MASSIVELY PARALLEL SIMULATIONS OF ASTROPHYSICAL PLASMAS: STATUS AND PERSPECTIVES OF THE COAST PROJECT

B. Thooris, E. Audit, A. S. Brun, Y. Fidaali, F. Masset, D. Pomarède, R. Teyssier
Institut de Recherche sur les Lois Fondamentales de l'Univers
DSM/IRFU CEA/Saclay 91191 Gif-sur-Yvette France
email: Bruno.Thooris@cea.fr
http://irfu.cea.fr/Projets/COAST

KEYWORDS
Large scale computing, parallel computing, astrophysics, plasmas simulation, visualization.

ABSTRACT

The COAST (for Computational Astrophysics) project is a program of massively parallel numerical simulations in astrophysics involving astrophysicists and software engineers from CEA/IRFU Saclay. The scientific objective is the understanding of the formation of structures in the Universe, including the study of large-scale cosmological structures and galaxy formation, turbulence in interstellar medium, stellar magnetohydrodynamics and protoplanetary systems. The simulations of astrophysical plasmas are performed on massively parallel mainframes (MareNostrum Barcelona, CCRT CEA France), using 3-D magnetohydrodynamics and N-body parallelized codes developed locally. We present in this paper an overview of the software codes and tools developed and some results of such simulations. We also describe the Saclay SD*vision* graphical interface, implemented in the framework of IDL Object graphics, our 3-D visualization tool for analysis of the computation results.

1. Introduction

The COAST project [1,2] is dedicated to high performance computing in astrophysics. The goal is the understanding of the formation of structures in the Universe, by developing advanced techniques in parallel computing and in applied mathematics to model galaxy formation and predict their observational signatures, as a function of physical parameters. Astrophysicists and software engineers collaborate to rationalize and optimize the development of simulation programs by creating a core of common specific modules and using common software tools for data handling, post-treatment, visualization, numerical methods, parallelization and optimization.

2. Overview of the simulation programs

Four major numerical simulation programs are used to cover different physics scales:

- The RAMSES code

RAMSES [3,4,5,6] is a hybrid, N-body and hydrodynamical 3-D code which solves the interplay of the dark matter component and the baryon gas for studying the structure and the distribution of galaxy clusters starting for the initial conditions of the Big Bang. The code is based on the Adaptive Mesh Refinement (AMR) technique, written in FORTRAN90 and parallelized with the MPI library [7]. Current developments focus on solving the full MHD set of equations.

- The HERACLES code

HERACLES [8,9,10,11,12] is a 3-D code which solves the equations of radiative transfer coupled to hydrodynamics. It studies thermal condensation in molecular clouds in the Interstellar Medium, radiative shocks, molecular jets of young stars and proto-planetary disks. It is written in FORTRAN90, parallelized with MPI and implemented in cartesian, cylindrical and spherical coordinates with regular mesh grids.

- The ASH [13,14] code

ASH (for Anelastic Spherical Harmonic) performs 3-D magnetohydrodynamics simulations in spherical geometry for the study of the turbulence and magnetic dynamo process in solar and stellar interiors. ASH, unlike the others codes presented which are completely developed in CEA/Saclay, is jointly developed at Saclay and at the University of Boulder.

- The JUPITER [15,16,17] code

JUPITER is a mutidimensional astrophysical hydrocode. It is based on a Godunov method, written in C and parallelized with MPI. The mesh geometry can either be

Proceedings of the 2008 High Performance
Computing & Simulation Conference ©ECMS
Waleed W. Smari (Ed.)
ISBN: 978-0-9553018-7-2 / ISBN: 978-0-9553018-6-5 (CD)

cartesian, cylindrical or spherical. It allows mesh refinement and includes special features adapted to the description of planets embedded in disks.

3. Computing facilities

The COAST team uses local resources for development and post-treatment: the 256 cores 2.6 GHz opteron cluster DAPHPC, with an infiniband interface and four visualization stations with 16 to 32 Gb RAM, ~1Tb disk, 4 processors, 1Gb memory graphics cards, 30 inches screens.
Massive simulations are performed at CCRT (CEA National Supercomputing Center) on Platine, ranking 26th in the TOP500 world supercomputer list (November 2007): 7456 Itanium cores, total 23 Tb memory, 47.7 Teraflops (4 Mhrs computation in 2007).
Other resources for massive simulations (2 Mhrs for 2007) can be accessed on MareNostrum at the Barcelona Supercomputing Center, ranking 13th in the TOP500 world supercomputer list (November 2007): 10240 IBM PowerPC 2.3 GHz cores with 94.2 Teraflops, 20Tb of main memory.
Furthermore, the project will have access in 2008 to the IDRIS (French CNRS Supercomputing Center) new Blue Gene/P system with 40000 cores.

