OpenLB User Guide
Associated to Release 0.8 of the Code

openlb.net

Contents

1 Preface

Aims of the user guide

2 Introduction

2.1 Fluid Flow Simulations

2.2 Lattice Boltzmann Methods

2.3 The OpenLB Project

2.3.1 What is OpenLB?

OpenLB is a numerical framework for lattice Boltzmann simulations, created by students and researchers with different background 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. To please the first audience, OpenLB offers a neat interface through which it is possible to set up a simulation with little effort. For the second audience, 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. Basically, generic programming is intended 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.

2.3.2 How to get 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 on the forum at 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 like buildtype and mode of parallelization found in `Makefile.inc` can serve useful information, too.

2.3.3 How to compile 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. In 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 edited (it is easy to understand!) 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.

2.3.4 What features are currently implemented?

Lattice Boltzmann models

BGK model for fluids	Section 6.1.3	Reference [2]
Regularized model for fluids	Section 6.1.3	Reference [3]
Multiple relaxation times (MRT)	Section 6.1.3	References [4, 5]
Entropic Lattice Boltzmann	Section 6.1.3	Reference [6]
BGK with adjustable speed of sound	Section 6.1.3	References [7, 8]
BGK and MRT with Smagorinsky model	Section 6.1.3	References [9]
Porous media model	Section 6.1.3	

Multiphysics coupling

Shan-Chen two-component fluid	Section 6.6	Reference [10]
Thermal fluid with Boussinesq approximation	Section 6.6	Reference [11]

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. [12])
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 [13]

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. Those advanced data structures generate a patchwork consisting of many `BlockLatticeXD` structures that are presented behind a unified interface. Applications of these structures are MPI-parallelism 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.

2.3.5 Participants

In 2006 Vincent Heuveline, Mathias J. Krause and Jonas Latt initiated the OpenLB project. Between 2006 and 2008 Jonas Latt was the project coordinator. Since 2009 Mathias J. Krause is coordinating the project. Since 2006 the following persons have contributed source code to OpenLB:

Lukas Baron: (Parallel) console output, time and performance measurement, porous media model, functors concept.

Jonas Fietz: Configure file parsing based on XML, heuristic load balancer, octree STL reader interface to CVMLCPP.

Thomas Henn: Pre-processing: Voxelizer interface based on STL and CVMLCPP.

Jonas Kratzke: Unit converter, GUI interface based on description files and OpenGPI, Bouzidi boundary condition.

Mathias J. Krause: Hybrid-parallelization approach, OpenMP parallelization, cuboid data structure for MPI parallelization, load balancing, Makefile environment for compilation, functors, Bouzidi boundary condition, integration and maintenance of added components (since 2008).

Jonas Latt: Development of the OpenLB core, integration and maintenance of added components (2006-2008).

Orestis Malaspinas: Alternative boundary conditions (Inamuro, Zou/He, Nonlinear FD), alternative LB models (Entropic LB, MRT).

Albert Mink: Functor arithmetic.

Patrick Nathen: BGK and MRT Smagorinsky model.

Bernd Stahl: 3D extension to MultiBlock structure for MPI parallelization, parallel version of (scalar or tensor-valued) data fields, VTK output of data.

Peter Weisbrod: Examples structure.

Gilles Zahnd: Rotating frame functors.

Simon Zimny: Pre-processing: Automated setting of boundary conditions.

3 Using OpenLB for Applications

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

1st Step: Initialization

2nd Step: Prepare Geometry It first gets the geometry from another file (a stl file here). Then it creates a mesh from that, and prepares the geometry required. 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. For these tasks the function `prepareGeometry` is called. Some examples and applications which use a rather simple geometry skip this step.

3rd Step: Prepare Lattice Then the converter between physics and lattice is set and the dynamic too afterwards. The kinds of dynamics are chosen between the different implementations. These possibilities depend on force acting or not, the single relaxation time (BGK) used or the multi relaxation time (MRT), the simulation dimension (it can also be a 2D model), a compressible or incompressible fluid considered, and the number of neighbouring voxel chosen. The boundary condition initialisation is done in order to enable any kind of them. The lattice is then defined in the function `prepareLattice`, with the boundary condition choices for every material number, and for which materials number dynamics are applied. It only defines the kind of boundary (like Bouzidi, bounce-back, velocity, or pressure) but 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` eventually starts the simulation during which the functions `setBoundaryValues`, `collideAndStream` and `getResults` (step 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 fuctions 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, which performs the collision and the streaming step. If more than one lattice is used, the function is called for each of them seperately.

7th Step: Computation and Output of the 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 structuration is the very same in every OpenLB simulation, only the choices made are changing 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 to the OpenLBs implementation, so new models might need changes or adds in the source code. For example, the classes defined in the code are always issued from a mother-class and have to match to the inputs functions, which may sometimes lead to unexpected issues to solve.

3.1 Lesson 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 the location of your choice. Create a Makefile in this directory according to the procedure explained in Section 2.3.3.

