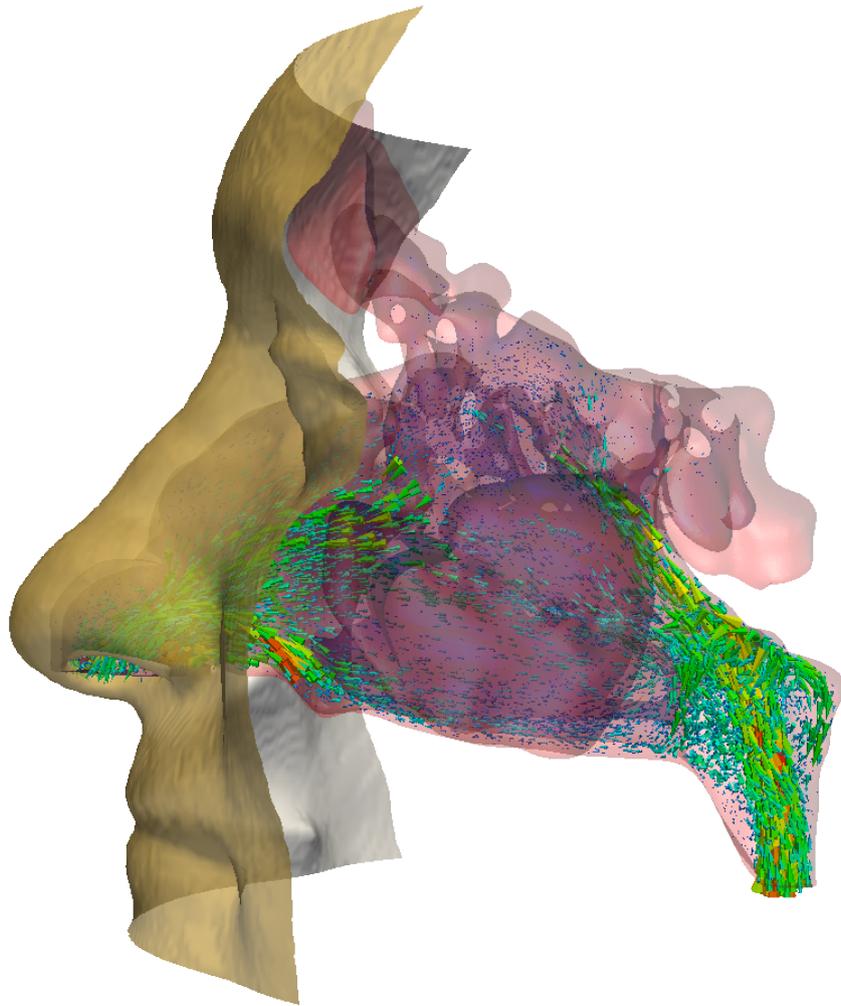


OpenLB User Guide
Associated to Release 1.0 of the Code



Copyright © 2006-2008 Jonas Latt
Copyright © 2008-2016 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”.

Contents

1	Introduction	6
1.1	Fluid Flow Simulations	6
1.2	Lattice Boltzmann Methods	6
1.3	The OpenLB Project	6
1.3.1	What is OpenLB?	6
1.3.2	Getting help with OpenLB	7
1.3.3	Which features are currently implemented?	8
1.3.4	Project participants	9
2	Using OpenLB for Applications	11
2.1	Lesson 2: Understand the <code>BlockLattice</code>	16
2.2	Lesson 3: Define and use boundary conditions	18
2.3	Lesson 4: Convert between lattice and physical units	22
2.4	Lesson 5: Extract data from a simulation	24
2.5	Lesson 6: Use an external force	26
2.6	Lesson 7: Understand genericity in OpenLB	26
2.7	Lesson 8: Use checkpointing for long duration simulations	28
2.8	Lesson 9: Save memory when domain boundaries are irregular	28
2.9	Lesson 10: Run your programs on a parallel machine	29
3	Compilation	31
3.1	Linux	31
3.2	Mac	31
3.3	Windows	31
4	Geometry	32
4.1	Material numbers	32
4.2	Indicator functions	32
4.3	Creating a Geometry	34
4.4	Excursion: Creating STL-files	35
5	Lattice Boltzmann Models and Core Data Structures	37
5.1	Concept – Data Organization	37
5.1.1	Cell – <code>BlockLattice</code> – <code>SuperLattice</code>	37
5.1.2	Descriptor	37
5.1.3	Dynamics	37
5.2	Classic BGK Model	39

5.3	MRT Model	39
5.4	Porous Media Model	39
5.5	Power Law Model	40
5.6	External Force	42
5.7	Multiphysics Couplings	42
5.7.1	The Shan-Chen Model	42
5.7.2	Implementation of Shan-Chen Two-phase Fluid	42
5.7.3	Implementation of Shan-Chen Two-component Fluid	43
5.7.4	Thermal Fluid with Boussinesq Approximation	44
5.8	Advection Diffusion Equation	44
6	Input / Output	47
6.1	The Output Data Structure in Parallel Simulations	47
6.2	Data Output to VTK File Format	48
6.3	Write Images in OpenLB	49
6.4	Console output	49
6.5	Read and write STL files	50
6.6	XML parameter files	51
7	Visualization with Paraview	53
7.1	Clip	53
7.2	Glyph	54
7.3	Stream Tracer	54
7.4	Transform	54
8	Functors – A General Concept For Input and Output of Data	55
8.1	Functors in OpenLB	55
8.2	How are functors used?	56
8.3	Functor arithmetic	61
9	Parallel program execution	64
9.1	Data-parallel structures	64
9.2	Duplicated data types	65
10	The example programs	66
10.1	aorta3d	66
10.2	bifurcation3d	66
10.2.1	Euler - Euler	66
10.2.2	Euler - Lagrange	67
10.3	bstep2d and bstep3d	69
10.4	cavity2d and cavity3d	69
10.5	cylinder2d and cylinder3d	69
10.6	multiComponent2d and multiComponent3d	69
10.7	nozzle3d	69

10.8 phaseSeparation2d and phaseSeparation3d	70
10.9 poiseuille2d	70
10.10thermal2d and thermal3d	70
10.11venturi3d	70
11 Bibliography	71
12 License	74

1 Introduction

1.1 Fluid Flow Simulations

1.2 Lattice Boltzmann Methods

This text is directed at people who want gain insight into *Lattice-Boltzmann Methods (LBM)*.

- The most recent publication that this documentation refers to, was written by **Erlend Magnus Vigen**. His PhD Thesis `THE LATTICE BOLTZMANN METHOD: FUNDAMENTALS AND ACOUSTICS` published in 2014, delivers a clear and complete introduction for beginners. Chapters 3 and 4 are particularly relevant, in which he develops the fundamentals, such as *theory of gas kinetics* and *the Boltzmann equation*.
- A concise introduction is given by **A. A. Mohamad**. In his book `LATTICE BOLTZMANN METHOD` [2011], he shows clearly, how one obtains macroscopic equations from *LBM* using *Chapman-Enskog expansion*.
- The readers who want insight into *Lattice-Gas Cellular Automatas* - the historical origin of *LBM* - may want to refer to **Dieter A. Wolf-Glodrow's** book `LATTICE-GAS CELLULAR AUTOMATA AND LATTICE BOLTZMANN MODELS` [2000]. Starting with *Cellular Automata*, he explores the beauty of *LBM*. A helpful interpretation of *LBM* is given in the beginning of the book.
- A quick overview of *LBM*, is obtained from the often cited paper of **S. Chen** and **G. D. Doolen** `LATTICE BOLTZMANN METHOD FOR FLUID FLOWS` published in 1998.

1.3 The OpenLB Project

1.3.1 What is OpenLB?

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 5.1.2).

1.3.2 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, are found on the website <http://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 `Makefile.inc` can provide useful information too.

subsectionCompiling OpenLB programs *Note: The framework for compiling OpenLB code is based on Makefiles and has so far been tested only on platforms of the Linux/Unix family, including Mac OS X and Cygwin. If you are working under Windows and want to get started quickly, you might consider installing the free Cygwin software [1], which efficiently emulates a Posix environment under Windows (a large part of OpenLB was developed under Cygwin).*

OpenLB consists of generic, template-based code, which needs to be included in the code of application programs, and precompiled 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 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 many OpenLB users presently tend to play around with the OpenLB code as well.

The file `Makefile.inc` 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, and a switch between normal/debug mode, and between sequential/openmp-parallel/mpi-parallel programs.

To compile your own OpenLB programs from an arbitrary directory, make a copy of a sample Makefile. Edit the `ROOT:=` entry to indicate the location of the OpenLB source, and the `OUTPUT:=` entry to explicit the name of your program, without file extension.

1.3.3 Which features are currently implemented?

Lattice Boltzmann models

BGK model for fluids	Section 5.1.3	Reference [15]
Regularized model for fluids	Section 5.1.3	Reference [25]
Multiple Relaxation Times (MRT)	Section 5.1.3	References [17, 34]
Entropic Lattice Boltzmann	Section 5.1.3	Reference [10]
BGK with adjustable speed of sound	Section 5.1.3	References [2, 16]
BGK and MRT with Smagorinsky model	Section 5.1.3	References [27]
Porous media model	Section 5.1.3	

Multiphysics coupling

Shan-Chen two-component fluid	Section 5.7	Reference [29]
Thermal fluid with Boussinesq approximation	Section 5.7	Reference [21]

Lattice structures

D2Q9	This lattice is available in the precompiled library
D3Q13	This lattice requires the use of a specific dynamics object (see also Ref. [18])
D3Q15	
D3Q19	This lattice is available in the precompiled library
D3Q27	

Boundary conditions for straight boundaries (including corners)

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

Boundary conditions for curved boundaries

Bouzidi	non-local	first order	References [12]
---------	-----------	-------------	-----------------

Data structures

The basic data structure used by an application programmer is the `BlockLatticeXD`. Here, the placeholder `X` stands for the number 2 or 3, depending on whether a 2D or 3D lattice is instantiated. A generalization of the `BlockLatticeXD` are the `CuboidStructureXD` and the `MultiBlockLatticeXD`, both of which have similar functionality but a slightly different scope. These advanced data structures generate a patchwork consisting of many `BlockLatticeXD` structures that are presented behind a unified interface. Applications of these structures are MPI-parallel and memory saving simulations that do not allocate memory in chosen subdomains of the numerical grid.

Input / Output

The basic mechanism behind I/O operations in OpenLB is the serialization and unserialization of a `BlockLatticeXD` and a `DataFieldXD`. 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.

1.3.4 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 has been coordinating the project. Since 2006 the following persons have contributed source code to OpenLB:

Armani Arfaoui: **core:** performance improvements for D3Q19 BGK collision operator

Saada Badie: **core:** performance improvements for D3Q19 BGK collision operator

Lukas Baron (active): **utilities:** (parallel) console output, time and performance measurement, **dynamics:** porous media model, **functors:** concept, div. functors implementation

Vojtech Cvrcek (active): **functors:** 2D adaptation, **dynamics:** power law, **examples:** updates

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 (active): **core:** super data implementation **io:** new serializer and serializable implementation, vti writer, new vti reader, **functors:** new discrete indicator

Thomas Henn (active): **io:** voxelizer interface based on STL, **particles:** particulate flows

Fabian Klemens (active): **functors:** flux

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

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

Marie-Luise Maier (active): **particles:** particulate flows, frame change

Orestis Malaspinas: **boundaries:** alternative boundary conditions (Inamuro, Zou/He, Nonlinear FD), **dynamics:** alternative LB models (Entropic LB, MRT) item[Cyril Masquelier:] **functors:** indicator, smooth indicator

Albert Mink (active): **functors:** arithmetic, **io:** parallel VTK interface

Patrick Nathen (active): **dynamics:** turbulence modelling (advanced subgrid-scale models), **examples:** nozzle3d

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)

Robin Trunk (active): **dynamics:** parallel thermal, advection diffusion models

Peter Weisbrod (active): **dynamics:** parallel multi phase/component, **examples:** structure and showcases, phaseSeparationXd

Gilles Zahnd: **functors:** rotating frame functors

Asher Zarth (active): **core:** vector implementation

Simon Zimny: **io:** pre-processing: automated setting of boundary conditions

2 Using OpenLB for Applications

The general way of functioning in OpenLB follows a generic path.

1st Step: Initialization The converter between physical and lattice units is set in this step. The parameters for the simulation setup are chosen here, too, if they have not already been set at the beginning.

2nd Step: Prepare geometry The geometry is acquired, either from another file (a `.stl` file here) or from defining indicator functions. A mesh is created from this information, and the required geometry is prepared. 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. Some examples and applications that use a rather simple geometry skip this step.

3rd Step: Prepare lattice The lattice dynamics are set here. The type of dynamics are selected from the different implementations. 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 neighbouring voxels chosen. The boundary condition initialization is done to enable any kind of boundaries. The lattice is then defined in the function `prepareLattice`, with the boundary condition choices for every material number, and for which material numbers, which corresponding dynamics are applied. Here, only the kind of boundary (like Bouzidi, bounce-back, velocity, or pressure) is defined, not the profile function itself.

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` (steps 5, 6 and 7 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 results are printed.

5th Step: Definition of initial and boundary conditions The first of the three important functions called during the loop, `setBoundaryValues`, puts into practice the boundary functions' values. In some applications, it needs to refresh them during each time step, in others they 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, `.gif` files or `.vtk` files of the results at certain timesteps.

