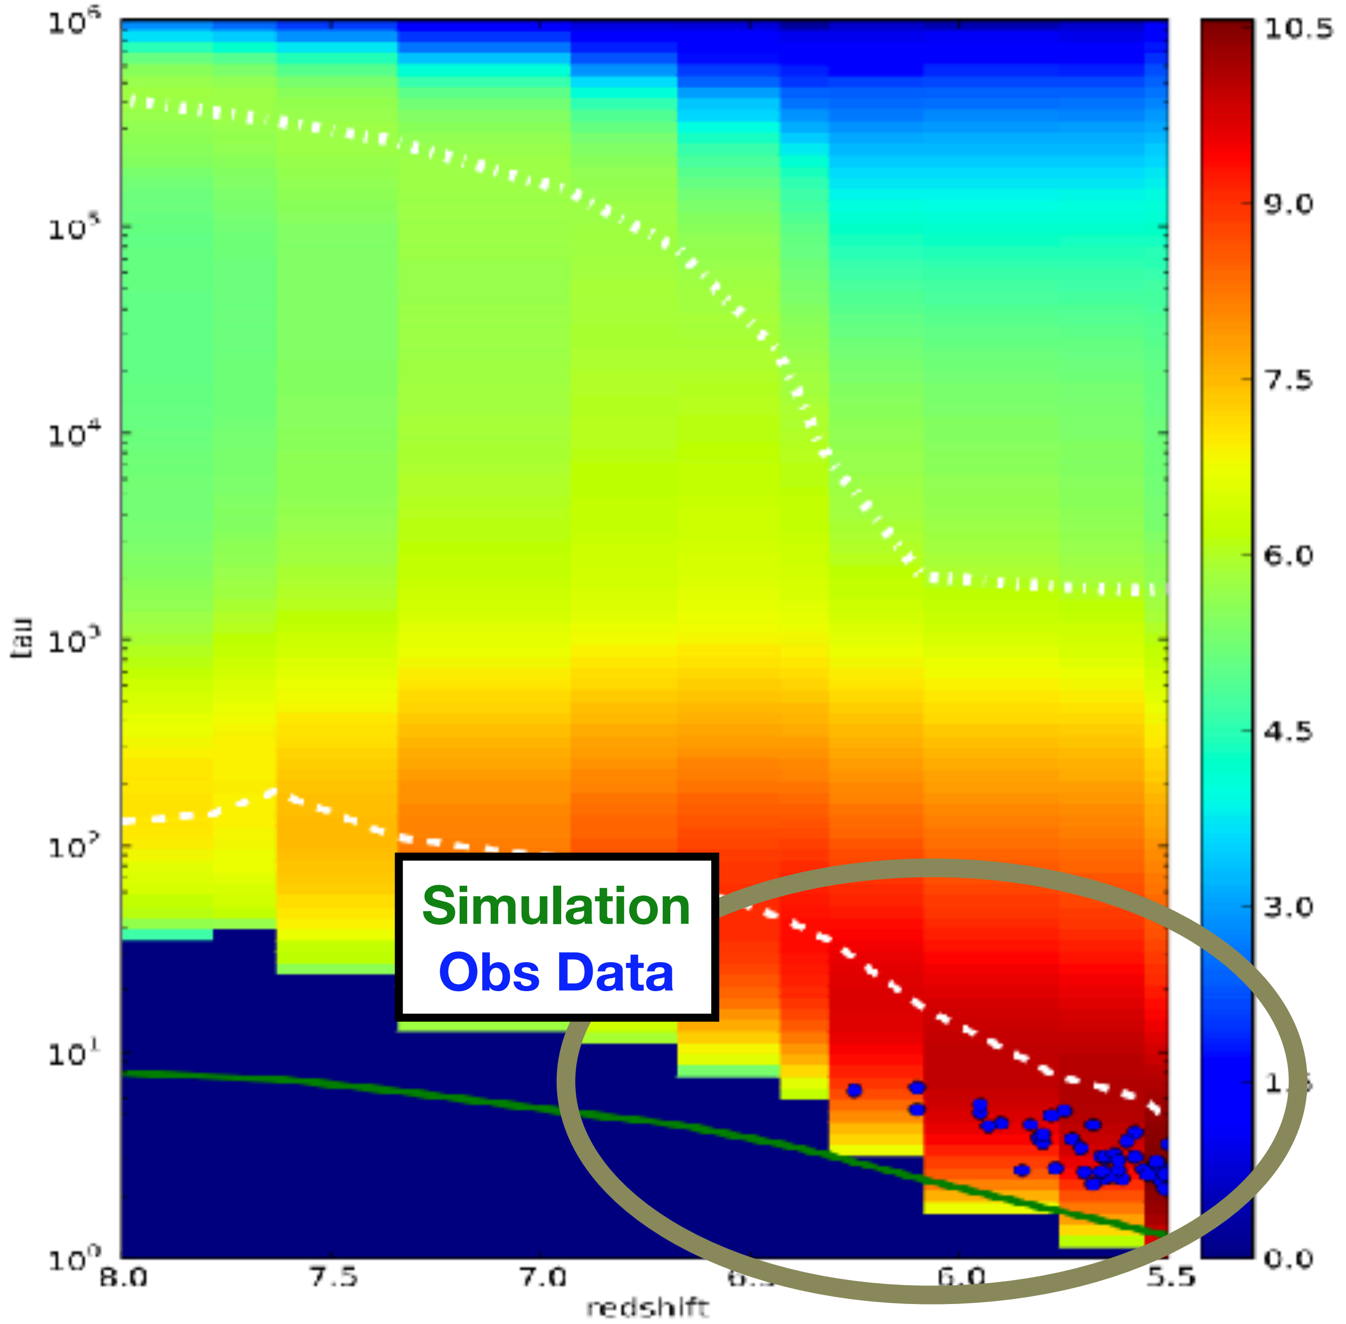

Simulation

Obs Data

Figure 21. Evolution of the effective GP optical depth τ_{eff} in the 100 Mpc h^{-1} box with high normalization clumping. Dots: measures of the effective optical depth $\tau_{\mathrm{eff}} = -\log T$ made by Fan et al. (2006). Green line: the average effective optical depth measure from our 100 Mpc h^{-1} simulation with subgrid clumping. Color map: probability distribution of $\tau \neq \tau_{\mathrm{eff}}$ measured in the same simulation, the scale being logarithmic. White dash-dotted line: the redshift evolution of $\langle\tau\rangle$. White dashed line: the redshift evolution of the maximum of the pdf of τ.

(A color version of this figure is available in the online journal.)

We can match xion(z) deduced from
the observations
without matching the observation
themselves…
(Aubert & Teyssier 2010)

Calibration of Reionization Simulations using LyA Forest

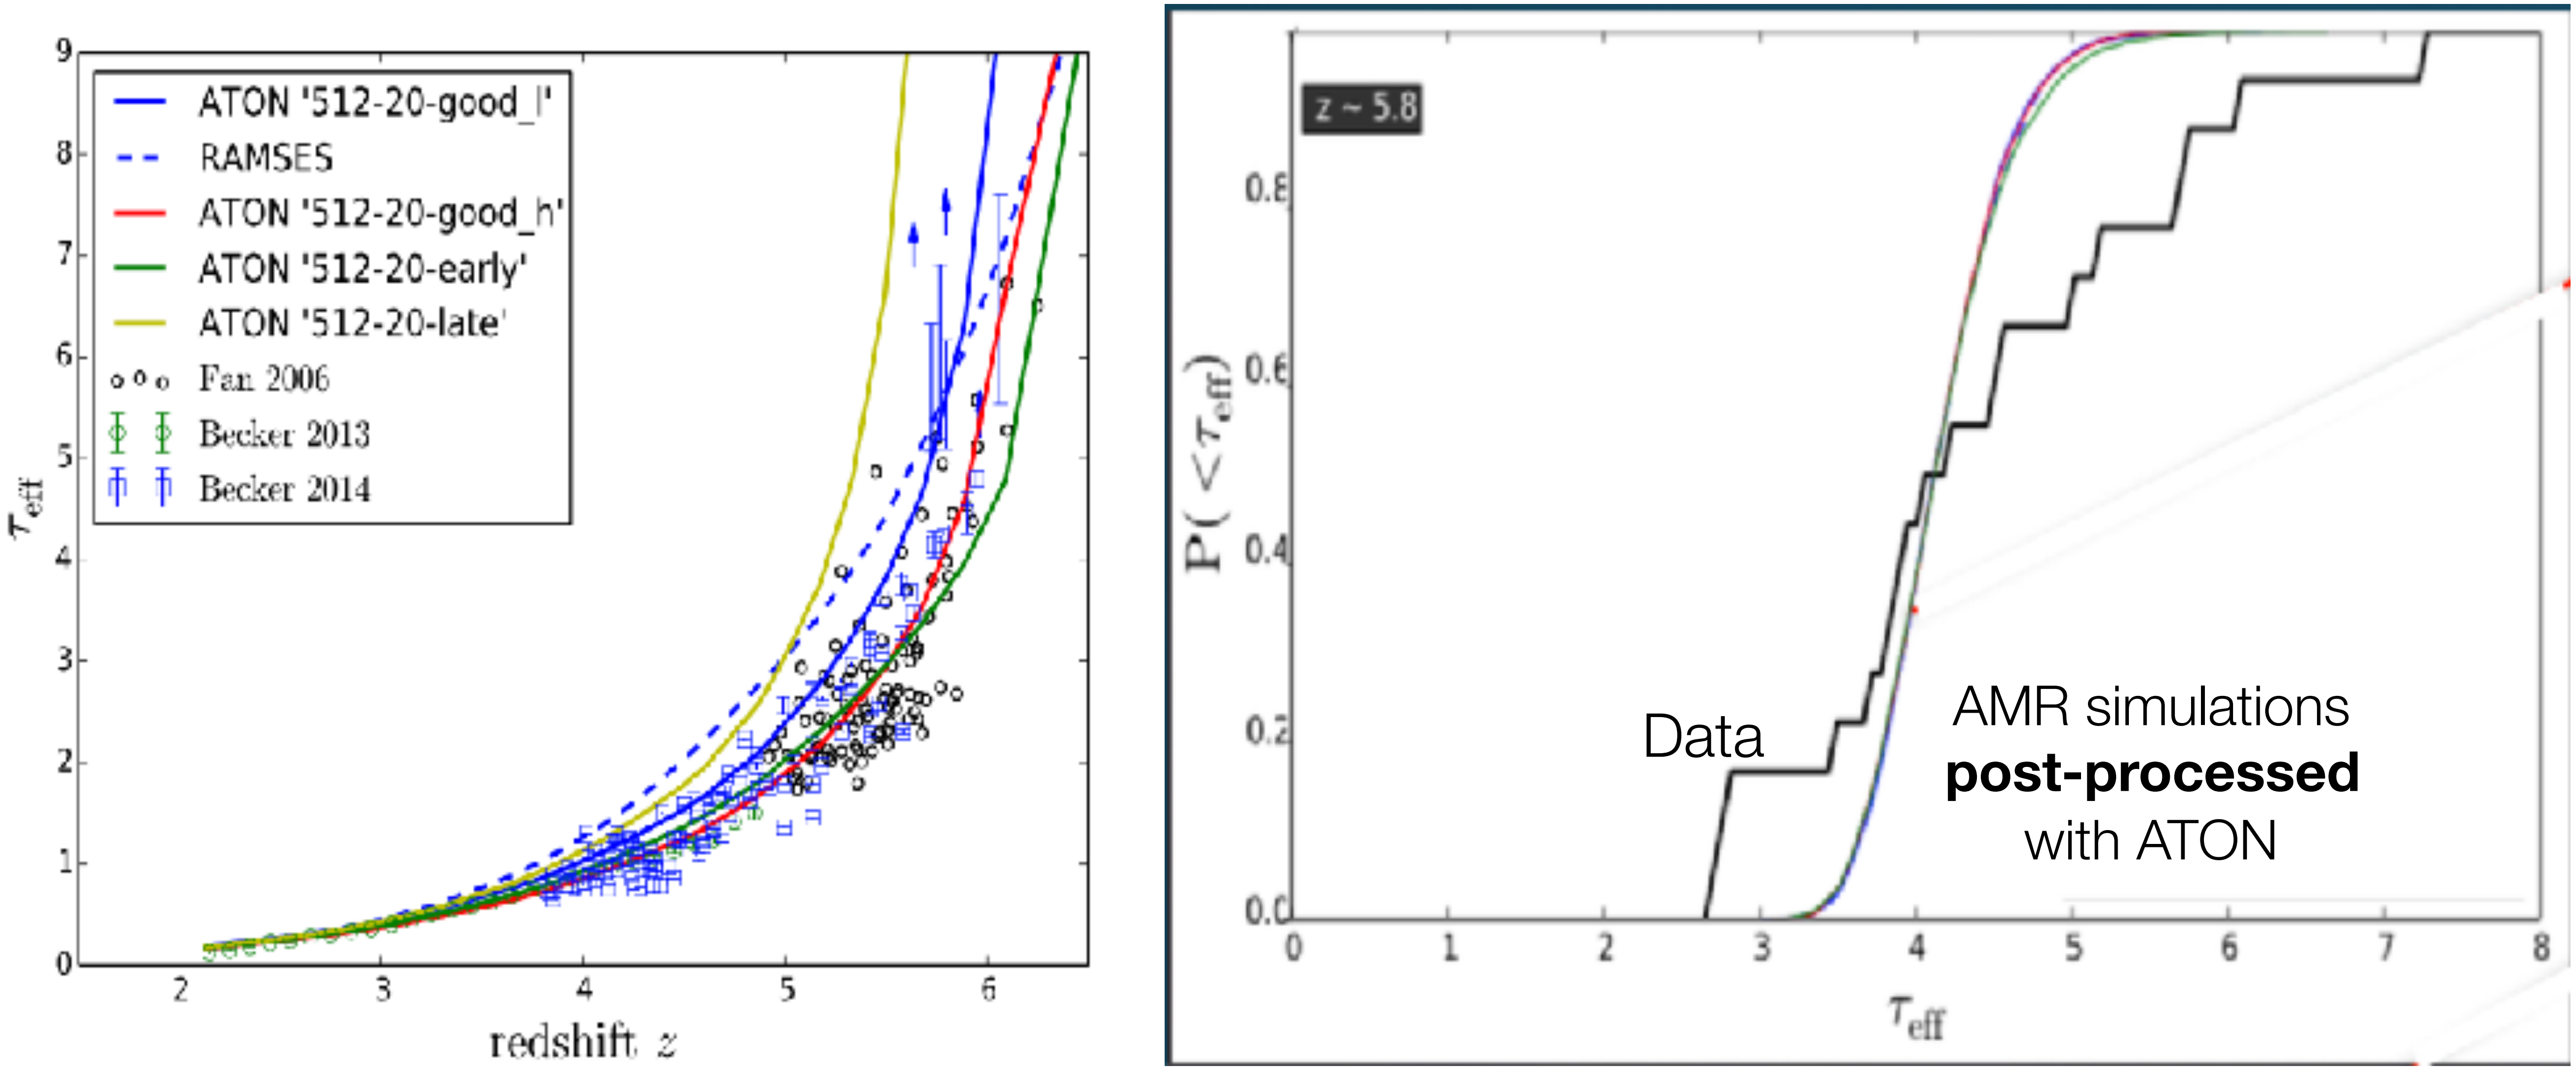

Data

AMR simulations **post-processed** with ATON

(a)

(b)

Ly-A data put constraints on reionization models: overlap by z~7, accelerated evolution of opacity at z>6, **large scatter of IGM opacities at late times !** **Probable significant contribution of rare and bright sources (quasars?) to the UV** background at z~6

Chardin +16

Indications of a late reionization ?