4. The recent COAST computational milestones

COAST members are participating in French or European founded collaborations (Horizon, Magnet, Sinerghy or STARS[2]).
In the context of the Horizon collaboration [18], COAST members succeeded in the HORIZON Grand Challenge Simulation at CEA/CCRT on Platine in September 2007, which is the largest ever N-body cosmological simulation performed. For the first time, have been performed a simulation of half the observable universe, with enough resolution to describe a Milky Way-like galaxy with more than 100 dark matter particles. The RAMSES code has been run on 6144 cores, 18 Tb RAM used for 2 months to simulate 4096^3 □70 billions particles. This is an improvement of about an order of magnitude with respect to previous experiments, as illustrated in Fig.1. This simulation has been chosen to simulate future weak-lensing surveys like DUNE or LSST [19,20].
Another challenge in computing in astrophysics in 2007 was the HORIZON "galaxy formation" simulation at MareNostrum. The characteristics of the run are the following: 1024^3 dark matter particles, 4 billions AMR cells, box size 50 Mpc/h, resolution in space 2 kpc. 2048 processors were needed for computing, 64 processors dedicated to I/O, 3 weeks of computations so far, down to z=1.9, 20 Tb of data generated and stored. The run performed simulations from large scale filaments to galactic discs.

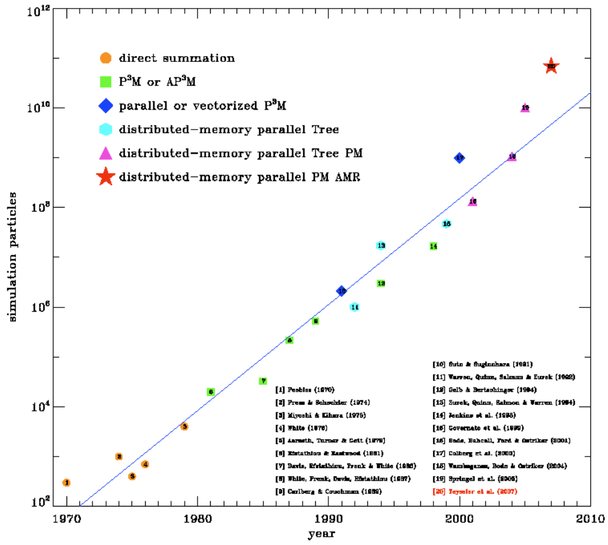

Fig.1 Size of the N-body simulations versus time (courtesy V. Springel [21]).

Two examples of MareNostrum data visualization are shown in Fig.2 and Fig.3, displaying the density distribution of the baryon gas at two different scales.

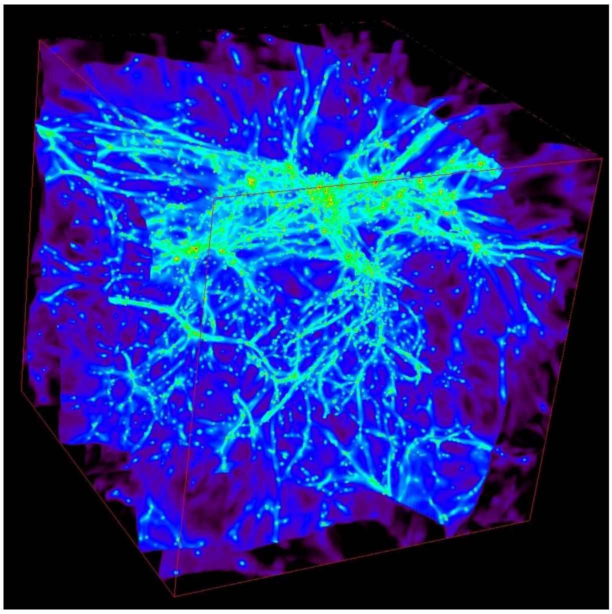

Fig.2 Density distribution of the baryon gas in the central region of simulation domain. Visualization made with SD*vision* software.

