Mem. S.A.It. Suppl. Vol. 1, 109
© SAIt 2003

Memorie della
Supplementi

A software interface between Parallel Tree- and AMR Hydrocodes.

V. Antonuccio-Delogu[1,2] ⋆, U. Becciani[1],
D. Ferro[3] and A. Romeo[3]

[1] INAF-Osservatorio Astrofisico di Catania - Via S. Sofia 78, 95125 Catania, ITALY e-mail: van@sunct.ct.astro.it

[2] Theoretical Astrophysics Center, Juliane Maries Vej 30, DK 2100 Copenhagen, DENMARK

[3] Dipartimento di Fisica e Astronomia, Universitá di Catania - Via S. Sofia 78, 95125 Catania, ITALY

Abstract. We discuss the problems arising when one tries to create a software interface between a parallel treecode (modelling the collisionless, Dark Matter component) and Adaptive Mesh Refinement (AMR) hydrodynamical schemes (which model the gaseous phases). Such an interface would allow one to perform N-body/Hydro simulations in those situations where the N-body and hydrodynamical codes are loosely coupled, and the gravitational influence of the dissipative, gaseous component on the DM component can be neglected. We discuss similarities and differences between the tree- and AMR hydrocode data structures, restricting ourselves to two particular examples which are widely used in these days: the oct-tree Barnes-Hut and the cartesian structured grid AMR schemes. Finally, we present a simple scheme for such a software interface which we are beginning to use in cosmological simulations.

Key words. Computational astrophyics – Numerical methods – Cosmology

1. Introduction

Parallel N-body codes represent widely used tools of contemporary theoretical cosmology. Their usage to simulate the origin and evolution of the Large-Scale Structure of the Universe dates back to about 15 years ago (Frenk, et al. 1988; Hockney & Eastwood 1988). Since those early achievements a few algorithmic and technological advancements have enabled the astrophysical community to arrive to develop present-day parallel codes which are able to deal with hundreds of million of particles within simulation boxes ranging from of a few to about a few thousands Megaparsecs, using

Send offprint requests to: V. Antonuccio-Delogu

⋆ TAC external associate scientist.

Correspondence to: V. Antonuccio-Delogu - INAF-Osservatorio Astrofisico di Catania - Via S. Sofia 78, 95125 Catania, ITALY

