

SIMULATION ON AMR GRIDS AT THE EXASCALE ERA: THE DYABLO CODE

M. Delorme¹, A. Durocher¹, D. Aubert², A. S. Brun³ and O. Marchal²

Abstract. Dyablo is a code developed at CEA since 2019 for the simulation

Keywords: Adaptive-Mesh-Refinement, Exascale, GPU, CFD

1 Introduction

With the upcoming launch of the German and French exascale supercomputers, European scientists have been developing and migrating HPC codes to fully support graphics processing units (GPUs). Due to the variety of constructors (Nvidia, AMD, Intel) and programming models (CUDA, HIP, SYCL), new codes tend to rely on performance portability libraries to outsource the specific tasks of memory allocation and parallel dispatch. Such libraries allow developers and physicists to implement computing kernels in an agnostic way that will be compiled and optimized on different backends for CPUs or GPUs. At the forefront of these libraries, Kokkos (Trott et al. 2022) (Carter Edwards et al. 2014) has now become a very relevant choice for the future exascale codes. In 2024, Kokkos became part of the newly founded High Performance Software Foundation*, a linux foundation dedicated to maintaining an open-source high-performance ecosystem on contemporary and future supercomputing architectures.

For a few years now, European teams have been benefitting from portability performance libraries to develop new codes in astrophysics. In France, the Idefix code (Lesur et al. 2023) has been the first to use Kokkos for astrophysics and to have reached a production stage for the simulation of MHD astrophysical fluids on fixed grids. Some simulations, however, require large ranges in temporal and spatial scales which cannot be reached with fixed grids. The use of adaptive mesh refinement (AMR) is thus a very promising solution. However, parallelism of AMR on GPU is a notoriously difficult problem due to the non-regular memory access and potentially heavy data structures.

In this communication, we present Dyablo, a code developed at CEA for the simulation of astrophysical fluids on AMR grids using Kokkos. The next section presents the key-concepts behind Dyablo, while section 3 presents two sample simulations and scalability plots.

2 The Dyablo Code

Dyablo is a finite-volume code aimed at leveraging performance from modern hardware and modern concepts in software engineering. The code is written in C++(17) and relies on Kokkos for intra-node performance portability and MPI for inter-node parallelism.

At its core, Dyablo relies on three pillars: separation of concerns, modularity and abstraction. *Separation of concerns* dictates that the code should be compartmentalized in such a way that physicists and mathematicians can implement modules without being able to impact the high-performance sections of the code. On the other hand, computer scientists should be able to change the various backends used by the code (e.g. the way the AMR is managed) without having an impact on the physics/mathematics side of the code. The second pillar,

¹ IRFU, CEA, Université Paris-Saclay, F-91191 Gif-Sur-Yvette, France

² Observatoire Astronomique de Strasbourg, Université de Strasbourg, CNRS UMR 7550, 11 rue de l'Université, 67000 Strasbourg, France

³ Université Paris-Diderot, AIM, Sorbonne Paris Cité, CEA, CNRS, F-91191 Gif-Sur-Yvette, France

*<https://hpsf.io>

modularity, implies that the code should be versatile enough so that we can activate or deactivate certain treatments without having to recompile the code and that the implementation of new treatments, solvers, or backends should be made easy by the code. Finally, *abstraction* is a way to provide the users high-level methods that will hide the complexity of the underlying data models, to allow mathematicians and physicists to have a code that is closer to their field of expertise.

These three pillars leads to the key concept of plugins in Dyablo. A plugin is an interchangeable piece of code that can be parametrized at runtime. Each aspect of Dyablo is a plugin: hyperbolic and parabolic solvers, time-step calculation, initialization, refinement or IOs are examples of different types of plugins in Dyablo. This allows us to implement a variety of solutions for each plugin, and let the user choose the combination at runtime without needing to recompile the code.

In order to streamline as much as possible calculations on GPU and avoid superfluous transfers and memory usage, we use a linear octree (see e.g. (Burstedde 2020)) stored on device for the AMR tree. Dyablo avoids having too many non-conformal interfaces by implementing block-based AMR where each entry of the linear octree stores a regular cartesian subgrid. Emphasis has been made on providing high-level abstractions to access the AMR grid. Hence, tree-traversal, neighbor finding and memory accessing are done via high-level methods.

Finally, all developments are made on a internal CEA gitlab with a custom continuous integration system that tests and validates each merge request on the code. A public mirror of the current development branch is pushed on github[†] and is widely accessible.

3 Simulations and Scaling

Simulations with Dyablo are now running on all French major super computers, including the Irene, Jean-Zay and Adastra computers, supporting both their CPU and GPU partitions. We now present two sample simulations as well as some scalability results.

3.1 Radiative cosmology




Fig. 1. Radiative cosmological simulation at half reionization. Refinement is triggered on the mass of the cells. **Left:** The ionization fraction is represented by the orange bubbles. **Right:** The distribution of dark matter particles in the box.

A first simulation has been done in the context of the GINEA (Aubert & Durocher 2021) research group and can be seen on Fig 1. This run simulates the evolution of matter in a Λ CDM universe with radiation. Dyablo solves the hydrodynamics equations, the Poisson equation for gravity using a preconditioned conjugate gradient solver, as well as radiation using an explicit M1 solver (Aubert & Teyssier 2008). The simulation has been

[†]<https://github.com/Dyablo-HPC/Dyablo>

computed on the AMD Rome CPU partition of the Irene supercomputer on 16 MPI ranks using the OpenMP backend of Kokkos. The simulation is started with a base resolution of 512^3 grid cells and 512^3 particles for a $128 Mpc/h$ comoving boxsize. The run allows for 3 active refinement levels, pushing the maximum resolution to 2048^3 . The refinement is triggered on the total mass inside a cell. The box is periodic along each direction.

