Adaptive Mesh Refinement :
application to cosmic structures formation

R. Teyssier

CEA Saclay - Service d'Astrophysique - Groupe Théorie et Modélisation

$\mathcal{NIC}$ Consortium

Collaborators :

Yann Rasera (PhD strudent, CEA Saclay)

Edouard Audit (CEA Saclay)

J-P. Chièze (CEA Saclay)

Stéphane Colombi (IAP)

RAMSES : a 3D N-body + Hydro Adaptive Mesh Refinement code

- Tree-based AMR (octree structure) : the cartesian mesh is recursively refined *on a cell by cell basis*.
- Full connectivity : each "oct" have direct access to neighboring parent cells and to children "octs". (memory overhead : 2 integers per cell).

$\longrightarrow$ Optimize the mesh adaptivity to complex geometries, but CPU overhead can be as large as 50%.

N body module :	Particle-Mesh method on AMR grids (similar to the ART code).
	Poisson equation solved using Conjugate Gradient and Multigrid.
Hydro module :	*Unsplit* second order Godunov method : Riemann solver with
	piecewise linear reconstruction (option : MUSCL or PLMDE).
Time integration :	Single time step or W cycle (fine levels subcycling)
Other	Cooling & UV heating, Zoom simulation technology
	MPI based parallel implementation $\longrightarrow$ *Space Filling Curves*

Data structure in RAMSES

Tree based : "Fully Threaded Tree" (Khokhlov 1998)

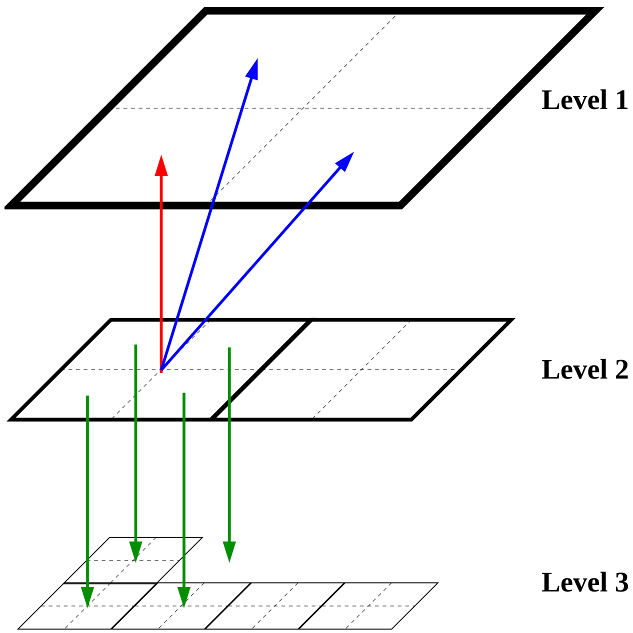

Basic Cartesian mesh

Recursive refinement on a *cell by cell basis*.

Fundamental objects : small grids of 8 cells or `oct`.

`octs` in the mesh are organized in a *linked list* for each level of refinement.

Full connectivity : each `oct` points towards

→ its "parent" cell

→ its 6 neighboring "parent" cells

→ its 8 "children" `octs`

2 disctinct types of cell :

- "leaf" cell or *active*.
- "split" cell *passive*.

Memory overhead :

- tree management : 2 integers per cell.
- passive cells : 14% in 3D, 33% in 2D, 50% in 1D

Adaptive refinement algorithm

Computing the refinement map : `flag=0 or 1`

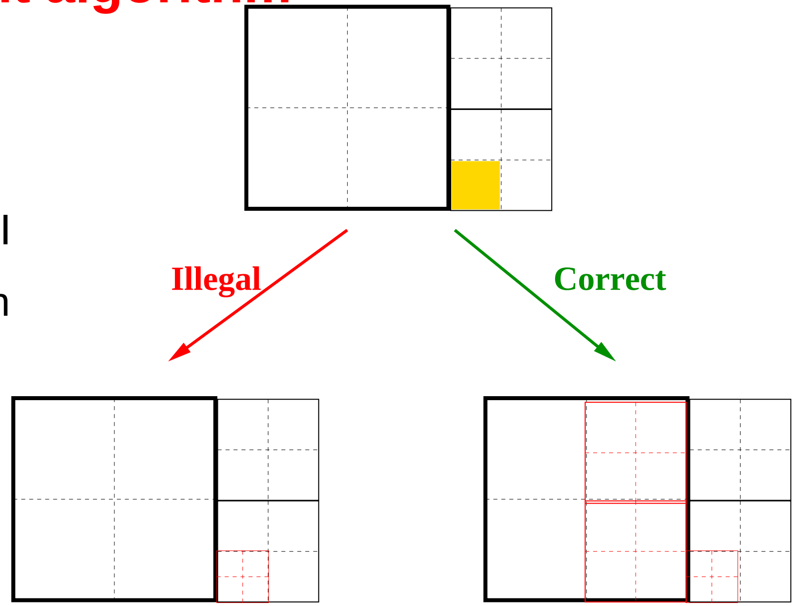

Step 1 : mesh consistency

if a split cell contains at least one split cell or one marked cell, then mark the cell with `flag=1`, and mark its 26 neighbors.

Step 2 : physical criteria

- high numerical dissipation : $\partial_x U$ or $\partial_x^2 U$
- geometrical constraints (zoom)
- quasi-Lagrangian evolution

Step 3 : mesh smoothing

apply a "dilatation" operator to the region marked with `flag=1` $\rightarrow$ *convex hull*.

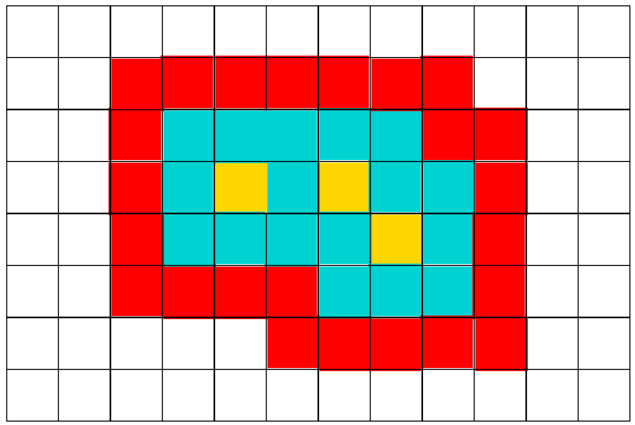

Hyperbolic systems of conservation laws

$\partial_t \mathbf{U} + \partial_x \mathbf{F}(\mathbf{U}) = 0$

Jacobian matrix $\mathbf{A}(\mathbf{U}) = \frac{\partial \mathbf{F}}{\partial \mathbf{U}}$

has real eigenvalues

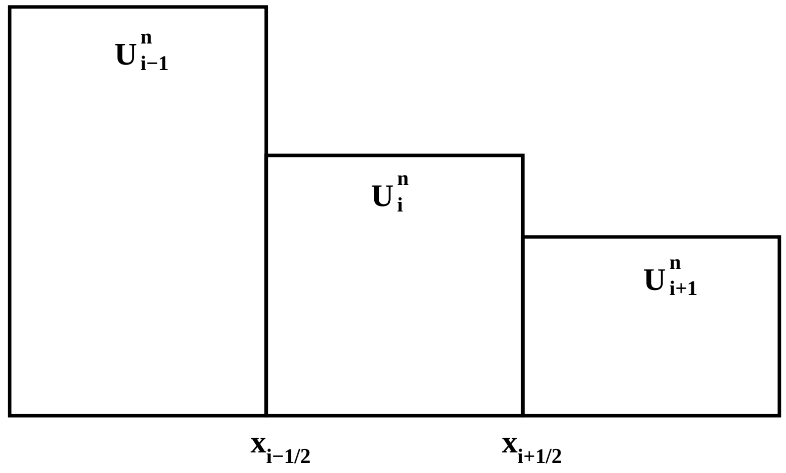

First order Godunov method :

Pieceweise constant initial data :

– *exact* Riemann problem at each interface

– compute new state as

$\mathbf{U}_i^{n+1} = < \mathbf{U}^{n+1}(x) >$

over interval $[x_{i-\frac{1}{2}}, x_{i+\frac{1}{2}}]$

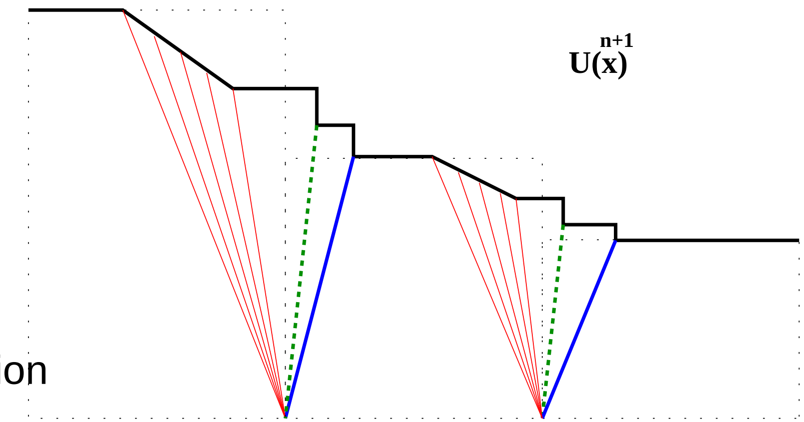

Alternative method :

$$\mathbf{U}_i^{n+1} = \mathbf{U}_i^n + \frac{\Delta t}{\Delta x}\left(\mathbf{F}^*_{i+\frac{1}{2}} - \mathbf{F}^*_{i-\frac{1}{2}}\right)$$

Higher order schemes :

Pieceweise linear (parabolic) TVD reconstruction

$\longrightarrow$ *approximate* second order schemes

Multigrid Godunov schemes

Prolongation (interpolation) to finer levels :

– fill buffer cells $\longrightarrow$ BC for fine levels

– refine $\longrightarrow$ create new `octs`

Restriction (averaging) to coarser levels :

– flux correction at coarse-fine boundaries

– de-refine $\longrightarrow$ average down the solution

Question : choice of interface variables ?

