



# Particle-resolved thermal lattice Boltzmann simulation using OpenACC on multi-GPUs

Ao Xu<sup>a,b,c,\*</sup>, Bo-Tao Li<sup>a</sup>

<sup>a</sup> School of Aeronautics, Northwestern Polytechnical University, Xi'an 710072, China

<sup>b</sup> Institute of Extreme Mechanics, Northwestern Polytechnical University, Xi'an 710072, China

<sup>c</sup> Key Laboratory of Icing and Anti/De-icing, China Aerodynamics Research and Development Center, Mianyang 621000, China

## ARTICLE INFO

### Keywords:

Particle-laden flow  
Thermal convection  
Lattice Boltzmann method  
GPU computing  
OpenACC

## ABSTRACT

We utilize the Open Accelerator (OpenACC) approach for graphics processing unit (GPU) accelerated particle-resolved thermal lattice Boltzmann (LB) simulation. We adopt the momentum-exchange method to calculate fluid-particle interactions to preserve the simplicity of the LB method. To address load imbalance issues, we extend the indirect addressing method to collect fluid-particle link information at each timestep and store indices of fluid-particle link in a fixed index array. We simulate the sedimentation of 4,800 hot particles in cold fluids with a domain size of  $4000^2$ , and the simulation achieves 1750 million lattice updates per second (MLUPS) on a single GPU. Furthermore, we implement a hybrid OpenACC and message passing interface (MPI) approach for multi-GPU accelerated simulation. This approach incorporates four optimization strategies, including building domain lists, utilizing request-answer communication, overlapping communications with computations, and executing computation tasks concurrently. By reducing data communication between GPUs, hiding communication latency through overlapping computation, and increasing the utilization of GPU resources, we achieve improved performance, reaching 10846 MLUPS using 8 GPUs. Our results demonstrate that the OpenACC-based GPU acceleration is promising for particle-resolved thermal lattice Boltzmann simulation.

## 1. Introduction

Particle-laden thermal convection occurs ubiquitously in nature and daily life [1,2]. For example, under the action of a wind field, sand particles of different sizes follow the airflow with suspended or leaping motion, and these particles may form dust storms by unstable thermal conditions during strong winds [3]. Another example is that under the influence of air conditioning and the ventilation system, atmospheric pollutant particles (PM10 and PM2.5) originating from dust and smoke can be suspended in the air for an extended period, and they may spread to a wider area with the aid of airflow [4].

Numerical models for simulating particle-laden flows can be generally classified into two categories [5]: the point-particle model and the particle-resolved model. In the point-particle model, solid particles are treated as discrete masses that are much smaller than the mesh size of the computational grid. Empirical correlations, such as Stokes' viscous drag law, are used to calculate forces exerted on the solid particles by the fluid. The advantage of the point-particle model is its relatively low computational load, making it suitable for tracking large numbers of

particles (but with limited particle volume fractions) [6–9]. In contrast, in the particle-resolved model, particle sizes are larger than the resolution of the computational grid (known as finite-size particles). They ensure the no-slip velocity boundary condition at the particle surface, and the forces and moments acting on the particles by the fluid are explicitly calculated by considering the interaction between the fluid phase and the solid phase. The advantage of the particle-resolved model is its ability to accurately describe fluid-particle interactions based on first principles. Popular particle-resolved model includes the arbitrary Lagrangian-Eulerian method [10], immersed boundary method [11], fictitious domain method [12,13], and lattice Boltzmann (LB) method [14]. The particle-resolve model can simulate dense suspension [15–18]; however, their high computational cost limits the number of particles that can be tracked.

Among these particle-resolved methods, the LB method is fascinating due to its ability to incorporate mesoscopic physical pictures while recovering macroscopic physical conservation laws with a relatively low computational cost. Open-source codes based on the LB method, including OpenLB [19], Palabos [20], and Sailfish [21], have facili-

\* Corresponding author at: School of Aeronautics, Northwestern Polytechnical University, Xi'an 710072, China.  
E-mail address: [axu@nwpu.edu.cn](mailto:axu@nwpu.edu.cn) (A. Xu).

tated the application of LB simulations in large-scale engineering problems [22]. Over the past decades, the advancement of general-purpose graphics processing units (GPUs) has significantly improved high-performance computing, enabling faster simulations of larger physical domains or higher computational resolutions [23,24]. Parallel computing frameworks utilizing GPU architectures can use Open Computing Language (OpenCL), Compute Unified Device Architecture (CUDA), and Open Accelerators (OpenACC) [25,26]. A detailed comparison of these programming standards can be found in our previous work [27,28]. Here, we highlight that, with improved data and task management, the OpenACC is promising for thermal LB simulation on GPUs [29–31]. As demonstrated in the most recent simulation of fluid flows and heat transfer in the side-heated convection cell [28], using OpenACC on a single GPU, the two-dimension (2D) simulation achieved 1.93 billion lattice updates per second (GLUPS) with a grid number of  $8193^2$ , and the three-dimension (3D) simulation achieved 1.04 GLUPS with a grid number of  $385^3$ , which is more than 76% of the theoretical maximum performance.

In this paper, we aim to extend our previous OpenACC accelerated LB simulation of single-phase thermal convection [27,28] to particle-laden thermal convection. In contrast to previous works [32,33] who adopted the immersed boundary method to evaluate the fluid-particle interaction force, we adopt the momentum-exchange method due to its simplicity and robustness [34–36]. To utilize the computing power of multi-node GPU clusters, we adopt a hybrid OpenACC and Message Passing Interface (MPI) approach, in which the OpenACC accelerates computation on a single GPU, and the MPI synchronizes the information between multiple GPUs. The rest of this paper is organized as follows. In Section 2, we introduce numerical details of the LB simulation for particle-laden thermal convection. In Section 3, we describe the implementation and optimization details for a single GPU. In Section 4, we describe the implementation and optimization for multi-GPUs. In Section 5, the main findings of this work are summarized.

## 2. Numerical method

### 2.1. The LB model for fluid flow and heat transfer

We simulate thermal convection based on the Boussinesq approximation. We assume the fluid flow is incompressible, and we treat the temperature as an active scalar that influences the velocity field through the buoyancy. We neglect viscous heat dissipation and compression work, and assume all the transport coefficients to be constants. Then, the governing equations can be written as

$$\nabla \cdot \mathbf{u} = 0 \quad (1a)$$

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\frac{1}{\rho_0} \nabla P + \nu \nabla^2 \mathbf{u} + g \beta_T (T - T_0) \hat{\mathbf{y}} \quad (1b)$$

$$\frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T = \alpha_T \nabla^2 T \quad (1c)$$

where  $\mathbf{u}$ ,  $P$  and  $T$  are the velocity, pressure, and temperature of the fluid, respectively.  $\rho_0$  and  $T_0$  are reference density and temperature, respectively.  $\nu$ ,  $\beta_T$  and  $\alpha_T$  denote the viscosity, thermal expansion coefficient, and thermal diffusivity of the fluid, respectively.  $\hat{\mathbf{y}}$  is the unit vector parallel to gravity. With the scaling

$$\mathbf{x}^* = \mathbf{x}/L_0, \quad t^* = t/(L_0^2/\alpha_T), \quad \mathbf{u}^* = \mathbf{u}/(\alpha_T/L_0), \quad (2)$$

$$P^* = P/(\rho_0 \alpha_T^2 / L_0^2), \quad T^* = (T - T_0)/\Delta_T$$

Then, Eq. (1) can be rewritten in dimensionless form as

$$\nabla \cdot \mathbf{u}^* = 0 \quad (3a)$$

$$\frac{\partial \mathbf{u}^*}{\partial t^*} + \mathbf{u}^* \cdot \nabla \mathbf{u}^* = -\nabla P^* + Pr \nabla^2 \mathbf{u}^* + Gr Pr^2 T^* \hat{\mathbf{y}} \quad (3b)$$

$$\frac{\partial T^*}{\partial t^*} + \mathbf{u}^* \cdot \nabla T^* = \nabla^2 T^* \quad (3c)$$

Here,  $L_0$  is the characteristic length and  $\Delta_T$  is the temperature difference. Two independent dimensionless parameters are the Prandtl number ( $Pr$ ) and the Grashof number ( $Gr$ ), which are defined as

$$Pr = \frac{\nu}{\alpha_T}, \quad Gr = \frac{g \beta_T \Delta_T L_0^3}{\nu^2} \quad (4)$$

Note a third dimensionless parameter of the Rayleigh number ( $Ra$ ) can be calculated as  $Ra = Pr \cdot Gr$ .

We adopt the double distribution function (DDF)-based LB model to simulate thermal convective flows with the Boussinesq approximation [37–40]. Specifically, we chose a D2Q9 discrete lattice in two-dimension (2D) for the Navier–Stokes equations to simulate fluid flows, and a D2Q5 discrete lattice in 2D for the energy equation to simulate heat transfer. To enhance the numerical stability, the multi-relaxation-time (MRT) collision operator is adopted in the evolution equations of both density and temperature distribution functions. The evolution equation of the density distribution function is written as

$$f_i(\mathbf{x} + \mathbf{e}_i \delta_t, t + \delta_t) - f_i(\mathbf{x}, t) = -(\mathbf{M}^{-1} \mathbf{S})_{ij} \left[ \mathbf{m}_j(\mathbf{x}, t) - \mathbf{m}_j^{(eq)}(\mathbf{x}, t) \right] + \delta_t F_i' \quad (5)$$

where  $f_i$  is the density distribution function.  $\mathbf{x}$  is the fluid parcel position,  $t$  is the time,  $\delta_t$  is the time step.  $\mathbf{e}_i$  is the discrete velocity along the  $i$ th direction. The macroscopic density  $\rho$  and velocity  $\mathbf{u}$  are obtained from  $\rho = \sum_{i=0}^{q-1} f_i$ ,  $\mathbf{u} = \left( \sum_{i=0}^{q-1} \mathbf{e}_i f_i + \frac{1}{2} \mathbf{F} \right) / \rho$ , where  $\mathbf{F} = \rho g \beta_T (T - T_0) \hat{\mathbf{y}}$  is the buoyancy force.

The evolution equation of the temperature distribution function is written as

$$g_i(\mathbf{x} + \mathbf{e}_i \delta_t, t + \delta_t) - g_i(\mathbf{x}, t) = -(\mathbf{N}^{-1} \mathbf{Q})_{ij} \left[ \mathbf{n}_j(\mathbf{x}, t) - \mathbf{n}_j^{(eq)}(\mathbf{x}, t) \right] \quad (6)$$

where  $g_i$  is the temperature distribution function. The macroscopic temperature  $T$  is obtained from  $T = \sum_{i=0}^{q-1} g_i$ . More numerical details of the thermal LB method can be found in our previous work [27,28,41–43].

### 2.2. Kinematic model of the solid particle

We consider the solid particle as a rigid body and its kinematics include translational and rotational motion. Specifically, we determine the translational motion of the solid particle using Newton's second law as

$$M_p \frac{d\mathbf{U}_c(t)}{dt} = \mathbf{F}_p(t) \quad (7)$$

where  $M_p$  is the mass of the particle,  $\mathbf{U}_c$  is the velocity of the particle center and  $\mathbf{F}_p$  is the total force exerted on the solid particle. The rotational motion of the solid particle is determined by Euler's second law as

$$\mathbf{I}_p \cdot \frac{d\boldsymbol{\Omega}(t)}{dt} + \boldsymbol{\Omega}(t) \times [\mathbf{I}_p \cdot \boldsymbol{\Omega}(t)] = \mathbf{T}_p(t) \quad (8)$$

where  $\mathbf{I}_p$  is the inertial tensor of the particle,  $\boldsymbol{\Omega}$  is the angular velocity, and  $\mathbf{T}_p$  is the torque exerted on the solid particle. In the simulation, we advance the fluid flows and the motion of the particles simultaneously. In other words, the time step used to update fluid and temperature fields is the same as that used for particles' kinematic.