This general structure is the maintained throughout every OpenLB simulation, it is only the choices that are made which determine the simulation: every real modification is done in the called functions, to prepare the geometry, the converter, the lattice, and the boundary profiles. Every change has to match OpenLB's implementation. Consequently, new models might need changes or additions to the source code. For example, the classes defined in the code are always issued from a parent-class and have to match to the functionality, which may sometimes lead to unexpected issues to solve.

sectionLesson 1: - A typical application program structure: Implement your first OpenLB program

Unpack the OpenLB tar-ball on your system, and compile one of the example programs. If this is successful, create a directory for this tutorial at a location of your choice. Create a Makefile in this directory, according to the procedure explained in Section 1.3.2.

A few lines are invariably the same from one OpenLB program to another:

Listing 2.1: Framework of an OpenLB program

```
1 #include "olb2D.h"
  #ifndef OLB_PRECOMPILED // Unless precompiled version is used,
    #include "olb2D.hh" // include full template code
  #endif

6 using namespace olb;
```

Some lines in this program deserve additional comments:

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 3: 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. Including all this code slows down compilation (2D codes may take around 10 seconds to compile, and 3D codes around 30 seconds). If this overhead becomes too annoying during frequent development-compilation cycles, the code can be precompiled for the required data types. Although this topic is not covered in the tutorial, this short explanation should clarify the meaning of the the cryptic `#ifndef OLB_PRECOMPILED`.

Line 6: All OpenLB code is contained in the namespace `olb`.

Furthermore, for the examples to compile, the following declarations need to be included into Listing 2.1 between Line 4 and 6:

```

#include <vector>      // Some C++ libraries which are
#include <cmath>      //   required for the following
#include <iostream>   //   examples
4 #include <iomanip>
#include <fstream>

using namespace olb;           // OpenLB namespaces which are
using namespace olb::descriptors; // accessed in the
9 using namespace olb::graphics; // examples
using namespace std;         // Namespace of standard C++
                               library

```

At this point, the code for the simulation of a fluid flow can be inserted at the place of line 10. The following simple example represents a fluid initially at rest with a slightly increased particle density within a disk around the center. The flow is modeled through the single relaxation-time BGK model, and it evolves in a system with periodic boundaries. *(It should be pointed out that this example is only used to illustrate programming issues. The chosen initial condition does not really represent a physically meaningful state of an incompressible fluid. The example “works” because the LB model is contrived into adopting a compressible regime. Interpreting the results of a BGK model under in the context of compressible flows, however, raises numerous issues of its own that cannot be covered here. Thus, use the code for learning purposes, but don’t attribute too much meaning to the numerical result.)*

Listing 2.2: to be inserted at Line 10 of Listing 1

```

#define LATTICE D2Q9Descriptor
typedef double T;
int nx = 20;
int ny = 30;
5 int numIter = 100;
T omega = 1.;
T r      = 5.;

int main(int argc, char* argv[]) {
10  olbInit(&argc, &argv);
    // Insert the central part of your code here
    BlockLattice2D<T, LATTICE> lattice(nx, ny);
    BGKdynamics<T, LATTICE> bulkDynamics (
        omega,
15  instances::getBulkMomenta<T, LATTICE>()
    );
    lattice.defineDynamics(0, nx-1, 0, ny-1, &bulkDynamics);

    for (int iX=0; iX<nx; ++iX) {
20     for (int iY=0; iY<ny; ++iY) {

```

```

        T rho=1., u[2] = {0.,0.};
        if ((iX-nx/2)*(iX-nx/2) + (iY-ny/2)*(iY-ny/2) < r*r) {
            rho = 1.01;
        }
25     lattice.get(iX,iY).iniEquilibrium(rho,u);
    }
}

    for (int iT=0; iT<numIter; ++iT) {
30     lattice.collide();
        lattice.stream(true);
    }

    ImageWriter<T> imageWriter("leeloo");
35     imageWriter.writeScaledGif (
        "lesson1",
        lattice.getDataAnalysis().getVelocityNorm() );
}

```

A few explanations are again in order:

Line 1: Choice of a lattice descriptor. Lattice descriptors specify not only which lattice you are going to use (for 2D simulations, the current OpenLB release gives you no choice but D2Q9 anyway), but also potentially the nature of additional scalars, such as an external force field, for which memory needs to be allocated on a grid cell.

Line 2: Choice of double precision floating point arithmetic. Any other floating point type can be used, including built-in types and user-defined types which are implemented through a C++ class.

Lines 3-7: Constants to specify the dimensions of the $n_x \times n_y$ lattice and the total number `numIter` of iteration steps. The relaxation parameter ω is the reciprocal of the relaxation time τ . It determines the value of the shear viscosity ν of the fluid.

Line 10: This line is gratuitous in sequential programs, but it is required in the context of MPI-parallelism (which is explained in Lesson 10). As a general rule, you will always want your program to be ready for both sequential and parallel executions. It is therefore good practice to include this line as a matter of routine, in all cases.

Line 12: Instantiation of a `BlockLattice2D` object. At this point, memory for the $n_x \times n_y \times 9$ particle populations is allocated. If additional memory has been requested for external scalars (this is not the case here), this memory is also allocated during the instantiation of the `BlockLattice2D`.

Lines 13-16: The `Dynamics` object determines the implementation of the collision step on grid nodes, in this case BGK [15]. Objects of type `BGKdynamics` can be customized by indicating how the momenta of distribution functions (particle density,

velocity, etc.) should be computed. By choosing a specific `Momenta` object, one can, for example, implement boundary conditions in which the dynamics are the same as in the bulk, but the momenta are computed differently because of missing particle populations. In the present example, a default implementation is chosen for the computation of the momenta.

Line 17: The previously instantiated dynamics are to be used on all lattice nodes. The domain on which to instantiate the dynamics is indicated explicitly, the x -index ranging from 0 to `nx-1`, and the y -index from 0 to `ny-1`.

Line 25: Initialize particle populations at an equilibrium distribution, with slightly increased density inside a circle of radius `r`.

Line 30: At each iteration step, the collision specified by the variable `bulkDynamics` is applied to each grid node.

Line 31: After collision follows the streaming step. The boolean argument `true` indicates that boundaries are periodic.

Line 34: The `ImageWriter` offers a means of producing 2D images in PPM format. If the package `ImageMagick` is installed on your machine, you can also get GIF images. Four colormaps are available for each of the four elements (“earth”, “water”, “air”, “fire”) and one for the fifth element “leeloo” (see Ref. [3]).

Line 37: An object of type `DataAnalysis2D` is instantiated to extract the norm of the velocity from the numerical result. From this, an image is created with help of the `ImageWriter`, by rescaling the colormap to the range of values adopted by the velocity norm in the numerical result.

You can easily observe that boundary conditions are periodic by playing around with the code and producing images at various time steps. Alternatively, no-slip walls are implemented by calling the method `BlockLattice2D::stream()` in line 28 with an argument `false`. This is the default argument, and the method can therefore be invoked with no argument at all:

Listing 2.3: Substitutions to replace periodic boundaries by no-slip walls

```
lattice.collide();  
2 lattice.stream();
```

These no-slip walls are obtained through a so-called *halfway bounceback* mechanism: particle populations on boundary cells, which would leave the computational domain during streaming, stay on the cell and their value is copied to the particle population with opposite velocity vector instead. After this, the usual collision step is executed. No efficiency overhead is incurred for the implementation of this mechanism, because it is an automatic side-effect of the algorithm in OpenLB for the streaming step [4].

2.1 Lesson 2: Understand the BlockLattice

This second lesson starts with a response to the scream of indignation you emitted in Lesson 1, when you learned that each cell of a `BlockLatticeXD` carries along its own `Dynamics` object, and collision is triggered by some dynamic run-time mechanism. How could the OpenLB developers favor object-oriented mumbojumbo over efficiency, right there in the core of the library?

The truth is that the overhead incurred by delegating collision to an object (instead of hard-coding collision somewhere inside the loop over grid nodes) is completely irrelevant. The efficiency loss is minimal on all platforms on which OpenLB has been tested so far, and it is negligible in face of other big-picture efficiency considerations.

One such consideration is about the separation of collision and streaming at Line 28 and 29 of Listing 2.2. The question to ask, instead of nitpicking over object-oriented vs. non-object-oriented issues, is whether it is really necessary to step through memory twice; once to execute collision and once to execute streaming. As a matter of fact, there are several ways of avoiding this time-consuming double access to memory, one of which is implemented in OpenLB and documented in Ref. [4]. For an OpenLB user, doing this is as easy as replacing the collision-streaming sequence by a call to the method `collideAndStream()`:

Listing 2.4: Collision and streaming in one step for improved efficiency

```
// collision-streaming cycles
// lattice.collide();
3 // lattice.stream(true);
  lattice.collideAndStream(true);
```

Using the method `collideAndStream` is, of course, only possible when you don't need to compute or modify anything between collision and streaming. When this is the case, the use of this method can however reduce by as much as 40% the execution time of your code, depending on your hardware.

The `BlockLattice2D<T, LATTICE>` is basically a `nx-by-ny-by-q` array of variables of type `T`. The following code for example is valid (although it is bad practice, as explained below):

Listing 2.5: Direct access to values in a `BlockLattice2D`

```
1 int nx, ny, someX, someY, someF;
// <...> some code to initialize nx, ny, someX and someY
BlockLattice2D<T, LATTICE> lattice(nx,ny); // instantiate
  BlockLattice
T value = lattice.get(someX,someY)[someF]; // read values
lattice.get(someX,someY)[someF] = 0.; // write values
```

The method `BlockLattice2D<T, LATTICE>::get()` delivers an object of type `Cell<LATTICE>`, which contains storage space for the particle populations and, if required by the `LATTICE`

template, for additional scalars. The `Cell` offers many methods to read and manipulate the data. You are much more likely to use those methods in practice, rather than accessing particle populations directly as in Listing 2.5:

Listing 2.6: Manipulation of data through methods of a `Cell`

```

int nx, ny, someX, someY, someF;
// <...> some code to initialize nx, ny, someX and someY
BlockLattice2D<T, LATTICE> lattice(nx,ny); // instantiate
    BlockLattice
// <...> some code to initialize dynamics objects of the lattice
5 T velocity[2];
lattice.get(someX,someY).computeU(velocity); // compute velocity
velocity[0] = 0.;
lattice.get(someX,someY).defineU(velocity); // modify velocity

```

In this example, the method `Cell<T>::computeU()` computes the velocity on a cell for you, using its dynamics object. Conversely, the method `Cell<T>::defineU()` modifies the velocity by translating the particle populations into space of moments, modifying the moment of the velocity, and leaving the others as they are.

In addition to being more convenient, the access to the data in Listing 2.6 has a distinct advantage to the approach of Listing 2.5. In Listing 2.5 the data inside a `Cell<T>` is accessed directly, whereas in Listing 2.6 it is accessed indirectly through the dynamics object of the cell. Although direct data access works in simple data structures, such as the present `BlockLattice2D`, only indirect data access can be used in complicated data structures. When the code is, for example executed in parallel, you cannot access the data directly, because it might not be found on your processor. The dynamics object, on the other hand, is smart enough to locate the data on the right processor, and to instantiate MPI communication to access it.

Generally speaking, the methods of a `Cell<T>` are separated into two groups, one for direct data access, and one for indirect data access through the dynamics object. When using OpenLB as an application programmer, it is strongly recommended that you only make use of methods in the second group, in order for your code to be extensible. Methods of the first group are used by programmers who wish to extend the OpenLB library, for example by writing a class to implement a new type of dynamics. Most of the subsequent lessons are written for application programmers, and the code is written with extensibility in mind, for example, by insisting on the possibility for it to be run in parallel with minimal changes.

The following list details some useful methods to access the data of a `Cell<T>` indirectly through the dynamics object:

void iniEquilibrium(T rho, const T u[Lattice<T>::d])

Initialize all particle populations at an equilibrium distribution with density `rho` and velocity `u`.

T computeRho() const

Compute the particle density on the cell.

void computeU(T u[Lattice<T>::d]) const

Compute the velocity on the cell.

void computeStress (T pi[util::TensorVal<Lattice<T>>::n]) const

Compute the off-equilibrium stress-tensor $\Pi^{(1)}$ on the cell.

void computePopulations(T* f) const

Retrieve the particle populations and store them in a q -element C-array.

void computeExternalField(int pos, int size, T* ext) const

Retrieve the external scalars and store them in a C-array.

void defineRho(T rho)

Modify the populations such that the density yields `rho` and the other moments are unchanged.

void defineU(const T u[Lattice<T>::d])

Modify the populations such that the velocity yields `u` and the other moments are unchanged.

void defineStress(const T pi[util::TensorVal<Lattice<T>>::n])

Modify the populations such that the tensor $\Pi^{(1)}$ yields `pi` and the other moments are unchanged.

void definePopulations(const T* f)

Attribute new values to all populations. The argument `f` is a C-array with q elements.

void defineExternalField(int pos, int size, const T* ext)

Attribute new values to all external scalars.