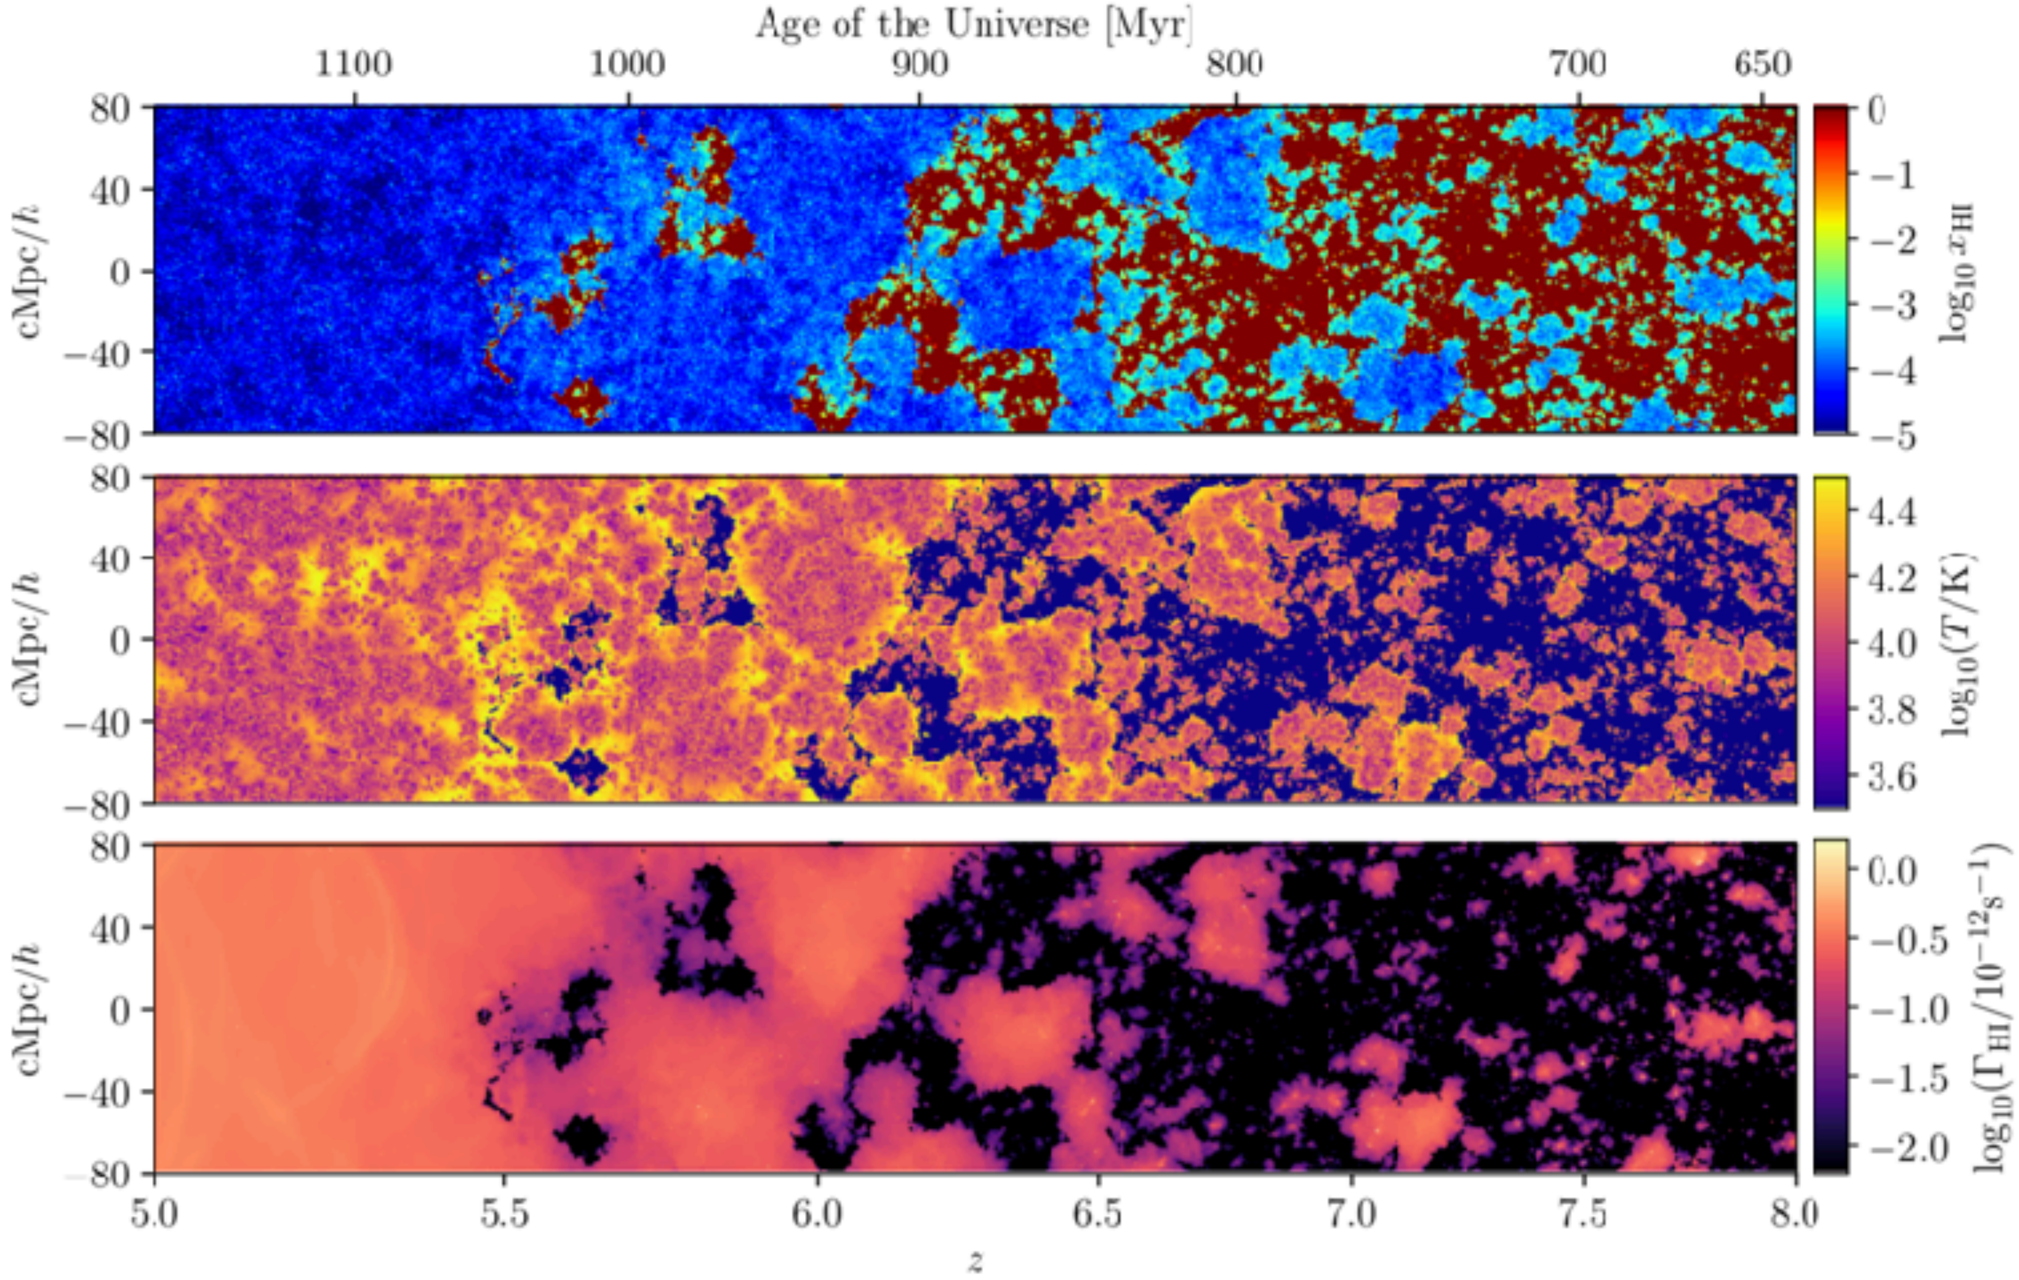

Kulkarni+ 2018

Radiative **Post-Processing** with cudATON (Aubert & Teyssier, 2010).
160 Mpc/h - 1024^3 - SPH-Gadget Sherwood Simulation

Indications of a late reionization ?

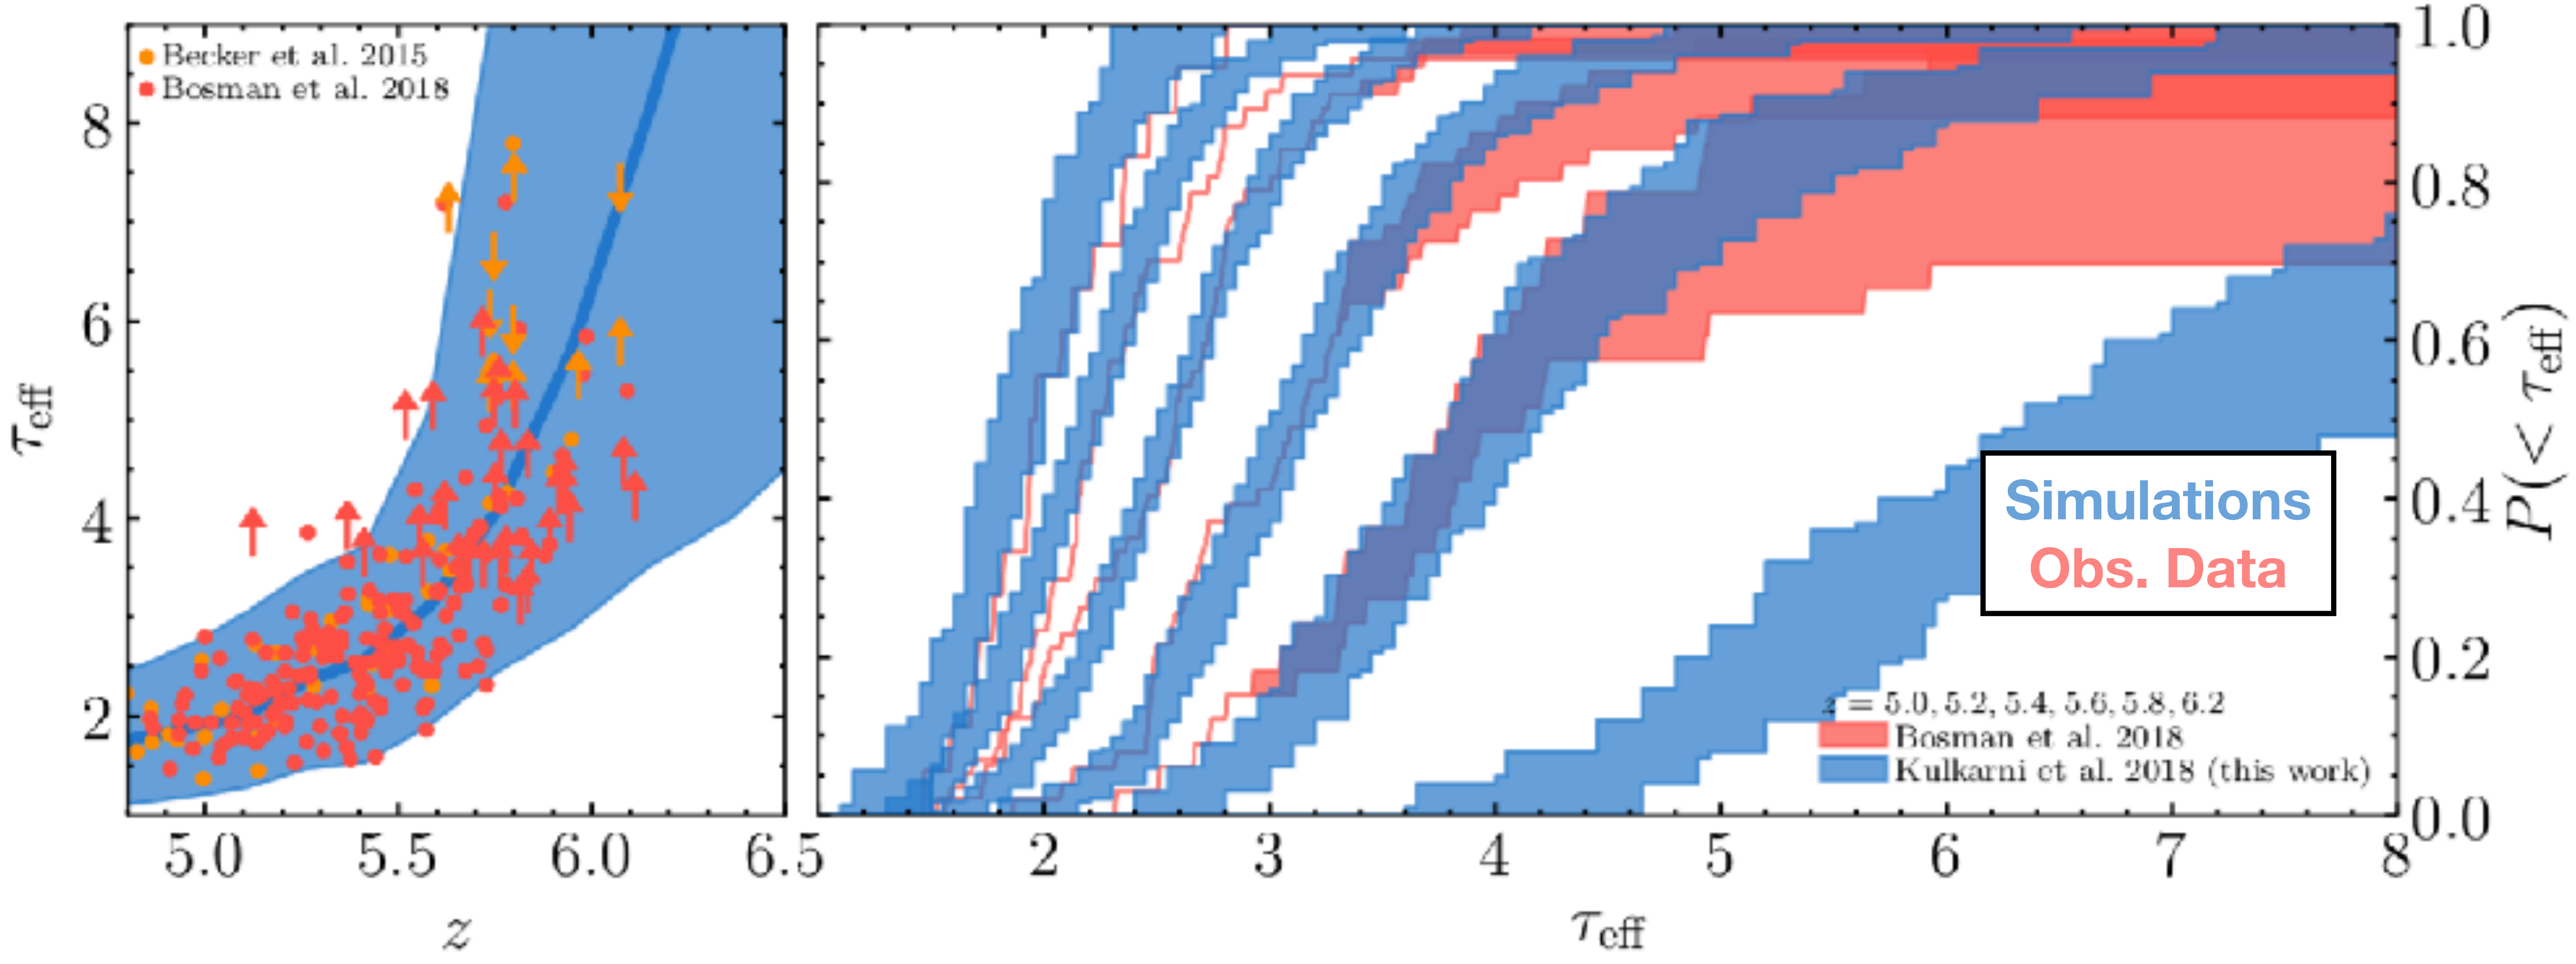

A much better fit to observed distributions of Intergalactic medium opacities, from Ly-Alpha Forest, is recovered with **a z~5.3 reionization (without quasars)**.
(« observed » xion(z) is likely to be a poor constraint of the reionization history)

Kulkarni+ 2018

	ATON (or CUDATON) (Aubert & Teyssier 2008,2010)	RAMSES-CUDATON 2013 Ocvirk, Gillet, Shapiro, Aubert et al. 2016	EMMA Aubert et al. 2015
Purpose	Rad. Post-Processing of a pre-existing hydro simulation	On the fly interaction of RAMSES (CPU) & CUDATON (GPU)	Multi-purpose cosmological simulation code with RT
Radiative hydrodynamics	🟥	🟩	🟩
Adaptive Mesh refinement	🟥	🟥	🟩
GPU	🟩 x 80 (RT only)	🟩 (x80) RT almost free	🟩 **x5** Overall (Poisson+Hydro +RT)

Setup: CoDall specs

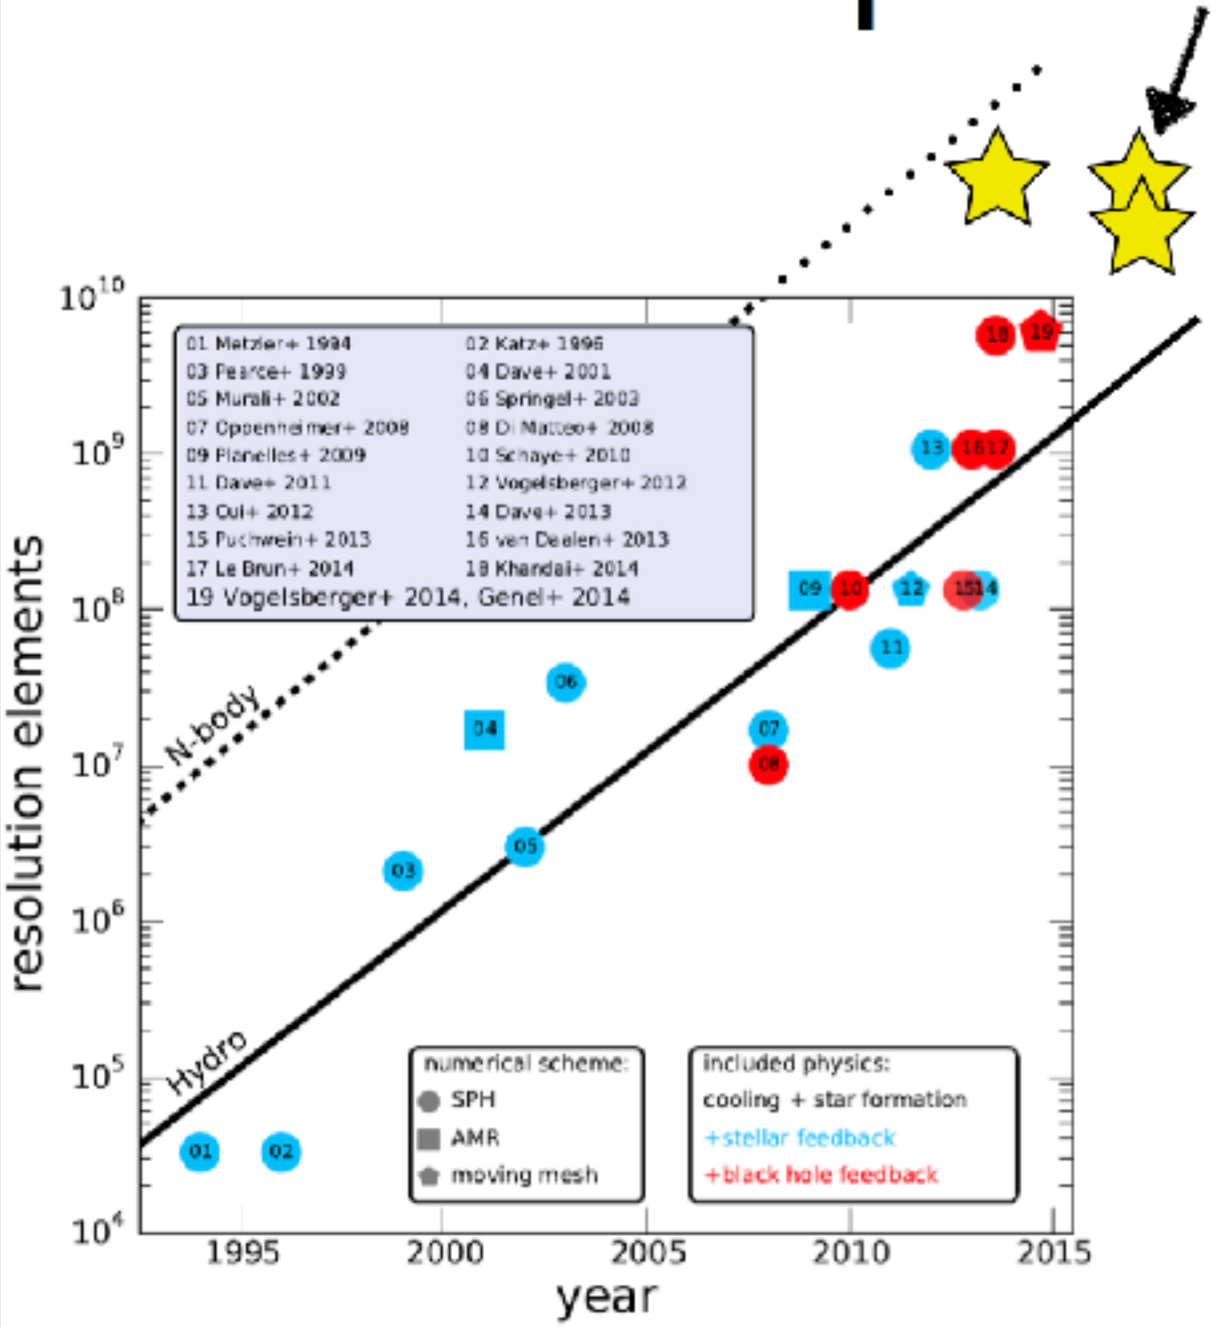

(taken from illustris website)