Constraint : $R^T P = I$

Several strategies :

– straight injection

– linear (parabolic) TVD reconstruction

Time integration : recursive sub-cycling

– froze coarse level during fine level solves

– average fluxes in time at coarse-fine boundaries

Solve for fine fluxes using buffer regions

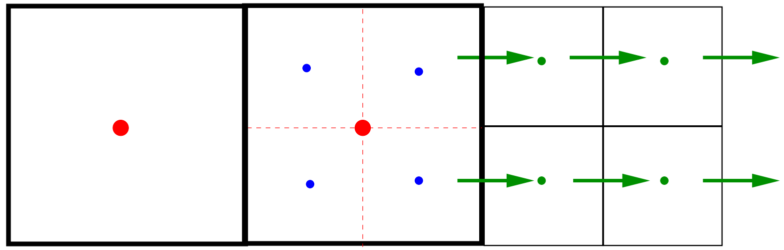

Coarse flux: time and space average of fine fluxes

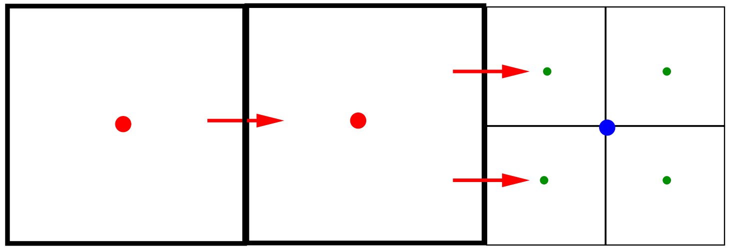

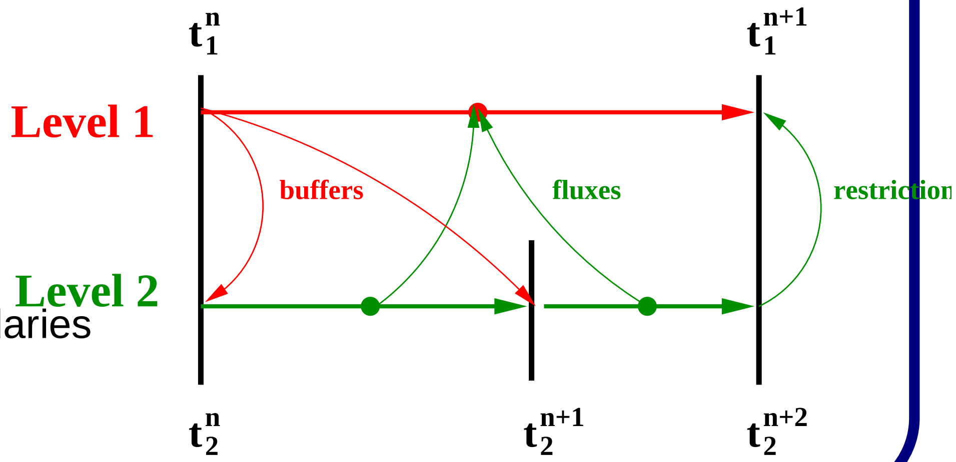

Adaptive time steps and 3D AMR schemes

High-order schemes use *non-linear* slope limiters to ensure stability and positivity

Robust high-order schemes (PPM, PLM) do exist in one space dimension on Cartesian mesh.

Generalisation of TVD properties to 3D AMR schemes ?

1. Directional splitting performs x-y-z sweeps to update conservative variables.

   TVD slopes are computed using new flow variables and therefore *depend on the chosen directional ordering*. Each directional sweep has to be performed for the whole hierarchy

   $\longrightarrow$ adaptive time steps are not permitted (ex : the FLASH code).

2. Consider *unsplit* schemes : Leveque's theorem applies : *"There are no fully 2D high-order scheme that is TVD in a 2D sense."*

   Solution : design multi-dimensional slope limiters that are *positivity preserving*.

   $\longrightarrow$ adaptive time steps are possible (ex : the RAMSES code).

Conclusion : adaptive time steps are robust for unsplit AMR schemes only.

N body and Poisson solver : the ART method (Kravtsov et al. 97)

Particle-Mesh scheme on AMR grids : cloud size $\rightarrow$ local mesh spacing

Poisson solver on the coarse gri

- Multigrid relaxation or FFT
- periodic or isolated

Poisson solver on the fine grids

- Conjugate Gradient relaxation
- Linear interpolation to fill up

  Dirichlet boundary conditions
- *One-way* recursive call

Future work :

