

Enhanced computational performance of the lattice Boltzmann model for simulating micron- and submicron-size particle flows and non-Newtonian fluid flows

Hakan Başağaoğlu^{a,*}, John R. Harwell^b, Hoa Nguyen^e, Sauro Succi^f

^a*Mechanical Engineering Division, Southwest Research Institute, San Antonio, TX 78238 USA*

^b*Defense Intelligence Solutions Division, Southwest Research Institute, San Antonio, TX 78238 USA*

^c*Department of Mathematics, Trinity University, San Antonio, TX 78212 USA*

^d*Istituto Applicazioni Calcolo, CNR-IAC Viale del Policlinico 137, 00161 Rome, Italy*

^e*Department of Mathematics, Trinity University, San Antonio, TX 78212 USA*

^f*Istituto Applicazioni Calcolo, CNR-IAC Viale del Policlinico 137, 00161 Rome, Italy*

Abstract

Significant improvements in the computational performance of the lattice-Boltzmann (LB) model, coded in FORTRAN90, were achieved through application of enhancement techniques. Applied techniques include optimization of array memory layouts, data structure simplification, random number generation outside the simulation thread(s), code parallelization via OpenMP, and intra- and inter-timestep task pipelining. Effectiveness of these optimization techniques was measured on three benchmark problems: (i) transient flow of multiple particles in a Newtonian fluid in a heterogeneous fractured porous domain, (ii) thermal fluctuation of the fluid at the sub-micron scale and the resultant Brownian motion of a particle, and (iii) non-Newtonian fluid flow in a smooth-walled channel. Application of the aforementioned optimization techniques resulted in an average $21\times$ performance improvement, which could significantly enhance practical uses of the LB models in diverse applications, focusing on the fate and transport of nano-size or micron-size particles in non-Newtonian fluids.

*Corresponding author

Email address: hasagaoglu@swri.org (Hakan Başağaoğlu)

Keywords: Computational methods in fluid dynamics, Hydrodynamics, Lattice-Boltzmann, FORTRAN, Optimization

1. Introduction

The lattice-Boltzmann (LB) method has emerged as a versatile computational fluid dynamics tool [1, 2] to simulate mesoscale simulations of flow and transport of nano- and micron-size particles [3, 4, 5, 6], multiphase and multicomponent flows [7, 8, 9, 10], laminar and turbulent flows [11, 12, 13], and non-Newtonian and viscoelastic flows [14, 15, 16] in geometrically-complex flow domains, with its potential applications in diverse fields extending from biomedical to the energy sector. The local nature of the calculations in the LB method makes the LB-based models amenable for code parallelization. Moreover, efficiencies in handling complex flow domain geometries and the explicit nature of the computational steps in the LB method facilitate its use in diverse applications. However, the computational efficiency of the LB-based numerical models need to be further improved to enhance their practical uses.

Improvements on the computational performance of three modules of the LB-based numerical model coded in FORTRAN90 are reported in this paper; however, these optimization techniques are not limited to FORTRAN90 coding. The first module was developed to simulate the flow of multiple circular-cylinder micron-size particles in pressure-driven Newtonian fluid flow in a heterogeneous fractured porous domain formed by an array of non-overlapping, non-uniform stationary solid obstacles. The second module focused on the flow of nano-size particles that exhibit Brownian motion induced by thermal fluctuations in the fluid. The third module was developed to simulate non-Newtonian fluid (pseudoplastic and dilatant) flows in a confined channel.

Several common LB performance enhancement techniques were initially applied to enhance the computational efficiency of the LB model. These techniques included the maximization of compile-time determinism, optimization of array memory layout for the streaming step, and OpenMP fine-tuning, which achieved

5-10× performance increases. Subsequently, additional techniques were applied to gain further performance increases, which involved generation of random numbers in a producer-consumer model and exploitation of task-level parallelism to perform task pipelining and output offloading. This yielded an additional 2-4× performance increase on simulations with thermal fluctuations and/or heavy I/O loads. Overall, an average of 21× computational performance increase was achieved through the combined application of these optimization techniques.

2. Lattice-Boltzmann (LB) Model

In the lattice-Boltzmann (LB) method, the mesodynamics of the incompressible, Newtonian fluid flow is described by [1, 2, 17] with a single relaxation time via the continuous Bhatnagar-Gross-Krook (BKG) equation [18]

$$f_i(\mathbf{r} + \mathbf{e}_i \Delta t, t + \Delta t) - f_i(\mathbf{r}, t) = \frac{\Delta t}{\tau} [f_i^{eq}(\mathbf{r}, t) - f_i(\mathbf{r}, t)], \quad (1)$$

where $f_i(\mathbf{r}, t)$ is the complete set of population densities of discrete velocities \mathbf{e}_i at position \mathbf{r} and discrete time t with a time increment of Δt , τ is the relaxation parameter, and f_i^{eq} is the local equilibrium [19]

$$f_i^{eq} = \omega_i \rho \left(1 + \frac{\mathbf{e}_i \cdot \mathbf{u}}{c_s^2} + \frac{(\mathbf{e}_i \cdot \mathbf{u})^2}{2c_s^4} - \frac{\mathbf{u} \cdot \mathbf{u}}{2c_s^2} \right), \quad (2)$$

where ω_i is the weight associated with \mathbf{e}_i and c_s is the speed of sound, $c_s = \Delta x / \sqrt{3} \Delta t$. The local fluid density, ρ , and velocity, \mathbf{u} , at the lattice node are given by $\rho = \sum_i f_i$ and $\rho \mathbf{u} = \sum_i f_i \mathbf{e}_i + \tau \rho \mathbf{g}$, where \mathbf{g} is the acceleration due to the force of gravity [20]. The left hand side of Eq. 1 describes the streaming of population densities, f_i , from a lattice node \mathbf{r} to the closest neighboring lattice node in the direction of \mathbf{e}_i on a regular lattice grid. The right hand side of Eq. 1 describes the local collision process. A D2Q9 (two-dimensional nine velocity vector) lattice [1] was used in numerical simulations in this paper. For a D2Q9 model, $e_0 = (0, 0)$, $e_i = \left[\cos\left(\frac{(i-1)\pi}{2}\right), \sin\left(\frac{(i-1)\pi}{2}\right) \right]$ for $i = 1-4$ (for the primitive lattice vectors of length unit), and $e_i = \sqrt{2} \left[\cos\left(\frac{(2i-9)\pi}{4}\right), \sin\left(\frac{(2i-9)\pi}{4}\right) \right]$ for $i = 5-8$ (for diagonal lattice vectors of length $\sqrt{2}$), $\omega_0 = \frac{4}{9}$, $\omega_i = \frac{1}{9}$ for

$i = 1-4$ and $\omega_i = \frac{1}{36}$ for $i = 5-8$. Through the Chapman-Enskog expansion, the LB method for a single-phase flow recovers the Navier-Stokes equation for weakly compressible fluids

$$\nabla \cdot \mathbf{u} = 0, \quad (3)$$

$$\partial_t \mathbf{u} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\frac{\nabla p}{\rho} + \nu \nabla^2 \mathbf{u} + \mathbf{g}, \quad (4)$$

with the fluid kinematic viscosity, $\nu = c_s^2 \Delta t (\tau - 0.5)$. τ is computed as a function of ν and determines how fast population densities approach the equilibrium distributions upon collision in Eq. 1. Pressure, P , is computed via the ideal gas relation, $P = c_s^2 \rho$.