- 16384 GPUs, 65536 CPUs
- $64\ h^{-1}$ Mpc side, 4096^3 grid
- $M_{halo_{min}} \sim 1 \times 10^8\ M_\odot$
- $\Delta x \sim 22$ kpc comoving (< 3.2 kpc physical)
- $z_{end}=5.8$
- ~ 6 days runtime, 2 PB data
- Planck 2013 cosmology
- New ICs: $M_{Virgo} = 2.e14$ Msun

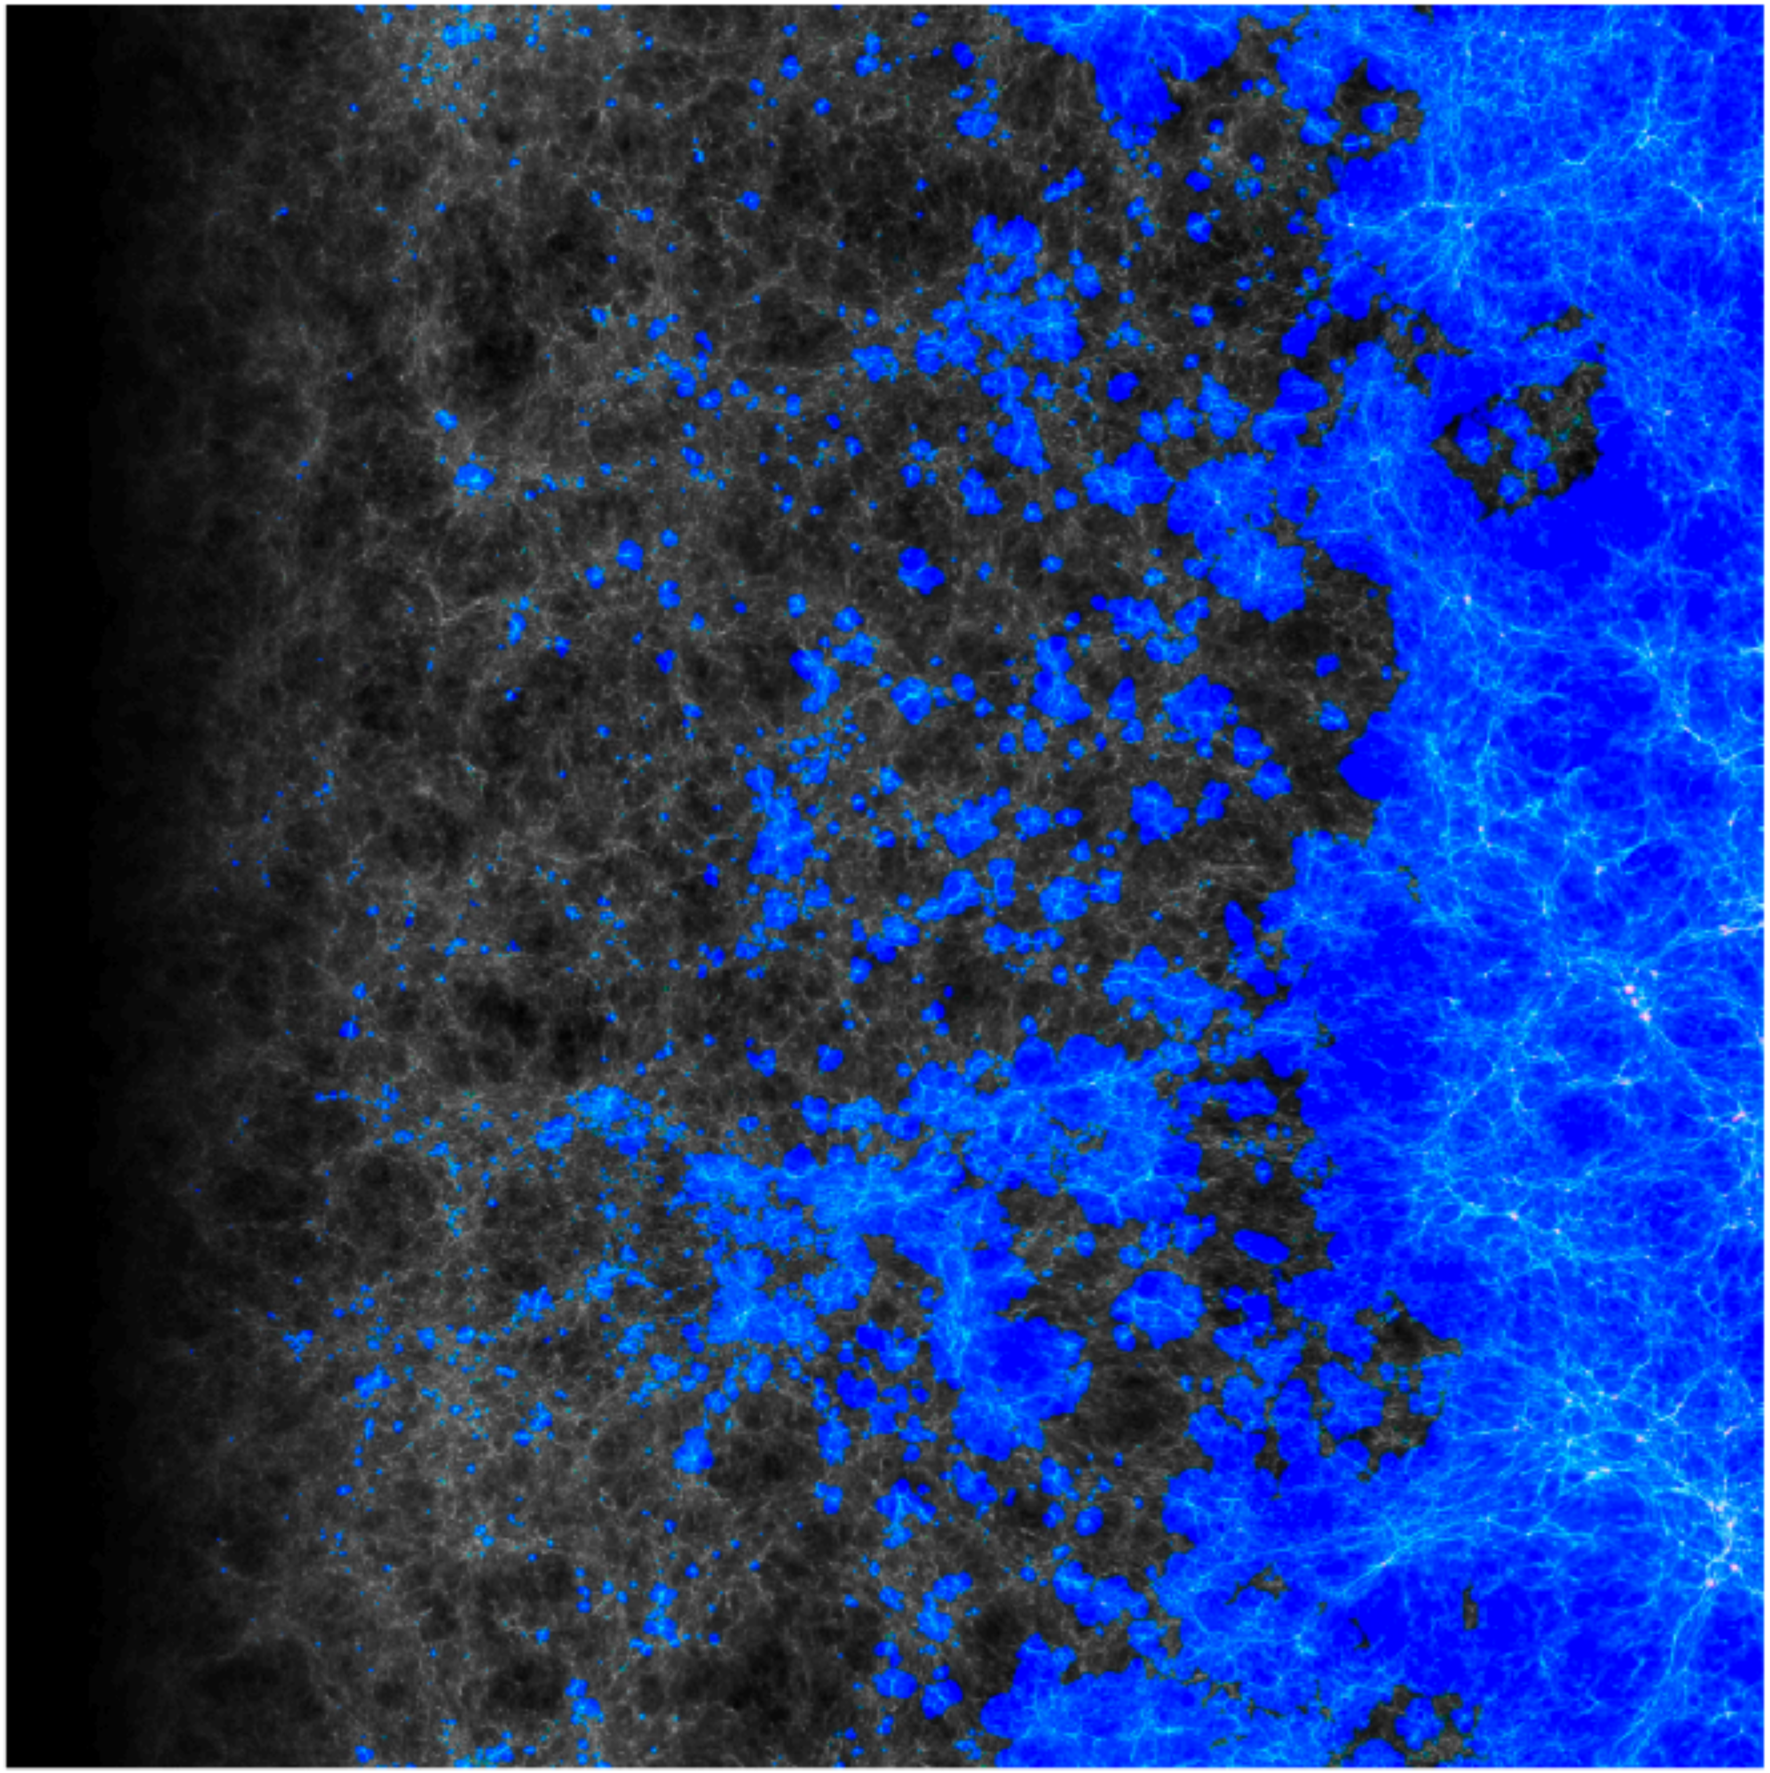

Figure 1. Illustration of the reionization process within the CoDaII simulation. The figure spans the full 94 Mpc in x and y, and spans from redshift z= 150 (left) to z= 5.8 (right) along the x axis. It is made from the concatenation of vertical, 4-cells-wide stripes taken from a series of ~ 1020 high-frequency CoDaII outputs. The color encodes temperature on a blue to red scale: blue regions are photo-heated, while bright red regions correspond to regions heated by supernovae feedback and accretion shocks. Brightness indicates the gas density contrast.

Coda I
(Ocvirk+ 16,
Dawoodbhoy+18)
& **Coda II**
(Ocvirk+ in prep.,
Lewis+ in prep.)

Galaxy population properties during the reionization with **radiative hydrodynamics Ramses-CUDATON**

Star Formation & suppression, environmental dependance
Photon budget & Escape fraction
Luminosity functions

(see **Ocvirk's** talk tomorrow

Next Stages :
on Piz Daint (CSSC, CH) &
Summit (ORNL, US)

	ATON (or CUDATON) (Aubert & Teyssier 2008,2010)	RAMSES-CUDATON 2013 Ocvirk, Gillet, Shapiro, Aubert et al. 2016	EMMA Aubert et al. 2015
Purpose	Rad. Post-Processing of a pre-existing hydro simulation	On the fly interaction of RAMSES (CPU) & CUDATON (GPU)	Multi-purpose cosmological simulation code with RT
Radiative hydrodynamics	🟥	🟩	🟩
Adaptive Mesh refinement	🟥	🟥	🟩
GPU	🟩 x 80 (RT only)	🟩 (x80) RT almost free	🟩 **x5** Overall (Poisson+Hydro +RT)

AMR Cosmological RT with **EMMA**

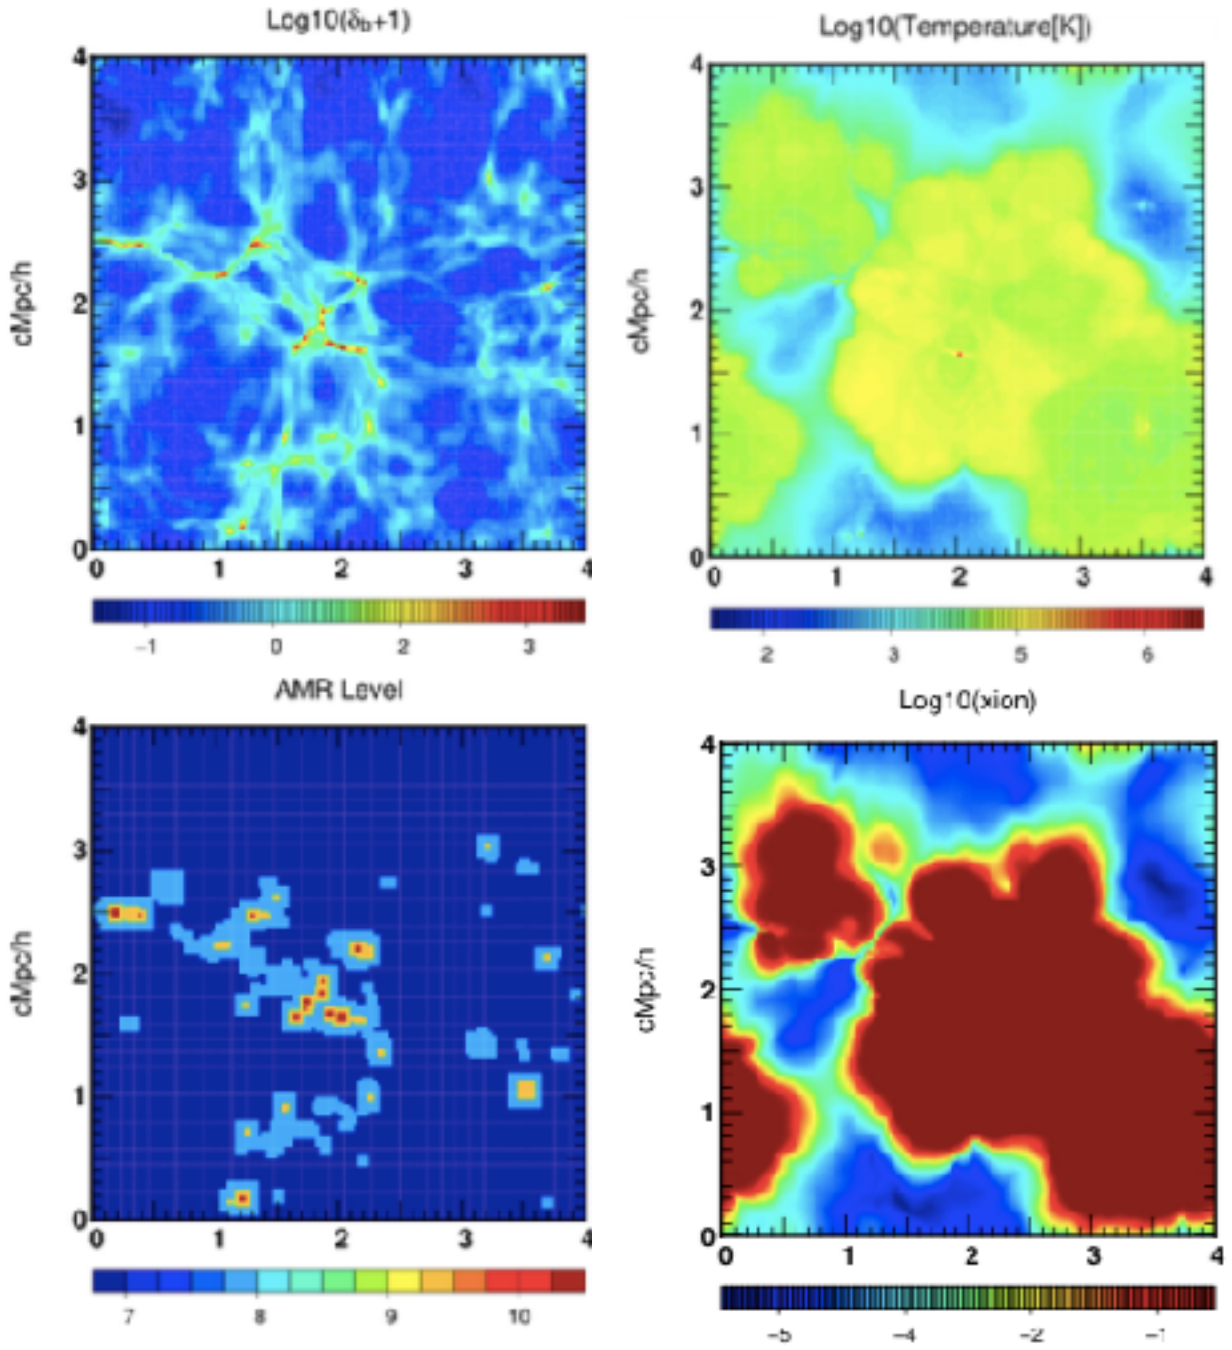

4 Mpc - 128^3 + 5 AMR levels

- **E**lectromagnétisme et **M**écanique sur **M**aille **A**daptative
- Full **standalone** cosmological code
- Collisionless Dynamics (PM)+ Hydro (Godunov/MUSCL/HLL) +Radiative trasnfer (M1)
- Full **AMR** radiative transport (like e.g. Ramses-RT (Rosdahl et al. 2013)) or restricted to the Coarse grid with thermo-chemistry on refined levels
- Star Formation + SN Feedback
- C+MPI Parallelisation (scales up to 2048 cores and 1024^3 coarse cells)
- **Optional GPU** (CUDA) acceleration for the Poisson , Hydro and RT solver

