

Parallel Multiphysics Simulations of Charged Particles in Microfluidic Flows

Dominik Bartuschat^{a,*}, Ulrich Rüde^a

^a*Lehrstuhl für Systemsimulation, Friedrich-Alexander Universität Erlangen-Nürnberg, Cauerstrasse 11, 91058 Erlangen, Germany*

Abstract

The article describes parallel multiphysics simulations of charged particles in microfluidic flows with the waLberla framework. To this end, three physical effects are coupled: rigid body dynamics, fluid flow modelled by a lattice Boltzmann algorithm, and electric potentials represented by a finite volume discretisation. For solving the finite volume discretisation for the electrostatic forces, a cell-centered multigrid algorithm is developed that conforms to the lattice Boltzmann meshes and the parallel communication structure of waLberla. The new functionality is validated with suitable benchmark scenarios. Additionally, the parallel scaling and the numerical efficiency of the algorithms are analysed on an advanced supercomputer.

© 2014 Published by Elsevier Ltd.

Keywords:

Parallel simulation; electrokinetic flow; fluid-particle interaction; cell-centered multigrid; MPI

1. Introduction

Computer simulations incorporating and coupling multiple physical effects rapidly gain importance in science and engineering. They can be used to predict and optimise the behaviour of processes or devices in engineering applications. However, the high computational complexity of such multiphysics simulations often requires the use of parallel supercomputers when realistic scenarios are studied.

In this paper, we develop a simulation model for charged objects in fluids that are additionally subject to electric fields. Systems of this type occur in a wide range of applications. For example, electrostatic filters can be designed, such that particles are removed from non-conductive liquids such as oil [1] or for filtering pollutants from exhaust gases. Medical scenarios may involve charged particles in aerosols for pulmonary drug delivery [2, 3] and in lab-on-a-chip systems where charged objects, such as cells or DNA, can be manipulated and separated by means of electric fields [4].

For simulating these scenarios, it is necessary to model the coupling between three system components: fluid flow, charged objects, and electric fields. This paper presents the design and implementation of efficient parallel algorithms for the simulation of such a coupled multiphysics constellation on advanced high performance computers. As example setup, we model a micro-channel flow of homogeneously charged rigid particles moving in a fluid, subject to an electrostatic field that is applied perpendicular to the flow direction. In this situation, the charged particles are transported with the fluid and are deflected by the electrostatic forces. We simulate two cases: the agglomeration of equally charged particles on an oppositely charged channel wall, and the separation of oppositely charged particles in a bifurcating micro-channel, as depicted in Fig. 1. Double-layer effects will not be considered in this paper, i. e., we treat the Hückel limit of electrophoresis.

The simulation scenario and the physical models are realised within WALBERLA [5, 6] (see Fig. 2). WALBERLA is a framework for massively parallel simulations of fluid flow applications that employ the lattice Boltzmann method (LBM) [7, 8]. The LBM can be used as an alternative to classical continuum mechanics approaches that has special advantages when simulating complex flow phenomena, such as moving objects. LBM is structurally an explicit

*Corresponding author

Email address: dominik.bartuschat@cs.fau.de (Dominik Bartuschat)

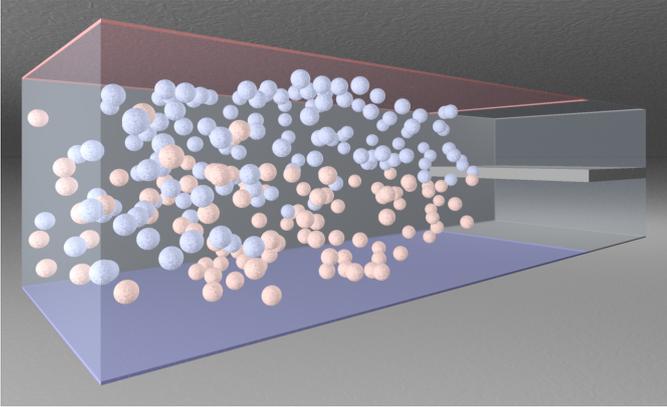


Figure 1. Separation of oppositely charged particles in a micro-channel fluid flow.

time stepping method and has local data dependencies. Thus, it can be parallelised with high efficiency even on large processor numbers. WALBERLA’s modular design permits the flexible extension with internal modules and the coupling with external libraries. For simulations involving the interaction of fluids and rigid bodies [9], WALBERLA can be coupled with the *pe* physics engine [10, 11]. The particles are treated as fully resolved rigid geometric objects. We use the *pe*’s parallel fast frictional dynamics algorithm [12] for simulating rigid body dynamics.

In the LBM, we use the two-relaxation-time (TRT) model of Ginzburg et al. [13]. The LBM and the rigid body simulation are coupled with a four-way interaction, as developed in Gtz et al. [9]. This uses the momentum exchange method of Yu et al. [14] similar to Ladd [15] and Aidun et al. [16]. Since the LBM per se can represent the forces between particles only correctly if the separation exceeds one lattice spacing [17], we employ a *lubrication correction* as introduced in Nguyen and Ladd [18].

In this paper, we describe how WALBERLA is augmented to additionally solve elliptic partial differential equations as they can be used to represent electric potentials, and how this is coupled to the LBM and the *pe* for particle transport. This requires the design and implementation of additional new functionality for boundary condition handling and efficiently solving large sparse linear systems of equations.

For discretising the electric potential equation, we choose a finite volume discretisation whose mesh spacing is conforming to the LBM grid. For the solution of the sparse linear systems, we develop a fast geometric multigrid (MG) algorithm. Fitting to WALBERLA’s parallelisation and data layout, a cell-centered MG method [19] is used. Galerkin coarsening is applied to provide a robust solver that is easy to use for different governing equations and boundary conditions (BCs).

In our work, achieving excellent computational performance is a primary goal. We gain additional speedups by using a representation of the linear system based on a *stencil* paradigm and using systematically that these sten-

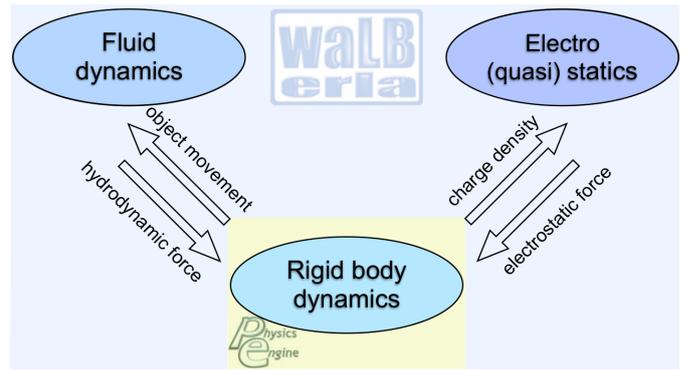


Figure 2. Coupled physical effects simulated with WALBERLA and the *pe*.

cils are almost uniformly identical throughout the simulation domain. Special stencils need to be used only near boundaries or material interfaces. For a discussion of various parallel implementation and performance optimisation approaches for MG methods, we refer to [20, 21].

The physical correctness of our simulations is validated for each of the coupled models. This includes fluid-particle interaction at low Reynolds numbers with lubrication correction and the electrostatic forces on the particles. Parameter studies show the impact of model parameters on physical accuracy.

Finally, we present performance results on the SuperMUC cluster of the Leibniz Supercomputing Centre¹ (LRZ) in Garching, Germany. Good scaling behaviour is shown for more than $68 \cdot 10^9$ lattice cells on up to 32,768 cores. In this case, more than seven million fully resolved charged objects that constitute for 9.4% solid volume fraction interact with the flow and exhibit electrostatic effects.

2. Numerical Modelling

2.1. Lattice Boltzmann Method

The LBM is based on kinetic theory for gases, describing molecular motion. However, in contrast to molecular dynamics, where individual molecules are considered, the LBM statistically describes ensembles of molecules. Particle² distribution functions (PDFs) represent the spatial and velocity distribution of molecules in *phase space* over time. The dynamics is described by the Boltzmann transport equation in terms of the temporal and spatial variation of PDFs, balanced by molecular collisions. The solution of the Boltzmann equation converges towards the Maxwell-Boltzmann distribution of molecules that are in local thermodynamic equilibrium. At any instant of time, macroscopic properties can be computed as moments of the PDFs [22, 23].

¹<http://www.lrz.de/services/compute/supermuc/>

²These virtual fluid particles in the LBM context are different from the simulated macroscopic particles

For the LBM, the phase space is discretised into a Cartesian lattice $\Omega_{\delta x} \subset \mathbb{R}^D$ of dimension D with spacing δx , and a finite set of Q discrete velocities $\vec{c}_q \in \mathbb{R}^D$, $q \in \{1, \dots, Q\}$. These are chosen such that within a time increment $\delta t = t_{n+1} - t_n$ with discrete time $T_{\delta t} = \{t_n : n = 0, 1, 2, \dots\} \subset \mathbb{R}_0^+$, molecules can move to adjacent lattice sites or stay. Associated with each \vec{c}_q is a PDF $f_q : \Omega_{\delta x} \times T_{\delta t} \mapsto \mathbb{R}$. Explicit discretisation in time and finite difference discretisation in space [24] result in the discrete lattice Boltzmann equation

$$f_q(\vec{x}_i + \vec{c}_q \delta t, t_n + \delta t) - f_q(\vec{x}_i, t_n) = \delta t \mathcal{C}_q + \delta t F_q. \quad (1)$$

This equation describes the streaming of the PDFs between neighbouring lattice sites and the subsequent collisions, represented by the collision operator \mathcal{C}_q . The external body-force term F_q will be disregarded in this section.

The LBM can be employed with different collision operators. The simplest collision operator is the BGK model [25] that linearly relaxes the PDFs towards equilibrium in velocity space, with a single dimensionless relaxation time τ (or collision frequency $\omega = \tau^{-1}$). However, the BGK has disadvantages regarding its stability and accuracy [26]. In particular, the wall locations for the standard implementation of no-slip BCs depend on τ : These BCs of first order enforce zero velocity at the wall by setting

$$f_{\bar{q}}(\vec{x}_f, t_n + \delta t) = \tilde{f}_{\bar{q}}(\vec{x}_f, t_n). \quad (2)$$

Consequently, PDFs of a fluid cell \vec{x}_f adjacent to a wall are reflected from the wall to direction \bar{q} opposite to their original direction q .

A more sophisticated operator with multiple relaxation times was presented in d’Humières [27]. It performs collisions in moment space, relaxing the different moments towards their equilibrium. This method is numerically stable and accurate [26], but computationally costly.

We use the stable, accurate, and computationally cheaper TRT collision operator [28, 13] for collisions in velocity space

$$\mathcal{C}_q = \lambda_e (f_q^e - f_q^{\text{eq},e}) + \lambda_o (f_q^o - f_q^{\text{eq},o}), \quad (3)$$

with two relaxation times, λ_e for even- and λ_o for odd-order non-conserved moments. The parameters λ_e and λ_o are eigenvalues of the TRT collision matrix \mathbf{S} that is related to the collision operator by $\vec{\mathcal{C}} = \mathbf{S}(\vec{f} - \vec{f}^{\text{eq}})$, with $\vec{f} := (f_0, \dots, f_{Q-1})^T \in \mathbb{R}^Q$ and \vec{f}^{eq} , $\vec{\mathcal{C}}$ defined analogously. $\lambda_e = -\tau^{-1}$ is related to the kinematic viscosity $\nu = (\tau - \frac{1}{2}) c_s^2 \delta t$, whereas λ_o can be chosen freely. To ensure stability, λ_o should be within the interval $] -2, 0[$. Ginzburg et al. [29] use the parameter $\Lambda = \left(\frac{1}{2} + \frac{1}{\lambda_e}\right) \left(\frac{1}{2} + \frac{1}{\lambda_o}\right)$. For bounce-back BCs (Eqn. (2)), walls are fixed midway between two lattice sites, if $\Lambda := \Lambda_{\text{mid}} = \frac{3}{16}$, independent of τ . For porous media, the optimal value is $\Lambda := \Lambda_{\text{por}} = \frac{1}{4}$.

TRT requires the decomposition $f_q = f_q^e + f_q^o$ into even and odd components

$$\begin{aligned} f_q^e &= \frac{1}{2}(f_q + f_{\bar{q}}) & \text{and} & & f_q^{\text{eq},e} &= \frac{1}{2}(f_q^{\text{eq}} + f_{\bar{q}}^{\text{eq}}) \\ f_q^o &= \frac{1}{2}(f_q - f_{\bar{q}}) & \text{and} & & f_q^{\text{eq},o} &= \frac{1}{2}(f_q^{\text{eq}} - f_{\bar{q}}^{\text{eq}}), \end{aligned} \quad (4)$$

with $\vec{c}_{\bar{q}} = -\vec{c}_q$. The local equilibrium distribution function for incompressible LBM, according to He and Luo [30], is then given for each lattice site by

$$\begin{aligned} f_q^{\text{eq},e} &= w_q \left(\rho_f - \frac{\rho_0}{2c_s^2} (\vec{u} \cdot \vec{u}) + \frac{\rho_0}{2c_s^4} (\vec{c}_q \cdot \vec{u})^2 \right) \\ f_q^{\text{eq},o} &= w_q \frac{\rho_0}{c_s^2} (\vec{c}_q \cdot \vec{u}), \end{aligned} \quad (5)$$

where ‘ \cdot ’ denotes the standard Euclidean scalar product. Both, macroscopic density $\rho_f = \rho_0 + \delta\rho$ with fluctuation $\delta\rho$, and velocity \vec{u} , can then be computed from moments of f

$$\begin{aligned} \rho_f(\vec{x}_i, t) &= \sum_q f_q(\vec{x}_i, t), \\ \vec{u}(\vec{x}_i, t) &= \frac{1}{\rho_0} \sum_q \vec{c}_q f_q(\vec{x}_i, t). \end{aligned} \quad (6)$$

We use the D3Q19 model of Qian et al. [31], where the thermodynamic speed of sound is $c_s = c/\sqrt{3}$, with $c = \delta x/\delta t$. Then, the weights w_q are: $w_1 = 1/3$, $w_{2,\dots,7} = 1/18$, and $w_{8,\dots,19} = 1/36$.