3. Microparticle Lattice-Boltzmann Module

The LB method described in Sec. 2 was extended to simulate flow of micron-size particle in Newtonian fluids in low - moderate Reynolds number flows by accommodating particle-fluid hydrodynamics [6, 21, 22], based on model formulations in Ref. [3, 4, 23, 24]. In these formulations, the population densities near particle surfaces were modified to account for particle-fluid hydrodynamic forces, $\mathbf{F}_{\mathbf{r}_b}$, which arise from momentum exchanges between the fluid outside the particle and the particle in motion [3, 23]. $\mathbf{F}_{\mathbf{r}_b}$ was calculated at particle boundary nodes located halfway between the intra-particle lattice node, \mathbf{r}_v , and extra-particle lattice node, $\mathbf{r}_v + \mathbf{e}_i$, (Fig. 1)

$$\mathbf{F}_{\mathbf{r}_b} = -2 \left[f'_i(\mathbf{r}_v + \mathbf{e}_i \Delta t, t^*) + \frac{\rho \omega_i}{c_s^2} (\mathbf{u}_{\mathbf{r}_b} \cdot \mathbf{e}_i) \right] \mathbf{e}_i, \quad (5)$$

where f'_i is the population density in the $-\mathbf{e}_i$ direction at the post-collision time t^* , and $\mathbf{u}_{\mathbf{r}_b}$ is the local particle velocity at the boundary node \mathbf{r}_b . Repulsive interaction forces between the particles and between the particles and stationary solid zones, including channel walls and inline obstacles, were expressed in terms

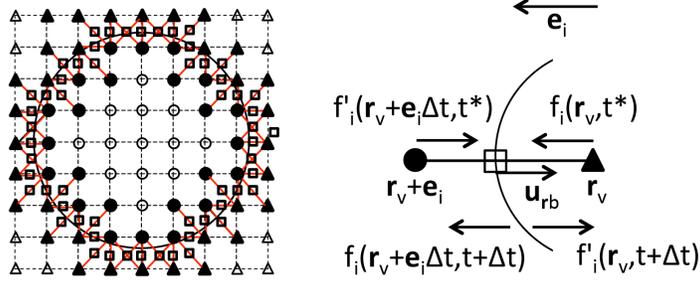


Figure 1: Representation of a particle in the LB model and the momentum exchanges between a particle and the bulk fluid [3, 21, 25]. Filled circles inside the particle represent intra-particle virtual fluid nodes closest to the particle surface, filled triangles outside the particle represent extra-particle bulk fluid nodes, and the unfilled squares represent the boundary nodes located half-way between the intra- and extra-particle nodes.

of mathematically simplified form of two-body Lennard-Jones (LJ) potentials [25]

$$\mathbf{F}_{\mathbf{r}_i} = -\psi \left(\frac{|\mathbf{r}_i|}{|\mathbf{r}_{it}|} \right)^{-13} \mathbf{n}, \quad (6)$$

where $|\mathbf{r}_i|$ is the distance between a particle surface node and the neighboring particle surface node ($\mathbf{r}_i = \mathbf{r}_{pp'}$) or between a particle surface node and the stationary solid node located on channel walls or inline obstacles ($\mathbf{r}_i = \mathbf{r}_{pw}$), p is the particle index, $|\mathbf{r}_{it}|$ is the repulsive threshold distance, \mathbf{n} is the unit vector along \mathbf{r}_i , and ψ is the stiffness parameter used to adjust the repulsive strength between the particles and between the particles and stationary solid zones (Fig. 2). ψ can be tuned to adjust repulsive interaction between neighboring particles or between a particle and stationary solid objects in near contact. The total hydrodynamic force, \mathbf{F}_T , exerted on the particle is

$$\mathbf{F}_T = \sum_{\mathbf{r}_b} \mathbf{F}_{\mathbf{r}_b} + \sum_{\mathbf{r}_b^{c,u}} \mathbf{F}_{\mathbf{r}_b^{c,u}} + \sum_{|\mathbf{r}_{pw}| \leq |\mathbf{r}_{it}|} \mathbf{F}_{\mathbf{r}_{pw}} + \sum_{|\mathbf{r}_{pp'}| \leq |\mathbf{r}_{it}|} \mathbf{F}_{\mathbf{r}_{pp'}}, \quad (7)$$

where $\mathbf{F}_{\mathbf{r}_b^{c,u}} = \pm \rho (\mathbf{u}_{\mathbf{r}_b^{c,u}} - \mathbf{U}_p) / \Delta t$ is the additional force induced by covered,

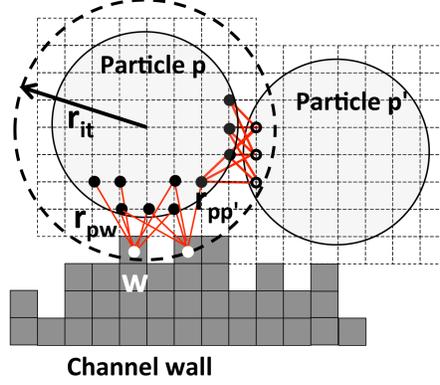


Figure 2: Inter-particle and particle-wall repulsive interactions [25]. Although a few interaction links (in red) are shown for simplicity, interaction potentials are defined over all intra-particle virtual nodes closest the particle surface (shown by black filled circles) and the wall nodes (denoted by w) and/or intra-particle virtual nodes of neighboring particles p' (shown by open circles) within a threshold interaction distance, r_{it} , of a particle p .

\mathbf{r}_b^c , and uncovered, \mathbf{r}_b^u , lattice nodes due to particle motion [4, 21, 23], and \mathbf{U}_p is the particle translational velocity. In Eq. 7, it was assumed that the fluid occupies the entire flow domain to ensure continuity in the flow field to avoid large artificial pressure gradients that may arise from compression and expansion of the fluid near the particle surface. Momentum exchanges occur between the mobile particle and the bulk fluid, following the method of Nguyen and Ladd [24]. The total torque on the particle, \mathbf{T}_T is defined as

$$\begin{aligned} \mathbf{T}_T = & \sum_{\mathbf{r}_b} (\mathbf{r}_b - \mathbf{r}_c) \times \mathbf{F}_{\mathbf{r}_b} + \sum_{\mathbf{r}_b^{c,u}} (\mathbf{r}_b^{c,u} - \mathbf{r}_c) \times \mathbf{F}_{\mathbf{r}_b^{c,u}} + \sum_{|\mathbf{r}_{pw}| \leq |\mathbf{r}_{it}|} (\mathbf{r}_{pw} - \mathbf{r}_c) \times \mathbf{F}_{\mathbf{r}_{pw}} \\ & + \sum_{|\mathbf{r}_{pp'}| \leq |\mathbf{r}_{it}|} (\mathbf{r}_{pp'} - \mathbf{r}_c) \times \mathbf{F}_{\mathbf{r}_{pp'}}. \end{aligned} \quad (8)$$

