LBcuda: a high-performance CUDA port of LBsoft for simulation of colloidal systems

Fabio Bonaccorso[a,b,*], Marco Lauricella[b,*], Andrea Montessori[c,b], Giorgio Amati[d], Massimo Bernaschi[b], Filippo Spiga[e], Adriano Tiribocchi[b], Sauro Succi[f,b,g]

[a]*Department of Physics and INFN, University of Rome Tor Vergata, Via della Ricerca Scientifica 1, 00133 Rome, Italy*
[b]*IAC-CNR, Via dei Taurini 19, 00185 Rome, Italy*
[c]*Dipartimento di Ingegneria, Università degli Studi Roma TRE, via Vito Volterra 62, Rome, 00146, Italy*
[d]*SCAI, SuperComputing Applications and Innovation Department, CINECA, Via dei Tizii, 6, Rome 00185, Italy*
[e]*NVIDIA Development UK Ltd, Milton Hall, Ely Rd, Milton, Cambridge CB24 6WZ, United Kingdom*
[f]*Center for Life Nano- & Neuro-Science, Fondazione Istituto Italiano di Tecnologia (IIT), 00161 Rome, Italy*
[g]*John A. Paulson School of Engineering and Applied Sciences, Harvard University, 33 Oxford St., Cambridge, MA 02138, USA*

Abstract

We present LBcuda, a GPU accelerated version of LBsoft, our open-source MPI-based software for the simulation of multi-component colloidal flows. We describe the design principles, the optimization and the resulting performance as compared to the CPU version, using both an average cost GPU and high-end NVidia GPU cards (V100 and the latest A100). The results show a substantial acceleration for the fluid solver reaching up to 200 GLUPS (Giga Lattice Updates Per Second) on a cluster made of 512 A100 NVIDIA cards simulating a grid of eight billion lattice points. These results open attractive prospects for the computational design of new materials based on colloidal particles.

[*]Corresponding authors

Email addresses: `fabio.bonaccorso@roma2.infn.it` (Fabio Bonaccorso), `marco.lauricella@cnr.it` (Marco Lauricella)

PROGRAM SUMMARY

Program Title: LBcuda

CPC Library link to program files: (to be added by Technical Editor)

Developer's repository link: https://github.com/copmat/LBcuda

Licensing provisions: 3-Clause BSD License

Programming language: CUDA Fortran

Nature of problem: Hydro-dynamics of colloidal multi-component systems and Pickering emulsions.

Solution method: Lattice-Boltzmann method solving the Navier-Stokes equations for the fluid dynamics within an Eulerian description. Particle solver describing colloidal particles within a Lagrangian representation coupled to the fluid solver. The numerical solution of the coupling algorithm includes the back reaction effects for each force terms according to a fluid-particle multi-scale paradigm.

Keywords: Lattice-Boltzmann, Colloids, CUDA, GPU

---

1. Introduction

In the last two decades, soft-glassy materials (SGM) have gained growing attention due to their applications in several industrial sectors. In particular, emulsions and foams are employed to design novel soft mesoscale materials for chemical, food processing, manufacturing, and biomedical purposes [1, 2, 3]. Besides the technological relevance, their major significant theoretical interest stems from their intriguing non-equilibrium effects, including long-time relaxation, yield-stress behavior, and highly non-Newtonian dynamics.

In this context, computational fluid dynamics (CFD) provides a valuable tool to improve the knowledge of the underlying physics of SGM. To that purpose, a reliable SGM model alongside its software implementation is of apparent interest for the rational designing and shaping up of novel soft porous materials.

In this paper, we present and make available the CUDA Fortran code **LBcuda**, specifically designed to simulate on GPUs bi-continuous systems with colloidal particles under a variety of different conditions. **LBcuda** is a direct port of the LBsoft code [4], an open-source software for simulations of soft glassy emulsions originally developed for CPU-architectures, which successfully combines the lattice-Boltzmann method (LBM) [5, 6, 7] with a

Lagrangian solver to tackle the multi-scale coupling of fluids and particles [8].

Nowadays, the straightforward parallelization of LBM makes the lattice-Boltzmann algorithm an excellent candidate for high-performance CFD, especially on GPU-based architectures, given the relative simplicity and locality of its underlying algorithm. As a consequence, several LBM implementations have been developed for GPU architectures, both academic packages such as the GPU-enabled versions of WaLBerla [9, 10], Palabos [11], Ludwig [12], MUPHY [13], and commercially licensed software such as XFlow 2021 [14], to name a few.

As aforementioned, the model to describe colloidal particles is derived from the previous CPU-based LBSoft code, to which the reader is referred for further details [4]. Briefly, SGM modeling requires specific implementations of LBM and Lagrangian solvers to include the hydrodynamic interactions between solid particles and fluids, following several strategies reported in the literature [15, 16, 17, 18]. This extension has opened the possibility to simulate complex colloidal systems, also referred to as Pickering emulsions [19] which are of primary interest for the rational design of SGM [20, 21, 22, 23]. This intrinsically multi-scale approach can catch, for example, the dynamical transition from a bi-continuous interfacially jammed emulsion gel, also referred to as bijel (see Fig. 1), capturing the associated mechanical and spatial properties [24].

The paper is structured as follows. In Section 2 we report a very brief description of the underlying method, referring our previous paper for a deeper explanation of the details. In Section 3 we describe the details of the data structures on the GPU, while in Section 4 we explain the parallelization strategy and its impact on performance. In Section 5 we report a set of tests used to validate the implementation and we investigate the performance against the reference CPU version (LBsoft). Finally, conclusion and outlook on future development directions are discussed.

2. Method

In this section, we briefly review the approach to the simulation of SGM implemented in LBcuda alongside the more significant algorithmic adaptations required by the GPU-based hardware. A more detailed illustration of the underlying algorithms can be found in Ref. [4].
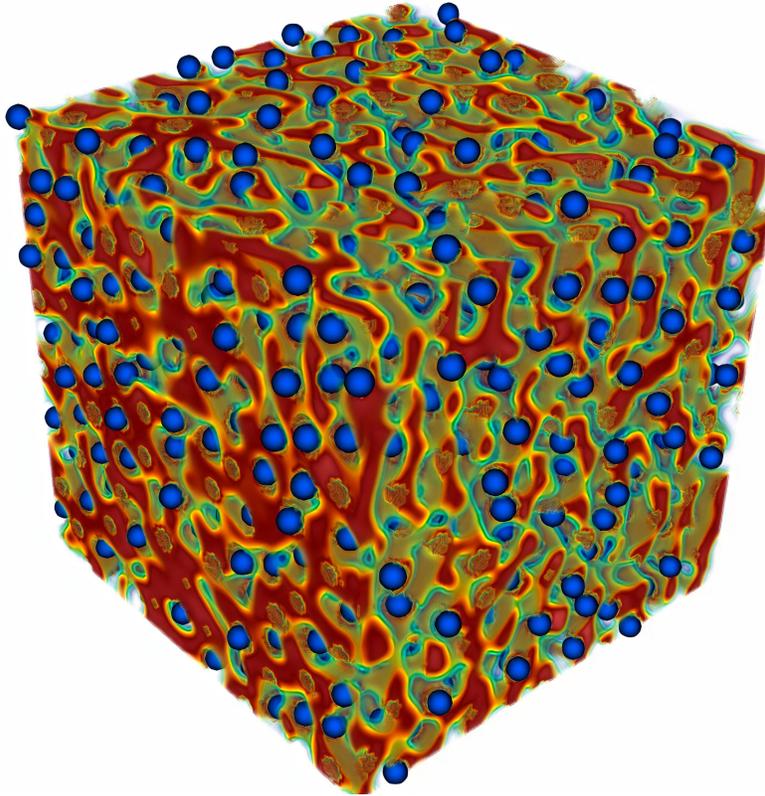
Figure 1: A typical bijel configuration with colloidal particles (blue spheres) entrapping the interfaces between the two fluids (red and transparent green).