- Fast Multipole Method

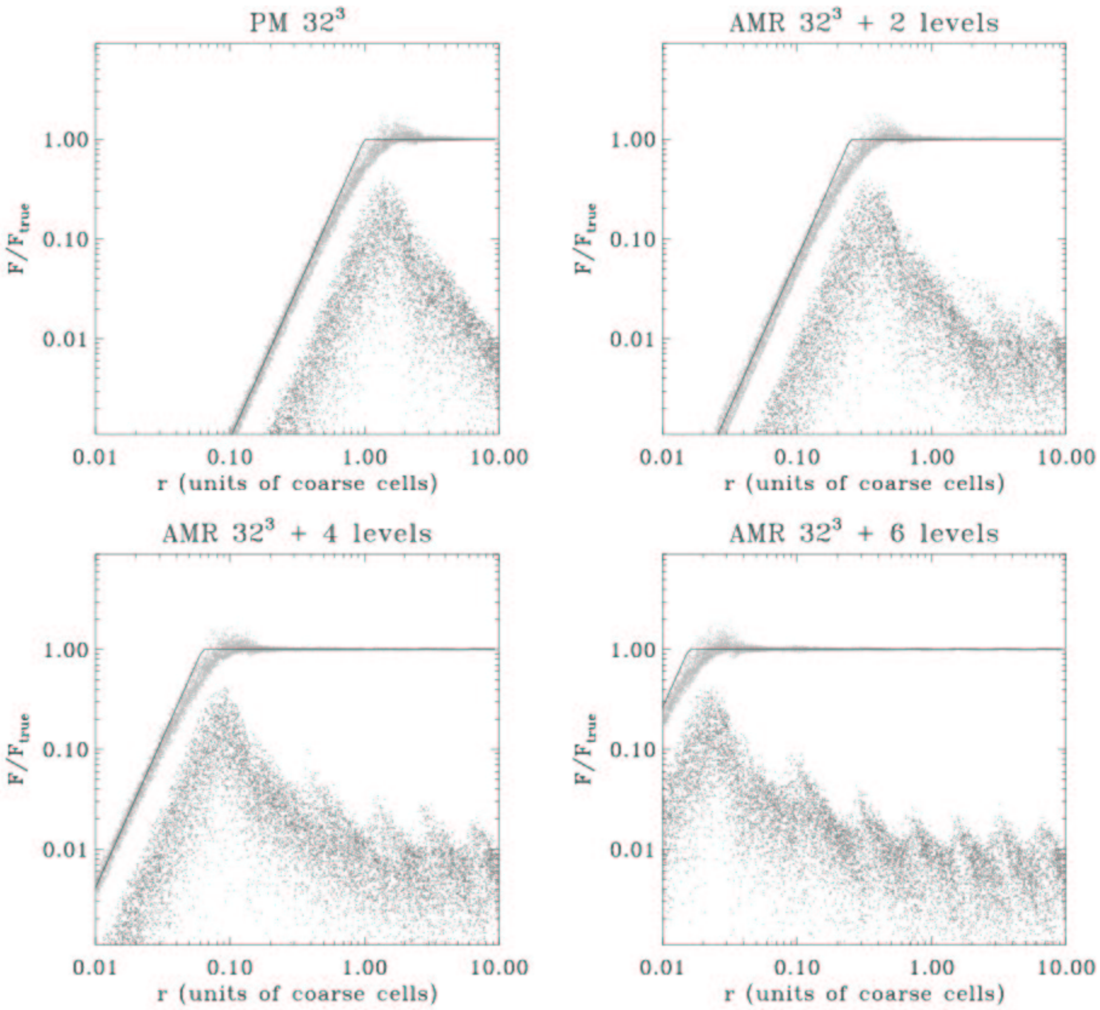

Shock tube test for SPH

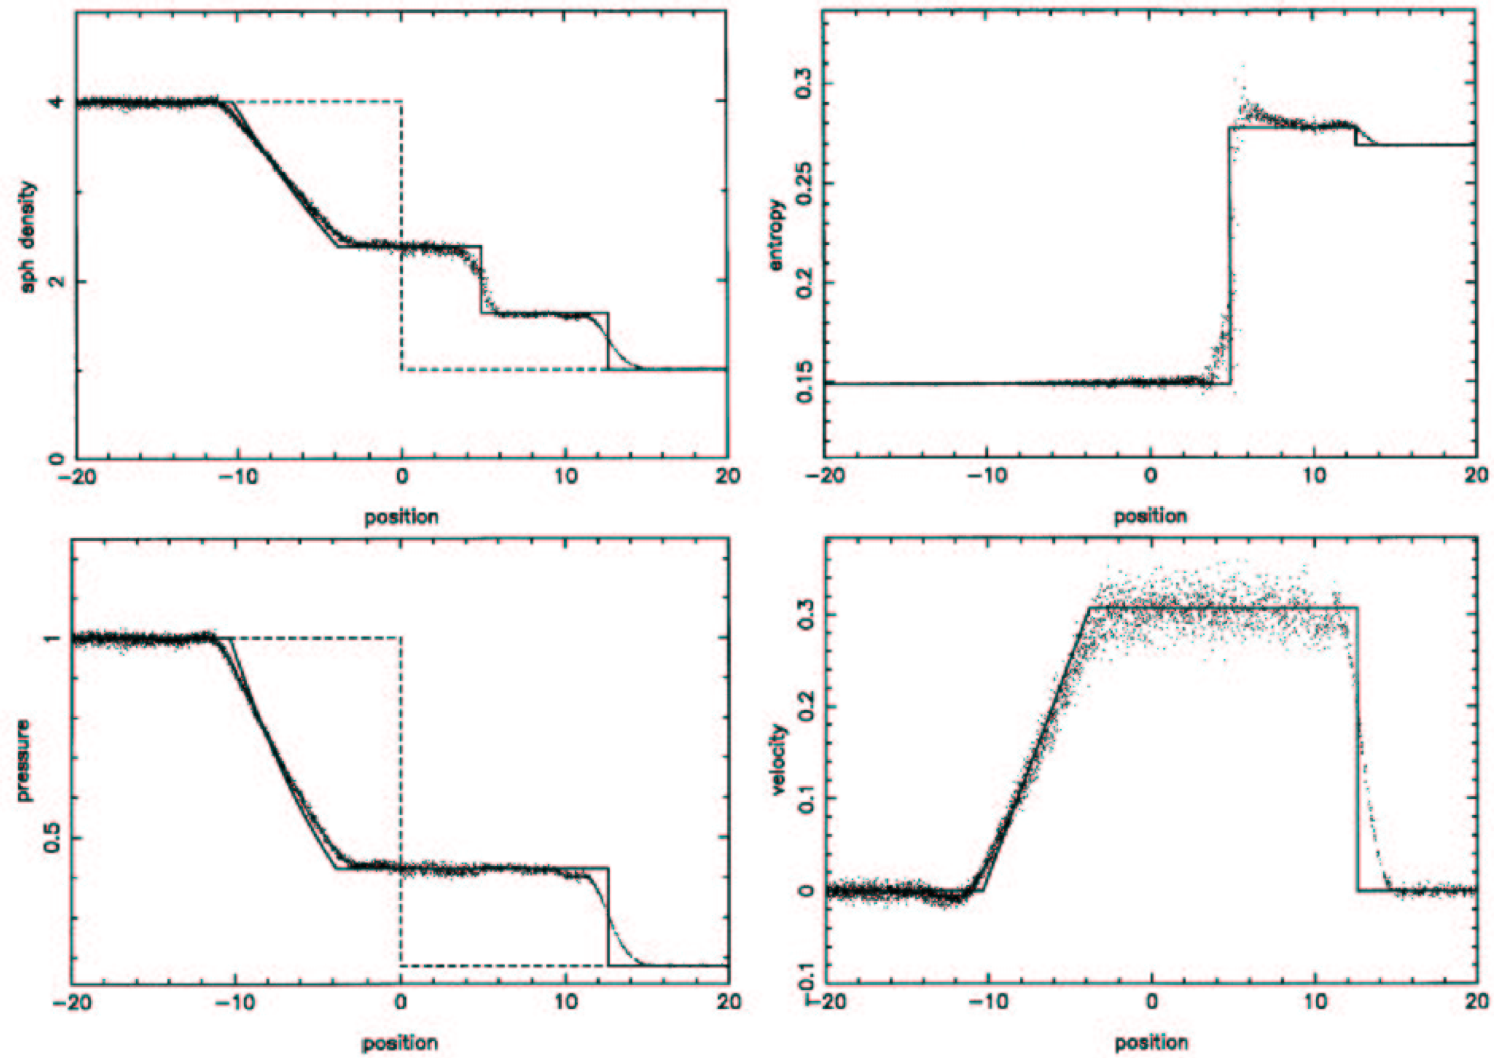

Ritchie & Thomas, MNRAS, 323, 743, (2001)