Aubert, Deparis & Ocvirk, MNRAS 2015

Global performances

x4-x15 acceleration (including non-GPU procedures and general overheads)

Acceleration can be obtained at the x15 level when strict floating point operations are required (gcc O2 or icc O2 -strict-fp operation)

CPU can be quite competitive if this restriction is lifted : x4 acceleration only (icc O2)

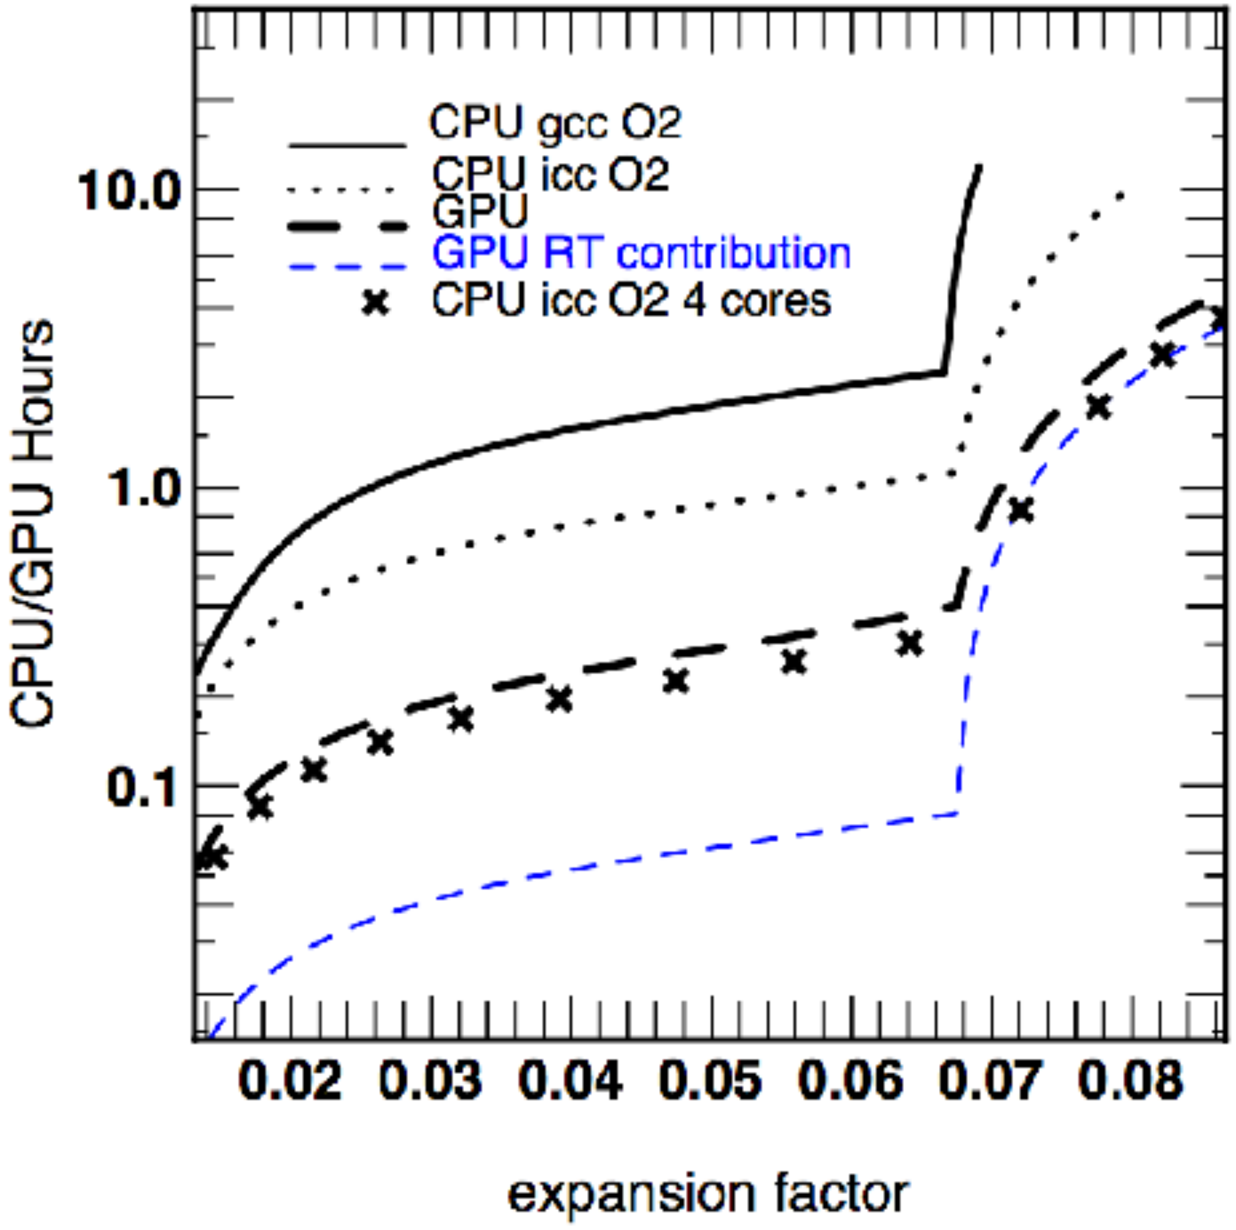

Figure 22. Comparison of the cumulative time spent to reach a given expansion factor for a 4 Mpc/h-128^3 cosmological simulation of the reionization. Times are given for a single computing device (i.e. 1 GPU or 1 CPU core). The thick black dashed line stands for the GPU run performed on a M2090 Nvidia GPU whereas the thin dashed blue line stands for the contribution of radiative transfer to this cost. The black solid (resp. dotted) line stands for a single CPU core (2.7 GHz Sandybridge Westemere) using the *gcc-O2* (resp. *icc-O2*) binary. The symbols stand for a 4-core CPU calculation using *icc-O2* on a Curie node.

Aubert, Deparis & Ocvirk, MNRAS 2015

Global design : Vectorisation

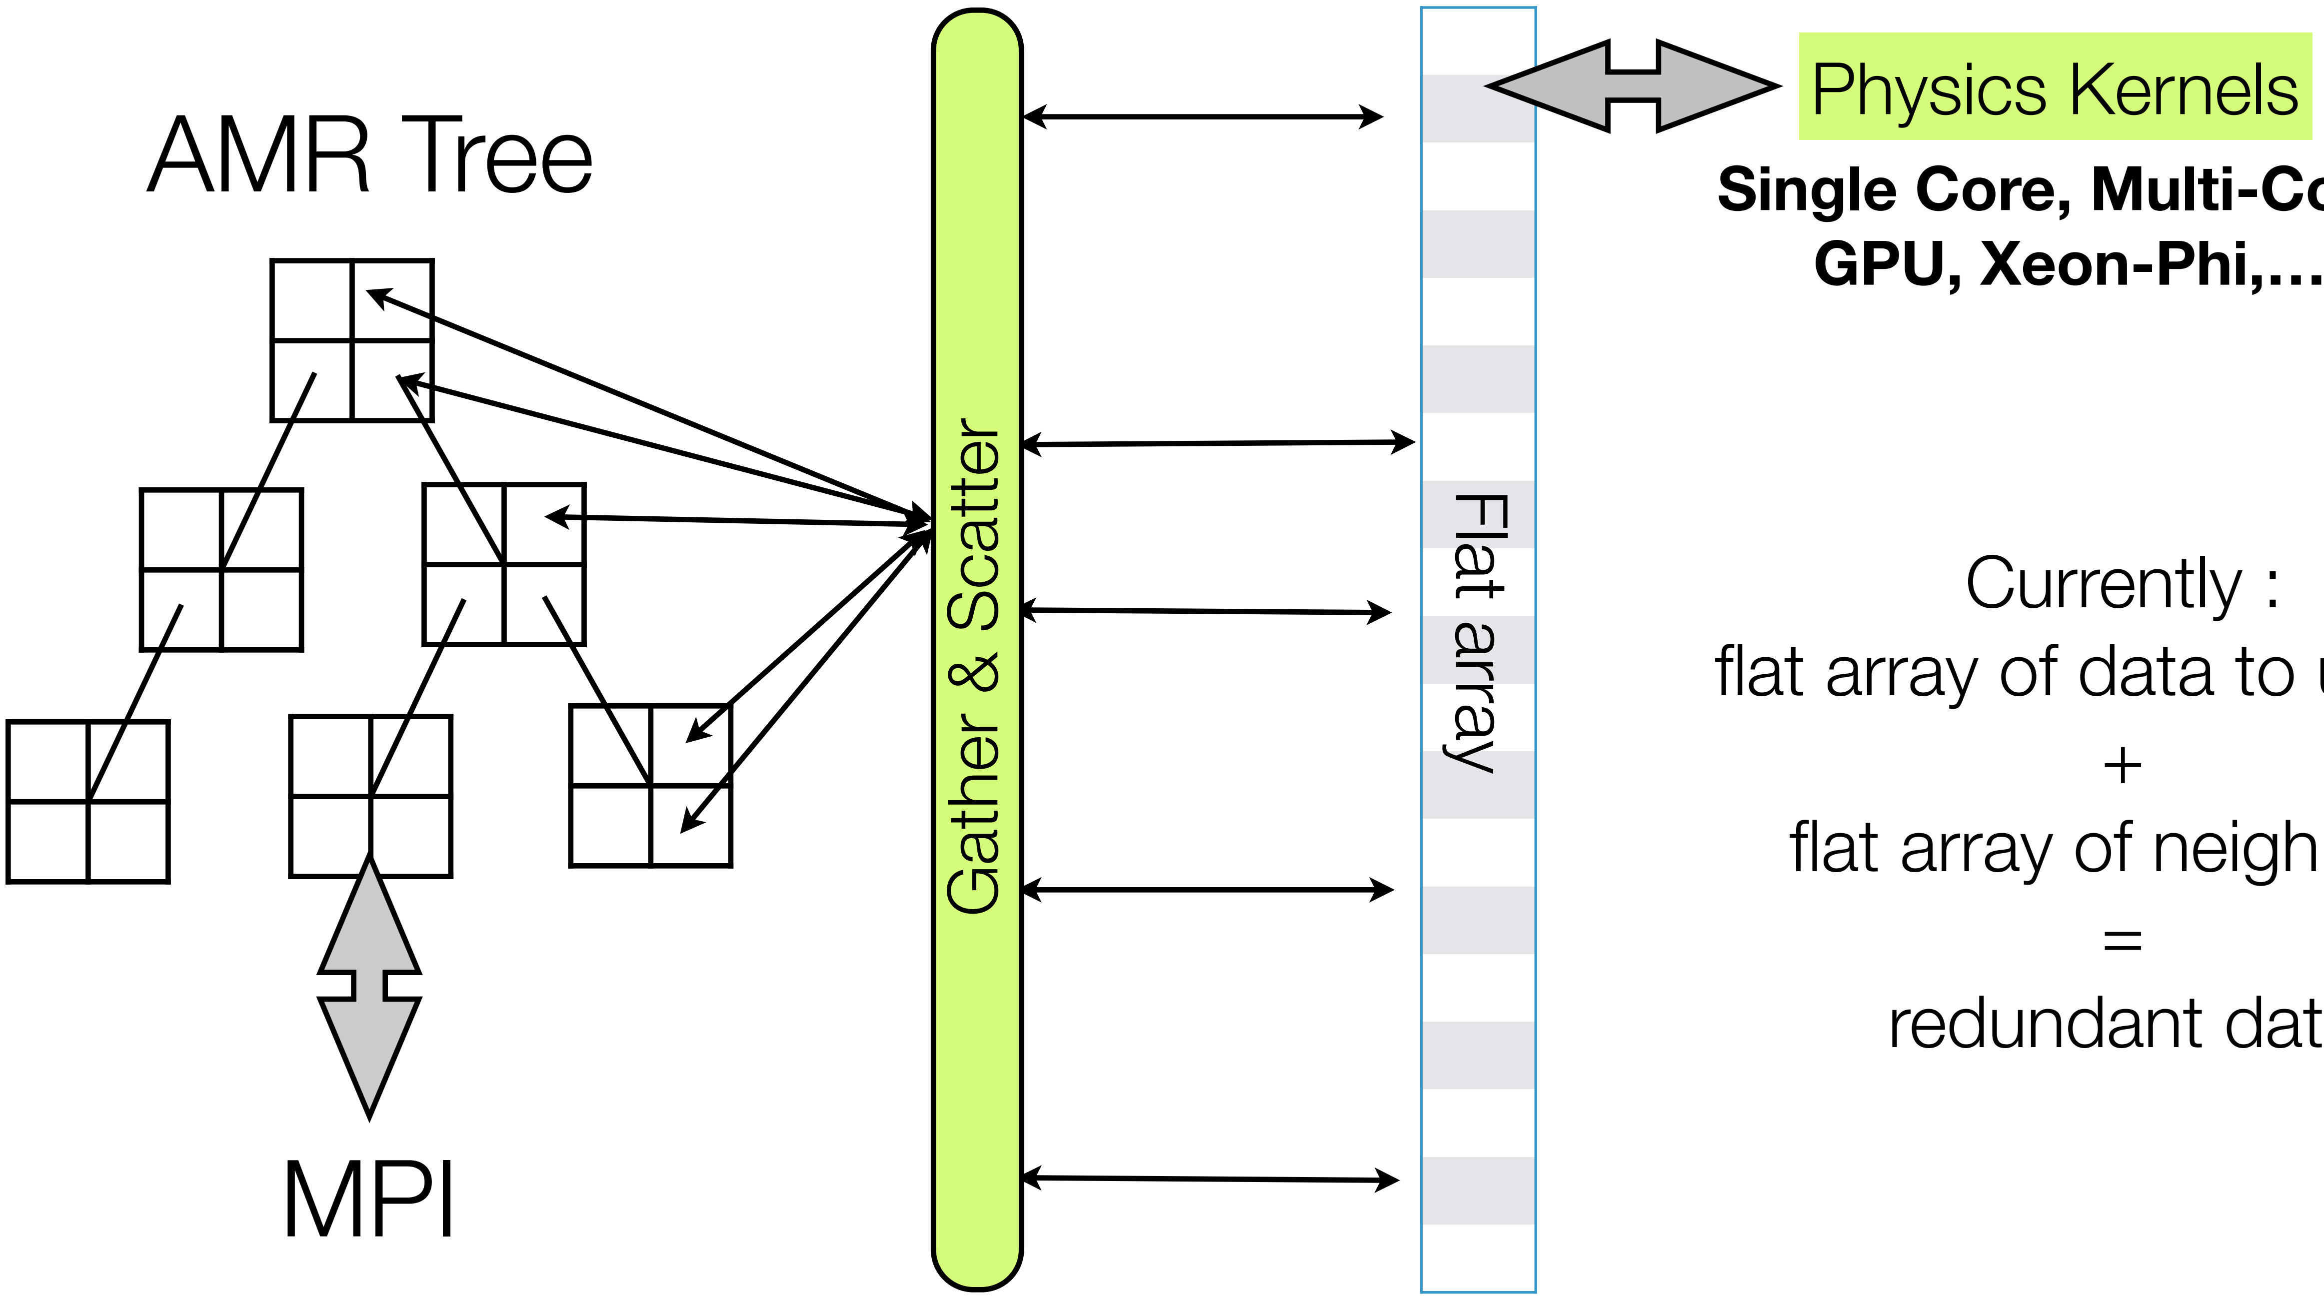

Single Core, Multi-Core, GPU, Xeon-Phi,…

Currently :
flat array of data to update
+
flat array of neighbors
=
redundant data