The code combines two different levels of description: the first exploits a continuum approach for the dynamics of immiscible fluids, whereas the second manages individual rigid bodies representing colloidal particles or other suspended species. The two levels exchange information at each step of the time integration scheme to describe the concurrent interaction among particles and surrounding fluids.

In the first level, the LBM exploits a fully discretized analog of the Boltzmann kinetic equation to model flows and hydrodynamic interactions in fluids.

In the LBM approach, the fundamental quantity is $f_i(\vec{r}; t)$, namely the probability of finding a "fluid particle" at the spatial point mesh $\vec{r}$ and at time

t with velocity $\vec{c}_i$ selected from a finite set of possible speeds. The LBcuda code implements the 3-dimension 19-speed cubic lattice scheme (D3Q19) with the discrete velocities $\vec{c}_i$ with $i \in [0, ..., 18]$ connecting mesh points with spacing Δx (length lattice unit) to first and second mesh neighbours, located at distance Δx and $\sqrt{2}\Delta x$, respectively (in other words, D3Q19 neglects 8 out of 27 possible velocities: those having distance $\sqrt{3}\Delta x$).

Denoted $\rho(\vec{r}; t)$ and $\vec{u}(\vec{r}; t)$ respectively the fluid density and the fluid velocity, the lattice-Boltzmann equation is implemented in single-relaxation time (Bhatnagar-Gross-Krook equation) as follows:

$$f_i(\vec{r} + \vec{c}_i; t + 1) = (1 - \omega)f_i(\vec{r}; t) + \omega f_i^{eq}(\rho(\vec{r}; t), \vec{u}(\vec{r}; t)) \qquad (1)$$

where f^{eq} is the lattice local equilibrium, basically the local Maxwell-Boltzmann distribution (see Appendix A), and ω is a frequency tuning the relaxation towards the local equilibrium on a timescale $\tau = 1/\omega$. The relaxation frequency ω controls the kinematic viscosity of the fluids according to the relation:

$$\nu = c_s^2 \frac{\Delta x^2}{\Delta t}\left(\frac{1}{\omega} - \frac{1}{2}\right), \qquad (2)$$

where Δx and Δt are the physical length and time of the correspondent counterparts in lattice units. Note that the positivity of the kinematic viscosity requires the condition $0 < \omega < 2$.

In order to model a two component systems we adopted, a color gradient (CG) algorithm, which enforces a diffuse interface between the two fluids [25]. In short, in the update phase of the populations, the CG collision contains three sub-steps: a plain BGK collision, a perturbation operator, and a final recoloring step. It is worth stressing that the last two sub-steps act only near the interface between the two fluids. Further details are reported in Appendix.

The second level of description involves a Lagrangian solver for the particle evolution, where each particle (colloid) is represented by a closed surface $\mathcal{S}$, taken, for simplicity, as a rigid sphere in the following.

The LBcuda code adopts the formulation given by Jansen and Harting [23], where only the exterior regions are filled with fluid, whereas the interior parts of the particles are solid nodes. The solid–fluid interaction is managed via a simple generalization of the bounce-back rule including the correction due to the relative motion of the solid particle with respect to the surrounding fluid medium.

Hence, the particle position, speed $\vec{v}_p$ and angular momentum $\vec{\omega}_p$ are updated according to Newton's equations of motion:

$$\begin{cases} \frac{d\vec{r}_p}{dt} = \vec{v}_p, \\ m_p \frac{d\vec{v}_p}{dt} = \vec{F}_p, \\ I_p \frac{d\vec{\omega}_p}{dt} = \vec{T}_p, \end{cases} \tag{3}$$

where m_p and I_p are the particle mass and moment of inertia, respectively.

Following Ladd's seminal works [18, 26], we advance in time eq. 3 with a leap-frog scheme, which is second order accurate in time. This set of equations considers the full many-body hydrodynamic interactions since the forces and torques are computed with the actual flow, as dictated by the presence of all N particles simultaneously.

3. Implementation

The code is implemented in CUDA Fortran, using modules to minimize code cluttering. The LBcuda code requires no external libraries besides the CUDA runtime and compiles using a simple Makefile. The code is written for the nvfortran compiler, with the GPU kernels confined in *cuf* extension files, whereas the I/O part and the main are coded in standard FORTRAN files.

The code is composed of 6 files:

- **dimension.cuf**, which sets constants for the LB algorithm and the physical values of the simulation

- **kernels_fluid.cuf**, containing all GPU variables

- **kernels_fluid_CG.cuf**, containing the color-gradient GPU code

- **kernels_fluid_PART.cuf**, containing the particles GPU code

- **write_output.f90**, which outputs the VTK and VTI files for external visualization by graphical programs (e.g, ParaView)

- **main.f90**, finally contains the driving code of all the subroutines.

Most of the input for the simulation is defined by setting Fortran **parameters** in dimension.cuf. In contrast, other runtime parameters, such as print

frequency of VTK output files and average statistical quantities, can be set up without recompilation in a plain-text input file.

The data for each fluid component are organized in a five dimension matrix having x,y,z, then the population index, and finally two possible values for *switching* between old and new values during the collide-stream phases of the LB algorithm, also referred to as one-step two-grid algorithm [27].

All data residing on the GPU are defined in kernels_fluid.cuf, whereas writing output files requires just a few memory passages from GPU to CPU at the printing frequency for fluid densities, flow field, and particle positions.

In order to solve the lattice-Boltzmann equation with particle dynamics, the algorithm proceeds executing the following sequence of subroutine calls:

- Each thread of the GPU device computes fluid and particle quantities at a single spatial point, say located at the (i, j, k) node;

- In each node, the code proceeds according to three different cases:

  1. If the node contains fluid far away from any particles, the thread will only advance the LB algorithm;
  2. If at that point a fluid touches a particle, the thread computes its part for the LB algorithm, then it computes its contribute for the force/toque integral. Note that when there are touching particles, a point can contribute to more than one particle.
  3. If the point is inside a particle, no computation is performed and the following steps are skipped;

- Apply the collision step of Eq. 1;

- Apply the halfway bounce-back rule at particle surface and the relative force terms on particles;

- Evolve position and angular velocity of particles (if present);

- Apply the stream step of Eq. 1.

It is worth stressing that the net force and torque exerted from the fluid on the particle center of mass is obtained by summing over all the particle surface nodes.

4. Parallelization strategy: CUDA and MPI

For the LB part of the algorithm, the CUDA porting decomposes the global domain according to a 3D block distribution among the CUDA threads (see Fig. 2). The selection of the block distribution is fixed at compile-time, and it can be tuned to obtain the best performance given the global grid dimension and the compute capability of the GPU device. Usually, a 3D decomposition that emphasizes the x-axis dimension achieves the best performance since it exploits the high memory bandwidth due to the data continuity in the column-major order of the FORTRAN language. Each thread will be responsible for only one grid point of the fluid box in each CUDA block. In the LB part, the thread iterates over all the fluid populations in most kernels. On the other hand, the thread computes the contribution of the fluid grid point to the force and torque of the overlapping particles, defined as particles whose surface overlaps the fluid node owned by the thread.

