Interactive 3D simulation for fluid-structure interactions using dual coupled GPUs

Bob Zigon · Luoding Zhu · Fengguang
Song

Abstract The scope of this work involves the integration of high speed parallel computation with interactive, 3D visualization of the lattice-Boltzmann-based immersed boundary method for fluid-structure-interaction. An NVIDIA Tesla K40c is used for the computations while an NVIDIA Quadro K5000 is used for 3D vector field visualization. The simulation can be paused at any time step so that the vector field can be explored. The density and placement of streamlines and glyphs are adjustable by the user, while panning and zooming is controlled by the mouse. The simulation can then be resumed. Unlike most scientific applications in computational fluid dynamics where visualization is performed after the computations, our software allows for real-time visualizations of the flow fields while the computations take place. To the best of our knowledge, such a tool on GPUs for FSI does not exist. Our software can facilitate debugging, enable observation of detailed local fields of flow and deformation while computing, and expedite identification of 'correct' parameter combinations in parametric studies for new phenomenon. Therefore, our software is expected to shorten the 'time to solution' process and expedite the scientific discoveries via scientific computing.

Bob Zigon
Beckman Coulter, Indianapolis, IN 46268 Tel.: +1-317-508-6905
E-mail: robert.zigon@beckman.com

Luoding Zhu
Department of Mathematical Sciences, Indiana University-Purdue University Indianapolis,
IN 46202

Fengguang Song
Department of Computer Science, Indiana University-Purdue University Indianapolis, IN
46202

1 Introduction

The fluid structure interaction (FSI) problem involves the interplay of a solid structure with a surrounding fluid flow. Such problems are ubiquitous and include examples like the deformation of a fish's fin while swimming [2], the bending of the cilium in the kidney lumen in response to shear flow [15], or the strong wind effect on a skyscraper [16]. Because the interaction of an elastic solid and a viscous fluid in nature is non-linear, multi-physics, and multi-scale, an analytical solution is very rare. Instead, a computational approach is practically viable.

To date, many computational methods exist for numerical studies of problems involving flow-structure interactions, probably because of the complexity and diversity of real-world FSI problems and the limitations of mathematics and computer resources. Each method has its strength and weakness. Some of them are comparable. The choice is problem dependent. We shall not attempt to compare these methods here; instead we list some of them for readers' reference. These methods include immersed boundary (IB) methods [5,20], the immersed interface methods [21], blob-projection [22], immersed continuum [23], and immersed finite element [24], the Arbitrary Lagrangian Eulerian (ALE) [27], the fictitious domain method [25], the material point method [26], the level set method [72], and the front tracking method [71].

We choose to use the popular IB method originated by Peskin [28] for the fluid-structure interactions. The reason is that it is well tested, efficient, and allows a variety of fluid and solid solvers to be combined. Within the IB method, there too exists different versions. Examples include the original versions [4], the vortex-method version [29], the volume-conserved version [30], the adaptive mesh refinement version [31], the (formally) second-order versions [33,34], the multigrid version [35], the penalty version [36], the implicit versions [37–39,42,43], the generalized version for a thick rod [44], the stochastic version [45], the porous media version, the lattice-Boltzmann version [3,48,49, 60,61,53–56,59,63], the fluid-solute-structure interaction version [50], and the variable viscosity version [52].

In this work, we strive to realize the *lattice-Boltzmann based immersed boundary (LB-IB)* method developed by Zhu et al. [3] for general fluid-deformable-structure interactions. The lattice Boltzmann (LB) method [62,64–70] is a widely used alternative to traditional numerical methods for flow problems. It employs a meso-scale description and incorporates a velocity distribution function that obeys an approximate Boltzmann equation. Compared to conventional approaches for solving the flow problem, the LB method is relatively simpler to use, easier to handle complex rigid boundaries (e.g. porous media), and more convenient to incorporate additional physics into a model to simulate new flow phenomena, particularly in three dimensions. Therefore, our software can be easily extended to other situations such as FSI involving non-Newtonian fluids.

Another reason for choosing the LB-IB method is the inherent parallelism in both the lattice Boltzmann (LB) method and the immersed boundary

method, which makes them good candidates for parallel computing on GPUs [58]. In the LB method, with each node in the computing domain acting independently of its neighbors, the streaming and collision of the fluid particles maps elegantly on to the thousands of cores present on a GPU. In a similar manner, in the IB method, the computation of the forces and subsequent application to the immersed object is also a procedure where adjacent domain members can be computed in parallel.

The literature shows the existence of some work implementing the lattice Boltzmann and immersed boundary methods on GPUs. Valero, et al. [7] demonstrated the performance of the 2D algorithm on Intel CPU's and NVIDIA GPU's. The goal of their work was to investigate optimization strategies for heterogeneous architectures on two-dimensional domains. In the work of Mawson et al. [8], they too developed a GPU library. Like Valero, they focused on implementation and performance, but also investigated the application to 3D domains. Although Mawson et al. stated that real time 3D simulations are possible with GPU acceleration, they also found it was very difficult to identify the z-depth for the object they placed in the field. They did not perform visualization, nor did they address how complex internal obstacles could be handled in their work. In 2016, Wu [57] developed a GPU accelerated LB-IB simulator for a three dimensional ellipsoidal membrane. They focused on the creation of efficient code for computation on a single GPU. This is in contrast to our approach which distributed the computation and visualization across two distinct GPUs.

Our work explores the use of GPUs not only to accelerate computation, but also interactively visualize computational results in real time in three dimensions. We have created a number of C++ classes that simplify the implementation on GPUs. To compare and demonstrate the efficiency of using the GPU, we also implement the algorithms on CPUs using OpenMP. We model a 3D viscous flow past a deformable mesh fixed at its midline behind a circular rigid cylinder as an example of our work, however, our software implementation is generic and can be used for other FSI problems.

Visualization has become an essential part in engineering, research and business workflows. The current practice for a researcher in computational fluid dynamics is as follows: one executes a simulation and saves the data to storage, waits for the simulation to terminate, and then loads the data from storage for visualization. For large-scale real-world FSI problems, a simulation may take days or even weeks to finish on modern parallel computers. Furthermore, computational studies frequently perform many series of simulations with different combinations of problem-specific parameters (i.e. *parametric studies*). Due to the essential *non-linearity* of FSI problems, the "correct" choices that may lead to new phenomenon or discovery are typically not known beforehand. The real-time visualization may substantially help in this regard by identifying the uninteresting or incorrect combinations long before simulations are completed. For this reason, we have made it easy for the researcher to rotate and zoom the simulation during the computation. We also allow the user to view the vector field using three dimensional vectors or streamlines, as

well as changing their density and placement, to better observe the behavior of the flow.

Integrating visualization into the computing framework brings value in other areas. A domain with 256^3 nodes that simulates for 10^5 seconds would consume 12 terabytes of disk space if the simulation was saved. With our approach you simply rerun the simulation. Moreover, any disk I/O that is incurred during simulation would certainly slow down the overall execution time if every time step is written to disk. We've also found that the interactivity simplified debugging. When a logic error occurred, the source was frequently very obvious from the output.

Regarding visualization, software toolkits like the OpenGL Volumizer [17], ParaView [19] and VisIt [18] enable a user to interactively visualize the data after it has been computed. In our work, we integrate large-scale simulation with real-time visualization using two GPUs. An NVIDIA Tesla K40c is used for computations while an NVIDIA Quadro K5000 is used for 3D visualization by streamlines or vector glyphs.

To the best of our knowledge this paper makes the following three contributions. First, we present an efficient GPU implementation of the LB-IB method in three dimensions. Second, we create a set of software classes capable of supporting both CPUs and GPUs. Finally, we provide an integrated approach to realizing online FSI visualization using multiple GPUs that emphasizes human interaction with the simulator during computation. This capability facilitates code debugging, allows one to observe detailed local flow and deformation dynamics while computing, and expedites identification of 'correct' parameter combinations in parametric studies for new phenomenon. It is also expected to shorten the 'time to solution' process and expedite the scientific discoveries via scientific computing.

The remainder of the paper is organized as follows. Section 2 introduces the LB-IB method, including the mathematical formulation and its discretization. Section 3 describes the software design for the OpenMP and GPU hardware platforms. Section 4 gives implementation details. Section 5 presents our results and section 6 concludes with a discussion of future work.

2 The LB-IB method

2.1 The mathematical formulation

The lattice Boltzmann method originated from Boltzmann's kinetic theory of dilute gases. The fundamental concept is that fluids can be modeled as large collections of particles with random motions. The exchange of momentum and energy is achieved through particle collisions and particle streaming. The LB method is an alternative to traditional numerical methods such as the fast Fourier Transform, the projection method, and the particle in cell method for obtaining the solution to the viscous incompressible flow problem. In contrast to solving for macroscopic variables like velocity and pressure, the LB method

uses a mesoscopic approach that deals with a particle velocity distribution function $g(\mathbf{x}, \xi, t)$ defined on a Eulerian grid. Here $\mathbf{x}$ represents the spatial coordinate, ξ represents particle velocity, and t is time.

Different from the lattice Boltzmann method, the goal of the immersed boundary method is to model the interaction of a fluid with an elastic material. The elastic material is treated as part of the fluid in which additional forces are applied. The elastic material is tracked on a Lagrangian grid by following the material points. The configuration of these points is used to compute elastic forces which are applied to the nearby lattice points of the fluid.

