(Toward) Radiative transfer on AMR with GPUs

Dominique Aubert
Université de Strasbourg
Austin, TX, 14.12.12

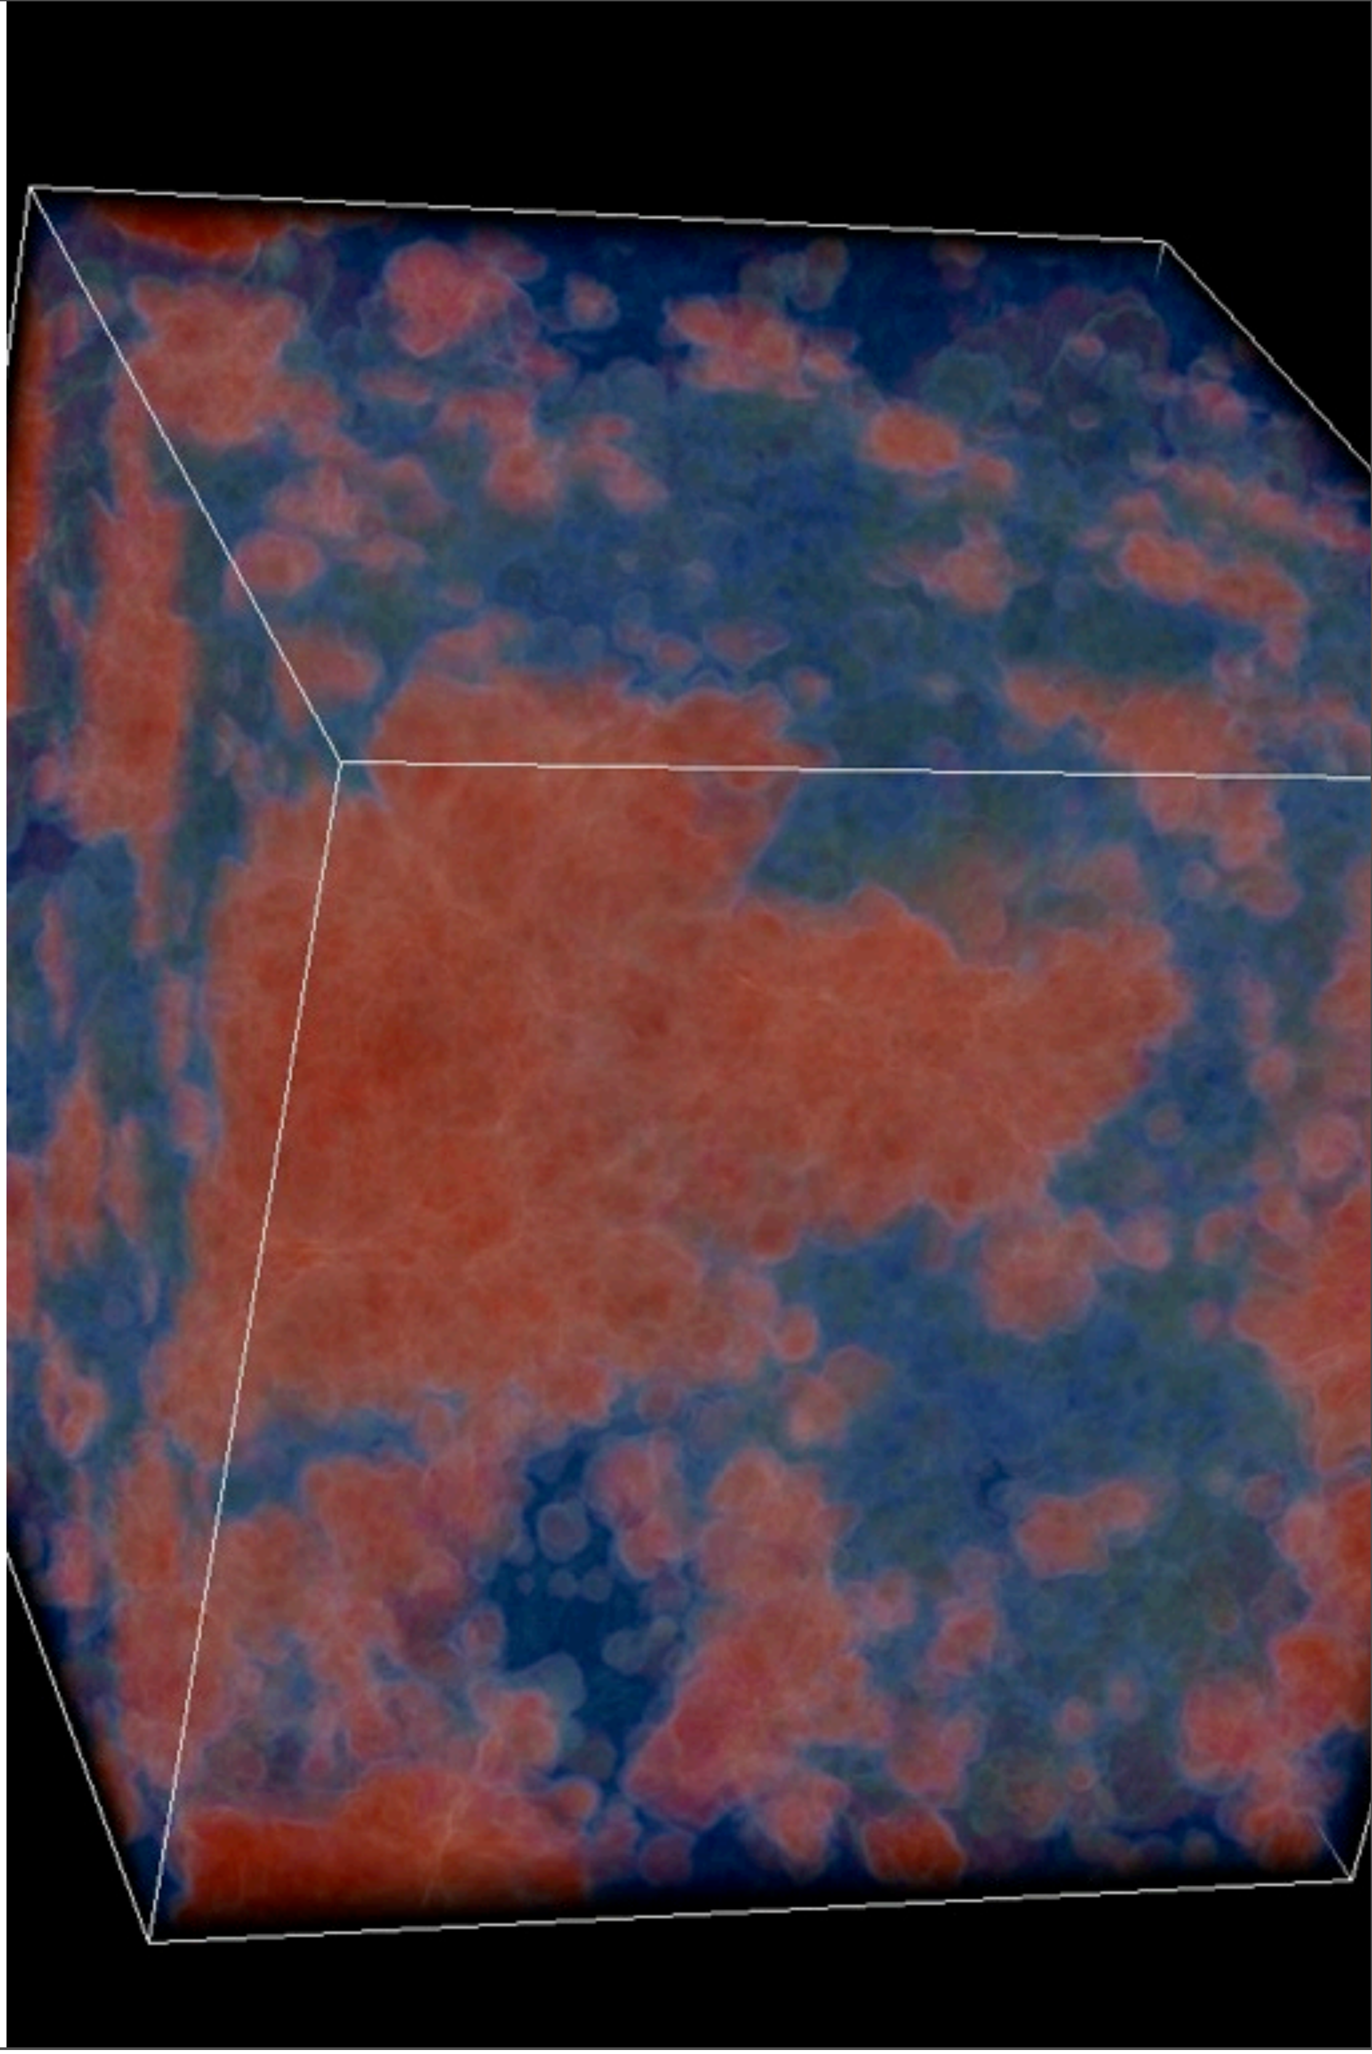

A few words about GPUs

- Cache and control replaced by calculation units
- Large number of Multiprocessors + Scheduler
- **High load + Independent + Non-Random Memory access**
- x10 to x100 compared to CPU
- High-level interface with **CUDA (C, Nvidia)**, OpenCL (Kronos)
- High-end GPUs ~1.5-2 kEuros, 4-6 GB RAM
- Tianhe (Changsha, 7168 GPUs), Titan (Oak Ridge 7000-18 000 GPUs by 2012), MareNostrum (???), in France: Titane (198 GPUs), Curie (268 GPUs)

Principle of GPU programming with CUDA

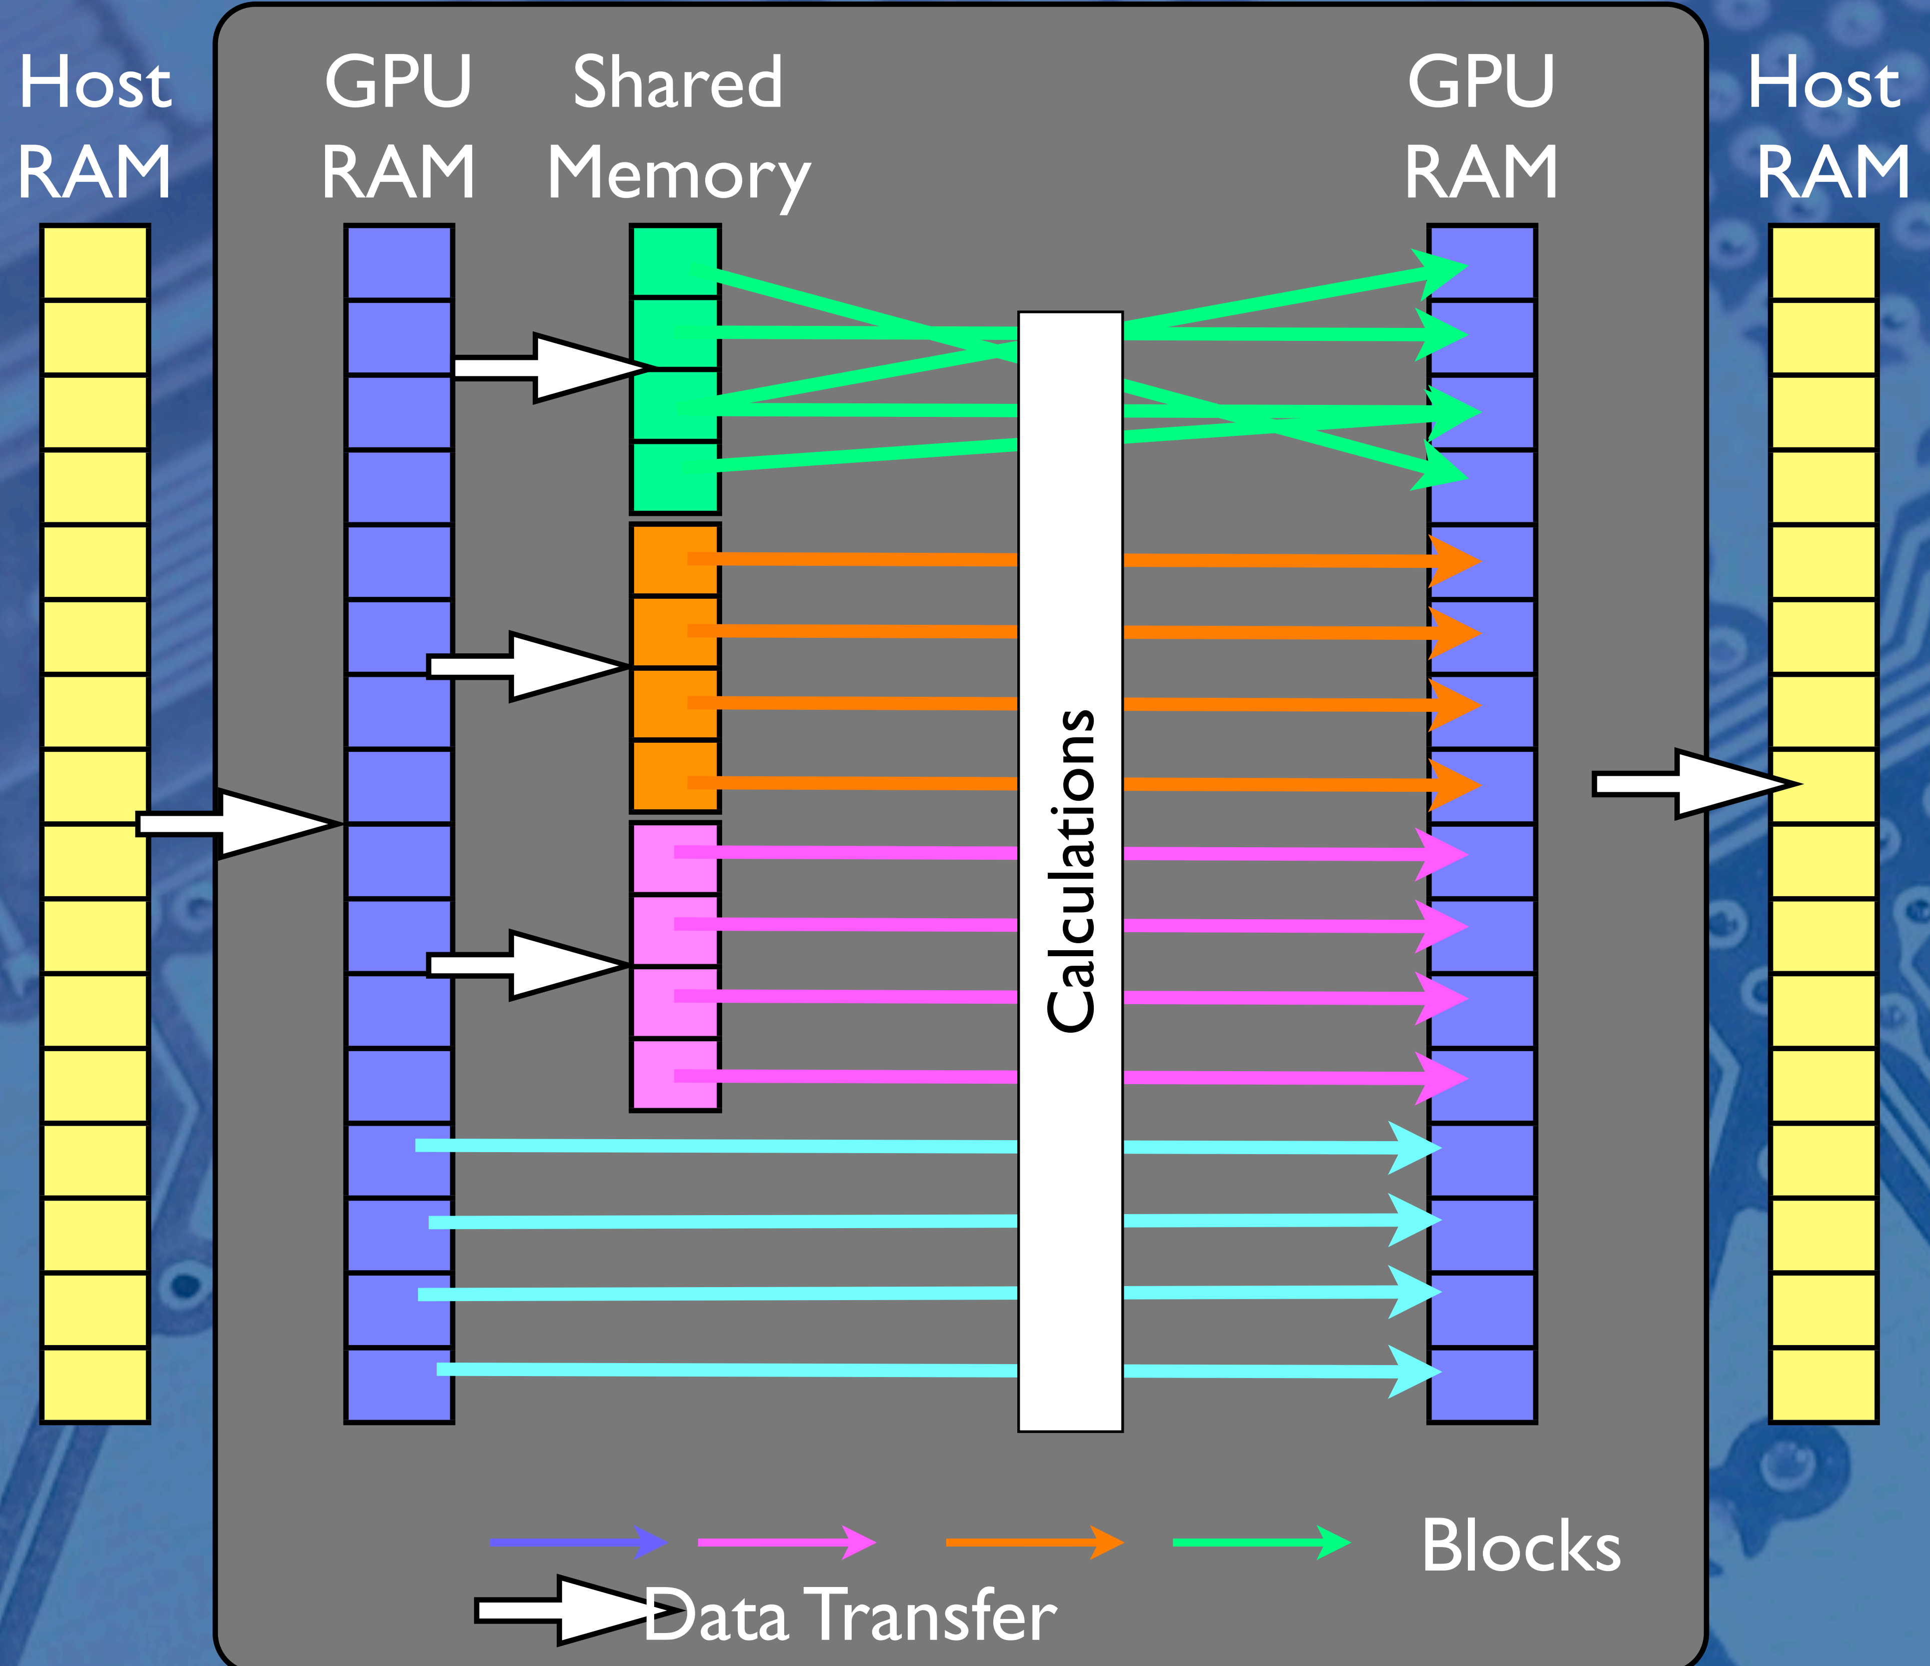

__If possible:__
independent &
identical
threads
+
High arithmetic
intensity
=
acceleration

1. Cosmological Radiative Transfer

Radiative Transfer equations : explicit solver

$$\frac{\partial E_\nu}{\partial t} + \nabla \mathbf{F}_\nu = -\kappa_\nu c E_\nu + S_\nu$$

First 2 moments of the RT equations + variable Eddington Tensor with M1 closure relation

Gonzales et al. 2008, Aubert & Teysier 2008, Rosdahl & Blaizot 2012

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}(\mathbf{U})}{\partial x} = S \Rightarrow \frac{\mathbf{U}^{\mathbf{p}+1} - \mathbf{U}^{\mathbf{p}}}{\Delta t} + \frac{\partial \mathbf{F}(\mathbf{U}^{\mathbf{p}})}{\partial x} = S$$

Explicit: CFL constrains

100 000 timesteps required to cover the reionization (z~5)