### 2.3. Boundary conditions at the fluid-particle interface

At the particle's curved surface, we adopt the interpolated bounce-back scheme to guarantee no-slip velocity boundary conditions. A parameter  $q = |\mathbf{x}_f - \mathbf{x}_w|/|\mathbf{x}_f - \mathbf{x}_b|$  is used to describe the fraction of fluid region in a grid spacing intersected by the solid surface, where  $\mathbf{x}_f$  is the fluid node near the boundary,  $\mathbf{x}_b$  is solid node near the boundary, and  $\mathbf{x}_w$  is wall interface. Based on the relative location of  $\mathbf{x}_w$  between  $\mathbf{x}_f$  and  $\mathbf{x}_b$ , we adopt a quadratic interpolation scheme of the density distribution function, which is given as [44]: for  $q \leq 0.5$

$$f_i(\mathbf{x}_f, t + \delta_t) = q(2q+1)f_i^+(\mathbf{x}_f, t) + (1-4q^2)f_i^+(\mathbf{x}_f - \mathbf{e}_i\delta_t, t) - q(1-2q)f_i^+(\mathbf{x}_f - 2\mathbf{e}_i\delta_t, t) + 2\omega_i\rho_0 \frac{\mathbf{e}_i \cdot \mathbf{u}_w}{c_s^2} \quad (9)$$

for  $q > 0.5$

$$f_i(\mathbf{x}_f, t + \delta_t) = \frac{1}{q(2q+1)}f_i^+(\mathbf{x}_f, t) + \frac{2q-1}{q}f_i^+(\mathbf{x}_f, t) - \frac{2q-1}{2q+1}f_i^+(\mathbf{x}_f - \mathbf{e}_i\delta_t, t) + \frac{1}{q(2q+1)}2\omega_i\rho_0 \frac{\mathbf{e}_i \cdot \mathbf{u}_w}{c_s^2} \quad (10)$$

where  $f_i$  is the distribution function associated with the velocity  $\mathbf{e}_i = -\mathbf{e}_i$ . It should be noted that Bouzidi's method requires information from the current fluid nodes as well as its adjacent nodes (i.e., information at  $\mathbf{x}_f$ ,  $\mathbf{x}_f - \mathbf{e}_i\delta_t$ , and  $\mathbf{x}_f - 2\mathbf{e}_i\delta_t$ ), which poses a challenge to the local computation property of the LB method. To address this issue, an alternative method is a single-node second-order curved boundary condition [45,46]. The accuracy, stability, and parallel efficiency of those single-node methods for particle-resolved simulation deserve further comprehensive investigation.

Meanwhile, we assume the fluids and particle temperatures are equal to a constant  $T_w$  at the surface of the particle. Then the bounce-back scheme for temperature distribution function at curved wall boundaries is given as [47,48]

$$g_i(\mathbf{x}_f, t + \delta_t) = \left[ c_{d1}g_i^+(\mathbf{x}_f, t) + c_{d2}g_i^+(\mathbf{x}_f - \mathbf{e}_i\delta_t, t) + c_{d3}g_i^+(\mathbf{x}_f, t) \right] + c_{d4}(2\omega_i T_w) \quad (11)$$

where  $g_i$  is the distribution function associated with the velocity  $\mathbf{e}_i = -\mathbf{e}_i$ , and  $g_i^+$  is the post-collision distribution function. The coefficients  $c_{d,1-4}$  are given as

$$c_{d1} = -1, \quad c_{d2} = \frac{2q-1}{2q+1}, \quad c_{d3} = \frac{2q-1}{2q+1}, \quad c_{d4} = \frac{2}{2q+1} \quad (12)$$

#### 2.4. Interaction between fluid and particle phases

We adopt the momentum-exchange method to calculate the force and torque exerted by the fluid on the solid particle due to its simplicity and robustness [34]. Specifically, the hydrodynamic force acting on the solid surface is obtained by summing up the local momentum exchange of the fluid parcels during the bounce-back process. Because the original momentum-exchange method proposed by Ladd [34] lacks local Galilean invariance [49], we employ a modified momentum-exchange method that simply introduces the relative velocity into the interfacial momentum transfer [36], then the total hydrodynamic force is calculated as

$$\mathbf{F} = \sum_{\mathbf{x}_f} \sum_{i_{bl}} [f_i^+(\mathbf{x}_f, t)(\mathbf{e}_i - \mathbf{u}_w) - f_i(\mathbf{x}_f, t + \delta_t)(\mathbf{e}_i - \mathbf{u}_w)] \quad (13)$$

and the total torque is calculated as

$$\mathbf{T} = \sum_{\mathbf{x}_f} \sum_{i_{bl}} (\mathbf{x}_{iw} - \mathbf{x}_c) \times [f_i^+(\mathbf{x}_f, t)(\mathbf{e}_i - \mathbf{u}_w) - f_i(\mathbf{x}_f, t + \delta_t)(\mathbf{e}_i - \mathbf{u}_w)] \quad (14)$$

A pair of particles will collide when their distance is small. We adopt the artificial repulsive force model to prevent overlap between the computational particles. Here, we choose the spring force model that generates a strong repulsive force pushing the two collision particles apart [50], then the repulsive force is given as

$$\mathbf{F}_{\text{repulsive}} = \begin{cases} 0 & \text{if } |\mathbf{x}_s| > s, \\ \frac{C}{\varepsilon_w} \left( \frac{|\mathbf{x}_s| - s}{s} \right)^2 \frac{\mathbf{x}_s}{|\mathbf{x}_s|} & \text{if } |\mathbf{x}_s| \leq s. \end{cases} \quad (15)$$

Here,  $\varepsilon_w$  represents the stiffness parameter and  $s$  represents the threshold distance. We choose  $\varepsilon_w = 0.001$  and  $s = 3$  l.u., where l.u. denotes the lattice length-unit in the LB simulation [51].  $C = \pi r_p^2(\rho_p - \rho_f)g$  is the force scale,  $\mathbf{x}_s$  denotes the vector with the smallest norm value that points from the wall to the particle.

#### 2.5. Refilling scheme to construct unknown distribution functions

When the particle moves relative to the fixed grids, a solid node may be uncovered by the particle and become a fluid node. To determine unknown density distribution functions for this new 'born' fluid node, we use the normal extrapolation refilling scheme with velocity constraint [52,53]. Specifically, we first determine the direction of a discrete velocity  $\mathbf{e}_c$  that maximizes the quantity  $\mathbf{n} \cdot \mathbf{e}_c$ , where  $\mathbf{n}$  is the outward normal vector of the wall at the newborn fluid node. Then the unknown density distribution functions at the newborn fluid node are obtained by a quadratic extrapolation

$$f_\alpha(\mathbf{x}_{new}, t + \delta_t) = 3f_\alpha(\mathbf{x}_{new} + \mathbf{e}_c\delta_t, t + \delta_t) - 3f_\alpha(\mathbf{x}_{new} + 2\mathbf{e}_c\delta_t, t + \delta_t) + f_\alpha(\mathbf{x}_{new} + 3\mathbf{e}_c\delta_t, t + \delta_t) \quad (16)$$

In the MRT framework, the velocity at the new fluid node can be constrained to the wall velocity via enforcing the momentum moments, which can reduce the fluctuations in the fluid-particle forces [53].

Meanwhile, to determine unknown temperature distribution functions for the new 'born' fluid node, we use the equilibrium refilling scheme [52,54] for simplify, which is

$$g_\alpha(\mathbf{x}_{new}, t + \delta_t) = g_\alpha^{eq}(T_w, \mathbf{u}_w) \quad (17)$$

#### 2.6. Validation of the particle-resolved LB model

The accuracy of the particle-resolved LB model described above has been validated in our previous work [55] via simulation of an elliptical particle settling in isothermal fluids, and a cold circular particle settling in a hot fluid. Furthermore, our results on an elliptical particle settling in thermal fluids have been confirmed by Walayat et al. [56] as well as Suzuki et al. [17]. Additionally, we compare the simulation results obtained from the multi-GPU simulation with those of the CPU simulation in Appendix A. Moreover, we validate the model for particle-particle interaction via simulating the draft-kissing-tumbling (DKT) process with thermal convection in Appendix B. In the following subsections, we will validate the model by simulating the sedimentation of multiple particles.


##### 2.6.1. Sedimentation of a group of 800 circular particles in isothermal fluids

Initially, the particles are placed in the upper part of a square cavity with dimensions of 5 cm  $\times$  5 cm. All four walls of the cavity are imposed with no-slip velocity boundary conditions. There are 20 lines of particles, each containing 40 particles, resulting in a total of 800 particles [see Fig. 1(a)]. Each particle has a density of  $\rho_p = 1.1$  g/cm<sup>3</sup> and a diameter of  $d_p = 1$  mm, resulting in a particle volume fraction of 25.1%. To mimic the working fluid of water, we set its viscosity as  $\nu_f = 10^{-6}$  m<sup>2</sup>/s and its density as  $\rho_f = 1$  g/cm<sup>3</sup>. Following Yu et al. [57], we adopt a reference velocity  $U_{ref} = \sqrt{g\pi d_p(\rho_p - \rho_f)/(2\rho_f)}$  to define the particle Reynolds number  $Re_p = U_{ref}d_p/\nu_f$ , which is related to the Archimedes number  $Ar = \sqrt{gd_p^3(\rho_p - \rho_f)/(\nu_f^2\rho_f)}$  as  $Re_p = \sqrt{\pi/2}Ar$ . In this case, we have  $Re_p = 12.41$  and  $Ar = 9.90$ . A detailed setting for simulation parameters is listed in Table 1. In Fig. 1, we present the contours of the vorticity field during particle sedimentation. The settling of particle clusters was initially relatively uniform, although particles nearer to the wall exhibited faster settling due to the influence of wall vorticity at  $t = 3.125$  s [see Fig. 1(b)]. Over time, the particles gradually assumed distinct shapes, with an umbrella-like formation appearing at  $t = 5$  s [see Fig. 1(c)], followed by the emergence of a bubble-like shape region in the center of the cavity's lower half at  $t = 6.25$  s [see Fig. 1(d)]. As the settling process continued, the particles began to converge in the center of the lower wall [see Fig. 1(e)], and eventually ruptured the bubble region of the cavity [see Figs. 1(f) and 1(g)]. Ultimately, the particles accumulated at the bottom of the cavity, and the fluid gradually

**Table 1**

Simulation parameters for the sedimentation of circular particles in isothermal fluids. Here, l.u. denotes the lattice length-unit and t.s. denotes the lattice time-step in the LB simulation [51]. The length unit conversion is  $l_* = 2.5 \times 10^{-5}$  m/l.u., and the time unit conversion is  $t_* = 3.125 \times 10^{-5}$  s/t.s.

	Physical system	LB system	Unit conversion
Domain size	$W \times H = 5 \text{ cm} \times 5 \text{ cm}$	$\overline{W} \times \overline{H} = 2000 \text{ l.u.} \times 2000 \text{ l.u.}$	$\mathbf{x} = \overline{\mathbf{x}} \cdot l_*$
Particle diameter	$d_p = 1 \text{ mm}$	$\overline{d}_p = 40 \text{ l.u.}$	$d_p = \overline{d}_p \cdot l_*$
Kinematic viscosity	$\nu_f = 10^{-6} \text{ m}^2/\text{s}$	$\overline{\nu}_f = 0.05 \text{ l.u.}^2/\text{t.s.}$	$\nu_f = \overline{\nu}_f \cdot l_*^2/t_*$
Gravity acceleration	$g = 9.8 \text{ m/s}^2$	$\overline{g} = 3.83 \times 10^{-4} \text{ l.u.}/\text{t.s.}^2$	$g = \overline{g} \cdot l_*/t_*^2$



**Fig. 1.** Contour of vorticity field during the sedimentation of 800 circular particles in isothermal fluids. (a)  $t = 0$  s, (b)  $t = 3.125$  s, (c)  $t = 5$  s, (d)  $t = 6.25$  s, (e)  $t = 6.875$  s, (f)  $t = 8.75$  s, (g)  $t = 10$  s, (h)  $t = 17.5$  s, (i)  $t = 25$  s.

returned to stationary state [see Fig. 1(h)]. After 25 s, the particles settled entirely and packed on the bottom wall [see Fig. 1(i)]. The overall patterns observed in the simulations are consistent with previous studies [15,50,58].