In each time step $t_n \in T_h$ the lattice Boltzmann method performs a *collide*- and a *stream step*

$$\begin{aligned} \tilde{f}_q(\vec{x}_i, t_n) &= f_q(\vec{x}_i, t_n) + \lambda_e [f_q^e(\vec{x}_i, t_n) - f_q^{\text{eq},e}(\vec{x}_i, t_n)] \\ &\quad + \lambda_o [f_q^o(\vec{x}_i, t_n) - f_q^{\text{eq},o}(\vec{x}_i, t_n)] \end{aligned} \quad (7)$$

$$f_q(\vec{x}_i + \vec{e}_q, t_n + \delta t) = \tilde{f}_q(\vec{x}_i, t_n), \quad (8)$$

where \tilde{f}_q denotes the post-collision state and $\vec{e}_q = \vec{c}_q \delta t$. In the collide step, the PDFs are relaxed towards equilibrium, whereas the stream step models mass transport.

Computationally, the collide step is compute-intensive while the stream step is memory-intensive. To increase efficiency, both steps are fused to a *stream-collide* step. Since PDFs are only streamed between adjacent lattice sites, LBM has a strictly local memory access pattern. Thus, it is parallelisable, and carefully designed implementations scale well on parallel architectures.

To ensure consistent units in the computations, WALBERLA converts physical entities to lattice units (LUs) at the beginning of the simulation. In the following, physical parameters in LUs are indicated with subscript L , e.g. δx_L . The parameters δx_L , δt_L , $\rho_{f,L}$ and the electric potential Φ_L are thus normalised to unity.

2.2. Fluid-Particle Interaction

Particles suspended in fluids get dragged along with a flow due to momentum transferred from the fluid. They in turn influence the flow by transferring momentum to the fluid. Methods based on the Navier-Stokes equation usually model the interaction by imposing the continuity of normal stresses and velocities at fluid-object interfaces.

Since the LBM is a mesoscopic method, momentum transferred from the fluid to objects can be directly computed from PDFs. The continuity of velocities can be imposed by no-slip BCs on the surface of the objects. This interaction is modelled in WALBERLA by the momentum exchange approach that was first presented in Ladd [15, 17], with particles treated as fluid-filled shells. Due to reduced stability, and because this method is limited to particles whose density is larger than the fluid density, Aidun et al. [16] suggest to treat the particles as solid objects.

Our method follows Yu et al. [14] and considers the particles as solid, rigid objects that are mapped to the lattice. Each cell whose center is overlapped by a particle is regarded as solid moving obstacle cell. The other cells are fluid cells, on which the LBM is applied with BCs modelling the moving objects. This method is efficient and easy to implement because the force acting on an object can be computed directly from PDFs in the surrounding cells. The solid particles come at the cost of fluid cells appearing and disappearing due to particle movement. Consequently, the PDFs of fluid cells must be reconstructed: we set the PDFs of the wake-cell to the equilibrium distribution based on the object surface velocity and the lattice fluid density $\rho_{f,L} = 1$.

Particles transfer momentum to the fluid by means of accelerating no-slip BCs introduced in Ladd [15]. The fluid velocity at a given fluid cell i at \vec{x}_i adjacent to the particle is matched to the local velocity \vec{u}_s at a particle surface cell s at \vec{x}_s by setting

$$f_{\vec{q}}(\vec{x}_i, t_n) = \tilde{f}_{\vec{q}}(\vec{x}_i, t_n) - 2\frac{\omega_{\vec{q}}}{c_s^2} \rho_f \vec{c}_{\vec{q}} \cdot \vec{u}_s. \quad (9)$$

Here, the PDF reflected from the particle surface opposite to the direction q pointing from \vec{x}_i to \vec{x}_s is modified accordingly. \vec{u}_s describes the object movement, including rotation and translation. Representing object boundaries by a *staircase approximation* is only first-order accurate, but it is computationally efficient and is sufficient for our applications as will be shown in Sec. 6. Higher-order BCs are either computationally more demanding, such as multi-reflection in Ginzburg et al. [32], or yield τ -dependent results, such as Bouzidi BCs [33]. More details on this viscosity dependence can be found in [34, 35].

The momentum transferred from fluid to the object can be computed from the difference in momentum density before and after the reflection of PDFs at the solid object. Summing up the momenta transferred from fluid cells i to neighbouring particle surface cells s , the hydrodynamic force on the particle can be computed following [36] as

$$\vec{F}_h = \sum_s \sum_{q \in D_s} \vec{c}_q \left[2\tilde{f}_{\vec{q}}(\vec{x}_i, t_n) - 2\frac{\omega_{\vec{q}}}{c_s^2} \rho_f \vec{c}_{\vec{q}} \cdot \vec{u}_s \right] \frac{\delta x^3}{\delta t}. \quad (10)$$

Here, $\vec{x}_i = \vec{x}_s + \vec{c}_{\vec{q}}$ and the set D_s of direction indices q , in which a given s is accessed from adjacent i .

The hydrodynamic force is summed up at the boundary sites of each solid object. These partial forces are then aggregated to the total force acting on the center of mass and the corresponding torque. In the subsequent position update for the objects, both the translation and rotation are computed. They in turn affect the fluid via \vec{u}_s and that way other particles. This interaction of fluid and particles alone corresponds to a two-way coupling. Since additionally the interaction of individual particles and of particles and walls is fully modelled in the *pe* algorithms, we achieve a full four-way coupling.

2.3. Lubrication Correction

Lubrication forces, defined here following Bossis and Brady [37] and Cox [38], are strong forces occurring when particles in a fluid move relative to each other in close proximity. In case of approaching particles, fluid in the gap between the particles must be displaced, and a high pressure arises. The corresponding force is increasing with larger relative velocity and decreasing gap width. Lubrication forces counteract the particle movement, i. e., they are attractive when particles move apart.

An insufficient handling of these lubrication forces can result in wrong particle interactions, such as particles unphysically sticking together in the simulation. The LBM with moving objects can accurately capture lubrication forces for particle separation distances larger than one lattice site for creeping flow [17]. For smaller gap sizes, the increasing force can not be recovered directly. Instead, the force stays approximately constant.

As shown in Ladd [17] for spheres of equal size, the lubrication force depends on $s = \frac{h_g}{R}$, the gap width normalised by the sphere radius. The normal lubrication force acting between the particle centers is proportional to s^{-1} . Tangential lubrication forces are proportional to $\ln(s^{-1})$. Due to the logarithmic dependency, their effect is much weaker for small gaps so that they are neglected in the present paper. We refer to results of Janoschek et al. [39], which show that for solid volume fractions below 30%, tangential lubrication correction can be neglected to accurately simulate shear-induced diffusion and suspension viscosity with spherical particles.

The lubrication correction for the LBM was introduced in Ladd [40] for the normal force between two equally sized spheres: By correcting the computed force on the spheres with the part of the lubrication force not captured by LBM, for distances below a threshold h_c in the order of one lattice spacing. It adds the difference between the analytical solution of the lubrication force [41, 38] in Stokes flow and the lubrication force at the threshold distance that is still captured by LBM, replacing $\frac{1}{s}$ by $(\frac{1}{s} - \frac{R}{h_c})$.

An extension for two differently sized spheres was presented in Ladd and Verberg [42]

$$\vec{F}_{ab}^{Lub} = \begin{cases} 6\pi\eta \frac{R_a^2 \cdot R_b^2}{(R_a + R_b)^2} \left(\frac{1}{h_g} - \frac{1}{h_c} \right) u_n^{rel} \hat{r}_{ab} & , \text{ if } h_g \leq h_c \\ 0 & , \text{ else} \end{cases} \quad (11)$$

with the relative particle velocity in direction of the contact normal $u_n^{rel} = \vec{u}_{ab} \cdot \hat{r}_{ab}$ given by the inner product of the velocity difference $\vec{u}_{ab} = \vec{u}_b - \vec{u}_a$ of both spheres and the normalised distance vector $\hat{r}_{ab} = \frac{\vec{r}_{ab}}{|\vec{r}_{ab}|}$. Here, $\vec{r}_{ab} = \vec{x}_b - \vec{x}_a$ is the vector connecting the centers of both spheres with radius R_a and R_b . For gaps h_g larger than the cut-off distance h_c , the lubrication correction force is not applied. Sphere-wall lubrication correction can be computed with this formula by letting e. g. $R_b \rightarrow \infty$ and for stationary walls setting $\vec{u}_b = 0$.

We perform normal lubrication force correction by adding the force in Eqn. (11) due to surrounding particles to the computed hydrodynamic force of each particle. Following Nguyen and Ladd [18], the lubrication correction term is added for each pair of particles. Nguyen and Ladd refer to Brady and Bossis [43], who demonstrate that this yields sufficiently accurate results for their method.

To ensure stability of the simulation, the following restrictions were added: The lubrication force is limited to a maximum value, preventing the system from getting too stiff, which would significantly restrict the time-step size. For this purpose, the gap width in the lubrication force computation is restricted to a minimum threshold value ($0.01 \cdot \delta x$ in all following simulations). This minimum gap is by a factor of R smaller than the distance where qualitatively important physics typically still occurs – at gaps down to $0.01R$ [18]. Another limiter was introduced for particles with a high separation velocity that can occur after the collision of particles. In case $u_n^{rel} > u_{max}^{sep}$, the normal lubrication force is limited by rescaling its magnitude F^{Lub} as $\vec{F}^{Lub} = \frac{\vec{F}^{Lub}}{F^{Lub}} (1 + \log_{10}(F^{Lub}))$. This limiter prevents problems that occur with the parallel fast frictional dynamics algorithm [12] as it is used in the *pe* and which resolves collisions by computing a post-collision velocity to push the particles apart. The lubrication force computed in the next time-step acts opposite to that movement, due to the tiny gap and high separation velocity. This leads to particles being pushed apart strongly, possibly resulting in stability problems.

All methods above use only the dominating singular terms of lubrication force components. Recently, an even more elaborate method for lubrication correction of a-spherical particles was presented in [39], taking into account higher-order singular terms. However, this is not incorporated in the present algorithm.

2.4. Electric Potential

The interaction of charged particles and of these particles with charged walls can be modelled by the electric potential and resulting Coulomb forces acting on the objects. The spatially varying electric potential $\Phi(\vec{x})$ produced by the charge density $\rho(\vec{x})$ is described by Poisson’s equation

$$-\Delta \Phi(\vec{x}) = \frac{\rho(\vec{x})}{\varepsilon} \quad (12)$$

for spatially constant permittivity ε . The charge density is zero in the fluid and nonzero at the particle locations.

Finite volume discretisation. In order to solve Eqn. (12), we apply a conservative finite volume [44, 45] scheme on the regular LBM lattice that subdivides the spatial domain into cells which act as control volumes. This includes volume integration over each cell and applying the divergence theorem. The resulting sum over the fluxes $\nabla \Phi = -\vec{E}$ across the cell surfaces then balances the volume integral over the right-hand side (RHS) of Eqn. (12), conserving charge. When the fluxes are approximated by central differences of Φ_i from the neighbouring and current cell, one obtains the seven-point stencil for each unknown Φ_i

$$\Delta_{\delta x} = \frac{-1}{\delta x^2} \left[\begin{pmatrix} 0 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \begin{pmatrix} 0 & -1 & 0 \\ -1 & 6 & -1 \\ 0 & -1 & 0 \end{pmatrix} \begin{pmatrix} 0 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 0 \end{pmatrix} \right] \quad (13)$$

for uniform lattice spacing δx . This scheme provides up to $\mathcal{O}(\delta x^2)$ accuracy. The unknowns are associated with the cell centers and represent the mean value over the cell. Consequences of this approach to the solver for the resulting linear system of equations are discussed in Sec. 3.

Boundary conditions. The boundary $\Gamma_{\delta x}$ of the discretised domain is formed by the outer faces of *near-boundary* cells. There, the following BCs can be used for the electric potential:

- Neumann BCs $\frac{\partial \Phi}{\partial \vec{n}} \Big|_{\Gamma_{\delta x}} = g_n$ define the gradient—or electric field—at the boundary in normal direction. Since this is the flux across a cell surface, these BCs are naturally treated in finite volume schemes.
- Dirichlet BCs $\Phi \Big|_{\Gamma_{\delta x}} = g_d$ impose an electric potential at the boundary. Since unknowns are defined at cell centers, boundary values are extrapolated to *ghost cells*, i. e., cells outside the physical domain.
- Periodic BCs cyclically extend the domain. They are realised algorithmically by accessing unknowns in cells at opposite sides of the domain.

Neumann and Dirichlet BCs for the cell-centered discretisation are incorporated in the stencil and the RHS at near-boundary cells. This is demonstrated for one dimension and spacing $\delta x = 1$. Consider an arbitrary 3-point stencil $[\alpha \ \beta \ \gamma]$ applied at a cell i : $\alpha \Phi_{i-1} + \beta \Phi_i + \gamma \Phi_{i+1} = f_i$.

For Dirichlet BCs, the boundary value $\Phi_i \Big|_{i=\frac{1}{2}} = g_d$ at cell $i = 1$ is linearly extrapolated to the ghost cell, yielding $\Phi_0 = 2g_d - \Phi_1$. Inserting this into the above discretisation, one obtains the stencil $[0 \ (\beta - \alpha) \ \gamma]$ and the RHS $f_1 - 2\alpha g_d$.

For Neumann BCs, the value $\frac{\partial \Phi_i}{\partial \vec{n}} \Big|_{i=\frac{1}{2}} = g_n$ is approximated by central differences at $i = \frac{1}{2}$ as $\Phi_0 - \Phi_1 = g_n$. Substituting this in the finite volume scheme, one obtains the stencil $[0 \ (\beta + \alpha) \ \gamma]$ and the RHS $f_1 - \alpha g_n$.

This treatment is preferred over the use of *ghost values* because it eliminates additional degrees of freedom.

Moreover, ghost values in the BC treatment of iterative solvers depend on values of a previous iteration, whereas our method implicitly uses the new values for the BCs.