Same test for Godunov : a fair comparison

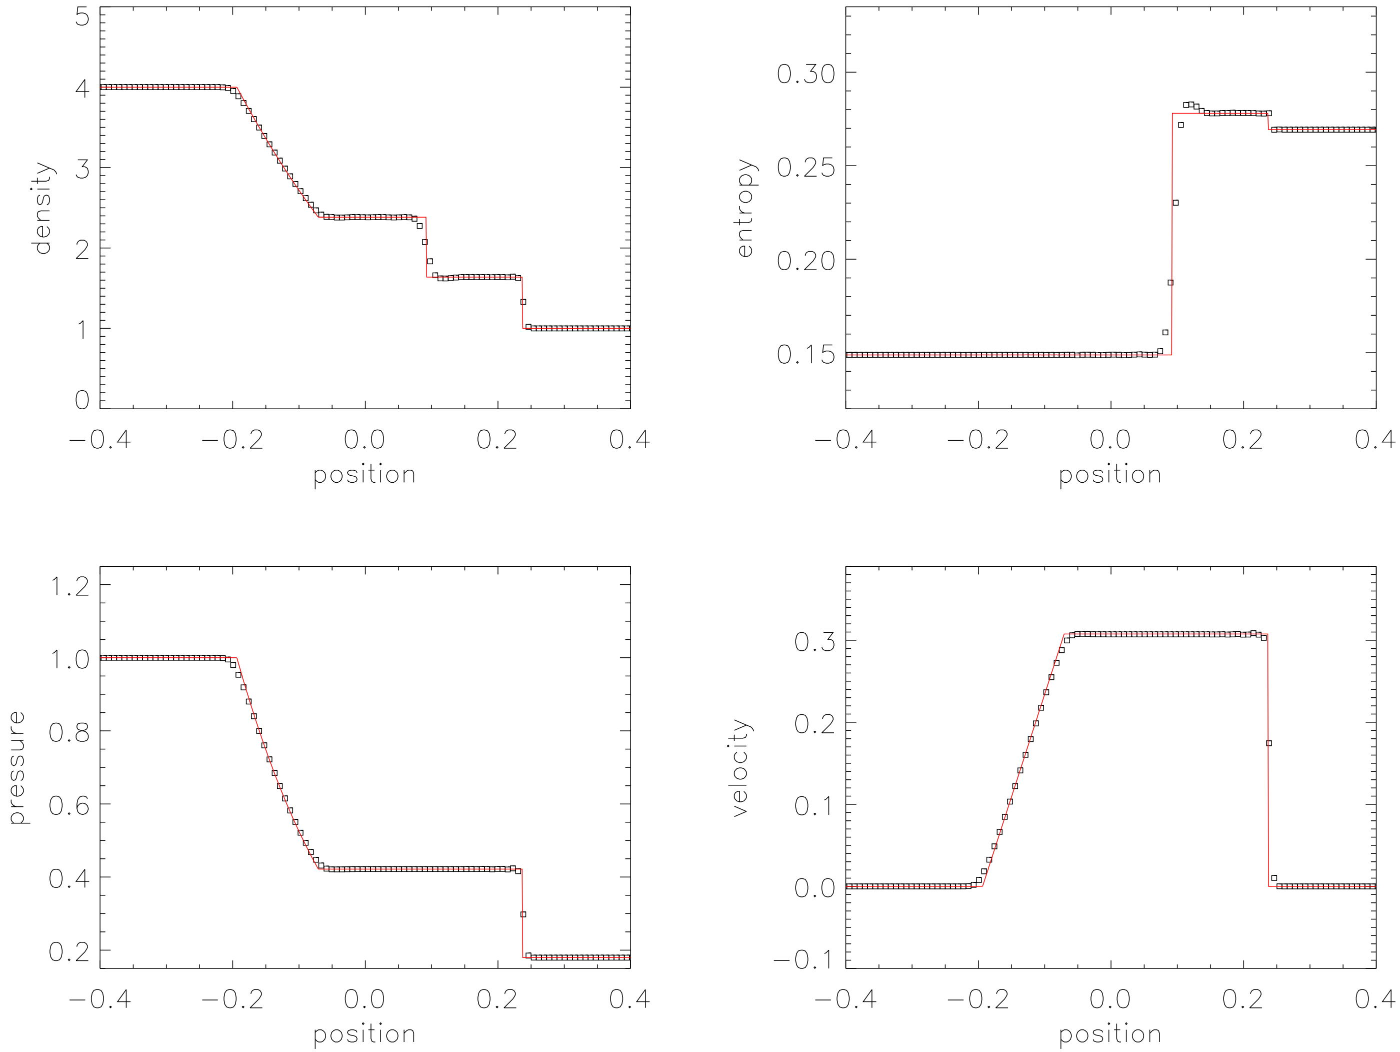

100 Eulerian mesh points with PLM

Same test for Godunov + AMR : an *unfair* comparison !

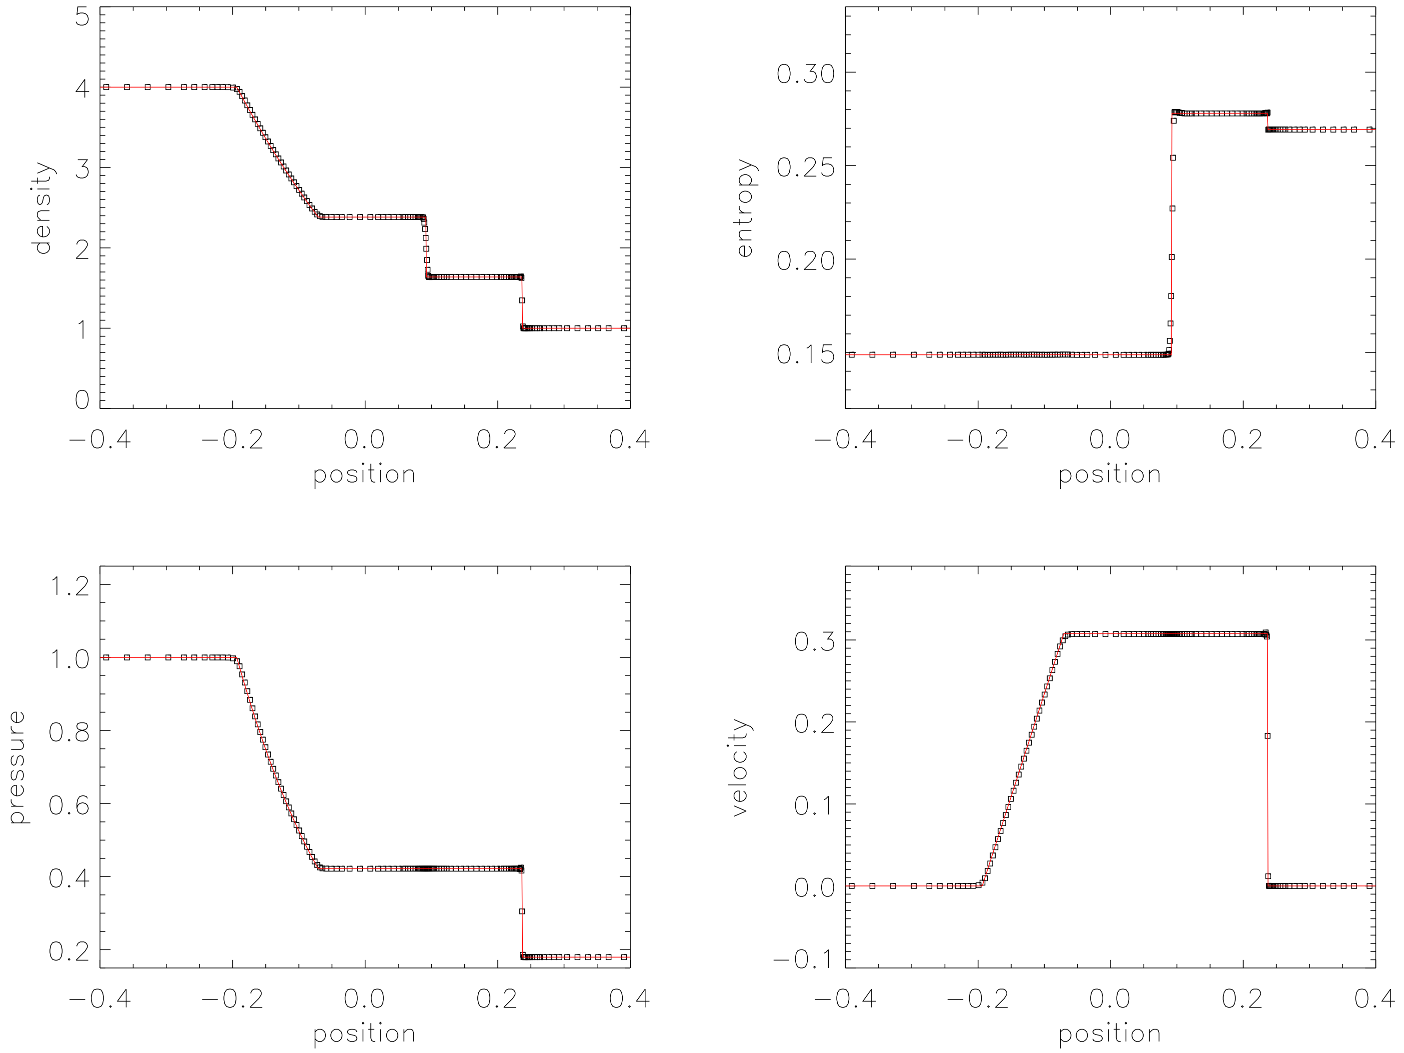

120 AMR mesh points with RAMSES

Rayleigh-Taylor instability : refining surface discontinuities

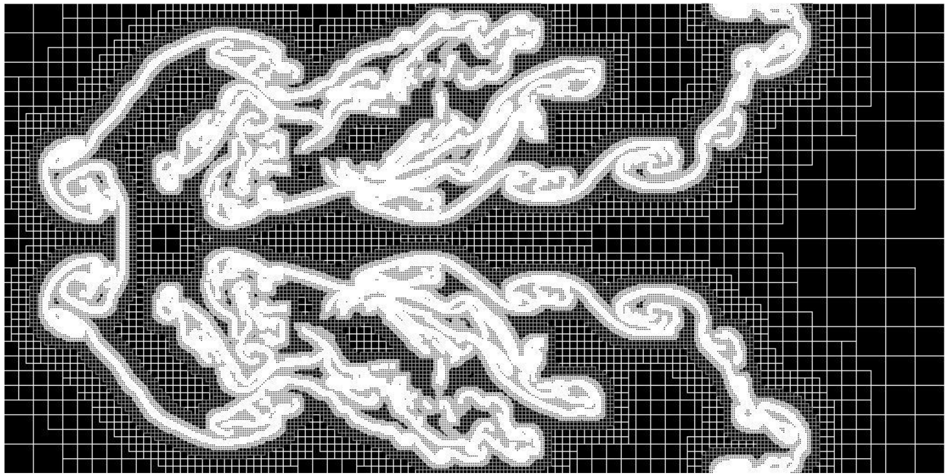

Maximum numerical dissipation occurs at the 2 fluids interface.

The optimal refinement strategy is based on density gradients.

The number of required cells is directly related to the *fractal exponent* n of the 2D surface.

$$N_{cell} \propto (\Delta x)^{-n}$$

Refinement strategy for cosmological simulations

- Quasi-Lagrangian mesh evolution (mimic SPH)

Do not refine any discontinuity (shocks) !

Compute for each cell

$$n = \frac{\rho_{DM}}{m_{DM}} + \frac{\rho_{gas}}{m_{gas}} + \frac{\rho_*}{m_*}$$

Trigger new refinements when $n > 10 - 40$.

- Re-simulation of a specific region : multiple mass particles and zooming AMR grid.
  - Initial conditions using the GRAFIC 2.0 package (Bertschinger 2001)
  - High mass particles in low resolution cells to describe large scale tidal fields
  - Low mass, high resolution particles around the chosen halo.
- Catastrophic refinement in fast cooling regions : need to carefully control the maximum refinement level
  - Compute the density thresholds dynamically according to density PDF
  - Add one level of refinement every now and then...