### 2.6.2. Sedimentation of a group of 172 circular or elliptical hot particles in cold fluids

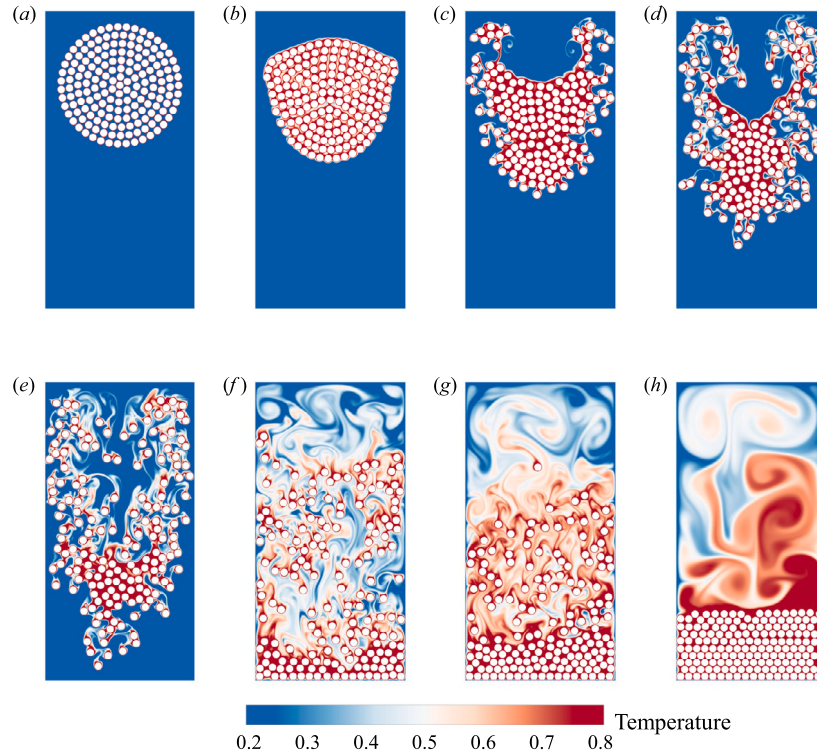
We first consider the sedimentation of circular hot particles in cold fluids. Initially, 172 particles are placed in the upper part of the rectangular cavity, forming a circular cluster. The size of the cavity is  $W \times H = 2 \text{ cm} \times 4 \text{ cm}$ , and all four walls of the cavity are imposed no-slip velocity boundary conditions. Each particle has a density of  $\rho_p = 1.5 \text{ g/cm}^3$  and a diameter of  $d_p = 1 \text{ mm}$ , resulting in a particle volume fraction of 16.9%. The position of the particle cluster is initialized as follows: the center of the cluster is occupied by one particle, located at

$(0.5W, 0.75H)$ , and seven layers of particles are then spread outward in a uniform ring shape, with the number of particles in each layer being 6, 12, 18, 24, 31, 37, 43 [see Fig. 2(a)]. To mimic the working fluid of water, we set its viscosity as  $\nu_f = 10^{-6} \text{ m}^2/\text{s}$  and its density as  $\rho_f = 1 \text{ g/cm}^3$ . We adopt a reference velocity  $U_{ref} = \sqrt{g\pi d_p(\rho_p - \rho_f)/(2\rho_f)}$  to define the particle Reynolds number  $Re_p = U_{ref}d_p/\nu_f$ , and the Archimedes number is  $Ar = \sqrt{gd_p^3(\rho_p - \rho_f)/(\nu_f^2\rho_f)}$ . In this case, we have  $Re_p = 87.7$  and  $Ar = 70.0$ . We set the surface temperature of the particles to be a constant high-temperature  $T = T_h$ , the wall of the square cavity to be a constant low-temperature  $T = T_c$ , and the initial temperature of the fluid also be a low-temperature  $T_c$ ; thus, the temperature difference is  $\Delta T = T_h - T_c$ . We choose the particle Grashof number  $Gr_p = g\beta_T\Delta T d_p^3/\nu_f^2 = 100$ , and the Prandtl number  $Pr = 5$ . A detailed setting for simulation parameters is listed in Table 2. In Fig. 2, we present the contours of the tempera-

**Table 2**

Simulation parameters for the sedimentation of circular hot particles in cold fluids. Here, t.u. denotes the temperature unit in the LB simulation [51]. The length unit conversion is  $l_* = 2.5 \times 10^{-5}$  m/l.u., the time unit conversion is  $t_* = 6.25 \times 10^{-6}$  s/t.s., and the temperature unit conversion is  $T_* = 48.6$  K/t.u.

	Physical system	LB system	Unit conversion
Domain size	$W \times H = 2 \text{ cm} \times 4 \text{ cm}$	$\overline{W} \times \overline{H} = 800 \text{ l.u.} \times 1600 \text{ l.u.}$	$\mathbf{x} = \overline{\mathbf{x}} \cdot l_*$
Particle diameter	$d_p = 1 \text{ mm}$	$\overline{d}_p = 40 \text{ l.u.}$	$d_p = \overline{d}_p \cdot l_*$
Kinematic viscosity	$\nu_f = 10^{-6} \text{ m}^2/\text{s}$	$\overline{\nu}_f = 0.01 \text{ l.u.}^2/\text{t.s.}$	$\nu_f = \overline{\nu}_f \cdot l_*^2/t_*$
Thermal diffusivity	$\alpha_T = 2 \times 10^{-5} \text{ m}^2/\text{s}$	$\overline{\alpha}_T = 0.002 \text{ l.u.}^2/\text{t.s.}$	$\alpha_T = \overline{\alpha}_T \cdot l_*^2/t_*$
Gravity acceleration	$g = 9.8 \text{ m/s}^2$	$\overline{g} = 1.53 \times 10^{-5} \text{ l.u./t.s.}^2$	$g = \overline{g} \cdot l_*/t_*^2$
Thermal expansion coefficient	$\beta = 2.1 \times 10^{-4} \text{ K}^{-1}$	$\overline{\beta} = 1.02 \times 10^{-2} \text{ t.u.}^{-1}$	$\beta = \overline{\beta}/T_*$
Temperature difference	$\Delta_T = 48.6 \text{ K}$	$\overline{\Delta}_T = 1 \text{ t.u.}$	$\Delta_T = \overline{\Delta}_T \cdot T_*$



**Fig. 2.** Contour of dimensionless temperature field  $T^* = (T - T_c)/\Delta_T$  during the sedimentation of 172 circular hot particles in cold fluids at the dimensionless time  $t^* = t/\sqrt{d_p/g}$  of (a) 0, (b) 7.42, (c) 14.85, (d) 22.27, (e) 29.70, (f) 61.87, (g) 84.15, (h) 180.67.

ture field during particle sedimentation. As the particle cluster settles, the particles closest to the vertical centerline experience a faster rate of settling than those located near the walls. The walls impede the settling of the particles, causing the cluster to contort into a pine cone shape [see Fig. 2(b)]. Subsequently, the particles close to the walls are swept upward, elongating the cluster into a crescent shape [see Figs. 2(c) and 2(d)]. As the sedimentation continues, the particle cluster becomes an irregular shape with the generation of numerous thermal plumes [see Figs. 2(e) and 2(f)], and these plumes play a significant role in pushing and mixing the particles. Eventually, most of the particles settle at the bottom of the cavity; however, due to buoyancy, some particles still float for an extended period before they can finally settle [see Figs. 2(g) and 2(h)]. The overall patterns observed in the simulations are consistent with previous studies [16,59,60].

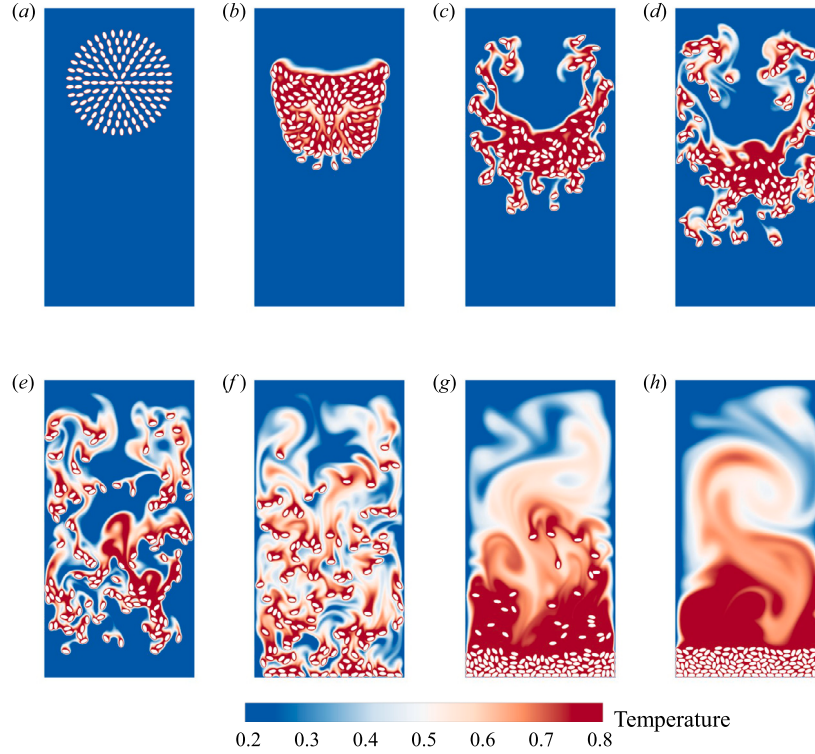
We then consider the sedimentation of elliptical hot particles in cold fluids. Initially, 129 particles are placed in the upper part of the rectangular cavity, forming a circular cluster. The size of the cavity is  $W \times H = 1 \text{ cm} \times 2 \text{ cm}$ , and all four walls of the cavity impose no-slip velocity boundary conditions. Each particle has a density of  $\rho_p = 1.5 \text{ g/cm}^3$ , a major axis of  $A = 0.5 \text{ mm}$  and a minor axis of  $B = 0.25 \text{ mm}$ , resulting in a particle volume fraction of 6.3%. The position of the particle cluster is initialized as follows: the center of the

cluster is occupied by one particle, located at  $(0.5W, 0.75H)$ , and six layers of particles are then spread outward in a uniform ring shape with the number of particles in each layer being 6, 12, 18, 24, 31, 37 [see Fig. 3(a)]. To mimic the working fluid of water, we set its viscosity as  $\nu_f = 10^{-6} \text{ m}^2/\text{s}$  and its density as  $\rho_f = 1 \text{ g/cm}^3$ . We adopt a reference velocity  $U_{ref} = \sqrt{g\pi A(\rho_p - \rho_f)/(2\rho_f)}$  to define the particle Reynolds number  $Re_p = U_{ref}d_p/\nu_f$ , and the Archimedes number is  $Ar = \sqrt{gd_p^3(\rho_p - \rho_f)/(v_f^2\rho_f)}$ . In this case, we have  $Re_p = 31.0$  and  $Ar = 24.7$ . We set the surface temperature of the particles to be a constant high-temperature  $T = T_h$ , the wall of the square cavity to be a constant low-temperature  $T = T_c$ , and the initial temperature of the fluid also be a low-temperature  $T_c$ ; thus, the temperature difference is  $\Delta_T = T_h - T_c$ . We choose the particle Grashof number  $Gr_p = g\beta_T\Delta_TA^3/v_f^2 = 10$ , and the Prandtl number  $Pr = 5$ . A detailed setting for simulation parameters is listed in Table 3. In Fig. 3, we present the contours of the temperature field during particle sedimentation. We can see that the collective motion of elliptical particles during sedimentation is qualitatively the same as that of circular particles: the initially placed cluster is contorted into a pine cone shape, followed by elongating into a crescent shape, and then an irregular shape. Previously, we identified an anomalous rolling mode and an inclined mode for a single hot elliptical particle settling in

**Table 3**

Simulation parameters for the sedimentation of elliptical hot particles in cold fluids. The length unit conversion is  $l_* = 10^{-5}$  m/l.u., the time unit conversion is  $t_* = 10^{-6}$  s/t.s., and the temperature unit conversion is  $T_* = 38.9$  K/t.u.