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

Listing 1: Framework of an OpenLB program

```
1   #include "olb2D.h"
2   #ifndef OLB_PRECOMPILED // Unless precompiled version is used,
3     #include "olb2D.hh"   // include full template code
4   #endif
5
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. It cannot be compiled in advance, and needs to be integrated verbatim 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 make clear what the cryptic `#ifndef OLB_PRECOMPILED` is about.

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

Furthermore, for the following examples to compile, the following declarations need to be included into Listing 1 between Line 4 and 6:

```
1   #include <vector>      // Some C++ libraries which are
2   #include <cmath>       //   required for the following
3   #include <iostream>    //   examples
4   #include <iomanip>
5   #include <fstream>
6
7   using namespace olb;  // OpenLB namespaces which are
8   using namespace olb::descriptors; // accessed in the
9   using namespace olb::graphics;    // examples
10  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 modelized 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 the light of compressible flows raises however numerous issues of its own that cannot be covered here. Thus, look at the code and learn your lesson, but don't attribute too much meaning to the numerical result.)*

Listing 2: to be inserted at Line 10 of Listing 1

```
1   #define LATTICE D2Q9Descriptor
2   typedef double T;
3   int nx = 20;
4   int ny = 30;
5   int numIter = 100;
6   T omega = 1.;
7   T r     = 5.;
8
9   int main(int argc, char* argv[]) {
10      olbInit(&argc, &argv);
11      // Insert the central part of your code here
12      BlockLattice2D<T, LATTICE> lattice(nx, ny);
13      BGKdynamics<T, LATTICE> bulkDynamics (
14          omega,
15          instances::getBulkMomenta<T,LATTICE>()
16      );
17      lattice.defineDynamics(0,nx-1,0,ny-1, &bulkDynamics);
18
19      for (int iX=0; iX<nx; ++iX) {
20          for (int iY=0; iY<ny; ++iY) {
```

```
21          T rho=1., u[2] = {0.,0.};
22          if ((iX-nx/2)*(iX-nx/2) + (iY-ny/2)*(iY-ny/2) < r*r) {
23              rho = 1.01;
24          }
25          lattice.get(iX,iY).iniEquilibrium(rho,u);
26      }
27  }
28
29  for (int iT=0; iT<numIter; ++iT) {
30      lattice.collide();
31      lattice.stream(true);
32  }
33
34  ImageWriter<T> imageWriter("leeloo");
35  imageWriter.writeScaledGif (
36      "lesson1",
37      lattice.getDataAnalysis().getVelocityNorm() );
38  }
```

A few explanations are again in order:

Line 1: Choice of a lattice descriptor. Lattice descriptors specify not only what 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 nx×ny lattice and the total number numIter of iteration steps. The relaxation parameter ω is the inverse 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 nx×ny×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 Block-Lattice2D.

Lines 13-16: The Dynamics object determines the implementation of the collision step on grid nodes, in this case BGK [2]. Objects of type BGKdynamics can be customized by indicating how the moments 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 is 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 is 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 of format `PPM`. 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. [14]).

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 3: Substitutions to replace periodic boundaries by no-slip walls

```
1    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 [15].

3.2 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 was tested so far, and it is negligible in face of other, big-picture efficiency considerations.

One such consideration is about the separation between collision and streaming at Line 28 and 29 of Listing 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. [15]. For an OpenLB user, doing this is as easy as replacing the collision-streaming sequence by a call to the method `collideAndStream()`:

Listing 4: Collision and streaming in one step for improved efficiency

```
1  //    collision-streaming cycles
2  //      lattice.collide();
3  //      lattice.stream(true);
4    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 5: Direct access to values in a BlockLattice2D

```
1    int nx, ny, someX, someY, someF;
2  // <...> some code to initialize nx, ny, someX and someY
3    BlockLattice2D<T, LATTICE> lattice(nx,ny); // instantiate BlockLattice
4    T value = lattice.get(someX,someY)[someF]; // read values
5    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 so 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 5:

Listing 6: Manipulation of data through methods of a Cell

```
1    int nx, ny, someX, someY, someF;
2  // <...> some code to initialize nx, ny, someX and someY
3    BlockLattice2D<T, LATTICE> lattice(nx,ny); // instantiate BlockLattice
4  // <...> some code to initialize dynamics objects of the lattice
5    T velocity[2];
6    lattice.get(someX,someY).computeU(velocity); // compute velocity
7    velocity[0] = 0.;
8    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. Inversely, 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.

Additionally to being more convenient, the access to the data in Listing 6 has a distinct advantage to the approach of Listing 5: in Listing 5 the data inside a `Cell<T>` is accessed directly,

whereas in Listing 6 it is accessed indirectly through the dynamics object of the cell. Although direct data access works in simple data structures 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 in 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 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 library OpenLB, for example by writing a class to implement a new type of dynamics. Most subsequent lessons are written for application programmers, and the code is written with extensibility in mind, insisting for example on the possibility to run it in parallel with minimal changes.

The following is a list of 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 7: Instantiation of a 3D lattice

```
1   #define D3Q19Descriptor LATTICE
```

```
2    int nx, ny, nz;
3  // <...> initialization of nx, ny, nz
4    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.

3.3 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 8: Instantiation of OnLatticeBoundaryCondition