Coulomb force. The electrostatic force acting on a particle is computed from the portions acting on each cell b of a rigid body

$$\vec{F}_e = -\delta x^3 \sum_b \nabla \Phi(\vec{x}_b) \rho(\vec{x}_b). \quad (14)$$

The gradient of the electric potential is computed by means of finite differences that provide $\mathcal{O}(\delta x^2)$ accuracy. Where possible, an isotropy-preserving D3Q19 stencil is used (cf. [46]) instead of a D3Q7 stencil. With the LBM D3Q19 stencil, the gradient can be computed using w_q -weighted differences of neighbouring values in 18 directions \vec{e}_q as

$$\nabla \Phi(\vec{x}_b) \approx \frac{1}{w_1} \sum_{q=2}^{19} w_q \Phi(\vec{x}_b + \vec{e}_q) \cdot \frac{\vec{e}_q}{\delta x^2} \quad (15)$$

Subsampling. When setting the RHS or computing the electrostatic force on the particles, the charge density of a particle at a cell is required. The naïve approach is to divide the particle charge by the particle volume and to assign the resulting value to each cell whose center is overlapped by the particle. However, due to the volume mapping, fluctuations occur and the actual charge of a particle may not be accurately represented. This leads to errors in the computations of electric potential and electrostatic force. Using a correction factor to adapt the charge density to the staircase-approximated volume is inefficient in parallel computations, as it requires additional communication. As a local alternative, we introduce a subsampling technique that computes the volume overlap ratio of the particle for each cell. It equidistantly subdivides a cell and counts the overlapped subvolumes. Usually used subsampling factors and their effects are provided in Sec. 6.3.

3. Cell-Centered Multigrid Method

Multigrid methods are highly efficient iterative algorithms for solving large linear systems of equations. They are based on a hierarchy of grids with different resolution levels. On each level, high-frequency components of the error w.r.t. the resolution are efficiently eliminated by a smoothing method. The smoothed error on a coarser grid can then be used to correct the solution on a finer grid. Between the grid levels, information is transferred by means of restriction (fine to coarse) and prolongation (vice versa). Coarsening is applied recursively, with a decreasing number of unknowns on coarse grids, which significantly reduces the computational effort. For the considered class of problems, only a fixed number of MG iterations is required to achieve a given accuracy. This leads to a time-to-solution linear in the number of unknowns. For a general review of MG methods we refer to [47, 48].

The geometric MG solver deployed in this paper has been developed for cell-centered discretisations because this permits a mesh setup conforming to the LBM. Consequently, we can seamlessly use the data structures and parallel communication routines of WALBERLA (see Sec. 5.1). Additionally, cell-centered discretisations lead to simpler coarsening schemes. They allow coarsening as long as the total number of cells per dimension on a level is even and the number of cells at the coarsest level for each block is a multiple of two. Node-based discretisations would require an odd number of cells and the introduction of additional nodes per block for coarsening. A review of cell-centered MG can be found in [19]. The coarsening scheme is chosen such that complex boundary conditions can be handled and a seven-point stencil results on all grid levels for the Poisson problem. The resulting solver is extensible for discontinuous dielectricity coefficients. This can be accomplished by Galerkin coarsening [49], with transfer operators based on averaging restriction and a nearest neighbour prolongation.

Since we violate the condition $m_p + m_r > M$ for the polynomial orders of prolongation m_p and restriction m_r , where M is the order of the differential operator, we must expect that convergence rates independent of the problem size cannot be achieved [50]. However, Braess [51] describes that the method can be efficient for the Poisson equation if an additional factor of about 2 is applied to *magnify* the coarse grid correction, i. e., we use an over-relaxation in the correction step. A more detailed analysis is given in [52]. We use this variant of cell-centered MG because it preserves the seven-point stencil on the coarse grids for Galerkin coarsening. This keeps the method simple and avoids parallel communication overheads. In the current implementation, V-cycles are applied. V(1,1)-cycles do not converge reliably in our test setup, but a higher number of pre- and post-smoothing steps leads to fast convergence. In Sec. 8 V(3,3)-cycles are used, and more details will be presented. The parallel cell-centered MG method is described in more detail in [53].

4. Algorithm for Charged Particles in Fluids

In this section, we present an overview of the coupled multiphysics simulation algorithm³. The coupled physical effects are illustrated in Fig. 2. Alg. 1 computes the motion of the charged particles in the fluid, including their interaction with other particles and charged planes. Different example setups are described in Sec. 7.

In each time step, the Poisson problem for the electrostatic potential is solved. For this, the RHS is set depending on the particle charge densities and adapted to the BCs as described in Sec. 2.4. Then, the multigrid cycles are performed until the termination criterion is met, i. e., the residual L_2 norm is sufficiently small. This includes

³An earlier version of this method has been presented in [53]

Algorithm 1: Charged Particles Algorithm.

```

foreach time step do
  // solve Poisson problem with particle charges:
  set RHS of Poisson's equation
  while residual  $\geq$  tol do
    | perform MG V-cycle to solve the Poisson
    | equation
  // solve lattice Boltzmann equation considering
  // particle velocities:
  begin
    | perform stream step
    | compute macroscopic variables
    | perform collide step
  // couple potential solver and LBM to pe:
  begin
    | apply hydrodynamic force to particles
    | perform lubrication correction
    | apply electrostatic force to particles
    | pe moves particles depending on forces
  
```

MPI communication of unknowns to *ghost layers* of neighbouring processes for each level.

The LBM is performed as a fused stream-collide step (see Sec. 2.1). The PDFs are communicated via MPI, and BCs including moving objects are treated. Then the stream step and collide step are executed as described in Sec. 2.1. Here the macroscopic velocity and density are required to compute the equilibrium distribution.

The hydrodynamic forces on the particles and electrostatic forces are computed as described in Sec. 2.2 and Sec. 2.4, respectively. For particles in close proximity, the *lubrication correction* of Sec. 2.3 is applied. All external forces are summed for all particles, so the time step for integrating the particle trajectories can be performed by invoking the *pe* rigid body dynamics engine. The *pe* employs the parallel fast frictional dynamics algorithm [12] that computes the new particle positions by translational and rotational motion and also resolves rigid body collisions, taking into account friction.

5. Implementation in WaLBerla for Multiphysics Simulations

WALBERLA (*widely applicable Lattice-Boltzmann from Erlangen*) is a software framework for massively parallel fluid simulations that implements the methods and algorithms described above. Due to its modular design, it can be flexibly extended [6], and is particularly suitable for coupled multiphysics simulations. The coupling strategy is based on accessing mutually dependent data structures (e. g. setting RHS of Eqn. (12) dependent on particles), and on a common BC handling concept (see Sec. 5.2).

In this section, we present an overview of the software concepts underlying WALBERLA and extend them to further support the development of multiphysics scenarios.

To this end, a novel approach to handle boundary conditions and a new solver module are integrated into WALBERLA.

5.1. WaLBerla Concepts

The software structure of WALBERLA comprises a *core* for sequence control and so-called *applications* that assemble functionality (i. e. callable objects) from existing *modules*. The core initialises data structures, performs the time stepping, and finalises the simulation. The central construct for time-dependent simulations is the *timeloop* that iterates through the time steps and executes the callable objects as they are specified by the application.

The WALBERLA *functionality management* [6] is based on dynamic application switches that select functionality for a given hardware at runtime. The most fundamental callable objects are called *kernels*. For different architectures, specific optimised kernels are provided.

WALBERLA provides a communication module using MPI for distributed memory parallelisation. Communication is specialised for simulations on uniform grids and supports various communication patterns. Moreover, WALBERLA conquers complexity by partitioning: it decomposes the domain spatially into *blocks* and functionality into *sweeps*. For more details, see [53].

Blocks. We use uniform rectangular grids for discretisation with cubic cells. Equally sized blocks of cells are distributed to the compute nodes. For parallelisation and the BC handling, a layer of ghost cells is introduced for each block. Block data include also logistic information, such as e. g. the location of the block within the domain or the MPI rank of the block. The blocks allow for a *heterogeneous* parallel execution, i. e., the code runs on clusters that consist of different types of nodes, by invoking architecture-specific kernels for each node.

Sweeps. The time stepping within a simulation or the repeated grid traversal of an iterative solver are organised via sweeps. Sweeps can be performed concurrently in parallel between different blocks. They can be concatenated and are then executed in the specified order. Nested sweeps support iterative solvers (see Sec. 5.3). In a MG setting smoothing, restriction, and prolongation are typical sweeps. Each sweep uses kernels and can have three phases: a preparatory subsweep, the actual subsweep executing the numerical algorithm, and possibly a post-processing subsweep. In particular, communication between neighbouring blocks via MPI or a time step of the particle movement via the *pe* are implemented within the pre- or post-processing subsweeps.

5.2. Handling Boundary Conditions

The handling of BCs is motivated by the specific LBM requirements: PDFs of ghost cells must be set such that the BCs will be fulfilled when streaming the PDFs into the

domain (cf. [54]). Thus, no special treatment is required when the PDFs in the ghost layer are accessed in a stream step. This concept is extended for governing equations other than the lattice Boltzmann equation.

For handling multiple physical field data in a multiphysics scenario, WALBERLA must deal with the correct BCs for each governing equation. Moreover, it must support individual boundaries for the different governing equations: In our simulation, a particle imposes a BC for Eqn. (1), but not for Eqn. (12).

Our new technique for handling BCs aims to satisfy the following criteria:

- Generality – applicability to different numerical methods.
- Modularity – different functionality can be provided in independent modules.
- Flexibility – different techniques for handling BCs can be provided.
- Configurability – all BCs can be specified in input files.

Flags are used to represent the state of a lattice site and to indicate for each boundary cell which BC treatment has to be performed. Different from statically determined BCs, this permits flexibility, e. g., the handling of moving boundaries. Cells that are neighbouring to a boundary cell are indicated by a *nearBC* flag and those in the interior by a *nonBC* flag. Each type of boundary condition has an individual *BC* flag.

Individual sets of *nonBC*, *nearBC* and *BC* flags are specified for each governing equation. Thus, the shape of the boundary and the BCs can be specified individually for each field. The actual boundary handling functionality is implemented in corresponding BC classes whose functions are executed when the associated *BC* flag is found.

All BCs are handled such that they are fulfilled when the corresponding cell is accessed in the subsequent sweep. In case of periodic BCs, MPI processes are configured in a periodic arrangement in the WALBERLA setup phase. Near boundary values are then copied to the periodic neighbour’s ghost layer by communication functions. All other BCs can be handled by either direct or direction-dependent BC treatment: Direct BC treatment directly sets the BC value at a boundary cell (as for Dirichlet BCs in node-based discretisations). Direction dependent BC treatment sets a BC value at a boundary cell depending on the value at a neighbouring cell, as for LBM no-slip BCs (Eqn. (2)), or the BCs in cell-centered discretisations from Sec. 2.4 whose handling is described in Sec. 5.3.

5.3. Linear Systems Solver Module

The large sparse linear systems of equations that arise in Sec. 2.4 from the discretisation of the potential equation can be solved in WALBERLA by means of the *lse_solver*

module that has been designed as an efficient and robust black-box solver on block-structured grids. For high performance, all operations are implemented as *compact stencil* operations. They are seamlessly integrated into WALBERLA to avoid the overhead of using an external library.

The application sets up the system matrix, right-hand side, and BC handling. When the solver sweep is added to the timeloop, the iteration is executed. An iterative solver requires a nested sweep that is executed until a specific convergence criterion is satisfied.

For the implicit boundary treatment and MG with Galerkin coarsening, spatially varying stencils are stored for each unknown. In order to reduce memory transfer, *quasi-constant* stencils were introduced: Constant stencils are used for equations with constant coefficients on the finest level. Only at *nearBC* cells, the stencils may vary for each cell and need to be loaded individually.

BC Handling for Scalar Potentials. The solver module uses and initiates the implicit BC handling from Sec. 2.4. The stencils are constructed in the initialisation phase to incorporate the BCs. For fixed geometries, when the boundaries do not move in time, this has to be performed only once. At the beginning of the solver sweep, the RHS is adapted to the BCs, independent of the actually used solver. For the MG, this is only necessary on the finest grid.

At a given cell, the stencil and RHS adaption can be applied independently for each direction, allowing different kinds of BCs in different directions, e. g., at a corner. Here we use a direction-dependent BC treatment: the stencil entry in direction of the boundary is moved to the central entry, and the BC value in that direction is brought to the RHS, both with a factor depending on the BC type.

Multigrid solver. When the multigrid solver is applied, it automatically allocates all coarse grid data structures and generates the coarse grid system by Galerkin coarsening (see Sec. 3). Pre- and post-smoothing is performed with a Red-Black Gauss-Seidel algorithm that enables straightforward parallelisation by rearranging the unknowns in a checker-board manner to decouple data dependencies. The problem on the coarsest grid is currently solved by a conjugate gradient (CG) solver. CG methods are more efficient than the Red-Black Gauss-Seidel, but still have non-optimal complexity. Nevertheless, the CG is an adequate choice since the problem size on the coarsest grid remains moderate even for high levels of parallelism.

5.4. Lubrication Correction Implementation

Based on the description in Sec. 2.3, a sweep has been implemented to perform the lubrication correction. It can be added to any moving obstacle algorithm in WALBERLA after the hydrodynamic forces have been computed. The lubrication correction term Eqn. (11) is added to the hydrodynamic force acting on each particle: The algorithm iterates over all non-fixed particles residing on the current

process. In order to obtain the correction term depending on neighbouring particles or walls, an inner loop iterates over all other objects on the current process to detect pairs. For positive gap sizes that are smaller than the threshold h_c , the correction term is computed, incorporating the limiters described in Sec. 2.3.

This algorithm, as implemented, has a complexity of $\mathcal{O}(N^2)$ for N particles. However, since N is usually many orders of magnitude smaller than the number of fluid cells, this has negligible impact on performance.

6. Physical Validation

This section presents quantitative validation results for the different components of the overall multiphysics algorithm. To validate the fluid-structure interaction at low Reynolds numbers, we first validate the drag force \vec{F}_d acting on fixed spheres arranged in a periodic array. We investigate the influence of volume mapping errors for different sphere radii. The corresponding small fluctuations can also be seen in the validation of lubrication forces. Finally, the electric potential simulation for charged particles is validated, followed by a validation of the resulting electrostatic forces acting on the particles.

6.1. Drag Force

We validate the drag force \vec{F}_d on spheres in an infinitely large domain. This approach prevents wall effects that arise for alternative drag force validations with a sphere in a channel of finite width with free-slip BCs, compared to the analytical solution in an infinitely large domain.

For equal-sized, fixed spheres arranged in a regular periodic array in Stokes flow, an analytical solution for the dimensionless drag K is derived in [55]. This extends earlier results in Hasimoto [56] for the drag on spheres in incompressible Stokes flow for higher solid volume fractions ϕ_s and different particle arrangements.

