

Implementation of a flux limiter into a fully-portable, algebra-based framework for heterogeneous computing

X. Álvarez*, N. Valle*, A. Gorobets**, F. X. Trias*
Corresponding author: xavier@cttc.upc.edu

* Heat and Mass Transfer Technological Center, Technical University of Catalonia,
C/ Colom 11, Terrassa (Barcelona), 08222, Spain

** Keldysh Institute of Applied Mathematics,
Miusskaya Sq. 4, Moscow, 125047, Russia

Abstract: During the last years, there has been a significant increment in the variety of hardware to overcome the power constraint in the context of the exascale challenge. This progress is leading to an increasing hybridisation of high-performance computing (HPC) systems, making the design of computing applications a rather complex problem, and is affecting most of the fields that rely on large-scale simulations. Many scientific computing applications have been partially ported, even rewritten entirely, to take advantage of the coprocessor devices (*e.g.* GPUs or MICs). This porting process usually becomes very costly for the traditional stencil-based implementations as they involve a large number of functions (*i.e.* every numerical method requires specific functions and data-sets). In contrast, algebra-based implementations have shown to provide with a high level of abstraction and portability because they simplify the calculations into universal algebraic kernels only. In this work, we present an algebraic formulation of a high-resolution scheme, a flux limiter, and detail its implementation into the fully-portable, algebra-based HPC² framework. As a result, the solution of the advection of a scalar field with sharp discontinuities relies on a reduced set of algebraic kernels. Therefore, the algebra-based approach combined with a multilevel MPI + OpenMP + OpenCL parallelisation naturally provides modularity and portability.

Keywords: High-resolution schemes, Flux limiter, Parallel CFD, Portability, Heterogeneous computing.

1 Introduction

Massively-parallel devices of various architectures are being incorporated to the modern supercomputers to overcome the power constraint in the context of the exascale challenge [1], increasing the variety of hardware significantly. This progress is leading to an increasing hybridisation of high-performance computing (HPC) systems and making the design of computing applications a rather complex problem. To take advantage of the most efficient HPC systems, the computing operations that form the algorithms, the so-called kernels, must be compatible with distributed- and shared-memory SIMD and MIMD parallelism and, more importantly, with stream processing (SP), which is a more restrictive parallel paradigm. Many scientific computing applications have been partially ported, even rewritten entirely, to take advantage of the coprocessor devices (*i.e.* GPUs or MICs). For instance, in [2] the reader can find the solution of incompressible two-phase flows on multi-GPU. Furthermore, examples of heterogeneous implementations of CFD algorithms for hybrid CPU+GPU supercomputations can be found in [3, 4], and an example of a petascale CFD simulation on 18.000 K20X GPUs in [5].

In this context of accelerated innovation, making an effort to design modular applications composed of a reduced number of independent and well-defined code blocks is worth it. On the one hand, this helps

to reduce the generation of errors and facilitates debugging. On the other hand, modular applications are user-friendly and more comfortable for porting to new architectures (the fewer the kernels of an application and its dependencies, the easier it is to provide portability). Furthermore, if the majority of computing kernels represent universal algebraic operations, then both the standard optimised libraries (*e.g.* ATLAS, cBLAST) and the specific in-house implementations can be used and easily switched.

In previous work, Oyarzun et al. [6] proposed a portable implementation model for direct numerical simulations (DNS) and large eddy simulations (LES) of incompressible turbulent flows on unstructured meshes. Roughly, the method consists of replacing traditional stencil data structures and sweeps by algebraic data structures and kernels. As a result, the algorithm of the time-integration phase relies on a reduced set of only three basic algebraic operations: the sparse matrix-vector product, the linear combination of vectors and the dot product. Consequently, this approach combined with a multilevel MPI + OpenMP + OpenCL parallelisation naturally provides modularity and portability.