$$c < \frac{\Delta x}{\Delta t} \Rightarrow \Delta t < \frac{\Delta x}{c}$$

with GPUs it's ok @ c=300 000 km/s

Aubert & Teyssier, 08,10

Post-Processed Radiative Transfer with ATON

gas density +sources

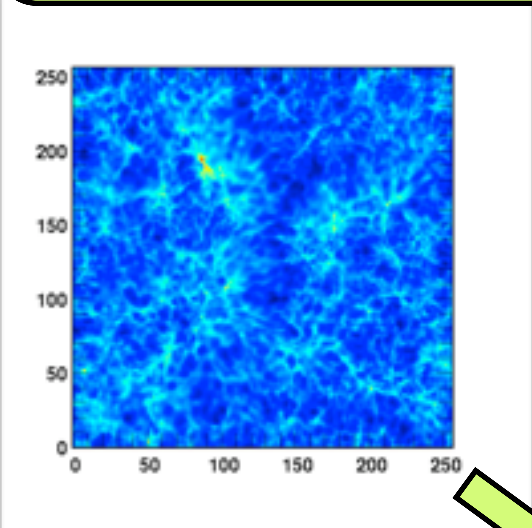

radiative energy

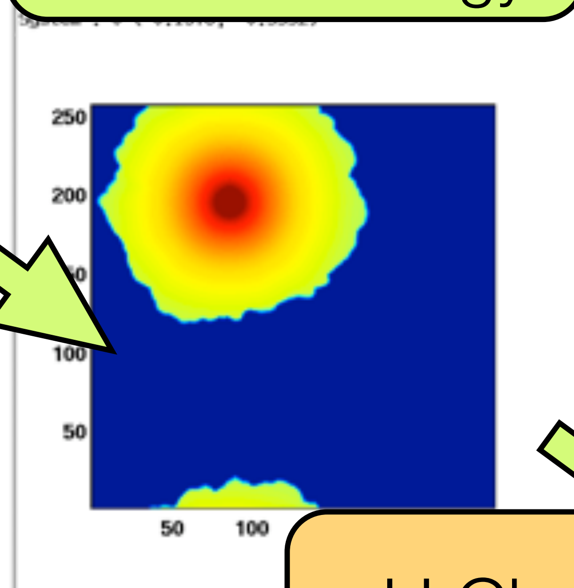

UV+X rad. transport

<u>Subcycled physics</u>
(almost) fixed
number of operations

Independent and high load

Ionisation state

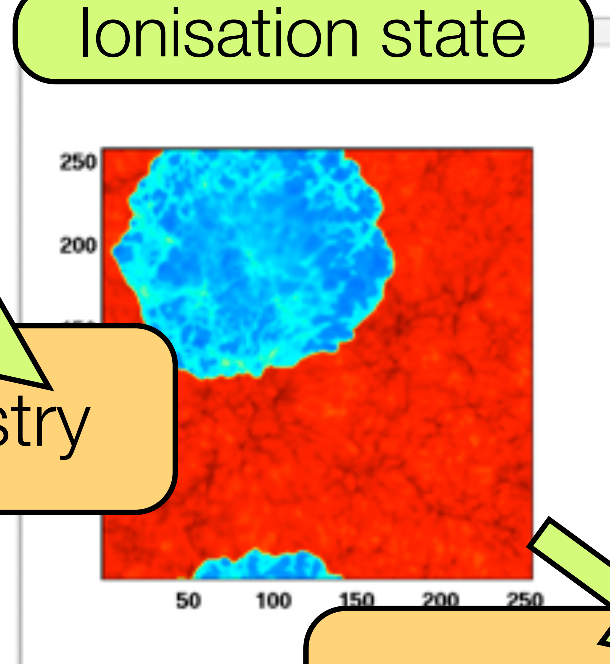

H Chemistry

Temperature

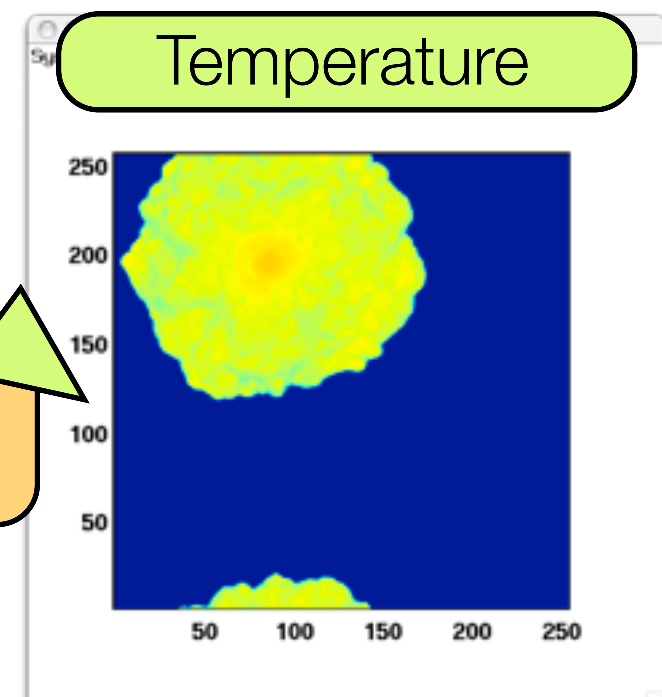

heating

<u>Conservative transport</u>
fixed & predictable number of operations

Regular 3D Grid

Independent and contiguous
calculations

Performances
GPUs VS CPUs

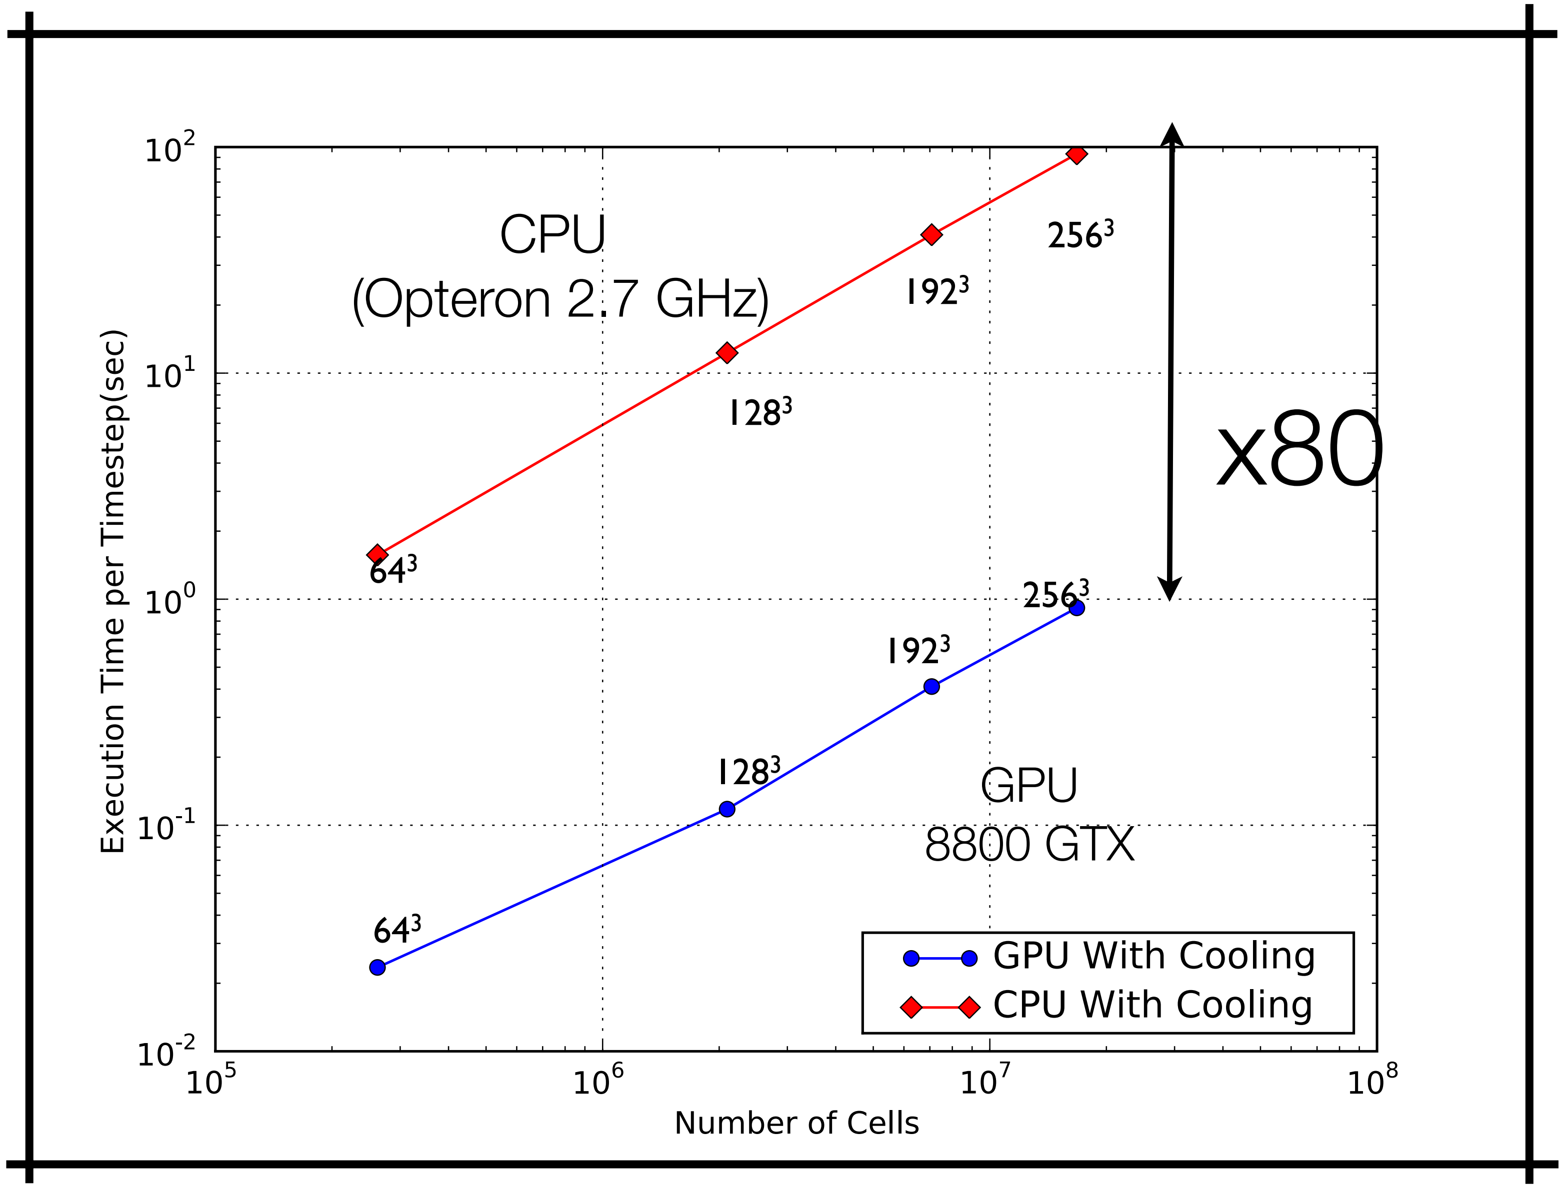

Aubert & Teyssier, 08,10

Multi-GPU with boundary layers

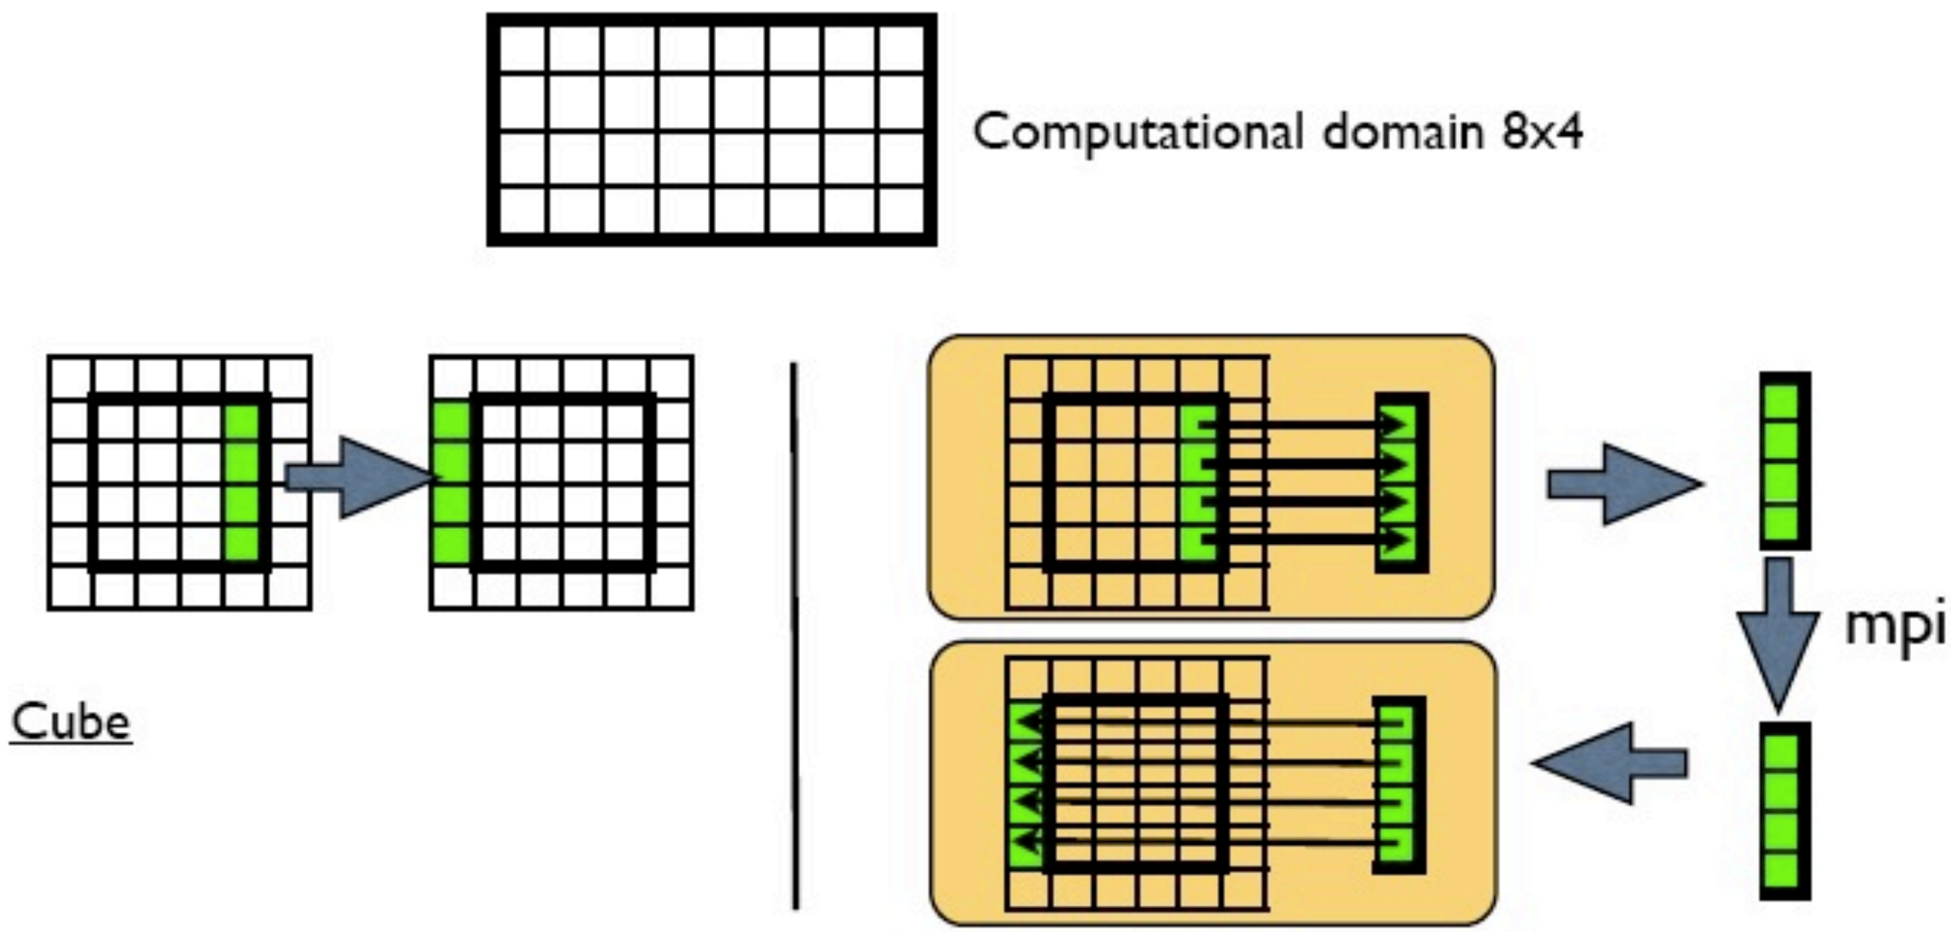

Aubert & Teyssier, 08,10

Applications :TRASH Project (**T**ransfert **RA**diatif **S**ur **H**ydrodynamique)

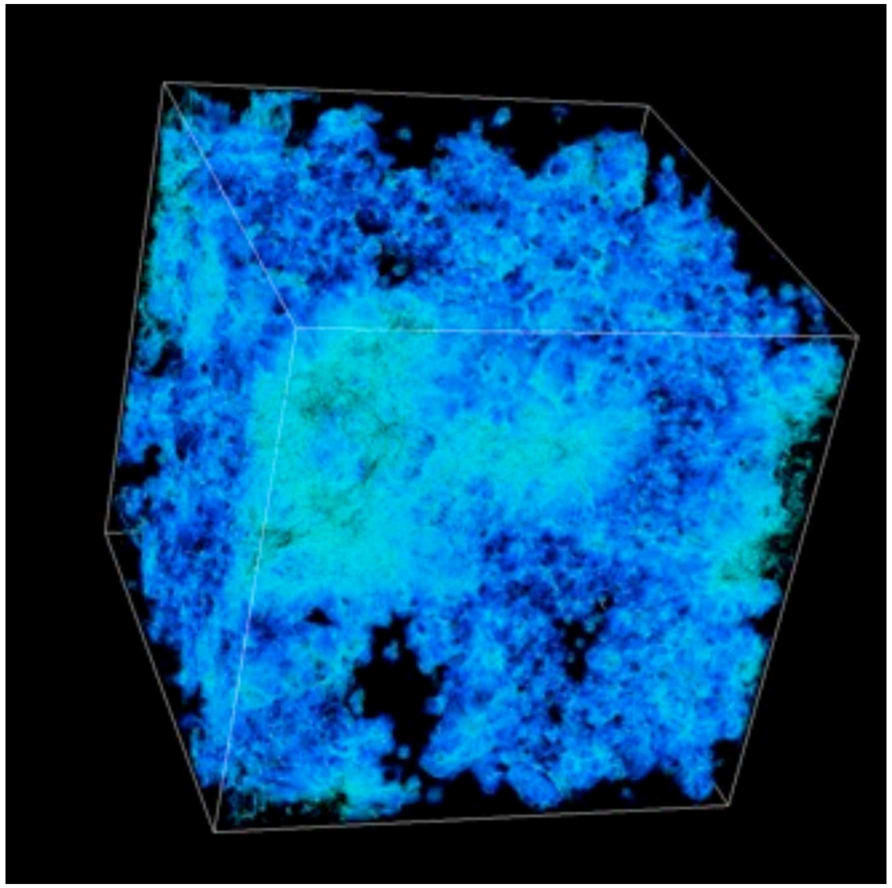

Gas and source distribution from the Mare Nostrum Hydro simulation
1024x1024x1024 cells + 2 refinement levels

Self-consistent stellar particles used as sources

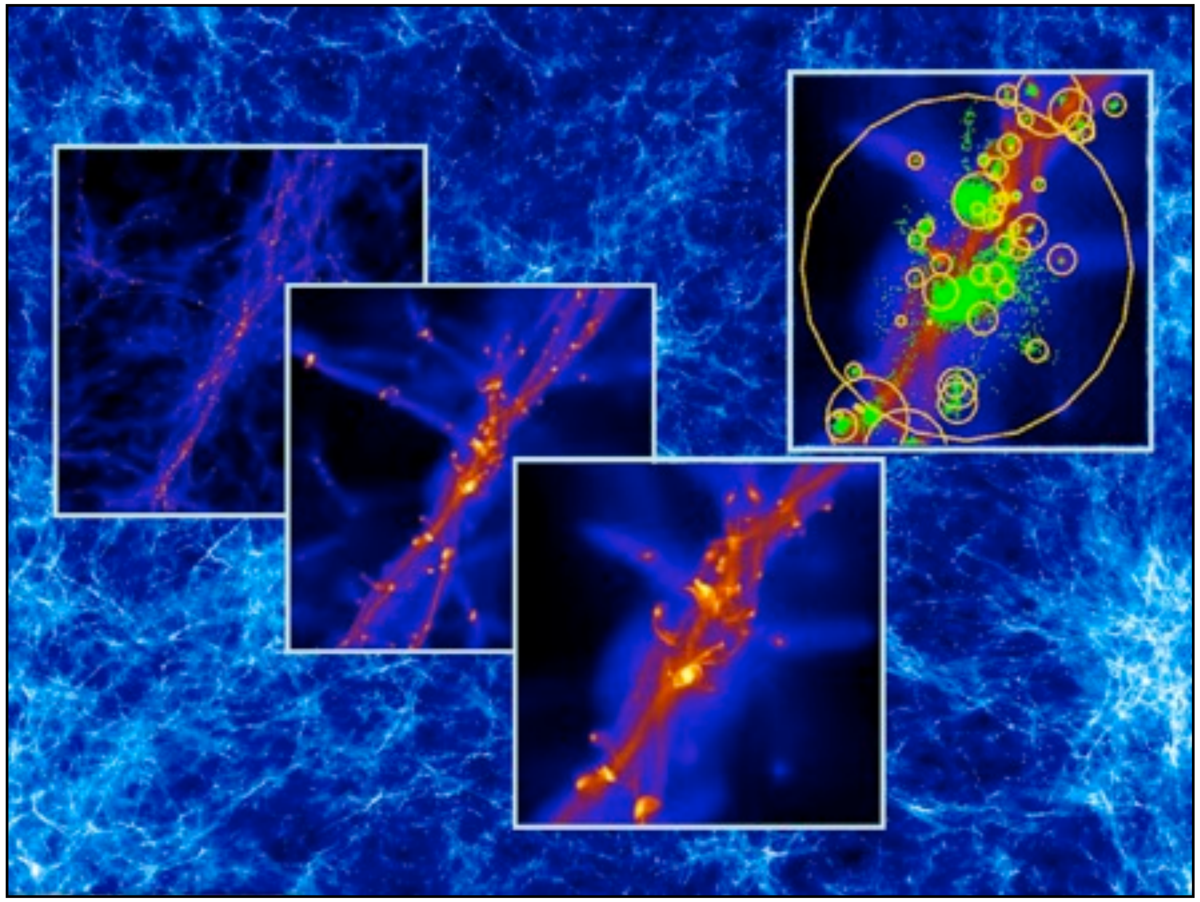

cudATON on TITANE-CCRT:
1024^3 grid
Cartesian domain decomposition
8x8x2
(128 GPUs - S1070 servers- Infiniband DDR)

~60 000 - 180 000 time steps
dt ~10 000 yrs over 1 Gyrs

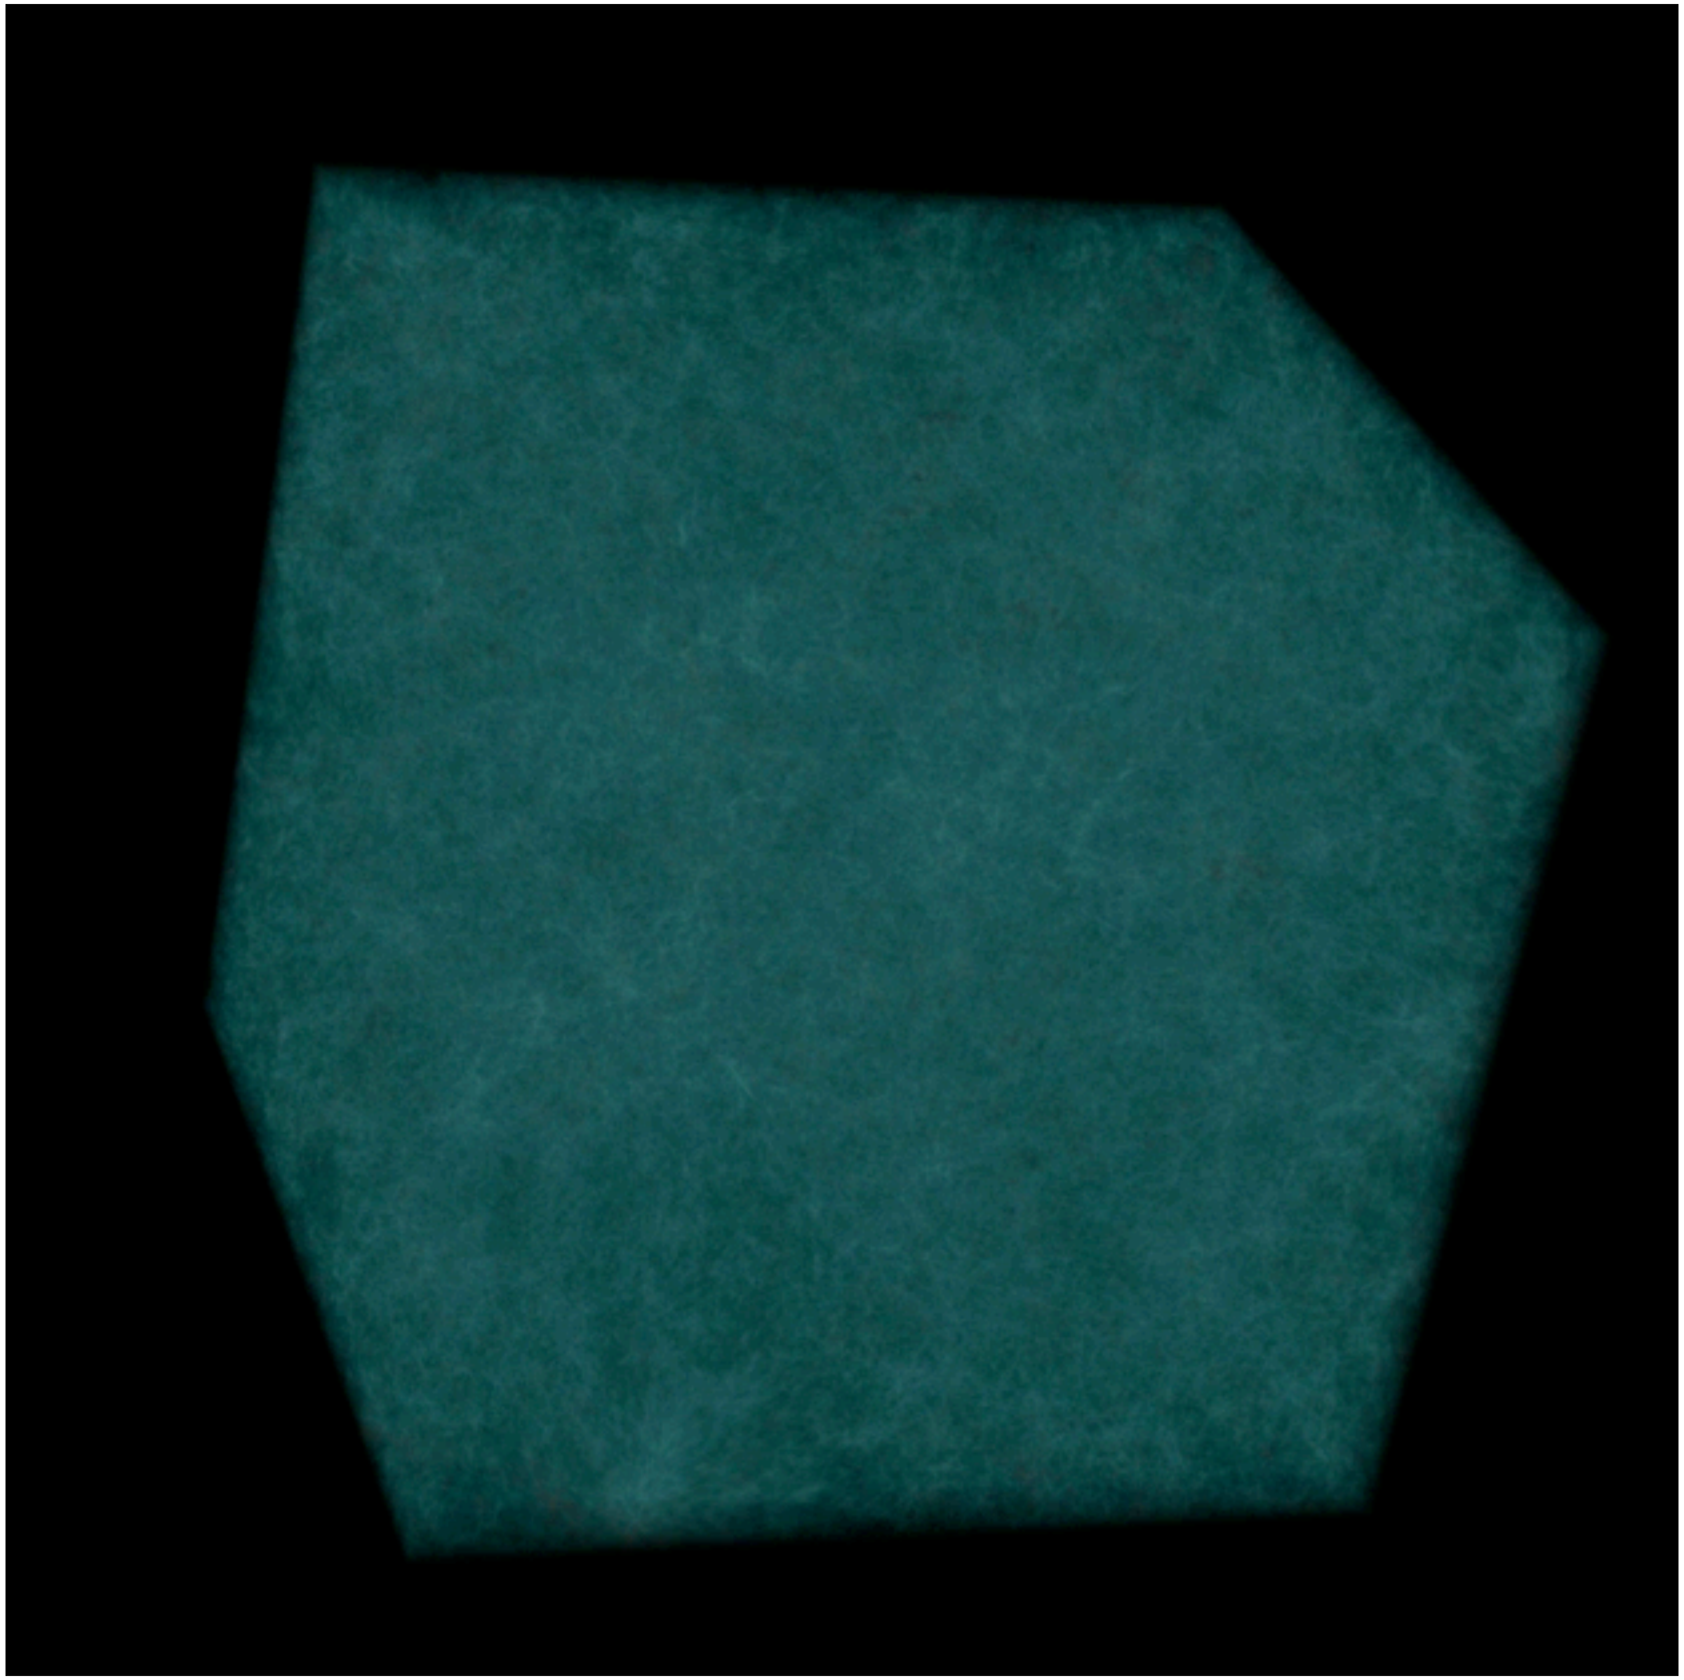

cudATON on TITANE-CCRT:
1024^3 grid
Cartesian domain decomposition
8x8x2
(128 GPUs - S1070 servers- Infiniband DDR)

~60 000 - 180 000 time steps
dt ~10 000 yrs over 1 Gyrs

Aubert & Teyssier 2010
Structure of the UV background
@ different resolution and sub-grid models