Shock propagation within the AMR grid

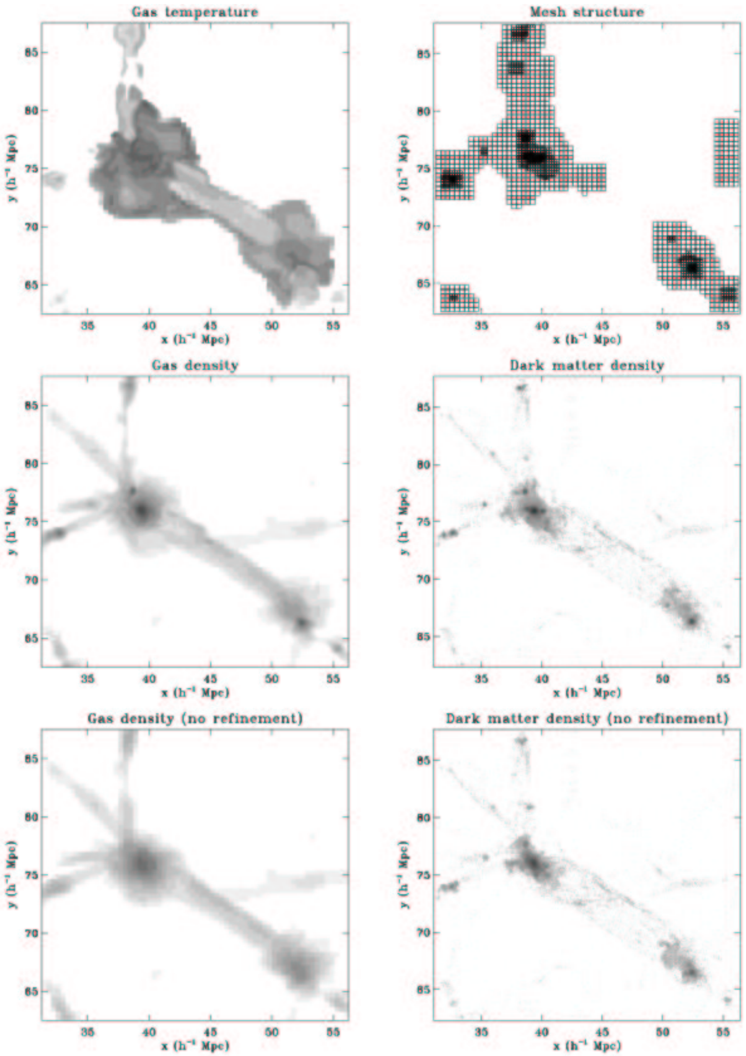

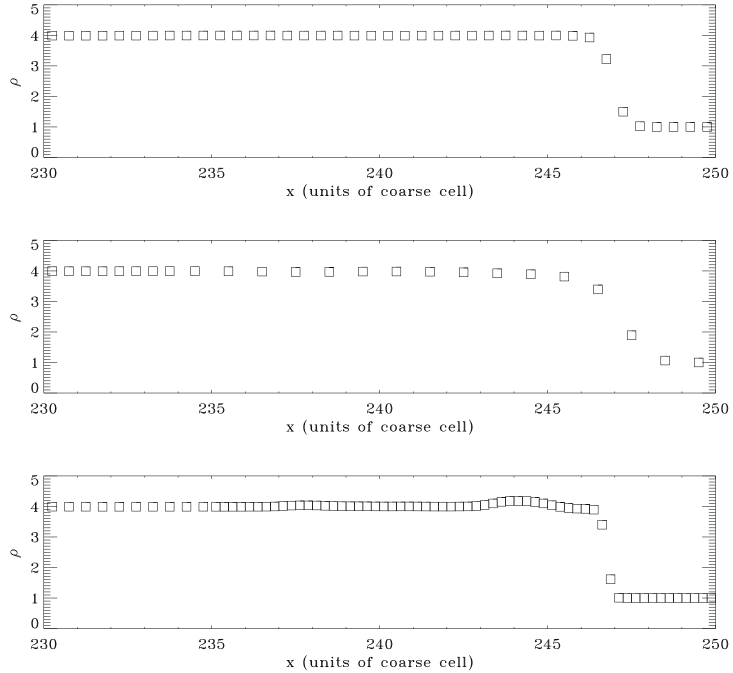

Strong shocks are unstable when travelling through *coarse-to-fine* interfaces.

Typical cosmological flows conspire to avoid this.

Weak shocks (within filaments) are unaffected.

Cosmology specific : stiffness

Euler equations with gravitational and cooling source terms :

$$\partial_t \rho + \partial_x \rho u = 0$$
$$\partial_t \rho u + \partial_x (\rho u^2 + P) = -\rho \partial_x \phi$$
$$\partial_t E + \partial_x u (E + P) = -\rho u \partial_x \phi - \rho \frac{\epsilon - \epsilon_{eq}}{t_{cool}}$$

Several characteristic time scales :

$$t_s = \Delta x / c_s \quad \text{versus} \quad t_{ff} = \sqrt{3\pi / 32 G \rho} \quad \text{versus} \quad t_{cool}$$

- Supersonic regime : $t_s \gg t_{ff}$.

Jeans length is underresolved : spurious heating in cold regions.

Solution : design clever fix for high-Mach flows.

- Overcooling regime : $t_s \gg t_{cool}$.

Implicit method for time integration.

Fragmentation length is underresolved : cooling catastrophe, spurious drag forces.

Solution : multiphase flow.

High Mach flow : hybrid scheme

For strong velocity gradients, fluid internal energy is dominated by truncation errors.

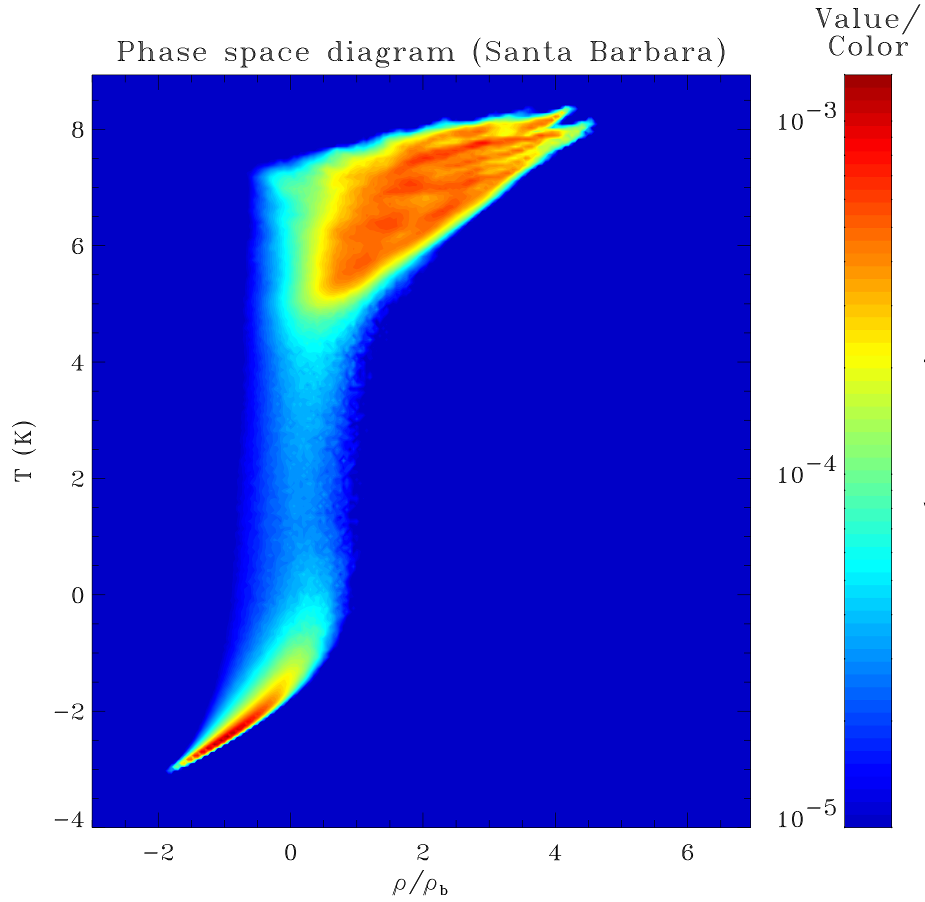

Conservative schemes may result into spurious (negative) temperature.

Switch-off shock heating, as long as internal energy lies below a *Galilean invariant* threshold : $\epsilon_c \simeq 0.5 \Delta u^2$ and use internal energy as new integration variable $\longrightarrow$ hybrid scheme

No shock heating below a given halo mass $M_{res} \simeq$ 50 particles.

Parallel computing with RAMSES

Domain decomposition using *mesh partionning* technics.