The translational velocity, \mathbf{U}_p , and the angular velocity of the particle, $\boldsymbol{\Omega}_p$, are advanced in time according to the discretized Newton's equations of motion

$$\mathbf{U}_p(t + \Delta t) \equiv \mathbf{U}_p(t) + \frac{\Delta t}{m_p} \mathbf{F}_T(t) + \frac{\Delta t}{\rho_p} (\rho_p - \rho) \mathbf{g}, \quad (9)$$

$$\boldsymbol{\Omega}_p(t + \Delta t) \equiv \boldsymbol{\Omega}_p(t) + \frac{\Delta t}{I_p} \mathbf{T}_T(t), \quad (10)$$

where m_p is the particle mass and $\mathbf{u}_b = \mathbf{U}_p + \boldsymbol{\Omega}_p \times (\mathbf{r}_b - \mathbf{r}_c)$. The new position of the particle is computed as $\mathbf{r}_c(t + \Delta t) = \mathbf{r}_c(t) + \mathbf{U}_p(t) \Delta t$. The population densities at \mathbf{r}_v and $\mathbf{r}_v + \mathbf{e}_i \Delta t$ are updated to account for momentum-exchange between the particle and bulk fluid in accordance with [3]

$$f'_i(\mathbf{r}_v, t + \Delta t) = f_i(\mathbf{r}_v, t^*) - \frac{2\rho\omega_i}{c_s^2} (\mathbf{u}_{\mathbf{r}_b} \cdot \mathbf{e}_i), \quad (11)$$

$$f_i(\mathbf{r}_v + \mathbf{e}_i \Delta t, t + \Delta t) = f'_i(\mathbf{r}_v + \mathbf{e}_i \Delta t, t^*) + \frac{2\rho\omega_i}{c_s^2} (\mathbf{u}_{\mathbf{r}_b} \cdot \mathbf{e}_i). \quad (12)$$

The microparticle LB model successfully simulated particle trajectories and velocities of micron-size particles in a microfluidic channel with staggered inline obstacles [6]. The model also successfully captured transition in flow migration modes ranging from steady equilibrium with a monotonic approach to strong oscillatory motion in a 2D smooth-walled flow channel with gradual increases in the flow Reynolds number [6, 21], consistent with the previously published numerical simulations based on the finite-element method [26].

4. Fluctuating LB-BGK (FLB-BKG) Module

If particles are of submicron-size, particle Brownian motion could affect particle trajectories. Particle kinetic energy associated with its Brownian motion should be in balance with the thermal energy of the fluid for numerical simulations to obey the fluctuation-dissipation theorem. Random noise (thermal fluctuations) introduced into the fluid randomly perturbs particle velocity as the particle exchanges momentum with the fluid. In the FLB-BGK formulation, thermal fluctuations are added to the post-collisional populations of the fluid [27]

$$f_i^*(\mathbf{r}, t^*) = f_i(\mathbf{r}, t^*) + \Delta t \zeta_i(\mathbf{r}, t), \quad (13)$$

where $f_i(\mathbf{r}, t^*)$ is the post-collisional population density and $\zeta_i(\mathbf{r}, t)$ is space/-time local distribution and acts at the level of the stress tensor and non-hydrodynamic modes, which can be written as [5]

$$\zeta_i(\mathbf{r}, t) = \sqrt{\frac{\rho k_B T \varpi (2 - \varpi)}{c_s^2}} \sum_{k=3}^8 \omega_i \chi_{ki} \Theta_k, \quad (14)$$

where k_B is the Boltzmann constant, T is the temperature, $\varpi = \Delta t/\tau$ for a Newtonian fluid. $\{\chi_{ki}\}_{k=0,8}$ is a set of nine lattice eigenvectors mutually orthonormal according to the scalar products $\sum_{i=0}^8 \omega_i \chi_{ki} \chi_{li} = \delta_{kl}$ and $\sum_{k=0}^8 \omega_i \chi_{ki} \chi_{ki'} = \delta_{ii'}$, and Θ_k is a set of six normal deviates. The eigenvectors for the D2Q9 lattice model correspond to the kinetic moments, in which $k=0$ is the mass density, $k=1-2$ components of the momentum, $k=3-5$ components of the deviatoric momentum flux, and the remaining $k=6-8$ eigenvectors correspond to the ghost modes [27, 28]. These eigenvectors can be expressed as $\chi_{ki} = e_{ki}/\sqrt{w_k}$, in which the vectors e_{ki} and their length, w_k , are listed in Table 1.

Table 1: Basis vectors of the D2Q9 model [29]. w_k is the length of the k th basis vector given by $w_k = \sum_i \omega_i e_{ki}^2$.

k	0	1	2	3	4	5	6	7	8
e_{ki}	1	\hat{e}_{ix}	\hat{e}_{iy}	$3\hat{e}_i^2 - 2$	$2\hat{e}_{ix}^2 - \hat{e}_i^2$	$\hat{e}_{ix}\hat{e}_{iy}$	$(3\hat{e}_i^2 - 4)\hat{e}_{ix}$	$(3\hat{e}_i^2 - 4)\hat{e}_{iy}$	$9\hat{e}_i^4 - 15\hat{e}_i^2 + 2$
w_k	1	1/3	1/3	4	4/9	1/9	2/3	2/3	16

Using the FLB model, we previously demonstrated (i) perfect thermalization of a fluctuating fluid in a confined channel in the presence of a suspended particle near the channel wall under no external force; (ii) a crossover from a ballistic regime to a diffusive regime at which particle velocity autocorrelation vanished while the ratio of the mean-squared displacement of particle positions to the elapsed time approached unity in the diffusive regime; and (iii) the particle obeyed the fluctuation-dissipation theorem (FDT) in a flowing fluid, when it is not in the vicinity of the walls and/or inertial forces are not large [5].

5. Non-Newtonian Fluid Flow Module

The fluid flow in the LB methods described in Sections 2 - 4 is Newtonian. Non-Newtonian fluid flow can be simulated by accommodating the nonlinear relations between the rate of strain and stress rate of the fluid, through which the local fluid kinematic viscosity dynamically evolves in time and space as a function of local velocity gradients.

Numerical modeling of early experimental studies of non-Newtonian fluids in pipe flows measuring the pressure gradient versus bulk fluid velocity for a wide range of fluids [30] suggested a power law relation between fluid viscosity and shear rate, $\dot{\gamma}$. Such power law relations can be built on, for example, via the Ostwald-de Waele model (Eq. 15)

$$\eta = \xi \dot{\gamma}^{n-1}. \quad (15)$$

where η is the dynamic viscosity of the fluid ($\eta = \rho\nu$), ξ is the consistency, n is the fluid-type identifier, $n < 1$, $n = 1$, and $n > 1$ correspond to pseudoplastic (shear-thinning), Newtonian, and dilatant (shear-thickening) fluids, respectively. The shear rate, $\dot{\gamma}$ can be computed from the the second invariant of the rate of strain tensor, Π_D [31]

$$\dot{\gamma} = 2\sqrt{\Pi_D}. \quad (16)$$

$$\Pi_D = \frac{1}{4} \left[\left(\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x} \right) \right]^2 - \frac{\partial u}{\partial x} \frac{\partial v}{\partial y}, \quad (17)$$

Eq. 17 was obtained from $\Pi_D = \frac{1}{2} \left([tr(\mathbf{D})]^2 - tr(\mathbf{D}^2) \right)$, where $\mathbf{D} = \frac{1}{2} (\nabla \mathbf{u}(\mathbf{x}) + (\nabla \mathbf{u}(\mathbf{x}))^T)$, $\mathbf{u} = (u, v)$, $\mathbf{x} = (x, y)$, and T is the transpose operator. From Eqs. 15 and 16 for incompressible fluids ($\rho = 1$ in lattice units),

$$\nu^* = \left[\kappa |\Pi_D|^{\frac{n-1}{2}} \right] \xi. \quad (18)$$

Different values were used for κ in Eq. 18 in the earlier numerical models based on the LB method. For example, $\kappa = 0.5^{n-1}$ in [32], $\kappa = 1$ in [33], and $\kappa = 2^{n-1}$

[16]. Our numerical simulation indicated that these κ values performed equally-well in calculating normalized steady fluid velocities in a smooth-walled flow channel. We adopted $\kappa = 2^{n-1}$ in the simulations used to test the computational performance enhancements in this paper.