The LBcuda code is designed to exploit multiple GPU devices. To that purpose, the code resorts to MPI having one GPU card associated to each MPI task. The LB domain is divided into sub-domains of equal size, whereas the variables related to the particles are replicated in all the MPI processes. Thus, the LB solver proceeds locally on each GPU device, with the extra computational cost due to the communication of border information among the local sub-domains of the neighbor MPI tasks.

Each MPI task computes the part (section) of the particles falling in its sub-domain in the particle solver. In this framework, particles evolution is crucial for achieving a high computational throughput by avoiding excessive communication (or memory conflicts) among threads while integrating the particle quantities. Thus, we have adopted a single particle list, which is stored on GPU devices. In particular, whenever the multi GPU is used, each MPI sub-domain has its list of owned particles on its GPU card. When a LB time step is completed, each thread in the sub-domain checks if it needs to compute the contribution of its fluid node to the computation of the surface integral for the force and torque that the fluid exerts on each particle surface node and vice versa. Hence, a global MPI reduction is used to compute the corresponding total force and torque acting on the center of mass of each particle, so that particle positions, orientations, and velocities can be advanced in time on all the MPI processes. The selection of the sub-domain particle list on each MPI process is made by a CUDA kernel, leveraging the parallel computing power of each GPU card that makes the
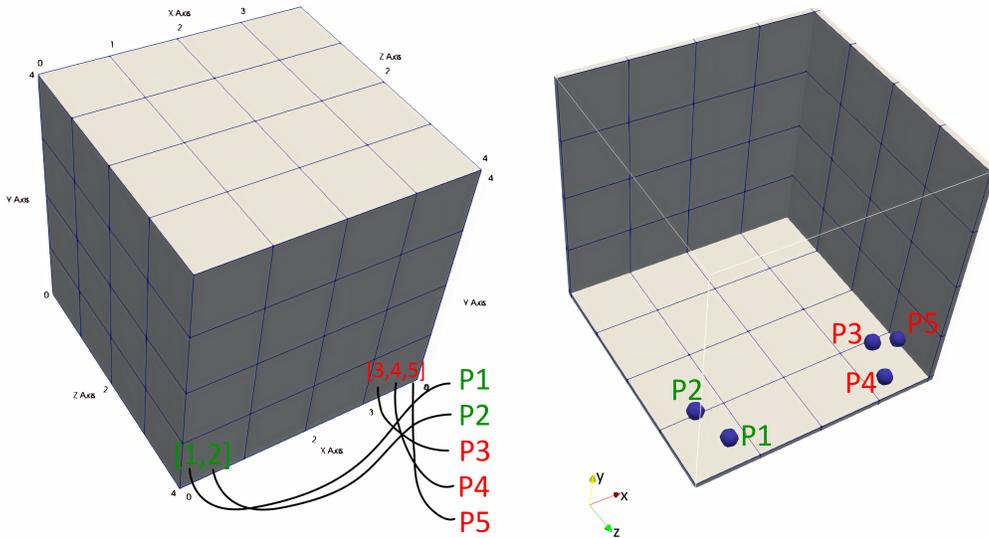
Figure 2: Sketch showing the domain decomposition strategy used for the particle data on the GPU device. On the left, each sub-domain (thread block) has a list of owned particles. On the right, the global GPU data vector stores all the particles in contiguous way over the sub-domains.

required computing time almost negligible.

It is worth highlighting that the overlap between particle and fluid node provides an unbalance in the work performed by the CUDA threads, notably wasting computational power at the surface particle node to treat the boundary conditions and the momentum exchange between the surrounding fluid and the particle. Nonetheless, we found that the overhead for the fluid solver is quite marginal, retaining an acceptable code scalability.

5. Performance results

In order to analyze the computational performance of LBcuda, we consider two cubic boxes of side 128 and 256 lattice nodes with periodic boundary along the three Cartesian axes. We limited the size to a 256 cubic box, which is the largest grid that fits in the memory of a single GPU device. We perform three different test cases. First, we examine the case of a single fluid in a cubic box with an initial density equal to one and zero velocity flow (*case* 1). As a second test (*case* 2), a bi-component system is considered where all the

fluid nodes are randomly filled with fluid mass density of the two fluids, red and blue component, to achieve the value of the order parameter, $\phi = \frac{\rho_r - \rho_b}{\rho_r + \rho_b}$, equal to 1 or -1 with zero velocity flow field with the subscripts r and b standing for "red" and "blue" fluids, respectively. The third test (*case* 3) checks the entire LBcuda algorithm with the two-component fluid combined with the particle solver describing the colloids in a rapid demixing emulsion. In particular, we defined three sub-setups with different numbers of particles to assess the performance of the particle solver. Hence, three values of the volume fraction occupied by particles are considered: 0.1%, 1.0% and 10%, labeled *case* 3*a*, *case* 3*b*, and *case* 3*c*, respectively.

To evaluate the performance of LBcuda, we compare the theoretical peak performance to the actual one achieved by our code. In particular, the roofline model [28] is used to rank the achievable computational performance in terms of Operational Intensity (OI), defined as the ratio between flops performed and data that need to be loaded/stored from/to memory. At low OI (say, < 10), the performance is limited by the memory bandwidth, whereas for higher OI values, the limitation comes from the availability of floating-point units. It is well known that LB is a bandwidth-limited numerical scheme, like most CFD models [29]. The OI index for LB schemes is around 0.7 for double-precision (DP) simulations using a D3Q19 lattice. As a matter of fact, for a single fluid, since the number of floating-point operations per lattice site and time step is $F \simeq 200 \div 250$ and the load/store burden in bytes is $B = 19 \times 2 \times 8 = 304$ (using double precision), the operational intensity is $F/B \sim 0.7$, whereas in single precision is $F/B \sim 1.4$ confirming that the code is bandwidth limited (see also Figure 3).

In the following, we assess the efficiency of the LBcuda code by means of the Giga Lattice Updates Per Second (GLUPS) metrics. In particular, the definition of GLUPS reads:

$$\text{GLUPS} = \frac{L_x L_y L_z}{10^9 t_s}, \tag{4}$$

where L_x, L_y, and L_z are the domain sizes in the $x-$, $y-$, and $z-$ axis, and t_s is the run (wall-clock) time (in seconds) per single time step iteration.

5.1. Single fluid

The *case* 1 with one fluid is tested using a plain BGK collision and a fused implementation (in which the collision and streaming step are performed simultaneously) of the LB time-advancing, the latter being the most popular

approach for major LB codes. The results are reported for two different GPUs: a V100 and a GeForce RTX 2060S in Table 1.

Size	GPU	Time	GLUPS	Approach
128^3	V100	1.09 ms/iter	1.923	fused
128^3	2060S	1.64 ms/iter	1.278	fused
256^3	V100	8.34 ms/iter	2.011	fused
256^3	2060S	15.6 ms/iter	1.075	fused
128^3	V100	2.09 ms/iter	1.003	plain
128^3	2060S	3.14 ms/iter	0.667	plain
256^3	V100	15.34 ms/iter	1.093	plain
256^3	2060S	30.1 ms/iter	0.557	plain

Table 1: Timings alongside with GLUPS of 128^3 and 256^3 cubic boxes for the *case* 1 using both the optimized fused and plain approach single fluid on a Tesla V100 and a GeForce RTX 2060S in single precision.