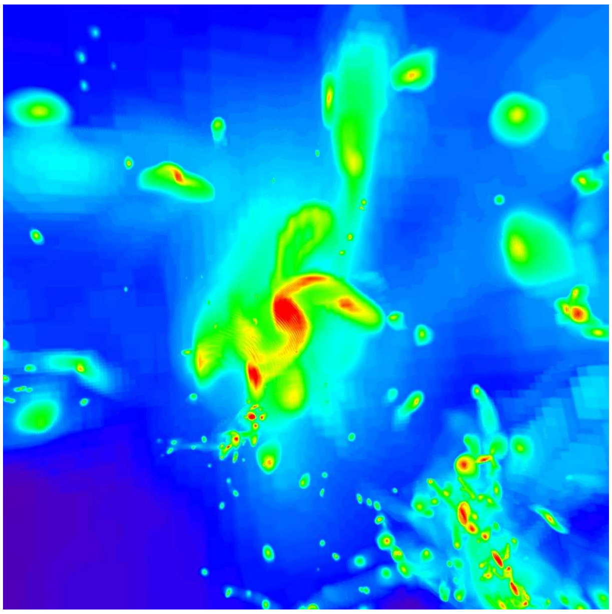

Fig.3 Zoom of galactic spirals, density distribution of the baryon gas. Visualization made with SD*vision* software

5. Data Handling

A unique format, HDF5 [22] (Hierarchical Data Format), has been chosen for the data produced by all our simulations codes. HDF5 is developed and maintained by NCSA (National Center for Supercomputing Applications). This library emphasizes storage and I/O efficiency in particular for data intensive computing environments. For instance, the HDF5 format can accommodate data in a variety of ways, such as compressed or chunked and the library is tuned and adapted to read and write data efficiently on parallel computing systems. NCSA maintains a suite of free, open source software, including the HDF5 I/O library and several utilities. The visualization tools developed is using IDL, which integrate a module able to read easily the HDF5 files.

6. Visualization

The visualization plays a very important role in the development of simulations codes. Fundamental aspects including domain decomposition, initial conditions, message passing and parallelization, treatment of boundary limits, can be controlled and evaluated qualitatively through visualization. Once in production phase, visualization is also used for the validation, the analysis and the interpretation of the results. A complete graphical interface named SD*vision* [23,24,25] has been developed in order to participate in the development of the simulation codes and visualize the large astrophysical simulation datasets produced in the context of the COAST program.

The SD*vision* graphical interface

The interface is implemented as a graphical widget providing interactive and immersive 3-dimensional navigation capabilities. The baseline technology is the object-oriented programming offered by IDL's Objects Graphics [26]. It benefits from hardware acceleration through its interface to the OpenGL libraries. An example of the widget displayed in its running state is shown in Fig. 4.

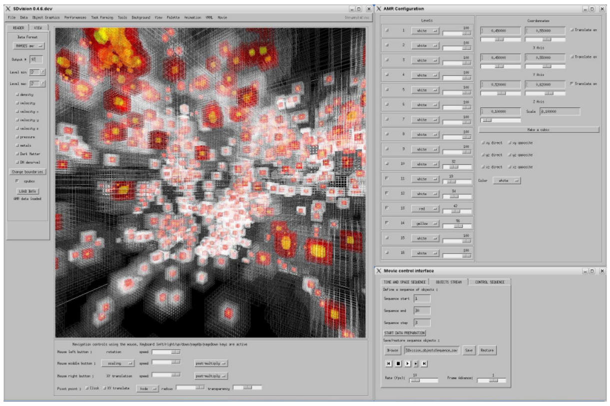

Fig.4 The SD*vision* widget used to visualize the levels 11,12,13 AMR cells of a galaxy cluster simulation mesh.

SD*vision* has been developed to visualize the huge amount of data produced by the codes RAMSES, HERACLES, JUPITER and ASH. It allows the display of complex scenes with scalar fields (volume projection, 3D isosurface, slices), vector fields (streamlines) and particle clouds.

Two examples of such visual representations are shown in Fig.5 and Fig.6.