In Eq. 16, $\nu^* \rightarrow \infty$ for $n < 1$ and $\nu^* \rightarrow 0$ for $n > 1$ as $\Pi_D \rightarrow 0$, which are both unphysical. To avoid unphysical values for ν^* , the lower and upper bounds for ν^* can be obtained by setting the relaxation parameter in the LBM $\tau \sim 0.5$ (to ensure numerical stability) and $\tau \sim 1.0$ (to ensure enhanced accuracy) [34]. In this formulation, the local kinematic viscosity, ν^* would make the relaxation parameter, τ^* local, through $\tau^* = 0.5 + 3\nu^*(\Delta t/\Delta x^2)$, where ν^* is introduced by Eq. 18. As a result, Eq. 1 for non-Newtonian fluid flow becomes

$$f_i(\mathbf{r} + \mathbf{e}_i \Delta t, t + \Delta t) - f_i(\mathbf{r}, t) = \frac{\Delta t}{\tau^*} [f_i^{eq}(\mathbf{r}, t) - f_i(\mathbf{r}, t)]. \quad (19)$$

and the momentum equation at each lattice node is computed through $\rho \mathbf{u} = f_i \mathbf{e}_i + \tau^* \rho \mathbf{g}$. The generalized analytic solution for the steady-state velocity profile of non-Newtonian fluid flows in a smooth-walled horizontal channel, is given by [32, 35]:

$$u(y) = u_o \left[1 - \left(\frac{2|y|}{W} \right)^{1+\frac{1}{n}} \right], \quad (20)$$

$$u_o = \left(\frac{1}{\xi} |\mathbf{g}| \right)^{\frac{1}{n}} \left(\frac{W}{2} \right)^{1+\frac{1}{n}} \left(\frac{n}{n+1} \right), \quad (21)$$

where W is the channel width and y is vertical distance (perpendicular to the main flow direction) from one of the channel walls. The upper and lower limits for ν^* were set to 10^{-5} and 0.1 to prevent unphysical values for ν^* . A successful validation of the non-Newtonian module against the analytic solution given in Eqs. 20 - 21 is shown in Fig. 3.

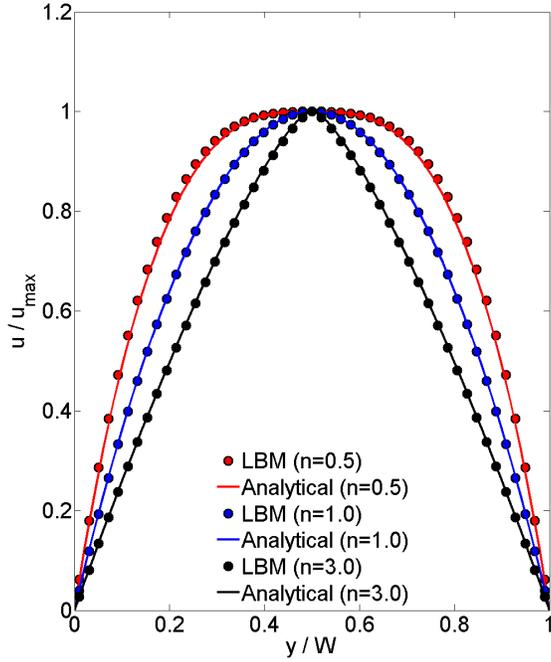


Figure 3: Normalized steady velocity profiles of gravity-driven pseudoplastic ($n = 0.5$), Newtonian ($n = 1$), and dilatant ($n = 3$) fluids in a horizontal, smooth-walled flow channel for $\kappa = 2^{n-1}$. ξ was set to 10^{-3} . Fluid velocities were normalized with respect to the maximum fluid velocity, u_{max} . LBM stands for the lattice-Boltzmann model simulations.

6. Benchmark Problems and Numerical Simulations

Three benchmark problems were set up to measure the computational efficiency of the combined implementation, when applicable, with the optimization techniques proposed in this paper. These benchmark problems involve (i) flow of initially closely-packed multiple particles in a Newtonian fluid flow driven by diagonally-aligned external force in a heterogeneous and disordered fractured-granular porous medium; (ii) nanoparticle flow and thermal fluctuation in the background fluid in a smooth-walled channel, and (iii) non-Newtonian fluid flow in a horizontal channel.

We report all the LB input parameters and computed variables in lattice

units in the subsequent section to facilitate ease of repeatability of the numerical simulations and results. $\Delta x = \Delta t = 1$ was imposed in all simulations discussed in this section. Translation of LB parameters from lattice units to physical dimensions, however, can be found in Refs. [1, 6].

Benchmark 1: Multi-particle Flows in Heterogeneous Porous Media

A numerical simulation was set up to simulate flow trajectories and velocities of cylindrical-circular particles in a spatially-heterogeneous, disordered fractured-porous media (Fig. 4). The fluid flow was assumed to be Newtonian, which was driven by a diagonally-oriented external force with a force strength of $|\mathbf{g}| = (10^{-6}, 10^{-5})$. No-flow condition was imposed along the lateral boundaries and periodic boundaries were implemented at the inlet and outlet. The domain size was set to $1,020 \times 1,620$. The kinematic viscosity of the fluid was set to 0.136, which results in $\tau = 0.636 < 1$ that met the numerical stability requirement of the LB method.

Once the steady-state flow field was obtained within 1%, 45 closely-packed cylindrical-circular particles were released into the flow field near the inlet. The particle radius was set to 6.5 lattice-spacing. Their release locations are shown in Fig. 4, in which surface-to-surface horizontal spacing between neighboring particles was set to 8.5 and the vertical surface-to-surface spacing between the top and bottom arrays of particles was set to 5.5. The LJ repulsive force, Eq. 6, was implemented between particles as well as between a particle and the closest wall with repulsive strength of 20, if the separation distance between them is less than or equal to 2.5. Final positions of the particles are shown in Fig. 5. The simulation ended when one of the particles reached the lower boundary, corresponding to 5.6×10^4 time-step after steady-flow field was established.

Benchmark 2: Submicron-size Particle Flow Simulations

For submicron-size particle simulations via the FLB, a lattice domain of $3,000 \times 41$ was used. External force strength, \mathbf{g} , was set to zero so that thermal fluctuations in the fluid were the only driving force for the particle motion. A

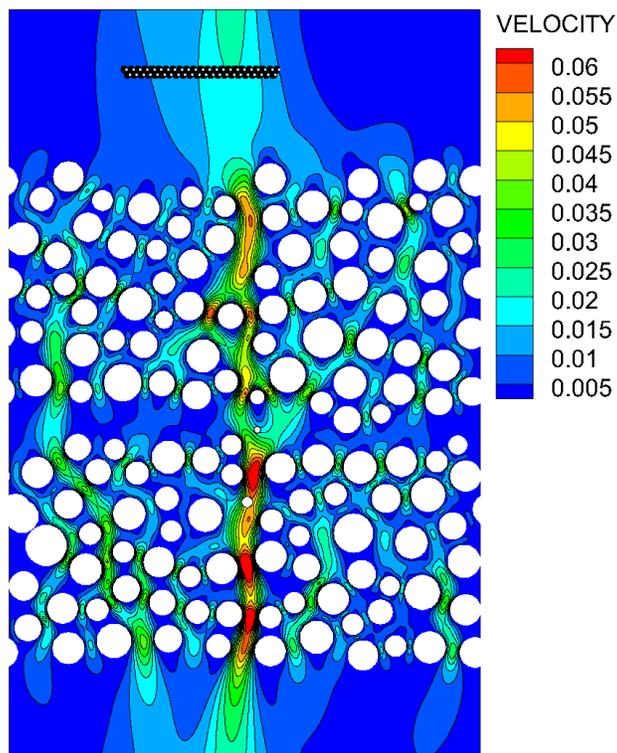


Figure 4: Release location of a total of 45 cylindrical-circular particle into a steady flow field in a heterogeneous, disordered, fractured granular-porous domain. Contours represent the flow velocity field, whose magnitudes are expressed in lattice units. Particles are shown by black-filled circles. Stationary solid grains are shown by white-filled circles.