$K = \frac{F_d}{6\pi\rho_f\nu\bar{u}R}$ is the drag force on a sphere in a periodic array immersed in an incompressible viscous fluid with average flow velocity \bar{u} , normalized w.r.t the drag force on an isolated sphere. In [55], K is computed for different values of $\chi = \sqrt[3]{\phi_s/\phi_{s,max}}$, where $\phi_{s,max}$ is the maximal solid volume fraction for a given packing. These results are shown to agree with values from literature within 0.5%. Moreover, a representation of the dimensionless drag is given as a power series in χ that allows to compute K for a wider range of ϕ_s than given in the paper. The relative error of this formulation w.r.t. the computed values are found to be below 1% for $\chi < 0.95$.

For validating our algorithm, we consider the simple cubic packing, where a fixed sphere is located at the center of a cubic cell with edge length L . In that case $\phi_{s,max} = \frac{\pi}{6}$ and $\phi_s = \frac{4\pi R^3}{3L^3}$ for spheres of radius R .

The fluid around the sphere is simulated with the incompressible TRT model described in Sec. 2.1. In all directions periodic BCs are applied. The fluid is driven by a uniform acceleration g_z in z -direction by means of a body-force term for Eqn. (1) that is added after the stream-collide step. We use the forcing term F_q according to Luo [57] as

$$F_q = w_q \left[\frac{(\vec{c}_q - \vec{u})}{c_s^2} + \frac{(\vec{c}_q \cdot \vec{u})}{c_s^4} \vec{c}_q \right] \cdot \vec{f}_{ext}, \quad (16)$$

where $\vec{f}_{ext} = \rho_f g_z$ is the external force-density. Since F_q affects the momentum density, the resulting macroscopic fluid velocity is given by [13]

$$\vec{u} = \frac{1}{\rho_0} \left(\sum_q f_q \vec{c}_q + \frac{\delta t}{2} \vec{f}_{ext} \right). \quad (17)$$

For the drag force validation, we compare the drag K^* from the simulation to the reference values of K . For that purpose, the hydrodynamic force \vec{F}^* on the sphere is computed in the simulation according to Eqn. (10). K takes into account the mean pressure-gradient [56]. Our simulations, however, do not include the corresponding pressure force \vec{F}_p from the fluid. Thus, the hydrodynamic force must be corrected by adding $\vec{F}_p = \rho_f g_z V_{sph}$, the force caused by a pressure gradient equivalent to $\rho_f g_z$.

To compute $K^* = \frac{F_L^* + F_{p,L}}{6\pi\rho_f\nu_L\bar{u}_L^*R_L}$ from the simulation results, the average fluid velocity \bar{u}^* in direction of the external forcing is computed over all fluid cells \vec{x}_f as $\bar{u}^* = \frac{1}{L^3} \sum_{\vec{x}_f} u_z(\vec{x})$. Each simulation is run until a steady state is reached, i. e., the relative change of \bar{u}^* between two time steps is close to machine accuracy. From the output values F_L^* and \bar{u}_L^* , the drag K^* is computed, together with the relative error $e_{r_K} = \frac{K^* - K}{K}$.

Additionally, the relative volume mapping error $e_{r_V} = \frac{V_{sim} - V_{sph}}{V_{sph}}$ of the staircase-approximated sphere volume V_{sim} w.r.t the theoretical volume $V_{sph} = \frac{4}{3}R_L^3\pi$ is computed. To indicate the influence of the volume mapping on the error, the relative error c_{r_K} of K^* corrected by the volume mapping error is computed. It is obtained by replacing F^* with $F_{cor}^* = \sqrt[3]{V_{sph}/V_{sim}} F^*$, as introduced in [58], based on the linear dependence of the Stokes drag force on the sphere radius.

In the simulations two different values of the TRT parameter Λ are used: Λ_{mid} (i. e. $\lambda_o = -8(2 - \omega)/(8 - \omega)$) and Λ_{por} (i. e. $\lambda_o = -(2 - \omega)$), cf. Sec. 2.1. The influence of τ is tested for values 1.7 and 3, corresponding to lattice viscosities $\nu_L = 0.4$ and $\nu_L = 0.83$, respectively.

Dependent on χ , the sphere radii are varied. To ensure small volume mapping errors, a cell with $L = 64$ is chosen. Consequently, the radii R_L are in the range 3.2 to 28.8 for the values of χ shown in Tab. 1. For simplicity, δx and δt are set to one. The acceleration in z -direction is chosen as $g_{z,L} = 5 \cdot 10^{-7}$, such that simulations are performed in the

Table 1. Drag force validation for single cubic cell of size 64^3 with $g_{L,z} = 5 \cdot 10^{-7}$ for different values of τ and Λ . Relative errors e_{r_K} and e_{r_K} marked with * are computed w.r.t. the power series values of K .

	χ	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.75	0.8	0.85	0.9
	R_L	3.2	6.4	9.6	12.8	16	19.2	22.4	24	25.6	27.2	28.8
	e_{r_V} [%]	-0.92	-0.92	-1.56	-0.46	0.58	-0.62	-0.07	-0.09	0.06	0.09	-0.04
$\tau = 1.7$	u_L^*	4.71e-3	1.95e-3	1.06e-3	626e-6	376e-6	227e-6	128e-6	93.6e-6	66.6e-6	46.2e-6	31.2e-6
	K^*	1.154	1.393	1.704	2.170	2.885	3.995	6.061	7.738	10.19	13.81	19.32
	e_{r_K} [%]	-0.95	0.34	0.23	0.84	1.51	0.53	0.95	1.04*	1.39	1.28	0.85
	c_{r_K} [%]	-0.64	0.65	0.75	0.99	1.33	0.72	0.97	1.06*	1.38	1.26	0.86
$\tau = 1.7$	u_L^*	4.77e-3	1.96e-3	1.07e-3	628e-6	377e-6	227e-6	128e-6	93.9e-6	66.9e-6	46.4e-6	31.4e-6
	K^*	1.140	1.386	1.698	2.163	2.877	3.984	6.043	7.714	10.15	13.76	19.24
	e_{r_K} [%]	-2.13	-0.18	-0.13	0.54	1.23	0.25	0.65	0.72*	1.04	0.90	0.43
	c_{r_K} [%]	-1.83	0.12	0.38	0.69	1.05	0.44	0.67	0.74*	1.02	0.88	0.44
$\tau = 3$	u_L^*	2.29e-3	941e-6	512e-6	302e-6	181e-6	109e-6	61.9e-6	45.3e-6	32.3e-6	22.5e-6	15.2e-6
	K^*	1.139	1.385	1.697	2.162	2.873	3.975	6.022	7.679	10.09	13.65	19.04
	e_{r_K} [%]	-2.16	-0.21	-0.19	0.45	1.10	0.04	0.31	0.27*	0.44	0.10	-0.63
	c_{r_K} [%]	-1.86	0.09	0.33	0.60	0.92	0.23	0.32	0.29*	0.43	0.08	-0.62

Stokes regime. The particle Reynolds numbers $Re_p = \frac{\bar{u} \cdot 2R}{\nu}$ have the maximal value of 0.075 at $\chi = 0.1$.

All parameters and the simulation results are shown in Tab. 1. The volume mapping error e_{r_V} is low for all sphere sizes, with maximal value -1.56% for $R_L = 9.6$. Moreover, e_{r_K} never exceeds 2.2% , which arises for Λ_{por} at the smallest solid volume fraction considered. For all other cases, e_{r_K} is approximately 1% , with largest values for $\tau = 1.7$ and Λ_{mid} . Using Λ_{por} there, results in a smaller error for a constant value of τ . Increasing τ to 3 reduces the error further. The increased viscosity results in higher average velocities. As expected, the volume mapping correction leads to relative errors c_{r_K} that are always smaller than e_{r_K} , with underestimated drag for the smallest and in one case the largest value of χ .

The results are in accordance with the drag evaluation in [35] for periodic regular (simple cubic) arrays of spheres, which examines the drag for TRT with $\tau = 3$ and Λ_{mid} . Different from this article, [35] focuses on flows through arrays of randomly arranged spheres for low and moderate Reynolds numbers.

6.2. Lubrication Correction

We show validation results for the lubrication correction in WALBERLA according to Eqn. (11). In order to validate the sphere-sphere and sphere-wall lubrication correction, two scenarios as in [59] are chosen: two equally sized spheres approaching each other with the same constant velocity u_{sph} and a sphere vertically approaching a wall with u_{sph} .

We compare the resulting corrected lubrication forces F_{Luc}^* in normal direction obtained from LBM simulations to the analytical solution that is presented for Stokes flow in [38]. In case of two spheres or a sphere and a fixed wall, the force is given by

$$F_n = \frac{3\pi}{2} \eta \frac{1}{h_g \lambda^2} u_n^{rel} + \mathcal{O}(\ln(h_g)), \quad (18)$$

with local surface curvature parameter $\lambda = \frac{1}{2R_a} + \frac{1}{2R_b}$. For comparison of the results, the forces are normalised as

$$f_{Norm} = \frac{F_{Luc}^*}{4R\eta u_{\text{sph}}} \approx \frac{3\pi}{4} \frac{1}{h_g \lambda}, \quad (19)$$

similar to [59]. The approximation for the analytical solution is derived from Eqn. (18) and is valid for small gaps. Eqn. (18) holds for sphere-sphere lubrication where $u_n^{rel} = 2u_{\text{sph}}$ and for sphere-wall lubrication where $u_n^{rel} = u_{\text{sph}}$.

All simulations are performed in the Stokes regime with fixed particle Reynolds number $Re_p = 0.096$. To see the influence of τ on the results, different velocities u_{sph} or sphere sizes are chosen to keep Re_p constant. TRT with only Λ_{por} is used, since there is no notable difference to Λ_{mid} in the results.

We uniformly use $h_c = \frac{2}{3}\delta x$ in our lubrication correction. In [18], this value is found to be suitable for normal lubrication force correction using BGK with $\tau = 1$. This value is close to $\tau = 0.93301$, where TRT and BGK yield the same results w.r.t. the wall locations [32]. Moreover, TRT exhibits no τ -dependency in this regard. Thus, in contrast to [18], who use BGK, considering the effective hydrodynamic radius is not necessary.

The spheres are placed in a rectangular channel with length L in x -direction and width W in the other dimensions. In all simulations, the spheres are centered at $W/2$ in y - and z -direction. In these directions, walls with free-slip BCs are modelled. W is chosen such that the wall effect on the force is negligible, i. e., about 100 times as large as the sphere diameters. The setups are sketched in Fig. 3.

For sphere-sphere lubrication, both spheres are initially placed with a given separation distance in x -direction and are moving towards each other with a fixed constant velocity u_{sph} each. Here, periodic BCs are applied in x -direction. (L is chosen such that the distance of the spheres

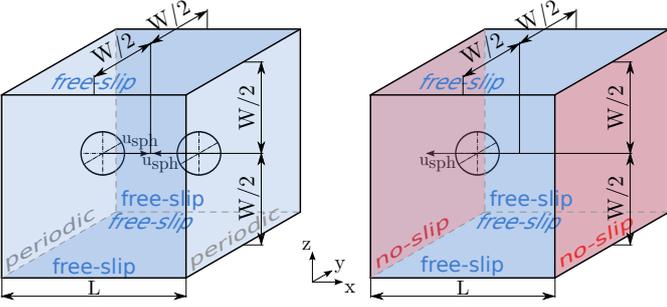


Figure 3. Setup for validation of sphere-sphere and sphere-wall lubrication correction.

in the periodically extended direction is at least 100 times the sphere diameter.)

For sphere-wall lubrication, the sphere is placed with a given distance from one of the no-slip walls in x -direction. The initial distances are chosen such that the hydrodynamic force is fully developed at positions where the lubrication force is measured. (L is chosen such that the distance to the wall opposite to the sphere movement is at least 100 times the sphere diameter.)

The sphere-sphere lubrication correction is evaluated for spheres with radius $R_L = 6$ that we frequently use in the simulations. Since the analytical solution for the lubrication force is valid only for small gap sizes compared to the radius of curvature (see error term in Eqn. (18)), the numerical solution is expected to coincide for small separation distances only. Thus, the simulations are also performed for larger spheres with radius $R_L = 48$. The domain size for $R_L = 6$ is set to $1216 \times 1216 \times 1216$ cells. For $R_L = 48$, it is increased to $1024 \times 2048 \times 2048$ cells. The values of τ and the corresponding lattice viscosities ν_L are shown in Tab.2, together with the velocity of each sphere. The following physical parameters

Table 2. Sphere-sphere lubrication correction parameters for spheres of $R_L = 6$ and $R_L = 48$ ($Re_p = 0.096$).

τ	ν_L	$u_{sph,L}$ ($R_L = 6$)	$u_{sph,L}$ ($R_L = 48$)
1.7	0.4	$1.6 \cdot 10^{-3}$	–
2.0	0.5	–	$0.25 \cdot 10^{-3}$
3.5	1.0	$4.0 \cdot 10^{-3}$	$0.5 \cdot 10^{-3}$
5.3	1.6	$6.4 \cdot 10^{-3}$	$0.8 \cdot 10^{-3}$
6.5	2.0	$8.0 \cdot 10^{-3}$	$1.0 \cdot 10^{-3}$

are chosen: $\rho_f = 10^3 \text{ kg/m}^3$, $\delta x = 10^{-3} \text{ m}$, and $\delta t = 0.3 \text{ s}$ —such that $\nu = 1.3 \cdot 10^{-6} \text{ m}^2/\text{s}$ for $\nu_L = 0.4$.

Sphere-wall lubrication correction is validated for spheres of different size moving with constant velocity $u_{sph,L} = 10^{-3}$. The domain size is set to $1216 \times 1216 \times 1216$ cells for $R_L = 6$, $1600 \times 1792 \times 1792$ for $R_L = 8$ and $R_L = 9$, and $1216 \times 2432 \times 2432$ for $R_L = 12$. The sphere radii are shown in Tab. 3, together with ν_L and the corresponding τ . Different from the sphere-sphere lubrication correction, we chose $\delta t = 0.125 \text{ s}$ —such that $\nu = 10^{-6} \text{ m}^2/\text{s}$ for $R_L = 6$.

All other physical parameters are kept identical.

Table 3. Sphere-wall lubrication correction parameters for different sphere radii ($Re_p = 0.096$).

