

OpenLB User Guide

Associated with Release 1.6 of the Code



Authors: Adrian Kummerländer, Samuel J. Avis, Halim Kusumaatmaja, Fedor Bukreev, Michael Crocoll, Davide Dapelo, Simon Großmann, Nicolas Hafen, Shota Ito, Julius Jeßberger, Eliane Kummer, Jan E. Marquardt, Johanna Mödl, Tim Pertz, František Prinz, Florian Raichle, Martin Sadric, Maximilian Schecher, Dennis Teutscher, Stephan Simonis, Mathias J. Krause



Copyright © 2006-2008 Jonas Latt
Copyright © 2008-2023 Mathias J. Krause
info@openlb.net

Permission is granted to copy, distribute and/or modify this document under the terms of the GNU Free Documentation License, Version 1.2 or any later version published by the Free Software Foundation; with no Invariant Sections, no Front-Cover Texts, and no Back-Cover Texts. A copy of the license is included in the Section entitled “GNU Free Documentation License” (Section [A.3](#)).

Abstract

OpenLB is an object-oriented implementation of LBM. It is the first implementation of a generic platform for LBM programming, which is shared with the open source community (GPLv2). Since the first release in 2007 [66], the code has been continuously improved and extended which is documented by thirteen releases [55–65, 67, 68] as well as the corresponding release notes which are available on the OpenLB website <https://www.openlb.net>. The OpenLB code is written in C++ and is used by application programmers as well as developers, with the ability to implement custom models OpenLB supports complex data structures that allow simulations in complex geometries and parallel execution using MPI, OpenMP and CUDA on high-performance computers. The source code uses the concepts of interfaces and templates, so that efficient, direct and intuitive implementations of the LBM become possible. The efficiency and scalability has been checked and proved by code reviews. This user manual and a source code documentation by DoxyGen are available on the OpenLB project website.

Contents

1	Introduction	9
1.1	Lattice Boltzmann Methods	9
1.2	OpenLB Project	10
1.2.1	Scope and Overview	10
1.2.2	What Makes OpenLB Special?	12
1.2.3	Which Features are Currently Implemented?	12
1.2.4	Project Participants	14
1.2.5	Getting Help with OpenLB	14
1.2.6	How to Cite OpenLB	15
2	Core Concepts	17
2.1	Descriptor	17
2.2	(Super,Block)Lattice	19
2.2.1	Context	19
2.2.2	BlockLattice	19
2.2.3	Communication	20
2.2.4	Cell	20
2.3	Dynamics	21
2.3.1	Collision Operators	22
2.4	Post Processors	23
2.5	Functors	25
2.5.1	Basic Functor Types	25
2.5.2	Functor Arithmetic	26
2.6	Parallelization	29
2.6.1	Supported Platforms	29
3	Geometry Creation and Meshing	31
3.1	Creating a Geometry	31
3.2	Setting Material Numbers	32
3.3	Building Simulation Domains with Geometry Primitive Functors	33
3.4	Reading STL-files	34

3.5	Excursus: Creating STL-files	35
4	Simulation Models	36
4.1	Non-dimensionalization and Choice of Simulation Parameters	36
4.2	Porous Media Model	36
4.3	Power Law Model	38
4.4	Multiphysics Couplings	39
4.4.1	Shan–Chen Model	39
4.4.2	Free Energy Model	41
4.4.3	Coupling Between Momentum and Energy Equations	44
4.5	Advection–Diffusion Equation	45
4.5.1	AdvectionDiffusion Boundary Conditions	47
4.5.2	Convergence Criterion	49
4.5.3	Creating an Application with AdvectionDiffusionDynamics	50
4.5.4	Obtaining Results in Thermal Simulations	54
4.5.5	Conduction Problems	55
4.6	Particles	59
4.6.1	Class ParticleSystem	62
4.6.2	Class SuperParticleSystem	62
4.6.3	Class ParticleManager	63
4.6.4	Resolved Lattice Interaction	63
4.6.5	Particle Descriptors	64
4.6.6	Particle Dynamics	64
4.6.7	Particle Functions	65
4.6.8	Discrete Contact Model for Surface Resolved Particles	66
4.6.9	Sub-grid Legacy Framework	68
5	Initial and Boundary Conditions	79
5.1	Define Boundary Method	79
5.1.1	Wet-node Method	81
5.1.2	Link-wise Method	81
5.1.3	AdvectionDiffusionBoundary	82
5.1.4	Additional Options	83
5.2	Define Initial Conditions	83
5.3	Define Boundary Values	83
6	Input and Output	85
6.1	Output Data Structure	85
6.2	Write Simulation Data to VTK File Format	86
6.3	CSV Writer	87

6.4	Write Images Instantaneously	87
6.4.1	GifWriter	89
6.4.2	Heatmap	89
6.5	Gnuplot Interface	90
6.5.1	Regression with Gnuplot	91
6.6	Console Output	92
6.7	Read and Write STL Files	94
6.8	XML Parameter Files	94
6.9	Visualization with ParaView	95
6.10	Clip	96
6.11	Glyph	96
6.12	Stream Tracer	97
6.13	Resample To Image	97
6.13.1	Application of Functors	98
6.13.2	Error Norm Functors	106
7	Flow Control and Optimization	107
7.1	Introduction	107
7.1.1	Iterative Approach – Line Search Algorithm	108
7.1.2	Descent Algorithms	108
7.1.3	Step Size Algorithms	109
7.1.4	Derivative Computation	110
7.2	Implementation	110
7.2.1	Optimizer Classes: Optimization Methods	110
7.2.2	OptiCase Classes: Gradient Computation	111
7.2.3	Parameter Explanation and Reading from XML	111
8	Examples	118
8.1	Example Overview	118
8.2	adsorption	121
8.2.1	adsorption3D	121
8.2.2	microMixer3D	121
8.3	advectionDiffusionReaction	121
8.3.1	advectionDiffusionReaction2d	121
8.3.2	reactionFiniteDifferences2d	121
8.3.3	advectionDiffusion1d	122
8.3.4	advectionDiffusion2d	123
8.3.5	advectionDiffusion3d	124
8.3.6	advectionDiffusionPipe3d	124
8.4	laminar	125

8.4.1	bstep2d and bstep3d	125
8.4.2	cavity2d, cavity2dSolver and cavity3d	125
8.4.3	cylinder2d and cylinder3d	125
8.4.4	poiseuille2d and poiseuille3d	125
8.4.5	poiseuille2dEOC and poiseuille3dEOC	125
8.4.6	powerLaw2d	127
8.4.7	testFlow3dSolver	127
8.5	optimization	130
8.5.1	domainIdentification3d	130
8.5.2	poiseuille2Opti	131
8.5.3	showcaseADf	132
8.5.4	showcaseRosenbrock	132
8.5.5	testFlowOpti3d	132
8.6	multiComponent	133
8.6.1	binaryShearFlow2d	133
8.6.2	contactAngle2d and contactAngle3d	133
8.6.3	fourRollMill2d	134
8.6.4	microFluidics2d	134
8.6.5	phaseSeparation2d and phaseSeparation3d	134
8.6.6	rayleighTaylor2d and rayleighTaylor3d	134
8.6.7	youngLaplace2d and youngLaplace3d	134
8.7	particles	135
8.7.1	bifurcation3d	135
8.7.2	dkt2d	136
8.7.3	magneticParticles3d	136
8.7.4	settlingCube3d	137
8.8	porousMedia	137
8.8.1	porousPoiseuille2d and porousPoiseuille3d	137
8.9	reaction	137
8.9.1	advectionDiffusionReaction2d	137
8.10	thermal	139
8.10.1	galliumMelting2d	139
8.10.2	porousPlate2d, porousPlate3d and porousPlate3dSolver	139
8.10.3	rayleighBenard2d and rayleighBenard3d	140
8.10.4	squareCavity2d and squareCavity3d	141
8.10.5	stefanMelting2d	142
8.11	turbulent	143
8.11.1	aorta3d	143
8.11.2	channel3d	143

8.11.3	nozzle3d	143
8.11.4	tgvd3d	145
8.11.5	venturi3d	147
8.12	freeSurface	148
8.12.1	breakingDam2d and breakingDam3d	149
8.12.2	fallingDrop2d and fallingDrop3d	149
8.12.3	deepFallingDrop2d	149
8.12.4	rayleighInstability3d	150
9	Building and Running	151
9.1	Install Dependencies	151
9.1.1	Linux	152
9.1.2	Mac	152
9.1.3	Windows	152
9.2	Compiling OpenLB Programs	152
9.2.1	Using NVIDIA GPUs in OpenLB	153
10	Step by Step: Using OpenLB for Applications	160
10.1	Lesson 1: Getting Started - Sketch of Application	161
10.2	Lesson 2: Define and Use Boundary Conditions	163
10.3	Lesson 3: UnitConverter - Lattice and Physical Units	165
10.4	Lesson 4: Extract Data From a Simulation	166
10.5	Lesson 5: Convergence Check	168
10.6	Lesson 6: Use an External Force	169
10.7	Lesson 7: Understand Genericity in OpenLB	169
10.8	Lesson 8: Use Checkpointing for Long Duration Simulations	171
10.9	Lesson 9: Run Your Programs on a Parallel Machine	171
10.10	Lesson 10: Work with Indicators	172
10.11	Alternative Approach: Using a Solver Class	173
10.11.1	Structure of an OpenLB Simulation in Solver Style	173
10.11.2	Set up an Application in Solver Style	174
10.11.3	Parameter Explanation and Reading from XML	174
	References	180
A	Appendix	190
A.1	Q&A	190
A.2	List of Project Participants	195
A.3	GNU Free Documentation License	198

1. Introduction

1.1. Lattice Boltzmann Methods

In general, the lattice Boltzmann method (LBM) can be interpreted as a bottom-up procedure to numerically approximate a given partial differential equation (PDE). In contrast to the conventional top-down discretization of the target equation (TEQ), the LBM emerges from the space and time discretization of a discrete velocity Boltzmann-type equation. Subsequently, the solution to the TEQ is recovered in a limiting process. Different PDEs can be approximated by altering certain features of the method. Exemplary LBM building blocks which induce the recovery of the aimed at transport equation are discussed in the following sections. Thorough derivations of specific LBM approaches are summarized more comprehensively for example in [22, 100].

The LB algorithm is typically divided into two steps: collision and streaming. During the collision step every distribution function f_i at a grid node $\mathbf{x} \in \Omega_{\Delta x} \subseteq \mathbb{R}^d$ receives a collision term J_i , i.e.

$$f_i^*(t, \mathbf{x}) = f_i(t, \mathbf{x}) + \Delta t J_i(t, \mathbf{x}) , \quad (1.1)$$

where Δt is the time step size and $t \in I_{\Delta t} \subseteq \mathbb{R}_{\geq 0}$. The most prominent collision operator introduced by Bhatnager, Gross and Krook (BGK) [78] reads

$$J_i(t, \mathbf{x}) = -\frac{1}{\tau} [f_i(t, \mathbf{x}) - f_i^{eq}(t, \mathbf{x})] . \quad (1.2)$$

The relaxation time $\tau > 0$ determines the speed at which the populations approach equilibrium and can for example be related to the viscosity in the hydrodynamic limit. The local equilibrium function f_i^{eq} approximates the Maxwell–Boltzmann distribution. At the streaming step all populations are shifted to the next grid point

$$f_i(t + \Delta t, \mathbf{x} + \Delta t \mathbf{c}_i) = f_i^*(t, \mathbf{x}) , \quad (1.3)$$

where $i = 0, 1, \dots, q-1$ enumerates the discrete velocities \mathbf{c}_i of dimension d contained in the underlying set $DdQq$. For the purpose of illustration, the $D2Q9$ and the $D3Q19$ discrete velocity sets are shown in Figures 1.1a and 1.1b, respectively. As a linkage from the LBM to the TEQ, the macroscopic conservative variables are obtained by respective moment summation over the populations and typically yield an approximation up to second order in space and time.




Figure 1.1.: Exemplary discrete velocity sets used for the descriptors implemented in OpenLB.

The following references are suggested for further insight into LBM.

- The book *The Lattice Boltzmann Method: Principles and Practice* [2017] by Krüger *et al.* [100] delivers a clear and complete introduction for beginners.
- A concise introduction is given by Mohamad [110]. The book *Lattice Boltzmann Method* [2011] documents for example the formal derivation of macroscopic limit equations using Chapman–Enskog expansion.
- An LBM predecessor—the Lattice-Gas Cellular Automata (LGCA)—is extensively described in Wolf-Gladrow’s book *Lattice-Gas Cellular Automata and Lattice Boltzmann Models* [2000]. Starting with the derivation of LGCA, the author derives the LBM step by step. Furthermore, a helpful interpretation of LBM is given in the beginning of the book.
- A quick overview of LBM, can be obtained from the paper *Lattice Boltzmann Method for Fluid Flows* [1998] by Chen and Doolen [85].

1.2. OpenLB Project

1.2.1. Scope and Overview

OpenLB is a generic implementation of LBM written using the C++ programming language and is used by application programmers as well as method developers. It is the first implementation of a comprehensive platform for LBM, that is freely shared with the open source community. Since its conception, OpenLB is developed as open source under the GPL2 license. Therefore, everyone with access to the source code has the right to use, adapt and publish the software under the same license. Since the first release in 2007 [66], the code has been continuously improved and extended across thirteen major releases. User guides and source code documentation for developers are available on the project website. In summary, OpenLB is a technically

complex and feature-rich CFD software library with an extensive set of example cases also written in modern C++.




Figure 1.2.: Resolved simulation with 250 million grid cells of turbulent mixing and chemical reactions done by OpenLB on 24 GPUs on the HoreKa supercomputer funded by the Ministry of Science, Research and the Arts Baden-Württemberg and by the Federal Ministry of Education and Research.

OpenLB is actively used in both developing novel methodologies and their applications. Developers of OpenLB, especially researchers from mathematics and computer science, focus on the methodology-oriented approach. Here a selection of recent (after 2016) publications is given:

- General LBM: [10, 15, 31, 37–39]
- HPC LBM: [23, 50, 33]
- Particle LB – Methods: [2, 20, 25, 27, 45–47]
- Turbulence LB – Methods: [12, 13, 34, 39, 41]
- Optimal Control LB – Methods: [18, 19, 35, 48]

The application-oriented approach is chosen by esp. engineering researchers. At the top level, each simulation in OpenLB currently consists of a fully parallel mesh construction (without any need for external tools), physical parametrization, assignment of dynamics encoding the cell-specific models, configuration of operators for multi-physics coupling, non-local boundary treatment and a selection of post-processing functions for extracting the relevant flow features

for further evaluation [22]. OpenLB enables application in a broad set of areas (physical characterizations) such as incompressible Newtonian and non-Newtonian fluid flows, multi-phase and multi-component flows (cf. Figure 1.2), light and thermal radiation as well as thermal flows, particulate flows using both Euler–Euler and Euler–Lagrange with resolved and sub-grid scale models, turbulent flow models (LES models and wall function approaches) and porous media models [22]. A selection of recent (after 2016) application-oriented publications is listed here:

- Particle LB – Applications: [49, 1, 3, 7, 8, 14, 24, 26, 28, 30, 54, 43, 44]
- Turbulence LB – Applications: [4–6, 9, 11, 13, 32, 52, 42]
- Optimal Control LB – Application: [16, 17, 19, 36]

OpenLB is not only an open source code, but also a community that aims to share resources for software development, produce reproducible scientific results, and provide a good start for students in LBM. Through the years, there have been about over 150 scientific articles published, over 90 with and over 60 without the OpenLB developers’ participation (2022). Until January 2023, more than 50 individuals from Germany, Switzerland, France, United Kingdom, Brazil, Czech Republic and USA have contributed to OpenLB.

1.2.2. What Makes OpenLB Special?

OpenLB is a numerical framework for lattice Boltzmann simulations, created by students and researchers with different backgrounds in computational fluid dynamics. The code can be used by application programmers to implement specific flow geometries in a straightforward way, and by developers to formulate new models. For this first group of users, OpenLB offers a neat interface through which it is possible to set up a simulation with little effort. For the second group, the structure of the code is kept conceptually simple, implementing basic concepts of the lattice Boltzmann theory step-by-step. Thanks to this, the code is an excellent framework for programmers to develop pieces of reusable code that can be exchanged in the community.

One key aspect of the OpenLB code is genericity in its many facets. The core concept of generic programming is to offer a single code that can serve many purposes. On one hand, the code implements dynamic genericity through the use of object-oriented interfaces. One use of this is that the behavior of lattice sites can be modified during program execution, to distinguish for example between bulk and boundary cells, or to modify the fluid viscosity or the value of a body force dynamically. Furthermore, C++ templates are used to achieve static genericity. As a result, it is sufficient to write a single generic code for various 3D lattice structures, such as D3Q15, D3Q19, and D3Q27 (for more information on lattice structures, see Section 2.1).

1.2.3. Which Features are Currently Implemented?

An excerpt of the the current features is given below. Note that an extended list is provided in [22]. For detailed description of individual features, see the respective sections below.

1.2.3.1. Pre-processing

Simulation domain construction with volume meshes (VTI)	Section 3.1
Simulation domain construction with surface meshes (STL)	Section 3.1
Simulation domain construction with geometry primitives	Section 3.1
Automatic parallel meshing	Section 3.1
Automatic load balancing	
Segmentation for boundaries setting	Section 3.2
Assigning boundary and initial values	Section 2.5, 6.13.1

1.2.3.2. Post-processing

Interfaces to parallel VTK formats	Section 6.2
Interface CSV	Section 6.3
Writing images build-in	Section 6.4
Evaluating with Gnuplot	Section 6.5
Analytical functors for e.g. error norms, integrals, ...	Section 2.5, 6.13.1

1.2.3.3. Lattice Boltzmann Models

BGK model for fluids	Section 2.3	Reference [85]
Regularized model for fluids	Section 2.3	Reference [104]
Multiple Relaxation Times (MRT)	Section 2.3	References [87, 125]
Entropic Lattice Boltzmann	Section 2.3	Reference [76]
BGK with adjustable speed of sound	Section 2.3	References [69, 86]
BGK and MRT with Smagorinsky model	Section 2.3	References [53]
Porous media model	Section 2.3	
Power law model	Section 2.3	

1.2.3.4. Multiphysics Coupling

See Section 4.4.

Shan-Chen two-component fluid	Section 4.4	Reference [117]
Free energy model for multicomponent fluids	Section 4.4	Reference [116]
Thermal fluid with Boussinesq approximation	Section 4.4	Reference [93]

1.2.3.5. Lattice Structures

See Section 2.1. Exemplary discrete velocity sets (lattices) available in OpenLB are $D2Q5$, $D2Q9$, $D3Q7$, $D3Q13$, $D3Q15$, $D3Q19$, $D3Q27$.

1.2.3.6. Boundary Conditions for Straight Boundaries (Including Corners)

See Section 5.1.

Regularized	local	Default choice in examples
Finite difference (FD) velocity gradients	non-local	Default choice in examples
Inamuro	local	
Zou/He	local	
Non-linear FD velocity gradients	non-local	

1.2.3.7. Boundary Conditions for Curved Boundaries

See Section 5.1.

Bouzidi	non-local	References [80]
---------	-----------	-----------------

1.2.3.8. Input / Output

The basic mechanism behind I/O operations in OpenLB is the serialization and unserialization of a `BlockLatticeXD`. This mechanism is used to save the state of a simulation and to produce VTK output for data post-processing with external tools. In both cases, the data is saved in the binary Base64 format, which ensures compact and (relatively) platform-independent data storage. For operational details, see for example Section 6.1.

1.2.4. Project Participants

The OpenLB project was initiated in 2006. Between 2006 and 2008 Jonas Latt was the project coordinator. As of 2009, Mathias J. Krause is coordinating the project. Since 2006 more than 45 persons contributed to OpenLB. A list is provided in Section A.2 of the Appendix.

1.2.5. Getting Help with OpenLB

The following resources are available for OpenLB users:

Web site: Most recent releases of the code and documentation, including this user guide, can be found on the website <https://www.openlb.net/>.

Forum: If you experience troubles with OpenLB, you may wish to post your concerns to the Lattice Boltzmann community in the forum on the OpenLB homepage.

Bug reports: If you think you found a bug in OpenLB, we encourage you to submit a report to bug@openlb.net. Useful bug reports include the full source code of the program in question, a description of the problem, an explanation of the circumstances under which

the problem occurred, and a short description of the hardware and the compiler used. Moreover, other makefile switches, such as buildtype and mode of parallelization found in `config.mk` can provide useful information too.

Spring School: To lower the hurdle and the assistance of people new to this field is the motivation in organizing yearly spring schools, namely **Lattice Boltzmann Methods with OpenLB Software Lab**. The first spring school was organized in [2017 in Hammamet \(Tunisia\)](#), followed by [Karlsruhe in 2018](#), [Mannheim in 2019](#), [Berlin in 2020](#), [Krakow \(Poland\) in 2022](#) and [London \(UK\) in 2023](#). All spring schools are organized as open international workshops. The intention of creating an international platform with courses for beginners in LBM, which are researchers from academia and industry, was fulfilled with a response of 40 participants from 15 countries, averaged over the years. The event is based on the interlaced educational concept of comprehensive and methodology-oriented courses offered by the core developer team of OpenLB together with the local partner group as well as professional guest lecturers within the field of LBM. The first half of the week is dedicated to the theoretical fundamentals of LBM up to ongoing research on selected topics. Followed by mentored training on case studies using OpenLB in the second half, where the participants gain deep insights into LBM and its various applications in different disciplines. More information about the next spring school can be found on the OpenLB website.

Consortium: If you want more or deeper support or get involved in further development of OpenLB, become a member of the OpenLB consortium. If you are interested, please send an email to consortium@openlb.net. The consortium promotes and aims at continuous development and maintenance of OpenLB, especially for (1) improvement of the general usage according to GNU GPL2 License and regular publication, (2) maintenance of the executability on current standard HPC-Hardware and (3) conservation of the current status of the LBM research in OpenLB.

1.2.6. How to Cite OpenLB

For references to OpenLB in general, we recommend to cite the most recent overview paper [\[22\]](#).

```
@article{openlb-2020,  
  author    = "Krause, M. J. and Kummerl{\"a}nder, A. and Avis, S. J.  
              and Kusumaatmaja, H. and Dapelo, D. and Klemens, F.  
              and Gaedtke, M. and Hafen, N. and Mink, A.  
              and Trunk, R. and Marquardt, J. E. and Maier, M.-L.  
              and Haussmann, M. and Simonis, S.",  
  title     = "OpenLB--Open source lattice Boltzmann code",  
  year      = "2020",  
  journal   = "Computers & Mathematics with Applications",
```

```

    doi      = "10.1016/j.camwa.2020.04.033"
}

```

For references to a specific release of the code, each version is associated with a citable Zenodo publication. The latest release 1.6 [65] is covered by the present user guide and citable via:

```

@software{olbRelease16,
  title      = {{OpenLB Release 1.6: Open Source Lattice Boltzmann Code}},
  author      = {Kummerl\{a\}nder, Adrian
    and Avis, Sam
    and Kusumaatmaja, Halim
    and Bukreev, Fedor
    and Crocoll, Michael
    and Dapelo, Davide
    and Hafen, Nicolas
    and Ito, Shota
    and Je\ss\{o\}berger, Julius
    and Marquardt, Jan E.
    and M\{o\}dl, Johanna
    and Pertzel, Tim
    and Prinz, Franti\v{s}ek
    and Raichle, Florian
    and Schecher, Maximilian
    and Simonis, Stephan
    and Teutscher, Dennis
    and Krause, Mathias J.},
  year        = 2023,
  month        = apr,
  publisher    = {Zenodo},
  doi          = {10.5281/zenodo.7773497},
  url          = {https://doi.org/10.5281/zenodo.7773497},
  version      = {1.6}
}

```

Starting with release 1.6, metadata is also defined in the standardized *Citation File Format* and included as a `CITATION.cff` in the release tarball.

2. Core Concepts

The basic data structure of any *Lattice Boltzmann* code is the eponymous *Lattice* or grid. For the most widely accepted formulations, such a grid is a regular, homogeneous lattice Ω_h with equal spacing $h \in \mathbb{R}_{>0}$ in all directions. Each spatial location on this lattice can be referred to as a *Cell* which is the core unit of the LB algorithm.

One of the main aspects motivating the use of LBM is its unique suitability for massive parallel processing. This parallelization aspect naturally leads to the need for further spatial decomposition of the lattice data structure in order to be able to distribute the processing to multiple independent processing units. In OpenLB, this spatial decomposition into so called *Blocks* motivates the core naming convention of distinguishing between *Super*- and *Block*-level structures.

This approach respects the spirit of LBM well and leads to elegant and efficient implementations. For complex geometries, a multi-block approach provides another advantage: A given domain can be represented by a certain number of regular blocks, which delivers a good compromise between highly efficient block-local processing and sparse memory consumption.

For most practical applications, the lattice data structure not only manages the population values assigned to the individual cells and declared by the descriptor (cf. Section 2.1), but also associated per-cell data such as force fields or reaction coefficients.

2.1. Descriptor

Descriptors are the structure that is used to define and access LB model specific information such as the number of dimensions and discrete velocities as well as weights and declarations of additional fields. As such, a descriptor is a central choice in any OpenLB application and passed as a template argument throughout the code base.

```
1 using T = FLOATING_POINT_TYPE;
2 using DESCRIPTOR = descriptors::D2Q9<>;
3
4 // number of spatial dimensions
5 const int d = descriptors::d<DESCRIPTOR>(); // == 2
6 // number of discrete velocities
7 const int q = descriptors::q<DESCRIPTOR>(); // == 9
8
9 // second discrete velocity vector
10 const Vector<int,2> c1 = descriptors::c<DESCRIPTOR>(1); // == {-1,1}
11
12 // weight of the first discrete velocity
13 const T w = descriptors::t<T,DESCRIPTOR>(0); // == 4./9.
```

OpenLB provides a rich set of such descriptors, most of which are defined in `src/dynamics/latticeDescriptors.h`. Despite the central role of this concept, the concrete definitions of for example the `D2Q9` descriptor, are quite compact. To illustrate this point, Listing 2.1 provides the full definition of this descriptor including all of its data.

```

1  template <typename... FIELDS>
2  struct D2Q9 : public LATTICE_DESCRIPTOR<2,9,POPULATION,FIELDS...> {
3      D2Q9() = delete;
4  };
5
6  namespace data {
7
8      template <>
9      constexpr int vicinity<2,9> = 1;
10
11     template <>
12     constexpr int c<2,9>[9][2] = {
13         { 0, 0},
14         {-1, 1}, {-1, 0}, {-1,-1}, { 0,-1},
15         { 1,-1}, { 1, 0}, { 1, 1}, { 0, 1}
16     };
17
18     template <>
19     constexpr int opposite<2,9>[9] = {
20         0, 5, 6, 7, 8, 1, 2, 3, 4
21     };
22
23     template <>
24     constexpr Fraction t<2,9>[9] = {
25         {4, 9}, {1, 36}, {1, 9}, {1, 36}, {1, 9},
26         {1, 36}, {1, 9}, {1, 36}, {1, 9}
27     };
28
29     template <>
30     constexpr Fraction cs2<2,9> = {1, 3};
31
32 }

```

Listing 2.1.: Definition of the `D2Q9` descriptor

Many LBM based solutions to practical problems require the ability to store not just the populations of each cell but also additional data such as an external force (see e.g. Section 10.6). For this reason every descriptor may explicitly declare such fields.

```

1  using DESCRIPTOR = descriptors::D2Q9<descriptors::FORCE>;
2
3  // Check whether DESCRIPTOR contains the field FORCE
4  DESCRIPTOR::provides<descriptors::FORCE>(); // == true
5  // Get cell-local memory location of the FORCE field
6  const int offset = DESCRIPTOR::index<descriptors::FORCE>(); // == 9
7  // Get size of the descriptor's FORCE field
8  const int size = DESCRIPTOR::size<descriptors::FORCE>(); // == 2

```

A list of various predefined fields can be found in the `src/dynamics/descriptorField.h`

header. Note that one may add an arbitrary list of such fields to a given descriptor. Even more so it is fully supported to add fields on a per-app basis. One only needs to make sure that the type is defined prior to any custom descriptor that depends on it, i.e. a user could write inside of a case.

```
1 struct MY_CUSTOM_FIELD: public FIELD_BASE<42,0,0> { };
2 using DESCRIPTOR = D2Q9<FORCE,MY_CUSTOM_FIELD>;
```

This custom descriptor can then be used in the same way as any other descriptor. This might even be preferable for very specific fields that are only used by a small number of apps or specific user-provided features.

2.2. (Super,Block)Lattice

The `SuperLattice` is OpenLB's main class for representing and maintaining a concrete lattice with all associated data, cell local models and non-local operators. It is constructed from a load balanced geometry, i.e. given a `SuperGeometry` instance which is in turn obtained from a `CuboidGeometry` and its associated `LoadBalancer`.

2.2.1. Context

The `CuboidGeometry` decomposes a given indicator geometry into individual cuboids. By design, there is a bijection between the set of cuboids and the set of block lattices resp. block geometries. Given a cuboid geometry, the `LoadBalancer` assigns the contained cuboids to the available number of MPI processes. Based on this structure, each spatial cell location of each lattice location encompassed by the cuboid geometry is assigned a so called *material number* managed by the `SuperGeometry` (cf. Chapter 3 on meshing). Finally, a concrete `BlockLattice` is instantiated by the `SuperLattice` instance of the responsible `LoadBalancer`-assigned MPI process. This class maintains the actual data associated with all its lattice cells. As such, each process only has access to the data of *its own* blocks extended by a small overlap for inter-block communication. This communication (i.e. synchronization and overlap data) is performed by `SuperCommunicator` instances maintained in *stages* of the `SuperLattice`.

2.2.2. BlockLattice

The `BlockLattice` class is the high-level interface abstracting the platform-specific implementation details of the actual core structures managed by `ConcreteBlockLattice` class template specializations.

Specifically, there are specializations for each supported platform: `CPU_SISD`, `CPU_SIMD`, and `GPU_CUDA`. At its core, adding support for a new platform consists of specializing all platform-templated classes.

2.2.2.1. Core Data

The very core data structure of OpenLB is the per-field (cf. Section 2.1) column or *Structure-of-Arrays* (SoA) storage `FieldArrayD`. This `FieldArrayD` maintains `(Cyclic)Column` instances for each component of its `FIELD` template argument in a fashion specific to the concrete platform. In other words, the same `FIELD` is stored differently for each block / platform, while sharing the same interface of column-wise access. This fundamental structure allows for exchanging the entire core data structure implementation without touching the higher level interfaces against which for example cell specific models are implemented (cf. Section 2.3).

On the next level, all `FieldArrayD` instances of a block are maintained within the `Data` class. This class provides an on-demand allocating `getFieldArray` method for access to the field array specific to a given `FIELD`.

The details of rendering this most performance-critical structure efficient across all platforms exceeds the scope of this guide, but you can be assured that the developer team invests lots of effort in preserving and increasing OpenLB's performance.

2.2.3. Communication

While the cuboid decomposition maintained by the `CuboidGeometry` is disjoint, block geometries and lattices overlap by a configurable overlap layer. Synchronizing this overlap layer, that is communicating the data owned by a given block to its neighbor's overlap layer, is how any inter-block communication is handled in OpenLB.

The most frequent type of overlap communication is the synchronization of a single layer of populations during the *PostCollide* stage of the core *collideAndStream* iteration.

Communication of given stages can be manually triggered using the following.

```
1 sLattice.getCommunicator(PostCollide{}).execute();
```

Similarly, communication stages may be configured w.r.t. to their extent and included fields. For example, we can consider the setup of the default *PostCollide* stage.

```
1 auto& communicator = sLattice.getCommunicator(PostCollide{});
2 communicator.template requestField<descriptors::POPULATION>();
3 communicator.requestOverlap(1);
4 communicator.exchangeRequests();
```

2.2.4. Cell

Users familiar with previous versions of OpenLB may remember the original hierarchy of the `SuperLattice` containing `BlockLattice` instances containing a D-dimensional array of `Cell` instances containing a Q-dimensional array of population values.

This data-holding `Cell` was refactored by OpenLB 1.4 [63] to only provide a view of the data maintained in field arrays (cf. Section 2.2.2.1). This provided an essential ingredient for the

implementation of both SIMD and GPU support in OpenLB 1.5 [64] (cf. Talk [51] for a summary of the refactoring journey).

Today, there are multiple per-platform and use case versions of the cell interface. Core structures such as Dynamics and non-local operators are templated against the *concept of a cell* instead of a specific implementation thereof. This is essential for being able to instantiate the same abstract model implementations on diverse hardware, e.g. both in vectorized CPU collision loops and CUDA GPU kernel functions.

2.3. Dynamics

In the context of OpenLB, the term *Dynamics* refers to the implementation of cell-local LB models. i.e. equilibrium distribution, collision step and momenta computation. Each *cell* of a lattice is assigned a *dynamics* describing its local behavior. This way, it is easy to implement inhomogeneous fluids which use a different type of physics from one cell to another.

The non-local streaming step between cells is modeled independent of the choice of cell-local models and remains invariant. OpenLB’s block-local propagation is implemented using the *Periodic Shift* pattern [23] on all platforms.

This concept of per-cell dynamics is tied together in the `Dynamics` interface. Every dynamics implementation describes the various components required to model the local behavior of the assigned cell locations. Specifically, this includes the set of momenta, the equilibrium distribution as well as the actual collision operator. Fittingly, most dynamics are declared as a tuple of those three components.

```

1 template <typename T, typename DESCRIPTOR>
2 using TRTdynamics = dynamics::Tuple<
3   T, DESCRIPTOR,
4   momenta::BulkTuple,
5   equilibria::SecondOrder,
6   collision::TRT
7 >;

```

In this example declaration of the `TRTdynamics` we can tell at a glance that its assigned cells will expose bulk momenta reconstructed directly from the population values and relax towards the second order equilibrium using the TRT collision operator. The dynamics tuple concept is quite powerful, e.g. we may apply a forcing scheme by adding a single line.

```

1 template <typename T, typename DESCRIPTOR>
2 using ForcedTRTdynamics = dynamics::Tuple<
3   T, DESCRIPTOR,
4   momenta::BulkTuple,
5   equilibria::SecondOrder,
6   collision::TRT,
7   forcing::Guo
8 >;

```

This `forcing::Guo` *combination rule* is the fourth and optional component of the dynamics tuple. Combination rules may arbitrarily manipulate the previous momenta, equilibrium and collision components. In the case of Guo, forcing this means that the momenta argument is shifted via `momenta::Forced` and the TRT collision operator is wrapped to force the post collision populations according to the Guo scheme.

OpenLB offers a comprehensive library of momenta, equilibria, collision operators and combination rules that can be easily combined into many different dynamics tuples. See

- `src/dynamics/momenta/aliases.h`
- `src/dynamics/equilibrium.h`
- `src/dynamics/collision.h`
- `src/dynamics/forcing.h`

for some examples.

While this framework covers most collision steps, a fallback option is provided via `dynamics::CustomCollision`. This class enables the combination of momenta with arbitrary collision and equilibrium implementations. One example for this are the `CombinedRLBdynamics` that are used as the foundation for various boundary conditions.

The full definition of the interface is available in `src/dynamics/interface.h`. Note that due to recent large refactoring to support execution on GPUs and SIMD CPUs, `Dynamics` currently contains various legacy methods that will be deprecated in future releases. New dynamics implementations should be formulated either as a `dynamics::Tuple` or `dynamics::CustomCollision` template.

2.3.1. Collision Operators

A partial summary of the currently supported collision operators can be found in the recent publication on OpenLB [22]. Furthermore, five commonly used collision schemes are compared in [12] for a typical 3D benchmark case of homogeneous isotropic turbulence. For derivations, analysis and theoretical comparisons the interested reader is referred to [100].

2.3.1.1. Implementation in Dynamics

As was touched upon in Section 2.3, local collision operators are expressed as `Dynamics` in the context of OpenLB. Specifically, the common dynamics tuple concept expresses collision operators as reusable elements alongside equilibria and momenta.

```

1 struct BGK {
2     using parameters = typename meta::list<descriptors::OMEGA>;
3
4     static std::string getName() {
5         return "BGK";

```

```

6   }
7
8   template <typename DESCRIPTOR, typename MOMENTA, typename EQUILIBRIUM>
9   struct type {
10      using EquilibriumF = typename EQUILIBRIUM::template type<DESCRIPTOR,MOMENTA>;
11
12      template <typename CELL, typename PARAMETERS, typename V=typename CELL::value_t>
13      CellStatistic<V> apply(CELL& cell, PARAMETERS& parameters) any_platform {
14          V fEq[DESCRIPTOR::q] { };
15          const auto statistic = EquilibriumF().compute(cell, parameters, fEq);
16          const V omega = parameters.template get<descriptors::OMEGA>();
17          for (int iPop=0; iPop < DESCRIPTOR::q; ++iPop) {
18              cell[iPop] *= V{1} - omega;
19              cell[iPop] += omega * fEq[iPop];
20          }
21          return statistic;
22      };
23  };
24 };

```

This is the complete listing of the well known BGK collision operator that is used by many different dynamics. Each collision operator consists of three elements: A `parameters` type list of fields that is used to parametrize the collision, a `getName` method that is used to generate human readable names for dynamics tuples and a nested `type` template that contains the actual `apply` method specific to each operator.

The nested `type` template is used to enable composition into dynamics tuples and will be automatically instantiated for the required descriptor, momenta and equilibrium types. Additionally, this is used as a place for injecting partial specialization which enable usage of auto-generated CSE-optimized kernels.

As is the case for all other elements, the `apply` template method follows a fixed signature for all collision operators. Each call is provided an instance of some platform-specific implementation of the cell concept alongside a parameters structure containing all requested values. Using these two inputs the method can perform the local collision and return the computed density and velocity magnitudes for usage in lattice statistics. This pattern is repeated at various places of the library. Examples for other instances are equilibria and momenta elements as well as post processors. Any implementations of this style are usable on any of OpenLB's target platforms (currently this means the scalar and vectorized CPU code as well as GPU support). They are also amenable to automatic code generation.

For an introduction on how to write your own dynamics, see the FAQ in Section [A.1](#).

2.4. Post Processors

While the basic concept of dynamics assigned to cells in a block lattice is conceptually close to the theory of LBM, it is not sufficiently general to address all possible requirements arising in complex applications. As a case in point, some boundary conditions are non-local and need

to access neighboring nodes. Their execution is taken care of by a *post processing* step, which instead of traversing the entire lattice a second time, applies to selected cell locations only.

While collision steps are easily parallelized due to their local nature, this doesn't hold once neighborhood access is required. For this reason, two adjacent concepts are used to group post processors. Each post processor is assigned to a *stage* and a *priority* within this stage. For example both `OuterVelocityCornerProcessor3D` and `OuterVelocityEdgeProcessor3D` are processed in the `PostStream` stage, but the former is executed after the latter to avoid access conflicts. In turn, post processors within the same stage and priority may be executed in parallel depending on the execution platform. Note that both the number of priorities and stages may be freely customized – e.g. the free surface code introduces a number of additional stages to interleave post processing and custom communications steps.

Each post processor consists of scope and priority declarations in addition to an `apply` template method. For the aforementioned `OuterVelocityCornerProcessor3D` this is declared as follows (see `src/boundary/boundaryPostProcessors3D.hh` for the full implementation).

```

1  template<typename T, typename DESCRIPTOR,
2      int xNormal, int yNormal, int zNormal>
3  struct OuterVelocityCornerProcessor3D {
4      static constexpr OperatorScope scope = OperatorScope::PerCell;
5
6      int getPriority() const {
7          return 1;
8      }
9
10     template <typename CELL>
11     void apply(CELL& cell) any_platform;
12 };

```

Here, the `OperatorScope::PerCell` declares that the `apply` function will be provided a cell implementation with neighborhood access for each assigned location. Other scopes such as `OperatorScope::PerBlock` (used for example for statistics computation) enable different access patterns. In any case, the assigned cell locations are maintained by OpenLB's post processor framework. For example

```

1  sLattice.addPostProcessor<PostStream>(indicator, meta::id<SomePerCellPostProcessor>{});

```

schedules a post processor for application to all indicated cells during the `PostStream` stage.

Note that this describes the new post processor concept adopted by OpenLB 1.5. Many existing *legacy* post processors use a different paradigm. Specifically, they are derived from `PostProcessor(2,3)D` and override virtual methods such as

```
void PostProcessor(2,3)D::process(BlockLattice<T,DESCRIPTOR>&).
```

All these post processors will be ported to the new approach in time. For CPU targets, legacy post processors can be used without restrictions but they are not supported on the GPU platform. For an introduction on how to write your own post processors, see the FAQ in Section A.1.

2.5. Functors

In general, a *functor* is a class that behaves like a function. Objects of a functor class perform computations by overloading the `operator()` in a standardized fashion, allowing for arithmetic combination. One big advantage of functors over functions is that they allow the creation of a hierarchy and bundle “classes of functions”.

2.5.1. Basic Functor Types

The functor concept is a user friendly and efficient technique to process lattice data and extract relevant data for post-processing. In the meantime, OpenLB deploys the functors also for the geometry, which is a very intuitive and powerful choice.

Basically, functors are applications that operate either on the lattice \mathbb{N}^3 or more generally on \mathbb{R}^3 . The values of such a functor may be three dimensional, as for example for the velocity, where the mapping is defined as

$$\text{Functor} : \Omega \rightarrow \mathbb{R}^d, \quad \text{for } d \in \mathbb{N}. \quad (2.1)$$

The nomenclature is based on the dimension of the domain. Let’s say the functor acts on a $3d$ (super) lattice, the the functor is called `SuperLatticeF3D`. If the functor value is density, then this functor is called `SuperLatticeDensity3D`.

2.5.1.1. GenericF

The `GenericF` functor stands at the top of the hierarchy and is a virtual base class that provides interfaces. Template parameter `s` defines the input data type and template parameter `T`, the output. The essential interface is the unwritten (pure virtual function) `operator()`. Commonly, this `()`-operator is used as an evaluation of a certain functor, e.g. pressure at position x .

2.5.1.2. AnalyticalF

This a subclass of `GenericF`, for functions that live in SI-units, e.g. for setting velocities in m/s . Parts of this class are, for example, constant, linear, interpolation and random functors, which can be evaluated by the `()`-operator. There is an `AnalyticalCalc` class, which inherits from `AnalyticalF` and establishes arithmetic operations ($+$, $-$, \cdot , $/$) between every type of

$$\text{AnalyticalF3D} : \mathbb{R}^3 \rightarrow \mathbb{R}^d, \quad d \in \mathbb{N}. \quad (2.2)$$

2.5.1.3. IndicatorF

This is another subclass of `GenericF` that returns a vector with elements 0 or 1, i.e.

$$\text{IndicatorF3D} : \mathbb{R}^3 \rightarrow \{0, 1\} . \quad (2.3)$$

These are used to construct geometries, e.g. `IndicatorSphere3D` creates a sphere using an origin and radius. Evaluation returns 1, if the vector is inside the sphere and 0 otherwise. In analogy to the `AnalyticalF`, there are arithmetic operations as well, but with a slightly different definition. The returned object of an addition is the union, multiplication returns the intersection, and subtraction represents the relative complement.

2.5.1.4. SmoothIndicatorF

Smooth indicators define a *small* epsilon region around the object such that it has a smooth transition from 0 to 1. In general, the mapping is defined as

$$\text{SmoothIndicatorF3D} : \mathbb{R}^3 \rightarrow [0, 1] . \quad (2.4)$$

2.5.1.5. SuperLatticeF

These functors are defined on the lattice via

$$\text{SuperLatticeF3D} : \mathbb{N}^3 \rightarrow \mathbb{R}^d , \quad d \in \mathbb{N} , \quad (2.5)$$

and commonly represent the raw simulation data, e.g. macroscopic moments such as pressure and velocity. `SuperLattice` functors are part of the parallelism layer and they delegate the calculations to the corresponding `BlockLattice` functors.

2.5.1.6. InterpolationF

Interpolation functors establish conversion between the analytical and lattice functors. They are very important in setting analytical boundary conditions, by evaluating the given analytical function on the lattice points. The reverse direction – from lattice to analytical functors – is where this functor receives its name, as the conversion is achieved by interpolation between the lattice points.

2.5.2. Functor Arithmetic

Functor arithmetic expressions are based on functor instances wrapped in `std::shared_ptr<F>` smart pointers. This layer is necessary to enable automatic memory management for trees of interdependent functors.

```

1 std::shared_ptr<SuperF3D<T>> aF(
2     new SuperConst3D<T>(superStructure, {1.0, 2.0}));
3 std::shared_ptr<SuperF3D<T>> bF(
4     new SuperConst3D<T>(superStructure, {2.0, 1.0}));
5 std::shared_ptr<SuperF3D<T>> cF = aF + bF;
6 // cF->operator() returns {3.0, 3.0}

```

Listing 2.2.: Basic showcase for `std::shared_ptr` based functor arithmetic

Note that `cF` can be passed out of scope without any regard for `aF` and `bF`, as managed pointers are stored internally. At first glance this new functor arithmetic may seem unnecessarily verbose for basic usage such as simply adding two functors and directly using the result. As such legacy functor arithmetic is still available for basic use cases. Usage of `std::shared_ptr` functor arithmetic is supported by both `FunctorPtr` and a DSL to ease development of more complex functor compositions.

The `FunctorPtr` helper template is used throughout the functor codebase to transparently accept functors independently of how their memory is managed. This means that functors managed by `std::shared_ptr` are accepted as arguments in any place where raw functor references were used previously. As a nice benefit `FunctorPtr` transparently forwards any calls of its own operator function to the operator of the underlying functor.

```

1 T error(FunctorPtr<SuperF3D<T>>&& f, T reference) {
2     T output[1] = { };
3     const int origin[4] = {0,0,0,0};
4     f(output, origin);
5     return fabs(output[0] - reference);
6 }
7
8 std::shared_ptr<SuperF3D<T>> managedF(
9     new SuperConst3D<T>(superStructure, 1.0));
10 SuperConst3D<T> rawF(superStructure, 1.0);
11
12 // error(managedF, 1.1) == error(rawF, 1.1) == 0.1

```

Listing 2.3.: `FunctorPtr` and `std::shared_ptr` based functor arithmetic

Functor arithmetic expressions may also contain constants in addition to functors. Any scalar constant used in the context of managed functor arithmetic is implicitly wrapped into a `SuperConst(2,3)D` instance.

```

1 std::shared_ptr<SuperF3D<T>> aF(/*...*/);
2 auto bF = 0.5 * aF + 2.0; // scalar multiplication and addition

```

Listing 2.4.: Constant scalars in managed functor arithmetic

Constant vectors are also supported if they are explicitly passed to the `SuperConst(2,3)D` constructor. Note that arithmetic operations of equidimensional functors are performed componentwise (i.e. `aF * aF` is not the scalar product).

```

1 std::shared_ptr<SuperF3D<T>> vectorF(
2     new SuperConst3D<T>(superStructure, {1.0, 2.0, 3.0}));
3 auto cF = aF / vectorF; // componentwise division

```

Listing 2.5.: Constant vectors in managed functor arithmetic

2.5.2.1. Functor Composition

Composing multiple managed functors which in turn need multiple arguments by themselves such as when calculating error norms in a reusable fashion can quickly lead to expressions that are fairly hard to read. For this reason, the `functor_dsl` namespace offers a set of conveniently named helper functions in order to deobfuscate such functor compositions.

Consider for example the following snippet which constructs and evaluates a relative error functor based on the L2 norm.

```

1 using namespace functor_dsl;
2 // decltype(wantedF) == std::shared_ptr<AnalyticalF3D<double,double>>
3 // decltype(f) == std::shared_ptr<SuperF3D<double>>
4 // decltype(indicatorF) == std::shared_ptr<SuperIndicatorF3D<double>>
5 auto wantedLatticeF = restrictF(wantedF, sLattice);
6 auto relErrorNormF = norm<2>(wantedLatticeF - f, indicatorF)
7                       / norm<2>(wantedLatticeF, indicatorF);
8 const int input[4];
9 double result[1];
10 relErrorNormF->operator()(result, input);
11 std::cout << "Relative error: " << result[0] << std::endl;

```

Listing 2.6.: `functor_dsl` supported functor composition

Note that lines 5–7 contain the full implementation of the expression

$$\frac{\| \text{wantedF} - f \|_2}{\| \text{wantedF} \|_2}, \quad (2.6)$$

i.e. the L2-normed relative error of an arbitrary functor `f` as compared to the analytical solution `wantedF`. This simply allows for moving even basically one-off functor compositions into reusable and easily verifiable functors whose implementation is as close to the actual mathematical definition as is reasonably possible. Correspondingly a more developed version of Listing 2.6 can be found in `SuperRelativeErrorLpNorm3D` which is used extensively by the `poiseuille3d` example to compare simulated and analytical solutions.

```

1 template <typename T, typename W, int P>
2 template <template <typename U> class DESCRIPTOR>
3 SuperRelativeErrorLpNorm3D<T,W,P>::SuperRelativeErrorLpNorm3D(
4     SuperLattice<T,DESCRIPTOR>& sLattice,
5     FunctorPtr<SuperF3D<T,W>>&& f,
6     FunctorPtr<AnalyticalF3D<T,W>>&& wantedF,
7     FunctorPtr<SuperIndicatorF3D<T>>&& indicatorF)

```

```

8   : SuperIdentity3D<T,W>([&]())
9   {
10      using namespace functor_dsl;
11
12      auto wantedLatticeF = restrictF(wantedF.toShared(), sLattice);
13
14      return norm<P>(wantedLatticeF-f.toShared(), indicatorF.toShared())
15                  / norm<P>(wantedLatticeF, indicatorF.toShared());
16  }()
17  {
18      this->getName() = "relErrorNormL" + std::to_string(P);
19  }

```

Listing 2.7.: Functor composition in `SuperRelativeErrorLpNorm3D`'s constructor

Disregarding the addition of `FunctorPtr` as well as further templatization, lines 12–15 are equivalent to the ad-hoc error norm in Listing 2.6. Also note how the actual composition happens inside of a lambda expression and is then returned to be stored by `SuperIdentity3D`. This allows for assigning composed functors their own name and renders them indistinguishable from *primitive* functors.

2.6. Parallelization

As applications in computational fluid dynamics require a large amount of resources, it is essential to have the flexibility to switch to a parallel platform easily. This section concentrates on parallelism on distributed memory machines using MPI, as distributed memory is the most common model on large-scale, parallel machines. All example cases in the OpenLB distribution can be compiled with MPI and executed in parallel. Data which is spatially distributed, such as lattice fields, is handled through a data-parallel paradigm. The data space is partitioned into smaller regions that are distributed over the nodes of a parallel machine. In the following, these types of structures are referred to as data-parallel structures. Other data types that require a small amount of storage space are duplicated on every node. These are referred to as duplicated data. All native C++ data types are automatically duplicated by virtue of the Single-Program-Multiple-Data model of MPI. These types should be used to handle scalar values, such as the parameters of the simulation, or integral values over the solution (e.g. the average energy).

For output on the console it is recommended to use `OstreamManager` since it transparently manages output in case of parallel execution (cf Chapter 6.6).

2.6.1. Supported Platforms

OpenLB models heterogeneous parallelization using its predominant block decomposition architecture. Each individual block of a super lattice is assigned one of currently three possible target platforms: `Platform::CPU_SISD`, `Platform::CPU_SIMD` or `Platform::GPU_CUDA`. The availabil-

ity of each platform is controlled by adding its name to the `PLATFORMS` variable in the `config.mk` build configuration. Note that further system specific changes to the compiler settings are almost certainly required for `CPU_SIMD` and `GPU_CUDA`. See the `config/` folder for some examples.