Decoupling Physics from Data Logistics - Stream vectors to accelerators

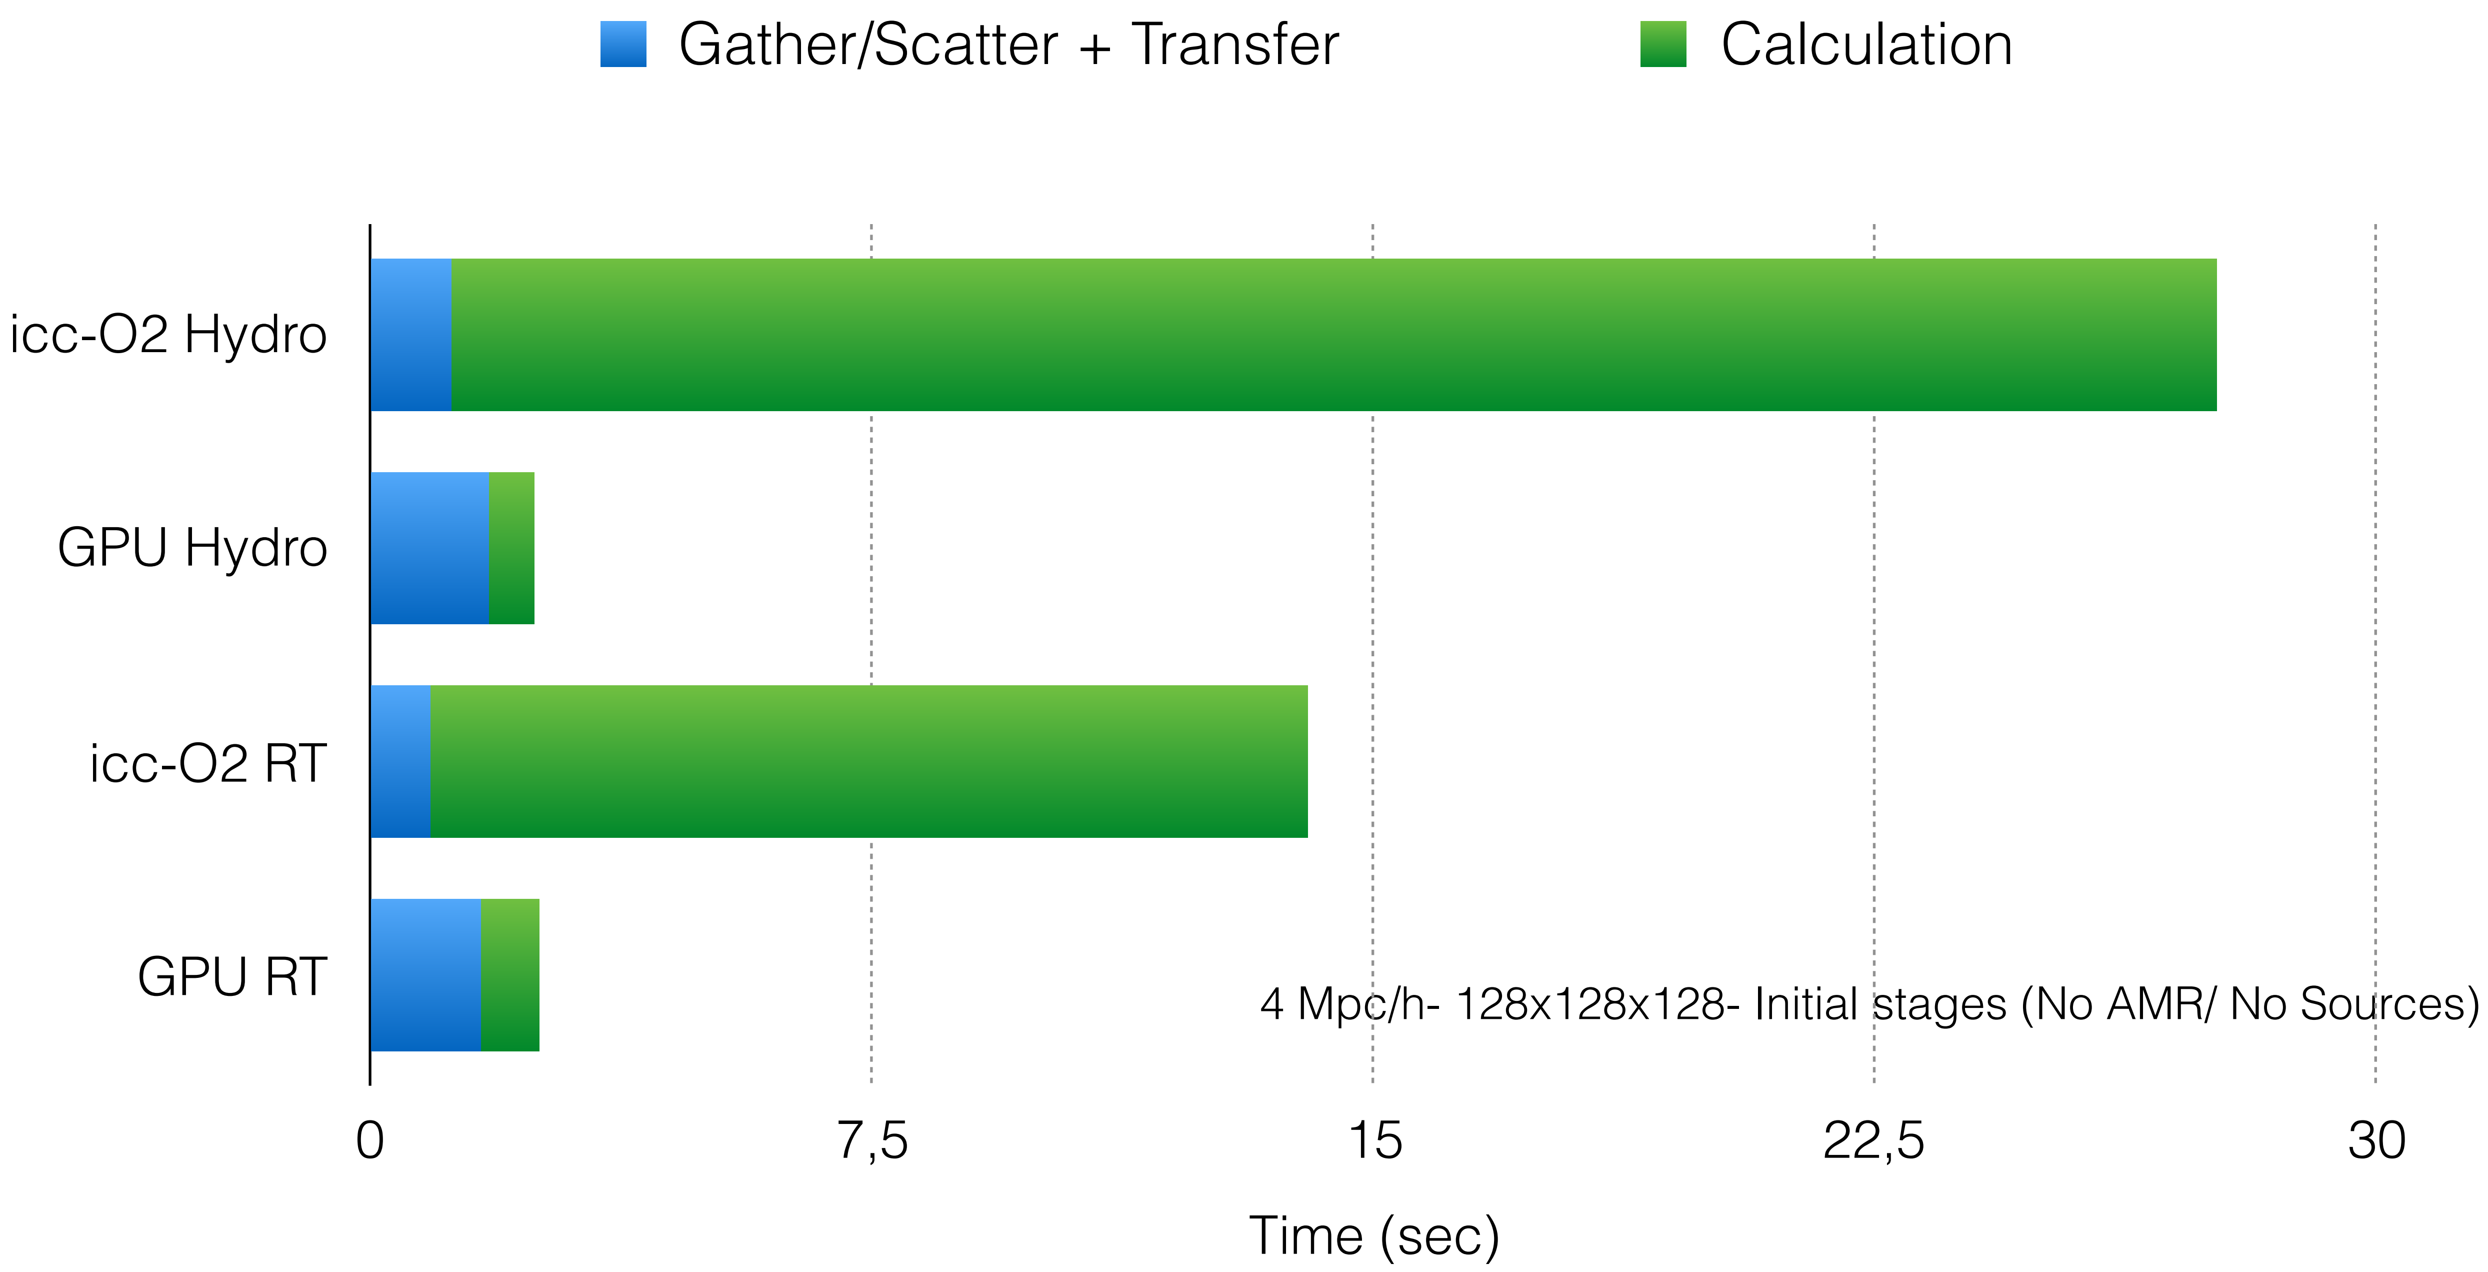

4 Mpc/h- 128x128x128- Initial stages (No AMR/ No Sources)

Data preparation and transfer currently limit the performance
Note that GS/Transfer x2 for GPUs

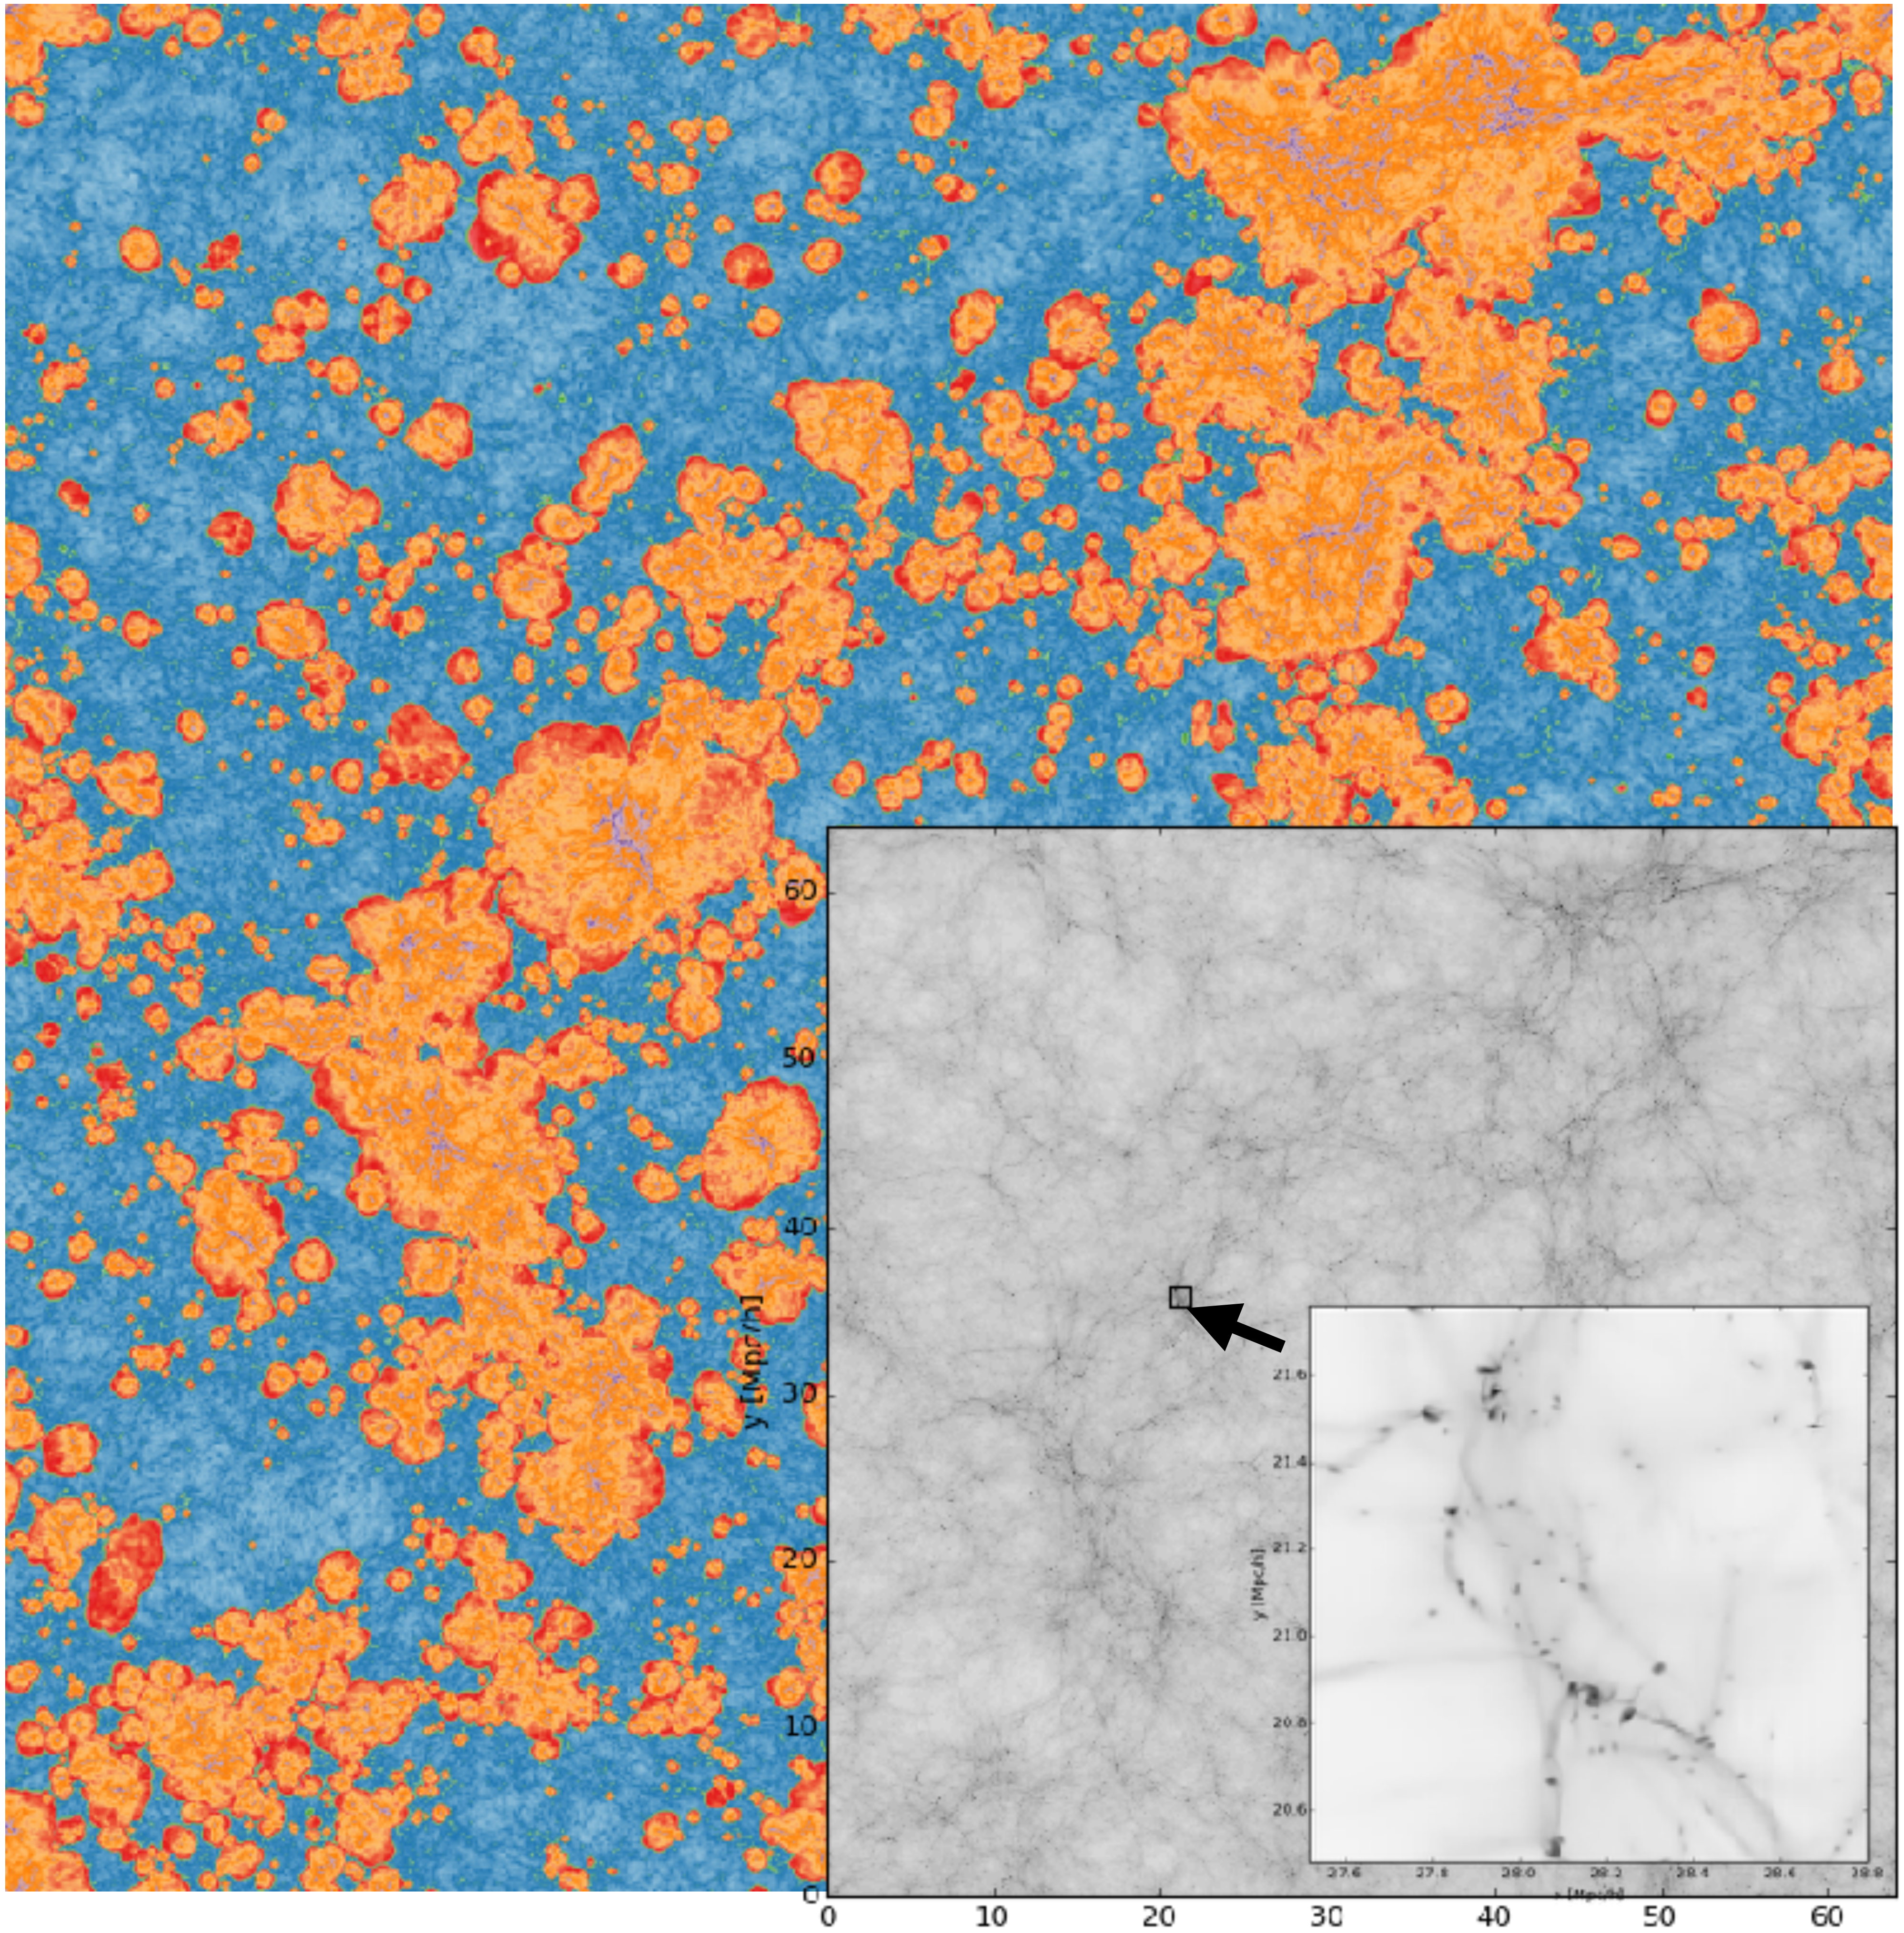

CODA I-AMR

91 Mpc/2048^3
16 billions resolution elements with AMR
@ z=6

32768 cores+4096 GPUs on Titan(DOE/ORNL) using **EMMA** simulation code (Aubert+ 15) (1.4 GPU acceleration rate)