Inspired by the compelling results of the algebraic implementation, we increased the level of abstraction and presented in [7] the HPC² (Heterogeneous Portable Code for HPC), a fully-portable, algebra-based framework with many potential applications in the fields of computational physics and mathematics. The strategies for the heterogeneous execution of the HPC² kernels were improved and detailed in [8] reporting satisfying strong scalability results on up to 32 nodes of a hybrid supercomputer equipped with a 14-core Intel E5-2697v3 CPU and an NVIDIA Tesla K40M GPU.

In this work, we aim at extending the scientific applications of the HPC² framework providing it with the computational requirements for the simulation of multiphase flows, which is a significant research line within our group [9, 10]. So far, the multiphase flow simulations of our laboratory have been carried out with stencil implementations on CPU supercomputers, and now we want to take a step forward. Specifically, our challenge is to implement a high-resolution scheme (specifically the SUPERBEE flux limiter in [11]) into our heterogeneous computing framework for solving hyperbolic problems in the presence of sharp discontinuities or shocks. For this purpose, the flux limiter must be somehow rewritten in an algebraic form so that stencil operations are avoided during the computation. Therefore, the algebraic formulation of a SUPERBEE flux limiter is described in Section 2. Such a formulation of the flux limiter becomes more challenging in comparison with that of a DNS because of its non-linearity. However, instead of being an inconvenience, this encourages us to demonstrate in Section 3 the high potential of our algebra-based implementation strategy again, showing that only the addition of a few simple algebraic kernels is required. Finally, we challenge in Section 4 the algebraic implementation of the flux limiter for the advection of scalar field with a sharp discontinuity using 2D structured grids in both CPU and GPU platforms.

2 High resolution schemes for the HPC²

Flux limiters are non-linear functions commonly used for solving hyperbolic problems in the presence of sharp discontinuities or shocks. Specifically, flux limiters are used to construct high-resolution schemes with the aim of obtaining second-order approximations and avoiding oscillations near the interface [12]. For instance, a flux limiter scheme is utilised in [13] for the evaluation of the convective terms in the simulation of two-phase flows using the conservative level-set method on unstructured grids.

Let us consider the typical form of a flux limiter for finite volume methods,

$$\theta_f = \theta_U + \Psi(r) \left(\frac{\theta_D - \theta_U}{2} \right), \quad (1)$$

where θ_f is the value of the scalar θ at the face of interest, θ_U and θ_D are upwind and downwind values of θ according to the velocity field u , and $\Psi(r)$ stands for the flux limiter function. The argument r , namely the discontinuity sensor, is chosen as the gradient ratio and is defined as

$$r_f = \frac{\Delta_U \theta}{\Delta_u \theta},$$

where $\Delta_U \theta$ is the gradient of θ at the upwind face and $\Delta_u \theta$ is the gradient at the face of interest. Finally, to facilitate the casting of the flux limiter into an algebraic form, we rewrite the Equation (1) in the less

common form:

$$\theta_f = \frac{\theta_U + \theta_D}{2} + \frac{\Psi(r) - 1}{2} (\theta_D - \theta_U). \quad (2)$$

The operator-based, finite volume discretisation of the Equation (2) is written as follows (the details of the mathematical background and the construction of most of the matrices below is out of the scope of this work, the reader is referred to Trias et al. [14], Valle et al. [15]):

$$\boldsymbol{\theta}_s = (\Pi_{c \rightarrow s} + \Omega(\mathbf{r}_s) \cdot \mathbf{Q}(\mathbf{u}_s) \cdot \Delta_{c \rightarrow s}) \boldsymbol{\theta}_c, \quad (3)$$