	Physical system	LB system	Unit conversion
Domain size	$W \times H = 1 \text{ cm} \times 2 \text{ cm}$	$\overline{W} \times \overline{H} = 1000 \text{ l.u.} \times 2000 \text{ l.u.}$	$\mathbf{x} = \overline{\mathbf{x}} \cdot l_*$
Particle size	$A \times B = 0.5 \text{ mm} \times 0.25 \text{ mm}$	$\overline{A} \times \overline{B} = 50 \text{ l.u.} \times 25 \text{ l.u.}$	$d_p = \overline{d}_p \cdot l_*$
Kinematic viscosity	$\nu_f = 10^{-6} \text{ m}^2/\text{s}$	$\overline{\nu}_f = 0.01 \text{ l.u.}^2/\text{t.s.}$	$\nu_f = \overline{\nu}_f \cdot l_*^2/t_*$
Thermal diffusivity	$\alpha_T = 2 \times 10^{-5} \text{ m}^2/\text{s}$	$\overline{\alpha}_T = 0.002 \text{ l.u.}^2/\text{t.s.}$	$\alpha_T = \overline{\alpha}_T \cdot l_*^2/t_*$
Gravity acceleration	$g = 9.8 \text{ m/s}^2$	$\overline{g} = 9.8 \times 10^{-7} \text{ l.u./t.s.}^2$	$g = \overline{g} \cdot l_*/t_*^2$
Thermal expansion coefficient	$\beta = 2.1 \times 10^{-4} \text{ K}^{-1}$	$\overline{\beta} = 8.16 \times 10^{-3} \text{ t.u.}^{-1}$	$\beta = \overline{\beta}/T_*$
Temperature difference	$\Delta_T = 38.9 \text{ K}$	$\overline{\Delta}_T = 1 \text{ t.u.}$	$\Delta_T = \overline{\Delta}_T \cdot T_*$



**Fig. 3.** Contour of dimensionless temperature field  $T^* = (T - T_c)/\Delta_T$  during the sedimentation of 129 elliptical hot particles in cold fluids at the dimensionless time  $t^* = t/\sqrt{A/g}$  of (a) 0, (b) 14.0, (c) 28.0, (d) 42.0, (e) 56.0, (f) 84.0, (g) 140.0, (h) 196.0.

cold fluids [55], suggesting particle shape plays a critical role in affecting the sedimentation; whether there is a new physical mechanism on multiple elliptical particles sedimentation deserves further study within a wide range of control parameters.

### 2.6.3. Sedimentation of 4,800 circular hot particles in cold fluids

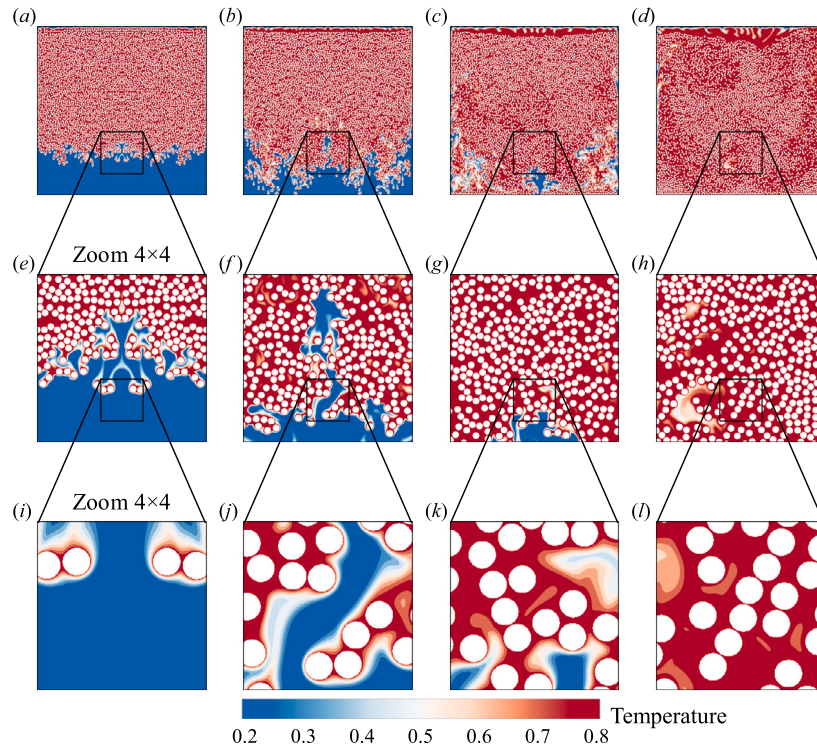
To better utilize the computing power of the GPUs, we also simulate the sedimentation of a large number of 4,800 circular hot particles in cold fluids. The size of the cavity is  $W \times H = 8 \text{ cm} \times 8 \text{ cm}$ , and all four walls of the cavity impose no-slip velocity boundary conditions. Each particle has a density of  $\rho_p = 1.1 \text{ g/cm}^3$  and a diameter of  $d_p = 0.8 \text{ mm}$ , resulting in a particle volume fraction of 37.7%. Initially, the particles are placed in the upper part of a square cavity, and there are 60 lines of particles with each line having 80 particles [see Fig. 4(a)]. To mimic the working fluid of water, we set its viscosity as  $\nu_f = 10^{-6} \text{ m}^2/\text{s}$  and its density as  $\rho_f = 1 \text{ g/cm}^3$ . In this case, we have  $Re_p = U_{ref} d_p / \nu_f = 28.1$  and  $Ar = \sqrt{g d_p^3 (\rho_p - \rho_f) / (\nu_f^2 \rho_f)} = 22.4$ . We set the surface temperature of the particles to be a constant high-temperature  $T = T_h$ , the wall of the square cavity to be a constant low-temperature  $T = T_c$ , and the initial temperature of the fluid also be a low-temperature  $T_c$ ; thus, the temperature difference is  $\Delta_T = T_h - T_c$ . We choose the particle Grashof number

$Gr_p = g \beta_T \Delta_T d_p^3 / \nu_f^2 = 60$ , and the Prandtl number  $Pr = 5$ . A detailed setting for simulation parameters is listed in Table 4. In Fig. 4, we present the contours of the temperature field during particle sedimentation. The particles in a fluid mixture tend to settle due to gravity, and they transfer heat to the surrounding fluid through convection, resulting in the formation of complex flow patterns. The Rayleigh-Taylor instability is a key mechanism that underlies this process, leading to the formation of finger-like patterns, bifurcation in flow patterns, coalescence, breaking of symmetry, and the development of the draft-kissing-tumbling phenomenon. This instability arises from the fact that the heavier particles tend to sink faster than the lighter fluid above them, thus, the interface between the two regions becomes unstable and complex vortices are formed. These vortices, which can vary in size and shape depending on the properties of the fluid mixture and the geometry of the container, play a key role in accelerating the mixing process. By pushing the particles back up to the top of the cavity, the particles are redistributed and the mixing efficiency is enhanced. However, as the particles become more evenly distributed, their settling velocity decreases, eventually leading to a stage of slow settling at the bottom of the container. In Section 3 and Section 4, we will use this problem as the benchmark to measure the parallel performance of the particle-resolved thermal LB simulation.

**Table 4**

Simulation parameters for the sedimentation of circular hot particles in cold fluids. The length unit conversion is  $l_* = 2 \times 10^{-5}$  m/l.u., the time unit conversion is  $t_* = 4 \times 10^{-6}$  s/t.s., and the temperature unit conversion is  $T_* = 56.9$  K/t.u.

	Physical system	LB system	Unit conversion
Domain size	$W \times H = 8 \text{ cm} \times 8 \text{ cm}$	$\overline{W} \times \overline{H} = 4000 \text{ l.u.} \times 4000 \text{ l.u.}$	$\mathbf{x} = \overline{\mathbf{x}} \cdot l_*$
Particle diameter	$d_p = 0.8 \text{ mm}$	$\overline{d}_p = 40 \text{ l.u.}$	$d_p = \overline{d}_p \cdot l_*$
Kinematic viscosity	$\nu_f = 10^{-6} \text{ m}^2/\text{s}$	$\overline{\nu}_f = 0.01 \text{ l.u.}^2/\text{t.s.}$	$\nu_f = \overline{\nu}_f \cdot l_*^2/t_*$
Thermal diffusivity	$\alpha_T = 2 \times 10^{-5} \text{ m}^2/\text{s}$	$\overline{\alpha}_T = 0.002 \text{ l.u.}^2/\text{t.s.}$	$\alpha_T = \overline{\alpha}_T \cdot l_*^2/t_*$
Gravity acceleration	$g = 9.8 \text{ m/s}^2$	$\overline{g} = 7.84 \times 10^{-6} \text{ l.u./t.s.}^2$	$g = \overline{g} \cdot l_*/t_*^2$
Thermal expansion coefficient	$\beta = 2.1 \times 10^{-4} \text{ K}^{-1}$	$\overline{\beta} = 1.20 \times 10^{-2} \text{ t.u.}^{-1}$	$\beta = \overline{\beta}/T_*$
Temperature difference	$\Delta_T = 56.9 \text{ K}$	$\overline{\Delta}_T = 1 \text{ t.u.}$	$\Delta_T = \overline{\Delta}_T \cdot T_*$



**Fig. 4.** Contour of temperature field during the sedimentation of 4,800 circular hot particles in cold fluids at the dimensionless time  $t^* = t/\sqrt{d_p/g}$  of (a) 70.8350, (b) 141.6700, (c) 212.5051, (d) 283.3401. The second row (e-h) and third row (i-l) show successive zooms of the area indicated in the black box.

### 3. Implementation and optimization on a single GPU

A naïve parallel approach for calculating fluid-particle interactions (i.e., implementing Eqs. (13) and (14) to obtain hydrodynamic force and torque, respectively) is to use one GPU thread per computational grid, with 128 threads working together to handle the calculation. However, this approach can lead to load imbalance issues because only fluid nodes near the particle surface perform the calculation. When 128 grids are processed simultaneously, the threads that are not assigned tasks to calculate fluid-particle interaction have to wait for all other threads to complete their calculations. To address this load imbalance issue, we adopt the indirect addressing method inspired by simulating flows in porous media [61,62]. Specifically, we collect the grid lines involved in calculating fluid-particle interaction into a contiguous array in memory, including the position indices (i.e.,  $i$  and  $j$  in the two-dimensional domain) of fluid nodes involved in the hydrodynamic force and torque computation, as well as the discrete velocity direction (i.e.,  $\alpha$ ) across the fluid-particle boundary. The advantage of this method is that only indexes of the participating computational variables are stored, and these indexed flow variables are all involved in the computation to avoid the load imbalance issue. In particle-laden flow, due to the motion of particles, the grid link connections may change at each timestep, and the

index information must be collected simultaneously. This is different from simulating flow in porous media, where the solid skeleton is stationary, and the indirectly addressed indexes only need to be calculated once throughout the simulation.

In addition, we establish a fixed mapping between the fluid-particle grid and the continuous memory array. As illustrated in Fig. 5, for a particle whose center is  $(x_c, y_c)$ , we only consider whether a fluid node possesses boundary links for  $x$ -coordinate ranges from  $INT(x_c - d_p/2) - 1$  to  $INT(x_c + d_p/2) + 2$  and  $y$ -coordinate ranges from  $INT(y_c - d_p/2) - 1$  to  $INT(y_c + d_p/2) + 2$ . Here,  $INT(d_p)$  is the largest integer not exceeding the magnitude of  $d_p$ . Previous results shown that this boundary link search algorithm is significantly faster than blindly comparing each node to each particle directly [63]. For a discrete velocity parallel to the grid line (e.g.,  $\mathbf{e}_1$ ), there will be  $N$  lines crossing the effective domain, with vertical coordinates  $j = 1, \dots, N$ . Here,  $N = INT(x_c + d_p/2) - INT(x_c - d_p/2) + 3$  denotes the length of the effective zone. Each line will have at most one fluid-particle interaction node pointing in the direction of  $\mathbf{e}_1$ . For a discrete velocity diagonal to the grid line (e.g.,  $\mathbf{e}_6$ ), there will be  $2N - 3$  lines crossing the effective domain and their coordinates meet the relation of  $i + j = 3, 4, \dots, 2N - 1$ . Similarly, each line will have at most one fluid-particle interaction node pointing in the direction of  $\mathbf{e}_6$ . Thus, based on the D2Q9 discrete veloc-




Fig. 5. Fixed mapping from solid-fluid grid linkage to continuous memory arrays. The gray area represents the effective zone with a size of  $N \times N$ , where  $N = INT(x_c + d_p/2) - INT(x_c - d_p/2) + 3$ .