Note that `CPU_SIMD` must always be enabled as some host-side data structures rely on this platform. In the absence of heterogeneous load balancing, the *most efficient* enabled platform is used by default for all blocks. Specifically, `CPU_SIMD` if only the two CPU platforms are enabled or `GPU_CUDA` (even if `CPU_SIMD` is also enabled at the same time).

2.6.1.1. SIMD

Modern CPUs offer vector instructions with a width of up to 512 bytes in the case of AVX-512. This means that 8 respectively 16 individual scalar values can be processed in a single instruction. In some situations this can significantly speed up the bulk collision step, which is why OpenLB supports this option for processing its `Dynamics`. This option is at its most powerful when combined with the `HYBRID` parallelization mode s.t. OpenMP is used to further parallelize the vectorized collision on each shared memory node. However, setting up the hybrid mode correctly on a HPC system is non trivial, which is why we suggest to stick to the MPI-only mode for users unexperienced in working with HPC systems.

Once enabled in the build configuration, vectorization is applied to the dominant collision of each block lattice transparently without requiring any additional code changes.

2.6.1.2. GPU

General purpose graphics processing units (GPGPUs) are an ideal platform for many LBM-based simulations due to their high memory bandwidth and high degree of parallelization. OpenLB currently supports transparent usage of Nvidia GPUs via CUDA for almost all dynamics and a core set of boundary post processors and coupling operators. Similarly to both other available platforms, enabling GPU support requires only adding `GPU_CUDA` to the `PLATFORMS` variable in `config.mk` and some system-specific updates to the compiler settings. MPI parallelization is fully supported for GPU blocks, enabling simulations on multi-GPU clusters.

Note that the resolution commonly needs to be increased significantly compared to the CPU-accommodating default values in order to take full advantage of GPU acceleration. Especially on desktop-grade GPUs changing the fundamental floating point type `T` to `float` is advisable. Also see Section 9.2.1 for guidance on how to install, compile and run GPU-aware OpenLB cases.

3. Geometry Creation and Meshing

This chapter presents how geometry data can be loaded in or created by OpenLB. Furthermore, the concept of material numbers is shown.

3.1. Creating a Geometry

OpenLB provides an interface for STL based geometry data and generates fully automated a structured regular mesh. On the other hand, geometries can be build out of primitive shapes such as cuboids, spheres and cylinders. By the implemented arithmetic that includes intersection, union and complement, those primitives can be assembled very generally. A computational domain such as the `SuperLattice` is created in 6 simple steps (see also Fig 3.1):

Step 1: Create an `Indicator3D` instance by

1. Reading an STL file, see example `aorta3d` 8.11.1 and section 3.4.
2. Pre-defined primitive shapes (cuboid, circle, cylinder, sphere) and their combinations (+, -, ·, /), as described in the example `venturi3d`, see Section 8.11.5.

Step 2: Construct `CuboidGeometry3D`. During construction, the geometry from step 1 is divided into the predefined number of cuboids that are thereafter automatically removed, shrunk and weighted for a good load balance. By weighting user can choose between *weight* and *volume* strategies. For complex shapes the last option is preferable. A larger cuboids number removes more unnecessary nodes, but implies higher communication costs.

Step 3: Construct `LoadBalancer` that assigns cuboids to threads. The standard option is the `HeuristicLoadBalancer`, whereby there are also other variants.

Step 4: Construct `SuperGeometry3D` that links material numbers to voxels (see Section 3.2).

Step 5: Set *material numbers* to different simulation space regions, where afterwards dynamics and boundaries are defined.

Step 6: Construct `SuperLattice` to perform stream and collide algorithm.

```
1 // Step 1: Create Indicator
2 STLreader<T> stlreader("filename.stl", voxelSize, stlSize, method);
3 IndicatorCuboid3D<T> indicator( extend, origin );
4 // Step 2: Construct cuboidGeometry.
```

```

5 CuboidGeometry3D<T> cuboidGeometry(stlReader / indicator, voxelSize, noOfCuboids,
   weightingStrategy);
6 // Step 3: Construct LoadBalancer.
7 HeuristicLoadBalancer<T> loadBalancer(cuboidGeometry);
8 // Step 4: Construct SuperGeometry.
9 SuperGeometry<T,3> superGeometry(cuboidGeometry, loadBalancer);
10 // Step 5: Set material numbers.
11 // set material number 2 for whole geometry
12 superGeometry.rename(0,2);
13 // change material number from 2 to 1 for inner (fluid) cells, so that only boundary
   cells have material number 2
14 superGeometry.rename(2,1,{1,1,1});
15 // or simply use an indicator that changes its lattices to one
16 superGeometry.rename(2,1,fluidIndicator);
17 // additional material numbers for other boundary conditions, the 3rd argument in the
   brackets is the material number which the boundary cells should face
18 superGeometry.rename(2,3,1,cylinderInFlow);
19 superGeometry.rename(2,4,1,outflowIndicator0);
20 superGeometry.rename(2,5,1,outflowIndicator1);
21 // Step 6: Construct SuperLattice.
22 SuperLattice<T,DESCRIPTOR> sLattice(superGeometry);

```

Listing 3.1.: Create geometry based on STL or primitive shapes. All six steps are presented briefly as source code.

The powerful application of the geometry generation of OpenLB can be demonstrated on the example `aorta3d`. This example is based on a very complex geometry and illustrates the highly user friendly and automated process from STL to computation grid the `SuperLattice`, see Figure 3.1.

3.2. Setting Material Numbers

OpenLB has a general concept for representation of a geometry. A specific number called the *material number* is assigned to each cell, defining whether that cell belongs to the boundary or to the fluid domain or whether it is superfluous in the computations. Figure 3.2 illustrates this using the example of a channel flow with an obstacle. The different collision and streaming steps on the boundary and the fluid are defined with respect to the material number. The benefit of using material numbers in CFD simulations is the automatic determination of streaming directions on boundary nodes, as this is not always practical by hand e.g. if material numbers of a complex geometry are obtained from a STL file.

Besides creating the domain, `IndicatorFXD` functions can be used to set *material numbers* with the help of one of the `rename` functions in `SuperGeometryXD`.

```

1 /// replace one material with another
2 void rename(int fromM, int toM);
3 /// replace one material that fulfills an indicator functor condition with another
4 void rename(int fromM, int toM, IndicatorF3D<bool,T>& condition);
5 /// replace one material with another respecting an offset (overlap)

```





Figure 3.1.: Six steps to create a Geometry. It starts by reading an STL file with the help of an STLreader and ends with the creation of a SuperLattice.

```

6 void rename(int fromM, int toM, { unsigned offsetX, unsigned offsetY, unsigned offsetZ
    });
7 /// renames all voxels of material fromM to toM if the number of voxels given by
    testDirection is of material testM
8 void rename(int fromM, int toM, int testM, std::vector<int> testDirection);
9 /// renames all voxels of material fromM to toM if two neighbour voxels in the
    direction of the discrete normal are voxels with material testM in the region
    where the indicator function is fulfilled
10 void rename(int fromM, int toM, int testM, IndicatorF3D<bool,T>& condition);

```

Listing 3.2.: Different rename functions to set material numbers.

3.3. Building Simulation Domains with Geometry Primitive Functors

For the purpose of setting up simulation domains (here called geometries), OpenLB provides several functors (see Section 2.5) for the creation of basic geometric primitives such as cuboids, circles, spheres, cones etc. These can be combined using the mathematical operators (+ union, − set difference, · intersection) to create more complex domains. It can be done in the application setup (see below) or in the XML interface (see Section 8.11.5).

```

1 Vector<T,3> C0(0,50,50);
2 Vector<T,3> C1(5,50,50);

```


should be set to the units of the STL part. If it is exported in meters, the scaling factor is 1, if in millimeters then it is 0.001. The reading methods can be chosen depending on the geometry complexity. For easy geometries the option 0 can be chosen, for the complex and possibly untight shapes the option 1 is to be set, whereby it is slower. The *verbose* argument can be true or false. It prints the information about the STL file in the terminal.

3.5. Excursus: Creating STL-files

The general process chain assumes that the geometry is already given in form of an STL file, if not created by the `IndicatorFXD`-functions. Simple geometries can be created using a CAD tool like FreeCAD (www.freecad.org). An introduction to modeling with FreeCAD can be found for example in <http://www.youtube.com/watch?v=geIrH1cOCzc>. The general procedure is mostly similar to the following description.

Firstly, a 2D sketch is created on a selected plane (e.g. the *xy*-plane) using different bi-dimensional shapes. In the next step the sketch is extruded in the third dimension. Several such 3d objects can be combined using operations like union, cut, intersection, rotation, trace, etc. to obtain the target geometry. Creating a square and a circle for the example `cylinder3d` in Figure 3.3 is an easy task. Also, complex geometries as that of a filter or a porous media can be set up easily with OpenLB's indicator approach.




Figure 3.3.: The geometry of the example `cylinder3d` from Section 8.4.3 opened in FreeCAD.

4. Simulation Models

4.1. Non-dimensionalization and Choice of Simulation Parameters

Basically, to describe a physical quantity you need a reference scale. By dividing a quantity by a variable of the same dimension, a dimensionless quantity is derived. The result is a number, which is called "the lattice value of the quantity" or the value of the amount "in lattice units" (denoted with \cdot_{LB}), if you consider the Lattice Boltzmann Method. The reference scale is named conversion factor between two reference systems. Furthermore, the conversion factor for the quantity ϕ is called C_ϕ and the non-dimensional quantity gets the symbol ϕ^* .

To get into this topic, the book *The Lattice Boltzmann Method* [100], especially Chapter 7.1 and 7.2 therein, and https://en.wikipedia.org/wiki/Dimensional_analysis are strongly recommended.

4.2. Porous Media Model

The permeability parameter K is a physical parameter that describes the macroscopic drag in a porous media model. For laminar flows it is defined by Darcy's law

$$K = -\frac{Q\mu L}{\Delta P}, \quad (4.1)$$

where $Q = UA$ is the flow rate, U is a characteristic velocity, A is a cross-sectional area, μ denotes the dynamic viscosity, L is a characteristic length, and ΔP specifies the pressure difference in between starting point and endpoint of the volume.

The porosity-value $d \in [0, 1]$ is a lattice-dependent value, where $d = 0$ means the medium is solid and $d = 1$ denotes a liquid. According to Brinkman Brinkman [82, 83], Borrvall and Petersson [79] and Pinget et al. [114], the Navier-Stokes equation is transformed (see Dornieden [88] and Stasius [121]). The discrete formulation of d describes a flow region by its permeability

$$d = 1 - h^{dim-1} \frac{\nu_{\text{LB}} \tau_{\text{LB}}}{K}, \quad (4.2)$$

where τ_{LB} is the relaxation time, ν_{LB} is the discrete kinematic viscosity and h is the length. Therefore $K \in [\nu_{\text{LB}} \tau_{\text{LB}} h^{dim-1}, \infty]$. To describe the porosity or permeability of a medium, a

descriptor for porosity must be used, such as the one below.

```
1 using DESCRIPTOR = descriptors::PorousD3Q19Descriptor;
```

Be aware that the porous media model only works in the generic compilation mode. In the function `prepareLattice`, dynamics for the corresponding number of the porous material are defined for example as follows.

```
1 void prepareLattice(..., Dynamics<T, DESCRIPTOR>& porousDynamics, ...){
2     /// Material=3 --> porous material
3     sLattice.defineDynamics(superGeometry, 3, &porousDynamics);
4     ...
5 }
```

In the function `setBoundaryValues`, the initial porosity value and external field are defined.

```
1 void setBoundaryValues(..., T physPermeability, int dim, ...){
2     // d in [0,1] is a lattice-dependent porosity-value
3     // depending on physical permeability K = physPermeability
4     T d = converter->latticePorosity(physPermeability);
5     AnalyticalConst3D<T,T> porosity(d);
6     sLattice.defineField<descriptors::POROSITY>(superGeometry, 3, porosity);
7     ...
8 }
```

In the `main` function, the required parameters as well as the porous media dynamics are defined.

```
1 int main(int argc, char* argv[]) {
2     ...
3     T physPermeability = 0.0003;
4     ...
5     PorousBGKdynamics<T, DESCRIPTOR> porousDynamics(converter->getOmega(),
6         instances::getBulkMomenta<T, DESCRIPTOR>());
7     ...
8 }
```

Additionally, the class `UnitConverter` in `src/core/units.h` provides useful functions for conversion between physical and lattice values:

```
1 /// converts a physical permeability K to a lattice-dependent porosity d
2 /// (a velocity scaling factor depending on Maxwellian distribution
3 /// function), needs PorousBGKdynamics
4 T latticePorosity(T K) const
5 { return 1 - pow(physLength(),getDim()-1)*getLatticeNu()*getTau()/K; }
6
7 /// converts a lattice-dependent porosity d (a velocity scaling factor
8 /// depending on Maxwellian distribution function) to a physical
9 /// permeability K, needs PorousBGKdynamics
10 T physPermeability(T d) const
11 { return pow(physLength(),getDim()-1)*getLatticeNu()*getTau()/(1-d); }
```

4.3. Power Law Model

The two most common deviations from Newton's Law observed in real systems are pseudo-plastic fluids and dilatant fluids. For pseudo-plastic fluids the viscosity of the system decreases as the shear rate is increased. On the other hand, as the shear rate by dilatant fluids is increased, the viscosity of the system also increases. The simplest model, that describes these two types of deviations, is called the (Ostwald–de Waele) Power Law model and defines the viscosity as

$$\mu = m\dot{\gamma}^{n-1}, \quad (4.3)$$

where m is the flow consistency index, $\dot{\gamma}$ is the shear rate, and n denotes the flow behavior index. Then

- $n < 1$: pseudoplastic fluids,
- $n = 1$: Newtonian fluids,
- $n > 1$: dilatant fluids.

To simulate a power law fluid, a descriptor for dynamic `omega` must be used, such as:

```
1 using DESCRIPTOR = descriptors::DynOmegaD2Q9Descriptor;
```

In the function `setBoundaryValues`, the initial same `omega`-argument is defined.

```
1 AnalyticalConst2D<T,T> omega0(converter.getOmega());
2 sLattice.defineField<descriptors::OMEGA>(
3   superGeometry.getMaterialIndicator({1,2,3,4}), omega0);
```

In the `main` function, the power law dynamics are defined.

```
1 int main(int argc, char* argv[]) {
2   ...
3   PowerLawBGKdynamics<T, DESCRIPTOR> bulkDynamics(converter.getOmega(), instances::
4     getBulkMomenta<T, DESCRIPTOR>(), m, n, converter.physTime());
5 }
```

To recover a nonconstant kinematic viscosity in the targeted PDE, also the `omega`-argument has to be no longer a constant. With using the power law model (4.3), the kinematic viscosity is computed in each step as

$$\nu = \frac{1}{\rho} m \dot{\gamma}^{n-1}. \quad (4.4)$$

The shear rate $\dot{\gamma}$ is computable via using the second invariant of the strain rate tensor D_{II} , i.e.

$$\dot{\gamma} = \sqrt{2D_{II}}, \quad (4.5)$$

where

$$D_{II} = \sum_{\alpha, \beta=1}^d E_{\alpha\beta} E_{\alpha\beta} , \quad (4.6)$$

and

$$E_{\alpha\beta} = - \left(1 - \frac{1}{\tau} \right) \frac{1}{2\rho\nu} \sum_{i=0}^{q-1} f_i^h \xi_{i\alpha} \xi_{i\beta} . \quad (4.7)$$

This concept is very significant, since $f_i^h \xi_{i\alpha} \xi_{i\beta}$ is usually computed during the collision process and therefore, in comparison to other CFD methods, adds almost no additional computational cost. The computation of a new `omega`-argument is done in `src/dynamics/powerLawBGKdynamics.h` via:

```
1 T computeOmega(T omega0_, T preFactor_, T rho_, T pi_[util::TensorVal<DESCRIPTOR >::n]
    ) ;
```

4.4. Multiphysics Couplings

4.4.1. Shan–Chen Model

For the simulation of both multiphase and multicomponent flow the Shan–Chen model is implemented in OpenLB. Since its first introduction [117], many variants of the model have been developed. In this implementation, there are several forcing schemes [92, 118] and interaction potentials to choose from.

4.4.1.1. Implementation of Shan–Chen Two-phase Fluid

The two phases can be simulated on the same lattice instance.

```
1 SuperLattice<T, DESCRIPTOR> sLattice(superGeometry);
```

Then the dynamics are chosen, which have to support external forces.

```
1 ForcedShanChenBGKdynamics<T, DESCRIPTOR> bulkDynamics1 (
2 omegal, instances::getExternalVelocityMomenta<T, DESCRIPTOR>() );
```

Possible choices for the dynamics are `Forced-BGK-dynamics` and `ForcedShanChen-BGK-dynamics`.

Then the interaction potential is chosen.

```
1 ShanChen93<T, T> interactionPotential;
```

Viable interaction potentials for one component multiphase flow are `ShanChen93`, `ShanChen94`, `CarnahanStarling` and `PengRobinson`. In this model `PsiEqualsRho` should not be used, because this would make all the mass gather in the same place.

To enable interaction between the fluid, they have to be coupled, so the kind of coupling has to be chosen (here: `ShanChenForcedSingleComponentGenerator3D`) and the material numbers to which it applies. Since in the case of single component flow there is only one lattice, it is coupled with itself.

```

1 const T G      = -120.;
2 ShanChenForcedSingleComponentGenerator3D<T,DESCRIPTOR> coupling(
3   G, rho0, interactionPotential);
4 sLattice.addLatticeCoupling(superGeometry, 1, coupling, sLattice);

```

The interaction strength G has to be negative and the correct choice depends on the chosen interaction potential. When using `PengRobinson` or `CarnahanStarling` interaction potential, G is canceled out during computation, so the result is not affected by it (though it still has to be negative).

Finally, during the main loop the lattices have to interact with each other (or in the case of only one fluid component the lattice with itself).

```

1 sLattice.communicate();
2 sLattice.executeCoupling();

```

These steps are placed immediately after the `collideAndStream` command.

Examples for the implementation of a LB simulation using the Shan–Chen model for two-phase flow are `examples/phaseSeparation2d` and `examples/phaseSeparation3d`.

4.4.1.2. Implementation of Shan–Chen Two-component Fluid

Two lattice instances are needed, one for each component (though operating on one geometry).

```

1 SuperLattice<T, DESCRIPTOR> sLatticeOne(superGeometry);
2 SuperLattice<T, DESCRIPTOR> sLatticeTwo(superGeometry);

```

Then the dynamics are chosen, which have to support external forces.

```

1 ForcedShanChenBGKdynamics<T, DESCRIPTOR> bulkDynamics1 (
2   omega1, instances::getExternalVelocityMomenta<T,DESCRIPTOR>() );
3 ForcedShanChenBGKdynamics<T, DESCRIPTOR> bulkDynamics2 (
4   omega2, instances::getExternalVelocityMomenta<T,DESCRIPTOR>() );

```

Possible choices for the dynamics are `ForcedBGKdynamics` and `ForcedShanChenBGKdynamics`. One should keep in mind that tasks like definition of dynamics, external fields and initial values and the collide and stream execution have to be carried out for each lattice instance separately. The same is true for data output. Then the interaction potential is chosen.

```

1 PsiEqualsRho<T,T> interactionPotential;

```

In the multicomponent case, the most frequently used interaction potential is `PsiEqualsRho`, but `ShanChen93` for example, would also be a viable choice.

To enable interaction between the fluid, they have to be coupled, so the kind of coupling has to be chosen (here: `ShanChenForcedGenerator3D`) and the material numbers to which it applies.

```

1 const T G      = 3.;
2 ShanChenForcedGenerator3D<T,DESCRIPTOR> coupling(
3   G, rho0, interactionPotential);
4 sLatticeOne.addLatticeCoupling(superGeometry, 1, coupling, sLatticeTwo);
5 sLatticeOne.addLatticeCoupling(superGeometry, 2, coupling, sLatticeTwo);

```

The interaction strength G has to be positive. If the chosen interaction potential is `PsiEqualsRho`, $G > 1$ is needed for separation of the fluids, but it should not be much higher than 3 for stability reasons.

Finally, during the main loop the lattices have to interact with each other.

```

1 sLatticeOne.communicate();
2 sLatticeTwo.communicate();
3 sLatticeOne.executeCoupling();

```

These steps are placed immediately after the `collideAndStream` command.

Examples for the implementation of a LB simulation using the Shan–Chen model for two-component flow are `examples/multiComponent2d` and `examples/multiComponent3d`.

4.4.2. Free Energy Model

As an alternative option for simulating multi-component flow, the free energy model has been implemented into OpenLB, and can be used for either two or three fluid components. Examples for the binary case are given in `youngLaplaceXd` and `contactAngleXd`, while an example of the ternary case with boundaries is provided in `microFluidics2d`. These are all contained within the `examples/multiComponent/` folder.

The approach taken in OpenLB is similar to that given in [116] and assumes equal densities and viscosities for each of the fluids. In the next sections the method will be outlined briefly for three components. The two component case is identical to taking the third fluid component to be zero and instead only uses two lattices.

4.4.2.1. Bulk Free Energy Model

Three lattices are required to track the density ρ , and order parameters ϕ and ψ . These are related to the individual component densities C_i by

$$\rho = C_1 + C_2 + C_3, \quad \phi = C_1 - C_2, \quad \psi = C_3, \quad (4.8)$$

respectively. By considering the free energy, a force is derived to drive the fluid towards the thermodynamic equilibrium. The density therefore obeys the Navier–Stokes equation with this added force. The equation of motion for the order parameters is the Cahn–Hilliard equation. The dynamics chosen for the first lattice must therefore include an external force, such as `ForcedBGKdynamics`, while for the second and third lattices `FreeEnergyBGKdynamics` is required.

```

1 ForcedBGKdynamics<T, DESCRIPTOR> bulkDynamics1 (
2     omega, instances::getBulkMomenta<T, DESCRIPTOR>() );
3 FreeEnergyBGKdynamics<T, DESCRIPTOR> bulkDynamics23 (
4     omega, gamma, instances::getBulkMomenta<T, DESCRIPTOR>() );

```

To compute the force, two lattice couplings are required. The first computes the chemical potentials for each lattice using the equations

$$\mu_\rho = A_1 + A_2 + \frac{\alpha^2}{4} [(\kappa_1 + \kappa_2)(\nabla^2\psi - \nabla^2\rho) + (\kappa_2 - \kappa_1)\nabla^2\phi] , \quad (4.9)$$

$$\mu_\phi = A_1 - A_2 + \frac{\alpha^2}{4} [(\kappa_2 - \kappa_1)(\nabla^2\rho - \nabla^2\psi) - (\kappa_1 + \kappa_2)\nabla^2\phi] , \quad (4.10)$$

$$\begin{aligned} \mu_\psi = & -A_1 - A_2 + \kappa_3\psi(\psi - 1)(2\psi - 1) + \frac{\alpha^2}{4} [(\kappa_1 + \kappa_2)\nabla^2\rho \\ & - (\kappa_2 - \kappa_1)\nabla^2\phi - (\kappa_1 + \kappa_2 + 4\kappa_3)\nabla^2\psi] , \end{aligned} \quad (4.11)$$

where A_1 and A_2 are defined as

$$A_1 = \frac{\kappa_1}{8}(\rho + \phi - \psi)(\rho + \phi - \psi - 2)(\rho + \phi - \psi - 1) , \quad (4.12)$$

$$A_2 = \frac{\kappa_2}{8}(\rho - \phi - \psi)(\rho - \phi - \psi - 2)(\rho - \phi - \psi - 1) , \quad (4.13)$$

respectively. The α the κ parameters are input parameters for the lattice coupling and can be used to tune the interfacial width and surface tensions. The interfacial width is given by α and the surface tensions are $\gamma_{mn} = \alpha(\kappa_m + \kappa_n)/6$.

The second lattice coupling then computes the force using

$$\mathbf{F} = -\rho\nabla\mu_\rho - \phi\nabla\mu_\phi - \psi\nabla\mu_\psi . \quad (4.14)$$

This depends non-locally upon the result of the first lattice coupling, and so the chemical potential must be communicated in between. To accommodate this, the first coupling is assigned to the ρ lattice and the force coupling is assigned to the ϕ lattice.

```

1 FreeEnergyChemicalPotentialGenerator3D<T, DESCRIPTOR> coupling1(
2     alpha, kappal, kappa2, kappa3 );
3 FreeEnergyForceGenerator3D<T, DESCRIPTOR> coupling2;
4
5 sLattice1.addLatticeCoupling<DESCRIPTOR>(
6     superGeometry, 1, coupling2, {&sLattice2, &sLattice3} );
7 sLattice2.addLatticeCoupling<DESCRIPTOR>(
8     superGeometry, 1, coupling3, {&sLattice1, &sLattice3} );

```

The following is then used in the main loop to calculate the force at each timestep.

```

1 sLattice1.executeCoupling();
2 sExternal1.communicate();
3 sExternal2.communicate();
4 sExternal3.communicate();

```

```
5 sLattice2.executeCoupling();
```

4.4.2.2. Boundaries in Free Energy Models

Using the `setFreeEnergyWallBoundary` function, bounce-back wall boundaries with controllable contact angles can be added via the following commands:

```
1 setFreeEnergyWallBoundary<T,DESCRIPTOR>(sLattice1, superGeometry, 2,
2   alpha, kappa1, kappa2, kappa3, h1, h2, h3, 1);
3 setFreeEnergyWallBoundary<T,DESCRIPTOR>(sLattice2, superGeometry, 2,
4   alpha, kappa1, kappa2, kappa3, h1, h2, h3, 2);
5 setFreeEnergyWallBoundary<T,DESCRIPTOR>(sLattice3, superGeometry, 2,
6   alpha, kappa1, kappa2, kappa3, h1, h2, h3, 3);
```

The final parameter, `latticeNumber`, is necessary to change each lattice differently. While the h_i parameters are related to the contact angles at the boundary. The contact angles θ_{mn} , are given by the following relation

$$\cos \theta_{mn} = \frac{(\alpha\kappa_n + 4h_n)^{3/2} - (\alpha\kappa_n - 4h_n)^{3/2}}{2(\kappa_m + \kappa_n)\sqrt{\alpha\kappa_n}} - \frac{(\alpha\kappa_m + 4h_m)^{3/2} - (\alpha\kappa_m - 4h_m)^{3/2}}{2(\kappa_m + \kappa_n)\sqrt{\alpha\kappa_m}}. \quad (4.15)$$

Notably, to set neutral wetting (90° angles), the values can be set to $h_i = 0$.

A demonstration of using these solid boundaries for a binary fluid case is provided in the `contactAngle(2,3)d` examples. The examples compare the simulated angles to those given by equation (4.15), respectively for dimensions $d = 2, 3$.

Open boundary conditions can also be implemented using the `setFreeEnergyInletBoundary` and `setFreeEnergyOutletBoundary` functions. These can be used to specify constant density or velocity boundaries. The first lattice is used to define the density or velocity boundary condition, while on the second and third lattices ϕ and ψ must instead be defined. For example, to set a constant velocity inlet, see the code snippet below.

```
1 setFreeEnergyInletBoundary(
2   sLattice1, omega, inletIndicator, "velocity", 1 );
3 setFreeEnergyInletBoundary(
4   sLattice2, omega, inletIndicator, "velocity", 2 );
5 setFreeEnergyInletBoundary(
6   sLattice3, omega, inletIndicator, "velocity", 3 );
7
8 sLattice1.defineU( inletIndicator, 0.002 );
9 sLattice2.defineRho( inletIndicator, 1. );
10 sLattice3.defineRho( inletIndicator, 0. );
```

However, this alone is insufficient to set a constant density outlet because ρ , ϕ , and ψ are redefined by a convective boundary condition on each time step. In this case an additional lattice coupling is required, using `FreeEnergyDensityOutletGeneratorXD`.

There are two additional requirements for open boundaries. The first is that the velocity must be coupled between the lattices using `FreeEnergyInletOutletGeneratorXD`, because this is

required for the collision step. The second is that the communication of the external field must now include two values. This ensures that ρ , ϕ , and ψ are properly set on block edges at the outlet. To see a full example of applying these boundary conditions, see the `microFluidics2d` example.

4.4.3. Coupling Between Momentum and Energy Equations

As explained in reference [111], there are different schemes to couple the momentum and energy equations by means of a buoyancy force (also called Boussinesq approximation). Some schemes add an extra force term to the collision term, other methods shift the velocity field according to Newton's second law, and others combine an extra force term and a velocity shift. The implementation applied in OpenLB belongs to this last group of schemes.

Once the boundary values for the velocity and temperature fields are set, collision and streaming functions are called. The dynamics with an external force \mathbf{F} used for the velocity calculation (e.g. `ForcedBGKdynamics`) shifts the velocity \mathbf{v} before executing the collision step. The shift follows equation (4.16),

$$\mathbf{u}_{\text{shift}} = \mathbf{u} + \frac{\mathbf{F}}{2} . \quad (4.16)$$

The code snippet responsible for this shift is defined in the collision function of the file `/dynamics/ics/dynamics.hh` for the class `ForcedBGKdynamics`.

```

1 this->momenta . computeRhoU( cell , rho , u ) ;
2 FieldPtr<T,DESCRIPTOR,FORCE> force = cell.template getFieldPointer<descriptors::FORCE>
   >();
3 for ( int iVe l=0; iVel<DESCRIPTOR >::d; ++iVel ) {
4     u[iVel] += force[iVel] / T{2};
5 }

```

Listing 4.1.: Velocity shift

After the corresponding collision step using the shifted velocity, the value of the density distribution functions f_i is modified by the external force with the call to the function.

```

1 lbm< Lattice >::addExternalForce( cell , u , omega )

```

This function follows

$$\bar{f}_i = f_i + \left(1 - \frac{\omega}{2}\right) w_i \left(\frac{\mathbf{c}_i - \mathbf{u}}{c_s^2} + \frac{\mathbf{c}_i \cdot \mathbf{u}}{c_s^4} \mathbf{c}_i \right) \cdot \mathbf{F} , \quad (4.17)$$

where \bar{f}_i represents the new distribution function (see reference [92] for the BGK model and [90] for MRT models), w_i are the weights of the discrete velocities \mathbf{c}_i , and $\omega = \tau^{-1}$ denotes the relaxation frequency. The coupling in the collision step for the temperature field is given by the use of the velocity from the isothermal field.

```

1 auto u = cell.template getFieldPointer<descriptors::VELOCITY>();

```

The equilibrium density distribution function for the temperature only has terms of first order (see e.g. [42]). After the collision step, the coupling function is called `NSlattice.executeCoupling()`, where the values of the external force in the `NSlattice` and of the advected velocity in the `ADlattice` are updated.

```

1
2 auto u = tPartner->get(iX, iY).template getFieldPointer<descriptors::VELOCITY>();
3 blockLattice.get(iX, iY).computeU(u);

```

Listing 4.2.: Velocity coupling

The new force is computed via the Boussinesq approximation

$$\mathbf{F} = \rho \frac{T - T_0}{\Delta T} \mathbf{g}. \quad (4.18)$$

The temperature T is obtained from the `ADlattice`, T_0 is the average temperature between the defined cold and hot temperatures, whereas ΔT is the difference between the hot and cold temperatures.

```

1
2 auto force = blockLattice.get(iX, iY).template getFieldPointer<descriptors::FORCE>();
3 T temperature = tPartner->get(iX, iY).computeRho();
4 T rho = blockLattice.get(iX, iY).computeRho();
5 for (unsigned iD = 0 ; iD < L :: d ; ++iD) {
6     force[iD] = gravity * rho * (temperature - T0) / deltaTemp * dir[iD];
7 }

```

Listing 4.3.: Computation of the Boussinesq force

4.5. Advection–Diffusion Equation

The advective and diffusive transport of a macroscopic density, energy or temperature is governed by the advection–diffusion equation

$$\frac{\partial c}{\partial t} = \nabla \cdot (D \nabla c) - \nabla \cdot (c \mathbf{v}) \quad \text{in } \Omega \times I, \quad (4.19)$$

where $c: \Omega \times I \rightarrow \mathbb{R}$, $(\mathbf{x}, t) \mapsto c(\mathbf{x}, t)$ is the considered physical quantity (temperature, particle density), $D > 0$ is the diffusion coefficient and \mathbf{v} is a velocity field affecting c . It is possible to approximate this equation with LBM by using an equilibrium distribution function different from the one for the Navier–Stokes equations [86, 37, 38]

$$g_i^{eq} = w_i \rho \left(1 + \frac{\mathbf{c}_i \cdot \mathbf{v}}{c_s^2} \right), \quad (4.20)$$

that takes the advective transport into account. In equation (4.20), w_i is a weighting factor, c_i a unit vector along the lattice directions, c_s the speed of sound, and i denotes the discrete velocity counter. To use this implementation the dynamics object has to be replaced by special advection–diffusion dynamics:

```
1 AdvectionDiffusionBGKdynamics<T, DESCRIPTOR> bulkDynamics (
2     converter.getOmega(),
3     instances::getBulkMomenta<T, DESCRIPTOR>());
```

Listing 4.4.: Advection diffusion dynamics object

Additionally, a different descriptor with fewer lattice velocities is used [96]:

```
1 using DESCRIPTOR = descriptors::AdvectionDiffusionD3Q7Descriptor;
```

Listing 4.5.: Advection diffusion descriptor

In OpenLB, the descriptors `D2Q5` and `D3Q7` are implemented for the advection–diffusion equation. Since the advection–diffusion equation simulates different physical conditions than the Navier–Stokes equations, another set of boundary conditions is needed. A Dirichlet condition for the density is already implemented, for example to simulate a boundary with a constant temperature.

```
1 ParticleAdvectionDiffusionBGKdynamics<T, ADDESCRIPTOR> bulkDynamicsAD ( omegaAD,
2     instances::getBulkMomenta<T, ADDESCRIPTOR>() );
```

Listing 4.6.: Advection diffusion dynamics object

Finally the boundary condition is set to the desired material number.

```
1 void prepareLattice(...) {
2     ...
3     /// Material=3 -> boundary with constant temperature
4     setAdvectionDiffusionTemperatureBoundary<T, ADDESCRIPTOR>(
5         sLatticeAD, omegaAD, superGeometry, 3);
6     ...
7 }
```

Listing 4.7.: Advection diffusion descriptor

To apply convective transport, a velocity vector has to be passed. This can either be done individually on each cell by using the following.

```
1 T velocity[3] = {vx, vy, vz};
2 ...
3 cell.defineField<descriptors::VELOCITY>(velocity);
```

Listing 4.8.: Add advective velocity on a cell

Alternatively, it can be passed to the whole `SuperLattice` using:

```
1 AnalyticalConst3D<T,T> velocity(vel);
2 ...
3 /// sets advective velocity for material 1
4 superLattice.defineField<descriptors::Velocity>(superGeometry, 1, velocity);
```

Listing 4.9.: Add advective velocity on a superlattice

Here, `vel` is a `std::vector<T>`.

4.5.1. AdvectionDiffusion Boundary Conditions

4.5.1.1. Dirichlet Boundary Condition

At the boundaries of a lattice, only the outgoing directions of the distribution functions are known, while those towards the domain need to be computed. Several types of implementations for Dirichlet boundary conditions are summarized in Section 5.1. At a Dirichlet boundary for the advection–diffusion equation, the observable, e.g. temperature, is set to a constant value. This boundary condition can be applied to flat walls, corners and edges (for three-dimensional domains). The algorithm to set a certain temperature on a wall is defined in the dynamics class `AdvectionDiffusionBoundariesDynamics` in the file `boundary/advectionDiffusionBoundaries.hh` and works as described below.

First, the index i of the unknown distribution function g_i incoming to the fluid domain is determined.

```
1 int missingNormal = 0 ;
2 constexpr auto missingDiagonal = util::subIndexOutgoing<L,direction,orientation>();
3 std::vector<int> knownIndexes = util::remainingIndexes<L>(missingDiagonal);
4 for (unsigned iPop = 0; iPop < missingDiagonal.size(); ++iPop)
5 {
6     int numOfNonNullComp = 0;
7     for (int iDim = 0 ; iDim < L::d; ++iDim)
8         numOfNonNullComp += abs (L::c[missingDiagonal[iPop]][iDim]);
9
10    if (numOfNonNullComp == 1)
11    {
12        missingNormal = missingDiagonal[iPop];
13        missingDiagonal.erase(missingDiagonal.begin()+iPop);
14        break;
15    }
16 }
```

Listing 4.10.: Collision step for a temperature boundary

Then, the sum of the rest of the populations is computed.

```
1 T sum = T( );
2 for (unsigned iPop = 0 ; iPop < knownIndexes.size(); ++iPop)
3 {
```

```

4   sum += cell[knownIndexes[iPop]];
5 }

```

The difference between the desired temperature value (given when setting the boundary condition) and this sum is the value assigned to the unknown distribution.

```

1 T temperature = this->momenta.computeRho( cell );
2 cell[missingNormal] = temperature - sum - (T)1;

```

After that, all distribution functions are determined and a regular collision step is performed.

```

1 boundaryDynamics.collide(cell, statistics);

```

As an example, take the case of a left wall in 2D. After the streaming step, all populations are known except for g_3 . Once the desired temperature T_{wall} at the wall is known, the value of the unknown distribution is computed via

$$T_{wall} = \sum_{i=0}^4 g_i \quad (4.21)$$

$$\Leftrightarrow g_3 = T_{wall} - (g_0 + g_1 + g_2 + g_4) . \quad (4.22)$$

4.5.1.2. Neumann Boundary Condition

For flat walls there is also an opportunity to prescribe the flux through the wall as a boundary condition which is

$$\nabla c \cdot \mathbf{n} = \text{flux} . \quad (4.23)$$

Here \mathbf{n} stands for the outer normal vector of the boundary. The flux multiplied with the size of the length discretization has to be written in the field `BOUNDARY`.

```

1 AnalyticalConst2D<T,T>flux_ (converter.getLatticeDensity(flux)*converter.getPhysDeltaX
   ());
2 sLatticeAD.defineField<descriptors::BOUNDARY>(boundaryMaterialIndicator,flux_);

```

The setter `setAdvectionDiffusionNeumannBoundary` for `AdvectionDiffusionNeumannBoundary` adds a preprocessing step which calculates according to the direction of the normal the Dirichlet value which should be at the wall in order to satisfy the flux condition. Therefore a simple backward or forward difference quotient is used as suggested in [100]. As an example, we take again a left wall in 2D and have that

$$c_{wall} = \Delta x \text{flux} + c(x_{wall} + \Delta x) . \quad (4.24)$$

```

1 T analyticalBoundaryValue = cell. template getField<descriptors::BOUNDARY>();
2
3     if constexpr(normal1 !=0 && normal2 == 0 && normal3 == 0){
4         if constexpr(normal1 >0){

```



```

5         // right boundary, difference quotient with x-1
6         neumannBoundary = analyticalBoundaryValue+cell.neighbor({-1,0,0}).
           computeRho();
7     }
8     else if constexpr(normal1<0){
9         //left boundary, difference quotient with x+1
10        neumannBoundary = cell.neighbor({1,0,0}).computeRho() -
           analyticalBoundaryValue;
11    }
12 }

```

Finally we have to set the calculated value for the cell with the `defineRho` command. After this precalculation step the Dirichlet Boundary Condition is applied.

4.5.1.3. Adiabatic Boundary Condition

Additionally, thanks to the simplified lattice velocity sets used in the thermal descriptors (`D2Q5` and `D3Q7`, cf. the discrete velocities in Figure 1.1 colored in red and orange), it is possible to implement an adiabatic boundary using bounce-back dynamics [108].

An adiabatic boundary condition requires no heat conduction in the normal direction of the boundary. In a general situation the adiabatic boundary is set on a solid wall, meaning that the normal velocity to the wall is zero. To implement an adiabatic wall, take a 2D south wall as example. The distribution function g_4 and the temperature at the wall are undetermined. The population g_4 can be computed from the distribution function in the opposite direction, in order to ensure that at the macroscopic level there is no heat conduction, i.e.

$$g_4 = g_2 . \quad (4.25)$$

This procedure corresponds to the bounce-back scheme. With all the distribution functions known, the temperature at the wall can be determined from its definition

$$T_{wall} = \sum_{i=0}^4 g_i . \quad (4.26)$$

4.5.2. Convergence Criterion

For thermal applications, the following convergence criterion can be applied to one of the computed fields or to both of them. Generally, a value tracer on the average energy is used, which is also available for any other quantity and in turn applicable for any TEQ approximated with LBM in OpenLB (see Section 10.5). Here, the average energy is defined proportional to the velocity squared, which makes it independent to use the `NSlattice` or the `ADlattice`, since both share the same macroscopic velocity field \mathbf{u} .

The parameters to initialize the tracer object are the characteristic velocity of the system `converter.getU()`, the characteristic length of the system `converter.getNy()`, and the desired

precision of the convergence `eps`. The listing 4.11 shows how the object is defined in the main function, and how its value is updated and checked at each time step.

```

1 util::ValueTracer<T> converge( converter.getU( ), converter.getNy( ), eps );
2 ...
3 for ( iT=0; iT<maxIter ; ++iT) {
4     converge.takeValue( ADlattice.getStatistics( ).getAverageEnergy( ), true );
5     ...
6     if ( converge.hasConverged() ) {
7         break;
8     }
9 }

```

Listing 4.11.: Convergence check

4.5.3. Creating an Application with AdvectionDiffusionDynamics

When creating a program e.g. for a thermal applications, the following points should be regarded.

4.5.3.1. Lattice Descriptors

Lattice descriptors used for the `AdvectionDiffusionDynamics` are `D2Q5` and `D3Q7`, which have less degrees of freedom for the velocity space than the classical discrete velocity sets for the Navier–Stokes equations (see Figure 1.1). With the Chapman–Enskog expansion it can be shown that approximating the advection–diffusion equation as a target does not require fourth order isotropic lattice tensors (see for example [126]), therefore descriptors with less discrete velocities can be used without loss of accuracy.

To approximate the Navier–Stokes velocity field u , a first descriptor should be defined that can include an external force, e.g. `ForcedD3Q19Descriptor`. Another descriptor is necessary to approximate the temperature T governed by an advection–diffusion equation, e.g. `AdvectionDiffusionD3Q7Descriptor`.

4.5.3.2. Preparing the Geometry

This step is similar to the isothermal procedure (without the second lattice for T). With help of indicators and STL files, the desired geometry can be created. Several material numbers are assigned to the cells that in turn can be implemented to specify several bulk, initial and boundary collision schemes and/or dynamics.

Reading STL files To conveniently read STL files, the OpenLB class `stlReader` is provided. There are differences, however, when compared to the isothermal case, due to the differing converter objects. An example of its use could be:

```
1 STLreader<T> nameIndicator( " fileName.stl ", converter.getDeltaX ( ), conversionFactor
    );
```

Listing 4.12.: Initialization of a STLreader object

The offsets between the STL file and the global geometry are much easier handled, if they are directly defined when creating the STL file, rather than trying to do it in the application code afterwards. Note that, if the geometry is adjusted for a grid resolution of $N = 100$ and a conversion factor of 1, and the resolution is increased to $N = 200$, the corresponding conversion factor also has to be increased by a factor 2, in order to keep the correct proportions of the model.

4.5.3.3. Preparing the Lattices

Recall that in a typical thermal application there are two independent lattices: one for the isothermal flow (usually referred to as `NSlattice`), and one for the thermal variables (e.g. the temperature, usually referred to as `ADlattice`). For each material number the desired dynamics behavior has to be defined. Commonly used possibilities are `instances::getNoDynamics` (do nothing), `instances::getBounceBack` (no slip), or `bulkDynamics` (previously set collision for the bulk). For a thermal lattice the we can define a boundary with a given temperature.

```
1 setAdvectionDiffusionTemperatureBoundary<T, TDESCRIPTOR>(
2     ADlattice, Tomega, superGeometry, 2);
```

Listing 4.13.: Definition of a temperature boundary

The chosen dynamics for a material number may differ between the isothermal and the thermal lattices, e.g. an obstacle with a given temperature inside a flow channel would have a no-slip behavior for the fluid part, but be part of the bulk and have a given temperature in the thermal lattice.

4.5.3.4. Initialization of the Lattices (iT=0)

NSlattice For all material numbers defined as `bulkDynamics`, an initial velocity and density has to be set (usually fluid flow at rest). Additionally, since the velocity field and the temperature are related by a force term, an external field has to be defined. The easiest way to do this is by the material number.

```
1 NSlattice.defineField<descriptors::FORCE>( superGeometry, 1, force );
```

Listing 4.14.: Initialization of an external force field

Here `force` is an element of type `AnalyticalF`, which can initially be set to zero.

ADlattice For the advection–diffusion lattice, an initial temperature is set (similar to the density variable on the Navier–Stokes lattice), as well as the distribution functions corresponding to this temperature value:

```

1 T Texample = 0.5;
2 T zeroVel[descriptors::d<T,DESCRIPTORS>()] = {0., 0.};
3 ConstAnalyticalF2D<T,T> Example( Texample );
4 std::vector<T> tEqExample(descriptors::q<T,DESCRIPTORS>() );
5 for ( int iPop = 0; iPop < descriptors::q<T,DESCRIPTORS>(); ++iPop )
6     {tEqExample [ iPop ] = advectionDiffusionLbHelpers<T,TDESCRIPTOR>::
7         equilibrium( iPop, Texample, zeroVel ); }
8 ConstAnalyticalF2D<T,T> EqExample( tEqExample );
9 ADlattice.defineRho( superGeometry, 1 ,Example );
10 ADlattice.definePopulations( superGeometry, 1, EqExample );

```

Listing 4.15.: Initialization of the temperature field

To apply convective transport, a velocity vector has to be passed, which can be also done by material number.

```

1 std::vector<T> zero ( 2, T( ) );
2 ConstAnalyticalF2D<T,T> velocity ( zero );
3 ADlattice.defineField<descriptors::VELOCITY>(superGeoemtry, 1,velocity);

```

The last step is to make the lattice ready for the simulation:

```

1 NSlattice.initialize( );
2 ADlattice.initialize( );

```

Listing 4.16.: Initialization of the lattices

4.5.3.5. Setting the Boundary Conditions

If the value of a boundary condition has to be updated during the simulation, e.g. via increasing the velocity at the inflow or changing the temperature of a boundary, this can be achieved following the same procedure as for the initial conditions (see Section 5.2 further below).

4.5.3.6. Getting the Results

The desired data is saved using the `VtkWriter` objects, which can write the value of functors in VTI files (VTK format used e.g. by ParaView [70]). The functors which are usually saved are the velocity field from the `NSlattice`, and the temperature field (referred to as density) of the `ADlattice`. Thermal and isothermal information must be saved in two different objects, since they have two different descriptors.

```

1 SuperLatticeVelocity2D<T,NSDESCRIPTOR> velocity( NSlattice );
2 SuperLatticeDensity2D<T,TDESCRIPTOR> density( ADlattice );
3 vtkWriterNS.addFunctor( velocity );
4 vtkWriterAD.addFunctor( density );

```

```

5 vtkWriterNS.write(iT);
6 vtkWriterAD.write(iT);

```

Listing 4.17.: Saving results in VTK files

It is important to emphasize that the data saved is in *lattice units*. Some conversions can be made in order to obtain physical magnitudes, which are described later in Section 4.5.4.

4.5.3.7. Structure of the Program

It is advisable to structure the main loop of the .cpp-file along the following steps:

- 1. Initialization** The converter between dimensionless and lattice units is set via e.g. N , Δt and the parameters for the simulation Ra , Pr , T_{cold} , T_{hot} , $L_{x,y,z}$.
- 2. Prepare geometry** The mesh is created and voxels are classified with different material numbers according to their behavior (inflow, outflow, etc.).
- 3. Prepare lattice** The lattice dynamics are set according to the material numbers assigned before. The boundary conditions are initialized. Since there are two different lattices, the definition of the dynamics and the kind of boundary conditions (though not the actual values yet) have to be made for each of them separately. At this point the coupling generator is initialized (usually on the `NSlattice`) and then it is indicated which material numbers are to be coupled with the `ADlattice`.
- 4. Main loop with timer** The functions `setBoundaryValues`, `collideAndStream`, and `getResults` are called repeatedly until the maximum of iterations is reached or the simulation has converged (if a convergence criteria is set).
- 5. Definition of initial and boundary conditions** The values for the boundary functions are set. In some applications the values are to be refreshed at each time step. Thermal and isothermal lattices are treated separately. As indicated before, velocity and density (`NSlattice`), as well as temperature (`ADlattice`) have to be defined. Additionally, the couplings, and external forces and velocities should be initialized and reused as required.
- 6. Collide and stream execution** The collision and the streaming steps are performed. This function is called for each of the lattices separately. After the streaming step, the coupling between the lattices (here, based on the Boussinesq approximation) is executed.
- 7. Computation and output of the results** Console and data outputs of the results at certain time steps are created.

4.5.4. Obtaining Results in Thermal Simulations

Here, the Rayleigh and the Prandtl numbers are the dimensionless numbers which control the physics of a convection problem. The Rayleigh number for a fluid is associated with buoyancy driven flow. When the Rayleigh number is below the critical value for that fluid, heat transfer is primarily in the form of conduction. When it exceeds the critical value, heat transfer is primarily in the form of convection. For natural convection, it is defined as

$$Ra = \frac{g\beta}{\nu\alpha} \Delta T L^3, \quad (4.27)$$

where g is the acceleration magnitude due to gravity, β is the thermal expansion coefficient, ν is the kinematic viscosity, α is the thermal diffusivity, ΔT is the temperature difference, and L denotes the characteristic length. The Prandtl number is defined as the ratio of momentum diffusivity ν to thermal diffusivity α

$$Pr = \frac{\nu}{\alpha}. \quad (4.28)$$

To handle differences between the converter objects for isothermal and thermal simulations some of the isothermal functions have been re-implemented for thermal simulations in a way that they only depend on lattice parameters and the Rayleigh and Prandtl numbers by modifying existing functors in the following files:

- `/functors/lattice/blockLatticeIntegralF3D.(h, hh)`
- `/functors/lattice/superLatticeIntegralF3D.(h, hh)`

4.5.4.1. Velocity

The resulting velocity magnitude v_{res} , independent of the lattice velocity `latticeU` selected, can be computed by

$$v_{res} = \frac{v_{LB}}{latticeU} \sqrt{Ra Pr} = \frac{v_{LB}}{N \Delta t} \sqrt{Ra Pr}. \quad (4.29)$$

The lattice velocity `latticeU` is obtained from the function `converter.getCharLatticeVelocity()` in the thermal converter object.

4.5.4.2. Pressure

The pressure in physical units is derived from the lattice pressure by using its definition from the isothermal converter object

$$p_{\text{phys}} = p_{\text{LB}} \frac{\text{physcForce}}{\text{physL}^{d-1}} \quad (4.30)$$

$$= p_{\text{LB}} \frac{\text{physL}^{d+1}}{\text{physT}^2} \frac{1}{\text{physL}^{d-1}} \quad (4.31)$$

$$= p_{\text{LB}} \left(\frac{\text{physL}}{\text{physT}} \right)^2 \quad (4.32)$$

$$= p_{\text{LB}} \left(\frac{\text{charU}}{\text{latticeU}} \right)^2 \quad (4.33)$$

$$= p_{\text{LB}} \frac{Ra Pr}{\text{latticeU}^2} . \quad (4.34)$$

The lattice pressure can easily be computed from the lattice density using

$$p_{\text{LB}} = \frac{\rho - 1}{3} . \quad (4.35)$$

The physical force can also be obtained from the computed lattice force

$$F_{\text{phys}} = F_{\text{LB}} \frac{\text{physL}^{d+1}}{\text{physT}^2} \quad (4.36)$$

$$= F_{\text{LB}} \frac{\left(\frac{\text{charL} \text{latticeL}}{\text{charL}} \right)^{d+1}}{\left(\frac{\text{charL} \text{latticeU} \text{latticeL}}{\text{charU} \text{charL}} \right)^2} \quad (4.37)$$

$$= F_{\text{LB}} \text{latticeL}^{d-1} \left(\frac{\text{charU}}{\text{latticeU}} \right)^2 \quad (4.38)$$

$$= F_{\text{LB}} \frac{\text{latticeL}^{d-1}}{\text{latticeU}^2} Ra Pr , \quad (4.39)$$

where d is the number of dimensions in the problem.