Inspired by parallel TREE codes (Zurek, Dubinsky, Springel...).

Local AMR tree surrounded by virtual buffer regions : *locally essential tree*.

Several cell ordering methods are implemented in RAMSES ⟶

1. column, row or plane major
2. Hilbert or Morton
3. user defined : angular, column + Hilbert...

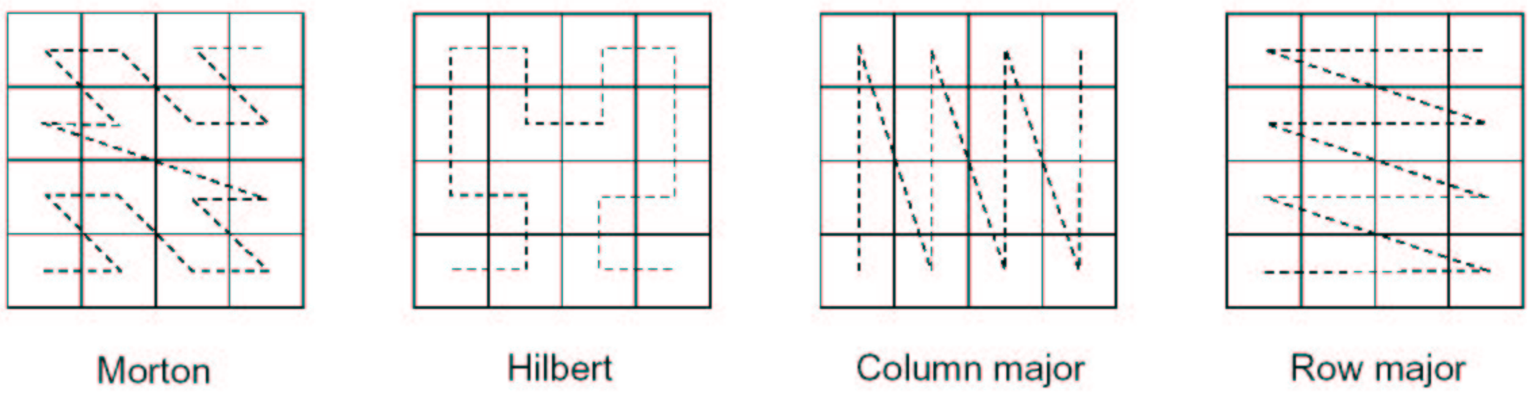

Dynamic partionning is performed by sorting each cell along chosen ordering.

3 different domain decomposition strategies

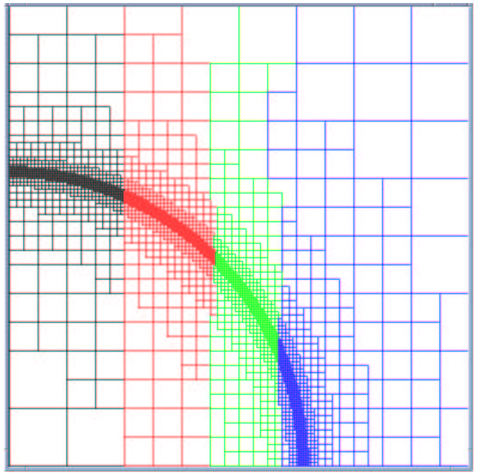

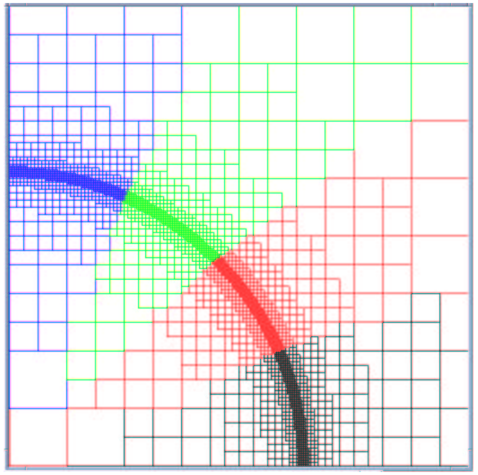

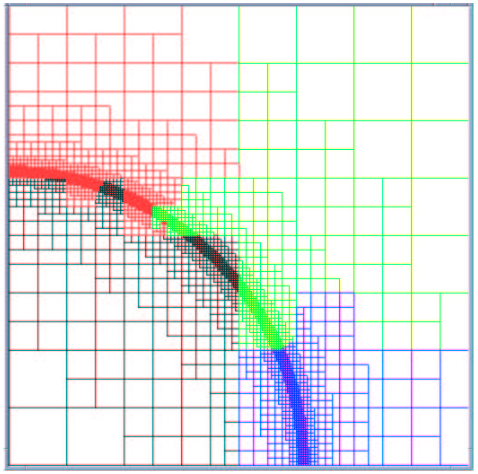

Hilbert space filling curve : an "optimal" ordering ?