single particle of radius 7.5 was placed into a horizontal channel and simulated for 8×10^5 time-step. $k_B T$ and ν were set to 5×10^{-5} and 0.1, respectively. Of the 28 available cores, 24 were used as OpenMP threads, and 4 dedicated to random number generation. Simulated trajectories of a Brownian particle released into Newtonian fluid in a confined channel are shown in Fig. 6. In these

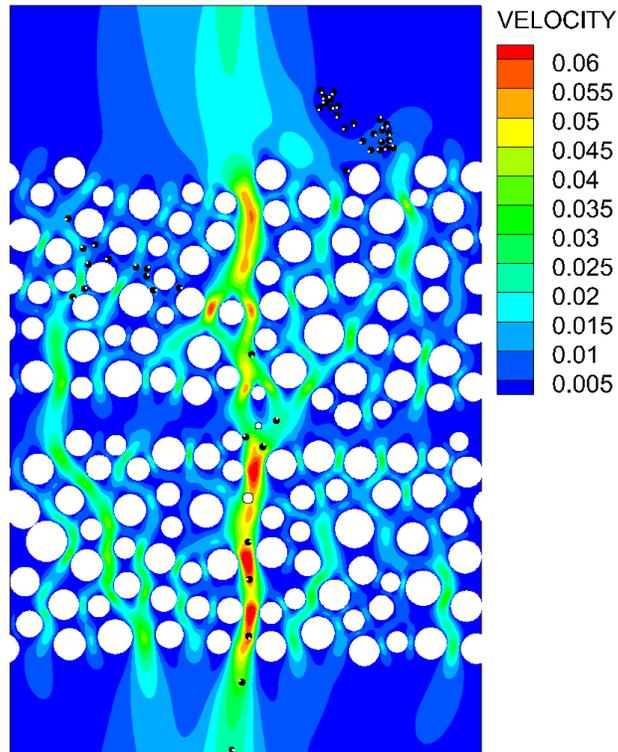


Figure 5: Final position of cylindrical-circular particles. Contours represent the flow velocity field, expressed in lattice units. Particles are shown by black-filled circles. White dots on the particles are used to show the angular rotation of the particles.

simulations, the fluid was assumed to be initially stagnant. Particle motion is due to combined effects of thermal fluctuations-induced particle Brownian motion (Eq. 13), particle-fluid hydrodynamics (Eq. 5), and the wall effects.

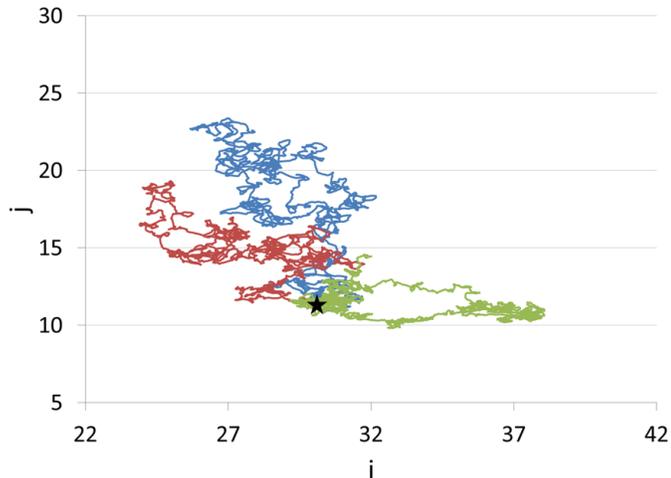


Figure 6: Three realizations of trajectories of a Brownian particle released into an initially stagnant Newtonian fluid in a narrow smooth-walled channel. Star marks the release location of the particle and (i,j) are the lattice indices.

Benchmark 3: Non-Newtonian Flow in a Smooth-Walled Channel

A benchmark problem that calculated the steady-state velocity profile of a pseudoplastic fluid flow with $n = 0.5$ in a smooth-walled horizontal channel (Fig. 3) was used. The domain size was set to 101×51 , and ξ was set to 10^{-3} . 2×10^6 time-step was used to obtain steady-state velocity profile of the non-Newtonian fluid flow.

7. Computational Enhancement Methods

Most computational enhancements to LB models typically focus on algorithmic optimizations, such as what data structures were used to hold the fluid/solid nodes. In this work, we extended this view with a strictly computational perspective to achieve significant performance improvements. We benchmarked our simulations on a 14 physical core/28 virtual core (via hyper-threading) Intel(R)

Xeon(R) CPU E5-2697 v3 2.6 GHz quad-processor machine with 64 GB RAM running Ubuntu 14.04. We used the Intel Fortran compiler v16. The model performance is shown in Table. 2, in which MLUPS (mega lattice-site updates per second) was computed via $MLUPS = (K_x K_y) / (10^6 T)$, where K_i is the number of grid points in the i th direction and T is the time taken to update the entire lattice grid [36].

Table 2: MLUPS for the original and optimized code. Increasing computational performance is shown from left to right as additional optimization techniques were applied. RNG represents MLUPS due to moving random number generation from a single thread to a producer-consumer model. KMP_AFFINITY corresponds to the MLUPS due to nuanced OpenMP thread placement to maximize data locality.

Benchmark	Original	RNG	Serial Optimizations	OpenMP	KMP_AFFINITY
1	2.82	-	6.62	37.27	47.27
2	3.42	5.76	15.00	82.75	138.20
3	5.89	-	12.26	57.10	85.85

The following optimization techniques were used to enhance the computational performance of the model:

Compiler Flags

Through examining the compiler, we found that the default suggested flags for ‘fast’ code (`-O3 -fast`) greatly reduced performance when combined with `-openmp`, because `-fast` sets `-xHost`. For modern machines, this means utilizing the AVX2 vector register set. If these vector register sets are used in any OpenMP loops, then the maximum number of active threads is limited by the number of available sets of AVX registers. We also found that `-O3 -ipo -openmp -fp-model fast=2 -mmodel=medium` produced the greatest performance increases across all simulations. Additional computational improvements were achieved by examining the optimization report provided by the compiler, addressing areas where it was not able to vectorize because of unknown loop bounds (i.e., variable or run-time constant), sub-optimal loop orderings, and

assumed vector/parallel dependencies. Consequently, we made all simulation-constant LB elements, such as the unit velocity vector (\mathbf{e}_i), weights (ω_i), geometric dimensions, and mathematical expressions compile-time constants. This required the use of FORTRAN03 language constructs, because initialization expressions and literal constants are not supported in FORTRAN95. This substantially increased the compiler’s ability to optimize, vectorize, and unroll complex execution flow. Overall we achieved a 2-4 \times serial performance increase through manipulation of compiler invocation and maximization of compile-time determinism, when compared with the un-optimized code (Table 2). While this technique was only applied to 2D simulations, we believe it will provide even greater performance increases when applied to 3D simulations, because the memory requirements for 3D simulations are much greater.