where $\boldsymbol{\theta}_s \in \mathbb{R}^m$ and $\boldsymbol{\theta}_c \in \mathbb{R}^n$ are the staggered and centred scalar fields respectively, $\mathbf{r}_s \in \mathbb{R}^m$ is the gradient ratio at the faces, and $\mathbf{u}_s = ((u_s)_1, (u_s)_2, \dots, (u_s)_m)^T \in \mathbb{R}^m$ is the auxiliary discrete staggered velocity which is related to the centered velocity field via a linear interpolation $\Gamma_{c \rightarrow s} \in \mathbb{R}^{m \times dn}$ such that $\mathbf{u}_s \equiv \Gamma_{c \rightarrow s} \mathbf{u}_c$. The size of these vectors, n and m , are the number of control volumes and faces on the computational domain respectively, and d is the number of dimensions of the simulation. The subindices c and s refer to whether the variables are cell-centred or staggered at the faces. The matrices $\Pi_{c \rightarrow s}$ and $\Delta_{c \rightarrow s}$ are constant and represent the scalar cell-to-face interpolator and the scalar cell-to-face difference operator respectively. The matrix $\mathbf{Q}(\mathbf{u}_s)$ is a variable and diagonal matrix which holds the sign of the velocity relative to the normal of the face \mathbf{u}_s . The elements in the diagonal of $\mathbf{Q}(\mathbf{u}_s)$ are recomputed in each time-step as

$$\mathbf{Q}(\mathbf{u}_s) = \text{diag}(\text{sign}(\mathbf{u}_s)). \quad (4)$$

The gradient ratio \mathbf{r}_s , which is the argument for computing $\Omega(\mathbf{r}_s)$, is measured as

$$\mathbf{r}_s(\boldsymbol{\theta}_c) = \frac{(\mathbf{Q}(\mathbf{u}_s)\text{UUD}_{c \rightarrow s} + \text{OUD}_{c \rightarrow s}) \boldsymbol{\theta}_c}{(\mathbf{Q}(\mathbf{u}_s)\Delta_{c \rightarrow s}) \boldsymbol{\theta}_c}. \quad (5)$$

The matrices $\text{OUD}_{c \rightarrow s}$ and $\text{UUD}_{c \rightarrow s}$ are the oriented and unoriented cell-to-face difference operators respectively [15]. The matrix $\Omega(\mathbf{r}_s)$ is a variable and diagonal which represents the term $(\Psi(r) - 1)/2$ of Equation (2). Then, considering the SUPERBEE flux limiter scheme [11], the elements in the diagonal of $\Omega(\mathbf{r}_s)$ become

$$\Omega(\mathbf{r}_s) = \text{diag} \left(\frac{\max(0, \max(\min(1, 2\mathbf{r}_s), \min(\mathbf{r}_s, 2))) - 1}{2} \right). \quad (6)$$

Finally, the complete algorithm for the time-integration of the advection equation using the algebraic formulation of a flux limiter is described in Algorithm 1.

Algorithm 1 Time-integration step of the advection with the SUPERBEE flux limiter

1. Compute the matrix $\mathbf{Q}(\mathbf{u}_s)$ as $\text{diag}(\mathbf{Q}(\mathbf{u}_s)) = \text{sign}(\mathbf{u}_s)$.
 2. Compute the vector $\mathbf{r}_s(\boldsymbol{\theta}_c) = ((\mathbf{Q}(\mathbf{u}_s)\text{UUD}_{c \rightarrow s} + \text{OUD}_{c \rightarrow s}) \boldsymbol{\theta}_c) / ((\mathbf{Q}(\mathbf{u}_s)\Delta_{c \rightarrow s}) \boldsymbol{\theta}_c)$.
 3. Compute the matrix $\Omega(\mathbf{r}_s)$ as $\text{diag}(\Omega(\mathbf{r}_s)) = (\max(0, \max(\min(1, 2\mathbf{r}_s), \min(\mathbf{r}_s, 2))) - 1) / 2$.
 4. Calculate $\boldsymbol{\theta}_c^{n+1}$ with 1st order Euler method: $\boldsymbol{\theta}_c^{n+1} = \boldsymbol{\theta}_c^n - dt \text{DIVU}_s (\Pi_{c \rightarrow s} + \Omega(\mathbf{r}_s)\mathbf{Q}(\mathbf{u}_s)\Delta_{c \rightarrow s}) \boldsymbol{\theta}_c$
-