For most applications the value of the force coefficients in the different coordinate directions can be of interest, which can be computed with

$$C_{F_i} = F_{i,\text{phys}} \frac{1}{\frac{1}{2} \text{charU}^2 \cdot \text{count}_i \cdot \text{latticeL}^{d-1}} = F_{i,\text{LB}} \frac{1}{\frac{1}{2} \text{latticeU}^2 \cdot \text{count}_i} , \quad (4.40)$$

where count_i is the number of cells in the surface perpendicular to the direction i of the force coefficient computed.

4.5.5. Conduction Problems

For heat conduction problems there is no velocity field that advects the temperature (see [109] and [98]). In absence of convection, radiation and heat generation, the energy equation for a

homogeneous medium is given by

$$\frac{\partial T}{\partial t} = \alpha \Delta T . \quad (4.41)$$

A conduction simulation can be executed using an independent advection–diffusion lattice, without any velocity field coupled. In the same way as in the convection–diffusion heat transfer, the temperature is obtained after summing the distribution functions over all directions. The equilibrium distribution function in this case with the BGK approximation is given by

$$g_i^{eq} = \omega_i \rho = \omega_i T , \quad (4.42)$$

which is equivalent to the one used in advection–diffusion simulations with the flow velocity set to zero. It means that conduction problems could be computed based on the available OpenLB code by only using a lattice with advection–diffusion dynamics and by setting the external velocity field to zero at any time.

4.5.5.1. Multiple-Relaxation-Time (MRT)

The implementation of the thermal lattice Boltzmann equation using the multiple-relaxation-time collision model is done similarly to the procedure used with the BGK collision model. A double MRT-LB is used, which consists of two sets of distribution functions: an isothermal MRT model for the mass-momentum equations, and a thermal MRT model for the temperature equation. Both sets are coupled by a force term according to the Boussinesq approximation. The macroscopic governing equation for the temperature is

$$\frac{\partial T}{\partial t} + v \nabla T = \alpha \Delta T , \quad (4.43)$$

where α is the thermal diffusivity coefficient.

The isothermal MRT model with an external force is already implemented in OpenLB for the `D2Q9` and `D3Q19` lattice models (dynamics class `ForcedMRTdynamics`). This means that only the thermal MRT counterparts for 2D and 3D have to be developed.

The computation of the force term in the MRT model in the `ForcedMRTdynamics` class uses the body force as described in [101]. It does not include however a velocity shift like the BGK model, due to negligible differences in benchmark tests.

D2Q5 thermal model The formulation for the `D2Q5` thermal MRT model is based on [107]. The temperature field distribution functions g_i are governed by the following equation

$$g_i(\mathbf{x} + \mathbf{c}_i \Delta t, t + \Delta t) - g_i(\mathbf{x}, t) = -N_i^{-1} \theta_i [\mathbf{n}(\mathbf{x}, t) - \mathbf{n}^{eq}(\mathbf{x}, t)] , \quad (4.44)$$

where \mathbf{g} and \mathbf{n} are column vectors with entries g_i and n_i for $i = 0, 1, \dots, q - 1$, and denote the distribution functions and the moments, respectively. The vectors N_i are the rows of the

orthogonal transformation matrix \mathbf{N} and θ_i are the entries of a non-negative, diagonal relaxation matrix. The macroscopic temperature T can be calculated by

$$T = \sum_{i=0}^4 g_i . \quad (4.45)$$

The weight coefficients for each lattice direction are given in equation (4.46) with

$$\omega_i = \begin{cases} \frac{3}{5}, & \text{if } i = 0 , \\ \frac{1}{10}, & \text{if } i = 1, 2, 3, 4 . \end{cases} \quad (4.46)$$

The transformation matrix \mathbf{N} maps the distribution functions for the temperature g_i to the corresponding moments n_i , i.e.

$$\mathbf{n} = \mathbf{N}\mathbf{g}. \quad (4.47)$$

The transformation matrix \mathbf{N} and its inverse matrix \mathbf{N}^{-1} are shown in equations (4.48) and (4.49). There are some differences in the order of the columns with respect to what is specified in the reference [107]. This is due to the different sequence used in numbering the velocity directions.

$$\mathbf{N} = \begin{pmatrix} \langle 1 | \\ \langle e_x | \\ \langle e_y | \\ \langle 5e^2 - 4 | \\ \langle e_x^2 - e_y^2 | \end{pmatrix} = \begin{bmatrix} 1 & 1 & 1 & 1 & 1 \\ 0 & -1 & 0 & 1 & 0 \\ 0 & 0 & -1 & 0 & 1 \\ -4 & 1 & 1 & 1 & 1 \\ 0 & 1 & -1 & 1 & -1 \end{bmatrix} , \quad (4.48)$$

$$\mathbf{N}^{-1} = \begin{pmatrix} \frac{1}{5} & 0 & 0 & -\frac{1}{5} & 0 \\ \frac{1}{5} & -\frac{1}{2} & 0 & \frac{1}{20} & \frac{1}{4} \\ \frac{1}{5} & 0 & -\frac{1}{2} & \frac{1}{20} & -\frac{1}{4} \\ \frac{1}{5} & \frac{1}{2} & 0 & \frac{1}{20} & -\frac{1}{4} \\ \frac{1}{5} & 0 & \frac{1}{2} & \frac{1}{20} & -\frac{1}{4} \end{pmatrix} . \quad (4.49)$$

The equilibrium moments \mathbf{n}^{eq} are defined as

$$\mathbf{n}^{eq} = (T, u_x T, u_y T, \varpi T, 0)^T , \quad (4.50)$$

where ϖ is a constant of the $D2Q5$ model, which we set to -2 . The diagonal relaxation matrix θ is defined by

$$\theta = \text{diag} (0, \zeta_a, \zeta_a, \zeta_e, \zeta_\nu) . \quad (4.51)$$

The first relaxation rate, corresponding to the temperature, is set to zero for simplicity, since the first moment is conserved. The relaxation rates ζ_e and ζ_ν are set to 1.5, whereas the relaxation rates ζ_a are functions of the thermal diffusivity α in (4.53). The speed of sound of the $D2Q5$

model is $c_s^2 = 0.2$.

$$\alpha = c_s^2 \left(\tau_a - \frac{1}{2} \right) = \frac{1}{5} \left(\tau_a - \frac{1}{2} \right) \quad (4.52)$$

$$\Rightarrow \zeta_a = \frac{1}{\tau_a} = \frac{1}{5\alpha + \frac{1}{2}}. \quad (4.53)$$

4.5.5.2. Particle Flows as advection–diffusion Problem

The quantity c in the advection–diffusion equation can be considered as a particle density, thereby giving a continuous ansatz for simulating particle flows. To solve for the particle distribution, an additional lattice is required with an appropriate descriptor and dynamics, which are only implemented for the 3D case.

```
1 using ADDESCRIPTOR = descriptors::ParticleAdvectionDiffusionD3Q7Descriptor;
```

Listing 4.18.: Advection diffusion descriptor for particle flows

The descriptor in Listing 4.18 allocates additional memory since for the computation of the particle velocity the velocity of the last time step has to be stored as well. These calculations are also non-local, therefore the communication of the additional data has to be ensured by an additional object, which is constructed according to line 1 of Listing 4.19 and communicates the data by a function as shown in line 2, which has to be called in the time loop.

```
1 SuperExternal3D<T, ADDESCRIPTOR, descriptors::VELOCITY> sExternal( superGeometry,
    sLatticeADE, sLatticeAD.getOverlap() );
2 sExternal.communicate();
```

Listing 4.19.: SuperExternal3D object for the communication of additional data

Although the same unit converter can be used for the advection–diffusion lattice, another relaxation parameter has to be handed to the dynamics, as shown in Listing 4.20. Additionally, some of the boundary conditions have to take the diffusion coefficient into account. Therefore a new ω_{ADE} is computed by

$$\omega_{ADE} = \left(4D \frac{U_{LB}}{L_{LB}U_C} + 0.5 \right)^{-1}. \quad (4.54)$$

with characteristic lattice velocity U_{LB} , characteristic velocity U_C , lattice length L_{LB} , as well as the desired diffusion coefficient D .

```
1 ParticleAdvectionDiffusionBGKdynamics<T, ADDESCRIPTOR> bulkDynamicsAD ( omegaAD,
    instances::getBulkMomenta<T, ADDESCRIPTOR>() );
```

Listing 4.20.: Dynamics for the simulation of particle flows via advection–diffusion equations

Applying the advection–diffusion equation to particle flow problems requires a new dynamics due to the handling of the particle velocity by the coupling processor of the two lattices, which differs for reasons of efficiency. When constructing the coupling post-processor as shown in Listing 4.21, forces acting on the particle can be added like the Stokes drag force as shown in line 2 and 3 of Listing 4.21. The implementation of new forces is straight forward, since only a new class which provides a function `applyForce(...)`, computing the force in a cell, needs to be written analogously to the existing `advDiffDragForce3D`. Finally the lattices are linked by line 4 of Listing 4.21, which needs to be applied to the Navier–Stokes lattice for reasons of accessibility.

```

1 AdvectionDiffusionParticleCouplingGenerator3D<T, NSDESCRIPTOR> coupling( ADDESCRIPTOR::
    index<descriptors::VELOCITY>());
2 advDiffDragForce3D<T, NSDESCRIPTOR> dragForce( converter, radius, partRho );
3 coupling.addForce( dragForce );
4 sLatticeNS.addLatticeCoupling( superGeometry, 1, coupling, sLatticeAD );

```

Listing 4.21.: Coupling of an advection–diffusion and a Navier–Stokes lattice for particle flow simulations

For the boundary conditions the same basic objects as for the advection–diffusion equation can be used, however there is an additional boundary condition shown on Listing 4.22 which has to be applied at all boundaries to ensure correctness of the finite differences scheme used to compute the particle velocity. Further information as well as results can be found in Trunk *et al.* [45] as well as in the examples section.

```

1 setExtFieldBoundary<T, ADDESCRIPTOR, descriptors::VELOCITY, descriptors::VELOCITY2>(
2     sLatticeAD, superGeometry.getMaterialIndicator({2, 3, 4, 5, 6}));

```

Listing 4.22.: Example of a boundary condition for the particle velocity for particle flow simulations

4.6. Particles

The present section summarizes OpenLB’s functionality regarding the consideration of discrete particles in a Lagrangian framework. This includes both *sub-grid* particles assuming spherical shapes and *surface resolved* particles with arbitrary shapes, which can be handled by a common particle framework. As the framework follows advances in the data concept of the lattice (cf. Section 2), it provides a dimension agnostic, flexible and easily extendable implementation. While abstract template meta functionality characterizes the data handling level, accessible high-level user-functions are provided for e.g. particle creation or coupling handling. In order to guarantee support for previously developed applications, the 3D-only sub-grid particle framework from the previous releases is included as *sub-grid (legacy) framework* (cf. Section 4.6.9) as well.

To get a good overview of the particle framework, the code of examples `settlingCube3d` (Section 8.7.4) and `bifurcation3d` (Section 8.7.1) is reviewed, focusing on the simulation of particles. The example `settlingCube3d` examines the settling of a cubical silica particle under the influence of gravity in surrounding water. It starts with integrating some libraries and namespaces, followed by the definition of different types (Listing 4.23), e.g. the *descriptor* and the *particle type*. Afterwards, some variables are set to a concrete value, used in the fluid and particle calculation (Listing 4.24). Particle settings include all the data to solve the equations of motion, such as the particle's starting position and density.

```

1  #include "olb3D.h"
2  #include "olb3D.hh"
3
4  using namespace olb;
5  using namespace olb::descriptors;
6  using namespace olb::graphics;
7  using namespace olb::util;
8  using namespace olb::particles;
9  using namespace olb::particles::dynamics;
10
11 using T = FLOATING_POINT_TYPE;
12 typedef D3Q19<POROSITY, VELOCITY_NUMERATOR, VELOCITY_DENOMINATOR> DESCRIPTOR;
13
14 //Define lattice type
15 typedef PorousParticleD3Q19Descriptor DESCRIPTOR;
16
17 //Define particleType
18 typedef ResolvedParticle3D PARTICLETYPE;

```

Listing 4.23.: Following the new particle system with following the example `settlingCube3d`

```

1  //Particle Settings
2  T centerX = lengthX*.5;
3  T centerY = lengthY*.5;
4  T centerZ = lengthZ*.9;
5  T const cubeDensity = 2500;
6  T const cubeEdgeLength = 0.0025;
7  Vector<T,3> cubeCenter = {centerX, centerY, centerZ};
8  Vector<T,3> cubeOrientation = {0., 15., 0.};
9  Vector<T,3> cubeVelocity = {0., 0., 0.};
10 Vector<T,3> externalAcceleration = {0., 0., -T(9.81) * (T(1) - physDensity /
    cubeDensity)};
11
12
13 // Characteristic Quantities
14 T const charPhysLength = lengthX;
15 T const charPhysVelocity = 0.15;    // Assumed maximal velocity

```

Listing 4.24.: Particle settings in example `settlingCube3d`

Like other simulations, particle flow simulations need basic, non particle-specific functions like `prepareGeometry` or `prepareLattice`. After initializing those functions, the main function

starts. The main section begins with initialization of physical units in the unit converter, which is explained in the Q&A in Section A.1. The unit converter is followed by the preparation of the geometry using the `prepareGeometry`-function and afterwards the `prepareLattice`-function. After those general simulation functions, the particle simulation starts. First, the `ParticleSystem` (Listing 4.25, explained in Section 4.6.1) is called followed by the calculation of the particle quantities like a smoothing factor and the extent of the particles. After those calculations, the particles are created. In the following lines, dynamics are assigned to the particles.

```

1  // Create ParticleSystem
2  ParticleSystem<T,PARTICLETYPE> particleSystem;
3
4  //Create particle manager handling coupling, gravity and particle dynamics
5  ParticleManager<T,DESCRIPTOR,PARTICLETYPE> particleManager(
6      particleSystem, superGeometry, sLattice, converter, externalAcceleration);
7
8  // Create and assign resolved particle dynamics
9  particleSystem.defineDynamics<
10      VerletParticleDynamics<T,PARTICLETYPE>>();
11
12  // Calculate particle quantities
13  T epsilon = 0.5*converter.getConversionFactorLength();
14  Vector<T,3> cubeExtend( cubeEdgeLength );
15
16  // Create Particle 1
17  creators::addResolvedCuboid3D( particleSystem, cubeCenter,
18                                cubeExtend, epsilon, cubeDensity, cubeOrientation );
19
20  // Create Particle 2
21  cubeCenter = {centerX,lengthY*T(0.51),lengthZ*T(.7)};
22  cubeOrientation = {0.,0.,15.};
23  creators::addResolvedCuboid3D( particleSystem, cubeCenter,
24                                cubeExtend, epsilon, cubeDensity, cubeOrientation );
25
26  // Check ParticleSystem
27  particleSystem.checkForErrors();

```

Listing 4.25.: Creation of particles and assigning dynamics

Before the main loop starts, Listing 4.26, we create a timer, Listing 4.26 and set initial values to the distribution functions by calling `setBoundaryValues`. After this, the following is processed at every time step. The fluid's influence on the particles is calculated by evaluating hydrodynamic forces acting on the particle surface. Afterwards, an external acceleration, e.g. gravity, is applied onto the particles (Listing 4.26) and the equations of motion are solved for each one. The back coupling from the particles to the fluid follows afterwards. Finally, the main loop ends with the `getResults`-function, which prints the results to the console and writes VTK data for post-processing with ParaView (Section 6.9) at previously defined time intervals.

```

1  /// === 4th Step: Main Loop with Timer ===
2  Timer<T> timer(converter.getLatticeTime(maxPhysT), superGeometry.getStatistics().
3      getNvoxel());
4  timer.start();

```

```

4
5
6  /// === 5th Step: Definition of Initial and Boundary Conditions ===
7  setBoundaryValues(sLattice, converter, 0, superGeometry);
8
9  clout << "MaxIT: " << converter.getLatticeTime(maxPhysT) << std::endl;
10
11  for (std::size_t iT = 0; iT < converter.getLatticeTime(maxPhysT)+10; ++iT) {
12
13      // Execute particle manager
14      particleManager.execute<
15          couple_lattice_to_particles<T, DESCRIPTOR, PARTICLETYPE>,
16          apply_gravity<T, PARTICLETYPE>,
17          process_dynamics<T, PARTICLETYPE>,
18          couple_particles_to_lattice<T, DESCRIPTOR, PARTICLETYPE>
19      >();
20
21      // Get Results
22      getResult(sLattice, converter, iT, superGeometry, timer, particleSystem);
23
24      // Collide and stream
25      sLattice.collideAndStream();
26  }
27
28  timer.stop();
29  timer.printSummary();

```

Listing 4.26.: Main Loop with Timer

As we followed the example for particle simulation `settlingCube3d`, some functions necessary for the simulations were introduced. Therefore, in the next chapters, the individual parts of the framework are examined.

4.6.1. Class ParticleSystem

The `ParticleSystem` stores all data concerning the particles in containers. Therefore, the class is used multiple times in a particle simulation. First, the `ParticleSystem` is created according to the desired `PARTICLETYPE` (Listing 4.25). However, the container of particles is empty. Therefore, we add two particles to it using creator functions and add dynamics via the `ParticleSystem`. Additionally, it is utilized in the `ParticleManager` (cf. Section 4.6.3) to access the particles and perform predefined operations on them.

One focus of the new particle system is the separation of data and operations according to the lattice framework (cf. Section 2). Therefore, only the data is stored in the `ParticleSystem`. For the operations, it is non-relevant and only used to store data of the calculations.

4.6.2. Class SuperParticleSystem

The example `bifurcation3d` (Section 8.7.1) makes use of OpenLB's domain decomposition approach. In order to use this, a `SuperParticleSystem` has to be created by passing the

`SuperGeometry`, which holds all information regarding the lattice decomposition:

```
1 SuperParticleSystem<T,PARTICLETYPE> superParticleSystem(superGeometry);
```

Listing 4.27.: Initialization of a SuperParticleSystem

4.6.3. Class ParticleManager

The `ParticleManager` can be used to encapsulate relevant reoccurring particle tasks as e.g. the particle-lattice-coupling. After its initial instantiation by providing the access to relevant particle, lattice and set-up specific data, its `execute()` method can be called with the respective tasks specified as template arguments in the desired order. The individual tasks (included in `particleTasks.h`) provide an `execute()` method as well and a parameter set specifying the coupling type and the potential embedding into a loop over all available particles. The `ParticleManager` also takes care of combining respective tasks into a single particle loop.

When using a domain decomposition, the particle core distribution has to be updated in every time step. The use of the `ParticleManager` in the `bifurcation3d` example (Section 8.7.1) respectively looks as follows:

```
1 ParticleManager<T,DESCRIPTOR,PARTICLETYPE> particleManager(  
2     superParticleSystem, superGeometry, superLattice, converter);
```

Listing 4.28.: Initialization of a ParticleManager

```
1 particleManager.execute<  
2     couple_lattice_to_particles<T,DESCRIPTOR,PARTICLETYPE>,  
3     process_dynamics<T,PARTICLETYPE>,  
4     update_particle_core_distribution<T,PARTICLETYPE>  
5 >();
```

Listing 4.29.: Execution of the ParticleManager

4.6.4. Resolved Lattice Interaction

In the directory `resolved`, all surface resolved specific functionality is bundled. The `blockLatticeInteraction.h` (only header file) and `blockLatticeInteraction.hh` files consist of five functions. All of those functions are needed to calculate and check the position of the particles inside the geometry. For example, the `checkSmoothIndicatorOutOfGeometry`-function checks if every part of the particle is inside the cell barrier. If the particle is partly outside of the geometry, the position needs to be changed. Another important functions is the `setBlockParticleField`, where a all cells, which are inside the particle are set as a particle field. Similar to the `blockLatticeInteraction.hh` also the `superLatticeInteraction.hh` exist. In this file the `setBlockParticleField` gets converted to the lattice structure with the

function `setSuperParticleField`. The file `momentumExchangeForce.h` provides functions to calculate hydrodynamic forces on the particle's surface via an adapted momentum exchange algorithm. The file (`smoothIndicatorInteraction.h`) is needed for the simulation of the area directly at the surface of the particles.

4.6.5. Particle Descriptors

The first file is the `particleDescriptorAlias.h` file, in which the alias' are given to different types of `PARTICLETYPES`. In the example `settlingCube3d`, right in the beginning, the `PARTICLETYPE ResolvedParticle3D` is chosen. After the choice of this alias, the dynamics and other main properties are set. Other possible particle types that can be chosen are `ResolvedParticle2D` OR `ResolvedSphere3D`.

For the `bifurcation3d` example (Section 8.7.1), a respective descriptor is chosen:

```
1 typedef D3Q19<> DESCRIPTOR;
2 typedef SubgridParticle3D PARTICLETYPE;
```

Listing 4.30.: Particle type and descriptor

4.6.6. Particle Dynamics

Another important part of the particle system are the dynamics. The files are used to define those properties for the chosen particle type. For example, in the `particleDynamics.h` and `particleDynamics.hh` all dynamic functions for the particle type are implemented. Therefore, all information for calculation of dynamic values can be found here e.g. acceleration or angular acceleration. Those functions get called in the main part of simulations, e.g. in the example `settlingCube3d` (`VerletParticleDynamics`), in Listing 4.25, as the dynamics are assigned to the particle type.

In the example `bifurcation3d` (Section 8.7.1), different capture methods can be used by choosing the respective setting beforehand:

```
1 //Set capture method:
2 // materialCapture: based on material number
3 // wallCapture:      based on more accurate stl description
4 typedef enum {materialCapture, wallCapture} ParticleDynamicsSetup;
5 const ParticleDynamicsSetup particleDynamicsSetup = wallCapture;
```

Listing 4.31.: Enum `ParticleDynamicsSetup` and example initialization

If the `wallCapture` is chosen, a `SolidBoundary` object has to be created and passed to the respective dynamics:

```
1 STLReader<T> stlReader( "../bifurcation3d.stl", converter.getConversionFactorLength() )
  ;
```

```

2   IndicatorLayer3D<T> extendedDomain( stlReader,          converter.
      getConversionFactorLength() );
3   // Create solid wall
4   const unsigned latticeMaterial = 2; //Material number of wall
5   const unsigned contactMaterial = 0; //Material identifier (only relevant for contact
      model)
6   SolidBoundary<T,3> wall( std::make_unique<IndicInverse<T,3>>(stlReader),
7                          latticeMaterial, contactMaterial );

```

Listing 4.32.: Solid wall creation

When using the `materialCapture` instead, a `MaterialIndicator` is necessary to identify material numbers that initialize the capture treatment.

```

1   std::vector<int> materials {2,4,5};
2   SuperIndicatorMaterial<T,3> materialIndicator (superGeometry, materials);

```

Listing 4.33.: Initialization of a SuperIndicatorMaterial

Both `SolidBoundary` and `MaterialIndicator` are then used in the function `prepareParticles` to define the chosen dynamics:

```

1   if (particleDynamicsSetup==wallCapture){
2       //Create verlet dynamics with material aware wall capture
3       superParticleSystem.defineDynamics<
4           VerletParticleDynamicsMaterialAwareWallCapture<T,PARTICLETYPE>>(
5           wall, materialIndicator);
6   } else {
7       //Create verlet dynamics with material capture
8       superParticleSystem.defineDynamics<
9           VerletParticleDynamicsMaterialCapture<T,PARTICLETYPE>>(materialIndicator);
10  }

```

Listing 4.34.: Creation of verlet dynamics

4.6.7. Particle Functions

In the functions-directory, additional free functions are defined. These functions are callable anywhere in the code. The first set of files including `particleCreatorFunctions.h`, `particleCreatorFunctions2D.h`, `particleCreatorFunctions3D.h` and `particleCreatorHelperFunctions.h` concentrate on functions concerning the creation of particles with different types of surface structures. These functions are therefore called first to create particles in the desired shape. In the example `settlingCube3d`, the function `addResolvedCuboid` is called and creates a particle in the shape of a cuboid. Also other geometries, like circles in 2D or cylinders in 3D, can be created. All of those functions are implemented in these files.

The file (`particleMotionfunctions.h`) concentrates on the main algorithms for solving the equations of motion. Two functions exist, using different integration-types, velocity Verlet algorithm (`velocityVerletIntegration`) or Euler-Integration (`eulerIntegration`). The former function is used in the `VerletParticleDynamics` class (Section 4.6.6) and are therefore called in the main part of the example. Often two functions for the same calculation exist as they need to match the dimension of a problem. Those are differentiated by partial template specialization.

The `particleDynamicsFunctions.h` also contains other important functions to simulate particle flows. Tasks are included in the `particleTasks.h` like e.g. the `couple_lattice_to_particles` or `couple_resolved_particles_to_lattice`. Both functions are used in the main loop of the example, to realize a two-way-coupling.

To sum up, many of the most important functions for the simulations of the particle flow are implemented in the `ParticleDynamicsFunctions.h`. Other functions, e.g. concerning the calculation of rotation of the particle body, are implemented in the `bodyMotionfunctions.h`. The `particleIoFunctions.h` contains functions to get the output of the calculation to the console. It consist of two important functions (`printResolvedParticleInfo` and `printResolvedParticleInfoSimple`), which are used in the `getResults`-function. The `getResults`-function is called at the end of the main part of every simulation.

4.6.8. Discrete Contact Model for Surface Resolved Particles

In order to simulate particulate flows accurately, it is often necessary to incorporate a contact model. The here used discrete contact model [29] allows for the treatment of particle-particle and particle-wall interactions, enabling the calculation of contact forces and their application to the particles. The discrete contact model consists of several steps that are integrated into the general algorithm. Let's discuss each step in detail:

- *Rough contact detection during coupling:* During the coupling stage, where particle information is transferred to the fluid lattice, a rough contact detection mechanism is employed. This step identifies potential contact regions between particles and the fluid. It determines potential contacts by identifying particles that couple to the same cell.
- *Communication of found contacts:* Once the potential contact regions are identified, the information regarding the found contact is communicated across all processes.
- *Correction of contact bounding box:* To improve the accuracy of the contact treatment, the contact bounding boxes are refined based on the information obtained during the communication step. This correction step helps in precisely defining the contact regions, ensuring that the subsequent calculations consider the actual contacts.
- *Determination of contact properties:* With the refined contact bounding boxes, the discrete contact model determines various contact properties. These properties include the contact volume, contact point, contact normal and other relevant parameters.

- *Calculation of contact force and application to particles:* Using the contact properties, the contact force is calculated from the parameters determined before and applied to the particles so that it's available when solving the equations of motion.
- *Removal of empty contact objects (optional):* After the contact forces have been determined and applied, empty contact objects, which no longer represent an existing contact, may be removed. This step helps in optimizing the computational efficiency by eliminating unnecessary iterations.

The usage within OpenLB is exemplified by *dkt2d*. First, we set types for the particle-particle and particle-wall interactions, which define how the contact is treated. This is represented in Listing 4.35).

```
1 typedef ParticleContactArbitraryFromOverlapVolume<T, DESCRIPTOR::d, true>
   PARTICLECONTACTTYPE;
2 typedef WallContactArbitraryFromOverlapVolume<T, DESCRIPTOR::d, true> WALLCONTACTTYPE;
```

Listing 4.35.: Contact types

Additionally, we define the walls for the contact treatment, as shown in Listing 4.36). Here, we create a `SolidBoundary` from an indicator and specify the minimal and maximal coordinates as well as the material number that represents the wall on the lattice and the material identifier, which defines the walls mechanical properties.

```
1 std::vector<SolidBoundary<T, DESCRIPTOR::d>> solidBoundaries;
2 solidBoundaries.push_back( SolidBoundary<T, DESCRIPTOR::d>(
3     std::make_unique<IndicInverse<T, DESCRIPTOR::d>>(
4         cuboid,
5         cuboid.getMin() - 5 * converter.getPhysDeltaX(),
6         cuboid.getMax() + 5 * converter.getPhysDeltaX(),
7         wallLatticeMaterialNumber,
8         wallContactMaterial));
```

Listing 4.36.: Solid boundaries

Similarly, we set a number that relates the particles to mechanical properties, see Listing 4.37.

```
1 for (std::size_t iP = 0; iP < particleSystem.size(); ++iP) {
2     auto particle = particleSystem.get(iP);
3     setContactMaterial(particle, particleContactMaterial);
4 }
```

Listing 4.37.: Particle material number

To store contact objects later on, we create an empty `ContactContainer` as shown in Listing 4.38).

```
1 ContactContainer<T, PARTICLECONTACTTYPE, WALLCONTACTTYPE> contactContainer;
```

Listing 4.38.: Contact container

By creating a lookup table `ContactProperties` that contains constant parameters which solely depend on the material combination in Listing 4.39), we save computational effort. Setting properties must be done for each material combination separately.

```

1 ContactProperties<T, 1> contactProperties;
2 contactProperties.set(particleContactMaterial,
3                     wallContactMaterial,
4                     evalEffectiveYoungModulus(
5                         youngsModulusParticle,
6                         youngsModulusWall,
7                         poissonRatioParticle,
8                         poissonRatioWall),
9                     coefficientOfRestitution,
10                    coefficientKineticFriction,
11                    coefficientStaticFriction);

```

Listing 4.39.: Contact properties

Finally, we process the contacts as shown in Listing 4.40).

```

1 processContacts<T, PARTICLETYPE, PARTICLECONTACTTYPE, WALLCONTACTTYPE,
2   ContactProperties<T, 1>>(
3   particleSystem, solidBoundaries,
4   contactContainer, contactProperties,
5   superGeometry, contactBoxResolutionPerDirection);

```

Listing 4.40.: Contact processing

4.6.9. Sub-grid Legacy Framework

In this Section the use of Lagrangian particles with the legacy framework is shown. Due to similar naming of classes and functions in the new common framework, it is worth noting that all terms are primarily referring to the naming convention used in the legacy framework itself and should not be mixed up with those of the new one.

Similar to the `BlockLattice` and `SuperLattice` structure, a `ParticleSystem3D` and `SuperParticleSystem3D` structure exists. In line 2 of Listing 4.41 the `SuperParticleSystem3D` is instantiated, taking a `SuperGeometry` as a parameter. In line 4 the `SuperParticleSysVtuWriter` is instantiated. It takes the `SuperParticleSystem3D`, a filename as string, and the wanted particle properties as arguments. Calling the function `SuperParticleSysVtuWriter.write(int timestep)` creates `.vtu` files of the particles positions for the given timestep. These files can be visualized with ParaView.

Line 10 of the listing instantiates an interpolation functor for the fluids velocity, which is used in line 13 during the instantiation of `StokesDragForce3D`. Particles need boundary conditions as well. In the listing, the simplest possible material boundary is presented. If a particle moves into a lattice node with material number 2, 4 or 5 its velocity is set to 0 and it is neglected during further computations, its state of activity is set to false. This `MaterialBoundary3D` is instantiated

in line 16. In lines 18 and 19 the force and boundary condition are added to and stored in the respective lists in the `SuperParticleSystem3D`.

The actual number crunching is then performed in line 25 which is positioned in the main loop of the program. The `supParticleSystem.simulate(T timeStep);` function integrates the particle trajectories by `timeStep`. Therefore all stored particle forces are computed and summed up. The particles are moved one step according to Newton's laws. Then all stored particle boundary conditions are applied. Parallelization of the particles is achieved automatically.

Results of this simulation are published in Henn et al. [14].

```

1  // SuperParticleSystems3D
2  SuperParticleSystem3D<T,PARTICLE> supParticleSystem(superGeometry);
3  // define which properties are to be written in output data
4  SuperParticleSysVtuWriter<T,PARTICLE> supParticleWriter(supParticleSystem, "particles
   ",
5      SuperParticleSysVtuWriter<T,PARTICLE>::particleProperties::velocity |
6      SuperParticleSysVtuWriter<T,PARTICLE>::particleProperties::mass |
7      SuperParticleSysVtuWriter<T,PARTICLE>::particleProperties::radius |
8      SuperParticleSysVtuWriter<T,PARTICLE>::particleProperties::active);
9
10 SuperLatticeInterpPhysVelocity3D<T,DESCRIPTOR> getVel(sLattice, converter);
11
12 auto stokesDragForce = make_shared<StokesDragForce3D<T,PARTICLE,DESCRIPTOR>> (getVel,
   converter);
13
14 // material numbers where particles should be reflected
15 std::set<int> boundMaterial = { 2, 4, 5};
16 auto materialBoundary = make_shared<MaterialBoundary3D<T, PARTICLE>> (superGeometry,
   boundMaterial);
17
18 supParticleSystem.addForce(stokesDragForce);
19 supParticleSystem.addBoundary(materialBoundary);
20 supParticleSystem.setOverlap(2. * converter.getPhysDeltaX());
21
22 \* ... *\
23
24 main loop {
25     supParticleSystem.simulate(converter.getPhysDeltaT());
26 }

```

Listing 4.41.: Usage of class `SuperParticleSystem3D`

4.6.9.1. Interpolation of Fluid Velocity

As the particle position $\mathbf{X} : I \rightarrow \Omega$ moves in the continuous domain Ω and information on the fluid velocity can only be computed on lattice nodes $\mathbf{x}_i \in \Omega_h$, interpolation of the fluid velocity is necessary every time fluid-particle forces are computed. Let $\mathbf{u}_i^F = \mathbf{u}^F(\mathbf{x}_i)$ be the computed solution of the Navier–Stokes Equation at lattice nodes \mathbf{x}_i . Let $p \in P_n$ be the interpolating polynomial of order n with $p(\mathbf{x}_i) = \mathbf{u}_i^F$ and $(\mathbf{x}_0, \dots, \mathbf{x}_n)$ the smallest interval containing all points in the brackets. Furthermore, let $C^n[a, b]$ be the vector space of continuous functions

that have continuous first n derivatives in $[a, b]$. Then the interpolation error of the polynomial interpolation is stated by the following theorem.

Theorem 1 (Interpolation error). *Let $\mathbf{u} \in C^{n+1}[a, b]$, $\mathbf{a}, \mathbf{b} \in \Omega$. Then for every $\mathbf{x} \in [a, b]$ there exists one $\hat{\mathbf{x}} \in (\mathbf{x}_0, \dots, \mathbf{x}_n, \mathbf{x})$, such that*

$$\mathbf{u}^F(\mathbf{x}) - p_n(\mathbf{x}) = \frac{d_{\mathbf{x}}^{n+1} \mathbf{u}^F(\hat{\mathbf{x}})}{(n+1)!} \prod_{j=0}^n (\mathbf{x} - \mathbf{x}_j) \quad (4.55)$$

holds.

Proof. See Rannacher [115, Satz 2.3]. □

Using linear ($n = 1$) interpolation for the fluid velocity between two neighboring lattice nodes $\mathbf{a} = \mathbf{x}_0 \in \Omega_h$, $\mathbf{b} = \mathbf{x}_1 \in \Omega_h$, $\|\mathbf{x}_1 - \mathbf{x}_0\|_2 = h$ clearly the following holds

$$f(\mathbf{x}) - p_1(\mathbf{x}) = \frac{1}{2} d_{\mathbf{x}}^2 \mathbf{u}^F(\hat{\mathbf{x}}) (\mathbf{x} - \mathbf{x}_0)(\mathbf{x} - \mathbf{x}_1) \quad (4.56)$$

$$\leq \frac{1}{2} d_{\mathbf{x}}^2 \mathbf{u}^F(\hat{\mathbf{x}}) h^2 \quad (4.57)$$

and the approximation error of the linear interpolation is of order $\mathcal{O}(h^2)$. In the following we give reason why this order of interpolation is sufficient.

Lets assume there exists an ideal error law of the form

$$\|\mathbf{u}_i^F - \mathbf{u}_i^{F*}\|_{L^2(\Omega_h)} = ch^\alpha,$$

for the discrete solution \mathbf{u}_h^F obtained by an LBM with lattice spacing h and the analytic solution \mathbf{u}^{F*} . Then $\alpha \in \mathbb{R}^+$ is the to be determined order of convergence. We further define the relative error

$$\text{Err}_h = \frac{\|\mathbf{u}_h^F - \mathbf{u}^{F*}\|_{L^2(\Omega_h)}}{\|\mathbf{u}^{F*}\|_{L^2(\Omega_h)}}.$$

The ratio of the error laws of two distinct lattice spacings h_i and h_j forms the EOC as

$$\text{EOC}_{i,j} = \frac{\ln(\text{Err}_{h_i}/\text{Err}_{h_j})}{\ln(h_i/h_j)}. \quad (4.58)$$

With this Krause [0, Chapter 2.3] determines an of EOC ≈ 2 for the discrete solution towards the analytic solution of a stationary flow in the unit cube governed by the incompressible NSE. Therefore the order of converge of the fluid velocity obtained by an LBM can be assumed to be $\mathcal{O}(h^2)$. This conclusion is backed up by the theoretical results obtained by [84]. This leads to the assumption that, each interpolation scheme of higher order than 2 would not be exhausted as the error of the incoming data is too large.

The interpolation is implemented as a trilinear interpolation using the eight nodes surrounding the particle. Let the point of interpolation $\hat{\mathbf{x}} \in [x_{(0,0,0)}, x_{(1,1,1)}]$ be in the cube spanned by




Figure 4.1.: Trilinear interpolation.

the lattice nodes $\mathbf{x}_{(0,0,0)}$ and $\mathbf{x}_{(1,1,1)}$, see Figure 4.1 for an illustration. We will denote by

$$\mathbf{d} = (d_0, d_1, d_2)^T = \hat{\mathbf{x}} - \mathbf{x}_{(0,0,0)} , \quad (4.59)$$

the distance of the particle to the next smaller lattice node. The fluid velocities at the eight corners are named accordingly $\mathbf{u}_{(i,j,k)}$, $i, j, k \in \{0, 1\}$. The trilinear interpolation is executed by three consecutive linear interpolations in the three different space directions. First we interpolate along the x -axis

$$u_{(d,0,0)} = u_{(0,0,0)}(h - d_0) + u_{(1,0,0)}d_0 , \quad (4.60)$$

$$u_{(d,1,0)} = u_{(0,1,0)}(h - d_0) + u_{(1,1,0)}d_0 , \quad (4.61)$$

$$u_{(d,0,1)} = u_{(0,0,1)}(h - d_0) + u_{(1,0,1)}d_0 , \quad (4.62)$$

$$u_{(d,1,1)} = u_{(0,1,1)}(h - d_0) + u_{(1,1,1)}d_0 , \quad (4.63)$$

followed by interpolation along the y -axis

$$u_{(d,d,0)} = u_{(d,0,0)}(h - d_1) + u_{(d,1,0)}d_1 , \quad (4.64)$$

$$u_{(d,d,1)} = u_{(d,0,1)}(h - d_1) + u_{(d,1,1)}d_1 , \quad (4.65)$$

and finally in direction of the z -axis

$$\mathbf{u}(\hat{\mathbf{x}}) = u_{(d,d,d)} = u_{(d,d,0)}(h - d_2) + u_{(d,d,1)}d_2 . \quad (4.66)$$

4.6.9.2. Class SuperParticleSystem3D

The implementation of the particle phase follows an hierarchical ansatz, similar to the `Cell` \rightarrow `BlockLattice3D` \rightarrow `SuperLattice` ansatz used for the implementation of the LBM. The equiv-

alent classes in the context of Lagrangian particles are `Particle3D` \rightarrow `ParticleSystem3D` \rightarrow `SuperParticleSystem3D`. The class `Particle3D` allocates memory for the variables of one single particle, such as its position, velocity, mass, radius and the force acting on it. It also provides the function `bool` `getActive()`, which returns the active state of the particle. Active particles' positions are updated during the simulation, in contrast to non-active particles, which are only used for particle-particle interaction. The class `Particle3D` is intended to be inherited from, in order to provide additional properties, such as electric or magnetic charge. The particles in the domain of a specific `BlockLattice3D` are combined in the class `ParticleSystem3D`. Finally the class `SuperParticleSystem3D` combines all `ParticleSystem3Ds`, and handles the transfer of particles between them.

The concept of the class `SuperParticleSystem3D` is to provide an easily adaptable framework for simulation of a large number of particles arranged in and interacting with a fluid. In this context *easily adaptable* means that simulated forces and boundary conditions are implemented in a modular manner, such that they are easily exchangeable. Development of new forces and boundary conditions can be readily done by inheritance of provided base classes. Particle-particle interaction can be activated if necessary and deactivated to decrease simulation time. The contact detection algorithm is interchangeable. This section introduces the `SuperParticleSystem3D` and the mentioned properties in more detail.

The class `SuperParticleSystem3D` is initialized by a call to the constructor simultaneously on all PUs.

```
1 SuperParticleSystem3D(CuboidGeometry3D<T>& cuboidGeometry, LoadBalancer<T>&
   loadBalancer, SuperGeometry<T,3>& superGeometry);
```

During the construction each PU instantiates one `ParticleSystem3D` for each local cuboid. Subsequently for each `ParticleSystem3D` a list of the ranks of PUs holding neighboring cuboids is created.

Particles can be added to the `SuperParticleSystem3D` by a call to one of the `addParticle()` functions.

```
1 /// Add a Particle to SuperParticleSystem
2 void addParticle(PARTICLETYPE<T> &p);
3 /// Add a number of identical Particles equally distributed in a given IndicatorF3D
4 void addParticle(IndicatorF3D<T>& ind, T mas, T rad, int no=1, std::vector<T> vel
   ={0.,0.,0.});
5 /// Add a number of identical Particles equally distributed in a given Material
   Number
6 void addParticle(std::set<int> material, T mas, T rad, int no=1, std::vector<T> vel
   ={0.,0.,0.});
7 /// Add Particles form a File. Save using saveToFile(std::string name)
8 void addParticlesFromFile(std::string name, T mass, T radius);
```

Currently there are four implementations of this class. The first adds single predefined particles, the second and third add a given number of equally distributed particles of the same

mass and radius in an area that can be defined by either a set of material numbers or an indicator function. The initial particle velocity can be set optionally. Finally particles can be added from an external file, containing their positions. In all cases the assignment to the correct ParticleSystem3D is carried out internally.

Particle forces and boundaries are implemented by the base classes Force3D and Boundary3D.

```

1  template<typename T, template<typename U> class PARTICLETYPE>
2  class Force3D {
3  public:
4      Force3D();
5      virtual void applyForce(typename std::deque<PARTICLETYPE<T> >::iterator p, int pInt
        , ParticleSystem3D<T, PARTICLETYPE>& psSys)=0;
6  }
```

Both classes are intended to be derived from in order to implement force and boundary specializations. The key function in both classes are applyForce() and applyBoundary(), which are called during each timestep of the main LBM loop. Force3D and Boundary3D specialisations are added to the SuperParticleSystem3D by passing a pointer to a class instantiation via a call to the respective function.

```

1  /// Add a force to system
2  void addForce(std::shared_ptr<Force3D<T, PARTICLETYPE> > f);
3  /// Add a boundary to system
4  void addBoundary(std::shared_ptr<Boundary3D<T, PARTICLETYPE> > b);
```

Both functions add the passed pointer to a list of forces and boundaries, which will be looped over during the simulation step. If necessary a contact detection algorithm can be added.

```

1  /// Set contact detection algorithm for particle-particle contact.
2  void setContactDetection(ContactDetection<T, PARTICLETYPE>& contactDetection);
```

A force based on contact between two particles is the contact force like described in the theory of Hertz and others and is named here as HertzMindlinDeresiewicz3D.

```

1  auto hertz = make_shared < HertzMindlinDeresiewicz3D<T, PARTICLE, DESCRIPTOR>
2      > (0.0003e9, 0.0003e9, 0.499, 0.499);
3  spSys.addForce(hertz);
```

Finally one timestep is computed by a call to the function simulate().

```

1  template<typename T, template<typename U> class PARTICLETYPE>
2  void SuperParticleSystem3D<T, PARTICLETYPE>::simulate(T dT)
3  {
4      for (auto pS : _pSystems) {
5          pS->_contactDetection->sort();
6          pS->simulate(dT);
7          pS->computeBoundary();
8      }
9      updateParticleDistribution();
10 }

```

This function contains a loop over the local `ParticleSystem3Ds` calling the local sorting algorithm and the functions `ParticleSystem3D::simulate()` and `ParticleSystem3D::computeBoundary()`. The sorting algorithm determines potential contact between particles according to the set `ContactDetection`.

```

1  inline void simulate(T dT) {
2      _pSys->computeForce();
3      _pSys->explicitEuler(dT);
4  }

```

The inline function `ParticleSystem3D::simulate()` first calls the local function `ParticleSystem3D::computeForce()`.

```

1  template<typename T, template<typename U> class PARTICLETYPE>
2  void ParticleSystem3D<T, PARTICLETYPE>::computeForce()
3  {
4      typename std::deque<PARTICLETYPE<T> >::iterator p;
5      int pInt = 0;
6      for (p = _particles.begin(); p != _particles.end(); ++p, ++pInt) {
7          if (p->getActive()) {
8              p->resetForce();
9              for (auto f : _forces) {
10                 f->applyForce(p, pInt, *this);
11             }
12         }
13     }
14 }

```

This function consists of a loop over all particles stored by the calling `ParticleSystem3D`. If the particle state is active, its force variable is reset to zero. Then the value computed by each previously added particle force is added to the particle's force variable. Finally, the particle velocity and position is updated by one step of an integration method.

Returning to the function `SuperParticleSystem3D::simulate(T dT)` the next command in the loop is a call of the function `ParticleSystem3D::computeBoundary()`, which has the same structure as the `ParticleSystem3D::computeForce()`. After executing the loop, the function `updateParticleDistribution()` is called, which redistributes the particles over the `ParticleSystem3Ds` according to their updated position. A detailed description of this function is provided at the end of the next section.

4.6.9.3. Implementation of the *Communication Optimal Strategy*

The *communication optimal strategy* is implemented in the function `SuperParticleSystem3D::updateParticleDistribution()` already mentioned above. The function has to be called after every update of the particle positions, in order to check if the particle remained in its current cuboid, as otherwise segmentation faults may occur during the computation of particle forces. The transfer is implemented using nonblocking operations of the MPI library.

```
1  template<typename T, template<typename U> class PARTICLETYPE>
2  void SuperParticleSystem3D<T, PARTICLETYPE>::updateParticleDistribution()
3  {
4      /* Find particles on wrong cuboid, store in relocate and delete */
5      //maps particles to their new rank
6      _relocate.clear();
7      for (unsigned int pS = 0; pS < _pSystems.size(); ++pS) {
8          auto par = _pSystems[pS]->_particles.begin();
9          while (par != _pSystems[pS]->_particles.end()) {
10             //Check if particle is still in his cuboid
11             if (checkCuboid(*par, 0)) {
12                 par++;
13             }
14             //If not --> find new cuboid
15             else {
16                 findCuboid(*par, 0);
17                 _relocate.insert(
18                     std::make_pair(this->_loadBalancer.rank(par->getCuboid()), (*par)));
19                 par = _pSystems[pS]->_particles.erase(par);
20             }
21         }
22     }
```

The function begins with with two nested loops. The outer loop is over all local `ParticleSystem3DS`, the inner loop over the `Particle3DS` of the current `ParticleSystem3D`. Each particle is checked if it remained in its cuboid during the last update, by the function `checkCuboid(*par, 0)`. The first parameter of `checkCuboid(*par, 0)` is the particle to be tested and the second parameter is an optional spatial extension of the cuboid. If the function returns `true` the counter is incremented and the next particle is tested. If the function returns `false` the particle together with the rank of its new cuboid are copied to the `std::multimap<int, PARTICLETYPE<T> > _relocate` for future treatment and removed from the `std::deque<PARTICLETYPE<T> > _particles` of particles.

```

1  /* Communicate number of Particles per cuboid*/
2  singleton::MpiNonBlockingHelper mpiNbHelper;
3
4  /* Serialise particles */
5  _send_buffer.clear();
6  T buffer[PARTICLETYPE<T>::serialPartSize];
7  for (auto rN : _relocate) {
8      rN.second.serialize(buffer);
9      _send_buffer[rN.first].insert(_send_buffer[rN.first].end(), buffer, buffer+
          PARTICLETYPE<T>::serialPartSize);
10 }

```

The function continues by instantiating the class `singleton::MpiNonBlockingHelper`, which handles memory for `MPI_Request` and `MPI_Status` messages. Then the particles buffered in `_relocate` are serialized. Meaning their data is written consecutively in memory and stored in a buffer `std::map<int, std::vector<double> > _send_buffer` in preparation for the transfer.

```

1  /*Send Particles */
2  int noSends = 0;
3  for (auto rN : _rankNeighbours) {
4      if (_send_buffer[rN].size() > 0) {
5          ++noSends;
6      }
7  }
8  mpiNbHelper.allocate(noSends);
9  for (auto rN : _rankNeighbours) {
10     if (_send_buffer[rN].size() > 0) {
11         singleton::mpi().iSend<double>(&_send_buffer[rN][0], _relocate.count(rN)*
            PARTICLETYPE<T>::serialPartSize, rN, &mpiNbHelper.get_mpiRequest()[k++],
            1);
12     }
13 }
14 singleton::mpi().barrier();

```

To find the number of send operations a loop over the ranks of neighboring cuboids is carried out, increasing the variable `count` each time data for a specific rank is available. Then the appropriate number of `MPI_Requests` is allocated. Finally the data is sent to the respective PUs via a nonblocking `MPI_Isend()` and all PUs wait until the send process is finished on each PU.

```

1  /*Receive and add particles*/
2  int flag = 0;
3  MPI_Iprobe(MPI_ANY_SOURCE, 1, MPI_COMM_WORLD, &flag, MPI_STATUS_IGNORE);
4  if (flag) {
5      for (auto rN : _rankNeighbours) {
6          MPI_Status status;
7          int flag = 0;
8          MPI_Iprobe(rN, 1, MPI_COMM_WORLD, &flag, &status);
9          if (flag) {
10             int amount = 0;
11             MPI_Get_count(&status, MPI_DOUBLE, &number_amount);
12             T recv_buffer[amount];
13             singleton::mpi().receive(recv_buffer, amount, rN, 1);
14             for (int iPar=0; iPar<amount; iPar+=PARTICLETYPE<T>::serialPartSize) {
15                 PARTICLETYPE<T> p;
16                 p.unserialize(&recv_buffer[iPar]);

```

```

17         if (singleton::mpi().getRank() == this->_loadBalancer.rank(p.getCuboid()))
18             {
19                 _pSystems[this->_loadBalancer.loc(p.getCuboid())]->addParticle(p);
20             }
21     }
22 }
23 }
24 if (noSends > 0) {
25     singleton::mpi().waitAll(mpiNbHelper);
26 }
27 }

```

On the receiving side the nonblocking routine `MPI_Iprobe()` checks whether an incoming transmission is available. The constant `MPI_ANY_SOURCE` indicates that messages from all ranks are accepted. If a message is awaiting reception the flag `flag` is set to a nonzero value and the following switch will be true. This query is not necessary, but the following loop can be entirely skipped if no particles are transferred, which is expected to be the case most of the time.

The subsequent loop tests for each single neighboring rank if a message awaits reception. If **true** the number of send `MPI_Doubles` is read from the `status` variable via an `MPI_Get_count()`. The appropriate memory is allocated and the message is received by wrapped call to `MPI_Recv()`, and written consecutively. Then new `Particle3DS` are instantiated, initialized with the received data and assigned to the respective `ParticleSystem3D` on the updated PU. Finally, a call to `MPI_Waitall()` makes sure, that all `MPI_Isend()`s have been processed by the recipients.