Table 5

Comparison of running time between direct and indirect addressing methods. The iteration steps are fixed as 6000. Data included in the brackets represent the percentage of time consumption. Note routines that account for less than 1% of the total time are neglected.

	Fluid-related	Particle-related	Overall
Direct addressing	41.9 s (52.8%)	37.1 s (46.7%)	79.4 s
Indirect addressing	43.1 s (78.6%)	11.3 s (20.6%)	54.9 s
Speed-up	1X	3.3X	1.4X

ity model, the maximum number of indexes collected across the entire influence domain is  $4 \times N + 4 \times (2N - 3) = 12N - 12$ . Although a few of these indexes may be unavailable (e.g., the line with  $j = 1$  does not intersect with the particle, resulting in two invalid index positions), it is unlikely to cause severe load imbalance issues.

In Table 5, we compare the running time between direct and indirect addressing methods using the NVIDIA Nsight Systems tool. Here, we simulate the sedimentation of 4,800 hot particles in cold fluids (see detailed settings in Section 2.6.3), and the iteration steps are fixed as 6000. The analysis involved summing the time consumed by all fluid-related computation steps and particle-related computation steps. Note subroutines that account for less than 1% of the total time are neglected. The results indicate that, although the indirect addressing method requires an additional step for collecting index information, it reduces the time consumption for particle-related computation by a factor of 3.3X compared to the direct addressing method. This leads to an overall improvement in code performance by 1.4X. We also adopt the Million Lattice Update Per Second (MLUPS) as the metric to characterize the parallel performance of the LB simulation, which is defined as [64]

$$MLUPS = \frac{\text{mesh size} \times \text{iteration steps}}{\text{running time} \times 10^6} \quad (18)$$

We obtain the overall parallel performance of 1209 MLUPS and 1750 MLUPS for the direct and indirect addressing methods, respectively.

#### 4. Implementation and optimization on multi-GPUs

A major limitation in GPU computing is the available device memory; for example, the state-of-art NVIDIA A100 GPU accelerator offers a maximum of 40 GB. A solution to the memory limitation is to use multiple GPUs, where the GPUs are distributed across multiple CPU




Fig. 6. Parallel performance in terms of (a) the MLUPS and (b) the parallel efficiency for a naïve implementation using hybrid OpenACC and MPI technique.

nodes, and MPI is used to coordinate the computational tasks. In this work, we conducted experiments on a GPU cluster where each node is equipped with four NVIDIA A100 GPUs. The network interconnects use 100 Gigabits per second (Gbps) Remote direct memory access over Converged Ethernet (RoCE). The inter-GPU-GPU communication within a node goes over the PCI-e. To automatically utilize the GPUDirect acceleration technologies, we adopt a CUDA-aware MPI implemented in OpenMPI. With GPUDirect technology, including Peer to Peer (P2P) and Remote Direct Memory Access (RDMA), the buffers can be directly sent from a GPU's memory to another GPU's memory or the network without touching the host memory [65].

##### 4.1. A naïve implementation using hybrid OpenACC and MPI technique

We adopt a mono-dimensional partitioning of the computational domain and decompose the domain along the  $y$ -direction. To facilitate communication between sub-domains, we add three ghost layers outside the boundary of each subdomain to exchange data with adjacent subdomains. In the particle-related calculation, it is essential to exchange the information of  $f^+$  and  $g^+$  at the boundary nodes after the collision step. When a particle is close to the boundary of a subdomain, the refilling calculation on the new 'born' fluid node requires information from adjunct subdomains to interpolate the distribution function. After the collision step, each subdomain exchanges the  $f_{1-8}^+$  at two layers of boundary nodes, and  $g_{1-4}^+$  at the boundary nodes with their neighbors. After calculating the fluid-particle interaction, the  $f_{0-8}$  at three layers of boundary nodes is exchanged with their neighbors. Although the computational domain is decomposed into slices and each GPU only accesses the fluid nodes of its assigned subdomain, the information of all particles is shared among all GPUs. This means that each particle-related loop requires iteration among all particles, leading to a loss of computational efficiency. The fluid nodes near the particle surface exchange momentum with the particles, and we use the MPI ALLREDUCE to obtain the combined forces on the particle in all subdomains.

Fig. 6 shows the parallel performance in terms of the MLUPS and parallel efficiency. Parallel efficiency is defined as  $\eta = T_1 / (T_n \cdot n)$ . Here,  $T_1$  denotes the running time using a single GPU, and  $T_n$  denotes the running time using  $n$  GPUs. We can see that with an increase in the number of GPUs, the MULPS generally increases. Using 8 GPUs, the simulation can achieve 5572 MLUPS, indicating a higher computational capacity. However, the parallel efficiency degrades when more GPUs are used. Notably, the parallel efficiency is only 39.8% using 8 GPUs, indicating the parallel code is not scalable. Two reasons may be responsible for the poor parallel performance. First, the presence of redundant data exchange, such as the exchange of distribution functions on three layers of ghost nodes, can increase communication overhead. Secondly, although the computational domain has been decomposed into slices and each GPU is only responsible for updating fluid node information within its subdomain, information on the particle group is required on every GPU, leading to additional communication overhead. To further boost the parallel performance using multi-GPUs, we describe some optimization strategies in the following subsections.






Fig. 7. Schematic illustration to decompose the particle groups and build the domain list.

#### 4.2. Building the domain list

In the previous subsection, we decomposed the flow field into slices and assigned each subdomain to one GPU. In this subsection, we further decompose the particles into subgroups and assign each subgroup to one GPU. In the particle-resolved LB method, the fluid-particle interaction within a timestep only affects fluid nodes near the particle surface. Thus, most of the particles in a subdomain do not have to share information with other subdomains, except for particles near the subdomain boundary. As illustrated in Fig. 7, we define an extended subdomain for a particle group compared to that for the flow field, which includes three parts: the upper and lower parts with a length of  $L$  (referred to as the top halo region and bottom halo region, respectively), and the central part similar as that of the sliced flow field (referred to as the exclusive region). Particles whose centers in the halo region share their information with the neighboring subdomains; while particles in the exclusive region do not communicate their information with other subdomains. Here,  $L$  is determined based on the farthest distance a particle can interact with the fluid node and another particle. The vertical length of the particle subdomain extends upward and downward by  $d_p + s$ , where  $d_p$  denotes the particle diameter and  $s$  denotes the threshold value for calculating the interparticle interaction force. Thus, the shared region has a length of  $L = 2(d_p + s)$ .

After decomposing the particle groups, we assign each subgroup to a GPU and build domain lists to store the region index of each particle. In this way, the particle-related calculations do not require iteration over all particles but only over the particles within the same list, thus reducing the size of the particle search. To synchronize the forces of the particles, only particles in the top halo region and bottom halo region need to be shared with the adjacent subdomains. Because the forces of the particles in the exclusive region do not require synchronization with other subdomains, we can reduce the size of messages that need to be sent. If a particle enters the exclusive region from the top (or bottom) halo region, it will disappear from the subdomain list of the neighboring subdomains. On the other hand, if a particle enters the top (or bottom) halo region from the exclusive region, information about this particle must be sent to the adjacent subdomain and added to their lists. However, updating the domain list introduces additional computational costs, which may not be worth it for particles with low-volume fractions.

Fig. 8 compares the performance of sharing particles' information among all GPUs and building the domain list. We can see that building domain list can improve parallel performance for all the cases. Specifically, using 8 GPUs, the simulation can achieve 6796 MLUPS with a




Fig. 8. Performance comparisons between sharing particles' information among all GPUs and building domain list in terms of (a) the MLUPS and (b) the parallel efficiency.

parallel efficiency of 48.5%. It should be noted that the idea of building domain list was inspired by the cell-linked list in smoothed particle hydrodynamics (SPH) simulations [66], where the number of particles can range from hundreds of thousands to millions, leading to a heavy computational load when searching neighboring particles and calculating particle interaction forces. To overcome this challenge, the simulation domain was divided into cells, and neighbor lists were created within each cell to limit the calculations to the same or adjacent cells. In the particle-resolved LB method, the particle size is much larger than the grid size, and the particle number is much smaller than that in SPH. Thus, it is unnecessary to further divide the subdomain assigned to each GPU into smaller ones; in other words, we only build one domain list within a subdomain. As the number of particles increases, the benefits gained from building domain lists will become more obvious.

#### 4.3. Utilizing request-answer communication

The refilling scheme used to construct unknown distribution functions (as described in subsection 2.5) can result in communication overhead if node information is required from neighboring subdomains. Due to the constant change in particle position, it is difficult to predict in advance which node information needs to be communicated. A straightforward solution is to exchange all possible information for extrapolation; however, this approach degrades the utilization efficiency of exchanged data, particularly for particles with a large specific area. To avoid the passing of excess data, we adopt the request-answer communication method between GPUs, which is similar to the communication method between CPUs [67]. The basic idea behind this method is to exchange only distribution functions that are needed for the extrapolation, rather than exchanging all distribution functions at boundary nodes of nearby subdomains. As illustrated in Fig. 9, when the particle surface moves from the dashed curve to the solid curve, the node  $x_{new}$  changes from a solid node to a fluid node. We assume that the extrapolation of the density distribution function requires information on the fluid nodes  $x_f$ ,  $x_{ff}$ ,  $x_{fff}$ , which are stored in the adjacent subdomain. To exchange this information, GPU1 sends a request to GPU0 to exchange the information at fluid nodes  $x_f$ ,  $x_{ff}$ , and  $x_{fff}$ . Upon receiving the request, GPU0 sends the information of distribution functions  $f_{0-8}(x_f, \delta_t)$ ,  $f_{0-8}(x_{ff}, \delta_t)$ , and  $f_{0-8}(x_{fff}, \delta_t)$  to GPU1 in the order they were requested. In this way, the fluid refilling calculation can be performed with only two communications. Practically, we pack discontinuous data into contiguous memory, exchange the information using MPI, and then unpack the synchronized information to the desired location. We allocate a sufficiently large array at the receiver to receive the entire message and use *MPI GET COUNT* to obtain its length. In the second communication, the distribution functions should be sent in the order that corresponds to the requested information, which makes it easier for the receiver to unpack the information.

Fig. 10 compares the performance of the above mentioned two approaches: exchanging all boundary node information and utilizing the request-answer communication approach. We can see the simulation can achieve 7746 MLUPS using 8 GPUs with a parallel efficiency of 55.3%. However, with a smaller number of GPUs, the parallel efficiency does not improve obviously. Upon further analysis of the time consump-

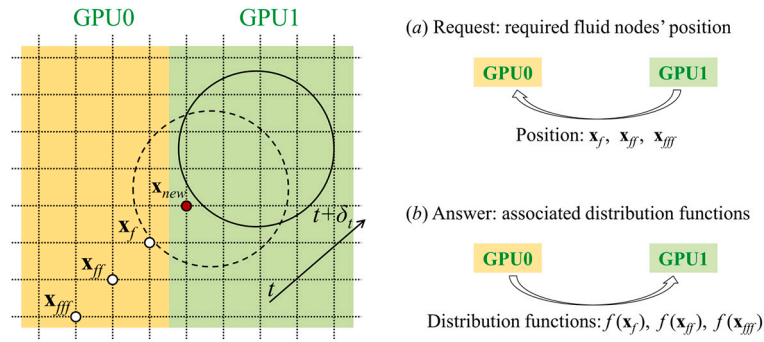


Fig. 9. Schematic illustration of request-answer communication: (a) sending the request to neighboring GPUs and (b) returning the requested distribution function.