The discussion of this lesson is also valid for 3D lattices, which are instantiated with the following instruction:

Listing 2.7: Instantiation of a 3D lattice

```
#define D3Q19Descriptor LATTICE
2 int nx, ny, nz;
// <...> initialization of nx, ny, nz
BlockLattice3D<T,LATTICE> lattice(nx,ny,nz);
```

The `BlockLattice2D` and the `BlockLattice3D` have different types, because they have distinct interfaces. The method `get()` for example requires 2 arguments in the 2D case and 3 arguments in 3D. The `Cell` class, an instance of which is delivered by the method `get()`, is however the same in 2D and 3D, although its template is instantiated with a different lattice descriptor (e.g. `D2Q9Descriptor` vs. `D3Q19Descriptor`). The above list of methods of the `Cell` is therefore valid in 3D as well.

2.2 Lesson 3: Define and use boundary conditions

The current OpenLB release offers five different 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. This instruction is the instantiation of a `OnLatticeBoundaryCondition` object:

Listing 2.8: Instantiation of `OnLatticeBoundaryCondition`

```

1 // Instantiate 2D boundary condition
  OnLatticeBoundaryCondition2D<T,D2Q9Descriptor>*
    boundaryCondition2D =
      createLocalBoundaryCondition2D(lattice);

// Instantiate 3D boundary condition
6 OnLatticeBoundaryCondition2D<T,D3Q19Descriptor>*
    boundaryCondition3D =
      createLocalBoundaryCondition3D(lattice);

```

Objects of type `OnLatticeBoundaryConditionXD` are used to attribute the role of boundary node to chosen nodes of the lattice. The following code configures a lattice in such a way that the rectangle following the lattice boundaries implements a boundary condition on the velocity.

Listing 2.9: Instantiation of velocity boundary condition along lattice boundaries

```

template<typename T>
void velocityBoundaryBox (
3   BlockLattice2D<T,D2Q9Descriptor>& lattice,
   OnLatticeBoundaryCondition2D<T,D2Q9Descriptor>& bc, T omega)
{
   int nx = lattice.getNx();
   int ny = lattice.getNy();
8   // top boundary
   bc.addVelocityBoundary1P(1,nx-2,ny-1,ny-1, omega);
   // bottom boundary
   bc.addVelocityBoundary1N(1,nx-2, 0, 0, omega);
   // left boundary
13  bc.addVelocityBoundary0N(0,0, 1, ny-2, omega);
   // right boundary
   bc.addVelocityBoundary0P(nx-1,nx-1, 1, ny-2, omega);

   // Corner nodes
18  bc.addExternalVelocityCornerNN(0,0, omega);
   bc.addExternalVelocityCornerNP(0,ny-1, omega);
   bc.addExternalVelocityCornerPN(nx-1,0, omega);
   bc.addExternalVelocityCornerPP(nx-1,ny-1, omega);

23 // Make the lattice ready for simulation

```

```

    lattice.initialize();
}

```

When boundary nodes are instantiated, it is necessary to specify the orientation of the boundary through the normal vector that points outside of the domain. The instruction `addVelocityBoundary1P` refers to a boundary whose normal is in positive y -direction (P stands for “positive”, and indexes are numbered as 0 for the x -index and 1 for the y -index). For external corners, the expression NN refers to any boundary vector whose opposite direction points inside the numerical domain. In this case, this boundary vector points in negative x -direction and negative y -direction. The term **External** in the method `addExternalVelocityCornerNN` refers to the fact that the domain boundaries are convex shaped. Corners of concave shaped boundaries are instantiated with methods of the form `addInternalVelocityCornerXX`, where X stands again for N or P and indicates the direction of a vector pointing outside the numerical domain.

Pressure boundaries are instantiated just as easily by replacing the word **Velocity** by **Pressure** in the methods of the `OnLatticeBoundaryCondition` object.

Things are slightly more complicated in 3D, where edges also need separate treatment. Edges are locations where two boundary surfaces that are orthogonal to each other meet. The following are typical instructions one may use in the 3D case. In 3D, the instruction `addVelocityBoundary0N` instantiates a plane boundary domain in negative x -direction (a left boundary). It takes 6 arguments, in addition to the `omega`-argument, in order to delimit the plane like a sub-volume with one degenerate space direction. The instruction `addExternalVelocityEdge0NP` instantiates an edge whose outward-pointing normal vector is in the 0-plane (in the plane in which $x = 0$) and which points in negative y - and positive z -direction. Counting of indexes is cyclic: the instruction `addExternalVelocityEdge1NP` denotes an edge with normal vector in the $y = 0$ -plane and with negative z - and positive x -direction. The **Edge** instructions also take 6+1 arguments, because they treat the edge like a sub-volume with two degenerate directions. In 3D, there are external and internal corners, and there are external and internal edges.

Although setting up the geometry of the numerical domain can be somewhat bothersome, especially in 3D, this is a one-time job. Once the setup is completed, specifying the required velocity and density on boundaries is straightforward. This is done through a call to the method `defineVelocity` or `defineDensity` of the corresponding cell. You may remember from LESSON 2, that on normal lattice Boltzmann nodes, these methods modify the value of particle populations in order to obtain the required velocity/density. On boundary nodes, the rules are different. Here, particle populations are *not modified* (which is necessary, as you may want to change the boundary velocity during a simulation, without tampering with the particle populations). On velocity boundaries, the method `defineVelocity` modifies the required velocity value for the boundary, whereas `defineDensity` has no effect. On pressure boundaries, the method `defineVelocity` has no effect and `defineDensity` picks out the required density value on the boundary. It should be pointed out that although the domain geometry is specified piece-wise (plane per plane, edge per edge, and corner per corner), the velocity/density can be adapted individually on every node. Furthermore, accessing parameters of the boundary on a per-

cell base is convenient, because it does not require the programmer to distinguish any more between plane boundaries, edges or corners. Finally, the choice of the velocity/-density value is not static: it can be altered at every time step to model time-dependent boundaries.

The following is a list of available boundary conditions. Instead of showing the actual class name of the boundary condition, the list indicates the names of functions that generate the boundary condition, as these are the ones you are likely to access as an end user. The `X` is a placeholder for 2 respectively 3, as all boundary conditions are implemented in 2D and 3D.

`createLocalBoundaryConditionXD`

This is the default local boundary condition. It implements a regularized boundary [25], which tends to be numerically stable in a last range of regimes.

`createInterpBoundaryConditionXD`

This is the default non-local boundary condition. It is based on the algorithm proposed by Skordos [31], and uses a finite difference scheme over adjacent neighbors to evaluate velocity gradients.

`createZouHeBoundaryConditionXD`

The local boundary condition introduced by Zou and He [35]. It is very accurate, especially in 2D simulations, but can have stability issues.

`createInamuroBoundaryConditionXD`

The local boundary condition by Inamuro Inamuro et al. [24]. It is very accurate in 2D and 3D, but can have stability issues. In 3D, it is slower than other boundary conditions, because it solves an implicit equation at every time step.

`createExtendedFdBoundaryConditionXD`

The approach is the same as in the boundary condition generated by `createInterpBoundaryConditionXD`, but this time, non-linear velocity terms of the Chapman-Enskog expansion are taken into account. This is rarely useful, but can make a difference in a very low Mach-number regime.

It should be clear by now, how powerful the abstraction mechanism of the “`OnLatticeBoundaryConditionXD`” objects is. With the help of this mechanism, 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.

This mechanism for the instantiation of boundary conditions is generic and easy to use, but it makes sense only in quite regular geometries. In irregular geometries, even if you agree on using a staircase approximation of domain boundaries, you will experience a hard time attributing the right boundary type to each cell. Although off-lattice boundaries are under investigation in the OpenLB project, they are not currently available. If your irregular domain boundaries implement a no-slip condition, your current best bet is

to implement them through a fullway bounce-back dynamics. In this approach, particle populations that are opposite to each other are swapped at each iteration step, and no additional collision is executed. The advantage of this procedure is that it is independent of the orientation of the domain. The following code implements for example a circular obstacle with no-slip walls in the center of a 2D domain:

Listing 2.10: Implementation of a bounce-back cylinder in the domain center

```

<...> definition of the types T and DESCRIPTOR
int nx, ny, r;
<...> initialization of nx and ny, r
BlockLattice2D<T,DESCRIPTOR> lattice(nx,ny);
5 <...> setup of the lattice
  for (int iX=0; iX<nx; ++iX) {
    for (int iY=0; iY<ny; ++iY) {
      if ((iX-nx/2)*(iX-nx/2) + (iY-ny/2)*(iY-ny/2) < r*r) {
10         lattice.defineDynamics(iX,iX,iY,iY,
           &instances::getBounceBack<T,D2Q9Descriptor>() );
      }
    }
  }
}

```

2.3 Lesson 4: Convert between lattice and physical units

Fluid flow problems are usually given in a system of metric units. For example consider a cylinder of diameter 3cm in a fluid channel with average inflow velocity of 4m/s . The fluid has a kinematic viscosity of $0,001\text{m}^2/\text{s}$. 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 unity. Therefore, when setting up a simulation, a conversion directive has to be defined that takes care of translating variables from physical units into lattice units and *vice versa*. In OpenLB, all these conversions are handled by a class called `LBconverter`. An instance of the `LBconverter` is generated with some reference values of the simulation and the desired discretization parameters. It provides a set of conversion functions, to enable a fast and easy way to convert between physical units 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. The converter internally builds a "dimensionless" system of units in which the characteristic values are one. The Reynolds number Re is an important parameter of this system. Furthermore, two discretization parameters $latticeL$ and $latticeU$ are provided to the converter. $latticeL$ is the discrete space intervall in physical units and from

this the dimensionless discretization parameter δ_x is determined: $\delta_x = \text{latticeL}/\text{charL}$. latticeU sets the relation between the discretization parameters for space δ_x and time δ_t in dimensionless units: $\text{latticeU} = \delta_u = \delta_t/\delta_x$. Instead of δ_t , those working with LBM often like to specify latticeU . One reason for this is that latticeU is proportional to the Mach number, and its choice is important to control compressibility effects.

Once the converter is initialized, its methods can be used to convert various quantities such as velocity, 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:

Listing 2.11: Use of Lbconverter in a 3D problem

```

<...> definition of type T
2  int dimension      = 3;
   T latticeL        = (T) 0.003;
   T latticeU        = (T) 0.02;
   T charNu          = (T) 0.001;
   T charL           = (T) 0.03;
7  T charU           = (T) 4.;
   T charRho         = (T) 1.;
   T pressureLevel   = (T) 0.;

Lbconverter<T> converter(
12  dimension, latticeL, latticeU,
    charNu, charL, charU, charRho, pressureLevel
);
writeLogFile(converter, "converterLog.dat");
cout << converter.getRe() << endl;
17 T omega = converter.getOmega();
<...> simulation
<...> evaluation of latticeRho at the back and the front of the
    cylinder
T pressureDrop =   converter.physPressure(latticeRhoFront)
                  - converter.physPressure(latticeRhoBack);

```

Line 2: Specify discretization parameters and characteristic values.

Line 11: Instantiate a Lbconverter object.

Line 15: Write simulation parameters and conversion factors in a logfile.

Line 16: Print the Reynolds number Re.

Line 17: The method `getOmega` computes first the viscosity in lattice units, and then the relaxation parameter ω .

Line 20: The converter automatically calculates the pressure values from the local density.

2.4 Lesson 5: 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. Instead, a `BlockLatticeXD` sums them up internally, and in this way keeps track of the average density, average energy (half the square of the velocity norm) and the maximum value of the velocity norm. 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.

One needs to take care to properly interpret the value of the discrete time to which those quantities correspond. Imagine your simulation is at a discrete time t . After execution of a collision and a streaming step, it is taken from time t to time $t + 1$. If after this you evaluate, for example the velocity at a point using the command `lattice.get(iX,iY).computeU(velocity)`, the computed quantity lives at a time $t + 1$ of the system. The values of the internal statistics, such as `lattice.getStatistics().getAverageEnergy()` correspond however to the discrete time t , because they were evaluated prior to the previous streaming step. This time shift between the state of the system and the value of the internal statistics can be confusing, and for this reason it would have made sense to avoid computing the statistics. On the other hand, keeping track of the statistics takes a negligibly small amount of time. This feature is therefore included in OpenLB out of efficiency considerations, and out of convenience, as it offers an easy means of monitoring the well behaving of a simulation.

Lattice cells whose dynamics is bounce-back, generated by

`instances::getBounceBack<T,LATTICE>()`,

and cells that don't execute any collision step, generated by

`instances::getNoDynamics<T,LATTICE>()`

don't contribute to the internal statistics of the lattice. The same holds for subdomains for which, by using the approach taught in Lesson 9, no memory is allocated.

Often, the information provided by the statistics of a lattice is not sufficient, and you would like to treat the numerical result more generally. To do this, you can extract data cell-by-cell from the `BlockLatticeXD` and store it into a scalar- or vector/tensor-valued matrix, named `ScalarFieldXD` in the first case and `TensorFieldXD` in the second. During parallel program execution, these matrices are parallelized, which makes it very efficient to analyze large data sets on a parallel machine. The data can then be further analyzed, for example by computing reductions, such as the average value. Alternatively, its content can be stored to disk in a binary VTK format for analysis with an external

tool. Extraction of numerical data from a `BlockLatticeXD` into a `ScalarFieldXD` / `VectorFieldXD` is taken care of by the `DataAnalysisXD` class.