Figure 4.2.: Overlap of the particle domains. Particles within a distance to of the sum of the two largest radii to a neighbor cuboid have to be transferred to this specific neighbor cuboid.

4.6.9.4. Shadow Particles

If particle collisions are considered, it may happen that particles P_m with center $X_m \in \tilde{\Omega}^j$ collide with particles P_n with center $X_n \in \tilde{\Omega}^k$ in a different cuboid, as illustrated in Figure 4.2.

Therefore P_n has to be known on $\mathbf{X}_m \in \tilde{\Omega}^j$ and so-called shadow particles are introduced. Shadow particles are static particles, whose positions and velocities are not explicitly computed during the update step. Particle collision across cuboid boundaries can only occur if the distance $d = \|\mathbf{X}_n - \mathbf{X}_m\|_2$ between the participating particles is less than the sum of the two largest radii of all particles in the system. Hence the width of the particle overlap has to be at least the sum of the two largest particle radii and all particles within this overlap have to be transferred to the neighbor cuboid after each update of the particle position by an additional communication step similar to the one introduced above.

5. Initial and Boundary Conditions

Each example located in `examples/` is typically structured along several function definitions and a `main` part. The function `setBoundaryValues` usually sets initial values for density and velocity (if not yet done in `prepareLattice`). Boundary values can be refreshed at certain time steps. In some applications, this function may be missing or empty, if there is no temporal change in the boundary values. An exemplary implementation can be found in `examples/laminar/poiseuille2d`, where a smooth start-up is used at the velocity inflow boundary as follows:

```
1 // No of time steps for smooth start-up
2 int iTmaxStart = converter.getLatticeTime( maxPhysT*0.4 );
3 int iTupdate = 5;
4
5 if (iT%iTupdate == 0 && iT <= iTmaxStart) {
6
7     // Smooth start curve, polynomial
8     PolynomialStartScale<T,T> StartScale(iTmaxStart, T(1));
9     // Creates and sets the Poiseuille // inflow profile using functors
10    T iTvec[1] = {T( iT )};
11    T frac[1] = {};
12    StartScale( frac,iTvec );
13    T maxVelocity = converter.getCharLatticeVelocity()*3./2.*frac;
14    T distance2Wall = L/2.;
15    Poiseuille2D<T>
16        poiseuilleU(superGeometry, 3,
17        maxVelocity, distance2Wall);
18    sLattice.defineU(superGeometry, 3,
19        poiseuilleU);
20 }
```

5.1. Define Boundary Method

There are two different types of boundaries, namely wet-node boundaries (on the nodes) and link-wise boundaries (in between the nodes). Examples of wet-node boundaries are `LocalVelocity`, `LocalPressure`, `InterpolatedPressure`, `InterpolatedVelocity`, `Zou-HePressure`, `ZouHeVelocity` and `AdvectionDiffusion`. Examples of link-wise boundaries are `Bouzidi`, `BouzidiVelocity`, `BouzidiZeroVelocity (+Interpolation)` and `YuPostProcessor (+Interpolation)`. The boundary declarations are usually designed as in the following code snippets:

```
1 setBounceBackBoundary(superLattice, superGeometry, materialNumber);
```

```

2 setLocalVelocityBoundary<T, DESCRIPTOR>(superLattice, omega, superGeometry,
   materialNumber);
3 setBouzidiBoundary(superLattice, superGeometry, materialNumber, indicator);

```

OpenLB offers implementations of several methods for approximating macroscopic boundary conditions. Some of the frequently used boundary methods are listed below and described in terms of locality of operations and general applicability. This list does not raise a claim to completeness. OpenLB includes the following boundary methods, which are callable with `set...Boundary` and the respective arguments.

1. **LocalPressure, LocalVelocity, LocalConvection, etc.:**

- local
- wet-node
- *Application:* boundary for fluid flows
- regularized boundary, re-computes all f_i from local momenta reconstructing off-equilibrium parts
- stable in most regimes
- implemented according to Latt and Chopard [105]

2. **InterpolatedPressure, Interpolated, etc.:**

- non-local
- wet-node
- *Application:* boundary for fluid flows (stable for higher Reynolds numbers)
- re-computes all f_i using a finite-difference scheme over adjacent cells for the velocity gradient
- implemented according to Skordos [119]

3. **setZouPressure, setZouHeVelocity, etc.:**

- local
- wet-node
- *Application:* boundary for fluid flows (low Reynolds number)
- computes missing f_i by applying symmetry conditions on the off-equilibrium part, enforcing velocity and pressure on the equilibrium
- highly accurate, but less stable
- implemented according to Zou and He [127]

4. **Bouzidi, BouzidiVelocity, etc.:**

- non-local
- link-wise
- *Application:* boundary for fluid flows (for curved-boundaries)
- computes missing f_i using a bounce-back rule which takes the distance between node and boundary into account
- second order accuracy in space

- implemented according to Bouzidi *et al.* [80]

5. `AdvectionDiffusionTemperature`, etc.:

- local and non-local
- wet-node
- *Application*: boundary for advection-diffusion problems (temperature, particle, ..)
- various boundary conditions for advection diffusion problems
- type of implementation differs for each condition
- for details see e.g. Trunk *et al.* [45]

The typical macroscopic effects that can be obtained with the wet-node-method are for example velocity and pressure Dirichlet boundary conditions and for the link-wise methods for example Dirichlet boundary conditions for a specific velocity, zero-velocity, convection and slip walls. With the `advectionDiffusion`-methods macroscopic boundary conditions in terms of Dirichlet boundaries with respect to temperature, convection, zero-distribution, or the external-field, are recoverable.

5.1.1. Wet-node Method

With the wet-node approach, for example Dirichlet-type boundaries for the velocity can be obtained. On a macroscopic level, this is used for example at inflow boundaries where the values for the inflow velocity are given.

```
1 setLocalVelocityBoundary<T, DESCRIPTOR> (superLattice, omega, superGeometry,
    materialNumber);
```

Moreover, the wet-node approach is applied for macroscopic boundaries for the pressure. This is used for example at outflow boundaries and fixes the values for the pressure in terms of a Dirichlet condition.

```
1 setLocalPressureBoundary<T, DESCRIPTOR> (superLattice, omega, superGeometry,
    materialNumber);
```

5.1.2. Link-wise Method

The link-wise method is applicable to recover macroscopic convection, which is used for outflow boundaries and approximates

$$\frac{\partial u}{\partial t} + u_{\text{average}} \frac{\partial u}{\partial n} = 0, \quad (5.1)$$

similar to a Neumann-type boundary for the velocity.

```
1 setInterpolatedConvectionBoundary<T, DESCRIPTOR> (superLattice, omega, superGeometry,
    materialNumber, averageVelocity);
```

Additionally, a slip boundary can be constructed that is used for solid boundaries and reflects outgoing f_i . The latter has the effect of zero velocity normal to the boundary and free flow tangential to the boundary.

```
1 setSlipBoundary<T, DESCRIPTOR> (superLattice, superGeometry, materialNumber);
```

Further, a macroscopic velocity used for curved boundaries with fixed values can be superimposed. It is realized as a Dirichlet boundary for the velocity and considers the boundary shape via an indicator functor. There are multiple `bulkMaterials` possible.

```
1 setBouzidiBoundary<T, DESCRIPTOR, BouzidiVelocityPostProcessor> (superLattice,
    superGeometry, material, indicator, bulkMaterialsList);
2 AnalyticalConstF3D<T> u;
3 setBouzidiVelocity(superLattice, superGeometry, material, u);
```

As a special case, the zero velocity boundary for curved boundaries which fixes the velocity to zero is callable with simplified syntax. It is a Dirichlet-type boundary for velocity and considers the shape by indicator. Here are also multiple `bulkMaterials` possible.

```
1 setBouzidiBoundary<T, DESCRIPTOR> (superLattice, superGeometry, material, indicator,
    bulkMaterialsList);
```

5.1.3. AdvectionDiffusionBoundary

An exemplary macroscopic effect of the `advectionDiffusion` boundary methods is the Dirichlet boundary condition for the temperature, which is used for inflow or wall boundaries and fixes the values for the temperature.

```
1 setAdvectionDiffusionTemperatureBoundary<T, DESCRIPTOR> (superLattice, omega,
    superGeometry, materialNumber);
```

In addition, convection can be modeled, which is used for outflow boundaries and interpolates incoming f_i from neighbors. It is similar to a Neumann-type temperature boundary.

```
1 setAdvectionDiffusionConvectionBoundary<T, DESCRIPTOR> (superLattice, superGeometry,
    materialNumber);
```

Further, the zero-distribution boundary is used for "sticky" boundaries and sets the incoming f_i to zero, such that particles touching the boundary are *trapped*.

```
1 setZeroDistributionBoundary<T, DESCRIPTOR> (superLattice, superGeometry, materialNumber
    );
```

At last also macroscopic external fields used for example in particle simulations can be supplied with boundary conditions. The following boundary method provides data on the boundary for particle calculations (velocity gradient).

```
1 setExtFieldBoundary<T, DESCRIPTOR, FIELD > (superLattice, superGeometry, materialNumber
    );
```

5.1.4. Additional Options

An additional option to supply boundary conditions in OpenLB is to choose LBM collision dynamics also at the boundary. One can add a specific dynamics object to a boundary material number, e.g.: `BulkDynamics` which executes the collision step as in a fluid node (e.g. with `BGKdynamics`).

```
1      using BulkDynamics = BGKdynamics<T,DESCRIPTOR>;
2      [...]
3      setLocalVelocityBoundary<T,DESCRIPTOR, BulkDynamics>(superLattice,omega,
                    superGeometry,materialNumber);
```

Another option is to choose an implementation via boundary dynamics. One can add boundary methods as dynamics, e.g. `BounceBack` (zero velocity), `BounceBackVelocity` (prescribes a nonzero velocity).

```
1      superLattice.defineDynamics<BounceBack> (superGeometry, materialNumber);
2      superLattice.defineDynamics <BounceBackVelocity>(superGeometry, materialNumber);
```

5.2. Define Initial Conditions

```
1      sLattice.defineRhoU(superGeometry, materialNumber, analyticalFunctor,
                    analyticalFunctor);
```

For each material number, the density (usually $\rho = 1$), the velocity (usually $u = 0$), and the distribution functions should be initialized. The functions expect lattice values, therefore physical values have to be converted, e.g.: via the function `getLatticeDensity(density)` of the `UnitConverter`. Note that instead of the `materialNumber` argument, any discrete indicator function can be used. Exemplary initializations are given below.

```
1  // Initial conditions
2  AnalyticalConst2D<T,T> rhoF(1);
3  std::vector<T> velocity(dim,T(0));
4  AnalyticalConst2D<T,T> uF(velocity);
5
6  // Initialize density
7  superLattice.defineRho(superGeometry, materialNumber, rhoF);
8
9  // Initialize velocity
10 superLattice.defineU(superGeometry, materialNumber, uF);
11
12 // Initialize distribution functions
13 // to local equilibrium
14 superLattice.iniEquilibrium(superGeometry, materialNumber,
15                             rhoF, uF);
```

5.3. Define Boundary Values

```
1 sLattice.defineU(superGeometry, materialNumber, analyticalFunctor);
```

Just like the initial conditions, boundary values can be set using `defineRho(...)` and `defineU(...)`. For a smooth start-up, the values can be scaled, e.g. according to a sinus-scale or polynomial-scale for a given startup-time.

```
1 // Smooth start curve, sinus
2 SinusStartScale<T, int>
3 StartScale(numerStartTimeSteps, maxValue);
4
5 //Smooth start curve, polynomial PolynomialStartScale<T, T>
6 StartScale(numerStartTimeSteps, maxValue);
7
8 // compute scale-factor "frac"
9 T iTvec[1] = {T(timestep)};
10 T frac[1] = {};
11 StartScale( frac, iTvec );
```

To apply a flow profile, one first has to update values in every n th time step, then initialize a functor and then set values using `defineRho(...)` and `defineU(...)`. These are the same functions as for the initial conditions. However, the time point when to call them is crucial.

```
1 if (timestep%updatePeriod==0 &&
2     timestep <= numberStartupTimesteps) {
3     Poiseuille2D<T> poiseuilleU(superGeometry,
4         materialNumber, maxVelocity*frac[0],
5         distance2Wall);
6     sLattice.defineU(superGeometry, materialNumber, poiseuilleU);
7 }
```

Further examples of 3D functors for this purpose are: rotating functors (linear for a rotating velocity field and quadratic for a rotating pressure field), Circle-Poiseuille, Elliptic-Poiseuille, and Rectangular-Poiseuille. The latter are only for on axis boundaries and can be constructed from points spanning a plane or a material number. These functors are summarized with arguments as follows:

```
1 RotatingLinear3D(axisPoint,
2     axisDirection, angularVelocity);
3
4 RotatingQuadratic1D(axisPoint, axisDirection, angularVelocity);
5
6 CirclePoiseuille3D(axisPoint,
7     axisDirection, maxVelocity, radius);
8
9 EllipticPoiseuille3D(center, semiPrincipalAxis1, semiPrincipalAxis2, maxVelocity);
10
11 RectanglePoiseuille3D(point1, point2,
12     point3, maxVelocity);
13
14 RectanglePoiseuille3D(superGeometry, materialNumber, maxVelocity, offsetX, offsetY,
15     offsetZ);
```

6. Input and Output

During development or during actual simulations, it might be necessary to parametrize your program. For this case, OpenLB provides an XML parser, which can read files produced by OpenMPI [71], thereby providing a user-friendly GUI. Details on the XML format and functions are given in Section 6.8.

The simulation data is stored in the VTK data format and can be further processed with ParaView. For output tasks that are performed only once during the simulation, it is recommended to write the data sequentially. Commonly, the geometry or cuboid information is one of these tasks. In contrast to the parallel version, it is easier to use and does not produce unnecessary data overhead. However, if the output is performed regularly in a parallel simulation, the performance may slow down using the sequential output method. Therefore, OpenLB has implemented a parallel data output functionality. At the lowest scope, every thread writes `.vti` files that contain the data. OpenLB writes a `.vti` file for every cuboid, to provided parallel data processing. Those `.vti` files are summarized and put together by the next hierarchy, the `.vtm` file. A `.vtm` file corresponds to the entire domain with respect to a certain time step. At the end, the different time steps are packed to a `.pvd` file, that is a collection of `.vtm` according to time steps.

The technical aspects are presented in Section 6.1, whereas the usage is demonstrated with an example in Section 6.2.

6.1. Output Data Structure

OpenLB simulation data is stored in file system according to the VTK data format [72]. This format has XML structure and the data therein is written as binary `Base64` code. Additionally, the simulation data is compressed by `zlib`, which allows to reduce data tremendously.

On the top level, the parallel output hierarchy contains a `.pvd` file, which consists of links to `.vtm` files. The `.vtm` files summarize the cuboids represented by `.vti` files.

```
<?xml version="1.0"?>
<VTKFile type="Collection" version="0.1" byte_order="LittleEndian">
  <Collection>
    <DataSet timestep="81920" group="" part="" file="data/VTM_iT0081920.vtm"/>
    <DataSet timestep="163840" group="" part="" file="data/VTM_iT0163840.vtm"/>
    <DataSet timestep="245760" group="" part="" file="data/VTM_iT0245760.vtm"/>
    <DataSet timestep="327680" group="" part="" file="data/VTM_iT0327680.vtm"/>
    <DataSet timestep="409600" group="" part="" file="data/VTM_iT0409600.vtm"/>
  </Collection>
</VTKFile>
```

```
</Collection>
</VTKFile>
```

Listing 6.1.: Example of a .pvd file that points for every time step to the corresponding .vtm file. Every time step is associated to a .vtm file.

```
<?xml version="1.0"?>
<VTKFile type="vtkMultiBlockDataSet" version="1.0" byte_order="LittleEndian">
<vtkMultiBlockDataSet>
<Block index="0" >
<DataSet index="0" file="VTM_iT0081920iC00000.vti"></DataSet>
</Block>
<Block index="1" >
<DataSet index="0" file="VTM_iT0081920iC00001.vti"></DataSet>
</Block>
<Block index="2" >
<DataSet index="0" file="VTM_iT0081920iC00002.vti"></DataSet>
</Block>
<Block index="3" >
<DataSet index="0" file="VTM_iT0081920iC00003.vti"></DataSet>
</Block>
</vtkMultiBlockDataSet>
</VTKFile>
```

Listing 6.2.: Example of a .vtm file that points to .vti files that hold data of a cuboids. Every cuboid writes its data to a .vti file, which are assembled by a .vtm file.

There is also a `BlockVTKWriter` that writes data sequentially. More details can be found in the source code and its documentation.

6.2. Write Simulation Data to VTK File Format

VTK data files can be visualized and postprocessed with the free software ParaView [70], which offers a graphical interface with extensive functionality. The following listing shows, on the one hand, how to write VTK files sequential for a geometry and cuboid functors. On the other hand, the usage of the parallel write-routine for velocity and pressure functors is shown.

```
1 // create VTK writer object
2 SuperVTMWriter3D<T> vtmWriter("FileNameGoesHere");
3 // write only the first iteration step
4 if (iT==0) {
5     SuperLatticeGeometry3D<T,DESCRIPTOR> geometry(sLattice, superGeometry);
6     SuperLatticeCuboid3D<T,DESCRIPTOR> cuboid(sLattice);
7     // writes the geometry and cuboids to file system, sequentially
8     vtmWriter.write(geometry);
9     vtmWriter.write(cuboid);
10    // mandatory to call the following write()-method
11    vtmWriter.createMasterFile();
12 }
13 // write every 2 sec (physical time scale)
14 if (iT%converter.getLatticeTime(2.)==0) {
15     // create functors that process data from SuperLattice
16     SuperLatticePhysVelocity3D<T,DESCRIPTOR> velocity(sLattice,
17                                                         converter);
```

```

18 SuperLatticePhysPressure3D<T,DESCRIPTOR> pressure(sLattice,
19                                                    converter);
20 vtmWriter.addFunctor( velocity );
21 vtmWriter.addFunctor( pressure );
22 // writes the added functors to file system, parallel
23 vtmWriter.write(iT);
24 }

```

Listing 6.3.: An exemplary code to write simulation data to file system.

Note that the function call `creatMasterFile()` in `iT == 0` is essential to write parallel VTK data.

6.3. CSV Writer

For some data analysis a CSV format of the data is necessary. In this case it is possible to use the CSV Writer to create these data files. The following lines show an application of the CSV Writer in the example `advectionDiffusion1d` (8.3.3). If one only wants to write in one data file, the filename can be given to the constructor of the CSV Writer. However the `plotFileName` parameter provides the possibility to set a new datafile with every call of this function. The `precision` parameter refers to the precision of the output data.

```

1 CSV<T> csv();
2 csv.writeDataFile(N, simulationAverage, "averageSimL2RelErr");

```

Listing 6.4.: Exemplary application of the CSV Writer

6.4. Write Images Instantaneously

OpenLB is able to output image data directly. This is helpful to get a brief overview of how the simulation is going on without using external visualization tools. Note that only 1D data or equivalent scalar-valued data can be represented by images. Hence, for vector-valued data, e.g. velocity, it is important to take an appropriate norm. This step transforms the vector into a scalar and the data becomes one dimensional as required.

For 2D applications it is straight forward to generate images, since every point of the computational grid represents a pixel. However, for 3D applications this assignment fails. OpenLB allows one to reduce the 3D grid to a 2D plane by parametrizing a hyperplane in 3D space. The resulting 2D block lattice represents the image by assigning lattice points to pixels.

An example of how to take a norm and how to reduce a plane is shown below.

```

1 // get the pointwise l2 norm of velocity
2 SuperEuklidNorm3D<T> normVel( velocity );
3 // reduce a hyperplane parametrized by normal (0,0,1) and centered in the mother
  geometry from the 3D data
4 BlockReduction3D2D<T> planeReduction( normVel, {0, 0, 1} );

```

Listing 6.5.: An exemplary code reducing a plane in $3D$

Note that internally the hyperplane is parametrized using the `Hyperplane3D` class. This example uses one of the helper constructors of `BlockReduction3D2D` to hide this detail for the common use case of parametrizing a hyperplane by a normal vector. There are further such helper constructors available if one wishes to for example define a hyperplane by two span vectors and its origin. However for full control over the hyperplane a `Hyperplane3D` instance may also be created by hand.

```
1 SuperEuklidNorm3D<T> normVel( velocity );
2 BlockReduction3D2D<T> planeReduction(
3     normVel,
4     // explicitly construct a 3D hyperplane
5     Hyperplane3D<T>()
6     .centeredIn(superGeometry.getCuboidGeometry().getMotherCuboid())
7     .spannedBy({1, 0, 0}, {0, 1, 0}));
8 BlockGifWriter<T> gifWriter;
9 gifWriter.write(planeReduction, iT, "vel" );
```

Listing 6.6.: Exemplary code to write images of an explicitly instantiated $3D$ hyperplane with

Both of these exemplary codes reduce a $3D$ hyperplane to a $2D$ lattice with 600 points on its longest side. It is possible to change this resolution either by providing it as a constructor argument to `BlockReduction3D2D` or by explicitly instantiating a `HyperplaneLattice3D`.

```
1 SuperEuklidNorm3D<T> normVel( velocity );
2 HyperplaneLattice3D<T> gifLattice(
3     superGeometry.getCuboidGeometry(),
4     Hyperplane3D<T>()
5     .centeredIn(superGeometry.getCuboidGeometry().getMotherCuboid())
6     .normalTo({0, -1, 0}),
7     // resolution (floating point values are used as grid spacing instead)
8     1000);
```

Listing 6.7.: Exemplary code using an explicitly instantiated $3D$ hyperplane lattice

In $2D$ the reduction of velocity data to a block can be achieved as follows.

```
1 SuperEuklidNorm2D<T, DESCRIPTOR> normVel( velocity );
2 BlockReduction2D2D<T> planeReduction( normVel );
```

Listing 6.8.: Exemplary code reducing data in $2D$

The resolution of 600 points on the longest side of the object is set as default, but can be altered similarly to the Listings 6.5, 6.6 and 6.7. There are two options of generating images of the processed values in $2D$ and $3D$.

6.4.1. GifWriter

In this example the constructor `gifWriter` generates automatically scaled images of the PPM data type which are scaled according to the minimum and maximum value of the desired value of the time step.

```
1 BlockReduction3D2D<T> planeReduction(normVel, gifLattice);
2 BlockGifWriter<T> gifWriter;
3 //gifWriter.write(planeReduction, 0, 0.7, iT, "vel"); //static scale
4 gifWriter.write( planeReduction, iT, "vel" ); // scaled
```

Listing 6.9.: Exemplary code using `gifWriter` to create PPM files

With `imagemagick`'s command `convert` the PPM files generated by `gifWriter` can be combined to an animated GIF file as follows.

```
convert tmp/imageData/*.ppm animation.gif
```

To reduce the GIF's file size you can use the options `fuzz` and `OptimizeFrame`, for example:

```
convert -fuzz 3% -layers OptimizeFrame tmp/imageData/*.ppm animation.gif
```

Even smaller files are possible with `ffmpeg` and conversion to MP4 video file. This could be done using a command like the following.

```
ffmpeg -pattern_type glob -i 'tmp2/imageData/*.ppm' animation.mp4
```

6.4.2. Heatmap

Whereas the `gifWriter` creates only automatically scaled PPM images, the functor `heatmap` has more options to adjust the JPEG files. For this purpose the variable `plotParam` can be created and the desired modifications, e.g. minimum and maximum values of the scale, can be passed on to the optional variable.

```
1 SuperEuklidNorm3D<T> normVel( velocity );
2 BlockReduction3D2D<T> planeReduction( normVel, {0, 0, 1} );
3 // write output as JPEG and changing properties
4 heatmap::plotParam<T> jpeg_Param;
5 jpeg_Param.contourlevel = 5; //setting the number of contour lines
6 jpeg_Param.colour = "rainbow"; //colour combination "grey", "pm3d", "blackbody" and "
    rainbow" can be chosen
7 heatmap::write(planeReduction, iT, jpeg_Param);
```

Listing 6.10.: Exemplary code using the functor `heatmap` with modified parameters

The exemplary code in listing 6.10 shows how to change the color set and number of contour lines in the generated images. All possible adjustments are listed and used in the example `venturi3D` (see Section 8.11.5).

6.5. Gnuplot Interface

Often, for the analysis of simulations a plot of the data is required. OpenLB offers an interface which uses `Gnuplot` to create plots. Furthermore, it is possible to see the particular data that was used for the plots in realtime and to use comparison data, which is directly used in the plot. An example for the usage from `examples/cylinder2d` is shown below.

```
1 // Gnuplot constructor (must be static!)
2 // for real-time plotting: gplot("name", true) // experimental!
3 static Gnuplot<T> gplot( "drag" );
4
5 ...
6
7 // set data for gnuplot: input={xValue, yValue(s),
8 // names (optional), position of key (optional)}
9 gplot.setData( converter.getPhysTime( iT ), {_drag[0], 5.58},
10 {"drag(openLB)", "drag(schaeferTurek)"}, "bottom right" );
11
12 // writes a png (or optional pdf) in one file for every timestep,
13 // if the png file is opened by an imageviewer it can be used as a "liveplot"
14 // optional for pdf output, use: gplot.writePDF()
15 gplot.writePNG();
16 }
```

Listing 6.11.: An exemplary code to plot simulation data.

The data `drag[0]` is calculated in the example and compared with the value 5.58. This is then plotted as shown in Figure 6.1.




Figure 6.1.: Gnuplot output of drag calculation in cylinder2d.

In order to have plots for different times, the following usage is recommended.

```

1  ...
2
3  // every (iT%vtkIter) write an png of the plot
4  if ( iT%( vtkIter ) == 0 ) {
5    // writes pngs: input={name of the files (optional),
6    // x range for the plot (optional)}
7    gplot.writePNG( iT, maxPhysT );

```

Listing 6.12.: Creating plots for different time steps.

6.5.1. Regression with Gnuplot

Moreover, Gnuplot can be used to create a linear regression of datasets. For instance, the analysis of the experimental order of convergence in a simulation can be executed as in the example `poiseuille2dEOC`.

The possible options are: Linear regression to the given data whereas it is possible to use a loglog-scaling (`loglogINVERTED` for inverting the x-axis). The implementation is done via the constructor of plot in the `.cpp` file itself as seen below.

```

1  static Gnuplot<T> gplot( "eoc", Gnuplot<T>::LOGLOG, Gnuplot<T>::LINREG );

```

The possible options for the scaling are: `LINEAR` (using the data as given), `LOGLOG` (using log of the x- and y-dataset) and `LOGLOGINVERTED` (using log of y-dataset and 1/log of x-dataset). For the regression type one can choose `LINREG` (linear Regression) and `OFF` (no regression).




Figure 6.2.: Example of using regression to analyze polynomial errors (left: old, right: new implementation)

6.6. Console Output

In OpenLB, there is an extension of the default `ostream`, which handles parallel output and prefixes each line with the name of the class that produced the output. Listing 6.13 is the output of the `bstep2d` example.

It is easy to determine which part of OpenLB has produced a specific message. This can be very helpful in the debugging process, as well as for quickly postprocessing console output or filtering out important information without any need to go into the code. Together with OpenLB's semi-CSV style output standard, it is possible to easily visualize any data imaginable with diagrams, such as convergence rates, data errors, or simple average mass density.

```
1 void MyClass::print() {
2   OstreamManager clout(std::cout, "MyClass");
3   ...
4   clout << "step=" << step << "; avRho=" << avRho
5       << "; maxU=" << maxU << std::endl;
6 }
```

Using the `OstreamManager` is easy and consists of two parts. First, an instance of the class `OstreamManager` is needed. The one created here in line 2 is called `clout` like all the other instances in OpenLB. This word consists of the two words class and output. Moreover, it is quite similar to standard `cout`. The constructor receives two arguments: one describing the ostream to use, the other one setting the prefix-text. In line 4 the usage of an instance of the `OstreamManager` is shown. There is not much difference in usage between a default `std::cout` and an instance of OpenLB's `OstreamManager`. The only thing to consider is that a normal `"\\n"` won't have the expected effect, so use `std::endl` instead.

In classes with many output producing functions however, you wouldn't like to instantiate `OstreamManager` for every single function, so a central instantiation is preferred. This is done by adding a mutable `OstreamManager` object as a private class member and initializing it in the initialization list of each defined constructor. An example implementation of this method can be found in `src/utilities/timer.hh`.

Another great benefit of `OstreamManager` is the reduction of output in parallel. Running a program using `cout` on multiple cores normally means getting one line of output for each process. `OstreamManager` will avoid this by default and display only the output of the first processor. If this behavior is unwanted in a specific case, it can be turned off for an instance named `clout` by `clout.setMultiOutput(true)`.

Further scenarios that are not yet implemented in OpenLB can make use of different streams like the ostream `std::cerr` for separate error output, file streams, or something completely different. In doing so, every stream needs its own instance.

```

$ ./bstep2d
...
[prepareGeometry] Prepare Geometry ...
[SuperGeometry2D] cleaned 0 outer boundary voxel(s)
[SuperGeometry2D] cleaned 0 inner boundary voxel(s)
[SuperGeometry2D] the model is correct!
[SuperGeometryStatistics2D] materialNumber=0; count=13846; minPhysR=(0,0); maxPhysR=(5,0.75)
[SuperGeometryStatistics2D] materialNumber=1; count=92865; minPhysR=(0.0166667,0.0166667); maxPhysR
    =(19.9833,1.48333)
[SuperGeometryStatistics2D] materialNumber=2; count=2448; minPhysR=(0,0); maxPhysR=(20,1.5)
[SuperGeometryStatistics2D] materialNumber=3; count=43; minPhysR=(0,0.783333); maxPhysR=(0,1.48333)
[SuperGeometryStatistics2D] materialNumber=4; count=89; minPhysR=(20,0.0166667); maxPhysR=(20,1.48333)
[prepareGeometry] Prepare Geometry ... OK
[prepareLattice] Prepare Lattice ...
[prepareLattice] Prepare Lattice ... OK
[main] starting simulation...
[SuperPlaneIntegralFluxVelocity2D] regionSize[m]=1.46667; flowRate[m^2/s]=0; meanVelocity[m/s]=0
[SuperPlaneIntegralFluxPressure2D] regionSize[m]=1.46667; force[N]=0; meanPressure[Pa]=0
[Timer] step=0; percent=0; passedTime=0.846; remTime=101519; MLUPs=0
[LatticeStatistics] step=0; t=0; uMax=1.49167e-154; avEnergy=0; avRho=1
[SuperPlaneIntegralFluxVelocity2D] regionSize[m]=1.46667; flowRate[m^2/s]=0; meanVelocity[m/s]=0
[SuperPlaneIntegralFluxPressure2D] regionSize[m]=1.46667; force[N]=0; meanPressure[Pa]=0
[Timer] step=300; percent=0.25; passedTime=2.503; remTime=998.697; MLUPs=17.2699
[LatticeStatistics] step=300; t=0.1; uMax=5.75006e-07; avEnergy=8.66459e-16; avRho=1
...

```

Listing 6.13.: Terminal output of example bstep2d.

6.7. Read and Write STL Files

OpenLB offers the possibility to read and write geometry data in the Standard Triangulation Language, STL for short. The OpenLB class `STLReader` provides the desired functionality. In the case that the STL file you want to read is too large, you can use ParaView's filter "Decimate" to reduce the number of facets.

The constructor of the class `STLReader` takes two necessary and three optional arguments.

```
1 STLReader(const std::string fName, T voxelSize, T stlSize=1,  
2           unsigned short int method = 2, bool verbose = false);
```

- `fName`: The filename of the STL file to be read.
- `voxelSize`: The intended spatial step size for the simulation in SI units (m).
- `stlSize`: Conversion factor if the STL file is not given in SI units. E.g.: For an STL file in cm, this factor is `stlSize = 0.01`.
- `method`: Switch between methods for determining inside and outside of geometry.
 - default: fast, less stable
 - 1: slow, more stable (for untight STLs)
- `verbose`: Switch to get more output.

Functionality: The STL file is read and stored in the class `STLmesh`. A class `Octree` is instantiated of side-length $\text{rad} = 2^{j-1} \cdot \text{voxelSize}$, $j \in \mathbb{N}$ with j such that a cube with diameter 2rad covers the entire STL. Intersections of triangles and the nodes of the `Octree` are computed and an index of the respective triangles is stored in each node. A node is a leaf if either $\text{rad} = \text{voxelSize}$ or if it does not contain any triangles.

In a second step, it is determined whether a leaf is inside the STL geometry by one of the following methods:

- (Default) One ray in Z-direction is defined for each voxel in XY-layer. All nodes are indicated on the fly (faster, less stable).
- Define three rays (X-, Y-, Z-direction) for each leaf and count intersections with STL for each ray. Odd number of intersection means inside. The final state is decided by a majority vote (slower, more stable).

6.8. XML Parameter Files

In OpenLB essential simulation parameters can be placed in a XML file. This is a useful feature, since once a program is compiled, the parameters can be changed through the XML file and

recompilation is redundant. As a consequence whenever parameter fitting or general simulations are wanted, this approach can help you since only editing the XML file is necessary. The parsing is implemented in the header file `io/xmlReader.h`.

The general format for the XML files is:

```
<Param>
  <Output>
    <Log>
      <VerboseLog> true </VerboseLog>
    </Log>
  </Output>
  <VisualizationImages>
    <Filename> image </Filename>
  </VisualizationImages>
</Param>
```

All parameters need to be wrapped in a `<Param>` tag. To open a config file, you just pass a string with the file name to the class constructor of `XMLreader`.

```
1 std::string fName("demo.xml");
2 XMLreader config(fName);
3
4 bool _verboseLog;
5 std::string imagename;
6 XMLreader outputConfig = config["Output"];
7
8 config.readOrWarn<bool>("Output", "Log", "VerboseLog", _verboseLog);
9 outputConfig.readOrWarn<bool>("Log", "VerboseLog", "", _verboseLog);
10 config.readOrWarn<std::string>("VisualizationImages", "Filename", "", imagename);
```

First, an `XMLreader` object `config` is created. There are multiple ways to access the configuration data. To select the tag you would like to read, you just use an associative array like syntax as shown above.

To get a specific value out of an XML parameter file, there are multiple methods. One is to pass a predefined variable to the method `readOrWarn`, which reads the respective value and prints a warning in case the data type is not matching or the value cannot be found. For large subtrees with lots of parameters, you can also create a subobject. For this, you just have to reassign your selected subtree to a new `XMLreader`-object as is done above for `Output`.

6.9. Visualization with ParaView

As already mentioned, there are several data formats that can be used in ParaView. Use 'File – Open' and choose the set of data you want to use. In regards to OpenLB it is enough to open the file with the ending `.pvd`, since it contains a reference to the `.vti` files. The chosen files should now be part of the 'Pipeline Browser', which should be on the left hand side (if any of the panels are missing you can add them in the 'View' menu on the top). Click on 'Apply' in the 'Properties' panel (usually located below the 'Pipeline Browser') after opening.

Your data should now be visible in the center window. From within the 'Properties' or in one of the top tool bars, you can change the 'Coloring' properties, which selects what shall be displayed (e.g. physical velocity, phys pressure), which part of this choice shall be displayed (e.g. magnitude, x-value) and the way it is colored.

Make sure that '3D' is part of the tool bar directly above the window where you can see your objects. If you cannot find it click on '2D' which should be written instead and change it to '3D' by doing this. The commands for moving your whole set of visible objects and thus changing the perspective are the following:

- Using the mouse wheel, you can zoom in and out.
- Using the right mouse button or 'Ctrl + left mouse button', you can move the object to the background or the foreground. In comparison to zooming in and out, this changes the level of the 3D-effect.
- Using the left mouse button allows you to turn the object.
- Clicking the mouse wheel allows you to move the object center.

Of course you can also stick to '2D', although in this case the mouse commands might change a bit.

You can visualize the temporal development of your simulation using the 'Play' button and the related buttons directly next to it. If you want to go to a certain time step, use the input field 'Time', which is also located here.

To manipulate your data in ParaView, numerous so-called 'Filters' are provided in the 'Filters' menu in the top bar.

6.10. Clip

With this filter, you can cut off parts of your objects, for example, to make it possible to look inside the geometry. There are several tool options to determine which part is cut off. You can choose between plane, box and sphere.

If the "wrong" side is cut off, check 'inside out' to make the other side visible.

6.10.0.1. Contour

Using 'Contour' you can show lines or planes of certain data values, which you can set.

6.11. Glyph

If you have a point data set, you can represent it as spheres using the filter 'Glyph' and choosing 'Sphere' as setting for 'Glyph Type'. Using the resolution settings, you can smooth the surface to make the sphere look more rounded.

There are alternative ways to represent the data. As an example, arrows can be used to show the direction of a velocity. Check ‘Glyph Type’ for further possibilities.

6.12. Stream Tracer

The Stream Tracer filter is a powerful tool that allows users to visualize the flow as streamlines, making for example turbulent flow more visible. By placing the seed of the stream tracer next to the area of interest, users can apply this filter. There are two possible forms of the seed: it can be set as a pointcloud or a line.

To use the Stream Tracer filter, simply select the seed type that best fits your needs and place it in the desired location. In the case of a pointcloud seed, the flow that goes through the sphere is transformed, while a line seed transforms the flow along the line. An example of the OpenLB example `nozzle3d` visualized with streamlines using a pointcloud seed can be seen in Figure 6.3.




Figure 6.3.: OpenLB example `nozzle3d` visualized with streamlines. In this case as seed, a pointcloud was used.

6.13. Resample To Image

Another filter for visualization purposes is `Resample To Image`. It applies a volumetric ray-marching algorithm on the data. The Figure 6.4 shows the result when this filter is applied.

Another useful filter for visualization purposes is the `Resample To Image` filter. This filter applies a volumetric raymarching algorithm to the data, allowing users to visualize the data in a new way. To use this filter, simply select the `Resample To Image` option and apply it to your data. The resulting visualization will be a 3D image that can be rotated and explored in real-time. An example of the OpenLB example `nozzle3d` visualized using the `Resample To Image` filter can be seen in Figure 6.4.

This visualization provides a unique perspective on the data and can help users to better understand the underlying patterns and structures in the data. It's important to note that the



Figure 6.4.: OpenLB example `nozzle3d` visualized with the help of the `Resample To Image` filter

`Resample To Image` filter requires a significant amount of computational resources to generate the 3D image. As such, it may not be suitable for use on larger datasets or on less powerful hardware. Additionally, the parameters of the filter can be adjusted to fine-tune the resulting visualization, but this requires some knowledge of the filter and its underlying algorithm.

6.13.1. Application of Functors

The concept of functors benefits from generality and therefore, they are used for many applications.

6.13.1.1. Extract Simulation Data

Velocity, pressure and other information can be extracted from the lattice using predefined functors, see Listing 6.14. All they need to know is a `SuperLatticeXD` and an `UnitConverter` - if dimension or physical units are wanted.

```

1 // Create functors
2 SuperLatticePhysVelocity3D<T,DESCRIPTOR> velocity(sLattice, converter);
3 SuperLatticePhysPressure3D<T,DESCRIPTOR> pressure(sLattice, converter);

```

Listing 6.14.: Code example for calculating velocity and pressure using functors.

6.13.1.2. Define Analytic Functions

Often the inflow velocity has Poiseuille profile which is defined analytically, by means of a function. OpenLB provides analytic functors to define e.g. a Poiseuille velocity profile, random values, linear and constant values.

```
1 Poiseuille2D<T> poiseuilleU(superGeometry, 3, maxVelocity, distance2Wall);
```

Listing 6.15.: Define a Poiseuille velocity profile for inflow boundary condition.

6.13.1.3. Interpolation

Another case for interpolation functors is the conversion of a given analytical functor, such as an analytical solution to a `SuperLattice` functor. Afterwards, the difference can be easily calculated with the help of the functor arithmetic, see Listing 6.17. Finally, specific norms implemented as functors facilitate analysis of convergence.

```
1 // define a analytic functor: R^3 -> R
2 AnalyticalConst3D<double, double> constAna(1.0);
3 // get analytic functor on the lattice: N^3 -> R
4 SuperLatticeFfromAnalyticalF3D<double, DESCRIPTOR> constLat(constAna,
5                                                                lattice);
```

Listing 6.16.: Transition from an analytical functor to a lattice functor.

Application of this is shown in the example `poiseuille2d`, which is discussed in Section 8.4.4

6.13.1.4. Arithmetic and Advanced Functor Usage

Functors can be added, subtracted, etc. which is a very useful and elegant method to treat data. Listing 6.17 shows how to compute the relative error over the whole three dimensional domain.

```
1 int input[1];
2 double normAnaSol[1], absErr[1], relErr[1];
3 // define analytical solution: R^3 -> R
4 // for snake of simplicity it is a constant function,
5 // however it may be any specialization of AnalyticalF3D
6 AnalyticalConst3D<double, double> dSol(1.0);
7 // get analytical solution on the lattice: N^3 -> R
8 SuperLatticeFfromAnalyticalF3D<double, DESCRIPTOR> dSolLattice(dSol, lattice);
9 // get density out of simulation data
10 SuperLatticeDensity3D<T, DESCRIPTOR> d(lattice);
11 // compute absolute error
12 SuperL2Norm3D<double> dL2Norm(dSolLattice - d, superGeometry, 1);
13 // compute norm of solution
14 SuperL2Norm3D<double> dSolL2Norm(dSolLattice, superGeometry, 1);
15 dL2Norm(absErr, input); // access absolute error
16 dSolL2Norm(normAnaSol, input); // access norm of the solution
17 relErr[0] = absErr[0] / normAnaSol[0];
18 clout << "density-L2-error(abs)=" << absErr[0] << ";";
19 << "density-L2-error(rel)=" << relErr[0] << std::endl;
```

Listing 6.17.: Computation of a relative error with respect to L^2 -norm.

For more detail, see the source code of example 8.4.4.

Assemble geometry with geometric primitives of type `IndicatorFXD`.

```

1 Vector<double,2> extendChannel(lx0, ly0);
2 Vector<double,2> originChannel;
3 IndicatorCuboid2D<double> channel(extendChannel, originChannel);
4 // setup step
5 Vector<double,2> extendStep(lx1, ly1);
6 Vector<double,2> originStep;
7 IndicatorCuboid2D<double> step(extendStep, originStep);
8 // remove step from channel
9 IndicatorIdentity2D<double> channelIdent(channel-step);

```

Listing 6.18.: Deploy functor arithmetic to build geometry data.

6.13.1.5. Setting Boundary Value

Boundary cells are marked by a certain material number in the `SuperGeometryXD`. Using a functor, velocities can be set simultaneously on all cells of this material. First, a vector that characterizes the maximum flow velocity and its directions is necessary. Then, a special functor uses this vector to initialize a Poiseuille profile. The direction can be extracted in the case of axis-parallel inflow regions automatically from the `SuperGeometryXD`. In the last step, the `SuperLattice` initializes all cells of a certain material given by the `SuperLatticeXD` with the velocities computed by the functor.

```

1 // Creates and sets the Poiseuille inflow profile using functors
2 double maxVel = converter.getCharLatticeVelocity();
3 CirclePoiseuille3D<double> poiseuilleU(superGeometry, 3, maxVel, distance2Wall);
4 sLattice.defineU(superGeometry, 3, poiseuilleU);

```

Listing 6.19.: Code example for setting a Poiseuille velocity profile and a constant pressure boundary in `cylinder3d`.

6.13.1.6. Flux Functor

The *flux* of a quantity is defined as the rate at which this quantity passes through a fixed boundary per unit time. As a *mathematical concept*, flux is represented by the surface integral of a vector field

$$\Phi = \int \mathbf{F} \cdot d\mathbf{A}, \quad (6.1)$$

where \mathbf{F} is a vector field, and $d\mathbf{A}$ is an area element of the surface A , in the direction of the surface normal \mathbf{n} .

Flux functors calculate the discrete flux

$$\Phi_h = h^2 \sum_i \mathbf{f}_i \cdot \mathbf{n} , \quad (6.2)$$

with h as the grid length of the surface and \mathbf{f}_i the vector of the quantity at grid point i .

As the grid of the area has to be independent from the lattice, the value of \mathbf{f}_i will be interpolated from the surrounding lattice points. In the general case this discrete value is calculated by `SuperPlaneIntegralF3D`. Note that the reduction of the relevant surface is performed by `BlockReduction3D2D` and that `SuperPlaneIntegralF3D` adds only the multiplication by the area unit as well as the normal vector for multidimensional \mathbf{f}_i . In turn specific flux functors such as `SuperPlaneIntegralFluxVelocity3D` only add functor instantiation and print methods. So, for the `SuperPlaneIntegralF3D` functor a surface needs to be defined. OpenLB currently supports using subsets of hyperplanes as the surfaces on which to calculate a flux.

Such a *hyperplane* can be defined by an origin and two span vectors, an origin and a normal vector or a 3D circle indicator. `BlockReduction3D2D` interpolates the full intersection of hyperplane and mother geometry. Optionally this maximal plane may be further restricted by arbitrary 2D indicators.

Note that `SuperPlaneIntegralF3D` as well as all specific flux functors provide a variety of constructors accepting various hyperplane parametrizations. For full control you may consider explicitly constructing a `Hyperplane3D` instance.

The *discretization* of a hyperplane parametrization (given by `Hyperplane3D`) into a discrete lattice is performed by `HyperplaneLattice3D`.

Step 1: Define the hyperplane by

a) origin and two span vectors

```
1 Vector<T,3> origin;
2 Vector<T,3> u, v;
```

b) origin and normal vector

```
1 Vector<T,3> origin;
2 Vector<T,3> normal;
```

c) normal vector (centered in mother cuboid)

```
1 Vector<T,3> normal;
```

d) circle indicator

```
1 IndicatorCircle3D<T> circleIndicator(center, normal, radius);
```

e) arbitrary hyperplane

```
1 // example parametrization of a hyperplane centered in the mother cuboid and normal to
  the Z-axis
```

```

2 Hyperplane3D<T> hyperplane()
3   .centeredIn(cuboidGeometry.getMotherCuboid())
4   .normalTo({0, 0, 1});

```

Step 1.1 (optional): Define the hyperplane discretization by
a) *grid length*

```

1 T h = converter.getLatticeL();
2 HyperplaneLattice3D<T> hyperplaneLattice(
3   cuboidGeometry,
4   Hyperplane3D<T>().originAt(origin).spannedBy(u, v),
5   h);

```

b) *grid resolution*

```

1 HyperplaneLattice3D<T> hyperplaneLattice(
2   cuboidGeometry,
3   Hyperplane3D<T>().originAt(origin).spannedBy(u, v),
4   600); // resolution

```

Step 1.2 (optional): Define the flux-relevant lattice points by
a) *list of material numbers*

```

1 std::vector<int> materials = {1, 2, 3};

```

a) *arbitrary indicator*

```

1 SuperIndicatorF3D<T> integrationIndicator...

```

Step 1.3 (optional): Restrict the discretized intersection of hyperplane and geometry by
a) *2D circle indicator (relative to hyperplane origin)*

```

1 T radius = 1.0;
2 IndicatorCircle2D<T> subplaneIndicator({0,0}, radius);

```

a) *arbitrary 2D indicator (relative to hyperplane origin)*

```

1 IndicatorF2D<T> subplaneIndicator...

```

Step 2: Create a SuperF3D functor for
a) *velocity flow*

```

1 SuperLatticePhysVelocity3D<T,DESCRIPTOR> f(sLattice, converter);

```

b) *pressure*

```

1 SuperLatticePhysPressure3D<T,DESCRIPTOR> f(sLattice, converter);

```

c) *any other SuperF3D functor*

```
1 SuperF3D<T> f...
```

Step 3: Instantiate `SuperPlaneIntegralF3D` functor depending on how the hyperplane was defined and discretized.

a) using origin, two span vectors and materials list

```
1 SuperPlaneIntegralF3D<T> fluxF(  
2   f, superGeometry, origin, u, v, materials);
```

b) using origin, normal vector and materials list

```
1 SuperPlaneIntegralF3D<T> fluxF(  
2   f, superGeometry, origin, normal, materials);
```

c) using normal vector and materials list

```
1 SuperPlaneIntegralF3D<T> fluxF(f, superGeometry, normal, materials);
```

d) using 3D circle indicator and materials list

```
1 SuperPlaneIntegralF3D<T> fluxF(  
2   f, superGeometry, circleIndicator, materials);
```

e) using arbitrary hyperplane and integration point indicator

```
1 SuperPlaneIntegralF3D<T> fluxF(  
2   f, superGeometry, hyperplane, integrationIndicator);
```

f) using arbitrary hyperplane, integration point indicator and subplane indicator

```
1 SuperPlaneIntegralF3D<T> fluxF(f, superGeometry, hyperplane, integrationIndicator,  
   subplaneIndicator);
```

g) using arbitrary hyperplane lattice, integration point indicator and subplane indicator

```
1 SuperPlaneIntegralF3D<T> fluxF(f, superGeometry, hyperplaneLattice,  
   integrationIndicator, subplaneIndicator);
```

Step 4: Get results using `operator()`

```
1 int input[1]; // irrelevant  
2 T output[5];  
3 fluxF(output, input);
```

- `output[0]`: flow rate or plane integral (if quantity has dimension 1)
- `output[1]`: size of the area
- `output[2..4]`: flow vector (ie. vector of summed quantities)

In many cases the functor argument is either the velocity or the pressure functor.

Thus **Step 2** and **Step 3** may be combined using `SuperPlaneIntegralFluxVelocity3D` respectively `SuperPlaneIntegralFluxPressure3D`. Their constructors are mostly identical to the ones provided by `SuperPlaneIntegralF3D`. In fact the only difference is that the first functor argument is replaced by references to `SuperLattice` and `UnitConverter`.

Step 2.1): Combined steps for velocity flux

```
1 SuperPlaneIntegralFluxVelocity3D<T> vFlux(superLattice, converter, ...);
```

Step 2.2): Combined steps for pressure flux

```
1 SuperPlaneIntegralFluxPressure3D<T> pFlux(superLattice, converter, ...);
```

Step 3.1): Output region size, volumetric flow rate and mean velocity

```
1 vFlux.print(std::string regionName,
2   std::string fluxSiScaleName, std::string meanSiScaleName);
```

- `fluxSiScaleName`: 'ml/s' or 'l/s' or '' (default= m^3/s)
- `meanSiScaleName`: 'mm/s' or '' (default= m/s)

Step 3.2): Output region size, force and mean pressure

```
1 pFlux.print(std::string regionName,
2   std::string fluxSiScaleName, std::string meanSiScaleName);
```

- `fluxSiScaleName`: 'MN' or 'kN' or '' (default= N)
- `meanSiScaleName`: 'mmHg' or '' (default= Pa)

6.13.1.7. Discrete Flux Functor

If a hyperplane is axis-aligned, flux functors may optionally be used in *discrete* mode. Passing `BlockDataReductionMode::Discrete` as the last argument to any plane integral or flux constructor instructs the internal `BlockReduction3D2D` instance to reduce the hyperplane by evaluating the underlying functor at the nearest lattice points instead of by interpolating physical positions.