```
1  // Instantiate 2D boundary condition
2    OnLatticeBoundaryCondition2D<T,D2Q9Descriptor>* boundaryCondition2D =
3        createLocalBoundaryCondition2D(lattice);
4
5  // Instantiate 3D boundary condition
6    OnLatticeBoundaryCondition2D<T,D3Q19Descriptor>* boundaryCondition3D =
7        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 9: Instantiation of velocity boundary condition along lattice boundaries

```
1    template<typename T>
2    void velocityBoundaryBox (
3        BlockLattice2D<T,D2Q9Descriptor>& lattice,
4        OnLatticeBoundaryCondition2D<T,D2Q9Descriptor>& bc, T omega)
5    {
6        int nx = lattice.getNx();
7        int ny = lattice.getNy();
8        // top boundary
9        bc.addVelocityBoundary1P(1,nx-2,ny-1,ny-1, omega);
10       // bottom boundary
11       bc.addVelocityBoundary1N(1,nx-2,   0,   0, omega);
12       // left boundary
13       bc.addVelocityBoundary0N(0,0, 1, ny-2, omega);
14       // right boundary
15       bc.addVelocityBoundary0P(nx-1,nx-1, 1, ny-2, omega);
```

```
16
17          // Corner nodes
18          bc.addExternalVelocityCornerNN(0,0, omega);
19          bc.addExternalVelocityCornerNP(0,ny-1, omega);
20          bc.addExternalVelocityCornerPN(nx-1,0, omega);
21          bc.addExternalVelocityCornerPP(nx-1,ny-1, omega);
22
23          // Make the lattice ready for simulation
24          lattice.initialize();
25      }
```

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 `addVelocity-Boundary1P` 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 seperate 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 `addVelocity-Boundary0N` instantiates a plane boundary domain in negative x-direction (a left boundary). It takes 6 arguments, additionally to the `omega`-argument 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 you are done with it, specifying the required velocity respectively 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* (that's necessary, because 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 was specified piece-wise (plane per plane, edge

per edge, and corner per corner), the velocity/density can be adapted individually on every node. Furthermore, acessing 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 adapted at every time step to modelize 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, because that's 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 [3], 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 [16], and uses a finite difference scheme over adjacent neighbors to evaluate velocity gradients.

createZouHeBoundaryConditionXD
The local boundary condition introduced by Zou and He [17]. It is very accurate, especially in 2D simulations, but can have stability issues.

createInamuroBoundaryConditionXD
The local boundary condition by Inamuro *et al.* [18]. 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 `createInterpBoundary-ConditionXD`, but this time, non-linear velocity terms of the Chapman-Enksog 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 their help, 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. It 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 gemoetries. 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 10: Implementation of a bounce-back cylinder in the domain center

```
1  <...> definition of the types T and DESCRIPTOR
2  int nx, ny, r;
3  <...> initialization of nx and ny, r
4  BlockLattice2D<T,DESCRIPTOR> lattice(nx,ny);
5  <...> setup of the lattice
6  for (int iX=0; iX<nx; ++iX) {
7      for (int iY=0; iY<ny; ++iY) {
8          if ((iX-nx/2)*(iX-nx/2) + (iY-ny/2)*(iY-ny/2) < r*r) {
9              lattice.defineDynamics(iX,iX,iY,iY,
10                      &instances::getBounceBack<T,D2Q9Descriptor>() );
11         }
12     }
13 }
```

3.4 Lesson 4: Conversion 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 $0,001m^2/s$. We are interested in 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, which 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 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 characteristical quantities of the fluid flow problem. In our 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 commited to the converter. *latticeL* is the discrete space intervall in physical units and from this the dimensionless discretization parameter δ_x is determined: $\delta_x = latticeL/charL$. *latticeU* sets the relation between the discretization parameters for space δ_x and time δ_t in dimensionless units: $latticeU = \delta_u = \delta_t/\delta_x$. Instead of δ_t, LB people 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 11: Use of LBconverter in a 3D problem

```
1  <...> definition of type T
2    int dimension   = 3;
```

```
 3     T latticeL       = (T) 0.003;
 4     T latticeU       = (T) 0.02;
 5     T charNu         = (T) 0.001;
 6     T charL          = (T) 0.03;
 7     T charU          = (T) 4.;
 8     T charRho        = (T) 1.;
 9     T pressureLevel = (T) 0.;
10
11     Lbconverter <T> converter(
12         dimension , latticeL , latticeU ,
13         charNu , charL , charU , charRho , pressureLevel
14     );
15     writeLogFile(converter , "converterLog.dat");
16     cout << converter.getRe() << endl;
17     T omega = converter.getOmega();
18 <...> simulation
19 <...> evaluation of latticeRho at the back and the front of the cylinder
20     T pressureDrop =    converter.physPressure(latticeRhoFront)
21                       - converter.physPressure(latticeRhoBack);
```

Line 2: Specify discretization parameters and characteristical 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.

3.5   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, the average energy (half the square of the velocity norm) and the maximum value of the velocity norm. Those values are accessed trough 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 be careful though 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 through 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 neglibibly 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 in 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, those 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 12: Produce a GIF image from 2D data

```
1  // <...> Create and initialize a variable lattice
2  //        of type BlockLattice2D<T,D2Q9Descriptor>
3    DataAnalysisBase2D<T,D2Q9Descriptor> const& analysis
4        = lattice.getDataAnalysis();
5    ImageWriter<T> imageWriter("earth");
6    imageWriter.writeScaledGif("vorticity", analysis.getVorticity,
7                               200, 200);
```

Line 3: Require an analysis object from the lattice. Alternatively, an instance of the class `Data-AnalysisXD` 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, distinguishing for example between sequential and parallel lattices.