The most straightforward way of visualizing the data is to produce a 2D snapshot of a scalar field. OpenLB creates images of format PPM. On a system of the Unix/Linux family with the package `ImageMagick` installed, it further supports automatic conversion into the more common GIF format (note that `ImageMagick` is open sourced, and that it is part of all major Linux distributions). The following example illustrates how a snapshot of the vorticity distribution in a 2D simulation is created:

Listing 2.12: Produce a GIF image from 2D data

```
// <...> Create and initialize a variable lattice
//           of type BlockLattice2D<T,D2Q9Descriptor>
DataAnalysisBase2D<T,D2Q9Descriptor> const& analysis
4   = lattice.getDataAnalysis();
ImageWriter<T> imageWriter("earth");
imageWriter.writeScaledGif("vorticity", analysis.getVorticity,
                           200, 200);
```

Line 3: Require an analysis object from the lattice. Alternatively, an instance of the class `DataAnalysisXD` could be prepared manually. The advantage of requiring it from the lattice is that, among different implementations of the class `DataAnalysisXD`, the most efficient one is automatically picked out for you. This distinguishes, for example, between sequential and parallel lattices.

Line 5: Prepare for creation of an image with the colormap "earth".

Line 6: Calculate the vorticity on every cell, and visualize it as a GIF image. The colormap is rescaled to fit the range of vorticity values. The dimension of the image is rescaled to fit into a 200×200 bounding box.

Producing 2D images is also useful in 3D simulations. In this case you can extract data on a plane orthogonal to one of the coordinate axes and produce an image from it. This is done through the `slice` methods of data fields:

Listing 2.13: Produce a GIF image from 3D data

```
// <...> Create and initialize a variable lattice
//           of type BlockLattice3D<T,D3Q19Descriptor>
3 DataAnalysisBase3D<T,D3Q19Descriptor> const& analysis
   = lattice.getDataAnalysis();
ImageWriter<T> imageWriter("earth");
// Extract a slice of the plane defined by z=0
int slicePos=0;
8 imageWriter.writeScaledGif (
   "vorticity", analysis.getVorticity.sliceZ(slicePos), 200, 200
   );
```

Although the computation of statistics and the production of 2D images are very useful, they are not always sufficient to extract all the required information from the simulation. When a detailed analysis is required, it makes sense to resort to an external tool that performs postprocessing of numerical data. For this, the data can be stored in a file in a VTK format. The function `writeVTKData3D` stores a scalar field and a vector field in the same VTK file:

Listing 2.14: Produce a VTK file from 3D data

```

1 // <...> Create and initialize a variable lattice
  //   of type BlockLattice3D<T,D3Q19Descriptor>
  DataAnalysisBase3D<T,D3Q19Descriptor> const& analysis
    = lattice.getDataAnalysis();
  writeVTKData3D( "lesson5",
6                 "vorticity", analysis.getVorticityNorm(),
                 "velocity", analysis.getVelocity(), 1., 1. );

```

The open source software `Paraview` [5], for example, is very useful for the visualization of 3D data contained in such a file.

2.5 Lesson 6: Use an external force

In simulations, the dynamics of a fluid is often driven by a force field (gravity, intermolecular interaction, etc.) which is space- and time-dependent, and which is possibly computed from an external source, independent of the LB simulation. In order to optimize memory access and to minimize cache-misses, the value of this force can be stored in a cell, adjacent to the particle populations. This is achieved by specifying external scalars in the lattice descriptor (see also Lesson 7). `OpenLB` offers, by default, the two descriptors `ForcedD2Q9Descriptor` and `ForcedD3Q19Descriptor`. The dynamics `ForcedBGKdynamics` accesses the force term defined by these descriptors, and implements a LB dynamics with body force. The algorithm is taken from Ref. [20] to guarantee second-order accuracy even when the force field is space and time dependent. An example for the implementation of a LB simulation with force term is found in the code `examples/poiseuille2d/forced`.

2.6 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 models. 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 flavours, most of which inherit from a common interface `BlockStructureXD`. The most common lattice is the regular `BlockLatticeXD`, which is replaced by the `MultiBlockLatticeXD` 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.

2.7 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:

Listing 2.15: Store and load the state of the simulation

```

int nx, ny;
<...> initialization of nx and ny
3 BlockLattice2D<T, DESCRIPTOR> lattice(nx, ny);
  // load data from a previous simulation
  lattice.load("simulation.checkpoint");
<...> run the simulation
  // save data for security, to be able to take up
8 // the simulation at this point later
  lattice.save("simulation.checkpoint");

```

Checkpointing is also illustrated in the example programs `bstep2D` and `bstep3D` (Section 10.3).

2.8 Lesson 9: Save memory when domain boundaries are irregular

It is possible in OpenLB to allocate several lattices of type `BlockLatticeXD` and hide them behind a common interface, to treat them as the components of a larger lattice.

This technique can be used to achieve parallelism, as it is described in the next lesson. Another application is the creation of lattices in which memory is allocated in selected subdomains only. This is useful for the simulation of flows with complicated domain boundaries, as no memory needs to be allocated outside the domain. An example program for this technique is under development, but is not yet available in the current release.

2.9 Lesson 10: Run your programs on a parallel machine

OpenLB programs can be executed on a parallel machine with distributed memory, based on MPI. The approach taught in this lesson uses `MultiBlockLatticeXD`, which inherits the interface of `BlockStructureXD`, and therefore behaves like a common, non-parallelized lattice. All techniques described in the previous lessons can be used with the `MultiBlockLatticeXD` as well, and thus work both in sequential and parallel programs. The only modification you are required to do, is to switch between `BlockLatticeXD` and `MultiBlockLatticeXD`. This can be achieved through a precompiler directive, as in the following code:

Listing 2.16: `MultiBlockLattice2D` for MPI-parallel programs

```

1  int nx, ny;
   <...> initialization of nx and ny
   #ifndef PARALLEL_MODE_MPI // sequential program execution
       BlockLattice2D<T, DESCRIPTOR> lattice(converter.getNx(),
                                           converter.getNy() );
6  #else // parallel program execution
       MultiBlockLattice2D<T, DESCRIPTOR> lattice (
           createRegularDataDistribution( converter.getNx(),
                                         converter.getNy() ) );
   #endif

```

To compile an OpenLB program for parallel execution using MPI, modify the file named `Makefile.inc`, found in the OpenLB root directory, by removing the hashes before the lines: `#CXX := mpic++`, and `#PARALLEL_MODE := MPI`. The modified lines are shown in Listing 2.17. Execute `make clean` and `make cleanbuild` within the desired program directory to eliminate previously compiled libraries, and recompile the program by typing 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, and in this case, is set to 2.

Listing 2.17: `Makefile.inc` edits for MPI-parallel programs

```

CXX           := g++
#CXX          := icpc -D__aligned__=ignored
#CXX          := mpiCC
CXX           := mpic++
5  ...

```

```
PARALLEL_MODE := OFF
PARALLEL_MODE := MPI
#PARALLEL_MODE := OMP
#PARALLEL_MODE := HYBRID
```

3 Compilation

3.1 Linux

This compilation starts with a clean Ubuntu 14.04 LTS system. Before installing any new software, run

```
sudo apt-get update
```

to update the package lists, so that the most recent versions of the packages will be installed.

Next, install the g++ compiler, which you will need to compile C++ programs:

```
sudo apt-get install g++
```

To benefit from the efficient parallelization, you will probably want to run the program on more than one core, so it is recommended to install Open-MPI:

```
sudo apt-get install openmpi-bin openmpi-doc libopenmpi-dev
```

For visualization purposes you can use, for example, the following open source software:

```
sudo apt-get install paraview
```

```
sudo apt-get install imagemagick
```

Paraview is an application built on top of the *Visualization Tool Kit* (VTK) libraries which can read VTI-files written by OpenLB. With *imagemagick*, OpenLB can directly produce gif-files during simulation.

Finally, go into the root folder of OpenLB and type

```
make
```

to compile the software library and all examples. If your system is set up correctly, you should see a lot compiler messages but no errors.

3.2 Mac

3.3 Windows

An installation guide for Windows using Cygwin can be found in the technical report TR4: Installing OpenLB in Windows/Cygwin [6].

4 Geometry

4.1 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 lies on the boundary or in the fluid domain or whether it is superfluous in the computations. Figure 4.1 illustrates this using the example of an external flow. The benefit of using material numbers in flow simulations is the automatic determination of fluid 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.

4.2 Indicator functions

OpenLB provides several functors (see Section 8) for the creation of basic geometric entities such as cuboids, circles, spheres, cones etc. All indicator functors inherit from `IndicatorFXD` and therefore, contain the following functions:

- `std::vector<T> operator()(std::vector<S> in)`: Takes a X-dimensional vector `in` in SI coordinates and returns either `true` if the point is inside the geometry, `false` otherwise.
- `std::vector<S>& getMin()`: Returns the lower corner of an axis aligned bounding box.
- `std::vector<S>& getMax()`: Returns the upper corner of an axis aligned bounding box.
- `bool distance(S& distance, const std::vector<S>& origin, const std::vector<S>& direction, int iC = -1)`: Stores the distance from the origin to the closest boundary in `direction` and returns `true` if found.

The geometries already implemented are:

- 2 Dimensions:
 - `IndicatorCuboid2D`
 - `IndicatorCircle2D`
- 3 Dimensions:
 - `IndicatorCuboid3D`


```

void rename(int fromM, int toM, int testM, std::vector<int>
    testDirection);
/// renames all boundary voxels of material fromBcMat to toBcMat
    if two neighbour voxels in the direction of the discrete normal
        are fluid voxels with material fluidM in the region where the
        indicator function is fulfilled
10 void rename(int fromBcMat, int toBcMat, int fluidMat, IndicatorF3D
    <bool,T>& condition);

```

4.3 Creating a Geometry

With the information in the last section, a computational domain is created in 6 simple steps (see also Fig 4.2):

1. Step: Create indicator by
 - a) Reading an STL-file.
 - b) Pre-defined geometric primitives, cf. Section 4.2.
 - c) Combinations of indicators (+, -, ·).
 - d) Other operators on indicators (e.g. extra layer for boundary).
2. Step: Construct cuboidGeometry. During construction cuboids will be automatically removed, shrunk and weighted for a good load balance.
3. Step: Construct loadBalancer.
4. Step: Construct superGeometry.
5. Step: Set *material numbers*.
6. Step: Construct superLattice.

```

/// 1. Step: Create indicator
STLreader<T> stlreader("filename.stl", voxelSize);
Cone3D<bool, T> cone(center1, center2, radius1, radius2);
Layer3D<bool, T> boundaryLayer(cone, voxelSize);
5
/// 2. Step: Construct cuboidGeometry.
CuboidGeometry3D<T> cuboidGeometry(indicator, voxelSize,
    noOfCuboids);

/// 3. Step: Construct loadBalancer.
10 HeuristicLoadBalancer<T> loadBalancer(cuboidGeometry);

/// 4. Step: Construct superGeometry.
SuperGeometry3D<T> superGeometry(cuboidGeometry, loadBalancer);

```

```

15  /// 5. Step: Set material numbers.
    /// set material number 2 for whole geometry
    superGeometry.rename(0,2,geometryIndicator);
    /// change material number from 2 to 1 for inner (fluid) cells, so
    /// that only boundary cells have material number 2
    superGeometry.rename(2,1,1,1,1); or
20  superGeometry.rename(2,1,fluidIndicator);
    /// additional material numbers for other boundary conditions
    superGeometry.rename(2,3,1,inflowIndicator);
    superGeometry.rename(2,4,1,outflowIndicator0);
    superGeometry.rename(2,5,1,outflowIndicator1);
25
    /// 6. Step: Construct superLattice.
    SuperLattice<T, DESCRIPTOR> sLattice(superGeometry);

```

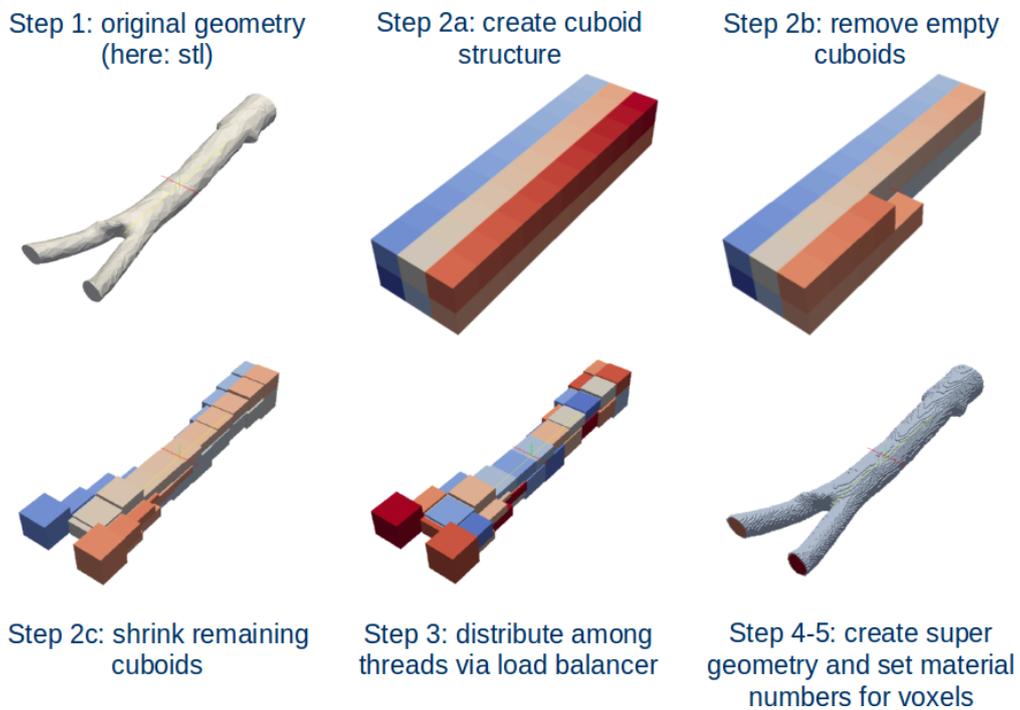


Figure 4.2: 6 steps to create a Geometry.