Note that this imposes restrictions on the accepted hyperplane and its lattice:

- `Hyperplane3D` normal must be orthogonal to a pair of unit vectors
- `HyperplaneLattice3D` spacing must equal the distance between lattice nodes

The restriction on the hyperplane lattice spacing is fulfilled implicitly when automatic lattice parametrization is used.

```

1 // discrete flux usage in examples/aorta3d
2 SuperPlaneIntegralFluxVelocity3D<T> vFluxInflow( sLattice, converter, superGeometry,
    inflow, materials, BlockDataReductionMode::Discrete );

```

6.13.1.8. Wall Shear Stress Functor

The *Wall Shear Stress* is defined as the parallel force per unit area exerted by a fluid on a wall. In the context of macroscopic fluid mechanics the *Wall Shear Stress* of a Newtonian fluid is given by

$$\tau_W = \mu \left. \frac{\partial \mathbf{u}}{\partial y} \right|_{y=0}, \quad (6.3)$$

where μ is the dynamic viscosity, \mathbf{u} is the velocity field and y the coordinate perpendicular to the wall. The *Wall Shear Stress Functor* calculates the discrete Wall Shear Stress

$$\tau_W = \boldsymbol{\sigma} \cdot \mathbf{n} - ((\boldsymbol{\sigma} \cdot \mathbf{n}) \cdot \mathbf{n}) \cdot \mathbf{n}, \quad (6.4)$$

where $\boldsymbol{\sigma}$ is the Cauchy stress tensor and \mathbf{n} the local unit normal vector of the surface. Since the lattice stress tensor $\boldsymbol{\Pi}$ is not defined on boundary cells, it is read out from an adjacent fluid cell in a discrete velocity direction associated with each boundary cell. The unit normal vector is obtained by a given `IndicatorF3D` instance, which is slightly increased in size. See `examples/poiseuille3d` for usage details. Due to the staircase approximation of the boundary, the wall shear stress calculation is first order accurate.




Figure 6.5.: Relative error of Wall Shear Stress L_p norm in `poiseuille3d`.

6.13.2. Error Norm Functors

While relative and absolute error norms may be calculated manually using functor arithmetic (see 6.13.1.4), they are also available as distinct functors. As such it is preferable to utilize `SuperRelativeErrorLpNormXD` and `SuperAbsoluteErrorLpNormXD` if one uses the common definition of relative and absolute error norms.

Let `wantedF` be the simulated solution functor and `f` the analytical solution, then

$$\text{SuperRelativeErrorLpNormXD implements } \frac{\| \text{wantedF} - f \|_p}{\| \text{wantedF} \|_p}, \quad (6.5)$$

$$\text{SuperAbsoluteErrorLpNormXD implements } \| \text{wantedF} - f \|_p. \quad (6.6)$$

An example of how to use these error norm functors in practice is given by the Poiseuille flow example as described in section 8.4.4.

```
1 Poiseuille2D<T> uSol(axisPoint, axisDirection, maxVelocity, radius);
2 SuperLatticePhysVelocity2D<T,DESCRIPTOR> u(sLattice, converter);
3 auto indicatorF = superGeometry.getMaterialIndicator(1);
4
5 SuperAbsoluteErrorL1Norm2D<T> absVelErrorNormL1(u, uSol, indicatorF);
6 absVelErrorNormL1(result, tmp);
7 clout << "velocity-L1-error(abs)=" << result[0];
8 SuperRelativeErrorL1Norm2D<T> relVelErrorNormL1(u, uSol, indicatorF);
9 relVelErrorNormL1(result, tmp);
10 clout << "; velocity-L1-error(rel)=" << result[0] << std::endl;
```

Listing 6.20.: L1 velocity error in `poiseuille2d`

Further implementation details are touched upon in section 2.5.2.1.

6.13.2.1. Grid Refinement Metric Functors

`SuperLatticeRefinementMetricKnudsen(2,3)D` implements an automatic block-level grid refinement criterion as described by Lagrava et al. in “Automatic grid refinement criterion for lattice Boltzmann method” [102]. This criterion uses the quality of the cell-local Knudsen number approximation as measured by `SuperLatticeKnudsen*D` to judge the adequacy of the block resolution.

7. Flow Control and Optimization

In almost every application, the optimization of a process regarding certain objectives is of general interest. In optimal flow control we are interested in optimizing the flow problem regarding a specific need, i.e. to find optimal flow parameters which lead to the desired flow properties. OpenLB provides a framework to solve such optimization problems with iterative, gradient-based methods. Basically, a flow simulation is encapsulated into the optimization framework which performs several flow simulations. In each step the input parameters are varied in order to approach the optimal solution. This chapter explains first how these optimal solutions are found and second how this is realized within OpenLB. Therefore, this section is structured as follows. Section 7.1 gives a brief overview of the theoretical optimization concepts and Section 7.2 deals with their implementation in OpenLB. Note that, in this chapter, specific notation is introduced where necessary and might not conform to the rest of the user guide. For examples, tensors and scalars are mostly printed in similar fonts below.

7.1. Introduction

This chapter is dedicated to the solution of optimization problems of the general form

$$\text{Find } \alpha, \text{ s.t. } J(f, \alpha) \text{ is minimized while } G(f, \alpha) = 0. \quad (7.1)$$

In this context, $\alpha \in \mathbb{R}^d$ is called the *control variable*, f is the state, J is the *objective functional* and G is the side condition. For simplicity, we assume in our code that the side condition $G(f, \alpha) = 0$ yields a unique solution f for any admissible control α , s.t. we have $f = f(\alpha)$ (unrestricted optimization approach).

Possible optimization problems are:

- Finding the minimum of an analytic function J w.r.t. an argument vector α , e.g. as a post-processing step.
- Optimization of flow simulation parameters: e.g., the performance of a mixer could be optimized w.r.t. simulation parameters. In this case, α contains the free simulation parameters, G are the governing fluid flow equations (mathematical model), f is the solution of the governing equations for a given α (e.g. the concentration field) and J measures the mixing quality.

- Parameter identification: sometimes, physical simulation parameters are not known, but with the help of additional data (e.g. NMR-data of the fluid velocity), we can reconstruct them: Here, we set α to be the unknown parameters, G are the governing equations, f is their solution and J measures the difference between simulated and "true" data (e.g. $J = \frac{1}{2} \|\mathbf{u}(f(\alpha)) - \mathbf{u}_{NMR}\|^2$). Using optimization, we can then find the unknown parameter α with the least difference to the "true" data.

The optimization framework provided in OpenLB is rather designed for large-scale applications like fluid flow simulations, where a time critical step (regarding computation time) is the function evaluation and not the optimization routine. For an application-oriented introduction, we also refer to the example `showcaseRosenbrock`.

7.1.1. Iterative Approach – Line Search Algorithm

The optimization methods in the OpenLB framework employ an iterative, gradient-based approach based on the line search concept. Starting with an initial guess for the control variables α_k , the next set of control variables α_{k+1} are found by proceeding along a direction d with a step size s . With each step the objective function should be minimized, i.e.

$$\alpha_{k+1} = \alpha_k + d s, \quad \text{with} \quad J(\alpha_{k+1}) < J(\alpha_k). \quad (7.2)$$

After each iteration step, either the objective function J or its derivatives regarding the current control variables $\frac{\partial J}{\partial \alpha}$ are evaluated. If they are below a user-defined threshold, the algorithm exits. The open questions at this point are listed together with their corresponding subsections giving answers to those questions:

1. How to choose the (descent) direction? (7.1.2)
2. How to choose the step size? (7.1.3)
3. How can we compute the derivatives? (7.1.4)

7.1.2. Descent Algorithms

The descent algorithm computes the descent direction d for our next iteration step, i.e. the direction where we likely approach the optimal solution. One possible approach is to use the local gradient of the objective function $\frac{\partial J}{\partial \alpha}$ as the descent direction. This method is referred to as the steepest descent approach. Since only first order derivatives are considered here, this method is relatively slow (it requires more iteration steps) compared to Newton or Quasi-Newton methods while its main advantage is its stability. Quasi-Newton methods such as LBFGS achieve higher convergence orders since these additionally include the approximated second derivatives. In the OpenLB framework, the steepest descent and LBFGS algorithm are provided for computing the descent direction.

7.1.3. Step Size Algorithms

The choice of suitable step sizes is a nontrivial task; if the step sizes are too large we may overshoot and miss the optimum. On the other side, if the steps are too small we will need a lot of iterations until the optimal solution is reached, resulting in longer computation duration. To find the optimal step size, step conditions like Armijo, normal or strong Wolfe-Powell rules can be applied.

The Armijo rule prevents choosing too large step sizes by checking if we achieve a certain minimal decrease in our objective function. The objective function in the next iteration step $k + 1$ can be approximated as

$$J(\alpha_{k+1}) \approx J(\alpha_k) + s_k \nabla J(\alpha_k)^T d, \quad (7.3)$$

where s_k is the current step size. First, relatively large step sizes are chosen and the objective function is evaluated at α_{k+1} . If the objective function is decreased sufficiently in comparison to the expected value for the objective function, the algorithm exits with the current step size; if not, a smaller step size is chosen. This condition can be written as

$$J(\alpha_{k+1}) = J(\alpha_k + s_k d) \leq J(\alpha_k) + \rho s_k \nabla J(\alpha_k)^T d \quad \text{with } \rho \in [0, 1]. \quad (7.4)$$

These steps are repeated until this condition is fulfilled. To add further flexibility a scalar ρ is introduced here allowing us to define what portion of the expected linear decrease we want to assure by adjusting the threshold. Note that the Armijo rule requires an evaluation of the objective function for each iteration step.

The Wolfe-Powell rule can be classified into the normal and strong formulation. Both contain the Armijo rule and consider an additional criterion which prevents us from choosing step sizes which are too small. For this, the directional derivative (scalar product of the proceeding direction and the derivative) on the old iteration is compared to the directional derivative on the new iteration point. Assuming that we are at the optimal solution, the directional derivative is zero. During our line search we are descending regarding the objective function which means that our directional derivative is always a negative value. If we move closer towards the optimum, the directional derivative must increase to get eventually to zero. Thus, we can check

$$\nabla J(\alpha_k + s_k d)^T d \geq \delta \nabla J(\alpha_k)^T d, \quad (7.5)$$

where δ is a tuning parameter. This is referred to as the normal Wolfe-Powell formulation. The strong formulation reads

$$|\nabla J(\alpha_k + s_k d)^T d| \leq -\delta \nabla J(\alpha_k)^T d. \quad (7.6)$$

The step size is varied until both the Armijo rule and the second condition of either the nor-

mal or strong Wolfe-Powell formulation are satisfied (or by exceeding the maximal number of attempts). Note that the Wolfe-Powell rule requires for each iteration step an evaluation of the objective function and its derivative at that position.

7.1.4. Derivative Computation

For almost all algorithms the evaluation of function gradients are necessary. For those there are also several approaches such as finite-difference schemes and algorithmic differencing. The following options are implemented in OpenLB:

- Usage of forward/ central difference quotients: evaluate J for neighboring values of α and compute the difference quotients. This is the simplest method and for a small number of control variables it is fast. However, it is the least accurate method.
- Forward automatic differentiation: evaluate J for operator-overloaded variables of type `ADf<T, n>`, where `T` is the underlying arithmetic type and `n` is the number of control variables. This method is a little slower than difference quotients, but it usually returns derivatives at full machine precision.
- Adjoint LBM: adjoint Lattice-Boltzmann equations are used, cf. [19, 21]. This method is perfect for distributed control problems since the computational expense remains constant for any number of control variables. However, this method requires the (theoretical) calculation of adjoint LB-equations which is a problem-specific task.

7.2. Implementation

The implementation is separated into two key steps: the `OptiCase` classes define how the objective functional J and its gradient are computed as functions of α . The `Optimizer` classes define optimization methods such as steepest descent or step size algorithms (independent of the question, how the function and gradient evaluations are performed).

7.2.1. Optimizer Classes: Optimization Methods

The common methods steepest descent, LBFGS and Barzilai–Borwein are implemented as children classes of the `Optimizer` classes. Their various free parameters (e.g. an upper limit for the number of optimization steps) can be passed via the constructor or read from an `xml` file. In many situations, the control variables have to remain in a certain range in order to guarantee physical meaningfulness and a stable simulation. E.g. for porosity optimization, the porosity has to lie between 0 and 1 in order to be physically meaningful. This can be achieved via bounding the control by user-defined thresholds (typically used for parameter optimization) or

by using a smooth projection map $p: \mathbb{R} \rightarrow [0, 1]$ (helpful for porosity optimization). Alternatively, if one has more than one control variables, a vector with the upper and lower bounds for each variable can be used to bound the control.

7.2.2. OptiCase Classes: Gradient Computation

For the numerical evaluation of the functional gradient $\frac{dJ}{d\alpha}$, the following four options are implemented:

- Forward/ central difference quotients. The method is implemented in the classes `OptiCaseFDQ` and `OptiCaseCDQ`, where the user (only) has to provide an expression for the evaluation of J at construction.
- Forward automatic differentiation. Evaluate J for operator-overloaded variables of type `ADf<T, n>`, where T is the underlying arithmetic type and n is the number of control variables. Therefore, the full source code has to be templated w.r.t. the arithmetic type T and one has to take care that the arithmetic operations are differentiable (which is both satisfied by the OpenLB code basis). One then passes two instances of J (one with type T , one with type `ADf<T, n>`) to the class `OptiCaseAD`, which then does the rest.
- Adjoint LBM. This method requires the (theoretical) calculation of adjoint LB-equations which is a problem-specific task. Because of that, the implementing class `OptiCaseDual` poses direct assumptions on the implementation of the LB simulation (e.g., it requires usage of the Solver framework and has so far only been implemented in the context of porosity and force optimization, cf. examples `DomainIdentification3d` and `TestFlowOpti3d`).

For difference quotients and forward automatic differentiation, any functions of type T (`const std::vector&`) (with or without LB simulation) are passed. Flow simulations have to be encapsulated by a suitable wrapper that accepts the control variables, runs the simulation and computes the objective. In the Solver app structure, the `getCallable` method fulfills this task. An introduction into forward automatic differentiation, whose usage is not restricted to the optimization context, is provided in the example `showcaseADf`.

7.2.3. Parameter Explanation and Reading from XML

In this short overview the relevant parameters for an app with optimization are listed with the respective names for using an `xml`-file for the input. The different parameters of each part in the `xml`-file are explained via a table of the following form:

Parameter	Name(type)	[declaration & definition] (&& and other usages)	Default value	Explanation: (if available, all) <i>possibilities</i>
-----------	------------	---	------------------	---

The explanation of each column is as follows:

- **Parameter:** Name of the parameter in the `xml`-file
- **Name (type):** Name of the parameter in the source code and its data type in brackets. Besides the common data types the abbreviations S and T are used for template parameters.
- **[declaration & definition]** (&& other usages): Location of declaration and definition (&& and sometimes some other important usages) of the parameter, e.g.:
`[solver.h & solver.hh] (&& example.cpp)`
 If the location of the declaration and definition is the same, only one location is indicated, e.g.
`[solver.hh]`
 If there is more than one location for the declaration and definition, it is indicated with an *and*, e.g.
`[solver.h & solver.hh] and [optiCaseDual.hh]`
 If there are different possibilities for the location of declaration and definition, it is indicated with curved brackets () and an *or*, e.g.
`([optimizerSteepestDescent.h] or [optimizerLBFGS.h] or [optimizerBarzilaiBorwein.h])`
- **Default value:** Is this parameter essential or optional?
 If a parameter is optional, it does not need to be defined. Then, the default value can be seen in this column.
 If a parameter is of such importance that without it the program has to exit, it is labeled as *EXIT*.
 Some parameters are indicated with **unused** which means, that the parameter is read but not used afterwards.
- **Explanation:** (if available, all) *possibilities*: Brief description and explanation of the parameter. Some parameters have different possibilities for their definition. In this case, all available possibilities are also offered in **bold** type letters, e.g.:
ad for *OptiCaseAD* or
dual for *OptiCaseDual* or
adTest for *OptiCaseADTest*

The arrangement of the parameters in the `xml`-file has the following structure:
example:

```
<Param>
  <Optimization>
    <MaxStepAttempts> 20 </MaxStepAttempts>
  </Optimization>
</Param>
```


Parameter	Name(type)	[declaration & definition] (&& and other usages)	Default value	Explanation: (if available, all) <i>possibilities</i>
<i>ControlType</i>	<code>_controlType</code> (std::string)	[optiCaseDual.h & optiCaseDual.hh]		Defines the control type to optimize <i>force</i> or <i>porosity</i>
<i>Control-Material</i>	<code>_control-Material</code> (int)	[optiSolverParameters.h] (&& optiCaseDual.h and optiCaseDual.hh)	0	Defines the number of control material, here 6 for porosity optimization problems
<i>Field-Dimension</i>	<code>_fieldDim</code> (int)	[optiSolverParameters.h] (&& optiCaseDual.h and optiCaseDual.hh)	0	Spatial dimension of controlled field
<i>Dimension-Control</i>	<code>_dimCtrl</code> (int)	(([optiCaseDual.h & optiCaseDual.hh]) and ([optimizer.h & optimizer.hh]) ([optimizerSteepestDescent.h] or [optimizerLBFGS.h] or [optimizerBarzilaiBorwein.h]) and [controller.h])	EXIT	Upper limit for number of control variables, so far not read by xml
<i>RegAlpha</i>	<code>regAlpha</code> (S)	[optiCaseDual.h & optiCaseDual.hh]	0	Weighting factor for regularization, so far unused
<i>Lambda</i>	<code>_lambda</code> (T)	(([optimizerSteepestDescent.h] or [optimizerLBFGS.h] or [optimizerBarzilaiBorwein.h]) (&& optimizerLineSearch.h))	1	Determines the initial step length lambda
<i>MaxIter</i>	<code>_maxIt</code> (int)	[optimizer.h & (optimizerSteepestDescent.h or optimizerLBFGS.h or optimizerBarzilaiBorwein.h)]	100	Maximal number of iterations of the optimizer

<i>MaxStep-Attempts</i>	<code>_maxStepAttempts</code> (int)	[optimizer.h & (optimizerSteepestDescent.h or optimizerLBFGS.h or optimizerBarzilaiBorwein.h)]	20	Maximal number of attempts at each optimization step for finding a suitable step size
<i>FailOnMaxIter</i>	<code>_failOnMaxIter</code> (bool)	[optimizer.h] (&& optimizer.hh or optimizerLBFGS.h or optimizerLineSearch.h)	1 (true)	if true, the optimization fails when reaching <code>_maxIt</code> and prints the warning: Optimization problem failed to converge within specified iteration limit of <code>_maxIt</code> iterations with tolerance of <code>_eps</code>
<i>Tolerance</i>	<code>_eps</code> (T)	[optimizer.h] (&& optimizer.hh)	1e-10	Tolerance of the optimizer. Optimizer stops if the norm of the vector of derivatives of the object functional is smaller than the tolerance of the optimizer
<i>L</i>	<code>_l</code> (int)	[optimizerLBFGS.h]	20	Maximal number of stored iteration steps for LBFGS algorithm.
<i>Verbose</i>	<code>_verboseOn</code> (bool)	[optimizer.h & optimizer.hh] (&& optimizerLineSearch.h)	1 (true)	Print Warnings and further output in the terminal, if true
<i>InputFile-Name</i>	<code>fname</code> (std::string)	((optimizerSteepestDescent.h] or [optimizerLBFGS.h] or [optimizerBarzilaiBorwein.h]) (&& optimizer.h and optimizerLineSearch.hh)	"control.dat"	Name of the file that contains the initial guess for the control values

<i>ControlTol- erance</i>	<code>_controlEps</code> (T)	<code>([optimizer.h & optim izerSteepestDescent.h and optimizerLBFGS. h and optim izerBarzilaiBorwein. h] (&& optimizer.hh)</code>	0	Optimization stops if the change of the control variables is less than this tolerance.
<i>StepCondi- tion</i>	<code>stepCondition</code> (std::string)	<code>([optimizerSteepest- Descent.h] or [optimizerLBFGS.h] or [optimizerBarzilai- Borwein.h]) (&& optimizerLineSearch.h)</code>	"Armijo" (SD), "Strong -Wolfe" (LBFGS), "None", (Barzilai- Borwein)	Defines the step condi- tion: None or Smaller or Armijo or Wolfe or StrongWolfe
<i>Vector- Bounds</i>	<code>_vectorBounds</code> (bool)	<code>[optimizer.h & optim izerLBFGS.h or optim izerBarzilaiBorwein. h] (&& optimizer.h and optimizerLineSearch. h)</code>	0 (false)	Determines, whether bounds on the control variables are applied component-wise
<i>Projection</i>	<code>_projection- Name</code> (std::string)	<code>[optiCaseDual.h & optiCaseDual.hh]</code>		mapping method be- tween the physical and computational control variables: Sigmoid or Rectifier or Softplus or Baron or Krause or Foerster or FoersterN or StasiusN or
<i>Reference- Solution</i>	<code>_compute- Reference</code> (bool)	<code>[optiCaseDual.h & optiCaseDual.hh]</code>	false	states if reference solu- tion is available
<i>StartValue</i>	<code>startValue</code> (T)	<code>(&& optimizer.h and optimizer.hh)</code>	0	Determines the initial guess for the optimiza- tion algorithm

<i>Upper-Bound</i>	<code>_upperBound</code> (T)	<code>([optimizer.h & optimizerSteepestDescent.h and optimizerLBFGS.h and optimizerBarzilaiBorwein.h] && optimizer.hh)</code>	false in all three optimizers	Sets an upper bound for the control values during the optimization process
<i>Lower-Bound</i>	<code>_lowerBound</code> (T)	<code>([optimizer.h & optimizerSteepestDescent.h and optimizerLBFGS.h and optimizerBarzilaiBorwein.h] && optimizer.hh)</code>	false in all three optimizers	Sets a lower bound for the control values during the optimization process
<i>Volume-Ratio</i>	<code>volumeRatio</code> (T)	<code>[optiCaseDual.hh]</code>	unused (//)	Determines the volume ratio

<i>VisualizationGnuplot:</i>				
<i>Visualized-Parameters</i>	gplotAnalysis-String (std::string)	[optimizerSteepestDescent.h] or [optimizerLBFGS.h] or [optimizerBarzilaiBorwein.h]		<p>Lists the parameters that will be plotted by Gnuplot during the optimization process. Possibilities:</p> <p>VALUE CONTROL DERIVATIVE ERROR NORM_DERIVATIVE</p>

8. Examples

8.1. Example Overview

A list of the currently included examples is given together with related keywords in the following Tables [8.1](#) and [8.2](#).

folder	example	turbulent	thermal	multi Component	multi Phase	particles	porous Media	transient flow	benchmark	showcase	STL geometry	geometry primitives	check- pointing
adsorption	adsorption3D												
	microMixer3D												
advectionDiffusion Reaction	advectionDiffusionReaction2d												
	reactionFiniteDifferences2d												
	advectionDiffusion1d												
	advectionDiffusion2d												
	advectionDiffusion3d												
	advectionDiffusionPipe2d												
laminar	bstep2d												
	bstep3d												
	cavity2d												
	cavity3d												
	cylinder2d												
	cylinder3d												
	poiseuille2d												
	poiseuille2dEOC												
	poiseuille3d												
	powerLaw2d												
	testFlow3dSolver												
multiComponent	binaryShearFlow2d												
	contactAngle2d												
	contactAngle3d												
	fourRollMill2d												
	microFluidics2d												
	phaseSeperation2d												
	phaseSeperation3d												
	rayleighTaylor2d												
	rayleighTaylor3d												
	youngLaplace2d												
	youngLaplace3d												



example includes relevant subject



example includes relevant subject and is recommended for beginning

Table 8.1.: Currently included examples in OpenLB (continued in Table 8.2).

folder	example	turbulent	thermal	multi Component	multi Phase	particles	porous Media	transient flow	benchmark	showcase	STL geometry	geometry primitives	check- pointing
optimization	domainIdentification3d												
	parameterIdentificationPoiseuille2d												
	showcaseADf												
	showcaseRosenbrock												
	testFlowOpti3d												
particles	bifurcation3d												
	dkt2d												
	magneticParticles3d												
	settlingCube3d												
porousMedia	porousPoiseuille2d												
	porousPoiseuille3d												
thermal	galliumMelting2d												
	porousPlate2d												
	porousPlate3d												
	porousPlate3dSolver												
	rayleighBernard2d												
	rayleighBernard3d												
	squareCavity2d												
	squareCavity3d												
	stefanMelting2d												
turbulent	aorta3d												
	channel3d												
	nozzle3d												
	tgV3d												
	venturi3d												
freeSurface	breakingDam2d												
	breakingDam3d												
	deepFallingDrop2d												
	fallingDrop2d												
	fallingDrop3d												
	rayleighInstability3d												



example includes relevant subject



example includes relevant subject and is recommended for beginning

Table 8.2.: Currently included examples in OpenLB (continuation of Table 8.1).

All the demo codes can be compiled with or without MPI, with or without OpenMP, and executed in serial or parallel.

8.2. adsorption

8.2.1. adsorption3D

This example shows the adsorption in a batch reactor using an Euler–Euler approach. The model is based on the linear driving force model and uses advection diffusion reaction lattices for particles, solute and particle loading. Different isotherms and mass transfer models can be used. An analytical solution is implemented when using the linear isotherm and surface diffusion.

8.2.2. microMixer3D

This example portrays the adsorption in a static mixing reactor using an Euler–Euler approach. Analogue to the example before, the model is based on the linear driving force model and uses advection diffusion reaction lattices for particles, solute and particle loading. Different isotherms and mass transfer models can be used.

8.3. advectionDiffusionReaction

8.3.1. advectionDiffusionReaction2d

This example illustrates a steady-state chemical reaction in a plug flow reactor. One can choose two types of reaction, $A \rightarrow C$ and $A \leftrightarrow C$. The concentration and analytical solution along the centerline of the rectangle domain is given in `./tmp/N<resolution>/gnuplotData` as well as the error plot for the concentration along the centerline. The default configuration executes three simulation runs and the average L^2 -error over the centerline is computed for each resolution. A plot of the resulting experimental order of convergence is provided in `./tmp/gnuplotData/`.

8.3.2. reactionFiniteDifferences2d

Similarly to the previous example, a simplified domain with no fluid motion and homogeneous species concentrations is simulated, but here with finite differences. The chemical reaction $|a|A \rightarrow |b|B$ is approximated, where the reaction rate $\nu = A/t_0$ is given, where t_0 is a time conversion factor. The initial conditions are set to $A(t=0) = 1$ and $B(t=0) = 0$ such that an analytical solution is possible via

$$A(t) = \exp\left(-\frac{|a|t}{t_0}\right), \quad (8.1)$$

$$B(t) = \left|\frac{b}{a}\right| \left[1 - \exp\left(-\frac{|a|t}{t_0}\right)\right]. \quad (8.2)$$

By default, the executable produces a plot in `./tmp/gnuplotData/` which is given in Figure 8.1 below.




Figure 8.1.: Solutions to concentration profiles in reactionFiniteDifferences2d.

8.3.3. advectionDiffusion1d

The `advectionDiffusion1d` example showcases second order mesh convergence of LBM for scalar linear one-dimensional advection–diffusion equations [37] of the form

$$\partial_t \chi + \partial_x F(\chi) - \mu \partial_{xx} \chi = 0, \quad (8.3)$$

where $\chi : \mathcal{X} \times \mathcal{I} \rightarrow \mathbb{R}$ is the conservative variable dependent on space $x \in \mathcal{X} \subseteq \mathbb{R}$ and time $t \in \mathcal{I} \subseteq \mathbb{R}_0^+$, $F \equiv u\chi$ is a linear function defined by the advection velocity $u \in \mathbb{R}$, and $\mu > 0$ denotes the diffusion coefficient. Hence, the LBM approximates the transport of the conservative variable χ along a one-dimensional line with periodic boundary conditions on $\mathcal{X} = [-1, 1]$. The initial pulse is defined by a sine profile, which is subsequently diffused and advected. An analytical solution at point x and time t is given by

$$\chi^*(x, t) = \sin[\pi(x - ut)] \exp(-\mu\pi^2 t). \quad (8.4)$$

In practice the simulation uses a two-dimensional square domain which is evaluated along a centerline to obtain the desired one-dimensional result. The domain is initialized with $\chi^*(x, t = 0)$.

Diffusive scaling is applied which results in the input parameters listed in Table 8.3. In the default setting, `advectionDiffusion1d` executes three simulation runs with increasing resolutions $N = 50, 100, 200$, respectively. Each simulation recovers $\mu = 1.5$ and a Péclet number of $Pe = 40/3$.

diffusive scaling $\triangle t = \triangle x^2$ for $Pe = 40/3$		
N	u_L	$\triangle x$
50	0.4	0.04
100	0.2	0.02
200	0.1	0.01

Table 8.3.: Default simulation parameters of `advectionDiffusion1d` with $\mu = 1.5$

The output of each simulation run is stored in the `tmp/N<number>` directory. At each simulation time step the average L2 relative error over the centerline is computed. Said average is then stored within the respective resolution directory `./gnuplotData/data/averageL2RelError.dat`. Additionally, the program averages the values in `averageL2RelError.dat` for each simulation run, which in turn is written to the global error file `tmp/gnuplotData/data/averageSimL2RelErr.dat`. For post-processing, a python3 script can be executed via

```
python3 advectionDiffusion1dPlot.py
```

The script requires the `matplotlib` python package which can be installed on any platform by issuing the following commands in a terminal:

```
python3 -m pip install -U pip
python3 -m pip install -U matplotlib
```

The script generates basic error plots for every file with the file extension `.dat` in `./tmp`. Finally, a global log-log error plot with reference curves is extracted from the data contained in `averageSimL2RelErr.dat`.

8.3.4. advectionDiffusion2d

The example `advectionDiffusion2d` acts as a mesh-convergence test for a solution to the scalar linear *two-dimensional* advection–diffusion equation

$$\partial_t \chi + \nabla_{\mathbf{x}} \mathbf{F}(\chi) - \mu \Delta_{\mathbf{x}} \chi = 0, \quad (8.5)$$

where $\chi : \mathcal{X} \times \mathcal{I} \rightarrow \mathbb{R}$ is the conservative variable dependent on space $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^2$ and time $t \in \mathcal{I} \subseteq \mathbb{R}_0^+$, $\mathbf{F} \equiv \mathbf{u}\chi$ is a linear function defined by the advection velocity $\mathbf{u} = (u_x, u_y)^T \in \mathbb{R}^2$, and $\mu > 0$ denotes the diffusion coefficient. Similarly, the analytical solution is given for any point $\mathbf{x} = (x, y)^T$ and time t as

$$\chi^*(x, y, t) = \sin[\pi(x - u_x t)] \sin[\pi(y - u_y t)] \exp(-2\mu\pi^2 t). \quad (8.6)$$

The simulation is executed on a square $\mathcal{X} = [-1, 1]^2$ which is periodically embedded in \mathbb{R}^2 . An error norm over the domain measures the deviation from the analytical solution up to the time step at which the initial pulse is diffused below 10%. For the default setting ($\mu = 0.05$ and $Pe = 100$), the outputs of three subsequent simulation runs are stored in a subfolder structure in `./tmp` and directly post-processed for visualization. A sequence of contour plots is generated with the highest computed resolution $N = 200$ and contained in `./tmp/N200/imageData`. Note that via issuing the command

```
python3 advectionDiffusion2dPlot.py
```

an error plot can be produced, which numerically validates the second order convergence in space.

8.3.5. advectionDiffusion3d

The example `advectionDiffusion3d` acts as a mesh-convergence test for a numerical solution to initial value problem

$$\begin{cases} \partial_t \chi(\mathbf{x}, t) + \nabla_{\mathbf{x}} \mathbf{F}(\chi(\mathbf{x}, t)) - \mu \\ bm \Delta_{\mathbf{x}} \chi(\mathbf{x}, t) = 0 & \text{in } \mathcal{X} \times \mathcal{I}, \\ \chi(\mathbf{x}, 0) \equiv \chi_0(\mathbf{x}) & \text{in } \mathcal{X}, \end{cases} \quad (8.7)$$

where $\chi : \mathcal{X} \times \mathcal{I} \rightarrow \mathbb{R}$ is the conservative variable dependent on space $\mathbf{x} \in \mathcal{X} \subseteq \mathbb{R}^3$ and time $t \in \mathcal{I} \subseteq \mathbb{R}_0^+$, $\mathbf{F} \equiv \mathbf{u}\chi$ is a linear function defined by the advection velocity $\mathbf{u} = (u_x, u_y, u_z)^T \in \mathbb{R}^3$, and $\mu > 0$ denotes the diffusion coefficient. Note that the domain $\mathcal{X} = [-1, 1]^3$ is periodic.

The example implements a smooth initial profile $\chi_0^s(\mathbf{x})$ and an unsmooth version $\chi_0^u(\mathbf{x})$. The former is a three-dimensional extrusion of the initial pulse in the `advectionDiffusion2d` example, such that the equation admits the analytical solution [2, 37, 38]

$$\chi^{*,s}(x, y, z, t) = \sin[\pi(x - u_x t)] \sin[\pi(y - u_y t)] \sin[\pi(z - u_z t)] \exp(-3\mu\pi^2 t). \quad (8.8)$$

: periodically connected in and outlets with forcing combined with different wall treatment. The latter comprises a Dirac delta at x_0 as initial pulse which induces a Dirac comb as super-positioned analytical solution [2]

$$\chi^{*,u}(x, y, z, t) = \frac{1}{\sqrt{4\pi\mu t}} \sum_{k \in \mathbb{Z}} \exp\left(\frac{-(x - x_0 - u_x t + 2k)^2}{4\mu t}\right) + 1. \quad (8.9)$$

For each case several error norms over the domain measure the deviations from the analytical solution. For the default setting the outputs of three subsequent simulation runs are stored in a subfolder structure in `./tmp` and directly post-processed for visualization. Via issuing the command

```
python3 advectionDiffusion3dPlot.py
```

an error plot is produced, which numerically validates the: periodically connected in- and outlets with forcing combined with different wall treatment second order convergence for both initializations, under the constraints on the grid Péclet number derived in [2].

8.3.6. advectionDiffusionPipe3d

This example implements a spreading Gaussian density package advecting within a square duct pipe with velocity \mathbf{U} . The precise description of the test-case can be found in [2]. Whereas the velocity is computed via approximating the incompressible Navier–Stokes equations with a *D3Q19* BGK LBM, the advection–diffusion equation for the density package is solved with finite differences (FD). Four different FD schemes can be employed within the example. The advantages of each of the schemes for a broad range of Pe are documented in [2].

8.4. laminar

8.4.1. bstep2d and bstep3d

This example implements the fluid flow over a backward facing step. Furthermore, it is shown how checkpointing is used to regularly save the state of the simulation. The 2D geometry corresponds to Armaly *et al.* [77].

8.4.2. cavity2d, cavity2dSolver and cavity3d

This example illustrates a flow in a cuboid, lid-driven cavity. The 2D version also shows how to use the XML parameter files and has an example description file for OpenMPI. This example is available in two different versions for sequential and parallel use. The `2dSolver`-version illustrates the use of the solver class concept (cf. Section 10.11) together with the XML parameter interface.

8.4.3. cylinder2d and cylinder3d

This example examines a steady flow past a cylinder placed in a channel. The cylinder is offset somewhat from the center of the flow to make the steady-state symmetrical flow unstable. At the inlet, a Poiseuille profile is imposed on the velocity, whereas the outlet implements a Dirichlet pressure condition set by $p = 0$, inspired by [123]. For high resolution, low `latticeU`, and enough time to converge, the results for pressure drop, drag and lift lie within the estimated intervals for the exact results. An unsteady flow with Karman vortex street can be created by changing the Reynolds number to $Re = 100$. The 3D version also shows the usage of the STL-reader. The model was created using the open source CAD tool FreeCAD [73].

8.4.4. poiseuille2d and poiseuille3d

For basic tests of boundary conditions, a comparison with analytical solutions is the easiest and most accurate approach. One of the fundamental applications of fluid dynamics is that of laminar flow of a Newtonian fluid in a circular pipe. This is known as Poiseuille flow. The analytical solution is easily found and is therefore a common benchmark case (see Figure 8.2). It is also one of the first examples in most fluid dynamics text books for the application of the principles of fluid dynamics. The extension of the Poiseuille flow in a round pipe from 2D to 3D is trivial, consequently it is also an ideal test case for curved boundaries in 3D as well.

8.4.5. poiseuille2dEOC and poiseuille3dEOC

The examples `poiseuille2dEOC` and `poiseuille3dEOC` are extensions of the `poiseuille2d` and `poiseuille3d` examples, respectively, focusing on the experimental order of convergence (EOC) of the simulation. The implementations therein run the Poiseuille flow simulations multiple times to subsequently analyze the different error norms of the simulations with an automated `gnuplot` output (cf. Section 6.5.1). Exemplary, the `poiseuille3dEOC` example is simulated for multiple configurations: the




Figure 8.2.: Geometry setup in example `poiseuille3d` with boundary patches and velocity profile.

periodic pipe flow with *forcing* and several boundary methods at the walls (keywords *interpolated*, *bouzidi*, and *bounce back* are used to specify the implementations in Section 5.1), as well as the a bounded domain case where a *velocity inlet* and a *pressure outlet* are prescribed and combined with a *bouzidi* boundary at the pipe wall. For all simulations, the Reynolds number of the flow problem is set to $Re = 10$ and the relaxation time to $\tau = 0.8$. The EOC is investigated by running four consecutive simulations with increasing grid resolution (number of lattice cells along the pipe diameter: $N_D = 21, 31, 41, 51$). A residual below $1e-5$ stops each simulation via the convergence criteria (see Section 4.5.2 and Section 10.5). The absolute error of the simulation from the analytic solution is evaluated on every lattice node whereby different norms for the error calculation are used. The L_1 -norm is defined as

$$E_{abs,L_1} = \sum_i^N \sum_d^D |\phi_{sim,i,d} - \phi_{ana,i,d}| \Delta x^D, \quad (8.10)$$

where N is the number of lattice nodes, D is the number of dimensions of the variable ϕ , ϕ_{sim} is the solution of a flow variable obtained by the simulation, ϕ_{ana} is the corresponding analytic solution, and Δx is the grid spacing. The L_2 -norm is defined as

$$E_{abs,L_2} = \sqrt{\sum_i^N \sum_d^D |\phi_{sim,i,d} - \phi_{ana,i,d}|^2 \Delta x^D}. \quad (8.11)$$

Finally, the L_∞ -norm is defined as

$$E_{abs,L_\infty} = \max_{i,d} |\phi_{sim,i,d} - \phi_{ana,i,d}|. \quad (8.12)$$

The relative error according to the L_1 -norm is defined as

$$E_{rel,L_1} = \frac{E_{abs,L_1}}{\sum_i^N \sum_d^D |\phi_{ana,i,d}| \Delta x^D}, \quad (8.13)$$

the L_2 -norm as

$$E_{rel,L_2} = \frac{E_{abs,L_2}}{\sqrt{\sum_i^N \sum_d^D |\phi_{ana,i,d}|^2 \Delta x^D}}, \quad (8.14)$$

and the L_∞ -norm as

$$E_{rel,L_\infty} = \frac{E_{abs,L_\infty}}{\max_{i,d} |\phi_{ana,i,d}|}. \quad (8.15)$$

The results of the EOC tests are plotted in Figures 8.3, 8.4, 8.5, 8.6, and 8.7.

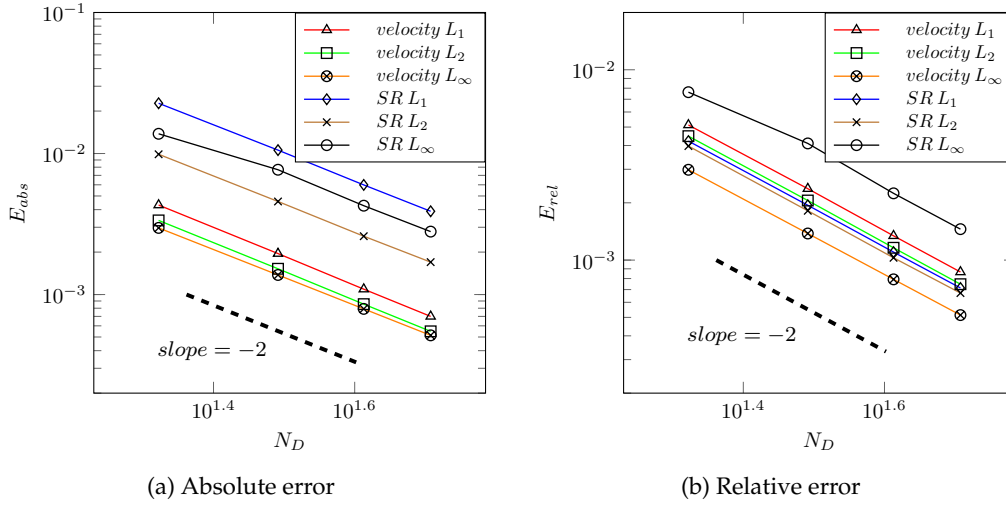


Figure 8.3.: EOC study for periodic pipe flow simulations (poiseuille3dEOC) with *forcing* and *interpolated* boundary walls. The absolute and relative error of the velocity and the strain-rate (SR) are plotted.

8.4.6. powerLaw2d

This example describe a steady non-Newtonian flow in a channel. At the inlet, a Poiseuille profile is imposed on the velocity, whereas the outlet implements a Dirichlet pressure condition set by $p = 0$.

8.4.7. testFlow3dSolver

This app implements a Navier–Stokes flow with an analytical solution [0]. The *standard* simulation as well as an EOC computation for various error norms are shown. The implementation makes use of the Solver framework, while simulation parameters are read from the corresponding parameter XML file. Among others, the following options can be selected:

- The flow domain can be either a cube or restricted to a sphere (node [Application][Domain])
- Different LBM-boundary conditions (all of them implement a Dirichlet fixed velocity, node [Application][BoundaryCondition])

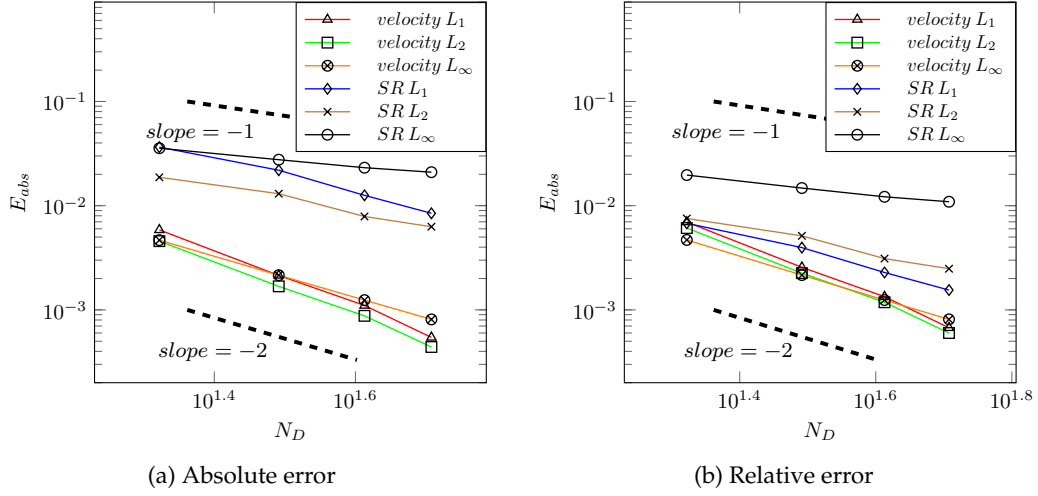


Figure 8.4.: EOC study for periodic pipe flow simulations (poiseuille3dEOC) with *forcing* and *bouzidi* boundary walls. The absolute and relative error of the velocity and the strain-rate (SR) are plotted.

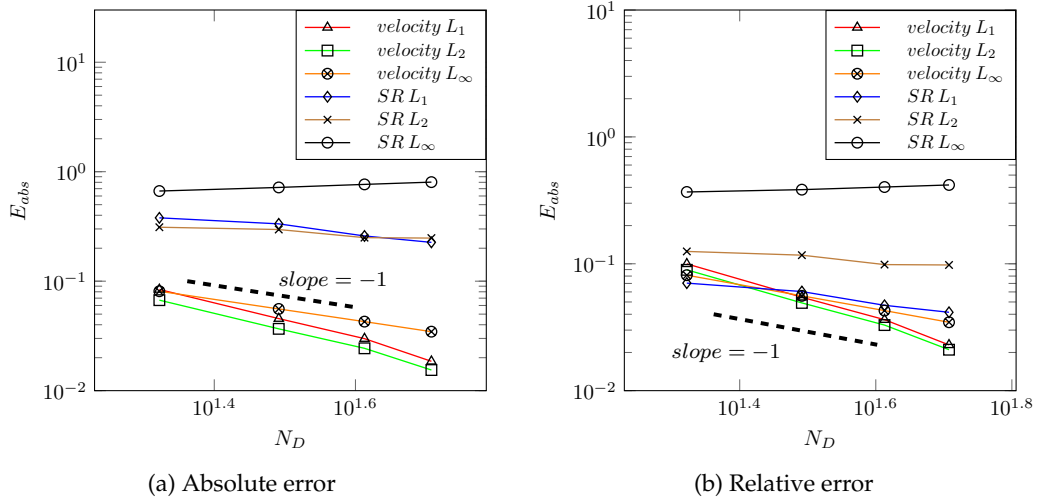




Figure 8.5.: EOC study for periodic pipe flow simulations (poiseuille3dEOC) with *forcing* and *bounce-back* boundary walls. The absolute and relative error of the velocity and the strain-rate (SR) are plotted.

Especially, the EOC computation can be used to benchmark various aspects of LBM implementation, e.g. the boundary conditions and the collision steps. The result on an example run is shown in Figure 8.8. The class structure is as follows: `TestFlowBase` implements the basic simulation details (geometry, boundary conditions etc.). Since it shall later be used for adjoint optimization, it inherits from `AdjointLbSolverBase`. The specification `TestFlowSolver` then implements the standard simulation. The distinction between `TestFlowBase` and `TestFlowSolver` is solely because `TestFlowBase` will also be used for optimization (cf.




(a) Absolute error




(b) Relative error

Figure 8.6.: EOC study for pipe flow simulations (poiseuille3dEOC) with *velocity inlet*, *pressure outlet* and *bouzidi* boundary walls. The absolute and relative error of the velocity and the strain-rate (SR) are plotted.



(a) Absolute error



(b) Relative error

Figure 8.7.: EOC study for pipe flow simulations (poiseuille3dEOC) with *velocity inlet*, *pressure outlet* and *bouzidi* boundary walls. The absolute and relative error of the pressure are plotted.

examples in Section 8.5.5).




Figure 8.8.: EOC study for the example `testFlow3dSolver` on a sphere, as it is returned via the `gnuplot` interface.

8.5. optimization

8.5.1. domainIdentification3d

In this example, a domain identification problem is solved: a cubic obstacle in the middle of a cubic fluid flow domain has to be identified given only the surrounding fluid flow, cf. [17]. Therefore, an optimization problem is set up, where the porosity field α in the design domain has to be found such that the error J between resulting velocity field \mathbf{u} and original velocity field \mathbf{u}^* is minimized. Hence, the objective is $J := \frac{1}{2} \int_{\tilde{\Omega}} |\mathbf{u} - \mathbf{u}^*|^2 dx$. The example employs adjoint Lattice Boltzmann methods for gradient computation. Results are shown in Figure 8.9.

The implementation makes use of the Solver framework and the XML interface for parameter reading, cf. Section 10.11. In addition to the typical parameters for simulation and optimization, the following geometric domains are read from the xml file:

- `<ObjectiveDomain>`: The domain $\tilde{\Omega}$, where the objective functional is computed.

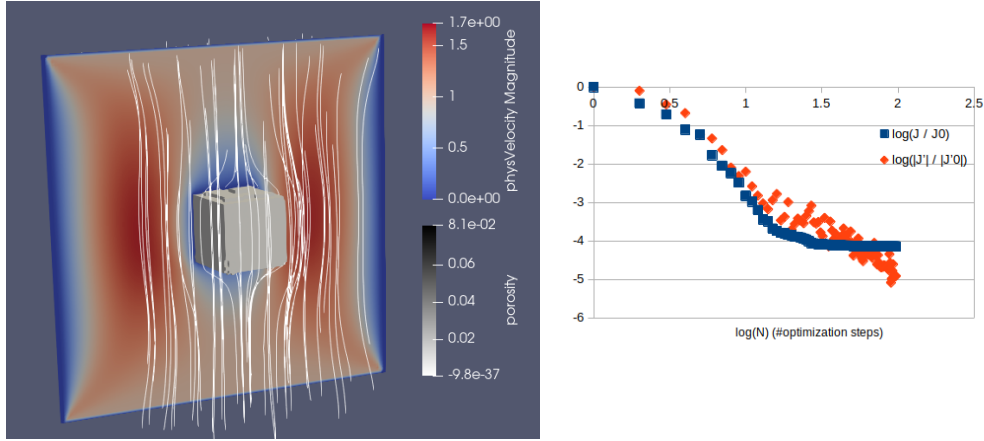


Figure 8.9.: Results of domain identification after 97 optimization steps. Left: identified object, surrounding velocity field and streamlines. Right: relative decrease of the objective functional and its derivative w.r.t. the design variables. We can see that the objective hardly changes after ca. 30 optimization steps.

- `<SimulationObject>`: The object that is to be identified by the example, it is used for the computation of the reference solution \mathbf{u}^* .
- `<DesignDomain>`: the domain, where the porosity field is computed by the optimization algorithm.

The OpenLB software provides the method *createIndicator* to create these domains or objects.

For example, a three-dimensional cuboid with length 1 in the origin can be created with

`createIndicatorF3D<T>(reading from xml-file)` with the input from the xml-file:

```
<IndicatorCuboid3D extend = "1 1 1" origin = "0 0 0"/>.
```

These domains are objects and depend on the specific optimization problem and have to be created or modified for each problem again.

The following optimization parameters are specific in the context of domain identification and set in the xml file:

- `ReferenceSolution`: decide whether the reference velocity \mathbf{u}^* is computed via a simulation
- `StartValueType`: select between Control, Permeability and Porosity

8.5.2. poiseuille2Opti

This is a simple showcase for a two-dimensional fluid flow optimization/ parameter identification problem: It is based on the simulation of a planar channel flow similar to the example `poiseuille2d`. By an optimization loop, the inlet pressure is determined s.t. a given mass flow rate is achieved. Hence, we solve the following optimization problem:

$$\operatorname{argmin}_{p_0} \frac{1}{2} (m(\mathbf{u}(p_0), p(p_0)) - m^*)^2, \quad (8.16)$$

where $(\mathbf{u}(p_0), p(p_0))$ is the solution of the Navier-Stokes equations corresponding to the inlet pressure p_0 , m is the mass flow rate corresponding to that solution, and m^* is the wanted mass flow rate.

The optimization is performed with the LBFGS method; the derivative of the objective w.r.t. the argument (inlet pressure) is computed with automatic differentiation.

8.5.3. showcaseADf

This app gives an introduction into the usage of forward automatic differentiation in order to compute derivatives of any numerical quantities in OpenLB. It is written in the style of literate programming as a tutorial that should allow for being read sequentially.

8.5.4. showcaseRosenbrock

This app gives an introduction into the usage of optimization functionalities in OpenLB in order to compute optima of functions or simulations. It is written in the style of literate programming as a tutorial that should allow for being read sequentially.

8.5.5. testFlowOpti3d

This app is built on top of the example laminar/testFlow3dSolver and solves parameter identification problems with an optimization approach [0]. It makes use of the Solver framework to create the following variants of application:

- Computing sensitivity of flow quantities w.r.t. the force field with Difference Quotients (DQ) or Automatic Differentiation (AD)
- Optimization with AD: scale the force field s.t. the velocity or dissipation error w.r.t. a given flow is minimized. 3 control variables scale the components of the force field. The objective is $J = \frac{1}{2} \int_{\Omega} |\mathbf{u} - \mathbf{u}^*|^2 dx$, where \mathbf{u}^* is the reference solution field (for control = (1, 1, 1)) and \mathbf{u} is the simulated solution (velocity or dissipation) with "estimated" control parameters.
- Optimization with adjoint LBM: identify the force field s.t. the velocity or dissipation error w.r.t. a given flow is minimized. The objective is as above, but the control is the (spatially) distributed field – three components at each mesh point.