Our overall approach for the LB-IB formulation follows that of [3]. The dimensionless form is formulated as follows:

$$\frac{\partial g(\mathbf{x}, \xi, t)}{\partial t} + \xi \cdot \frac{\partial g(\mathbf{x}, \xi, t)}{\partial \mathbf{x}} + \mathbf{f}_{ib}(\mathbf{x}, t) \cdot \frac{\partial g(\mathbf{x}, \xi, t)}{\partial \xi} = -\frac{1}{\tau}(g(\mathbf{x}, \xi, t) - g^{(0)}(\mathbf{x}, \xi, t)). \tag{1}$$

Bhatnagar-Gross-Krook (BGK) [6] are attributed to equation (1) which describes the motion of both the fluid and the immersed boundary in the context of the LB method. The quantity $g(\mathbf{x}, \xi, t) \, \mathbf{dx} \, \mathbf{d\xi}$ represents the probability of finding a particle at time t, located the interval $[\mathbf{x}, \mathbf{x} + \mathbf{dx}]$, while moving with velocity in the interval $[\xi, \xi + \mathbf{d\xi}]$. The term

$$-\frac{1}{\tau}(g(\mathbf{x}, \xi, t) - g^{(0)}(\mathbf{x}, \xi, t)) \tag{2}$$

in (1) is the BGK approximation to the complex collision operator in the Boltzmann equation, where τ is the relaxation time and $g^{(0)}(\mathbf{x}, \xi, t)$ is the Maxwellian distribution. The term $\mathbf{f}_{ib}(\mathbf{x}, t)$ is the force imparted by the immersed boundary to the fluid. This term is largely responsible for the unification of the LB and IB methods. As a result, there is no need to explicitly remesh the immersed boundary because the two methods are coupled by way of $\mathbf{f}_{ib}(\mathbf{x}, t)$.

The LB method requires the macroscopic variables fluid mass density, $\rho(\mathbf{x}, t)$, and the momentum, $(\rho\mathbf{u})(\mathbf{x}, t)$, which are defined in (3) and (4) as functions of the velocity distribution function $g(\mathbf{x}, \xi, t)$.

$$\rho(\mathbf{x}, t) = \int g(\mathbf{x}, \xi, t) \, d\xi \tag{3}$$

$$(\rho\mathbf{u})(\mathbf{x}, t) = \int g(\mathbf{x}, \xi, t)\xi \, d\xi \tag{4}$$

The Eulerian force density $\mathbf{f}_{ib}(\mathbf{x}, t)$ defined on the fixed Eulerian lattice is calculated from the Lagrangian force density $\mathbf{F}_{ib}(\boldsymbol{\alpha}, t)$ defined on the Lagrangian grid by equation (5),

$$\mathbf{f}_{ib}(\mathbf{x}, t) = \int \mathbf{F}_{ib}(\boldsymbol{\alpha}, t)\delta(\mathbf{x} - \mathbf{X}(\boldsymbol{\alpha}, t)) \, d\boldsymbol{\alpha} \tag{5}$$

where the function $\delta(\mathbf{x})$ is the Dirac δ-function. The Lagrangian force density $\mathbf{F}_{ib}$ is computed as follows:

$$\mathbf{F}_{ib}(\boldsymbol{\alpha}, t) = -\frac{\partial \mathcal{E}}{\partial \mathbf{X}} = -\frac{\partial(\mathcal{E}_s + \mathcal{E}_b)}{\partial \mathbf{X}} \tag{6}$$

In equation (6) the elastic potential energy density $\mathcal{E}$ consists of a stretching/compression component $\mathcal{E}_s$ and a bending component $\mathcal{E}_b$. These last two quantities are defined by equations (7) and (8) respectively.

$$\mathcal{E}_s = \frac{1}{2} K_s \int \int d\alpha_2 \, d\alpha_3 \int \left(\left| \frac{\partial \mathbf{X}(\boldsymbol{\alpha}, t)}{\partial \alpha_1} \right| - 1 \right)^2 d\alpha_1 \tag{7}$$

$$\mathcal{E}_b = \frac{1}{2} K_b \int \int d\alpha_2 \, d\alpha_3 \int \left(\left| \frac{\partial^2 \mathbf{X}(\boldsymbol{\alpha}, t)}{\partial \alpha_1^2} \right| \right)^2 d\alpha_1 \tag{8}$$

The variables α_1, α_2, α_3 are the three components of the Lagrangian variable $\boldsymbol{\alpha}$. In the case of an immersed surface, such as the flexible membrane in section 5, α_2 may be used to denote a fiber, α_1 to denote the arc length along the fiber, and α_3 is not used. K_s is the stretching/compression coefficient and K_b is the bending coefficient. Both constants are related to Young's modulus of the membrane.

The motion of the flexible membrane is described by a system of first order ordinary differential equations. Equation (9) describes the system.

$$\frac{\partial \mathbf{X}}{\partial t}(\boldsymbol{\alpha}, t) = \mathbf{U}(\boldsymbol{\alpha}, t) \tag{9}$$

$\mathbf{X}(\boldsymbol{\alpha}, t)$ is the Eulerian coordinate of the immersed membrane at time t whose Lagrangian coordinate is $\boldsymbol{\alpha}$. The immersed boundary velocity $\mathbf{U}(\boldsymbol{\alpha}, t)$ is interpolated from the fluid velocity $\mathbf{u}(\mathbf{x}, t)$ by using the same δ-function to apply the boundary force to the fluid. Equation (10) describes the immersed boundary velocity.

$$\mathbf{U}(\boldsymbol{\alpha}, t) = \int \mathbf{u}(\mathbf{x}, t) \delta(\mathbf{x} - \mathbf{X}(\boldsymbol{\alpha}, t)) \, d\mathbf{x} \tag{10}$$

2.2 Discretization

Equations 1 through 9 in the previous section are discretized on a uniform fixed Eulerian lattice for the fluid with a mesh width of h (the number of grid nodes is N_x, N_y and N_z in the x, y and z directions respectively). There is also a moving Lagrangian grid for the immersed boundary with an initial mesh width $\Delta \alpha_1 = \Delta \alpha_2 = h/2$.

The D3Q19 model seen in figure 1 is used to discretize the BGK equation. In this model, particles can enter and exit each lattice node along eighteen different directions. The nineteenth direction represents the particles remaining at rest at the node. The particle velocity space $\boldsymbol{\xi}$ is then discretized by a set of 19 velocities (see figure 1).

Fig. 1: D3Q19 model

Let $g_j(\mathbf{x}, t)$ be the distribution function along $\boldsymbol{\xi}_j$. A second order space and time discretization, equation (11), in a Lagrangian coordinate system is applied to derive the lattice Boltzmann equation that advances $g_j(\mathbf{x}, t)$ forward by one step.

$$g_j(\mathbf{x}+\boldsymbol{\xi}_j, t+1) = g_j(\mathbf{x}, t) - \frac{1}{\tau}\left[g_j(\mathbf{x}, t) - g_j^0(\mathbf{x}, t)\right] + \left[1 - \frac{1}{2\tau}\right]w_j\left[\frac{\boldsymbol{\xi}_j - \mathbf{u}}{c_s^2} + \frac{\boldsymbol{\xi}_j \cdot \mathbf{u}}{c_s^4}\boldsymbol{\xi}_j\right]\cdot\mathbf{f}_{ib}$$

(11)

Here w_j is a weight for direction $\boldsymbol{\xi}_j$.

The constant $c_s = 1/\sqrt{3}$ is the speed of sound for the model. The relaxation time τ is related to the dimensionless fluid viscosity ν by the equation $\nu = \frac{2\tau-1}{6}$. The fluid velocity $\mathbf{u}$ and the force $\mathbf{f}_{ib}$ are evaluated at time t.

The density $\rho(\mathbf{x}, t)$ and momentum $(\rho\mathbf{u})(\mathbf{x}, t)$ are related to $g_j(\mathbf{x}, t)$ at each node by

$$\rho(\mathbf{x}, t) = \sum_j g_j(\mathbf{x}, t),$$

(12)

$$(\rho\mathbf{u})(\mathbf{x}, t) = \sum_j \xi_j g_j(\mathbf{x}, t) + \frac{\mathbf{f}_{ib}(\mathbf{x}, t)}{2},$$

(13)

and the equilibrium distribution function g_j^0 is given by

$$g_j^0(\mathbf{x}, t) = \rho(\mathbf{x}, t)w_j\left[1 + 3\boldsymbol{\xi}_j \cdot \mathbf{u}(\mathbf{x}, t) + \frac{9}{2}(\boldsymbol{\xi}_j \cdot \mathbf{u}(\mathbf{x}, t))^2 - \frac{3}{2}\mathbf{u}(\mathbf{x}, t) \cdot \mathbf{u}(\mathbf{x}, t)\right]. \quad (14)$$

Assume the duration of the time step is set to 1. Let n be the time step index so that: $g^n = g(\mathbf{x}, \boldsymbol{\xi}, n)$, $\mathbf{X}^n(\boldsymbol{\alpha}) = \mathbf{X}(\boldsymbol{\alpha}, n)$, $\mathbf{u}^n = \mathbf{u}(\mathbf{x}, n)$, $p^n = p(\mathbf{x}, n)$ and $\rho^n = \rho(\mathbf{x}, n)$.

Let the flexible membrane be represented by a discrete collection of fibers whose Lagrangian coordinate is α_2. Let $\alpha_2 = q\Delta\alpha_2$, where q is an integer.