4.4 Excursion: 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 [7]. An introduction to modelling with FreeCAD can be found for example in <http://www.youtube.com/watch?v=6RxHCR7TLtI>. The general

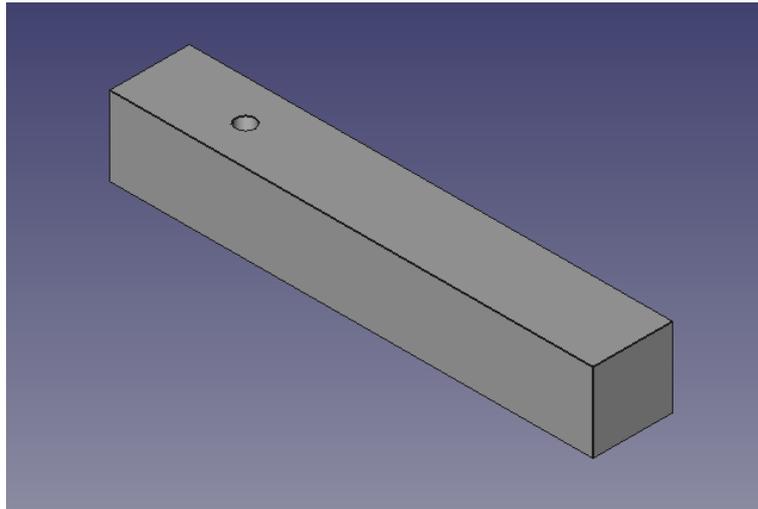


Figure 4.3: The geometry for the example `cylinder3d` from Section 10.5 opened in FreeCAD.

procedure is mostly similar to the following description.

Firstly, a 2D drawing is created on a selected plane (e. g. the xy plane) using circles and polygons. In the next step a “height” is assigned to it 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 4.3 is not very difficult, the more complex geometry of a formula one car, however, can be a challenging and time consuming task.

5 Lattice Boltzmann Models and Core Data Structures

5.1 Concept – Data Organization

5.1.1 Cell – BlockLattice – SuperLattice

LBM, in its most widely accepted formulation, is executed on a regular, homogeneous lattice Ω_h with equal grid spacing $h \in \mathbb{R}_{>0}$ in all directions. When numerical constraints require that a given problem is solved on an inhomogeneous grid, it is common to adopt a so-called *multi-block* approach: the computational domain is partitioned into sub-grids with different levels of resolution, and the interface between those sub-grids is handled appropriately. This approach appears to respect the spirit of LBM well and leads to implementations that are both elegant and efficient, since the execution on a set of regular blocks is relatively fast compared to an unstructured grid representation of the whole geometry. For complex domains, a multi-block approach provides another advantage. A given domain can be represented by a certain number of regular blocks, which leads to cheap execution times on the one hand, and to a sparse memory consumption on the other hand. Furthermore, it encourages a particularly efficient form of *data parallelism*, in which an array is cut into regular pieces and distributed over the nodes of a parallel machine. As a result, LB applications can even be run on large parallel machines with a particularly satisfying gain of speed.

The same spirit is adopted in the OpenLB package, in which the basic datastructure is a **BlockLattice**, which represents a regular array of **Cells**. In each **Cell**, the q variables for the storage of the discrete velocity distribution functions f_i ($i = 0, 1, \dots, q - 1$) are contiguous in memory. This data structure is encapsulated by a higher level, object-oriented layer. The purpose of this layer is to handle groups of **BlockLattice**, and to build higher level software constructs in a transparent way. Those constructs are called **SuperLattices**.

5.1.2 Descriptor

5.1.3 Dynamics

The core of OpenLB consists of a simple and efficient array-like construct called a **BlockLattice**. This object executes an LB algorithm in a very traditional sense, i.e.

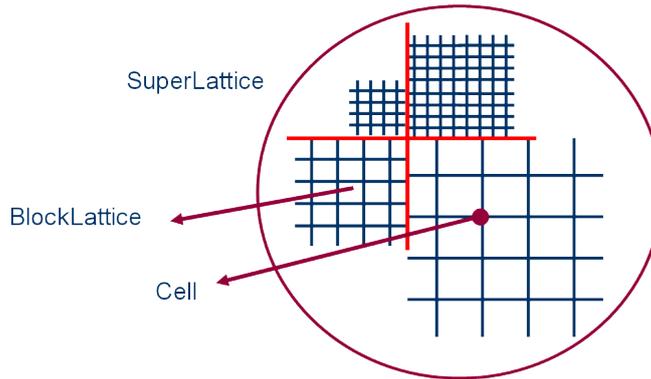


Figure 5.1: Data structures in OpenLB: A number of `BlockLattices` build a `SuperLattice` to adopt higher level software constructs like multi-block, grid refined lattices and parallelised lattices

the lattice Boltzmann equation is split into two equations, namely the *collision step*:

$$\tilde{f}_i^h(t, \vec{r}) = f_i^h(t, \vec{r}) - \frac{1}{3\nu + 1/2} \left(f_i^h(t, \vec{r}) - M_{f_i^h}^{eq}(t, \vec{r}) \right) \quad \text{in } I_h \times \Omega_h \times Q \quad (5.1)$$

and the *streaming step (propagation step)*:

$$f_i^h(t + h^2, \vec{r} + h^2 \vec{v}_i) = \tilde{f}_i^h(t, \vec{r}) \quad \text{in } I_h \times \Omega_h \times Q . \quad (5.2)$$

All `Cells` of the `BlockLattice` are iteratively parsed, and a local collision step is executed, followed by a non-local streaming step. The streaming step is independent of the choice of lattice Boltzmann dynamics and remains invariant. On the other hand, the collision step determines the physics of the model and can be configured by the user, by attributing a fully configurable dynamics object to each `Cell`. In this way, it is easy to implement inhomogeneous fluids which use a different type of physics from one `Cell` to another. For each time step the collision and streaming step can be executed separately in two loops over all `Cells` or only in one. Both versions are implemented in OpenLB. Yet, for many applications the method fusing the two loops is preferable.

Although this concept of a `BlockLattice` is neat and should please the programmer by being conceptually close to the theory of LBM, it is not sufficiently general to address all possible issues arising in real life. As a case in point, some boundary conditions are non-local and need to access neighbouring nodes. Therefore, their implementation does not fit into the framework of a `BlockLattice` explained previously. The philosophy of OpenLB takes for granted that such situations, although they arise, take place in spatially confined areas only, for example the domain boundaries. They may therefore be implemented by slightly less efficient means, without spoiling the overall efficiency of the code. Their execution is taken care of by a post-processing step, which, instead of

stepping over the whole lattice a second time, parses selected `Cells` only.

5.2 Classic BGK Model

5.3 MRT Model

5.4 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:

$$K = -\frac{Q\mu L}{\Delta P}, \quad (5.3)$$

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

The porosity-value $d \in [0, 1]$ is a lattice-dependent value, $d = 0$ means the medium is solid and $d = 1$ denotes a liquid. According to Brinkman [13, 14], Borrvall and Petersson [11] and Pinggen et al. [28], the Navier-Stokes equation is transformed (see Dornieden [19] and Stasius [32]). The discrete formulation of d describes a flow region by its permeability:

$$d = 1 - h^{dim-1} \frac{\nu_{LB} \tau_{LB}}{K} \quad (5.4)$$

τ_{LB} is the relaxation time, ν_{LB} is the discrete kinematic viscosity and h is the length. Therefore $K \in [\nu_{LB} \tau_{LB} h^{dim-1}, \infty]$. To describe the porosity or permeability of a medium, a descriptor for porosity must be used, such as:

```
#define DESCRIPTOR 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:

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

In function `setBoundaryValues`, the initial porosity value and external field is defined:

```
void setBoundaryValues(..., T physPermeability, int dim, ...) {
    // d in [0,1] is a lattice-dependent porosity-value
    // depending on physical permeability K = physPermeability
    T d = converter->latticePorosity(physPermeability);
}
```

```

5 AnalyticalConst3D<T,T> porosity(d);
  sLattice.defineExternalField(superGeometry, 3,
    DESCRIPTOR<T>::ExternalField::porosityIsAt, 1, porosity);
  ...
}

```

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

```

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

```

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

```

// converts a physical permeability K to a lattice-dependent
  porosity d
2 // (a velocity scaling factor depending on Maxwellian
  distribution
  // function), needs PorousBGKdynamics
  T latticePorosity(T K) const
  { return 1 - pow(physLength(),getDim()-1)*getLatticeNu()*getTau()/
    K; }

7 // converts a lattice-dependent porosity d (a velocity scaling
  factor
  // depending on Maxwellian distribution function) to a physical
  // permeability K, needs PorousBGKdynamics
  T physPermeability(T d) const
  { return pow(physLength(),getDim()-1)*getLatticeNu()*getTau()/(1-d
    ); }

```

5.5 Power Law Model

The two most common deviation from Newton's Law observed in real systems are pseudo-plastic fluids and dilatant fluids. By 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 this two type of deviations, was proposed by *de Waele* and *Ostwald* and is called the Power Law model that is defined by the viscosity as

$$\mu = m\dot{\gamma}^{n-1} . \quad (5.5)$$

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

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

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

```
#define DESCRIPTOR DynOmegaD2Q9Descriptor
```

In function `setBoundaryValues`, the initial same ω -argument is defined:

```
AnalyticalConst2D<T,T> omega0(converter.getOmega());
  sLattice.defineExternalField(superGeometry, 1, DESCRIPTOR<T
    >::ExternalField::omegaBeginsAt, 1, omega0);
  sLattice.defineExternalField(superGeometry, 2, DESCRIPTOR<T
    >::ExternalField::omegaBeginsAt, 1, omega0);
4  sLattice.defineExternalField(superGeometry, 3, DESCRIPTOR<T
    >::ExternalField::omegaBeginsAt, 1, omega0);
  sLattice.defineExternalField(superGeometry, 4, DESCRIPTOR<T
    >::ExternalField::omegaBeginsAt, 1, omega0);
```

In the main function, the power law dynamics is defined:

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

In 5.1 the kinematic viscosity is not more constant and then also the ω -argument is not more constant. With using the power law model 5.5 the kinematic viscosity is computed in each step as

$$\nu = \frac{1}{\rho} m \dot{\gamma}^{n-1} . \quad (5.6)$$

The shear rate $\dot{\gamma}$ is possible to compute with using the second invariant of the strain rate tensor D_{II}

$$\dot{\gamma} = \sqrt{2D_{II}} , \quad (5.7)$$

where

$$D_{II} = \sum_{\alpha,\beta=1}^d E_{\alpha\beta} E_{\alpha\beta} , \quad (5.8)$$

where

$$E_{\alpha\beta} = - \left(1 - \frac{1}{\tau} \right) \frac{1}{2\rho\nu} \sum_{i=0}^{q-1} f_i^h \xi_{i\alpha} \xi_{i\beta} . \quad (5.9)$$

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

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

5.6 External Force

In simulations, the dynamics of a fluid is often driven by a force field (gravity, intermolecular interaction, etc.) which is space- and time-dependent, and which is possibly computed from an external source, independent of the LB simulation. In order to optimize memory access and to minimize cache-misses, the value of this force can be stored in a cell, adjacent to the particle populations. This is achieved by specifying external scalars in the lattice descriptor (see also Lesson 7). OpenLB offers by default the two descriptors `ForcedD2Q9Descriptor` and `ForcedD3Q19Descriptor`. The dynamics `ForcedBGKdynamics` accesses the force term defined by these descriptors, and implements a LB dynamics with body force. The algorithm is taken from Ref. [20] to guarantee second-order accuracy even when the force field is space and time dependent. An example for the implementation of a LB simulation with force term is found in the code `examples/poiseuille2d/forced`.

As an alternative, the velocity shift forcing scheme developed by Shan and Chen [29] and improved by Shan and Doolen [30] is also implemented and can be accessed using `ForcedShanChenBGKdynamics`.

5.7 Multiphysics Couplings

5.7.1 The Shan-Chen Model

For the simulation of both multiphase and multicomponent flow the Shan-Chen model is implemented in OpenLB. Since its first introduction [29], many variants of the model have been developed. In this implementation, there are several forcing schemes [20, 30] and interaction potentials to choose from.

5.7.2 Implementation of Shan-Chen Two-phase Fluid

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

```
SuperLattice3D<T, DESCRIPTOR> sLattice(superGeometry);
```

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

```
ForcedShanChenBGKdynamics<T, DESCRIPTOR> bulkDynamics1 (
omega1, instances::getExternalVelocityMomenta<T, DESCRIPTOR>() );
```

Possible choices for the dynamics are `ForcedBGKdynamics` and `ForcedShanChenBGKdynamics`.

Then the interaction potential is chosen:

```
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.

```
const T G = -120.;
ShanChenForcedSingleComponentGenerator3D<T, DESCRIPTOR> coupling(
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();
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`.

5.7.3 Implementation of Shan-Chen Two-component Fluid

Two lattice instances are needed – one for each component (though there is still only one geometry):

```
SuperLattice3D<T, DESCRIPTOR> sLatticeOne(superGeometry);
SuperLattice3D<T, DESCRIPTOR> sLatticeTwo(superGeometry);
```

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

```
ForcedShanChenBGKdynamics<T, DESCRIPTOR> bulkDynamics1 (
omega1, instances::getExternalVelocityMomenta<T, DESCRIPTOR>() );
```

```

3 ForcedShanChenBGKdynamics<T, DESCRIPTOR> bulkDynamics2 (
  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.

```

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

```

The interaction strength G has to be positive. If the cosen 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:

```

sLatticeOne.communicate();
sLatticeTwo.communicate();
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`.

5.7.4 Thermal Fluid with Boussinesq Approximation

5.8 Advection Diffusion Equation

Transport of a macroscopic density, energy or temperature is governed by the Advection-Diffusion-Equation

$$\frac{\partial c}{\partial t} = \nabla \cdot (D \nabla c) - \nabla \cdot (\vec{v} c), \quad (5.10)$$

where c is the considered physical quantity (temperature, particle density), D is the diffusion coefficient and v is a velocity field affecting c . It is possible to solve this

equation in terms of LBM by using an equilibrium distribution function different from the one for the Navier-Stokes Equation [26]

$$g_i^{eq} = w_i \rho \left(1 + \frac{c_i \cdot \vec{v}}{c_s^2} \right), \quad (5.11)$$

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

Listing 5.1: advection diffusion dynamics object

```

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

```

Additionally, a different descriptor with fewer lattice velocities is used [22]:

Listing 5.2: advection diffusion descriptor

```

#define DESCRIPTOR AdvectionDiffusionD3Q7Descriptor

```

In OpenLB D2Q5 and D3Q7 descriptors are implemented for the Advection-Diffusion Equation. Since the Advection-Diffusion Equation simulates different physical conditions than the Navier-Stokes-Equation, 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. To apply this condition, firstly a `sOnLatticeBoundaryCondition3D` object for the advection-diffusion boundaries has to be constructed:

Listing 5.3: advection diffusion dynamics object

```

sOnLatticeBoundaryCondition3D<T, DESCRIPTOR>
  sBoundaryConditionAD(sLattice);
int nC = sBoundaryConditionAD.get_sLattice().getLoadBalancer().
  size();
4 sBoundaryConditionAD.set_overlap(1);
  for (int iC = 0; iC < nC; iC++) {
    OnLatticeAdvectionDiffusionBoundaryCondition3D<T, DESCRIPTOR>*
      ADblockBC =
        createAdvectionDiffusionBoundaryCondition3D<T, DESCRIPTOR,
9      AdvectionDiffusionBGKdynamics<T, DESCRIPTOR> >
        (sBoundaryConditionAD.get_sLattice().getExtendedBlockLattice(
          iC));
    sBoundaryConditionAD.get_CDblockBCs().push_back(ADblockBC);
  }

```

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

Listing 5.4: advection diffusion descriptor

```

void prepareLattice(..., SuperLattice3D<T, DESCRIPTOR>& sLattice,
                  sOnLatticeBoundaryCondition3D<T, DESCRIPTOR>& bc,
3                 SuperGeometry3D<T>& superGeometry, T omega,...) {
    ...
    /// Material=3 -> boundary with constant temerature
    bc.addTemperatureBoundary(superGeometry, 3, omega);
    ...
8 }

```

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

Listing 5.5: add advective velocity on a cell

```

T velocity[3] = {vx,vy,vz};
2 ...
cell.defineExternalField (
    DESCRIPTOR<T>::ExternalField::velocityBeginsAt,
    DESCRIPTOR<T>::ExternalField::sizeOfVelocity,
    velocity);

```

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

Listing 5.6: add advective velocity on a superlattice

```

ConstAnalyticalF3D<T,T> velocity(vel);
...
/// sets advective velocity for material 1
4 superLattice.defineExternalField(superGeometry, 1,
    DESCRIPTOR<T>::ExternalField::velocityBeginsAt,
    DESCRIPTOR<T>::ExternalField::sizeOfVelocity,
    velocity)

```

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

6 Input / Output

During development or even during actual simulation, it might be necessary to parametrize your program. For this case, OpenLB provides an XML parser, which can read files produced by OpenMPI [8], thereby providing a nice GUI, if you are so inclined. Details on the XML format and functions are given in Section 6.6.

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. Every thread writes the data of its cuboids in one VTI file. One ends up with several VTI files, that contain partial simulation data. The key is to use an appropriate PVD file, which links the VTI files to produce a meaningful output. The technical aspects are presented in Section 6.1, whereas the usage is demonstrated with an example in Section 6.2.

6.1 The Output Data Structure in Parallel Simulations

OpenLB simulation data is stored in the VTK data format [9]. This format has XML structure and its data can be written either in human readable ASCII or binary Base64 code. The implemented parallel output structure contains a PVD file, which consists of links to the real data written by the threads and stored in VTI files. An example of a PVD file is shown in Figure 6.1.

For certain time steps, every thread writes its data to a VTI file. Since they do their work independently of each other, the writing happens in parallel. These data files are linked in the aforementioned PVD file and the hierarchical structure of the simulation data is built.

It is still possible to write VTK files sequentially. This method does not need the hierarchical overhead and is much easier to use. Commonly, data like geometry, rank and cuboid is written only once during the simulation. In order to prevent big data overhead and complicated structures, the sequential routine is preferred for those functions.

```

<?xml version="1.0"?>
<VTKFile type="Collection" version="0.1" byte_order="LittleEndian">
<Collection>
<DataSet timestep="360" group="" part=" 0" file="cylinder3D iT0000360iC00000.vti"/>
<DataSet timestep="360" group="" part=" 1" file="cylinder3D iT0000360iC00001.vti"/>
<DataSet timestep="360" group="" part=" 2" file="cylinder3D iT0000360iC00002.vti"/>
<DataSet timestep="360" group="" part=" 3" file="cylinder3D iT0000360iC00003.vti"/>
<DataSet timestep="360" group="" part=" 4" file="cylinder3D iT0000360iC00004.vti"/>
<DataSet timestep="360" group="" part=" 5" file="cylinder3D iT0000360iC00005.vti"/>
<DataSet timestep="360" group="" part=" 6" file="cylinder3D iT0000360iC00006.vti"/>
</Collection>
</VTKFile>

```

Figure 6.1: The PVD file consists of links to the simulation data, which is stored in VTI files. Since every thread writes its data to a single VTI file, this program was executed by 7 threads.

6.2 Data Output to VTK File Format

VTK data files can be visualized and postprocessed with the free software Paraview [5], which offers a nice graphical interface. 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.

```

// binary data format is default
SuperVTKwriter3D<T> vtkWriter("FileNameGoesHere");
3 // ASCII data format is obtained by
// SuperVTKwriter3D<T> vtkWriter("FileNameGoesHere",false);
SuperLatticePhysVelocity3D<T, DESCRIPTOR> velocity(sLattice,
    converter);
SuperLatticePhysPressure3D<T, DESCRIPTOR> pressure(sLattice,
    converter);
vtkWriter.addFunctor( velocity );
8 vtkWriter.addFunctor( pressure );

if (iT==0) {
    SuperLatticeGeometry3D<T, DESCRIPTOR> geometry(sLattice,
        superGeometry);
    SuperLatticeCuboid3D<T, DESCRIPTOR> cuboid(sLattice);
13 // writes the geometry and cuboid no. to a single VTI file
    // sequentially
    vtkWriter.write(geometry);
    vtkWriter.write(cuboid);
    // mandatory to call the following write()-method
    vtkWriter.createMasterFile();
18 }

if (iT%converter.numTimeSteps(.3)==0) {

```

```

23 } // writes the added functors in parallel to vtk files
    vtkWriter.write(iT);

```

6.3 Write Images in OpenLB

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

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 application 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 to reduce the 3D grid with the help of a plane. The resulting plane represents then the image by assigning plane points to pixels.

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

```

2 // get the pointwise l2 norm of velocity
  SuperEuklidNorm3D<T, DESCRIPTOR> normVel( velocity );
  // put a plane with normal (0,0,1) in the 3 dimensional data
  BlockLatticeReduction3D<T, DESCRIPTOR> planeReduction( normVel, 0,
    0, 1 );
  BlockGifWriter<T> gifWriter;
  //gifWriter.write(planeReduction, 0, 0.7, iT, "vel"); //static
  scale
7 gifWriter.write( planeReduction, iT, "vel" ); // scaled

```

6.4 Console output

In OpenLB, there is an extension of default ostreams, which handles parallel output and prefixes each line with the name of the class that produced the output. Shown below is the output of one of the example programs from Section 10:

```

$ ./cylinder3d
[main] Nx=252; Ny=43; Nz=43
[BlockGeometry3D] the model is correct!
[BlockGeometry3D] wrote vti-File
[BlockGeometryStatistics3D] materialNumber=0; count=1892
[BlockGeometryStatistics3D] materialNumber=1; count=416970
[main] step=0; t=0; avEnergy=0; avRho=1; uMax=0
[reIniGeometry] step=0; scalingFactor=3.37314e-12
[main] step=50; t=2.5; avEnergy=6.5764e-08; avRho=0.999936; uMax=0.00507172

```

```
[computeResults] deltaP=-2.0906e-10
```

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 in diagrams, such as convergence rate, data errors, or simple average mass density.

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

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.{h,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, of course, needs its own instance.

6.5 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 2 necessary and 3 optional arguments.

```
STLreader(const std::string fName, T voxelSize, T stlSize=1,
          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. STL file in cm \rightarrow 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).

6.6 XML parameter files

In OpenLB essential simulation parameter can be placed in a XML. This is a useful feature, since once a program is compiled the parameter 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 editing only the XML file. The parsing is implemented in the the header file `io/xmlReader.h`.

The general format for the XML files is:

```

<Param>
  <Mesh>
3     <lx>1</lx>
      <ly>3</ly>
    </Mesh>
    <VisualizationImages>
      <Filename>image</Filename>
8    </VisualizationImages>
</Param>

```

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

```

1 string fName("demo.xml");
  XMLreader config(fName);

  int lx, ly;
  std::string imagename;
6 config["Mesh"]["lx"].get(lx);
  XMLreader meshconfig = config["Mesh"];

  ly = config["Mesh"]["ly"].get<int>();
  config["VisualizationImages"]["Filename"].get(Filename);

```

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 `get()`, which automatically converts the string in the config file to the correct type, if it is one of the basic C++ types. The other method is to call `get` without a parameter but with the needed type as a template parameter, like `get<int>()`. 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 `Mesh`.

7 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. If there is a plus in front of the file name, choose this file to open the numbered collection of single 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 centre.

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.

7.1 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.

Contour

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

7.2 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.

7.3 Stream Tracer

Using the Stream Tracer allows you to draw flow lines.

Temporal Interpolator

Using this filter, you can interpolate between sets of data.

7.4 Transform

Using ‘Transform’ you can change the position and orientation of your objects, as well as the scale.

8 Functors – A General Concept For Input and Output of Data

Roughly speaking, a functor is a class that behaves like a function. Objects of a functor class perform computations by overloading the `operator()`. One big advantage of functors over functions is that they allow the creation of a hierarchy and bundle "classes of functions". Moreover, parameters that are constant over several function evaluations only need to be passed once during instantiation.

8.1 Functors in OpenLB

In OpenLB, functors are used for a wide variety of tasks. They are divided by the unit system they are working in, making extensive use of inheritance, templates and other advantages that comes with C++.

GenericF 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*.

```
template <typename T, typename S>
class GenericF {
protected:
    GenericF(int targetDim, int sourceDim);
5   std::string _name;
private:
    int _n;
    int _m;
public:
10  std::shared_ptr< GenericF<T,S> > _ptrCalcC;
    // has to be implemented for 'every' derived class
    virtual bool operator() (T output[], const S input[]) = 0;
};
```

AnalyticalF is a subclass of `GenericF` for functions that lives 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 a `AnalyticalCalc` class, which inherits from `AnalyticalF` and establishes arithmetic operations (`+`, `-`, `*`, `/`) between every type of `AnalyticalF`.

IndicatorF is an other subclass of `GenericF` that returns a vector with elements 0 or 1. 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.

BlockLatticeF/SuperLatticeF is just an other subclass of `GenericF`. These functors are defined on the lattice and commonly represent the raw simulation data, e.g. pressure, velocity. `SuperLattice` functors are part of the parallelism layer and they delegate the calculations to the corresponding `BlockLattice` functors.

Examples are `SuperLatticeDensity3D : $\Omega \rightarrow \mathbb{R}$` and `SuperLatticeVelocity3D : $\Omega \rightarrow \mathbb{R}^3$` , for a domain $\Omega \subset \mathbb{N}^3$

InterpolationF 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.

8.2 How are functors used?

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

Data output / data extraction Velocity, pressure, cuboids and other information can be extracted from the lattice using predefined functors. All they need to know is a `SuperLattice` and converter; if dimensions are wanted.