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

Line 6: Calculate 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 13: Produce a GIF image from 3D data

```
1  // <...> Create and initialize a variable lattice
2  //       of type BlockLattice3D<T,D3Q19Descriptor>
3    DataAnalysisBase3D<T,D3Q19Descriptor> const& analysis
4        = lattice.getDataAnalysis();
5    ImageWriter<T> imageWriter("earth");
6  // Extract a slice of the plane defined by z=0
7    int slicePos=0;
8    imageWriter.writeScaledGif (
9        "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 14: Produce a VTK file from 3D data

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

The open source software `Paraview` [19] for example is very useful for the visualization of 3D data contained in such a file.

3.6 Lesson 6: Use an external force

In simulations, the dynamics of a fluid is often driven by a force field (gravity, inter-molecular 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 minimze 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 `Forced-D3Q19Descriptor`. 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 `forced-Poiseuille`.

3.7   Lesson 7: Understand in what sense OpenLB is generic

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 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 in face of 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. It 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 is implemented through a mechanism of dynamic (run-time) genericity. In this way, the dynamics can be different from one cell to another, and it 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 `addPost-ProcessorXD`. 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.

3.8 Lesson 8: Use checkpointing in long-lasting 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 to 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. The basic commands for saving and loading data are `saveData` and `loadData`. They take as first argument the object to be serialized resp. unserialized, and as second argument the filename:

Listing 15: Store and load the state of the simulation

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

Checkpointing is also illustrated in the example programs `bstep2D` (Section 10.2) and `bstep3D` (Section 10.3).

3.9 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.

3.10 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 swith between `BlockLatticeXD` and `MultiBlockLatticeXD`. This can be achieved through a precompiler directive, as in the following code:

Listing 16: MultiBlockLattice2D for MPI-parallel programs

```
1  int nx, ny;
2  <...> initialization of nx and ny
3  #ifndef PARALLEL_MODE_MPI  // sequential program execution
4      BlockLattice2D<T, DESCRIPTOR> lattice(converter.getNx(),
5                                            converter.getNy() );
6  #else                      // parallel program execution
7      MultiBlockLattice2D<T, DESCRIPTOR> lattice (
8          createRegularDataDistribution( converter.getNx(),
9                                         converter.getNy() ) );
10 #endif
```

In a shared-memory environment, OpenMP is an alternative to MPI for parallelism. To parallelize OpenLB with OpenMP, no code needs to be changed at all. Just modify a flag in the Makefile as described below.

To obtain parallel versions of the example programs, modify the flags `CXX` and `PARALLEL_MODE` in the file `Makefile.inc` in the OpenLB root directory. Then, enter the directory of the desired example, eliminate previously compiled libraries (`make clean; make cleanbuild`), and recompile the example by typing the command `make`.

4 Compilation

4.1 Linux

Let us start with a clean Ubuntu 12.04 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.

Then install the g++ compiler which you will need to compile c++ programs:
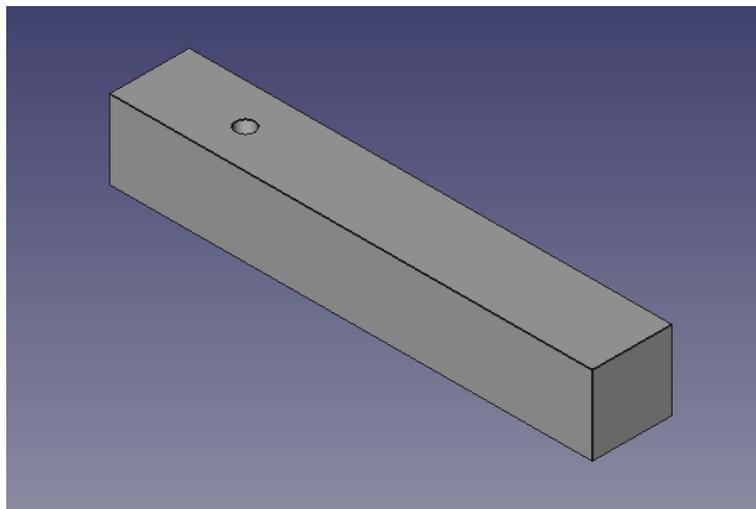
   `sudo apt-get install g++`

2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4
3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4
3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4
3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4
3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4
3	1	1	1	1	1	5	5	5	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4
3	1	1	1	1	5	5	0	5	5	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4
3	1	1	1	1	5	0	0	0	5	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4
3	1	1	1	1	5	5	0	5	5	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4
3	1	1	1	1	1	5	5	5	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4
3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4
3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4
3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4
3	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	4
2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2

Figure 1: Material numbers for a 2d channel flow similar to the example `cylinder2d` from Section 10.7 (1=fluid, 2=no-slip boundary, 3=velocity boundary, 4=constant pressure boundary, 5=curved boundary, 0=do nothing).

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 writen 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.

4.2   Mac

4.3   Windows

5   Geometry

5.1   Material numbers

In OpenLB exists a general concept for representation of a geometry. A specific number called the *material number* is 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 1 illustrates this at the example of an external flow. The reward of using material numbers in flow simulations is to determine the fluid directions on boundary nodes automatically, since this is not always practical by hand e. g. if material numbers of a complex geometry are obtained from a `stl` file.

Figure 2: The geometry for the example `cylinder3d` from Section 10.8 opened in FreeCAD.

5.2 Automated preprocessing

5.3 Creating a geometry

The general process chain assumes that the geometry is already given in form of an `stl` file. Simple geometries can be created using a CAD tool like FreeCAD [21]. An introduction to modelling with FreeCAD can be found for example in `http://www.youtube.com/watch?v=6RxHCR7TLtI`. The general 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 2 is not very difficult, the more complex geometry of a formula one car however can be a challenging and time consuming task!

6 Lattice Boltzmann Models

6.1 Concept – Data Organization

6.1.1 Cell – BlockLattice – SuperLattice

Collide and Stream

7 Input / Output

7.1 Overview

Let's have a look at output first. Imagine you need to extract the values of the velocity field from a simulation and write them to a file. The purpose of this could be to produce images, or to analyze the data. A spontaneous attempt to reach this goal might lead to a code akin to the following:

Listing 17: Manual approach to writing the velocity norm into a text file