The key differences between the two GPUs are: the V100 has 5132 Cuda cores offering a peak performance of 14TFlop/s and a memory bandwidth of 900GB/s, whereas the GeForce RTX 2060S has 2176 cores for a peak of 6.4TFlop/s and a memory bandwidth of 448GB/s. On the other hand, we observe that the obtained 2.011 GLUPS for a cubic box of side 256 is comparable with the state-of-the-art represented, for instance, by the highly optimized code by G. Falcucci et al. in Ref. [30] which reaches 3.406 GLUPS on a single V100 with the fused implementation on the same cubic box size.

Although the fused approach reduces of a factor two the number of memory accesses, it is worth highlighting that the particle solver requires the mandatory use of the plain approach, showing a decrease of the performance of about a factor two for *case* 1. Indeed, the particle boundaries require using the plain Lattice-Boltzmann algorithm, where the collision is first computed for all the fluid lattice nodes. Then, the boundary conditions are applied (internal walls or particles), and finally the streaming of the populations is carried out (see Section 3).

From an in-depth analysis using the NVidia dedicated tool (NSight), we observe that on the GeForce RTX 2060S the main kernel (the LB time-stepping before the streaming substep)128x1x1 achieves better performance achieves almost $\sim 520 GFlop/s$ with an arithmetic intensity of 2.0 in single precision (see Figure 3). This result shows that we are far from an intensity ~ 15.0 which should give the peak performance, and in the *case* 1 we at-
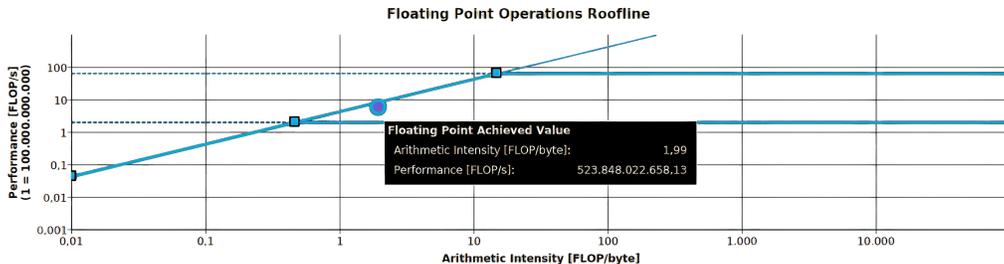
Figure 3: Roofline model for single fluid using a GeForce RTX 2060S, as measured by NVIDIA NSight in the plain LB approach and in single precision. Bandwidth and Float point computation limits were obtained performing *memory-stream* and High Performance Computing Linpack benchmark. Note that the LBcuda code lies on the left part of the plot showing that it is bandwidth limited.

tain about 60% of memory bandwidth utilization due to its non-optimal use following the plain approach.

5.2. Two fluid test

The *case* 2 is related to the simulation of a two-component system by the color-gradient model (see Section 2). We remark that the time integration is implemented using the plain approach with a standard collide-stream 2-pass algorithm. Table 2 highlights the measured performance for the 128^3 and 256^3 cubic box. In particular, we observe an increase of about a factor 3 with respect to the previous *case* 1, which is mainly due to the larger number of operations, more than doubled, in the color gradient collision operator containing three steps (see Section 2) instead of the single step of the plain BGK single fluid case.

12

Size	GPU	CUDA Block	Time	GLUPS
128^3	V100	8x4x4	4.5 ms/iter	0.466
128^3	2060S	8x4x4	11 ms/iter	0.190
256^3	V100	8x4x4	34.1 ms/iter	0.492
256^3	2060S	8x4x4	82.3 ms/iter	0.204
128^3	V100	128x1x1	3.96 ms/iter	0.554
128^3	2060S	128x1x1	8.96 ms/iter	0.245
256^3	V100	128x1x1	27.9 ms/iter	0.601
256^3	2060S	128x1x1	64.5 ms/iter	0.261

Table 2: Timings alongside with GLUPS for 128^3 and 256^3, 2 fluids with CG using a Tesla V100 and a GeForce RTX 2060S with the plain approach in single precision with two different decompositions of threads in CUDA block. Note that the CUDA block configuration 128x1x1 achieves better performance exploiting the contiguous data over x in the population arrays.

5.3. Particles

The entire LBcuda algorithm is evaluated with the two-component colour gradient method (see Section 2) combined with the particle solver to model the colloids in a rapid demixing emulsion. To assess the performance of the particle solver, we prepared three simulation setups with different numbers of particles with radius equal to 5.5 lattice units in a cubic box of side 256 lattice points. The three cases, in the following labelled *case* 3a, 3b, and 3c, differ in the ratio between the volume occupied by the particles compared to the box volume equal to $\varphi = 0.1\%$, 1.0% and 10%, respectively, in order to study the impact of the particle evolution on the simulation time. The performance impact of having particles goes from negligible in *case* 3a to be comparable with the LB computation in *case* 3c in which the time for each iteration almost doubles, as shown in table 3.

case	GPU	φ	Time	GLUPS
3a	V100	0.1 %	34.6 ms/iter	0.484
3a	2060S	0.1 %	84 ms/iter	0.199
3b	V100	1.0 %	37.6 ms/iter	0.446
3b	2060S	1.0 %	88 ms/iter	0.190
3c	V100	10 %	56 ms/iter	0.299
3c	2060S	10 %	113 ms/iter	0.148

Table 3: Timings alongside with GLUPS for the three cases with particle volume fraction, φ, equal to 0.1%, 1.0%, and 10%, respectively, in a cubic box of 256^3 lattice nodes in single precision.

It is worth highlighting that LBcuda code implements a double precision accumulator because of floating point accuracy problems related to the momentum transfer from the fluid to each particle. In particular, the particle force and torque computation suffer from floating accuracy problems due to strong cancellation between addends of alternating sign over the nodes of the two-fluids interface.

The benchmark results can be also compared to the corresponding box size without particles reported in Table 2. For the test in *case* 3 the particle radius is equal to 5.5 lattice units, and the particle positions are randomly distributed in the box. The initial particle velocity is zero in all runs. The particle wettability is tuned to set an angle equal to 90° with respect to the axis $\vec{x}^{\star}$ in the local reference frame of each particle. The impenetrability among particles is avoided by an hertzian repulsive contact force computed by means of neighbor's lists (see Section 4). The lubrication force is also considered by adding an extra force term whenever two particles are located at a mutual distance lower than 2/3 lattice unit, as reported in previous simulations [23]. The particle mass was estimated as the weight corresponding to a particle made of silica [31].