Listing 8.1: Code example for calculating velocity and pressure using functors.

```

// Create the data-reading functors...
2 SuperLatticePhysVelocity3D<T, DESCRIPTOR> velocity(&sLattice, &
  converter);
SuperLatticePhysPressure3D<T, DESCRIPTOR> pressure(&sLattice, &
  converter);
// geometries are often constructed by simple geometries
IndicatorSphere3D<bool, T> mySphere(origin, 1);

```

Interpolation Interpolation is necessary in order to start a simulation smoothly or to obtain velocities between those computed on the lattice points. For the start of a simulation, the inflow velocity is smoothly increased from 0 to the desired velocity using a variable called `frac`. It is clear that `frac` should be 0 at the beginning of the simulation and 1 after a certain number of time steps `iTmaxStart`.

Listing 8.2: Code example for smoothly starting the inflow velocity in `cylinder3d` with a x^5 curve.

```
PolynomialStartScale<T, int> nPolynomialStartScale(iTmaxStart, T(1)
);
std::vector<int> iTvec(1, iT);
T frac = nPolynomialStartScale(iTvec)[0];
```

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. Finally, specific norms implemented as functors facilitate analysis of convergence. Application of this is shown in the example `poiseuille2d`, which is discussed in 10.9

Setting boundary values Boundary cells are marked by a certain material number in the `SuperGeometry`. 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 `SuperGeometry`. In the last step, the `SuperLattice` initializes all cells of a certain material given by the `SuperLattice` with the velocities computed by the functor.

Listing 8.3: Code example for setting a Poiseuille velocity profile and a constant pressure boundary in `cylinder3d`.

```
// Creates and sets the Poiseuille inflow profile using functors
2 std::vector<T> maxVelocity(3,0);
maxVelocity[0] = 2.25*frac*converter.getLatticeU();
SquarePoiseuilleInflow3D<T> poiseuilleU(superGeometry, 3,
maxVelocity);
sLattice.defineU(superGeometry, 3, poiseuilleU);
```

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,

$$\Phi = \int \vec{F} \cdot d\vec{A}$$

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

The *flux functor* calculates the discrete flux

$$\Phi_h = h^2 \sum_i \vec{f}_i \cdot \vec{n}$$

with h as the grid length of the surface and \vec{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 \vec{f}_i will be interpolated from the surrounding lattice points.

So, for the SuperLatticeFlux functor a surface needs to be defined, here a plane, and an SuperLatticeF functor.

The *plane* can be defined by a circle indicator, a starting point and a normal, or a starting point and two vectors. Optionally, you can set a *radius* for the plane. The *grid length* of the area can be defined. The default for this value is the lattice length. Another optional feature is a *material list*, so that only the points with the predefined material numbers are used for calculation (the default material number is 1). Next is a *SuperLatticeF functor*, which defines the quantity you want to measure.

Step 1: Define the plane by

a) *a circle indicator*

```
IndicatorCircle3D<T,T> circleInd(center1, center2, center3,
    normal1, normal2, normal3, radius);
```

b) *a normal, a starting point and, optionally, a radius*

```
std::vector<T> startingPoint, planeNormal;
T radius;
```

c) *two vectors, a starting point and, optionally, a radius*

```
std::vector<T> startingPoint, planeVectorU, planeVectorV;
T radius;
```

Step 2 (optional): Define the grid length of the plane

```
T h = converter.getLatticeL();
```

Step 3 (optional): Define the material list

```
std::list<int> materials;
```

Step 4: Create a SuperLatticeF functor

a) *for velocity flow*

```
SuperLatticePhysVelocity3D<T, DESCRIPTOR> vel(sLattice, converter)
;
```

b) *for pressure*

```
SuperLatticePhysPressure3D<T, DESCRIPTOR> press(sLattice,  
converter);
```

c) *or any other SuperLatticeF functor*

```
SuperLatticeF3D<T, DESCRIPTOR> ...;
```

Step 5: create a SuperLatticeFlux functor (depending on how the plane was defined)

a) *circle indicator*

```
SuperLatticeFlux3D(SuperLatticeF3D<T, DESCRIPTOR>& f,  
SuperGeometry3D<T>& sg, IndicatorCircle3D<bool, T>& circle,  
std::list<int> materials, T h = T());
```

b) *normal and startingPoint*

```
2 SuperLatticeFlux3D(SuperLatticeF3D<T, DESCRIPTOR>& f,  
SuperGeometry3D<T>& sg, std::vector<T>& n, std::vector<T>  
A,  
std::list<int> materials, T radius = T(), T h = T());
```

c) *two vectors and startingPoint*

```
2 SuperLatticeFlux3D(SuperLatticeF3D<T, DESCRIPTOR>& f,  
SuperGeometry3D<T>& sg, std::vector<T>& u, std::vector<T>&  
v,  
std::vector<T> A, std::list<int> materials, T radius = T()  
,  
T h = T());
```

In addition to the arguments for the plane, the constructor takes 2 **necessary** and 3 **optional** arguments.

- *f*: the functor defined in **Step 4**
- *sg*: the SuperGeometry3D object
- *materials*: **default** is material number 1
- *radius*: **default** is the diameter of the geometry
- *h*: **default** is the lattice length

Step 6: Get results by using the operator()

```
1 int input [3];  
T output [5];  
flux(output, input);
```

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

Because, in general, the SuperLattice functor is either the velocity functor or the pressure functor, **Step 4** and **Step 5** can be combined. The constructors, depending on how the plane is defined, are identical to the ones used for SuperLatticeFlux, only the SuperLatticeF3D<T, DESCRIPTOR> argument is replaced by the two arguments SuperLattice3D<T, DESCRIPTOR> and LBconverter<T>.

Step 4.1): Combined steps for velocity flow

```

SuperLatticePhysVelocityFlux3D<T, DESCRIPTOR>
2   vFlux(SuperLattice3D<T, DESCRIPTOR> sLattice, LBconverter<
      T> converter,
          ...);

```

Step 4.2): Combined steps for pressure

```

SuperLatticePhysPressureFlux3D<T, DESCRIPTOR>
2   pFlux(SuperLattice3D<T, DESCRIPTOR> sLattice, LBconverter<T>
      converter,
          ...);

```

For these two functors there is a `print()` function.

Step 5.1): Output for velocity functor (region size[m²], volumetric flow rate and mean velocity)

```

vFlux.print(std::string fluxSiScale, std::string meanSiScale);

```

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

Step 5.2): output for pressure functor (region size[m²], force and pressure)

```

pFlux.print(std::string fluxSiScale, std::string meanSiScale);

```

- `fluxSiScale`: 'MN' or 'kN' or '' (default= N)

- *meanSiScale*: 'mmHg' or '' (default=*Pa*)

Shown below are two code examples for the implementation of the flux functor in cylinder3d.

Example 1: *circle indicator, material list and SuperLatticeFlux3D*

Listing 8.4: Code example for getting the volumetric flow rate of the velocity flow in cylinder3d.

```

std::list<int> materials;
materials.push_back(1);
materials.push_back(6);
4 IndicatorCircle3D<bool,T> circleInd(2., 0.205, 0.205, 1., 1., 0.,
    2.);
SuperLatticePhysVelocity3D<T, DESCRIPTOR> vel(sLattice, converter)
;
SuperLatticeFlux3D<T, DESCRIPTOR> flux(vel, superGeometry,
    circleInd);
9 clout << "flowRate=" << flux(input)[0];
clout << "regionSize=" << flux(input)[1] << endl;

```

Example 2: *normal, startingPoint and SuperLatticePhysPressureFlux3D*

Listing 8.5: Code example for getting the pressure on a area in cylinder3d.

```

std::vector<T> A(3, T()), n(3, T());
A[0]=2.;A[1]=0.205;A[2]=0.205;
n[0]=1.;n[1]=1.;n[2]=0.;
5 SuperLatticePhysPressureFlux3D<T, DESCRIPTOR> pFlux(sLattice,
    converter,
        superGeometry, n, A);
pFlux.print();

```

8.3 Functor arithmetic

Simulation data often needs heavy post-processing, in order to get relevant data. With the functor arithmetic OpenLB provides a very user friendly tool to process simulation data during simulation time. E.g. it facilitates the computation of relative errors.

Listing 8.6: Basic showcase for arithmetic operations for AnalyticalF2D.

```

AnalyticalConst2D<T, T> one(1.);

```

```

AnalyticalConst2D<T, T> two(2.);
3 AnalyticalIdentity2D<T, T> tmp(one + two);
  // or equivalent
AnalyticPlus2D<T,T> aPlus(one,two);
AnalyticalIdentity2D<T, T> tmp2(aPlus);

```

To access the computed data, the interface `operator()` is used. Obviously, this is a very rudimentary example. However, it extends the functor concept in a natural way by arithmetic operations.

Listing 8.7: Computation of a relative error with respect to L^2 -norm.

```

int input [1];
T normAnaSol [1], absErr [1], relErr [1];
// define analytical solution: R^3 -> R
4 // for snake of simplicity it is a constant function,
  // however it may be any specialization of AnalyticalF3D<T,T>
AnalyticalConst3D<T, T> dSol(1.);
// get analytical solution on the lattice: N^3 -> R
SuperLatticeFfromAnalyticalF3D<T, DESCRIPTOR> dSolLattice( dSol,
  lattice );
9 // get density out of simulation data
SuperLatticeDensity3D<T, DESCRIPTOR> d( lattice );

SuperL2Norm3D<T, DESCRIPTOR> dL2Norm( dSolLattice - d,
  superGeometry, 1 );
SuperL2Norm3D<T, DESCRIPTOR> dSolL2Norm( dSolLattice,
  superGeometry, 1 );
14 dL2Norm( absErr, input );
dSolL2Norm( normAnaSol, input );
relErr [0] = absErr [0] / normAnaSol [0];
clout << "denstity-L2-error(abs)=" << absErr [0] << ";"
      << "denstity-L2-error(rel)=" << relErr [0] << std::endl;

```

For more detail, see the source code of example 10.9.

The following explains the memory management of the functor arithmetic in OpenLB. It is strongly based on the example shown in Listing 8.6 and in particular on its third line. First, the `operator+()` declared in `AnalyticalF2<T,S>` is called by the object `one`, as shown in Figure 8.1. Its implementation is realized in the file `analyticCalc2D.hh`. Basically, there happens two things. A new object of type `AnalyticalPlus2D<T,S>` will be created and a `shared_ptr` to it, is stored into a variable of the object `one`. The `shared_ptr` is used to free the memory allocated by the new object. By now, object `one` cares about the arithmetic operation. However, if `one` is used for other arithmetic operations, its `shared_ptr` may be overwritten, which can causes runtime errors. It would be more intuitive if `tmp` cared about memory management. As a consequence, `tmp` should hold the `shared_ptr`, which is achieved in two steps. First, constructing an `AnalyticalPlus2D<T,S>` object, moves the `shared_ptr` from object `one` to `AnalyticalPlus2D<T,S>`. Then by constructing `tmp` the `shared_ptr` moves once again to

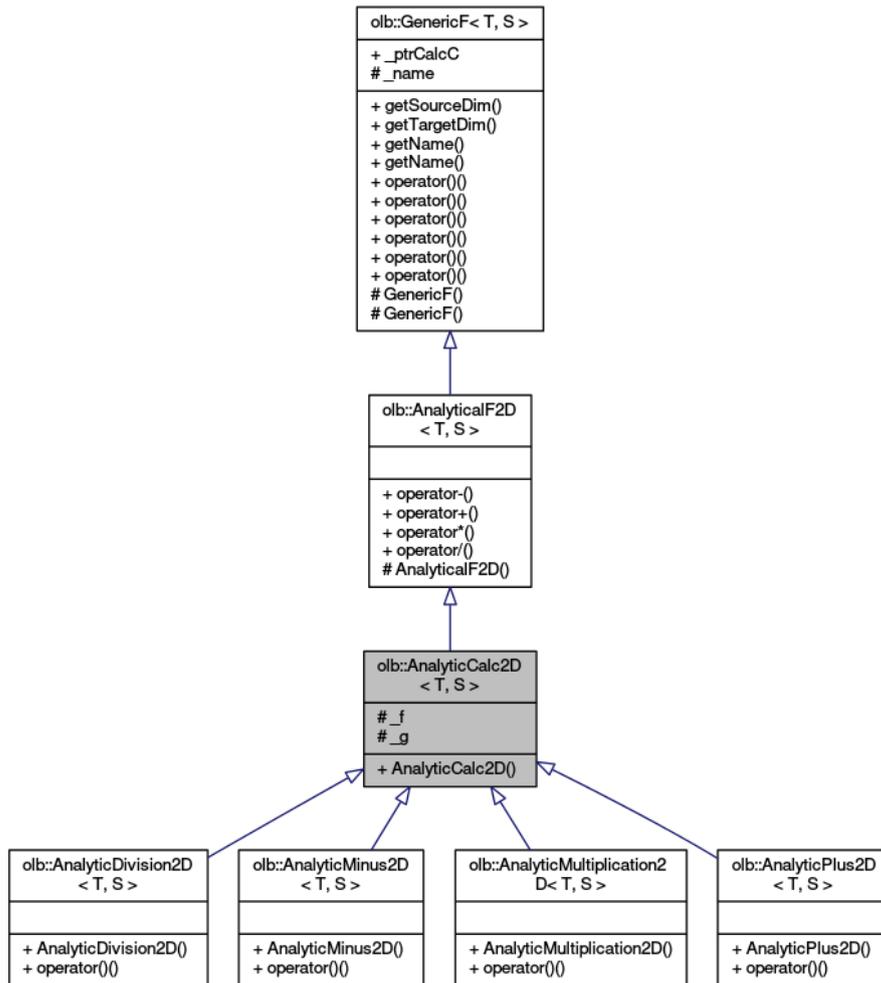


Figure 8.1: Inheritance for AnalyticCalc2D is shown.

the created AnalyticalIdentity2D<T,S>. Finally, tmp holds the shared_ptr and thus is responsible for the memory management.