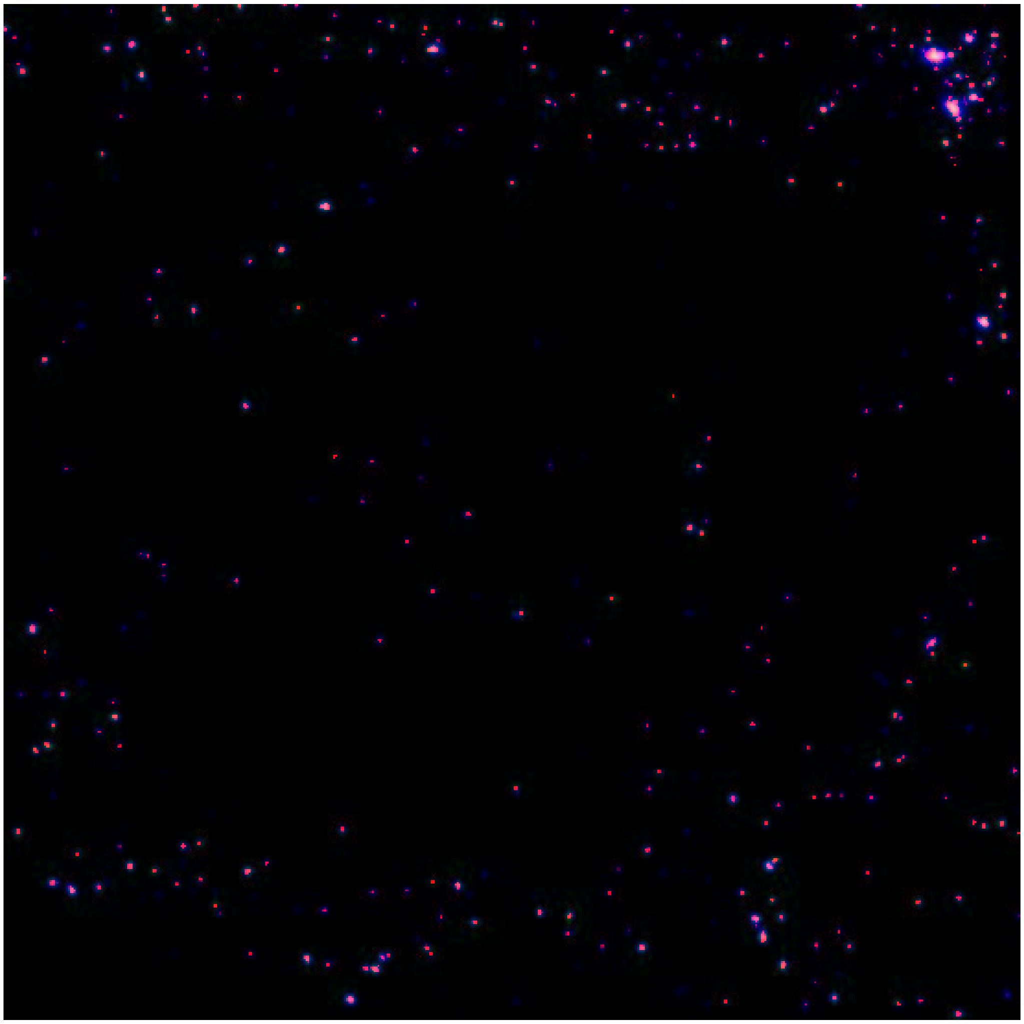

Fig. 1. Simulation of galaxy formation within underdense regions. Colour scale corresponds to density, ranging from $2 \times 10^3 \delta_c$ (*blue*) to $6 \times 10 \delta_c$, where $\delta_c(\Omega)$ is the critical overdensity for collapse. This snapshot shows a region of $12.5 h^{-1}$ Mpc.

a softening length small enough to allow the generation of galaxy-sized Dark Matter (DM) halos within cosmologically representative volumes (Evrard et al. 2002; Diaferio et al. 1999; Moore et al. 1999). These purely collisionless simulations are still important to address some crucial problems, like the evolution and merger histories of individual halos of subgalactic sizes within cluster and small groups (Governato et al. 2001). The problem of the overproduction of subgalactic halos, and of the apparent lack of these objects in our Local Group neighborhood is a good example of the fritful interaction between theoretical hints originating from cosmological simulations and observational results. Although N-body cosmological simulations do not include the baryonic (gaseous and stellar) component, they can still be very useful to

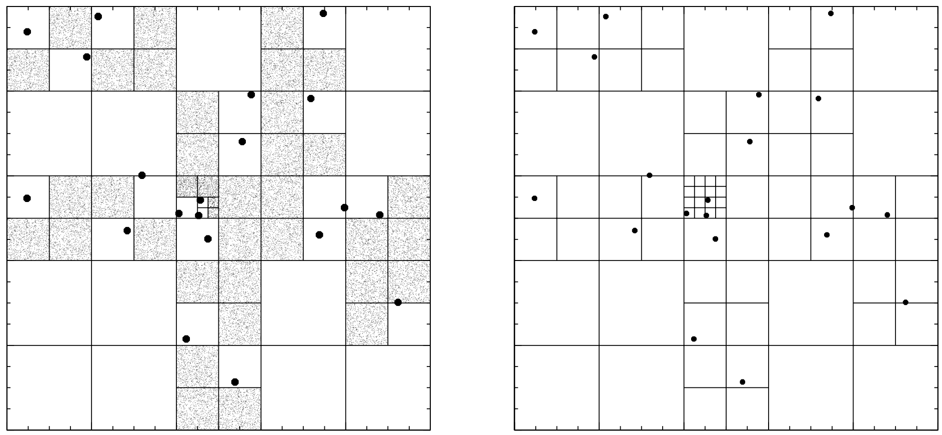

Fig. 2. *Left panel:* Quad-tree Barnes-Hut decomposition for a simple N-body configuration. Only cells containing at least one particle are included in the tree. We have shaded those cells at level n-1 and n which are not included into the tree, because they do not contain any particle. *Right panel:* AMR decomposition of the same region as done in the PARAMESH library. Note that now all cells at level n are included in the tree structure.

study some problems which would otherwise be very difficult to study observationally.