20+ millions
cpu hours
Jan.-Mar 2017

INCITE CODA
(PI : Shapiro)

WMAP5-CLUES ICs
spatial res.=500 pc
mass res~2e6 Msol
stellar res ~70 000 Msol
reduced speed of light (10%)

Reionization times of z=0 haloes

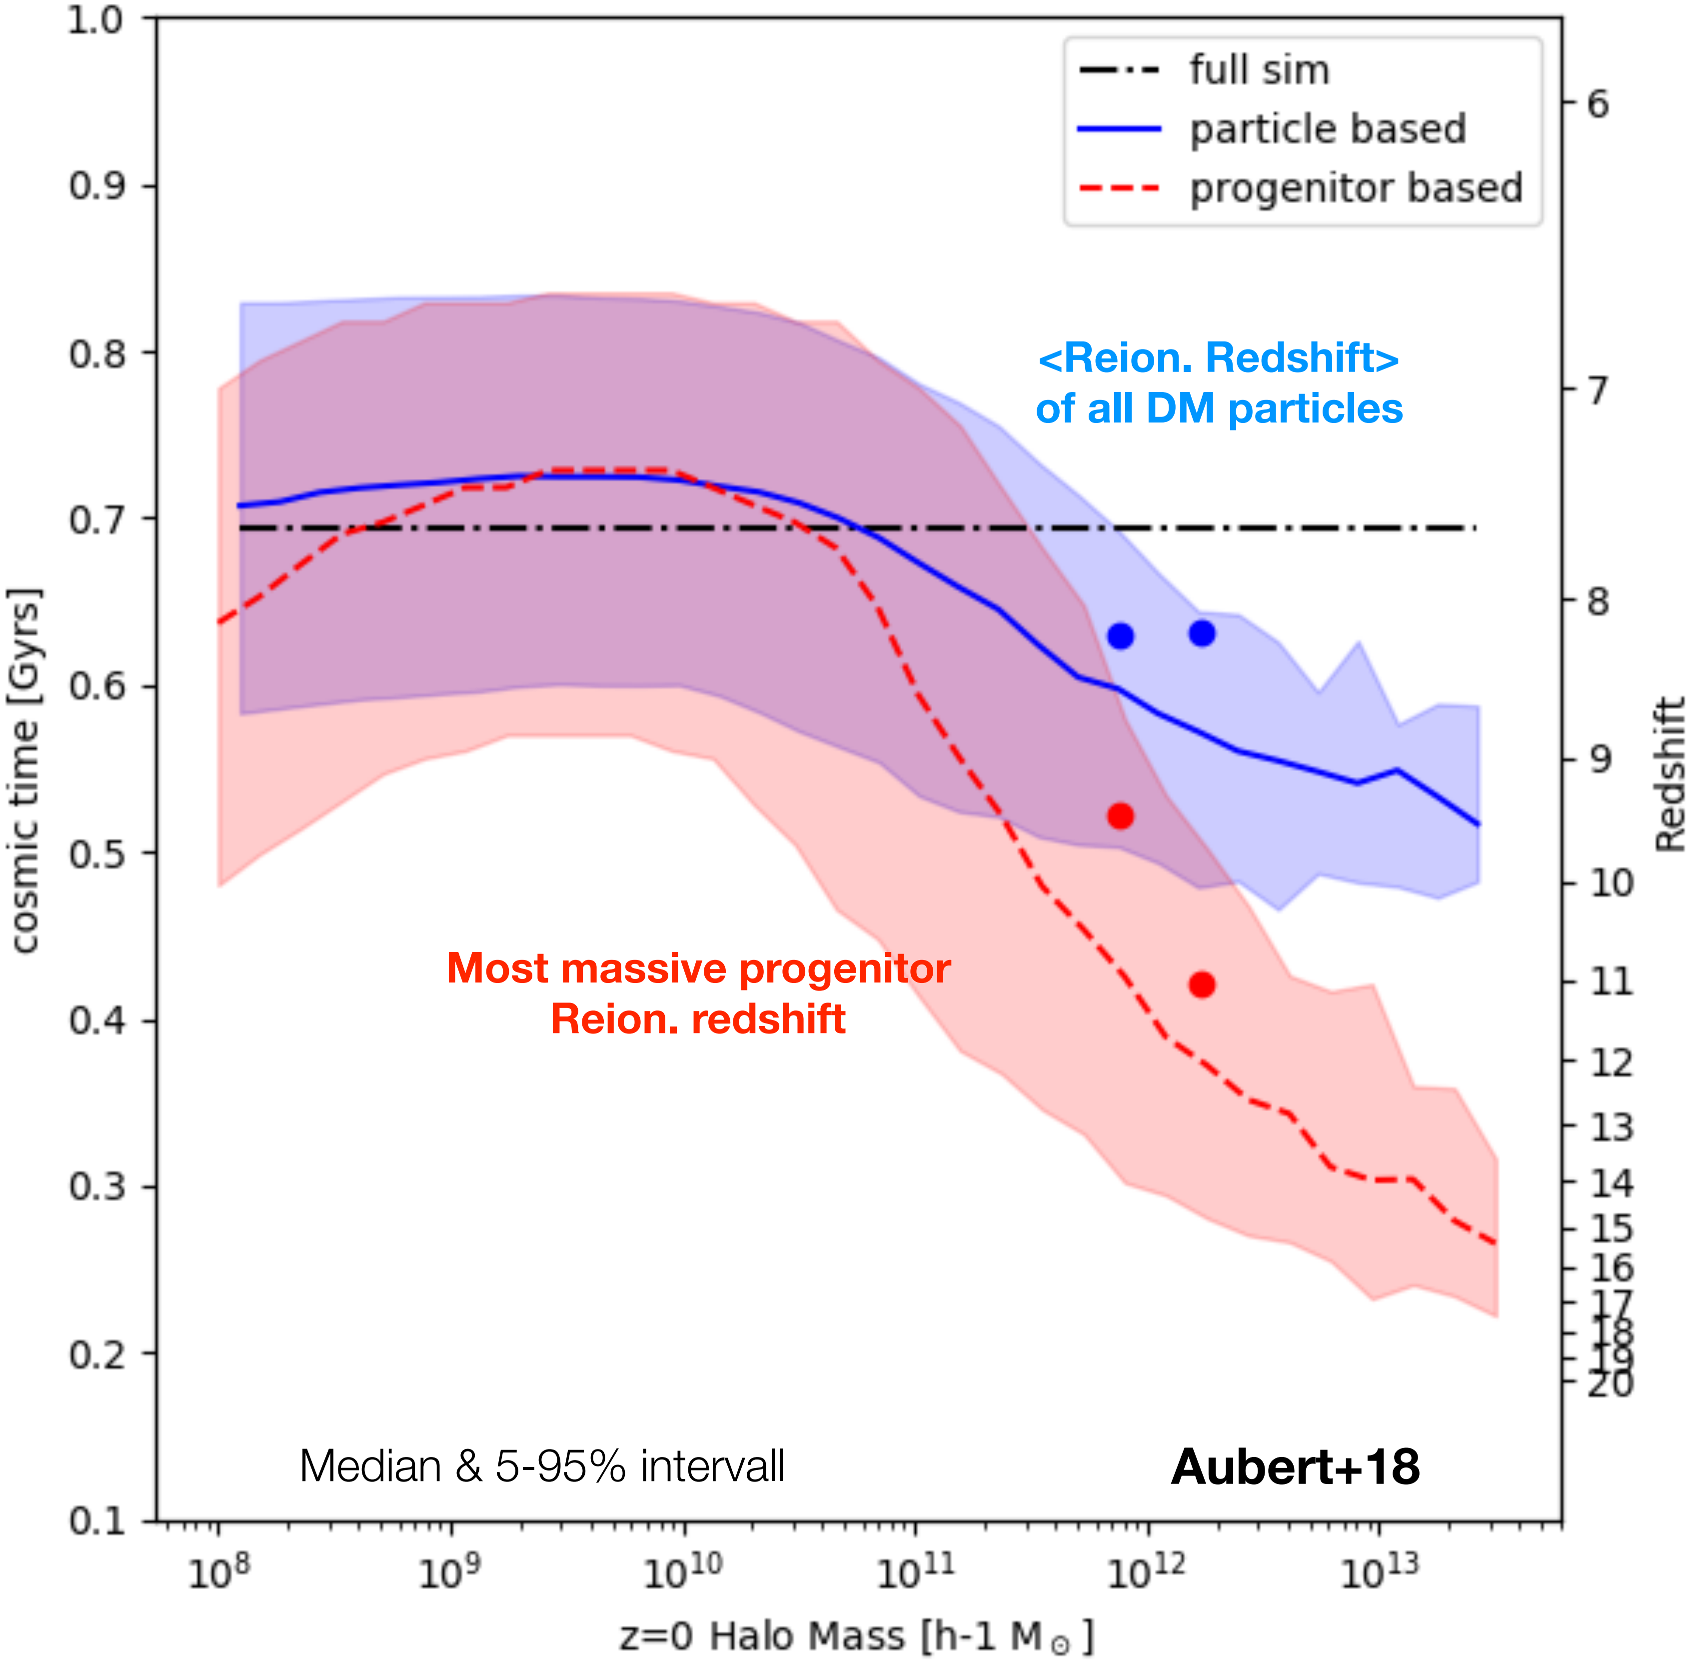

<Reion. Redshift> of all DM particles

Most massive progenitor Reion. redshift

Median & 5-95% intervall

Aubert+18

- Massive objects reionize early on
- First sources in most massive z=0 halos

- Light Objects reionize like the full volume
- The scatter is significant (~200 Myrs)

- Light Objects have median reionization times at later times than the full box (late reionization of faint objects ?)

- At low-mass, progenitor-based predictions are biased

- Particle-based predictions tend to predict later reionization times (sensitive to the full reionization history of the object, even diffuse matter)

The case of the Local Group (I)

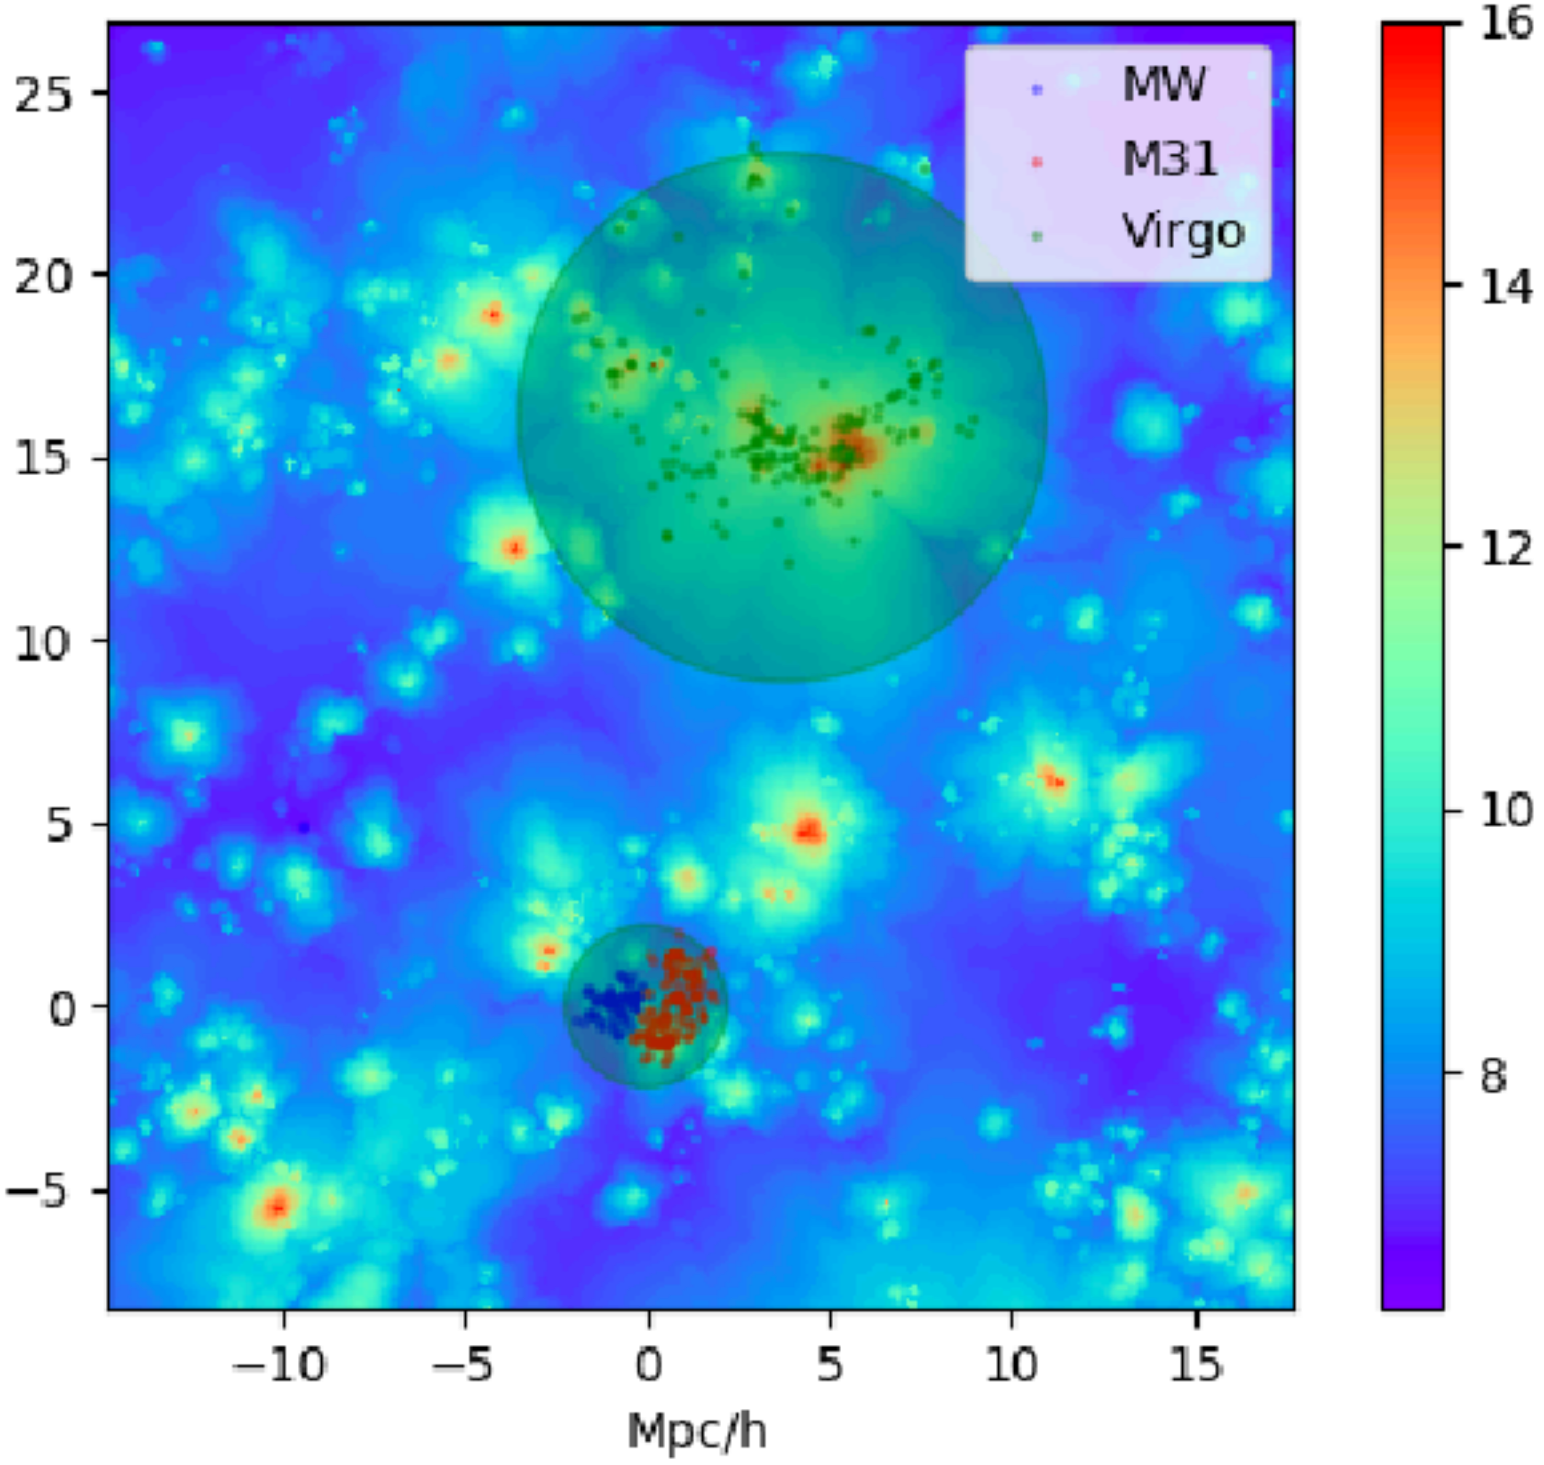

CLUES Initial Conditions
produces
a MW and M31 pair
in proper cluster environment
(Virgo & Fornax)