R_L	6	8	9	12
ν_L	1/8	1/6	3/16	1/4
τ	0.875	1.0	1.0625	1.25

For each simulation, the computed force F_{Luc}^* on a sphere including lubrication correction is normalised according to Eqn. (19). The results for sphere-sphere lubrication at different gap sizes $h_{g,L}$ are depicted in Fig. 4 and Fig. 5 for $R_L = 6$ and $R_L = 48$, respectively. The results for sphere-wall lubrication are depicted in Fig. 6 at different normalised gap sizes $\frac{h_g}{2R}$. In each figure, the normalised approximation to the analytical solution in Eqn. (19) is plotted for reference.

Additionally, the normalised hydrodynamic force without lubrication correction is shown in Fig. 4 and in Fig. 6, indicated by ‘nC’. In Fig. 4, it stays approximately constant below one lattice spacing, as already reported in [17]. In Fig. 6, it also stays constant for gaps smaller than one lattice spacing, but then decreases for $h_{g,L} < 0.5$. The reason for the decrease is that a lattice cell whose center is overlapped by the particle is no longer considered a fluid cell. Consequently, in the cell at minimum distance between sphere and wall, the pressure is no longer represented correctly. This leads there to an attractive force and an overall decreasing force on the sphere. In contrast, the spheres of the sphere-sphere lubrication validation move symmetrically towards the center of the cell between the spheres, and thus this cell center is never overlapped.

The lubrication correction method reproduces the analytical solution for decreasing gap sizes in all simulations (see Fig. 4 to Fig. 6). For larger gap sizes, the forces from the simulations and the analytical solution deviate increasingly, especially for $R_L = 6$ (see Fig. 4). This difference is

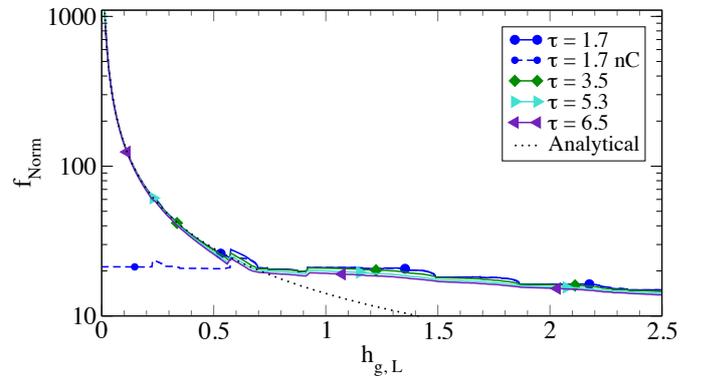


Figure 4. Validation of sphere-sphere lubrication correction for spheres of radius $R_L = 6$.

reduced in Fig. 5, where the simulation results agree better with the analytical solution also for larger gap sizes.

Overall, the forces in Fig. 5 coincide very well for large

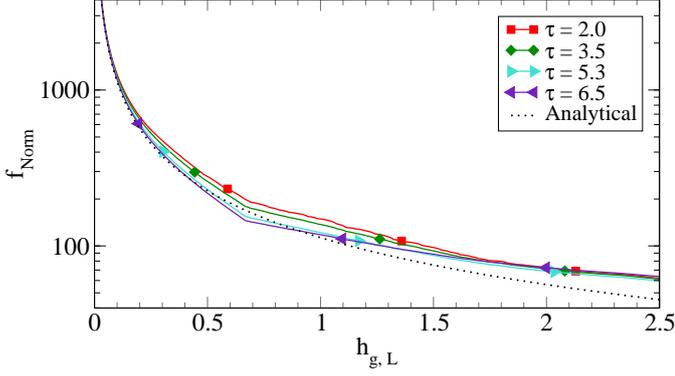


Figure 5. Validation of sphere-sphere lubrication correction for spheres of radius $R_L = 48$.

and small gap sizes. Only for $h_{g,L}$ smaller than approximately two the forces start diverging, depending on u_{sph} and τ . Nevertheless, all simulated forces lie closely around the analytical solution in that range of $h_{g,L}$.

For the case of smaller spheres, as evaluated in Fig. 4 and Fig. 6, the volume mapping error becomes visible. Naturally, this error decreases with increasing sphere size, as indicated by the ripple in Fig. 6. The influence of τ and u_{sph} on the force fluctuations due to volume mapping effects for a constant sphere size can be seen in Fig. 4, where the jumps decrease with increasing τ and u_{sph} . In both cases, this error increases for $h_{g,L} < 1$, but is insignificant for small gap sizes because the lubrication correction dominates.

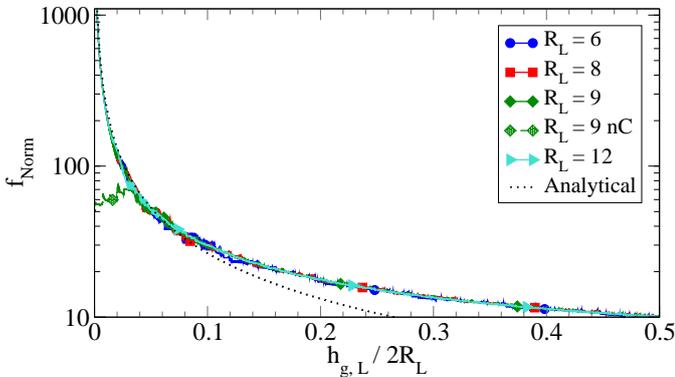


Figure 6. Validation of sphere-wall lubrication correction for different sphere sizes.

Overall, the forces in Fig. 4 differ only slightly. The hydrodynamic force marginally decreases with increasing τ and u_{sph} . Tests with higher Re_p have shown that for a fixed τ but increasing u_{sph} , both, the normalised hydrodynamic forces and the volume mapping errors, decrease slightly. Thus, these effects can be mainly attributed to u_{sph} . Also the forces in Fig. 6 differ only slightly. No systematic dependence on R_L and τ values could be determined. The reason is that u_{sph} is constant, reflecting the minor influence of τ on the force.

The lubrication force correction formula is often also applied for larger Reynolds numbers than the ones presented here. This is justified because, due to the small gaps, one can still approximate the flow by Stokes flow [39].

6.3. Electric Potential and Electrostatic Force

To validate the electric potential computation, we simulate a homogeneously charged sphere in a large domain. The numerical solution is then compared to the analytical solution for the potential distribution $\Phi(\vec{r})$, which in case of an infinitely large domain is

$$\Phi(\vec{r}) = \begin{cases} \frac{1}{4\pi\epsilon} \frac{Q}{|\vec{r}|} & \text{if } |\vec{r}| \geq R \\ \frac{1}{4\pi\epsilon} \frac{Q}{2R} \left(3 - \left(\frac{|\vec{r}|}{R} \right)^2 \right) & \text{if } |\vec{r}| < R, \end{cases} \quad (20)$$

with the distance from the sphere center \vec{r} , the total charge Q of the sphere and its radius R . Simulating a domain so large that the electric potential decays to zero at the boundary is not feasible. Thus, we set the boundary values to the analytical solution in our validation experiments.

For a sphere with radius $R_L = 6$ placed at the center of a domain of size 256^3 , both, analytical (Φ) and numerical (Φ^*) solutions, are depicted in Fig. 7 along a line through the domain center. The simulation parameters are $Q = 8000 e$ and $\epsilon = 78.5 \cdot \epsilon_0$, with vacuum permittivity ϵ_0 . Additionally, a subsampling factor of three (i. e. subdivision of each cell in three intervals for each dimension) is used. Clearly, both solutions are in good agreement.

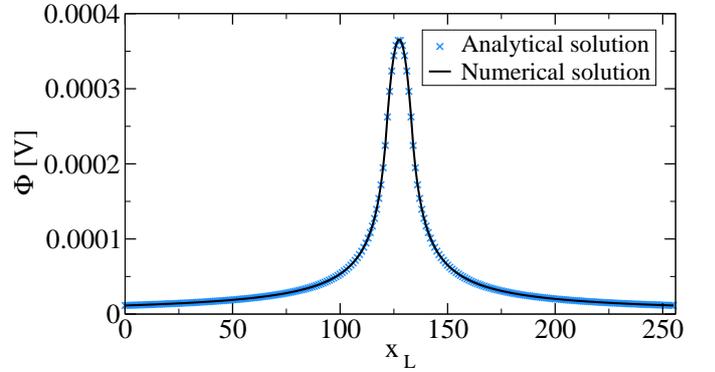


Figure 7. Analytical and numerical solution for electric potential from homogeneously charged sphere.

A quantitative assessment of the maximal errors in the electric potential computation is performed for different particle radii and subsampling factors s_f . First, the particle position with maximal volume mapping error e_{r_V} (see Sec. 6.1) is determined. To this end, the sphere is placed at the domain center and then repeatedly shifted by a small distance in all dimensions (up to $\pm 0.51 \delta x / s_f$). From the errors e_{r_V} , the maximal error $e_{r_V \text{max}}$ is obtained. Additionally, the average volume mapping error $\overline{|e_{r_V}|}$ is computed from the absolute values of e_{r_V} . For the position of $e_{r_V \text{max}}$, the relative error $e_{r_\Phi} = \frac{\Phi^* - \Phi}{\Phi}$ of the numerical solution Φ^* is computed in the whole domain. In

Tab. 4, the L_2 norm and the maximum norm of that error is shown, together with $e_{r_V \max}$ and $\overline{|e_{r_V}|}$ for different sphere sizes and subsampling factors.

The average error caused by the volume mapping ex-

Table 4. Electric potential validation for homogeneously charged sphere in domain of size 256^3 for different radii and subsampling factors. Comparison of relative volume mapping errors e_{r_V} and of relative electric potential errors e_{r_Φ} at the position of maximal volume mapping error.

		pot. subsampling (s_{f_Φ})			
		1	2	3	4
$R_L = 4$	$ e_{r_V} $	1.14%	0.284%	0.145%	0.081%
	$e_{r_V \max}$	7.06%	2.58%	-1.43%	0.72%
	$\ e_{r_\Phi}\ _{L_2}$	1.87%	0.622%	0.568%	0.192%
	$\ e_{r_\Phi}\ _\infty$	10.9%	3.37%	-2.06%	1.13%
$R_L = 6$	$ e_{r_V} $	0.638%	0.140%	0.065%	0.039%
	$e_{r_V \max}$	-2.74%	-1.43%	-0.33%	-0.25%
	$\ e_{r_\Phi}\ _{L_2}$	0.927%	0.568%	0.290%	0.273%
	$\ e_{r_\Phi}\ _\infty$	-4.48%	-2.19%	-0.80%	-0.61%
$R_L = 8$	$ e_{r_V} $	0.285%	0.081%	0.040%	0.021%
	$e_{r_V \max}$	2.58%	0.72%	-0.25%	-0.14%
	$\ e_{r_\Phi}\ _{L_2}$	0.622%	0.192%	0.273%	0.251%
	$\ e_{r_\Phi}\ _\infty$	3.37%	0.97%	-0.67%	-0.48%
$R_L = 9$	$ e_{r_V} $	0.258%	0.066%	0.028%	0.018%
	$e_{r_V \max}$	1.91%	-0.33%	0.32%	-0.11%
	$\ e_{r_\Phi}\ _{L_2}$	0.444%	0.290%	0.185%	0.244%
	$\ e_{r_\Phi}\ _\infty$	2.48%	-0.82%	0.44%	-0.39%
$R_L = 12$	$ e_{r_V} $	0.146%	0.040%	0.018%	0.009%
	$e_{r_V \max}$	-1.43%	-0.25%	-0.11%	-0.10%
	$\ e_{r_\Phi}\ _{L_2}$	0.567%	0.273%	0.244%	0.243%
	$\ e_{r_\Phi}\ _\infty$	-2.02%	-0.58%	-0.39%	-0.39%

ceeds 1% only for radius 4 when no subsampling is used. For spheres with radii of at least 6, even the maximal error is below 3% without subsampling. Additional subsampling reduces the volume mapping further. As expected, $\overline{|e_{r_V}|}$ is approximately halved when incrementing s_f by one. The relative error of the potential over the whole domain, measured in the L_2 norm, is well below 2% for all considered sphere sizes, even if no subsampling is used. Only the maximal value of e_{r_Φ} may become higher than 1% without subsampling. With higher subsampling and larger spheres, all shown errors decrease until the lower bound of the discretisation error is reached. The results also show that the errors for same $s_f R_L$ match very well.

The Coulomb force computation according to Sec. 2.4 is validated by simulating a charged sphere in a domain with homogeneous electric field. For different sphere sizes and subsampling factors, the sphere is placed at the position of maximal volume mapping error near the domain center. From the simulated electrostatic force F_C^* and the analytical solution F_C , the relative error is computed by $e_{r_{F_C}} = \frac{F_C^* - F_C}{F_C}$. The domain size and physical properties of the spheres are the same as in the electric potential validation. At the left and right boundary, Dirichlet conditions with different values are applied ($\Phi_W = 0$, $\Phi_E = -10V$).

In all other dimensions, homogeneous Neumann BCs are applied. From the potential difference across the domain and the particle charge, the reference value of the Coulomb force is obtained.

The sphere is first placed at positions with maximal volume mapping error occurring for a given subsampling factor of the Coulomb force. Then $e_{r_{F_C}}$ is equal to $e_{r_V \max}$ for same subsampling factors of potential and electrostatic force. For unequal subsampling factors, $e_{r_{F_C}}$ changes only slightly. I. e., the smaller volume mapping error for the electric potential computation has only minor effect on the accuracy of the force.

For the positions with $e_{r_V \max}$ of the electric potential subsampling factor, errors are shown in Tab. 5. Applying

Table 5. Coulomb force validation in domain of size 256^3 , relative errors $e_{r_{F_C}}$ for different radii and subsampling factors. Particles are located at the position of maximal volume mapping error for potential subsampling.

		force subsampling ($s_{f_{F_C}}$)			R_L
		1	2	3	
pot. subsampling (s_{f_Φ})	4	1	7.06%	-0.49%	-0.25%
		2	-5.27%	2.58%	-0.64%
		3	-5.68%	1.48%	-1.43%
	6	1	-2.74%	0.27%	-0.13%
		2	-0.11%	-1.43%	-0.03%
		3	0.18%	-0.41%	-0.33%
	8	1	2.58%	0.20%	0.05%
		2	-1.80%	0.72%	-0.19%
		3	-0.44%	-0.04%	-0.25%
9	1	1.91%	0.26%	0.02%	
	2	0.27%	-0.33%	0.01%	
	3	1.91%	0.14%	0.32%	
12	1	-1.43%	-0.02%	-0.08%	
	2	-0.08%	-0.25%	0.02%	
	3	0.09%	-0.01%	-0.11%	

the the same subsampling to potential and electrostatic force computation results in errors equal to e_{r_V} . For different combinations of subsampling factors, the error decreases. Then the position corresponds to $e_{r_V \max}$ of the potential subsampling where e_{r_V} of the force subsampling is smaller.