In figure 1 we show an example of a simulation aimed at studying DM halos within underdense regions (Antonuccio-Delogu et al. 2002). The underlying cosmological model is a ΛCDM with $\Omega_\Lambda = 0.7, \Omega_m = 0.3$. The simulation box was $50h^{-1}$ Mpc wide, and within this region a total number of 256^3 particles were initially arranged in a constrained configuration according to the Hoffman-Ribak algorithm (Bertschinger 2001). The constraints were chosen in such a way that large undedense region extending over two octants of the simulation volume will form. The picture shows only a smaller section of the simulation box occupied by this "Void". From the picture and a topological analysis (Antonuccio-Delogu et al., in preparation) one finds that DM halos are not randomly distributed within the Void, but tend to form within filamentary structures having a coherence length comparable to the size of the Void. The first unbiased observational attempts at characterizing the properties of galaxy populations in Voids seem to confirm these predictions (Grogin & Geller 1999, 2000). It is interesting to note that subgalactic halos are also abundantly produced in this simulation (which has a softening length of only 5 h^{-1} Kpc) and they seem to have a very different, more uniform distribution, in qualitative agreement with recent work on the spatial distribution of Ly-α clouds in Voids (Manning 2002).

These examples show how collisionless N-body simulations can still be very useful to study some crucial cosmological problems. However, radiation emitting matter is made of baryons which are described in terms of dissipative equations, and it is this the component which is accessible to observations. It is then very important to be able to perform N-body/Hydrodynamical simulations, where the gaseous component falls into the potential wells traced by the Dark Matter Large-scale Structure, and to be able to follow its radiative cooling and all the other complex phenomena which determine its final state. In general, this is a formidable problem because radiative and/or hydrodynamical phenomena exhibit a wide range of spatial and temporal scales, making a realistic numerical treatment of galaxy formation and evolution within a cluster almost impossible to be afforded even with the most ambitious computational systems. An example can be offered by supernovae explosion. The stellar environments where SN form have typical sizes of few parsecs, yet their winds can expand for more than 10 Kpc, thus heating the interstellar medium, and if the source lies at the boundary of an elliptical galaxy, they can also rapidly expand in the lesser dense Intergalactic Medium, over much larger distances. These events usually take place on time scales of the order of a few million years, but winds can continue to expand for 10 or 20 million years (depending on the density of the ambient medium), gradually releasing their kinetic energy.

2. Data Structures

The main data structure of the Barnes & Hut octal-tree algorithm is veny similar to the data structures adopted by many Adaptive Mesh Refinement (AMR) eulerian algorithms, particularly by those making use of *Structured Cartesian Grids* (hereafter SCG). The latter are decompositions of the spatial domain obtained by recursively decomposing the simulation box with planes parallel to the orthogonal axes of a cartesian domain. SCGs are easy to build and maintain, and have already been adopted in software libraries explicitly designed for AMR algorithms, for instance by the PARAMESH library (MacNeice et al. 2000). From now on we will restrict ourselves to SCGs, but most of the considerations we do still hold for more general types of AMR decompositions. The structure of the AMR refinement is constrained by the numerical accuracy of the solution. As is evident from a comparison between the left and right panels of figure 2, both a Barnes-Hut and a AMR SCG produced by PARAMESH are more structured in those regions where the density field becomes more structured. However, the SCG refinement in general includes more cells than the Barnes-Hut, because at the edges between neighbouring cells of different sizes, a uniform distribution of grid points engendered by the projection of the smaller and deeper cells is required by interpolation and extrapolation schemes which are used to refine and de-refine the solution between different cell levels.

3. Mapping a tree onto an AMR grid

3.1. General scheme

Our strategy in this paper is to explore the possibility of creating a general interface which could “link” together two already existing codes.

The collisionless N-body part models the DM and the stellar components, while the hydrodynamical code is designed to model the evolution of the diffuse baryonic components. The dominant coupling between these two components is through the gravitational field, and in most cases of interest to cosmological simulations the characteristic temporal scales over which the DM component evolves are larger by 2-3 orders of magnitudes than those of the gaseous component (cf. Table 1). This implies that also the typical timesteps of the gaseous component are smaller by the same order of magnitude, so that many more integration timesteps of the hydro component of the code are required, implying in turn that the two components can be considered to be computationally *loosely coupled*.

Moreover, in typical cosmological simulations the contribution to the total gravitational potential from the gaseous component is significant only on galactic and subgalactic scales at any epoch, so that if one is interested in modelling phenom-

Table 1. Characteristic time scales. Times in the last row are in Myrs.