```
1    ofstream ofile("velocity.dat");
2    for (int iX=0; iX<nx; ++iX) {
3        for (int iY=0; iY<ny; ++iY) {
4            T u[2];
5            lattice.get(iX,iY).computeVelocity(u);
6            ofile << sqrt(u[0]*u[0]+u[1]*u[1]) << " ";
7        }
8    }
```

This code has several flaws, and is actually quite likely not to work in a MPI-parallel program. Indeed, in such a case the above code is executed on every node. In particular, every node opens a file `velocity.dat` and executes the loop over space directions. The instruction `lattice.get(iX,iY).computeVelocity(u)` is actually properly executed only on the node on which cell (`iX,iY`) is defined, and the result of the computation is broadcasted to every node. If the file `velocity.dat` is located on a hard disk space which is local to each node, it ends up with the expected content. It is however unnecessarily duplicated over each node of the parallel machine. If on the other hand `velocity.dat` is on a shared disk space (a network drive), all nodes write to the same file, which then has a non-deterministic and unusable content.

Technically speaking, this program can be fixed by using the type `olb_ofstream` instead of `ofstream`. In that case, only the main node with identifier 0 creates a file and writes into it, which guarantees a correct content of `velocity.dat`. The code is however horribly unefficient, because every single value of the velocity `u` keeps being unnecessarily broadcasted to every node. As a matter of fact, the type `olb_ofstream` is intended for occasional file access, such as for keeping track of the time evolution of the kinetic energy, averaged over space. To save spatially extended data to a file, dedicated functions should be used instead. The following command for example can be used to get an ASCII output of the velocity norm:

```
1    saveAsciiData(lattice.getAnalysis().getVelocityNorm(),
2                  "velocity.dat");
```

This one-line instruction produces an efficient code which works both in serial and parallel. Furthermore, it is arguably easier to read and to understand than the code at the beginning of this section.

Section 9 explains the difference between large distributed data, and small duplicated data in MPI-parallel OpenLB applications. This distinction has a direct consequence on the recommended choice of I/O operations. Section 7.2 explains how to handled small data volumes, such as the parameters of the simulation or the time-evolution of average quantities. Section 7.3 lists the available options for I/O on distributed data sets. Finally, Section 7.4 contains a few recommendations on how to post-process the data.

During development or even during actual simulation, it might be neccessary to parametrize your program. For this case, OpenLB provides an XML parser, which can read files produced by OpenGPI [22], thereby even providing a nice GUI if you are so inclined. Details on the XML format and functions are given in Section 7.6.

7.2 Console output

In OpenLB exists an extension of default ostreams that handles parallel output and prefixes each line with the name of the class that produced the output. Here 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 as well in debugging process as in 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 easily possible to visualize every imaginable data like convergence rate, data errors, or simple average mass density in diagrams.

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

Using the `OstreamManager` is easy and consists of two parts. First, an instance of the class `OstreamManager` is needed, the one created here in Line 2 is called `clout` like all the other instances in OpenLB. This word consists one the one hand of the two words class and output, on the other hand it is quite similar to standard `cout`. The constructor gets two arguments, one describing the

ostream to use, the other one setting the prefix-text. In line 4 is shown the usage of an instance of the `OstreamManager`. 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 prefered. This is done by adding a `mutable OstreamManager` object as 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 benefits of `OstreamManager` is the reduction of output in parallel. Running a program using `cout` on multiple cores means normally to get one output line for each process. OstreamManager will avoid this by default and display only the output of the first processor. If that behavior is not wanted in a specific case, it can be turned off for an instance named `clout` by `clout.setMultiOutput(true)`.

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

7.3 I/O of spatially extended data and system checkpointing

OpenLB supports I/O of spatially extended data with two formats. In the ASCII format, all numbers are written on a single line, separated by a space charachter. In the binary format, the data is written in ASCII characters using Base64 encoding.

The command `saveAsciidata` is used to write ASCII files:

```
1   // Compute the velocity norm in each point and write an ASCII file
2     saveAsciiData(lattice.getDataAnalysis().getVelocityNorm(),
3                   "velocityNorm.dat");
4   // Compute the pressure norm in each point and write an ASCII file
5     saveAsciiData(lattice.getDataAnalysis().getPressure(),
6                   "pressure.dat");
```

For Base64 encoded binary files, use the command `saveData`:

```
1   // Save the velocity file into a Base64 encoded binary file
2     saveData(lattice.getDataAnalysis().getVelocity(), "velocity.dat");
3   // Save all values of the lattice, including external scalars,
4   //   into a binary file
5     saveData(lattice, "lattice.checkpoint")
```

Unlike the ASCII format, which loses accuracy, the binary format saves an exact match of the floating point numbers. It can therefore be used to save the state of the system, and restore it later:

```
1   // Restore the state of a previously saved lattice, note that you
2   //   need to create the lattice first, and that you are solely
3   //   responsible for allocating the right amount of memory.
4     loadData(lattice, "lattice.checkpoint")
```

7.4 Producing images and post-processing the data

OpenLB also supports saving data in a binary VTK format, using the class `VtkImageOutput2D` for 2D data, and `VtkImageOutput3D` for 3D data. This corresponds essentially to the binary Base64 format explained in the previous section, with some additional information indicated in an XML structure. This additional information is used to specify the name of data fields, and the physical units of space dimensions. The following listing shows how to write a VTK file, containing the pressure and velocity field of a simulation:

```
1    // dx is the space step between two grid points
2    double dx = converter.getDeltaX();
3    // dt is the time interval between two iterations
4    double dt = converter.getDeltaT();
5    // Create a VTK file, and specify the value of grid spacing
6    VtkImageOutput3D<double> vtkOut("vtkOutput.dat", dx);
7    // Write the pressure, and restore its physical units
8    vtkOut.writeData<float>(analysis.getPressure(), "pressure", dx*dx/dt*dt);
9    // Write the velocity, and restore its physical units
10   vtkOut.writeData<3,float>(analysis.getVelocity(), "velocity", dx/dt);
```

In this example, it is assumed that the simulation is run with double precision floating point arithmetics. This high precision is often required for the simulation, but not for the postprocessing of the data. To divide storage requirements by a factor two, the data is therefore converted to single precision in this example, by use of the template parameter `float` to the function `writeData`.

VTK data files can for example be visualized and postprocessed with the free software Paraview [19], which offers a nice graphical interface.

7.5 Read and write .stl-files

OpenLB enables the possibility to read and write geometry data in the Standard Triangulation Language, short: stl. 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.

7.6 XML parameter files

As explained in the introduction, OpenGPI provides an API to access configuration data for your application. This might come in handy to avoid multiple recompilations while searching for optimal parameters or in general development. The parsing is implemented in the the header tile `io/xmlReader.h`.

The general format for the XML files is

```
1   <Param>
2       <Mesh>
3           <lx>1</lx>
4           <ly>3</ly>
5       </Mesh>
6       <VisualizationImages>
7           <Filename>image</Filename>
8       </VisualizationImages>
9   </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");
2   XMLreader config(fName);
3
4   int lx, ly;
5   std::string imagename;
6   config["Mesh"]["lx"].get(lx);
7   XMLreader meshconfig = config["Mesh"];
8
9   ly = config["Mesh"]["ly"].get<int>();
10  config["VisualizationImages"]["Filename"].get(Filename);
```

Let us examine the code above. 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, there are now multiple methods. One is to pass a predefined variable to 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 paramenter, 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`.

8   Functors – a general concept for input and output of data

8.1   What are functors?

Functors are generally spoken classes that behave like functions. Instances of a functor perform some computations by overloading the `operator()`. The big advantage of functors over functions is that they allow to create a hierarchy and bundle "classes of functions". Moreover, parameters that are constant over several function evaluations need to be passed only once during instantiation.

8.2   Functors in OpenLB

In OpenLB, functors are used for a wide variety of tasks. They are divided by the unit system they are working on, making excessive use of heritage, templates and other stuff that C++ provides. Figure 3 gives a visual overview over the different classes of functors in OpenLB and the relation between them.

GenericF   stands on top of the hierarchy. It defines only the template parameters `S` for the input data type, `T` for output and the existence of an unwritten `operator()` as interface. `GenericF` combines `AnalyticalF` – functors acting on the SI level – with `Block/SuperLatticeF` acting on the specific lattice-Boltzmann units. Arithmetic operations need to be implemented only on this level and only once. By using the conversion functions of the two subclasses it is possible to compute the difference between an analytically given, exact solution and a computed discrete solution.
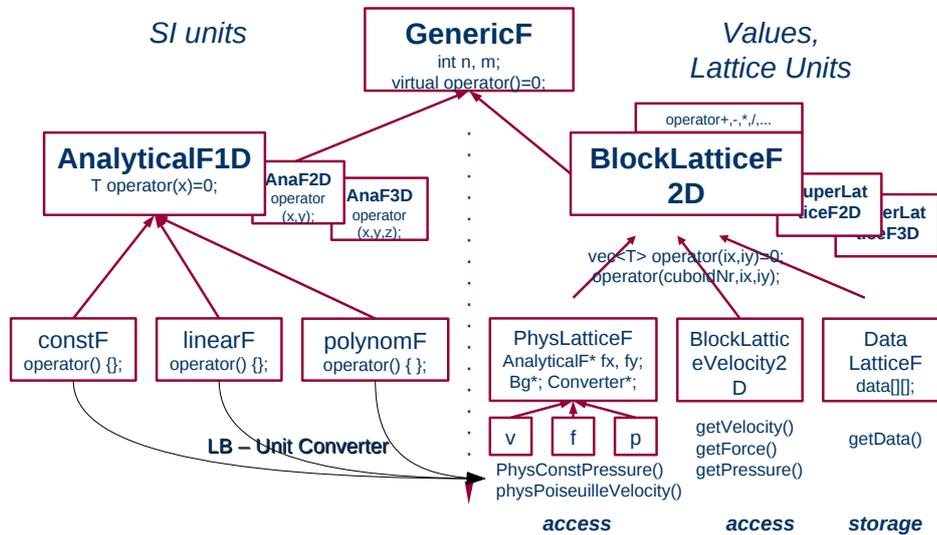
Figure 3: Overview of functors in OpenLB.

AnalyticalF A subclass of GenericF for functions that live in SI-units, e.g. for setting velocities in m/s. The input parameter type is fixed to `int`. Part of this class are e.g. constant functions, interpolation functions of linear and higher order that expect a number of points, and topological indicator functions for spheres for spheres..

Block/SuperLatticeF A subclass of GenericF for functions that live in the world of lattice-Boltzmann units, expecting discrete integer coordinates and returning lattice-Boltzmann units.

8.3 How to use these functors?

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 for Code example for writing velocity, pressure, cuboids and other information in the example cylinder3d from section 10.8.

Listing 18: Code example for writing velocity and pressure using functors.