3 Implementation of the flux limiter into the HPC²

In our previous works [6, 8], we proposed a portable implementation model for direct numerical simulations (DNS) and large eddy simulations (LES) of incompressible turbulent flows. As a result, the algorithm for the time-integration relies on a reduced set of only three basic algebraic operations: the sparse matrix-vector product (`SpMV`), the linear combination of vectors (`axpy`) and the dot product (`ddot`). However, it can be deduced from the Equations (4), (5) and (6) that some new kernels are required to perform element-wise operations over the vectors (*e.g.* an element-wise division is required for computing the gradient ratio as in Equation (5)). Nevertheless, instead of being an inconvenience, this encourages us to demonstrate the

high potential of our algebra-based implementation strategy again, showing that only the addition of six simple algebraic kernels of element-wise operations over the vectors is required to implement the algebraic flux limiter into our fully-portable, algebra-based framework. These new kernels are described below.

$$\begin{aligned}
y = \text{axy}(y, x, a) &\longrightarrow y_i = ay_i/x_i, \\
y = \text{shft}(y, a) &\longrightarrow y_i = y_i - a, \\
y = \text{scal}(y, a) &\longrightarrow y_i = ay_i, \\
y = \text{vmax}, \text{vmin}(y, x) &\longrightarrow y_i = \max, \min(y_i, x_i), \\
y = \text{smax}, \text{smin}(y, a) &\longrightarrow y_i = \max, \min(y_i, a), \\
y = \text{sign}(x) &\longrightarrow y_i = \{-1 \text{ if } x_i < 0, 1 \text{ otherwise}\}.
\end{aligned}$$

The six kernels above do not show appreciable differences regarding their computational behaviour compared with the `axy`. On the one hand, they are simple element-wise operations over the vectors; hence there is no need for communications in distributed-memory parallelisation. Besides, they provide a uniform aligned memory access with coalescing of memory transactions which suit the stream processing paradigm perfectly. On the other hand, the arithmetic intensity of this new kernels (*i.e.* the number of FLOP per byte) is very low, likewise to that of the `axy`; thus their performance is also memory-bounded. Therefore, having already efficient OpenMP, OpenCL and CUDA implementations of `axy`, implementing the six new kernels above is straightforward.

In comparison with the stencil-based, the algebra-based implementation is very similar regarding performance. The distributed-memory parallelisation remains the same since the size of the *halo* (*i.e.* the adjacent elements owned by different MPI processes) depends directly on the partitioning of the computational domain. On the other hand, the stencil-based approach could appear to be slightly more efficient regarding the use of data in some numerical schemes since they can discriminate some values with conditional statements, for instance when locating upwind values. However, these conditional statements may harm the memory access and the coalescing of memory transactions, which is very important in some parallel paradigms such as stream processing. Nevertheless, the data oversize affect only a few operators by a small percentage in the algebra-based implementation.

4 Numerical tests

In this section, the algebraic implementation of the flux limiter is tested for a canonical case. We consider the simulation of the advection of a scalar field with a sharp discontinuity using a high-resolution scheme (see Algorithm 1). Namely, the marker function is initialised in a 2D domain with the shape of a rhodonea [16] for three different structured grids of 32x32, 128x128 and 512x512 cells. Besides, two different fixed velocity fields are evaluated. Namely, a rotating field $\mathbf{u} = (y, -x)$, together with Dirichlet boundary conditions in all the boundaries, and a flat field $\mathbf{u} = (0, 1)$, combined with periodic boundary conditions on top and bottom.