	Cluster	Filament	Void
δ_{DM}	$\approx 10^3$	$5 \times 10^1 - 10^2$	-10^{-1}
δ_{gas}	$10^3 - 5 \times 10^3$	$10^2 - 10^3$	$-10^{-1}\Omega_b$
T_{gas}	$10^7 - 10^8$ K	$10^5 - 7 \times 10^6$ K	$10^4 - 10^5$ K
τ_{DM}	130	590-417	1.320

ena on larger scales one can simply include only the DM component as the source of the gravitational field. As a result of these two circumstances, the software interface between the N-body and the hydro components should perform two tasks:

– Map the gravitational potential of the DM component onto the AMR grid of the hydro code;
– Update the global gravitational potential when either of the two components changes siginificantly.

In view of the fact that the gravitational action of the gas is negligible, the DM component can be evolved independently of the gaseous one. The hydrodynamic code requires as an input the DM gravitational potential, but being the two components loosely coupled the update of the potential in the hydrodynamic code must not be performed at each timestep, but only when the configuration has significantly evolved.

Before discussing an actual implementation of this scheme, we would like to point out a few important restrictions of this scheme. The time scale τ_{DM} in practice is not always the best choice to get an estimate of the characteristic interval over which the DM gravitational potential changes significantly. After all, actual N-body simulations are performed with a finite number of particles, and altough this number can seem quite large, it is many orders of magnitude smaller than the actual number of DM particles! Discreteness (or *"graininess"*) means that few outliers can always exist in an actual simulation. These are for instance particles which acquire an anomalously high velocity, crosssing the entire simulation box in lesser than a dynamical time, thus modifying the gravitational potential of the DM configuration significantly along the path of their trajectories. Although the timestep of a N-body simulation is chosen in such a way to make the leapfrog integrator stable (Efstathiou, Davis, White, & Frenk 1985; Springel, Yoshida, & White 2001), discreteness effects can always induce few particles to violate this condition for a few timesteps (the same discreteness effects induce then an artificial relaxation of these outliers through dynamical friction). However, these outliers can represent a problem only if they would cluster and move coherently even for a few timesteps through the simulation box, but in reality they are always isolated and do not affect significantly the environment through which they move (Antonuccio-Delogu et al., in preparation).

A second restriction to this scheme comes from the AMR grid. In many codes the maximum level of refinement is a free parameter, and a large value of this parameter implies a higher spatial resolution in those zones where refinement is requires (e.g. at shocks). The maximum level of refinement of the AMR grid should however be chosen in such a way that finest cells of the AMR grid are always 4-5 times *larger* than the *terminal* cells of the N-body tree (i.e. those cells containing only one particle), because on these scales the accuracy of the tree in approximating the DM potential dimisnishes significantly. This can be a serious problem particularly in the underdense regions at some inetermediate epochs, when clustering on galactic scales has not yet had time to develop, and there are few particles within a large volume, thus making the terminal cells rather large.

3.2. An implemementation of the interface

We have implemented the scheme described in this section as a F90 module within the AMR code FLASH (Fryxell et al. 2000; Calder et al. 2002). The implementation is relatively straightforward. *At each timestep* the module performs the following operations:

1. Check whether the gravitational potential needs to be updated: if the answer is "NO" then go to step 3);
2. If step 1) is "YES" then read the input positions and velocities from FLY and compute the grav.pot.;
3. Evolve the system of one timestep;
4. Check whether simulation should be stopped: if the answer is "NO" go to step 1), else STOP

We will now look slightly in more detail at these steps.

1. At start a parameter `CheckPotential` is checked. If its value is `0` an output file from FLY with the same redshift as the current redshift is searched. If this file does not exist, the two output files whose redshifts define the smallest interval within which the current redshfit is contained are taken to be the extremes over which an interpolation is performed. The interpolated particle positions are read by FLASH.
   This checkpointing operation is performed in addition to other checkpointing perations performed by FLASH at startup.
2. The particles positions are mapped onto the FLASH AMR grid structure using a procedure similar to that of the module `MapParticlesTo Mesh`, which is already included in FLASH (we are currently referring to version 2.2). Then the gravtitational potential on the mesh is computed using a method similar to that used by FLY, i.e. adpting a multipole expansion of the potential.
3. This is done simply by including the module `gravity` of FLASH.
4. The simulation is stopped either when one of the conditions for this to happen already present in FLASH are met, or if one of the two conditions quoted at the end of the previous subsection are met.