All runs were simulated on both the CPU and GPU architectures using the previous LBSoft code and the corresponding GPU ported version, LBcuda. In all the *cases* 3 we observe only a small deviation in the position always lower than 10^{-4} in lattice units, mainly due to the aforementioned floating point accuracy problem. Indeed, the order of the addends over the particle surface nodes is completely random on the GPU device. In the *case* 3c we observe the arrest of the phase demixing process with particles located at the fluid-fluid interface entrapping the demixing process into a
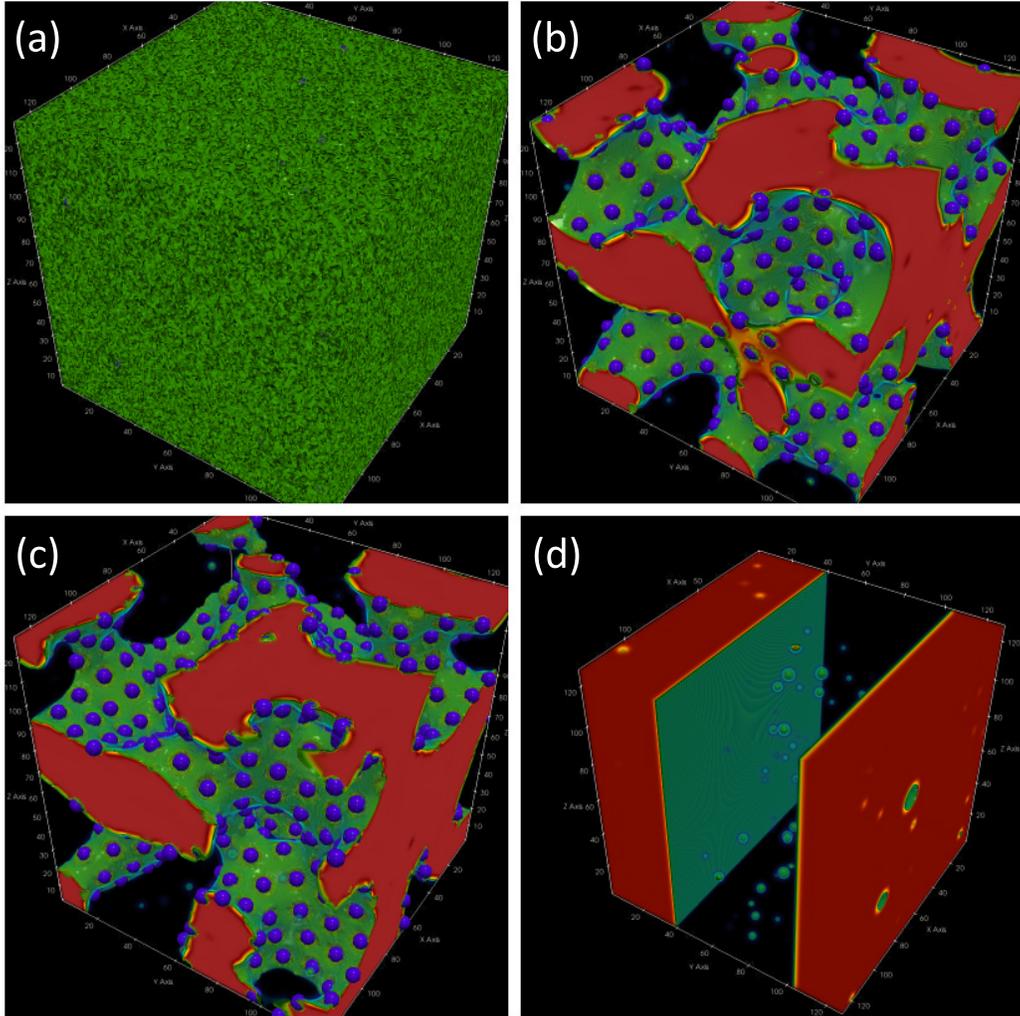
Figure 4: Renderings for 256^3 simulation without (plot d) and with particles (plot a,b,c). From top-left to bottom-right) a) Initial condition for particle simulation. b) Density field after 50k iterations with 10% volume fraction. c) Density field after 100k iterations (10% vol. fraction). d) Density field after 100k iterations without particles starting from a mixed bi-component fluid system (similar to plot a).

metastable state, the bi-continuous jammed gel state [32]. In Figure 4 we show two of these numerical experiments, in which a random mixture of two fluids evolves in a very different way in the presence of a high number of particles (*case* 3c) and without particles (*case* 2). Indeed, the particles stop

15

the complete spinodal decomposition of the two fluids corresponding to the condition of minimal energy and minimal interface between them. Whenever a high volume fraction of particles is in the simulation box, the separation surface becomes way more "corrugated" showing the formation of the bi-continuous jammed gel state. Figure 4 shows the initial condition, the fluids after 50k (with particles), and final configurations of both simulations after 100k iterations (*case* 3*c* and *case* 2).

5.4. Multi GPUs Performance

The LBcuda code resorts to the Message Passing Interface (MPI) library to exchange data among GPU devices running in parallel. The performance obtained using multiple GPUs show a good scaling behavior as reported in Tables 4 and 5. The benchmark in Table 4 was carried out on Marconi100 at CINECA, a cluster of V100 cards, each with 16 GB of global memory using a different number of GPU devices, for three cubic boxes: 256^3, 512^3, and 1024^3. The cluster is made of nodes, each endowed with four GPU cards so that the MPI communication does not incur network latency unless the job is using more than 4 GPU cards. The benchmark in Table 5 ran on a cluster made of NVIDIA DGX A100. Each NVIDIA DGX A100 is equipped with 8 A100 NVIDIA GPU with 80 GB of RAM interconnected intra-node through the NVswitch and 8 NVIDIA Infiniband HDR (one NIC for each GPU) for multi-node scaling.

Grid	1	2	4	8	16	32
$256^3 V100\ 16GB$	0.60	1.17	2.24	2.93	2.71	3.15
$512^3 V100\ 16GB$	x	x	2.67	4.57	6.23	9.26
$1024^3 V100\ 16GB$	x	x	x	x	x	12.48

Table 4: Performance for 256^3, 512^3, 1024^3, measured in GLUPS running on multiple NVIDIA Volta V100 GPUs (each card equipped with 16 GB of RAM). Note that the symbol (x) denotes a grid size too big to fit in the local GPU memory. For the 256^3, performance degrade dramatically when using more than 8 V100 GPUs and more than 64 A100 because each local domain becomes too small (256x256x8) to offset the cost of scheduling the GPU task.

All benchmarks reported in Tables 4 and 5 were carried out on a bi-component fluid system without particles (*case* 2). The data for the cubic box sizes 256^3, 512^3 are also plotted in Figure 5 with an evident decrease in the scalability whenever the code starts to run on more than one node

Grid	8	16	32	64	128	256	512
$256^3 A100\ 80GB$	6.17	9.42	12.38	15.38			
$512^3 A100\ 80GB$	7.63	14.13	23.13	36.05			
$1024^3 A100\ 80GB$	11.65	20.57	28.59	49.11	76.91		
$2048^3 A100\ 80GB$	x	x	x	55.22	102.3	160.9	204.5

Table 5: Performance for 256^3, 512^3, 1024^3, 2048^3 measured in GLUPS running on multiple NVIDIA Volta A100 GPUs (each card equipped with 80 GByte of RAM). Note that the symbol (x) denotes a grid size too big to fit in the local GPU memory.

(more than 4 GPUs). Hence, the three tests, *case 3a*, *3b*, and *3c*, with different values in the particle volume fraction, φ, were performed to probe the efficiency of the particles solver as a function of the particles number in the system. All benchmarks were carried out on eight GPU cards: both V100 cards with a 16 GB RAM and A100 cards with a 40 GB RAM. In Table 6, we observe a clear communication penalty as the particles number increases that is mainly due to the replicated data strategy used for the particle solver parallelization. Indeed, the replicated data parallel approach replicates the physical quantities of all particles across the MPI processes, performing local updating and global MPI sum reductions in order to advance the system in time, which decreases the performance as the quantity of particles data increases. Nonetheless, the analysis of the performance as a function of the number of particles shows that the code is able to reach about 3.80 GLUPS in a system with 10% in the particle volume fraction.