Aubert & Teyssier, ApJ, 2010

Timings on Titane

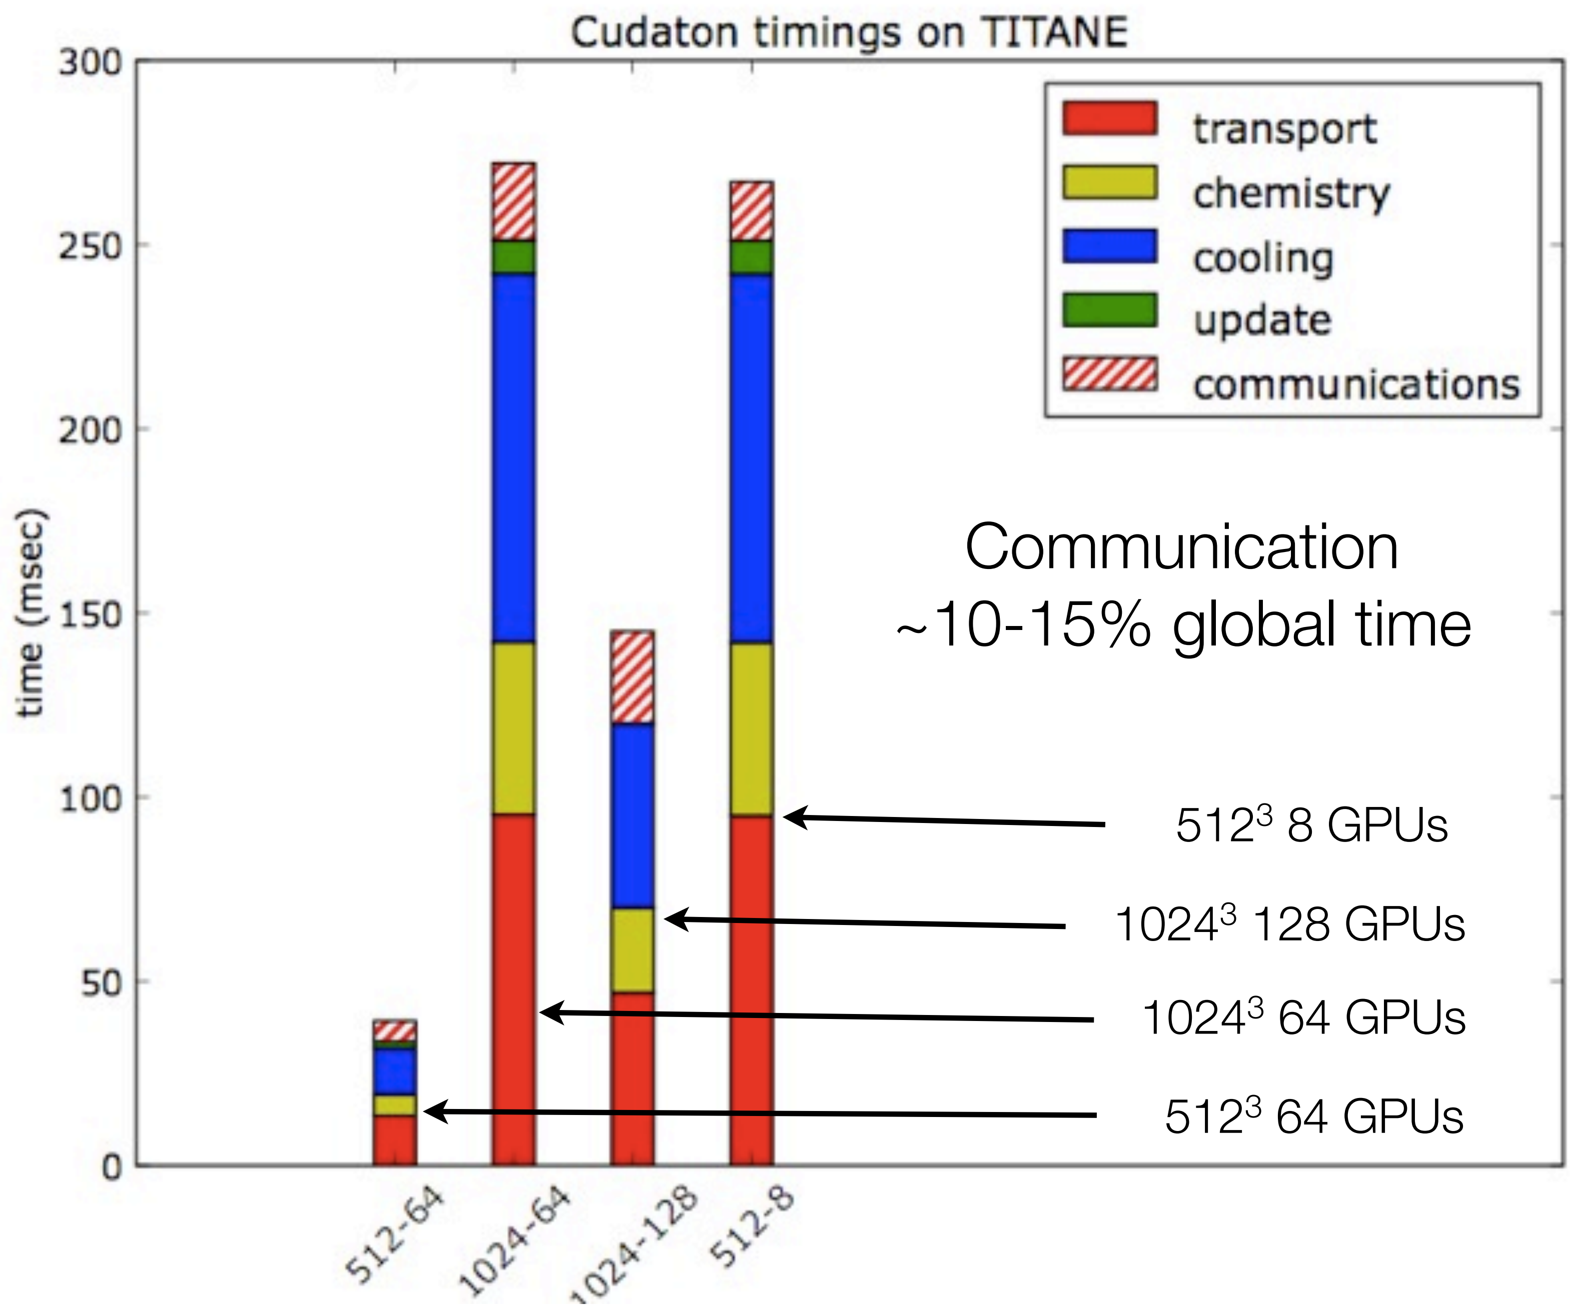

Aubert & Teyssier, ApJ, 2010

Small scale effects

100 Mpc/h -1024^3 box
clumping C(delta) extracted from a 12.5 h/Mpc -1024^3

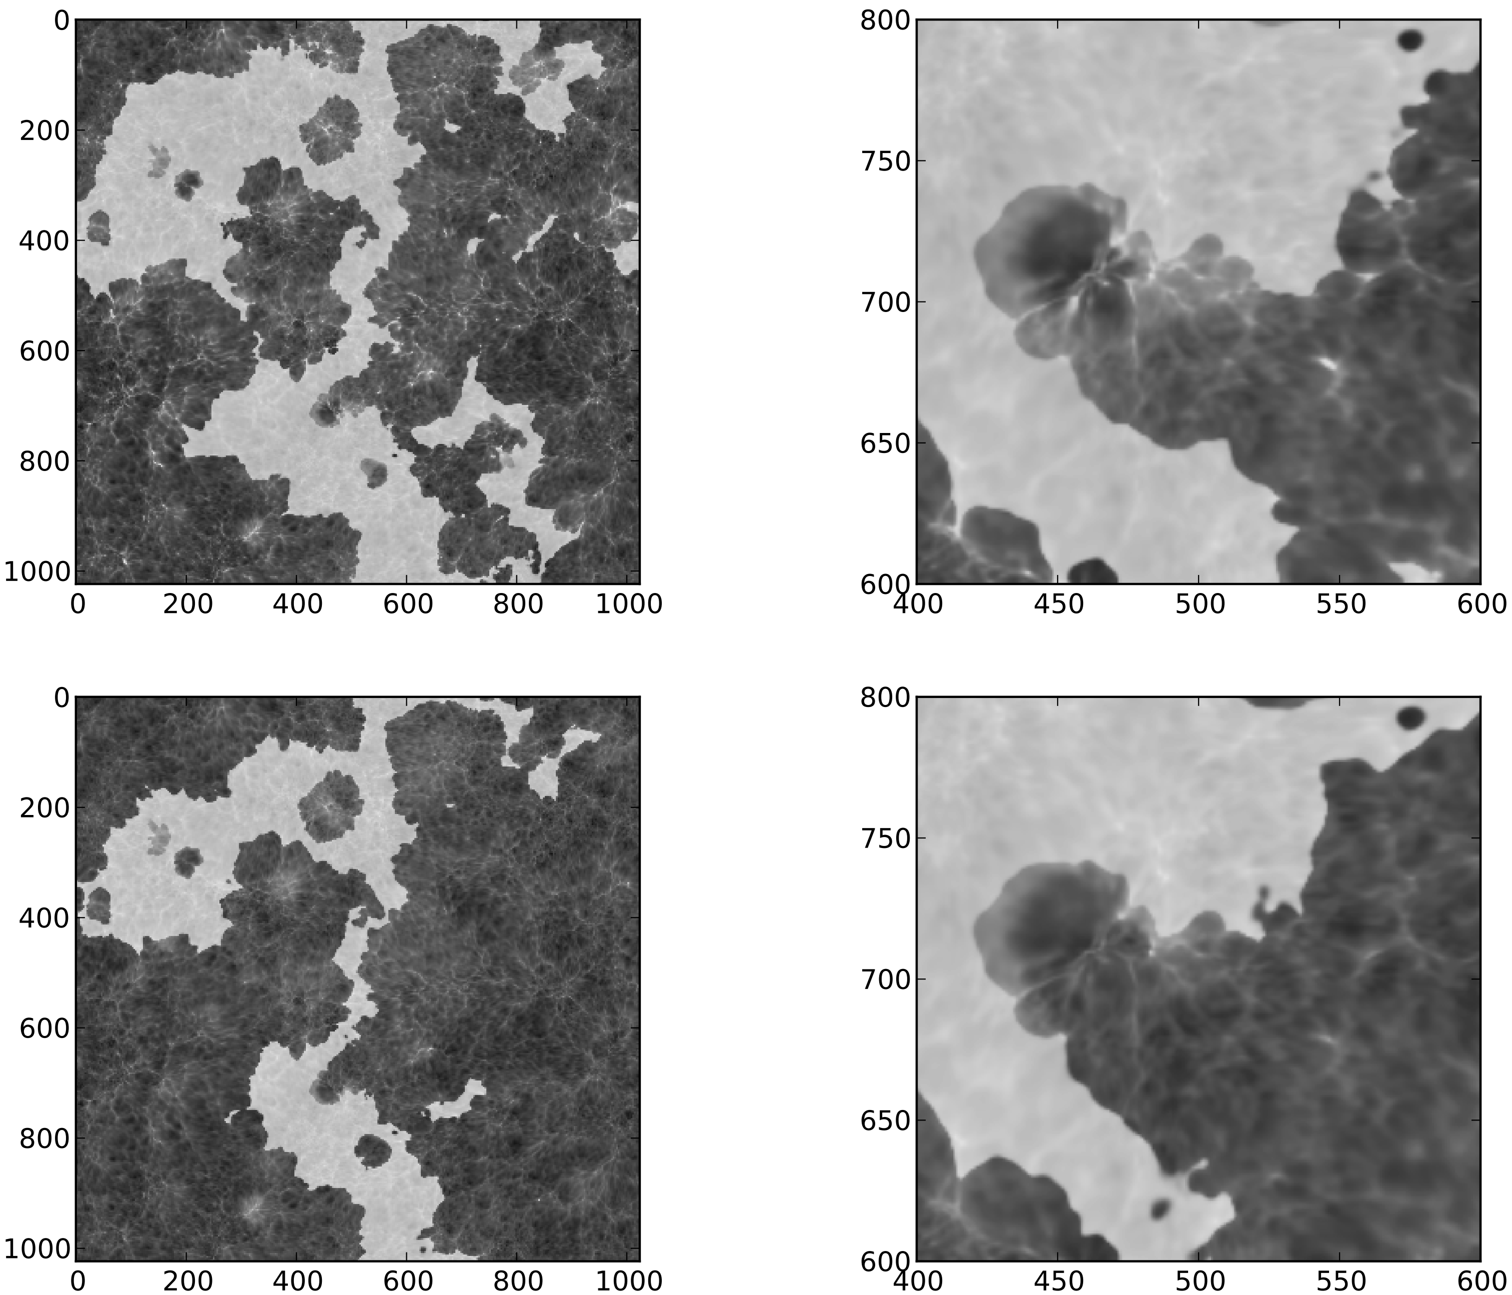

with subgrid clumping

without subgrid clumping

Aubert & Teyssier, ApJ, 2010

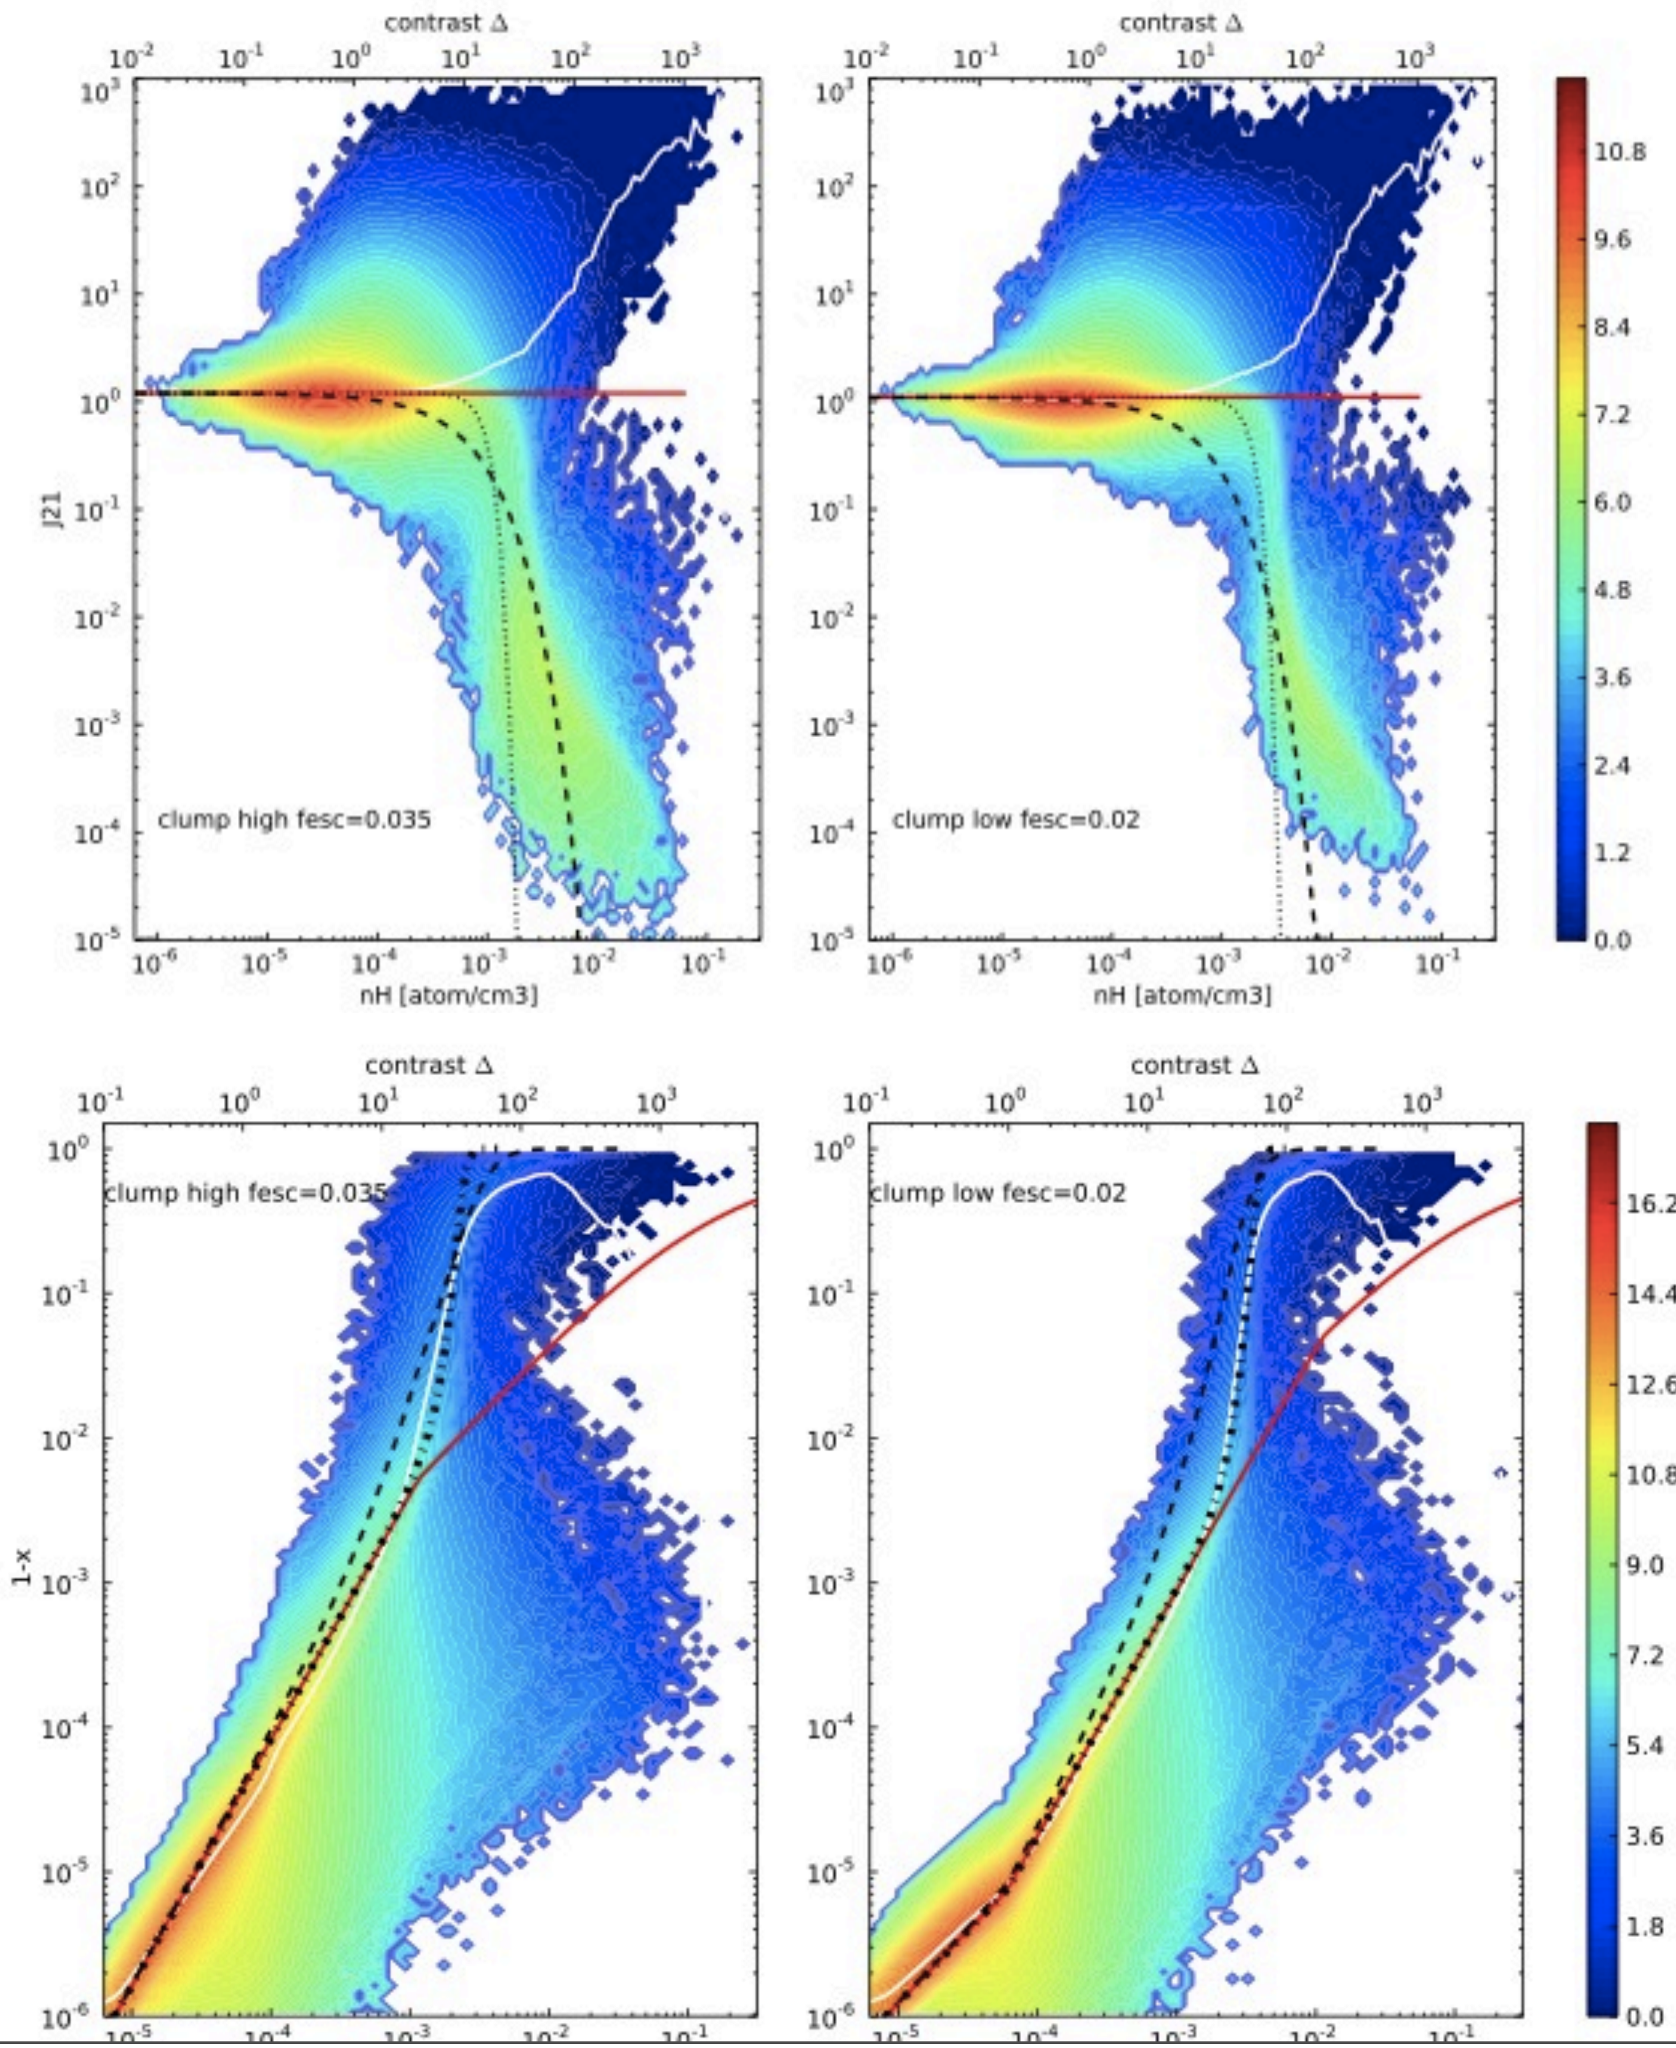

J21 Vs nH

x Vs nH

Aubert & Teyssier, ApJ, 2010

Residual Neutral Fraction and J21

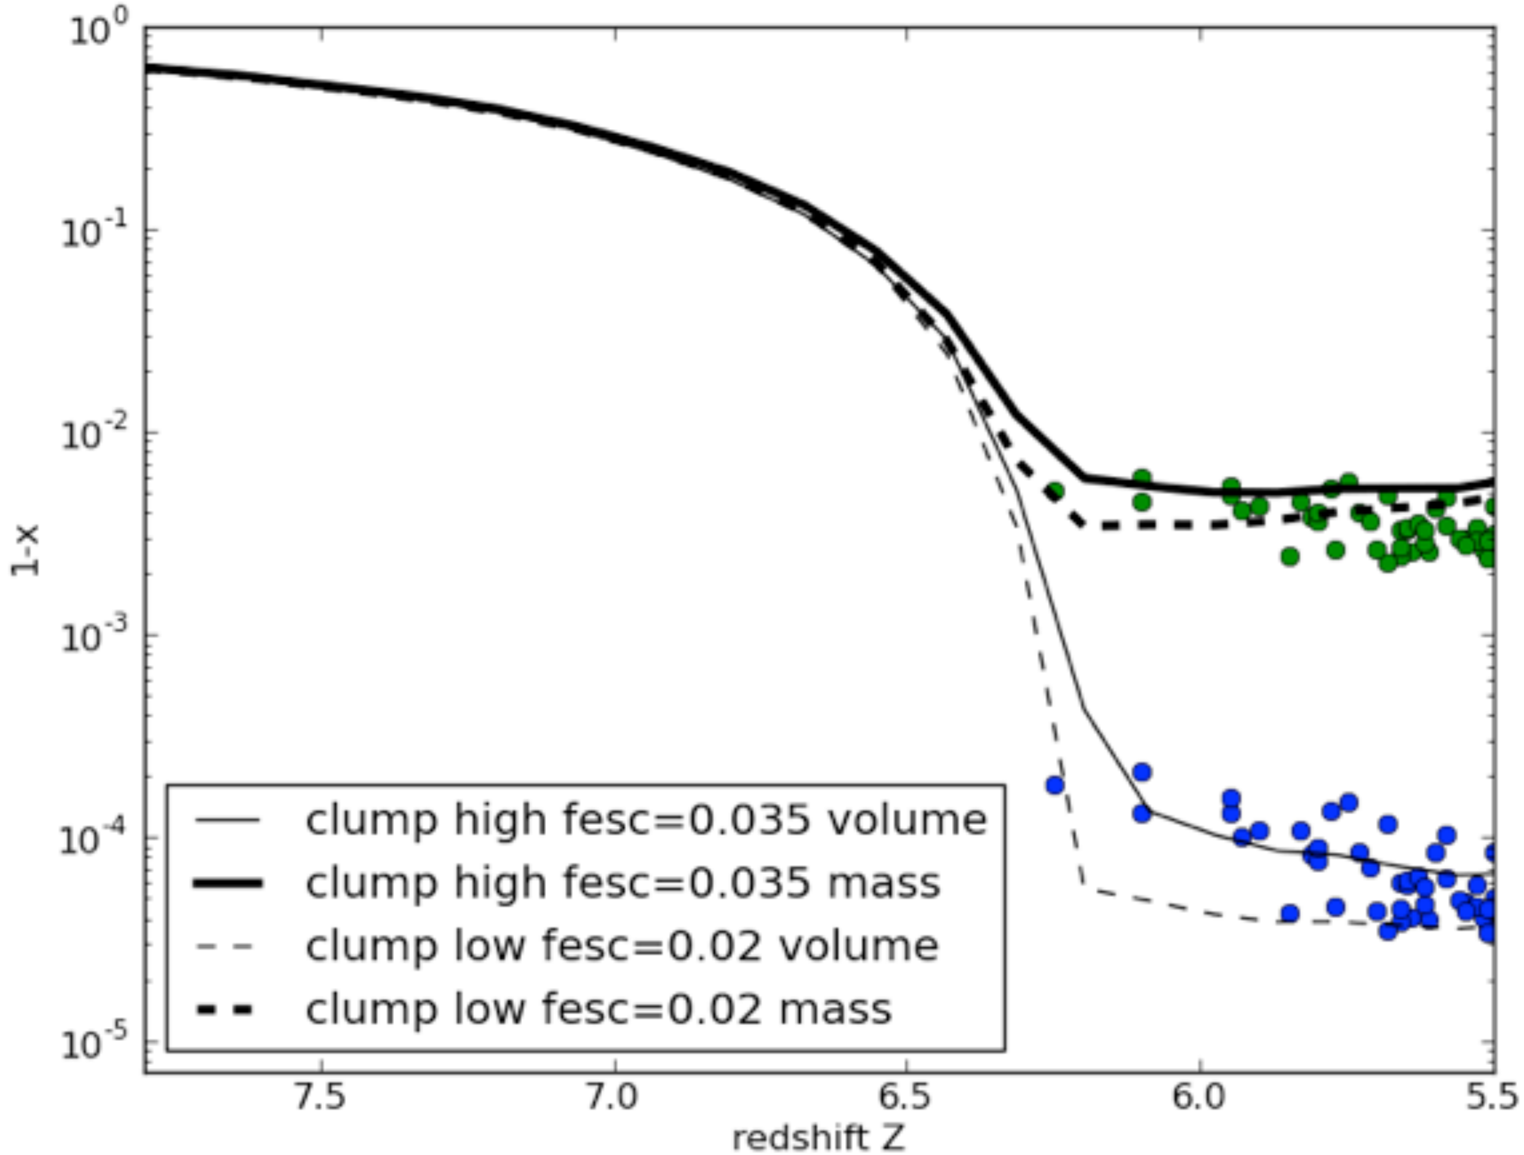

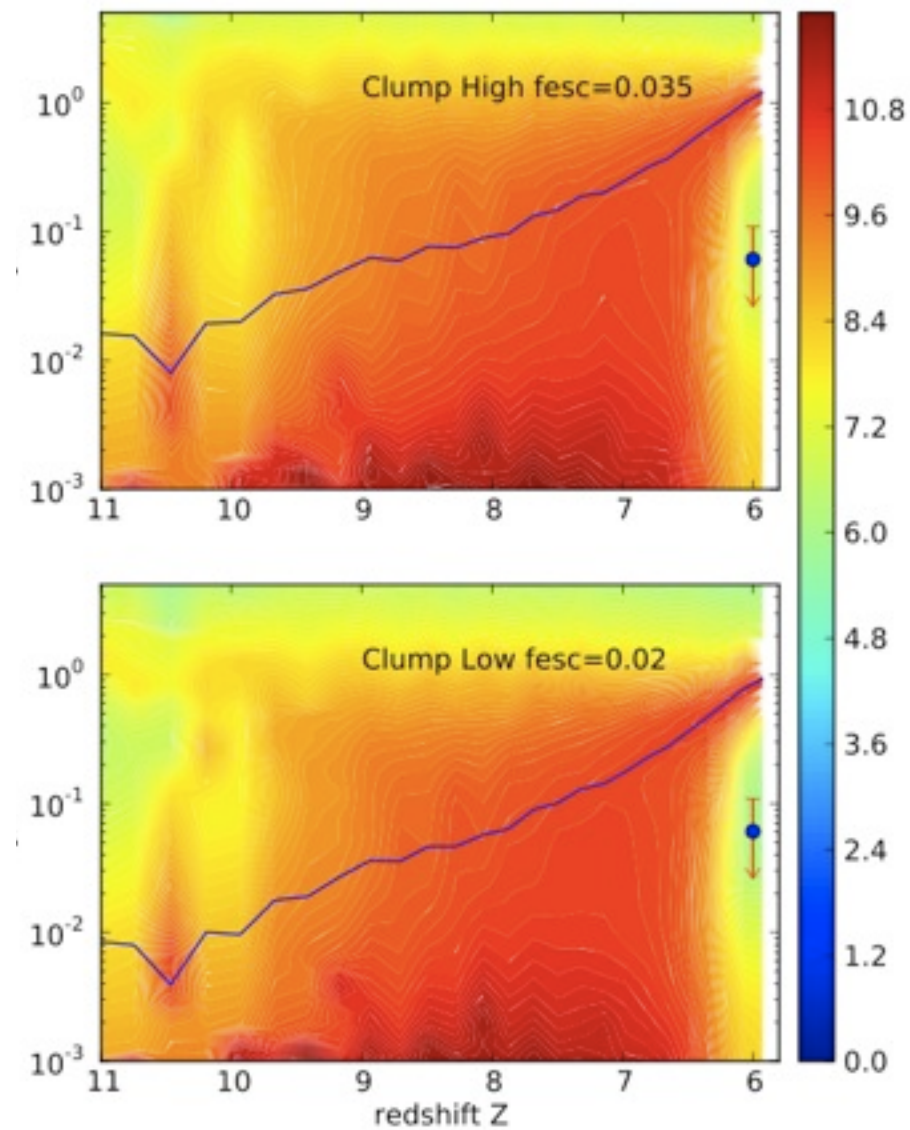

~100 runs @ 1024^3 resolution