Fig. 10. Performance comparisons between exchanging all boundary node information and utilizing request-answer communication in terms of (a) the MLUPS and (b) the parallel efficiency.

tion using the NVIDIA Nsight Systems tool, we find that even though the request-answer communication approach reduced the amount of data transfer between GPUs, the total time consumption (including both communication initiation and data transfer duration) increased. This increase was likely due to the increased latency on the CPU side, resulting from multiple communication initiations and calls to the *MPI GET COUNT* function. When dealing with 8 or more GPUs, inter-GPU communication encounters limitations within the PCI-e interface, leading to a decrease in data transfer speeds and an increase in data transfer duration. This is the scenario where the request-answer communication approach enhances parallel performance and overall efficiency. In the 2D simulation, due to low communication overhead, the increased latency on the CPU side outweighs the benefit of reduced communications between GPUs. For this reason, we do not use the request-answer approach to optimize the interpolation of the distribution function at the curved surface of the particle, as the additional communication load from the low percentage of invalid information was acceptable for high-volume fractions of particles. We deduce that in three-dimensional simulations when the overhead between GPU communications is more intense, the use of a request-answer communication approach would be more effective.


#### 4.4. Overlapping communications with computations

Previous studies have demonstrated that hiding communication overhead behind the kernel runtime can increase the parallel performance by a factor of around 1.3X for LB simulation of single-phase flow [28,68,69]. In Fig. 11, we further illustrate the overlap of communications with computations for LB simulation of particle-laden thermal flows. Specifically, we first carry out the collision step to update the density distribution function ( $f$ ) and temperature distribution function ( $g$ ) at boundary nodes, followed by building domain lists in each subdomain. If a particle in a subdomain transition from an exclusive state to a shared state, namely it moves from the exclusive region to a top (or bottom) halo region, we use MPI communications to synchronize the particles' information, including the position ( $\mathbf{r}$ ), the velocity ( $\mathbf{U}$ ), the orientation angle ( $\theta$ ), and the angular velocity ( $\Omega_z$ ). The synchronization of the particle's information can be overlapped with updating the temperature distribution function ( $g$ ) at inner nodes. After that,

(a) Request: required fluid nodes' position



(b) Answer: associated distribution functions



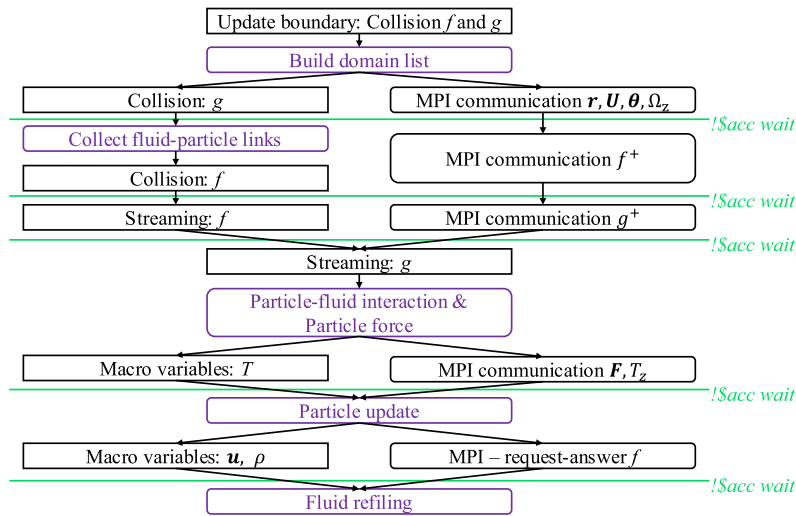
we collect the particle-fluid grid linkage and carry the collision step to update the density distribution function ( $f$ ) at inner nodes; meanwhile, we synchronize the post-collision density distribution function  $f^+$  to hide the kernel runtime. Similarly, the synchronization of post-collision temperature distribution function  $g^+$  can be overlapped with the streaming of density distribution function ( $f$ ). After the streaming of the temperature distribution function ( $g$ ) and the calculation of fluid-particle interactions, the particles in the top (or bottom) halo region of the subdomain lists need to exchange the information of fluid-particle interaction force ( $\mathbf{F}$ ) and torque ( $T_z$ ) with the neighboring subdomains, which can be overlapped with the computation of macroscopic temperature ( $T$ ). Before the fluid refilling calculation, we synchronize the density distribution function ( $f$ ) using the request-answer communication approach, which can be overlapped with the computation of macroscopic velocity ( $\mathbf{u}$ ) and density ( $\rho$ ).

Fig. 12 compares the performance between non-overlapping and overlapping communication. We can see that the MLUPS and parallel efficiency improved for all the cases when the communications and computations overlapped, and the advantage of using the overlapping mode becomes more pronounced with an increase in the number of GPUs. Notably, the performance increases by 1.22X when using 8 GPUs (i.e., increase from 7746 MLUPS to 9466 MLUPS, and a parallel efficiency from 55.3% to 67.6%), which shows similar gains to that for LB simulation of single-phase flow, suggesting that hiding communication overhead behind the kernel runtime is an effective approach for optimizing LB simulations on multi-GPUs [28,68,69].


#### 4.5. Executing computation tasks concurrently

In the heterogeneous CPU-GPU architecture, the CPU acts as the controller, which offloads data and computational tasks to the GPU, and retrieves the results when the computation is complete, as illustrated in Fig. 13(a). This non-concurrent computation can lead to performance bottlenecks if the CPU is unable to keep up with the demands of the GPU. A solution is to take advantage of the parallel processing power of the GPU and execute computation tasks concurrently. In the OpenACC, task parallelism can be achieved via the `!$acc async(n)` handle to concurrently execute independent tasks on a single GPU. As illustrated in Fig. 13(b), each task is processed by a separate stream of GPU commands, and the concurrent execution of computation tasks reduces synchronization between CPU and GPU. By processing multiple tasks in parallel, the GPU can be kept busy, and utilization can be maximized.

Fig. 14 compares the performance between non-concurrent and concurrent computation. We can see that the MLUPS and parallel efficiency slightly improved for all the cases if the computations are executed concurrently. With 8 GPUs, the simulation achieved 10846 MLUPS, yielding a parallel efficiency of 77.5%. We further analyzed the runtime of concurrent computation using the NVIDIA Nsight Systems tool and found that each GPU has three distinct work queues that simultaneously execute the computational tasks.



**Fig. 11.** Schematic illustration of overlapping communications with computations for LB simulation of particle-laden thermal flows. The purple rectangular represents particle-related computations.



**Fig. 12.** Performance comparisons between non-overlapping and overlapping communication with computation in terms of (a) the MLUPS and (b) the parallel efficiency.

## 5. Conclusions

In this work, we have utilized OpenACC-based GPU computing to perform particle-resolved thermal LB simulations, in which the momentum-exchange method was adopted to calculate particle-fluid interactions. We extended the indirect addressing method to collect fluid-particle link information at each timestep and store indices of fluid-particle links in a fixed index array. This mapping of the index array helps solve the issue of load imbalance by ensuring fluid-particle interactions are only calculated at indexed positions. Using this approach, the simulation of 4,800 hot particles settling in cold fluids with a domain size of  $4000^2$  achieved 1750 MLUPS on a single GPU.

We also implemented a hybrid approach combining OpenACC and MPI for multi-GPU accelerated simulation. This approach incorporates four optimization strategies to enhance parallel performance. First, we build the domain list and optimize the fluid-particle interactions by considering only those within the same domain or adjacent domains, thereby avoiding the need to loop over all particles. Next, we utilize request-answer communication and exchange only the necessary distribution functions, rather than exchanging information for all boundary nodes. To further improve performance, we overlap communications with computations. This allows us to hide communication latency behind the consumed computational time, resulting in significant gains for multi-GPU simulations. Additionally, we maximize the utilization of GPU resources by executing computational tasks concurrently, enhancing parallel efficiency by ensuring efficient use of available processing power. Overall, using 8 GPUs, these optimizations lead to a parallel performance increase from 5572 MLUPS to 10846 MLUPS, with a corresponding improvement in parallel efficiency from 39.8% to 77.5%. To ensure the correctness of the code utilizing the hybrid OpenACC and MPI approach, we recommend an incremental approach utilize the above four optimization strategies to accelerate the code.

In the future, we plan to extend these optimization strategies to three-dimensional particle-resolved thermal flows, where the computational load is intense and the overhead to launch the kernel is relatively lower. In the three-dimension simulation, to reduce memory requirements, the distribution function can be reconstructed from available hydrodynamic variables instead of storing the full set of discrete populations [70].

## CRediT authorship contribution statement

**Ao Xu:** Conceptualization, Formal analysis, Funding acquisition, Project administration, Resources, Software, Supervision, Writing – original draft, Writing – review & editing. **Bo-Tao Li:** Data curation, Formal analysis, Investigation, Software, Validation, Visualization, Writing – original draft.

## Declaration of competing interest

We confirm that the manuscript has been read and approved by all named authors and that there are no other persons who satisfied the criteria for authorship but are not listed. We further confirm that the order of authors listed in the manuscript has been approved by all of us.

## Data availability

Data will be made available on request.

## Acknowledgements

This work was supported by the National Natural Science Foundation of China (NSFC) through Grant Nos. 12272311 and 11902268, and the Open Fund of Key Laboratory of Icing and Anti/De-icing (Grant No. IADL20200301). The authors acknowledge the Beijing Beilong Super Cloud Computing Co., Ltd for providing HPC resources that have contributed to the research results reported within this paper (URL: <http://www.blsc.cn/>).

## Appendix A. An elliptical cold particle settling in hot fluids

We compare the trajectory and orientation of a single elliptical cold particle settling in hot fluids using both the multi-GPU code and the corresponding CPU code. The contour of the temperature field during

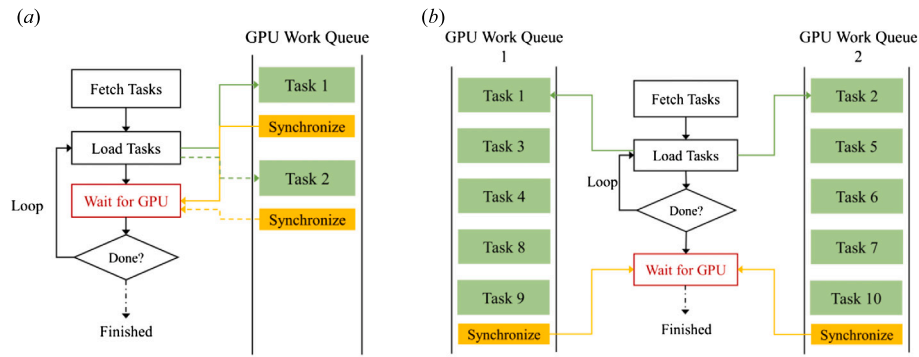


Fig. 13. Schematic illustration of (a) non-concurrent and (b) concurrent execution of computational tasks on a single GPU.




Fig. 14. Performance comparisons between non-concurrent and concurrent computation in terms of (a) the MLUPS and (b) the parallel efficiency.

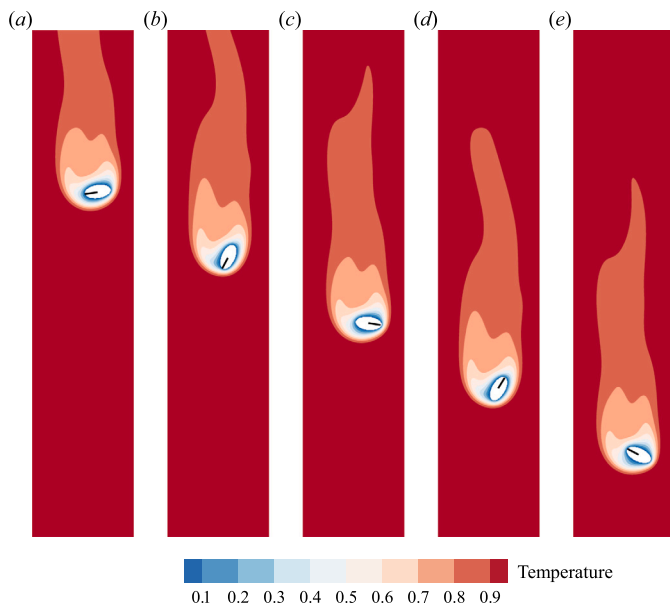


Fig. 15. Contour of dimensionless temperature field  $T^* = (T - T_c) / \Delta T$  during the sedimentation of an elliptical cold particle in a hot fluid at the dimensionless time  $t^* = t\nu / A^2$  of (a) 6.875, (b) 7.5, (c) 8.125, (d) 8.75, (e) 9.375. Note that only heights between 1.875 cm and 3.875 cm are shown for better visualization.

the sedimentation is shown in Fig. 15. The simulation setting is similar to our previous work [55]. However, instead of using the moving domain technique to mimic an infinitely long channel, we now adopt an alternative approach of utilizing a closed cavity with a very small width-to-height aspect ratio [17] to minimize the end effect of top and bottom boundaries. Specifically, the size of the cavity is  $W \times H = 0.4 \text{ cm} \times 8 \text{ cm}$ , and all four walls of the cavity are imposed no-slip velocity boundary conditions.