Improved Performance through Data Reduction

Through memory bandwidth analysis, we determined that the large arrays used in our LB code for maintaining (u_x, u_y) in non-Newtonian fluid flow simulations dramatically decreased performance. This was due to the increased cache pressure on the small, fast L1 and L2 CPU caches, resulting in large numbers of cache misses for data read/writes, and consequent increased usage of the slower main memory (RAM) to service those read/write requests [37]. As the u_x/u_y properties of a fluid node are independent of the u_x/u_y values of its neighbors, we were able to eliminate the large ($X \times Y$) array for maintaining this property, and use scalars instead. We achieved an additional performance increase of up to 10 \times from application of this technique due to decreased cache pressure and increased loop fusion/vectorization. However, this array minimization approach cannot be applied unilaterally, as we observed cases in which switching from large arrays to a few scalars minimally increases or even decreases performance. This can occur in loops featuring heavy floating point and vector register usage (such as during the collision step in the LB model), where reserving otherwise available registers to hold scalar subroutine results hinders the

compiler optimization capabilities.

Random Number Generation for FLB-BGK model

The addition of thermal fluctuations to the LB-BGK model for submicron-size particle flow simulations required the generation of large chunks of random numbers each timestep. Our initial approach was to generate the random numbers in a large chunk immediately prior to usage (single thread); this had the side effect of constantly flushing large amounts of hot data from the caches. We addressed this computational bottleneck by shifting our generation strategy to a producer-consumer model, in which a C++ wrapper generates random numbers in large chunks in dedicated threads (separate from simulation thread(s)), and places them into a queue for consumption by the simulation. Each timestep, the simulation would consume one chunk of random numbers from the queue, greatly improving the cache hit/miss ratio. The effect of this technique is shown in Table 2 as part of the optimized parallel execution time for Benchmark 2, where the greatest overall increase in MLUPS was achieved. Table 2 reveals that moving the random number generator to separate threads prior to implementation of other optimization techniques resulted in $1.68\times$ speed-up.

Code Parallelization via OpenMP

We utilized the Intel OpenMP runtime environment for distributing the computational workload across the cores of a system, per standard LB optimization. However, our use of compile-time determinism discussed above enabled the use of the `$omp schedule(static)` directive for loop parallelization, greatly reducing run-time overhead in comparison to `$omp schedule(dynamic)` that is mostly commonly applied. We further optimized the performance of OpenMP threads by disallowing floating between cores by setting

```
KMP_AFFINITY=compact,granularity=fine,1,0.
```

In a dual processor system, for example, threads 0 and 1 (which would likely be operating on adjacent sections of an array) are precluded from being assigned to separate processors, which greatly reduces inter-processor commu-

nication. Such communication is several orders of magnitude slower than intra-processor core communication. Overall, average performance increased by $8\times$ beyond the optimized serial code, as shown in Table 2. In our simulations, the collision/streaming steps constituted the bulk of the computational workload, even after the particles were released; therefore, we decided not to parallelize the particle-particle or particle-wall interactions beyond the auto-parallelization performed by the compiler.

Table 3: Simulation speed-ups.

Benchmark	Serial Speed-up	Parallel Speed-up	Overall Speed-up
1	2.34	7.15	21.77
2	4.38	9.21	40.41
3	2.08	7.0	14.58

Data Structure Arrangements

For 2D simulations, data structure can be organized in memory in either a (i, j, k) or a (k, i, j) format [38], with (i, j) representing the (x,y) coordinates, and k representing the direction. (i, j, k) format optimizes the memory usage patterns in the streaming step, where all execution time is spent on data read/writes, whereas (k, i, j) format optimizes for the memory access patterns in the collision step where the majority of execution time is spent waiting on the results of floating point calculations. This data structure arrangement provided up to a 40% performance increase, and is part of the optimized serial performance shown in Table 2.

The fused streaming-collision method, common in LB optimization for single-phase flow simulations, reduces the required memory bandwidth for a node update in LB simulations, in which post-collision populations computed on the lattice grid layout are immediately streamed to the next neighboring nodes and stored on the second grid layout (i.e., on the second copy) and then these two grid layouts are swapped. Our code, however, does not incorporate the

fused streaming-collision approach common, as our current implementation is not amenable to such an algorithmic change due to post-collision operations prior to streaming the populations (Eqs. 11 - 12 for Benchmark 1 and Eqs. 18 - 19 for Benchmark 3), but future work may address this.

Task-level Pipelining

Building the rank ordering optimization of the main LB data structure, we added an additional dimension (time) to the population distribution rank (i.e. (i, j, k, t)), similar to what is described in [36, 37, 39]. This eliminated the need for a (t) and $(t + 1)$ population distribution structure in our implementation, and enabled a further performance enhancement: efficient task-level parallelism/pipelining (Fig 7), similar to the method in Ref. [40], but within a single computational node. By codifying the data dependence graph below with Intel’s Threading Building Blocks (TBB) library, we were able to achieve parallelism between many aspects of the LB model. For example, pipelining the collision (t) and streaming $(t + 1)$ actions increased performance $1.25\times$ above the OpenMP code, and performing the particle interaction (t) calculations with the streaming $(t + 1)$ operation yielded additional increases of 10-20% for simulations with large numbers of particles. In addition, by offloading simulation reporting from the OpenMP simulation threads (similar to the random number generate approach above), we were able to achieve an additional $2\times$ increase beyond the OpenMP code for simulations with high reporting volumes. However, we also observed 10-20% slowdowns on other simulations when we applied pipelining and/or offloading, demonstrating that the effectiveness of this technique is dependent on simulation parameters. Therefore, fruitful runs utilizing this technique are not included in Table 2.

8. Summary and Conclusions

A combination of different optimization techniques, when applicable, on the computational performance of the LB model for simulating three different problems were investigated. The problems studied involved (i) microparticle flows in