Aubert & Teyssier, ApJ, 2010

Residual Neutral Fraction and J21

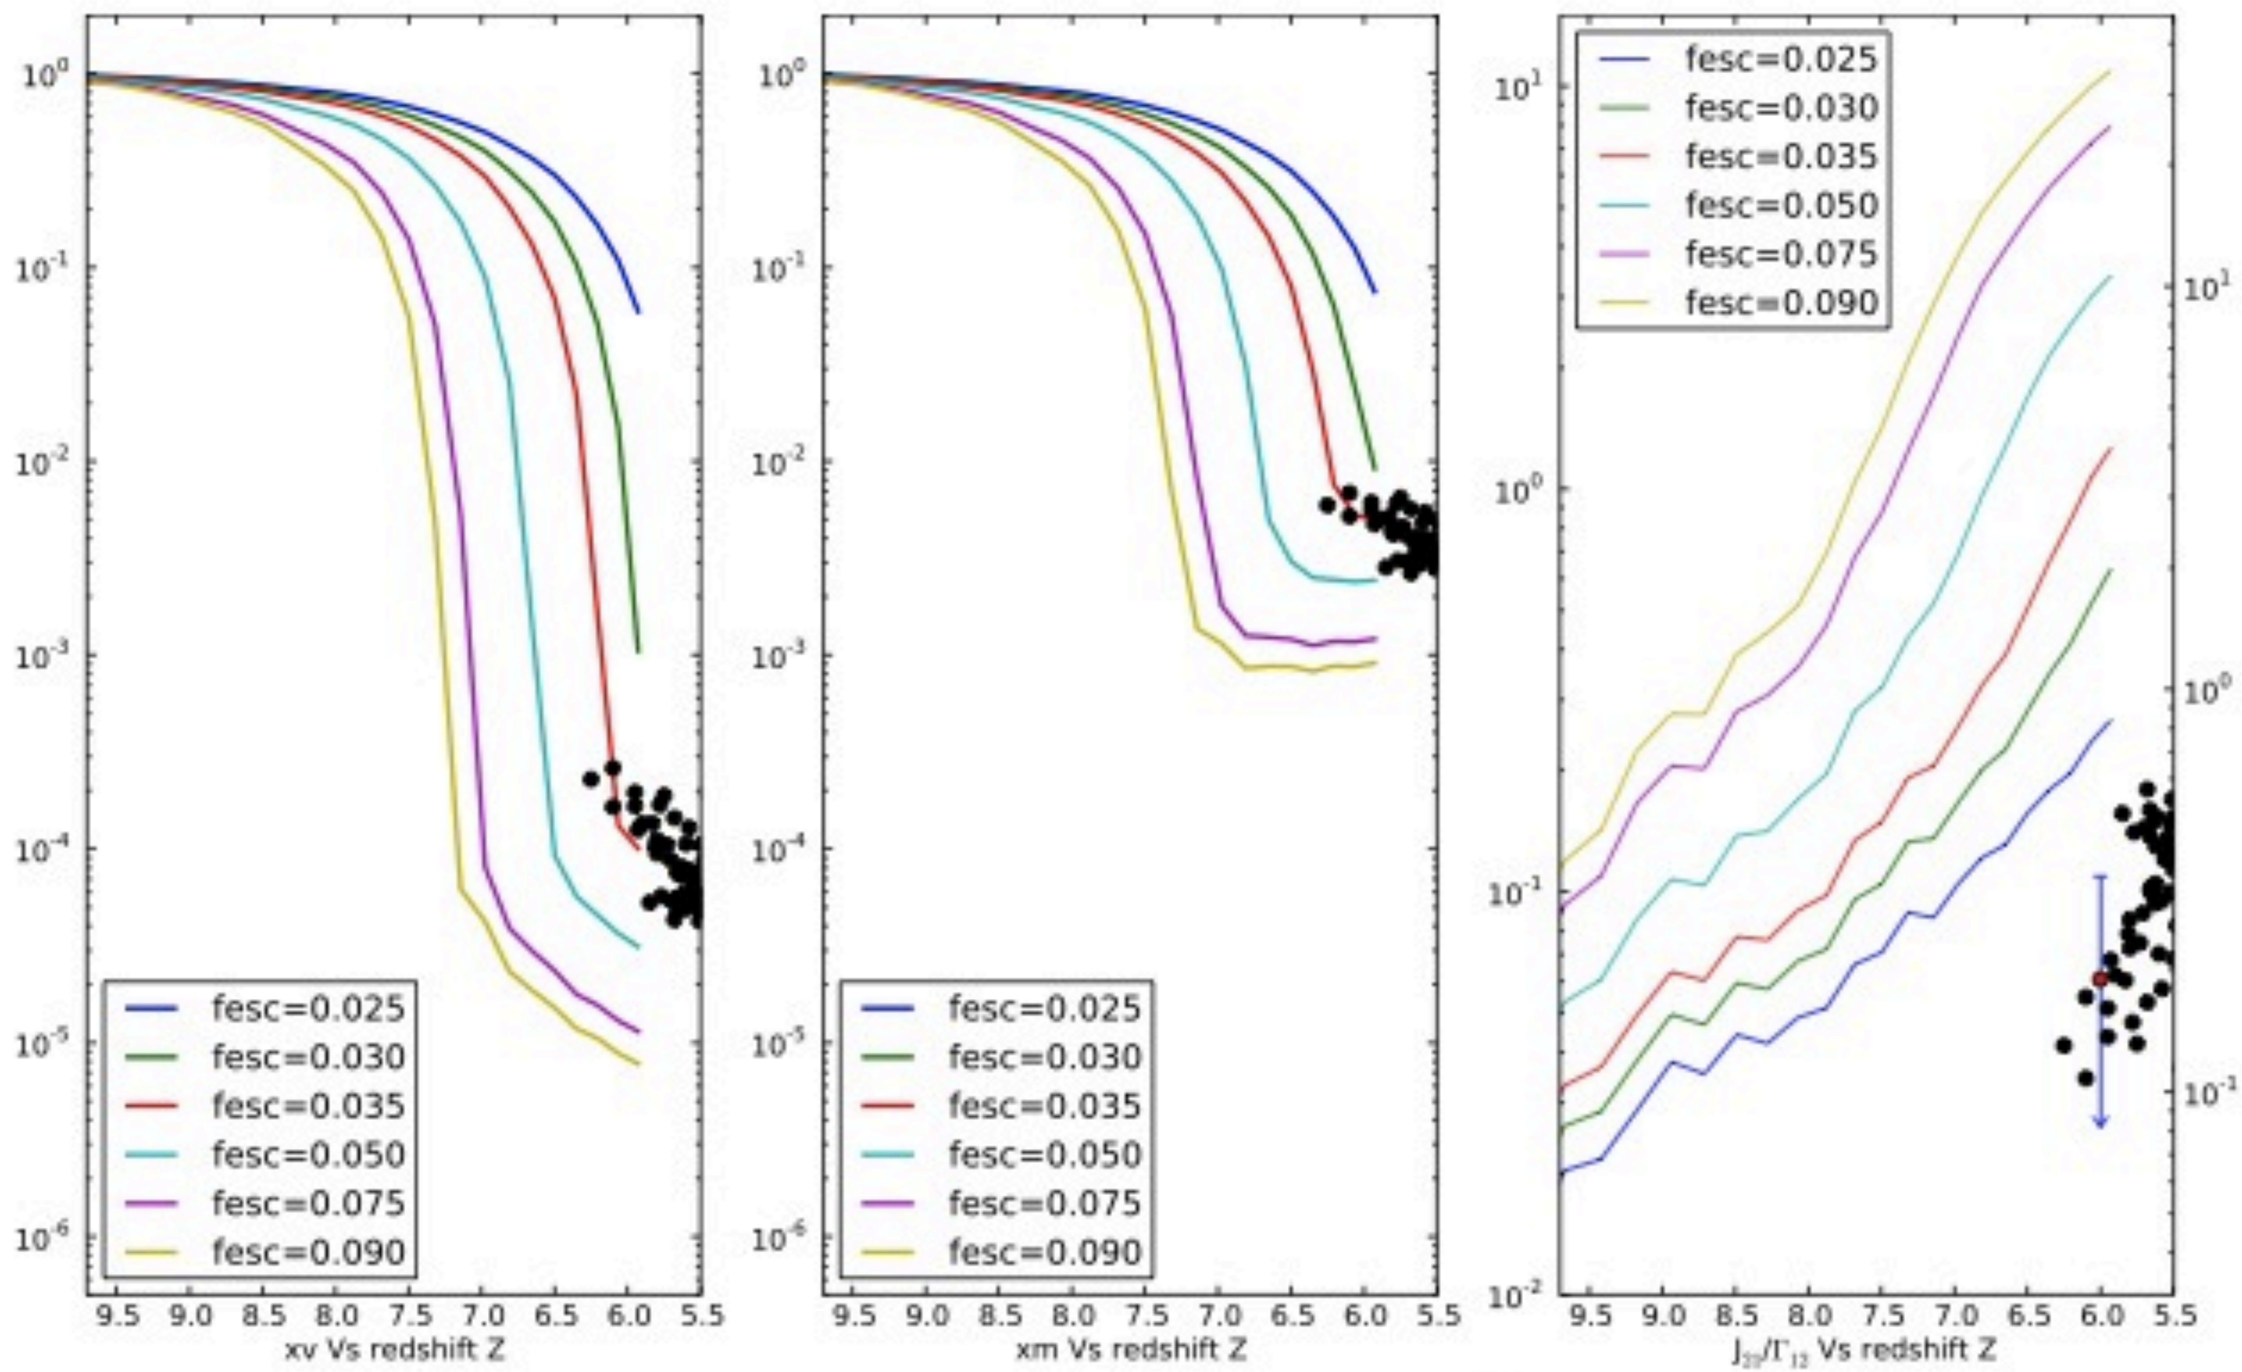

100 Mpc/h -1024^3

Aubert & Teyssier, ApJ, 2010

Application : Local Group Reionisation (with P. Ocvirk)

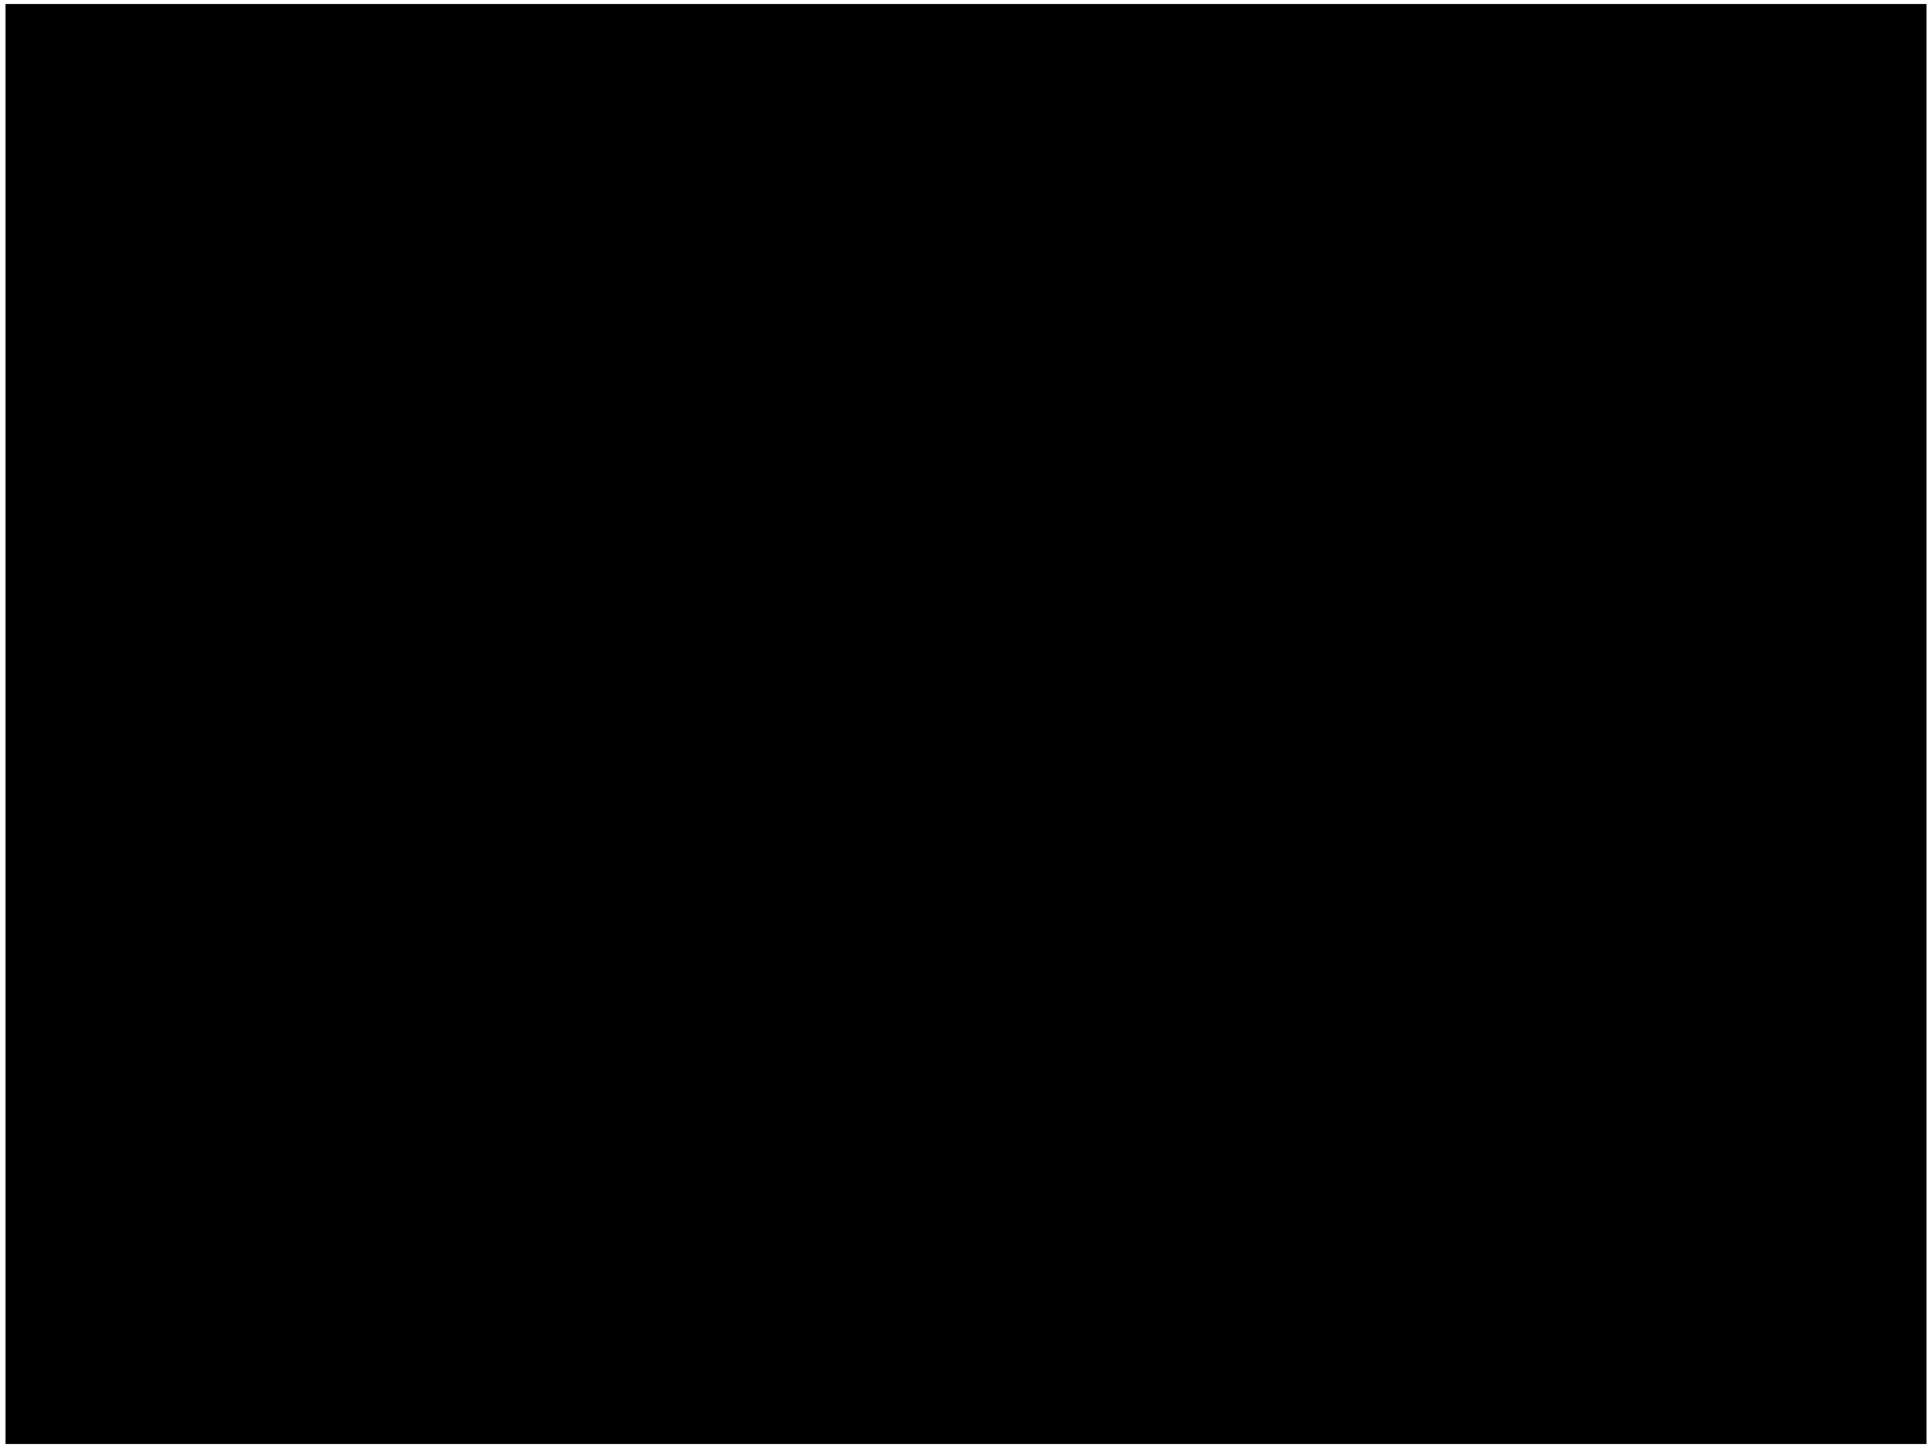

CLUES zoom on the local Group

Timing of the local reionisation ?

Ocvirk et al.2012a,b (submitted+in prep.)

Application: Merger Trees of HII regions during overlap (with J. Chardin)

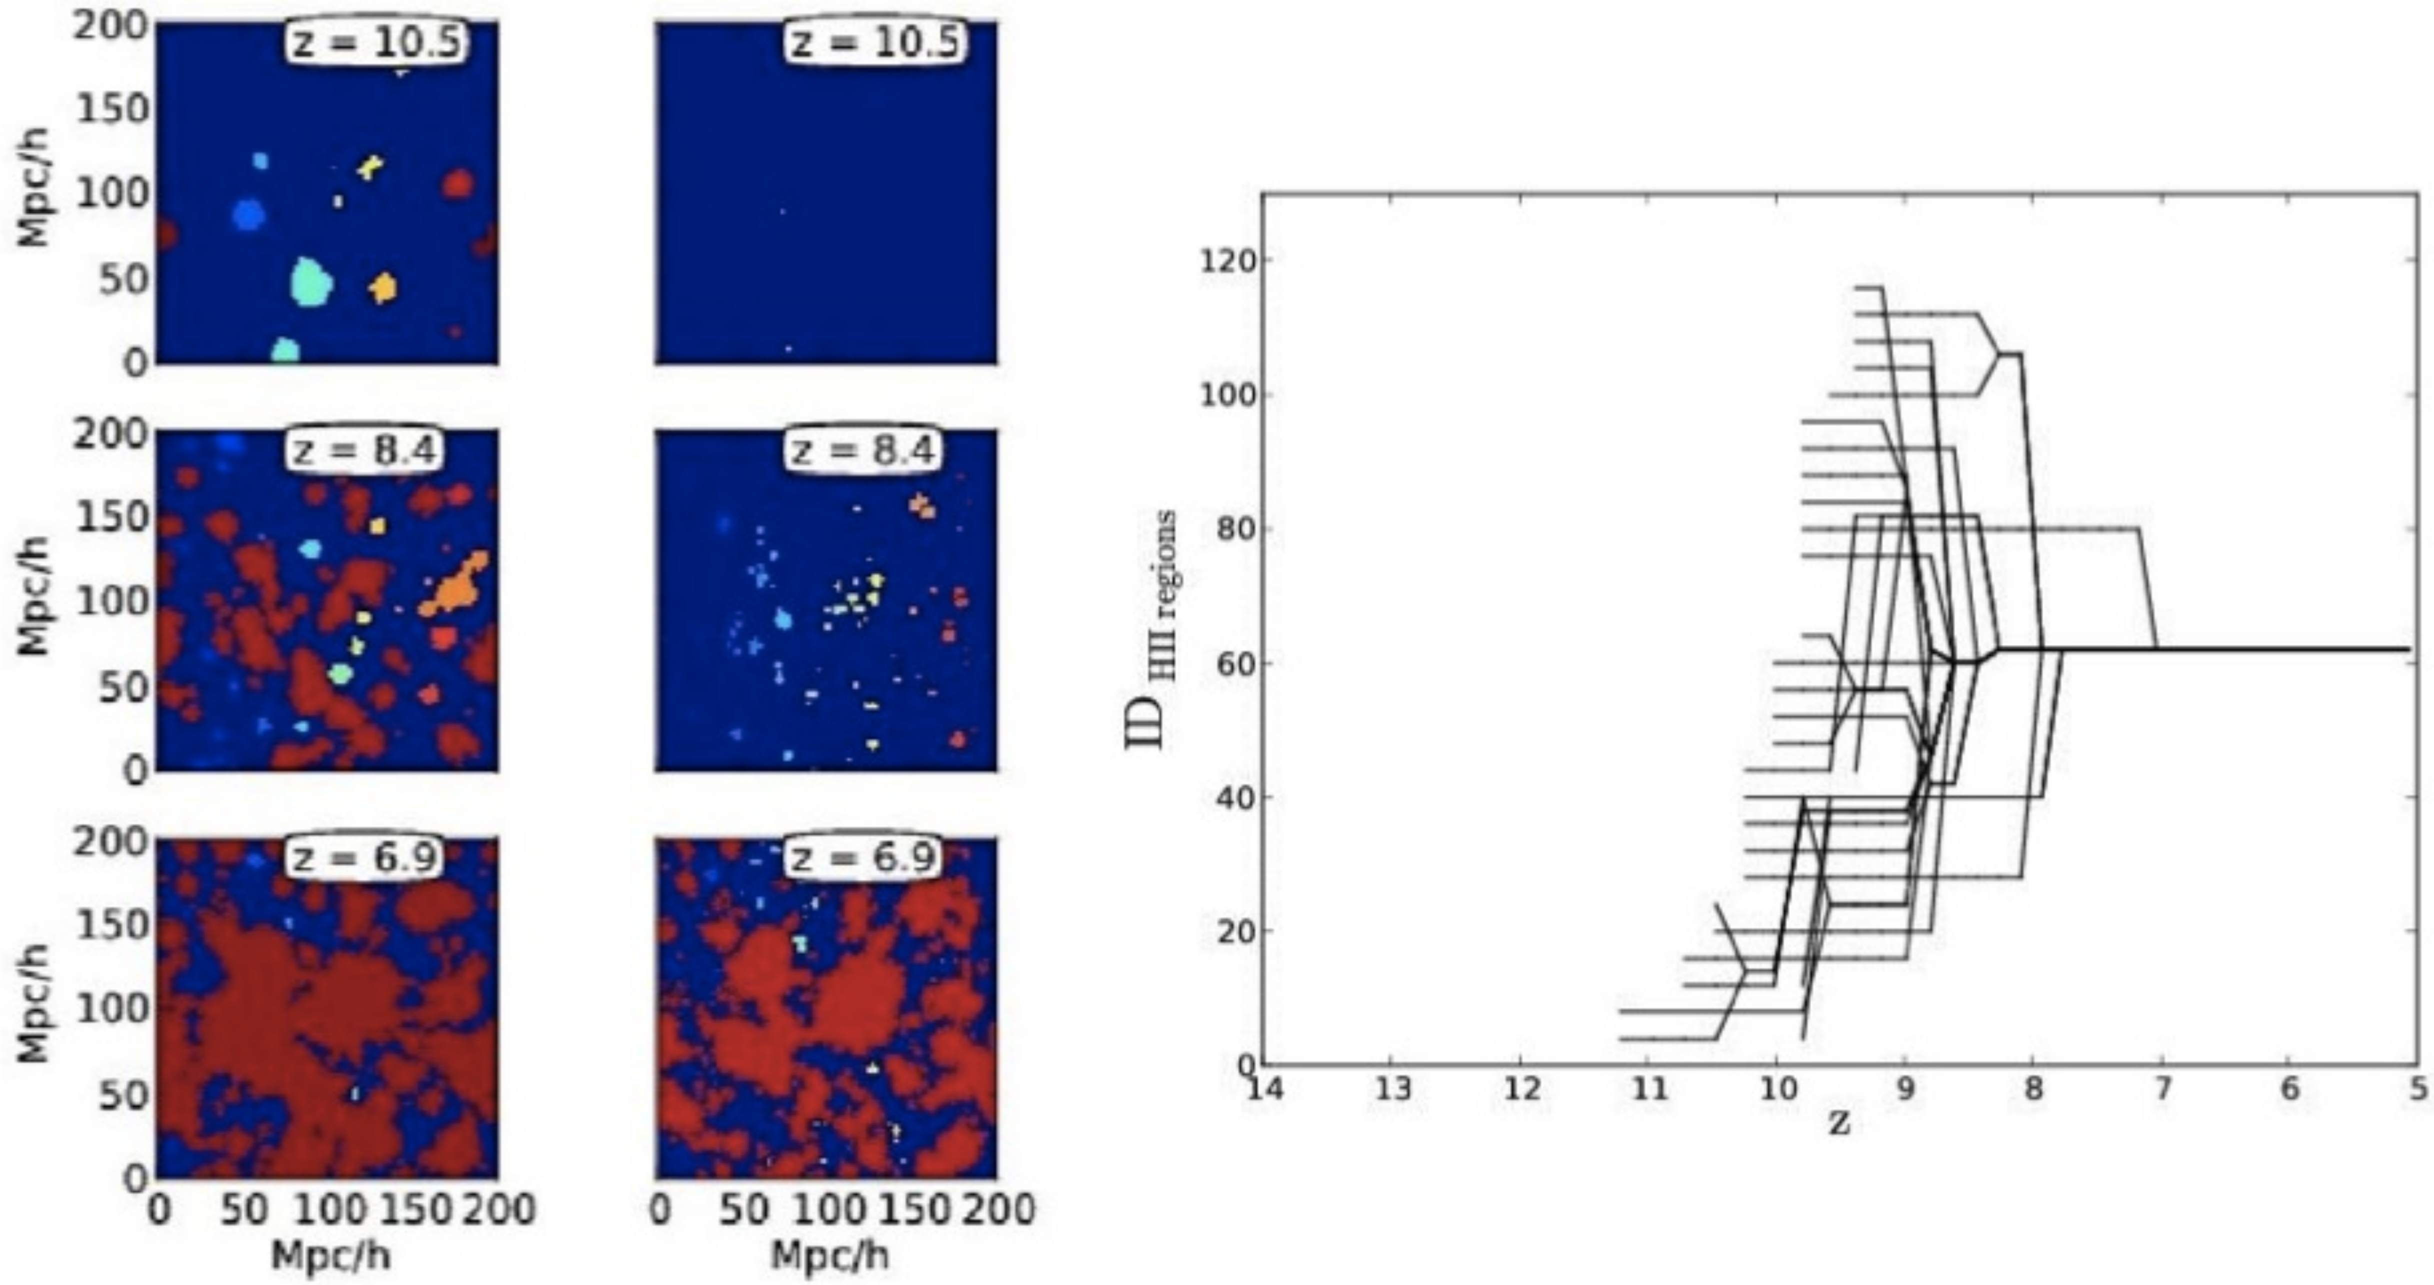

Chardin, Aubert & Ocvirk, A&A, 2012

Large Volumes for 21cm forecast (with B. Semelin)