In the simulation, each particle has a density of  $\rho_p = 1.001 \text{ g/cm}^3$ , a major axis of  $A = 1 \text{ mm}$ , and a minor axis of  $B = 0.5 \text{ mm}$ . The particle is released at  $(0.5W, 0.75H)$  with an initial angle of  $60^\circ$  between the particle's major axis and the horizontal direction. To mimic




Fig. 16. Time series of (a, b) horizontal and vertical positions of the particle center, and (b, c) the angle  $\theta$  between the particle major axis and the horizontal direction in terms of  $\cos \theta$  and  $\sin \theta$ , respectively.

the working fluid of water, we set its viscosity as  $\nu_f = 10^{-6} \text{ m}^2/\text{s}$  and its density as  $\rho_f = 1 \text{ g/cm}^3$ . In this case, we have  $Re_p = U_{ref} d_p / \nu_f = 3.92$  and  $Ar = \sqrt{g A^3 (\rho_p - \rho_f) / (\nu_f^2 \rho_f)} = 3.13$ . Meanwhile, we choose  $Gr_p = g \beta_T \Delta T A^3 / \nu_f^2 = 200$  and  $Pr = 7$ . From Fig. 16, we can see that the simulation using the moving domain technique [55] gives the same results as that adopting a sufficiently large domain; in addition, the multi-GPU simulation provides the same results to those obtained from the CPU-based simulation.


## Appendix B. Two circular hot particles settling in cold fluids

We conduct a comparison between the velocity and position of two circular hot particles settling in cold fluids to validate the particle-particle interactions, and the contour of the temperature field during the sedimentation is shown in Fig. 17. The simulation setting is similar to the one used in Tao et al. [60], known as draft-kissing-tumbling (DKT) with convection. Specifically, the size of the cavity is  $W \times H = 2 \text{ cm} \times 6 \text{ cm}$ , and all four walls of the cavity are imposed no-slip velocity boundary conditions. Each particle has a density of  $\rho_p = 1.01 \text{ g/cm}^3$  and a diameter of  $d_p = 2 \text{ mm}$ . The lower particle is released at  $(0.5W - 0.005d_p, 0.8H)$  and the upper particle is released at  $(0.5W, 0.8H + 2d_p)$ . The lower particle was deliberately offset from channel centerline to induce tumbling, because previous calculations showed that otherwise both particles remained perfectly aligned for a long time after catching up, and the offset prevents prolonged stable alignment.

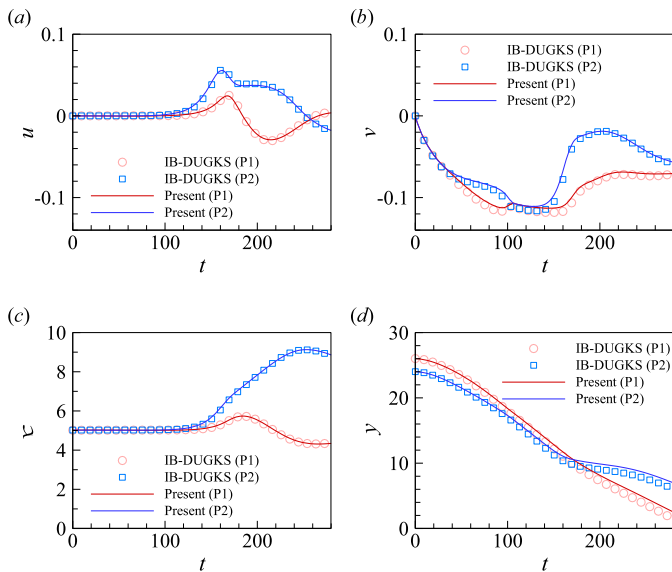
**Table 6**

Simulation parameters for the sedimentation of two circular hot particles in cold fluids. The length unit conversion is  $l_* = 5 \times 10^{-5}$  m/l.u., the time unit conversion is  $t_* = 2.5 \times 10^{-5}$  s/t.s., and the temperature unit conversion is  $T_* = 6.07$  K/t.u.

	Physical system	LB system	Unit conversion
Domain size	$W \times H = 2 \text{ cm} \times 6 \text{ cm}$	$\overline{W} \times \overline{H} = 400 \text{ l.u.} \times 1200 \text{ l.u.}$	$\mathbf{x} = \overline{\mathbf{x}} \cdot l_*$
Particle diameter	$d_p = 2 \text{ mm}$	$\overline{d}_p = 40 \text{ l.u.}$	$d_p = \overline{d}_p \cdot l_*$
Kinematic viscosity	$\nu_f = 10^{-6} \text{ m}^2/\text{s}$	$\overline{\nu}_f = 0.01 \text{ l.u.}^2/\text{t.s.}$	$\nu_f = \overline{\nu}_f \cdot l_*^2/t_*$
Thermal diffusivity	$\alpha_T = 10^{-6} \text{ m}^2/\text{s}$	$\overline{\alpha}_T = 0.01 \text{ l.u.}^2/\text{t.s.}$	$\alpha_T = \overline{\alpha}_T \cdot l_*^2/t_*$
Gravity acceleration	$g = 9.8 \text{ m/s}^2$	$\overline{g} = 1.23 \times 10^{-4} \text{ l.u./t.s.}^2$	$g = \overline{g} \cdot l_*/t_*^2$
Thermal expansion coefficient	$\beta = 2.1 \times 10^{-4} \text{ K}^{-1}$	$\overline{\beta} = 1.28 \times 10^{-3} \text{ t.u.}^{-1}$	$\beta = \overline{\beta}/T_*$
Temperature difference	$\Delta_T = 6.07 \text{ K}$	$\overline{\Delta}_T = 1 \text{ t.u.}$	$\Delta_T = \overline{\Delta}_T \cdot T_*$



**Fig. 17.** Contour of dimensionless temperature field  $T^* = (T - T_c)/\Delta_T$  during the sedimentation of two circular hot particles settling in cold fluids at the dimensionless time  $t^* = t/\sqrt{d_p/g}$  of (a) 0, (b) 70, (c) 140, (d) 280.



**Fig. 18.** Time series of (a, b) velocity and (c, d) position of the two circular hot particles settling in cold fluids. The time is normalized as  $t^* = t/\sqrt{d_p/g}$ , the velocity is normalized as  $\mathbf{u}^* = \mathbf{u}/\sqrt{gd_p}$ , and the position is normalized as  $\mathbf{x}^* = \mathbf{x}/d_p$ . Here, “P1” denotes the upper particle, and “P2” denotes the lower particle. Data shown in (b) are from Tao et al. [60], while data shown in (a,c,d) are from private communication with S. Tao.

We set the viscosity of the fluid as  $\nu_f = 10^{-6} \text{ m}^2/\text{s}$  and its density as  $\rho_f = 1 \text{ g/cm}^3$ . In this case, we have  $Re_p = U_{ref} d_p/\nu_f = \sqrt{gd_p}(\rho_p - \rho_f)/(2\rho_f) = 35.09$  and  $Ar = \sqrt{gd_p^3(\rho_p - \rho_f)/(\nu_f^2 \rho_f)} = 28.0$ . Meanwhile, we choose  $Gr_p = g\beta_T \Delta_T A^3/\nu_f^2 = 100$  and  $Pr = 1$ . A detailed setting for simulation parameters is listed in Table 6. From Fig. 18, we can see that the present results show good agreement with Tao et al. [60] using a sharp interface immersed boundary-discrete unified gas kinetic scheme (IB-DUGKS), thus validating the code’s ability to simulate particle-particle interactions.

## References

- [1] V. Mathai, D. Lohse, C. Sun, Bubbly and buoyant particle-laden turbulent flows, *Annu. Rev. Condens. Matter Phys.* 11 (2020) 529–559.
- [2] L. Brandt, F. Coletti, Particle-laden turbulence: progress and perspectives, *Annu. Rev. Fluid Mech.* 54 (2022) 159–189.
- [3] P. Wang, Y. Lei, Z. Zhu, X. Zheng, Drag model of finite-sized particle in turbulent wall-bound flow over sediment bed, *J. Fluid Mech.* 964 (2023), A9.
- [4] D. Norbäck, C. Lu, Y. Zhang, B. Li, Z. Zhao, C. Huang, X. Zhang, H. Qian, Y. Sun, J. Wang, et al., Sources of indoor particulate matter (PM) and outdoor air pollution in China in relation to asthma, wheeze, rhinitis and eczema among pre-school children: synergistic effects between antibiotics use and PM10 and second hand smoke, *Environ. Int.* 125 (2019) 252–260.
- [5] M. Maxey, Simulation methods for particulate flows and concentrated suspensions, *Annu. Rev. Fluid Mech.* 49 (2017) 171–193.
- [6] V. Patočka, E. Calzavarini, N. Tosi, Settling of inertial particles in turbulent Rayleigh-Bénard convection, *Phys. Rev. Fluids* 5 (11) (2020) 114304.
- [7] V. Patočka, N. Tosi, E. Calzavarini, Residence time of inertial particles in 3D thermal convection: implications for magma reservoirs, *Earth Planet. Sci. Lett.* 591 (2022) 117622.
- [8] W. Yang, Y.-Z. Zhang, B.-F. Wang, Y. Dong, Q. Zhou, Dynamic coupling between carrier and dispersed phases in Rayleigh-Bénard convection laden with inertial isothermal particles, *J. Fluid Mech.* 930 (2022) A24.
- [9] W. Yang, Z.-H. Wan, Q. Zhou, Y. Dong, On the energy transport and heat transfer efficiency in radiatively heated particle-laden Rayleigh-Bénard convection, *J. Fluid Mech.* 953 (2022) A35.
- [10] H.H. Hu, N.A. Patankar, M. Zhu, Direct numerical simulations of fluid–solid systems using the arbitrary Lagrangian–Eulerian technique, *J. Comput. Phys.* 169 (2) (2001) 427–462.
- [11] C.S. Peskin, The immersed boundary method, *Acta Numer.* 11 (2002) 479–517.
- [12] R. Glowinski, T.-W. Pan, T.I. Hesla, D.D. Joseph, J. Periaux, A fictitious domain approach to the direct numerical simulation of incompressible viscous flow past moving rigid bodies: application to particulate flow, *J. Comput. Phys.* 169 (2) (2001) 363–426.
- [13] Z. Yu, X. Shao, A. Wachs, A fictitious domain method for particulate flows with heat transfer, *J. Comput. Phys.* 217 (2) (2006) 424–452.
- [14] C.K. Aidun, J.R. Clausen, Lattice-Boltzmann method for complex flows, *Annu. Rev. Fluid Mech.* 42 (2010) 439–472.
- [15] S. Tao, H. Zhang, Z. Guo, L.-P. Wang, A combined immersed boundary and discrete unified gas kinetic scheme for particle–fluid flows, *J. Comput. Phys.* 375 (2018) 498–518.
- [16] K. Walayat, Z. Wang, K. Usman, M. Liu, An efficient multi-grid finite element fictitious boundary method for particulate flows with thermal convection, *Int. J. Heat Mass Transf.* 126 (2018) 452–465.
- [17] K. Suzuki, T. Kawasaki, T. Asaoka, M. Yoshino, Numerical simulations of solid–liquid and solid–solid interactions in ice slurry flows by the thermal immersed boundary–lattice Boltzmann method, *Int. J. Heat Mass Transf.* 157 (2020) 119944.
- [18] A.D. Demou, M.N. Ardekani, P. Mirbod, L. Brandt, Turbulent Rayleigh-Bénard convection in non-colloidal suspensions, *J. Fluid Mech.* 945 (2022) A6.
- [19] M.J. Krause, A. Kummerländer, S.J. Avis, H. Kusumaatmaja, D. Dapelo, F. Klemens, M. Gaedte, N. Hafen, A. Mink, R. Trunk, et al., OpenLB—open source lattice Boltzmann code, *Comput. Math. Appl.* 81 (2021) 258–288.