GPU	No particles	$\varphi = 0.1\%$	$\varphi = 1.0\%$	$\varphi = 10\%$
8 V100@16	30 ms (4.47)	40 ms (3.35)	51 ms (2.63)	95 ms (1.41)
8 A100@40	18 ms (7.55)	19 ms (7.06)	22 ms (6.10)	35 ms (3.83)

Table 6: Time per single iteration alongside with GLUPS in parenthesis for 512^3 without and with particles (at different particle volume fraction φ) on 2 different machines: using eight V100 with 16 GB of RAM on 2 nodes (connected by InfiniBand) and using eight A100 with 40 GB in a single node (without the latency time due to the InfiniBand communication).

5.5. Comparing LBcuda with LBsoft

For the sake of completeness, we probe the gain provided by the CUDA port reported in the present article. Thus, the *case 2* bi-component system
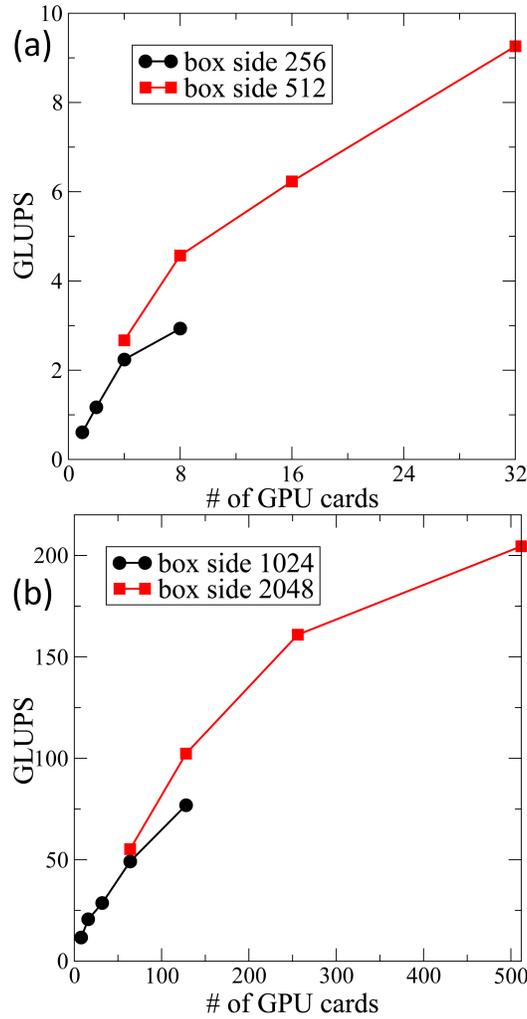
Figure 5: Top panel: Measured GLUPS for different numbers of V100 cards with a 16 GB RAM, running two cubic boxes of side 256 and 512, respectively. Bottom panel: Measured GLUPS for different numbers of A100 cards with a 80 GB RAM, running two cubic boxes of side 1024 and 2048, respectively.

was initialized with the same values in density and flow field in both LBcuda and LBsoft code. The cubic box size is equal to 512^3 lattice points. Although it is difficult to compare two completely different computing architectures, we measure the wall-clock time obtained on a GPU cluster made of V100 cards with 16 GB RAM and a CPU cluster containing two Intel Cascade Lake 8260

18

CPUs at 2.40 GHz with 48 cores and 384 GB of RAM per node. We note that the code produces the same results in single precision unless a slight difference of 10^{-4} order of magnitude in the particle positions due to floating point accuracy problems (mainly due to different order of summation). The wall-clock time for iteration results in 12 ms on 32 V100 cards of LBcuda versus 88 ms on 528 cores of LBSoft, confirming the clear advantage in running the CUDA ersion on a GPU HPC cluster.

6. Conclusion

We have presented LBcuda, a CUDA port of LBsoft, an open-source software aimed at simulating specifically colloidal systems. LBcuda is written in CUDA Fortran and permits to simulate large system sizes running on multiple GPU devices by exploiting an efficient parallel domain decomposition implementation.

In particular, the code shows good scaling behavior of the fluid solver achieving the performance peak of 200 GLUPS on 512 NVIDIA A100 cards with a grid of eight billion lattice points. On the other hand, the particle solver combined with the LB approach shows a very satisfactory performance in terms of scalability in both system size and number of processing cores, especially using the Nvidia Ampere A100 cards.

In this work, the main structure of LBcuda has been outlined along with the key steps of its implementation. Furthermore, several cases have been introduced to test the code over typical problems that the LBcuda code can deal with. In particular, the simulations with particles demonstrate the capabilities of the present code to reproduce the complex dynamics of bi-jel systems in a rapid de-mixing emulsion.

The LBcuda code is open source and completely accessible at the public repository GitHub, which is in line with the spirit of open-source software, mainly to promote the contribution of independent developers.

7. Acknowledgments

19

the ISCRA project "porting LBSOft in CUda on multi-node GPUs (LB-SOCU)".

Appendix

For the details of the color gradient (GC) model of the Lattice Boltzmann method employed in the bi-component systems [25], we recall some notions.

In the color gradient LB for two-component flows, two sets of distribution functions are defined to track the evolution of the two fluid components, which occurs via a streaming-collision algorithm:

$$f_i^k\left(\vec{x} + \vec{c}_i\Delta t,\, t + \Delta t\right) = f_i^k\left(\vec{x},\, t\right) + \Omega_i^k(f_i^k\left(\vec{x},\, t\right)), \tag{5}$$

where f_i^k is the discrete distribution function, representing the probability of finding a particle of the k^{th} component at position $\vec{x}$ and time t with discrete velocity $\vec{c}_i$.

In the last Eq. i is the index running over the lattice discrete directions $i = 0, ..., b$, where $b = 18$ for a three dimensional 19 speed lattice (D3Q19) implemented in LBcuda. The lattice time step Δt has been taken as 1 (in lattice units) for convenience. The density ρ^k of the k^{th} component is given by the zeroth moment of the distribution functions:

$$\rho^k\left(\vec{x},\, t\right) = \sum_i f_i^k\left(\vec{x},\, t\right), \tag{6}$$

while the total fluid density is assessed as $\rho = \sum_k \rho^k$, and the total momentum of the mixture is given as the sum of the linear momentum of the two components:

$$\rho\vec{u} = \sum_k \sum_i f_i^k\left(\vec{x},\, t\right)\vec{c}_i. \tag{7}$$

The collision operator in the CG model is made of three parts:

$$\Omega_i^k = \left(\Omega_i^k\right)^{(3)}\left[\left(\Omega_i^k\right)^{(1)} + \left(\Omega_i^k\right)^{(2)}\right]. \tag{8}$$

In the above, $\left(\Omega_i^k\right)^{(1)}$ stands for the standard collisional relaxation which reads:

$$\left(\Omega_i^k\right)^{(1)} = \omega(f_i^{k,eq} - f_i^k), \tag{9}$$

where $\omega = 2/(6\bar{\nu} - 1)$ is the effective relaxation parameter being $\bar{\nu}$ the mean viscosity of the bi-component system computed as $\frac{1}{\bar{\nu}} = \frac{\rho_1}{(\rho_1 + \rho_2)}\frac{1}{\nu_1} + \frac{\rho_2}{(\rho_1 + \rho_2)}\frac{1}{\nu_2}$