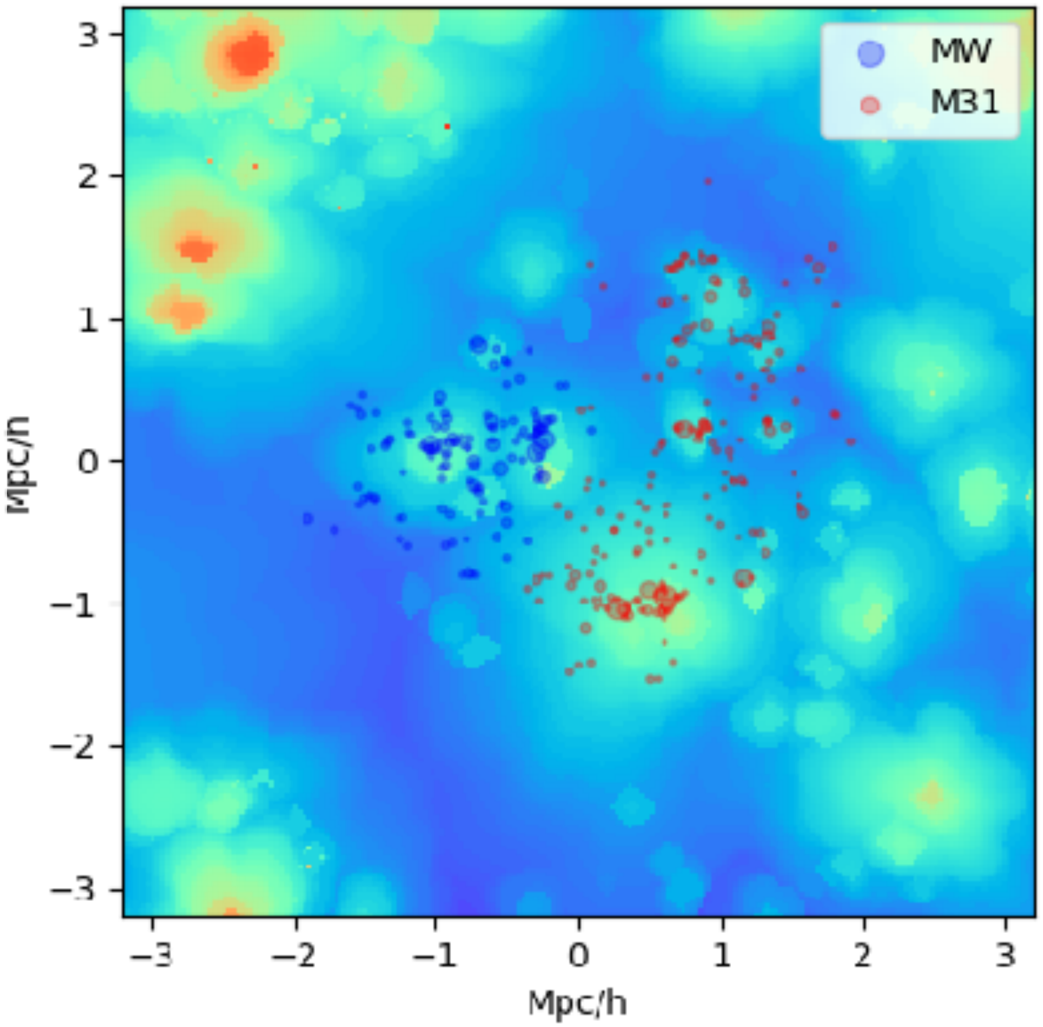

Ionization front speeds during reionization

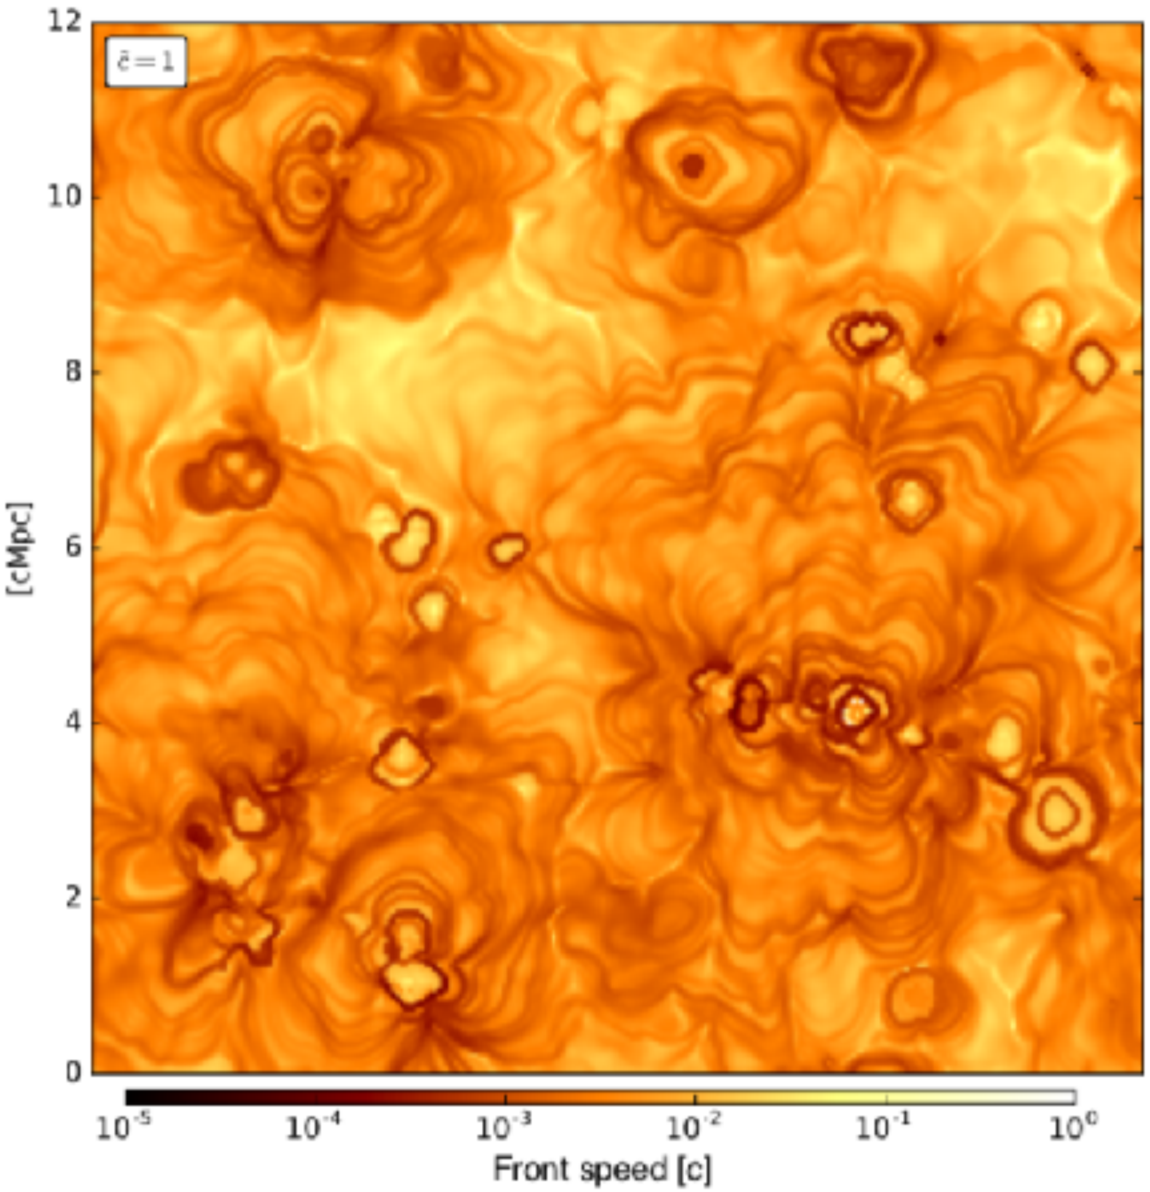

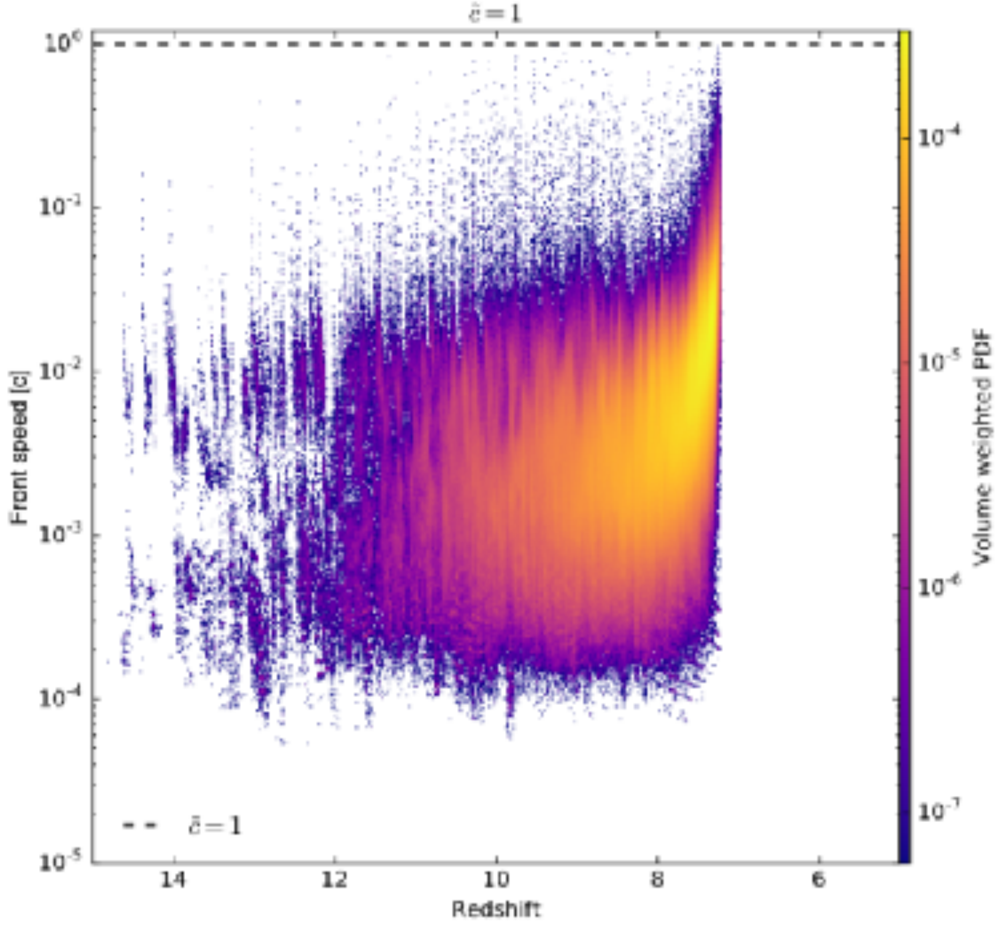

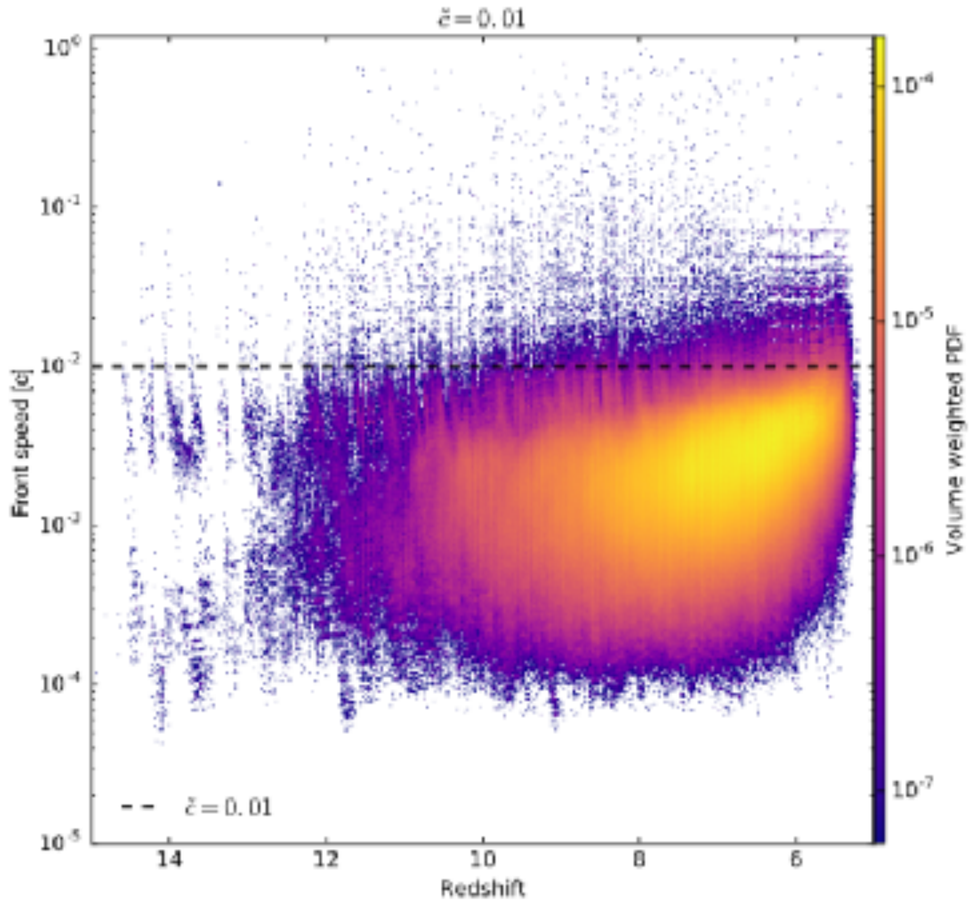

reduced speed-of-light impact front speeds during the overlap stage

Deparis+ 18

Cosmic variance of ionization front speeds and median reionization times

Deparis+ 18

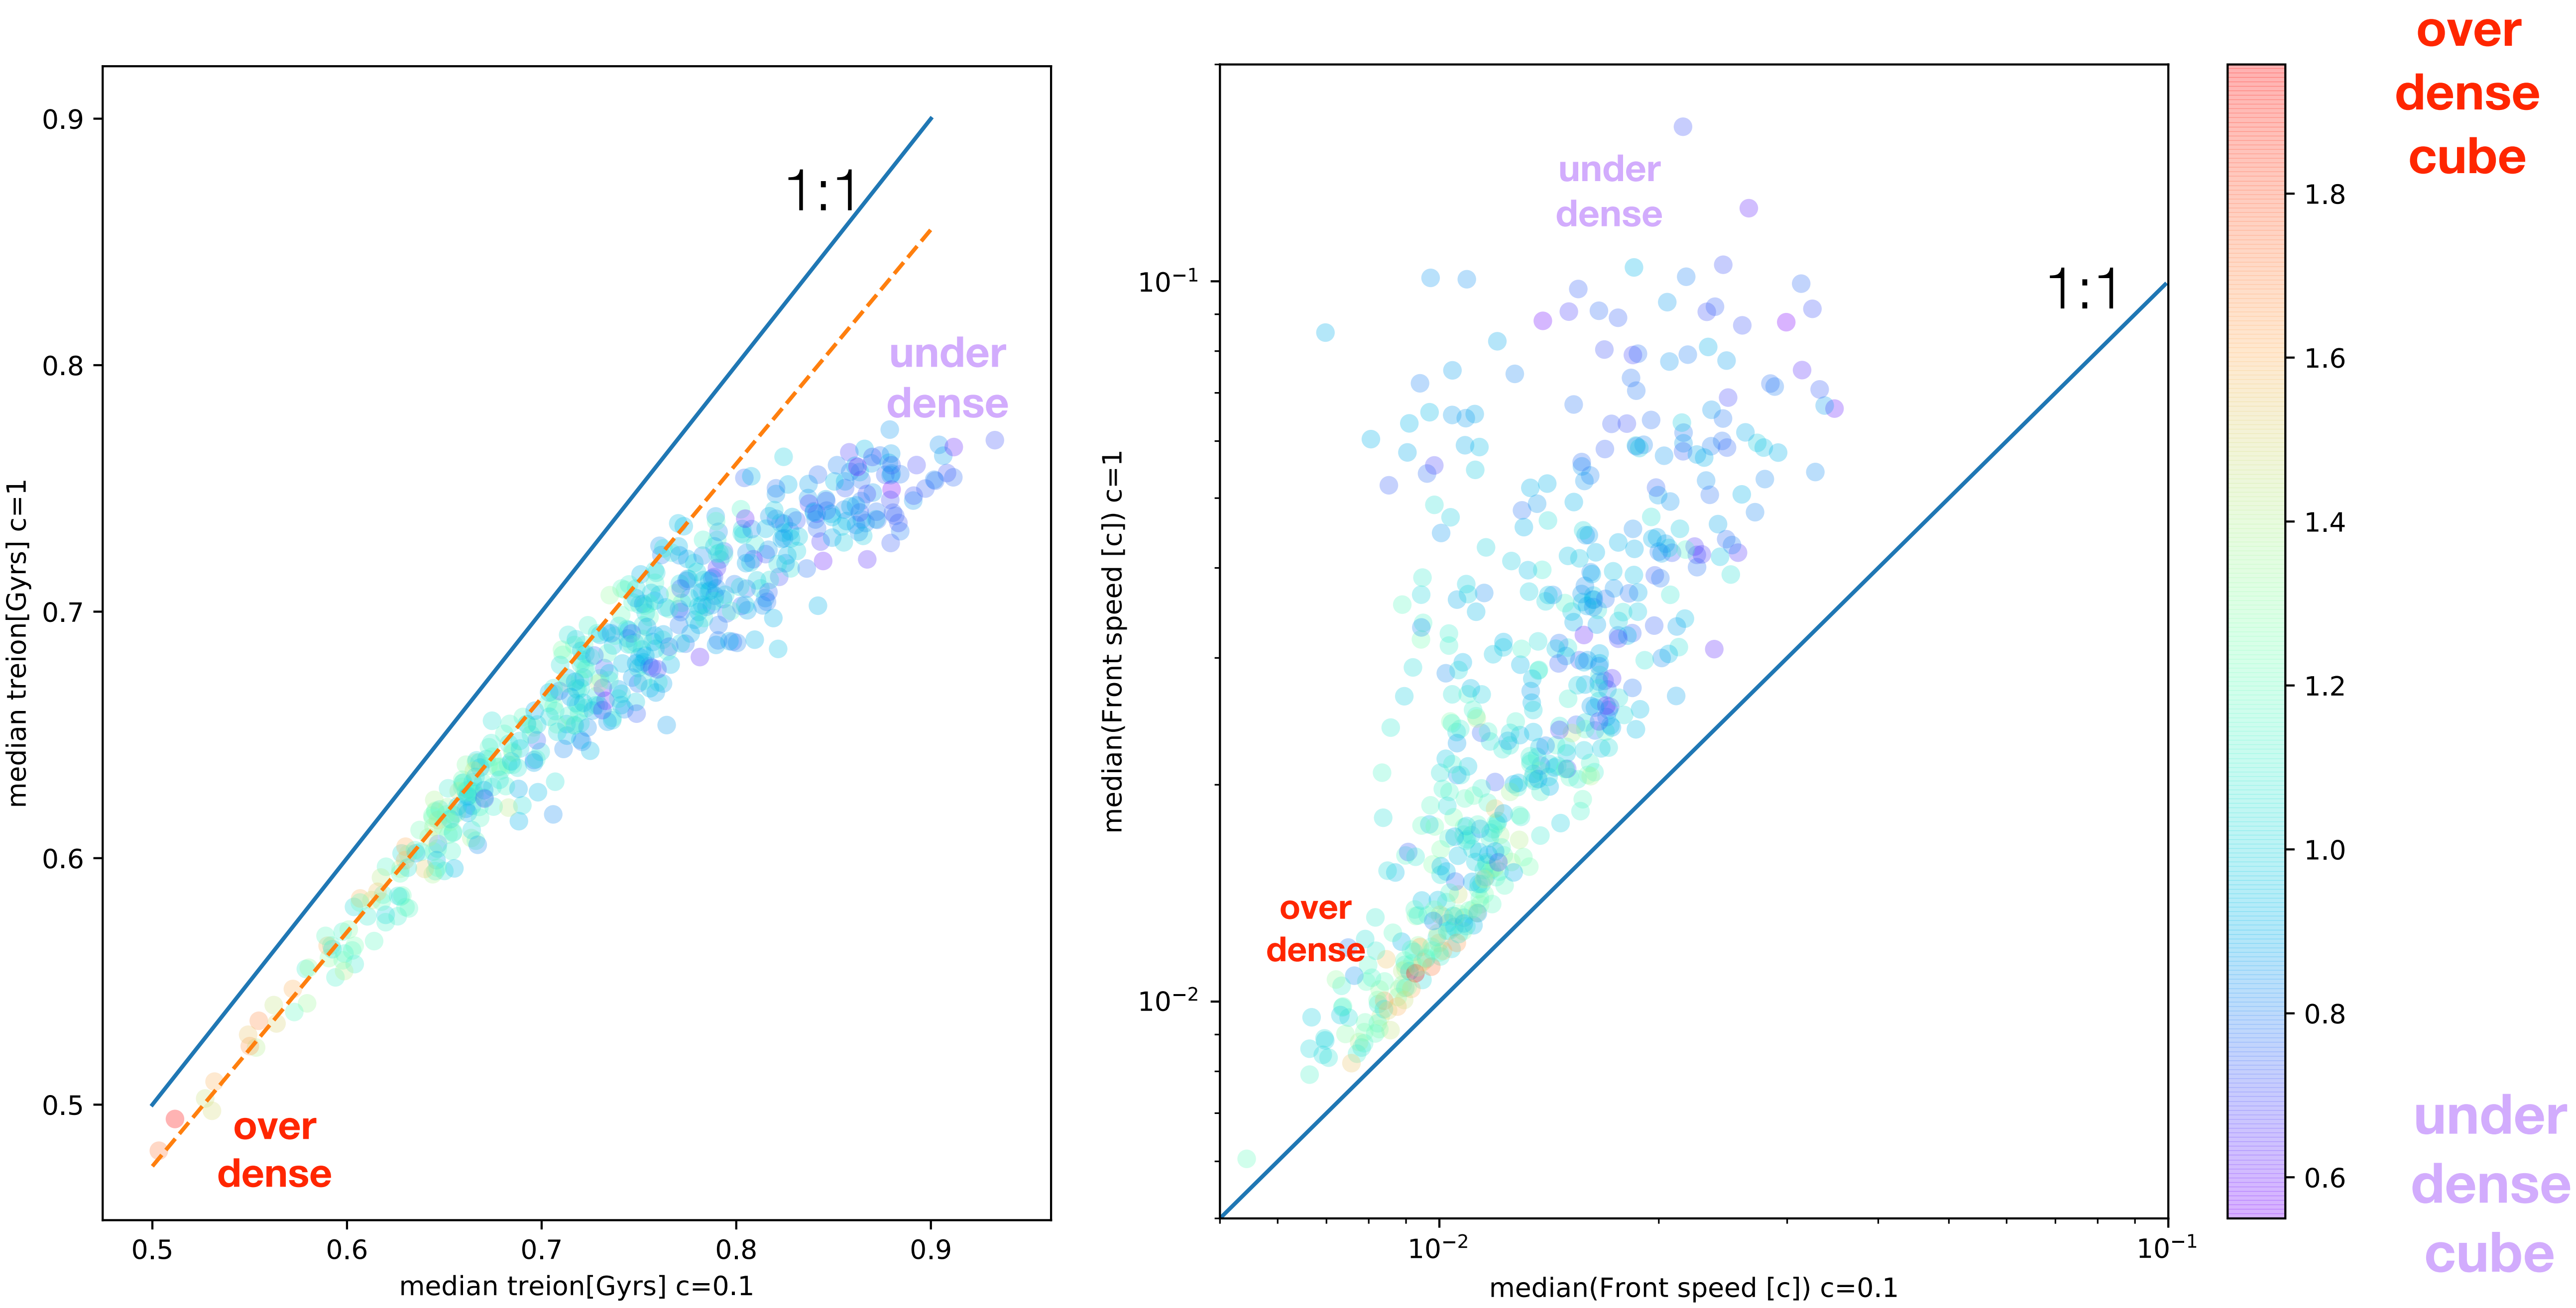

512 x 12 Mpc boxes
in a 91 Mpc large volume

choice of c results in
greater relative discrepancies in underdense regions

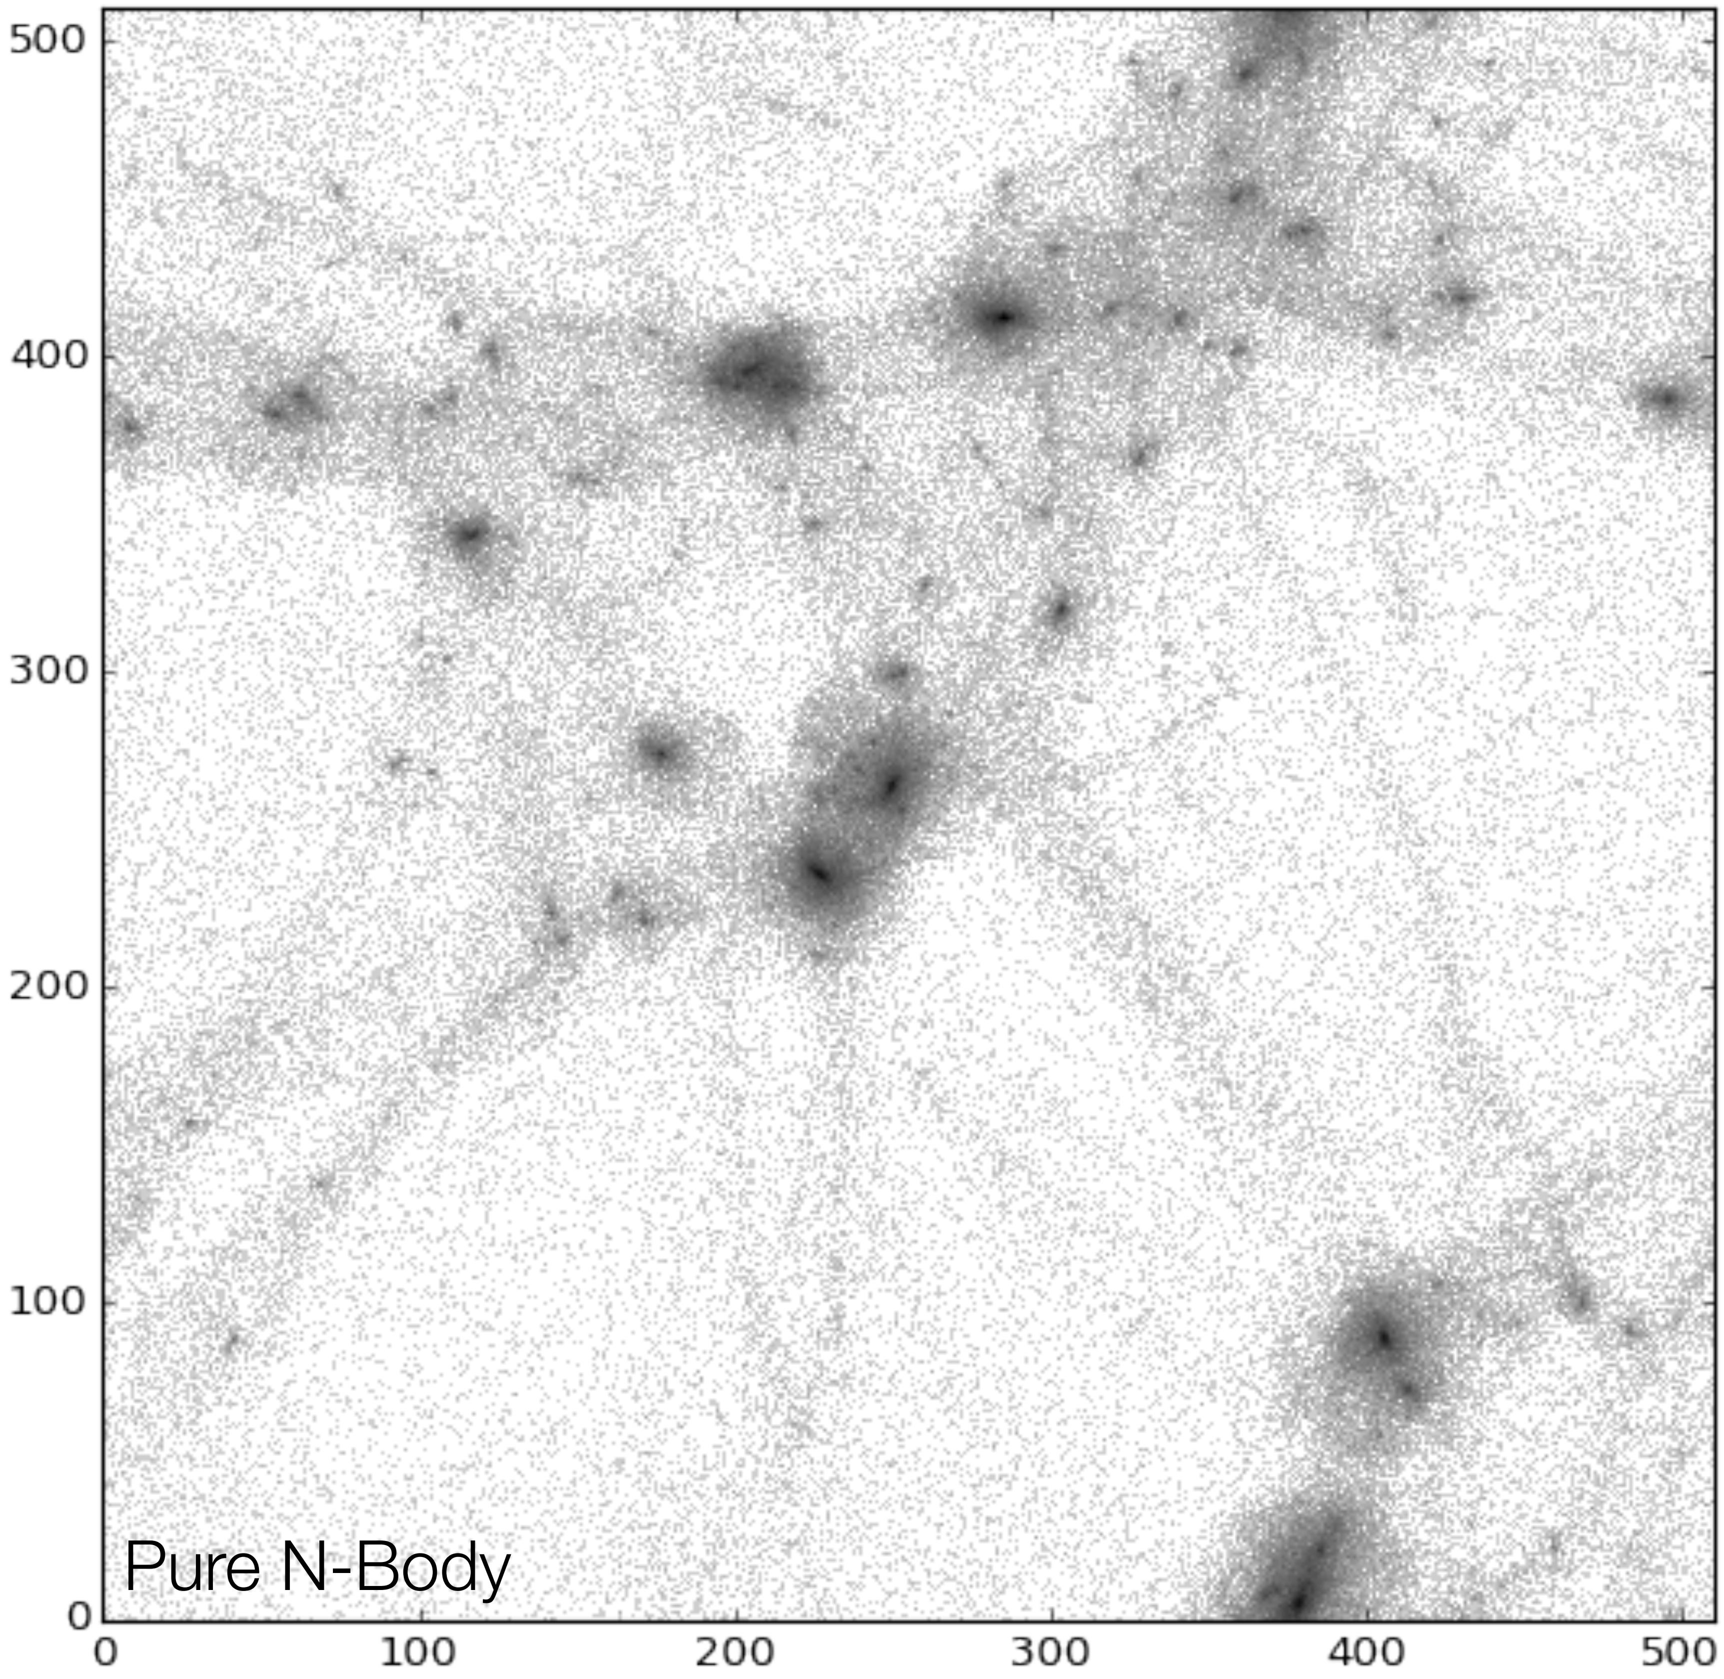

EMMA v2.0
codename ‘Kereon’

we trash the fully threaded AMR tree and replace it with Morton multi-level ordering