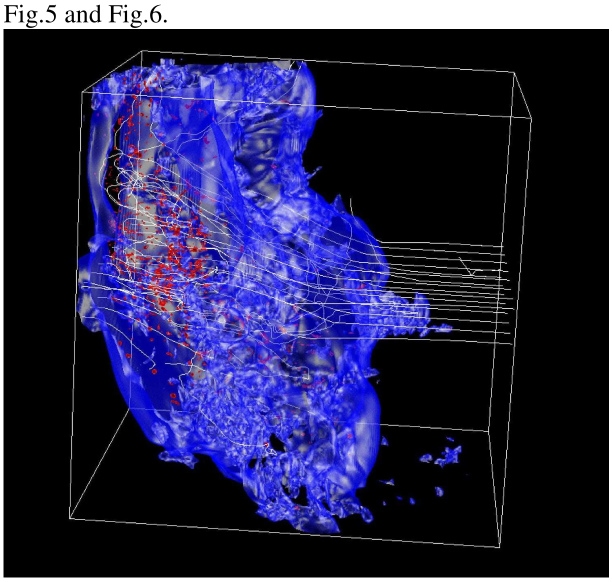

Fig.5 Visualization of the iso-density surfaces obtained in a high-resolution 1200^3 HERACLES 256-processors simulation of turbulences in the interstellar medium.

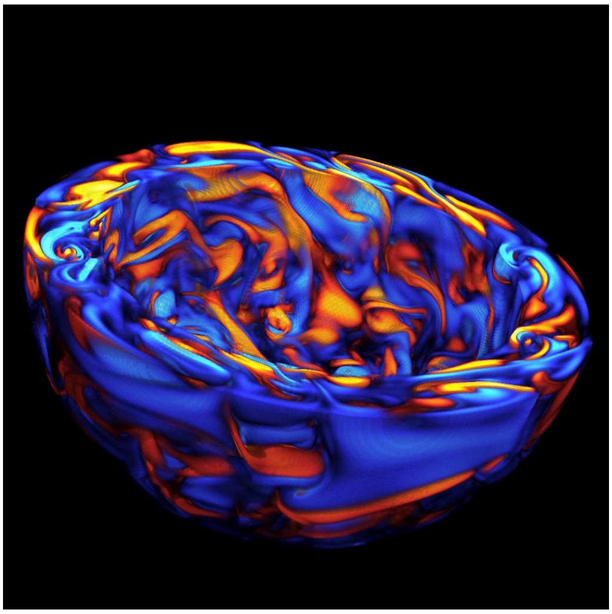

Fig.6 SD*vision* display of a 500x500x500 ASH simulation. This view of the azimuthal component of the magnetic field in the solar convection zone (positive in red, negative in blue) are obtained by the volume rendering technique with clipping.

7. Perspectives and conclusion

To be able to follow the improvements of the computers techniques, the COAST team dedicates a part of his human resources in training on software tools such as optimization techniques on recent mainframes (IDRIS Blue Gene) and R&D on new promising technologies for computing (GPU, Cells,…). Future challenges include the improvement by an order of magnitude in the N-body simulation of cosmological structure formation to improve the resolution. Using the next generation of high-performance centers, the objective is to simulate 8192^3 □550 billions particles

Computational astrophysics has a bright future, lying on the ever increasing performances of massively parallel mainframes. To achieve his ambitions, the project relies on a synergy between astrophysicists and software developers, local computing resources and access to supercomputers. Other major projects and initiatives are currently following the same approach: e.g. the FLASH Center at the University of Chicago [27], the ASTROSIM European Network for Computational Astrophysics [28], or the VIRGO consortium for Cosmological Supercomputer Simulations [29].

8. References

[1] http://irfu.cea.fr/Projets/COAST

[2] D. Pomarède, B. Thooris, E. Audit, R. Teyssier. Numerical Simulations of Astrophysical Plasmas. *Proceedings of the 6th IASTED International Conference on Modelling, Simulations, and Optimization (MSO2006)*, Gaborone, Botswana, September 11-13, 2006, ed. H. Nyongesa, 507-058, Acta Press, ISBN:0-88986-618-X

[3] R.Teyssier, Cosmological hydrodynamics with adaptive mesh refinement - A new high resolution code called RAMSES, Astronomy and Astrophysics, 385 (2002) 337-364

[4] S. Fromang, P. Hennebelle, R. Teyssier. A high order Godunov scheme with constrained transport and adaptive mesh refinement for astrophysical magnetohydrodynamics, A&A 457 (2006) 371