The basic simulation setup is implemented in the base class `TestFlowBase`. Further optimization-specific implementations (which are used in all optimization variants, e.g. objective computation) are done in `TestFlowOptiBase`. For DQ and AD, `TestFlowSolverDirectOpti` implements the case-specific features, e.g. that the vector of control variables acts as a scaling factor for the force field. For adjoint LBM, case-specific features are implemented in `TestFlowSolverOptiAdjoint`. The structure of parameter classes corresponds to that of the solver classes.

Simulation and optimization parameters are read from the corresponding parameter XML file. The following optimization parameters are specific for this example and set in the XML file:

- `ReferenceSolution`: decide whether reference solution \mathbf{u}^* is computed via a simulation

- `TestFlowOptiMode`: decide, whether velocity or dissipation is compared in the objective functional
- `OptiReferenceMode`: use analytical or discrete (simulated) solution as \mathbf{u}^*
- `ControlType (force)`: type of control parameters in the context of adjoint optimization
- `CuboidWiseControl`: if true, then the domain is decomposed into an arbitrary number of cuboids and one control variable is set as a scaling factor for each cuboid (this is meaningful if the sensitivities are computed with finite difference quotients or automatic differentiation).

The results of example runs with AD and adjoint LBM are shown in Figure 8.10.




Figure 8.10.: Decrease of the objective for increasing optimization steps for different resolutions (left: AD, right: adjoint LBM). The trends differ since the number of control variables is fixed in the AD case while it increases with the resolution in the adjoint optimization case.

8.6. multiComponent

The examples in this folder demonstrate the use of specific LBM models (e.g. Section 4.4.2) for multiphase and multicomponent flows.

8.6.1. binaryShearFlow2d

A circular domain of one fluid phase is immersed in a rectangle filled with another fluid phase. The top and bottom walls are moving in opposite directions, such that the droplet shaped phase is exposed to shear flow and deforms accordingly. The default parameter setting is taken from [99] and injected into the more general ternary free energy model from [116]. Both scenarios, breakup and steady state of the initial droplet, are implemented and visualized as `.vtk` output. Reference simulations are provided in [40].

8.6.2. contactAngle2d and contactAngle3d

In this example a semi-spherical droplet of fluid is initialized within a different fluid at a solid boundary. The contact angle is measured as the droplet comes to equilibrium. This is compared with the analyti-

cal angle predicted by the parameters set for the boundary (100 degrees for preset values). This example demonstrates how to use the solid wetting boundaries for the free-energy model with two fluid components.

8.6.3. fourRollMill2d

Here, a spherical domain filled with one fluid phase is immersed in a square filled with another phase of equal density and viscosity. Four circle structures which represent roller sections are equidistantly distributed in the corners of the domain. The bottom left and top right cylinders begin to spin in counterclockwise direction. Whereas the top left and bottom right cylinders spin in clock-wise direction. A velocity field of extensional type deforms the initial droplet accordingly. Dependent on the non-dimensional parameter setting in the example header, the droplet reaches steady state or breaks up. Reference simulations are provided in [40].

8.6.4. microFluidics2d

This example shows a microfluidic channel creating droplets of two fluid components. Poiseuille velocity profiles are imposed at the various channel inlets, while a constant density outlet is imposed at the end of the channel to allow the droplets to exit the simulation. This example demonstrates the use of three fluid components with the free energy model. It also shows the use of open boundary conditions, specifically velocity inlet and density outlet boundaries.

8.6.5. phaseSeparation2d and phaseSeparation3d

In these examples the simulation is initialized with a given density plus small, random variation over the domain. This condition is unstable and leads to liquid-vapor phase separation. Boundaries are assumed to be periodic. These examples show the usage of multiphase flow.

8.6.6. rayleighTaylor2d and rayleighTaylor3d

This example demonstrates Rayleigh–Taylor instability in 2D and 3D, generated by a heavy fluid penetrating a light one. The multicomponent fluid model by X. Shan and H. Chen is used [117]. These examples show the usage of multicomponent flow and periodic boundaries.

8.6.7. youngLaplace2d and youngLaplace3d

In this example the two-component free energy model is used in its simplest configuration to perform a Young–Laplace pressure test. A circular or spherical domain of a fluid with radius R is immersed in another fluid. A diffusive interface forms and the pressure difference across the interface, Δp , is calculated

and compared to that given by the Young–Laplace equation,

$$\Delta p = \frac{\gamma}{R} = \frac{\alpha}{6R}(\kappa_1 + \kappa_2) \quad \text{for 2D ,} \quad (8.17)$$

$$\Delta p = \frac{2\gamma}{R} = \frac{\alpha}{3R}(\kappa_1 + \kappa_2) \quad \text{for 3D .} \quad (8.18)$$

The parameters α and κ_i are input parameters to the simulation which define the interfacial width and surface tension, γ , respectively. The pressure difference is calculated between a point in the middle of the circular domain and a point furthest away from it in the computational domain.

8.7. particles

8.7.1. bifurcation3d

The `bifurcation3d` example simulates particulate flow through an exemplary bifurcation of the human bronchial system. The geometry is a splitting pipe, with one inflow and two outflows. The fluid is transporting micrometer scale particles and the escape and capture rate is computed. There are two implementations of the problem. The first one is a Euler–Euler ansatz, meaning that the fluid phase as well as the particle phase are modeled as continua. The second is an Euler–Lagrange ansatz, where the particles are modeled as discrete objects.

8.7.1.1. eulerEuler

In this example the particles are viewed as a continuum and described by a advection–diffusion equation. This is done similar to the thermal examples, where the temperature is the considered quantity. For particles however, inertia has to be taken into account. This is achieved by applying the Stokes drag force to the velocity field. Since for this computations also the velocity of the previous time step is required, the new descriptor `ParticleAdvectionDiffusionD3Q7Descriptor` has to be used, that is capable of saving 2 velocity fields. Besides an extra lattice for the advection–diffusion equation, a `SuperExternal3D` structure is required to manage the communication for parallel execution.

```

1 SuperExternal3D<T,ADDESCRIPTOR,descriptors::VELOCITY> sExternal(
2     superGeometry,
3     sLatticeAD,
4     sLatticeAD.getOverlap());
5
6     ...
7
8 sExternal.communicate();
```

The function `communicate()` is called in the time loop and handles the communication analogue to the lattices.

Furthermore the new dynamics object `ParticleAdvectionDiffusionBGKdynamics` is required to access the saved velocity fields correctly and use them in an efficient way. For information on the coupling of the

lattices we refer to the section on the advection–diffusion equation for particle flow problems 4.5.5.2. In this example only the Stokes drag is applied by

```
1 advDiffDragForce3D<T, NSDESCRIPTOR> dragForce( converter, radius, partRho );
```

For the simulation of particles as a continuum, also new boundary conditions are required. Here `setZeroDistributionBoundary` represents an unidirectional outflow condition, that removes particle concentrations that cross a boundary. For the usual outflow at the bottom of the bifurcation a new `AdvectionDiffusionConvectionBoundary` for advection–diffusion lattices can be applied, that approximates a Neumann boundary condition, for further reference see [45]. Since non-local computations (gradient is required) are performed on the the external field, also a Neumann boundary condition is required that is here implemented as `setExtFieldBoundary`.

8.7.1.2. eulerLagrange

The main task of his example is to show the using of Lagrangian particles with OpenLB. As this example is used to show the application of the particle framework (8.7), implementation specifics can be found there.

8.7.2. dkt2d

OpenLB provides an alternative approach to conventional resolved particle simulation methods, referred to as the homogenized lattice Boltzmann method (HLBM). It was introduced in "Particle flow simulations with homogenized lattice Boltzmann methods" by Krause *et al.* [20] and extended for the simulation of 3D particles in Trunk *et al.* [45]. It was eventually revisited in [47]. In this approach the porous media model, introduced into LBM by Spaid and Phelan [120], is extended by enabling the simulation of moving porous media. In order to avoid pressure fluctuations, the local porosity coefficient is used as a smoothing parameter.

The example `dkt2d` employs said approach for the sedimentation of two particles under gravity in a water-like fluid in 2D. The rectangular domain is limited by no-slip boundary conditions. This setup is usually referred to as a drafting–kissing–tumbling (DKT) phenomenon and is widely used as a reference setup for the simulation of particle dynamics submerged in a fluid. The benchmark case is described e.g. in [89, 124]. For the calculation of forces a DNS approach is chosen which also leads to a back-coupling of the particle on the fluid, inducing a flow. The example demonstrates the usage of HLBM in the OpenLB framework as well as the utilization of the `gnuplot-writer` to print simulation results.

8.7.3. magneticParticles3d

Warning: This example can currently only be run sequentially! High-gradient magnetic separation is a method to separate ferromagnetic particles from a suspension. The simulation shows the deposition of magnetic particles on a single magnetized wire and models the magnetic separation step of the complete process.

8.7.4. settlingCube3d

The case examines the settling of a cubical silica particle under gravity in a surrounding fluid. The rectangular domain is limited by no-slip boundary conditions. For the calculation of forces a DNS approach is chosen which also leads to a back-coupling of the particle on the fluid, inducing a flow. The example demonstrates the usage of HLBM in the OpenLB framework as well as the utilization of the `gnuplot`-writer to print simulation results (Section 8.7.2).

8.8. porousMedia

8.8.1. porousPoiseuille2d and porousPoiseuille3d

This example simulates Poiseuille flow through porous media. The implementation reproduces the benchmark *Example A* by Guo and Zhao [95]. The theoretical maximum velocity is calculated as in Equation 21, and the velocity profile as in Equation 23 of [95]. For a schematic simulation setup see Figure 8.11.




Figure 8.11.: Geometry used in the example `porousPoiseuille2d` with boundary patches and velocity profile.

8.9. reaction

8.9.1. advectionDiffusionReaction2d

This example illustrates a forward or a reversible reaction of a substance into another one ($A \rightarrow C$ or $A \rightleftharpoons C$). The concentration of each substance is modeled with a one-dimensional advection diffusion equation. We consider a steady state in a plug flow reactor. This means that it is time independent and we have a constant velocity field u . The chemical reaction is modeled with a linear but coupled source term. The

equations read

$$\begin{cases} u\partial_x c_A &= D\partial_x^2 c_A - k_H c_A + k_R c_C, \\ u\partial_x c_C &= D\partial_x^2 c_C + k_H c_A - k_R c_C, \end{cases} \quad (8.19)$$

with the diffusion coefficient $D > 0$, forth reaction rate coefficient $k_H > 0$ and backwards reaction rate coefficient $k_R > 0$ ($= 0$ in case of $A \rightarrow C$). This LBM models the transport of the species concentration along a one-dimensional line on $[0, 10]$. In practice the simulation uses a two-dimensional rectangular domain which is evaluated along a centerline to obtain the desired one-dimensional result. The height of the domain depends on the resolution which holds the number of voxels for the height constant.

On the bottom and the top of the rectangular periodic boundaries and at the inlet and outlet the concentrations from the analytical solution are set. The solution is given by

$$c_A(x) = c_{A,0} e^{\lambda x} + \frac{k_R}{k_H + k_R} c_{A,0} (1 - e^{\lambda x}), \quad (8.20)$$

$$c_C(x) = c_{A,0} - c_A(x), \quad (8.21)$$

with $\lambda = \frac{u - \sqrt{u^2 + 4(k_H + k_R)D}}{2D}$.

Diffusive scaling is applied and a physical diffusivity of $D = 0.1$ and a flow rate of $u = 0.5$ which leads to a Péclet number of $Pe = 100$.

Every species has its own lattice and stored in a vector. In the `simulate` method we iterate over every element of the vector `adlattices`.

In the default setting, `advectionDiffusionReaction2d` executes three simulation runs with increasing resolutions $N = 200, 250, 300$, respectively. The output of each simulation run is stored in the `tmp/N<number>` directory. It contains a plot `centerConcentrations.pdf` of the concentrations and the analytical solution along the centerline and an error plot of the numerical and analytical solution along the centerline `ErrorConcentration.pdf`. After each simulation has converged the average L2 relative Error over the centerline is computed. Said average is then stored within `tmp/gnuplotData/data/averageL2RelError.dat`. The order of convergence can be seen in the log-log error plot in `tmp/gnuplotData/concentration_eoc.png`. The EOC plot is only done for one species (A or C) and the species can be chosen in the return statement of the method `errorOverLine`.

One can select the `reactionType` `a2c` or `a2cAndBack` which automatically provides the data for modeling the reaction. It contains the number of reactions, the reaction rate coefficients `physReactionCoeff[numReactions]` (k_H, k_R), the number of species `numComponents` and their names, the stoichiometric coefficients `stochCoeff` which are sorted according to the number of reactions and inside each reaction block according to the species number. Finally we assume that the reaction rate satisfies a power law depending on the concentration of the species. The exponent is given by the reaction order `reactionOrders` which is sorted in the same way as `stochCoeff`. In the example cases these exponents are always 1. The chemical reaction itself is represented as a source term for each Advection Diffusion equation. This source term is calculated in the `ConcentrationAdvectionDiffusionCouplingGenerator` for every species which can handle arbitrary number of species and reactions and stored in the field `SOURCE`.

8.10. thermal

8.10.1. galliumMelting2d

The solution for the melting problem (solid-liquid phase change) coupled with natural convection is found using the lattice Boltzmann method after Huang and Wu [97]. The equilibrium distribution function for the temperature is modified in order to deal with the latent-heat source term. That way, iteration steps or solving a group of linear equations is avoided, which results in enhanced efficiency. The phase interface is located by the current total enthalpy, and its movement is considered by the immersed moving boundary scheme after Noble and Torczynski [112]. This method was validated by comparison with experimental values (e.g. Gau and Viskanta [91]).

8.10.2. porousPlate2d, porousPlate3d and porousPlate3dSolver

The porous plate problem is implemented as described in [94] and [113]. To test the coupled model's accuracy and to determine its EOC, we use a numerical simulation the porous plate problem including a temperature gradient and natural convection in a square cavity. The porous plate problem describes a channel flow, where the upper cool plate moves with a constant velocity, and through the bottom warm plate a constant normal flow is injected and withdrawn at the same rate from the upper plate. At the left and right hand side of the domain a periodic boundary condition is applied and constant velocity and temperature boundary conditions are applied to the top and bottom plates according to Figure 8.12. An




Figure 8.12.: Schematic representation of the porous plate's simulation setup including the boundary conditions.

analytical solution for the given steady state problem is given for the velocity and temperature distribu-

tions by

$$u_x(y) = u_{x,0} \left(\frac{e^{Re \cdot y/L} - 1}{e^{Re} - 1} \right), \quad (8.22)$$

$$T(y) = T_0 + \Delta T = \left(\frac{e^{Pr \cdot Re \cdot y/L} - 1}{e^{Pr \cdot Re} - 1} \right). \quad (8.23)$$

Here $u_{x,0}$ is the upper plate's velocity, $Re = \frac{u_{y,0}L}{\nu}$ the Reynolds number depending on the injected velocity $u_{y,0}$, the fluid's viscosity ν and the channel length L . The temperature difference between the hot and cold plate is given by $\Delta T = T_h - T_c$. First we implement a couple of simulations to scale the velocity and temperature profiles for a range of the Reynolds number and Prandtl number. The relative global error is computed via [94]

$$E = \frac{\sqrt{\sum_i |T(x_i) - T_a(x_i)|^2}}{\sqrt{\sum_i |T_a(x_i)|^2}}, \quad (8.24)$$

where the summation is over the entire system, T_a is the analytical solution (8.23). The `porousPlate3dSolver` example implements the same simulation, but additionally illustrates the application of the solver class concept.

8.10.3. rayleighBenard2d and rayleighBenard3d

The Rayleigh–Bénard convection is a typical case of natural convection, where the lower boundary is heated and a regular pattern of convection cells is developed. This is a suitable test platform for thermal algorithms, since the driving force is a coupling between momentum and energy equations by means of a buoyancy force, which is function of the temperature, and the temperature varies spatially inside the domain. This example demonstrates Rayleigh–Bénard convection rolls in 2D and 3D, simulated with the thermal LB model by Guo *et al.* [93], between a hot plate at the bottom and a cold plate at the top.

8.10.3.1. Setup

The case considered has an aspect ratio ($AR = Lx/Ly$) of 2, which enhances the appearance of unstable modes. The lower wall is heated with a constant temperature ($T = 1$), and the upper wall is isothermal and cold ($T = 0$). The vertical walls are set to be periodic.

Among the example programs implemented in OpenLB, a demo code for the Rayleigh–Bénard convection in 2D and 3D is provided. This code is taken as a base for the development of most of the thermal applications. For the simulation of the Rayleigh–Bénard convection only one modification is made to the code regarding the initial conditions: to enhance the appearance of the convection cells, an instability in the domain is introduced. The available code initializes a small area near the lower boundary with a slightly higher temperature, introducing a perturbation in the system, whereas the rest of the domain is initialized with the cold temperature. In the modified code there is no local perturbation, but the initial temperature at the domain is dependent on the space coordinates. The domain is initialized with zero

velocity and a temperature field by using a functor according to

$$T(x, y, t = 0) = T_{max} \left[\left(1 - \frac{y}{L_y} \right) + 0.1 \cos \left(2\pi \frac{x}{L_x} \right) \right]. \quad (8.25)$$

The files created to help with the initialization of the temperature field are called `tempField.h` and `tempField2.h`. Listing 8.1 shows the corresponding usage. The first file computes the temperature at every point of the lattice, as a function of its macroscopic position, and then this value is applied on the lattice as the density (line 3). The second file calculates the equilibrium distribution functions for every node corresponding to the given temperature and zero velocity. Next, the populations are defined for the desired material number in line 4.

```

1 TemperatureField2D<T,T> Initial( converter );
2 TemperatureFieldPop2D<T,T> EqInitial( converter );
3 ADlattice.defineRho( superGeometry, 1, Initial );
4 ADlattice.definePopulations( superGeometry, 1, EqInitial );

```

Listing 8.1.: Initialization of the temperature field

In (8.25), the y -dependent part of the equation matches the stationary solution of the problem, corresponding to a case where there is no fluid movement and the heat transfer only occurs by conduction. The cosine term introduces a disturbance in the system, which enhances the appearance of the convection cells.

8.10.3.2. Simulation Parameters

Computations are run for a range of different Rayleigh ($3 \cdot 10^3$, $6 \cdot 10^5$) and Prandtl numbers (0.3, 1). The spatial resolution was fixed to 100 cells in the y -direction, and the time discretization was switched between 10^{-3} and 10^{-4} , which give lattice velocities of 0.1 and 0.01 respectively. The convergence criterion is applied on the average energy, and it is set to a precision of 10^{-5} .

8.10.4. squareCavity2d and squareCavity3d

A common application for the validation of thermal models is the numerical simulation of the natural convection in a square cavity. For this configuration there is an extensive database in a wide range of Rayleigh numbers, which allows to verify the accuracy of the thermal model.

8.10.4.1. Setup

The problem considered is shown schematically in Figure 8.13. The horizontal walls of the cavity are adiabatic, while the vertical walls are kept isothermal, with the left wall at high temperature ($T_{hot} = 1$) and the right wall at low temperature ($T_{cold} = 0$).

The dynamics chosen for the velocity field is `ForcedBGKdynamics`, and for the temperature field `AdvectionDiffusionBGKdynamics`.




Figure 8.13.: Schematic diagram of the simulation domain for the example `squareCavity2d`.

8.10.4.2. Simulation Parameters

Taking air at 293K as working fluid, the value of the Prandtl number is $Pr = 0.71$ and is kept constant. The Rayleigh number ranges from 10^3 to 10^6 . Different spatial resolutions are tested for each Rayleigh number, in order to study the grid convergence. The time-step size is adjusted so that the lattice velocity stays at the value 0.02. This ensures that the Mach number is kept at incompressible levels. The convergence criterion is set by a standard deviation of 10^{-6} in the kinetic energy.

8.10.4.3. MRT

The new implemented MRT model for thermal applications is first examined on the 2-dimensional cavity. The only setup differences to the BGK model are the lattice descriptors (`ForcedMRTD2Q9Descriptor` and `AdvectionDiffusionMRTD2Q5Descriptor`) and the dynamics objects selected, which are now specialized for the MRT dynamics (`ForcedMRTdynamics` and `AdvectionDiffusionMRTdynamics`). This simulation is used as a test for different important aspects of the implementation. First, the formulation of the MRT model, particularly the values of the transformation matrix, the relaxation times and the sound speed of the lattice are based on [106], but it shows variations over 10% with respect to the BGK model. A second formulation [107] is selected, which shows much closer results to the BGK model. No special treatment is required to make use of the available boundary conditions. The number of iterations required to achieve the desired precision, that is, the number of time steps until the steady-state solution is reached, is found to be usually higher for the MRT simulations. Furthermore, the execution time is between 4 and 8 times longer when compared to the BGK simulations.

8.10.5. stefanMelting2d

The solution for the melting problem (solid-liquid phase change) is computed using the LBM from Huang and Wu [97]. The equilibrium distribution function for the temperature is modified in order to deal with

the latent-heat source term. That way, iteration steps or solving a group of linear equations is avoided, which results in enhanced efficiency. The phase interface is located by the current total enthalpy, and its movement is considered by the immersed moving boundary scheme after Noble and Torczynski [112]. Huang and Wu validated this method by the problem of conduction-induced melting in a semi-infinite space, comparing its results to analytical solutions.

8.11. turbulent

8.11.1. aorta3d

In this example, the fluid flow through a bifurcation is simulated. The geometry is obtained from a mesh in STL format. With Bouzidi boundary conditions, the curved boundary is adequately mapped and initialized entirely automatically. A Smagorinsky LES BGK model is used for the dynamics to stabilize the turbulent flow simulation for low resolutions. The output is the flux computed at the inflow and outflow region. The results have been validated through comparison with other results obtained with FEM and FVM.

8.11.2. channel3d

This example features the application of wall functions in a bi-periodic, fully developed turbulent channel flow for friction Reynolds numbers of $Re_\tau = 1000$ and $Re_\tau = 2000$. For the published results and further reference see [9].

8.11.3. nozzle3d

On the one hand this example describes building a cylindrical 3D geometry in OpenLB, on the other hand it examines turbulent flow in a nozzle injection tube using different turbulence models and Reynolds numbers.

For characterization different physical parameters have to be set. The resolution N defines most physical parameters such as the velocity `charU`, the kinematic viscosity ν and two characteristic lengths `charL` and `latticeL`. Physical length `charL` is used to characterize the geometry and the Reynolds number. Lattice length `latticeL` defines the mesh size and is calculated as $\text{latticeL} = \text{charL}/N$. More information about the parameter definitions are located in the file `units.h`. Figure 8.14 illustrates the geometry and the nozzle's size as a function of the characteristic length `charL`. The nozzle consists of two circular cylinders. The inflow (red) is located left in the *inletCylinder*. The outflow (green) is at the right end of the *injectionTube*. At the main inlet, either a block profile or a power 1/7 profile is imposed as a Dirichlet velocity boundary condition, whereas at the outlet a Dirichlet pressure condition is set by $p = 0$ (i.e. $\rho = 1$). Two vectors, `origin` and `extend`, describe the center and normal direction of the cylinder's circular start (origin) and end (extend) plane. The radius is defined in the function.

As mentioned before, this example simulates turbulent fluid flow. The flow behavior in the inlet is characterized by the Reynolds number. The following turbulence models are based on large eddy simulation

Table 8.4.: Preset simulation parameters of the `nozzle3d` example.

parameter	value
<code>charL</code>	1m
<code>latticeL</code>	$\frac{1}{3}$ m
<code>charU</code>	$1 \frac{\text{m}}{\text{s}}$
ν	$0.00002 \frac{\text{m}^2}{\text{s}}$
Re_{inlet}	5000
turbulence model	Smagorinsky




Figure 8.14.: Cross section of a 3D geometry of `nozzle3d` in dependency of characteristic length `charL`.

(LES). The idea behind LES is to simulate only eddies larger than a certain grid filter length, while smaller eddies are modeled. Several models are currently implemented, e.g.:

- The **Smagorinsky model** reduces the turbulence to a so called eddy viscosity. This viscosity depends on the Smagorinsky constant, which has to be defined. This model has certain disadvantages at the wall.
- The **Shear-improved Smagorinsky model (SISM)** is based on the Smagorinsky model. Compared to the original model, the SISM works at the wall very well. Similarly, a model specific constant has to be defined.

The following code shows the model selection. A model is selected, when the correlate line is uncom-

mented. Below, the model specific constants are defined. In this case the Smagorinsky model is selected. Smagorinsky constant is set to 0.15.

```

1  /// Choose your turbulent model of choice
2
3  #define Smagorinsky
4
5  ...
6
7  #elif defined(Smagorinsky)
8  bulkDynamics = new SmagorinskyBGKdynamics<T, DESCRIPTOR>(converter.getOmega(),
    instances::getBulkMomenta<T, DESCRIPTOR>(),
9  0.04, converter.getLatticeL(), converter.physTime());

```

As an example, Figure 8.15 shows the results with preset parameters. The simulations strongly depends




Figure 8.15.: Physical velocity field after 200 seconds with preset parameters (Smagorinsky Model, $C_S = 0.15$, $\text{latticeL} = \frac{1}{3}\text{m}$, $Re_{inlet} = 5000$).

on the Smagorinsky constant's value, used in the turbulence model. However, the constant is not a general calculable value and valid for one model. It could be a function of the Reynolds number and/or another dimensionless parameter. Thus, a physically useful value has to be found by trial and error, or chosen as an educated guess in the beginning. Generally, if the constant's value is chosen too small, the simulation results will become unstable and/or unphysical. If the value is too large, the model will introduce too much artificial viscosity and smooth the results.

8.11.4. tgv3d

The Taylor–Green vortex (TGV) is one of the simplest configurations to investigate the generation of small scale structures and the resulting turbulence. The cubic domain $\Omega = (2\pi)^3$ with periodic boundaries and the single mode initialization contribute to the model's simplicity. In consequence, the TGV is a common




Figure 8.16.: Physical velocity magnitude plotted on the cross section for two Smagorinsky constants. Increasing the Smagorinsky constant visibly smoothens the results and straightens the appearing turbulence.




Figure 8.17.: Isosurfaces of vorticity for the Taylor–Green vortex at $t = 12s$.

benchmark case for direct numerical simulation (DNS) as well as large eddy simulation (LES). This example demonstrates the usage of different sub-grid models and visualizes their effects on global turbulence quantities. The molecular dissipation rate, the eddy dissipation rate and the effective dissipation rate are calculated and plotted over the simulation time. The results can be compared with a DNS solution published by Brachet *et al.* [81].

8.11.5. venturi3d

This example examines a steady flow in a Venturi tube. A Venturi tube is a cylindrical tube, which has a reduced cross-section in the middle part. At this constriction is an injection tube. As a result of the accelerating fluid in the constriction, the static pressure decreases and the injection tube's fluid is pumped in the main tube. The overall geometry is built with adding together single bodies. Each body's geometry is defined by certain points (position vectors) in the coordinate system and their radius. A cone-shaped cylinder needs the center of the start and end circle as well as the radii. Following code builds the geometry and shows the semantics.

```

1  /// Definition of the geometry of the venturi
2
3  //Definition of the cross-sections' centers
4  Vector<T,3> C0(0,50,50);
5  Vector<T,3> C1(5,50,50);
6  Vector<T,3> C2(40,50,50);
7  Vector<T,3> C3(80,50,50);
8  Vector<T,3> C4(120,50,50);
9  Vector<T,3> C5(160,50,50);
10 Vector<T,3> C6(195,50,50);
11 Vector<T,3> C7(200,50,50);
12 Vector<T,3> C8(190,50,50);
13 Vector<T,3> C9(115,50,50);
14 Vector<T,3> C10(115,25,50);
15 Vector<T,3> C11(115,5,50);
16 Vector<T,3> C12(115,3,50);
17 Vector<T,3> C13(115,7,50);
18
19 //Definition of the radii
20 T radius1 = 10 ; // radius of the tightest part
21 T radius2 = 20 ; // radius of the widest part
22 T radius3 = 4 ; // radius of the small exit
23
24 //Building the cylinders and cones
25 IndicatorCylinder3D<T> inflow(C0, C1, radius2);
26 IndicatorCylinder3D<T> cyl1(C1, C2, radius2);
27 IndicatorCone3D<T> co1(C2, C3, radius2, radius1);
28 IndicatorCylinder3D<T> cyl2(C3, C4, radius1);
29 IndicatorCone3D<T> co2(C4, C5, radius1, radius2);
30 IndicatorCylinder3D<T> cyl3(C5, C6, radius2);
31 IndicatorCylinder3D<T> outflow0(C7, C8, radius2);
32 IndicatorCylinder3D<T> cyl4(C9, C10, radius3);
33 IndicatorCone3D<T> co3(C10, C11, radius3, radius1);
34 IndicatorCylinder3D<T> outflow1(C12, C13, radius1);
35
36 //Addition of the cylinders to overall geometry

```

```
37 IndicatorIdentity3D<T> venturi(cyl1 + cyl2 + cyl3 + cyl4 + co1 + co2 + co3);
```

Figure 8.18 visualizes the defined point positions and Figure 8.19 shows the computational geometry. At




Figure 8.18.: Schematic diagram visualizing the defined point positions for `venturi3d`.




Figure 8.19.: Built geometry as used in simulation of example `venturi3d`.

the main inlet, a Poiseuille profile is imposed as a Dirichlet velocity boundary condition, whereas at the outlet and the minor inlet, a Dirichlet pressure condition is set by $p = 0$ (i.e. $\rho = 1$). Figure 8.20 visualizes the computed velocity magnitude in the Venturi tube geometry.

8.12. freeSurface

The free surface approach [122] is a numerical simulation of two phases, where one of the phases does not have to be simulated thoroughly. In the provided examples of this category, these two phases are usually water and air, with the air phase being the one that is handled in a simplified manner.




Figure 8.20.: Simulation of the the example `venturi3d` after 200 simulated time steps.

8.12.1. `breakingDam2d` and `breakingDam3d`

The `breakingDam2d` example is based on a physical experiment conducted by LaRocque *et al.* [103]. An enclosed box contains an area of fluid in the lower left corner. With the start of the simulation, the fluid spreads throughout the box, with a visible wave forming after the fluid reaches the right-hand side wall.




Figure 8.21.: Initial breaking dam setup of `breakingDam2d` with the fluid in gray.

The example `breakingDam3d` extends the 2D example to 3D. A full simulation run of the `breakingDam3d` example, visualized in ParaView, can be found on the OpenLB YouTube page (<https://youtu.be/X8yeLCKUldQ>).

8.12.2. `fallingDrop2d` and `fallingDrop3d`

This example type simulates a drop falling into a pool of the same liquid. Figure 8.22 shows an excerpt of the first couple of time steps. On the right half of these steps, the forming of a so-called crown can be seen.

8.12.3. `deepFallingDrop2d`

A variation of the `fallingDrop2d` example type with a deeper pool of liquid that a drop will fall into. This adapted pool depth allows changes to the droplet properties, such as size, density or velocity, to be more apparent in the simulation.



Figure 8.22.: Setup and some exemplary steps of the example `fallingDrop2d`.

8.12.4. `rayleighInstability3d`

This example covers Plateau-Rayleigh instabilities. Figure 8.23 shows the initial setup as well as multiple steps throughout a simulation of this example. Each of these steps after the initial setup was captured after 1800 additional simulation steps. The perturbation was made by setting the fill level on the border of the main cylinder according to a sine wave with a wave length set to $\delta = 9R$ with R being the radius of the cylinder. This initial sine wave can be seen in the setup: Red cells are full, blue cells are empty.




Figure 8.23.: Initial setup and different steps of the Plateau-Rayleigh Instability in `rayleighInstability3d`.

9. Building and Running

OpenLB is developed for high performance computing. As such, Linux-based systems are the first class target platform. The reason for this is that most HPC clusters and, in fact, all of the 500 fastest supercomputers in the world (cf. <https://www.top500.org/>) run some kind of Linux distributions.

9.1. Install Dependencies

GNU Make in addition to a reasonably current C++ compiler supporting C++17 is all that is needed to compile and run non-parallelized OpenLB applications. OpenLB is able to utilize vectorization (AVX2/AVX-512) on x86 CPUs and NVIDIA GPUs for block-local processing. CPU targets may additionally utilize OpenMP for shared memory parallelization while any communication between individual processes is performed using MPI. It has been successfully employed for simulations on computers ranging from low-end smartphones up to supercomputers.

The present release 1.6 has been explicitly tested in the following environments:

- NixOS 22.11 and unstable (Nix Flake provided)
- Ubuntu 20.04, 22.04
- Red Hat Enterprise Linux 8.x (HoreKa, BwUniCluster2)
- Windows 10, 11 (WSL)
- MacOS 13

as well as compilers:

- GCC 9 and later
- Clang 13 and later
- Intel C++ 2021.4 and later
- NVIDIA CUDA 11.4 and later
- NVIDIA HPC SDK 21.3 and later
- MPI libraries OpenMPI 3.1, 4.1 (CUDA-awareness required for Multi-GPU); Intel MPI 2021.3.0 and later

Other CPU targets are also supported, e.g. common Smartphone ARM CPUs and Apple M1/M2.

9.1.1. Linux

It is recommended to work on a Linux-based machine. Please ensure to have the above-mentioned dependencies installed (Section 9.1). Further description is provided below in Section 9.1.3.

9.1.2. Mac

Configuring OpenLB on MacOS is explained for the release 1.5 using MacOS 11.6 in the technical report [TR6: Configuring OpenLB on MacOS](#) [74].

9.1.3. Windows

The preferable approach is to use the Windows Subsystem for Linux (WSL) introduced in Windows 10. A guide can be found in the technical report [TR5: Installing OpenLB in Windows 10](#) [75].

9.2. Compiling OpenLB Programs

OpenLB consists of generic, template-based code, which needs to be included in the code of application programs, and dependency libraries that are to be linked with the program. The installation process is light and does not require an explicit precompilation and installation of libraries. Instead, it is sufficient to unpack the source code into an arbitrary directory. Compilation of libraries is handled on-demand by the `Makefile` of an application/example program.

To get familiar with OpenLB, new users are encouraged to have a look at programs in the `examples` directory. Once inside one of the example directories, entering the command `make` will first produce libraries and then the end-user example program. This close relationship between the production of libraries and end-user programs reflects the fact that using OpenLB presently translates to writing a C++ program using the OpenLB library functions.

The file `config.mk` in the root directory can be easily edited to modify the compilation process. Available options include the choice of the compiler (GNU `g++` is the default), optimization flags, a switch between normal/debug mode, between sequential/openmp-parallel/mpi-parallel programs, and between (un)vectorized CPU and GPU platforms.

Example configuration files for common build types and systems are included in the `config/` directory of the release tarball.

To compile your own OpenLB programs from an arbitrary directory, make a copy of a sample `Makefile` contained in a default example folder. Edit the `OLB_ROOT` entry to indicate the location of the OpenLB source, and the `EXAMPLE` entry to explicit the name of your program, without file extension.

A minimal but perfectly sufficient development environment for OpenLB consists of a supported C++ compiler, a plain text editor of choice and the GNU core utilities including `make`. This means that many Linux distributions either already include everything one needs to get started or at least include everything in their default package repositories for convenient installation. One may of course also use more involved text editors with e.g. debugger integration or even full integrated development environments.

For compiling OpenLB applications outside of the provided build system one only needs to make the `src` folder available for inclusion and define the compiler and linker flags as they are visible in the printout of the default build. However, simply calling the appropriate `make` target from inside the text editor/IDE may be more convenient.

9.2.1. Using NVIDIA GPUs in OpenLB

The following section is a quick guide on how to install the CUDA functionality for Nvidia graphics cards on both Windows or Linux. The first two sections describe how to install CUDA on Windows via WSL or Linux, respectively. The third section discusses how to set up OpenMPI, a CUDA-aware MPI implementation, and the fourth and final section explains how to configure OpenLB to make use of the installed functionalities.

9.2.1.1. CUDA on Windows with WSL

As mentioned in the chapter about install dependencies (Section 9.1.3), the preferred approach for OpenLB on Windows is to use the Windows Subsystem for Linux (WSL). The following was written with the assumption that OpenLB has been successfully set up on WSL with Ubuntu.

The following specifications are needed to get CUDA running via WSL:

- Windows 10 version 21H2 or higher
- CUDA compatible Nvidia graphics card
- WSL 2 with a glibc-based distribution (e.g. Ubuntu)

To find out which Windows version exactly you're using, open up the `run` dialog box in Windows and type in the command `winver`, which will display a pop-up window similar to the one below:

In order to find out what graphics card you have and whether it is compatible with CUDA, open the up the Windows `run` dialog and type in the command `dxdiag`, which will open the `DirectX Diagnostic Tool`. Under the tab `Render`, it will display the information regarding your graphics card. In the example picture of the `DirectX Diagnostic Tool` below, the graphics card is a `NVIDIA GeForce GTX 1650`. NVIDIA provides the information on which graphics card is compatible with CUDA on their website (<https://developer.nvidia.com/cuda-gpus>).

CUDA is only supported on version 2 of the Windows Subsystem for Linux (WSL). To confirm which version of WSL is installed, open the Windows PowerShell with administrator rights and type in the command

```
wsl --list --verbose
```

This will display which Linux distribution and which version of WSL is currently installed. The output should look similar to the following:

```
PS C:\Windows> wsl --list --verbose
  NAME      STATE          VERSION
* Ubuntu    Stopped        1
```



Figure 9.1.: Pop-up window displaying the exact version and build of Windows.

In this example the distribution that is installed is `Ubuntu` and the WSL version is 1. Upgrading to the necessary version 2 can be done by typing

```
wsl --set-version Ubuntu 2
```

into the `PowerShell` terminal. Note that when using a different distribution for WSL, the command has to be adjusted accordingly.

An error might occur claiming that a certain hard-link target does not exist. This means that there is software installed on WSL that collides with the update. The error message will provide the path of the non-existing hard-link, which will be a hint onto which package causes this error. In the WSL terminal, the command

```
sudo apt list --installed
```

will give an overview over all the installed packages. The conflicting package can then be removed with

```
sudo apt-get remove [PACKAGE-NAME]
```

Once the package has been removed, WSL can be upgraded. On a successful upgrade, we should receive a message that the conversion is complete and we can verify the version with the

```
wsl --list --verbose
```

command. The conflicting package can then be reinstalled.

In order for WSL to have access to the GPU hardware, virtual GPU needs to be enabled on Windows. This can be done by installing an appropriate driver on Windows. It should not be necessary to install any




Figure 9.2.: The Render tab of the DirectX Diagnostic Tool.

device drivers on WSL itself. It is even highly suggested not to do so, since any installation of a driver on WSL itself might override the functionality provided by the driver that is installed onto Windows. As of the writing of this guide, the most recent NVIDIA drivers automatically support virtual GPU for WSL. The newest driver can be directly downloaded from the NVIDIA website (<https://www.nvidia.com/download/index.aspx>). The website offers drop down lists to specify what product type, device, operating system, etc. the driver is needed for. Once the most recent driver is installed, we can install the CUDA toolkit on WSL.

The following commands typed into the WSL terminal will install the Nvidia CUDA toolkit on WSL (Ubuntu):

```
sudo apt-key del 7fa2af80

wget https://developer.download.nvidia.com/compute/cuda/repos/
wsl-ubuntu/x86_64/cuda-wsl-ubuntu.pin

sudo mv cuda-wsl-ubuntu.pin /etc/apt/preferences.d/
cuda-repository-pin-600
```

```

sudo apt-key adv --fetch-keys https://developer.download.nvidia.com/
compute/cuda/repos/wsl-ubuntu/
x86_64/3bf863cc.pub

sudo add-apt-repository 'deb https://developer.download.nvidia.com/
compute/cuda/repos/wsl-ubuntu/x86_64/ /'

sudo apt-get update

sudo apt-get -y install cuda

```

If the NVIDIA CUDA compiler is correctly installed, the command

```
nvcc --version
```

will reply with a message similar to the following:

```

1  nvcc: NVIDIA (R) Cuda compiler driver
2  Copyright (c) 2005-2022 NVIDIA Corporation
3  Built on Mon_Oct_24_19:12:58_PDT_2022
4  Cuda compilation tools, release 12.0, V12.0.76
5  Build cuda_12.0.r12.0/compiler.31968024_0

```

Listing 9.1.: Version details of an installed Cuda compiler

To check the versions of CUDA and the driver, the command

```
nvidia-smi
```

will respond with the NVIDIA System Management Interface, displaying various information about the installed GPUs (see Figure 9.3). The CUDA toolkit should now be properly installed and working.

9.2.1.2. CUDA on Linux

Before installing the CUDA toolkit on Linux, typing the command

```
lspci | grep -i nvidia
```

can confirm that the GPU is CUDA-capable.

To install the CUDA toolkit on Linux, visit the the NVIDIA website and choose the fitting operating system, architecture, distribution, as well as the preferred installation type for your system (<https://developer.nvidia.com/cuda-toolkit>). The website will then provide you with the correct commands with which you can install the CUDA toolkit on your Linux system.

After the installation of the toolkit, the environment variables need to be set:

```
export PATH=/usr/local/cuda-12.0/bin${PATH:+:${PATH}}
```

```
Thu Jan 26 15:00:16 2023
```

NVIDIA-SMI 527.92.01				Driver Version: 528.02		CUDA Version: 12.0	
GPU	Name	Persistence-M	Bus-Id	Disp.A	Volatile	Uncorr. ECC	
Fan	Temp	Perf	Pwr:Usage/Cap	Memory-Usage	GPU-Util	Compute M.	
						MIG M.	
0	NVIDIA GeForce ...	On	00000000:01:00.0	Off		N/A	
N/A	38C	P8	1W / 50W	0MiB / 4096MiB	0%	Default	N/A

Processes:							
GPU	GI	CI	PID	Type	Process name	GPU Memory	
	ID	ID				Usage	
0	N/A	N/A				N/A	

Figure 9.3.: The NVIDIA System Management Interface

If the installation was done with a run file, the `LD_LIBRARY_PATH` variable has to be set, as well. The following command sets this variable on a 64-bit system. The command for 32-bit systems is almost identical: `lib64` has to be exchanged for `lib`:

```
export LD_LIBRARY_PATH=/usr/local/cuda-12.0/lib64\
    ${LD_LIBRARY_PATH:+:${LD_LIBRARY_PATH}}
```

If a different install path or version of the CUDA toolkit has been chosen during the installation process, both commands above have to be altered accordingly. To confirm that the installation has been successful, use the commands

```
nvcc --version
```

and

```
nvidia-smi
```

If the CUDA toolkit has been installed correctly, an output similar to those shown in Listing 9.1 and Figure 9.3 respectively.

9.2.1.3. OpenMPI

To have the functionality of MPI in combination with CUDA, there are several CUDA-aware MPI implementations available. In this guide we will describe the installation of the open-source implementation OpenMPI in four steps:

1. Download the desired OpenMPI version from the website (<https://www.open-mpi.org/software/>). As of the writing of this section, the most current stable release version was `openmpi-4.1.5.tar.bz2`.

2. In your Linux (or WSL for Windows) terminal, move to the folder where the file was saved to and extract the downloaded package via the command

```
tar -jxf openmpi-4.1.5.tar.bz2
```

3. Change into this new directory to configure, compile and install OpenMPI with the following commands:

```
./configure --prefix=$HOME/opt/openmpi
              --with-cuda=/usr/local/cuda-12.0/include
make all
make install
```

Note that the path following `-prefix=` is the path we wish to install OpenMPI in and the path following `-with-cuda=` is the location of the `include` folder of your CUDA installation. These paths might be different depending on the users choices.

4. Change the environment variables with the following two commands in the Linux or WSL terminal:

```
echo "export PATH=$PATH:$HOME/opt/openmpi/bin" >> $HOME/.bashrc
echo "export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:
      $HOME/opt/openmpi/lib" \ >> $HOME/.bashrc
```

Once again the path for OpenMPI might be different, depending on where the software was installed.

To see whether the installation of OpenMPI was successful, we can enter the command

```
mpi_info --parsable -l 9
      --all | grep mpi_built_with_cuda_support:value
```

If the installation was done successfully, the terminal should respond with the output `true`.

9.2.1.4. Utilizing CUDA in OpenLB

The root directory contains a folder named `config`, in which several build config examples can be found. The `config.mk` makefile of the root directory can be replaced with the makefile that suits the current needs (e.g. using only the GPU, using the GPU with MPI, using CPU with MPI, etc.). Each example makefile also includes instructions. Make a backup of the current `config.mk` in the root directory and replace it with a copy of the makefile `gpu_only` found in the `config` folder. After renaming `gpu_only` to `config.mk`, we open the file and check the value of `CUDA_ARCH`: This value might have to be changed, depending on your graphics card and its architecture. The file `rules.mk` in the root directory contains a table that shows which architecture goes with which value:

##	CUDA Architecture	Version	
##	-----+-----		
##	Fermi	20	
##	Kepler	30, 35, 37	
##	Maxwell	50, 52, 53	
##	Pascal	60, 61, 62	
##	Volta	70, 72	
##	Turing	75	
##	Ampere	80, 86, 87	

Another table on the internet (<https://en.wikipedia.org/wiki/CUDA>) shows which graphics card corresponds to which architecture. This guide used the GTX 1650 as an example for the graphics card. The following picture shows that the GTX 1650 corresponds to the Turing architecture, so the value of `CUDA_ARCH` has to be set to 75 in both `config.mk` and `rules.mk` files. After saving the changes of

7.5	Turing	TU102, TU104, TU106, TU116, TU117	NVIDIA TITAN RTX, GeForce RTX 2080 Ti, RTX 2080 Super, RTX 2080, RTX 2070 Super, RTX 2070, RTX 2060 Super, RTX 2060 12GB, RTX 2060, GeForce GTX 1660 Ti, GTX 1660 Super, GTX 1660, GTX 1650 Super, GTX 1650, MX550, MX450

Figure 9.4.: Table containing Nvidia GPUs with the Turing Microarchitecture.

`CUDA_ARCH` in both `config.mk` and `rules.mk`, the `config.mk` can be compiled via the command `make clean; make` in your WSL (or Linux) terminal. It is now possible to compile and execute one of the GPU-enabled OpenLB examples with CUDA support.

10. Step by Step: Using OpenLB for Applications

The general way of functioning in OpenLB follows a generic path. The following structure is maintained throughout every OpenLB application example, to provide a common structure and guide beginners.

1st Step: Initialization The converter between physical and lattice units is set in this step. It is also defined, where the simulation data is stored and which lattice type is used.

2nd Step: Prepare geometry The geometry is acquired, either from another file (a `.stl` file) or from defining indicator functions. Then, the mesh is created and initialized based on the given geometry. This consists of classifying voxels with material numbers, according to the kind of voxels they are: an inner voxel containing fluid ruled by the fluid dynamics will have a different number than a voxel on the inflow with conditions on its velocity. The function `prepareGeometry` is called for these tasks. Further, the mesh is distributed over the threads to establish good scaling properties.

3rd Step: Prepare lattice According to the material numbers of the geometry, the lattice dynamics are set here. This step characterizes the collision model and boundary behavior. The choices depend on whether a force is acting or not, the use of single relaxation time (BGK) or multiple relaxation times (MRT), the simulation dimension (it can also be a 2D model), whether compressible or incompressible fluid is considered, and the number of neighboring voxels chosen. By the creation of a computing grid, the `SuperLattice`, the allocation of the required data is done as well.

4th Step: Main loop with timer The timer is initialized and started, then a loop over all time steps `it` starts the simulation, during which the functions `setBoundaryValues`, `collideAndStream` and `getResults` (the 5th, 6th, and 7th step, respectively) are called repeatedly until a maximum of iterations is reached, or the simulation has converged. At the end, the timer is stopped and the summary is printed to the console.

5th Step: Definition of initial and boundary conditions The first of the three important functions called during the loop, `setBoundaryValues`, sets the slowly increasing inflow boundary condition. Since the boundary is time dependent, this happens in the main loop. In some applications, the boundaries stay the same during the whole simulation and the function doesn't need to do anything after the very first iteration.

6th Step: Collide and stream execution Another function `collideAndStream` is called each iteration step, to perform the collision and the streaming step. If more than one lattice is used, the function is called for each lattice separately.

7th Step: Computation and output of results At the end of each iteration step, the function `getResults` is called, which creates console output, `.ppm` files or `.vti` files of the results at certain time steps. The ideal is to get the relevant simulation data with functors and thus facilitate the post processing significantly. By passing the converter and the time step, the frequency of writing or displaying data can be chosen easily. In many applications, the console output is required more often than the VTK data.

10.1. Lesson 1: Getting Started - Sketch of Application

This section presents example `bstep2d` that can be found in the recent release of OpenLB. This example simulates a flow over a backward-facing step and serves as an illustration of OpenLB and its features. In order to execute the simulation and get some results, download and unpack OpenLB on (preferably) a Linux system, see Section 9.1.1. Then, generate a executable file by compiling the program through the command `make`. Finally, launch the simulation by `./bstep2d` and observe the terminal output, see Section 6.6.

A few lines are invariably the same for all OpenLB applications, see Listing 10.1.

```
1 #include "olb2D.h"
2 #include "olb2D.hh"
3
4 using namespace olb;           // OpenLB namespaces
5 using namespace olb::descriptors; //
6
7 using T = FLOATING_POINT_TYPE;
8 using DESCRIPTOR = D2Q9<>;
```

Listing 10.1.: Framework of an OpenLB program. Fundamental properties of the simulation are defined here.

Line 1: The header file `olb2D.h` includes definitions for the whole 2D code present in the release. In the same way, access to 3D code is obtained by including the file `olb3D.h`.

Line 2: Most OpenLB code depends on template parameters. Therefore, it cannot be compiled in advance, and needs to be integrated “as is” into your programs via the file `olb2D.hh` or `olb3D.hh` respectively.

Line 4: All OpenLB code is contained in the namespace `olb`. The descriptors have an own namespace and define the lattice arrangement, e.g. `D2Q9` or `D3Q19`.

Line 7: Choice of precision for floating point arithmetic. The default type `FLOATING_POINT_TYPE` is defined in the `config.mk` file and usually equals `float` or `double`. Any other floating point type can be used, including built-in types and user-defined types which are implemented through a C++ class.

Line 8: Choice of a lattice descriptor. Lattice descriptors specify not only which lattice (see Figure 1.1 for exemplary velocity sets) are employed, but is also are used to compute the size of various dependent

fields such as force vectors.