Table 1: Number of times that each kernel is executed per time-step

Step of Algorithm 1	SpMV	axy	axy	shft	scal	vmax, vmin	smax, smin	sign
1 – Compute matrix $\mathbf{Q}(\mathbf{u}_s)$	0	0	0	0	0	0	0	1
2 – Compute gradient ratio	5	1	1	0	0	0	0	0
3 – Compute matrix $\Omega(\mathbf{r}_s)$	0	0	0	1	2	1	3	0
4 – 1st order Euler	6	2	0	0	0	0	0	0
Total number of executions	11	3	1	1	2	2	2	1

To demonstrate the effectiveness of the algebra-based approach, we show in Table 4 the number of times that each algebraic kernel is called every time-step in the numerical Algorithm 1. Then, the comparison of the relative time spent in each operation in both CPU and GPU is shown in Figure 4 (for simplicity, the vector kernels have been grouped). This comparison demonstrates that our implementation model relies almost

entirely on the algebraic kernels. The 96% and 88% of the computational time are employed for running kernels in the CPU and GPU respectively. Furthermore, the performance wasted in the *others* group is due to operations that are not directly involved with the algorithm such as the printing of simulation outputs; and given that the CPU must download the GPU data, and the PCIe bus is slower than the main memory bandwidth, it is consistent that the simulation on GPU spends more time in the *others* group than the CPU.

The flux limiter implementation has been tested on a single node equipped with an Intel i5-2300 and an Nvidia GTX 590, and the results in Figure 4 confirm again that the overall performance is only depending on the performance of the kernels. Hence, the cost estimation of large-scale simulations of multiphase flows on hybrid supercomputers can be extrapolated from the performance study in our previous work [8]. Therefore, a nonsense waste of computational resources in repetitive performance study is avoided.

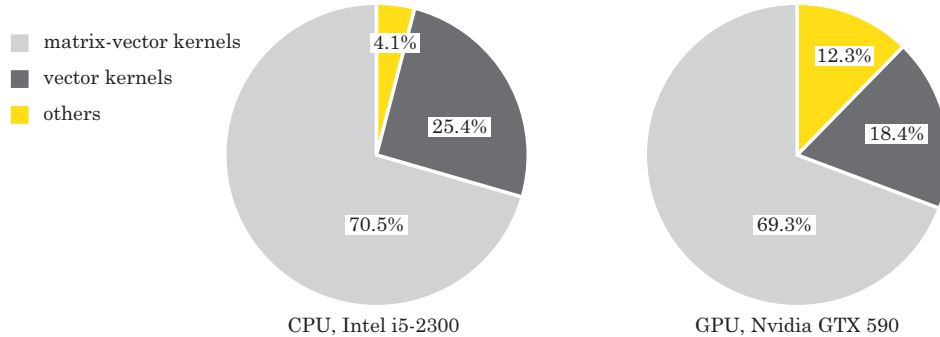


Figure 1: Comparison of the computational cost of the operations in one time step.

The Figure 4 shows various plots of the marker function for the 32x32 grid (top) and the 128x128 (bottom). On the left hand side, the initial state of θ_c is shown. In the center, the final state of θ_c after one cycle under the flat velocity field with periodic boundary conditions. On the right hand side, the final state of θ_c after one complete revolution under the rotating velocity field.