Now let each fiber be represented by a discrete collection of points whose Lagrangian coordinate is α_1. Let $\alpha_1 = m\Delta\alpha_1$, where m is an integer. The

"half integer" points are given by $\alpha_1 = (m + 1/2)\Delta\alpha_1$. For any function $\phi(\alpha)$, define operator $D_\alpha\phi$ be the centered difference operator with respect to α.

The stretching energy and corresponding force are discretized as,

$$\mathcal{E}_s = \frac{1}{2}K_s \sum_m (|D_{\alpha_1}\mathbf{X}| - 1)^2 \Delta\alpha_1 \tag{15}$$

$$= \frac{1}{2}K_s \sum_{m=1}^{n_f-1} \left(\frac{|\mathbf{X}_{m+1} - \mathbf{X}_m|}{\Delta\alpha_1} - 1 \right)^2 \Delta\alpha_1 \tag{16}$$

and

$$(\mathbf{F}_s)_l = \frac{K_s}{\Delta\alpha_1^2} \sum_{m=1}^{n_f-1} (|\mathbf{X}_{m+1} - \mathbf{X}_m| - \Delta\alpha_1) \frac{\mathbf{X}_{m+1} - \mathbf{X}_m}{|\mathbf{X}_{m+1} - \mathbf{X}_m|} (\delta_{m,l} - \delta_{m+1,l}). \tag{17}$$

Here $(\mathbf{F}_s)_l, l = 1, 2 \ldots, n_f$ is the Lagrangian force density $\mathbf{F}_s$ associated with node l. In a similar manner, the bending energy and corresponding force are discretized as,

$$\mathcal{E}_b = \frac{1}{2}K_b \sum_m |D_{\alpha_1} D_{\alpha_1}\mathbf{X}|^2 \Delta\alpha_1 \tag{18}$$

$$= \frac{1}{2}K_b \sum_{m=2}^{n_f-1} \left[\frac{|\mathbf{X}_{m+1} + \mathbf{X}_{m-1} - 2\mathbf{X}_m|^2}{(\Delta\alpha_1)^4} \right] \Delta\alpha_1 \tag{19}$$

and

$$(\mathbf{F}_b)_l = \frac{K_b}{\Delta\alpha_1^3} \sum_{m=2}^{n_f-1} (\mathbf{X}_{m+1} + \mathbf{X}_{m-1} - 2\mathbf{X}_m)(2\delta_{m,l} - \delta_{m+1,l} - \delta_{m-1,l}). \tag{20}$$

Here $(\mathbf{F}_b)_l, l = 1, 2 \ldots, n_f$ is the Lagrangian force density $\mathbf{F}_b$ associated with node l. n_f is the total number of grid points on the flexible membrane and $\delta_{k,l}$ is the Kronecker symbol.

The total Lagrangian force density is $\mathbf{F}(\boldsymbol{\alpha}, t) = \mathbf{F}_s(\boldsymbol{\alpha}, t) + \mathbf{F}_b(\boldsymbol{\alpha}, t)$. The two integral relations for equations (5) and (10) can now be discretized as

$$\mathbf{f}_{ib}^n(\mathbf{x}) = \sum_{\boldsymbol{\alpha}} \mathbf{F}^n(\boldsymbol{\alpha})\delta_h(\mathbf{x} - \mathbf{X}^n(\boldsymbol{\alpha})) \, \Delta\boldsymbol{\alpha} \tag{21}$$

and

$$\mathbf{U}^{n+1}(\boldsymbol{\alpha}) = \sum_{\mathbf{x}} \mathbf{u}^{n+1}(\mathbf{x})\delta_h(\mathbf{x} - \mathbf{X}^n(\boldsymbol{\alpha})) \, h^3. \tag{22}$$

Here the notation $\sum_{\boldsymbol{\alpha}}$ means that the sum with respect to $\boldsymbol{\alpha}$ is taken over all of the discrete collection of points. Similarly, $\sum_{\mathbf{x}}$ means that the sum with respect to $\mathbf{x}$ is taken over all discrete points of the form $\mathbf{x} = (ih, jh, kh)$. δ_h is an approximation of the Dirac δ-function. In the IB method, δ_h has the form,

$$\delta_h(\mathbf{x}) = h^{-3} \psi\left(\frac{x}{h}\right) \psi\left(\frac{y}{h}\right) \psi\left(\frac{z}{h}\right) \tag{23}$$

where h is the mesh spacing, $\mathbf{x} = (x, y, z)$.

See [5] for details regarding the choice of $\psi(r)$. With $\mathbf{U}^{n+1}(\boldsymbol{\alpha})$ known from equation (10), the flexible membrane motion equation is

$$\frac{\mathbf{X}^{n+1}(\boldsymbol{\alpha}) - \mathbf{X}^n(\boldsymbol{\alpha})}{\Delta t} = \mathbf{U}^{n+1}(\boldsymbol{\alpha}) \tag{24}$$

or

$$\mathbf{X}^{n+1}(\boldsymbol{\alpha}) = \mathbf{X}^n(\boldsymbol{\alpha}) + \mathbf{U}^{n+1}(\boldsymbol{\alpha}) \cdot \Delta t. \tag{25}$$

3 Software Design

In order to understand the benefits of using GPUs, we also implement the LB-IB method using OpenMP for comparison. Since the application is written in C++, some classes were designed to help with code reuse between the GPU and CPU implementations. Two major classes help the software to identify the natural hardware specific implementation points. Figures 2 and 3 convey the essence of those two classes. The Solver abstract base class in figure 2 shows seven primary methods, five of which are abstract, that must be implemented to solve the lattice Boltzmann equation (the remaining two dozen methods are helper functions that are omitted from the class for clarity). A similar approach is used for the ImmersedBoundary abstract base class in figure 3.
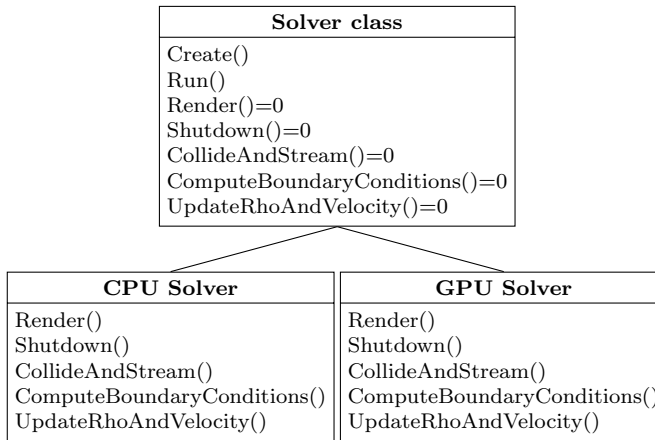

Fig. 2: Class diagram for the solver relationships

The overall program structure follows in listing 1. On lines 1 and 2, the Solver and ImmersedBoundary objects are first given an opportunity to initialize and create necessary data structures. The subclassing of the Solver allows
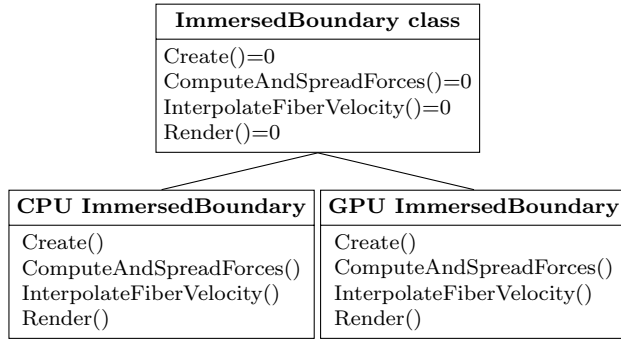
ImmersedBoundary class
Create()=0
ComputeAndSpreadForces()=0
InterpolateFiberVelocity()=0
Render()=0

CPU ImmersedBoundary	GPU ImmersedBoundary
Create()	Create()
ComputeAndSpreadForces()	ComputeAndSpreadForces()
InterpolateFiberVelocity()	InterpolateFiberVelocity()
Render()	Render()

Fig. 3: Class diagram for the immersed boundary relationships

```
1    Solver.Create()
2    IB.Create()
3
4    While Not Done
5    {
6        Solver.Run()−Compute field at time Tn
7            1. IB.Compute & Spread Forces              3%
8            2. Solver.Stream & Collide               61%
9            3. Solver.Compute Boundary Conditions      2%
10           4. Solver.Update Fluid Density & Velocity 16%
11           5. IB.Interpolate Fiber Velocity           2%
12           6. Solver.Compute Inlet & Outlet bc's     11%
13
14       Solver.Render()−Visualize field at time Tn     5%
15   }
```

Listing 1: Pseudo code for main program logic

the GPU to allocate the two large 4-dimensional matrices that store the node density values, $g(N_x, N_y, N_z, j)$, as well as the velocity and density fields, in RAM that is local to the GPU. In a similar fashion, the CPU solver allocates its matrices in host RAM. From that point forward, the Solver and Immersed-Boundary objects pass pointers to their large data structures when they need to be operated on.

After this initialization, the Run() method computes the field at time T_n on the Tesla card (line 6), while the Render() method displays and visualizes the generated field (line 14). The 6 steps associated with computing the field in lines 7 through 12 are implemented with the subclassed form of the IB (Immersed Boundary) class or the Solver (lattice Boltzmann) class.

The use of the abstract base class makes it easier to reason about the two different hardware platforms. Common variables like the dimensions of the domain, the Reynold number, or the current time step are naturally allocated in the base classes while hardware specific details are in the subclasses. In fact, this class design would make it fairly straightforward to evaluate the

performance of the LB-IB algorithm on a distinct accelerator such as the Xeon Phi from Intel.