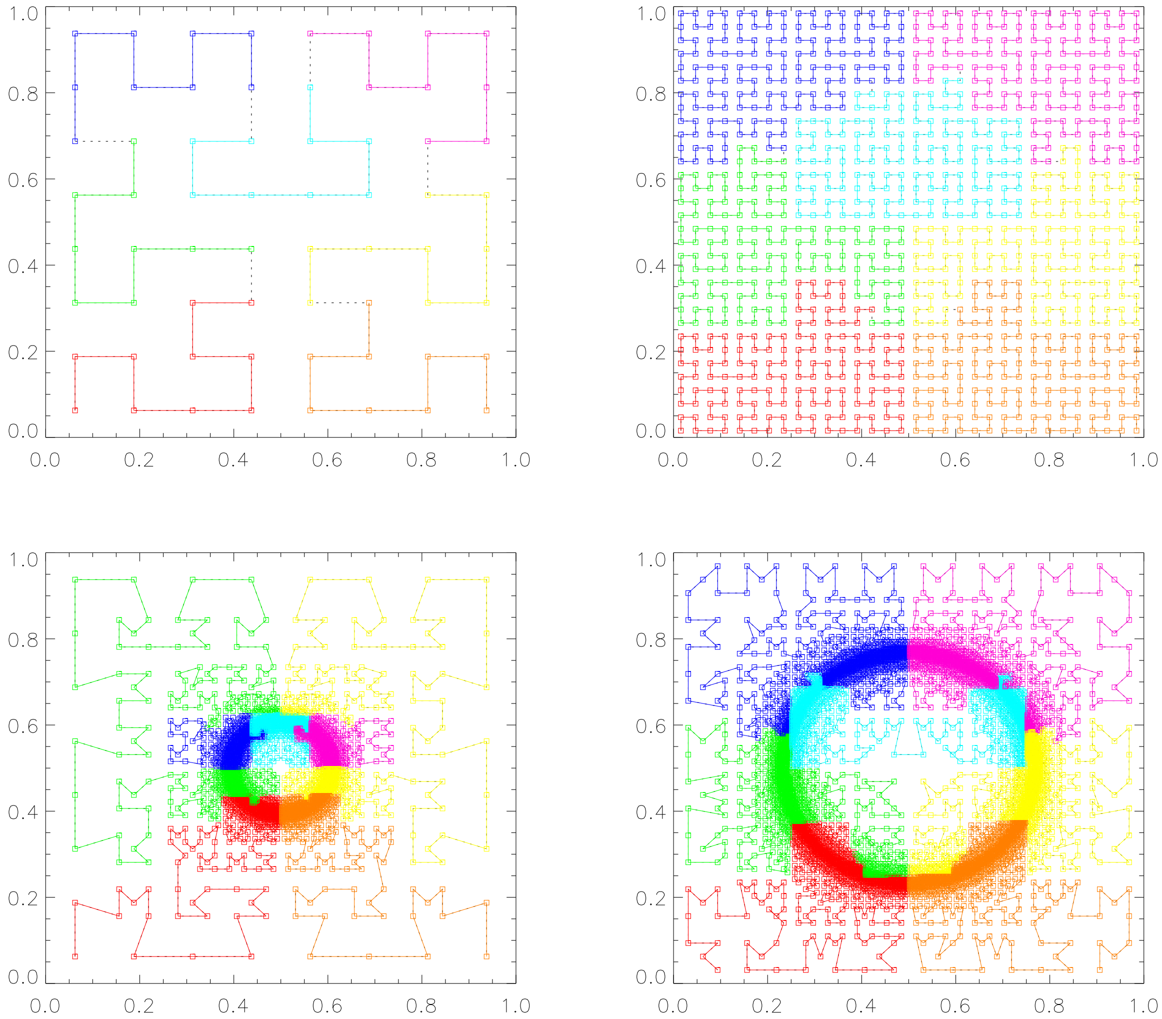

1024^3 particles pure N body 400 h^{-1}Mpc ΛCDM at z=1

256 HP ES45 processors

220 Go used memory

150 h up to z=0

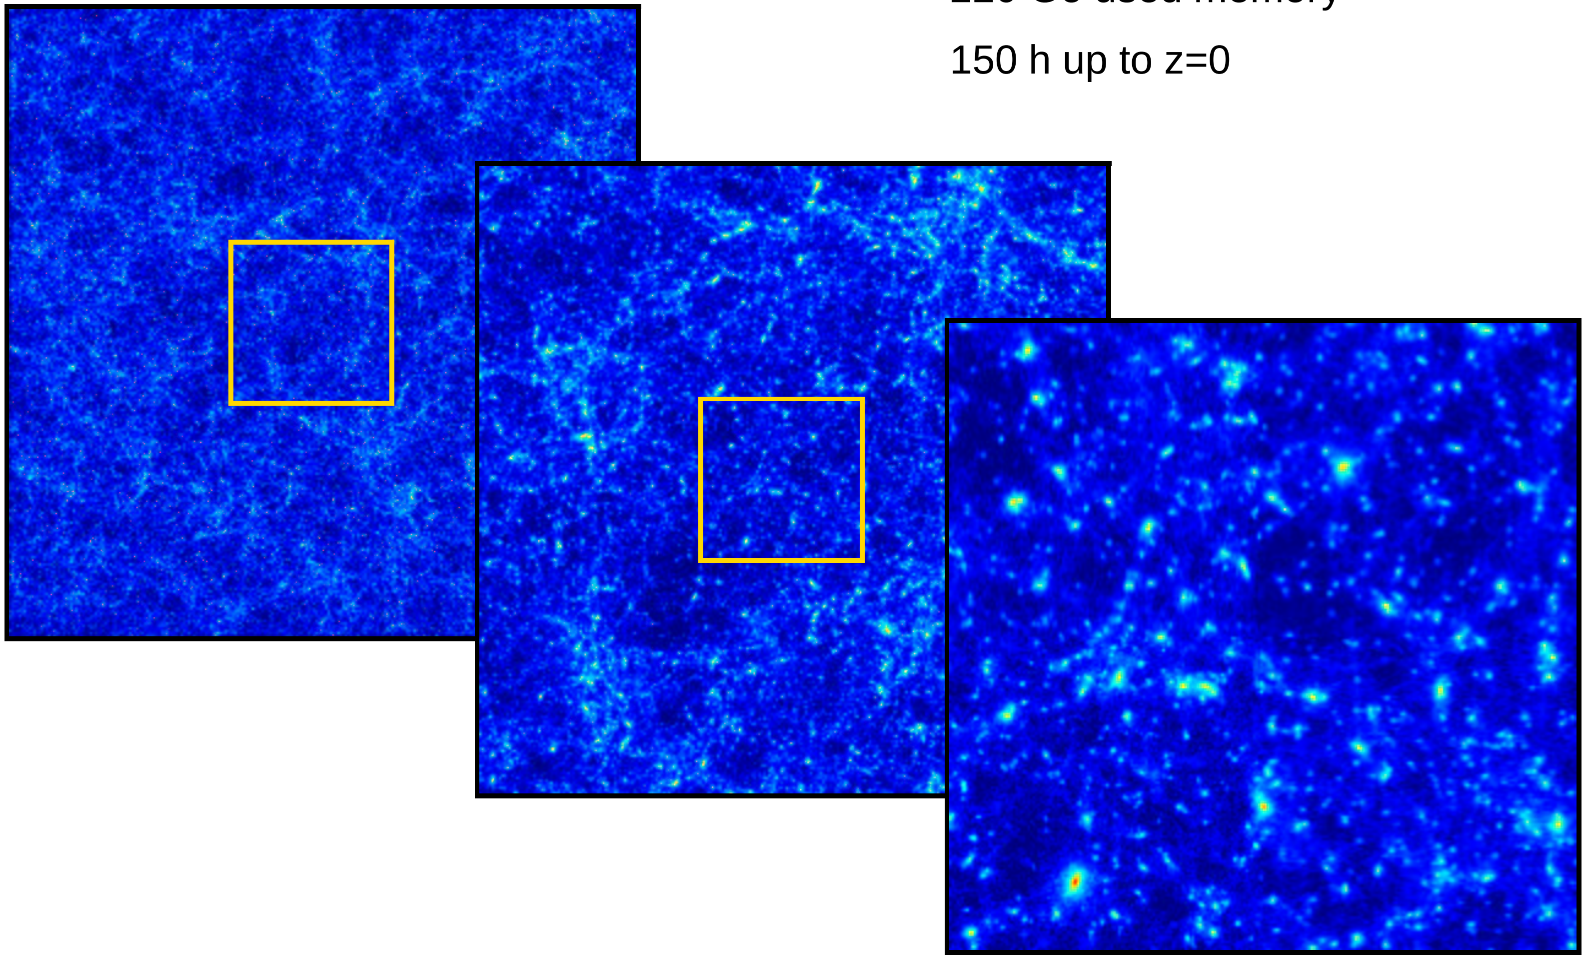

Good load balancing and memory control

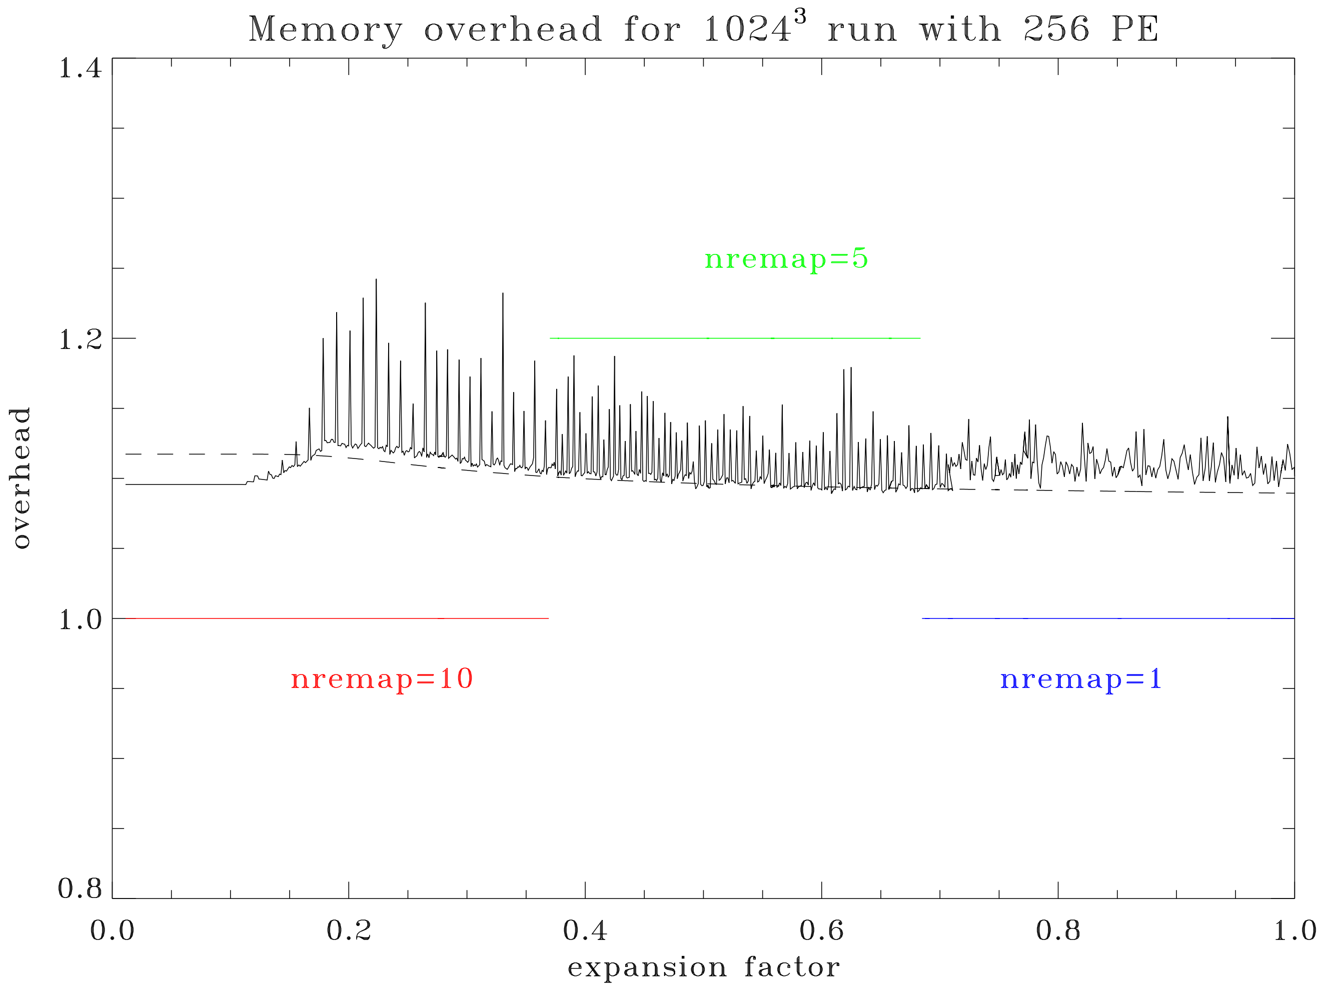

Performance overview using Santa Barbara cluster

256^3 particles

40×10^6 AMR cells

256^3 base grid

5 levels of refinement

32 processors (HP ES45)