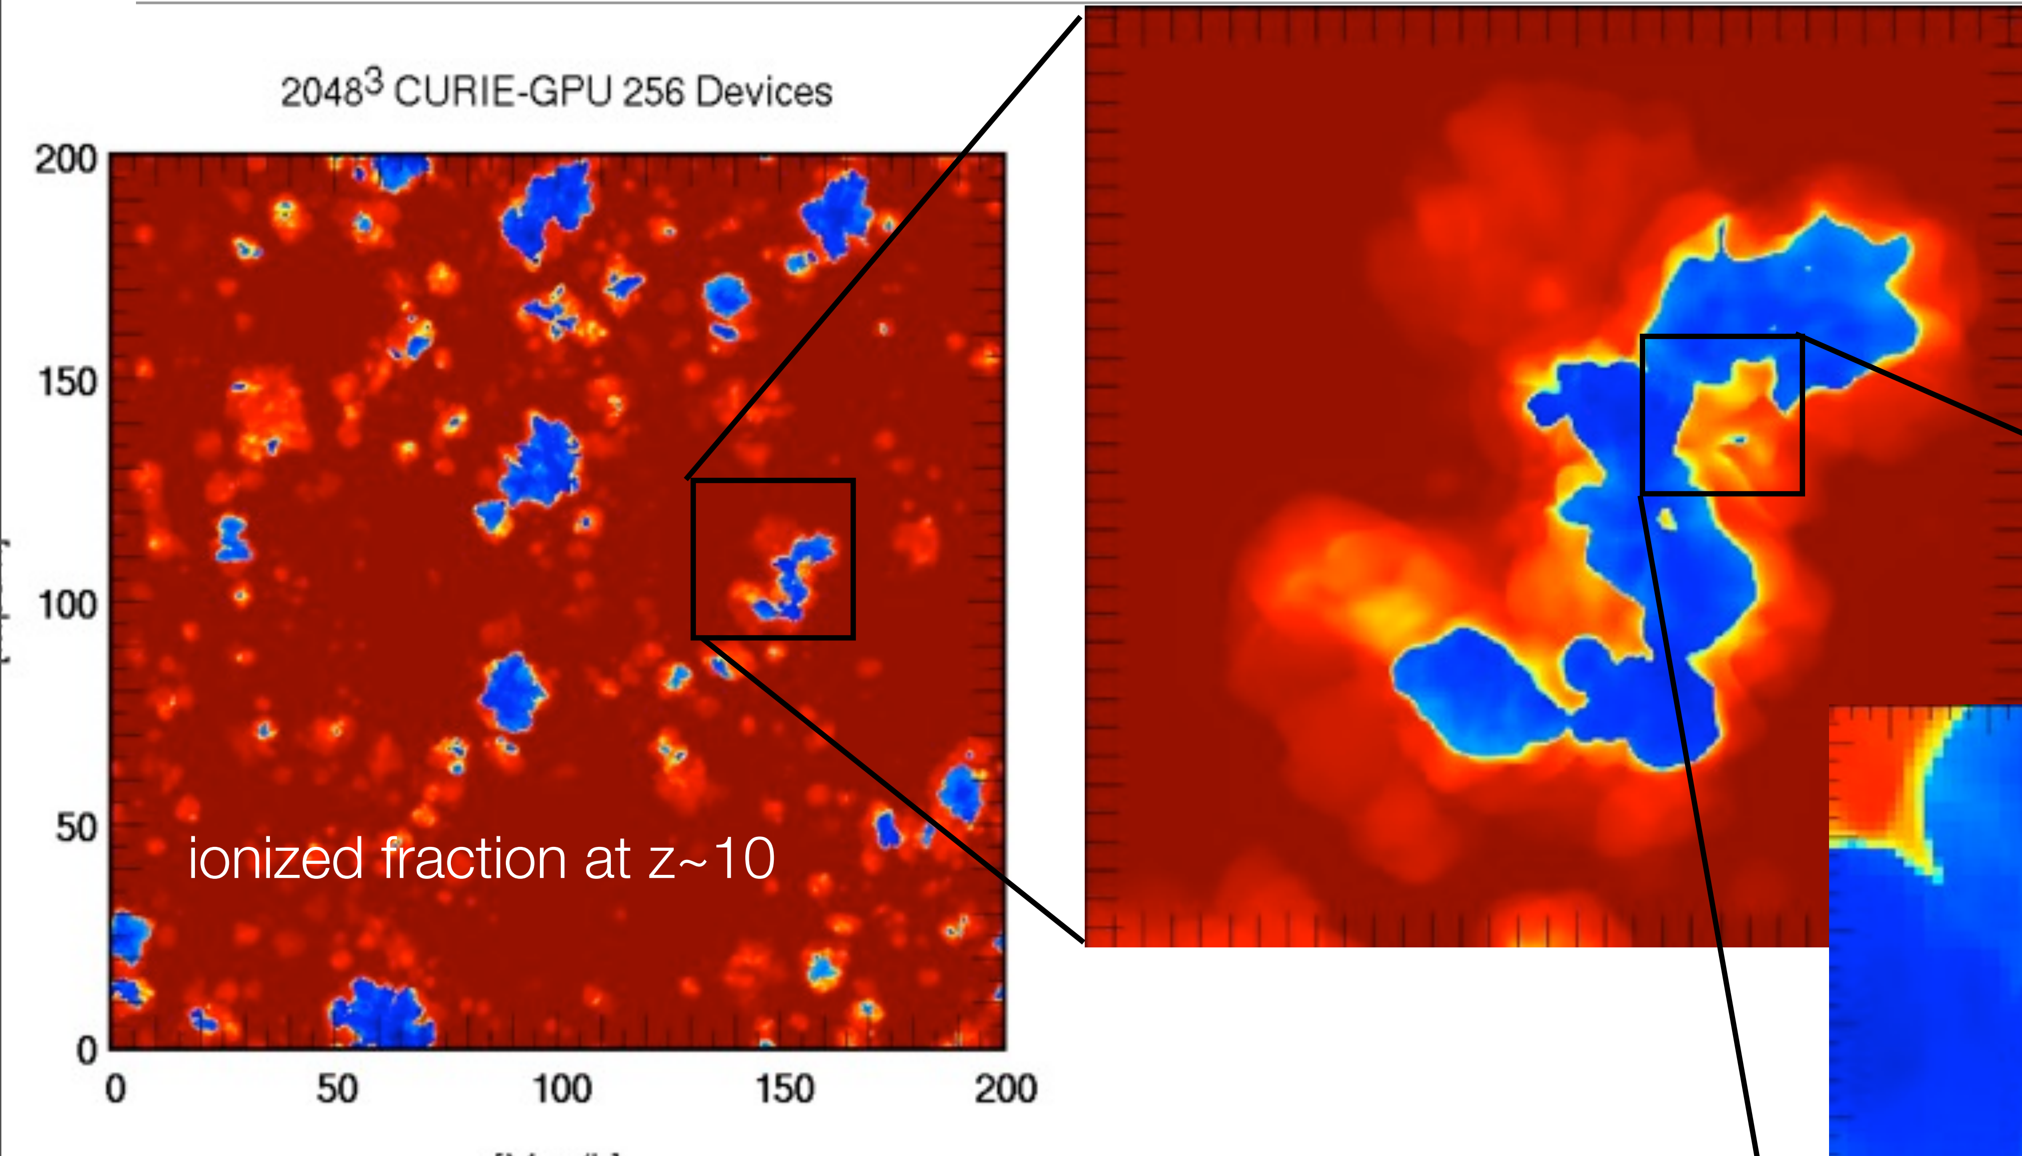

Grand Challenge
Curie-CCRT
256 GPUs
2048x2048x2048
60 000 pdt -15h

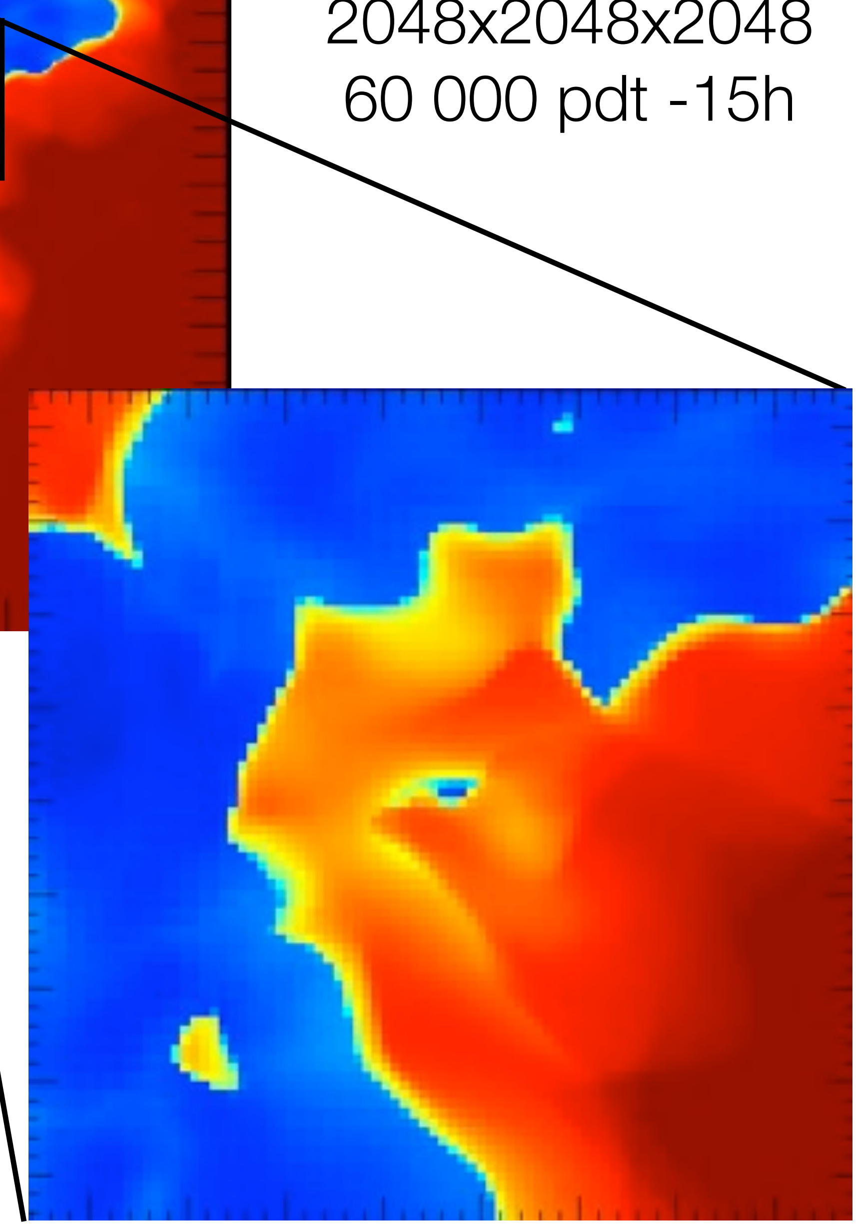

RAMSES-RT (with T. Stranex & R. Teyssier, Zurich)

RAMSES & ATON are coupled

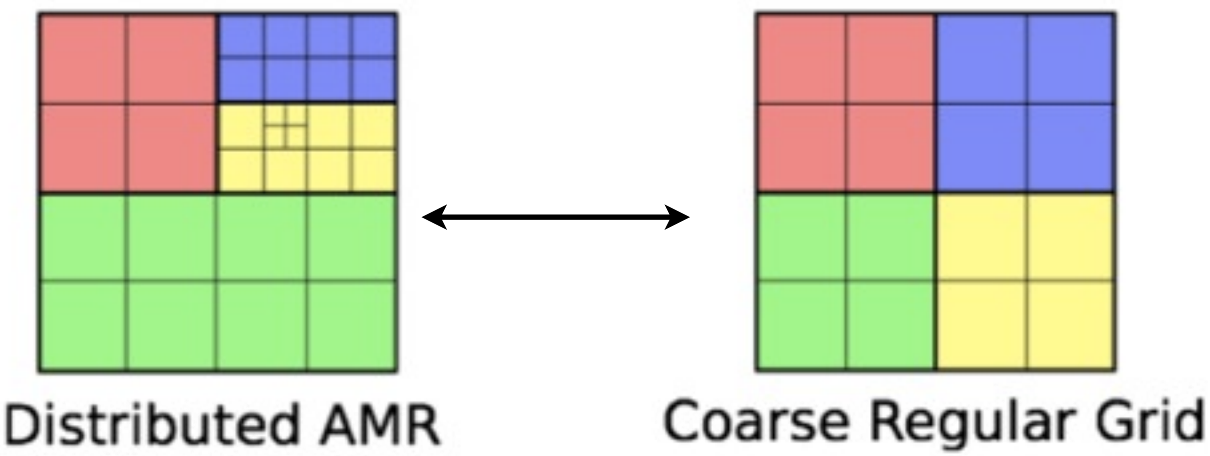

RAMSES
(DYNAMICS)

ATON
(RT)

256^3 cosmo simulation on Titane
16 MPI processes, 2 GPUs per machine, 1 MPI process per GPU
Timing per coarse time step:

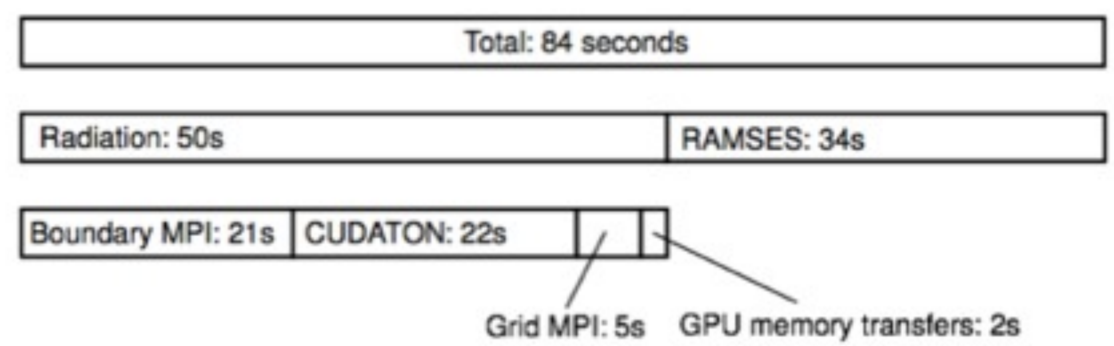

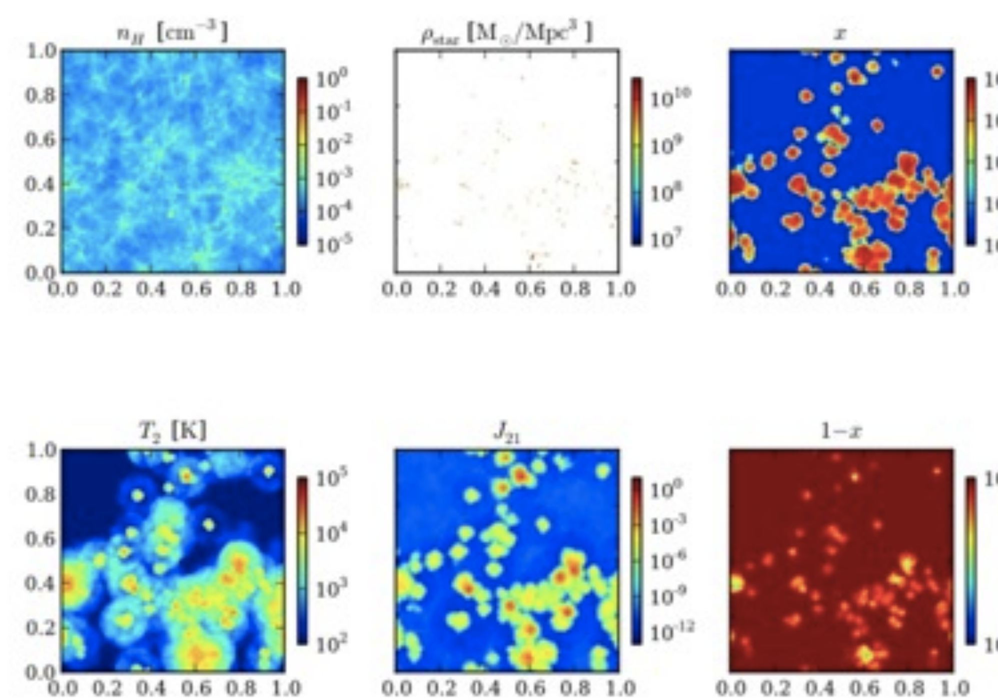

UNIGRID version will be used on Titan for the INCITE project

courtesy T. Stranex

3. Towards Multi-Fluid AMR

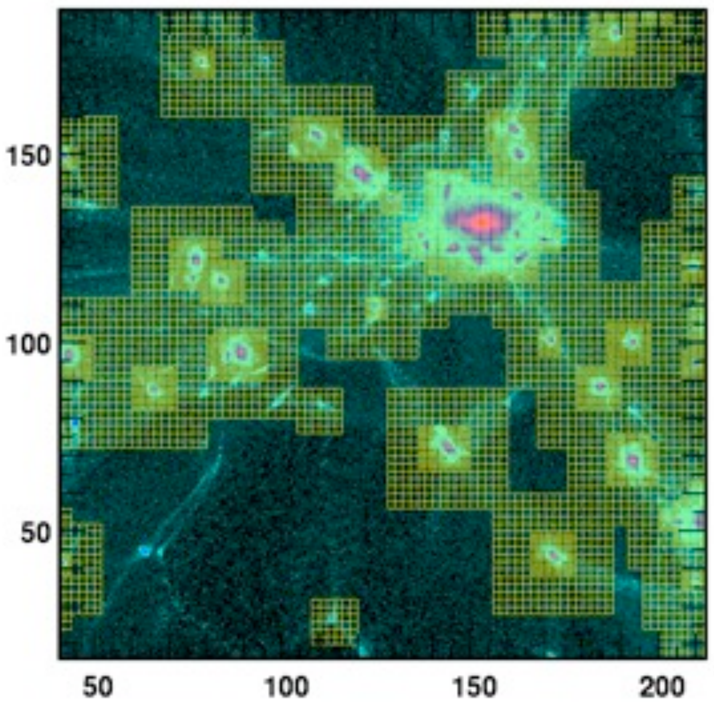

N-Body
AMR+GPU+Multi ok
~x10 w.r.t. CPU
Aubert et al. 2009

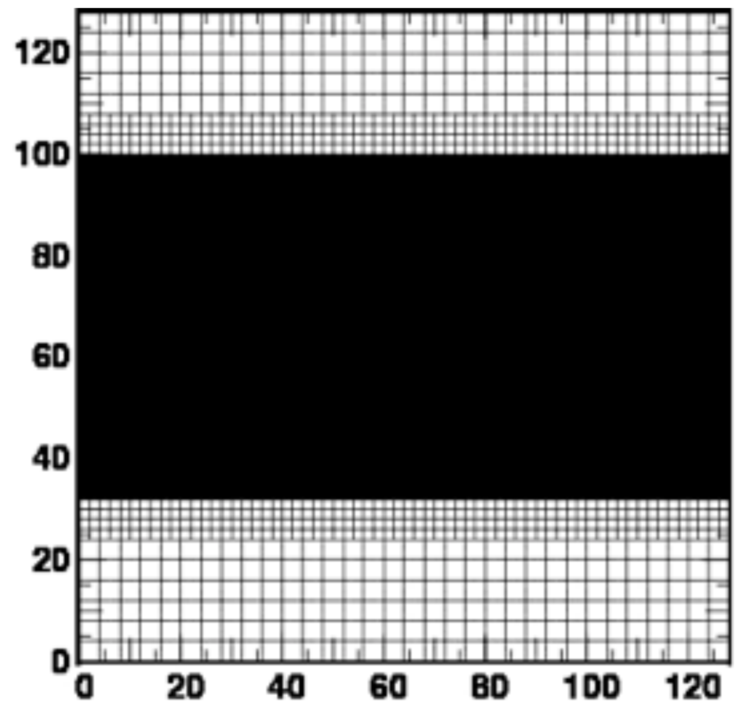

Hydro
AMR+GPU ok
~x15 w.r.t. CPU

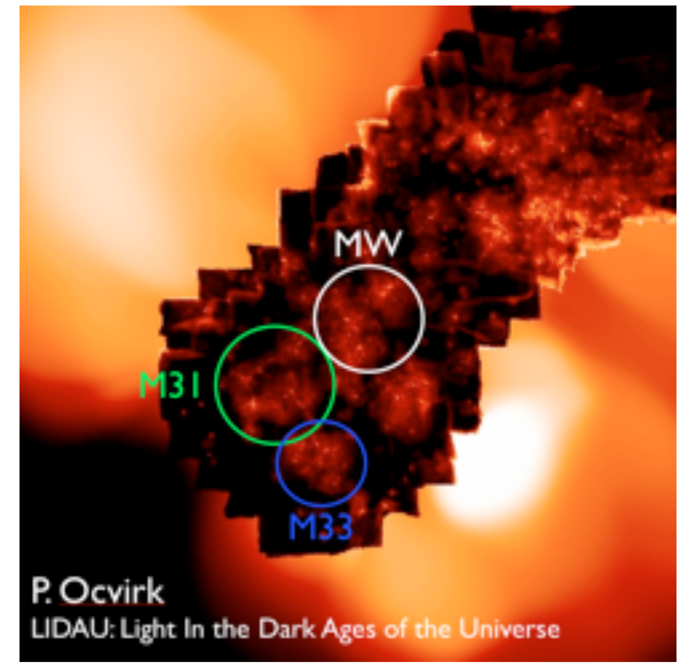

Radiative Transfer
GPU+ Multi ok
x 30-40 ? w.r.t. CPU
Aubert & Teyssier, 2008,2010

EMMA Project:
-3 fluids coupled on an AMR structure with Hardware Acceleration, with e.g. GPUs

Multi-GPU PM

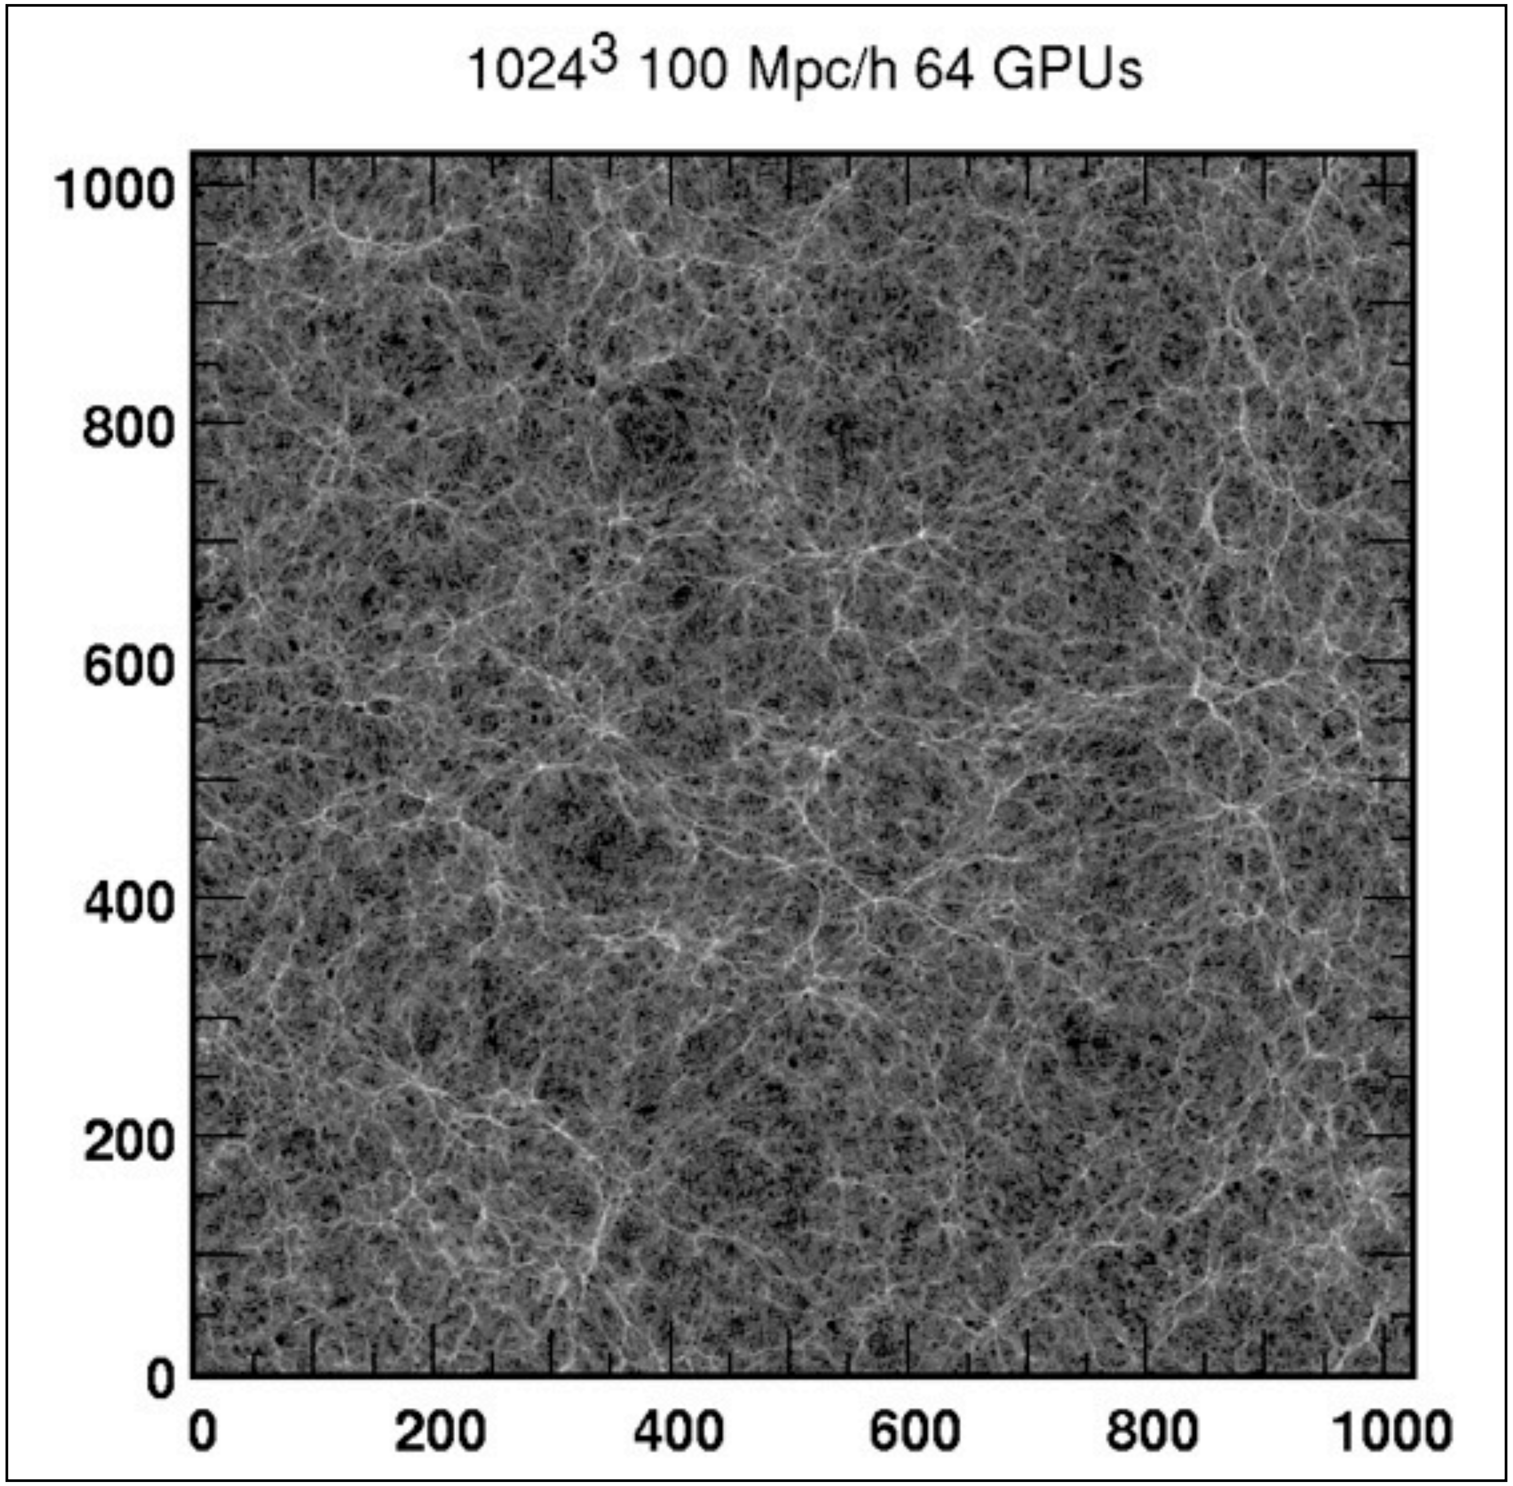

1.2 billions particles
(1024^3 real particles
+2 10^8 ghosts)
8 sec/tstep on 64 Teslas
with 25 % spent in
communications

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

asynchronous coms ?