The next code presents a brief overview about the structure of an OpenLB application, see Listing 10.2. It aims rather to introduce and guidelines the beginners, than explain the classes and methods in depth. Details on the shown functions can be found in the source code, this means in the `bstep2d.cpp` file, as well as in the following chapters.

```
1 SuperGeometry<T,2> prepareGeometry(LBconverter<T> const& converter)
2 {
3     // create Cuboids and assign them to threads
4     // create SuperGeometry
5     // set material numbers
6     return superGeometry;
7 }
8 void prepareLattice(...)
9 {
10    // set dynamics for fluid and boundary lattices
11    // set initial values, rho and u
12 }
13 void setBoundaryValues(...)
14 {
15    // set Poiseuille velocity profile at inflow
16    // increase inflow velocity slowly over time
17 }
18 void getResults(...)
19 {
20    // write simulation data do vtk files and terminal
21 }
22
23 int main(int argc, char* argv[])
24 {
25    // === 1st Step: Initialization ===
26    olbInit( &argc, &argv );
27    singleton::directories().setOutputDir( "./tmp/" ); // set output directory
28    OstreamManager clout( std::cout, "main" );
29
30    UnitConverterFromResolutionAndRelaxationTime<T, DESCRIPTOR> converter(
31        (T)    N,                // resolution
32        (T)    relaxationTime,    // relaxation time
33        (T)    charL,            // charPhysLength: reference length of simulation geometry
34        (T)    1.,               // charPhysVelocity: maximal/highest expected velocity
35                                // during simulation in __m / s__
36        (T)    1./19230.76923,    // physViscosity: physical kinematic viscosity in __m^2 /
37                                // s__
38        (T)    1.                // physDensity: physical density in __kg / m^3__
39    );
40
41    // Prints the converter log as console output
42    converter.print();
43    // Writes the converter log in a file
44    converter.write("bstep2d");
45
46    // === 2nd Step: Prepare Geometry ===
47    // Instantiation of a superGeometry
48    SuperGeometry<T,2> superGeometry( prepareGeometry(converter) );
```

```

48 // === 3rd Step: Prepare Lattice ===
49 SuperLattice<T,DESCRIPTOR> sLattice( superGeometry );
50 BGKdynamics<T,DESCRIPTOR> bulkDynamics (
51     converter.getLatticeRelaxationFrequency(),
52     instances::getBulkMomenta<T,DESCRIPTOR>()
53 );
54
55 //prepare Lattice and set boundaryConditions
56 prepareLattice( converter, sLattice, bulkDynamics, superGeometry );
57
58 // instantiate reusable functors
59 SuperPlaneIntegralFluxVelocity2D<T> velocityFlux( sLattice,
60     converter,
61     superGeometry,
62     {lengthStep/2., heightInlet / 2.},
63     {0., 1.} );
64
65 SuperPlaneIntegralFluxPressure2D<T> pressureFlux( sLattice,
66     converter,
67     superGeometry,
68     {lengthStep/2., heightInlet / 2. },
69     {0., 1.} );
70
71 // === 4th Step: Main Loop with Timer ===
72 clout << "starting simulation..." << std::endl;
73 Timer<T> timer( converter.getLatticeTime( maxPhysT ), superGeometry.getStatistics().
74     getNvoxel() );
75 timer.start();
76
77 for ( std::size_t iT = 0; iT < converter.getLatticeTime( maxPhysT ); ++iT ) {
78     // === 5th Step: Definition of Initial and Boundary Conditions ===
79     setBoundaryValues( converter, sLattice, iT, superGeometry );
80     // === 6th Step: Collide and Stream Execution ===
81     sLattice.collideAndStream();
82     // === 7th Step: Computation and Output of the Results ===
83     getResults( sLattice, converter, iT, superGeometry, timer, velocityFlux,
84         pressureFlux );
85 }
86 timer.stop();
87 timer.printSummary();
88 }

```

Listing 10.2.: A brief overview of a typically OpenLB application, `bstep2d`. Details on the specific functions can be found in the following chapters.

10.2. Lesson 2: Define and Use Boundary Conditions

The current OpenLB release offers a wide range of boundary conditions for the implementation of pressure and velocity boundaries. They support boundaries that are aligned with the numerical grid, and also implement proper corner nodes in 2D and 3D, and edge nodes that connect two plane boundaries in 3D. The choice of a boundary condition is conceptually separated from the definition of the location of

boundary nodes. It is therefore possible to modify the choice of the boundary condition by changing a single instruction in a program. An overview of the available boundary conditions is given by [22].

The new boundary condition system utilizes free floating functions and doesn't require a class structure. Consequently the following classes are obsolete in the current release:

```
sOn/OffLatticeBoundaryConditionXD,  
OnLatticeBoundaryConditionXD,  
(Off)BoundaryConditionInstantiatorXD and  
Regularized/InterpolationBoundaryManagerXD.
```

Key Features of the new system are:

- Free floating design that allows for general functions like "setBoundary" to be used in multiple boundary Conditions.
- Overall slimmer design with fewer function calls and fewer loops through the block domain.
- `MomentaVector` and `dynamicsVector` is stored in `/src/core/blockLatticeStructure3D.h`
- Uncluttered function call design, which makes it easier to create new boundary conditions

The new boundaries are similarly named as in `addSlipBoundary` becomes `setSlipBoundary` in the new system. For example, if you want to use a slip boundary condition:

Define dynamics Keep in mind that `sLattice.defineDynamics` is the same for the old and new boundary system.

Define your boundary type In this case it is the `slipBoundary`. To set the boundary, call the fitting `setBoundaryCondition` function in the following manner inside your `prepareLattice` function:

- `setSlipBoundary<T, DESCRIPTOR> ("superLattice", "superGeometry", "MaterialNumber");`
- The difference between the old and the new system is that every `boundaryCondition` needs to have the `SuperLattice` as an argument. The `latticeRelaxationFrequency omega` is usually called as the second argument.

Define initial conditions

- on-lattice: `sLattice.defineRhoU(..)`
- off-lattice: `sLattice.defineRho(..), sLattice.defineUBouzidi(..)`

Define boundary values

- on-lattice: `sLattice.defineU(..)`
- off-lattice: `sLattice.defineUBouzidi(..)`

With the help of this system, one can treat local and non-local boundary conditions the same way. Furthermore, they can be used both for sequential and parallel program execution, as it is shown in Lesson 10. The mechanism behind this is explained in Lesson 7. The bottom line is that both local and non-local boundary conditions instantiate a special dynamics object and assign it to boundary cells. Non-local boundaries additionally instantiate post-processing objects which take care of non-local aspects of the algorithm.

10.3. Lesson 3: UnitConverter - Lattice and Physical Units

Fluid flow problems are usually given in a system of metric units. For example consider a cylinder of diameter 3 cm in a fluid channel with average inflow velocity of 4 m. The fluid has a kinematic viscosity of $0.001 \text{ m}^2 \text{ s}^{-1}$. The value of interest is the pressure difference measured in Pa at the front and the back of the cylinder (with respect to the flow direction). However, the variables used in a LB simulation live in a system of lattice units, in which the distance between two lattice cells and the time interval between two iteration steps are chosen to be unity. Therefore, when setting up a simulation, a conversion directive has to be defined that takes care of scaling variables from physical units into lattice units and *vice versa*.

In OpenLB, all these conversions are handled by a class called `UnitConverter`, see Listing 10.3. An instance of the `UnitConverter` is generated with desired discretization parameters and reference values in SI units. It provides a set of conversion functions to enable a fast and easy way to convert between physical and lattice units. In addition, it gives information about the parameters of the fluid flow simulation, such as the Reynolds number or the relaxation parameter ω .

Let's have a closer look at the input parameters: The reference values represent characteristic quantities of the fluid flow problem. In this example, it is suitable to choose the cylinder's diameter as characteristic length and the average inflow speed as characteristic velocity. Furthermore, two discretization parameters, namely the grid size Δx in m and time step size Δt in s are provided to the converter. From these reference values and discretization parameters, all the conversion factors and the relaxation time τ are calculated.

Due to the fact, that there are stability bounds for the relaxation time and the maximum occurring lattice velocity, one does not usually chose Δx and Δt , but sets stable and accurate values for any two out of resolution, relaxation time or characteristic (maximum) lattice velocity. To make that easily available for the user of OpenLB, there are different constructors for the `UnitConverter` class:

```
UnitConverterFromRelaxationTimeAndLatticeVelocity,  
UnitConverterFromResolutionAndLatticeVelocity,  
UnitConverterFromResolutionAndRelaxationTime.
```

Once the converter is initialized, its methods can be used to convert various quantities such as velocity, time, force or pressure. The function for the latter helps us to evaluate the pressure drop in our example problem, as shown in the the following code snippet:

```
1 UnitConverterFromResolutionAndRelaxationTime<T, DESCRIPTOR> const converter(  
2     int {N}, // resolution: number of voxels per charPhysL  
3     (T) 0.53, // latticeRelaxationTime: relaxation time, have to be greater  
           than 0.5!  
4     (T) 0.1, // charPhysLength: reference length of simulation geometry  
5     (T) 0.2, // charPhysVelocity: maximal/highest expected velocity during  
           simulation in __m / s__  
6     (T) 0.2*2.*0.05/Re, // physViscosity: physical kinematic viscosity in __m^2 / s__  
7     (T) 1.0, // physDensity: physical density in __kg / m^3__  
8 );  
9 // Prints the converter log as console output  
10 converter.print();  
11 // Writes the converter log in a file  
12 converter.write("converterLogFile");  
13 // conversion from seconds to iteration steps and vice-versa
```

```

14 int iT = converter.getLatticeTime(maxPhysT);
15 T sec = converter.getPhysTime(iT);
16 <...> simulation
17 <...> evaluation of latticeRho at the back and the front of the cylinder
18 T latticePressureFront = latticeRhoFront / descriptors::invCs2<T,DESCRIPTOR>();
19 T latticePressureBack = latticeRhoBack / descriptors::invCs2<T,DESCRIPTOR>();
20 T pressureDrop = converter.getPhysPressure(latticePressureFront)
21                 - converter.getPhysPressure(latticePressureBack);

```

Listing 10.3.: Use of UnitConverter in a 3D problem.

Line 1–7: Instantiate an `UnitConverter` object and specify discretization parameters as well as characteristic values.

Line 9: Write simulation parameters and conversion factors to terminal.

Line 11: Write simulation parameters and conversion factors in a logfile.

Line 13 and 14: The conversion from physical units (seconds) to discrete ones (time steps) is managed by the converter.

Line 17–20: The converter automatically calculates the pressure values from the local density.

10.4. Lesson 4: Extract Data From a Simulation

When the collision step is executed, the value of the density and the velocity are computed internally, in order to evaluate the equilibrium distribution. Those macroscopic variables are however interesting for the OpenLB end-user as well, and it would be a shame to simply neglect their value after use. These values are accessed through the method `getStatistics()` of a `BlockLattice`:

T `lattice.getStatistics().getAverageRho()` Returns average density evaluated during the previous collision step.

T `lattice.getStatistics().getAverageEnergy()` Returns half the average velocity norm evaluated during the previous collision step.

T `lattice.getStatistics().getMaxU()` Returns maximum value of the velocity norm evaluated during the previous collision step.

Often, the information provided by the statistics of a lattice is not sufficient, and more generally numerical results are required. To do this, you can get data cell-by-cell from the `BlockLatticeXD` and `SuperLatticeXD` through functors, see Chapter 2.5. Functors act on the underlying lattice and process its data to relevant macroscopic units, e.g. density, velocity, stress, flux, pressure and drag. Functors provide an `operator()` that instead of access stored data, computes every time it is called the data. Since OpenLB version 0.8, the concept of functors unfolds not only for postprocessing, but also for boundary conditions and the generation of geometry, see Chapter 2.5. In Listing 10.4 it is shown, how to extract data out of a `SuperLattice` named `sLattice` and an `SuperGeometry3D` named `sGeometry`. The data format is a legal `.vtk` file, that can be processed further with ParaView.

```

1 // generate the writer object
2 SuperVTMWriter3D<T> vtmWriter("bstep3d");
3 // write every 0.2 seconds
4 if (iT==converter.getLatticeTime(0.2)) {
5     // create functors
6     SuperLatticeGeometry3D<T,DESCRIPTOR> geometry(sLattice, sGeometry);
7     SuperLatticeCuboid3D<T,DESCRIPTOR> cuboid(sLattice);
8     SuperLatticeRank3D<T,DESCRIPTOR> rank(sLattice);
9     // write functors to file system, vtk formata
10    vtmWriter.write(geometry);
11    vtmWriter.write(cuboid);
12    vtmWriter.write(rank);
13 }

```

Listing 10.4.: Extract simulation data to `vtk` file format.

As before mentioned, OpenLB provides functors for a bunch of data, see Listing 10.5. More details about writing simulation data can be found in Chapter 6.

```

1 // Create the functors by only passing lattice and converter
2 SuperLatticePhysVelocity3D<T,DESCRIPTOR> velocity(sLattice, converter);
3 SuperLatticePhysPressure3D<T,DESCRIPTOR> pressure(sLattice, converter);
4 // Create functor that corresponds to material numbers
5 SuperLatticeGeometry3D<T,DESCRIPTOR> geometry(sLattice, superGeometry);

```

Listing 10.5.: Code example for creating velocity, pressure and geometry functors.

The most straightforward and convenient way of visualizing simulation data is to produce a 2D snapshot of a scalar valued functor. This is done through the `BlockReduction3D2D`, which puts a plane into arbitrary 3D functors. Afterwards, this plane can be easily written to a image file. OpenLB creates images of PPM format as shown in Listing 10.6.

```

1 // velocity is an application:  $R^3 \rightarrow R^3$ 
2 // an image in its very basic sense is an application:  $R^2 \rightarrow R$ 
3
4 // transformation of data is presented below
5 // get velocity functor
6 SuperLatticePhysVelocity3D<T,DESCRIPTOR> velocity(sLattice, converter);
7 // get scalar valued functor by applying the point wise l2 norm
8 SuperEuklidNorm3D<T> normVel( velocity );
9 // put a plane with normal (0,0,1) in the 3 dimensional data
10 BlockReduction3D2D<T> planeReduction( normVel, {0, 0, 1} );
11 BlockGifWriter<T> gifWriter;
12 // write ppm image to file system
13 gifWriter.write( planeReduction, iT, "vel" );

```

Listing 10.6.: Create a PPM image out of a 3D velocity functor.

This image writer provides *in situ* visualization which, in contrast to the `VTkWriter`, produces smaller data sets that can be interpreted immediately without requiring other software.

10.5. Lesson 5: Convergence Check

The class `ValueTracer` checks for time-convergence of a given scalar ϕ . The convergence is reached when the standard deviation σ of the monitored value ϕ is smaller than a given residuum ϵ times the expected value $\bar{\phi}$.

$$\sigma(\phi) = \sqrt{\frac{1}{N+1} \sum_{i=0}^N (\phi_i - \bar{\phi})^2} < \epsilon \bar{\phi} \quad (10.1)$$

The expected value ϕ is the average over the last N time steps with $\phi_i := \phi(t^* - i\Delta t)$ where t^* is the current time step and Δt denotes the time step size.

$$\bar{\phi} = \frac{1}{N+1} \sum_{i=0}^N \phi_i \quad (10.2)$$

The value N should be chosen as a problem specific time period. As an example `charT = charL/charU` and `N = converter.getLatticeTime(charT)`. To initialize a `ValueTracer` object use:

```
1 util::ValueTracer<T> converge( numberTimeSteps, residuum );
```

Listing 10.7.: Create a PPM image out of a 3D velocity functor.

For example, to check for convergence with a residuum of $\epsilon = 10^{-5}$ every physical second:

```
1 util::ValueTracer<T> converge( converter.getLatticeTime(1.0), 1e-5 );
```

Listing 10.8.: Create a PPM image out of a 3D velocity functor.

It is required to pass the monitored value to the `ValueTracer` object every time steps by:

```
1 for (iT = 0; iT < maxIter; ++iT) {
2     ...
3     converge.takeValue( monitoredValue, isVerbose );
4     ..
5 }
```

Listing 10.9.: Create a PPM image out of a 3D velocity functor.

If you like to print average value and its standard derivation every number of time steps chosen during initialization set `isVerbose` to true otherwise choose false. It is good idea to choose average energy as monitored value:

```
1 converge.takeValue(SLattice.getStatistics().getAverageEnergy(), true);
```

Listing 10.10.: Create a PPM image out of a 3D velocity functor.

Do something like the following in the time loop:

```
1 if (converge.hasConverged()) {
2     cout << "Simulation converged." << std::endl;
3     break;
```

4 }

Listing 10.11.: Create a PPM image out of a 3D velocity functor.

10.6. Lesson 6: Use an External Force

In simulations the dynamics of a fluid are often driven by some kind of externally imposed force field. In order to optimize memory access and to minimize cache-misses, the value of this force can be stored right alongside the cell's population values. This is achieved by specifying additional fields in the lattice descriptor (see Sections 2.1 and 4.4).

At this point we want to consider a time- and space-independent external force as a basic example. Listing 10.12 shows how such an external force can be defined for all cells of a certain material number.

```
1 // Define constant force
2 AnalyticalConst2D<T,T> force(
3     8.0 * converter.getLatticeViscosity()
4     * converter.getCharLatticeVelocity()
5     / ( Ly*Ly ), // x-component of the force
6     0.0);        // y-component of the force
7
8 // Initialize force for materials 1 and 2
9 superLattice.defineField<FORCE>(
10    superGeometry.getMaterialIndicator({1, 2}), force);
```

Listing 10.12.: Define a constant external force

This code was adapted from `examples/laminar/poiseuille2d` where just such a constant force is used to drive the channel flow. Note that the underlying `D2Q9` descriptor's field list must contain a `FORCE` field in order for `defineField<FORCE>` to work.

The command `SuperLattice::defineField` provides just one of many template methods one can use to work with descriptor fields. For example, we can read and write a cell's force field as follows:

```
1 // Get a reference to the memory location of a cell's force vector
2 FieldPtr<T,DESCRIPTOR,FORCE> force = cell.template getFieldPointer<FORCE>();
3 // Read a cell's force vector as an OpenLB vector value
4 Vector<double,2> force = cell.template getField<FORCE>();
5 // Set a cell's force vector to zero
6 cell.template setField<FORCE>(Vector<double,2>(0.0, 0.0));
```

Note that these methods work the same way for any other field that might be declared by a cell's specific descriptor.

10.7. Lesson 7: Understand Genericity in OpenLB

OpenLB is a framework for the implementation of lattice Boltzmann algorithms. Although most of the code shipped with the distribution is about fluid dynamics, it is open to various types of physical mod-

els. Generally speaking, a model which makes use of OpenLB must be formulated in terms of the “local collision followed by nearest-neighbor streaming” philosophy. A current restriction to OpenLB is that the streaming step can only include nearest neighbors: there is no possibility to include larger neighborhoods within the modular framework of the library, i.e. without tampering with OpenLB source code. Except for this restriction, one is completely free to define the topology of the neighborhood of cells, to implement an arbitrary local collision step, and to add non-local corrections for the implementation of, say, a boundary condition.

To reach this level of genericity, OpenLB distinguishes between non-modifiable core components, which you’ll always use as they are, and modular extensions. As far as these extensions are concerned, you have the choice to use default implementations that are part of OpenLB or to write your own. As a scientific developer, concentrating on these, usually quite short, extensions means that you can concentrate on the physics of your model instead of technical implementation details. By respecting this concept of modularity, you can automatically take advantage of all structural additions to OpenLB. In the current release, the most important addition is parallelism: you can run your code in parallel without (or almost without) having to care about parallelism and MPI.

The most important non-modifiable components are the lattice and the cell. You can configure their behavior, but you are not expected to write a new class which inherits from or replaces the lattice or the cell. Lattices are offered in different flavors, most of which inherit from a common interface `BlockStructureXD`. The most common lattice is the regular `BlockLatticeXD`, which is replaced by the `SuperLatticeXD` for parallel applications and for memory-saving applications when faced with irregular domain boundaries. An alternative choice for parallelism and memory savings is the `CuboidStructureXD`, which does not inherit from `BlockStructureXD`, but instead allows for more general constructs.

The modular extensions are classes that customize the behavior of core-components. An important extension of this kind is the lattice descriptor. This specifies the number of particle populations contained in a cell, and defines the lattice constants and lattice velocities, which are used to specify the neighborhood relation between a cell and its nearest neighbors. The lattice descriptor can also be used to require additional allocation of memory on a cell for external scalars, such as a force field. The integration of a lattice descriptor in a lattice happens via a template mechanism of C++. This mechanism takes place statically, i.e. before program execution, and avoids the potential efficiency loss of a dynamic, object-oriented approach. Furthermore, template specialization is used to optimize the OpenLB code specifically for some types of lattices. Because of the template-based approach, a lattice descriptor needs not inherit from some interface. Instead, you are free to simply implement a new class, inspired from the default descriptors in the files `core/latticeDescriptors.h` and `core/latticeDescriptor.hh`.

The dynamics executed by a cell are implemented through a mechanism of dynamic (run-time) genericity. In this way, the dynamics can be different from one cell to another, and can change during program execution. There are two mechanisms of this type in OpenLB, one to implement local dynamics, and one for non-local dynamics. To implement local dynamics, one needs to write a new class which inherits the interface of the abstract class `Dynamics`. The purpose of this class is to specify the nature of the collision step, as well as other important information (for example, how to compute the velocity moments on a

cell). For non-local dynamics, a so-called post-processor needs to be implemented and integrated into a `BlockLatticeXD` through a call to the method `addPostProcessorXD`. This terminology can be somewhat confusing, because the term “post-processing” is used in the CFD community in the context of data analysis at the end of a simulation. In OpenLB, a post-processor is an operator which is applied to the lattice after each streaming step. Thus, the time-evolution of an OpenLB lattice consists of three steps: (1) local collision, (2) nearest-neighbor streaming, and (3) non-local postprocessing. Implementing the dynamics of a cell through a postprocessor is usually less efficient than when the mechanism of the `Dynamics` classes is used. It is therefore important to respect the spirit of the lattice Boltzmann method and to express the collision as a local operation whenever possible.

10.8. Lesson 8: Use Checkpointing for Long Duration Simulations

All types of data in OpenLB can be stored in a file or loaded from a file. This includes the data of a `BlockLatticeXD` and the data of a `ScalarFieldXD` or a `TensorFieldXD`. All these classes implement the interface `Serializable<T>`. This guarantees that they can transform their content into a data stream of type `T`, or read from such a stream. Serialization and unserialization of data is mainly used for file access, but it can be applied to different aims, such as copying data between two objects of different type. The data is stored in the ascii-based binary format `Base64`. Although `Base64`-encoded data requires 25% more storage space than when a pure binary format is used, this approach was chosen in OpenLB to enhance compatibility of the code between platforms. Saving and loading data is invoked by calling the `save` and `load` method on the object to be serialized. These methods take the filename as an optional (but recommended) argument, as shown below:

```

1 int nx, ny;
2 <...> initialization of nx and ny
3 BlockLattice<T,DESCRIPTOR> lattice(nx, ny);
4 // load data from a previous simulation
5 lattice.load("simulation.checkpoint");
6 <...> run the simulation
7 // save data for security, to be able to take up
8 // the simulation at this point later
9 lattice.save("simulation.checkpoint");

```

Listing 10.13.: Store and load the state of the simulation.

Checkpointing is also illustrated in the example programs `bstep2D` and `bstep3D` (Section 8.4.1).

10.9. Lesson 9: Run Your Programs on a Parallel Machine

OpenLB programs can be executed on a parallel machine with distributed memory, based on MPI. To compile an OpenLB program for parallel execution, modify the file named `config.mk`, found in the OpenLB

root directory, by changing the respective lines: `CXX := mpic++`, and `PARALLEL_MODE := MPI`. The modified lines are shown in Listing 10.14. Execute `make clean` and `make cleanbuild` within the desired program directory to eliminate previously compiled libraries, and recompile the program by executing the `make` command. To run the program in parallel, use the command `mpirun -np 2 ./cavity2d`. Here `-np 2` specifies the number of processors to be used.

```

1 CXX          := mpic++
2 ...
3 PARALLEL_MODE := MPI

```

Listing 10.14.: Edited config.mk for MPI-parallel programs.

10.10. Lesson 10: Work with Indicators

Many of the methods covered up until this point accepted geometry and material number arguments to define their working domain. This can lead to repetition and code that is harder to read than necessary. An alternative to e.g. setting up bulk dynamics using raw material numbers is available in the form of indicator functors.

In fact most of the material number accepting operations we have covered so far use generic lattice indicators under the hood, specifically `SuperIndicatorMaterial3D`.

```

1 superLattice.defineDynamics(superGeometry, 1, &bulkDynamics);
2 superLattice.defineDynamics(superGeometry, 3, &bulkDynamics);
3 superLattice.defineDynamics(superGeometry, 4, &bulkDynamics);
4 // is equivalent to:
5 SuperIndicatorMaterial3D<T> bulkIndicator({1, 3, 4});
6 superLattice.defineDynamics(bulkIndicator, &bulkDynamics);

```

Listing 10.15.: Indicator usage example

This can be further abstracted using `SuperGeometry3D`'s indicator factory:

```

1 auto bulkIndicator = superGeometry.getMaterialIndicator({1, 3, 4});
2 superLattice.defineDynamics(bulkIndicator, &bulkDynamics);

```

Listing 10.16.: Indicator usage in `bstep3d`

The advantage of this pattern is that we explicitly named materials 1, 3 and 4 as bulk materials and can reuse the indicator whenever we operate on bulk cells:

```

1 superLattice.defineRhoU(bulkIndicator, rho, u);
2 superLattice.iniEquilibrium(bulkIndicator, rho, u);

```

Listing 10.17.: Indicator reuse in `bstep3d`

This way the bulk material domain is defined in a central place which will come in handy should we need to change them in the future.

Note that for one-off usage this can be written even more compactly:

```
1 superLattice.defineDynamics(  
2   superGeometry.getMaterialIndicator({1, 3, 4}), &bulkDynamics);
```

Listing 10.18.: Inline indicator usage

This pattern of using indicators instead of raw material numbers is available for all material number accepting methods of `SuperGeometryXD`.

The methods themselves support arbitrary `SuperIndicatorFXD` instances and as such are not restricted to material indicators.

10.11. Alternative Approach: Using a Solver Class

Quite a lot of program components are similar for each OpenLB application: e.g. the collide and stream loop is part of every simulation. The concept of a solver class is meant to perform such steps automatically, s.t. the user only has to define those steps which are specific for his/her application. Moreover, a generic interface shall be given for other programs (e.g. launch from python scripts or execution of optimization routines). For both purposes, this is work in progress and more improvements and functionalities are under development. In the following, the parts of an OpenLB program in solver style are explained. These steps are also illustrated by the examples `cavity2dSolver` and `porousPlate3dSolver`.

10.11.1. Structure of an OpenLB Simulation in Solver Style

10.11.1.1. Parameter handling

In order to allow flexible interfaces to other programs, all parameters which are needed for simulation and interface are stored publicly in structs. For different groups of parameters (e.g. simulation/ output/ stationarity), different structs are used.

For `Simulation`, `Output` and `Stationarity`, basic versions containing the essential parameters are given by `SimulationBase`, `OutputBase`, `StationarityBase`, respectively. These can be supplemented by inheritance.

More parameter structs could be added for individualization. For instance, a `Results` struct could be used to save simulation results.

10.11.1.2. List parameter structs and lattices

A map of parameter structs with corresponding names is defined as a `meta::map`. Similarly, a map of lattice names and descriptors is defined. Some typical names are provided at `src/solver/names.h`; a list which is intended to be extended for individualization. The two maps are then given to the solver class as template parameters.

10.11.1.3. Definition of a solver class

Many standard routines for simulation are implemented in the existing class `LbSolver`. It is templated w.r.t. maps of parameters and lattices and should therefore fit to most application cases. However, some steps (like the definition of the geometry) depend on the application and have to be defined for each application. Therefore, an application-specific solver class is created as a child class of `LbSolver`. It has to implement the methods `prepareGeometry`, `prepareLattices`, `setInitialValues` and `setBoundaryValues`, similar to the classical app structure. Moreover, methods `getResults`, `computeResults`, `writeImages`, `writeVTK` and `writeGnuplot` can be defined if such output is desired. They are all called automatically during construction/ simulation. The access to the parameter structs works with the tags defined above: e.g., `this->parameters (Simulation()).maxTime` gives the maximal simulation time (which is a member of the struct `SimulationBase`). Similarly, we find access to super geometry and super lattices via `this->geometry()` and `this-> lattice(LatticeName())`, respectively. An automatic check, whether the simulation became stationary, is executed if a parameter struct with tag `Stationarity` is available (and the corresponding struct inherits from `StationarityBase`).

10.11.1.4. Main method

First, instances of the parameter structs and the solver class are constructed. This can be done classically, using the constructors (cf. example `porousPlate3dSolver`), or, if XML-reading has been implemented for all parameter structs, with the create-from XML-interface (cf. example `cavity2dSolver`).

Secondly, the `solve()` method of the solver instance is called in order to run the simulation.

10.11.2. Set up an Application in Solver Style

In order to set up your own OpenLB application in solver style, the following steps should be followed:

- Select parameter structs. You can use existing ones or inherit from them/ define them completely new. The simulation parameters are expected to inherit from `SimulationParameters` and provide a unit converter. The output parameters should inherit from `OutputParameters`. You are free to add more parameter structs (e.g. for simulation results) yourself.
- Define a solver class. It should inherit publicly from the `LbSolver` class and implement the missing virtual methods like `prepareGeometry`.
- Define the main method. Either construct instances of the parameter structs and the solver class or use the create-from-xml interface. Then call the `solve()` method and possibly perform postprocessing.

10.11.3. Parameter Explanation and Reading from XML

Introduction

In this short overview the relevant parameters for an app in solver style are listed with the respective names for using an xml-file for the input. The documentation is divided into two different subsections,

which are the different main parts of the xml-file:

10.11.3.1. Application

10.11.3.2. Output

In each subsection, the different parameters are explained via a table of the following form:

Parameter	Name(type)	Class (file)	Default value	Explanation: (if available, all) <i>possibilities</i>
-----------	------------	--------------	---------------	---

The explanation of each column is as follows:

- **Parameter:** Name of the parameter in the xml-file
- **Name (type):** Name of the parameter in the source code and its data type in brackets. Besides the common data types the abbreviations S and T are used for template parameters.
- **Class (file):** name of the class in which the parameter is stored and name of the file in which the class is defined.
- **Default value:** Is this parameter essential or optional or unused?
If a parameter is optional, it is not needed to be defined. Then, the default value can be seen in this column.
If a parameter is so important that without it the program has to exit, it is labeled as *EXIT*.
Some parameters are indicated with **unused** which means, that the parameter is read but not used afterwards.
- **Explanation:** (if available, all) *possibilities*: Brief description and explanation of the parameter. Some parameters have different possibilities for their definition. In this case, all available possibilities are also offered in **bold** type letters, e.g.:
ad for *OptiCaseAD* or
dual for *OptiCaseDual* or
adTest for *OptiCaseADTest*

The arrangement of the parameters in the xml-file has the following structures:

```
<Param>
  <name of the subsection in this documentation>
    <superordinate term of the parameter>
      <name of the parameter> value of the parameter </name of the parameter>
    </superordinate term of the parameter>
  </name of the subsection in this documentation>
</Param>
```

example:

```
<Param>
  <Application>
    <Discretization>
      <Resolution> 128 </Resolution>
    </Discretization>
  </Application>
</Param>
```

10.11.3.1. Application

In the following, all parameters for the general setup of the application are explained.

Parameter	Name(type)	Class (file)	Default value	Explanation: (if available, all) <i>possibilities</i>
<i>Name</i>	name (std::string)	OutputGeneral (solverParameters.h)	"unnamed"	Output name <i>Name.dat</i> for information of Unit-Converter in the <i>tmp</i> folder
<i>OlbDir</i>	olbDir (std::string)	OutputGeneral (solverParameters.h)	"../..../"	Defines the trail for the OpenLB directory
<i>Pressure-Filter</i>	pressureFilter-On (bool)	SimulationBase (solverParameters.h)	false	Weights for moments computing, if (this-> pressureFilterOn): _lattice-> stripeOffDensityOffset (_lattice-> getStatistics(). getAverageRho() - (T) 1) in lbSolver.hh
<i>Discretization:</i>				
<i>Resolution</i>	_resolution (int)	UnitConverter (unitConverter.hh)		Defines the resolution number of the simulation area
<i>Lattice-Relaxation-Time</i>	_lattice-Relaxation-Time (T)	UnitConverter (unitConverter.hh)		Defines the lattice relaxation time and thereby the <i>physicalDeltaT</i> , <i>_conversionTime</i> respectively
<i>PhysParameters:</i>				
<i>CharPhys-Length</i>	_charPhys-Length (T)	UnitConverter (unitConverter.hh)		Defines the characteristic physical length
<i>CharPhys-Velocity</i>	_charPhys-Velocity (T)	UnitConverter (unitConverter.hh)		Defines the characteristic physical velocity
<i>PhysDensity</i>	_physDensity (T)	UnitConverter (unitConverter.hh)		Defines the physical density
<i>CharPhys-Pressure</i>	_charPhys-Pressure (T)	UnitConverter (unitConverter.hh)		Defines the characteristic physical pressure
<i>Phys-Viscosity</i>	_phys-Viscosity (T)	UnitConverter (unitConverter.hh)		Defines the physical viscosity

<i>PhysMax-Time</i>	maxTime (S)	SimulationBase (solverParameters.h)		Defines the maximal simulation time in seconds
<i>StartUp-Time</i>	startUpTime (S)	SimulationBase (solverParameters.h)		Defines the start time until which time the simulation is started up. From then on, the <i>convergence criterion</i> is checked (in solver3D.hh) and determines the boundaries for <i>defineU</i> in primalMode.
<i>Boundary-Value-Update-Time</i>	_phys-Boundary-ValueUpdate-Time (S)	SimulationBase (solverParameters.h)	PhysMax-Time/100	Defines the boundary value update time
<i>mesh:</i>				
<i>noCuboids-PerProcess</i>	noC (int)	SimulationBase (solverParameters.h)	1, unused	Defines the number of cuboids per process
<i>ConvergenceCheck:</i>				
<i>Type</i>	convergence-Type[0] (std::string)	Stationarity (solverParameters.h)	"Max-Lattice-Velocity"	Define which quantity is regarded when checking the convergence. Alternative values: AverageEnergy, AverageRho
<i>Interval</i>	phys-Interval[0] (BaseType<T>)	Stationarity (solverParameters.h)		Time interval in physical in which the values are checked for convergence
<i>Residuum</i>	epsilon[0] (BaseType<T>)	Stationarity (solverParameters.h)		sensitivity for the convergence check

10.11.3.2. Output

In the following, all parameters for the output produced by the application are explained.

Parameter	Name(type)	Class (file)	Default value	Explanation: (if available, all) <i>possibilities</i>
<i>Multi-Output</i>	Used as (bool)	testDomain3d.cpp		Choose 0 or 1 for clout.setMultiOutput, output for all processors. Needs to be manually added to the olbInit call
<i>OutputDir</i>	outputDir (std::string)	OutputGeneral (solverParameters.h)	"/temp/"	Choose your output directory
<i>PrintLog-Converter</i>	printLog-Converter (bool)	OutputGeneral (solverParameters.h)	true	If true, call writeLog-Converter(), i.e. print the converter
<i>Log:</i>				
<i>SaveTime</i>	logT (BaseType<T>)	OutputGeneral (solverParameters.h)	EXIT	Choose physical time, write the data in the terminal each x seconds until PhysMaxTime; printLog(iT) in lbSolver.hh
<i>VerboseLog</i>	verbose (bool)	OutputGeneral (solverParameters.h)	1 (true)	Show statements in the terminal, if true
<i>VisualizationVTK, VisualizationImages and VisualizationGnuplot:</i>				
<i>Output</i>	out (bool)	OutputPlot (solverParameters.h)	false	States whether Output or not
<i>FileName</i>	filename (std::string)	OutputPlot (solverParameters.h)	"unnamed"	name of the outputfile
<i>SaveTime</i>	saveTime (S)	OutputPlot (solverParameters.h)	EXIT	Choose physical time, write the data each x seconds until PhysMaxTime

<i>Timer:</i>				
<i>PrintMode</i>	timerPrint- Mode (int)	OutputGeneral (solverParameters.h)	2	mode of the display style passed to printStep() of Timer instance, 0 for single-line layout, usable for data extraction as csv, 1 for single-line layout, (not conform with out- put rules), 2 for double-line layout, (not conform with out- put rules), 3 for performance output only

References

References Involving OpenLB (Articles)

- [1] F. Bukreev, F. Raichle, H. Nirschl, and M. J. Krause. “Simulation of adsorption processes on moving particles based on an Euler–Euler description using a lattice Boltzmann discretization”. In: *Chemical Engineering Science* 270 (2023), p. 118485. DOI: <https://doi.org/10.1016/j.ces.2023.118485>.
- [2] D. Dapelo, S. Simonis, M. J. Krause, and J. Bridgeman. “Lattice–Boltzmann coupled models for advection–diffusion flow on a wide range of Peclet numbers”. In: *Journal of Computational Science* 51 (2021), p. 101363. DOI: <https://doi.org/10.1016/j.jocs.2021.101363>.
- [3] R. Ditscherlein, O. Furat, E. Löwer, R. Mehnert, R. Trunk, T. Leissner, M. J. Krause, V. Schmidt, and U. Peuker. “PARROT: A Pilot Study on the Open Access Provision of Particle-Discrete Tomographic Datasets”. In: *Microscopy and Microanalysis* 28.2 (2022), pp. 350–360. DOI: [10.1017/S143192762101391X](https://doi.org/10.1017/S143192762101391X).
- [4] M. Gaedtke, T. Hoffmann, V. Reinhardt, G. Thäter, H. Nirschl, and M. J. Krause. “Flow and heat transfer simulation with a thermal large eddy lattice Boltzmann method in an annular gap with an inner rotating cylinder”. In: *International Journal of Modern Physics C* 30.02n03 (2019), p. 1950013. DOI: [10.1142/S012918311950013X](https://doi.org/10.1142/S012918311950013X).
- [5] M. Gaedtke, S. Wachter, S. Kunkel, S. Sonnick, M. Rädle, H. Nirschl, and M. J. Krause. “Numerical study on the application of vacuum insulation panels and a latent heat storage for refrigerated vehicles with a large Eddy lattice Boltzmann method”. In: *Heat and Mass Transfer* (2019), pp. 1–13. DOI: [10.1007/s00231-019-02753-4](https://doi.org/10.1007/s00231-019-02753-4).
- [6] M. Gaedtke, S. Wachter, M. Rädle, H. Nirschl, and M. J. Krause. “Application of a lattice Boltzmann method combined with a Smagorinsky turbulence model to spatially resolved heat flux inside a refrigerated vehicle”. In: *Computers & Mathematics with Applications* 76.10 (Nov. 2018), pp. 2315–2329. DOI: <https://doi.org/10.1016/j.camwa.2018.08.018>.
- [7] N. Hafen, A. Dittler, and M. J. Krause. “Simulation of particulate matter structure detachment from surfaces of wall-flow filters applying lattice Boltzmann methods”. In: *Computers & Fluids* 239 (2022), p. 105381. DOI: <https://doi.org/10.1016/j.compfluid.2022.105381>.
- [8] N. Hafen, J. Thieringer, J. Meyer, M. J. Krause, and A. Dittler. “Numerical investigation of detachment and transport of particulate structures in wall-flow filters using lattice Boltzmann methods”. In: *Journal of Fluid Mechanics* 956 (2023), A30. DOI: [10.1017/jfm.2023.35](https://doi.org/10.1017/jfm.2023.35).

- [9] M. Haussmann, A. Claro Berreta, G. Lipeme Kouyi, N. Riviere, H. Nirschl, and M. J. Krause. "Large-eddy simulation coupled with wall models for turbulent channel flows at high Reynolds numbers with a lattice Boltzmann method –Application to Coriolis mass flowmeter". In: *Computers & Mathematics with Applications* 78.10 (2019), pp. 3285–3302. DOI: <https://doi.org/10.1016/j.camwa.2019.04.033>.
- [10] M. Haussmann, N. Hafen, F. Raichle, R. Trunk, H. Nirschl, and M. J. Krause. "Galilean invariance study on different lattice Boltzmann fluid–solid interface approaches for vortex-induced vibrations". In: *Computers & Mathematics with Applications* 80.5 (2020), pp. 671–691. DOI: <https://doi.org/10.1016/j.camwa.2020.04.022>.
- [11] M. Haussmann, P. Reinshaus, S. Simonis, H. Nirschl, and M. J. Krause. "Fluid–Structure Interaction Simulation of a Coriolis Mass Flowmeter Using a Lattice Boltzmann Method". In: *Fluids* 6.4 (2021). DOI: [10.3390/fluids6040167](https://doi.org/10.3390/fluids6040167).
- [12] M. Haussmann, S. Simonis, H. Nirschl, and M. J. Krause. "Direct numerical simulation of decaying homogeneous isotropic turbulence – numerical experiments on stability, consistency and accuracy of distinct lattice Boltzmann methods". In: *International Journal of Modern Physics C* 30.09 (2019), p. 1950074. DOI: [10.1142/S0129183119500748](https://doi.org/10.1142/S0129183119500748).
- [13] M. Haussmann et al. "Evaluation of a Near-Wall-Modeled Large Eddy Lattice Boltzmann Method for the Analysis of Complex Flows Relevant to IC Engines". In: *Computation* 8.43 (2020). DOI: [10.3390/computation8020043](https://doi.org/10.3390/computation8020043).
- [14] T. Henn, G. Thäter, W. Dörfler, H. Nirschl, and M. J. Krause. "Parallel dilute particulate flow simulations in the human nasal cavity". In: *Computers & Fluids* 124 (2016), pp. 197–207. DOI: <http://dx.doi.org/10.1016/j.compfluid.2015.08.002>.
- [15] S. Höcker, R. Trunk, W. Dörfler, and M. J. Krause. "Towards the simulations of inertial dense particulate flows with a volume-averaged lattice Boltzmann method". In: *Computers & Fluids* 166 (2018), pp. 152–162. DOI: <https://doi.org/10.1016/j.compfluid.2018.02.011>.
- [16] J. Jeßberger, J. E. Marquardt, L. Heim, J. Mangold, F. Bukreev, and M. J. Krause. "Optimization of a Micromixer with Automatic Differentiation". In: *Fluids* 7.5 (2022). DOI: [10.3390/fluids7050144](https://doi.org/10.3390/fluids7050144).
- [17] F. Klemens, S. Schuhmann, R. Balbierer, G. Guthausen, H. Nirschl, G. Thäter, and M. J. Krause. "Noise reduction of flow MRI measurements using a lattice Boltzmann based topology optimisation approach". In: *Computers & Fluids* 197 (2020), p. 104391. DOI: <https://doi.org/10.1016/j.compfluid.2019.104391>.
- [18] F. Klemens, S. Schuhmann, G. Guthausen, G. Thäter, and M. J. Krause. "CFD-MRI: A coupled measurement and simulation approach for accurate fluid flow characterisation and domain identification". In: *Computers & Fluids* 166 (2018), pp. 218–224. DOI: <https://doi.org/10.1016/j.compfluid.2018.02.022>.
- [19] K. Klemens, B. Förster, M. Dorn, G. Thäter, and M. J. Krause. "Solving fluid flow domain identification problems with adjoint lattice Boltzmann methods". In: *Computers & Mathematics with Applications* 79.1 (2020), pp. 17–33. DOI: <https://doi.org/10.1016/j.camwa.2018.07.010>.
- [20] M. J. Krause, F. Klemens, T. Henn, R. Trunk, and R. Nirschl. "Particle flow simulations with homogenised lattice Boltzmann methods". In: *Particuology* 34 (2017), pp. 1–13. DOI: <http://doi.org/10.1016/j.partic.2016.11.001>.

- [21] M. J. Krause, G. Thäter, and V. Heuveline. “Adjoint-based Fluid Flow Control and Optimisation with Lattice Boltzmann Methods”. In: *Computers & Mathematics with Applications* 65.6 (2013), pp. 945–960. DOI: [10.1016/j.camwa.2012.08.007](https://doi.org/10.1016/j.camwa.2012.08.007).
- [22] M. J. Krause et al. “OpenLB–Open source lattice Boltzmann code”. In: *Computers & Mathematics with Applications* 81 (2021), pp. 258–288. DOI: <https://doi.org/10.1016/j.camwa.2020.04.033>.
- [23] A. Kummerländer, M. Dorn, M. Frank, and M. J. Krause. “Implicit propagation of directly addressed grids in lattice Boltzmann methods”. In: *Concurrency and Computation: Practice and Experience* 35.8 (Mar. 2023), e7509. DOI: <https://doi.org/10.1002/cpe.7509>.
- [24] L. de Luca Xavier Augusto, J. Ross-Jones, G. Cantarelli Lopes, P. Tronville, J. Silveira Gonçalves, M. Rädle, and M. J. Krause. “Microfiber Filter Performance Prediction using a Lattice-Boltzmann Method”. In: *Communications in Computational Physics* 23.4 (2018), pp. 910–931. DOI: [10.4208/cicp.OA-2016-0180](https://doi.org/10.4208/cicp.OA-2016-0180).
- [25] M.-L. Maier, T. Henn, G. Thäter, H. Nirschl, and M. J. Krause. “Towards Validated Multiscale Simulation with a Two-Way Coupled LBM and DEM”. In: *Chemical Engineering & Technology* 40.9 (2017), pp. 1591–1598. DOI: [10.1002/ceat.201600547](https://doi.org/10.1002/ceat.201600547).
- [26] M.-L. Maier, S. Milles, S. Schuhmann, G. Guthausen, H. Nirschl, and M. J. Krause. “Fluid flow simulations verified by measurements to investigate adsorption processes in a static mixer”. In: *Computers & Mathematics with Applications* 76.11 (2018), pp. 2744–2757. DOI: <https://doi.org/10.1016/j.camwa.2018.08.066>.
- [27] M.-L. Maier, R. A. Patel, N. I. Prasianakis, S. V. Churakov, H. Nirschl, and M. J. Krause. “Coupling of multiscale lattice Boltzmann discrete–element method for reactive particle fluid flows”. In: *Phys. Rev. E* 103 (3 Mar. 2021), p. 033306. DOI: [10.1103/PhysRevE.103.033306](https://doi.org/10.1103/PhysRevE.103.033306).
- [28] J. E. Marquardt, C. Arlt, R. Trunk, M. Franzreb, and M. J. Krause. “Numerical and experimental examination of the retention of magnetic nanoparticles in magnetic chromatography”. In: *Computers & Mathematics with Applications* 89 (2021), pp. 34–43. DOI: <https://doi.org/10.1016/j.camwa.2021.02.010>.
- [29] J. E. Marquardt, U. J. Römer, H. Nirschl, and M. J. Krause. “A discrete contact model for complex arbitrary-shaped convex geometries”. In: *Particuology* 80 (2023), pp. 180–191. DOI: <https://doi.org/10.1016/j.partic.2022.12.005>.
- [30] A. Mink, K. Schediwy, C. Posten, H. Nirschl, S. Simonis, and M. J. Krause. “Comprehensive Computational Model for Coupled Fluid Flow, Mass Transfer, and Light Supply in Tubular Photobioreactors Equipped with Glass Sponges”. In: *Energies* 15.20 (2022). DOI: [10.3390/en15207671](https://doi.org/10.3390/en15207671).
- [31] A. Mink, G. Thäter, H. Nirschl, and M. J. Krause. “A 3D Lattice Boltzmann Method for Light Simulation in Participating Media”. In: *Journal of Computational Science* 17, Part 2 (2016), pp. 431–437. DOI: [http://dx.doi.org/10.1016/j.jocs.2016.03.014](https://doi.org/10.1016/j.jocs.2016.03.014).
- [32] H. Mirzaee, T. Henn, M. J. Krause, L. Goubergrits, C. Schumann, M. Neugebauer, T. Kuehne, T. Preusser, and A. Hennemuth. “MRI-based computational hemodynamics in patients with aortic coarctation using the lattice Boltzmann methods: Clinical validation study”. In: *Journal of Magnetic Resonance Imaging* 45.1 (2016), pp. 139–146. DOI: [10.1002/jmri.25366](https://doi.org/10.1002/jmri.25366).

- [33] M. Mohrhard, G. Thäter, J. Bludau, B. Horvat, and M. J. Krause. “An Auto-Vecotorization Friendly Parallel Lattice Boltzmann Streaming Scheme for Direct Addressing”. In: *Computers & Fluids* 181 (2019), pp. 1–7. DOI: <https://doi.org/10.1016/j.compfluid.2019.01.001>.
- [34] P. Nathen, D. Gaudlitz, M. J. Krause, and N. Adams. “On the Stability and Accuracy of the BGK, MRT and RLB Boltzmann Schemes for the Simulation of Turbulent Flows”. In: *Communications in Computational Physics* 23.3 (2018), pp. 846–876. DOI: [10.4208/cicp.OA-2016-0229](https://doi.org/10.4208/cicp.OA-2016-0229).
- [35] F. Reinke, N. Hafen, M. Haussmann, M. Novosel, M. J. Krause, and A. Dittler. “Applied Geometry Optimization of an Innovative 3D-Printed Wet-Scrubber Nozzle with a Lattice Boltzmann Method”. In: *Chemie Ingenieur Technik* 94.3 (2022), pp. 348–355. DOI: <https://doi.org/10.1002/cite.202100151>.
- [36] J. Ross-Jones, M. Gaedtke, S. Sonnick, M. Rädle, H. Nirschl, and M. J. Krause. “Conjugate heat transfer through nano scale porous media to optimize vacuum insulation panels with lattice Boltzmann methods”. In: *Computers & Mathematics with Applications* 77 (2019), pp. 209–221. DOI: <https://doi.org/10.1016/j.camwa.2018.09.023>.
- [37] S. Simonis, M. Frank, and M. J. Krause. “On relaxation systems and their relation to discrete velocity Boltzmann models for scalar advection–diffusion equations”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 378.2175 (2020), p. 20190400. DOI: [10.1098/rsta.2019.0400](https://doi.org/10.1098/rsta.2019.0400).
- [38] S. Simonis, M. Frank, and M. J. Krause. “Constructing relaxation systems for lattice Boltzmann methods”. In: *Applied Mathematics Letters* 137 (2023), p. 108484. DOI: <https://doi.org/10.1016/j.aml.2022.108484>.
- [39] S. Simonis, M. Haussmann, L. Kronberg, W. Dörfler, and M. J. Krause. “Linear and brute force stability of orthogonal moment multiple–relaxation–time lattice Boltzmann methods applied to homogeneous isotropic turbulence”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 379.2208 (2021), p. 20200405. DOI: [10.1098/rsta.2020.0405](https://doi.org/10.1098/rsta.2020.0405).
- [40] S. Simonis, J. Nguyen, S. J. Avis, W. Dörfler, and M. J. Krause. “Binary fluid flow simulations with free energy lattice Boltzmann methods”. In: *Discrete and Continuous Dynamical Systems - S* (2023). DOI: [10.3934/dcdss.2023069](https://doi.org/10.3934/dcdss.2023069).
- [41] S. Simonis, D. Oberle, M. Gaedtke, P. Jenny, and M. J. Krause. “Temporal large eddy simulation with lattice Boltzmann methods”. In: *Journal of Computational Physics* 454 (2022), p. 110991. DOI: <https://doi.org/10.1016/j.jcp.2022.110991>.
- [42] M. Siodlaczek, M. Gaedtke, S. Simonis, M. Schweiker, M. Homma, and M. J. Krause. “Numerical evaluation of thermal comfort using a large eddy lattice Boltzmann method”. In: *Building and Environment* 192 (2021), p. 107618. DOI: <https://doi.org/10.1016/j.buildenv.2021.107618>.
- [43] J. Thieringer, N. Hafen, J. Meyer, M. J. Krause, and A. Dittler. “Investigation of the Rearrangement of Reactive–Inert Particulate Structures in a Single Channel of a Wall–Flow Filter”. In: *Separations* 9.8 (2022). DOI: [10.3390/separations9080195](https://doi.org/10.3390/separations9080195).
- [44] R. Trunk, C. Bretl, G. Thäter, H. Nirschl, M. Dorn, and M. J. Krause. “A Study on Shape-Dependent Settling of Single Particles with Equal Volume Using Surface Resolved Simulations”. In: *Computation* 9.4 (2021). DOI: [10.3390/computation9040040](https://doi.org/10.3390/computation9040040).