We conclude that subsampling is particularly important when computing the electrostatic force, as its value is directly proportional to the particle charge. To obtain errors in the electrostatic force comparable to those of the fluid-particle interaction, a subsampling factor of two is sufficient for spheres with radius $6\delta x$ or $8\delta x$. Spheres with smaller radius require higher subsampling and those with larger radius require no subsampling.

7. Towards Realistic Multiphysics Simulations

We proceed to present two showcases that demonstrate the proper operation of the coupled multiphysics simu-

lation. For the showcases, animations are available via permalinks:

1. The attraction of homogeneously charged particles in fluid flow by an oppositely charged surface and their deposition on that surface³.
2. The separation of oppositely, homogeneously charged particles in fluid flow in a bisecting micro-channel, due to an electric field imposed by charged surfaces⁴.

Both showcases simulate a micro-channel with fluid flow in longitudinal direction. Showcase 1 with a channel size of $2.56 \text{ mm} \times 5.76 \text{ mm} \times 2.56 \text{ mm}$ is depicted in Fig. 8. The flow results from an inflow BC (Eqn. (9)) with velocity 1 mm/s at both, front and back side. Showcase 2 is depicted in Fig. 9. The channel of size $2.56 \text{ mm} \times 7.68 \text{ mm} \times 1.92 \text{ mm}$ is split into two halves after 5.76 mm by a fixed beam of height $160 \mu\text{m}$ with no-slip BCs that is attached to the side walls. The flow is imposed here by the inflow BC with velocity 1 mm/s at the front side. At the back side zero pressure outflow is modelled with a *pressure anti-bounce-back* BC from [13]. On all other boundaries, no-slip BCs are applied.

The LBM parameters are identical for both showcases. The fluid has a density of 10^3 kg/m^3 and a kinematic viscosity of $10^{-6} \text{ m}^2/\text{s}$. TRT with $\Lambda = \Lambda_{\text{mid}}$ is chosen. In order to sufficiently resolve the particles, the lattice spacing is $\delta x = 10 \mu\text{m}$. For the relaxation time $\tau = 1.7$, this results in a time increment of $\delta t = 4 \cdot 10^{-5} \text{ s}$. The density of the particles is only 14% higher than the fluid density, and thus gravity is neglected. Moreover, the lubrication correction according to Sec. 2.3 is applied for showcase 2.

For showcase 1, the homogeneously charged spheres with $8000 e$ (elementary charge e) have a radius of $60 \mu\text{m}$. The particles are initially arranged as four sets of 6^3 equidistantly placed spheres with center-to-center distance $120 \mu\text{m}$ in each dimension, as can be seen in Fig. 8(a). The electric field is perpendicular to the flow, i. e., top and bottom plane impose the field via Dirichlet BCs for the potential. In the simulation, the potential is set to $\Phi|_{\Gamma_B} = -100 \text{ V}$ at the (red) bottom plane and to $\Phi|_{\Gamma_T} = 0 \text{ V}$ at the (blue) top plane. All other boundaries are modelled as insulators, i. e., homogeneous Neumann BCs $\frac{\partial \Phi}{\partial \vec{n}}|_{\Gamma \setminus (\Gamma_B \cup \Gamma_T)} = 0$. The parameters for this showcase, and the sets of spheres, are also used for the scaling experiments of Sec. 8.

For showcase 2, the homogeneously charged spheres with radius $80 \mu\text{m}$ are inserted into the domain at random positions near the front wall. Their charge is set to $\pm 40\,000 e$ with randomly generated sign. The insertion frequency is chosen such that 4% solid volume fraction is reached, under the assumption that the spheres move with the average fluid inflow velocity. A homogeneous electric field is imposed in vertical direction via Dirichlet BCs at

the channel front section. The (blue) bottom plane is negatively charged with $\Phi|_{\Gamma_{B,f}} = -76.8 \text{ V}$, and the (red) top plane positively with $\Phi|_{\Gamma_{T,f}} = 76.8 \text{ V}$. To reduce particle agglomeration at the charged planes, homogeneous Neumann BCs are modelled at the top and bottom plane after the channel is split in two parts. These BCs are again applied to all other boundaries.

In Fig. 8, the initial setup is shown for the agglomeration simulation, together with a later state. The Poiseuille

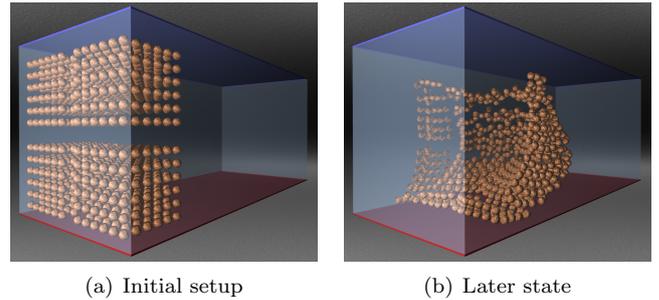


Figure 8. Initial and a later state (time step 35 200) of the charged particle agglomeration simulation.

flow profile gets visible in Fig. 8(b) from the positions of the equally charged particles that are attracted by the bottom surface while they repel each other.

In Fig. 9, simulation results are shown for the separation showcase. Fig. 9(a) shows the particle distribution after 75 300 time steps. The particles move with the flow in longitudinal direction and are separated by the electric field. While the positively charged (red) particles move towards the bottom plane and the lower half of the separated channel, the negatively charged (blue) particles move in direction of the upper half.

In order to illustrate the particle movement, pathlines of each particle are depicted in the color indicating their charge. To visualise the trajectories of all particles, and to evaluate how well particles are separated, all pathlines at the end of the simulation are shown in Fig. 9(b). Here, one can see that almost all particles move to the correct subchannel intended for their charge, independent of their initial positions. The particles that are not sorted correctly, i. e. positively charged spheres entering the upper half of the separated channel and negatively charged particles entering the lower half, were hindered by other particles moving in opposite vertical direction.

The simulations were performed on 144 processes of the LiMa cluster of the computing center RRZE in Erlangen, Germany. Showcase 1 comprises 128 000 time steps that are carried out in 10.8 h runtime. For showcase 2, 210 000 time steps are executed within 15.7 h. This performance, however, is degraded by writing visualisation data to disc, and by additional global communication required for the insertion and deletion of particles in showcase 2. Thus, the computational performance is further evaluated in Sec. 8.

³<https://www10.cs.fau.de/permalink/igahque9ie>

⁴<https://www10.cs.fau.de/permalink/kai4Iepha3>

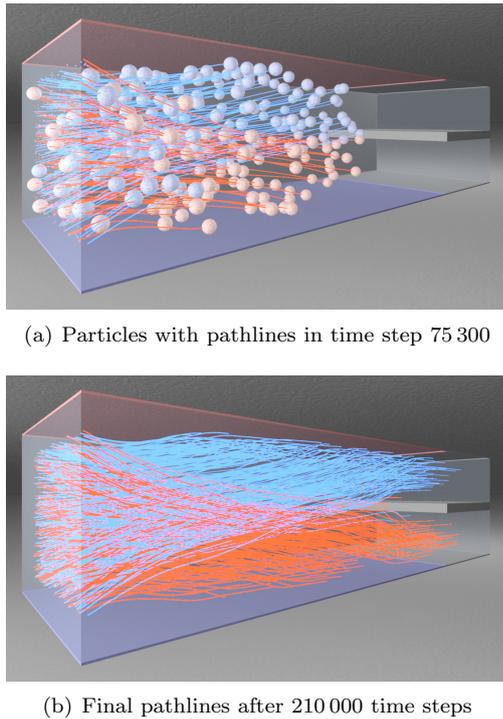


Figure 9. Particle distribution and pathlines for the charged particle separation simulation.

8. Performance Evaluation

8.1. Simulation Environment and Setup

The computational performance is evaluated on the SuperMUC supercomputer at LRZ. This cluster comprises 18 thin islands with 512 compute nodes each, connected by a high speed InfiniBand FDR10 interconnect. Each node contains two Intel Xeon E5-2680 ‘‘Sandy Bridge-EP’’ octa-core processors that are running at 2.5 GHz, and 32 GB DDR3 RAM. The code is built with the Intel 13.1 compiler, IBM MPI 1.3, and Boost 1.49. We initially show that the maximal overall performance is achieved with 16 MPI processes per node. This number of processes is used also in the later experiments that employ several nodes.

For the scaling measurements, the parameters described in Sec. 7 are used. A block of size 128^3 cells with one set of spheres of radius $R = 6\delta x$ is assigned to each process. For the single-node performance evaluation, the problem is extended once in x-direction and then only in y-direction until all 16 cores are occupied. For the weak scaling experiments, the problem size per node is kept constant. The setup on a fully occupied node with $256 \times 512 \times 256$ lattice cells and 3456 particles is shown in Fig. 10. The weak scaling problem size is extended successively in all three dimensions, retaining maximal extension in longitudinal direction (to keep the aspect ratio of the domain comparable). The maximum problem size on 2048 nodes comprises $68.7 \cdot 10^9$ lattice cells and $7.08 \cdot 10^6$ particles.

About 9.4% of the domain is covered by moving obstacle cells. These cells are not updated by the LBM. Nev-

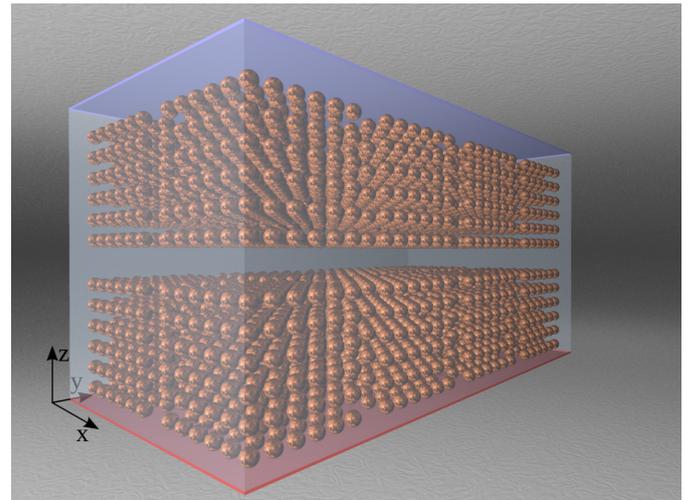


Figure 10. Charged particle simulation setup on a single node for weak scaling experiments.

ertheless, they are used for the potential computation. We define the measure *million lattice updates per second* (MLUPS) as the number of cells the Poisson problem can be solved for within one second. This metric indicates the performance to obtain the solution up to a given accuracy. In the case of the LBM we refer to *million fluid lattice updates per second* (MFLUPS) [60] considering only fluid cells.

Appropriate to the sphere radius, we use a subsampling factor of two for both, setting the RHS of Poisson’s equation and the electrostatic force computation. To find a suitable residual threshold for the MG solver, the L_2 norms of the residual and error are compared for several V(2,2)-cycles in Fig. 11. The error is strongly reduced up to a residual norm $1.5 \cdot 10^{-8}$. Once the residual norm is smaller than 10^{-9} , it is hardly reduced at all. In the scaling

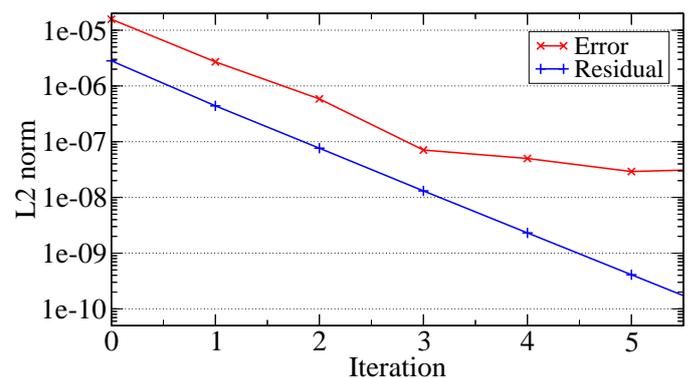


Figure 11. Error and residual after several V(2,2)-cycles

experiments, V(3,3)-cycles are performed on seven levels, with 2^3 cells per block on the coarsest level. At the first time step, an initial solution is computed, requiring five V(3,3)-cycles. Later steps only update the previous solution, requiring one V(3,3)-cycle. The corresponding average L_2 residual norm over all time steps is about $3 \cdot 10^{-9}$.

This norm is monitored and never exceeds $1.5 \cdot 10^{-8}$.

8.2. Single-Node Performance

Before evaluating the scaling behaviour on several islands, we examine the single-node weak scaling performance. An increasing number of processes is allocated to the processors in a round-robin fashion, and 240 time steps are performed on the problem described in Sec. 8.1 (128^3 cells per block). The problem on the coarsest grid with up to 128 cells (16 blocks, 7 MG levels) is solved with 6 CG iterations. This number of iterations corresponds to the diameter of the maximal problem size on the coarsest grid.

The total runtime increases from 258s on one core to 294s on 16 cores. This is equivalent to 88% parallel efficiency of the overall algorithm. Fig. 12 shows the speedup of the whole algorithm and its main parts for up to 16 MPI processes. The LBM scales almost ideally and

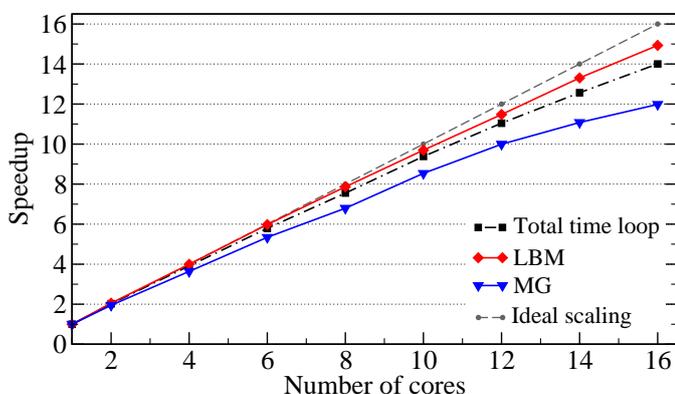


Figure 12. Single-node weak scaling performance of MG and LBM sweep for 240 time steps on SuperMUC.

achieves 93% parallel efficiency. Its performance is restricted only by computation-related data transfers, with negligible communication between the processes. MG performance is additionally reduced by up to 10% relative share of intra-node communication on the MG runtime, resulting in 75% parallel efficiency on a full node.