9 Parallel program execution

Whenever possible, an OpenLB application should be written in such a way that it works well on both serial and parallel platforms. 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. Furthermore, MPI parallelism has become an important option even on simple desktop computers, which quite often possess multi-core processors. In this case, you will often find that MPI is actually more efficient and/or easier to obtain in a non-commercial compiler setting than OpenMP. Fortunately, it is straightforward to write parallelizable applications with OpenLB if a few basic concepts are respected. As a matter of fact, all example programs in the OpenLB distribution can be compiled with MPI and executed in parallel.

To achieve parallelism with programs that have the look and feel of serial applications, OpenLB distinguishes two classes of data. Data which is spatially distributed, such as the lattice and scalar- or vector-valued data 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 strongly recommended to use OpenLB's `OstreamManager` since it can help reducing output in case of parallel execution (cf Chapter 6.4).

9.1 Data-parallel structures

Obtaining data-parallelism in OpenLB is as easy as using the `MultiBlockLatticeXD` instead of a `BlockLatticeXD`, a `MultiScalarFieldXD` instead of a `ScalarFieldXD`, and a `MultiVectorFieldXD` instead of a `VectorFieldXD`. In most common situations, only the case of the `BlockLatticeXD` actually needs to be treated explicitly, and this is handled in a single line of code, as it is shown in Lesson 10 (Section 2.9). Scalar- and vector-valued fields are usually generated automatically, as in the following expression:

```
// This yields an object of type ScalarFieldXD in serial,  
2 // and an object of type MultiScalarFieldXD in parallel  
lattice.getDataAnalysis().getVelocity();
```

The difference between the serial and the parallel case is handled transparently by addressing the data fields through the virtual base `ScalarFieldBaseXD` and `VectorFieldBaseXD`, which is the same for the serial and the parallel data type:

```
// The following instruction works for in serial as well as
2 //   in parallel, because ScalarFieldBase2D is an abstract
//   base to both ScalarField2D and MultiScalarField2D
  ScalarFieldBase2D<T,Lattice> const& velocity
    = lattice.getDataAnalysis().getVelocity();
```

The most important rule to respect when handling data-parallel types in application programs is to never implement explicit loops over space dimensions. Although the resulting code does yield the expected result, it is likely to run very slowly. The reason for this is that the loops cannot be parallelized, and the code therefore runs at the speed of a single processor, or even slower because of the implied MPI communications. An example is given in Section 6, where it is shown how to use predefined functions for I/O operations on data-parallel structures, instead of explicit space loops.

9.2 Duplicated data types

The rule for duplicated data types is simple: all data types except for the data-parallel ones mentioned in the previous section are duplicated. The three following rules need to be respected to ensure that the value from some input is properly duplicated over processors:

1. The call to `olbInit` at the beginning of a program ensures distribution of input from the command-line.
2. The use of `cin` ensures distribution of input from the terminal.
3. The use of `olb ifstream` instead of `fstream` ensures distribution of input from a data file.

10 The example programs

All the demo codes can be compiled with or without MPI, with or without OpenMP, and executed in serial or parallel.

10.1 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 turbulent BGK model is used for the dynamics to stabilize the 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.

10.2 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 exist two implemetations of the problem. The first one is a Euler-Euler ansatz, meaning that the fluid phase as well as the particle phase are modelled as continua. The second is an Euler-Lagrange ansatz, where the particles are modelled as discrete objects.

10.2.1 Euler - Euler

In this example the particles are viewed as a continuum and described by a advection-diffusion equation. This is done similar to the example `thermal3d` 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 `StokesDragAdvectionDiffusionD3Q7Descriptor` 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.

```
SuperExternal3D<T, ADDESCRIPTOR> sExternal(  
    superGeometry, sLatticeAD,  
    ADDESCRIPTOR<T>::ExternalField::velocityBeginsAt,
```

```

5         ADDESCRIPTOR<T>::ExternalField::numScalars,
           sLatticeAD.getOverlap());

...

sExternal.communicate();

```

The function `communicate()` is called in the time loop and handles the communication analogue to the lattices.

Furthermore the new dynamics object `StokesDragAdvectionDiffusionBGKdynamics` is required to access the saved velocity fields correctly and use them in an efficient way. The coupling of the lattices is done by

```

1 StokesDragCouplingGenerator3D<T, NSDESCRIPTOR> coupling(
    converter, radius, partRho,
    ADDESCRIPTOR<T>::ExternalField::velocityBeginsAt);
sLatticeNS.addLatticeCoupling(superGeometry, 1, coupling,
    sLatticeAD);

```

This object also handles the application of the drag force to the velocity field, therefore it takes the particle radius `radius` and density `partRho` as arguments.

For the simulation of particles as a continuum, also new boundary conditions are required. Here `addZeroDistributionBoundary` 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 `ConvectionBoundary` for advection-diffusion lattices can be applied, that approximates a Neumann boundary condition. 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 `addExtFieldBoundary`.

10.2.2 Euler - Lagrange

The main task of his example is to show the using of Lagrangian particles with OpenLB. Similar to the `BlockLattice` and `SuperLattice` structure a `ParticleSystem` and `SuperParticleSystem` structure exists. In line 2 of Listing 10.1 the `SuperParticleSystem` is instantiated. It takes a `SuperGeometry` and `LBConverter` as parameters. In line 4 the `SuperParticleSysVtuWriter` is instantiated. It takes the `SuperParticleSystem`, a filename as `string`, and the wanted particle properties as arguments. Calling the function `SuperParticleSysVtuWriter.write(int timestep)` does create `*.vtu` files of the particles positions for the given timestep. These files can be visualized with Paraview. Line 10 instantiates an interpolation functor for the fluids velocity, which is used in line 13 during the instantiation of `StokesDragForce`. Particles need boundary conditions also. In this example the simplest possible boundary is chosen. 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. This `MaterialBoundary` 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 `SuperParticleSystem`.

The actual number crunching is then done 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 done automatically.

Results of this simulation are published in Henn et al. [23].

Listing 10.1: Usage of class SuperParticleSystem

```

// SuperParticleSystems3D
SuperParticleSystem3D<T, PARTICLE> supParticleSystem(
    superGeometry, converter);
// define which properties are to be written in output data
SuperParticleSysVtuWriter<T, PARTICLE> supParticleWriter(
    supParticleSystem, "particles",
5   SuperParticleSysVtuWriter<T, PARTICLE>::particleProperties::
    velocity |
    SuperParticleSysVtuWriter<T, PARTICLE>::particleProperties::
    mass |
    SuperParticleSysVtuWriter<T, PARTICLE>::particleProperties::
    radius |
    SuperParticleSysVtuWriter<T, PARTICLE>::particleProperties::
    active);

10  SuperLatticeInterpPhysVelocity3D<T, DESCRIPTOR> getVel(sLattice,
    converter);

    auto stokesDragForce = make_shared<StokesDragForce3D<T, PARTICLE
    , DESCRIPTOR>> (getVel, converter);

// material numbers where particles should be reflected
15  std::set<int> boundMaterial = { 2, 4, 5};
    auto materialBoundary = make_shared<MaterialBoundary3D<T,
    PARTICLE>> (superGeometry, boundMaterial);

    supParticleSystem.addForce(stokesDragForce);
    supParticleSystem.addBoundary(materialBoundary);
20  supParticleSystem.setOverlap(2. * converter.getLatticeL());

/* ... */

main loop {
25  supParticleSystem.simulate(converter.physTime());
}

```

Besides the particles the examples uses the save feature of the SuperLattice. By

```
sLattice.save("fluidSolution")
```

and

```
sLattice.load("fluidSolution")
```

the current state of the `SuperLattice` can be saved and loaded again. Using this feature the startup phase for the fluid has to be computed only once.

10.3 `bstep2d` and `bstep3d`

This example implements a backward facing step. Furthermore, it is shown how checkpointing is used to regularly save the state of the simulation.

10.4 `cavity2d` 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 OpenGPI. This example is available in two different versions for sequential and parallel use.

10.5 `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 [33]. 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 [7].

10.6 `multiComponent2d` and `multiComponent3d`

This example demonstrates Rayleigh-Taylor instability in 2D and 3D, generated by a heavy fluid penetrating a light one. The multi-component fluid model by X. Shan and H. Chen is used [29]. These examples show the usage of multicomponent flow and periodic boundaries.

10.7 `nozzle3d`

This example examines a turbulent flow in a nozzle injection tube. 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. $\rho=1$). The example shows the usage of turbulence models.

10.8 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.

10.9 poiseuille2d

This example examines a 2D Poiseuille flow. Computation of error norms via functors is also shown. `bgkPoiseuille2d` and `mrtPoiseuille2d` use a velocity or pressure boundary at the inlet/outlet. In `forcedPoiseuille2d` the boundaries are periodic between the inlet and outlet. As the flow is driven by a body force, it illustrates both the use of a body force and periodic boundaries. In addition to different flavors of BGK [15] and the regularized LB model [25], OpenLB offers implementations of entropic and multiple-relaxation-time (MRT) models. `mrtPoiseuille2d` illustrates the use of MRT. An example program for the entropic model is not yet available.

10.10 thermal2d and thermal3d

This example demonstrates Rayleigh-Bénard convection rolls in 2D and 3D, simulated with the thermal LB model by Guo et al. [21], between a hot plate at the bottom and a cold plate at the top.

10.11 venturi3d

This example examines a steady flow in a venturi tube. At 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. $\rho=1$). The example shows the usage of the Indicator functors to build up a geometry and explains how to set boundary conditions automatically.

11 Bibliography

- [1] The Cygwin project. <http://www.cygwin.com/>.
- [2] LB model with adjustable speed of sound. Technical report. <http://www.lbmethod.org/openlb/techreports.html>.
- [3] The Fifth Element. http://en.wikipedia.org/wiki/The_Fifth_Element.
- [4] How to implement your DdQq dynamics with only q variables per node. Technical report. <http://www.lbmethod.org/openlb/techreports.html>.
- [5] The Paraview project. <http://www.paraview.org>.
- [6] Installing OpenLB in Windows/Cygwin. Technical Report. <http://optilb.org/openlb/tech-reports>.
- [7] FreeCAD: An Open Source parametric 3D CAD modeler. <http://free-cad.sourceforge.net/>.
- [8] The OpenGPI project. <http://www.opengpi.org>.
- [9] The VTK data format documentation. <http://www.vtk.org/VTK/img/file-formats.pdf>.
- [10] S. Ansumali. “Minimal kinetic modeling of hydrodynamics”. PhD thesis. Swiss Federal Institute of Technology Zurich, 2004.
- [11] 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.
- [12] 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.
- [13] 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.
- [14] 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.
- [15] S. Chen and G. D. Doolen. “Lattice Boltzmann Method for Fluid Flows”. In: *Ann. Rev. Fluid Mech.* 30 (1998), pp. 329–364.
- [16] 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.

- [17] 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.
- [18] D. d’Humières, M. Bouzidi, and P. Lallemand. “Thirteen-velocity three-dimensional lattice Boltzmann model”. In: *Phys. Rev. E* 63 (2001), p. 066702. DOI: 10.1103/PhysRevE.63.066702.
- [19] T. Dornieden. “Optimierung von Strömungsgebieten mit adjungierten Lattice Boltzmann Methoden”. Diplomarbeit. Karlsruhe Institute of Technology (KIT), 2013.
- [20] 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.
- [21] 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.
- [22] 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).
- [23] 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.
- [24] T. Inamuro, M. Yoshina, and F. Ogino. “A non-slip boundary condition for lattice Boltzmann simulations”. In: *Phys. Fluids* 7 (1995), pp. 2928–2930.
- [25] J. Latt and B. Chopard. “Lattice Boltzmann Method with regularized non-equilibrium distribution functions”. In: *Math. Comp. Sim.* 72 (2006), pp. 165–168.
- [26] A. A. Mohamad. *Lattice Boltzmann Method - Fundamentals and Engineering Applications with Computer Codes*. Springer-Verlag, 2011.
- [27] 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*. American Institute of Aeronautics and Astronautics. 2013. DOI: doi:10.2514/6.2013-2573.
- [28] G. Pinggen, 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.
- [29] 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.
- [30] X. Shan and G. Doolen. “Multicomponent lattice-Boltzmann model with interparticle interaction”. In: *Journal of Statistical Physics* 81 (1995), pp. 379–393.
- [31] P. A. Skordos. “Initial and boundary conditions for the lattice Boltzmann method”. In: *Phys. Rev. E* 48 (1993), pp. 4824–4842.

- [32] S. Stasius. “Identifikation von Strömungsgebieten mit adjungierten Lattice Boltzmann Methoden (ALBM)”. Diplomarbeit. Karlsruhe Institute for Technology (KIT), 2014.
- [33] 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.
- [34] 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.
- [35] 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.

12 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 non-commercially. 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.