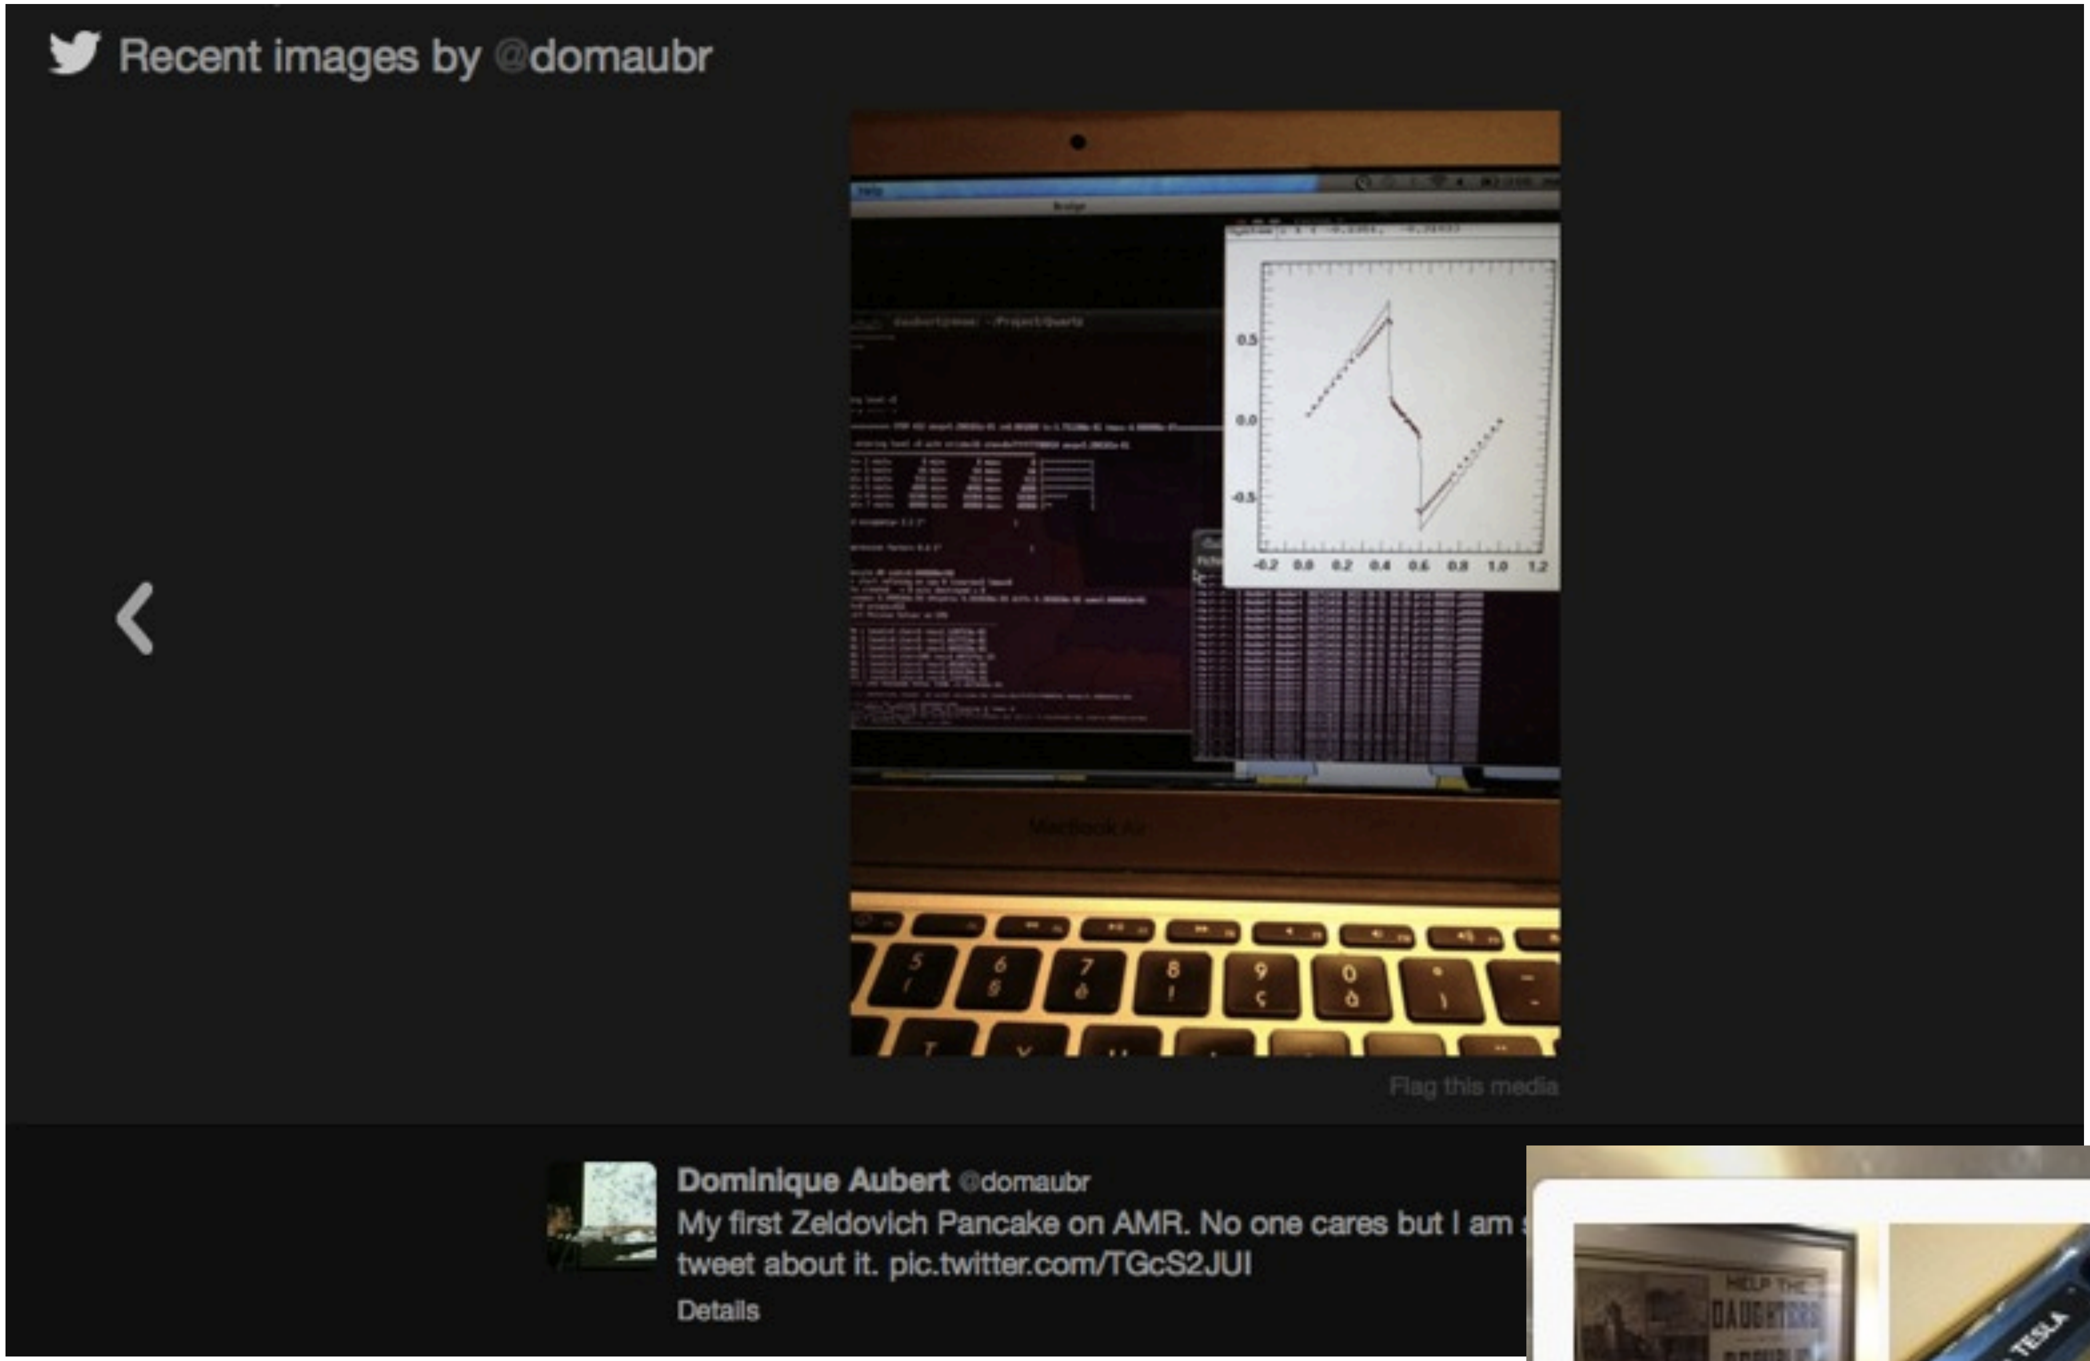

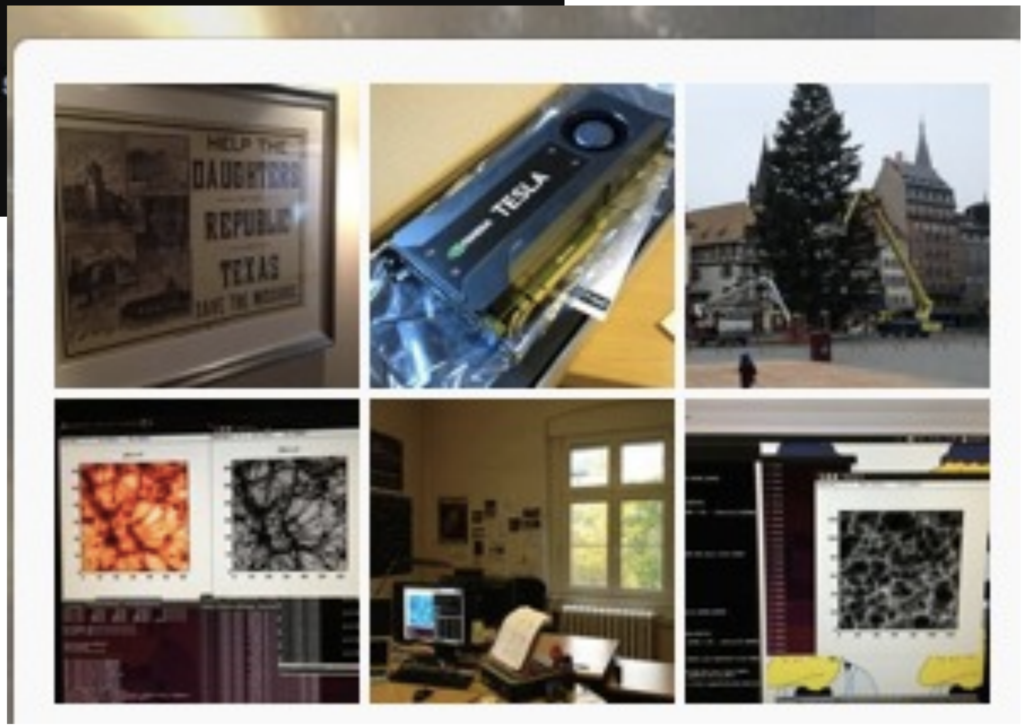

Under Heavy development

Quartz

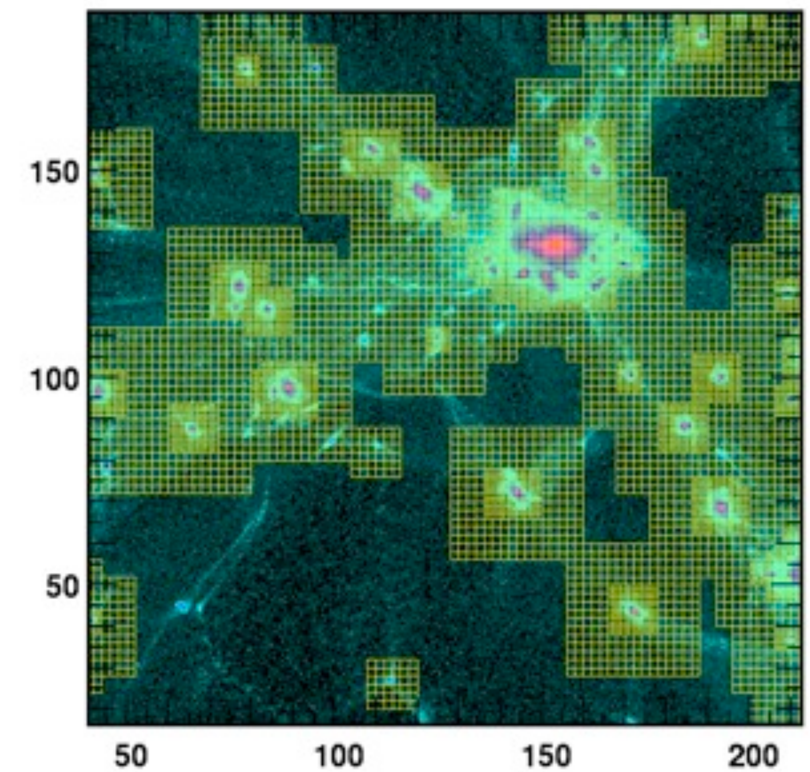

- Written in C+CUDA+MPI
- Parallel (Space-Filling Curve + essential Tree domain decomposition)
- AMR, with FTT data structure
- N-Body + Hydro only (for the moment)
- MG Poisson Solver on GPU+ MUSCL-Hancock Godunov Hydo Solver on GPU + Data Logistics on CPU
- Hopefully will become EMMA (**E**lectro**M**agnetism and **M**echanics on **A**MR) for gravity+hydro **+radiation**

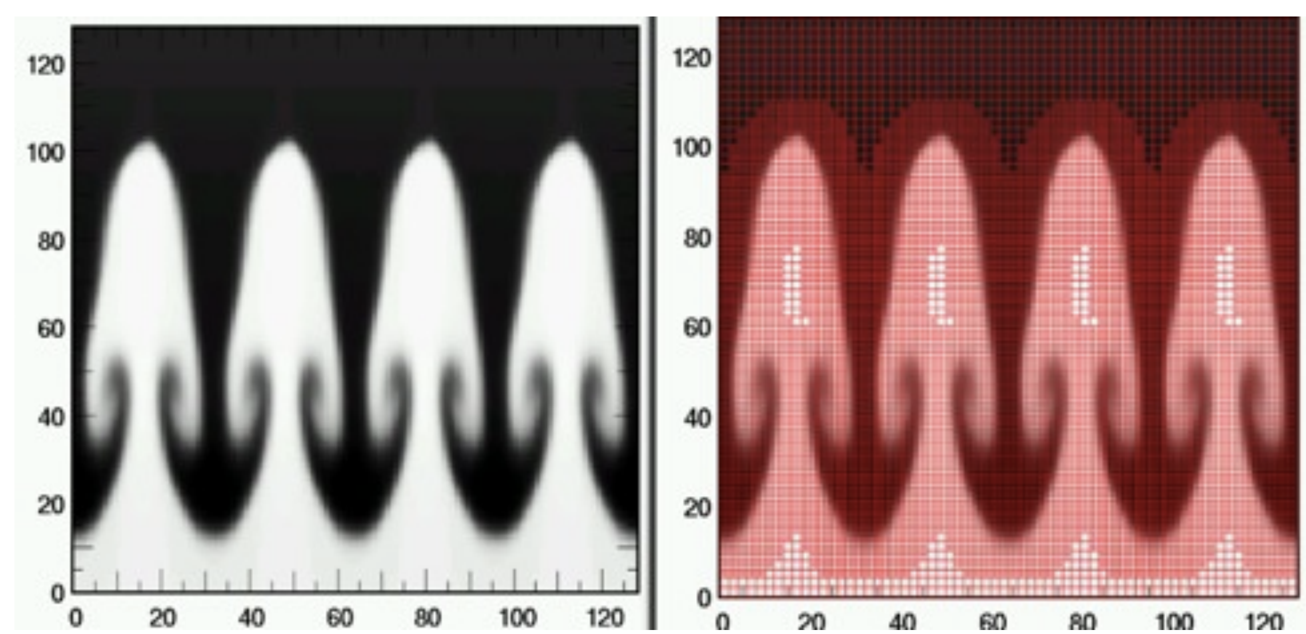

Fully Threaded Tree (Khokhlov 1997)
(aka «Pointer Party»)

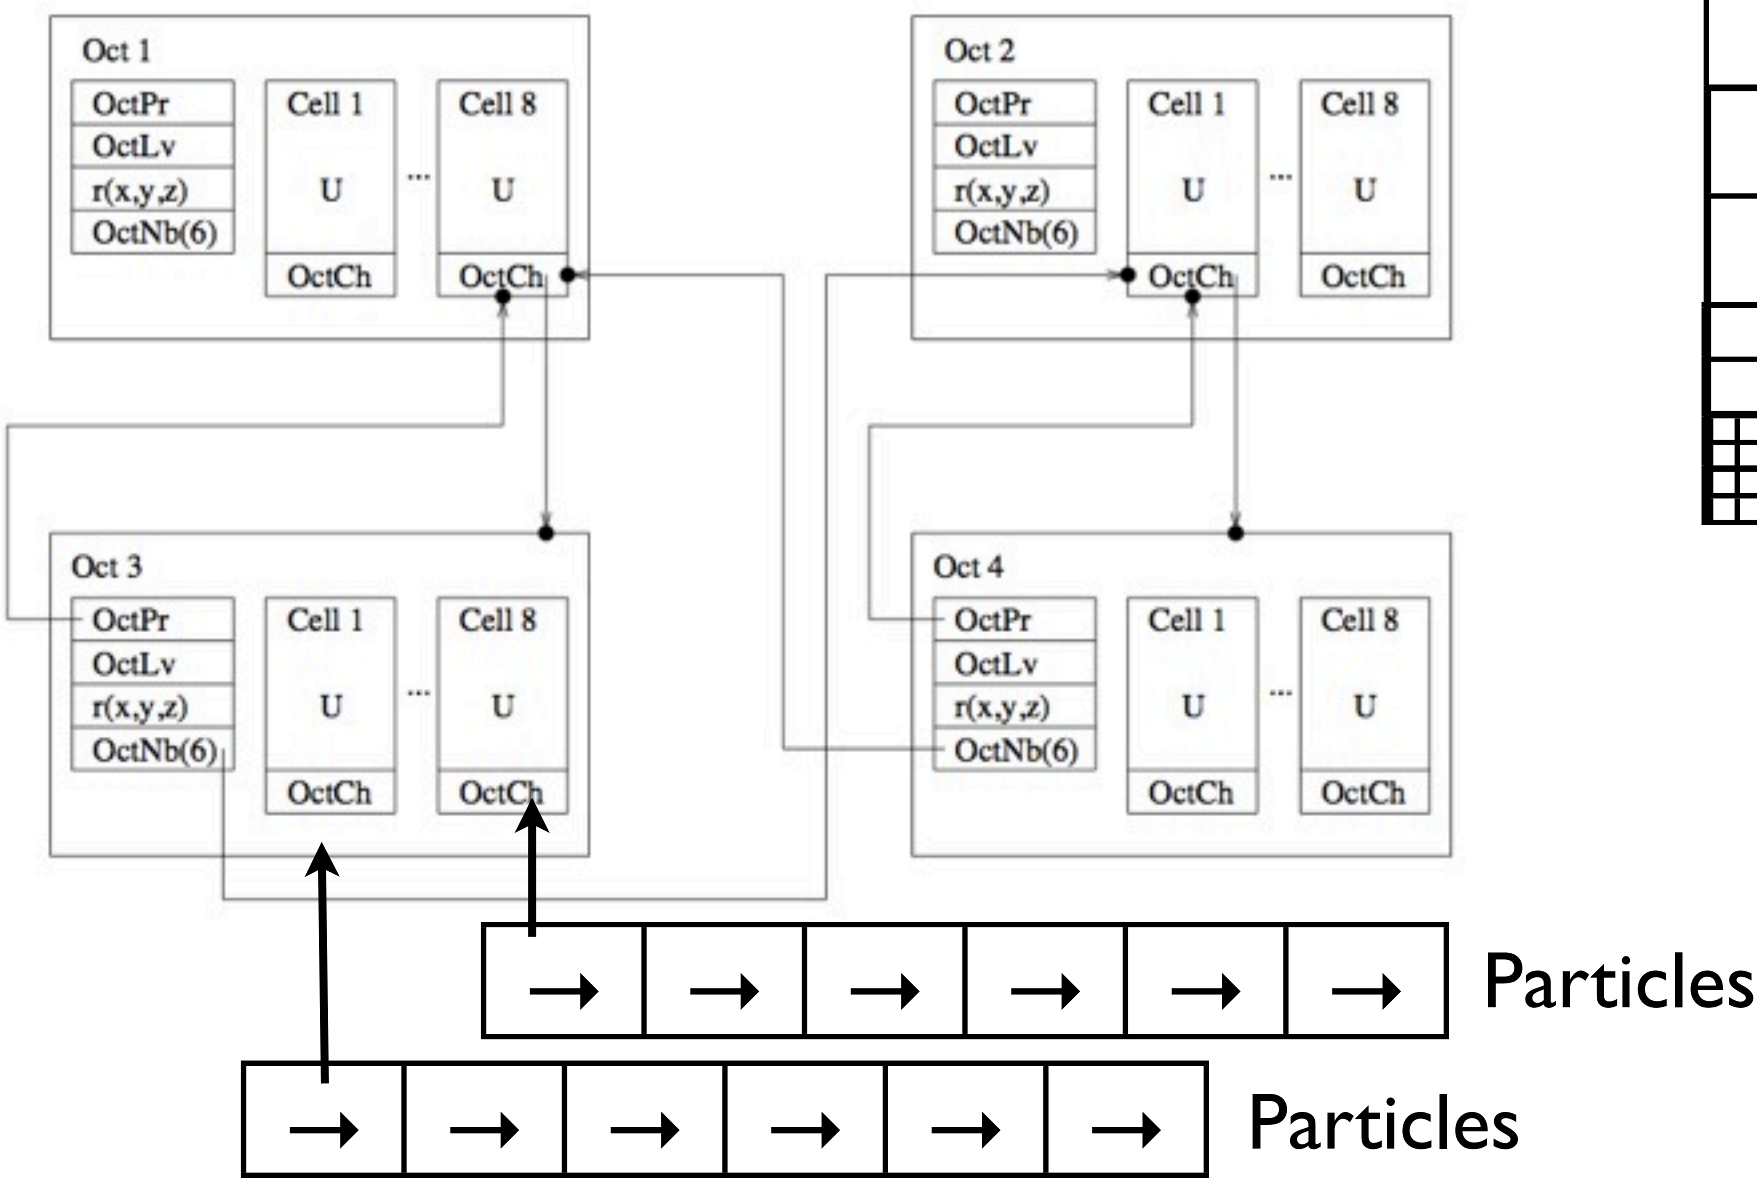

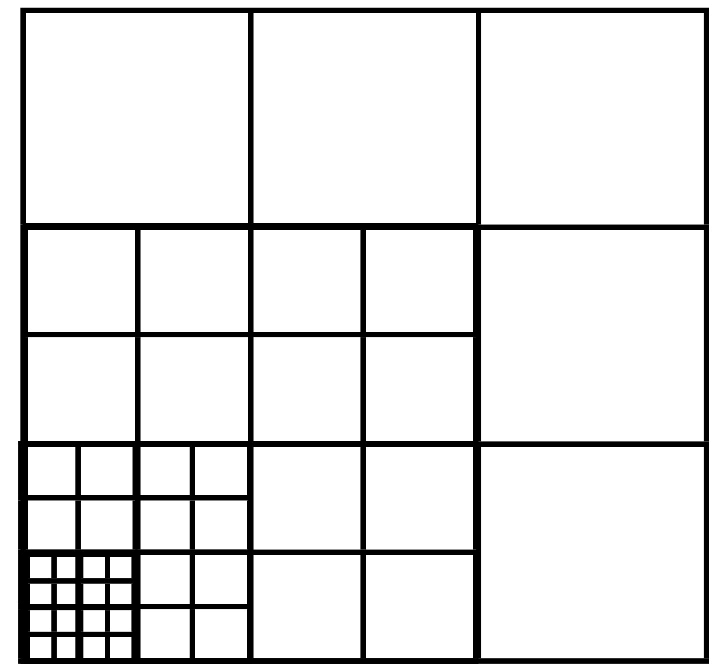

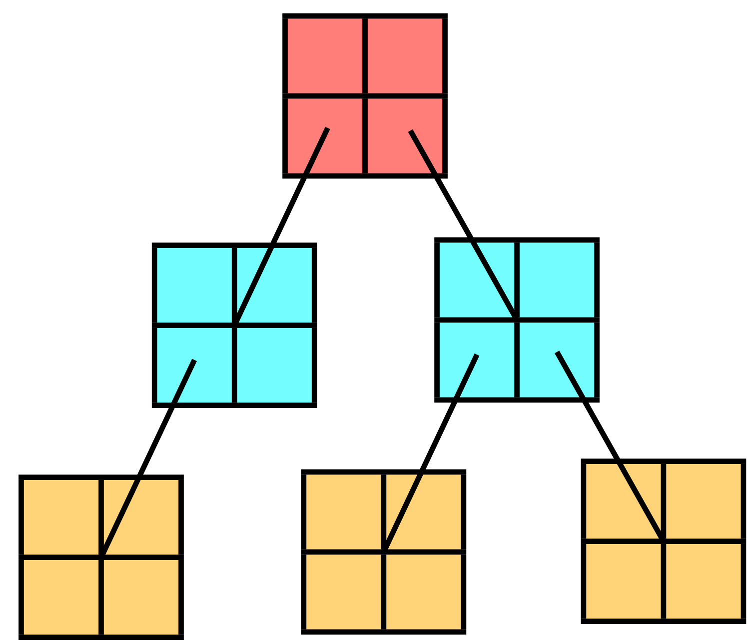

ART (Kravtsov et al. 1997) RAMSES (Teyssier 2001)