The single-node performance of the LBM sweep is 51.4 MFLUPS, and of the MG sweep 91.8 MLUPS. These values will serve as base performance in the following section. We employ 16 processes per node because this leads to the highest overall performance.

Considering the complicated physics simulated in the present paper, both the LBM and the MG performance compare favourably to the results in [61] and [62], respectively.

About 2×78 M(F)LUPS are achieved in [61] for a pure LBM kernel on a SuperMUC compute node, utilizing SIMD and running at 2.7 GHz. While the pure LBM can fully exploit the main memory, the multiphysics simulations have to fit data structures for various fields into memory. The resulting reduced block size leads to a slight deterioration of the performance (cf. [9]). Moreover, the moving obstacles degrade the performance further as they cause

a non-consecutive structure of fluid cells, leading to an irregular memory access pattern.

The highly optimized SuperMUC single-node performance reported for the finite element MG solver in [62] corresponds to 170 MLUPS. This performance is measured for V(3,3)-cycles and is applied to a Poisson problem with Dirichlet BCs.

8.3. Weak Scaling

We perform weak scaling experiments for up to 2048 nodes on SuperMUC for 240 time steps. The problem size is successively doubled in all dimensions, as shown in Tab. 6. The number of CG iterations required to solve

Table 6. Number of required CG coarse grid iterations for different problem sizes.

#n.	size	#iter.	#n.	size	#iter.
1	$2 \times 2 \times 4$	6	64	$8 \times 8 \times 16$	26
2	$2 \times 2 \times 8$	10	128	$8 \times 8 \times 32$	52
4	$4 \times 2 \times 8$	12	256	$16 \times 8 \times 32$	54
8	$4 \times 4 \times 8$	12	512	$16 \times 16 \times 32$	54
16	$4 \times 4 \times 16$	24	1024	$32 \times 16 \times 64$	114
32	$8 \times 4 \times 16$	26	2048	$32 \times 32 \times 64$	116

the coarsest grid problem is depicted for different problem sizes. When doubling the domain in all three dimensions, the number of CG iterations approximately doubles. This corresponds to the expected behaviour that the required number of iterations scales with the diameter of the problem size [20], according to the growth in the condition number [63]. However, when doubling the problem size, CG iterations sometimes stay constant or have to be increased. This results from different shares of Neumann and Dirichlet BCs on the boundary. Whenever the relative proportion of Neumann BCs increases, convergence deteriorates and more CG iterations are necessary.

The runtimes of all parts of the algorithm are shown in Fig. 13 for different problem sizes, indicating their shares on the total runtime. This diagram is based on the maximal (for MG, LBM, pe) or average (others) runtimes of the different sweeps among all processes. The upper part of the diagram shows the cost of fluid-simulation related sweeps, such as LBM, moving obstacle mapping (Map), hydrodynamic force computation (HydrF), and lubrication correction (Lubr) sweep. In the middle, the cost of the pe sweep is shown. Below, the costs of sweeps related to the electric potential computation are displayed. These include MG, setting the RHS of Poisson's equation (SetRHS), communication of the electric potential before the gradient computation (PtCm), and the sweep computing the electrostatic force (ElectF).

For a more precise evaluation, the exact figures are shown in Tab. 7 for one node and 2048 nodes. The total runtime (Whl) is less than the sum of the individual sweeps, since different sweeps are slow on different processes. Sweeps whose runtimes depend on the problem size—mainly due

Table 7. Time of the whole algorithm and its sweeps for 240 time steps on a single node and on 2048 nodes.

#n.	Whl	LBM	MG	<i>pe</i>	PtCm	Oth
	t[s]	t[s] ([%])	t[s] ([%])	t[s]	t[s]	t[s]
1	294	143 (48)	88 (30)	2	3	61
2048	353	157 (41)	136 (35)	27	7	60

to increasing MPI communication—are LBM, MG, *pe*, and PtCm. Overall, LBM and MG take up more than 75% of the total time, w.r.t. the runtimes of the individual sweeps.

The sweeps that scale perfectly—HydrF, LubrC, Map, SetRHS, and ElectF—are summarised as ‘Oth’. For longer simulation times the particles attracted by the bottom wall are no longer evenly distributed, possibly causing load imbalances. However, they hardly affect the overall performance. For the simulation of showcase 2, the relative share of the lubrication correction is below 0.1%, and each other sweep of ‘Oth’ is well below 4% of the total runtime.

Overall, the coupled multiphysics algorithm achieves 83% parallel efficiency on 2048 nodes. Since most time is spent to execute LBM and MG, we will now turn to analyse them in more detail. Fig. 14 displays the parallel performance for different numbers of nodes. On 2048 nodes, MG executes 121,083 MLUPS, corresponding to a parallel efficiency of 64%. The LBM performs 95,372 MFLUPS, with 91% parallel efficiency.

Finally, the average execution times of different MG sub-sweeps are presented in Fig. 15. The parallel performance is degraded in particular by communication (Comm), but also by the CG coarse grid solver (CrsCG) due to the increasing number of required iterations. Clearly, the CG-based coarsest grid solver is a non-scalable part of the solver that will ultimately dominate the performance. However, for computations up to the size presented in this paper, the overhead is still in an acceptable range.

All other sweeps show ideal scaling. Among these sweeps, most time is spent for smoothing (Smoot) and residual computation (Resid). All other parts require only a small portion of the overall MG time. These are the sweeps for

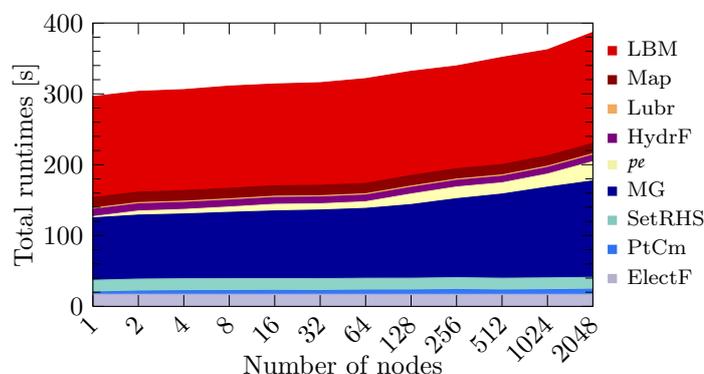


Figure 13. Runtimes of charged particle algorithm sweeps for 240 time steps for an increasing number of nodes.

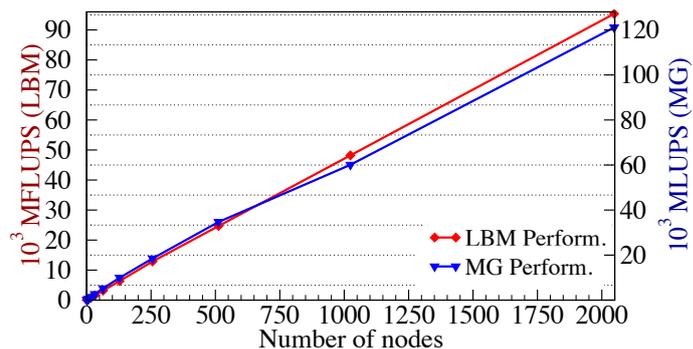


Figure 14. Weak scaling performance of MG and LBM sweep for 240 time steps.

prolongation (Proton, 2.7 s), restriction (Restr, 1.2 s), and the sweeps summarised as ‘Other’, i. e., those for adapting the RHS to the BCs (about 1 s) and setting the solution on the following coarser grid with zero (0.2 s). The sweeps for checking the termination criterion and adapting the stencils to the BCs at the beginning of the simulation are negligible (< 0.01 s).

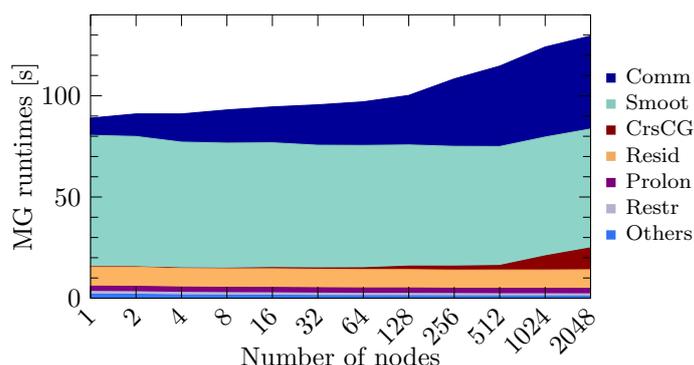


Figure 15. Average runtime of different MG sweeps for 240 time steps for an increasing number of nodes.

9. Conclusion

In this article we demonstrate the potential of coupled multiphysics simulations using the LBM, a physics engine, and a scalar potential solver. These components are integrated into the parallel software framework WALBERLA to simulate fully resolved charged particles in microfluidic flows. We describe the software design, validate the physical correctness for different model parameters, and present the excellent computational performance and scalability of the algorithms. The results show that the multiphysics algorithm allows the physically correct simulation of several millions of interacting charged particles, as required for advanced realistic multiphysics simulations.

For the presented simulations, WALBERLA is extended by a parallel cell-centered MG method to compute the electrostatic potential including an efficient handling of the boundary conditions. The code is carefully optimised by

exploiting that the stencils on the finest level are (almost) constant and thus the main memory access can be reduced, while retaining the flexibility of Galerkin coarsening.

The fluid simulation and the fluid-particle coupling are validated for low Reynolds number flows for spheres arranged in a periodic array. The largest relative error for the drag is shown to be below 2.2% for various sphere resolutions, lattice viscosities, and values of the ‘magic’ TRT parameter Λ . For commonly used parameters and sphere resolutions, the relative error is around 1% or even below. We further observe that the value of the parameter Λ that is optimised for porous media yields slightly better results than the alternative value which is chosen to set the boundary mid-way between two lattice sites. For the higher examined lattice viscosity the error can be reduced further. The error analysis is extended to an evaluation of volume mapping errors, and estimates of the drag that is corrected taking these errors into account. Furthermore, the effect and the correctness of the lubrication correction method are shown for different sphere resolutions, sphere velocities, and lattice viscosities at low Reynolds numbers. The influence of these parameters on the lubrication correction is analysed carefully in the situation of sphere-sphere and sphere-wall interaction. This confirms that for the TRT model the lattice viscosity has only a minor impact.

To validate the interaction with electrostatic forces, the computed electric potential and the resulting force on a single charged sphere are analysed for different sphere resolutions. A subsampling of the sphere geometry is used to reduce volume mapping errors and thus to improve the accuracy of the electrostatic force computation, while keeping the mesh conforming to the LBM and without increasing the systems that must be solved. The relative error in the computed electric potential decreases with increasing sphere resolutions and with more levels of subsampling, however, the error in the electric potential stagnates once the finite volume discretisation error is reached. The largest relative error of the potential is well below 2% for the resolutions presented. Subsampling is found to have more influence on the electrostatic force than on the electric potential. The subsampling for the electrostatic force computation is found to be necessary only for sphere radii below nine lattice spacings to obtain errors that are of comparable magnitude to those arising from the fluid-particle interaction.

The parallel performance of the multiphysics algorithm is examined on SuperMUC. The overall algorithm has a parallel efficiency of 83% on 32,768 cores, simulating more than seven million charged particles. Most time of the algorithm is spent for the LBM and the MG solver. For both, the single-node performance compares favourably to optimised reference implementations of pure LBM and finite element MG. The MG scales with 64% parallel efficiency, achieving more than $121 \cdot 10^9$ cell updates per second to solve Poisson’s equation with mixed boundary conditions. Here, the reduced scalability results from an increasing number of coarsest grid solver iterations and in-

creasing MPI communication, with many small messages on the coarse grids. The LBM for fluid-particle interaction scales almost perfectly with 91% parallel efficiency, achieving more than $95 \cdot 10^9$ fluid cell updates per second.

The example scenarios and their animations demonstrate the qualitative correctness of the overall algorithm and hint at future applications. The present algorithm can be applied to find design parameters, e. g. for optimal separation efficiency of oppositely charged particles as in the separation scenario, depending on fluid and particle properties. Moreover, the algorithm allows future simulations of large numbers of charged particles in more complex geometries, such as electrostatic filters or pulmonary drug delivery, also for higher Reynolds numbers.

The MG solver allows to simulate multiphysics applications beyond those in this article, such as simulations comprising temperature fields described by Poisson’s equation. Also jumping dielectricity coefficients can be realised easily to simulate dielectrophoresis. The solver module and BC treatment can be applied to more complicated elliptic partial differential equations, such as the Poisson-Boltzmann equation. This equation allows to simulate charged particles in electrolyte solutions, including double-layer effects that are of relevance for simulating lab-on-a-chip systems.

Acknowledgements

The authors would like to thank Simon Bogner, Tobias Preclik, Daniel Ritter, Ehsan Fattahi, and Bjrn Gmeiner for valuable discussions, and Christian Kuschel, Markus Huber, and Gaby Fleig for support in the correction process. The authors are grateful to the RRZE and LRZ for providing the computational resources on LiMa and SuperMUC, respectively.