The values on the right of each line represent the execution duration as a percentage of the total time needed to generate one frame for a domain size of 256^3. Our test case deals with the deformation of a flexible mesh that is composed of an orthogonal collection of fibers. Listing 1 suggests that the streaming and collision step in the lattice Boltzmann method is the most expensive function to implement. For our problem, the computation of the fluid forces and the movement of the fibers required just 3% of *the frame duration*. The frame duration is the amount of time required for the computing GPU to advance the LB-IB algorithm, plus the amount of time required for the visualization GPU render the three dimensional scene. We expect more complex immersed objects will consume more of the computing time. For example, in an extreme case, if the number of nodes in the immersible structure is comparable to the number of fluid grid nodes, the computation of the interactive forces may then become comparable to or even dominate the execution time.

4 Implementation of LB-IB and visualization on CPU and GPU

4.1 OpenMP details on CPU

Given the challenges associated with writing and debugging GPU code, we decided to implement the OpenMP version first. This approach helped us identify and solve race conditions and performance issues before we moved to the GPU. It also helped us recognize the two base classes in figures 2 and 3.

The primary difference between the two implementations lies in the data parallel approach used on the GPU. On the GPU we typically launched one thread for each node in the 3D domain that required computing. The equivalent implementation with OpenMP on CPU requires 3 nested for-loops. However, the outer most loop (for the Z axis) is preceded by the "#pragma omp parallel" directive to request distribution of the work across multiple CPU cores.

Many of the lessons learned while implementing the OpenMP code could be applied to the GPU code. For example, the GPU code, like the OpenMP code, merged the collision and streaming steps into a single step, thereby eliminating a large block of memory reads and writes. In addition, the AOS (Array of Structures) implementation of the fluid velocity field that negatively impacted the OpenMP code also affected the GPU code. The cache friendly solution was to switch to an SOA (Structure of Arrays) approach which minimized cache line reloads on the CPU. On the GPU, the SOA approach effectively reduced the number of memory transactions by a factor of 3. For example, when a warp of 32 cores generated 32 memory addresses to access the X component of the fluid velocity, the SOA approach guaranteed that those 256 bytes were physically contiguous in RAM. In contrast, the AOS approach distributed

those same 32 memory addresses across 768 bytes because the stride between logically adjacent X components was now 24 instead of 8.

Although C++ natively supports three dimensional arrays, we wrapped our arrays in a class so that subscript checking could be enabled in the debug build of the application. This greatly simplified the search for errant logic that occasionally indexed before the beginning or past the end of an array. Given that we consciously accessed the domain in Z-Y-X axis order, we could then insure that accesses to sequential nodes were physically contiguous in memory which again further minimized cache line reloads on the CPU.

4.2 GPU details

The GPU used for implementing the LB-IB algorithms and computing the 3D vector field was the Tesla K40c. This card features 2880 CUDA cores, 1.4 TFLOPS double precision (DP), and 12 GB of RAM. The development environment consisted of CUDA 7.5 and Visual Studio 2013 under Windows 7/64.

Computation of the bending and stretching forces (that accompany the spreading of the forces) in line 7 of listing 1 was a straight forward implementation of equations (17) and (20). A data parallel approach was taken where one GPU thread is dedicated to each point on the Lagrangian grid with synchronization primitives inserted as appropriate.

The subsequent spreading of the forces from the fibers to the fluid in equation (21) is described in kernel listing 2. First, notice in lines 20 through 25 how all of the accesses to global memory on the GPU are initiated as early as possible in the kernel. This aids in filling the memory controller pipeline and minimizes stalls later in the code due to unavailable operands.

Next, from the host computer's perspective, we simply launch as many GPU threads as there are fibers and points per fiber. Each thread will then be responsible for accessing the points from the Eulerian grid and accumulating them from the neighboring fiber nodes on the Lagrangian grid. The potential problem with this approach is that a race condition can occur. Fortunately, modern GPUs have an atomic add instruction that makes this accumulation indivisible as shown on lines 58 through 60 of listing 2.

Finally, equation (21) initially concerned us from a performance standpoint because the Eulerian force $\mathbf{f_{ib}}$ is computed by way of a smoothing function $\delta_h(\mathbf{x})$ that accesses the $4 \times 4 \times 4$ cube of values around each node of the Lagrangian grid. With $\delta_h(\mathbf{x})$ defined by equations (23), it becomes apparent that the non-coalesced nature of the 64 coefficients may negatively impact the function. However, with line 7 of listing 1 consuming 3% of the execution time, we deferred further analysis until the number of Lagrangian grid points significantly increased.

The lattice Boltzmann method consists of a collision step, a streaming step, and a boundary computation step. Our first implementation on the GPU followed this sequence. However, the latency associated with reading or writing

```
1   __global__
2   void SpreadForcesKernel(
3
4     double *fx, double *fy, double *fz,
5     double *filxn1, double *filyn1, double *filzn1,
6     double *ffx, double *ffy, double *ffz,
7     double dx, double dy, double dz,
8     int fffwidth,   // width of ff? 3D matrices
9     int fffheight,  // height of ff? 3D matrices
10    int NumberOfFibers,
11    int PointsPerFiber)
12  {
13  const int kf = blockIdx.x*blockDim.x + threadIdx.x;
14  const int jf = blockIdx.y*blockDim.y + threadIdx.y;
15
16  if (jf >= NumberOfFibers) return;
17  if (kf >= PointsPerFiber) return;
18
19  const int kfjfinx    = kf + jf*PointsPerFiber;
20  const double filxn1t = filxn1[kfjfinx];
21  const double filyn1t = filyn1[kfjfinx];
22  const double filzn1t = filzn1[kfjfinx];
23  const double fxt     = fx[kfjfinx];
24  const double fyt     = fy[kfjfinx];
25  const double fzt     = fz[kfjfinx];
26
27  const double cx = 3.1415926535  / (2.0 * dx);
28  const double cy = 3.1415926535  / (2.0 * dy);
29  const double cz = 3.1415926535  / (2.0 * dz);
30  const double cf = 1.0 / (64.0 * dx*dy*dz);
31  const double K0 = cf;
32
33  const int istart = static_cast<int>(floor(filxn1t/dx-2)+1);
34  const int jstart = static_cast<int>(floor(filyn1t/dy-2)+1);
35  const int kstart = static_cast<int>(floor(filzn1t/dz-2)+1);
36
37  const int istop  = istart + 4;
38  const int jstop  = jstart + 4;
39  const int kstop  = kstart + 4;
40
41  for (int i = istart; i < istop; i++)
42  {
43    const double rx = dx*static_cast<double>(i)-filxn1t;
44    const double K1 = (1.0 + COS(cx*rx)) * K0;
45
46    for (int j = jstart; j < jstop; j++)
47    {
48      const double ry   = dy*static_cast<double>(j)-filyn1t;
49      const double K1K2 = K1*(1.0 + COS(cy*ry));
50
51      for (int k = kstart; k < kstop; k++)
52      {
53        const double rz = dz*static_cast<double>(k)-filzn1t;
54        const double K3 = ((REAL)1.0 + COS(cz*rz))*K1K2;
55
56        const int inx  = i + j*fffwidth + k*fffwidth*fffheight;
57
58        atomicAdd(ffx+inx, fxt*K3);
59        atomicAdd(ffy+inx, fyt*K3);
60        atomicAdd(ffz+inx, fzt*K3);
61      }
62    }
63  }
64  }
```

Listing 2: CUDA kernel for spreading forces

GPU memory can take between 400 and 600 clock cycles. As such, it is in our interest to minimize redundant reads and writes. In our second GPU implementation, we merged the collision and streaming steps into a single step, thereby eliminating the extra reading and writing of $8 \times 19 \times 256^3$ bytes. This resulted in a *two fold increase* in the execution speed of the method, which further reinforced our belief that the LB method is memory bound and not arithmetic bound.

Line 10 of listing 1 shows that updating the fluid density and velocity consumed about 16% of the GPU time for one frame. Equations (12) and (13) describe the process. Implementing the two equations directly would result in inefficient GPU code. The GPU will block only when an operand isn't available. The better solution is to read the 19 distribution values into an array to fill the memory controller's read pipeline. The final CUDA kernel is described in listing 3. Lines $28 - 33$ focus on performing all memory accesses. Lines $40 - 51$ focus on the actual computation which are not likely to stall because the memory based operands should now be in registers.

Line 11 of the pseudo code in listing 1 describes the interpolation of the fiber velocity in terms of equation (22). When combined with equation (25), the fibers are repositioned to their new points in space. Equation (22), like equation (21), also concerned us from a performance standpoint because the Lagrangian velocity field $\mathbf{U}$ is computed by way of $\delta_h(\mathbf{x})$ that accesses the $4 \times 4 \times 4$ cube of values around each node of the Eulerian grid. The access pattern is nearly identical to that used in lines 41 through 63 of the SpreadForcesKernel in listing 2. However, with line 11 of Listing 1 consuming 2% of the execution time, we again deferred optimization until the number of Lagrangian grid points significantly increased.