will we be able to take advantage of :

- shared GPU/CPU memory ?
- GPU/GPU comm ?
- Zerocopy/ Shared GPU/CPU memory ?
- tensor operations ?

For sure : a much much simpler code

What's next for GPU-driven simulations?

The elephant in the room :

Machine Learning

Machine learning may push hardware : we should be able to use the same hardware
(even intrinsics for e.g. tensor operations or hard-wired upscaling/downscaling)
= GPUs **during** simulation

Cognitive networks might actually replace physics modules
= GPUs **before** simulations
for greater efficiency
(and greater modularity ?)

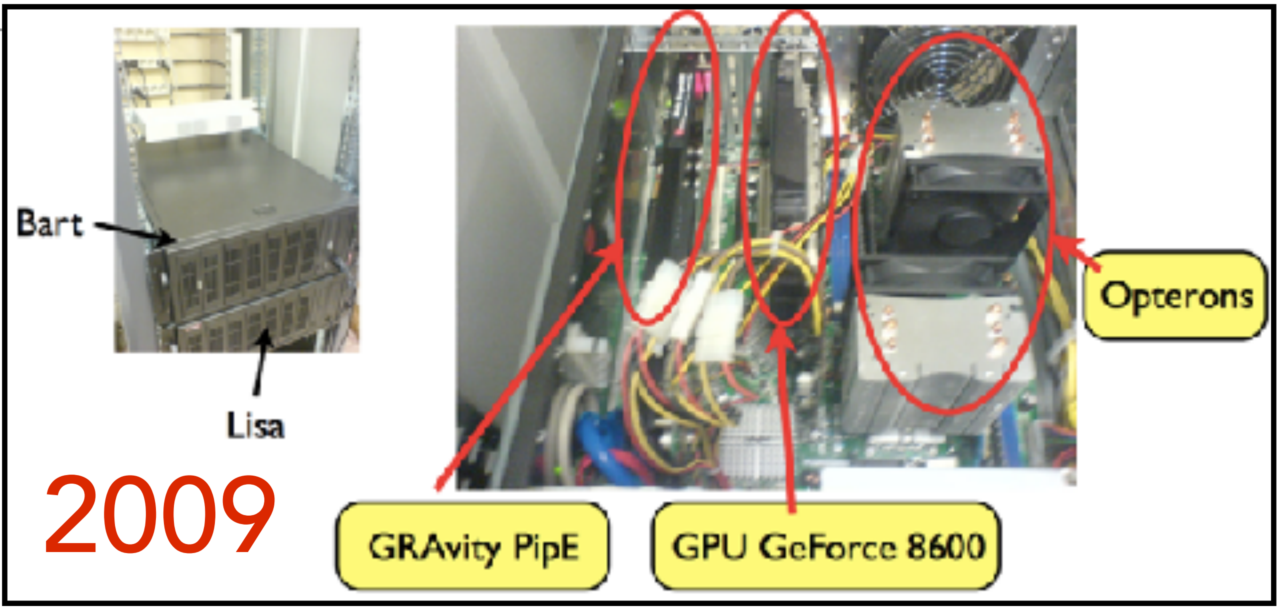

2009

2015

?

2020+