[5] R. Teyssier, S. Fromang, and E. Dormy. Kinematic dynamos using constrained transport with high order Godunov schemes and adaptive mesh refinement , J. Comp. Phys. 218 (2006) 44

[6] R. Osmont, D. Pomarède, V. Gautard, R. Teyssier, B. Thooris. Monitoring and Control of the RAMSES Simulation Program, *Proceedings of the CCP2007 Conference on Computational Physics*, Brussels, Belgium, September 5-8, 2007

[7] The Message Passing Interface (MPI) Standard, http://www-unix.mcs.anl.gov/mpi/

[8] M.Gonzalez, E.Audit, P.Huynh. HERACLES : a Three Dimensional Radiation Hydrodynamics Code. *Astronomy and Astrophysics,* 464 2 (2007) 429-435.

[9] E. Audit and P. Hennebelle, Thermal Condensation in a Turbulent Atomic Hydrogen Flow, *Astronomy and Astrophysics*, 433, 2005, 1-13.

[10] M. Gonzalez and P. Velarde. Radiative Shocks and Jets Simulated with the ARWEN and HERACLES codes. *Proceedings of the IGPP/DAPNIA International Conference on Numerical Modeling of Space Plasma Flows, ASTRONUM2007, Paris, France, June 11-15, 2007, to appear in the Astronomical Society of the Pacific Conference Series.*

[11] E. Audit. Fragmentation in the Interstellar Medium. *Proceedings of the IGPP/DAPNIA International Conference on Numerical Modeling of Space Plasma Flows, ASTRONUM2007, Paris, France, June 11-15, 2007, to appear in the Astronomical Society of the Pacific Conference Series.*

[12] E. Audit, V. Gautard, D. Pomarède, B. Thooris. Enabling Tools and Techniques for the Optimization of the HERACLES Simulation Program, *Proceedings of the 6th EUROSIM Congress*, Ljubljana, Slovenia, September 9-13, 2007

[13] J. Ballot, A.S. Brun, and S. Turck-Chieze. Simulations of turbulent convection in rotating young solarlike stars: differential rotation and meridional circulation. ApJ 669 (2007) 1190

[14] L. Jouve and A.S. Brun. On the role of meridional flows in flux transport dynamo models. A&A 474 (2007) 239

[15] F. Masset, A. Morbidelli, and A. Crida. Disk surface density transitions as protoplanet traps. ApJ 642 (2006) 478

[16] A. Crida, A. Morbidelli, and F. Masset. Simulating planet migration in globally evolving disks. A&A 461 (2007) 1173

[17] http://www.maths.qmul.ac.uk/~masset/index.html

[18] http://www.projet-horizon.fr/

[19] http://www.dune-mission.net/

[20] http://www.lsst.org/lsst_home.shtml

[21] V. Springel, et al. Simulating the joint evolution of quasars, galaxies and their large-scale distribution. Astro-ph/0504097 (2005).

[22] http://www.hdfgroup.org

[23] D. Pomarède, E. Audit, R. Teyssier, B. Thooris. Visualization of large astrophysical simulations datasets. *Proceedings of the Conference on Computational Physics 2006, CCP2006, Gyeongju, Republic of Korea, aug.29-sept.1 2006, ed. J.S. Kim, Computer Physics Communication*, 177 (2007) 263, doi:10.1016/j.cpc.2007.02.065

[24] D. Pomarède, Y. Fidaali, E. Audit, A.S. Brun, F. Masset, R. Teyssier. Interactive visualization of astrophysical plasma simulations with SDvision, *Proceedings of the IGPP/DAPNIA International Conference on Numerical Modeling of Space Plasma Flows, ASTRONUM2007, Paris, France, June 11-15, 2007, to appear in the Astronomical Society of the Pacific Conference Series.*

[25] D. Pomarède, E. Audit, A.S. Brun, V. Gautard, F. Masset, R. Teyssier, B. Thooris. Visualization of astrophysical simulations using IDL Object Graphics, *Proceedings of the Computer Graphics Imaging and Visualization 2007 Conference*, Bangkok, Thailand, August 14-17, 2007, IEEE Computer Society ISBN 0-7695-2928-3 p471-480

[26] http://www.ittvis.com/idl/

[27] http://flash.uchicago.edu/

[28] http://www.astrosim.net/

[29] http://www.virgo.dur.ac.uk/