4.3 Visualization details

In our software implementation, the Tesla GPU card generates the velocity and pressure fields at time T_n and then the Quadro GPU card uses custom GLSL shaders [9] to generate the display. The code was explicitly designed to separate the two functions. From Listing 1, it seemed that with visualization taking only 5% of the frame time, the dual GPU approach probably wasn't necessary for this example. However, as our immersed objects become more complex and more numerous, we have an architecture in place *that will allow us to overlap the computing and visualization phases*. This will amount to the Tesla card computing the field at time T_n while the Quadro visualizes time T_{n-1}. In addition, as we anticipate to implement isosurfaces or stream ribbons to visualize vortices [12] in the future, these calculations will be performed exclusively on the Quadro card.

All of the 3D graphics were implemented using OpenGL. OpenGL is a cross platform API for rendering 2D and 3D vector graphics. The original API is called the Direct Mode API. It is simple to use if one is familiar with 3D graphics concepts. Unfortunately, the Direct Mode API is poorly matched

```
1   __global__
2   void UpdateRhoAndVelocityKernel(
3      double *fin,
4      double *UField, double *VField, double *WField,
5      double *Rho,
6      double *ffx, double *ffy, double *ffz,
7      double dt,
8      double gl,
9      int width,
10     int height,
11     int zdim)
12  {
13      const int x = GetXIndex();
14      const int y = GetYIndex();
15      const int z = GetZIndex();
16
17      if (x < 2 || x >= width-2) return;
18      if (y < 2 || y >= height-2) return;
19      if (z < 2 || z >= zdim-2) return;
20
21      const int inx     = x + width*(y + z*height);
22      const int NumCells= width*height*zdim;
23      double F[19];
24
25      for (int i = 0; i < 19; i++)
26          F[i] = fin[inx + i*NumCells];
27
28      double ffxt = ffx[inx];
29      double ffyt = ffy[inx];
30      double ffzt = ffz[inx];
31      double SumF = 0.0;
32      double SumX = 0.0;
33      double SumY = 0.0;
34      double SumZ = 0.0;
35
36      for (int i = 0; i < 19; i++)
37      {
38          double Q;
39          SumF += (Q = F[i]);
40          SumX += Xi[i].x*Q;
41          SumY += Xi[i].y*Q;
42          SumZ += Xi[i].z*Q;
43      }
44
45      double X1   = (SumX+0.5*dt*ffxt)/SumF;
46      double Y1   = (SumY+0.5*dt*ffyt)/SumF;
47      double Z1   = (SumZ+0.5*dt*(ffzt+SumF*gl))/SumF;
48      Rho[inx]    = SumF;
49      UField[inx] = X1;
50      VField[inx] = Y1;
51      WField[inx] = Z1;
52  }
```

Listing 3: CUDA kernel for updating fluid density and velocity

for modern GPU hardware. Modern hardware prefers to be handed blocks of thousands or millions of vectors at time. As a result, starting with OpenGL 3.0, the Direct Mode API was deprecated in preference for the new API. The new API is more challenging to use. However, some of our simple tests showed that for a given GPU card like a Quadro K5000, the drawing rate was minimally 10 times faster than the Direct Mode API. In this work, we chose to implement the graphics using as much of the new API as possible. As mentioned earlier, we took this implementation path in anticipation of using a large number of complex, immersed boundary objects in our future work.

We follow [1] on visualization of vector fields. The author discussed how vector glyphs can be used as trajectories of imaginary particles that are released into the vector field over a short period of time δt. This was one of the two visualization techniques that we have realized. The author then describes a broader set of tools known as *stream objects* whose purpose is to utilize those same trajectories over a longer period of time.

The other visualization technique we implemented is the *streamline* visualization. For a time independent vector field, a streamline is a curved path starting from a given point $\mathbf{x_0}$ which is tangent to $\mathbf{v}$, the vector field. If a streamline is modeled as a parametric function $S(\tau) = \mathbf{x}(\tau)$, where τ represents the arc-length coordinate along the curve, then a streamline obeys the equation

$$\frac{d\mathbf{x}(\tau)}{d\tau} \times \mathbf{v}(\mathbf{x}(\tau)) = 0. \tag{26}$$

This can also be expressed as the following ODE,

$$\frac{d\mathbf{x}(\tau)}{d\tau} = \frac{\mathbf{v}(\mathbf{x}(\tau))}{|\mathbf{v}(\mathbf{x}(\tau))|} \tag{27}$$

with the initial condition $\mathbf{x}(s = 0) = \mathbf{x_0}$ and the constraint $s \in [0, S_{max}]$. When equation (27) is integrated over τ from 0 to s, we have the equation

$$\mathbf{x}(s) = \mathbf{x}(0) + \int_0^s \frac{\mathbf{v}(\mathbf{x}(\tau))}{|\mathbf{v}(\mathbf{x}(\tau))|} \, d\tau, \tag{28}$$

with $\mathbf{x}(s = 0) = \mathbf{x_0}$.

Equation (27) can be implemented using Euler's method as follows,

$$\int_0^T \frac{\mathbf{v}(\mathbf{x}(t))}{|\mathbf{v}(\mathbf{x}(t))|} \, dt \approx \sum_{i=1}^{N=T/\delta t} \frac{\mathbf{v}(x_i)}{|\mathbf{v}(x_i)|} \delta t, \quad \mathbf{x}_i = \mathbf{x}_{i-1} + \mathbf{v}_{i-1}\delta t. \tag{29}$$

However, the global error of Euler's method, $\mathcal{O}(\delta t)$, suggests we pursue a better integrator. The Runge-Kutta 2 (i.e., RK2) and the RK4 integrators have global errors of $\mathcal{O}(\delta t^2)$ and $\mathcal{O}(\delta t^4)$ respectively. Therefore, we have implemented both in the event we suspected drift in the placement of the streamlines.

Our visualization subsystem is capable of generating the 3D vector field using glyphs or streamlines [10–14]. The glyph based approach draws an arrow at a point in 3D space that is oriented with the vector flow as shown in

figure 4. This approach is easy to implement but suffers from the problem of visual clutter. Even with the ability to rotate and zoom our 3D domain during simulation it can be difficult to discern details of the underlying field.
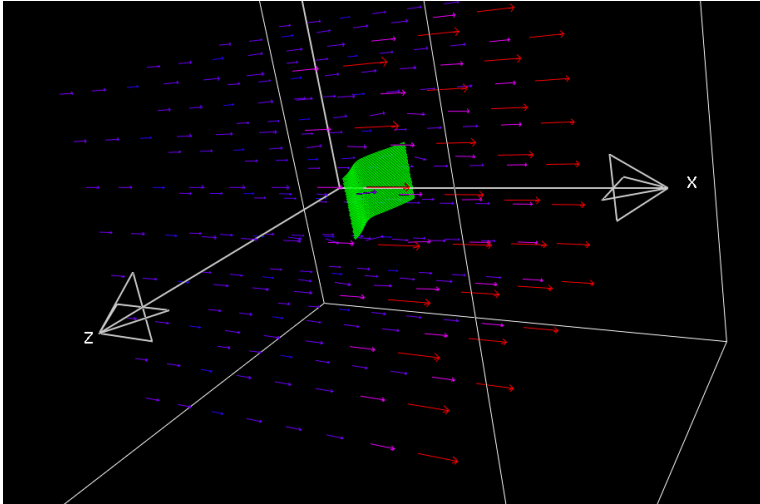

Fig. 4: User interface with glyphs

Listing 4 describes the algorithm. Lines 6 through 10 scan the velocity field and search for the longest vector that is a fluid node. Lines 13 through 19 then loop one more time over the velocity field to find the *head* and *tail* of each fluid node vector. After the *head* is normalized an arrow is drawn.

Streamlines, on the other hand, show where the vector flow has come from and where it is going to. By changing the length of the streamlines during the computation we can more easily accentuate features such as vortices. The streamlines are always tangent to the vector field, and fluid never crosses a streamline. Figure 5 shows a screen shot of the user interface while the simulation is running and the field is viewed with streamlines.

Listing 5 describes the streamline algorithm. During the initialization of the application a number of 3D points called *seeds* are calculated. These seeds represent the starting point for the stream lines. The seeds are uniformly distributed across the inlet plane of the 3D domain. When the *DrawStreamLines* procedure is called, the logic selects a seed and then performs an RK2 integration over *PathLineLength* points. This essentially implements equation (28). When the *TraceRK2* procedure returns, the *StreamLine* variable contains a collection 3D points that trace the path. Line 6 then calls the *Draw* procedure to display the path.