20

(ν_1 and ν_2 are the kinematic viscosities of the two pure components in the bulk). The equilibrium distribution function of the k^{th} component $f_i^{k,eq}$ is given by a low-Mach, second-order, expansion of a local Maxwellian, namely:

$$f_i^{k,eq} = w_i \rho^k (1 + \frac{\vec{c}_i \cdot \vec{u}}{c_s^2} + \frac{(\vec{c}_i \cdot \vec{u})^2}{2c_s^4} - \frac{\vec{u} \cdot \vec{u}}{2c_s^2}), \tag{10}$$

where $c_s = 1/\sqrt{3}$ is the sound speed of the model. The symbol $\left(\Omega_i^k\right)^{(2)}$ denotes the perturbation step, which contributes to the build up of an interfacial tension. Finally, $\left(\Omega_i^k\right)^{(3)}$ is the recoloring step, which promotes the segregation between species, so as to minimize their mutual diffusion.

In order to reproduce the correct form of the stress tensor, the perturbation operator can be constructed by exploiting the concept of the continuum surface force. Firstly, the perturbation operator must satisfy the following conservation constraints:

$$\sum_i \left(\Omega_i^k\right)^{(2)} = 0 \tag{11}$$

$$\sum_k \sum_i \left(\Omega_i^k\right)^{(2)} \vec{c}_i = 0 \tag{12}$$

By performing a Chapman-Enskog expansion, it can be shown that the hydrodynamic limit of Eq.5 is represented by a set of equations for the conservation of mass and linear momentum:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot \rho \vec{u} = 0 \tag{13}$$

$$\frac{\partial \rho \vec{u}}{\partial t} + \nabla \cdot \rho \vec{u} \vec{u} = -\nabla p + \nabla \cdot [\rho \nu (\nabla \vec{u} + \nabla \vec{u}^T)] + \nabla \cdot \boldsymbol{\Sigma} \tag{14}$$

where $p = \sum_k p_k$ is the pressure and $\nu = c_s^2 (\tau - 1/2)$ is the kinematic viscosity of the mixture, being τ the single relaxation time.

The stress tensor in the momentum equation is given by:

$$\boldsymbol{\Sigma} = -\tau \sum_i \sum_k \left(\Omega_i^k\right)^{(2)} \vec{c}_i \vec{c}_i \tag{15}$$

Since the perturbation operator is responsible for generating interfacial tension, the following relation must hold:

21

$$\nabla \cdot \mathbf{\Sigma} = \vec{F} \tag{16}$$

Denoting $\Theta = (\rho^1 - \rho^2)/(\rho^1 + \rho^2)$ the phase field, by choosing the second operator as:

$$\left(\Omega_i^k\right)^{(2)} = \frac{A_k}{2}|\nabla\Theta| \left[w_i \frac{(\vec{c}_i \cdot \nabla\Theta)^2}{|\nabla\Theta|^2} - B_i \right], \tag{17}$$

substituting it into Eqs 11 and 16 and by imposing that the set B_i must satisfy the following isotropy constraints

$$\sum_i B_i = \frac{1}{3} \; ; \sum_i B_i \vec{c}_i = 0 \; ; \sum_i B_i \vec{c}_i \vec{c}_i = \frac{1}{3}\mathbf{I}, \tag{18}$$

we obtain an equation for the surface tension of the model:

$$\sigma = \frac{2}{9}(A_1 + A_2)\frac{1}{\omega} = \frac{4}{9}A\frac{1}{\omega}. \tag{19}$$

The above relation shows a direct link between the surface tension and the parameter A with $A_1 = A_2$. In actual practice, after choosing the viscosity of the two components and the surface tension of the model, at each time step, one locally computes the A coefficient by using the formula reported in eq. (19).

As pointed out above, the perturbation operator generates an interfacial tension in compliance with the capillary-stress tensor of the Navier-Stokes equations for a multicomponent fluid system.

Nonetheless, the perturbation operator alone does not guarantee the immiscibility of different fluid components. For this reason, a further step is needed (i.e. the recoloring step) to minimize the mutual diffusion between components.

Following the work of Latva-Kokko and Rothman, the recoloring operator for the two sets of distributions takes the following form:

$$\left(\Omega_i^1\right)^{(3)} = \frac{\rho^1}{\rho}f_i^* + \beta\frac{\rho^1\rho^2}{\rho^2}\cos\phi_i f_i^{eq,0} \tag{20}$$

$$\left(\Omega_i^2\right)^{(3)} = \frac{\rho^2}{\rho}f_i^* - \beta\frac{\rho^1\rho^2}{\rho^2}\cos\phi_i f_i^{eq,0}, \tag{21}$$

D3Q19	$\{i : |c|^2 = 0\}$	$\{i : |c|^2 = 1\}$	$\{i : |c|^2 = 2\}$
w_i	1/3	1/18	1/36
B_i	-2/9	1/54	1/27

D3Q19 lattice velocity and weights [25].

where $f_i^* = \sum_k f_i^{k,*}$ denotes the set of post-perturbation distributions, $\rho = \rho^1 + \rho^2$, $\cos \phi_i$ is the angle between the phase field gradient and the i^{th} lattice vector and $f_i^{eq,0} = f_i(\rho, \vec{u} = 0)^{eq} = \sum_k f_i^k(\rho, \vec{u} = 0)^{eq}$ is the total zero-velocity equilibrium distribution function. In Eq. 21, the coefficient β is a free parameter which tunes the interface width, thus playing the role of an inverse diffusion length scale. The coefficients used in Eqs 10 and 17 are reported in Table.

For the details of the particle time evolution, a full explanation is in Ref. [4]. For the sake of completeness, the main key steps are outlined in the following.

The velocity at a boundary node of p-th particle is given by:

$$\vec{u}_b = \vec{v}_p + \vec{r}_b \times \vec{\omega}_p, \tag{22}$$

where, $r_b = \vec{r}_s + \frac{1}{2}\vec{c}_i$ is the location of the moving wall along the i-th link connecting the solid node $\vec{r}_s$ to the fluid node $\vec{r}_i = \vec{r}_s + \vec{c}_i$. All coordinates are relative to the center of the p-th particle, located at position $\vec{r}_p$ and moving with translation and angular velocities $\vec{v}_p$ and $\vec{\omega}_p$, respectively.

The timestep is made unit for simplicity. This velocity sets the bias between colliding pairs:

$$f_i(\vec{r} + \vec{c}_i, t + 1) = f_{\bar{i}}(\vec{r} + \vec{c}_i, t') + 2\rho w_i u_{bi}, \tag{23}$$

$$f_{\bar{i}}(\vec{r}, t + 1) = f_i(\vec{r}, t') - 2\rho w_i u_{bi} \tag{24}$$

where t' denotes post-collisional states and we have set

$$u_{bi} = \frac{\vec{u}_b \cdot \vec{c}_i}{c_s^2}$$

Note that these rules reduce to the usual bounce-back conditions for a solid at rest, $\vec{v}_p = \vec{\omega}_p = 0$.