We have implicitly assumed that the hydrodynamical simulations is performed *after* the N-body one, using the N-body positions at a few selected outputs as an input to compute the gravitational potential needed by FLASH to evolve the configuration. It would of course be possible to perform concurrent runs of the two codes, either on the same computing system or on two different partitions of the same system, which could also be a GRID of large-scale facilities. The condition of loose coupling would then translates into the exchange of a signifcant amount of data (of the order of a few GBytes for each output of a state-of-the-art N-body simulation) over some widely spaced time intervals. This is a situation ideal for current possibilities of planned GRID facilities, making this kind of software interface an ideal one to perform large simulations of cosmological models.

4. Conclusions

We have shortly described the general characteristics of a software interface between N-body tree- and AMR hydrocodes. There are many similarities between these two types of data structures, making a map between the two a possible task. *Loose coupling* of the two parts (the collisionless and the hydrodynamical) is the key which makes actual implementation of this scheme a practical tool to perform simulation over large-scale computational grids. Using this mapping to perform cosmological N-body/Hydrodinamical simulations poses further problems, however, which restrict the physical parameter space over which simulations can be performed. The latter is however large enough to encompass, for instance, simulations of the evolution of the Intergalactic Medium over regions encompassing a few clusters of galaxies. Other problems which require high resolution over in some small regions embedded within the Large-Scale Structure of the Universe arise for instance in the study of the propapagation of Ultra-High-Energy Cosmic Rays, where a detailed knowledge of the intergalactic magnetic field over regions of the size of our Local Group of galaxies is important to make quantittative estimates of the diffusion of charged cosmic rays, and ultimately to determine more precisely the extent of the maximum region from which these particles could originate (Isola & Sigl 2002). The inclusion of feedback effects from galaxies is possible, but in order to fully exploit the possibilities of spatial and temporal resolution offered by the AMR class one should be able to further push the resolution of the gravitational potential provided by the N-body part in those regions where the hydrocode requires it. Numerical simulations which include these effects have been attempted recently (Tassis et al. 2002). We have recemtly begun a program to realize a more asbtract and general software interface than that described here, using the encapsulation possibilities offered by C++. A more detailed analysis will be soon presented elsewhere (Antonuccio-Delogu & Becciani, in preparation).

Acknowledgements. Part of this work was supported by the Theoretical Astrophysics Center and by the Italian Space Agency (ASI).

References

Antonuccio-Delogu, V., Becciani, U., van Kampen, E., Pagliaro, A., Romeo, A. B., Colafrancesco, S., Germaná, A., & Gambera, M. 2002, MNRAS, 332, 7

Bertschinger, E. 2001, ApJS, 137, 1

Calder, A. C. et al. 2002, ApJS, 143, 201

Diaferio, A., Kauffmann, G., Colberg, J. M., & White, S. D. M. 1999, MNRAS, 307, 537

Efstathiou, G., Davis, M., White, S. D. M., & Frenk, C. S. 1985, ApJS, 57, 241

Evrard, A. E. et al. 2002, ApJ, 573, 7

Frenk, C. S., White, S. D. M., Davis, M., & Efstathiou, G. 1988, ApJ, 327, 507

Fryxell, B. et al. 2000, ApJS, 131, 273

Governato, F., Ghigna, S., Moore, B., Quinn, T., Stadel, J., & Lake, G. 2001, ApJ, 547, 555

Grogin, N. A. & Geller, M. J. 1999, AJ, 118, 2561

Grogin, N. A. & Geller, M. J. 2000, AJ, 119, 32

Hockney, R. W. & Eastwood, J. W. 1988, Bristol: Hilger, 1988,

Isola, C. & Sigl, G. 2002, Phys. Rev. D, 66, 83002

Manning, C. V. 2002, ApJ, 574, 599

Moore, B., Ghigna, S., Governato, F., Lake, G., Quinn, T., Stadel, J., & Tozzi, P. 1999, ApJ, 524, L19

MacNeice, P, Olson,K. M., Mobarry, C., deFainchtein, R. & Packer, C. 2000, Computer Physics Communications, 126, 330-354

Springel, V., Yoshida, N., & White, S. D. M. 2001, New Astronomy, Volume 6, Issue 2, p. 79-117., 6, 79

Tassis, K., Abel, T., Bryan, N. & Norman, M.L. 2002, astro-ph/0212457