The *TraceRK2* procedure performs the actual integration. The procedure keeps track of the *PathLineLength*. Line 14 traverses the vector field until the *PathLineLength* becomes too long. Otherwise line 15 calls *InterpolateVelocity*

```
1   DrawGlyphs(UData, VData, WData, IsBoundaryNode, Width, Height, ZDim)
2   {
3     MaxLen = 0
4
5     // Find the length of the largest vector
6     for Z = 1 to ZDim
7       for Y = 1 to Height
8         for X = 1 to Width
9           if IsBoundaryNode(X,Y,Z) = 0
10            MaxLen = Max(MaxLen, VectorLength(X,Y,Z))
11
12    // Draw the vectors properly scaled
13    for Z = 1 to ZDim
14      for Y = 1 to Height
15        for X = 1 to Width
16          if IsBoundaryNode(X,Y,Z) = 0
17            Head = [UData(X,Y,Z), VData(X,Y,Z), WData(X,Y,Z)] / MaxLen
18            Tail = X,Y,Z
19            DrawArrow(Head,Tail)
20  }
```

Listing 4: Pseudo code to draw vector glyphs
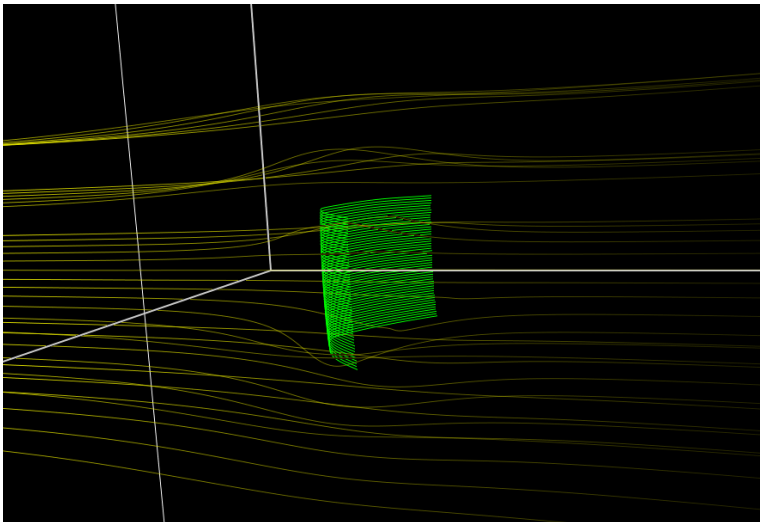
Fig. 5: User interface with streamlines

which performs trilinear interpolation of the vector field at the current point. The interpolated value is appended to the *StreamLine* on line 18. Line 20 checks the length of the interpolated velocity and exits the procedure if the magnitude is smaller than 10^{-5}. The algorithm updates the *CurrentPoint* which advances it on the field. Finally the *PathLineLength* is updated by the step size.

```
1   DrawStreamLines(}
2   {
3     for  i  =  1  to  NumberOfSeeds
4       aSeed  =  Seeds[i]
5       TraceRK2(aSeed ,  StreamLine )
6       Draw( StreamLine )
7   }
8
9   TraceRK2 (aSeed ,  StreamLine )
10  {
11    PathLineLength  =  0
12    CurrentPoint    =  aSeed
13
14    for  i =0, PathLineLength<MaxPathLineLength ,  i++
15        bool  InSide  =  InterpolateVelocity ( CurrentPoint ,  vA)
16        if  not  InSide  return
17
18        StreamLine . add ( CurrentPoint )
19
20        if  Length (vA)  <  1E−5  return
21
22        vecB  =  CurrentPoint  +  vA∗StepSize
23        bool  InSide  =  InterpolateVelocity ( vecB ,  vB)
24
25        if  Length (vB)  <  1E−5  return
26
27        v  =  (vB  +  vA)  ∗  0.5
28
29        CurrentPoint    =  CurrentPoint  +  v∗StepSize
30        PathLineLength  =  PathLineLength  +  StepSize
31  }
```

Listing 5: Pseudo code to draw stream lines

This visualization technique has several useful properties. First, it is easy to implement which makes it amenable to implementation on the host CPU or the graphics GPU. Second, if the hardware permits, all of the stream lines can be computed in parallel.

As mentioned earlier, the placement and density of the glyphs and stream lines can be changed while the simulation is running. In addition, one key stroke will pause the simulation while another key stroke will toggle between the two visualization techniques. This type of functionality supports the exploratory nature of the application.

5 Results

The results of our real-time GPU implementation of the 3D LB-IB method is illustrated by the example problem shown in figure 6: a 3D viscous incompressible fluid flows around a circular rigid cylinder with a tethered flexible mesh placed behind. The fluid enters from the left, flows around a cylinder,

interacts with the flexible mesh, and then exits on the right. The remaining four faces of the domain implement a no-slip boundary condition. The flexible mesh is tethered in space along a vertical line that divides the mesh in half. The flexible mesh consists of 52 fibers oriented vertically, 52 fibers oriented horizontally and 103 points per fiber. The simulation shows how the mesh folds in half under interaction with the flow field and the complicated flow patterns behind the mesh.


Fig. 6: Cartoon of the 3D field flowing around a cylinder and interacting with a flexible mesh.

To measure the performance we started with 3 configurations (a single CPU core, 8-CPU cores under OpenMP, and a single GPU). Each configuration was executed for our example problem in figure 6 with a grid size of 128, 160, 192 and 256. Each test ran for 100 time steps. We then computed the average duration per time step by dividing the execution time by 100.

The execution times for the double precision implementations are in figure 7. For a given grid size, the GPU is approximately 10 times faster than 8 cores using OpenMP, and the 8 core OpenMP version is approximately 8 times faster than a single core. As shown earlier, the rendering performed by the Quadro GPU consumed only 5% of the total compute duration for each frame. In this example, we explored how the flow around a cylinder affects the flexible mesh. However, our framework can be adapted to explore almost all FSI problems involving a viscous incompressible fluid and an immersed elastic structure, including the flow around a sphere, a cube or a torus.

The simulator has numerous parameters that control the behavior of the LB-IB method. These parameters include the domain dimensions N_x, N_y, N_z, the number of fibers n_f, initial fluid field velocity, the Reynolds number Re, fiber stretching coefficient K_s and the fiber bending rigidity K_b. All of these

---

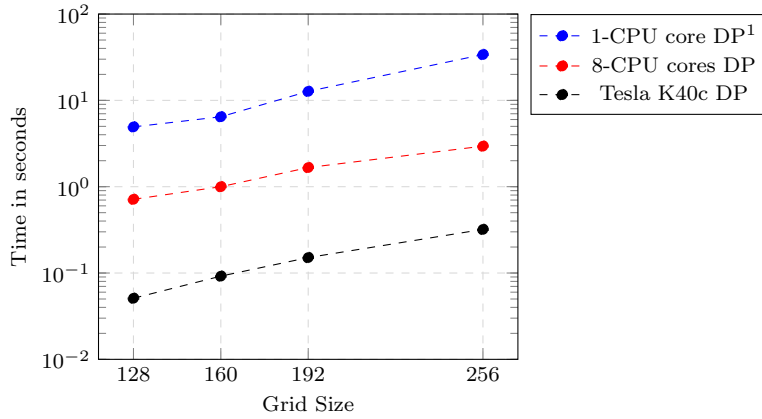[1] Lenovo D30, 8 core E5-2609@2.4GHz, 32GB RAM, Windows 7/64

Fig. 7: Average time step execution times for 104 fibers and 103 points/fiber

parameters can be modified for experimentation purposes or for parametric studies. Table 1 shows the parameters for the four simulations we performed.

Time step	N_x	N_y	N_z	n_f	Initial velocity	Re	K_s	K_b
100000	256	256	256	104	0.03	150	0.0004	0.0005
100000	256	256	256	104	0.03	150	0.0004	0.0015
100000	256	256	256	104	0.03	150	0.0004	0.0032
100000	256	256	256	104	0.03	150	0.0004	0.0050

Table 1: Simulation parameters for figures 9, 10, 11 and 12.

Figures 9, 10, 11 and 12 demonstrate the mesh final position, the mesh shape, and the flow field (by streamlines) at time 100,000 (in lattice Boltzmann units). These figures reveal that with a larger bending coefficient the mesh is deformed less at the final equilibrium state, and the flow patterns behind the mesh become more and more chaotic and complicated.

To highlight the interactive features of our software, figures 13, 14, 15 and 16 display the time evolution of mesh deformation and the flow field with some of the user interface elements present. The bounding box is drawn to help the user understand the extent of the domain. Arrows for the X, Y and Z axis are drawn to better understand the orientation of the simulation. Finally, a menu of options (figure 8) is presented (along with the current state of the simulation) so that the user can interactively change

- the vector glyphs or the stream lines
- the use of an RK2 vs RK4 stream line integrator
- the ability to pause the simulation on a time step
- the ability to single step the simulation.

The view point of the user is changed by simply dragging the mouse in the X and Y direction so that the bounding box is rotated. The simulation parameters are shown in table 2.

Validation of numerical codes like LB-IB can be performed through the use of convergence checking. In this approach, a series of gradually refined grids ensures that the results are reliable and accurate. In our case, we needed to prove that our solutions were valid *and* that the CPU and GPU versions were sufficiently identical. Given that floating point arithmetic is not associative, it can be challenging to produce CPU results that match the GPU. Since we had access to the code produced by Zhu [3] (that had been previously validated with convergence checking), we constantly compared our results with theirs.

Our first step was to develop our 1-core OpenMP code and configure it with parameters identical to [3]. Then, as each new version of the code was produced, we would compare all 507,904 values in our $124 \times 64 \times 64$ test domain with that of Zhu's at time step 1, 100, and 1,000. Our criteria was that the results needed to match to three decimal places (the codes were implemented using double precision floating point). Since our OpenMP implementation was carefully thought through, our 1-core, $4-$core and 8-core versions were identical. This was relatively easy to achieve because the code accessed and operated on its parameters in the same order. In fact, when the three versions didn't match it was usually due to a logic error or race condition.

With these three versions in place, we then implemented the GPU version. Again, with each new version of the GPU code, we would compare the values from our $124 \times 64 \times 64$ domain to those produced by the OpenMP version. We again used a criteria that our results needed to match to three decimal places.