```
1  // Create the data-reading functors...
2  SuperLatticePhysVelocity3D<T, DESCRIPTOR> velocity(&sLattice, &converter);
3  SuperLatticePhysPressure3D<T, DESCRIPTOR> pressure(&sLattice, &converter);
4  // ...and pass them to the generic write function
5  CuboidVTKout3D<T>::write<DESCRIPTOR>(velocity, iT);
6  CuboidVTKout3D<T>::write<DESCRIPTOR>(pressure, iT);
```

Interpolation Interpolation is necessary to smoothly start a simulation or to obtain velocities between the computed ones 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`.

Code example particles...

Listing 19: Code example for smoothly starting the inflow velocity in cylinder3d with a x^5 curve.

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

Setting of boundary values   Boundary cells are marked by a certain material number in SuperGeometry. Using a functor, velocities can be set at once on all cells of this material. First a vector is necessary that characterizes the maximum flow velocity and its directions. 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 SuperGeometry. In the last step, SuperLattice initializes all cells of a certain material given by SuperLattice with the velocities computed by the functor.

Listing 20: Code example for setting a Poiseuille velocity profile and a constant pressure boundary in cylinder3d.

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

9   Parallel program execution

Whenever possible, an OpenLB application should be written in such a way that it works well on both serial and a parallel platforms. Indeed, as applications in computational fluid dynamics require a large amount of resources, it is important to be able to switch to a parallel platform in a flexible way. 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 that case, you will often find that MPI is actually more efficient and/or easier to obtain in a non-commercial compiler setting than OpenMP. It is fortunately straightforward to write parallelizable application 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 which have the look and fell of serial applications, OpenLB distinguishes two classes of data. Data which is spatially distributed, such as the lattice, or scalar-respectively vector-valued data fields, is handled through a data-parallel paradigm. The data space is partioned into smaller regions that are distributed over the nodes of a parallel machine. In the following, this type of structures are referred to as data-parallel strucures. Other data types which require a small amount of storage space are duplicated on every node, and they are referred to

in the following 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 7.2).

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 point is handled in a single line in the code, as it is for example shown in Lesson 10 (Section 3.10). Scalar- and vector-valued fields are usually generated automatically, as in the following expression:

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

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

```
1  // The following instruction works for in serial as well as
2  //   in parallel, because ScalarFieldBase2D is an abstract
3  //   base to both ScalarField2D and MultiScalarField2D
4    ScalarFieldBase2D<T,Lattice> const& velocity
5      = 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 7.3, 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, and 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 fully automatically. As dynamics a Smagorinsky turbulent BGK model is used to stabilize the simulation for low resolutions. As output the pressure drop is computed. The results has been validated by comparison with other results obtained with FEM and FVM.

10.2 bstep2d

In the implementation of a backward facing step, the definition of the domain geometry is slightly more complicated than previously: additionally to edges and exterior corners, the boundary contains an interior node. This example shows how to define boundaries in such a case. It is furthermore shown how to use checkpointing and save the state of the simulation regularly.

10.3 bstep3d

This is the 3D version of the backward facing step. It illustrates the use of various types of boundary nodes in 3D, including surfaces, edges, and corners. The output of the simulation includes 2D images of the velocity-norm along slices, 3D VTK data that can be visualized for example with Paraview [19], and binary data that stores the state of the simulation in regular time intervals.

10.4 cavity2d

This 2D example illustrates a flow in a rectangular, lid-driven cavity. It also shows how to use the XML parameter files and has an example description file for OpenGPI.

10.5 cavity3d

This is a 3D extension of the lid-driven cavity.

10.6 cavity3d_cuboid_mpi

The 3D lid-driven cavity, parallelized with the `CuboidStructure3D` instead of `MultiBlockStructure3D`.

10.7 cylinder2d

This example is only slightly more complicated than the previous one. A cylindrical obstacle, implemented via bounce-back nodes, is placed in the flow and leads to von Karman instability.

10.8   cylinder3d

This is a 3D extension of the cylindrical obstacle. The geometry is taken from a benchmark [23] where researchers from different groups apply different methods in order to compare their results. The model was created using the open source CAD tool FreeCAD [21], the project file is also available. As results the pressure drop between directly before and behind the cylinder as well at the drag and lift is computed. They converge and in the range presented in [23].

10.9   forcedPoiseuille2d

In this implementation of a Poiseuille flow, the boundaries are periodic between inlet and outlet. The flow is driven by a body force instead of a pressure gradient. This simple example illustrates the use of a body force in OpenLB.