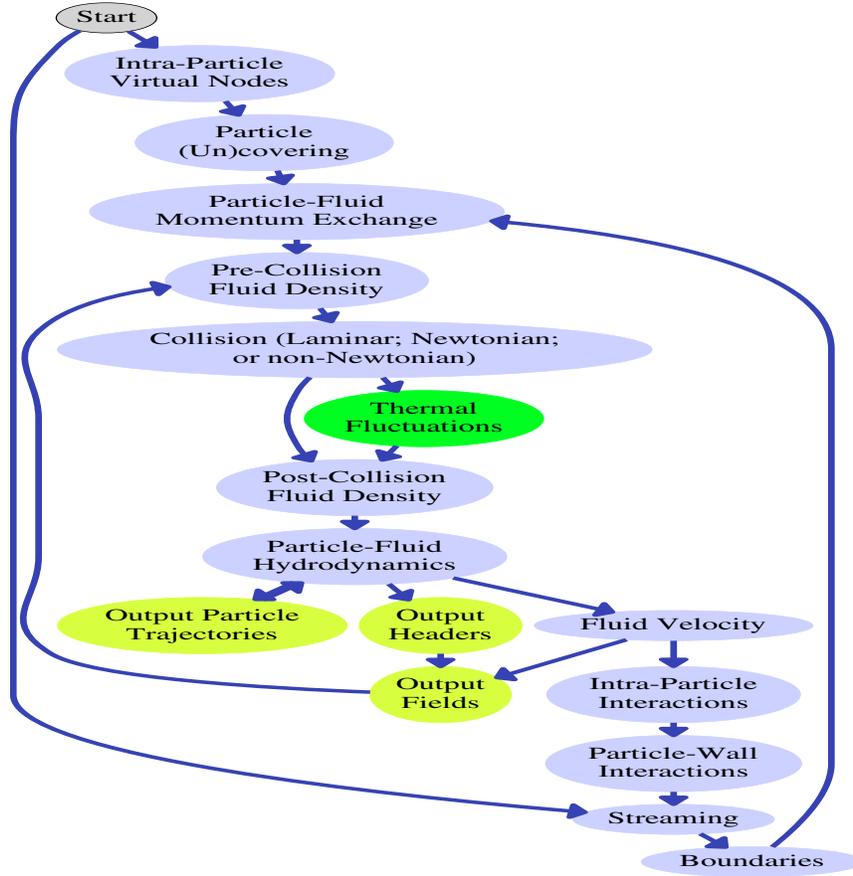


Figure 7: Task level pipelining displaying the flow of data in the LB code. Modules shown in green were only used for some simulations. Modules shown in yellow handle simulation reporting.

a heterogeneous fracture granular-porous domain, (ii) fluid thermalization and the resultant Brownian particle motion important for submicron-scale particle flow simulations, and (iii) pseudoplastic fluid flow in a horizontal channel. In this paper, we applied these techniques to each benchmark problem separately, rather than to a problem involving all the modeling features simultaneously (e.g., simulating the flow of nano-size particle in a non-Newtonian fluid in a geometrically complex domain) to isolate the resultant speed-ups associated with

each module.

Applied techniques include optimization of array memory layouts, data structure simplification, random number generation outside the simulation thread(s), code parallelization via OpenMP, and intra- and inter-timestep task pipelining. As a result of implementation of these optimization techniques, we achieved significant speed-ups in the range of 14.58 (for the third benchmark problem) and 40.41 (for the second benchmark problem), with an average of $21\times$ performance improvement in LB simulations.

When Kuznik *et al.* [41] ran their LB single phase flow simulation on 32 threads (closer to 28 threads were available to us) for a flow domain size of 256×256 , 512×512 , and $1,024\times 1,024$ lattice nodes (comparable to the number of lattice nodes in our Benchmark problems), they accomplished 86-89 MLUPS. Because Kuznik *et al.* simulated only a single-phase Newtonian fluid flow in a flow domain with no internal obstacles (i.e., there were no particles, thermal fluctuations, or non-Newtonian fluid flows), their reported speed-ups should be envisioned as an upper-bound for computational speed-ups on a single GPU. Our optimized LB code with computationally-involved thermal fluctuation calculations (Benchmark 2), however, performed $\sim 1.57\times$ better than a single phase LB simulation on a single GPU. Similarly, our optimized code with non-Newtonian fluid flow simulations (computationally-demanding calculations for local velocity-dependent fluid kinematic viscosity in Benchmark 3) displayed nearly equal ($\sim 1.02\times$ longer) computational performance with the single phase LB flow simulation on a single GPU. Moreover, when 45 particles were released into a flow domain that consists of internal obstacles, our optimized code resulted in a only $\sim 1.82\times$ longer simulation, when compared to the LB simulation on a single GPU for a single-phase particle-free flow in a flow domain with no internal obstacles. Thus, the computationally-demanding LB simulations (advanced modeling features and/or complex flow domain geometries) in our benchmark problems on a CPU performed comparably or better than the single phase LB flow simulations on a single GPU reported in Ref. [41].

In regards to computational improvements through the use of GPUs in LB

simulations of more complex flow problems, GPU/CPU speed-ups ranging from 2 to 12 were reported in LB simulations of multicomponent LB equations with multirange fluid-fluid interactions in flow domains of 128×128 and $1,024 \times 1,024$ lattice nodes [42], which are comparable to the flow domain sizes in our benchmark problems, as reported in Table 3. In addition, the computational performance of 53 MLUPS from a 3D LB simulation of particle-fluid dynamics in an irregular domain using a MPI-based parallelized code on C870 GPU [43] was at the lower end of the computational performances reported in Table 2). Thus, our optimized CPU-based LB code performs comparably or better than previously reported multicomponent (fluid-fluid or fluid-particle) flow simulations on a single GPU. However, nearly an order of magnitude additional speed-up could be possible by changing the GPU cards, for example, moving from Fermi GPU to Pascal GPU without changing the code. Moreover, the authors in Ref. [43] reported 18x further speed-ups with the use of multiple GPUs as they move their simulations from C870 GPU to 8GT200 GPU, demonstrating the important role that state-of-the-art hardware platforms play in achieving maximum performance, and that cutting-edge GPUs can still outstrip cutting-edge CPU performance. This suggests that future work should investigate hybrid CPU-GPU systems in order to achieve maximum throughput and maintain scalability as problem size increases.

The aforementioned significant performance improvements ($14.58\text{-}40.41 \times$) could enhance the use of the LB method in simulating natural or targeted flows of nano-size and micron-size biotic (e.g., bacteria) or abiotic (engineered vectors or colloids) particles in geometrically complex flow domains. Applications may involve, for example, engineered targeted drug delivery to targeted tumor cells for cancer treatment, or health and safety assessments that utilize fate and transport of radionuclides bound to colloids in the vicinity of nuclear power plants or nuclear waste disposal sites.

In such cases, the computational performance and effectiveness of the aforementioned optimization techniques should be assessed on 3D simulations. The exponentially greater computational and memory requirements for such simu-

lations means that the application of the techniques discussed in this paper is not straightforward, and will likely require modification in order to be successfully applied. For example, use of multi-processor system may become more feasible as memory requirements increase, necessitating revisions to our current OpenMP model. Furthermore, to target large 3D problems, multiple computational nodes will have to be utilized, requiring investigation into techniques of distributed memory sharing and task allocation on an inter-node level.

Acknowledgement

The work of H. Başağaoğlu was supported by Internal Research Project 18-R8305. The work of J. Harwell was supported in part by Internal Research Project 10-R8629. The authors thank Phil M. Westhart for his help with simulations and Miriam R. Juckett for reviewing the manuscript.

References

- [1] S. Succi, *The Lattice-Boltzmann Equation for Fluid Dynamics and Beyond*, Oxford University Press, New York, 2001.
- [2] D. A. Wolf-Gladrow, *A Lattice Gas Cellular Automata and Lattice Boltzmann Model*, Springer-Verlag, Berlin Heidelberg, 2000.
- [3] 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.
- [4] E-J Ding, C.K. Aidun, Extension of the lattice-Boltzmann method for direct simulation of suspended particles near contact, *J Stat Phys* 112 (2003) 685–708.
- [5] H. Başağaoğlu, S. Melchionna, S. Succi, V. Yakhot, Fluctuation-dissipation relations from a FLB-BGK model, *Europhys Lett* 99 (2012) 64001.