Elapsed time : 100 h

Used memory : 12 Go

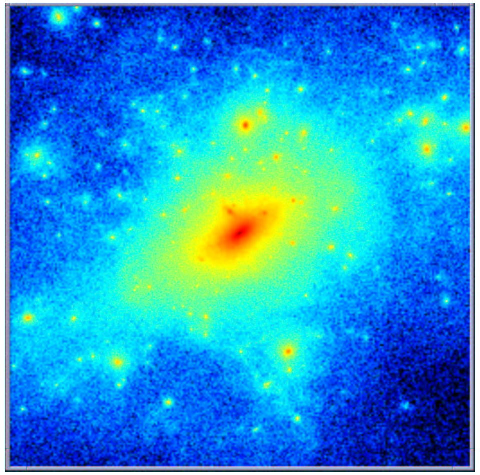

Dark matter density

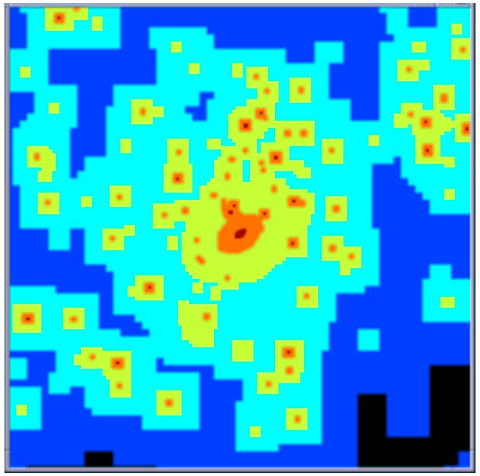

Refinement levels

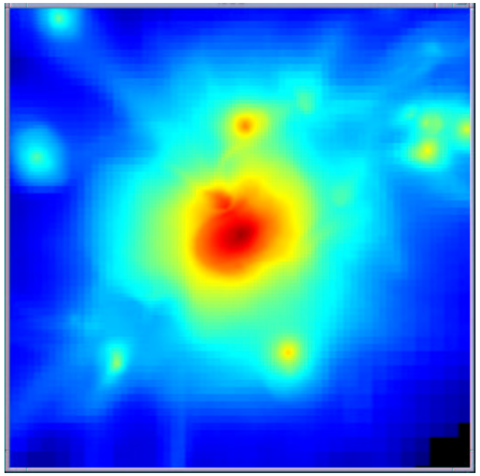

X-ray luminosity

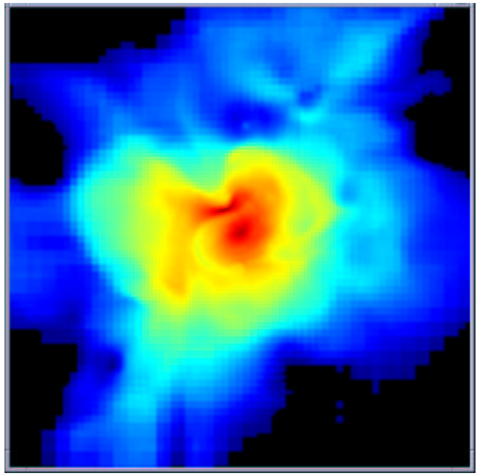

Gas temperature

A Coma-like cluster with increasing resolution : towards galactic scales ?

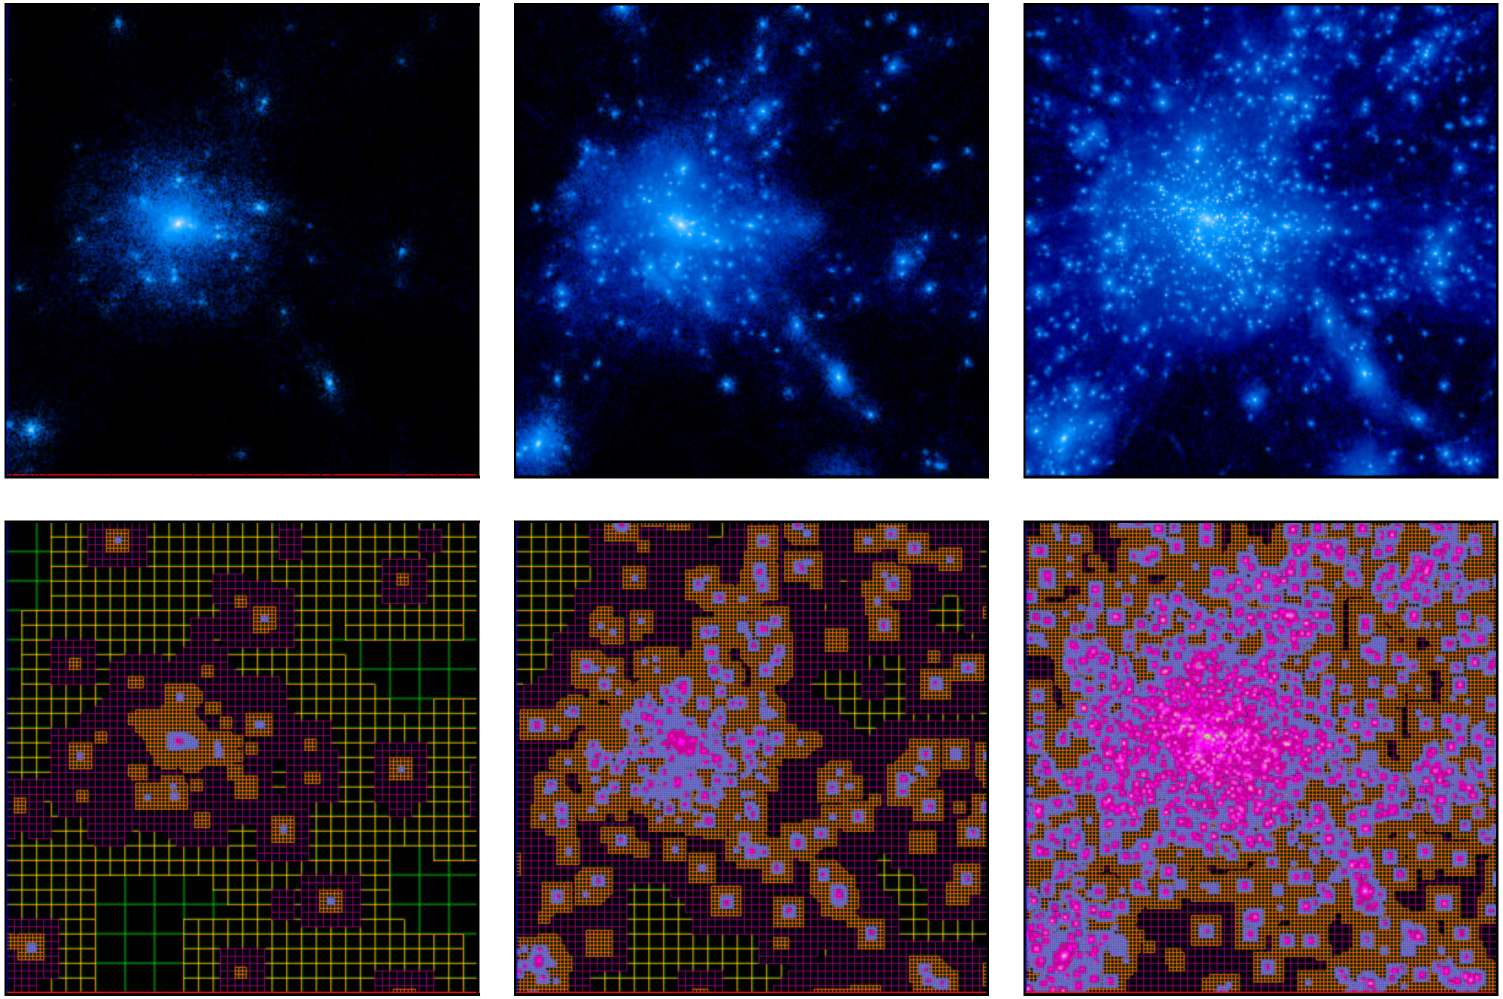

N $\simeq 5 \times 10^4$ N $\simeq 4 \times 10^5$ N $\simeq 3 \times 10^6$

Image size 12 h^{-1} Mpc

512^3 particles, hydro, cooling & star formation 10 h^{-1}Mpc ΛCDM at z=3

256 HP ES45 processors

350 h up to z=2.8