10.10   mrt2d

Additionally to different flavors of BGK [2] and the regularized LB model [3], OpenLB offers implementations of entropic and multiple-relaxation-time (MRT) models. The present example illustrates the use of MRT. An example program for the entropic model is not yet available.

10.11   multiComponent2d

Rayleigh-Taylor instability in 2D, generated by a heavy fluid penetrating a light one. The multi-component fluid model by X. Shan and H. Chen is used [10].

10.12   multiComponent3d

Rayleigh-Taylor instability in 3D, generated by a heavy fluid penetrating a light one. The multi-component fluid model by X. Shan and H. Chen is used [10].

10.13   poiseuille2d

This implementation of a 2D Poiseuille flow demonstrates the setup of a basic simulation in OpenLB. The demo offers the choice of local and non-local boundaries, and the choice of implementing the inlet/outlet with a velocity or pressure boundary condition. It further shows how to extract statistics on the state of the simulation (for example the average energy), and how to produce a snapshot of the velocity field in terms of a `GIF` image.

10.14   thermal2d

Rayleigh-Bénard convection rolls in 2D, simulated with the thermal LB model by Z. Guo e.a. [11].

10.15   thermal3d

Rayleigh-Bénard convection rolls in 3D, simulated with the thermal LB model by Z. Guo e.a. [11].

References

[1] The Cygwin project. `http://www.cygwin.com/`.

[2] S. Chen and G. D. Doolen. Lattice Boltzmann method for fluid flows. *Ann. Rev. Fluid Mech.*, 30:329–364, 1998.

[3] J. Latt and B. Chopard. Lattice boltzmann method with regularized non-equilibrium distribution functions. *Math. Comp. Sim.*, 72:165–168, 2006.

[4] D. d'Humières, I. Ginzburg, M. Krafczyk, P. Lallemand, and L.-S. Luo. Multiple-relaxation-time lattice Boltzmann models in three dimensions. *Phil. Trans. R. Soc. Lond. A*, 360:437–451, 2002.

[5] D. Yu, R. Mei, L.-S. Luo, , and W. Shyy. Viscous flow computations with the method of lattice Boltzmann equation. *Prog. Aerosp. Science*, 39:329–367, 2003.

[6] Santosh Ansumali. *Minimal kinetic modeling of hydrodynamics*. PhD thesis, Swiss Federal Institute of Technology Zurich, 2004.

[7] LB model with adjustable speed of sound. Technical report. `http://www.lbmethod.org/openlb/techreports.html`.

[8] Bastien Chopard, Alexandre Dupuis, Alexandre Masselot, and Pascal Luthi. Cellular automata and lattice Boltzmann techniques: an approach to model and simulate complex systems. *Adv. Compl. Sys.*, 5:103–246, 2002.

[9] 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.

[10] Xiaowen Shan and Hudong Chen. Lattice Boltzmann model for simulating flows with multiple phases and components. *Phys. Rev. E*, 47:1815–1819, 1993.

[11] Zhaoli Guo, Baochang Shi, and Chuguang Zheng. A coupled lattice BGK model for the Boussinesq equations. *Int. J. Num. Meth. Fluids*, 39:325–342, 2002.

[12] Dominique d'Humières, M'hamed Bouzidi, and Pierre Lallemand. Thirteen-velocity three-dimensional lattice Boltzmann model. *Phys. Rev. E*, 63:066702, 2001.

[13] M'hamed Bouzidi, Mouaouia Firdaouss, and Pierre Lallemand. Momentum transfer of a Boltzmann-lattice fluid with boundaries. *Physics of Fluids*, 13(11):3452–3459, 2001.

[14] The Fifth Element. `http://en.wikipedia.org/wiki/The_Fifth_Element`.

[15] How to implement your DdQq dynamics with only q variables per node. Technical report. `http://www.lbmethod.org/openlb/techreports.html`.

[16] P. A. Skordos. Initial and boundary conditions for the lattice Boltzmann method. *Phys. Rev. E*, 48:4824–4842, 1993.

[17] Q. Zou and X. He. On pressure and velocity boundary conditions for the lattice Boltzmann BGK model. *Phys. Fluids*, 9:1592–1598, 1997.

[18] T. Inamuro, M. Yoshina, and F. Ogino. A non-slip boundary condition for lattice Boltzmann simulations. *Phys. Fluids*, 7:2928–2930, 1995.

[19] The Paraview project. `http://www.paraview.org`.

[20] Z. Guo, C. Zheng, and B. Shi. Discrete lattice effects on the forcing term in the lattice Boltzmann method. *Phys. Rev. E*, 65:046308, 2002.

[21] FreeCAD: An Open Source parametric 3D CAD modeler. `http://free-cad.sourceforge.net/`.

[22] The OpenGPI project. `http://www.opengpi.org`.

[23] S. Turek and M. Schäfer. Benchmark computations of laminar flow around cylinder. In *Flow Simulation with High-Performance Computers II*, volume 52 of *Notes on Numerical Fluid Mechanics*, pages 547–566. Vieweg, January 1996.

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

Preamble

The purpose of this License is to make a manual, textbook, or other functional and useful document "free" in the sense of freedom: to assure everyone the effective freedom to copy and redistribute it, with or without modifying it, either commercially or noncommercially. Secondarily, 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.