- [20] J. Latt, O. Malaspinas, D. Kontaxakis, A. Parmigiani, D. Lagrafa, F. Brogi, M.B. Belgacem, Y. Thorimbert, S. Leclaire, S. Li, et al., Palabos: parallel lattice Boltzmann solver, *Comput. Math. Appl.* 81 (2021) 334–350.
- [21] M. Januszewski, M. Kostur, Sailfish: a flexible multi-GPU implementation of the lattice Boltzmann method, *Comput. Phys. Commun.* 185 (9) (2014) 2350–2368.
- [22] G. Amati, S. Succi, P. Fanelli, V.K. Kravets, G. Falcucci, Projecting LBM performance on exascale class architectures: a tentative outlook, *J. Comput. Sci.* 55 (2021) 101447.
- [23] Z. Liu, X. Chu, X. Lv, H. Meng, S. Shi, W. Han, J. Xu, H. Fu, G. Yang, Sunwaylb: enabling extreme-scale lattice Boltzmann method based computing fluid dynamics simulations on sunway taihulight, in: 2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS), IEEE, 2019, pp. 557–566.
- [24] G. Falcucci, G. Amati, P. Fanelli, V.K. Kravets, G. Polverino, M. Porfiri, S. Succi, Extreme flow simulations reveal skeletal adaptations of deep-sea sponges, *Nature* 595 (7868) (2021) 537–541.
- [25] O. Navarro-Hinojosa, S. Ruiz-Loza, M. Alencastre-Miranda, Physically based visual simulation of the Lattice Boltzmann method on the GPU: a survey, *J. Supercomput.* 74 (2018) 3441–3467.
- [26] K.E. Niemeyer, C.-J. Sung, Recent progress and challenges in exploiting graphics processors in computational fluid dynamics, *J. Supercomput.* 67 (2014) 528–564.
- [27] A. Xu, L. Shi, T. Zhao, Accelerated lattice Boltzmann simulation using GPU and OpenACC with data management, *Int. J. Heat Mass Transf.* 109 (2017) 577–588.
- [28] A. Xu, B.-T. Li, Multi-GPU thermal lattice Boltzmann simulations using OpenACC and MPI, *Int. J. Heat Mass Transf.* 201 (2023) 123649.
- [29] E. Calore, J. Kraus, S.F. Schifano, R. Tripiccione, Accelerating lattice Boltzmann applications with OpenACC, in: Euro-Par 2015: Parallel Processing, Springer, 2015, pp. 613–624.
- [30] S. Blair, C. Albing, A. Grund, A. Jocksch, Accelerating an MPI lattice Boltzmann code using OpenACC, in: Proceedings of the Second Workshop on Accelerator Programming Using Directives, 2015, pp. 1–9.
- [31] E. Calore, A. Gabbana, J. Kraus, S.F. Schifano, R. Tripiccione, Performance and portability of accelerated lattice Boltzmann applications with OpenACC, *Concurr. Comput., Pract. Exp.* 28 (12) (2016) 3485–3502.
- [32] Q. Xiong, B. Li, G. Zhou, X. Fang, J. Xu, J. Wang, X. He, X. Wang, L. Wang, W. Ge, et al., Large-scale DNS of gas–solid flows on Mole-8.5, *Chem. Eng. Sci.* 71 (2012) 422–430.
- [33] K. Ma, M. Jiang, Z. Liu, Accelerating fully resolved simulation of particle-laden flows on heterogeneous computer architectures, *Particuology* 81 (2023) 25–37.
- [34] A.J. Ladd, Numerical simulations of particulate suspensions via a discretized Boltzmann equation. Part 1. Theoretical foundation, *J. Fluid Mech.* 271 (1994) 285–309.
- [35] Y. Chen, Q. Cai, Z. Xia, M. Wang, S. Chen, Momentum-exchange method in lattice Boltzmann simulations of particle–fluid interactions, *Phys. Rev. E* 88 (1) (2013) 013303.
- [36] B. Wen, C. Zhang, Y. Tu, C. Wang, H. Fang, Galilean invariant fluid–solid interfacial dynamics in lattice Boltzmann simulations, *J. Comput. Phys.* 266 (2014) 161–170.
- [37] H. Yoshida, M. Nagaoka, Multiple-relaxation-time lattice Boltzmann model for the convection and anisotropic diffusion equation, *J. Comput. Phys.* 229 (20) (2010) 7774–7795.
- [38] Z. Chai, T. Zhao, Lattice Boltzmann model for the convection-diffusion equation, *Phys. Rev. E* 87 (6) (2013) 063309.
- [39] J. Wang, D. Wang, P. Lallemand, L.-S. Luo, Lattice Boltzmann simulations of thermal convective flows in two dimensions, *Comput. Math. Appl.* 65 (2) (2013) 262–286.
- [40] D. Contrino, P. Lallemand, P. Asinari, L.-S. Luo, Lattice-Boltzmann simulations of the thermally driven 2D square cavity at high Rayleigh numbers, *J. Comput. Phys.* 275 (2014) 257–272.
- [41] A. Xu, L. Shi, H.-D. Xi, Lattice Boltzmann simulations of three-dimensional thermal convective flows at high Rayleigh number, *Int. J. Heat Mass Transf.* 140 (2019) 359–370.
- [42] A. Xu, X. Chen, H.-D. Xi, Tristable flow states and reversal of the large-scale circulation in two-dimensional circular convection cells, *J. Fluid Mech.* 910 (2021) A33.
- [43] A. Xu, B.-R. Xu, H.-D. Xi, Wall-sheared thermal convection: heat transfer enhancement and turbulence relaminarization, *J. Fluid Mech.* 960 (2023) A2.
- [44] M. Bouzidi, M. Firdaouss, P. Lallemand, Momentum transfer of a Boltzmann-lattice fluid with boundaries, *Phys. Fluids* 13 (11) (2001) 3452–3459.
- [45] W. Zhao, W.-A. Yong, Single-node second-order boundary schemes for the lattice Boltzmann method, *J. Comput. Phys.* 329 (2017) 1–15.
- [46] S. Tao, Q. He, B. Chen, X. Yang, S. Huang, One-point second-order curved boundary condition for lattice Boltzmann simulation of suspended particles, *Comput. Math. Appl.* 76 (7) (2018) 1593–1607.
- [47] L. Li, R. Mei, J.F. Klausner, Boundary conditions for thermal lattice Boltzmann equation method, *J. Comput. Phys.* 237 (2013) 366–395.
- [48] L. Li, R. Mei, J.F. Klausner, Lattice Boltzmann models for the convection-diffusion equation: D2Q5 vs D2Q9, *Int. J. Heat Mass Transf.* 108 (2017) 41–62.
- [49] J.R. Clausen, C.K. Aidun, Galilean invariance in the lattice-Boltzmann method and its effect on the calculation of rheological properties in suspensions, *Int. J. Multiph. Flow* 35 (4) (2009) 307–311.
- [50] Z.-G. Feng, E.E. Michaelides, The immersed boundary-lattice Boltzmann method for solving fluid–particles interaction problems, *J. Comput. Phys.* 195 (2) (2004) 602–628.
- [51] H. Huang, M. Sukop, X. Lu, Multiphase Lattice Boltzmann Methods: Theory and Application, John Wiley & Sons, 2015.
- [52] P. Lallemand, L.-S. Luo, Lattice Boltzmann method for moving boundaries, *J. Comput. Phys.* 184 (2) (2003) 406–421.
- [53] C. Peng, Y. Teng, B. Hwang, Z. Guo, L.-P. Wang, Implementation issues and benchmarking of lattice Boltzmann method for moving rigid particle simulations in a viscous flow, *Comput. Math. Appl.* 72 (2) (2016) 349–374.
- [54] T. Rosemann, B. Kravets, S. Reinecke, H. Kruggel-Emden, M. Wu, B. Peters, Comparison of numerical schemes for 3D lattice Boltzmann simulations of moving rigid particles in thermal fluid flows, *Powder Technol.* 356 (2019) 528–546.
- [55] A. Xu, L. Shi, T. Zhao, Thermal effects on the sedimentation behavior of elliptical particles, *Int. J. Heat Mass Transf.* 126 (2018) 753–764.
- [56] K. Walayat, Z. Zhang, K. Usman, J. Chang, M. Liu, Dynamics of elliptic particle sedimentation with thermal convection, *Phys. Fluids* 30 (10) (2018) 103301.
- [57] Z. Yu, N. Phan-Thien, Y. Fan, R.I. Tanner, Viscoelastic mobility problem of a system of particles, *J. Non-Newton. Fluid Mech.* 104 (2–3) (2002) 87–124.
- [58] D. Wan, S. Turek, Direct numerical simulation of particulate flow via multigrid FEM techniques and the fictitious boundary method, *Int. J. Numer. Methods Fluids* 51 (5) (2006) 531–566.
- [59] Z.-G. Feng, E.E. Michaelides, Inclusion of heat transfer computations for particle laden flows, *Phys. Fluids* 20 (4) (2008) 040604.
- [60] S. Tao, L. Wang, Q. He, J. Chen, J. Luo, A sharp interface immersed boundary-discrete unified gas kinetic scheme for fluid–solid flows with heat transfer, *Int. Commun. Heat Mass Transf.* 139 (2022) 106424.
- [61] C. Pan, J.F. Prins, C.T. Miller, A high-performance lattice Boltzmann implementation to model flow in porous media, *Comput. Phys. Commun.* 158 (2) (2004) 89–105.
- [62] C. Huang, B. Shi, Z. Guo, Z. Chai, Multi-GPU based lattice Boltzmann method for hemodynamic simulation in patient-specific cerebral aneurysm, *Commun. Comput. Phys.* 17 (4) (2015) 960–974.
- [63] H. Gao, H. Li, L.-P. Wang, Lattice Boltzmann simulation of turbulent flow laden with finite-size particles, *Comput. Math. Appl.* 65 (2) (2013) 194–210.
- [64] P. Bailey, J. Myre, S.D. Walsh, D.J. Lilja, M.O. Saar, Accelerating lattice Boltzmann fluid flow simulations using graphics processors, in: 2009 International Conference on Parallel Processing, IEEE, 2009, pp. 550–557.
- [65] C.-C. Ye, P.-J.-Y. Zhang, Z.-H. Wan, R. Yan, D.-J. Sun, Accelerating CFD simulation with high order finite difference method on curvilinear coordinates for modern GPU clusters, *Adv. Aerodyn.* 4 (1) (2022) 7.
- [66] J. Domínguez, A. Crespo, M. Gómez-Gesteira, J. Marongiu, Neighbour lists in smoothed particle hydrodynamics, *Int. J. Numer. Methods Fluids* 67 (12) (2011) 2026–2042.
- [67] N. Geneva, C. Peng, X. Li, L.-P. Wang, A scalable interface-resolved simulation of particle-laden flow using the lattice Boltzmann method, *Parallel Comput.* 67 (2017) 20–37.
- [68] W. Xian, A. Takayuki, Multi-GPU performance of incompressible flow computation by lattice Boltzmann method on GPU cluster, *Parallel Comput.* 37 (9) (2011) 521–535.
- [69] P.-Y. Hong, L.-M. Huang, L.-S. Lin, C.-A. Lin, Scalable multi-relaxation-time lattice Boltzmann simulations on multi-GPU cluster, *Comput. Fluids* 110 (2015) 1–8.
- [70] A. Tiribocchi, A. Montessori, G. Amati, M. Bernaschi, F. Bonaccorso, S. Orlandini, S. Succi, M. Lauricella, Lightweight lattice Boltzmann, *J. Chem. Phys.* 158 (10) (2023) 104101.