Fully **THREADED** Tree

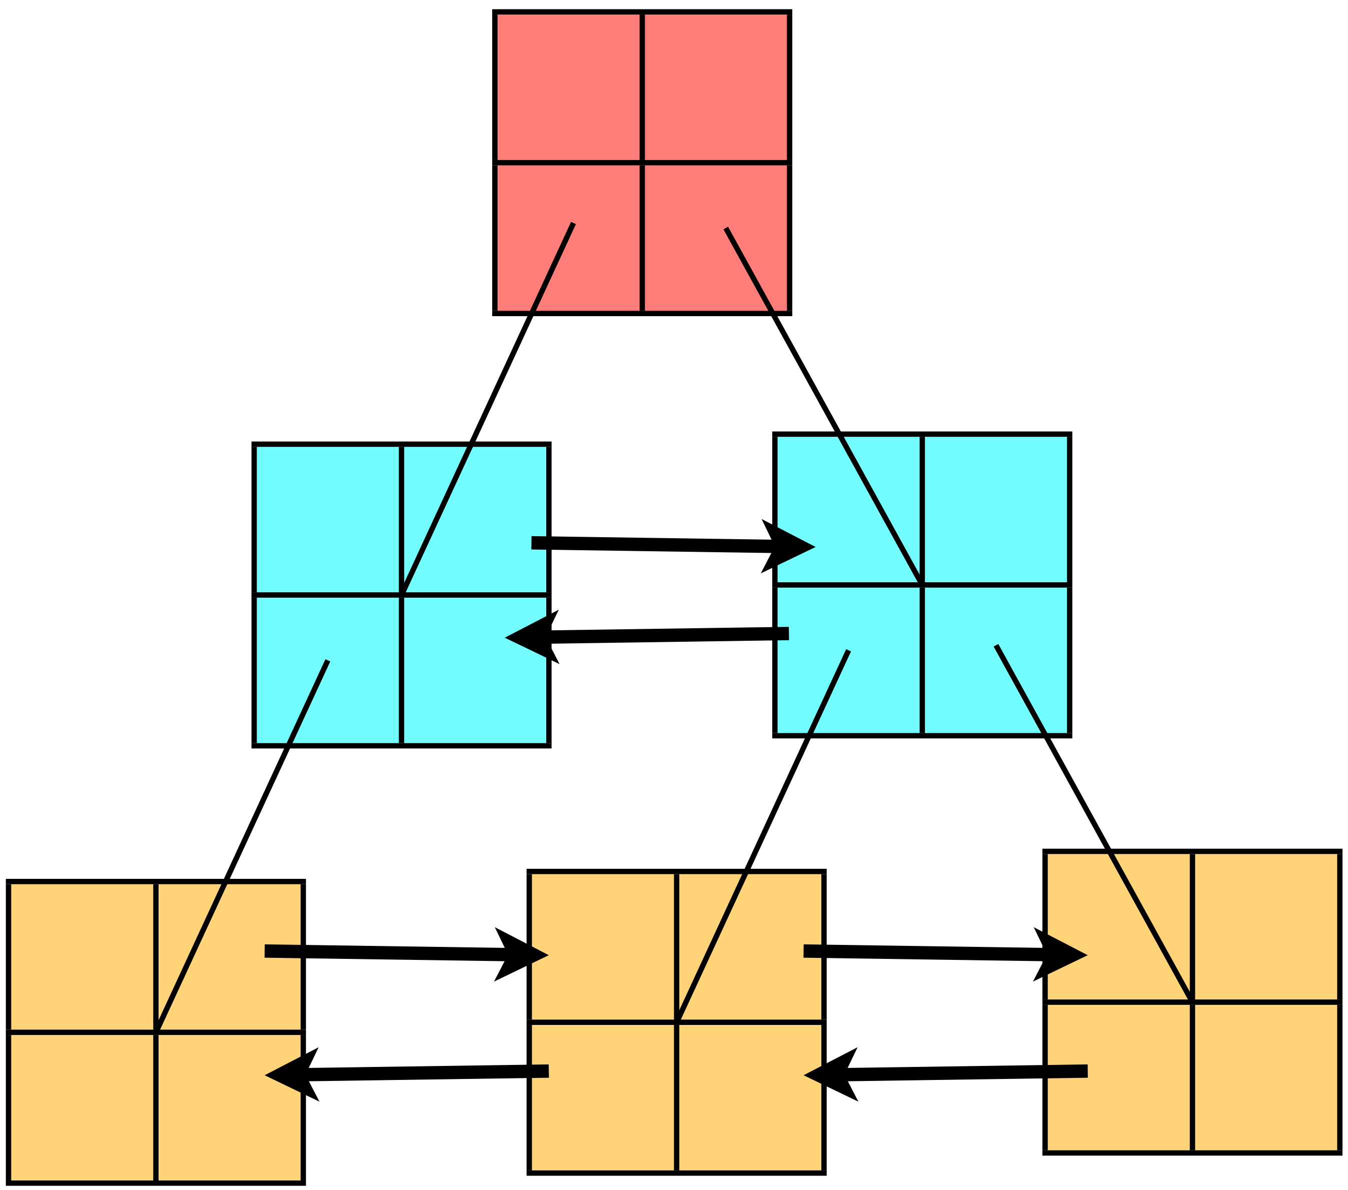

In a lot of cases, the tree is explored **Horizontally**
Level by Level (with some +/-1 level interactions at boundaries)

Even CIC can be considered level by level

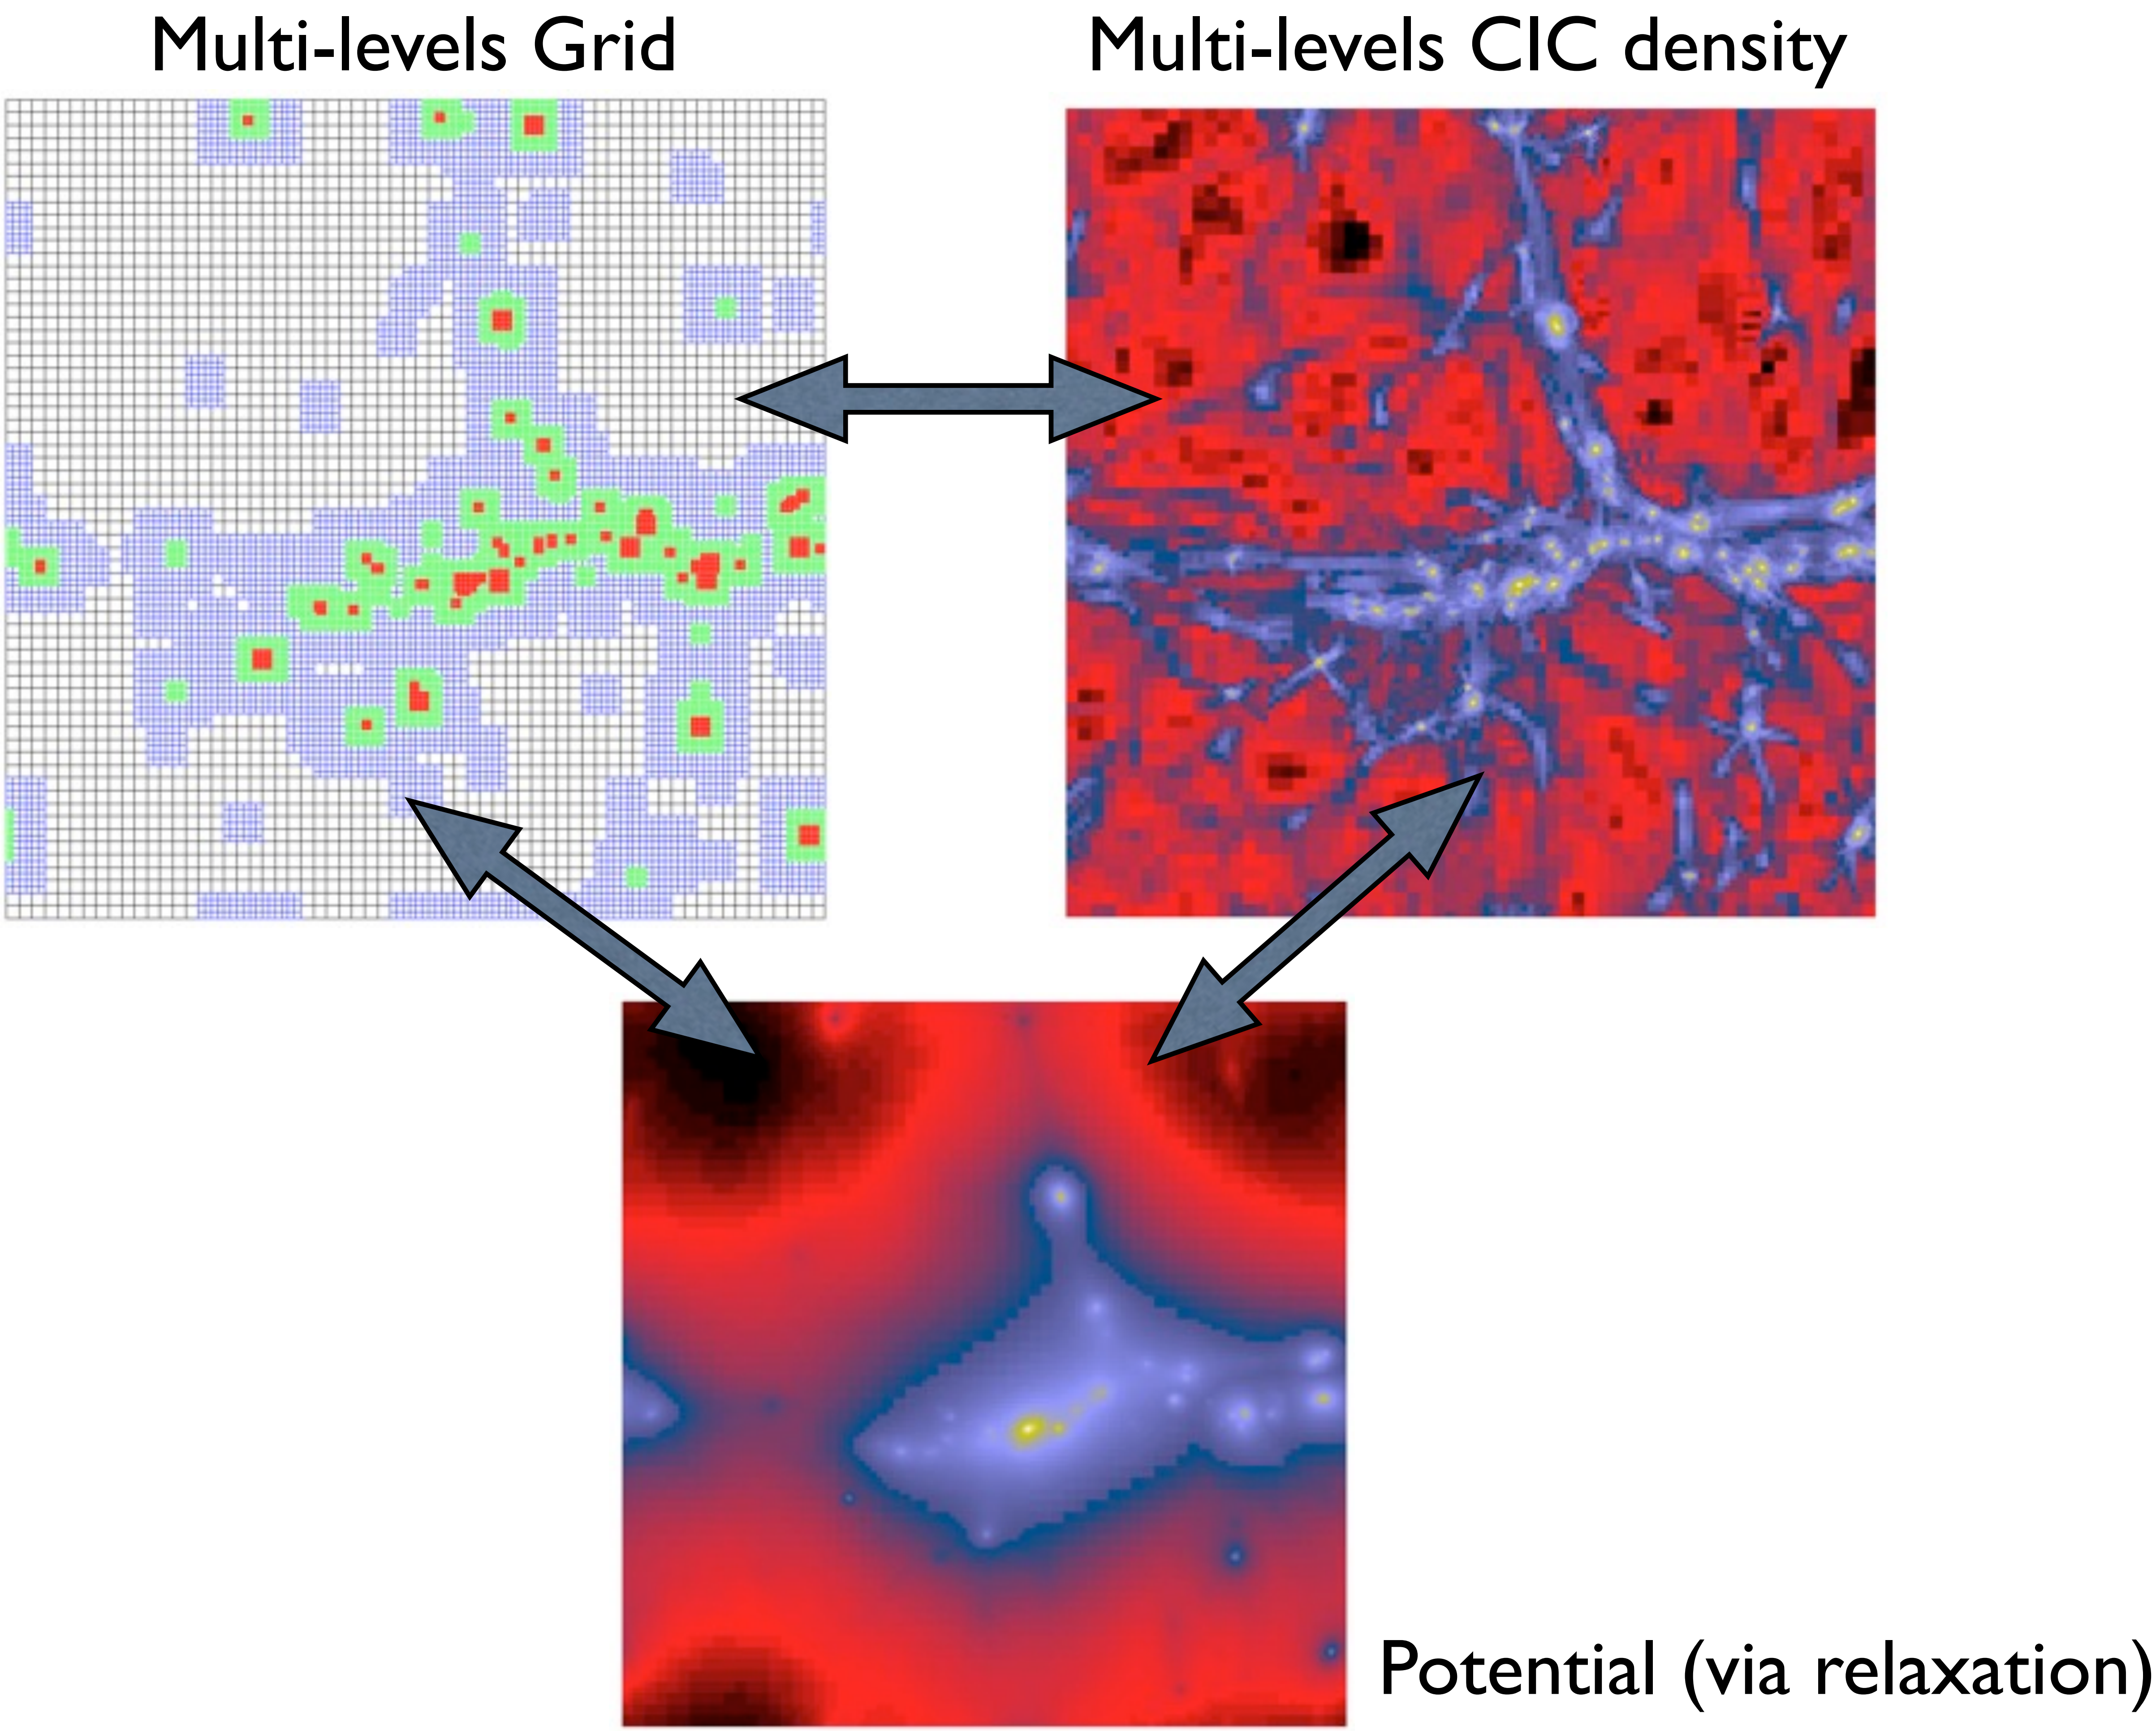
Multi-levels Grid
Multi-levels CIC density
Potential (via relaxation)

«Vectorization»

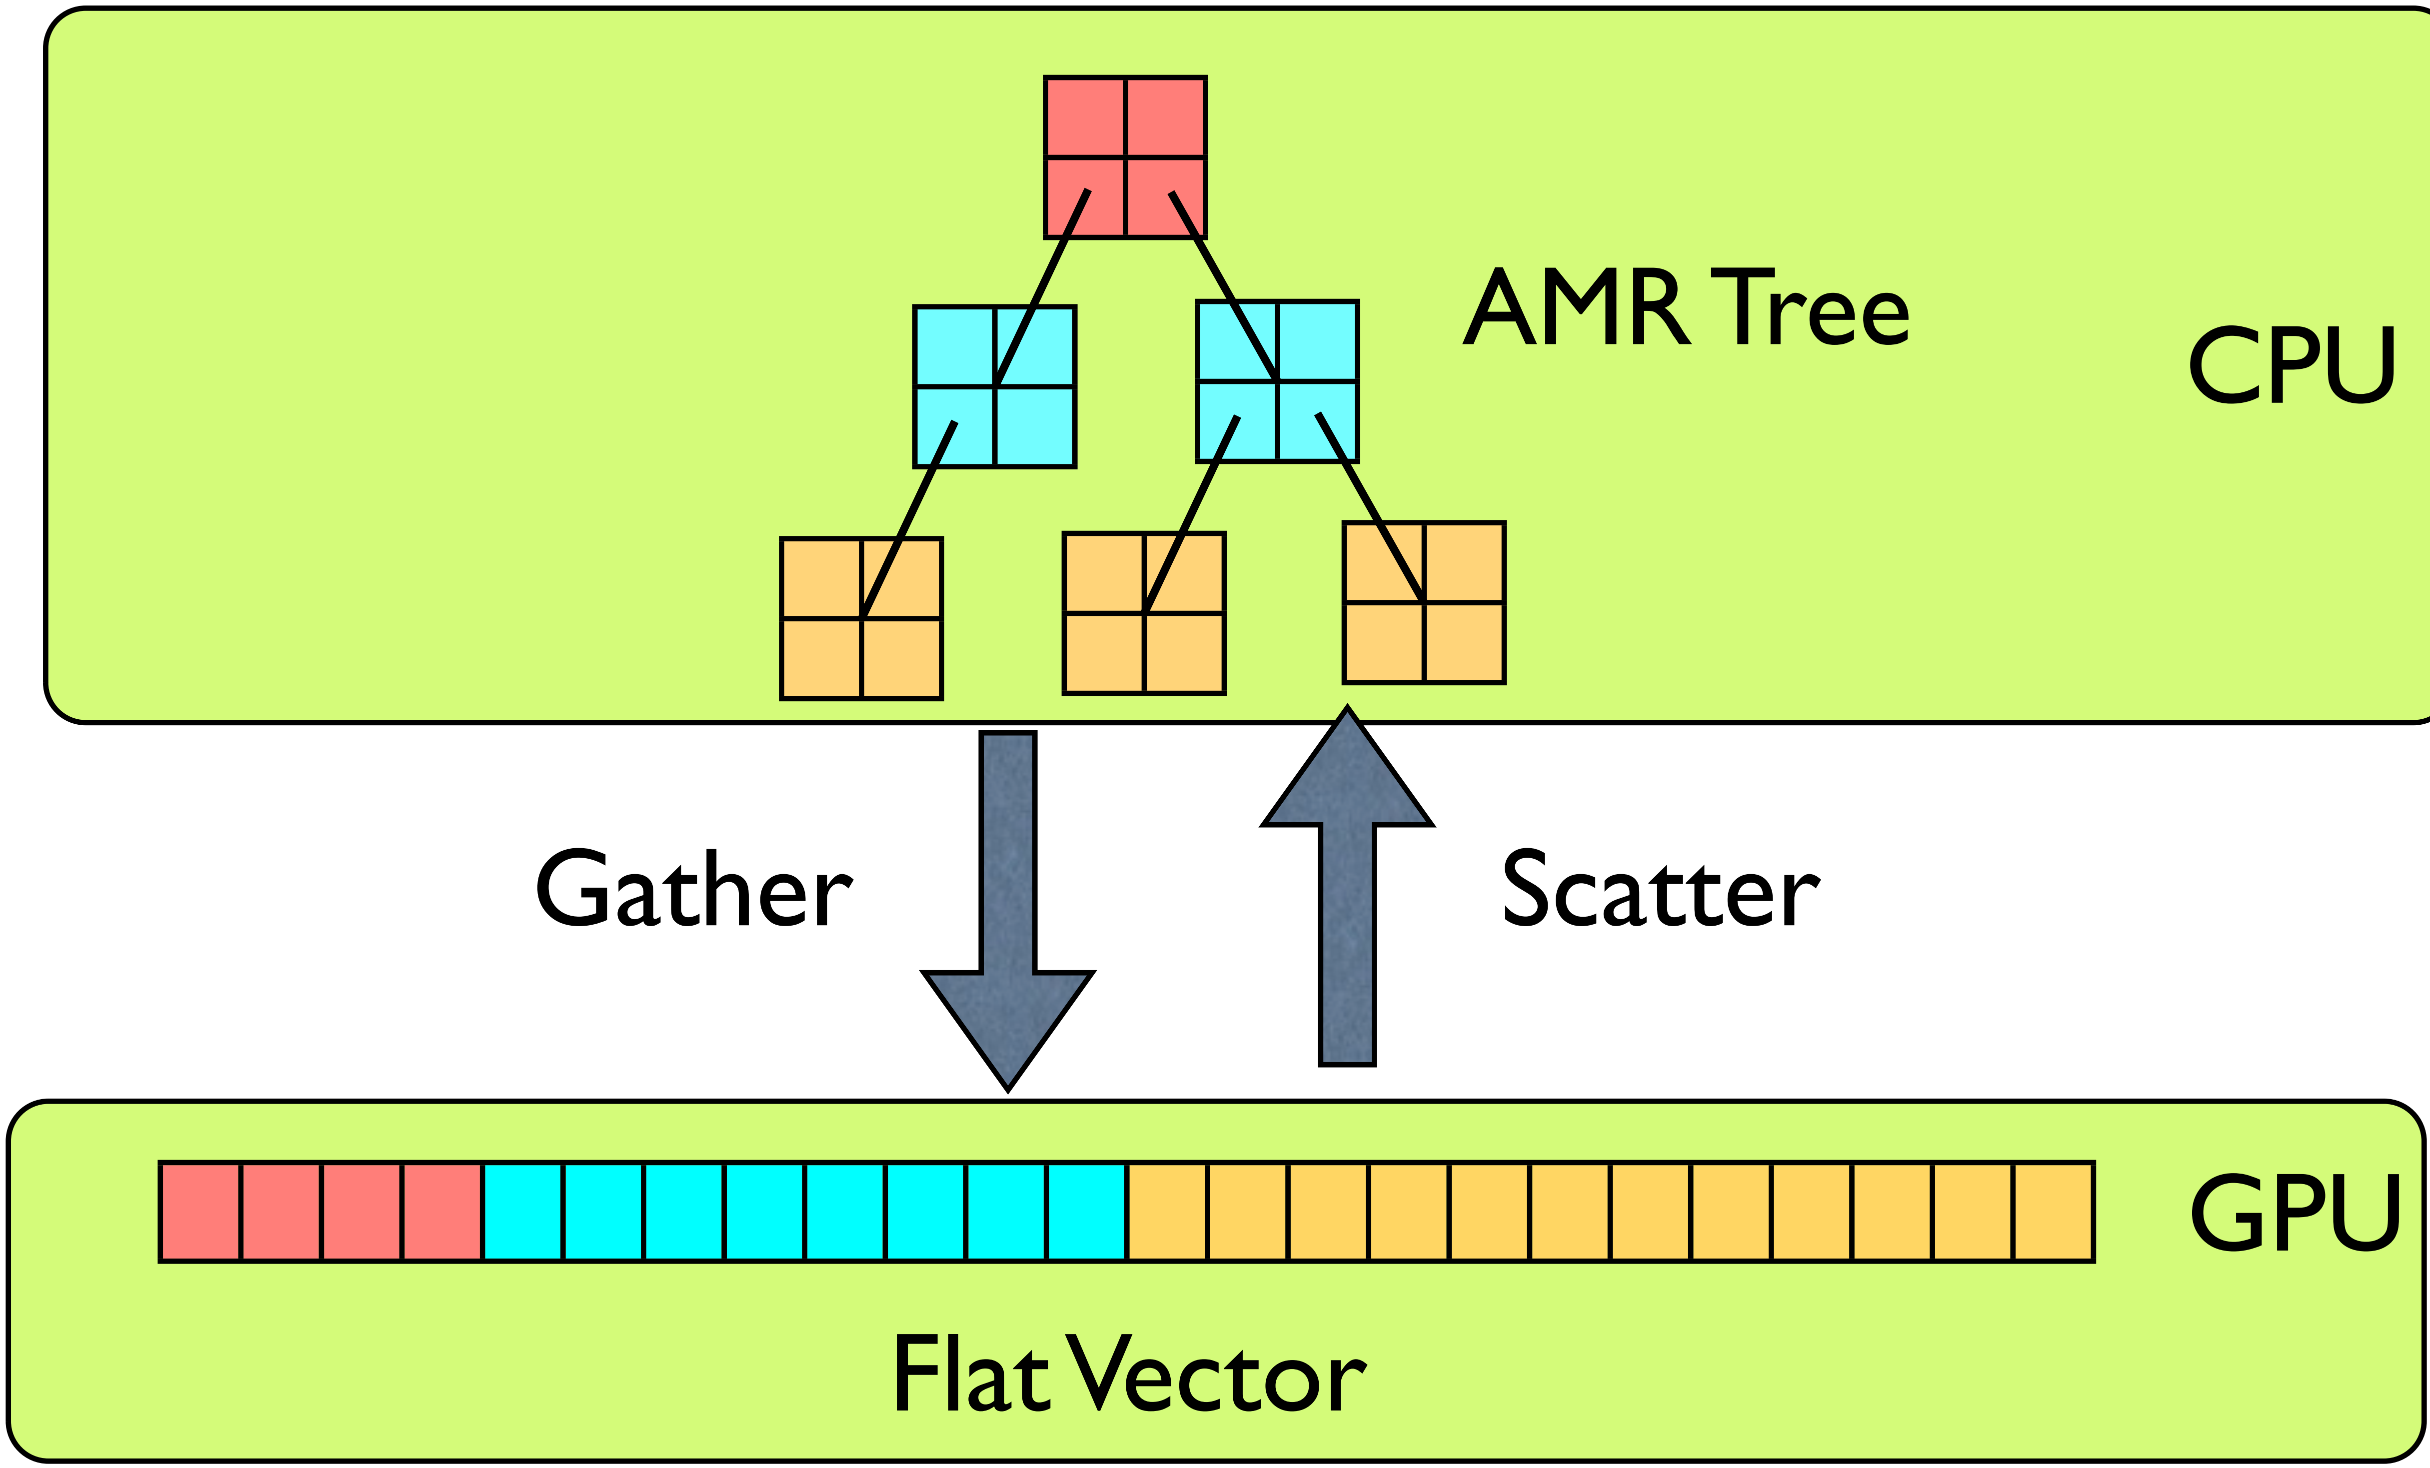

Leads to a bottle neck. Patch based AMR may be more appropriate
(see e.g. Schive et al. 2009)

How do we vectorize ?