- [45] R. Trunk, T. Henn, W. Dörfler, H. Nirschl, and M. J. Krause. “Inertial Dilute Particulate Fluid Flow Simulations with an Euler-Euler Lattice Boltzmann Method”. In: *Journal of Computational Science* 17, Part 2 (2016), pp. 438–445. DOI: <http://dx.doi.org/10.1016/j.jocs.2016.03.013>.
- [46] R. Trunk, J. Marquardt, G. Thäter, H. Nirschl, and M. J. Krause. “Towards the Simulation of arbitrarily shaped 3D particles using a homogenised lattice Boltzmann method”. In: *Computers & Fluids* 172 (2018), pp. 621–631. DOI: <https://doi.org/10.1016/j.compfluid.2018.02.027>.
- [47] R. Trunk, T. Weckerle, N. Hafen, G. Thäter, H. Nirschl, and M. J. Krause. “Revisiting the Homogenized Lattice Boltzmann Method with Applications on Particulate Flows”. In: *Computation* 9.2 (2021). DOI: [10.3390/computation9020011](https://doi.org/10.3390/computation9020011).
- [48] A. Zarth, F. Klemens, G. Thäter, and M. J. Krause. “Towards shape optimisation of fluid flows using lattice Boltzmann methods and automatic differentiation”. In: *Computers & Mathematics with Applications* 90 (2021), pp. 46–54. DOI: <https://doi.org/10.1016/j.camwa.2021.02.016>.

References Involving OpenLB (Proceedings)

- [49] C. Bretl, R. Trunk, H. Nirschl, G. Thäter, M. Dorn, and M. J. Krause. “Preliminary Study of Particle Settling Behaviour by Shape Parameters via Lattice Boltzmann Simulations”. In: *High Performance Computing in Science and Engineering 20*. Ed. by W. E. Nagel, D. H. Kröner, and M. M. Resch. Cham: Springer International Publishing, 2021, pp. 245–259.
- [50] A. Kummerländer, F. Bukreev, S. Berg, M. Dorn, and M. J. Krause. “Advances in Computational Process Engineering using Lattice Boltzmann Methods on High Performance Computers for Solving Fluid Flow Problems”. In: *High Performance Computing in Science and Engineering '21*. accepted (to appear). Springer, 2023.
- [51] A. Kummerländer and M. J. Krause. “Research Software Engineering in OpenLB: Refactoring a Legacy Code to State-Of-The-Art Performance”. In: *deRSE23 Conference for Research Software Engineering in Germany*. Zenodo, Feb. 2023. DOI: [10.5281/zenodo.7662082](https://doi.org/10.5281/zenodo.7662082).
- [52] N. Nadim, T. Chandratilleke, and M. J. Krause. “LBM-LES Modelling of Low Reynolds Number Turbulent Flow Over NACA0012 Aerofoil”. English. In: *Fluid-Structure-Sound Interactions and Control*. Ed. by Y. Zhou, A. Lucey, Y. Liu, and L. Huang. Lecture Notes in Mechanical Engineering. Springer Berlin Heidelberg, 2016, pp. 205–210. DOI: [10.1007/978-3-662-48868-3_33](https://doi.org/10.1007/978-3-662-48868-3_33).
- [53] P. Nathen, D. Gaudlitz, M. J. Krause, and J. Kratzke. “An extension of the Lattice Boltzmann Method for simulating turbulent flows around rotating geometries of arbitrary shape”. In: *21st AIAA Computational Fluid Dynamics Conference, San Diego*. American Institute of Aeronautics and Astronautics, 2013. DOI: [doi:10.2514/6.2013-2573](https://doi.org/10.2514/6.2013-2573).
- [54] S. Simonis and M. J. Krause. “Forschungsnahe Lehre unter Pandemiebedingungen”. In: *Mitteilungen der Deutschen Mathematiker-Vereinigung* 30.1 (2022), pp. 43–45. DOI: [doi:10.1515/dmvm-2022-0015](https://doi.org/10.1515/dmvm-2022-0015).

References Involving OpenLB (Software)

- [55] M. J. Krause, T. Henn, L. Baron, J. Kratzke, J. Fietz, and T. Dornieden. *OpenLB Release 0.7: Open Source Lattice Boltzmann Code*. Version 0.7. Feb. 2012. DOI: [10.5281/zenodo.3625936](https://doi.org/10.5281/zenodo.3625936).
- [56] M. J. Krause, T. Henn, L. Baron, A. Mink, P. Weisbrod, P. Nathen, and G. Zahnd. *OpenLB Release 0.8: Open Source Lattice Boltzmann Code*. Version 0.8. Nov. 2013. DOI: [10.5281/zenodo.3625938](https://doi.org/10.5281/zenodo.3625938).
- [57] M. J. Krause, T. Henn, A. Mink, R. Trunk, P. Weisbrod, P. Nathen, F. Klemens, and M.-L. Maier. *OpenLB Release 0.9: Open Source Lattice Boltzmann Code*. Version 0.9. Mar. 2015. DOI: [10.5281/zenodo.3625941](https://doi.org/10.5281/zenodo.3625941).
- [58] M. J. Krause, T. Henn, A. Mink, R. Trunk, P. Weisbrod, P. Nathen, F. Klemens, and M.-L. Maier. *OpenLB Release 1.0: Open Source Lattice Boltzmann Code*. Version 1.0. Mar. 2016. DOI: [10.5281/zenodo.3625943](https://doi.org/10.5281/zenodo.3625943).
- [59] M. J. Krause, A. Mink, R. Trunk, F. Klemens, M.-L. Maier, M. Mohrhard, A. Claro Barreto, M. Haußmann, M. Gaedtke, and J. Ross-Jones. *OpenLB Release 1.2: Open Source Lattice Boltzmann Code*. Version 1.2. Feb. 2018. DOI: [10.5281/zenodo.3625960](https://doi.org/10.5281/zenodo.3625960).
- [60] M. J. Krause, S. Zimny, T. Henn, and J. Fietz. *OpenLB Release 0.6: Open Source Lattice Boltzmann Code*. Version 0.6. May 2011. DOI: [10.5281/zenodo.3625929](https://doi.org/10.5281/zenodo.3625929).
- [61] M. J. Krause et al. *OpenLB Release 1.1: Open Source Lattice Boltzmann Code*. Version 1.1. Apr. 2017. DOI: [10.5281/zenodo.3625955](https://doi.org/10.5281/zenodo.3625955).
- [62] M. J. Krause et al. *OpenLB Release 1.3: Open Source Lattice Boltzmann Code*. Version 1.3. May 2019. DOI: [10.5281/zenodo.3625967](https://doi.org/10.5281/zenodo.3625967).
- [63] M. J. Krause et al. *OpenLB Release 1.4: Open Source Lattice Boltzmann Code*. Version 1.4. Nov. 2020. DOI: [10.5281/zenodo.4279263](https://doi.org/10.5281/zenodo.4279263).
- [64] A. Kummerländer et al. *OpenLB Release 1.5: Open Source Lattice Boltzmann Code*. Version 1.5. Apr. 2022. DOI: [10.5281/zenodo.6469606](https://doi.org/10.5281/zenodo.6469606).
- [65] A. Kummerländer et al. *OpenLB Release 1.6: Open Source Lattice Boltzmann Code*. Version 1.6. Apr. 2023. DOI: [10.5281/zenodo.7773497](https://doi.org/10.5281/zenodo.7773497).
- [66] J. Latt and M. J. Krause. *OpenLB Release 0.3: Open Source Lattice Boltzmann Code*. Version 0.3. July 2007. DOI: [10.5281/zenodo.3625765](https://doi.org/10.5281/zenodo.3625765).
- [67] J. Latt and M. J. Krause. *OpenLB Release 0.4: Open Source Lattice Boltzmann Code*. Version 0.4. Jan. 2008. DOI: [10.5281/zenodo.3625909](https://doi.org/10.5281/zenodo.3625909).
- [68] J. Latt, M. J. Krause, O. Malaspinas, and B. Stahl. *OpenLB Release 0.5: Open Source Lattice Boltzmann Code*. Version 0.5. May 2008. DOI: [10.5281/zenodo.3625925](https://doi.org/10.5281/zenodo.3625925).

Other References

- [69] LB model with adjustable speed of sound. Technical report. <http://www.openlb.net/tech-reports>.
- [70] The Paraview project. <http://www.paraview.org>.
- [71] The OpenMPI project. <http://www.openmpi.org>.
- [72] The VTK data format documentation. <http://www.vtk.org/VTK/img/file-formats.pdf>.

- [73] FreeCAD: An Open Source parametric 3D CAD modeler. <https://www.freecadweb.org/>.
- [74] Configuring OpenLB on MacOS. Technical Report. <http://www.openlb.net/tech-reports>.
- [75] Installing OpenLB in Windows 10/Ubuntu bash. Technical Report. <http://www.openlb.net/tech-reports>.
- [76] S. Ansumali. "Minimal kinetic modeling of hydrodynamics". PhD thesis. Swiss Federal Institute of Technology Zurich, 2004.
- [77] B. F. Armaly, F. Durst, J. C. F. Pereira, and B. Schönung. "Experimental and theoretical investigation of backward-facing step flow". In: *Journal of Fluid Mechanics* 127 (1983), pp. 473–496. DOI: [10.1017/S0022112083002839](https://doi.org/10.1017/S0022112083002839).
- [78] P. L. Bhatnagar, E. P. Gross, and M. Krook. "A model for collision processes in gases. I. Small amplitude processes in charged and neutral one-component systems". In: *Physical Review* 94.3 (1954), pp. 511–525. DOI: [10.1103/PhysRev.94.511](https://doi.org/10.1103/PhysRev.94.511).
- [79] T. Borrvall and J. Petersson. "Topology optimization of fluids in Stokes flow". In: *International Journal for Numerical Methods in Fluids* 41.1 (2003), pp. 77–107. DOI: [10.1002/flid.426](https://doi.org/10.1002/flid.426).
- [80] M. Bouzidi, M. Firdaouss, and P. Lallemand. "Momentum transfer of a Boltzmann-lattice fluid with boundaries". In: *Physics of Fluids* 13.11 (2001), pp. 3452–3459. DOI: [10.1063/1.1399290](https://doi.org/10.1063/1.1399290).
- [81] M. E. Brachet, D. I. Meiron, S. A. Orszag, B. Nickel, R. H. Morf, and U. Frisch. "Small-scale structure of the Taylor–Green vortex". In: *Journal of Fluid Mechanics* 130 (1983), pp. 411–452.
- [82] H. Brinkman. "A calculation of the viscous force exerted by a flowing fluid on a dense swarm of particles". English. In: *Applied Scientific Research* 1.1 (1949), pp. 27–34. DOI: [10.1007/BF02120313](https://doi.org/10.1007/BF02120313).
- [83] H. Brinkman. "On the permeability of media consisting of closely packed porous particles". English. In: *Applied Scientific Research* 1.1 (1949), pp. 81–86. DOI: [10.1007/BF02120318](https://doi.org/10.1007/BF02120318).
- [84] A. Caiazzo and M. Junk. "Asymptotic analysis of lattice Boltzmann methods for flow-rigid body interaction". In: *Progress in Computational Physics* 3 (2013), p. 91.
- [85] S. Chen and G. D. Doolen. "Lattice Boltzmann Method for Fluid Flows". In: *Ann. Rev. Fluid Mech.* 30 (1998), pp. 329–364.
- [86] B. Chopard, A. Dupuis, A. Masselot, and P. Luthi. "Cellular Automata and Lattice Boltzmann techniques: an approach to model and simulate complex systems". In: *Adv. Compl. Sys.* 5 (2002), pp. 103–246. DOI: [10.1142/S0219525902000602](https://doi.org/10.1142/S0219525902000602).
- [87] D. d’Humières, I. Ginzburg, M. Krafczyk, P. Lallemand, and L.-S. Luo. "Multiple-relaxation-time lattice Boltzmann models in three dimensions". In: *Phil. Trans. R. Soc. Lond. A* 360 (2002), pp. 437–451.
- [88] T. Dornieden. "Optimierung von Strömungsgebieten mit adjungierten Lattice Boltzmann Methoden". Diplomarbeit. Karlsruhe Institute of Technology (KIT), 2013.
- [89] Z.-G. F. and E. E. M. "The immersed boundary-lattice Boltzmann method for solving fluid–particles interaction problems". In: *Journal of Computational Physics* 195.2 (2004), pp. 602–628. DOI: [10.1016/j.jcp.2003.10.013](https://doi.org/10.1016/j.jcp.2003.10.013).

- [90] A. Fakhari and M. H. Rahimian. "Phase-field modeling by the method of lattice Boltzmann equations". In: *Physical Review E* 81.3 (2010), p. 036707.
- [91] C. Gau and R. Viskanta. "Melting and Solidification of a Pure Metal on a Vertical Wall". In: *Journal of Heat Transfer* 108.1 (1986), pp. 174–181.
- [92] Z. Guo, C. Zheng, and B. Shi. "Discrete lattice effects on the forcing term in the lattice Boltzmann method". In: *Phys. Rev. E* 65 (2002), p. 046308.
- [93] Z. Guo, B. Shi, and C. Zheng. "A coupled lattice BGK model for the Boussinesq equations". In: *Int. J. Num. Meth. Fluids* 39 (2002), pp. 325–342. DOI: [10.1002/flid.337](https://doi.org/10.1002/flid.337).
- [94] Z. Guo, B. Shi, and C. Zheng. "A coupled lattice BGK model for the Boussinesq equations". In: *International Journal for Numerical Methods in Fluids* 39.4 (2002), pp. 325–342.
- [95] Z. Guo and T. S. Zhao. "Lattice Boltzmann model for incompressible flows through porous media". In: *Phys. Rev. E* 66 (3 2002), p. 036304. DOI: [10.1103/PhysRevE.66.036304](https://doi.org/10.1103/PhysRevE.66.036304).
- [96] X.-Y. L. H-B Huang and M. C. Sukop. "Numerical study of lattice Boltzmann methods for a convection-diffusion equation coupled with Navier-Stokes equations". In: *J. Phys. A: Math. Theor.* 44.5 (2011).
- [97] R. Huang and H. Wu. "Phase interface effects in the total enthalpy-based lattice Boltzmann model for solid-liquid phase change". In: *Journal of Computational Physics* 294 (2015), pp. 346–362.
- [98] G. Kałuża. "The numerical solution of the transient heat conduction problem using the lattice Boltzmann method". In: *Scientific Research of the Institute of Mathematics and Computer Science* 11.1 (2012), pp. 23–30.
- [99] A. Komrakova, O. Shardt, D. Eskin, and J. Derksen. "Lattice Boltzmann simulations of drop deformation and breakup in shear flow". In: *International Journal of Multiphase Flow* 59 (2014), pp. 24–43. DOI: <https://doi.org/10.1016/j.ijmultiphaseflow.2013.10.009>.
- [100] T. Krüger, H. Kusumaatmaja, A. Kuzmin, O. Shardt, G. Silva, and E. M. Viggen. *The Lattice Boltzmann Method*. Springer, 2017.
- [101] A. Ladd and R. Verberg. "Lattice-Boltzmann simulations of particle-fluid suspensions". In: *Journal of Statistical Physics* 104.5-6 (2001), pp. 1191–1251.
- [102] D. Lagrava, O. Malaspinas, J. Latt, and B. Chopard. "Automatic grid refinement criterion for lattice Boltzmann method". In: *ArXiv e-prints* (July 2015).
- [103] L. Larocque, J. Imran, and M. Chaudhry. "Experimental and Numerical Investigations of Two-Dimensional Dam-Break Flows". In: *Journal of Hydraulic Engineering* 139 (June 2013), pp. 569–579. DOI: [10.1061/\(ASCE\)HY.1943-7900.0000705](https://doi.org/10.1061/(ASCE)HY.1943-7900.0000705).
- [104] J. Latt and B. Chopard. "Lattice Boltzmann Method with regularized non-equilibrium distribution functions". In: *Math. Comp. Sim.* 72 (2006), pp. 165–168.
- [105] J. Latt, B. Chopard, O. Malaspinas, M. Deville, and A. Michler. "Straight velocity boundaries in the lattice Boltzmann method". In: *Phys. Rev. E* 77 (5 2008), p. 056703. DOI: [10.1103/PhysRevE.77.056703](https://doi.org/10.1103/PhysRevE.77.056703).
- [106] L. Li, R. Mei, and J. F. Klausner. "Boundary conditions for thermal lattice Boltzmann equation method". In: *Journal of Computational Physics* 237 (2013), pp. 366–395.

- [107] Q. Liu and Y.-L. He. "Double multiple-relaxation-time lattice Boltzmann model for solid-liquid phase change with natural convection in porous media". In: *Physica A: Statistical Mechanics and its Applications* 438 (2015), pp. 94–106.
- [108] A. Mezrhab, M. A. Moussaoui, M. Jami, H. Naji, and M. Bouzidi. "Double MRT thermal lattice Boltzmann method for simulating convective flows". In: *Physics Letters A* 374.34 (2010), pp. 3499–3507.
- [109] S. C. Mishra and H. K. Roy. "Solving transient conduction and radiation heat transfer problems using the lattice Boltzmann method and the finite volume method". In: *Journal of Computational Physics* 223.1 (2007), pp. 89–107.
- [110] A. A. Mohamad. *Lattice Boltzmann Method - Fundamentals and Engineering Applications with Computer Codes*. Springer-Verlag, 2011.
- [111] A. Mohamad and A. Kuzmin. "A critical evaluation of force term in lattice Boltzmann method, natural convection problem". In: *International Journal of Heat and Mass Transfer* 53.5 (2010), pp. 990–996.
- [112] D. Noble and J. Torczynski. "A lattice-Boltzmann method for partially saturated computational cells". In: *Int. J. Modern Phys. C* 9.8 (1998), pp. 1189–1202.
- [113] Y. Peng, C. Shu, and Y. Chew. "Simplified thermal lattice Boltzmann model for incompressible thermal flows". In: *Physical Review E* 68.2 (2003), p. 026701.
- [114] G. Pingen, A. Evgrafov, and K. Maute. "Topology optimization of flow domains using the lattice Boltzmann method". English. In: *Structural and Multidisciplinary Optimization* 34.6 (2007), pp. 507–524. DOI: [10.1007/s00158-007-0105-7](https://doi.org/10.1007/s00158-007-0105-7).
- [115] R. Rannacher. *Einfuehrung in die Numerische Mathematik (Numerik 0)*. Vorlesungsskriptum SS 2005. Universitaet Heidelberg, 2006.
- [116] C. Semperebon, T. Krüger, and H. Kusumaatmaja. "A Ternary Free Energy Lattice Boltzmann Model with Tunable Surface Tensions and Contact Angles". In: *Physical Review E* 93.3 (2016), p. 033305.
- [117] X. Shan and H. Chen. "Lattice Boltzmann model for simulating flows with multiple phases and components". In: *Phys. Rev. E* 47 (1993), pp. 1815–1819. DOI: [10.1103/PhysRevE.47.1815](https://doi.org/10.1103/PhysRevE.47.1815).
- [118] X. Shan and G. Doolen. "Multicomponent lattice-Boltzmann model with interparticle interaction". In: *Journal of Statistical Physics* 81 (1995), pp. 379–393.
- [119] P. A. Skordos. "Initial and boundary conditions for the lattice Boltzmann method". In: *Phys. Rev. E* 48 (1993), pp. 4824–4842.
- [120] M. A. A. Spaid and F. R. Phelan. "Lattice Boltzmann methods for modeling microscale flow in fibrous porous media". In: *Physics of Fluids* 9.9 (1997), pp. 2468–2474. DOI: [10.1063/1.869392](https://doi.org/10.1063/1.869392).
- [121] S. Stasius. "Identifikation von Strömungsgebieten mit adjungierten Lattice Boltzmann Methoden (ALBM)". Diplomarbeit. Karlsruhe Institute for Technology (KIT), 2014.
- [122] N. Thurey. "Physically based animation of free surface flows with the lattice Boltzmann method". In: *Ph. D. Thesis, University of Erlangen* (2007).
- [123] S. Turek and M. Schäfer. "Benchmark computations of laminar flow around cylinder". In: *Flow Simulation with High-Performance Computers II*. Vol. 52. Notes on Numerical Fluid Mechanics. Vieweg, Jan. 1996, pp. 547–566.

- [124] L. Wang, Z. L. Guo, and J. C. Mi. "Drafting, kissing and tumbling process of two particles with different sizes". In: *Computers & Fluids* 96 (2014), pp. 20–34. DOI: [10.1016/j.compfluid.2014.03.005](https://doi.org/10.1016/j.compfluid.2014.03.005).
- [125] D. Yu, R. Mei, L.-S. Luo, and W. Shyy. "Viscous flow computations with the method of lattice Boltzmann equation". In: *Progress in Aerospace Sciences* 39.5 (2003), pp. 329–367.
- [126] H. Zheng, C. Shu, and Y.-T. Chew. "A lattice Boltzmann model for multiphase flows with large density ratio". In: *Journal of Computational Physics* 218.1 (2006), pp. 353–371.
- [127] Q. Zou and X. He. "On pressure and velocity boundary conditions for the lattice Boltzmann BGK model". In: *Phys. Fluids* 9 (1997), pp. 1592–1598.

A. Appendix

A.1. Q&A

In this Q&A part, some potential questions concerning the code are answered:

What do I need the unit converter for?

The unit converter (Listing A.1) is used in every simulation done with OpenLB. In this class, the physical units, like length or mass, are converted to lattice units and vice versa. This step is necessary to get a result in the correct physical dimensions and units.

```
1 UnitConverterFromResolutionAndLatticeVelocity<T, DESCRIPTOR> converter(  
2     (int) res, //resolution  
3     (T) charLatticeVelocity, //charLatticeVelocity  
4     (T) charPhysLength, //charPhysLength  
5     (T) charPhysVelocity, //charPhysVelocity  
6     (T) physViscosity, //physViscosity  
7     (T) physDensity //physDensity  
8 );  
9 converter.print();
```

Listing A.1.: UnitConverter

For a closer look, also checkout the respective example in Section 10.3.

How can I write my own dynamics?

Dynamics are the classes that model the cell-specific computation of momenta, equilibria and collision operator. Ideally, new dynamics are constructed using OpenLB's flexible *dynamics tuple* system which allows the composition of dynamics as a tuple of momenta, equilibria, and collision operator in addition to a optional *combination rule* for declaring e.g. forcing schemes. This is the approach used for most dynamics in OpenLB (also see Section 2.3).

```
1 template <typename T, typename DESCRIPTOR>  
2 using ForcedTRTDynamics = dynamics::Tuple<  
3     T, DESCRIPTOR,  
4     momenta::BulkTuple,  
5     equilibria::SecondOrder,  
6     collision::TRT,  
7     forcing::Guo  
8 >;
```

Modifying this example to use the BGK collision operator without forcing is as simple as writing:

```

1 template <typename T, typename DESCRIPTOR>
2 using BGKdynamics = dynamics::Tuple<
3   T, DESCRIPTOR,
4   momenta::BulkTuple,
5   equilibria::SecondOrder,
6   collision::BGK
7 >;

```

In most cases that are not yet covered by the extensive library of dynamics tuples, it should be sufficient to write e.g. a new collision operator to be plugged into this framework. As a fallback, fully custom dynamics can be implemented using `dynamics::CustomCollision` following this basic scaffold:

```

1 template <typename T, typename DESCRIPTOR, typename MOMENTA=momenta::BulkTuple>
2 struct MyCustomDynamics final : public dynamics::CustomCollision<T,DESCRIPTOR,MOMENTA>
3 {
4   using MomentaF = typename MOMENTA::template type<DESCRIPTOR>;
5
6   // Declare list of parameter fields (can be empty)
7   using parameters = meta::list<...>;
8
9   // Allow exchanging the momenta, used for example to construct boundary dynamics
10  template <typename M>
11  using exchange_momenta = MyCustomDynamics<T,DESCRIPTOR,DYNAMICS,M>;
12
13  template <typename CELL, typename PARAMETERS, typename V=typename CELL::value_t>
14  CellStatistic<V> apply(CELL& cell, PARAMETERS& parameters) any_platform {
15    // Implement custom collision here
16  };
17
18  T computeEquilibrium(int iPop, T rho, const T u[DESCRIPTOR::d]) const override
19    any_platform {
20    // Implement custom equilibrium computation here
21  };
22
23  std::type_index id() override {
24    return typeid(MyCustomDynamics);
25  };
26
27  AbstractParameters<T,DESCRIPTOR>& getParameters(BlockLattice<T,DESCRIPTOR>& block)
28    override {
29    return block.template getData<OperatorParameters<MyCustomDynamics>>>();
30  }
31
32  // Return human readable name
33  std::string getName() const override {
34    return "MyCustomDynamics<" + MomentaF().getName() + ">";
35  };
36 };

```

How can I write my own post processor?

A non-local operator, also referred to as a *post processor* in OpenLB, is any class that provides a scope, a priority and an `apply` method template (also see Section 2.4). The scope declares how the `apply` method is

to be called, the priority is used to sort the execution sequence of multiple operators assigned to the same *stage* and the `apply` method template implements the actual instructions to be performed.

```

1 struct MyCustomPostProcessor {
2     // One of OperatorScope::(PerCell,PerCellWithParameters,PerBlock)
3     static constexpr OperatorScope scope = OperatorScope::PerCell;
4
5     int getPriority() const {
6         return 0;
7     }
8
9     template <typename CELL>
10    void apply(CELL& cell) any_platform {
11        // custom non-local code here
12        // access neighbors via 'cell.neighbor(c_i)'
13    }
14 };

```

Listing A.2.: Simple post processor implementation

This new post processor can be assigned to cells of the lattice using the various overloads of `SuperLattice::addPostProcessor`:

```

1 // Assign MyCustomPostProcessor to all cells
2 sLattice.addPostProcessor(meta::id<MyCustomPostProcessor>{});
3 // Assign MyCustomPostProcessor to indicated cells
4 sLattice.addPostProcessor(indicatorF,
5                             meta::id<MyCustomPostProcessor>{});

```

Listing A.3.: Simple post processor assignment

If the operator depends on non-cell-specific parameters, they can be declared by changing the `scope` to `OperatorScope::PerCellWithParameters` and modifying the `apply` template arguments.

```

1 struct MyCustomPostProcessor {
2     static constexpr OperatorScope scope = OperatorScope::PerCellWithParameters;
3
4     using parameters = meta::list<OMEGA>;
5
6     int getPriority() const {
7         return 0;
8     }
9
10    template <typename CELL, typename PARAMETERS>
11    void apply(CELL& cell, PARAMETERS& parameters) any_platform {
12        // access parameter via 'parameters.template get<OMEGA>()'
13    }
14 };

```

Listing A.4.: Simple post processor implementation with parameters

The parameters are set in the same way as dynamics parameters.

```

1 // Set OMEGA parameter of all dynamics and post processors to 0.6
2 sLattice.setParameter<OMEGA>(0.6);

```

How can I write my own coupling operator?

A coupling operator is any class that provides a scope and an `apply` method template. Different from single-lattice non-local operators we receive not one cell but a named tuple of them. Consider for illustration a basic coupling between NSE and ADE lattices in Listing A.5.

```
1 struct NavierStokesAdvectionDiffusionCoupling {
2     // Declare that we want cell-wise coupling with some global parameters
3     static constexpr OperatorScope scope = OperatorScope::PerCellWithParameters;
4
5     // Declare the two parameters custom to this coupling operator
6     struct FORCE_PREFACTOR : public descriptors::FIELD_BASE<0,1> { };
7     struct T0 : public descriptors::FIELD_BASE<1> { };
8
9     // Declare which parameters are required
10    using parameters = meta::list<FORCE_PREFACTOR,T0>;
11
12    template <typename CELLS, typename PARAMETERS>
13    void apply(CELLS& cells, PARAMETERS& parameters) any_platform
14    {
15        using V = typename CELLS::template value_t<names::NavierStokes>::value_t;
16        using DESCRIPTOR = typename CELLS::template value_t<names::NavierStokes>::
            descriptor_t;
17
18        // Get the cell of the NavierStokes lattice
19        auto& cellNSE = cells.template get<names::NavierStokes>();
20        // Get the cell of the Temperature lattice
21        auto& cellADE = cells.template get<names::Temperature>();
22
23        // Computation of the Boussinesq force
24        auto force = cellNSE.template getFieldPointer<descriptors::FORCE>();
25        auto forcePrefactor = parameters.template get<FORCE_PREFACTOR>();
26        V temperatureDifference = cellADE.computeRho() - parameters.template get<T0>();
27        for (unsigned iD = 0; iD < DESCRIPTOR::d; ++iD) {
28            force[iD] = forcePrefactor[iD] * temperatureDifference;
29        }
30        // Velocity coupling
31        V u[DESCRIPTOR::d] { };
32        cellNSE.computeU(u);
33        cellADE.template setField<descriptors::VELOCITY>(u);
34    }
35
36 };
```

Listing A.5.: Full implementation of basic NSE-ADE coupling

Coupling operators are instantiated using the `SuperLatticeCoupling` template by providing a list of names and assigned lattices. For the NSE-ADE coupling this will look similar to (see e.g. `thermal/rayleighBenard(2,3)d` for a practical example):

```
1 SuperLatticeCoupling coupling(
2     NavierStokesAdvectionDiffusionCoupling{},
3     names::NavierStokes{}, NSlattice,
4     names::Temperature{}, ADlattice);
5 coupling.setParameter<NavierStokesAdvectionDiffusionCoupling::T0>(...);
6 coupling.setParameter<NavierStokesAdvectionDiffusionCoupling::FORCE_PREFACTOR>(...);
```

As we can see the parameters are set for the specific coupling instance.

What are aliases used for?

Aliases are used to keep the code simpler for users and allow them to get a faster overview over the code. Aliases aren't used for all parts of the program. Especially if the user likes to change the code for a special problem, alias-functions may not be available and therefore need to be created by the user himself or normal functions can be used. To sum up, *alias*-functions aren't necessary for the simulation to work, but allow to simplify the code for a better overview for the user.

Why do we use "constexpr" to define some functions?

Constexpr allows you to get the output of the function at the time of compilation. The value returned is set to a constant expression and as a result the runtime can be reduced, due to the constant value. If the return value is part of an if-loop, it won't be checked more than once, as the result is already clear, after the first check. In the example below the rotation of a 2D-particle is calculated and therefore the output is set to a concrete value at compilation time, and this value stays constant.

```
1      static constexpr Vector<T,2> execute( Vector<T,2> input, Vector<T,4>
      rotationMatrix, Vector<T,2> rotationCenter = Vector<T,2>(0.,0.) )
2      {
3          Vector<T,2> dist = input - rotationCenter;
4          return Vector<T,2>(
5              dist[0]*rotationMatrix[0] +
6              dist[1]*rotationMatrix[2],
7              dist[0]*rotationMatrix[1] +
8              dist[1]*rotationMatrix[3] );
9      }
```

Listing A.6.: constexpr

7. What is the difference between files ending with .h and ending with .hh?

The files ending with .h like particleDynamics.h are header-files, where classes and functions are only declared. The specification (definition) happens in the .hh files. In the following two code-snippets of the same function, the differences in terms of specification can be compared.

```
1      template<typename T, typename DESCRIPTOR>
2      class VerletParticleDynamics : public ParticleDynamics<T,DESCRIPTOR> {
3      public:
4          /// Constructor
5          VerletParticleDynamics( T timeStepSize );
6          /// Procesisng step
7          void process (Particle<T,DESCRIPTOR>& particle) override;
8      private:
9          T _timeStepSize;
10     };
```

Listing A.7.: particleDynamics.h

```

1      template<typename T, typename PARTICLETYPE>
2      VerletParticleDynamics<T,PARTICLETYPE>::VerletParticleDynamics ( T timeStepSize
3      )
4      : _timeStepSize( timeStepSize )
5      {
6          this->getName() = "VerletParticleDynamics";
7      }
8
9      template<typename T, typename PARTICLETYPE>
10     void VerletParticleDynamics<T,PARTICLETYPE>::process (
11     Particle<T,PARTICLETYPE>& particle )
12     {
13         //Calculate acceleration
14         auto acceleration = getAcceleration<T,PARTICLETYPE>( particle );
15         //Check for angular components
16         if constexpr ( providesAngle<PARTICLETYPE>() ) {
17             //Calculate angular acceleration
18             auto angularAcceleration = getAngAcceleration<T,PARTICLETYPE>(
19             particle );
20             //Verlet algorithm
21             particles::dynamics::velocityVerletIntegration<T, PARTICLETYPE
22             >(
23             particle, _timeStepSize , acceleration, angularAcceleration );
24             //Update rotation matrix
25             updateRotationMatrix<T,PARTICLETYPE>( particle );
26         } else {
27             //Verlet algorithm without rotation
28             particles::dynamics::velocityVerletIntegration<T, PARTICLETYPE
29             >(
30             particle, _timeStepSize , acceleration );
31         }
32     }

```

Listing A.8.: particleDynamics.hh

A.2. List of Project Participants

Since 2006 the following persons have contributed source code to OpenLB:

Armani Arfaoui: core: performance improvements for D3Q19 BGK collision operator

Sam Avis (active): dynamics: multicomponent free energy model

Saada Badie: core: performance improvements for D3Q19 BGK collision operator

Lukas Baron: utilities: (parallel) console output, time and performance measurement, dynamics: porous media model, functors: concept, div. functors implementation

Fedor Bukreev (active): reaction: adsorption and reaction models, examples: adsorption examples, organization: testing

Vojtech Cvrcek: dynamics: power law, examples: power law, updates, functors: 2D adaptation

Davide Dapelo (active): **core:** power-law unit converter, **dynamics:** Guo–Zhao porous, contributions on power-law, contributions on HLBM, **examples:** reactionFiniteDifferences2d, advectionDiffusion3d, advectionDiffusionPipe3d, **functors:** contributions on indicator and smooth indicator

Tim Dornieden: **functors:** smooth start scaling, **io:** vti writer

Jonas Fietz: **io:** configure file parsing based on XML, octree STL reader interface to CVMLCPP (< release 0.9), **communication:** heuristic load balancer

Benjamin Förster: **core:** super data implementation **io:** new serializer and serializable implementation, vti writer, new vti reader, **functors:** new discrete indicator

Max Gaedtke: **core:** unit converter, **dynamics:** thermal, **examples:** thermal

Simon Großmann: **example:** poiseuille2dEOC, **io:** csv and gnuplot interface, **postprocessing:** eoc analysis

Nicolas Hafen (active): **particles:** core framework, surface resolved particles, coupling, dynamics, creator-functions, **dynamics:** moving porous media (HLBM), **examples:** surface resolved particle simulations, **particles:** particle framework refactoring, sub-grid scale refactoring

Marc Haussmann: **dynamics:** turbulence modeling, **examples:** tgv3d, **io:** gnuplot heatmap

Thomas Henn: **io:** voxelizer interface based on STL, **particles:** particulate flows

Claudius Holeksa : **postProcessor:** free surface, **example:** free surface

Anna Husfeldt : **functors:** signed distance surface framework

Jonathan Jeppener-Haltenhoff: **functors:** wall shear stress, **examples:** channel3d, poiseuille3d, **core:** contributions to define field, **documentation:** user guide

Julius Jeßberger (active): **core:** solver, template momenta concept, optimization, **examples:** poiseuille2d, cavity2dSolver, porousPlate3dSolver, testFlow3dSolver, optimization examples, **postprocessing:** error analysis, **utilities:** algorithmic differentiation

Fabian Klemens: **functors:** flux, indicator-based functors **io:** gnuplot interface

Jonas Kratzke: **core:** unit converter, **io:** GUI interface based on description files and OpenGPI, **boundaries:** Bouzidi boundary condition

Mathias J. Krause (active): **core:** hybrid-parallelization approach, super structure, **communication:** OpenMP parallelization, cuboid data structure for MPI parallelization, load balancing, **general:** makefile environment for compilation, integration and maintenance of added components (since 2008), **boundaries:** Bouzidi boundary condition, convection, **geometry:** concept, parallelization, statistics, **io** new serializer and serializable concept, **functors:** concept, div. functors implementation, **examples:** venturi3d, aorta3d, optimization (-2020), **optimization:** (-2020), **organization:** integration and maintenance of added components (2008-2017), project management (2006-)

Louis Kronberg: **core:** ade unit converter, **dynamics:** KBC, entropic LB, Cumulant, **examples:** advectionDiffusion1d, advectionDiffusion2d, bstep2d

Eliane Kummer: **documentation:** user guide

Adrian Kummerländer (active): **core:** SIMD CPU support, CUDA GPU support, population and field data structure, propagation pattern, vector hierarchy, cell interface, field data interface, meta descriptors, automatic code generation, **dynamics:** new dynamics concept, dynamics tuple, momenta concept **communication:** block propagation, communication, **functors:** lp-norm, flux, reduction, lattice indicator, error norms, refinement quality criterion, composition, **boundaries:** new post processor concept, water-tightness testing and post-processor priority, **general:** CI maintenance, Nix environment

Jonas Latt: **core:** basic block structure, **communication:** basic parallel block lattice approach (< release 0.9), **io:** vti writer, **general:** integration and maintenance of added components (2006-2008), **boundaries:** basic boundary structure, **dynamics:** basic dynamics structure, **examples:** numerous examples, which have been further developed in recent years, **organization:** integration and maintenance of added components (2006-2008), project management (2006-2008)

Marie-Luise Maier: **particles:** particulate flows, frame change

Orestis Malaspinas: **boundaries:** alternative boundary conditions (Inamuro, Zou/He, Nonlinear FD), **dynamics:** alternative LB models (Entropic LB, MRT)

Jan E. Marquardt (active): **particles:** surface resolved particles, coupling, creator-functions, **dynamics:** Euler-Euler particle dynamics, **functors:** signed distance surface framework, **utilities:** algorithmic differentiation

Cyril Masquelier: **functors:** indicator, smooth indicator

Albert Mink: **functors:** arithmetic, **io:** parallel VTK interface3, zLib compression for VTK data, GifWriter, **dynamics:** radiative transport, **boundary:** diffuse reflective boundary

Markus Mohrhard: **general:** makefile environment for parallel compilation, **organization:** integration and maintenance of added components (2018-)

Johanna Mödl: **core:** convection diffusion reaction dynamics, **examples:** advectionDiffusionReaction2d

Patrick Nathen: **dynamics:** turbulence modeling (advanced subgrid-scale models), **examples:** nozzle3d

Aleksandra Pachalieva: **dynamics:** thermal (MRT model), **examples:** thermal (MRT model)

Maximilian Schecher (active) : **postProcessor:** free surface, **example:** free surface

Stephan Simonis (active): **core:** ade unit converter **examples:** advectionDiffusion1d, advectionDiffusion2d, advectionDiffusion3d, advectionDiffusionPipe3d, binaryShearFlow2d, fourRollMill2d, **documentation:** user guide, **dynamics:** MRT, KBC, Cumulant, entropic LB, free energy model

Lukas Springmann: **particles:** user-guide, unit tests

Bernd Stahl: **communication:** 3D extension to MultiBlock structure for MPI parallelization (< release 0.9), **core:** parallel version of (scalar or tensor-valued) data fields (< release 0.9), **io:** VTK output of data (< release 0.9)

Dennis Teutscher (active) : **visualisation, organization:** user guide

Robin Trunk: **dynamics:** parallel thermal, advection diffusion models, 3D HLBm, Euler-Euler particle, multicomponent free energy model

Peter Weisbrod: **dynamics:** parallel multi phase/component, **examples:** structure and showcases, phaseSeparationXd

Gilles Zahnd: **functors:** rotating frame functors

Asher Zarth: **core:** vector implementation

Simon Zimny: **io:** pre-processing: automated setting of boundary conditions

A.3. GNU Free Documentation License

GNU Free Documentation License

Version 1.2, November 2002

Copyright © 2000,2001,2002 Free Software Foundation, Inc.

51 Franklin St, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The purpose of this License is to make a manual, textbook, or other functional and useful document “free” in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondly, this License preserves for the author and publisher a way to get credit for their work, while not being considered responsible for modifications made by others.

This License is a kind of “copyleft”, which means that derivative works of the document must themselves be free in the same sense. It complements the GNU General Public License, which is a copyleft license designed for free software.

We have designed this License in order to use it for manuals for free software, because free software needs free documentation: a free program should come with manuals providing the same freedoms that the software does. But this License is not limited to software manuals; it can be used for any textual work, regardless of subject matter or whether it is published as a printed book. We recommend this License principally for works whose purpose is instruction or reference.

1. Applicability and definitions

This License applies to any manual or other work, in any medium, that contains a notice placed by the copyright holder saying it can be distributed under the terms of this License. Such a notice grants a world-wide, royalty-free license, unlimited in duration, to use that work under the conditions stated herein. The “**Document**”, below, refers to any such manual or work. Any member of the public is a licensee, and is addressed as “**you**”. You accept the license if you copy, modify or distribute the work in a way requiring permission under copyright law.

A “**Modified Version**” of the Document means any work containing the Document or a portion of it, either copied verbatim, or with modifications and/or translated into another language.

A “**Secondary Section**” is a named appendix or a front-matter section of the Document that deals exclusively with the relationship of the publishers or authors of the Document to the Document’s overall subject (or to related matters) and contains nothing that could fall directly within that overall subject. (Thus, if the Document is in part a textbook of mathematics, a Secondary Section may not explain any mathematics.) The relationship could be a matter of historical connection with the subject or with related matters, or of legal, commercial, philosophical, ethical or political position regarding them.

The “**Invariant Sections**” are certain Secondary Sections whose titles are designated, as being those of Invariant Sections, in the notice that says that the Document is released under this License. If a section does not fit the above definition of Secondary then it is not allowed to be designated as Invariant. The Document may contain zero Invariant Sections. If the Document does not identify any Invariant Sections then there are none.

The “**Cover Texts**” are certain short passages of text that are listed, as Front-Cover Texts or Back-Cover Texts, in the notice that says that the Document is released under this License. A Front-Cover Text may be at most 5 words, and a Back-Cover Text may be at most 25 words.

A “**Transparent**” copy of the Document means a machine-readable copy, represented in a format whose specification is available to the general public, that is suitable for revising the document straightforwardly with generic text editors or (for images composed of pixels) generic paint programs or (for drawings) some widely available drawing editor, and that is suitable for input to text formatters or for automatic translation to a variety of formats suitable for input to text formatters. A copy made in an otherwise Transparent file format whose markup, or absence of markup, has been arranged to thwart or discourage subsequent modification by readers is not Transparent. An image format is not Transparent if used for any substantial amount of text. A copy that is not “Transparent” is called “**Opaque**”.

Examples of suitable formats for Transparent copies include plain ASCII without markup, Texinfo input format, LaTeX input format, SGML or XML using a publicly available DTD, and standard-conforming simple HTML, PostScript or PDF designed for human modification. Examples of transparent image formats include PNG, XCF and JPG. Opaque formats include proprietary formats that can be read and edited only by proprietary word processors, SGML or XML for which the DTD and/or processing tools are not generally available, and the machine-generated HTML, PostScript or PDF produced by some word processors for output purposes only.

The “**Title Page**” means, for a printed book, the title page itself, plus such following pages as are needed to hold, legibly, the material this License requires to appear in the title page. For works in formats which do not have any title page as such, “Title Page” means the text near the most prominent appearance of the work’s title, preceding the beginning of the body of the text.

A section “**Entitled XYZ**” means a named subunit of the Document whose title either is precisely XYZ or contains XYZ in parentheses following text that translates XYZ in another language. (Here XYZ stands for a specific section name mentioned below, such as “**Acknowledgements**”, “**Dedications**”, “**Endorsements**”, or “**History**”.) To “**Preserve the Title**” of such a section when you modify the Document means that it remains a section “Entitled XYZ” according to this definition.

The Document may include Warranty Disclaimers next to the notice which states that this License applies to the Document. These Warranty Disclaimers are considered to be included by reference in this License, but only as regards disclaiming warranties: any other implication that these Warranty Disclaimers may have is void and has no effect on the meaning of this License.

2. Verbatim Copying

You may copy and distribute the Document in any medium, either commercially or noncommercially, provided that this License, the copyright notices, and the license notice saying this License applies to the Document are reproduced in all copies, and that you add no other conditions whatsoever to those of this License. You may not use technical measures to obstruct or control the reading or further copying of the copies you make or distribute. However, you may accept compensation in exchange for copies. If you distribute a large enough number of copies you must also follow the conditions in section 3.

You may also lend copies, under the same conditions stated above, and you may publicly display copies.

3. Copying in quantity

If you publish printed copies (or copies in media that commonly have printed covers) of the Document, numbering more than 100, and the Document's license notice requires Cover Texts, you must enclose the copies in covers that carry, clearly and legibly, all these Cover Texts: Front-Cover Texts on the front cover, and Back-Cover Texts on the back cover. Both covers must also clearly and legibly identify you as the publisher of these copies. The front cover must present the full title with all words of the title equally prominent and visible. You may add other material on the covers in addition. Copying with changes limited to the covers, as long as they preserve the title of the Document and satisfy these conditions, can be treated as verbatim copying in other respects.

If the required texts for either cover are too voluminous to fit legibly, you should put the first ones listed (as many as fit reasonably) on the actual cover, and continue the rest onto adjacent pages.

If you publish or distribute Opaque copies of the Document numbering more than 100, you must either include a machine-readable Transparent copy along with each Opaque copy, or state in or with each Opaque copy a computer-network location from which the general network-using public has access to download using public-standard network protocols a complete Transparent copy of the Document, free of added material. If you use the latter option, you must take reasonably prudent steps, when you begin distribution of Opaque copies in quantity, to ensure that this Transparent copy will remain thus accessible at the stated location until at least one year after the last time you distribute an Opaque copy (directly or through your agents or retailers) of that edition to the public.

It is requested, but not required, that you contact the authors of the Document well before redistributing any large number of copies, to give them a chance to provide you with an updated version of the Document.

4. Modifications

You may copy and distribute a Modified Version of the Document under the conditions of sections 2 and 3 above, provided that you release the Modified Version under precisely this License, with the Modified Version filling the role of the Document, thus licensing distribution and modification of the Modified Version to whoever possesses a copy of it. In addition, you must do these things in the Modified Version:

- A. Use in the Title Page (and on the covers, if any) a title distinct from that of the Document, and from those of previous versions (which should, if there were any, be listed in the History section of the Document). You may use the same title as a previous version if the original publisher of that version gives permission.

- B. List on the Title Page, as authors, one or more persons or entities responsible for authorship of the modifications in the Modified Version, together with at least five of the principal authors of the Document (all of its principal authors, if it has fewer than five), unless they release you from this requirement.
- C. State on the Title page the name of the publisher of the Modified Version, as the publisher.
- D. Preserve all the copyright notices of the Document.
- E. Add an appropriate copyright notice for your modifications adjacent to the other copyright notices.
- F. Include, immediately after the copyright notices, a license notice giving the public permission to use the Modified Version under the terms of this License, in the form shown in the Addendum below.
- G. Preserve in that license notice the full lists of Invariant Sections and required Cover Texts given in the Document's license notice.
- H. Include an unaltered copy of this License.
- I. Preserve the section Entitled "History", Preserve its Title, and add to it an item stating at least the title, year, new authors, and publisher of the Modified Version as given on the Title Page. If there is no section Entitled "History" in the Document, create one stating the title, year, authors, and publisher of the Document as given on its Title Page, then add an item describing the Modified Version as stated in the previous sentence.
- J. Preserve the network location, if any, given in the Document for public access to a Transparent copy of the Document, and likewise the network locations given in the Document for previous versions it was based on. These may be placed in the "History" section. You may omit a network location for a work that was published at least four years before the Document itself, or if the original publisher of the version it refers to gives permission.
- K. For any section Entitled "Acknowledgements" or "Dedications", Preserve the Title of the section, and preserve in the section all the substance and tone of each of the contributor acknowledgements and/or dedications given therein.
- L. Preserve all the Invariant Sections of the Document, unaltered in their text and in their titles. Section numbers or the equivalent are not considered part of the section titles.
- M. Delete any section Entitled "Endorsements". Such a section may not be included in the Modified Version.
- N. Do not retitle any existing section to be Entitled "Endorsements" or to conflict in title with any Invariant Section.
- O. Preserve any Warranty Disclaimers.

If the Modified Version includes new front-matter sections or appendices that qualify as Secondary Sections and contain no material copied from the Document, you may at your option designate some or all of these sections as invariant. To do this, add their titles to the list of Invariant Sections in the Modified Version's license notice. These titles must be distinct from any other section titles.

You may add a section Entitled "Endorsements", provided it contains nothing but endorsements of your Modified Version by various parties—for example, statements of peer review or that the text has been approved by an organization as the authoritative definition of a standard.

You may add a passage of up to five words as a Front-Cover Text, and a passage of up to 25 words as a Back-Cover Text, to the end of the list of Cover Texts in the Modified Version. Only one passage of Front-Cover Text and one of Back-Cover Text may be added by (or through arrangements made by) any one entity. If the Document already includes a cover text for the same cover, previously added by you or by arrangement made by the same entity you are acting on behalf of, you may not add another; but you may replace the old one, on explicit permission from the previous publisher that added the old one.

The author(s) and publisher(s) of the Document do not by this License give permission to use their names for publicity for or to assert or imply endorsement of any Modified Version.

5. Combining documents

You may combine the Document with other documents released under this License, under the terms defined in section 4 above for modified versions, provided that you include in the combination all of the Invariant Sections of all of the original documents, unmodified, and list them all as Invariant Sections of your combined work in its license notice, and that you preserve all their Warranty Disclaimers.

The combined work need only contain one copy of this License, and multiple identical Invariant Sections may be replaced with a single copy. If there are multiple Invariant Sections with the same name but different contents, make the title of each such section unique by adding at the end of it, in parentheses, the name of the original author or publisher of that section if known, or else a unique number. Make the same adjustment to the section titles in the list of Invariant Sections in the license notice of the combined work.

In the combination, you must combine any sections Entitled “History” in the various original documents, forming one section Entitled “History”; likewise combine any sections Entitled “Acknowledgements”, and any sections Entitled “Dedications”. You must delete all sections Entitled “Endorsements”.

6. Collections of documents

You may make a collection consisting of the Document and other documents released under this License, and replace the individual copies of this License in the various documents with a single copy that is included in the collection, provided that you follow the rules of this License for verbatim copying of each of the documents in all other respects.

You may extract a single document from such a collection, and distribute it individually under this License, provided you insert a copy of this License into the extracted document, and follow this License in all other respects regarding verbatim copying of that document.

7. Aggregation with independent works

A compilation of the Document or its derivatives with other separate and independent documents or works, in or on a volume of a storage or distribution medium, is called an “aggregate” if the copyright resulting from the compilation is not used to limit the legal rights of the compilation’s users beyond what the individual works permit. When the Document is included in an aggregate, this License does not apply to the other works in the aggregate which are not themselves derivative works of the Document.

If the Cover Text requirement of section 3 is applicable to these copies of the Document, then if the Document is less than one half of the entire aggregate, the Document’s Cover Texts may be placed on covers that bracket the Document within the aggregate, or the electronic equivalent of covers if the Document is in electronic form. Otherwise they must appear on printed covers that bracket the whole aggregate.

8. Translation

Translation is considered a kind of modification, so you may distribute translations of the Document under the terms of section 4. Replacing Invariant Sections with translations requires special permission from their copyright holders, but you may include translations of some or all Invariant Sections in addition to the original versions of these Invariant Sections. You may include a translation of this License, and all the license notices in the Document, and any Warranty Disclaimers, provided that you also include the original English version of this License and the original versions of those notices and disclaimers. In case of a disagreement between the translation and the original version of this License or a notice or disclaimer, the original version will prevail.

If a section in the Document is Entitled “Acknowledgements”, “Dedications”, or “History”, the requirement (section 4) to Preserve its Title (section 1) will typically require changing the actual title.

9. Termination

You may not copy, modify, sublicense, or distribute the Document except as expressly provided for under this License. Any other attempt to copy, modify, sublicense or distribute the Document is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

10. Future revisions of this license

The Free Software Foundation may publish new, revised versions of the GNU Free Documentation License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. See <http://www.gnu.org/copyleft/>.

Each version of the License is given a distinguishing version number. If the Document specifies that a particular numbered version of this License “or any later version” applies to it, you have the option of following the terms and conditions either of that specified version or of any later version that has been published (not as a draft) by the Free Software Foundation. If the Document does not specify a version number of this License, you may choose any version ever published (not as a draft) by the Free Software Foundation.