- [6] H. Başağaoğlu, S. Allwein, S. Succi, H. Dixon, J.T. Carrola Jr., S. Stothoff, Two- and three-dimensional lattice-Boltzmann simulations of particle migration in microchannels, *Microfluid Nanofluid* 15 (6) (2013) 785–796.
- [7] X. Shan, H. Chen, Lattice-Boltzmann model for simulation flows with multiple phases and components, *Phys Rev E* 47 (1993) 1815–1819.
- [8] L. Amaya-Bower, T. Lee, Numerical simulation of single bubble rising in vertical and inclined square channel using lattice Boltzmann method, *Chem Eng Sci* 66 (2011) 935–952.
- [9] J. Bao, L. Schaefer, Lattice Boltzmann equation model for multi-component multi-phase flow with high density ratios, *Appl Math Model* 37 (2013) 1860–1871.
- [10] M.R. Kamali, J.J.J. Gillissen, H.E. Van den Akker, S. Sundaresan, Lattice-Boltzmann-based two-phase thermal model for simulating phase change, *Phys Rev E* 88 (2013) 033302.
- [11] H. Chen, S. Kandasany, S. Orszag, R. Shock, S. Succi, V. Yakhot, Extended Boltzmann kinetic equation for turbulent flow, *Nature* 301 (2003) 633–636.
- [12] H. Başağaoğlu, P. Meakin, S. Succi, Energy dissipation measures in three-dimensional porous media, *Phys Rev E* 72 (2005) 046705.
- [13] C. Bartlett, H. Chen, I. Staroselsky, J. Wanderer, V. Yakhot, Lattice Boltzmann two-equation model for turbulence simulations: High-reynolds number flow past circular cylinder, *Int J Heat Fluid Flow* 42 (2013) 1–9.
- [14] J. Onishi, Y. Chen, H. Ohashi, Dynamic simulation of multi-component viscoelastic fluids using the lattice-Boltzmann method, *Phys A* 362 (2006) 84–92.
- [15] O. Malaspinas, N. Fiétier, M. Deville, Lattice Boltzmann method for the simulation of viscoelastic fluid flows, *J Non-Newtonian Fluid Mech* 165 (2010) 1637–1653.

- [16] A.A. Delouei, M. Nazari, M.H. Kayhani, S. Succi, Non-Newtonian unconfined flow and heat transfer over a heated cylinder using the direct-forcing immersed boundary-thermal lattice Boltzmann method, *Phys Rev E* 89 (2014) 053312.
- [17] R. Benzi, S. Succi, M. Vergassola, The lattice-Boltzmann equation: Theory and applications, *Phys Rep* 222 (1992) 145–197.
- [18] P.L. Bhatnagar, E.P. Gross, M. Krook, A model for collision process in gases. I. Small amplitude processes in charged and neutral one-component systems, *Phys Rev* 94 (3) (1954) 511–525.
- [19] Y.H. Qian, D. D’Humières, P. Lallemand, Lattice BGK models for Navier-Stokes equation, *Europhys Lett* 17 (1992) 479–484.
- [20] J.M. Buick, C. A. Greated, Gravity in a lattice Boltzmann model, *Phys Rev E* 61 (2000) 5307–5320.
- [21] H. Başağaoğlu, P. Meakin, S. Succi, G.R. Redden, T.R. Ginn, Two-dimensional lattice-Boltzmann simulation of colloid migration in rough-walled narrow flow channels, *Phys Rev E* 77 (2008) 031405.
- [22] H. Başağaoğlu, J.T. Carrola Jr., C.J. Freitas, B. Başağaoğlu, S. Succi, Lattice Boltzmann simulations of vortex entrapment of particles in a microchannel with curved and flat edges, *Microfluid Nanofluid* 18 (5–6) (2015) 1165–1175.
- [23] C.K. Aidun, Y. Lu, E-J Ding, Direct analysis of particulate suspensions with inertia using the discrete Boltzmann equation, *J Fluid Mech* 373 (1998) 287–311.
- [24] N-Q Nguyen, A.J.C. Ladd, Lubrication corrections for lattice-Boltzmann simulations of particle suspensions, *Phys Rev E* 66 (2002) 046708.
- [25] H. Başağaoğlu, S. Succi, Lattice-Boltzmann simulations of repulsive particle-particle and particle-wall interactions: Coughing and choking, *J Chem Phys* 132 (5–6) (2010) 134111.

- [26] J. Feng, H. Hu, D.D. Joseph, Direct simulation of initial value problems for the motion of solid bodies in a Newtonian fluid part 1. Sedimentation, *J Fluid Mech* 261 (95–134) 046705.
- [27] R. Adhikari, K. Stratford, J.E. Cates, A.J. Wagner, Fluctuating lattice Boltzmann, *Europhys Lett* 71 (2005) 473–410.
- [28] R. Adhikari, S. Succi, Duality in matrix lattice Boltzmann models, *Phys Rev E* 78 (2005) 066701.
- [29] D.B. Dünweg, U.D. Schiller, A.J.C. Ladd, Statistical mechanics of the fluctuating lattice Boltzmann equation, *Phys Rev E* 76 (2007) 036704.
- [30] D. W. Dodge, A. B. Metzner, Turbulent flow of non-Newtonian systems, *AIChEJ* 5 (1959) 189–204.
- [31] A. Nejat, V. Abdollahi, J. Vahidkhah, Lattice Boltzmann simulation of non-Newtonian flows past confined cylinders, *J Fluid Mech* 166 (12) (2011) 689–697.
- [32] J. Psihogios, M.E. Kainourgiakis, A.G. Yiotis, A. Th. Papaioannous, A. K. Stubos, Lattice Boltzmann of non-Newtonian flow in digitally reconstructed porous domain, *Transp Porous Med* 70 (2007) 279–292.
- [33] H. Hamed, M.H. Rahimian, Numerical simulation of non-Newtonian pseudo-plastic fluid in a micro-channel using the lattice-Boltzmann method, *World J Mech* 1 (2011) 231–242.
- [34] S. Gabbanelli, G. Drazer, J. Koplik, Lattice Boltzmann method for non-Newtonian (power-law) fluids, *Phys Rev E* 71 (2005) 046312.
- [35] S. Whitaker, *Introduction to Fluid Mechanics*, Krieger Pub Com, Florida, 1968.
- [36] A.G. Shet, S. Sorathiya, S. Krithivasan, A. Deshpande, B. Kaul, S. Sherlekar, S. Ansumali, Data structure and movement for lattice-based simulations, *Phys Rev E* 88 (2013) 013314.

- [37] F. Massaioli, G. Amati, The Fourth European Workshop on OpenMP, in: EWOMP'02, 2002.
- [38] T. Pohl, M. Kowarschik, J. Wilke, K. Iglberger, U. Rüde, Optimization and profiling of the cache performance of parallel lattice-Boltzmann codes, *Parallel Process Lett* 13 (2003) 549–560.
- [39] L. Peng, K. Nomura, T. Oyakawa, R. Kalia, A. Nakano, P. Vashishta, Parallel lattice Boltzmann flow simulation on emerging multi-core platforms, *LNCS* 5168 (2008) 763–777.
- [40] D. Mudigere, S. Sherlekar, S. Ansumali, Delayed difference scheme for large scale scientific simulations, *Phys Rev Lett* 113 (2014) 218701.
- [41] F. Kuznik, C. Obrecht, G. Rusaouen, J.-J. Roux, LBM based flow simulation using GPU computing processor, *Comp Math Appl* 59 (2010) 2380–2392.
- [42] M. Bernaschi, L. Rossi, R. Benzi, M. Sbragaglia, S. Succi, Graphic processing unit implementation of lattice Boltzmann models for flowing soft systems, *Phys Rev E* 80 (2010) 066707.
- [43] M. Bernaschi, M. Massimiliano, S. Melchionna, S. Succi, E. Kaxiaras, A flexible high-performance lattice Boltzmann GPU code for the simulations of fluid flows in complex geometries, *Concurrency Computat: Pract Exper* (2009) DOI: 10.1002/cpe.1466.