Timestep	N_x	N_y	N_z	n_f	Initial velocity	Re	K_s	K_b
4000	256	256	256	104	0.03	150	0.0004	0.00008
8000	256	256	256	104	0.03	150	0.0004	0.00008
12000	256	256	256	104	0.03	150	0.0004	0.00008
16000	256	256	256	104	0.03	150	0.0004	0.00008

Table 2: Simulation parameters for figures 13, 14, 15 and 16

6 Summary and Future Work

This paper presents a dual GPU interactive implementation of the lattice Boltzmann based immersed boundary method for fluid-structre-interaction problems in three dimensions. The implementation is demonstrated on a 3D, viscous, incompressible flow past a deformable mesh behind a rigid cylinder of circular section. But our software works for almost all FSI problems invovling an incompressibel viscous fluid and an elastic structure (rigid or deformable). To the best of our knowledge, our software is the first GPU tool integrating

Fig. 8: The menu in the interface

visualization and computing in CFD. It is expected to shorten time-to-solution and speed up scientific discoveries in the FSI field.

Our single-CPU core, 8-CPU core and Tesla implementations are compared from a performance perspective using double precision arithmetic. Our simulations demonstrate an 80 fold improvement of the single GPU over the single CPU core. Our base class design for the LB-IB algorithm greatly simplifies the hardware specific implementations. A single application was created that addressed both hardware platforms for the purpose of performance monitoring and comparison.

Our GPU implementation of the smoothed Dirac delta function for identifying influence and dependent domains of a Lagrangian point on the structure (for force spreading and velocity interpolation) may need to be optimized when the immersed body becomes so complex that the number of Lagrangian structure grid points are comparable to the Eulerian fluid grid points. Otherwise the software performance may deteriorate. This is a nice future work.

Our dual GPU approach shows that there is a large number of visualization cycles available for more complex immersed objects. As another future work, we plan to explore distributing the computation across multiple GPUs, expanding the visualization capabilities to include isosurfaces of quantities such as vorticity, volume rendering on the visualization GPU, and overlapping the computation with the rendering.
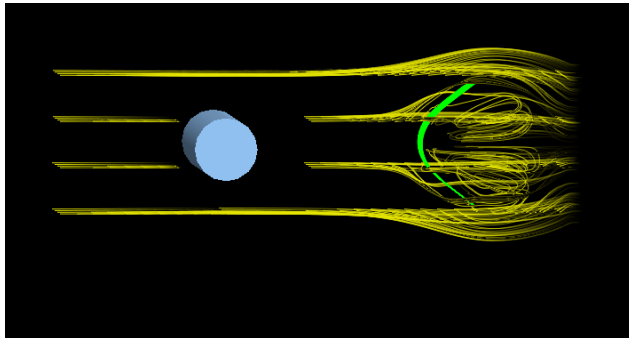

Fig. 9: $K_b = 0.0005$ at time step $100,000$

Fig. 10: $K_b = 0.0015$ at time step $100,000$

Fig. 11: $K_b = 0.0032$ at time step $100,000$

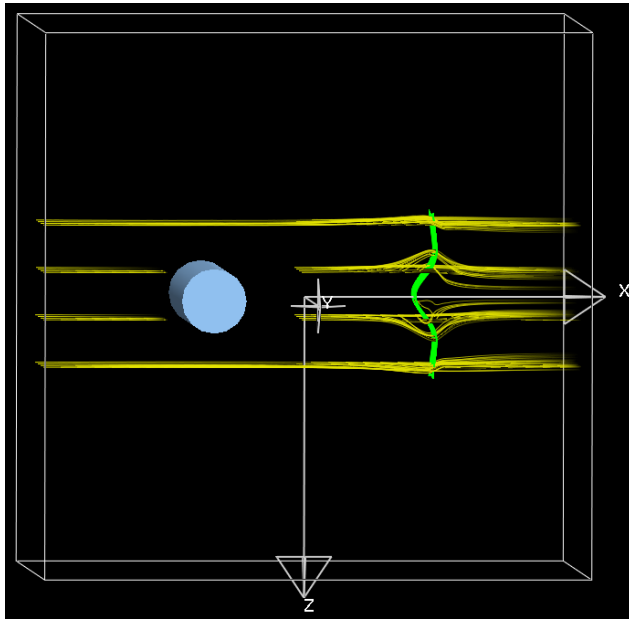
Fig. 12: $K_b = 0.0050$ at time step $100,000$

Fig. 13: $K_b = 0.0005$ at time step $4,000$

Fig. 14: $K_b = 0.0005$ at time step $8,000$

Fig. 15: $K_b = 0.0005$ at time step $12,000$

Fig. 16: $K_b = 0.0005$ at time step $16,000$

References

1. Alexandru C. Telea. *Data Visualization Principles and Practice, 2nd edition.* CRC Press, 2015.
2. Fang-Bao Tian, Haoxiang Luo, Luoding Zhu, Xi Yun Lu, *Interaction between a flexible filament and a downstream rigid body*, Physical Review E **82**, (2010), 026301.
3. Luoding Zhu, Guowei He, Shizhao Wang, Laura Miller, Xing Zhang, Quian You, Shiaofen Fang, *An immersed boundary method based on the lattice Boltzmann approach in three dimensions with application*, Computers and Mathematics with Applications **61**, (2011), 3506–3518.
4. C.S. Peskin, *Flow patterns around heart valves; a numerical method*, J. Comput. Phys., **25**, (1977), 220.
5. C.S. Peskin, *The immersed boundary method*, Acta Numer., **11**, (2002), 409.
6. P.L. Bhatnagar, E.P. Gross, M. Krook, *A model for collision processes in gases, I; small amplitude process in charged and neutral one-component system*, Phys. Rev., **94**, (1954), 511.
7. Pedro Valero-Lara, Francisco D. Igual, Manuel Prieto-Matas, Alfredo Pinelli, Julien Favier, *Accelerating fluid-solid simulations (Lattice-Boltzmann & Immersed-Boundary) on heterogeneous architectures*, Journal of Computational Science, Elsevier, **10**, (2015), 249–261.
8. Mark Mawson, Pedro Valero Lara, Julien Favier, Alfredo Pinelli, Alistair Revell, *Fast fluid-structure interaction using lattice Boltzmann and Immersed Boundary Methods*, NVIDIA GPU Conference, 2013.
9. Mike Bailey, Steve Cunningham, *Graphics Shaders Theory and Practice, 2nd Edition*, CRC Press, 2012.
10. Hongfeng Yu, Chaoli Wang, Kwan-Liu Ma, *Parallel Hierarchical Visualization of Large Time Varying 3D Vector Fields*, Supercomputing 2007.
11. Chenyang Xu, Jerry L. Prince, *Snakes, Shapes, and Gradient Vector Flow*, IEEE Trans. Image Process., **7**, (1998), 359–369.
12. Daniel Weiskopf, *GPU Based Interactive Visualization Techniques*, Springer, 2006.
13. Benjamin Spencer, Robert S. Laramee, Guoning Chen, Eugene Zhang, *Evenly Space Streamlines for Surfaces: An Image Based Approach*, Computer Graphics Forum, **28**, (2009), 1618–1631.
14. Nelson Max, Barry Becker, Roger Crawfis, *Flow Volumes for Interactive Vector Field Visualization*, Proceedings Visualization '93, (1993), 19–24.
15. Lina C. Espinha, David A. Hoey, Paulo R. Fernandes, Helder C. Rodrigues, Christopher R. Jacobs, *Oscillatory Fluid Flow Influences Primary Cilia and Microtubule Mechanics*, Cytoskeleton (Hoboken, N.J.) **71**, (2014) 435-445.
16. Shenghong Huang, Rong Li, Q.S. Li, *Numerical simulation on fluid-structure interaction of wind around super-tall building at high reynolds number conditions*, Structure Engineering and Mechanics, An International Journal, **46**, (2013) 197–212.
17. P. Bhaniramka, Y. Demange, *OpenGL Volumizer: A toolkit for high quality volume rendering of large data sets*, In 2002 Symposium on Volume Visualization and Graphics, 2002, 45–53.
18. Hank Childs, Eric Brugger, Brad Whitlock, Jeremy Meredith, Sean Ahern, David Pugmire, Kathleen Biagas, Mark Miller, Cyrus Harrison, Gunther H. Weber, Hari Krishnan, Thomas Fogal, Allen Sanderson, Christoph Garth, E. Wes Bethel, David Camp, Oliver Rübel, Marc Durant, Jean M. Favre, Paul Navrátil, *VisIt: An End-User Tool For Visualizing and Analyzing Very Large Data*, High Performance Visualization–Enabling Extreme-Scale Scientific Insight, 2012, 357–372.
19. James Ahrens, Berk Geveci, Charles Law, *ParaView:An End User Tool for Large Data Visualization*, Visualization Handbook, Elsevier, 2005, ISBN-13: 978-0123875822.
20. R. Mittal and G. Iaccarino, *Immersed boundary methods*, Annu. Rev. Fluid Mech., **37**, (2005), 239–261.
21. R.J. LeVeque and Z.L. Li, *Immersed interface methods for Stokes flows with elastic boundaries or surface tension* SIAM J. Sci. Comput., **18**, (1997), 709–735.
22. R. Cortez, *A vortex/impulse method for immersed boundary motion in high Reynolds number flows*, J. Comput. Phys., **160**, (2000), 385–400.