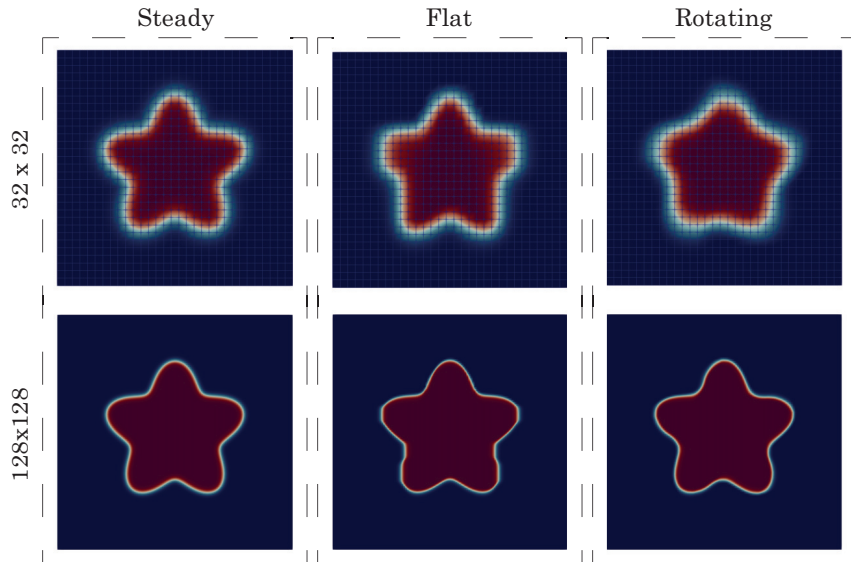


Figure 2: Plots of the marker function θ_c for different grids and cases.

Finally, a study of the evolution of the error with the grid size of both the flat and the rotating simulations is shown in Figure 4 for the structured grids of 32x32, 128x128 and 512x512 cells. The error has been computed as the norm of the difference between the final and initial values of the marker function, normalised

by the norm of the initial value,

$$\epsilon = \frac{\|\boldsymbol{\theta}_f - \boldsymbol{\theta}_i\|}{\|\boldsymbol{\theta}_i\|}. \quad (7)$$

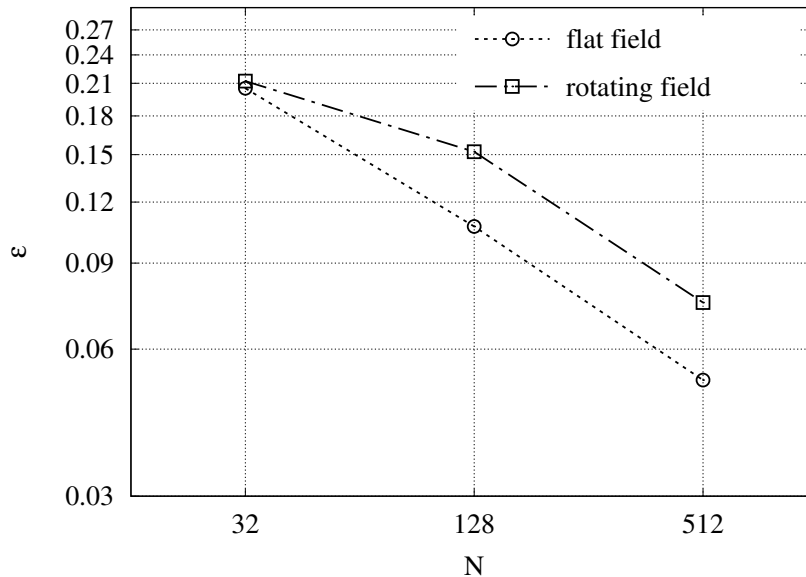


Figure 3: Error versus number of grid divisions per dimension.

5 Conclusion

An algebraic formulation of a high resolution scheme has been presented. A flux limiter has been implemented into the HPC² framework using only simple vector operations. We have shown that the addition of six simple algebraic kernels is sufficient to implement high-resolution, non-linear schemes into our framework. The advection of a scalar field with a sharp discontinuity has been simulated using different 2D structured grids, velocity fields and boundary conditions, and the simulations have been run on both CPU and GPU. Hence, our fully-portable, algebra-based framework for heterogeneous computing has shown a great potential for the simulation of multiphase flows on hybrid supercomputers.

Acknowledgments

The work has been financially supported by the *Ministerio de Economía y Competitividad*, Spain (ENE2017-88697-R and ENE2015-70672-P). X. Á. is supported by a *FI AGAUR* predoctoral contract (2018FI_B1_00081). N. V. is supported by a *FI AGAUR* predoctoral contract (2018FI_B1_000109). F. X. T. is supported by a *Ramón y Cajal* postdoctoral contract (RYC-2012-11996).