These collision rules produce a net momentum transfer between the fluid and the solid site:

$$\vec{F}_i(\vec{r}_b, t + \frac{1}{2}) = 2\vec{c}_i[f_i(\vec{r}, t') - f_{\bar{i}}(\vec{r}_i, t') - 2\rho w_i u_{bi}] \tag{25}$$

The net force acting upon particle p is obtained by summing over all boundary sites $\vec{r}_b$ and associated interacting links, namely:

$$\vec{F}_p = \sum_{\vec{r}_b, i \in \Sigma_p} \vec{F}_i(\vec{r}_b). \tag{26}$$

Similarly, the total torque $\vec{T}_p$ is computed as

$$\vec{T}_p = \sum_{\vec{r}_b, i \in \Sigma_p} \vec{F}_i(\vec{r}_b) \times \vec{r}_b, \tag{27}$$

References

[1] A. Fernandez-Nieves, A. M. Puertas, Fluids, Colloids and Soft Materials: An Introduction to Soft Matter Physics, volume 7, John Wiley & Sons, 2016.

[2] R. Piazza, Soft matter: the stuff that dreams are made of, Springer Science & Business Media, 2011.

[3] R. Mezzenga, P. Schurtenberger, A. Burbidge, M. Michel, Understanding foods as soft materials, Nature materials 4 (2005) 729.

[4] F. Bonaccorso, A.Montessori, A.Tiribocchi, G.Amati, M.Bernaschi, M. Lauricella, S.Succi, Lbsoft: A parallel open-source software for simulation of colloidal systems, Computer Physics Communications 256 (2020) 107455.

[5] S. Succi, The lattice Boltzmann equation: for complex states of flowing matter, Oxford University Press, 2018.

[6] T. Krüger, H. Kusumaatmaja, A. Kuzmin, O. Shardt, G. Silva, E. M. Viggen, The lattice boltzmann method, Springer International Publishing 10 (2017) 4–15.

[7] R. Benzi, S. Succi, M. Vergassola, The lattice boltzmann equation: theory and applications, Physics Reports 222 (1992) 145–197.

[8] M. Bernaschi, S. Melchionna, S. Succi, Mesoscopic simulations at the physics-chemistry-biology interface, Reviews of Modern Physics 91 (2019) 025004.

[9] M. Bauer, S. Eibl, C. Godenschwager, N. Kohl, M. Kuron, C. Rettinger, F. Schornbaum, C. Schwarzmeier, D. Thönnes, H. Köstler, et al., walberla: A block-structured high-performance framework for multiphysics simulations, Computers & Mathematics with Applications 81 (2021) 478–501.

[10] M. Holzer, M. Bauer, H. Köstler, U. Rüde, Highly efficient lattice boltzmann multiphase simulations of immiscible fluids at high-density ratios on cpus and gpus through code generation, The International Journal of High Performance Computing Applications (2021) 10943420211016525.

[11] J. Latt, O. Malaspinas, D. Kontaxakis, A. Parmigiani, D. Lagrava, F. Brogi, M. B. Belgacem, Y. Thorimbert, S. Leclaire, S. Li, et al., Palabos: parallel lattice boltzmann solver, Computers & Mathematics with Applications 81 (2021) 334–350.

[12] J.-C. Desplat, I. Pagonabarraga, P. Bladon, Ludwig: A parallel lattice-boltzmann code for complex fluids, Computer Physics Communications 134 (2001) 273–290.

[13] M. Bernaschi, S. Melchionna, S. Succi, M. Fyta, E. Kaxiras, J. K. Sircar, Muphy: A parallel multi physics/scale code for high performance bio-fluidic simulations, Computer Physics Communications 180 (2009) 1495–1502.

[14] D. M. Holman, R. M. Brionnaud, Z. Abiza, Solution to industry benchmark problems with the lattice-boltzmann code xflow, in: Proceeding in the European Congress on Computational Methods in Applied Sciences and Engineering (ECCOMAS).

[15] A. J. Ladd, Lattice-boltzmann methods for suspensions of solid particles, Molecular Physics 113 (2015) 2531–2537.

[16] A. Ladd, R. Verberg, Lattice-boltzmann simulations of particle-fluid suspensions, Journal of statistical physics 104 (2001) 1191–1251.

[17] C. K. Aidun, Y. Lu, E.-J. Ding, Direct analysis of particulate suspensions with inertia using the discrete boltzmann equation, Journal of Fluid Mechanics 373 (1998) 287–311.

[18] A. J. Ladd, Numerical simulations of particulate suspensions via a discretized boltzmann equation. part 1. theoretical foundation, Journal of fluid mechanics 271 (1994) 285–309.

[19] S. U. Pickering, Cxcvi.—emulsions, Journal of the Chemical Society, Transactions 91 (1907) 2001–2021.

[20] Q. Xie, G. B. Davies, J. Harting, Direct assembly of magnetic janus particles at a droplet interface, ACS nano 11 (2017) 11232–11239.

[21] H. Liu, Q. Kang, C. R. Leonardi, S. Schmieschek, A. Narváez, B. D. Jones, J. R. Williams, A. J. Valocchi, J. Harting, Multiphase lattice boltzmann simulations for porous media applications, Computational Geosciences 20 (2016) 777–805.

[22] S. Frijters, F. Günther, J. Harting, Effects of nanoparticles and surfactant on droplets in shear flow, Soft Matter 8 (2012) 6542–6556.

[23] F. Jansen, J. Harting, From bijels to pickering emulsions: A lattice boltzmann study, Physical Review E 83 (2011) 046707.

[24] Z. Sun, X. Yan, Y. Xiao, L. Hu, M. Eggersdorfer, D. Chen, Z. Yang, D. A. Weitz, Pickering emulsions stabilized by colloidal surfactants: Role of solid particles, Particuology (2021).

[25] S. Leclaire, A. Parmigiani, O. Malaspinas, B. Chopard, J. Latt, Generalized three-dimensional lattice boltzmann color-gradient method for immiscible two-phase pore-scale imbibition and drainage in porous media, Physical Review E 95 (2017) 033306.

[26] A. J. Ladd, Numerical simulations of particulate suspensions via a discretized boltzmann equation. part 2. numerical results, Journal of fluid mechanics 271 (1994) 311–339.

[27] M. Wittmann, T. Zeiser, G. Hager, G. Wellein, Comparison of different propagation steps for lattice boltzmann methods, Computers & Mathematics with Applications 65 (2013) 924–935.

[28] S. Williams, A. Waterman, D. Patterson, Roofline: an insightful visual performance model for multicore architectures, Communications of the ACM 52 (2009) 65–76.

[29] S.Succi, G.Amati, F. Bonaccorso, M. Lauricella, M.Bernaschi, A.Montessori, A.Tiribocchi, Towards exascale design of soft mesoscale materials, Journal of Computational Science 46 (2020) 101175.

[30] G. Falcucci, G. Amati, P. Fanelli, V. K. Krastev, G. Polverino, M. Porfiri, S. Succi, Extreme flow simulations reveal skeletal adaptations of deep-sea sponges, Nature 595 (2021) 537–541.

[31] E. M. Herzig, K. White, A. B. Schofield, W. C. Poon, P. S. Clegg, Bicontinuous emulsions stabilized solely by colloidal particles, Nature materials 6 (2007) 966–971.

[32] K. Stratford, R. Adhikari, I. Pagonabarraga, J.-C. Desplat, M. E. Cates, Colloidal jamming at interfaces: A route to fluid-bicontinuous gels, Science 309 (2005) 2198–2201.