23. XS Wang, *From Immersed Boundary Method to Immersed Continuum Method*, International Journal for Multiscale Computational Engineering, **4**, (2006), 127–145.
24. L. Zhang, A. Gersternberger, X. Wang, and W.K. Liu, *Immersed finite element method*, Comput. Methods Appl. Mech. Eng., **193**, (2004), 2051.
25. R. Glowinski, T. Pan, J. Periaux, *A fictitious domain method for Dirichlet problem and applications*, Comp. Methods in Appl. Mech. and Eng., **111**, (1994).
26. D.Sulsky, Z. Chen and H.L. Schreyer, *A particle method for history-dependent materials*, Comput. Mech. Appl. Mech. Eng., **118**, (1994), 179–197.
27. Hughes, Thomas JR and Liu, Wing Kam and Zimmermann, Thomas K, *Lagrangian-Eulerian finite element formulation for incompressible viscous flows*, Computer methods in applied mechanics and engineering, **29**, (1981), 329–349.
28. C.S. Peskin, *Flow patterns around heart valves: a digital computer method for solving the equations of motion. PhD thesis.*, Physiol., Albert Einstein Coll. Med, Univ. Microfilms. 378:72-30. (1972).
29. M. F. McCracken and C.S. Peskin, *A vortex method for blood flow through heart valves.* J. Comput. Phys., **35**, (1980), 183–205.
30. M. E. Rosar and C.S. Peskin, *Fluid flow in collapsible elastic tubes: A three-dimensional numerical model*, New York J. Math., **7**, (2001), 281–302.
31. A. M. Roma, C.S. Peskin and M. J. Berger(1999), *An adaptive version of the immersed boundary method*, J. Comput. Phys., **153**, (1999), 509–534.
32. BE Griffith, RD Hornung, DM McQueen, CS Peskin, *An adaptive, formally second order accurate version of the immersed boundary method*, J. Comput. Phys.,, **223**, (2007), 10–49.
33. M.C. Lai and C.S. Peskin, *An immersed boundary method with formal second order accuracy and reduced numerical viscosity*, J. Comput. Phys., **160**,(2000), 705.
34. B.E. Griffith, C.S. Peskin, *On the order of accuracy of the immersed boundary method: Higher order convergence rates for sufficient smooth problems*, J. Comput. Phys., **208** (2015), 75–105.
35. L. Zhu, C.S. Peskin, *Simulation of a flexible flapping filament in a flowing soap film by the immersed boundary method*, J. Comput. Phys., **179**, (2002), 452–468.
36. Y. Kim and C.S. Peskin, *Penalty immersed boundary method for an elastic boundary with mass*, Phys. of Fluids, **19**, (2007).
37. L.J. Fauci and A.L. Fogelson, *Truncated Newton methods and the modeling of complex elastic structures*, Comm. Pure Appl. Math., **46**, (1993), 787.
38. K. Taira and T. Colonius, *The immersed boundary method: a projection approach*, J. Comput. Phys., **225**, (2007), 2118–2137.
39. Y. Mori and C.S. Peskin, *Implicit second-order immersed boundary method with boundary mass*, Computer Methods in Applied Mechanics and Engineering, **197**, (2008), 2049–2067.
40. TY Hou, Z Shi, *An efficient semi-implicit immersed boundary method for the Navier-Stokes equations*, J. Comput. Phys., **227**, (2008), 8968–8991.
41. EP Newren, AL Fogelson, RD Guy, RM Kirby, *A comparison of implicit solvers for the Immersed Boundary equations*, Comput. Methods Appl. Mech. Engrg., **197**, (2008), 2290–2304.
42. Jian Hao and Luoding Zhu, *A lattice Boltzmann based implicit immersed boundary method for fluid-structure-interaction*, Computers and Mathematics with Applications, **59**, (2010), 185–193.
43. Jian Hao and Luoding Zhu, *A 3D implicit immersed boundary method with application*, Theoretical and Applied Mechanics Letter, **1**, (2011).
44. S. Lim, A. Ferent, X.S. Wang, and C.S. Peskin, *Dynamics of a closed rod with twist and bend in fluid*, Siam Journal on Scientific Computing, **31**, (2008), 273–302.
45. P. J. Atzberger, P. R. Kramer, and C. S. Peskin, *A stochastic immersed boundary method for biological fluid dynamics at microscopic length scale*, J. Comput. Phys., **224**, (2006), 1255–1292.
46. P. J. Atzberger and P. R. Kramer, *Error analysis of a stochastic immersed boundary method incorporating thermal fluctuations*, Mathematics and Computers in Simulation, **79**, (2008), 379–408.
47. Z.G. Feng and E.E. Michaelides, *The immersed boundary-lattice Boltzmann method for solving fluid-particles interaction problems*, J. Comput. Phys., **95**, (2004), 602–628.

48. Z.G. Feng and E.E. Michaelides, *Proteus: a direct forcing method in the simulations of particulate flows*, J. Comput. Phys., **202**, (2005), 20–51.
49. FB Tian, H Luo, L Zhu, X Lu, JC Liao, *An efficient immersed boundary-lattice Boltzmann method for the fluid-structure-interaction*, submitted to J. Comput. Phys., July 2010.
50. P Lee, BE Griffith, CS Peskin, *The immersed boundary method for advection-electrodiffusion with implicit timestepping and local mesh refinement*, Journal of Computational Physics, **229**), (2010), 5208–5227.
51. Fai, Thomas G and Griffith, Boyce E and Mori, Yoichiro and Peskin, Charles S, *Immersed boundary method for variable viscosity and variable density problems using fast constant-coefficient linear solvers I: Numerical method and results*, SIAM Journal on Scientific Computing, **35**, (2013), B1132–B1161.
52. Fai, Thomas G and Griffith, Boyce E and Mori, Yoichiro and Peskin, Charles S, *Immersed boundary method for variable viscosity and variable density problems using fast constant-coefficient linear solvers II: theory*, SIAM Journal on Scientific Computing, **36**, (2014), B589–B621.
53. Zhang, Chunze and Cheng, Yongguang and Zhu, Luoding and Wu, Jiayang, *Accuracy improvement of the immersed boundary–lattice Boltzmann coupling scheme by iterative force correction*, Computers & Fluids, **124**, (2016), 246–260.
54. Wu, J and Shu, Chang, *Implicit velocity correction-based immersed boundary-lattice Boltzmann method and its applications*, Journal of Computational Physics, **228**, (2009), 1963–1979.
55. Niu, XD and Shu, C and Chew, YT and Peng, Y, *A momentum exchange-based immersed boundary-lattice Boltzmann method for simulating incompressible viscous flows*, Physics Letters A, **354**, (2006), 173–182.
56. Wu, J and Shu, C and Zhang, YH, *Simulation of incompressible viscous flows around moving objects by a variant of immersed boundary-lattice Boltzmann method*, International Journal of Numerical Methods for Heat and Fluid Flow, **62**, (2010), 327–354.
57. Wu J, Cheng Y, Zhou W, Zhang C, Diao W, *GPU acceleration of FSI Simulations by the immersed boundary-lattice Boltzmann coupling scheme*, Computers & Mathematics with Applications, 2016.
58. Jiri Kraus, *Optimizing a LBM code for compute clusters with Kepler GPUs*, http://on-demand.gputechconf.com/gtc/2014/presentations/S4186-optimizing-lbm-code-compute-clusters-kepler-gpus.pdf, (2014).
59. Cheng, Yongguang and Zhu, Luoding and Zhang, Chunze, *Numerical study of stability and accuracy of the immersed boundary method coupled to the lattice Boltzmann BGK model*, Commun. Comput. Phys, **16**, (2014), 136–168.
60. Cheng, Yongguang and Zhang, Hui, *Immersed boundary method and lattice Boltzmann method coupled FSI simulation of mitral leaflet flow*, Computers & Fluids, **39**, (2010), 871–881.
61. Shu, Chang and Liu, Ningyu and Chew, Yong-Tian, *A novel immersed boundary velocity correction–lattice Boltzmann method and its application to simulate flow past a circular cylinder*, Journal of Computational Physics, **226**, (2007), 1607–1622.
62. Huang, Haibo and Sukop, Michael and Lu, Xiyun, Multiphase lattice Boltzmann methods: Theory and application, John Wiley & Sons, 2015.
63. Liu, Nansheng and Peng, Yan and Liang, Youwen and Lu, Xiyun, *Flow over a traveling wavy foil with a passively flapping flat plate*, Physical Review E, **85**, (2012), 056316.
64. Guo, Zhaoli and Shu, Chang, Lattice Boltzmann method and its applications in engineering, World Scientific, 2013.
65. Y. H. Qian, *Lattice gas and lattice kinetic theory applied to the Navier-Stokes equations", PhD thesis.*, University Pierre et Marie Curie, Paris, (1990).
66. S. Hou, Q. Zou, S. Chen, G. Doolen, and A. Cogley, *Simulation of cavity flow by the lattice Boltzmann method*, J. Compt. Phys., **118**, (1995), 329.
67. X. He, S. Chen, R. Zhang, *A lattice Boltzmann scheme for incompressible multiphase flow and its application in simulation of Rayleigh-Taylor instability*, J. Comput. Phys., **152**, (1999), 642–663.
68. D. A. Wolf-Gladrow, *Lattice-gas cellular automata and lattice Boltzmann Models – an introduction*, Springer, Berlin, 2000.

69. S. Succi, *The lattice Boltzmann equation*, Oxford Univ Press, Oxford, 2001.
70. L. S. Luo, *Unified Theory of the lattice Boltzmann models for nonideal gases*, Phys. Rev. Lett., **81**, (1998), 1618.
71. Kim, Joung-Dong and Li, Yan and Li, Xiaolin, *Simulation of parachute FSI using the front tracking method*, Journal of Fluids and Structures, **37**, (2013), 100–119.
72. Cottet, Georges-Henri and Maitre, Emmanuel, *A level set method for fluid-structure interactions with immersed surfaces*, Mathematical models and methods in applied sciences, **16**, (2006), 415–438.