References

- [1] Jack Dongarra et al. The International Exascale Software Project roadmap. *Int. J. High Perform. Comput. Appl.*, 25(1):3–60, feb 2011.
- [2] Peter Zaspel and Michael Griebel. Solving incompressible two-phase flows on multi-GPU clusters. *Comput. Fluids*, 80(1):356–364, 2013.
- [3] Andrey Gorobets, F. Xavier Trias, and Assensi Oliva. A parallel MPI+OpenMP+OpenCL algorithm for hybrid supercomputations of incompressible flows. *Comput. Fluids*, 88:764–772, dec 2013.

- [4] Chuanfu Xu, Xiaogang Deng, Lilun Zhang, Jianbin Fang, Guangxue Wang, Yi Jiang, Wei Cao, Yonggang Che, Yongxian Wang, Zhenghua Wang, Wei Liu, and Xinghua Cheng. Collaborating CPU and GPU for large-scale high-order CFD simulations with complex grids on the TianHe-1A supercomputer. *J. Comput. Phys.*, 278(1):275–297, dec 2014.
- [5] Peter E. Vincent, Freddie Witherden, Brian Vermeire, Jin Seok Park, and Arvind Iyer. Towards Green Aviation with Python at Petascale. In *SC16 Int. Conf. High Perform. Comput. Networking, Storage Anal.*, number November, pages 1–11. IEEE, nov 2016.
- [6] Guillermo Oyarzun, Ricard Borrell, Andrey Gorobets, and Assensi Oliva. Portable implementation model for CFD simulations. Application to hybrid CPU/GPU supercomputers. *Int. J. Comput. Fluid Dyn.*, 31(9):396–411, oct 2017.
- [7] Xavier Álvarez, Andrey Gorobets, F. Xavier Trias, Ricard Borrell, and Guillermo Oyarzun. HPC²—A fully-portable, algebra-based framework for heterogeneous computing. Application to CFD. *Comput. Fluids (published online)*, 2018.
- [8] Xavier Álvarez, Andrey Gorobets, and F. Xavier Trias. Strategies for the heterogeneous execution of large-scale simulations on hybrid supercomputers. In *7th Eur. Conf. Comput. Fluid Dyn.*, 2018.
- [9] Néstor Balcázar, Oriol Lehmkuhl, Lluís Jofre, and Assensi Oliva. Level-set simulations of buoyancy-driven motion of single and multiple bubbles. *Int. J. Heat Fluid Flow*, 56:91–107, dec 2015.
- [10] Enrique Gutiérrez, Néstor Balcázar, Eduard Bartrons, and Joaquim Rigola. Numerical study of Taylor bubbles rising in a stagnant liquid using a level-set/moving-mesh method. *Chem. Eng. Sci.*, 164:158–177, jun 2017.
- [11] P. K. Sweby. High Resolution Schemes Using Flux Limiters for Hyperbolic Conservation Laws. *SIAM J. Numer. Anal.*, 21(5):995–1011, 1984.
- [12] Charles Hirsch. *Numerical Computation of Internal and External Flows. Volume 1*. John Wiley & Sons, 2007.
- [13] Néstor Balcázar, Lluís Jofre, Oriol Lehmkuhl, Jesús Castro, and Joaquim Rigola. A finite-volume/level-set method for simulating two-phase flows on unstructured grids. *Int. J. Multiph. Flow*, 64:55–72, sep 2014.
- [14] F. Xavier Trias, Oriol Lehmkuhl, Assensi Oliva, C. D. Pérez-Segarra, and R. W. C. P. Verstappen. Symmetry-preserving discretization of Navier–Stokes equations on collocated unstructured grids. *J. Comput. Phys.*, 258:246–267, feb 2014.
- [15] Nicolas Valle, Xavier Alvarez, F. Xavier Trias, Jesus Castro, and Assensi Oliva. Algebraic implementation of a flux limiter for heterogeneous computing. In *Tenth Int. Conf. Comput. Fluid Dyn.*, 2018.
- [16] M. Oevermann and R. Klein. A Cartesian grid finite volume method for elliptic equations with variable coefficients and embedded interfaces. *J. Comput. Phys.*, 219(2):749–769, dec 2006.