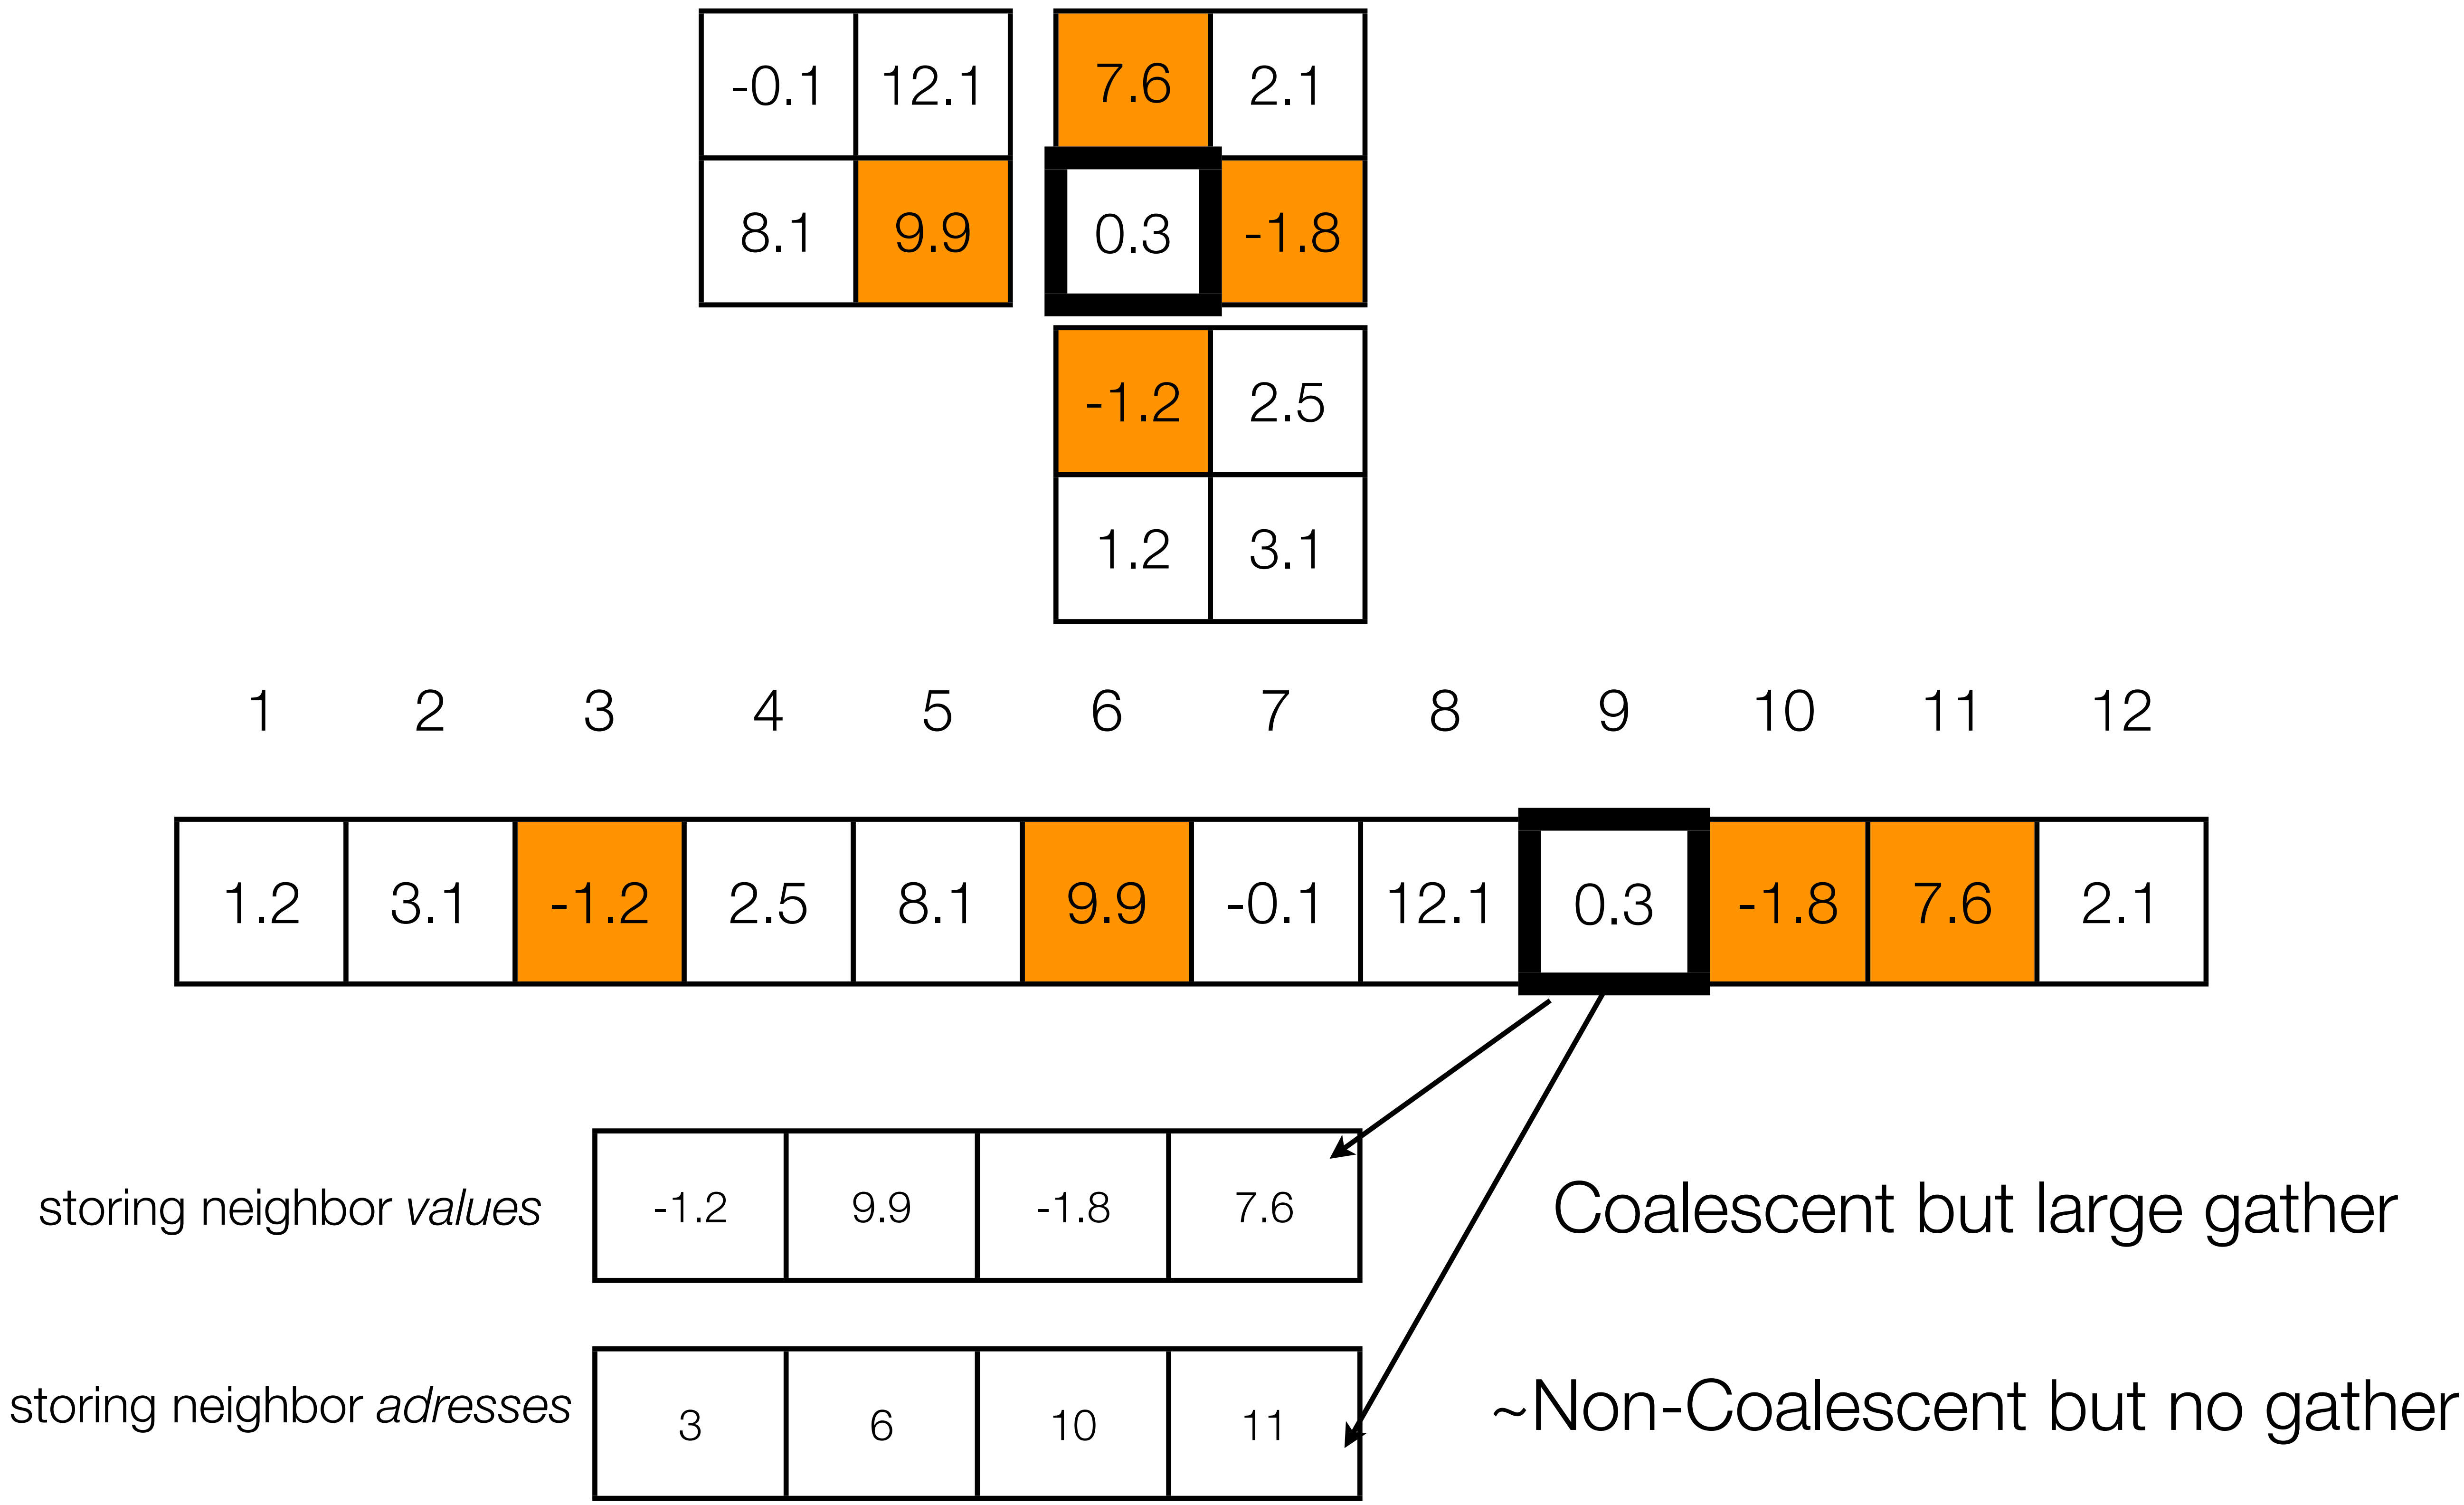

AMR issues with Explicit formulation

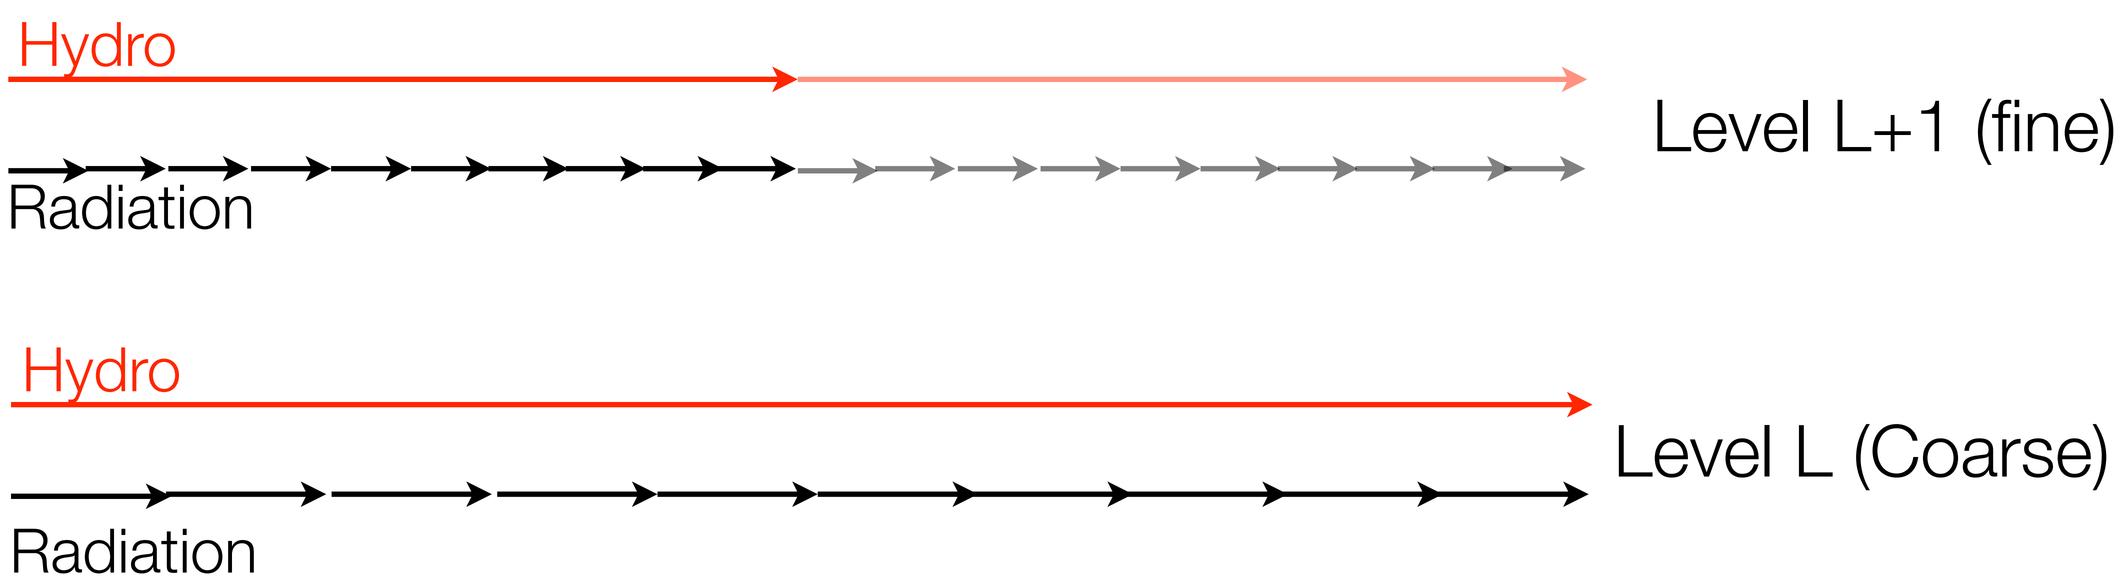

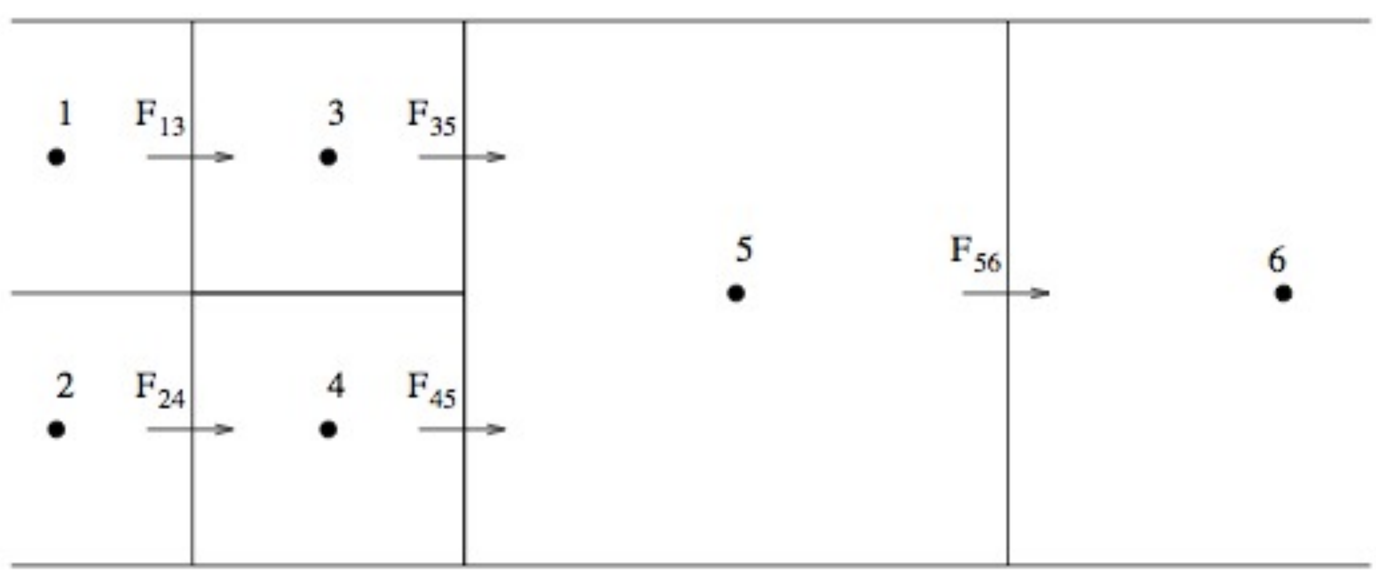

Subcycling induce problematic inter-levels interaction

It forces the hydro to be synchronized with radiation
E.g Rosdahl & Blaizot reduces the speed of light by 10-100 and synchronize the hydro on a small radiation timestep

Current Status

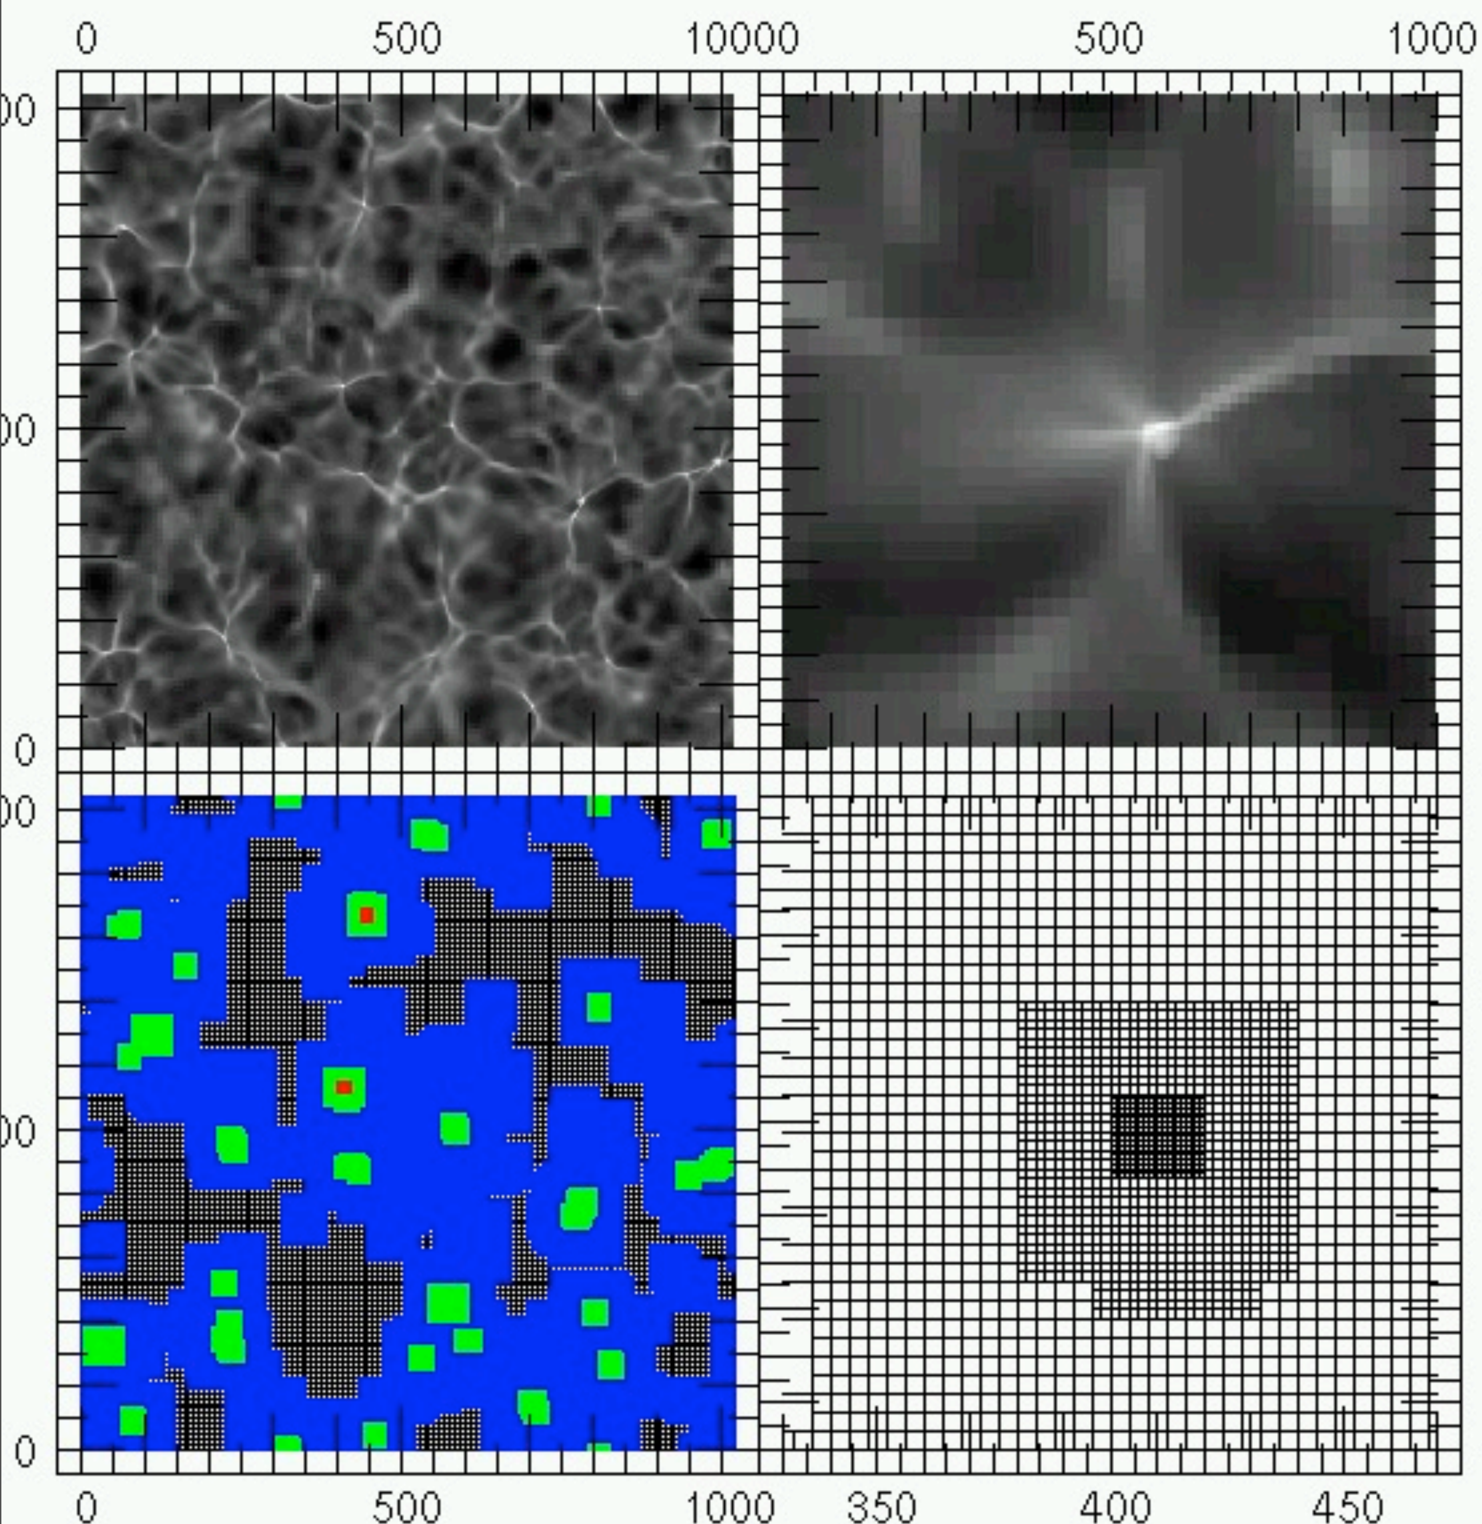

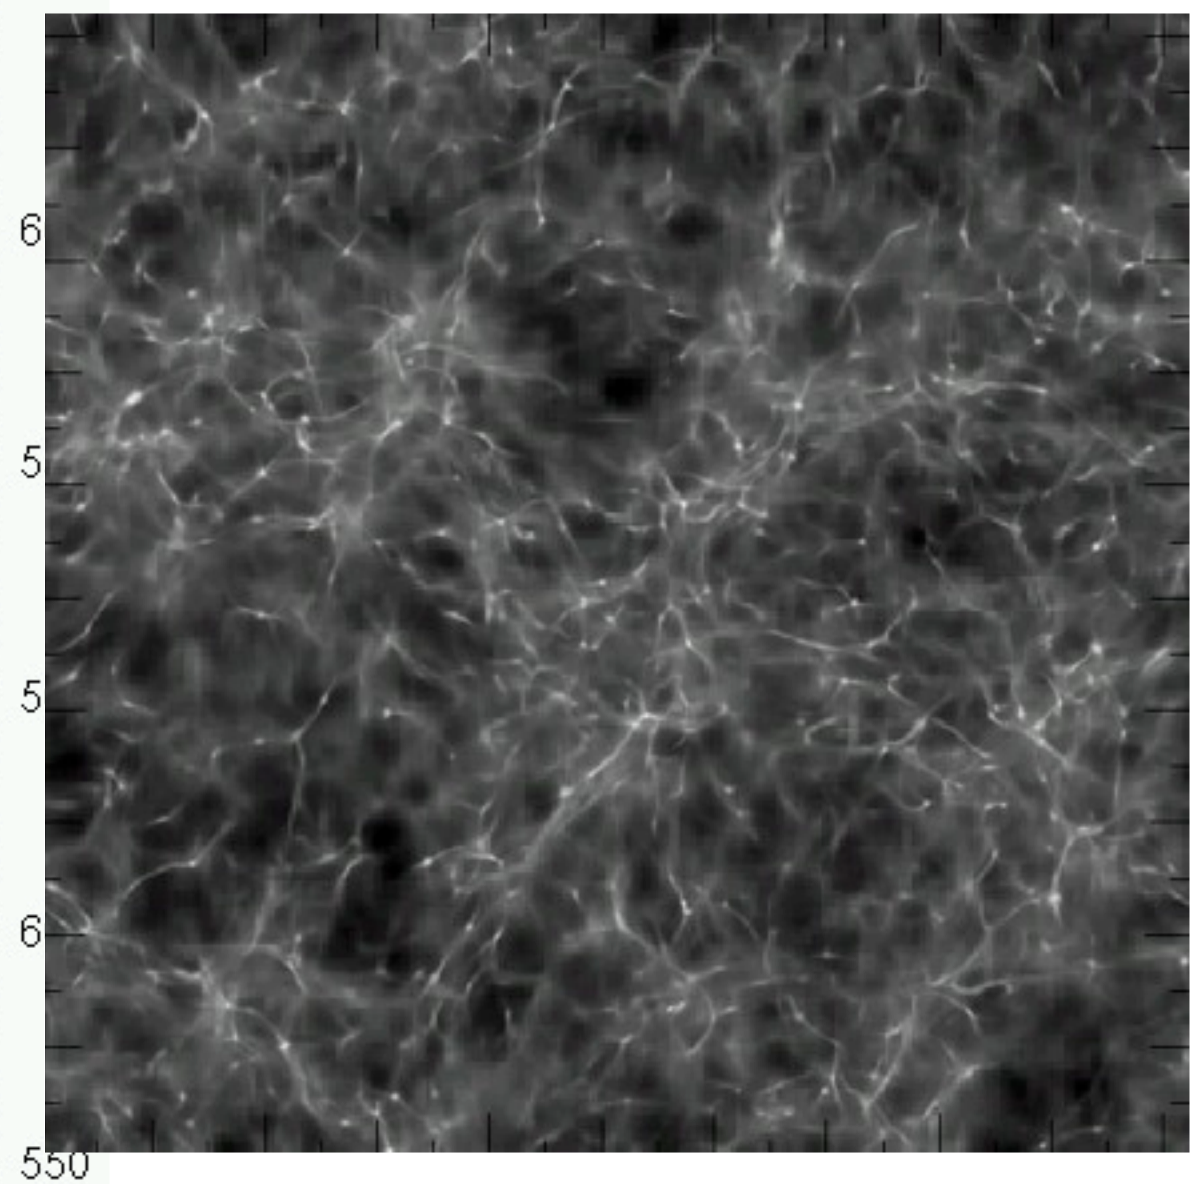

Without optimizations ~X10-15 (DP) compared to CPU for Hydro. RT might kill it or increase it...