References

References

- [1] H. Yanada, K. D. Tran, Fundamental Investigation of Charge Injection Type of Electrostatic Oil Filter, *J. Adv. Mech. Des. Syst.* 2 (1) (2008) 119–132. doi:10.1299/jamdsm.2.119.
- [2] A. G. Bailey, A. H. Hashish, T. J. Williams, Drug delivery by inhalation of charged particles, *J. Electrostat.* 44 (1–2) (1998) 3–10. doi:10.1016/S0304-3886(98)00017-5.
- [3] P. W. Longest, L. T. Holbrook, In silico models of aerosol delivery to the respiratory tract — Development and applications, *Adv. Drug Deliv. Rev.* 64 (4) (2012) 296–311. doi:10.1016/j.addr.2011.05.009.
- [4] Y. Kang, D. Li, Electrokinetic motion of particles and cells in microchannels, *Microfluid Nanofluidics* 6 (4) (2009) 431–460. doi:10.1007/s10404-009-0408-7.
- [5] S. Donath, J. Gtz, S. Bergler, C. Feichtinger, K. Iglberger, U. Rde, waLBerla: The Need for Large-Scale Super Computers, in: *High Performance Computing in Science and Engineering. Garching/Munich 2007*, Springer, 2008, pp. 459–473.
- [6] C. Feichtinger, S. Donath, H. Kstler, J. Gtz, U. Rde, WaLBerla: HPC software design for computational engineering simulations, *J. Comput. Sci.* 2 (2) (2011) 105–112. doi:10.1016/j.jocs.2011.01.004.

- [7] S. Chen, G. D. Doolen, Lattice Boltzmann Method for Fluid Flows, *Ann. Rev. Fluid Mech.* 30 (1) (1998) 329–364. doi:10.1146/annurev.fluid.30.1.329.
- [8] C. K. Aidun, J. R. Clausen, Lattice-Boltzmann method for complex flows, *Ann. Rev. Fluid Mech.* 42 (1) (2010) 439–472. doi:10.1146/annurev-fluid-121108-145519.
- [9] J. Götz, K. Iglberger, C. Feichtinger, S. Donath, U. Rüde, Coupling multibody dynamics and computational fluid dynamics on 8192 processor cores, *Parallel Comput.* 36 (2) (2010) 142–141. doi:10.1016/j.parco.2010.01.005.
- [10] J. Götz, K. Iglberger, M. Stürmer, U. Rüde, Direct Numerical Simulation of Particulate Flows on 294912 Processor Cores, in: *Proc. 2010 ACM/IEEE Proc. Int. Conf. for High Performance Computing, Networking, Storage and Analysis, SC '10, IEEE, 2010*, pp. 1–11.
- [11] K. Iglberger, U. Rüde, Massively Parallel Rigid Body Dynamics Simulation, *OSRD* 23 (3–4) (2009) 159–167. doi:10.1007/s00450-009-0066-8.
- [12] K. Iglberger, Software design of a massively parallel rigid body framework, Ph.D. thesis, University of Erlangen-Nürnberg (2010).
- [13] I. Ginzburg, F. Verhaeghe, D. d’Humières, Two-relaxation-time lattice Boltzmann scheme: About parametrization, velocity, pressure and mixed boundary conditions, *Commun. Comput. Phys.* 3 (2) (2008) 427–478.
- [14] D. Yu, R. Mei, L.-S. Luo, W. Shyy, Viscous flow computations with the method of lattice Boltzmann equation, *Prog. Aerosp. Sci.* 39 (5) (2003) 329–367. doi:10.1016/S0376-0421(03)00003-4.
- [15] A. J. C. Ladd, Numerical simulations of particulate suspensions via a discretized Boltzmann equation. Part 1. Theoretical foundation, *J. Fluid Mech.* 271 (1994) 285–309. doi:10.1017/S0022112094001771.
- [16] C. K. Aidun, Y. Lu, E.-J. Ding, Direct analysis of particulate suspensions with inertia using the discrete Boltzmann equation, *J. Fluid Mech.* 373 (1) (1998) 287–311. doi:10.1017/S0022112098002493.
- [17] A. J. C. Ladd, Numerical simulations of particulate suspensions via a discretized Boltzmann equation. Part 2. Numerical results, *J. Fluid Mech.* 271 (1994) 311–339. doi:10.1017/S0022112094001783.
- [18] N. Q. Nguyen, A. J. C. Ladd, Lubrication corrections for lattice-Boltzmann simulations of particle suspensions, *Phys. Rev. E* 66 (4) (2002) 046708. doi:10.1103/PhysRevE.66.046708.
- [19] M. Mohr, R. Wienands, Cell-centred Multigrid Revisited, *Comp. Vis. Sci.* 7 (3–4) (2004) 129–140. doi:10.1007/s00791-004-0137-0.
- [20] B. Gmeiner, H. Köstler, M. Stürmer, U. Rüde, Parallel multi-grid on hierarchical hybrid grids: a performance study on current high performance computing clusters, *Concurrency Computat.: Pract. Exper.* 26 (1) (2014) 217–240. doi:10.1002/cpe.2968.
- [21] F. Hülsemann, M. Kowarschik, M. Mohr, U. Rüde, Parallel Geometric Multigrid, in: *Numerical Solution of Partial Differential Equations on Parallel Computers*, Vol. 51 of *Lect. Notes Comput. Sci. Eng.*, Springer, 2006, Ch. 5, pp. 165–208. doi:10.1007/3-540-31619-1_5.
- [22] S. Succi, *The Lattice Boltzmann Equation: For Fluid Dynamics and Beyond*, Numerical Mathematics and Scientific Computation, Oxford University Press, 2001.
- [23] D. Hänel, *Molekulare Gasdynamik*, Springer, 2004.
- [24] D. A. Wolf-Gladrow, *Lattice-Gas Cellular Automata and Lattice Boltzmann Models: An Introduction*, Springer, 2000.
- [25] P. L. Bhatnagar, E. P. Gross, M. Krook, A Model for Collision Processes in Gases. I. Small Amplitude Processes in Charged and Neutral One-Component Systems, *Phys. Rev.* 94 (3) (1954) 511–525. doi:10.1103/PhysRev.94.511.
- [26] L.-S. Luo, W. Liao, X. Chen, Y. Peng, W. Zhang, Numerics of the lattice Boltzmann method: Effects of collision models on the lattice Boltzmann simulations, *Phys. Rev. E* 83 (5) (2011) 056710. doi:10.1103/PhysRevE.83.056710.
- [27] D. d’Humières, Generalized lattice-Boltzmann equations, in: *Rarefied Gas Dynamics: Theory and Simulations*, Vol. 159 of *Prog. Astronaut. Aeronaut.*, 1992, pp. 450–458.
- [28] I. Ginzburg, J.-P. Carlier, C. Kao, Lattice Boltzmann approach to Richards’ equation, in: *Computational Methods in Water Resources*, Vol. 55 of *Developments in Water Science*, Elsevier, 2004, pp. 583–595. doi:10.1016/S0167-5648(04)80083-2.
- [29] I. Ginzburg, F. Verhaeghe, D. d’Humières, Study of simple hydrodynamic solutions with the two-relaxation-times lattice Boltzmann scheme, *Commun. Comput. Phys.* 3 (3) (2008) 519–581.
- [30] X. He, L.-S. Luo, Lattice Boltzmann Model for the Incompressible Navier–Stokes Equation, *J. Stat. Phys.* 88 (3) (1997) 927–944. doi:10.1023/B:J0SS.0000015179.12689.e4.
- [31] Y. H. Qian, D. d’Humières, P. Lallemand, Lattice BGK Models for Navier–Stokes Equation, *Europhys. Lett.* 17 (6) (1992) 479. doi:10.1209/0295-5075/17/6/001.
- [32] I. Ginzburg, D. d’Humières, Multireflection boundary conditions for lattice Boltzmann models, *Phys. Rev. E* 68 (6) (2003) 066614. doi:10.1103/PhysRevE.68.066614.
- [33] M. Bouzidi, M. Firdaouss, P. Lallemand, Momentum transfer of a Boltzmann-lattice fluid with boundaries, *Phys. Fluids* 13 (11) (2001) 3452–3459. doi:10.1063/1.1399290.
- [34] D. d’Humières, I. Ginzburg, Viscosity Independent Numerical Errors for Lattice Boltzmann Models: From Recurrence Equations to “Magic” Collision Numbers, *Comput. Math. Appl.* 58 (5) (2009) 823–840. doi:10.1016/j.camwa.2009.02.008.
- [35] S. Bogner, S. Mohanty, U. Rüde, Drag correlation for dilute and moderately dense fluid-particle systems using the lattice Boltzmann method, arXiv preprint arXiv:1401.2025.
- [36] R. Mei, D. Yu, W. Shyy, L.-S. Luo, Force evaluation in the lattice Boltzmann method involving curved geometry, *Phys. Rev. E* 65 (4) (2002) 041203. doi:10.1103/PhysRevE.65.041203.
- [37] G. Bossis, J. F. Brady, Dynamic simulation of sheared suspensions. i. general method, *J. Chem. Phys.* 80 (10) (1984) 5141–5154. doi:10.1063/1.446585.
- [38] R. G. Cox, The motion of suspended particles almost in contact, *Int. J. Multiphase Flow* 1 (2) (1974) 343–371. doi:10.1016/0301-9322(74)90019-6.
- [39] F. Janoschek, J. Harting, F. Toschi, Accurate lubrication corrections for spherical and non-spherical particles in discretized fluid simulations, arXiv preprint arXiv:1308.6482.
- [40] A. J. C. Ladd, Sedimentation of homogeneous suspensions of non-Brownian spheres, *Phys. Fluids* 9 (3) (1997) 491–499. doi:10.1063/1.869212.
- [41] I. L. Claeys, J. F. Brady, Lubrication singularities of the grand resistance tensor for two arbitrary particles, *Phys. Chem. Hydrodyn.* 11 (3) (1989) 261–293.
- [42] A. J. C. Ladd, R. Verberg, Lattice-Boltzmann Simulations of Particle-Fluid Suspensions, *J. Stat. Phys.* 104 (5) (2001) 1191–1251. doi:10.1023/A:1010414013942.
- [43] J. F. Brady, G. Bossis, Stokesian dynamics, *Ann. Rev. Fluid Mech.* 20 (1) (1988) 111–157. doi:10.1146/annurev.fl.20.010188.000551.
- [44] C. Hirsch, *Numerical Computation of Internal and External Flows: The Fundamentals of Computational Fluid Dynamics*, Vol. 1, Butterworth-Heinemann, 2007.
- [45] R. Eymard, T. Gallouët, R. Herbin, Finite volume methods, in: P. Ciarlet, J. Lions (Eds.), *Techniques of Scientific Computing (Part 3)*, Vol. 7 of *Handbook of Numerical Analysis*, Elsevier, 2000, pp. 713–1018. doi:10.1016/S1570-8659(00)07005-8.
- [46] R. Ramadugu, S. P. Thampi, R. Adhikari, S. Succi, S. Ansumali, Lattice differential operators for computational physics, *Europhys. Lett.* 101 (5) (2013) 50006. doi:10.1209/0295-5075/101/50006.
- [47] A. Brandt, Multi-Level Adaptive Solutions to Boundary-Value Problems, *Math. Comput.* 31 (138) (1977) 333–390. doi:10.1090/S0025-5718-1977-0431719-X.
- [48] U. Trottenberg, C. Oosterlee, A. Schüller, *Multigrid*, Academic Press, 2001.
- [49] R. E. Alcouffe, A. Brandt, J. E. Dendy, J. W. Painter, The

- Multi-Grid Method for the Diffusion Equation with Strongly Discontinuous Coefficients, *SIAM J. Sci. and Stat. Comp* 2 (4) (1981) 430–454. doi:10.1137/0902035.
- [50] W. Hackbusch, *Multi-Grid Methods and Applications*, Vol. 4 of Springer Ser. Comput. Math., Springer, 1985.
- [51] D. Braess, Towards Algebraic Multigrid for Elliptic Problems of Second Order, *Computing* 55 (4) (1995) 379–393. doi:10.1007/BF02238488.
- [52] J. Bramble, R. Ewing, J. Pasciak, J. Shen, The analysis of multi-grid algorithms for cell centered finite difference methods, *Adv. Comput. Math.* 5 (1) (1996) 15–29. doi:10.1007/BF02124733.
- [53] D. Bartuschat, D. Ritter, U. Rüde, Parallel multigrid for electrokinetic simulation in particle-fluid flows, in: *High Performance Computing and Simulation (HPCS)*. Madrid 2012, IEEE, 2012, pp. 374–380. doi:10.1109/HPCSim.2012.6266940.
- [54] C. Feichtinger, Design and Performance Evaluation of a Software Framework for Multi-Physics Simulations on Heterogeneous Supercomputers, Ph.D. thesis, University of Erlangen-Nürnberg (2012).
- [55] A. S. Sangani, A. Acrivos, Slow flow through a periodic array of spheres, *Int. J. Multiphase Flow* 8 (4) (1982) 343–360. doi:10.1016/0301-9322(82)90047-7.
- [56] H. Hasimoto, On the periodic fundamental solutions of the Stokes equations and their application to viscous flow past a cubic array of spheres, *J. Fluid Mech.* 5 (2) (1959) 317–328. doi:10.1017/S0022112059000222.
- [57] L.-S. Luo, Unified Theory of Lattice Boltzmann Models for Nonideal Gases, *Phys. Rev. Lett.* 81 (8) (1998) 1618–1621. doi:10.1103/PhysRevLett.81.1618.
- [58] J. Götz, Massively Parallel Direct Numerical Simulation of Particulate Flows, Ph.D. thesis, University of Erlangen-Nürnberg (2012).
- [59] E. J. Ding, C. K. Aidun, Extension of the lattice-Boltzmann method for direct simulation of suspended particles near contact, *J. Stat. Phys.* 112 (3) (2003) 685–708. doi:10.1023/A:1023880126272.
- [60] G. Wellein, T. Zeiser, G. Hager, S. Donath, On the single processor performance of simple lattice Boltzmann kernels, *Comput. Fluids* 35 (8) (2006) 910–919. doi:10.1016/j.compfluid.2005.02.008.
- [61] C. Godenschwager, F. Schornbaum, M. Bauer, H. Köstler, U. Rüde, A Framework for Hybrid Parallel Flow Simulations with a Trillion Cells in Complex Geometries, in: *Proc. Int. Conf. on High Performance Computing, Networking, Storage and Analysis, SC '13*, ACM, 2013, pp. 35:1–35:12. doi:10.1145/2503210.2503273.
- [62] B. Gmeiner, U. Rüde, Peta-Scale Hierarchical Hybrid Multigrid Using Hybrid Parallelization, in: *Large-Scale Scientific Computing, Lecture Notes in Computer Science*, Springer, 2014, pp. 439–447. doi:10.1007/978-3-662-43880-0_50.
- [63] J. R. Shewchuk, An Introduction to the Conjugate Gradient Method Without the Agonizing Pain, Tech. rep., Pittsburgh, PA, USA (1994).