3.2 Solar Convection




Fig. 2. Solar convection simulation. **Left:** Vertical velocity. We see in blue the hot matter rising through the box and in red the rapid downflows of cold matter. **Right:** Refinement level of the box. The low-resolution regions are in blue while high-resolution are in purple. The refinement is triggered on the strong downflows, allowing the precise resolution of the edges of the convection cells.

The second simulation we present is one of solar convection and can be seen on Fig 2. We initialize a cartesian slab in hydrostatic equilibrium using a polytropic profile akin to the one defined in (Cattaneo et al. 1991). The precise setup can be found in (Delorme et al. 2022). In this simulation, Dyablo solves the hydrodynamics equations with static gravity and explicit thermal conduction and viscosity. The simulation ran on four Nvidia a100 on the Storm cluster at CEA. The initial resolution is set to $64 \times 64 \times 16$ and allows for 4 active refinement levels, meaning the maximum resolution would reach $1024 \times 1024 \times 256$. Boundary conditions are periodic horizontally and impenetrable, stress-free walls in hydrostatic equilibrium for the vertical boundaries.

3.3 Weak-Scaling

To evaluate scaling of Dyablo on super-computers and in particular on GPUs we use the solar-convection setup described in the previous section. Knowing the box is periodic along the horizontal direction we build a weak-scaling test where we take a single realization of a convection run, and we tile it horizontally having one domain per MPI process. This allows us to grow the size of the problem perfectly with the number of processes which is perfectly suited for load-balancing. To evaluate the performance of the computation kernels we deactivate IOs. Since all domains are tiled, the box is perfectly load-balanced at all times, and thus we also deactivate the well-balancing process. We proceed by having Dyablo compute a hundred iterations, with an AMR cycle every iteration.

We ran that scalability test on CPUs and GPUs on the Jean-Zay and Adastra French supercomputers. The results can be seen on Fig 3 for the GPUs. Scalability is indeed good, and in the case of Adastra we reach 2048 GPUs, having a total domain with approximately 62 billion cells. The computational kernels in blue, orange and green scale perfectly while the AMR cycle tends to increase time consumption slightly due to an increase of usage in communication bandwidth. Optimization is currently ongoing to improve the timing of the "other" bars that also tend to puff-up when increasing the number of processes. Finally we can see a discrepancy between the times-to-solution on a100 and mi250 that should be roughly the same. We believe this discrepancy comes from a difficulty for the ROCm compiler to optimize the HIP code generated by Kokkos due to the high level of abstractions in Dyablo.

4 Conclusions

Dyablo is nearing its first production runs. We have shown that the methodology employed for the development of the code allowed physicists from very different communities to implement their schemes and models easily.




Fig. 3. Weak scalability of the convection run on GPUs for French Supercomputers. **Left:** Scalability on Jean-Zay. Each node has 8 Nvidia a100 GPUs. **Right:** Scalability on Adastra. Each node has 4 AMD MI250 GPUs. Each process is bound to one MI250 GCD, so one GPU holds 2 MPI-processes and one node 8.

Thanks to the use of Kokkos, Dyablo is natively compatible with all French Tier-0 and Tier-1 supercomputers. The code is now able to run and scale on a large number of GPUs on the largest French supercomputer. In the coming years, the work will be focused on optimisation, IO formats and on improving current runs while adding more physics.

We acknowledge the financial support of the ERC Whole Sun Synergy grant #810218 and CNES Solar Orbiter supports. We are thankful to GENCI via project 1623. We also acknowledge the support of INSU and its programs PNST, PNCG, PNHE, PNPS and PCMI.

References

- Aubert, D. & Durocher, A. 2021, in SF2A-2021: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics, ed. A. Siebert, K. Baillié, E. Lagadec, N. Lagarde, J. Malzac, J. B. Marquette, M. N'Diaye, J. Richard, & O. Venot, 473–475
- Aubert, D. & Teyssier, R. 2008, *Monthly Notices of the Royal Astronomical Society*, 387, 295
- Burstedde, C. 2020, *ACM Transactions on Mathematical Software*, 46, 1
- Carter Edwards, H., Trott, C. R., & Daniel, S. 2014, *Journal of Parallel and Distributed Computing*, 74, 3202
- Cattaneo, F., Brummell, N. H., Toomre, J., Malagoli, A., & Hurlburt, N. E. 1991, *ApJ*, 370, 282
- Delorme, M., Durocher, A., Brun, A. S., & Strugarek, A. 2022, in SF2A-2022: Proceedings of the Annual meeting of the French Society of Astronomy and Astrophysics, ed. J. Richard, A. Siebert, E. Lagadec, N. Lagarde, O. Venot, J. Malzac, J. B. Marquette, M. N'Diaye, & B. Briot, 209–213
- Lesur, G. R. J., Baghdadi, S., Wafflard-Fernandez, G., et al. 2023, *A&A*, 677, A9
- Trott, C. R., Lebrun-Grandie, D., Arndt, D., et al. 2022, *IEEE Transactions on Parallel and Distributed Systems*, 33, 805